SAN hosts and cloud clients

SAN hosts and cloud clients

NetApp
March 29, 2024

Table of Contents

SAN hosts and cloud clients

SAN host configuration overview

This content provides host and version-specific best practices to quickly configure SAN hosts connected to ONTAP storage.

Your SAN hosts and ONTAP storage can reside either on-premises, in the cloud, or both.

> In a cloud environment, hosts are typically referred to as clients. All host-specific information in this content also applies to cloud clients.

For example, you can configure on-premises ONTAP storage to connect to on-premises SAN hosts or to connect to SAN cloud clients. You could also configure Cloud Volumes ONTAP or Amazon FSx for ONTAP to connect to on-premises SAN hosts or to connect to SAN cloud clients.

Correct configuration is important for best performance and successful failover.

Related information

- The ONTAP SAN Configuration for your version of ONTAP
- The ONTAP SAN Administration Guide for your version of ONTAP
- The ONTAP Release Notes for your version of ONTAP
- The E-Series documentation to find SANtricity related documentation.

Install SAN host utilities

Overview

Discover the latest information on SAN host utility releases and access the installation procedure for your host configuration.

> For reliable operation after installation, use the NetApp Interoperability Matrix Tool to verify that your host supports the complete NVMe over Fabrics (including NVMe over TCP and NVMe over Fibre Channel), iSCSI, FC, or FCoE configuration.

AIX Host Utilities

AIX Host Utilities 6.1 Release Notes

The release notes describe new features and enhancements, issues fixed in the current release, known problems and limitations, and important cautions related to configuring and managing your specific AIX host with your ONTAP storage system.

For specific information about the operating system versions and updates that the Host Utilities support, see the NetApp Interoperability Matrix Tool.

What's new

The AIX Host Utilities 6.1 release contains the following new features and enhancements:

- AIX Host Utilities 6.1 added support for the memory fault issue that occurred in earlier versions of the AIX host OS. With AIX Host Utilities 6.1, only the sanlun binary has changed. The MPIO and related ODM remain unchanged.

Fixed in this release

BugID	Title	Description
872113	`sanlun lun show -p` command might cause a memory fault on some versions of AIX host OS	Intermittent instances of AIX coredump are reported while running the `sanlun lun show -p` command. Sanlun's `lun show -p` option provides the multipathing information for all the LUNs discovered on a host. It arranges this information to present which SCSI device is sourced from which LUN, the path state (primary or secondary), and other details. However, on some AIX hosts running the `sanlun lun show -p` command might cause a memory fault. This issue is observed only when you run the sanlun command with the `-p` option.

Known problems and limitations

You should be aware of the following known problems and limitations that might impact performance on your specific host.

Bug ID	Title	Description
1069147	AIX HU Sanlun reports incorrect HBA speed	Instances of sanlun displaying incorrect HBA speeds are reported while running the `sanlun fcp show adapter -v` command. The `sanlun fcp show adapter -v` command displays the HBA cards information, such as supported and negotiated speeds for the adapters. This seems to be a reporting issue only. To identify the actual speed, use the `fcstat fcsx` command.

NetApp Bugs Online provides complete information for most known issues, including suggested workarounds where possible. Some keyword combinations and bug types that you might want to use include the following:

- FCP General: Displays FC and HBA bugs that are not associated with a specific host.
- FCP - AIX

Install AIX Host Utilities 6.1

The AIX Unified Host Utilities assists you to manage NetApp ONTAP storage attached to an AIX host.

AIX Host Utilities support the following protocols:

- FC
- FCoE
- iSCSI

AIX Host Utilities support the following environments:

- AIX MPIO (Native OS)
- PowerVM

  For more information about PowerVM, see the IBM PowerVM Live Partition Mobility Red Book.

What you'll need

- For reliable operation, verify that your entire iSCSI, FC, or FCoE configuration is supported.

  You can use the NetApp Interoperability Matrix Tool to verify your configuration.

- Dynamic tracking must be enabled for all FC and FCoE initiators.

> The NetApp AIX Host Utilities software package is available on the NetApp Support Site in a compressed tar.gz file. You must install the AIX Host Utilities kit while using AIX MPIO with NetApp ONTAP Storage.

Steps

1. Log in to your host.
   - On an AIX host, log in as **root**.
   - On a PowerVM host, log in as **padmin**, and then enter the `oem_setup_env` command to become root.
2. Go to the NetApp Support Site and download the compressed file containing the Host Utilities to a directory on your host.
3. Go to the directory containing the download.
4. Decompress the file and extract the SAN toolkit software package.

   ```
   tar -xvf ntap_aix_host_utilities_6.1.tar.gz
   ```

   The following directory is created when you decompress the file: `ntap_aix_host_utilities_6.1`. This directory will have one of the following subdirectories: MPIO, NON_MPIO, or SAN_Tool_Kit.

5. Install the AIX MPIO:

   ```
   installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/MPIO
   NetApp.MPIO_Host_Utilities_Kit
   ```

6. Install the SAN toolkit:

   ```
   installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/SAN_Tool_Kit
   NetApp.SAN_toolkit
   ```

7. Reboot the host.
8. Verify the installation:

```
`sanlun version`
```

SAN Toolkit

AIX Host Utilities is a NetApp host software that provides a command line toolkit on your IBM AIX host. The toolkit is installed when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns the LUN information.

```
#sanlun lun show all
```

Example output

```
controller(7mode)/ device host lun

vserver(Cmode) lun-pathname filename adapter protocol size mode
----------------------------------------------------------------
data_vserver   /vol/vol1/lun1 hdisk0 fcs0    FCP      60g C
data_vserver   /vol/vol2/lun2 hdisk0 fcs0    FCP      20g C
data_vserver   /vol/vol3/lun3 hdisk11 fcs0   FCP      20g C
data_vserver   /vol/vol4/lun4 hdisk14 fcs0   FCP      20g C
```

This toolkit is common across all Host Utilities configurations and protocols. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance. The SAN toolkit is supported on AIX and PowerVM/VIOS OS versions.

AIX Host Utilities 6.1 sample command reference

You can use the AIX Host Utilities 6.1 sample command reference for an end-to-end validation of the NetApp storage configuration using the host utilities tool.

List all host initiators mapped to host

You can retrieve a list of host initiators mapped to a host.

```
# sanlun fcp show adapter -v
```

Example output

```
bash-3.2# sanlun fcp show adapter -v
adapter name: fcs0
WWPN: 100000109b22e143
WWNN: 200000109b22e143
driver name: /usr/lib/drivers/pci/emfcdd
model: df1000e31410150
model description: FC Adapter
serial number: YA50HY79S117
hardware version: Not Available
driver version: 7.2.5.0
firmware version: 00012000040025700027
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 16 GBit/sec
negotiated speed: Unknown
OS device name: fcs0
adapter name: fcs1
WWPN: 100000109b22e144
WWNN: 200000109b22e144
driver name: /usr/lib/drivers/pci/emfcdd
model: df1000e31410150
model description: FC Adapter
serial number: YA50HY79S117
hardware version: Not Available
driver version: 7.2.5.0
firmware version: 00012000040025700027
Number of ports: 1
port type: Fabric
port state: Operational
supported speed: 16 GBit/sec
negotiated speed: Unknown
OS device name: fcs1
bash-3.2#
```

List all LUNs mapped to host

You can retrieve a list of all LUNs mapped to a host.

```
# sanlun lun show -p -v all
```

Example output

```
ONTAP Path: vs_aix_clus:/vol/gpfs_205p2_207p1_vol_0_8/aix_205p2_207p1_lun
LUN: 88
LUN Size: 15g
Host Device: hdisk9
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   fcs1    fc_aix_2     1
up     secondary path2   fcs0    fc_aix_3     1
up     secondary path3   fcs1    fc_aix_4     1
```

List all LUNs mapped to host from a given SVM

You can retrieve a list of all LUNs mapped to a host from a specified SVM.

```
# sanlun lun show -p -v sanboot_unix
```

Example output

```
ONTAP Path: sanboot_unix:/vol/aix_205p2_boot_0/boot_205p2_lun
LUN: 0
LUN Size: 80.0g
Host Device: hdisk85
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    sanboot_1   1
up      primary    path1 fcs1    sanboot_2   1
up      secondary  path2 fcs0    sanboot_3   1
up      secondary  path3 fcs1    sanboot_4   1
```

List all attributes of a given LUN mapped to host

You can retrieve a list of all attributes of a specified LUN mapped to a host.

```
# sanlun lun show -p -v
vs_aix_clus:/vol/gpfs_205p2_207p1_vol_0_8/aix_205p2_207p1_lun
```

Example output

```
ONTAP Path: vs_aix_clus:/vol/gpfs_205p2_207p1_vol_0_8/aix_205p2_207p1_lun
LUN: 88
LUN Size: 15g
Host Device: hdisk9
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 fcs1    fc_aix_2   1
up       secondary path2 fcs0    fc_aix_3   1
up       secondary path3 fcs1    fc_aix_4   1
```

List ONTAP LUN attributes by host device filename

You can retrieve a list of ONTAP LUN attributes by specifying a host device filename.

```
#sanlun lun show -d /dev/hdisk1
```

Example output

```
controller(7mode)/
device host lun
vserver(Cmode)     lun-pathname
--------------------------------------------------------------------------
---
vs_aix_clus       /vol/gpfs_205p2_207p1_vol_0_0/aix_205p2_207p1_lun

filename adapter protocol size mode
-----------------------------------
hdisk1    fcs0    FCP     15g  C
```

List all SVM target LIF WWPNs attached to host

You can retrieve a list of all SVM target LIF WWPNs attached to a host.

```
# sanlun lun show -wwpn
```

Example output

```
controller(7mode)/
target device host lun
vserver(Cmode)          wwpn            lun-pathname
--------------------------------------------------------------------------
------

vs_aix_clus          203300a098ba7afe
/vol/gpfs_205p2_207p1_vol_0_0/aix_205p2_207p1_lun
vs_aix_clus          203300a098ba7afe
/vol/gpfs_205p2_207p1_vol_0_9/aix_205p2_207p1_lun
vs_aix_clus          203300a098ba7afe
/vol/gpfs_205p2_207p1_vol_en_0_0/aix_205p2_207p1_lun_en
vs_aix_clus          202f00a098ba7afe
/vol/gpfs_205p2_207p1_vol_en_0_1/aix_205p2_207p1_lun_en

filename     adapter    size  mode
-----------------------------------
hdisk1       fcs0       15g    C
hdisk10      fcs0       15g    C
hdisk11      fcs0       15g    C
hdisk12      fcs0       15g    C
```

HP-UX Host Utilities

HP-UX Host Utilities 6.0 Release Notes

The release notes describe new features and enhancements, issues fixed in the current release, known problems and limitations, and important cautions related to configuring and managing your specific HP-UX host with your ONTAP storage system.

The HP-UX Host Utilities 6.0 continues to support the following versions:

- HP-UX 11iv2
- HP-UX 11iv3

There are no new features, enhancements, known limitations, or cautions for the HP-UX host utilities 6.0 release.

HP-UX Host Utilities 6.0

The HP-UX Host Utilities enable you to connect an HP-UX host to NetApp storage.

The HP-UX Host Utilities supports multiple protocols and the following environments:

- Native MPIO
- Veritas Dynamic Multipathing (DMP)

> ⓘ To indicate which environment is being used, this document sometimes specifies "DMP" for the Veritas DMP environment and "MPIO" for the HP-UX native environment. In some cases, the commands you use might vary depending on which drivers you are using. In those cases, both the environment and driver types are specified.

What you'll need

- For reliable operation, verify that your entire iSCSI, FC, or FCoE configuration is supported.

You can use the NetApp Interoperability Matrix Tool to verify your configuration.

About this task

The NetApp HP-UX Host Utilities software package is available on the NetApp Support Site in a compressed file. After you download the file, you must decompress it before installation.

Steps

1. Log in to your host.
2. Download the HP-UX Host Utilities file `netapp_hpux_host_utilities_6.0_ia_pa.depot.gz` from the NetApp Support Site to your HP-UX host.
3. Decompress the `netapp_hpux_host_utilities_6.0_ia_pa.depot.gz` file:

   ```
   # gunzip netapp_hpux_host_utilities_6.0_ia_pa.depot.gz
   ```

   The system places the extracted software in the directory where you uncompressed the depot file.
4. Install the software:

```
# swinstall -s /depot_path
```

`depot_path` provides the path and name of the depot file.

The `swinstall` command runs an installation script that verifies the status of your HP-UX setup. If your system meets the requirements, this script installs the `sanlun` utility and diagnostic scripts in the `/opt/NetApp/santools/bin` directory.

5. Verify the installation:

```
sanlun version
```

SAN Toolkit

HP-UX Host Utilities is a NetApp host software that provides a command line tool kit on your HP-UX host. The toolkit is installed when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility which helps you manage the LUNs and host bus adapters. (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns the LUN information.

```
# sanlun lun show all

controller(7mode)/ device host lun
vserver(Cmode)           lun-pathname                        filename
adapter    protocol   size   mode
--------------------------------------------------------------------------
----------------------------------
sanboot_unix         /vol/hpux_boot/boot_hpux_lun      /dev/rdsk/c34t0d0
fclp1      FCP        150g    C
sanboot_unix         /vol/hpux_boot/boot_hpux_lun      /dev/rdsk/c23t0d0
fclp1      FCP        150g    C
sanboot_unix         /vol/hpux_boot/boot_hpux_lun      /dev/rdsk/c12t0d0
fclp0      FCP        150g    C
sanboot_unix         /vol/hpux_boot/boot_hpux_lun      /dev/rdsk/c81t0d0
fclp0      FCP        150g    C
```

This toolkit is common across all Host Utilities configurations and protocols. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance.

HP-UX Host Utilities 6.0 command reference

You can use the HP-UX Unified Host Utilities 6.0 sample command reference for an end-to-end validation of the NetApp storage configuration using the host utilities tool.

List all host initiators mapped to host

You can retrieve a list of all host initiators mapped to a host.

```
# sanlun fcp show adapter -v
```

Example output

```
adapter name:      fclp2
WWPN:              10000000c985ef92
WWNN:              20000000c985ef92
driver name:       fclp
model:             AJ763-63001
model description: HP 8Gb Dual Channel PCI-e 2.0 FC HBA
serial number:     MY19034N9U
hardware version:  3
driver version:    @(#) FCLP: PCIe Fibre Channel driver (FibrChanl-02),
B.11.31.1805, Feb  5 2018, FCLP_IFC (3,2)
firmware version:  2.02X2 SLI-3 (U3D2.02X2)
Number of ports:   1 of 2
port type:         Unknown
port state:        Link Down
supported speed:   8 GBit/sec
negotiated speed:  Speed not established
OS device name:    /dev/fclp2

adapter name:      fclp3
WWPN:              10000000c985ef93
WWNN:              20000000c985ef93
driver name:       fclp
model:             AJ763-63001
model description: HP 8Gb Dual Channel PCI-e 2.0 FC HBA
serial number:     MY19034N9U
hardware version:  3
driver version:    @(#) FCLP: PCIe Fibre Channel driver (FibrChanl-02),
B.11.31.1805, Feb  5 2018, FCLP_IFC (3,2)
firmware version:  2.02X2 SLI-3 (U3D2.02X2)
Number of ports:   2 of 2
port type:         Unknown
port state:        Link Down
supported speed:   8 GBit/sec
negotiated speed:  Speed not established
OS device name:    /dev/fclp3
```

List all LUNs mapped to host

You can retrieve a list of all LUNs mapped to a host.

```
# sanlun lun show -p -v all
```

Example output

```
\
                    ONTAP Path:
vs_hp_cluster:/vol/chathpux_217_vol_en_1_10/hp_en_217_lun
                           LUN: 55
                      LUN Size: 15g
                   Host Device: /dev/rdisk/disk718
                          Mode: C
                            VG: /dev/vg_data
              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/c37t6d7     fclp0   hpux_7
0
up       primary    /dev/dsk/c22t6d7     fclp1   hpux_8
0
up       secondary  /dev/dsk/c36t6d7     fclp0   hpux_5
1
up       secondary  /dev/dsk/c44t6d7     fclp1   hpux_6
1
```

List all LUNs mapped to host from a given SVM

You can retrieve a list of all LUNs mapped to host from a certain SVM.

```
# sanlun lun show -p -v vs_hp_cluster
```

Example output

```
                    ONTAP Path:
vs_hp_cluster:/vol/chathpux_217_vol_en_1_10/hp_en_217_lun
                           LUN: 55
                      LUN Size: 15g
                   Host Device: /dev/rdisk/disk718
                          Mode: C
                            VG: /dev/vg_data
              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/c37t6d7     fclp0   hpux_7
0
up       primary    /dev/dsk/c22t6d7     fclp1   hpux_8
0
up       secondary  /dev/dsk/c36t6d7     fclp0   hpux_5
1
up       secondary  /dev/dsk/c44t6d7     fclp1   hpux_6
1
```

List all attributes of a given LUN mapped to host

You can retrieve a list of all attributes of a specified LUN mapped to a host.

```
# sanlun lun show -p -v
vs_hp_cluster:/vol/chathpux_217_vol_en_1_5/hp_en_217_lun
```

Example output

```
                    ONTAP Path:
vs_hp_cluster:/vol/chathpux_217_vol_en_1_5/hp_en_217_lun
                           LUN: 49
                      LUN Size: 15g
                   Host Device: /dev/rdisk/disk712
                          Mode: C
                            VG: /dev/vg_data
              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/c37t6d1     fclp0   hpux_7
0
up       primary    /dev/dsk/c22t6d1     fclp1   hpux_8
0
up       secondary  /dev/dsk/c36t6d1     fclp0   hpux_5
1
up       secondary  /dev/dsk/c44t6d1     fclp1   hpux_6
1
```

List ONTAP LUN attributes by host device filename

You can retrieve a list of ONTAP LUN attributes by a specified host device filename.

```
#sanlun lun show -dv /dev/rdisk/disk716
```

Example output

```
                                                                    device
host                  lun
vserver              lun-pathname                                   filename
adapter    protocol   size    mode
--------------------------------------------------------------------------
--------------------------------------------
vs_hp_cluster        /vol/chathpux_217_vol_en_1_14/hp_en_217_lun
/dev/rdisk/disk716 0         FCP        15g     C
             LUN Serial number: 80D7l?NiNP5U
         Controller Model Name: AFF-A800
          Vserver FCP nodename: 208400a098ba7afe
          Vserver FCP portname: 207e00a098ba7afe
              Vserver LIF name: hpux_5
            Vserver IP address: 10.141.54.30
                                10.141.54.35
                                10.141.54.37
                                10.141.54.33
                                10.141.54.31
           Vserver volume name: chathpux_217_vol_en_1_14
MSID::0x00000000000000000000000080915935
         Vserver snapshot name:
```

List all SVM target LIF WWPNs attached to host

You can retrieve a list of all SVM target LIF WWPNs attached to a host.

```
# sanlun lun show -wwpn
```

Example output

```
controller(7mode)/
vserver(Cmode)          target wwpn         lun-pathname
device filename
---------------------------------------------------------------------------
------------------------------
vs_hp_cluster           208300a098ba7afe
/vol/chathpux_217_vol_en_1_10/hp_en_217_lun    /dev/rdsk/c22t6d7
vs_hp_cluster           208100a098ba7afe
/vol/chathpux_217_vol_en_1_10/hp_en_217_lun    /dev/rdsk/c44t6d7
vs_hp_cluster           208200a098ba7afe
/vol/chathpux_217_vol_en_1_10/hp_en_217_lun    /dev/rdsk/c37t6d7
vs_hp_cluster           207e00a098ba7afe
/vol/chathpux_217_vol_en_1_10/hp_en_217_lun    /dev/rdsk/c36t6d7
vs_hp_cluster           207d00a098ba7afe   /vol/chathpux_217_os/hp_217_os
/dev/rdsk/c18t7d4
vs_hp_cluster           207f00a098ba7afe   /vol/chathpux_217_os/hp_217_os
/dev/rdsk/c42t7d4

host adapter     lun size     mode
---------------------------------
fclp1            15g          C
fclp1            15g          C
fclp0            15g          C
fclp0            15g          C
fclp1            30g          C
fclp0            30g          C
```

Linux Unified Host Utilities

Linux Unified Host Utilities 7.1 Release Notes

The release notes describe new features and enhancements, known problems and limitations, and important cautions for configuring and managing your specific host with your ONTAP storage system.

For specific information about the operating system versions and updates that the Host Utilities support, see the NetApp Interoperability Matrix Tool.

What's New

The Linux Host Utilities 7.1 release contains the following new features and enhancements:

- Linux Host Utilities is now called *Linux Unified Host Utilities* because it supports NetApp E-Series storage systems running SANtricity as well as AFF, FAS, and ASA systems running ONTAP.

> Any mention of Host Utilities or Linux Host Utilities in this document refers to Linux Unified Host Utilities.

- The following operating systems are now supported:
  - SUSE Linux Enterprise Server 15 series
  - Oracle VM 3.2 series
  - Oracle Linux 6 and 7 series
  - Red Hat Enterprise Linux 6 and 7 series
  - SUSE Linux Enterprise Server 11 SP4
  - KVM and XEN, RHEV 6.4 and 6.5
  - Citrix XenServer
- On Red Hat Enterprise Linux (RHEL) 6 and RHEL 7 hosts, a tuned package for setting server profiles is now supported. You can use the `tuned-adm` command to set different profiles, depending on the environment. For example, you can also use the virtual guest profile as a guest virtual machine and you can use the enterprise storage profile for configurations where LUNs from enterprise storage arrays are used. Using these tuned packages can result in improvement in throughput and latency in ONTAP.
- Adds support for 32GB FC adapters from Broadcom Emulex and Marvell Qlogic.

> NetApp continues to work with the Host Utilities to add support for features after the initial release. For latest information about the features that are supported and the new features that have been added, see the NetApp Interoperability Matrix Tool.

Fixed in this release

The intermittent host OS failure issue that occurs when running the `sanlun lun show -p` command in SLES12SP1, OL7.2, RHEL7.2, and RHEL 6.8 is fixed in this release.

Known problems and limitations

The Linux Host Utilities 7.1 release has the following known problems and limitations.

NetApp Bug ID	Title	Description
1457017	`sanlun` installation issues warning messages related to `libdevmapper.so` and `libnl.so` libraries. These warnings do not affect the functionality of `sanlun` kit.	When you execute the Linux Unified 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 Refer to the public report 1508554.

NetApp Bugs Online provides complete information for most known issues, including suggested workarounds where possible.

Install Linux Unified Host Utilities 7.1

The Linux Unified Host Utilities (LUHU) assists you to manage NetApp ONTAP storage attached to a Linux host. NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

The following Linux distributions are supported:

- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server
- Oracle Linux
- Oracle VM
- Citrix XenServer

What you'll need

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit or 64-bit .rpm file.

- For reliable operation, you must verify that your entire iSCSI, FC, or FCoE configuration is supported.

  You can use the NetApp Interoperability Matrix Tool to verify your configuration.

- You must install the host bus adapter (HBA) management packages available on the vendor support site.

  The management software enables the SAN toolkit commands to gather information about the FC HBAs, such as their WWPNs. For the `sanlun fcp show adapter` command to work, verify that the following packages are correctly installed:

  - Marvell QLogic HBA – QConvergeConsole CLI
  - Broadcom Emulex HBA - OneCommand Manager core application CLI
  - Marvell Brocade HBA – Brocade Command Utility CLI

- RPM Packages "libhbaapi" and "libhbalinux" available for each Linux distribution should be installed on the host OS.

> ⓘ Linux Unified Host Utilities software does not support NVMe over Fibre Channel (NVMe/FC) and NVMe over TCP (NVMe/TCP) host protocols.

Steps

1. If you have a version of Linux Unified Host Utilities currently installed, use the following command to remove it:

```
rpm -e netapp_linux_unified_host_utilities-7-1
```

For older versions of Linux Host Utilities, go to the directory where the host utility software is installed and enter the uninstall command to remove the installed package.

2. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
3. Go to the directory to which you downloaded the software package and use the following command to install it:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_xx.rpm
```

Example output

```
Verifying...                          #################################
[100%]
Preparing...                          #################################
[100%]
Updating / installing...
   1:netapp_linux_unified_host_utiliti#################################
[100%]
```

4. Verify the installation:

```
sanlun version
```

Example output

```
sanlun version 7.1.386.1644
```

Recommended driver settings with Linux kernel

When you configure an FC environment that uses native inbox drivers that are bundled with the Linux kernel, you can use the default values for the drivers.

SAN Toolkit

Linux Unified Host Utilities is a NetApp host software that provides a command line tool kit on your Linux host.

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                  lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

- This toolkit is common across all configurations and protocols of the Host Utilities. As a result, some of its contents apply to one configuration, but not another. Having unused components does not affect your system performance.
- The SAN toolkit is not supported on Citrix XenServer, Oracle VM, and Red Hat Enterprise Virtualization Hypervisor.

Linux Unified Host Utilities 7.1 command reference

You can use the Linux Unified Host Utilities 7.1 sample command reference for an end-to-end validation of the NetApp storage configuration using the host utilities tool.

List all host initiators mapped to host

You can retrieve a list of all host initiators mapped to a host.

```
# sanlun fcp show adapter -v
```

Example output

```
adapter name:      host15
WWPN:              10000090fa022736
WWNN:              20000090fa022736
driver name:       lpfc
model:             LPe16002B-M6
model description: Emulex LPe16002B-M6 PCIe 2-port 16Gb Fibre Channel
Adapter
serial number:     FC24637890
hardware version:  0000000b 00000010 00000000
driver version:    12.8.0.5; HBAAPI(I) v2.3.d, 07-12-10
firmware version:  12.8.340.8
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec, 8 GBit/sec, 16 GBit/sec
negotiated speed:  16 GBit/sec
OS device name:    /sys/class/scsi_host/host15

adapter name:      host16
WWPN:              10000090fa022737
WWNN:              20000090fa022737
driver name:       lpfc
model:             LPe16002B-M6
model description: Emulex LPe16002B-M6 PCIe 2-port 16Gb Fibre Channel
Adapter
serial number:     FC24637890
hardware version:  0000000b 00000010 00000000
driver version:    12.8.0.5; HBAAPI(I) v2.3.d, 07-12-10
firmware version:  12.8.340.8
Number of ports:   1
port type:         Fabric
port state:        Operational
supported speed:   4 GBit/sec, 8 GBit/sec, 16 GBit/sec
negotiated speed:  16 GBit/sec
OS device name:    /sys/class/scsi_host/host16
```

List all LUNs mapped to host

You can retrieve a list of all LUNs mapped to a host.

```
# sanlun lun show -p -v all
```

Example output

```
                    ONTAP Path: vs_sanboot:/vol/sanboot_169/lun
                           LUN: 0
                      LUN Size: 150g
                       Product: cDOT
                   Host Device: 3600a0980383143393124515873683561
              Multipath Policy: service-time 0
                DM-MP Features: 3 queue_if_no_path pg_init_retries 50
              Hardware Handler: 1 alua
            Multipath Provider: Native
----------- --------- ---------- ------- ------------ ----------
----------
            host      vserver           host:
dm-mp       path      path       /dev/   chan:        vserver    major:
state       state     type       node    id:lun       LIF        minor
----------- --------- ---------- ------- ------------ ----------
----------
active      up        primary    sdq    15:0:5:0     lif_18      65:0
active      up        primary    sds    16:0:5:0     lif_17      65:32
active      up        primary    sdac   16:0:7:0     lif_25      65:192
active      up        primary    sdad   15:0:7:0     lif_26      65:208
active      up        secondary  sdt    15:0:4:0     lif_20      65:48
active      up        secondary  sdr    15:0:6:0     lif_19      65:16
active      up        secondary  sdad   16:0:4:0     lif_27      66:96
active      up        secondary  sdan   16:0:6:0     lif_28      66:112
```

List all LUNs mapped to host from a given SVM

You can retrieve a list of all LUNs mapped to a host from a specific storage VM (SVM).

```
# sanlun lun show -p -v vs_sanboot
```

Example output

```
                    ONTAP Path: vs_sanboot:/vol/sanboot_169/lun
                           LUN: 0
                      LUN Size: 160g
                       Product: cDOT
                   Host Device: 3600a0980383143393124515873683561
              Multipath Policy: service-time 0
                DM-MP Features: 3 queue_if_no_path pg_init_retries 50
              Hardware Handler: 1 alua
            Multipath Provider: Native
----------- --------- ---------- ------- ------------ ---------------
----------
            host      vserver            host:
dm-mp       path      path       /dev/   chan:        vserver
major:
state       state     type       node    id:lun       LIF
minor
----------- --------- ---------- ------- ------------ ---------------
----------
active      up        primary    sdce    15:0:5:0     lif_16g_5
69:32
active      up        primary    sdfk    16:0:5:0     lif_16g_7
130:96
active      up        primary    sdfm    16:0:7:0     lif_16g_8
130:128
active      up        primary    sdcg    15:0:7:0     lif_16g_6
69:64
active      up        secondary  sdcd    15:0:4:0     lif_16g_1
69:16
active      up        secondary  sdcf    15:0:6:0     lif_16g_2
69:48
active      up        secondary  sdfj    16:0:4:0     lif_16g_3
130:80
active      up        secondary  sdfl    16:0:6:0     lif_16g_4
130:112
```

List all attributes of a given LUN mapped to host

You can retrieve a list of all attributes of a specified LUN mapped to a host.

```
# sanlun lun show -p -v vs_sanboot:/vol/sanboot_169/lun
```

Example output

```
                   ONTAP Path: vs_sanboot:/vol/sanboot_169/lun
                          LUN: 0
                     LUN Size: 160g
                      Product: cDOT
                  Host Device: 3600a0980383143393124515873683561
             Multipath Policy: service-time 0
               DM-MP Features: 3 queue_if_no_path pg_init_retries 50
             Hardware Handler: 1 alua
           Multipath Provider: Native
----------- --------- ---------- ------- ------------ -----------------
----------
            host      vserver            host:
dm-mp       path      path       /dev/   chan:        vserver
major:
state       state     type       node    id:lun       LIF
minor
----------- --------- ---------- ------- ------------ -----------------
----------
active      up        primary    sdce    15:0:5:0     lif_16g_5
69:32
active      up        primary    sdfk    16:0:5:0     lif_16g_7
130:96
active      up        primary    sdfm    16:0:7:0     lif_16g_8
130:128
active      up        primary    sdcg    15:0:7:0     lif_16g_6
69:64
active      up        secondary  sdcd    15:0:4:0     lif_16g_1
69:16
active      up        secondary  sdcf    15:0:6:0     lif_16g_2
69:48
active      up        secondary  sdfj    16:0:4:0     lif_16g_3
130:80
active      up        secondary  sdfl    16:0:6:0     lif_16g_4
130:112
```

List the ONTAP SVM identity from which a given LUN is mapped to host

You can retrieve a list of ONTAP SVM identity from which a specific LUN is mapped to a hist.

```
# sanlun lun show -m -v vs_sanboot:/vol/sanboot_169/lun
```

Example output

```
                                                             device
host                  lun
vserver                       lun-pathname                   filename
adapter    protocol   size    product
---------------------------------------------------------------------------
-------------------------------------
vs_sanboot                    /vol/sanboot_169/lun           /dev/sdfm
host16     FCP        160g    cDOT
             LUN Serial number: 81C91$QXsh5a
         Controller Model Name: AFF-A400
          Vserver FCP nodename: 2008d039ea1308e5
          Vserver FCP portname: 2010d039ea1308e5
              Vserver LIF name: lif_16g_8
            Vserver IP address: 10.141.12.165
                                10.141.12.161
                                10.141.12.163
           Vserver volume name: sanboot_169
MSID::0x00000000000000000000000809E7CC3
         Vserver snapshot name:
```

List ONTAP LUN attributes by host device filename

You can retrieve a list of ONTAP LUN attributes by a host device filename.

```
# sanlun lun show -d /dev/sdce
```

Example output

```
controller(7mode/E-Series)/                          device        host
lun
vserver(cDOT/FlashRay)      lun-pathname           filename      adapter
protocol   size    product
---------------------------------------------------------------------------
--------------------------
vs_sanboot                 /vol/sanboot_169/lun    /dev/sdce     host15
FCP        160g    cDOT
[root@sr630-13-169 ~]#
```

List all SVM target LIF WWPNs attached to host

You can retrieve a list of all SVM target LIF WWPNs attached to a host.

```
# sanlun lun show -wwpn
```

Example output

```
controller(7mode/E-Series)/   target
device          host       lun
vserver(cDOT/FlashRay)        wwpn                 lun-pathname
filename        adapter    size    product
---------------------------------------------------------------------------
---------------------------------------------
vs_169_16gEmu                 202cd039ea1308e5   /vol/VOL_8g_169_2_8/lun
/dev/sdlo       host18     10g     cDOT
vs_169_16gEmu                 202cd039ea1308e5   /vol/VOL_8g_169_2_9/lun
/dev/sdlp       host18     10g     cDOT
vs_169_16gEmu                 202cd039ea1308e5   /vol/VOL_8g_169_2_7/lun
/dev/sdln       host18     10g     cDOT
vs_169_16gEmu                 202cd039ea1308e5   /vol/VOL_8g_169_2_5/lun
/dev/sdll       host18     10g     cDOT
```

List ONTAP LUNs seen on host by a given SVM target LIF WWPN

You can retrieve a list of ONTAP LUNs noticed on a host by a specified SVM target LIF WWPN.

```
# sanlun lun show -wwpn 2010d039ea1308e5
```

Example output

```
controller(7mode/E-Series)/   target
device         host       lun
vserver(cDOT/FlashRay)        wwpn                 lun-pathname
filename       adapter    size    product
---------------------------------------------------------------------------
-------------------------------------
vs_sanboot                    2010d039ea1308e5   /vol/sanboot_169/lun
/dev/sdfm      host16     160g    cDOT
```

Solaris Host Utilities

Solaris Host Utilities 6.2 Release Notes

The release notes describe new features and enhancements, issues fixed in the current release, known problems and limitations, and important cautions related to configuring and managing your specific Solaris host with your ONTAP storage system..

For specific information about the operating system versions and updates that the Host Utilities support, see the NetApp Interoperability Matrix Tool.

Known problems and limitations

You should be aware of the following known problems and limitations that might impact performance on your specific host.

Bug ID	Title	Description
1385189	Solaris 11.4 FC driver binding changes required in HUK 6.2	Solaris 11.4 and HUK recommendations: FC driver binding is changed from `ssd(4D)` to `sd(4D)`. Move configuration that you have in `ssd.conf` to `sd.conf` as mentioned in Oracle (Doc ID 2595926.1). The behavior varies across newly installed Solaris 11.4 systems and systems upgraded from 11.3 or lower versions.

NetApp Bugs Online provides complete information for most known issues, including suggested workarounds where possible. Some keyword combinations and bug types that you might want to use include the following:

- FCP General: Displays FC and host bus adapter (HBA) bugs that are not associated with a specific host.
- FCP - Solaris

Install Solaris Host Utilities 6.2

The Solaris Unified Host Utilities assists you in managing NetApp ONTAP storage attached to a Solaris host.

The Solaris Host Utilities 6.2 supports several Solaris environments and multiple protocols. The primary host utilities environments are:

- Native OS with MPxIO with either the Fibre Channel (FC) or iSCSI protocol on a system using either a SPARC processor or an x86/64 processor.
- Veritas Dynamic Multipathing (DMP) with either the FC or iSCSI protocol on a system using a SPARC processor, or the iSCSI protocol on a system using an x86/64 processor.

The Solaris Unified Host Utilities 6.2 continues to support the following versions of Solaris:

- Solaris 11.x series
- Solaris 10.x series

What you'll need

- For reliable operation, verify that your entire iSCSI,FC or FCoE configuration is supported.

You can use the NetApp Interoperability Matrix Tool to verify your configuration.

The NetApp Solaris Host Utilities software package is available on the NetApp Support Site in a compressed file format for your processor. You can download the Host Utilities software package for your environment from the Support site.

Steps

1. Login to your host as root.
2. Download a copy of the compressed file containing the Host Utilities from the NetApp Support Site to a directory on your host.

   At the time this documentation was prepared, the compressed files were called:

   - SPARC CPU: `netapp_solaris_host_utilities_6_2_sparc.tar.gz`
   - x86/x64 CPU: `netapp_solaris_host_utilities_6_2_amd.tar.gz`
3. Go to the directory containing the download.
4. Unzip the file using the `gunzip` command:

   ```
   # gunzip netapp_solaris_host_utilities_6_2_sparc.tar.gz
   ```
5. Unzip the file. You can use the `tar xvf` command to do this.

   ```
   # tar xvf netapp_solaris_host_utilities_6_2_sparc.tar
   ```
6. Add the packages that you extracted from tar file to your host. You can use the `pkgadd` command to do this.

   The packages are added to the `/opt/NTAP/SANToolkit/bin` directory. The following example uses the `pkgadd` command to install the Solaris installation package:

   ```
   # pkgadd -d ./NTAPSANTool.pkg
   ```
7. Confirm that the toolkit was successfully installed by using the `pkginfo` command or the `ls - al` command.

```
# ls -alR /opt/NTAP/SANToolkit
/opt/NTAP/SANToolkit:
total 1038
drwxr-xr-x   3 root     sys            4 Jul 22  2019 .
drwxr-xr-x   3 root     sys            3 Jul 22  2019 ..
drwxr-xr-x   2 root     sys            6 Jul 22  2019 bin
-r-xr-xr-x   1 root     sys       432666 Sep 13  2017 NOTICES.PDF

/opt/NTAP/SANToolkit/bin:
total 7962
drwxr-xr-x   2 root     sys            6 Jul 22  2019 .
drwxr-xr-x   3 root     sys            4 Jul 22  2019 ..
-r-xr-xr-x   1 root     sys      2308252 Sep 13  2017 host_config
-r-xr-xr-x   1 root     sys          995 Sep 13  2017 san_version
-r-xr-xr-x   1 root     sys      1669204 Sep 13  2017 sanlun
-r-xr-xr-x   1 root     sys          677 Sep 13  2017 vidpid.dat

# (cd /usr/share/man/man1; ls -al host_config.1 sanlun.1)
-r-xr-xr-x   1 root     sys        12266 Sep 13  2017 host_config.1
-r-xr-xr-x   1 root     sys         9044 Sep 13  2017 sanlun.1
```

8. After you finish, configure the host parameters for your environment using the `/opt/NTAP/SANToolkit/bin/host_config` command:
   - MPxIO
   - Veritas DMP
9. Verify the installation:

```
sanlun version
```

SAN Toolkit

Solaris Host Utilities is a NetApp host software that provides a command line toolkit on your Oracle Solaris host. The toolkit is installed when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns the LUN information.

```
#sanlun lun show all
controller(7mode)/ device host lun
vserver(Cmode)                     lun-pathname       filename
adapter protocol size mode
---------------------------------------------------------------------------
---------------------------------------------------------
data_vserver                     /vol/vol1/lun1
/dev/rdsk/c0t600A098038304437522B4E694E49792Dd0s2 qlc3   FCP       10g
cDOT
data_vserver                     /vol/vol0/lun2
/dev/rdsk/c0t600A098038304437522B4E694E497938d0s2 qlc3   FCP       10g
cDOT
data_vserver                     /vol/vol2/lun3
/dev/rdsk/c0t600A098038304437522B4E694E497939d0s2 qlc3   FCP       10g
cDOT
data_vserver                     /vol/vol3/lun4
/dev/rdsk/c0t600A098038304437522B4E694E497941d0s2 qlc3   FCP       10g
cDOT
```

This toolkit is common across all Host Utilities configurations and protocols. As a result, some toolkit content might apply to one configuration, but not another. Having unused components does not affect your system performance.

Solaris Host Utilities 6.2 command reference

You can use the Solaris Host Utilities 6.2 sample command reference for an end-to-end validation of the NetApp storage configuration using the host utilities tool.

List all host initiators mapped to host

You can retrieve a list of all host initiators mapped to a host.

```
# sanlun fcp show adapter -v
```

Example output

```
adapter name:      qlc3
WWPN:              21000024ff17a301
WWNN:              20000024ff17a301
driver name:       qlc
model:             7335902
model description: 7115462, Oracle Storage Dual-Port 32 Gb Fibre Channel
PCIe HBA
serial number:     463916R+1720333838
hardware version:  Not Available
driver version:    210226-5.10
firmware version:  8.08.04
Number of ports:   1 of 2
port type:         Fabric
port state:        Operational
supported speed:   8 GBit/sec, 16 GBit/sec, 32 GBit/sec
negotiated speed:  32 GBit/sec
OS device name:    /dev/cfg/c7

adapter name:      qlc2
WWPN:              21000024ff17a300
WWNN:              20000024ff17a300
driver name:       qlc
model:             7335902
model description: 7115462, Oracle Storage Dual-Port 32 Gb Fibre Channel
PCIe HBA
serial number:     463916R+1720333838
hardware version:  Not Available
driver version:    210226-5.10
firmware version:  8.08.04
Number of ports:   2 of 2
port type:         Fabric
port state:        Operational
supported speed:   8 GBit/sec, 16 GBit/sec, 32 GBit/sec
negotiated speed:  16 GBit/sec
OS device name:    /dev/cfg/c6
```

List all LUNs mapped to host

You can retrieve a list of all LUNs mapped to a host.

```
# sanlun lun show -p -v all
```

Example output

```
                    ONTAP Path: data_vserver:/vol1/lun1
                           LUN: 1
                      LUN Size: 10g
                   Host Device:
/dev/rdsk/c0t600A0980383044485A3F4E694E4F775Ad0s2
                          Mode: C
            Multipath Provider: Sun Microsystems
              Multipath Policy: Native
```

List all LUNs mapped to host from a given SVM/ List all attributes of a given LUN mapped to host

You can retrieve a list of all LUNs mapped to a host from a certain SVM.

```
# sanlun lun show -p -v sanboot_unix`
```

Example output

```
ONTAP Path: sanboot_unix:/vol/sol_boot/sanboot_lun
                           LUN: 0
                      LUN Size: 180.0g
```

List ONTAP LUN attributes by host device filename

You can retrieve a list of all ONTAP LUN attributes by specifying a host device filename.

```
# sanlun lun show all
```

Example output

```
controller(7mode/E-Series)/                                         device
vserver(cDOT/FlashRay)       lun-pathname
filename
--------------------------------------------------------------------------
-------------------------------------
sanboot_unix                 /vol/sol_193_boot/chatsol_193_sanboot
/dev/rdsk/c0t600A098038304437522B4E694E4A3043d0s2

host adapter    protocol lun size   product
---------------------------------------------
qlc3            FCP      180.0g     cDOT
```

Windows Unified Host Utilities

Install Windows Unified Host Utilities 7.2

The Windows Unified Host Utilities (WUHU) enable you to connect a Windows host computer to NetApp storage systems.

The Windows Unified Host Utilities supports the following versions of Windows:

- Windows 2022
- Windows 2019
- Windows 2016
- Windows 2012R2
- Windows 2012

Windows Unified Host Utilities includes an installation program that sets the required Windows registry and Host Bus Adapter (HBA) parameters so that a Windows host can correctly handle the storage system behaviors for NetApp ONTAP and E-Series platforms.

When you install the Host Utilities software, the installer sets the required Windows registry and HBA parameters.

The following programs and files are installed on the Windows host computer. The default directory is `C:\Program Files\NetApp\Windows Host Utilities`.

Program	Purpose
`emulexhba.reg`	Troubleshooting program; run this program only if instructed to do so by technical support personnel.
`\NetAppQCLI\fcconfig.exe`	Used by the installation program to set the HBA parameters.
`\NetAppQCLI\fcconfig.ini`	Used by the installation program to set the HBA parameters.
`\NetAppQCLI*.*`	Used by the installation program to set the QLogic FC HBA parameters.
`san_version.exe`	Displays the version of the Host Utilities and FC HBAs.

The host utilities support different Windows host configurations, protocols, and multipathing options. For more information, see the NetApp Interoperability Matrix Tool.

Verify your host and storage system configuration

Before you install the host utilities, you must verify that the host utilities version supports your host and storage system configuration so that the software installs correctly.

Steps

1. Check the supported configuration in the NetApp Interoperability Matrix Tool.

2. Check the hotfixes required for the respective host in the SAN host Windows documentation.

   > The Using Windows server 2022 with ONTAP document provides instructions on installing Windows hotfixes for Windows server 2022. Refer to the Windows documents in the host configurations category to find the relevant hotfix information for earlier versions of Windows server.

3. Add the iSCSI, FCP, or NVMe-oF license and start the target service.

   > The FC and iSCSI protocols do not require licenses on E-Series storage systems using the SANtricity Storage Manager.

4. Verify your cabling.

   See the SAN configuration reference documentation for your version of ONTAP or E-Series hardware cabling for detailed cabling and configuration information.

Configure FC HBAs and switches

Install and configure one or more supported FC host bus adapters (HBAs) for FC connections to the storage system.

The Windows Host Utilities installer sets the required FC HBA settings.

> Do not change the HBA settings manually.

Steps

1. Install one or more supported FC HBAs according to the instructions provided by the HBA vendor.
2. Obtain the supported HBA drivers and management utilities and install them according to the instructions provided by the HBA vendor.
3. Connect the HBAs to your FC switches or directly to the storage system.
4. Create zones on the FC switch according to your FC switch documentation.
5. For ONTAP, zone the switch by the WWPN. Be sure to use the WWPN of the logical interfaces (LIFs) and not the WWPN of the physical ports on the storage controllers. See the SAN configuration reference documentation for more information.

Install the Host Utilities

The installation program installs the host utilities package and sets the Windows registry and HBA settings.

You must specify whether to include multipathing support when you install the Windows Unified Host Utilities software package. The installer prompts you for the following options. You can also run a quiet (unattended) installation from a Windows command prompt.

Multipathing support

- Choose `MPIO` if you have more than one path from the Windows host or virtual machine to the storage system.
- Choose `no MPIO` only if you are using a single path to the storage system.

The MPIO selection is not available for Windows XP and Windows Vista systems; multipath I/O is not

supported on these guest operating systems.
For Hyper-V guests, raw (pass-through) disks do not appear in the guest OS if you choose multipathing support. You can either use raw disks, or you can use MPIO, but you cannot use both in the guest OS.

You can install the host utilities interactively or using the command line. The new Host Utilities installation package must be in a path that is accessible by the Windows host. Follow the instructions for installing the Host Utilities interactively or from the Windows command line.

Install interactively

To install the Host Utilities software package interactively, you must run the host utilities installation program and follow the prompts.

Steps

1. Download the executable file from the NetApp Support Site.
2. Change to the directory where you downloaded the executable file.
3. Run the `netapp_windows_host_utilities_7.2_x64` file and follow the instructions on the screen.
4. Reboot the Windows host when prompted.

Install from a command line

You can perform a quiet (unattended) installation of the host utilities by entering the appropriate commands at the Windows command prompt. The system automatically reboots when the installation is complete.

Steps

1. Enter the following command at the Windows command prompt:

```
msiexec /i installer.msi /quiet MULTIPATHING= {0 | 1}
[INSTALLDIR=inst_path]
```

   - `installer` is the name of the `.msi` file for your CPU architecture.
   - MULTIPATHING specifies whether MPIO support is installed. The allowed values are "0" for no and "1" for yes.
   - `inst_path` is the path where the host utilities files are installed. The default path is `C:\Program Files\NetApp\Windows Host Utilities\`.

To see the standard Microsoft Installer (MSI) options for logging and other functions, enter `msiexec /help` at the Windows command prompt. For example, the `msiexec /i install.msi /quiet /l*v <install.log> LOGVERBOSE=1` command displays logging information.

Install Windows Unified Host Utilities 7.1

The Windows Unified Host Utilities (WUHU) enable you to connect a Windows host computer to NetApp storage systems.

The Windows Unified Host Utilities supports the following versions of Windows:

- Windows 2022
- Windows 2019
- Windows 2016
- Windows 2012R2
- Windows 2012

Windows Unified Host Utilities includes an installation program that sets the required Windows registry and Host Bus Adapter (HBA) parameters so that a Windows host can correctly handle the storage system behaviors for NetApp ONTAP and E-Series platforms.

When you install the Host Utilities software, the installer sets the required Windows registry and Host Bus Adapter (HBA) parameters.

The following programs and files are installed on the Windows host computer. The default directory is `C:\Program Files\NetApp\Windows Host Utilities`.

Program	Purpose
`emulexhba.reg`	Troubleshooting program; run this program only if instructed to do so by technical support personnel.
`\NetAppQCLI\fcconfig.exe`	Used by the installation program to set HBA parameters.
`\NetAppQCLI\fcconfig.ini`	Used by the installation program to set HBA parameters.
`\NetAppQCLI*.*`	Used by the installation program to set QLogic FC HBA parameters.
`san_version.exe`	Displays the version of the Host Utilities and FC HBAs.

The Host Utilities support different Windows host configurations, protocols, and multipathing options. See the NetApp Interoperability Matrix Tool for the most current list of supported configurations.

Verify your host and storage system configurations

Before you install the Host Utilities, you must verify that the Host Utilities version supports your host and storage system configuration so that the software installs correctly.

Steps

1. Check the supported configuration in the NetApp Interoperability Matrix Tool.
2. Check the hotfixes required for the respective host on the SAN host Windows documentation.

   > ⓘ The Using Windows server 2022 with ONTAP document provides instructions on installing Windows hotfixes for Windows server 2022. Refer to the Windows documents in the host configurations category to find the relevant hotfix information for earlier versions of Windows server.

3. Add the iSCSI or FCP license and start the target service.

The FC and iSCSI protocols do not require licenses on E-Series storage systems using the SANtricity Storage Manager.

4. Verify your cabling

   See the SAN configuration reference documentation for your version of ONTAP or E-Series hardware cabling for detailed cabling and configuration information.

Configure FC HBAs and switches

Install and configure one or more supported FC host bus adapters (HBAs) for FC connections to a storage system.

The Windows Host Utilities installer sets the required FC HBA settings.

Do not change the HBA settings manually.

Steps

1. Install one or more supported FC HBAs according to the instructions provided by the HBA vendor.
2. Obtain the supported HBA drivers and management utilities and install them according to the instructions provided by the HBA vendor.
3. Connect the HBAs to your FC switches or directly to the storage system.
4. Create zones on the FC switch according to your FC switch documentation.
5. For ONTAP, zone the switch by worldwide port name (WWPN). Be sure to use the WWPN of the LIFs and not of the physical ports on the storage controllers. See the SAN configuration reference documentation for more information.

Install the Host Utilities

The installation program installs the Host Utilities package and sets the Windows registry and HBA settings.

You must specify whether to include multipathing support when you install the Windows Unified Host Utilities software package. The installer prompts you for the following choice. You can also run a quiet (unattended) installation from a Windows command prompt.

Multipathing support

- Choose `MPIO` if you have more than one path from the Windows host or virtual machine to the storage system.
- Choose `no MPIO` only if you are using a single path to the storage system.

The MPIO selection is not available for Windows XP and Windows Vista systems; multipath I/O is not supported on these guest operating systems.
For Hyper-V guests, raw (pass-through) disks do not appear in the guest OS if you choose multipathing support. You can either use raw disks, or you can use MPIO, but you cannot use both in the guest OS.

You can install the host utilities interactively or using the command line. The new Host Utilities installation package must be in a path that is accessible by the Windows host. Follow the instructions for installing the Host Utilities interactively or from the Windows command line.

Install interactively

Steps

To install the Host Utilities software package interactively, you must run the Host Utilities installation program and follow the prompts.

Steps

1. Download the executable file from the NetApp Support Site.
2. Change to the directory from which you downloaded the executable file.
3. Run the `netapp_windows_host_utilities_7.1_x64` file and follow the instructions on the screen.
4. Reboot the Windows host when prompted.

Install from a command line

You can perform a quiet (unattended) installation of the Host Utilities by entering the appropriate commands at a Windows command prompt. The system automatically reboots when the installation is complete.

Steps

1. Enter the following command at a Windows command prompt:

```
msiexec /i installer.msi /quiet MULTIPATHING= {0 | 1}
[INSTALLDIR=inst_path]
```

   - `installer` is the name of the `.msi` file for your CPU architecture
   - MULTIPATHING specifies whether MPIO support is installed. Allowed values are "0" for no, "1" for yes
   - `inst_path` is the path where the Host Utilities files are installed. The default path is `C:\Program Files\NetApp\Windows Host Utilities\`.

> To see the standard Microsoft Installer (MSI) options for logging and other functions, enter `msiexec /help` at a Windows command prompt. For example, the`msiexec /i install.msi /quiet /l*v <install.log> LOGVERBOSE=1` command displays logging information.

Upgrade the Windows Unified Host Utilities

The new Host Utilities installation package must be in a path that is accessible by the Windows host. Follow the instructions for installing the Host Utilities interactively or from the Windows command line in order to upgrade the installation package.

Upgrade interactively

To upgrade the Host Utilities software package interactively, you must run the Host Utilities installation program and follow the prompts.

Steps

1. Change to the directory where you downloaded the executable file.
2. Run the executable file and follow the instructions on the screen.
3. Reboot the Windows host when prompted.
4. After the reboot completes, check the host utility version:
   a. Open **Control Panel**.
   b. Go to **Program and features** and check the host utility version.

Upgrade from a command line

You can perform a quiet (unattended) upgrade of the new Host Utilities by entering the appropriate commands at the Windows command prompt.

Steps

1. Enter the following command at the Windows command prompt:

```
msiexec /i installer.msi /quiet MULTIPATHING= {0 | 1}
[INSTALLDIR=inst_path]
```

   - `installer` is the name of the `.msi` file for your CPU architecture.
   - MULTIPATHING specifies whether MPIO support is installed. The allowed values are "0" for no and "1" for yes.
   - `inst_path` is the path where the Host Utilities files are installed. The default path is `C:\Program Files\NetApp\Windows Host Utilities\`.

> To see the standard Microsoft Installer (MSI) options for logging and other functions, enter `msiexec /help` at the Windows command prompt. For example, the `msiexec /i install.msi /quiet /l*v <install.log> LOGVERBOSE=1` command displays logging information.

The system automatically reboots when the installation is complete.

Repair and remove the Windows Unified Host Utilities

You can use the **Repair** option of the Host Utilities installation program to update the Host bus adapter (HBA) and Windows registry settings. You can also remove the Host Utilities entirely, either interactively or from the Windows command line.

Repair or remove interactively

The **Repair** option updates the Windows registry and FC HBAs with the required settings. You can also remove the Host Utilities entirely.

Steps

1. Open Windows **Programs and Features** (Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, and Windows 2022).
2. Select **NetApp Windows Unified Host Utilities**.
3. Select **Change**.
4. Select **Repair** or **Remove**, as needed.
5. Follow the instructions on the screen.

Repair or remove from command line

The **Repair** option updates the Windows registry and FC HBAs with the required settings. You can also remove the Host Utilities entirely from a Windows command line.

Steps

1. Enter the following command on the Windows command line to repair Windows Host Utilities:

   ```
   msiexec /f installer.msi [/quiet]
   ```

   - `/f` repairs the installation.
   - `installer.msi` is the name of the Windows Host Utilities installation program on your system.
   - `/quiet` suppresses all feedback and reboots the system automatically without prompting when the command completes.

Configure registry settings

The Host Utilities require certain registry and parameter settings to verify that the Windows host correctly handles the storage system behavior.

Windows Host Utilities set the parameters that affect how the Windows host responds to a delay or loss of data. The particular values have been selected to verify that the Windows host correctly handles events such as the failover of one controller in the storage system to its partner controller.

Not all values apply to the device-specific module (DSM) for SANtricity Storage Manager; however, any overlap of values set by the Host Utilities and those set by the DSM for SANtricity Storage Manager do not result in conflicts.

FC, NVMe/FC, and iSCSI HBAs also have parameters that you must set to ensure best performance and to successfully handle storage system events.

The installation program supplied with Windows Unified Host Utilities sets the Windows, FC, and NVMe/FC HBA parameters to the supported values.

You must manually set the iSCSI HBA parameters.

The installer sets different values depending on whether you specify multipath I/O (MPIO) support when running the installation program.

You should not change these values unless NetApp technical support directs you to do so.

Registry values set by Windows Unified Host Utilities 7.2

The Windows Unified Host Utilities installer automatically sets registry values that are based on the choices that you make during the installation. You should be aware of these registry values and the operating system version.

The following values are set by the Windows Unified Host Utilities installer. All values are in decimal unless otherwise stated.

HKLM is the abbreviation for `HKEY_LOCAL_MACHINE`.

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters \DsmMaximumRetryTimeDuringStateTransition	120	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016 and Windows 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet \Services\msdsm\Parameters \DsmMaximumStateTransitionTime	120	When MPIO support is specified and your server is Windows Server 2012 R2, or Windows Server 2016, Windows 2019 or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters\DsmSupportedDeviceList	"NETAPP LUN", "NETAPP LUN C-Mode" "NVMe NetApp ONTAO Con"	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Control \Class\ {iSCSI_driver_GUID}\ instance_ID \Parameters \IPSecConfigTimeout	60	Always
HKLM\SYSTEM\CurrentControlSet\Control \Class\ {iSCSI_driver_GUID}\ instance_ID \Parameters \LinkDownTime	10	Always
HKLM\SYSTEM\CurrentControlSet\Services \ClusDisk \Parameters\ManageDisksOnSystemBuses	1	Always

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Control \Class\ {iSCSI_driver_GUID}\ instance_ID \Parameters \MaxRequestHoldTime	120	When no MPIO support is selected
HKLM\SYSTEM\CurrentControlSet\Control \Class\ {iSCSI_driver_GUID}\ instance_ID \Parameters \MaxRequestHoldTime	30	Always
HKLM\SYSTEM\CurrentControlSet\Control \MPDEV\MPIOSupportedDeviceList	"NETAPP LUN", "NETAPP LUN C-Mode", "NVMe NetApp ONTAO Con"	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\PathRecoveryInterval	30	When your server is Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\PathVerifyEnabled	1	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters\PathVerifyEnabled	1	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services \vnetapp\Parameters\PathVerifyEnabled	0	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services \mpio\Parameters\PDORemovePeriod	130	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services\msdsm \Parameters\PDORemovePeriod	130	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016 , Windows Server 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services\vnetapp \Parameters\PDORemovePeriod	130	When MPIO support is specified, except if Data ONTAP DSM is detected

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\RetryCount	6	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services\msdsm\Parameters\RetryCount	6	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\RetryInterval	1	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services\msdsm\Parameters\RetryInterval	1	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Windows Server 2022
HKLM\SYSTEM\CurrentControlSet\Services\vnetapp\Parameters\RetryInterval	1	When MPIO support is specified
HKLM\SYSTEM\CurrentControlSet\Services\disk\TimeOutValue	120	When no MPIO support is selected
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\UseCustomPathRecoveryInterval	1	When MPIO support is specified and your server is Windows Server 2012 R2, Windows Server 2016, Windows Server 2019, or Windows Server 2022

NVMe parameters

The following NVMe Emulex driver parameters are updated when installing Windows Unified Host Utilities 7.2:

- EnableNVMe = 1
- NVMEMode = 0
- LimTransferSize=1

Registry values set by Windows Unified Host Utilities 7.1

The Windows Unified Host Utilities installer automatically sets registry values that are based on the choices that you make during installation. You should be aware of these registry values, the operating system version.

The following values are set by the Windows Unified Host Utilities installer. All values are in decimal unless otherwise noted.

`HKLM` is the abbreviation for `HKEY_LOCAL_MACHINE`.

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters \DsmMaximumRetryTimeDuringStateTransition	120	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters \DsmMaximumStateTransitionTime	120	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\msdsm \Parameters\DsmSupportedDeviceList	"NETAPP LUN"	When MPIO support is specified
	"NETAPP LUN", "NETAPP LUN C-Mode"	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Control\Class {iSCSI_driver_GUID}\ instance_ID\Parameters \IPSecConfigTimeout	60	Always, except when Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Control \Class{iSCSI_driver_GUID} \ instance_ID\Parameters\LinkDownTime	10	Always
HKLM\SYSTEM\CurrentControlSet\Services\ClusDisk \Parameters\ManageDisksOnSystemBuses	1	Always, except when Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Control \Class{iSCSI_driver_GUID} \ instance_ID\Parameters\MaxRequestHoldTime	120	When no MPIO support is selected
	30	Always, except when Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet \Control\MPDEV\MPIOSupportedDeviceList	"NETAPP LUN"	When MPIO support is specified
	"NETAPP LUN", "NETAPP LUN C-Mode"	When MPIO is support-specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\PathRecoveryInterval	40	When your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 only
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\PathVerifyEnabled	0	When MPIO support is specified, except if Data ONTAP DSM is detected

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Services\msdsm \Parameters\PathVerifyEnabled	0	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters\PathVerifyEnabled	0	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \msiscdsm\Parameters\PathVerifyEnabled	0	When MPIO support is specified and your server is Windows Server 2003, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\vnetapp \Parameters\PathVerifyEnabled	0	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\PDORemovePeriod	130	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\msdsm \Parameters\PDORemovePeriod	130	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\msiscdsm \Parameters\PDORemovePeriod	130	When MPIO support is specified and your server is Windows Server 2003, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \vnetapp \Parameters\PDORemovePeriod	130	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \mpio\Parameters\RetryCount	6	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\msdsm \Parameters\RetryCount	6	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \msiscdsm\Parameters\RetryCount	6	When MPIO support is specified and your server is Windows Server 2003, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \vnetapp\Parameters\RetryCount	6	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \mpio\Parameters\RetryInterval	1	When MPIO support is specified, except if Data ONTAP DSM is detected

Registry key	Value	When set
HKLM\SYSTEM\CurrentControlSet\Services \msdsm\Parameters\RetryInterval	1	When MPIO support is specified and your server is Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services \vnetapp\Parameters\RetryInterval	1	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet \Services\disk\TimeOutValue	120	When no MPIO support is selected, except if Data ONTAP DSM is detected
	60	When MPIO support is specified, except if Data ONTAP DSM is detected
HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\UseCustomPathRecoveryInterval	1	When your server is Windows Server 2008,Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 only

See the Microsoft documents for the registry parameter details.

FC HBA values set by Windows Host Utilities

On systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

For Emulex FC HBAs, the installer sets the following parameters:

When MPIO is selected

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

When MPIO is not selected

Property type	Property value
LinkTimeOut	30
NodeTimeOut	120

For QLogic Fibre Channel HBAs, the installer sets the following parameters:

When MPIO is selected

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

When MPIO is not selected

Property type	Property value
LinkDownTimeOut	30
PortDownRetryCount	120

> The names of the parameters might vary slightly depending on the program.
> For example, in the QLogic QConvergeConsole program, the parameter is displayed as `Link Down Timeout`.
> The Host Utilities `fcconfig.ini` file displays this parameter as either `LinkDownTimeOut` or `MpioLinkDownTimeOut`, depending on whether MPIO is specified. However, all of these names refer to the same HBA parameter. See Emulex or QLogic to learn more about the timeout parameters.

Understand the Host Utilities changes to FC HBA driver settings

During the installation of the required Emulex or QLogic HBA drivers on an FC system, several parameters are checked and, in some cases, modified.

The Host Utilities set values for the following parameters if MS DSM for Windows MPIO is detected:

- LinkTimeOut – defines the length of time in seconds that the host port waits before resuming I/O after a physical link is down.
- NodeTimeOut – defines the length of time in seconds before the host port recognizes that a connection to the target device is down.

When troubleshooting HBA issues, check to make sure these settings have the correct values. The correct values depend on two factors:

- The HBA vendor
- Whether you are using multipathing software (MPIO)

You can correct the HBA settings by running the Repair option of the Windows Host Utilities installer.

Emulex HBA drivers

If you have an FC system, you must verify the Emulex HBA driver settings. These settings must exist for each port on the HBA.

Steps

1. Open OnCommand Manager.
2. Select the appropriate HBA from the list and click the **Driver Parameters** tab.

   The driver parameters appear.

   a. If you are using MPIO software, ensure that you have the following driver settings:
      - LinkTimeOut - 1
      - NodeTimeOut - 10

   b. If you are not using MPIO software, ensure that you have the following driver settings:
      - LinkTimeOut - 30
      - NodeTimeOut - 120

QLogic HBA drivers

On FC systems, you must verify the QLogic HBA driver settings. These settings must exist for each port on the HBA.

Steps

1. Open QConvergeConsole, and then click **Connect** on the toolbar.

   The **Connect to Host** dialog box appears.

2. Select the appropriate host from the list, and then select **Connect**.

   A list of HBAs appears in the FC HBA pane.

3. Select the appropriate HBA port from the list, and then select the **Settings** tab.
4. Select **Advanced HBA Port Settings** from the **Select Settings** section.
5. If you are using MPIO software, verify that you have the following driver settings:
   - Link Down Timeout (linkdwnto) - 1
   - Port Down Retry Count (portdwnrc) - 10
6. If you are not using MPIO software, verify that you have the following driver settings:
   - Link Down Timeout (linkdwnto) - 30
   - Port Down Retry Count (portdwnrc) - 120

Troubleshoot

You can use the general troubleshooting techniques for Windows Host Utilities. Be sure to check the latest Release Notes for known issues and solutions.

The following is a list of the different areas you can investigate for potential interoperability issues:

- To identify potential interoperability issues, confirm that the Host Utilities support your combination of host operating system software, host hardware, ONTAP software, and storage system hardware. See the NetApp Interoperability Matrix Tool for more information.
- Verify that you have the correct iSCSI configuration.
- If iSCSI LUNs are not available after a reboot, verify that the target is listed as persistent on the **Persistent Targets** tab of the Microsoft iSCSI initiator GUI.
- If applications using the LUNs display errors on startup, verify that the applications are configured to depend on the iSCSI service.
- For FC paths to storage controllers running ONTAP, verify that the FC switches are zoned using the WWPNs of the target LIFs, not the WWPNs of the physical ports on the node.
- Review the Release Notes for Windows Host Utilities to check for known issues. The Release Notes include a list of known issues and limitations.
- Review the troubleshooting information in the SAN Administration Guide for your version of ONTAP.
- Search NetApp Bugs Online for recently discovered issues.
    - In the Bug Type field under Advanced Search, select **iSCSI - Windows** and then select **Go**. You should repeat the search for Bug Type **FCP -Windows**.
- Collect information about your system.
- Record any error messages that are displayed on the host or storage system console.
- Collect the host and storage system log files.
- Record the symptoms of the problem and any changes made to the host or storage system just before the problem appeared.
- If you are unable to resolve the problem, contact NetApp technical support for assistance.

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, and then you verify the multipath settings for the NetApp ONTAP LUNs.

AIX and PowerVM/VIOS

Use IBM AIX 7.2 and/or PowerVM (VIOS 3.1) with ONTAP

You can use the ONTAP SAN host configuration settings to configure IBM AIX 7.2 and/or PowerVM (VIOS 3.1) with ONTAP as the target.

Install the AIX/VIOS host utilities

You must install the AIX Host Utilities Kit while using AIX MPIO with NetApp ONTAP storage.

You can download the compressed file containing the Host Utilities software packages from the NetApp Support Site. After you have the file, you must decompress it to get the two software packages you need to install the host utilities.

NetApp AIX Host Utilities 6.1 is the latest release. This release addresses the memory leak issue that was reported in the previous releases. Refer to release notes section for additional information.

Steps

1. Login to your host.
   - On an AIX host, log in as **root**.
   - On a PowerVM host, log in as **padmin**, and then enter the `oem_setup_env` command to become root.
2. Download a copy of the compressed file containing the Host Utilities from NetApp Support Site to a directory on your host.
3. Go to the directory containing the download.
4. Uncompress the file and extract the SAN Toolkit software package.

   ```
   tar -xvf ntap_aix_host_utilities_6.1.tar.gz
   ```

   The following directory is created when you decompress the file: `ntap_aix_host_utilities_6.1`. This directory will have one of the following subdirectories: MPIO, NON_MPIO, or SAN_Tool_Kit.
5. Install the AIX MPIO:

   ```
   installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/MPIO
   NetApp.MPIO_Host_Utilities_Kit
   ```

6. Install the SAN Toolkit:

   ```
   installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/SAN_Tool_Kit
   NetApp.SAN_toolkit
   ```

7. Reboot the host.

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
#sanlun lun show

controller(7mode)/                    device     host       lun
vserver(Cmode)     lun-pathname      filename   adapter    protocol   size
mode
---------------------------------------------------------------------------
------
data_vserver       /vol/vol1/lun1    hdisk0     fcs0       FCP        60g
C
data_vserver       /vol/vol2/lun2    hdisk0     fcs0       FCP        20g
C
data_vserver       /vol/vol3/lun3    hdisk11    fcs0       FCP        20g
C
data_vserver       /vol/vol4/lun4    hdisk14    fcs0       FCP        20g
C
```

SAN booting

What you'll need

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

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

Multipathing

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

For AIX, Path Control Module (PCM) is responsible for controlling multiple paths. PCM is a storage vendor supplied code that handles path management. This gets installed and enabled as part of the Host Utilities installation.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

> All SAN Arrays (ASA) configurations are supported beginning in ONTAP 9.8 for AIX Hosts.

```
# 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
```

Recommended Settings

Following are some recommended parameter settings for ONTAP LUNs. The critical parameters for ONTAP LUNs are set automatically after installing the NetApp Host Utilities Kit.

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
rw_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

Parameter	Environment	Value for AIX	Note
num_cmd_elems	MPIO / non-MPIO	1024 for AIX	FC EN0G

Recommended settings for MetroCluster

By default, the AIX operating system enforces a shorter I/O timeout when no paths to a LUN are available. This might occur in configurations including single-switch SAN fabric and MetroCluster configurations that experience unplanned failovers. For additional information and recommended changes to default settings, refer to NetApp KB1001318

AIX support with SM-BC

Beginning with ONTAP 9.11.1, AIX is supported with SM-BC. With an AIX configuration, the primary cluster is the "active" cluster.

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

To configure AIX for SM-BC, refer to the Knowledge Base article How to configure an AIX host for SnapMirror Business Continuity (SM-BC).

Known issues

The IBM AIX 7.2 and/or PowerVM (VIOS 3.1) with ONTAP 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

NetApp Bug ID	Title	Description	Partner ID
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	Seeing I/O issues on VIOS 3.1.1.10 during SFO faults and straight I/O	On VIOS 3.1.1 IO failures may be seen on NPIV client disk which are backed by 16/32Gb FC adapters. Also, a `vfchost` driver may get into a state where it stops processing I/O requests from the client. Applying IBM APAR IJ22290 IBM APAR IJ23222 will fix the issue.	NA

Use IBM AIX 7.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure IBM AIX 7.1 with ONTAP as the target.

Install the AIX Host Utilities

You must install the AIX Host Utilities Kit while using AIX MPIO with NetApp ONTAP storage.

You can download the compressed file containing the Host Utilities software packages from the NetApp Support Site. After you have the file, you must extract it to get the two software packages you need to install the Host Utilities.

Steps

1. Login to your host.
   - On an AIX host, log in as **root**.
2. Download a copy of the compressed file containing the Host Utilities from NetApp Support Site to a directory on your host.
3. Go to the directory containing the download.
4. Decompress the file and extract the SAN Toolkit software package.

```
tar -xvf ntap_aix_host_utilities_6.1.tar.tgz
```

The following directory is created when you decompress the file: `ntap_aix_host_utilities_6.1`. This directory will have one of the following subdirectories: MPIO, NON_MPIO, or SAN_Tool_Kit.

5. Install the AIX MPIO:

```
installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/MPIO
NetApp.MPIO_Host_Utilities_Kit
```

6. Install the SAN Toolkit:
```
installp -aXYd /var/tmp/ntap_aix_host_utilities_6.1/SAN_Tool_Kit
NetApp.SAN_toolkit
```
7. Reboot the host.

SAN Toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
#sanlun lun show

controller(7mode)/                       device      host        lun
vserver(Cmode)      lun-pathname        filename    adapter     protocol    size
mode
--------------------------------------------------------------------------
------
data_vserver        /vol/vol1/lun1      hdisk0      fcs0        FCP         60g
C
data_vserver        /vol/vol2/lun2      hdisk0      fcs0        FCP         20g
C
data_vserver        /vol/vol3/lun3      hdisk11     fcs0        FCP         20g
C
data_vserver        /vol/vol4/lun4      hdisk14     fcs0        FCP         20g
C
```

SAN Booting

What you'll need

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

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an AIX host. You can set up a SAN boot LUN to work in an AIX MPIO environment that is running the AIX Host Utilities with either the FC or FCoE protocol. The method you use for creating a SAN boot LUN and installing a new OS image on it in an AIX MPIO environment can vary, depending on which protocol you are using.

Multipathing

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

For AIX, Path Control Module (PCM) is responsible for controlling multiple paths. PCM is a storage vendor-supplied code that handles path management. This gets installed and enabled as part of the Host Utilities installation.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

All SAN Arrays (ASA) configurations are supported beginning in ONTAP 9.8 for AIX Hosts.

```
# 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
```

Recommended Settings

Following are some recommended parameter settings for ONTAP LUNs. The critical parameters for ONTAP LUNs are set automatically after installing the NetApp Host Utilities Kit.

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

Parameter	Environment	Value for AIX	Note
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

Recommended Settings for MetroCluster

By default, the AIX operating system enforces a shorter I/O timeout when no paths to a LUN are available. This might occur in configurations including single-switch SAN fabric and MetroCluster configurations that experience unplanned failovers. For additional information and recommended changes to default settings, refer to NetApp KB1001318

AIX support with SM-BC

Beginning with ONTAP 9.11.1, AIX is supported with SM-BC. With an AIX configuration, the primary cluster is the "active" cluster.

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

To configure AIX for SM-BC, refer to the Knowledge Base article How to configure an AIX host for SnapMirror Business Continuity (SM-BC).

Known issues

There are no known issues.

CentOS

Release notes

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 does not 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.

CentOS 8

Use CentOS 8.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 8.5 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 8.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> ⓘ Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

CentOS 8.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no know issues for the CentOS 8.5 with ONTAP release.

Use CentOS 8.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                 lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 8.4 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 8.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

CentOS 8.4 is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the CentOS 8.4 with ONTAP release.

Use CentOS 8.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host               lun
vserver(cDOT/FlashRay)     lun-pathname     filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1     /dev/sdb   host16    FCP
120.0g   cDOT
data_vserver          /vol/vol1/lun1     /dev/sdc   host15    FCP
120.0g   cDOT
data_vserver          /vol/vol2/lun2     /dev/sdd   host16    FCP
120.0g   cDOT
data_vserver          /vol/vol2/lun2     /dev/sde   host15    FCP
120.0g   cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 8.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 8.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 8.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.3.

Use CentOS 8.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do

not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                lun
vserver(cDOT/FlashRay)     lun-pathname    filename   adapter    protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp

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 is successful.

Multipathing

For CentOS 8.2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 8.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 8.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.2.

Use CentOS 8.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 8.1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 8.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
| |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 8.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.1.

Use CentOS 8.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 8.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do

not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp

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 is successful.

Multipathing

For CentOS 8.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 8.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> ⓘ Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 8.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Parameter	Setting
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.0.

CentOS 7

Use CentOS 7.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host                lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.9 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.9 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.8 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.8 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.7 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.7 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

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

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.6 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname filename   adapter  protocol size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383134765 7244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.5 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> ⓘ You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the

LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313476572 44e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The CentOS 7.4 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use CentOS 7.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the CentOS 7.3 with ONTAP release.

Use CentOS 7.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the CentOS 7.2 with ONTAP release.

Use CentOS 7.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do

not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the CentOS 7.1 with ONTAP release.

Use CentOS 7.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 7.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                lun
vserver(cDOT/FlashRay)      lun-pathname     filename   adapter   protocol  size
Product
----------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 7.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. CentOS 7.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 7.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the CentOS 7.0 with ONTAP release.

CentOS 6

Use CentOS 6.10 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.10 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.10 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.10 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.10 is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.10.

Use CentOS 6.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.9 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.9.

Use CentOS 6.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host             lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.8 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.8 is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

touch /etc/multipath.conf

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.8.

Use CentOS 6.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.7 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.7.

Use CentOS 6.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                 lun
vserver(cDOT/FlashRay)     lun-pathname   filename   adapter   protocol   size
Product
---------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.6 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs. To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.6.

Use CentOS 6.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                 device      host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename    adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver           /vol/vol1/lun1   /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver           /vol/vol1/lun1   /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2   /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2   /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.5 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.5.

Use CentOS 6.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure CentOS 6.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device      host                lun
vserver(cDOT/FlashRay)     lun-pathname   filename    adapter   protocol   size
Product
--------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For CentOS 6.4 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CentOS 6.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   CentOS 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The CentOS 6.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

For CentOS (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.4.

Citrix

Use Citrix Hypervisor with ONTAP

You can configure ONTAP SAN host configuration settings for Citrix Hypervisor 8 series

OS releases with FC, FCoE and iSCSi protocols.

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Citrix Hypervisor (CH) 8.x the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. CH 8.x is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
You can use the `/sbin/mpathutil` status command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA personas.

All SAN Array (ASA) Configuration

For All SAN Array (ASA) configuration there should be one group of paths with single priorities. All the paths are Active/Optimized, which means they are serviced by the controller and I/O is sent on all the active paths.

Example

The following example displays the correct output for an ONTAP LUN with four Active/Optimized paths:

```
# mpathutil status
3600a098038303344674244a357579386a dm-13 NETAPP  ,LUN C-Mode
size=30G 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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Non-ASA Configuration

For non-ASA configuration there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# mpathutil status
3600a09803830344674244a357579386a dm-13 NETAPP  ,LUN C-Mode
size=30G 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
 |- 1:0:0:11   sde   8:64    active ready running
`- 12:0:8:11  sdua  66:544  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
   |- 1:0:9:11   sddo  71:96   active ready running
   `- 12:0:26:11 sdyt  129:720 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Citrix Hypervisor 8.x OS is compiled with all settings required to recognize and correctly manage ONTAP LUNs. For Citrix Hypervisor 8.x, an empty zerobyte `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file.

Enable the host multipath service from the **Xencenter Management Portal** and verify that the multipath service is enabled and running.

```
# systemct1 status multipathd
multipathd.service - Device-Mapper Multipath Device Controller
  Loaded:  load (/usr/lib/systemd/system/multipathd.service; enabled;
vendor preset: enabled)
 Drop-In:  /etc/systemd/system/multipathd.service.d
            slice.config
  Active:  active (running) since Fri YYYY-MM-DD 00:00:26 IST; 1 month 9
days ago
Main PID:  3789 (multipathd)
  CGroup:  /control.slice/multipathd.service
            3789  /sbin/multipathd
```

There is no requirement to append content to the `/etc/multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults. You can add the following syntax to the multipath.conf file to exclude the unwanted devices.

```
# cat /etc/multipath.conf
blacklist {
     wwid    <DevId>
     devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
     devnode "^hd[a-z]"
     devnode "^cciss.*"
}
```

Replace the **<DevID>** with the WWID string of the device you want to exclude.

Example

In this example for Citrix Hypervisor 8.x, `sda` is the local SCSI disk that we need to add to the blacklist.

1. Run the following command to determine the WWID:

```
# lib/udev/scsi_id -gud /dev/sda
3600a098038303458772450714535317a
```

2. Add this WWID to the blacklist stanza in the `/etc/multipath.conf`:

```
#cat /etc/multipath.conf
blacklist {
  wwid    3600a098038303458772450714535317a
  devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9*]"
  devnode "^hd[a-z]"
  devnode "^cciss.*"
}
```

Refer to the multipath parameter runtime configuration by using the `$multipathd show config` command. You should always check your running configuration for legacy settings that might be overriding default settings, especially in the defaults section.

The following table shows the critical **multipathd** parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they need to be corrected by later stanzas in **multipath.conf** that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. The following defaults should only be overridden in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

Parameter	Setting
`detect_prio`	yes
`dev_loss_tmo`	"infinity"

Parameter	Setting
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"
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

Example

The following example illustrates how to correct an overridden default. In this case, the **multipath.conf** file defines values for **path_checker** and **detect_prio** that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
# cat /etc/multipath.conf
defaults {
  path_checker readsector0
  detect_prio no
}
devices{
        device{
                vendor "NETAPP "
                product "LUN.*"
                path_checker tur
                detect_prio yes
        }
}
```

Citrix Hypervisor recommends use of Citrix VM tools for all Linux and Windows based guest VMs for a supported configuration.

Known issues

The Citrix Hypervisor with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Citrix Tracker ID
1242343	Kernel disruption on Citrix Hypervisor 8.0 with QLogic QLE2742 32GB FC during storage failover operations	Kernel disruption might occur during storage failover operations on Citrix Hypervisor 8.0 kernel (4.19.0+1) with QLogic QLE2742 32GB HBA. This issue prompts a reboot of the operating system and causes application disruption. If kdump is configured, the kernel disruption generates a vmcore file under the /var/crash/ directory. You can use the vmcore file to understand the cause of the failure. After the kernel disruption, you can recover the operating system by rebooting the host operating system and restarting the application.	NETAPP-98

Use Citrix XenServer with ONTAP

You can configure ONTAP SAN host configuration settings for Citrix XenServer 7 series OS releases with FC, FCoE, and iSCSi protocols.

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

Multipath support in Citrix XenServer is based on the Device Mapper Multipathd components. Device mapper nodes are not automatically created for all LUNs presented to the XenServer and are only provisioned when LUNs are actively used by the Storage Management Layer (API). Citrix XenServer Storage Manager API plugin handles activating and deactivating multipath nodes automatically.

Due to incompatibilities with the Integrated Multipath Management architecture, Citrix recommends that you use the Citrix XenCenter application for managing the storage configuration. If it is necessary to query the status of Device Mapper tables manually, or list active device mapper multipath nodes on the system, you can use the `/sbin/mpathutil status` command to verify the settings for your ONTAP LUNs. For more information refer to the standard vendor documentation for Citrix XenServer.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# mpathutil status
show topology
3600a098038303458772450714535317a dm-0 NETAPP  , LUN C-Mode
size=80G 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
| |- 2:0:2:0   sdc   8:32    active ready running
| |- 12:0:5:0  sdn   8:208   active ready running
| |- 2:0:6:0   sdg   8:96    active ready running
| `- 12:0:0:0  sdi   8:128   active ready running
|-+- policy='service-time 0' prio=10 status=enabled
  |- 2:0:0:0   sda   8:0     active ready running
  |- 2:0:1:0   sdb   8:16    active ready running
  |- 12:0:3:0  sd1   8:176   active ready running
  `-12:0:6:0   sdo   8:224   active ready running
[root@sanhost ~]#
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

All SAN Array Configuration

In All SAN Array (ASA) configurations, all paths to a given Logical Unit (LUN) are active and optimized. This means I/O can be served through all paths at the same time, thereby enabling better performance.

Example

The following example displays the correct output for an ONTAP LUN with all four Active/Optimized paths:

```
# mpathutil status
show topology
3600a098038303458772450714535317a dm-0 NETAPP  , LUN C-Mode
size=80G 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
| |- 2:0:2:0   sdc   8:32    active ready running
| |- 12:0:5:0  sdn   8:208   active ready running
| |- 2:0:6:0   sdg   8:96    active ready running
| `- 12:0:0:0  sdi   8:128   active ready running
[root@sanhost ~]#
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Citrix XenServer 7.x OS is compiled with all settings required to recognize and correctly manage ONTAP LUNs. For Citrix XenServer 7.x, an empty zerobyte `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file.

Enable the host multipath service from the **Xencenter Management Portal** and verify that the multipath service is enabled and running.

```
# systemct1 status multipathd
multipathd.service - Device-Mapper Multipath Device Controller
  Loaded:  load (/usr/lib/systemd/system/multipathd.service; enabled;
vendor preset: enabled)
 Drop-In:  /etc/systemd/system/multipathd.service.d
             slice.config
  Active:  active (running) since Fri YYYY-MM-DD 00:00:26 IST; 1 month 9
days ago
Main PID:  3789 (multipathd)
  CGroup:  /control.slice/multipathd.service
            3789  /sbin/multipathd
```

There is no requirement to append content to the `/etc/multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults. You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices.

```
# cat /etc/multipath.conf
blacklist {
     wwid    <DevId>
     devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
     devnode "^hd[a-z]"
     devnode "^cciss.*"
}
```

Replace the **<DevID>** with the WWID string of the device you want to exclude.

Example

In this example for Citrix XenServer 7.x, `sda` is the local SCSI disk that we need to add to the blacklist.

1. Run the following command to determine the WWID:

```
# lib/udev/scsi_id -gud /dev/sda
3600a098038303458772450714535317a
```

2. Add this WWID to the blacklist stanza in the `/etc/multipath.conf`:

```
#cat /etc/multipath.conf
blacklist {
  wwid    3600a098038303458772450714535317a
  devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9*]"
  devnode "^hd[a-z]"
  devnode "^cciss.*"
}
```

Refer to the multipath parameter runtime configuration by using the `$multipathd show config` command. You should always check your running configuration for legacy settings that might be overriding default settings, especially in the defaults section.

The following table shows the critical **multipathd** parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they need to be corrected by later stanzas in **multipath.conf** that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. The following defaults should only be overridden in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

Parameter	Setting
`detect_prio`	yes
`dev_loss_tmo`	"infinity"
`failback`	immediate

Parameter	Setting
fast_io_fail_tmo	5
features	"3 queue_if_no_path pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
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

Example

The following example illustrates how to correct an overridden default. In this case, the **multipath.conf** file defines values for **path_checker** and **detect_prio** that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
# cat /etc/multipath.conf
defaults {
  path_checker readsector0
  detect_prio no
}
devices{
        device{
                vendor "NETAPP "
                product "LUN.*"
                path_checker tur
                detect_prio yes
        }
}
```

Citrix XenServer recommends use of Citrix VM tools for all Linux and Windows based guest VMs for a supported configuration.

Known issues

There are no known issues for the Citrix XenServer with ONTAP release.

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

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is 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 NetApp ONTAP storage, `VMW_SATP_ALUA` plugin is used by default with `VMW_PSP_RR` as a path selection policy (PSP). You can confirm by using the following command:

```
`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 the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are 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.600a0980383148693724545244395855
```

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.600a098038314869372454524439585 5
   Runtime Name: vmhba4:C0:T3:L11
   Device: naa.600a0980383148693724545244395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
   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.600a0980383148693724545244395855
   Runtime Name: vmhba3:C0:T3:L11
   Device: naa.600a0980383148693724545244395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
   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.
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
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 user interface, 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 NetApp 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 ONTAP tools for VMware vSphere documentation, TR-4597, and TR-4400.

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.600a098038304731783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
   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

- TR-4597-VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror Business Continuity (SM-BC) 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

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is 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 NetApp ONTAP storage, VMW_SATP_ALUA plugin is used by default with VMW_PSP_RR as a path selection policy (PSP). This can be confirmed by using the below command.

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 the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are 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.
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

`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 user interface, 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 NetApp 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, please refer to ONTAP tools for VMware vSphere documentation as well as TR-4597-VMware vSphere with ONTAP and TR-4400.

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.

Settings	Default	ONTAP Recommended	Description
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.
```

Validating the Disk IOPs setting

Use the following command to verify the IOPs setting.

esxcli storage nmp device list -d naa.600a098038304731783f506670553355

```
naa.600a098038304731783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
   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

- TR-4597-VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror Business Continuity (SM-BC) 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

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is 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 NetApp ONTAP storage, VMW_SATP_ALUA plugin is used by default with `VMW_PSP_RR` as a path selection policy (PSP). This can be confirmed by using the below command:

```
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
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

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 the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are 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 user interface, 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 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 does not recommend changing any vVols settings from default at this time.

Refer to the NetApp 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, please refer to ONTAP tools for VMware vSphere documentation as well as TR-4597 and TR-4400.

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.
```

Validating the Disk IOPs setting

Use the following command to verify the IOPs setting:

esxcli storage nmp device list -d naa.600a098038304731783f506670553355

```
naa.600a098038304731783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304731783f506670553355)
   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

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

- TR-4597-VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror Business Continuity (SM-BC) with VMware vSphere Metro Storage Cluster (vMSC)

HP-UX

Use HP-UX 11i v3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure HP-UX 11i v3 with ONTAP as the target.

Install the HP-UX Host Utilities

You can download the compressed file containing the Host Utilities software packages from the NetApp Support Site. After you have the file, you must uncompress it to get the software packages you need to install the Host Utilities.

Steps

1. Download a copy of the compressed file containing the Host Utilities from the NetApp Support Site to a directory on your host.
2. Go to the directory containing the download.
3. Uncompress the file.

   ```
   gunzip netapp_hpux_host_utilities_6.0_ia_pa.depot.gz
   ```

4. Enter the following command to install the software:

   ```
   swinstall -s /netapp_hpux_host_utilities_6.0_ia_pa.depot NetApp_santoolkit
   ```

5. Reboot the host.

SAN Toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
#sanlun lun show

controller(7mode)/                                device
host      lun
vserver(Cmode)     lun-pathname                  filename
adapter protocol  size  mode
--------------------------------------------------------------------------
------
sanboot_unix      /vol/hpux_215_boot_en_0/goot_hpux_215_lun
/dev/rdsk/c11t0d0 fcd0   FCP      150g   C
sanboot_unix      /vol/hpux_215_boot_en_0/goot_hpux_215_lun
/dev/rdsk/c24t0d0 fcd1   FCP      150g   C
sanboot_unix      /vol/hpux_215_boot_en_0/goot_hpux_215_lun
/dev/rdsk/c21t0d0 fcd1   FCP      150g   C
sanboot_unix      /vol/hpux_215_boot_en_0/goot_hpux_215_lun
/dev/rdsk/c12t0d0 fcd0   FCP      150g   C
```

SAN Booting

What you'll need

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

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.

Multipathing

Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues on the remaining paths. For a host to have multiple paths to a LUN, multipathing must be enabled. The HP-UX Host Utilities support different multipathing solutions based on your configuration. The following is for the Native Multipathing solution.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

All SAN Array (ASA) configurations are supported beginning in ONTAP 9.8 for HP-UX 11iv3

```
# 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
```

Recommended Settings

Following are some recommended parameter settings for HPUX 11i v3 and NetApp ONTAP LUNs. NetApp uses the default settings for HP-UX.

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

Known issues

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

NetApp Bug ID	Title	Description	Partner ID
1447287	AUFO event on the isolated master cluster in SM-BC 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 Business Continuity (SM-BC) 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 1 MB 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 1 MB. Refer HP-UX vendor documentation for details.	NA

Oracle Linux

Release notes

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 does not 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.

OL 9

Use Oracle Linux 9.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 9.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                 lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver           /vol/vol1/lun1    /dev/sdb    host16    FCP        80.0g
cDOT
data_vserver           /vol/vol1/lun1    /dev/sdc    host15    FCP        80.0g
cDOT
data_vserver           /vol/vol2/lun2    /dev/sdd    host16    FCP        80.0g
cDOT
data_vserver           /vol/vol2/lun2    /dev/sde    host15    FCP        80.0g
cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For OL 9.2 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. OL 9.2 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA configurations.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 9.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example demonstrates how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 9.2 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 9.2.

KVM settings

You can also use the recommended settings to configure Kernel-based Virtual Machine (KVM). There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The Oracle Linux 9.2 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1508554	SAN LUN utility with Emulex HBA needs symbolic links from library packages	When you execute the Linux Unified 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`	Not Applicable

Use Oracle Linux 9.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 9.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16    FCP        80.0g
cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15    FCP        80.0g
cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16    FCP        80.0g
cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15    FCP        80.0g
cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For OL 9.1 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. OL 9.1 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA configurations.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383036347ffb4d59646c4436 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 9.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example demonstrates how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 9.1 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 9.1.

KVM settings

You can also use the recommended settings to configure Kernel-based Virtual Machine (KVM). There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The Oracle Linux 9.1 with NetApp ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1508554	SAN LUN utility with Emulex HBA needs symbolic links from library packages	When you execute the Linux Unified 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 a 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`	Not Applicable

Use Oracle Linux 9.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 9.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
---------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP       80.0g
cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP       80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP       80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP       80.0g
cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 9.0 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. OL 9.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383036347ffb4d59646c4436 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 9.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

> To configure Oracle Linux 9.0 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 9.0.

KVM settings

You can also use the recommended settings to configure Kernel-based Virtual Machine (KVM). There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The Oracle Linux 9.0 with NetApp ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID

1508554	SAN LUN utility with Emulex HBA needs symbolic links from library packages	When you execute the Linux Unified 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 a 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`
			Not Applicable

OL 8

Use Oracle Linux 8.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                   device     host                 lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP        80.0g
cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP        80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP        80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP        80.0g
cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.8 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. OL 8.8 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath outputs for a LUN mapped to ASA and non-ASA configurations.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The OL 8.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example demonstrates how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

> To configure Oracle Linux 8.8 RedHat Enterprise Kernel, use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.8.

KVM settings

You can also use the recommended settings to configure a Kernel-based Virtual Machine (KVM). There are no changes required to configure a KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Oracle Linux 8.8 with ONTAP release.

Use Oracle Linux 8.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP       80.0g
cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP       80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP       80.0g
cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP       80.0g
cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.7, the `/etc/multipath.conf` file must exist. You do not need to make specific changes to the file because OL 8.7 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to All SAN Array (ASA) and non-ASA configurations.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. You should require no more than four paths. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383036347ffb4d59646c4436 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 8.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	uniform
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

ⓘ To configure Oracle Linux 8.7 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.7.

KVM settings

You can also use the recommended settings to configure the Kernel-based Virtual Machine (KVM). There are no changes required to configure the KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Oracle Linux 8.7 with ONTAP release.

Use Oracle Linux 8.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit

.rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host                lun
vserver(cDOT/FlashRay)     lun-pathname    filename   adapter   protocol  size
Product
-------------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.6 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. OL 8.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped non-ASA personas.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	infinity
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

> To configure Oracle Linux 8.6 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.6.

KVM Settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Oracle Linux 8.6 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues section in the corresponding Red Hat Enterprise Linux release documentation.

Use Oracle Linux 8.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 8.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped non-ASA personas.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 8.5 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.5.

KVM Settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Oracle Linux 8.5 with ONTAP release.

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues section in the corresponding Red Hat Enterprise Linux release documentation.

Use Oracle Linux 8.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.

2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 8.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped non-ASA personas.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	uniform
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 8.4 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.4.

KVM Settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Oracle Linux 8.4 with ONTAP release.

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues section in the corresponding Red Hat Enterprise Linux release documentation.

Use Oracle Linux 8.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 8.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped non-ASA personas.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
  |- 11:0:7:6   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
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 8.3 Red Hat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.3.

Known issues

There are no known issues for the Oracle Linux 8.3 with ONTAP release.

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues section in the corresponding Red Hat Enterprise Linux release documentation.

Use Oracle Linux 8.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 8.2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 8.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 8.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	uniform
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

> To configure Oracle Linux 8.2 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.2.

Known issues

There are no known issues for the Oracle Linux 8.2 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.2.

Use Oracle Linux 8.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 8.1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 8.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
| |- 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
```

> ⓘ Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

> To configure Oracle Linux 8.1 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.1.

Known issues

There are no known issues for the Oracle Linux 8.1 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.1.

Use Oracle Linux 8.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 8.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 8.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 8.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
| |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 8.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
 }
}
```

To configure Oracle Linux 8.0 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.0.

Known issues

There are no known issues for the Oracle Linux 8.0 with ONTAP release.

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 8.0.

OL 7

Use Oracle Linux 7.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                 lun
vserver(cDOT/FlashRay)     lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver           /vol/vol1/lun1    /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver           /vol/vol1/lun1    /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2    /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2    /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 7.9 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 7.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034587724507145354l5a dm-15 NETAPP  ,LUN C-Mode
size=40G 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
| |- 11:0:5:7   sdbg 67:160  active ready running
| `- 12:0:13:7  sdlg 67:480  active ready running
|-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:8:7   sdck 69:128  active ready running
  |- 11:0:12:7  sddy 128:0   active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Oracle Linux 7.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 7.9 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.9.

Known issues

The Oracle Linux 7.9 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host	When you set the `disable_changed_wwids` multipath configuration parameter to YES, it disables access to the path device in the event of a worldwide identifier (WWID) change. Multipath disables access to the path device until the WWID of the path is restored to the WWID of the multipath device. See the NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7 for more information.	Not applicable

Use Oracle Linux 7.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability

Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)       lun-pathname   filename   adapter    protocol    size
Product
-------------------------------------------------------------------------------------------
data_vserver            /vol/vol1/lun1     /dev/sdb    host16     FCP
120.0g   cDOT
data_vserver            /vol/vol1/lun1     /dev/sdc    host15     FCP
120.0g   cDOT
data_vserver            /vol/vol2/lun2     /dev/sdd    host16     FCP
120.0g   cDOT
data_vserver            /vol/vol2/lun2     /dev/sde    host15     FCP
120.0g   cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux (OL) 7.8 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. OL 7.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped non-ASA personas.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038303634722b4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=80G features='4 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.8 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.8.

Known issues

The Oracle Linux 7.8 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

NetApp Bug ID	Title	Description	Bugzilla ID
1311575	IO delays observed due to Read/Write operations failed to switch through secondary paths during storage failover with Qlogic QLE2672(16G)	I/O operations might fail to resume through secondary paths during storage failover operations on Oracle Linux 7.7 kernel (5.4.17-2011.0.7.el7uek.x86_6) with QLogic QLE2672 16G HBA. If I/O progress stops due to blocked primary paths during storage failover, the I/O operation might not resume through secondary paths causing an I/O delay. The I/O operation resumes only after primary paths come online after the completion of the storage failover giveback operation.	17171
1311576	IO delays observed due to Read/Write operation failing to switch through secondary paths during storage failover with Emulex LPe16002(16G)	I/O operations might fail to resume through secondary paths during storage failover operations on Oracle Linux 7.7 kernel (5.4.17-2011.0.7.el7uek.x86_6) with Emulex LPe16002 16G HBA. If I/O progress stops due to blocked primary paths during storage failover, the I/O operation might not resume through secondary paths causing an I/O delay. The I/O operation resumes only after primary paths come online after the completion of the storage failover giveback operation.	17172

NetApp Bug ID	Title	Description	Bugzilla ID
1246134	IO delays observed and reports are moving to blocked, NOT PRESENT state during storage failover with Emulex LPe16002(16G)	During storage failover operations on the Oracle Linux 7.6 with the UEK5U2 kernel running with an Emulex LPe16002B-M6 16G Fibre Channel (FC) host bus adapter (HBA), I/O progress might stop due to reports getting blocked. The storage failover operation reports change from "online" state to "blocked" state, causing a delay in read and write operations. After the operation is completed successfully, the reports fail to move back to "online" state and continue to remain in "blocked" state.	16852

NetApp Bug ID	Title	Description	Bugzilla ID
1246327	IO delays observed and Rports are moving to blocked, NOT PRESENT state during storage failover with Qlogic QLE2672(16G) and QLE2742(32G)	Fibre Channel (FC) remote ports might be blocked on Red Hat Enterprise Linux (RHEL) 7.6 with the QLogic QLE2672 16G host during storage failover operations. Because the logical interfaces go down when a storage node is down, the remote ports set the storage node status to blocked. IO progress might stop due to the blocked ports if you are running both a QLogic QLE2672 16G host and a QLE2742 32GB Fibre Channel (FC) host bus adapter (HBA). When the storage node returns to its optimal state, the logical interfaces also come up and the remote ports should be online. However, the remote ports might still be blocked. This blocked state registers as failed faulty to LUNS at the multipath layer. You can verify the state of the remote ports with the following command: # cat /sys/class/fc_remote_ports/rport-*/port_stat You should see the following output: Blocked Blocked Blocked Blocked Online Online	16853

Use Oracle Linux 7.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.7 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.7 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.7.

Known issues

The Oracle Linux 7.7 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A
1311575	IO delays observed due to Read/Write operations failed to switch through secondary paths during storage failover with Qlogic QLE2672(16G)	I/O operations might fail to resume through secondary paths during storage failover operations on Oracle Linux 7.7 kernel (5.4.17-2011.0.7.el7uek.x86_6) with QLogic QLE2672 16G HBA. If I/O progress stops due to blocked primary paths during storage failover, the I/O operation might not resume through secondary paths causing an I/O delay. The I/O operation resumes only after primary paths come online after the completion of the storage failover giveback operation.	17171

NetApp Bug ID	Title	Description	Bugzilla ID
1311576	IO delays observed due to Read/Write operation failing to switch through secondary paths during storage failover with Emulex LPe16002(16G)	I/O operations might fail to resume through secondary paths during storage failover operations on Oracle Linux 7.7 kernel (5.4.17-2011.0.7.el7uek.x86_6) with Emulex LPe16002 16G HBA. If I/O progress stops due to blocked primary paths during storage failover, the I/O operation might not resume through secondary paths causing an I/O delay. The I/O operation resumes only after primary paths come online after the completion of the storage failover giveback operation.	17172
1246134	IO delays observed and reports are moving to blocked, NOT PRESENT state during storage failover with Emulex LPe16002(16G)	During storage failover operations on the Oracle Linux 7.6 with the UEK5U2 kernel running with an Emulex LPe16002B-M6 16G Fibre Channel (FC) host bus adapter (HBA), I/O progress might stop due to reports getting blocked. The storage failover operation reports change from "online" state to "blocked" state, causing a delay in read and write operations. After the operation is completed successfully, the reports fail to move back to "online" state and continue to remain in "blocked" state.	16852

NetApp Bug ID	Title	Description	Bugzilla ID
1246327	IO delays observed and Rports are moving to blocked, NOT PRESENT state during storage failover with Qlogic QLE2672(16G) and QLE2742(32G)	Fibre Channel (FC) remote ports might be blocked on Red Hat Enterprise Linux (RHEL) 7.6 with the QLogic QLE2672 16G host during storage failover operations. Because the logical interfaces go down when a storage node is down, the remote ports set the storage node status to blocked. IO progress might stop due to the blocked ports if you are running both a QLogic QLE2672 16G host and a QLE2742 32GB Fibre Channel (FC) host bus adapter (HBA). When the storage node returns to its optimal state, the logical interfaces also come up and the remote ports should be online. However, the remote ports might still be blocked. This blocked state registers as failed faulty to LUNS at the multipath layer. You can verify the state of the remote ports with the following command: # cat /sys/class/fc_remote_ports/rport-*/port_stat You should see the following output: Blocked Blocked Blocked Blocked Online Online	16853

Use Oracle Linux 7.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                    device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.6 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 7.6 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.6.

Known issues

The Oracle Linux 7.6 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A
1202736	LUNs might not be available during host discovery due to "Not Present" state of remote ports on a OL7U6 host with QLogic QLE2742 adapter	During host discovery, the status of Fibre Channel (FC) remote ports on a OL7U6 host with a QLogic QLE2742 adapter might enter into "Not Present" state. Remote ports with a "Not Present" state might cause paths to LUNs to become unavailable. During storage failover, the path redundancy might be reduced and result in an I/O outage. You can check the remote port status by entering the following command: # cat /sys/class/fc_remote_ports/rport-*/port_state The following is an example of the output that is displayed: Online Online Not Present Online Online	16613

NetApp Bug ID	Title	Description	Bugzilla ID
1204078	Kernel disruption occurs on Oracle Linux 7.6 running with Qlogic(QLE2672) 16GB FC HBA during storage failover operations	During storage failover operations on the Oracle Linux 7.6 with a Qlogic QLE2672 Fibre Channel (FC) host bus adapter (HBA), a kernel disruption occurs due to a panic in the kernel. The kernel panic causes Oracle Linux 7.6 to reboot, which leads to an application disruption. If the kdump mechanism is enabled, the kernel panic generates a vmcore file located in the /var/crash/ directory. You can analyze the vmcore file to determine the cause of the panic. After the kernel disruption, you can reboot the host OS and recover the operating system, and then you can restart any applications as required.	16606
1204351	Kernel disruption might occur on Oracle Linux 7.6 running with Qlogic(QLE2742) 32GB FC HBA during storage failover operations	During storage failover operations on the Oracle Linux 7.6 with a Qlogic QLE2742 Fibre Channel (FC) host bus adapter (HBA), a kernel disruption might occur due to a panic in the kernel. The kernel panic causes Oracle Linux 7.6 to reboot, which leads to an application disruption. If the kdump mechanism is enabled, the kernel panic generates a vmcore file located in the /var/crash/ directory. You can analyze the vmcore file to determine the cause of the panic. After the kernel disruption, you can reboot the host OS and recover the operating system, and then you can restart any applications as required.	16605

NetApp Bug ID	Title	Description	Bugzilla ID
1204352	Kernel disruption might occur on Oracle Linux 7.6 running with Emulex (LPe32002-M2)32GB FC HBA during storage failover operations	During storage failover operations on the Oracle Linux 7.6 with an Emulex LPe32002-M2 Fibre Channel (FC) host bus adapter (HBA), a kernel disruption might occur due to a panic in the kernel. The kernel panic causes Oracle Linux 7.6 to reboot, which leads to an application disruption. If the kdump mechanism is enabled, the kernel panic generates a vmcore file located in the /var/crash/ directory. You can analyze the vmcore file to determine the cause of the panic. After the kernel disruption, you can reboot the host OS and recover the operating system, and then you can restart any applications as required.	16607
11246134	No I/O progress on Oracle Linux 7.6 with UEK5U2 kernel, running with an Emulex LPe16002B-M6 16G FC HBA during storage failover operations	During storage failover operations on the Oracle Linux 7.6 with the UEK5U2 kernel running with an Emulex LPe16002B-M6 16G Fibre Channel (FC) host bus adapter (HBA), I/O progress might stop due to reports getting blocked. The storage failover operation reports change from an "online" state to a "blocked" state, causing a delay in read and write operations. After the operation has completed successfully, the reports fail to move back to an "online" state and continue to remain in a "blocked" state.	16852

NetApp Bug ID	Title	Description	Bugzilla ID
1246327	Remote port status on QLogic QLE2672 16G host blocked during storage failover operations	Fibre Channel (FC) remote ports might be blocked on Red Hat Enterprise Linux (RHEL) 7.6 with the QLogic QLE2672 16G host during storage failover operations. Because the logical interfaces go down when a storage node is down, the remote ports set the storage node status to blocked. IO progress might stop due to the blocked ports if you are running both a QLogic QLE2672 16G host and a QLE2742 32GB Fibre Channel (FC) host bus adapter (HBA). When the storage node returns to its optimal state, the logical interfaces also come up and the remote ports should be online. However, the remote ports might still be blocked. This blocked state registers as failed faulty to LUNS at the multipath layer. You can verify the state of the remote ports with the following command: # cat /sys/class/fc_remote_ports/rport-*/port_stat You should see the following output: Blocked Blocked Blocked Blocked Online Online	16853

Use Oracle Linux 7.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   sda is the local SCSI disk that we need to add it to the blacklist.

2. Add the WWID to the blacklist stanza in /etc/multipath.conf:

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

You should always check your /etc/multipath.conf file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical multipathd parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the multipath.conf file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.5 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.5.

Known issues

The Oracle Linux 7.5 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A
1177239	Kernel disruption observed on OL7.5 with Qlogic QLE2672 16G FC during storage failover operations	During storage failover operations on Oracle Linux 7 (OL7.5) with kernel 4.1.12-112.16.4.el7uek.x86_64 and the Qlogic QLE2672 HBA, you might observe kernel disruption. This prompts a reboot of the operating system which causes an application disruption. If kdump is configured, the kernel disruption creates a vmcore file in the /var/crash/ directory. This disruption can be observed in the module “kmem_cache_alloc+118,” which is logged in the vmcore file and identified with the string "exception RIP: kmem_cache_alloc+118." After a kernel disruption, you can recover by rebooting the host operating system and restarting the application.	

Use Oracle Linux 7.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 7.4 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.4.

Known issues

The Oracle Linux 7.4 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A
1109468	Firmware dumps observed on an OL7.4 Hypervisor with QLE8362 card	During storage failover operations on an OL7.4 Hypervisor with QLE8362 card, the firmware dumps are observed occasionally. The firmware dumps might result in an I/O outage on the host, which might go up to 500 seconds. After the adapter completes the firmware dump, the I/O operation resumes in the normal manner. No further recovery procedure is required on the host. To indicate the firmware dump, the following message is displayed in the /var/log/message file: qla2xxx [0000:0c:00.3]-d001:8: Firmware dump saved to temp buffer (8/ffffc90008901000), dump status flags (0x3f)	16039

Use Oracle Linux 7.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.3 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.3.

Known issues

There are no known issues for the Oracle Linux 7.3 with ONTAP release.

Use Oracle Linux 7.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

Oracle Linux 7.2 supports Unbreakable Enterprise Kernel (UEK) R3 and UEK R4. The OS boots with UEK R3 kernel by default.

Oracle Linux 7.2 UEK R3 Configuration

For Oracle Linux 7.2 UEK R3, create an empty multipath.conf file. The settings for Oracle Linux 7.2 UEK with and without ALUA update automatically by default. To Enable ALUA Handler, perform the following steps:

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `dracut -f` command to recreate the initrd-image.
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

Oracle Linux 7.2 UEK R4 Configuration

For Oracle Linux 7.2 UEK R4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.2 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.2.

Known issues

There are no known issues for the Oracle Linux 7.2 with ONTAP release.

Use Oracle Linux 7.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

Oracle Linux 7.1 supports Unbreakable Enterprise Kernel (UEK) R3 and UEK R4. The OS boots with UEK R3 kernel by default.

Oracle Linux 7.1 UEK R3 Configuration

For Oracle Linux 7.1 UEK R3, create an empty multipath.conf file. The settings for Oracle Linux 7.1 UEK with and without ALUA update automatically by default. To Enable ALUA Handler, perform the following steps:

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `dracut -f` command to recreate the initrd-image.
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

Oracle Linux 7.1 UEK R4 Configuration

For Oracle Linux 7.1 UEK R4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 7.1 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.1.

Known issues

There are no known issues for the Oracle Linux 7.1 with ONTAP release.

Use Oracle Linux 7.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 7.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host              lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 7.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 7.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs. To Enable ALUA Handler, perform the following steps:

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Recreate the initrd-image with the `dracut -f` command.
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
|-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 7.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 7.0 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 7.0.

Known issues

The Oracle Linux 7.0 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
901558	OL7.0 : Host loses all paths to the lun and hangs due to "RSCN timeout" error on OL 7.0 UEK r3U5 Beta on Emulex 8G(LPe12002) host	You might observe that the Emulex 8G(LPe12002) host hangs and there is a high I/O outage during storage failover operations with I/O. You might observe paths not recovering, which is a result of the RSCN timeout, due to which the host loses all the paths and hangs. Probability of hitting this issue is high.	14898
901557	OL 7.0: High IO outage observed on QLogic 8G FC (QLE2562) SAN host during storage failover operations with IO	You might observe high IO outage on QLogic 8G FC (QLE2562) host during storage failover operations with IO. Aborts and Device resets manifests as IO outage on the host. Probability of hitting this IO outage is high.	14894
894766	OL7.0: Dracut fails to include scsi_dh_alua.ko module in initramfs on UEKR3U5 alpha	The scsi_dh_alua module might not load even after adding the parameter "rdloaddriver=scsi_dh_alua" in the kernel command line and creating Dracut. As a result, ALUA is not enabled for NetApp LUNs as recommended.	14860

NetApp Bug ID	Title	Description	Bugzilla ID
894796	Anaconda displays an iSCSI login failure message although logins are successful during OL 7.0 OS installation	When you are installing OL 7.0, the anaconda installation screen displays that iSCSI login to multiple target IPs have failed though the iSCSI logins are successful. Anaconda displays following error message: "Node Login Failed" You will observe this error only when you select multiple target IPs for iSCSI login. You can continue the OS installation by clicking the "ok" button. This bug does not hamper either the iSCSI or the OL 7.0 OS installation.	14870
894771	OL7.0 : Anaconda does not add bootdev argument in kernel cmd line to set IP address for iSCSI SANboot OS install	Anaconda does not add a bootdev argument in the kernel command line where you set the IPv4 address during the OL 7.0 OS installation on an iSCSI multipath'd LUN. Owing to this, you cannot assign IP addresses to any of the Ethernet interfaces that were configured to establish iSCSI sessions with the storage subsystem during the OL 7.0 boot. Since iSCSI sessions are not established, the root LUN is not discovered when the OS boots and hence the OS boot fails.	14871
916501	Qlogic 10G FCoE (QLE8152) host kernel crash observed during storage failover operations with IO	You may observe a kernel crash in Qlogic driver module on 10G FCoE Qlogic (QLE8152) host. The crash occurs during storage failover operations with IO. Probability of hitting this crash is high which leads to longer IO outage on the host.	15019

OL 6

Use Oracle Linux 6.10 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.10 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.10 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.10 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
root=/dev/mapper/vg_ibmx3550m421096-lv_root
rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.10 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 6.10 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.10.

Known issues

There are no known issues for the Oracle Linux 6.10 with ONTAP release.

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.10.

Use Oracle Linux 6.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                   device      host                  lun
vserver(cDOT/FlashRay)     lun-pathname       filename    adapter   protocol    size
Product
-------------------------------------------------------------------------
data_vserver              /vol/vol1/lun1     /dev/sdb    host16    FCP
120.0g   cDOT
data_vserver              /vol/vol1/lun1     /dev/sdc    host15    FCP
120.0g   cDOT
data_vserver              /vol/vol2/lun2     /dev/sdd    host16    FCP
120.0g   cDOT
data_vserver              /vol/vol2/lun2     /dev/sde    host15    FCP
120.0g   cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.9 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

To configure Oracle Linux 6.9 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.9.

Known issues

The Oracle Linux 6.9 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1082780	Firmware dumps are observed occasionally on OL6.9 hypervisor with the QLE8362 card	During storage failover operations on OL6.9 hypervisor with QLE8362 card, the firmware dumps are observed occasionally. The firmware dumps might result in an I/O outage on the host which might go up to a thousand seconds. After the adapter completes the firmware dump, the I/O operation resumes in the normal manner. No further recovery procedure is required on the host. To indicate the firmware dump, the following message is displayed in the /var/log/message file: qla2xxx [0000:0c:00.3]-d001:3: Firmware dump saved to temp buffer (3/ffffc90008901000), dump status flags (0x3f).	16039

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.9.

Use Oracle Linux 6.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.8 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs. To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
root=/dev/mapper/vg_ibmx3550m421096-lv_root
rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 6.8 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.8.

Known issues

There are no known issues for the Oracle Linux 6.8 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.8.

Use Oracle Linux 6.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.7 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 6.7 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.7.

Known issues

There are no known issues for the Oracle Linux 6.7 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.7.

Use Oracle Linux 6.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.6 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 6.6 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.6.

Known issues

There are no known issues for the Oracle Linux 6.6 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.6.

Use Oracle Linux 6.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 6.5 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.5.

Known issues

There are no known issues for the Oracle Linux 6.5 with ONTAP release.

> For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.5.

Use Oracle Linux 6.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Oracle Linux 6.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Oracle Linux 6.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. Oracle Linux 6.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-3.8.13-68.1.2.el6uek.x86_64 ro
   root=/dev/mapper/vg_ibmx3550m421096-lv_root
   rd_NO_LUKSrd_LVM_LV=vg_ibmx3550m421096/lv_root LANG=en_US.UTF-8
   rd_NO_MDSYSFONT=latarcyrheb-sun16 crashkernel=256M KEYBOARDTYPE=pc
   KEYTABLE=us rd_LVM_LV=vg_ibmx3550m421096/lv_swap rd_NO_DM rhgb quiet
   rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   Oracle 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.
   You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
   There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, which means they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 0:0:26:37 sdje 8:384   active ready running
| |- 0:0:25:37 sdik 135:64  active ready running
|-+- policy='round-robin 0' prio=10 status=enabled
  |- 0:0:18:37 sdda 70:128  active ready running
  |- 0:0:19:37 sddu 71:192  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The Oracle Linux 6.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

```
# chkconfig multipathd on
# /etc/init.d/multipathd start
```

- There is no requirement to add anything directly to the `multipath.conf` file unless you have devices that you do not want multipath to manage or you have existing settings that override defaults.
- You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices:
  - Replace the <DevId> with the WWID string of the device you want to exclude:

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

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

2. Add this WWID to the "blacklist" stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they must be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. You should only override these defaults in consultation with NetApp and/or the OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs.
If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

> To configure Oracle Linux 6.4 RedHat Enterprise Kernel (RHCK), use the recommended settings for Red Hat Enterprise Linux (RHEL) 6.4.

Known issues

The Oracle Linux 6.4 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
713555	QLogic adapter resets are seen on OL6.4 and OL5.9 with UEK2 on controller faults such as takeover/giveback, and reboot	QLogic adapter resets are seen on OL6.4 hosts with UEK2 (kernel-uek-2.6.39-400.17.1.el6uek) or OL5.9 hosts with UEK2 (kernel-uek-2.6.39 400.17.1.el5uek) when controller faults happen (such as takeover, giveback, and reboots). These resets are intermittent. When these adapter resets happen, a prolonged I/O outage (sometimes, more than 10 minutes) might occur until the adapter resets succeed and the paths' status are updated by dm-multipath. In /var/log/messages, messages similar to the following are seen when this bug is hit: kernel: qla2xxx [0000:11:00.0]-8018:0: ADAPTER RESET ISSUED nexus=0:2:13. This is observed with the kernel version: On OL6.4: kernel-uek-2.6.39-400.17.1.el6uek On OL5.9: kernel-uek-2.6.39-400.17.1.el5uek	13999

NetApp Bug ID	Title	Description	Bugzilla ID
715217	Delay in path recovery on OL6.4 or OL5.9 hosts with UEK2 may result in delayed I/O resumption on controller or fabric faults	When a controller fault (storage failover or giveback, reboots and so on) or a fabric fault (FC port disable or enable) occurs with I/O on Oracle Linux 6.4 or Oracle Linux 5.9 hosts with UEK2 Kernel, the path recovery by DM-Multipath takes a long time (4mins. to 10 mins). Sometimes, during the paths recovering to active state, the following lpfc driver errors are also seen: kernel: sd 0:0:8:3: [sdlt] Result: hostbyte=DID_ERROR driverbyte=DRIVER_OK Due to this delay in path recovery during fault events, the I/O resumption also delays. OL 6.4 Versions: device-mapper-1.02.77-9.el6 device-mapper-multipath-0.4.9-64.0.1.el6 kernel-uek-2.6.39-400.17.1.el6uek OL 5.9 Versions: device-mapper-1.02.77-9.el5 device-mapper-multipath-0.4.9-64.0.1.el5 kernel-uek-2.6.39-400.17.1.el5uek	14001

NetApp Bug ID	Title	Description	Bugzilla ID
709911	DM Multipath on OL6.4 & OL5.9 iSCSI with UEK2 kernel takes long time to update LUN path status after storage faults	On systems running Oracle Linux 6 Update4 and Oracle Linux 5 Update9 iSCSI with Unbreakable Enterprise Kernel Release 2 (UEK2), a problem has been seen during storage fault events where DM Multipath (DMMP) takes around 15 minutes to update the path status of Device Mapper (DM) devices (LUNs). If you run the "multipath -ll" command during this interval, the path status is shown as "failed ready running" for that DM device (LUN). The path status is eventually updated as "active ready running." This issue is seen with following version: Oracle Linux 6 Update 4: UEK2 Kernel: 2.6.39-400.17.1.el6uek.x86_64 Multipath: device-mapper-multipath-0.4.9-64.0.1.el6.x86_64 iSCSI: iscsi-initiator-utils-6.2.0.873-2.0.1.el6.x86_64 Oracle Linux 5 Update 9: UEK2 Kernel: 2.6.39-400.17.1.el5uek Multipath: device-mapper-multipath-0.4.9-64.0.1.el5.x86_64 iSCSI: iscsi-initiator-utils-6.2.0.872-16.0.1.el5.x86_64	13984

NetApp Bug ID	Title	Description	Bugzilla ID
739909	The SG_IO ioctl system call fails on dm-multipath devices after an FC fault on OL6.x and OL5.x hosts with UEK2	A problem is seen on Oracle Linux 6.x hosts with UEK2 kernel and Oracle Linux 5.x hosts with UEK2 kernel. The sg_* commands on a multipath device fail with EAGAIN error code (errno) after a fabric fault that makes all the paths in the active path group go down. This problem is seen only when there is no I/O occurring to the multipath devices. The following is an example: # sg_inq -v /dev/mapper/3600a098041764937303f436c75324370 inquiry cdb: 12 00 00 00 24 00 ioctl(SG_IO v3) failed with os_err (errno) = 11 inquiry: pass through os error: Resource temporarily unavailable HDIO_GET_IDENTITY ioctl failed: Resource temporarily unavailable [11] Both SCSI INQUIRY and fetching ATA information failed on /dev/mapper/3600a098041764937303f436c75324370 # This problem occurs because the path group switchover to other active groups is not activated during ioctl() calls when no I/O is occurring on the DM-Multipath device. The problem has been observed on the following versions of the kernel-uek and device-mapper-multipath packages: OL6.4 versions:	14082

For Oracle Linux (Red Hat compatible kernel) known issues, see the known issues for Red Hat Enterprise Linux (RHEL) 6.4.

RHEL

Release notes

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 does not 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.

RHEL 9

Use Red Hat Enterprise Linux 9.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 9.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

SAN Tool Kit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
---------------------------------------------------------------------------
----------
vs_147_32glpe            /vol/vol1/lun   /dev/sdb   Host11   FCP        10g
cDOT
vs_147_32glpe            /vol/vol1/lun   /dev/sdx   Host11   FCP        10g
cDOT
vs_147_32glpe            /vol/vol2/lun   /dev/sdbt  host12  FCP        10g
cDOT
vs_147_32glpe            /vol/vol2/lun   /dev/sdax  host12   FCP        10g
cDOT
```

SAN booting

What you'll need

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

Multipathing

For Red Hat Enterprise Linux (RHEL) 9.3, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 9.3 is compiled with all the settings that are required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038314359725d516c69733471 dm-22 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
  |- 11:0:3:0  sdau 66:224 active ready running
  |- 12:0:4:0  sdco 69:192 active ready running
  |- 12:0:0:0  sdav 66:240 active ready running
  `- 11:0:2:0  sdat 66:208 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831497832245 44d334a644d dm-10 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
| |- 12:0:0:18 sdbj 67:208 active ready running
| `- 11:0:1:18 sdan 66:112 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:18 sdt  65:48  active ready running
  `- 12:0:3:18 sdcf 69:48  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 9.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

* There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
* To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be

corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 9.3 with ONTAP release has the following known issues:

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

Use Red Hat Enterprise Linux 9.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 9.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

SAN Tool Kit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 9.2, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 9.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 9.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 9.2 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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.	Not Applicable
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.	2173947

Use Red Hat Enterprise Linux 9.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 9.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

Install the Linux Unified Host Utilities is strongly recommended by NetApp, but not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it and use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:
   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> ⓘ You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 9.1 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 9.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 9.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 9.1 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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.	N/A

Use Red Hat Enterprise Linux 9.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 9.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 9.0 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 9.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 9.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 9.0 with ONTAP release.

RHEL 8

Use Red Hat Enterprise Linux 8.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux (RHEL) 8.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN Tool Kit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
---------------------------------------------------------------------------
-------------------
vs_161_32gLpe          /vol/vol19/lun    /dev/sdcd  host15     FCP
10g      cDOT
vs_161_32gLpe          /vol/vol20/lun    /dev/sdce  host15     FCP
10g      cDOT
vs_161_32gLpe          /vol/vol18/lun    /dev/sdcc  host15     FCP
10g      cDOT
vs_161_32gLpe          /vol/vol17/lun    /dev/sdcb  host15     FCP
10g      cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For RHEL 8.9, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.9 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038314778375d53694b536e53 dm-16 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
  |- 14:0:0:0   sda  8:0    active ready running
  |- 15:0:8:0   sdcf 69:48  active ready running
  |- 15:0:0:0   sdaq 66:160 active ready running
  `- 14:0:9:0   sdv  65:80  active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038314837352453694b542f4a 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 8.9 OS recognizes ONTAP LUNs and automatically sets all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for RHEL 8.9.

Use Red Hat Enterprise Linux 8.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN Tool Kit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and host bus adapters (HBAs). The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
vs_163_32gQlc          /vol/vol1/lun1   /dev/sdb   host14   FCP
10.0g  cDOT
vs_163_32gQlc          /vol/vol1/lun1   /dev/sdc   host15   FCP
10.0g  cDOT
vs_163_32gQlc          /vol/vol2/lun2   /dev/sdd   host14   FCP
10.0g  cDOT
vs_163_32gQlc          /vol/vol2/lun2   /dev/sde   host15   FCP
10.0g  cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.8, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.8 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G    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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383148373524536 94b542f4a 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 8.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.8 with ONTAP release.

Use Red Hat Enterprise Linux 8.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.7 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.7 with ONTAP release.

Use Red Hat Enterprise Linux 8.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.6 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.6 with ONTAP release.

Use Red Hat Enterprise Linux 8.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.5 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> ⓘ Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.5 with ONTAP release.

Use Red Hat Enterprise Linux 8.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.4 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.4 with ONTAP release.

Use Red Hat Enterprise Linux 8.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you’ll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 8.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.3 with ONTAP release.

Use Red Hat Enterprise Linux 8.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.2 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. RHEL 8.2 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The RHEL 8.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 8.2 with ONTAP release.

Use Red Hat Enterprise Linux 8.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 8.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
| |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 8.1 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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	1760819

NetApp Bug ID	Title	Description	Bugzilla ID
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.	1744082
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.	1758504

Use Red Hat Enterprise Linux 8.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 8.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 8.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 8.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 8.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they

will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 8.0 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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.	1710009
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.	1690356
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.	1709995

NetApp Bug ID	Title	Description	Bugzilla ID
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.	1709995
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	1702005
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.	1705573

RHEL 7

Use Red Hat Enterprise Linux 7.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.9 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

touch /etc/multipath.conf

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.9 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use Red Hat Enterprise Linux 7.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.8 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.8 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use Red Hat Enterprise Linux 7.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   `rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64`

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host                 lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.7 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

touch /etc/multipath.conf

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.7 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

NetApp Bug ID	Title	Description	Bugzilla ID
1258856	Remote ports transit to a blocked state on RHEL7U7 with Emulex LPe16002 16GB FC during storage failover operations	Remote ports might transit to a blocked state on a RHEL 7.7 host with a LPe16002 16GB FC adapter 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.	1743667
1261474	Remote ports transit to blocked state on RHEL7U7 with Emulex LPe32002 32GB FC	Remote ports might transit to a blocked state on a RHEL 7.7 host with LPe32002 32GB FC adapter 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.	1745995

Use Red Hat Enterprise Linux 7.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

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

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.6 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

NetApp Bug ID	Title	Description	Bugzilla ID
1186754	Remote ports status on RHEL7U6 with QLogic QLE2742 host might be in blocked during host discovery	During host discovery, FC remote port status on RHEL7U6 host with a QLogic QLE2742 adapter might enter a blocked state. These blocked remote ports might result in the paths to LUNs becoming unavailable. During storage failover, the path redundancy might be reduced and result in I/O outage. You can check the remote port status by entering the following command:# cat /sys/class/fc_remote_ports/rport-*/port_state	1628039
1190698	Remote port status on RHEL7U6 with QLogic QLE2672 host might be in blocked during storage failover operations	FC remote ports might be blocked on Red Hat Enterprise Linux (RHEL) 7U6 with the QLogic QLE2672 host during storage failover operations. Because the logical interfaces go down when a storage node is down, the remote ports set the storage node status to blocked. When the storage node returns to its optimal state, the logical interfaces also come up and the remote ports should be online. However, the remote portsmight still be blocked. This blocked state registers as failed faulty to LUNS at the multipath layer. You can verify the remote ports state with the following command:# cat /sys/class/fc_remote_ports/rport-*/port_state	1643459

Use Red Hat Enterprise Linux 7.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device      host                 lun
vserver(cDOT/FlashRay)    lun-pathname  filename    adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383131347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383131347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.5 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

NetApp Bug ID	Title	Description	Bugzilla ID
1139053	Kernel disruption occurs on RHEL7.5 with QLogic QLE2672 16GB FC during storage failover operations	During storage failover operations on the RHEL7U5 kernel with QLogic QLE2672 16GB fibre channel host bus adapter, the kernel disruption occurs due to a panic in the kernel. The kernel panic causes RHEL 7.5 to reboot, which leads to an application disruption. The kernel panic generates the vmcore file under the /var/crash/directory if kdump is configured. The vmcore file is used to understand the cause of the failure. In this case, the panic was observed in the "get_next_timer_interrupt+440" module which is logged in the vmcore file with the following string: " [exception RIP: get_next_timer_interrupt+440]" After the kernel disruption, you can recover the operating system by rebooting the host operating system and restarting the application as required.	1542564

NetApp Bug ID	Title	Description	Bugzilla ID
1138536	Kernel disruption occurs on RHEL7U5 with QLogic QLE2742 32GB FC during storage failover operations	During storage failover operations on the Red Hat Enterprise Linux (RHEL) RHEL7U5 kernel with QLogic QLE2742 HBA, kernel disruption occurs due to a panic in the kernel. The kernel panic leads to a reboot of the operating system, causing an application disruption. The kernel panic generates the vmcore file under the /var/crash/ directory if kdump is configured. When the kernel panics, you can use the vmcore file to investigate the reason for the failure. The following example shows a panic in the bget_next_timer_interrupt +440b module. The panic is logged in the vmcore file with the following string: " [exception RIP: get_next_timer_interrupt+ 440]" You can recover the operating system by rebooting the host OS and restarting the application as required.	1541972

NetApp Bug ID	Title	Description	Bugzilla ID
1148090	Kernel disruption occurs on RHEL 7.5 with QLogic QLE2742 32GB FC HBA during storage failover operations	During storage failover operations on the Red Hat Enterprise Linux (RHEL) 7.5 kernel with a QLogic QLE2742 Fibre Channel (FC) host bus adapter (HBA), a kernel disruption occurs due to a panic in the kernel. The kernel panic causes RHEL 7.5 to reboot, which leads to an application disruption. If the kdump mechanism is enabled, the kernel panic generates a vmcore file located in the /var/crash/ directory. You can analyze the vmcore file to determine the cause of the panic. In this instance, when storage failover with the QLogic QLE2742 HBA event occurs, the "native_queued_spin_lock_slowpath+464" module is affected. You can locate the event in the vmcore file by finding the following string: " [exception RIP: native_queued_spin_lock_slowpath+464]" After the kernel disruption, you can reboot the Host OS and recover the operating system, and then you can restart the applications as required.	1559050

NetApp Bug ID	Title	Description	Bugzilla ID
1146898	Kernel disruption occurs on RHEL 7.5 with Emulex HBAs during storage failover operations	During storage failover operations on a Red Hat Enterprise Linux (RHEL) 7.5 system with Emulex LPe32002-M2 32-GB FC host bus adapters (HBAs), a disruption in the kernel occurs. The kernel disruption causes a reboot of the operating system, which in turn causes an application disruption. If you configure kdump, the kernel disruption generates the vmcore file under the /var/crash/ directory. You can use the vmcore file to determine the cause of the failure. In the following example, you can see the disruption in the "lpfc_hba_clean_txcmplq+368" module. This disruption is logged in the vmcore file with the following string: " [exception RIP: lpfc_hba_clean_txcmplq+368]" After the kernel disruption, reboot the host OS to recover the operating system. Restart the application as required.	1554777

Use Red Hat Enterprise Linux 7.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                 device     host                lun
vserver(cDOT/FlashRay)      lun-pathname    filename   adapter   protocol  size
Product
-------------------------------------------------------------------------
data_vserver            /vol/vol1/lun1     /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver            /vol/vol1/lun1     /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver            /vol/vol2/lun2     /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver            /vol/vol2/lun2     /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.4 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1440718	If you unmap or map a LUN without performing a SCSI rescan, it might lead to data corruption on the host.	When you set the 'disable_changed_wwids' multipath configuration parameter to YES, it disables access to the path device in the event of a WWID change. Multipath will disable access to the path device until the WWID of the path is restored to the WWID of the multipath device. To learn more, see NetApp Knowledge Base: The filesystem corruption on iSCSI LUN on the Oracle Linux 7.	N/A

Use Red Hat Enterprise Linux 7.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the

LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver           /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver           /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver           /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.

The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 7.3 with ONTAP release.

Use Red Hat Enterprise Linux 7.2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383134765724 4e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 7.2 with ONTAP release.

Use Red Hat Enterprise Linux 7.1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp

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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the <DevId> with the WWID string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the multipath.conf file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   sda is the local SCSI disk that we need to add it to the blacklist.

2. Add the WWID to the blacklist stanza in /etc/multipath.conf:

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

You should always check your /etc/multipath.conf file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical multipathd parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the multipath.conf file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.1 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
799323	Emulex FCoE (OCe10102-FX-D) host hang or path failures observed during I/O with storage failover operations	You might observe a host hang or path failures on Emulex 10G FCoE host (OCe10102-FX-D) during I/O with storage failover operations. In such scenarios, you might see the following message: "driver's buffer pool is empty, IO busied and SCSI Layer I/O Abort Request Status"	1061755
836875	IP addresses are not always assigned during the boot of a RHEL 7.0 OS installed on an iSCSI multipath'd LUN	When you install the root(/) on a iSCSI multipath'd LUN, the IP address for the Ethernet interfaces are specified in the kernel command line so that the IP addresses are assigned before the iSCSI service starts. However, dracut cannot assign IP addresses to all the Ethernet ports during the boot, before the iSCSI service starts. This causes the iSCSI login to fail on interfaces without IP addresses. You will see the iSCSI service attempt to login numerous times, which will cause a delay in the OS boot time.	1114966

Use Red Hat Enterprise Linux 7.0 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 7.0 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 7.0 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 7.0 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G  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
  |- 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj  8:144  active ready running
| |- 11:0:2:0 sdr  65:16  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:0:0 sdb  8:i6   active ready running
  |- 12:0:0:0 sdz  65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 7.0 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 7.0 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
844417	Emulex 16G FC (LPe16002B-M6) host crashes during I/O with storage failover operations	You might observe a 16G FC Emulex (LPe16002B-M6) host crash during I/O with storage failover operations.	1131393

NetApp Bug ID	Title	Description	Bugzilla ID
811587	Emulex 16G FC (LPe16002B-M6) host crashes during I/O with storage failover operations	You might observe a 16G FC Emulex (LPe16002B-M6) host crash during I/O with storage failover operations.	1079735
803071	Emulex 16G FC (LPe16002B-M6) host crashes during I/O with storage failover operations	You might observe a 16G FC Emulex (LPe16002B-M6) host crash during I/O with storage failover operations.	1067895
820163	QLogic host hang or path failures observed during I/O with storage failover operations	You might observe a host hang or path failures on QLogic host during I/O with storage failover operations. In such scenarios, you might see the following message: "Mailbox cmd timeout occurred, cmd=0x54, mb[0]=0x54 and Firmware dump saved to temp buffer" messages which leads to host hung/path failure.	1090378
799323	Emulex FCoE (OCe10102-FX-D) host hang or path failures observed during I/O with storage failover operations	You might observe a host hang or path failures on Emulex 10G FCoE host (OCe10102-FX-D) during I/O with storage failover operations. In such scenarios, you might see the following message: "driver's buffer pool is empty, IO busied and SCSI Layer I/O Abort Request Status" messages which leads to host hung/path failures.	1061755

NetApp Bug ID	Title	Description	Bugzilla ID
849212	Emulex 16G FC (LPe16002B-M6) host hang or path failures are observed during I/O with storage failover operations	You might observe a host hang or path failures on Emulex 16G FC (LPe16002B-M6) host during I/O with storage failover operations. In such scenarios, you might see the following message: "RSCN timeout Data and iotag x1301 is out of range: max iotag" messages which leads to host hung/path failures.	1109274
836800	Anaconda displays an iSCSI login failure message although logins are successful during RHEL 7.0 OS installation	When you install the root(/) on a iSCSI multipath'd LUN, the IP address for the Ethernet interfaces are specified in the kernel command line so that the IP addresses are assigned before the iSCSI service starts. However, dracut cannot assign IP addresses to all the Ethernet ports during the boot, before the iSCSI service starts. This causes the iSCSI login to fail on interfaces without IP addresses. You will see the iSCSI service attempt to login numerous times, which will cause a delay in the OS boot time.	1114966

NetApp Bug ID	Title	Description	Bugzilla ID
836875	IP addresses are not always assigned during the boot of a RHEL 7.0 OS installed on an iSCSI multipath'd LUN	When you are installing RHEL 7.0, the anaconda installation screen displays that iSCSI login to multiple target IPs have failed though the iSCSI logins are successful. Anaconda displays following error message: "Node Login Failed" You will observe this error only when you select multiple target IPs for iSCSI login. You can continue the OS installation by clicking the "ok" button. This bug does not hamper either the iSCSI or the RHEL 7.0 OS installation.	1114820
836657	Anaconda does not add bootdev argument in kernel cmd line to set IP address for RHEL 7.0 OS installed on iSCSI multipath'd LUN	Anaconda does not add a bootdev argument in the kernel command line where you set the IPv4 address during the RHEL 7.0 OS installation on an iSCSI multipath'd LUN. This prevents assigning of IP addresses to any of the Ethernet interfaces that were configured to establish iSCSI sessions with the storage subsystem during the RHEL 7.0 boot. Since iSCSI sessions are not established, the root LUN is not discovered when the OS boots and hence the OS boot fails.	1114464

RHEL 6

Use Red Hat Enterprise Linux 6.10 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.10 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability

Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname   filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.10 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.10 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.10 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

* There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
* To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 6.10 with ONTAP release.

Use Red Hat Enterprise Linux 6.9 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.9 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.9 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.9 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.9 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 6.9 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1067272	Remote port status on EMULEX LPe32002 host might be in 'Blocked' state during storage failover operations	During storage failover operations, certain remote port status on RHEL 6.9 host with LPe32002 adapter might get into 'Blocked' state. Because the logical interfaces go down when a storage node is down, the remote port sets the storage node status to "Blocked" state. However, when the storage node comes back to optimal state, the logical interfaces also comes up and the remote port state is expected to be 'Online'. But, on certain occasion the remote port continues to be in 'Blocked' state. This state manifests as 'failed faulty' to LUNS at multipath layer.	427496

NetApp Bug ID	Title	Description	Bugzilla ID
1076584	Firmware dumps occur on Red Hat Enterprise Linux 6.9 QLogic QE8362 HBA during storage failover operations	Firmware dumps can occur during storage failover operations on Red Hat Enterprise Linux (RHEL) 6.9 hosts with QLogic QLE8362 host bus adapters (HBA), firmware dumps are observed occasionally. The firmware dumps might manifest as an I/O outage on the host that can last as long as 1200 seconds. After the adapter completes dumping the firmware cores, the I/O operation resumes normally. No further recovery procedure is required on the host. To indicate the firmware dump, the following message is displayed in /var/log/ message file: kernel: qla2xxx [0000:0c:00.3]-d001:3: Firmware dump saved to temp buffer (3/ffffc90018b01000), dump status flags (0x3f)	1438711

Use Red Hat Enterprise Linux 6.8 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.8 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.

2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.8 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.8 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.8 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 6.8 with ONTAP release.

Use Red Hat Enterprise Linux 6.7 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.7 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.7 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.7 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.7 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuartion.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the RHEL 6.7 with ONTAP release.

Use Red Hat Enterprise Linux 6.6 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.6 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host             lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.6 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.6 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work: `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.6 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 6.6 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
863878	Kernel crash occurs with RHEL 6U6 host during storage failures	You might observe kernel crash on RHEL 6U6 host during storage/fabric.	1158363
1076584	IO stall up to 300 sec seen with QLogic 16G FC (QLE2672) host during storage failures in RHEL 6U4	You might observe an IO stall up to 300 sec on QLogic 16G FC (QLE2672) host during storage/fabric failures.	1135962
795684	RHEL6 U5 multipathd incorrectly group multipath maps during MoD and storage failover fault operations	You might observe an incorrect path grouping on LUNs during LUN move on demand operation along with storage faults. During LUN move operation multipath path priorities will change and multipath is unable to reloads the device table due to device failure caused by storage fault. This leads to incorrect path grouping.	1151020

Use Red Hat Enterprise Linux 6.5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-----------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

Example

```
kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the <DevId> with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"

Parameter	Setting
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 6.5 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
760515	Path failures or host hangs were observed in RHEL 6.5 8G Qlogic FC SAN host during storage failover operations	Path failures or host hangs were observed in RHEL 6.5 8G Qlogic FC SAN host during storage failover operations.	1033136

NetApp Bug ID	Title	Description	Bugzilla ID
758271	bnx2 firmware fails to load when booting with custom initrd (dracut -f)	Broadcom NetXtreme II Gigabit controller ports will not ping due to bnx2 firmware fails to load during boot with custom initrd.	1007463
799394	RHEL 6U5: Emulex 16G FC (LPe16002B-M6) host crash is seen during I/O with storage failover operations	16G FC Emulex (LPe16002B-M6) host crash is seen during I/O with storage failover operations.	1063699
786571	QLogic FCoE host hangs/path failures observed in RHEL 6.5 during I/O with storage failover operations	QLogic FCoE (QLE8242) host hangs/path failures are observed in RHEL 6.5 during I/O with storage failover operations. In such scenarios, you might see the following message: "Mailbox cmd timeout occurred, cmd=0x54, mb[0]=0x54. Scheduling ISP abort" messages which leads to host hung/path failures.	1068619
801580	QLogic 16G FC host hangs or path failures observed in RHEL 6.5 during I/O with storage failover operations	The I/O delays of more than 600 seconds are observed with QLogic 16G FC host (QLE2672) during storage failover operations. In such scenarios, the following message is displayed: "Failed mbx[0]=54, mb[1]=0, mb[2]=76b9, mb[3]=5200, cmd=54"	1068622

Use Red Hat Enterprise Linux 6.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Red Hat Enterprise Linux 6.4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Red Hat Enterprise Linux (RHEL) 6.4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. RHEL 6.4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.
To Enable ALUA Handler, perform the following steps:

Steps

1. Create a backup of the initrd-image.
2. Append the following parameter value to the kernel for ALUA and non-ALUA to work:
   `rdloaddriver=scsi_dh_alua`

   Example

   ```
   kernel /vmlinuz-2.6.32-358.6.1.el6.x86_64 ro root=/dev/mapper/
   vg_ibmx355021082-lv_root rd_NO_LUKS rd_LVM_LV=vg_ibmx355021082/ lv_root
   LANG=en_US.UTF-8 rd_LVM_LV=vg_ibmx355021082/lv_swap rd_NO_MD
   SYSFONT=latarcyrheb-sun16 crashkernel=auto KEYBOARDTYPE=pc KEYTABLE=us
   rd_NO_DM rhgb quiet rdloaddriver=scsi_dh_alua
   ```

3. Use the `mkinitrd` command to recreate the initrd-image.
   RHEL 6x and later versions use either:
   The command: `mkinitrd -f /boot/ initrd-"uname -r".img uname -r`
   Or
   The command: `dracut -f`
4. Reboot the host.
5. Verify the output of the `cat /proc/cmdline` command to ensure that the setting is complete.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
|- 1:0:8:1  sdb 8:16 active ready running
|- 2:0:8:1  sdd 8:48 active ready running
|- 1:0:9:1  sdc 8:32 active ready running
|- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

The RHEL 6.4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

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

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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	"round-robin 0"
polling_interval	5
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

The RHEL 6.4 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
673009	Creating ext4 file system on LV striped across 15 or more discard-enabled, thinly provisioned multipath devices triggers "request botched" kernel errors	"Request blotched" kernel errors have been seen when users attempt to create an ext4 file system on discard-enabled, thinly provisioned multipath devices. As a result, creating the ext4 file system might take longer to complete and occasional disruption occurs. This issue has occurred only when users attempt to create the ext4 file system on a LV striped across 15 or more discard-enabled multipath devices on systems running Red Hat Enterprise Linux 6.x and Data ONTAP 8.1.3 and later operating in 7-Mode. The issue happens because the kernel erroneously attempts to merge discard requests, which is not supported on Red Hat Enterprise Linux 6.x at this time. When this issue occurs, multiple instances of the following message are written to syslog (/var/log/messages): kernel: blk: request botched. As a result, file system creation might take longer time to complete than expected.	907844

Solaris

Use Solaris 11.4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Solaris 11.4 with ONTAP as the target.

Install the Solaris Host Utilities

You can download the compressed file containing the Host Utilities software packages from the NetApp Support Site. After you download the file, you must extract the zip file to get the software packages you need to install the Host Utilities.

Steps

1. Download a copy of the compressed file containing the Host Utilities from the NetApp Support Site to a directory on your host.
2. Go to the directory containing the download.
3. Decompress the file.

   The following example decompresses files for a SPARC system. For x86-64 platforms, use the `x86/x64` package.

   ```
   gunzip netapp_solaris_host_utilities_6_2N20170913_0304_sparc.tar.gz
   ```

4. Use the `tar xvf` command to extract the file.

   ```
   tar xvf netapp_solaris_host_utilities_6_2N20170913_0304_sparc.tar
   ```

5. Add the packages that you extracted from the .tar file to your host.

   ```
   pkgadd -d NTAPSANTool.pkg
   ```

   The packages are added to the `/opt/NTAP/SANToolkit/bin` directory.

   To complete the installation, you must configure the host parameters for your environment (Oracle Solaris I/O Multipathing or MPxIO in this case) by using the `host_config` command.

   The `host_config` command has the following format:

   ```
   /opt/NTAP/SANToolkit/bin/host_config ←setup> ←protocol fcp|iscsi|mixed>
   ←multipath mpxio|dmp| non> [-noalua] [-mcc 60|90|120]
   ```

   The `host_config` command does the following:

   - Changes the FC and SCSI driver settings for x86 and SPARC systems
   - Provides SCSI timeout settings for both MPxIO configurations
   - Sets the VID/PID information
   - Enables or disables ALUA
   - Configures the ALUA settings used by MPxIO and the SCSI drivers for both x86 and SPARC systems

6. Reboot the host.

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
#sanlun lun show

controller(7mode)/                  device
host              lun
vserver(Cmode)      lun-pathname    filename
adapter protocol  size  mode
---------------------------------------------------------------------------
----------------------------
data_vserver       /vol/vol1/lun1
/dev/rdsk/c0t600A098038314362692451465A2F4F39d0s2  qlc1  FCP       60g   C
data_vserver       /vol/vol2/lun2
/dev/rdsk/c0t600A098038314362705D51465A626475d0s2  qlc1  FCP       20g   C
```

SAN booting

What you'll need

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

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 using the FC protocol and running Solaris Host Utilities. The method you use to set up a SAN boot LUN can vary depending on your volume manager and file system. See Install Solaris Host Utilities for details on a SAN boot LUNs in a Solaris MPIO (Multipath I/O) environment.

Multipathing

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 or MPxIO is enabled by default for Solaris 11.4. The default setting in `/kernel/drv/fp.conf` changes to mpxio-disable="no".

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

The path priorities are displayed against the **Access State** section for each LUN in the OS native `mpathadm show lu <LUN>` command.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

The output for the `sanlun` command is the same for ASA and non-ASA configurations.

The path priorities are displayed against the **Access State** section for each LUN in the OS native `mpathadm show lu <LUN>` command.

```
#sanlun lun show -pv sparc-s7-16-49:/vol/solaris_vol_1_0/solaris_lun

                    ONTAP Path: sparc-s7-16-
49:/vol/solaris_vol_1_0/solaris_lun
                           LUN: 0
                      LUN Size: 30g
                   Host Device:
/dev/rdsk/c0t600A098038314362692451465A2F4F39d0s2
                          Mode: C
            Multipath Provider: Sun Microsystems
              Multipath Policy: Native
```

All SAN Arrays (ASA) configurations are supported beginning ONTAP 9.8 for Solaris hosts.

Recommended settings

NetApp recommends using the following parameter settings for Solaris 11.4 SPARC and x86_64 with NetApp ONTAP LUNs. These parameter values are set by Host Utilities. For additional Solaris 11.4 system settings, see Oracle DOC ID: 2595926.1.

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

All Solaris OS versions (including Solaris 10.x and Solaris 11.x) support Solaris HUK 6.2.

- For Solaris 11.4, the FC driver binding is changed from `ssd` to `sd`. The following configuration files get partially updated during the HUK 6.2 installation process:
  - `/kernel/drv/sd.conf`

    - `/etc/driver/drv/scsi_vhci.conf`

- For Solaris 11.3, the FC driver binding uses `ssd`. The following configuration files get partially updated during the HUK 6.2 installation process:
    - `/kernel/drv/ssd.conf`
    - `/etc/driver/drv/scsi_vhci.conf`
- For Solaris 10.x, the following configuration files get fully updated during the HUK 6.2 installation process:
    - `/kernel/drv/sd.conf`
    - `/kernel/drv/ssd.conf`
    - `/kernel/drv/scsi_vhci.conf`

To resolve any configuration issues, see the Knowledge Base article What are the Solaris host recommendations for supporting HUK 6.2.

NetApp recommends the following for a successful 4KB aligned I/O with zpools using NetApp LUNs:

- Verify that you are running a recent enough Solaris OS to ensure that all Solaris features supporting 4KB I/O size alignment are available.
- Verify that the Solaris 10 update 11 is installed with latest kernel patches and Solaris 11.4 with the latest Support Repository Update (SRU).
- The NetApp logical unit must have `lun/host-type` as `Solaris` regardless of the LUN size.

Recommended settings for MetroCluster

By default, the Solaris OS will fail 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 appropriate for standard ONTAP clusters. However, in MetroCluster configurations the value of `fcp_offline_delay` must be increased to **120s** to ensure that I/O does not prematurely time out during operations including unplanned fail overs. For additional information and recommended changes to default settings, 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 have been re-branded generally as "Oracle Virtual Machines" despite the fact that they are based on different architectures.
- In some cases, multiple options can be used together such as a Solaris Container within a particular Solaris Logical Domain.
- NetApp generally supports the use of these virtualization technologies where the overall configuration is supported by Oracle and any partition with direct access to LUNs is listed on the NetApp Interoperability Matrix 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`, do not need specific qualifications as they do not have direct access to NetApp LUNs. Only the partition or virtual machine that has direct access to the underlying LUN, such as an LDOM I/O domain, must be found in the NetApp Interoperability Matrix Tool.

Recommended settings for virtualization

When LUNs are used as virtual disk devices within an LDOM, the source of the LUN is masked by virtualization and the LDOM will not properly detect the block sizes. To prevent this issue, the LDOM OS must be patched for *Oracle Bug 15824910* and a `vdc.conf` file must be created that sets the block size of the virtual disk to `4096`. See Oracle DOC: 2157669.1 for more information.

To verify the patch do the following:

Steps

1. Create a zpool.
2. Run `zdb -C` against the zpool and verify that the value of **ashift** is `12`.

   If the value of **ashift** is not `12`, verify that the correct patch was installed and recheck the contents of `vdc.conf`.

   Do not proceed until **ashift** shows a value of `12`.

> Patches are available for Oracle bug 15824910 on various versions of Solaris. Contact Oracle if assistance is required in determining the best kernel patch.

Recommended settings for SnapMirror Business Continuity

In order to verify that the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in a SnapMirror Business Continuity (SM-BC) environment, you must configure the following setting on the Solaris 11.4 host. This setting overrides the failover module `f_tpgs` to prevent the execution of the code path that detects the contradiction.

> Beginning with ONTAP 9.9.1, SM-BC setting configurations are supported in the Solaris 11.4 host.

Follow the instructions to configure the override parameter:

Steps

1. Create the configuration file `/etc/driver/drv/scsi_vhci.conf` with an entry similar to the following for the NetApp storage type connected to the host:

   ```
   scsi-vhci-failover-override =
   "NETAPP  LUN","f_tpgs"
   ```

2. Use the `devprop` and `mdb` commands to verify that the override parameter has been successfully applied:

   ```
   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 Business Continuity (SM-BC) configuration.

Known issues

The Solaris 11.4 with ONTAP release 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)
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

Use Solaris 11.3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Solaris 11.3 with ONTAP as the target.

Install the Solaris Host Utilities

You can download the compressed file containing the Host Utilities software packages from the NetApp Support Site. After you have the file, you must extract it to get the software packages you need to install the Host Utilities.

Steps

1. Download a copy of the compressed file containing the Host Utilities from the NetApp Support Site to a directory on your host.
2. Go to the directory containing the download.
3. Extract the file.

   The following example decompresses files for a SPARC system. For x86-64 platforms, use the x86/x64 package.

   ```
   gunzip netapp_solaris_host_utilities_6_2N20170913_0304_sparc.tar.gz
   ```

4. Use the `tar xvf` command to unzip the file.

   ```
   tar xvf netapp_solaris_host_utilities_6_2N20170913_0304_sparc.tar
   ```

5. Add the packages that you extracted from the tar file to your host.

   ```
   pkgadd -d NTAPSANTool.pkg
   ```

   The packages are added to the `/opt/NTAP/SANToolkit/bin` directory.

   To complete the installation, you must configure the host parameters for your environment (MPxIO in this case) by using the `host_config` command.

   The `host_config` command has the following format:

   ```
   /opt/NTAP/SANToolkit/bin/host_config <-setup> <-protocol fcp|iscsi|mixed> <-
   multipath mpxio|dmp| non> [-noalua] [-mcc 60|90|120]
   ```

   The `host_config` command does the following:

   - Changes the Fibre Channel and SCSI driver settings for the X86 and SPARC systems
   - Provides SCSI timeout settings for both the MPxIO configurations
   - Sets the VID/PID information
   - Enables or disables ALUA
   - Configures the ALUA settings used by MPxIO and the SCSI drivers for both X86 and SPARC systems.

6. Reboot the host.

SAN toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
#sanlun lun show

controller(7mode)/                 device
host              lun
vserver(Cmode)      lun-pathname     filename
adapter protocol   size  mode
---------------------------------------------------------------------------
----------------------------
data_vserver       /vol/vol1/lun1
/dev/rdsk/c0t600A098038314362692451465A2F4F39d0s2  qlc1  FCP       60g   C
data_vserver       /vol/vol2/lun2
/dev/rdsk/c0t600A098038314362705D51465A626475d0s2  qlc1  FCP       20g   C
```

SAN booting

What you'll need

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

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 using the FC protocol and running the Solaris Host Utilities. The method you use to set up a SAN boot LUN can vary depending on your volume manager and file system. See Install Solaris Host Utilities for details on SAN Booting LUNs in a Solaris MPIO environment.

Multipathing

Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues on the remaining paths.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

Active/Non-Optimized paths:

The path priorities are displayed against the **Access State** section for each LUN in the OS native `mpathadm show lu <LUN>` command.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

The output for the `sanlun` command is the same for ASA and non-ASA configurations.

The path priorities are displayed against the **Access State** section for each LUN in the OS native `mpathadm show lu <LUN>` command.

```
#sanlun lun show -pv sparc-s7-16-49:/vol/solaris_vol_1_0/solaris_lun

                    ONTAP Path: sparc-s7-16-
49:/vol/solaris_vol_1_0/solaris_lun
                           LUN: 0
                      LUN Size: 30g
                   Host Device:
/dev/rdsk/c0t600A0980383143626924514654A2F4F39d0s2
                          Mode: C
            Multipath Provider: Sun Microsystems
              Multipath Policy: Native
```

All SAN Arrays (ASA) configurations are supported beginning in ONTAP 9.8 for Solaris Hosts.

Recommended settings

Following are some parameter settings that are recommended for Solaris 11.3 SPARC and x86_64 with NetApp ONTAP LUNs. These parameter values are set by Host Utilities.

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

Recommended settings for MetroCluster

By default, the Solaris operating system will fail I/Os after 20 seconds 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 appropriate for standard ONTAP clusters. However, in MetroCluster configurations, the value of `fcp_offline_delay` must be increased to **120s** to ensure that I/O does not prematurely time out during operations including unplanned failovers. For additional information and recommended changes to default settings, 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 have been rebranded generally as "Oracle Virtual Machines" despite the fact that they are based on very different architectures.
- In some cases, multiple options can be used together such as a Solaris Container within a particular Solaris Logical Domain.
- NetApp generally supports the use of these virtualization technologies where the overall configuration is supported by Oracle and any partition with direct access to LUNs is listed on the NetApp Interoperability Matrix in a supported configuration. This includes root containers, LDOM IO domains, and LDOM's using NPIV to access LUNs.
- Partitions and/or virtual machines which use only virtualized storage resources, such as a `vdsk`, do not need specific qualification as they do not have direct access to NetApp LUNs. Only the partition/VM that has direct access to the underlying LUN, such as an LDOM IO domain, must be found in the NetApp Interoperability Matrix.

Recommended settings for virtualization

When LUNs are used as virtual disk devices within an LDOM, the source of the LUN is masked by virtualization and the LDOM will not properly detect the block sizes. To prevent this issue, the LDOM operating system must be patched for Oracle Bug 15824910 and a `vdc.conf` file must be created that sets the block size of the virtual disk to 4096. See Oracle Doc 2157669.1 for more information.

To verify the patch, do the following:

Steps

1. Create a zpool.
2. Run `zdb -C` against the zpool and verify that the value of **ashift** is 12.

   If the value of **ashift** is not 12, verify that the correct patch was installed and recheck the contents of vdc.conf.

   Do not proceed until **ashift** shows a value of 12.

> Patches are available for Oracle bug 15824910 on various versions of Solaris. Contact Oracle if assistance is required in determining the best kernel patch.

Recommended settings for SnapMirror Business Continuity

In order to verify that the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in a SnapMirror Business Continuity (SM-BC) environment, you must configure the following setting on the Solaris 11.3 host. This setting overrides the failover module `f_tpgs` to prevent the execution of the code path that detects the contradiction.

> Beginning with ONTAP 9.9.1, SM-BC setting configurations are supported in the Solaris 11.3 host.

Follow the instructions to configure the override parameter:

Steps

1. Create the configuration file `/etc/driver/drv/scsi_vhci.conf` with an entry similar to the following for the NetApp storage type connected to the host:

```
scsi-vhci-failover-override =
"NETAPP  LUN","f_tpgs"
```

2. Use the `devprop` and `mdb` commands to verify that the override parameter has been successfully applied:

```
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 KB article Solaris Host support recommended settings in SnapMirror Business Continuity (SM-BC) configuration.

Known issues

The Solaris 11.3 with ONTAP release 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

NetApp Bug ID	Title	Description	Oracle ID
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

SLES

Release notes

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 does not 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.

SLES 15

Use SUSE Linux Enterprise Server 15 SP5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 SP5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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.

   > ℹ️ 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 SP5, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. The SUSE Linux Enterprise Server 15 SP5 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313476572 44e527766394e 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

SUSE Linux Enterprise Server 15 SP5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}
devices {
   device {
      vendor          "NETAPP  "
      product          "LUN.*"
      no_path_retry     queue
      path_checker      tur
   }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 15 SP5 with ONTAP release.

Use SUSE Linux Enterprise Server 15 SP4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 SP4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                device     host                lun
vserver(cDOT/FlashRay)   lun-pathname   filename   adapter   protocol   size
Product
--------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 SP4 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 15 SP4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e 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
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

SUSE Linux Enterprise Server 15 SP4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the command:
`touch /etc/multipath.conf`.

The first time you create this file, you might need to enable and start the multipath services.

There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.

You can add the following syntax to the `multipath.conf` file to exclude the unwanted devices.

Replace `<DevId>` with the `WWID` string of the device you want to exclude. Use the following command to determine the `WWID`:

Example

In this example, `sda` is the local SCSI disk that we need to add to the blacklist.

Steps

1. Run the following command to determine the `WWID`:

   ```
   # /usr/lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

2. Add the `WWID` value to the blacklist stanza in the `/etc/multipath.conf` file:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding default settings.

The table below shows the critical multipathd parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in `multipath.conf` that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If these parameters cannot be removed because other SAN arrays are still attached to the host, they can instead be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}
devices {
   device {
      vendor         "NETAPP  "
      product         "LUN.*"
      no_path_retry     queue
      path_checker      tur
   }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 15 SP4 with ONTAP release.

Use SUSE Linux Enterprise Server 15 SP3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 SP3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the

LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 SP3 the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 15 SP3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038313477657244e527766394e 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 15 SP3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configuration.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}
devices {
   device {
      vendor         "NETAPP  "
      product         "LUN.*"
      no_path_retry     queue
      path_checker      tur
   }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 15 SP3 with ONTAP release.

Use SUSE Linux Enterprise Server 15 SP2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 SP2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)   lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 SP2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 15 SP2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs. Use the `multipath -ll` command verify the settings for your ONTAP LUNs.

There should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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=enabled
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 15 SP2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}
devices {
   device {
      vendor          "NETAPP  "
      product          "LUN.*"
      no_path_retry     queue
      path_checker      tur
   }
}
```

Known issues

The SLES 15 SP2 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID

1308744	iSCSI boot from SAN fails to boot with a static IP configuration after completing an SLES15SP2 OS installation	iSCSI sanbooted LUN failed to boot up after completing an SLES 15 SP2 OS installation with a static IP configuration. Bootup failure occurs every time with the static IP configuration. This leads to the server refusing to continue the boot up process with the following error message:
		`dracut-cmdline[241]: warning: Empty autoconf values default to dhcp`
		
		`dracut: FATAL: FATAL: For argument ip=eth4:static, setting client-ip does not make sense for dhcp`
		
		`dracut: Refusing to continue`
		
		`reboot: System halted`
			1167494

Use SUSE Linux Enterprise Server 15 SP1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 SP1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 SP1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 15 SP1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
 |- 11:0:0:0 sdb 8:i6 active ready running
 |- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 15 SP1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

Known issues

The SLES 15 SP1 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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 ----	1139137

Use SUSE Linux Enterprise Server 15 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 15 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability

Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)      lun-pathname     filename   adapter    protocol   size
Product
-------------------------------------------------------------------------
data_vserver            /vol/vol1/lun1    /dev/sdb    host16    FCP
120.0g   cDOT
data_vserver            /vol/vol1/lun1    /dev/sdc    host15    FCP
120.0g   cDOT
data_vserver            /vol/vol2/lun2    /dev/sdd    host16    FCP
120.0g   cDOT
data_vserver            /vol/vol2/lun2    /dev/sde    host15    FCP
120.0g   cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 15 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 15 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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=enabled
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
 |- 11:0:0:0 sdb 8:i6 active ready running
 |- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 15 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5

Parameter	Setting
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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry      fail
}

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

Known issues

The SLES 15 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
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)`	1104173

SLES 12

Use SUSE Linux Enterprise Server 12 SP5 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 SP5 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 SP5 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 SP5 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
#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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 SP5 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected

specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
}
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
  }
}
```

Known issues

The SLES 12 SP5 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1284293	Kernel disruption occurs on SLES12 SP5 with QLogic QLE2562 8GB FC HBA during storage failover operations	Kernel disruption occurs during storage failover operations on the SLES12 SP5 kernel with a QLogic QLE2562 Fibre Channel (FC) host bus adapter (HBA). The kernel disruption causes SLES12 SP5 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. Check the vmcore file to determine the cause of the disruption. A storage failover with a QLogic QLE2562 HBA event affects the "THREAD_INFO: ffff8aedf723c2c0" module. Locate this event in the vmcore file by finding the following string: " [THREAD_INFO: ffff8aedf723c2c0]". After the kernel disruption, reboot the host OS to enable it to recover. Then restart the applications.	1157966

Use SUSE Linux Enterprise Server 12 SP4 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 SP4 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you’ll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/             device     host               lun
vserver(cDOT/FlashRay)     lun-pathname  filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 SP4 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 SP4 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
#multipath -ll
3600a0980383134765724e4e527766394e 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 SP4 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

`sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
}
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
  }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 12 SP4 with ONTAP release.

Use SUSE Linux Enterprise Server 12 SP3 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 SP3 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host               lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 SP3 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 SP3 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handler' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 SP3 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 no_path_retry fail
}
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 no_path_retry queue
 path_checker tur
  }
}
```

Known issues

The SLES 15 SP3 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1089555	Kernel disruption observed on kernel version SLES12 SP3 with Emulex LPe16002 16GB FC during storage failover operation	A kernel disruption might occur during storage failover operations on kernel version SLES12 SP3 with Emulex LPe16002 HBA. The kernel disruption prompts a reboot of the operating system, which in turn causes an application disruption. If the kdump is configured, the kernel disruption generates a vmcore file under /var/crash/directory. You can investigate the cause of the failure in the vmcore file. Example: In the observed case, the kernel disruption was observed in the module “lpfc_sli_ringtxcmpl_put+51” and is logged in the vmcore file – exception RIP: lpfc_sli_ringtxcmpl_put+51. Recover the operating system after the kernel disruption by rebooting the host operating system and restarting the application.	1042847

NetApp Bug ID	Title	Description	Bugzilla ID
1089561	Kernel disruption observed on kernel version SLES12 SP3 with Emulex LPe32002 32GB FC during storage failover operations	A kernel disruption might occur during storage failover operations on kernel version SLES12 SP3 with Emulex LPe32002 HBA. The kernel disruption prompts a reboot of the operating system, which in turn causes an application disruption. If the kdump is configured, the kernel disruption generates a vmcore file under /var/crash/directory. You can investigate the cause of the failure in the vmcore file. Example: In the observed case, the kernel disruption was observed in the module “lpfc_sli_free_hbq+76” and is logged in the vmcore file – exception RIP: lpfc_sli_free_hbq+76. Recover the operating system after the kernel disruption by rebooting the host operating system and restarting the application.	1042807

NetApp Bug ID	Title	Description	Bugzilla ID
1117248	Kernel disruption observed on SLES12SP3 with QLogic QLE2562 8GB FC during storage failover operations	During storage failover operations on the Sles12sp3 kernel (kernel-default-4.4.82-6.3.1) with QLogic QLE2562 HBA, the kernel disruption was observed due to a panic in the kernel. The kernel panic leads to a reboot of the operating system, causing an application disruption. The kernel panic generates the vmcore file under the /var/crash/ directory if kdump is configured. Upon the kernel panic, the vmcore file can be used to understand the cause of the failure. Example: In this case, the panic was observed in the “blk_finish_request+289” module. It is logged in the vmcore file with the following string: "exception RIP: blk_finish_request+289" After the kernel disruption, you can recover the operating system by rebooting the Host OS. You can restart the application as required.	1062496

NetApp Bug ID	Title	Description	Bugzilla ID
1117261	Kernel disruption observed on SLES12SP3 with Qlogic QLE2662 16GB FC during storage failover operations	During storage failover operations on Sles12sp3 kernel (kernel-default-4.4.82-6.3.1) with Qlogic QLE2662 HBA, you might observe kernel disruption. This prompts a reboot of the operating system causing application disruption. The kernel disruption generates a vmcore file under /var/crash/ directory if kdump is configured. The vmcore file can be used to understand the cause of the failure. Example: In this case the Kernel disruption was observed in the module "unknown or invalid address" and is logged in vmcore file with the following string - exception RIP: unknown or invalid address. After kernel disruption, the operating system can be recovered by rebooting the host operating system and restarting the application as required.	1062508

NetApp Bug ID	Title	Description	Bugzilla ID
1117274	Kernel disruption observed on SLES12SP3 with Emulex LPe16002 16GB FC during storage failover operations	During storage failover operations on Sles12sp3 kernel (kernel-default-4.4.87-3.1) with Emulex LPe16002 HBA, you might observe kernel disruption. This prompts a reboot of the operating system causing application disruption. The kernel disruption generates a vmcore file under the /var/crash/ directory if kdump is configured. The vmcore file can be used to understand the cause of the failure. Example: In this case kernel disruption was observed in the module "raw_spin_lock_irqsave+30" and is logged in the vmcore file with the following string: – exception RIP: _raw_spin_lock_irqsave+30. After kernel disruption, the operating system can be recovered by rebooting the host operating system and restarting the application as required.	1062514

Use SUSE Linux Enterprise Server 12 SP2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 SP2 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)     lun-pathname      filename   adapter   protocol    size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1     /dev/sdb    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1     /dev/sdc    host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2     /dev/sdd    host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2     /dev/sde    host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 SP2 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 SP2 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 SP2 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"

Parameter	Setting
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 12 SP2 with ONTAP release.

Use SUSE Linux Enterprise Server 12 SP1 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 SP1 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you’ll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/                  device     host                  lun
vserver(cDOT/FlashRay)    lun-pathname    filename   adapter   protocol   size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1    /dev/sdb   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1    /dev/sdc   host15    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sdd   host16    FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2    /dev/sde   host15    FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 SP1 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 SP1 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

> Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

Active/Non-Optimized paths:

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 SP1 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN.*

Parameter	Setting
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

Known issues

There are no known issues for the SUSE Linux Enterprise Server 12 SP1 with ONTAP release.

Use SUSE Linux Enterprise Server 12 with ONTAP

You can use the ONTAP SAN host configuration settings to configure SUSE Linux Enterprise Server 12 with ONTAP as the target.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 32-bit and 64-bit .rpm file. If you do not know which file is right for your configuration, use the NetApp Interoperability Matrix Tool to verify which one you need.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed you should upgrade it or, you should remove it and use the following steps to install the latest version.

1. Download the 32-bit or 64-bit Linux Unified Host Utilities software package from the NetApp Support Site Site to your host.
2. Use the following command to install the software package:

```
rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
```

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

SAN Toolkit

The toolkit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following example, the `sanlun lun show` command returns LUN information.

```
# sanlun lun show all
```

Example output:

```
controller(7mode/E-Series)/            device     host              lun
vserver(cDOT/FlashRay)    lun-pathname  filename   adapter  protocol  size
Product
-------------------------------------------------------------------------
data_vserver          /vol/vol1/lun1   /dev/sdb   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol1/lun1   /dev/sdc   host15   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sdd   host16   FCP
120.0g  cDOT
data_vserver          /vol/vol2/lun2   /dev/sde   host15   FCP
120.0g  cDOT
```

SAN Booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For SUSE Linux Enterprise Server 12 the /etc/multipath.conf file must exist, but you do not need to make specific changes to the file. SUSE Linux Enterprise Server 12 is compiled with all settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a0980383034466b2b4a3775474859 dm-3 NETAPP,LUN C-Mode
size=20G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=50 status=active
| |- 1:0:8:1  sdb 8:16 active ready running
| `- 2:0:8:1  sdd 8:48 active ready running
`-+- policy='round-robin 0' prio=10 status=enabled
  |- 1:0:9:1  sdc 8:32 active ready running
  `- 2:0:9:1  sde 8:64 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038313477657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G 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
| |- 11:0:1:0 sdj 8:144 active ready running
| |- 11:0:2:0 sdr 65:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
|- 11:0:0:0 sdb 8:i6 active ready running
|- 12:0:0:0 sdz 65:144 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

SUSE Linux Enterprise Server 12 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly.
The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

1. Run the following command to determine the WWID:

```
# /lib/udev/scsi_id -gud /dev/sda
360030057024d0730239134810c0cb833
```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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"
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

Parameter	Setting
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `detect_prio` that are not compatible with ONTAP LUNs. If they cannot be removed because of other SAN arrays still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
 path_checker readsector0
 detect_prio no
 }
devices {
 device {
 vendor "NETAPP "
 product "LUN.*"
 path_checker tur
 detect_prio yes
 }
}
```

Known issues

The SLES 12 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
873555	scsi_dh_alua module is not loaded during multipathd startup on local boot	scsi_dh_alua is a Linux ALUA device handler module. This is is not loaded during multipathd startup on local boot. Due to this device handler will not be loaded though ALUA is enabled on target side.	908529
863584	The message "conflicting device node '/dev/mapper/360xx' found" appears on the screen when you create a DM device on SLES12	You might observe a failure in creating a link to DM devices under /dev/mapper dir in SLES 12 and see the messages "conflicting device node '/dev/mapper/360xx' found".	903001

NetApp Bug ID	Title	Description	Bugzilla ID
847490	Multipath daemon shows path failures on SLES 12	You might observe path failures on the SLES12 multipath daemon during I/O with storage or fabric faults.	890854

Ubuntu

Use Ubuntu 22.04 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Ubuntu 22.04 with ONTAP as the target.

> NetApp Linux Unified Host Utilities software package is not available for Ubuntu 22.04 OS.

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Ubuntu 22.04, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. Ubuntu 22.04 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038314559533f524d6c652f62 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038314c4c715d5732674e6141 dm-0 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
| |- 11:0:1:0 sda 8:0  active ready running
| `- 12:0:2:0 sdd 8:48 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 11:0:2:0 sdb 8:16 active ready running
  `- 12:0:1:0 sdc 8:32 active ready running
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Ubuntu 22.04 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

  In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

  Steps

  1. Run the following command to determine the WWID:

     ```
     # /lib/udev/scsi_id -gud /dev/sda
     360030057024d0730239134810c0cb833
     ```

     `sda` is the local SCSI disk that we need to add it to the blacklist.

  2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Ubuntu 22.04 with ONTAP release.

Use Ubuntu 20.04 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Ubuntu 20.04 with ONTAP as the target.

NetApp Linux Unified Host Utilities software package is not available for Ubuntu 20.04 OS.

SAN booting

What you'll need

If you decide to use SAN booting, it must be supported by your configuration. You can use the NetApp 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 is successful.

Multipathing

For Ubuntu 20.04, the `/etc/multipath.conf` file must exist, but you do not need to make specific changes to the file. Ubuntu 20.04 is compiled with all the settings required to recognize and correctly manage ONTAP LUNs.

You can use the `multipath -ll` command to verify the settings for your ONTAP LUNs.
The following sections provide sample multipath output for a LUN mapped to ASA and non-ASA personas.

All SAN Array configurations

In All SAN Array (ASA) configurations, all paths to a given LUN are active and optimized. This improves performance by serving I/O operations through all paths at the same time.

Example

The following example displays the correct output for an ONTAP LUN:

```
# multipath -ll
3600a098038314559533f524d6c652f62 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# multipath -ll
3600a098038314837352453694b542f4a 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
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

The Ubuntu 20.04 OS is compiled to recognize ONTAP LUNs and automatically set all configuration parameters correctly for both ASA and non-ASA configurations. You can further optimize performance for your host configuration with the following recommended settings.

The `multipath.conf` file must exist for the multipath daemon to start, but you can create an empty, zero-byte file by using the following command:

`touch /etc/multipath.conf`

The first time you create this file, you might need to enable and start the multipath services:

```
# systemctl enable multipathd
# systemctl start multipathd
```

- There is no requirement to add anything directly to the `multipath.conf` file, unless you have devices that you do not want to be managed by multipath or you have existing settings that override defaults.
- To exclude unwanted devices, add the following syntax to the `multipath.conf` file .

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

  Replace the `<DevId>` with the `WWID` string of the device you want to exclude.

  Example

In this example, we are going to determine the WWID of a device and add to the `multipath.conf` file.

Steps

1. Run the following command to determine the WWID:

   ```
   # /lib/udev/scsi_id -gud /dev/sda
   360030057024d0730239134810c0cb833
   ```

   `sda` is the local SCSI disk that we need to add it to the blacklist.

2. Add the `WWID` to the blacklist stanza in `/etc/multipath.conf`:

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

You should always check your `/etc/multipath.conf` file for legacy settings, especially in the defaults section, that might be overriding the default settings.

The following table demonstrates the critical `multipathd` parameters for ONTAP LUNs and the required values. If a host is connected to LUNs from other vendors and any of these parameters are overridden, they will need to be corrected by later stanzas in the `multipath.conf` file that apply specifically to ONTAP LUNs. If this is not done, the ONTAP LUNs might not work as expected. These defaults should only be overridden in consultation with NetApp and/or an OS vendor and only when the impact is fully understood.

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

Parameter	Setting
prio	"ontap"
product	LUN.*
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

Example

The following example shows how to correct an overridden default. In this case, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that are not compatible with ONTAP LUNs. If they cannot be removed because other SAN arrays are still attached to the host, these parameters can be corrected specifically for ONTAP LUNs with a device stanza.

```
defaults {
   path_checker      readsector0
   no_path_retry     fail
}

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

KVM settings

You can use the recommended settings to configure Kernel-based Virtual Machine (KVM) as well. There are no changes required to configure KVM as the LUN is mapped to the hypervisor.

Known issues

There are no known issues for the Ubuntu 20.04 with ONTAP release.

Veritas

Use Veritas Infoscale 8 for Linux with ONTAP

You can use ONTAP SAN host configuration settings for the Veritas Infoscale storage foundation 8 series release for Red Hat Enterprise Linux and Oracle Linux (RHCK based) platforms with FC, FCoE, and iSCSI protocols.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN Toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following illustration, the `sanlun show` command returns LUN information.

```
# sanlun show -p -v SFRAC:/vol/fen1/lun1

              ONTAP Path: SFRAC:/vol/fen1/lun1
                     LUN: 0
                LUN Size: 10g
                 Product: cDOT
                DMP NODE: sfrac0_47
      Multipath Provider: Veritas
---------- ------ --------- ------
------------------------------------
Veritas     host   vserver    host:
path        path   path     /dev/   chan:     vserver           major:
state       state  type     node    id:lun    LIF                   minor
---------- ------ --------- ------
------------------------------------
enabled     up     active/non-optimized sdea   14:0:1:0   lif_10
128:32
enabled (a) up     active/optimized     sdcj   14:0:0:0   lif_2
69:112
enabled (a) up     active/optimized     sdb    13:0:0:0   lif_1
8:16
enabled     up     active/non-optimized sdas   13:0:1:0   lif_9
66:192
```

SAN Booting

What you'll need

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

Refer to Veritas Support Portal (Product Matrix, Platform Lookup, HCL Matrix) to verify SAN Boot configuration supportability and known caveats.

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 is successful.

Multipathing

You must verify that your configuration meets the system requirements. For more information, see the NetApp Interoperability Matrix Tool and the Veritas HCL Matrix.

Example

In this example, the `vxdmpadm` command is used to verify that the VxDMP Multipath has the ONTAP target array attached.

```
# 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
```

With Veritas Dynamic Multipathing (VxDMP), you must perform configuration tasks to claim NetApp LUNs as Veritas Multipath Devices. You must have the Array Support Library (ASL) and the Array Policy Module (APM) packages installed that Veritas provides for NetApp storage systems. While the Veritas Software Installation loads the default ASL APM packages along with the product, it is recommended to use the latest supported packages listed on Veritas support portal.

Example

The following example displays the Veritas Support Library (ASL) and the Array Policy Module (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
```

All SAN Array Configuration

In All SAN Array (ASA) configurations, all paths to a given Logical Unit (LUN) are active and optimized. This means I/O can be served through all paths at the same time, thereby enabling better performance.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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     -
```

> ⓘ Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Non-ASA Configuration

For non-ASA configuration there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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     -
-
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Recommended Settings

Settings for Veritas Multipath

The following Veritas VxDMP tunables are recommended by NetApp for optimum system configuration in storage failover operations.

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

DMP tunables are set online by using the `vxdmpadm` command as follows:

```
# vxdmpadm settune dmp_tunable=value
```

The values of these tunable can be verified dynamically by using `#vxdmpadm gettune`.

Example

The following example shows the effective VxDMP tunables on the SAN host.

```
# 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
```

Settings by protocol

- For FC/FCoE only: Use the default timeout values.
- For iSCSI only: 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. Setting the value of `replacement_timeout` to 120 in the iSCSI configuration file is recommended.

Example

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

Settings by OS platforms

For Red Hat Enterprise Linux 7 and 8 series, you must configure `udev rport` values to support the Veritas Infoscale environment in storage failover scenarios. Create 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'"
```

For all other settings specific to Veritas, refer to the standard Veritas Infoscale product documentation.

Multipath Coexistence

If you have a heterogenous multipath environment including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, please refer to the Veritas Product Administration guide for configuration settings.

Known issues

There are no known issues for the Veritas Infoscale 8 for Linux with ONTAP release.

Use Veritas Infoscale 7 for Linux with ONTAP

You can use ONTAP SAN host configuration settings for the Veritas Infoscale storage foundation 7 series release for Red Hat Enterprise Linux & Oracle Linux (RHCK based) platforms with FC, FCoE and iSCSI protocols.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN Toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following illustration, the `sanlun show` command returns LUN information.

```
# sanlun show -p -v SFRAC:/vol/fen1/lun1

              ONTAP Path: SFRAC:/vol/fen1/lun1
                     LUN: 0
                LUN Size: 10g
                 Product: cDOT
                DMP NODE: sfrac0_47
      Multipath Provider: Veritas
----------  ------  ---------  ------
-------------------------------------
Veritas     host    vserver    host:
path        path    path      /dev/   chan:     vserver           major:
state       state   type      node    id:lun    LIF                   minor
----------  ------  ---------  ------
-------------------------------------
enabled     up      active/non-optimized sdea   14:0:1:0   lif_10
128:32
enabled (a) up      active/optimized     sdcj   14:0:0:0   lif_2
69:112
enabled (a) up      active/optimized     sdb    13:0:0:0   lif_1
8:16
enabled     up      active/non-optimized sdas   13:0:1:0   lif_9
66:192
```

SAN Booting

What you'll need

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

Refer to Veritas Support Portal (Product Matrix, Platform Lookup, HCL Matrix) to verify SAN Boot configuration supportability and known caveats.

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 is successful.

Multipathing

You must verify that your configuration meets the system requirements. For more information, see the NetApp Interoperability Matrix Tool and the Veritas HCL Matrix.

Example

In this example, the `vxdmpadm` command is used to verify that the VxDMP Multipath has the ONTAP target array attached.

```
# 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
```

With Veritas Dynamic Multipathing (VxDMP), you must perform configuration tasks to claim NetApp LUNs as Veritas Multipath Devices. You must have the Array Support Library (ASL) and the Array Policy Module (APM) packages installed that Veritas provides for NetApp storage systems. While the Veritas Software Installation loads the default ASL APM packages along with the product, it is recommended to use the latest supported packages listed on Veritas support portal.

Example

The following example displays the Veritas Support Library (ASL) and the Array Policy Module (APM) configuration.

```
# vxdmpadm list dmpnode dmpnodename=sfrac0_47 | grep asl
asl         = libvxnetapp.so
# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so              vm-7.4-rev-1    6.1

# 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
```

All SAN Array Configuration

In All SAN Array (ASA) configurations, all paths to a given Logical Unit (LUN) are active and optimized. This means I/O can be served through all paths at the same time, thereby enabling better performance.

Example

The following example displays the correct output for an ONTAP LUN:

```
# 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     -
```

Do not use an excessive number of paths to a single LUN. No more than 4 paths should be required. More than 8 paths might cause path issues during storage failures.

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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     -
-
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended Settings

Settings for Veritas Multipath

The following Veritas VxDMP tunables are recommended by NetApp for optimum system configuration in storage failover operations.

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

DMP tunables are set online by using the `vxdmpadm` command as follows:

```
# vxdmpadm settune dmp_tunable=value
```

The values of these tunable can be verified dynamically by using `#vxdmpadm gettune`.

Example

The following example shows the effective VxDMP tunables on the SAN host.

```
# 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
```

Settings by protocol

- For FC/FCoE only: Use the default timeout values.
- For iSCSI only: 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. Setting the value of `replacement_timeout` to 120 in the iSCSI configuration file is recommended.

Example

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

Settings by OS platforms

For Red Hat Enterprise Linux 7 and 8 series, you must configure `udev rport` values to support the Veritas Infoscale environment in storage failover scenarios. Create 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'"
```

For all other settings specific to Veritas, refer to the standard Veritas Infoscale product documentation.

Multipath Coexistence

If you have a heterogenous multipath environment including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, please refer to the Veritas Product Administration guide for configuration settings.

Known issues

There are no known issues for the Veritas Infoscale 7 for Linux with ONTAP release.

Use Veritas Storage Foundation 6 for Linux with ONTAP

You can use ONTAP SAN host configuration settings for the Veritas Storage Foundation 6 series release for Red Hat Enterprise Linux & Oracle Linux (RHCK based) platforms with FC, FCoE and iSCSI protocols.

Install the Linux Unified Host Utilities

The NetApp Linux Unified Host Utilities software package is available on the NetApp Support Site in a 64-bit .rpm file.

NetApp strongly recommends installing the Linux Unified Host Utilities, but it is not mandatory. The utilities do not change any settings on your Linux host. The utilities improve management and assist NetApp customer support in gathering information about your configuration.

What you'll need

If you have a version of Linux Unified Host Utilities currently installed, you should upgrade or remove it, and then use the following steps to install the latest version.

Steps

1. Download the 64-bit Linux Unified Host Utilities software package from the NetApp Support Site to your host.
2. Install the software package:

   ```
   rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64
   ```

SAN Toolkit

The tool kit is installed automatically when you install the NetApp Host Utilities package. This kit provides the `sanlun` utility, which helps you manage LUNs and HBAs. The `sanlun` command returns information about the LUNs mapped to your host, multipathing, and information necessary to create initiator groups.

Example

In the following illustration, the `sanlun show` command returns LUN information.

```
# sanlun show -p -v SFRAC:/vol/fen1/lun1

              ONTAP Path: SFRAC:/vol/fen1/lun1
                     LUN: 0
                LUN Size: 10g
                 Product: cDOT
                DMP NODE: sfrac0_47
      Multipath Provider: Veritas
----------  ------  ---------  ------
-------------------------------------
Veritas     host    vserver    host:
path        path    path      /dev/   chan:     vserver           major:
state       state   type      node    id:lun    LIF                   minor
----------  ------  ---------  ------
-------------------------------------
enabled     up      active/non-optimized sdea   14:0:1:0   lif_10
128:32
enabled (a) up      active/optimized     sdcj   14:0:0:0   lif_2
69:112
enabled (a) up      active/optimized     sdb    13:0:0:0   lif_1
8:16
enabled     up      active/non-optimized sdas   13:0:1:0   lif_9
66:192
```

SAN Booting

What you'll need

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

Refer to Veritas Support Portal (Product Matrix, Platform Lookup, HCL Matrix) to verify SAN Boot configuration supportability and known caveats.

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 is successful.

Multipathing

You must verify that your configuration meets the system requirements. For more information, see the NetApp Interoperability Matrix Tool and the Veritas HCL Matrix.

Example

In this example, the `vxdmpadm` command is used to verify that the VxDMP Multipath has the ONTAP target array attached.

```
# 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
```

With Veritas Dynamic Multipathing (VxDMP), you must perform configuration tasks to claim NetApp LUNs as Veritas Multipath Devices. You must have the Array Support Library (ASL) and the Array Policy Module (APM) packages installed that Veritas provides for NetApp storage systems. While the Veritas Software Installation loads the default ASL APM packages along with the product, it is recommended to use the latest supported packages listed on Veritas support portal.

Example

The following example displays the Veritas Support Library (ASL) and the Array Policy Module (APM) configuration.

```
# vxdmpadm list dmpnode dmpnodename=sfrac0_47 | grep asl
asl         = libvxnetapp.so
```

```
# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so              vm-7.4-rev-1    6.1

# 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
```

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located. The paths with the lower priorities are active but are non-optimized because they are served from a different controller. The non-optimized paths are only used when no optimized paths are available.

Example

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

```
# 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     -
-
```

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Settings for Veritas Multipath

The following Veritas VxDMP tunables are recommended by NetApp for optimum system configuration in storage failover operations.

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

DMP tunables are set online by using the `vxdmpadm` command as follows:

```
# vxdmpadm settune dmp_tunable=value
```

The values of these tunable can be verified dynamically by using `#vxdmpadm gettune`.

Example

The following example shows the effective VxDMP tunables on the SAN host.

```
# 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
```

Settings by protocol

- For FC/FCoE only: Use the default timeout values.
- For iSCSI only: 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. Setting the value of `replacement_timeout` to 120 in the iSCSI configuration file is recommended.

Example

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

Settings by OS platforms

For Red Hat Enterprise Linux 7 and 8 series, you must configure `udev rport` values to support the Veritas Infoscale environment in storage failover scenarios. Create 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'"
```

> For all other settings specific to Veritas, refer to the standard Veritas Infoscale product documentation.

Multipath Coexistence

If you have a heterogenous multipath environment including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, please refer to the Veritas Product Administration guide for configuration settings.

Known issues

There are no known issues for the Veritas Storage Foundation 6 for Linux with ONTAP release.

Windows

Using Windows Server 2022 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Windows server 2022 with ONTAP as the target.

Booting the OS

There are two options for booting the operating system: by using either local boot or SAN boot. For local booting, you install the OS on the local hard disk (SSD, SATA, RAID, and so on). For SAN booting, see instructions below.

SAN booting

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

1. Map the SAN boot LUN to the host.
2. Verify multiple paths are available. Remember, multiple paths will only be 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 the boot is successful.

Install Windows hotfixes

NetApp recommends that the **latest cumulative update** is installed on the server.

> Go to the Microsoft Update Catalog 2022 website to obtain and install the required Windows hotfixes for your version of Windows.

1. Download hotfixes from the Microsoft support site.

> Some hotfixes are not available for direct download. In these cases, you will need to request a given hotfix from Microsoft support personnel.

2. Follow the instructions provided by Microsoft to install the hotfixes.

> Many hotfixes require a reboot of your Windows host, but you can opt to wait to reboot the host until *after* you install or upgrade the Host Utilities.

Install the Windows Unified Host Utilities

The Windows Unified Host Utilities (WUHU) are a set of software programs with documentation that enables you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installation of the latest utility kit. For WUHU configuration information and instructions, refer to the Windows Unified Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Multipathing

You must install MPIO software and have multipathing set up if your Windows host has more than one path to the storage system. Without MPIO software, the operating system might see each path as a separate disk, which can lead to data corruption. The MPIO software presents a single disk to the operating system for all paths, and a device-specific module (DSM) manages path failover.

On a Windows system, the two main components to any MPIO solution are a DSM and the Windows MPIO. MPIO is not supported for Windows XP or Windows Vista running in a Hyper- V virtual machine.

> When you select MPIO support, the Windows Unified Host Utilities enables the included MPIO feature of Windows Server 2022.

SAN configuration

Non-ASA configuration

For non-ASA configuration there should be two groups of paths with different priorities.

The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located.

The paths with the lower priorities are active but are non-optimized because they are served from a different controller.

The non-optimized paths are only used when no optimized paths are available.

Example

The following example displays the correct output for an ONTAP LUN with two active/optimized paths and two active/non-optimized paths.

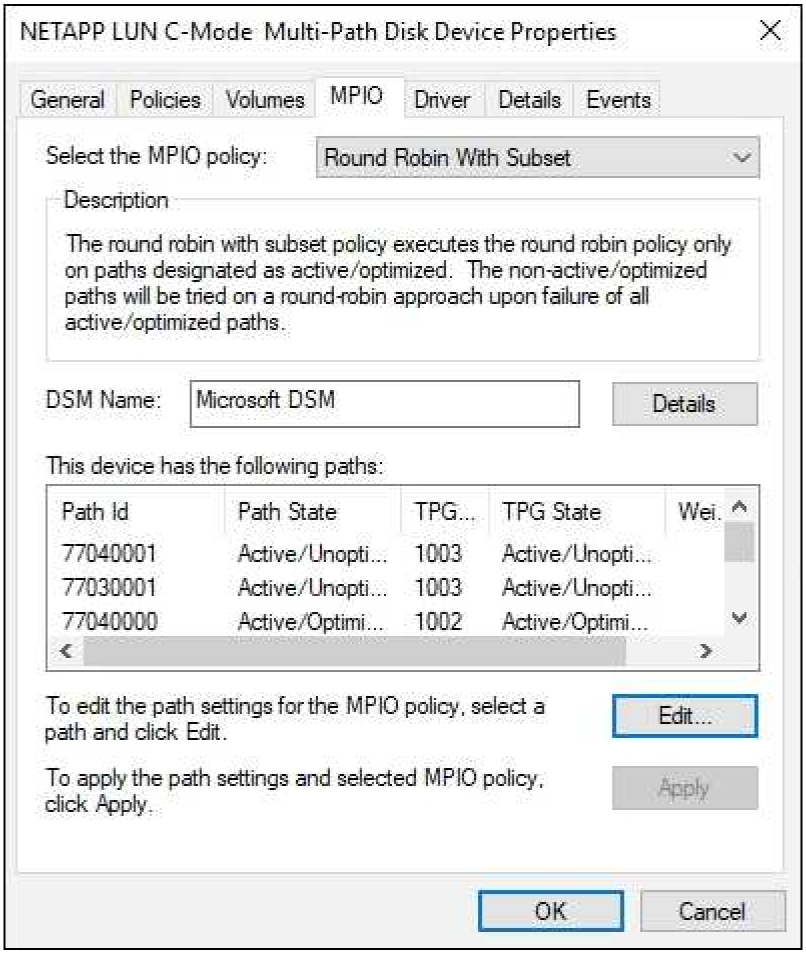

All SAN array configuration

For All SAN Array (ASA) configuration, there should be one group of paths with single priorities. All paths are active/optimized; that is, they are serviced by the controller and that the I/O is sent on all the active paths.

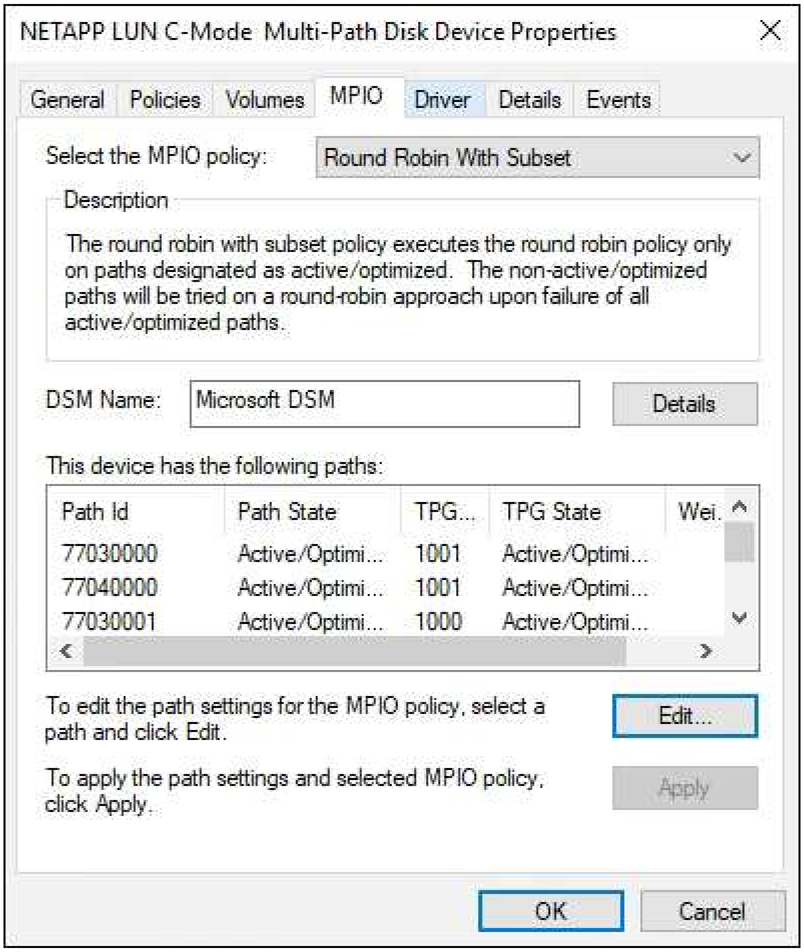

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

On systems using FC, the following timeout values for Emulex and QLogic FC HBAs are required when MPIO is selected.

For Emulex Fibre Channel HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

For QLogic Fibre Channel HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

Windows Unified Host Utility will set these values. For detailed recommended settings, refer to the Windows Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Known issues

There are no known issues for the Windows Server 2022 with ONTAP release.

Using Windows Server 2019 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Windows server 2019 with ONTAP as the target.

Booting the OS

There are two options for booting the operating system: by using either local boot or SAN boot. For local booting, you install the OS on the local hard disk (SSD, SATA, RAID, and so on). For SAN booting, see instructions below.

SAN booting

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

1. Map the SAN boot LUN to the host.
2. Verify multiple paths are available. Remember, multiple paths will only be 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 the boot is successful.

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Install Windows hotfixes

NetApp recommends that the **latest cumulative update** is installed on the server.

Go to the Microsoft Update Catalog 2019 website to obtain and install the required Windows hotfixes for your version of Windows.

1. Download hotfixes from the Microsoft support site.

   Some hotfixes are not available for direct download. In these cases, you will need to request a given hotfix from Microsoft support personnel.

2. Follow the instructions provided by Microsoft to install the hotfixes.

Many hotfixes require a reboot of your Windows host, but you can opt to wait to reboot the host until *after* you install or upgrade the Host Utilities.

Install the Windows Unified Host Utilities

The Windows Unified Host Utilities (WUHU) are a set of software programs with documentation that enables you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installation of the latest utility kit. For WUHU configuration information and instructions, refer to the Windows Unified Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Multipathing

You must install MPIO software and have multipathing set up if your Windows host has more than one path to the storage system. Without MPIO software, the operating system might see each path as a separate disk, which can lead to data corruption. The MPIO software presents a single disk to the operating system for all paths, and a device-specific module (DSM) manages path failover.

On a Windows system, the two main components to any MPIO solution are a DSM and the Windows MPIO. MPIO is not supported for Windows XP or Windows Vista running in a Hyper- V virtual machine.

When you select MPIO support, the Windows Unified Host Utilities enables the included MPIO feature of Windows Server 2019.

SAN configuration

Non-ASA configuration

For non-ASA configuration there should be two groups of paths with different priorities.

The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located.

The paths with the lower priorities are active but are non-optimized because they are served from a different controller.

The non-optimized paths are only used when no optimized paths are available.

Example

The following example displays the correct output for an ONTAP LUN with two active/optimized paths and two active/non-optimized paths.

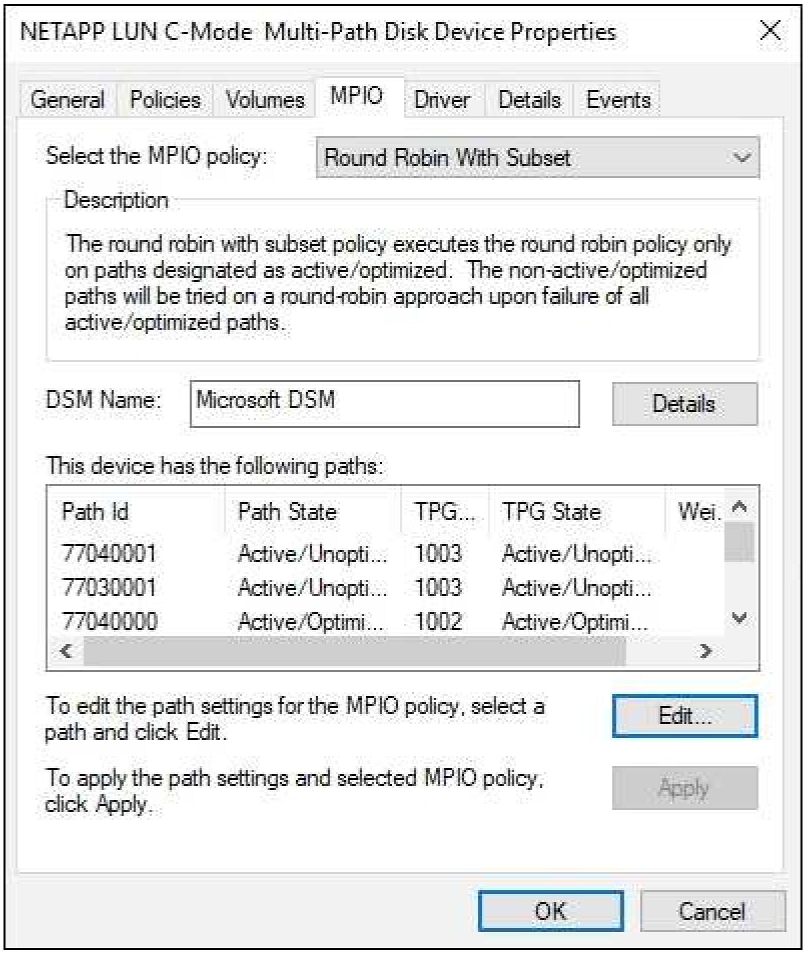

All SAN array configuration

For All SAN Array (ASA) configuration, there should be one group of paths with single priorities. All paths are active/optimized; that is, they are serviced by the controller and that the I/O is sent on all the active paths.

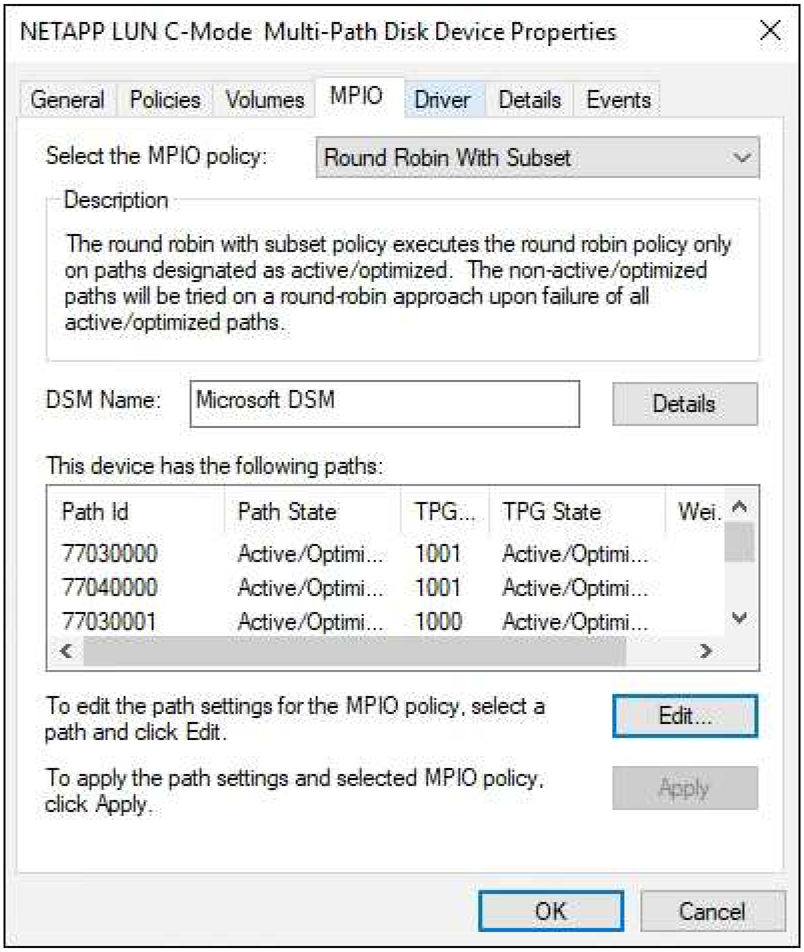

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

On systems using FC, the following timeout values for Emulex and QLogic FC HBAs are required when MPIO is selected.

For Emulex Fibre Channel HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

For QLogic Fibre Channel HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

Windows Unified Host Utility will set these values. For detailed recommended settings, refer to the Windows Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Known issues

There are no known issues for the Windows Server 2019 with ONTAP release.

Using Windows Server 2016 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Windows server 2016 with ONTAP as the target.

Booting the OS

There are two options for booting the operating system: by using either local boot or SAN boot. For local booting, you install the OS on the local hard disk (SSD, SATA, RAID, and so on). For SAN booting, see instructions below.

SAN booting

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

1. Map the SAN boot LUN to the host.
2. Verify multiple paths are available. Remember, multiple paths will only be 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 the boot is successful.

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Install Windows hotfixes

NetApp recommends that the **latest cumulative update** is installed on the server.

Go to the Microsoft Update Catalog 2016 website to obtain and install the required Windows hotfixes for your version of Windows.

1. Download hotfixes from the Microsoft support site.

   Some hotfixes are not available for direct download. In these cases, you will need to request a given hotfix from Microsoft support personnel.

2. Follow the instructions provided by Microsoft to install the hotfixes.

Many hotfixes require a reboot of your Windows host, but you can opt to wait to reboot the host until *after* you install or upgrade the Host Utilities.

Install the Windows Unified Host Utilities

The Windows Unified Host Utilities (WUHU) are a set of software programs with documentation that enables you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installation of the latest utility kit. For WUHU configuration information and instructions, refer to the Windows Unified Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Multipathing

You must install MPIO software and have multipathing set up if your Windows host has more than one path to the storage system. Without MPIO software, the operating system might see each path as a separate disk, which can lead to data corruption. The MPIO software presents a single disk to the operating system for all paths, and a device-specific module (DSM) manages path failover.

On a Windows system, the two main components to any MPIO solution are a DSM and the Windows MPIO. MPIO is not supported for Windows XP or Windows Vista running in a Hyper- V virtual machine.

When you select MPIO support, the Windows Unified Host Utilities enables the included MPIO feature of Windows Server 2016.

SAN configuration

Non-ASA configuration

For non-ASA configuration there should be two groups of paths with different priorities.

The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located.

The paths with the lower priorities are active but are non-optimized because they are served from a different controller.

The non-optimized paths are only used when no optimized paths are available.

Example

The following example displays the correct output for an ONTAP LUN with two active/optimized paths and two active/non-optimized paths.

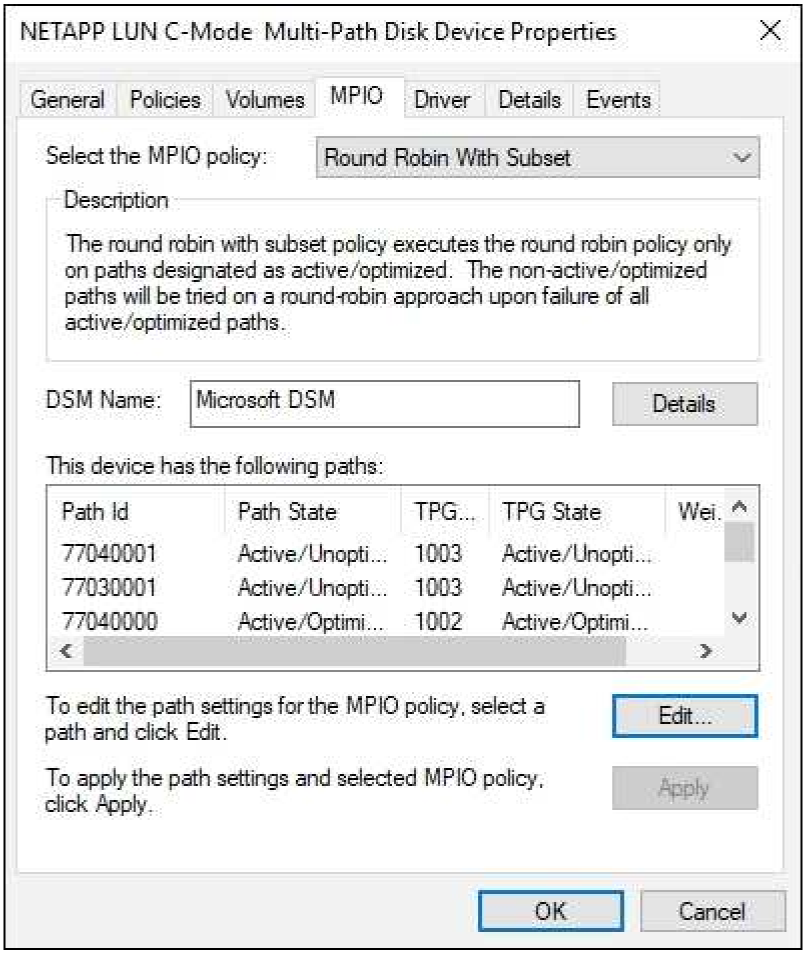

All SAN array configuration

For All SAN Array (ASA) configuration, there should be one group of paths with single priorities. All paths are active/optimized; that is, they are serviced by the controller and that the I/O is sent on all the active paths.

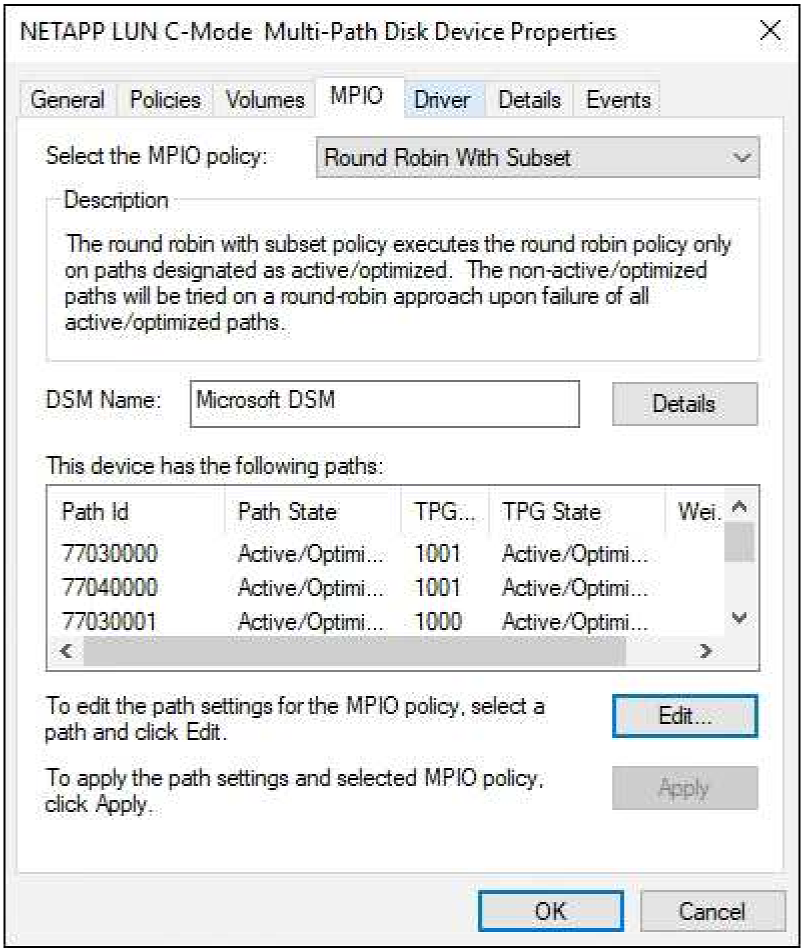

> ⓘ Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Recommended settings

On systems using FC, the following timeout values for Emulex and QLogic FC HBAs are required when MPIO is selected.

For Emulex Fibre Channel HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

For QLogic Fibre Channel HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

Windows Unified Host Utility will set these values. For detailed recommended settings, refer to the Windows Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Known issues

There are no known issues for the Windows Server 2016 with ONTAP release.

Using Windows Server 2012 R2 with ONTAP

You can use the ONTAP SAN host configuration settings to configure Windows server 2012 R2 with ONTAP as the target.

Booting the OS

There are two options for booting the operating system: by using either local boot or SAN boot. For local booting, you install the OS on the local hard disk (SSD, SATA, RAID, and so on). For SAN booting, see instructions below.

SAN booting

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

1. Map the SAN boot LUN to the host.
2. Verify multiple paths are available. Remember, multiple paths will only be 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 the boot is successful.

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Install Windows hotfixes

NetApp recommends that the **latest cumulative update** is installed on the server.

Go to the Microsoft Update Catalog 2012 R2 website to obtain and install the required Windows hotfixes for your version of Windows.

1. Download hotfixes from the Microsoft support site.

Some hotfixes are not available for direct download. In these cases, you will need to request a given hotfix from Microsoft support personnel.

2. Follow the instructions provided by Microsoft to install the hotfixes.

Many hotfixes require a reboot of your Windows host, but you can opt to wait to reboot the host until *after* you install or upgrade the Host Utilities.

Install the Windows Unified Host Utilities

The Windows Unified Host Utilities (WUHU) are a set of software programs with documentation that enables you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installation of the latest utility kit. For WUHU configuration information and instructions, refer to the Windows Unified Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Multipathing

You must install MPIO software and have multipathing set up if your Windows host has more than one path to the storage system. Without MPIO software, the operating system might see each path as a separate disk, which can lead to data corruption. The MPIO software presents a single disk to the operating system for all paths, and a device-specific module (DSM) manages path failover.

On a Windows system, the two main components to any MPIO solution are a DSM and the Windows MPIO. MPIO is not supported for Windows XP or Windows Vista running in a Hyper- V virtual machine.

When you select MPIO support, the Windows Unified Host Utilities enables the included MPIO feature of Windows Server 2012 R2.

SAN configuration

Non-ASA configuration

For non-ASA configuration there should be two groups of paths with different priorities.

The paths with the higher priorities are Active/Optimized, meaning they are serviced by the controller where the aggregate is located.

The paths with the lower priorities are active but are non-optimized because they are served from a different controller.

The non-optimized paths are only used when no optimized paths are available.

Example

The following example displays the correct output for an ONTAP LUN with two active/optimized paths and two active/non-optimized paths.

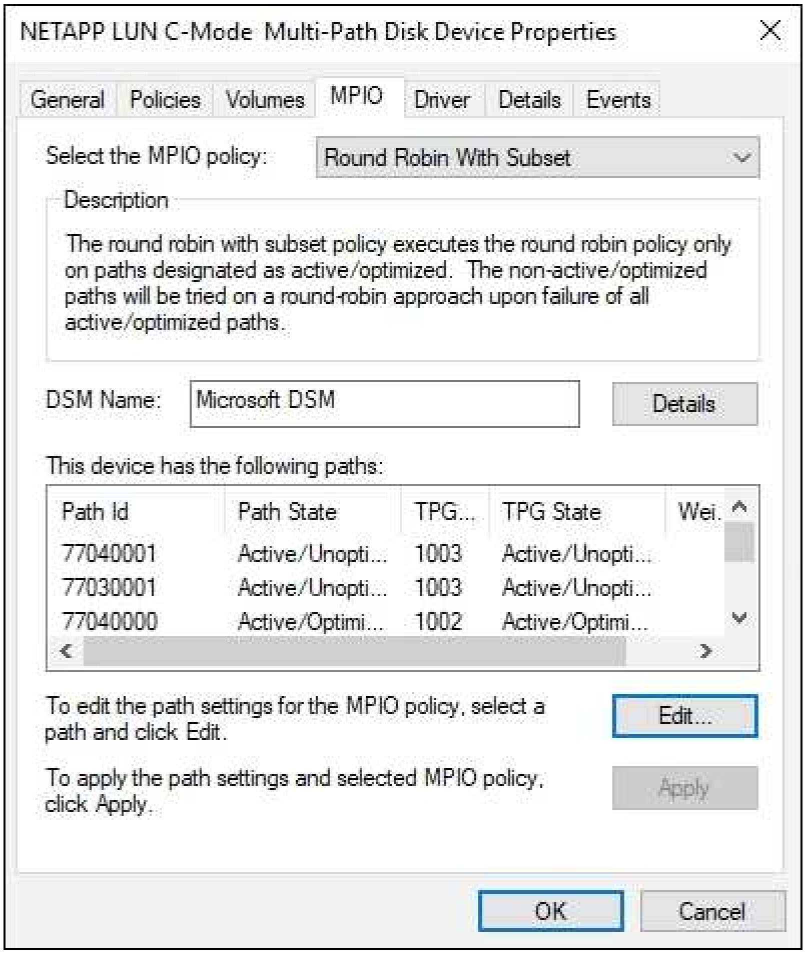

All SAN array configuration

For All SAN Array (ASA) configuration, there should be one group of paths with single priorities. All paths are active/optimized; that is, they are serviced by the controller and that the I/O is sent on all the active paths.

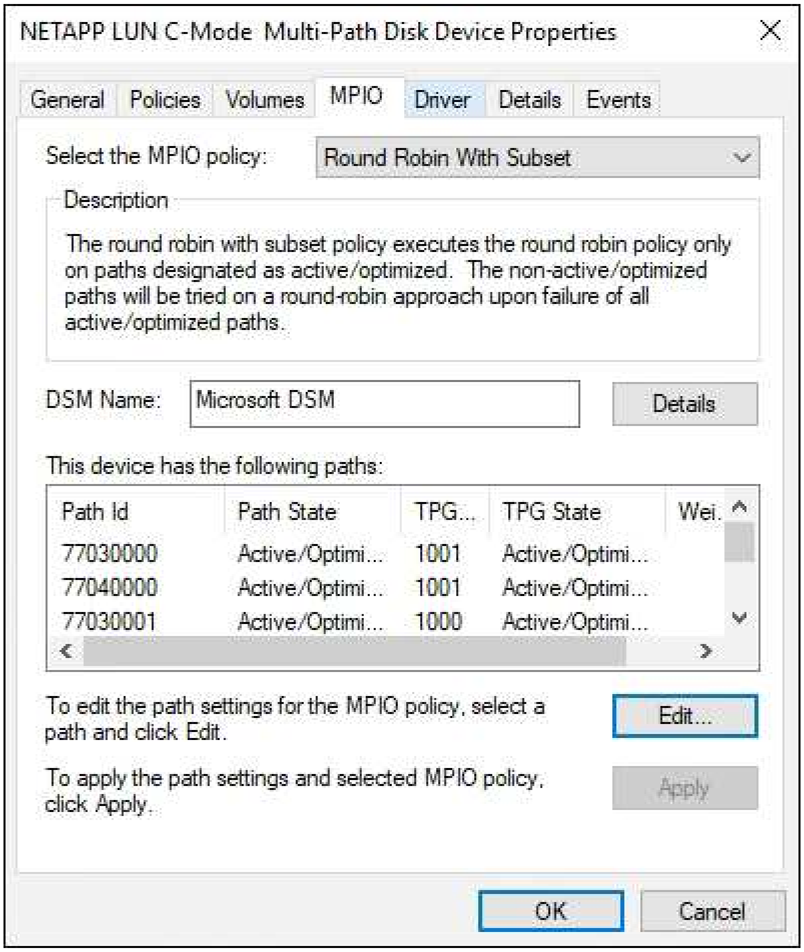

Do not use an excessive number of paths to a single LUN. No more than four paths should be required. More than eight paths might cause path issues during storage failures.

Hyper-V VHD requires alignment for best performance

If the data block boundaries of a disk partition do not align with the block boundaries of the underlying LUN, the storage system often has to complete two block reads or writes for every operating system block read or write. The additional block reads and writes caused by the misalignment might create serious performance problems.

Misalignment is caused by the location of the starting sector for each partition defined by the master boot record.

Partitions created by Windows Server 2016 should be aligned by default.

Use the `Get-NaVirtualDiskAlignment` cmdlet in the ONTAP PowerShell Toolkit to check whether partitions are aligned with underlying LUNs. If the partitions are incorrectly aligned, use the `Repair-NaVirtualDiskAlignment` cmdlet to create a new VHD file with the correct alignment. This cmdlet copies all of the partitions to the new file. The original VHD file is not modified or deleted. The virtual machine must be shut down while the data is copied.

You can download the ONTAP PowerShell Toolkit at NetApp Communities. You must unzip the `DataONTAP.zip` file into the location specified by the environment variable `%PSModulePath%` (or use the `Install.ps1` script to do it for you). Once you have completed the installation, use the `Show-NaHelp` cmdlet to get help for the cmdlets.

The PowerShell Toolkit supports only fixed-size VHD files with MBR-type partitions. VHDs using Windows dynamic disks or GPT partitions are not supported. In addition, the PowerShell Toolkit requires a minimum partition size of 4 GB. Smaller partitions cannot be correctly aligned.

> For Linux virtual machines using the GRUB boot loader on a VHD, you need to update the boot configuration after running the PowerShell Toolkit.

Reinstall GRUB for Linux guests after correcting MBR alignment with PowerShell Toolkit

After running `mbralign` on disks for correcting MBR alignment with PowerShell Toolkit on Linux guest operating systems using the GRUB boot loader, you must reinstall GRUB to ensure that the guest operating system boots correctly.

The PowerShell Toolkit cmdlet has completed on the VHD file for the virtual machine. This topic applies only to Linux guest operating systems using the GRUB boot loader and `SystemRescueCd`.

1. Mount the ISO image of Disk 1 of the installation CDs for the correct version of Linux for the virtual machine.
2. Open the console for the virtual machine in Hyper-V Manager.
3. If the VM is running and hung at the GRUB screen, click in the display area to make sure it is active, then click the **Ctrl-Alt-Delete** toolbar icon to reboot the VM. If the VM is not running, start it, and then immediately click in the display area to make sure it is active.
4. As soon as you see the VMware BIOS splash screen, press the **Esc** key once. The boot menu displays.
5. At the boot menu, select **CD-ROM**.
6. At the Linux boot screen, enter: `linux rescue`
7. Take the defaults for Anaconda (the blue/red configuration screens). Networking is optional.
8. Launch GRUB by entering: `grub`
9. If there is only one virtual disk in this VM, or if there are multiple disks, but the first is the boot disk, run the following GRUB commands:

```
root (hd0,0)
setup (hd0)
quit
```

If you have multiple virtual disks in the VM, and the boot disk is not the first disk, or you are fixing GRUB by booting from the misaligned backup VHD, enter the following command to identify the boot disk:

```
find /boot/grub/stage1
```

Then run the following commands:

```
root (boot_disk,0)
setup (boot_disk)
quit
```

Note that `boot_disk`, above, is a placeholder for the actual disk identifier of the boot disk.

10. Press **Ctrl-D** to log out.

Linux rescue shuts down and then reboots.

Recommended settings

On systems using FC, the following timeout values for Emulex and QLogic FC HBAs are required when MPIO is selected.

For Emulex Fibre Channel HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

For QLogic Fibre Channel HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

Windows Unified Host Utility will set these values. For detailed recommended settings, refer to the Windows Host Utilities documentation and select the installation procedure for your Windows Unified Host Utilities version.

Known issues

There are no known issues for the Windows Server 2012 R2 with ONTAP release.

Configure hosts with NVMe-oF

Overview

You can configure certain SAN hosts for the NVMe over Fabrics (NVMe-oF) protocol, which includes NVMe over Fibre Channel (NVMe/FC) and NVMe over TCP (NVMe/TCP), with ONTAP as the target. Depending on your host operating system and ONTAP version, you configure and validate the NVMe/FC or NVMe/TCP protocol, or both on the host.

NVMe/FC Host Configuration for AIX with ONTAP

You can enable NVMe over Fibre Channel (NVMe/FC) on IBM AIX and VIOS/PowerVM hosts using ONTAP storage as the target. For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

The following support is available for NVMe/FC host configuration for an AIX host with ONTAP:

- Beginning with ONTAP 9.13.1, NVMe/FC support is added for IBM AIX 7.2 TL5 SP6, AIX 7.3 TL1 SP2, and VIOS 3.1.4.21 releases with SAN boot support for both physical and virtual stacks. See the IBM documentation for more information on setting up SAN boot support.
- NVMe/FC is supported with Power9 and Power10 IBM servers.
- No separate PCM (Path Control Module), such as Host Utilities for AIX SCSI Multipath I/O (MPIO) support, is required for NVMe devices.
- Virtualization support with NetApp (VIOS/PowerVM) is introduced with VIOS 3.1.4.21. This is *only* supported through NPIV (N_PortID Virtualization) storage virtualization mode using the Power10 IBM server.

What you'll need

- Verify that you have 32GB FC Emulex adapters (EN1A, EN1B, EN1L, EN1M) or 64GB FC adapters (EN1N, EN1P) with adapter firmware 12.4.257.30 and later versions.
- If you have a MetroCluster configuration, NetApp recommends changing the AIX NVMe/FC default APD (All Path Down) time for supporting MetroCluster unplanned switchover events to avoid the AIX operating system enforcing a shorter I/O timeout. For additional information and the recommended changes to default settings, refer to public report 1553249.
- By default, the Asymmetric Namespace Access Transition Timeout (ANATT) value for the AIX host OS is 30 seconds. IBM provides an Interim Fix (ifix) which caps the ANATT value at 60 seconds; you need to install an ifix from the IBM website to ensure that all the ONTAP workflows are non-disruptive.

  For NVMe/FC AIX support, you must install an ifix on the GA versions of AIX OS. This is not required for the VIOS/PowerVM OS.

  The ifix details are as follows:

  - For AIX level 72-TL5-SP6-2320, install the `IJ46710s6a.230509.epkg.Z` package.
  - For AIX level 73-TL1-SP2-2320, install the `IJ46711s2a.230509.epkg.Z` package.

    For more information on managing ifixes, see Managing Interim Fixes on AIX.

> ⓘ You need to install the ifixes on an AIX version with no previously installed ifixes related to `devices.pciex.pciexclass.010802.rte` on the system. If these ifixes are present, they will conflict with the new installation.

The following table demonstrates HBAs assigned to the AIX LPAR (AIX Logical Partition) or the physical stack:

Host OS	Power Arch	Power FW version	Mode	Comments
AIX 7.2 TL5 SP6	Power9	FW 950 or later	Physical stack	ifix available through TS012877410.
	Power10	FW 1010 or later	Physical stack	SAN booting is supported. ifix available through TS012877410.
AIX 7.3 TL1 SP2	Power9	FW 950 or later	Physical stack	ifix available through TS012877410.
	Power10	FW 1010 or later	Physical and virtual stack	ifix available through TS012877410.

The following table demonstrates HBAs assigned to the VIOS with NPIV-enabled support in a virtualized mode:

Host OS	Power Arch	Power FW version	Mode	Comments
VIOS/PowerVM 3.1.4.21	Power10	FW 1010 or later	Virtual stack	Support starts from AIX 7.3 TL1 SP2 for VIOC

Known limitations

NVMe/FC host configuration for AIX with ONTAP has the following known limitations:

- Qlogic/Marvel 32G FC HBAs on an AIX host does not support NVMe/FC.
- SAN boot is not supported for NVMe/FC devices using Power9 IBM server.

Multipathing

IBM MPIO (Multi Path I/O), used for NVMe multipathing, is provided by default when you install the AIX OS.

You can verify that NVMe multipathing is enabled for an AIX host by using the `lsmpio` command:

```
#[root@aix_server /]: lsmpio -l hdisk1
```

Example output

```
name     path_id  status   path_status  parent  connection
hdisk1  8         Enabled  Sel,Opt       nvme12  fcnvme0, 9
hdisk1  9         Enabled  Sel,Non       nvme65  fcnvme1, 9
hdisk1  10        Enabled  Sel,Opt       nvme37  fcnvme1, 9
hdisk1  11        Enabled  Sel,Non       nvme60  fcnvme0, 9
```

Configure NVMe/FC

You can use the following procedure to configure NVMe/FC for Broadcom/Emulex adapters.

Steps

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.
2. By default, the NVMe/FC protocol support is enabled in the physical FC; however, the NVMe/FC protocol support is disabled in the Virtual Fibre Channel (VFC) on Virtual I/O Server (VIOS).

   Retrieve a list of virtual adapters:

   ```
   $ lsmap -all -npiv
   ```

   Example output

   ```
   Name          Physloc                            ClntID ClntName
   ClntOS
   ------------- ---------------------------------- ------ --------------
   -------
   vfchost0      U9105.22A.785DB61-V2-C2                 4 s1022-iop-mcc-
   AIX
   Status:LOGGED_IN
   FC name:fcs4                  FC loc code:U78DA.ND0.WZS01UY-P0-C7-T0
   Ports logged in:3
   Flags:0xea<LOGGED_IN,STRIP_MERGE,SCSI_CLIENT,NVME_CLIENT>
   VFC client name:fcs0          VFC client DRC:U9105.22A.785DB61-V4-C2
   ```

3. Enable support for the NVMe/FC protocol on an adapter by running the `ioscli vfcctrl` command on the VIOS:

   ```
   $  vfcctrl -enable -protocol nvme -vadapter vfchost0
   ```

   Example output

```
The "nvme" protocol for "vfchost0" is enabled.
```

4. Verify that the support has been enabled on the adapter:

```
# lsattr -El vfchost0
```

Example output

```
alt_site_wwpn        WWPN to use - Only set after migration   False
current_wwpn  0      WWPN to use - Only set after migration   False
enable_nvme   yes    Enable or disable NVME protocol for NPIV True
label                User defined label                       True
limit_intr    false  Limit NPIV Interrupt Sources             True
map_port      fcs4   Physical FC Port                         False
num_per_nvme  0      Number of NPIV NVME queues per range     True
num_per_range 0      Number of NPIV SCSI queues per range     True
```

5. Enable NVMe/Fc protocol for all current adapters or selected adapters:
   a. Enable the NVMe/Fc protocol for all adapters:
      i. Change the `dflt_enabl_nvme` attribute value of `viosnpiv0` pseudo device to `yes`.
      ii. Set the `enable_nvme` attribute value to `yes` for all the VFC host devices.

      ```
      # chdev -l viosnpiv0 -a dflt_enabl_nvme=yes
      ```

      ```
      # lsattr -El viosnpiv0
      ```

      Example output

```
bufs_per_cmd    10  NPIV Number of local bufs per cmd
True
dflt_enabl_nvme yes Default NVME Protocol setting for a new NPIV
adapter True
num_local_cmds  5   NPIV Number of local cmds per channel
True
num_per_nvme    8   NPIV Number of NVME queues per range
True
num_per_range   8   NPIV Number of SCSI queues per range
True
secure_va_info  no  NPIV Secure Virtual Adapter Information
True
```

b. Enable the NVMe/Fc protocol for selected adapters by changing the `enable_nvme` value of the VFC host device attribute to `yes`.

6. Verify that `FC-NVMe Protocol Device` has been created on the server:

```
# [root@aix_server /]: lsdev |grep fcnvme
```

Exmaple output

```
fcnvme0       Available 00-00-02    FC-NVMe Protocol Device
fcnvme1       Available 00-01-02    FC-NVMe Protocol Device
```

7. Record the host NQN from the server:

```
# [root@aix_server /]: lsattr -El fcnvme0
```

Example output

```
attach     switch
How this adapter is connected  False
autoconfig available
Configuration State            True
host_nqn   nqn.2014-08.org.nvmexpress:uuid:64e039bd-27d2-421c-858d-
8a378dec31e8 Host NQN (NVMe Qualified Name) True
```

```
[root@aix_server /]: lsattr -El fcnvme1
```

Example output

```
attach     switch
How this adapter is connected  False
autoconfig available
Configuration State            True
host_nqn   nqn.2014-08.org.nvmexpress:uuid:64e039bd-27d2-421c-858d-
8a378dec31e8 Host NQN (NVMe Qualified Name) True
```

8. Check the host NQN and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_s922-55-lpar2
```

Example output

```
Vserver         Subsystem                Host NQN
------- ---------
----------------------------------------------------------
vs_s922-55-lpar2 subsystem_s922-55-lpar2 nqn.2014-
08.org.nvmexpress:uuid:64e039bd-27d2-421c-858d-8a378dec31e8
```

9. Verify that the initiator ports are up and running and you can see the target LIFs.

Validate NVMe/FC

You need to verify that the ONTAP namespaces correctly reflect on the host. Run the following command to do so:

```
# [root@aix_server /]: lsdev -Cc disk |grep NVMe
```

Example output

```
hdisk1  Available 00-00-02 NVMe 4K Disk
```

You can check the multipathing status:

```
#[root@aix_server /]: lsmpio -l hdisk1
```

Example output

```
name     path_id  status   path_status  parent  connection
hdisk1  8        Enabled  Sel,Opt      nvme12  fcnvme0, 9
hdisk1  9        Enabled  Sel,Non      nvme65  fcnvme1, 9
hdisk1  10       Enabled  Sel,Opt      nvme37  fcnvme1, 9
hdisk1  11       Enabled  Sel,Non      nvme60  fcnvme0, 9
```

Known issues

The NVMe/FC host configuration for AIX with ONTAP has the following known issues:

Burt ID	Title	Description
1553249	AIX NVMe/FC default APD time to be modified for supporting MCC Unplanned Switchover events	By default, AIX operating systems use an all path down (APD) timeout value of 20sec for NVMe/FC. However, ONTAP MetroCluster automatic unplanned switchover (AUSO) and TieBreaker initiated switchover workflows might take a little longer than the APD timeout window, causing I/O errors.
1546017	AIX NVMe/FC caps ANATT at 60s, instead of 120s as advertised by ONTAP	ONTAP advertises the ANA(asymmetric namespace access) transition timeout in controller identify at 120sec. Currently, with ifix, AIX reads the ANA transition timeout from controller identify, but effectively clamps it to 60sec if it is over that limit.
1541386	AIX NVMe/FC hits EIO after ANATT expiry	For any storage failover (SFO) events, if the ANA(asymmetric namespace access) transitioning exceeds the ANA transition timeout cap on a given path, the AIX NVMe/FC host fails with an I/O error despite having alternate healthy paths available to the namespace.
1541380	AIX NVMe/FC waits for half/full ANATT to expire before resuming I/O after ANA AEN	IBM AIX NVMe/FC does not support some Asynchronous notifications (AENs) that ONTAP publishes. This sub-optimal ANA handling will result in sub optimal performance during SFO operations.

Troubleshoot

Before troubleshooting any NVMe/FC failures, verify that you are running a configuration that is compliant to the Interoperability Matrix Tool (IMT) specifications. If you are still facing issues, contact NetApp support for further triage.

ESXi

NVMe-oF Host Configuration for ESXi 8.x with ONTAP

You can configure NVMe over Fabrics (NVMe-oF) on initiator hosts running ESXi 8.x and ONTAP as the target.

Supportability

- Beginning with ONTAP 9.10.1, NVMe/TCP protocol is supported for ONTAP.
- Beginning with ONTAP 9.9.1 P3, NVMe/FC protocol is supported for ESXi 8 and later.

Features

- ESXi initiator hosts can run both NVMe/FC and FCP traffic through the same adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. See the NetApp Interoperability Matrix Tool for the most current list of supported configurations and versions.
- For ESXi 8.0 and later releases, HPP (high performance plugin) is the default plugin for NVMe devices.

Known limitations

- RDM mapping is not supported.

Enable NVMe/FC

NVMe/FC is enabled by default in vSphere releases.

Verify host NQN

You must check the ESXi host NQN string and verify that it matches with the host NQN string for the corresponding subsystem on the ONTAP array.

```
# esxcli nvme info get
```

Example output:

```
Host NQN: nqn.2014-08.org.nvmexpress:uuid:62a19711-ba8c-475d-c954-
0000c9f1a436
```

```
# vserver nvme subsystem host show -vserver nvme_fc
```

Example output:

```
Vserver Subsystem Host NQN
------- ---------
----------------------------------------------------------
nvme_fc nvme_ss  nqn.2014-08.org.nvmexpress:uuid:62a19711-ba8c-475d-c954-
0000c9f1a436
```

If the host NQN strings do not match, you should use the `vserver nvme subsystem host add` command to update the correct host NQN string on your corresponding ONTAP NVMe subsystem.

Configure Broadcom/Emulex and Marvell/Qlogic

The `lpfc` driver and the `qlnativefc` driver in vSphere 8.x have the NVMe/FC capability enabled by default.

See NetApp Interoperability Matrix Tool to check whether the configuration is supported with the driver or firmware.

Validate NVMe/FC

You can use the following procedure to validate NVMe/FC.

Steps

1. Verify that the NVMe/FC adapter is listed on the ESXi host:

   ```
   # esxcli nvme adapter list
   ```

   Example output:

   ```
   Adapter  Adapter Qualified Name           Transport Type  Driver
   Associated Devices
   -------  -------------------------------  --------------  ----------
   ------------------
   vmhba64  aqn:lpfc:100000109b579f11        FC              lpfc
   vmhba65  aqn:lpfc:100000109b579f12        FC              lpfc
   vmhba66  aqn:qlnativefc:2100f4e9d456e286  FC              qlnativefc
   vmhba67  aqn:qlnativefc:2100f4e9d456e287  FC              qlnativefc
   ```

2. Verify that the NVMe/FC namespaces are correctly created:

   The UUIDs in the following example represent the NVMe/FC namespace devices.

```
# esxcfg-mpath -b
uuid.116cb7ed9e574a0faf35ac2ec115969d : NVMe Fibre Channel Disk
(uuid.116cb7ed9e574a0faf35ac2ec115969d)
   vmhba64:C0:T0:L5 LUN:5 state:active fc Adapter: WWNN:
20:00:00:24:ff:7f:4a:50 WWPN: 21:00:00:24:ff:7f:4a:50  Target: WWNN:
20:04:d0:39:ea:3a:b2:1f WWPN: 20:05:d0:39:ea:3a:b2:1f
   vmhba64:C0:T1:L5 LUN:5 state:active fc Adapter: WWNN:
20:00:00:24:ff:7f:4a:50 WWPN: 21:00:00:24:ff:7f:4a:50  Target: WWNN:
20:04:d0:39:ea:3a:b2:1f WWPN: 20:07:d0:39:ea:3a:b2:1f
   vmhba65:C0:T1:L5 LUN:5 state:active fc Adapter: WWNN:
20:00:00:24:ff:7f:4a:51 WWPN: 21:00:00:24:ff:7f:4a:51  Target: WWNN:
20:04:d0:39:ea:3a:b2:1f WWPN: 20:08:d0:39:ea:3a:b2:1f
   vmhba65:C0:T0:L5 LUN:5 state:active fc Adapter: WWNN:
20:00:00:24:ff:7f:4a:51 WWPN: 21:00:00:24:ff:7f:4a:51  Target: WWNN:
20:04:d0:39:ea:3a:b2:1f WWPN: 20:06:d0:39:ea:3a:b2:1f
```

In ONTAP 9.7, the default block size for an NVMe/FC namespace is 4K. This default size is not compatible with ESXi. Therefore, when creating namespaces for ESXi, you must set the namespace block size as **512B**. You can do this using the `vserver nvme namespace create` command.

Example,

```
vserver nvme namespace create -vserver vs_1 -path
/vol/nsvol/namespace1 -size 100g -ostype vmware -block-size 512B
```

Refer to the ONTAP 9 Command man pages for additional details.

3. Verify the status of the individual ANA paths of the respective NVMe/FC namespace devices:

```
# esxcli storage hpp path list -d uuid.df960bebb5a74a3eaaa1ae55e6b3411d

fc.20000024ff7f4a50:21000024ff7f4a50-
fc.2004d039ea3ab21f:2005d039ea3ab21f-
uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Runtime Name: vmhba64:C0:T0:L3
   Device: uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Device Display Name: NVMe Fibre Channel Disk
(uuid.df960bebb5a74a3eaaa1ae55e6b3411d)
   Path State: active unoptimized
   Path Config: {ANA_GRP_id=4,ANA_GRP_state=ANO,health=UP}

fc.20000024ff7f4a51:21000024ff7f4a51-
fc.2004d039ea3ab21f:2008d039ea3ab21f-
uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Runtime Name: vmhba65:C0:T1:L3
   Device: uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Device Display Name: NVMe Fibre Channel Disk
(uuid.df960bebb5a74a3eaaa1ae55e6b3411d)
   Path State: active
   Path Config: {ANA_GRP_id=4,ANA_GRP_state=AO,health=UP}

fc.20000024ff7f4a51:21000024ff7f4a51-
fc.2004d039ea3ab21f:2006d039ea3ab21f-
uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Runtime Name: vmhba65:C0:T0:L3
   Device: uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Device Display Name: NVMe Fibre Channel Disk
(uuid.df960bebb5a74a3eaaa1ae55e6b3411d)
   Path State: active unoptimized
   Path Config: {ANA_GRP_id=4,ANA_GRP_state=ANO,health=UP}

fc.20000024ff7f4a50:21000024ff7f4a50-
fc.2004d039ea3ab21f:2007d039ea3ab21f-
uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Runtime Name: vmhba64:C0:T1:L3
   Device: uuid.df960bebb5a74a3eaaa1ae55e6b3411d
   Device Display Name: NVMe Fibre Channel Disk
(uuid.df960bebb5a74a3eaaa1ae55e6b3411d)
   Path State: active
   Path Config: {ANA_GRP_id=4,ANA_GRP_state=AO,health=UP}
```

Configure NVMe/TCP

In ESXi 8.x, the required NVMe/TCP modules are loaded by default. To configure the network and the

NVMe/TCP adapter, refer to the VMware vSphere documentation.

Validate NVMe/TCP

You can use the following procedure to validate NVMe/TCP.

Steps

1. Verify the status of the NVMe/TCP adapter:

   ```
   esxcli nvme adapter list
   ```

   Example output:

   ```
   Adapter  Adapter Qualified Name           Transport Type  Driver
   Associated Devices
   -------  -------------------------------  --------------  -------
   ------------------
   vmhba65  aqn:nvmetcp:ec-2a-72-0f-e2-30-T  TCP             nvmetcp
   vmnic0
   vmhba66  aqn:nvmetcp:34-80-0d-30-d1-a0-T  TCP             nvmetcp
   vmnic2
   vmhba67  aqn:nvmetcp:34-80-0d-30-d1-a1-T  TCP             nvmetcp
   vmnic3
   ```

2. Retrieve a list of NVMe/TCP connections:

   ```
   esxcli nvme controller list
   ```

   Example output:

```
Name                                                        Controller Number
Adapter  Transport Type  Is Online  Is VVOL
---------------------------------------------------------------------------------
---------------------------------  -----------------  -------
nqn.2014-08.org.nvmexpress.discovery#vmhba64#192.168.100.166:8009  256
vmhba64  TCP                     true    false
nqn.1992-
08.com.netapp:sn.89bb1a28a89a11ed8a88d039ea263f93:subsystem.nvme_ss#vmhb
a64#192.168.100.165:4420 258  vmhba64  TCP  true    false
nqn.1992-
08.com.netapp:sn.89bb1a28a89a11ed8a88d039ea263f93:subsystem.nvme_ss#vmhb
a64#192.168.100.168:4420 259  vmhba64  TCP  true    false
nqn.1992-
08.com.netapp:sn.89bb1a28a89a11ed8a88d039ea263f93:subsystem.nvme_ss#vmhb
a64#192.168.100.166:4420 260  vmhba64  TCP  true    false
nqn.2014-08.org.nvmexpress.discovery#vmhba64#192.168.100.165:8009  261
vmhba64  TCP                     true    false
nqn.2014-08.org.nvmexpress.discovery#vmhba65#192.168.100.155:8009  262
vmhba65  TCP                     true    false
nqn.1992-
08.com.netapp:sn.89bb1a28a89a11ed8a88d039ea263f93:subsystem.nvme_ss#vmhb
a64#192.168.100.167:4420 264  vmhba64  TCP  true    false
```

3. Retrieve a list of the number of paths to an NVMe namespace:

```
esxcli storage hpp path list -d uuid.f4f14337c3ad4a639edf0e21de8b88bf
```

Example output:

```
tcp.vmnic2:34:80:0d:30:ca:e0-tcp.192.168.100.165:4420-
uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Runtime Name: vmhba64:C0:T0:L5
   Device: uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Device Display Name: NVMe TCP Disk
(uuid.f4f14337c3ad4a639edf0e21de8b88bf)
   Path State: active
   Path Config: {ANA_GRP_id=6,ANA_GRP_state=AO,health=UP}

tcp.vmnic2:34:80:0d:30:ca:e0-tcp.192.168.100.168:4420-
uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Runtime Name: vmhba64:C0:T3:L5
   Device: uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Device Display Name: NVMe TCP Disk
(uuid.f4f14337c3ad4a639edf0e21de8b88bf)
   Path State: active unoptimized
   Path Config: {ANA_GRP_id=6,ANA_GRP_state=ANO,health=UP}

tcp.vmnic2:34:80:0d:30:ca:e0-tcp.192.168.100.166:4420-
uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Runtime Name: vmhba64:C0:T2:L5
   Device: uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Device Display Name: NVMe TCP Disk
(uuid.f4f14337c3ad4a639edf0e21de8b88bf)
   Path State: active unoptimized
   Path Config: {ANA_GRP_id=6,ANA_GRP_state=ANO,health=UP}

tcp.vmnic2:34:80:0d:30:ca:e0-tcp.192.168.100.167:4420-
uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Runtime Name: vmhba64:C0:T1:L5
   Device: uuid.f4f14337c3ad4a639edf0e21de8b88bf
   Device Display Name: NVMe TCP Disk
(uuid.f4f14337c3ad4a639edf0e21de8b88bf)
   Path State: active
   Path Config: {ANA_GRP_id=6,ANA_GRP_state=AO,health=UP}
```

Known issues

The NVMe-oF host configuration for ESXi 8.x with ONTAP has the following known issues:

NetApp Bug ID	Title	Description
1420654	ONTAP node non-operational when NVMe/FC protocol is used with ONTAP version 9.9.1	ONTAP 9.9.1 has introduced support for the NVMe "abort" command. When ONTAP receives the "abort" command to abort an NVMe fused command that is waiting for its partner command, an ONTAP node disruption occurs. The issue is noticed only with hosts that use NVMe fused commands (for example, ESX) and Fibre Channel (FC) transport.
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

TR-4597-VMware vSphere with ONTAP
VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
VMware vSphere 6.x and 7.x support with NetApp® SnapMirror® Business Continuity (SM-BC)

NVMe-oF Host Configuration for ESXi 7.x with ONTAP

Supportability

- Beginning with ONTAP 9.7, NVMe over Fibre Channel (NVMe/FC) support is added for VMWare vSphere releases.
- Beginning with 7.0U3c, NVMe/TCP feature is supported for ESXi Hypervisor.
- Beginning with ONTAP 9.10.1, NVMe/TCP feature is supported for ONTAP.

Features

- ESXi initiator host can run both NVMe/FC and FCP traffic through the same adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. See the NetApp Interoperability Matrix for the most current list of supported configurations and versions.
- Beginning with ONTAP 9.9.1 P3, NVMe/FC feature is supported for ESXi 7.0 update 3.
- For ESXi 7.0 and later releases, HPP (high performance plugin) is the default plugin for NVMe devices.

Known limitations

The following configurations are not supported:

- RDM mapping

* VVols

Enable NVMe/FC

1. Check the ESXi host NQN string and verify that it matches with the host NQN string for the corresponding subsystem on the ONTAP array:

```
# esxcli nvme  info get
Host NQN: nqn.2014-08.com.vmware:nvme:nvme-esx

# vserver nvme subsystem host show -vserver vserver_nvme
  Vserver Subsystem             Host NQN
  ------- ------------------- ----------------------------------------
  vserver_nvme ss_vserver_nvme nqn.2014-08.com.vmware:nvme:nvme-esx
```

Configure Broadcom/Emulex

1. Check whether the configuration is supported with required driver/firmware by referring to NetApp Interoperability Matrix.
2. Set the lpfc driver parameter `lpfc_enable_fc4_type=3` for enabling NVMe/FC support in the `lpfc` driver and reboot the host.

> Starting with vSphere 7.0 update 3, the `brcmnvmefc` driver is no longer available. Therefore, the `lpfc` driver now includes the NVMe over Fibre Channel (NVMe/FC) functionality previously delivered with the `brcmnvmefc` driver.

> The `lpfc_enable_fc4_type=3` parameter is set by default for the LPe35000-series adapters. You must perform the following command to set it manually for LPe32000-series and LPe31000-series adapters.

```
# esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3

#esxcli system module parameters list  -m lpfc | grep lpfc_enable_fc4_type
lpfc_enable_fc4_type              int     3      Defines what FC4 types
are supported

#esxcli storage core adapter list
HBA Name  Driver   Link State  UID
Capabilities         Description
--------  -------  ----------  ------------------------------------
-------------------  -----------
vmhba1    lpfc     link-up     fc.200000109b95456f:100000109b95456f
Second Level Lun ID  (0000:86:00.0) Emulex Corporation Emulex LPe36000
Fibre Channel Adapter    FC HBA
vmhba2    lpfc     link-up     fc.200000109b954570:100000109b954570
Second Level Lun ID  (0000:86:00.1) Emulex Corporation Emulex LPe36000
Fibre Channel Adapter    FC HBA
vmhba64   lpfc     link-up     fc.200000109b95456f:100000109b95456f
(0000:86:00.0) Emulex Corporation Emulex LPe36000 Fibre Channel Adapter
NVMe HBA
vmhba65   lpfc     link-up     fc.200000109b954570:100000109b954570
(0000:86:00.1) Emulex Corporation Emulex LPe36000 Fibre Channel Adapter
NVMe HBA
```

Configure Marvell/QLogic

Steps

1. Check whether configuration is supported with required driver/firmware by referring to NetApp Interoperability Matrix.
2. Set the `qlnativefc` driver parameter `ql2xnvmesupport=1` for enabling NVMe/FC support in the `qlnativefc` driver and reboot the host.

   ```
   # esxcfg-module -s 'ql2xnvmesupport=1' qlnativefc
   ```

   > The `qlnativefc` driver parameter is set by default for the Qle 277x-series adapters. You must perform the following command to set it manually for Qle 277x series adapters.

   ```
   esxcfg-module -l | grep qlnativefc
   qlnativefc                4    1912
   ```

3. Check whether nvme is enabled on the adapter:

```
  #esxcli storage core adapter list
HBA Name  Driver      Link State  UID
Capabilities         Description
--------  ----------  ----------  ------------------------------------
-------------------  -----------
 vmhba3    qlnativefc  link-up     fc.20000024ff1817ae:21000024ff1817ae
Second Level Lun ID  (0000:5e:00.0) QLogic Corp QLE2742 Dual Port 32Gb
Fibre Channel to PCIe Adapter    FC Adapter
vmhba4    qlnativefc  link-up     fc.20000024ff1817af:21000024ff1817af
Second Level Lun ID  (0000:5e:00.1) QLogic Corp QLE2742 Dual Port 32Gb
Fibre Channel to PCIe Adapter FC Adapter
vmhba64   qlnativefc  link-up     fc.20000024ff1817ae:21000024ff1817ae
(0000:5e:00.0) QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe
Adapter  NVMe FC Adapter
vmhba65   qlnativefc  link-up     fc.20000024ff1817af:21000024ff1817af
(0000:5e:00.1) QLogic Corp QLE2742 Dual Port 32Gb Fibre Channel to PCIe
Adapter  NVMe FC Adapter
```

Validate NVMe/FC

1. Verify that NVMe/FC adapter is listed on the ESXi host:

```
# esxcli nvme adapter list

Adapter  Adapter Qualified Name           Transport Type  Driver
Associated Devices
-------  -------------------------------  --------------  ----------
------------------
vmhba64  aqn:qlnativefc:21000024ff1817ae  FC              qlnativefc
vmhba65  aqn:qlnativefc:21000024ff1817af  FC              qlnativefc
vmhba66  aqn:lpfc:100000109b579d9c        FC              lpfc
vmhba67  aqn:lpfc:100000109b579d9d        FC              lpfc
```

2. Verify that the NVMe/FC namespaces are properly created:

The UUIDs in the following example represent the NVMe/FC namespace devices.

```
# esxcfg-mpath -b
uuid.5084e29a6bb24fbca5ba076eda8ecd7e : NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
   vmhba65:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:69 WWPN: 21:00:34:80:0d:6d:72:69  Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:2f:00:a0:98:df:e3:d1
   vmhba65:C0:T1:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:69 WWPN: 21:00:34:80:0d:6d:72:69  Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:1a:00:a0:98:df:e3:d1
   vmhba64:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:68 WWPN: 21:00:34:80:0d:6d:72:68  Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:18:00:a0:98:df:e3:d1
   vmhba64:C0:T1:L1 LUN:1 state:active fc Adapter: WWNN:
20:00:34:80:0d:6d:72:68 WWPN: 21:00:34:80:0d:6d:72:68  Target: WWNN:
20:17:00:a0:98:df:e3:d1 WWPN: 20:19:00:a0:98:df:e3:d1
```

In ONTAP 9.7, the default block size for a NVMe/FC namespace is 4K. This default size is not compatible with ESXi. Therefore, when creating namespaces for ESXi, you must set the namespace block size as 512b. You can do this using the `vserver nvme namespace create` command.

Example

`vserver nvme namespace create -vserver vs_1 -path /vol/nsvol/namespace1 -size 100g -ostype vmware -block-size 512B`

Refer to the ONTAP 9 Command man pages for additional details.

3. Verify the status of the individual ANA paths of the respective NVMe/FC namespace devices:

```
esxcli storage hpp path list -d uuid.5084e29a6bb24fbca5ba076eda8ecd7e
fc.200034800d6d7268:210034800d6d7268-
fc.201700a098dfe3d1:201800a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Runtime Name: vmhba64:C0:T0:L1
   Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
   Path State: active
   Path Config: {TPG_id=0,TPG_state=AO,RTP_id=0,health=UP}

fc.200034800d6d7269:210034800d6d7269-
fc.201700a098dfe3d1:201a00a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Runtime Name: vmhba65:C0:T1:L1
   Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
   Path State: active
   Path Config: {TPG_id=0,TPG_state=AO,RTP_id=0,health=UP}

fc.200034800d6d7269:210034800d6d7269-
fc.201700a098dfe3d1:202f00a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Runtime Name: vmhba65:C0:T0:L1
   Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
   Path State: active unoptimized
   Path Config: {TPG_id=0,TPG_state=ANO,RTP_id=0,health=UP}

fc.200034800d6d7268:210034800d6d7268-
fc.201700a098dfe3d1:201900a098dfe3d1-
uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Runtime Name: vmhba64:C0:T1:L1
   Device: uuid.5084e29a6bb24fbca5ba076eda8ecd7e
   Device Display Name: NVMe Fibre Channel Disk
(uuid.5084e29a6bb24fbca5ba076eda8ecd7e)
   Path State: active unoptimized
   Path Config: {TPG_id=0,TPG_state=ANO,RTP_id=0,health=UP}
```

Configure NVMe/TCP

Starting 7.0U3c, the required NVMe/TCP modules will be loaded by default. For configuring the network and the NVMe/TCP adapter, refer to the VMware vSphere documentation.

Validate NVMe/TCP

Steps

1. Verify the status of the NVMe/TCP adapter.

```
[root@R650-8-45:~] esxcli nvme adapter list
Adapter    Adapter Qualified Name
--------- -------------------------------
vmhba64    aqn:nvmetcp:34-80-0d-30-ca-e0-T
vmhba65    aqn:nvmetc:34-80-13d-30-ca-e1-T
list
Transport Type   Driver   Associated Devices
---------------  -------  ------------------
TCP              nvmetcp    vmnzc2
TCP              nvmetcp    vmnzc3
```

2. To list the NVMe/TCP connections, use the following command:

```
[root@R650-8-45:~] esxcli nvme controller list
Name
-----------
nqn.1992-
08.com.netapp:sn.5e347cf68e0511ec9ec2d039ea13e6ed:subsystem.vs_name_tcp_
ss#vmhba64#192.168.100.11:4420
nqn.1992-
08.com.netapp:sn.5e347cf68e0511ec9ec2d039ea13e6ed:subsystem.vs_name_tcp_
ss#vmhba64#192.168.101.11:4420
Controller Number  Adapter   Transport Type   IS Online
----------------- ---------  ---------------  ---------
1580              vmhba64    TCP              true
1588              vmhba65    TCP              true
```

3. To list the number of paths to an NVMe namespace, use the following command:

```
[root@R650-8-45:~] esxcli storage hpp path list -d
uuid.400bf333abf74ab8b96dc18ffadc3f99
tcp.vmnic2:34:80:0d:30:ca:eo-tcp.unknown-
uuid.400bf333abf74ab8b96dc18ffadc3f99
   Runtime Name: vmhba64:C0:T0:L3
   Device: uuid.400bf333abf74ab8b96dc18ffadc3f99
   Device Display Name: NVMe TCP Disk
(uuid.400bf333abf74ab8b96dc18ffadc3f99)
   Path State: active unoptimized
   Path config: {TPG_id=0,TPG_state=ANO,RTP_id=0,health=UP}

tcp.vmnic3:34:80:0d:30:ca:el-tcp.unknown-
uuid.400bf333abf74ab8b96dc18ffadc3f99
   Runtime Name: vmhba65:C0:T1:L3
   Device: uuid.400bf333abf74ab8b96dc18ffadc3f99
   Device Display Name: NVMe TCP Disk
(uuid.400bf333abf74ab8b96dc18ffadc3f99)
   Path State: active
   Path config: {TPG_id=0,TPG_state=AO,RTP_id=0,health=UP}
```

Known issues

The NVMe-oF host configuration for ESXi 7.x with ONTAP has the following known issues:

NetApp Bug ID	Title	Workaround
1420654	ONTAP node non-operational when NVMe/FC protocol is used with ONTAP version 9.9.1	Check and rectify any network issues in the host fabric. If this does not help, upgrade to a patch that fixes this issue.

Related information

TR-4597-VMware vSphere with ONTAP
VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
VMware vSphere 6.x and 7.x support with NetApp® SnapMirror® Business Continuity (SM-BC)

Oracle Linux

OL 9

NVMe-oF Host Configuration for Oracle Linux 9.2 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Oracle Linux (OL) 9.2 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for OL 9.2 with ONTAP:

* Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
* Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

* Oracle Linux 9.2 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported OL 9.2 software versions.

Steps

1. Install OL 9.2 GA on the server. After the installation is complete, verify that you are running the specified OL 9.2 GA kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-101.103.2.1.el9uek.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-2.2.1-2.el9.x86_64
   ```

3. Install the `libnvme` package:

   ```
   #rpm -qa|grep libnvme
   ```

   Example output

```
libnvme-1.2-2.el9.x86_64
```

4. On the Oracle Linux 9.2 host, check the `hostnqn` string at `/etc/nvme/hostnqn`:

   ```
   # cat /etc/nvme/hostnqn
   ```

   Example output:

   ```
   nqn.2014-08.org.nvmexpress:uuid:bc59d14c-47f3-11eb-b93c-3a68dd48673f
   ```

5. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

   ```
   ::> vserver nvme subsystem host show -vserver vs_ol_nvme
   ```

   Example output:

   ```
   Vserver     Subsystem          Host NQN
   ----------- ---------------
   ----------------------------------------------------------
   vs_nvme207   nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:bc59d14c-
   47f3-11eb-b93c-3a68dd48673f
   ```

   If the `hostnqn` strings do not match, you can use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex adapters or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   14.2.455.11, sli-4:2:c
   14.2.455.11, sli-4:2:c
   ```

   ```
   # cat /sys/module/lpfc/version
   0:14.2.0.5
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b3c081f
0x100000109b3c0820
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the OL 9.2 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.12.00 DVR:v10.02.08.100-k
QLE2742 FW:v9.12.00 DVR:v10.02.08.100-k
```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

```
# nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.22

Discovery Log Number of Records 8, Generation counter 18
=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.166.23
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  1
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.166.22
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  2
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.167.23
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
# nvme discover -t tcp -w 192.168.8.1 -a 192.168.8.48
# nvme discover -t tcp -w 192.168.8.1 -a 192.168.8.49
# nvme discover -t tcp -w 192.168.9.1 -a 192.168.9.48
# nvme discover -t tcp -w 192.168.9.1 -a 192.168.9.49
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.8.1 -a 192.168.8.48 -l 1800
# nvme connect-all -t tcp -w 192.168.8.1 -a 192.168.8.49 -l 1800
# nvme connect-all -t tcp -w 192.168.9.1 -a 192.168.9.48 -l 1800
# nvme connect-all -t tcp -w 192.168.9.1 -a 192.168.9.49 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify the following NVMe/FC settings on the OL 9.2 host:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller



Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                   85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_
ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
```

NVMe/TCP

```
nvme list-subsys /dev/nvme1n22
```

Example output

```
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.68c036aaa3cf11edbb95d039ea243511:subsystem.tcp
\
 +- nvme2 tcp
traddr=192.168.8.49,trsvcid=4420,host_traddr=192.168.8.1 live
optimized
 +- nvme3 tcp
traddr=192.168.8.48,trsvcid=4420,host_traddr=192.168.8.1 live
optimized
 +- nvme6 tcp
traddr=192.168.9.49,trsvcid=4420,host_traddr=192.168.9.1 live non-
optimized
 +- nvme7 tcp
traddr=192.168.9.48,trsvcid=4420,host_traddr=192.168.9.1 live non-
optimized
```

4. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device          Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns




NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

There are no known issues.

NVMe-oF Host Configuration for Oracle Linux 9.1 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Oracle Linux (OL) 9.1 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for OL 9.1 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- Oracle Linux 9.1 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported OL 9.1 software versions.

Steps

1. Install OL 9.1 GA on the server. After the installation is complete, verify that you are running the specified OL 9.1 GA kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-3.60.5.1.el9uek.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-2.0-4.el9.x86_64
   ```

3. Install the `libnvme` package:

   ```
   #rpm -qa|grep libnvme
   ```

Example output

```
libnvme-1.0-5.el9.x86_64.rpm
```

4. On the Oracle Linux 9.1 host, check the `hostnqn` string at `/etc/nvme/hostnqn`:

```
# cat /etc/nvme/hostnqn
```

Example output:

```
nqn.2014-08.org.nvmexpress:uuid:bc59d14c-47f3-11eb-b93c-3a68dd48673f
```

5. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_ol_nvme
```

Example output:

```
Vserver     Subsystem           Host NQN
----------- ---------------
----------------------------------------------------------
vs_ol_nvme  nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:bc59d14c-
47f3-11eb-b93c-3a68dd48673f
```

> If the `hostnqn` strings do not match, you can use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex adapters or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   12.8.614.23, sli-4:2:c
   12.8.614.23, sli-4:2:c


   # cat /sys/module/lpfc/version
   0:14.0.0.1
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b3c081f
0x100000109b3c0820

# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the OL 9.1 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.18.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.18.02 DVR:v10.02.00.106-k
```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
#  nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
Discovery Log Number of Records 6, Generation counter 8
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 0
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.17
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 1
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.5.17
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 2
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.15
sectype: none
=====Discovery Log Entry 3======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
```

```
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:subsystem.host_95
traddr: 192.168.6.17
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.15
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.17
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.17
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.15 -l 1800
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.17 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.15 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.17 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify the following NVMe/FC settings on the OL 9.1 host:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_
ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
```

NVMe/TCP

```
nvme list-subsys /dev/nvme1n22
```

Example output

```
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.68c036aaa3cf11edbb95d039ea243511:subsystem.tcp
\
+- nvme2 tcp
traddr=192.168.8.49,trsvcid=4420,host_traddr=192.168.8.1 live
optimized
+- nvme3 tcp
traddr=192.168.8.48,trsvcid=4420,host_traddr=192.168.8.1 live
optimized
+- nvme6 tcp
traddr=192.168.9.49,trsvcid=4420,host_traddr=192.168.9.1 live non-
optimized
+- nvme7 tcp
traddr=192.168.9.48,trsvcid=4420,host_traddr=192.168.9.1 live non-
optimized
```

4. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device         Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns


NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 9.1 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1536937	`nvme list-subsys` command prints repeated NVMe controllers for a subsystem	The `nvme list-subsys` command should return a unique list of NVMe controllers associated with a given subsystem. In Oracle Linux 9.1, the `nvme list-subsys` command returns NVMe controllers with the respective asymmetric namespace access (ANA) state for all namespaces that belong to a given subsystem. However, it would be useful to display unique NVMe controller entries with the path state if you list the subsystem command syntax for a given namespace because the ANA state is a per-namespace attribute.	17998
1539101	Oracle Linux 9.1 NVMe-oF hosts fail to create a persistent discovery controller	On Oracle Linux 9.1 NVMe-oF hosts, you can use the `nvme discover -p` command to create Persistent Discovery Controllers (PDCs). When this command is used, one PDC should be created per initiator-target combination. However, if you are running Oracle Linux 9.1 on an NVMe-oF host, PDC creation fails when the `nvme discover -p` command is executed.	18196

NVMe/FC Host Configuration for Oracle Linux 9.0 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe/FC and other transports, is supported with Oracle Linux (OL) 9.0 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

* Oracle Linux 9.0 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can validate the minimum supported OL 9.0 software versions using the following procedure.

Steps

1. Install OL 9.0 GA on the server. After the installation is complete, verify that you are running the specified OL 9.0 GA kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-0.30.19.el9uek.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-1.16-3.el9.x86_64
   ```

3. On the Oracle Linux 9.0 host, check the `hostnqn` string at `/etc/nvme/hostnqn`:

   ```
   # cat /etc/nvme/hostnqn
   ```

   Example output:

   ```
   nqn.2014-08.org.nvmexpress:uuid:4c4c4544-0032-3310-8033-b8c04f4c5132
   ```

4. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_ol_nvme
```

Example output:

```
Vserver     Subsystem          Host NQN
----------- ---------------
-----------------------------------------------------------
vs_ol_nvme  nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:4c4c4544-
0032-3310-8033-b8c04f4c5132
```

> If the `hostnqn` strings do not match, you can use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model.

```
# cat /sys/class/scsi_host/host*/modelname
```

Example output:

```
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
```

Example output:

```
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.11, sli-4:2:c
14.0.505.11, sli-4:2:c

# cat /sys/module/lpfc/version
0:12.8.0.11
```

For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the OL 9.0 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
```

2. Verify that `ql2xnvmeenable` is set which enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
#  nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
Discovery Log Number of Records 6, Generation counter 8
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 0
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.17
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 1
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.5.17
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 2
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.15
sectype: none
=====Discovery Log Entry 3======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:subsystem.host_95
```

```
traddr: 192.168.6.17
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

   ```
   # nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.15
   # nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.17
   # nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
   # nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.17
   ```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

   ```
   nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
   ```

   Example output:

   ```
   # nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.15 -l 1800
   # nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.17 -l 1800
   # nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.15 -l 1800
   # nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.17 -l 1800
   ```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify the following NVMe/FC settings on the OL 9.0 host:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status:

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device         Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns


NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for Oracle Linux 9.0 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 9.0 NVMe-oF Hosts create duplicate Persistent Discovery Controllers	On Oracle Linux 9.0 NVMe over Fabrics (NVMe-oF) hosts, you can use the `nvme discover -p` command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Oracle Linux 9.0 with an NVMe-oF host, a duplicate PDC is created each time `nvme discover -p` is executed. This leads to unnecessary usage of resources on both the host and the target.	18118

OL 8

NVMe-oF Host Configuration for Oracle Linux 8.8 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Oracle Linux (OL) 8.8 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for OL 8.8 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native `nvme-cli` package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Both NVMe and SCSI traffic can be run on the same co-existent host. Therefore, for SCSI LUNs, you can configure dm-multipath for SCSI mpath devices, whereas you might use NVMe multipath to configure NVMe-oF namespace devices on the host.
- There is no sanlun support for NVMe-oF. Therefore, there is no host utility support for NVMe-oF on an OL 8.8 host. You can rely on the NetApp plug-in included in the native `nvme-cli` package for all NVMe-oF transports.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

Oracle Linux 8.8 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported OL 8.8 software versions.

Steps

1. Install OL 8.8 GA on the server. After the installation is complete, verify that you are running the specified OL 8.8 GA kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-101.103.2.1.el8uek.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-1.16-7.el8.x86_64
   ```

3. On the Oracle Linux 8.8 host, check the `hostnqn` string at `/etc/nvme/hostnqn`:

   ```
   # cat /etc/nvme/hostnqn
   ```

   Example output:

   ```
   nqn.2014-08.org.nvmexpress:uuid:bc59d14c-47f3-11eb-b93c-3a68dd48673f
   ```

4. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

   ```
   ::> vserver nvme subsystem host show -vserver vs_ol_nvme
   ```

   Example output:

```
Vserver     Subsystem          Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme207   nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:bc59d14c-
47f3-11eb-b93c-3a68dd48673f
```

If the `hostnqn` strings do not match, you can use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

5. Reboot the host.

If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and `dm-multipath` for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from `dm-multipath` to prevent `dm-multipath` from claiming these namespace devices. You can add the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf

defaults {
  enable_foreign     NONE
}
```

Restart the multipathd daemon by running a `systemctl restart multipathd` command. This allows the new setting to take effect.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex adapters or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   14.2.539.16, sli-4:2:c
   14.2.539.16, sli-4:2:c
   ```

   ```
   # cat /sys/module/lpfc/version
   0:14.2.0.5
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b3c081f
0x100000109b3c0820
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the OL 8.8 GA kernel has the latest fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.12.00 DVR:v10.02.08.100-k
QLE2742 FW:v9.12.00 DVR:v10.02.08.100-k
```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
#  nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
Discovery Log Number of Records 6, Generation counter 8
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 0
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.17
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 1
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.5.17
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: unrecognized
treq: not specified
portid: 2
trsvcid: 8009
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
traddr: 192.168.6.15
sectype: none
=====Discovery Log Entry 3======
trtype: tcp
adrfam: ipv4
```

```
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:subsystem.host_95
traddr: 192.168.6.17
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.15
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.17
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.17
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.15 -l 1800
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.17 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.15 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.17 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller



Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                   85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_
ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
```

NVMe/TCP

```
nvme list-subsys /dev/nvme1n22
```

Example output

```
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.68c036aaa3cf11edbb95d039ea243511:subsystem.tcp
\
+- nvme2 tcp
traddr=192.168.8.49,trsvcid=4420,host_traddr=192.168.8.1 live non-
optimized
+- nvme3 tcp
traddr=192.168.8.48,trsvcid=4420,host_traddr=192.168.8.1 live non-
optimized
+- nvme6 tcp
traddr=192.168.9.49,trsvcid=4420,host_traddr=192.168.9.1 live
optimized
+- nvme7 tcp
traddr=192.168.9.48,trsvcid=4420,host_traddr=192.168.9.1 live
optimized
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device         Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns




NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 8.8 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 8.8 NVMe-oF hosts create duplicate PDCs	On OL 8.8 NVMe-oF hosts, Persistent Discovery Controllers (PDCs) are created by passing the `-p` option to the `nvme discover` command. For a given initiator-target combination, only one PDC is expected to be created with each invocation of the `nvme discover` command. However, beginning with OL 8.x, NVMe-oF hosts end up creating duplicate PDCs with each invocation of the `nvme discover` command with the `-p` option. This wastes resources on both the host and the target.	18118

NVMe-oF Host Configuration for Oracle Linux 8.7 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Oracle Linux (OL) 8.7 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe/FC host configuration for OL 8.7 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native `nvme-cli` package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- OL 8.7 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported OL 8.7 software versions.

Steps

1. Install OL 8.7 GA on the server. After the installation is complete, verify that you are running the specified OL 8.7 GA kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-3.60.5.1.el8uek.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-1.16-5.el8.x86_64
   ```

3. On the Oracle Linux 8.7 host, check the `hostnqn` string at `/etc/nvme/hostnqn`:

   ```
   # cat /etc/nvme/hostnqn
   ```

   Example output:

   ```
   nqn.2014-08.org.nvmexpress:uuid:791c54eb-545d-4ed3-8d41-91a0a53d4b24
   ```

4. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

   ```
   ::> vserver nvme subsystem host show -vserver vs_ol_nvme
   ```

   Example output:

```
Vserver     Subsystem          Host NQN
----------- ---------------
----------------------------------------------------------
vs_ol_nvme  nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:791c54eb-
545d-4ed3-8d41-91a0a53d4b24
```

> If the `hostnqn` strings do not match, you can use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

5. Reboot the host.

> If you intend to run both NVMe and SCSI traffic on the same Oracle Linux 8.7 co-existent host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This also means the ONTAP namespaces should be blacklisted in dm-multipath to prevent dm-multipath from claiming these namespace devices. You can do this by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:
>
> ```
> #cat /etc/multipath.conf
> defaults {
>     enable_foreign  NONE
> }
> ```
>
> Restart the multipathd daemon by running the `systemctl restart multipathd` command to apply the new settings.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model.

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   12.8.614.23, sli-4:2:c
   12.8.614.23, sli-4:2:c


   # cat /sys/module/lpfc/version
   0:14.0.0.1
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b3c081f
0x100000109b3c0820
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b3c081f WWNN x200000109b3c081f DID
x060300 ONLINE
NVME RPORT WWPN x2010d039ea2c3e2d WWNN x200fd039ea2c3e2d DID x061f0e
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2011d039ea2c3e2d WWNN x200fd039ea2c3e2d DID x06270f
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000a71 Cmpl 0000000a71 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000558611c6 Issue 000000005578bb69 OutIO
fffffffffff2a9a3
abort 0000007a noxri 00000000 nondlp 00000447 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000a8e Err 0000e2a8
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b3c0820 WWNN x200000109b3c0820 DID
x060200 ONLINE
NVME RPORT WWPN x2015d039ea2c3e2d WWNN x200fd039ea2c3e2d DID x062e0c
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2014d039ea2c3e2d WWNN x200fd039ea2c3e2d DID x06290f
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000a69 Cmpl 0000000a69 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000055814701 Issue 0000000055744b1c OutIO
fffffffffff3041b
abort 00000046 noxri 00000000 nondlp 0000043f qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000a89 Err 0000e2f3
```

Marvell/Qlogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the OL 8.7 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

   ```
   # cat /sys/class/fc_host/host*/symbolic_name
   ```

   Example output

   ```
   QLE2742 FW:v9.10.11 DVR:v10.02.06.200-k
   QLE2742 FW:v9.10.11 DVR:v10.02.06.200-k
   ```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

   ```
   # cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
   1
   ```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

   ```
   # cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
   256
   ```

> ⓘ This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

   ```
   #  nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
   Discovery Log Number of Records 6, Generation counter 8
   =====Discovery Log Entry 0======
   trtype: tcp
   adrfam: ipv4
   subtype: unrecognized
   treq: not specified
   portid: 0
   trsvcid: 8009
   subnqn: nqn.1992-
   08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
   traddr: 192.168.6.17
   sectype: none
   =====Discovery Log Entry 1======
   trtype: tcp
   adrfam: ipv4
   subtype: unrecognized
   treq: not specified
   portid: 1
   trsvcid: 8009
   subnqn: nqn.1992-
   08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
   traddr: 192.168.5.17
   sectype: none
   =====Discovery Log Entry 2======
   trtype: tcp
   adrfam: ipv4
   subtype: unrecognized
   treq: not specified
   portid: 2
   trsvcid: 8009
   subnqn: nqn.1992-
   08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:discovery
   ```

```
traddr: 192.168.6.15
sectype: none
=====Discovery Log Entry 3======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.1c6ac66338e711eda41dd039ea3ad566:subsystem.host_95
traddr: 192.168.6.17
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.15
# nvme discover -t tcp -w 192.168.5.13 -a 192.168.5.17
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.15
# nvme discover -t tcp -w 192.168.6.13 -a 192.168.6.17
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.15 -l 1800
# nvme connect-all -t tcp -w 192.168.5.13 -a 192.168.5.17 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.15 -l 1800
# nvme connect-all -t tcp -w 192.168.6.13 -a 192.168.6.17 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled by checking:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_
ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
optimized
```

NVMe/TCP

```
# nvme list-subsys /dev/nvme1n40
```

Example output:

```
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.68c036aaa3cf11edbb95d039ea243511:subsystem.tcp
\
+- nvme2 tcp
traddr=192.168.8.49,trsvcid=4420,host_traddr=192.168.8.1 live non-
optimized
+- nvme3 tcp
traddr=192.168.8.48,trsvcid=4420,host_traddr=192.168.8.1 live non-
optimized
+- nvme6 tcp
traddr=192.168.9.49,trsvcid=4420,host_traddr=192.168.9.1 live
optimized
+- nvme7 tcp
traddr=192.168.9.48,trsvcid=4420,host_traddr=192.168.9.1 live
optimized
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device        Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns


NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 8.7 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 8.7 NVMe-oF Hosts create duplicate Persistent Discovery Controllers	On OL 8.7 NVMe-oF hosts, Persistent Discovery Controllers (PDCs) are created by passing the `-p` option to the `nvme discover` command. For a given initiator-target combination, only one PDC is expected to be created with each invocation of the `nvme discover` command. However, beginning with OL 8.x, NVMe-oF hosts end up creating duplicate PDCs with each invocation of the `nvme discover` command with the `-p` option. This wastes resources on both the host and the target.	18118

NVMe/FC Host Configuration for Oracle Linux 8.6 with ONTAP

Supportability

NVMe over Fabrics or NVMe-oF (including NVMe/FC and NVMe/TCP) is supported with Oracle Linux 8.6 with Asymmetric Namespace Access (ANA) that is required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. This document contains the details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on Oracle Linux 8.6 and ONTAP as the target.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Features

- Oracle Linux 8.6 has in-kernel NVMe multipath enabled by default for NVMe namepsaces.
- With Oracle Linux 8.6, `nvme-fc auto-connect` scripts are included in the native `nvme-cli` package. You can rely on these native auto-connect scripts instead of installing external vendor provided outbox auto-connect scripts.
- With Oracle Linux 8.6, a native `udev` rule is provided as part of the `nvme-cli` package which enables round-robin load balancing for NVMe multipath. Therefore, you need not manually create this rule anymore.
- With Oracle Linux 8.6, both NVMe and SCSI traffic can be run on the same co-existent host. In fact, that is expected to be the commonly deployed host configuration. Therefore, you can configure dm-multipath as usual for SCSI LUNs resulting in mpath devices whereas NVMe multipath can be used to configure NVMe-oF multipath devices (for example, `/dev/nvmeXnY`) on the host.

- With Oracle Linux 8.6, the NetApp plug-in in the native `nvme-cli` package is capable of displaying ONTAP details as well as ONTAP namespaces.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Configuration requirements

Refer to the NetApp Interoperability Matrix for exact details regarding supported configurations.

Enable NVMe/FC with Oracle Linux 8.6

Steps

1. Install Oracle Linux 8.6 GA on the server. After the installation is complete, verify that you are running the specified Oracle Linux 8.6 GA kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

```
# uname -r
5.4.17-2136.307.3.1.el8uek.x86_64
```

2. Install the `nvme-cli` package:

```
# rpm -qa|grep nvme-cli
nvme-cli-1.14-3.el8.x86_64
```

3. On the Oracle Linux 8.6 host, check the `hostnqn` string at `/etc/nvme/hostnqn` and verify that it matches the hostnqn string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:4c4c4544-0032-3310-8033-b8c04f4c5132
::> vserver nvme subsystem host show -vserver vs_ol_nvme
Vserver     Subsystem          Host NQN
----------- ---------------
----------------------------------------------------------
vs_ol_nvme  nvme_ss_ol_1    nqn.2014-08.org.nvmexpress:uuid:9ed5b327-
b9fc-4cf5-97b3-1b5d986345d1
```

> If the `hostnqn` strings do not match, you should use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match to the `hostnqn` string from `/etc/nvme/hostnqn` on the host:

4. Reboot the host.

If you intend to run both NVMe and SCSI traffic on the same Oracle Linux 8.6 co-existent host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This also means the ONTAP namespaces should be blacklisted in dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
#cat /etc/multipath.conf
defaults {
    enable_foreign  NONE
}
```

Restart the multipathd daemon by running the `systemctl restart multipathd` command to let the new setting take effect.

Configure Broadcom FC adapter for NVMe/FC

Steps

1. Verify that you are using the supported adapter. For the most current list of supported adapters see the NetApp Interoperability Matrix:

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver. For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix:

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.11, sli-4:2:c
14.0.505.11, sli-4:2:c

# cat /sys/module/lpfc/version
0:12.8.0.11
```

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Enable 1MB I/O size

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure the Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. Verify that you are running the supported adapter driver and firmware versions. The native inbox qla2xxx driver included in the OL 8.6 GA kernel has the latest upstream fixes essential for ONTAP support:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
```

2. Verify that `ql2xnvmeenable` is set which enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Similarly, verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.
Example,

```
#nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Now run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. MAke sure you pass a longer `ctrl_loss_tmo` period (such as, say 30 minutes, which can be set through `-l 1800`) during the `connect-all` so that it would retry for a longer period in the event of a path loss. For example,

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe/FC

Steps

1. Verify the following NVMe/FC settings on the Oracle Linux 8.6 host:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
2                 85.90 GB / 85.90 GB  24 KiB + 0 B  FFFFFFFF
3                 85.90 GB / 85.90 GB  4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status:

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

```
# nvme netapp ontapdevices -o column

Device         Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1   vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2   vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3   vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns


NSID       UUID                                   Size
------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 8.6 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 8.6 NVMe-oF Hosts create duplicate Persistent Discovery Controllers	On Oracle Linux 8.6 NVMe over Fabrics (NVMe-oF) hosts, you can use the `nvme discover -p` command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Oracle Linux 8.6 with an NVMe-oF host, a duplicate PDC is created each time `nvme discover -p` is executed. This leads to unnecessary usage of resources on both the host and the target.	18118

NVMe/FC Host Configuration for Oracle Linux 8.5 with ONTAP

Supportability

NVMe over Fabrics or NVMe-oF (including NVMe/FC and NVMe/TCP) is supported with Oracle Linux 8.5 with Asymmetric Namespace Access (ANA) that is required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. This document contains the details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on Oracle Linux 8.5 and ONTAP as the target.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Features

- Oracle Linux 8.5 has in-kernel NVMe multipath enabled by default for NVMe namepsaces.
- With Oracle Linux 8.5, `nvme-fc auto-connect` scripts are included in the native `nvme-cli` package. You can rely on these native auto-connect scripts instead of installing external vendor provided outbox auto-connect scripts.
- With Oracle Linux 8.5, a native `udev` rule is provided as part of the `nvme-cli` package which enables round-robin load balancing for NVMe multipath. Therefore, you do not need to manually create this rule anymore.
- With Oracle Linux 8.5, both NVMe and SCSI traffic can be run on the same co-existent host. In fact, that is expected to be the commonly deployed host configuration. Therefore, you can configure dm-multipath as usual for SCSI LUNs resulting in mpath devices whereas NVMe multipath can be used to configure NVMe-oF multipath devices (for example, `/dev/nvmeXnY`) on the host.
- With Oracle Linux 8.5, the NetApp plugin in the native `nvme-cli` package is capable of displaying ONTAP

details as well as ONTAP namespaces.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Configuration requirements

Refer to the NetApp Interoperability Matrix for exact details regarding supported configurations.

Enable NVMe/FC with Oracle Linux 8.5

Steps

1. Install Oracle Linux 8.5 General Availability (GA) on the server. After the installation is complete, verify that you are running the specified Oracle Linux 8.5 GA kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

   ```
   # uname -r
   5.4.17-2136.309.4.el8uek.x86_64
   ```

2. Install the `nvme-cli` package.

   ```
   # rpm -qa|grep nvme-cli
   nvme-cli-1.14-3.el8.x86_64
   ```

3. On the Oracle Linux 8.5 host, check the `hostnqn` string at `/etc/nvme/hostnqn` and verify that it matches the `hostnqn` string for the corresponding subsystem on the ONTAP array.

   ```
   # cat /etc/nvme/hostnqn
   nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
   ::> vserver nvme subsystem host show -vserver vs_ol_nvme

   Vserver     Subsystem       Host NQN
   ---------------------------------------------
   vs_ol_nvme nvme_ss_ol_1   nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-
   4cf5-97b3-1b5d986345d1
   ```

   If the `hostnqn` strings do not match, you should use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match to the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

4. Reboot the host.

If you intend to run both NVMe and SCSI traffic on the same Oracle Linux 8.5 co-existent host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This also means the ONTAP namespaces should be blacklisted in dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
#cat /etc/multipath.conf
defaults {
    enable_foreign  NONE
}
```

Restart the `multipathd` daemon by running the `systemctl restart multipathd` command to let the new setting take effect.

Configure the Broadcom FC adapter for NVMe/FC

Steps

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.11, sli-4:2:c
14.0.505.11, sli-4:2:c

# cat /sys/module/lpfc/version
0:12.8.0.5
```

3. Verify that `lpfc_enable_fc4_type` is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b213a00
0x100000109b2139ff
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b213a00 WWNN x200000109b213a00 DID
x031700     ONLINE
NVME RPORT WWPN x208cd039ea243510 WWNN x208bd039ea243510 DID x03180a
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2090d039ea243510 WWNN x208bd039ea243510 DID x03140a
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000000079efc Issue 0000000000079eeb OutIO
ffffffffffffffef
abort 00000002 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err   00000000
FCP CMPL: xb 00000002 Err 00000004

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b2139ff WWNN x200000109b2139ff DID
x031300 ONLINE
NVME RPORT WWPN x208ed039ea243510 WWNN x208bd039ea243510 DID x03230c
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2092d039ea243510 WWNN x208bd039ea243510 DID x03120c
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000000029ba0 Issue 0000000000029ba2 OutIO
0000000000000002
abort 00000002 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000002 Err 00000004
```

Enable 1MB I/O size

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure the Marvell/QLogic FC adapter for NVMe/FC

Steps

1. Verify that you are running the supported adapter driver and firmware versions. The native inbox qla2xxx driver included in the OL 8.5 GA kernel has the latest upstream fixes essential for ONTAP support:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as an NVMe/FC initiator.

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port is able to fetch discovery log page data across the supported NVMe/TCP LIFs.

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none

...
```

2. Similarly, verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.

Example,

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Now run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Make sure you provide a longer `ctrl_loss_tmo` timer period (such as say 30 minutes, which can be set adding `-l 1800`) during `connect-all` so that it would retry for a longer period in the event of a path loss.
   Example:

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe/FC

Steps

1. Verify the following NVMe/FC settings on the Oracle Linux 8.5 host.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discovered on the host.

```
# nvme list
Node         SN                    Model
---------------------------------------------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
/dev/nvme0n2 814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
/dev/nvme0n3 814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller

Namespace Usage  Format                 FW            Rev
--------------------------------------------------------------
1                85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
2                85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
3                85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify the NetApp plug-in displays correct values for each ONTAP namespace device.

```
# nvme netapp ontapdevices -o column
Device       Vserver  Namespace Path
-----------------------------------
/dev/nvme0n1  vs_ol_nvme  /vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2  vs_ol_nvme  /vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3  vs_ol_nvme  /vol/ol_nvme_vol_1_1_1/ol_nvme_ns

NSID    UUID                                   Size
-----------------------------------------------------------
1       72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2       04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3       264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 8.5 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 8.5 NVMe-oF Hosts create duplicate Persistent Discovery Controllers	On Oracle Linux 8.5 NVMe over Fabrics (NVMe-oF) hosts, you can use the `nvme discover -p` command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Oracle Linux 8.5 with an NVMe-oF host, a duplicate PDC is created each time `nvme discover -p` is executed. This leads to unnecessary usage of resources on both the host and the target.	18118

NVMe/FC Host Configuration for Oracle Linux 8.4 with ONTAP

Supportability

NVMe over Fabrics or NVMe-oF (including NVMe/FC and NVMe/TCP) is supported with Oracle Linux 8.4 with Asymmetric Namespace Access (ANA), which is required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. This topic details how to enable NVMe-oF with in-kernel NVMe Multipath using ANA on Oracle Linux 8.4 with ONTAP as the target.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Features

- Oracle Linux 8.4 has in-kernel NVMe multipath enabled by default for NVMe namepsaces.
- With Oracle Linux 8.4, `nvme-fc auto-connect` scripts are included in the native `nvme-cli` package. You can rely on these native auto-connect scripts instead of installing external vendor provided outbox auto-connect scripts.
- With Oracle Linux 8.4, a native `udev` rule is provided as part of the `nvme-cli` package which enables round-robin load balancing for NVMe multipath. Therefore, you do not need to manually create this rule anymore.
- With Oracle Linux 8.4, both NVMe and SCSI traffic can be run on the same co-existent host. In fact, that is expected to be the commonly deployed host configuration. Therefore, you can configure dm-multipath as usual for SCSI LUNs resulting in mpath devices whereas NVMe multipath can be used to configure NVMe-oF multipath devices (for example, `/dev/nvmeXnY`) on the host.
- With Oracle Linux 8.4, the NetApp plugin in the native `nvme-cli` package is capable of displaying ONTAP

details as well as ONTAP namespaces.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Configuration requirements

Refer to the NetApp Interoperability Matrix (IMT) for exact details on supported configurations.

Enable NVMe/FC

Steps

1. Install Oracle Linux 8.4 GA on the server. After the installation is complete, verify that you are running the specified Oracle Linux 8.4 GA kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

```
# uname -r
5.4.17-2102.206.1.el8uek.x86_64
```

2. Install the `nvme-cli` package.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.12-3.el8.x86_64
```

3. On the Oracle Linux 8.4 host, check the hostnqn string at `/etc/nvme/hostnqn` and verify that it matches the hostnqn string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:8b43c7c6-e98d-4cc7-a699-d66a69aa714e
::> vserver nvme subsystem host show -vserver vs_coexistance_2

Vserver          Subsystem Host NQN
-------
--------------------------------------------------------------------
vs_coexistance_2 nvme_1    nqn.2014-08.org.nvmexpress:uuid:753881b6-3163-
46f9-8145-0d1653d99389
```

> If the hostnqn strings do not match, you should use the `vserver modify` command to update the hostnqn string on your corresponding ONTAP array subsystem to match to the hostnqn string from `/etc/nvme/hostnqn` on the host.

4. Reboot the host.

If you intend to run both NVMe and SCSI traffic on the same Oracle Linux 8.4 co-existent host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This also means the ONTAP namespaces should be blacklisted in dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
#cat /etc/multipath.conf
defaults {
    enable_foreign  NONE
}
```

Restart the multipathd daemon by running the `systemctl restart multipathd` command to let the new setting take effect.

Configuring the Broadcom FC adapter for NVMe/FC

Steps

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.11, sli-4:2:c
14.0.505.11, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:12.8.0.5
```

3. Verify that `lpfc_enable_fc4_type` is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b213a00
0x100000109b2139ff

# cat /sys/class/fc_host/host*/port_state
Online
Online

# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b213a00 WWNN x200000109b213a00 DID
x031700     ONLINE
NVME RPORT WWPN x208cd039ea243510 WWNN x208bd039ea243510 DID x03180a
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2090d039ea243510 WWNN x208bd039ea243510 DID x03140a
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000000079efc Issue 0000000000079eeb OutIO
ffffffffffffffef
abort 00000002 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err   00000000
FCP CMPL: xb 00000002 Err 00000004

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b2139ff WWNN x200000109b2139ff DID
x031300 ONLINE
NVME RPORT WWPN x208ed039ea243510 WWNN x208bd039ea243510 DID x03230c
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x2092d039ea243510 WWNN x208bd039ea243510 DID x03120c
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000000029ba0 Issue 0000000000029ba2 OutIO
0000000000000002
abort 00000002 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000002 Err 00000004
```

Enabling 1MB I/O size

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure the Marvell/QLogic FC adapter for NVMe/FC

Steps

1. Verify that you are running the supported adapter driver and firmware versions. The native inbox qla2xxx driver included in the OL 8.4 GA kernel has the latest upstream fixes essential for ONTAP support.

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.08.02 DVR:v10.02.00.103-k
QLE2742 FW:v9.08.02 DVR:v10.02.00.103-k
```

2. Verify that the `ql2xnvmeenable` parameter is set which enables the Marvell adapter to function as an NVMe/FC initiator.

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port is able to fetch discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Similarly, verify that other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.
Example,

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
#nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Now run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Make sure you provide a longer `ctrl_loss_tmo` timer period (30 minutes or more, which can be set adding `-l 1800`) during `connect-all` so that it would retry for a longer period in the event of a path loss.
   Example:

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe/FC

Steps

1. Verify the following NVMe/FC settings on the Oracle Linux 8.4 host:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discoverd on the host:

```
# nvme list
Node                   SN                Model
Namespace
------------------------------------------------------------------------------------
-----------------
/dev/nvme0n1     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
1
/dev/nvme0n2     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
2
/dev/nvme0n3     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
3

Usage      Format         FW Rev
------------------------------------------------------
85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
```

3. Verify that the controller state of each path is live and has the correct ANA status.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live non-
optimized
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify that the NetApp plug-in displays correct values for each ONTAP namespace devices.

```
# nvme netapp ontapdevices -o column
Device                  Vserver          Namespace Path
----------------------- ------------------------------
-----------------------------------------------------------------------
--------- --
/dev/nvme0n1      vs_ol_nvme
/vol/ol_nvme_vol_1_1_0/ol_nvme_ns
/dev/nvme0n2      vs_ol_nvme
/vol/ol_nvme_vol_1_0_0/ol_nvme_ns
/dev/nvme0n3      vs_ol_nvme
/vol/ol_nvme_vol_1_1_1/ol_nvme_ns

NSID        UUID                                  Size
--------------------------------------------------------------
1          72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
2          04bf9f6e-9031-40ea-99c7-a1a61b2d7d08   85.90GB
3          264823b1-8e03-4155-80dd-e904237014a4   85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Known issues

The NVMe-oF host configuration for OL 8.4 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1517321	Oracle Linux 8.4 NVMe-oF hosts create duplicate Persistent Discovery Controllers	On Oracle Linux 8.4 NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Oracle Linux 8.4 with an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	18118

NVMe/FC Host Configuration for Oracle Linux 8.3 with ONTAP

Supportability

NVMe over Fabrics or NVMe-oF (including NVMe/FC) is supported with Oracle Linux 8.3 with Asymmetric Namespace Access (ANA) required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the ALUA equivalent in the NVMe-oF environment and is currently implemented with in-kernel NVMe Multipath. This document contains the details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on OL 8.3 and ONTAP as the target.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Features

- Oracle Linux 8.3 has in-kernel NVMe multipath enabled by default for NVMe namespaces.
- With Oracle Linux 8.3, `nvme-fc auto-connect` scripts are included in the native nvme-cli package. You can rely on these native auto-connect scripts instead of installing external vendor provided outbox auto-connect scripts.
- With Oracle Linux 8.3, a native `udev` rule is provided as part of the `nvme-cli` package which enables round-robin load balancing for NVMe multipath. Therefore, you do not need to manually create this rule anymore.
- With Oracle Linux 8.3, both NVMe and SCSI traffic can be run on the same co-existent host. In fact, that is expected to be the commonly deployed host configuration. So for SCSI, you can configure dm-multipath as usual for SCSI LUNs resulting in mpath devices whereas NVMe multipath can be used to configure NVMe-oF multipath devices (for example, `/dev/nvmeXnY`) on the host.
- With Oracle Linux 8.3, the NetApp plugin in the native `nvme-cli` package is capable of displaying ONTAP details as well as ONTAP namespaces.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Configuration requirements

Refer to the NetApp Interoperability Matrix Tool (IMT) for the current list of supported configurations.

Enable NVMe/FC with Oracle Linux 8.3

Steps

1. Install Oracle Linux 8.3 GA on the server. After the installation is complete, verify that you are running the specified Oracle Linux 8.3 GA kernel. See the NetApp Interoperability Matrix Tool for the most current list of supported versions.

```
# uname -r
5.4.17-2011.7.4.el8uek.x86_64
```

2. Install the `nvme-cli` package.

```
# rpm -qa|grep nvme-cli

nvme-cli-1.12-2.el8.x86_64_
```

3. On the Oracle Linux 8.3 host, check the hostnqn string at `/etc/nvme/hostnqn` and verify that it matches the `hostnqn` string for the corresponding subsystem on the ONTAP array.

```
#cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:446c21ab-f4c1-47ed-9a8f-1def96f3fed2

::> vserver nvme subsystem host show -vserver vs_coexistance_2
Vserver   Subsystem     Host NQN
------- ---------
----------------------------------------------------------
vs_coexistance_2 nvme_1 nqn.2014-08.org.nvmexpress:uuid:446c21ab-f4c1-
47ed-9a8f-1def96f3fed2
```

> If the hostnqn strings do not match, you should use the `vserver modify` command to update the hostnqn string on your corresponding ONTAP array subsystem to match to the hostnqn string from `/etc/nvme/hostnqn` on the host.

4. Reboot the host.

If you intend to run both NVMe and SCSI traffic on the same Oracle Linux 8.3 co-existent host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and `dm-multipath` for ONTAP LUNs respectively. This also means the ONTAP namespaces should be blacklisted in `dm-multipath` to prevent `dm-multipath` from claiming these namespace devices. This can be done by adding the *enable_foreign* setting to the `/etc/multipath.conf` file:

```
#cat /etc/multipath.conf
defaults {
   enable_foreign  NONE
}
```

Restart the multipathd daemon by running the *systemctl restart multipathd* command to let the new setting take effect.

Configure the Broadcom FC adapter for NVMe/FC

Steps

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.

```
#cat /sys/class/scsi_host/host*/modelname
LPe36002-M2
LPe36002-M2
```

```
#cat /sys/class/scsi_host/host*/modeldesc
Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. For the most current list of supported adapter drivers and firmware versions, see the NetApp Interoperability Matrix Tool.

```
#cat /sys/class/scsi_host/host*/fwrev
12.8.351.49, sli-4:6:d
12.8.351.49, sli-4:6:d
```

```
#cat /sys/module/lpfc/version
0:12.6.0.3
```

3. Verify that the `lpfc_enable_fc4_type` parameter is set to 3.

```
#cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs.

```
#cat /sys/class/fc_host/host*/port_name
0x100000109bf0447b
0x100000109bf0447c
```

```
#cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
#cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109bf0447b WWNN x200000109bf0447b DID
x022400 ONLINE
NVME RPORT WWPN x20e1d039ea243510 WWNN x20e0d039ea243510 DID x0a0314
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x20e4d039ea243510 WWNN x20e0d039ea243510 DID x0a0713
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 00000003b6 Cmpl 00000003b6 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000be1425e8 Issue 00000000be1425f2 OutIO
000000000000000a
abort 00000251 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000c5b Err 0000d176

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109bf0447c WWNN x200000109bf0447c DID
x021600 ONLINE
NVME RPORT WWPN x20e2d039ea243510 WWNN x20e0d039ea243510 DID x0a0213
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x20e3d039ea243510 WWNN x20e0d039ea243510 DID x0a0614
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000419 Cmpl 0000000419 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000be37ff65 Issue 00000000be37ff84 OutIO
000000000000001f
abort 0000025a noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000c89 Err 0000cd87
```

Enable 1MB I/O size

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure the Marvell/QLogic FC adapter for NVMe/FC

Steps

1. Verify that you are running the supported adapter driver and firmware versions. The native inbox qla2xxx driver included in the OL 8.3 GA kernel has the latest upstream fixes essential for ONTAP support.

```
#cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.10.11 DVR:v10.01.00.25-k
QLE2742 FW:v9.10.11 DVR:v10.01.00.25-k
```

2. Verify that the `ql2xnvmeenable` parameter is set which enables the Marvell adapter to function as an NVMe/FC initiator.

```
#cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Validate NVMe/FC

Steps

1. Verify the following NVMe/FC settings on the Oracle Linux 8.3 host.

```
#cat /sys/module/nvme_core/parameters/multipath
Y

#cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

#cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and correctly discoverd on the host.

```
# nvme list
Node         SN                   Model                   Namespace Usage
Format FW Rev
---------------- --------------------
---------------------------------------- --------- ------------------
/dev/nvme0n1 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 1        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n10 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 10      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n11 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 11      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n12 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 12      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n13 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 13      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n14 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 14      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n15 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 15      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n16 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 16      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n17 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 17      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n18 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 18      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n19 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 19      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n2 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 2        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n20 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 20      37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n3 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 3        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n4 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 4        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n5 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 5        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n6 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 6        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n7 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 7        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n8 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 8        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
/dev/nvme0n9 81Ec-JRMlkL9AAAAAAAB NetApp ONTAP Controller 9        37.58
GB / 37.58 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify that the controller state of each path is live and has correct ANA status.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.b79f5c6e4d0911edb3a0d039ea243511:subsystem.nvme_1
\ +
+- nvme214 fc traddr=nn-0x20e0d039ea243510:pn-0x20e4d039ea243510
host_traddr=nn-0x200000109bf0447b:pn-0x100000109bf0447b live non-
optimized
+- nvme219 fc traddr=nn-0x20e0d039ea243510:pn-0x20e2d039ea243510
host_traddr=nn-0x200000109bf0447c:pn-0x100000109bf0447c live optimized
+- nvme223 fc traddr=nn-0x20e0d039ea243510:pn-0x20e1d039ea243510
host_traddr=nn-0x200000109bf0447b:pn-0x100000109bf0447b live optimized
+- nvme228 fc traddr=nn-0x20e0d039ea243510:pn-0x20e3d039ea243510
host_traddr=nn-0x200000109bf0447c:pn-0x100000109bf0447c live non-
optimized
```

4. Verify the NetApp plug-in displays correct values for each ONTAP namespace devices.

```
#nvme netapp ontapdevices -o column
Device      Vserver        Namespace Path             NSID UUID
Size
---------------- -------------------------
-------------------------------------------------- ---- ---------
/dev/nvme0n1 LPE36002_ASA_BL /vol/fcnvme_1_0_0/fcnvme_ns 1 ae10e16d-
1fa4-49c2-8594-02bf6f3b1af1 37.58GB
/dev/nvme0n10 LPE36002_ASA_BL /vol/fcnvme_1_0_9/fcnvme_ns 10 2cf00782-
e2bf-40fe-8495-63e4501727cd 37.58GB
/dev/nvme0n11 LPE36002_ASA_BL /vol/fcnvme_1_1_9/fcnvme_ns 11 fbefbe6c-
90fe-46a2-8a51-47bad9e2eb95 37.58GB
/dev/nvme0n12 LPE36002_ASA_BL /vol/fcnvme_1_1_0/fcnvme_ns 12 0e9cc8fa-
d821-4f1c-8944-3003dcded864 37.58GB
/dev/nvme0n13 LPE36002_ASA_BL /vol/fcnvme_1_1_1/fcnvme_ns 13 31f03b13-
aaf9-4a3f-826b-d126ef007991 37.58GB
/dev/nvme0n14 LPE36002_ASA_BL /vol/fcnvme_1_1_8/fcnvme_ns 14 bcf4627c-
5bf9-4a51-a920-5da174ec9876 37.58GB
/dev/nvme0n15 LPE36002_ASA_BL /vol/fcnvme_1_1_7/fcnvme_ns 15 239fd09d-
11db-46a3-8e94-b5ebe6eb2421 37.58GB
/dev/nvme0n16 LPE36002_ASA_BL /vol/fcnvme_1_1_2/fcnvme_ns 16 1d8004df-
f2e8-48c8-8ccb-ce45f18a15ae 37.58GB
/dev/nvme0n17 LPE36002_ASA_BL /vol/fcnvme_1_1_3/fcnvme_ns 17 4f7afbcf-
3ace-4e6c-9245-cbf5bd155ef4 37.58GB
/dev/nvme0n18 LPE36002_ASA_BL /vol/fcnvme_1_1_4/fcnvme_ns 18 b022c944-
6ebf-4986-a28c-8d9e8ec130c9 37.58GB
/dev/nvme0n19 LPE36002_ASA_BL /vol/fcnvme_1_1_5/fcnvme_ns 19 c457d0c7-
bfea-43aa-97ef-c749d8612a72 37.58GB
/dev/nvme0n2 LPE36002_ASA_BL /vol/fcnvme_1_0_1/fcnvme_ns 2 d2413d8b-
e82e-4412-89d3-c9a751ed7716 37.58GB
/dev/nvme0n20 LPE36002_ASA_BL /vol/fcnvme_1_1_6/fcnvme_ns 20 650e0d93-
967d-4415-874a-36bf9c93c952 37.58GB
/dev/nvme0n3 LPE36002_ASA_BL /vol/fcnvme_1_0_2/fcnvme_ns 3 09d89d9a-
7835-423f-93e7-f6f3ece1dcbc 37.58GB
/dev/nvme0n4 LPE36002_ASA_BL /vol/fcnvme_1_0_3/fcnvme_ns 4 d8e99326-
a67c-469f-b3e9-e0e4a38c8a76 37.58GB
/dev/nvme0n5 LPE36002_ASA_BL /vol/fcnvme_1_0_4/fcnvme_ns 5 c91c71f9-
3e04-4844-b376-30acab6311f1 37.58GB
/dev/nvme0n6 LPE36002_ASA_BL /vol/fcnvme_1_0_5/fcnvme_ns 6 4e8b4345-
e5b1-4aa4-ae1a-adf0de2879ea 37.58GB
/dev/nvme0n7 LPE36002_ASA_BL /vol/fcnvme_1_0_6/fcnvme_ns 7 ef715a16-
a946-4bb8-8735-74f214785874 37.58GB
/dev/nvme0n8 LPE36002_ASA_BL /vol/fcnvme_1_0_7/fcnvme_ns 8 4b038502-
966c-49fd-9631-a17f23478ae0 37.58GB
/dev/nvme0n9 LPE36002_ASA_BL /vol/fcnvme_1_0_8/fcnvme_ns 9 f565724c-
992f-41f6-83b5-da1fe741c09b 37.58GB
```

```
#nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
{
"Device" : "/dev/nvme0n1",
"Vserver" : "LPE36002_ASA_BL",
"Namespace_Path" : "/vol/fcnvme_1_0_0/fcnvme_ns",
"NSID" : 1,
"UUID" : "ae10e16d-1fa4-49c2-8594-02bf6f3b1af1",
"Size" : "37.58GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 9175040
},
{
"Device" : "/dev/nvme0n10",
"Vserver" : "LPE36002_ASA_BL",
"Namespace_Path" : "/vol/fcnvme_1_0_9/fcnvme_ns",
"NSID" : 10,
"UUID" : "2cf00782-e2bf-40fe-8495-63e4501727cd",
"Size" : "37.58GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 9175040
},
{
"Device" : "/dev/nvme0n11",
"Vserver" : "LPE36002_ASA_BL",
"Namespace_Path" : "/vol/fcnvme_1_1_9/fcnvme_ns",
"NSID" : 11,
"UUID" : "fbefbe6c-90fe-46a2-8a51-47bad9e2eb95",
"Size" : "37.58GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 9175040
},
{
"Device" : "/dev/nvme0n12",
"Vserver" : "LPE36002_ASA_BL",
"Namespace_Path" : "/vol/fcnvme_1_1_0/fcnvme_ns",
"NSID" : 12,
"UUID" : "0e9cc8fa-d821-4f1c-8944-3003dcded864",
"Size" : "37.58GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 9175040
},
{
"Device" : "/dev/nvme0n13",
```

```
"Vserver" : "LPE36002_ASA_BL",
"Namespace_Path" : "/vol/fcnvme_1_1_1/fcnvme_ns",
"NSID" : 13,
"UUID" : "31f03b13-aaf9-4a3f-826b-d126ef007991",
"Size" : "37.58GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 9175040
},
```

Known issues

The NVMe-oF host configuration for OL 8.3 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Oracle Bugzilla
1517321	Oracle Linux 8.3 NVMe-oF Hosts create duplicate Persistent Discovery Controllers	On Oracle Linux 8.3 NVMe over Fabrics (NVMe-oF) hosts, you can use the `nvme discover -p` command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Oracle Linux 8.3 with an NVMe-oF host, a duplicate PDC is created each time `nvme discover -p` is executed. This leads to unnecessary usage of resources on both the host and the target.	18118

NVMe/FC Host Configuration for Oracle Linux 8.2 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.6 or later for Oracle Linux 8.2. Oracle Linux 8.2 host can run both NVMe/FC and FCP traffic through the same fibre channel (FC) initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations see the NetApp Interoperability Matrix.

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC

1. Install Oracle Linux 8.2 on the server.
2. After the installation is complete, verify that you are running the supported Unbreakable Enterprise kernel. See the NetApp Interoperability Matrix.

```
# uname -r
5.4.17-2011.1.2.el8uek.x86_64
```

3. Upgrade the nvme-cli package. The native nvme-cli package contains the NVMe/FC auto-connect scripts, ONTAP udev rule which enables round-robin load balancing for NVMe Multipath as well as the NetApp plug-in for ONTAP namespaces.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.9-5.el8.x86_64
```

4. On the Oracle Linux 8.2 host, check the host NQN string at /etc/nvme/hostnqn and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
```

```
::> vserver nvme subsystem host show -vserver vs_ol_nvme
Vserver  Subsystem Host NQN
----------- ---------------
----------------------------------------------------------
vs_ol_nvme
              nvme_ss_ol_1
                                 nqn.2014-
08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
```

If the hostnqn strings do not match, you should use the vserver modify command to update the host NQN string on your corresponding ONTAP array subsystem to match to host NQN string from etc/nvme/hostnqn on the host.

Configure the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. NVMe support in lpfc is already enabled by default:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

   Newer lpfc drivers (both inbox and outbox) have lpfc_enable_fc4_type default set to 3. Therefore, you do not need to set this explicitly in the /etc/modprobe.d/lpfc.conf.

3. Verify that the NVMe/FC initiator ports are enabled and able to see the target ports, and all are up and running.

   In the example below, only a single initiator port has been enabled and connected with two target LIFs as seen in the below output:

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Validate NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

In the above example, two namespaces are mapped to the Oracle Linux 8.2 ANA host. These are visible through four target LIFs: two local node LIFs and two other partner/remote node LIFs. This setup shows as two ANA Optimized and two ANA Inaccessible paths for each namespace on the host.

2. Verify that the namespaces are created.

```
# nvme list
Node                  SN
Model                                               Namespace Usage
Format         FW Rev
--------------------   ---------------------------------------
----------------------------------------------- -----------------
------------------------------- ---------------- ---------------
/dev/nvme0n1     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
1                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n2     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
2                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n3     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
3                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_ss_ol_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device                 Vserver                          Namespace Path
NSID    UUID
Size
----------------------- ------------------------------
-----------------------------------------------------------------------
--------- ---------------------------------------------------------
---------
/dev/nvme0n1      vs_ol_nvme
/vol/ol_nvme_vol_1_1_0/ol_nvme_ns              1          72b887b1-5fb6-
47b8-be0b-33326e2542e2   85.90GB
/dev/nvme0n2      vs_ol_nvme
/vol/ol_nvme_vol_1_0_0/ol_nvme_ns              2          04bf9f6e-9031-
40ea-99c7-a1a61b2d7d08    85.90GB
/dev/nvme0n3      vs_ol_nvme
/vol/ol_nvme_vol_1_1_1/ol_nvme_ns              3          264823b1-8e03-
4155-80dd-e904237014a4  85.90GB

# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_1_0/ol_nvme_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
```

```
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_ol_nvme",
        "Namespace_Path" : "/vol/ol_nvme_vol_1_0_0/ol_nvme_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_ol_nvme",
         "Namespace_Path" : "/vol/ol_nvme_vol_1_1_1/ol_nvme_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
}
```

Enable 1MB I/O size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

NVMe/FC Host Configuration for Oracle Linux 8.1 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.6 or later for Oracle Linux 8.1. Oracle Linux 8.1 host can run both NVMe and SCSI traffic through the same fibre channel (FC) initiator adapter ports. Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same FC adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations see the NetApp Interoperability Matrix.

You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

- Native NVMe/FC auto-connect scripts are not available in the nvme-cli package. Use the HBA vendor provided external auto-connect scripts.
- By default, round-robin load balancing is not enabled in NVMe Multipath. You must write a udev rule to enable this functionality. Steps are provided in the section on Enabling NVMe/FC on Oracle Linux 8.1.
- There is no sanlun support for NVMe/FC and, as a consequence, no Linux Unified Host Utilities (LUHU) support for NVMe/FC on Oracle Linux 8.1. Use the ONTAP command output available as part of the NetApp plug-in included in the native nvme-cli.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC

1. Install Oracle Linux 8.1 on the server.
2. After the installation is complete, verify that you are running the supported Unbreakable Enterprise kernel. See the NetApp Interoperability Matrix.

```
# uname -r
5.4.17-2011.0.7.el8uek.x86_64
```

3. Upgrade the nvme-cli package.

```
# rpm -qa | grep nvmefc
nvmefc-connect-12.6.61.0-1.noarch
```

4. Add the string below as a separate udev rule at /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules. This enables round-robin load balancing for NVMe multipath.

```
# cat /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules
# Enable round-robin for NetApp ONTAP
ACTION=="add", SUBSYSTEM=="nvme-subsystem", ATTR{model}=="NetApp ONTAP
Controller", ATTR{iopolicy}="round-robin"
```

5. On the Oracle Linux 8.1 host, check the host NQN string at /etc/nvme/hostnqn and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

```
*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
------- ---------
----------------------------------------------------------
Oracle Linux_141_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

If the hostnqn strings do not match, you should use the vserver modify command to update the host NQN string on your corresponding ONTAP array subsystem to match to host NQN string from etc/nvme/hostnqn on the host.

6. Reboot the host.

Configure the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. NVMe support in lpfc is already enabled by default:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

Newer lpfc drivers (both inbox and outbox) have lpfc_enable_fc4_type default set to 3. Therefore, you do not need to set this explicitly in the /etc/modprobe.d/lpfc.conf.

3. Next, install the recommended lpfc auto-connect scripts:

```
# rpm -ivh nvmefc-connect-12.6.61.0-1.noarch.rpm
```

4. Verify that the auto-connect scripts are installed.

```
# rpm -qa | grep nvmefc
nvmefc-connect-12.6.61.0-1.noarch
```

5. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62

# cat /sys/class/fc_host/host*/port_state
Online
Online
```

6. Verify that the NVMe/FC initiator ports are enabled and able to see the target ports, and all are up and running.

   In the example below, only a single initiator port has been enabled and connected with two target LIFs as seen in the below output:

```
# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2947 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
```

Validate NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

In the above example, two namespaces are mapped to the Oracle Linux 8.1 ANA host. These are visible through four target LIFs: two local node LIFs and two other partner/remote node LIFs. This setup shows as two ANA Optimized and two ANA Inaccessible paths for each namespace on the host.

2. Verify that the namespaces are created.

```
# nvme list
Node                SN                                           Model
Namespace Usage                              Format          FW Rev
-------------------- ---------------------------------------
---------------------------------------- ----------------
------------------------------  ----------------  -------------
/dev/nvme0n1  814vWBNRwfBCAAAAAAAB NetApp ONTAP Controller        2
107.37 GB / 107.37 GB  4 KiB + 0 B   FFFFFFFF
/dev/nvme0n2  814vWBNRwfBCAAAAAAAB NetApp ONTAP Controller        3
107.37 GB / 107.37 GB  4 KiB + 0 B   FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5a32407351c711eaaa4800a098df41bd:subsystem.test
\
+- nvme0 fc traddr=nn-0x207300a098dfdd91:pn-0x207400a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live
inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device   Vserver  Namespace Path             NSID   UUID   Size
-------  --------- -------------------------  ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/ol_157_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/ol_157_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

Enable 1MB I/O size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> ⓘ This is not applicable to Qlogic NVMe/FC hosts.

OL 7

NVMe/FC Host Configuration for Oracle Linux 7.9 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.6 or later for Oracle Linux 7.9. Oracle Linux 7.9 host can run both NVMe and SCSI traffic through the same fibre channel (FC) initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations see the NetApp Interoperability Matrix.

> ⓘ You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

- Native NVMe/FC auto-connect scripts are not available in the `nvme-cli` package. Use the HBA vendor provided external auto-connect scripts.
- By default, round-robin load balancing is not enabled in NVMe Multipath. You must write a udev rule to enable this functionality. Steps are provided in the section on Enabling NVMe/FC on Oracle Linux 7.9.
- There is no sanlun support for NVMe/FC and, as a consequence, no Linux Unified Host Utilities (LUHU) support for NVMe/FC on Oracle Linux 7.9. Use the ONTAP command output available as part of the NetApp plug-in included in the native nvme-cli.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC

1. Install Oracle Linux 7.9 on the server.
2. After the installation is complete, verify that you are running the supported Unbreakable Enterprise kernel. See the NetApp Interoperability Matrix.

```
# uname -r
5.4.17-2011.6.2.el7uek.x86_64
```

3. Upgrade the `nvme-cli` package.

```
# rpm -qa | grep nvme-cli
nvme-cli-1.8.1-3.el7.x86_64
```

4. Add the string below as a separate udev rule at `/lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules`. This enables round-robin load balancing for NVMe multipath.

```
# cat /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules
# Enable round-robin for NetApp ONTAP
ACTION=="add", SUBSYSTEMS=="nvme-subsystem", ATTRS{model}=="NetApp ONTAP
Controller          ", ATTR{iopolicy}="round-robin"
```

5. On the Oracle Linux L 7.9 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:497ad959-e6d0-4987-8dc2-a89267400874
```

```
*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
------- --------- -------------------------------------- -----------
ol_157_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:497ad959-e6d0-4987-8dc2-a89267400874
```

If the `hostnqn` strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP array subsystem to match to host NQN string from `etc/nvme/hostnqn` on the host.

6. Reboot the host.

Configure the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. NVMe support in lpfc is already enabled by default:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

   Newer lpfc drivers (both inbox and outbox) have lpfc_enable_fc4_type default set to 3. Therefore, you do not need to set this explicitly in the `/etc/modprobe.d/lpfc.conf`.

3. Next, install the recommended lpfc auto-connect scripts:

```
# rpm -ivh nvmefc-connect-12.8.264.0-1.noarch.rpm
```

4. Verify that the auto-connect scripts are installed.

```
# rpm -qa | grep nvmefc
nvmefc-connect-12.8.264.0-1.noarch
```

5. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62

# cat /sys/class/fc_host/host*/port_state
Online
Online
```

6. Verify that the NVMe/FC initiator ports are enabled and able to see the target ports, and all are up and running.

   In the example below, only a single initiator port has been enabled and connected with two target LIFs as seen in the below output:

```
# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2947 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
```

Validate NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

In the above example, two namespaces are mapped to the Oracle Linux 7.9 ANA host. These are visible through four target LIFs: two local node LIFs and two other partner/remote node LIFs. This setup shows as two ANA Optimized and two ANA Inaccessible paths for each namespace on the host.

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- -------------------- -----------------------
/dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
53.69 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 – NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.ol_157_nvme_
ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device   Vserver  Namespace Path             NSID   UUID   Size
-------  -------- ------------------------   ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/ol_157_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/ol_157_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

Enable 1MB I/O size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> ⓘ This is not applicable to Qlogic NVMe/FC hosts.

NVMe/FC Host Configuration for Oracle Linux 7.8 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.6 or later for Oracle Linux 7.8. Oracle Linux 7.8 host can run both NVMe and SCSI traffic through the same fibre channel (FC) initiator adapter ports. Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same FC adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations see the NetApp Interoperability Matrix.

> ⓘ You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

- Native NVMe/FC auto-connect scripts are not available in the nvme-cli package. Use the HBA vendor provided external auto-connect scripts.
- By default, round-robin load balancing is not enabled in NVMe Multipath. You must write a udev rule to enable this functionality. Steps are provided in the section on Enabling NVMe/FC on Oracle Linux 7.8.
- There is no sanlun support for NVMe/FC and, as a consequence, no Linux Unified Host Utilities (LUHU) support for NVMe/FC on Oracle Linux 7.8. Use the ONTAP command output available as part of the NetApp plug-in included in the native nvme-cli.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enabling NVMe/FC

1. Install Oracle Linux 7.8 on the server.
2. After the installation is complete, verify that you are running the supported Unbreakable Enterprise kernel. See the NetApp Interoperability Matrix.

```
# uname -r
4.14.35-1902.9.2.el7uek
```

3. Upgrade the nvme-cli package.

```
# rpm -qa | grep nvme-cli
nvme-cli-1.8.1-3.el7.x86_64
```

4. Add the string below as a separate udev rule at /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules. This enables round-robin load balancing for NVMe multipath.

```
# cat /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules
# Enable round-robin for NetApp ONTAP
ACTION=="add", SUBSYSTEM=="nvme-subsystem", ATTR{model}=="NetApp ONTAP
Controller", ATTR{iopolicy}="round-robin"
```

5. On the Oracle Linux L 7.8 host, check the host NQN string at /etc/nvme/hostnqn and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

```
*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
------- --------- -------------------------------------- -----------
ol_157_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

If the hostnqn strings do not match, you should use the vserver modify command to update the host NQN string on your corresponding ONTAP array subsystem to match to host NQN string from etc/nvme/hostnqn on the host.

6. Reboot the host.

Configuring the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters, see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. NVMe support in lpfc is already enabled by default:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

   Newer lpfc drivers (both inbox and outbox) have lpfc_enable_fc4_type default set to 3. Therefore, you do not need to set this explicitly in the /etc/modprobe.d/lpfc.conf.

3. Next, install the recommended lpfc auto-connect scripts:

```
# rpm -ivh nvmefc-connect-12.4.65.0-1.noarch.rpm
```

4. Verify that the auto-connect scripts are installed.

```
# rpm -qa | grep nvmefc
nvmefc-connect-12.4.65.0-1.noarch
```

5. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62

# cat /sys/class/fc_host/host*/port_state
Online
Online
```

6. Verify that the NVMe/FC initiator ports are enabled and able to see the target ports, and all are up and running.

   In the example below, only a single initiator port has been enabled and connected with two target LIFs as seen in the below output:

```
# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2947 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
```

Validating NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

In the above example, two namespaces are mapped to the Oracle Linux 7.8 ANA host. These are visible through four target LIFs: two local node LIFs and two other partner/remote node LIFs. This setup shows as two ANA Optimized and two ANA Inaccessible paths for each namespace on the host.

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- -------------------- -----------------------
/dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
53.69 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 – NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.ol_157_nvme_
ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device   Vserver  Namespace Path             NSID   UUID   Size
-------  -------- -------------------------  ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/ol_157_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/ol_157_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

Enabling 1MB I/O size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

NVMe/FC Host Configuration for Oracle Linux 7.7 with ONTAP

Supportability

NVMe/FC is supported on ONTAP 9.6 or later for the following versions of Oracle Linux

- OL 7.7

  OL 7.7 host can run both NVMe & SCSI traffic through the same fibre channel initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

  For the most current list of supported configurations see the NetApp Interoperability Matrix.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

- Native NVMe/FC auto-connect scripts are not available in the nvme-cli package.
  You can use the HBA vendor provided external auto-connect scripts.
- By default, round-robin load balancing is not enabled.
  You must write a udev rule to enable this functionality. Steps are provided in the section on Enabling NVMe/FC on OL 7.7.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enabling NVMe on OL 7.7

1. Ensure the default Oracle Linux 7.7 kernel is installed.
2. Reboot the host and verify that it boots into specified OL 7.7 kernel.

```
# uname -r
4.14.35-1902.9.2.el7uek
```

3. Upgrade to the nvme-cli-1.8.1-3.el7 package.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.8.1-3.el7.x86_64
```

4. Add the string below as a separate udev rule at `/lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules`. This enables round-robin load balancing for NVMe multipath.

```
# Enable round-robin for NetApp ONTAP
ACTION==”add”, SUBSYSTEM==”nvme-subsystem”, ATTR{model}==”NetApp ONTAP
Controller”, ATTR{iopolicy}=”round-robin
```

5. On the OL 7.7 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

```
*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
------- --------- -------------------------------------- -----------
ol_157_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

If the host NQN strings do not match, you should use the vserver modify command to update the host NQN string on your corresponding ONTAP array subsystem to match to host NQN string from `/etc/nvme/hostnqn` on the host.

1. Reboot the host.

Configuring the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Copy and install the Broadcom outbox auto-connect scripts package.

```
# rpm -ivh nvmefc-connect-12.4.65.0-1.noarch.rpm
```

3. Reboot the host.
4. Verify that you are using the recommended Broadcom lpfc firmware, native inbox driver & outbox auto-connect package versions. For a list of supported versions, see the NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/fwrev
12.4.243.17, sil-4.2.c
12.4.243.17, sil-4.2.c

# cat /sys/module/lpfc/version
0:12.0.0.10

# rpm -qa | grep nvmefc
nvmefc-connect-12.4.65.0-1.noarch
```

5. Verify that lpfc_enable_fc4_type is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

6. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

7. Verify that the NVMe/FC initiator ports are enabled, running and able to see the target LIFs.

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2977 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
NVME Statistics
…
```

Validating NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y

# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- -------------------- -----------------------
/dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
53.69 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 – NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.ol_157_nvme_
ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device   Vserver  Namespace Path             NSID   UUID   Size
-------  -------- ------------------------   ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/ol_157_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/ol_157_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

Enabling 1MB I/O size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

LPFC verbose logging

Set the lpfc driver for NVMe/FC.

Steps

1. Set the `lpfc_log_verbose` driver setting to any of the following values to log NVMe/FC events.

```
#define LOG_NVME 0x00100000 /* NVME general events. */
#define LOG_NVME_DISC 0x00200000 /* NVME Discovery/Connect events. */
#define LOG_NVME_ABTS 0x00400000 /* NVME ABTS events. */
#define LOG_NVME_IOERR 0x00800000 /* NVME IO Error events. */
```

2. After setting the values, run the `dracut-f` command and reboot the host.
3. Verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose 15728771
```

RHEL

RHEL 9

NVMe-oF host configuration for RHEL 9.3 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Red Hat Enterprise Linux (RHEL) 9.3 with Asymmetric

Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for NVMe-oF host configuration for RHEL 9.3 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA)without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

RHEL 9.3 has in-kernel NVMe multipath enabled for NVMe namespaces by default; therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported RHEL 9.3 software versions.

Steps

1. Install RHEL 9.3 on the server. After the installation is complete, verify that you are running the specified RHEL 9.3 kernel:

   ```
   # uname -r
   ```

   Example output:

   ```
   5.14.0-362.8.1.el9_3.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

   Example output:

   ```
   nvme-cli-2.4-10.el9.x86_64
   ```

3. Install the `libnvme` package:

```
#rpm -qa|grep libnvme
```

Example output

```
libnvme-1.4-7.el9.x86_64
```

4. On the RHEL 9.3 host, check the hostnqn string at `/etc/nvme/hostnqn`:

```
# cat /etc/nvme/hostnqn
```

Example output

```
nqn.2014-08.org.nvmexpress:uuid:060fd513-83be-4c3e-aba1-52e169056dcf
```

5. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_nvme147
```

Example output:

```
Vserver     Subsystem          Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme147   rhel_147_LPe32002    nqn.2014-
08.org.nvmexpress:uuid:060fd513-83be-4c3e-aba1-52e169056dcf
```

> If the `hostnqn` strings do not match, use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   14.2.539.16, sli-4:2:c
   14.2.539.16, sli-4:2:c

   # cat /sys/module/lpfc/version
   0:14.2.0.12
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b3c081f
0x100000109b3c0820
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b3c081f WWNN x200000109b3c081f DID
x062300 ONLINE
NVME RPORT       WWPN x2143d039ea165877 WWNN x2142d039ea165877 DID
x061b15 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2145d039ea165877 WWNN x2142d039ea165877 DID
x061115 TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 000000040b Cmpl 000000040b Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000001f5c4538 Issue 000000001f58da22 OutIO
fffffffffffc94ea
abort 00000630 noxri 00000000 nondlp 00001071 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000630 Err 0001bd4a
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b3c0820 WWNN x200000109b3c0820 DID
x062c00 ONLINE
NVME RPORT       WWPN x2144d039ea165877 WWNN x2142d039ea165877 DID
x060215 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2146d039ea165877 WWNN x2142d039ea165877 DID
x061815 TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 000000040b Cmpl 000000040b Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000001f5c3618 Issue 000000001f5967a4 OutIO
fffffffffffd318c
abort 00000629 noxri 00000000 nondlp 0000044e qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000629 Err 0001bd3d
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the RHEL 9.3 GA kernel has the latest fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

   ```
   # cat /sys/class/fc_host/host*/symbolic_name
   ```

   Example output

   ```
   QLE2772 FW:v9.10.11 DVR:v10.02.08.200-k
   QLE2772 FW:v9.10.11 DVR:v10.02.08.200-k
   ```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

   ```
   # cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
   1
   ```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

   ```
   # cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
   256
   ```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

   ```
   # nvme discover -t tcp -w 192.168.167.1 -a 192.168.167.16

   Discovery Log Number of Records 8, Generation counter 10
   =====Discovery Log Entry 0======
   trtype:  tcp
   adrfam:  ipv4
   subtype: current discovery subsystem
   treq:    not specified
   portid:  0
   trsvcid: 8009
   subnqn:  nqn.1992-
   08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
   traddr:  192.168.166.17
   eflags:  explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 1======
   trtype:  tcp
   adrfam:  ipv4
   subtype: current discovery subsystem
   treq:    not specified
   portid:  1
   trsvcid: 8009
   subnqn:  nqn.1992
   08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
   traddr:  192.168.167.17
   eflags:  explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 2======
   trtype:  tcp
   adrfam:  ipv4
   subtype: current discovery subsystem
   treq:    not specified
   portid:  2
   ```

```
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
traddr:  192.168.166.16
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 3======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  3
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
traddr:  192.168.167.16
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
...
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
#nvme discover -t tcp -w 192.168.166.5 -a 192.168.166.22
#nvme discover -t tcp -w 192.168.166.5 -a 192.168.166.23
#nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.22
#nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.23
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
#   nvme    connect-all -t  tcp -w  192.168.166.1   -a  192.168.166.16
-l   1800
#   nvme    connect-all -t  tcp -w  192.168.166.1   -a  192.168.166.17
-l   1800
#   nvme    connect-all -t  tcp -w  192.168.167.1   -a  192.168.167.16
-l   1800
#   nvme    connect-all -t  tcp -w  192.168.167.1   -a  192.168.167.17
-l   1800
```

Validate NVMe-oF

You can use the following procedure to validate NVME-oF.

Steps

1. Verify that the in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, model set to NetApp ONTAP Controller and load balancing iopolicy set to round-robin) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node          SN                   Model
---------------------------------------------------------
/dev/nvme5n21 81CYrNQlis3WAAAAAAAB  NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 21.47 GB / 21.47 GB   4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme5n21
```

Example output:

```
nvme-subsys4 - NQN=nqn.1992-
08.com.netapp:sn.e80cc121ca6911ed8cbdd039ea165590:subsystem.rhel_
147_LPE32002
\
 +- nvme2 fc traddr=nn-0x2142d039ea165877:pn-
0x2144d039ea165877,host_traddr=nn-0x200000109b3c0820:pn-
0x100000109b3c0820 live optimized
 +- nvme3 fc traddr=nn-0x2142d039ea165877:pn-
0x2145d039ea165877,host_traddr=nn-0x200000109b3c081f:pn-
0x100000109b3c081f live non-optimized
 +- nvme4 fc traddr=nn-0x2142d039ea165877:pn-
0x2146d039ea165877,host_traddr=nn-0x200000109b3c0820:pn-
0x100000109b3c0820 live non-optimized
 +- nvme6 fc traddr=nn-0x2142d039ea165877:pn-
0x2143d039ea165877,host_traddr=nn-0x200000109b3c081f:pn-
0x100000109b3c081f live optimized
```

NVMe/TCP

```
# nvme list-subsys /dev/nvme1n1
```

Example output:

```
nvme-subsys1 - NQN=nqn.1992- 08.com.netapp:sn.
bbfb4ee8dfb611edbd07d039ea165590:subsystem.rhel_tcp_95
+- nvme1 tcp
traddr=192.168.167.16,trsvcid=4420,host_traddr=192.168.167.1,src_add
r=192.168.167.1 live
+- nvme2 tcp
traddr=192.168.167.17,trsvcid=4420,host_traddr=192.168.167.1,src_add
r=192.168.167.1 live
+- nvme3 tcp
traddr=192.168.167.17,trsvcid=4420,host_traddr=192.168.166.1,src_add
r=192.168.166.1 live
+- nvme4 tcp
traddr=192.168.166.16,trsvcid=4420,host_traddr=192.168.166.1,src_add
r=192.168.166.1 live
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device        Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1 vs_tcp           /vol/vol1/ns1



NSID       UUID                                   Size
------------------------------------------------------------
1          6fcb8ea0-dc1e-4933-b798-8a62a626cb7f 21.47GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output


```
{

"ONTAPdevices" : [
{

"Device" : "/dev/nvme1n1",
"Vserver" : "vs_tcp_95",
"Namespace_Path" : "/vol/vol1/ns1",
"NSID" : 1,
"UUID" : "6fcb8ea0-dc1e-4933-b798-8a62a626cb7f",
"Size" : "21.47GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 5242880
},

]
}
```

Known issues

There are no known issues for the NVMe-oF host configuration for RHEL 9.3 with ONTAP release.

NVMe-oF host configuration for RHEL 9.2 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Red Hat Enterprise Linux (RHEL) 9.2 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for RHEL 9.2 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- RHEL 9.2 has in-kernel NVMe multipath enabled for NVMe namespaces by default, therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported RHEL 9.2 software versions.

Steps

1. Install RHEL 9.2 on the server. After the installation is complete, verify that you are running the specified RHEL 9.2 kernel.

   ```
   # uname -r
   ```

   Example output:

   ```
   5.14.0-284.11.1.el9_2.x86_64
   ```

2. Install the `nvme-cli` package:

   ```
   # rpm -qa|grep nvme-cli
   ```

Example output:

```
nvme-cli-2.2.1-2.el9.x86_64
```

3. Install the `libnvme` package:

```
#rpm -qa|grep libnvme
```

Example output

```
libnvme-1.2-2.el9.x86_64
```

4. On the RHEL 9.2 host, check the hostnqn string at `/etc/nvme/hostnqn`:

```
# cat /etc/nvme/hostnqn
```

Example output

```
nqn.2014-08.org.nvmexpress:uuid:4c4c4544-0032-3310-8033-b8c04f4c5132
```

5. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_ol_nvme
```

Example output:

```
Vserver     Subsystem          Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme207   rhel_207_LPe32002    nqn.2014-
08.org.nvmexpress:uuid:325e7554-1f9b-11ec-8489-3a68dd61a4df
```

> If the `hostnqn` strings do not match, use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model.

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver.

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   14.0.639.18, sli-4:2:c
   14.0.639.18, sli-4:2:c

   # cat /sys/module/lpfc/version
   0:12.8.0.11
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the RHEL 9.2 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

   ```
   # cat /sys/class/fc_host/host*/symbolic_name
   ```

   Example output

   ```
   QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
   QLE2742 FW:v9.08.02 DVR:v10.02.00.106-k
   ```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

   ```
   # cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
   1
   ```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

   ```
   # cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
   256
   ```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

```
# nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.22

Discovery Log Number of Records 8, Generation counter 18
=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.166.23
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  1
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.166.22
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  2
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:discovery
traddr:  192.168.167.23
eflags:  explicit discovery connections, duplicate discovery information
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data.

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
#nvme discover -t tcp -w 192.168.166.5 -a 192.168.166.22
#nvme discover -t tcp -w 192.168.166.5 -a 192.168.166.23
#nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.22
#nvme discover -t tcp -w 192.168.167.5 -a 192.168.167.23
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
#   nvme    connect-all -t  tcp -w  192.168.166.5   -a  192.168.166.22
-l  1800
#   nvme    connect-all -t  tcp -w  192.168.166.5   -a  192.168.166.23
-l  1800
#   nvme    connect-all -t  tcp -w  192.168.167.5   -a  192.168.167.22
-l  1800
#   nvme    connect-all -t  tcp -w  192.168.167.5   -a  192.168.167.23
-l  1800
```

Validate NVMe-oF

You can use the following procedure to validate NVME-oF.

Steps

1. Verify that the in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, model set to NetApp ONTAP Controller and load balancing iopolicy set to round-robin) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 81CZ5BQuUNfGAAAAAAAB   NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 21.47 GB / 21.47 GB   4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys4 - NQN=nqn.1992-
08.com.netapp:sn.8763d311b2ac11ed950ed039ea951c46:subsystem.rhel_207
_LB \
+- nvme1 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a7d039ea954d17,host_traddr=nn-0x200000109b1b95ef:pn-
0x100000109b1b95ef live optimized
+- nvme2 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a8d039ea954d17,host_traddr=nn-0x200000109b1b95f0:pn-
0x100000109b1b95f0 live optimized
+- nvme3 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20aad039ea954d17,host_traddr=nn-0x200000109b1b95f0:pn-
0x100000109b1b95f0 live non-optimized
+- nvme5 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a9d039ea954d17,host_traddr=nn-0x200000109b1b95ef:pn-
0x100000109b1b95ef live non-optimized
```

NVMe/TCP

```
# nvme list-subsys /dev/nvme1n1
```

Example output:

```
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.c680f5bcae1411ed8639d039ea951c46:subsystem.rhel_tcp
97 \
+- nvme1 tcp
traddr=192.168.167.23,trsvcid=4420,host_traddr=192.168.167.5 live
non-optimized
+- nvme2 tcp
traddr=192.168.167.22,trsvcid=4420,host_traddr=192.168.167.5 live
non-optimized
+- nvme3 tcp
traddr=192.168.166.23,trsvcid=4420,host_traddr=192.168.166.5 live
optimized
+- nvme4 tcp
traddr=192.168.166.22,trsvcid=4420,host_traddr=192.168.166.5 live
optimized
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device        Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1 vs_tcp           /vol/vol1/ns1



NSID       UUID                                   Size
------------------------------------------------------------
1          79c2c569-b7fa-42d5-b870-d9d6d7e5fa84 21.47GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{

"ONTAPdevices" : [
{

"Device" : "/dev/nvme0n1",
"Vserver" : "vs_tcp79",
"Namespace_Path" : "/vol/vol1/ns1",
"NSID" : 1,
"UUID" : "79c2c569-b7fa-42d5-b870-d9d6d7e5fa84",
"Size" : "21.47GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 5242880
},

]
}
```

Known issues

There are no known issues.

NVMe-oF host configuration for RHEL 9.1 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and NVMe/TCP) is supported with RHEL 9.1 with Asymmetric Namespace Access (ANA) that is required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. This document contains the details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 9.1 and ONTAP as the target.

The following support is available for the NVMe-oF host configuration for RHEL 9.1 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

Refer to the NetApp Interoperability Matrix Tool for accurate details regarding supported configurations.

Features

RHEL 9.1 includes support for in-kernel NVMe multipath for NVMe namespaces enabled by default, without the need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 9.1 on the server.
2. After the installation is complete, verify that you are running the specified RHEL 9.1 kernel. See the NetApp Interoperability Matrix Tool for the most current list of supported versions.

   Example:

   ```
   # uname -r
    5.14.0-162.6.1.el9_1.x86_64
   ```

3. Install the `nvme-cli` package:

   Example:

```
# rpm -qa|grep nvme-cli
nvme-cli-2.0-4.el9.x86_64
```

4. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. Example:

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:325e7554-1f9b-11ec-8489-3a68dd61a4df


::> vserver nvme subsystem host show -vserver vs_nvme207
Vserver     Subsystem       Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme207 rhel_207_LPe32002     nqn.2014-
08.org.nvmexpress:uuid:325e7554-1f9b-11ec-8489-3a68dd61a4df
```

> If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

5. Reboot the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix Tool for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2

# cat /sys/class/scsi_host/host*/modeldesc

Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. See the NetApp Interoperability Matrix Tool for the most current list of supported adapter driver and firmware versions.

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.11, sli-4:2:c
14.0.505.11, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:14.2.0.5
```

3. Verify that `lpfc_enable_fc4_type` is set to 3

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1b95ef
0x100000109b1b95f0
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1b95ef WWNN x200000109b1b95ef DID
x061700 ONLINE
NVME RPORT       WWPN x2035d039ea1308e5 WWNN x2082d039ea1308e5 DID
x062f05 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2083d039ea1308e5 WWNN x2082d039ea1308e5 DID
x062407 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000000001df6c Issue 000000000001df6e OutIO
0000000000000002
        abort 00000000 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000000 Err 00000004

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1b95f0 WWNN x200000109b1b95f0 DID
x061400 ONLINE
NVME RPORT       WWPN x2036d039ea1308e5 WWNN x2082d039ea1308e5 DID
x061605 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2037d039ea1308e5 WWNN x2082d039ea1308e5 DID
x062007 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000000001dd28 Issue 000000000001dd29 OutIO
0000000000000001
        abort 00000000 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000000 Err 00000004
```

Marvell/QLogic FC adapter for NVMe/FC

The native inbox `qla2xxx` driver included in the RHEL 9.1 kernel has the latest fixes which are essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions using the following command:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2772 FW:v9.08.02 DVR:v10.02.07.400-k-debug
QLE2772 FW:v9.08.02 DVR:v10.02.07.400-k-debug
```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as an NVMe/FC initiator using the following command:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51

Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Verify that the other NVMe/TCP initiator-target LIF combos can successfully fetch discovery log page data. For example:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Make sure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) while running the `connect-all` command so that it would retry for a longer period of time in the event of a path loss. For example:

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is indeed enabled by checking:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example:

```
# nvme list
Node           SN                     Model                     Namespace
------------   ---------------------- ----------------------------------
/dev/nvme0n1   81CZ5BQuUNfGAAAAAAAB   NetApp ONTAP Controller   1

Usage                 Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example:

   Example (a):

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys10 - NQN=nqn.1992-
08.com.netapp:sn.82e7f9edc72311ec8187d039ea14107d:subsystem.rhel_131_QLe
2742
\
 +- nvme2 fc traddr=nn-0x2038d039ea1308e5:pn-
0x2039d039ea1308e5,host_traddr=nn-0x20000024ff171d30:pn-
0x21000024ff171d30 live non-optimized
 +- nvme3 fc traddr=nn-0x2038d039ea1308e5:pn-
0x203cd039ea1308e5,host_traddr=nn-0x20000024ff171d31:pn-
0x21000024ff171d31 live optimized
 +- nvme4 fc traddr=nn-0x2038d039ea1308e5:pn-
0x203bd039ea1308e5,host_traddr=nn-0x20000024ff171d30:pn-
0x21000024ff171d30 live optimized
 +- nvme5 fc traddr=nn-0x2038d039ea1308e5:pn-
0x203ad039ea1308e5,host_traddr=nn-0x20000024ff171d31:pn-
0x21000024ff171d31 live non-optimized
```

   Example (b):

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.bf0691a7c74411ec8187d039ea14107d:subsystem.rhel_tcp_133
\
 +- nvme1 tcp
traddr=192.168.166.21,trsvcid=4420,host_traddr=192.168.166.5 live non-
optimized
 +- nvme2 tcp
traddr=192.168.166.20,trsvcid=4420,host_traddr=192.168.166.5 live
optimized
 +- nvme3 tcp
traddr=192.168.167.21,trsvcid=4420,host_traddr=192.168.167.5 live non-
optimized
 +- nvme4 tcp
traddr=192.168.167.20,trsvcid=4420,host_traddr=192.168.167.5 live
optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device.

   Example (a):

```
# nvme netapp ontapdevices -o column
Device        Vserver          Namespace Path
---------     -------
--------------------------------------------------
/dev/nvme0n1 vs_tcp79     /vol/vol1/ns1

NSID  UUID                                   Size
----  ------------------------------         ------
1     79c2c569-b7fa-42d5-b870-d9d6d7e5fa84  21.47GB


# nvme netapp ontapdevices -o json
{

  "ONTAPdevices" : [
  {

      "Device" : "/dev/nvme0n1",
      "Vserver" : "vs_tcp79",
      "Namespace_Path" : "/vol/vol1/ns1",
      "NSID" : 1,
      "UUID" : "79c2c569-b7fa-42d5-b870-d9d6d7e5fa84",
      "Size" : "21.47GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 5242880
    },

]

}
```

Example (b):

```
# nvme netapp ontapdevices -o column

Device           Vserver                   Namespace Path
---------------- -------------------------
------------------------------------
/dev/nvme1n1     vs_tcp_133                /vol/vol1/ns1

NSID UUID                                   Size
-------------------------------------------------------
1    1ef7cb56-bfed-43c1-97c1-ef22eeb92657   21.47GB

# nvme netapp ontapdevices -o json
{
  "ONTAPdevices":[
    {
      "Device":"/dev/nvme1n1",
      "Vserver":"vs_tcp_133",
      "Namespace_Path":"/vol/vol1/ns1",
      "NSID":1,
      "UUID":"1ef7cb56-bfed-43c1-97c1-ef22eeb92657",
      "Size":"21.47GB",
      "LBA_Data_Size":4096,
      "Namespace_Size":5242880
    },
  ]

}
```

Known issues

The NVMe-oF host configuration for RHEL 9.1 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1503468	`nvme list-subsys` command returns repeated nvme controller list for a given subsystem	The `nvme list-subsys` command should return a unique list of nvme controllers associated to a given subsystem. In RHEL 9.1, the `nvme list-subsys` command returns nvme controllers with its respective ANA state for all namespaces that belong to a given subsystem. However, the ANA state is a per-namespace attribute therefore, it would be ideal to display unique nvme controller entries with the path state if you list the subsystem command syntax for a given namespace.	2130106

NVMe-oF host configuration for RHEL 9.0 with ONTAP

NVMe-oF (including NVMe/FC and NVMe/TCP) is supported with RHEL 9.0 with Asymmetric Namespace Access (ANA) required for surviving storage failovers (SFOs) on the ONTAP array. ANA is the ALUA equivalent in the NVM-oF environment, and is currently implemented with in-kernel NVMe Multipath. This document contains the details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 9.0 and ONTAP as the target.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- Beginning with RHEL 9.0, NVMe/TCP is no longer a technology preview feature (unlike RHEL 8) but a fully supported enterprise feature itself.
- Beginning with RHEL 9.0, in-kernel NVMe multipath is enabled for NVMe namespaces by default, without the need for explicit settings (unlike RHEL 8).

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe Multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 9.0 on the server.
2. After the installation is complete, verify that you are running the specified RHEL 9.0 kernel. See NetApp Interoperability Matrix for the most current list of supported versions.

```
# uname -r
5.14.0-70.13.1.el9_0.x86_64
```

3. Install the `nvme-cli` package.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.16-3.el9.x86_64
```

4. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. For example,

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
```

```
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
Vserver     Subsystem Host     NQN
----------- ---------------
-----------------------------------------------------------
vs_fcnvme_14 nvme_141_1 nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-
4cf5-97b3-1b5d986345d1
```

> If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string from `/etc/nvme/hostnqn` on the host.

5. Reboot the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. For additional details on supported adapters, see the NetApp Interoperability Matrix Tool.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. For the most current list of supported adapter driver and firmware versions, see NetApp Interoperability Matrix.

```
# cat /sys/class/scsi_host/host*/fwrev
12.8.351.47, sli-4:2:c
12.8.351.47, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:14.0.0.4
```

3. Verify that `lpfc_enable_fc4_type` is set to `3`.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and you are able to see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic

The native inbox qla2xxx driver included in the RHEL 9.0 kernel has the latest fixes, essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.06.02 DVR:v10.02.00.200-k
QLE2742 FW:v9.06.02 DVR:v10.02.00.200-k
```

1. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as a NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port is able to fetch discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51

Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Similarly, verify that the other NVMe/TCP initiator-target LIF combos are able to successfully fetch the discovery log page data. For example,

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it would retry for a longer period of time in the event of a path loss. For example,

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is indeed enabled by checking:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMf settings (for example, model set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host.

   Example (a):

```
# nvme list
Node         SN                   Model                    Namespace
Usage
------      ---------------------------------------
------------------------
/dev/nvme0n1 814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller  1
85.90 GB / 85.90 GB

Format         FW Rev
---------------------
4 KiB + 0 B   FFFFFFFF
```

Example (b):

```
# nvme list
Node           SN                   Model                    Namespace
Usage
----------------------------------------------------
------------------------------------
/dev/nvme0n1   81CZ5BQuUNfGAAAAAAAB NetApp ONTAP Controller   1
85.90 GB / 85.90 GB

Format         FW Rev
-----------------------
4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has a proper ANA status.

   Example (a):

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_141_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

Example (b):

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
\
+- nvme0 tcp traddr=192.168.1.51 trsvcid=4420 host_traddr=192.168.1.8
live optimized
+- nvme10 tcp traddr=192.168.2.56 trsvcid=4420 host_traddr=192.168.2.9
live optimized
+- nvme15 tcp traddr=192.168.2.57 trsvcid=4420 host_traddr=192.168.2.9
live non-optimized
+- nvme5 tcp traddr=192.168.1.52 trsvcid=4420 host_traddr=192.168.1.8
live non-optimized
```

5. Verify the NetApp plug-in displays proper values for each ONTAP namespace device.

   Example (a):

```
# nvme netapp ontapdevices -o column
Device       Vserver        Namespace Path
NSID
----------------------- ------------------------------
-------------------------
/dev/nvme0n1  vs_fcnvme_141  /vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns   1

UUID                                   Size
--------------------------------------------
72b887b1-5fb6-47b8-be0b-33326e2542e2   85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_fcnvme_141",
        "Namespace_Path" : "/vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    }
  ]
}
```

Example (b):

```
# nvme netapp ontapdevices -o column
Device               Vserver                   Namespace Path
--------------------- -------------------------
------------------------------------
/dev/nvme0n1         vs_tcp_118
/vol/tcpnvme_118_1_0_0/tcpnvme_118_ns

NSID   UUID                               Size
-------------------------------------------------------
1     4a3e89de-b239-45d8-be0c-b81f6418283c 85.90GB
```

```
# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
     "Device" : "/dev/nvme0n1",
      "Vserver" : "vs_tcp_118",
      "Namespace_Path" : "/vol/tcpnvme_118_1_0_0/tcpnvme_118_ns",
      "NSID" : 1,
      "UUID" : "4a3e89de-b239-45d8-be0c-b81f6418283c",
      "Size" : "85.90GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 20971520
    },
  ]

}
```

Known issues

The NVMe-oF host configuration for RHEL 9.0 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1479047	RHEL 9.0 NVMe-oF hosts create duplicate Persistent Discovery Controllers	On NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Red Hat Enterprise Linux (RHEL) 9.0 with an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	2087000

RHEL 8

NVMe-oF host configuration for RHEL 8.9 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Red Hat Enterprise Linux (RHEL) 8.9 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for NVMe-oF host configuration for RHEL 8.9 with ONTAP:

* Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Known limitations

* In-kernel NVMe multipath is disabled by default for RHEL 8.9 NVMe-oF hosts. Therefore, you need to enable it manually.
* On RHEL 8.9 hosts, NVMe/TCP is a technology preview feature due to open issues.
* SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel multipath

You can use the following procedure to enable in-kernel multipath.

Steps

1. Install RHEL 8.9 on the host server.
2. After the installation is complete, verify that you are running the specified RHEL 8.9 kernel:

   ```
   # uname -r
   ```

   Example output

   ```
   4.18.0-513.5.1.el8_9.x86_64
   ```

3. Install the nvme-cli package:

   ```
   rpm -qa|grep nvme-cli
   ```

   Example output

   ```
   nvme-cli-1.16-9.el8.x86_64
   ```

4. Enable in -kernel NVMe multipath:

   ```
   # grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
   4.18.0-513.5.1.el8_9.x86_64
   ```

5. On the host, check the host NQN string at `/etc/nvme/hostnqn`:

```
# cat /etc/nvme/hostnqn
```

Example output

```
nqn.2014-08.org.nvmexpress:uuid:4c4c4544-0032-3410-8035-b8c04f4c5132
```

6. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
```

Example output

```
Vserver     Subsystem       Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme101 rhel_101_QLe2772    nqn.2014-08.org.nvmexpress:
uuid:4c4c4544-0032-3410-8035-b8c04f4c5132
```

> If the host NQN strings do not match, you can use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

7. Reboot the host.

> If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This should exclude the ONTAP namespaces from dm-multipath and prevent dm-multipath from claiming these namespace devices. You can do this by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
  enable_foreign  NONE
}
```

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

```
# cat /sys/class/scsi_host/host*/modelname
```

Example output:

```
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
```

Example output:

```
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

```
# cat /sys/class/scsi_host/host*/fwrev
14.2.539.16, sli-4:2:c
14.2.539.16, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:14.0.0.21
```

For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec88
0x10000090fae0ec89
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec88 WWNN x20000090fae0ec88 DID
x0a1300 ONLINE
NVME RPORT       WWPN x2049d039ea36a105 WWNN x2048d039ea36a105 DID
x0a0c0a TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000024 Cmpl 0000000024 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000000001aa Issue 00000000000001ab OutIO
0000000000000001
        abort 00000002 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000002 Err 00000003
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x10000090fae0ec89 WWNN x20000090fae0ec89 DID
x0a1200 ONLINE
NVME RPORT       WWPN x204ad039ea36a105 WWNN x2048d039ea36a105 DID
x0a080a TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000024 Cmpl 0000000024 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000000001ac Issue 00000000000001ad OutIO
0000000000000001
        abort 00000002 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000002 Err 00000003
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the RHEL 8.9 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware

versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
```

Example output

```
QLE2742 FW: v9.10.11 DVR: v10.02.08.200-k
QLE2742 FW: v9.10.11 DVR: v10.02.08.200-k
```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

   ```
   # nvme discover -t tcp -w 192.168.111.79 -a 192.168.111.14 -l 1800

   Discovery Log Number of Records 8, Generation counter 18
   =====Discovery Log Entry 0======
   trtype:  tcp
   adrfam:  ipv4
   subtype: unrecognized
   treq:    not specified.
   portid:  0
   trsvcid: 8009
   subnqn:  nqn.1992-08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:
   discovery
   traddr:  192.168.211.15
   sectype: none
   =====Discovery Log Entry 1======
   trtype:  tcp
   adrfam:  ipv4
   subtype: unrecognized
   treq:    not specified.
   portid:  1
   trsvcid: 8009
   subnqn:  nqn.1992-08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:
   discovery
   traddr:  192.168.111.15
   sectype: none ..........
   ```

2. Verify that the other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

```
# nvme  discovery   -t   tcp    -w  192.168.111.79   -a 192.168.111.14
# nvme  discovery   -t   tcp    -w  192.168.111.79   -a 192.168.111.15
# nvme  discovery   -t   tcp    -w  192.168.211.79   -a 192.168.211.14
# nvme  discovery   -t   tcp    -w  192.168.211.79   -a 192.168.211.15
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme  connect-all -t  tcp -w  192.168.111.79  -a  192.168.111.14  -l
1800
# nvme  connect-all -t  tcp -w  192.168.111.79  -a  192.168.111.15  -l
1800
# nvme  connect-all -t  tcp -w  192.168.211.79  -a  192.168.211.14  -l
1800
# nvme  connect-all -t  tcp -w  192.168.211.79  -a  192.168.211.15  -l
1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that the in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 81Gx7NSiKSQqAAAAAAAB   NetApp ONTAP Controller


Namespace Usage    Format             FW            Rev
-----------------------------------------------------------
1                 21.47 GB / 21.47 GB   4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme3n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.8e501f8ebafa11ec9b99d039ea359e4b:subsystem.rhel_163
_Qle2742
+- nvme0 fc traddr=nn-0x204dd039ea36a105:pn-0x2050d039ea36a105
host_traddr=nn-0x20000024ff7f4994:pn-0x21000024ff7f4994 live non-
optimized
+- nvme1 fc traddr=nn-0x204dd039ea36a105:pn-0x2050d039ea36a105
host_traddr=nn-0x20000024ff7f4994:pn-0x21000024ff7f4994 live non-
optimized
+- nvme2 fc traddr=nn-0x204dd039ea36a105:pn-0x204fd039ea36a105
host_traddr=nn-0x20000024ff7f4995:pn-0x21000024ff7f4995 live
optimized
+- nvme3 fc traddr=nn-0x204dd039ea36a105:pn-0x204ed039ea36a105
host_traddr=nn-0x20000024ff7f4994:pn-0x21000024ff7f4994 live
optimized
```

NVMe/TCP

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp
_165\
+- nvme0 tcp traddr=192.168.111.15 trsvcid=4420
host_traddr=192.168.111.79 live non-optimized
+- nvme1 tcp traddr=192.168.111.14 trsvcid=4420
host_traddr=192.168.111.79 live optimized
+- nvme2 tcp traddr=192.168.211.15 trsvcid=4420
host_traddr=192.168.211.79 live non-optimized
+- nvme3 tcp traddr=192.168.211.14 trsvcid=4420
host_traddr=192.168.211.79 live optimized
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device         Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1 vs_tcp79           /vol/vol1/ns


NSID       UUID                                   Size
------------------------------------------------------------
1          aa197984-3f62-4a80-97de-e89436360cec 21.47GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
  "ONTAPdevices”: [
    {
      "Device”: "/dev/nvme0n1",
      "Vserver”: "vs_tcp79",
      "Namespace Path”: "/vol/vol1/ns",
      "NSID”: 1,
      "UUID”: "aa197984-3f62-4a80-97de-e89436360cec",
      "Size”: "21.47GB",
      "LBA_Data_Size”: 4096,
      "Namespace Size" : 5242880
    },
]

}
```

Known issues

The NVMe-oF host configuration for RHEL 8.9 with ONTAP release has the following known issue:

NetApp Bug ID	Title	Description	Bugzilla ID
1479047	RHEL 8.9 NVMe-oF hosts create duplicate persistent discovery controllers	On NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running Red Hat Enterprise Linux (RHEL) 8.9 on an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	2087000

NVMe-oF host configuration for RHEL 8.8 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Red Hat Enterprise Linux (RHEL) 8.8 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for RHEL 8.8 with ONTAP:

- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Known limitations

- In-kernel NVMe multipath is disabled by default for RHEL 8.8 NVMe-oF hosts. Therefore, you need to enable it manually.
- On RHEL 8.8 hosts, NVMe/TCP is a technology preview feature due to open issues.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel multipath

You can use the following procedure to enable in-kernal multipath.

Steps

1. Install RHEL 8.8 on the host server.
2. After the installation is complete, verify that you are running the specified RHEL 8.8 kernel.

```
# uname -r
```

Example output

```
4.18.0-477.10.1.el8_8.x86_64
```

3. Install the nvme-cli package:

```
rpm -qa|grep nvme-cli
```

Example output

```
nvme-cli-1.16-7.el8.x86_64
```

4. Enable in -kernel NVMe multipath:

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-477.10.1.el8_8.x86_64
```

5. On the host, check the host NQN string at `/etc/nvme/hostnqn`:

```
# cat /etc/nvme/hostnqn
```

Example output

```
nqn.2014-08.org.nvmexpress:uuid:f6517cae-3133-11e8-bbff-7ed30aef123f
```

6. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

```
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
```

Example output

```
Vserver     Subsystem       Host NQN
----------- ---------------
----------------------------------------------------------
vs_nvme161 rhel_161_LPe32002    nqn.2014-
08.org.nvmexpress:uuid:f6517cae-3133-11e8-bbff-7ed30aef123f
```

If the host NQN strings do not match, you can use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

7. Reboot the host.

If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the `enable_foreign` setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
  enable_foreign  NONE
}
```

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model:

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver:

   ```
   # cat /sys/class/scsi_host/host*/fwrev
   14.0.639.18, sli-4:2:c
   14.0.639.18, sli-4:2:c
   ```

   ```
   # cat /sys/module/lpfc/version
   0:14.0.0.18
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec88 WWNN x20000090fae0ec88 DID
x0a1300 ONLINE
NVME RPORT       WWPN x2049d039ea36a105 WWNN x2048d039ea36a105 DID
x0a0c0a TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x204bd039ea36a105 WWNN x2048d039ea36a105 DID
x0a100a TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000134 Cmpl 0000000134 Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000825e567 Issue 00000000825d7ed OutIO
fffffffffffff286
abort 0000027c noxri 00000000 nondlp 00000a02 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000782 Err 000130fa

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x10000090fae0ec89 WWNN x20000090fae0ec89 DID
x0a1200 ONLINE
NVME RPORT       WWPN x204ad039ea36a105 WWNN x2048d039ea36a105 DID
x0a080a TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x204cd039ea36a105 WWNN x2048d039ea36a105 DID
x0a090a TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000134 Cmpl 0000000134 Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 00000000826ced5 Issue 00000000826c226 OutIO
fffffffffffff351
        abort 0000029d noxri 00000000 nondlp 000008df qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000821 Err 00012fcd
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the RHEL 8.8 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

   ```
   # cat /sys/class/fc_host/host*/symbolic_name
   ```

   Example output

   ```
   QLE2772 FW:v9.10.11 DVR:v10.02.07.900-k-debug
   QLE2772 FW:v9.10.11 DVR:v10.02.07.900-k-debug
   ```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

   ```
   # cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
   1
   ```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

   ```
   # cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
   256
   ```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

```
# nvme discover -t tcp -w 192.168.111.79 -a 192.168.111.14

Discovery Log Number of Records 8, Generation counter 10
=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.211.15
sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  1
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.111.15
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  2
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.211.14
sectype: none
..........
```

2. Verify that the other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w host-traddr -a traddr
```

Example output:

```
# nvme  discovery   -t   tcp    -w  192.168.111.79   -a 192.168.111.14
# nvme  discovery   -t   tcp    -w  192.168.111.79   -a 192.168.111.15
# nvme  discovery   -t   tcp    -w  192.168.211.79   -a 192.168.211.14
# nvme  discovery   -t   tcp    -w  192.168.211.79   -a 192.168.211.15
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
# nvme  connect-all -t  tcp -w  192.168.111.79  -a  192.168.111.14  -l
1800
# nvme  connect-all -t  tcp -w  192.168.111.79  -a  192.168.111.15  -l
1800
# nvme  connect-all -t  tcp -w  192.168.211.79  -a  192.168.211.14  -l
1800
# nvme  connect-all -t  tcp -w  192.168.211.79  -a  192.168.211.15  -l
1800
```

Validate NVMe-oF

You can use the following procedure to validate NVME-oF.

Steps

1. Verify that the in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme3n1 81Gx7NSiKSQeAAAAAAAB   NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 21.47 GB / 21.47 GB   4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme3n1
```

Example output:

```
nvme-subsys3 - NQN=nqn.1992-
08.com.netapp:sn.ab4fa6a5ba8b11ecbe3dd039ea359e4b:subsystem.rhel_161
_Lpe32002
\
 +- nvme0 fc traddr=nn-0x2048d039ea36a105:pn-0x204cd039ea36a105
host_traddr=nn-0x20000090fae0ec89:pn-0x10000090fae0ec89 live non-
optimized
 +- nvme1 fc traddr=nn-0x2048d039ea36a105:pn-0x204ad039ea36a105
host_traddr=nn-0x20000090fae0ec89:pn-0x10000090fae0ec89 live
optimized
 +- nvme2 fc traddr=nn-0x2048d039ea36a105:pn-0x204bd039ea36a105
host_traddr=nn-0x20000090fae0ec88:pn-0x10000090fae0ec88 live non-
optimized
 +- nvme4 fc traddr=nn-0x2048d039ea36a105:pn-0x2049d039ea36a105
host_traddr=nn-0x20000090fae0ec88:pn-0x10000090fae0ec88 live
optimized
```

NVMe/TCP

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp
_165
\
 +- nvme0 tcp traddr=192.168.111.15 trsvcid=4420
host_traddr=192.168.111.79 live non-optimized
 +- nvme1 tcp traddr=192.168.111.14 trsvcid=4420
host_traddr=192.168.111.79 live optimized
 +- nvme2 tcp traddr=192.168.211.15 trsvcid=4420
host_traddr=192.168.211.79 live non-optimized
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device        Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1 vs_tcp           /vol/vol1/ns1



NSID       UUID                                   Size
------------------------------------------------------------
1          338d73ce-b5a8-4847-9cc9-b127c75d8855 21.47GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
  "ONTAPdevices" : [
    {
      "Device" : "/dev/nvme0n1",
      "Vserver" : "vs_tcp79",
      "Namespace_Path" : "/vol/vol1/ns1",
      "NSID" : 1,
      "UUID" : "338d73ce-b5a8-4847-9cc9-b127c75d8855",
      "Size" : "21.47GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 5242880
    },
]

}
```

Known issues

The NVMe-oF host configuration for RHEL 8.8 with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1479047	RHEL 8.8 NVMe-oF hosts create duplicate persistent discovery controllers	On NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running Red Hat Enterprise Linux (RHEL) 8.8 on an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	2087000

NVMe-oF host configuration for RHEL 8.7 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and other transports) is supported with Red Hat Enterprise Linux (RHEL) 8.7 with ANA (Asymmetric Namespace Access). ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. During this procedure, you enable NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 8.7 and ONTAP as the target.

See the NetApp Interoperability Matrix Tool for accurate details regarding supported configurations.

Features

RHEL 8.7 includes support for NVMe/TCP (as a Technology Preview feature) in addition to NVMe/FC. The NetApp plugin in the native nvme-cli package is capable of displaying ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

Known limitations

- For RHEL 8.7, in-kernel NVMe multipath remains disabled by default. Therefore, you need to enable it manually.
- NVMe/TCP on RHEL 8.7 remains a Technology Preview feature due to open issues. Refer to the RHEL 8.7 release notes for details.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe Multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 8.7 on the server.
2. After the installation is complete, verify that you are running the specified RHEL 8.7 kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

Example:

```
# uname -r
4.18.0-425.3.1.el8.x86_64
```

3. Install the `nvme-cli` package:

   Example:

```
# rpm -qa|grep nvme-cli
nvme-cli-1.16-5.el8.x86_64
```

4. Enable in-kernel NVMe multipath:

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-425.3.1.el8.x86_64
```

5. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. Example:

```
# cat /etc/nvme/hostnqn

          nqn.2014-08.org.nvmexpress:uuid:a7f7a1d4-311a-11e8-b634-
7ed30aef10b7

::> vserver nvme subsystem host show -vserver vs_nvme167
Vserver     Subsystem       Host NQN
----------- --------------- ----------------
vs_nvme167 rhel_167_LPe35002  nqn.2014-08.org.nvmexpress:uuid: a7f7a1d4-
311a-11e8-b634-7ed30aef10b7
```

   If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

6. Reboot the host.

If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from dm-multipath to prevent dm-multipath from claiming these namespace devices. You can do this by adding the enable_foreign setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
        enable_foreign     NONE
}
```

Restart the multipathd daemon by running a `systemctl restart multipathd` command to allow the new setting to take effect.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe35002-M2
LPe35002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe35002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe35002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. See the NetApp Interoperability Matrix for the most current list of supported adapter driver and firmware versions.

```
# cat /sys/class/scsi_host/host*/fwrev
14.0.505.12, sli-4:6:d
14.0.505.12, sli-4:6:d
# cat /sys/module/lpfc/version
0:14.0.0.15
```

3. Verify that `lpfc_enable_fc4_type` is set to 3

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b95467c
0x100000109b95467b
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b95467c WWNN x200000109b95467c DID
x0a1500 ONLINE
NVME RPORT       WWPN x2071d039ea36a105 WWNN x206ed039ea36a105 DID
x0a0907 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2072d039ea36a105 WWNN x206ed039ea36a105 DID
x0a0805 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 00000001c7 Cmpl 00000001c7 Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000004909837 Issue 0000000004908cfc OutIO
fffffffffffff4c5
abort 0000004a noxri 00000000 nondlp 00000458 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00000061 Err 00017f43

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b95467b WWNN x200000109b95467b DID
x0a1100 ONLINE
NVME RPORT       WWPN x2070d039ea36a105 WWNN x206ed039ea36a105 DID
x0a1007 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x206fd039ea36a105 WWNN x206ed039ea36a105 DID
x0a0c05 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 00000001c7 Cmpl 00000001c7 Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000004909464 Issue 0000000004908531 OutIO
fffffffffffff0cd
abort 0000004f noxri 00000000 nondlp 00000361 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000006b Err 00017f99
```

Marvell/QLogic FC adapter for NVMe/FC

The native inbox `qla2xxx` driver included in the RHEL 8.7 kernel has the latest fixes which are essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions using the following command:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2772 FW:v9.08.02 DVR:v10.02.07.400-k-debug
QLE2772 FW:v9.08.02 DVR:v10.02.07.400-k-debug
```

2. Verify `ql2xnvmeenable` is set, which enables the Marvell adapter to function as a NVMe/FC initiator using the following command:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

ⓘ This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.211.5 -a 192.168.211.14

Discovery Log Number of Records 8, Generation counter 10

=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:
nqn.199208.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.211.15
sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  1
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.111.15
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
portid:  2
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.211.14
sectype: none
=====Discovery Log Entry 3======
trtype:  tcp
adrfam:  ipv4
subtype: unrecognized
treq:    not specified
```

```
portid:  3
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:discovery
traddr:  192.168.111.14
sectype: none
=====Discovery Log Entry 4======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp_165
traddr:  192.168.211.15
sectype: none
=====Discovery Log Entry 5======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  1
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp_165
traddr:  192.168.111.15
sectype: none
=====Discovery Log Entry 6======

trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  2
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp_165
traddr:  192.168.211.14
sectype: none

=====Discovery Log Entry 7======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
```

```
   portid:  3

trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp_165
traddr:  192.168.111.14
sectype: none
[root@R650-13-79 ~]#
```

2. Verify that other NVMe/TCP initiator-target LIF combos can successfully fetch discovery log page data. For example:

```
# nvme discover -t tcp -w 192.168.211.5 -a 192.168.211.14
# nvme discover -t tcp -w 192.168.211.5 -a 192.168.211.15
# nvme discover -t tcp -w 192.168.111.5 -a 192.168.111.14
# nvme discover -t tcp -w 192.168.111.5 -a 192.168.111.15
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it would retry for a longer period of time in the event of a path loss. For example:

```
# nvme connect-all -t tcp -w 192.168.211.5-a 192.168.211.14 -l 1800
# nvme connect-all -t tcp -w 192.168.211.5 -a 192.168.211.15 -l 1800
# nvme connect-all -t tcp -w 192.168.111.5 -a 192.168.111.14 -l 1800
# nvme connect-all -t tcp -w 192.168.111.5 -a 192.168.111.15 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is indeed enabled by checking:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example:

```
# nvme list
Node            SN                      Model                       Namespace
------------    --------------------- ---------------------------------
/dev/nvme0n1    81Gx7NSiKSRNAAAAAAAB    NetApp ONTAP Controller    1

Usage                   Format          FW Rev
-------------------   -----------     --------
21.47  GB /  21.47  GB  4 KiB + 0 B     FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example:

```
# nvme list-subsys /dev/nvme1n1

nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.154a5833c78c11ecb069d039ea359e4b:subsystem.rhel_tcp_165

\

 +- nvme0 tcp traddr=192.168.211.15 trsvcid=4420
host_traddr=192.168.211.5 live non-optimized

 +- nvme1 tcp traddr=192.168.211.14 trsvcid=4420
host_traddr=192.168.211.5 live optimized

 +- nvme2 tcp traddr=192.168.111.15 trsvcid=4420
host_traddr=192.168.111.5 live non-optimized

 +- nvme3 tcp traddr=192.168.111.14 trsvcid=4420
host_traddr=192.168.111.5 live optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device. For example:

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme0n1 vs_tcp79     /vol/vol1/ns1

NSID  UUID                                   Size
----  ------------------------------         ------
1     79c2c569-b7fa-42d5-b870-d9d6d7e5fa84  21.47GB


# nvme netapp ontapdevices -o json
{

  "ONTAPdevices" : [
  {

      "Device" : "/dev/nvme0n1",
      "Vserver" : "vs_tcp79",
      "Namespace_Path" : "/vol/vol1/ns1",
      "NSID" : 1,
      "UUID" : "79c2c569-b7fa-42d5-b870-d9d6d7e5fa84",
      "Size" : "21.47GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 5242880
    },

]

}
```

Known issues

The NVMe-oF host configuration for RHEL 8.7 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1479047	RHEL 8.7 NVMe-oF hosts create duplicate Persistent Discovery Controllers	On NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Red Hat Enterprise Linux (RHEL) 8.7 with an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	2087000

NVMe-oF host configuration for RHEL 8.6 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and other transports) is supported with Red Hat Enterprise Linux (RHEL) 8.6 with ANA (Asymmetric Namespace Access). ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. During this procedure, you enable NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 8.6 and ONTAP as the target

See the NetApp Interoperability Matrix Tool for accurate details regarding supported configurations.

Features

- RHEL 8.6 includes support for NVMe/TCP (as a Technology Preview feature) in addition to NVMe/FC. The NetApp plugin in the native nvme-cli package is capable of displaying ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

Known limitations

- For RHEL 8.6, in-kernel NVMe multipath remains disabled by default. Therefore, you need to enable it manually.
- NVMe/TCP on RHEL 8.6 remains a Technology Preview feature due to open issues. Refer to the RHEL 8.6 Release Notes for details.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe Multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 8.6 on the server. After the installation is complete, verify that you are running the specified RHEL 8.6 kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.
2. After the installation is complete, verify that you are running the specified RHEL 8.6 kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

Example:

```
# uname -r
4.18.0-372.9.1.el8.x86_64
```

3. Install the `nvme-cli` package:

   Example:

```
# rpm -qa|grep nvme-cli
nvme-cli-1.16-3.el8.x86_64
```

4. Enable in-kernel NVMe multipath:

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-372.9.1.el8.x86_64
```

5. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. Example:

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
Vserver     Subsystem       Host NQN
----------- ---------------
----------------------------------------------------------
vs_fcnvme_14 nvme_141_1     nqn.2014-08.org.nvmexpress:uuid:9ed5b327-
b9fc-4cf5-97b3-1b5d986345d1
```

   If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

6. Reboot the host.

If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the enable_foreign setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
        enable_foreign     NONE
}
```

Restart the multipathd daemon by running a `systemctl restart multipathd` command to allow the new setting to take effect.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. See the NetApp Interoperability Matrix for the most current list of supported adapter driver and firmware versions.

```
# cat /sys/class/scsi_host/host*/fwrev
12.8.351.47, sli-4:2:c
12.8.351.47, sli-4:2:c
# cat /sys/module/lpfc/version
0:14.0.0.4
```

3. Verify that `lpfc_enable_fc4_type` is set to 3

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC adapter for NVMe/FC

The native inbox `qla2xxx` driver included in the RHEL 8.6 kernel has the latest upstream fixes which are essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.06.02 DVR:v10.02.00.200-k
QLE2742 FW:v9.06.02 DVR:v10.02.00.200-k
```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as a NVMe/FC initiator using the following command:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Verify that other NVMe/TCP initiator-target LIF combos can successfully fetch discovery log page data. For example:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it would retry for a longer period of time in the event of a path loss. For example:

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example:

```
# nvme list
Node           SN                    Model                        Namespace
------------   --------------------- ----------------------------------
/dev/nvme0n1   814vWBNRwf9HAAAAAAAB   NetApp ONTAP Controller   1

Usage                Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example:

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_141_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device. For example:

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme0n1 vs_fcnvme_141    /vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns

NSID  UUID                                   Size
----  ------------------------------         ------
1     72b887b1-5fb6-47b8-be0b-33326e2542e2  85.90GB


# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_fcnvme_141",
        "Namespace_Path" : "/vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    }
  ]
}
```

Known issues

The NVMe-oF host configuration for RHEL 8.6 with ONTAP has the following known issues:

NetApp Bug ID	Title	Description	Bugzilla ID
1479047	RHEL 8.6 NVMe-oF hosts create duplicate Persistent Discovery Controllers	On NVMe over Fabrics (NVMe-oF) hosts, you can use the "nvme discover -p" command to create Persistent Discovery Controllers (PDCs). When this command is used, only one PDC should be created per initiator-target combination. However, if you are running ONTAP 9.10.1 and Red Hat Enterprise Linux (RHEL) 8.6 with an NVMe-oF host, a duplicate PDC is created each time "nvme discover -p" is executed. This leads to unnecessary usage of resources on both the host and the target.	2087000

NVMe-oF host configuration for RHEL 8.5 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and other transports) is supported with Red Hat Enterprise Linux (RHEL) 8.5 with ANA (Asymmetric Namespace Access). ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. During this procedure, you enable NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 8.5 and ONTAP as the target.

See the NetApp Interoperability Matrix Tool for accurate details regarding supported configurations.

Features

RHEL 8.5 includes support for NVMe/TCP (as a Technology Preview feature) in addition to NVMe/FC. The NetApp plugin in the native nvme-cli package can display ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

Known limitations

- For RHEL 8.5, in-kernel NVMe multipath remains disabled by default. Therefore, you need to enable it manually.
- NVMe/TCP on RHEL 8.5 remains a Technology Preview feature due to open issues. Refer to the RHEL 8.5 Release Notes for details.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe Multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 8.5 GA on the server. After the installation is complete, verify that you are running the specified RHEL 8.5 GA kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

   Example:

   ```
   # uname -r
   4.18.0-348.el8.x86_64
   ```

2. Install the `nvme-cli` package:

   Example:

   ```
   # rpm -qa|grep nvme-cli
   nvme-cli-1.14-3.el8.x86_64
   ```

3. Enable in-kernel NVMe multipath:

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-348.el8.x86_64
```

4. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. Example:

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
Vserver     Subsystem       Host NQN
----------- ---------------
----------------------------------------------------------
vs_fcnvme_14 nvme_141_1     nqn.2014-08.org.nvmexpress:uuid:9ed5b327-
b9fc-4cf5-97b3-1b5d986345d1
```

> If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

5. Reboot the host.

> If you intend to run both NVMe and SCSI co-existent traffic on the same host, NetApp recommends using in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from dm-multipath to prevent dm-multipath from claiming these namespace devices. You can do this by adding the enable_foreign setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
        enable_foreign     NONE
}
```

> Restart the multipathd daemon by running a `systemctl restart multipathd` command to allow the new setting to take effect.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. See the NetApp Interoperability Matrix for the most current list of supported adapter driver and firmware versions.

```
# cat /sys/class/scsi_host/host*/fwrev
12.8.351.47, sli-4:2:c
12.8.351.47, sli-4:2:c
# cat /sys/module/lpfc/version
0:12.8.0.10
```

3. Verify that `lpfc_enable_fc4_type` is set to 3

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and that you can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205

# cat /sys/class/fc_host/host*/port_state
Online
Online

# cat /sys/class/scsi_host/host*/nvme_info

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic

The native inbox `qla2xxx` driver included in the RHEL 8.5 GA kernel has the latest fixes which are essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as a NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

> This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Verify that other NVMe/TCP initiator-target LIF combos can successfully fetch discovery log page data. For example:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it retries for a longer period of time in the event of a path loss. For example:

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and `load balancing iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example:

```
# nvme list
Node           SN                    Model                        Namespace
------------   --------------------- ----------------------------------
/dev/nvme0n1    814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller   1

Usage                Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example:

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_141_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device. For example:

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme0n1 vs_fcnvme_141  vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns

NSID  UUID                                   Size
----  ------------------------------         ------
1     72b887b1-5fb6-47b8-be0b-33326e2542e2  85.90GB


# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_fcnvme_141",
        "Namespace_Path" : "/vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    }
  ]
}
```

Known issues

There are no known issues.

NVMe-oF Host Configuration for RHEL 8.4 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and other transports) is supported with Red Hat Enterprise Linux (RHEL) 8.4 with ANA (Asymmetric Namespace Access). ANA is the asymmetric logical unit access (ALUA) equivalent in the NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. You can enable NVMe-oF with in-kernel NVMe Multipath using ANA on RHEL 8.4 and ONTAP as the target.

Features

There are no new features in this release.

Known limitations

- For RHEL 8.4, in-kernel NVMe multipath is disabled by default. Therefore, you need to enable it manually.
- NVMe/TCP on RHEL 8.4 remains a Technology Preview feature due to open issues. Refer to the RHEL 8.4 Release Notes for details.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe multipath

You can use the following procedure to enable in-kernel NVMe multipath.

Steps

1. Install RHEL 8.4 GA on the server.
2. After the installation is complete, verify that you are running the specified RHEL 8.4 kernel. See the NetApp Interoperability Matrix for the most current list of supported versions.

   Example:

   ```
   # uname -r
   4.18.0-305.el8.x86_64
   ```

3. Install the `nvme-cli` package:

   Example:

   ```
   # rpm -qa|grep nvme-cli
   nvme-cli-1.12-3.el8.x86_64
   ```

4. Enable in-kernel NVMe multipath:

   ```
   # grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
   4.18.0-305.el8.x86_64
   ```

5. On the host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. Example:

   ```
   # cat /etc/nvme/hostnqn
   nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
   ::> vserver nvme subsystem host show -vserver vs_fcnvme_141
   Vserver     Subsystem       Host NQN
   ----------- ---------------
   ----------------------------------------------------------
   vs_fcnvme_14 nvme_141_1     nqn.2014-08.org.nvmexpress:uuid:9ed5b327-
   b9fc-4cf5-97b3-1b5d986345d1
   ```

If the host NQN strings do not match, you should use the `vserver modify` command to update the host NQN string on your corresponding ONTAP NVMe subsystem to match the host NQN string `/etc/nvme/hostnqn` on the host.

6. Reboot the host.

If you intend to run both NVMe & SCSI co-existent traffic on the same host, it is recommended to use in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. This means that the ONTAP namespaces should be excluded from dm-multipath to prevent dm-multipath from claiming these namespace devices. This can be done by adding the enable_foreign setting to the `/etc/multipath.conf` file:

```
# cat /etc/multipath.conf
defaults {
        enable_foreign     NONE
}
```

Restart the multipathd daemon by running a `systemctl restart multipathd` command to allow the new setting to take effect.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix Tool for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and inbox driver. See the NetApp Interoperability Matrix Tool for the most current list of supported adapter driver and firmware versions.

```
# cat /sys/class/scsi_host/host*/fwrev
12.8.340.8, sli-4:2:c
12.8.340.8, sli-4:2:c
# cat /sys/module/lpfc/version
0:12.8.0.5
```

3. Verify that `lpfc_enable_fc4_type` is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running, and you are able to see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

Marvell/QLogic FC adapter for NVMe/FC

The native inbox `qla2xxx` driver included in the RHEL 8.4 GA kernel has the latest fixes which are essential for ONTAP support.

Steps

1. Verify that you are running the supported adapter driver and firmware versions using the following command:

   ```
   # cat /sys/class/fc_host/host*/symbolic_name
   QLE2742 FW:v9.06.02 DVR:v10.02.00.104-k
   QLE2742 FW:v9.06.02 DVR:v10.02.00.104-k
   ```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as a NVMe/FC initiator using the following command:

   ```
   # cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
   1
   ```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

   ```
   # cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
   256
   ```

> ⓘ This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Verify that other NVMe/TCP initiator-target LIF combos are able to successfully fetch discovery log page data. For example,

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it would retry for a longer period of time in the event of a path loss. For example,

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and load balancing `iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example,

   Example (a):

```
# nvme list
Node           SN                    Model                        Namespace
------------   --------------------- ---------------------------------
/dev/nvme0n1   81CZ5BQuUNfGAAAAAAB  NetApp ONTAP Controller   1

Usage                Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

Example (b):

```
# nvme list
Node           SN                    Model                        Namespace
------------   --------------------- ---------------------------------
/dev/nvme0n1   81CYrBQuTHQFAAAAAAC  NetApp ONTAP Controller   1

Usage                Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example,

   Example (a):

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.04ba0732530911ea8e8300a098dfdd91:subsystem.nvme_145_1
\
+- nvme2 fc traddr=nn-0x208100a098dfdd91:pn-0x208200a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live non-
optimized
+- nvme3 fc traddr=nn-0x208100a098dfdd91:pn-0x208500a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live non-
optimized
+- nvme4 fc traddr=nn-0x208100a098dfdd91:pn-0x208400a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live optimized
+- nvme6 fc traddr=nn-0x208100a098dfdd91:pn-0x208300a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live optimized
```

   Example (b):

```
#nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.37ba7d9cbfba11eba35dd039ea165514:subsystem.nvme_114_tcp
_1
\
+- nvme0 tcp traddr=192.168.2.36 trsvcid=4420 host_traddr=192.168.1.4
live optimized
+- nvme1 tcp traddr=192.168.1.31 trsvcid=4420 host_traddr=192.168.1.4
live optimized
+- nvme10 tcp traddr=192.168.2.37 trsvcid=4420 host_traddr=192.168.1.4
live non-optimized
+- nvme11 tcp traddr=192.168.1.32 trsvcid=4420 host_traddr=192.168.1.4
live non-optimized
+- nvme20 tcp traddr=192.168.2.36 trsvcid=4420 host_traddr=192.168.2.5
live optimized
+- nvme21 tcp traddr=192.168.1.31 trsvcid=4420 host_traddr=192.168.2.5
live optimized
+- nvme30 tcp traddr=192.168.2.37 trsvcid=4420 host_traddr=192.168.2.5
live non-optimized
+- nvme31 tcp traddr=192.168.1.32 trsvcid=4420 host_traddr=192.168.2.5
live non-optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device. For example,

   Example (a):

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme1n1 vserver_fcnvme_145 /vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns

NSID  UUID                                   Size
-------  ------------------------------     --------------------
1      23766b68-e261-444e-b378-2e84dbe0e5e1  85.90GB


# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
     {
       "Device" : "/dev/nvme1n1",
       "Vserver" : "vserver_fcnvme_145",
       "Namespace_Path" : "/vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns",
       "NSID" : 1,
       "UUID" : "23766b68-e261-444e-b378-2e84dbe0e5e1",
       "Size" : "85.90GB",
       "LBA_Data_Size" : 4096,
       "Namespace_Size" : 20971520
     }
  ]
}
```

Example (b):

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme0n1 vs_tcp_114       /vol/tcpnvme_114_1_0_1/tcpnvme_114_ns

NSID  UUID                                   Size
-------  ------------------------------         ---------------
1      a6aee036-e12f-4b07-8e79-4d38a9165686  85.90GB


# nvme netapp ontapdevices -o json
{
     "ONTAPdevices" : [
     {
          "Device" : "/dev/nvme0n1",
           "Vserver" : "vs_tcp_114",
          "Namespace_Path" : "/vol/tcpnvme_114_1_0_1/tcpnvme_114_ns",
          "NSID" : 1,
          "UUID" : "a6aee036-e12f-4b07-8e79-4d38a9165686",
          "Size" : "85.90GB",
          "LBA_Data_Size" : 4096,
          "Namespace_Size" : 20971520
       }
  ]

}
```

Known issues

There are no known issues.

NVMe/FC host configuration for RHEL 8.3 with ONTAP

NVMe/FC is supported on ONTAP 9.6 or later for Red Hat Enterprise Linux (RHEL) 8.3. The RHEL 8.3 host runs both NVMe and SCSI traffic through the same FC initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

See the NetApp Interoperability Matrix Tool for the most current list of supported configurations.

Features

There are no new features in this release.

Known limitations

- For RHEL 8.3, in-kernel NVMe multipath is disabled by default. You can enable it manually.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC on RHEL 8.3

You can use the following procedure to enable NVMe/FC.

Steps

1. Install Red Hat Enterprise Linux 8.3 GA on the server.
2. If you are upgrading from RHEL 8.2 to RHEL 8.3 using the `yum update/upgrade` command, your `/etc/nvme/host*` files might be lost. To avoid file loss, use the following procedure:

   Steps

   a. Backup your `/etc/nvme/host*` files.
   b. If you have a manually edited `udev` rule, remove it:

   ```
   /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules
   ```

   c. Perform the upgrade.
   d. After the upgrade is complete, run the following command:

   ```
   yum remove nvme-cli
   ```

   e. Restore the host files at `/etc/nvme/`.

   ```
   yum install nvmecli
   ```

   f. Copy the original `/etc/nvme/host*` contents from the backup to the actual host files at `/etc/nvme/`.
3. After the installation is complete, verify that you're running the specified RHEL kernel:

   ```
   # uname -r
   4.18.0-240.el8.x86_64
   ```

   See the NetApp Interoperability Matrix Tool for the most current list of supported versions.
4. Install the nvme-cli package:

```
# rpm -qa|grep nvme-cli
nvme-cli-1.12-2.el8.x86_64
```

5. Enable in-kernel NVMe multipath.

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-240.el8.x86_64
```

6. On the RHEL 8.3 host, check the host NQN string at `/etc/nvme/hostnqn` verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array:

```
# cat /etc/nvme/hostnqn
```

Example output:

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
```

7. Verify that the `hostnqn` string matches the hostnqn string for the corresponding subsystem on the ONTAP array:

```
vserver nvme subsystem host show -vserver vs_fcnvme_141
```

Example output

```
::> vserver nvme subsystem host show -vserver vs_fcnvme_141
Vserver          Subsystem        Host           NQN
-----------      --------------- ----------- ---------------
vs_fcnvme_141     nvme_141_1                 nqn.2014-
08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1
```

If the host NQN strings do not match, use the `vserver modify` command to update the host NQN string on the corresponding ONTAP array subsystem to match to host NQN string from `/etc/nvme/hostnqn` on the host.

8. Reboot the host.
9. Optionally, update the `enable_foreign` setting.

If you intend to run both NVMe and SCSI traffic on the same RHEL 8.3 co-existent host, NetApp recommends that you use in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs, respectively. You should also blacklist the ONTAP namespaces in dm-multipath to prevent dm-multipath from claiming these namespace devices. You can do this by adding the `enable_foreign` setting to the /etc/multipath.conf, as shown below:

```
# cat /etc/multipath.conf
defaults {
   enable_foreign NONE
}
```

Restart the multipathd daemon by running a `systemctl restart multipathd`.

Validate NVMe/FC

You can use the following procedure to validate NVMe/FC.

Steps

1. Verify the following NVMe/FC settings:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created and properly discovered on the host.

```
/dev/nvme0n1     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
1                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n2     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
2                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n3     814vWBNRwf9HAAAAAAAB  NetApp ONTAP Controller
3                  85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.5f5f2c4aa73b11e9967e00a098df41bd:subsystem.nvme_141_1
\
+- nvme0 fc traddr=nn-0x203700a098dfdd91:pn-0x203800a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme1 fc traddr=nn-0x203700a098dfdd91:pn-0x203900a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x203700a098dfdd91:pn-0x203a00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x203700a098dfdd91:pn-0x203d00a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
```

4. Verify the NetApp plug-in for ONTAP devices:

Column

```
# nvme netapp ontapdevices -o column
```

Example output

```
Device                Vserver             Namespace Path
NSID                       UUID                         Size
--------------- ---------------
------------------------------------------------  --------
------------------------------------------  ---------
/dev/nvme0n1      vs_fcnvme_141
/vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns        1      72b887b1-5fb6-
47b8-be0b-33326e2542e2    85.90GB
/dev/nvme0n2      vs_fcnvme_141
/vol/fcnvme_141_vol_1_0_0/fcnvme_141_ns        2      04bf9f6e-9031-
40ea-99c7-a1a61b2d7d08    85.90GB
/dev/nvme0n3      vs_fcnvme_141
/vol/fcnvme_141_vol_1_1_1/fcnvme_141_ns        3      264823b1-8e03-
4155-80dd-e904237014a4    85.90GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{
"ONTAPdevices" : [
    {
        "Device" : "/dev/nvme0n1",
        "Vserver" : "vs_fcnvme_141",
        "Namespace_Path" :
"/vol/fcnvme_141_vol_1_1_0/fcnvme_141_ns",
        "NSID" : 1,
        "UUID" : "72b887b1-5fb6-47b8-be0b-33326e2542e2",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
    },
    {
        "Device" : "/dev/nvme0n2",
        "Vserver" : "vs_fcnvme_141",
        "Namespace_Path" :
"/vol/fcnvme_141_vol_1_0_0/fcnvme_141_ns",
        "NSID" : 2,
        "UUID" : "04bf9f6e-9031-40ea-99c7-a1a61b2d7d08",
        "Size" : "85.90GB",
        "LBA_Data_Size" : 4096,
        "Namespace_Size" : 20971520
      },
      {
         "Device" : "/dev/nvme0n3",
         "Vserver" : "vs_fcnvme_141",
         "Namespace_Path" :
"/vol/fcnvme_141_vol_1_1_1/fcnvme_141_ns",
         "NSID" : 3,
         "UUID" : "264823b1-8e03-4155-80dd-e904237014a4",
         "Size" : "85.90GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 20971520
       },
  ]
```

Configure the Broadcom FC adapter for NVMe/FC

You can use the following procedure to configure a Broadcom FC adapter.

For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.

Steps

1. Verify that you are using the supported adapter.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that `lpfc_enable_fc4_type` is set to "**3**".

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

3. Verify that the initiator ports are up and running and can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

4. Enable 1 MB I/O size *(optional)*.

   The `lpfc_sg_seg_cnt` parameter needs to be set to 256 for the lpfc driver to issue I/O requests up to 1 MB in size.

   ```
   # cat /etc/modprobe.d/lpfc.conf
   options lpfc lpfc_sg_seg_cnt=256
   ```

5. Run a `dracut -f` command and then reboot the host.
6. After the host boots up, verify that lpfc_sg_seg_cnt is set to 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

7. Verify that you are using the recommended Broadcom lpfc firmware as well as the inbox driver:

```
# cat /sys/class/scsi_host/host*/fwrev
12.8.340.8, sli-4:2:c
12.8.340.8, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:12.8.0.1
```

NVMe/FC host configuration for RHEL 8.2 with ONTAP

NVMe/FC is supported on ONTAP 9.6 or later for Red Hat Enterprise Linux (RHEL) 8.2. The RHEL 8.2 host runs both NVMe and SCSI traffic through the same fibre channel (FC) initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

See the NetApp Interoperability Matrix Tool for the most current list of supported configurations.

Features

- Beginning with RHEL 8.2, `nvme-fc auto-connect` scripts are included in the native `nvme-cli` package. You can rely on these native auto-connect scripts instead of having to install the external vendor provided outbox auto-connect scripts.
- Beginning with RHEL 8.2, a native `udev` rule is already provided as part of the `nvme-cli` package which enables round-robin load balancing for NVMe multipath. You need not manually create this rule any more (as was done in RHEL 8.1).
- Beginning with RHEL 8.2, both NVMe and SCSI traffic can be run on the same co-existent host. In fact, this is the expected deployed host configuration. Therefore, for SCSI, you can configure `dm-multipath` as usual for SCSI LUNs resulting in `mpath` devices, whereas NVMe multipath can be used to configure NVMe-oF multipath devices on the host.
- Beginning with RHEL 8.2, the NetApp plug-in in the native `nvme-cli` package is capable of displaying ONTAP details for ONTAP namespaces.

Known limitations

- For RHEL 8.2, in-kernel NVMe multipath is disabled by default. Therefore, you need to enable it manually.
- SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC

You can use the following procedure to enable NVMe/FC.

Steps

1. Install Red Hat Enterprise Linux 8.2 GA on the server.
2. If you are upgrading from RHEL 8.1 to RHEL 8.2 using `yum update/upgrade`, your `/etc/nvme/host*` files might be lost. To avoid the file loss, do the following:
   a. Backup your `/etc/nvme/host*` files.
   b. If you have a manually edited `udev` rule, remove it:

      ```
      /lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules
      ```

   c. Perform the upgrade.
   d. After the upgrade is complete, run the following command:

      ```
      yum remove nvme-cli
      ```

   e. Restore the host files at `/etc/nvme/`.

      ```
      yum install nvmecli
      ```

   f. Copy the original `/etc/nvme/host*` contents from the backup to the actual host files at `/etc/nvme/`.
3. After the installation is complete, verify that you are running the specified Red Hat Enterprise Linux kernel.

   ```
   # uname -r
   4.18.0-193.el8.x86_64
   ```

   See the NetApp Interoperability Matrix Tool for the most current list of supported versions.
4. Install the nvme-cli package.

   ```
   # rpm -qa|grep nvme-cli
   nvme-cli-1.9.5.el8.x86_64
   ```

5. Enable in-kernel NVMe multipath.

   ```
   # grubby -args=nvme_core.multipath=Y -update-kernel /boot/vmlinuz-
   4.18.0-193.el8.x86_64
   ```

6. On the RHEL 8.2 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array.

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-1b5d986345d1


::> vserver nvme subsystem host show -vserver vs_fcnvme_141
Vserver      Subsystem        Host           NQN
----------- --------------- ----------- ---------------
  vs_fcnvme_141
    nvme_141_1
        nqn.2014-08.org.nvmexpress:uuid:9ed5b327-b9fc-4cf5-97b3-
1b5d986345d1
```

If the host NQN strings do not match, use the `vserver modify` command to update the host NQN string on the corresponding ONTAP array subsystem to match to host NQN string from `/etc/nvme/hostnqn` on the host.

7. Reboot the host.
8. Update the `enable_foreign` setting *(optional)*.

   If you intend to run both NVMe and SCSI traffic on the same RHEL 8.2 co-existent host, NetApp recommends using in-kernel NVMe multipath for ONTAP namespaces and dm-multipath for ONTAP LUNs respectively. You should also blacklist the ONTAP namespaces in dm-multipath to prevent dm-multipath from claiming these namespace devices. You can do this by adding the `enable_foreign` setting to the `/etc/multipath.conf`, as shown below.

```
# cat /etc/multipath.conf
defaults {
   enable_foreign NONE
}
```

9. Restart the multipathd daemon by running a `systemctl restart multipathd`.

Configure the Broadcom FC adapter for NVMe/FC

You can use the following procedure to configure a Broadcom FC adapter.

For the most current list of supported adapters, see the NetApp Interoperability Matrix Tool.

Steps

1. Verify that you are using the supported adapter.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that `lpfc_enable_fc4_type` is set to "**3**".

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

3. Verify that the initiator ports are up and running and can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

4. Enable 1 MB I/O size *(optional)*.

   The `lpfc_sg_seg_cnt` parameter needs to be set to 256 for the lpfc driver to issue I/O requests up to 1 MB in size.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

5. Run a `dracut -f` command and then reboot the host.
6. After the host boots up, verify that lpfc_sg_seg_cnt is set to 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

7. Verify that you are using the recommended Broadcom lpfc firmware as well as the inbox driver.

```
# cat /sys/class/scsi_host/host*/fwrev
12.6.182.8, sli-4:2:c
12.6.182.8, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:12.6.0.2
```

8. Verify that `lpfc_enable_fc4_type` is set to "**3**".

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

9. Verify that the initiator ports are up and running and can see the target LIFs.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b1c1204
0x100000109b1c1205
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b1c1204 WWNN x200000109b1c1204 DID
x011d00 ONLINE
NVME RPORT WWPN x203800a098dfdd91 WWNN x203700a098dfdd91 DID x010c07
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203900a098dfdd91 WWNN x203700a098dfdd91 DID x011507
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000f78 Cmpl 0000000f78 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002fe29bba Issue 000000002fe29bc4 OutIO
000000000000000a
abort 00001bc7 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001e15 Err 0000d906
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b1c1205 WWNN x200000109b1c1205 DID
x011900 ONLINE
NVME RPORT WWPN x203d00a098dfdd91 WWNN x203700a098dfdd91 DID x010007
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203a00a098dfdd91 WWNN x203700a098dfdd91 DID x012a07
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000fa8 Cmpl 0000000fa8 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000002e14f170 Issue 000000002e14f17a OutIO
000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

10. Enable 1 MB I/O size *(optional)*.

    The `lpfc_sg_seg_cnt` parameter needs to be set to 256 for the lpfc driver to issue I/O requests up to 1 MB in size.

    ```
    # cat /etc/modprobe.d/lpfc.conf
    options lpfc lpfc_sg_seg_cnt=256
    ```

11. Run a `dracut -f` command and then reboot the host.
12. After the host boots up, verify that lpfc_sg_seg_cnt is set to 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

Validate NVMe/FC

You can use the following procedure to validate NVMe/FC.

Steps

1. Verify the following NVMe/FC settings.

   ```
   # cat /sys/module/nvme_core/parameters/multipath
   Y
   ```

   ```
   # cat /sys/class/nvme-subsystem/nvme-subsys*/model
   NetApp ONTAP Controller
   NetApp ONTAP Controller
   ```

   ```
   # cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
   round-robin
   round-robin
   ```

2. Verify that the namespaces are created.

   ```
   # nvme list
   Node SN Model Namespace Usage Format FW Rev
   ---------------- -------------------- -----------------------
   /dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
   53.69 GB 4 KiB + 0 B FFFFFFFF
   ```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 – NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.rhel_141_nvm
e_ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device    Vserver  Namespace Path             NSID   UUID   Size
-------   --------- -------------------------  ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/rhel_141_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/rhel_141_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

NVMe/FC host configuration for RHEL 8.1 with ONTAP

NVMe/FC is supported on ONTAP 9.6 or later for Red Hat Enterprise Linux (RHEL) 8.1. A RHEL 8.1 host can run both NVMe and SCSI traffic through the same FC initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

See the NetApp Interoperability Matrix Tool for the most current list of supported configurations.

Known limitations

* Native NVMe/FC auto-connect scripts are not available in the `nvme-cli` package. You can use the host bus adapter (HBA) vendor-provided external auto-connect script.
* NVMe multipath is disabled by default. Therefore, you need to enable it manually.
* By default, round-robin load balancing is not enabled. You can enable this functionality by writing a `udev` rule.
* SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC

You can use the following procedure to enable NVMe/FC.

Steps

1. Install Red Hat Enterprise Linux 8.1 on the server.
2. After the installation is complete, verify that you are running the specified RHEL kernel:

```
# uname -r
4.18.0-147.el8.x86_64
```

   See the NetApp Interoperability Matrix Tool for the most current list of supported versions.

3. Install the `nvme-cli-1.8.1-3.el8` package:

```
# rpm -qa|grep nvme-cli
nvme-cli-1.8.1-3.el8.x86_64
```

4. Enable in-kernel NVMe multipath:

```
# grubby –args=nvme_core.multipath=Y –update-kernel /boot/vmlinuz-
4.18.0-147.el8.x86_64
```

5. Add the following string as a separate udev rule at `/lib/udev/rules.d/71-nvme-iopolicy-netapp-ONTAP.rules`. This enables round-robin load balancing for NVMe multipath:

```
# Enable round-robin for NetApp ONTAP
ACTION==”add”, SUBSYSTEM==”nvme-subsystem”, ATTR{model}==”NetApp ONTAP
Controller”, ATTR{iopolicy}=”round-robin
```

6. On the RHEL 8.1 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array:

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

```
*> vserver nvme subsystem host show -vserver vs_nvme_10
Vserver Subsystem Host NQN
------- --------- -------------------------------------- -----------
rhel_141_nvme_ss_10_0
nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
```

> If the host NQN strings do not match, use the `vserver modify` command to update the host NQN string on your corresponding ONTAP array subsystem to match with the host NQN string from `/etc/nvme/hostnqn` on the host.

7. Reboot the host.

Configure the Broadcom FC adapter for NVMe/FC

You can use the following procedure to configure a Broadcom FC adapter.

Steps

1. Verify that you are using the supported adapter. See the NetApp Interoperability Matrix Tool for the most current list of supported adapters.

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Copy and install the Broadcom lpfc outbox driver and auto-connect scripts:

```
# tar -xvzf elx-lpfc-dd-rhel8-12.4.243.20-ds-1.tar.gz
# cd elx-lpfc-dd-rhel8-12.4.2453.20-ds-1
# ./elx_lpfc_install-sh -i -n
```

> The native drivers that are bundled with the OS are called the inbox drivers. If you download the outbox drivers (drivers that are not included with an OS release), an auto-connect script is included in the download and should be installed as part of the driver installation process.

3. Reboot the host.

4. Verify that you are using the recommended Broadcom lpfc firmware, outbox driver, and auto-connect package versions:

```
# cat /sys/class/scsi_host/host*/fwrev
12.4.243.20, sil-4.2.c
12.4.243.20, sil-4.2.c
```

```
# cat /sys/module/lpfc/version
0:12.4.243.20
```

```
# rpm -qa | grep nvmefc
nvmefc-connect-12.6.61.0-1.noarch
```

5. Verify that `lpfc_enable_fc4_type` is set to 3:

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

6. Verify that the initiator ports are up and running:

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

7. Verify that the NVMe/FC initiator ports are enabled, running, and you can see the target LIFs:

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2977 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
NVME Statistics
…
```

Enable 1MB I/O Size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Validate NVMe/FC

You can use the following procedure to validate NVMe/FC.

Steps

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- -------------------- -----------------------
/dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
53.69 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 – NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.rhel_141_nvm
e_ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device    Vserver   Namespace Path             NSID   UUID   Size
-------  --------- -------------------------  ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/rhel_141_vol_10_0/rhel_141_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/rhel_141_vol_10_0/rhel_141_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

SLES

SLES 15

NVMe-oF host configuration for SUSE Linux Enterprise Server 15 SP5 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with SUSE Linux Enterprise Server (SLES) 15 SP5 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FCP environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for SLES 15 SP5 with ONTAP:

- Both NVMe and SCSI traffic can be run on the same co-existent host. Therefore, for SCSI LUNs, you can configure dm-multipath for SCSI mpath devices, whereas you might use NVMe multipath to configure NVMe-oF namespace devices on the host.
- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native `nvme-cli` package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- Support for NVMe secure, in-band authentication

* Support for persistent discovery controllers (PDCs) using a unique discovery NQN

Known limitations

* SAN booting using the NVMe-oF protocol is currently not supported.
* There's no `sanlun` support for NVMe-oF. Therefore, the host utility support isn't available for NVMe-oF on an SLES 15 SP5 host. You can rely on the NetApp plug-in included in the native nvme-cli package for all NVMe-oF transports.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex FC or Marvell/Qlogic FC adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the recommended adapter model:

```
cat /sys/class/scsi_host/host*/modelname
```

Example output:

```
LPe32002 M2
LPe32002-M2
```

2. Verify the adapter model description:

```
cat /sys/class/scsi_host/host*/modeldesc
```

Example output:

```
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

3. Verify that you are using the recommended Emulex host bus adapter (HBA) firmware versions:

```
cat /sys/class/scsi_host/host*/fwrev
```

Example output:

```
14.0.639.20, sli-4:2:c
14.0.639.20, sli-4:2:c
```

4. Verify that you are using the recommended LPFC driver version:

```
cat /sys/module/lpfc/version
```

Example output:

```
0:14.2.0.13
```

5. Verify that you can view your initiator ports:

```
cat /sys/class/fc_host/host*/port_name
```

Example output:

```
0x100000109b579d5e
0x100000109b579d5f
```

6. Verify that your initiator ports are online:

```
cat /sys/class/fc_host/host*/port_state
```

Example output:

```
Online
Online
```

7. Verify that the NVMe/FC initiator ports are enabled and that the target ports are visible:

```
cat /sys/class/scsi_host/host*/nvme_info
```

Example output:

In this example, one initiator port is enabled and connected with two target LIFs.

```
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b579d5e WWNN x200000109b579d5e DID
x011c00 ONLINE
NVME RPORT WWPN x208400a098dfdd91 WWNN x208100a098dfdd91 DID x011503
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208500a098dfdd91 WWNN x208100a098dfdd91 DID x010003
TARGET DISCSRVC *ONLINE

NVME Statistics
LS: Xmt 0000000e49 Cmpl 0000000e49 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003ceb594f Issue 000000003ce65dbe OutIO
fffffffffffb046f
abort 00000bd2 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 000014f4 Err 00012abd

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b579d5f WWNN x200000109b579d5f DID
x011b00 ONLINE
NVME RPORT WWPN x208300a098dfdd91 WWNN x208100a098dfdd91 DID x010c03
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208200a098dfdd91 WWNN x208100a098dfdd91 DID x012a03
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000e50 Cmpl 0000000e50 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003c9859ca Issue 000000003c93515e OutIO
fffffffffffaf794
abort 00000b73 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000159d Err 000135c3
```

8. Reboot the host.

Marvell/QLogic

Steps

1. The native inbox qla2xxx driver included in the SLES 15 SP5 kernel has the latest fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
cat /sys/class/fc_host/host*/symbolic_name
```

Example output:

```
QLE2742 FW:v9.12.01 DVR: v10.02.08.300-k
QLE2742 FW:v9.12.01 DVR: v10.02.08.300-k
```

2. Verify that the `ql2xnvmeenable` parameter is set to 1:

```
cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Enable NVMe services

There are two NVMe/FC boot services included in the `nvme-cli` package, however, *only* `nvmefc-boot-connections.service` is enabled to start during system boot; `nvmf-autoconnect.service` is not enabled. Therefore, you need to manually enable `nvmf-autoconnect.service` to start during system boot.

Steps

1. Enable `nvmf-autoconnect.service`:

```
# systemctl enable nvmf-autoconnect.service
Created symlink /etc/systemd/system/default.target.wants/nvmf-
autoconnect.service → /usr/lib/systemd/system/nvmf-autoconnect.service.
```

2. Reboot the host.
3. Verify that `nvmf-autoconnect.service` and `nvmefc-boot-connections.service` are running after the system boots up:

   Example output:

```
# systemctl status nvmf-autoconnect.service
nvmf-autoconnect.service - Connect NVMe-oF subsystems automatically
during boot
Loaded: loaded (/usr/lib/systemd/system/nvmf-autoconnect.service;
enabled; vendor preset: disabled)
Active: inactive (dead) since Thu 2023-05-25 14:55:00 IST; 11min
ago
Process: 2108 ExecStartPre=/sbin/modprobe nvme-fabrics (code=exited,
status=0/SUCCESS)
Process: 2114 ExecStart=/usr/sbin/nvme connect-all (code=exited,
status=0/SUCCESS)
Main PID: 2114 (code=exited, status=0/SUCCESS)

systemd[1]: Starting Connect NVMe-oF subsystems automatically during
boot...
nvme[2114]: traddr=nn-0x201700a098fd4ca6:pn-0x201800a098fd4ca6 is
already connected
systemd[1]: nvmf-autoconnect.service: Deactivated successfully.
systemd[1]: Finished Connect NVMe-oF subsystems automatically during
boot.

# systemctl status nvmefc-boot-connections.service
nvmefc-boot-connections.service - Auto-connect to subsystems on FC-NVME
devices found during boot
Loaded: loaded (/usr/lib/systemd/system/nvmefc-boot-
connections.service; enabled; vendor preset: enabled)
Active: inactive (dead) since Thu 2023-05-25 14:55:00 IST; 11min ago
Main PID: 1647 (code=exited, status=0/SUCCESS)

systemd[1]: Starting Auto-connect to subsystems on FC-NVME devices found
during boot...
systemd[1]: nvmefc-boot-connections.service: Succeeded.
systemd[1]: Finished Auto-connect to subsystems on FC-NVME devices found
during boot.
```

Configure NVMe/TCP

You can use the following procedure to configure NVMe/TCP.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w <host-traddr> -a <traddr>
   ```

   Example output:

   ```
   # nvme discover -t tcp -w 192.168.1.4 -a 192.168.1.31

   Discovery Log Number of Records 8, Generation counter 18
   =====Discovery Log Entry 0====== trtype: tcp
   adrfam: ipv4
   subtype: current discovery subsystem treq: not specified
   portid: 0
   trsvcid: 8009 subnqn: nqn.1992-
   08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
   192.168.2.117
   eflags: explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 1====== trtype: tcp
   adrfam: ipv4
   subtype: current discovery subsystem treq: not specified
   portid: 1
   trsvcid: 8009 subnqn: nqn.1992-
   08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
   192.168.1.117
   eflags: explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 2====== trtype: tcp
   adrfam: ipv4
   subtype: current discovery subsystem treq: not specified
   portid: 2
   trsvcid: 8009 subnqn: nqn.1992-
   08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
   192.168.2.116
   eflags: explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 3====== trtype: tcp
   adrfam: ipv4
   subtype: current discovery subsystem treq: not specified
   portid: 3
   trsvcid: 8009 subnqn: nqn.1992-
   ```

```
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
192.168.1.116
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 4====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 0
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.2.117 eflags: not specified sectype: none
=====Discovery Log Entry 5====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 1
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.1.117 eflags: not specified sectype: none
=====Discovery Log Entry 6====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 2
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.2.116 eflags: not specified sectype: none
=====Discovery Log Entry 7====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 3
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.1.116 eflags: not specified sectype: none
```

2. Verify that all other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w <host-traddr> -a <traddr>
```

Example output:

```
# nvme discover -t tcp -w 192.168.1.4 -a 192.168.1.32
# nvme discover -t tcp -w 192.168.2.5 -a 192.168.2.36
# nvme discover -t tcp -w 192.168.2.5 -a 192.168.2.37
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l
<ctrl_loss_timeout_in_seconds>
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.1.4 -a 192.168.1.31 -l -1
# nvme connect-all -t tcp -w 192.168.1.4 -a 192.168.1.32 -l -1
# nvme connect-all -t tcp -w 192.168.2.5 -a 192.168.1.36 -l -1
# nvme connect-all -t tcp -w 192.168.2.5 -a 192.168.1.37 -l -1
```

NetApp recommends setting the `ctrl-loss-tmo` option to `-1` so that the NVMe/TCP initiator attempts to reconnect indefinitely in the event of a path loss.

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the host has the correct controller model for the ONTAP NVMe namespaces:

```
cat /sys/class/nvme-subsystem/nvme-subsys*/model
```

Example output:

```
NetApp ONTAP Controller
NetApp ONTAP Controller
```

3. Verify the NVMe I/O policy for the respective ONTAP NVMe I/O controller:

```
cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
```

Example output:

```
round-robin
round-robin
```

4. Verify that the ONTAP namespaces are visible to the host:

```
nvme list -v
```

Example output:

```
Subsystem        Subsystem-NQN
Controllers
----------------
------------------------------------------------------------------------
------------ -----------------------
nvme-subsys0     nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_dhcha
p   nvme0, nvme1, nvme2, nvme3


Device   SN                   MN
FR       TxPort Asdress        Subsystem    Namespaces
-------- -------------------- ----------------------------------------
-------- ---------------------------------------------
nvme0    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.2.214,trsvcid=4420,host_traddr=192.168.2.14 nvme-subsys0
nvme0n1
nvme1    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.2.215,trsvcid=4420,host_traddr=192.168.2.14 nvme-subsys0
nvme0n1
nvme2    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.1.214,trsvcid=4420,host_traddr=192.168.1.14 nvme-subsys0
nvme0n1
nvme3    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.1.215,trsvcid=4420,host_traddr=192.168.1.14 nvme-subsys0
nvme0n1


Device       Generic      NSID       Usage                 Format
Controllers
------------ ------------ ----------
-------------------------------------------------------------
/dev/nvme0n1 /dev/ng0n1   0x1     1.07  GB /   1.07  GB    4 KiB +  0 B
nvme0, nvme1, nvme2, nvme3
```

5. Verify that the controller state of each path is live and has the correct ANA status:

```
nvme list-subsys /dev/<subsystem_name>
```

NVMe/FC

Example output

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.04ba0732530911ea8e8300a098dfdd91:subsystem.nvme_145
_1
\
+- nvme2 fc traddr=nn-0x208100a098dfdd91:pn-
0x208200a098dfdd91,host_traddr=nn-0x200000109b579d5f:pn-
0x100000109b579d5f live optimized
+- nvme3 fc traddr=nn-0x208100a098dfdd91:pn-
0x208500a098dfdd91,host_traddr=nn-0x200000109b579d5e:pn-
0x100000109b579d5e live optimized
+- nvme4 fc traddr=nn-0x208100a098dfdd91:pn-
0x208400a098dfdd91,host_traddr=nn-0x200000109b579d5e:pn-
0x100000109b579d5e live non-optimized
+- nvme6 fc traddr=nn-0x208100a098dfdd91:pn-
0x208300a098dfdd91,host_traddr=nn-0x200000109b579d5f:pn-
0x100000109b579d5f live non-optimized
```

NVMe/TCP

Example output

```
# nvme list-subsys
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
hostnqn=nqn.2014-08.org.nvmexpress:uuid:e58eca24-faff-11ea-8fee-
3a68dd3b5c5f
iopolicy=round-robin

 +- nvme0 tcp
traddr=192.168.2.214,trsvcid=4420,host_traddr=192.168.2.14 live
 +- nvme1 tcp
traddr=192.168.2.215,trsvcid=4420,host_traddr=192.168.2.14 live
 +- nvme2 tcp
traddr=192.168.1.214,trsvcid=4420,host_traddr=192.168.1.14 live
 +- nvme3 tcp
traddr=192.168.1.215,trsvcid=4420,host_traddr=192.168.1.14 live
```

6. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

`nvme netapp ontapdevices -o column`

Example output:

```
Device           Vserver                   Namespace Path
NSID UUID                                   Size
---------------- -------------------------
---------------------------------------------------------------------
---------------------------
/dev/nvme0n1     vs_CLIENT114
/vol/CLIENT114_vol_0_10/CLIENT114_ns10       1    c6586535-da8a-
40fa-8c20-759ea0d69d33   1.07GB
```

JSON

`nvme netapp ontapdevices -o json`

Example output:

```
{
"ONTAPdevices":[
{
"Device":"/dev/nvme0n1",
"Vserver":"vs_CLIENT114",
"Namespace_Path":"/vol/CLIENT114_vol_0_10/CLIENT114_ns10",
"NSID":1,
"UUID":"c6586535-da8a-40fa-8c20-759ea0d69d33",
"Size":"1.07GB",
"LBA_Data_Size":4096,
"Namespace_Size":262144
}
]
}
```

Create a persistent discovery controller

Beginning with ONTAP 9.11.1, you can create a persistent discovery controller (PDC) for your SLES 15 SP5 host by using the following procedure. A PDC is required to automatically detect NVMe subsystem add or remove scenarios and changes to the discovery log page data.

Steps

1. Verify that the discovery log page data is available and can be retrieved through the initiator port and target LIF combination:

```
nvme discover -t <trtype> -w <host-traddr> -a <traddr>
```

Example output:

```
Discovery Log Number of Records 16, Generation counter 14
=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.1.214
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.1.215
eflags:  explicit discovery connections, duplicate discovery
information
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.2.215
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 3======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
```

```
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.2.214
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 4======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 5======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 6======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 7======
```

```
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.2.214
eflags:  none
sectype: none
=====Discovery Log Entry 8======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 9======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 10======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
```

```
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 11======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.2.214
eflags:  none
sectype: none
=====Discovery Log Entry 12======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 13======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 14======
trtype:  tcp
```

```
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 15======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.2.214
eflags:  none
sectype: none
```

2. Create a PDC for the discovery subsystem:

```
nvme discover -t <trtype> -w <host-traddr> -a <traddr> -p
```

Example output:

```
nvme discover -t tcp -w 192.168.1.16 -a 192.168.1.116 -p
```

3. From the ONTAP controller, verify that the PDC has been created:

```
vserver nvme show-discovery-controller -instance -vserver vserver_name
```

Example output:

```
vserver nvme show-discovery-controller -instance -vserver vs_nvme175
Vserver Name: vs_CLIENT116 Controller ID: 00C0h
Discovery Subsystem NQN: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery Logical
Interface UUID: d23cbb0a-c0a6-11ec-9731-d039ea165abc Logical Interface:
CLIENT116_lif_4a_1
Node: A400-14-124
Host NQN: nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-be09-
74362c0c1afc
Transport Protocol: nvme-tcp
Initiator Transport Address: 192.168.1.16
Host Identifier: 59de25be738348f08a79df4bce9573f3 Admin Queue Depth: 32
Header Digest Enabled: false Data Digest Enabled: false
Vserver UUID: 48391d66-c0a6-11ec-aaa5-d039ea165514
```

Set up secure in-band authentication

Beginning with ONTAP 9.12.1, secure in-band authentication is supported over NVMe/TCP and NVMe/FC between your SLES 15 SP5 host and your ONTAP controller.

To set up secure authentication, each host or controller must be associated with a `DH-HMAC-CHAP` key, which is a combination of the NQN of the NVMe host or controller and an authentication secret configured by the administrator. To authenticate its peer, an NVMe host or controller must recognize the key associated with the peer.

You can set up secure in-band authentication using the CLI or a config JSON file. If you need to specify different dhchap keys for different subsystems, you must use a config JSON file.

CLI

Steps

1. Obtain the host NQN:

```
cat /etc/nvme/hostnqn
```

2. Generate the dhchap key for the SLES15 SP5 host:

```
nvme gen-dhchap-key -s optional_secret -l key_length {32|48|64} -m
HMAC_function {0|1|2|3} -n host_nqn

•   -s secret key in hexadecimal characters to be used to initialize
the host key
•   -l length of the resulting key in bytes
•   -m HMAC function to use for key transformation
0 = none, 1- SHA-256, 2 = SHA-384, 3=SHA-512
•   -n host NQN to use for key transformation
```

In the following example, a random dhchap key with HMAC set to 3 (SHA-512) is generated.

```
# nvme gen-dhchap-key -m 3 -n nqn.2014-
08.org.nvmexpress:uuid:d3ca725a- ac8d-4d88-b46a-174ac235139b
DHHC-
1:03:J2UJQfj9f0pLnpF/ASDJRTyILKJRr5CougGpGdQSysPrLu6RW1fGl5VSjbeDF1n
1DEh3nVBe19nQ/LxreSBeH/bx/pU=:
```

3. On the ONTAP controller, add the host and specify both dhchap keys:

```
vserver nvme subsystem host add -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn> -dhchap-host-secret
<authentication_host_secret> -dhchap-controller-secret
<authentication_controller_secret> -dhchap-hash-function {sha-
256|sha-512} -dhchap-group {none|2048-bit|3072-bit|4096-bit|6144-
bit|8192-bit}
```

4. A host supports two types of authentication methods, unidirectional and bidirectional. On the host, connect to the ONTAP controller and specify dhchap keys based on the chosen authentication method:

```
nvme connect -t tcp -w <host-traddr> -a <tr-addr> -n <host_nqn> -S
<authentication_host_secret> -C <authentication_controller_secret>
```

5. Validate the `nvme connect authentication` command by verifying the host and controller dhchap keys:

   a. Verify the host dhchap keys:

   ```
   $cat /sys/class/nvme-subsystem/<nvme-subsysX>/nvme*/dhchap_secret
   ```

   Example output for unidirectional configuration:

   ```
   # cat /sys/class/nvme-subsystem/nvme-subsys1/nvme*/dhchap_secret
   DHHC-
   1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
   wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
   DHHC-
   1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
   wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
   DHHC-
   1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
   wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
   DHHC-
   1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
   wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
   ```

   b. Verify the controller dhchap keys:

   ```
   $cat /sys/class/nvme-subsystem/<nvme-
   subsysX>/nvme*/dhchap_ctrl_secret
   ```

   Example output for bidirectional configuration:

```
# cat /sys/class/nvme-subsystem/nvme-
subsys6/nvme*/dhchap_ctrl_secret
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
```

JSON file

You can use the `/etc/nvme/config.json` file with the `nvme connect-all` command when multiple NVMe subsystems are available on the ONTAP controller configuration.

You can generate the JSON file using `-o` option. Refer to the NVMe connect-all man pages for more syntax options.

Steps

1. Configure the JSON file:

```
# cat /etc/nvme/config.json
[
 {
    "hostnqn":"nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-
be09-74362c0c1afc",
    "hostid":"3ae10b42-21af-48ce-a40b-cfb5bad81839",
    "dhchap_key":"DHHC-
1:03:Cu3ZZfIz1WMlqZFnCMqpAgn/T6EVOcIFHez215U+Pow8jTgBF2UbNk3DK4wfk2E
ptWpna1rpwG5CndpOgxpRxh9m41w=:"
 },

 {
    "hostnqn":"nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-
be09-74362c0c1afc",
    "subsystems":[
        {
            "nqn":"nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_C
LIENT116",
            "ports":[
```

```
                {
                     "transport":"tcp",
                     "traddr":"192.168.1.117",
                     "host_traddr":"192.168.1.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.1.116",
                     "host_traddr":"192.168.1.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.2.117",
                     "host_traddr":"192.168.2.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.2.116",
                     "host_traddr":"192.168.2.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                }
            ]
        }
    ]
  }
]

[NOTE]
In the preceding example, `dhchap_key` corresponds to
`dhchap_secret` and `dhchap_ctrl_key` corresponds to
`dhchap_ctrl_secret`.
```

2. Connect to the ONTAP controller using the config JSON file:

```
nvme connect-all -J /etc/nvme/config.json
```

Example output:

```
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
```

3. Verify that the dhchap secrets have been enabled for the respective controllers for each subsystem:
   a. Verify the host dhchap keys:

   ```
   # cat /sys/class/nvme-subsystem/nvme-subsys0/nvme0/dhchap_secret
   ```

   Example output:

   ```
   DHHC-1:01:NunEWY7AZlXqxITGheByarwZdQvU4ebZg9HOjIr6nOHEkxJg:
   ```

   b. Verify the controller dhchap keys:

   ```
   # cat /sys/class/nvme-subsystem/nvme-
   subsys0/nvme0/dhchap_ctrl_secret
   ```

   Example output:

   ```
   DHHC-
   1:03:2YJinsxa2v3+m8qqCiTnmgBZoH6mIT6G/6f0aGO8viVZB4VLNLH4z8CvK7pV
   YxN6S5fOAtaU3DNi12rieRMfdbg3704=:
   ```

Known issues

There are no known issues for the SLES 15 SP5 with ONTAP release.

NVMe-oF host configuration for SUSE Linux Enterprise Server 15 SP4 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with SUSE Linux Enterprise Server (SLES) 15 SP4 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FCP environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for SLES 15 SP4 with ONTAP:

- Both NVMe and SCSI traffic can be run on the same co-existent host. Therefore, for SCSI LUNs, you can configure dm-multipath for SCSI mpath devices, whereas you might use NVMe multipath to configure NVMe-oF namespace devices on the host.
- Support for NVMe over TCP (NVMe/TCP) in addition to NVMe/FC. The NetApp plug-in in the native nvme-cli package displays ONTAP details for both NVMe/FC and NVMe/TCP namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

- Support for NVMe secure, in-band authentication
- Support for persistent discovery controllers (PDCs) using a unique discovery NQN

Known limitations

- SAN booting using the NVMe-oF protocol is currently not supported.
- There’s no sanlun support for NVMe-oF. Therefore, the host utility support isn’t available for NVMe-oF on an SLES15 SP5 host. You can rely on the NetApp plug-in included in the native nvme-cli package for all NVMe-oF transports.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex FC adapters or Marvell/Qlogic FC adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the recommended adapter model:

   ```
   cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe32002 M2
   LPe32002-M2
   ```

2. Verify the adapter model description:

   ```
   cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

3. Verify that you are using the recommended Emulex host bus adapter (HBA) firmware versions:

   ```
   cat /sys/class/scsi_host/host*/fwrev
   ```

   Example output:

   ```
   12.8.351.47, sli-4:2:c
   12.8.351.47, sli-4:2:c
   ```

4. Verify that you are using the recommended LPFC driver version:

   ```
   cat /sys/module/lpfc/version
   ```

   Example output:

   ```
   0:14.2.0.6
   ```

5. Verify that you can view your initiator ports:

```
cat /sys/class/fc_host/host*/port_name
```

Example output:

```
0x100000109b579d5e
0x100000109b579d5f
```

6. Verify that your initiator ports are online:

```
cat /sys/class/fc_host/host*/port_state
```

Example output:

```
Online
Online
```

7. Verify that the NVMe/FC initiator ports are enabled and that the target ports are visible:

```
cat /sys/class/scsi_host/host*/nvme_info
```

Example output:

In this example, one initiator port is enabled and connected with two target LIFs.

```
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b579d5e WWNN x200000109b579d5e DID
x011c00 ONLINE
NVME RPORT WWPN x208400a098dfdd91 WWNN x208100a098dfdd91 DID x011503
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208500a098dfdd91 WWNN x208100a098dfdd91 DID x010003
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000e49 Cmpl 0000000e49 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003ceb594f Issue 000000003ce65dbe OutIO
fffffffffffb046f
abort 00000bd2 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 000014f4 Err 00012abd

NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b579d5f WWNN x200000109b579d5f DID
x011b00 ONLINE
NVME RPORT WWPN x208300a098dfdd91 WWNN x208100a098dfdd91 DID x010c03
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208200a098dfdd91 WWNN x208100a098dfdd91 DID x012a03
TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 0000000e50 Cmpl 0000000e50 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003c9859ca Issue 000000003c93515e OutIO
ffffffffffaf794
abort 00000b73 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000159d Err 000135c3
```

8. Reboot the host.

Marvell/QLogic

Steps

1. The native inbox qla2xxx driver included in the SLES 15 SP4 kernel has the latest fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
cat /sys/class/fc_host/host*/symbolic_name
```

Example output:

```
QLE2742 FW:v9.08.02 DVR:v10.02.07.800-k QLE2742 FW:v9.08.02
DVR:v10.02.07.800-k
```

2. Verify that the `ql2xnvmeenable` parameter is set to 1:

```
cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

ⓘ This is not applicable to Qlogic NVMe/FC hosts.

Enable NVMe services

There are two NVMe/FC boot services included in the `nvme-cli` package, however, *only* `nvmefc-boot-connections.service` is enabled to start during system boot; `nvmf-autoconnect.service` is not enabled. Therefore, you need to manually enable `nvmf-autoconnect.service` to start during system boot.

Steps

1. Enable `nvmf-autoconnect.service`:

```
# systemctl enable nvmf-autoconnect.service
Created symlink /etc/systemd/system/default.target.wants/nvmf-
autoconnect.service → /usr/lib/systemd/system/nvmf-autoconnect.service.
```

2. Reboot the host.
3. Verify that `nvmf-autoconnect.service` and `nvmefc-boot-connections.service` are running after the system boots up:

   Example output:

```
# systemctl status nvmf-autoconnect.service
   nvmf-autoconnect.service - Connect NVMe-oF subsystems automatically
during boot
     Loaded: loaded (/usr/lib/systemd/system/nvmf-autoconnect.service;
enabled; vendor preset: disabled)
     Active: inactive (dead) since Thu 2023-05-25 14:55:00 IST; 11min
ago
    Process: 2108 ExecStartPre=/sbin/modprobe nvme-fabrics (code=exited,
status=0/SUCCESS)
    Process: 2114 ExecStart=/usr/sbin/nvme connect-all (code=exited,
status=0/SUCCESS)
   Main PID: 2114 (code=exited, status=0/SUCCESS)

   systemd[1]: Starting Connect NVMe-oF subsystems automatically during
boot...
   nvme[2114]: traddr=nn-0x201700a098fd4ca6:pn-0x201800a098fd4ca6 is
already connected
   systemd[1]: nvmf-autoconnect.service: Deactivated successfully.
   systemd[1]: Finished Connect NVMe-oF subsystems automatically during
boot.

# systemctl status nvmefc-boot-connections.service
nvmefc-boot-connections.service - Auto-connect to subsystems on FC-NVME
devices found during boot
   Loaded: loaded (/usr/lib/systemd/system/nvmefc-boot-
connections.service; enabled; vendor preset: enabled)
   Active: inactive (dead) since Thu 2023-05-25 14:55:00 IST; 11min ago
 Main PID: 1647 (code=exited, status=0/SUCCESS)

systemd[1]: Starting Auto-connect to subsystems on FC-NVME devices found
during boot...
systemd[1]: nvmefc-boot-connections.service: Succeeded.
systemd[1]: Finished Auto-connect to subsystems on FC-NVME devices found
during boot.
```

Configure NVMe/TCP

You can use the following procedure to configure NVMe/TCP.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
nvme discover -t tcp -w <host-traddr> -a <traddr>
```

Example output:

```
# nvme discover -t tcp -w 192.168.1.4 -a 192.168.1.31

Discovery Log Number of Records 8, Generation counter 18
=====Discovery Log Entry 0====== trtype: tcp
adrfam: ipv4
subtype: current discovery subsystem treq: not specified
portid: 0
trsvcid: 8009 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
192.168.2.117
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 1====== trtype: tcp
adrfam: ipv4
subtype: current discovery subsystem treq: not specified
portid: 1
trsvcid: 8009 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
192.168.1.117
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 2====== trtype: tcp
adrfam: ipv4
subtype: current discovery subsystem treq: not specified
portid: 2
trsvcid: 8009 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
192.168.2.116
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 3====== trtype: tcp
adrfam: ipv4
subtype: current discovery subsystem treq: not specified
portid: 3
trsvcid: 8009 subnqn: nqn.1992-
```

```
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery traddr:
192.168.1.116
eflags: explicit discovery connections, duplicate discovery information
sectype: none
=====Discovery Log Entry 4====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 0
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.2.117 eflags: not specified sectype: none
=====Discovery Log Entry 5====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 1
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.1.117 eflags: not specified sectype: none
=====Discovery Log Entry 6====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 2
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.2.116 eflags: not specified sectype: none
=====Discovery Log Entry 7====== trtype: tcp
adrfam: ipv4
subtype: nvme subsystem treq: not specified portid: 3
trsvcid: 4420 subnqn: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_CLIEN
T116
traddr: 192.168.1.116 eflags: not specified sectype: none
```

2. Verify that all other NVMe/TCP initiator-target LIF combinations can successfully fetch discovery log page data:

```
nvme discover -t tcp -w <host-traddr> -a <traddr>
```

Example output:

```
# nvme discover -t tcp -w 192.168.1.4 -a 192.168.1.32
# nvme discover -t tcp -w 192.168.2.5 -a 192.168.2.36
# nvme discover -t tcp -w 192.168.2.5 -a 192.168.2.37
```

3. Run the `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l
<ctrl_loss_timeout_in_seconds>
```

Example output:

```
# nvme connect-all -t tcp -w 192.168.1.4 -a 192.168.1.31 -l -1
# nvme connect-all -t tcp -w 192.168.1.4 -a 192.168.1.32 -l -1
# nvme connect-all -t tcp -w 192.168.2.5 -a 192.168.1.36 -l -1
# nvme connect-all -t tcp -w 192.168.2.5 -a 192.168.1.37 -l -1
```

> NetApp recommends setting the `ctrl-loss-tmo` option to `-1` so that the NVMe/TCP initiator attempts to reconnect indefinitely in the event of a path loss.

Validate NVMe-oF

You can use the following procedure to validate NVMe-oF.

Steps

1. Verify that in-kernel NVMe multipath is enabled:

```
cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the host has the correct controller model for the ONTAP NVMe namespaces:

```
cat /sys/class/nvme-subsystem/nvme-subsys*/model
```

Example output:

```
NetApp ONTAP Controller
NetApp ONTAP Controller
```

3. Verify the NVMe I/O policy for the respective ONTAP NVMe I/O controller:

```
cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
```

Example output:

```
round-robin
round-robin
```

4. Verify that the ONTAP namespaces are visible to the host:

```
nvme list -v
```

Example output:

```
Subsystem        Subsystem-NQN
Controllers
----------------
------------------------------------------------------------------------
------------ -----------------------
nvme-subsys0     nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_dhcha
p    nvme0, nvme1, nvme2, nvme3


Device   SN                   MN
FR       TxPort Asdress        Subsystem    Namespaces
-------- -------------------- ----------------------------------------
-------- ---------------------------------------------
nvme0    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.2.214,trsvcid=4420,host_traddr=192.168.2.14 nvme-subsys0
nvme0n1
nvme1    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.2.215,trsvcid=4420,host_traddr=192.168.2.14 nvme-subsys0
nvme0n1
nvme2    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.1.214,trsvcid=4420,host_traddr=192.168.1.14 nvme-subsys0
nvme0n1
nvme3    81LGgBUqsI3EAAAAAAAE NetApp ONTAP Controller   FFFFFFFF tcp
traddr=192.168.1.215,trsvcid=4420,host_traddr=192.168.1.14 nvme-subsys0
nvme0n1


Device       Generic      NSID       Usage                Format
Controllers
------------ ------------ ----------
-------------------------------------------------------------
/dev/nvme0n1 /dev/ng0n1   0x1     1.07  GB /   1.07  GB    4 KiB +  0 B
nvme0, nvme1, nvme2, nvme3
```

5. Verify that the controller state of each path is live and has the correct ANA status:

```
nvme list-subsys /dev/<subsystem_name>
```

NVMe/FC

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.04ba0732530911ea8e8300a098dfdd91:subsystem.nvme_145
_1
\
+- nvme2 fc traddr=nn-0x208100a098dfdd91:pn-
0x208200a098dfdd91,host_traddr=nn-0x200000109b579d5f:pn-
0x100000109b579d5f live optimized
+- nvme3 fc traddr=nn-0x208100a098dfdd91:pn-
0x208500a098dfdd91,host_traddr=nn-0x200000109b579d5e:pn-
0x100000109b579d5e live optimized
+- nvme4 fc traddr=nn-0x208100a098dfdd91:pn-
0x208400a098dfdd91,host_traddr=nn-0x200000109b579d5e:pn-
0x100000109b579d5e live non-optimized
+- nvme6 fc traddr=nn-0x208100a098dfdd91:pn-
0x208300a098dfdd91,host_traddr=nn-0x200000109b579d5f:pn-
0x100000109b579d5f live non-optimized
```

NVMe/TCP

```
# nvme list-subsys
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
hostnqn=nqn.2014-08.org.nvmexpress:uuid:e58eca24-faff-11ea-8fee-
3a68dd3b5c5f
iopolicy=round-robin

 +- nvme0 tcp
traddr=192.168.2.214,trsvcid=4420,host_traddr=192.168.2.14 live
 +- nvme1 tcp
traddr=192.168.2.215,trsvcid=4420,host_traddr=192.168.2.14 live
 +- nvme2 tcp
traddr=192.168.1.214,trsvcid=4420,host_traddr=192.168.1.14 live
 +- nvme3 tcp
traddr=192.168.1.215,trsvcid=4420,host_traddr=192.168.1.14 live
```

6. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

`nvme netapp ontapdevices -o column`

Example output:

```
Device           Vserver                   Namespace Path
NSID UUID                                   Size
---------------- ------------------------
--------------------------------------------------------------------
---------------------------
/dev/nvme0n1     vs_CLIENT114
/vol/CLIENT114_vol_0_10/CLIENT114_ns10       1    c6586535-da8a-
40fa-8c20-759ea0d69d33   1.07GB
```

JSON

`nvme netapp ontapdevices -o json`

Example output:

```
{
  "ONTAPdevices":[
    {
      "Device":"/dev/nvme0n1",
      "Vserver":"vs_CLIENT114",
      "Namespace_Path":"/vol/CLIENT114_vol_0_10/CLIENT114_ns10",
      "NSID":1,
      "UUID":"c6586535-da8a-40fa-8c20-759ea0d69d33",
      "Size":"1.07GB",
      "LBA_Data_Size":4096,
      "Namespace_Size":262144
    }
  ]
}
```

Create a persistent discovery controller

Beginning with ONTAP 9.11.1, you can create a persistent discovery controller (PDC) for your SLES 15 SP4 host by using the following procedure. A PDC is required to automatically detect NVMe subsystem add or remove scenarios and changes to the discovery log page data.

Steps

1. Verify that the discovery log page data is available and can be retrieved through the initiator port and target

LIF combination:

```
nvme discover -t <trtype> -w <host-traddr> -a <traddr>
```

Example output:

```
Discovery Log Number of Records 16, Generation counter 14
=====Discovery Log Entry 0======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.1.214
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 1======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.1.215
eflags:  explicit discovery connections, duplicate discovery
information
sectype: none
=====Discovery Log Entry 2======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.2.215
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 3======
trtype:  tcp
adrfam:  ipv4
subtype: current discovery subsystem
treq:    not specified
portid:  0
```

```
trsvcid: 8009
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:discovery
traddr:  192.168.2.214
eflags:  explicit discovery connections, duplicate discovery
information sectype: none
=====Discovery Log Entry 4======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 5======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 6======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 7======
```

```
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_n
one
traddr:  192.168.2.214
eflags:  none
sectype: none
=====Discovery Log Entry 8======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 9======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 10======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
```

```
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 11======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.subsys_C
LIENT114
traddr:  192.168.2.214
eflags:  none
sectype: none
=====Discovery Log Entry 12======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.1.214
eflags:  none
sectype: none
=====Discovery Log Entry 13======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.1.215
eflags:  none
sectype: none
=====Discovery Log Entry 14======
trtype:  tcp
```

```
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.2.215
eflags:  none
sectype: none
=====Discovery Log Entry 15======
trtype:  tcp
adrfam:  ipv4
subtype: nvme subsystem
treq:    not specified
portid:  0
trsvcid: 4420
subnqn:  nqn.1992-
08.com.netapp:sn.0501daf15dda11eeab68d039eaa7a232:subsystem.unidir_d
hchap
traddr:  192.168.2.214
eflags:  none
sectype: none
```

2. Create a PDC for the discovery subsystem:

```
nvme discover -t <trtype> -w <host-traddr> -a <traddr> -p
```

Example output:

```
nvme discover -t tcp -w 192.168.1.16 -a 192.168.1.116 -p
```

3. From the ONTAP controller, verify that the PDC has been created:

```
vserver nvme show-discovery-controller -instance -vserver vserver_name
```

Example output:

```
vserver nvme show-discovery-controller -instance -vserver vs_nvme175
Vserver Name: vs_CLIENT116 Controller ID: 00C0h
Discovery Subsystem NQN: nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:discovery Logical
Interface UUID: d23cbb0a-c0a6-11ec-9731-d039ea165abc Logical Interface:
CLIENT116_lif_4a_1
Node: A400-14-124
Host NQN: nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-be09-
74362c0c1afc
Transport Protocol: nvme-tcp
Initiator Transport Address: 192.168.1.16
Host Identifier: 59de25be738348f08a79df4bce9573f3 Admin Queue Depth: 32
Header Digest Enabled: false Data Digest Enabled: false
Vserver UUID: 48391d66-c0a6-11ec-aaa5-d039ea165514
```

Set up secure in-band authentication

Beginning with ONTAP 9.12.1, secure, in-band authentication is supported over NVMe/TCP and NVMe/FC between your SLES 15 SP4 host and your ONTAP controller.

To set up secure authentication, each host or controller must be associated with a `DH-HMAC-CHAP` key, which is a combination of the NQN of the NVMe host or controller and an authentication secret configured by the administrator. To authenticate its peer, an NVMe host or controller must recognize the key associated with the peer.

You can set up secure in-band authentication using the CLI or a config JSON file. If you need to specify different dhchap keys for different subsystems, you must use a config JSON file.

CLI

Steps

1. Obtain the host NQN:

```
cat /etc/nvme/hostnqn
```

2. Generate the dhchap key for the SLES15 SP4 host:

```
nvme gen-dhchap-key -s optional_secret -l key_length {32|48|64} -m
HMAC_function {0|1|2|3} -n host_nqn

•   -s secret key in hexadecimal characters to be used to initialize
the host key
•   -l length of the resulting key in bytes
•   -m HMAC function to use for key transformation
0 = none, 1- SHA-256, 2 = SHA-384, 3=SHA-512
•   -n host NQN to use for key transformation
```

+
In the following example, a random dhchap key with HMAC set to 3 (SHA-512) is generated.

```
# nvme gen-dhchap-key -m 3 -n nqn.2014-08.org.nvmexpress:uuid:d3ca725a-
ac8d-4d88-b46a-174ac235139b
DHHC-
1:03:J2UJQfj9f0pLnpF/ASDJRTyILKJRr5CougGpGdQSysPrLu6RW1fGl5VSjbeDF1n1DE
h3nVBe19nQ/LxreSBeH/bx/pU=:
```

1. On the ONTAP controller, add the host and specify both dhchap keys:

```
vserver nvme subsystem host add -vserver <svm_name> -subsystem
<subsystem> -host-nqn <host_nqn> -dhchap-host-secret
<authentication_host_secret> -dhchap-controller-secret
<authentication_controller_secret> -dhchap-hash-function {sha-
256|sha-512} -dhchap-group {none|2048-bit|3072-bit|4096-bit|6144-
bit|8192-bit}
```

2. A host supports two types of authentication methods, unidirectional and bidirectional. On the host, connect to the ONTAP controller and specify dhchap keys based on the chosen authentication method:

```
nvme connect -t tcp -w <host-traddr> -a <tr-addr> -n <host_nqn> -S
<authentication_host_secret> -C <authentication_controller_secret>
```

3. Validate the `nvme connect authentication` command by verifying the host and controller dhchap keys:
   a. Verify the host dhchap keys:

      ```
      $cat /sys/class/nvme-subsystem/<nvme-subsysX>/nvme*/dhchap_secret
      ```

      Example output for unidirectional configuration:

      ```
      SR650-14-114:~ # cat /sys/class/nvme-subsystem/nvme-
      subsys1/nvme*/dhchap_secret
      DHHC-
      1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
      wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
      DHHC-
      1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
      wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
      DHHC-
      1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
      wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
      DHHC-
      1:03:je1nQCmjJLUKD62mpYbzlpuw0OIws86NB96uNO/t3jbvhp7fjyR9bIRjOHg8
      wQtye1JCFSMkBQH3pTKGdYR1OV9gx00=:
      ```

   b. Verify the controller dhchap keys:

      ```
      $cat /sys/class/nvme-subsystem/<nvme-
      subsysX>/nvme*/dhchap_ctrl_secret
      ```

      Example output for bidirectional configuration:

```
SR650-14-114:~ # cat /sys/class/nvme-subsystem/nvme-
subsys6/nvme*/dhchap_ctrl_secret
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
DHHC-
1:03:WorVEV83eYO53kV4Iel5OpphbX5LAphO3F8fgH3913tlrkSGDBJTt3crXeTU
B8fCwGbPsEyz6CXxdQJi6kbn4IzmkFU=:
```

JSON file

You can use the `/etc/nvme/config.json` file with the `nvme connect-all` command when multiple NVMe subsystems are available on the ONTAP controller configuration.

You can generate the JSON file using `-o` option. Refer to the NVMe connect-all man pages for more syntax options.

Steps

1. Configure the JSON file:

```
# cat /etc/nvme/config.json
[
 {
    "hostnqn":"nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-
be09-74362c0c1afc",
    "hostid":"3ae10b42-21af-48ce-a40b-cfb5bad81839",
    "dhchap_key":"DHHC-
1:03:Cu3ZZfIz1WMlqZFnCMqpAgn/T6EVOcIFHez215U+Pow8jTgBF2UbNk3DK4wfk2E
ptWpna1rpwG5CndpOgxpRxh9m41w=:"
 },

 {
    "hostnqn":"nqn.2014-08.org.nvmexpress:uuid:12372496-59c4-4d1b-
be09-74362c0c1afc",
    "subsystems":[
        {
            "nqn":"nqn.1992-
08.com.netapp:sn.48391d66c0a611ecaaa5d039ea165514:subsystem.subsys_C
LIENT116",
            "ports":[
```

```
                {
                     "transport":"tcp",
                     "traddr":"192.168.1.117",
                     "host_traddr":"192.168.1.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.1.116",
                     "host_traddr":"192.168.1.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.2.117",
                     "host_traddr":"192.168.2.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                },
                {
                     "transport":"tcp",
                     "traddr":"192.168.2.116",
                     "host_traddr":"192.168.2.16",
                     "trsvcid":"4420",
                     "dhchap_ctrl_key":"DHHC-
1:01:0h58bcT/uu0rCpGsDYU6ZHZvRuVqsYKuBRS0Nu0VPx5HEwaZ:"
                }
            ]
        }
    ]
 }
]

[NOTE]
In the preceding example, `dhchap_key` corresponds to
`dhchap_secret` and `dhchap_ctrl_key` corresponds to
`dhchap_ctrl_secret`.
```

2. Connect to the ONTAP controller using the config JSON file:

```
nvme connect-all -J /etc/nvme/config.json
```

Example output:

```
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.116 is already connected
traddr=192.168.1.116 is already connected
traddr=192.168.2.117 is already connected
traddr=192.168.1.117 is already connected
```

3. Verify that the dhchap secrets have been enabled for the respective controllers for each subsystem:
   a. Verify the host dhchap keys:

   ```
   # cat /sys/class/nvme-subsystem/nvme-subsys0/nvme0/dhchap_secret
   ```

   Example output:

   ```
   DHHC-1:01:NunEWY7AZlXqxITGheByarwZdQvU4ebZg9HOjIr6nOHEkxJg:
   ```

   b. Verify the controller dhchap keys:

   ```
   # cat /sys/class/nvme-subsystem/nvme-
   subsys0/nvme0/dhchap_ctrl_secret
   ```

   Example output:

   ```
   DHHC-
   1:03:2YJinsxa2v3+m8qqCiTnmgBZoH6mIT6G/6f0aGO8viVZB4VLNLH4z8CvK7pV
   YxN6S5fOAtaU3DNi12rieRMfdbg3704=:
   ```

Known issues

There are no known issues for the SLES 15 SP4 with ONTAP release.

NVMe-oF Host Configuration for SUSE Linux Enterprise Server 15 SP3 with ONTAP

NVMe over Fabrics or NVMe-oF (including NVMe/FC and other transports) is supported with SUSE Linux Enterprise Server 15 SP3 (SLES15 SP3) with ANA (Asymmetric Namespace Access). ANA is the ALUA equivalent in NVMe-oF environment, and is currently implemented with in-kernel NVMe Multipath. The details for enabling NVMe-oF with in-kernel NVMe Multipath using ANA on SLES15 SP3 and ONTAP as the target has been documented here.

Refer to the NetApp Interoperability Matrix for accurate details regarding supported configurations.

Features

- SLES15 SP3 supports NVMe/FC and other transports.
- There is no sanlun support for NVMe-oF. Therefore, there is no LUHU support for NVMe-oF on SLES15 SP3. You can rely on the NetApp plug-in included in the native nvme-cli for the same instead. This should work for all NVMe-oF transports.
- Both NVMe and SCSI traffic can be run on the same co-existent host. In fact, that is expected to be the commonly deployed host config for customers. Therefore, for SCSI, you may configure `dm-multipath` as usual for SCSI LUNs resulting in mpath devices, whereas NVMe multipath might be used to configure NVMe-oF multipath devices on the host.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable in-kernel NVMe Multipath

In-kernel NVMe multipath is already enabled by default on SLES hosts such as SLES15 SP3. Therefore, no additional setting is required here. Refer to the NetApp Interoperability Matrix for accurate details regarding supported configurations.

NVMe-oF initiator packages

Refer to the NetApp Interoperability Matrix for accurate details regarding supported configurations.

1. Verify that you have the requisite kernel & nvme-cli MU packages installed on the SLES15 SP3 MU host.

   Example:

   ```
   # uname -r
   5.3.18-59.5-default

   # rpm -qa|grep nvme-cli
   nvme-cli-1.13-3.3.1.x86_64
   ```

   The above nvme-cli MU package now includes the following:

- **NVMe/FC auto-connect scripts** - Required for NVMe/FC auto-(re)connect when underlying paths to the namespaces are restored as well as during the host reboot:

```
# rpm -ql nvme-cli-1.13-3.3.1.x86_64
/etc/nvme
/etc/nvme/hostid
/etc/nvme/hostnqn
/usr/lib/systemd/system/nvmefc-boot-connections.service
/usr/lib/systemd/system/nvmefc-connect.target
/usr/lib/systemd/system/nvmefc-connect@.service
...
```

- **ONTAP udev rule** - New udev rule to ensure NVMe multipath round-robin loadbalancer default applies to all ONTAP namespaces:

```
# rpm -ql nvme-cli-1.13-3.3.1.x86_64
/etc/nvme
/etc/nvme/hostid
/etc/nvme/hostnqn
/usr/lib/systemd/system/nvmefc-boot-connections.service
/usr/lib/systemd/system/nvmf-autoconnect.service
/usr/lib/systemd/system/nvmf-connect.target
/usr/lib/systemd/system/nvmf-connect@.service
/usr/lib/udev/rules.d/70-nvmf-autoconnect.rules
/usr/lib/udev/rules.d/71-nvmf-iopolicy-netapp.rules
...
# cat /usr/lib/udev/rules.d/71-nvmf-iopolicy-netapp.rules
# Enable round-robin for NetApp ONTAP and NetApp E-Series
ACTION=="add", SUBSYSTEM=="nvme-subsystem", ATTR{model}=="NetApp
ONTAP Controller", ATTR{iopolicy}="round-robin"
ACTION=="add", SUBSYSTEM=="nvme-subsystem", ATTR{model}=="NetApp E-
Series", ATTR{iopolicy}="round-robin"
```

- **NetApp plug-in for ONTAP devices** - The existing NetApp plug-in has now been modified to handle ONTAP namespaces as well.

2. Check the hostnqn string at `/etc/nvme/hostnqn` on the host and ensure that it properly matches with the hostnqn string for the corresponding subsystem on the ONTAP array. For example,

```
# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:3ca559e1-5588-4fc4-b7d6-5ccfb0b9f054
::> vserver nvme subsystem host show -vserver vs_fcnvme_145
Vserver     Subsystem      Host NQN
-------     ---------      ----------------------------------
vs_nvme_145 nvme_145_1 nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-
41a6-a1fd-cf8262c8713f
            nvme_145_2 nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-
41a6-a1fd-cf8262c8713f
            nvme_145_3 nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-
41a6-a1fd-cf8262c8713f
            nvme_145_4 nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-
41a6-a1fd-cf8262c8713f
            nvme_145_5 nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-
41a6-a1fd-cf8262c8713f
5 entries were displayed.
```

Proceed with the below steps depending on the FC adapter being used on the host.

Configure NVMe/FC

Broadcom/Emulex

1. Verify that you have the recommended adapter and firmware versions. For example,

```
# cat /sys/class/scsi_host/host*/modelname
LPe32002-M2
LPe32002-M2
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
# cat /sys/class/scsi_host/host*/fwrev
12.8.340.8, sli-4:2:c
12.8.840.8, sli-4:2:c
```

   - The newer lpfc drivers (both inbox and outbox) already have lpfc_enable_fc4_type default set to 3, therefore, you no longer need to set this explicitly in the `/etc/modprobe.d/lpfc.conf`, and recreate the `initrd`. The `lpfc nvme` support is already enabled by default:

     ```
     # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
     3
     ```

   - The existing native inbox lpfc driver is already the latest and compatible with NVMe/FC. Therefore, you do not need to install the lpfc oob driver.

```
# cat /sys/module/lpfc/version
0:12.8.0.10
```

2. Verify that the initiator ports are up and running:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b579d5e
0x100000109b579d5f
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

3. Verify that the NVMe/FC initiator ports are enabled and you are able to see the target ports, and all are up and running. In this example, only 1 initiator port is enabled and connected with two target LIFs as seen in the output:

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b579d5e WWNN x200000109b579d5e DID
x011c00 ONLINE
NVME RPORT WWPN x208400a098dfdd91 WWNN x208100a098dfdd91 DID x011503
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208500a098dfdd91 WWNN x208100a098dfdd91 DID x010003
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000e49 Cmpl 0000000e49 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003ceb594f Issue 000000003ce65dbe OutIO
fffffffffffb046f
abort 00000bd2 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 000014f4 Err 00012abd
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b579d5f WWNN x200000109b579d5f DID
x011b00 ONLINE
NVME RPORT WWPN x208300a098dfdd91 WWNN x208100a098dfdd91 DID x010c03
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208200a098dfdd91 WWNN x208100a098dfdd91 DID x012a03
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000e50 Cmpl 0000000e50 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003c9859ca Issue 000000003c93515e OutIO
fffffffffffaf794
abort 00000b73 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000159d Err 000135c3
```

4. Reboot the host.

Enable 1MB I/O Size (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size should be up to 1 MB. However, to issue I/O requests of size 1 MB for the Broadcom NVMe/FC host, the lpfc parameter `lpfc_sg_seg_cnt` should also be bumped up to 256 from the default value of 64. Use the following instructions to do so:

1. Append the value 256 in the respective `modprobe lpfc.conf` file:

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. After reboot, verify that the above setting has been applied by checking the corresponding sysfs value:

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

Now the Broadcom NVMe/FC host should be able to send up 1MB I/O requests on the ONTAP namespace devices.

Marvell/QLogic

The native inbox qla2xxx driver included in the newer SLES15 SP3 MU kernel has the latest upstream fixes, essential for ONTAP support.

1. Verify that you are running the supported adapter driver and firmware versions, for example:

```
# cat /sys/class/fc_host/host*/symbolic_name
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
QLE2742 FW:v9.06.02 DVR:v10.02.00.106-k
```

2. Verify `ql2xnvmeenable` is set which enables the Marvell adapter to function as a NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Configure NVMe/TCP

Unlike NVMe/FC, NVMe/TCP has no auto-connect functionality. This manifests two major limitations on the Linux NVMe/TCP host:

- **No auto-reconnect after paths get reinstated** NVMe/TCP cannot automatically reconnect to a path that is reinstated beyond the default `ctrl-loss-tmo` timer of 10 minutes following a path down.
- **No auto-connect during host bootup** NVMe/TCP cannot automatically connect during host bootup as well.

You should set the retry period for failover events to at least 30 minutes to prevent timeouts. You can increase the retry period by increasing the value of the ctrl_loss_tmo timer. Following are the details:

Steps

1. Verify whether the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.51
Discovery Log Number of Records 10, Generation counter 119
=====Discovery Log Entry 0======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.2.56
sectype: none
=====Discovery Log Entry 1======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 1
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_1
traddr: 192.168.1.51
sectype: none
=====Discovery Log Entry 2======
trtype: tcp
adrfam: ipv4
subtype: nvme subsystem
treq: not specified
portid: 0
trsvcid: 4420
subnqn: nqn.1992-
08.com.netapp:sn.56e362e9bb4f11ebbaded039ea165abc:subsystem.nvme_118_tcp
_2
traddr: 192.168.2.56
sectype: none
...
```

2. Verify that other NVMe/TCP initiator-target LIF combos are able to successfully fetch discovery log page data. For example,

```
# nvme discover -t tcp -w 192.168.1.8 -a 192.168.1.52
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.56
# nvme discover -t tcp -w 192.168.2.9 -a 192.168.2.57
```

3. Run `nvme connect-all` command across all the supported NVMe/TCP initiator-target LIFs across the nodes. Ensure you set a longer `ctrl_loss_tmo` timer retry period (for example, 30 minutes, which can be set through `-l 1800`) during the connect-all so that it would retry for a longer period of time in the event of a path loss. For example,

```
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.51 -l 1800
# nvme connect-all -t tcp -w 192.168.1.8 -a 192.168.1.52 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.56 -l 1800
# nvme connect-all -t tcp -w 192.168.2.9 -a 192.168.2.57 -l 1800
```

Validate NVMe-oF

1. Verify that in-kernel NVMe multipath is indeed enabled by checking:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, `model` set to `NetApp ONTAP Controller` and `load balancing iopolicy` set to `round-robin`) for the respective ONTAP namespaces properly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller

# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the ONTAP namespaces properly reflect on the host. For example,

```
# nvme list
Node           SN                    Model                          Namespace
------------   --------------------- ----------------------------------
/dev/nvme0n1   81CZ5BQuUNfGAAAAAAAB  NetApp ONTAP Controller   1

Usage                 Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

Another example:

```
# nvme list
Node           SN                    Model                          Namespace
------------   --------------------- ----------------------------------
/dev/nvme0n1   81CYrBQuTHQFAAAAAAAC  NetApp ONTAP Controller   1

Usage                 Format         FW Rev
-------------------  -----------    --------
85.90 GB / 85.90 GB  4 KiB + 0 B    FFFFFFFF
```

4. Verify that the controller state of each path is live and has proper ANA status. For example,

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.04ba0732530911ea8e8300a098dfdd91:subsystem.nvme_145_1
\
+- nvme2 fc traddr=nn-0x208100a098dfdd91:pn-0x208200a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live non-
optimized
+- nvme3 fc traddr=nn-0x208100a098dfdd91:pn-0x208500a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live non-
optimized
+- nvme4 fc traddr=nn-0x208100a098dfdd91:pn-0x208400a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live optimized
+- nvme6 fc traddr=nn-0x208100a098dfdd91:pn-0x208300a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live optimized
```

Another example:

```
#nvme list-subsys /dev/nvme0n1
nvme-subsys0 - NQN=nqn.1992-
08.com.netapp:sn.37ba7d9cbfba11eba35dd039ea165514:subsystem.nvme_114_tcp
_1
\
+- nvme0 tcp traddr=192.168.2.36 trsvcid=4420 host_traddr=192.168.1.4
live optimized
+- nvme1 tcp traddr=192.168.1.31 trsvcid=4420 host_traddr=192.168.1.4
live optimized
+- nvme10 tcp traddr=192.168.2.37 trsvcid=4420 host_traddr=192.168.1.4
live non-optimized
+- nvme11 tcp traddr=192.168.1.32 trsvcid=4420 host_traddr=192.168.1.4
live non-optimized
+- nvme20 tcp traddr=192.168.2.36 trsvcid=4420 host_traddr=192.168.2.5
live optimized
+- nvme21 tcp traddr=192.168.1.31 trsvcid=4420 host_traddr=192.168.2.5
live optimized
+- nvme30 tcp traddr=192.168.2.37 trsvcid=4420 host_traddr=192.168.2.5
live non-optimized
+- nvme31 tcp traddr=192.168.1.32 trsvcid=4420 host_traddr=192.168.2.5
live non-optimized
```

5. Verify that the NetApp plug-in displays proper values for each ONTAP namespace device. For example,

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme1n1 vserver_fcnvme_145 /vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns

NSID  UUID                                   Size
----  ------------------------------         ------
1      23766b68-e261-444e-b378-2e84dbe0e5e1  85.90GB


# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
     {
       "Device" : "/dev/nvme1n1",
       "Vserver" : "vserver_fcnvme_145",
       "Namespace_Path" : "/vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns",
       "NSID" : 1,
       "UUID" : "23766b68-e261-444e-b378-2e84dbe0e5e1",
       "Size" : "85.90GB",
       "LBA_Data_Size" : 4096,
       "Namespace_Size" : 20971520
     }
  ]
}
```

Another example:

```
# nvme netapp ontapdevices -o column
Device       Vserver          Namespace Path
---------    -------
--------------------------------------------------
/dev/nvme0n1 vs_tcp_114       /vol/tcpnvme_114_1_0_1/tcpnvme_114_ns

NSID  UUID                                   Size
----  ------------------------------         ------
1      a6aee036-e12f-4b07-8e79-4d38a9165686  85.90GB


# nvme netapp ontapdevices -o json
{
     "ONTAPdevices" : [
     {
          "Device" : "/dev/nvme0n1",
           "Vserver" : "vs_tcp_114",
          "Namespace_Path" : "/vol/tcpnvme_114_1_0_1/tcpnvme_114_ns",
          "NSID" : 1,
          "UUID" : "a6aee036-e12f-4b07-8e79-4d38a9165686",
          "Size" : "85.90GB",
          "LBA_Data_Size" : 4096,
          "Namespace_Size" : 20971520
       }
  ]

}
```

== Known issues

There are no known issues.

NVMe/FC Host Configuration for SUSE Linux Enterprise Server 15 SP2 with ONTAP

NVMe/FC is supported on ONTAP 9.6 and above with SLES15 SP2. SLES15 SP2 host can run both NVMe/FC, & FCP traffic through the same fibre channel initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

> ⓘ You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC on SLES15 SP2

1. Upgrade to the recommended SLES15 SP2 MU kernel version.
2. Upgrade the native nvme-cli package.

   This native nvme-cli package contains the NVMe/FC auto-connect scripts, ONTAP udev rule which enables round-robin load balancing for NVMe Multipath as well as the NetApp plug-in for ONTAP namespaces.

   ```
   # rpm -qa|grep nvme-cli
   nvme-cli-1.10-2.38.x86_64
   ```

3. On the SLES15 SP2 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. For example:

   ```
   # cat /etc/nvme/hostnqn
   nqn.2014-08.org.nvmexpress:uuid:3ca559e1-5588-4fc4-b7d6-5ccfb0b9f054
   ```

   ```
   ::> vserver nvme subsystem host show -vserver vs_fcnvme_145
   Vserver Subsystem Host NQN
   ------- ---------
   ----------------------------------------------------------
   vs_fcnvme_145
   nvme_145_1
   nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-41a6-a1fd-cf8262c8713f
   nvme_145_2
   nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-41a6-a1fd-cf8262c8713f
   nvme_145_3
   nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-41a6-a1fd-cf8262c8713f
   nvme_145_4
   nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-41a6-a1fd-cf8262c8713f
   nvme_145_5
   nqn.2014-08.org.nvmexpress:uuid:c7b07b16-a22e-41a6-a1fd-cf8262c8713f
   5 entries were displayed.
   ```

4. Reboot the host.

Configure the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters see the NetApp Interoperability Matrix.

   ```
   # cat /sys/class/scsi_host/host*/modelname
   LPe32002-M2
   LPe32002-M2
   ```

```
# cat /sys/class/scsi_host/host*/modeldesc
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
```

2. Verify that you are using the recommended Broadcom lpfc firmware and native inbox driver versions.

```
# cat /sys/class/scsi_host/host*/fwrev
12.6.240.40, sli-4:2:c
12.6.240.40, sli-4:2:c
```

```
# cat /sys/module/lpfc/version
0:12.8.0.2
```

3. Verify that lpfc_enable_fc4_type is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x100000109b579d5e
0x100000109b579d5f
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

5. Verify that the NVMe/FC initiator ports are enabled, running and able to see the target LIFs.

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109b579d5e WWNN x200000109b579d5e DID
x011c00 ONLINE
NVME RPORT WWPN x208400a098dfdd91 WWNN x208100a098dfdd91 DID x011503
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208500a098dfdd91 WWNN x208100a098dfdd91 DID x010003
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000e49 Cmpl 0000000e49 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003ceb594f Issue 000000003ce65dbe OutIO
fffffffffffb046f
abort 00000bd2 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 000014f4 Err 00012abd
NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109b579d5f WWNN x200000109b579d5f DID
x011b00 ONLINE
NVME RPORT WWPN x208300a098dfdd91 WWNN x208100a098dfdd91 DID x010c03
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x208200a098dfdd91 WWNN x208100a098dfdd91 DID x012a03
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000e50 Cmpl 0000000e50 Abort 00000000
LS XMIT: Err 00000000 CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000003c9859ca Issue 000000003c93515e OutIO
fffffffffffaf794
abort 00000b73 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000159d Err 000135c3
```

Validate NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
```

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- --------------------
---------------------------------------- ---------
-------------------------- ---------------- --------
/dev/nvme1n1 814vWBNRwfBGAAAAAAAB NetApp ONTAP Controller 1 85.90 GB /
85.90 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys /dev/nvme1n1
nvme-subsys1 - NQN=nqn.1992-
08.com.netapp:sn.04ba0732530911ea8e8300a098dfdd91:subsystem.nvme_145_1
\
+- nvme2 fc traddr=nn-0x208100a098dfdd91:pn-0x208200a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live
inaccessible
+- nvme3 fc traddr=nn-0x208100a098dfdd91:pn-0x208500a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live
inaccessible
+- nvme4 fc traddr=nn-0x208100a098dfdd91:pn-0x208400a098dfdd91
host_traddr=nn-0x200000109b579d5e:pn-0x100000109b579d5e live optimized
+- nvme6 fc traddr=nn-0x208100a098dfdd91:pn-0x208300a098dfdd91
host_traddr=nn-0x200000109b579d5f:pn-0x100000109b579d5f live optimized
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device Vserver Namespace Path NSID UUID Size
---------------- -------------------------
-------------------------------------------------- ----
-------------------------------------- ---------
/dev/nvme1n1 vserver_fcnvme_145 /vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns
1 23766b68-e261-444e-b378-2e84dbe0e5e1 85.90GB

# nvme netapp ontapdevices -o json
{
"ONTAPdevices" : [
     {
       "Device" : "/dev/nvme1n1",
       "Vserver" : "vserver_fcnvme_145",
       "Namespace_Path" : "/vol/fcnvme_145_vol_1_0_0/fcnvme_145_ns",
       "NSID" : 1,
       "UUID" : "23766b68-e261-444e-b378-2e84dbe0e5e1",
       "Size" : "85.90GB",
       "LBA_Data_Size" : 4096,
       "Namespace_Size" : 20971520
     },
  ]
}
```

== Known issues
There are no known issues.

Enable 1MB I/O Size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

LPFC Verbose Logging

Set the lpfc driver for NVMe/FC.

Steps

1. Set the `lpfc_log_verbose` driver setting to any of the following values to log NVMe/FC events.

```
#define LOG_NVME 0x00100000 /* NVME general events. */
#define LOG_NVME_DISC 0x00200000 /* NVME Discovery/Connect events. */
#define LOG_NVME_ABTS 0x00400000 /* NVME ABTS events. */
#define LOG_NVME_IOERR 0x00800000 /* NVME IO Error events. */
```

2. After setting the values, run the `dracut-f` command and reboot the host.
3. Verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose 15728771
```

NVMe/FC Host Configuration for SUSE Linux Enterprise Server 15 SP1 with ONTAP

You can configure NVMe over Fibre Channel (NVMe/FC) on hosts running SUSE Linux Enterprise Server 15 SP1 and ONTAP as the target.

NVMe/FC is supported on ONTAP 9.6 or later for the following versions of SLES:

- SLES15 SP1

  SLES15 SP1 host can run both NVMe/FC, & FCP traffic through the same fibre channel initiator adapter ports. See the Hardware Universe for a list of supported FC adapters and controllers.

  For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

- Native NVMe/FC auto-connect scripts are included in the nvme-cli package. You can use the native inbox lpfc driver on SLES15 SP1.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Enable NVMe/FC on SLES15 SP1

1. Upgrade to the recommended SLES15 SP2 MU kernel
2. Upgrade to the recommended nvme-cli MU version.

   This nvme-cli package contains the native NVMe/FC auto-connect scripts, so you do not need to install the external NVMe/FC auto-connect scripts provided by Broadcom on the SLES15 SP1 host. This package also includes the ONTAP udev rule which enables round-robin load balancing for NVMe multipath, and the NetApp plug-in for ONTAP devices.

   ```
   # rpm -qa | grep nvme-cli
   nvme-cli-1.8.1-6.9.1.x86_64
   ```

3. On the SLES15 SP1 host, check the host NQN string at `/etc/nvme/hostnqn` and verify that it matches the host NQN string for the corresponding subsystem on the ONTAP array. For example:

   ```
   # cat /etc/nvme/hostnqn
   nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
   ```

   ```
   *> vserver nvme subsystem host show -vserver vs_nvme_10
   Vserver Subsystem Host NQN
   ------- --------- -------------------------------------- -----------
   sles_117_nvme_ss_10_0
   nqn.2014-08.org.nvmexpress:uuid:75953f3b-77fe-4e03-bf3c-09d5a156fbcd
   ```

4. Reboot the host.

Configure the Broadcom FC Adapter for NVMe/FC

1. Verify that you are using the supported adapter. For the most current list of supported adapters see the NetApp Interoperability Matrix.

   ```
   # cat /sys/class/scsi_host/host*/modelname
   LPe32002-M2
   LPe32002-M2
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom lpfc firmware and native inbox driver versions.

```
# cat /sys/class/scsi_host/host*/fwrev
12.4.243.17, sil-4.2.c
12.4.243.17, sil-4.2.c
```

```
# cat /sys/module/lpfc/version
0:12.6.0.0
```

3. Verify that lpfc_enable_fc4_type is set to 3.

```
# cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
3
```

4. Verify that the initiator ports are up and running.

```
# cat /sys/class/fc_host/host*/port_name
0x10000090fae0ec61
0x10000090fae0ec62
```

```
# cat /sys/class/fc_host/host*/port_state
Online
Online
```

5. Verify that the NVMe/FC initiator ports are enabled, running and able to see the target LIFs.

```
# cat /sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 NVME 2947 SCSI 2977 ELS 250
NVME LPORT lpfc0 WWPN x10000090fae0ec61 WWNN x20000090fae0ec61 DID
x012000 ONLINE
NVME RPORT WWPN x202d00a098c80f09 WWNN x202c00a098c80f09 DID x010201
TARGET DISCSRVC ONLINE
NVME RPORT WWPN x203100a098c80f09 WWNN x202c00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
NVME Statistics
…
```

Validate NVMe/FC

1. Verify the following NVMe/FC settings.

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

2. Verify that the namespaces are created.

```
# nvme list
Node SN Model Namespace Usage Format FW Rev
---------------- -------------------- -----------------------
/dev/nvme0n1 80BADBKnB/JvAAAAAAAC NetApp ONTAP Controller 1 53.69 GB /
53.69 GB 4 KiB + 0 B FFFFFFFF
```

3. Verify the status of the ANA paths.

```
# nvme list-subsys/dev/nvme0n1
Nvme-subsysf0 - NQN=nqn.1992-
08.com.netapp:sn.341541339b9511e8a9b500a098c80f09:subsystem.sles_117_nvm
e_ss_10_0
\
+- nvme0 fc traddr=nn-0x202c00a098c80f09:pn-0x202d00a098c80f09
host_traddr=nn-0x20000090fae0ec61:pn-0x10000090fae0ec61 live optimized
+- nvme1 fc traddr=nn-0x207300a098dfdd91:pn-0x207600a098dfdd91
host_traddr=nn-0x200000109b1c1204:pn-0x100000109b1c1204 live
inaccessible
+- nvme2 fc traddr=nn-0x207300a098dfdd91:pn-0x207500a098dfdd91
host_traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live optimized
+- nvme3 fc traddr=nn-0x207300a098dfdd91:pn-0x207700a098dfdd91 host
traddr=nn-0x200000109b1c1205:pn-0x100000109b1c1205 live inaccessible
```

4. Verify the NetApp plug-in for ONTAP devices.

```
# nvme netapp ontapdevices -o column
Device   Vserver  Namespace Path             NSID   UUID   Size
-------  -------- ------------------------  ------ ----- -----
/dev/nvme0n1   vs_nvme_10       /vol/sles_117_vol_10_0/sles_117_ns_10_0
1        55baf453-f629-4a18-9364-b6aee3f50dad   53.69GB

# nvme netapp ontapdevices -o json
{
   "ONTAPdevices" : [
   {
        Device" : "/dev/nvme0n1",
        "Vserver" : "vs_nvme_10",
        "Namespace_Path" : "/vol/sles_117_vol_10_0/sles_117_ns_10_0",
         "NSID" : 1,
         "UUID" : "55baf453-f629-4a18-9364-b6aee3f50dad",
         "Size" : "53.69GB",
         "LBA_Data_Size" : 4096,
         "Namespace_Size" : 13107200
    }
]
```

== Known issues

There are no known issues.

Enable 1MB I/O Size for Broadcom NVMe/FC

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

LPFC Verbose Logging

Set the lpfc driver for NVMe/FC.

Steps

1. Set the `lpfc_log_verbose` driver setting to any of the following values to log NVMe/FC events.

```
#define LOG_NVME 0x00100000 /* NVME general events. */
#define LOG_NVME_DISC 0x00200000 /* NVME Discovery/Connect events. */
#define LOG_NVME_ABTS 0x00400000 /* NVME ABTS events. */
#define LOG_NVME_IOERR 0x00800000 /* NVME IO Error events. */
```

2. After setting the values, run the `dracut-f` command and reboot the host.
3. Verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose 15728771
```

Ubuntu

NVMe-oF host configuration for Ubuntu 22.04 with ONTAP

NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Ubuntu 22.04 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

The following support is available for the NVMe-oF host configuration for Ubuntu 22.04 with ONTAP:

- The NetApp plug-in in the native nvme-cli package displays ONTAP details for NVMe/FC namespaces.
- Use of NVMe and SCSI co-existent traffic on the same host on a given host bus adapter (HBA), without the explicit dm-multipath settings to prevent claiming NVMe namespaces.

For additional details on supported configurations, see the NetApp Interoperability Matrix Tool.

Features

Ubuntu 22.04 has in-kernel NVMe multipath enabled for NVMe namespaces by default. Therefore, there is no need for explicit settings.

Known limitations

SAN booting using the NVMe-oF protocol is currently not supported.

Validate software versions

You can use the following procedure to validate the minimum supported Ubuntu 22.04 software versions.

Steps

1. Install Ubuntu 22.04 on the server. After the installation is complete, verify that you are running the specified Ubuntu 22.04 kernel:

   ```
   # uname -r
   ```

   Example output:

   ```
   5.15.0-101-generic
   ```

2. Install the `nvme-cli` package:

   ```
   # apt list | grep nvme
   ```

   Example output:

   ```
   nvme-cli/jammy-updates,now 1.16-3ubuntu0.1 amd64
   ```

3. On the Ubuntu 22.04 host, check the hostnqn string at `/etc/nvme/hostnqn`:

   ```
   # cat /etc/nvme/hostnqn
   ```

   Example output

   ```
   nqn.2014-08.org.nvmexpress:uuid:063a9fa0-438a-4737-b9b4-95a21c66d041
   ```

4. Verify that the `hostnqn` string matches the `hostnqn` string for the corresponding subsystem on the ONTAP array:

   ```
   ::> vserver nvme subsystem host show -vserver vs_106_fc_nvme
   ```

   Example output:

```
Vserver     Subsystem           Host NQN
----------- ---------------
----------------------------------------------------------
vs_106_fc_nvme ub_106 nqn.2014-08.org.nvmexpress:uuid:c04702c8-e91e-
4353-9995-ba4536214631
```

> If the `hostnqn` strings do not match, use the `vserver modify` command to update the `hostnqn` string on your corresponding ONTAP array subsystem to match the `hostnqn` string from `/etc/nvme/hostnqn` on the host.

Configure NVMe/FC

You can configure NVMe/FC for Broadcom/Emulex or Marvell/Qlogic adapters.

Broadcom/Emulex

Steps

1. Verify that you are using the supported adapter model.

   ```
   # cat /sys/class/scsi_host/host*/modelname
   ```

   Example output:

   ```
   LPe36002-M64
   LPe36002-M64
   ```

   ```
   # cat /sys/class/scsi_host/host*/modeldesc
   ```

   Example output:

   ```
   Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
   Emulex LPe36002-M64 2-Port 64Gb Fibre Channel Adapter
   ```

2. Verify that you are using the recommended Broadcom `lpfc` firmware and inbox driver.

   ```
   # cat /sys/class/scsi_host/host*/fwrev

   14.2.673.40, sli-4:6:d
   14.2.673.40, sli-4:6:d

   # cat /sys/module/lpfc/version
   0: 14.0.0.4
   ```

   For the most current list of supported adapter driver and firmware versions, see the NetApp Interoperability Matrix Tool.

3. Verify that `lpfc_enable_fc4_type` is set to `3`:

   ```
   # cat /sys/module/lpfc/parameters/lpfc_enable_fc4_type
   3
   ```

4. Verify that the initiator ports are up and running and that you can see the target LIFs:

```
# cat /sys/class/fc_host/host*/port_name
0x100000109bf0447c
0x100000109bf0447b
# cat /sys/class/fc_host/host*/port_state
Online
Online
# cat /sys/class/scsi_host/host*/nvme_info
        NVME Initiator Enabled
XRI Dist lpfc1 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc1 WWPN x100000109bf0447c WWNN x200000109bf0447c DID
x022300 ONLINE
NVME RPORT       WWPN x200cd039eaa8138b WWNN x200ad039eaa8138b DID
x021509 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x2010d039eaa8138b WWNN x200ad039eaa8138b DID
x021108 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 0000000000005238 Issue 000000000000523a OutIO
0000000000000002
        abort 00000000 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000000 Err 00000000

NVME Initiator Enabled
XRI Dist lpfc0 Total 6144 IO 5894 ELS 250
NVME LPORT lpfc0 WWPN x100000109bf0447b WWNN x200000109bf0447b DID
x022600 ONLINE
NVME RPORT       WWPN x200bd039eaa8138b WWNN x200ad039eaa8138b DID
x021409 TARGET DISCSRVC ONLINE
NVME RPORT       WWPN x200fd039eaa8138b WWNN x200ad039eaa8138b DID
x021008 TARGET DISCSRVC ONLINE

NVME Statistics
LS: Xmt 000000000e Cmpl 000000000e Abort 00000000
LS XMIT: Err 00000000  CMPL: xb 00000000 Err 00000000
Total FCP Cmpl 000000000000523c Issue 000000000000523e OutIO
0000000000000002
        abort 00000000 noxri 00000000 nondlp 00000000 qdepth
00000000 wqerr 00000000 err 00000000
FCP CMPL: xb 00000000 Err 00000000
```

Marvell/QLogic FC Adapter for NVMe/FC

Steps

1. The native inbox qla2xxx driver included in the Ubuntu 22.04 GA kernel has the latest upstream fixes essential for ONTAP support. Verify that you are running the supported adapter driver and firmware versions:

```
# cat /sys/class/fc_host/host*/symbolic_name
```

Example output

```
QLE2872 FW: v9.14.02 DVR: v10.02.06.200-k
QLE2872 FW: v9.14.02 DVR: v10.02.06.200-k
```

2. Verify that `ql2xnvmeenable` is set. This enables the Marvell adapter to function as an NVMe/FC initiator:

```
# cat /sys/module/qla2xxx/parameters/ql2xnvmeenable
1
```

Enable 1MB I/O (Optional)

ONTAP reports an MDTS (Max Data Transfer Size) of 8 in the Identify Controller data which means the maximum I/O request size can be up to 1MB. However, to issue I/O requests of size 1 MB for a Broadcom NVMe/FC host, you must increase the `lpfc` value of the `lpfc_sg_seg_cnt` parameter to 256 from the default value of 64.

Steps

1. Set the `lpfc_sg_seg_cnt` parameter to 256.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_sg_seg_cnt=256
```

2. Run a `dracut -f` command, and reboot the host.
3. Verify that `lpfc_sg_seg_cnt` is 256.

```
# cat /sys/module/lpfc/parameters/lpfc_sg_seg_cnt
256
```

This is not applicable to Qlogic NVMe/FC hosts.

Configure NVMe/TCP

NVMe/TCP does not have auto-connect functionality. Therefore, if a path goes down and is not reinstated within the default time out period of 10 minutes, NVMe/TCP cannot automatically reconnect. To prevent a time out, you should set the retry period for failover events to at least 30 minutes.

Steps

1. Verify that the initiator port can fetch the discovery log page data across the supported NVMe/TCP LIFs:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

   Example output:

   ```
   # nvme discover -t tcp -w 10.10.11.47-a 10.10.10.122

   Discovery Log Number of Records 8, Generation counter 10
   =====Discovery Log Entry 0======
   trtype:  tcp
   adrfam:  ipv4
   subtype: current discovery subsystem
   treq:    not specified
   portid:  0
   trsvcid: 8009
   subnqn:  nqn.1992-
   08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
   traddr:  10.10.10.122
   eflags:  explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 1======
   trtype:  tcp
   adrfam:  ipv4
   subtype: current discovery subsystem
   treq:    not specified
   portid:  1
   trsvcid: 8009
   subnqn:  nqn.1992
   08.com.netapp:sn.bbfb4ee8dfb611edbd07d039ea165590:discovery
   traddr:  10.10.10.124
   eflags:  explicit discovery connections, duplicate discovery information
   sectype: none
   =====Discovery Log Entry 2======
   trtype:  tcp
   ```

2. Verify that the other NVMe/TCP initiator-target LIF combinations are able to successfully fetch discovery log page data:

   ```
   nvme discover -t tcp -w host-traddr -a traddr
   ```

Example output:

```
#nvme discover -t tcp -w 10.10.10.47 -a 10.10.10.122
#nvme discover -t tcp -w 10.10.10.47 -a 10.10.10.124
#nvme discover -t tcp -w 10.10.11.47 -a 10.10.11.122
#nvme discover -t tcp -w 10.10.11.47 -a 10.10.11.
```

3. Run the nvme connect-all command across all the supported NVMe/TCP initiator-target LIFs across the nodes, and set the controller loss timeout period for at least 30 minutes or 1800 seconds:

```
nvme connect-all -t tcp -w host-traddr -a traddr -l 1800
```

Example output:

```
#   nvme    connect-all -t  tcp -w  10.10.10.47 -a  10.10.10.122 -l 1800
#   nvme    connect-all -t  tcp -w  10.10.10.47 -a  10.10.10.124 -l 1800
#   nvme    connect-all -t  tcp -w  10.10.11.47 -a  10.10.11.122 -l 1800
#   nvme    connect-all -t  tcp -w  10.10.11.47 -a  10.10.11.124 -l 1800
```

Validate NVMe-oF

You can use the following procedure to validate NVME-oF.

Steps

1. Verify that the in-kernel NVMe multipath is enabled:

```
# cat /sys/module/nvme_core/parameters/multipath
Y
```

2. Verify that the appropriate NVMe-oF settings (such as, model set to NetApp ONTAP Controller and load balancing iopolicy set to round-robin) for the respective ONTAP namespaces correctly reflect on the host:

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/model
NetApp ONTAP Controller
NetApp ONTAP Controller
```

```
# cat /sys/class/nvme-subsystem/nvme-subsys*/iopolicy
round-robin
round-robin
```

3. Verify that the namespaces are created and correctly discovered on the host:

```
# nvme list
```

Example output:

```
Node         SN                   Model
---------------------------------------------------------
/dev/nvme0n1 81CZ5BQuUNfGAAAAAAAB   NetApp ONTAP Controller


Namespace Usage    Format             FW             Rev
-----------------------------------------------------------
1                 21.47 GB / 21.47 GB   4 KiB + 0 B   FFFFFFFF
```

4. Verify that the controller state of each path is live and has the correct ANA status:

NVMe/FC

```
# nvme list-subsys /dev/nvme0n1
```

Example output:

```
nvme-subsys4 - NQN=nqn.1992-
08.com.netapp:sn.8763d311b2ac11ed950ed039ea951c46:subsystem. ub_106
\
+- nvme1 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a7d039ea954d17,host_traddr=nn-0x200000109b1b95ef:pn-
0x100000109b1b95ef live optimized
+- nvme2 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a8d039ea954d17,host_traddr=nn-0x200000109b1b95f0:pn-
0x100000109b1b95f0 live optimized
+- nvme3 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20aad039ea954d17,host_traddr=nn-0x200000109b1b95f0:pn-
0x100000109b1b95f0 live non-optimized
+- nvme5 fc traddr=nn-0x20a6d039ea954d17:pn-
0x20a9d039ea954d17,host_traddr=nn-0x200000109b1b95ef:pn-
0x100000109b1b95ef live non-optimized
```

NVME/TCP

```
# nvme list-subsys /dev/nvme1n1
```

Example output:

```
nvme-subsys1 - NQN=nqn.1992- 08.com.netapp:sn.
bbfb4ee8dfb611edbd07d039ea165590:subsystem.rhel_tcp_95
+- nvme1 tcp
traddr=10.10.10.122,trsvcid=4420,host_traddr=10.10.10.47,src_addr=10
.10.10.47 live
+- nvme2 tcp
traddr=10.10.10.124,trsvcid=4420,host_traddr=10.10.10.47,src_addr=10
.10.10.47 live
+- nvme3 tcp
traddr=10.10.11.122,trsvcid=4420,host_traddr=10.10.11.47,src_addr=10
.10.11.47 live
+- nvme4 tcp
traddr=10.10.11.124,trsvcid=4420,host_traddr=10.10.11.47,src_addr=10
.10.11.47 live
```

5. Verify that the NetApp plug-in displays the correct values for each ONTAP namespace device:

Column

```
# nvme netapp ontapdevices -o column
```

Example output:

```
Device        Vserver   Namespace Path
----------------------- ------------------------------
/dev/nvme0n1 co_iscsi_tcp_ubuntu /vol/vol1/ns1



NSID       UUID                                   Size
------------------------------------------------------------
1          79c2c569-b7fa-42d5-b870-d9d6d7e5fa84 21.47GB
```

JSON

```
# nvme netapp ontapdevices -o json
```

Example output

```
{

"ONTAPdevices" : [
{

"Device" : "/dev/nvme0n1",
"Vserver" : "co_iscsi_tcp_ubuntu",
"Namespace_Path" : "/vol/nvmevol1/ns1",
"NSID" : 1,
"UUID" : "79c2c569-b7fa-42d5-b870-d9d6d7e5fa84",
"Size" : "21.47GB",
"LBA_Data_Size" : 4096,
"Namespace_Size" : 5242880
},

]
}
```

Known issues

Theere are no known issues for the NVMe-oF host configuration for Ubuntu 22.04 with ONTAP release.

Windows

NVMe/FC Host Configuration for Windows Server 2022 with ONTAP

You can configure NVMe over Fibre Channel (NVMe/FC) on hosts running Windows Server 2022 using ONTAP as the target.

NVMe/FC is supported on ONTAP 9.7 or later for Windows Server 2022.

Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same 32G FC adapter ports. For FCP and FC/NVMe, use MSDSM as the Microsoft Multipath I/O (MPIO) option.

See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

Known limitations

Windows Failover Cluster (WFC) is not supported with ONTAP NVMe/FC because ONTAP does not currently support persistent reservations with NVMe/FC.

> The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI NVMe driver. This translational overhead does not necessarily impact performance, but it does negate the performance benefits of NVMe/FC. Thus, on Windows servers, NVMe/FC and FCP performance is the same, unlike on other operating systems such as Linux, where NVMe/FC performance is significantly better than that of FCP.

Enable NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

Steps

1. Install OneCommand Manager utility on the Windows host.
2. On each of the HBA initiator ports, set the following HBA driver parameters:
    - EnableNVMe = 1
    - NVMEMode = 0
    - LimTransferSize=1
3. Reboot the host.

Configure the Broadcom FC adapter in Windows for NVMe/FC

With the Broadcom adapter for FC/NVMe in a Windows environment, a `hostnqn` is associated with each host bus adapter (HBA) port. The `hostnqn` is formatted as follows.

Example

```
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9765
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9766
```

Enable MPIO for NVMe devices on the Windows host

1. Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
2. Open the MPIO properties.
3. From the **Discover Multi-Paths** tab, add the device ID listed for NVMe.

   MPIO becomes aware of the NVMe devices, which are visible under disk management.

4. Open **Disk Management** and go to **Disk Properties**.
5. From the **MPIO** tab, click **Details**.
6. Set the following MSDSM settings:
   - PathVerifiedPeriod: **10**
   - PathVerifyEnabled: **Enable**
   - RetryCount: **6**
   - RetryInterval: **1**
   - PDORemovedPeriod: **130**
7. Select the MPIO Policy **Round Robin with Subset**.
8. Change the registry values:

   ```
   HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInter
   val DWORD -> 30

   HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\
   UseCustomPathRecoveryInterval  DWORD-> 1
   ```

9. Reboot the host.

   The NVMe configuration is now complete on the Windows host.

Validate NVMe/FC

1. Validate that the Port Type is FC+NVMe.

   Now that NVMe is enabled, you should see the `Port Type` listed as `FC+NVMe`, as follows.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd listhba

Manageable HBA List

Port WWN        : 10:00:00:10:9b:1b:97:65
Node WWN        : 20:00:00:10:9b:1b:97:65
Fabric Name     : 10:00:c4:f5:7c:a5:32:e0
Flags           : 8000e300
Host Name       : INTEROP-57-159
Mfg             : Emulex Corporation
Serial No.      : FC71367217
Port Number     : 0
Mode            : Initiator
PCI Bus Number  : 94
PCI Function    : 0
Port Type       : FC+NVMe
Model           : LPe32002-M2

Port WWN        : 10:00:00:10:9b:1b:97:66
Node WWN        : 20:00:00:10:9b:1b:97:66
Fabric Name     : 10:00:c4:f5:7c:a5:32:e0
Flags           : 8000e300
Host Name       : INTEROP-57-159
Mfg             : Emulex Corporation
Serial No.      : FC71367217
Port Number     : 1
Mode            : Initiator
PCI Bus Number  : 94
PCI Function    : 1
Port Type       : FC+NVMe
Model           : LPe32002-M2
```

2. Validate that NVMe/FC subsystems have been discovered.

   The `nvme-list` command lists the NVMe/FC discovered subsystems.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:65

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:65

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:09:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0180
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:06:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0181
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available
Note: At present Namespace Management is not supported by NetApp Arrays.
```

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:66

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:66

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:07:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0140
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:08:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0141
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.
```

3. Validate that namespaces have been created.

   The `nvme-list-ns` command lists the namespaces for a specified NVMe target that lists the namespaces connected to the host.

```
PS C:\Program Files\Emulex\Util\OCManager> .\HbaCmd.exe nvme-list-ns
10:00:00:10:9b:1b:97:66 20:08:d0:39:ea:14:11:04 nq
.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159 0


Active Namespaces (attached to controller 0x0141):

                                       SCSI           SCSI
SCSI
   NSID           DeviceName        Bus Number    Target Number     OS
LUN
-----------  --------------------  ------------  ----------------
---------
0x00000001   \\.\PHYSICALDRIVE9         0              1            0
0x00000002   \\.\PHYSICALDRIVE10        0              1            1
0x00000003   \\.\PHYSICALDRIVE11        0              1            2
0x00000004   \\.\PHYSICALDRIVE12        0              1            3
0x00000005   \\.\PHYSICALDRIVE13        0              1            4
0x00000006   \\.\PHYSICALDRIVE14        0              1            5
0x00000007   \\.\PHYSICALDRIVE15        0              1            6
0x00000008   \\.\PHYSICALDRIVE16        0              1            7
```

NVMe/FC Host Configuration for Windows Server 2019 with ONTAP

You can configure NVMe over Fibre Channel (NVMe/FC) on hosts running Windows Server 2019 using ONTAP as the target.

NVMe/FC is supported on ONTAP 9.7 or later for Windows Server 2019.

Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same 32G FC adapter ports. For FCP and FC/NVMe, use MSDSM as the Microsoft Multipath I/O (MPIO) option.

See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

Windows Failover Cluster (WFC) is not supported with ONTAP NVMe/FC because ONTAP does not currently support persistent reservations with NVMe/FC.

> The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI NVMe driver. This translational overhead does not necessarily impact performance, but it does negate the performance benefits of NVMe/FC. Thus, on Windows servers, NVMe/FC and FCP performance is the same, unlike on other operating systems such as Linux, where NVMe/FC performance is significantly better than that of FCP.

Enable NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

Steps

1. Install OneCommand Manager utility on the Windows host.
2. On each of the HBA initiator ports, set the following HBA driver parameters:
   - EnableNVMe = 1
   - NVMEMode = 0
   - LimTransferSize=1
3. Reboot the host.

Configure the Broadcom FC adapter in Windows for NVMe/FC

With the Broadcom adapter for FC/NVMe in a Windows environment, a `hostnqn` is associated with each host bus adapter (HBA) port. The `hostnqn` is formatted as follows.

Example

```
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9765
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9766
```

Enable MPIO for NVMe devices on the Windows host

1. Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
2. Open the MPIO properties.
3. From the **Discover Multi-Paths** tab, add the device ID listed for NVMe.

   MPIO becomes aware of the NVMe devices, which are visible under disk management.
4. Open **Disk Management** and go to **Disk Properties**.
5. From the **MPIO** tab, click **Details**.
6. Set the following MSDSM settings:
   - PathVerifiedPeriod: **10**
   - PathVerifyEnabled: **Enable**
   - RetryCount: **6**
   - RetryInterval: **1**
   - PDORemovedPeriod: **130**
7. Select the MPIO Policy **Round Robin with Subset**.

8. Change the registry values:

```
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInter
val DWORD -> 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\
UseCustomPathRecoveryInterval  DWORD-> 1
```

9. Reboot the host.

   The NVMe configuration is now complete on the Windows host.

Validate NVMe/FC

1. Validate that the Port Type is FC+NVMe.

   Now that NVMe is enabled, you should see the `Port Type` listed as `FC+NVMe`, as follows.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd listhba

Manageable HBA List

Port WWN       : 10:00:00:10:9b:1b:97:65
Node WWN       : 20:00:00:10:9b:1b:97:65
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 0
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 0
Port Type      : FC+NVMe
Model          : LPe32002-M2

Port WWN       : 10:00:00:10:9b:1b:97:66
Node WWN       : 20:00:00:10:9b:1b:97:66
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 1
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 1
Port Type      : FC+NVMe
Model          : LPe32002-M2
```

2. Validate that NVMe/FC subsystems have been discovered.

   The `nvme-list` command lists the NVMe/FC discovered subsystems.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:65

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:65

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:09:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0180
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:06:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0181
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available
Note: At present Namespace Management is not supported by NetApp Arrays.
```

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:66

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:66

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:07:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0140
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:08:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0141
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.
```

3. Validate that namespaces have been created.

   The `nvme-list-ns` command lists the namespaces for a specified NVMe target that lists the namespaces connected to the host.

```
PS C:\Program Files\Emulex\Util\OCManager> .\HbaCmd.exe nvme-list-ns
10:00:00:10:9b:1b:97:66 20:08:d0:39:ea:14:11:04 nq
.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159 0


Active Namespaces (attached to controller 0x0141):

                                       SCSI           SCSI
SCSI
   NSID           DeviceName        Bus Number    Target Number     OS
LUN
-----------  --------------------  ------------  ----------------
---------
0x00000001   \\.\PHYSICALDRIVE9         0               1             0
0x00000002   \\.\PHYSICALDRIVE10        0               1             1
0x00000003   \\.\PHYSICALDRIVE11        0               1             2
0x00000004   \\.\PHYSICALDRIVE12        0               1             3
0x00000005   \\.\PHYSICALDRIVE13        0               1             4
0x00000006   \\.\PHYSICALDRIVE14        0               1             5
0x00000007   \\.\PHYSICALDRIVE15        0               1             6
0x00000008   \\.\PHYSICALDRIVE16        0               1             7
```

NVMe/FC Host Configuration for Windows Server 2016 with ONTAP

You can configure NVMe over Fibre Channel (NVMe/FC) on hosts running Windows Server 2016 using ONTAP as the target.

NVMe/FC is supported on ONTAP 9.7 or later for Windows Server 2016.

Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same 32G FC adapter ports. For FCP and FC/NVMe, use MSDSM as the Microsoft Multipath I/O (MPIO) option.

See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

Windows Failover Cluster (WFC) is not supported with ONTAP NVMe/FC because ONTAP does not currently support persistent reservations with NVMe/FC.

> The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI NVMe driver. This translational overhead does not necessarily impact performance, but it does negate the performance benefits of NVMe/FC. Thus, on Windows servers, NVMe/FC and FCP performance is the same, unlike on other operating systems such as Linux, where NVMe/FC performance is significantly better than that of FCP.

Enable NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

Steps

1. Install OneCommand Manager utility on the Windows host.
2. On each of the HBA initiator ports, set the following HBA driver parameters:
   - EnableNVMe = 1
   - NVMEMode = 0
   - LimTransferSize=1
3. Reboot the host.

Configure the Broadcom FC adapter in Windows for NVMe/FC

With the Broadcom adapter for FC/NVMe in a Windows environment, a `hostnqn` is associated with each host bus adapter (HBA) port. The `hostnqn` is formatted as follows.

Example

```
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9765
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9766
```

Enable MPIO for NVMe devices on the Windows host

1. Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
2. Open the MPIO properties.
3. From the **Discover Multi-Paths** tab, add the device ID listed for NVMe.

   MPIO becomes aware of the NVMe devices, which are visible under disk management.
4. Open **Disk Management** and go to **Disk Properties**.
5. From the **MPIO** tab, click **Details**.
6. Set the following MSDSM settings:
   - PathVerifiedPeriod: **10**
   - PathVerifyEnabled: **Enable**
   - RetryCount: **6**
   - RetryInterval: **1**
   - PDORemovedPeriod: **130**
7. Select the MPIO Policy **Round Robin with Subset**.

8. Change the registry values:

```
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInter
val DWORD -> 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\
UseCustomPathRecoveryInterval  DWORD-> 1
```

9. Reboot the host.

   The NVMe configuration is now complete on the Windows host.

Validate NVMe/FC

1. Validate that the Port Type is FC+NVMe.

   Now that NVMe is enabled, you should see the `Port Type` listed as `FC+NVMe`, as follows.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd listhba

Manageable HBA List

Port WWN       : 10:00:00:10:9b:1b:97:65
Node WWN       : 20:00:00:10:9b:1b:97:65
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 0
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 0
Port Type      : FC+NVMe
Model          : LPe32002-M2

Port WWN       : 10:00:00:10:9b:1b:97:66
Node WWN       : 20:00:00:10:9b:1b:97:66
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 1
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 1
Port Type      : FC+NVMe
Model          : LPe32002-M2
```

2. Validate that NVMe/FC subsystems have been discovered.

   The `nvme-list` command lists the NVMe/FC discovered subsystems.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:65

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:65

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:09:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0180
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:06:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0181
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available
Note: At present Namespace Management is not supported by NetApp Arrays.
```

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:66

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:66

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:07:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0140
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:08:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0141
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.
```

3. Validate that namespaces have been created.

   The `nvme-list-ns` command lists the namespaces for a specified NVMe target that lists the namespaces connected to the host.

```
PS C:\Program Files\Emulex\Util\OCManager> .\HbaCmd.exe nvme-list-ns
10:00:00:10:9b:1b:97:66 20:08:d0:39:ea:14:11:04 nq
.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159 0


Active Namespaces (attached to controller 0x0141):

                                       SCSI           SCSI
SCSI
   NSID           DeviceName        Bus Number    Target Number     OS
LUN
-----------  --------------------  ------------  ---------------
---------
0x00000001   \\.\PHYSICALDRIVE9         0              1              0
0x00000002   \\.\PHYSICALDRIVE10        0              1              1
0x00000003   \\.\PHYSICALDRIVE11        0              1              2
0x00000004   \\.\PHYSICALDRIVE12        0              1              3
0x00000005   \\.\PHYSICALDRIVE13        0              1              4
0x00000006   \\.\PHYSICALDRIVE14        0              1              5
0x00000007   \\.\PHYSICALDRIVE15        0              1              6
0x00000008   \\.\PHYSICALDRIVE16        0              1              7
```

NVMe/FC Host Configuration for Windows Server 2012 R2 with ONTAP

You can configure NVMe over Fibre Channel (NVMe/FC) on hosts running Windows Server 2012 R2 using ONTAP as the target.

NVMe/FC is supported on ONTAP 9.7 or later for Windows Server 2012.

Note that the Broadcom initiator can serve both NVMe/FC and FCP traffic through the same 32G FC adapter ports. For FCP and FC/NVMe, use MSDSM as the Microsoft Multipath I/O (MPIO) option.

See the Hardware Universe for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the NetApp Interoperability Matrix.

> You can use the configuration settings provided in this document to configure cloud clients connected to Cloud Volumes ONTAP and Amazon FSx for ONTAP.

Known limitations

Windows Failover Cluster (WFC) is not supported with ONTAP NVMe/FC because ONTAP does not currently support persistent reservations with NVMe/FC.

> The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI □ NVMe driver. This translational overhead does not necessarily impact performance, but it does negate the performance benefits of NVMe/FC. Thus, on Windows servers, NVMe/FC and FCP performance is the same, unlike on other operating systems such as Linux, where NVMe/FC performance is significantly better than that of FCP.

Enable NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

Steps

1. Install OneCommand Manager utility on the Windows host.
2. On each of the HBA initiator ports, set the following HBA driver parameters:
   - EnableNVMe = 1
   - NVMEMode = 0
   - LimTransferSize=1
3. Reboot the host.

Configure the Broadcom FC adapter in Windows for NVMe/FC

With the Broadcom adapter for FC/NVMe in a Windows environment, a `hostnqn` is associated with each host bus adapter (HBA) port. The `hostnqn` is formatted as follows.

Example

```
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9765
nqn.2017-01.com.broadcom:ecd:nvmf:fc:100000109b1b9766
```

Enable MPIO for NVMe devices on the Windows host

1. Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
2. Open the MPIO properties.
3. From the **Discover Multi-Paths** tab, add the device ID listed for NVMe.

   MPIO becomes aware of the NVMe devices, which are visible under disk management.
4. Open **Disk Management** and go to **Disk Properties**.
5. From the **MPIO** tab, click **Details**.
6. Set the following MSDSM settings:
   - PathVerifiedPeriod: **10**
   - PathVerifyEnabled: **Enable**
   - RetryCount: **6**
   - RetryInterval: **1**
   - PDORemovedPeriod: **130**
7. Select the MPIO Policy **Round Robin with Subset**.

8. Change the registry values:

```
HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInter
val DWORD -> 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\
UseCustomPathRecoveryInterval  DWORD-> 1
```

9. Reboot the host.

   The NVMe configuration is now complete on the Windows host.

Validate NVMe/FC

1. Validate that the Port Type is FC+NVMe.

   Now that NVMe is enabled, you should see the `Port Type` listed as `FC+NVMe`, as follows.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd listhba

Manageable HBA List

Port WWN       : 10:00:00:10:9b:1b:97:65
Node WWN       : 20:00:00:10:9b:1b:97:65
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 0
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 0
Port Type      : FC+NVMe
Model          : LPe32002-M2

Port WWN       : 10:00:00:10:9b:1b:97:66
Node WWN       : 20:00:00:10:9b:1b:97:66
Fabric Name    : 10:00:c4:f5:7c:a5:32:e0
Flags          : 8000e300
Host Name      : INTEROP-57-159
Mfg            : Emulex Corporation
Serial No.     : FC71367217
Port Number    : 1
Mode           : Initiator
PCI Bus Number : 94
PCI Function   : 1
Port Type      : FC+NVMe
Model          : LPe32002-M2
```

2. Validate that NVMe/FC subsystems have been discovered.

   The `nvme-list` command lists the NVMe/FC discovered subsystems.

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:65

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:65

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:09:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0180
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:06:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0181
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available
Note: At present Namespace Management is not supported by NetApp Arrays.
```

```
PS C:\Program Files\Emulex\Util\OCManager> .\hbacmd nvme-list
10:00:00:10:9b:1b:97:66

Discovered NVMe Subsystems for 10:00:00:10:9b:1b:97:66

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:07:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0140
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

NVMe Qualified Name     :  nqn.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159
Port WWN                :  20:08:d0:39:ea:14:11:04
Node WWN                :  20:05:d0:39:ea:14:11:04
Controller ID           :  0x0141
Model Number            :  NetApp ONTAP Controller
Serial Number           :  81CGZBPU5T/uAAAAAAAB
Firmware Version        :  FFFFFFFF
Total Capacity          :  Not Available
Unallocated Capacity    :  Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.
```

3. Validate that namespaces have been created.

   The `nvme-list-ns` command lists the namespaces for a specified NVMe target that lists the namespaces connected to the host.

```
PS C:\Program Files\Emulex\Util\OCManager> .\HbaCmd.exe nvme-list-ns
10:00:00:10:9b:1b:97:66 20:08:d0:39:ea:14:11:04 nq
.1992-
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_int
erop-57-159 0


Active Namespaces (attached to controller 0x0141):

                                       SCSI           SCSI
SCSI
   NSID           DeviceName        Bus Number    Target Number     OS
LUN
-----------  --------------------  ------------  ---------------
---------
0x00000001   \\.\PHYSICALDRIVE9         0               1              0
0x00000002   \\.\PHYSICALDRIVE10        0               1              1
0x00000003   \\.\PHYSICALDRIVE11        0               1              2
0x00000004   \\.\PHYSICALDRIVE12        0               1              3
0x00000005   \\.\PHYSICALDRIVE13        0               1              4
0x00000006   \\.\PHYSICALDRIVE14        0               1              5
0x00000007   \\.\PHYSICALDRIVE15        0               1              6
0x00000008   \\.\PHYSICALDRIVE16        0               1              7
```

Troubleshoot

Before troubleshooting any NVMe-oF failures for RHEL, OL, and SLES hosts, verify that you are running a configuration that is compliant to the Interoperability Matrix Tool (IMT) specifications and then proceed with the next steps to debug any host side issues.

> ⓘ The troubleshooting instructions are not applicable for AIX, Windows, and ESXi hosts.

Enable verbose logging

If you have an issue with your configuration, verbose logging can provide essential information for troubleshooting.

The procedure to set verbose logging for Qlogic (Qla2xxx) is different from the procedure to set LPFC verbose logging.

LPFC

Set the lpfc driver for NVMe/FC.

Steps

1. Set the `lpfc_log_verbose` driver setting to any of the following values to log NVMe/FC events.

```
#define LOG_NVME 0x00100000 /* NVME general events. */
#define LOG_NVME_DISC 0x00200000 /* NVME Discovery/Connect events.
*/
#define LOG_NVME_ABTS 0x00400000 /* NVME ABTS events. */
#define LOG_NVME_IOERR 0x00800000 /* NVME IO Error events. */
```

2. After setting the values, run the `dracut-f` command and reboot the host.
3. Verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf options lpfc
lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose 15728771
```

Qla2xxx

There is no specific qla2xxx logging for NVMe/FC similar to that for the `lpfc` driver. Instead, set the general qla2xxx logging level.

Steps

1. Append the `ql2xextended_error_logging=0x1e400000` value to the corresponding `modprobe qla2xxx conf` file.
2. Execute the `dracut -f` command and then reboot the host.
3. After reboot, verify that the verbose logging has been enabled:

```
# cat /etc/modprobe.d/qla2xxx.conf
```

Example output:

```
options qla2xxx ql2xnvmeenable=1
ql2xextended_error_logging=0x1e400000
# cat /sys/module/qla2xxx/parameters/ql2xextended_error_logging
507510784
```

Common nvme-cli errors and workarounds

The errors displayed by `nvme-cli` during `nvme discover`, `nvme connect`, or `nvme connect-all` operations and the workarounds are shown in the following table:

Error message	Probable cause	Workaround
`Failed to write to /dev/nvme-fabrics: Invalid argument`	Incorrect syntax	Verify that you are using the correct syntax for the `nvme discover`, `nvme connect`, and `nvme connect-all` commands.

<table>
<tr><th>Error message</th><th>Probable cause</th><th>Workaround</th></tr>
<tr><td><code>Failed to write to /dev/nvme-fabrics: No such file or directory</code></td><td>Multiple issues can trigger this, for example, providing wrong arguments to the NVMe commands is one of the common causes.</td><td>

* Verify that you have passed the correct arguments (such as, correct WWNN string, WWPN string, and more) to the commands.
* If the arguments are correct, but you still see this error, check whether the `/sys/class/scsi_host/host*/nvme_info` command output is correct, the NVMe initiator is displayed as `Enabled`, and the NVMe/FC target LIFs are correctly displayed under the remote ports sections. Example:

```
# cat
/sys/class/scsi_host/host*/nvme_info
NVME Initiator Enabled
NVME LPORT lpfc0 WWPN
x10000090fae0ec9d WWNN
x20000090fae0ec9d DID x012000 ONLINE
NVME RPORT WWPN x200b00a098c80f09
WWNN x200a00a098c80f09 DID x010601
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000000000006 Cmpl
0000000000000006
FCP: Rd 0000000000000071 Wr
0000000000000005 IO 0000000000000031
Cmpl 00000000000000a6 Outstanding
0000000000000001
NVME Initiator Enabled
NVME LPORT lpfc1 WWPN
x10000090fae0ec9e WWNN
x20000090fae0ec9e DID x012400 ONLINE
NVME RPORT WWPN x200900a098c80f09
WWNN x200800a098c80f09 DID x010301
TARGET DISCSRVC ONLINE
NVME Statistics
LS: Xmt 0000000000000006 Cmpl
0000000000000006
FCP: Rd 0000000000000073 Wr
0000000000000005 IO 0000000000000031
Cmpl 00000000000000a8 Outstanding
0000000000000001
```

* If the target LIFs are not displayed as above in the `nvme_info` command output, check the `/var/log/messages` and `dmesg` command outputs for any suspicious NVMe/FC failures, and report or fix accordingly.

</td></tr>
<tr><td></td><td></td><td></td></tr>
</table>

Error message	Probable cause	Workaround
`No discovery log entries to fetch`	Generally observed when the `/etc/nvme/hostnqn` string has not been added to the corresponding subsystem on the NetApp array or an incorrect `hostnqn` string has been added to the respective subsystem.	Verify that the exact `/etc/nvme/hostnqn` string is added to the corresponding subsystem on the NetApp array (verify using the `vserver nvme subsystem host show` command).
`Failed to write to /dev/nvme-fabrics: Operation already in progress`	Observed when the controller associations or specified operation is already created or in the process of being created. This could happen as part of the auto-connect scripts installed above.	None. Try running the `nvme discover` command again after some time. For `nvme connect` and `connect-all`, run the `nvme list` command to verify that the namespace devices are already created and displayed on the host.

When to contact technical support

If you are still facing issues, collect the following files and command outputs and contact NetApp support for further triage:

```
cat /sys/class/scsi_host/host*/nvme_info
/var/log/messages
dmesg
nvme discover output as in:
nvme discover --transport=fc --traddr=nn-0x200a00a098c80f09:pn
-0x200b00a098c80f09 --host-traddr=nn-0x20000090fae0ec9d:pn
-0x10000090fae0ec9d
nvme list
nvme list-subsys /dev/nvmeXnY
```

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

http://www.netapp.com/us/legal/copyright.aspx

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

http://www.netapp.com/us/legal/netapptmlist.aspx

Patents

A current list of NetApp owned patents can be found at:

https://www.netapp.com/us/media/patents-page.pdf

Privacy policy

https://www.netapp.com/us/legal/privacypolicy/index.aspx

Open source

Provides information about third-party copyright and licenses used in this product.

- Notice for AIX
- Notice for Linux
- Notice for Solaris
- Notice for Windows