

HITACHI

Solution Implementation Guide

Configuring FC-NVMe on VMware ESXi Host Using VSP One Block High End Storage Systems

MK-SL-467-00

Configuring FC-NVMe on VMware ESXi Host Using VSP One Block High End Storage Systems

Table of Contents

Notices and Disclaimer	3
About This Document.....	4
Introduction.....	4
Intended Audience.....	4
Value Proposition	4
Document Revisions	5
Comments.....	5
Overview	6
Key Components	6
Connectivity Diagram.....	6
VMware ESXi Systems and Multipathing.....	7
NVMe on ESXi Host.....	7
ESXi Multipathing	7
VSP One Block High End Storage System	8
Hardware Requirements	8
Software Requirements.....	9
CCI Configuration.....	9
Storage Side Configuration.....	10
Configuring the Storage Port.....	10
Configuring the NVM Subsystem	16
Registering the NVM Subsystem Port	16

Configuring the Namespace Security Settings..... 17

Creating a Namespace by Registering LDEV 18

Host-Side Configuration..... 21

Part 1..... 21

Part 2 22

Notices and Disclaimer

© 2026 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara (collectively, "Hitachi"). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. "Materials" mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara at https://support.HitachiVantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara.

By using this software, you agree that you are responsible for:

- 1) Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals to access relevant data; and
- 2) Verifying that data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

EXPORT CONTROLS - Licensee will comply fully with all applicable export laws and regulations of the United States and other countries, and Licensee shall not export, or allow the export or re-export of, the Software, API, or Materials in violation of any such laws or regulations. By downloading or using the Software, API, or Materials, Licensee agrees to the foregoing and represents and warrants that Licensee is not located in, under the control of, or a national or resident of any embargoed or restricted country.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, Flash Copy, IBM, Lotus, MVS, OS/390, PowerPC, RS6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z/VM, BCPi™ and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, the Microsoft Corporate Logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or web site are properties of their respective owners.

About This Document

Introduction

This document describes how to configure FC-NVMe on a VMware ESXi system using the Hitachi Virtual Storage Platform One Block High End (VSP One BHE) family storage systems.

Intended Audience

This document is intended for a technical audience who needs to configure FC-NVMe for VSP One BHE series storage systems.

Value Proposition

VSP One Block 85 delivers enterprise-class performance and proven interoperability optimized for modern virtualized and consolidated environments. When deployed with Fibre Channel NVMe (FC-NVMe), the platform provides consistently low latency, high throughput, and predictable performance required for business-critical workloads.

FC-NVMe on the VSP One Block 85 storage system leverages existing Fibre Channel investments while unlocking next-generation NVMe performance, offering a seamless migration path from SCSI-based FC environments without operational disruption. This approach preserves the reliability, determinism, and operational familiarity of Fibre Channel networks while significantly improving I/O efficiency and scalability.

Tightly integrated with the VMware ecosystems, VSP One Block 85 supports broad compatibility across leading server, HBA, and switch vendors, simplifying deployment and lifecycle management in heterogeneous data centers. Organizations benefit from Hitachi's industry-leading data availability, built-in cyber resilience, and energy-efficient architecture – enabling secure, resilient, and cost-effective modernization of their storage infrastructure.

Key benefits include:

- Certified VMware interoperability ensures predictable and reliable FC-NVMe deployments
- Ultra-low latency and high throughput enabled by FC-NVMe for performance-intensive workloads
- Seamless reuse of existing Fibre Channel infrastructure, reducing migration risk and cost
- Deterministic performance and high availability inherent to Fibre Channel architecture
- Broad compatibility with leading servers, HBAs, and Fibre Channel switches
- Deep VMware integration for simplified operations and lifecycle management
- Enterprise-class data availability with built-in cyber resilience
- Energy-efficient, eco-conscious design supporting sustainability and cost-control objectives.

For detailed, component-level compatibility information, see the Product Compatibility Guide (PCG).
<https://compatibility.hitachivantara.com/products/interop-matrix>

Document Revisions

Revision Number	Date	Details
v1.0	May 2026	Initial release

Comments

Send any comments on this document to Docs-Feedback@hitachivantara.com. Include the document title, including the revision level, and refer to specific sections and paragraphs whenever possible. All comments become the property of Hitachi Vantara LLC.

Overview

This guide provides a practical introduction to configuring NVMe over Fibre Channel (NVMe/FC) on VMware ESXi 8.0 Update 3. It is designed to help administrators understand both the conceptual foundations and the hands-on steps required to deploy NVMe over Fabrics (NVMe-oF) in a modern vSphere environment.

The document begins with an overview of NVMe/FC architecture, key concepts, and system requirements. It then walks through a step-by-step configuration process for enabling and validating NVMe/FC connectivity on ESXi hosts.

Additionally, this guide demonstrates how to integrate a Hitachi VSP One Block 85 (B85) storage array with VMware ESXi, including detailed procedures for:

- Configuring NVMe subsystems
- Setting port parameters
- Creating and managing namespaces
- Using Hitachi RAID Manager for storage configuration
- Leveraging the VMware High Performance Plugin (HPP) for optimal performance

By following this guide, administrators will gain the foundational knowledge and practical skills needed to deploy NVMe-oF solutions reliably and efficiently within a vSphere data center.

Key Components

- VMware ESXi systems (version 8.0 Update 3)
- VSP One Block High End (B85) storage system
- Fiber Channel (FC) SAN Infrastructure
- Hitachi RAID Manager

Connectivity Diagram

For an NVMe over Fibre Channel network, you must install an NVMe over Fibre Channel adapter on the ESXi host to connect to Fibre Channel Storage using the NVMe protocol. In this document, an Emulex FC-NVMe capable HBA has been used to connect to VSP One Block High End storage through two FC switches.

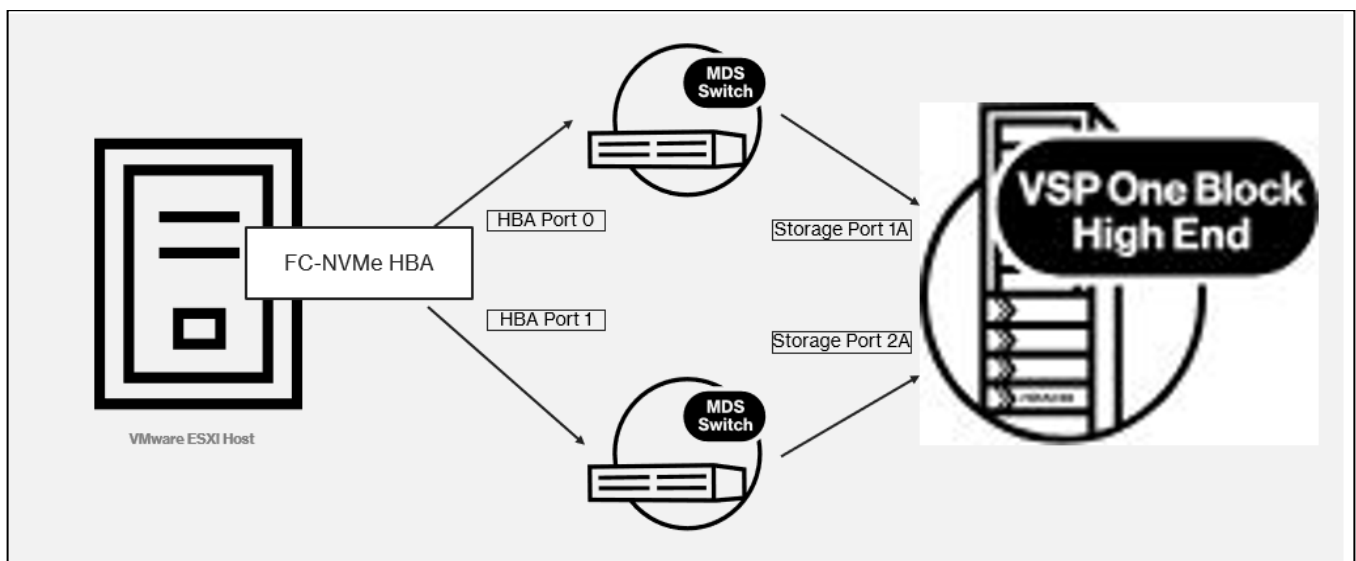


Figure 1 Port 1A and Port 2A of a VSP One Block 85 storage system connected to ESXi host through two Fibre Channel switches

VMware ESXi Systems and Multipathing

NVMe on ESXi Host

The rapid adoption of all-flash storage and latency-sensitive workloads has exposed the performance limitations of traditional SCSI-based storage protocols. To address these challenges, NVMe (Non-Volatile Memory Express) was developed as a modern storage access protocol optimized for high-speed, low-latency solid-state media. NVMe significantly improves throughput, reduces I/O latency, and scales more efficiently than legacy interfaces such as SATA or SAS, making it ideal for enterprise and cloud-scale environments.

In an ESXi environment, the host communicates with NVMe storage using a compatible PCIe storage adapter or Fibre Channel HBA. After it has been installed and properly configured, the ESXi host automatically discovers available NVMe namespaces, which appear as storage devices in the vSphere Client. These devices can then be formatted with VMFS datastores or used for other supported storage purposes.

Successful deployment of NVMe-oF in ESXi depends on several critical factors, including:

- ESXi version and update level compatibility
- Certified NVMe-oF-capable HBAs
- Supported drivers and firmware levels
- Proper fabric zoning and configuration

When these prerequisites are met, NVMe-oF enables ESXi hosts to achieve near-local NVMe performance while benefiting from the scalability and resiliency of shared storage.

ESXi Multipathing

In VMware ESXi environments, multipathing plays a critical role in delivering high availability and efficient use of storage resources by maintaining multiple physical paths between the host and storage targets. In recent ESXi releases, this functionality is enhanced by the High-Performance Plug-in (HPP), which serves as the default multipathing plug-in for NVMe and select high-speed SCSI devices.

- The High-Performance Plug-in (HPP) in ESXi is part of the Pluggable Storage Architecture (PSA) and is specifically designed to handle modern high-speed storage devices such as NVMe.
- Within PSA, ESXi uses claim rules to assign a multipathing plug-in (such as NMP or HPP) to each storage device, and after HPP claims a device, it becomes responsible for path management, failover, and load balancing across all available paths.
- HPP introduces more advanced and flexible path selection schemes compared to legacy plug-ins, allowing administrators to control behavior using parameters such as IOPS limits, bandwidth (bytes), and latency evaluation intervals through ESXCLI, enabling dynamic and performance-aware path switching rather than static policies. It supports multiple policies such as FIXED and load-balancing strategies (for example, IOPS-based).
- For NVMe and NVMe-oF, HPP works closely with NVMe multipathing concepts such as Asymmetric Namespace Access (ANA). It understands optimized versus non-optimized paths and prioritizes I/O intelligently while maintaining fast, seamless failover. This makes HPP particularly effective for NVMe over Fibre Channel and other low-latency transports where parallelism and rapid path decisions are critical.

For additional information on the VMware High Performance Plug-in, see: [Broadcom Technical Documentation](#)

The following is a list of `esxcli` commands that you can use to get more information on the ESXi host for NVMe devices:

- To list the paths currently claimed by HPP:

```
[root@localhost:/vmfs/volumes/61271846-fd2d484c-10d9-d4f5ef4f1008] esxcli storage hpp path list
```

```
fc.200034800d6f9624:210034800d6f9624-fc.50060e8008753832:50060e8008753832-
eui.00507538000000000060e80875380b00

Runtime Name: vmhba64:C0:T0:L0
Device: eui.00507538000000000060e80875380b00
Device Display Name: NVMe Fibre Channel Disk (eui.00507538000000000060e80875380b00)
Path State: dead
Path Config:
```

- To list the devices currently controlled by HPP:

```
[root@hag3-61200:~] esxcli storage hpp device list
eui.009111a5000000000060e83411a50e11
Device Display Name: NVMe Fibre Channel Disk (eui.009111a5000000000060e83411a50e11)
Path Selection Scheme: LB-RR
Path Selection Scheme Config: {iops=1000,bytes=10485760;}
Current Path: vmhba64:C0:T0:L10
Working Path Set: vmhba64:C0:T0:L10, vmhba65:C0:T0:L10
Is SSD: true
Is Local: false
Paths: vmhba64:C0:T0:L10, vmhba65:C0:T0:L10
Use ANO: false
```

VSP One Block High End Storage System

VSP One Block High-End (VSP One BHE) is Hitachi Vantara's flagship enterprise storage platform, designed to deliver uncompromising performance, scalability, and resiliency for mission-critical workloads. Built for large-scale, highly consolidated environments, VSP One BHE provides a unified block storage architecture that supports modern NVMe and NVMe-over-Fabrics deployments while maintaining the reliability and data services expected from a high-end enterprise array. With advanced capabilities such as end-to-end NVMe support, active-active controller architecture, and built-in data protection features, VSP One BHE enables organizations to run latency-sensitive applications, large databases, and virtualization platforms with consistent performance and maximum availability. Its tight integration with VMware vSphere and support for standards-based protocols make it a robust foundation for modern data centers adopting high-performance storage technologies such as FC-NVMe.

Starting from SVOS 10.5.3, VSP One BHE supports NVMe over Fibre Channel networks delivering comprehensive, enterprise-class NVMe-oF support designed to maximize performance, scalability, and availability for mission-critical workloads.

Hardware Requirements

The following lists the hardware used in this setup:

- A Fibre Channel storage system that supports NVMe, in this guide, VSP One BHE is used.
- A compatible ESXi host. For the operating system and server support, see: <https://compatibilityguide.broadcom.com/>
- A hardware NVMe adapter (Fibre Channel HBA that supports NVMe). For the Emulex HBA support and driver download, see: <https://www.broadcom.com/support/emulex> and [Broadcom | VMware | Hardware Compatibility Guide](#)
- An NVMe controller. After the required hardware NVMe adapter is installed, it automatically connects to all targets and controllers that are reachable at that moment.

Software Requirements

The following lists the software used in this setup:

- Operating system install media: ESXi 8.0U3g or later: VMware ESXi 8.0 Update 3g Release Notes
- CCI Software: Command Control Interface version 01-87-00/06 or later versions.

CCI Configuration

The following are the details for the CCI server used for this illustration. For detailed information on installation, configuration, and CCI commands, see the Virtual Storage Platform Family CCI Installation and Configuration Guide on the [Hitachi Vantara Documentation Portal](#).

.CCI version 01-87-03/03 supports VSP One Block 85 storage system from microcode A0-05-21-00/00 or later.

This section describes how to configure CCI.

1. Update the CCI management server to the latest release

```
[root@ilab-cci ~]# raidqry -h
Model   : RAID-Manager/Linux/x64
Ver&Rev: 01-87-03/03
-----
```

2. Configure the HORCM configuration file for the NVMe storage, as shown in the following example:

```
[root@ilab-cci etc]# cat horcm7079.conf
HORCM_MON
#ip address      service          poll(10ms)      timeout(10ms)
127.0.0.1        horcm7079        1000             3000

HORCM_CMD
\\.\IPCMD-10.20.30.40-31001 \\.\IPCMD-10.20.30.41-31002
[root@ilab-cci etc]#
[root@local-boot ~]# horcmstart.sh 7079
starting HORCM inst 7079
HORCM inst 7079 starts successfully.
```

In this example, `horcm7079.conf` is the storage HORCM configuration file. 10.20.30.40 is the AMC-0 Service IP and 10.20.30.41 is the AMC-1 IP address.

3. Log in to the storage system by running the following command:

```
[root@local-boot ~]# raidcom -login <username> <password> -IH7079
[root@ilab-cci etc]# ps -ef | grep horcm
-
-
root          17581          1  0 Mar16 ?           00:01:38 horcmd_07079
```

```
root      1112614  174985  0 07:30 pts/16    00:00:00 grep --color=auto horcm
[root@ilab-cci etc]#
```

Storage Side Configuration

This section describes the process for configuring the VSP One BHE storage NVM Subsystem with RAID Manager (CCI). The configuration procedure for VSP One Block 85 is similar to that of the VSP One Block 20 series storage system. Currently, with SVOS 10.5.3, VSP One BHE supports FC-NVMe on its Gen2 CHBs that is, Prism CHBs only.

Configuring the Storage Port

This section describes how to configure the storage port.

1. Verify the current settings of the storage ports and remove any FC-SCSI settings that exist on the ports.
2. Verify the port mode using the RAID Manager command as follows:

```
[root@ilab-cci etc]# raidcom get port -key detail -IH7079
```

PORT	TYPE	ATTR	SPD	LPID	FAB	CONN	SSW	SL	Serial#	WWN	PHY_PORT
PORT_MODE	RSGID	SECURITY									
CL1-B SCSI	FIBRE	TAR 0 -	AUT	D9	Y	PtoP	Y	0	9****39****3	50060e803411a501	-
CL1-B SCSI	FIBRE	MCU 0 -	AUT	D9	Y	PtoP	Y	0	9****3	50060e803411a501	-
CL1-B SCSI	FIBRE	RCU 0 -	AUT	D9	Y	PtoP	Y	0	9****3	50060e803411a501	-
CL1-B SCSI	FIBRE	ELUN 0 -	AUT	D9	Y	PtoP	Y	0	9****3	50060e803411a501	-
CL1-D SCSI	FIBRE	TAR 0 -	AUT	D3	Y	PtoP	Y	0	9****3	50060e803411a503	-
CL1-E SCSI	FIBRE	TAR 0 -	AUT	B2	Y	PtoP	Y	0	9****3	50060e803411a504	-
CL1-F ISCSI	ISCSI	TAR 0 -	25G	00	N	UNKN	Y	0	9****3	-	-
CL1-F ISCSI	ISCSI	MCU 0 -	25G	00	N	UNKN	Y	0	9****3	-	-
CL1-F ISCSI	ISCSI	RCU 0 -	25G	00	N	UNKN	Y	0	9****3	-	-
CL1-F ISCSI	ISCSI	ELUN 0 -	25G	00	N	UNKN	Y	0	9****3	-	-
CL1-H NVME TCP	NVME_TCP	TAR 0 -	100G	00	N	UNKN	N	0	9****3	-	-
CL2-B SCSI	FIBRE	TAR 0 -	AUT	C3	Y	PtoP	Y	0	9****3	50060e803411a511	-
CL2-B SCSI	FIBRE	MCU 0 -	AUT	C3	Y	PtoP	Y	0	9****3	50060e803411a511	-

CL2-B	FIBRE	RCU	AUT	C3	Y	PtoP	Y	0	9****3	50060e803411a511	-
SCSI		0 -									
CL2-B	FIBRE	ELUN	AUT	C3	Y	PtoP	Y	0	9****3	50060e803411a511	-
SCSI		0 -									
CL2-D	FIBRE	TAR	AUT	B6	Y	PtoP	Y	0	9****3	50060e803411a513	-
SCSI		0 -									
CL2-E	FIBRE	TAR	AUT	98	Y	PtoP	Y	0	9****3	50060e803411a514	-
SCSI		0 -									
CL2-F	ISCSI	TAR	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL2-F	ISCSI	MCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL2-F	ISCSI	RCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL2-F	ISCSI	ELUN	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL2-H	NVME_TCP	TAR	100G	00	N	UNKN	N	0	9****3	-	-
NVME TCP		0 -									
CL3-B	FIBRE	TAR	AUT	D6	Y	PtoP	Y	0	9****3	50060e803411a521	-
SCSI		0 -									
CL3-B	FIBRE	MCU	AUT	D6	Y	PtoP	Y	0	9****3	50060e803411a521	-
SCSI		0 -									
CL3-B	FIBRE	RCU	AUT	D6	Y	PtoP	Y	0	9****3	50060e803411a521	-
SCSI		0 -									
CL3-B	FIBRE	ELUN	AUT	D6	Y	PtoP	Y	0	9****3	50060e803411a521	-
SCSI		0 -									
CL3-D	FIBRE	TAR	AUT	D2	Y	PtoP	Y	0	9****3	50060e803411a523	-
SCSI		0 -									
CL3-E	FIBRE	TAR	AUT	B1	Y	PtoP	Y	0	9****3	50060e803411a524	-
SCSI		0 -									
CL3-F	ISCSI	TAR	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL3-F	ISCSI	MCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL3-F	ISCSI	RCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL3-F	ISCSI	ELUN	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL3-H	NVME_TCP	TAR	100G	00	N	UNKN	N	0	9****3	-	-
NVME TCP		0 -									
CL4-B	FIBRE	TAR	AUT	BC	Y	PtoP	Y	0	9****3	50060e803411a531	-
SCSI		0 -									
CL4-D	FIBRE	TAR	AUT	B5	Y	PtoP	Y	0	9****3	50060e803411a533	-
SCSI		0 -									
CL4-E	FIBRE	TAR	AUT	97	Y	PtoP	Y	0	9****3	50060e803411a534	-
SCSI		0 -									

CL4-F	ISCSI	TAR	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL4-F	ISCSI	MCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL4-F	ISCSI	RCU	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL4-F	ISCSI	ELUN	25G	00	N	UNKN	Y	0	9****3	-	-
ISCSI		0 -									
CL4-H	NVME_TCP	TAR	100G	00	N	UNKN	N	0	9****3	-	-
NVME TCP		0 -									
CL5-B	FIBRE	TAR	AUT	D5	Y	PtoP	Y	0	9****3	50060e803411a541	-
SCSI		0 -									
CL5-B	FIBRE	MCU	AUT	D5	Y	PtoP	Y	0	9****3	50060e803411a541	-
SCSI		0 -									
CL5-B	FIBRE	RCU	AUT	D5	Y	PtoP	Y	0	9****3	50060e803411a541	-
SCSI		0 -									
CL5-B	FIBRE	ELUN	AUT	D5	Y	PtoP	Y	0	9****3	50060e803411a541	-
SCSI		0 -									
CL5-D	FIBRE	TAR	AUT	D1	Y	PtoP	Y	0	9****3	50060e803411a543	-
SCSI		0 -									
CL5-E	FIBRE	TAR	AUT	AE	Y	PtoP	Y	0	9****3	50060e803411a544	-
SCSI		0 -									
CL6-B	FIBRE	TAR	AUT	BA	Y	PtoP	Y	0	9****3	50060e803411a551	-
SCSI		0 -									
CL6-B	FIBRE	MCU	AUT	BA	Y	PtoP	Y	0	9****3	50060e803411a551	-
SCSI		0 -									
CL6-B	FIBRE	RCU	AUT	BA	Y	PtoP	Y	0	9****3	50060e803411a551	-
SCSI		0 -									
CL6-B	FIBRE	ELUN	AUT	BA	Y	PtoP	Y	0	9****3	50060e803411a551	-
SCSI		0 -									
CL6-D	FIBRE	TAR	AUT	B4	Y	PtoP	Y	0	9****3	50060e803411a553	-
SCSI		0 -									
CL6-E	FIBRE	TAR	AUT	90	Y	PtoP	Y	0	9****3	50060e803411a554	-
SCSI		0 -									
CL7-B	FIBRE	TAR	AUT	D4	Y	PtoP	Y	0	9****3	50060e803411a561	-
SCSI		0 -									
CL7-B	FIBRE	MCU	AUT	D4	Y	PtoP	Y	0	9****3	50060e803411a561	-
SCSI		0 -									
CL7-B	FIBRE	RCU	AUT	D4	Y	PtoP	Y	0	9****3	50060e803411a561	-
SCSI		0 -									
CL7-B	FIBRE	ELUN	AUT	D4	Y	PtoP	Y	0	9****3	50060e803411a561	-
SCSI		0 -									
CL7-D	FIBRE	TAR	AUT	CE	Y	PtoP	Y	0	9****3	50060e803411a563	-
NVME		0 -									
CL7-E	FIBRE	TAR	AUT	AD	Y	PtoP	Y	0	9****3	50060e803411a564	-
SCSI		0 -									

```

CL8-B  FIBRE  TAR    AUT    B9  N    PtoP  Y    0    9****3  50060e803411a571  -
SCSI      0  -
CL8-B  FIBRE  MCU    AUT    B9  N    PtoP  Y    0    9****3  50060e803411a571  -
SCSI      0  -
CL8-B  FIBRE  RCU    AUT    B9  N    PtoP  Y    0    9****3  50060e803411a571  -
SCSI      0  -
CL8-B  FIBRE  ELUN   AUT    B9  N    PtoP  Y    0    9****3  50060e803411a571  -
SCSI      0  -
CL8-D  FIBRE  TAR    AUT    B3  Y    PtoP  Y    0    9****3  50060e803411a573  -
NVME      0  -
CL8-E  FIBRE  TAR    AUT    8F  Y    PtoP  Y    0    9****3  50060e803411a574  -
SCSI      0  -
[root@ilab-cci etc]#

```

3. Change the port mode of the Port CL5-E and Port CL6-E using the RAID Manager command as follows:

```

[root@ilab-cci etc]# raidcom modify port -port CL5-E -port_mode nvme -request_id auto -
IH7079
REQID : 121
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom modify port -port CL6-E -port_mode nvme -request_id auto -
IH7079
REQID : 122
[root@ilab-cci etc]#

```

4. Verify that the port mode is changed to NVMe as follows:

```

[root@ilab-cci etc]# raidcom get port -key detail -IH7079 | grep CL5-E
CL5-E  FIBRE  TAR    AUT    AE  Y    PtoP  Y    0    9****3  50060e803411a544  -
NVME      0  -
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get port -key detail -IH7079 | grep CL6-E
CL6-E  FIBRE  TAR    AUT    90  Y    PtoP  Y    0    9****3  50060e803411a554  -
NVME      0  -
[root@ilab-cci etc]#

```

5. Configure the topology settings.

- a. Set the port topology:

- For direct connectivity between the HBA and the storage port, set the port topology to Fabric OFF and **Point to Point**.
- For FC connectivity between the HBA and the storage port through the SAN switch, set the port topology to Fabric ON and **Point to Point**.

- b. Here both Port CL5-E and Port CL6-E are connected to storage through Fibre Channel switch, so, set port topology as Fabric ON and PtoP:

```
[root@ilab-cci etc]# raidcom get port -key detail -IH7079 | grep CL5-E
CL5-E  FIBRE TAR    AUT    AE  Y    PtoP  Y      0    9****3  50060e803411a544  -
NVME          0  -

[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get port -key detail -IH7079 | grep CL6-E
CL6-E  FIBRE TAR    AUT    90  Y    PtoP  Y      0    9****3  50060e803411a554  -
NVME          0  -

[root@ilab-cci etc]#
```

6. Configure the WWN security settings.

- a. Set the port security switch to enable and register the HBA WWN with the storage ports.
- b. Enable the WWN security switch for Port CL5-E and Port CL6-E:

```
[root@ilab-cci etc]#raidcom modify port -port CL5-E -security_switch y -I7079
[root@ilab-cci etc]#raidcom modify port -port CL6-E -security_switch y -I7079
```

- c. Verify whether host groups exist on the storage ports. If present, delete them.

```
[root@ilab-cci etc]# raidcom get host_grp -port CL5-E -IH7079
PORT  GID  GROUP_NAME                Serial#  HMD          HMO_BITS
CL5-E   0  5E-G00                    9****3  LINUX/IRIX
CL5-E   1  SN1700Q-TN                9****3  LINUX/IRIX   68

[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get host_grp -port CL6-E -IH7079
PORT  GID  GROUP_NAME                Serial#  HMD          HMO_BITS
CL6-E   0  6E-G00                    9****3  LINUX/IRIX
CL6-E   1  SN1700Q-TN                9****3  LINUX/IRIX   68

[root@ilab-cci etc]# raidcom delete host_grp -port CL5-E-1 -fx -IH7079
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom delete host_grp -port CL6-E-1 -fx -IH7079
[root@ilab-cci etc]#
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get host_grp -port CL6-E -IH7079
```

```

PORT   GID  GROUP_NAME                Serial# HMD          HMO_BITS
CL6-E   0   6E-G00                    9****3 LINUX/IRIX
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get host_grp -port CL5-E -IH7079
PORT   GID  GROUP_NAME                Serial# HMD          HMO_BITS
CL5-E   0   5E-G00                    9****3 LINUX/IRIX
[root@ilab-cci etc]#

```

- d. Check the WWN of the HBA from the host or the Fibre Channel switch and register the HBA WWN on these ports to allow access with the RAID Manager:

```

[root@ilab-cci etc]# raidcom add hba_wnn -port CL6-E -hba_wnn 10005cba2cfcae45 -IH7079
[root@ilab-cci etc]# raidcom add hba_wnn -port CL5-E -hba_wnn 10005cba2cfcae46 -IH7079

```

- e. Verify the HBA WWN with Port CL5-E and Port CL6-E:

```

[root@ilab-cci etc]# raidcom get hba_wnn -port CL6-E -IH7079
PORT   GID  GROUP_NAME                HWWN   Serial# NICK_NAME
CL6-E   0   6E-G00                    10005cba2cfcae45  9****3  -
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get hba_wnn -port CL5-E -IH7079
PORT   GID  GROUP_NAME                HWWN   Serial# NICK_NAME
CL5-E   0   5E-G00                    10005cba2cfcae46  9****3  -
[root@ilab-cci etc]#

```

Configuring the NVM Subsystem

Creating an NVM Subsystem

1. Verify the existing NVM subsystem information of the storage using the following command:

```
[root@ilab-cci etc]# raidcom get nvm_subsystem -key opt -IH7079
[root@ilab-cci etc]#
```

2. Create an NVM Subsystem on the storage system using the RAID Manager command, based on the operating system host mode option. First, verify whether any NVM Subsystems already exist, and then create a new NVM Subsystem using an available ID.

In the following example, the NVM Subsystem TN-B85-32G-V2-FC-NVMe (with host mode as VMWARE_EX) is created with NVM Subsystem ID 1:

```
[root@ilab-cci etc]# raidcom add nvm_subsystem -nvm_subsystem_id 1 -nvm_subsystem_name TN-
B85-32G-V2-FC-NVMe -host_mode VMWARE_EX -request_id auto -IH7079
REQID : 123
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get nvm_subsystem -key opt -IH7079
NVMSS_ID  NVMSS_NAME                NVMSS_NQN
          1  TN-B85-32G-V2-FC-NVMe      nqn.1994-04.jp.co.hitachi:nvme:storage-subsystem-sn.9-
70053-nvmssid.00001
[root@ilab-cci etc]#
```

3. Verify the details of the newly created NVM Subsystem ID 1 as follows:

```
[root@ilab-cci etc]# raidcom get nvm_subsystem -nvm_subsystem_id 1 -IH7079
NVMSS_ID  RGID  NVMSS_NAME                SECURITY T10PI    HMD          HMO_BITS
V_NVMSS_ID
          1    0  TN-B85-32G-V2-FC-NVMe      ENABLE  DISABLE  VMWARE EX   -
1
[root@ilab-cci etc]#
```

4. Note the NQN of NVM Subsystem ID 1 for future use.

```
nqn.1994-04.jp.co.hitachi:nvme:storage-subsystem-sn.9-70053-nvmssid.00001
```

Registering the NVM Subsystem Port

1. Register the Storage FC Ports that are connected to the ESXi hosts with the NVM Subsystem.

In the following example, Port CL5-E and Port CL6-E are registered with NVM Subsystem ID 1.

```
[root@ilab-cci etc]# raidcom add nvm_subsystem_port -nvm_subsystem_id 1 -port CL6-E -
request_id auto -IH7079
REQID : 124
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom add nvm_subsystem_port -nvm_subsystem_id 1 -port CL5-E -
request_id auto -IH7079
REQID : 125
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get nvm_subsystem_port -nvm_subsystem_id 1 -IH7079
PORT    NVMS_ID  NVMS_NAME
CL5-E      1  TN-B85-32G-V2-FC-NVMe
CL6-E      1  TN-B85-32G-V2-FC-NVMe
[root@ilab-cci etc]#
```

Configuring the Namespace Security Settings

1. Verify the Host NQN and note it for the Host NQN settings in the NVM Subsystem.

Host NQN for VMware host:

```
esxcli nvme info get
nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
```

5. Set the Host NQN to allow access to the NVM Subsystem using the RAID Manager command as follows:

```
[root@ilab-cci etc]# raidcom get host_nqn -nvm_subsystem_id 1 -IH7079
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom add host_nqn -nvm_subsystem_id 1 -host_nqn nqn.2017-
01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae45 -request_id auto -IH7079
REQID : 126
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom add host_nqn -nvm_subsystem_id 1 -host_nqn nqn.2017-
01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46 -request_id auto -IH7079
REQID : 127
[root@ilab-cci etc]#
[root@ilab-cci etc]# raidcom get host_nqn -nvm_subsystem_id 1 -IH7079
NVMS_ID  NVMS_NAME                HOST_NQN
      1  TN-B85-32G-V2-FC-NVMe    nqn.2017-
01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
      1  TN-B85-32G-V2-FC-NVMe    nqn.2017-
01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae45
```

```
[root@ilab-cci etc]#
```

Creating a Namespace by Registering LDEV

After registering the LDEV with the NVM Subsystem, a namespace is created and assigned a Namespace ID

1. In VSP One Block 85 storage systems, creating an LDEV (DRS Volume) from a pool is a mandatory step for storage provisioning.
2. Use the following command to create a DRS Volume from an HDP pool with capacity saving disabled:

```
raidcom add ldev -pool <pool_id> -ldev_id <virtual_volume_id> -capacity
<virtual_volume_capacity> -capacity_saving disable -drs -request_id auto -
IH<horcm_instance>
```

6. Use the following command to create a DRS volume from an HDP pool with capacity saving and compression acceleration enabled:

```
raidcom add ldev -pool <pool_id> -ldev_id <virtual_volume_id> -capacity
<virtual_volume_capacity> -capacity_saving deduplication_compression -
compression_acceleration enable -drs -request_id auto -IH<horcm_instance>
```

In the following example, four LDEVs (ID 0E:02 to 0E:05) are created with a capacity of 10 GB from pool ID 7, and then formatted.

7. Create LDEVs by running the following command:

```
raidcom add ldev -pool 7 -ldev_id 0E:02 -capacity 10G -capacity_saving
deduplication_compression -compression_acceleration enable -drs -request_id auto -I7079
raidcom add ldev -pool 7 -ldev_id 0E:03 -capacity 10G -capacity_saving
deduplication_compression -compression_acceleration enable -drs -request_id auto -I7079
raidcom add ldev -pool 7 -ldev_id 0E:04 -capacity 10G -capacity_saving
deduplication_compression -compression_acceleration enable -drs -request_id auto -I7079
raidcom add ldev -pool 7 -ldev_id 0E:05 -capacity 10G -capacity_saving
deduplication_compression -compression_acceleration enable -drs -request_id auto -I7079
```

8. Format LDEVs by running the following command:

```
raidcom initialize ldev -operation fmt -ldev_id 0E:02 -IH7079

raidcom initialize ldev -operation fmt -ldev_id 0E:03 -IH7079
raidcom initialize ldev -operation fmt -ldev_id 0E:04 -IH7079
raidcom initialize ldev -operation fmt -ldev_id 0E:05 -IH7079
```

For more information about LDEV specifications and latest storage provisioning recommendations, see the *Storage Provision Guide* for the relevant version on the [Hitachi Vantara Documentation Portal](#).

9. Create namespaces from the newly created LDEVs and register them with the NVM Subsystem ID, as follows:

```
[root@ilab-cci etc]# raidcom add namespace -nvm_subsystem_id 1 -ns_id auto -ldev_id 0E:02
-request_id auto -IH7079
REQID : 13a

[root@ilab-cci etc]# raidcom add namespace -nvm_subsystem_id 1 -ns_id auto -ldev_id 0E:03
-request_id auto -IH7079
REQID : 13b

[root@ilab-cci etc]# raidcom add namespace -nvm_subsystem_id 1 -ns_id auto -ldev_id 0E:04
-request_id auto -IH7079
REQID : 13c

[root@ilab-cci etc]# raidcom add namespace -nvm_subsystem_id 1 -ns_id auto -ldev_id 0E:05
-request_id auto -IH7079
REQID : 13d
```

10. Verify that the namespaces are created for the LDEVs in the NVM Subsystem:

```
[root@ilab-cci etc]# raidcom get namespace -nvm_subsystem_id 1 -IH7079
NVMSS_ID  NVMSS_NAME                NSID  LDEVID  CAPACITY(BLK)
      1  TN-B85-32G-V2-FC-NVMe        1    3585    524288000
      1  TN-B85-32G-V2-FC-NVMe        2    3586    20971520
      1  TN-B85-32G-V2-FC-NVMe        3    3587    20971520
      1  TN-B85-32G-V2-FC-NVMe        4    3588    20971520
      1  TN-B85-32G-V2-FC-NVMe        5    3589    20971520
```

11. Configure the namespace path settings.
12. Register the host NQN with the namespace ID to allow ESXi host access to the namespace.
13. Set the host NQN-namespace path using the RAID Manager command, as follows:

```
[root@ilab-cci etc]# raidcom add namespace_path -nvm_subsystem_id 1 -ns_id 2 -host_nqn
nqn.2014-08.com.vmware:nvme:hag3-61200 -request_id auto -IH7079
REQID : 144

[root@ilab-cci etc]# raidcom add namespace_path -nvm_subsystem_id 1 -ns_id 3 -host_nqn
nqn.2014-08.com.vmware:nvme:hag3-61200 -request_id auto -IH7079
REQID : 145

[root@ilab-cci etc]# raidcom add namespace_path -nvm_subsystem_id 1 -ns_id 4 -host_nqn
nqn.2014-08.com.vmware:nvme:hag3-61200 -request_id auto -IH7079
REQID : 146

[root@ilab-cci etc]# raidcom add namespace_path -nvm_subsystem_id 1 -ns_id 5 -host_nqn
nqn.2014-08.com.vmware:nvme:hag3-61200 -request_id auto -IH7079
```

```
REQID : 147
```

14. Verify the namespace path information using the following command:

```
[root@ilab-cci etc]# raidcom get namespace_path -nvm_subsystem_id 1 -IH7079
NVMS_ID  NVMS_NAME                NSID  LDEV#  HOST_NQN
      1  TN-B85-32G-V2-FC-NVMe        1    3585  nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
      1  TN-B85-32G-V2-FC-NVMe        1    3585  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        2    3586  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        3    3587  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        4    3588  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        5    3589  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        6    3590  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        7    3591  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        8    3592  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        9    3593  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe       10    3600  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe       11    3601  nqn.2014-08.com.vmware:nvme:hag3-61200
      1  TN-B85-32G-V2-FC-NVMe        1    3585  nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae45
[root@ilab-cci etc]#
```

15. Verify the host NQN login status on the storage ports:

```
[root@ilab-cci etc]# raidcom get port -port CL5-E -key login_host_nqn -IH7079
PORT      LOGIN_STATUS  HOST_NQN
CL5-E     LOGOUT        nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
CL5-E     LOGIN         nqn.2014-08.com.vmware:nvme:hag3-61200
[root@ilab-cci etc]#
```

Note: The host NQN might initially appear with a LOGOUT status. Restarting the ESXi host will change the status to LOGIN if connectivity and configuration are correct. After restarting the host, the login status should appear as follows:

```
[root@ilab-cci etc]# raidcom get port -port CL5-E -key login_host_nqn -IH7079
PORT      LOGIN_STATUS  HOST_NQN
CL5-E     LOGIN         nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
CL5-E     LOGIN         nqn.2014-08.com.vmware:nvme:hag3-61200
[root@ilab-cci etc]#
```

Host-Side Configuration

The host-side configuration is carried out in two distinct phases. In the first phase, the ESXi host and its associated HBA ports must be configured to enable and support the NVMe protocol. This includes ensuring that the necessary NVMe settings and transport configurations are properly applied to the host. After this initial configuration is complete, the host generates and provides an NVMe Qualified Name (NQN).

This NQN is required to proceed with the storage-side configuration. Using this NQN, the required NVMe subsystems, host access permissions, and namespace mappings are configured on the storage array.

After completing the storage-side setup, the second phase of the host-side configuration is carried out to enable discovery of NVMe namespaces on the ESXi host.

Part 1

This section describes how to configure and prepare the ESXi server for FC-NVMe.

1. Complete the connectivity by installing an NVMe-supported dual-port FC HBA into a supported server. Both HBA ports are connected to a VSP One BHE storage system ports through the FC SAN switch. In this scenario, VMware ESXi 8.0 update 3 was installed on an NVMe-supported server, and an Emulex FC-HBA with an NVMe-supported adapter was installed on the ESXi host.
2. Install the NVMe-supported HBA driver on the ESXi 8.0 Update3 host. For Emulex HBA, the ESXi 8.0 update 3 bundled lpfc driver supports NVMe.
3. Configure Fibre Channel zoning with the HBA and storage port WWNs on the Fibre Channel switch.
4. Enable the NVMe module on the ESXi host:

To change the FC4 type to FCP and NVMe for an Emulex HBA, run the following command:

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

The valid values for type are:

- n 0 = Register FCP only (for LPe32000-series adapters); Register both FCP and NVMe (for LPe35000-series adapters).
- n 1 = Register FCP only.
- n 3 = Register both FCP and NVMe. This parameter requires a system reboot or quick boot.

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

```
[root@hag3-61200:~] esxcli system module parameters list -m lpfc | grep lpfc_enable_fc4_type
lpfc_enable_fc4_type          int          Defines what FC4 types are supported
[root@hag3-61200:~]
[root@hag3-61200:~] esxcli system module parameters set -p lpfc_enable_fc4_type=3 -m lpfc
[root@hag3-61200:~]
[root@hag3-61200:~] esxcli system module parameters list -m lpfc | grep lpfc_enable_fc4_type
lpfc_enable_fc4_type          int          3          Defines what FC4 types are supported
[root@hag3-61200:~]
```

5. To generate the NVMe Qualified Name (NQN) of the adapter, run the following command:

```
esxcfg-module -s 'vmknvme_hostnqn_format=0' vmknvme
```

FRestart the ESXi host for the changes to take effect.

6. Note the Host NQN information by running the following command:

```
esxcli nvme info get
```

The output will be similar to the following:

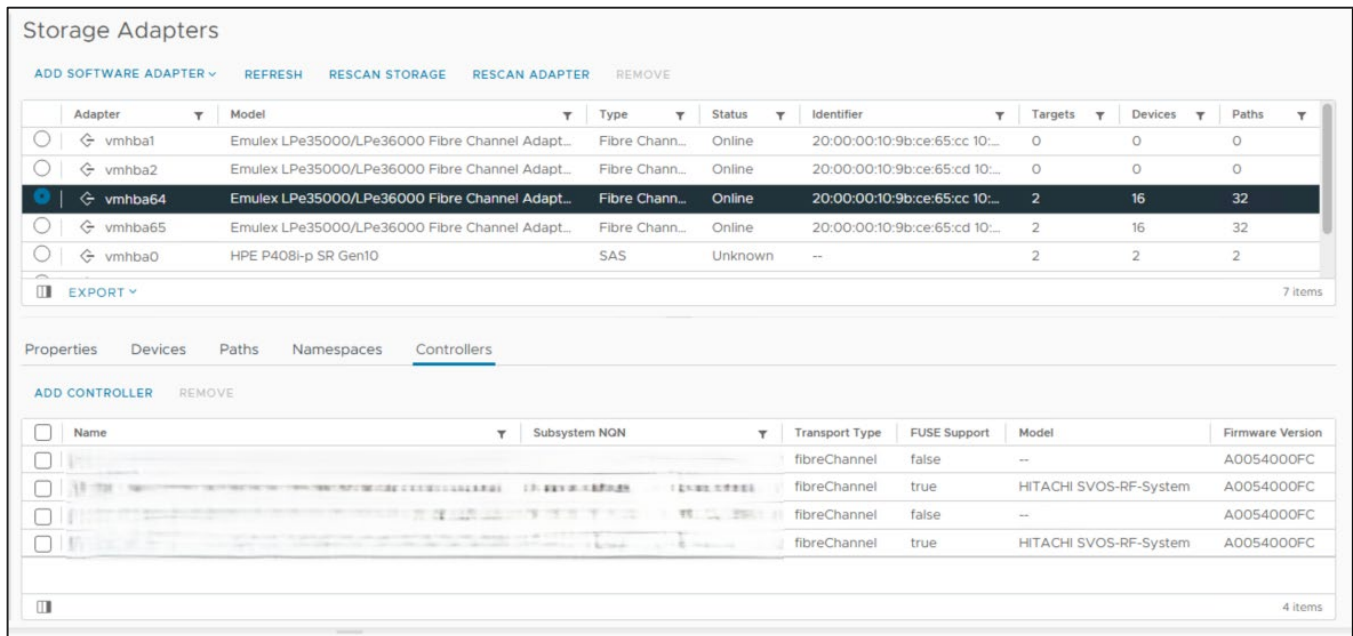
```
Host NQN:
nqn.2017-01.com.broadcom:ecd:nvmf:fc:10005cba2cfcae46
```

Note: The Host adapter NQN does not change with host restarts or driver updates; however, it changes with a new host OS installation.

The NQN uniquely identifies the host and is a mandatory input for performing the storage-side configuration. Using this NQN, the required NVMe subsystems, host access permissions, and namespace mappings are configured on the storage array.

Part 2

After completing the storage side configuration through CCI and restarting the host, the NVMe controller should appear in the storage adapter tab of the ESXi host. See the following illustrative example image.



After the NVMe controller is added to the ESXi host, the associated namespaces with that NVMe controller should appear as follows:

```
[root@hag3-61200:~] esxcfg-mpath -b
eui.009111a5000000000060e83411a50e11 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e11)
    vmhba64:C0:T0:L10 LUN:10 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54
    vmhba65:C0:T0:L10 LUN:10 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e01 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e01)
    vmhba64:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54
    vmhba65:C0:T0:L0 LUN:0 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e02 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e02)
    vmhba64:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54

    vmhba65:C0:T0:L1 LUN:1 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e03 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e03)
    vmhba64:C0:T0:L2 LUN:2 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54
    vmhba65:C0:T0:L2 LUN:2 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e04 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e04)
    vmhba64:C0:T0:L3 LUN:3 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54
    vmhba65:C0:T0:L3 LUN:3 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e05 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e05)
    vmhba64:C0:T0:L4 LUN:4 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPN: 50:06:0e:80:34:11:a5:54
```

```
vmhba65:C0:T0:L4 LUN:4 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e06 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e06)

vmhba64:C0:T0:L5 LUN:5 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPNN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPNN: 50:06:0e:80:34:11:a5:54

vmhba65:C0:T0:L5 LUN:5 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e07 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e07)

vmhba64:C0:T0:L6 LUN:6 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPNN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPNN: 50:06:0e:80:34:11:a5:54

vmhba65:C0:T0:L6 LUN:6 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e08 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e08)

vmhba64:C0:T0:L7 LUN:7 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPNN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPNN: 50:06:0e:80:34:11:a5:54

vmhba65:C0:T0:L7 LUN:7 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e09 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e09)

vmhba64:C0:T0:L8 LUN:8 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPNN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPNN: 50:06:0e:80:34:11:a5:54

vmhba65:C0:T0:L8 LUN:8 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44

eui.009111a5000000000060e83411a50e10 : NVMe Fibre Channel Disk
(eui.009111a5000000000060e83411a50e10)

vmhba64:C0:T0:L9 LUN:9 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:45 WWPNN:
10:00:5c:ba:2c:fc:ae:45 Target: WWNN: 50:06:0e:80:34:11:a5:54 WWPNN: 50:06:0e:80:34:11:a5:54

vmhba65:C0:T0:L9 LUN:9 state:active fc Adapter: WWNN: 20:00:5c:ba:2c:fc:ae:46 WWPNN:
10:00:5c:ba:2c:fc:ae:46 Target: WWNN: 50:06:0e:80:34:11:a5:44 WWPNN: 50:06:0e:80:34:11:a5:44
```