

NVMe Configuration over Fibre Channel for Lenovo ONTAP 9.5 or greater

Introduction

NVMe is a data storage protocol that delivers the fastest response times for business-critical enterprise applications. Its broader version NVMe over Fabrics (NVMe-oF) encompasses the entire data path, from server to network to storage system.

This is a step-by-step guide for you to configure NVMe over fibre channel (NVMe/FC). The main procedures are:

1. Configure the initiator.
2. Configure the DM storage target.

Configure SUSE® Linux Enterprise Server as Initiator

Before you begin

1. You have a DM7000F with 32 GB dual port Fibre Channel target cards (4XC7A14396) in each controller.
2. You have a ThinkSystem™ DB class switch running Fabric OS 8.2.0 or higher.
3. You have a server running SUSE Linux Enterprise Server SLES 15 (kernel 4.12.14-25.25.1).
4. You have a server with the ThinkSystem Emulex LPe 32002 or LPe 35002 dual port Fibre Channel adapter.
5. You have upgraded the LPe 32002 or LPe 35002 Dual Port HBA adapter firmware. The latest versions is available at: <https://datacentersupport.lenovo.com>.

Note:

All Fibre Channel Zones must be configured so that each host port can see at least one port from each controller. For example:

Host port	Controller and port
HBA port 1	Controller A port 1, Controller B port 1
HBA port 2	Controller A port 2, Controller B port 2

6. You have downloaded and installed the compatible Light Pulse Fibre Channel (LPFC) driver on your host. The latest version is also available at: <https://datacentersupport.lenovo.com>.
7. You have downloaded the command line interface tool for NVMe (nvme-cli 1.5-7.19.1).

Procedure

Step 1. Enable SLES 15 as the initiator for NVMe/FC.

```
# cat /etc/modprobe.d/lpfc.conf
options lpfc lpfc_enable_fc4_type=3
```

Step 2. Reload the LPFC driver.

Step 3. Rebuild the initramfs to enable NVMe on start up.

```
dracut -force
```

Step 4. Reboot the host.

After you finish

Query the NVMe Qualified Name (NQN) for the host and record it for future use.

```
# cat /etc/nvme/hostnqn
```

Configure the DM storage target

The following are steps for configuring NVMe/FC. Use ThinkSystem Storage Manager for DM Series for the configurations.

Before you begin

1. You have deployed and installed Lenovo Data ONTAP software (9.5 or greater).
2. SLES 15 is installed on the initiator for NVMe/FC.
3. At least one non-root aggregate is available.

Procedure

- Step 1. Create and configure the Storage Virtual Machine (SVM) that will host the NVMe subsystem and namespaces.
- a. Create the SVM.
 1. Select **Storage → SVMs** from the left pane. The SVMs page is displayed.
 2. Click **Create**. The SVM creation wizard shows up.
 - b. Continue the popped-up wizard to set up the SVM.
 - ① Configure basic details for the SVM.
 1. Provide a name for the SVM.
 2. Select **NVMe** as the data protocol.
 3. Select the appropriate IPspace, security style and root aggregate.
 4. Click **Submit & Continue**.
 - ② Configure NVMe protocol.
 1. Specify the network interface names and adapter port information as desired.
 2. Provision a subsystem by specifying its name and Host OS. Provide the Host NQN you queried in the previous section.
 3. Provision an NVMe namespace by specifying the size, block size, volume, and aggregate information. Click **+** to provision more namespaces. A new volume will be created for each namespace specified on the specified aggregate.
 4. Click **Submit & Continue**.
 - ③ (Optional) Specify SVM administrator details.
 1. Specify the username and password for the administrator.
 2. Select the “Create a new LIF for SVM management” checkbox if SVM management is desired and enter the details for the management LIF. For more information, see [Configuring the administration details of an SVM](#).
 3. Click **Submit & Continue** to apply changes or **Skip**.
 - ④ Acknowledge SVM creation.

Click **OK**. The SVM is set up along with the new namespace(s).
- Step 2. If you need to manually create network interfaces for NVMe, perform the following steps.
- a. Create a new network interface for NVMe.

1. Select **Network** → **Network interfaces** from the left pane. The Network Interfaces page is displayed.
 2. Click **Create**. The Create Network Interface window is displayed.
- b. Configure basic details for the interface.
1. Provide a name for the network interface.
 2. Select **Serves Data** as the interface role.
 3. Assign the interface to the SVM created previously by selecting the SVM from the drop-down list.
 4. Select **NVMe** for protocol access.
 5. Choose the home node and port.
 6. Click **Create**. An interface for NVMe is created.
- c. Repeat step **a** and **b** for each network interface desired.

Step 3. Verify the namespace mapping.

- a. Select **Storage** → **NVMe** → **NVMe Namespaces** from the left pane. The NVMe namespace page is displayed.
- b. Select the SVM you created from the drop-down list next to the page title. Verify that all the created namespaces appear and they're associated with appropriate subsystems.
- c. Select **Storage** → **NVMe** → **Subsystems** from the left pane. The NVMe subsystem page is displayed.
- d. Highlight one subsystem from the list and click **Edit**. Verify that the correct host NQNs and mapped namespaces show up.

Step 4. Connect namespaces to the initiator.

- a. Ensure each host port can discover the NVMe target.

```
nvme discover --transport=fc --traddr=nn-<Target Port WWNN>:pn-<Target Port WWPN> --host-
traddr=nn-<Initiator Port WWNN>:pn-<Initiator Port WWPN>
```

If the discovery is successful, the output should look like the following:

```
Discovery Log Number of Records 1, Generation counter 4
====Discovery Log Entry 0====
trtype: fc
adrfam: fibre-channel
subtype: nvme subsystem
trreq: not specified
portid: 0
trsvcid: none
subnqn: nqn.1992-08.com.netapp:sn.XXXXXXXXXX:subsystem.<NVMe_subsystem_name>
traddr: nn-0XXXXXXXXXXXXXXXXX:pn-0XXXXXXXXXXXXXXXXX
```

- b. Connect all available namespaces.

```
nvme connect-all --transport=fc --traddr=nn-<Target Port WWNN>:pn-<Target Port WWPN> --host-
traddr=nn-<Initiator Port WWNN>:pn-<Initiator Port WWPN>
```

The messages log should show entries similar to the following:

```
Jul 02 03:06:23 host-XXX kernel: nvme nvme1: queue_size 128 > ctrl maxcmd 32, reducing to queue_size
Jul 02 03:06:23 host-XXX kernel: nvme nvme1: NVME-FC{1}: controller connect complete
Jul 02 03:06:23 host-XXX kernel: nvme nvme1: NVME-FC{1}: new ctrl: NQN "nqn.2014-08.org.nvmexpress.discovery"
Jul 02 03:06:24 host-XXX kernel: nvme nvme1: Removing ctrl: NQN "nqn.2014-08.org.nvmexpress.discovery"
Jul 02 03:06:25 host-XXX kernel: nvme nvme1: ANA group 1: optimized.
Jul 02 03:06:25 host-XXX kernel: nvme nvme1: ANA group 2: inaccessible.
Jul 02 03:06:25 host-XXX kernel: nvme nvme1: ANA group 3: optimized
```

- c. Repeat step **a** and **b** for each set of ports.

Note: Query WWPNs and WWNNs for the target and host. They should begin with "0x".

```
# cat /sys/class/scsi_host/host*/nvme_info
```

Step 5. Reboot the host.

After you finish

Re-connect target namespaces using command referenced in Step 4b. Ensure that they are attached properly by checking either of the following:

1. Whether the devices appear in `/dev/nvme0n<Namespace ID>`.
2. Whether messages indicating the connection(s) were successful appear in the system log file or in `dmesg`.
3. (If you are using the NVMe Command Line Tool) Whether all namespaces appear after you execute the command `nvme list`.

Support for asymmetric namespace access (ANA)

ANA is a protocol for monitoring and communicating path states to the host operating system's Multipath I/O (MPIO) or multipath stack, which uses information communicated through ANA to select and manage multiple paths between the initiator and target.

SLES 15 is the only host OS that supports ANA as part of its initiator stack. In its latest version, round-robin is set as the default multipathing policy for better use of the multiple active paths. For cases where round-robin is not the default multipathing policy or it's not even supported on your host, update SLES 15 to kernel 4.12.14-25.25.1 and install the `nvme-cli` package (1.5-7.19.1) mentioned in the prerequisites.

Trademarks

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2019 Lenovo.

Fourth Edition (September 2019)

© Copyright Lenovo 2018, 2019.

Lenovo and the Lenovo logo are trademarks of Lenovo in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.