



# **NVMe/FC Host Configurations**

## **ONTAP SAN Host**

NetApp  
March 29, 2021

# Table of Contents

- NVMe/FC Host Configurations ..... 1
  - RHEL ..... 1
  - SUSE ..... 24
  - Oracle Linux ..... 35
  - ESXi ..... 56

# NVMe/FC Host Configurations

## RHEL

### NVMe/FC Host Configuration for RHEL 8.1 with ONTAP

#### Supportability

NVMe/FC is supported on ONTAP 9.6 or later for the following versions of RHEL:

- RHEL 8.1

RHEL 8.1 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](#).

#### 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 script.
- By default, NVMe multipath is disabled. It must be manually enabled. Steps are provided in the section on Enabling NVMe/FC on RHEL 8.1.
- 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 RHEL 8.1.

#### Enabling NVMe/FC on RHEL 8.1

1. Install Red Hat Enterprise Linux 8.1 on the server.
2. After the installation is complete, verify that you are running the specified Red Hat Enterprise Linux kernel. See the [NetApp Interoperability Matrix](#) for the most current list of supported versions.

```
# uname -r
4.18.0-147.el8.x86_64
```

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 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
```

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-09d5a156fbc
```

```
*> 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-09d5a156fbc
```



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.

7. 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 `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 come bundled with the OS are called inbox drivers. If you download the outbox drivers (drivers that are not included with the 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 & 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 nvmeofc
nvmeofc-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 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.rhel_141_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/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
    }
  ]
}
```

### Enabling 1MB I/O Size for Broadcom NVMe/FC

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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
```

## LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, 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 RHEL 8.2 with ONTAP

### Supportability

NVMe/FC is supported on ONTAP 9.6 or later for 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. For the most current list of supported configurations and versions, see the [NetApp Interoperability Matrix](#).

### Known limitations

For RHEL 8.2, in-kernel NVMe multipath remains disabled by default. Therefore, you need to enable it manually. Steps for doing so are provided in the next section, "Enabling NVMe/FC on RHEL 8.2".



## Enabling NVMe/FC on RHEL 8.2

1. Install Red Hat Enterprise Linux 8.2 GA on the server.

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 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/`.

2. 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](#) for the most current list of supported versions.

3. Install the `nvme-cli` package.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.9.5.el8.x86_64
```

4. Enable in-kernel NVMe multipath.

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-193.el8.x86_64
```

5. 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.

6. Reboot the host.
7. 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, we recommend 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 do this by adding the `enable_foreign` setting to the `/etc/multipath.conf`, as shown below.

```
# cat /etc/multipath.conf
defaults {
    enable_foreign NONE
}
```

8. Restart the multipathd daemon by running a `systemctl restart multipathd`.

## Configuring the Broadcom FC adapter for NVMe/FC

For the most current list of supported adapters see the see the [NetApp Interoperability Matrix](#).

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
0000000000000000a
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
0000000000000000a
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
0000000000000000a
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
0000000000000000a
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
```

## 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/JvAAAAAAC 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_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/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
    }
  ]
}

```

### LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, 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 RHEL 8.3 with ONTAP

### Supportability

NVMe/FC is supported on ONTAP 9.6 or later for RHEL 8.3. The RHEL 8.3 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. For the most current list of supported configurations and versions, see the [NetApp Interoperability Matrix](#).

### Known limitations

For RHEL 8.3, in-kernel NVMe multipath remains disabled by default. Therefore, you need to enable it manually. Steps for doing so are provided in the next section, "Enabling NVMe/FC on RHEL 8.3".

### Enabling NVMe/FC on RHEL 8.3

1. Install Red Hat Enterprise Linux 8.3 GA on the server.

If you are upgrading from RHEL 8.2 to RHEL 8.3 using `yum update/upgrade`, your `/etc/nvme/host*` files might be lost. To avoid 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/`.

2. After the installation is complete, verify that you're running the specified Red Hat Enterprise Linux kernel.

```
# uname -r
4.18.0-240.el8.x86_64
```

See the [NetApp Interoperability Matrix](#) for the most current list of supported versions.

3. Install the `nvme-cli` package.

```
# rpm -qa|grep nvme-cli
nvme-cli-1.12-2.el8.x86_64
```

4. Enable in-kernel NVMe multipath.

```
# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-
4.18.0-240.el8.x86_64
```

5. On the RHEL 8.3 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

::> vservers nvme subsystem host show -vservers vs_fc_nvme_141

::> vservers nvme subsystem host show -vservers vs_fc_nvme_141
Vserver          Subsystem      Host          NQN
-----
vs_fc_nvme_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 NQM string from `/etc/nvme/hostnqn` on the host.

6. Reboot the host.

7. Update the `enable_foreign` setting (*optional*).



If you intend to run both NVMe and SCSI traffic on the same RHEL 8.3 co-existent host, we recommend 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 do this by adding the `enable_foreign` setting to the `/etc/multipath.conf`, as shown below.

```
# cat /etc/multipath.conf
defaults {
    enable_foreign NONE
}
```

8. Restart the multipathd daemon by running a `systemctl restart multipathd`.

## 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 and properly discovered on the host.

```

/dev/nvme0n1      814vWBNRwf9HAAAAAAAAAB NetApp ONTAP Controller
1                85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n2      814vWBNRwf9HAAAAAAAAAB NetApp ONTAP Controller
2                85.90 GB / 85.90 GB     4 KiB + 0 B   FFFFFFFF
/dev/nvme0n3      814vWBNRwf9HAAAAAAAAAB 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.

```

# nvme netapp ontapdevices -o column
Device          Vserver          Namespace Path
NSID            UUID              Size
-----
-----
-----
/dev/nvme0n1    vs_fc_nvme_141
/vol/fc_nvme_141_vol_1_1_0/fc_nvme_141_ns    1    72b887b1-5fb6-
47b8-be0b-33326e2542e2    85.90GB
/dev/nvme0n2    vs_fc_nvme_141
/vol/fc_nvme_141_vol_1_1_0_0/fc_nvme_141_ns    2    04bf9f6e-9031-
40ea-99c7-a1a61b2d7d08    85.90GB
/dev/nvme0n3    vs_fc_nvme_141
/vol/fc_nvme_141_vol_1_1_1/fc_nvme_141_ns    3    264823b1-8e03-
4155-80dd-e904237014a4    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
    },
    {
      "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
    }
  ]
}
```

## Configuring the Broadcom FC adapter for NVMe/FC

For the most current list of supported adapters see the see the [NetApp Interoperability Matrix](#).

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
0000000000000000a
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
0000000000000000a
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
```

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
0000000000000000a
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
0000000000000000a
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
```

## LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose
15728771
```

# SUSE

## NVMe/FC Host Configuration for SUSE Linux Enterprise Server 15 SP1 with ONTAP

### Supportability

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](#).

### Known limitations

None. Native NVMe/FC auto-connect scripts are included in the `nvme-cli` package. You can use the native `inbox lpfc` driver on SLES15 SP1.

### Enabling 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.

### 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. Verify that you are using the recommended Broadcom lpc 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
...
```

## 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 80BADBKkB/JvAAAAAAC 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
    }
  ]
}

```

## Enabling 1MB I/O Size for Broadcom NVMe/FC

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

## LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, 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 SP2 with ONTAP

### Supportability

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](#).

### Known limitations

None.

### Enabling 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_fc_nvme_145
Vserver Subsystem Host NQN
-----
-----
vs_fc_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.

```

#### 4. 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. 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
ffffffffffffb046f
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
ffffffffffffaf794
abort 00000b73 noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 0000159d Err 000135c3
```

## 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
```

```
# 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 814vWBNRwfbGAAAAAAB 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
    },
  ],
}

```

## Enabling 1MB I/O Size for Broadcom NVMe/FC

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

## LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose
15728771
```

## Oracle Linux

### 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](#).

#### Known limitations

None.

#### Enabling 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
```

```
::> vservers 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.

## 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. 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
0000000000000000a
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
0000000000000000a
abort 000016bb noxri 00000000 nondlp 00000000 qdepth 00000000 wqerr
00000000 err 00000000
FCP CMPL: xb 00001f50 Err 0000d9f8
```

## 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 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
Format        FW Rev      Namespace Usage
-----
-----
-----
/dev/nvme0n1  814vWBNRwf9HAAAAAAAAAB NetApp ONTAP Controller
1              85.90 GB / 85.90 GB    4 KiB + 0 B  FFFFFFFF
/dev/nvme0n2  814vWBNRwf9HAAAAAAAAAB NetApp ONTAP Controller
2              85.90 GB / 85.90 GB    4 KiB + 0 B  FFFFFFFF
/dev/nvme0n3  814vWBNRwf9HAAAAAAAAAB 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_1_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
  },
]
}

```

### Enabling 1MB I/O Size for Broadcom NVMe/FC

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

# 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](#).

## 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.

## Enabling 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 nvmecli
nvmecli-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.

## 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 nvme-fc-connect-12.6.61.0-1.noarch.rpm
```

4. Verify that the auto-connect scripts are installed.

```
# rpm -qa | grep nvme_fc
nvme_fc-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
```

## 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 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  814vWBNRwfBCAAAAAAB NetApp ONTAP Controller      2
107.37 GB / 107.37 GB  4 KiB + 0 B  FFFFFFFF
/dev/nvme0n2  814vWBNRwfBCAAAAAAB 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
    }
  ]
}

```

## Enabling 1MB I/O Size for Broadcom NVMe/FC

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

## 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](#).

## 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.

## 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 nvme-fc-connect-12.4.65.0-1.noarch.rpm
```

4. Verify that the auto-connect scripts are installed.

```
# rpm -qa | grep nvme-fc
nvme-fc-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/JvAAAAAAC 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

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

## 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](#).

## 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.

## 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 nvme-fc-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 nvmeofc
nvmeofc-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/JvAAAAAAC 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

The `lpfc_sg_seg_cnt` parameter must be set to 256 in order for the host to issue 1MB size I/O.

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

```

### LPFC Verbose Logging

1. You can 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 any of these values, run `dracut-f` and reboot host.
3. After rebooting, verify the settings.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_log_verbose=0xf00083

# cat /sys/module/lpfc/parameters/lpfc_log_verbose
15728771
```

## ESXi

### NVMe/FC Host Configuration for ESXi 7.0 with ONTAP

#### Supportability

NVMe/FC is supported on ONTAP 9.7 or later for ESXi 7.0.

ESXi initiator host can run both NVMe/FC & FCP traffic through the same 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](#).

#### Known limitations

The following are not supported:

- RDM mapping
- VVols

#### Enabling NVMe/FC with ANA

1. Disable the `HppManageDegradedPaths` parameter for improved interoperability with ONTAP:

```
# esxcfg-advcfg -s 0 /Misc/HppManageDegradedPaths
```

2. Reboot the host.
3. After reboot, verify that the `HppManageDegradedPaths` parameter is now disabled:

```
# esxcfg-advcfg -g /Misc/HppManageDegradedPaths
Value of HppManageDegradedPaths is 0
```

4. Check the ESXi host NQN string and verify that it matches with the host NQN string for the corresponding subsystem on the ONTAP array.

#### Example

```
# esxcli nvme info get
Host NQN: nqn.2014-08.com.vmware:nvme:chat-54-113

*> vserver nvme subsystem host show -vserver co_nv_fc_esx
Vserver Subsystem Host NQN
-----
-----
co_nv_fc_esx
      subsys_chat_54_113_nvme
              nqn.2014-08.com.vmware:nvme:chat-54-113
```

### Configuring the Broadcom FC adapter for NVMe/FC

1. Install the recommended lpfc driver by copying it to a temporary folder and then executing the following command:

```
# esxcli software vib install -d /tmp/t/Emulex-FCoE-FC-lpfc-12.4.224.0-
offline-bundle-13621872.zip --no-sig-check
Installation Result
  Message: The update completed successfully, but the system needs to
be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: EMU_bootbank_lpfc_12.4.224.0-1OEM.688.0.0.13621872
  VIBs Removed: EMU_bootbank_lpfc_12.4.211.6-1OEM.688.0.0.13621872
  VIBs Skipped:
```

2. If necessary, set the lpfc driver parameter `lpfc_enable_fc4_type=3` for enabling NVMe/FC support in the lpfc driver:



This parameter is set by default for the LPe35000-series adapters. You must perform the following step to set it manually for LPe32000-series & LPe31000-series adapters.

```
# esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3
```

3. Use the `elxmgmt` utility to upgrade the Broadcom FC adapter firmware to the recommended version:

```
# esxcli software vib install -d /tmp/t/Emulex-elxmgmt-6.8.7-12.4.211.7.zip --no-sig-check
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: EMU_bootbank_emu-esx-elxmgmt_12.4.211.7-01
  VIBs Removed:
  VIBs Skipped:
...
```

4. Reboot the host.
5. After reboot, verify that the recommended lpfc driver and adapter firmware versions have applied and the initiator ports are online:

```
# esxcli storage san fc list
Adapter: vmhba3
  Port ID: 010600
  Node Name: 20:00:00:90:fa:e0:ec:8e
  Port Name: 10:00:00:90:fa:e0:ec:8e
  Speed: 32 Gbps
  Port Type: NPort
  Port State: ONLINE
  Model Description: Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre
Channel Adapter
  Hardware Version: 0000000c
  OptionROM Version: 12.4.217.2
  Firmware Version: 12.4.217.2
  Driver Name: lpfc
  DriverVersion: 12.4.224.0

Adapter: vmhba4
  Port ID: 010F00
  Node Name: 20:00:00:90:fa:e0:ec:8f
  Port Name: 10:00:00:90:fa:e0:ec:8f
  Speed: 32 Gbps
  Port Type: NPort
  Port State: ONLINE
  Model Description: Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre
Channel Adapter
  Hardware Version: 0000000c
  OptionROM Version: 12.4.217.2
  Firmware Version: 12.4.217.2
  Driver Name: lpfc
  DriverVersion: 12.4.224.0
```

## Validating NVMe/FC

1. Verify that the ONTAP target NVMe/FC controllers are properly discovered on the ESXi host:

```
# esxcli nvme controller list
```

```
Name
```

```
Controller Number  Adapter  Transport Type  Is Online
```

```
-----  
-----  
-----  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_01#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                259  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_09#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                263  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_11#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                267  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_10#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                265  
vmhba32  FC                false  
nqn.1992-  
08.com.netapp:sn.e7f89c2c245d11e9975300a098dfce55:subsystem.interop_57_v  
m_02#vmhba32#204900a098dfe3d1:204a00a098dfe3d1                261  
vmhba32  FC                false
```

## 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.0d12b7cd97344be8a53b7913f8f72f04 : NVMe Fibre Channel Disk
(uuid.0d12b7cd97344be8a53b7913f8f72f04)
  vmhba65:C0:T9:L30 LUN:30 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4d:00:a0:98:df:e3:d1
  vmhba64:C0:T9:L30 LUN:30 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4c:00:a0:98:df:e3:d1
  vmhba64:C0:T5:L30 LUN:30 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4a:00:a0:98:df:e3:d1
  vmhba65:C0:T0:L30 LUN:30 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:49:00:a0:98:df:e3:d1 WWPN: 20:4b:00:a0:98:df:e3:d1

uuid.49de7683950d47c9898f51443d893910 : NVMe Fibre Channel Disk
(uuid.49de7683950d47c9898f51443d893910)
  vmhba65:C0:T12:L39 LUN:39 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:27:00:a0:98:df:e3:d1
  vmhba65:C0:T13:L39 LUN:39 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8f WWPN: 10:00:00:90:fa:e0:ec:8f Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:29:00:a0:98:df:e3:d1
  vmhba64:C0:T12:L39 LUN:39 state:active fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:3b:00:a0:98:df:e3:d1
  vmhba64:C0:T13:L39 LUN:39 state:standby fc Adapter: WWNN:
20:00:00:90:fa:e0:ec:8e WWPN: 10:00:00:90:fa:e0:ec:8e Target: WWNN:
20:3a:00:a0:98:df:e3:d1 WWPN: 20:28:00: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 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
```

```
fc.20000090fae0ec8f:10000090fae0ec8f-
fc.204900a098dfe3d1:204d00a098dfe3d1-
uuid.1aa669c5376240a28ae47d8d549586ea
  Runtime Name: vmhba65:C0:T9:L33
  Device: uuid.1aa669c5376240a28ae47d8d549586ea
  Device Display Name: NVMe Fibre Channel Disk
(uuid.1aa669c5376240a28ae47d8d549586ea)
  Path State: active
```

```
fc.20000090fae0ec8e:10000090fae0ec8e-
fc.204900a098dfe3d1:204a00a098dfe3d1-
uuid.1aa669c5376240a28ae47d8d549586ea
  Runtime Name: vmhba64:C0:T5:L33
  Device: uuid.1aa669c5376240a28ae47d8d549586ea
  Device Display Name: NVMe Fibre Channel Disk
(uuid.1aa669c5376240a28ae47d8d549586ea)
  Path State: standby
```

```
:leveloffset: -1
```

```
= Windows
```

```
:leveloffset: +1
```

```
[[IDe33bc71c427003a80b912e5b177d6b73]]
```

```
= NVMe/FC Host Configuration for Windows Server 2019 with ONTAP
```

```
:toc: macro
```

```
:hardbreaks:
```

```
:toclevels: 1
```

```
:nofooter:
```

```
:icons: font
```

```
:linkattrs:
```

```
:imagesdir: ./media/
```

```
== Supportability
```

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 link:[https://hwu.netapp.com/Home/Index\[Hardware Universe\]](https://hwu.netapp.com/Home/Index[Hardware Universe]) for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the



link:[https://mysupport.netapp.com/matrix/\[NetApp Interoperability Matrix\]](https://mysupport.netapp.com/matrix/[NetApp%20Interoperability%20Matrix])].

#### == Known limitations

Windows Failover Cluster (WFC) is not supported with ONTAP NVMe/FC because ONTAP does not currently support persistent reservations with NVMe/FC.

NOTE: The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI  $\square$  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.

#### == Enabling NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

- . Install OneCommand Manager utility on the Windows host.
- . On each of the HBA initiator ports, set the following HBA driver parameters:
  - +
    - \* EnableNVMe = 1
    - \* NVMEMode = 0
    - \* LimTransferSize=1
  - +
    - . Reboot the host.

#### == Configuring 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
```

#### === Enabling MPIO for NVMe devices on the Windows host

- . Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
- . Open the MPIO properties.

```

. 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.
+
. Open *Disk Management* and go to *Disk Properties*.
. From the *MPIO* tab, click *Details*.
. Set the following MSDSM settings:
+
* PathVerifiedPeriod: *10*
* PathVerifyEnabled: *Enable*
* RetryCount: *6*
* RetryInterval: *1*
* PDORemovedPeriod: *130*
+
. Select the MPIO Policy *Round Robin with Subset*.
. Change the registry values:
+

```

HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInterval DWORD → 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio \Parameters\ UseCustomPathRecoveryInterval DWORD→ 1

```

+
. Reboot the host.
+
The NVMe configuration is now complete on the Windows host.

== Validating NVMe/FC

. 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> .\hbcamd 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

```
. 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> .\hbcmd 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_interop-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/uAAAAAAB  
Firmware Version : FFFFFFFF  
Total Capacity : Not Available  
Unallocated Capacity : Not Available
```

```
NVMe           Qualified           Name           :           nqn.1992-  
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_interop-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/uAAAAAAB  
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> .\hbcmd 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_interop-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/uAAAAAAB
Firmware Version : FFFFFFFF
Total Capacity : Not Available
Unallocated Capacity : Not Available
```

```
NVMe           Qualified           Name           :           nqn.1992-08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_interop-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/uAAAAAAB
Firmware Version : FFFFFFFF
Total Capacity : Not Available
Unallocated Capacity : Not Available
```

Note: At present Namespace Management is not supported by NetApp Arrays.

```
. 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_interop-57-159 0
```

Active Namespaces (attached to controller 0x0141):

NSID	DeviceName	SCSI Bus Number	SCSI Target Number	SCSI 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

```
[[IDfc64a5a7647082a204b1ac0ea28c07c1]]
```

```
= NVMe/FC Host Configuration for Windows Server 2016 with ONTAP
```

```
:toc: macro
```

```
:hardbreaks:
```

```
:toclevels: 1
```

```
:nofooter:
```

```
:icons: font
```

```
:linkattrs:
```

```
:imagesdir: ./media/
```

```
== Supportability
```

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 link:[https://hwu.netapp.com/Home/Index\[Hardware Universe\]](https://hwu.netapp.com/Home/Index[Hardware Universe]) for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the link:[https://mysupport.netapp.com/matrix/\[NetApp Interoperability Matrix\]](https://mysupport.netapp.com/matrix/[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.

NOTE: The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI  $\square$  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.

== Enabling NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

- . Install OneCommand Manager utility on the Windows host.
  
- . On each of the HBA initiator ports, set the following HBA driver parameters:
  - +
    - \* EnableNVMe = 1
    - \* NVMEMode = 0
    - \* LimTransferSize=1
  - +
    - . Reboot the host.

== Configuring 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
```

=== Enabling MPIO for NVMe devices on the Windows host

- . Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.
- . Open the MPIO properties.
- . 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.
- +
  - . Open *\*Disk Management\** and go to *\*Disk Properties\**.
  - . From the *\*MPIO\** tab, click *\*Details\**.
  - . Set the following MSDSM settings:
    - +
      - \* PathVerifiedPeriod: *\*10\**
      - \* PathVerifyEnabled: *\*Enable\**
      - \* RetryCount: *\*6\**

```
* RetryInterval: *1*
* PDORemovedPeriod: *130*
+
. Select the MPIO Policy *Round Robin with Subset*.
. Change the registry values:
+
```

HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInterval DWORD → 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\UseCustomPathRecoveryInterval DWORD → 1

```
+
. Reboot the host.
+
The NVMe configuration is now complete on the Windows host.

== Validating NVMe/FC

. 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

```
. 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> .\hbcmd 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_interop-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/uAAAAAAB  
Firmware Version : FFFFFFFF  
Total Capacity : Not Available  
Unallocated Capacity : Not Available
```

```
NVMe           Qualified           Name           :           nqn.1992-  
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_interop-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/uAAAAAAB  
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> .\hbcmd 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_interop-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/uAAAAAAB  
 Firmware Version : FFFFFFFF  
 Total Capacity : Not Available  
 Unallocated Capacity : Not Available

NVMe Qualified Name : nqn.1992-08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win\_nvme\_interop-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/uAAAAAAB  
 Firmware Version : FFFFFFFF  
 Total Capacity : Not Available  
 Unallocated Capacity : Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.

```
. 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_interop-57-159 0
```

Active Namespaces (attached to controller 0x0141):

NSID	DeviceName	SCSI Bus Number	SCSI Target Number	SCSI 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

```
[[ID4fa061668eeebab4a7f62a27c1a18c55]]
= NVMe/FC Host Configuration for Windows Server 2012 R2 with ONTAP
:toc: macro
:hardbreaks:
```

```
:toclevels: 1
:nofooter:
:icons: font
:linkattrs:
:imagesdir: ./media/
```

#### == Supportability

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 link:[https://hwu.netapp.com/Home/Index\[Hardware Universe\]](https://hwu.netapp.com/Home/Index[Hardware Universe]) for a list of supported FC adapters and controllers. For the most current list of supported configurations & versions, see the link:[https://mysupport.netapp.com/matrix/\[NetApp Interoperability Matrix\]](https://mysupport.netapp.com/matrix/[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.

NOTE: The external driver shipped by Broadcom for Windows NVMe/FC is not a true NVMe/FC driver but a translational SCSI  $\square$  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.

#### == Enabling NVMe/FC on a Windows initiator host

Follow these steps to enable FC/NVMe on the Windows initiator host:

- . Install OneCommand Manager utility on the Windows host.
  
- . On each of the HBA initiator ports, set the following HBA driver parameters:
  - +
  - \* EnableNVMe = 1
  - \* NVMEMode = 0
  - \* LimTransferSize=1
  - +
- . Reboot the host.

## == Configuring 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
```

## === Enabling MPIO for NVMe devices on the Windows host

. Install Windows Host Utility Kit 7.1 to set the driver parameters that are common to both FC and NVMe.

. Open the MPIO properties.

. 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.

+

. Open *\*Disk Management\** and go to *\*Disk Properties\**.

. From the *\*MPIO\** tab, click *\*Details\**.

. Set the following MSDSM settings:

+

\* PathVerifiedPeriod: \*10\*

\* PathVerifyEnabled: \*Enable\*

\* RetryCount: \*6\*

\* RetryInterval: \*1\*

\* PDORemovedPeriod: \*130\*

+

. Select the MPIO Policy *\*Round Robin with Subset\**.

. Change the registry values:

+

HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\PathRecoveryInterval DWORD → 30

HKLM\SYSTEM\CurrentControlSet\Services\mpio\Parameters\UseCustomPathRecoveryInterval DWORD → 1

```
+
. Reboot the host.
+
The NVMe configuration is now complete on the Windows host.

== Validating NVMe/FC

. 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> .\hba.cmd 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
```

```
. 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> .\hbcmd 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_interop-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/uAAAAAAB  
Firmware Version : FFFFFFFF  
Total Capacity : Not Available  
Unallocated Capacity : Not Available
```

```
NVMe           Qualified           Name           :           nqn.1992-  
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_interop-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/uAAAAAAB  
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> .\hbcmd 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_interop-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/uAAAAAAB  
Firmware Version : FFFFFFFF  
Total Capacity : Not Available  
Unallocated Capacity : Not Available
```

```
NVMe           Qualified           Name           :           nqn.1992-  
08.com.netapp:sn.a3b74c32db2911eab229d039ea141105:subsystem.win_nvme_interop-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/uAAAAAAB  
 Firmware Version : FFFFFFFF  
 Total Capacity : Not Available  
 Unallocated Capacity : Not Available

Note: At present Namespace Management is not supported by NetApp Arrays.

```

. 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_interop-57-159 0
  
```

Active Namespaces (attached to controller 0x0141):

NSID	DeviceName	SCSI Bus Number	SCSI Target Number	SCSI 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

```

:leveloffset: -1

:leveloffset: -1

:leveloffset: -1

<<<
*Copyright Information*
  
```

Copyright © 2021 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice.

NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and

Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

\*Trademark Information\*

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.