

# **Cisco and Hitachi Adaptive Solutions with Cisco UCS X-Series Modular System and Hitachi Virtual Storage Platform**

---

Reference Architecture Guide

© 2023 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 copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (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 LLC at [https://support.hitachivantara.com/en\\_us/contact-us.html](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 LLC.

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; and
2. Verifying that your 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.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z15, z16, z/VM, 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, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge 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 website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at <https://www.hitachivantara.com/en-us/company/legal.html> or [https://knowledge.hitachivantara.com/Documents/Open\\_Source\\_Software](https://knowledge.hitachivantara.com/Documents/Open_Source_Software).

## Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@HitachiVantara.com](mailto:SolutionLab@HitachiVantara.com). To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

### Revision history

Changes	Date
Added iSCSI information.	April 18, 2023
Removed iSCSI information.	March 30, 2023
Initial release.	March 6, 2023

---

## Reference Architecture Guide

This reference architecture introduces Cisco UCS X-Series into Cisco Hitachi Adaptive solutions and validates the seamless integration of the Hitachi Virtual Storage Platform (VSP) with VMware solutions. The Cisco and Hitachi Adaptive Solution is the latest converged infrastructure that is managed exclusively by Cisco Intersight. It is designed to meet the requirements of modern applications and improve operational efficiency, agility, and scale with its modular architecture.

With this modular architecture, you have the flexible capabilities of both blade and rack servers by offering compute density, storage capacity, and expandability in a single system, and this enables a greater range of workloads in your data center. Regardless of the workload, Cisco UCS X-Series with Intersight provides administrators ease of deployment and management across the entire UCS ecosystem by offering effortless modernization of your cloud operations.

Instantly consolidate any app onto an efficient modular, scalable, cloud-based infrastructure management for the entire stack. VSP storage systems truly enable a simplified approach to managing the datacenter. Cisco UCS X-Series backed by VSP provides customers a future-proof converged infrastructure stack backed by one of the most reliable enterprise storage systems which guarantees 100% data availability.

### Solution design

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure is a validated reference architecture targeting Virtual Server Infrastructure (VSI) implementations. The architecture is built around the Cisco UCS and the Hitachi VSP connected by Cisco MDS Multilayer SAN switches, and further enabled with Cisco Nexus switches. These components come together to form a powerful and scalable design, built on the best practices of both companies to create an ideal environment for virtualized systems.

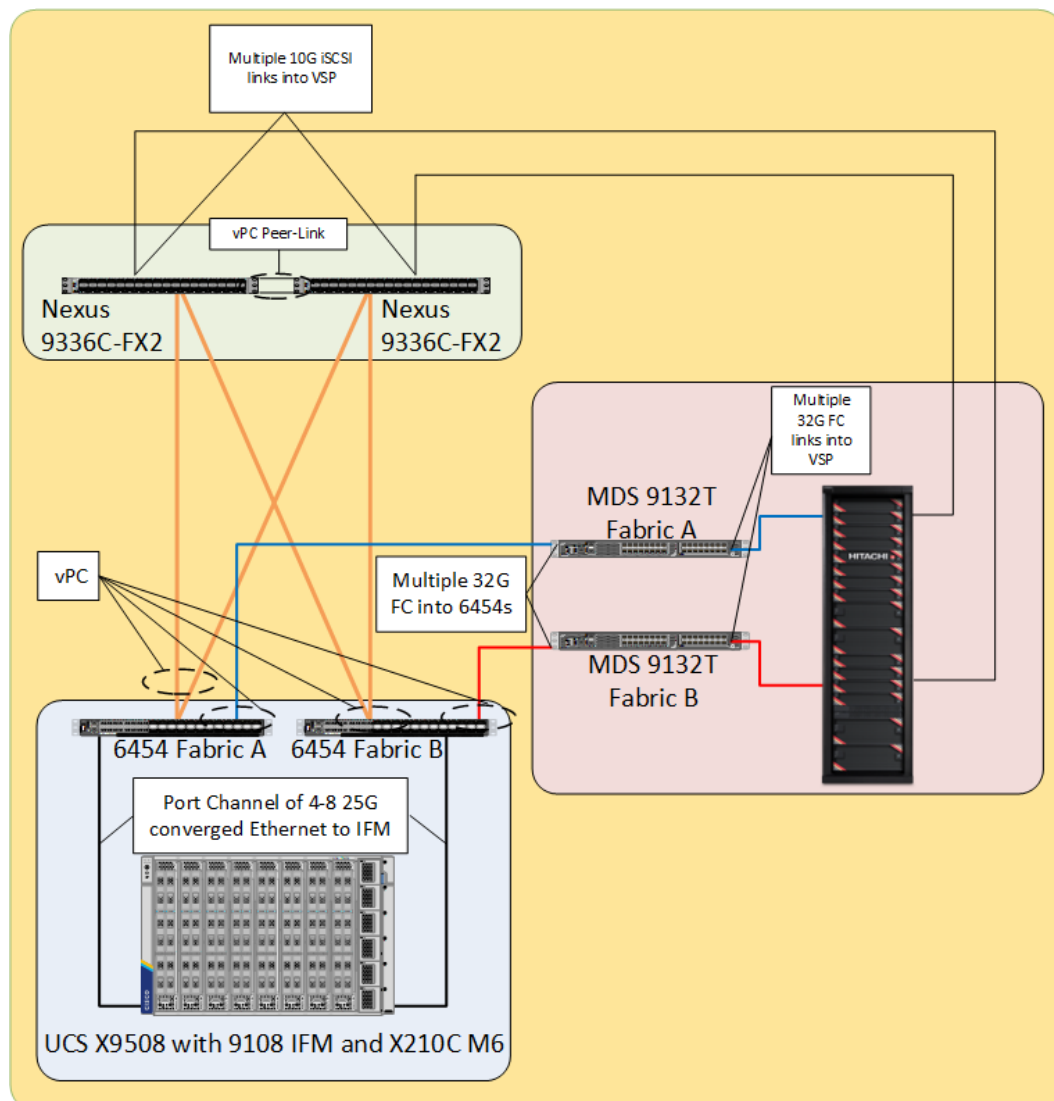
The solution is built and validated for a topology featuring a pair of Cisco UCS Fabric Interconnects as well as a Hitachi VSP storage system, with both using the same MDS and Nexus switching infrastructure.

The topology shown in the following figure leverages these products:

- Cisco Nexus 9336C-FX2 — 100 Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6454 Fabric Interconnect — Unified management of UCS compute, and the compute networks access to storage and networks.
- Cisco UCS x210c M6 Compute Node — High powered, versatile blade server, conceived for virtual computing.
- Cisco MDS 9132T — 32 Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- Hitachi VSP 5000 Series — Enterprise, high-performance storage system.



**Note:** Any component mentioned in this reference architecture guide can be supplemented with comparable products as long as they are listed in both Cisco and Hitachi hardware compatibility listings.



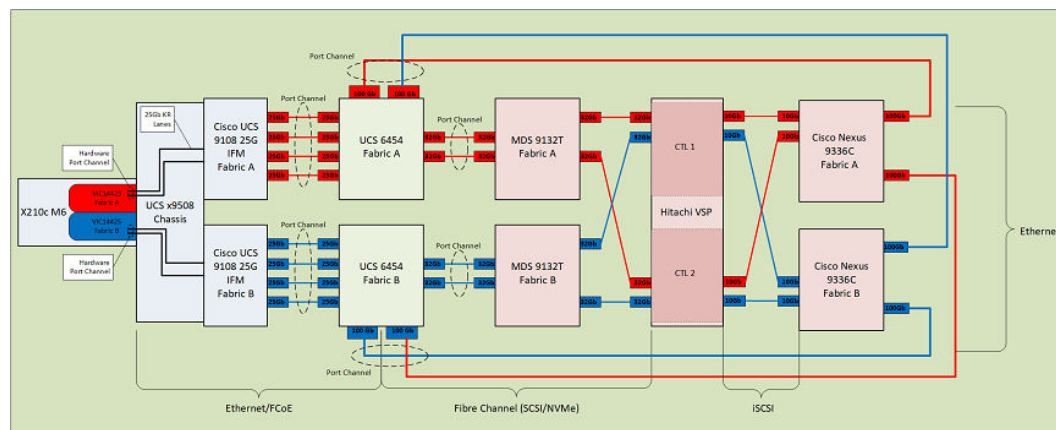
The Cisco UCS x210C M6 compute nodes in this topology are hosted within a Cisco UCS x9508 chassis and connect into the fabric interconnects (FIs) from the chassis using the Cisco UCS 9108 Intelligent Fabric Module (IFM). The 9108 IFM supports 25 G connections into the 10/25 G ports of the Cisco UCS 6454 fabric interconnects, delivering a high port availability that could fit well in a branch office setting.

Management components for this architecture also include:

- Cisco Intersight (required) — Comprehensive unified visibility across UCS domains as well as datacenter ecosystem in conjunction with automation workflows, along with proactive alerts and enablement of expedited Cisco technical assistance center (TAC) communications.
- Cisco Data Center Network Manager/Nexus Dashboard Fabric Controller (optional) — Multi-layer network configuration and monitoring.

The topology has been validated for vSphere 7.0 U3 to accommodate a larger range of expected customer deployments. Previous and newer versions of vSphere, as well as other vendor hypervisors, might be supported. These additional hypervisors must be within the compatibility and interoperability matrices listed at the Cisco <https://ucshcltool.cloudapps.cisco.com/public/> site as well as the Hitachi <https://compatibility.hitachivantara.com/> site.

The following illustration shows the end-to-end data path, where the Hitachi VSP is connected to the Cisco MDS 9132T at 32 Gb Fibre Channel speeds, in conjunction with being connected to the Cisco Nexus at 10 Gb speeds for iSCSI links. Both Fibre Channel and Ethernet connections offer coverage at the 6454 FIs and those connections are then port channeled into the Cisco UCS 9108 intelligent fabric modules to feed into the Cisco Virtual Interface Card 14425.



## Hardware and software versions

The following table lists the validated hardware and software versions used for this solution. Component and software version substitution from what is listed is considered acceptable within this reference architecture, but substitutions will need to comply with the hardware and software compatibility matrices from both Cisco and Hitachi.

- Cisco UCS Hardware Compatibility Matrix: <https://ucshcltool.cloudapps.cisco.com/public/>
- Cisco Nexus and MDS Interoperability Matrix: <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx/Matrix1.html>
- Cisco Nexus Recommended Releases for Nexus 9K: [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended\\_release/b\\_Minimum\\_and\\_Recommended\\_Cisco\\_NX-OS\\_Releases\\_for\\_Cisco\\_Nexus\\_9000\\_Series\\_Switches.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html)
- Cisco MDS Recommended Releases: [https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b\\_MDS\\_NX-OS\\_Recommended\\_Releases.html](https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html)
- Hitachi Vantara Interoperability: [https://support.hitachivantara.com/en\\_us/interoperability.html](https://support.hitachivantara.com/en_us/interoperability.html) sub-page -> (VSP 5x00, G1X00, F1500, E-series, Gxx0, Fxx0, VSP, HUS VM VMWare Support Matrix)

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	NX-OS 7.0(3)I7(9)
Compute	Cisco UCS Fabric Interconnect 6454	4.2(1m)
	Cisco UCS 9108 IFM	4.2(1j)
	Cisco UCS x210C M6	5.0(2b)
	VMware vSphere ESXi	VMware-ESXi-7.0.3-Custom-Cisco-4.2.1-a
	ESXi 7.0 U3 nenic	1.0.42.0-1
	ESXi 7.0 U3 nfnic	5.0.0.34-1
	VM Virtual Hardware Version	19
Storage	Hitachi VSP 5000 series	SVOS 90-08-61-00/00
	Cisco MDS 9132T	9.2(2)

## Physical cabling for the UCS 6454 with Hitachi Virtual Storage Platform

This section explains the cabling examples used for the validation of the topologies used in the lab environment. To make connectivity clear, the tables include both the local and remote port locations.

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. The upstream network from the Nexus 9336C-FX2 switches is out of scope of this document, with only the assumption that these switches will connect to the upstream switch or switches with a virtual Port Channel (vPC).



**Note:** Hitachi Virtual Storage Platform controller connections can either be iSCSI, FC-SCSI, or FC-NVMe based on environment needs. Always verify that correct Hitachi Virtual Storage Platform channel boards are used for controller connections.

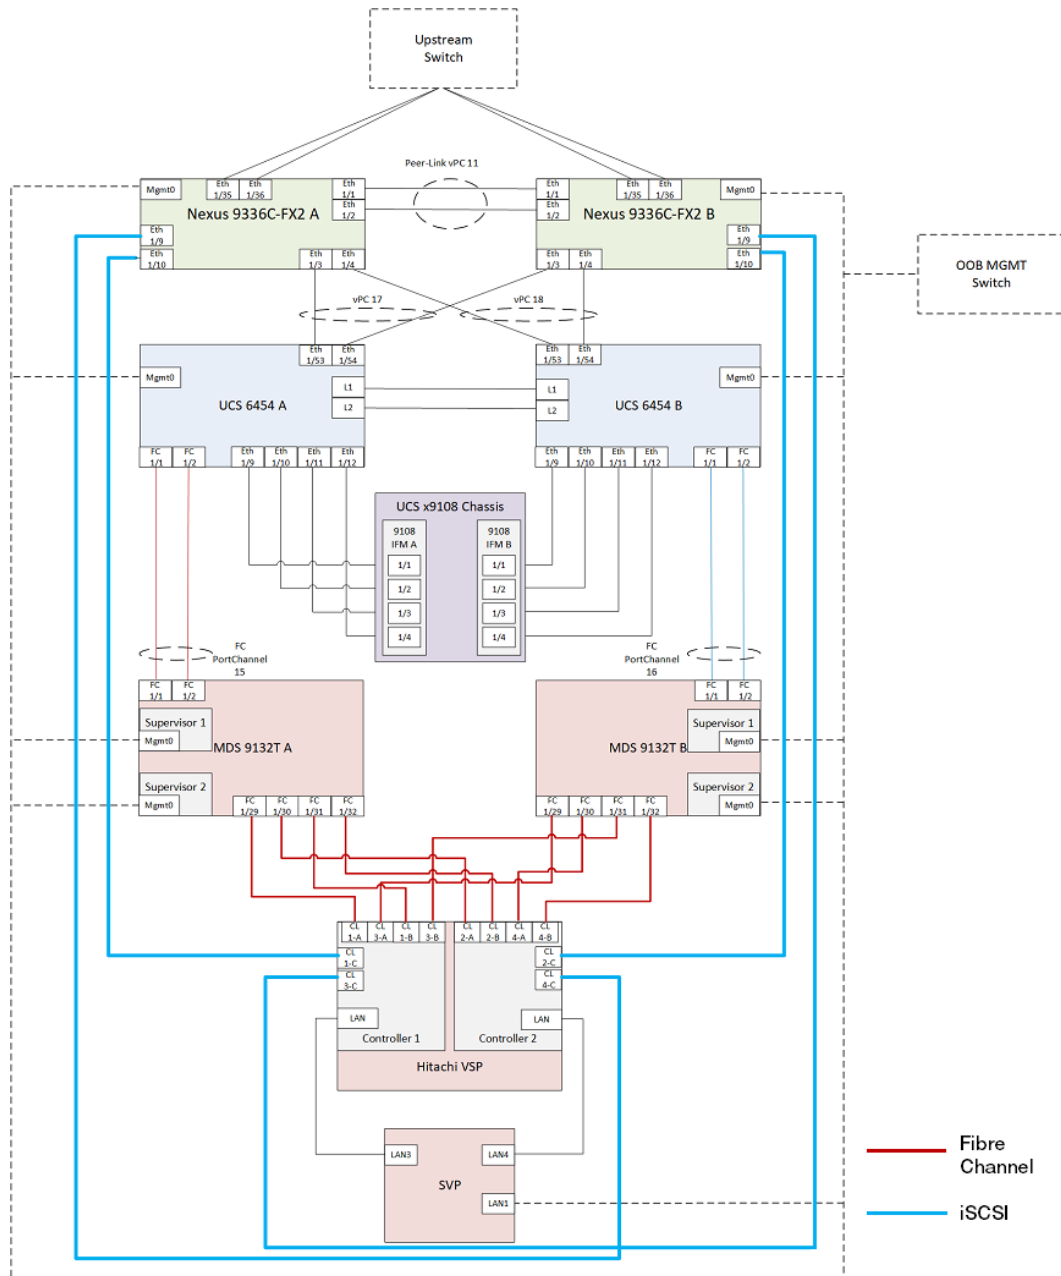


**Note:** When using the SaaS Intersight version the network segment that the system is on must have outbound Internet access.



## Physical cabling for the UCS 6454 with Hitachi Virtual Storage Platform

The following figure shows the cabling configuration used in the design featuring the Cisco UCS 6454 with Hitachi Virtual Storage Platform. Connections stemming from the VSP controller in red are Fibre Channel connections while blue indicates iSCSI.



The following tables list the specific port connections with the cables used in the deployment of the Cisco UCS 6454, Cisco UCS x9108 chassis, Cisco MDS 9132T, Cisco Nexus 9336C-FX2 and the Hitachi Virtual Storage Platform.

The following table lists Cisco Nexus 9336C-FX2 A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336C-FX2 A	Eth1/1	40 GbE	Cisco Nexus 9336C-FX2 B	Eth1/1
	Eth1/2	40 GbE	Cisco Nexus 9336C-FX2 B	Eth1/2
	Eth1/3	100 GbE	Cisco UCS 6454 FI A	Eth 1/53
	Eth1/4	100 GbE	Cisco UCS 6454 FI B	Eth 1/53
	Eth1/9	10 Gb	Hitachi VSP Controller 1 -iSCSI	CL4-C
	Eth1/10	10 Gb	Hitachi VSP Controller 2 -iSCSI	CL1-C
	Eth1/35	40 GbE or 100 GbE	Upstream Network Switch	Any
	Eth1/36	40 GbE or 100 GbE	Upstream Network Switch	Any
	MGMT0	GbE	GbE management switch	Any

The following table lists Cisco Nexus 9336C-FX2 B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9336C-FX2 B	Eth1/1	40 GbE	Cisco Nexus 9336C-FX2 A	Eth1/1
	Eth1/2	40 GbE	Cisco Nexus 9336C-FX2 A	Eth1/2
	Eth1/3	100 GbE	Cisco UCS 6454 FI A	Eth 1/54
	Eth1/4	100 GbE	Cisco UCS 6454 FI B	Eth 1/54
	Eth1/9	10 Gb	Hitachi VSP Controller 1 -iSCSI	CL3-C
	Eth1/10	10 Gb	Hitachi VSP Controller 2 -iSCSI	CL2-C
	Eth1/35	40GbE or 100GbE	Upstream Network Switch	Any
	Eth1/36	40GbE or 100GbE	Upstream Network Switch	Any
	MGMT0	GbE	GbE management switch	Any

The following table lists Cisco UCS 6454 A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6454 FI A	FC 1/1	32 Gb FC	MDS 9132T A	FC 1/1
	FC 1/2	32 Gb FC	MDS 9132T A	FC 1/2
	Eth1/9	25 GbE	Cisco UCS Chassis 9108 A	IFM 1/1
	Eth1/10	25 GbE	Cisco UCS Chassis 9108 A	IFM 1/2
	Eth1/11	25 GbE	Cisco UCS Chassis 9108 A	IFM 1/3
	Eth1/12	25 GbE	Cisco UCS Chassis 9108 A	IFM 1/4
	Eth1/53	40 GbE	Cisco Nexus 9336C-FX2 A	Eth1/3
	Eth1/54	40 GbE	Cisco Nexus 9336C-FX2 B	Eth1/3
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI B	L1
	L2	GbE	Cisco UCS 6454 FI B	L2

The following table lists: Cisco UCS 6454 B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS 6454 FI B	FC 1/1	32 Gb FC	MDS 9132T B	FC 1/1
	FC 1/2	32 Gb FC	MDS 9132T B	FC 1/2
	Eth1/9	10 GbE	Cisco UCS Chassis 9108 B	IFM 1/1
	Eth1/10	10 GbE	Cisco UCS Chassis 9108 B	IFM 1/2
	Eth1/11	10 GbE	Cisco UCS Chassis 9108 B	IFM 1/3
	Eth1/12	10 GbE	Cisco UCS Chassis 9108 B	IFM 1/4
	Eth1/53	40 GbE	Cisco Nexus 9336C-FX2 A	Eth1/4
	Eth1/54	40 GbE	Cisco Nexus 9336C-FX2 B	Eth1/4
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS 6454 FI A	L1
	L2	GbE	Cisco UCS 6454 FI A	L2

The following table lists Cisco MDS 9132T A cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T A	FC 1/1	32 Gb FC	Cisco UCS 6454 FI A	FC 1/1
	FC 1/1	32 Gb FC	Cisco UCS 6454 FI A	FC 1/2
	FC 1/29	32 Gb FC	Hitachi VSP Controller 1	CL1-A
	FC 1/30	32 Gb FC	Hitachi VSP Controller 2	CL2-A
	FC 1/31	32 Gb FC	Hitachi VSP Controller 1	CL1-B
	FC 1/32	32 Gb FC	Hitachi VSP Controller 2	CL2-B
	Sup1 MGMT0	GbE	GbE management switch	Any
	Sup2 MGMT0	GbE	GbE management switch	Any

The following table lists Cisco MDS 9132T B cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9132T B	FC 1/1	32 Gb FC	Cisco UCS 6454 FI B	FC 1/1
	FC 1/1	32 Gb FC	Cisco UCS 6454 FI B	FC 1/2
	FC 1/29	32 Gb FC	Hitachi VSP Controller 1	CL3-A
	FC 1/30	32 Gb FC	Hitachi VSP Controller 2	CL4-A
	FC 1/31	32 Gb FC	Hitachi VSP Controller 1	CL3-B
	FC 1/32	32 Gb FC	Hitachi VSP Controller 2	CL4-B
	Sup1 MGMT0	GbE	GbE management switch	Any
	Sup2 MGMT0	GbE	GbE management switch	Any

The following table lists Hitachi VSP cabling information.

Local Device	Local Port	Connection	Remote Device	Remote Port
Hitachi VSP	CL1-A	32 Gb FC	Cisco MDS 9132T A	FC 1/29
	CL2-A	32 Gb FC	Cisco MDS 9132T A	FC 1/30
	CL3-A	32 Gb FC	Cisco MDS 9132T B	FC 1/29
	CL4-A	32 Gb FC	Cisco MDS 9132T B	FC 1/30
	CL1-B	32 Gb FC	Cisco MDS 9132T A	FC 1/31
	CL2-B	32 Gb FC	Cisco MDS 9132T A	FC 1/32
	CL3-B	32 Gb FC	Cisco MDS 9132T B	FC 1/31
	CL4-B	32 Gb FC	Cisco MDS 9132T B	FC 1/32
	CL1-C	10 Gb iSCSI	Cisco Nexus 9336C-FX2 A	Eth 1/10
	CL2-C	10 Gb iSCSI	Cisco Nexus 9336C-FX2 B	Eth 1/10
	CL3-C	10 Gb iSCSI	Cisco Nexus 9336C-FX2 B	Eth 1/9
	CL4-C	10 Gb iSCSI	Cisco Nexus 9336C-FX2 A	Eth 1/9
	Cont1 LAN	GbE	SVP	LAN3
	Cont2 LAN	GbE	SVP	LAN4

## Configuration

This section describes configuration details for this reference architecture.

### Cisco Nexus switch configuration

The Nexus switch configuration explains the basic L2 and L3 functionality for the application environment used in the validation environment hosted by the UCS domains. The application gateways are hosted by the pair of Nexus switches, but primary routing is passed on to an existing router that is upstream of the converged infrastructure. This upstream router needs to be aware of any networks created on the Nexus switches, but configuration of an upstream router is beyond the scope of this guide.

See *Initial Nexus Configuration Dialogue* in [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_hitachi\\_adaptivesolutions\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html) for the Cisco Nexus 9336-FX2 configuration steps. Similar steps are required for other Cisco Nexus switches that may be used for a deployment. Additionally, see [Cisco Nexus iSCSI switch configuration \(on page 26\)](#) for information about Cisco Nexus iSCSI configuration. When making changes to the design, comply with the compatibility matrices of Cisco and Hitachi. Consult the configuration documents of the differing equipment to confirm the correct implementation steps.

## Cisco MDS configuration

See *Initial MDS Configuration Dialogue* in [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_hitachi\\_adaptivesolutions\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html) for the Cisco MDS 9132T configuration steps. The Cisco MDS 9706 was used, but similar steps are required for other Cisco MDS 9000 series switches that might be used for a deployment. When making changes to the design, comply with the compatibility matrices of Cisco and Hitachi. Consult the configuration documents of the differing equipment to confirm the correct implementation steps.

## Configure Fibre Channel ports on Hitachi VSP

For Hitachi VSP Fibre Channel ports to be exposed correctly to the MDS and Cisco UCS components, ports must be modified from their default values. Before beginning, ensure that you have credentials on the Hitachi Virtual Storage Platform that have at least the Administrator role permissions within Hitachi Storage Navigator. Your partner or Hitachi services personnel provide credentials to your Hitachi Virtual Storage Platform after initial setup and configuration of the storage system. See *Configuring Fibre Channel Ports on Hitachi Virtual Storage Platform* in [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_hitachi\\_adaptivesolutions\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html) for the configuration steps for Hitachi VSP Fibre Channel ports.



**Note:** For FC-NVMe deployments, port configuration must be done using CCI RAIDCOM. See [Appendix A Fibre Channel NVMe CCI RAIDCOM configuration \(on page 30\)](#) for additional commands for FC-NVMe configuration.

## Intersight Managed Mode (IMM) configuration for Cisco UCS

This section covers configuration requirements to manage Cisco UCS backed by Hitachi VSP using Intersight. See the Cisco UCS X-Series Quick Start Guide at <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x-series-quick-start-guide.html> for detailed configuration information.



## Fabric interconnects

The latest Cisco UCS chassis must use interconnects in IMM mode. This requires administrators to identify the deployment method from the CLI and define the configuration that allows claiming to Intersight.

```
Note: Intersight management mode setup available through console based
configuration method alone.

^C
Type 'reboot' to abort configuration and reboot system
or Type 'X' to cancel GUI configuratuion and go back to console or Pre
ss any other key to see the installation progress from GUI (reboot/X) ?

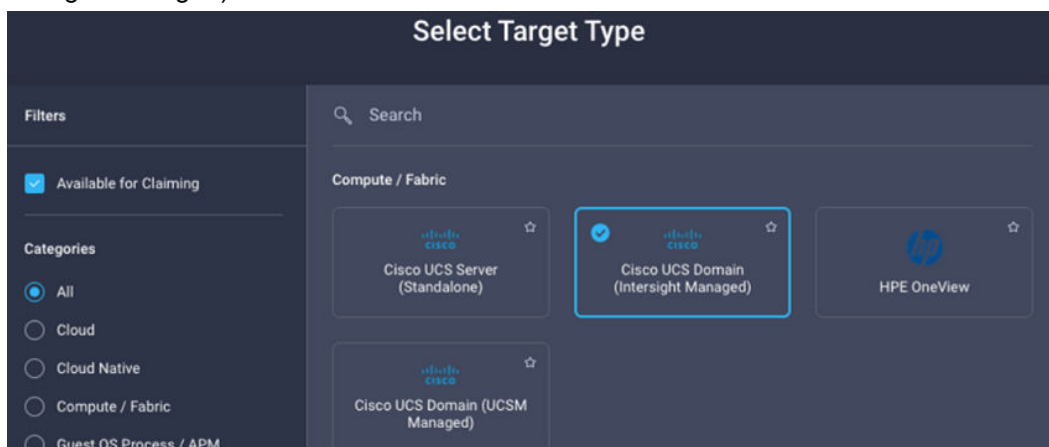
Enter the configuration method. (console/gui) ? console
Enter the management mode. (ucsm/intersight)? intersight

You have chosen to setup a new Fabric interconnect in "intersight" mana
ged mode. Continue? (y/n) y
Enforce strong password? (y/n) [y]: y
Enter the password for "admin":
Confirm the password for "admin": Internal CLI error: Invalid argument

Enter the switch fabric (A/B) [] A
Enter the system name: █
```

## Claim Fabric Interconnects to Intersight

After Fabric Interconnects are configured, administrators must claim them from Intersight.com. This allows all configurations to be made from Intersight and allows all profile and policy management to be pushed to your on-premises domain. When Fabric Interconnects are configured, a device connector tab is presented that provides a device ID and a claim code used to claim the domain from the target type Cisco UCS Domain (Intersight Managed).



## UCS domain profile

A domain profile is required to enable chassis discovery. Within a domain profile administrators configure both port channels for Nexus and MDS connections, as well as upstream VLANs for management, application, and vMotion they may have within their environment as well as port roles for server connections. Before configuration of a UCS domain profile, verify that configuration steps in [Cisco Nexus switch configuration \(on page 15\)](#) and [Cisco MDS configuration \(on page 16\)](#) have been completed.

The screenshot displays the 'Policies' section for 'Fabric Interconnect A' in the 'Configured' state. It shows the 'Port Configuration' tab with sub-tabs for 'General', 'Identifiers', and 'Connectivity'. The 'Port' configuration is set to 'F1\_A\_Port\_Policy'. Below this is a visual representation of the port channels and their members. At the bottom, there are four summary tables:

Port Type		Port Channel Type	
FC	8	FC Uplink	1
Ethernet	46	Ethernet Uplink	1

Port Role		Port Channel Role	
Server	2	FC Uplink	4
Unconfigured	46	Ethernet Uplink	2

## UCS chassis profile

A UCS chassis profile provides the ability for administrators to define power policies related to physical deployment, as well as a thermal policy that controls the degree of fan control among the units. The chassis profile also defines Intersight Managed Mode (IMC) access policies that define both in-band and out-of-band communication configuration, as well as SNMP resources that might interact across the domain with other products such as Nexus Dashboard Fabric Controller. Default values were used for Power, SNMP, and Thermal policies for the purpose of this document.

The screenshot shows the configuration page for 'SCHQ\_Chassis\_Profile' under 'UCS Chassis Profiles'. The status is 'OK'. The profile details are as follows:

Details	Details
Status: OK	IMC Access Policy: IMC
Name: SCHQ_Chassis_Profile	Power: Power_Default
Chassis: F16454xA-1	SNMP: SNMP_Default
Last Update: Apr 4, 2022 1:54 PM	Thermal: Thermal_Default
Description: -	
Organization: default	

## UCS server profile template

Create a server profile in the Cisco Intersight template format to allow seamless provisioning of server identities to physical nodes within a UCS environment. Use profile templates to define compute, management, and network configuration. See the Cisco UCS X-Series Quick Start Guide at <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x-series-quick-start-guide.html> for detailed configuration steps.

### Compute configuration

Within the compute configuration portion of the template, you can define the UUID pool, BIOS policy, boot order policy, and virtual media policy.



**Note:** The UUID pool uses 0000-540000000001 where 54 is the pattern for the domain that uses 6454 fabric interconnect.

#### BIOS policy

A BIOS policy automates the configuration of BIOS settings on managed devices. You can create one or more BIOS policies that contain a specific grouping of BIOS settings. If you do not specify a BIOS policy for a server, the BIOS settings remain as they are. If a BIOS policy is specified, the values in the policy replace any previously configured values on a server (including bare metal server configuration settings). You must restart the server to apply the BIOS policy settings.

#### Boot order policy

A boot order policy configures the linear ordering of devices and enables you to change the boot order and boot mode. You can also add multiple devices under various device types, rearrange the boot order, and set parameters for each boot device type. The inventory view lists the actual boot order configured on a server. The boot order displays details that include device name, device type, and configuration details such as Boot Mode (Legacy or UEFI) and Secure Boot Mode (Enabled or Disabled).

For SAN Boot device configuration in legacy mode, provide the boot target Logical Unit Number (LUN), device slot ID, interface name, and target WWPN. This information can be obtained from Hitachi VSP Storage Navigator. The following image represents a boot from SAN to a VSP target WWN and LUN ID.

The screenshot displays a configuration interface for SAN Boot. It features four distinct configuration sections, each for a different fabric and primary/secondary status. Each section includes a title bar with a collapse/expand icon, a status indicator (a green toggle switch labeled 'Enabled'), and a delete icon. The fields are as follows:

- SAN Boot (Fabric-A-Primary):** Device Name: Fabric-A-Primary; LUN: 0; Slot: (empty); Interface Name: vHBA-Fabric-A; Target WWPN: 50:06:0E:80:12:9C:CA:00.
- SAN Boot (Fabric-A-Secondary):** (Fields are not visible in the screenshot).
- SAN Boot (Fabric-B-Primary):** Device Name: Fabric-B-Primary; LUN: 0; Slot: (empty); Interface Name: vHBA-Fabric-B; Target WWPN: 50:06:0E:80:12:9C:CA:44.
- SAN Boot (Fabric-B-Secondary):** (Fields are not visible in the screenshot).

Following the best practices defined in *Cisco and Hitachi Adaptive Solution for Converged Infrastructure - Cisco Validated Design Guide* at [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_hitachi\\_adaptivesolutions\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html), there will be 4 targets within the boot from SAN configuration Fabric A (Primary and Secondary) as well as Fabric B (Primary and Secondary). When creating the configuration verify that the correct interface is selected based on the communication protocol used. Before boot device configuration your VSP storage system must be configured and online.

### Virtual media policy

A virtual media policy enables you to install an operating system on the server using the KVM console and virtual media, mount files to the host from a remote file share, and enable virtual media encryption. You can create one or more virtual media policies that contain virtual media mappings for different OS images, and configure up to two virtual media mappings, one for installation files, and the other for image files.

## Management configuration

Within the management configuration portion of the template, you can define IMC access, local users, SNMP, and virtual KVM.

## IMC access

An IMC access policy enables you to configure and manage your network by mapping IP pools to the chassis profile. Use this policy to configure a VLAN and associate it with an IP address using the IP address pool.

Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (FI-Attached) | UCS Chassis

A minimum of one configuration must be enabled. Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Check here for more info, [Help Centre](#)

In-Band Configuration  Enabled

VLAN ID \*  
29    
4 - 4093

IPv4 address configuration

IPv6 address configuration

IP Pool \*

## Local users

Local user policies automate the configuration of local user preferences. You can create one or more local user policies that contain a list of local users that need to be configured.

## Simple Network Management Protocol (SNMP)

An SNMP policy configures the SNMP settings for sending fault and alert information by SNMP traps from the managed devices. Any existing SNMP users or SNMP traps configured previously on managed devices are removed and replaced with users or traps that you configure in this policy. If you have not added any users or traps in the policy, the existing users or traps on the server are removed but not replaced.

## Virtual KVM

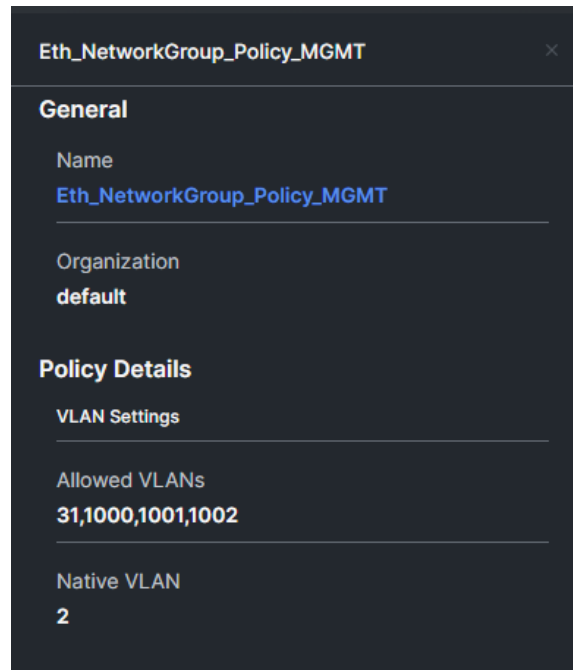
A virtual KVM policy enables specific grouping of virtual KVM properties. This policy specifies the number of allowed concurrent KVM sessions, port information, and video encryption options.

## Network configuration

LAN and SAN connectivity policies are defined within the network configuration portion of the template.

## LAN connectivity policy

A LAN connectivity policy is used to create an Ethernet network group policy, Ethernet network control policy, Ethernet adapter policy, and Ethernet QoS policy. Default values can be taken for the Ethernet network control policy, Ethernet adapter policy, and Ethernet QoS policies. For the Ethernet network group policy, users will create a policy per vNIC use case to pass the appropriate VLANs as defined on the ToR Nexus to the corresponding vNIC. For example, an Ethernet adapter group policy for application, and management will be required where, for example, the Ethernet adapter group policy for management passes VLAN 31 (OOB mgmt.), 1000 (vMotion), 1001 (iSCSI A), and 1002 (iSCSI B).



Alternatively, the other Ethernet network group policy for application will allow VLANs 30, along with any other VLANs planned for use within the ESXi distributed port group.

When you attach a LAN connectivity policy to a server profile, the addresses of the MAC address pool, or the static MAC address, are automatically assigned. MAC address pools can be created when configuring the LAN connectivity policy.



**Note:** For MAC pool configuration there will be a fabric A and B pool; the recommendation is to place 54 for the 4th octet to represent the 6454 FI being used. It is also recommended that 0A is placed in the 5th octet of the starting MAC address to identify all the MAC addresses as fabric A addresses, that is, 00:25:B5:54:0A:00. Similarly, 0B is placed next to the last octet to identify all MAC addresses as fabric B addresses, that is, 00:25:B5:54:0B:00.

Following best practices, there must be redundant connectivity for both A and B fabrics. The following image represents two virtual network interfaces per application: management, and iSCSI connectivity.

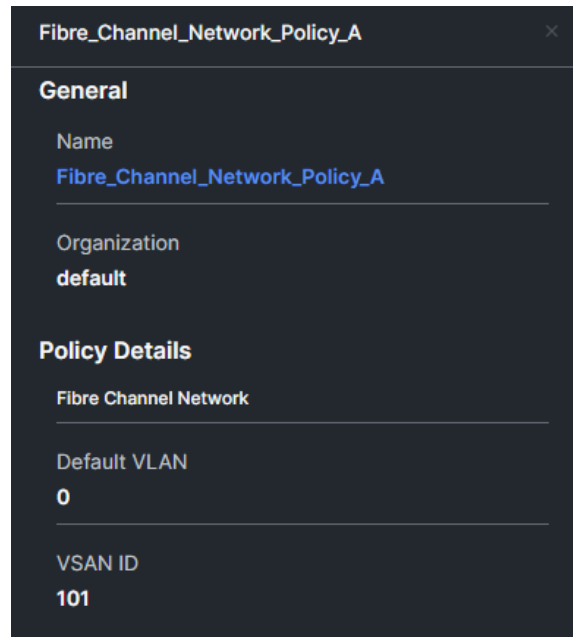
The screenshot displays the 'vNIC Configuration' interface. At the top, there are two tabs: 'Manual vNICs Placement' and 'Auto vNICs Placement'. Below the tabs, a blue banner contains a note: 'For auto placement option the vNICs will be automatically distributed between adaptors during profile deployment. Learn more at [Help Center](#)'. A blue 'Add vNIC' button is located below the banner. The main area features a table with columns for Name, Switch ID, Failover, Pin Group, and a lightning bolt icon. The table lists seven vNICs, each with a checkbox in the first column and a three-dot menu in the last column.

<input type="checkbox"/>	Name	Switch ID	Failover	Pin Group	⚡
<input type="checkbox"/>	vNIC-ISCSI-A	A	Enabled	-	...
<input type="checkbox"/>	vNIC00-APP-A	A	Enabled	-	...
<input type="checkbox"/>	vNIC00-MGMT-A	A	Enabled	-	...
<input type="checkbox"/>	vNIC-ISCSL_B	B	Enabled	-	...
<input type="checkbox"/>	vNIC01-APP-B	B	Enabled	-	...
<input type="checkbox"/>	vNIC01-MGMT-B	B	Enabled	-	...

After a server profile has been assigned, the MAC addresses of the vNICs can be verified from the server inventory.

## SAN connectivity policy

A SAN connectivity policy is used to create a Fibre Channel network policy, Fibre Channel adapter policy, and Fibre Channel QoS policy. Default values for the Fibre Channel adapter policy and Fibre Channel QoS policy can be used. For the Fibre Channel network policy, 2 policies must be created, one for fabric A and one for fabric B to identify the VSANs as defined in the Cisco MDS. The following example shows the Fibre Channel adapter policy for fabric A vHBAs.



**Note:** When configuring the Fibre Channel network policy, verify that you are assigning the correct VSAN to fabric A and fabric B vHBAs. That is, VSAN 101 for fabric A and VSAN 102 for fabric B.

When you attach a SAN Connectivity policy to a server profile, the addresses of the WWPN and WWNN pools, or the static WWPN and WWNN addresses, are automatically assigned. WWNN pools and WWPN pools can be created when creating the SAN connectivity policy.



**Note:** For WWNN the 6th octet was changed from 00 to 54 to represent identifying information for the 6454 Cisco UCS domain, that is, 20:00:00:25:B5:54:00:00.



**Note:** For WWPN two pools will be created, and the same pattern as the WWNN will be used. Additionally, the 7th octet will be 0A or 0B to identify fabrics A or B. That is, 20:00:00:25:B5:54:0A:00 and 20:00:00:25:B5:54:0B:00.

Following best practices, there must be redundant connectivity for both A and B fabrics. The following image represents two vHBA interfaces per FC-SCSI or FC-NVMe vHBA used.



**Note:** When creating vHBA for FC-SCSI, a vHBA type fc-initiator is used. Alternatively, if you are creating vHBA for FC-NVMe, a vHBA type fc-initiator-nvme is used.



Step 2  
**Policy Details**  
Add policy details

Manual vHBAs Placement    **Auto vHBAs Placement**

WWNN Address

**Pool**    Static

WWNN Address Pool \*

Selected Pool    WWNN\_Pool    |

For auto placement option the vHBAs will be automatically distributed between adaptors during profile deployment. Learn more at [Help Center](#)

**Add vHBA**

<input type="checkbox"/>	Name	Switch ID	Pin Group	
<input type="checkbox"/>	vHBA-Fabric-A	A	-	...
<input type="checkbox"/>	vHBA-NVMe-A	A	-	...
<input type="checkbox"/>	vHBA-NVMe-B	B	-	...
<input type="checkbox"/>	vHBA-Fabric-B	B	-	...

After a server profile has been assigned, the WWNs of vHBAs can be verified from the server inventory.

## Hitachi VSP storage configuration

After a server profile template is assigned to the respective chassis node, you must configure the boot from SAN configuration with Hitachi Virtual Storage Platform. See *Configure Host Connectivity and Presentation of Storage on Hitachi Virtual Storage Platform* in the *Cisco and Hitachi Adaptive Solution for Converged Infrastructure - Cisco Validated Design Guide* at [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_hitachi\\_adaptivesolutions\\_ci.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci.html) for information about how to configure VSP storage systems with VMware using best practices.

## Hitachi VSP iSCSI configuration

To use iSCSI connectivity with the Hitachi VSP backed by Cisco UCS you must confirm channel board availability and make respective connections into the top-of-rack Nexus switch. The following are the best practices to consider when enabling iSCSI with Cisco UCS backed by Hitachi VSP running VMware.

- If jumbo frames are required for iSCSI traffic, verify that all ports are configured with the correct MTU value (server, switches, and Hitachi VSP).
- On VMware ESXi, verify the correct MTU setting on vSphere virtual switches.
- Disable flow control features on the ports that are used.
- Verify that different iSCSI subnets do not transmit traffic between each other.

See *Configuring iSCSI ports* at [https://knowledge.hitachivantara.com/Documents/Management\\_Software/SVOS/9.2/Volume\\_Management\\_-\\_VSP\\_E990/Provisioning/15\\_Configuring\\_iSCSI\\_ports](https://knowledge.hitachivantara.com/Documents/Management_Software/SVOS/9.2/Volume_Management_-_VSP_E990/Provisioning/15_Configuring_iSCSI_ports) for more information about iSCSI configuration with Hitachi VSP. See *Configuring iSCSI and iSER Adapters and Storage* at <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.storage.doc/GUID-C476065E-C02F-47FA-A5F7-3B3F2FD40EA8.html> for more information about iSCSI configuration with VMware ESXi.

## Cisco Nexus iSCSI switch configuration

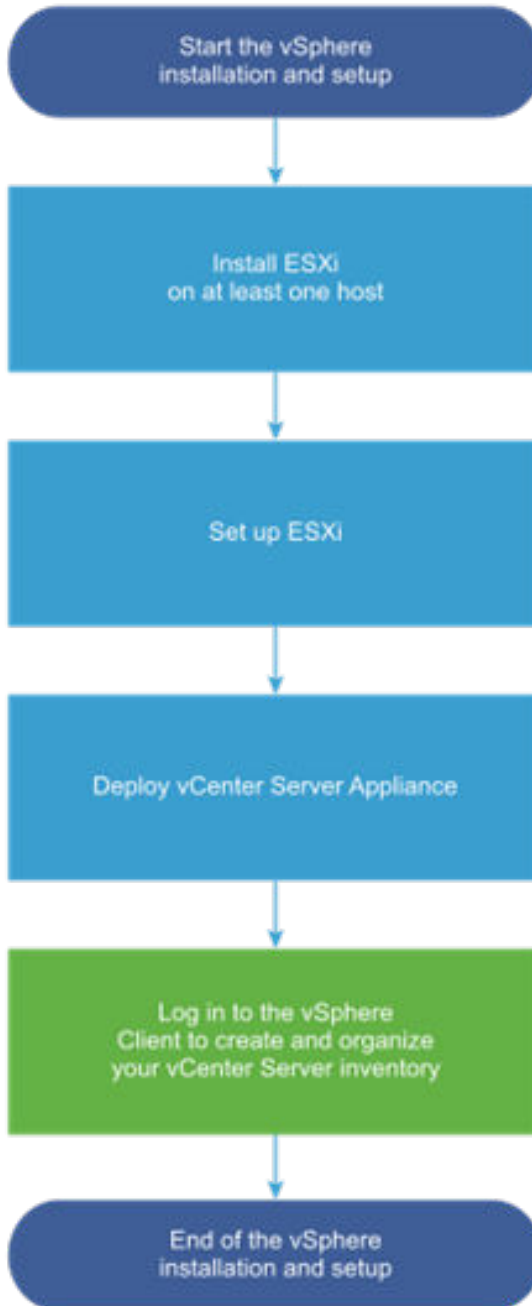
For iSCSI Cisco Nexus switch configuration provide dedicated VLANs for fabric A and B traffic respectively. The following configuration can be used when creating a dedicated VLAN for iSCSI connectivity on Cisco Nexus switches using jumbo frames. In this example VLAN 1001 and 1002 are used, respectively.

```
9336-A # configure
9336-A (config)# vlan 1002
9336-A (config-vlan)# name iSCSI_B
9336-A (config)# exit
9336-A (config)# interface ethernet 1/9
9336-A (config)#switchport access vlan 1002
9336-A (config)#spanning-tree port type edge
9336-A (config)#mtu 9216
9336-A (config)#no shutdown
9336-A (config)#exit
9336-B # configure
9336-B (config)# vlan 1001
9336-B (config-vlan)# name iSCSI_A
9336-B (config)# exit
9336-B (config)# interface ethernet 1/9
9336-B (config)#switchport access vlan 1001
9336-B (config)#spanning-tree port type edge
9336-B (config)#mtu 9216
9336-B (config)#no shutdown
9336-B (config)#exit
```

For any additional iSCSI connections from Hitachi VSP in Cisco Nexus, switch ports must be provided access to the appropriate VLAN.

## Operating system installation

After server profiles have been assigned, the ESXi operating system must be installed. The following image provides a workflow for deploying VMware ESXi and vCenter.



See *Install an operating system with software repository* as well as *Install an operating system with KVM-attached virtual media* within the *Cisco UCS X-Series Quick Start Guide* at <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x-series-quick-start-guide.html> for detailed procedures for mounting the ISO during installation.

### Download the Cisco VMware ESXi custom image

The VMware Cisco custom image will need to be downloaded for use during installation by manual access to the UCS KVM vMedia, or from the repository as explained in the previous section. The custom image can be downloaded from the VMware software repository at [https://customerconnect.vmware.com/en/downloads/info/slug/datacenter\\_cloud\\_infrastructure/vmware\\_vsphere/7\\_0](https://customerconnect.vmware.com/en/downloads/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/7_0).



**Note:** ESXi version 7.0U3 was used to validate this solution, but other ESXi operating systems can be used as long as listings comply with the hardware and software compatibility matrices from both Cisco and Hitachi as outlined in the Hardware and Software Versions section of this document.

See *VMware ESXi Installation and setup* in <https://docs.vmware.com/en/VMware-vSphere/index.html> for detailed ESXi configuration steps.

### VMware vCenter

VMware vCenter is used to manage ESXi hosts within the Cisco UCS domain, and users have the option to install the vCenter appliance within the UCS cluster or use a pre-existing infrastructure outside of the deployed converged infrastructure.



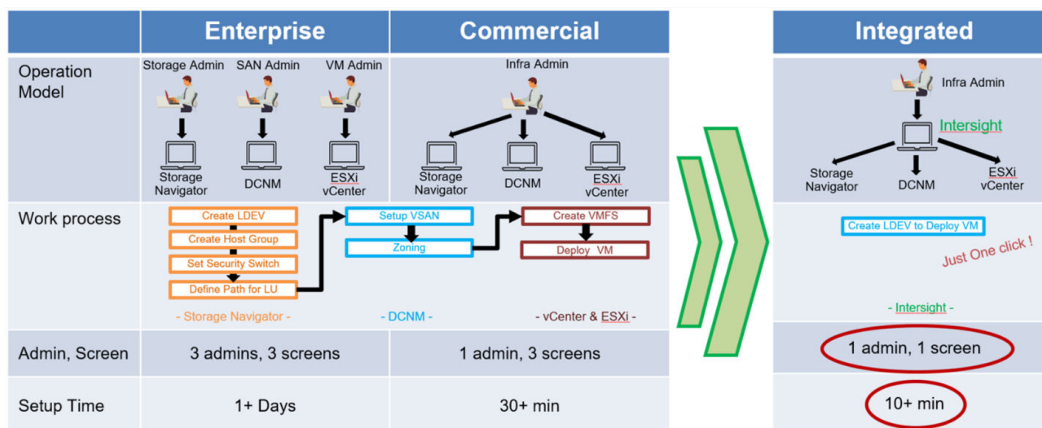
**Note:** The first ESXi host will need to go through the initial configuration using the VMware host client if a vCenter appliance is being installed to the UCS cluster. Subsequent hosts can be configured directly to the vCenter server after it is installed to the first ESXi host, or all hosts can be configured directly within the vCenter if a pre-existing server is used that is outside of the deployed converged infrastructure.

The VMware vCenter virtual appliance OVA can be download from the VMware customer portal at [https://customerconnect.vmware.com/en/downloads/info/slug/datacenter\\_cloud\\_infrastructure/vmware\\_vsphere/7\\_0](https://customerconnect.vmware.com/en/downloads/info/slug/datacenter_cloud_infrastructure/vmware_vsphere/7_0). After downloaded, vCenter server installation as well as configuration best practices can be found in VMware product documentation at <https://docs.vmware.com/en/VMware-vSphere/index.html>.

## Additional Hitachi VSP capabilities from Cisco Intersight

After the Cisco and Hitachi Adaptive Solution has been deployed, you can directly onboard the Hitachi VSP to Cisco Intersight for advanced single pane of glass management and automation. When the Hitachi VSP is onboarded, you will have access to all native VSP information as presented in Storage Navigator as well as the capability to use automation workflows from Intersight Cloud Orchestrator (ICO) to orchestrate end to end compute, SAN, and network configuration.

The following illustration shows the potential time savings for using Hitachi VSP and Cisco Intersight.



See the *Integrating Hitachi Virtual Storage Platform with Cisco Intersight Quick Start Guide* at <https://www.hitachivantara.com/en-us/pdf/architecture-guide/integrating-virtual-storage-platform-with-cisco-intersight.pdf> for step-by-step directions on how to onboard the Hitachi VSP to Cisco Intersight.

Additionally, see *Leveraging Hitachi Storage with Cisco Intersight for Consolidated Management and Automation Reference Architecture Guide* at <https://www.hitachivantara.com/en-us/pdf/architecture-guide/cisco-adaptive-solutions-leveraging-storage-with-cisco-intersight-for-consolidated-management-automation.pdf> for example automation workflows and processes.

---

## Appendix A: Appendix A Fibre Channel NVMe CCI RAIDCOM configuration

Before implementing this subset of directions, you must deploy a command control interface (CCI) server to issue the following commands. See the *Command Control Interface Installation and Configuration Guide* at [https://knowledge.hitachivantara.com/Documents/Management\\_Software/Command\\_Control\\_Interface](https://knowledge.hitachivantara.com/Documents/Management_Software/Command_Control_Interface) for more information. Additionally, see [Appendix B – Example HORCM file with virtual command device \(on page 48\)](#) for an example HORCM file.

Run the following commands to change VSP ports to NVMe mode on the VSP.

```
C:\HORCM\etc>raidcom modify port -port CL1-B -port_mode nvme -request_id auto
REQID : 134
C:\HORCM\etc>raidcom modify port -port CL2-B -port_mode nvme -request_id auto
REQID : 135
C:\HORCM\etc>raidcom modify port -port CL3-B -port_mode nvme -request_id auto
REQID : 136
C:\HORCM\etc>raidcom modify port -port CL4-B -port_mode nvme -request_id auto
REQID : 137
```

Run this command to verify that required VSP ports are in NVMe mode (the PHY\_PORT attribute will be NVMe).

```
C:\HORCM\etc>raidcom get port -key detail
PORT    TYPE  ATTR  SPD  LPID  FAB  CONN  SSW  SL  Serial#  WWN
PHY_PORT PORT_MODE
CL1-A   FIBRE TAR  AUT  EF   Y    PtoP Y     0   540016  50060e80089c5000
-       SCSI
CL1-B   FIBRE TAR  AUT  D9   Y    PtoP N     0   540016  50060e80089c5001
-       NVME
CL1-C   FIBRE TAR  AUT  E1   Y    PtoP Y     0   540016  50060e80089c5002
-       SCSI
CL1-D   FIBRE TAR  AUT  D3   Y    PtoP Y     0   540016  50060e80089c5003
-       SCSI
CL1-E   FIBRE TAR  AUT  B2   Y    PtoP N     0   540016  50060e80089c5004
-       SCSI
CL1-F   FIBRE TAR  AUT  A7   Y    PtoP N     0   540016  50060e80089c5005
-       SCSI
CL1-G   FIBRE TAR  AUT  AC   Y    PtoP N     0   540016  50060e80089c5006
-       SCSI
CL1-H   FIBRE TAR  AUT  9F   Y    PtoP N     0   540016  50060e80089c5007
-       SCSI
CL2-A   FIBRE TAR  AUT  CD   Y    PtoP Y     0   540016  50060e80089c5010
-       SCSI
CL2-B   FIBRE TAR  AUT  C3   Y    PtoP N     0   540016  50060e80089c5011
-       NVME
CL2-C   FIBRE TAR  AUT  C9   Y    PtoP Y     0   540016  50060e80089c5012
-       SCSI
CL2-D   FIBRE TAR  AUT  B6   Y    PtoP Y     0   540016  50060e80089c5013
-       SCSI
CL2-E   FIBRE TAR  AUT  98   Y    PtoP N     0   540016  50060e80089c5014
-       SCSI
CL2-F   FIBRE TAR  AUT  80   Y    PtoP N     0   540016  50060e80089c5015
-       SCSI
CL2-G   FIBRE TAR  AUT  88   Y    PtoP N     0   540016  50060e80089c5016
-       SCSI
CL2-H   FIBRE TAR  AUT  76   Y    PtoP N     0   540016  50060e80089c5017
-       SCSI
CL3-A   FIBRE TAR  AUT  E8   Y    PtoP Y     0   540016  50060e80089c5020
-       SCSI
CL3-B   FIBRE TAR  AUT  D6   Y    PtoP N     0   540016  50060e80089c5021
-       NVME
CL3-C   FIBRE TAR  AUT  E0   Y    PtoP Y     0   540016  50060e80089c5022
-       SCSI
CL3-D   FIBRE TAR  AUT  D2   Y    PtoP Y     0   540016  50060e80089c5023
-       SCSI
CL3-E   FIBRE TAR  AUT  B1   Y    PtoP N     0   540016  50060e80089c5024
-       SCSI
CL3-F   FIBRE TAR  AUT  A6   Y    PtoP N     0   540016  50060e80089c5025
-       SCSI
CL3-G   FIBRE TAR  AUT  AB   Y    PtoP N     0   540016  50060e80089c5026
```

-	SCSI									
CL3-H	FIBRE	TAR	AUT	9E	Y	PtoP	N	0	540016	50060e80089c5027
-	SCSI									
CL4-A	FIBRE	TAR	AUT	CC	Y	PtoP	Y	0	540016	50060e80089c5030
-	SCSI									
CL4-B	FIBRE	TAR	AUT	BC	Y	PtoP	N	0	540016	50060e80089c5031
-	NVME									
CL4-C	FIBRE	TAR	AUT	C7	Y	PtoP	Y	0	540016	50060e80089c5032
-	SCSI									
CL4-D	FIBRE	TAR	AUT	B5	Y	PtoP	Y	0	540016	50060e80089c5033
-	SCSI									
CL4-E	FIBRE	TAR	AUT	97	Y	PtoP	N	0	540016	50060e80089c5034
-	SCSI									
CL4-F	FIBRE	TAR	AUT	7C	Y	PtoP	N	0	540016	50060e80089c5035
-	SCSI									
CL4-G	FIBRE	TAR	AUT	84	Y	PtoP	N	0	540016	50060e80089c5036
-	SCSI									
CL4-H	FIBRE	TAR	AUT	75	Y	PtoP	N	0	540016	50060e80089c5037
-	SCSI									
CL5-A	FIBRE	TAR	AUT	E4	Y	PtoP	Y	0	540016	50060e80089c5040
-	SCSI									
CL5-B	FIBRE	TAR	AUT	D5	Y	PtoP	N	0	540016	50060e80089c5041
-	SCSI									
CL5-C	FIBRE	TAR	AUT	DC	Y	PtoP	N	0	540016	50060e80089c5042
-	SCSI									
CL5-D	FIBRE	TAR	AUT	D1	Y	PtoP	N	0	540016	50060e80089c5043
-	SCSI									
CL5-E	FIBRE	TAR	AUT	AE	Y	PtoP	N	0	540016	50060e80089c5044
-	SCSI									
CL5-F	FIBRE	TAR	AUT	A5	Y	PtoP	N	0	540016	50060e80089c5045
-	SCSI									
CL5-G	FIBRE	TAR	AUT	AA	Y	PtoP	N	0	540016	50060e80089c5046
-	SCSI									
CL5-H	FIBRE	TAR	AUT	9D	Y	PtoP	N	0	540016	50060e80089c5047
-	SCSI									
CL6-A	FIBRE	TAR	AUT	CB	Y	PtoP	N	0	540016	50060e80089c5050
-	SCSI									
CL6-B	FIBRE	TAR	AUT	BA	Y	PtoP	N	0	540016	50060e80089c5051
-	SCSI									
CL6-C	FIBRE	TAR	AUT	C6	Y	PtoP	N	0	540016	50060e80089c5052
-	SCSI									
CL6-D	FIBRE	TAR	AUT	B4	Y	PtoP	N	0	540016	50060e80089c5053
-	SCSI									
CL6-E	FIBRE	TAR	AUT	90	Y	PtoP	N	0	540016	50060e80089c5054
-	SCSI									
CL6-F	FIBRE	TAR	AUT	7A	Y	PtoP	N	0	540016	50060e80089c5055
-	SCSI									
CL6-G	FIBRE	TAR	AUT	82	Y	PtoP	N	0	540016	50060e80089c5056
-	SCSI									
CL6-H	FIBRE	TAR	AUT	74	Y	PtoP	N	0	540016	50060e80089c5057
-	SCSI									

Appendix A: Appendix A Fibre Channel NVMe CCI RAIDCOM configuration



CL7-A	FIBRE	TAR	AUT	E2	Y	PtoP	N	0	540016	50060e80089c5060
-	-	SCSI								
CL7-B	FIBRE	TAR	AUT	D4	Y	PtoP	N	0	540016	50060e80089c5061
-	-	SCSI								
CL7-C	FIBRE	TAR	AUT	DA	Y	PtoP	N	0	540016	50060e80089c5062
-	-	SCSI								
CL7-D	FIBRE	TAR	AUT	CE	Y	PtoP	N	0	540016	50060e80089c5063
-	-	SCSI								
CL7-E	FIBRE	TAR	AUT	AD	Y	PtoP	N	0	540016	50060e80089c5064
-	-	SCSI								
CL7-F	FIBRE	TAR	AUT	A3	Y	PtoP	N	0	540016	50060e80089c5065
-	-	SCSI								
CL7-G	FIBRE	TAR	AUT	A9	Y	PtoP	N	0	540016	50060e80089c5066
-	-	SCSI								
CL7-H	FIBRE	TAR	AUT	9B	Y	PtoP	N	0	540016	50060e80089c5067
-	-	SCSI								
CL8-A	FIBRE	TAR	AUT	CA	Y	PtoP	N	0	540016	50060e80089c5070
-	-	SCSI								
CL8-B	FIBRE	TAR	AUT	B9	Y	PtoP	N	0	540016	50060e80089c5071
-	-	SCSI								
CL8-C	FIBRE	TAR	AUT	C5	Y	PtoP	N	0	540016	50060e80089c5072
-	-	SCSI								
CL8-D	FIBRE	TAR	AUT	B3	Y	PtoP	N	0	540016	50060e80089c5073
-	-	SCSI								
CL8-E	FIBRE	TAR	AUT	8F	Y	PtoP	N	0	540016	50060e80089c5074
-	-	SCSI								
CL8-F	FIBRE	TAR	AUT	79	Y	PtoP	N	0	540016	50060e80089c5075
-	-	SCSI								
CL8-G	FIBRE	TAR	AUT	81	Y	PtoP	N	0	540016	50060e80089c5076
-	-	SCSI								
CL8-H	FIBRE	TAR	AUT	73	Y	PtoP	N	0	540016	50060e80089c5077
-	-	SCSI								

Run the following commands to enable security switch on VSP NVMe ports.

```
C:\HORCM\etc>raidcom modify port -port CL1-B -security_switch y
C:\HORCM\etc>raidcom modify port -port CL2-B -security_switch y
C:\HORCM\etc>raidcom modify port -port CL3-B -security_switch y
C:\HORCM\etc>raidcom modify port -port CL4-B -security_switch y
```

Run the following command to verify required VSP ports in NVMe have security switch enabled (the SSW attribute will be Y).

```
C:\HORCM\etc>raidcom get port -key detail
PORT    TYPE  ATTR  SPD  LPID  FAB  CONN  SSW  SL  Serial#  WWN
PHY_PORT PORT_MODE
CL1-A   FIBRE TAR  AUT  EF   Y    PtoP Y     0   540016  50060e80089c5000
-       SCSI
CL1-B   FIBRE TAR  AUT  D9   Y    PtoP Y     0   540016  50060e80089c5001
-       NVME
CL1-C   FIBRE TAR  AUT  E1   Y    PtoP Y     0   540016  50060e80089c5002
-       SCSI
CL1-D   FIBRE TAR  AUT  D3   Y    PtoP Y     0   540016  50060e80089c5003
-       SCSI
CL1-E   FIBRE TAR  AUT  B2   Y    PtoP N     0   540016  50060e80089c5004
-       SCSI
CL1-F   FIBRE TAR  AUT  A7   Y    PtoP N     0   540016  50060e80089c5005
-       SCSI
CL1-G   FIBRE TAR  AUT  AC   Y    PtoP N     0   540016  50060e80089c5006
-       SCSI
CL1-H   FIBRE TAR  AUT  9F   Y    PtoP N     0   540016  50060e80089c5007
-       SCSI
CL2-A   FIBRE TAR  AUT  CD   Y    PtoP Y     0   540016  50060e80089c5010
-       SCSI
CL2-B   FIBRE TAR  AUT  C3   Y    PtoP Y     0   540016  50060e80089c5011
-       NVME
CL2-C   FIBRE TAR  AUT  C9   Y    PtoP Y     0   540016  50060e80089c5012
-       SCSI
CL2-D   FIBRE TAR  AUT  B6   Y    PtoP Y     0   540016  50060e80089c5013
-       SCSI
CL2-E   FIBRE TAR  AUT  98   Y    PtoP N     0   540016  50060e80089c5014
-       SCSI
CL2-F   FIBRE TAR  AUT  80   Y    PtoP N     0   540016  50060e80089c5015
-       SCSI
CL2-G   FIBRE TAR  AUT  88   Y    PtoP N     0   540016  50060e80089c5016
-       SCSI
CL2-H   FIBRE TAR  AUT  76   Y    PtoP N     0   540016  50060e80089c5017
-       SCSI
CL3-A   FIBRE TAR  AUT  E8   Y    PtoP Y     0   540016  50060e80089c5020
-       SCSI
CL3-B   FIBRE TAR  AUT  D6   Y    PtoP Y     0   540016  50060e80089c5021
-       NVME
CL3-C   FIBRE TAR  AUT  E0   Y    PtoP Y     0   540016  50060e80089c5022
-       SCSI
CL3-D   FIBRE TAR  AUT  D2   Y    PtoP Y     0   540016  50060e80089c5023
-       SCSI
CL3-E   FIBRE TAR  AUT  B1   Y    PtoP N     0   540016  50060e80089c5024
-       SCSI
CL3-F   FIBRE TAR  AUT  A6   Y    PtoP N     0   540016  50060e80089c5025
-       SCSI
CL3-G   FIBRE TAR  AUT  AB   Y    PtoP N     0   540016  50060e80089c5026
```

-	SCSI									
CL3-H	FIBRE	TAR	AUT	9E	Y	PtoP	N	0	540016	50060e80089c5027
-	SCSI									
CL4-A	FIBRE	TAR	AUT	CC	Y	PtoP	Y	0	540016	50060e80089c5030
-	SCSI									
CL4-B	FIBRE	TAR	AUT	BC	Y	PtoP	Y	0	540016	50060e80089c5031
-	NVME									
CL4-C	FIBRE	TAR	AUT	C7	Y	PtoP	Y	0	540016	50060e80089c5032
-	SCSI									
CL4-D	FIBRE	TAR	AUT	B5	Y	PtoP	Y	0	540016	50060e80089c5033
-	SCSI									
CL4-E	FIBRE	TAR	AUT	97	Y	PtoP	N	0	540016	50060e80089c5034
-	SCSI									
CL4-F	FIBRE	TAR	AUT	7C	Y	PtoP	N	0	540016	50060e80089c5035
-	SCSI									
CL4-G	FIBRE	TAR	AUT	84	Y	PtoP	N	0	540016	50060e80089c5036
-	SCSI									
CL4-H	FIBRE	TAR	AUT	75	Y	PtoP	N	0	540016	50060e80089c5037
-	SCSI									
CL5-A	FIBRE	TAR	AUT	E4	Y	PtoP	Y	0	540016	50060e80089c5040
-	SCSI									
CL5-B	FIBRE	TAR	AUT	D5	Y	PtoP	N	0	540016	50060e80089c5041
-	SCSI									
CL5-C	FIBRE	TAR	AUT	DC	Y	PtoP	N	0	540016	50060e80089c5042
-	SCSI									
CL5-D	FIBRE	TAR	AUT	D1	Y	PtoP	N	0	540016	50060e80089c5043
-	SCSI									
CL5-E	FIBRE	TAR	AUT	AE	Y	PtoP	N	0	540016	50060e80089c5044
-	SCSI									
CL5-F	FIBRE	TAR	AUT	A5	Y	PtoP	N	0	540016	50060e80089c5045
-	SCSI									
CL5-G	FIBRE	TAR	AUT	AA	Y	PtoP	N	0	540016	50060e80089c5046
-	SCSI									
CL5-H	FIBRE	TAR	AUT	9D	Y	PtoP	N	0	540016	50060e80089c5047
-	SCSI									
CL6-A	FIBRE	TAR	AUT	CB	Y	PtoP	N	0	540016	50060e80089c5050
-	SCSI									
CL6-B	FIBRE	TAR	AUT	BA	Y	PtoP	N	0	540016	50060e80089c5051
-	SCSI									
CL6-C	FIBRE	TAR	AUT	C6	Y	PtoP	N	0	540016	50060e80089c5052
-	SCSI									
CL6-D	FIBRE	TAR	AUT	B4	Y	PtoP	N	0	540016	50060e80089c5053
-	SCSI									
CL6-E	FIBRE	TAR	AUT	90	Y	PtoP	N	0	540016	50060e80089c5054
-	SCSI									
CL6-F	FIBRE	TAR	AUT	7A	Y	PtoP	N	0	540016	50060e80089c5055
-	SCSI									
CL6-G	FIBRE	TAR	AUT	82	Y	PtoP	N	0	540016	50060e80089c5056
-	SCSI									
CL6-H	FIBRE	TAR	AUT	74	Y	PtoP	N	0	540016	50060e80089c5057
-	SCSI									

CL7-A	FIBRE	TAR	AUT	E2	Y	PtoP	N	0	540016	50060e80089c5060
-	SCSI									
CL7-B	FIBRE	TAR	AUT	D4	Y	PtoP	N	0	540016	50060e80089c5061
-	SCSI									
CL7-C	FIBRE	TAR	AUT	DA	Y	PtoP	N	0	540016	50060e80089c5062
-	SCSI									
CL7-D	FIBRE	TAR	AUT	CE	Y	PtoP	N	0	540016	50060e80089c5063
-	SCSI									
CL7-E	FIBRE	TAR	AUT	AD	Y	PtoP	N	0	540016	50060e80089c5064
-	SCSI									
CL7-F	FIBRE	TAR	AUT	A3	Y	PtoP	N	0	540016	50060e80089c5065
-	SCSI									
CL7-G	FIBRE	TAR	AUT	A9	Y	PtoP	N	0	540016	50060e80089c5066
-	SCSI									
CL7-H	FIBRE	TAR	AUT	9B	Y	PtoP	N	0	540016	50060e80089c5067
-	SCSI									
CL8-A	FIBRE	TAR	AUT	CA	Y	PtoP	N	0	540016	50060e80089c5070
-	SCSI									
CL8-B	FIBRE	TAR	AUT	B9	Y	PtoP	N	0	540016	50060e80089c5071
-	SCSI									
CL8-C	FIBRE	TAR	AUT	C5	Y	PtoP	N	0	540016	50060e80089c5072
-	SCSI									
CL8-D	FIBRE	TAR	AUT	B3	Y	PtoP	N	0	540016	50060e80089c5073
-	SCSI									
CL8-E	FIBRE	TAR	AUT	8F	Y	PtoP	N	0	540016	50060e80089c5074
-	SCSI									
CL8-F	FIBRE	TAR	AUT	79	Y	PtoP	N	0	540016	50060e80089c5075
-	SCSI									
CL8-G	FIBRE	TAR	AUT	81	Y	PtoP	N	0	540016	50060e80089c5076
-	SCSI									
CL8-H	FIBRE	TAR	AUT	73	Y	PtoP	N	0	540016	50060e80089c5077
-	SCSI									

Run the following commands to create NVMe host groups.

```
C:\HORCM\etc>raidcom add host_grp -port CL1-B-1 -host_grp_name VSI_5600-0_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL2-B-1 -host_grp_name VSI_5600-0_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL3-B-1 -host_grp_name VSI_5600-0_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL4-B-1 -host_grp_name VSI_5600-0_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL1-B-2 -host_grp_name VSI_5600-1_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL2-B-2 -host_grp_name VSI_5600-1_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL3-B-2 -host_grp_name VSI_5600-1_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL4-B-2 -host_grp_name VSI_5600-1_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL1-B-3 -host_grp_name VSI_5600-2_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL2-B-3 -host_grp_name VSI_5600-2_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL3-B-3 -host_grp_name VSI_5600-2_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL4-B-3 -host_grp_name VSI_5600-2_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL1-B-4 -host_grp_name VSI_5600-3_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL2-B-4 -host_grp_name VSI_5600-3_Fab_A
C:\HORCM\etc>raidcom add host_grp -port CL3-B-4 -host_grp_name VSI_5600-3_Fab_B
C:\HORCM\etc>raidcom add host_grp -port CL4-B-4 -host_grp_name VSI_5600-3_Fab_B
```

Run the following commands to verify that NVMe host groups have been created.

```
C:\HORCM\etc>raidcom get host_grp -port CL1-B
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL1-B  0    1B-G00                                   540016  LINUX/IRIX
CL1-B  1    VSI_5600-0_Fab_A                       540016  LINUX/IRIX
CL1-B  2    VSI_5600-1_Fab_A                       540016  LINUX/IRIX
CL1-B  3    VSI_5600-2_Fab_A                       540016  LINUX/IRIX
CL1-B  4    VSI_5600-3_Fab_A                       540016  LINUX/IRIX
C:\HORCM\etc>raidcom get host_grp -port CL2-B
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL2-B  0    2B-G00                                   540016  LINUX/IRIX
CL2-B  1    VSI_5600-0_Fab_A                       540016  LINUX/IRIX
CL2-B  2    VSI_5600-1_Fab_A                       540016  LINUX/IRIX
CL2-B  3    VSI_5600-2_Fab_A                       540016  LINUX/IRIX
CL2-B  4    VSI_5600-3_Fab_A                       540016  LINUX/IRIX
C:\HORCM\etc>raidcom get host_grp -port CL3-B
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL3-B  0    3B-G00                                   540016  LINUX/IRIX
CL3-B  1    VSI_5600-0_Fab_B                       540016  LINUX/IRIX
CL3-B  2    VSI_5600-1_Fab_B                       540016  LINUX/IRIX
CL3-B  3    VSI_5600-2_Fab_B                       540016  LINUX/IRIX
CL3-B  4    VSI_5600-3_Fab_B                       540016  LINUX/IRIX
C:\HORCM\etc>raidcom get host_grp -port CL4-B
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL4-B  0    4B-G00                                   540016  LINUX/IRIX
CL4-B  1    VSI_5600-0_Fab_B                       540016  LINUX/IRIX
CL4-B  2    VSI_5600-1_Fab_B                       540016  LINUX/IRIX
CL4-B  3    VSI_5600-2_Fab_B                       540016  LINUX/IRIX
CL4-B  4    VSI_5600-3_Fab_B                       540016  LINUX/IRIX
```

Run the following commands to allocate UCS server WWNs to host groups from VSP ports and host group ID. UCS server WWNs can be obtained from Intersight after the correct SAN connectivity policy with a dedicated NVMe vHBA has been applied to the server.

```
C:\HORCM\etc>raidcom add hba_wnn -port CL1-B-1 -hba_wnn 20000025,B5850A0B
C:\HORCM\etc>raidcom add hba_wnn -port CL2-B-1 -hba_wnn 20000025,B5850A0B
C:\HORCM\etc>raidcom add hba_wnn -port CL3-B-1 -hba_wnn 20000025,B5850B0B
C:\HORCM\etc>raidcom add hba_wnn -port CL4-B-1 -hba_wnn 20000025,B5850B0B
C:\HORCM\etc>raidcom add hba_wnn -port CL1-B-2 -hba_wnn 20000025,B5850A0A
C:\HORCM\etc>raidcom add hba_wnn -port CL2-B-2 -hba_wnn 20000025,B5850A0A
C:\HORCM\etc>raidcom add hba_wnn -port CL3-B-2 -hba_wnn 20000025,B5850B0A
C:\HORCM\etc>raidcom add hba_wnn -port CL4-B-2 -hba_wnn 20000025,B5850B0A
C:\HORCM\etc>raidcom add hba_wnn -port CL1-B-3 -hba_wnn 20000025,B5850A09
C:\HORCM\etc>raidcom add hba_wnn -port CL2-B-3 -hba_wnn 20000025,B5850A09
C:\HORCM\etc>raidcom add hba_wnn -port CL3-B-3 -hba_wnn 20000025,B5850B09
C:\HORCM\etc>raidcom add hba_wnn -port CL4-B-3 -hba_wnn 20000025,B5850B09
C:\HORCM\etc>raidcom add hba_wnn -port CL1-B-4 -hba_wnn 20000025,B5850A08
C:\HORCM\etc>raidcom add hba_wnn -port CL2-B-4 -hba_wnn 20000025,B5850A08
C:\HORCM\etc>raidcom add hba_wnn -port CL3-B-4 -hba_wnn 20000025,B5850B08
C:\HORCM\etc>raidcom add hba_wnn -port CL4-B-4 -hba_wnn 20000025,B5850B08
```

Run the following commands to set standard port host groups to host mode LINUX and set the host mode option to 13.

```
C:\HORCM\etc>raidcom modify host_grp -port CL1-B-0 -host_mode LINUX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL2-B-0 -host_mode LINUX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL3-B-0 -host_mode LINUX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL4-B-0 -host_mode LINUX -
set_host_mode_opt 13
```

Run the following commands to set user created host groups to host mode VMware and set the host mode option to 13.

```
C:\HORCM\etc>raidcom modify host_grp -port CL1-B-1 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL1-B-2 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL1-B-3 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL1-B-4 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL2-B-1 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL2-B-2 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL2-B-3 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL2-B-4 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL3-B-1 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL3-B-2 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL3-B-3 -host_mode VMWARE_EX -
set_host_mode_opt 13

C:\HORCM\etc>raidcom modify host_grp -port CL3-B-4 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL4-B-1 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL4-B-2 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL4-B-3 -host_mode VMWARE_EX -
set_host_mode_opt 13
C:\HORCM\etc>raidcom modify host_grp -port CL4-B-4 -host_mode VMWARE_EX -
set_host_mode_opt 13
```

Run the following commands to verify host mode and host mode options.

```
C:\HORCM\etc>raidcom get host_grp -port CL1-B-1
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL1-B  0    1B-G00                                   540016  LINUX/IRIX   13
CL1-B  1    VSI_5600-0_Fab_A                       540016  VMWARE_EX   13
CL1-B  2    VSI_5600-1_Fab_A                       540016  VMWARE_EX   13
CL1-B  3    VSI_5600-2_Fab_A                       540016  VMWARE_EX   13
CL1-B  4    VSI_5600-3_Fab_A                       540016  VMWARE_EX   13
C:\HORCM\etc>raidcom get host_grp -port CL2-B-1
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL2-B  0    2B-G00                                   540016  LINUX/IRIX   13
CL2-B  1    VSI_5600-0_Fab_A                       540016  VMWARE_EX   13
CL2-B  2    VSI_5600-1_Fab_A                       540016  VMWARE_EX   13
CL2-B  3    VSI_5600-2_Fab_A                       540016  VMWARE_EX   13
CL2-B  4    VSI_5600-3_Fab_A                       540016  VMWARE_EX   13
C:\HORCM\etc>raidcom get host_grp -port CL3-B-1
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL3-B  0    3B-G00                                   540016  LINUX/IRIX   13
CL3-B  1    VSI_5600-0_Fab_B                       540016  VMWARE_EX   13
CL3-B  2    VSI_5600-1_Fab_B                       540016  VMWARE_EX   13
CL3-B  3    VSI_5600-2_Fab_B                       540016  VMWARE_EX   13
CL3-B  4    VSI_5600-3_Fab_B                       540016  VMWARE_EX   13
C:\HORCM\etc>raidcom get host_grp -port CL4-B-1
PORT  GID  GROUP_NAME                               Serial#  HMD          HMO_BITS
CL4-B  0    4B-G00                                   540016  LINUX/IRIX   13
CL4-B  1    VSI_5600-0_Fab_B                       540016  VMWARE_EX   13
CL4-B  2    VSI_5600-1_Fab_B                       540016  VMWARE_EX   13
CL4-B  3    VSI_5600-2_Fab_B                       540016  VMWARE_EX   13
CL4-B  4    VSI_5600-3_Fab_B                       540016  VMWARE_EX   13
```

Run the following commands to create an NVMe subsystem and apply VMware host mode.

```
C:\HORCM\etc>raidcom add nvm_subsystem -nvm_subsystem_id 2 -nvm_subsystem_name
UCS_nvm_subsystem -host_mode VMWARE_EX -request_id auto
REQID : 138
```

Run the following commands to verify that the subsystem has been created.

```
C:\HORCM\etc>raidcom get nvm_subsystem -nvm_subsystem_id 2
NVSS_ID  RGID  NVSS_NAME                               SECURITY  T10PI  HMD          HMO_BITS
2        0    UCS_nvm_subsystem                       ENABLE    DISABLE VMWARE_EX   -
```

Run the following commands to verify VSP NQN for the NVMe subsystem.

```
C:\HORCM\etc>raidcom get nvm_subsystem -nvm_subsystem_id 2 -key opt
NVSS_ID  NVSS_NAME                               NVSS_NQN
2        UCS_nvm_subsystem                       nqn.1994-04.jp.co.hitachi:nvme:storage-
subsystem-sn.5-40016-nvssid.00002
```



Run the following commands to add VSP ports to the NVMe subsystem.

```
C:\HORCM\etc>raidcom add nvm_subsystem_port -nvm_subsystem_id 2 -port CL1-B -
request_id auto
REQID : 139
C:\HORCM\etc>raidcom add nvm_subsystem_port -nvm_subsystem_id 2 -port CL2-B -
request_id auto
REQID : 13a
C:\HORCM\etc>raidcom add nvm_subsystem_port -nvm_subsystem_id 2 -port CL3-B -
request_id auto
REQID : 13b
C:\HORCM\etc>raidcom add nvm_subsystem_port -nvm_subsystem_id 2 -port CL4-B -
request_id auto
REQID : 13c
```

Run the following commands to verify that ports have been allocated to the NVMe subsystem.

```
C:\HORCM\etc>raidcom get nvm_subsystem_port -nvm_subsystem_id 2
PORT    NVMSS_ID  NVMSS_NAME
CL1-B   2        UCS_nvm_subsystem
CL2-B   2        UCS_nvm_subsystem
CL3-B   2        UCS_nvm_subsystem
CL4-B   2        UCS_nvm_subsystem
```

Run the following commands to add the NQN to NVMe subsystem after obtaining the ESXi host NQN.

```
C:\HORCM\etc>raidcom add host_nqn -nvm_subsystem_id 2 -host_nqn nqn.2014-
08.local.hvlab.vsi:nvme:esxi-0 -request_id auto
REQID : 13d
C:\HORCM\etc>raidcom add host_nqn -nvm_subsystem_id 2 -host_nqn nqn.2014-
08.local.hvlab.vsi:nvme:esxi-1 -request_id auto
REQID : 13e
C:\HORCM\etc>raidcom add host_nqn -nvm_subsystem_id 2 -host_nqn nqn.2014-
08.local.hvlab.vsi:nvme:esxi-2 -request_id auto
REQID : 13f
C:\HORCM\etc>raidcom add host_nqn -nvm_subsystem_id 2 -host_nqn nqn.2014-
08.local.hvlab.vsi:nvme:esxi-3 -request_id auto
REQID : 140
```

Run the following command to verify that the ESXi host NQN has been added to the NVMe subsystem.

```
C:\HORCM\etc>raidcom get host_nqn -nvm_subsystem_id 2
NVMSS_ID  NVMSS_NAME                HOST_NQN
2         UCS_nvm_subsystem        nqn.2014-08.local.hvlab.vsi:nvme:esxi-2
2         UCS_nvm_subsystem        nqn.2014-08.local.hvlab.vsi:nvme:esxi-3
2         UCS_nvm_subsystem        nqn.2014-08.local.hvlab.vsi:nvme:esxi-0
2         UCS_nvm_subsystem        nqn.2014-08.local.hvlab.vsi:nvme:esxi-1
```

LDEVs must be created on the VSP before running the following commands. In this example LDEV ID 14, 15, 16, and 17 are used. Use the following commands to verify LDEV attributes.

```
C:\HORCM\etc>raidcom get ldev -ldev_id 14
```

```
Serial# : 540016
LDEV : 14
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 4294967296
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
CMP : -
EXP_SPACE : -
B_POOLID : 0
LDEV_NAMING : VSI-VMFS_NVME-01
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 3
SSID : 0004
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
CSV_Mode : DISABLED
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Disable
RSGID : 0
PWSV_S : -
CL_MIG : N
```

```
C:\HORCM\etc>raidcom get ldev -ldev_id 15
```

```
Serial# : 540016
LDEV : 15
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 4294967296
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
CMP : -
EXP_SPACE : -
B_POOLID : 0
LDEV_NAMING : VSI-VMFS_NVME-02
```

```

STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 0
SSID : 0004
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
CSV_Mode : DISABLED
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Disable
RSGID : 0
PWSV_S : -
CL_MIG : N

C:\HORCM\etc>raidcom get ldev -ldev_id 16
Serial# : 540016
LDEV : 16
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 4294967296
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
CMP : -
EXP_SPACE : -
B_POOLID : 0
LDEV_NAMING : VSI-VMFS_NVME-03
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 1
SSID : 0004
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
CSV_Mode : DISABLED
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Disable
RSGID : 0
PWSV_S : -
CL_MIG : N

```

```

C:\HORCM\etc>raidcom get ldev -ldev_id 17
Serial# : 540016
LDEV : 17
SL : 0
CL : 0
VOL_TYPE : OPEN-V-CVS
VOL_Capacity(BLK) : 4294967296
NUM_PORT : 0
PORTs :
F_POOLID : NONE
VOL_ATTR : CVS : HDP
CMP : -
EXP_SPACE : -
B_POOLID : 0
LDEV_NAMING : VSI-VMFS_NVME-04
STS : NML
OPE_TYPE : NONE
OPE_RATE : 100
MP# : 2
SSID : 0004
Used_Block(BLK) : 0
FLA(MB) : Disable
RSV(MB) : 0
CSV_Status : DISABLED
CSV_PROGRESS(%) : -
CSV_Mode : DISABLED
CSV_PROCESS_MODE : -
DEDUPLICATION_DATA : DISABLED
ALUA : Disable
RSGID : 0
PWSV_S : -
CL_MIG : N

```

Run the following commands to allocate LDEVs to namespaces within the NVMe subsystem.

```

C:\HORCM\etc>raidcom add namespace -nvm_subsystem_id 2 -ns_id auto -ldev_id 14 -
request_id auto
REQID : 141
C:\HORCM\etc>raidcom add namespace -nvm_subsystem_id 2 -ns_id auto -ldev_id 15 -
request_id auto
REQID : 142
C:\HORCM\etc>raidcom add namespace -nvm_subsystem_id 2 -ns_id auto -ldev_id 16 -
request_id auto
REQID : 143
C:\HORCM\etc>raidcom add namespace -nvm_subsystem_id 2 -ns_id auto -ldev_id 17 -
request_id auto
REQID : 144

```

Run the following command to verify LDEV allocation to namespaces.

```
C:\HORCM\etc>raidcom get namespace -nvm_subsystem_id 2
NVMSS_ID  NVMSS_NAME                NSID  LDEVID  CAPACITY (BLK)
      2  UCS_nvm_subsystem          1     14     4294967296
      2  UCS_nvm_subsystem          2     15     4294967296
      2  UCS_nvm_subsystem          3     16     4294967296
      2  UCS_nvm_subsystem          4     17     4294967296
```

Run the following commands to create paths between the NVMe namespaces and ESXi hosts.

```
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 1 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-0 -request_id auto
REQID : 145
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 2 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-0 -request_id auto
REQID : 146
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 3 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-0 -request_id auto
REQID : 147
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 4 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-0 -request_id auto
REQID : 148
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 1 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-1 -request_id auto
REQID : 149
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 2 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-1 -request_id auto
REQID : 14a
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 3 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-1 -request_id auto
REQID : 14b
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 4 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-1 -request_id auto
REQID : 14c
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 1 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-2 -request_id auto
REQID : 14d
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 2 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-2 -request_id auto
REQID : 14e
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 3 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-2 -request_id auto
REQID : 14f
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 4 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-2 -request_id auto
REQID : 150
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 1 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-3 -request_id auto
REQID : 151
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 2 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-3 -request_id auto
REQID : 152
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 3 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-3 -request_id auto
REQID : 153
C:\HORCM\etc>raidcom add namespace_path -nvm_subsystem_id 2 -ns_id 4 -host_nqn
nqn.2014-08.local.hvlab.vsi:nvme:esxi-3 -request_id auto
```

Run the following command to verify data paths of hosts and NVMe namespaces and the NVMe subsystem.

```
C:\HORCM\etc>raidcom get namespace_path -nvm_subsystem_id 2
NVMS_ID NVMS_NAME NSID LDEV# HOST_NQN
2 UCS_nvm_subsystem 1 14 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-2
2 UCS_nvm_subsystem 2 15 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-2
2 UCS_nvm_subsystem 3 16 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-2
2 UCS_nvm_subsystem 4 17 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-2
2 UCS_nvm_subsystem 1 14 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-3
2 UCS_nvm_subsystem 2 15 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-3
2 UCS_nvm_subsystem 3 16 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-3
2 UCS_nvm_subsystem 4 17 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-3
2 UCS_nvm_subsystem 1 14 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-0
2 UCS_nvm_subsystem 2 15 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-0
2 UCS_nvm_subsystem 3 16 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-0
2 UCS_nvm_subsystem 4 17 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-0
2 UCS_nvm_subsystem 1 14 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-1
2 UCS_nvm_subsystem 2 15 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-1
2 UCS_nvm_subsystem 3 16 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-1
2 UCS_nvm_subsystem 4 17 nqn.2014-
08.local.hvlab.vsi:nvme:esxi-1
```

After these commands are completed, NVMe datastores can be onboarded from vCenter.

---

## Appendix B: Appendix B – Example HORCM file with virtual command device

Virtual command devices can be used for out-of-band (OOB) communication to the Hitachi VSP. For enterprise systems such as VSP 5000, VSP 1500, or VSP 1000, the IP command device will use the SVP IP address. For midrange systems such as VSP E, F, or G series, the controller GUM address is used. See the *Command Control Interface Installation and Configuration Guide* at [https://knowledge.hitachivantara.com/Documents/Management\\_Software/Command\\_Control\\_Interface](https://knowledge.hitachivantara.com/Documents/Management_Software/Command_Control_Interface) for additional information about the command control interface (CCI).

The following is an example HORCM file for Hitachi VSP OOB communication. The HORCM\_MON parameter defines host 10.76.1.10 listening on UDP port 21002 for HORCM commands. HORCM\_CMD defines the IP address of the SVP or controller depending on the VSP model. In this example HORCM\_CMD leverages a virtual command device pointing to SVP IP 172.25.47.113 using default UDP communication port 31001.

```
HORCM_MON
#ip_address      service      poll(10ms)    timeout(10ms)
10.76.1.10      21002        1000          3000

HORCM_CMD
#dev_name        dev_name      dev_name
\\.\IPCMD-172.25.47.113-31001
```



**Hitachi Vantara**

Corporate Headquarters  
2535 Augustine Drive  
Santa Clara, CA 95054 USA



[HitachiVantara.com/contact](https://HitachiVantara.com/contact)