



NetApp Verified Architecture

# **FlexPod SAN Solution with Cisco UCS X-Series Direct and NetApp ASA NVA deployment**

Jyh-shing Chen, NetApp  
October 2025 | NVA-1178-DEPLOY

In partnership with



## **Abstract**

This NetApp Verified Architecture details the implementation of a FlexPod SAN solution using Cisco UCS X-Series Direct with NetApp's new ASA storage system. It covers the configuration of X-Series Direct managed by Cisco Intersight, ASA storage, and VMware virtual infrastructure deployment. The NVA validates the resiliency of FlexPod infrastructure for mission-critical Oracle and Microsoft SQL databases.

## TABLE OF CONTENTS

<b>Program summary</b>	<b>5</b>
<b>Introduction</b>	<b>5</b>
Audience	5
FlexPod overview	5
FlexPod SAN solution overview	6
Medium-scale FlexPod SAN solution with Cisco UCS X-Series Direct and NetApp ASA in a direct-attached / switchless configuration	8
<b>Solution overview</b>	<b>10</b>
Solution components	10
Solution topologies	15
<b>Deployment procedures</b>	<b>17</b>
Hardware and software revisions	17
Physical connectivity	18
Design and implementation considerations for networking and IP-based SAN	21
Uplink switch configurations	26
ONTAP cluster setup	28
Storage configuration for ESXi hosts	43
Cisco Intersight configuration	45
VMware ESXi 8.0U3 installation	103
Initial ESXi host configuration	107
VMware vCenter 8.0U3 installation and VMware cluster configurations	115
NetApp ONTAP tools for VMware vSphere deployment	134
Microsoft SQL 2022 database server with direct iSCSI LUN access configuration	144
Oracle 21c database server with direct iSCSI LUN access configuration	161
<b>Solution verification</b>	<b>170</b>
Ecosystem interoperability validation	170
Microsoft SQL database testing	179
Oracle RAC database testing	195
Solution availability and infrastructure life-cycle management	203
<b>Conclusion</b>	<b>227</b>
<b>Appendix</b>	<b>227</b>
Appendix A: Resolve issues encountered during solution validation	227
Appendix B: Install SQL server and SQL Server Management Studio on Windows	234



Appendix C: Install and configure bare-metal Oracle Linux 8 .....	241
Appendix D: Oracle Grid Infrastructure installation .....	247
Appendix E: Oracle RAC database installation .....	256
Appendix F: Configure ONTAP storage for client NVMe/TCP access .....	261
Appendix G: Configure VMware cluster and hosts for NVMe/TCP protocol access to storage .....	264
Appendix H: Configure bare-metal Oracle Linux for NVMe/TCP protocol access to storage .....	277
Appendix I: Infrastructure configuration updates for direct-attached FC-based SAN storage support .....	282
Appendix J: Installation and configuration of FC SAN-booted ESXi host .....	311
Appendix K: Configuration updates for FC SAN Booted ESXi host to access NVMe/FC storage .....	318
Appendix L: Configuration updates for the iSCSI SAN-booted ESXi hosts to access FC and NVMe/FC storage .....	328
Appendix M: Configuration updates for iSCSI SAN-booted Oracle Linux host to access FC and NVMe/FC storage .....	337
<b>Where to find additional information .....</b>	<b>353</b>
<b>Acknowledgement .....</b>	<b>355</b>
<b>Version history .....</b>	<b>355</b>

## LIST OF TABLES

Table 1 Key ASA A-Series technical specifications .....	10
Table 2 Hardware and Software versions .....	17
Table 3 UCS X-Series Direct and ASA storage Ethernet connectivity .....	19
Table 4 Cisco UCS X-Series Direct and Cisco Nexus uplink switch connectivity .....	20
Table 5 VLAN information .....	21
Table 6 Server virtual NIC configurations .....	22
Table 7 ONTAP installation and configuration information .....	29
Table 8 Server virtual NIC placement and configurations .....	71
Table 9 LAN connectivity for vNICs .....	82
Table 10 MAC address pools .....	84
Table 11 Ethernet network group policies .....	86
Table 12 Ethernet adapter policy association to vNICs .....	88
Table 13 Non-HA small OTV deployment resource requirements and limits .....	135
Table 14 UCS X-Series Direct and ASA storage Fibre Channel connections .....	283
Table 15 vHBA configuration information .....	294

## LIST OF FIGURES

Figure 1 FlexPod datacenter solution .....	5
--------------------------------------------	---

Figure 2 Small scale FlexPod SAN infrastructure.....	7
Figure 3 Large scale FlexPod SAN infrastructure.....	8
Figure 4 Medium-scale FlexPod SAN with UCS X-Series Direct and NetApp ASA .....	9
Figure 5 ASA A50 front and back views .....	11
Figure 6 UCS X-Series Direct front and rear views .....	13
Figure 7 Cisco UCS Fabric Interconnect 9108 100G (UCSX-S9108-100G).....	13
Figure 8 Cisco UCS X215 M8 compute node .....	14
Figure 9 Cisco UCS VIC 15230 .....	15
Figure 10 Cisco Nexus 93600CD-GX switch.....	15
Figure 11 UCS X-Series Direct and direct-attached ASA topology for IP-based SAN .....	16
Figure 12 UCS X-Series Direct and direct-attached ASA topology for IP-based and FC-based SAN.....	16
Figure 13 IP-based SAN topology .....	17
Figure 14 X-Series Direct and ASA storage Ethernet connectivity .....	20
Figure 15 SAN multipathing .....	23
Figure 16 Virtual NIC configuration for direct iSCSI storage access .....	24
Figure 17 iSCSI protocol physical access to storage.....	25
Figure 18 Igroup mappings of LUNs to initiators.....	26
Figure 19 UCS X-Series Direct and ASA storage Fibre Channel connectivity diagram .....	283

## Program summary

The FlexPod® SAN solution with Cisco® UCS X-Series Direct and NetApp® ASA is a predesigned, best practice architecture utilizing the Cisco Unified Computing System (UCS), the Cisco Nexus switches, and NetApp All-flash SAN Array (ASA) storage technologies.

To provide an optimal infrastructure foundation for remote and branch offices, small to midsize businesses, and dedicated application deployments, this FlexPod SAN solution uses UCS X-Series Direct with integrated Fabric Interconnects to reduce solution components. It also employs a new NetApp ASA A-Series storage system, directly attaching the storage controllers to the Fabric Interconnects without requiring Nexus switches for IP-based storage connection to the UCS X-Series compute nodes. Existing Nexus switches are used for network uplinks access.

The solution remains flexible and can be scaled up or out. Various compute, network, and storage components can be chosen to meet different requirements. Additionally, compute and storage resources can be added post-deployment to enhance performance and capacity.

## Introduction

### Audience

The intended audience of this NetApp Verified Architecture (NVA) includes but is not limited to IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of a resilient FlexPod SAN infrastructure built to deliver IT efficiency and enable IT innovation.

It is assumed that the reader has the following background knowledge:

- Understanding of SAN concepts
- Familiarity with the administration of compute, network, and storage systems

### FlexPod overview

FlexPod® is a best practice converged infrastructure data center architecture that includes the following components from Cisco® and NetApp®:

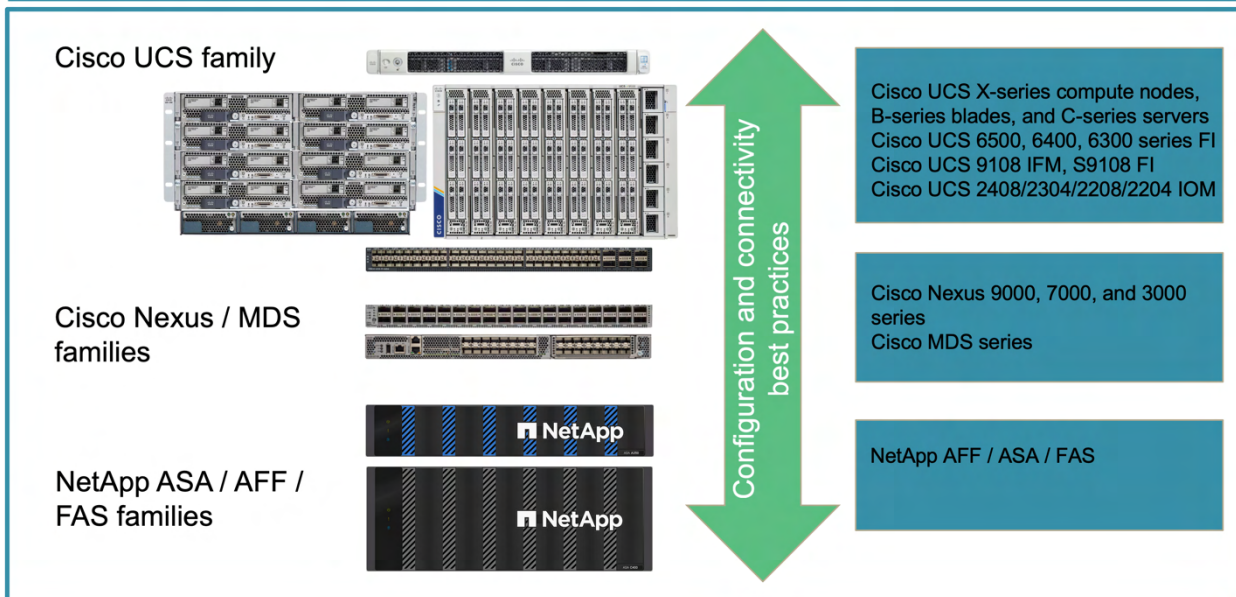
- Cisco Unified Computing System (Cisco UCS®)
- Cisco Nexus and MDS families of switches
- NetApp Fabric-Attached Storage (FAS), All-Flash FAS (AFF), and All-flash SAN Array (ASA).

Shown in Figure 1 are some of the components utilized for creating FlexPod solutions. These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence.

---

**Figure 1 FlexPod datacenter solution**

## FlexPod Datacenter Solution



Each of the FlexPod component families shown (Cisco UCS, Cisco Nexus/MDS switches, and NetApp storage) provides platform and resource options to scale the infrastructure up or down as per application requirement, while supporting the features and functionalities that are required under the configuration and connectivity best practices of FlexPod. FlexPod solutions can also be replicated for environments that require multiple consistent deployments by rolling out additional FlexPod stacks.

All FlexPod components have been integrated so you can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlexPod is its ability to maintain consistency at scale. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. The reference architecture reinforces the wire-once strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect. Furthermore, with template-based and profile-based server management, you can easily enforce configuration consistency for your servers and simplify server life-cycle management such as server replacement while maintaining server identities.

### FlexPod SAN solution overview


Storage Area Network (SAN) provides a reliable and scalable infrastructure for deploying virtual infrastructure like VMware and business critical databases applications running on Microsoft SQL and Oracle databases. The building blocks for a SAN solution are compute, network, and storage components. As there are a variety of components to choose from, understanding your requirements thoroughly will be important for selecting the right solution architecture and components.

### Small-scale FlexPod SAN deployment

When you have a remote office which needs virtual servers to provide an infrastructure reliable for database applications, you might consider a small-scale FlexPod solution design for cost considerations.

Shown in Figure 2 is a small FlexPod SAN infrastructure design using two Cisco UCS C225 M8 rack servers, two Cisco Nexus 93600CD-GX switches, and a highly available entry-level NetApp All-flash SAN Array (ASA) A20 dual-controller storage system.

**Figure 2 Small scale FlexPod SAN infrastructure**



This design utilizes redundant components and connectivity to provide resilient and highly available infrastructure. The two network switches provide dual fabric for IP based SAN deployment. Each server is connected to both SAN fabric A and B. Similarly, the two storage controllers in the ASA A20 system each has connectivity to the two SAN fabric as well.

Despite being a small-scale solution, the design can be scaled out by adding additional servers to provide more compute resources. On the storage side, adding additional ASA A20 HA pair, or picking a different ASA model, can provide higher storage performance and adding external disk shelves can provide additional storage capacity.


## Large-scale FlexPod SAN deployment

For enterprise datacenters, deploying a large-scale SAN solution infrastructure to provide secure multi-tenant access for various departments, or customers in the service provider use cases, requires lots of compute and storage resources, ease of management, and the ability to implement zero-trust framework for security considerations.

FlexPod datacenter solutions are designed with security as the foundation and provide validated designs for business-critical applications so you can deploy the solutions with confidence. To see the latest available Cisco Validated Designs (CVDs) for FlexPod, please check out the [FlexPod Design Guide](#) web site.

Shown in Figure 3 is a FlexPod datacenter SAN infrastructure design using Cisco UCS X-Series Modular System with X9508 Chassis, 9108 Intelligent Fabric Modules (IFMs), X215c M8 Compute Nodes, UCS 6536 Fabric Interconnects (FIs), Nexus 9336C-FX2 switches, MDS 9132T switches, and a highly available high-end NetApp ASA A90 dual-controller storage system.

**Figure 3 Large scale FlexPod SAN infrastructure**



This large scale FlexPod datacenter infrastructure is designed with scalability in mind and includes redundant components and connectivity for resiliency and high availability. The Cisco UCS X-Series X9508 Chassis, the UCS 9108 IFM and the X215c M8 AMD CPU based compute node along with the UCS 6536 FI constitute the compute block which can be scaled massively by adding additional X9508 Chassis and compute nodes. Additional disk shelves and ASA A90 HA pairs can be added to increase storage capacity and storage performance to a massive scale.

The above solution infrastructure design includes both Nexus Ethernet switches as well as MDS Fibre Channel switches to support multiple storage protocols, including FC, iSCSI, Non-volatile Memory Express (NVMe) over FC (NVMe/FC), and NVMe over TCP (NVMe/TCP).

You can choose to implement IP-based SAN only, FC-based SAN only, or implement both to adopt different protocols for workload latency and throughput requirements. The NVMe over fabric protocols offers lower latency and improved input/output operations per second (IOPs) due to more efficient device connections, command sets, and reduced CPU utilization.

Furthermore, you can enable AI and accelerate VDI workload running on the FlexPod solution infrastructure by incorporating Cisco UCS X-Fabric Technology with the UCS 9416 X-Fabric Module and adding GPU nodes. By integrating the X440p PCIe GPU Nodes with your X215c, X210c, or X410c compute nodes through the UCS X9416 X-Fabric Module, you can enhance VDI workload performance and begin utilizing AI capabilities within your FlexPod infrastructure. FlexPod infrastructure with AI capabilities enable deeper data analysis and the development of AI workflows, ultimately aiming to improve business profitability.

### **Medium-scale FlexPod SAN solution with Cisco UCS X-Series Direct and NetApp ASA in a direct-attached / switchless configuration**


Some customer use cases and applications might require a medium solution scale where the compute and storage resources and performances might start small but requires limited amount of scaling as the business grows. There might also be mandates for the solution to be deployed cost effectively for multiple offices or for multiple applications.



The FlexPod SAN solution with Cisco UCS X-Series Direct and the mid-range NetApp ASA A50 is a great combination for this target scale. Like its UCS Mini predecessor, the X-Series Direct integrates two Cisco UCS Fabric Interconnect 9108 100G, in the back of X9508 Chassis, thus reducing the UCS rack space requirements and costs for building medium-scale FlexPod SAN solutions.

By configuring some of the FI ports as appliance ports, you can directly attach the new NetApp ASA storage systems to the FIs without requiring additional Ethernet switches in between UCS compute and NetApp storage. Figure 4 shows the medium-scale FlexPod IP-based SAN solution with X-Series Direct and NetApp ASA storage in a direct-attached / switchless configuration.

**Figure 4 Medium-scale FlexPod SAN with UCS X-Series Direct and NetApp ASA**



**Note:** Other compute nodes supported on the X-Series Direct can also be used. Similarly, other ASA storage systems can also be used based on the storage requirements.

**Note:** You can also build a FlexPod FC-based SAN solution without additional external MDS FC switches by directly attaching NetApp ASA storage system to the FIs' appliance ports which support FC breakout transceivers and with the FIs' FC Switching Mode configured for switch mode to provide FC switching services. Please consult [Cisco UCS X-Series Direct Fabric Interconnect 9108 100G installation and service guide](#) for alternative SAN topology examples for the X-Series Direct.

The X-Series Direct with direct-attached ASA storage design allows you to scale your compute and storage resources to a medium scale. You can start with just a few compute nodes in the chassis. When you need more compute resource as your business grows, you can simply add additional compute nodes without adding any additional cables until you have fully populated the chassis. You can add additional storage capacity by connecting external disk shelves to the storage controller. You can also add additional storage performance by expanding the storage cluster size with additional storage controller HA pairs and attaching those additional storage controllers to the FIs' appliance ports. Like the large scale FlexPod SAN deployment, you can also utilize GPUs to accelerate VDI or enable AI workloads on the X-Series Direct.

**Note:** For a storage cluster with more than two nodes, a pair of cluster switches will be needed for storage cluster communication.

**Note:** When the FIs run out of appliance ports, the solution topology can be updated to utilize switch-attached storage to further expand the storage cluster size by following FlexPod datacenter-based solution designs.

## Solution overview

To understand the FlexPod IP-based SAN solution with Cisco UCS X-Series Direct and NetApp ASA, we will go over the solution components, solution topology and connectivity, and some solution design considerations. For solution validation, we will provide information on ecosystem interoperability validation, database testing configuration for Microsoft SQL server and Oracle RAC, solution availability, and some life-cycle management activities. In the Appendix, we also include information for deploying Oracle database on bare-metal server and configuring NVMe/TCP protocol access for the VMware infrastructure and bare-metal Oracle Linux server.

## Solution components

### NetApp All-flash SAN Array (ASA)

The new NetApp All-flash SAN Array (ASA) all-flash scale-out storage systems are simple, powerful, optimized for block deployments, and support advanced data management and protection features. With six-nines availability and symmetric active-active multipathing, the new NetApp ASA systems can be used to modernize SAN infrastructure, simplify storage management, and accelerate business critical applications.

The NetApp ASA family includes A-Series models for performance-demanding workloads and the C-Series models optimized for cost-effective large-capacity and general-purpose applications. The new ASA A-Series family, which is the focus of this NVA, includes ASA A20, A30, A50, A70, A90, and A1K. They offer a wide range of performance and capacity to meet the diverse customer requirements for their SAN solutions. The ASA A-Series platforms support IP-based and FC-based SAN protocols, including iSCSI, NVMe/TCP, FC, and NVMe/FC. Table 1 below highlights some of the key technical specifications of the new ASA A-Series models.

**Table 1 Key ASA A-Series technical specifications**

Specifications	ASA A1K	ASA A90	ASA A70	ASA A50	ASA A30	ASA A20
Form factor	2 x 2U	4U	4U	2U	2U	2U
Max cluster size	12 Nodes	12 Nodes	12 Nodes	12 Nodes	8 Nodes	6 Nodes
Max raw capacity per HA pair	2.67 PB	2.67 PB	2.67 PB	1.8 PB	1.1 PB	734 TB
Max raw capacity per cluster	16 PB	16 PB	16 PB	11 PB	4.4 PB	2.2 PB
PCIe expansion slots per HA pair	18	18	18	8	8	8
Max FC speed	64 Gbps	64 Gbps	64 Gbps	64 Gbps	64 Gbps	64 Gbps
Max Ethernet speed	200 Gbps	200 Gbps	200 Gbps	100 Gbps	100 Gbps	100 Gbps
Minimum ONTAP version	9.16.0 GA	9.16.0 GA	9.16.0 GA	9.16.1	9.16.1	9.16.1

**Note:** For this FlexPod SAN solution validation, we are utilizing an ASA A50 HA pair. Please see [NetApp ASA datasheet](#) and [NetApp Hardware Universe](#) for detailed technical specifications and the various supported configurations and cluster limits of these new ASA platforms.

The NetApp new ASA A-series storage systems offer the following important benefits:

- High data availability with dual-controller architecture, symmetric active-active multipathing access to storage, non-disruptive firmware upgrade, and six-nines data availability guarantee.




- Scale-out storage performance with clustering for up to 6 nodes for ASA A20, up to 8 nodes for ASA A30, and up to 12 nodes for the remaining mid-range and high-end ASA A-Series storage systems.
- Always-on storage efficiency with data compression, deduplication, and compaction to efficiently store your data with four to one storage efficiency guarantee.
- Secure storage access and administration with multi-factor authentication, role-based access control, multi-admin verification to safeguard sensitive storage operations, secure multi-tenancy to minimize attack surface, and FIPS 140-2 compliant mode for regulatory compliance.
- Simple to deploy, manage, and upgrade the storage system. Easy to provision SAN storage units (LUNs and NVMe namespaces) and add data protection from ONTAP System Manager. The integration with vCenter allows storage provisioning directly from the NetApp ONTAP tools for VMware vSphere plugin in vCenter.

To get started on using the new NetApp ASA A-Series storage systems to protect your critical business data, review the solution information and deployment procedures in this NVA and check out the additional information available on the [ASA documentation](#) site.

## NetApp ASA A50

The NetApp ASA A50 comes in 2U form factor per HA pair. An HA pair includes two storage controllers and a total of 48 CPU cores, 256GB of RAM, 32GB of NVRAM, 8 expansion slots, and up to 24 drives in the embedded disk shelf. The ASA A50 platform supports FC, iSCSI, NVMe/FC, and NVMe/TCP protocols. The ASA A50 storage system can be scaled out to a maximum cluster size of 12 nodes (6 HA pairs), and a max raw capacity of 11PB per cluster at max cluster size.

**Figure 5 ASA A50 front and back views**



## NetApp ONTAP

NetApp® ONTAP® is an enterprise data-management platform with native capabilities that includes industry leading storage efficiency capabilities and could be clustered up to 12 nodes for SAN using mid-range and high-end platforms. NetApp Snapshot™ technology which is an integral part of ONTAP enables instantaneous backups of critical datasets and cloning of datasets and offers comprehensive disaster recovery capabilities.

NetApp storage systems also offer a hybrid cloud foundation for customers to take advantage of the seamless data mobility enabled by NetApp intelligent data infrastructure which can easily get data from the edge where it is generated to the core where it is utilized and to the cloud to take advantage of the on-demand elastic compute and AI / ML capabilities to gain actionable business insights from your valuable

data. Consult [NetApp Hardware Universe](#) and [ASA documentation](#) sites for the latest updates on the supported limits and features for the new ASA storage systems.

NetApp ONTAP System Manager is a web service included with ONTAP. You can use a web browser to manage ONTAP storage system and storage units (SCSI LUNs and NVMe namespaces). You can create consistency group for related storage units used by an application and then create Snapshot for the consistency group to back up your application. In ONTAP System Manager, you can view information regarding important alerts and notifications, see capacity utilization and storage efficiency, configure network and host access to the storage system, and perform ONTAP updates.

If you are already familiar with NetApp AFF, FAS, and prior ASA systems, the new ASA systems have the following main differences:

- When a cluster is created, a default data SVM, svm1, is created with the SAN protocols enabled. IP data LIFs support iSCSI and NVMe/TCP protocols and use the default-data-blocks service policy by default.
- Instead of user-managed aggregates and volumes, ONTAP manages the common pool of storage with storage availability zone and automatically creates or deletes the associated volume when a storage unit is created or deleted.
- Storage units are always thinly provisioned, and they are automatically placed and rebalanced as needed.
- Temperature-sensitive storage efficiency is not applicable, and data compression begins without waiting for data to become cold.

Due to the differences, ONTAP CLI, ONTAP System Manager, and ONTAP REST API endpoints have changed accordingly to simplify user experience.

## Cisco Intersight

Cisco Intersight® is an IT operations platform for infrastructure lifecycle management. It is delivered as a software-as-a-service (SaaS) platform and is also available in a connected or private virtual appliance format. Intersight helps IT teams manage and automate their Cisco UCS infrastructure across data centers, and remote and branch offices.

The Cisco Intersight managed mode manages Cisco UCS fabric interconnect-attached systems through a Redfish-based standard model. Its template- and policy-based management simplify server deployment and enforce configuration consistency. Intersight is used to manage the Cisco UCS X-Series Direct and the Cisco UCS X215c M8 Compute Nodes used in this FlexPod SAN solution.


## Cisco UCS X-Series Direct

Cisco UCS X-Series Modular System with the X9508 Chassis is engineered to be adaptable and future-ready. With an I/O midplane-free design, I/O connectivity for the X9508 Chassis is accomplished with front-loading vertically oriented computing nodes that intersect with horizontally oriented I/O connectivity modules in the rear of the chassis. Cisco UCS X-Series is powered by Cisco Intersight, making it simple to deploy and manage at scale.

Cisco UCS X-Series Direct integrates the UCS Fabric Interconnect 9108 100G directly into the UCS X9508 chassis at the top rear. The FlexPod SAN solution created using UCS X-Series Direct requires fewer components, less rack space, and reduced power and cooling requirements.

Figure 6 presents the front and rear views of the UCS X-Series Direct. The front view displays four X215c M8 Compute Nodes, while the rear view shows two integrated fabric interconnects positioned at the top of the chassis.

**Figure 6 UCS X-Series Direct front and rear views**



The Cisco UCS X-Series Direct provides the following important features:

- The seven-rack-unit chassis has eight front-facing vertical slots. These can house a combination of compute and GPU nodes, e.g. X215c M8, X210c M7, and X410c M7, and X440p PCIe GPU nodes.
- Two Cisco UCS Fabric Interconnect 9108 100G provide uplink network switch connectivity and direct-attached storage connectivity to support both IP-based SAN and FC-based SAN.
- At the bottom of the chassis in the rear are optional slots for Cisco UCS 9416 X-Fabric Modules that can be used to support UCS X440p PCIe GPU nodes' connectivity.

You can scale the number of compute nodes deployed in the X-Series Direct chassis from just a few to the maximum of eight compute nodes as you grow your business. In addition, Intersight managed UCS X-Series Direct gives you the visibility into all your UCS X-Series Direct deployments in remote and branch offices and the same wire-once, policy-based UCS management benefits.

## Cisco UCS Fabric Interconnect 9108 100G

Fabric Interconnects are an important part of the Cisco UCS solution for UCS management and network connectivity. The Cisco UCS Fabric Interconnect 9108 100G is designed exclusively for the UCS X-Series Direct as it plugs directly into the UCS X9508 chassis at the top of the rear. See Figure 7 for the available connections of the UCS Fabric Interconnect 9108 100G.

**Figure 7 Cisco UCS Fabric Interconnect 9108 100G (UCSX-S9108-100G)**




The UCS Fabric Interconnect 9108 100G has the following important features:

- Supports 1/10/25/40/100-Gbps Ethernet or Fibre Channel over Ethernet (FCoE).
- Supports up to 1.6 Tbps throughput for servers with 8 x 100Gbps per FI and two FIs per chassis.
- Supports up to 8 x 32Gbps Fibre Channel connectivity through breakout cables on unified ports (1 and 2).
- Supports connecting the chassis to upstream network switches on Ethernet uplink ports and can directly attach storage appliances using IP-based or FC-based SAN connectivity on appliance ports.

## Cisco X215c M8 compute node

Cisco UCS X215c M8 Compute Node (Figure 8) is integrated into the Cisco X-Series Modular System. Up to eight Cisco UCS X215c M8 Compute Nodes can reside in the 7-rack-unit Cisco X-Series Direct's X9508 chassis, offering one of the highest densities of compute, and IO per rack unit in the industry.

**Figure 8 Cisco UCS X215 M8 compute node**



The Cisco UCS X215c M8 Compute Node offers the following:

- Up to two 4<sup>th</sup> / 5<sup>th</sup> Gen. AMD EPYC™ CPUs with up to 128 / 160 cores per processor.
- Up to 24 x 256GB DDR5-5600 / DDR5-6400 DIMMs, in a 2-socket configuration with 4<sup>th</sup> / 5<sup>th</sup> Gen. AMD EPYC™ CPUs
- Optional front mezzanine GPU module to support up to two half-height half-length GPUs.
- Modular LAN on Motherboard (mLOM) support for Cisco UCS Virtual Interface Card (VIC) 15230 for 100Gbps connectivity and server secure boot, or VIC 15420 for up to 50Gbps (2 x 25Gbps) per server per fabric interconnect.
- Optional PCIe mezzanine card for X-Fabric can also be installed in the server's mezzanine slot at the bottom rear of the chassis to connect to UCS X-Fabric modules and enable connectivity to the X440p PCIe node.
- Mezzanine slot support for VIC 15422 to add to VIC 15420 connectivity for a total of 100Gbps (4 x 25) per server per fabric interconnect. VIC 15422 also links to X-Fabric technology.
- The server security feature includes secure boot silicon root of trust FPGA, ACT2 anti-counterfeit provisions, and optional Trusted Platform Model (TPM).

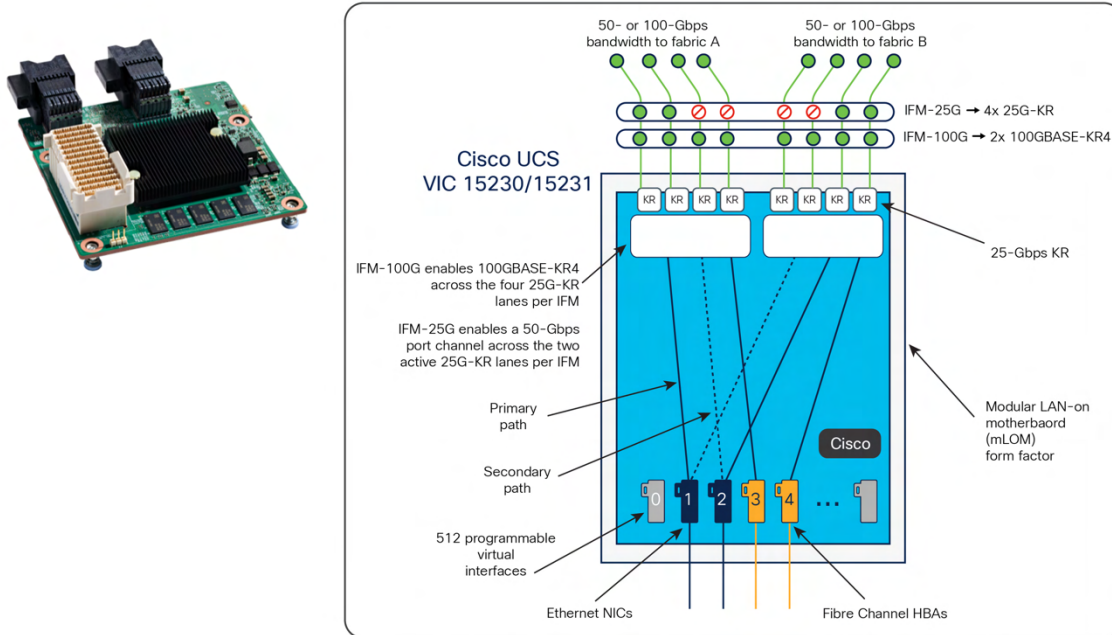
## Cisco UCS VIC 15230

The Cisco UCS VIC 15230 includes secure boot technology and replaces the otherwise equivalent VIC 15231 which does not have secure boot support. Cisco UCS VIC 15230 / 15231 (Figure 9) are 2 x 100Gbps Ethernet/FCoE-capable modular LAN on motherboard (mLOM) adapters designed for the Cisco UCS X-Series compute nodes.

The Cisco UCS VIC 15230 (Figure 9) adapters used in this solution validation enable policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs in the server profile.



**Figure 9 Cisco UCS VIC 15230**



## Cisco Nexus 93600CD-GX switches

The Cisco Nexus 9000 Series Switches offer both modular and fixed 1/10/25/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of nonblocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

The Cisco Nexus 93600CD-GX Switch (Figure 10) is a 1RU switch that supports 12 Tbps of bandwidth and 4.0 bpps across 28 fixed 40/100G QSFP-28 ports and 8 fixed 10/25/40/50/100/200/400G QSFP-DD ports. The 28 ports support 10/25-Gbps.

**Figure 10 Cisco Nexus 93600CD-GX switch**



For the FlexPod IP-based SAN solution with Cisco UCS X-Series Direct and NetApp ASA, the storage cluster is directly connected to the UCS Fabric Interconnect 9108 100G in the X-Series Direct without using a pair of Nexus switches in between. The Nexus switches is used to provide in-band management access to the solution for solution deployment and monitoring.

## Solution topologies

This section describes two direct-attached storage topologies for building a FlexPod SAN solution with Cisco UCS X-Series Direct and NetApp ASA. The first topology (Figure 11) supports IP-based SAN protocols and is the topology deployed for this solution validation.


If your solution also requires FC-based protocols, then you can deploy the second solution topology (Figure 12) which supports both IP-based and FC-based SAN protocols with the direct-attached ASA storage.

## Topology for IP-based SAN

The topology for a FlexPod IP-based SAN solution is very simple when directly attaching a NetApp ASA storage system to the appliance ports of the integrated FIs in the UCS X-Series Direct as shown in Figure

11. The solution uses redundant components and connectivity to ensure data availability when encountering single-point-of-failure scenarios and for life-cycle management activities like rebooting a component after software / firmware upgrade.

**Figure 11 UCS X-Series Direct and direct-attached ASA topology for IP-based SAN**



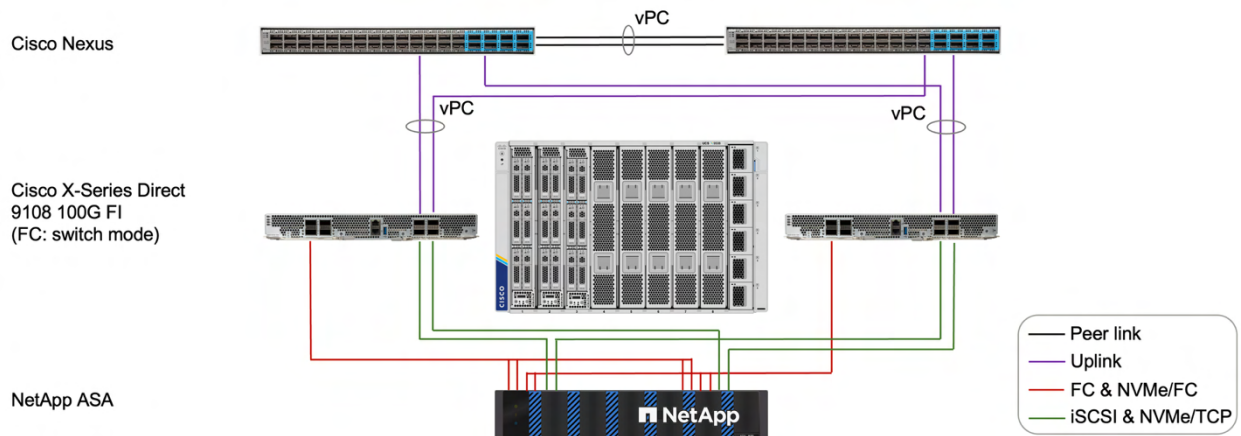
For the storage connections to the FIs in the X-Series Direct, these connections are being treated as individual links without link aggregation configuration because FIs do not support virtual port channel (vPC). As a result, interface groups are also not configured on the ASA storage system for those connections. Link-level failure resiliency is provided by SAN multipathing available in the solution infrastructure and configured in the deployed operating systems.

Even though Cisco Nexus switches are shown in the topology above, they are not involved in the SAN data paths and are only used for FI network uplinks so the solution environment can be accessed from the outside world. As a result, existing top-of-rack Ethernet switches can also be utilized if they support VLAN and vPC or equivalent technologies.

## Topology for both IP-based and FC-based SAN

When your solution requires both IP-based and FC-based SAN protocols, you can add additional FC connectivity directly from the FIs to the FC adapters in the NetApp ASA storage system as shown in Figure 12.

**Figure 12 UCS X-Series Direct and direct-attached ASA topology for IP-based and FC-based SAN**



For proper FC SAN operations, the FC Switching Mode for the FIs will need to be configured as switch mode to provide FC related services. Each FI is considered as a separate FC fabric and the FC best practices applicable to FC switches can be applied on the FIs. The deployment information below focuses on IP-based SAN configurations and validations. Please refer to appendices for FC-based SAN configuration details.


**Note:** Only port 1 and 2 on the FIs can be configured as FC storage port for FC connectivity to the ASA storage controllers using DS-SFP-4x32G-SW transceiver for breakout. Please consult [Cisco UCS X-Series Direct fabric interconnect 9108 100G installation and service guide](#) for details.

## Deployment procedures

For this FlexPod solution validation, we will focus on deploying Intersight managed UCS X-Series Direct, ASA A50 storage system, VMware virtual infrastructure, ONTAP tools for VMware vSphere, Microsoft SQL server, and Oracle RAC database using iSCSI protocol in a direct-attached storage configuration.

We will cover end-to-end iSCSI configuration shown in Figure 13 for the VMware virtual infrastructure and demonstrate the direct iSCSI protocol access to storage from the SQL and Oracle virtual machines. To provide external access to the solution, we will also cover network uplink configurations to an existing Nexus switch pair.

Figure 13 IP-based SAN topology



## Hardware and software revisions

FlexPod solutions can be built by using any supported hardware components and software versions that are listed in the [NetApp Interoperability Matrix Tool](#), [Cisco UCS Hardware and Software Compatibility List](#), and [Broadcom Compatibility Guide](#). Table 2 shows an example of the hardware and software revisions that can be used for the FlexPod solution validation.

Table 2 Hardware and Software versions

Component	Product	Version
Compute	Cisco UCS Fabric Interconnects UCSX-S9108-100G (Intersight managed mode)	4.3(5.240191)
	Cisco UCSX-215c-M8	5.3(0.250001)
	CPU	AMD EPYC 9534 64-Core CPU 2.45GHz

	Cisco VIC 15230 UCSX-ML-V5D200GV2	5.3(4.84)
Network	Cisco Nexus 93600CD-GX NX-OS	10.4(4)M
Storage	NetApp ASA A50 ONTAP	9.16.1P7
Software	Cisco Intersight	SaaS
	VMware vSphere	vSphere 8 update 3
	Cisco VIC nenic driver	2.0.15.0
	Cisco VIC nfnic driver	5.0.0.45
	VMware vCenter	vCenter 8 update 3
	NetApp ONTAP tools for VMware vSphere	10.4
	Oracle Database 21c Grid Infrastructure	21.3
	Oracle Database 21c Enterprise Edition	21.3
	SLOB	2.5.4.0
	Microsoft Windows Server 2022 Datacenter	10.0.20348
	Microsoft SQL Server	2022
	Microsoft SQL Server Management Studio	20.2.1
	HammerDB	4.12

**Note:** Additional software versions are needed for life-cycle management procedure validations.





## Physical connectivity

The following sections describe the physical connectivity needed for the solution deployment. Specifically, the connectivity in the following three areas: ASA A50 cluster connectivity, UCS X-Series Direct and ASA storage Ethernet connectivity, and UCS X-Series Direct FI network uplink connectivity.

### ASA A50 cluster connectivity

For the two-node ASA A50 cluster deployed for this solution validation, the controllers in the HA pair are connected directly together with cables without using cluster switches for cost savings. Please refer to the Supported Adapter Cards and Supported Connections sections in [NetApp Hardware Universe](#) (HWU) for your specific controller model and cluster interconnect configurations if you are not using the ASA A50 model for your solution deployment.

In the ASA A50 controllers, they have two dual-port 40/100G Ethernet Controller, X60130A, cards for Cluster/HA and IP-based SAN. The information from HWU Supported Adapter Cards section for ASA A50 indicates that X60130A is used for Cluster/HA and the Priority Slot Assignment for these two adapters are 4,2. So, these two cards should be installed in slots 4 and 2.

Supported Adapter Cards - ASA A50 with ONTAP 9.16.1														 Adapter Card Help Guide
Priority	Category	Bus Type	Mktg Part No	Images	LED	Mfg Part No	Description	Plug Type	Optical Module Included?	Cables	Supported Protocol(s)	Min ONTAP	Max Qty <sup>[1]</sup>	Priority Slot Assignment
1	Cluster/ HA	IO Module	X60130A		-	111-05341	2p, 40G/100G Ethernet Controller CX6-DX	QSFP28	No	<a href="#">View</a>	Ethernet 40/100 Gb, RoCEv2	9.16.1	2	4,2
2	NVMe Storage	IO Module	X60130A <sup>[2]</sup>		-	111-05341	2p, 40G/100G Ethernet Controller CX6-DX	QSFP28	No	<a href="#">View</a>	Ethernet 40/100 Gb, RoCEv2	9.16.1	2	3,1
3	Block Access (Target)	IO Module	X60143A <sup>[3]</sup>		<a href="#">View</a>	111-05342	4P,64Gb FC,Target-Init,NO SFP	SFP+	Yes	<a href="#">View</a>	Fibre Channel 16/32/64 Gb, NVMe/FC 32/64 Gb	9.16.1	3	2,1,3

From the Cluster Interconnect tab of the Supported Connections table, we can find the 2-node Switchless Cluster/HA 100GbE (option 2) with Dual NIC configuration. Based on the information indicated, the e4a



and e2a ports of the X60130A adapters in the two controllers should be connected for cluster interconnect. (i.e. e4a to e4a and e2a to e2a) This leaves e2b and e4b available for IP-based SAN usage.

NetApp | Hardware Universe

Products ▾ Utilities ▾ Toolbox ▾ Inform

Specifications

Specifications

ASA A50

Click here

Supp

R

Supported Connections - ASA A50 with ONTAP 9.16.1

Cluster Interconnect

Storage Connection

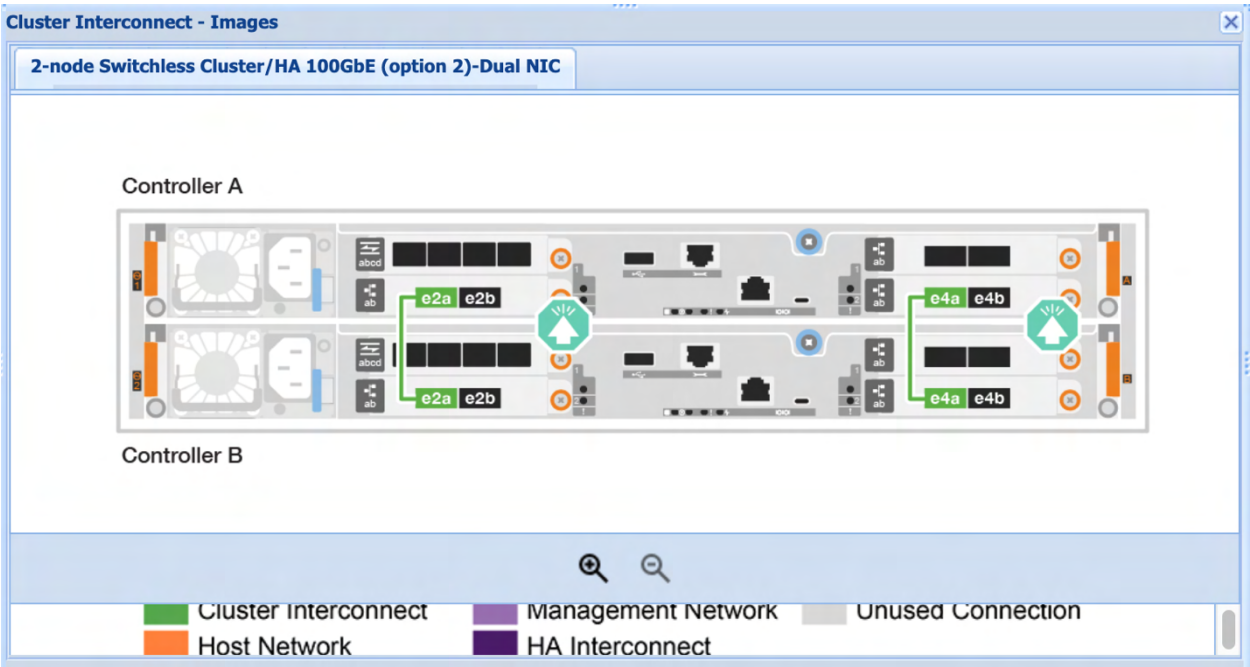
Shared Cluster/Storage Connection

Configuration Type	Port Configuration	Diagram	Cluster Ports per Node	Switch Ports per Node
<input type="radio"/> 2-node Switchless Cluster/HA 100GbE (option 1)	Default, Single NIC	<a href="#">View</a>	2	None
<input checked="" type="radio"/> 2-node Switchless Cluster/HA 100GbE (option 2)	Dual NIC	<a href="#">View</a>	2	None
<input type="radio"/> Switched 100GbE or 40GbE Cluster/HA (Option 1)	Default, Single NIC	<a href="#">View</a>	2	2
<input type="radio"/> Switched 100GbE or 40GbE Cluster/HA (Option 2)	Dual NIC	<a href="#">View</a>	2	2
<input type="radio"/> Switched 10/25GbE Cluster (Option 1)	Supported		4	4
<input type="radio"/> Switched 10/25GbE Cluster (Option 2)	Supported		4	4

Connection Configuration: 2-node Switchless Cluster/HA 100GbE (option 2) [Dual NIC]

Port Location	Port	Target	Target Ports
Adapter Card (X60130A)	e4a	Partner Node	e4a
	e2a	Partner Node	e2a

Clicking on the View link under the Diagram column in the HWU screenshot above shows an image of the cluster interconnect cabling.



UCS X-Series Direct and ASA storage Ethernet connectivity


The IP-based SAN connectivity between ASA A50 and the UCS Fabric Interconnect 9108 100G are shown in Table 3 and Figure 14. The two ASA A50 storage controllers, A and B, have connectivity to the two independent IP SAN fabric provided by FI A and FI B.

Table 3 UCS X-Series Direct and ASA storage Ethernet connectivity

X-Series Direct	Port	NetApp ASA Storage	Port
Cisco UCS Fabric Interconnect 9108 100G A	eth 1/3	NetApp ASA A50 A	e2b
	eth 1/4	NetApp ASA A50 B	e2b

Cisco UCS Fabric Interconnect 9108 100G B	eth 1/3	NetApp ASA A50 A	e4b
	eth 1/4	NetApp ASA A50 B	e4b

**Figure 14 X-Series Direct and ASA storage Ethernet connectivity**



## UCS X-Series Direct FI network uplink connectivity

For the IP-based SAN operations within the solution, the uplink Nexus switches are not involved in the SAN data paths. SAN IO goes directly from the UCS compute nodes to ASA storage using their connections to FI A or FI B.

As a result, existing top-of-rack switches that supports VLAN and vPC can be utilized to provide external network access to the FlexPod SAN solution environment such as for the in-band management access to the ESXi hosts and the virtual machines running in the VMware cluster. Provides information on how the uplinks can be connected.

**Table 4 Cisco UCS X-Series Direct and Cisco Nexus uplink switch connectivity**

X-Series Direct	Port	Cisco Nexus Uplink Switch	Port
Cisco UCS Fabric Interconnect 9108 100G A	eth 1/5	Cisco Nexus A	eth uplink port for FI A
	eth 1/6	Cisco Nexus B	eth uplink port for FI A
Cisco UCS Fabric Interconnect 9108 100G B	eth 1/5	Cisco Nexus A	eth uplink port for FI B
	eth 1/6	Cisco Nexus B	eth uplink port for FI B

## Physical connectivity validation

To confirm physical connectivity of the solution after configuration, the Cisco Discovery Protocol (CDP) enabled on the switches and storage can be utilized to report the identity of the connected devices. This validation should be performed shortly after the connections are made and after the ONTAP cluster is created.

The CDP neighbor information from the uplink Nexus switches can be used to confirm the connected X-Series Direct FI devices and ports. Use the example command below and replace the <FI-uplink-ports>

with the ports that you used for FI connectivity to display CDP neighbor information from your uplink Nexus switches.

```
# show cdp neighbors interface eth<FI-uplink-ports>
```

In addition, the outputs from the network device-discovery show command in ONTAP can be used to confirm both ONTAP cluster interconnect connections (e2a, e4a) and the connections from the storage controllers (e2b, e4b) to the X-Series Direct FIs.

```
fpsa-a50-u0909:> network device-discovery show -node fpsa-a50-u0909-0* -protocol cdp -port
e2*,e4*
Node/      Local   Discovered
Protocol   Port    Device (LLDP: ChassisID)  Interface      Platform
-----
fpsa-a50-u0909-02/cdp
  e2a      fpsa-a50-u0909-01      e2a      ASA-A50
  e2b      fpsa-x9508-u0901-fi-A.nva.local (*****
            Ethernet1/4          UCSX-S9108-100G
  e4a      fpsa-a50-u0909-01      e4a      ASA-A50
  e4b      fpsa-x9508-u0901-fi-B.nva.local (*****
            Ethernet1/4          UCSX-S9108-100G
fpsa-a50-u0909-01/cdp
  e2a      fpsa-a50-u0909-02      e2a      ASA-A50
  e2b      fpsa-x9508-u0901-fi-A.nva.local (*****
            Ethernet1/3          UCSX-S9108-100G
  e4a      fpsa-a50-u0909-02      e4a      ASA-A50
  e4b      fpsa-x9508-u0901-fi-B.nva.local (*****
            Ethernet1/3          UCSX-S9108-100G
8 entries were displayed.
```

**Note:** The serial numbers of the FIs are replaced with asterisks in the outputs above.

## Design and implementation considerations for networking and IP-based SAN

In this section, we briefly discuss a few networking and IP-based SAN design aspects to provide the background information for the deployment configurations.

### Network segmentation

For security and ease of network traffic management, VLANs are utilized to segregate the different network traffic types. Table 5 lists the VLANs defined and used for the solution validation. You should adapt the information below to a configuration that is suitable for your environment.

Also, depending on your needs, you may or may not need all the VLANs listed below. For example, if you are only interested in using iSCSI protocol, then you don't need the two VLANs defined for use with NVMe/TCP protocol.

**Table 5 VLAN information**

VLAN Name	ID	Remarks
Native VLAN	2	Native VLAN
OOB-Mgmt VLAN	2271	Out-of-band Management VLAN
IB-Mgmt VLAN	2272	In-band Management VLAN
iSCSI-A VLAN	2273	iSCSI A VLAN
iSCSI-B VLAN	2274	iSCSI B VLAN
vMotion VLAN	2275	vMotion VLAN
VM-Traffic VLAN	2276	VM Traffic VLAN
NVMe-TCP-A VLAN	2277	NVMe/TCP A VLAN

NVMe-TCP-B VLAN	2278	NVMe/TCP B VLAN
-----------------	------	-----------------

**Note:** For this validation, the out-of-band management VLAN for the solution components' management IPs and the solution components' console connections are using an existing network infrastructure which are not included in the solution topology diagrams.

## Server virtual NICs

On Cisco UCS, templates and profiles are used to simplify server management as well as enforce configuration consistency. For the template configuration used to derive the ESXi host profiles, there are six virtual NICs defined for the server, which is the same as other FlexPod solutions with IP-based storage protocol access design. Table 6 lists the virtual NICs configurations for the servers. All vNICs are pinned to either fabric A or fabric B without enabling failover. The vNIC failover for in-band management, VM traffic, and vMotion network are accomplished by having multiple vNICs connected to the virtual switch (vSwitch) or virtual distributed switch (vDS) in VMware with failover configurations. For iSCSI and NVMe/TCP traffic, SAN multipathing is the mechanism for path management and component or path failure resiliency.

**Table 6 Server virtual NIC configurations**

vNIC Name	Switch ID	Failover	Usage
00-vSwitch0-A	A	Disabled	In-band Management
01-vSwitch0-B	B	Disabled	In-band Management
02-vDS0-A	A	Disabled	VM-Traffic, vMotion
03-vDS0-B	B	Disabled	VM-Traffic, vMotion
04-iSCSI-A	A	Disabled	iSCSI-A, NVMe-TCP-A
05-iSCSI-B	B	Disabled	iSCSI-B, NVMe-TCP-B


## IP-based SAN multipathing

For IP-based SAN multipathing, two independent SAN fabric are available in the solution: one using FI A and the other using FI B. The VIC 15230 card in the X215c M8 compute node provides the server's internal connectivity to both FI A and FI B using vNIC paths pinned to the respective fabrics.

The ASA A50 dual controller storage system has two controllers: controller A and controller B. Each controller is connected to both FI A and FI B for IP-based SAN connectivity as shown in Figure 15. With symmetric active-active multipathing support in ASA, an iSCSI LUN can be accessed from both controllers. When one iSCSI LIF per fabric is created on each storage controller, there should be four total active/optimized paths to each iSCSI LUN.

Thanks to the design of having two separate fabrics, dual controllers, and having access to a LUN through multiple paths, IO can utilize remaining available paths during various single-point-of-failure scenarios and life-cycle management procedures such as having a single cable failure or rebooting a fabric interconnect or storage controller after software / firmware upgrade.

**Figure 15 SAN multipathing**

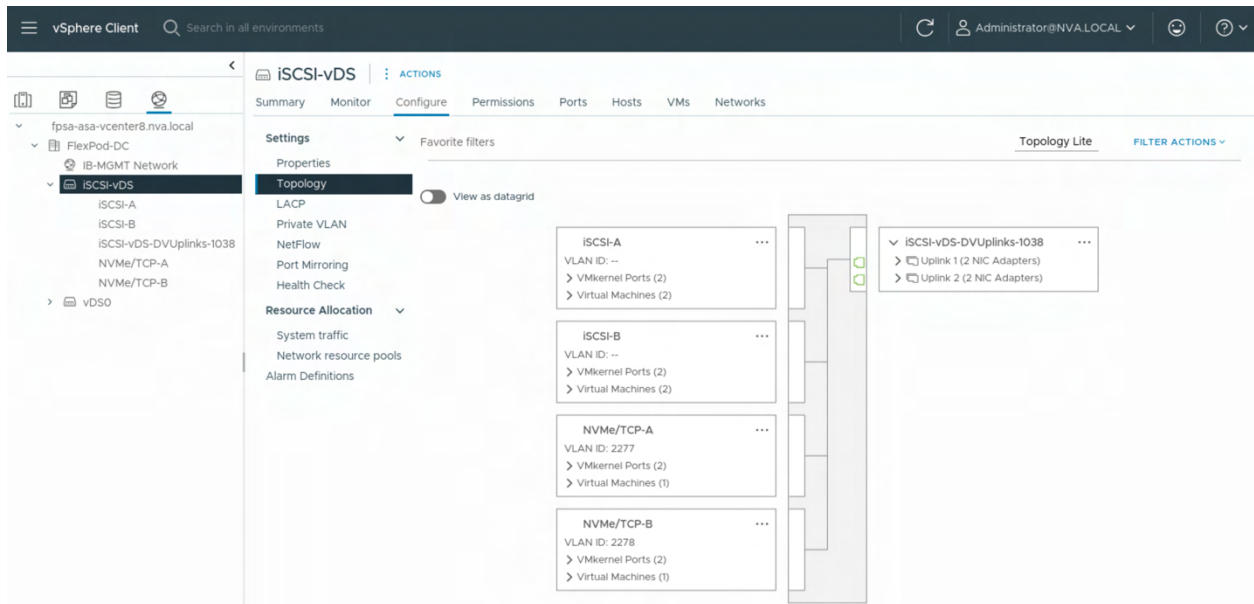


**Note:** SAN multipathing works similarly when using NVMe/TCP protocol.

## VMware distributed port groups for iSCSI

To provide iSCSI network access, a distributed virtual switch, iSCSI-vDS, is defined in the VMware cluster with two uplink adapters for accessing iSCSI-A and iSCSI-B network respectively via FI A and FI B. iSCSI-A and iSCSI-B distributed port groups are created for iSCSI traffic and NVMe/TCP-A and NVMe/TCP-B port groups are created for NVMe/TCP traffic.

When direct iSCSI storage access is needed by using the software iSCSI initiator within a VM, the VM needs its vNICs properly connected to the respective distributed iSCSI port groups.



**Note:** iSCSI-A and iSCSI-B VLANs are configured as the native VLANs for the respective uplink adapters in UCS for iSCSI SAN boot. As a result, the VLAN ID for the iSCSI port groups are not configured in VMware. However, when using NVMe/TCP protocol on the same iSCSI vNICs, appropriate NVMe/TCP VLANs need to be configured for the NVMe/TCP port groups.

**Note:** If it is desirable to have separate vNICs for NVMe/TCP traffic, then the NVMe/TCP network design can mimic the iSCSI network design and additional NVMe-vDS can be created with two NVMe/TCP vNIC uplinks. In that case, the NVMe/TCP-A and NVMe/TCP-B distributed port groups

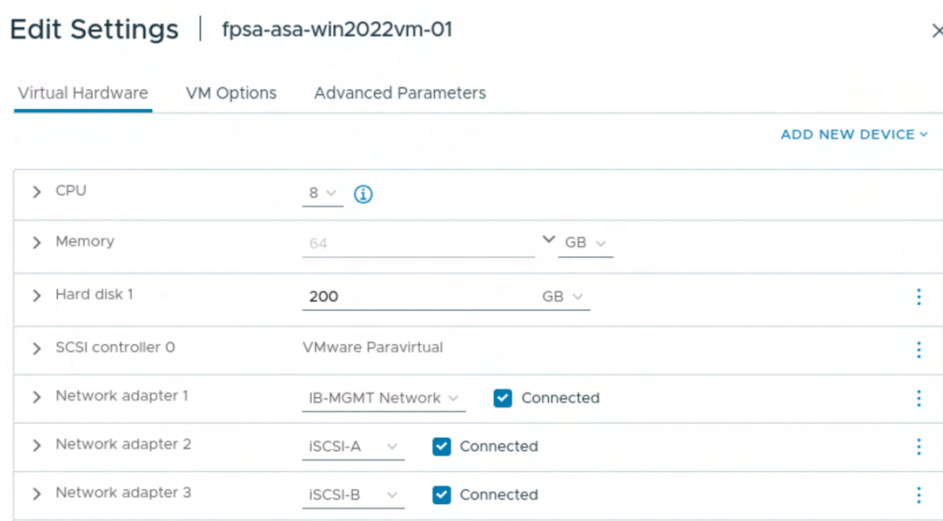
will be connected to the NVMe-vDS to logically separate NVMe/TCP traffic from iSCSI traffic from vNIC and vDS perspectives.

## Virtual machine vNICs for direct iSCSI storage access

For VMs to directly access the ASA storage using the iSCSI initiator inside the VM, attentions are required to properly connect its vNICs to the appropriate port groups configured in vCenter for VMware networking configuration to provide the underlying physical iSCSI network access.

For example, Figure 16 shows the vNIC configuration for the Windows 2022 VM with SQL server 2022. It includes three network adapters, one for in-band management access and two for direct iSCSI storage access from within the VM. The vNICs used for iSCSI traffic (adapters 2 and 3) are connected to the two distributed iSCSI port groups to access FI A and FI B, respectively, for iSCSI A and iSCSI B network connectivity.

**Figure 16 Virtual NIC configuration for direct iSCSI storage access**



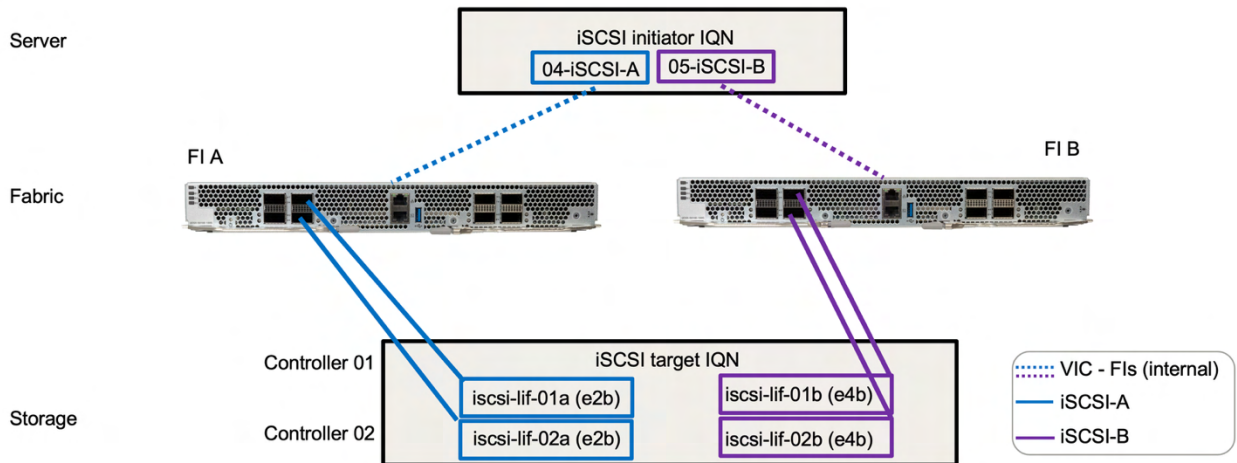
## Enable client access to storage

For the ASA storage systems, it is very easy to provision storage units for client access. The process is similar whether you are using SCSI LUNs with iSCSI protocol or NVMe namespaces with NVMe/TCP protocol. We will provide a brief overview of enabling client access to SCSI LUNs here. Please refer to Appendix F for discussions on client access to NVMe namespaces.

Figure 17 shows the physical paths for the server's iSCSI initiator to communicate through the two iSCSI vNICs with the iSCSI target portals on the storage system. FI A and FI B provides the two SAN fabrics for multipathing access to storage. On the storage cluster, logical interfaces (LIFs) are created on the respective VLAN ports of the underlying physical connections (e2b and e4b ports). The creation of vNICs on the server is accomplished with server profile derived from server profile template. The creation of the iSCSI target LIFs on the ASA storage can be accomplished by using the ONTAP System Manager or ONTAP CLI.



**Figure 17 iSCSI protocol physical access to storage**



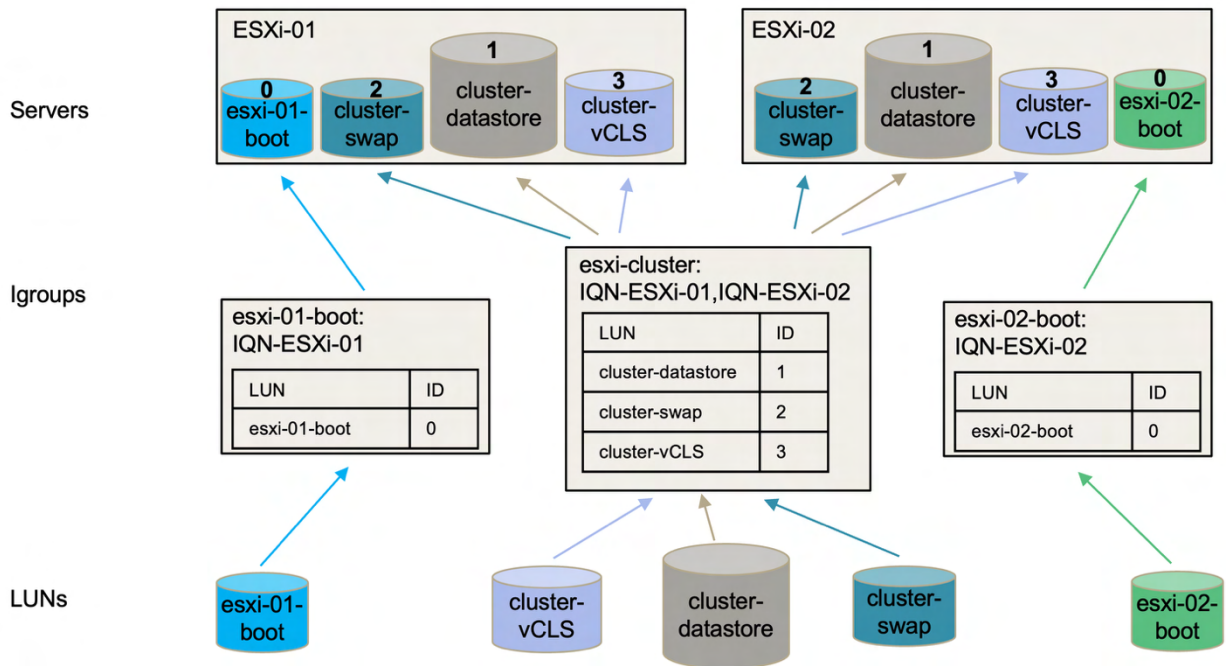
**Note:** An iSCSI LIF on a controller port can be configured to automatically failover to the controller's HA partner when the port goes down or when controller failover happens. LIF failover provides enhanced IO availability for these failure scenarios. The default IP LIF, which supports both iSCSI and NVMe/TCP protocols, does not support LIF failover. For the solution validation, we will configure iSCSI-only LIFs for iSCSI protocol usage as failover is supported for iSCSI LIFs. Due to utilizing different VLANs for traffic separation, we will also configure separate NVMe/TCP-only LIFs, which do not support failover, for NVMe/TCP protocol usage.

To ensure proper physical connectivity, network ping tool with jumbo frame setting can be used to verify connectivity and consistent jumbo frame configuration of the solution components. Configurations of the vNICs and iSCSI initiator are needed in the operating system, along with multipathing support, to discover iSCSI storage target and the available LUNs provisioned for the client.

After LUNs are created by using System Manager or ONTAP CLI, they can be mapped for client iSCSI initiator access using initiator groups (igroups). Both ONTAP System Manager and ONTAP CLI can be used to create initiator groups.

For a VMware cluster with SAN booted ESXi hosts, host specific igroups are created to map their respective SAN boot LUNs as LUN 0. The shared datastore LUNs are mapped to all hosts in the VMware cluster by using a cluster igroup which includes the initiator IQNs of all hosts in the cluster. The igroup mapping of LUNs are illustrated in Figure 18.

**Figure 18 Igroup mappings of LUNs to initiators**



**Note:** ONTAP tools for VMware vSphere which integrates ONTAP storage with VMware cluster creates igroup it uses to allow storage provisioning directly from vCenter.

**Note:** For the new ASA storage systems, operations and configurations have been simplified and optimized for SAN. The cluster initialization process creates a default storage virtual machine svm1 and users can start to create LUNs without needing to configure other objects such as aggregates and volumes after the cluster is initialized. Please refer to [ASA documentation](#) for additional details.

## Uplink switch configurations

This FlexPod IP-based SAN solution validation utilizes a configuration which directly attaches NetApp ASA A50 storage to the UCS X-Series Direct without having Nexus switches in between the X-Series Direct compute and the ASA A50 storage system. To enable access to the FlexPod environment, the fabric interconnects in the UCS X-Series Direct are connected to a pair of existing Nexus switches for uplink. If you are deploying a new pair of Nexus switches, please refer to FlexPod Datacenter CVDs such as the example below for reference on Nexus switch deployment procedures.

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_base\\_imm\\_manual\\_deploy.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_base_imm_manual_deploy.html)

## Check uplink Nexus switch features and configurations

For compatibility, please check that the following features and configurations are enabled and implemented on the uplink Nexus switches.

```
feature nxapi
feature udd
feature interface-vlan
feature lacp
feature vpc
feature lldp
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
```



```
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management - Repeat this command to add additional NTP
servers
clock timezone <timezone> <hour-offset> <minute-offset>
clock summer-time <timezone> <start-week> <start-day> <start-month> <start-time> <end-week> <end-
day> <end-month> <end-time> <offset-minutes>
```

## Configure VLANs

For external access to VLANs configured in the solution, such as for the in-band management IP addresses of the ESXi host and the VMs deployed on the VMware solution, the FI to Nexus switch uplinks will need to have those VLANs configured.

From the global configuration mode, run the VLAN commands on both uplink Nexus switches to define the VLANs. Use Table 5 for reference and adapt it for your specific deployment environment.

```
vlan <vlan-id>
name <VLAN-NAME>
```

## Configure port channel and vPC for FI ports

For the port channel (PC) configurations, FI A utilizes port channel 21 and FI B utilizes port channel 22 in the validation environment. Perform the port channel and vPC related configuration in global configuration mode on both Nexus switches.

```
interface port-channel21
description <ucs-cluster-name>-a
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <allowed-vlan-ranges>
spanning-tree port type edge trunk
mtu 9216
vpc 21

interface port-channel22
description <ucs-cluster-name>-b
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <allowed-vlan-ranges>
spanning-tree port type edge trunk
mtu 9216
vpc 22
```

**Note:** For the allowed VLAN ranges, we configured it to allow all non-native VLANs shown in Table 5 used in the solution except for the OOB-Mgmt VLAN which utilizes other existing switching infrastructure. Configure VLANs and port channel IDs as appropriate for your deployment.

## Configure ports and enable UDLD for UCS interfaces

Configure Cisco UCS Fabric Interconnect 9108 100G cluster name <ucs-clustername>, interfaces, and port channels in the uplink switches as appropriate for your deployment in global configuration mode.

```
Switch A:

interface Eth<FI-A-uplink-port>
description <ucs-cluster-name>-a:1/5
udld enable
channel-group 21 mode active
no shutdown
interface Eth<FI-B-uplink-port>
description <ucs-cluster-name>-b:1/5
udld enable
channel-group 22 mode active
no shutdown
```

Switch B:

```
interface Eth<FI-A-uplink-port>
  description <ucs-cluster-name>-a:1/6
  udld enable
  channel-group 21 mode active
  no shutdown
interface Eth<FI-B-uplink-port>
  description <ucs-cluster-name>-b:1/6
  udld enable
  channel-group 22 mode active
  no shutdown
```

## Add NTP distribution interface if needed

From global configuration mode, configure both uplink Nexus switches for NTP distribution if NTP is not already configured and no existing NTP servers are available for use.

Switch A:

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-b-mgmt0-ip> use-vrf management
```

Switch B:

```
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
ntp peer <switch-a-mgmt0-ip> use-vrf management
```

## Save configuration

Save the running configuration as startup configuration on both switches.

```
copy running-config startup-config
```

## Check switch configurations

The following commands can be used to check for correct switch configuration.

```
show run
show vpc
show port-channel summary
show ntp peer-status
show cdp neighbors
show lldp neighbors
show udld neighbors
show run int
show int
show int status
```

**Note:** Some of these commands need to be run again after further configuration of the FlexPod components are complete to see complete results.

## ONTAP cluster setup

### Pre-requisites

- Cluster interconnect cabled based on supported options in NetApp Hardware Universe
- ONTAP license

- Network adapters appropriate for iSCSI and NVMe/TCP protocols

## Create ONTAP cluster

Before configuring ONTAP nodes, review the [installation and setup workflow](#) for the new ASA storage systems to review the hardware requirements, prepare your site, install and cable the hardware components, and setup your ONTAP cluster. Please refer to Table 7 for a list of ONTAP cluster installation and configuration information.

**Table 7 ONTAP installation and configuration information**

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01_mgmt_ip>
Cluster node 01 netmask	<node01_mgmt_mask>
Cluster node 01 gateway	<node01_mgmt_gateway>
Cluster node 02 IP address	<node02_mgmt_ip>
Cluster node 02 netmask	<node02_mgmt_mask>
Cluster node 02 gateway	<node02_mgmt_gateway>
ONTAP 9.16.1 URL	<url_boot_software>
Name for cluster	<clustername>
Cluster administrator password	<clustermgmt_password>
Cluster management IP address	<clustermgmt_ip>
Cluster management gateway	<clustermgmt_gateway>
Cluster management netmask	<clustermgmt_mask>
ONTAP cluster license	<license_file>
Domain name	<domain_name>
DNS server IP (you can enter more than one)	<dns_server_ip>
NTP server IP (you can enter more than one)	<ntp_server_ip>
Controller location	<controller_location>

**Note:** The ASA A50 system used in this FlexPod solution validation is set up in a two-node switchless cluster configuration.

To initialize node 01 (controller A) and node 02 (controller B), use two serial console port program sessions to communicate with the storage controller A and controller B, respectively.

## Initialize node 01

To initialize node 01, complete the following steps:

Connect to the storage system console port. You should see a LOADER-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

Allow the system to boot up.

```
autoboot
```

Press Ctrl-C to enter the Boot menu.

**Note:** If ONTAP 9.16.1P3 or later is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.16.1P3 or later is the version being booted, continue with step 14. It is recommended that the latest ONTAP release and patch available from [NetApp support site](#) is used.

To install new software, select option 7.

Enter `y` to perform an upgrade.

Select `e0M` for the network port you want to use for the download.

Enter `y` to reboot now.

Enter the IP address, network mask, and default gateway for `e0M` in their respective places.

```
Enter the IP address for port e0M: <node01-mgmt-ip>
Enter the netmask for port e0M: <node01-mgmt-mask>
Enter the IP address of the default gateway: <node01-mgmt-gateway>
```

Enter the URL where the software can be found.

**Note:** This web server must be pingable from the node.

```
<url_boot_software>
```

Press Enter for the username, indicating no username.

Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

Enter `y` to reboot the node.

**Note:** When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the loader prompt. If these actions occur, the system might deviate from this procedure.

Press Ctrl-C to enter the Boot menu.

Select option 4 to Initialize and configure system.

Enter `y` to zero disks, reset config, and install a new file system.

Enter `y` to erase all the data on the disks.

**Note:** The initialization can take several minutes. When initialization is complete, the storage system reboots. You can continue with the node 02 configuration while the disks for node 01 are zeroing.

While node 01 is initializing, begin the initializing procedures for node 02.

## Initialize node 02

To initialize node 02, complete the following steps:

1. Connect to the storage system console port. You should see a LOADER-B prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

Allow the system to boot.

```
autoboot
```

Press Ctrl-C to enter the Boot menu.

**Note:** If ONTAP 9.16.1P3 or later is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.16.1P3 or later is the version being booted, continue with step 14. It is recommended that the latest ONTAP release and patch available from [NetApp support site](#) is used.

To install new software, select option 7.

Enter `y` to perform an upgrade.

Select `e0M` for the network port you want to use for the download.

Enter `y` to reboot now.

Enter the IP address, network mask, and default gateway for `e0M` in their respective places.

```
Enter the IP address for port e0M: <node02-mgmt-ip>
Enter the netmask for port e0M: <node02-mgmt-mask>
Enter the IP address of the default gateway: <node02-mgmt-gateway>
```

Enter the URL where the software can be found.

**Note:** This web server must be pingable.

```
<url_boot_software>
```

Press Enter for the username, indicating no username.

Enter *y* to set the newly installed software as the default to be used for subsequent reboots.

Enter *y* to reboot the node.

**Note:** When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-B prompt. If these actions occur, the system might deviate from this procedure.

Press Ctrl-C to enter the Boot menu.

Select option 4 for Initialize and configure system.

Enter *y* to zero disks, reset config, and install a new file system.

Enter *y* to erase all the data on the disks.

**Note:** The initialization can take several minutes.

## Initialize cluster

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP boots on the node for the first time.

**Note:** If the Initialize node step above stopped automatically at LOADER prompt, enter autoboot to boot the node.

To set up the first cluster node, follow these steps:

1. Follow the prompts to set up node 01.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
```

```
This system will send event messages and periodic reports to NetApp Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
```

```
Enabling AutoSupport can significantly speed problem determination
and resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <node01-mgmt-ip>
Enter the node management interface netmask: <node01-mgmt-mask>
Enter the node management interface default gateway: <node01-mgmt-gateway>
A node management interface on port e0M with IP address <node01-mgmt-ip> has been created.
```

Use your web browser to complete cluster setup by accessing `https://<node01-mgmt-ip>`

Otherwise, press Enter to complete cluster setup using the command line interface:

To complete cluster setup, open a web browser and navigate to `https://<node01-mgmt-ip>`.

**Note:** Cluster setup can also be performed using the CLI. This document describes the cluster setup using the NetApp ONTAP System Manager guided setup.

Create a strong password for the admin account and click Continue.

The screenshot shows the NetApp ONTAP System Manager interface. On the left, the 'Health' section indicates 'Found 2 healthy nodes' and shows a diagram of two nodes labeled 'ASA-A50'. The main panel is titled 'Initialize storage system'. It contains a 'Password' section with a prompt to 'Create a strong password for the admin account.' and two password input fields, followed by a 'Continue' button. Below this are three expandable sections: 'Network addresses', 'Network services', and 'Encryption'. At the bottom of the main panel is an 'Initialize' button.

In the Network addresses section, provide the storage system name, management address, subnet mask, gateway, node 02 IP, and click Continue.

In the Network services section, check the boxes for Resolve host names using DNS and Synchronize times using NTP, and click Add to add the respective DNS domain, name server(s), and NTP server(s) information, and then click Continue.

**Note:** For redundancy, specify multiple servers for each service.

In the Encryption section, encrypt data at rest (hardware encryption) is enabled with onboard key manager by default. You can optionally specify an external key manager instead. Enter and repeat the passphrase for the onboard key manager, and then click Continue.

**Note:** Be sure to save the encryption passphrase information for the key manager securely for future recovery needs.

Click Initialize at the bottom of the Initialize storage system screen.

**Note:** The cluster configuration process might take several minutes to complete. Don't reload the browser while it is being configured to avoid misconfigurations.

## Apply ONTAP license

You can install license via NetApp License File (NLF) using NetApp ONTAP System Manager and the following steps.

1. Login to ONTAP System Manager with the storage/cluster management IP.

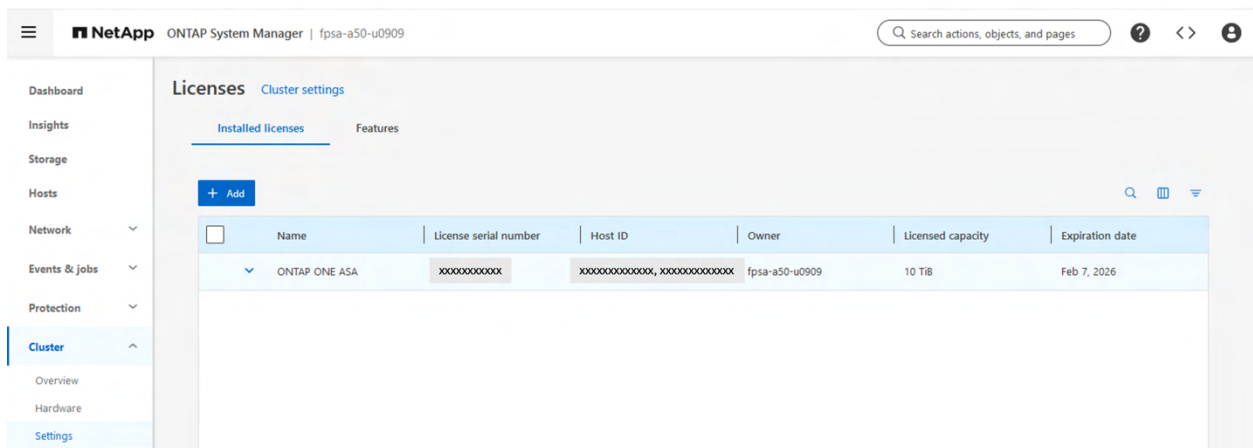
Select Cluster > Settings.

Under Licenses, click the arrow icon to go into the Licenses page.

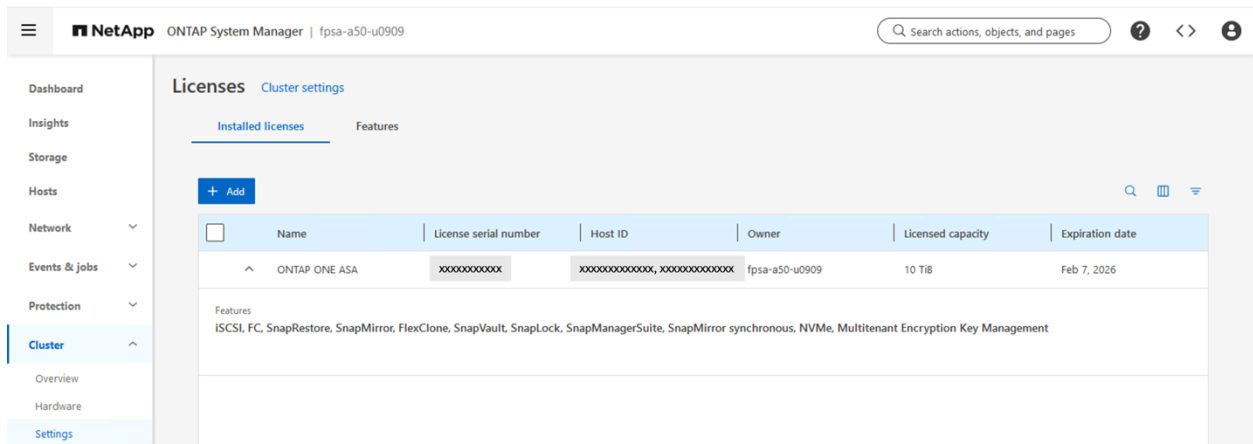
Click +Add icon.

Select Browse to select the NLF you downloaded.

The installed License shows up along with the controller SN, cluster name, licensed capacity, and expiration date.



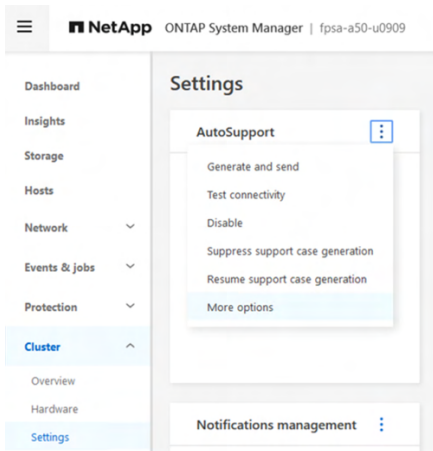
Click on the Down arrow next to the license name to see the list of licensed features.



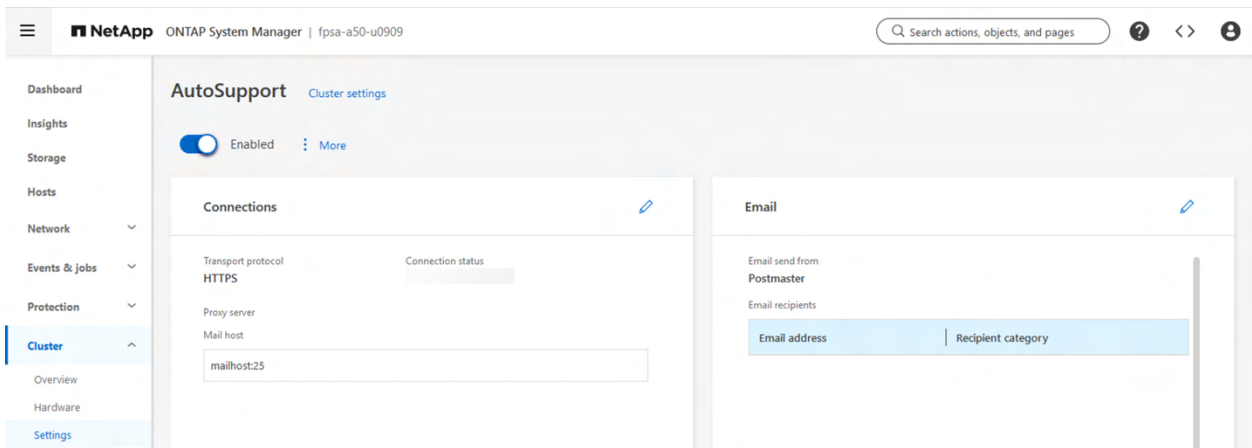
## Enable AutoSupport

To enable AutoSupport, follow the steps below.

1. Select the Settings menu under the Cluster menu.
2. Click the ellipsis in the AutoSupport tile and select More options.

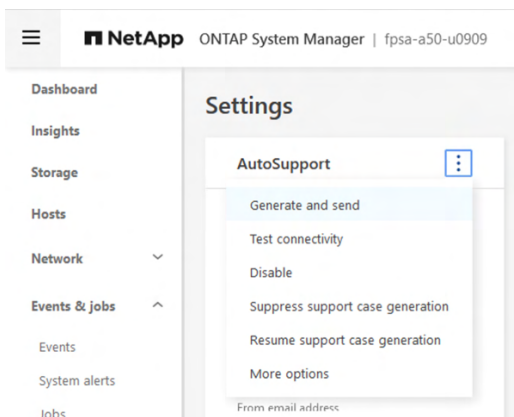


3. In the Connections tile, click Edit icon to change the transport protocol, add a proxy server address and a mail host as needed.

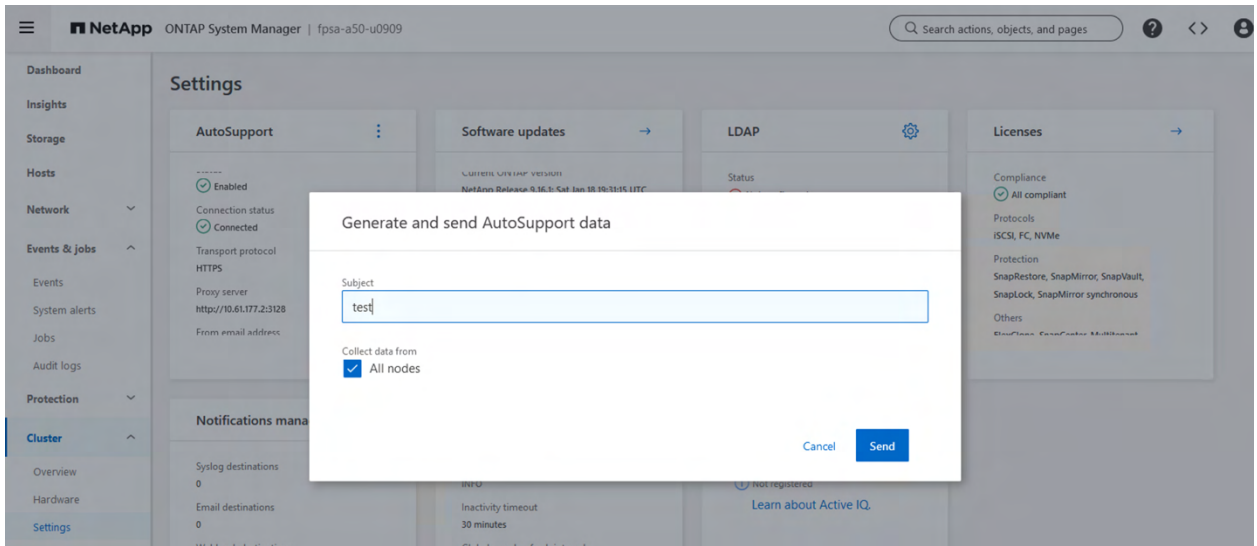


4. Click Save.
5. In the Email tile, click Edit icon to enter the desired email sender, email recipient, and select a recipient category from the drop-down list.
6. Click Save.

**Note:** You can test the AutoSupport configuration by invoking the Generate and Send action from the AutoSupport tile. In the dialog box for Generate and send AutoSupport data, provide a Subject and click Send. After a few minutes, check for an email from NetApp with the subject of NetApp Automated AutoSupport Acknowledgement to confirm the proper AutoSupport configuration.

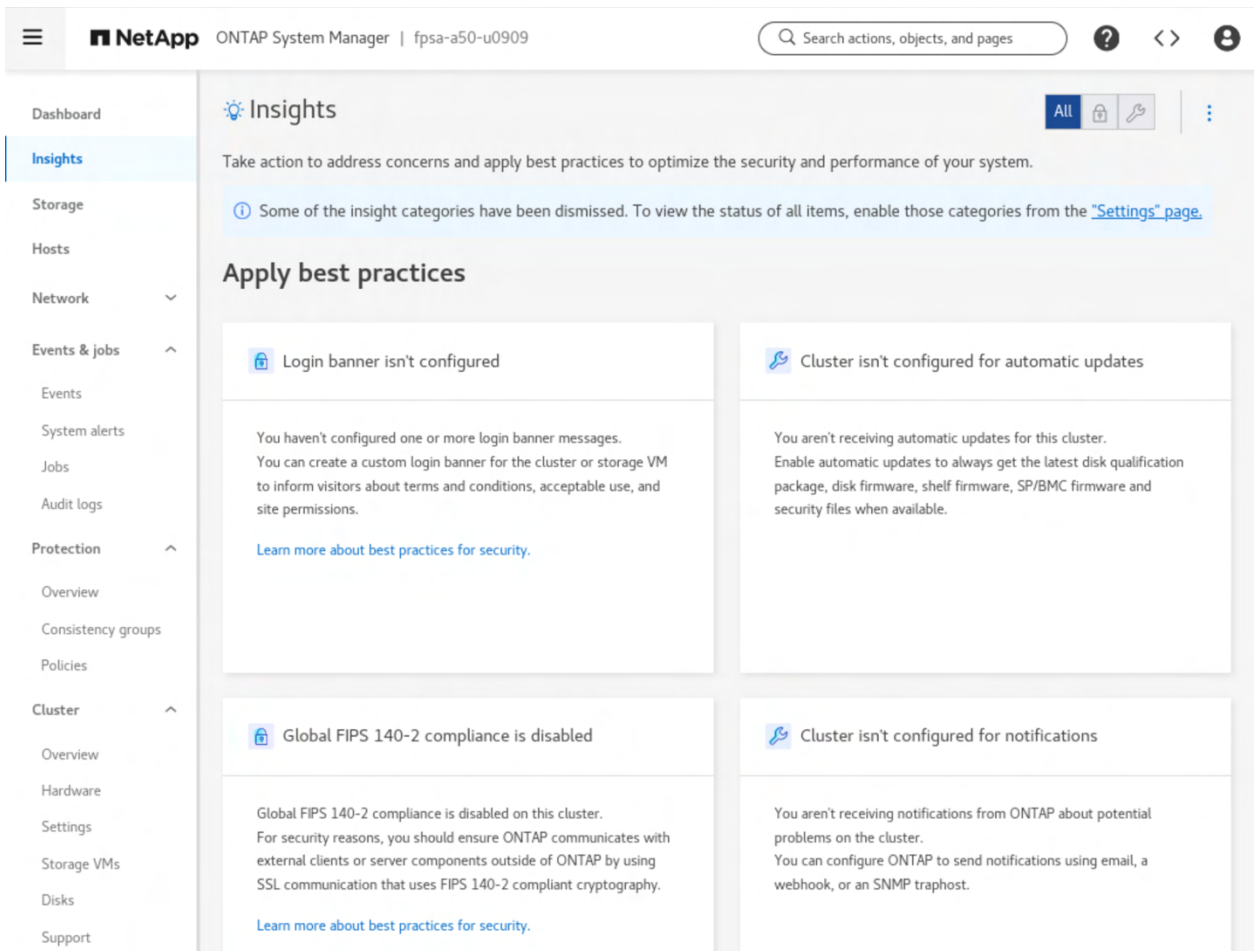




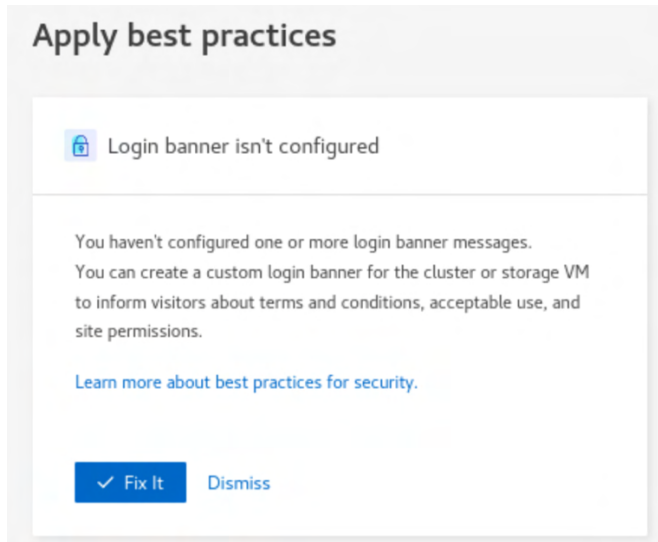


## Apply ONTAP best practices

After the ONTAP cluster is initialized, you should take further actions to apply best practices highlighted in the Insights pane as shown in the screenshot below to optimize your ONTAP system deployment.



For example, you can hover over the Login banner isn't configured issue and click on the Fix It button to address the issue.



After clicking on the Fix It icon, the Login banner message dialog appears for you to enter a login banner message. You can choose to apply the entered message to the cluster login and the automatically created storage VM svm1 login. See the screenshot below for an example.

Login banner message

---

Login banners allow an organization to present operators, administrators, and even miscreants with terms and conditions of acceptable use, and they indicate who is permitted access to the system. This approach helps establish expectations for access and use of the system.

Access restricted to authorized users ONLY!

☒ Apply to cluster login

☒ Apply to storage VM login

[View all storage VMs](#)

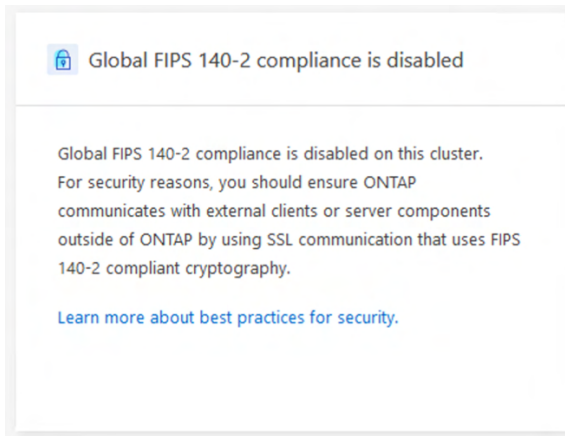
☒ Storage VMs

☒ svm1

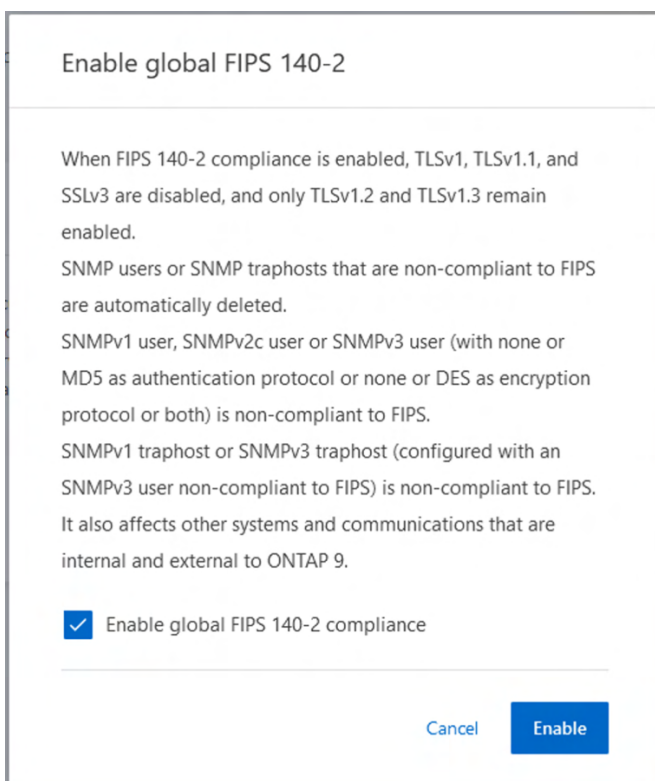
Cancel

Enable

For security reasons, you should ensure ONTAP communicates with external clients or server components by using SSL communication that uses FIPS 140-2 compliant cryptography.



You can hover over the Global FIPS 140-2 compliance is disabled issue and click on the Fix It button to address this issue. Check the box to enable Global FIPS 140-2 compliance after reviewing the messages in the dialog and noting the caveats and automatic actions ONTAP will be taking to remove non-compliant SNMP users or SNMP traphosts when FIPS 140-2 is enabled. Click Enable to proceed.



For additional information on hardening your FlexPod solution and deploying a FlexPod solution utilizing zero-trust framework, please refer to the [FlexPod Security Hardening](#) technical report and [FlexPod Datacenter Zero Trust Framework CVD](#).

## Configure Timezone

To configure time synchronization on the cluster, follow these steps:

1. Set the timezone for the cluster.

```
timezone -timezone <timezone>
```

**Note:** For example, in the eastern United States, the time zone is America/New\_York.

Check the timezone setting.

```
date
```

## Configure service processor network interface

1. To assign a static IPv4 address to the Service Processor on each node, run the following commands:

```
system service-processor network modify -node <node01> -address-family IPv4 -enable true -dhcp none -ip-address <node01-sp-ip> -netmask <node01-sp-mask> -gateway <node01-sp-gateway>
```

```
system service-processor network modify -node <node02> -address-family IPv4 -enable true -dhcp none -ip-address <node02-sp-ip> -netmask <node02-sp-mask> -gateway <node02-sp-gateway>
```

**Note:** The service processor IP addresses should be in the same subnet as the node management IP addresses.

Check the service process configuration.

```
system service-processor show
```

## Disable flow control on data ports

1. Run the following command to disable flow control on 25/100G data ports on node01 and node 02:

```
network port modify -node <node01> -port e2b,e4b -flowcontrol-admin none
network port modify -node <node02> -port e2b,e4b -flowcontrol-admin none
```

**Note:** Disable flow control only on ports that are used for data traffic. Adjust the list of ports to match your deployment environment.

## Enable network discovery protocols

1. To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command:

```
node run -node * options cdpd.enable on
```

To enable the Link-layer Discovery Protocol (LLDP) on the NetApp storage controllers, run the following command:

```
node run * options lldp.enable on
```

## Create broadcast domains

IPspaces are distinct IP address spaces that can be utilized for secure multi-tenant deployment to support overlapping IP addresses from the same IP address range. Broadcast domains are intended to group network ports that belong to the same layer 2 network. The ports in the group can then be used by a storage virtual machine (SVM) for data or management traffic.

During cluster initialization, the Cluster IPspace, Cluster broadcast domain, Default IPspace, and Default broadcast domain are created automatically with ports used for cluster communication and cluster management as shown in the example output below.

```
fpsa-a50-u0909:> broadcast-domain show
(network port broadcast-domain show)
IPspace Broadcast
Name      Domain Name      MTU   Port List      Update
-----
Cluster Cluster      9000
```

			fpsa-a50-u0909-02:e2a	complete
			fpsa-a50-u0909-02:e4a	complete
			fpsa-a50-u0909-01:e2a	complete
			fpsa-a50-u0909-01:e4a	complete
Default	Default	1500		
			fpsa-a50-u0909-02:e0M	complete
			fpsa-a50-u0909-01:e0M	complete
2 entries were displayed.				

For this single-tenant deployment, we are creating additional broadcast domains in the Default IPspace for iSCSI data and NVMe/TCP data traffic.

1. To create iSCSI-A and iSCSI-B data broadcast domains for iSCSI data traffic with a maximum transmission unit (MTU) of 9000, run the following commands:

```
network port broadcast-domain create -ipspace Default -broadcast-domain iSCSI-A -mtu 9000
network port broadcast-domain create -ipspace Default -broadcast-domain iSCSI-B -mtu 9000
```

To create NVMe-TCP-A and NVMe-TCP-B data broadcast domains for NVMe/TCP data traffic with a maximum transmission unit (MTU) of 9000, run the following commands:

```
network port broadcast-domain create -ipspace Default -broadcast-domain NVMe-TCP-A -mtu 9000
network port broadcast-domain create -ipspace Default -broadcast-domain NVMe-TCP-B -mtu 9000
```

**Note:** Please refer to the [FlexPod Datacenter Zero Trust Framework CVD](#) for details on secure multi-tenant deployment using separate IPspaces and VLANs for each tenant to create separate administrative network domains for tenants to access their data stored in the cluster,

## Remove default broadcast domains

By default, all network ports are included in separate default broadcast domain. Network ports used for data services (for example, e2b, e4b, and so on) should be removed from their default broadcast domain and that broadcast domain should be deleted.

```
fpsa-a50-u0909::> broadcast-domain show
(network port broadcast-domain show)
IPspace Broadcast
Name      Domain Name      MTU    Port List
-----
Cluster   Cluster           9000   fpsa-a50-u0909-02:e2a
                                     fpsa-a50-u0909-02:e4a
                                     fpsa-a50-u0909-01:e2a
                                     fpsa-a50-u0909-01:e4a
                                     complete
                                     complete
                                     complete
                                     complete
Default   Default           1500   fpsa-a50-u0909-02:e0M
                                     fpsa-a50-u0909-01:e0M
                                     complete
                                     complete
          Default-1       9000   fpsa-a50-u0909-01:e2b
                                     complete
          Default-2       9000   fpsa-a50-u0909-02:e2b
                                     complete
          Default-3       9000   fpsa-a50-u0909-01:e4b
                                     complete
          Default-4       9000   fpsa-a50-u0909-02:e4b
                                     complete
...
```

1. To perform this task, run the following commands:

```
broadcast-domain delete -broadcast-domain <Default-N> -ipspace Default
broadcast-domain show
```

**Note:** Delete the Default broadcast domains with Network ports (Default-1, Default-2, and so on). This does not include Cluster ports and management ports.

```
fpsa-a50-u0909::> broadcast-domain delete -broadcast-domain Default-1,Default-2,Default-3,Default-4
(network port broadcast-domain delete)
4 entries were deleted.
```

## Configure VLANs on data ports

You can use VLANs in ONTAP to provide logical segmentation of networks by creating separate broadcast domains that are defined on a port. Here we create separate VLANs for iSCSI and NVMe/TCP data traffic.

1. To create the iSCSI-A and iSCSI-B VLAN ports and add them to their respective broadcast domains, run the following commands:

```
vlan create -node <ndoe01> -vlan-name e2b-<iscsi-a-vlan-id>
vlan create -node <node01> -vlan-name e4b-<iscsi-b-vlan-id>
vlan create -node <ndoe02> -vlan-name e2b-<iscsi-a-vlan-id>
vlan create -node <node02> -vlan-name e4b-<iscsi-b-vlan-id>

broadcast-domain add-ports -broadcast-domain iSCSI-A -ports <node01>:e2b-<iscsi-a-vlan-id>,<node02>:e2b-<iscsi-a-vlan-id>
broadcast-domain add-ports -broadcast-domain iSCSI-B -ports <node01>:e4b-<iscsi-b-vlan-id>,<node02>:e4b-<iscsi-b-vlan-id>

fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-01 -vlan-name e2b-2277
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-01 -vlan-name e4b-2278
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-02 -vlan-name e2b-2277
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-02 -vlan-name e4b-2278

fpsa-a50-u0909::> broadcast-domain add-ports -broadcast-domain iSCSI-A -ports fpsa-a50-u0909-01:e2b-2273,fpsa-a50-u0909-02:e2b-2273
(network port broadcast-domain add-ports)
fpsa-a50-u0909::> broadcast-domain add-ports -broadcast-domain iSCSI-B -ports fpsa-a50-u0909-01:e4b-2274,fpsa-a50-u0909-02:e4b-2274
(network port broadcast-domain add-ports)
```

- To create the NVMe-TCP-A and NVMe-TCP-B VLAN ports and add them to their respective NVMe-TCP-A and NVMe-TCP-B broadcast domains, run the following commands:

```
vlan create -node <ndoe01> -vlan-name e2b-<nvme-tcp-a-vlan-id>
vlan create -node <node01> -vlan-name e4b-<nvme-tcp-b-vlan-id>
vlan create -node <ndoe02> -vlan-name e2b-<nvme-tcp-a-vlan-id>
vlan create -node <node02> -vlan-name e4b-<nvme-tcp-b-vlan-id>

broadcast-domain add-ports -broadcast-domain NVMe-TCP-A -ports <node01>:e2b-<nvme-tcp-a-vlan-id>,<node02>:e2b-<nvme-tcp-a-vlan-id>
broadcast-domain add-ports -broadcast-domain NVMe-TCP-B -ports <node01>:e4b-<nvme-tcp-b-vlan-id>,<node02>:e4b-<nvme-tcp-b-vlan-id>

fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-01 -vlan-name e2b-2277
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-01 -vlan-name e4b-2278
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-02 -vlan-name e2b-2277
fpsa-a50-u0909::> vlan create -node fpsa-a50-u0909-02 -vlan-name e4b-2278

fpsa-a50-u0909::> broadcast-domain add-ports -broadcast-domain NVMe-TCP-A -ports fpsa-a50-u0909-01:e2b-2277,fpsa-a50-u0909-02:e2b-2277
(network port broadcast-domain add-ports)
fpsa-a50-u0909::> broadcast-domain add-ports -broadcast-domain NVMe-TCP-B -ports fpsa-a50-u0909-01:e4b-2278,fpsa-a50-u0909-02:e4b-2278
(network port broadcast-domain add-ports)
```

**Note:** For the above commands, adjust the physical port number based on your specific deployment. Here is an example from the deployment.

```
fpsa-a50-u0909::> broadcast-domain show
(network port broadcast-domain show)
IPspace Broadcast
Name      Domain Name      MTU  Port List
Update
Status Details
```



Cluster	Cluster	9000	fpsa-a50-u0909-02:e2a	complete
			fpsa-a50-u0909-02:e4a	complete
			fpsa-a50-u0909-01:e2a	complete
			fpsa-a50-u0909-01:e4a	complete
Default	Default	1500	fpsa-a50-u0909-02:e0M	complete
			fpsa-a50-u0909-01:e0M	complete
	NVMe-TCP-A	9000	fpsa-a50-u0909-02:e2b-2277	complete
			fpsa-a50-u0909-01:e2b-2277	complete
	NVMe-TCP-B	9000	fpsa-a50-u0909-02:e4b-2278	complete
			fpsa-a50-u0909-01:e4b-2278	complete
	iSCSI-A	9000	fpsa-a50-u0909-02:e2b-2273	complete
			fpsa-a50-u0909-01:e2b-2273	complete
	iSCSI-B	9000	fpsa-a50-u0909-02:e4b-2274	complete
			fpsa-a50-u0909-01:e4b-2274	complete

6 entries were displayed.

## Verify protocol support by the storage virtual machine

1. To verify that the storage protocols needed are already enabled on the automatically created vservers svm1, run the following command:

```
vserver show-protocols -vserver svm1
vserver iscsi show
vserver fcp show
```

**Note:** By default, all the block storage protocols are supported by the storage virtual machine svm1 which is created during cluster initialization. See below for example outputs.

```
fpsa-a50-u0909:> vserver show-protocols -vserver svm1

Vserver: svm1
Protocols: fcp, iscsi, nvme

fpsa-a50-u0909:> vserver iscsi show
Vserver      Target      Target      Status
Name         Alias      Admin
-----
svm1         iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2  svm1 up

fpsa-a50-u0909:> vserver fcp show
Vserver      Target Name      Status
Admin
-----
svm1         20:00:d0:39:ea:c6:a7:94  up
```

**Note:** The iSCSI target name will be needed for iSCSI SAN boot configuration in the UCS server profile.

## Create iSCSI LIFs

An iSCSI logical network interface (LIF) has an IP address and is associated with a physical or logical port in the storage controller for iSCSI communication. The default-data-blocks service-policy for IP LIFs enables the created LIFs to serve both iSCSI and NVMe/TCP protocols. However, since we have separate VLANs for iSCSI and NVMe/TCP protocols, separate LIFs are created to serve iSCSI and NVMe/TCP protocols and to configure iSCSI LIF failover for iSCSI-only LIFs. Here we are covering only iSCSI LIFs. Please see Appendices F, G, and H for ONTAP, VMware client, and Oracle Linux client NVMe/TCP configuration examples. To create iSCSI-only LIFs with LIF failover enabled, follow the steps below.

1. Create a custom iSCSI-only service-policy with core-data and core-iscsi services using advanced privilege as shown in the example below.

```
fpsa-a50-u0909::> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when directed to do so
by NetApp personnel.
Do you want to continue? {y|n}: y

fpsa-a50-u0909::*> network interface service-policy create -vserver svml -policy custom-data-iscsi -service data-core,data-iscsi

To verify:

fpsa-a50-u0909::*> network interface service-policy show -policy custom-data-iscsi
Vserver      Policy                Service: Allowed Addresses
-----
svml
              custom-data-iscsi          data-core: 0.0.0.0/0
                                      data-iscsi: 0.0.0.0/0

fpsa-a50-u0909::*> set -privilege admin
```

To create four iSCSI LIFs (two on each node, one for fabric A and the other for fabric B), follow the example below.

```
network interface create -vserver svml -lif iscsi-lif-01a -service-policy custom-data-iscsi -
home-node <node01> -home-port e2b-<iscsi-a-vlan-id> -address <node01-iscsi-a-ip> -netmask <iscsi-
a-mask> -failover-policy sfo-partner-only -status-admin up
network interface create -vserver svml -lif iscsi-lif-01b -service-policy custom-data-iscsi -
home-node <node01> -home-port e4b-<iscsi-b-vlan-id> -address <node01-iscsi-b-ip> -netmask <iscsi-
b-mask> -failover-policy sfo-partner-only -status-admin up
network interface create -vserver svml -lif iscsi-lif-02a -service-policy custom-data-iscsi -
home-node <node02> -home-port e2b-<iscsi-a-vlan-id> -address <node02-iscsi-a-ip> -netmask <iscsi-
a-mask> -failover-policy sfo-partner-only -status-admin up
network interface create -vserver svml -lif iscsi-lif-02b -service-policy custom-data-iscsi -
home-node <node02> -home-port e4b-<iscsi-b-vlan-id> -address <node02-iscsi-b-ip> -netmask <iscsi-
b-mask> -failover-policy sfo-partner-only -status-admin up
```

Example:

```
fpsa-a50-u0909::> network interface create -vserver svml -lif iscsi-lif-01a -service-policy
custom-data-iscsi -home-node fpsa-a50-u0909-01 -home-port e2b-2273 -address 172.22.73.101 -
netmask 255.255.255.0 -failover-policy sfo-partner-only -status-admin up
fpsa-a50-u0909::> network interface create -vserver svml -lif iscsi-lif-01b -service-policy
custom-data-iscsi -home-node fpsa-a50-u0909-01 -home-port e4b-2274 -address 172.22.74.101 -
netmask 255.255.255.0 -failover-policy sfo-partner-only -status-admin up
fpsa-a50-u0909::> network interface create -vserver svml -lif iscsi-lif-02a -service-policy
custom-data-iscsi -home-node fpsa-a50-u0909-02 -home-port e2b-2273 -address 172.22.73.102 -
netmask 255.255.255.0 -failover-policy sfo-partner-only -status-admin up
fpsa-a50-u0909::> network interface create -vserver svml -lif iscsi-lif-02b -service-policy
custom-data-iscsi -home-node fpsa-a50-u0909-02 -home-port e4b-2274 -address 172.22.74.102 -
netmask 255.255.255.0 -failover-policy sfo-partner-only -status-admin up
```

To verify:

```
fpsa-a50-u0909::> net int show -lif iscsi*
(network interface show)
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	iscsi-lif-01a	up/up	172.22.73.101/24	fpsa-a50-u0909-01	e2b-2273	true
	iscsi-lif-01b	up/up	172.22.74.101/24	fpsa-a50-u0909-01	e4b-2274	true
	iscsi-lif-02a	up/up	172.22.73.102/24	fpsa-a50-u0909-02	e2b-2273	true
	iscsi-lif-02b	up/up	172.22.74.102/24	fpsa-a50-u0909-02	e4b-2274	true

4 entries were displayed.

## Storage configuration for ESXi hosts

### Create initiator groups for iSCSI storage access

The initiator groups designate which hosts have permission to access specific LUNs on the storage system. Two categories of initiator groups are created. For iSCSI SAN boot, multiple initiator groups are created with one host iSCSI Qualified Name (IQN) in each group to map its SAN boot LUN. For shared LUN access, an initiator group is configured to include all IQNs for the hosts in the VMware cluster.

1. Run the following command on NetApp to create iscsi initiator groups (igroups).

```
lun igroup create -igroup <igroup-name> -protocol iscsi -ostype vmware -initiator <vm-host-iqn>

fpsa-a50-u0909:> lun igroup create -igroup FlexPod-ASA-esxi-01-boot-iscsi -protocol iscsi -
ostype vmware -initiator iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
fpsa-a50-u0909:> lun igroup create -igroup FlexPod-ASA-esxi-02-boot-iscsi -protocol iscsi -
ostype vmware -initiator iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2

fpsa-a50-u0909:> lun igroup create -igroup FlexPod-ASA-esxi-cluster-iscsi -protocol iscsi -
ostype vmware -initiator iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1,iqn.2010-
11.com.flexpod:flexpod-asa-ucshost:2

To verify:

fpsa-a50-u0909:> igroup show -protocol iscsi
Vserver   Igroup          Protocol OS Type  Initiators
-----
svm1      FlexPod-ASA-esxi-01-boot-iscsi iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
svm1      FlexPod-ASA-esxi-02-boot-iscsi iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
svm1      FlexPod-ASA-esxi-cluster-iscsi iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
                                                iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2

3 entries were displayed.
```

**Note:** With the new ASA storage system, the SVM parameter can be omitted from many CLI commands for the default svm1 SVM. You can optionally specify the SVM with the -vserver svm1 option or use another manually created SVM for secure multi-tenant deployments. For this deployment validation, we are utilizing the default svm1 SVM for simplicity.

### Create LUNs for hosts with ONTAP System Manager

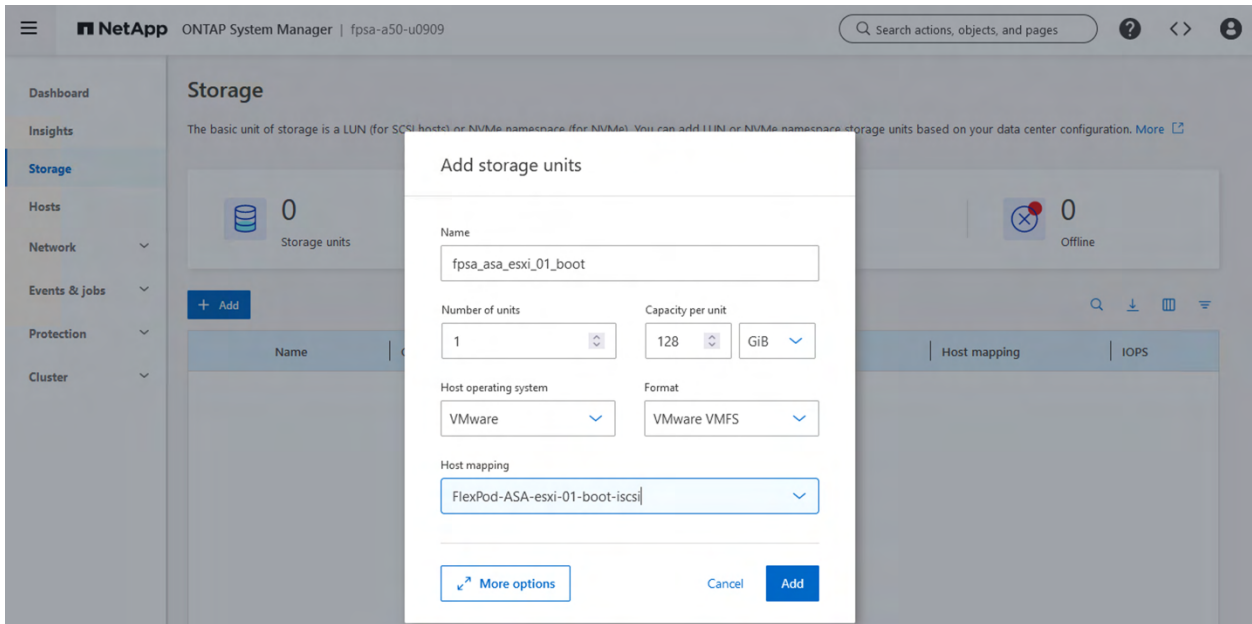
A traditional SAN solution uses LUNs for clients to access storage. With ONTAP System Manager, you can easily provision LUNs as LUN is the default storage unit. In addition, you can optionally provision NVMe namespaces and utilize NVMe protocol to access the ASA storage system. To create LUNs for hosts with ONTAP System Manager, follow the steps below.

1. Login to ONTAP System Manager.

Select Storage from menu.

On the Storage page, click +Add to add storage unit(s).

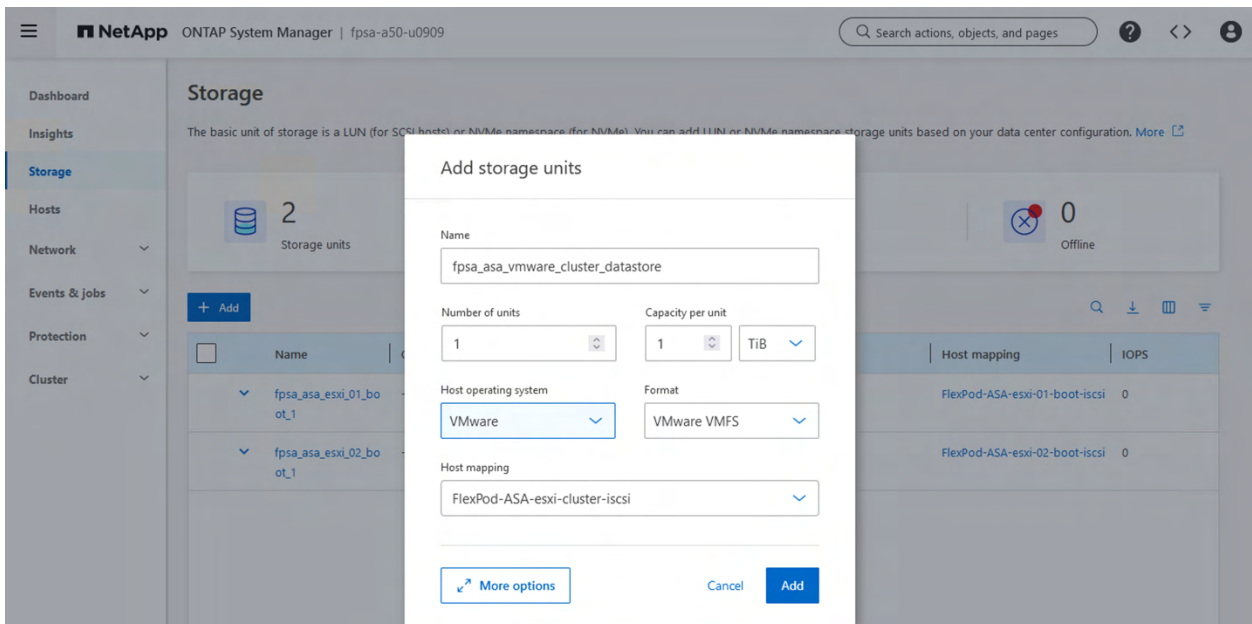
In the Add Storage Unit dialog, provide the name, number of units, capacity, host operating system, format, and select the correct SAN boot igroup for the host mapping.



Click Add to add the storage unit.

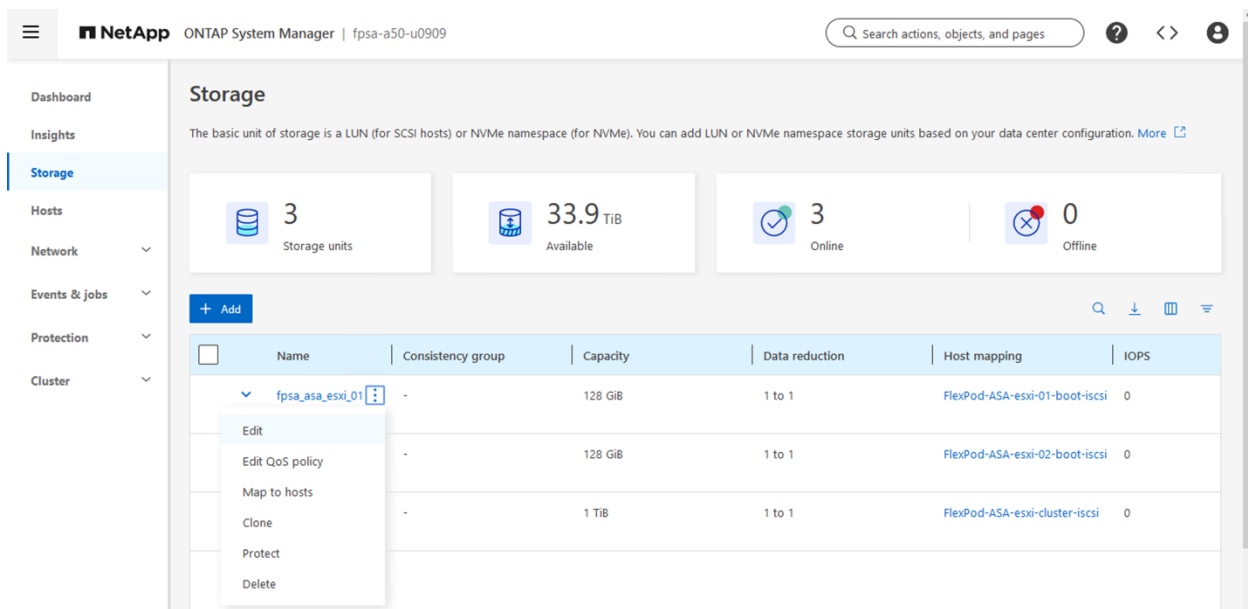
Repeat steps 3-5 to create the boot LUNs for all the ESXi hosts one by one. For each ESXi host, select the appropriate SAN boot igroup for the host mapping field.

Repeat steps 3-5 to create a shared LUN for the ESXi cluster datastore. For the host mapping, select the igroup which includes all the ESXi hosts. Create additional shared datastores as needed by either specifying the number of units when the datastore sizes are the same or repeat the step to specify additional LUNs with different LUN capacities.



**Note:** LUN is the default type for a new storage unit in the basic add storage unit dialog. Additional options for the storage unit creation are available when you click on More options in the dialog, including specifying quality of service (QoS) setting, selecting NVMe protocol, and configuring data protection settings.

**Note:** Clicking on the three dots next to the name of the storage unit as shown below opens a menu with additional configurations or actions that can be performed on the storage unit.



**Note:** The created storage units are listed on the Storage page after creation. You can also perform the LUN creation and host initiator group mapping by using ONTAP CLI commands. The following example shows the created storage-unit, LUN, and LUN mapping using CLI.

```
fpsa-a50-u0909::> storage-unit show
Vserver Name                                     Type      Size
-----
svml
  fpsa_asa_esxi_01_boot_1                       lun       128GB
  fpsa_asa_esxi_02_boot_1                       lun       128GB
  fpsa_asa_vmware_cluster_datastore_1 lun       1TB
3 entries were displayed.

fpsa-a50-u0909::> lun show
Vserver  Path                                     State  Mapped  Type      Size
-----
svml     fpsa_asa_esxi_01_boot_1                 online mapped  vmware    128GB
svml     fpsa_asa_esxi_02_boot_1                 online mapped  vmware    128GB
svml     fpsa_asa_vmware_cluster_datastore_1     online mapped  vmware    1TB
3 entries were displayed.

fpsa-a50-u0909::> lun show -m
Vserver  Path                                     Igroup   LUN ID  Protocol
-----
svml     fpsa_asa_esxi_01_boot_1                 FlexPod-ASA-esxi-01-boot-iscsi  0  iscsi
svml     fpsa_asa_esxi_02_boot_1                 FlexPod-ASA-esxi-02-boot-iscsi  0  iscsi
svml     fpsa_asa_vmware_cluster_datastore_1     FlexPod-ASA-esxi-cluster-iscsi  1  iscsi
3 entries were displayed.
```

## Cisco Intersight configuration

### Set up Cisco Intersight Managed Mode on Cisco UCS Fabric Interconnects

To support Intersight management of the X-Series Direct, the Cisco UCS Fabric Interconnects (FIs) need to be initialized in Intersight managed mode. To perform initialization of the FIs, connect to the console ports of the FIs and provide configuration information in the Basic System Configuration Dialog using the examples below as references.

You should start with the FI A that is located at the top rear of the chassis and configure it with FI A configurations.

```

----- Basic System Configuration Dialog -----

This setup utility will guide you through the basic configuration of
the system. Only minimal configuration including IP connectivity to
the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.
To back track or make modifications to already entered values,
complete input till end of section and answer no when prompted
to apply configuration.

Enter the configuration method. (console/gui) ? console

Enter the management mode. (ucsm/intersight)? intersight

The Fabric interconnect will be configured in the intersight managed mode. Choose (y/n) to
proceed: y

Enforce strong password? (y/n) [y]:

Enter the password for "admin":
Confirm the password for "admin":

Enter the switch fabric (A/B) []: A

Enter the system name: <ucs-cluster-name>

Physical Switch Mgmt0 IP address : <ucsa-mgmt-ip>

Physical Switch Mgmt0 IPv4 netmask : <ucsa-mgmt-mask>

IPv4 address of the default gateway : <ucs-mgmt-gateway>

DNS IP address : <dns-server-1-ip>

Configure the default domain name? (yes/no) [n]: yes

Default domain name : <dns-domain-name>

Following configurations will be applied:

Management Mode=intersight
Switch Fabric=A
System Name=<ucs-cluster-name>
Enforced Strong Password=yes
Physical Switch Mgmt0 IP Address=<ucsa-mgmt-ip>
Physical Switch Mgmt0 IP Netmask=<ucsa-mgmt-mask>
Default Gateway=<ucs-mgmt-gateway>
DNS Server=<dns-server-1-ip>
Domain Name=<dns-domain-name>

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.

```

**Note:** Replace the configuration information above with information appropriate for your deployment environment. Be sure to enter intersight for Intersight management mode during the configuration.

**Note:** Configuring fabric interconnects to use Cisco Intersight managed mode is a disruptive process, and existing configuration information will be lost. Customers are encouraged to make a backup of their existing configuration if the fabric interconnects were utilized for another purpose previously.

When FI A is set up and available, FI B setup process will automatically discover FI A as shown in the example console output below. FI B will prompt you to enter the admin password for FI A. Provide the management IP address for FI B and apply the configuration.

```

----- Basic System Configuration Dialog -----

```



This setup utility will guide you through the basic configuration of the system. Only minimal configuration including IP connectivity to the Fabric interconnect and its clustering mode is performed through these steps.

Type Ctrl-C at any time to abort configuration and reboot system.  
To back track or make modifications to already entered values, complete input till end of section and answer no when prompted to apply configuration.

Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y

Enter the admin password of the peer Fabric interconnect:

Connecting to peer Fabric interconnect... done

Retrieving config from peer Fabric interconnect... done

Peer Fabric interconnect management mode : intersight

Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>

Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucs-mgmt-mask>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : <ucsb-mgmt-ip>

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

Applying configuration. Please wait.

## Set up Cisco Intersight Account

1. Go to Cisco Intersight SaaS portal, <https://intersight.com>, and click Create an account.

Select a Region that is appropriate for your account and click Next.

Read and accept the license agreement. Click Next.

Provide an Account Name and click Create.

You can register smart licensing to start using your purchased license tiers for Cisco Intersight services or register to start a trial to evaluate Intersight services.

**Note:** You can also choose to add the Cisco UCS X-Series Direct FIs to an existing Cisco Intersight account.

## Set up Cisco Intersight Licensing

When setting up a new Cisco Intersight account, the account needs to be enabled for Cisco Smart Software Licensing. Please use the following steps as reference to set up Intersight licensing.

1. Log into the Cisco Smart Licensing portal: <https://software.cisco.com/software/smart-licensing/alerts>.

Verify that the correct virtual account is selected.

Under Inventory > General, generate a new token for product registration.

Copy this newly created token.

In Cisco Intersight click Select Service > System, then click Administration > Licensing.

Under Actions, click Register.

Enter the copied token from the Cisco Smart Licensing portal. Click Next.

Drop-down the pre-selected Default Tier \* and select the license type.

Select Move All Servers to Default Tier.

Click Register, then click Register again.

When the registration is successful, which takes a few minutes, the information about the associated Cisco Smart account and default licensing tier selected in the last step is displayed.

## Set Up Cisco Intersight Resource Group

For this procedure, a single Cisco Intersight resource group is created where resources will be logically grouped. Depending on your requirements, you can choose to create multiple resource groups for granular control of the resources.

1. Login to Intersight.

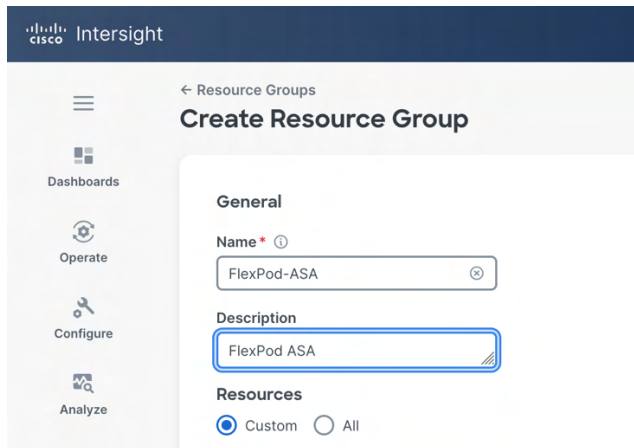
On the left, click on the Navigation Menu icon to expand or collapse the menu. Collapse the menu and select System near the bottom of the list to open the System menu.

Select Resource Groups under the System menu.

Click Create Resource Group in the top-right corner.

Provide a Name and a Description for the Resource Group.

Select Custom for custom resource assignment and click Create.

The screenshot shows the Cisco Intersight web interface. On the left is a navigation menu with icons for Dashboards, Operate, Configure, and Analyze. The main content area is titled 'Resource Groups' and 'Create Resource Group'. It contains a 'General' section with a 'Name' field (containing 'FlexPod-ASA') and a 'Description' field (containing 'FlexPod ASA'). Below this is a 'Resources' section with two radio buttons: 'Custom' (selected) and 'All'.

## Set Up Cisco Intersight Organization

In this step, an Intersight organization is created where all Cisco Intersight managed X-Series Direct configurations, including policies, are defined.

1. Login to Intersight.

Navigate to System > Organizations.

Click Create Organization in the top-right corner.

Provide a name and a description for the organization and click Next.

Select the Resource Group created in the previous procedure and click Next.

Review the information and click Create.

← Organizations  
**Create Organization**

General  
Configuration  
**Summary**

### Summary

Verify the details of the Organization and create

General

Name  
FlexPod-ASA
Description  
FlexPod ASA

Resource Groups

Search
Filters 1 results

Name	Used Organizations	Description
FlexPod-ASA	-	FlexPod ASA

Rows per page 10 < 1 >

Cancel
Back Create

## Claim Cisco UCS Fabric Interconnects in Cisco Intersight

Make sure the initial configurations for both UCS X-Series Direct FIs have been completed before claiming them in Intersight. To claim UCS FIs in Cisco Intersight, follow the steps below.

1. Use a web browser to access the IP address of FI A.

Login to the device using the previously configured username and password.

Under the Device Connector tab, the current device status will show “Not claimed”. Copy the Device ID and Claim Code information on the right to claim the device in Cisco Intersight.

Device Console | fpsa-x9508-u0901-fi

System Information **Device Connector** Inventory Diagnostic Data

The Device Connector management controller enables secure infrastructure management through Cisco Intersight. [Learn about Configuring Device Connector.](#)

Device Connector
Internet
Intersight

ACCESS MODE ALLOW CONTROL

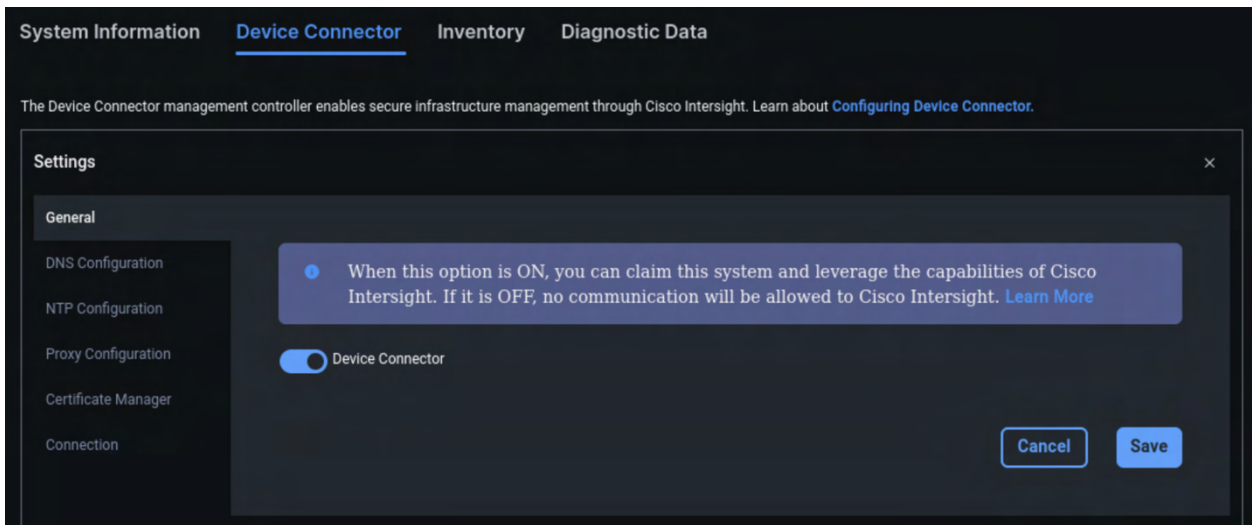
Not Claimed

The connection to the Cisco Intersight Portal is successful, but device is still not claimed. To claim the device open Cisco Intersight, create a new account and follow the guidance or go to the Targets page and click Claim a New Device for existing account.
Open Intersight

Device ID
Claim Code

**Note:** The Claim Code information is only valid for a limited time and new code will be generated when the old one expires.

**Note:** Under the Device Connector Settings menu, you can configure additional network settings appropriate for your environment, including DNS, NTP, and network proxy server information.

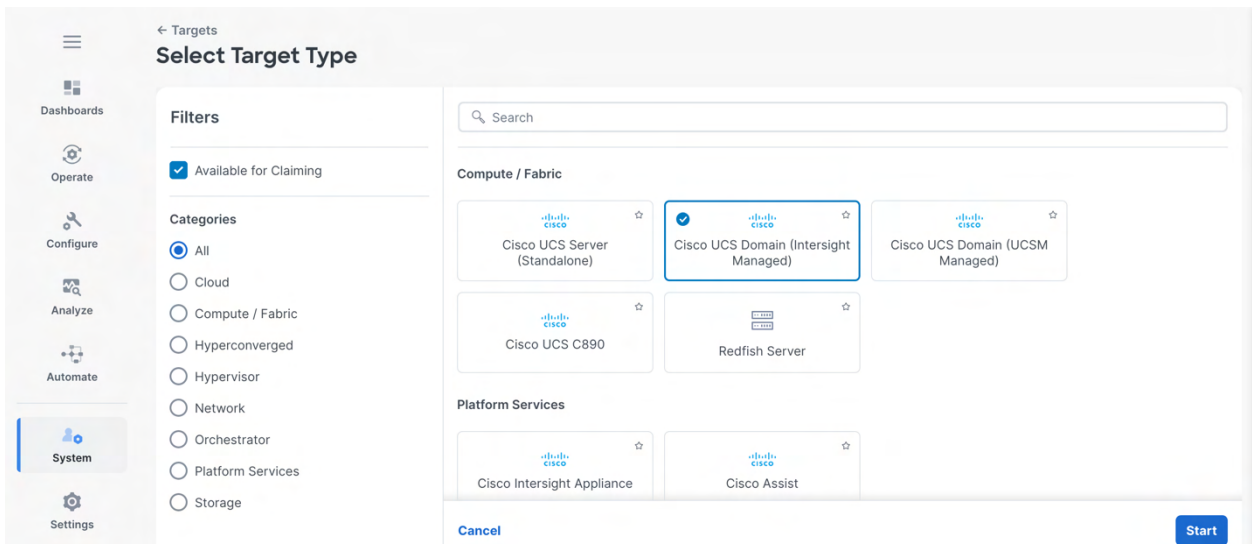


Login to Intersight.

Navigate to System > Targets.

Click Claim a New Target in the top-right corner.

Select Cisco UCS Domain (Intersight Managed) and click Start.



Copy and paste the Device ID and Claim Code obtained from the Cisco UCS FI to Intersight.

Select the previously created Resource Group and click Claim.

← Targets  
**Claim a New Target**

Device ID \*

Claim Code \*

Device ID

Claim Code

Required Required

**Resource Groups**

Select resource groups, if required. This is not mandatory, since by default, the claimed target will be added to "All" type resource groups.

[Export](#)

<input checked="" type="checkbox"/>	Name	Usage	Description
<input checked="" type="checkbox"/>	FlexPod-ASA	FlexPod-ASA	FlexPod ASA

Selected 1 of 11   [Show Selected](#)   [Unselect All](#)   Rows per page 10 < 1 2 >

With a successful device claim, Cisco UCS FIs should appear as targets in Cisco Intersight.

← Targets  
**fpsa-x9508-u0901-fi** Healthy

**Details**

Health Healthy

Status Connected

Name  
fpsa-x9508-u0901-fi

Type  
Intersight Managed Domain

Vendor  
Cisco Systems, Inc.

FI A  
fpsa-x9508-u0901-fi FI-A

FI B  
fpsa-x9508-u0901-fi FI-B

**Sub Targets**

Search Filters 4 results

<input type="checkbox"/>	Name	Status
<input type="checkbox"/>	C1-B1-UCSX-215C-M8	Connected
<input type="checkbox"/>	C1-B2-UCSX-215C-M8	Connected
<input type="checkbox"/>	C1-B3-UCSX-215C-M8	Connected
<input type="checkbox"/>	C1-B4-UCSX-215C-M8	Connected

Rows per page 10 < 1 >

**Events**

Alarms No Alarms

Active Acknowledged

No Alarms

In addition, the Device Connector tab in the FI accessed via the web browser should also indicate that the device is claimed with the associated Intersight Account which claimed the device.

Device Console | fpsa-x9508-u0901-fi

System Information **Device Connector** Inventory Diagnostic Data

The Device Connector management controller enables secure infrastructure management through Cisco Intersight. [Learn about Configuring Device Connector.](#)

**Device Connector** [Settings](#) [Refresh](#)

ACCESS MODE: ALLOW CONTROL

Device ID

Claimed to Account

[Unclaim](#)

Claimed

## Create a Cisco UCS Domain Profile

A Cisco UCS domain profile configures a fabric interconnect pair through reusable policies to configure ports, port channels, VLANs, and VSANs in the network. It defines the characteristics of the configured ports on fabric interconnects. One Cisco UCS domain profile can be assigned to one fabric interconnect domain. The domain-related policies can be attached to the profile either at the time of creation or later.

1. Login to Intersight.

Navigate to Configure > Profiles.

In the main window, select UCS Domain Profiles and click Create UCS Domain Profile.

The screenshot shows the 'Create UCS Domain Profile' screen in the Intersight interface. The left sidebar contains navigation options: Dashboards, Operate, Configure (selected), Analyze, Automate, System, and Settings. The main content area is titled 'Create UCS Domain Profile' and shows a 'UCS Domain Assignment' step. A message at the top indicates 'caused by inconsistent configuration.' Below this is a large gear and wrench icon. The text 'UCS Domain Assignment' is followed by 'Create a Fabric Interconnect pair and assign to a domain profile immediately or later.' There are three dots indicating the current step. A link 'About UCS Domain Profile Creation' is visible, along with a checkbox 'Do not show this page again'. At the bottom, there are 'Cancel' and 'Start' buttons.

On the Create UCS Domain Profile screen, click Start.

Select the organization from the drop-down list.

Provide a name for the domain profile and optionally set Tags and provide a Description.

The screenshot shows the 'Create UCS Domain Profile' screen in the Intersight interface, specifically the 'General' configuration step. The left sidebar contains navigation options: Dashboards, Operate, Configure (selected), Analyze, Automate, System, and Settings. The main content area is titled 'Create UCS Domain Profile' and shows a list of steps: 1 General (selected), 2 UCS Domain Assignment, 3 VLAN & VSAN Configuration, 4 Ports Configuration, 5 UCS Domain Configuration, and 6 Summary. The 'General' step is expanded, showing fields for 'Organization' (FlexPod-ASA), 'Name' (FlexPod-ASA-Domain-Profile), 'Set Tags' (Enter a tag in the key:value format.), and 'Description' (Description). The 'Next' button is highlighted.

Click Next.



## Assign UCS domain to domain profile

1. Assign a Cisco UCS domain to this new domain profile by clicking Assign Now and selecting the previously added Cisco UCS domain.

The screenshot shows the 'Create UCS Domain Profile' interface. On the left is a sidebar with navigation options: Dashboards, Operate, Configure (selected), Analyze, Automate, System, and Settings. The main panel has a breadcrumb '← UCS Domain Profiles' and a title 'Create UCS Domain Profile'. Below the title is a list of steps: 1. General, 2. UCS Domain Assignment (selected), 3. VLAN & VSAN Configuration, 4. Ports Configuration, 5. UCS Domain Configuration, and 6. Summary. The 'UCS Domain Assignment' section has a heading 'UCS Domain Assignment' and a sub-heading 'Choose to assign a fabric interconnect pair to the profile now or later.' Below this are two buttons: 'Assign Now' (selected) and 'Assign Later'. A blue information box states: 'Choose to assign a Fabric Interconnect pair now or later. If you choose Assign Now, select a pair that you want to assign and click Next. If you choose Assign Later, click Next to proceed to policy selection.' There is a 'Show Assigned' toggle switch. Below is a search bar and a table with 1 result. The table has columns for 'Domain Name', 'Fabric Interconnect A' (with sub-columns Model, Serial, Bundle Ve...), and 'Fabric Interconnect B' (with sub-columns Model, Serial, Bundle Ve...). The row shows 'fpsa-x9508-u0901-fi' assigned to 'Fabric Interconnect A' (Model: UCSX-S9..., Serial: FCH2824..., Bundle Ve...: 4.3(5.240...)). Below the table are 'Selected 1 of 1', 'Show Selected', 'Unselect All', and 'Rows per page' (10). At the bottom are 'Close', 'Back', and 'Next' buttons.

Click Next.

## VLAN configuration

A single VLAN policy is created for both fabric interconnects for IP-based storage.

1. Click Select Policy next to VLAN Configuration under Fabric Interconnect A.

The screenshot shows the 'Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)' interface. The sidebar is the same as in the previous screenshot. The main panel has a breadcrumb '← UCS Domain Profiles' and a title 'Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)'. Below the title is a list of steps: 1. General, 2. UCS Domain Assignment, 3. VLAN & VSAN Configuration (selected), 4. Ports Configuration, 5. UCS Domain Configuration, and 6. Summary. The 'VLAN & VSAN Configuration' section has a heading 'VLAN & VSAN Configuration' and a sub-heading 'Create or select a policy for the Fabric Interconnect pair.' Below this are two sections for 'Fabric Interconnect A' and 'Fabric Interconnect B', each showing '0 of 2 Policies Configured'. Each section has two rows: 'VLAN Configuration' and 'VSAN Configuration', each with a 'Select Policy' button. At the bottom are 'Close', 'Back', and 'Next' buttons.

Click Create New in the pane on the right.

Verify the correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags and provide a Description.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Create VLAN

**1 General**

Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-VLAN

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
VLAN policy for both Fis  
24 / 1024

[Cancel](#) [Next](#)

Click Next.

Click Add VLANs.

Provide a name and VLAN ID for the native VLAN.

Make sure Auto Allow on Uplinks is enabled.

To create the required Multicast policy, click Select Policy.

In the window on the right, Click Create New to create a new Multicast Policy.

Provide a Name for the Multicast Policy and optionally provide Tags and Description and click Next.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Create VLAN

## Create Multicast Policy

**1 General**

Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-Multicast-Policy

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description  
0 / 1024

[Cancel](#) [Next](#)

Leave the Snooping State and Source IP Proxy State selected and click Create.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Create VLAN

## Create Multicast Policy

1 General

2 Policy Details

### Policy Details

Add policy details.

#### Multicast Policy

☒ Snooping State ⓘ  
☐ Querier State ⓘ  
☒ Source IP Proxy State ⓘ

< [Cancel](#) [Back](#) [Create](#)

Click Add to add the VLAN.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Create VLAN

1 General

2 Policy Details

### Add VLANs

Add VLANs to the policy

VLANs should have one Multicast policy associated to it

#### Configuration

Prefix \* ⓘ  ⓘ
 VLAN IDs \* ⓘ  ⓘ

☒ Auto Allow On Uplinks ⓘ  
☐ Enable VLAN Sharing ⓘ

#### Multicast Policy \*

Selected Policy ⓘ [Edit Selection](#) ⓘ

[Cancel](#) [Add](#)

Select Set Native VLAN ID and enter the VLAN number under VLAN ID.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Create VLAN

General

2 Policy Details

**Add VLANs**

☐ Show VLAN ID Ranges

**Filters** 2 results

<input type="checkbox"/>	VLAN ID	Name	Sharing Ty...	Primary VL...	Multicast Policy	Auto A	
<input type="checkbox"/>	1	default	None			Yes	...
<input type="checkbox"/>	2	Native-VLAN_2	None		FlexPod-ASA-M...	Yes	...

Rows per page 10 < 1 >

☒ Set Native VLAN ID

VLAN ID

[Cancel](#) [Back](#) [Create](#)

Add the remaining VLANs for FlexPod by clicking Add VLANs and entering the VLANs one by one.

Reuse the previously created multicast policy for all the VLANs.

The VLANs created during this validation are shown below.

**VLANs**

**Add VLANs**

☐ Show VLAN ID Ranges

**Filters** 9 results

<input type="checkbox"/>	VLAN ID	Name	Sharing Type	Pr	Multicast Policy	Auto Allow On Uplin	
<input type="checkbox"/>	1	default	None			Yes	...
<input type="checkbox"/>	2	Native-VLAN	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2272	IB-MGMT	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2273	iSCSI-A	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2274	iSCSI-B	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2275	vMotion	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2276	VM-Network	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2277	NVMe-TCP-A	None		FlexPod-ASA-Multicast-Policy	Yes	...
<input type="checkbox"/>	2278	NVMe-TCP-B	None		FlexPod-ASA-Multicast-Policy	Yes	...

**Note:** The NVMe-TCP VLANs shown are only needed when NVMe/TCP protocol will be used in the environment.

Click Create at bottom right to finish creating the VLAN policy and associated VLANs.

Click Select Policy next to VLAN Configuration for Fabric Interconnect B, select the same VLAN policy, and click Next.

## Port configuration for FI A

Using two separate port policies for the two fabric interconnects provide flexibility when port configuration (port numbers, speed, port channel ID) differs between the two FIs.

1. Click Select Policy for Fabric Interconnect A.

Click Create New to define a new port configuration policy.

Verify correct organization is selected from the drop-down list and provide a name.

Select the UCSX-S9108-100G Switch Model.

Optionally set Tags and provide a Description.

The screenshot shows the 'Create Port' configuration page within the 'UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)' section. The page has a left sidebar with navigation options: Dashboards, Operate, Configure (selected), Analyze, Automate, System, and Settings. The main content area is titled 'Create Port' and contains a 'General' tab. The 'General' tab includes fields for 'Organization' (FlexPod-ASA), 'Name' (FlexPod-ASA-Port-Policy-A), 'Switch Model' (UCSX-S9108-100G), 'Set Tags' (a text input field with a placeholder 'Enter a tag in the key:value format.'), and 'Description' (a text input field with a placeholder 'Description' and a character count '0 / 1024'). At the bottom right, there is a 'Next' button.

Click Next.

Click Next and Next again to skip the Unified Port and Breakout Options if you are not using FC and don't have need to use breakout options.

Click the Port Channels tab and click Create Port Channel.

The screenshot shows the 'Edit' configuration page for 'FlexPod-ASA-Port-Policy-A' within the 'Policies > Port' section. The page has a left sidebar with navigation options: Dashboards, Operate, Configure (selected), Analyze, Automate, System, and Settings. The main content area is titled 'Edit' and contains a 'Port Roles' tab. The 'Port Roles' tab includes a 'Create Port Channel' button and a list of port roles. Below the list, there is a photograph of a Cisco UCS X-Series server rack.

Select Ethernet Uplink Port Channel as the role, provide a port-channel ID, and select a value for Admin Speed from drop-down list.

Policies > Port > FlexPod-ASA-Port-Policy-A

## Edit

### Create Port Channel

Configuration

**i** The combined maximum number of Ethernet Uplink, FCoE Uplink, and Appliance port channels permitted is 8 and the maximum number of FC port channels permitted is 4.

Role  
 Ethernet Uplink Port Channel

Port Channel ID \* ① 21 1 - 256

Admin Speed ① 100Gbps

FEC ① CI91

Ethernet Network Group ①

Select Policies

Flow Control

Select Policy

Link Aggregation

Select Policy

**Note:** You can create Ethernet Network Group, Flow Control, Link Aggregation for defining disjoint Layer-2 domain or fine tune port-channel parameters. These policies were not used in this deployment and system default values were utilized.

Under Link Control, click Select Policy and click Create New in the upper right hand corner.

Verify the correct organization is selected from the drop-down list, provide a name for the policy, and click Next.

Policies > Port > FlexPod-ASA-Port-Policy-A > Edit

## Create Link Control

1 General

2 Policy Details

**General**

Add a name, description, and tag for the policy.

Organization \* FlexPod-ASA

Name \* FlexPod-ASA-UDLD-Link-Control

Set Tags

Enter a tag in the key:value format.

Description

Description 0 / 1024

Leave the default values selected and click Create.

Policies > Port > FlexPod-ASA-Port-Policy-A > Edit

## Create Link Control

1 General

2 Policy Details

**Policy Details**

Add policy details.

Configuration

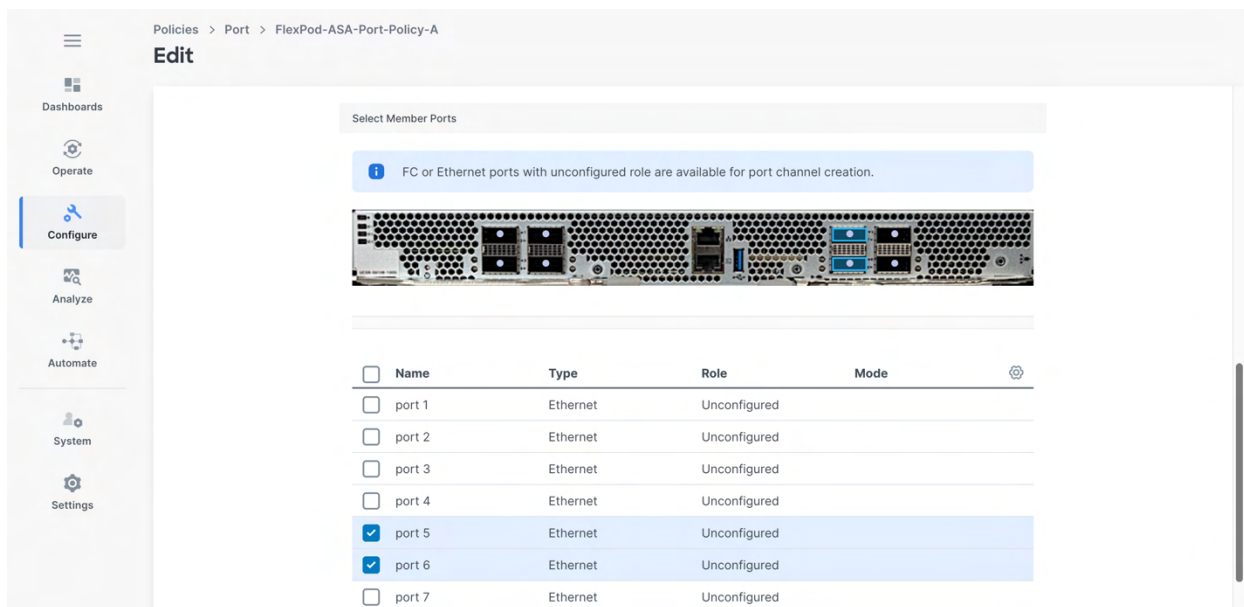
UDLD Admin State ①

UDLD Mode ①

Normal Aggressive

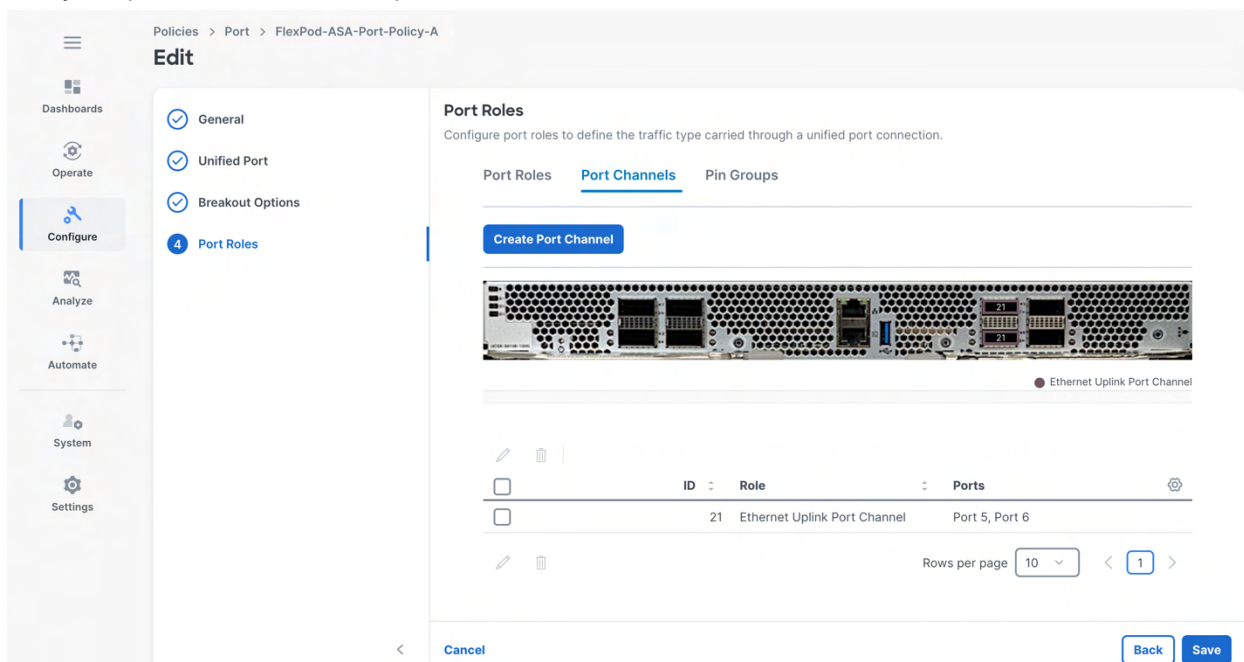
Under Select Member Ports, select the ports that are connected to the Nexus uplink switches. Adjust the port selections to match your deployment environment.





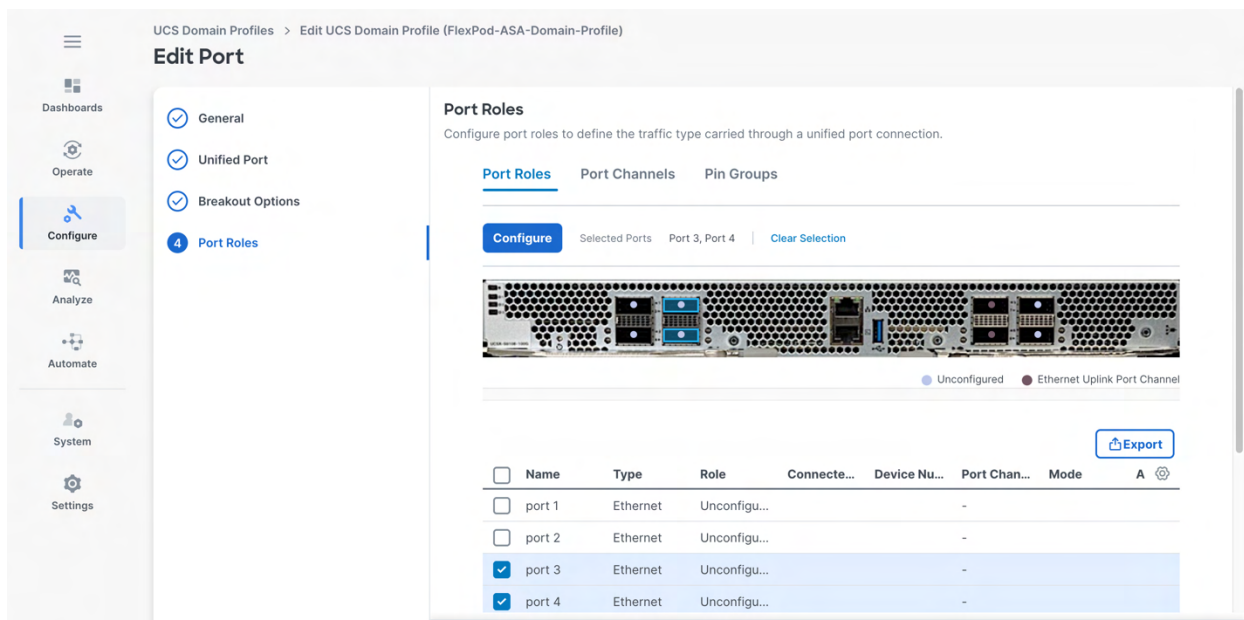
Click Save.

Verify the port-channel ID, Role, ports and click Save.

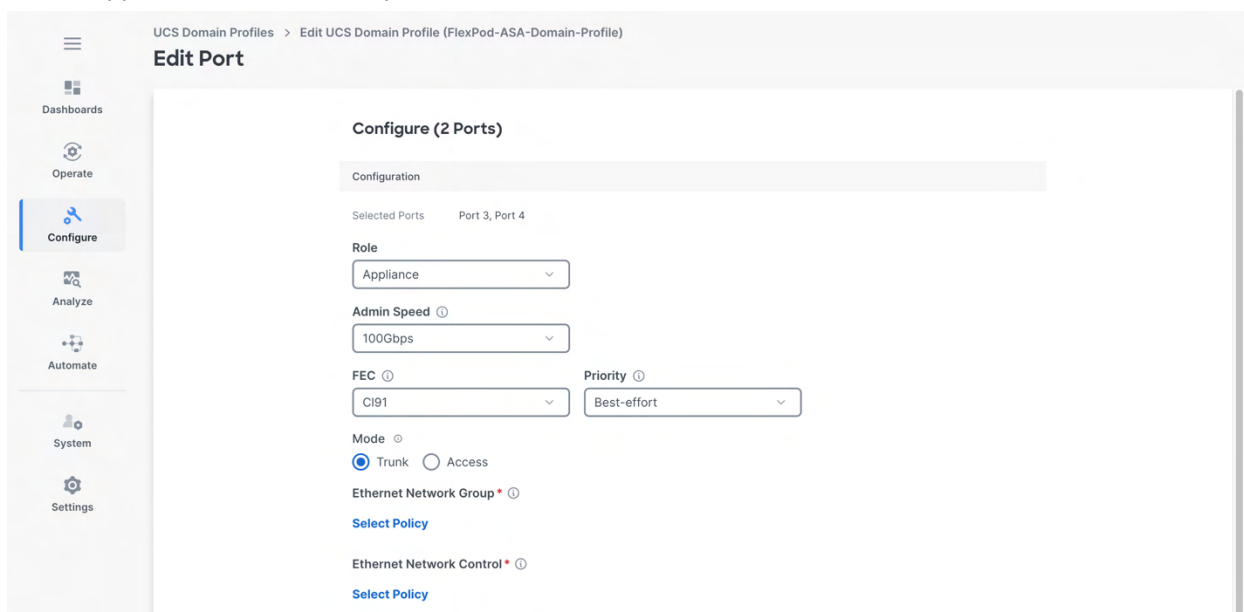


Select Port Roles tab.

Under Port Roles tab, select the ports where NetApp ASA storage controllers are connected to and click Configure.



Select Appliance Role, Admin Speed, and leave the others as default.



Under Ethernet Network Group, click Select Policy.

On the Select Policy page, click Create New.

Make sure the Organization is selected.

Provide a Name and optionally set Tags add a Description. Click Next.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Edit Port

## Create Ethernet Network Group

1 General

2 Policy Details

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-Network-Group-Storage-A

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Under Policy Details, click Add VLANs and select Enter Manually.

Enter the iSCSI-A and NVMe-TCP-A VLANs in the list that will be going to the storage from FI A and click Enter.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Edit Port

## Create Ethernet Network Group

1 General

2 Policy Details

**Policy Details**  
Manage policy settings and allowed VLANs.

☐ Enable QinQ (802.1Q-in-802.1Q) Tunneling on the vNIC

[Add VLANs](#)

☐ Show VLAN ID Ranges

**Enter Manually**

VLANs \* ①  
2273,2277

Must be between 1 - 4093

[Cancel](#) [Enter](#)

Click Create.

Under Ethernet Network Control, click Select Policy.

On the Select Policy page, click Create New.

Make sure the Organization is selected.

Provide a Name and optionally set Tags add a Description. Click Next.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Edit Port

## Create Ethernet Network Control

1 General

2 Policy Details

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-Network-Control-Storage

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Under Policy Details, Enable CDP, LLDP Transmit, and LLDP Receive.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile) > Edit Port

## Create Ethernet Network Control

Dashboards
Operate
Configure
Analyze
Automate
System
Settings

General

2 Policy Details

### Policy Details

Add policy details.

**i** This policy is applicable only for UCS Servers (FI-Attached)

☒ Enable CDP ⓘ

MAC Register Mode ⓘ  
☒ Only Native VLAN ☐ All Host VLANs

Action on Uplink Fail ⓘ  
☒ Link Down ☐ Warning

**!** Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.

### MAC Security

Forge ⓘ  
☒ Allow ☐ Deny

### LLDP

☒ Enable Transmit ⓘ

☒ Enable Receive ⓘ

[Cancel](#) [Back](#) [Create](#)

Click Save and click Save again.

Configure and change the port role for the select ports to Appliance.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Edit Port

Dashboards
Operate
Configure
Analyze
Automate
System
Settings

General

Unified Port

Breakout Options


4 Port Roles

### Port Roles

Configure port roles to define the traffic type carried through a unified port connection.

**Port Roles** | Port Channels | Pin Groups

[Configure](#) Selected Ports Port 3, Port 4 [Clear Selection](#)



● Unconfigured ● Ethernet Uplink Port Channel

[Export](#)

<input type="checkbox"/>	Name	Type	Role	Connecte...	Device Nu...	Port Chan...	Mode	A	⚙
<input type="checkbox"/>	port 1	Ethernet	Unconfigu...			-			
<input type="checkbox"/>	port 2	Ethernet	Unconfigu...			-			
<input checked="" type="checkbox"/>	port 3	Ethernet	Unconfigu...			-			
<input checked="" type="checkbox"/>	port 4	Ethernet	Unconfigu...			-			

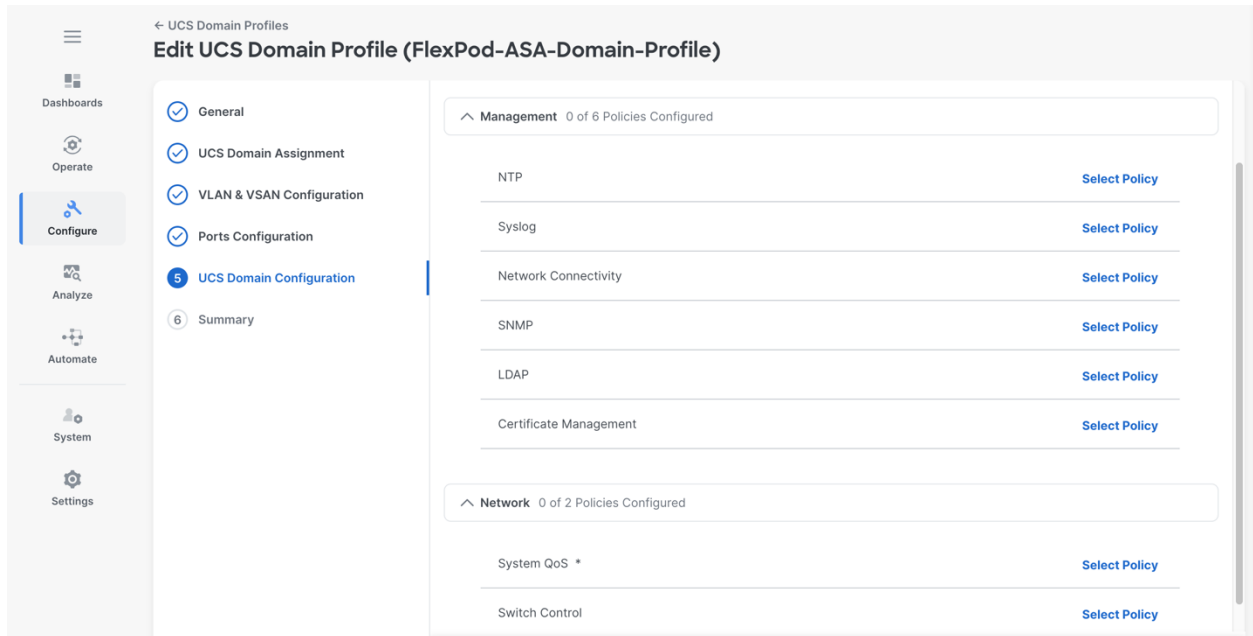
## Port configuration for FI B

1. From the Navigate menu, navigate to Edit the Ports Configuration for the previously saved UCS Domain Profile to create port configuration for FI B.
2. Repeat the steps in Port configuration for FI A to create the port policy for Fabric Interconnect B including the Ethernet port-channel.
3. Apply the same Link Control policy created during port configuration for FI A for the port channel creation.
4. Utilize the following information for the configuration.

- Name of the port policy: FlexPod-ASA-Port-Policy-B
- Ethernet Port-Channel ID: 22
- Network Storage Group: FlexPod-ASA-Network-Group-Storage-B with iSCSI-B and NVMe-TCP-B VLANs

## UCS Domain Configuration

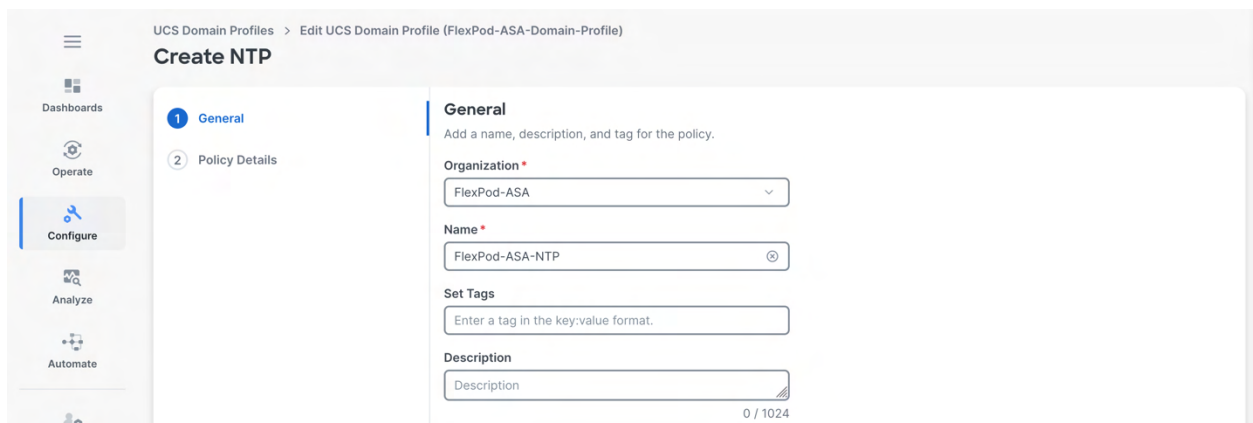
Under UCS domain configuration, additional policies can be configured to setup NTP, Syslog, DNS settings, SNMP, QoS and UCS operating mode (end host or switch mode). For this deployment, four policies (NTP, Network Connectivity, and System QoS) will be configured, as shown below:



## Configure NTP Policy

1. Click Select Policy next to NTP and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.



Click Next.

Enable NTP, provide the first NTP server IP address, and select the time zone from the drop-down list. Add a second NTP server by clicking + next to the first NTP server IP address.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

### Create NTP

**General** (selected) | **Policy Details**

Add policy details.

☒ Enable NTP ⓘ

**NTP Servers \*** ⓘ

172.21.62.121 ⓘ

**NTP Servers \*** ⓘ

172.21.62.122 ⓘ

**Timezone** ⓘ

America/New\_York

**Note:** When the NTP server IP addresses are the uplink Nexus switch management IPs, the NTP distribution should be configured in the Cisco Nexus switches.

Click Create.

## Configure Network Connectivity Policy

1. Click Select Policy next to Network Connectivity and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

### Create Network Connectivity

**General** (selected) | **Policy Details**

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-Network-Connectivity ⓘ

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

Click Next.

Provide DNS server IP addresses for Cisco UCS. Click Create.

## System QoS Policy

1. Click Select Policy next to System QoS\* and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy.



UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Create System QoS

**1 General**

Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-QoS

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Click Next.

Change the MTU for Best Effort class to 9216.

Keep the default selections or change the parameters if necessary.

Click Create.

UCS Domain Profiles > Edit UCS Domain Profile (FlexPod-ASA-Domain-Profile)

## Create System QoS

**2 Policy Details**

Add policy details.

**Configure Priorities**

☐ Platinum

☐ Gold

☐ Silver

☐ Bronze

☒ Best Effort

☐ Fibre Channel

**CoS** **Weight** **Allow Packet Drops** **MTU**

Any 5 0 - 10 9216 1500 - 9216

3 5 0 - 6 0 - 10 2240 1500 - 9216

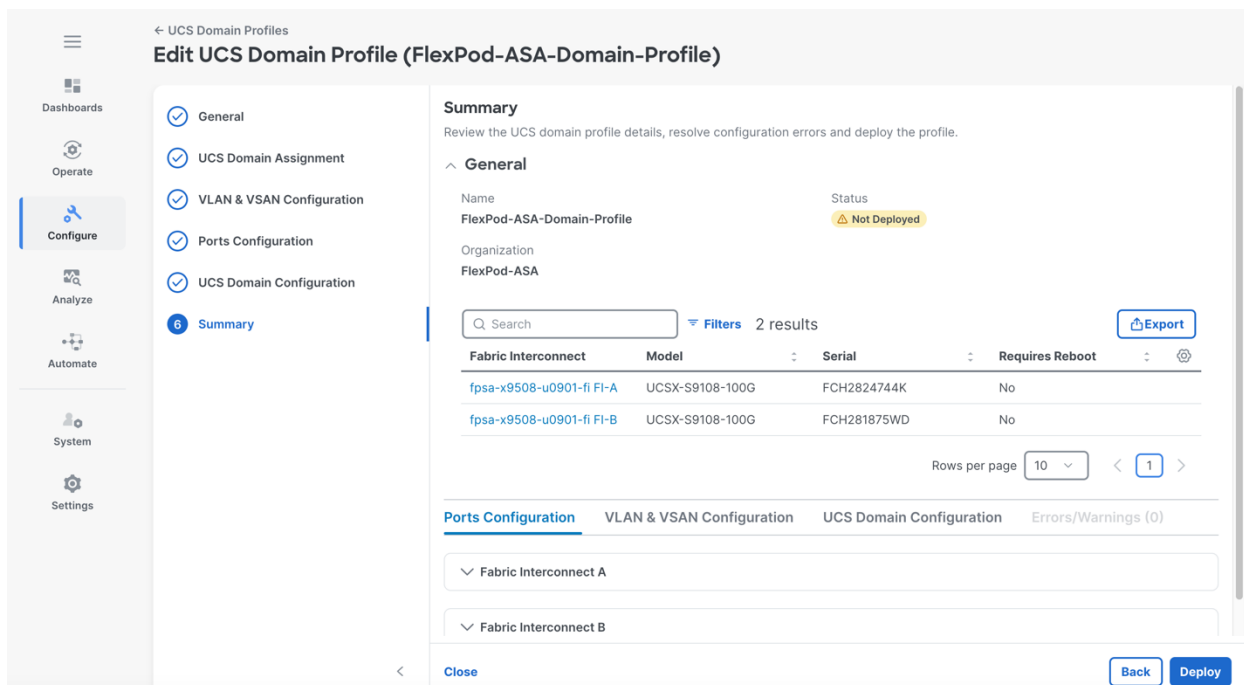
Cancel Back Create

Click Next.

## UCS Domain Profile Summary

1. Verify all settings by expanding on the settings to check and make sure that the configurations are correct.





## Deploy the Cisco UCS Domain Profile

1. From the UCS domain profile Summary view, Click Deploy.

Acknowledge any warnings and click Deploy again.

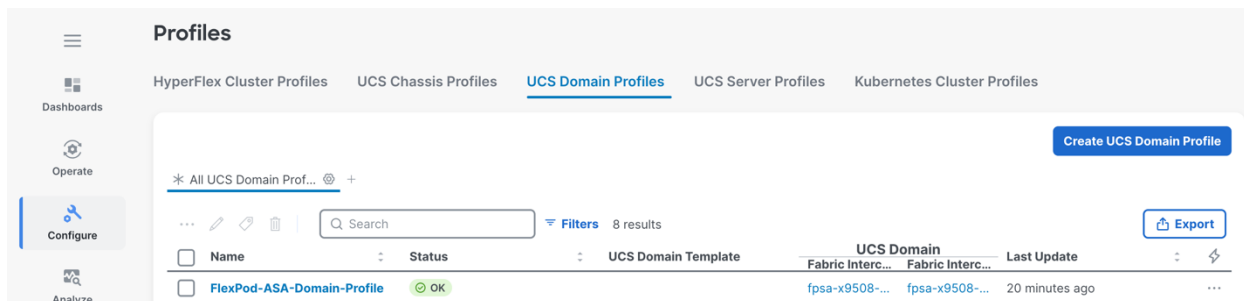
**Note:** The system will take some time to validate and configure the settings on the fabric interconnects. Log into the console ports of the FIs to monitor when the Cisco UCS fabric interconnects have finished configuration and are successfully rebooted.

## Verify Cisco UCS Domain Profile Deployment

When the Cisco UCS domain profile has been successfully deployed, the Cisco UCS chassis and the compute nodes should be successfully discovered. It takes a while to discover the compute nodes for the first time. Watch the number of outstanding requests in Intersight.

1. Navigate to Configure > Profiles > UCS Domain Profiles.

Verify that the domain profile has been successfully deployed.



Verify that the chassis has been discovered and is visible under Operate > Chassis.

Requests								
<input type="checkbox"/>	Blade Discovery	Success	system@inter...	Blade Server	a few second...	fpsa-x9508-u0901-fi-1-3	18 s	67ac3689696... ⋮
<input type="checkbox"/>	Blade Discovery	Success	system@inter...	Blade Server	a few second...	fpsa-x9508-u0901-fi-1-4	19 s	67ac3689696... ⋮
<input type="checkbox"/>	Chassis Inventory	Success	Jyh-shing.Ch...	Chassis	a few second...	fpsa-x9508-u0901-fi-1	6 s	67ac3686696... ⋮
<input type="checkbox"/>	Chassis Inventory	Success	Jyh-shing.Ch...	Chassis	a few second...	fpsa-x9508-u0901-fi-1	6 s	67ac3686696... ⋮
<input type="checkbox"/>	Rediscover Primary Chassis	Success	Jyh-shing.Ch...	Fabric Interco...	a few second...	fpsa-x9508-u0901-fi-1, switch B	1 s	67ac3684696... ⋮
<input type="checkbox"/>	Rediscover Primary Chassis	Success	Jyh-shing.Ch...	Fabric Interco...	a few second...	fpsa-x9508-u0901-fi-1, switch A	1 s	67ac3684696... ⋮
<input type="checkbox"/>	Deploy Domain Profile	Success	Jyh-shing.Ch...	Fabric Interco...	24 minutes ago	fpsa-x9508-u0901-fi FI-B	25 s	67ac30f4696... ⋮
<input type="checkbox"/>	Deploy Domain Profile	Success	Jyh-shing.Ch...	Fabric Interco...	24 minutes ago	fpsa-x9508-u0901-fi FI-A	1 m 28 s	67ac30f4696... ⋮

## Configure Cisco UCS chassis profile

Cisco UCS chassis profile configures a UCS X9508 chassis through reusable policies. It defines the characteristics of power distribution and fan configuration in the chassis. One Cisco UCS chassis profile can be assigned to one chassis.

1. Navigate to Configure > Profiles.

In the main window, select UCS Chassis Profiles tab and click Create UCS Chassis Profile.

From the Create UCS Chassis Profile screen, click Start.

## UCS chassis profile general configuration

1. Select the organization from the drop-down list.

Provide a name for the chassis profile.

Optionally set Tags and provide a Description.

Click Next.

## UCS chassis profile chassis assignment

1. Assign the Cisco UCS chassis to this new chassis profile by clicking Assign Now and selecting a Cisco UCS chassis.

Click Next.

## UCS chassis profile chassis configuration

1. Click Select Policy next to Power.

Click Create New to create a new power policy.

Make sure the correct Organization is selected.

Enter a Name for the policy.

Optionally set Tags and provide a Description.

Click Next.

Select All Platforms.

UCS Chassis Profiles > Create UCS Chassis Profile

## Create Power Policy

General

2 Policy Details

Policy Details

Add policy details.

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

Configuration

☒ Power Profiling ⓘ

Power Priority ⓘ Low

Power Restore ⓘ Always Off

Power Redundancy ⓘ Grid

Processor Package Power Lim... ⓘ Default

☒ Power Save Mode ⓘ

☒ Dynamic Power Rebalancing ⓘ

☒ Extended Power Capacity ⓘ

Power Allocation (Watts) ⓘ 0

**Note:** It is recommended to leave all settings at their defaults, but the settings can be adjusted later according to performance and sustainability requirements.

Click Create to create the power policy.

## Create and Apply Thermal Policy

1. Click Select Policy next to Thermal.

Click Create New to create a new policy.

Make sure the correct Organization is selected.

Enter a Name for the policy.

Optionally set Tags and provide a Description.

UCS Chassis Profiles > Create UCS Chassis Profile

## Create Thermal Policy

1 General

2 Policy Details

General

Add a name, description, and tag for the policy.

Organization \* FlexPod-ASA

Name \* FlexPod-ASA-Thermal ⓘ

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

Click Next.

**Note:** It is recommended to leave all settings at their defaults, but the settings can be adjusted later according to performance and sustainability requirements.

Click **Create** to create the thermal policy.

Click **Next**.

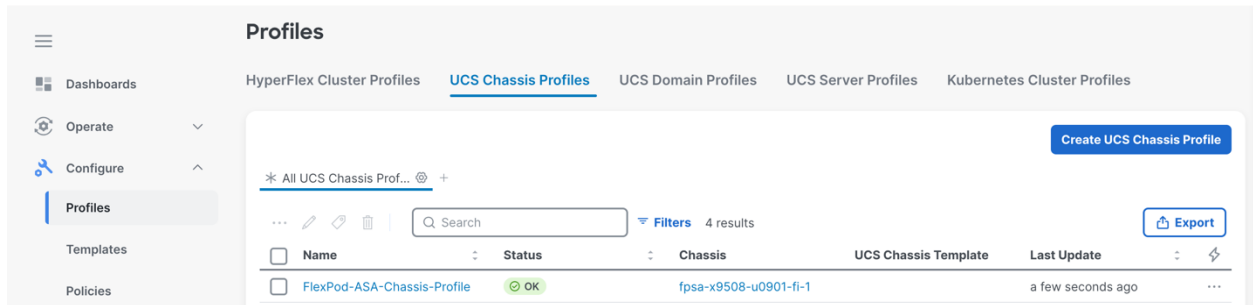
## Review and deploy UCS chassis profile

1. Review the UCS Chassis Profile Summary.

Click **Deploy**.

Click **Deploy** again to deploy the profile.

When deployment is completed, the profile Status should show OK.



**Note:** The above chassis profile creation procedures can be used to create profiles for additional chassis. In these additional chassis profiles, the power and thermal policies can be reused as needed.

## Configure server profile template

In the Cisco Intersight platform, a server profile enables resource management by simplifying policy alignment and server configuration. The server profiles are derived from a server profile template. A server profile template and its associated policies can be created using the server profile template wizard. After creating server the profile template, customers can derive multiple consistent server profiles from the template. The server profile templates captured in this validation supports Cisco UCS X215c M8 compute nodes with 5th Generation VICs and 4th Generation and 5th Generation AMD EPYC processors.

## vNIC placement for server profile template

In this deployment, a server profile template is created for VMware hosts with iSCSI connected storage, including SAN boot. Six vNICs are configured with two vNICs to support iSCSI boot from SAN using two separate SAN fabrics. These vNICs are manually placed as listed in Table 8.

**Table 8 Server virtual NIC placement and configurations**

PCI Order	vNIC Name	Switch ID	Failover
0	00-vSwitch0-A	A	Disabled
1	01-vSwitch0-B	B	Disabled
2	02-vDS0-A	A	Disabled
3	03-vDS0-B	B	Disabled
4	04-iSCSI-A	A	Disabled
5	05-iSCSI-B	B	Disabled

**Note:** NVMe-TCP traffic will share with iSCSI traffic on the iSCSI vNICs when NVMe/TCP is being used. If separating NVMe/TCP traffic from iSCSI traffic is desirable, you can add additional vNICs for the NVMe/TCP traffic.

## Server profile template creation

1. Login to Intersight.

Navigate to Configure > Templates.

In the main window, click Create UCS Server Profile Template.

## Server profile template general configuration

1. Select the organization from the drop-down list

Provide a name for the server profile template.

Select UCS Server (FI-Attached).  
Optionally set Tags and provide a Description.

The screenshot shows the 'Create UCS Server Profile Template' page. On the left is a navigation sidebar with 'Templates' selected. The main content area has a breadcrumb 'UCS Server Profile Templates' and a title 'Create UCS Server Profile Template'. Below the title is a list of steps: 1. General, 2. Compute Configuration, 3. Management Configuration, 4. Storage Configuration, 5. Network Configuration, and 6. Summary. The 'General' step is active. The form fields include: 'Organization' (FlexPod-ASA), 'Name' (FlexPod-ASA-AMD-ISCSI-Boot), 'Target Platform' (radio buttons for 'UCS Server (Standalone)' and 'UCS Server (FI-Attached)', with the latter selected), 'Set Tags' (a text input for key-value tags), and 'Description' (a text input). A character count '0 / 1024' is at the bottom right.

Click Next.

## Server profile template compute configuration – UUID Pool

1. Click Select Pool under UUID Pool and then click Create New in the pane on the right.  
Verify correct organization is selected from the drop-down list and provide a name for the UUID Pool.  
Optionally set Tags or provide a Description.

The screenshot shows the 'Create UUID' page. The breadcrumb is 'UCS Server Profile Templates > Create UCS Server Profile Template'. The title is 'Create UUID'. The left sidebar is the same as the previous screenshot. The main content area has a list of steps: 1. General, 2. Pool Details. The 'General' step is active. The form fields include: 'Organization' (FlexPod-ASA), 'Name' (FlexPod-ASA-UUID-Pool), 'Set Tags' (a text input for key-value tags), and 'Description' (a text input). A character count '0 / 1024' is at the bottom right.

Click Next.

Provide a UUID Prefix.

Add a UUID block with starting UUID and size.



UCS Server Profile Templates > Create UCS Server Profile Template

### Create UUID

General

Pool Details

Pool Details  
Collection of UUID suffix Blocks.

Configuration

Prefix \*

AAA00000-0000-0001

UUID Blocks

From

AAA0-00000000000001

Size

64

1 - 1024

Click Create.

## Server profile template compute configuration – BIOS Policy

1. Click Select Policy next to BIOS and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description.

Templates > Create UCS Server Profile Template

### Create BIOS Policy

General

Policy Details

General  
Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-AMD-M8-Virt-BIOS

Set Tags

Enter a tag in the key:value format.

Description

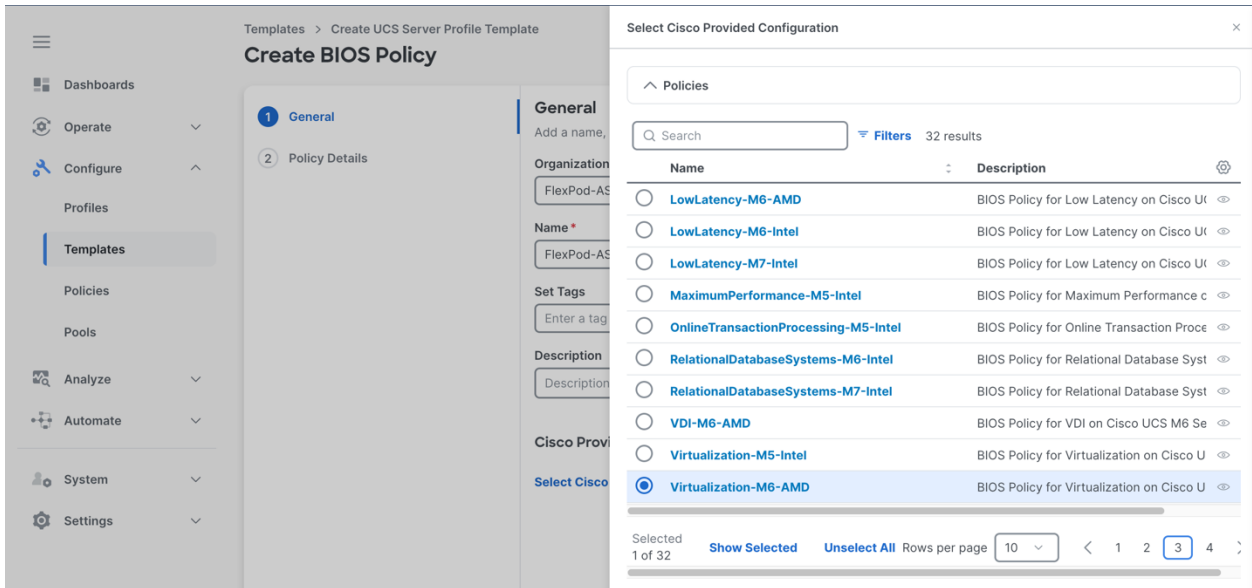
Description

0 / 1024

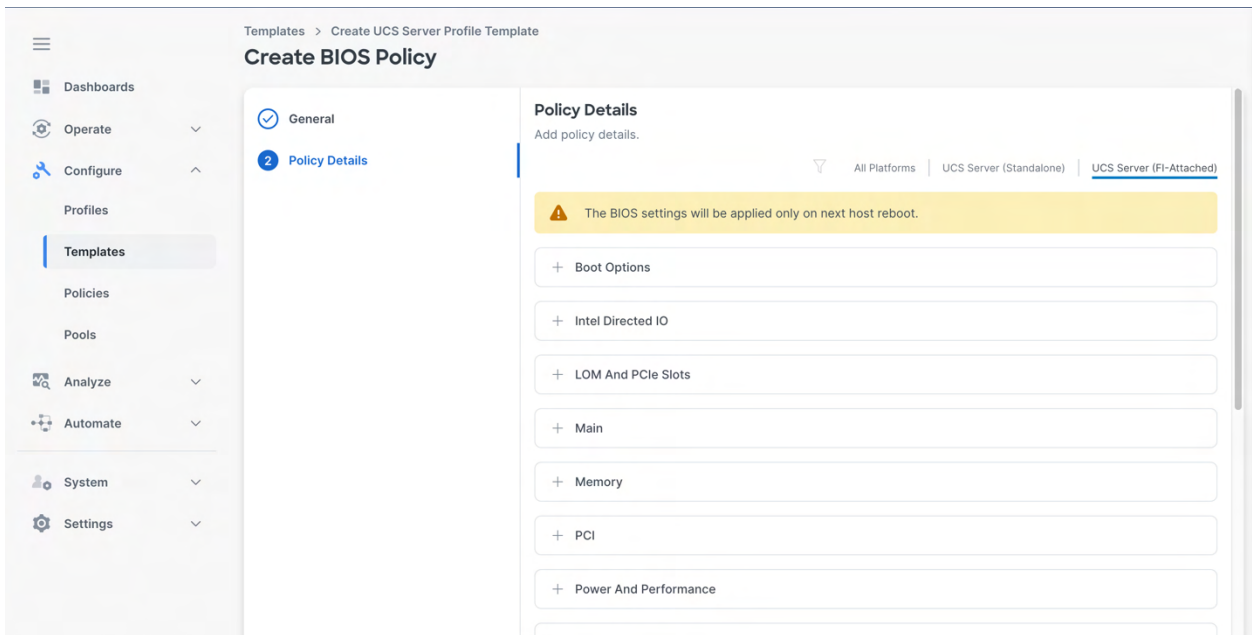
Cisco Provided BIOS Configuration

Select Cisco Provided Configuration

**Note:** Here you have an option to click and Select Cisco Provided Configuration to review and select one that is suitable for your environment.



Select Virtualization-M6-AMD and click Next.



On the Policy Details screen, select appropriate values for the BIOS settings. Click on the + under Power And Performance and set CPPC to enabled.

**Note:** Except for CPPC under Power and Performance, the BIOS values were left as platform default based on virtualization workload recommendations in [Performance Tuning for Cisco UCS M8 Platforms with 4<sup>th</sup> Gen and 5<sup>th</sup> Gen AMD EPYC Processors](#).

Click Create.

## Server profile template compute configuration – Boot Order

1. Click Select Policy next to Boot Order and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description.

Click Next.

For Configured Boot Mode, select Unified Extensible Firmware Interface (UEFI).

Turn on Enable Secure Boot.

Click Add Boot Device drop-down list and select Virtual Media.  
Provide a device name and select KVM Mapped DVD for sub-type.

From the Add Boot Device drop-down list, select iSCSI Boot.

Provide the Device Name: iSCSI-A-Boot and the exact name of the interface used for iSCSI boot under Interface Name: 04-iSCSI-A.

**Note:** The device names (iSCSI-A-Boot and iSCSI-B-Boot) are being defined here and will be used in the later steps of the iSCSI configuration.

From the Add Boot Device drop-down list, select iSCSI Boot.

Provide the Device Name: iSCSI-B-Boot and the exact name of the interface used for iSCSI boot under Interface Name: 05-iSCSI-B.

From the Add Boot Device drop-down list, select Virtual Media.

Add Device Name CIMC-Mapped-DVD and select the subtype CIMC MAPPED DVD.

Verify the order of the boot policies and adjust the boot order to the following using arrows next to the Delete button.

Click Create.

## Server profile template compute configuration – Virtual Media

1. Click Select Policy next to Virtual Media and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Optionally set Tags or provide a Description.

The screenshot shows the 'Create Virtual Media Policy' form in the General tab. The left sidebar contains a navigation menu with 'Templates' selected. The main content area has two tabs: 'General' (active) and 'Policy Details'. The 'General' tab contains the following fields:

- Organization:** A dropdown menu with 'FlexPod-ASA' selected.
- Name:** A text input field with 'FlexPod-ASA-KVM-Mount-Media' entered.
- Set Tags:** A text input field with the placeholder 'Enter a tag in the key:value format.'
- Description:** A text input field with 'Description' entered.

The bottom right corner of the form shows '0 / 1024' characters.

Click Next.

Turn on Enable Virtual Media, Enable Virtual Media Encryption, and Enable Low Power USB.

Do not Add Virtual Media at this time, but the policy can be modified and used to map an ISO for a CIMC Mapped DVD.

The screenshot shows the 'Create Virtual Media Policy' form in the Policy Details tab. The left sidebar is the same as the previous screenshot. The main content area has two tabs: 'General' and 'Policy Details' (active). The 'Policy Details' tab contains the following elements:

- Configuration:** Three toggle switches, all turned on: 'Enable Virtual Media', 'Enable Virtual Media Encryption', and 'Enable Low Power USB'.
- Add Virtual Media:** A blue link.
- Table:** A table with columns 'Name', 'Type', 'Protocol', and 'File Location'. The table is empty, and the text 'NO ITEMS AVAILABLE' is displayed below it.

Click Create.

Click Next to move to Management Configuration.

## Server profile template management configuration

Four policies will be added to the management configuration:

- IMC Access to define a pool of IP addresses for compute node KVM access
- IPMI Over LAN to allow Intersight to manage IPMI messages
- Local User to provide local administrator to access KVM
- Virtual KVM to allow the Tunneled KVM

## Server profile template management configuration – IMC Access

1. Click Select Policy next to IMC Access and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Optionally, set Tags or provide a Description.

The screenshot shows the 'Create IMC Access Policy' form in the UCS Manager interface. The left sidebar contains navigation options: Dashboards, Operate, Configure, Profiles, Templates (selected), Policies, Pools, Analyze, and Automate. The main content area has a breadcrumb trail: Templates > Create UCS Server Profile Template > Create IMC Access Policy. Below the breadcrumb is a progress indicator with two steps: 1 General (selected) and 2 Policy Details. The 'General' section includes a description: 'Add a name, description, and tag for the policy.' It contains four input fields: 'Organization' (a dropdown menu with 'FlexPod-ASA' selected), 'Name' (a text field with 'FlexPod-ASA-IMC-Access' and a copy icon), 'Set Tags' (a text field with placeholder 'Enter a tag in the key:value format.'), and 'Description' (a text area with 'Description' and a character count '0 / 1024').

Click Next.

You can select in-band management access to the compute node using an in-band management VLAN or out-of-band management access via the Mgmt0 interfaces of the FIs. KVM Policies like SNMP, vMedia and Syslog are currently not supported via Out-Of-Band and will require an In-Band IP to be configured. Below we are using out-of-band access in the configuration.

Click UCS Server (FI-Attached).

Disable In-Band Configuration and enable Out-Of-Band Configuration

Under IP Pool, click Select IP Pool and then click Create New in the pane on the right.

Verify the correct organization is selected from the drop-down list and provide a name for the policy.

Optionally set Tags or provide a Description.

The screenshot shows the 'Create IP Pool' form in the UCS Manager interface. The left sidebar contains navigation options: Fabric Interconnects, Networking, HyperFlex Clusters, Storage, Virtualization, Integrated Systems, Configure (selected), Profiles, Templates, Policies, and Pools. The main content area has a breadcrumb trail: Templates > Create UCS Server Profile Template > Create IMC Access Policy > Create IP Pool. Below the breadcrumb is a progress indicator with three steps: 1 General (selected), 2 IPv4 Pool Details, and 3 IPv6 Pool Details. The 'General' section includes a description: 'Pool represents a collection of IPv4 and/or IPv6 addresses that can be allocated to other configuration entities like server profiles.' It contains four input fields: 'Organization' (a dropdown menu with 'FlexPod-ASA' selected), 'Name' (a text field with 'FlexPod-ASA-OOB-MGMT-IP-Pool' and a copy icon), 'Set Tags' (a text field with placeholder 'Enter a tag in the key:value format.'), and 'Description' (a text area with 'Description' and a character count '0 / 1024'). At the bottom, there is a checkbox labeled 'Configure Subnet at Block Level' with a question mark icon.

Click Next.

Select Configure IPv4 Pool and provide the information to define a pool for KVM IP address assignment including an IP Block.

**Note:** The management IP pool subnet should be accessible from the host that is trying to open the KVM connection. In the example shown here, the hosts trying to open a KVM connection would need to be able to route to the 172.22.71.0/24 subnet.

Click Next.

Deselect Configure IPv6 Pool or provide proper IPv6 Pool information if required.

Click Create to finish configuring the IP address pool.

Click Create to finish configuring the IMC access policy.

## Server profile template management configuration – IPMI Over LAN

1. Click Select Policy next to IPMI Over LAN and then click Create New in the pane on the right.

Verify the correct organization is selected from the drop-down list and provide a name for the policy.

Optionally set Tags or provide a Description.

Click Next

On the right, select UCS Server (FI-Attached)

Turn on Enable IPMI Over LAN.

From the Privilege Level drop-down list, select admin.

Provide an encryption key with even number of hexadecimal characters less than 40 in length for IPMI communication.



The screenshot shows the 'Create IPMI Over LAN Policy' configuration page. The left sidebar contains a navigation menu with 'Templates' selected. The main content area has two tabs: 'General' (selected) and 'Policy Details'. The 'Policy Details' tab is active, showing a toggle for 'Enable IPMI Over LAN' (checked), a 'Privilege Level' dropdown set to 'admin', and an 'Encryption Key' field with a 'Show' button. At the top right, there are platform filters: 'All Platforms', 'UCS Server (Standalone)', and 'UCS Server (FI-Attached)' (selected).

Click Create.

## Server profile template management configuration – Local User

1. Click Select Policy next to Local User and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description.

The screenshot shows the 'Create Local User Policy' configuration page. The left sidebar contains a navigation menu with 'Templates' selected. The main content area has two tabs: 'General' (selected) and 'Policy Details'. The 'General' tab is active, showing fields for 'Organization' (FlexPod-ASA), 'Name' (FlexPod-ASA-Local-User), 'Set Tags' (a text input with a placeholder 'Enter a tag in the key:value format.'), and 'Description' (a text input with a placeholder 'Description'). At the bottom right of the description field, it shows '0 / 1024' characters.

Click Next.

Verify that UCS Server (FI-Attached) is selected.

Verify that Enforce Strong Password is selected.

This screenshot is identical to the one above, showing the 'Create Local User Policy' configuration page with the 'General' tab active. It displays the same fields for Organization, Name, Tags, and Description.

Click Add New User.

Provide a username, select a role for the user, and provide a password.

The screenshot shows the 'Create Local User Policy' configuration page. The left sidebar contains navigation options: Dashboards, Operate, Configure, Profiles, Templates (selected), Policies, Pools, Analyze, Automate, System, and Settings. The main content area has two tabs: 'General' and 'Policy Details' (selected). Under 'Policy Details', there are sections for 'Password Properties' and 'Local Users'.

**Password Properties:**

- Enforce Strong Password:** Enabled (toggle switch).
- Enable Password Expiry:** Disabled (toggle switch).
- Password History:** Set to 5 (range 0 - 5).
- Always Send User Password:** Disabled (toggle switch).

**Local Users:**

An information box states: "This policy will remove existing user accounts other than the ones configured with this policy. However, the default admin user account is not deleted from the endpoint device. You can only enable/disable or change account password for the admin account by creating a user with the user name and role as 'admin'. If there are no users in the policy, only the admin user account will be available on the endpoint device. By default, IPMI support is enabled for all users."

**Add New User:**

A user entry for 'flexadmin (admin)' is shown with the 'Enable' toggle switch turned on. Below this, the 'Add New User' form is visible with the following fields:

- Username:** flexadmin
- Role:** admin
- Password:** (masked with dots, with a 'Show' button)
- Password Confirmation:** (masked with dots, with a 'Show' button)

**Note:** The username and password combination defined here will be used as an alternate to log in to KVMs and can be used for IPMI.

Click Create to finish configuring the user.

## Server profile template management configuration – Virtual KVM

1. Click Select Policy next to Virtual KVM and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description.

The screenshot shows the 'Create Virtual KVM Policy' configuration page. The left sidebar is the same as the previous screenshot. The main content area has two tabs: 'General' (selected) and 'Policy Details'. The 'General' tab contains the following fields:

- Organization:** FlexPod-ASA (selected from a dropdown menu).
- Name:** FlexPod-ASA-Virtual-KVM (with a copy icon).
- Set Tags:** Enter a tag in the key-value format.
- Description:** Description (with a character count of 0 / 1024).

Click Next.

Verify that UCS Server (FI-Attached) is selected.

Turn on "Allow Tunneled vKVM."

Click Create.

**Note:** To fully enable Tunneled KVM after the Server Profile Template has been created, go to System > Settings > Security and Privacy and click Configure. Turn on “Allow Tunneled vKVM Launch” and “Allow Tunneled vKVM Configuration.”

Click Next and Next again to skip Storage Configuration and move to Network Configuration.

## Server profile template network configuration – LAN Connectivity

The LAN connectivity policy defines the connections and network communication resources between the server and the LAN. This policy uses pools to assign MAC addresses to servers and to identify the vNICs that the servers use to communicate with the network. For iSCSI hosts, this policy also defines an IQN address pool. The vNICs for the iSCSI SAN booted hosts are manually placed by using Advanced placement with 6 vNICs configured as listed in Table 9.

**Table 9 LAN connectivity for vNICs**

vNIC Name	Switch ID	PCI Order	Slot ID	VLANs
00-vSwitch0-A	A	0	MLOM	IB-MGMT
01-vSwitch0-B	B	1	MLOM	IB-MGMT
02-vDS0-A	A	2	MLOM	VM-Traffic, vMotion
03-vDS0-B	B	3	MLOM	VM-Traffic, vMotion
04-iSCSI-A	A	4	MLOM	iSCSI-A
05-iSCSI-B	B	5	MLOM	iSCSI-B

1. Click Select Policy next to LAN Connectivity and then click Create New in the pane on the right.

Verify the correct organization is selected from the drop-down list and provide a name for the policy.

Select UCS Server (FI-Attached).

Optionally, set Tags or provide a Description.

Templates > Create UCS Server Profile Template

### Create LAN Connectivity Policy

1 General

2 Policy Details

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-iSCSI-Boot-LAN-Connectivity

**Target Platform**   
☐ UCS Server (Standalone)
 ☒ UCS Server (FI-Attached)

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Click Next.

Under IQN, select Pool.

Click Select Pool under IQN Pool and then click Create New in the pane on the right.

Verify the correct organization is selected from the drop-down list and provide a name for the IQN Pool.

Optionally set Tags or provide a Description.

Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy

### Create Iqnpool:Pool

1 General

2 Pool Details

**General**  
Pool represents a collection of iSCSI Qualified Names (IQNs) for use as initiator identifiers by iSCSI vNICs.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-IQN-Pool

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Click Next.

Provide the values for Prefix and IQN Block to create the IQN pool.

Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy

### Create Iqnpool:Pool

✓ General

2 Pool Details

**Pool Details**  
Collection of IQN Blocks.

**Configuration**

**Prefix \***   
iqn.2010-11.com.flexpod

**IQN Blocks**

Suffix	From	Size
FlexPod-ASA-ucshost	1	32

>= 0 1 - 1024

Click Create.

Under vNIC Configuration, select Manual vNICs Placement.

Click Add and select vNIC.

## Server profile template network configuration – MAC Address Pool

When creating the first vNIC, the MAC address pool has not been defined yet therefore a new MAC address pool will need to be created. Two separate MAC address pools are configured, one for each Fabric. MAC-Pool-A will be reused for all Fabric-A vNICs, and MAC-Pool-B will be reused for all Fabric-B vNICs.

**Table 10 MAC address pools**

Pool Name	Starting MAC Address	Size	vNICs
MAC-Pool-A	00:25:B5:AA:A0:00	256*	01-vSwitch0-A, 03-VDS0-A, 05-ISCSI-A
MAC-Pool-B	00:25:B5:AA:B0:00	256*	02-vSwitch0-B, 04-VDS0-B, 06-ISCSI-B

**Note:** Each server requires 3 MAC addresses from each pool. Adjust the size of the pool according to your requirements.

1. Click Select Pool under MAC Pool and then click Create New in the pane on the right.

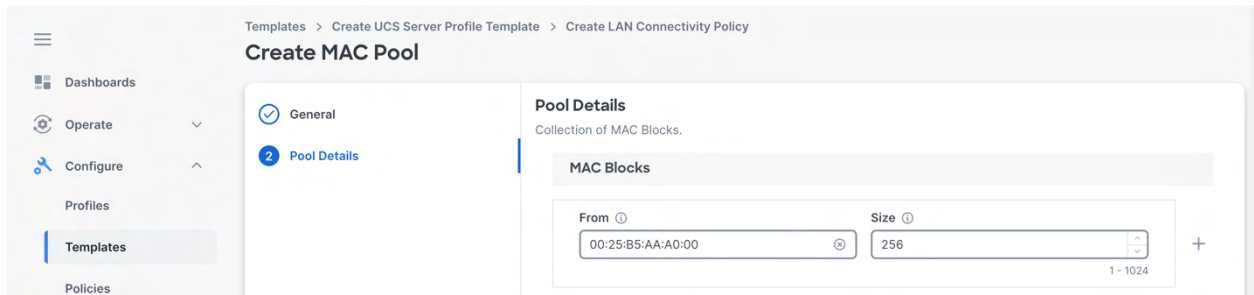
Verify the correct organization is selected from the drop-down list and provide a name for the pool from Table 10 depending on the vNIC being created.

Optionally set Tags or provide a Description.

Click Next.

Provide the starting MAC address from Table 10.

Provide the size of the MAC address pool from Table 10.

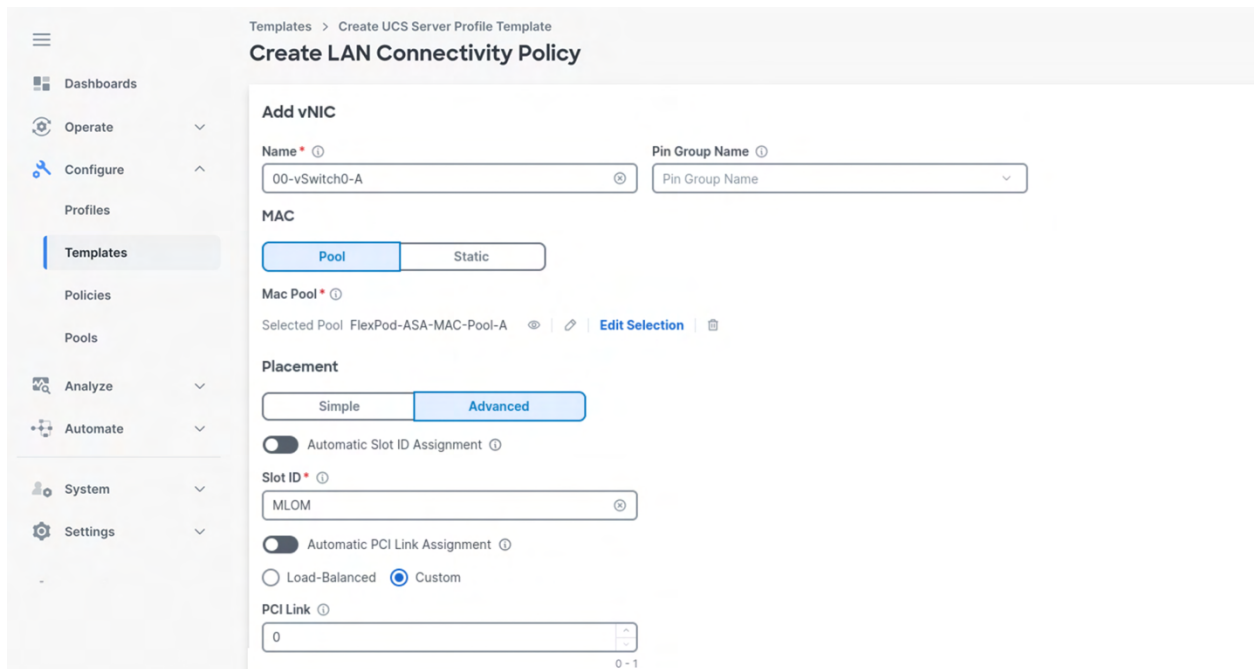


**Note:** For ease of troubleshooting FlexPod, some information can be coded into the MAC address pool to help identify the rack ID in the data center and the fabric interconnect identity. For example, in the starting address 00:25:B5:AA:A0:00 can be used to identify rack AA, fabric interconnect A.

Click Create to finish creating the MAC address pool.

## Server profile template network configuration – Add vNIC

1. From the Add vNIC screen, provide vNIC Name, Switch ID, and PCI Order information from Table 9 using Advanced placement.



For Consistent Device Naming (CDN), from the drop-down list, select vNIC Name.

Verify that Failover is disabled because the failover will be provided by attaching multiple NICs to the VMware vSwitch and vDS.

## Server profile template network configuration – Ethernet Network Group

Ethernet Network Group policies will be created and reused on applicable vNICs as covered below. The ethernet network group policy defines the VLANs allowed for a particular vNIC, therefore multiple network group policies will be defined for this deployment as listed in Table 11.



**Table 11 Ethernet network group policies**

Policy Name	vNICs	VLANs	Native VLAN
FlexPod-ASA-vSwitch0-Network-Group	01-vSwitch0-A 02-vSwitch0-B	Native VLAN IB-MGMT	Native VLAN (2)
FlexPod-ASA-vDS0-Network-Group	03-vDS0-A 04-vDS0-B	Native VLAN VM Traffic vMotion	Native VLAN (2)
FlexPod-ASA-iSCSI-A-Network-Group	05-iSCSI-A	iSCSI-A NVMe-TCP-A	iSCSI-A (2273)
FlexPod-ASA-iSCSI-B-Network-Group	06-iSCSI-B	iSCSI-B NVMe-TCP-B	iSCSI-B (2274)

**Note:** Add the NVMe-TCP VLANs to the iSCSI network group policies when using NVMe-TCP with the iSCSI vNICs.

1. Click Select Policy under Ethernet Network Group Policy and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy from Table 11.

Optionally, set Tags or provide a Description.

Click Next.

Enter the allowed VLANs and the native VLAN ID from Table 11.



Click Create.

**Note:** When ethernet group policies are shared between two vNICs, the ethernet group policy only needs to be defined for the first vNIC. For subsequent vNIC policy mapping, click Select Policy and pick the previously defined ethernet group policy from the list.

## Server profile template network configuration – Ethernet Network Control

The Ethernet Network Control Policy is used to enable Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) for the vNICs. A single policy will be created here and reused for all the vNICs.

1. Click Select Policy under Ethernet Network Control Policy and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Optionally set Tags or provide a Description.

The screenshot shows the 'Create Ethernet Network Control' form in the General tab. The left sidebar contains navigation links: Dashboards, Operate, Configure, Profiles, Templates (selected), Policies, Pools, Analyze, and Automate. The main content area has a breadcrumb trail: Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy. The form title is 'Create Ethernet Network Control'. The General tab is active, showing fields for Organization (FlexPod-ASA), Name (FlexPod-ASA-Network-Control), Set Tags (a text input with a placeholder 'Enter a tag in the key-value format.'), and Description (a text input with a placeholder 'Description' and a character count '0 / 1024').

Click Next.

Enable Cisco Discovery Protocol and both Enable Transmit and Enable Receive under LLDP.

The screenshot shows the 'Create Ethernet Network Control' form in the Policy Details tab. The left sidebar is the same as the previous screenshot. The main content area has a breadcrumb trail: Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy. The form title is 'Create Ethernet Network Control'. The Policy Details tab is active, showing a blue information banner: 'This policy is applicable only for UCS Servers (FI-Attached)'. Below this are several configuration options: 'Enable CDP' (checked), 'MAC Register Mode' (radio buttons for 'Only Native VLAN' and 'All Host VLANs', with 'Only Native VLAN' selected), 'Action on Uplink Fail' (radio buttons for 'Link Down' and 'Warning', with 'Link Down' selected), and a yellow warning box: 'Important! If the Action on Uplink is set to Warning, the switch will not fail over if uplink connectivity is lost.' Below the warning box are two sections: 'MAC Security' with 'Forge' (radio buttons for 'Allow' and 'Deny', with 'Allow' selected) and 'LLDP' with 'Enable Transmit' and 'Enable Receive' (both checked).

Click Create to finish creating Ethernet network control policy.

## Server profile template network configuration – Ethernet QoS

The Ethernet QoS policy is used to enable jumbo maximum transmission units (MTUs) for all the vNICs. A single policy will be created and reused for all the vNICs.

1. Click Select Policy under Ethernet QoS and click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description.

The screenshot shows the 'Create Ethernet QoS' form in the UCS Manager interface. The 'General' tab is active, showing fields for Organization (FlexPod-ASA), Name (FlexPod-ASA-Ethernet-QoS), Set Tags (empty), and Description (empty). The left sidebar shows the navigation menu with 'Templates' selected.

Click Next.

Change the MTU, Bytes value to 9000.

The screenshot shows the 'Create Ethernet QoS' form in the UCS Manager interface. The 'Policy Details' tab is active, showing QoS Settings. The MTU, Bytes value is set to 9000. The Rate Limit, Mbps value is set to 0. The Burst value is set to 10240. The Priority is set to Best-effort. The left sidebar shows the navigation menu with 'Templates' selected.

Click Create to finish setting up the Ethernet QoS policy.

## Server profile template network configuration – Ethernet Adapter

The ethernet adapter policy is used to set the interrupts and the send and receive queues. The values are set according to the best-practices guidance for the operating system in use. Cisco Intersight provides default VMware Ethernet Adapter policy for typical VMware deployments.

You can optionally configure a tweaked ethernet adapter policy for additional hardware receive queues handled by multiple CPUs in scenarios where there is a lot of vMotion traffic and multiple flows. In this deployment, a modified ethernet adapter policy, FlexPod-ASA-VMware-High-Traffic, is created and attached to the 03-vDS0-A and 04-vDS0-B interfaces which handle vMotion.

**Table 12 Ethernet adapter policy association to vNICs**

Policy Name	vNICs
FlexPod-ASA-Ethernet-Adapter-VMware	00-vSwitch0-A

	01-vSwitch0-B
FlexPod-ASA-Ethernet-Adapter-VMware-High-Traffic	02-vDS0-A 03-vDS0-B
FlexPod-ASA-Ethernet-Adapter-16RXQs-5G	04-iSCSI-A 05-iSCSI-B

1. Click Select Policy under Ethernet Adapter and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy. Optionally set Tags or provide a Description. Click Select Cisco Provided Configuration, find and select VMware policy.

The screenshot displays the 'Create Ethernet Adapter' configuration page in the Cisco UCS Manager. The left sidebar shows the navigation menu with 'Templates' selected. The main content area has two tabs: 'General' (active) and 'Policy Details'. The 'General' tab contains the following fields:

- Organization \***: FlexPod-ASA
- Name \***: FlexPod-ASA-Ethernet-Adapter-VMware
- Set Tags**: Enter a tag in the key:value format.
- Description**: Description (0 / 1024 characters)

Below the description field, the 'Cisco Provided Ethernet Adapter Configuration' section shows 'Selected Cisco Provided Configuration' as 'VMware'. There are links for 'Edit Selection' and a trash icon.

Click Next.

For the FlexPod-ASA-Ethernet-Adapter-VMware, click Create and skip the rest of the steps in this Ethernet Adapter policy creation section.

For the FlexPod-ASA-Ethernet-Adapter-VMware-High-Traffic policy (for vDS0 interfaces), make the following modifications after selecting the VMware policy:

- Increase Interrupts to 11
- Increase Receive Queue Count to 8
- Increase Receive Ring Size to 4096
- Increase Transmit Ring Size to 4096
- Increase Completion Queue Count to 9
- Enable Receive Side Scaling

Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy

## Create Ethernet Adapter

1 General

2 Policy Details

**General**

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-Ethernet-Adapter-VMware-High-Traffic

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

**Cisco Provided Ethernet Adapter Configuration**

Selected Cisco Provided Configuration VMware [Edit Selection](#)

Templates > Create UCS Server Profile Template > Create LAN Connectivity Policy

## Create Ethernet Adapter

General

2 Policy Details

**Interrupt Settings**

**Interrupts** **Interrupt Mode** **Interrupt Timer, us**

11 MSix 125

1 - 1024 0 - 65535

**Interrupt Coalescing Type**

Min

**Receive**

**Receive Queue Count** **Receive Ring Size**

8 4096

1 - 1000 64 - 16384

**Transmit**

**Transmit Queue Count** **Transmit Ring Size**

1 4096

1 - 1000 64 - 16384

**Completion**

**Completion Queue Count** **Completion Ring Size**

9 1

1 - 2000 1 - 256

For the FlexPod-ASA-Ethernet-Adapter-16RXQs-5G policy for iSCSI interfaces with 5<sup>th</sup> Generation VICs, make the following modifications after selecting the VMware policy:

- Increase Interrupts to 19
- Increase Receive Queue Count to 16
- Increase Receive Ring Size to 16384
- Increase Transmit Ring Size to 16384
- Increase Completion Queue Count to 17
- Enable Receive Side Scaling

Click Add.

**Note:** For all the non-iSCSI vNIC, skip the iSCSI-A and iSCSI-B policy creation sections below.

## Server profile template network configuration – iSCSI Boot A

The iSCSI-Boot-A policy is only applied to vNICs 05-iSCSI-A and should not be created for other vNICs. The iSCSI-Boot-B policy creation is explained next.

To create this policy, the following iSCSI target and LIF information will be needed from NetApp storage:

```
fpsa-a50-u0909::> iscsi show
```

Vserver	Target Name	Target Alias	Status Admin
svml	iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2	svml	up

```
fpsa-a50-u0909::> network interface show -data-protocol iscsi
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	iscsi-lif-01a	up/up	172.22.73.101/24	fpsa-a50-u0909-01	e2b-2273	true
	iscsi-lif-01b	up/up	172.22.74.101/24	fpsa-a50-u0909-01	e4b-2274	true
	iscsi-lif-02a	up/up	172.22.73.102/24	fpsa-a50-u0909-02	e2b-2273	true
	iscsi-lif-02b	up/up	172.22.74.102/24	fpsa-a50-u0909-02	e4b-2274	true

4 entries were displayed.

1. Click Select Policy under iSCSI Boot and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy.

Click Next.

Select Static under Configuration.

- Click Select Policy under Primary Target and then click Create New in the pane on the right. Verify correct organization is selected from the drop-down list and provide a name for the policy.



Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

### Create iSCSI Static Target

**1 General**

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-iSCSI-Boot-A-Primary-Target

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

Click Next.

Provide the Target Name captured from NetApp storage, IP Address of iscsi-lif-01a, Port 3260 and Lun ID of 0.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

### Create iSCSI Static Target

**2 Policy Details**

Add policy details.

This policy is applicable only for UCS Servers (FI-Attached)

**Configuration**

**Target Name \***

38ea2911ef9608d039eac6a795:vs.2

**IP Address \***

172.22.73.101

**Port \***

3260

**Lun ID \***

0

1 - 65535

Click Create.

Click Select Policy under Secondary Target and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

### Create iSCSI Static Target

**1 General**

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-iSCSI-Boot-A-Secondary-Target

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

Click Next.

Provide the Target Name captured from NetApp storage, IP Address of iscsi-lif-02a, Port 3260 and Lun ID of 0.



Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Static Target

Dashboards
Operate
Configure
Analyze
Automate

General

2 Policy Details

### Policy Details

Add policy details.

**1** This policy is applicable only for UCS Servers (FI-Attached)

#### Configuration

Target Name \*

IP Address \*

Port \*

38ea2911ef9608d039eac6a795:vs.2

172.22.73.102

3260

Lun ID \*

0

Click Create.

Click Select Policy under iSCSI Adapter and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Adapter

Dashboards
Operate
Configure
Analyze
Automate

1 General

2 Policy Details

### General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-iSCSI-Adapter

Set Tags

Enter a tag in the key:value format.

Description

Description

Click Next.

Accept the default policies. Customers can adjust the timers if necessary.

Click Create.

Scroll down to Initiator IP Source and make sure Pool is selected.

Click Select Pool under IP Pool and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the pool.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create IP Pool

Dashboards
Operate
Configure
Analyze
Automate
System

1 General

2 IPv4 Pool Details

3 IPv6 Pool Details

### General

Pool represents a collection of IPv4 and/or IPv6 addresses that can be allocated to other configuration entities like server profiles.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-iSCSI-A-IP-Pool

Set Tags

Enter a tag in the key:value format.

Description

Description

☐ Configure Subnet at Block Level

Click Next.

Make sure Configure IPv4 Pool is selected. Enter the IP pool information for iSCSI-A subnet.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

### Create IP Pool

**General**

**IPv4 Pool Details**  
Network interface configuration data for IPv4 interfaces.

☒ Configure IPv4 Pool

**Configuration**

Netmask \* ①: 255.255.255.0

Gateway ①: 0.0.0.0

Primary DNS ①: Primary DNS

Secondary DNS ①: Secondary DNS

**IP Blocks**

[Add IP Blocks](#)

IP Block
<div>From ①: 172.22.73.30</div> <div>Size ①: 16</div> <div>To ①: 172.22.73.47</div>

1 - 1024

[Close](#) [Back](#) [Next](#)

**Note:** Since the iSCSI network is not routable but the Gateway parameter is required, enter 0.0.0.0 for the Gateway. This will result in a gateway not being set for the interface.

Click Next.

Disable Configure IPv6 Pool.

Click Create.

Verify all the policies and pools are correctly mapped for the iSCSI-Boot-A policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity)

### Create iSCSI Boot

**General**

**Policy Details**

☐ Auto ☒ Static

**Primary iSCSI Static Target \*** ①

Selected Policy: FlexPod-ASA-iSCSI-Boot-A-Pri... [Edit Selection](#)

**Secondary iSCSI Static Target** ①

Selected Policy: FlexPod-ASA-iSCSI-Boot-A-Se... [Edit Selection](#)

**iSCSI Adapter** ①

Selected Policy: FlexPod-ASA-iSCSI-Adapter [Edit Selection](#)

**Authentication**

☐ CHAP ①

☐ Mutual CHAP ①

**Initiator IP Source**

☒ Pool ☐ DHCP ☐ Static

**IP Pool \*** ①

Selected Pool: FlexPod-ASA-iSCSI-A-IP-Pool [Edit Selection](#)

[Cancel](#) [Back](#) [Create](#)

Click Create.

## Server profile template network configuration – iSCSI Boot B

The iSCSI-Boot-B policy is only applied to vNIC 06-iSCSI-B and should not be created for other vNICs. The iSCSI Target and LIF information obtained from NetApp storage in the previous section will be utilized.

1. Click Select Policy under iSCSI Boot and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

The screenshot shows the 'Create iSCSI Boot' configuration page. The breadcrumb trail is 'Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity)'. The left sidebar has 'Configure' selected. The main panel has two tabs: 'General' (active) and 'Policy Details'. The 'General' tab contains the following fields:

- Organization \***: A dropdown menu with 'FlexPod-ASA' selected.
- Name \***: A text input field with 'FlexPod-ASA-iSCSI-Boot-B' entered.
- Set Tags**: A text input field with the placeholder 'Enter a tag in the key:value format.'
- Description**: A text input field with 'Description' entered.

A character count '0 / 1024' is visible at the bottom right of the description field.

Click Next.

Select Static under Configuration.

The screenshot shows the 'Create iSCSI Boot' configuration page with the 'Policy Details' tab selected. The breadcrumb trail is the same as the previous screenshot. The left sidebar has 'Configure' selected. The main panel shows the 'Policy Details' tab with the following content:

- Policy Details**: A section header with the instruction 'Add policy details.'
- Configuration**: A section header with a note: 'This policy is applicable only for UCS Servers (FI-Attached)'.
- Configuration**: A section with two tabs: 'Auto' and 'Static' (selected).
- Primary iSCSI Static Target \***: A section with a 'Select Policy' link.
- Secondary iSCSI Static Target**: A section with a 'Select Policy' link.
- iSCSI Adapter**: A section with a 'Select Policy' link.
- Authentication**: A section header.

Click Select Policy under Primary Target and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-ISCST-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-ISCST-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Static Target

1

General

2

Policy Details

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-ISCST-Boot-B-Primary-Target

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Click Next.

Provide the Target Name captured from NetApp storage, IP Address of iscsi-lif-01b, Port 3260 and Lun ID of 0.

Edit UCS Server Profile (FlexPod-ASA-AMD-ISCST-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-ISCST-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Static Target

✓

General

2

Policy Details

**Policy Details**  
Add policy details.

**Configuration**

**Target Name \*** 38ea2911ef9608d039eac6a795:vs.2

**IP Address \*** 172.22.74.101

**Port \*** 3260

**Lun ID \*** 0

1 - 65535

Click Create.

Click Select Policy under Secondary Target and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-ISCST-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-ISCST-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Static Target

1

General

2

Policy Details

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-ISCST-Boot-B-Secondary-Target

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description

0 / 1024

Click Next.

Provide the Target Name captured from NetApp storage, IP Address of iscsi-lif-02b, Port 3260 and Lun ID of 0.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create iSCSI Static Target

Dashboards
Operate
Configure
Analyze
Automate

General

Policy Details

Add policy details.

This policy is applicable only for UCS Servers (FI-Attached)

### Configuration

Target Name \* ⓘ 38ea2911ef9608d039eac6a795:vs.2 ⓘ

IP Address \* ⓘ 172.22.74.102 ⓘ

Port \* ⓘ 3260 ⓘ

Lun ID ⓘ 0 ⓘ

1 - 65535

Click Create.

Click Select Policy under iSCSI Adapter and then select the previously configured adapter policy FlexPod-ASA-iSCSI-Adapter in the pane on the right.

Scroll down to Initiator IP Source and make sure Pool is selected.

Click Select Pool under IP Pool and then click Create New in the pane on the right.

Verify correct organization is selected from the drop-down list and provide a name for the pool.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create IP Pool

Dashboards
Operate
Configure
Analyze
Automate
System

General

IPv4 Pool Details

IPv6 Pool Details

### General

Pool represents a collection of IPv4 and/or IPv6 addresses that can be allocated to other configuration entities like server profiles.

Organization \* FlexPod-ASA ⓘ

Name \* FlexPod-ASA-iSCSI-B-IP-Pool ⓘ

Set Tags Enter a tag in the key:value format. ⓘ

Description Description ⓘ

0 / 1024

☐ Configure Subnet at Block Level ⓘ

Click Next.

Make sure Configure IPv4 Pool is selected. Enter the IP pool information for iSCSI-B subnet.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity) > Create iSCSI Boot

## Create IP Pool

✓ General

2 IPv4 Pool Details

3 IPv6 Pool Details

### IPv4 Pool Details

Network interface configuration data for IPv4 interfaces.

☒ Configure IPv4 Pool

#### Configuration

Netmask \* ⓘ

Gateway ⓘ

Primary DNS ⓘ

Secondary DNS ⓘ

#### IP Blocks

[Add IP Blocks](#)

— IP Block

From ⓘ

Size ⓘ

1 - 1024

[Close](#) [Back](#) [Next](#)

**Note:** Since the iSCSI network is not routable but the Gateway parameter is required, enter 0.0.0.0 for the Gateway. This will result in a gateway not being set for the interface.

Click Next.

Disable Configure IPv6 Pool.

Click Create.

Verify all the policies and pools are correctly mapped for the iSCSI-Boot-B policy.

Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot) > Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity)

## Create iSCSI Boot

✓ General

2 Policy Details

Auto

Static

Primary iSCSI Static Target \* ⓘ

Selected Policy FlexPod-ASA-iSCSI-Boot-B-Pri... ⓘ | [Edit Selection](#) | [Delete](#)

Secondary iSCSI Static Target ⓘ

Selected Policy FlexPod-ASA-iSCSI-Boot-B-Se... ⓘ | [Edit Selection](#) | [Delete](#)

iSCSI Adapter ⓘ

Selected Policy FlexPod-ASA-iSCSI-Adapter ⓘ | [Edit Selection](#) | [Delete](#)

#### Authentication

☐ CHAP ⓘ

☐ Mutual CHAP ⓘ

#### Initiator IP Source

☒ Pool ☐ DHCP ☐ Static

IP Pool \* ⓘ

Selected Pool FlexPod-ASA-iSCSI-B-IP-Pool ⓘ | [Edit Selection](#) | [Delete](#)

[Cancel](#) [Back](#) [Create](#)

Click Create.

Click Create to finish creating the vNIC.

Go back to the Ethernet Adapter section and repeat the vNIC creation for all six vNICs.

**Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot)**

**Edit LAN Connectivity Policy (FlexPod-ASA-iSCSI-Boot-LAN-Connectivity)**

General | **Policy Details**

IQN Pool ⓘ

Selected Pool FlexPod-ASA-IQN-Pool ⓘ Edit Selection

**vNIC Configuration**

Manual vNICs Placement | Auto vNICs Placement

Add ▾

Graphic vNICs Editor

Q Search Filters 6 results Export

Name	Slot ID	Switch ID	PCI Order	Failover	MAC Pool
00-vSwitch0-A	MLOM	A	0	Disabled	FlexPod-ASA-MAC-Pool-A
01-vSwitch0-B	MLOM	B	1	Disabled	FlexPod-ASA-Mac-Pool-B
02-vDS0-A	MLOM	A	2	Disabled	FlexPod-ASA-MAC-Pool-A
03-vDS0-B	MLOM	B	3	Disabled	FlexPod-ASA-Mac-Pool-B
04-iSCSI-A	MLOM	A	4	Disabled	FlexPod-ASA-MAC-Pool-A
05-iSCSI-B	MLOM	B	5	Disabled	FlexPod-ASA-Mac-Pool-B

Click Create to finish creating the LAN Connectivity policy for iSCSI hosts.

## Derive Server Profiles

Follow the steps below to derive server profiles from the created server profile template.

1. From the Server profile template Summary screen, click Derive Profiles.

**Edit UCS Server Profile (FlexPod-ASA-AMD-iSCSI-Boot)**

General | Compute Configuration | Management Configuration | Storage Configuration | Network Configuration | **Summary**

**Summary**

Verify details of the template and the policies, resolve errors and deploy.

General

Name FlexPod-ASA-AMD-iSCSI-Boot Organization FlexPod-ASA

Target Platform UCS Server (FI-Attached)

Compute Configuration | Management Configuration | Storage Configuration | Network Configuration | Errors/Warnings (0)

BIOS	FlexPod-ASA-AMD-M8-Virt-BIOS
Boot Order	FlexPod-ASA-iSCSI-Boot-Order
UUID	FlexPod-ASA-UUID-Pool
Virtual Media	FlexPod-ASA-KVM-Mount-Media

Close Back Derive Profiles

**Note:** This action can also be performed later by navigating to Templates, clicking the dots next to the template name and selecting Derive Profiles.

Under the Server Assignment, select Assign Now and select Cisco UCSX-215C-M8 server(s). Customers can select one or more servers depending on the number of profiles to be deployed.



UCS Server Profile Templates > FlexPod-ASA-AMD-iSCSI-Boot

## Derive

1 General

2 Details

3 Summary

FlexPod-ASA-AMD-iSCSI-Boot

FlexPod-ASA

Target Platform

UCS Server (FI-Attached)

Server Assignment

Assign Now

From a Resource Pool

Chassis Slot Location

Serial Number

Assign Later

Listed servers have been fully discovered and are available for assignment. Learn about [Configuring UCS Server Profiles](#).

Search

Filters 4 results

Export

<input type="checkbox"/>	Name	User La...	Health	Model	UCS Domain	Ser
<input checked="" type="checkbox"/>	fpsa-x9508-u0901-fi-1-1		Healthy	UCSX-215C-M8	fpsa-x9508-u	FCI
<input checked="" type="checkbox"/>	fpsa-x9508-u0901-fi-1-2		Healthy	UCSX-215C-M8	fpsa-x9508-u	FCI
<input type="checkbox"/>	fpsa-x9508-u0901-fi-1-3		Healthy	UCSX-215C-M8	fpsa-x9508-u	FCI
<input type="checkbox"/>	fpsa-x9508-u0901-fi-1-4		Healthy	UCSX-215C-M8	fpsa-x9508-u	FCI

Selected 2 of 4

Show Selected

Unselect All

Rows per page 10

1

Cancel

Next

Click Next.

**Note:** Cisco Intersight will fill in the default information based on the number of servers selected.

Adjust the fields as needed. It is recommended to use the server hostname for the Server Profile name.

UCS Server Profile Templates > FlexPod-ASA-AMD-iSCSI-Boot

## Derive

1 General

2 Details

3 Summary

FlexPod ASA iSCSI SAN Boot

26 / 1024

Set Tags

Enter a tag in the key:value format.

Derive

Profile Name Prefix

Digits Count

Start Index for Suffix

fpsa-asa-esxi-0

1

1

1 Name \*

Organization \*

Assigned Server

fpsa-asa-esxi-01

FlexPod-ASA

fpsa-x9508-u0901-fi-1-1

2 Name \*

Organization \*

Assigned Server

fpsa-asa-esxi-02

FlexPod-ASA

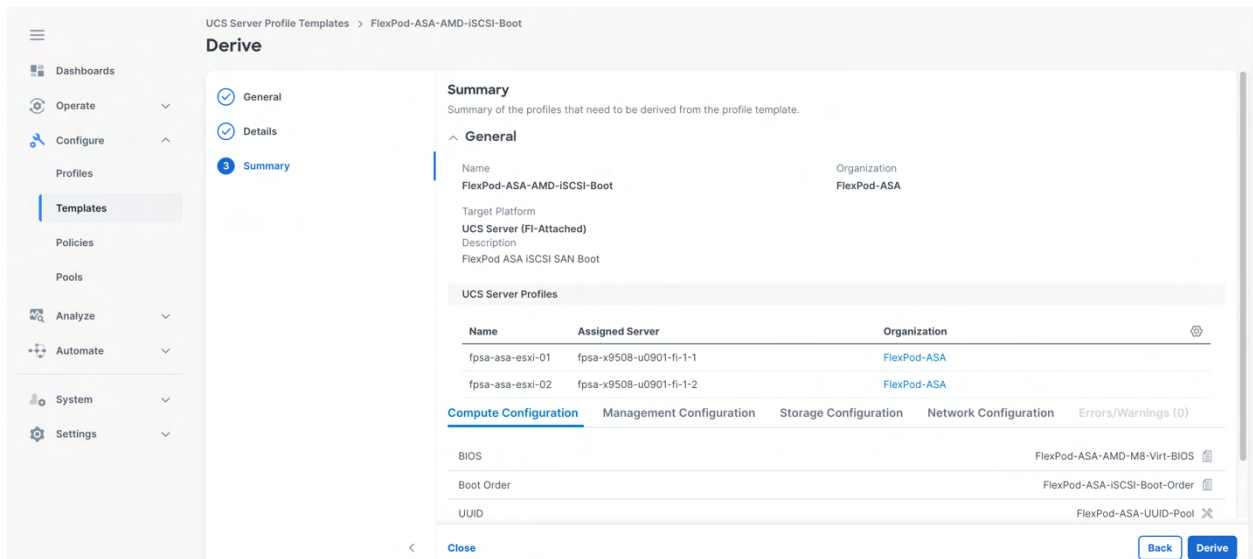
fpsa-x9508-u0901-fi-1-2

Close

Back

Next

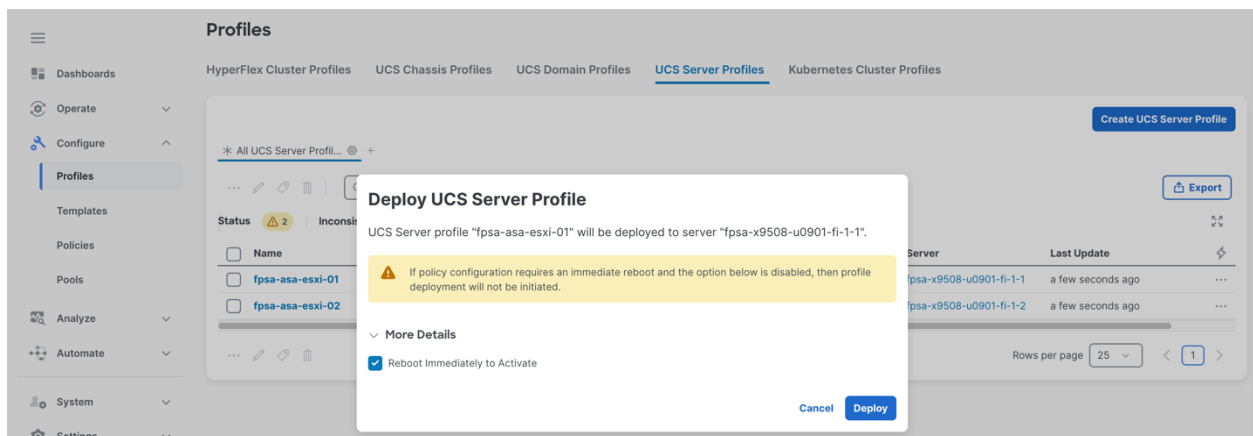
Click Next.



Verify the information and click Derive to create the Server Profile(s).

Navigate to Configure > Profiles > UCS Server Profiles, select the profile(s) just created and click the dots at the top of the column and select Deploy.

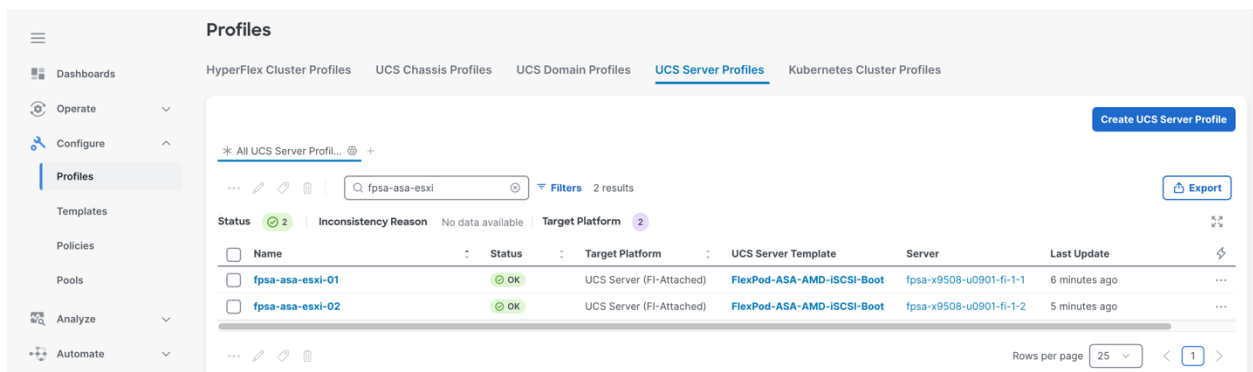
Check the box for Reboot Immediately to Activate and click Deploy to confirm.



Cisco Intersight will start deploying the server profile(s) and it will take some time to apply all the policies.

Click the Requests icon at the top right-hand corner of the window to see the request and monitor its progress.

When the Server Profile(s) are deployed successfully, they will appear under the Server Profiles with the status of OK.



**Note:** You can enter part of the Server Profile name in the search box to filter the results as shown in the screenshot above.

## VMware ESXi 8.0U3 installation

This section provides instructions for installing VMware ESXi 8.0U3 in a FlexPod environment. On successful completion of these steps, multiple ESXi hosts will be provisioned and ready to be added to VMware vCenter.

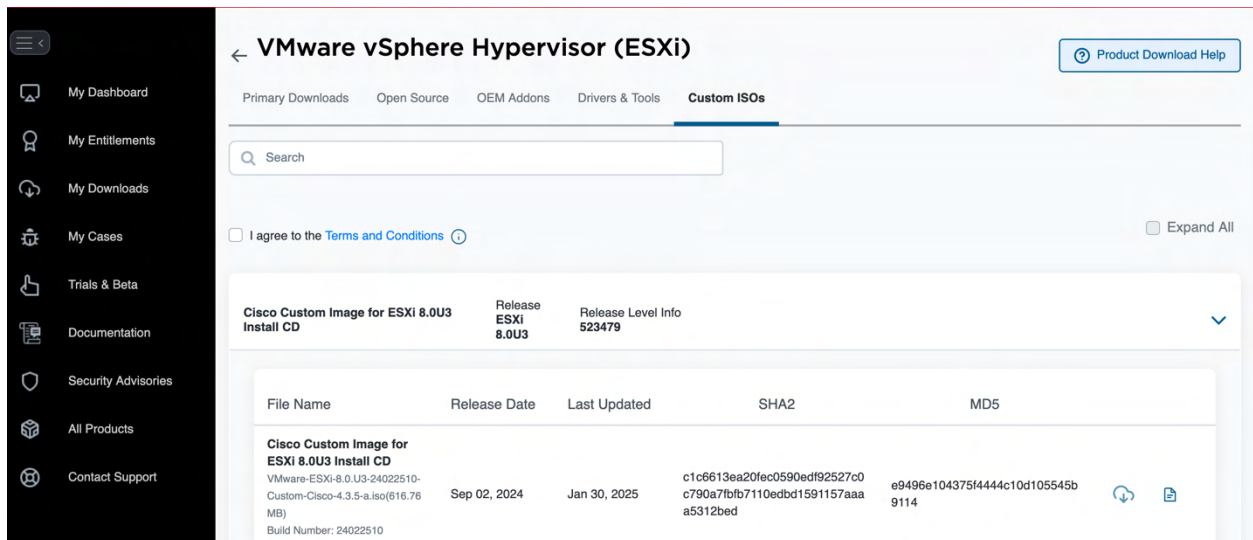
Several methods exist for installing ESXi in a VMware environment. These procedures focus using the built-in keyboard, video, mouse (KVM) console and virtual media features in Intersight to map installation media to the individual servers.

## Download ESXi 8.0U3 from Broadcom

1. Login to Broadcom support portal download site, <https://support.broadcom.com/group/ecx/downloads>.

**Note:** You will need a Broadcom user id and password.

Look for VMware vSphere Hypervisor (ESXi) Custom ISOs.



The screenshot shows the Broadcom support portal interface for downloading VMware vSphere Hypervisor (ESXi) Custom ISOs. The page title is "VMware vSphere Hypervisor (ESXi)". The navigation bar includes "Primary Downloads", "Open Source", "OEM Addons", "Drivers & Tools", and "Custom ISOs". A search bar is present. Below the search bar, there is a checkbox for "I agree to the Terms and Conditions" and a link to "Expand All". The main content area displays the "Cisco Custom Image for ESXi 8.0U3 Install CD" with release information: "Release ESXi 8.0U3" and "Release Level Info 523479". A table lists the download links with columns: "File Name", "Release Date", "Last Updated", "SHA2", and "MD5".

File Name	Release Date	Last Updated	SHA2	MD5
Cisco Custom Image for ESXi 8.0U3 Install CD				
VMware-ESXi-8.0.U3-24022510-Custom-Cisco-4.3.5-a.iso(616.76 MB)	Sep 02, 2024	Jan 30, 2025	c1c6613ea20fec0590edf92527c0c790a7fbfb7110edbd1591157aaa5312bed	e9496e104375f4444c10d105545b9114
Build Number: 24022510				

Download the Cisco Custom Image for ESXi 8.0U3 Install CD.

**Note:** For this solution validation, the VMware-ESXi-8.0.U3-244022510-Custom-Cisco-4.3.5-a.iso image was used.

## Login to Intersight and access KVM

The KVM access available in Intersight enables administrators to begin the installation of the operating system (OS) through mounted virtual media. It is necessary to log into the Cisco Intersight to access KVM.

1. Login to Intersight.

Navigate to Operate > Servers.

Find the Server with the desired Server Profile assigned and click the three dots on the right to see more options.

Click launch vKVM.

**Note:** Since the Cisco Custom ISO image will be mapped to the vKVM, it is important to use the standard vKVM and not the Tunneled vKVM. Be sure that Intersight is being accessed from a subnet that has direct access to the subnet where the vKVM IPs are provisioned on.

Follow the prompts to ignore certificate warnings (if any) and launch the HTML5 KVM console.  
Repeat steps 1 - 5 to launch the HTML5 KVM console for all the ESXi servers.

## Set up VMware ESXi installation

Follow these steps on each ESXi host to prepare the server for OS Installation.

1. In the KVM window, navigate to Virtual Media > vKVM-Mapped vDVD.
2. Browse and select the ESXi installer ISO image file downloaded previously.
3. Click Map Drive.
4. Select Boot Device > vKVM-Mapped vDVD.
5. Click Confirm.
6. Select Power > Reset System and Confirm to reboot the Server if the server is showing shell prompt. If the server is shutdown, select Power > Power On System.
7. Monitor the server boot process in the KVM. The server should find the mapped LUN during boot and begin to load the ESXi installer.

**Note:** If the ESXi installer fails to load because the software certificates cannot be validated, reset the server, and when prompted, press F2 to go into BIOS and set the system time and date to current. The ESXi installer should load properly.

## Install VMware ESXi onto the mapped boot LUN

Repeat the following steps for each ESXi host.

1. After the ESXi installer finishes loading (from the last step), press Enter to continue with the installation.

Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

**Note:** You can navigate to Tools > Keyboard to open virtual keyboard for easy access to function keys.

Select the NetApp boot LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

Select the appropriate keyboard layout and press Enter to continue.

Enter and confirm the root password and press Enter to continue.

The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.

After the installation is complete, press Enter to reboot the server. The ISO image will be unmapped automatically.

## Set up management networking for ESXi hosts

To configure ESXi host with access to the management network, follow these steps on each ESXi host.

1. After the server has finished rebooting, press F2 to customize VMware ESXi in the KVM console.

Log in as root, enter the password configured during installation, and press Enter to log in.

Use the down arrow key to select Troubleshooting Options and press Enter.

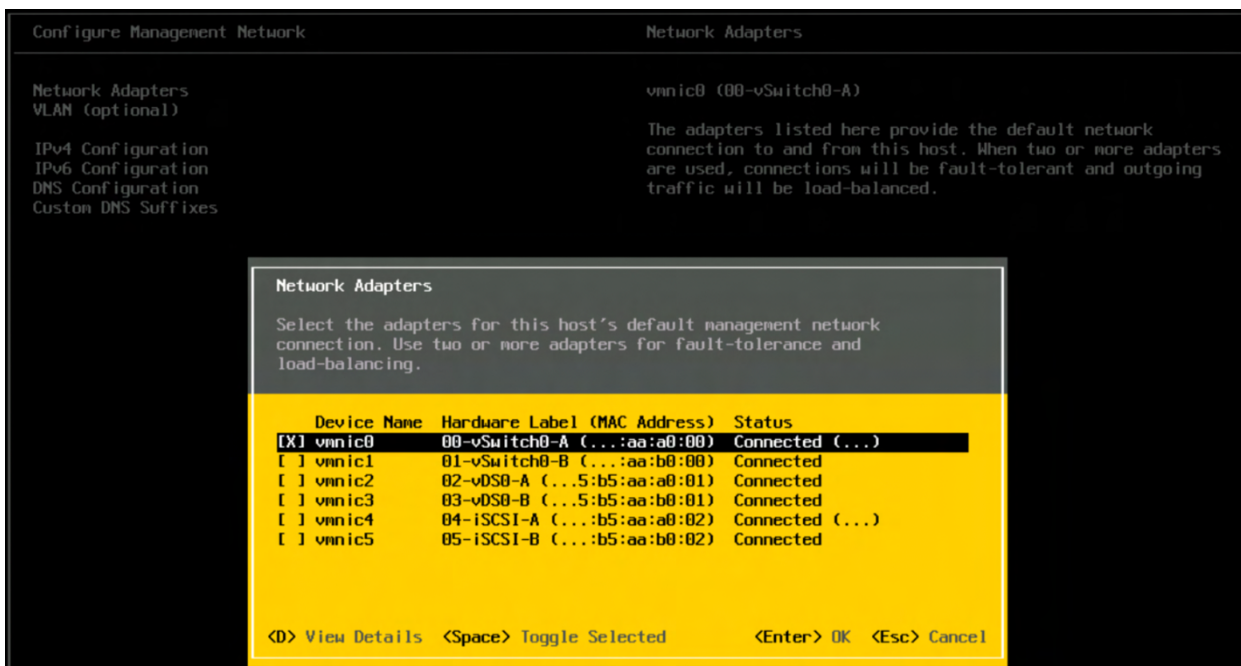
Select Enable ESXi Shell and press Enter.

Select Enable SSH and press Enter.

Press Esc to exit the Troubleshooting Options menu.

Select the Configure Management Network option and press Enter.

Select Network Adapters and press Enter. Ensure the vmnic numbers align with the numbers under the Hardware Label (for example, vmnic0 and 00-vSwitch0-A).



Press Enter.

Select the VLAN option and press Enter.

Enter the IB-MGMT VLAN and press Enter.

Select IPv4 Configuration and press Enter.

**Note:** When using DHCP to set the ESXi host networking configuration, manual IP address configuration is not required.

Select the Set static IPv4 address and network configuration option by using the arrow keys and space bar.

Enter the IP address, subnet mask, and default gateway for managing the ESXi host.

Press Enter to accept the changes to the IP configuration.

Select the IPv6 Configuration option and press Enter.

Using the spacebar, select Disable IPv6 (restart required) and press Enter.

Select the DNS Configuration option and press Enter.

**Note:** If the IP address is configured manually, the DNS information must be provided.

Using the spacebar, select Use the following DNS server addresses and hostname:

Under Primary DNS Server, enter the IP address of the primary DNS server.

Optionally provide the Alternate DNS Server IP address.

Under Hostname, enter the fully qualified domain name (FQDN) for the ESXi host.

Press Enter to accept the changes to the DNS configuration.

Press Escape key to exit the Configure Management Network submenu.

Press Y to confirm the changes and reboot the ESXi host.

## Reset VMware ESXi host VMkernel port vmk0 MAC address (optional)

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port on which it is placed. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will occur because vmk0 retains the assigned MAC

address unless the ESXi system configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, complete the following steps:

1. From the ESXi console menu main screen, select Macros > Static Macros > Ctrl + Alt + F's > Ctrl + Alt + F1 to access the VMware console command line interface.

Login as root.

Enter the `esxcfg-vmknic` command below to get a detailed listing of interface vmk0. vmk0 should be a part of the Management Network port group. Note the IP address and network mask of vmk0.

```
esxcfg-vmknic -l
```

To remove vmk0, enter the following command:

```
esxcfg-vmknic -d "Management Network"
```

To add vmk0 again with a random MAC address, enter the following command:

```
esxcfg-vmknic -a -i <vmk0_ip> -n <vmk0_netmask> "Management Network".
```

Verify that vmk0 has been added again with a random MAC address:

```
esxcfg-vmknic -l
```

Tag vmk0 as the management interface:

```
esxcli network ip interface tag add -i vmk0 -t Management
```

When vmk0 was re-added, if a message popped up saying vmk1 was marked as the management interface, remove it by the following command:

```
esxcli network ip interface tag remove -i vmk1 -t Management
```

Enter exit to logout of the ESXi console.

Select Macros > Static Macros > Ctrl + Alt + F's > Ctrl + Alt + F2 to return to the VMware ESXi menu.

## Update Cisco VIC drivers (optional)

During ESXi 8.0U3 installation, the Cisco VIC nenic version 2.0.15.0 from the Cisco Custom ISO for VMware ESXi version 8.0U3 is installed on the system.

Consult the [Cisco UCS Hardware Compatibility List](#) to determine the latest supported firmware and driver and take the following steps if the updates are needed for the VIC nenic Ethernet driver.

1. Download the Cisco VIC driver ISO and extract the component bundle.

Using an SCP program to copy the component bundle to the /tmp directory on each ESXi host.

SSH to each VMware ESXi host and login as root.

Run the `esxcli` commands on each host to apply the component bundle. See below for an example command format.

```
esxcli software component apply -d /tmp/Cisco-nenic_xxx.zip
```

Reboot the host.

After reboot, SSH back into each host and use the following commands to ensure the correct version is installed:

```
esxcli software component list | grep nenic
```

**Note:** If updates to the VIC nfnic Fibre Channel driver is needed, download the driver bundle and follow procedures like the above to apply the update. When checking for the updated nfnic driver, grep for nfnic instead.

## Initial ESXi host configuration

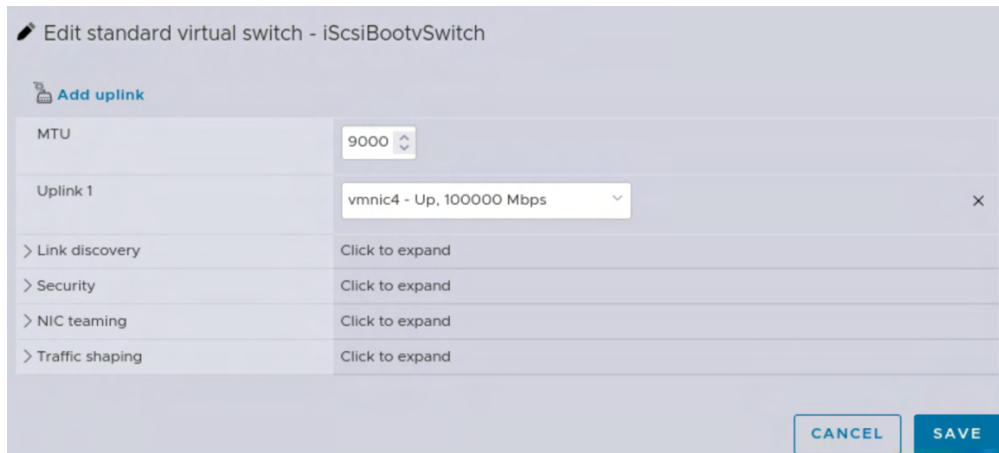
In this procedure, you're only setting up the first ESXi host. The remaining hosts will be added to vCenter and configured from vCenter.

### Login to the first ESXi host using the VMware Host Client


1. Open a web browser and navigate to the first ESXi server's management host name or IP address. For the security risk warning, select Advanced and Accept the Risk and continue. Provide the username and password for the root user and click LOGIN. Decide whether to join the VMware Customer Experience Improvement Program or not and click OK.

### Set Up iSCSI VMkernel Ports and Virtual Switch

1. From the Web Navigator, click Networking. In the center pane, select the Virtual switches tab. Right-click the iScsiBootvSwitch link under the Name column and select Edit settings. Change the MTU to 9000. Click Save to save the changes to iScsiBootvSwitch.



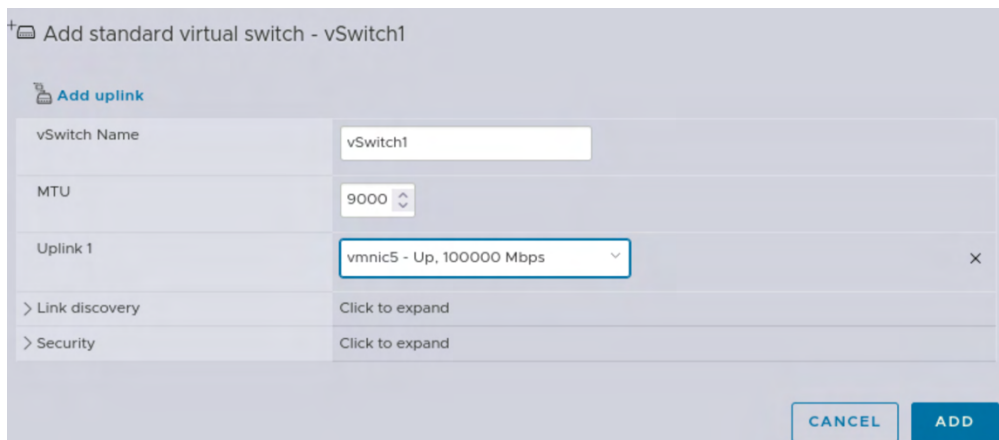
Edit standard virtual switch - iScsiBootvSwitch

 Add uplink


MTU	9000
Uplink 1	vmnic4 - Up, 100000 Mbps
> Link discovery	Click to expand
> Security	Click to expand
> NIC teaming	Click to expand
> Traffic shaping	Click to expand

CANCEL SAVE

- Select Add standard virtual switch.
- Name the switch vSwitch1.
- Change the MTU to 9000.
- From the drop-down list select vmnic5 for Uplink 1.



+ Add standard virtual switch - vSwitch1

 Add uplink

vSwitch Name	vSwitch1
MTU	9000
Uplink 1	vmnic5 - Up, 100000 Mbps
> Link discovery	Click to expand
> Security	Click to expand

CANCEL ADD



Click ADD to add vSwitch1.

In the center pane, select the VMkernel NICs tab.

Right-click the iScsiBootPG link and select Edit settings.

Change the MTU to 9000.

Expand IPv4 Settings and enter a unique IP address in the iSCSI-A subnet but outside of the Cisco Intersight FlexPod-ASA-iSCSI-A-IP-Pool.

**Note:** It is important to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments in Cisco UCS.

Uncheck Management under Services.

Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.22.73.11
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

CANCEL SAVE

Click SAVE to save the changes to iScsiBootPG VMkernel NIC.

Select Add VMkernel NIC.

For New port group, enter iScsiBootPG-B.

For Virtual switch, select vSwitch1 from the drop-down list.

Change the MTU to 9000.

For IPv4 settings, select Static.

Expand IPv4 Settings and enter a unique IP address and Subnet mask in the iSCSI-B subnet but outside of the Cisco Intersight FlexPod-ASA-iSCSI-B-IP-Pool.

**Add VMkernel NIC**

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	vSwitch1
VLAN ID	0
MTU	9000
IP version	IPv4 only
IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	172.22.74.11
Subnet mask	255.255.255.0
TCP/IP stack	Default TCP/IP stack
Services	<input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication

**CANCEL CREATE**

**Note:** Leave the VLAN ID as 0 since the vNIC already has iSCSI-B network configured as native VLAN in UCS.

Click **CREATE** to complete creating the VMkernel NIC.

In the center pane, select the Port groups tab.

Right-click the iScsiBootPG link and Edit settings.

Change the Name to iScsiBootPG-A.

Click **SAVE** to complete editing the port group name.

## Configure software iSCSI

1. On the left select Storage, then in the center pane select the Adapters tab.

Select Software iSCSI to configure software iSCSI for the host.

In the Configure iSCSI window, under Dynamic targets, click Add dynamic target.

Select Click to add address and enter the IP address of iscsi-lif-01a from svm1. Press Enter.

Repeat steps 3-4 to add the IP addresses for iscsi-lif-02a, iscsi-lif-01b, and iscsi-lif-02b.

Click **SAVE CONFIGURATION**.

Click Software iSCSI again open configuration window for iSCSI software adapter.

Verify that four static targets and four dynamic targets are listed for the host.

Configure iSCSI - vmhba64

ISCSI enabled ☐ Disabled ☒ Enabled

> Name and alias `iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2 (iscsi_vmk)`

> CHAP authentication `Do not use CHAP`

> Mutual CHAP authentication `Do not use CHAP`

> Advanced settings `Click to expand`

Network port bindings

Add port binding Remove port binding

VMkernel NIC	Port group	IPv4 address
No port bindings		

Static targets

Add static target Remove static target Edit settings

Target	Address	Port
<code>iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2</code>	<code>172.22.73.101</code>	<code>3260</code>
<code>iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2</code>	<code>172.22.73.102</code>	<code>3260</code>
<code>iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2</code>	<code>172.22.74.102</code>	<code>3260</code>
<code>iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2</code>	<code>172.22.74.101</code>	<code>3260</code>

Dynamic targets

Add dynamic target Remove dynamic target Edit settings

Address	Port
<code>172.22.73.101</code>	<code>3260</code>
<code>172.22.73.102</code>	<code>3260</code>
<code>172.22.74.101</code>	<code>3260</code>
<code>172.22.74.102</code>	<code>3260</code>

Click Cancel to close the window.

**Note:** If the host shows an alarm stating that connectivity with the boot disk was lost, place the host in Maintenance Mode and reboot the host.

## Set up in-band management VMkernel Ports and Virtual Switch

1. From the Host Client Navigator, select Networking.

In the center pane, select the Virtual switches tab.

Right-click the vSwitch0 link and then select Edit settings.

Leave the MTU at 1500.


**Note:** The vSwitch0 is only used for in-band management with 1500 MTU.


Click Add uplink.

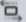

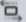

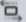

If vmnic1 is not selected for Uplink 2, then use the pulldown to select vmnic1.

Expand NIC teaming.

In the Failover order section, if vmnic1 does not have a status of Active, select vmnic1 and click Mark active and verify that vmnic1 now has a status of Active.

 Edit standard virtual switch - vSwitch0

 [Add uplink](#)

MTU	1500									
Uplink 1	vmnic0 - Up, 100000 Mbps <span>✕</span>									
Uplink 2	vmnic1 - Up, 100000 Mbps <span>✕</span>									
> Link discovery	Click to expand									
> Security	Click to expand									
▼ NIC teaming										
Load balancing	Route based on originating port ID									
Network failover detection	Link status only									
Notify switches	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failback	<input checked="" type="radio"/> Yes <input type="radio"/> No									
Failover order	<div> <span>✕ Mark standby</span> <span>↑ Move up</span> <span>↓ Move down</span> </div> <table border="1"> <thead> <tr> <th>Name</th> <th>Speed</th> <th>Status</th> </tr> </thead> <tbody> <tr> <td> vmnic0</td> <td>100000 Mbps, full duplex</td> <td>Active</td> </tr> <tr> <td> vmnic1</td> <td>100000 Mbps, full duplex</td> <td>Active</td> </tr> </tbody> </table>	Name	Speed	Status	 vmnic0	100000 Mbps, full duplex	Active	 vmnic1	100000 Mbps, full duplex	Active
Name	Speed	Status								
 vmnic0	100000 Mbps, full duplex	Active								
 vmnic1	100000 Mbps, full duplex	Active								
> Traffic shaping	Click to expand									

CANCEL SAVE

Click SAVE.

Select Networking and then select the Port groups tab.

In the center pane, right-click VM Network and select Edit settings.

Name the port group IB-MGMT Network and set the proper VLAN ID for the IB-MGMT VLAN.

Click SAVE to finalize the edits for the IB-MGMT Network port group.

Edit port group - VM Network

Name	IB-MGMT Network
VLAN ID	2272
Virtual switch	vSwitch0
> Security	Click to expand
> NIC teaming	Click to expand
> Traffic shaping	Click to expand

CANCEL
SAVE

Select the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 should look similar to the following screenshot.

**vSwitch0**

Type: Standard vSwitch  
 Port groups: 2  
 Uplinks: 2

<b>vSwitch Details</b>	
MTU	1500
Ports	8570 (8553 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1
<b>NIC teaming policy</b>	
Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes
<b>Security policy</b>	
Allow promiscuous mode	No
Allow forged transmits	No
Allow MAC changes	No
<b>Shaping policy</b>	
Enabled	No

**vSwitch topology**

IB-MGMT Network  
VLAN ID: 2272

Management Network  
VLAN ID: 2272  
VMkernel ports (1)  
vmmic0: 172.22.72.11

Physical adapters  
vmmic1: 100000 Mbps, Full  
vmmic0: 100000 Mbps, Full

## Create iSCSI datastore

To create a new iSCSI datastores, complete the following steps:

1. From the Host Client Navigator, select Storage.

In the center pane, select the Datastore tab and click New datastore.

Right-click on the cluster and select New Datastore under the Storage menu.

Select Create new VMFS datastore for the creation type and click Next.

Provide a datastore name and select a device that is unclaimed.

New datastore - iscsi\_datastore\_1

1 Select creation type
2 **Select device**
3 Select partitioning options
4 Ready to complete

### Select device

Select a device on which to create a new VMFS partition

Name

The following devices are unclaimed and can be used to create a new VMFS datastore

Name	Type	Capacity	Free space
NETAPP iSCSI Disk (naa.600a098038323448723f5...	Disk (SSD)	1,024 GB	1,024 GB

1 items

Click NEXT.

Keep the partition option of Use full disk and VMFS 6 and click NEXT.

New datastore - iscsi\_datastore\_1

1 Select creation type
2 Select device
3 **Select partitioning options**
4 Ready to complete

### Ready to complete

Summary

Name	iscsi_datastore_1
Disk	NETAPP iSCSI Disk (naa.600a098038323448723f5877434a5250)
Partitioning	Use full disk
VMFS version	6

VMFS (1,024 GB)

CANCEL BACK NEXT FINISH

Review the summary and click Finish to create the datastore iscsi\_datastore\_1.

Select YES for the warning of disk content about to be erased and replaced with specified configuration.

Review the datastore information from ESXi Host Client > Storage > Datastores.

vm ESXi Host Client

root@fpga-asa-esxi-01 | Help | Search

Navigation: Host, Virtual Machines, Storage, Networking

Storage: 0 Virtual Machines, 1 Storage, 1 Networking

fpga-asa-esxi-01.nva.local - Storage

Datastores Adapters Devices Persistent Memory

+ New datastore + Increase capacity + Register a VM + Datastore browser + Refresh + Actions

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
iscsi_datastore_1	SSD	1,023.75 GB	1.43 GB	1,022.32 GB	VMFS6	Supported	Single

1 items



## Configure NTP Servers

1. From the ESXi Host Client Navigator, select Manage.
2. In the center pane, click System > Time & date.
3. Click Edit NTP Settings.
4. Select Use Network Time Protocol (enable NTP client).
5. Use the drop-down list to select Start and stop with host.
6. Enter the NTP server IP addresses separated by commas in the NTP servers.

**Note:** Use the IP addresses of the in-band management NTP Distribution Interfaces configured in the Nexus switches if applicable.

**Edit NTP Settings**

Specify how the date and time of this host should be set.

☐ Manually configure the date and time on this host

03/04/2025 3:47 PM

☒ Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 172.21.62.121, 172.21.62.122

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

**CANCEL** **SAVE**

Click **SAVE** to save the configuration changes.

Select the **Services** tab.

Right-click **ntpd** and click **Start**.

Navigate to **Manage > System > Time & date** and the NTP service status should now show a status of **Running**.

**ESXi Host Client**

**Navigator**

- Host
- Manage
- Monitor
- Virtual Machines (0)
- Storage (1)
- Networking (1)
  - vSwitch0
  - More networks...

**Advanced settings**

- Autostart
- Swap
- Time & date**

**Time & date**

- Edit NTP Settings** | **Edit PTP Settings** | **Refresh** | **Actions**
- Current date and time: Tuesday, March 04, 2025, 20:53:25 UTC
- NTP service status: Running
- NTP servers: 172.21.62.121, 172.21.62.122
- PTP client: Disabled
- PTP service status: Stopped
- > Network interface: --



## Configure host power policy on the first ESXi host

To configure host power policy for the first ESXi host, follow the steps below. This policy can be adjusted based on customer requirements.

1. From the ESXi Host Client Navigator, click Manage.
2. In the center pane, click Hardware > Power Management.
3. Click Change policy.
4. Select High performance and click OK.

## VMware vCenter 8.0U3 installation and VMware cluster configurations

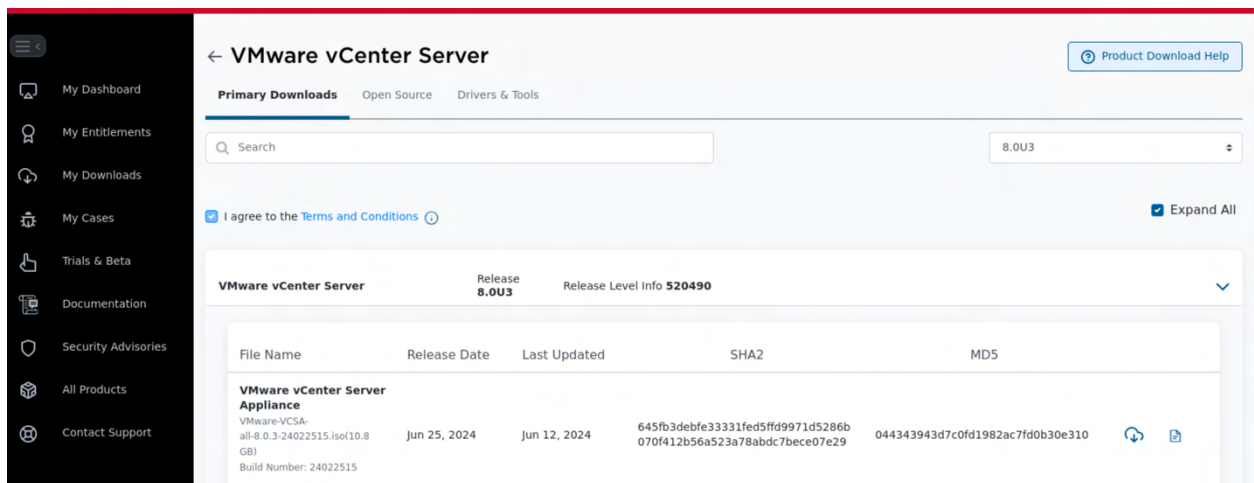
The procedures in the following sections provide instructions for installing VMware vCenter Server Appliance 8.0U3 in a FlexPod environment and finish the configurations of the ESXi hosts.

### Download VMware vCenter Server Appliance 8.0U3 ISO

1. Login to Broadcom support portal download site, <https://support.broadcom.com/group/ecx/downloads>.

**Note:** You will need a Broadcom user id and password.

2. Look for VMware vCenter Server Appliance.



Download the VMware vCenter Server Appliance 8.0U3 installation ISO.

**Note:** For this solution validation, the VMware-VCSA-all-8.0.3-244022515.iso image is used.

### Install the VMware vCenter Server Appliance

The VCSA deployment consists of 2 stages: installation and configuration.

1. Locate and copy the downloaded vCenter Server Appliance 8.0U3 ISO file to the desktop of the management workstation / jump host.

Mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).

In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.

Click Install to start the vCenter Server Appliance deployment wizard.

Click NEXT in the Introduction section.

Read and accept the End user license agreement and click NEXT.

In the vCenter Server deployment target window, enter the FQDN or IP address of the destination host, Username and Password. Click NEXT.

**Note:** Installation of vCenter on a separate existing management infrastructure vCenter is recommended. If a separate management infrastructure is not available, customers can choose the recently configured first ESXi host as an installation target. The recently configured ESXi host is used for this solution deployment.

Click YES to accept the certificate.

Enter the Appliance VM name and password details shown in the Set up vCenter Server VM section. Click NEXT.

In the Select deployment size section, select the Deployment size and Storage size. For example, select Small and Default. Click NEXT.

Select the datastore for VM storage. Click NEXT.

In the Configure network settings section, configure the following settings:

Select a Network: (for example, IB-MGMT Network)

**Note:** When the vCenter is running on the FlexPod, it is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and not be moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, trying to bring up vCenter on a different host than the one it was running on before the shutdown will cause problems with the network connectivity. With the vDS, for a virtual machine to move from one host to another, vCenter must be up and running to coordinate the move of the virtual ports on the vDS. If vCenter is down, the port move on the vDS cannot occur correctly. Moving vCenter to a different host on vSwitch0 does not require vCenter to already be up and running.

Provide IP version, IP assignment method, FQDN, IP address, subnet mask, default gateway, and DNS servers.

Click NEXT.

Review all values and click FINISH to complete the installation.

**Note:** The vCenter Server appliance installation will take several minutes to complete.

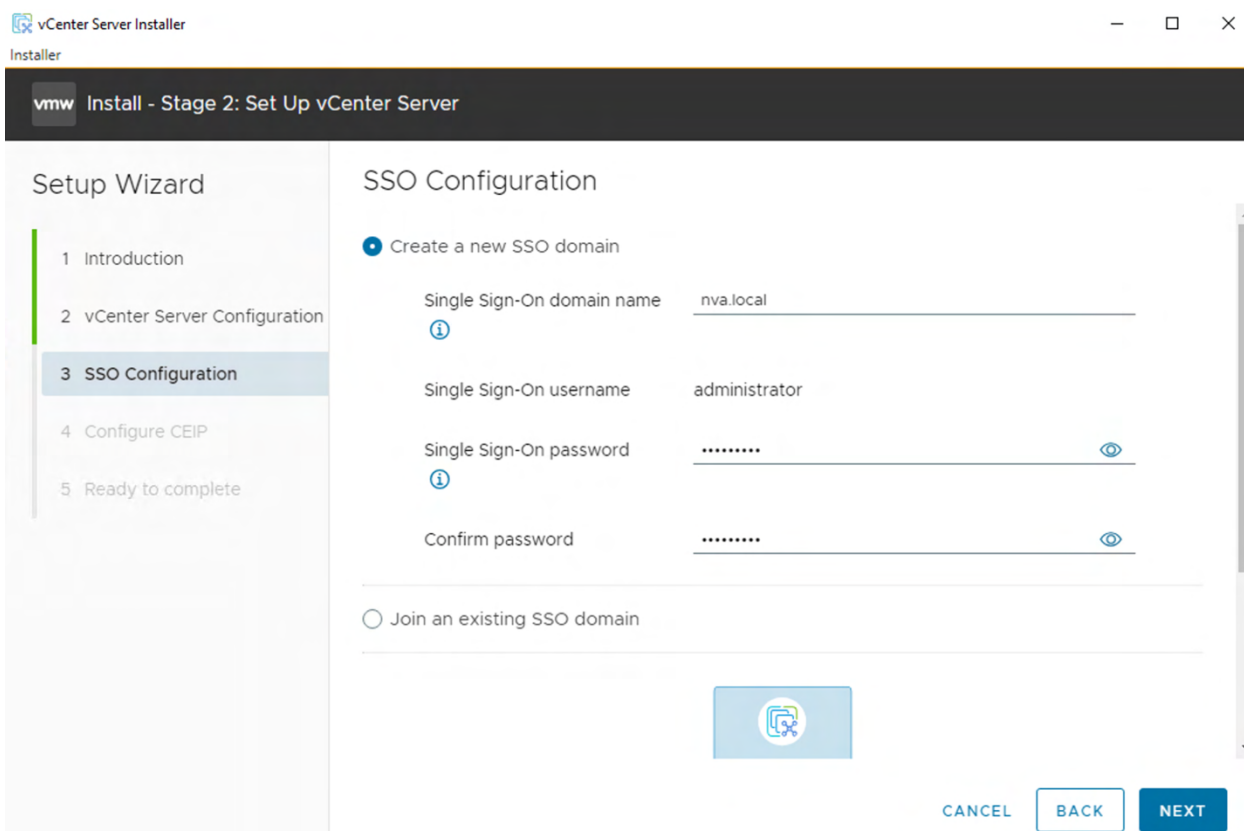
When Stage 1, Deploy vCenter Server, is completed, Click CONTINUE to proceed with stage 2.

Click NEXT.

In the vCenter Server Configuration window, configure Time synchronization mode, NTP servers, and SSH access.

Click NEXT.

Complete the SSO configuration according to your organization's security policies.



Click NEXT.

Decide whether to join VMware's Customer Experience Improvement Program (CEIP). Click NEXT.

Review the configuration and click FINISH.

Click OK.

**Note:** vCenter Server setup will take several minutes to complete Install – Stage 2.

Click CLOSE.

Eject or unmount the VCSA installer ISO.

## Verify and adjust vCenter appliance CPU configuration

When a vCenter deployment size of Small or larger was selected in the vCenter setup, it is possible that the VCSA's CPU setup does not match the Cisco UCS server CPU hardware configuration.

Cisco UCS X215c M8 servers are 2-socket servers. During this validation, the Small deployment size was selected and vCenter VM was setup for a 4-socket server with 1 core per socket. This setup can cause issues in the VMware ESXi cluster admission control. To verify and adjust CPU configuration, use the following steps.

1. Open a web browser on the management workstation and navigate to the vCenter or ESXi server where the vCenter appliance was deployed and login.

Under Virtual Machines, select the vCenter VM, right-click and click Edit settings.

In the Edit settings window, expand CPU and check the value of Sockets.

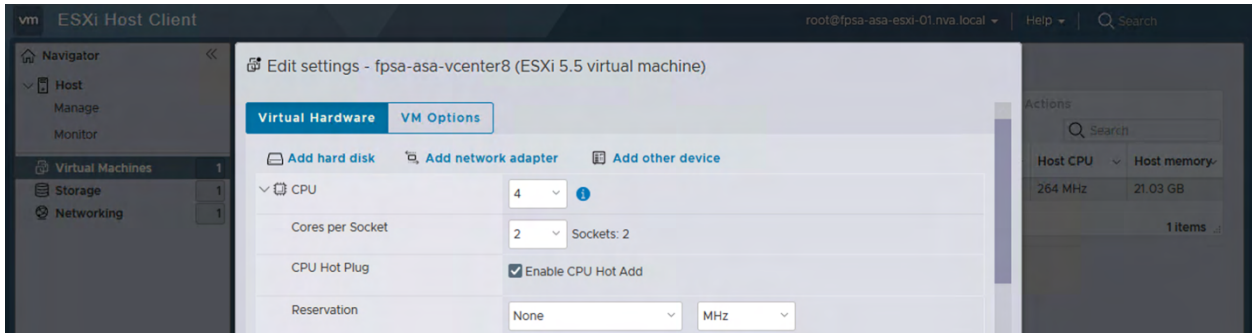
If the number of Sockets match the server configuration, click Cancel.

If the number of Sockets does not match the server configuration, it will need to be adjusted.

Right-click the vCenter VM and click Guest OS > Shut down. Click Yes on the confirmation.

When vCenter is shut down, right-click the vCenter VM and click Edit settings.

In the Edit settings window, expand CPU section and adjust the Cores per Socket value so the Sockets value matches the server configuration.



Click SAVE.

Right-click the vCenter VM and click Power > Power on. Wait approximately 10 minutes for vCenter to come up.

## vCenter Server Management interface configuration

1. Using a web browser, navigate to <https://<vcenter-ip-address>:5480>. Navigate and accept security warnings.

Login to the VMware vCenter Server Management interface as root with the root password set during the vCenter installation.

In the menu on the left, click Time.

Click EDIT to the right of Time zone.

Select the appropriate Time zone and click SAVE.

In the menu on the left select Administration.

According to your Security Policy, adjust the settings for the root user and password.

In the menu on the left click Update.

Follow the prompts to stage and install any available vCenter updates.

In the upper right-hand corner of the screen, click root > Logout to logout of the Appliance Management interface.

## Create data center in vCenter and enable HA and DRS

1. Using a web browser, navigate to <https://<vcenter-fqdn>> and navigate through security screens.
2. With VMware vCenter 7.0 and above, you must use the vCenter's FQDN.
3. Select LAUNCH VSPHERE CLIENT.
4. Login using the Single Sign-On username and password created during the vCenter installation. Dismiss the Licensing warning.
5. In the center pane, click ACTIONS > New Datacenter.
6. Type FlexPod-DC in the Datacenter name field and click OK.
7. Expand the vCenter on the left navigation pan.
8. Right-click the datacenter FlexPod-DC in the list in the left pane. Click New Cluster.
9. Provide a name for the cluster.
10. Turn on DRS and vSphere HA. Do not turn on vSAN.
11. Select Import image from a new host to set up the cluster's image.

## New Cluster

- 1 Basics
- 2 Image
- 3 Review

### Basics

Name	FlexPod
Location	FlexPod-DC
vSphere DRS	<input checked="" type="checkbox"/>
vSphere HA	<input checked="" type="checkbox"/>
vSAN	<input type="checkbox"/> Enable vSAN ESA

☒ Manage all hosts in the cluster with a single image

Choose how to set up the cluster's image

☐ Compose a new image  
☐ Import image from an existing host in the vCenter inventory  
☒ Import image from a new host

☐ Manage configuration at a cluster level

CANCEL NEXT

Click NEXT.

In the Image section, provide ESXi host name, root credential, and click FIND HOST.

## New Cluster

- 1 Basics
- 2 Image
- 3 Review

### Image

Enter the host details

fpsa-asa-esxi-01.nva.local root

FIND HOST

For the Security Alert, click YES to connect to the ESXi host.

Leave the box checked for Also move select host to cluster and click NEXT.

Review the information for cluster creation and click FINISH to proceed.

## New Cluster

- 1 Basics
- 2 Image
- 3 Review

### Review

Review the details before the cluster is created

Name	FlexPod
Location	FlexPod-DC
vSphere DRS	Enabled
vSphere HA	Enabled
vSAN	Disabled
vSAN ESA	Disabled
Selected image for cluster	Enabled
Cluster-level configuration management	Disabled

Image setup  
Imported from fpsa-asa-esxi-01.nva.local

ESXi 8.0 U3 - 24022510

Cisco Cisco-UCS-Addon-ESXi 4.3.5-a

0 additional components

SHOW COMPONENTS

CANCEL BACK FINISH

## Add remaining ESXi hosts to the cluster

1. Right-click the cluster and select Add Hosts.
2. In the Host name or IP address field, enter either the IP address or the FQDN of the additional ESXi host and click NEXT.
3. Provide username and password and click NEXT.
4. For the Security Alert, click YES to replace the host certificate with a new certificate signed by the VMware Certificate Server.
5. Review Host Summary and click YES.
6. Click NEXT on Assign license screen.
7. Select the appropriate Lockdown mode for your security requirement and click NEXT.
8. For the Host lifecycle, select Use the current image on host and click NEXT.
9. Use the default for the host's discovered virtual machine location and click NEXT.
10. Review host add summary information and click FINISH to proceed.

**Note:** The added ESXi host(s) will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed. The host will also have a TPM Encryption Recovery Key Backup Alarm which will be addressed later.

## Configure in-band management virtual switch

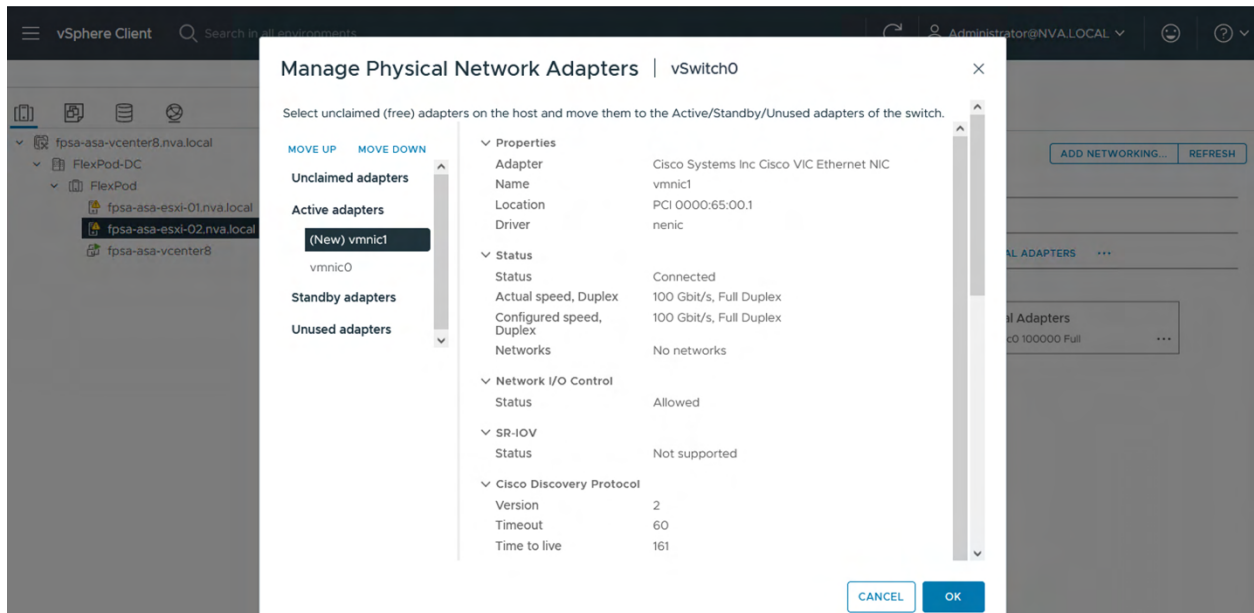
1. From the Host Client Navigator, select the newly added host under the FlexPod cluster.

In the center pane, select Configure tab.

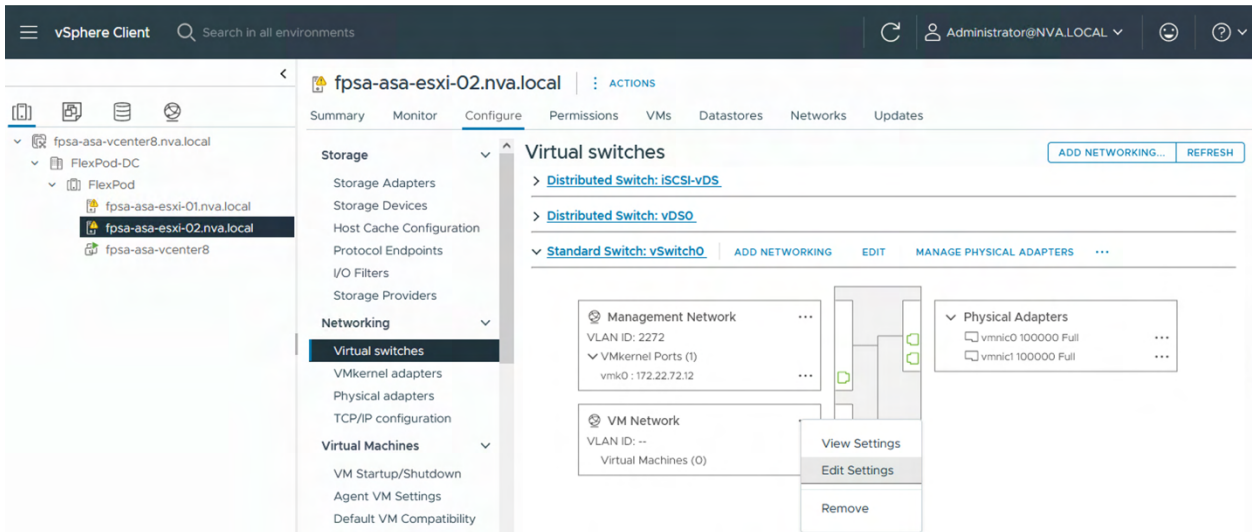
Select Virtual switches under Networking.

Expand Standard Switch: vSwitch0 and click on MANAGE PHYSICAL ADAPTERS.

Select vmnic1 under Unclaimed adapters and click MOVE DOWN to move it under Active adapters.



Click OK. The added vmnic1 is now also listed under the Physical Adapters.



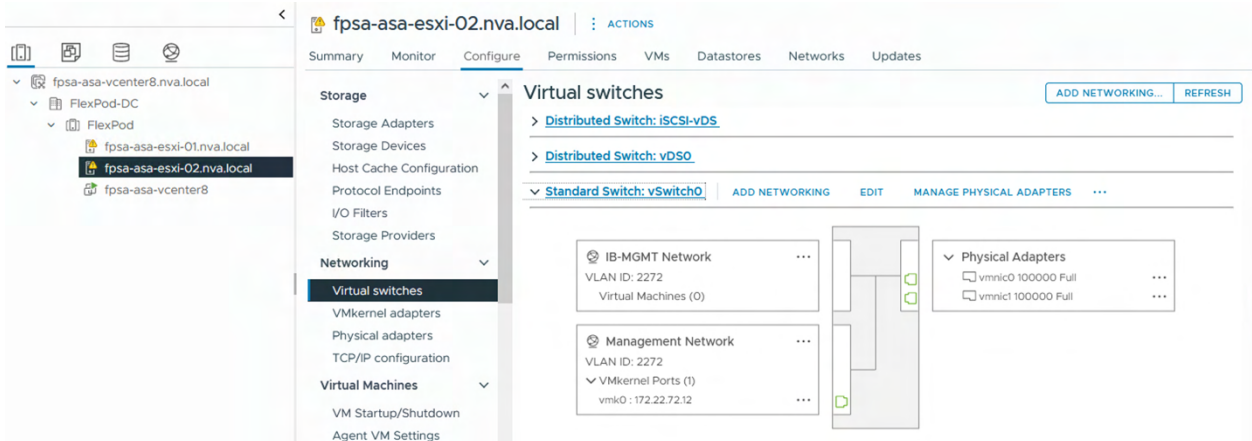
Click on the dots next to the VM Network port group and select Edit Settings.

Name the port group IB-MGMT Network and set the proper VLAN ID for the IB-MGMT VLAN.

#### VM Network - Edit Settings | fpga-asa-esxi-02.nva.local

Properties	
Network label	IB-MGMT Network
VLAN ID	2272
Security	
Traffic shaping	
Teaming and failover	

Click OK.



**Note:** The vSwitch0 is only used for in-band management with 1500 MTU.

## Create vSphere Distributed Switch (vDS) in vCenter

This section provides procedures for setting up VMware vDS in vCenter. Based on the VLAN configuration in Intersight, a vMotion, and a VM-Traffic port group will be added to the vDS in two port groups.

1. Login to vCenter.



Select Inventory under the top-level menu.

Click the fourth icon at the top to go to Networking.

Expand the vCenter and right-click the FlexPod-DC datacenter and click Distributed Switch > New Distributed Switch.

2. Give the Distributed Switch a descriptive name (for example, vDS0) and click NEXT.
3. Make sure version 8.0.3 - ESXi 8.0.3 and later is selected and click NEXT.
4. Select None for Network offloads compatibility. Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control enabled. Otherwise, disable Network I/O Control. Enter VM-Traffic for the Port group name. Click NEXT.
5. Skip DPU Failover Configuration and click NEXT.
6. Review the information and click FINISH to complete creating the vDS.

**New Distributed Switch**

1 Name and location

2 Select version

3 Configure settings

4 DPU Failover Configuration

**5 Ready to complete**

**Ready to complete** X

Review your settings selections before finishing the wizard.

Name	vDS0
Version	8.0.3
Network Offloads compatibility	None
Number of uplinks	2 (2 Active mode, 0 Standby mode)
Network I/O Control	Disabled
Default port group	VM-Traffic

✓ Suggested next actions

- New Distributed Port Group
- Add and Manage Hosts

① These actions will be available in the Actions menu of the new distributed switch.

7. Expand the FlexPod-DC datacenter and the newly created vDS.
8. Right-click the VM-Traffic port group and click Edit Settings.
9. Select VLAN.
10. Select for VLAN type and enter the VM-Traffic VLAN ID (for example, 2276). Click OK.
11. Right-click the vDS and click Settings > Edit Settings.
12. In the Edit Settings window, click the Advanced tab.
13. Change the MTU to 9000. Keep the Cisco Discovery Protocol as the Discovery Protocol type and optionally set the Operation to Both. Click OK.

## Distributed Switch - Edit Settings | vDS0 ×

General **Advanced** Uplinks

MTU (Bytes)

Multicast filtering mode

Discovery protocol

Type

Operation

14. To create the vMotion port group, right-click the vDS, select Distributed Port Group > New Distributed Port Group.
15. Enter vMotion as the name and click NEXT.
16. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion (for example, 2275), and click NEXT.
17. Review the information and click FINISH.

## Add ESXi hosts to vDS0

1. Right-click the vDS and click Add and Manage Hosts.
2. Make sure Add hosts is selected and click NEXT.
3. Click SELECT ALL to select all ESXi hosts. Click NEXT.
4. If all hosts had alignment in the ESXi console screen between vmnic numbers and vNIC numbers, leave Adapters on all hosts selected. To the right of vmnic2, use the pulldown to select Uplink 1. To the right of vmnic3, use the pulldown to select Uplink 2. Click NEXT. If the vmnic numbers and vNIC numbers did not align, select Adapters per host and select vDS uplinks individually on each host.

### vDS0 - Add and Manage Hosts

- 1 Select task
- 2 Select hosts
- 3 Manage physical adapters**
- 4 Manage VMkernel adapters
- 5 Migrate VM networking
- 6 Ready to complete

### Manage physical adapters ×

Add or remove physical network adapters to this distributed switch.

Adapters on all hosts **Adapters per host**

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

Physical network adapters	In use by switch	Assign uplink
>>  vmnic0	2 hosts / 2 switches	None
>>  vmnic1	2 hosts / 1 switch	None
>>  vmnic2	This switch	Uplink 1
>>  vmnic3	This switch	Uplink 2
>>  vmnic4	2 hosts / 2 switches	None
>>  vmnic5	2 hosts / 1 switch	None

5. Do not migrate any VMkernel ports and click NEXT.
6. Do not migrate any virtual machine networking ports. Click NEXT.
7. Click FINISH to complete adding the ESXi host to the vDS.

## Configure vmkernel port for vMotion

1. Select Hosts and Clusters and select the first ESXi host. In the center pane, select the Configure tab.
2. In the list under Networking, select VMkernel adapters.
3. Select ADD NETWORKING.
4. In the Add Networking window, ensure that VMkernel Network Adapter is selected and click NEXT.
5. Ensure that Select an existing network is selected.
6. Select vMotion from the list and click NEXT.
7. From the MTU drop-down list, select Custom and ensure the MTU is set to 9000.
8. From the TCP/IP stack drop-down list, select vMotion. Click NEXT.

**Add Networking**

1 Select connection type  
2 Select target device  
**3 Port properties**  
4 IPv4 settings  
5 Ready to complete

**Port properties**

Specify VMkernel port settings.

Network label: vMotion (vDS0)

MTU: Custom 9000

TCP/IP stack: vMotion

**Available services**

Enabled services

<input checked="" type="checkbox"/> vMotion	<input type="checkbox"/> vSphere Replication NFC	<input type="checkbox"/> NVMe over RDMA
<input type="checkbox"/> Provisioning	<input type="checkbox"/> vSAN	
<input type="checkbox"/> Fault Tolerance logging	<input type="checkbox"/> vSAN Witness	
<input type="checkbox"/> Management	<input type="checkbox"/> vSphere Backup NFC	
<input type="checkbox"/> vSphere Replication	<input type="checkbox"/> NVMe over TCP	

9. Select Use static IPv4 settings and fill in the IPv4 address and Subnet mask for the first ESXi host's vMotion IPv4 address and Subnet mask. Click NEXT.

Review the information and click FINISH to complete adding the vMotion VMkernel port to the first ESXi host.

Repeat the steps in this section for all other configured ESXi hosts.

## Configure iSCSI vDS in vCenter

1. Select Inventory under the top-level menu.

Click the fourth icon at the top, to go to Networking.

Expand the vCenter and right-click the FlexPod-DC datacenter and click Distributed Switch > New Distributed Switch.

Give the Distributed Switch a descriptive name (for example, iSCSI-vDS) and click NEXT.

Make sure version 8.0.3 - ESXi 8.0.3 and later is selected and click NEXT.

Change the Number of uplinks to 2. If VMware Network I/O Control is to be used for Quality of Service, leave Network I/O Control enabled. Otherwise, disable Network I/O Control. Uncheck Create a default port group. Click NEXT.

Review the information and click FINISH to complete creating the vDS.

Expand the FlexPod-DC datacenter and the newly created vDS. Click the newly created vDS.

Right-click the new vDS and click Settings > Edit Settings.

In the Edit Settings window, click the Advanced tab.

Change the MTU to 9000. Keep the Cisco Discovery Protocol as the Discovery Protocol type and optionally set the Operation to Both. Click OK.

## Distributed Switch - Edit Settings | iSCSI-vDS ×

General	Advanced	Uplinks
MTU (Bytes)	9000	
Multicast filtering mode	IGMP/MLD snooping	
Discovery protocol		
Type	Cisco Discovery Protocol	
Operation	Both	

To create the iSCSI-A port group, right-click the vDS, select Distributed Port Group > New Distributed Port Group.

Enter iSCSI-A as the name and click NEXT.

Leave the VLAN type set to None, check the box for Customize default policies configuration, and click NEXT.

Leave the Security options set to Reject and click NEXT.

Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.

Select Uplink 2 from the list of Active uplinks and click MOVE DOWN twice to place Uplink 2 in the list of Unused uplinks. This will pin all iSCSI-A traffic to UCS Fabric Interconnect A.

### New Distributed Port Group

- Name and location
- Configure settings
- Security
- Traffic shaping
- Teaming and failover**
- Monitoring
- Miscellaneous

### Teaming and failover

Notify switches Yes

Fallback Yes

Failover order ⓘ

MOVE UP MOVE DOWN

Active uplinks

Uplink 1

Standby uplinks

Unused uplinks

Uplink 2

Click NEXT.

Leave NetFlow disabled and click NEXT.

Leave Block all ports set as No and click NEXT.

Confirm the options and click FINISH to create the port group.

To create the iSCSI-B port group, right-click the vDS, select Distributed Port Group > New Distributed Port Group.

Enter iSCSI-B as the name and click NEXT.

Leave the VLAN type set to None, check the box for Customize default policies configuration, and click NEXT.

Leave the Security options set to Reject and click NEXT.

Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.

Select Uplink 1 from the list of Active uplinks and click MOVE DOWN three times to place Uplink 1 in the list of Unused uplinks. This will pin all iSCSI-B traffic to UCS Fabric Interconnect B.

The screenshot shows the 'New Distributed Port Group' configuration wizard, specifically the 'Teaming and failover' step. On the left, a sidebar lists the steps: 1 Name and location, 2 Configure settings, 3 Security, 4 Traffic shaping, 5 Teaming and failover (selected), 6 Monitoring, 7 Miscellaneous, and 8 Ready to complete. The main area is titled 'Teaming and failover' and contains several sections: 'Load balancing' with a dropdown set to 'Route based on originating virtual port'; 'Network failure detection' with a dropdown set to 'Link status only'; 'Notify switches' with a dropdown set to 'Yes'; and 'Failback' with a dropdown set to 'Yes'. Below these is the 'Failover order' section, which includes 'MOVE UP' and 'MOVE DOWN' buttons. Under 'Active uplinks', 'Uplink 2' is listed. Under 'Standby uplinks', there are no items. Under 'Unused uplinks', 'Uplink 1' is listed and highlighted with a dark background. At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'NEXT'.

Click NEXT.

Leave NetFlow disabled and click NEXT.

Leave Block all ports set as No and click NEXT.

Confirm the options and click FINISH to create the port group.

## Add hosts to iSCSI vDS

1. Right-click the iSCSI-vDS and click Add and Manage Hosts.

Make sure Add hosts is selected and click NEXT.

Select all configured iSCSI-booted hosts and click NEXT.

If all hosts had alignment in the ESXi console screen between vmnic numbers and vNIC numbers, leave Adapters on all hosts selected. To the right of vmnic5, use the pulldown to select Uplink 2. Click NEXT. If the vmnic numbers and vNIC numbers did not align, select Adapters per host and select vDS uplinks individually on each host.

## iSCSI-vDS - Add and Manage Hosts

- 1 Select task
- 2 Select hosts
- 3 **Manage physical adapters**
- 4 Manage VMkernel adapters
- 5 Migrate VM networking
- 6 Ready to complete

### Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Adapters on all hosts    Adapters per host

To associate a physical network adapter with an uplink, use "Assign uplink". This assignment would be applied to all the hosts that have the same physical network adapter available.

Physical network adapters	In use by switch	Assign uplink
>> vmnic0	2 hosts / 2 switches	None
>> vmnic1	2 hosts / 1 switch	None
>> vmnic2	2 hosts / 2 switches	None
>> vmnic3	2 hosts / 2 switches	None
>> vmnic4	2 hosts / 2 switches	None
>> vmnic5	This switch	Uplink 2

**Note:** It is important to assign the uplink as shown above. This allows the port groups to be pinned to the appropriate Cisco UCS Fabric and iSCSI network connectivity to be maintained.

To the right of vmk2, click ASSIGN PORT GROUP.

## iSCSI-vDS - Add and Manage Hosts

- 1 Select task
- 2 Select hosts
- 3 Manage physical adapters
- 4 **Manage VMkernel adapters**
- 5 Migrate VM networking
- 6 Ready to complete

### Manage VMkernel adapters

Manage and assign VMkernel network adapters to the distributed switch.

Adapters on all hosts    Adapters per host

To assign vmkernel network adapter to port group, click on the arrow or "Assign port group" button. This assignment would be applied to all the hosts that have the same vmkernel network adapter available.

Name	Related hosts	Assign port group
>> vmk0		
>> vmk1		
<< vmk2		
>> vmk3		

Name	NSX port group ID	Distributed switch	Actions
iSCSI-A	--	iSCSI-vDS	ASSIGN
iSCSI-B	--	iSCSI-vDS	UNASSIGN

To the right of iSCSI-B, click ASSIGN. Click NEXT.

Do not migrate any virtual machine networking ports. Click NEXT.

Click FINISH to complete adding the ESXi host(s) to the vDS.

## Configure iSCSI VMkernel port

1. Select Hosts and Clusters and select the first ESXi host added to the iSCSI-vDS.
2. In the center pane, select the Configure tab.

In the list under Networking, select Virtual switches.

Expand Standard Switch: vSwitch1. To the right of vSwitch1, select ... > Remove. Click YES to confirm the removal of vSwitch1.

Expand Standard Switch: iScsiBootvSwitch. To the right of iScsiBootvSwitch, select ... > Remove. Click YES to confirm the removal of iScsiBootvSwitch.

Expand Distributed Switch: iSCSI-vDS. To the right of Distributed Switch: iSCSI-vDS, click MANAGE PHYSICAL ADAPTERS.

Click the drop-down list to select Uplink 1 for vmnic4 and click OK.

In the center pane under Networking, select VMkernel adapters. Click ADD NETWORKING.

In the Add Networking window, ensure that VMkernel Network Adapter is selected and click NEXT.



Ensure that Select an existing network is selected, then select iSCSI-A, and click NEXT.

From the MTU drop-down list, select Custom and ensure the MTU is set to 9000. Click NEXT.

Select Use static IPv4 settings and fill in the IPv4 address and Subnet mask for the ESXi host's iSCSI-A IPv4 address and Subnet mask. Click NEXT.

Review the information and click FINISH to complete adding the iSCSI-A VMkernel port.

fpsa-asa-esxi-01.nva.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage > VMkernel adapters

Networking > ADD NETWORKING... REFRESH

	Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
⋮ >>	vmk0	Management Network	vSwitch0	172.22.72.11	Default	Management
⋮ >>	vmk1	iSCSI-A	iSCSI-vDS	172.22.73.11	Default	--
⋮ >>	vmk2	iSCSI-B	iSCSI-vDS	172.22.74.11	Default	--
⋮ >>	vmk3	vMotion	vDS0	172.22.75.11	vMotion	vMotion

Repeat the steps in this section for additional iSCSI booted hosts.

**Note:** The additional ESXi hosts may not have the Standard Switch: vSwitch1. When that is the case, skip the removal of that switch.

**Note:** When configuring additional ESXi hosts, be sure to also update or re-create VMkernel adapters for the iSCSI-B and vMotion port groups as needed. As an example, the following is the updated VMkernel adapters for the second host.

fpsa-asa-esxi-02.nva.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage > VMkernel adapters

Networking > ADD NETWORKING... REFRESH

	Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
⋮ >>	vmk0	Management Network	vSwitch0	172.22.72.12	Default	Management
⋮ >>	vmk1	iSCSI-A	iSCSI-vDS	172.22.73.12	Default	--
⋮ >>	vmk2	iSCSI-B	iSCSI-vDS	172.22.74.12	Default	--
⋮ >>	vmk3	vMotion	vDS0	172.22.75.12	vMotion	vMotion

## Configure iSCSI software adapter for additional hosts

1. Select Hosts and Clusters and select one of the remaining iSCSI SAN booted ESXi hosts.

In the center pane under Storage, click Storage Adapters.

Select the iSCSI Software Adapter and in the window below, click the Dynamic Discovery tab.

Click ADD.

Enter the IP address of the storage controller's LIF iscsi-lif-01a and click OK.

Repeat this process to add the IPs for iscsi-lif-02a, iscsi-lif-01b, and iscsi-lif-02b.

Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.

Under Static Discovery, four static targets should now be listed.



**Storage Adapters**

ADD SOFTWARE ADAPTER ▾ REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

Adapter	Model	Type	Status	Identifier	Targets	Devices
vmhba64	ISCSI Software Adapter	ISCSI	Online	iscsi_vmk(iqn.2010-11.com.flexpod:flexpod-asa-ucs-host:1)	4	2

Manage Columns Export ▾ 1 item

Properties Devices Paths **Dynamic Discovery** Static Discovery Network Port Binding Advanced Options

ADD REMOVE AUTHENTICATION ADVANCED...

2. Under Paths, four paths with the “Active (I/O)” Status should now be listed for each LUN.

**Storage Adapters**

ADD SOFTWARE ADAPTER ▾ REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba64	ISCSI Software Adapter	ISCSI	Online	iscsi_vmk(iqn.2010-11.com.flexpod:flexpod-asa-ucs-host:1)	4	2	8

Manage Columns Export ▾ 1 item

Properties Devices **Paths** Dynamic Discovery Static Discovery Network Port Binding Advanced Options

ENABLE DISABLE

Runtime Name	Target	LUN	Status
vmhba64:C0:T0:L0	iqn.1992-08.com.netapp.psn.57cd8838ea2911ef9.608d039eac6a795:vs.2:1.72.22.73.101:3260	0	Active (I/O)
vmhba64:C3:T0:L0	iqn.1992-08.com.netapp.psn.57cd8838ea2911ef9.608d039eac6a795:vs.2:1.72.22.74.101:3260	0	Active (I/O)
vmhba64:C2:T0:L0	iqn.1992-08.com.netapp.psn.57cd8838ea2911ef9.608d039eac6a795:vs.2:1.72.22.74.102:3260	0	Active (I/O)
vmhba64:C1:T0:L0	iqn.1992-08.com.netapp.psn.57cd8838ea2911ef9.608d039eac6a795:vs.2:1.72.22.73.102:3260	0	Active (I/O)
vmhba64:C0:T0:L1	iqn.1992-08.com.netapp.psn.57cd8838ea2911ef9.608d039eac6a795:vs.2:1.72.22.73.102:3260	1	Active (I/O)

Manage Columns Export ▾ 8 items

3. Repeat the steps in this section on additional iSCSI SAN booted ESXi hosts.

## Mount shared iSCSI datastore on additional hosts

1. Select Hosts and Clusters and select one of the remaining iSCSI SAN booted ESXi hosts.
2. Right-click on the host, select Storage > Rescan Storage, and click OK.
3. Select Inventory under the top-level menu.

Click the third icon at the top to go to Storage.

Expand FlexPod-DC and select the shared iSCSI datastore.

4. In the center pane, click the Hosts tab to confirm that the shared datastore is now also mounted on the select host.

## Create additional storage for VM swap and vSphere cluster services

Login to ONTAP cluster using command line or from ONTAP System Manager to create additional storage units to host the swap space for VMs and for vSphere Cluster Services (vCLS) VMs. To use the command line tool, follow the steps below.

1. Login to ONTAP.
2. Create a 200G swap LUN and a 100G vCLS LUN.

```
lun create fpsa_asa_vmware_swap_1 -size 200G -ostype vmware
lun create fpsa_asa_vmware_vcls_1 -size 100G -ostype vmware
```

3. Map the LUNs to the VMware cluster igroup.

```
lun map -path fpsa_asa_vmware_swap_1 -igroup FlexPod-ASA-esxi-cluster-iscsi -lun-id 2
lun map -path fpsa_asa_vmware_vcls_1 -igroup FlexPod-ASA-esxi-cluster-iscsi -lun-id 3
```

## Configure VMware cluster and ESXi host VM swap file location

To create the datastore for VM swap file usage, follow the steps below.

1. Navigate to Inventory > Storage.

Right-click on FlexPod cluster and select Storage > Rescan Storage.

Right-click on FlexPod cluster and select Storage > New Datastore.

Specify VMFS datastore type and click NEXT.

Provide a datastore name, select a host from the drop-down list to view available devices, select the LUN with 200G capacity, and click NEXT.

**New Datastore**

1 Type

**2 Name and device selection**

3 VMFS version

4 Partition configuration

5 Ready to complete

**Name and device selection**

Specify datastore name and a disk/LUN for provisioning the datastore.

Name vmware\_swap

The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

Select a host fpsa-asa-esxi-01.nva.local

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Clu VM Sup
<input type="radio"/>	NETAPP ISCSI Disk (na a.600a09803832344872 3f5877434a5253)	3	100.00 G B	Supported	Flash	512e	Nc
<input checked="" type="radio"/>	NETAPP ISCSI Disk (na a.600a09803832344867 24587338696c57)	2	200.00 G B	Supported	Flash	512e	Nc

Manage Columns Export

2 items

CANCEL BACK NEXT

Use the default VMFS 6 for the VMFS version and click NEXT.

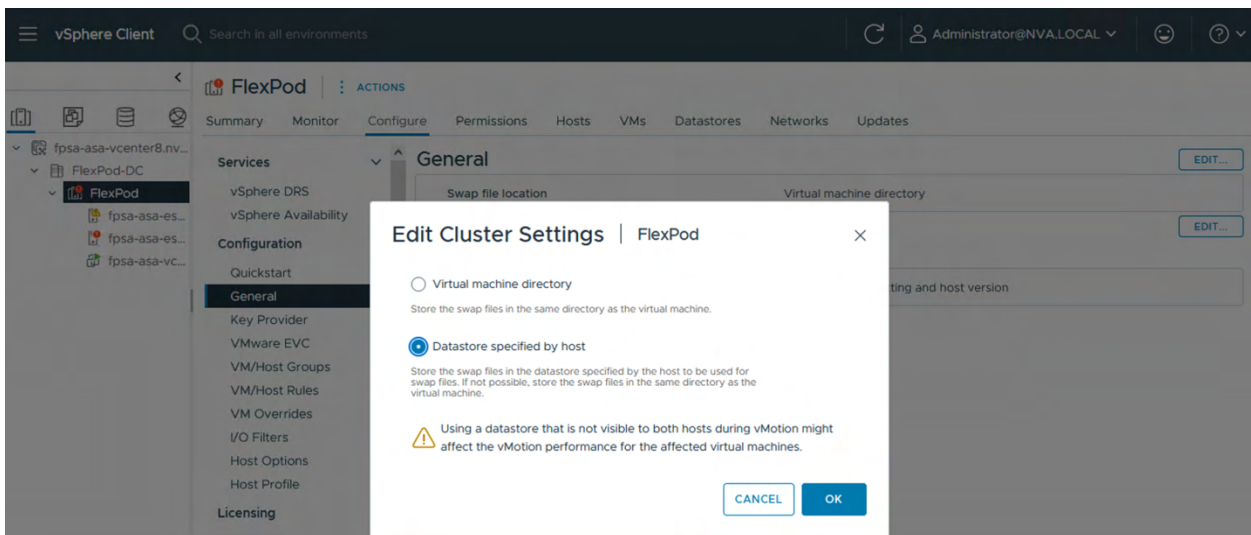
Review the disk layout and partition configuration and click NEXT.

Review the new datastore configuration information and click FINISH to create the new datastore.

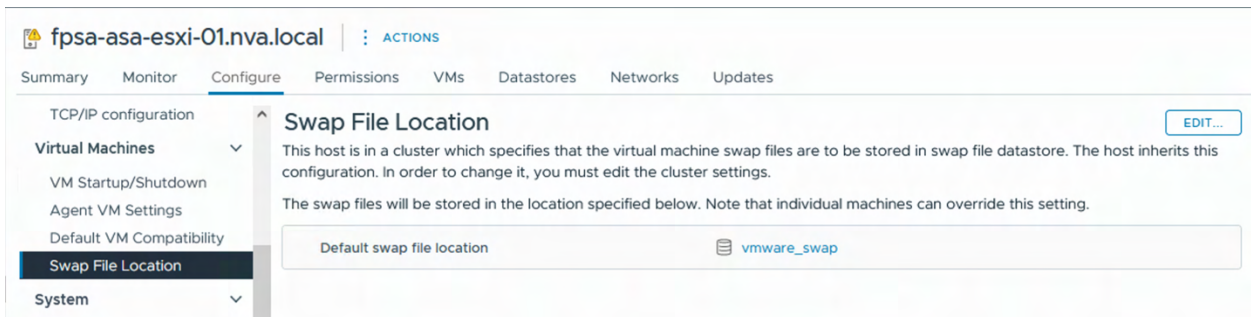
To configure vCenter to use the datastore specified by host for the swap file location, follow the steps below.

1. From the Navigator, select Inventory.

2. Right-click the FlexPod cluster and click Settings.
3. Click Configuration > General in the list located on the left and click EDIT to the right of General.
4. Select Datastore specified by host for the Swap file location and click OK.



5. From the Inventory list, select an ESXi host.
- Right-click the ESXi host and click Settings.
- In the center pane under Virtual Machines, click Swap File Location.
- On the right, click EDIT.
- Select vmware\_swap datastore and click OK.



Repeat steps 5-9 for each ESXi host to configure the Swap File Location.

## Configure datastore for vSphere Cluster Services (vCLS)

To create the datastore for vSphere Cluster Services (vCLS), follow the steps below.

1. Navigate to Inventory > Storage.
- Right-click on FlexPod cluster and select Storage > Rescan Storage.
- Right-click on FlexPod cluster and select Storage > New Datastore.
- Specify VMFS datastore type and click NEXT.
- Provide a datastore name, select a host from the drop-down list to view available devices, select the LUN with 100G capacity created for this purpose, and click NEXT.

New Datasore

1 Type

2 Name and device selection

3 VMFS version

4 Partition configuration

5 Ready to complete

Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

Name

vmware\_vcls

1

The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

Select a host

fpsa-asa-esxi-01.nva.local

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Clu VM Sup
	NETAPP iSCSI Disk (naa.600a098038323448723f5877434a5253)	3	100.00 GB	Supported	Flash	512e	Nd

Manage Columns

Export

1 item

CANCEL

BACK

NEXT

## Configure vCenter to utilize vCLS datastore

1. From the Navigator, select Inventory.
2. Right-click the FlexPod cluster and select Settings.

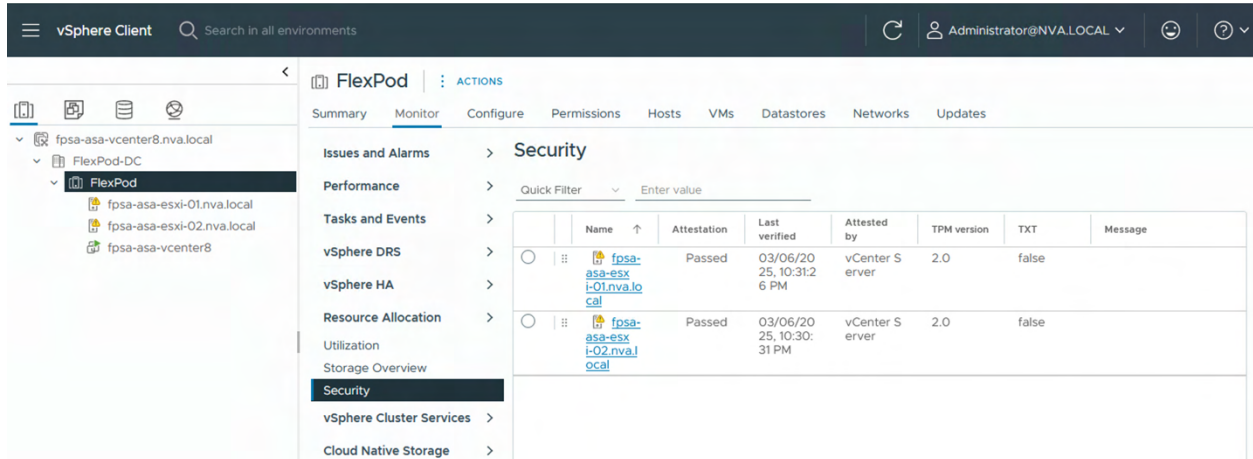
In the center of the window, click **ADD**. Select the `vmware_vcls` datastore and click **ADD**.

## vCenter Trusted Platform Module (TPM) attestation



For this validation, UEFI secure boot was enabled in the boot order policy. A server can boot with UEFI Secure Boot with or without a TPM 2.0 module. If it has a TPM, VMware vCenter can attest that the server booted with UEFI Secure Boot. To verify the vCenter TPM attestation, follow these steps:

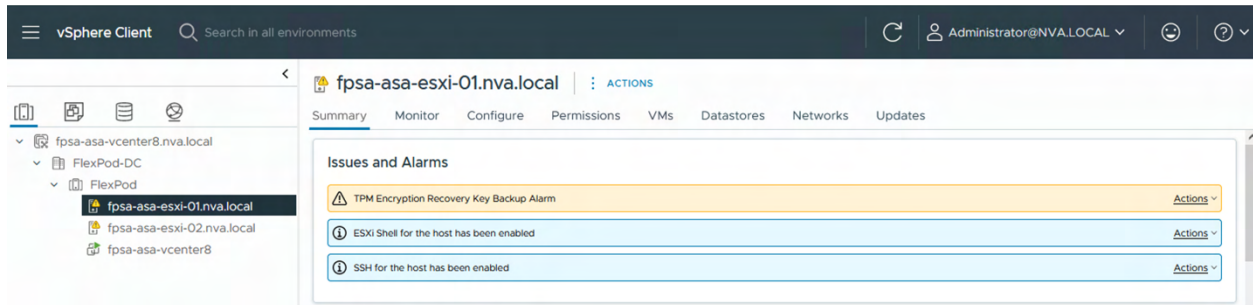
1. From the Navigator, select Inventory.
2. Right-click on the FlexPod cluster and select Settings.
3. In the center pane, click the Monitor tab.
4. Click Security. The TPM attestation status is shown for the hosts in the cluster.



**Note:** It may be necessary to disconnect and reconnect or reboot a host from vCenter to get it to pass attestation the first time

## Issues and alarms

After the initial ESXi 8.0U3 and vCenter 8.0U3 installation, the hosts may report several issues and alarms like in the screenshot below.



For the ESXi Shell for the host has been enabled and SSH for the host has been enabled messages, you can choose the Suppress Warning from the Actions drop down list to suppress them.

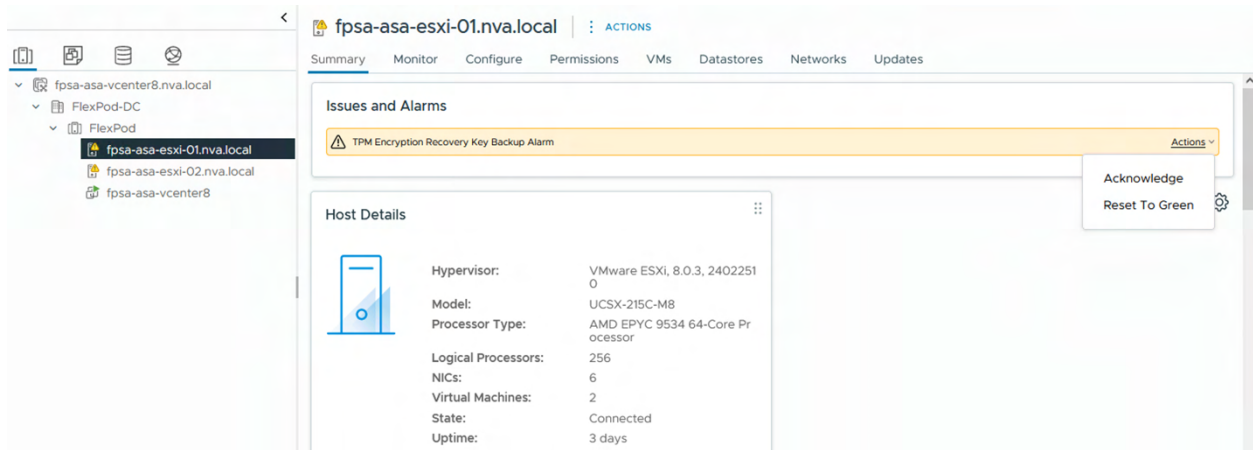
To address the TPM Encryption Recovery Key Backup Alarm, perform the following steps on each ESXi host to save the recovery key.

1. Login to ESXi host using SSH.
2. Gather the recovery key information with the esxcli command as show in the example below.

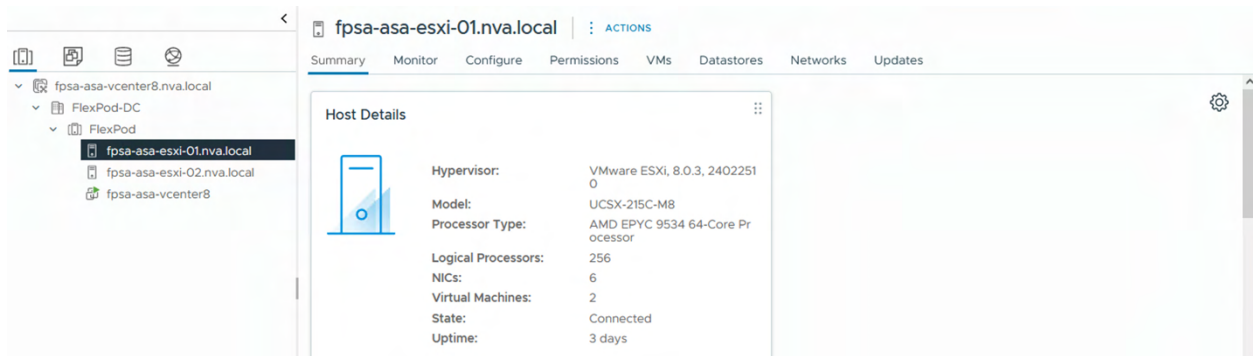
```
[root@fpsa-asa-esxi-01:~] esxcli system settings encryption recovery list
Recovery ID                                     Key
-----
{891883AB-5F2D-43EA-9C43-4CDDCCD5070E} 072028-485541-453114-715828-109630-338805-063025-060991-338038-391972-150028-612887-704056-181304-000828-041826
```

Store the keys from all hosts in a safe location.

Click on the Action drop-down list and select Reset to Green.



Repeat the steps above on all ESXi hosts to resolve the TPM warning so no issues and alarms show up on the hosts' Summary tab.



**Note:** Please refer to the solution validation section for an example of how to apply the encryption recovery key after a server profile is assigned to a different physical server to resolve the boot issue.

**Note:** Additional Issues and Alarms such as datastore usage warning might show up as the usage on the datastore increases due to virtual machine deployment. Please refer to Appendix A for the steps to take to increase datastore size.

## NetApp ONTAP tools for VMware vSphere deployment

ONTAP tools for VMware vSphere (OTV) is a set of tools for virtual machine lifecycle management. It is a collection of scalable, event-driven, microservices deployed as a virtual appliance and it integrates the VMware ecosystem and ONTAP storage to provide the following functionalities.

- Virtual machine functionality like basic protection and disaster recovery
- VASA Provider for VM granular management
- Storage policy-based management
- Storage Replication Adapter (SRA)

## ONTAP tools for VMware vSphere 10.4 pre-installation considerations

Before deploying, please review the [prerequisites](#) information to determine the specific deployment configuration that suits your needs. For example:

- Select thin-provisioned or thick-provisioned storage

- Select a deployment type: non-HA (small / medium) or HA (small / medium / large)

See the table in the prerequisite page linked above for the CPU, memory, storage requirements for the different deployment types and the number of ESXi hosts and vVols they support. For this validation, we are using Non-HA Small deployment type. Table 13 highlights the associated resources requirements and supported limits.

**Table 13 Non-HA small OTV deployment resource requirements and limits**

Resources	Requirement / Limit
CPU requirement	9
Memory requirement (GB)	18
Disk requirement (GB) thick provisioned	350
Number of vVols supported	~12,000
Number of ESXi hosts supported	32

**Note:** The effective number of vVols that can be supported also depends on the storage platform deployed. Please refer to [NetApp Hardware Universe](#) for the various storage platform support limits.

To facilitate proper operations and communications for the deployment, please refer to the network port requirements listed on the prerequisites page under the port requirements section. Please ensure that the necessary network configurations are in place to permit the needed traffic within your network for the associated services to function correctly.

The pre-deployment check section highlights the need to already have vCenter deployed and the login information for deploying the tool into vCenter. In addition, the three IP addresses for load-balancer, Kubernetes control plane and worker node in the Non-HA Small deployment type need to be already configured in DNS.

## Deploy ONTAP tools for VMware vSphere

1. Download the ONTAP tools for VMware vSphere 10 OVA from <https://support.netapp.com>.

The screenshot shows the NetApp Support website. The main navigation bar includes 'PRODUCTS', 'DOCS & KNOWLEDGE BASE', 'COMMUNITY', 'DOWNLOADS', and 'CASES'. The breadcrumb trail indicates the path: 'Products > All Products > ONTAP tools for VMware vSphere 10 (Downloads)'. The page title is 'ONTAP tools for VMware vSphere 10'. Below the title, there are tabs for 'Documentation', 'Community', and 'Downloads', with 'Downloads' being the active tab. A search bar is located on the right side of the page. The main content area features a 'Download Latest Release [10.4]' button and a dropdown menu to 'Select listed version or enter another'. A 'Go' button is also present. A sidebar on the right contains links for 'Elio', 'Chat', 'Case', and 'Phone'. A description of the ONTAP tools for VMware vSphere 10 is provided, stating it is based on next generation architecture which supports native high availability and scalability of the VASA Provider (for iSCSI and NFS vVols). Simplifies management of multiple VMware vCenter servers and ONTAP clusters.

**Note:** You can use md5sum or sha256sum tool to generate a checksum against the downloaded image and compare it with the checksum information listed on the download page to ensure image integrity.

Login to vCenter.

Navigate to Inventory.



Select Actions for the FlexPod-DC datacenter and select Deploy OVF Template.

Browse to the downloaded ONTAP tools OVA file, select the file, and click NEXT.

Enter the VM name and select a datacenter or folder to deploy the VM and click NEXT.

Select a compute resource and click NEXT.

Verify the template details, click Ignore for the certificate is not trusted warning, and click NEXT.

Check the box to accept the license agreements and click NEXT.

Select a datastore for VM storage, select Thin Provision option for the virtual disk format, and click NEXT.

On the Select networks screen, browse to select a destination network, such as IB-MGMT Network, click OK, and click NEXT.

On the Customize template screen, enter the required information for System Configuration, Deployment Configuration, and Node Configuration network details and click NEXT.

On the Customize hardware screen, customize virtual hardware for the VM if needed and click NEXT.

Review the configuration details entered and click FINISH to proceed with the deployment of NetApp ONTAP tools VM.

The screenshot shows the 'Deploy OVF Template' wizard in a two-pane layout. The left pane, titled 'Deploy OVF Template', contains a vertical list of steps: 1. Select an OVF template, 2. Select a name and folder, 3. Select a compute resource, 4. Review details, 5. License agreements, 6. Select storage, 7. Select networks, 8. Customize template, and 9. Ready to complete. Step 9 is highlighted with a dark background and white text. The right pane, titled 'Ready to complete', contains a summary of the selections. It starts with 'Review your selections before finishing the wizard' and lists several expandable sections: 'Select a name and folder' (Name: fpsa-asa-otv10, Template name: netapp-ontap-tools-for-vmware-vsphere-10.4-1744554600, Folder: FlexPod-DC), 'Select a compute resource' (Resource: FlexPod), 'Review details' (Download size: 7.1 GB), 'Select storage' (Size on disk: 11.1 GB, Storage mapping: 1, All disks: Datastore: iscsi\_datastore\_1; Format: Thin provision), 'Select networks' (Network mapping: nat, IB-MGMT Network; IP allocation settings: IP protocol: IPv4, IP allocation: Static - Manual), and 'Customize template' (Properties: Administrator username\* = admin, NTP Servers = 10.61.177.2, ONTAP tools IP address\* = 172.22.72.62, ONTAP tools virtual IP address\* = 172.22.72.63, HostName\* = fpsa-asa-otv10-node, Primary DNS\* = 10.61.177.2, Secondary DNS\* = 10.61.177.4, Search domains\* = nva.local). At the bottom right of the right pane are three buttons: 'CANCEL', 'BACK', and 'FINISH'.

Ready to complete	
Review your selections before finishing the wizard	
▼ Select a name and folder	
Name	fpsa-asa-otv10
Template name	netapp-ontap-tools-for-vmware-vsphere-10.4-1744554600
Folder	FlexPod-DC
▼ Select a compute resource	
Resource	FlexPod
▼ Review details	
Download size	7.1 GB
▼ Select storage	
Size on disk	11.1 GB
Storage mapping	1
All disks	Datastore: iscsi_datastore_1; Format: Thin provision
▼ Select networks	
Network mapping	1
nat	IB-MGMT Network
IP allocation settings	
IP protocol	IPv4
IP allocation	Static - Manual
▼ Customize template	
Properties	Administrator username* = admin NTP Servers = 10.61.177.2 ONTAP tools IP address* = 172.22.72.62 ONTAP tools virtual IP address* = 172.22.72.63 HostName* = fpsa-asa-otv10-node Primary DNS* = 10.61.177.2 Secondary DNS* = 10.61.177.4 Search domains* = nva.local

Open the VM console and wait for the deployment to complete and the node login prompt to appear.

## Add vCenter Server instance in ONTAP tool manager

By integrating with vCenter, ONTAP tools enables you to perform storage tasks like provisioning, snapshots, and data protection directly from the vSphere client, reducing the needs to switch to separate storage management consoles.

1. Use a web browser and open the ONTAP tools Manger URL, <https://<ONTAPtoolsIP>:8443/virtualization/ui/>, as shown in the node's console.

```

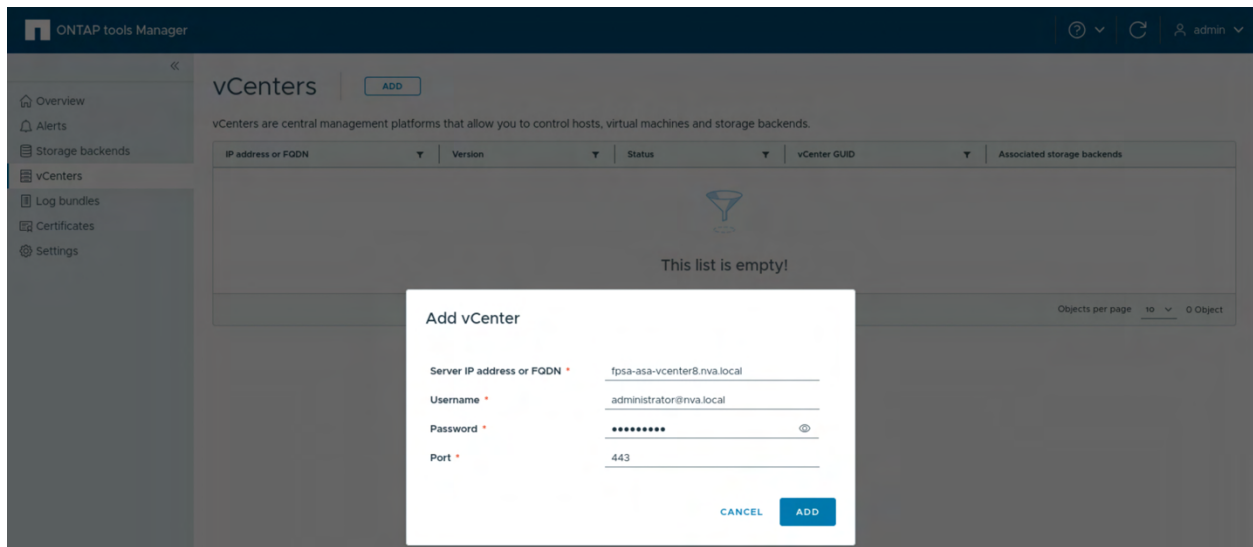
ONTAP tools for VMware vSphere
Node IP address:
IPv4 address: 172.22.72.64
APPLICATION STATUS:
ONTAP tools for VMware vSphere is running.
ManagerURL: https://172.22.72.62:8443/virtualization/ui
fpsa-asa-otv10-node login: _

```

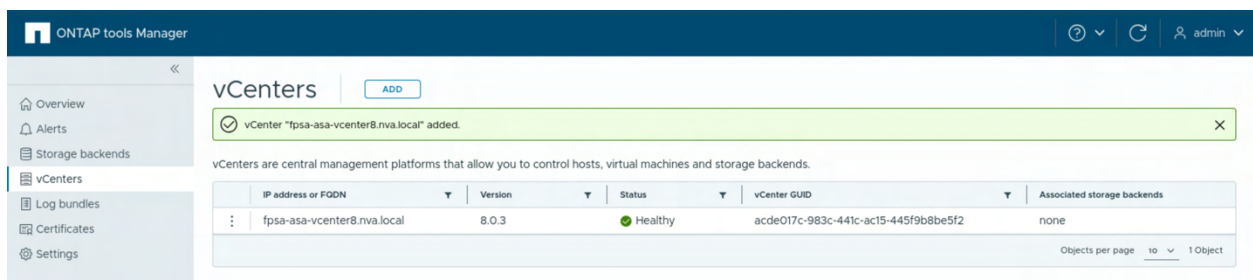
Login with the administrator credentials you provided during the deployment configuration.

On the Getting Started screen, select go to vCenters for the vCenter menu, and click ADD on the top of center pane.

In the Add vCenter dialog, provide vCenter IP/name and credentials and click ADD.



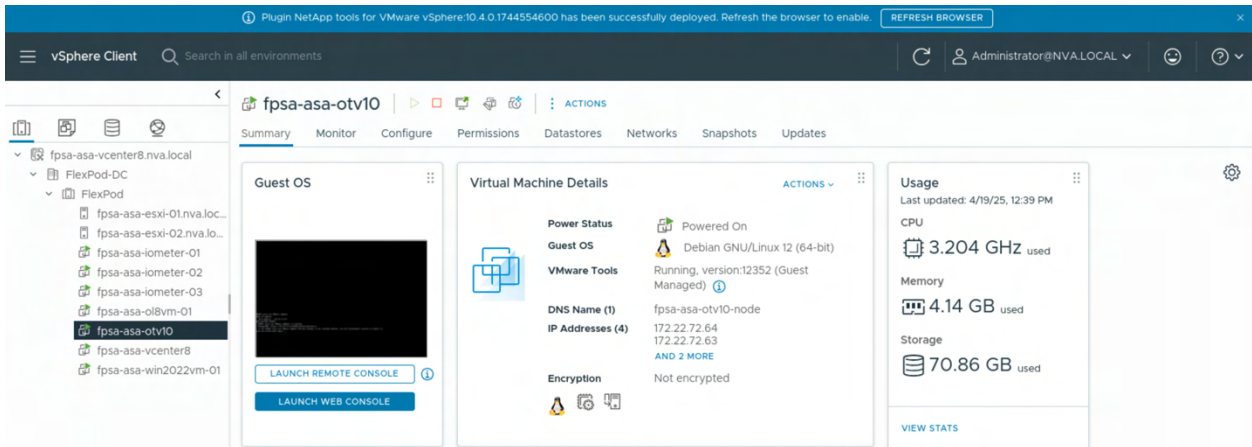
The added vCenter shows up in the vCenters screen.



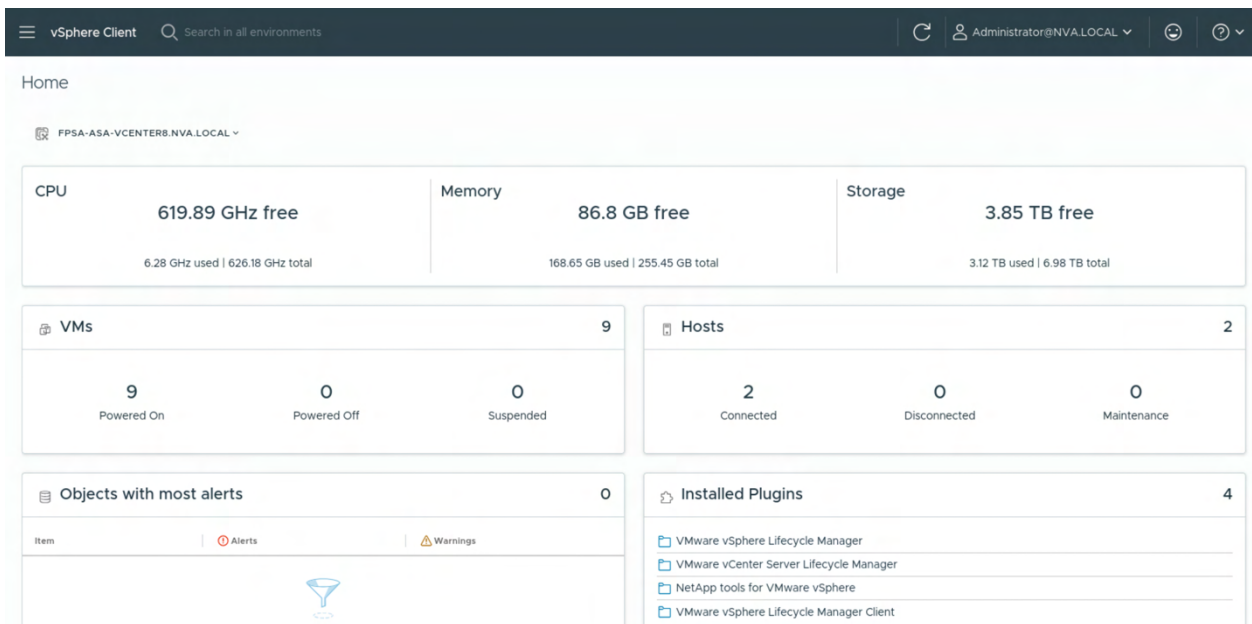
**Note:** Adding a vCenter Server instance to ONTAP tools automatically triggers the following actions:

- The NetApp ONTAP tools plug-in is registered as a remote plug-in in vCenter.
- Custom privileges for the plug-ins and APIs are applied to the vCenter Server instance.
- Custom roles are created to manage the users.
- The plug-in appears as a shortcut on the vSphere user interface.

On your vCenter GUI, there should be a banner message asking you to refresh the browser to enable the NetApp tools for VMware vSphere Plugin after it has been successfully deployed. Click REFRESH BROWSER.



The Home screen of the vSphere Client Installed Plugins section will include NetApp tools for VMware vSphere, and the menu will include a selection for NetApp ONTAP tools.



## Add a storage backend from ONTAP tools plugin

Adding a storage backend enables you to onboard an ONTAP cluster. To add a storage backend to a locally scoped cluster, add your ONTAP systems directly using the ONTAP tools plug-in in vCenter.

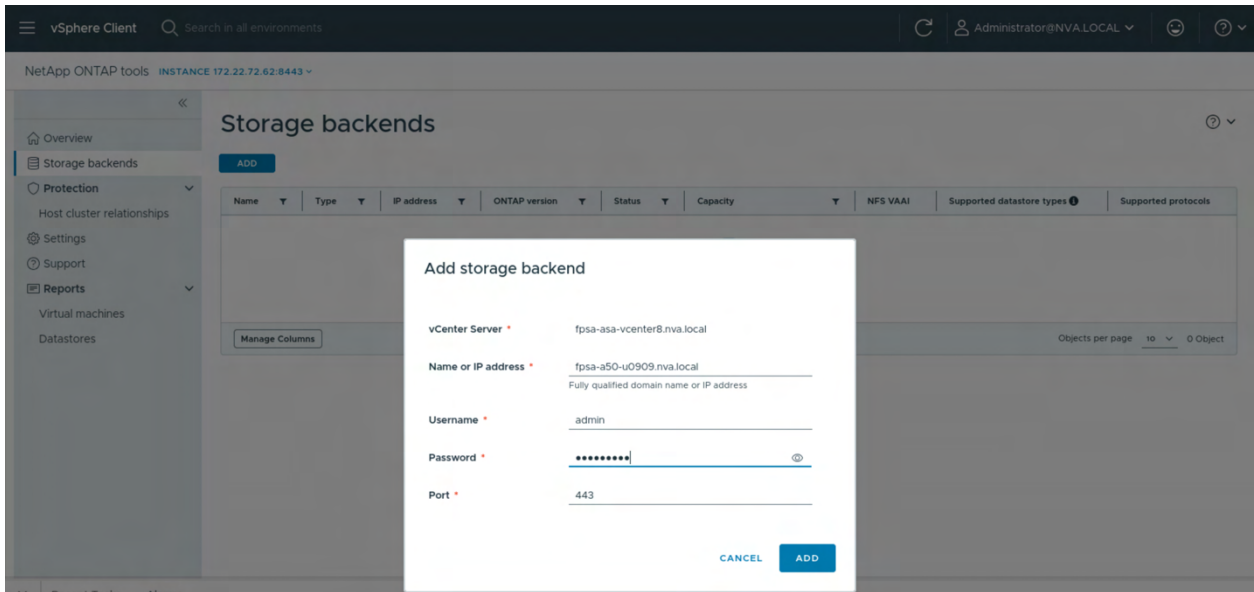
**Note:** If you will be using a multi-tenant storage deployment for different vCenter to utilize different Storage Virtual Machine (SVM) from the same storage, please refer to [ONTAP tools for VMware vSphere documentation](#) for additional steps.

1. Login to vCenter.

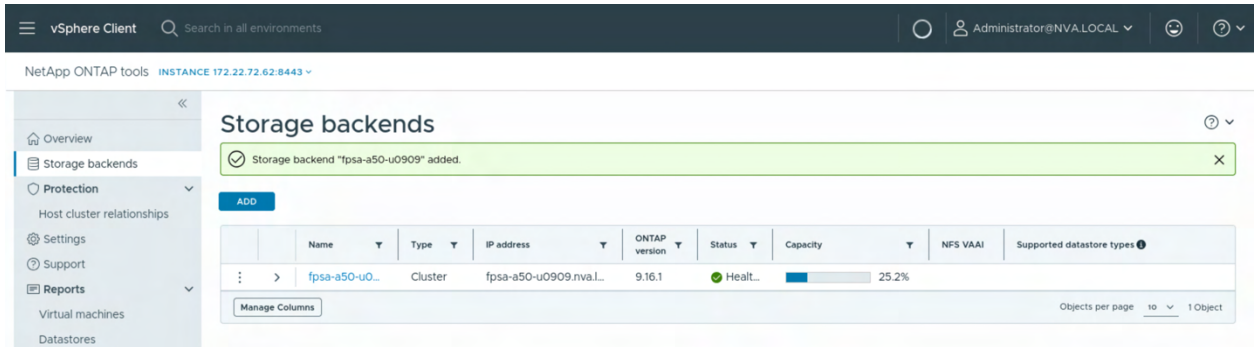
In the shortcuts page, select **NetApp ONTAP tools** under the Plugins section.

**Select Storage backends** from the left menu and click Add in the center pane.

In the Add storage backend dialog, provide the name or IP address of the storage backend, username, password, port, and click ADD.



The newly added storage backend shows up in the list.

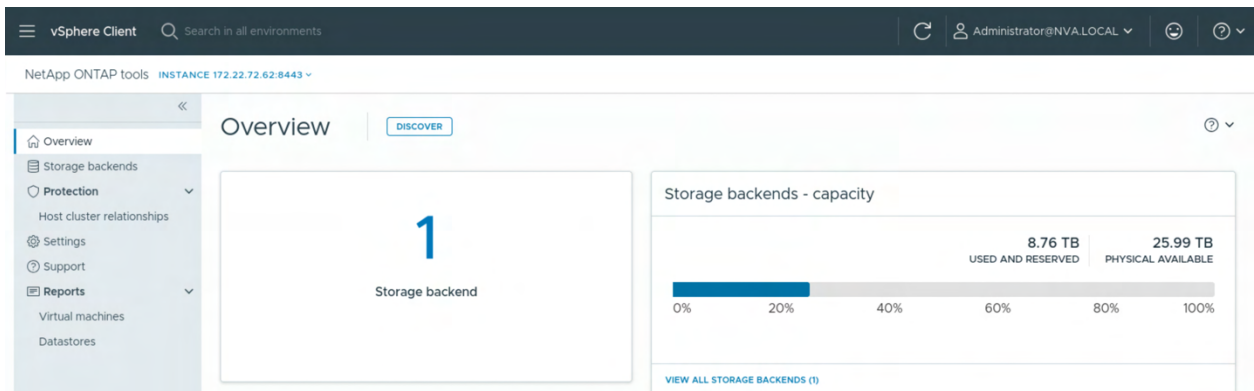


## Initiate a discovery after adding a storage backend

After adding a storage backend, initiate a discovery operation to discover any newly added/updated storage backend, hosts, datastores, and protection status/relationship on the vCenter.

1. In the shortcuts page, select **NetApp ONTAP tools** under the Plugins section.

Select **Overview** from menu and click **DISCOVER** in the center pane.



Click **START** to start the discovery.

Monitor the discovery status from the Recent Tasks in vCenter.

Recent Tasks		Alarms						
Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
ONTAP tools Discover cluster protection	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>	Queued	Discover cluster protection initiated with job id 777	NVA.LOCAL\Administrator	3 ms	04/19/2025, 12:42:28 PM		<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>
ONTAP tools Datastore Discovery	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>	Queued	Datastore Discovery initiated with job id 776	NVA.LOCAL\Administrator	2 ms	04/19/2025, 12:42:28 PM		<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>
ONTAP tools Discover hosts	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>	Completed	Discover hosts initiated with job id 775	NVA.LOCAL\Administrator	3 ms	04/19/2025, 12:42:28 PM	04/19/2025, 12:42:55 PM	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>
ONTAP tools Discover storage backend	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>	25%	Discover storage backend initiated with job id 770	NVA.LOCAL\Administrator	3 ms	04/19/2025, 12:42:28 PM		<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>
ONTAP tools Discovery	<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>	25%	Discovery initiated with job id 769	NVA.LOCAL\Administrator	9 ms	04/19/2025, 12:42:27 PM		<a href="#">fpga-asa-vcenter</a> <a href="#">8.nva.local</a>

**Note:** The discovery tasks will take several minutes time to complete. The discovered information for the storage backend, virtual machines, and datastores are shown in the Overview page as well as under the respective menu pages for Storage backends, Virtual machines, and Datastores.

**Note:** During the discovery process, ONTAP tools performs some igroup related operations on storage to help facilitate datastore creation integration from ONTAP tools in the future.

- Here is the igroup information before the ONTAP tools discovery process was initiated. The three individual host igroups were created for SAN boot LUN mapping usage. The cluster igroup, which includes all hosts was created for mapping shared datastore LUNs.

```
fpga-a50-u0909:>> igroup show -igroup FlexPod-ASA-esxi*
Vserver      Igroup      Protocol OS Type  Initiators
-----
svml         FlexPod-ASA-esxi-01-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
svml         FlexPod-ASA-esxi-02-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
svml         FlexPod-ASA-esxi-03-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3
svml         FlexPod-ASA-esxi-cluster-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3
4 entries were displayed.
```

- ONTAP tools renames existing ESXi cluster igroup which includes all the ESXi hosts in the cluster by adding a prefix of otv\_ to the igroup name. As a result, the original igroup name FlexPod-ASA-esxi-cluster-iscsi becomes otv\_FlexPod-ASA-esxi-cluster-iscsi.
- Then, ONTAP tools creates a parent igroup of the existing ESXi cluster igroup using the original name of the cluster igroup but with mixed protocol igroup type. As shown in the example below, ONTAP tools created FlexPod-ASA-esxi-cluster-iscsi igroup with mixed protocol type and it has the otv\_FlexPod-ASA-esxi-cluster-iscsi igroup as its child igroup.
- When creating a new datastore from ONTAP tools directly in the future, you can select FlexPod-ASA-esxi-cluster-iscsi as the custom igroup for the datastore mapping.

```
fpga-a50-u0909:>> igroup show -igroup *FlexPod-ASA-esxi*
Vserver      Igroup      Protocol OS Type  Initiators
-----
svml         FlexPod-ASA-esxi-01-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
svml         FlexPod-ASA-esxi-02-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
svml         FlexPod-ASA-esxi-03-boot-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3
svml         FlexPod-ASA-esxi-cluster-iscsi  mixed  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3
svml         otv_FlexPod-ASA-esxi-cluster-iscsi  iscsi  vmware  iqn.2010-11.com.flexpod:flexpod-asa-ucshost:1
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:2
                                                    iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3
5 entries were displayed.
```

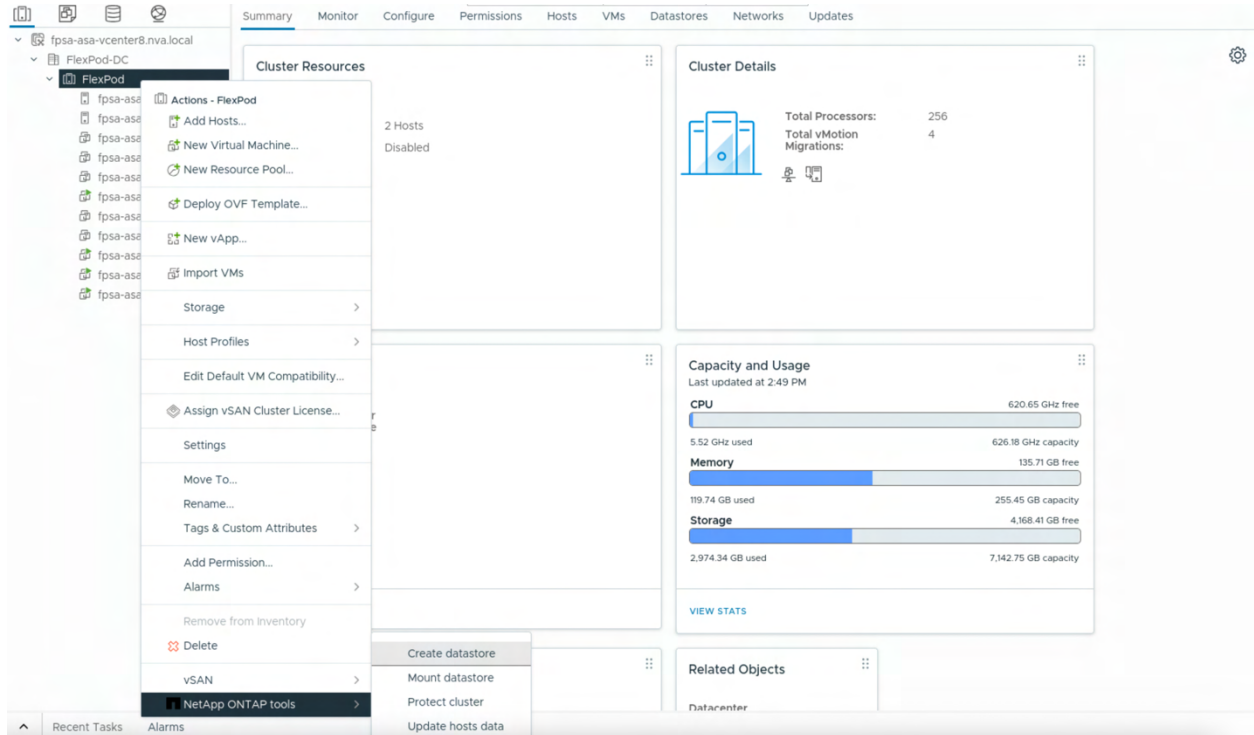


## Using ONTAP tools to provision a datastore

By using ONTAP tools which integrates VMware vSphere with ONTAP storage, you can provision a datastore directly from vCenter. ONTAP tools plugin will coordinate the tasks of creating a new storage unit, which can be a LUN or a NVMe namespace, invoking device rescan on the hosts, creating a datastore based on the newly created storage unit and making the datastore available to all the hosts in the cluster. This integration greatly simplifies datastore creation process for a VMware environment.

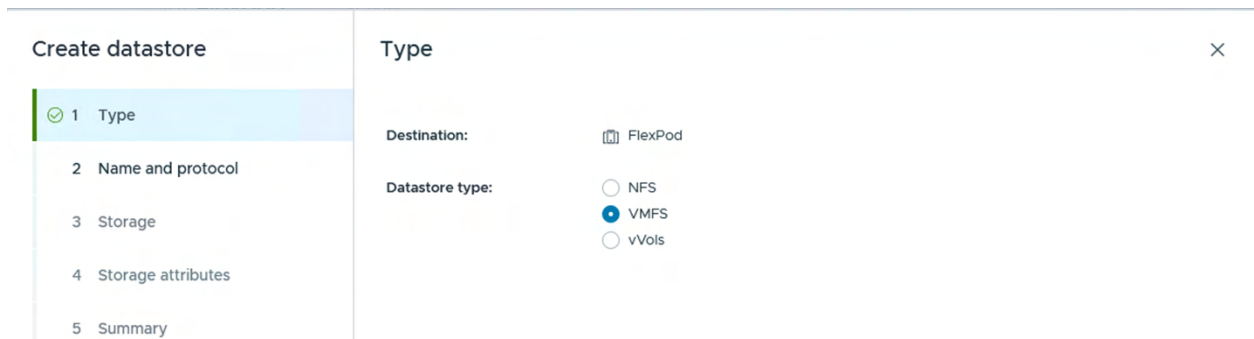
### 1. Login to vCenter.

Right-click on the FlexPod cluster in the vSphere Client inventory view and scroll down to the bottom of the menu to see NetApp ONTAP tools.



Click NetApp ONTAP tools and then select Create datastore.

In the Create datastore dialog, select VMFS Datastore type. Click NEXT.



Provide datastore name, size, and select protocol. Click NEXT.

Create datastore

1 Type

2 Name and protocol

3 Storage

4 Storage attributes

5 Summary

Name and protocol

Datastore name \*

otv\_vmfs\_iscsi\_1

Size \*

500

GB

Minimum supported size is 2 GB.

Protocol \*

ISCSI

Advanced options

Datastore cluster

**Note:** Under Advanced options, you can select a Datastore cluster if you have it.

The ASA cluster should already be selected if it is the only backend storage.

Expand Advanced options and select the initiator group (igroup) from the drop-down list. Click NEXT.

Create datastore

1 Type

2 Name and protocol

3 Storage

4 Storage attributes

5 Summary

Storage

Select a storage VM for the datastore.

Storage VM name	Tier	Platform	QoS configured
fpsa-a50-u0909 / svm1	Performance	ASA r2	No

Manage Columns

1 Storage VMs

Advanced options

Custom initiator group name

FlexPod-ASA-esxi-cluster-iscsi

Choose an existing initiator group or give a new name to the default initiator group.

You can optionally Enable QoS and specify minimum and maximum IOPs. Click NEXT.

Create datastore

1 Type

2 Name and protocol

3 Storage

4 Storage attributes

5 Summary

Storage attributes

Specify the storage details for provisioning the datastore.

Storage unit:

A new storage unit will be automatically created.  
LUNs are used for iSCSI and FC protocol and namespaces are used for NVMe/TCP and NVMe/FC protocol.

Space reserve:

Thin

Enable QoS

☒

Review Summary information and click FINISH to create the datastore.

142

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture



Create datastore

1 Type

2 Name and protocol

3 Storage

4 Storage attributes

5 Summary

Summary

A new datastore will be created with these settings.

Type

Destination

FlexPod

Datastore type

VMFS

Name and protocol

Datastore name

otv\_vmfs\_iscsi\_1

Size

500 GB

Protocol

iSCSI

Storage

Storage VM

fpsa-a50-u0909 / svm1

Custom initiator group name

FlexPod-ASA-esxi-cluster-iscsi

CANCEL

BACK

FINISH

Click OK on the pop-up dialog and track the datastore creation under the Recent Tasks in vCenter.

Task Name	Target	Status	Details	Initiator	Queued For	Start Time	Completion Time	Server
Process VMFS datastore updates	fpsa-asa-esxi-02.nva.local	Completed		System	1 ms	04/19/2025, 11:40:26 PM	04/19/2025, 11:40:26 PM	fpsa-asa-vcenter8.nva.local
Create VMFS datastore	fpsa-asa-esxi-01.nva.local	Completed		NVALocal\Administrator	1 ms	04/19/2025, 11:40:26 PM	04/19/2025, 11:40:26 PM	fpsa-asa-vcenter8.nva.local
Refresh host storage system	fpsa-asa-esxi-01.nva.local	Completed		NVALocal\Administrator	1 ms	04/19/2025, 11:40:25 PM	04/19/2025, 11:40:25 PM	fpsa-asa-vcenter8.nva.local
Rescan HBA	fpsa-asa-esxi-01.nva.local	Completed		NVALocal\Administrator	2 ms	04/19/2025, 11:40:25 PM	04/19/2025, 11:40:25 PM	fpsa-asa-vcenter8.nva.local
Refresh host storage system	fpsa-asa-esxi-02.nva.local	Completed		NVALocal\Administrator	2 ms	04/19/2025, 11:40:22 PM	04/19/2025, 11:40:23 PM	fpsa-asa-vcenter8.nva.local
Refresh host storage system	fpsa-asa-esxi-01.nva.local	Completed		NVALocal\Administrator	2 ms	04/19/2025, 11:40:22 PM	04/19/2025, 11:40:23 PM	fpsa-asa-vcenter8.nva.local
Rescan HBA	fpsa-asa-esxi-01.nva.local	Completed		NVALocal\Administrator	2 ms	04/19/2025, 11:40:22 PM	04/19/2025, 11:40:22 PM	fpsa-asa-vcenter8.nva.local
Rescan HBA	fpsa-asa-esxi-02.nva.local	Completed		NVALocal\Administrator	3 ms	04/19/2025, 11:40:22 PM	04/19/2025, 11:40:22 PM	fpsa-asa-vcenter8.nva.local
ONTAP tools Create Datastore	fpsa-asa-vcenter8.nva.local	Completed	Create datastore initiated with job id 1325	NVALocal\Administrator	6 ms	04/19/2025, 11:40:11 PM	04/19/2025, 11:40:55 PM	fpsa-asa-vcenter8.nva.local

After the tasks have been completed, check the Storage view in vCenter to confirm the datastore creation.

vSphere Client

Search in all environments

Administrator@NVALocal

otv\_vmfs\_iscsi\_1

Summary Monitor Configure Permissions Files Hosts VMs

Details

Type

VMFS 6

Hosts

2

Virtual machines

VM templates

Location

ds:///vmfs/volumes/68046c5f-4a194b85-77e2-0025b5aaa000/

Capacity and Usage

Last updated at 11:43 PM

Storage

498.33 GB free

1.42 GB used

499.75 GB capacity

VIEW STATS REFRESH

At this point, your FlexPod virtual infrastructure installation is complete and ready for operations. By using ONTAP tools plugin which integrates ASA storage system with vCenter, you can easily create additional datastores for your virtual infrastructure.

Next, we will provide procedures for you to configure the iSCSI software initiator in virtual servers to have direct access to iSCSI LUNs from the ASA storage system for your Microsoft SQL and Oracle database servers. The direct iSCSI LUN access enables you to extend application and storage best practices from physical servers into virtual servers.

## Microsoft SQL 2022 database server with direct iSCSI LUN access configuration

This section provides information on configuring direct iSCSI LUN access for SQL server, including installing NetApp Windows Host Utilities. Please refer to Appendix B for high-level installation process for Microsoft SQL Server 2022 and Microsoft SQL Server Manager.

Many recommendations and best-practices guides are available for most SQL server settings. But the recommendations vary depending on your needs. Therefore, you should thoroughly test and validate your specific database environment.

### Configure direct iSCSI LUN access for SQL server

Configuring direct iSCSI LUN access for the Windows virtual machine allows the SQL database server to perform IO directly to those LUNs. In addition, it also allows QoS to be applied directly on those LUNs to fine tune performance requirements.

The direct iSCSI access is made possible by creating vNICs for the VM to enable multipath access from both iSCSI-A and iSCSI-B networks, followed by LUN and igroup creation and mapping in the storage, iSCSI initiator configuration in the VM, and utilizing multipathing to access mapped storage LUNs.

### Add vNICs to the Windows VM for iSCSI network access

To add vNICs to the Windows VM, use the steps below.

1. Login to vCenter.
2. Right-click on the VM and click Edit Settings.
3. In the Edit Settings dialog, click Add New Device, select Network Adapter under Network.
4. Click Add New Device again to add a second Network Adapter.
5. For the New Network \* adapter, click port group drop-down list to Browse and select iSCSI-A port group and click OK.
6. For the New Network 2 \* adapter, click port group drop-down list to Browse and select iSCSI-B port group and click OK.

ADD NEW DEVICE ▾

> CPU	8 ▾ ⓘ	
> Memory	64	GB ▾
> Hard disk 1	200	GB ▾
> SCSI controller 0	VMware Paravirtual	
> Network adapter 1	IB-MGMT Network ▾	<input checked="" type="checkbox"/> Connected
> New Network *	iSCSI-A ▾	<input checked="" type="checkbox"/> Connected
> New Network 2 *	iSCSI-B ▾	<input checked="" type="checkbox"/> Connected
> CD/DVD drive 1 *	Datastore ISO File ▾	<input type="checkbox"/> Connected
> USB xHCI controller	USB 3.2	
> Video card	Specify custom settings ▾	
> SATA controller 0	AHCI	
> Security Devices	Not Configured	
> Other	Additional Hardware	

CANCEL

OK

- Click OK at the bottom of the dialog to save the settings.
- Open the VM settings again to confirm.

## Configure vNICs and confirm iSCSI fabric connectivity

To configure and confirm the vNIC and iSCSI fabric virtual connectivity, perform the following steps.

- Gather the VM's MAC addresses from iSCSI-A and iSCSI-B port group in vCenter Network view under the iSCSI-vDS.

vSphere Client									
Search in all environments									
Administrator@NVA.LOCAL ▾									
iSCSI-A ACTIONS									
Summary Monitor Configure Permissions Ports Hosts VMs									
	Port ID	Name	Connectee	Runtime MAC Address	Port Group	State	VLAN ID	VIF ID	
⋮	0	--	fpsa-asa-esxi-01.nva.local - vmk1	00:50:56:60:5f:e4	iSCSI-A	Link Up	VLAN access: 0		
⋮	1	--	fpsa-asa-esxi-02.nva.local - vmk1	00:50:56:67:c8:e1	iSCSI-A	Link Up	VLAN access: 0		
⋮	2	--	fpsa-asa-ol8vm-01	00:50:56:be:48:09	iSCSI-A	Link Up	VLAN access: 0		
⋮	3	--	--	--	iSCSI-A	--	VLAN access: 0		
⋮	4	--	fpsa-asa-win2022vm-01	00:50:56:be:78:af	iSCSI-A	Link Up	VLAN access: 0		

Port ID	Name	Connectee	Runtime MAC Address	Port Group	State	VLAN ID	VIF ID
8	--	fpga-asa-esxi-02.nva.local - v mk2	00:50:56:6f:13:cd	iSCSI-B	Link U p	VLAN acces s: 0	
9	--	fpga-asa-esxi-01.nva.local - v mk2	00:50:56:6c:04:04	iSCSI-B	Link U p	VLAN acces s: 0	
10	--	fpga-asa-ol8vm-01	00:50:56:be:b3:dc	iSCSI-B	Link U p	VLAN acces s: 0	
11	--	--	--	iSCSI-B	--	VLAN acces s: 0	
12	--	fpga-asa-win2022vm-01	00:50:56:be:0d:07	iSCSI-B	Link U p	VLAN acces s: 0	

Configure the VM's iSCSI network interfaces with appropriate iSCSI IP address and netmask in Ethernet Adapter TCP/IP properties.

Control Panel > Network and Internet > Network Connections

Organize Disable this network device Diagnose this connection Rename this connection View status of this connection Change settings of this connection

Ethernet0 Network vmxnet3 Ethernet Adapter

Ethernet2 Unidentified network vmxnet3 Ethernet Adapter #3

Ethernet1 Unidentified network vmxnet3 Ethernet Adapter #2

Ethernet1 Status

General

Connection

IPv4 Connectivity:

IPv6 Connectivity:

Media State:

Duration:

Speed:

Details...

Activity

Sent

Bytes: 0

Properties Disable

3 items 1 item selected

Ethernet1 Properties

Networking Sharing

Connect using:

vmxnet3 Ethernet Adapter #2

Configure...

This connection uses the following items:

- ☒ Client for Microsoft Networks
- ☒ File and Printer Sharing for Microsoft Networks
- ☒ QoS Packet Scheduler
- ☒ Internet Protocol Version 4 (TCP/IPv4)
- ☐ Microsoft Network Adapter Multiplexor Protocol
- ☒ Microsoft LLDP Protocol Driver
- ☐ Internet Protocol Version 6 (TCP/IPv6)

Install... Uninstall Properties

Description

Transmission Control Protocol/Internet Protocol. The default wide area network protocol that provides communication across diverse interconnected networks.

OK Cancel

Internet Protocol Version 4 (TCP/IPv4) Properties

General

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

☐ Obtain an IP address automatically

☒ Use the following IP address:

IP address: 172 . 22 . 73 . 84

Subnet mask: 255 . 255 . 255 . 0

Default gateway: . . .

☐ Obtain DNS server address automatically

☒ Use the following DNS server addresses:

Preferred DNS server: . . .

Alternate DNS server: . . .

☐ Validate settings upon exit

Advanced...

OK Cancel

```

C:\Users\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Ethernet0:

    Connection-specific DNS Suffix  . : nva.local
    IPv4 Address. . . . . : 172.22.72.84
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . : 172.22.72.1

Ethernet adapter Ethernet1:

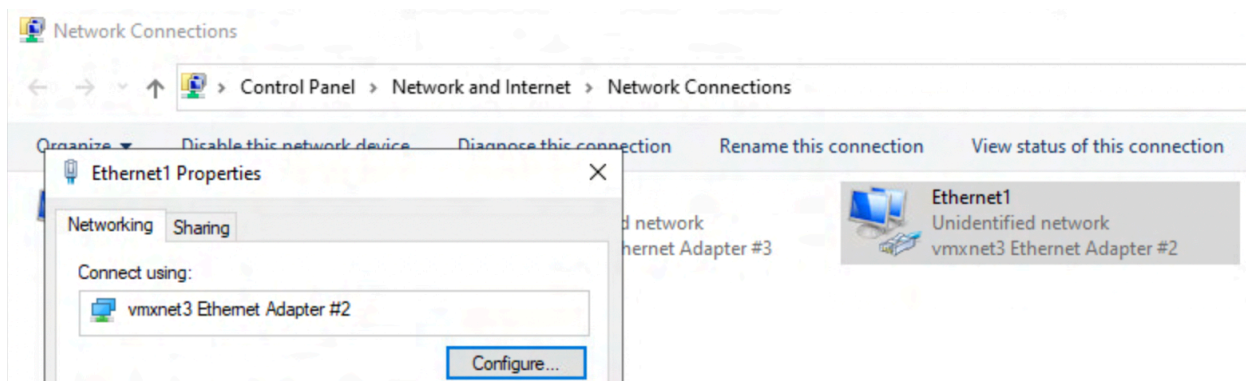
    Connection-specific DNS Suffix  . :
    IPv4 Address. . . . . : 172.22.73.84
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :

Ethernet adapter Ethernet2:

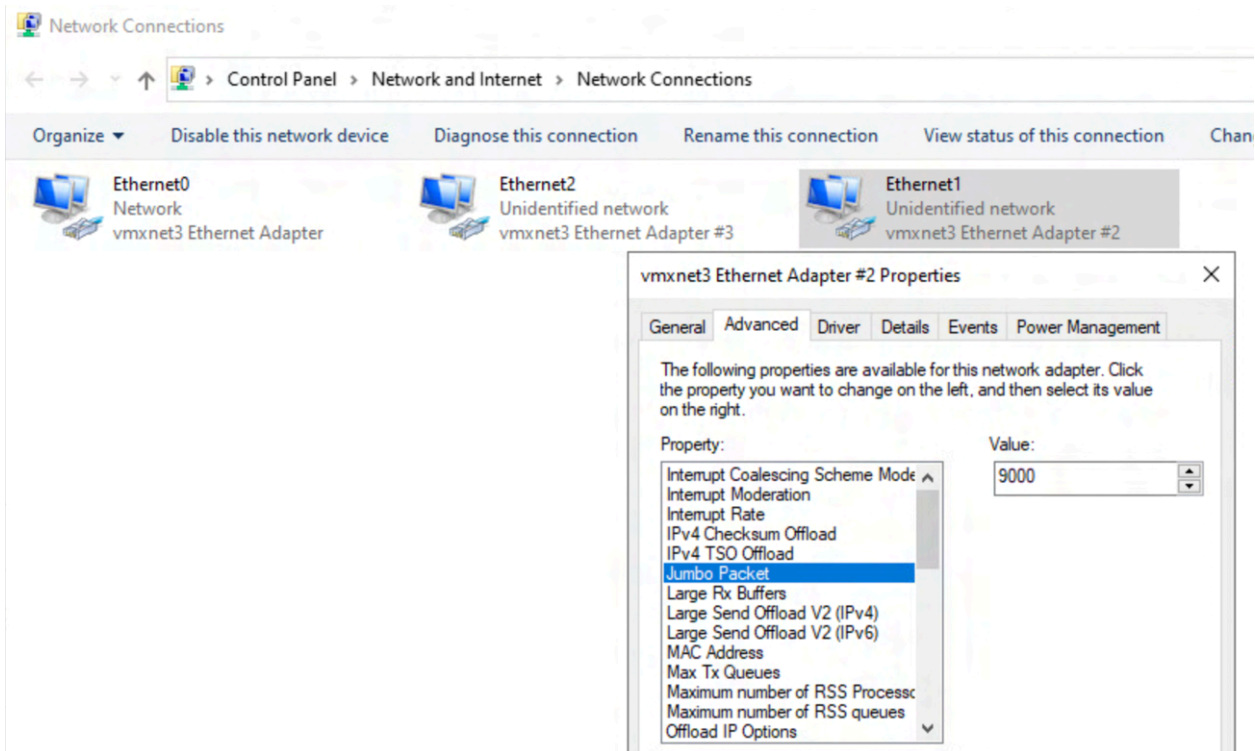
    Connection-specific DNS Suffix  . :
    IPv4 Address. . . . . : 172.22.74.84
    Subnet Mask . . . . . : 255.255.255.0
    Default Gateway . . . . . :

```

In the vmxnet3 Ethernet Adapter properties' Networking tab, click on Configure, go to Advanced tab, and set Jumbo Packet property value to 9000.







## 2. Login to the storage cluster and obtain the storage controllers' iSCSI LIFs.

```
fpsa-a50-u0909::> network interface show -role data
```

Vserver	Logical Interface	Status	Network Address/Mask	Current Node	Current Port	Is Home
svml	iscsi-lif-01a	up/up	172.22.73.101/24	fpsa-a50-u0909-01	e2b-2273	true
	iscsi-lif-01b	up/up	172.22.74.101/24	fpsa-a50-u0909-01	e4b-2274	true
	iscsi-lif-02a	up/up	172.22.73.102/24	fpsa-a50-u0909-02	e2b-2273	true
	iscsi-lif-02b	up/up	172.22.74.102/24	fpsa-a50-u0909-02	e4b-2274	true

4 entries were displayed.

## 3. Check the VM's iSCSI network connectivity by pinging all the iSCSI LIFs configured in the storage controllers with jumbo frame packet size.



```

C:\Users\Administrator>ping -l 9000 -n 1 172.22.73.101

Pinging 172.22.73.101 with 9000 bytes of data:
Reply from 172.22.73.101: bytes=9000 time<1ms TTL=64

Ping statistics for 172.22.73.101:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Administrator>ping -l 9000 -n 1 172.22.73.102

Pinging 172.22.73.102 with 9000 bytes of data:
Reply from 172.22.73.102: bytes=9000 time<1ms TTL=64

Ping statistics for 172.22.73.102:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Administrator>ping -l 9000 -n 1 172.22.74.101

Pinging 172.22.74.101 with 9000 bytes of data:
Reply from 172.22.74.101: bytes=9000 time<1ms TTL=64

Ping statistics for 172.22.74.101:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Administrator>ping -l 9000 -n 1 172.22.74.102

Pinging 172.22.74.102 with 9000 bytes of data:
Reply from 172.22.74.102: bytes=9000 time<1ms TTL=64

Ping statistics for 172.22.74.102:
    Packets: Sent = 1, Received = 1, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

```

## Install pre-requisites for NetApp Windows Host Utilities installation

NetApp recommends installing the **latest hotfixes**, **cumulative updates**, and security updates that are available from Microsoft.

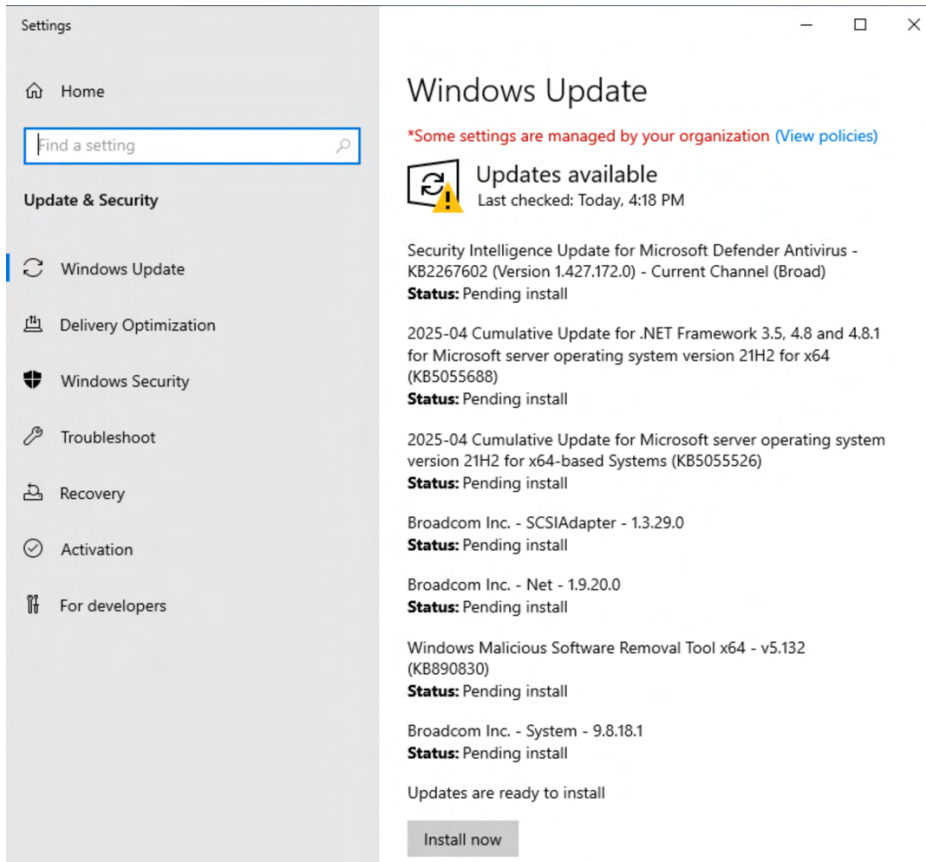
1. Download the updates from the Microsoft Update Catalog 2022 and Microsoft Security Update links below.

<https://www.catalog.update.microsoft.com/Search.aspx?q=update%20%22windows%20server%202022%22>

<https://msrc.microsoft.com/update-guide/>

Follow the instructions provided by Microsoft to install the hotfixes and updates.

Enable Windows Updates as needed from your host and install available updates.



Many hotfixes and updates require a reboot of your Windows host. Reboot the host if instructed.

## Install NetApp Windows Host Utilities

1. Before you install the Host Utilities, you should verify that your host and storage system configuration are supported. Please refer to the information in the following page.


[https://docs.netapp.com/us-en/ontap-sanhost/hu\\_wuhu\\_72.html](https://docs.netapp.com/us-en/ontap-sanhost/hu_wuhu_72.html)


Download NetApp Host Utilities Version 7.2 for Windows from NetApp support site download link below.

<https://mysupport.netapp.com/site/products/all/details/hostutilities/downloads-tab/download/61343/7.2/downloads>

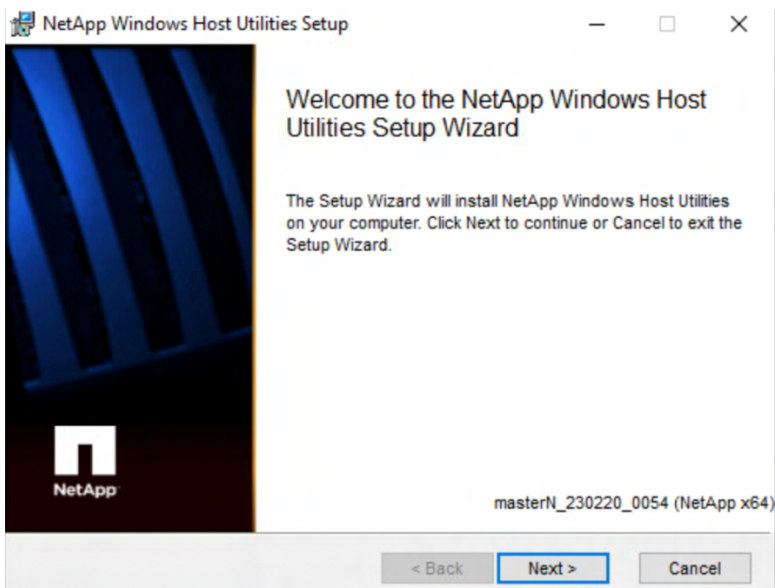
Launch Windows Host Utilities installer.

If you see a message indicating that Windows Defender SmartScreen prevented the application from starting, click the More info link and then click the Run anyway button to start the installation.

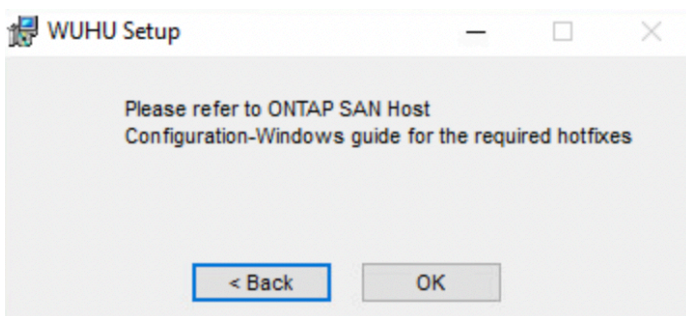




When the NetApp Windows Unified Host Utilities (WUHU) setup wizard is launched. Click Next.



Click OK on the WUHU Setup screen to confirm that the required hotfixes and updates have been installed.



Check the box on the End-User License Agreement screen and click Next.  
Select Yes, install support for Multipath I/O and click Next.




Accept the default destination folder and click Next.

Click Next to accept the default installation location or click Browse to specify an alternative location.

Click Install to start the installation of the host utilities.

Click OK for the following pop-up message which indicates that no updates were made to FC settings.



After the installation is completed, click Finish.

Click Yes to restart the virtual machine.

After reboot, use Windows PowerShell to confirm that Multipath-IO feature is installed.

```
PS C:\Users\Administrator> Get-WindowsFeature -Name Multipath-IO

Display Name                                     Name                                     Install State
-----
[X] Multipath I/O                               Multipath-IO                             Installed

PS C:\Users\Administrator> _
```

Enable Microsoft Device Specific Module (MSDSM) to automatically claim SAN disks for Microsoft Multipath I/O (MPIO) for the iSCSI bus type.

```
PS C:\Users\Administrator> Enable-MSDSMAutomaticClaim -BusType iSCSI

VendorId ProductId
-----
MSFT2005 iSCSIBusType_0x9
False

PS C:\Users\Administrator> Get-MSDSMAutomaticClaimSettings

Name Value
----
iSCSI True
SAS   False
```

## Enable Microsoft iSCSI initiator service

To enable the iSCSI service on Windows, follow the steps below.

1. Start the MSiSCSI service with Windows PowerShell Start-Service command.
2. Set service startup type to automatic.
3. Retrieve the iSCSI initiator qualified node name for later storage configuration.

```
PS C:\Users\Administrator> Start-Service MSiSCSI
PS C:\Users\Administrator> Set-Service MSiSCSI -StartupType Automatic
PS C:\Users\Administrator> (Get-InitiatorPort).NodeAddress
iqn.1991-05.com.microsoft:win2022vm-01
```

## Create LUNs and provide access to the Windows iSCSI initiator

Follow the steps below to create LUNs and provide Windows iSCSI initiator access to them.

1. Create a set of LUNs for the Windows SQL database usage.

```
fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_log_1 -size 200g -ostype windows
fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_log_2 -size 200g -ostype windows

fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_data_1 -size 500g -ostype windows
fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_data_2 -size 500g -ostype windows
fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_data_3 -size 500g -ostype windows
fpsa-a50-u0909:> lun create -path fpsa_asa_win2022vm_01_sql_data_4 -size 500g -ostype windows
```

**Note:** Here we are creating two LUNs for log and four LUNs for SQL data. Adjust the number of LUNs and their sizes to suit your database requirements.

2. Create initiator group and include the Windows software iSCSI initiator in the igroup.

```
fpsa-a50-u0909:> igroup create -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -protocol iscsi -ostype windows -initiator iqn.1991-05.com.microsoft:win2022vm-01
```

3. Map the created LUNs to the initiator group.

```
fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_log_1 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 1
fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_log_2 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 2

fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_data_1 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 11
fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_data_2 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 12
fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_data_3 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 13
fpsa-a50-u0909:> lun map -path fpsa_asa_win2022vm_01_sql_data_4 -igroup FlexPod-ASA-fpsa-asa-win2022vm-01 -lun-id 14

fpsa-a50-u0909:> lun show -m -igroup FlexPod-ASA-fpsa-asa-win2022vm-01
```



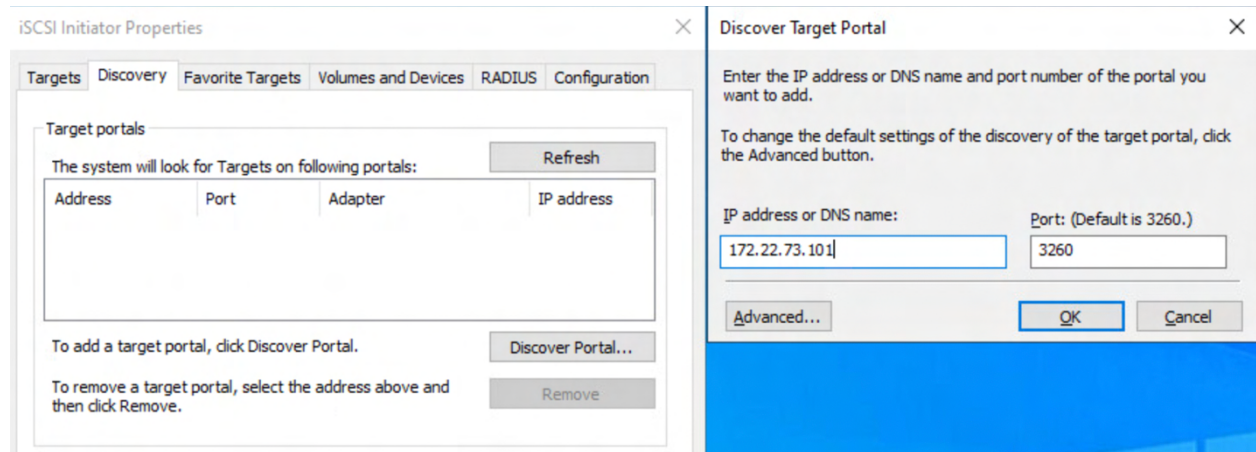
Vserver	Path	Igroup	LUN ID	Protocol
svml	fpsa_asa_win2022vm_01_sql_data_1	FlexPod-ASA-fpsa-asa-win2022vm-01	11	iscsi
svml	fpsa_asa_win2022vm_01_sql_data_2	FlexPod-ASA-fpsa-asa-win2022vm-01	12	iscsi
svml	fpsa_asa_win2022vm_01_sql_data_3	FlexPod-ASA-fpsa-asa-win2022vm-01	13	iscsi
svml	fpsa_asa_win2022vm_01_sql_data_4	FlexPod-ASA-fpsa-asa-win2022vm-01	14	iscsi
svml	fpsa_asa_win2022vm_01_sql_log_1	FlexPod-ASA-fpsa-asa-win2022vm-01	1	iscsi
svml	fpsa_asa_win2022vm_01_sql_log_2	FlexPod-ASA-fpsa-asa-win2022vm-01	2	iscsi

6 entries were displayed.

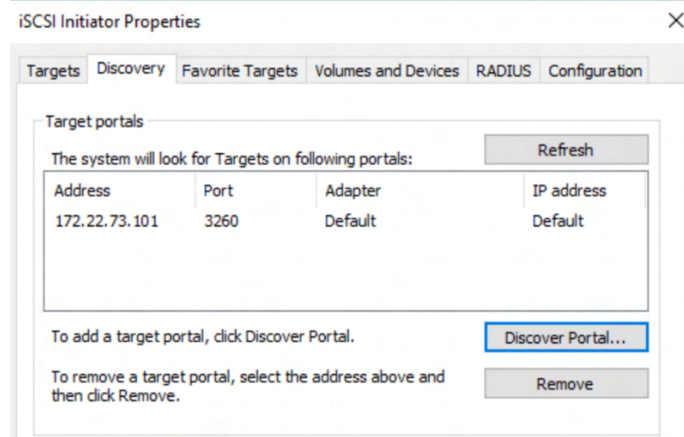
## Add iSCSI LIF addresses in iSCSI Initiator properties

To add all the iSCSI LIF addresses for Microsoft iSCSI Initiator to discover, follow the steps below.

1. Start Server Manager.
2. Select iSCSI Initiator from the Tools menu in Server Manager to open the iSCSI Initiator Properties.
3. Select Discovery tab and click Discover Portal to add a target portal.
4. Enter the first iSCSI LIF address configured in storage into the IP address field in the Discover Target Portal dialog.

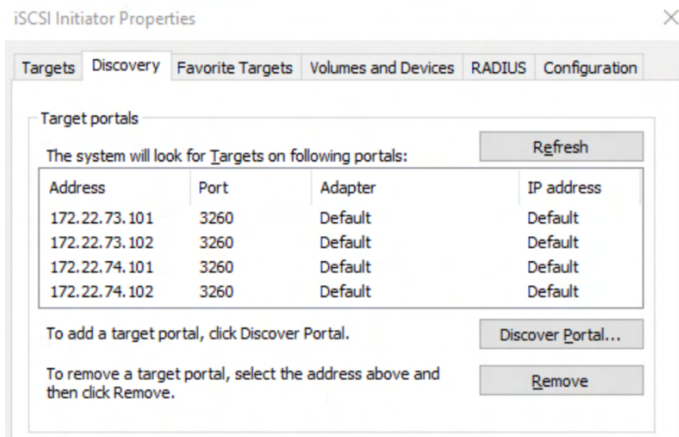


5. Click OK. The iSCSI LIF is added to the target portals list.



6. Repeat the steps above to add the remaining iSCSI LIFs configured in storage for them to all show up in the target portals list.

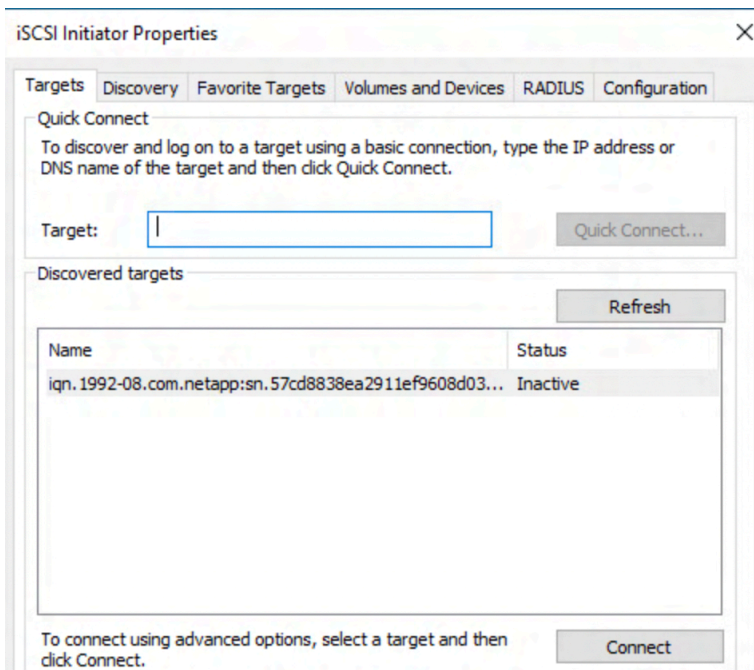




## Connect to iSCSI target


To connect the iSCSI initiator to target, follow the steps below.

1. Open iSCSI Initiator Properties from Server Manager Tools menu.
2. Select Targets tab.

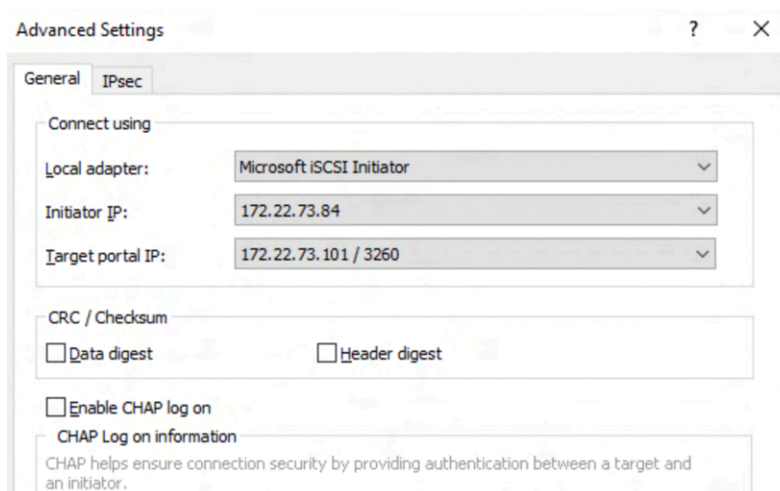


**Note:** At this point, the iSCSI target IQN from storage should show up under the Discovered targets list with an Inactive status.

3. Highlight the target, then click on Connect.
4. Check Enable multi-path box and click Advanced button in the Connect to Target dialog.

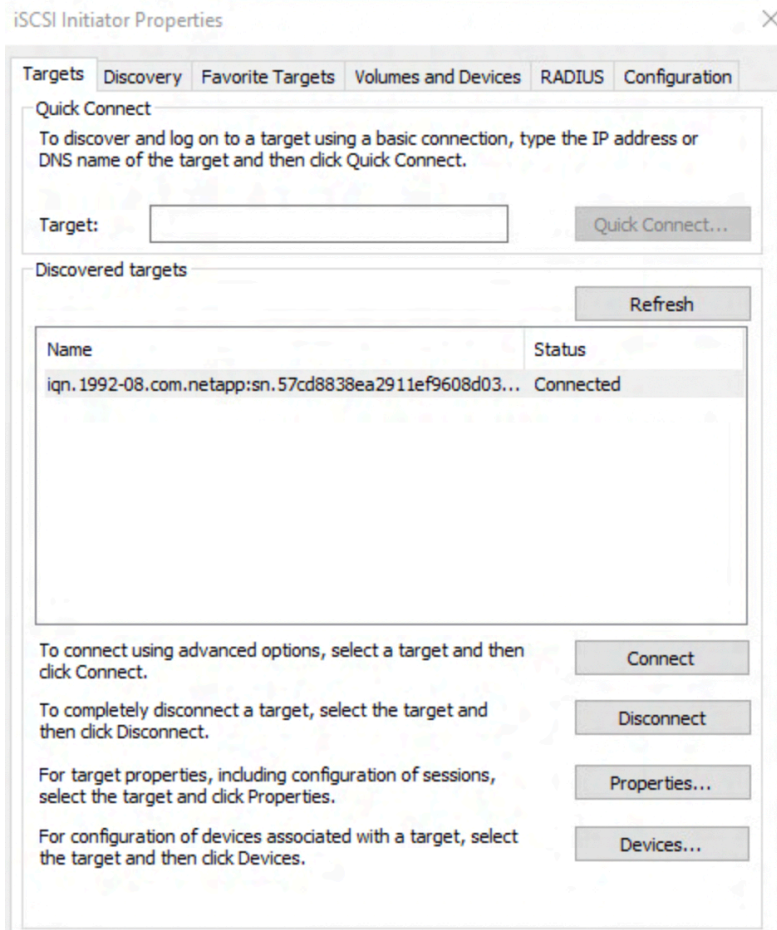


5. In the Advanced Settings dialog, select Microsoft iSCSI Initiator from the Local adapter drop-down list, select the iSCSI initiator IP from the iSCSI-A fabric, select the first iSCSI target LIF in the iSCSI-A fabric.

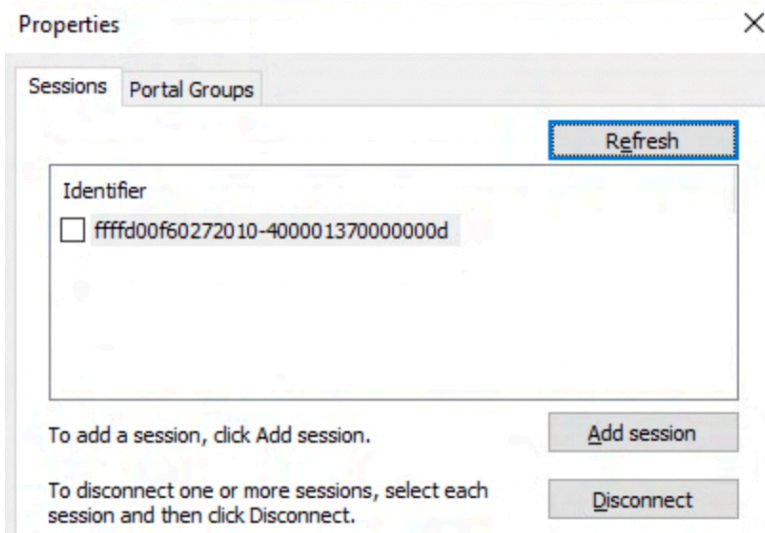


6. Click OK on the Advanced Settings dialog and then click OK in the Connect to Target dialog to complete the connection configuration of the first path.

**Note:** The status of the target should become Connected.




7. Click the Properties button with the target selected to bring up the Properties view. Under the Sessions tab, it should list an identifier for a connected session.



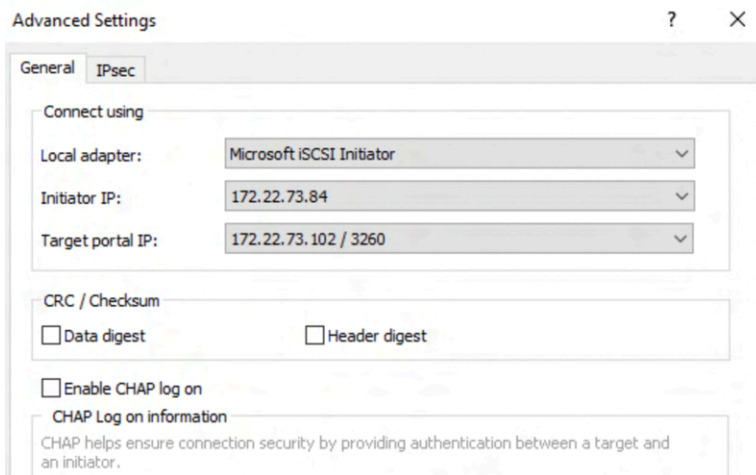
**Note:** Click Add session to open the Connect to Target dialog. The remaining paths can be added by using Add session in the Properties view of the target.

8. Click on Properties and select Add session.

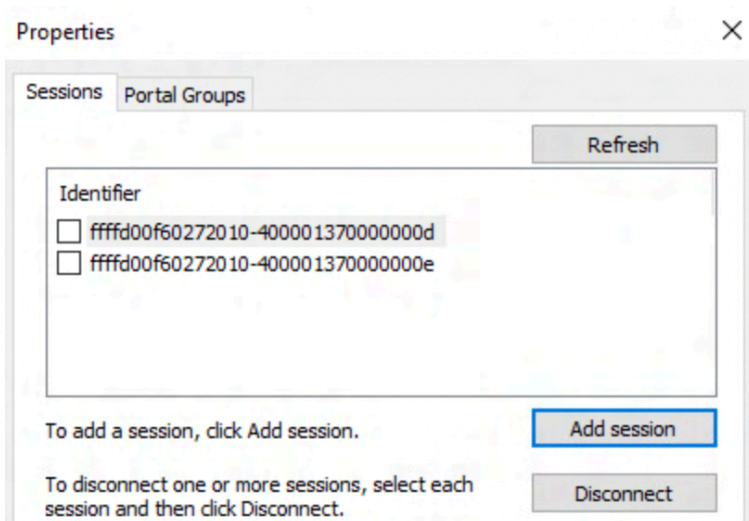
9. Check Enable multi-path and click Advanced.



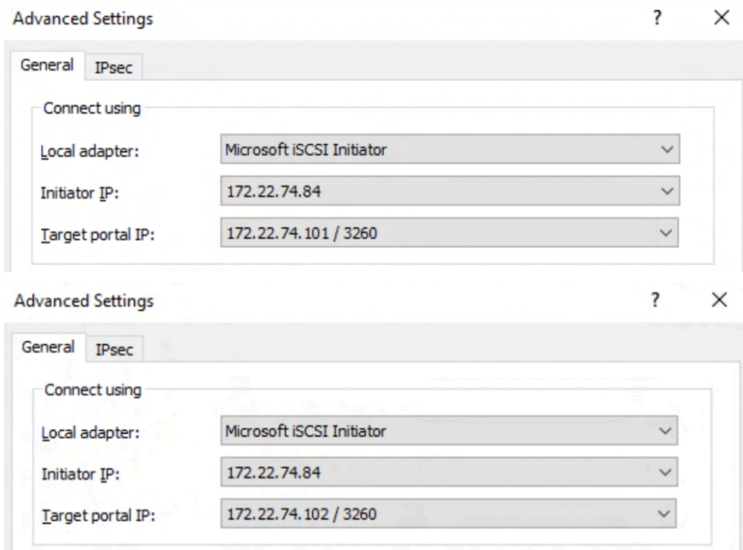
10. In the Advanced Settings dialog, select Microsoft iSCSI Initiator from the Local adapter drop-down list, select the iSCSI initiator IP from the iSCSI-A fabric, select the second iSCSI target LIF in the iSCSI-A fabric.



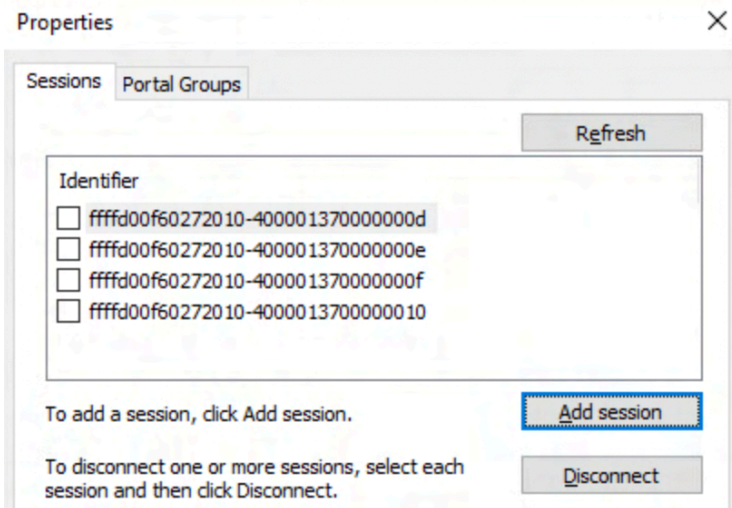
11. Click OK on the Advanced Settings dialog and then click OK in the Connect to Target dialog to complete the connection configuration of the second session.



12. Repeat steps 8 to 12 for the two remaining sessions, using initiator IP address in iSCSI-B fabric and the two iSCSI target LIFs in the iSCSI-B fabric for the two sessions.



13. Afterwards, there should be four sessions listed in the Properties dialog.



14. Click OK on the Properties dialog for the target.

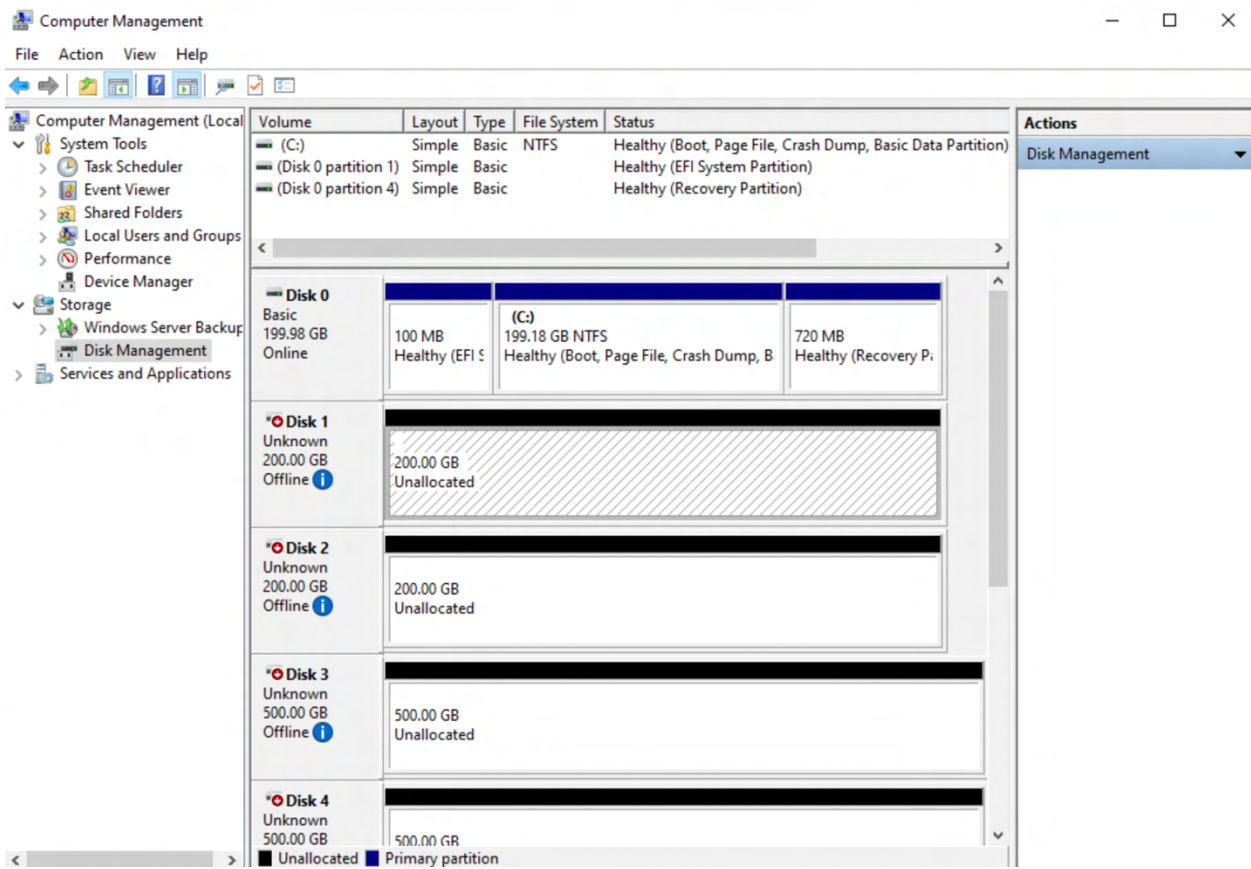
15. Click OK on the iSCSI Initiator Properties dialog.

## Check for the discovered LUNs and confirm multipath access

To check for the LUNs discovered by the Microsoft iSCSI initiator, follow the steps below.

1. From Server Manager Tools menu, select Computer Management.
2. On the left pane, click Disk Management under Storage and you should see all the LUNs mapped to the host which are discovered by the Microsoft iSCSI initiator.

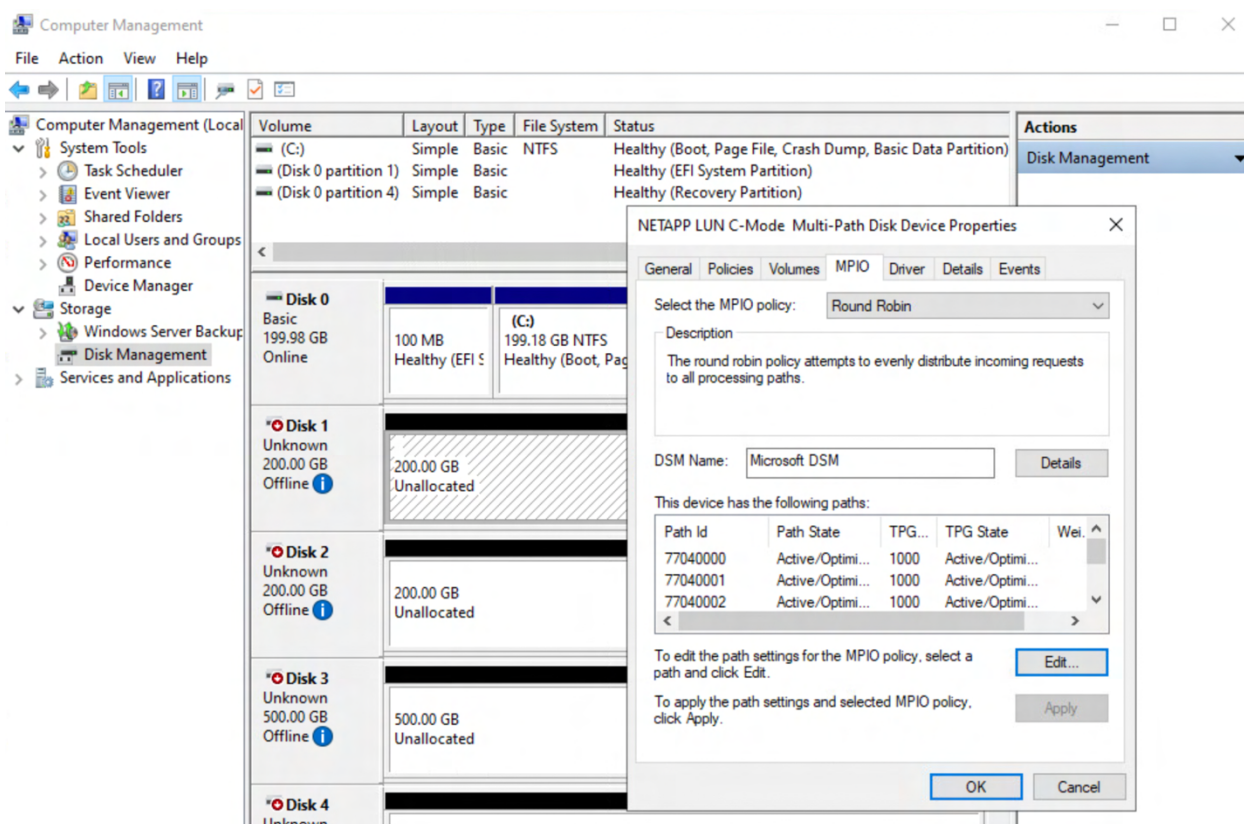




Right-click on one of the discovered disks near the Disk # and select Properties.

Click on the MPIO tab. The MPIO tab lists the MPIO setting information as well as the paths associated with the LUN / disk. You should see 4 total paths listed with target port group (TPG) State of Active/Optimized.





## Oracle 21c database server with direct iSCSI LUN access configuration

FlexPod converged infrastructure with Cisco UCS X-Series Direct and NetApp ASA storage provides a highly available SAN solution infrastructure for enterprise Oracle Real Application Clusters (RAC) database environment created on Oracle Grid Infrastructure and using Oracle Automatic Storage Management (ASM). During the solution validation, we deployed a single node Oracle 21c instance in both virtual machine and bare-metal host formats.

Please refer to the Appendices D and E for highlights on Oracle Grid Infrastructure and Oracle database installation steps used for this validation which is applicable to both the Oracle database virtual machine and bare-metal Oracle database server. In addition, we also included highlights of bare-metal Oracle Linux 8 server configuration for Oracle deployment.

While Oracle database deployment procedures can vary between environments and when using different protocols, there are still a lot of similarities between deploying it on a virtual machine versus on a bare-metal server or when using iSCSI versus using FC protocol. The following sections highlight some of the configuration aspects for configuring direct iSCSI LUN access for Oracle 21c database deployment on an Oracle Linux 8 virtual machine.

**Note:** It is not within the scope of a FlexPod solution documentation to cover Oracle deployment procedures. Please refer to Oracle documentation links listed in the References section for details on the deployment requirements and installation and configuration of Oracle Grid Infrastructure and Oracle RAC.

## Configure direct iSCSI LUN access for Oracle Linux 8 virtual machine

In this section, we provide example procedures for configuring Oracle Linux 8 virtual machine to directly access iSCSI LUNs from the storage. The access is made possible by creating vNICs for the VM to enable multi-path access from both iSCSI-A and iSCSI-B networks, followed by LUN and igroup creation

and mapping in the storage, iSCSI initiator configuration in the VM, and utilizing multipath to access storage.

## Add vNICs to VM for iSCSI network access

To add vNICs to the VM, use the steps below.

1. Login to vCenter.
2. Right-click on the VM and click Edit Settings.
3. In the Edit Settings dialog, click Add New Device, select Network Adapter under Network.
4. Click Add New Device again to add a second Network Adapter.
5. For the New Network \* adapter, click port group drop-down list to Browse and select iSCSI-A port group and click OK.
6. For the New Network 2 \* adapter, click port group drop-down list to Browse and select iSCSI-B port group and click OK.
7. Click OK at the bottom of the dialog to save the settings.
8. Open the VM settings again to confirm.

**Edit Settings** | fpsa-asa-ol8vm-01 ×

Virtual Hardware | VM Options | Advanced Parameters

ADD NEW DEVICE ▾

> CPU	4 ▾ ⓘ	
> Memory	64 ▾	GB ▾
> Hard disk 1	200 ▾	GB ▾ ⋮
> SCSI controller 0	VMware Paravirtual	⋮
> Network adapter 1	IB-MGMT Network ▾	<input checked="" type="checkbox"/> Connected ⋮
> Network adapter 2	iSCSI-A ▾	<input checked="" type="checkbox"/> Connected ⋮
> Network adapter 3	iSCSI-B ▾	<input checked="" type="checkbox"/> Connected ⋮
> CD/DVD drive 1	Datastore ISO File ▾	<input checked="" type="checkbox"/> Connected ⋮
> Video card	Specify custom settings ▾	
> SATA controller 0	AHCI	⋮
> Security Devices	Not Configured	
> Other	Additional Hardware	

CANCEL

OK

## Configure and confirm vNIC and iSCSI fabric connectivity

To configure and confirm the vNIC and iSCSI fabric virtual connectivity, perform the following steps.

1. Gather the VM's MAC addresses from iSCSI-A and iSCSI-B port group in vCenter Network view under the iSCSI-vDS.

The top screenshot shows the vSphere Client interface for iSCSI-A. The 'Ports' tab is selected, displaying a table with columns: Port ID, Name, Connectee, Runtime MAC Address, Port Group, State, VLAN ID, VIF ID, and Segm ID. The table lists three ports: Port 0 connected to fpsa-asa-esxi-01.nva.local - vmk1, Port 1 connected to fpsa-asa-esxi-02.nva.local - vmk1, and Port 2 connected to fpsa-asa-ol8vm-01. All ports are in a 'Link Up' state.

The bottom screenshot shows the vSphere Client interface for iSCSI-B. The 'Ports' tab is selected, displaying a table with columns: Port ID, Name, Connectee, Runtime MAC Address, Port Group, State, VLAN ID, VIF ID, and Segm ID. The table lists three ports: Port 8 connected to fpsa-asa-esxi-02.nva.local - vmk2, Port 9 connected to fpsa-asa-esxi-01.nva.local - vmk2, and Port 10 connected to fpsa-asa-ol8vm-01. All ports are in a 'Link Up' state.

2. Configure the VM's iSCSI network interfaces with appropriate IP address, netmask, and MTU using your preferred network configuration tool and confirm the interface configurations.

```
ens35: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    inet 172.22.73.85 netmask 255.255.255.0 broadcast 172.22.73.255
    inet6 fe80::250:56ff:febe:4809 prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:be:48:09 txqueuelen 1000 (Ethernet)
    RX packets 232 bytes 25394 (24.7 KiB)
    RX errors 0 dropped 74 overruns 0 frame 0
    TX packets 102 bytes 11614 (11.3 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

ens36: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    inet 172.22.74.85 netmask 255.255.255.0 broadcast 172.22.74.255
    inet6 fe80::250:56ff:febe:b3dc prefixlen 64 scopeid 0x20<link>
    ether 00:50:56:be:b3:dc txqueuelen 1000 (Ethernet)
    RX packets 235 bytes 25766 (25.1 KiB)
    RX errors 0 dropped 74 overruns 0 frame 0
    TX packets 101 bytes 11524 (11.2 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Login to the storage cluster and obtain the storage controllers' iSCSI LIFs.

```
fpsa-a50-u0909:> network interface show -role data
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svm1	iscsi-lif-01a	up/up	172.22.73.101/24	fpsa-a50-u0909-01	e2b-2273	true
	iscsi-lif-01b	up/up	172.22.74.101/24	fpsa-a50-u0909-01	e4b-2274	true
	iscsi-lif-02a	up/up	172.22.73.102/24	fpsa-a50-u0909-02	e2b-2273	true
	iscsi-lif-02b	up/up	172.22.74.102/24	fpsa-a50-u0909-02	e4b-2274	true

4 entries were displayed.

4. Check the VM's physical iSCSI network connectivity by pinging all the iSCSI LIFs configured in the storage controllers with jumbo frame packet size.

```
[admin@fpsa-asa-ol8vm-01 ~]$ ping -c 2 172.22.73.101 -s 9000
PING 172.22.73.101 (172.22.73.101) 9000(9028) bytes of data.
9008 bytes from 172.22.73.101: icmp_seq=1 ttl=64 time=0.265 ms
9008 bytes from 172.22.73.101: icmp_seq=2 ttl=64 time=0.410 ms

--- 172.22.73.101 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1055ms
rtt min/avg/max/mdev = 0.265/0.337/0.410/0.074 ms
```

```
[admin@fpsa-asa-ol8vm-01 ~]$ ping -c 2 172.22.73.102 -s 9000
PING 172.22.73.102 (172.22.73.102) 9000(9028) bytes of data.
9008 bytes from 172.22.73.102: icmp_seq=1 ttl=64 time=0.224 ms
9008 bytes from 172.22.73.102: icmp_seq=2 ttl=64 time=0.269 ms

--- 172.22.73.102 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1032ms
rtt min/avg/max/mdev = 0.224/0.246/0.269/0.027 ms

[admin@fpsa-asa-ol8vm-01 ~]$ ping -c 2 172.22.74.101 -s 9000
PING 172.22.74.101 (172.22.74.101) 9000(9028) bytes of data.
9008 bytes from 172.22.74.101: icmp_seq=1 ttl=64 time=0.285 ms
9008 bytes from 172.22.74.101: icmp_seq=2 ttl=64 time=0.388 ms

--- 172.22.74.101 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1030ms
rtt min/avg/max/mdev = 0.285/0.336/0.388/0.054 ms

[admin@fpsa-asa-ol8vm-01 ~]$ ping -c 2 172.22.74.102 -s 9000
PING 172.22.74.102 (172.22.74.102) 9000(9028) bytes of data.
9008 bytes from 172.22.74.102: icmp_seq=1 ttl=64 time=0.481 ms
9008 bytes from 172.22.74.102: icmp_seq=2 ttl=64 time=0.309 ms

--- 172.22.74.102 ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1059ms
rtt min/avg/max/mdev = 0.309/0.395/0.481/0.086 ms
```

## Create LUNs and provide access to the Oracle Linux 8 initiator

Follow the steps below to create LUNs and provide Oracle Linux 8 initiator access to them.

1. Create a set of LUNs for the Oracle Linux 8 database usage. Adjust the number of LUNs and their sizes to suit your database requirements.

```
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_redolog_1 -size 200g -ostype linux
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_redolog_2 -size 200g -ostype linux

fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_ocrvote_1 -size 100g -ostype linux
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_ocrvote_2 -size 100g -ostype linux

fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_slobdata_1 -size 500g -ostype linux
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_slobdata_2 -size 500g -ostype linux
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_slobdata_3 -size 500g -ostype linux
fpsa-a50-u0909:> lun create -path fpsa_asa_ol8vm_01_oracle_slobdata_4 -size 500g -ostype linux
...
```

**Note:** The two OCRVOTE LUNs are used for Oracle cluster registry and voting disks. The two REDOLOG LUNs are for logs and the eight SLOBDATA LUNs are for SLOB data.

2. Create initiator group (igroup) and include the Oracle Linux 8 software initiator to the igroup.

```
fpsa-a50-u0909:> igroup create -igroup FlexPod-ASA-fpsa-asa-ol8vm-01 -protocol iscsi -ostype
linux -initiator iqn.2010-11.com.flexpod:fpsa-asa-ol8vm-01
```

**Note:** Provide a unique iSCSI IQN appropriate for your environment and use the same IQN for your initiator configuration later.

3. Map the created LUNs to the initiator group.

```
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_redolog_1 -igroup FlexPod-ASA-fpsa-asa-
ol8vm-01 -lun-id 1
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_redolog_2 -igroup FlexPod-ASA-fpsa-asa-
ol8vm-01 -lun-id 2

fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_ocrvote_1 -igroup FlexPod-ASA-fpsa-asa-
ol8vm-01 -lun-id 3
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_ocrvote_2 -igroup FlexPod-ASA-fpsa-asa-
ol8vm-01 -lun-id 4
```

```
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_slobdata_1 -igroup FlexPod-ASA-fpsa-asa-ol8vm-01 -lun-id 11
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_slobdata_2 -igroup FlexPod-ASA-fpsa-asa-ol8vm-01 -lun-id 12
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_slobdata_3 -igroup FlexPod-ASA-fpsa-asa-ol8vm-01 -lun-id 13
fpsa-a50-u0909:> lun map -path fpsa_asa_ol8vm_01_oracle_slobdata_4 -igroup FlexPod-ASA-fpsa-asa-ol8vm-01 -lun-id 14
...
```

## Configure Oracle Linux 8 software iSCSI initiator

1. Check to see if the `iscsi-initiator-utils` package is already installed and install it if needed.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo dnf list installed | grep iscsi-initiator-utils
[sudo] password for admin:
iscsi-initiator-utils.x86_64                               6.2.1.4-8.git095f59c.0.1.el8
@anaconda
iscsi-initiator-utils-iscsiuio.x86_64                     6.2.1.4-8.git095f59c.0.1.el8
@anaconda
```

**Note:** If the package is not already installed, install it with the `dnf install iscsi-initiator-util` command.

Edit the `/etc/iscsi/initiatorname.iscsi` file to provide a unique `InitiatorName` parameter appropriate for your environment. Be sure that it matches the IQN used in the created storage group above.

```
InitiatorName=iqn.2010-11.com.flexpod:fpsa-asa-ol8vm-01
```

Restart the `iscsid` service and check its status.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo systemctl restart iscsid

[admin@fpsa-asa-ol8vm-01 ~]$ sudo systemctl status iscsid
• iscsid.service - Open-iSCSI
  Loaded: loaded (/usr/lib/systemd/system/iscsid.service; disabled; vendor preset: disabled)
  Active: active (running) since Sun 2025-03-16 01:17:49 EDT; 5s ago
    Docs: man:iscsid(8)
           man:iscsiuio(8)
           man:iscsiadm(8)
  Main PID: 4071 (iscsid)
    Status: "Ready to process requests"
     Tasks: 1 (limit: 409313)
  Memory: 1.9M
  CGroup: /system.slice/iscsid.service
          └─4071 /usr/sbin/iscsid -f -d2

Mar 16 01:17:49 fpsa-asa-ol8vm-01.nva.local systemd[1]: Starting Open-iSCSI...
Mar 16 01:17:49 fpsa-asa-ol8vm-01.nva.local iscsid[4071]: iscsid: InitiatorName=iqn.2010-11.com.flexpod:fpsa-asa-ol8vm-01
Mar 16 01:17:49 fpsa-asa-ol8vm-01.nva.local iscsid[4071]: iscsid: InitiatorAlias=fpsa-asa-ol8vm-01.nva.local
Mar 16 01:17:49 fpsa-asa-ol8vm-01.nva.local iscsid[4071]: iscsid: Max file limits 1024 262144
Mar 16 01:17:49 fpsa-asa-ol8vm-01.nva.local systemd[1]: Started Open-iSCSI.
```

Discover target with the `iscsiadm` tool and specify discovery type of `sendtarget` and one of the target portal IP addresses.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo iscsiadm -m discovery -t sendtargets -p 172.22.73.101
172.22.73.101:3260,1027 iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2
172.22.74.102:3260,1030 iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2
172.22.73.102:3260,1029 iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2
172.22.74.101:3260,1028 iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2
```

Login to the target with the target IQN.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo iscsiadm -m node -T iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2 -l
Logging in to [iface: default, target: iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.73.101,3260]
```

```

Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.74.102,3260]
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.73.102,3260]
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.74.101,3260]
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.73.101,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.74.102,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.73.102,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2, portal: 172.22.74.101,3260] successful.

```

Enable automatic login upon reboot.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo iscsiadm -m node -L automatic
```

## Rescan and discover LUNs

To perform a rescan to discover LUNs from the connected target, perform the following steps.

1. Issue the `rescan-scsi-bus.sh` command to rescan the SCSI bus.

```

[admin@fpsa-asa-ol8vm-01 ~]$ sudo rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 0 0 0 0 ...
  OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 00
        Vendor: VMware   Model: Virtual disk   Rev: 2.0
        Type:   Direct-Access                     ANSI SCSI revision: 06
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
...

```

## Enable multipathing for LUNs

Having multiple paths to storage devices provides path redundancy and the ability to load-balance I/O commands between paths. The multipathing tool consolidates multiple I/O paths to a LUN as a single multipath device. To check and enable multipathing for storage LUNs, follow the steps below.

1. Check to see if multipathing is enabled.

```

[admin@fpsa-asa-ol8vm-01 ~]$ sudo multipath -ll
[sudo] password for admin:
Mar 16 09:54:05 | /etc/multipath.conf does not exist, blacklisting all devices.
Mar 16 09:54:05 | You can run "/sbin/mpathconf --enable" to create
Mar 16 09:54:05 | /etc/multipath.conf. See man mpathconf(8) for more details
Mar 16 09:54:05 | DM multipath kernel driver not loaded

```

**Note:** As shown in the output above, multipath is not yet enabled because `/etc/multipath.conf` does not exist.

Enable multipath configuration.

```

[admin@fpsa-asa-ol8vm-01 ~]$ sudo /sbin/mpathconf --enable

[admin@fpsa-asa-ol8vm-01 ~]$ ls -l /etc/multipath.conf
-rw----- 1 root root 421 Mar 16 09:55 /etc/multipath.conf

```

Restart multipath daemon `multipathd` and check its status.

```

[admin@fpsa-asa-ol8vm-01 ~]$ sudo systemctl restart multipathd

[admin@fpsa-asa-ol8vm-01 ~]$ sudo systemctl status multipathd
• multipathd.service - Device-Mapper Multipath Device Controller
  Loaded: loaded (/usr/lib/systemd/system/multipathd.service; enabled; vendor preset: enabled)
  Active: active (running) since Sun 2025-03-16 09:56:36 EDT; 13s ago

```



```

Process: 45679 ExecStartPre=/sbin/multipath -A (code=exited, status=0/SUCCESS)
Process: 45677 ExecStartPre=/sbin/modprobe -a scsi_dh_alua scsi_dh_emc scsi_dh_rdac dm-
multipath (code=exited, status=0/SUCCESS)
Main PID: 45681 (multipathd)
Status: "up"
Tasks: 7
Memory: 14.9M
CGroup: /system.slice/multipathd.service
└─45681 /sbin/multipathd -d -s

Mar 16 09:56:36 fpsa-asa-ol8vm-01.nva.local multipathd[45681]: 3600a0980383234486724587338696c66:
load table [0 1048576000 multi>
...

```

Check the multipath device information.

```

[admin@fpsa-asa-ol8vm-01 ~]$ sudo multipath -ll
[sudo] password for admin:
3600a0980383234486724587338696c58 dm-3 NETAPP,LUN C-Mode
size=200G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
   |- 33:0:0:1 sdi 8:128 active ready running
   |- 34:0:0:1 sdb 8:16 active ready running
   |- 35:0:0:1 sdc 8:32 active ready running
   `-- 36:0:0:1 sdd 8:48 active ready running
...

3600a0980383234486724587338696c61 dm-4 NETAPP,LUN C-Mode
size=500G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
   |- 33:0:0:11 sds 65:32 active ready running
   |- 34:0:0:11 sdj 8:144 active ready running
   |- 35:0:0:11 sdh 8:112 active ready running
   `-- 36:0:0:11 sdk 8:160 active ready running
...

```

**Note:** The output above provides a list of LUNs claimed by multipathd as well as the device information such as LUN size and ALUA status for the available paths. For the first device dm-3 listed above, the multipath device dm-3 includes four device paths: sdi, sdb, sdc, and sdd.

## Install NetApp Linux Host Utilities

1. Download the NetApp Linux Unified Host Utilities 7.1 tool from <https://support.netapp.com>.

The screenshot shows the NetApp Support website. At the top, there is a search bar and navigation tabs for PRODUCTS, SYSTEMS, DOCS & KNOWLEDGE BASE, COMMUNITY, DOWNLOADS, TOOLS & SECURITY, CASES & PARTS, and PARTNERS. Below the navigation, there is a section for "Discover GenAI Search" with a video player. The main content area shows the "Host Utilities 7.1" page, dated 20-Feb-2025. Under the heading "Linux Host Utilities", there are two download buttons: "DOWNLOAD NETAPP\_LINUX\_UNIFIED\_HOST\_UTILITIES-7.1.I386.RPM [569.10 KB]" and "DOWNLOAD NETAPP\_LINUX\_UNIFIED\_HOST\_UTILITIES-7.1.X86\_64.RPM [575.34 KB]".

**Note:** You will need a NetApp support account and login with your username and password.

1. Install the NetApp Linux Unified Host Utilities tool and verify the version.

```

[admin@fpsa-asa-ol8vm-01 ~]$ cd /tmp
[admin@fpsa-asa-ol8vm-01 tmp]$ ls netapp*
netapp_linux_unified_host_utilities-7-1.x86_64.rpm

```

```
[admin@fpsa-asa-ol8vm-01 tmp]$ sudo rpm -ivh netapp_linux_unified_host_utilities-7-1.x86_64.rpm
[sudo] password for admin:
Verifying... ##### [100%]
Preparing... ##### [100%]
Warning: libnl.so library not found, some sanlun commands may not work. Refer Linux Host
Utilities Installation and Set up Guide for more details
Warning: libHBAAPI.so library not found, some sanlun commands may not work. Refer Linux Host
Utilities Installation and Setup Guide for more details
Updating / installing...
 1:netapp_linux_unified_host_utiliti##### [100%]

[admin@fpsa-asa-ol8vm-01 tmp]$ sudo sanlun version
sanlun version 7.1.386.1644
```

**Note:** You must have the root privilege to run the sanlun tool.

## 2. Use the sanlun tool to list the discovered NetApp LUNs.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo sanlun lun show -p

          ONTAP Path: svm1:fpsa_asa_ol8vm_01_oracle_redolog_1
            LUN: 1
        LUN Size: 200g
        Product: cDOT
    Host Device: 3600a0980383234486724587338696c58
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host      vserver
path      path
state     type      /dev/  host      vserver
          node    adapter  LIF
-----
up        primary    sdi     host33    iscsi-lif-01a
up        primary    sdb     host34    iscsi-lif-02b
up        primary    sdc     host35    iscsi-lif-02a
up        primary    sdd     host36    iscsi-lif-01b

...

          ONTAP Path: svm1:fpsa_asa_ol8vm_01_oracle_slobdata_1
            LUN: 11
        LUN Size: 500g
        Product: cDOT
    Host Device: 3600a0980383234486724587338696c61
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host      vserver
path      path
state     type      /dev/  host      vserver
          node    adapter  LIF
-----
up        primary    sds     host33    iscsi-lif-01a
up        primary    sdj     host34    iscsi-lif-02b
up        primary    sdh     host35    iscsi-lif-02a
up        primary    sdk     host36    iscsi-lif-01b

...
```

**Note:** In the partial sanlun lun show -p output above, you can see the LUN information such as ONTAP path, LUN ID, LUN size, as well as the multipath information. The path type information indicates that all paths for the same LUN are primary, or active/optimized in SCSI Asymmetric Logical Unit Access (ALUA) terms.

## Enable persistent device naming for multipath devices

The Linux udev rules provide a mechanism for mapping the devices corresponding to redolog and slobdata LUNs to persistent device names. Having persistent device name makes it easier to identify devices for creating disk groups for Oracle database.

## 1. Use the sanlun tool to gather Host Device and ONTAP Path information for devices.

```
[admin@fpsa-asa-ol8vm-01 ~]$ sudo sanlun lun show -p

ONTAP Path: svm1:fpsa_asa_ol8vm_01_oracle_redolog_1
LUN: 1
LUN Size: 200g
Product: cDOT
Host Device: 3600a0980383234486724587338696c58
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host      vserver
path      path      /dev/      host      vserver
state     type     node     adapter   LIF
-----
up        primary  sdi      host33    iscsi-lif-01a
up        primary  sdb      host34    iscsi-lif-02b
up        primary  sdc      host35    iscsi-lif-02a
up        primary  sdd      host36    iscsi-lif-01b
...

```

Create udev rules to map Host Devices to device aliases for consistent device names.

```
[admin@fpsa-asa-ol8vm-01 rules.d]$ pwd
/etc/udev/rules.d

[admin@fpsa-asa-ol8vm-01 rules.d]$ cat 71-iscsi-redolog.rules
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696c58",
SYMLINK+="iscsiredolog1", OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696c59",
SYMLINK+="iscsiredolog2", OWNER="grid", GROUP="oinstall", MODE="0660"

```

**Note:** With the udev rule above, the LUN with ONTAP Path: svm1:fpsa\_asa\_ol8vm\_01\_oracle\_redolog\_1 and Host Device: 3600a0980383234486724587338696c58 will be mapped to device alias /dev/iscsiredolog1 by the udev rule.

**Note:** For easier management, we created three udev rules, 71-iscsi-redolog.rules, 72-iscsi-slobdata.rules, and 73-iscsi-ocrvote.rules for redolog, slobdata, and ocrvote devices, respectively.

**Note:** Adjust the OWNER and GROUP based on your Oracle deployment user and group information as needed. While the device alias links will still be owned by root:root, the dm devices will be updated to the specified OWNER and GROUP.

Update and apply the udev rules.

```
[admin@fpsa-asa-ol8vm-01 rules.d]$ sudo udevadm control --reload-rules
[admin@fpsa-asa-ol8vm-01 rules.d]$ sudo udevadm trigger --type=devices --action=change

```

Check for the created device aliases.

```
[admin@fpsa-asa-ol8vm-01 rules.d]$ sudo ls -l /dev | grep redolog
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsiredolog1 -> dm-3
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsiredolog2 -> dm-4

[admin@fpsa-asa-ol8vm-01 rules.d]$ sudo ls -l /dev/ | grep ocrvote
lrwxrwxrwx. 1 root root          5 Mar 16 11:47 iscsiocrvote1 -> dm-13
lrwxrwxrwx. 1 root root          5 Mar 16 11:47 iscsiocrvote2 -> dm-14

[admin@fpsa-asa-ol8vm-01 rules.d]$ sudo ls -l /dev | grep slobdata
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsislobdata1 -> dm-7
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsislobdata2 -> dm-5
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsislobdata3 -> dm-9
lrwxrwxrwx. 1 root root          5 Mar 16 11:47 iscsislobdata4 -> dm-11
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsislobdata5 -> dm-6
lrwxrwxrwx. 1 root root          4 Mar 16 11:47 iscsislobdata6 -> dm-8
lrwxrwxrwx. 1 root root          5 Mar 16 11:47 iscsislobdata7 -> dm-10
lrwxrwxrwx. 1 root root          5 Mar 16 11:47 iscsislobdata8 -> dm-12

```

## Solution verification

The FlexPod SAN solution with Cisco UCS X-Series Direct and NetApp ASA uses configurations already supported by NetApp Interoperability Matrix Tool (IMT), Cisco Hardware and Software Compatibility List (HCL), and Broadcom Compatible Guide (BCG). After the solution is deployed, a variety of test cases are conducted for solution verifications. The following sections provides information on ecosystem interoperability validation, Microsoft SQL and Oracle RAC database testing, solution availability and infrastructure resiliency testing, and some life-cycle management tasks.

## Ecosystem interoperability validation

FlexPod solutions are built on top of the ecosystem interoperability foundations established by the NetApp interoperability team and partners like Cisco and VMware to validate ecosystem interoperability and provide interoperability information to our joint customers. The following sections highlight some of the tools that can be used to check for supported configurations.

### NetApp Interoperability Matrix Tool (IMT)

To check for SAN interoperability information for FlexPod solution in NetApp Interoperability Matrix Tool (IMT): <https://support.netapp.com/matrix> , follow the steps below.

1. Go to [NetApp IMT site](https://support.netapp.com/matrix).
2. Use the Advanced Search tool to search for FlexPod.

**NetApp Interoperability Matrix Tool**

HOME ADMIN SUBSCRIPTIONS NOTIFICATIONS SAVED SEARCH CAN'T FIND CONFIG? TOOLBOX ▼ REPORTS ▼ PREFERENCE ▼ HELP ▼ TAKE A TOUR New JYH-SHING ▼

### Advanced Search

FlexPod

Solutions - Top matches [Show All Matches](#)

- FlexPod MAX Data
- FlexPod SAN Simplified
- FlexPod Switch

Keywords - Top matches

- citrix6.2\_flexpod

**Filters**

- ☐ Do not show Limited Support components
- ☐ Do not show EOVS components

**Search Assistant**

You may start by

- Adding criteria from the Search box
- Selecting components from these commonly used categories

[ONTAP OS](#) [Protocol](#) [Host OS](#)

[SnapManager](#) [Host-Multipath](#) [SnapDrive](#)

**Search Criteria** [Clear All](#)

None Selected

3. Highlight the FlexPod SAN Simplified solution and click Add next to it to add FlexPod SAN Simplified solution to the search criteria.

### Advanced Search

FlexPod

Solutions - Top matches [Show All Matches](#)

- FlexPod MAX Data
- FlexPod SAN Simplified [Add](#)
- FlexPod Switch

Keywords - Top matches

- citrix6.2\_flexpod

4. Click on Next >> Refine Search Criteria link in the middle to bring up the search refinement dialog.

NetApp.com NetApp Support Contact NetApp

## NetApp Interoperability Matrix Tool

HOME ADMIN SUBSCRIPTIONS NOTIFICATIONS SAVED SEARCH CAN'T FIND CONFIG? TOOLBOX ▾ REPORTS ▾ PREFERENCE ▾ HELP ▾ TAKE A TOUR New JYH-SHING ▾

### Advanced Search

✕
Q

**Solutions - Top matches** Show All Matches

- FlexPod MAX Data
- ✓ FlexPod SAN Simplified
- FlexPod Switch

**Keywords - Top matches**

- citrix6.2\_flexpod

**Filters**

☐ Do not show Limited Support components

☐ Do not show EOVS components

**Search Assistant**

You have chosen the following

FlexPod SAN Simplified - 319 results found

[Next >> Refine Search Criteria](#)  
[Skip To Supported Configurations](#)

**Search Criteria** Clear All

**Solution**

FlexPod SAN Simplified

- In the Refine Search Criteria window, select the desired ONTAP version, storage protocol, and host operating system for the FlexPod solution. Click the arrow on the right to view the supported configurations.

NetApp.com NetApp Support Contact NetApp

## NetApp Interoperability Matrix Tool

HOME ADMIN SUBSCRIPTIONS NOTIFICATIONS SAVED SEARCH CAN'T FIND CONFIG? TOOLBOX ▾ REPORTS ▾ PREFERENCE ▾ HELP ▾ TAKE A TOUR New JYH-SHING ▾

**Search Criteria** Include EOVS and LS Dynamically Change Apply Changes Save Criteria Change Criteria

**Refine Search Criteria** Sort Policies & Guidelines Generate URL Filters Clear All

**FlexPod SAN Simplified (1)**

**ONTAP OS (NETAPP OS)**

Search Components

- ☐ ONTAP 9.10.1
- ☐ ONTAP 9.11.1
- ☐ ONTAP 9.12.1
- ☐ ONTAP 9.13.1
- ☐ ONTAP 9.14.1
- ☐ ONTAP 9.15.1
- ✓ ☒ ONTAP 9.16.1
- ☐ ONTAP 9.9.1

**Protocol**

Search Components

- ✓ ☒ iSCSI
- ☐ FCP
- ☐ FCoE
- ☐ NVMe/FC
- ☐ NVMe/TCP

**Host OS (OS)**

Search Components

- ✓ ☒ VMware ESXi 8.0 (Update 3)
- ☐ Canonical Ubuntu Server 20.04
- ☐ Canonical Ubuntu Server 22.04
- ☐ Canonical Ubuntu Server 24.04
- ☐ Citrix Hypervisor 8.2
- ☐ Microsoft Windows Server 2016
- ☐ Microsoft Windows Server 2016 Hyper-V
- ☐ Microsoft Windows Server 2019

➤
View Supported Configurations

- The supported configuration is shown along with notes related to the supported configuration.

NetApp.com NetApp Support Contact NetApp

## NetApp Interoperability Matrix Tool

HOME ADMIN SUBSCRIPTIONS NOTIFICATIONS SAVED SEARCH CAN'T FIND CONFIG? TOOLBOX ▾ REPORTS ▾ PREFERENCE ▾ HELP ▾ TAKE A TOUR New JYH-SHING ▾

[Configuration Component Compare](#)
[Policies & Guidelines](#)
[Filters](#)
[Export](#)

**FlexPod SAN Simplified (1)**

⏪
⏩
Page 1 of 1
⏪
⏩
Displaying configurations 1 - 1 of 1

➤
0240626-103936182 (Supported)
☆ What If

Details

**ONTAP OS (NETAPP OS)**

ONTAP 9.16.1

ASA (2)

FabricPool (1)

OS (5)

SM active sync (4)

**Protocol**

iSCSI

Generic (4)

**Host OS (OS)**

VMware ESXi 8.0 (Update 3)

HBA, Driver and Firmware (1)

MCC (1)

ONTAP (1)

**Note:** There is a note for HBA, Driver and Firmware which indicates that NetApp supports all OEM and third-party adapters from Broadcom, Cisco, and Marvell.

**10420:**

NetApp supports all OEM and third-party (rebranded) Fibre Channel HBA/Mezzanine adapters that are based on models from Broadcom, Cisco and Marvell.

NetApp recommends using the latest Driver/Firmware that is listed in the VMware Compatibility Guide (VCG). <https://www.vmware.com/resources/compatibility/search.php?deviceCategory=io>

**Note:** For Cisco VIC support details, please refer to the section below on checking Cisco UCS hardware and software compatibility list.

## Cisco UCS hardware and software compatibility (HCL)

To check for the compatibility of Cisco UCS with operating systems and storage, go to the Cisco UCS hardware and software compatibility list (HCL): <https://ucshcltool.cloudapps.cisco.com/public/>. Follow the steps below to search for interoperability.

1. Go to [Cisco HCL site](#).
2. Select the category of searches by picking the category from the Search by drop-down list. Here we are searching by Servers.
3. Select the appropriate information from the drop-down lists for the search. Click Search.

The screenshot shows the Cisco UCS Hardware and Software Compatibility (HCL) search tool. The browser address bar displays the URL <https://ucshcltool.cloudapps.cisco.com/public/>. The page title is "UCS Hardware and Software Compatibility". The search criteria are as follows:

- Search By :** Servers (selected from a dropdown menu)
- Server Type :** X-Series (selected from a dropdown menu)
- Server Model :** Cisco UCS X215c M8 (selected from a dropdown menu)
- Processor Version :** 4th Gen AMD EPYC 9004 Series Processors (selected from a dropdown menu)
- Operating System :** VMware (selected from a dropdown menu)
- Operating System Version :** ESXi 8.0 U3 (selected from a dropdown menu)

At the bottom right, there are two buttons: "Reset All" and "Search".

4. Refine the search results by selecting the product category, UCS server firmware, adapters, and storage.



Search Type	Servers
Server Type	X-Series (UCSM)
Server Model	Cisco UCS X215c M8
Processor Version	4th Gen AMD EPYC 9004 Series Processors
Operating System	VMware
Operating System Version	ESXi 8.0 U3

Advisories : None

Search Results :

Refine by  
[Select All](#) | [Clear All](#)

Product Category

- ☒ Adapters
- ☐ SSD
- ☒ Storage
- ☐ Switch

UCS Server Firmware

- ☒ 5.3(0)
- ☐ 4.3(5)

Expand All Collapse All

Export Excel Export PDF

Component	Details	Documents
5.3(0) last published 2025-04-09 ( <a href="#">change log</a> )	Firmware Bundle <a href="#">Driver ISO</a>	<a href="#">View Notes</a> <a href="#">Release Notes</a> <a href="#">Install &amp; Upgrade Guides</a>
Adapters: CNA × Storage: Array ×		
Adapters		
> CNA		
Storage		
> Array		

**Note:** Here we selected only adapters and storage under the product category. For the UCS server firmware, we only selected 5.3(0) series. For the Adapters and Storage in the middle, we picked CNA and Array.

- Click on CNA under Adapters in the middle to expand the information. Look for the adapter which you are using to see a list of drivers.

Component	Details	Documents
5.3(0) last published 2025-04-09 ( <a href="#">change log</a> )	Firmware Bundle <a href="#">Driver ISO</a>	<a href="#">View Notes</a> <a href="#">Release Notes</a> <a href="#">Install &amp; Upgrade Guides</a>
Adapters: CNA × Storage: Array ×		
Adapters		
> CNA		
Cisco UCSX-ME-V5Q50G: Cisco UCS VIC 15422	Firmware Version: 5.3(4)-48 Driver Version: 1.0.9.0-1OEM.800.1.0.20613240 nenic-ens ⓘ Adapter BIOS: <none> Notes: <none>	
Cisco UCSX-ME-V5Q50G: Cisco UCS VIC 15422	Firmware Version: 5.3(4)-48 Driver Version: 2.0.15.0-1OEM.800.1.0.20613240 nenic ⓘ Adapter BIOS: <none> Notes: <none>	
Cisco UCSX-ME-V5Q50G: Cisco UCS VIC 15422	Firmware Version: 5.3(4)-48 Driver Version: 5.0.0.45-1OEM.803.0.0.24022510 nfnic ⓘ Adapter BIOS: <none> Notes: <none>	
Cisco UCSX-ML-V5D200GV2: Cisco UCS VIC 15230	Firmware Version: 5.3(4)-48 Driver Version: 1.0.9.0-1OEM.800.1.0.20613240 nenic-ens ⓘ Adapter BIOS: <none> Notes: <none>	
Cisco UCSX-ML-V5D200GV2: Cisco UCS VIC 15230	Firmware Version: 5.3(4)-48 Driver Version: 2.0.15.0-1OEM.800.1.0.20613240 nenic ⓘ Adapter BIOS: <none> Notes: <none>	
Cisco UCSX-ML-V5D200GV2: Cisco UCS VIC 15230	Firmware Version: 5.3(4)-48 Driver Version: 5.0.0.45-1OEM.803.0.0.24022510 nfnic ⓘ Adapter BIOS: <none> Notes: <none>	

**Note:** For the VIC 15230 (UCSX-ML-V5D200GV2) adapter which we are using, the supported Ethernet (nenic) firmware / driver versions are: 5.3(4)-48 / 2.0.15.0-1OEM 800.1.0.20613240. (See the highlighted cell in the screenshot above.

- Click on Array under Storage to expand on the storage array support and scroll down to look for the NetApp ASA A50 array.

NetApp ASA A30	Firmware Version	Refer to storage vendor's support list, <a href="https://mysupport.netapp.com/">https://mysupport.netapp.com/</a>
	Driver Version	<none>
	Adapter BIOS	<none>
	Notes	1,2,7,9,21,23,25,27,28
NetApp ASA A50	Firmware Version	Refer to storage vendor's support list, <a href="https://mysupport.netapp.com/">https://mysupport.netapp.com/</a>
	Driver Version	<none>
	Adapter BIOS	<none>
	Notes	1,2,7,9,21,23,25,27,28

**Note:** As indicated in the output above, refer to NetApp IMT for information on the ONTAP version support.

- Click on the notes link to see a list of notes associated with the configuration.

**Note:** For example, the partial note output screenshot below indicates that both iSCSI SAN boot and iSCSI data LUN are supported for our direct-attached storage configuration, where the ASA A50 is attached to the fabric interconnects' appliance ports.

27 : iSCSI Boot is supported. Both Appliance and Uplink ports are supported unless otherwise noted

28 : iSCSI data LUN access if the storage array is directly connected to UCS Fabric Interconnects using Appliance Ports is supported

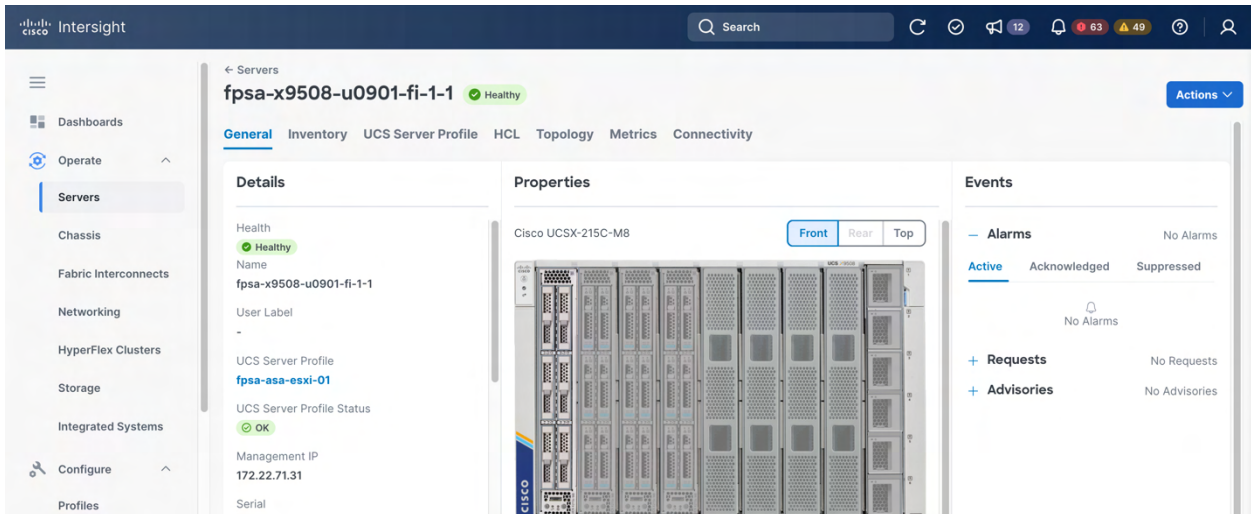
## Cisco Intersight HCL validation

Cisco Intersight can perform server validation checks against the Cisco UCS hardware and software compatibility list. To perform server validation checks against Cisco HCL, follow the steps below.

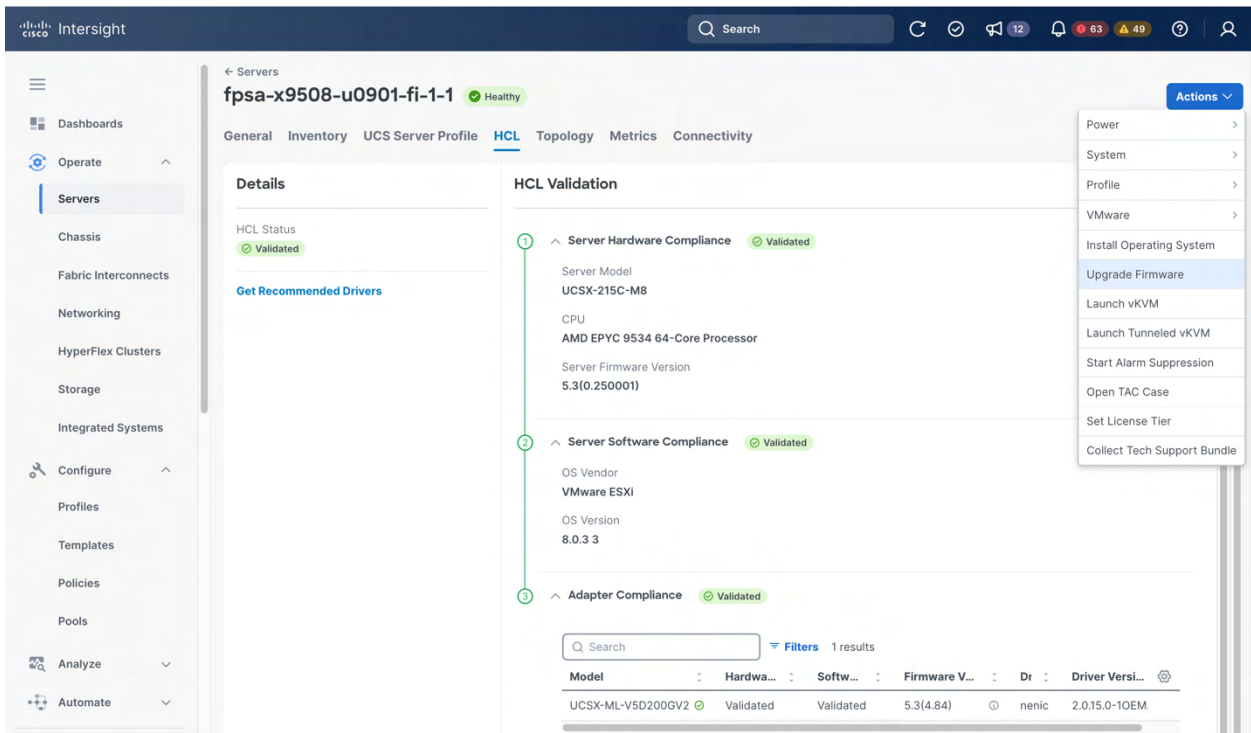
- Login to the Cisco Intersight account under which the Cisco UCS X-Series is being managed.
- Navigate to Operate > Servers and search for the servers that you are using by applying a filter to narrow down the server list if needed.

The screenshot shows the Cisco Intersight interface. On the left is a navigation menu with options like Dashboards, Operate, Servers, Chassis, Fabric Interconnects, Networking, HyperFlex Clusters, Storage, Integrated Systems, and Configure. The main area is titled 'Servers' and shows a search bar with 'fpga-asal' and 'Filters 3 results'. Below this are several summary cards: Health (3 Healthy), Power (On 3), HCL Status (Incomplete 1, Validated 2), Bundle Version (5.3(0.250001) 3), Utility Storage (No 3), and Firmware Version (5.3(0.250001) 3). At the bottom is a table with columns: Name, Health, Model, Mem..., UCS Domain, Server Profile, and Bundle Vers... The table lists three servers, all with a 'Healthy' status.

- Select a server by clicking on the server link under the Name column.

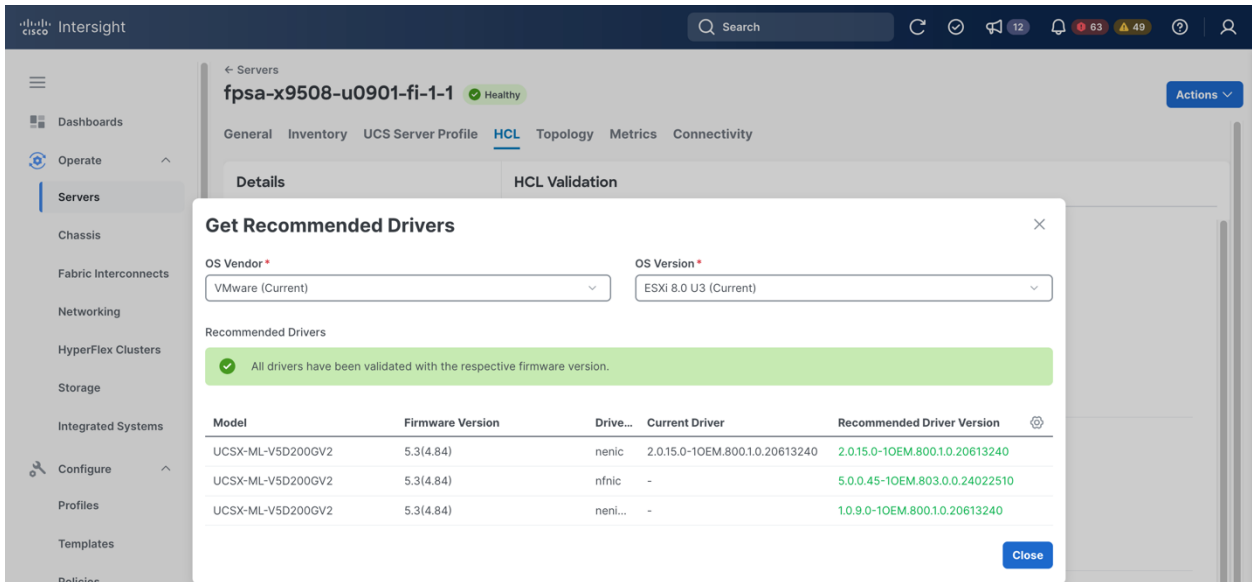


4. Click on the HCL tab in the center pane.
5. Intersight validates the configuration against Cisco HCL and summarizes the compliance statuses for server hardware, server software, and VIC adapter.



**Note:** Under the Actions menu, there are many options available for server management, including server firmware upgrade.

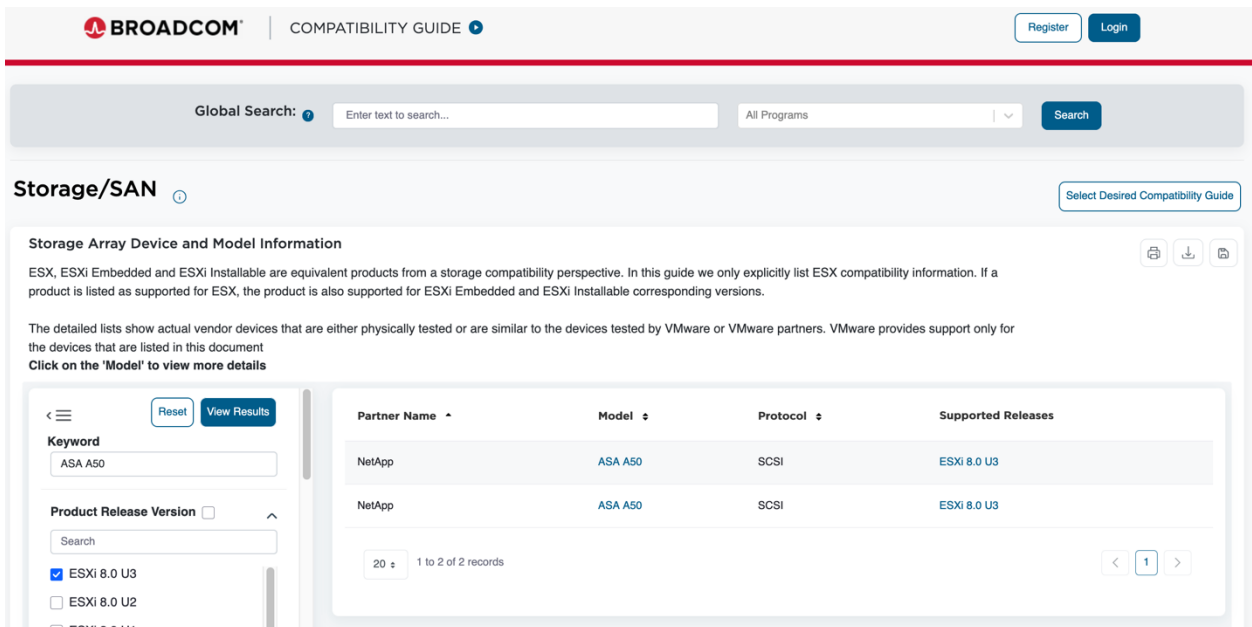
6. If you click on the Get Recommended Drivers link under Details, you can see additional details for the recommended driver versions. In addition, you can also select different OS Vendor and OS Version combinations from the drop-down list to see the recommended drivers for the selected configuration.



## Broadcom compatibility guide (BCG)

To check for the compatibility of the VMware solution ecosystem, go to Broadcom Compatibility Guide (BCG) website: <https://compatibilityguide.broadcom.com/>. Follow the steps below to search for compatibility of NetApp storage and Cisco VIC with VMware configurations.

1. Go to [Broadcom compatibility guide](https://compatibilityguide.broadcom.com/).
2. Select Storage/SAN under the Storage/Availability category.
3. On the Storage/SAN page, select appropriate configuration information. Here we specified ASA A50 in keyword field, ESXi 8.0 U3 product version, NetApp partner, SCSI protocol, iSCSI transport, and then clicked View Results.



4. Select a result to view the supported configuration details. In the example below, it shows both FC and iSCSI support as well as the supported configuration details.

Model Release Details							
VMware Product Name:		Attention:					
ESXi 8.0 U3		Storage partners using ESX 4.0 or later may recommend VMW_PSP_RR for path failover policy for certain storage array models. If desired, contact the storage array manufacturer for recommendation and instruction to set VMW_PSP_RR appropriately.					
Firmware Version	Transport_Type	Test Configuration	MPP Plugin	SATP Plugin	PSP Plugin	Note	Features
ONTAP 9.16.1	FC	Switch	NMP	VMW_SATP_ALUA	VMW_PSP_RR	-	<a href="#">View</a>
ONTAP 9.16.1	FC	Switch	NMP	VMW_SATP_DEFAULT_AA	VMW_PSP_RR	-	<a href="#">View</a>
ONTAP 9.16.1	iSCSI	SW iSCSI	NMP	VMW_SATP_ALUA	VMW_PSP_RR	-	<a href="#">View</a>
ONTAP 9.16.1	FC	Switch	NMP	ft-satp-eternus-ahx/1.0.0	VMW_PSP_RR	-	<a href="#">View</a>

**Note:** The iSCSI row above indicates ESXi 8.0U3 support for ASA A50 running ONTAP 9.16.1 with iSCSI transport, iSCSI software initiator, Native Multipathing Plugin (NMP), VMW\_SATP\_ALUA ALUA Storage Array Type Plugin (SATP), and VMW\_PSP\_RR round-robin Path Selection Policy Plugin (PSP).

- Click View at the end of a supported row to view the supported feature information.

ONTAP 9.16.1	iSCSI	SW iSCSI	NMP	VMW_SATP_ALUA	VMW_PSP_RR	-	<a href="#">Hide</a>
Feature Category	Features		Plugin	Plugin_Version			
VAAI-Block	UNMAP		-	-			
	Full Copy		-	-			
	Block Zero		-	-			
	HW Assisted Locking		-	-			
	Thin Provisioning Stun		-	-			

**Note:** The feature category VAAI-Block indicates that the ASA A50 supports the VMware vSphere API for Array Integration (VAAI) and the supported commands are listed under the Features column.

Go back to [Broadcom compatibility guide](#) .

Select Platform & Compute category and select IO Devices.

- On the IO Devices page, select appropriate configuration information. Here we specified VIC 15230 in keyword field, ESXi 8.0 U3 product version, Cisco brand, Network device type, and clicked View Results.

COMPATIBILITY GUIDE

Register
Login

Global Search:
All Programs
Search

IO Devices

Select Desired Compatibility Guide

Click here to Read Important Support Information

**I/O Device and Model Information**

The detailed lists show actual vendor devices that are either physically tested or are similar to the devices tested by VMware or VMware partners. VMware provides support only for the devices that are listed in this document.

Click on the 'Model' to view more details

Reset
View Results

Keyword

Product Release Version

☒ ESXi 8.0 U3
☐ ESXi 8.0 U2
☐ ESXi 8.0 U1

Brand Name	Model	Device Type	Supported Releases
Cisco	UCS VIC 15230	Network	ESXi 8.0 U3

20
1 to 1 of 1 records

Click on the UCS VIC 15230 model link to see details.

COMPATIBILITY GUIDE

Register
Login

Back to Search Page

**Model Details**

<b>Model:</b> UCS VIC 15230  <b>Number of Ports:</b> 2  <b>SVID:</b> 1137	<b>Device Type:</b> Network  <b>VID:</b> 1137  <b>SSID:</b> 02df	<b>Brand Name:</b> Cisco  <b>DID:</b> 0043  <b>Notes:</b> Firmware versions listed are the minimum supported versions. Refer to <a href="http://kb.vmware.com/kb/2030818">http://kb.vmware.com/kb/2030818</a> for additional information on other supported driver and firmware combinations
------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**Model Release Details**

ESXi Versions:

Release	Driver Name	Driver Version	Firmware Version	Additional Firmware Version	Type	Notes	Features
ESXi 8.0 U3	nenic	2.0.4.0	5.2(3)	-	Partner Async, native	Visit <a href="https://knowledge.broadcom.com/external/article?articleId=366755">https://knowledge.broadcom.com/external/article?articleId=366755</a> for driver download instructions Copy the component version here: 2.0.4.0 Go to <a href="https://knowledge.broadcom.com/external/article?articleId=366755">https://knowledge.broadcom.com/external/article?articleId=366755</a> for driver download instructions	View
ESXi 8.0 U3	nenic	2.0.15.0-1OEM	5.3(4)	-	Partner Async, native	Visit <a href="https://knowledge.broadcom.com/external/article?articleId=366755">https://knowledge.broadcom.com/external/article?articleId=366755</a> for driver download instructions Copy the component version here: 2.0.15.0-1OEM Go to <a href="https://knowledge.broadcom.com/external/article?articleId=366755">https://knowledge.broadcom.com/external/article?articleId=366755</a> for driver download instructions	View

**Note:** The above is a partial screenshot of the supported list of driver/firmware versions, which includes the 2.0.15.0-1OEM driver and 5.3(4) firmware that we are using.

**Note:** Based on [Broadcom KB 315219](#), I/O adapters in BCG list drivers at the time of certification. Customers should refer to OEM vendor matrix for the recommendations on latest combination of supported driver and firmware.



## Microsoft SQL database testing

Be sure that the Microsoft SQL database server and NetApp Host Utilities for Windows had already been installed. In addition, LUNs needed for the SQL server testing had been created in storage and in-guest iSCSI initiator had been configured with proper multipathing support for the LUNs.

For SQL database testing, we deployed HammerDB, which is a free open-source database benchmarking application that supports various databases servers, including Microsoft SQL Server, IBM Db2, Oracle, PostgreSQL, MariaDB, and MySQL.

For HammerDB testing of the SQL database workload, we created two striped volumes for databases files using the following layout:

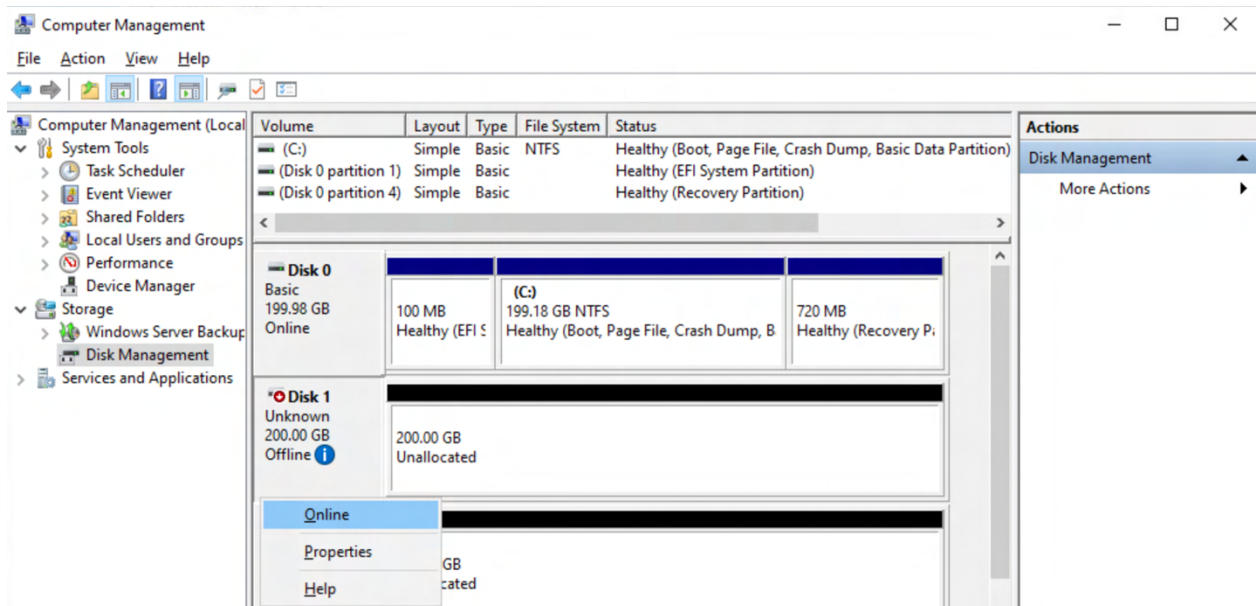
- 4 x 500GB LUNs are used for SQL data
- 2 x 200GB LUNs are used for SQL log

## Create striped volumes

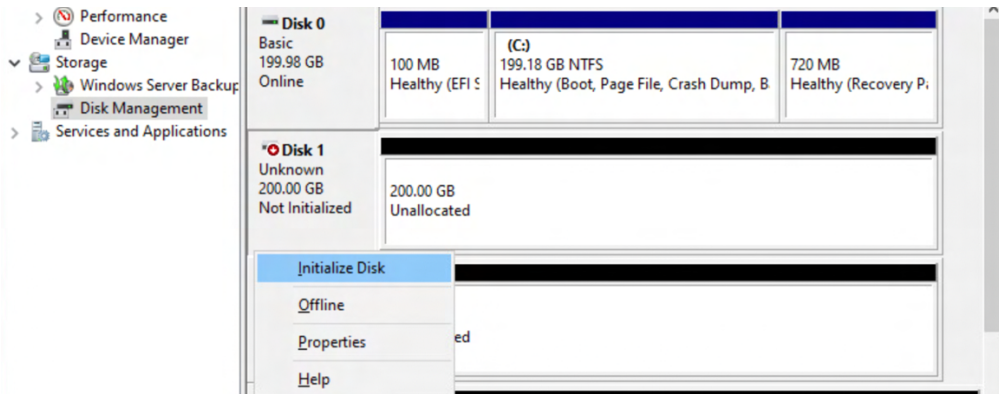
After completing the Microsoft SQL server installation and direct iSCSI LUN access from the virtual machine with multipathing access to storage, you can use the steps below to create two striped volumes, one for data and one for log, for the SQL database.

**Note:** Striped volumes allow disk IO to be distributed to multiple disks to improve IO performance for a volume.

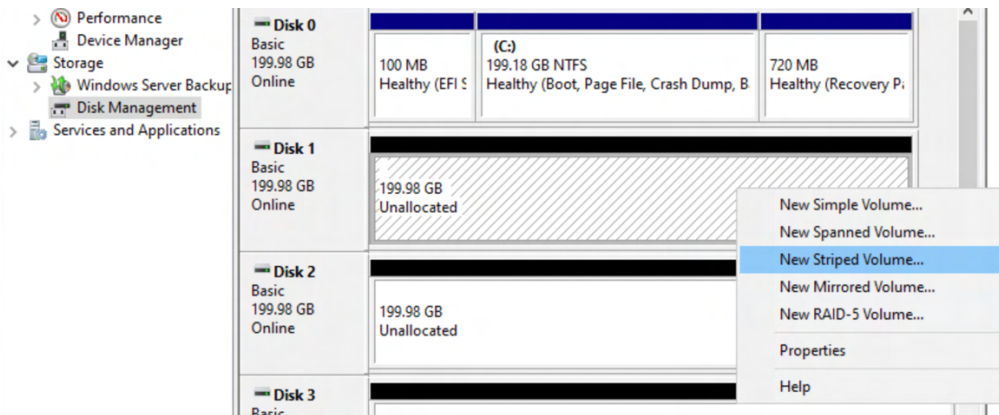
1. Login to Windows.
2. Open Server Manager.
3. Select Computer Management from the Tools menu.
4. Click on Disk Management under Storage.
5. Right-click on the first 200G log disk, Disk 1, and select Online to bring it online.



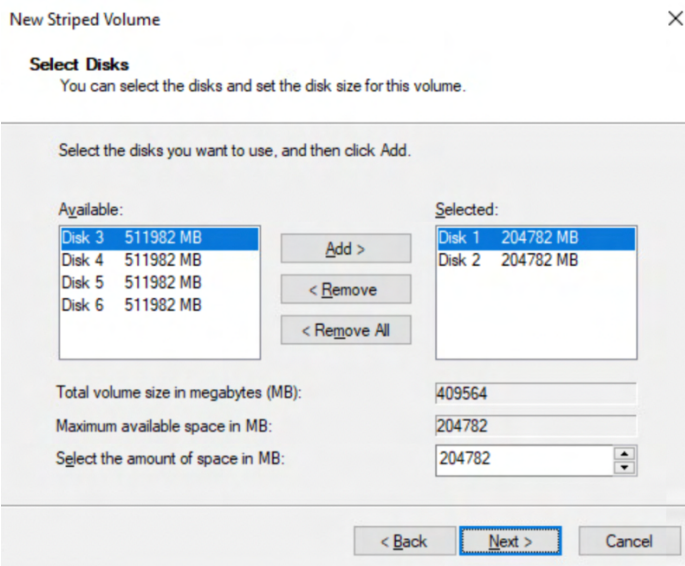
6. Right-click on the disks again and select initialize.



7. Click OK on the Initialize Disk dialog to initialize.
8. Follow similar steps above to online and initialize the remaining disks created for SQL data and log.
9. Right-click on graph indicating Unallocated area for Disk 1 and select New Striped Volume.



10. Click Next on the new striped volume creation dialog.
11. In the Add disks screen, select Disk 2, click Add to add Disk 2 to the striped volume, and click Next.



12. Assign Drive Letter and click Next.
13. Provide a Volume Label, select Perform a quick format, and click Next.

New Striped Volume ✕

**Format Volume**  
To store data on this volume, you must format it first.

Choose whether you want to format this volume, and if so, what settings you want to use.

☐ Do not format this volume  
☒ Format this volume with the following settings:

File system: NTFS  
 Allocation unit size: Default  
 Volume label: SQL Log


☒ Perform a quick format  
☐ Enable file and folder compression

< Back Next > Cancel

14. Click Finish to initiate the volume creation.

15. Review the Disk Management warning and click Yes to proceed.

Disk Management ✕

 The operation you selected will convert the selected basic disk(s) to dynamic disk(s). If you convert the disk(s) to dynamic, you will not be able to start installed operating systems from any volume on the disk(s) (except the current boot volume). Are you sure you want to continue?

Yes No

16. After the stripe volume has been created, the two disks properly reflect the same drive letter and volume label.

<b>Disk 1</b> Dynamic 199.98 GB Online	<b>SQL Log (E:)</b> 199.98 GB NTFS Healthy
<b>Disk 2</b> Dynamic 199.98 GB Online	<b>SQL Log (E:)</b> 199.98 GB NTFS Healthy

17. Follow steps 9 - 16 to create a new striped volume for SQL data and select the SQL data disks.

**New Striped Volume** [X]

**Select Disks**  
You can select the disks and set the disk size for this volume.

Select the disks you want to use, and then click Add.

Available:	Selected:
	Disk 3 511982 MB
	Disk 4 511982 MB
	Disk 5 511982 MB
	Disk 6 511982 MB

Buttons: Add > < Remove < Remove All

Total volume size in megabytes (MB): 2047928

Maximum available space in MB: 511982

Select the amount of space in MB: 511982

Navigation: < Back Next > Cancel

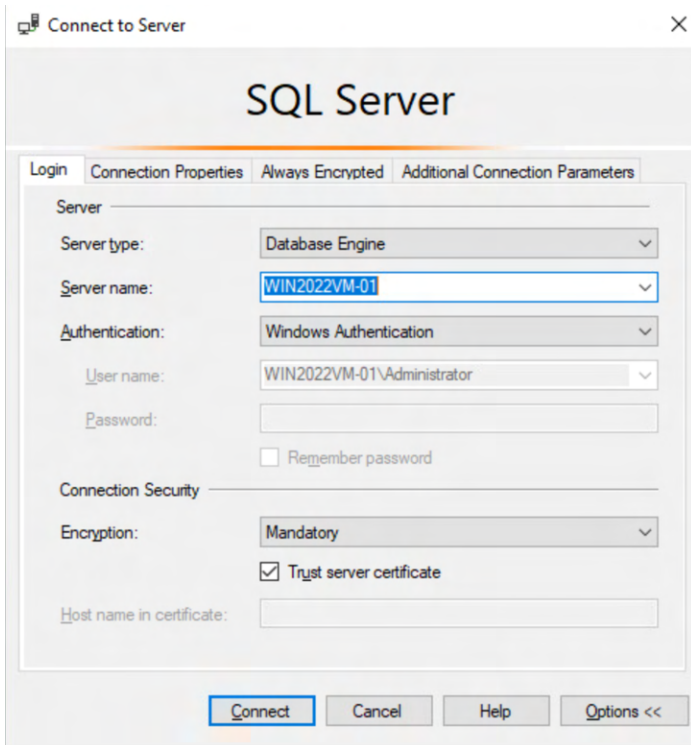
<b>Disk 3</b> Dynamic 499.98 GB Online	<b>SQL Data (F:)</b> 499.98 GB NTFS Healthy
<b>Disk 4</b> Dynamic 499.98 GB Online	<b>SQL Data (F:)</b> 499.98 GB NTFS Healthy
<b>Disk 5</b> Dynamic 499.98 GB Online	<b>SQL Data (F:)</b> 499.98 GB NTFS Healthy
<b>Disk 6</b> Dynamic 499.98 GB Online	<b>SQL Data (F:)</b> 499.98 GB NTFS Healthy

**Note:** Depending on the size of the disks, it might take some time for the disks to be formatted and the striped volume to be created.

## Create test SQL database with SQL Server Management Studio

1. Start SQL Server Management Studio.

Check Trust server certificate and click Connect.



Right-click the Database folder in SQL Server Manager Studio and then choose New Database.

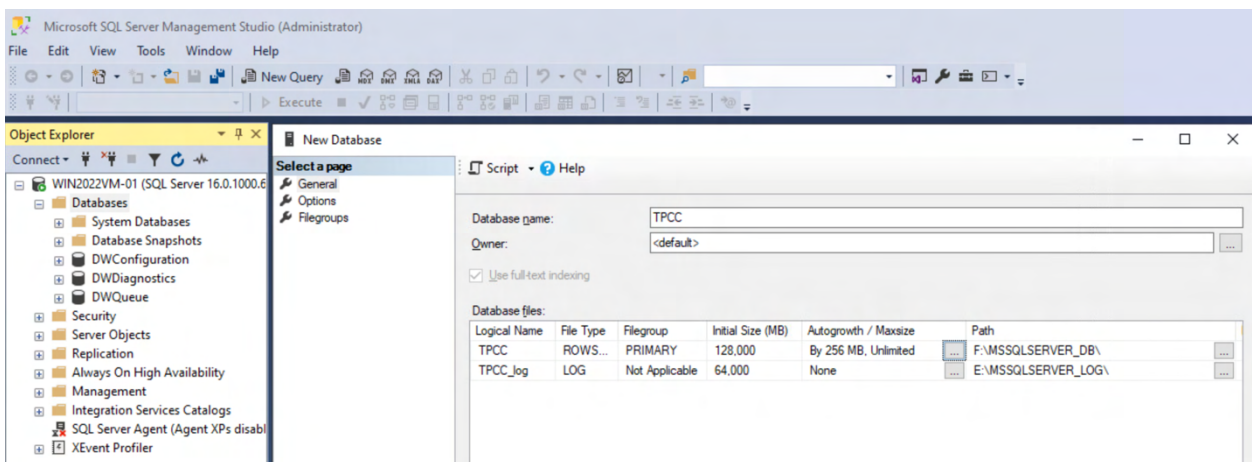
Name the new database TPCC.

Set the initial size of the data file to 128,000 MB and the log file to 64,000 MB.

Adjust the Autogrowth limits by clicking the ellipsis buttons and modify the value in the pop-up dialog.

Set the data file to grow by 256MB to unlimited size.

Set the log file to disable auto-growth and click OK.



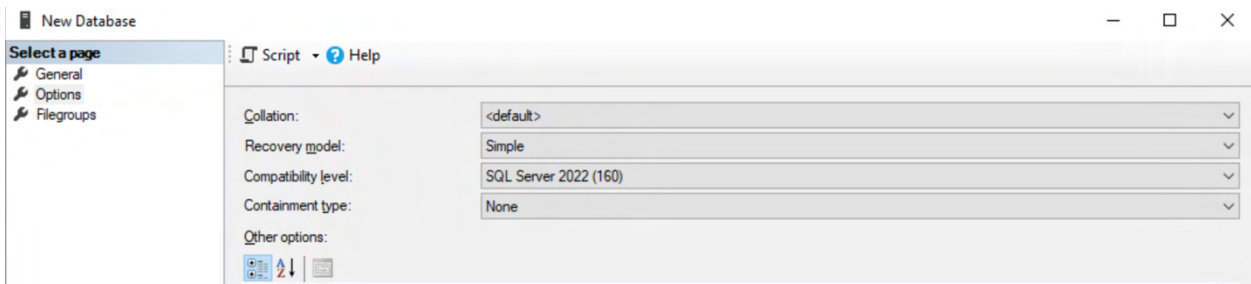
**Note:** The database file path should already be set to the locations on the striped volume for SQL server which was specified during the SQL server installation. Update the paths if required.

Choose **Options** page in the **New Database** dialog.

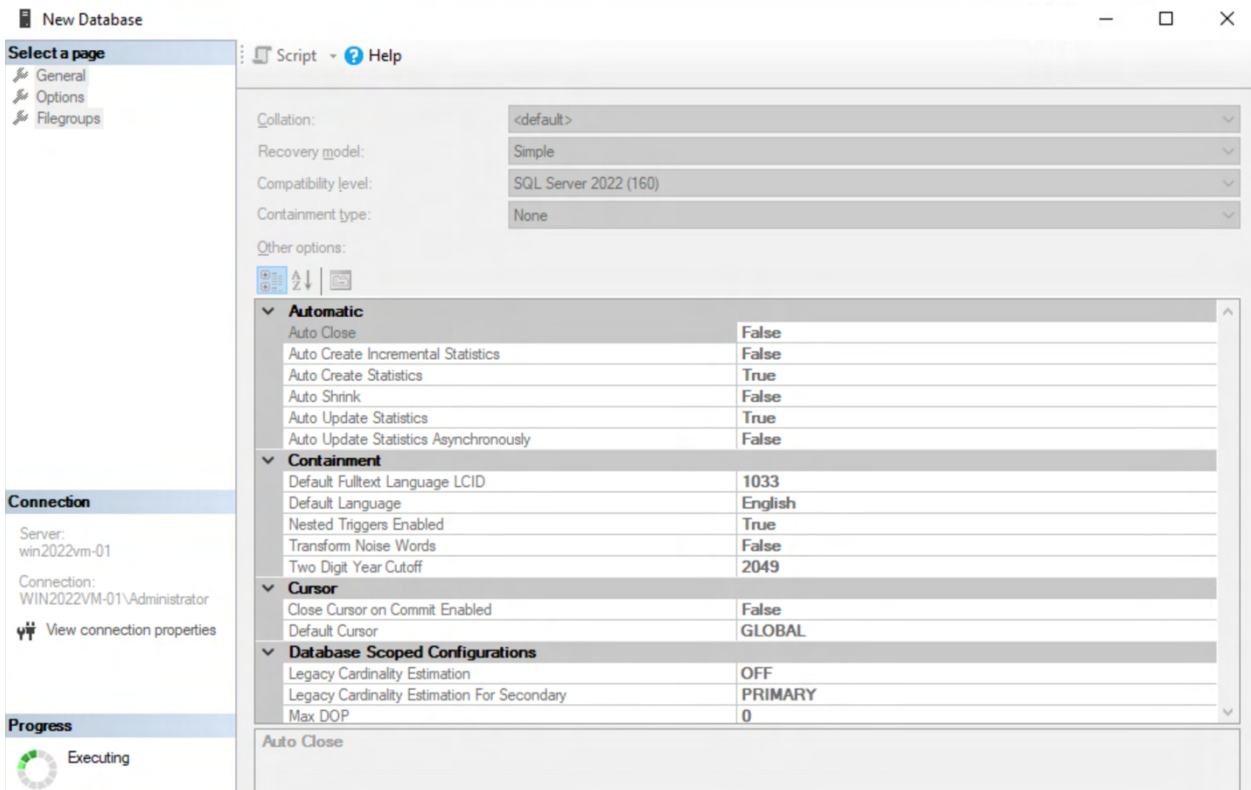
Set **Compatibility level** to **SQL Server 2022 (160)**.

Set the **Recovery model** to **Simple**, so that the loading doesn't fill up the transaction logs.





Click OK to create the TPCC database, which can take a few minutes to complete.



**Note:** You will see the spinning wheel under the Progress section in the lower left-hand corner of the New Database dialog.

## Install HammerDB database IO tool

HammerDB is a database IO generation tool that you can download to exercise the Microsoft SQL server database configuration. To install HammerDB tool, follow the steps below.

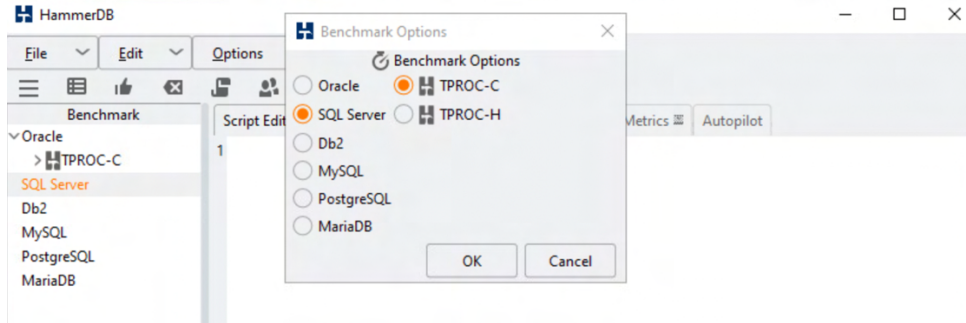
1. Download the HammerDB Windows 64-bit installer from the HammerDB download site below.  
<https://www.hammerdb.com/download.html>
2. Start the downloaded HammerDB setup wizard.
3. Click I accept the agreement on the License Agreement screen and click Next.
4. Specify installation directory and click Next.
5. Click Next again to begin installing HammerDB on the machine.
6. Click Finish after the installation is completed.



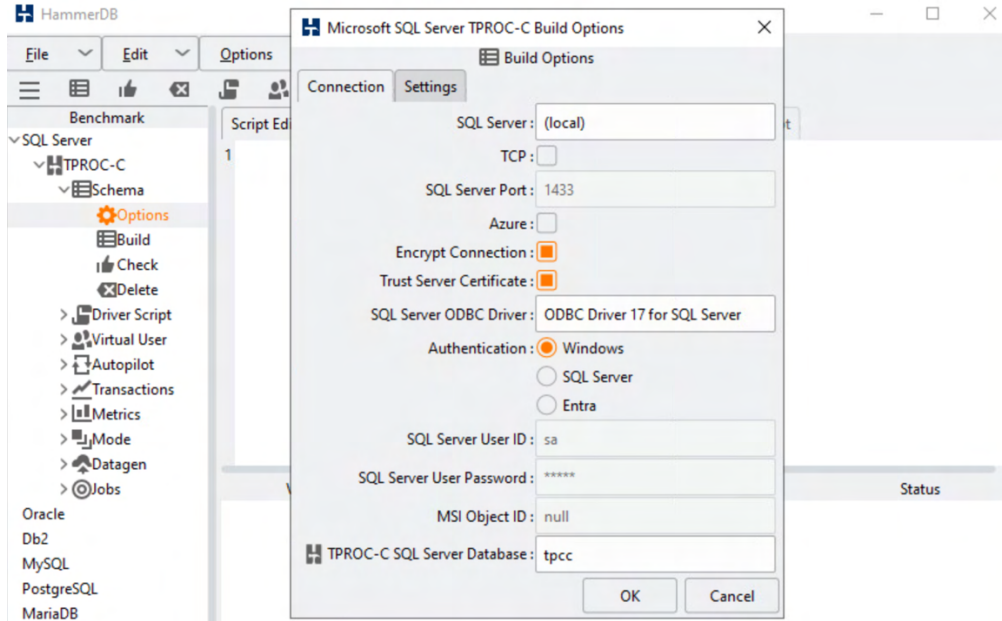
## Create the connection and schema

1. Start HammerDB application.
2. Double-click SQL Server in the Benchmark panel.

Select TPROC-C from the HammerDB site and click OK and then click OK again to confirm.



Expand TPROC-C, expand Schema, and double-click Options.



Update settings in the Connection tab as needed.

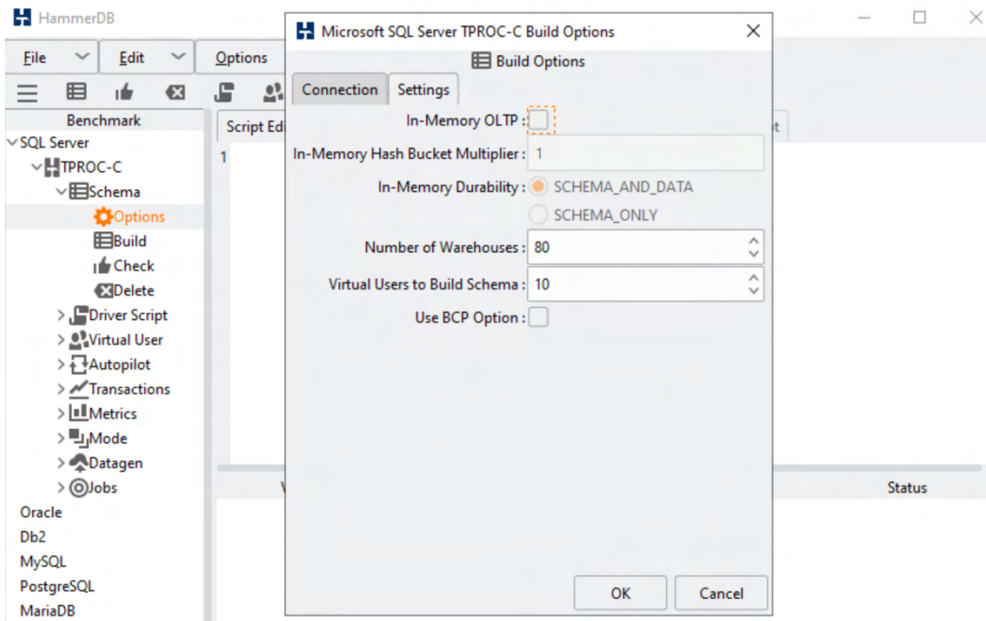
Click Settings tab, configure Number of Warehouse.

**Note:** Some guidelines suggest 10 to 100 warehouses per CPU. For this validation, this value is set to 10 times the number of cores in the VM, e.g. 80 for an 8-core instance.

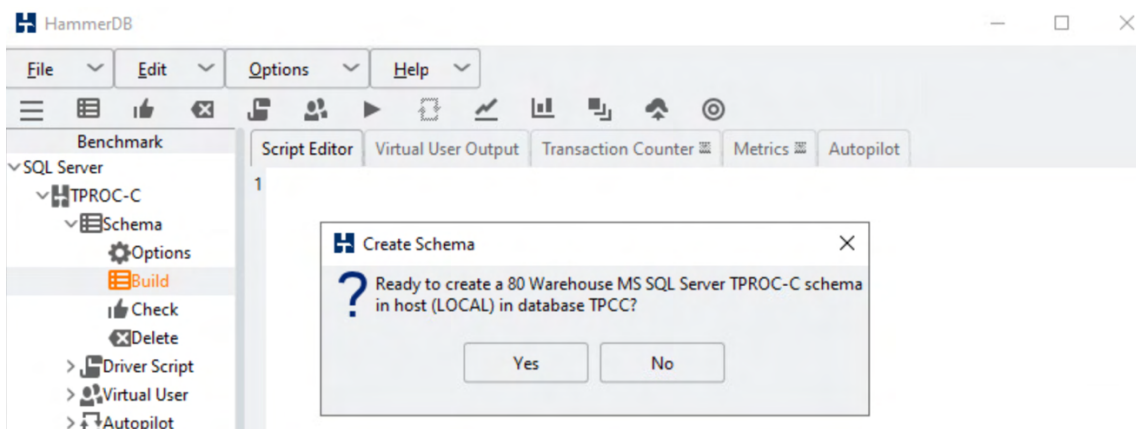
Configure Virtual Users to Build Schema, e.g. 10 for an 8-core instance.

**Note:** Choose a number that is between 1- and 2-times the client vCPU count.

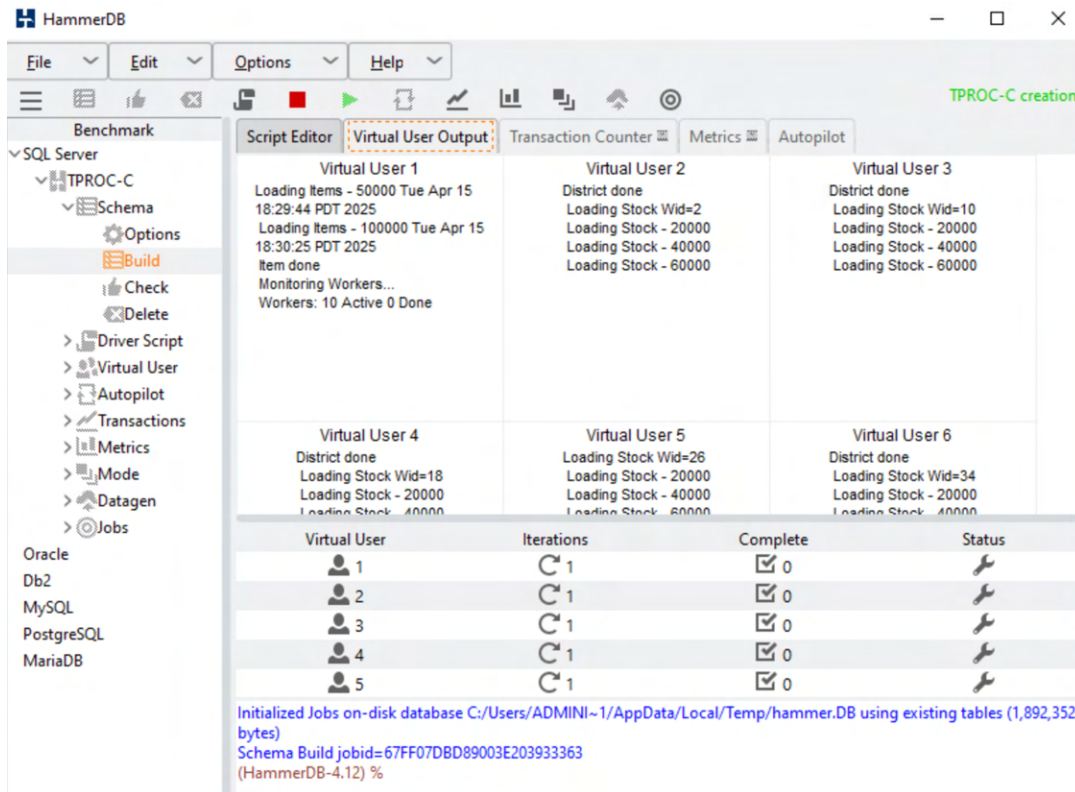
Clear the Use BPC Option and click OK.



Double-click Build under the Schema section and then click Yes to create the specified warehouse and schema.



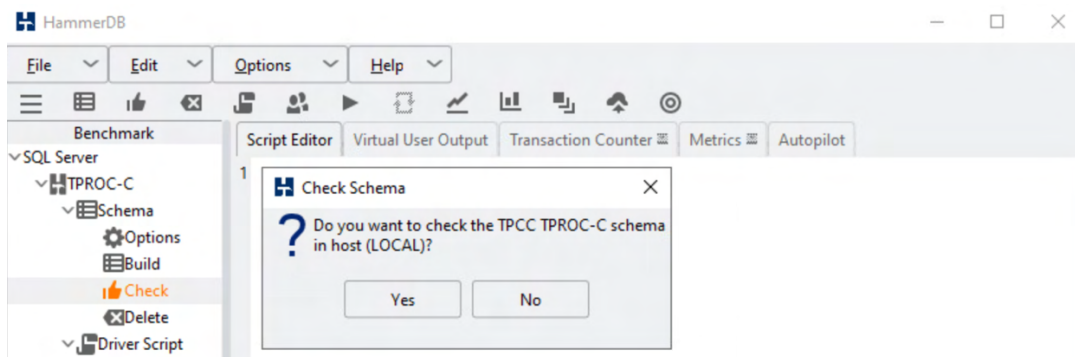
When the build is completed, click the red block icon near the top center of the screen to destroy the virtual user.



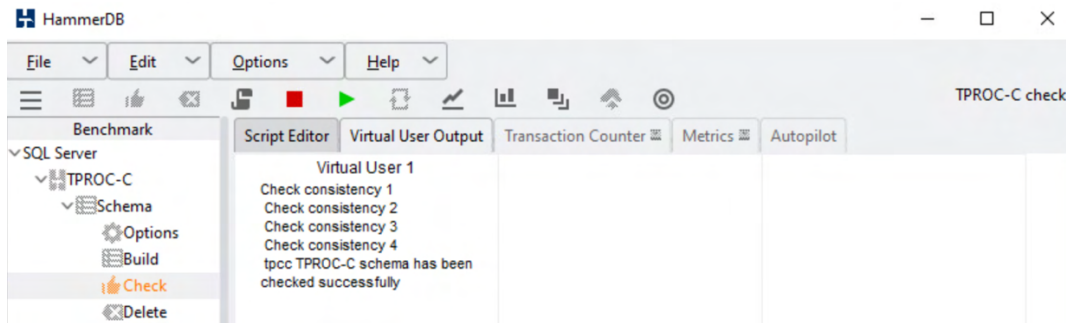
## Check the Schema

To check whether the Schema has been built successfully and if the data is consistent, follow the steps below.

1. Go to the **Benchmark** panel and expand SQL Server > TPROC-C > Schema.
2. Double-click Check and click Yes to check the Schema.



3. Review the output in Script Editor to confirm that tpcc TPROC-C schema has been checked successfully.

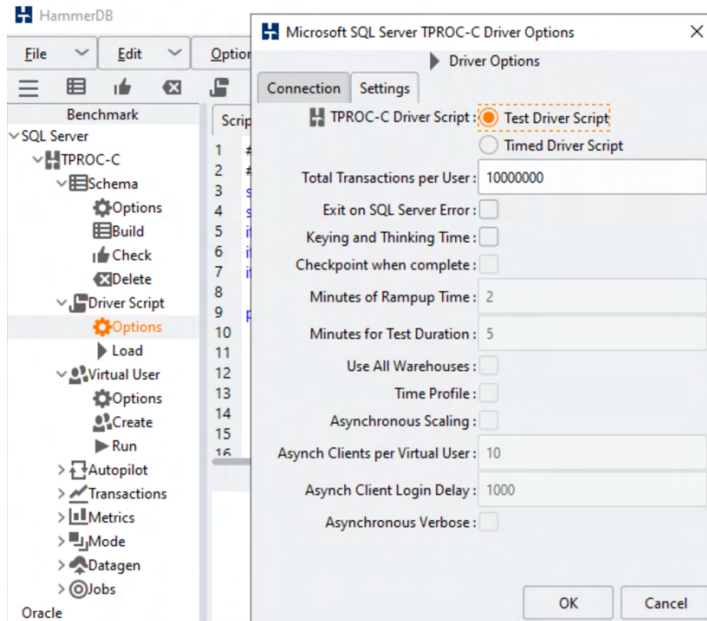


Click the red stop button when done.

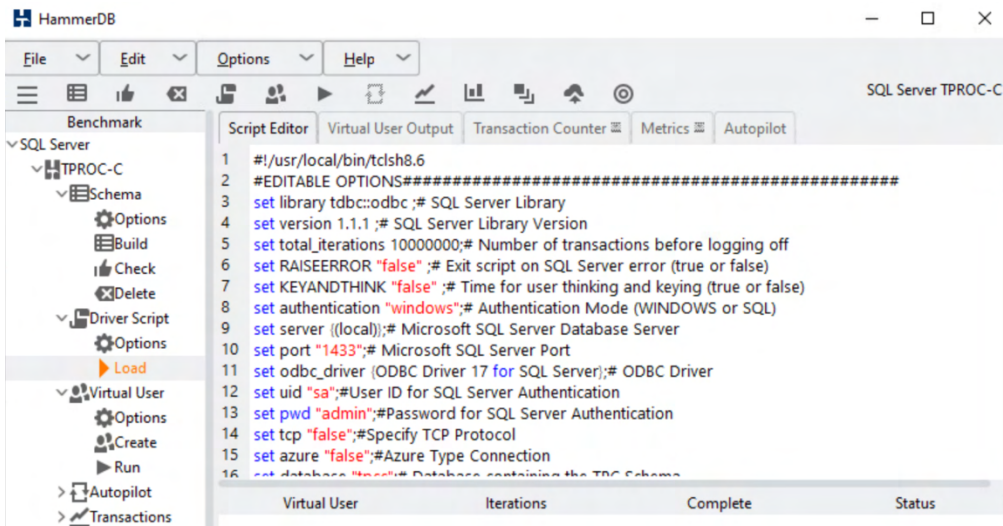
## Run a quick test workload to verify database connectivity

To configure a test workload that can be used to validate database connectivity, follow the steps below.

1. Go to the **Benchmark** panel and expand SQL Server > TPROC-C > Driver Script.
2. Double-click Options and select Settings tab.
3. Select Test Driver Script and then click OK.

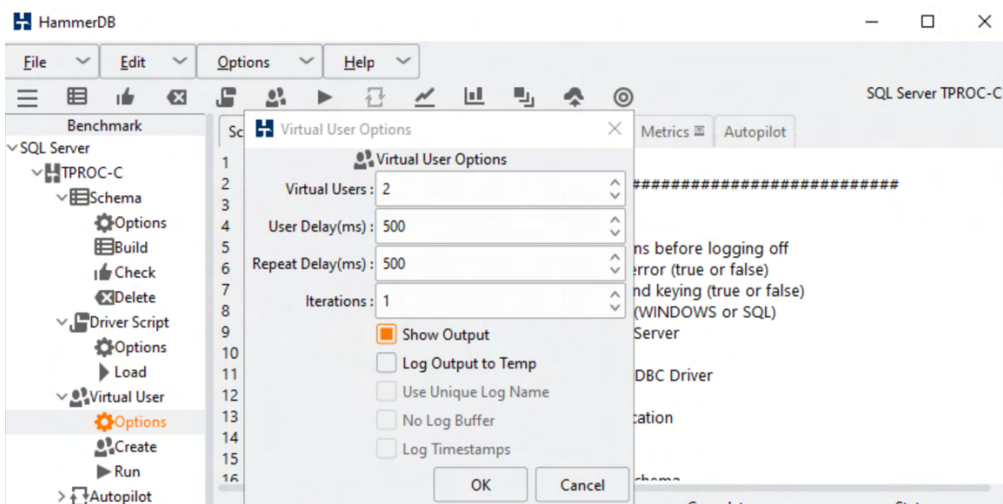


4. Check the Script Editor tab to confirm that the Driver Script is loaded.



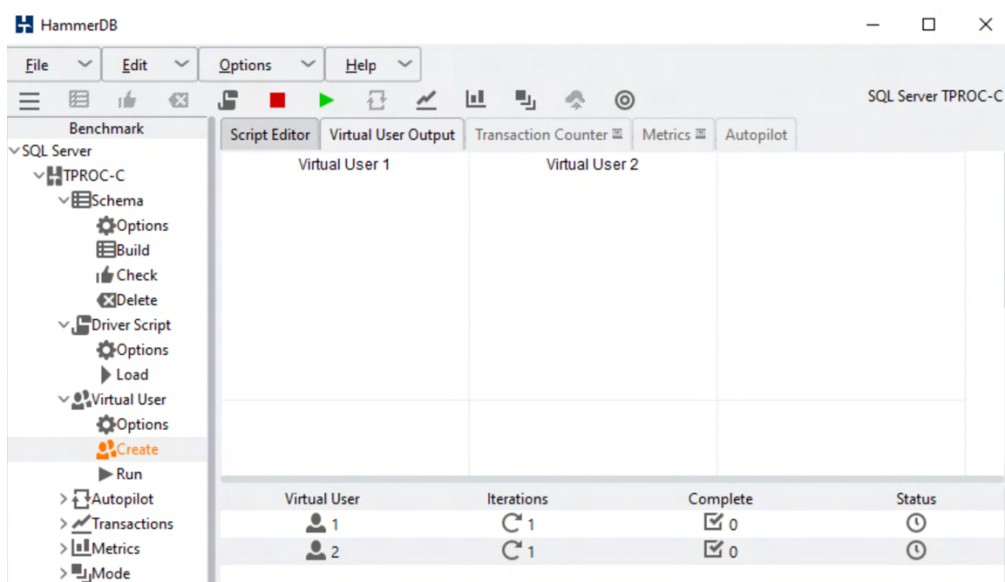
Expand Virtual User on the menu and double-click Options.

Select 2 for the number of Virtual Users and click OK.

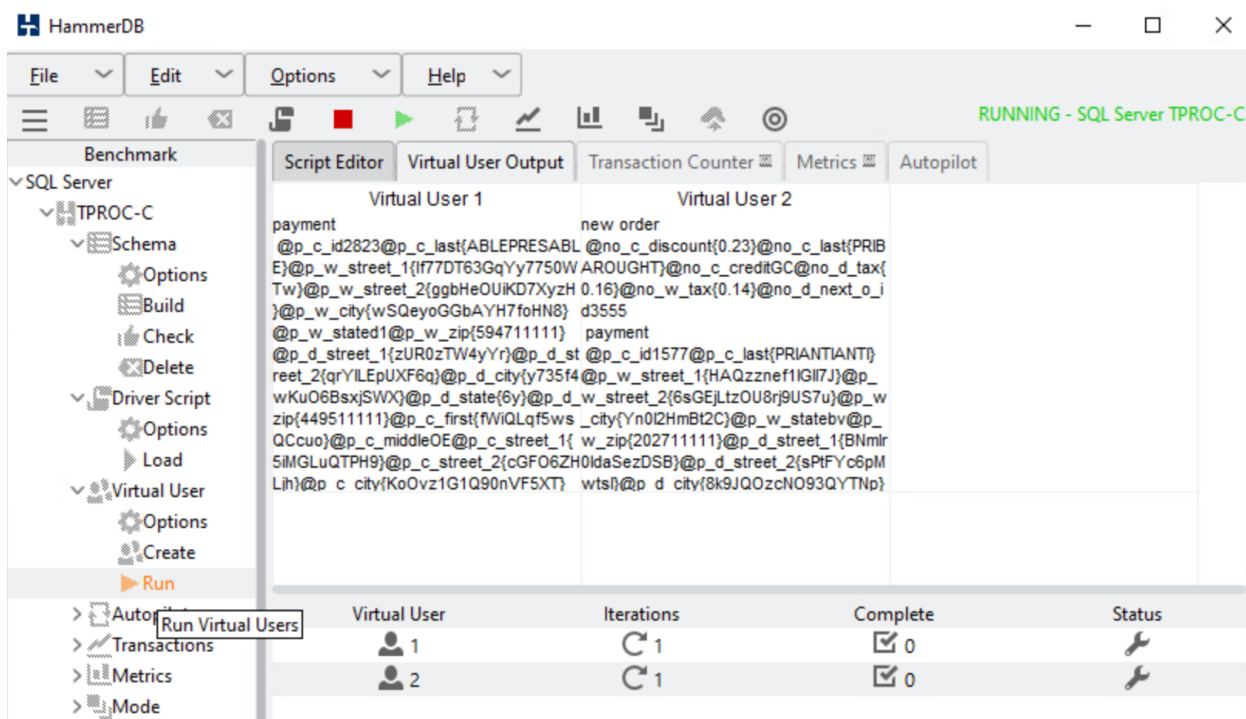


Double-click on Create under the Virtual User menu to create the Virtual Users.





Double-click on Run under Virtual User to observe the workload generated by the Virtual Users.



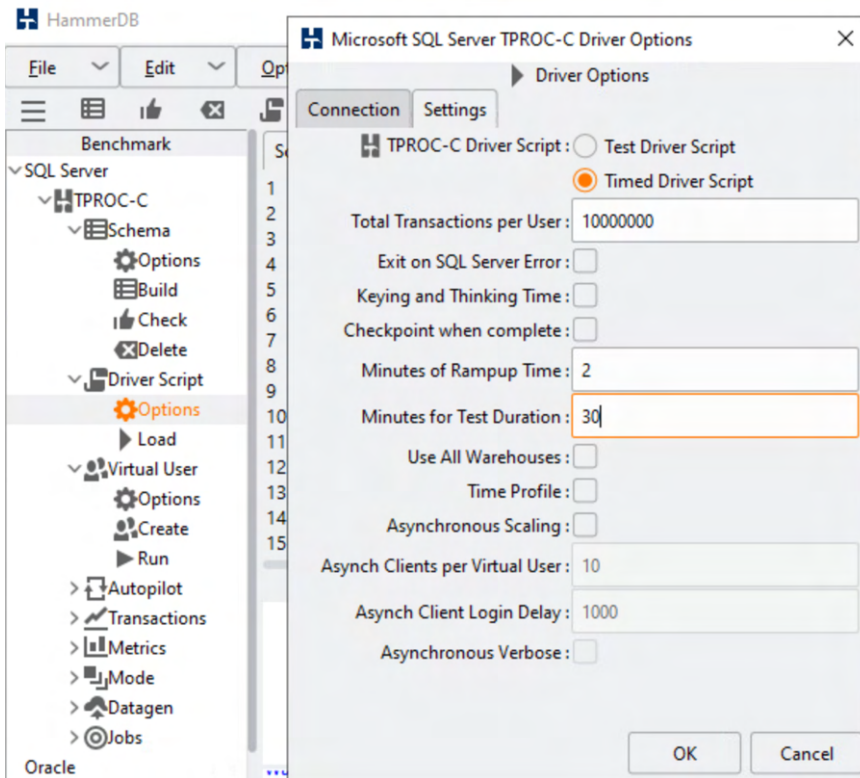
Click the red stop button to stop the workload when you are done with it.

## Run timed workload for performance testing

To conduct performance test, invoke Timed Workload by following the steps below.

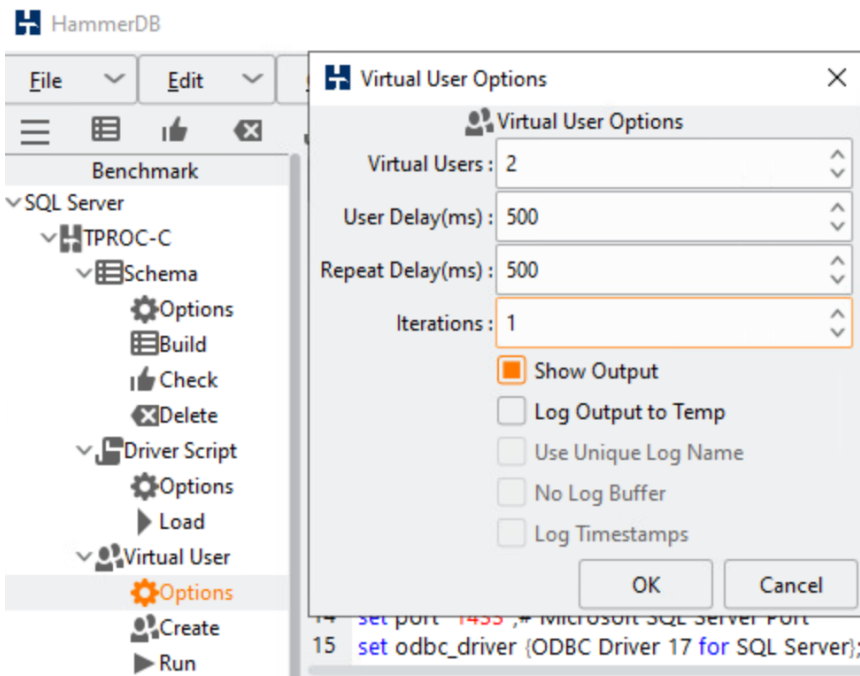
1. Go to the **Benchmark** panel and expand SQL Server > TPROC-C > Driver Script.
2. Double-click Options and select Settings tab.
3. Select Timed Driver Script.
4. Configure desired Minutes of Rampup Time and Minutes for Test Duration and click OK.



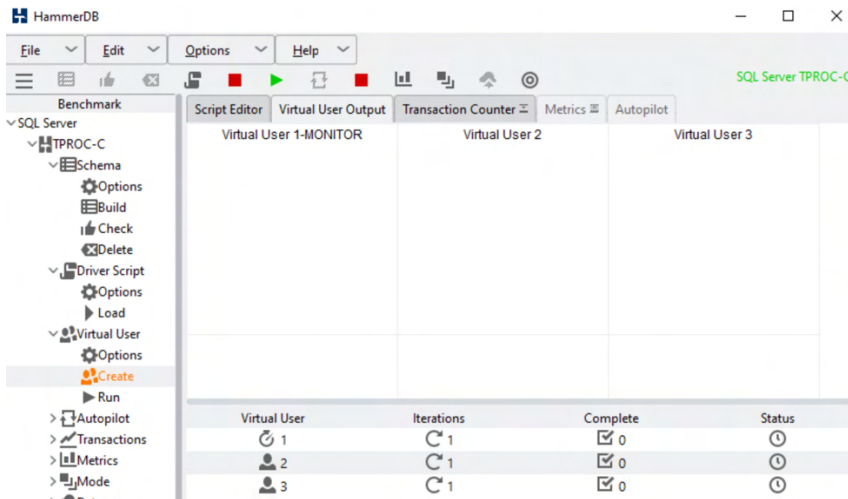


Go to Virtual User and double-click Options to update options if needed.

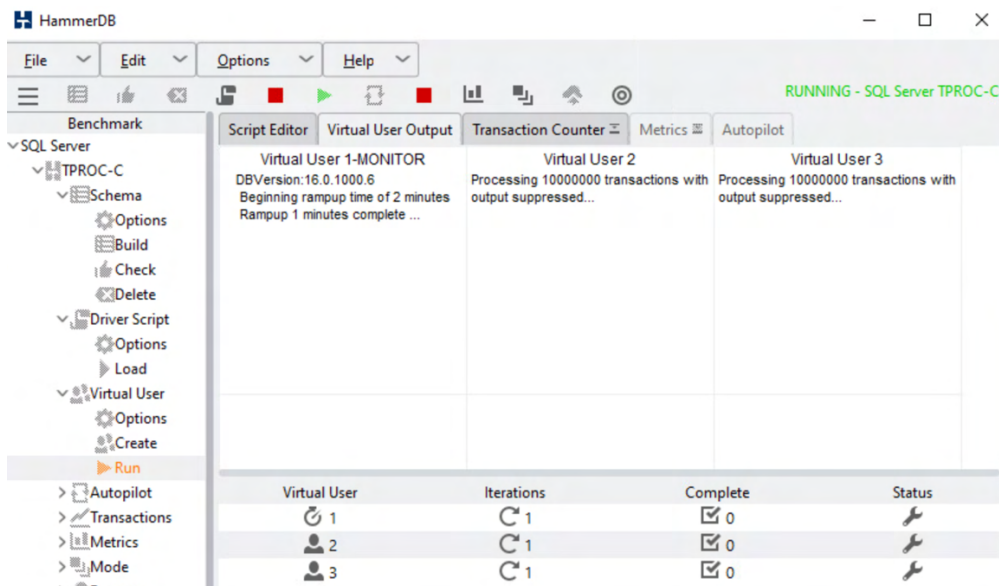
Select the **Show Output** checkbox to enable error messages in the console. Click OK.



Double-click Create under Virtual Users and observe in the Virtual User Output tab that an additional Monitor user has been created for monitoring and timing.

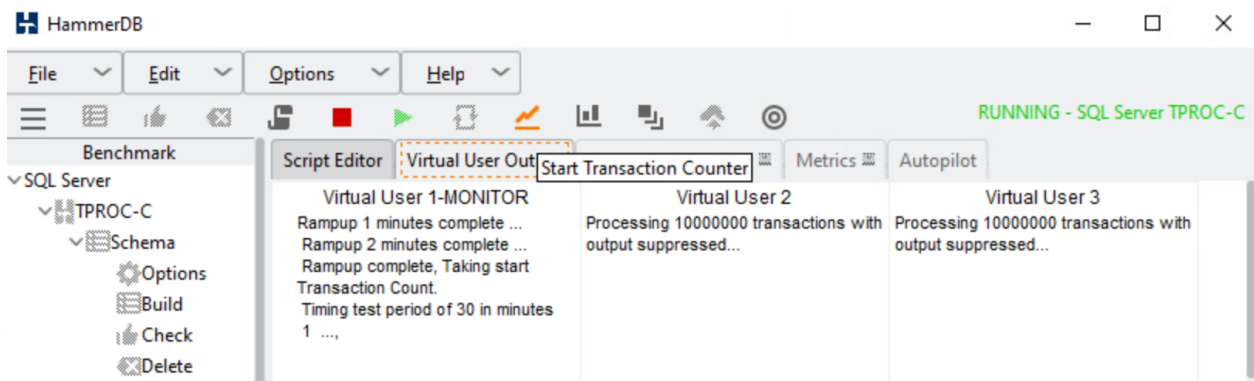


Double-click Run under Virtual User.

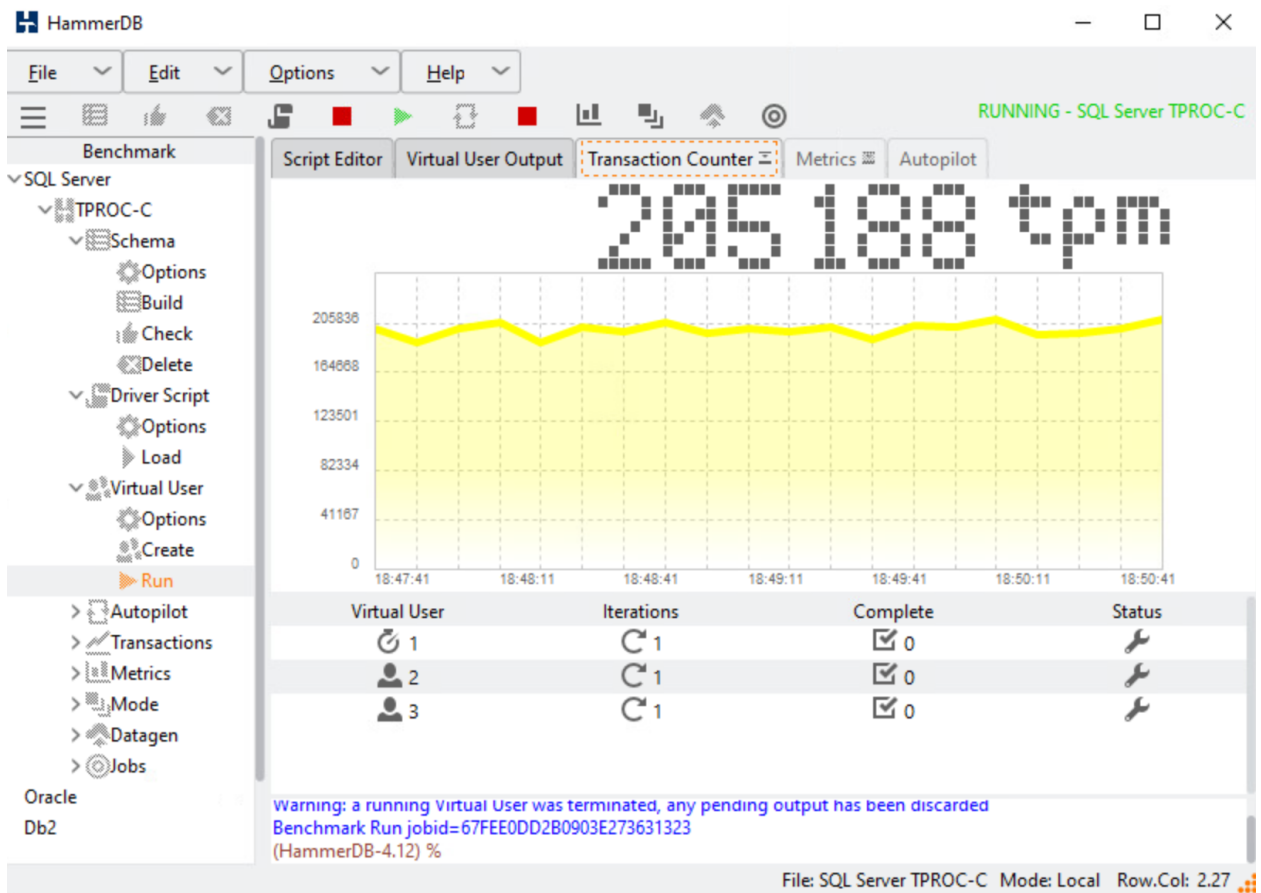


**Note:** The workload begins but the Virtual User output is suppressed. The workload progress is being reported under the Monitor Virtual User as shown above.

After a workload is started, press the Start Transaction Counter button to start the transaction counter reporting of transactions per minutes (TPM).



The transaction counters will be reported under the Transaction Counter tab.



At the end of the testing, the Monitor Virtual User output includes results for New Orders per Minute (NOPM) and Transactions per Minute (TPM) from the testing.

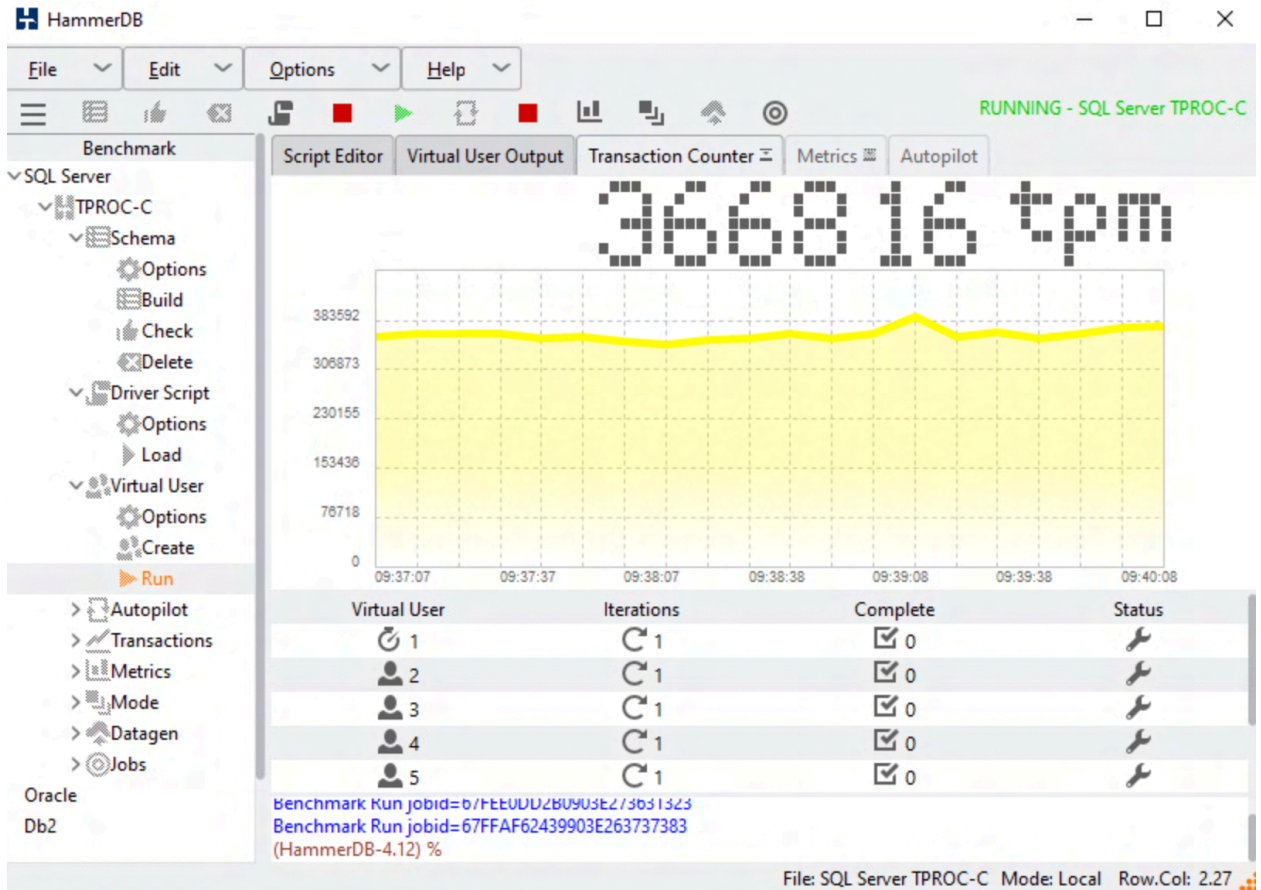
**Benchmark**

- SQL Server
  - TPROC-C
    - Schema
      - Options
      - Build
      - Check
      - Delete
    - Driver Script
      - Options
      - Load
    - Virtual User
      - Options
      - Create
    - Run
    - Autopilot
    - Transactions
    - Metrics
    - Mode
    - Datagen
    - Jobs

**Virtual User Output**

Virtual User	Iterations	Complete	Status
1	1	1	✓
2	1	1	✓
3	1	1	✓

**Note:** The NOPM and TPM results will depend on the specific SQL server solution environment and HammerDB test configurations. For example, in this validation environment, when the number of virtual users was increased from two to six, the monitor user reported 153858 NOPM and 357466 TPM. The following is a Transaction Counter snapshot captured during that test.



## Oracle RAC database testing

The Silly Little Oracle Benchmark (SLOB) is a comprehensive toolkit designed for generating and testing I/O operations within an Oracle database. SLOB is highly effective in evaluating the I/O subsystem using authentic Oracle SGA-buffered physical I/O. It facilitates the testing of physical random single-block reads (db file sequential read) and random single-block writes (DBWR flushing capability). Typically, SLOB issues single-block reads for the read workload, which are generally 8K in size, corresponding to the database block size.

For testing the SLOB workload, we created a pluggable database named ISCSIDBP with following layout:

- 8 x 800GB LUNs are used for SLOB data
- 2 x 200GB LUNs are used for redo log
- Two ASM disk groups were created, one for SLOB data and the other for redo log and both with external redundancy

The SLOB data disk group provided the necessary storage to create one big file tablespaces for the SLOB database. We loaded the SLOB schema onto the SLOB data disk group, with a total size of up to 3 TB.

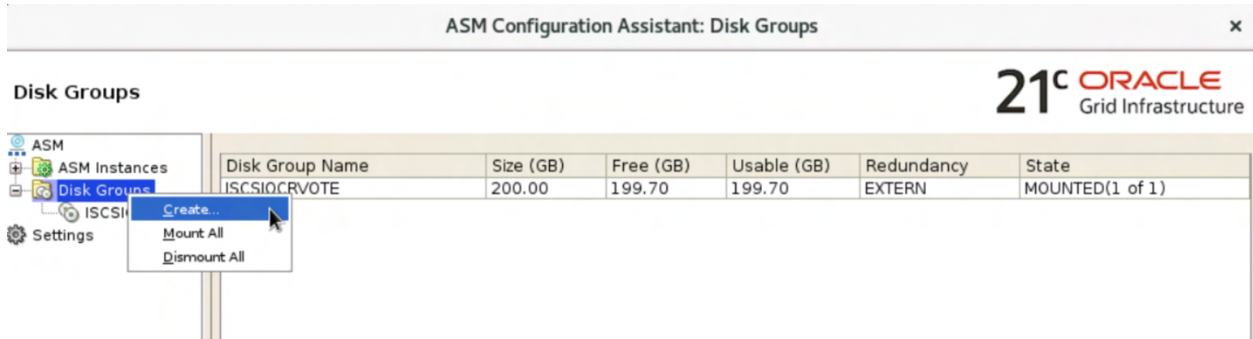
We utilized SLOB to generate OLTP workload on the database server which applied the workload to the Oracle database, log, and temp files. Tests with different numbers of users and database update percentages were conducted and the generated Oracle AWR reports were reviewed for metrics such as IOPs and latencies.



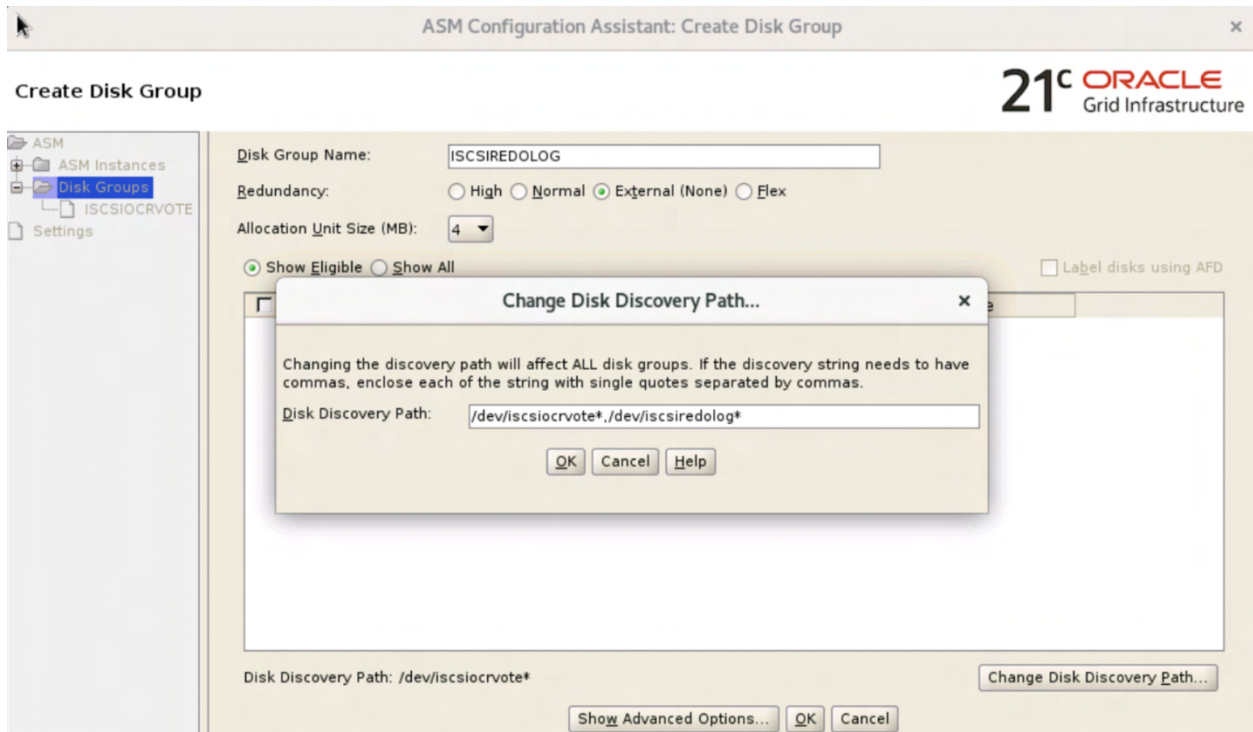
## Create ASM disk groups

After successfully installed Oracle Grid Infrastructure and Oracle RAC Database software, you can use the steps below to create two ASM disk groups for Oracle database usage.

1. Login as grid user and run ASM configuration assistance tool asmca.
2. Right-click on Disk Groups in the left menu and select Create.

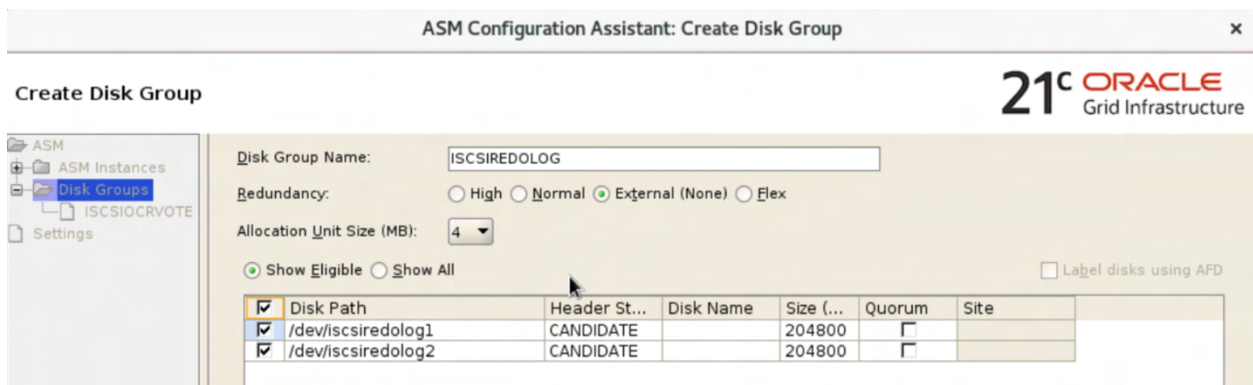


3. Enter Disk Group name "ISCSIREDOLOG", click Change Disk Discovery Path and update it include /dev/iscsiredolog\* to find the ISCSIREDOLOG disk group LUNs and then click OK.

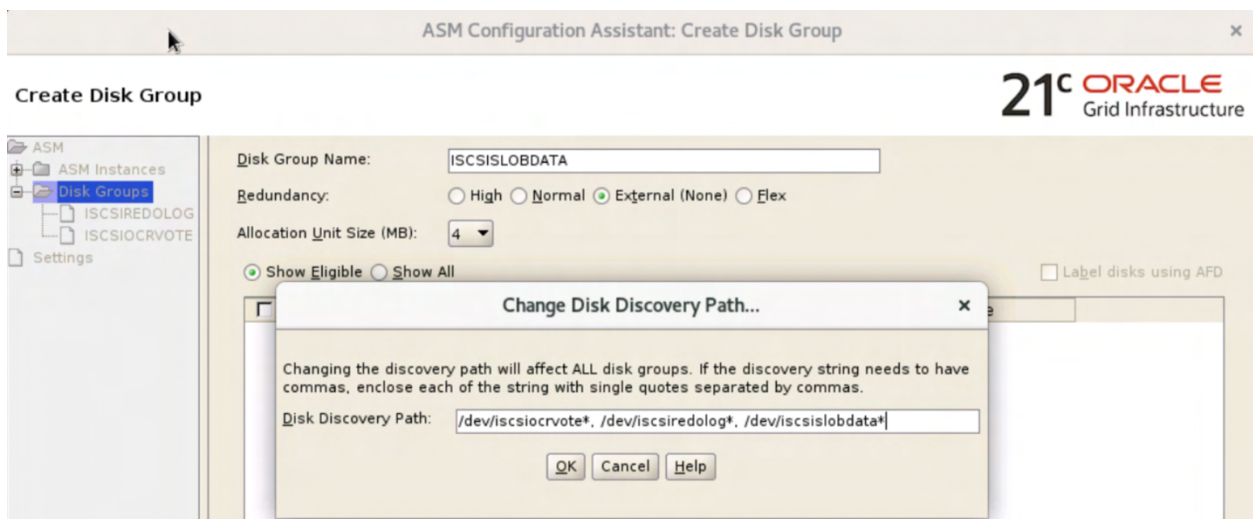


4. Select both /dev/iscsiredolog\* paths, External (None) redundancy and then click OK to create the ISCSIREDOLOG disk group.

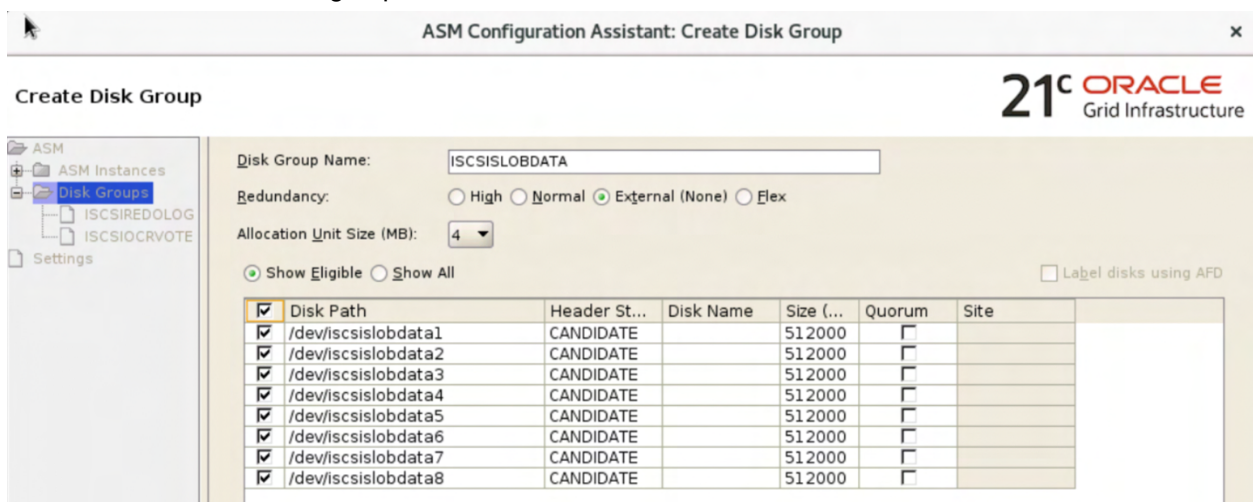




- Right-click on Disk Groups again in the left menu and select Create.
- Enter Disk Group name "ISCSISLOBDATA", click Change Disk Discovery Path and update it include /dev/iscsislobdata\* to find the ISCSISLOBDATA disk group LUNs and then click OK.



- Select all /dev/iscsislobdata\* paths, External (None) redundancy and then click OK to create the ISCSISLOBDATA disk group.



- Check to confirm the two newly created disks groups are listed along with the ISCSIOCRVOTE disk group created during the Oracle Grid Infrastructure installation.

ASM Configuration Assistant: Disk Groups x

### Disk Groups

ASM
ASM Instances
**Disk Groups**
ISCSISLOBD...
ISCSIR...
ISCSIOCRVOTE
Settings

Disk Group Name	Size (GB)	Free (GB)	Usable (GB)	Redundancy	State
ISCSISLOBDATA	4000.00	3999.82	3999.82	EXTERN	MOUNTED(1 of 1)
ISCSIREDOLOG	400.00	399.89	399.89	EXTERN	MOUNTED(1 of 1)
ISCSIOCRVOTE	200.00	199.70	199.70	EXTERN	MOUNTED(1 of 1)

21<sup>c</sup>

**ORACLE**  
Grid Infrastructure

## Create test Oracle database

After the ASM disk groups are created, you can proceed to create a test Oracle database.

1. Login as oracle user and invoke database configuration tool dbca to create a new database.

Select Create a database operation and click Next.

Database Configuration Assistant - Application - Step 1 of 14 x

### Select Database Operation

Database Operation
Creation Mode
Deployment Type
Database Identification
Storage Option
Fast Recovery Option
Database Options
Configuration Options

Select the operation that you want to perform.

☒ Create a database

☐ Configure an existing database

☐ Delete database

☐ Manage templates

☐ Manage pluggable databases

☐ Oracle RAC database instance management

21<sup>c</sup>

**ORACLE**  
Database

Select Typical configuration and enter ISCSIRACDB as the Global database name. Use the default ASM storage type. Browse to select ISCSISLOBDATA disk group as the database file location and ISCSIREDOLOG disk group as the Fast Recovery Area. Enter and confirm the Administrator password, enter ISCSIRACDBP as the pluggable database name, and click Next.

Database Configuration Assistant - Create a database - Step 2 of 14

## Select Database Creation Mode

**21c ORACLE Database**

- Database Operation
- Creation Mode**
- Deployment Type
- Database Identification
- Storage Option
- Fast Recovery Option
- Database Options
- Configuration Options
- Management Options
- User Credentials
- Creation Option
- Summary
- Progress Page
- Finish

☒ Typical configuration

Global database name: ISCSIDB

Storage type: Automatic Storage Management (ASM)

Database files location: +ISCSISLOBDATA/{DB\_UNIQUE\_NAME} [Browse...](#)

Fast Recovery Area (FRA): +ISCSIREDOLOG [Browse...](#)

Database character set: AL32UTF8 - Unicode UTF-8 Universal character set

Administrative password: .....

Confirm password: .....

☒ Create as Container database

Pluggable database name: ISCSIDBP

☐ Advanced configuration

The dbca tool goes through the prerequisite checks and provides a summary.

Database Configuration Assistant - Create 'ISCSIDB' database - Step 4 of 6

## Summary

**21c ORACLE Database**

- Database Operation
- Creation Mode
- Prerequisite Checks
- Summary**
- Progress Page
- Finish

**Database Configuration Assistant**

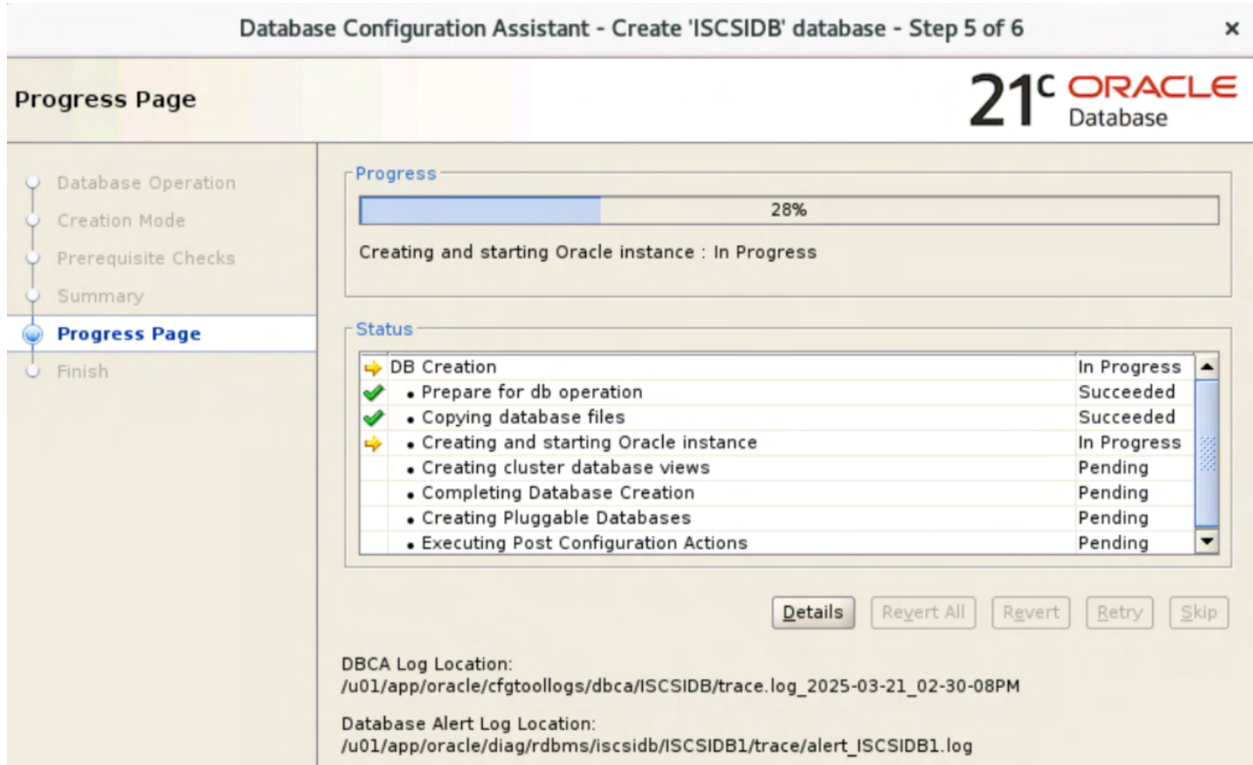
- Global Settings**
  - Global database name: ISCSIDB
  - Node List: fpsa-asa-linux-01
  - Configuration type: Oracle Real Application Cluster (RAC) database - Admin Managed
  - SID: ISCSIDB
  - Create as Container database: Yes
  - Pluggable Database Name: ISCSIDBP
  - Number of Pluggable Databases: 1
  - Use Local Undo tablespace for PDBs: Yes
  - Database Files Storage Type: Automatic Storage Management (ASM)
  - Memory Configuration Type: Automatic Shared Memory Management
  - Template name: General Purpose
- Initialization Parameters**
  - audit\_file\_dest: {ORACLE\_BASE}/admin/{DB\_UNIQUE\_NAME}/adump
  - audit\_trail: db
  - cluster\_database: true
  - compatible: 21.0.0
  - db\_block\_size: 8192 BYTES
  - db\_create\_file\_dest: +ISCSISLOBDATA/{DB\_UNIQUE\_NAME}/
  - db\_name: ISCSIDB
  - db\_recovery\_file\_dest: +ISCSIREDOLOG

Optionally click Save Response File and provide a filename for the response file.

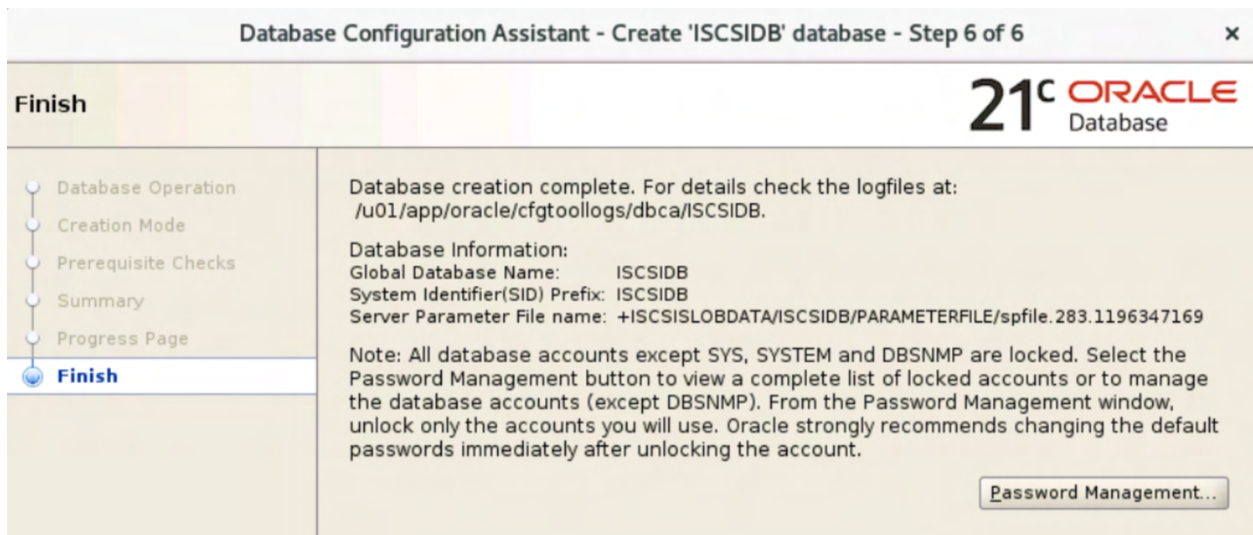
Click Finish and wait for the database creation to complete.



**Note:** The database alert log as well as the DBCA log locations are provided on the progress screen. You can monitor those files for the progress on database creation details.



When the database creation completes, you can see the logfile and database information in the Finish screen. Click Close to exit dbca tool.



## Set environment variables

To access database properly, some environmental variables need to be configured and activated, including ORACLE\_BASE, ORACLE\_HOME, ORACLE\_SID, etc. See below for an example .bash\_profile for the oracle user.

```
[oracle@fpsa-asa-linux-01 ~]$ cat .bash_profile
```

```
# .bash_profile

# Get the aliases and functions
if [ -f ~/.bashrc ]; then
    . ~/.bashrc
fi

# User specific environment and startup programs
umask 022
ulimit -u 16384 -n 65536
export SRVM_DISABLE_MTTRANS=true
export CV_ASSUME_DISTID=OL8
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=$ORACLE_BASE/product/21.3.0/dbhome_1
export PATH=$PATH:$ORACLE_HOME/bin
export ORACLE_SID=ISCSIDB1
export TMP=/tmp
export TMPDIR=$TMP

export LD_LIBRARY_PATH=$ORACLE_HOME/lib:/lib:/usr/lib
export CLASSPATH=$ORACLE_HOME/jre:$ORACLE_HOME/jlib:$ORACLE_HOME/rdbms/jlib

export GRID_HOME=/u01/app/21.3.0/grid
export PATH=$PATH:$GRID_HOME/bin
export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:$GRID_HOME/lib

alias sql='sqlplus / as sysdba'
```

## Download and install SLOB tool

The SLOB tool can be downloaded from the link on the SLOB resource page:

<https://kevinclosson.net/slob/> or directly from the GitHub page below:

[https://github.com/therealkevinc/SLOB\\_2.5.4/blob/main/2021.05.12.slob\\_2.5.4.0.tar.gz](https://github.com/therealkevinc/SLOB_2.5.4/blob/main/2021.05.12.slob_2.5.4.0.tar.gz)

After downloading the SLOB tar archive, it can be extracted to a location that you prefer. To gain some understanding of the SLOB tool and usage example, refer to the information available on the [SLOB Deployment – A Picture Tutorial](#) site and the [SLOB 2 documentation](#).

## Create database tablespace and temp files

### 1. Prepare commands for creating tablespace and temp files.

```
$ cat create_ts_iscsi.sql
CREATE BIGFILE TABLESPACE SLOB DATAFILE '+ISCSISLOBDATA/slob.dbf' SIZE 3072 G;

alter tablespace temp add tempfile '+ISCSISLOBDATA/temp02.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp03.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp04.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp05.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp06.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp07.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp08.dbf' size 30G ;
alter tablespace temp add tempfile '+ISCSISLOBDATA/temp09.dbf' size 30G ;
```

### 2. Connect to database with the SQL\*Plus tool, check database connection, and set session.

```
$ sql

SQL*Plus: Release 21.0.0.0.0 - Production on Fri Mar 21 14:55:50 2025
Version 21.3.0.0.0

Copyright (c) 1982, 2021, Oracle. All rights reserved.

Connected to:
Oracle Database 21c Enterprise Edition Release 21.0.0.0.0 - Production
Version 21.3.0.0.0

SQL> select name,open_mode from v$pdb;
```

```

NAME
-----
OPEN_MODE
-----
PDB$SEED
READ ONLY

ISCSIDBP
READ WRITE

SQL> alter session set container=iscsidbp;

Session altered.

```

**Note:** If the pluggable database open\_mode is MOUNTED, open the database first.

```
SQL> alter pluggable database iscsidbp open;
```

### 3. Run SQL commands from file.

```

SQL> @create_ts_iscsi

Tablespace created.

Tablespace altered.

Tablespace altered.

Tablespace altered.

Tablespace altered.

Tablespace altered.

Tablespace altered.

Tablespace altered.

Tablespace altered.

```

## Run SLOB workload

1. Update the slob.conf and set the desired database update percentage and appropriate access credential.

```

$ cat slob.conf
PDBNAME="ISCSIDBP"
SCAN_PCT=0
UPDATE_PCT=20
RUN_TIME=3600
WORK_LOOP=0
SCALE=35G
SCAN_TABLE_SZ=1M
WORK_UNIT=64
REDO_STRESS=LITE
LOAD_PARALLEL_DEGREE=16
THREADS_PER_SCHEMA=1
DATABASE_STATISTICS_TYPE=awr
OBFUSCATE_COLUMNS=TRUE
DBA_PRIV_USER="system"
SYSDBA_PASSWD="xxxxxxxxxx"
EXTERNAL_SCRIPT=''
DO_HOTSPOT=FALSE
HOTSPOT_MB=8
HOTSPOT_OFFSET_MB=16
HOTSPOT_FREQUENCY=3
HOT_SCHEMA_FREQUENCY=0
THINK_TM_FREQUENCY=
THINK_TM_MIN=.1

```



```
THINK_TM_MAX=.5
```

**Note:** The above configuration calls for 20% update percentage and runs the workload for one hour (3600 seconds).

2. Invoke the workload with the desired user count.

```
./runit.sh 64
```

**Note:** After a test was completed, we reviewed the generated Oracle AWR workload repository report `awr_rac.txt` for the various statistics such as latency and IOPs. For example, the User I/O db file sequential read average wait time, or latency, was 0.483 milliseconds with 20% update and 32 users from one of the runs. The corresponding IOPs, sum of physical reads and writes, observed was around 200,000 IOPs. By performing multiple SLOB workload runs at different % update and user scales, you can get an idea of how the database might perform under various conditions.

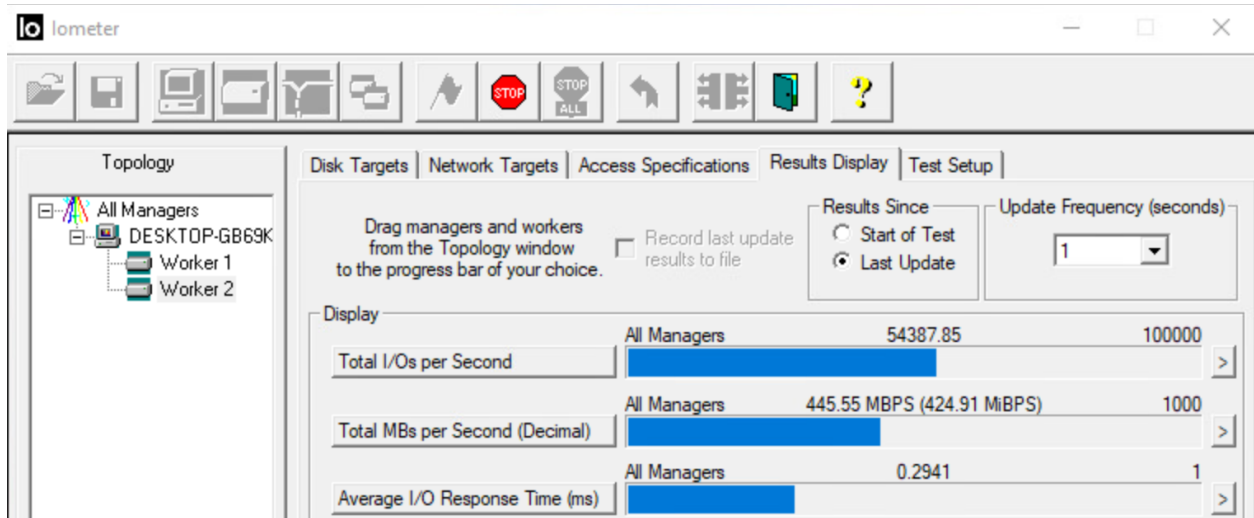
## Solution availability and infrastructure life-cycle management

FlexPod SAN solution infrastructure is designed to provide a highly resilient SAN environment for your mission-critical applications. The infrastructure resiliency is the result of ensuring no single-point-of-failure scenarios by deploying the solution infrastructure with redundant components and redundant connectivity between components.

Due to this redundant infrastructure design, a solution infrastructure provides multiple available paths in the SAN ecosystem to connect to the underlying storage whether the storage solution is using SCSI LUNs with iSCSI protocol or using NVMe namespaces with NVMe/TCP protocol.

The following sections highlights some of the tests performed to ensure infrastructure resiliency and solution availability when there is a component failure and for life-cycle management tasks such as component software / firmware updates.

During the validation, HammerDB tool was used to generate database I/O on the Microsoft SQL server like previously discussed previously. In addition, IOMeter tool was configured to run 8K I/O with 75% read to two virtual disks on another Windows VM with 1 second display updates for real-time observations of IOPS, throughput, and average I/O latency as shown below.



Except for the scenario where a server is powered off which leads to application being stopped and the VMs rely on VMware HA to be restarted on a surviving host, IO should continue to run during other life-cycle management and failure scenario testing due to highly available SAN infrastructure design and SAN multipathing.

## ONTAP storage controller failover

Storage controller failover happens in situations when you would like to add I/O adapters to the storage cluster or upgrade the storage cluster software one node at a time. It is a good test to confirm that you have multipathing configured properly throughout your SAN ecosystem.

To exercise storage failover, you can shut down or reboot one of the controller nodes from ONTAP System Manager using the steps below. We will confirm the SAN paths for a datastore LUN to validate SAN path availability and iSCSI LIF failover.

1. Login to vCenter.

Select one of the ESXi hosts in the data center Inventory view.

Go to Configure tab, select Storage Devices, select one of the storage devices, and then click on the Paths tab below the device list to see the paths.

The screenshot shows the vCenter configuration interface for a host named 'fpsa-asa-esxi-01.nva.local'. The 'Configure' tab is selected, and the 'Storage Devices' section is expanded. The 'Paths' tab is active, showing a table of iSCSI paths. The table has columns: Runtime Name, Status, Target, Transport, Name, and Preferred. There are four rows, all with a status of 'Active (I/O)'.

Runtime Name	Status	Target	Transport	Name	Preferred
vmhba64:C0:T0:L1	Active (I/O)	iqn.1992-08.com.netap.p:sn.57cd8838ea2911ef9.608d039eac6a795:vs.2.1.72.22.73.101:3260	iSCSI	vmhba64:C0:T0:L1	No
vmhba64:C3:T0:L1	Active (I/O)	iqn.1992-08.com.netap.p:sn.57cd8838ea2911ef9.608d039eac6a795:vs.2.1.72.22.74.101:3260	iSCSI	vmhba64:C3:T0:L1	No
vmhba64:C2:T0:L1	Active (I/O)	iqn.1992-08.com.netap.p:sn.57cd8838ea2911ef9.608d039eac6a795:vs.2.1.72.22.74.102:3260	iSCSI	vmhba64:C2:T0:L1	No
vmhba64:C1:T0:L1	Active (I/O)	iqn.1992-08.com.netap.p:sn.57cd8838ea2911ef9.608d039eac6a795:vs.2.1.72.22.73.102:3260	iSCSI	vmhba64:C1:T0:L1	No

Confirm that all four paths to the LUN are showing Active(I/O) status in normal condition.

Login to ONTAP CLI.

Check iSCSI LIF information to confirm that each node has two iSCSI LIFs located in their respective VLAN ports as shown in the example below.

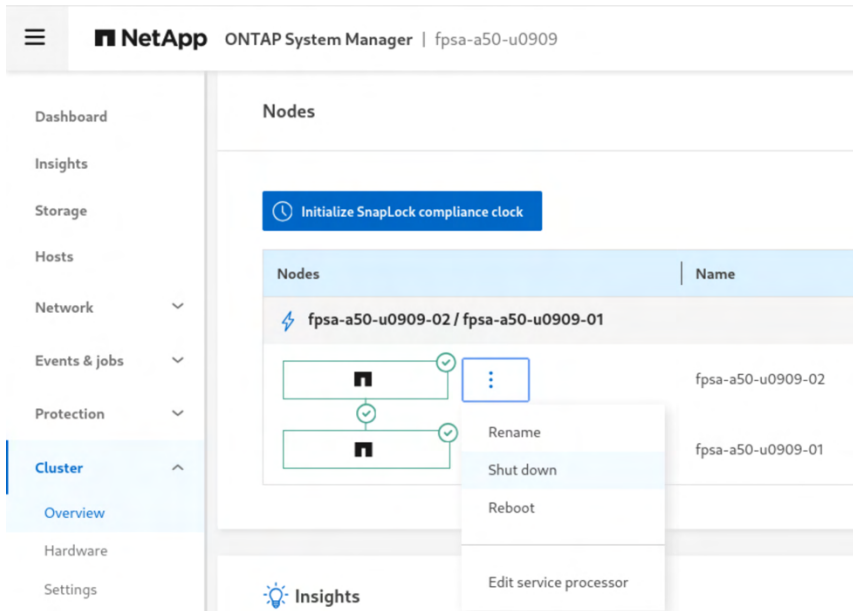
```
fpsa-a50-u0909:>> net int show -lif iscsi*
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
svm1
  iscsi-lif-01a up/up       172.22.73.101/24 fpsa-a50-u0909-01 e2b-2273 true
  iscsi-lif-01b up/up       172.22.74.101/24 fpsa-a50-u0909-01 e4b-2274 true
  iscsi-lif-02a up/up       172.22.73.102/24 fpsa-a50-u0909-02 e2b-2273 true
  iscsi-lif-02b up/up       172.22.74.102/24 fpsa-a50-u0909-02 e4b-2274 true
4 entries were displayed.
```

Login to ONTAP System Manager.

Select Cluster > Overview from the left menu to navigate to the Cluster Overview page.

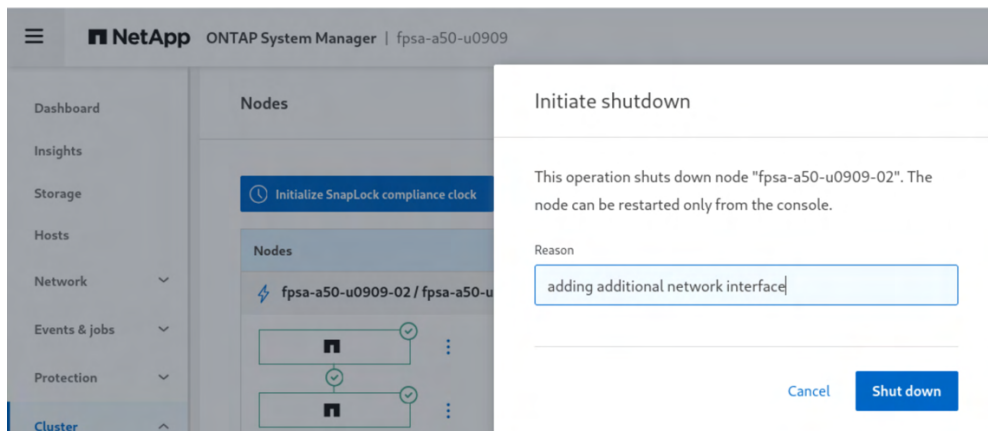
Scroll down to the Nodes section.

Click on the three dots next to a node and select either Shut down or Reboot from the menu.



**Note:** When a node is rebooted or shut down, the partner controller will automatically takeover the workload previously served by the node which was getting rebooted or shut down.

Provide a reason for performing the operation. For example, shutting down a node to add an adapter for additional connectivity.



**Note:** After a shutdown, the node can be started only from the console or remote connection to the node's service processor.

ONTAP System Manager will indicate that a takeover happened. In this case, node 02 was shut down and node 01 has taken over node 02.

NetApp

ONTAP System Manager | fpsa-a50-u0909

Dashboard

Insights

Storage

Hosts

Network

Events & jobs

Protection

Cluster

Overview

Nodes

Initialize SnapLock compliance clock

Nodes

Name

fpsa-a50-u0909-02 / fpsa-a50-u0909-01

Node "fpsa-a50-u0909-01" has taken over node "fpsa-a50-u0909-02".

fpsa-a50-u0909-02

fpsa-a50-u0909-01

Double check HammerDB transaction counter view and IOMeter IO and both should indicate that IO is still running.

Go back to vCenter.

Select the same ESXi host from the data center Inventory view.

Go to Configure tab, select Storage Devices, select one of the storage devices, and then click on the Paths tab below the device list to see the paths.

Confirm that all paths to the same LUN are still showing Active(I/O) status which indicates that iSCSI LIF migration happened.

**Note:** If you see two of the paths showing Dead status, please go back to the iSCSI LIF configuration section for steps on properly configuring iSCSI-only LIFs with failover enabled.

Go back to ONTAP CLI.

Check iSCSI LIF information to confirm that the LIFs originally located in node 02 were migrated to node 01 as shown in the example below. The migrated LIFs should show false for the Is Home status.

```

fpsa-a50-u0909::> net int show -lif iscsi*
(network interface show)

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svm1	iscsi-lif-01a	up/up	172.22.73.101/24	fpsa-a50-u0909-01	e2b-2273	true
	iscsi-lif-01b	up/up	172.22.74.101/24	fpsa-a50-u0909-01	e4b-2274	true
	iscsi-lif-02a	up/up	172.22.73.102/24	fpsa-a50-u0909-01	e2b-2273	false
	iscsi-lif-02b	up/up	172.22.74.102/24	fpsa-a50-u0909-01	e4b-2274	false

4 entries were displayed.

When you are ready to boot node 02 up again, connect to the console of node 02 and issue autoboot command from LOADER prompt.

```
LOADER-B> autoboot
```

After a few minutes, the cluster should be back to normal status again. You should check the paths to devices on the ESXi hosts to confirm that all paths are in Active(I/O) status. In addition, check iSCSI LIF information with ONTAP CLI to confirm that the migrated iSCSI LIFs are now back to their home node.

206

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture

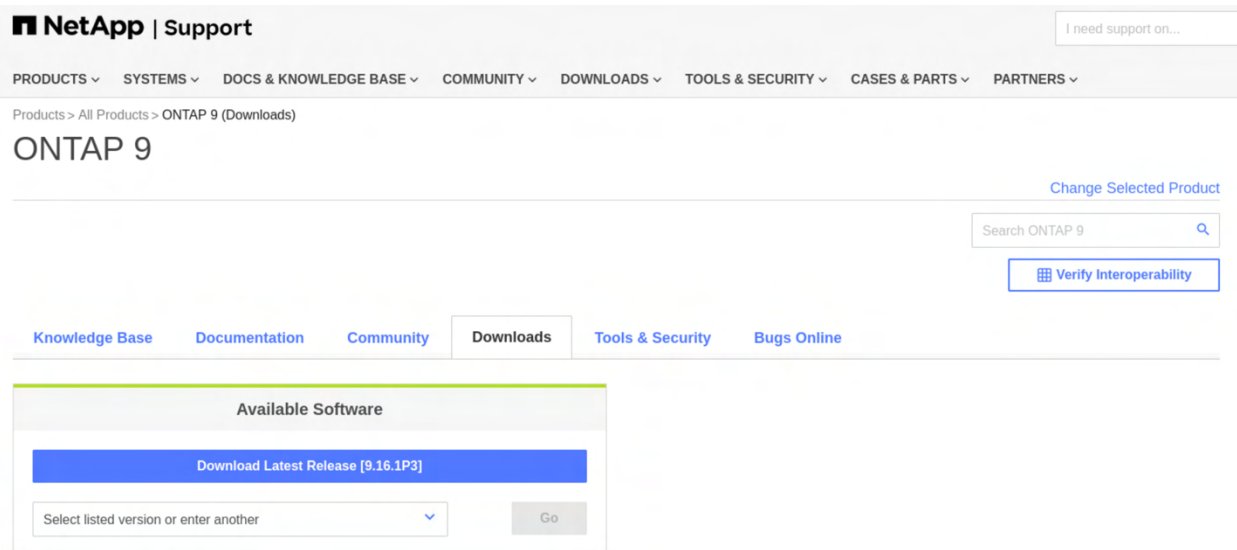
**Note:** For the use case of adding additional network interface cards into the storage cluster, you can then proceed to go through the process of shutting down the other node to add the interface card in that node before booting it up again.

## ONTAP software upgrade

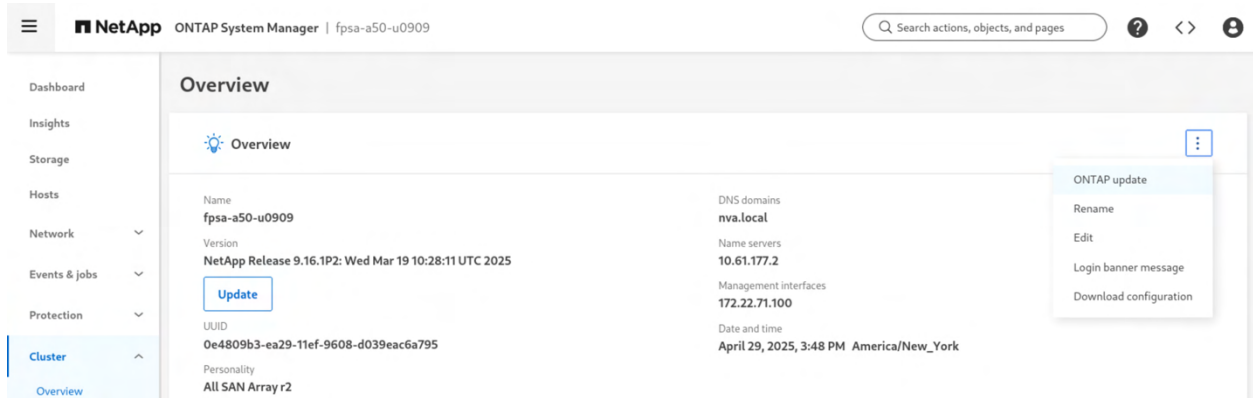
ONTAP storage cluster software and firmware upgrades utilize the storage failover mechanism to upgrade the nodes one at a time. Before performing an upgrade, be sure to confirm the ecosystem interoperability for the upgrade destination combination using the process discussed previously.

To perform ONTAP upgrade, follow the steps below.

1. Login to NetApp support site <https://support.netapp.com> with your credential.
2. Go to the [ONTAP 9 download link](#).
3. Click on the Download Latest Release link.



4. Review the Caution / MustRead information and End User License Agreement. Click on the box below the information to acknowledge and accept. Click ACCEPT & CONTINUE.
5. Review the information regarding Restrictions on Encryption Technology and select the appropriate ONTAP 9 image to download.
6. Login to ONTAP System Manager.
7. Expand the Cluster menu on the left and select Overview page.
8. Click on the Update button under ONTAP version or click on the three dots on the right and select ONTAP update.



- On the Software updates screen, click on +Add image and download the image from the server or browse to select the downloaded ONTAP image for upload from the local client.

#### Select an ONTAP image to update

Either add an image from your local machine or an HTTP or FTP server, or select an existing image below.

+ Add image

Download from the server

Upload from the local client

Wednesday, Mar 19, 2025, 6:28 AM

- After the ONTAP image has been uploaded, click to select the uploaded image and click Update at the bottom of the page.

#### Select an ONTAP image to update

Either add an image from your local machine or an HTTP or FTP server, or select an existing image below.

+ Add image

Available software image

☐

9.16.1P2 (Image build time: Wednesday, Mar 19, 2025, 6:28 AM)

☒

9.16.1P3 (Image build time: Wednesday, Apr 23, 2025, 10:50 PM)

Node	Phase	Status	Approximate time elapsed
▼ fpsa-a50-u0909-01	ONTAP updates	Completed	30 mins 2 secs
▼ fpsa-a50-u0909-02	ONTAP updates	Completed	30 mins 31 secs
1 - 2 of 2 << < 1 > >>			

✓ The post update checks were completed.

Update

Validate

- A series of pre-update checks will be performed. Afterwards, and review the resulting messages from the pre-update checks.

Pre-update check	Status	Message
▼ Manual checks that can be done using Upgrade ONTAP documentation	⚠	Manual validation checks need to be performed. Refer to ...o can result in an update failure or an I/O disruption.
▼ Nodes to update list	⚠	List of nodes to be upgraded: "fpsa-a50-u0909-01, fpsa-a50-u0909-02."
▼ ONTAP API to REST transition warning	⚠	NetApp ONTAP API has been used on this cluster for O...NetApp ONTAP API is approaching end of availability.
▼ SAN compatibility for manual configurability check	⚠	Since this cluster is configured for SAN, manually confirm that the SAN configuration is fully supported.
1 - 4 of 4 << < 1 > >>		

Update with warnings

Validate

Back

- Address any errors or warnings and click Validate again if needed or click Update with warnings to proceed with update.
- During the ONTAP update process, you can click on the down arrow in front of the controller node names to expand the details of the update status.



Cluster version  
**NetApp Release 9.16.1P2: Wed Mar 19 10:28:11 UTC 2025**  
Updated version  
**9.16.1P3**

Updating ...

Estimated time:1 hrs 42 mins

Time elapsed:33 mins 35 secs

Node	Phase	Status	Approximate time elapsed
▼ fpga-a50-u0909-01	ONTAP updates	In progress	3 mins 21 secs
▲ fpga-a50-u0909-02	ONTAP updates	Completed	30 mins 24 secs

Subtask	Status	Start time	End time	Message	Advice
Download	Completed	Tuesday, Apr 29, 2025, 4:19 PM	Tuesday, Apr 29, 2025, 4:20 PM	Image update complete	-
Failover	Completed	Tuesday, Apr 29, 2025, 4:20 PM	Tuesday, Apr 29, 2025, 4:29 PM	Takeover complete	-
Giveback	Completed	Tuesday, Apr 29, 2025, 4:29 PM	Tuesday, Apr 29, 2025, 4:49 PM	Giveback complete	-

1 - 2 of 2 << < 1 > >>

Pause

**Note:** For a single HA pair configuration, the updates for controllers happen sequentially for the nodes through a takeover and giveback process after completing image update for the new image to take effect. Client IO should continue during the non-disruptive ONTAP upgrade process, assuming SAN multipathing had been properly configured as described in this NVA.

14. Wait for the upgrade process to complete and optionally delete the prior ONTAP image.

Software updates

Cluster settings

ONTAP updates

All other updates

The operation installs a new version of ONTAP. Storage services remain online during the upgrade. [Reasons to update ONTAP.](#)

[View update history](#)

Cluster version

NetApp Release 9.16.1P3: Thu Apr 24 02:50:10 UTC 2025

Select an ONTAP image to update

Either add an image from your local machine or an HTTP or FTP server, or select an existing image below.

+ Add image

Available software image

9.16.1P3 (Image build time: Wednesday, Apr 23, 2025, 10:50 PM)

Node	Phase	Status	Approximate time elapsed
▼ fpga-a50-u0909-01	ONTAP updates	Completed	30 mins 21 secs
▼ fpga-a50-u0909-02	ONTAP updates	Completed	30 mins 24 secs

1 - 2 of 2 << < 1 > >>

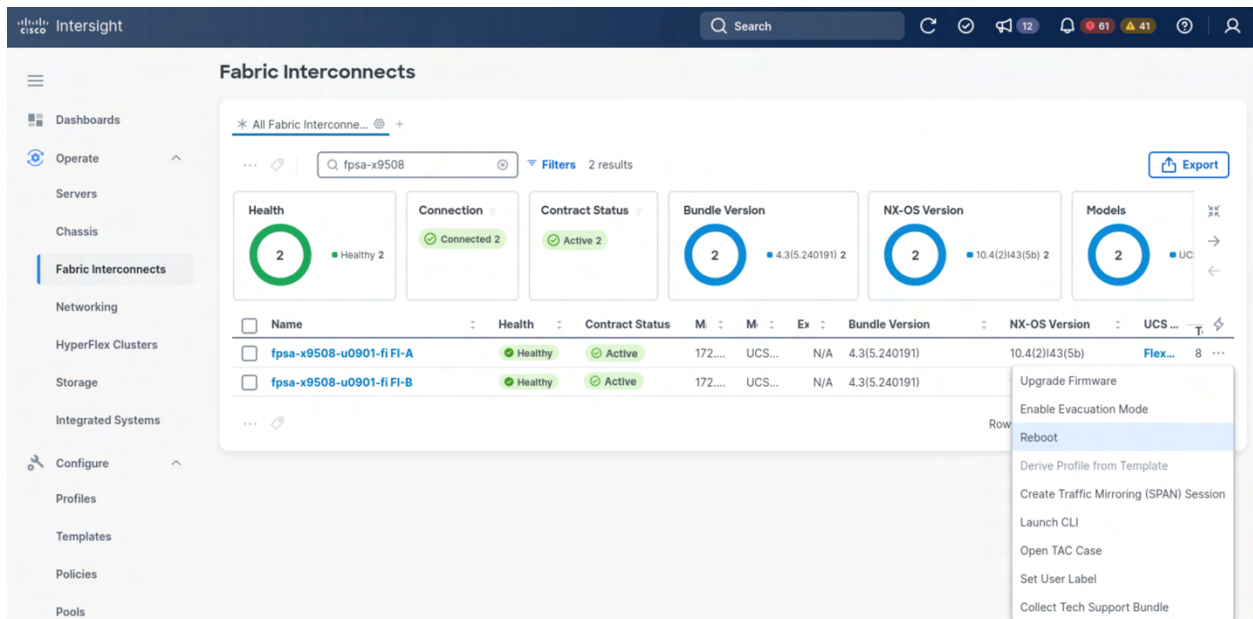
The post update checks were completed.

## UCS fabric interconnect reboot

Fabric interconnect reboot is needed after firmware upgrade. Testing the impact of fabric interconnect reboot is a good way to make sure that the solution is behaving properly in preparation for future fabric interconnect firmware upgrades. It is also a good test to confirm that you have multipathing configured properly throughout your SAN ecosystem.

To exercise fabric interconnect reboot, follow the steps below.


1. Login to Intersight.
2. Navigate to Operate > Fabric Interconnects.
3. Put a partial name of the FIs into Filters to narrow down the search if needed.
4. Click on the three dots near the right to open the actions menu to select Reboot.



5. Review the information in Reboot Fabric Interconnect dialog and then click Reboot.

### Reboot Fabric Interconnect

The selected Fabric Interconnect 'fpsa-x9508-u0901-fi FI-A' will reboot.

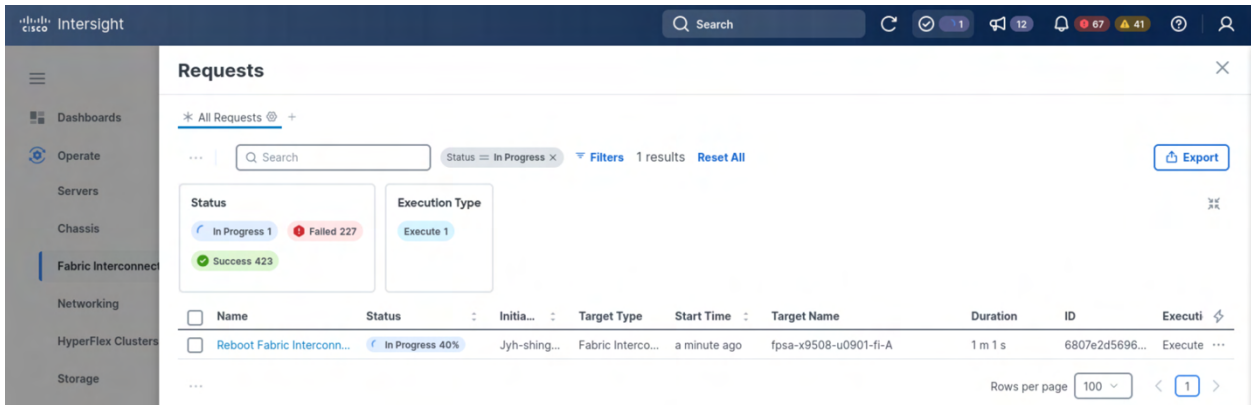
 All traffic in the Fabric Interconnect will stop. Traffic will failover to the peer Fabric Interconnect for failover vNICs.

☒ Evacuate Fabric Interconnect traffic before reboot.

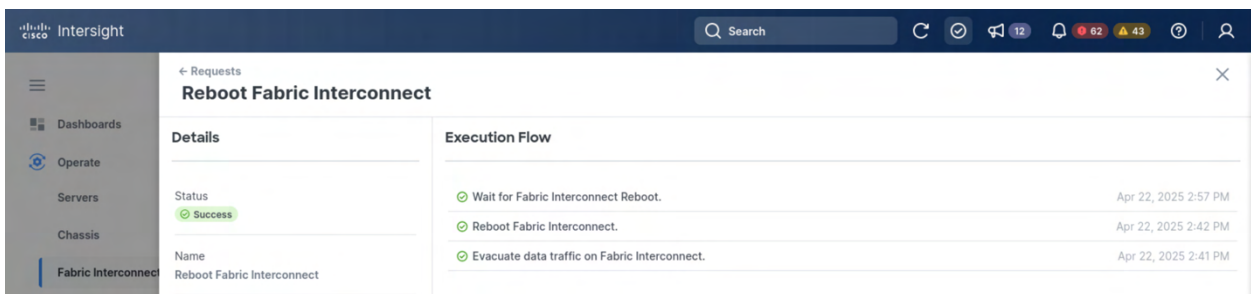
☐ Force Reboot ⓘ

[Cancel](#) [Reboot](#)

6. Click on the spinning Requests icon at the top to open the Requests view and monitor the progress.



7. Check on IOMeter IO and HammerDB IO to confirm that storage IO in the VMs are still going.
8. Check on ESXi host LUN paths to confirm the paths going through fabric interconnect A show Dead status while the paths going through fabric interconnect B are still showing Active(I/O) status.
9. Wait for Intersight Reboot Fabric Interconnect task to complete.



10. Perform a rescan of the iSCSI adapter on the ESXi hosts and check LUN paths again to confirm that all paths are showing Active(I/O) status.

## UCS fabric interconnect upgrade

If your fabric interconnects are not already running firmware bundle 4.3(5.240191) (NX-OS version 10.4(2)143(5b)), upgrade them to 4.3(5.240191). To perform the firmware upgrade, follow the steps below.

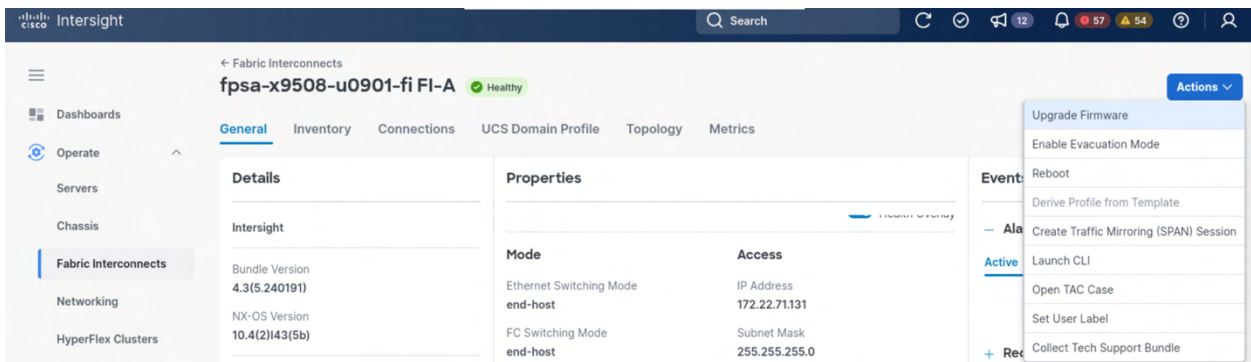
1. Login to Intersight.

Navigate to Operate > Fabric Interconnects.

Find the Fabric interconnect you wish to upgrade and click on its link under the Name column.

Check the software bundle version under the Details section under the General tab.

Click Actions drop-down box and select Upgrade Firmware.



Click Start at the bottom of the Upgrade firmware screen.

Click Next in the General screen.

Select the desired firmware version listed in the Version screen and click Next.

The screenshot shows the Cisco Intersight 'Upgrade Firmware' interface. The left sidebar contains navigation options: Dashboards, Operate (selected), Servers, Chassis, Fabric Interconnects (highlighted), Networking, HyperFlex Clusters, Storage, Integrated Systems, and Configure. The main panel has three tabs: General, Version (selected), and Summary. The 'Version' tab displays a table of available firmware bundles. A blue information box at the top states: 'The selected firmware bundle will be downloaded from intersight.com. By default, the upgrade enables Fabric Interconnect traffic evacuation. Use Advanced Mode to exclude Fabric Interconnect traffic evacuation.' Below this is a toggle for 'Advanced Mode' which is currently off. A search bar and a 'Filters' button (showing 5 results) are present. The table lists three versions, with the second one selected.

Version	Size	Release Date	Description
<input type="radio"/> 4.3(6.250094)	2.64 GiB	Apr 30, 2025 11:59 ...	Cisco Intersight Infrastructure Bundle
<input checked="" type="radio"/> 4.3(5.250033)	2.64 GiB	Apr 10, 2025 12:02 ...	Cisco Intersight Infrastructure Bundle
<input type="radio"/> 4.3(5.240191)	2.64 GiB	Jan 20, 2025 5:07 ...	Cisco Intersight Infrastructure Bundle

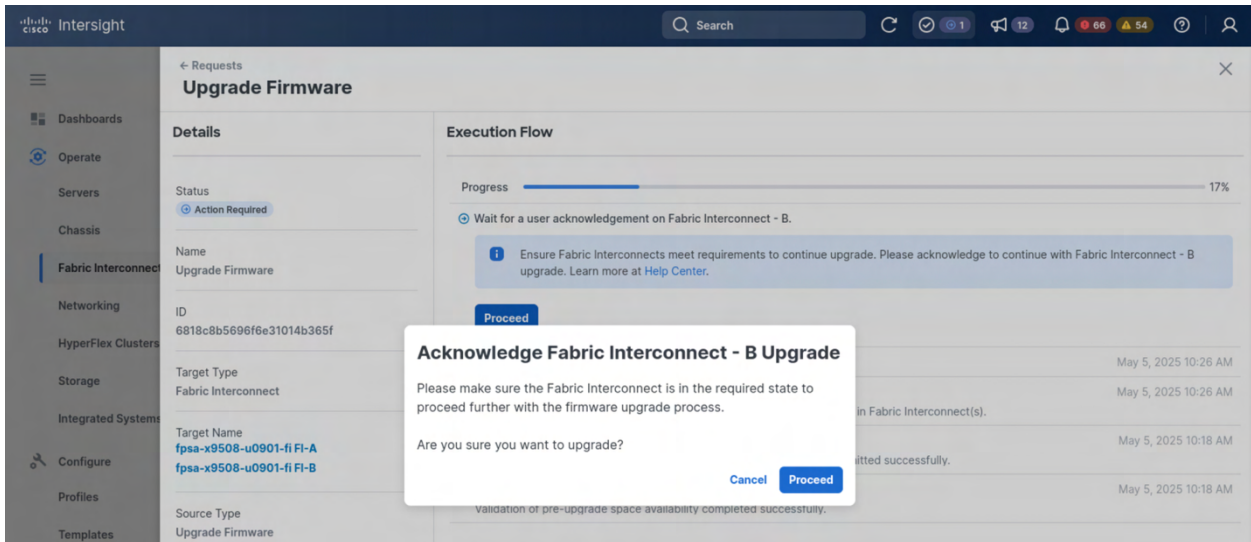
**Note:** Here we are selecting the 4.3(5.250033) firmware bundle for the purpose of demonstrating the upgrade procedures.

Confirm the firmware bundle in the Summary screen and Click Upgrade.

Click on the Requests at the top to view Requests and click the Firmware upgrade request to monitor its progress.

The screenshot shows the 'Requests' tab in the 'Upgrade Firmware' section. The left sidebar is the same as the previous image. The main panel has two tabs: Details and Execution Flow. The 'Details' tab shows the status of the request as 'In Progress'. It lists the Name as 'Upgrade Firmware', the ID as '6818c8b5696f6e31014b365f', the Target Type as 'Fabric Interconnect', and the Target Name as 'fpa-x9508-u0901-fi FI-A' and 'fpa-x9508-u0901-fi FI-B'. The 'Execution Flow' tab shows a progress bar at 8% and a list of steps: 'Wait for image download to complete in endpoint.' (0% completed), 'Initiate image download to the endpoint.' (Download request for intersight-ucs-x-direct-infra.4.3.5.250033.bin submitted successfully, May 5, 2025 10:18 AM), and 'Validate the requirements for the endpoint.' (Validation of pre-upgrade space availability completed successfully, May 5, 2025 10:18 AM).

During the upgrade process, Intersight requires user acknowledgement to perform upgrade. Click Proceed when the Acknowledge dialog shows up.



**Note:** The Fabric Interconnect upgrades are performed serially and requires acknowledgement before each one is upgraded.

Wait for the upgrades to complete on both FIs.

## Virtual machine migration

For a VMware virtual infrastructure, the mobility of virtual machines is an important functionality. Being able to move the virtual machine to a different host with vMotion allows a host to be serviced for situations like adding additional memory or replacing the server when it goes out of support period.

With ONTAP storage, you can have differentiated Quality of Service (QoS) for your different datastores that are based on LUNs or NVMe namespaces. With storage vMotion, VMs can be migrated to different datastores for management or performance reasons.

To have seamless migration of a virtual machine, the virtual infrastructure needs to be properly configured as illustrated in this solution. Network connectivity as well as network configurations need to be consistent across the cluster of ESXi hosts. In addition, the datastores need to be mapped consistently across the cluster of ESXi hosts.

For this validation, we performed VM migration for the IOMeter VM from one host to another as well as migrating the backing datastore from one to another using the steps below.

1. Login to vCenter.
2. Right-click on the VM to be migrated and select Migrate from the pop-up menu.
3. For the migration type, select both Change both compute resource and storage to exercise both vMotion and storage vMotion. Click VM ORIGIN to see the VM information. Click NEXT.



Migrate | fpsa-asa-iometer-01

1 Select a migration type
2 Select a compute resource
3 Select storage
4 Select networks
5 Select vMotion priority
6 Ready to complete

### Select a migration type

Change the virtual machines' compute resource, storage, or both.

☐ Change compute resource only  
Migrate the virtual machines to another host or cluster.
☐ Change storage only  
Migrate the virtual machines' storage to a compatible datastore or datastore group.
☒ Change both compute resource and storage  
Migrate the virtual machines to a specific host or cluster and their storage to a compatible datastore or datastore group.
☐ Cross vCenter Server export  
Migrate the virtual machines to a vCenter Server not linked to the current SSC.

VM ORIGIN ⓘ

Cluster
FlexPod
Host
fpsa-asa-esxi-01.nva.local
Networks
IB-MGMT Network
Storage
iscsi\_datastore\_1
vmware\_swap

4. For the compute resource, expand the FlexPod cluster and select the other host. click NEXT.

Migrate | fpsa-asa-iometer-01

1 Select a migration type
2 Select a compute resource
3 Select storage
4 Select networks
5 Select vMotion priority
6 Ready to complete

### Select a compute resource

Select a cluster, host, vApp or resource pool to run the virtual machines.

VM ORIGIN ⓘ

fpsa-asa-vcenter8.nva.local
FlexPod-DC
FlexPod
fpsa-asa-esxi-01.nva.local
fpsa-asa-esxi-02.nva.local

5. For storage, select a different datastore to host the VM. Click NEXT.

Migrate | fpsa-asa-iometer-01

1 Select a migration type
2 Select a compute resource
3 Select storage
4 Select networks
5 Select vMotion priority
6 Ready to complete

### Select storage

Select the destination storage for the virtual machine migration.

VM ORIGIN ⓘ

BATCH CONFIGURE
CONFIGURE PER DISK

Select virtual disk format Same format as source
VM Storage Policy Keep existing VM storage policies

	Name	Storage Compatibility	Capacity	Provisioned	Free	Type
<input type="radio"/>	iscsi_datastore_1	--	5 TB	3.66 TB	2.04 TB	VMFS
<input checked="" type="radio"/>	otv_vmfs_iscsi_1	--	499.75 GB	201.42 GB	298.33 GB	VMFS
<input type="radio"/>	vmware_swap	--	399.75 GB	130.01 GB	269.74 GB	VMFS
<input type="radio"/>	vmware_vcls	--	99.75 GB	1.41 GB	98.34 GB	VMFS

Manage Columns
Items per page 10 4 items

**Note:** The dedicated VM swap datastore will not be changed since all hosts are configured the same.

6. For network, leave the VM network unchanged. Click NEXT.



## Migrate | fpsa-asa-iometer-01

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select networks
- 5 Select vMotion priority
- 6 Ready to complete

### Select networks

Select destination networks for the virtual machine migration.

VM ORIGIN ⓘ

Migrate VM networking by selecting a new destination network for all VM network adapters attached to the same source network.

	Source Network ▼	Used By ▼	Destination Network ▼
»	IB-MGMT Network	1 VMs / 1 Network adapters	IB-MGMT Network ▾
1 item			

7. Select the recommended Schedule vMotion with high priority and click NEXT.

## Migrate | fpsa-asa-iometer-01

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select networks
- 5 Select vMotion priority
- 6 Ready to complete

### Select vMotion priority

Protect the performance of your running virtual machines by prioritizing the allocation of CPU resources.

VM ORIGIN ⓘ

- ☒ Schedule vMotion with high priority (recommended)  
vMotion receives higher CPU scheduling preference relative to normal priority migrations. vMotion might complete more quickly.
- ☐ Schedule normal vMotion  
vMotion receives lower CPU scheduling preference relative to high priority migrations. You can extend vMotion duration.

8. Review the information on the Ready to complete page and click FINISH.

## Migrate | fpsa-asa-iometer-01

- 1 Select a migration type
- 2 Select a compute resource
- 3 Select storage
- 4 Select networks
- 5 Select vMotion priority
- 6 Ready to complete

### Ready to complete

Verify that the information is correct and click Finish to start the migration.

VM ORIGIN ⓘ

Migration Type	Change compute resource and storage
Virtual Machine	fpsa-asa-iometer-01
Cluster	FlexPod
Host	fpsa-asa-esxi-02.nva.local
vMotion Priority	High
Networks	No network reassignments
Storage	otv_vmfs_iscsi_1
Disk Format	Same format as source

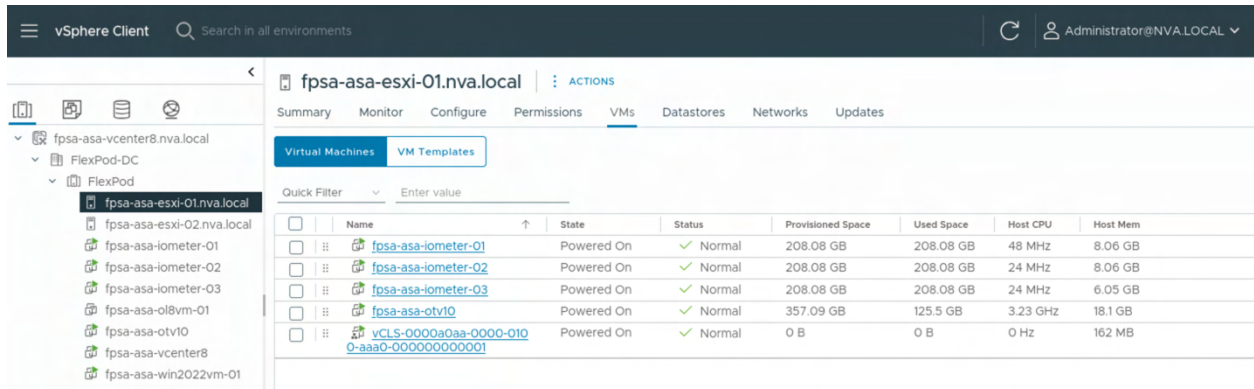
9. Check the VM migration status in Recent Tasks.
10. After migration is completed, confirm that the VM IO is still running and the VM is on the correct host and datastore.

## VMware high availability (HA)

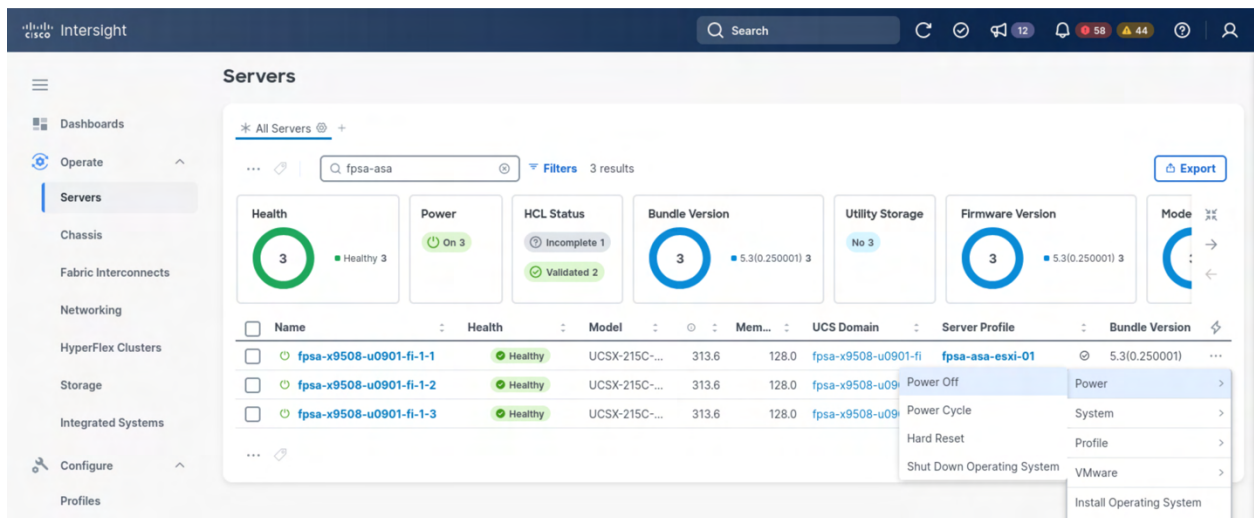
VMware vSphere high availability (HA) feature monitors ESXi host and VMs to detect failures. When a server outage is detected, VMs are restarted on a surviving host in the cluster. In addition, when it detects an operating system failure, the VM is also restarted.

To exercise VMware HA feature and confirm VMs running on a failed host is restarted, follow the steps below.

1. Login to vCenter.
2. Select a host which we will introduce a failure manually by powering off the server from Intersight.
3. Click on the host and go to its VMs tab.
4. Take note of the VMs that are running on the host.



5. Login to Intersight.
6. Navigate to Operate > Servers.
7. Provide a filter string to narrow down the list of servers.
8. Find the server that we will power off and click on the three dots on the right to select Power > Power Off.



9. Click Power Off to confirm that the selected server will be powered off.

## Power Off Server

Server 'fpsa-x9508-u0901-fi-1-1' will be Powered Off.

Cancel Power Off

- Go back to vCenter to check on the host and VM status. The powered off host shows not responding and the VMs running on that host become disconnected.

The screenshot shows the vSphere Client interface with the host 'fpsa-asa-esxi-01.nva.local' selected. The 'Virtual Machines' tab is active, displaying a list of VMs. The VMs are in a 'Disconnected' state, indicated by the 'Status' column showing a green checkmark and the word 'Normal'.

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
fpsa-asa-esxi-02.nva.local	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	0 B
fpsa-asa-iometer-01 (disconnected)	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	0 B
fpsa-asa-iometer-02 (disconnected)	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	0 B
fpsa-asa-iometer-03 (disconnected)	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	0 B
fpsa-asa-ol8vm-01	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	0 B
fpsa-asa-otv10 (disconnected)	Powered On	✓ Normal	357.09 GB	125.5 GB	0 Hz	0 B
fpsa-asa-vcenter8	Powered On	✓ Normal	0 B	0 B	0 Hz	0 B
fpsa-asa-win2022vm-01	Powered On	✓ Normal	0 B	0 B	0 Hz	0 B

- After VMware HA kicks in, the previously disconnected VMs are running again on the surviving host. The vCLS VM on the host which went down will not be restarted.

The screenshot shows the vSphere Client interface with the host 'fpsa-asa-esxi-02.nva.local' selected. The 'Virtual Machines' tab is active, displaying a list of VMs. The VMs are in a 'Powered On' state, indicated by the 'Status' column showing a green checkmark and the word 'Normal'.

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
fpsa-asa-esxi-01.nva.local	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	72 MB
fpsa-asa-iometer-01	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	72 MB
fpsa-asa-iometer-02	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	76 MB
fpsa-asa-iometer-03	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	76 MB
fpsa-asa-ol8vm-01	Powered Off	✓ Normal	264.29 GB	200 GB	0 Hz	0 B
fpsa-asa-otv10	Powered On	✓ Normal	357.09 GB	125.51 GB	0 Hz	76 MB
fpsa-asa-vcenter8	Powered On	✓ Normal	722.89 GB	722.89 GB	391 MHz	20.97 GB
fpsa-asa-win2022vm-01	Powered On	✓ Normal	264.09 GB	264.09 GB	195 MHz	51.15 GB
vCLS-0000a0aa-0000-0100-aaa0-000000000001 (disconnected)	Powered On	✓ Normal	0 B	0 B	0 Hz	164 MB

The screenshot shows the vSphere Client interface with the host 'fpsa-asa-esxi-01.nva.local' selected. The 'Virtual Machines' tab is active, displaying a list of VMs. The VMs are in a 'Powered On' state, indicated by the 'Status' column showing a green checkmark and the word 'Normal'.

Name	State	Status	Provisioned Space	Used Space	Host CPU	Host Mem
fpsa-asa-esxi-02.nva.local	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	72 MB
fpsa-asa-iometer-01	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	72 MB
fpsa-asa-iometer-02	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	76 MB
fpsa-asa-iometer-03	Powered On	✓ Normal	208.08 GB	208.08 GB	0 Hz	76 MB
fpsa-asa-ol8vm-01	Powered Off	✓ Normal	264.29 GB	200 GB	0 Hz	0 B
fpsa-asa-otv10	Powered On	✓ Normal	357.09 GB	125.51 GB	0 Hz	76 MB
fpsa-asa-vcenter8	Powered On	✓ Normal	722.89 GB	722.89 GB	391 MHz	20.97 GB
fpsa-asa-win2022vm-01	Powered On	✓ Normal	264.09 GB	264.09 GB	195 MHz	51.15 GB
vCLS-0000a0aa-0000-0100-aaa0-000000000001 (disconnected)	Powered On	✓ Normal	0 B	0 B	0 Hz	164 MB

- Go back to Cisco Intersight Server list and invoke the action to power on the server which was brought down earlier.

13. After the server is back up and connected to vCenter, some VMs are migrated back to the server by Distributed Resource Scheduler (DRS).

The steps above verifies that VMs which went down due to host issue is properly restarted on another host in the cluster by VMware HA. Detecting host and VM failure as well as restarting VM do take time. For applications that require higher availability not provided by VM restart, additional resiliency features should be implemented using operating system and application specific mechanisms.

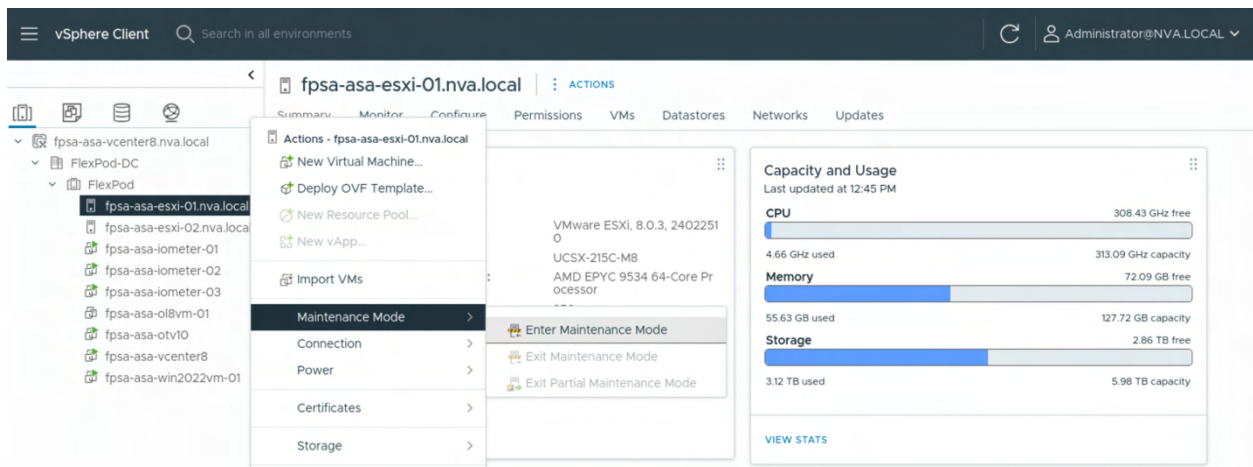
For example, for Microsoft SQL server, a virtual multi-node Windows Failover Cluster can be implemented to help with resiliency. For Oracle database, a virtual multi-node Oracle Grid Infrastructure along with Real Application Clusters (RAC) database configuration can be leveraged to further enhance availability. In these cases, proper VM anti-affinity rules should be implemented to distribute VMs among the available hosts to maximize application availability.

## UCS server firmware upgrade

The firmware of a UCS server in the VMware cluster can be upgraded without impacting the running VMs by placing the server in maintenance mode and evacuate the application VMs off the host. Follow the steps below to upgrade UCS server firmware.

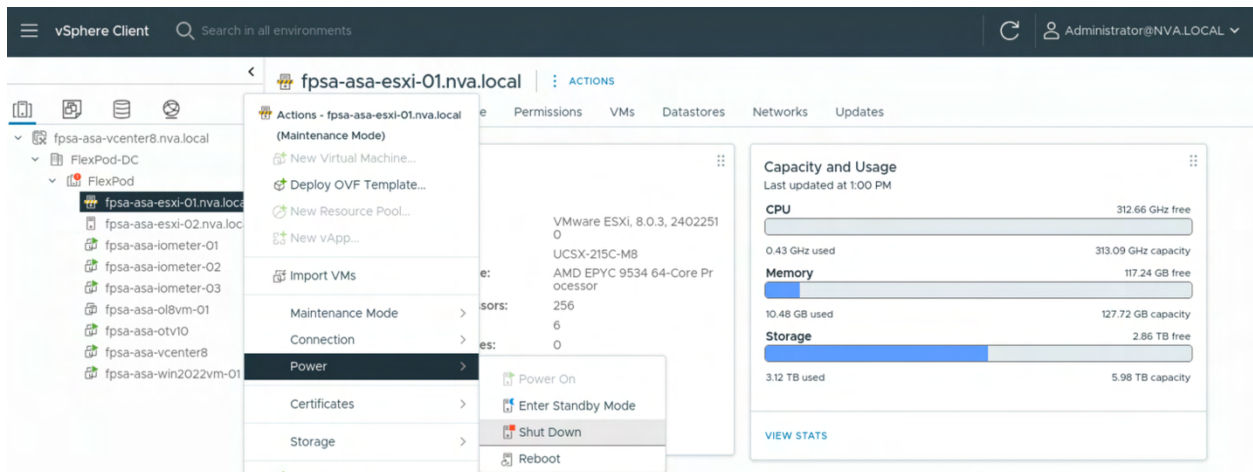
**Note:** You should check for ecosystem interoperability first before upgrading server firmware.

1. Login to vCenter.
2. From the Inventory view, right-click on the server which requires firmware upgrade. Select Maintenance Mode from the menu and then select Enter Maintenance Mode.

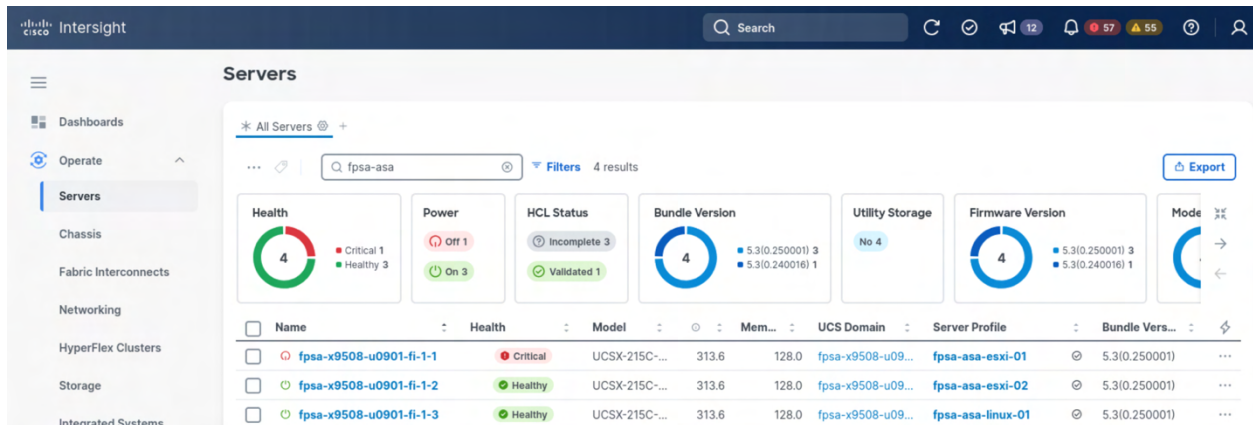


3. Click OK in the Enter Maintenance Mode dialog which will need the powered-on VMs to be migrated off the host.
4. Click OK again to proceed.
5. Migrate the VMs running on the host to other hosts in the cluster for the host to enter Maintenance Mode.
6. Right-click on the host in the Inventory view again and select Power > Shutdown Down.

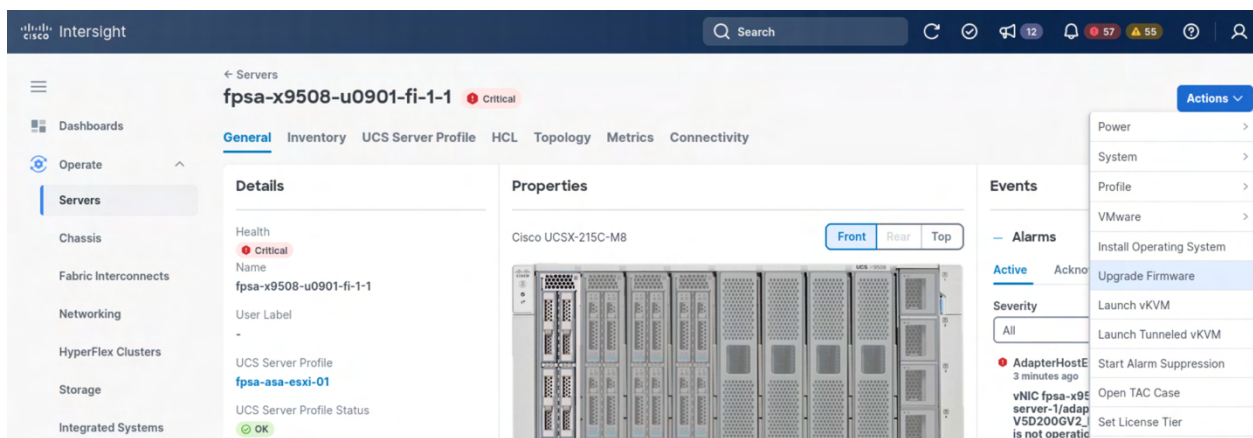




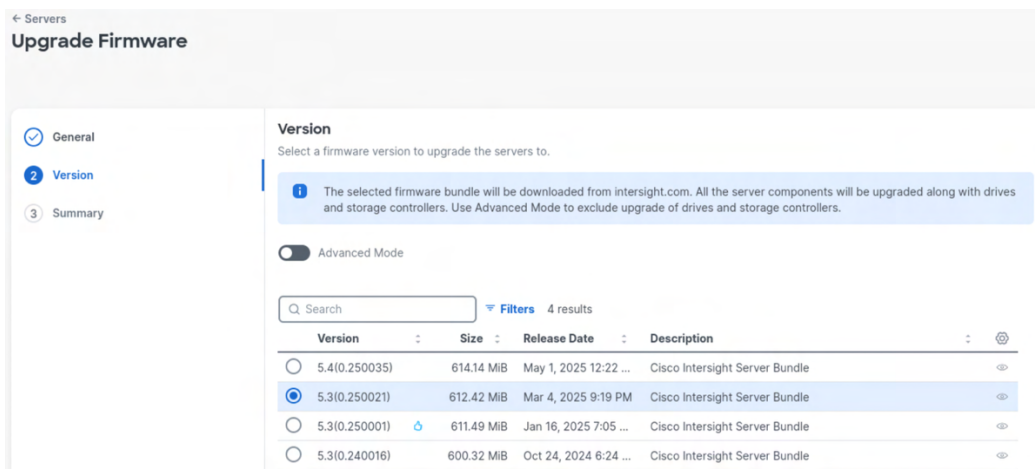
7. Provide a reason in the Shut Down Host dialog and click OK to shut down the selected host.
8. Login to Intersight.
9. Navigate to Operate > Servers and provide a filter to narrow down the list of servers as needed.



10. Click on the server link under the Name column to go to the server view.
11. Open the Actions drop-down list and select Upgrade Firmware.

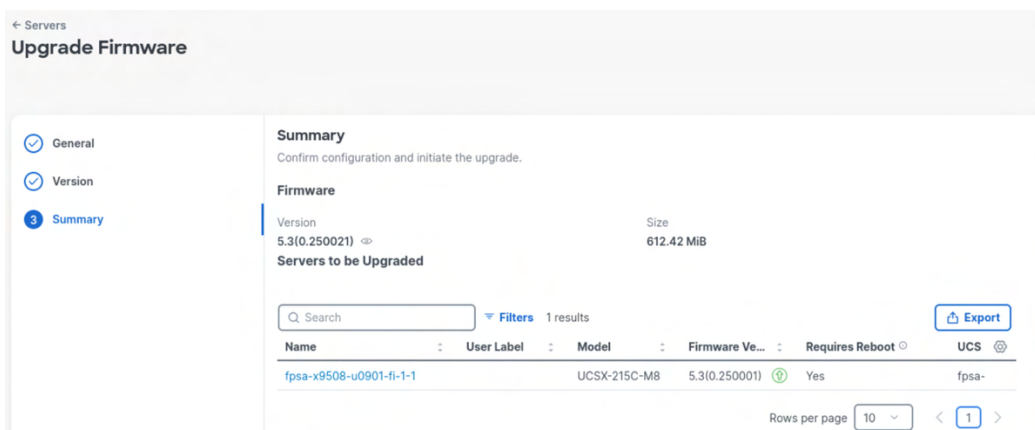


12. Click Start at the bottom of the Upgrade Firmware page.
13. Click Next to see a list of firmware versions to select from.
14. Select the desired firmware version and click Next.



**Note:** In this case, the recommended version, indicated with the thumbs up icon, is the current running firmware. We are selecting a newer firmware bundle for the purpose of documenting the firmware upgrade process.

15. Click Upgrade at the bottom of the Summary screen.



16. Select Reboot Immediately to Begin Upgrade and then click Upgrade.

## Upgrade Firmware

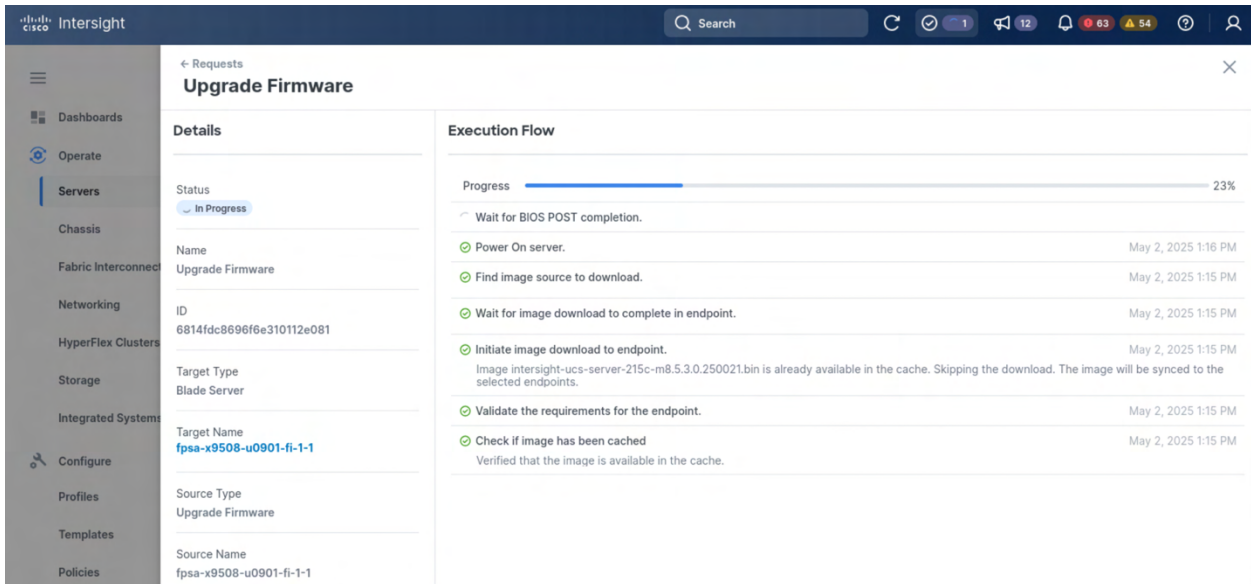
Firmware will be installed on next boot. To reboot immediately, please enable the option below.

☒ Reboot Immediately to Begin Upgrade

[Cancel](#) [Upgrade](#)

17. Wait for the firmware upgrade to complete. You can click on the Requests list at the top and click on the Upgrade Firmware request to view the upgrade progress and execution flow.





**Note:** It took around 30 minutes for the above firmware upgrade to complete.

18. After firmware upgrade is completed, go back to the server's Action drop-down list and select Power > Power On to power on the server if the server is not already powered on.
19. Login to vCenter.
20. After the server boots up and connected back to vCenter, right-click on the server in the Inventory view and select Maintenance Mode > Exit Maintenance Mode for the server to exit out of maintenance mode.
21. After DRS kicks in, some of the VMs will be migrated back onto the server.

## UCS compute node replacement and server profile move

Cisco UCS supports stateless compute where a server profile can be moved from one physical server to another due to hardware upgrade or replacement. Hosts in a FlexPod solution are typically configured to boot from SAN, which makes deploying a server profile onto a different server simple while preserving the server configurations.

For UEFI secure boot configuration with a TPM 2.0 module, an additional process of updating ESXi boot configuration is needed due to booting with a different server and TPM module. The following highlights the steps of deploying a server profile to a different physical server when using the UEFI secure boot configuration.

Before proceeding with the procedures below, it is a good idea to login to vCenter and review your running VMs and shut down VMs that are not essential, especially when the remaining hosts in the cluster are already highly utilized.

In the following example, we are moving fpsa-asa-esxi-02 server profile deployed on server fpsa-x9508-u0901-fi-1-2 to fpsa-x9508-u0901-fi-1-4.

The screenshot shows the Cisco Intersight web interface. The left sidebar has a menu with 'Operate' expanded, showing 'Servers', 'Chassis', 'Fabric Interconnects', 'Networking', 'HyperFlex Clusters', and 'Storage'. The main content area is titled 'Profiles' and has tabs for 'HyperFlex Cluster Profiles', 'UCS Chassis Profiles', 'UCS Domain Profiles', and 'UCS Server Profiles'. The 'UCS Server Profiles' tab is active. At the top right, there's a 'Create UCS Server Profile' button. Below the tabs, there's a search bar with 'fpsa-asa-esxi' and a 'Filters' button showing '3 results'. Below the search bar, there are status filters: 'Status' (1 OK, 2 OK), 'Inconsistency Reason' (No data available), 'Template Sync Status' (3 OK), and 'Target Platform' (3). A table lists the profiles:

Name	Status	Target Platform	UCS Server Template	Server	Last Update
fpsa-asa-esxi-01	OK	UCS Server (FI-Attached)	FlexPod-ASA-AMD-iSCSI-Boot	fpsa-x9508-u0901-fi-1-1	3 hours ago
fpsa-asa-esxi-02	OK	UCS Server (FI-Attached)	FlexPod-ASA-AMD-iSCSI-Boot	fpsa-x9508-u0901-fi-1-2	Apr 3, 2025 10:49 PM

The screenshot shows the Cisco Intersight web interface. The left sidebar has a menu with 'Operate' expanded, showing 'Servers', 'Chassis', 'Fabric Interconnects', 'Networking', 'HyperFlex Clusters', and 'Storage'. The 'Servers' tab is active. At the top right, there's a search bar with 'fpsa-x9508-u0901' and a 'Filters' button showing '4 results'. Below the search bar, there are several summary cards: 'Health' (4 Healthy), 'Power' (Off 1, On 3), 'HCL Status' (Incomplete 2, Validated 2), 'Bundle Version' (4, 5.3(0.250001)), 'Utility Storage' (No 4), 'Firmware Version' (4, 5.3(0.250001)), and 'Mode'. Below the cards, a table lists the servers:

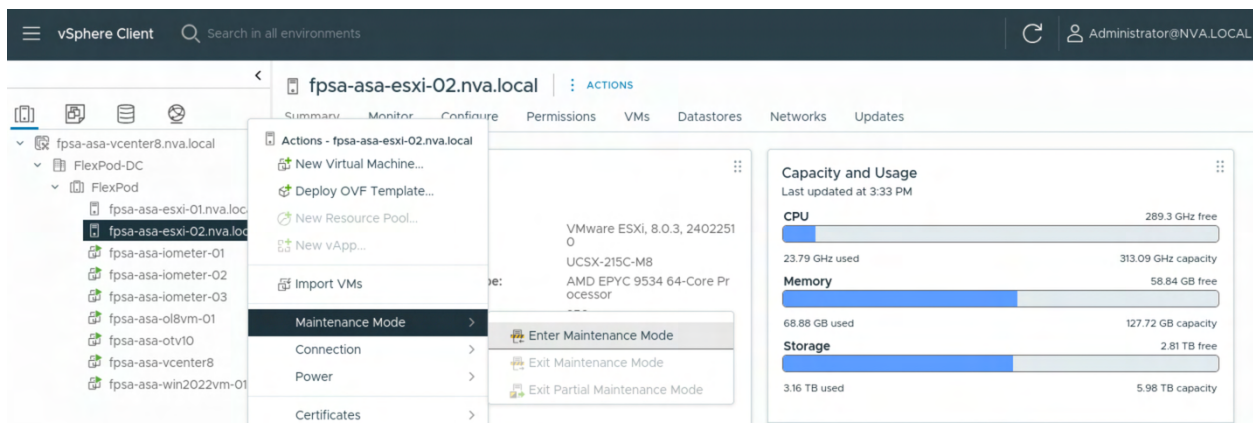
Name	Health	Model	Mem...	UCS Domain	Server Profile	Bundle Vers...
fpsa-x9508-u0901-fi-1-1	Healthy	UCSX-215C-...	313.6	128.0	fpsa-x9508-u09...	fpsa-asa-esxi-01
fpsa-x9508-u0901-fi-1-2	Healthy	UCSX-215C-...	313.6	128.0	fpsa-x9508-u09...	fpsa-asa-esxi-02
fpsa-x9508-u0901-fi-1-3	Healthy	UCSX-215C-...	313.6	128.0	fpsa-x9508-u09...	fpsa-asa-linux-01
fpsa-x9508-u0901-fi-1-4	Healthy	UCSX-215C-...	313.6	128.0	fpsa-x9508-u09...	5.3(0.250001)

## Prepare for server profile move

For this example, we are assuming that the reason behind a server profile move is to replace the underlying physical server, and the existing server is still in working condition. In that case, we can put the host into maintenance mode first to migrate workloads off the host before shutting it down. Before you put the host into maintenance mode, be sure you have the recovery key information handy or login to your ESXi host to issue the command below to obtain the information.

```
esxcli system settings encryption recovery list
```

1. Login to vCenter.
2. Right-click on the server in the Inventory and click Maintenance Mode from the menu.
3. Select Enter maintenance Mode and click OK on the Enter Maintenance Mode dialog.



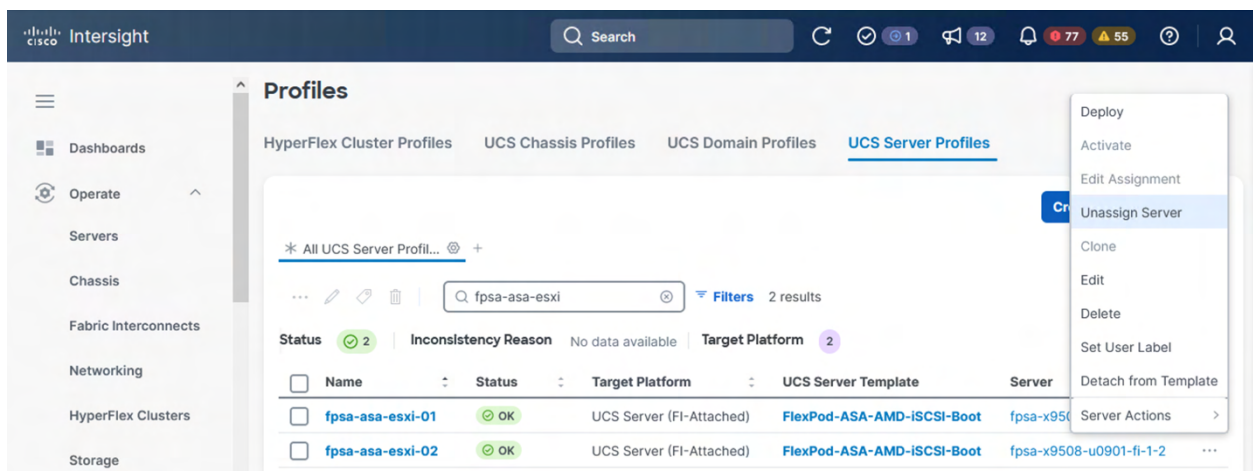
Migrate the workloads and VMs off the server and the server for it to transition into Maintenance mode. Right-click on the server in the Inventory view, select Power > Shut Down, provide a reason in the Shut Down Host dialog, and then click Ok to shut down the server.

**Note:** In the case of a failed server hardware, it is not possible to put the host into maintenance mode. However, the workload should have already been migrated to the remaining hosts by the VMware HA feature configured in the FlexPod solution.

## Update server profile assignment

After confirming that the server has Power Off status in Intersight, we are ready to update the server profile to unassign the existing server and then assign the profile to a new server.

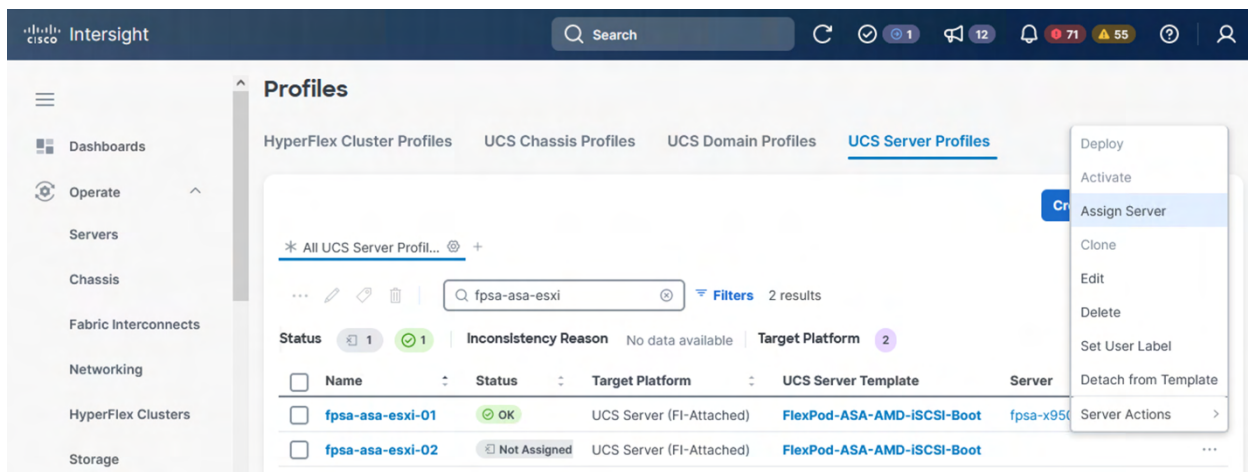
1. Login to Intersight.
2. Navigate to Profiles and optionally enter partial profile name in the Filters to narrow down the profile search.
3. Click on the dots to the right-hand side of the profile to open the action menu.



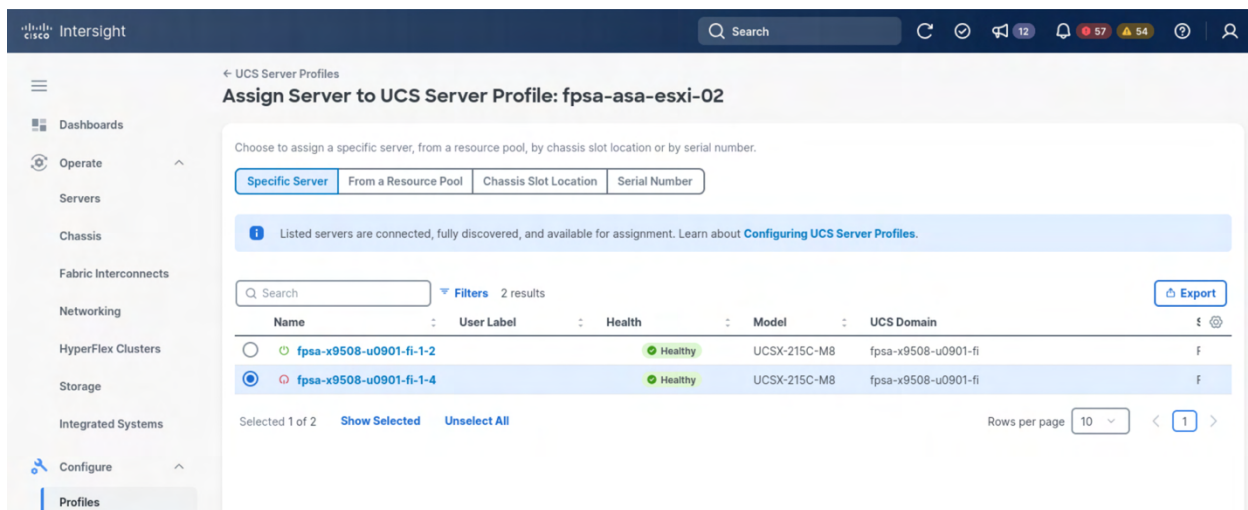
4. Select Unassign Server and confirm the action by clicking the Unassign button in the dialog.

**Note:** It will take time for the server to be unconfigured.

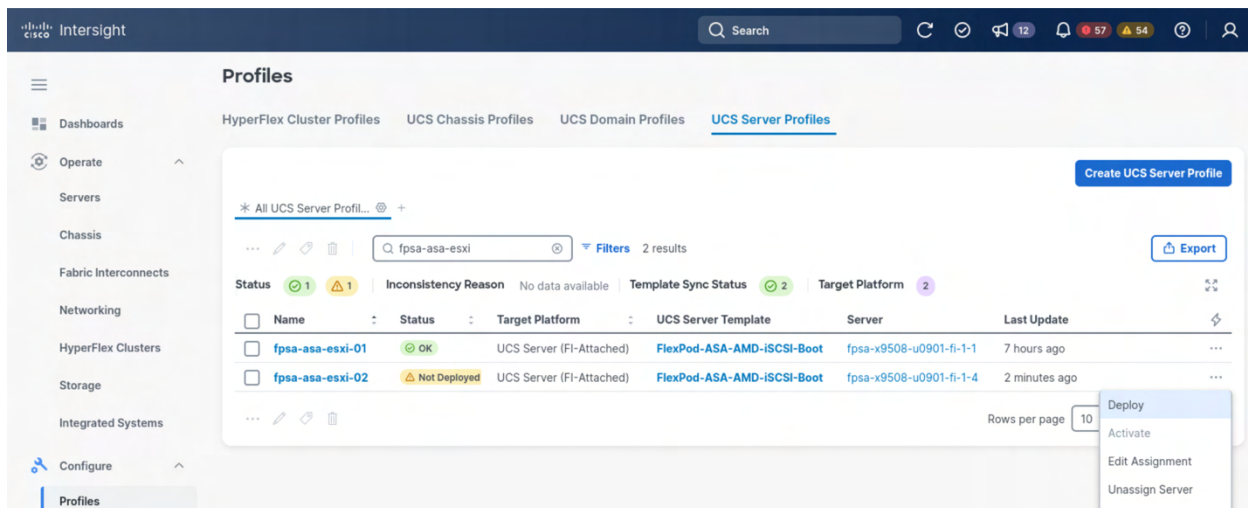
5. After unconfiguring the existing server, the profile status changes to Not Assigned.



6. Click on the dots to the right of the profile and select Assign Server.
7. Select the new server from the list and click Assign at the bottom.



8. After the server profile status changes from Validating to Not Deployed, click the dots to the right of the profile and select Deploy.



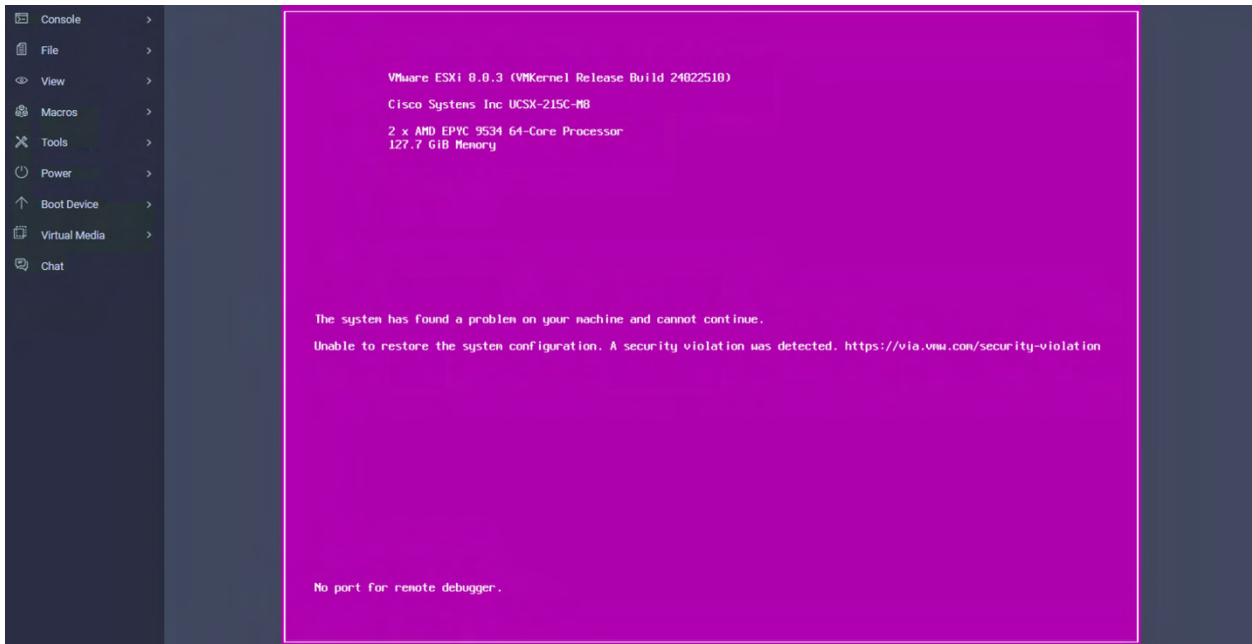
9. Check the box for Reboot Immediately to Activate in the Deploy UCS Server Profile dialog and click Deploy.

10. It will take some time for the server profile to be deployed. The Status of the server profile will go through Validating and Activating before it changes to OK.
11. Launch the server's console by clicking the dots to the right of the profile, select Server Actions > Launch vKVM.

## Resolve server profile boot issue

After a server profile is moved from one physical server to another in the following conditions, the ESXi host runs into purple screen of death (PSOD) and the ESXi host will fail to boot:

- TPM present in the node (Cisco UCS M5 family servers or above)
- Host installed with ESXi 7.0 U2 or above
- UEFI secure boot mode is in effect



The following error message is seen on the PSOD screen.

```
The system has found a problem on your machine and cannot continue.  
Unable to restore system configuration. A security violation was detected.  
https://via.vmw.com/security-violation
```

To resolve the server boot issue, follow the steps below to apply the previously saved encryption recovery information.

1. Select Power on the vKVM menu, select Reset System, click Confirm for the action.
2. Stop the ESXi boot sequence by pressing Shift + O when you see the ESXi boot screen.





3. Add the recovery key by adding the following boot option, and press Enter to continue the boot process:

```
encryptionRecoveryKey=<recovery_key>
```

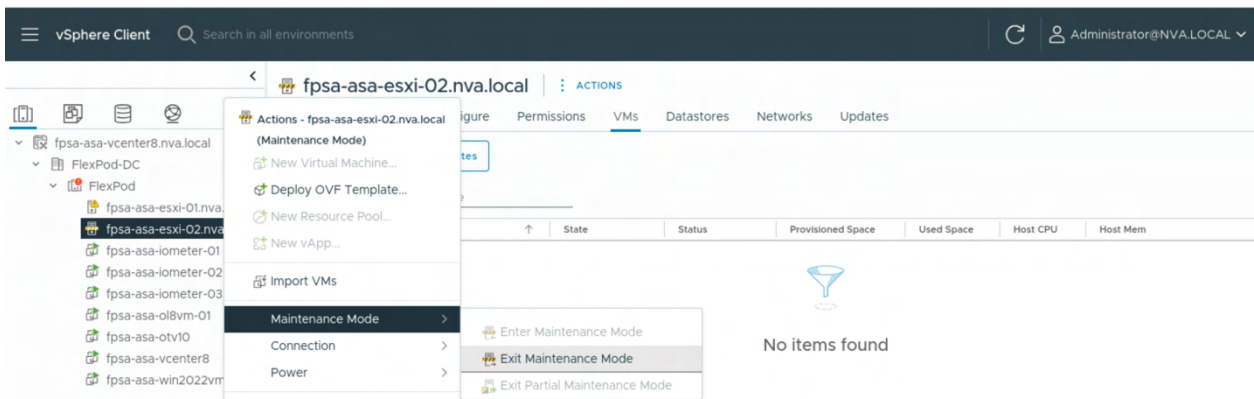
After the ESXi host is booted, login as root user.

Issue the following command to persist the encryption recovery key boot option for future boot to succeed without needing to provide it again.

```
/sbin/auto-backup.sh
```

Login to vCenter.

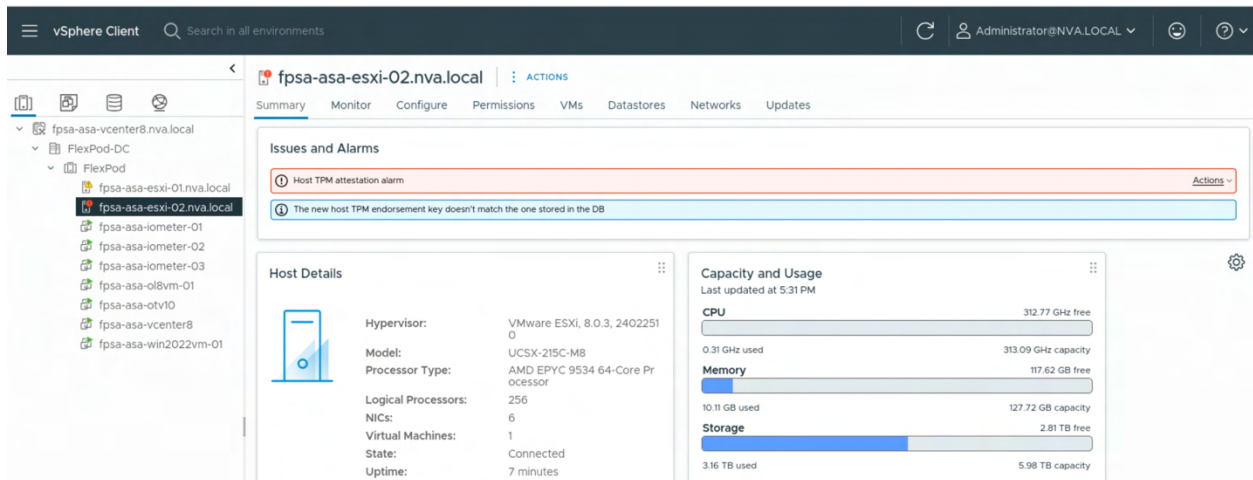
Right-click on the server in the Inventory, select Maintenance Mode > Exit Maintenance Mode.



vCenter reports the following two issues in the Summary page of the host.

- Host TPM attestation alarm
- The new host TPM endorsement key doesn't match the one stored in the DB.





**Note:** Refer to Appendix A which includes sections that follow the two Broadcom KB articles below to resolve the issues.

[Broadcom KB 369525 - Host TPM Attestation Alarm present in the vSphere UI](#)

[Broadcom KB 316512 - The new host TPM endorsement key doesn't match the one stored in the DB](#)

## Conclusion

FlexPod is a predesigned, best-practice, data center architecture that is built on the cloud-scale technologies offered by Cisco and NetApp, enabling management synergies across the complete IT infrastructure environment. FlexPod is an optimal platform for both virtual infrastructure and bare-metal deployments of enterprise applications and databases.

FlexPod SAN solutions with Cisco UCS X-Series Direct and the new NetApp ASA A-Series storage systems deliver medium-scale SAN solutions ideal for VMware virtual infrastructure and highly transactional mission-critical applications such as Oracle and Microsoft SQL Server databases.

NetApp ASA A-series systems offer simplicity, reliability, availability, and serviceability to keep your data always available. They also provide comprehensive data management and data protection capabilities for your enterprise applications with industry-leading SAN-optimized ONTAP software.

The validated FlexPod SAN solution with Cisco UCS X-Series Direct and NetApp ASA storage provides comprehensive instructions for deploying infrastructure, configuring and testing Microsoft SQL Server and Oracle databases. Additionally, it verifies critical use cases, failure scenarios, and life-cycle management activities to ensure the robustness of the solution and maintain the FlexPod SAN solution's currency.

## Appendix

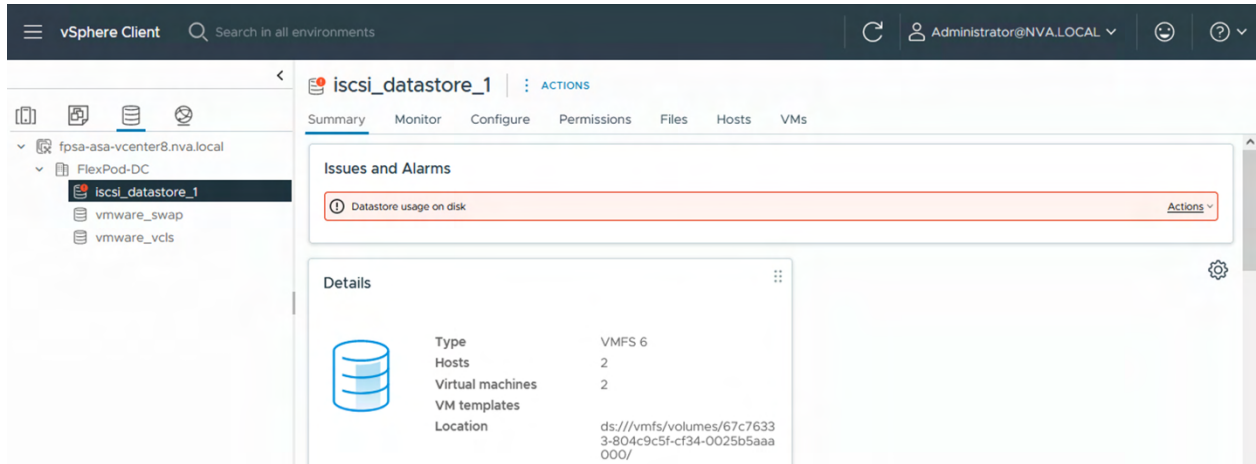
### Appendix A: Resolve issues encountered during solution validation

This appendix provides information to address the below issues which were encountered during the solution validation.

- High usage on datastore issue
- Host TPM attestation alarm in vCenter issue
- The new host TPM endorsement key doesn't match the one stored in the DB issue

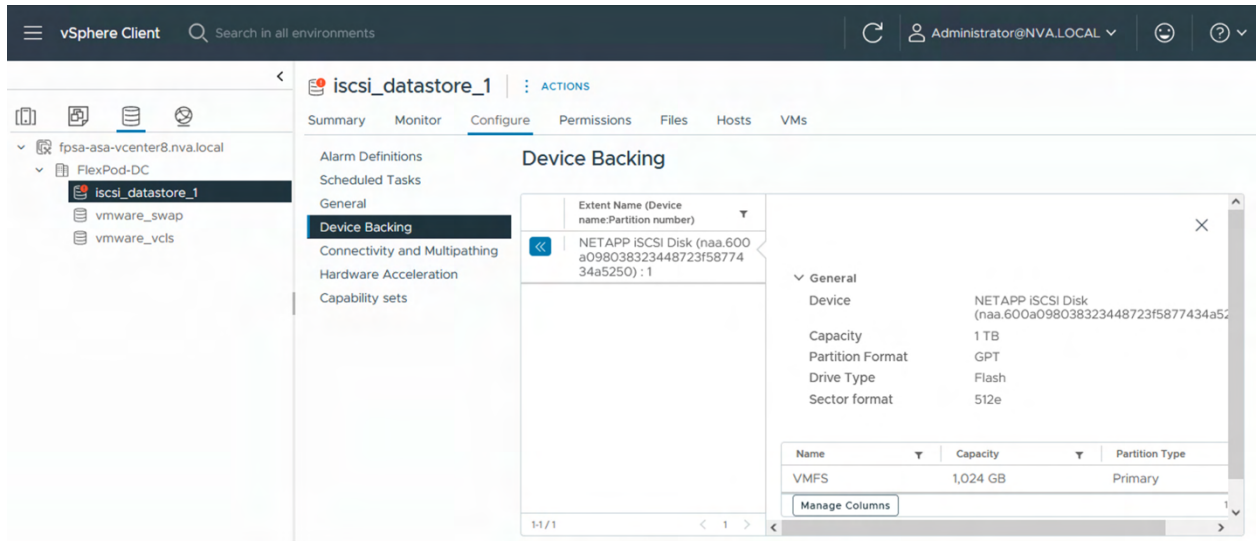
## High usage on datastore issue

When a datastore utilization gets too high, vCenter will report an issue for the datastore usage on disk as shown in the screenshot below.



To address this issue, follow the steps below to increase the LUN size in storage and extend the datastore size in vCenter.

1. From vCenter Storage view, select the datastore reporting the issue.
2. In the center pane, click Configure tab and select Device Backing to see the storage disk information.



3. Login to ONTAP cluster CLI.
4. Use the lun command and lun mapping command to correlate and identify the LUN which needs to be resized. In this case, it is the `fpsa_asa_vmware_cluster_datastore_1` LUN.

```
fpsa-a50-u0909::> lun show
Vserver  Path                                     State  Mapped  Type      Size
-----
svml1    fpsa_asa_esxi_01_boot_1                online mapped  vmware    128GB
svml1    fpsa_asa_esxi_02_boot_1                online mapped  vmware    128GB
svml1    fpsa_asa_vmware_cluster_datastore_1    online mapped  vmware    1TB
svml1    fpsa_asa_vmware_swap_1                 online mapped  vmware    200GB
svml1    fpsa_asa_vmware_vcls_1                 online mapped  vmware    100GB
5 entries were displayed.
```

```
fpsa-a50-u0909:> lun show -m
```

Vserver	Path	Igroup	LUN ID	Protocol
svml	fpsa_asa_esxi_01_boot_1	FlexPod-ASA-esxi-01-boot-iscsi	0	iscsi
svml	fpsa_asa_esxi_02_boot_1	FlexPod-ASA-esxi-02-boot-iscsi	0	iscsi
svml	fpsa_asa_vmware_cluster_datastore_1	FlexPod-ASA-esxi-cluster-iscsi	1	iscsi
svml	fpsa_asa_vmware_swap_1	FlexPod-ASA-esxi-cluster-iscsi	2	iscsi
svml	fpsa_asa_vmware_vcls_1	FlexPod-ASA-esxi-cluster-iscsi	3	iscsi

5 entries were displayed.

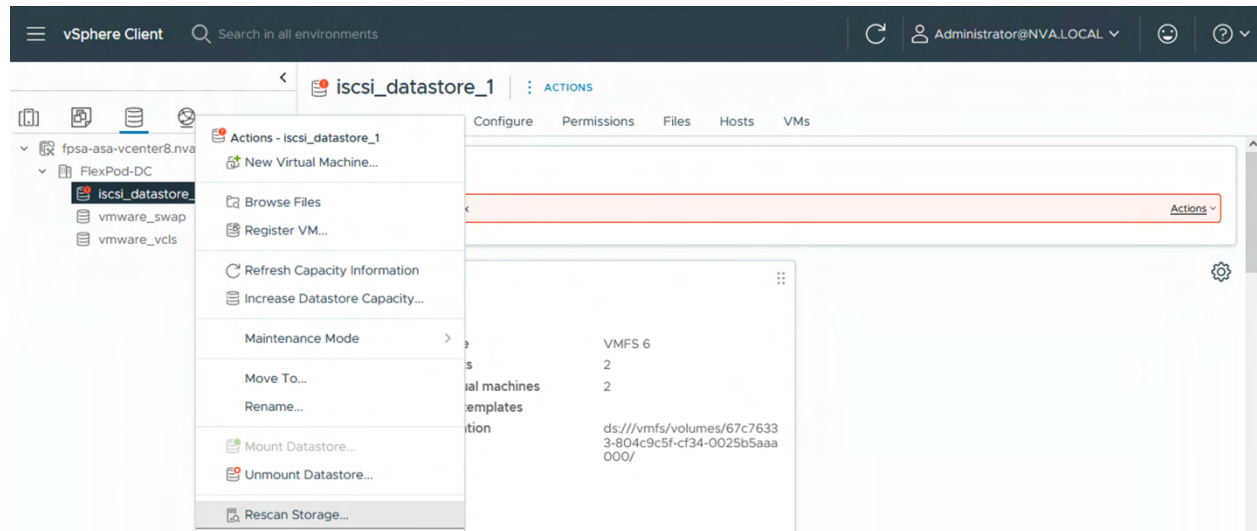
- Re-size the LUN with the lun resize command in the example to increase the LUN size and check the LUN size afterwards.

```
fpsa-a50-u0909:> lun resize -path fpsa_asa_vmware_cluster_datastore_1 -size 5T
```

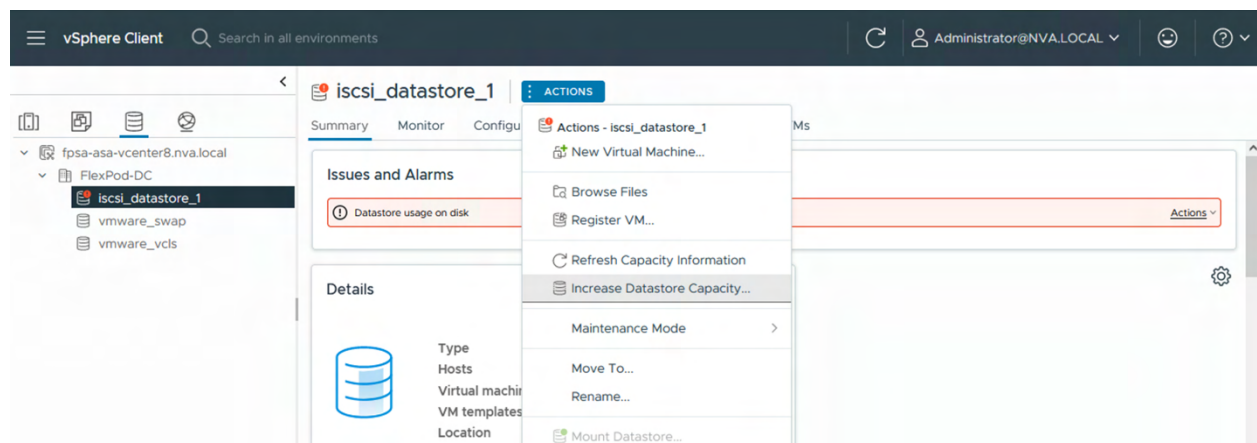
```
fpsa-a50-u0909:> lun show -path fpsa_asa_vmware_cluster_datastore_1
```

Vserver	Path	State	Mapped	Type	Size
svml	fpsa_asa_vmware_cluster_datastore_1	online	mapped	vmware	5TB

- Login to vCenter.
- Go to Storage view.
- Right-click on the datastore, click Rescan storage, and click OK to proceed.



- In the center pane, click ACTIONS, and select Increase Datastore Capacity.



- Select the device listed, confirm the increased capacity, and click NEXT.

## Increase Datastore Capacity

### 1 Select Device

### 2 Specify Configuration

### 3 Ready To Complete

## Select Device

Select disk/LUN to be used to increase size of the datastore.

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Exp
	NETAPP ISCSI Disk (naa.600a098038323448723f5877434a5250)	1	5.00 TB	Supported	Flash	512e	Yes

Manage Columns

Export

1 item

11. Review the disk layout change and click NEXT.

12. Review the settings for increasing the size of the datastore and click FINISH.

13. Confirm that the Datastore usage on disk warning is automatically removed after a few seconds.

## Host TPM attestation alarm in vCenter issue

After assigning an existing server profile to a different physical server with TPM, the following message appears in the vCenter host summary page and the Issues and Alarms for the host under its Monitor tab.

Host TPM attestation alarm

To resolve the host TPM attestation alarm, following the steps described in [Broadcom KB 369525: Host TPM Attestation Alarm present in the vsphere UI](#).

1. Right-click on the host in the Inventory view and select Connection > Disconnect.

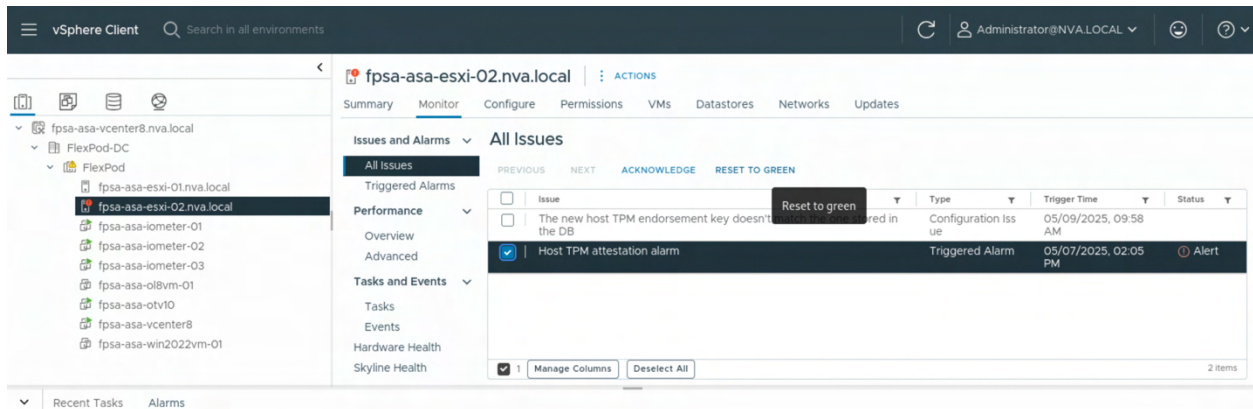
Wait for the host to disconnect fully and show disconnected status.

Right-click on the host again and select Connection > Connect.

Wait for the host to reconnect.

Got to the Monitor tab for the host and select All Issues.

Locate and select the TPM attestation alarm and click Reset to Green.



## The new host TPM endorsement key doesn't match the one stored in the DB issue

After assigning an existing server profile to a different physical server with TPM, the following message appears in the vCenter host summary page and the Issues and Alarms for the host under its Monitor tab.

The new host TPM endorsement key doesn't match the one stored in the DB

To resolve this issue, follow the steps described in [Broadcom KB 316512: The new host TPM endorsement key doesn't match the one stored in the DB](#).

1. Take an offline snapshot of the vCenter server appliance (vCSA).

Right click on the affected host and place it into Maintenance Mode. Migrate VMs off the host as needed. SSH to the vCSA as root user.

Launch BASH shell after login.

```
(root@fpsa-asa-vcenter8) Password:
Connected to service

* List APIs: "help api list"
* List Plugins: "help pi list"
* Launch BASH: "shell"

Command> shell
Shell access is granted to root
root@fpsa-asa-vcenter8 [ ~ ]#
```

Stop the vpxd service.

```
root@fpsa-asa-vcenter8 [ ~ ]# service-control --stop vmware-vpxd
Operation not cancellable. Please wait for it to finish...
Performing stop operation on service vpxd...
Successfully stopped service vpxd
```

Backup the VPX\_HOST table.

```
root@fpsa-asa-vcenter8 [ ~ ]# /opt/vmware/vpostgres/current/bin/pg_dump -U postgres -t VPX_HOST
VCDB > /var/core/VPX_HOST.sql
```

Run the command below to get the specific host ID information. Update the parameter for dns\_name to the FQDN host name of the host with the issue.

```
root@fpsa-asa-vcenter8 [ ~ ]# /opt/vmware/vpostgres/current/bin/psql -U postgres -d VCDB -h
localhost -c "select id,dns_name,endorsement_key,attestation_identity_key from VPX_HOST WHERE
dns_name = 'fpsa-asa-esxi-02.nva.local';"
```

id	dns_name
endorsement_key	
	attestation_identity_key
-----+-----+	
-----	
-----	
-----	
-----	
-----+-----	
-----	
-----	
-----	
-----	
1022   fpsa-asa-esxi-02.nva.local	
AToAAQALAAmAsgAgg3GXZ0SEs/gakMyNRqXXJP1S124GUgtk8qHaGzMUaaOABgCAAEMAEagAAAAAAEApbTBQdG5b9L4BUtlgxyRgd8zGLZFHLsCb	
ercIoYjPfrSdlMH7bhKljtbOVu+RvQpwyIM+qZNcS7rAcEjb/Oo/yejsfAdvzh7d+nmwzsVDMTOxWstsEO7ePIX4PfoVef28vzqSlSrKfcbre4QAGApJP8ilu7q9u9g2GXr8kl0tSqHVG9hDKSK7	
dq9NSGUlKY4EziXRCaJEQnddIZ4K/8ZefXvhaUdv52C4Tf8HuM0Pvf4me+QSlyJzx/gCNal/5u9tJ4FKyNABzxo4sLBkHcmLCgkNEv0meEWR9ZwRLAVDhKTXWi4Q3ZxBN8T3jrHMU/VLUtUoalGS7f	
9/2lyy//Q==	
ARGAAQALAAUAacgAABAAFAALCAAAAQABAQC31+tRH2a3exfv+vXdMRqRN95p94aHTSy25jpLoJWvpkBhI+EQNekZ3t/107ZQd4J2dY/Ijtk6OOSf/puoN31F/o/LOejquinRNBzf	
+V+0231atLTuDxziv18kugYI5LoPuYYtwG1/jCjQ0ZK6kUwPY7kiJVoeYyy0hgQ0s+dtahsBudOKkuRTF0LSgrFGU80f7YJdWQmAkvKgcfqGRShI0guZl4O5xtLkEKSLNeUpHiXTpjdwncTgyFkI88	
WdjAvfaDhe0hWHbEkHeVa60VEKahxJO49kJJGEawf8R7otDPIn8tN0apYkU3X3l7UUozpxMcoPp3yHMCpVGWfwQjEP	
(1 row)	

Run the command below to clear the key value information. Replace the id with the id you obtained from the above step for your host.

UPDATE 1

Confirm the key value for the affected host has been cleared. Again, replace the id parameter below with the one you obtained for your host.

id	dns_name	endorsement_key	attestation_identity_key
1022	fpsa-asa-esxi-02.nva.local		

```
root@fpasa-asa-vcenter8 [ ~ ]# service-control --start vmware-vpxd
Operation not cancellable. Please wait for it to finish...
Performing start operation on service vpxd...
Successfully started service vpxd
```



Disconnect and then reconnect the affected host in vCenter.

Ensure that the new key value information is stored in the column of the affected host.

```
root@fpsa-asa-vcenter8 [ ~ ]# /opt/vmware/vpostgres/current/bin/psql -U postgres -d VCDB -h localhost -c "select ID,DNS_NAME,endorsement_key,attestation_identity_key from VPX_HOST where id = 1022;"
```

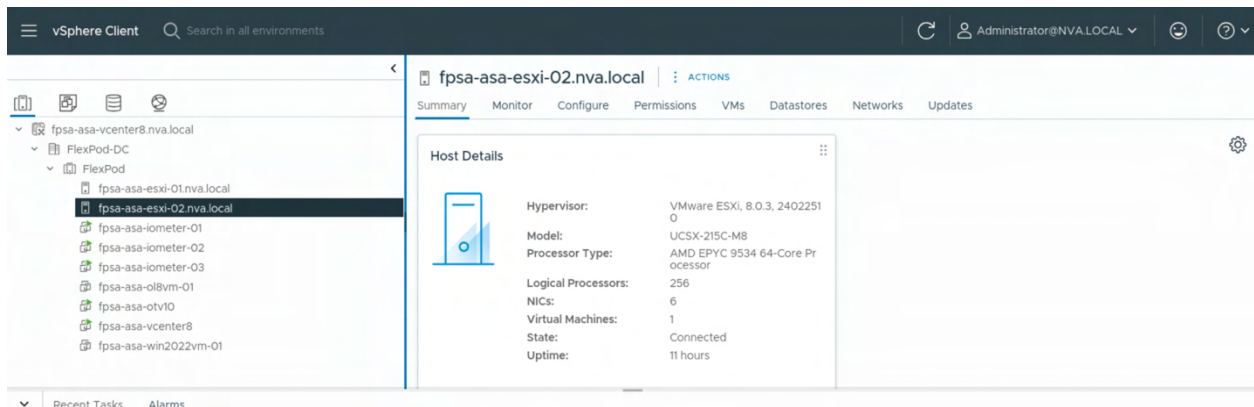
Example output:

id	dns_name	endorsement_key	attestation_identity_key
1022	fpsa-asa-esxi-02.nva.local	AToAAQALAAAMsgAgg3GXZ0SEs/gakMyNRqXXJP1S124GUgtk8qHaGzMUaaoABgCAAEAEAgAAAAAAEAtxj7U09GhChyExrUbwQihntjIzEXch6/ulm21Xe2q5BPWBj/VcORCD5yH8h0Rn+fwf7icxR4/QLBliQklvN0sXpw/Dd9Cb9j7XbRoACM8Dc9wxUQ70ozq2VZOx8wAC4YIXqtTzfCKUusyZISfzD023hjIsZbtitS4wyI0VMQqNPJwop54dPHsmknWBWPpyx8V20fk4jGH0Xjy3RsuudbXdtYdvjc51Zr0DxxZNC8041HuhHu0gMp912rGqeyV1iyQURYgMh4adcfv+w04BW1WCRZqNP12popqXYGwsMavTml22W/nkfjUOpAFWwYtidoEJul4Yj7j00BX2/djc8EaDQ==	ARgAAQALAAUAUcgAAABAAFAALCAAAAQABAQD16lWferVEFZ67cr1ggQ2LftuaA4PNNSoOiQjCBYBPExYHq63Ig3+PF4jKutgLaGzC6MGDPqrmT3Xdgr7GqbdWZTheZI348gVguw98yUlk+Co9oqirmpZTxi5DRH2HmLkFaOJsqbSGjbrH5DtIFjxYHtbAnibTiJBRfA04PwQgUKY/1fti0i5xGNXkT6Bf+wHiXEiQATBa4KjoL4q/66QI/ARUmaxVFapb/qhKHY0HC2vpjB3n4AkZFOatsy11v27Jq+3Yi05D0pBvyaE8yWVpcjveMT2Uf/RrdPNj/+MMaNybOGowdZUwdqlD/h3qd6ppz+vGuCLyhQ5wHJ/sTW6X

(1 row)

Right-click on the host in Inventory view and select Maintenance Mode > Exit Maintenance Mode to bring the host out of maintenance mode.

Check to confirm that the alarm is no longer present in the Summary view of the host.



SSH into the ESXi host to list the encryption recovery key and save it for future usage.

```
[root@fpsa-asa-esxi-02:~] esxcli system settings encryption recovery list
```

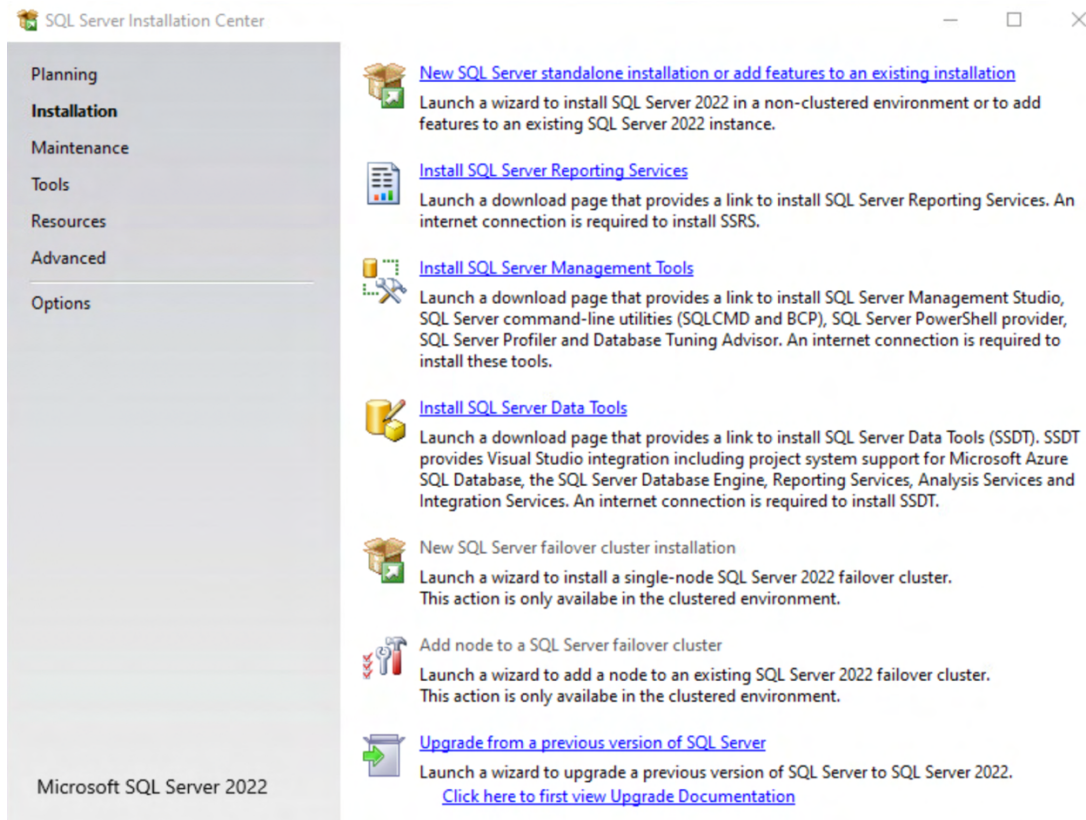
Recovery ID	Key
-----	---

## Appendix B: Install SQL server and SQL Server Management Studio on Windows

### Install SQL Server 2022

To install SQL Server 2022 on Windows Server 2022, follow the steps below.

1. Edit the Windows Server 2022 VM setting and browse to select the downloaded ISO image for the CD/DVD device.
2. From the Windows VM, click on the mapped DVD drive and run the setup program.



3. Review the linked information under the Planning menu to make sure that the hardware and software requirements are met.
4. Select New SQL Server standalone installation listed under the Installation menu to start the installation process.
5. Select the Edition that you would like to install and provide license if required and click Next.

**Note:** For the validation, we installed the Enterprise edition.

6. Check the box to accept the license terms and then click Next.
7. Check the box to use Microsoft Update to check for updates and then click Next.
8. Select the required features, e.g. Select All, and use default installation location, and then click Next.
9. Configure instance, e.g. use the default Instance ID, and then click Next.

SQL Server 2022 Setup

## Instance Configuration

Specify the name and instance ID for the instance of SQL Server. Instance ID becomes part of the installation path.

Edition  
License Terms  
Global Rules  
Microsoft Update  
Product Updates  
Install Setup Files  
Install Rules  
Feature Selection  
Feature Rules  
**Instance Configuration**  
PolyBase Configuration  
Server Configuration  
Database Engine Configuration  
Analysis Services Configuration  
Integration Services Scale Out ...

☒ Default instance

☐ Named instance: \*

Instance ID:

SQL Server directory: C:\Program Files\Microsoft SQL Server\MSSQL16.MSSQLSERVER

Analysis Services directory: C:\Program Files\Microsoft SQL Server\MSAS16.MSSQLSERVER

Installed instances:

Instance Name	Instance ID	Features	Edition	Version
---------------	-------------	----------	---------	---------

10. Accept the default ports for PolyBase configuration and click Next.

11. For Server Configuration, check to Grant Perform Volume Maintenance Tasks privilege to SQL Server Database Engine Service. Click Next.

SQL Server 2022 Setup

## Server Configuration

Specify the service accounts and collation configuration.

Edition  
License Terms  
Global Rules  
Microsoft Update  
Product Updates  
Install Setup Files  
Install Rules  
Feature Selection  
Feature Rules  
Instance Configuration  
PolyBase Configuration  
**Server Configuration**  
Database Engine Configuration  
Analysis Services Configuration  
Integration Services Scale Out ...  
Integration Services Scale Out ...  
Feature Configuration Rules  
Ready to Install  
Installation Progress  
Complete

Service Accounts Collation

Microsoft recommends that you use a separate account for each SQL Server service.

Service	Account Name	Password	Startup Type
SQL Server Agent	NT Service\SQLSERVERA...		Manual
SQL Server Database Engine	NT Service\MSSQLSERVER		Automatic
SQL Server Analysis Services	NT Service\MSSQLServe...		Automatic
SQL Server Integration Services 16.0	NT Service\MsDtsServer...		Automatic
SQL Server Integration Services Sc...	NT Service\SSISScaleOut...		Automatic
SQL Server Integration Services Sc...	NT Service\SSISScaleOut...		Automatic
SQL Server PolyBase Engine	NT AUTHORITY\NETWO...		Automatic
SQL Server PolyBase Data Movem...	NT AUTHORITY\NETWO...		Automatic
SQL Server Launchpad	NT Service\MSSQLLaun...		Automatic
SQL Full-text Filter Daemon Launc...	NT Service\MSSQLFDLa...		Manual
SQL Server Browser	NT AUTHORITY\LOCAL ...		Disabled

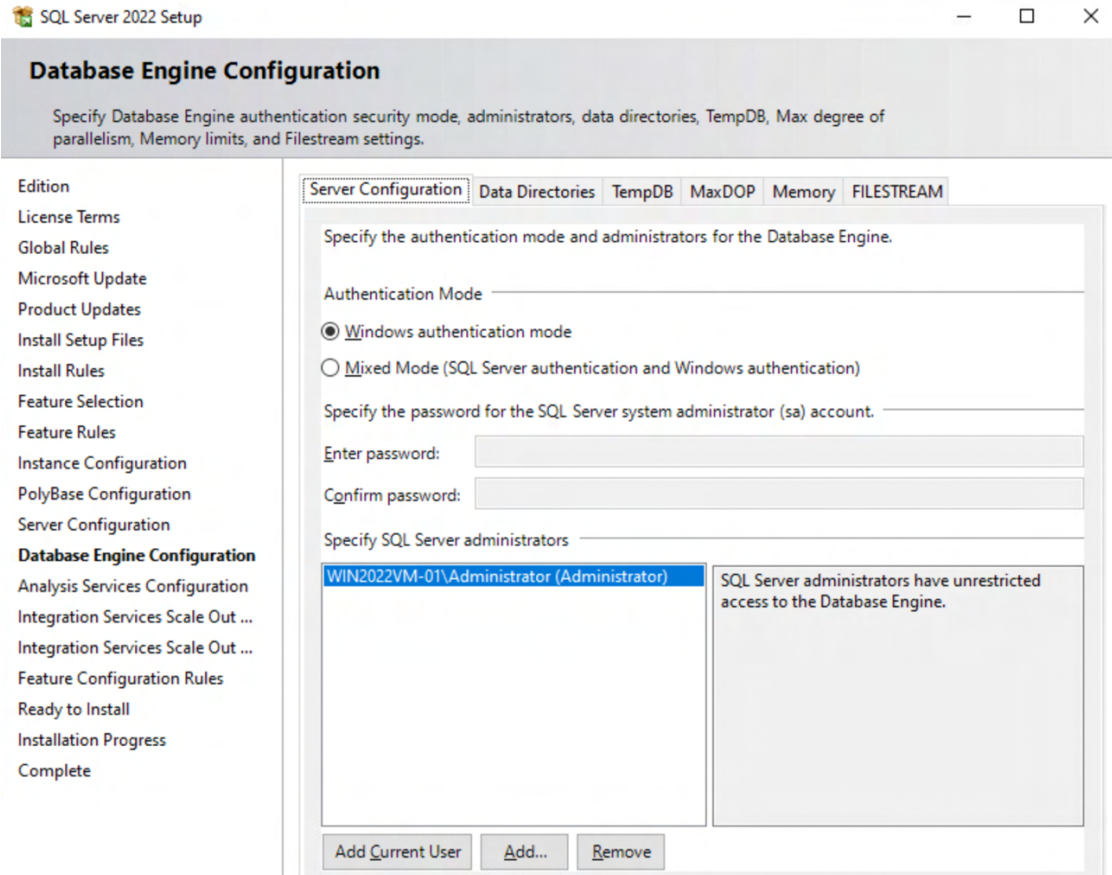
☒ Grant Perform Volume Maintenance Tasks privilege to SQL Server Database Engine Service

This privilege enables instant file initialization by avoiding zeroing of data pages. This may lead to information disclosure by allowing deleted content to be accessed.

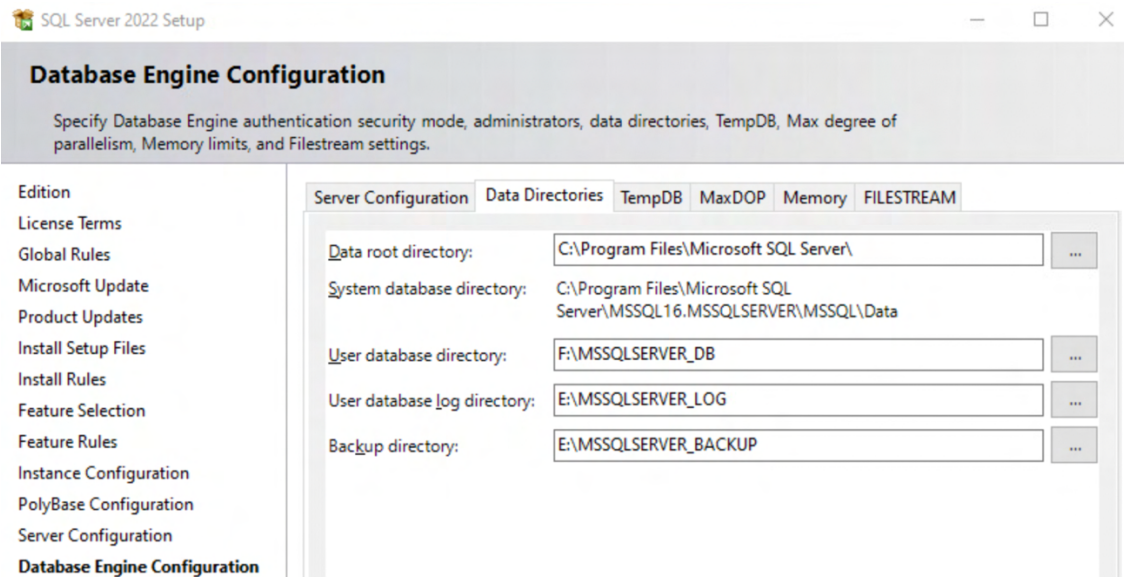
[Click here for details](#)

**Note:** This privilege enables instance file initialization by avoiding zeroing of data pages.

12. In Database Engine Configuration dialog, provide the SQL Database administrator by Add Current User or Add in the Server Configuration tab.



13. Under the Data Directories tab, provide the database, log, and backup directories.



14. Under the TempDB tab, you can update the number of files to a higher number later if you notice contention on the TempDB resource later. Click Next.



SQL Server 2022 Setup

## Database Engine Configuration

Specify Database Engine authentication security mode, administrators, data directories, TempDB, Max degree of parallelism, Memory limits, and Filestream settings.

Edition  
License Terms  
Global Rules  
Microsoft Update  
Product Updates  
Install Setup Files  
Install Rules  
Feature Selection  
Feature Rules  
Instance Configuration  
PolyBase Configuration  
Server Configuration  
**Database Engine Configuration**  
Analysis Services Configuration  
Integration Services Scale Out ...  
Integration Services Scale Out ...  
Feature Configuration Rules  
Ready to Install  
Installation Progress  
Complete

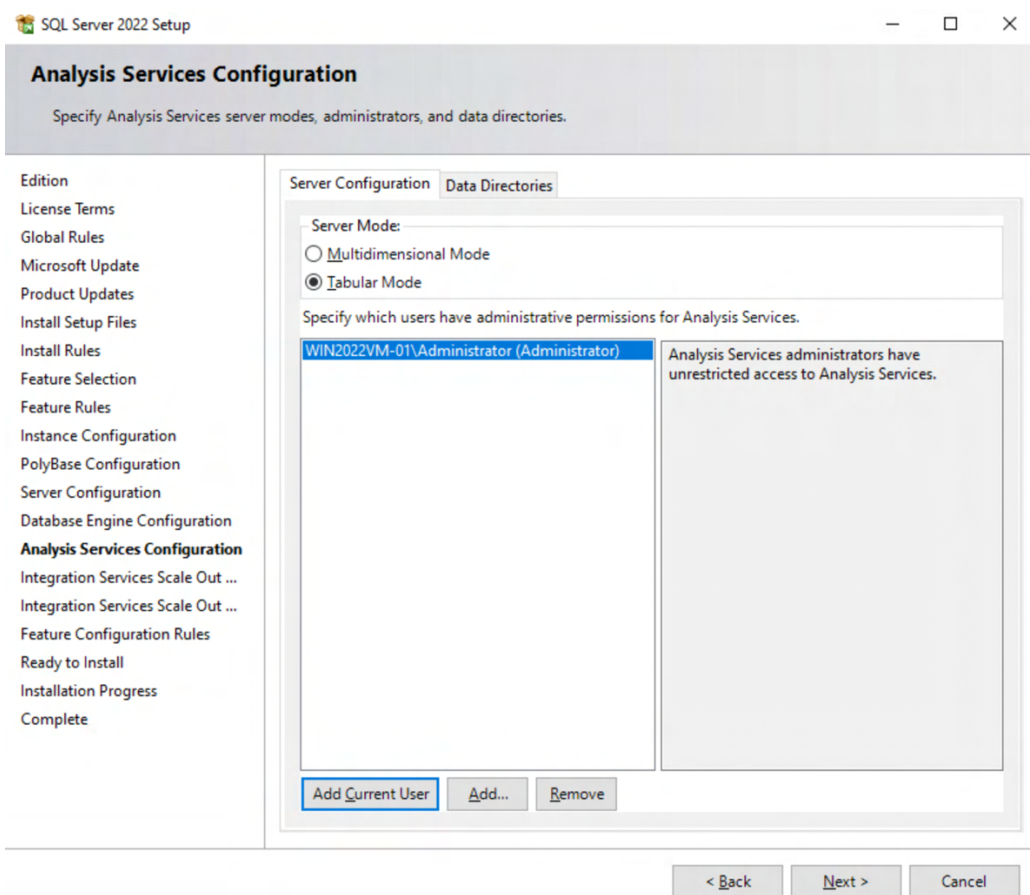
Server Configuration
Data Directories
TempDB
MaxDOP
Memory
FILESTREAM

TempDB data files: tempdb.mdf, tempdb\_mssql\_#.ndf  
Number of files: 8  
Initial size (MB): 8 Total initial size (MB): 64  
Autogrowth (MB): 64 Total autogrowth (MB): 512  
Data directories: C:\Program Files\Microsoft SQL Server\MSSQL16.MSSQLSERVER  
TempDB log file: templog.ldf  
Initial size (MB): 8 Setup could take longer with large initial size.  
Autogrowth (MB): 64  
Log directory: C:\Program Files\Microsoft SQL Server\MSSQL16.MSSQLSERVER

Add...  
Remove

< Back
Next >
Cancel

15. For the Analysis service configuration, provide the user(s) that should have the administrative permissions for the Analysis service. Click Next.



Click Next on the Integration Services Scale Out configuration – Master Node screen.



SQL Server 2022 Setup

## Integration Services Scale Out Configuration - Master Node

Specify the port number and security certificate for the Scale Out Master node.

Edition  
License Terms  
Global Rules  
Microsoft Update  
Product Updates  
Install Setup Files  
Install Rules  
Feature Selection  
Feature Rules  
Instance Configuration  
PolyBase Configuration  
Server Configuration  
Database Engine Configuration  
Analysis Services Configuration  
**Integration Services Scale Ou...**  
Integration Services Scale Out ...  
Feature Configuration Rules  
Ready to Install  
Installation Progress  
Complete

Specify a port number that the master node uses to communicate with the worker nodes.

Port Number:

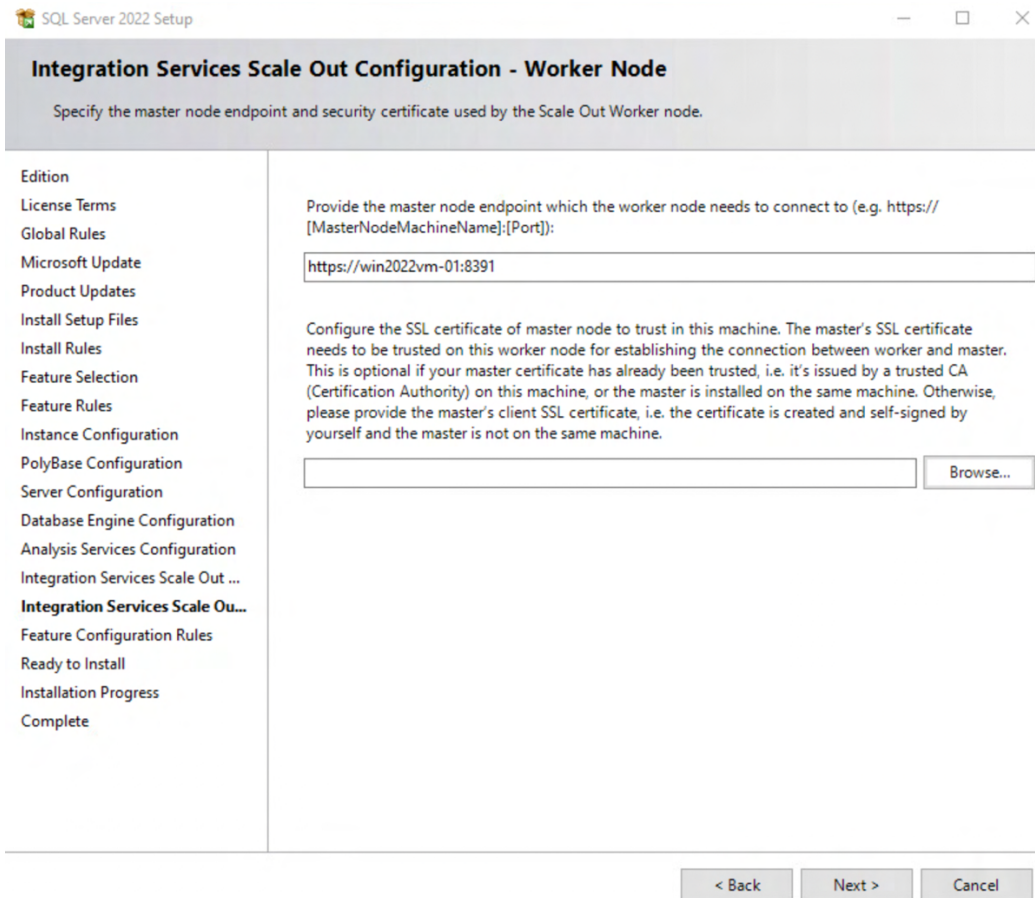
Select a SSL certificate that is used for the communication between the master node and worker nodes in the scale out topology. A default self-signed certificate is created if you choose to create a new SSL certificate.

☒ Create a new SSL certificate

CNs in the certificate:

☐ Use an existing SSL certificate

Click Next on the Integration Services Scale Out configuration – Worker Node screen.



Review the information on the Ready to Install screen and click Install.

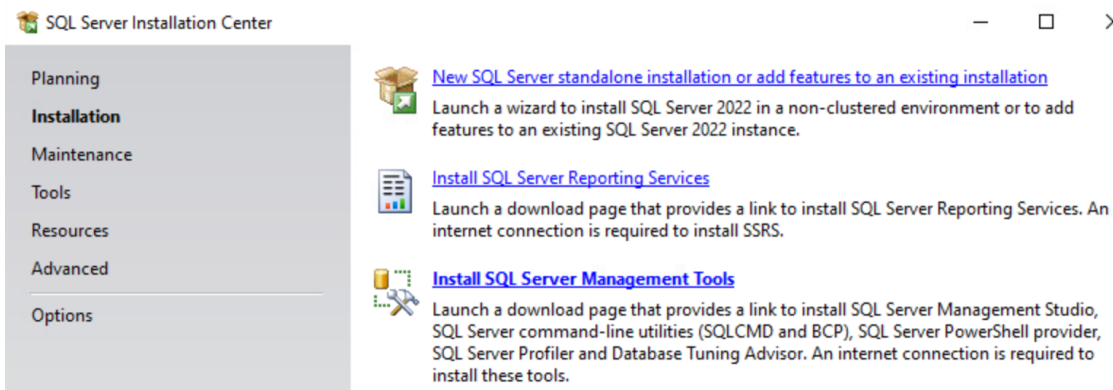
Review the installation results when the installation is completed. Click Close to exit the installer.

**Note:** The installation might take several minutes to complete.

## Install Microsoft SQL Server Management Studio 20.2.1

To install Microsoft SQL Server Management Studio, follow the steps below.

1. From SQL Server Installation Center, select Installation, and then click on Install SQL Server Management Tools to launch the SQL Server Management Studio download page.

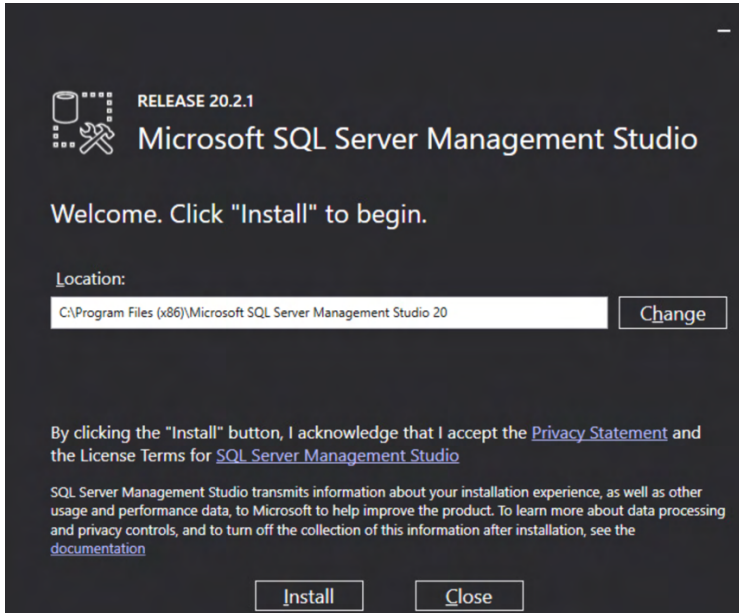


2. Review the hardware and software requirements.


3. Go to the Installation media section for links to download SQL Server Management Studio (SSMS) and cumulative patches.

**Note:** For this validation, we are using release 20.2.1.

4. Click on the downloaded setup application to start the installer.



5. Click Install to accept privacy statement and license terms and start the installation process.



6. Click Close when the setup is completed.

## Appendix C: Install and configure bare-metal Oracle Linux 8

This appendix outlines some of the procedures for installing and configuring bare-metal Oracle Linux 8 to access iSCSI LUNs from storage. The UCS server profile template for Linux was cloned from the VMware template for ESXi hosts with modifications to utilize non-VMware Ethernet adapter policy for all vNICs. The LAN connectivity for vNICs listed in Table 9 applies to the bare-metal Linux host as well.

The iSCSI vNICs in the server profile enable multipathing access to storage via both iSCSI-A and iSCSI-B networks. Network bonding interfaces are utilized for failover access to IB-MGMT and VM-Traffic using their respective vNIC pairs pinned to fabric A and fabric B.

## SAN boot LUN creation and mapping

Perform the steps below to install Oracle Linux 8 on a mapped SAN boot LUN.

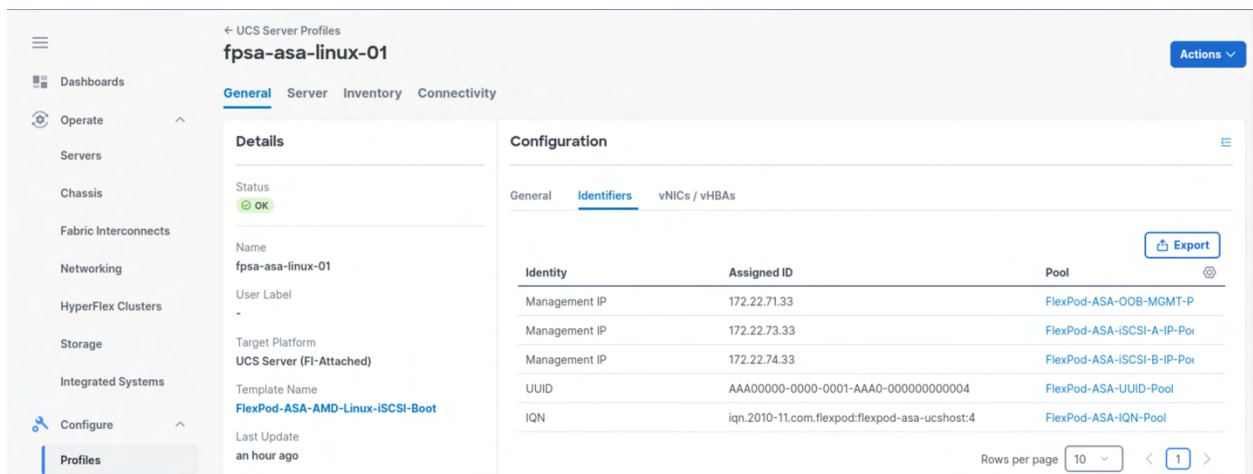
1. Create a boot LUN for installation in storage.

```
fpsa-a50-u0909::> lun create -path fpsa_asa_linux_01_ol8_boot_1 -size 200g -ostype linux
```

2. Create a boot igroup and include the server's iSCSI IQN.

```
fpsa-a50-u0909::> igroup create -igroup FlexPod-ASA-linux-01-boot-iscsi -protocol iscsi -ostype linux -initiator iqn.2010-11.com.flexpod:flexpod-asa-ucshost:4
```

**Note:** Retrieve server's iSCSI IQN information from server profile's Identifiers tab under Configuration.



Identity	Assigned ID	Pool
Management IP	172.22.71.33	FlexPod-ASA-OOB-MGMT-P
Management IP	172.22.73.33	FlexPod-ASA-iSCSI-A-IP-Por
Management IP	172.22.74.33	FlexPod-ASA-iSCSI-B-IP-Por
UUID	AAA00000-0000-0001-AAA0-000000000004	FlexPod-ASA-UUID-Pool
IQN	iqn.2010-11.com.flexpod:flexpod-asa-ucshost:4	FlexPod-ASA-IQN-Pool

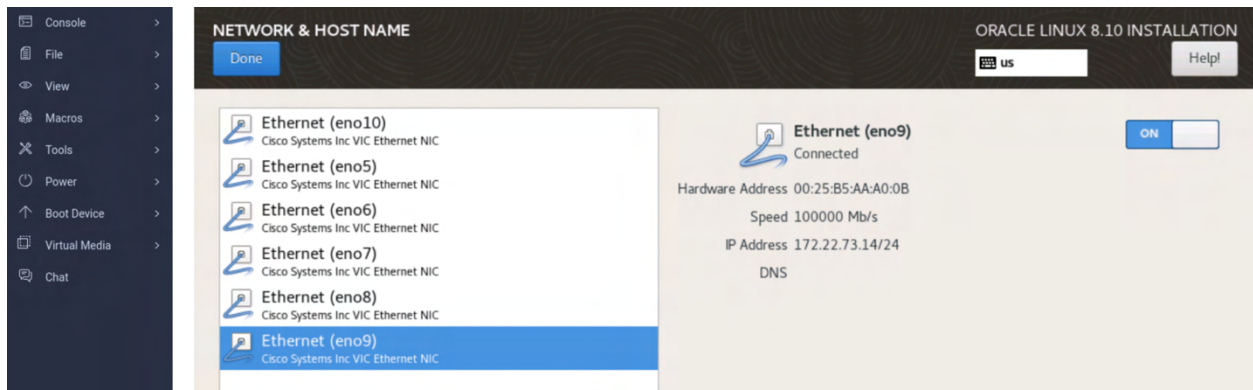
3. Map the boot LUN to the igroup as LUN 0.

```
fpsa-a50-u0909::> lun map -path fpsa_asa_linux_01_ol8_boot_1 -igroup FlexPod-ASA-linux-01-boot-iscsi -lun-id 0
```

## Oracle Linux 8 installation

1. Launch vKVM for the server, map OL8 ISO image from virtual media as vKVM-Mapped DVD, set boot device to vKVM-Mapped DVD, and reset or power on host.

Configure IP addresses, netmask, and MTU for the NICs assigned for iSCSI-A and iSCSI-B networking under the Network & Host Name section from the Installation Summary screen.



**NETWORK & HOST NAME**

Done

ORACLE LINUX 8.10 INSTALLATION

us Help

**Ethernet (eno10)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno5)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno6)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno7)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno8)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno9)**  
Cisco Systems Inc VIC Ethernet NIC

**Ethernet (eno9)**  
Connected

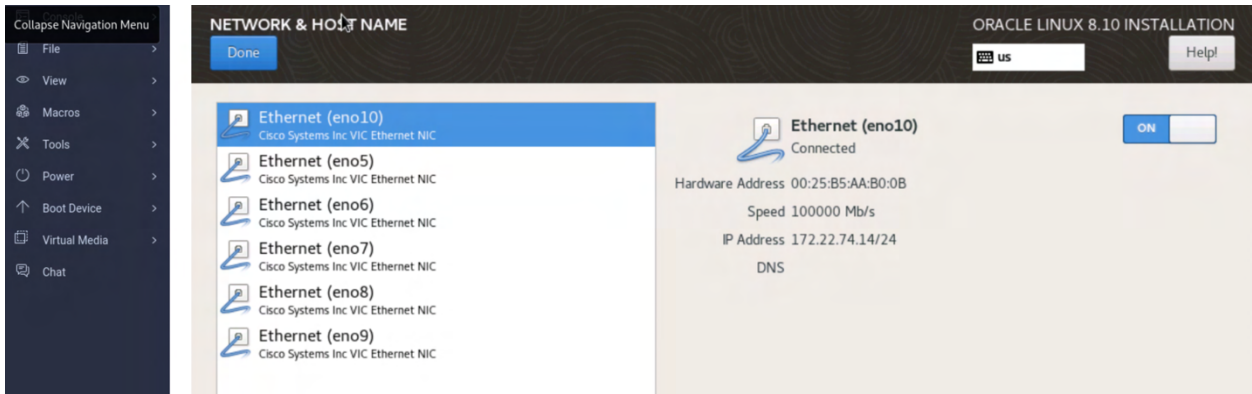
Hardware Address 00:25:B5:AA:A0:0B

Speed 100000 Mb/s

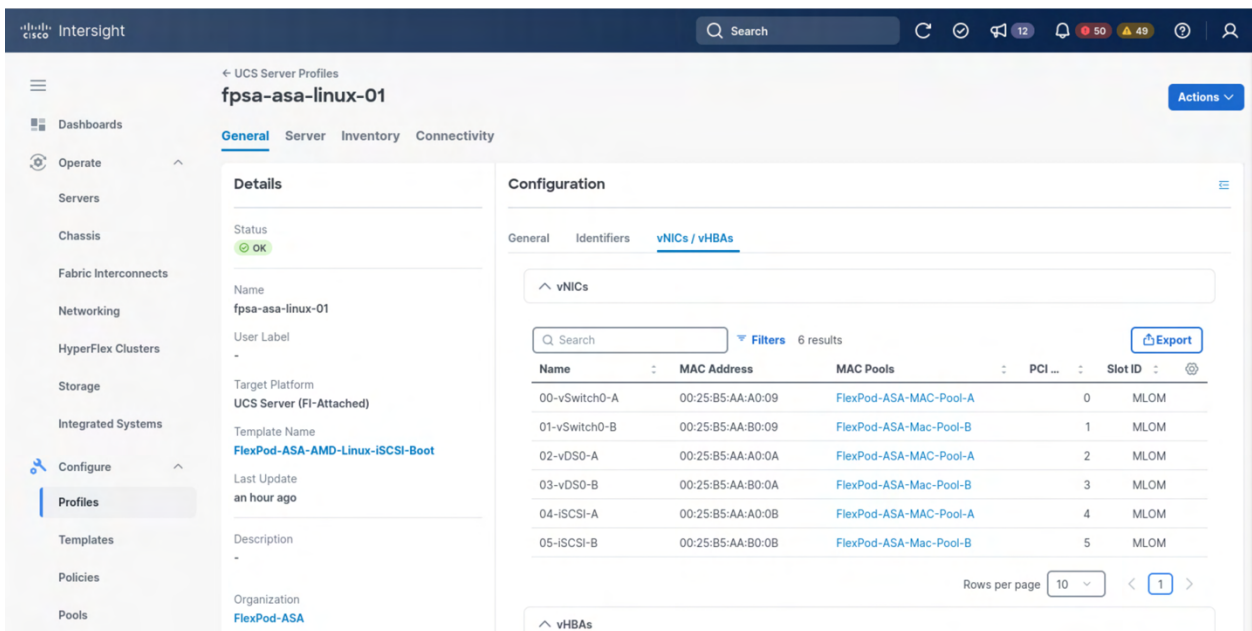
IP Address 172.22.73.14/24

DNS

ON



**Note:** Retrieve the vNIC MAC address information from the Intersight vNICs/vHBAs configuration for the server profile to identify the NICs appropriate for iSCSI-A and iSCSI-B network.



Click Add a disk under INSTALLATION DESTINATION from the Installation Summary screen.  
Click Add iSCSI Target, provide one of the iSCSI LIF's IP address, and click Start Discovery.




**ADD iSCSI STORAGE TARGET**

To use iSCSI disks, you must provide the address of your iSCSI target and the iSCSI initiator name you've configured for your host.

Target IP Address:

iSCSI Initiator Name:

 **Example:** iqn.2012-09.com.example:diskarrays-sn-a8675309

Discovery Authentication Type:

☐ Bind targets to network interfaces This may take a moment...

Check all the nodes/portals on the list and click Log In.



**ADD iSCSI STORAGE TARGET**

The following nodes were discovered using the iSCSI initiator **iqn.2010-11.com.flexpod:flexpod-asa-ucshost:4** using the portal IP address **172.22.73.101**. Please select which nodes you wish to log into:

	Node Name	Interface	Portal
<input checked="" type="checkbox"/>	iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2		172.22.
<input checked="" type="checkbox"/>	iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2		172.22.
<input checked="" type="checkbox"/>	iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2		172.22.
<input checked="" type="checkbox"/>	iqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:vs.2		172.22.

Node Login Authentication Type: No credentials (discovery authentication disabled)

Log In

Cancel OK

Click on the Multipath tab and select the discovered NETAPP LUN as the installation destination, then click DONE.

**INSTALLATION DESTINATION** ORACLE LINUX 8.10 INSTALLATION

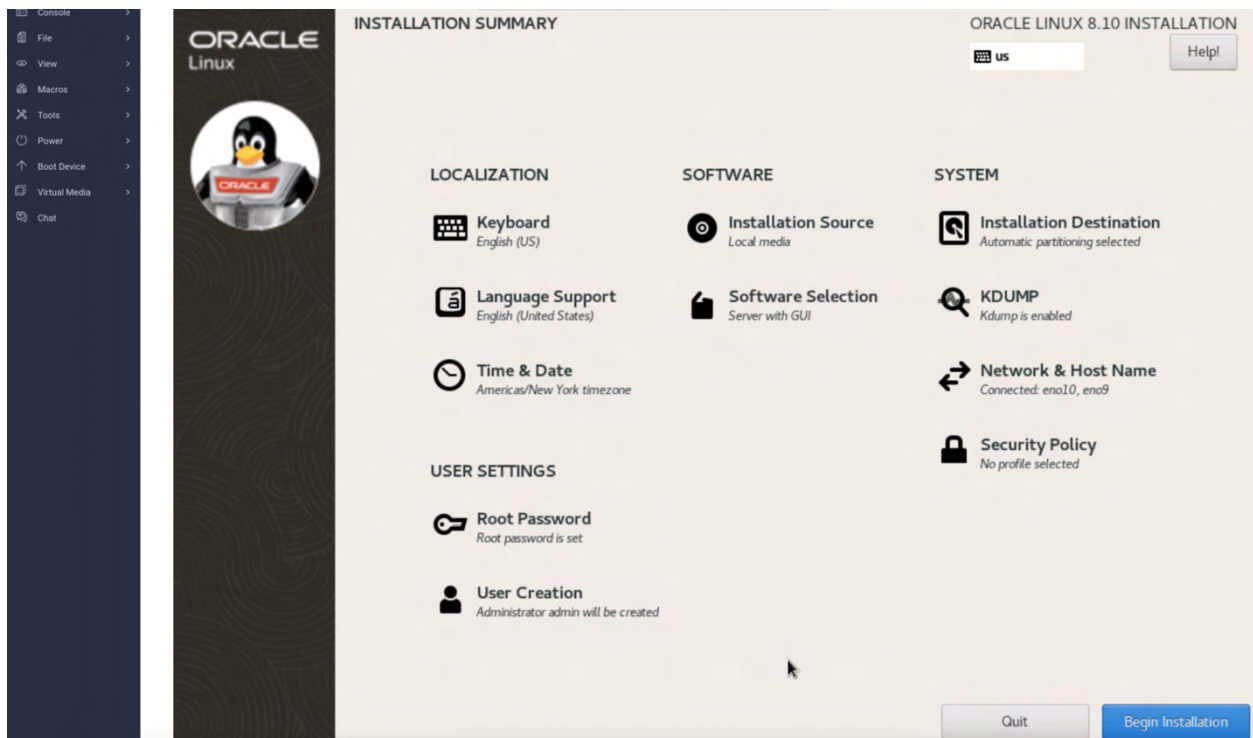
Done Help

Search Multipath Devices Other SAN Devices NVDIMM Devices

Filter By: None

WWID	Capacity	Vendor	Interconnect	Paths
<input checked="" type="checkbox"/>	200 GiB	NETAPP		sdf sde

With appropriate NICs configured for iSCSI networking and SAN boot LUN discovered and selected, root password configured, and user creation information provided, the OS installation can then begin.



## Server management address configuration

After the installation is completed, we can create a management IP on a bonding interface built on top of the two network interfaces that are pinned to the two network fabrics for IB-MGMT network access.

1. Login to the server console.
2. Identify the NIC devices that will be configured with the bonding interface.

```
[admin@localhost ~]$ ifconfig
eno5: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    ether 00:25:b5:aa:a0:09 txqueuelen 1000 (Ethernet)
    RX packets 65 bytes 18944 (18.5 KiB)
    RX errors 0 dropped 40 overruns 0 frame 0
    TX packets 2 bytes 180 (180.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eno6: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 9000
    ether 00:25:b5:aa:b0:09 txqueuelen 1000 (Ethernet)
    RX packets 65 bytes 18944 (18.5 KiB)
    RX errors 0 dropped 40 overruns 0 frame 0
    TX packets 1 bytes 90 (90.0 B)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Use the nmcli tool to create a bonding interface with the desired bonding options.

```
nmcli con add type bond ifname ibmgmt \
bond.options "mode=active-backup,miimon=100" \
ipv4.method disabled \
ipv6.method ignore
```

**Note:** The IPv4 address will be configured in a later step.

4. Add the two physical interfaces to the bond.

```
nmcli con add type ethernet ifname eno5 master ibmgmt
nmcli con add type ethernet ifname eno6 master ibmgmt
```

5. **Create a VLAN interface on top of the bond and specify IPv4 address, DNS server, and Domain.**

```
nmcli con add type vlan con-name ibmgmt2272 \
```

```
dev ibmgtm \
id 2272 \
ip4 172.22.72.14/24 \
gw4 172.22.72.1 \
ipv4.dns 10.61.177.2 \
ipv4.dns-search "nva.local"
```

**Note:** Here is a list of connections for reference.

```
[admin@fpsa-asa-linux-01 ~]$ nmcli con show
```

NAME	UUID	TYPE	DEVICE
ibmgtm2272	355e8084-3ce7-441e-97ea-94f0e2ce235f	vlan	ibmgtm.2272
bond-ibmgtm	c685d6e2-2132-46ef-8a7c-1c3d5099ef66	bond	ibmgtm
eno10	04004d06-3b1c-428f-9c83-6438dfa54354	ethernet	eno10
eno9	9edffb12-5286-429b-afe0-d258941a84af	ethernet	eno9
virbr0	51ea5f16-c374-4855-a46e-b069e13eea67	bridge	virbr0
bond-slave-eno5	d01de8fb-d516-415e-9705-6dd18ca9bcde	ethernet	eno5
bond-slave-eno6	64c1d853-85ab-41d5-b2f3-b5698329ff31	ethernet	eno6
eno5	ca42a2c2-06c3-4a66-8210-164313c4d4a1	ethernet	--
eno6	f2a90143-6eb0-4a37-9150-84e2bfcadb9a	ethernet	--
eno7	62220475-4c0d-4757-96ea-768137f0a0c5	ethernet	--
eno8	e10526a9-454a-4d8c-a0fb-d7a9ca01cc95	ethernet	--
iSCSI-A	517489b0-ff24-470e-be84-93ce27dccfe7	ethernet	--
iSCSI-B	fb06d7d-1ed7-4295-be7d-b27750793116	ethernet	--

## Additional server and storage configuration

Additional configurations for the bare-metal server are like the steps already included in the Oracle Linux VM configuration section to configure the following:

- Create LUNs and provide access to the initiator
- Configure software initiator and login to iSCSI target
- Rescan and discover LUNs
- Enable multipathing for LUNs
- Install NetApp Linux Host Utilities
- Enable persistent device naming for multipath devices using udev rules
- Configure proxy server if needed for additional software installation
- Configure bonding on eno7 and eno8 and private network address for Oracle installation by using the server management address configuration above as an example.

## Appendix D: Oracle Grid Infrastructure installation

This section highlights some of the steps taken for the single node Oracle 21c Grid Infrastructure installation on Oracle Linux. Deployment of a multi-mode cluster follows similar process, but with additional information for additional nodes where needed.

**Note:** It is not within the scope of this document to include the specifics of an Oracle Grid Infrastructure installation. Please click this link for the [Oracle 21c Grid Infrastructure Installation and Upgrade Guide](#) to see detailed information for installation.

## Configure OS prerequisite for Oracle software

To configure the operating system prerequisites for your Oracle deployment, here are highlights of some relevant steps.

1. Login as root user and install the "oracle-database-preinstall-21c" rpm package to simplify operating system configuration in preparation for Oracle software installations.

```
# dnf install oracle-database-preinstall-21c
```

Set the secure Linux to permissive by editing the "/etc/selinux/config" file, making sure the SELINUX flag is set as follows:

```
SELINUX=permissive
```

Check the status of the firewall by running following commands. (The status displays as active (running) or inactive (dead)). If the firewall is active / running, run the command to stop it:

```
# systemctl status firewalld
# systemctl stop firewalld
```

Check and add any required groups and users for Oracle database. For example:

```
# groupadd -g 54421 oinstall
# groupadd -g 54422 dba
# groupadd -g 54423 oper
# groupadd -g 54424 backupdba
# groupadd -g 54425 dgdba
# groupadd -g 54426 kmdba
# groupadd -g 54327 asmdba
# groupadd -g 54328 asmoper
# groupadd -g 54329 asmadmin

# useradd -u 54321 oracle -g oinstall
# useradd -u 54322 grid -g oinstall
# passwd oracle
# passwd grid

# usermod -a -G dba,oper,backupdba,dgdba,kmdba,racdba,asmadmin,asmdba oracle
# usermod -a -G dba,backupdba,dgdba,kmdba,racdba,asmadmin,asmdba,asmoper grid
```

Configure Public IP, Virtual IP, Private IP, and SCAN IP Addresses needed for the nodes and cluster in DNS and /etc/hosts file. For example, the following checks for the Single Client Access Name (SCAN) address lookup. See below for an example from our bare-metal server test instance.

```
# cat /etc/hosts
127.0.0.1    localhost localhost.localdomain localhost4 localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6

### Public IP ###
172.22.72.14  fpsa-asa-linux-01.nva.local          fpsa-asa-linux-01

### Virtual IP ###
172.22.72.24  fpsa-asa-linux-01-vip.nva.local        fpsa-asa-linux-01-vip

### Private IP ###
172.22.76.14  fpsa-asa-linux-01-priv.nva.local      fpsa-asa-linux-01-priv

### SCAN IP ###
172.22.72.51  fpsa-asa-linux-scan.nva.local            fpsa-asa-linux-scan
172.22.72.52  fpsa-asa-linux-scan.nva.local            fpsa-asa-linux-scan
172.22.72.53  fpsa-asa-linux-scan.nva.local            fpsa-asa-linux-scan

# nslookup fpsa-asa-linux-scan
Server:      10.61.177.2
Address:     10.61.177.2#53

Name:   fpsa-asa-linux-scan.nva.local
Address: 172.22.72.51
Name:   fpsa-asa-linux-scan.nva.local
Address: 172.22.72.52
Name:   fpsa-asa-linux-scan.nva.local
Address: 172.22.72.53
```

Create the directory structure according to your environment requirements. For example:

```
# mkdir -p /u01/app/21.3.0/grid
# mkdir -p /u01/app/grid
# mkdir -p /u01/app/oracle
# chown -R grid:oinstall /u01
# chown oracle:oinstall /u01/app/oracle
```

```
# chmod -R 775 /u01/
```

Login as the grid user, download the Oracle Grid Infrastructure image files and extract the files into the Grid home:

```
cd /u01/app/21.3.0/grid  
unzip -q <download_location>/LINUX.X64_213000_grid_home.zip
```

Run the cluster validation script named “runcluvfy.sh” as follows:

```
./runcluvfy.sh stage -pre crsinst -n <list_of_nodes> -verbose
```

**Note:** Resolve any issues reported by the validation script, e.g. insufficient swap space size, missing NTP time sync, missing group membership for users, etc. Look for the following success message near the end of the output to confirm that the cluster services set up was successful before moving on to the next step.

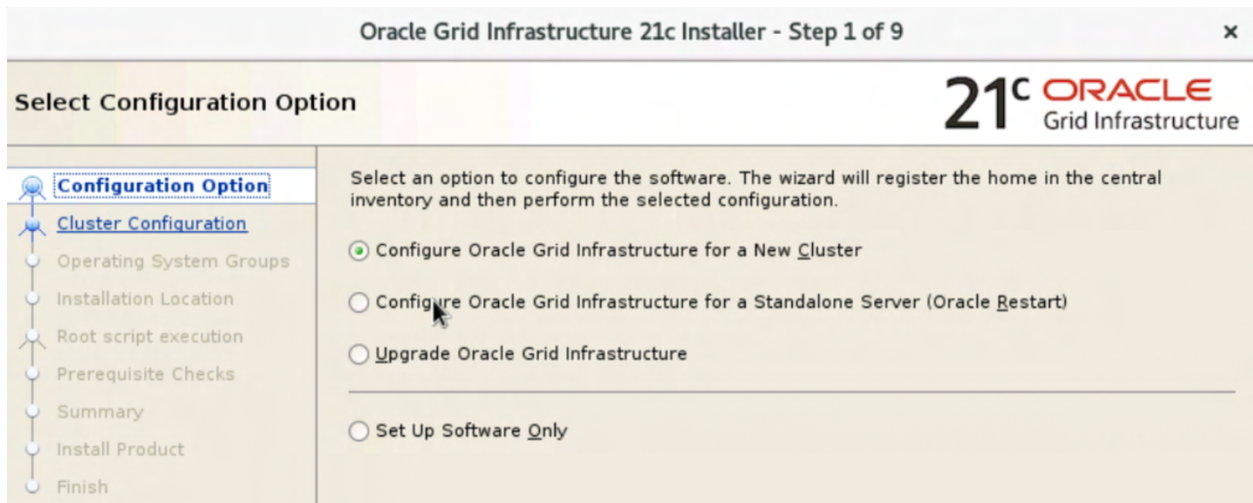
```
Pre-check for cluster services setup was successful.
```

## Oracle Grid Infrastructure setup

1. As the grid user, launch the gridSetup.sh program from the Grid home where the Oracle 21c Grid Infrastructure software binaries reside.

```
./gridSetup.sh
```

Select Configure Oracle Grid Infrastructure for a New Cluster and click Next.



For the Cluster Configuration, select Configure an Oracle Standalone Cluster and click Next.



Oracle Grid Infrastructure 21c Installer - Step 2 of 9

## Select Cluster Configuration

21<sup>c</sup> ORACLE Grid Infrastructure

[Configuration Option](#)

**Cluster Configuration**

[Operating System Groups](#)

Installation Location

Root script execution

Prerequisite Checks

Summary

Install Product

Finish

Choose the required cluster configuration.

☒ Configure an Oracle Standalone Cluster

☐ Configure an Oracle Domain Services Cluster

Oracle Extended clusters are special purpose clusters that constitute nodes which span across multiple sites. Specify a minimum of 3 site names and a maximum of 5 (e.g., siteA, siteB, siteC).

☐ Configure as an Oracle Extended cluster

Site names:

In next window, enter the names for your cluster and cluster SCAN name that are unique throughout your entire enterprise network. Click Next.

Oracle Grid Infrastructure 21c Installer - Step 3 of 17

## Grid Plug and Play Information

21<sup>c</sup> ORACLE Grid Infrastructure

[Configuration Option](#)

[Cluster Configuration](#)

**Grid Plug and Play**

[Cluster Node Information](#)

Network Interface Usage

Storage Option

GIMR Option

GIMR Storage Option

Create ASM Disk Group

ASM Password

Single Client Access Name (SCAN) allows clients to use one name in connection strings to connect to the cluster as a whole. Client connect requests to the SCAN name can be handled by any

☒ Create Local SCAN

Cluster Name:

SCAN Name:

SCAN Port:

☐ Use Shared SCAN

SCAN Client Data:

☐ Configure GNS

☐ Configure nodes Virtual IPs as assigned by the Dynamic Networks

**Note:** You can also select to Configure GNS if you have configured your domain name server (DNS) to send to the GNS virtual IP address name resolution requests.

In the Cluster Node Information window, the node where the gridSetup.sh program was launched from should already be added. click Add under the list to add additional Public Hostname and Virtual Hostname for the remaining cluster nodes if needed.



Oracle Grid Infrastructure 21c Installer - Step 4 of 16

## Cluster Node Information

21c ORACLE Grid Infrastructure

Provide the list of nodes to be managed by Oracle Grid Infrastructure with their Public Hostname and Virtual Hostname.

Public Hostname	Virtual Hostname
fpsa-asa-linux-01.nva.local	fpsa-asa-linux-01-vip.nva.local

SSH connectivity... Use Cluster Configuration File... Add... Edit... Remove

OS Username:  OS Password:

☐ Reuse private and public keys existing in the user home

Test Setup

Click the SSH Connectivity. Enter the operating system username and password for the Oracle software owner (grid). Click Setup to establish password-less SSH connectivity.

A message window may appear, indicating that it might take several minutes to configure SSH connectivity between the nodes when there are many nodes. After some time, another message window appears indicating that password-less SSH connectivity has been established between the cluster nodes. Click OK and then click Next to continue.

In the Network Interface Usage screen, select the usage type for each network interface for Public and Private Network Traffic and click Next.

Oracle Grid Infrastructure 21c Installer - Step 5 of 17

## Specify Network Interface Usage

21c ORACLE Grid Infrastructure

Private interfaces are used by Oracle Grid Infrastructure for internode traffic.

Interface Name	Subnet	Use for
eno9	172.22.73.0	Do Not Use
eno10	172.22.74.0	Do Not Use
oracle.2276	172.22.76.0	ASM & Private
ibmgmt.2272	172.22.72.0	Public

Note: While configuring an Oracle Member Cluster for Databases using the Grid Naming Service (GNS), only networks that have dynamic host configuration protocol (DHCP) assigned addresses can be designated as 'Public'.

**Note:** Select Do Not Use for the interfaces that should not be used. In this example, eno9 and eno10 are used for iSCSI-A and iSCSI-B storage network and should not be used for Oracle.

In the storage option, select Use Oracle Flex ASM for storage and click Next.

Oracle Grid Infrastructure 21c Installer - Step 6 of 17

Storage Option Information

Configuration Option

Cluster Configuration

Grid Plug and Play

Cluster Node Information

Network Interface Usage

**Storage Option**

GIMR Option

GIMR Storage Option

Create ASM Disk Group

ASM Password

Operating System Groups

21c ORACLE

Grid Infrastructure

You can place Oracle Cluster Registry (OCR) files and voting disk files on Oracle ASM storage, or on a file system. Oracle ASM can be configured on this cluster or can be an existing ASM on a storage server cluster.

☒ Use Oracle Flex ASM for storage

Choose this option to configure OCR and voting disks on ASM storage. ASM instance will be configured on reduced number of cluster nodes.

☐ Configure as ASM Client Cluster

Choose this option to store OCR and Voting disk files on Oracle ASM Storage configured on a storage server cluster.

ASM Client Data:

☐ Use Shared File System

Choose this option to configure OCR and voting disk files on an existing shared file system.

Select a configuration option for Grid Infrastructure Management Repository (GIMR). For this deployment, the Do Not use a GIMR database option was selected. Click Next.

Oracle Grid Infrastructure 21c Installer - Step 7 of 17

Create GIMR Option

Configuration Option

Cluster Configuration

Grid Plug and Play

Cluster Node Information

Network Interface Usage

Storage Option

**GIMR Option**

GIMR Storage Option

Create ASM Disk Group

ASM Password

Operating System Groups

Installation Location

21c ORACLE

Grid Infrastructure

The Grid Infrastructure Management Repository(GIMR) is an essential component for complete operation of the Autonomous Health Framework, that offers enhanced real time diagnostics and performance management, and Fleet Patching and Provisioning. The components that depend on the repository in whole or in part are Cluster Health Advisor, Cluster Health Monitor, QoS Management, Fleet Patching and Provisioning and Cluster Activity Log. It is best practice to install this option and failure to do so could compromise timely resolution of issues as well as available functionality for patching.

Select one of the GIMR configuration options

☐ Use a Local GIMR database

The GIMR database will have to be configured later in a separate RAC Database Oracle Home that is installed on all cluster nodes.

☐ Use an existing remote GIMR database

Specify a credential file:

☒ Do Not use a GIMR database

In the Create ASM Disk Group window, click Change Discovery Path to filter persistent device names for OCRVOTE disks, select the “iscsiocrvote1” & “iscsiocrvote2” disks to store OCR and Voting disk files. Enter the name of disk group “ISCSIOCRVOTE” and select appropriate external redundancy options as shown in the example below, then click Next.

252

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture



Oracle Grid Infrastructure 21c Installer - Step 8 of 16

Create ASM Disk Group

21c ORACLE  
Grid Infrastructure

[Configuration Option](#)  
[Cluster Configuration](#)  
[Grid Plug and Play](#)  
[Cluster Node Information](#)  
[Network Interface Usage](#)  
[Storage Option](#)  
[GIMR Option](#)  
**Create ASM Disk Group**  
[ASM Password](#)  
[Operating System Groups](#)  
[Installation Location](#)  
[Root script execution](#)  
[Prerequisite Checks](#)  
[Summary](#)  
[Install Product](#)  
[Finish](#)

OCR and Voting disk data will be stored in the following ASM Disk group. Select disks and characteristics of this Disk group.

Disk group name

Redundancy ☐ Flex ☐ High ☐ Normal ☒ External

Allocation Unit Size  MB

Select Disks Show Candidate/Provisioned Disks ▾

<input checked="" type="checkbox"/>	Disk Path	Size (in MB)	Status
<input checked="" type="checkbox"/>	/dev/iscsiocrvote1	102400	Candidate
<input checked="" type="checkbox"/>	/dev/iscsiocrvote2	102400	Candidate

Disk Discovery Path: '/dev/iscsiocrvote\*'

[Change Discovery Path...](#)

☐ Configure Oracle ASM Filter Driver

Select this option to configure ASM Filter Driver (AFD) to simplify configuration and management of disk devices by Oracle ASM.

**Note:** For this solution, we did not configure Oracle ASM Filter Driver.

Specify the password for the Oracle ASM account, then click Next.

Oracle Grid Infrastructure 21c Installer - Step 9 of 16

Specify ASM Password

21c ORACLE  
Grid Infrastructure

[Configuration Option](#)  
[Cluster Configuration](#)  
[Grid Plug and Play](#)  
[Cluster Node Information](#)  
[Network Interface Usage](#)  
[Storage Option](#)  
[GIMR Option](#)  
[Create ASM Disk Group](#)  
**ASM Password**  
[Operating System Groups](#)  
[Installation Location](#)  
[Root script execution](#)  
[Prerequisite Checks](#)

The new Oracle Automatic Storage Management (Oracle ASM) instance requires its own SYS user with SYSASM privileges for administration. Oracle recommends that you create a less privileged ASMSNMP user with SYSDBA privileges to monitor the ASM instance.

Specify the password for these user accounts.

☐ Use different passwords for these accounts

	Password	Confirm Password
SYS	<input type="text"/>	<input type="text"/>
ASMSNMP	<input type="text"/>	<input type="text"/>

☒ Use same passwords for these accounts

Specify Password:  Confirm Password:

For this solution, “Do not use Intelligent Platform Management Interface (IPMI)” was selected. Click Next.

Oracle Grid Infrastructure 21c Installer - Step 10 of 18

Failure Isolation Support

21c ORACLE  
Grid Infrastructure

Choose one of the following Failure Isolation Support options.

☐ Use Intelligent Platform Management Interface (IPMI)

To ensure successful installation with IPMI enabled, ensure your IPMI drivers are properly installed and enabled.

Path (ipmiutil):

User Name :

Password :

☒ Do not use Intelligent Platform Management Interface (IPMI)

[Configuration Option](#)

[Cluster Configuration](#)

[Grid Plug and Play](#)

[Cluster Node Information](#)

[Network Interface Usage](#)

[Storage Option](#)

[GIMR Option](#)

[Create ASM Disk Group](#)

[ASM Password](#)

**Failure Isolation**

You can configure to have this instance of the Oracle Grid Infrastructure and Oracle Automatic Storage Management to be managed by Enterprise Manager Cloud Control. For this solution, this option was not selected. Click Next.

Oracle Grid Infrastructure 21c Installer - Step 11 of 18

Specify Management Options

21c ORACLE  
Grid Infrastructure

You can configure to have this instance of Oracle Grid Infrastructure and Oracle Automatic Storage Management to be managed by Enterprise Manager Cloud Control. Specify the details of the Cloud Control configuration to perform the registration.

☐ Register with Enterprise Manager (EM) Cloud Control

OMS host:

OMS port:

EM Admin User Name:

EM Admin Password:

[Configuration Option](#)

[Cluster Configuration](#)

[Grid Plug and Play](#)

[Cluster Node Information](#)

[Network Interface Usage](#)

[Storage Option](#)

[GIMR Option](#)

[Create ASM Disk Group](#)

[ASM Password](#)

[Failure Isolation](#)

**Management Options**

Select the appropriate operating system group names for Oracle ASM according to your environment.

Oracle Grid Infrastructure 21c Installer - Step 12 of 18

Privileged Operating System Groups

21c ORACLE  
Grid Infrastructure

Select the name of the operating system group, that you want to use for operating system authentication to Oracle Automatic Storage Management.

Oracle ASM Administrator (OSASM) Group

Oracle ASM DBA (OSDBA for ASM) Group

Oracle ASM Operator (OSOPER for ASM) Group (Optional)

[Configuration Option](#)

[Cluster Configuration](#)

[Grid Plug and Play](#)

[Cluster Node Information](#)

[Network Interface Usage](#)

[Storage Option](#)



Specify the Oracle base and inventory directory to use for the Oracle Grid Infrastructure installation and then click Next.

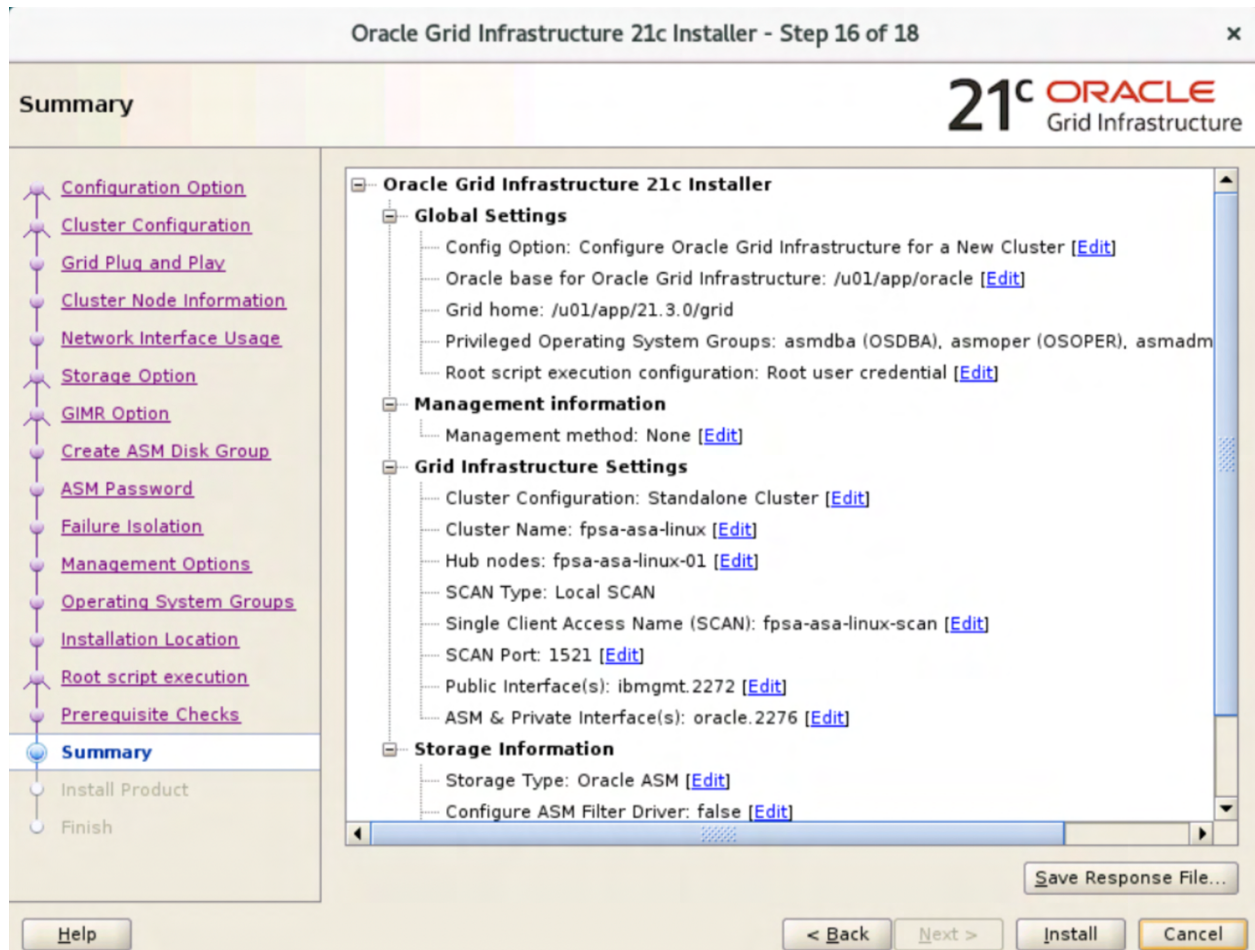
Click Automatically run configuration scripts to run scripts automatically and enter the relevant root user credentials. Click Next.

Wait while the prerequisite checks complete.

If you see any issues, click the "Fix & Check Again." If any of the checks have a status of Failed and are not fixable, then you must manually correct these issues. After you have fixed the issue, you can click

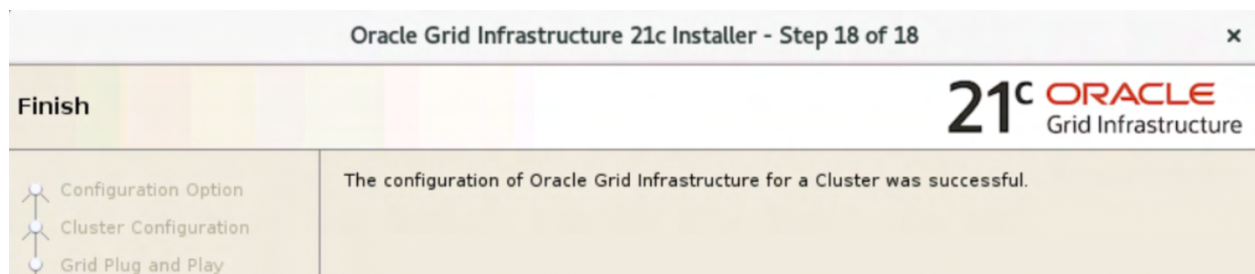
Check Again to have the installer check the requirement and update the status. Repeat as needed until all the checks have a status of Succeeded. Click Next.

Review the contents of the Summary window and then click Install. The installer displays a progress indicator enabling you to monitor the installation process.



Click Yes in the pop-up window to allow the installer generated script to run as privileged user.

At the end of the installation, you will see a confirmation message indicating that the installation was successful. Click Close to finish the Oracle Grid Infrastructure setup.



## Appendix E: Oracle RAC database installation

After Oracle Grid Infrastructure is installed successfully, you can proceed to install Oracle Database 21c software. The following highlights the steps taken for this solution validation. You can create databases using the DBCA tool afterwards.



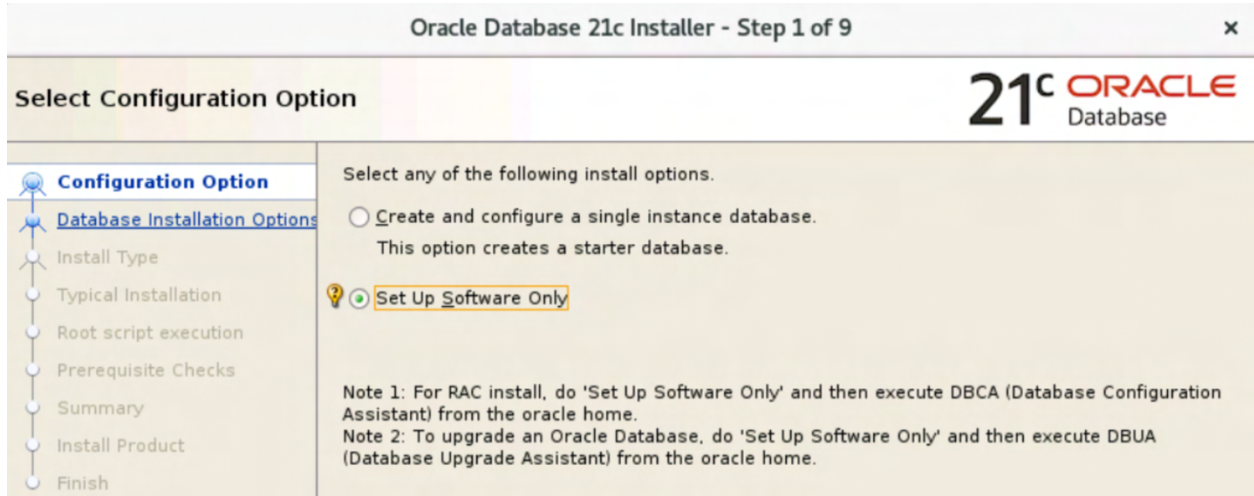
**Note:** It is not within the scope of this document to include the specifics of an Oracle Real Application Cluster (RAC) database installation. Please refer to the [Oracle database installation documentation](#) for specific installation instructions for your environment.

1. Login as oracle user to perform the Oracle database software installation.
2. Create directory structure for the software installation.

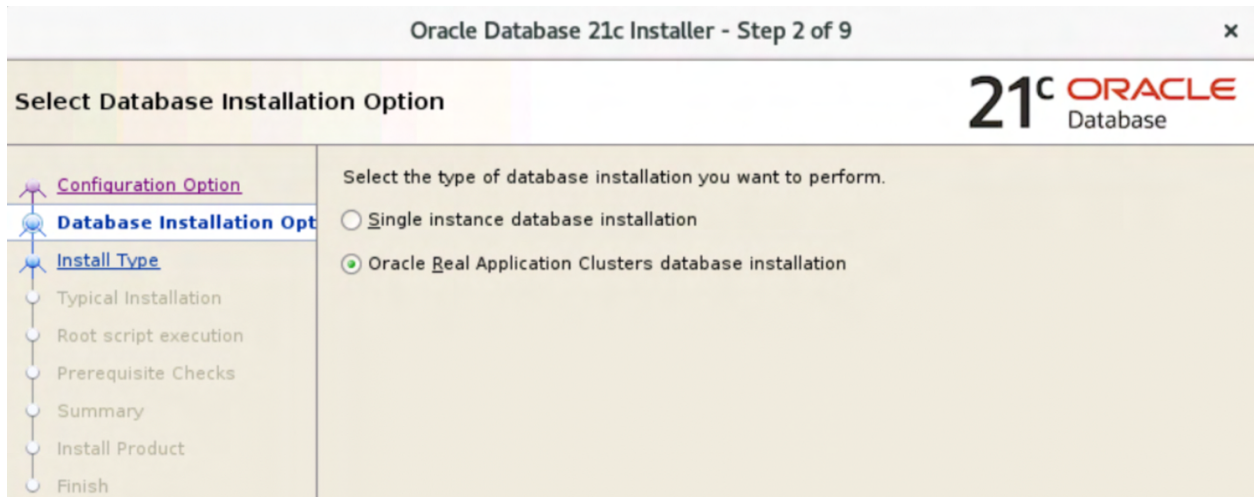
```
$ mkdir -p /u01/app/oracle/product/21.3.0/dbhome_1
$ cd /u01/app/oracle/product/21.3.0/dbhome_1
$ unzip -q <download_location>/LINUX.X64_213000_db_home.zip
```

Issue `./runInstaller` command from the Oracle Database 21c installation location.

3. For Configuration Option, select Set Up Software Only and click Next.



4. Select Oracle Real Application Clusters database installation and click Next.



5. The local node where the installer is running should already be selected. For this setup, we are installing a single node cluster.

Oracle Database 21c Installer - Step 3 of 10

**Select List of Nodes**

21<sup>c</sup> ORACLE Database

Configuration Option  
Database Installation Options  
**Nodes Selection**  
Install Type  
Typical Installation  
Root script execution  
Prerequisite Checks  
Summary  
Install Product  
Finish

Select nodes (in addition to the local node) in the cluster where the installer should install Oracle RAC or Oracle RAC One.

	Node name
<input checked="" type="checkbox"/>	1 fpsa-asa-linux-01

SSH connectivity... Select all Deselect all

OS Username: oracle OS Password: .....

☐ Reuse private and public keys existing in the user home

Test Setup

Click SSH connectivity and enter the password for the "oracle" user. Click Setup to configure password-less SSH connectivity and click Test to test it when it is complete. When the test is complete, click Next.

Select the Database Edition Options according to your environments and click Next.

Oracle Database 21c Installer - Step 4 of 11

**Select Database Edition**

21<sup>c</sup> ORACLE Database

Configuration Option  
Database Installation Options  
Nodes Selection  
**Database Edition**  
Installation Location  
Operating System Groups  
Root script execution  
Prerequisite Checks

Which database edition do you want to install?

☒ Enterprise Edition  
Oracle Database 21c Enterprise Edition is a self-managing database that has the scalability, performance, high availability, and security features required to run the most demanding, mission-critical applications.

☐ Standard Edition 2  
Oracle Database 21c Standard Edition 2 is a full-featured data management solution ideally suited to the needs of medium-sized businesses.

Enter the appropriate Oracle base location for the installation and click Next.

Oracle Database 21c Installer - Step 5 of 11

Specify Installation Location

21c ORACLE Database

[Configuration Option](#)  
[Database Installation Options](#)  
[Nodes Selection](#)  
[Database Edition](#)  
**[Installation Location](#)**  
[Operating System Groups](#)  
[Root script execution](#)

Specify a path to place all Oracle software and configuration-related files installed by this installation owner. This location is the Oracle base directory for the installation owner.

Oracle base:

This software directory is the Oracle Database home directory.

Software location: /u01/app/oracle/product/21.3.0/dbhome\_1

Select the desired operating system groups and click Next.

Oracle Database 21c Installer - Step 6 of 11

Privileged Operating System groups

21c ORACLE Database

[Configuration Option](#)  
[Database Installation Options](#)  
[Nodes Selection](#)  
[Database Edition](#)  
[Installation Location](#)  
**[Operating System Groups](#)**  
[Root script execution](#)  
[Prerequisite Checks](#)  
[Summary](#)  
[Install Product](#)  
[Finish](#)

SYS privileges are required to create a database using operating system (OS) authentication. Membership in OS Groups grants the corresponding SYS privilege, eg. membership in OSDBA grants the SYSDBA privilege.

Database Administrator (OSDBA) group:

Database Operator (OSOPER) group (Optional):

Database Backup and Recovery (OSBACKUPDBA) group:

Data Guard administrative (OSDGDBA) group:

Encryption Key Management administrative (OSKMDBA) group:

Real Application Cluster administrative (OSRACDBA) group:

Select the option Automatically run configuration script from the Root script execution configuration menu, provide root password, and click Next.



Oracle Database 21c Installer - Step 7 of 11

## Root script execution configuration

**21c ORACLE Database**

During the software configuration, certain operations have to be performed as "root" user. You can choose to have the installer perform these operations automatically by specifying inputs for one of the options below. The input specified will also be used by the installer to perform additional prerequisite checks.

☒ Automatically run configuration scripts

☒ Use "root" user credential

Password :

☐ Use sudo

Program path :

User name :

Password :

Configuration Option  
Database Installation Options  
Nodes Selection  
Database Edition  
Installation Location  
Operating System Groups  
**Root script execution**  
Prerequisite Checks  
Summary  
Install Product  
Finish

Wait for the prerequisite check to complete. If there are any problems, click "Fix & Check Again" or try to fix those by checking and manually resolving the issue.

Verify the Oracle Database summary information and then click Install.

Oracle Database 21c Installer - Step 9 of 11

## Summary

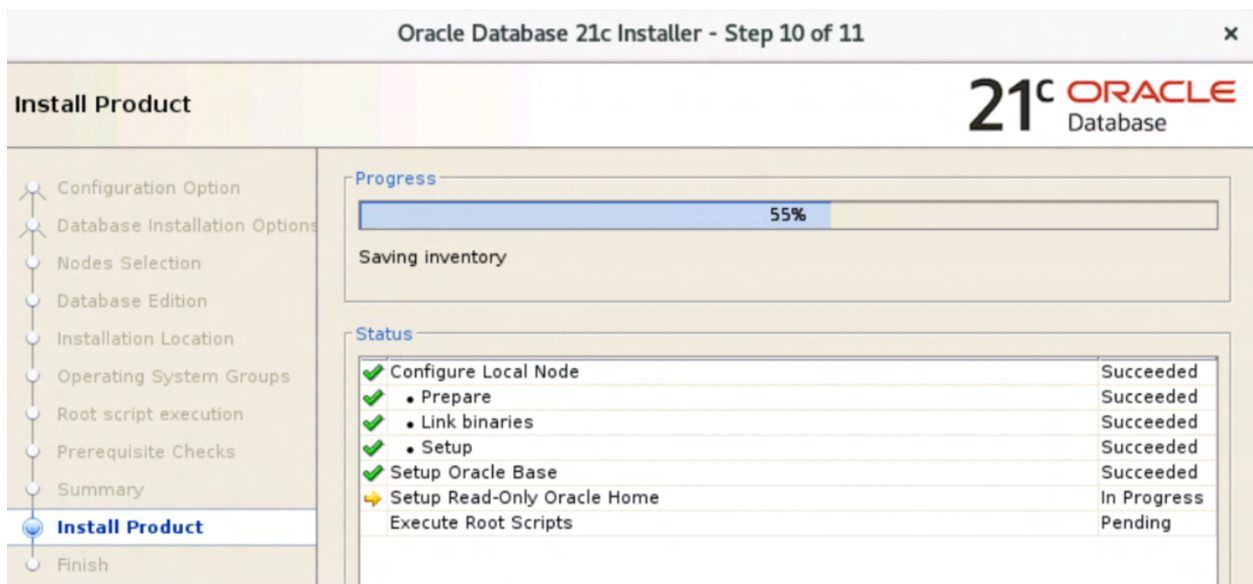
**21c ORACLE Database**

Configuration Option  
Database Installation Options  
Nodes Selection  
Database Edition  
Installation Location  
Operating System Groups  
Root script execution  
Prerequisite Checks  
**Summary**  
Install Product  
Finish

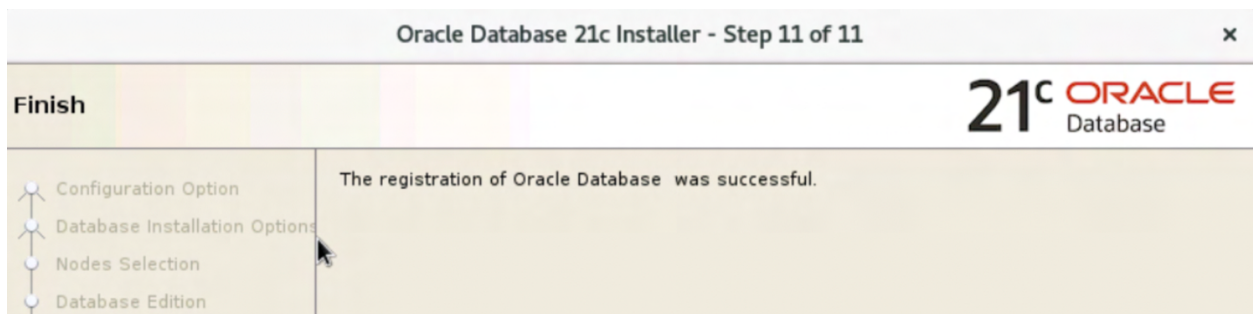
**Oracle Database 21c Installer**

- Global settings**
  - Database edition: Enterprise Edition (Set Up Software Only) [\[Edit\]](#)
  - Oracle base: /u01/app/oracle [\[Edit\]](#)
  - Software location: /u01/app/oracle/product/21.3.0/dbhome\_1
  - Privileged Operating System groups: dba (OSDBA), oper (OSOPER), backupdba (OSBACK)
  - Root script execution configuration: Root user credential [\[Edit\]](#)
- Grid Options**
  - Cluster Nodes: fpsa-asa-linux-01 [\[Edit\]](#)

The progress bar shows the installation progress. Click Yes to allow the configuration script generated by the Installer to run as privileged user in the pop-up window when prompted.



Wait for the installation of Oracle Database to finish and then click Close to exit the installer.



## Appendix F: Configure ONTAP storage for client NVMe/TCP access

### An introduction to NVMe terminologies

The nonvolatile memory express (NVMe) protocol is a transport protocol used for accessing nonvolatile storage media. NVMe over Fabrics (NVMeoF) enables NVMe-based communication over fabric connections. NVMe/TCP utilizes Ethernet fabric and NVMe/FC utilizes FC fabric.

An NVMe namespace is a quantity of non-volatile memory that can be formatted into logical blocks. Namespaces are the equivalent of LUNs for FC and iSCSI protocols.

An NVMe subsystem includes one or more NVMe controllers, namespaces, NVM subsystem ports, an NVM storage medium, and an interface between the controller and the NVM storage medium. Before clients can utilize NVMe namespaces, they need to be mapped to an NVMe subsystem.

From operations perspective, NVMe subsystems are the equivalent of igroups for FC and iSCSI protocols. You will need to add NVMe Qualified Names (NQN) of the hosts to the NVMe subsystem to allow those hosts to access the NVMe namespaces mapped to the NVMe subsystem.

The namespace IDs are identifiers for the controllers to access the namespaces. The role of the NVMe namespace IDs is like the LUN IDs for FC and iSCSI protocols. A namespace should only be mapped to a single host group.

## Configure network for NVMe VLAN

Check to ensure VLAN ports and broadcast-domains for NVMe/TCP traffic have already been configured. In the example below the broadcast-domain NVMe-TCP-A includes the two VLAN ports from both controllers for VLAN 2277. Similarly, the broadcast-domain NVMe-TCP-B includes the two VLAN ports from both controllers for VLAN 2278.

```
fpsa-a50-u0909:> broadcast-domain show
(network port broadcast-domain show)
IPspace Broadcast
Name      Domain Name      MTU    Port List      Update
-----
Cluster Cluster      9000
          fpsa-a50-u0909-02:e2a    complete
          fpsa-a50-u0909-02:e4a    complete
          fpsa-a50-u0909-01:e2a    complete
          fpsa-a50-u0909-01:e4a    complete
Default Default      1500
          fpsa-a50-u0909-02:e0M    complete
          fpsa-a50-u0909-01:e0M    complete
          NVMe-TCP-A      9000
          fpsa-a50-u0909-02:e2b-2277 complete
          fpsa-a50-u0909-01:e2b-2277 complete
          NVMe-TCP-B      9000
          fpsa-a50-u0909-02:e4b-2278 complete
          fpsa-a50-u0909-01:e4b-2278 complete
          iSCSI-A         9000
          fpsa-a50-u0909-02:e2b-2273 complete
          fpsa-a50-u0909-01:e2b-2273 complete
          iSCSI-B         9000
          fpsa-a50-u0909-02:e4b-2274 complete
          fpsa-a50-u0909-01:e4b-2274 complete
6 entries were displayed.
```

Please refer to the earlier ONTAP storage configuration section for relevant configurations before proceeding.

## Create NVMe/TCP LIFs

A NVMe/TCP logical network interface (LIF) is an IP address associated with a physical or logical port in the storage controller for NVMe/TCP communication. The default-data-blocks service-policy for IP LIFs enables the created IP LIFs to serve both iSCSI and NVMe/TCP protocols. However, since we have separate VLANs for iSCSI and NVMe/TCP protocols, separate LIFs are created to serve iSCSI and NVMe/TCP protocols. To create NVMe/TCP-only LIFs, follow the steps below.

1. Create a custom NVMe/TCP-only service-policy with core-data and core-nvme-tcp services using advanced privilege as shown in the example below.

```
fpsa-a50-u0909:> set -privilege advanced

Warning: These advanced commands are potentially dangerous; use them only when directed to do so
by NetApp personnel.
Do you want to continue? {y|n}: y

fpsa-a50-u0909:> network interface service-policy create -vserver svml -policy custom-data-
nvme-tcp -service data-core,data-nvme-tcp

To verify:

fpsa-a50-u0909:> network interface service-policy show -policy custom-data-nvme-tcp
Vserver Policy Service: Allowed Addresses
-----
svml
      custom-data-nvme-tcp data-core: 0.0.0.0/0
                           data-nvme-tcp: 0.0.0.0/0

fpsa-a50-u0909:> set -privilege admin
```



To create four NVMe/TCP LIFs (two on each node, one for fabric A and the other for fabric B), follow the example below.

```
network interface create -vserver svml -lif nvmetcp-lif-01a -service-policy custom-data-nvme-tcp
-home-node <node01> -home-port e2b-<nvmetcp-a-vlan-id> -address <node01-nvmetcp-a-ip> -netmask
<nvmetcp-a-mask> -status-admin up

network interface create -vserver svml -lif nvmetcp-lif-01b -service-policy custom-data-nvme-tcp
-home-node <node01> -home-port e4b-<nvmetcp-b-vlan-id> -address <node01-nvmetcp-b-ip> -netmask
<nvmetcp-b-mask> -status-admin up

network interface create -vserver svml -lif nvmetcp-lif-02a -service-policy custom-data-nvme-tcp
-home-node <node02> -home-port e2b-<nvmetcp-a-vlan-id> -address <node02-nvmetcp-a-ip> -netmask
<nvmetcp-a-mask> -status-admin up

network interface create -vserver svml -lif nvmetcp-lif-02b -service-policy custom-data-nvme-tcp
-home-node <node02> -home-port e4b-<nvmetcp-b-vlan-id> -address <node02-nvmetcp-b-ip> -netmask
<nvmetcp-b-mask> -status-admin up
```

Example:

```
fpsa-a50-u0909::> network interface create -vserver svml -lif nvmetcp-lif-01a -service-policy
custom-data-nvme-tcp -home-node fpsa-a50-u0909-01 -home-port e2b-2277 -address 172.22.77.101 -
netmask 255.255.255.0 -status-admin up

fpsa-a50-u0909::> network interface create -vserver svml -lif nvmetcp-lif-01b -service-policy
custom-data-nvme-tcp -home-node fpsa-a50-u0909-01 -home-port e4b-2278 -address 172.22.78.101 -
netmask 255.255.255.0 -status-admin up

fpsa-a50-u0909::> network interface create -vserver svml -lif nvmetcp-lif-02a -service-policy
custom-data-nvme-tcp -home-node fpsa-a50-u0909-02 -home-port e2b-2277 -address 172.22.77.102 -
netmask 255.255.255.0 -status-admin up

fpsa-a50-u0909::> network interface create -vserver svml -lif nvmetcp-lif-02b -service-policy
custom-data-nvme-tcp -home-node fpsa-a50-u0909-02 -home-port e4b-2278 -address 172.22.78.102 -
netmask 255.255.255.0 -status-admin up
```

To verify:

```
fpsa-a50-u0909::> net int show -lif nvmetcp*
(network interface show)
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	nvmetcp-lif-01a	up/up	172.22.77.101/24	fpsa-a50-u0909-01	e2b-2277	true
	nvmetcp-lif-01b	up/up	172.22.78.101/24	fpsa-a50-u0909-01	e4b-2278	true
	nvmetcp-lif-02a	up/up	172.22.77.102/24	fpsa-a50-u0909-02	e2b-2277	true
	nvmetcp-lif-02b	up/up	172.22.78.102/24	fpsa-a50-u0909-02	e4b-2278	true

4 entries were displayed.

**Note:** Two NVMe LIFs should be created on each controller for the two SAN fabrics for load-balancing and fabric / connectivity failure resiliency. NVMe/TCP LIFs do not support LIF failover.

## Create NVMe subsystem

Follow the example below to create a subsystem named FlexPod-ASA-esxi-cluster-nvme with vmware host type and show the created NVMe subsystem.

```
fpsa-a50-u0909::> nvme subsystem create -subsystem FlexPod-ASA-esxi-cluster-nvme -ostype vmware
(vserver nvme subsystem create)

fpsa-a50-u0909::> nvme subsystem show
(vserver nvme subsystem show)
```

Vserver	Subsystem	Target NQN
svml		

```
FlexPod-ASA-esxi-cluster-nvme nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-ASA-esxi-cluster-nvme
```

**Note:** The nvme subsystem show command also shows the NQN of the target.

## Add host NQNs to NVMe subsystem host

Follow the example below to add the NQNs from the hosts to the NVMe subsystem host group.

```
fpsa-a50-u0909:> nvme subsystem host add -subsystem FlexPod-ASA-esxi-cluster-nvme -host-nqn
nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01
(vserver nvme subsystem host add)

fpsa-a50-u0909:> nvme subsystem host add -subsystem FlexPod-ASA-esxi-cluster-nvme -host-nqn
nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-02
(vserver nvme subsystem host add)

fpsa-a50-u0909:> nvme subsystem host show
(vserver nvme subsystem host show)
Vserver Subsystem Priority Host NQN
-----
svml1 FlexPod-ASA-esxi-cluster-nvme
      regular nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01
      regular nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-02
2 entries were displayed.
```

## Create NVMe namespace and map it to NVMe subsystem host

To create an NVMe namespace and map it to the NVMe subsystem for host access, follow the steps below.

1. Create an NVMe namespace with the desired host like by following the example below.

```
fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_vmware_cluster_datastore_nvme_1 -size 1TB
-ostype vmware
(vserver nvme namespace create)
[Job 4078] Job succeeded.

fpsa-a50-u0909:> nvme namespace show
(vserver nvme namespace show)
Vserver Path State Size Subsystem NSID
-----
svml1 fpsa_asa_vmware_cluster_datastore_nvme_1 online 1TB -
```

2. Map an NVMe namespace to a subsystem by following the example below.

```
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-esxi-cluster-nvme -path
fpsa_asa_vmware_cluster_datastore_nvme_1
(vserver nvme subsystem map add)

fpsa-a50-u0909:> nvme subsystem map show
(vserver nvme subsystem map show)
Vserver Subsystem NSID Namespace Path
-----
svml1 FlexPod-ASA-esxi-cluster-nvme
      00000001h fpsa_asa_vmware_cluster_datastore_nvme_1
```

## Appendix G: Configure VMware cluster and hosts for NVMe/TCP protocol access to storage

### Create NVMe/TCP distributed port groups in VMware cluster

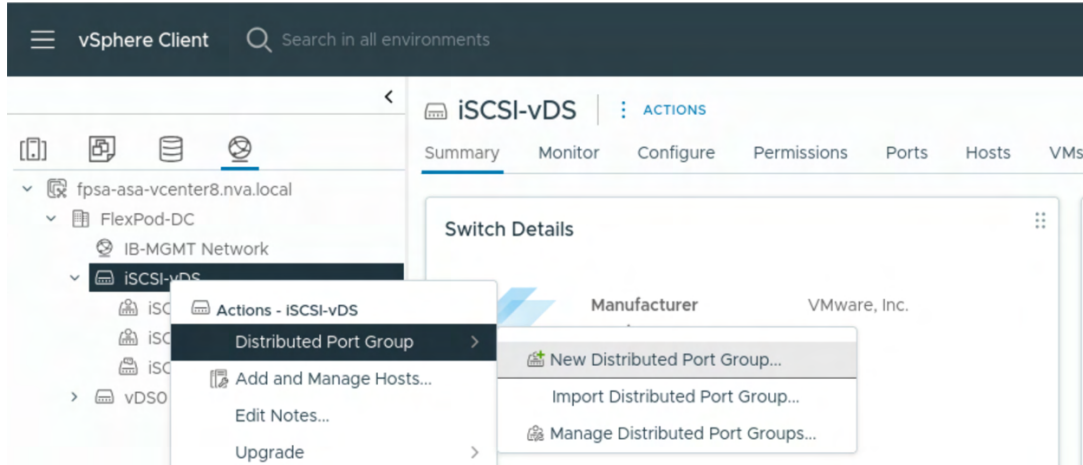
The same vDS switch and physical adapters that are utilized for iSCSI storage traffic will also be used for NVMe/TCP storage traffic, although the NVMe/TCP storage traffic will be using its VLANs and not use the native iSCSI VLANs configured for those adapters.

To add the two NVMe/TCP port groups, one for each SAN fabric, follow the steps below.

Login to vCenter.

Click the Networking icon in the Inventory view.

Right-click on the iSCSI-vDS switch, select Distributed Port Group > New Distributed Port Group.



Enter the new distributed port group name, e.g. NVMe/TCP-A and click NEXT.

The screenshot shows the 'New Distributed Port Group' wizard. On the left, a sidebar lists the steps: 1 Name and location (active), 2 Configure settings, and 3 Ready to complete. The main area is titled 'Name and location' and contains the instruction 'Specify distributed port group name and location.' There are two fields: 'Name' with the value 'NVMe/TCP-A' and 'Location' with a switch icon and the value 'iSCSI-vDS'.

Click on the drop-down list under VAN to select VLAN type, enter the VLAN ID for NVMe/TCP-A VLAN and click NEXT.

The screenshot shows the 'New Distributed Port Group' wizard, Step 2: Configure settings. The sidebar shows steps 1 Name and location, 2 Configure settings (active), and 3 Ready to complete. The main area is titled 'Configure settings' and contains the instruction 'Set general properties of the new port group.' There are four settings: 'Port binding' set to 'Static binding', 'Port allocation' set to 'Elastic', 'Number of ports' set to '8', and 'Network resource pool' set to '(default)'. Below these is a section titled 'VLAN' with 'VLAN type' set to 'VLAN' and 'VLAN ID' set to '2277'. At the bottom, there is an 'Advanced' section with a checkbox 'Customize default policies configuration' which is unchecked.

Review the information and click FINISH to complete.

New Distributed Port Group

1 Name and location

2 Configure settings

3 Ready to complete

Ready to complete

Review the changes before proceeding.

Distributed port group name	NVMe/TCP-A
Port binding	Static binding
Number of ports	8
Port allocation	Elastic
Network resource pool	(default)
VLAN ID	2277

Repeat the steps above to create the NVMe/TCP-B port group with its VLAN ID.

New Distributed Port Group

1 Name and location

2 Configure settings

3 Ready to complete

Ready to complete

Review the changes before proceeding.

Distributed port group name	NVMe/TCP-B
Port binding	Static binding
Number of ports	8
Port allocation	Elastic
Network resource pool	(default)
VLAN ID	2278

Click on the iSCSI-vDS switch and then click on the Networks tab to review and confirm the VLAN ID information for the newly created NVMe/TCP port groups.

vSphere Client

Search in all environments

fsa-asa-vcenter8.nva.local

FlexPod-DC

IB-MGMT Network

iSCSI-vDS

iSCSI-A

iSCSI-B

iSCSI-vDS-DVUplinks-1038

NVMe/TCP-A

NVMe/TCP-B

iSCSI-vDS

ACTIONS

SummaryMonitorConfigurePermissionsPortsHostsVMsNetworks

Distributed Port Groups

Uplink Port Groups

Quick FilterEnter value

	Name	VLAN ID	NSX Port Group ID
<input type="checkbox"/>	iSCSI-A	VLAN access: 0	
<input type="checkbox"/>	iSCSI-B	VLAN access: 0	
<input type="checkbox"/>	NVMe/TCP-A	VLAN access: 2277	
<input type="checkbox"/>	NVMe/TCP-B	VLAN access: 2278	

Click on NVMe/TCP-A distributed port group, click Properties, and then click EDIT at the upper-right hand corner to edit its properties.

Select Teaming and failover page, click on Uplink 2 under the Active uplinks and click MOVE DOWN twice to move Uplink 2 down to the Unused uplinks section and click OK.

## Distributed Port Group - Edit Settings | NVMe/TCP-A



General	Load balancing	Route based on originating virtual port ▾
Advanced	Network failure detection	Link status only ▾
VLAN	Notify switches	Yes ▾
Security	Fallback	Yes ▾
Traffic shaping		
<b>Teaming and failover</b>	<b>Failover order</b> ⓘ <a href="#">MOVE UP</a> <a href="#">MOVE DOWN</a> <a href="#">RESTORE DEFAULTS</a> <b>Active uplinks</b> <div>Uplink 1</div> <b>Standby uplinks</b> <b>Unused uplinks</b> <div>Uplink 2</div>	
Monitoring		
Miscellaneous		

**Note:** This configures NVMe/TCP-A port group to utilize only uplink 1 which is connected to SAN fabric A.

Similarly, configure NVMe/TCP-B port group to utilize only uplink 2 which is connected to SAN fabric B and move Uplink 1 down to the Unused uplinks section and then click OK.

## Distributed Port Group - Edit Settings | NVMe/TCP-B



General	Load balancing	Route based on originating virtual port ▾
Advanced	Network failure detection	Link status only ▾
VLAN	Notify switches	Yes ▾
Security	Fallback	Yes ▾
Traffic shaping		
<b>Teaming and failover</b>	<b>Failover order</b> ⓘ <a href="#">MOVE UP</a> <a href="#">MOVE DOWN</a> <a href="#">RESTORE DEFAULTS</a> <b>Active uplinks</b> <div>Uplink 2</div> <b>Standby uplinks</b> <b>Unused uplinks</b> <div>Uplink 1</div>	
Monitoring		
Miscellaneous		

## Add vmkernel adapters to ESXi hosts for NVMe/TCP storage traffic

To add the vmkernel adapters, one for each SAN fabric, for NVMe/TCP storage traffic, follow the steps below and perform these steps on each host in the cluster.

1. Select a host listed under the vCenter Inventory view, select Configure tab. Select VMkernel adapters, click ADD NETWORKING to open the Add Networking dialog, Select VMkernel Network Adapter as the connection type and click NEXT.

Add Networking

1 Select connection type
2 Select target device
3 Port properties
4 IPv4 settings
5 Ready to complete

Select connection type

Select a connection type to create.

☒ **VMkernel Network Adapter**  
The VMkernel TCP/IP stack handles traffic for ESXi services such as vSphere vMotion, iSCSI, NFS, FCoE, Fault Tolerance, vSAN, host management and etc.

☐ **Virtual Machine Port Group for a Standard Switch**  
A port group handles the virtual machine traffic on standard switch.

☐ **Physical Network Adapter**  
A physical network adapter handles the network traffic to other hosts on the network.

For the target device, select an existing network, and then select NVMe/TCP-A, and then click NEXT.

Add Networking

1 Select connection type
2 Select target device
3 Port properties
4 IPv4 settings
5 Ready to complete

Select target device

Select a target device for the new connection.

☒ Select an existing network
☐ Select an existing standard switch
☐ New standard switch

Quick Filter

	Name	NSX Port Group ID	Distributed Switch
<input type="radio"/>	iSCSI-A	--	iSCSI-vDS
<input type="radio"/>	iSCSI-B	--	iSCSI-vDS
<input checked="" type="radio"/>	NVMe/TCP-A	--	iSCSI-vDS
<input type="radio"/>	NVMe/TCP-B	--	iSCSI-vDS
<input type="radio"/>	VM-Traffic	--	VDS0
<input type="radio"/>	vMotion	--	VDS0

For the Enabled services under port properties, select NVMe over TCP and then click NEXT.

Add Networking

1 Select connection type
2 Select target device
3 Port properties
4 IPv4 settings
5 Ready to complete

Port properties

Specify VMkernel port settings.

Network label

MTU

TCP/IP stack

Available services

Enabled services

☐ vMotion
☐ vSphere Replication NFC
☐ NVMe over RDMA

☐ Provisioning
☐ vSAN

☐ Fault Tolerance logging
☐ vSAN Witness

☐ Management
☐ vSphere Backup NFC

☐ vSphere Replication
☒ NVMe over TCP

Provide the IP address and subnet mask for the IPv4 settings and click NETXT.

268

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture



### Add Networking

- Select connection type
- Select target device
- Port properties
- IPv4 settings**
- Ready to complete

### IPv4 settings

Specify VMkernel IPv4 settings.

☐ Obtain IPv4 settings automatically  
☒ Use static IPv4 settings

IPv4 address	172.22.77.11
Subnet mask	255.255.255.0
Default gateway	<input type="checkbox"/> Override default gateway for this adapter 172.22.72.1
DNS server addresses	10.61.177.2

Review the selection and then click FINISH.

### Add Networking

- Select connection type
- Select target device
- Port properties
- IPv4 settings
- Ready to complete**

### Ready to complete

Review your selections before finishing the wizard

- Select target device**

Distributed port group	NVMe/TCP-A
Distributed switch	iSCSI-vDS
- Port properties**

New port group	NVMe/TCP-A (iSCSI-vDS)
MTU	9000
vMotion	Disabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled
vSAN Witness	Disabled
vSphere Backup NFC	Disabled
NVMe over TCP	Enabled
NVMe over RDMA	Disabled
- IPv4 settings**

IPv4 address	172.22.77.11 (static)
Subnet mask	255.255.255.0

Repeat the steps above to create the second VMkernel port for the NVMe/TCP-B port group and be sure to select NVMe/TCP-B port group and assign proper IP address for the NVMe/TCP-B VLAN.

## Add Networking

- 1 Select connection type
- 2 Select target device
- 3 Port properties
- 4 IPv4 settings
- 5 Ready to complete

## Ready to complete

Review your selections before finishing the wizard

### ▼ Select target device

Distributed port group	NVMe/TCP-B
Distributed switch	iSCSI-vDS

### ▼ Port properties

New port group	NVMe/TCP-B (iSCSI-vDS)
MTU	9000
vMotion	Disabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled
vSAN Witness	Disabled
vSphere Backup NFC	Disabled
NVMe over TCP	Enabled
NVMe over RDMA	Disabled

### ▼ IPv4 settings

IPv4 address	172.22.78.11 (static)
Subnet mask	255.255.255.0

Review the VMkernel adapter information to confirm.

fpsa-asa-esxi-01.nva.local ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage > VMkernel adapters

Networking > ADD NETWORKING... REFRESH

	Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
Virtual switches						
VMkernel adapters	vmk0	Management Network	vSwitch0	172.22.72.11	Default	Management
Physical adapters	vmk1	iSCSI-A	iSCSI-vDS	172.22.73.11	Default	--
TCP/IP configuration	vmk2	iSCSI-B	iSCSI-vDS	172.22.74.11	Default	--
Virtual Machines	vmk3	vMotion	vDS0	172.22.75.11	vMotion	vMotion
System	vmk4	NVMe/TCP-A	iSCSI-vDS	172.22.77.11	Default	NVMe over TCP
Hardware	vmk5	NVMe/TCP-B	iSCSI-vDS	172.22.78.11	Default	NVMe over TCP
Virtual Flash						

Repeat the above steps on the remaining hosts in the cluster.

fpsa-asa-esxi-02.nva.local ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage > VMkernel adapters

Networking > ADD NETWORKING... REFRESH

	Device	Network Label	Switch	IP Address	TCP/IP Stack	Enabled Services
Virtual switches						
VMkernel adapters	vmk0	Management Network	vSwitch0	172.22.72.12	Default	Management
Physical adapters	vmk1	iSCSI-A	iSCSI-vDS	172.22.73.12	Default	--
TCP/IP configuration	vmk2	iSCSI-B	iSCSI-vDS	172.22.74.12	Default	--
Virtual Machines	vmk3	vMotion	vDS0	172.22.75.12	vMotion	vMotion
System	vmk4	NVMe/TCP-A	iSCSI-vDS	172.22.77.12	Default	NVMe over TCP
Hardware	vmk5	NVMe/TCP-B	iSCSI-vDS	172.22.78.12	Default	NVMe over TCP
Virtual Flash						

Confirm host accessibility to NVMe/TCP LIFs on the target by login into target and ping the hosts' vmkernel ports as shown in the example below.

```
fpsa-a50-u0909::> network ping -lif nvmetcp-lif-01a -vserver svml -destination 172.22.77.11
172.22.77.11 is alive

fpsa-a50-u0909::> network ping -lif nvmetcp-lif-01a -vserver svml -destination 172.22.77.12
172.22.77.12 is alive

fpsa-a50-u0909::> network ping -lif nvmetcp-lif-01b -vserver svml -destination 172.22.78.11
172.22.78.11 is alive

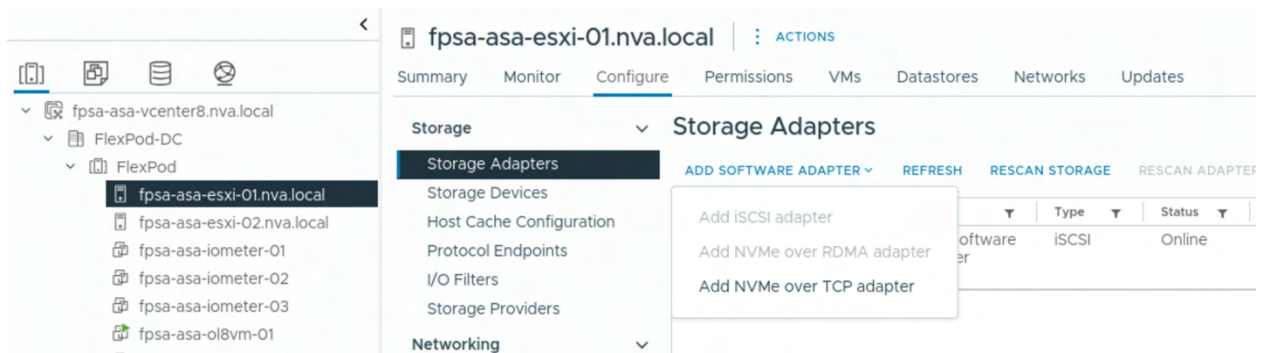
fpsa-a50-u0909::> network ping -lif nvmetcp-lif-01b -vserver svml -destination 172.22.78.12
172.22.78.12 is alive
```

**Note:** You can use any NVMe/TCP LIFs connected to the SAN fabric A to ping the hosts' vmkernel port configured for the NVMe/TCP-A port group. Similarly, ping the hosts' vmkernel port configured for the NVMe/TCP-B port group by using a LIF connected to the SAN fabric B.

## Add NVMe/TCP storage adapter on ESXi hosts

To add the NVMe over TCP adapters on the ESXi hosts, follow the steps below and perform the steps on each host in the cluster.

1. Login to vCenter.
2. Select a host listed under the vCenter Inventory view, select Configure tab.
3. Select Storage Adapters, open the drop-down list for ADD SOFTWARE ADAPTER, and select Add NVMe over TCP adapter.



4. Select the vmnic4 adapter configured for NVMe/TCP-A traffic.



5. Repeat the steps above and select vmnic5 adapter configured for NVMe/TCP-B traffic.

## Add Software NVMe over TCP adapter

fpsa-asa-esxi-01.nva.local

Enable software NVMe adapter on the selected physical network adapter.

Physical Network Adapter

vmnic5/nenic

CANCEL

OK

- Select the first NVMe over TCP adapter added in the upper pane, click the Controllers tab in the lower pane and then click ADD CONTROLLER.

The screenshot shows the NetApp FlexPod configuration interface. On the left, a tree view shows the hierarchy: **fpsa-asa-esxi-01.nva.local** > **FlexPod-DC** > **FlexPod** > **fpsa-asa-esxi-01.nva.local**. The main panel is titled **fpsa-asa-esxi-01.nva.local** and has tabs for **Summary**, **Monitor**, **Configure** (selected), **Permissions**, **VMs**, **Datastores**, **Networks**, and **Updates**. Under the **Configure** tab, there are sections for **Storage**, **Networking**, and **Virtual Machines**. The **Storage** section is expanded, showing **Storage Adapters** and **Storage Devices**. The **Storage Adapters** table is as follows:

Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba64	ISCSI Software Adapter	ISCSI	Online	iscsi_vmk(jqn.2010-11.com:flexpod:flexpod-asa-ucs-host:2)	4	5	20
vmhba65	VMware NVMe over TCP Storage Adapter	NVMe over TCP	Online	--	0	0	0
vmhba66	VMware NVMe over TCP Storage Adapter	NVMe over TCP	Online	--	0	0	0

Below the table, there are buttons for **Manage Columns** and **Export**. The **Controllers** tab is selected, showing an **ADD CONTROLLER** button and a table for adding new controllers:

Name	Subsystem NQN	Transport Type	FUSE Support	Model	Firmware Version
------	---------------	----------------	--------------	-------	------------------

- The NVMe Qualified Name (NQN) for the host is listed in the Add controller dialog. The host NQN will need to be added to the host of the NVMe subsystem created in storage to allow host access to the mapped namespaces.
- Enter the NVMe/TCP LIF created for the first storage controller to access NVMe/TCP-A fabric. Enter 8009 for the port number and then click DISCOVER CONTROLLERS.



## Add controller | vmhba66



Automatically

Manually

Host NQN

nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01



IP

172.22.78.101

Enter IPv4 / IPv6 address

☐ Central discovery controller

Port Number

8009

Range more from 0

Digest parameter

☐ Header digest

☐ Data digest

DISCOVER CONTROLLERS

Select which controller to connect

<input type="checkbox"/>	Id	Subsystem NQN	Transport Type	IP	Port Number
I-CLUSTER-NVME					
<input checked="" type="checkbox"/>	65535	nqn.1992-08.com.netap: p:sn.57cd8838ea2911ef9 608d039eac6a795:subs ystem.FlexPod-ASA-esx i-cluster-nvme	nvm	172.22.78.101	4420
<input type="checkbox"/>	65535	nqn.1992-08.com.netap	nvm	172.22.77.101	4420
<input checked="" type="checkbox"/> 2	Manage Columns		Deselect All		8 items

12. Select the two storage controllers with IPs that are in the NVMe-TCP-B fabric and click OK.

13. Repeat the steps above for each host in the cluster.

## Check for the mapped NVMe namespace

To check for the NVMe namespaces available to the host, perform the steps below.

1. Click to select one of the hosts from the vCenter Inventory view and click on the Configure tab.
2. Click Storage Devices to see the list of discovered mapped devices and click to select one of the devices with name starting with NVMe TCP DISK, assuming an NVMe namespace has already been configured in ONTAP and mapped to the hosts in the VMware cluster.



Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type	Transport
NETAPP SCSI Disk (naa.600a098038323448723f5877434a5253)	3	disk	100.00 GB	vmware vcs	Attached	Supported	Flash	ISCSI
NVMe TCP Disk (uuid.8d94d8de0d5f1f0a707d039eac6a795)	0	disk	1.00 TB	Not Consumed	Attached	Supported	Flash	TCPTTRANSPORT

Runtime Name	Status	Target	Transport	Name	Preferred
vmhba66:CO:TL0	Active (I/O)		TCPTTRANSPORT	vmhba66:CO:TL0	No
vmhba66:CO:TO:LO	Active (I/O)		TCPTTRANSPORT	vmhba66:CO:TO:LO	No
vmhba65:CO:TL0	Active (I/O)		TCPTTRANSPORT	vmhba65:CO:TL0	No
vmhba65:CO:TO:LO	Active (I/O)		TCPTTRANSPORT	vmhba65:CO:TO:LO	No

- Click on the Paths tab in the lower pane to see the four different paths that can be used to access the device. There are two paths through each NVMe over TCP adapter for the connections to the two storage controllers. All four paths should show Active (I/O) Status because the NetApp ASA controllers provide symmetric active/active for SCSI LUNs and NVMe namespaces.

## Create a datastore from the mapped NVMe namespace

To creation a datastore from a discovered NVMe namespace, follow the steps below.

- Select the Storage view in vCenter, right-click the FlexPod-DC, select Storage, and click New Datastore.

- Select VMFS datastore type and click NEXT.

### New Datastore

- Type
- Name and device selection
- Partition configuration
- Ready to complete

### Type

Specify datastore type.

☒ VMFS  
Create a VMFS datastore on a disk/LUN.

☐ NFS  
Create an NFS datastore on an NFS share over the network.

☐ vVol  
Create a Virtual Volumes datastore on a storage container connected to a storage provider.

- Provide a name for the datastore. Select a host from the drop-down list to view and select a device and then click NEXT.

New Datastore

1 Type

2 Name and device selection

3 VMFS version

4 Partition configuration

5 Ready to complete

Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

Name

nvme\_datastore\_1

1

The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

×

Select a host

fpsa-asa-esxi-01.nva.local

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Cluster VM Support
	NVMe TCP Disk (uuid:8d94d8de0d1511f0a707d039eac6a795)	0	1.00 TB	Supported	Flash	512e	No

Manage Columns

Export

1 item

- Keep the default VMFS 6 for VMFS version and click NEXT.

New Datastore

1 Type

2 Name and device selection

3 VMFS version

VMFS version

Specify the VMFS version for the datastore.

•

VMFS 6

VMFS 6 enables advanced format (512e) and automatic space reclamation support.

○

VMFS 5

VMFS 5 enables 2+TB LUN support.

- Review the disk partition configuration information and click NEXT.

New Datastore

1 Type

2 Name and device selection

3 VMFS version

4 Partition configuration

5 Ready to complete

Partition configuration

Review the disk layout and specify partition configuration details.

Partition Configuration

Use all available partitions

Datastore Size

1024

GB

Block size

1 MB

Space Reclamation Granularity

1 MB

Space Reclamation Priority

Low

Empty: 1.0 TB

Free Space:

1024GB

Usage on selected partition:

1TB

- Review the selections and click FINISH.

276

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture

### New Datastore

- Type
- Name and device selection
- VMFS version
- Partition configuration
- Ready to complete

### Ready to complete

Review your selections before finishing the wizard

▼ Name and device selection

Datastore namenvme\_datastore\_1  
Disk/LUNNVMe TCP Disk (uuid.8d94d8de0d151f0a707d039eac6a795)

▼ VMFS version

VersionVMFS 6

▼ Partition configuration

Datastore size1.00 TB  
Partition formatGPT  
Block size1 MB  
Space reclamation granularity1 MB  
Space reclamation priorityLow: Deleted or unmapped blocks are reclaimed on the LUN at low priority

7. Confirm the datastore creation by check the datastores from the storage view.

FlexPod-DC

Summary

Monitor

Configure

Permissions

Hosts & Clusters

VMs

**Datastores**

Networks

Updates

iscsi\_datastore\_1

nvme\_datastore\_1

vmware\_swap

vmware\_vcls

win2022vm\_sql\_datastore

Quick Filter

Enter value

<input type="checkbox"/>	Name	Status	Type	Datastore Cluster	Capacity	Free
<input type="checkbox"/>	iscsi_datastore_1	✓ al	Norm VMFS 6		5 TB	2.5 TB
<input type="checkbox"/>	nvme_datastore_1	✓ al	Norm VMFS 6		1,023.75 GB	1,022.32 GB
<input type="checkbox"/>	vmware_swap	✓ al	Norm VMFS 6		399.75 GB	247.06 GB
<input type="checkbox"/>	vmware_vcls	✓ al	Norm VMFS 6		99.75 GB	98.34 GB
<input type="checkbox"/>	win2022vm_sql_datastore	✓ al	Norm VMFS 6		499.75 GB	310.73 GB

## Appendix H: Configure bare-metal Oracle Linux for NVMe/TCP protocol access to storage

Oracle Linux 8.10 includes in-kernel NVMe multipathing support for NVMe namespaces. The OS native `nvme-cli` command can be used to display the mapped NVMe namespaces. The procedures below highlight the steps to configure bare-metal Oracle Linux for NVMe/TCP protocol access to the mapped NVMe namespaces from storage.

### Check and configure NVMe support on Oracle Linux

To check for NVMe support on Oracle Linux, follow the steps below.

1. Check for kernel version.

```
[admin@fpsa-asa-linux-01 ~]$ sudo uname -r
5.15.0-206.153.7.1.el8uek.x86_64
```

Check to see if the nvme related kernel modules are running. (nvme, nvme\_core, nvme\_common, nvme-tcp, etc.)

```
[admin@fpsa-asa-linux-01 ~]$ sudo lsmod | grep nvme
```

Start the needed nvme related kernel modules if they are not already running.

```
[admin@fpsa-asa-linux-01 ~]$ sudo modprobe nvme
[admin@fpsa-asa-linux-01 ~]$ sudo modprobe nvme-tcp
[admin@fpsa-asa-linux-01 ~]$ lsmod | grep nvme
```

```
nvme_tcp          57344  0
nvme_fabrics      36864  1 nvme_tcp
nvme              61440  0
nvme_core         208896  3 nvme_tcp,nvme,nvme_fabrics
nvme_common       24576  1 nvme_core
t10_pi            16384  2 sd_mod,nvme_core
```

Verify that in-kernel NVMe multipath is enabled.

```
[admin@fpsa-asa-linux-01 ~]$ sudo cat
/sys/module/nvme_core/parameters/multipath
Y
```

Check for installed nvme-cli package.

```
[admin@fpsa-asa-linux-01 ~]$ sudo dnf list installed | grep nvme-cli
nvme-cli.x86_64          1.16-9.el8
@anaconda
```

Check for the host NQN in /etc/nvme/hostnqn file.

```
[admin@fpsa-asa-linux-01 ~]$ sudo cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:0000a0aa-0000-0100-aaa0-000000000004
```

Configure dm-multipath configuration file /etc/multipath.conf to exclude NVMe namespaces so the in-kernel multipathing is used for NVMe namespaces by adding enable\_foreign parameter and setting it to NONE.

```
[admin@fpsa-asa-linux-01 ~]$ sudo cat /etc/multipath.conf
defaults {
    find_multipaths yes
    user_friendly_names yes
    enable_foreign NONE
}

blacklist {
}
```

Restart multipathd to use the updated configuration.

```
[admin@fpsa-asa-linux-01 ~]$ sudo systemctl restart multipathd
```

## Configure network interfaces for NVMe VLAN access

1. Create NVMe-TCP-A and NVMe-TCP-B connections with their specific VLANs on the same physical Ethernet interfaces used for iSCSI-A and iSCSI-B access, respectively.

```
[admin@fpsa-asa-linux-01 ~]$ sudo nmcli con add type vlan con-name NVMe-TCP-A dev eno9 id 2277
ip4 172.22.77.14/24
Connection 'NVMe-TCP-A' (9d06ca0b-4b2d-4245-b160-f8331d0b4e86) successfully added.

[admin@fpsa-asa-linux-01 ~]$ sudo nmcli con add type vlan con-name NVMe-TCP-B dev eno10 id 2278
ip4 172.22.78.14/24
Connection 'NVMe-TCP-B' (ca91b80e-c9b6-4998-8260-da229b34289c) successfully added.
```

Check the physical connectivity by pinging the target NVMe-TCP LIFs with jumbo frame.

```
[admin@fpsa-asa-linux-01 ~]$ ping -c 1 172.22.77.101 -s 9000
PING 172.22.77.101 (172.22.77.101) 9000(9028) bytes of data.
9008 bytes from 172.22.77.101: icmp_seq=1 ttl=64 time=0.144 ms

--- 172.22.77.101 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.144/0.144/0.144/0.000 ms

[admin@fpsa-asa-linux-01 ~]$ ping -c 1 172.22.77.102 -s 9000
PING 172.22.77.102 (172.22.77.102) 9000(9028) bytes of data.
9008 bytes from 172.22.77.102: icmp_seq=1 ttl=64 time=0.137 ms

--- 172.22.77.102 ping statistics ---
```

```

1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.137/0.137/0.137/0.000 ms

[admin@fpsa-asa-linux-01 ~]$ ping -c 1 172.22.78.101 -s 9000
PING 172.22.78.101 (172.22.78.101) 9000(9028) bytes of data.
9008 bytes from 172.22.78.101: icmp_seq=1 ttl=64 time=0.208 ms

--- 172.22.78.101 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.208/0.208/0.208/0.000 ms

[admin@fpsa-asa-linux-01 ~]$ ping -c 1 172.22.78.102 -s 9000
PING 172.22.78.102 (172.22.78.102) 9000(9028) bytes of data.
9008 bytes from 172.22.78.102: icmp_seq=1 ttl=64 time=0.214 ms

--- 172.22.78.102 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.214/0.214/0.214/0.000 ms

```

## Create NVMe namespaces in storage for Oracle database

1. To create NVMe namespaces in storage for later Oracle ASM disks group creation, follow the example below.

```

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_redolog_nvme_1 -size 200g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_redolog_nvme_2 -size 200g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_1 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_2 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_3 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_4 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_5 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_6 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_7 -size 800g
-ostype linux

fpsa-a50-u0909:> nvme namespace create -path fpsa_asa_linux_01_oracle_slobdata_nvme_8 -size 800g
-ostype linux

```

Check and confirm the namespace creations.

```

fpsa-a50-u0909:> nvme namespace show -path fpsa_asa_linux_01*
(vserver nvme namespace show)
Vserver Path                               State      Size Subsystem      NSID
-----
svml
  fpsa_asa_linux_01_oracle_redolog_nvme_1 online 200GB -
  fpsa_asa_linux_01_oracle_redolog_nvme_2 online 200GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_1 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_2 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_3 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_4 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_5 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_6 online 800GB -
  fpsa_asa_linux_01_oracle_slobdata_nvme_7 online 800GB -

```

```
fpsa_asa_linux_01_oracle_slobdata_nvme_8 online 800GB -  
10 entries were displayed.
```

## Create NVMe subsystem for the host

### 1. Create an NVMe subsystem with Linux OS type.

```
fpsa-a50-u0909:> nvme subsystem create -subsystem FlexPod-ASA-linux-cluster-nvme -ostype linux  
(vserver nvme subsystem create)
```

Check to confirm the NVMe subsystem creation.

```
fpsa-a50-u0909:> nvme subsystem show -subsystem FlexPod-ASA-linux-cluster-nvme  
(vserver nvme subsystem show)  
Vserver Subsystem      Target NQN  
-----  
svml  
      FlexPod-ASA-linux-cluster-nvme nqn.1992-  
08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-ASA-linux-cluster-nvme
```

## Add host NQN to the NVMe subsystem host

To add the host NQN to NVMe subsystem, follow the steps below.

### 2. Look up the host NQN from /etc/nvme/hostnqn file.

```
[admin@fpsa-asa-linux-01 ~]$ sudo cat /etc/nvme/hostnqn  
nqn.2014-08.org.nvmexpress:uuid:0000a0aa-0000-0100-aaa0-000000000004
```

### 3. Add the host NQN to the NVMe subsystem host.

```
fpsa-a50-u0909:> nvme subsystem host add -subsystem FlexPod-ASA-linux-cluster-nvme -host-nqn  
nqn.2014-08.org.nvmexpress:uuid:0000a0aa-0000-0100-aaa0-000000000004
```

### 4. Check to confirm that the host NQN has been added.

```
fpsa-a50-u0909:> nvme subsystem host show -subsystem FlexPod-ASA-linux-cluster-nvme  
(vserver nvme subsystem host show)  
Vserver Subsystem Priority Host NQN  
-----  
svml      FlexPod-ASA-linux-cluster-nvme  
          regular      nqn.2014-08.org.nvmexpress:uuid:0000a0aa-0000-0100-aaa0-000000000004
```

## Map the namespaces to the subsystem

To map the namespaces the NVMe subsystem which includes the host NQN, follow the steps below.

### 1. Add the name spaces one by one to the NVMe subsystem.

```
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_redolog_nvme_1  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_redolog_nvme_2  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_slobdata_nvme_1  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_slobdata_nvme_2  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_slobdata_nvme_3  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_slobdata_nvme_4  
  
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path  
fpsa_asa_linux_01_oracle_slobdata_nvme_5
```



```
fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path
fpsa_asa_linux_01_oracle_slobdata_nvme_6

fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path
fpsa_asa_linux_01_oracle_slobdata_nvme_7

fpsa-a50-u0909:> nvme subsystem map add -subsystem FlexPod-ASA-linux-cluster-nvme -path
fpsa_asa_linux_01_oracle_slobdata_nvme_8
```

## 2. Check to confirm that the NVMe namespaces have been mapped.

```
fpsa-a50-u0909:> nvme subsystem map show -subsystem FlexPod-ASA-linux-cluster-nvme
(vserver nvme subsystem map show)
Vserver      Subsystem      NSID Namespace Path
-----
svml1        FlexPod-ASA-linux-cluster-nvme
              00000001h fpsa_asa_linux_01_oracle_redolog_nvme_1
              00000002h fpsa_asa_linux_01_oracle_redolog_nvme_2
              00000003h fpsa_asa_linux_01_oracle_slobdata_nvme_1
              00000004h fpsa_asa_linux_01_oracle_slobdata_nvme_2
              00000005h fpsa_asa_linux_01_oracle_slobdata_nvme_3
              00000006h fpsa_asa_linux_01_oracle_slobdata_nvme_4
              00000007h fpsa_asa_linux_01_oracle_slobdata_nvme_5
              00000008h fpsa_asa_linux_01_oracle_slobdata_nvme_6
              00000009h fpsa_asa_linux_01_oracle_slobdata_nvme_7
              0000000Ah fpsa_asa_linux_01_oracle_slobdata_nvme_8
10 entries were displayed.
```

## Configure host to access the mapped namespaces with consistent device name

To configure the host to discover and connect to the mapped NVMe namespaces and to create udev rules under /etc/udev/rules.d to map the namespaces consistently with aliases for Oracle ASM disk group creation, follow the procedures below.

### 1. Run the nvme discover command with all initiator-target LIF combinations to discover subsystems.

```
[admin@fpsa-asa-linux-01 ~]$ sudo nvme discover -t tcp -w 172.22.77.14 -a 172.22.77.101
[admin@fpsa-asa-linux-01 ~]$ sudo nvme discover -t tcp -w 172.22.77.14 -a 172.22.77.102
[admin@fpsa-asa-linux-01 ~]$ sudo nvme discover -t tcp -w 172.22.78.14 -a 172.22.78.101
[admin@fpsa-asa-linux-01 ~]$ sudo nvme discover -t tcp -w 172.22.78.14 -a 172.22.78.102
```

Run the nvme connect-all command across all initiator-target LIFs to connect the NVMe namespaces.

```
[admin@fpsa-asa-linux-01 ~]$ sudo nvme connect-all -t tcp -w 172.22.77.14 -a 172.22.77.101 -l -1
[admin@fpsa-asa-linux-01 ~]$ sudo nvme connect-all -t tcp -w 172.22.77.14 -a 172.22.77.102 -l -1
[admin@fpsa-asa-linux-01 ~]$ sudo nvme connect-all -t tcp -w 172.22.78.14 -a 172.22.78.101 -l -1
[admin@fpsa-asa-linux-01 ~]$ sudo nvme connect-all -t tcp -w 172.22.78.14 -a 172.22.78.102 -l -1
```

**Note:** NetApp recommends setting the ctrl-loss-tmo option to -1 so that the NVMe/TCP initiator attempts to reconnect indefinitely in the event of a path loss.

### 2. Run the nvme tool with ONTAP tools option to obtain namespace UUID and path mapping information.

```
[admin@fpsa-asa-linux-01 ~]$ sudo nvme netapp ontapdevices -o json
[sudo] password for admin:
{
  "ONTAPdevices" : [
    {
      "Device" : "/dev/nvme0n1",
      "Vserver" : "svml1",
      "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_redolog_nvme_1/blocks",
      "NSID" : 1,
      "UUID" : "704f70ff-0d75-11f0-a707-d039eac6a795",
      "Size" : "214.75GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 52428800
    }
  ],
}
```

```
{
  "Device" : "/dev/nvme0n2",
  "Vserver" : "svml",
  "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_redolog_nvme_2/blocks",
  "NSID" : 2,
  "UUID" : "76dde16c-0d75-11f0-a707-d039eac6a795",
  "Size" : "214.75GB",
  "LBA_Data_Size" : 4096,
  "Namespace_Size" : 52428800
},
...
```

3. Create udev rules for the redolog and slobdata NVMe namespace devices to map them to persistent device aliases using the UUID information. See the following for partial examples.

```
[admin@fpsa-asa-linux-01 rules.d]$ pwd
/etc/udev/rules.d

[admin@fpsa-asa-linux-01 rules.d]$ sudo cat 74-nvme-redolog.rules
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="704f70ff-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeredolog1", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="76dde16c-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeredolog2", OWNER="grid", GROUP="oinstall", MODE="0660"

[admin@fpsa-asa-linux-01 rules.d]$ sudo cat 75-nvme-slobdata.rules
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="8c541a42-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata1", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="90b68f3d-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata2", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="97b0b2f0-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata3", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="9dc68622-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata4", OWNER="grid", GROUP="oinstall", MODE="0660"
...
```

Reload and activate the new udev rules to create aliases for the NVMe devices.

```
[admin@fpsa-asa-linux-01 rules.d]$ sudo udevadm control --reload
[admin@fpsa-asa-linux-01 rules.d]$ sudo udevadm trigger --type=devices --action=change
```

Check and confirm the NVMe device persistent naming.

```
[admin@fpsa-asa-linux-01 rules.d]$ sudo ls -l /dev/nvmeredolog*
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeredolog1 -> nvme0n1
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeredolog2 -> nvme0n2

[admin@fpsa-asa-linux-01 rules.d]$ sudo ls -l /dev/nvmeslobdata*
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata1 -> nvme0n3
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata2 -> nvme0n4
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata3 -> nvme0n5
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata4 -> nvme0n6
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata5 -> nvme0n7
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata6 -> nvme0n8
lrwxrwxrwx 1 root root 7 Mar 30 22:23 /dev/nvmeslobdata7 -> nvme0n9
lrwxrwxrwx 1 root root 8 Mar 30 22:23 /dev/nvmeslobdata8 -> nvme0n10
```


## Appendix I: Infrastructure configuration updates for direct-attached FC-based SAN storage support

The following subsections provide the infrastructure configuration updates necessary to enable FC-based SAN for the UCS X-Series Direct with direct-attached ASA FC storage. Please refer to additional appendices for details on the installation and configuration of FC SAN booted ESXi host, adding NVMe/FC access for the FC SAN booted ESXi host, adding FC and NVMe/FC access for iSCSI SAN booted ESXi host, and adding FC and NVMe/FC access for the bare-metal Oracle Linux host.

## UCS X-Series Direct and ASA storage Fibre Channel connectivity

Ports 1 and 2 of the Cisco UCS S9108 Fabric Interconnect support FC breakout connections with DS-SFP-4x32G-SW transceiver. By configuring port 1 of both FIs to support 4x32G FC breakout and connecting the UCS X-Series Direct and ASA storage as described in the connectivity diagram Figure 19 using connections detailed in Table 14 below, the ASA A50 storage system will have FC connectivity to the two independent FC SAN fabrics provided by FI A and FI B.

**Figure 19 UCS X-Series Direct and ASA storage Fibre Channel connectivity diagram**



**Table 14 UCS X-Series Direct and ASA storage Fibre Channel connections**

UCS X-Series Direct	Port	ASA Storage	Port
Cisco UCS Fabric Interconnect 9108 100G A	FC 1/1	NetApp ASA A50 A	1a
	FC 1/2	NetApp ASA A50 A	1b
	FC 1/3	NetApp ASA A50 B	1a
	FC 1/4	NetApp ASA A50 B	1b
Cisco UCS Fabric Interconnect 9108 100G B	FC 1/1	NetApp ASA A50 A	1c
	FC 1/2	NetApp ASA A50 A	1d
	FC 1/3	NetApp ASA A50 B	1c
	FC 1/4	NetApp ASA A50 B	1d

## Create UCS domain switch control policy

To support direct-attached NetApp ASA FC storage, the FI FC Switching Mode setting should be configured as Switch mode to provide FC services. Follow the steps below to configure the UCS domain switch control policy and select Switch mode for FC and End Host mode for Ethernet.

1. Login to Cisco Intersight.
2. Navigate to Configure > Policies and click Create Policy.
3. Select UCS Domain for Platform Type, select Switch Control for Policy Type, and click Start.
4. Select the Organization, provide a Name for the policy, and optionally set Tags or provide a Description, then click Next.
5. In the Policy Details section, select End Host for Ethernet Switching Mode, select Switch for FC Switching Mode, and leave the VLAN configuration unchanged.

General

2 Policy Details

### Policy Details

Add policy details.

**This policy is applicable only for UCS Domains**

#### Switching Mode

Ethernet ☐ FC ☐

☒ End Host ☐ Switch ☐ End Host ☒ Switch

#### VLAN Port Count

☐ Enable VLAN Port Count Optimization ⓘ

#### System Reserved VLANs

**By default, the reserved VLAN Start ID is 3915 and a range of 128 VLANs starting from this ID is blocked for internal use. This range of VLAN IDs cannot be used in any VLAN policy. To change the default range, configure a new Reserved VLAN Start ID.**

Reserved VLAN Start ID \* ⓘ  Reserved VLAN End ID

6. Click Create to create the policy.

## Create UCS WWNN and WWPN pools

To create WWNN and WWPN pools in Intersight, follow the steps below.

1. Navigate to Configure > Pools and click Create Pool.
2. Select option WWNN and click Start.
3. Provide a Name for the pool and optionally set Tags or provide a Description, then click Next.
4. In the WWNN Pool Create screen, select FlexPod-ASA organization, enter FlexPod-ASA-WWNN-Pool for Name. Click Next.

1 General

2 Pool Details

### General

Pool represents a collection of WWN addresses that can be allocated to VHBAs of a Server Profile

Organization \*

Name \*

Set Tags

Description

0 / 1024

5. Add WWNN block From and Size information for the pool according to your environment requirements and click Create to create the WWNN pool policy.

General

2 Pool Details

### Pool Details

Block of WWNN Identifiers.

#### WWNN Blocks

From ⓘ	Size ⓘ
<input type="text" value="20:00:00:25:B5:15:00:00"/>	<input type="text" value="512"/>

1 - 1024

6. Navigate to Configure > Pools and click Create Pool.

7. Select option WWPN and click Start.
8. In the WWPN Pool Create screen, select FlexPod-ASA organization, enter FlexPod-ASA-WWPN-Pool for Name. Click Next.

**1 General**

**2 Pool Details**

**General**  
Pool represents a collection of WWN addresses that can be allocated to VHBA's of a Server Profile

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-WWPN-Pool

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description  
0 / 1024

9. Add WWPN block From and Size information for the pool according to your environment requirements and click Create to create the WWPN pool policy.

**General**

**2 Pool Details**

**Pool Details**  
Block of WWPN Identifiers.

**WWPN Blocks**

From	Size
20:00:00:25:B5:09:00:00	512

1 - 1024

## Create UCS VSAN policies

Two VSAN policies are needed for FC protocol as different FC fabrics use different VSAN IDs. For the X-Series Direct with UCS domain's FC Switching Mode set to Switch, FI-A and FI-B serve as the two independent FC fabrics. Follow the steps below to configure VSAN Policies for the Cisco UCS Domain profile.

1. Navigate to go to Configure > Policies and click Create Policy.
2. Select UCS Domain for Platform Type, select VSAN for Policy Type, and click Start.
3. In the VSAN Create General section, select the Organization from drop-down list, provide policy Name, optionally set Tags or provide a Description, and click Next.

**1 General**

**2 Policy Details**

**General**  
Add a name, description, and tag for the policy.

**Organization \***  
FlexPod-ASA

**Name \***  
FlexPod-ASA-VSAN-A

**Set Tags**  
Enter a tag in the key:value format.

**Description**  
Description  
0 / 1024

4. In the Policy Details section, select Add VSAN.

5. In the Add VSAN dialog, provide VSAN Name, VSAN Scope, VSAN ID, and FCoE VLAN ID as appropriate for your environment for FI A. Click Add.

**Add VSAN**

Name \* ⓘ  
VSAN A ⓘ

VSAN Scope ⓘ  
☐ Storage & Uplink ⓘ ☒ Storage ⓘ ☐ Uplink ⓘ

VSAN ID \* ⓘ  
121 ⓘ  
1 - 4093

FCoE VLAN ID \* ⓘ  
121 ⓘ

Cancel Add

**Note:** For this direct-attached storage design, we are selecting the Storage VSAN Scope. Storage VSAN scope allows you to connect and configure direct-attached storage, with the fabric interconnect FC Switching Mode configured as Switch. You can configure local zones on this VSAN using FC Zone policies and attaching them to the corresponding vHBAs.

**Policy Details**  
Add policy details.

*i* This policy is applicable only for UCS Domains

☐ Uplink Trunking ⓘ

**Add VSAN**

Name	VSAN ID	VSAN Scope	FCoE VLAN ID ⓘ
VSAN A	121	Storage	121

Rows per page 10 < 1 >

6. Click Create to create the VSAN policy.
7. Repeat steps 1 – 6 to create another VSAN policy for FI-B.

**General**  
Add a name, description, and tag for the policy.

Organization \*  
FlexPod-ASA

Name \*  
FlexPod-ASA-VSAN-B ⓘ

Set Tags  
Enter a tag in the key:value format.

Description  
Description

0 / 1024



General

2 Policy Details

### Policy Details

Add policy details.

This policy is applicable only for UCS Domains

Uplink Trunking

Add VSAN

Name	VSAN ID	VSAN Scope	FCoE VLAN ID
VSAN B	122	Storage	122

Rows per page 10 < 1 >

## Create UCS FI port policies to support both IP and FC traffic

Instead of creating new FI port policies from scratch to support both IP and FC traffic, we will clone the IP port policies created previously and update them to also support FC traffic.

1. Navigate to Configure > Policies and search for the FlexPod ASA Port Polices.
2. Clone FlexPod-ASA-Port-Policy-A to FlexPod-ASA-Port-Policy-IP-FC-A and clone FlexPod-ASA-Port-Policy-B to FlexPod-ASA-Port-Policy-IP-FC-B. Be sure to select the FlexPod-ASA organization when cloning.
3. Edit port policy FlexPod-ASA-Port-Policy-IP-FC-A.
4. Under Unified Port step, move the slider to configure port 1 as FC port and click Next.

General

2 Unified Port

3 Breakout Options

4 Port Roles

### Unified Port

Configure the port modes to carry FC or Ethernet traffic.

Move slider to configure unified ports.

Configuring unified port will result in the deletion of previously configured breakout ports, port roles and port channels roles.

Fibre Channel Ports
 

1 Fibre Channel Ports (Port 1)

FC

Ports 1-1

Ethernet

Ports 2-8

5. Under Breakout Options, select Fiber Channel tab. Then, select Port 1 and click Configure.

- ✓ General
- ✓ Unified Port
- 3 Breakout Options
- 4 Port Roles

### Breakout Options

Configure breakout ports on FC or Ethernet.

Ethernet

Fibre Channel

Configure

Selected Ports

Port 1

Clear Selection



FC Ethernet Port Modes

<input checked="" type="checkbox"/> Port	Type	Speed	Breakout Ports	
<input checked="" type="checkbox"/> Port 1	FC	16G	Port 1/1, Port 1/2, Port 1/...	

Selected 1 of 1

Show Selected

Unselect All

6. Select the breakout port speed, click Set, and then click Next.

- ✓ General
- ✓ Unified Port
- 3 Breakout Options
- 4 Port Roles

### Breakout Options

Configure breakout ports on FC or Ethernet.

Ethernet

Fibre Channel

#### Set Breakout



Modifying the speed of an existing FC breakout port, will result in the deletion of previously configured port roles and port channel roles.

Selected Ports

Port 1

☐ 4x8G

☐ 4x16G

☒ 4x32G

Cancel

Set

Selected 1 of 1

Show Selected

Unselect All

7. Select all the configured FC breakout ports and click Configure.
8. Select FC Storage from the drop-down list for Role and provide the VSAN ID.

### Configure (4 Ports)

Configuration

Selected Ports

Port 1/1, Port 1/2, Port 1/3, Port 1/4

Role

FC Storage

Admin Speed ⓘ

32Gbps

VSAN ID \* ⓘ

121

1 - 4093

9. Click Save to save the port role configuration and click Save again to save the policy.

- ✓ General
- ✓ Unified Port
- ✓ Breakout Options
- 4 Port Roles

## Port Roles Port Channels Pin Groups

Configure

Selected Ports -



FC Storage Unconfigured Appliance Ethernet Uplink Port Channel

Export

<input type="checkbox"/>	Name	Type	Role	Connected ...	Device Nu...	Port Channe	
<input type="checkbox"/>	port 1/1	FC	FC Storage			-	
<input type="checkbox"/>	port 1/2	FC	FC Storage			-	
<input type="checkbox"/>	port 1/3	FC	FC Storage			-	
<input type="checkbox"/>	port 1/4	FC	FC Storage			-	
<input type="checkbox"/>	port 2	Ethernet	Unconfigur...			-	
<input type="checkbox"/>	port 3	Ethernet	Appliance			-	
<input type="checkbox"/>	port 4	Ethernet	Appliance			-	

- Repeat steps 3 to 9 above to modify the cloned FlexPod-ASA-Port-Policy-IP-FC-B policy for fabric B with fabric B VSAN ID.

## Create UCS Fibre Channel Network Policy

To configure the Fibre Channel Network Policy for the UCS Server profile, follow the steps below.

- Navigate to Configure > Policies and click Create Policy.
- Select UCS Server for the Platform Type, select Fibre Channel Network Policy Type, and click Start.

**Note:** For this solution, we configured two Fibre Channel network policy as “FlexPod-ASA-FC-Network-A” and “FlexPod-ASA-FC-Network-B” to carry the two VSAN traffic with ID 121 and 122, respectively on Fabric Interconnects A and B.

- In the Fibre Channel Network Create General section, select the Organization from drop-down list, provide policy Name, optionally set Tags or provide a Description, and click Next.

### 1 General

### 2 Policy Details

## General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-FC-Network-A

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

- In the Policy Details section, click to select UCS Server (FI-Attached) platform, provide the VSAN ID for FC fabric A.

✓ General

2 Policy Details

### Policy Details

Add policy details.

All Platforms
UCS Server (Standalone)
UCS Server (FI-Attached)

#### Fibre Channel Network

VSAN ID ⓘ

121

^
v

1 - 4094

- Click Create to create the Fibre Channel Network policy for fabric A.
- Repeat steps 1 to 5 and create the Fibre Channel Network policy for fabric B with the appropriate policy name and VSAN ID for fabric B.

1 General

2 Policy Details

### General

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-FC-Network-B

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

✓ General

2 Policy Details

### Policy Details

Add policy details.

All Platforms
UCS Server (Standalone)
UCS Server (FI-Attached)

#### Fibre Channel Network

VSAN ID ⓘ

122

^
v

1 - 4094

## Create UCS Fibre Chanel QoS Policy

To configure the Fibre Channel QoS Policy for the UCS Server profile, follow the steps below.

- Navigate to Configure > Polices and click Create Policy.
- Select UCS Server for the Platform Type, select Fibre Channel QoS Policy Type, and click Start.
- In the Fibre Channel QoS Create General section, select the Organization from drop-down list, provide policy Name, optionally set Tags or provide a Description, and click Next.

- 1 General
- 2 Policy Details

### General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-FC-QoS

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

4. In the Policy Details section, click to select UCS Server (FI-Attached) platform and use the default FC QoS configuration.

- ✓ General
- 2 Policy Details

### Policy Details

Add policy details.



All Platforms

UCS Server (Standalone)

UCS Server (FI-Attached)

#### Fibre Channel QoS

Rate Limit, Mbps ⓘ

0

0 - 100000

Maximum Data Field Size, Bytes ⓘ

2112

256 - 2112

Burst ⓘ

10240

1 - 1000000

Priority ⓘ

FC

5. Click Create to create the Fibre Channel QoS policy.

## Create UCS Fibre Channel Adapter Policies

Fibre Channel Adapter Policy defines whether the adapter is used for FC or NVMe/FC traffic and the operating system environment. Follow the steps below to create the Fibre Channel Adapter Policies for use with VMware and Linux UCS server profiles.

1. Navigate to Configure > Policies and click Create Policy.
2. Select UCS Server for the Platform Type, select Fibre Channel Adapter Policy Type, and click Start.
3. In the Fibre Channel Adapter Create General section, select the Organization from drop-down list, provide policy Name, optionally set Tags or provide a Description.

1 General

2 Policy Details

## General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-FC-Adapter-VMware

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

Cisco Provided Fibre Channel Adapter Configuration ⓘ

[Select Cisco Provided Configuration](#)

4. Click on Select Cisco Provided Configuration to select from a list of cisco provided Fibre Channel adapter configurations.

Select Cisco Provided Configuration

Q Search Filters 9 results

Name	Description
<input type="radio"/> FCNVMInitiator	
<input type="radio"/> FCNVMTarget	
<input type="radio"/> Initiator	
<input type="radio"/> Linux	
<input type="radio"/> Solaris	
<input type="radio"/> Target	
<input checked="" type="radio"/> VMWare	
<input type="radio"/> Windows	
<input type="radio"/> WindowsBoot	

Selected 1 of 9 [Show Selected](#) [Unselect All](#) Rows per page 10 < 1 >

5. Select VMWare from the list, click Select, and click Next.



- ✓ General
- 2 Policy Details

## Policy Details

Add policy details.



All Platforms

UCS Server (Standalone)

UCS Server (FI-Attached)

### Error Recovery

☐ FCP Error Recovery ⓘ

Port Down Timeout, ms ⓘ

10000

0 - 240000

Link Down Timeout, ms ⓘ

30000

0 - 240000

I/O Retry Timeout, Seconds ⓘ

5

1 - 59

Port Down IO Retry, ms ⓘ

30

0 - 255

### Error Detection

Error Detection Timeout ⓘ

2000

1000 - 100000

### Resource Allocation

6. In the Policy Details section, click to select UCS Server (FI-Attached) platform and use the default FC adapter configuration.
7. Click Create to create the Fibre Channel Adapter policy using Cisco provided configuration for VMware.
8. Repeat steps 1 to 7 above to create another FC Adapter policy using Cisco provided configuration for Linux for UCS Server (FI-Attached) platform and default settings.

- 1 General
- 2 Policy Details

## General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-FC-Adapter-Linux

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

### Cisco Provided Fibre Channel Adapter Configuration ⓘ

Selected Cisco Provided Configuration Linux



[Edit Selection](#)



- Repeat steps 1 to 7 above to create an NVMe/FC initiator adapter policy using Cisco provided configuration for FCNVMeInitiator for UCS Server (FI-Attached) platform and use default settings.

1 General

2 Policy Details

General

Add a name, description, and tag for the policy.

Organization \*  
FlexPod-ASA

Name \*  
FlexPod-ASA-FC-Adapter-FCNVMeInitiator

Set Tags  
Enter a tag in the key:value format.

Description  
Description  
0 / 1024

Cisco Provided Fibre Channel Adapter Configuration ⓘ

Selected Cisco Provided Configuration FCNVMeInitiator ⓘ | [Edit Selection](#) |

**Note:** For the NVMe/FC initiator, the policy configuration is OS agnostic and does not require separate policies for VMware and Linux.

## Create UCS FC SAN Connectivity Policy

For FC SAN connectivity, we configured 4 vHBAs: two for FC traffic and two for NVMe/FC traffic. Additional vHBAs can be created for either FC or NVMe/FC traffic as needed.

For FC SAN boot, the first two FC vHBAs (vHBA0 and vHBA1) are configured to boot from the two independent FC SAN fabrics. vHBA0 was configured to carry FC traffic on VSAN 121 through FI-A, while vHBA1 was configured to carry FC traffic on VSAN 122 through FI-B.

The other two vHBAs (vHBA2 and vHBA3) were configured for NVMe/FC traffic on the same set of VSANs as FC. vHBA2 was configured to carry the NVMe/FC network traffic on VSAN 121 and FI-A, while vHBA3 was configured to carry the NVMe/FC network traffic on VSAN 122 and FI-B.

The server profiles were updated to include four vHBAs as listed in Table 15.

**Table 15 vHBA configuration information**

vHBA Name	vHBA Type	FI ID	PCI Order	Fibre Channel Network
vHBA0	fc-initiator	FI-A	6	FlexPod-ASA-FC-Network-A
vHBA1	fc-initiator	FI-B	7	FlexPod-ASA-FC-Network-B
vHBA2	fc-nvme-initiator	FI-A	8	FlexPod-ASA-FC-Network-A
vHBA3	fc-nvme-initiator	FI-B	9	FlexPod-ASA-FC-Network-B

To create Fibre Channel SAN Connectivity policy for the UCS Server profile, follow the steps below.

- Navigate to Configure > Policies and click Create Policy.
- Select UCS Server for the Platform Type, select SAN Connectivity Policy Type, and click Start.
- In the SAN Connectivity Create General section, select the Organization from drop-down list, provide policy Name, select UCS Server (FI-Attached) Target Platform, optionally set Tags or provide a Description, and click Next.

1

General

2

Policy Details

## General

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-SAN-Connectivity

**Target Platform** ⓘ

☐ UCS Server (Standalone)
 ☒ UCS Server (FI-Attached)

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

- In the Policy Details section, select Manual vHBAs Placement, click Select Pool to select the WWNN Pool, click to select the previously created WWNN pool, then click Select.

✓

General

2

Policy Details

## Policy Details

Add policy details.

Manual vHBAs Placement

Auto vHBAs Placement

WWNN

Pool

Static

**WWNN Pool \*** ⓘ

Selected Pool

FlexPod-ASA-WWNN-Pool

👁

✎

Edit Selection

🗑

- Click Add and click vHBA from the drop-down list.
- In the Add vHBA section, enter vHBA0 for Name and select fc-initiator vHBA Type.
- For the WWPN Pool, select the previously created FlexPod-ASA-WWPN-Pool.

### Add vHBA

**Name \*** ⓘ

vHBA0

**vHBA Type** ⓘ

fc-initiator

**Pin Group Name** ⓘ

Pin Group Name

WWPN

Pool

Static

**WWPN Pool \*** ⓘ

Selected Pool

FlexPod-ASA-WWPN-Pool

👁

✎

Edit Selection

🗑

- Select Advanced Placement option, provide MLOM for Slot ID, 0 for PCI Link, A for Switch ID, and 6 for PCI Order based on the information in Table 15 corresponding to vHBA0.

## Placement

Simple

Advanced

☒ Automatic Slot ID Assignment ⓘ

Slot ID ⓘ  

MLOM

☒ Automatic PCI Link Assignment ⓘ

☐ Load-Balanced ☒ Custom

PCI Link ⓘ  

0

0 - 1

Switch ID \* ⓘ  

A

PCI Order ⓘ  

6

Persistent LUN Bindings

☒ Persistent LUN Bindings ⓘ

9. Select FlexPod-ASA-FC-Network-A for Fibre Channel Network Policy, select FlexPod-ASA-FC-QoS for Fibre Channel QoS Policy, select FlexPod-ASA-FC-Adapter-VMware for Fibre Channel Adapter Policy, and click Add.

Fibre Channel Network Policy \* ⓘ  
Selected Policy FlexPod-ASA-FC-Network-A ⓘ | ⓘ | [Edit Selection](#) | ⓘ

Fibre Channel QoS Policy \* ⓘ  
Selected Policy FlexPod-ASA-FC-QoS ⓘ | ⓘ | [Edit Selection](#) | ⓘ

Fibre Channel Adapter Policy \* ⓘ  
Selected Policy FlexPod-ASA-FC-Adapter-VMware ⓘ | ⓘ | [Edit Selection](#) | ⓘ

FC Zone ⓘ  
[Select Policies](#)

10. Repeat steps 5 to 9 to add vHBA1 and use B for Switch ID, 7 for PCI Order, and FlexPod-ASA-FC-Network-B for Fibre Channel Network Policy.
11. Repeat steps 5 to 9 two more times to add vHBA2 and vHBA3. For vHBA Type, use fc-nvme-initiator for both. For vHBA2, use A for Switch ID, 8 for PCI Order, and FlexPod-ASA-FC-Network-A for Fibre Channel Network Policy. For vHBA3, use B for Switch ID, 9 for PCI Order, and FlexPod-ASA-FC-Network-B for Fibre Channel Network Policy.
12. After adding the four vHBAs, confirm that the Switch ID, PCI Order, and the Fibre Channel Network Policy are using the appropriate values identified in Table 15 above.

**Policy Details**  
Add policy details.

Manual vHBAs Placement | Auto vHBAs Placement

WWNN

Pool | Static

WWNN Pool \* FlexPod-ASA-WWNN-Pool [Edit Selection](#)

For manual placement option you need to specify placement for each vHBA. Learn more at [Help Center](#)

[Add](#) [Graphic vHBAs Editor](#)

Q Search [Filters](#) 4 results

<input type="checkbox"/>	Na...	Slot ID	Switch ID	PCI Order	WWPN Pool	Fibre Channel Network ...	T...	Fibre Channel Qo...	Fibre Channel Adapter Policy
<input type="checkbox"/>	vHBA0	ML0M	A	6	FlexPod-ASA-WWPN-Pool	FlexPod-ASA-FC-Network-A	-	FlexPod-ASA-FC-QoS	FlexPod-ASA-FC-Adapter-VMwar
<input type="checkbox"/>	vHBA1	ML0M	B	7	FlexPod-ASA-WWPN-Pool	FlexPod-ASA-FC-Network-B	-	FlexPod-ASA-FC-QoS	FlexPod-ASA-FC-Adapter-VMwar
<input type="checkbox"/>	vHBA2	ML0M	A	8	FlexPod-ASA-WWPN-Pool	FlexPod-ASA-FC-Network-A	-	FlexPod-ASA-FC-QoS	FlexPod-ASA-FC-Adapter-VMwar
<input type="checkbox"/>	vHBA3	ML0M	B	9	FlexPod-ASA-WWPN-Pool	FlexPod-ASA-FC-Network-B	-	FlexPod-ASA-FC-QoS	FlexPod-ASA-FC-Adapter-VMwar

Rows per page 10 < 1 >

[Cancel](#) [Back](#) [Create](#)

13. Click Create to create the policy.

**Note:** The created SAN Connectivity policy above is to be used for a VMware environment. For the bare-metal Linux environment, you can clone the above SAN Connectivity policy and modify the Fibre Channel Adapter Policy to use FlexPod-ASA-FC-Adapter-Linux.

## Update UCS domain profile to include policies for FC connectivity

To support FC connectivity in the UCS domain, the UCS domain profile needs to be updated to incorporate VSAN policies, FI port policy updates with FC port configuration, and set FC Switching mode to Switch mode. To update the existing UCS domain profile created for IP SAN to support FC SAN, follow the steps below.

1. Shut down the VMs in the VMware cluster, shutdown vCenter Server appliance, and put the ESXi hosts in maintenance mode, and then shutdown those ESXi hosts.

**Note:** If you are also using bare-metal Linux hosts in the solution environment, shutdown those Linux hosts also for this UCS domain profile update.

2. Navigate to Configure > Profiles, click UCS Domain Profiles tab, select FlexPod-ASA-Domain-Profile, click on the three dots on the right and select Edit to edit the domain profile.

**Note:** You can also create a Clone of the domain profile before modification to save a copy of the existing domain profile.

3. Click Next twice for the VLAN & VSAN Configuration section.
4. Select FlexPod-ASA-VSAN-A policy for the VSAN Configuration under Fabric Interconnect A and click Select.
5. Select FlexPod-ASA-VSAN-B policy for the VSAN Configuration under Fabric Interconnect B and click Select.

- ✓ General
- ✓ UCS Domain Assignment
- 3 VLAN & VSAN Configuration**
- 4 Ports Configuration
- 5 UCS Domain Configuration
- 6 Summary

## VLAN & VSAN Configuration

Create or select a policy for the Fabric Interconnect pair.

^ Fabric Interconnect A 2 of 2 Policies Configured

VLAN Configuration 🗑️ | ✎️ | 👁️ | ● FlexPod-ASA-VLAN

VSAN Configuration 🗑️ | ✎️ | 👁️ | ● FlexPod-ASA-VSAN-A

^ Fabric Interconnect B 2 of 2 Policies Configured

VLAN Configuration 🗑️ | ✎️ | 👁️ | ● FlexPod-ASA-VLAN

VSAN Configuration 🗑️ | ✎️ | 👁️ | ● FlexPod-ASA-VSAN-B

6. Click Next for the Ports Configuration section.
7. Under Fabric Interconnect A, click the trash bin icon to clear the current policy, and then select the updated port policy FlexPod-ASA-Port-Policy-IP-FC-A for FI A.

- ✓ General
- ✓ UCS Domain Assignment
- ✓ VLAN & VSAN Configuration
- 4 Ports Configuration**
- 5 UCS Domain Configuration
- 6 Summary

## Ports Configuration

Create or select a port policy for the Fabric Interconnect pair.

**i** Configure ports by creating or selecting a policy.

^ Fabric Interconnect A Configured

Ports Configuration Selected Policy ● FlexPod-ASA-Port-Policy-IP-FC-A 👁️ | ✎️ | 🗑️



Ports | Port Channels

● FC Storage ● Unconfigured ● Appliance ● Ethernet Uplink Port Channel

8. Under Fabric Interconnect B, click the trash bin icon to clear the current policy, and then select the updated port policy FlexPod-ASA-Port-Policy-IP-FC-B for FI B.

- ✓ General
- ✓ UCS Domain Assignment
- ✓ VLAN & VSAN Configuration
- 4 Ports Configuration**
- 5 UCS Domain Configuration
- 6 Summary

^ Fabric Interconnect B Configured

Ports Configuration Selected Policy ● FlexPod-ASA-Port-Policy-IP-FC-B 👁️ | ✎️ | 🗑️



Ports | Port Channels

● FC Storage ● Unconfigured ● Appliance ● Ethernet Uplink Port Channel

9. Click Next for UCS Domain Configuration. Under Switch Control, click Select Policy, select the FlexPod-ASA-Switch-Control policy, and click Select.
10. Click Next, click Deploy, and click Deploy again to confirm the deployment of the updated UCS domain profile.



**Note:** It will take some time to deploy the updated UCS domain profile. You can monitor the status of the deployment under UCS Domain Profiles tab and see the deployment execution flow under Intersight Requests.

The first screenshot shows the 'Profiles' section in Intersight, specifically the 'UCS Domain Profiles' tab. It displays a table with one profile, 'FlexPod-ASA-Domain-Profile', which is in a 'Configuring' state. The second screenshot shows the 'Requests' section, filtered to show 'In Progress' requests. Two requests for 'Deploy Domain Profile' are visible, both in progress. The third screenshot provides a detailed view of the 'Deploy Domain Profile' request, showing its status as 'In Progress' and a progress bar at 53%. The execution flow includes steps like 'Wait for Peer Fabric Interconnect to come up after reboot', 'Deploy Fiber Channel and Ethernet Breakout Ports', and 'Deploy System QoS Policy'.

## Create FC logical interfaces in storage

A Fibre Channel (FC) logical interface (LIF) is assigned a World Wide Port Name (WWPN) and linked to a physical FC port on the storage controller for Fibre Channel communications. To create four FC LIFs (two on each node, one for fabric A and the other for fabric B) and four NVMe/FC LIFs, follow the command examples below.

```
network interface create -vserver <svm-name> -lif fc-lif-01a -data-protocol fcp -home-node
<node01> -home-port 1a -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-01b -data-protocol fcp -home-node
<node01> -home-port 1c -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-02a -data-protocol fcp -home-node
```

```

<node02> -home-port 1a -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-02b -data-protocol fcp -home-node
<node02> -home-port 1c -status-admin up

network interface create -vserver <svm-name> -lif fc-lif-01a -data-protocol fc-nvme -home-node
<node01> -home-port 1b -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-01b -data-protocol fc-nvme -home-node
<node01> -home-port 1d -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-02a -data-protocol fc-nvme -home-node
<node02> -home-port 1b -status-admin up
network interface create -vserver <svm-name> -lif fc-lif-02b -data-protocol fc-nvme -home-node
<node02> -home-port 1d -status-admin up

```

Example:

```

fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-lif-01a -data-protocol fcp -
home-node fpsa-a50-u0909-01 -home-port 1a -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-lif-01b -data-protocol fcp -
home-node fpsa-a50-u0909-01 -home-port 1c -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-lif-02a -data-protocol fcp -
home-node fpsa-a50-u0909-02 -home-port 1a -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-lif-02b -data-protocol fcp -
home-node fpsa-a50-u0909-02 -home-port 1c -status-admin up

```

```

fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-nvme-lif-01a -data-protocol fc-
nvme -home-node fpsa-a50-u0909-01 -home-port 1b -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-nvme-lif-01b -data-protocol fc-
nvme -home-node fpsa-a50-u0909-01 -home-port 1d -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-nvme-lif-02a -data-protocol fc-
nvme -home-node fpsa-a50-u0909-02 -home-port 1b -status-admin up
fpsa-a50-u0909::> network interface create -vserver svm1 -lif fc-nvme-lif-02b -data-protocol fc-
nvme -home-node fpsa-a50-u0909-02 -home-port 1d -status-admin up

```

To verify:

```

fpsa-a50-u0909::> net int show -data-protocol fcp
(network interface show)

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svm1	fc-lif-01a	up/up	20:05:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1a	true
	fc-lif-01b	up/up	20:02:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1c	true
	fc-lif-02a	up/up	20:06:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1a	true
	fc-lif-02b	up/up	20:08:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1c	true

4 entries were displayed.

```

fpsa-a50-u0909::> net int show -data-protocol fc-nvme
(network interface show)

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svm1	fc-nvme-lif-01a	up/up	20:0b:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1b	true
	fc-nvme-lif-01b	up/up	20:0c:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1d	true
	fc-nvme-lif-02a	up/up	20:0d:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1b	true
	fc-nvme-lif-02b	up/up	20:1c:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1d	true

4 entries were displayed.

## FC LIF failover

The LIFs for FC protocol communications will failover automatically to its partner by default to help improve data availability during failover situations. As a result, after a storage failover, the FC LIFs on the node which went down will be migrated automatically to the surviving node as shown in the example below.

Check LIF status and location:

```

fpsa-a50-u0909::> net int show -lif fc-lif*

```

```

(network interface show)
Vserver      Logical   Status   Network      Current      Current Is
Interface    Admin/Oper Address/Mask Node          Port        Home
-----
svml
    fc-lif-01a  up/up    20:05:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1a  true
    fc-lif-01b  up/up    20:02:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1c  true
    fc-lif-02a  up/up    20:06:d0:39:ea:c6:a7:94  fpsa-a50-u0909-02  1a  true
    fc-lif-02b  up/up    20:08:d0:39:ea:c6:a7:94  fpsa-a50-u0909-02  1c  true
4 entries were displayed.

Perform a storage failover to confirm FC LIF failover:

fpsa-a50-u0909::> storage failover takeover -ofnode fpsa-a50-u0909-02

Warning: A takeover will be initiated. When the partner node reboots, a giveback will be
automatically initiated. Do you want to continue? {y|n}: y

Check LIF status and location again after storage failover:

fpsa-a50-u0909::> net int show -lif fc-lif-*
(network interface show)
Vserver      Logical   Status   Network      Current      Current Is
Interface    Admin/Oper Address/Mask Node          Port        Home
-----
svml
    fc-lif-01a  up/up    20:05:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1a  true
    fc-lif-01b  up/up    20:02:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1c  true
    fc-lif-02a  up/up    20:06:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1a  false
    fc-lif-02b  up/up    20:08:d0:39:ea:c6:a7:94  fpsa-a50-u0909-01  1c  false
4 entries were displayed.

```

**Note:** After a storage failover operation, the two FC LIFs which were hosted on the second node automatically failed over to the corresponding ports on the surviving first node. As a result, all FC LIFs remain available, and the host would continue to see all paths.

**Note:** With ONTAP version 9.16.1, NVMe/FC LIF automatic failover is not supported. Please refer to ONTAP documentation for LIF failover support details when you are deploying a solution with a newer ONTAP release.

## Create a UCS Server Profile Boot Order Policy with FC SAN boot support

The VMware ESXi and Linux servers can be configured to boot from FC SAN instead of booting from iSCSI SAN. Booting from SAN offers many benefits, including easier server replacement, disaster recovery, reduced power and cooling requirement without local disks, better performance, etc. Implementing SAN boot helps to realize the full benefits of Cisco UCS stateless computing features, including server profile mobility benefit for server upgrade or replacement.

To enable FC SAN boot in Boot Order Policy, you need to enter the WWPN of FC LIFs created on the NetApp storage controllers, the LUN ID to be used for SAN boot, and the appropriate device name to access the target. Follow the example steps below to clone the iSCSI SAN boot policy and modify the cloned copy to create a UCS Server Profile boot order policy with FC SAN boot support.

1. Login to Cisco Intersight.
2. Navigate to Configure > Policies.
3. Select the FlexPod-ASA-iSCSI-Boot-Order policy, click on the three dots on the right, and select Clone to clone the policy.
4. Update the Policy Name of the boot order policy to FlexPod-ASA-iSCSI-FC-Boot-Order since we are adding FC SAN boot. Click Clone.
5. After cloning is completed, select the newly create FlexPod-ASA-iSCSI-FC-Boot-Order, click on the three dots on the right-hand side, and click Edit to edit the policy.

- Click Add Boot Device drop-down list and select SAN Boot to add a new SAN boot device. We are adding four total FC SAN boot devices which correspond to the four paths to the FC SAN boot LUN.
- Enter the appropriate device name, LUN ID, Interface Name, and Target WWPN based on the information from your environment. See the following table and screenshots as example.

Device Name	LUN ID	Interface Name	Target WWPN
FC-LIF-01A	0	vHBA0	20:05:d0:39:ea:c6:a7:94
FC-LIF-01B	0	vHBA1	20:02:d0:39:ea:c6:a7:94
FC-LIF-02A	0	vHBA0	20:06:d0:39:ea:c6:a7:94
FC-LIF-02B	0	vHBA1	20:08:d0:39:ea:c6:a7:94

**Note:** Replace the Target WWPN in the table above with the ones from your environment when creating the boot order policy.

General

2 Policy Details

SAN Boot (FC-LIF-01A) Enabled

Device Name \* ?

FC-LIF-01A

LUN ?

0

0 - 255

Slot ?

Slot

Interface Name \* ?

vHBA0

Target WWPN \* ?

20:05:d0:39:ea:c6:a7:94

Bootloader Name ?

Bootloader Name

Bootloader Description ?

Bootloader Description

Bootloader Path ?

Bootloader Path

General

2 Policy Details

SAN Boot (FC-LIF-01B) Enabled

Device Name \* ?

FC-LIF-01B

LUN ?

0

0 - 255

Slot ?

Slot

Interface Name \* ?

vHBA1

Target WWPN \* ?

20:02:d0:39:ea:c6:a7:94

Bootloader Name ?

Bootloader Name

Bootloader Description ?

Bootloader Description

Bootloader Path ?

Bootloader Path

General

2 Policy Details

SAN Boot (FC-LIF-02A) Enabled

Device Name \* ?

FC-LIF-02A

LUN ?

0

0 - 255

Slot ?

Slot

Interface Name \* ?

vHBA0

Target WWPN \* ?

20:06:d0:39:ea:c6:a7:94

Bootloader Name ?

Bootloader Name

Bootloader Description ?

Bootloader Description

Bootloader Path ?

Bootloader Path

SAN Boot (FC-LIF-02B) Enabled

Device Name \* ?

FC-LIF-02B

LUN ?

0

0 - 255

Slot ?

Slot

Interface Name \* ?

vHBA1

Target WWPN \* ?

20:08:d0:39:ea:c6:a7:94

Bootloader Name ?

Bootloader Name

Bootloader Description ?

Bootloader Description

Bootloader Path ?

Bootloader Path

- After the four SAN boot device paths are added, click on the Up or Down arrow at the right-hand side of the SAN boot devices to rearrange the SAN boot device order. See the screenshot below for an example where the FC SAN boot devices are ordered below the iSCSI boot devices.

General

2 Policy Details

Add Boot Device ?

+ Virtual Media (KVM-Mapped-DVD) Enabled

+ iSCSI Boot (iSCSI-A-Boot) Enabled

+ iSCSI Boot (iSCSI-B-Boot) Enabled

+ SAN Boot (FC-LIF-01A) Enabled

+ SAN Boot (FC-LIF-01B) Enabled

+ SAN Boot (FC-LIF-02A) Enabled

+ SAN Boot (FC-LIF-02B) Enabled

+ Virtual Media (CIMC-Mapped-DVD) Enabled

- Click Save and click Save again to save the updated policy.

303

FlexPod SAN Solution with Cisco UCS X-Series  
Direct and NetApp ASA

© 2025 NetApp, Inc. All rights reserved. NetApp Verified Architecture

**Note:** We will be deploying a server profile that supports both iSCSI and FC SAN boot after creating a new template with the updated boot order and SAN connectivity policies.

**Note:** A SAN boot LUN should be configured for either iSCSI SAN boot or FC SAN boot, but not both. For a particular server, decide whether the server will use iSCSI boot or FC boot and configure the SAN boot LUN on the storage for access through iSCSI or FC protocol accordingly.

## Create a Server Profile template with the updated boot order policy and SAN connectivity

To create a server profile template with the updated boot order policy to boot from FC SAN, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Templates, click on UCS Server Profile Templates tab.
3. Select FlexPod-ASA-AMD-iSCSI-Boot template from the list, click on the three dots on the right, and click Clone to clone the template.
4. Click Next in the General section.
5. In the Details section, update Clone Name to FlexPod-ASA-AMD-iSCSI-FC-Boot and click Clone.
6. After cloning is completed, select the cloned FlexPod-ASA-AMD-iSCSI-FC-Boot template, click the three dots on the right, and click Edit.
7. Click Next for the Compute Configuration section.
8. Highlight the Boot Order row, click x to detach the existing iSCSI boot order. Then, click Select Policy and select FlexPod-ASA-iSCSI-FC-Boot-Order. Click Select.

← UCS Server Profile Templates

### Edit UCS Server Profile Template (FlexPod-ASA-AMD-iSCSI-FC-Boot)

- General
- 2 Compute Configuration**
- 3 Management Configuration
- 4 Storage Configuration
- 5 Network Configuration
- 6 Summary

#### Compute Configuration

Create or select existing Compute policies that you want to associate with this template.

##### UUID Assignment

UUID Pool ⓘ

Selected Pool FlexPod-ASA-UUID-Pool ⓘ | [Edit Selection](#) | ⓘ

BIOS	● FlexPod-ASA-AMD-M8-Virt-BIOS ⓘ
Boot Order	● FlexPod-ASA-iSCSI-FC-Boot-Order ⓘ
Firmware	ⓘ

9. Click Next three times for the Network Configuration section.
10. Click on the SAN Connectivity row, click Select Policy, select FlexPod-ASA-SAN-Connectivity on the list, and click Select.

General

Compute Configuration

Management Configuration

Storage Configuration

**5 Network Configuration**

6 Summary

### Network Configuration

Create or select existing Network Configuration policies that you want to associate with this template.

LAN Connectivity	● FlexPod-ASA-iSCSI-Boot-LAN-Connectivity ⓘ
SAN Connectivity	● FlexPod-ASA-SAN-Connectivity ⓘ

11. Click Next and review the Summary information.



- Click Derive Profiles if you are ready to derive server profiles from the template. Otherwise, click Close.

## Derive a UCS server profile with both iSCSI and FC boot order policy

- Login to Cisco Intersight.
- Navigate to Configure > Templates, and click on UCS Server Profile Templates tab.
- Select FlexPod-ASA-AMD-iSCSI-FC-Boot template from the list, click on the three dots on the right, and click Derive Profiles.
- In the General section, you can select server(s) for the derived server profile(s) or select Assign Later to only derive server profile(s) without assigning server(s). Click Next.

**Note:** We are deriving only one server profile in the above example and will configure it for FC SAN boot and add the host to the existing VMware cluster.

- In the Details section, update the profile name and organization information as needed.

- Click Next.
- Review the information in the Summary section. Click Derive.
- Navigate to Configure > Profiles.

9. The newly derived server profile has a status of Not Deployed. Select the profile, click on the three dots at the right-hand side, click Deploy.
10. In the Deploy UCS Server Profiles dialog, select Reboot immediately to activate, and click Deploy.

### Deploy UCS Server Profile

UCS server profile "fpsa-asa-esxi-03" will be deployed to server "fpsa-x9508-u0901-fi-1-4".

**⚠️** If policy configuration requires an immediate reboot and the option below is disabled, then profile deployment will not be initiated.

▼ More Details

☒ Reboot immediately to activate.

Cancel
Deploy

**Note:** It can take several minutes for the server profile to be deployed. You can click on the spinning request icon next to the check mark to review the request progress in detail.

The screenshot shows the Cisco Intersight interface. The top navigation bar includes the Cisco logo, 'Intersight' text, a search bar, and various icons for navigation and notifications. The main content area is titled 'Deploy Server Profile' and is divided into two panels.

**Details Panel:**

- Status:** In Progress (indicated by a blue bar)
- Name:** Deploy Server Profile
- ID:** 686c1305696f6e31010faf6c
- Target Type:** Blade Server
- Target Name:** fpsa-x9508-u0901-fi-1-4
- Source Type:** Server Profile

**Execution Flow Panel:**

A progress bar at the top of this panel shows 10% completion. Below it, a list of steps in the deployment process is shown, each with a green checkmark icon and a timestamp of 'Jul 7, 2025 2:33 PM':

- Wait For BIOS POST Completion
- Power on server
- Validate user access to the network policies
- Validate user access to the compute and management policies
- Validate user access to the profile
- Prepare Server Profile Deploy

## Create FC Zone policy to enable server FC protocol access to storage

FC zoning provides access for initiators and targets based on their World Wide Port Names (WWPNs). For a server configured with FC SAN boot information, appropriate single initiator single target FC zones are automatically created in the respective Fabric Interconnects when the FC Switching Mode is configured as Switch mode. This allows for FC communication with the specified target WWPNs listed in the FC SAN boot configuration.

However, for an iSCSI SAN booted host without FC SAN boot information configured, FC zone policy information will need to be added to the associated vHBAs for the host to communicate with the FC target LIFs.

**Note:** For NVMe/FC communication, appropriate FC Zone information with NVMe/FC LIFs from storage will be required for the NVMe/FC vHBAs. You can follow similar steps as below to configure FC Zones on NVMe/FC vHBAs for NVMe/FC communications.

Follow the steps below to obtain the WWPNs of the ONTAP FC target LIFs and create two FC Zone policies for accessing the targets from the two FC fabrics that can be assigned to the FC vHBAs based on their assigned fabric.

1. Login to ONTAP CLI.
2. Use the network interface show command to list the FC target LIF information as show in the example below.

```
fpsa-a50-u0909::> net int show -lif fc-lif-*
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
svml
    fc-lif-01a  up/up       20:05:d0:39:ea:c6:a7:94 fpsa-a50-u0909-01 1a true
    fc-lif-01b  up/up       20:02:d0:39:ea:c6:a7:94 fpsa-a50-u0909-01 1c true
    fc-lif-02a  up/up       20:06:d0:39:ea:c6:a7:94 fpsa-a50-u0909-02 1a true
    fc-lif-02b  up/up       20:08:d0:39:ea:c6:a7:94 fpsa-a50-u0909-02 1c true
4 entries were displayed.
```

3. Login to Cisco Intersight.
4. Navigate to Configure > Policies, click on Create Policy.
5. Select UCS Server Platform Type, select FC Zone, and click Start.
6. In the FC Zone Create General section, select the Organization from drop-down list, provide policy Name for fabric A, optionally set Tags or provide a Description. Click Next.

1 General

2 Policy Details

### General

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-FC-Zone-A

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

7. In the Policy Details section, select Single Initiator Multiple Target for the FC Target Zoning Type.
8. Click Add Target.
9. Provide the information from the first FC LIF in fabric A. Use FC LIF name from storage for Name, provide FC LIF WWPN, select the switch ID for the fabric, provide the VSAN ID, and click Add.

✓ General

2 Policy Details

### Add Target

Name \*

fc-lif-01a

WWPN \* ⓘ

20:05:d0:39:ea:c6:a7:94

Switch ID ⓘ

☒ A
☐ B

VSAN ID \* ⓘ

121

1 - 4093

Cancel Add

Initiator Multiple Target ☐ None

VSAN ID Switch ID

ITEMS AVAILABLE

10. Repeat steps 8 and 9 above to add the second FC LIF in the same fabric A as a target.

✓ General

2 Policy Details

### Add Target

Name \*

fc-lif-02a

WWPN \* ⓘ

20:06:d0:39:ea:c6:a7:94

Switch ID ⓘ

☒ A
☐ B

VSAN ID \* ⓘ

121

1 - 4093

Cancel Add

Initiator Multiple Target ☐ None

VSAN ID Switch ID

121 A

ITEMS AVAILABLE

11. Review the target information in the list.

✓ General

2 Policy Details

### Policy Details

Add policy details.

FC Target Zoning Type ⓘ

☐ Single Initiator Single Target
☒ Single Initiator Multiple Target
☐ None

Add Target

✎

🗑

⋮

<input type="checkbox"/>	Name	WWPN	VSAN ID	Switch ID
<input type="checkbox"/>	fc-lif-01a	20:05:d0:39:ea:c6:a7:94	121	A
<input type="checkbox"/>	fc-lif-02a	20:06:d0:39:ea:c6:a7:94	121	A

12. Click Create to create the FC Zone policy for fabric A.

13. Repeat steps 4 to 12 to create another FC Zone policy for fabric B and add the FC target LIFs from fabric B to the list of targets.

- 1 General
- 2 Policy Details

## General

Add a name, description, and tag for the policy.

Organization \*

FlexPod-ASA

Name \*

FlexPod-ASA-FC-Zone-B

Set Tags

Enter a tag in the key:value format.

Description

Description

0 / 1024

- ✓ General
- 2 Policy Details

## Add Target

Name \*

fc-lif-01b

WWPN \* ⓘ

20:02:d0:39:ea:c6:a7:94

Switch ID ⓘ

☐ A ☒ B

VSAN ID \* ⓘ

122

1 - 4093

Initiator Multiple Target ☐ None

VSAN ID Switch ID

ITEMS AVAILABLE

- ✓ General
- 2 Policy Details

## Add Target

Name \*

fc-lif-02b

WWPN \* ⓘ

20:08:d0:39:ea:c6:a7:94

Switch ID ⓘ

☐ A ☒ B

VSAN ID \* ⓘ

122

1 - 4093

Initiator Multiple Target ☐ None

VSAN ID Switch ID

122 B

Rows per page 10

< 1 >

- ✓ General
- 2 Policy Details

## Policy Details

Add policy details.

FC Target Zoning Type ⓘ

☐ Single Initiator Single Target ☒ Single Initiator Multiple Target ☐ None

Add Target

<input type="checkbox"/>	Name	WWPN	VSAN ID	Switch ID	
<input type="checkbox"/>	fc-lif-01b	20:02:d0:39:ea:c6:a7:94	122	B	
<input type="checkbox"/>	fc-lif-02b	20:08:d0:39:ea:c6:a7:94	122	B	

## Update SAN connectivity policy with the FC zone policies

To update the SAN connectivity policy with the FC zoning information, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Policies.
3. Select the SAN Connectivity policy to be updated, click on the three dots on the right, and select Edit.
4. Click Next to go to the Policy Details section.
5. Select vHBA0 that is connected to switch A, click on the three dots on the right, and select Edit.
6. On the Edit vHBA section, scroll down to the bottom for FC Zone selection, click on Select Policies.
7. Select the FC Zone policy created for fabric A which includes target LIFs in fabric A. Click Select.

Policies > SAN Connectivity > FlexPod-ASA-FC-Net

### Edit

A

PCI Order ⓘ  
6

**Persistent LUN Bindings**  
☐ Persistent LUN Bindings ⓘ

**Fibre Channel Network Policy \*** ⓘ  
Selected Policy FlexPod-ASA-FC-Net

**Fibre Channel QoS Policy \*** ⓘ  
Selected Policy FlexPod-ASA-FC-QoS

**Fibre Channel Adapter Policy \*** ⓘ  
Selected Policy FlexPod-ASA-FC-Ada

**FC Zone** ⓘ  
[Select Policies](#)

Create Policy

Search

Filters 2 results

<input type="checkbox"/>	Name	FC Target Zoning Type	Last Update	
<input type="checkbox"/>	FlexPod-ASA-FC-Zone-B	Single Initiator Multiple Target	19 minutes ago	
<input checked="" type="checkbox"/>	FlexPod-ASA-FC-Zone-A	Single Initiator Multiple Target	23 minutes ago	

Selected 1 of 2

Show Selected Unselect All

Rows per page 10

< 1 >

8. Click Update.
9. Select vHBA1 that is connected to switch B, click on the three dots on the right, and select Edit.
10. On the Edit vHBA section, scroll down to the bottom for FC Zone selection, click on Select Policies.



11. Select the FC Zone policy created for fabric B which includes target LIFs in fabric B. Click Select.

The screenshot shows the 'Select Policies' window. On the left, a sidebar lists various policy categories: Persistent LUN Bindings, Fibre Channel Network Policy, Fibre Channel QoS Policy, Fibre Channel Adapter Policy, and FC Zone. The 'FC Zone' section is expanded, showing a 'Select Policies' button. The main area displays a table with two policies: 'FlexPod-ASA-FC-Zone-B' (selected) and 'FlexPod-ASA-FC-Zone-A'. Both policies have a 'Single Initiator Multiple Target' zoning type. The 'Last Update' column shows '22 minutes ago' for the selected policy and '26 minutes ago' for the other. At the bottom, there are controls for 'Selected 1 of 2', 'Show Selected', 'Unselect All', and a pagination bar showing '1' of 1 page.

12. Click Update.

13. Click Save & Deploy.

14. Click Save & Proceed to redeploy the associated server profile.

15. Check all the boxes on the Deploy Server Profiles screen and click Deploy to deploy and activate the updated server profile.

**Note:** You can monitor the Requests page to see the deployment progress.

**Note:** If you are using a secure multi-tenant solution design with multiple storage virtual machines (SVMs), you can create additional server profile templates and FC zone policies using the appropriate FC LIFs from the respective SVMs to derive server profiles for the servers to communicate with a specific set of FC LIFs in a particular SVM.

## Appendix J: Installation and configuration of FC SAN-booted ESXi host

The following subsections provide information for the installation and configuration of an FC SAN-booted ESXi host. Wherever the configuration procedures are the same as those for iSCSI SAN-booted ESXi host, they are noted as such. Please refer to the corresponding iSCSI configuration sections in the NVA for detailed information.

### Create FC igroups on storage

For FC SAN boot, multiple initiator groups are created. One boot igroup per FC SAN-booted host for SAN boot LUN mapping. Both FC initiator WWPNs from a server should be included in that server-specific igroup for SAN boot LUN access.

The other igroup is for mapping the shared FC LUNs for shared VMware cluster datastore access. For this shared group, all WWPNs from all FC-capable hosts in the VMware cluster should be added for shared FC datastore access.

The WWPNs of the initiators/vHBAs can be retrieved from the deployed server profile in InterSight. To create FC igroups an ESXi host with FC vHBAs, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Profiles.
3. Select the server profile for a particular server, go to General tab, click on vNICs / vHBAs tab under Configuration.
4. Scroll down to the bottom of the page to see a list of vHBAs and their associated WWPNs.

← UCS Server Profiles

## fpsa-asa-esxi-03

Actions ▾

General Server Inventory Connectivity

**Details**

Status  
OK

**General** Edit

Name  
fpsa-asa-esxi-03

Description  
-

User Label  
-

**Configuration**

Policies Identifiers **vNICs / vHBAs** Errors/Warnings (0)

^ vHBAs

Q Search Filters 4 results Export

Name	WWPN ID	WWPN Pools	PCI Ord...
vHBA0	20:00:00:25:B5:09:00:00	FlexPod-AS...	6
vHBA1	20:00:00:25:B5:09:00:01	FlexPod-AS...	7
vHBA2	20:00:00:25:B5:09:00:02	FlexPod-AS...	8
vHBA3	20:00:00:25:B5:09:00:03	FlexPod-AS...	9

**Note:** For this solution, we are using vHBA0 / vHBA1 for FC protocol and vHBA2 / vHBA3 for NVMe/FC protocol. The WWPNs needed for storage igroup creation for FC SAN are listed in the WWPN ID column for the vHBA0 / vHBA1 adapters.

5. Login to ONTAP.
6. Follow the commands in the example below to create initiator groups for FC LUN mapping and access.

```
lun igroup create -igroup <igroup-name> -protocol fcp -ostype vmware -initiator <WWPN1,WWPN2>
```

Example:

```
fpsa-a50-u0909:> lun igroup create -vserver svml -igroup FlexPod-ASA-esxi-03-boot-fc -protocol fcp -ostype vmware -initiator 20:00:00:25:B5:09:00:00,20:00:00:25:B5:09:00:01
```

```
fpsa-a50-u0909:> lun igroup create -vserver svml -igroup FlexPod-ASA-esxi-cluster-fc -protocol fcp -ostype vmware -initiator 20:00:00:25:B5:09:00:00,20:00:00:25:B5:09:00:01
```

To verify:

```
fpsa-a50-u0909:> igroup show -protocol fcp -vserver svml
```

Vserver	Igroup	Protocol	OS	Type	Initiators
svml	FlexPod-ASA-esxi-03-boot-fc	fcp	vmware		20:00:00:25:b5:09:00:00 20:00:00:25:b5:09:00:01
svml	FlexPod-ASA-esxi-cluster-fc	fcp	vmware		20:00:00:25:b5:09:00:00 20:00:00:25:b5:09:00:01

2 entries were displayed.

**Note:** For additional FC SAN-boot ESXi hosts with FC vHBAs, create additional boot igroups to map their respective boot LUNs. In addition, add the WWPNs from the additional hosts to the VMware cluster FC igroup to share LUNs mapped for FC datastores.

## Create and map LUNs for FC SAN boot and FC datastore

Login to ONTAP cluster using command line or from ONTAP System Manager to create LUNs for FC SAN boot and FC datastore. To use the command line interface, follow the steps below.

1. Login to ONTAP.
2. Create a 128G FC boot LUN and a 1 TB LUN for FC datastore. Update the LUN sizes based on your requirements accordingly.

```
lun create -vserver <svm-name> <lun-path> -size <lun-size> -ostype vmware  
or  
lun create -vserver <svm-name> -path <lun-path> -size <lun-size> -ostype vmware  
Example:  
lun create -vserver svml fpsa_asa_esxi_03_boot_fc_1 -size 128G -ostype vmware  
lun create -vserver svml -path fpsa_asa_vmware_cluster_datastore_fc_1 -size 1T -ostype vmware
```

3. Map the LUNs to the host SAN boot and shared VMware cluster igroups.

```
lun map -vserver <svm-name> -path <lun-path> -igroup <igroup-name> -lun-id <lun-id>  
Example:  
lun map -vserver svml -path fpsa_asa_esxi_03_boot_fc_1 -igroup FlexPod-ASA-esxi-03-boot-fc -lun-id 0  
lun map -vserver svml -path fpsa_asa_vmware_cluster_datastore_fc_1 -igroup FlexPod-ASA-esxi-cluster-fc -lun-id 1  
To verify:  
fpsa-a50-u0909::> lun show -m -igroup *fc  
Vserver      Path  
-----  
svml         fpsa_asa_esxi_03_boot_fc_1  
svml         fpsa_asa_vmware_cluster_datastore_fc_1  
2 entries were displayed.
```

Igroup	LUN ID	Protocol
FlexPod-ASA-esxi-03-boot-fc	0	fc
FlexPod-ASA-esxi-cluster-fc	1	fc

## FC SAN-booted ESXi host installation

The procedures for installing ESXi on the FC SAN boot LUN are like those for the iSCSI SAN booted ESXi host. The installation process is started by booting the host with the ESXi installation CD and then picking the LUN to be used for SAN boot as the installation destination.

The mapped FC SAN boot LUNs is discovered during the UCS server boot process as show in the following screenshot.



Copyright (c) 2024 Cisco Systems, Inc.

Press <F2> Setup : <F6> Boot Menu : <F12> Network Boot  
Bios Version : X215M8.4.3.5c.0.1202241033  
Platform ID : X215M8

Processor(s) AMD EPYC 9534 64-Core Processor

Total Memory = 128 GB Effective Memory = 128 GB  
Memory Operating Speed 4800 Mhz  
Cisco VIC Fibre Channel Driver Version 2.2(1i)  
Cisco VIC Simple Network Protocol Driver Version 2.2(1i)  
(C) 2013 Cisco Systems, Inc.

SAN	Storage	20:05:D0:39:EA:C6:A7:94	128.00 GB
SAN	Storage	20:06:D0:39:EA:C6:A7:94	128.00 GB
SAN	Storage	20:02:D0:39:EA:C6:A7:94	128.00 GB
SAN	Storage	20:08:D0:39:EA:C6:A7:94	128.00 GB

92

**Note:** As mentioned previously, FC SAN boot configuration leads to automatic zoning in the FIs when their FC Switching Mode is configured as Switch mode. To check for FC zoning configurations, you can login to the FIs, connect to nxos, and check the active zoneset information. There were two automatically created single-initiator-single-target zones on each FI for a FC SAN-booted host as shown in the example below. The two zones in FI-A and FI-B shown below provide the four total paths seen in the BIOS boot up screen for the SAN boot LUN with the four FC target LIFs listed.

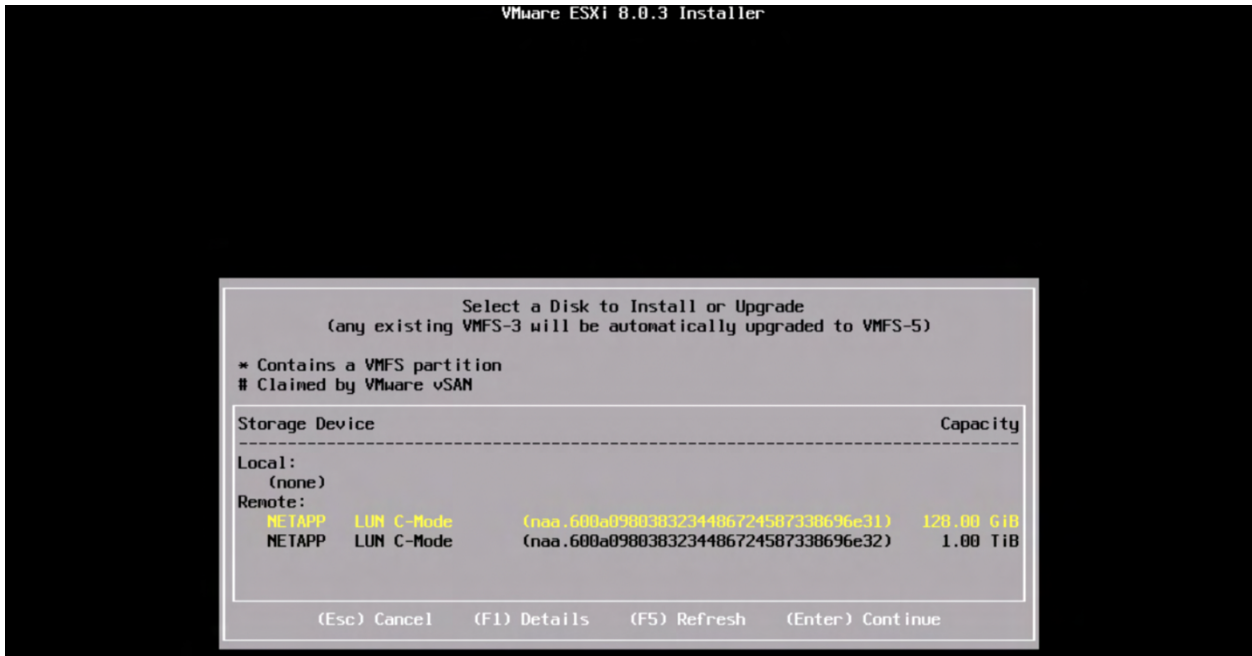
```
fpsa-x9508-u0901-fi-A# connect nxos
...
fpsa-x9508-u0901-fi-A(nx-os)# show zoneset active
zoneset name ucs-fpsa-x9508-u0901-fi-vsan-121-zoneset vsan 121
  zone name ucs_fpsa-x9508-u0901-fi_A_1_fpsaasaesxi03_vHBA0 vsan 121
    * fcid 0x230080 [pwwn 20:00:00:25:b5:09:00:00]
    * fcid 0x230001 [pwwn 20:06:d0:39:ea:c6:a7:94]

  zone name ucs_fpsa-x9508-u0901-fi_A_2_fpsaasaesxi03_vHBA0 vsan 121
    * fcid 0x230080 [pwwn 20:00:00:25:b5:09:00:00]
    * fcid 0x230021 [pwwn 20:05:d0:39:ea:c6:a7:94]
...
```

```
fpsa-x9508-u0901-fi-B# connect nxos
...
fpsa-x9508-u0901-fi-B(nx-os)# show zoneset active
zoneset name ucs-fpsa-x9508-u0901-fi-vsan-122-zoneset vsan 122
  zone name ucs_fpsa-x9508-u0901-fi_B_1_fpsaasaesxi03_vHBA1 vsan 122
    * fcid 0x1c0080 [pwwn 20:00:00:25:b5:09:00:01]
    * fcid 0x1c0001 [pwwn 20:02:d0:39:ea:c6:a7:94]

  zone name ucs_fpsa-x9508-u0901-fi_B_2_fpsaasaesxi03_vHBA1 vsan 122
    * fcid 0x1c0080 [pwwn 20:00:00:25:b5:09:00:01]
    * fcid 0x1c0041 [pwwn 20:08:d0:39:ea:c6:a7:94]
...
```

After the server is booted with the ESXi installation CD VMware-ESXi-8.0.U3-24022510-Custom-Cisco-4.3.5-a.iso and server firmware 5.3(0.250001), the ESXi installer shows the discovered LUNs.



Select the LUN to be used for ESXi FC SAN boot disk, the 128G LUN for this example, and hit Enter to start the installation process.

**Note:** Please refer to the subsections under VMware ESXi 8U3 installation in the NVA for additional details on installing VMware ESXi onto the boot LUN, setting up management network, resetting VMkernel port vmk0 MAC address, and updating Cisco VIC drivers.

**Note:** To check the versions of the nenic and nfnic drivers, used respectively for Ethernet and FC communications, you can login to the ESXi host as root and examine the versions by following the examples below.

```
[root@fpsa-asa-esxi-03:~] vmkload_mod -s nenic | grep Version
Version: 2.0.15.0-10EM.800.1.0.20613240

[root@fpsa-asa-esxi-03:~] vmkload_mod -s nfnic | grep Version
Version: 5.0.0.45-10EM.800.1.0.20613240
```

## vCenter configurations for the new FC host

The general vCenter configurations for the newly added FC SAN booted ESXi host is very similar to those for the iSCSI SAN booted ESXi hosts. Instead of duplicating the procedures, please go over the following subsections available under the VMware cluster configuration procedures in the NVA to configure and check the newly created FC SAN booted host in vCenter. We have also set up iSCSI SAN access for the FC SAN booted ESXi host, so it can utilize the shared iSCSI datastores.

- Add remaining ESXi host to the cluster
- Configure in-band management virtual switch
- Add ESXi host to vDS0 (assign Uplink1 to vmnic2 and Uplink2 to vmnic3)
- Configure VMkernel port for vMotion
- Add host to iSCSI vDS (assign Uplink1 to vmnic4 and Uplink2 to vmnic5)

- Configure iSCSI VMkernel port (skip steps 4 - 7)
- Configure iSCSI software adapter for additional hosts
- Mount shared iSCSI datastore on additional hosts
- Configure VMware cluster and ESXi host VM swap file location (steps 5 - 9)
- vCenter Trusted Platform Module (TPM) attestation

**Note:** The IQN of the host will need to be added to the ONTAP iSCSI igroup for the VMware cluster to access the shared iSCSI datastores.

**Note:** It is very important to save the ESXi host's encryption recovery information when secure boot is configured. This information is needed for recovery when performing a server replacement as documented in the NVA. You can obtain the encryption recovery information by executing the following command as the root user and save the output in a secure location for later use.

```
[root@fpsa-asa-esxi-03:~] esxcli system settings encryption recovery list
Recovery ID                               Key
-----
{xxxxx...}                               xxxxxxxx-...
```

## Create FC datastore

To add FC datastore to the newly added FC host, follow the steps below.

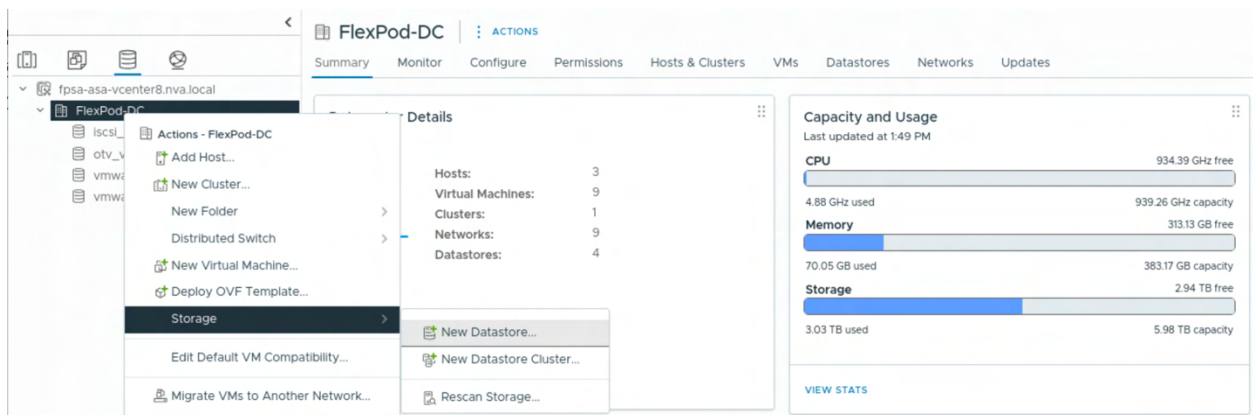
1. Login to vCenter.
2. Click to select the FC host from the vCenter Inventory view and click on the Configure tab.
3. Click Storage Devices to see the list of discovered mapped devices and select the 1 TB FC device which was created in ONTAP storage in an earlier section.
4. Click on the Paths tab in the lower pane to see the four different paths that can be used to access the device. There are two paths through each FC vHBA for the connections to the two storage controllers. All four paths should show Active (I/O) Status because the NetApp ASA controllers provide symmetric active/active access for SCSI LUNs using FC protocol.

The screenshot shows the vSphere Client interface for a host named 'fpsa-asa-esxi-03.nva.local'. The 'Storage' tab is selected, and the 'Storage Devices' section is expanded. A table lists the storage devices, including a 'NETAPP Fibre Channel Disk' with a capacity of 1.00 TB. The 'Paths' tab is selected, showing four paths with 'Active (I/O)' status.

Runtime Name	Status	Target	Transport	Name	Preferred
vmhba0:C0:T1L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94	Fibre Channel	vmhba0:C0:T1L1	No
vmhba1:C0:T2L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94	Fibre Channel	vmhba1:C0:T2L1	No
vmhba1:C0:T1L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94	Fibre Channel	vmhba1:C0:T1L1	No
vmhba0:C0:T2L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94	Fibre Channel	vmhba0:C0:T2L1	No

5. Select the Storage view in vCenter, right-click the FlexPod-DC, select Storage, and click New Datastore.





6. Select VMFS datastore type and click NEXT.
7. Provide a name for the datastore. Select a host from the drop-down list to view and select a device and then click NEXT.

### New Datastore

- Type
- Name and device selection**
- VMFS version
- Partition configuration
- Ready to complete

### Name and device selection

Specify datastore name and a disk/LUN for provisioning the datastore.

Name

The datastore will be accessible to all the hosts that are configured with access to the selected disk/LUN. If you do not find the disk/LUN that you are interested in, it might not be accessible to that host. Try changing the host or configure accessibility of that disk/LUN.

Select a host

Select a host to view its accessible disks/LUNs:

	Name	LUN	Capacity	Hardware Acceleration	Drive Type	Sector Format	Chk VM Su
<input type="radio"/>	NETAPP ISCSI Disk (naa.600a0980383234486724587338696d4a)	4	500.00 GB	Supported	Flash	512e	No
<input checked="" type="radio"/>	NETAPP Fibre Channel Disk (naa.600a0980383234486724587338696e32)	1	1.00 TB	Supported	Flash	512e	Yes

8. Keep the default VMFS 6 for VMFS version and click NEXT.
9. Review the disk partition configuration information and click NEXT.
10. Review the selections and click FINISH.

### New Datastore

- Type
- Name and device selection
- VMFS version
- Partition configuration
- Ready to complete**

### Ready to complete

Review your selections before finishing the wizard

**Name and device selection**

Datastore name: fc\_datastore\_1

Disk/LUN: NETAPP Fibre Channel Disk (naa.600a0980383234486724587338696e32)

**VMFS version**

Version: VMFS 6

**Partition configuration**

Datastore size: 1.00 TB

Partition format: GPT

Block size: 1 MB

Space reclamation granularity: 1 MB

Space reclamation priority: Low: Deleted or unmapped blocks are reclaimed on the LUN at low priority

11. Confirm the FC datastore creation in the storage view.

	Name	Status	Type	Datastore Cluster	Capacity	Free
<input type="checkbox"/>	fc_datastore_1	✓ Normal	VMFS 6		1,023.75 GB	1,022.32 GB
<input type="checkbox"/>	iscsi_datastore_1	✓ Normal	VMFS 6		5 TB	2.2 TB
<input type="checkbox"/>	otv_vmfs_iscsi_1	✓ Normal	VMFS 6		499.75 GB	298.32 GB
<input type="checkbox"/>	vmware_swap	✓ Normal	VMFS 6		399.75 GB	365.16 GB
<input type="checkbox"/>	vmware_vcls	✓ Normal	VMFS 6		99.75 GB	98.34 GB

## Appendix K: Configuration updates for FC SAN Booted ESXi host to access NVMe/FC storage

Follow the steps in the subsections below to provide NVMe/FC protocol access to storage for the FC SAN booted ESXi host.

### Gather host NVMe/FC vHBA WWPN information

To gather the World Wide Port Name (WWPN) information for NVMe/FC storage access configuration of the host, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Profiles.
3. Select the server profile for the server, go to General tab, and click Identifiers under Configuration.

Identity	Assigned ID
ISCSI Boot IP/04-iscsi-A	172.22.73.32
ISCSI Boot IP/05-iscsi-B	172.22.74.32
Out-of-Band Management IP	172.22.71.32
WWNN	20:00:00:25:B5:15:00:00
UUID	AAA00000-0000-0001-AAA0-000000000003
IQN	iqn.2010-11.com.flexpod:flexpod-asa-ucshost:3

4. The World Wide Node Name (WWNN) of the server is listed under the Assigned ID of the WWNN Identity.
5. Click on vNICs / vHBAs tab under Configuration.

6. Scroll down to the bottom of the page to see the associated WWPNs for vHBA2 and vHBA3 that will be used for NVMe/FC storage access.

UCS Server Profiles

## fpsa-asa-esxi-03

General Server Inventory Connectivity

Status: OK

General

Name: fpsa-asa-esxi-03

Description: -

User Label: -

Configuration

Policies Identifiers **vNICs / vHBAs** Errors/Warnings (0)

^ vHBAs

Search Filters 4 results Export

Name	WWPN ID	WWPN Pools	PCI Ord...
vHBA0	20:00:00:25:B5:09:00:00	FlexPod-AS...	6
vHBA1	20:00:00:25:B5:09:00:01	FlexPod-AS...	7
vHBA2	20:00:00:25:B5:09:00:02	FlexPod-AS...	8
vHBA3	20:00:00:25:B5:09:00:03	FlexPod-AS...	9

**Note:** For this solution, vHBA0 and vHBA1 are used for FC protocol access.

## Retrieve host NQN from vCenter

To obtain the NVMe/FC host's NVMe Qualified Name (NQN), follow the steps below.

1. Login to vCenter.
2. Select the FC host in the vCenter Inventory.
3. Click on Configure tab, select Storage Adapters in the center pane under Storage.
4. Select vHBA2 under the Storage Adapter view.
5. Click on the Controllers tab under the storage adapter list.
6. Click ADD CONTROLLER.

vSphere Client

fpsa-asa-esxi-03.nva.local

Summary Monitor Configure Permissions VMs Datastores Networks Updates

Storage

Storage Adapters

ADD SOFTWARE ADAPTER REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

Adapter Model Type Status Identifier

Host NQN nqn.2014-08.com.vmware:nvme:fpsa-asa-esxi-... COPY

World Wide Node Name Hexadecimal digits grouped as 2-8 pairs Central discovery controller

World Wide Port Name Hexadecimal digits grouped as 2-8 pairs

DISCOVER CONTROLLERS

Select which controller to connect

Id	Subsystem NQN	Transport Type	World Wide Node Name	World Wide Port Name
0 items				

7. The host NQN is listed in the Host NQN field.
8. Click Cancel to exit the view.

## Set Up NVMe/FC storage for host access

To setup storage for host access to NVMe namespaces via NVMe/FC protocol, follow the steps in the subsections below to create a NVMe subsystem, add host NQN to the subsystem, create NVMe namespace, and map the NVMe name space to the NVMe subsystem for host access.

1. Login to ONTAP.
2. Create an NVMe subsystem with VMware OS type.

```
nvme subsystem create -vserver <svm-name> -subsystem <subsystem-name> -ostype vmware
```

Example:

```
fpsa-a50-u0909::> nvme subsystem create -vserver svml -subsystem FlexPod-ASA-esxi-cluster-nvme-fc
-ostype vmware
(vserver nvme subsystem create)
```

To verify:

```
fpsa-a50-u0909::> nvme subsystem show -subsystem FlexPod-ASA-esxi-cluster-nvme-fc
(vserver nvme subsystem show)
```

Vserver	Subsystem	Target NQN
svml	FlexPod-ASA-esxi-cluster-nvme-fc	nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-ASA-esxi-cluster-nvme-fc

3. Add the host NQN retrieved from vCenter to the NVMe subsystem host.

```
nvme subsystem host add -vserver <svm-name> -subsystem <subsystem-name> -host-nqn <host-nqn>
```

Example:

```
fpsa-a50-u0909::> nvme subsystem host add -vserver svml -subsystem FlexPod-ASA-esxi-cluster-nvme-fc
-host-nqn nqn.2014-08.com.vmware:nvme:fpsa-asa-esxi-03
(vserver nvme subsystem host add)
```

To verify:

```
fpsa-a50-u0909::> nvme subsystem host show -subsystem FlexPod-ASA-esxi-cluster-nvme-fc
(vserver nvme subsystem host show)
```

Vserver	Subsystem	Priority	Host NQN
svml	FlexPod-ASA-esxi-cluster-nvme-fc	regular	nqn.2014-08.com.vmware:nvme:fpsa-asa-esxi-03

4. Create a NVMe namespace with the desired capacity and VMware OS type.

```
nvme namespace create -vserver <svm-name> -path <namespace-name> -size <size> -ostype vmware
```

Example:

```
fpsa-a50-u0909::> nvme namespace create -vserver svml -path
fpsa_asa_vmware_cluster_datastore_nvme_fc_1 -size 1000g -ostype vmware
(vserver nvme namespace create)
[Job 14573] Job succeeded.
```

To verify:

```
fpsa-a50-u0909::> nvme namespace show -path fpsa_asa_vmware_cluster_datastore_nvme_fc*
(vserver nvme namespace show)
```

Vserver	Path	State	Size	Subsystem	NSID
svml	fpsa_asa_vmware_cluster_datastore_nvme_fc_1	online	1000GB	-	-

## 5. Add the created NVMe namespace to the NVMe subsystem map.

```
nvme subsystem map add -vserver <svm-name> -subsystem <subsystem-name> -path <namespace-name>
```

Example:

```
fpsa-a50-u0909::> nvme subsystem map add -vserver svml -subsystem FlexPod-ASA-esxi-cluster-nvme-  
fc -path fpsa_asa_vmware_cluster_datastore_nvme_fc_1  
(vserver nvme subsystem map add)
```

To verify:

```
fpsa-a50-u0909::> nvme subsystem map show -subsystem FlexPod-ASA-esxi-cluster-nvme-fc  
(vserver nvme subsystem map show)
```

Vserver	Subsystem	NSID	Namespace	Path
svml	FlexPod-ASA-esxi-cluster-nvme-fc	00000001h	fpsa_asa_vmware_cluster_datastore_nvme_fc_1	

## Create UCS FC Zone policies for NVMe/FC vHBAs

FC zoning provides initiator access to targets based on World Wide Port Names (WWPNs). To enable NVMe/FC protocol access to storage for a server, FC zoning for the NVMe/FC vHBAs and NVMe/FC target LIFs are needed for the two FC fabrics.

Follow the steps below to obtain the WWPNs of the ONTAP target NVMe/FC LIFs and create two FC Zone policies for the server NVMe/FC vHBAs to access those respective NVMe/FC target LIFs.

1. Login to ONTAP CLI.
2. Use the network interface show command to list the NVMe/FC LIF information as shown in the example below.

```
fpsa-a50-u0909::> net int show -vserver svml -lif fc-nvme*  
(network interface show)
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	fc-nvme-lif-01a	up/up	20:0b:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1b	true
	fc-nvme-lif-01b	up/up	20:0c:d0:39:ea:c6:a7:94	fpsa-a50-u0909-01	1d	true
	fc-nvme-lif-02a	up/up	20:0d:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1b	true
	fc-nvme-lif-02b	up/up	20:1c:d0:39:ea:c6:a7:94	fpsa-a50-u0909-02	1d	true

4 entries were displayed.

3. Login to Cisco Intersight.
4. Navigate to Configure > Policies, click on Create Policy.
5. Select UCS Server Platform Type, select FC Zone, and click Start.
6. In the FC Zone Create General section, select the Organization from drop-down list, provide policy Name for fabric A, optionally set Tags or provide a Description. Click Next.

**1 General**

**2 Policy Details**

## General

Add a name, description, and tag for the policy.

**Organization \***

FlexPod-ASA

**Name \***

FlexPod-ASA-FC-NVMe-Zone-A

**Set Tags**

Enter a tag in the key:value format.

**Description**

Description

0 / 1024

7. In the Policy Details section, select Single Initiator Multiple Target for the FC Target Zoning Type.
8. Click Add Target.
9. Provide the information from the first NVMe/FC target LIF in fabric A. Use NVMe/FC LIF name from storage for Name, provide NVMe/FC LIF WWPN, select the switch ID for the fabric, provide the VSAN ID, and click Add.

**Policy Details**

**Add Target**

**Name \***

fc-nvme-lif-01a

**WWPN \***

20:0b:d0:39:ea:c6:a7:94

**Switch ID**

☒ A ☐ B

**VSAN ID \***

121

1 - 4093

Cancel Add

10. Repeat steps 8 and 9 above to add the second NVMe/FC target LIF in the same fabric A.



11. Review the target information in the list.

	Name	WWPN	Switch ID	VSAN ID
<input type="checkbox"/>	fc-nvme-lif-01a	20:0b:d0:39:ea:c6:a7:94	A	121
<input type="checkbox"/>	fc-nvme-lif-02a	20:0d:d0:39:ea:c6:a7:94	A	121

12. Click Create to create the NVMe/FC Zone policy for fabric A.

13. Repeat steps 4 to 12 to create another NVMe/FC Zone policy for fabric B and add the NVMe/FC target LIFs from fabric B to the list of targets.

General

2 Policy Details

### Policy Details

Add Target

Name \*

fc-nvme-lif-01b

WWPN \* ⓘ

20:0c:d0:39:ea:c6:a7:94

Switch ID ⓘ

☐ A
☒ B

VSAN ID \* ⓘ

122

1 - 4093

Cancel

Add

General

2 Policy Details

### Policy Details

Add Target

Name \*

fc-nvme-lif-02b

WWPN \* ⓘ

20:1c:d0:39:ea:c6:a7:94

Switch ID ⓘ

☐ A
☒ B

VSAN ID \* ⓘ

122

1 - 4093

Cancel

Add

General

2 Policy Details

### Policy Details

Add policy details.

FC Target Zoning Type ⓘ

☐ Single Initiator Single Target
☒ Single Initiator Multiple Target
☐ None

Add Target

	Name	WWPN	Switch ID
<input type="checkbox"/>	fc-nvme-lif-01b	20:0c:d0:39:ea:c6:a7:94	B
<input type="checkbox"/>	fc-nvme-lif-02b	20:1c:d0:39:ea:c6:a7:94	B

## Assign NVMe/FC Zone policies to ESXi FC server profile

To assign NVMe/FC zone policies to the ESXi FC server profile template for NVMe/FC protocol access to storage and then deploy the updated server profile, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Policies, click on Create Policy.
3. Select UCS Server Platform Type, select FC Zone, find, select, and edit the FlexPod ASA SAN connectivity policy for FC.
4. Click Next in the General section.
5. In the Policy Details section, select vHBA2 for NVMe/FC access via fabric A, click on the three dots on the right and select Edit.
6. Scroll down to the bottom of the Edit screen, click Select Policies under FC Zone.
7. Select FlexPod-ASA-FC-NVMe-Zone-A policy, click Select.
8. Check to confirm FlexPod-ASA-FC-NVMe-Zone-A policy is listed under FC Zone. Click Update.

Policies > SAN Connectivity > FlexPod-ASA-SAN-Connectivity

## Edit

PCI Order ⓘ

8

### Persistent LUN Bindings

☒ Persistent LUN Bindings ⓘ

### Fibre Channel Network Policy \* ⓘ

Selected Policy FlexPod-ASA-FC-Network-A ⓘ | ✎ | [Edit Selection](#) | 🗑️

### Fibre Channel QoS Policy \* ⓘ

Selected Policy FlexPod-ASA-FC-QoS ⓘ | ✎ | [Edit Selection](#) | 🗑️

### Fibre Channel Adapter Policy \* ⓘ

Selected Policy FlexPod-ASA-FC-Adapter-VMware ⓘ | ✎ | [Edit Selection](#) | 🗑️

### FC Zone ⓘ

^ Selected Policy 1 Policy [Edit Selection](#) | 🗑️

FlexPod-ASA-FC-NVMe-Zone-A

ⓘ | ✎ | 🗑️

9. Repeat steps 5 to 8 for vHBA3 for NVMe/FC access via fabric B and select the FlexPod-ASA-FC-NVMe-Zone-B policy and update the policy.

### FC Zone ⓘ

^ Selected Policy 1 Policy [Edit Selection](#) | 🗑️

FlexPod-ASA-FC-NVMe-Zone-B


ⓘ | ✎ | 🗑️

10. Click Save & Deploy.
11. Confirm policy save by clicking Save & Proceed.
12. Check the appropriate boxes in Deploy Server Profiles screen and click Deploy.

## Deploy Server Profiles



This policy is associated with 1 UCS server profile. Policy changes will only take effect after associated profiles are deployed. Only policies edited here will be deployed, along with any dependent policies. To deploy all policies, select **Deploy all associated policies whether modified or not**.

 The server is powered on. Deploying and activating the profile may cause a reboot and disruption. If deploying policy configurations requires an immediate reboot, check **Reboot immediately to activate**.

More Details

- ☒ Reboot immediately to activate.
- ☒ Deploy all associated policies whether modified or not.
- ☒ I understand that potential disruption may occur during profile deployment. \*

Cancel Deploy

**Note:** To avoid disruptions to the host, put the host into maintenance mode and shut it down from vCenter first before deploying the updated server profile.

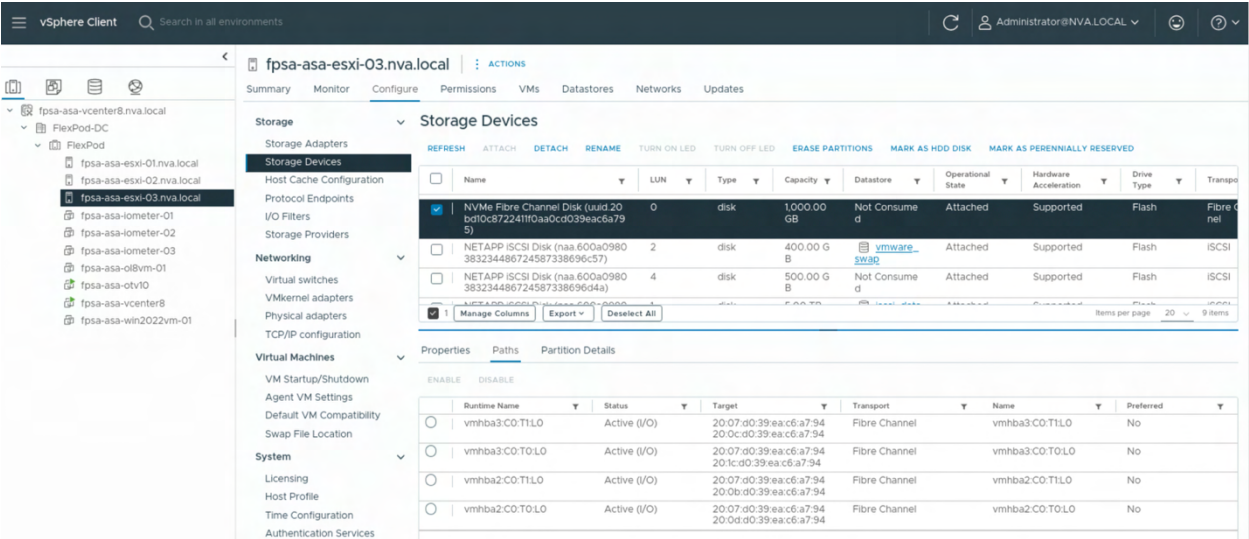
## Create NVMe/FC datastore

To create NVMe/FC datastore on the FC host, follow the steps in the subsections below.

### Check for the mapped NVMe namespace

To check for the NVMe namespaces available to the host, perform the steps below.

1. Login to vCenter.
2. Click to select one of the hosts from the vCenter Inventory view. Click on the Configure tab.
3. Click Storage Devices to see the list of discovered devices. Click to select one of the devices with name starting with NVMe Fibre Channel Disk.



Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type	Transport
NVMe Fibre Channel Disk (uuid:20-bd10c8722411f0aa0cd039eac6a795)	0	disk	1,000.00 GB	Not Consumed	Attached	Supported	Flash	Fibre Channel
NETAPP (iSCSI) Disk (naa:600a0980383234486724587338696d57)	2	disk	400.00 GB	vmware:storage	Attached	Supported	Flash	iSCSI
NETAPP (iSCSI) Disk (naa:600a0980383234486724587338696d4a)	4	disk	500.00 GB	Not Consumed	Attached	Supported	Flash	iSCSI

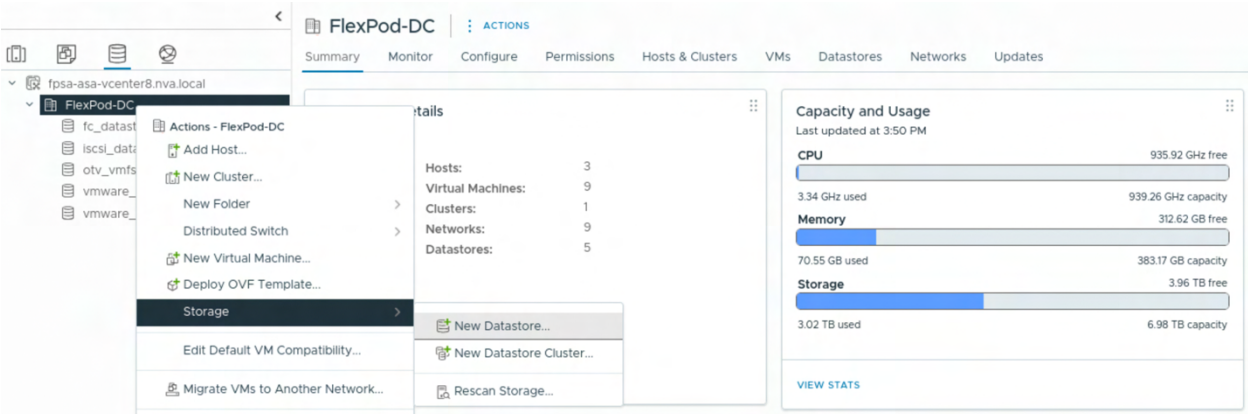
Runtime Name	Status	Target	Transport	Name	Preferred
vmhba3:C0:T1L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0c:d0:39:ea:c6:a7:94	Fibre Channel	vmhba3:C0:T1L0	No
vmhba3:C0:T0L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:1c:d0:39:ea:c6:a7:94	Fibre Channel	vmhba3:C0:T0L0	No
vmhba2:C0:T1L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0b:d0:39:ea:c6:a7:94	Fibre Channel	vmhba2:C0:T1L0	No
vmhba2:C0:T0L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0d:d0:39:ea:c6:a7:94	Fibre Channel	vmhba2:C0:T0L0	No

4. Click on the Paths tab in the lower pane to see the four different paths that can be used to access the device. There are two paths through each of the NVMe/FC adapter for the connections to the two storage controllers. All four paths should show Active (I/O) Status because the NetApp ASA controllers provide symmetric active/active access for NVMe namespaces.

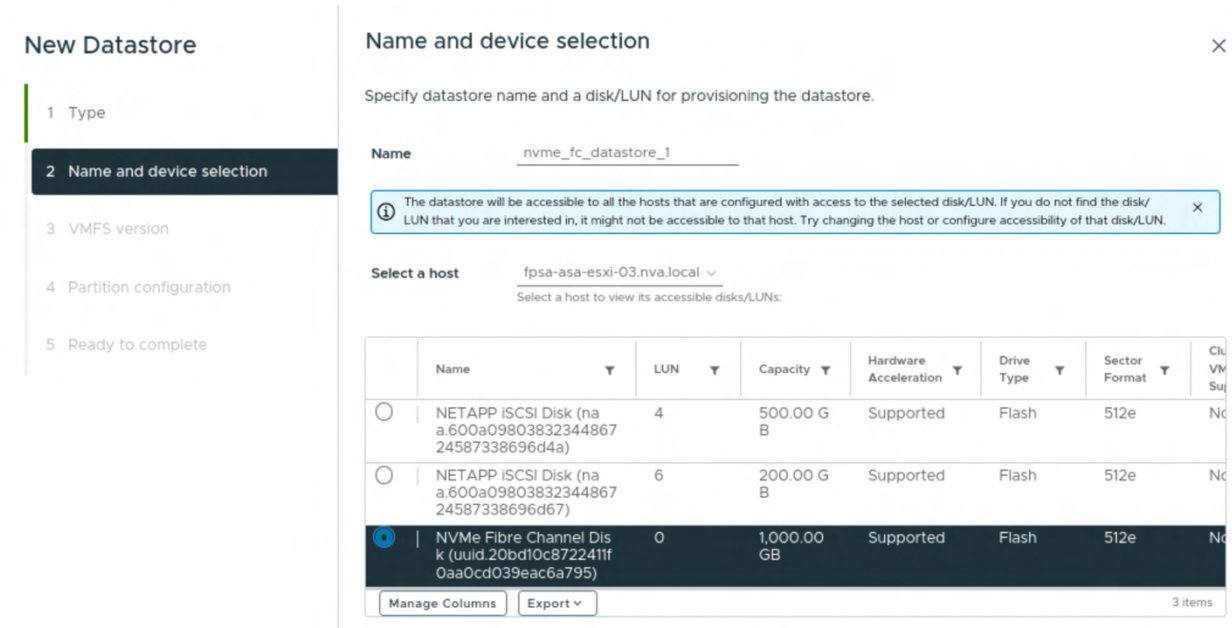
# Create a datastore from the mapped NVMe namespace

To creation a datastore from a discovered NVMe namespace, follow the steps below.

1. Select the Storage view in vCenter, right-click the FlexPod-DC, select Storage, and click New Datastore.



2. Select VMFS datastore type and click NEXT.
3. Provide a name for the datastore. Select the FC host from the drop-down list to view and select NVMe/FC device and then click NEXT.



4. Keep the default VMFS 6 for VMFS version and click NEXT.
5. Review the disk partition configuration information and click NEXT.
6. Review the selections and click FINISH.

### New Datastore

- Type
- Name and device selection
- VMFS version
- Partition configuration
- Ready to complete

### Ready to complete

Review your selections before finishing the wizard

Name and device selection

Datastore name  
nvme\_fc\_datastore\_1
Disk/LUN  
NVMe Fibre Channel Disk (uuid.20bd10c8722411f0aa0cd039eac6a795)

VMFS version

Version  
VMFS 6

Partition configuration

Datastore size  
1,000.00 GB
Partition format  
GPT
Block size  
1 MB
Space reclamation granularity  
1 MB
Space reclamation priority  
Low: Deleted or unmapped blocks are reclaimed on the LUN at low priority

- Confirm the NVMe/FC datastore creation in the storage view.

FlexPod-DC

ACTIONS

Summary

Monitor

Configure

Permissions

Hosts & Clusters

VMs

Datastores

Networks

Updates

Datastores

Datastore Clusters

Datastore Folders

Quick Filter

Enter value

<input type="checkbox"/>	Name	Status	Type	Datastore Cluster	Capacity	Free
<input type="checkbox"/>	fc_datastore_1	✓ Normal	VMFS 6		1,023.75 GB	1,022.32 GB
<input type="checkbox"/>	iscsi_datastore_1	✓ Normal	VMFS 6		5 TB	2.21 TB
<input type="checkbox"/>	nvme_fc_datastore_1	✓ Normal	VMFS 6		999.75 GB	998.32 GB
<input type="checkbox"/>	otv_vmfs_iscsi_1	✓ Normal	VMFS 6		499.75 GB	298.32 GB
<input type="checkbox"/>	vmware_swap	✓ Normal	VMFS 6		399.75 GB	365.16 GB
<input type="checkbox"/>	vmware_vcls	✓ Normal	VMFS 6		99.75 GB	98.34 GB

## Appendix L: Configuration updates for the iSCSI SAN-booted ESXi hosts to access FC and NVMe/FC storage

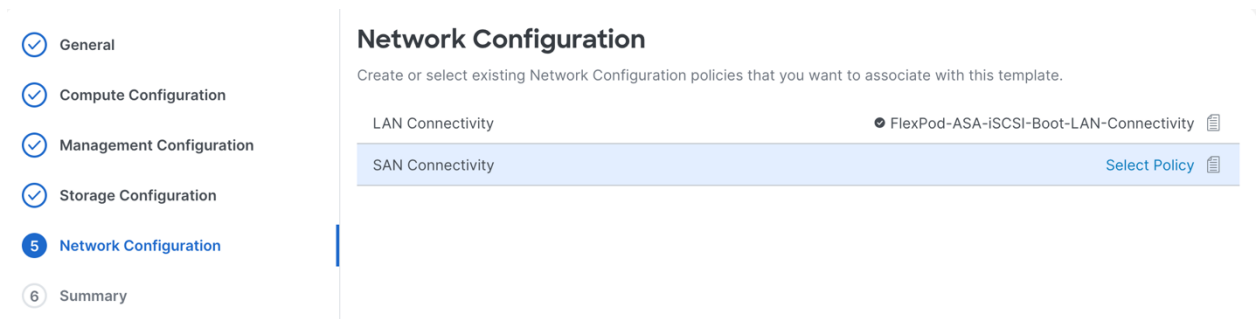
With UCS X-Series Direct and direct-attached ASA storage, you can easily deploy IP-based storage protocols as demonstrated in this NVA. Afterwards, you can add additional FC-based storage protocol access if it is desirable. To accomplish that, the IP SAN-based server profile template and server profiles can be updated to incorporate FC connectivity and FC SAN related configurations and vHBAs for FC and NVMe/FC storage access.

Follow the steps in the subsections below to perform configuration updates for the iSCSI SAN-booted ESXi hosts to also access storage presented via FC and NVMe/FC protocols.

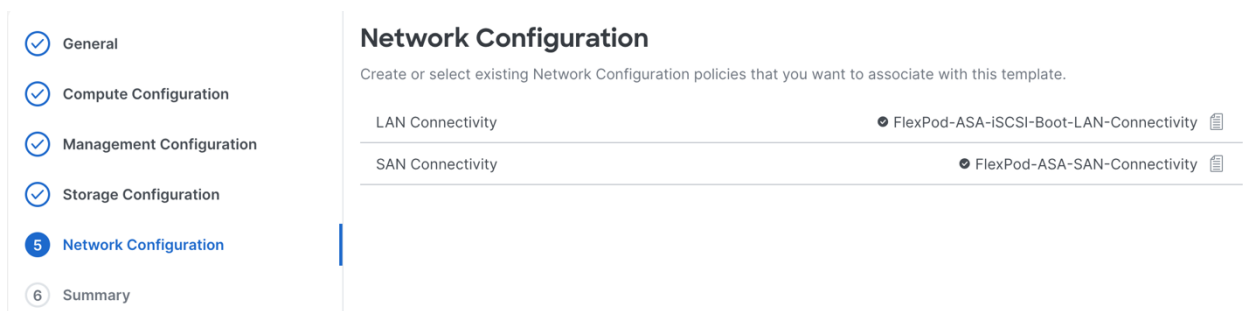
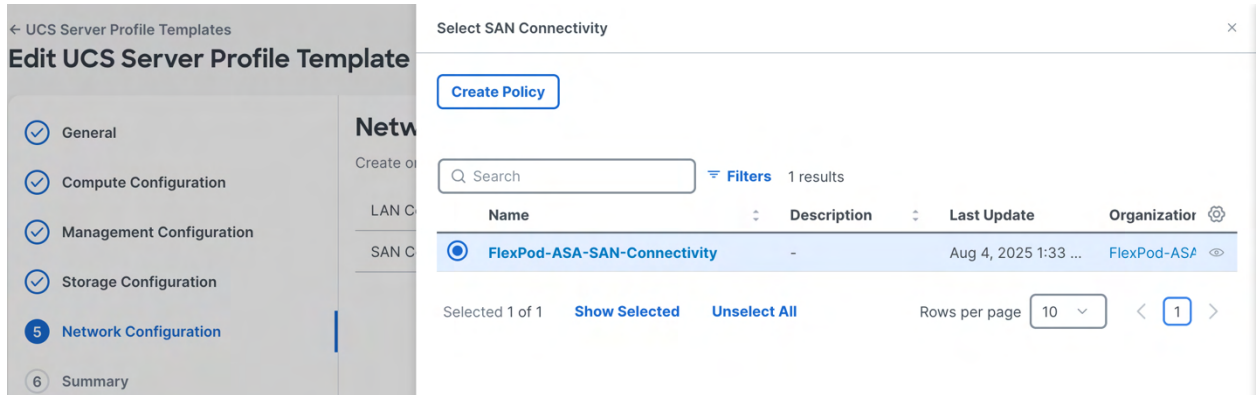
### Update server profile to include FC and NVMe/FC vHBAs

- Login to Cisco Intersight.
- Navigate to Configure > Templates, click on UCS Server Profile Templates tab.
- Select FlexPod-ASA-AMD-iSCSI-Boot template from the list, click on the three dots on the right, and click Edit.
- In the General section, click Next to get to Compute Configuration section.
- Since we are keeping iSCSI SAN boot configuration, we are not changing the boot order.
- Click Next three times to get to the Network Configuration section.
- In the Network configuration section, hover your mouse pointer near the right-hand side of the SAN Connectivity policy icon and click on Select Policy after it becomes available.



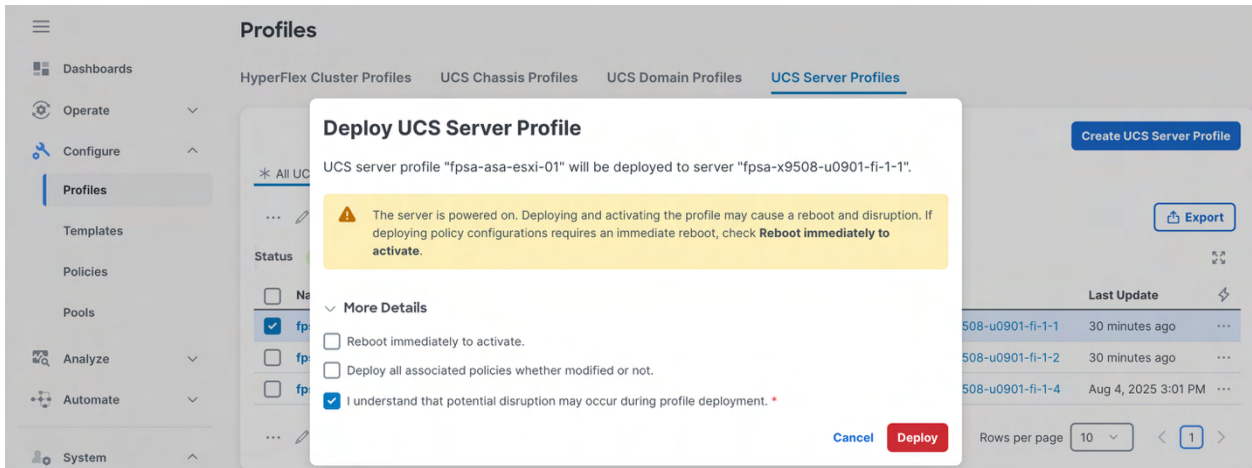


8. Select the previously created FlexPod-ASA-SAN-Connectivity policy and click Select.



9. Click Next for the Summary section.





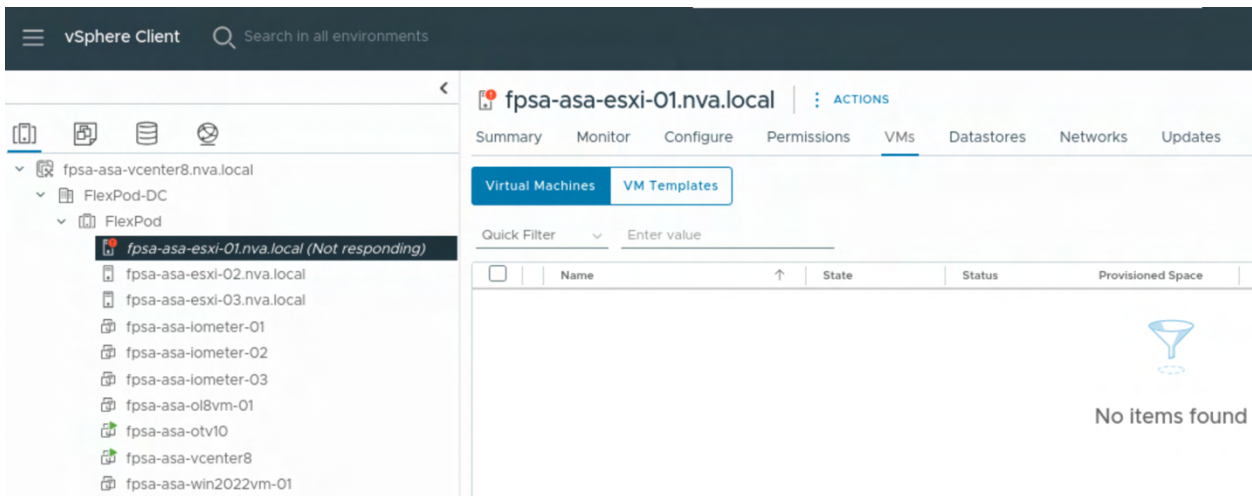
**Note:** Here we are not checking the option to Reboot immediately to activate. Instead, we will go to vCenter to transition the server that will be updated into maintenance mode first and shut down the server before activating the updated server profile.

6. Login to vCenter.

7. In the Inventory view, right-click on the host to be updated and place it into maintenance mode.

**Note:** It might take some time for the VMs running on the host to be migrated off to another host. Once VMs have been migrated off, the host enters maintenance mode.

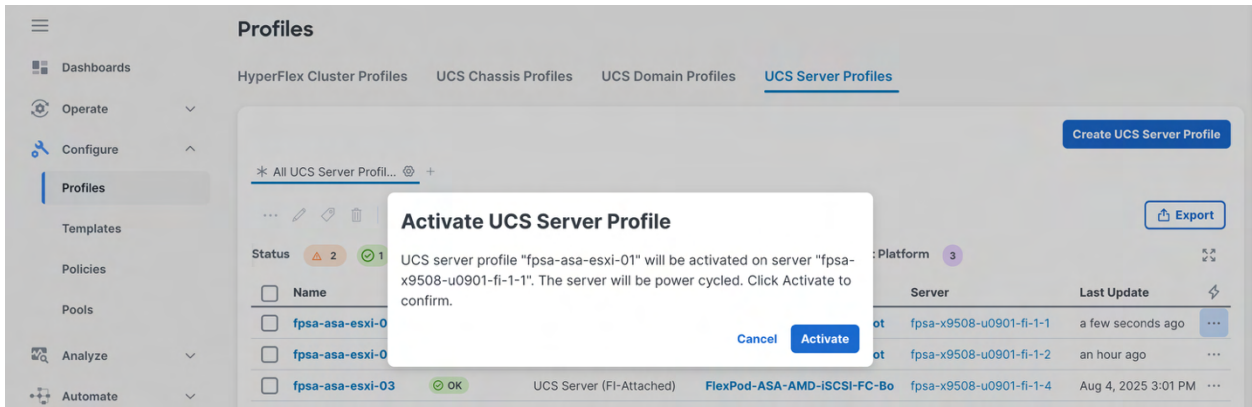
8. Right-click on the host again and select Power > shutdown. Provide a reason for the shutdown and click OK to proceed. Wait till the server shows Not Responding status in vCenter to proceed.



9. Login to Cisco Intersight.

10. Locate the server profile to be activated, click on the three dots on the right, and select Activate.

11. Click Activate to confirm the server profile activation and server reboot.

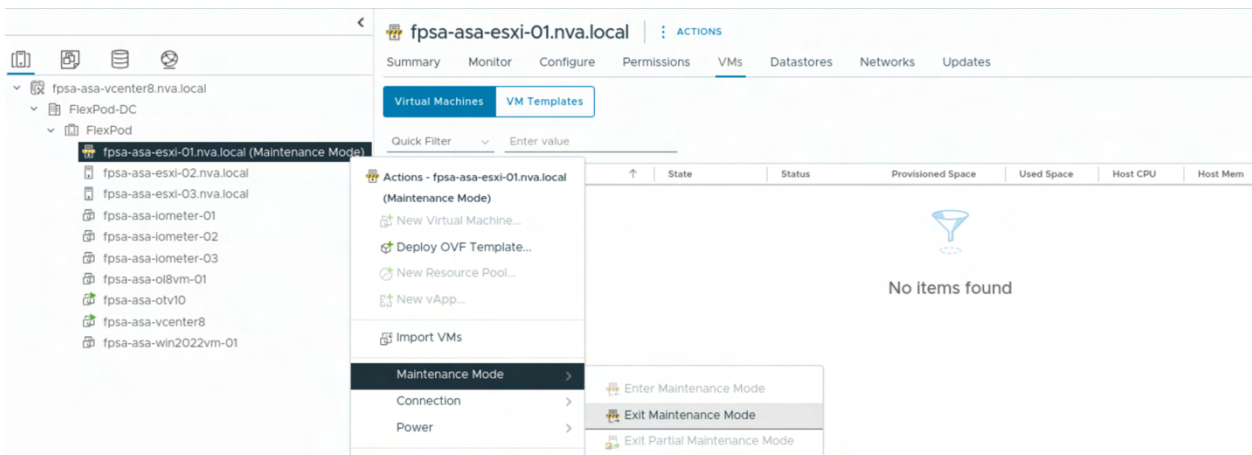


**Note:** You can launch the vKVM available under Server Actions menu for the server profile to monitor the server reboot process. The iSCSI SAN-booted ESXi host should be able to boot from iSCSI SAN as before.

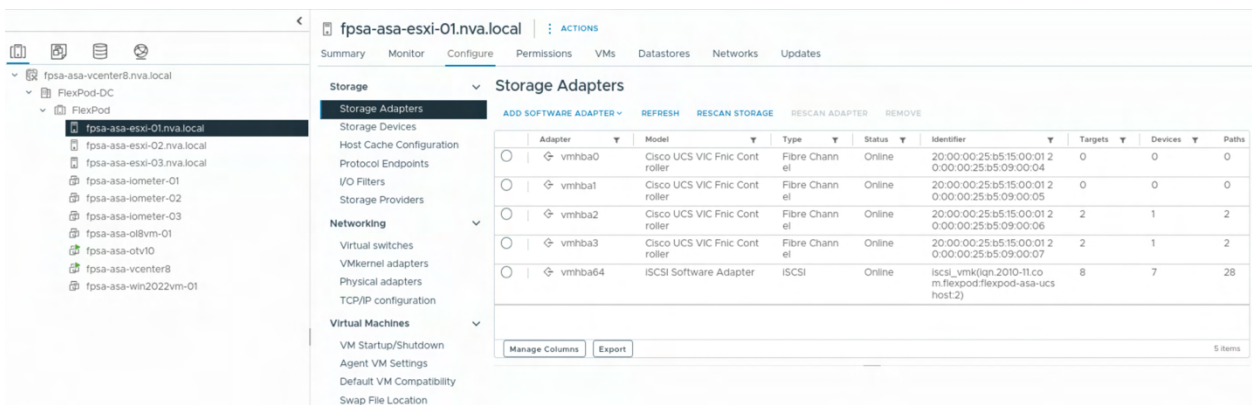
12. Go back to vCenter.

13. Monitor the server status until it shows that the server is connected and in Maintenance Mode.

14. Right-click on the server and select Maintenance Mode > Exit Maintenance Mode from the menu.



15. Select the server, click on Configure tab, and then click on Storage Adapters under Storage. The server should now show four new FC adapters in addition to the iSCSI Software Adapter configured before.



16. Repeat the steps in this section on the remaining iSCSI SAN-booted ESXi hosts to add four additional vHBAs to provide FC and NVMe/FC protocol access to storage.

## Update storage mapping to allow iSCSI SAN-booted ESXi host FC and NVMe/FC storage access

To allow the iSCSI SAN-booted ESXi hosts access to the FC and NVMe/FC storage, add WWPNs of their FC vHBAs into the ESXi cluster igroup for FC and add the host NQNs into the NVMe subsystem host list. Follow the steps below for each ESXi host to provide FC and NVMe/FC storage access to the host.

### Configuration update for FC storage access

1. Login to vCenter.
2. Select an ESXi host, click on Configure tab, click on Storage Adapters under the Storage menu.

Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba0	Cisco UCS VIC Fibre Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:04	0	0	0
vmhba1	Cisco UCS VIC Fibre Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:05	0	0	0

3. Under the Identifier column, it lists both WWNN and WWPN of the adapter separated by a space. Note down the WWPNs for vmhba0 and vmhba1 configured for FC storage access.

**Note:** The WWPNs are also available in the vHBAs' configuration details of the server profile in Intersight, as shown below.

Name	WWPN ID	WWPN Pools	PCI Ord...	Slot ID	PCI Li
vHBA0	20:00:00:25:B5:09:00:04	FlexPod-AS...	6	MLOM	
vHBA1	20:00:00:25:B5:09:00:05	FlexPod-AS...	7	MLOM	
vHBA2	20:00:00:25:B5:09:00:06	FlexPod-AS...	8	MLOM	
vHBA3	20:00:00:25:B5:09:00:07	FlexPod-AS...	9	MLOM	

4. Login to ONTAP.
5. Add the two WWPN of the host vHBAs into the ESXi cluster igroup for shared cluster LUN / datastore access.

```
igroup add -vserver <svm-name> -igroup <igroup-name> -initiator <WWPN1,WWPN2>
```

Example:

```
fpsa-a50-u0909:> igroup add -vserver svm1 -igroup FlexPod-ASA-esxi-cluster-fc -initiator 20:00:00:25:b5:09:00:04,20:00:00:25:b5:09:00:05
```

To verify:

```
fpsa-a50-u0909:> igroup show -igroup FlexPod-ASA-esxi-cluster-fc
Vserver  Igroup          Protocol OS Type  Initiators
```



```

-----
svm1      FlexPod-ASA-esxi-cluster-fc fcp vmware 20:00:00:25:b5:09:00:00
                                                20:00:00:25:b5:09:00:01
                                                20:00:00:25:b5:09:00:04
                                                20:00:00:25:b5:09:00:05

```

6. Repeat the steps above to add FC protocol access for any additional iSCSI SAN booted ESXi hosts.

## Configuration update for NVMe/FC storage access

1. Login to vCenter.
2. Select an ESXi host, click on Configure tab, click on Storage Adapters under the Storage menu.
3. Select vmhba2 from the list of storage adapters.
4. Click on the Controllers tab under the list of storage adapters.
5. Click on ADD Controller.

Storage Adapters configuration page for host **fpsa-asa-esxi-01.nva.local**. The page shows a list of storage adapters. The selected adapter is **vmhba2**, which is a Cisco UCS VIC Fnic Controller. The table below shows the details of the adapters.

Adapter	Model	Type	Status	Identifier
vmhba0	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:04
vmhba1	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:05
vmhba2	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:06
vmhba3	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 20:00:00:25:b5:09:00:07

Below the table, the 'Controllers' tab is selected, showing an 'ADD CONTROLLER' button.

6. Note down the host NQN identifier information.

Add controller | vmhba2

Automatically | Manually

Host NQN: nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01

COPY

**Note:** The iSCSI SAN booted host seems to have a different NQN format compared to that of the FC SAN booted ESXi host.

7. Login to ONTAP.
8. Add the host NQN information to the NVMe subsystem host list for the storage subsystem created for the ESXi cluster NVMe/FC storage access.

```
nvme subsystem host add -subsystem <subsystem-name> -vserver <svm-name> -host-nqn <host-nqn>
```

Example:

```

fpsa-a50-u0909:> nvme subsystem host add -subsystem FlexPod-ASA-esxi-cluster-nvme-fc -vserver
svm1 -host-nqn nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01
(vserver nvme subsystem host add)

```



To verify:

```
fpsa-a50-u0909:> nvme subsystem host show -subsystem FlexPod-ASA-esxi-cluster-nvme-fc
(vserver nvme subsystem host show)
Vserver Subsystem Priority Host NQN
-----
svml     FlexPod-ASA-esxi-cluster-nvme-fc
         regular      nqn.2014-08.com.vmware:nvme:fpsa-asa-esxi-03
         regular      nqn.2014-08.local.nva:nvme:fpsa-asa-esxi-01
2 entries were displayed.
```

6. Repeat the steps above to add NVMe/FC protocol access for any additional iSCSI SAN booted ESXi hosts.

## Confirm FC and NVMe/FC storage access in vCenter

After FC and NVMe/FC storage access have been provided to an iSCSI SAN booted host, follow the steps below to confirm their access in vCenter.

1. Login to vCenter.
2. Select the ESXi host and its Configure tab.
3. Select Storage Adapters under the Storage menu.
4. The FC and NVMe/FC targets and devices should be listed under the FC and NVMe/FC vHBAs respectively.

Adapter	Model	Type	Status	Identifier	Targets	Devices	Paths
vmhba0	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 2 0:00:00:25:b5:09:00:04	2	1	2
vmhba1	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 2 0:00:00:25:b5:09:00:05	2	1	2
vmhba2	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 2 0:00:00:25:b5:09:00:06	2	1	2
vmhba3	Cisco UCS VIC Fnic Controller	Fibre Channel	Online	20:00:00:25:b5:15:00:01 2 0:00:00:25:b5:09:00:07	2	1	2

**Note:** If the targets and devices are not yet available in the Targets and Devices columns, click on the RESCAN STORAGE action above the list of storage adapters to initiate a storage rescan. In this example, only one shared FC LUN and one shared NVMe/FC namespace are presented. So, you should see 1 device each under each adapter with 2 targets/paths.

5. Click on Storage Devices under the Storage menu, select the FC device from the list of Storage Devices. Click on the Paths tab under the storage device list and you should see four total paths with Active(I/O) status indicating the FC LUN presented from the ASA storage has symmetric active/active access from all four available paths.

fpsa-asa-esxi-01.nva.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

**Storage**

- Storage Adapters
- Storage Devices**
- Host Cache Configuration
- Protocol Endpoints
- I/O Filters
- Storage Providers

**Networking**

- Virtual switches
- VMkernel adapters
- Physical adapters
- TCP/IP configuration

**Virtual Machines**

- VM Startup/Shutdown
- Agent VM Settings
- Default VM Compatibility
- Swap File Location

**System**

- Licensing
- Host Profile
- Time Configuration
- Authentication Services
- Certificate
- Power Management

**Storage Devices**

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS HDD DISK MARK AS LOCAL ...

	Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type	Transport
<input checked="" type="checkbox"/>	NETAPP Fibre Channel Disk (naa.600a0980383234486724587338696c32)	1	disk	1.00 TB	fc_datastore_1	Attached	Supported	Flash	Fibre Channel
<input type="checkbox"/>	NETAPP ISCSI Disk (naa.600a0980383234486724587338696c57)	2	disk	400.00 GB	vmware_swap	Attached	Supported	Flash	ISCSI
<input type="checkbox"/>	NETAPP ISCSI Disk (naa.600a0980383234486724587338696d4a)	4	disk	500.00 GB	Not Consumed	Attached	Supported	Flash	ISCSI
<input type="checkbox"/>	NETAPP ISCSI Disk (naa.600a0980383234486724587338696d4a)	0	disk	128.00 GB	Not Consumed	Attached	Supported	Flash	ISCSI

Manage Columns Export Deselect All Items per page 20 9 items

Properties Paths Partition Details

ENABLE DISABLE

	Runtime Name	Status	Target	Transport	Name	Preferred
<input type="radio"/>	vmhba0:C0:T1:L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94 20:06:d0:39:ea:c6:a7:94	Fibre Channel	vmhba0:C0:T1:L1	No
<input type="radio"/>	vmhba1:C0:T2:L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94 20:08:d0:39:ea:c6:a7:94	Fibre Channel	vmhba1:C0:T2:L1	No
<input type="radio"/>	vmhba0:C0:T3:L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94 20:05:d0:39:ea:c6:a7:94	Fibre Channel	vmhba0:C0:T3:L1	No
<input type="radio"/>	vmhba1:C0:T1:L1	Active (I/O)	20:00:d0:39:ea:c6:a7:94 20:02:d0:39:ea:c6:a7:94	Fibre Channel	vmhba1:C0:T1:L1	No

**Note:** NVMe/FC namespace is also presented as symmetric active/active with four total paths as shown in the example below.

fpsa-asa-esxi-01.nva.local | ACTIONS

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

**Storage**

- Storage Adapters
- Storage Devices**
- Host Cache Configuration
- Protocol Endpoints
- I/O Filters
- Storage Providers

**Networking**

- Virtual switches
- VMkernel adapters
- Physical adapters
- TCP/IP configuration

**Virtual Machines**

- VM Startup/Shutdown
- Agent VM Settings
- Default VM Compatibility
- Swap File Location

**System**

- Licensing
- Host Profile
- Time Configuration
- Authentication Services

**Storage Devices**

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS HDD DISK ...

	Name	LUN	Type	Capacity	Datastore	Operational State	Hardware Acceleration	Drive Type	Transport
<input checked="" type="checkbox"/>	NVMe Fibre Channel Disk (uuld.c9d396e5359d11f0a8a6d039eac6a795)	0	disk	1.00 TB	Not Consumed	Attached	Supported	Flash	Fibre Channel
<input type="checkbox"/>	NETAPP ISCSI Disk (naa.600a098038323448723f5877434a5250)	1	disk	5.00 TB	iscsi_datastore_1	Attached	Supported	Flash	ISCSI

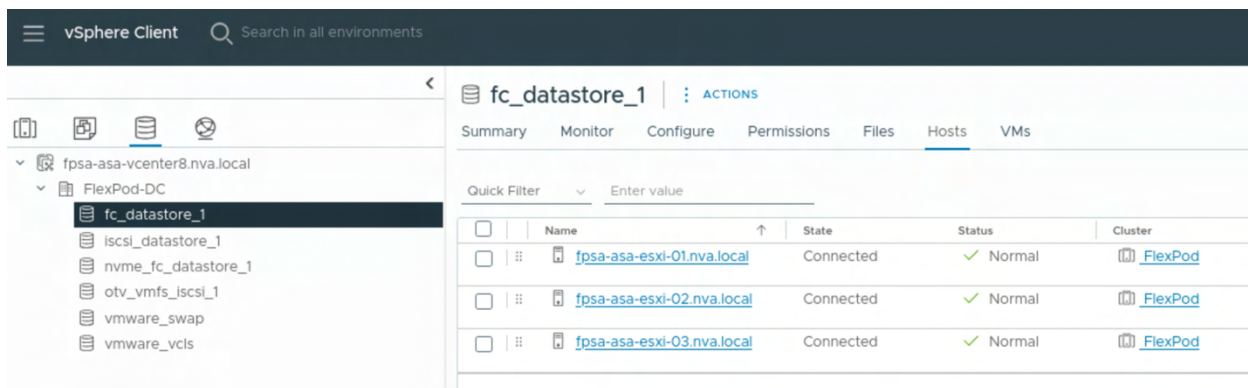
Manage Columns Export Deselect All Items per page 20 9 items

Properties Paths Partition Details

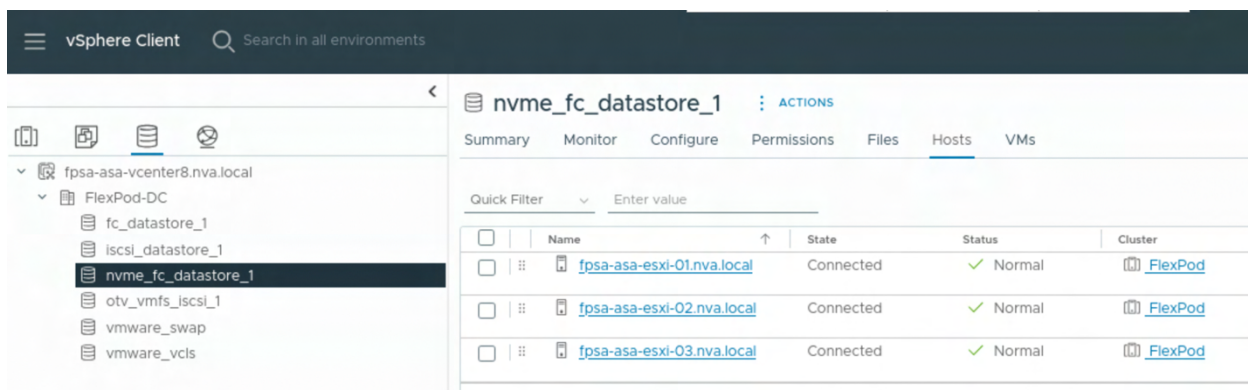
ENABLE DISABLE

	Runtime Name	Status	Target	Transport	Name	Preferred
<input type="radio"/>	vmhba3:C0:T1:L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0c:d0:39:ea:c6:a7:94	Fibre Channel	vmhba3:C0:T1:L0	No
<input type="radio"/>	vmhba3:C0:T0:L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:1c:d0:39:ea:c6:a7:94	Fibre Channel	vmhba3:C0:T0:L0	No
<input type="radio"/>	vmhba2:C0:T1:L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0b:d0:39:ea:c6:a7:94	Fibre Channel	vmhba2:C0:T1:L0	No
<input type="radio"/>	vmhba2:C0:T0:L0	Active (I/O)	20:07:d0:39:ea:c6:a7:94 20:0d:d0:39:ea:c6:a7:94	Fibre Channel	vmhba2:C0:T0:L0	No

6. Select the Storage icon in vCenter menu.
7. Expand the FlexPod data center and select the FC datastore.
8. Click on the Hosts tab to confirm that all hosts are connected to the datastore.



9. Select the NVMe/FC datastore on the left, click on the Hosts tab to confirm that all hosts are connected to the datastore.



## Appendix M: Configuration updates for iSCSI SAN-booted Oracle Linux host to access FC and NVMe/FC storage

With Cisco UCS X-Series Direct and direct-attached NetApp ASA storage, it is very simple to deploy IP-based storage protocols as demonstrated in this NVA. When additional FC-based storage protocol access is desirable, the IP storage-based server profile template and server profiles can be updated to incorporate vHBAs and FC SAN related configurations.

Follow the steps in the subsections below to perform configuration updates for the iSCSI SAN-booted Oracle Linux hosts to access storage using FC and NVMe/FC protocols. For validation, we will stop the NVMe/TCP protocol access to the database created on the NVMe storage, configure the iSCSI SAN-booted Oracle Linux host for FC and NVMe/FC storage access. Then, we will bring up the database already created on NVMe namespaces using NVMe/FC protocol instead.

**Note:** ONTAP does not support having a storage client utilizing both NVMe/TCP and NVMe/FC protocols to access namespaces in the same SVM. If a NVMe storage client requires both NVMe/TCP and NVMe/FC protocol access to two different sets of NVMe namespaces, two different SVMs should be utilized. One SVM can provide NVMe/TCP protocol access to a set of namespaces, while the other SVM can provide NVMe/FC protocol access to another set of namespaces. Since the NVMe storage provisioning is the same in ONTAP whether you use NVMe/TCP or NVMe/FC protocol, we will just utilize the NVMe namespaces which are already created and convert them for access from NVMe/TCP protocol to NVMe/FC protocol.

### Confirm access to database on NVMe storage via NVMe/TCP protocol

1. Login to Linux host as oracle user.
2. Check for nvmedb access.

```
[oracle@fpsa-asa-linux-01 ~]$ srvctl status database -db nvmedb
Instance NVMEDB1 is running on node fpsa-asa-linux-01
```

### 3. Check for NVMe device access via the nvme tool.

```
[oracle@fpsa-asa-linux-01 ~]$ sudo nvme netapp ontapdevices -o json
[sudo] password for oracle:
{
  "ONTAPdevices" : [
    {
      "Device" : "/dev/nvme0n1",
      "Vserver" : "svm1",
      "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_redolog_nvme_1/blocks",
      "NSID" : 1,
      "UUID" : "704f70ff-0d75-11f0-a707-d039eac6a795",
      "Size" : "214.75GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 52428800
    },
    {
      "Device" : "/dev/nvme0n10",
      "Vserver" : "svm1",
      "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_slobdata_nvme_8/blocks",
      "NSID" : 10,
      "UUID" : "b58774fb-0d75-11f0-a707-d039eac6a795",
      "Size" : "858.99GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 209715200
    }
  ],
  ...
}
```

### 4. Check for udev rules used to implement persistent and user-friendly aliases for the NVMe devices.

```
[oracle@fpsa-asa-linux-01 ~]$ sudo ls /etc/udev/rules.d/*nvme*.rules
/etc/udev/rules.d/74-nvme-redolog.rules  /etc/udev/rules.d/75-nvme-slobdata.rules

[oracle@fpsa-asa-linux-01 ~]$ sudo cat /etc/udev/rules.d/74-nvme-redolog.rules
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="704f70ff-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeredolog1", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="76dde16c-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeredolog2", OWNER="grid", GROUP="oinstall", MODE="0660"

[oracle@fpsa-asa-linux-01 ~]$ sudo cat /etc/udev/rules.d/75-nvme-slobdata.rules
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="8c541a42-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata1", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="90b68f3d-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata2", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="97b0b2f0-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata3", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="9dc68622-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata4", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="a4393880-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata5", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="a8afcd44-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata6", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="b0136cc1-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata7", OWNER="grid", GROUP="oinstall", MODE="0660"
KERNEL=="nvme*", SUBSYSTEM=="block", ATTR{uuid}=="b58774fb-0d75-11f0-a707-d039eac6a795",
SYMLINK+="nvmeslobdata8", OWNER="grid", GROUP="oinstall", MODE="0660"
```

### 5. Check for NVMe device persistent naming.

```
[oracle@fpsa-asa-linux-01 ~]$ sudo ls -l /dev/nvmeredolog*
lrwxrwxrwx 1 root root 7 Aug 18 15:52 /dev/nvmeredolog1 -> nvme0n1
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeredolog2 -> nvme0n2

[oracle@fpsa-asa-linux-01 ~]$ sudo ls -l /dev/nvmeslobdata*
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata1 -> nvme0n3
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata2 -> nvme0n4
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata3 -> nvme0n5
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata4 -> nvme0n6
```

```
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata5 -> nvme0n7
lrwxrwxrwx 1 root root 7 Aug 18 15:40 /dev/nvmeslobdata6 -> nvme0n8
lrwxrwxrwx 1 root root 7 Aug 18 15:37 /dev/nvmeslobdata7 -> nvme0n9
lrwxrwxrwx 1 root root 8 Aug 18 15:55 /dev/nvmeslobdata8 -> nvme0n10
```

## Stop Oracle database and disk groups on NVMe/TCP storage

1. Login to Linux host as oracle user.
2. Stop the Oracle database nvmedb that is running on NVMe/TCP storage.

```
[oracle@fpsa-asa-linux-01 ~]$ srvctl status database -db nvmedb
Instance NVMEDB1 is running on node fpsa-asa-linux-01

[oracle@fpsa-asa-linux-01 ~]$ srvctl stop database -db nvmedb

[oracle@fpsa-asa-linux-01 ~]$ srvctl status database -db nvmedb
Instance NVMEDB1 is not running on node fpsa-asa-linux-01
```

3. Login to Linux host as grid user.
4. Invoke asmcmd tool and list the configured disk groups.

```
[grid@fpsa-asa-linux-01 ~]$ asmcmd
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU  Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED    EXTERN    N       512      512             4096  4194304   204800   204468
0           204468      0              Y             ISCSIOCRVOOTE/
MOUNTED    EXTERN    N       512      512             4096  4194304   409600   408848
0           408848      0              N             ISCSIREDOLOG/
MOUNTED    EXTERN    N       512      512             4096  4194304   6553600  939044
0           939044      0              N             ISCSISLOBDATA/
MOUNTED    EXTERN    N       4096     4096            4096  4194304   409600   408848
0           408848      0              N             NVMEREDOLOG/
MOUNTED    EXTERN    N       4096     4096            4096  4194304   6553600  967128
0           967128      0              N             NVMESLOBDATA/
```

5. Umount the two NVMe disk groups and then exit the asmcmd tool.

```
ASMCMD> umount NVMEREDOLOG
ASMCMD> umount NVMESLOBDATA
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU  Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED    EXTERN    N       512      512             4096  4194304   204800   204468
0           204468      0              Y             ISCSIOCRVOOTE/
MOUNTED    EXTERN    N       512      512             4096  4194304   409600   408848
0           408848      0              N             ISCSIREDOLOG/
MOUNTED    EXTERN    N       512      512             4096  4194304   6553600  939044
0           939044      0              N             ISCSISLOBDATA/
ASMCMD> exit
```

## Discontinue host NVMe/TCP protocol access to namespaces

1. Login to Linux host as admin user.
2. List connected NVMe subsystem.

```
[admin@fpsa-asa-linux-01 ~]$ nvme list-subsys
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-
ASA-linux-cluster-nvme
\
+- nvme0 tcp traddr=172.22.78.102 trsvcid=4420 host_traddr=172.22.77.14 src_addr=172.22.77.14
live
+- nvme1 tcp traddr=172.22.77.102 trsvcid=4420 host_traddr=172.22.77.14 src_addr=172.22.77.14
live
+- nvme2 tcp traddr=172.22.78.101 trsvcid=4420 host_traddr=172.22.77.14 src_addr=172.22.77.14
live
+- nvme3 tcp traddr=172.22.77.101 trsvcid=4420 host_traddr=172.22.77.14 src_addr=172.22.77.14
live
```



```

+- nvme4 tcp traddr=172.22.78.102 trsvcid=4420 host_traddr=172.22.78.14 src_addr=172.22.78.14
live
+- nvme5 tcp traddr=172.22.77.102 trsvcid=4420 host_traddr=172.22.78.14 src_addr=172.22.78.14
live
+- nvme6 tcp traddr=172.22.78.101 trsvcid=4420 host_traddr=172.22.78.14 src_addr=172.22.78.14
live
+- nvme7 tcp traddr=172.22.77.101 trsvcid=4420 host_traddr=172.22.78.14 src_addr=172.22.78.14
live

```

### 3. Disconnect from the NVMe subsystem.

```

[admin@fpsa-asa-linux-01 ~]$ sudo nvme disconnect --nqn=nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-ASA-linux-cluster-nvme
NQN:nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-ASA-linux-cluster-nvme disconnected 8 controller(s)
[admin@fpsa-asa-linux-01 ~]$ nvme list-subsys
[admin@fpsa-asa-linux-01 ~]$

```

## Bring down ONTAP NVMe/TCP LIFs when no longer needed

1. Login to ONTAP storage.
2. Bring down NVMe/TCP LIFs.

```

fpsa-a50-u0909::> net int show nvmetcp-lif*
(network interface show)

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	nvmetcp-lif-01a	up/up	172.22.77.101/24	fpsa-a50-u0909-01	e2b-2277	true
	nvmetcp-lif-01b	up/up	172.22.78.101/24	fpsa-a50-u0909-01	e4b-2278	true
	nvmetcp-lif-02a	up/up	172.22.77.102/24	fpsa-a50-u0909-02	e2b-2277	true
	nvmetcp-lif-02b	up/up	172.22.78.102/24	fpsa-a50-u0909-02	e4b-2278	true

```

4 entries were displayed.

fpsa-a50-u0909::> net int modify -vserver svml nvmetcp-lif* -status-admin down
(network interface modify)
4 entries were modified.

fpsa-a50-u0909::> net int show nvmetcp-lif*
(network interface show)

```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
svml	nvmetcp-lif-01a	down/down	172.22.77.101/24	fpsa-a50-u0909-01	e2b-2277	true
	nvmetcp-lif-01b	down/down	172.22.78.101/24	fpsa-a50-u0909-01	e4b-2278	true
	nvmetcp-lif-02a	down/down	172.22.77.102/24	fpsa-a50-u0909-02	e2b-2277	true
	nvmetcp-lif-02b	down/down	172.22.78.102/24	fpsa-a50-u0909-02	e4b-2278	true

```

4 entries were displayed.

```

**Note:** Before bringing down the NVMe/TCP LIFs, ensure that no other clients are utilizing NVMe storage via those NVMe/TCP LIFs.

## Update server profile to include FC and NVMe/FC vHBAs

1. Login to Cisco Intersight.
2. Navigate to Configure > Templates, click on UCS Server Profile Templates tab.
3. Select FlexPod-ASA-AMD-Linux-iSCSI-Boot template from the list, click on the three dots on the right, and click Edit.
4. In the General section, click Next to get to Compute Configuration section.
5. Since we are keeping iSCSI SAN boot configuration, we are not changing the boot order.
6. Click Next three times to get to the Network Configuration section.



7. In the Network configuration section, hover near the right-hand side of the SAN Connectivity policy icon and click on Select Policy after it becomes available.

✓ General

✓ Compute Configuration

✓ Management Configuration

✓ Storage Configuration

5 Network Configuration

6 Summary

### Network Configuration

Create or select existing Network Configuration policies that you want to associate with this template.

LAN Connectivity

● FlexPod-ASA-Linux-iSCSI-Boot-LAN-Connectivity

SAN Connectivity

Select Policy

8. Select the previously created FlexPod-ASA-SAN-Connectivity policy and click Select.

← UCS Server Profile Templates

Edit UCS Server Profile Template

✓ General

✓ Compute Configuration

✓ Management Configuration

✓ Storage Configuration

5 Network Configuration

6 Summary

Select SAN Connectivity

Create Policy

Q Search

Filters

1 results

Name	Description	Last Update	Organization
<input type="radio"/> FlexPod-ASA-SAN-Connectivity	-	a few seconds ago	FlexPod-ASA

Rows per page 10 < 1 >

✓ General

✓ Compute Configuration

✓ Management Configuration

✓ Storage Configuration

5 Network Configuration

6 Summary

### Network Configuration

Create or select existing Network Configuration policies that you want to associate with this template.

LAN Connectivity

● FlexPod-ASA-Linux-iSCSI-Boot-LAN-Connectivity

SAN Connectivity

● FlexPod-ASA-SAN-Connectivity

9. Click Next for the Summary section.

✓ General

✓ Compute Configuration

✓ Management Configuration

✓ Storage Configuration

✓ Network Configuration

6 Summary

### Summary

Verify details of the template and the policies, resolve errors, and deploy.

^ General

Name

FlexPod-ASA-AMD-Linux-iSCSI-Boot

Organization

FlexPod-ASA

Target Platform

UCS Server (FI-Attached)

Compute Configuration	Management Configuration	Storage Configuration	Network Configuration	Errors/Warnings (0)
BIOS			FlexPod-ASA-AMD-M8-Virt-BIOS	
Boot Order			FlexPod-ASA-iSCSI-Boot-Order	
UUID			FlexPod-ASA-UUID-Pool	
Virtual Media			FlexPod-ASA-KVM-Mount-Media	

Close

Back

Derive Profiles

10. Review the configurations in the Summary section and click Close when done.

**Note:** Since we will be updating existing server profiles to deploy the changes, we are not selecting Derive Profiles here, which is for creating new server profiles for additional servers.

## Create a server boot order policy to disable secure boot

As part of the Oracle Linux host configuration update to support FC and NVMe/FC protocols, we will need to update the UCS VIC driver for FC and NVMe/FC. To be able to run the updated driver module, we will need to disable secure boot for Oracle Linux.

We will clone a copy of the iSCSI boot order for modification and disable secure boot in the cloned copy, to avoid affecting secure boot configured for ESXi hosts. To create a clone of iSCSI boot order and disable secure boot for Oracle Linux, follow the steps below.

1. Login to Cisco Intersight.
2. Navigate to Configure > Policies.
3. Narrow down the list of policies by providing a search string in the filter.
4. Select FlexPod-ASA-iSCSI-Boot-Order policy, click on the three dots on the right, and select Clone.
5. Edit the Policy Name as needed, select the FlexPod ASA organization, and click Clone.

Policies > Boot Order > FlexPod-ASA-iSCSI-Boot-Order

## Clone

Platform Type	Type
UCS Server	Boot Order

Tags	Description
-	-

**Clone Details**

**Policy Name \***  
FlexPod-ASA-Linux-iSCSI-Boot-Order

**Organization \***  
FlexPod-ASA

**Description**  
Description 0 / 1024

**Set Tags**  
Enter a tag in the key:value format.

6. Select the newly cloned iSCSI boot order policy, click the three dots on the right, and click Edit.
7. Click Next for the Policy Details section, unselect Enable Secure Boot, and click Save.

Policies > Boot Order > FlexPod-ASA-Linux-iSCSI-Boot-Order

## Edit

General

**2 Policy Details**

### Policy Details

Add policy details.

**Configured Boot Mode**

☒ Unified Extensible Firmware Interface (UEFI) ☐ Legacy

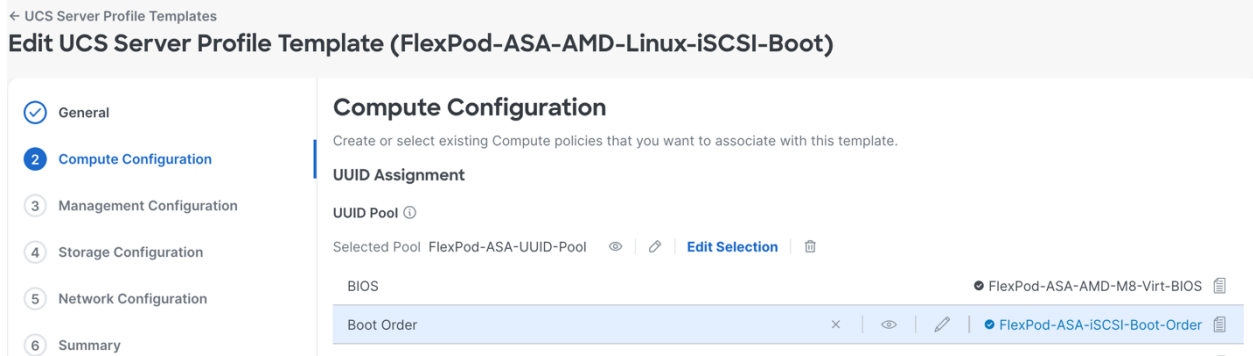
☐ Enable Secure Boot

The boot device name cannot be a reserved keyword. However, it can be used in conjunction with other letters, numbers, underscores, and hyphens (EU\_CDROM-3, etc.). Reserved words include: all, ALL, CDROM, EFI, EOD, FDD, HDD, HDDANY, HTTP, ISCSI, ISCSIAN, LOCALCDD, LOCALHDD, NULL, NVME, NVMEANY, PCHSTORAGE, PCHSTORANY, PXE, SAN, SANANY, SDANY, SDCARD, UEFISHELL, USB, USBOD, USBFDD, USBHDD, VMCIMCCD, VMCIMCHDD, VMEDIA, VMFDD, VMKVMCD, VMKVMHDD.

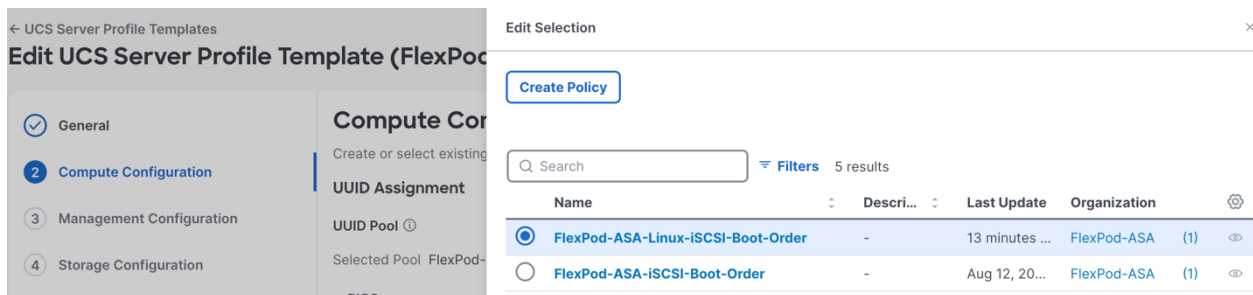
**Add Boot Device**

+ Virtual Media (KVM-Mapped-DVD)	<input checked="" type="checkbox"/> Enabled
+ iSCSI Boot (iSCSI-A-Boot)	<input checked="" type="checkbox"/> Enabled
+ iSCSI Boot (iSCSI-B-Boot)	<input checked="" type="checkbox"/> Enabled

8. Navigate to Configure > Templates, and click on UCS Server Profile Templates tab.
9. Narrow down search with a search string in the filter as needed to find the server profile template FlexPod-ASA-AMD-Linux-iSCSI-Boot. Click the three dots on the right and select Edit.
10. Click Next in the General section to get to Compute Configuration section.



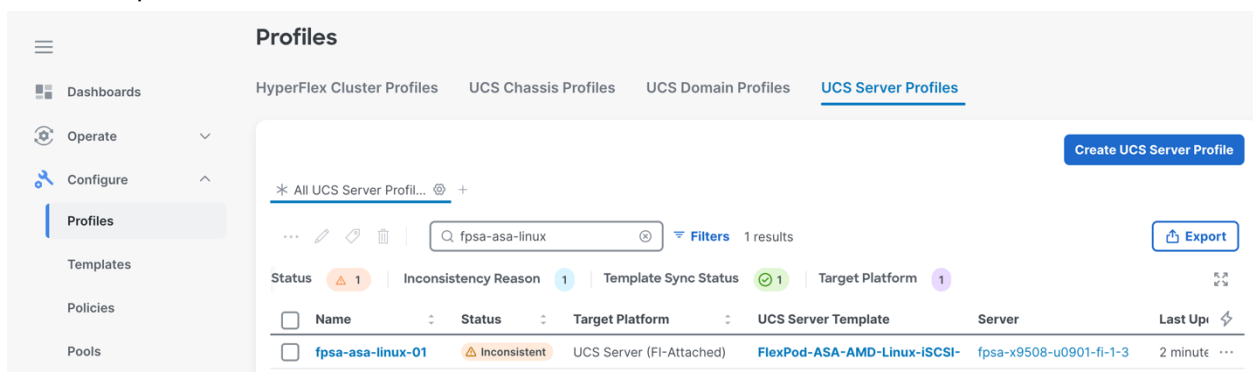
- Click on the policy name for Boot Order to open the dialog for policy selection and select the newly created FlexPod-ASA-Linux-iSCSI-Boot-Order policy.



- Click Save on the Boot Order policy selection screen.
- Click Close on the Compute Configuration section to exit the server profile template update.

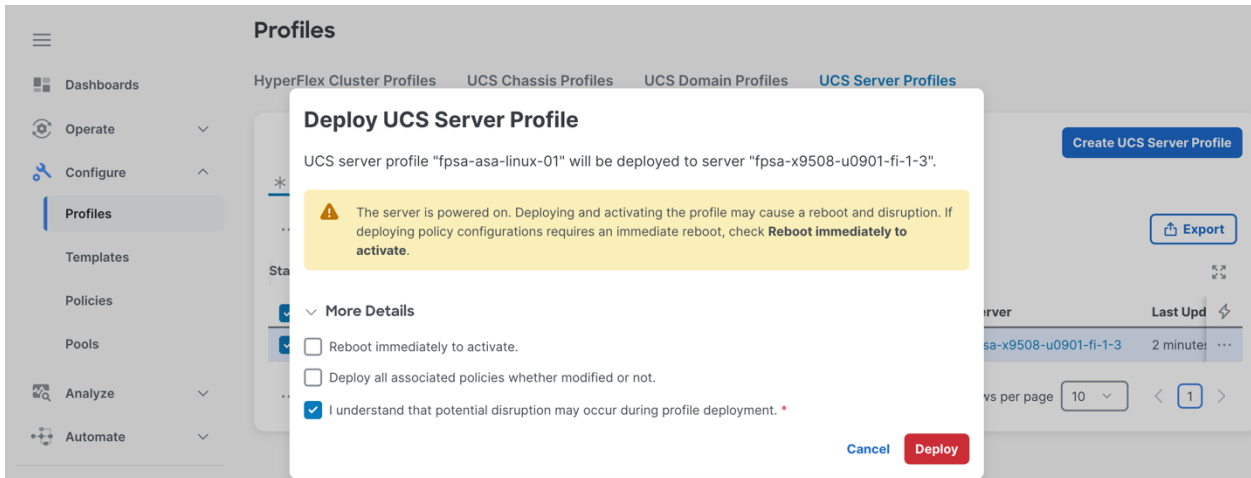
## Deploy updated UCS server profile

- Navigate to Configure > Profiles and click on UCS Server Profiles tab.
- In the Filter box, provide partial name of the server profiles to narrow down the display of relevant server profiles.



**Note:** Since the template used for deriving the iSCSI SAN-booted Linux host had changed, the Linux server profile showed a status of Inconsistent.

- Select the server profile to be re-deployed, click on the three dots on the right, and click Deploy.
- On the Deploy UCS Server Profile screen, read the warnings about deploying and activating the profile may cause a reboot and disruption. Click Deploy.



**Note:** Here we are not checking the option to Reboot immediately to activate. Instead, we will shut down the Oracle Linux host before activating the updated server profile.

5. From the server profile actions, select Power > Shutdown Operating System.
6. Launch server vKVM console, wait to confirm that the server had been shut down. Then, select Power > Power on System.

## Update UCS FC and NVMe/FC driver on Oracle Linux

To install or update the fnic driver to support FC and NVMe/FC protocols on Oracle Linux, follow the steps below.

1. Download the Linux driver release ISO from Cisco Software Download site, <https://software.cisco.com/download/home>, and transfer the driver ISO to the /tmp directory on the Oracle Linux server.

**Note:** For this example, the 4.3(5g) release was downloaded through searching for UCS X215c M8 Compute Node, select Unified Computing System (UCS) Drivers software type, and then expand the driver releases to find 4.3(5g) release, which is the UCSX-Drivers-linux.4.3.5g.iso.

2. Login to the Oracle Linux host as root.
3. Mount the ISO file under /mnt/cdrom.

```
[root@fpga-asa-linux-01 ~]# cd /tmp
[root@fpga-asa-linux-01 tmp]# mount -t iso9660 UCSX-drivers-linux.4.3.5g.iso /mnt/cdrom
mount: /mnt/cdrom: WARNING: device write-protected, mounted read-only.
```

4. Locate the Cisco VIC driver for the kernel version of the Oracle Linux. See below for an example of using UEK7U2 UEK kernel version.

```
[root@fpga-asa-linux-01 tmp]# uname -a
Linux fpga-asa-linux-01 5.15.0-206.153.7.1.el8uek.x86_64 #2 SMP Wed May 22 20:49:34 PDT 2024
x86_64 x86_64 x86_64 GNU/Linux

[root@fpga-asa-linux-01 UEK7U2]# cd /mnt/cdrom/Storage/Cisco/VIC/Oracle/UEK7U2/OL8.5/
[root@fpga-asa-linux-01 OL8.5]# ls
attr.json  fnic-2.0.0.100-oe18u5x_5.15.0_200.131.27-dd.iso  README.txt
attr.txt   fnic-2.0.0.100.tar.bz2
fcc        kmod-fnic-2.0.0.100-321.0.oe18u5x_5.15.0_200.131.27uek.x86_64.rpm
```

**Note:** Be sure to select a driver version which matches the major and minor versions of the running kernel for compatibility.

5. Follow the information in the driver README.txt file to install or update the fnic driver and reboot the host after installation / update.

```
[root@fpsa-asa-linux-01 OL8.5]# rpm -ivh kmod-fnic-2.0.0.100-321.0.oel8u5x_5.15.0_200.131.27uek.x86_64.rpm
Verifying... ##### [100%]
Preparing... ##### [100%]
Updating / installing...
 1:kmod-fnic-2.0.0.100-321.0.oel8u5x##### [100%]
[root@fpsa-asa-linux-01 OL8.5]# cd /tmp
[root@fpsa-asa-linux-01 tmp]# umount /mnt/cdrom
[root@fpsa-asa-linux-01 tmp]# reboot
```

6. After the host is rebooted, check to confirm that the fnic driver modules are running with nvme\_fc support.

```
[root@fpsa-asa-linux-01 ~]# lsmod | grep fnic
fnic                475136  0
nvme_fc              61440   1 fnic
scsi_transport_fc    90112   1 fnic
```

**Note:** If the fnic module is not running, use the modprobe command to try to load the fnic module. The following is the error you will see when trying to load fnic module when secure boot is enabled. Refer to the earlier subsection on disable secure boot in the iSCSI boot order policy update for the Linux host.

```
[root@fpsa-asa-linux-01 OL8.5]# modprobe fnic
modprobe: ERROR: could not insert 'fnic': Key was rejected by service
```

## FC storage access configurations

To allow the iSCSI SAN-booted Oracle Linux host access to storage via FC protocol, you need to create the needed LUNs and an igroup, add host vHBAs' WWPNs to the igroup, map the created LUNs to the igroup, and then issue discovery command on the host for it to see the newly mapped LUNs.

### Gather Linux host WWPN information

1. Login to Linux host as root.
2. Obtain host vHBA WWPNs with the fcc tool installed with the fnic driver.

```
[root@fpsa-asa-linux-01 ~]# which fcc
/usr/bin/fcc
[root@fpsa-asa-linux-01 ~]# fcc
FC HBAs:
HBA      Port Name      Port ID  State   Device
host0    20:00:00:25:b5:09:00:0c 23:01:40 Online  fnic0
host1    20:00:00:25:b5:09:00:0d 1c:01:40 Online  fnic1

host0 Remote Ports:
Path      Port Name      Port ID  State   Roles
0:0-0     20:06:d0:39:ea:c6:a7:94 23:00:01 Online  FCP Target
0:0-1     20:05:d0:39:ea:c6:a7:94 23:00:21 Online  FCP Target

host1 Remote Ports:
Path      Port Name      Port ID  State   Roles
1:0-0     20:02:d0:39:ea:c6:a7:94 1c:00:01 Online  FCP Target
1:0-1     20:08:d0:39:ea:c6:a7:94 1c:00:41 Online  FCP Target
```

**Note:** The host HBAs' WWPNs are the two Port Names listed next to the host0 and host1 HBAs.

### Create and map LUNs on storage

1. Login to ONTAP CLI.
2. Create a set of LUNs for the Oracle Linux database usage via FC protocol. Adjust the number of LUNs and their sizes to suit your database requirements.

```
lun create -vserver <svm-name> -path <path-name> -size <lun-size> -ostype linux
```

Example:



```
fpsa-a50-u0909::> lun create -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_1 -size 500g -ostype linux
fpsa-a50-u0909::> lun create -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_2 -size 500g -ostype linux
fpsa-a50-u0909::> lun create -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_3 -size 500g -ostype linux
fpsa-a50-u0909::> lun create -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_4 -size 500g -ostype linux
```

To verify:

```
fpsa-a50-u0909::> lun show -path fpsa_asa_linux_01_oracle_slobdata_fc*
Vserver Path State Mapped Type Size
-----
svml fpsa_asa_linux_01_oracle_slobdata_fc_1 online unmapped linux 500GB
svml fpsa_asa_linux_01_oracle_slobdata_fc_2 online unmapped linux 500GB
svml fpsa_asa_linux_01_oracle_slobdata_fc_3 online unmapped linux 500GB
svml fpsa_asa_linux_01_oracle_slobdata_fc_4 online unmapped linux 500GB
4 entries were displayed.
```

### 3. Create an igroup with the Linux host HBA WWPNs for FC LUN mapping.

```
igroup create -vserver <svm-name> -igroup <igroup-name> -protocol fcp -ostype linux -initiator <WWPN1,WWPN2>
```

Example:

```
fpsa-a50-u0909::> igroup create -vserver svml -igroup FlexPod-ASA-linux-cluster-fc -protocol fcp -ostype linux -initiator 20:00:00:25:b5:09:00:0c,20:00:00:25:b5:09:00:0d
```

To verify:

```
fpsa-a50-u0909::> igroup show -igroup FlexPod-ASA-linux-cluster-fc
Vserver Igroup Protocol OS Type Initiators
-----
svml FlexPod-ASA-linux-cluster-fc fcp linux 20:00:00:25:b5:09:00:0c
20:00:00:25:b5:09:00:0d
```

### 4. Map the FC LUNs to the Linux host igroup.

```
lun map -vserver <svm-name> -path <path-name> -igroup <igroup-name> -lun-id <lun-id>
```

Example:

```
fpsa-a50-u0909::> lun map -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_1 -igroup FlexPod-ASA-linux-cluster-fc -lun-id 1
fpsa-a50-u0909::> lun map -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_2 -igroup FlexPod-ASA-linux-cluster-fc -lun-id 2
fpsa-a50-u0909::> lun map -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_3 -igroup FlexPod-ASA-linux-cluster-fc -lun-id 3
fpsa-a50-u0909::> lun map -vserver svml -path fpsa_asa_linux_01_oracle_slobdata_fc_4 -igroup FlexPod-ASA-linux-cluster-fc -lun-id 4
```

To verify:

```
fpsa-a50-u0909::> lun show -m -igroup FlexPod-ASA-linux-cluster-fc
Vserver Path Igroup LUN ID Protocol
-----
svml fpsa_asa_linux_01_oracle_slobdata_fc_1 FlexPod-ASA-linux-cluster-fc 1 fcp
svml fpsa_asa_linux_01_oracle_slobdata_fc_2 FlexPod-ASA-linux-cluster-fc 2 fcp
svml fpsa_asa_linux_01_oracle_slobdata_fc_3 FlexPod-ASA-linux-cluster-fc 3 fcp
svml fpsa_asa_linux_01_oracle_slobdata_fc_4 FlexPod-ASA-linux-cluster-fc 4 fcp
4 entries were displayed.
```

## Discover the mapped LUNs on the host

1. Login to Linux host as root user.
2. Perform HBA reset with the fcc tool and then list the discovered LUNs.

```

[root@fpsa-asa-linux-01 ~]# fcc reset host0
reset host0
[root@fpsa-asa-linux-01 ~]# fcc reset host1
reset host1

[root@fpsa-asa-linux-01 ~]# fcc
FC HBAs:
HBA          Port Name          Port ID   State   Device
host0        20:00:00:25:b5:09:00:0c 23:01:40 Online  fnic0
host1        20:00:00:25:b5:09:00:0d 1c:01:40 Online  fnic1

host0 Remote Ports:
Path      Port Name          Port ID   State   Roles
0:0:0-0   20:06:d0:39:ea:c6:a7:94 23:00:01 Online  FCP Target
0:0:0-1   20:05:d0:39:ea:c6:a7:94 23:00:21 Online  FCP Target

host0 LUNs:
Path      Device      Size  Vendor      Model      State
0:0:0:1   sdba       536 GB NETAPP      LUN C-Mode running
0:0:0:2   sdbb       536 GB NETAPP      LUN C-Mode running
0:0:0:3   sdbc       536 GB NETAPP      LUN C-Mode running
0:0:0:4   sdbd       536 GB NETAPP      LUN C-Mode running
0:0:1:1   sdbe       536 GB NETAPP      LUN C-Mode running
0:0:1:2   sdbf       536 GB NETAPP      LUN C-Mode running
0:0:1:3   sdbg       536 GB NETAPP      LUN C-Mode running
0:0:1:4   sdbh       536 GB NETAPP      LUN C-Mode running

host1 Remote Ports:
Path      Port Name          Port ID   State   Roles
1:0:0-0   20:02:d0:39:ea:c6:a7:94 1c:00:01 Online  FCP Target
1:0:0-1   20:08:d0:39:ea:c6:a7:94 1c:00:41 Online  FCP Target

host1 LUNs:
Path      Device      Size  Vendor      Model      State
1:0:0:1   sdbm       536 GB NETAPP      LUN C-Mode running
1:0:0:2   sdbn       536 GB NETAPP      LUN C-Mode running
1:0:0:3   sdbo       536 GB NETAPP      LUN C-Mode running
1:0:0:4   sdbp       536 GB NETAPP      LUN C-Mode running
1:0:1:1   sdbi       536 GB NETAPP      LUN C-Mode running
1:0:1:2   sdbj       536 GB NETAPP      LUN C-Mode running
1:0:1:3   sdbk       536 GB NETAPP      LUN C-Mode running
1:0:1:4   sdbl       536 GB NETAPP      LUN C-Mode running

```

**Note:** Even though there are only 4 LUNs mapped in this example, both host0 and host1 HBA discovered 8 devices due to having two paths per fabric.

3. Use the previously installed sanlun tool from NetApp Linux Host Utilities to show the discovered NetApp LUNs and their attributes.

```

[root@fpsa-asa-linux-01 ~]# sanlun lun show -p
...
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_1
LUN: 1
LUN Size: 500g
Product: cDOT
Host Device: 3600a0980383234486724587338696e35
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host      vserver
path      /dev/      host      vserver
state     type       node     adapter   LIF
-----
up        primary   sdba     host0     fc-lif-02a
up        primary   sdbe     host0     fc-lif-01a
up        primary   sdbi     host1     fc-lif-02b
up        primary   sdbm     host1     fc-lif-01b

```

ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_2				
LUN: 2				
LUN Size: 500g				
Product: cDOT				
Host Device: 3600a0980383234486724587338696e36				
Multipath Policy: service-time 0				
Multipath Provider: Native				
-----				
host	vserver			
path	path	/dev/	host	vserver
state	type	node	adapter	LIF
-----				
up	primary	sdbb	host0	fc-lif-02a
up	primary	sdbf	host0	fc-lif-01a
up	primary	sdbj	host1	fc-lif-02b
up	primary	sdbn	host1	fc-lif-01b
...				

**Note:** The partial sanlun lun show -p outputs above show two of the four mapped FC LUNs, with four paths per LUN. In addition to the FC LUNs, the command also includes the iSCSI LUNs which were previously mapped to the host, but they are omitted from the outputs above.

**Note:** To utilize the FC LUNs mapped to the Linux host, follow the same processes used for iSCSI LUNs such as creating udev rules for consistent device alias names, etc.

### Use FC LUNs for Oracle database testing

Here is a list of mapped LUNs from storage, after the FC slobdata LUNs were resized from 500G to 800G and four additional LUNs for SLOB database and two additional LUNs for redolog were created.

fpsa-a50-u0909::*> lun show -m -igroup FlexPod-ASA-linux-cluster-fc				
Vserver	Path	Igroup	LUN ID	Protocol
-----				
svml	fpsa_asa_linux_01_oracle_redolog_fc_1	FlexPod-ASA-linux-cluster-fc	11	fc
svml	fpsa_asa_linux_01_oracle_redolog_fc_2	FlexPod-ASA-linux-cluster-fc	12	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_1	FlexPod-ASA-linux-cluster-fc	1	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_2	FlexPod-ASA-linux-cluster-fc	2	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_3	FlexPod-ASA-linux-cluster-fc	3	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_4	FlexPod-ASA-linux-cluster-fc	4	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_5	FlexPod-ASA-linux-cluster-fc	5	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_6	FlexPod-ASA-linux-cluster-fc	6	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_7	FlexPod-ASA-linux-cluster-fc	7	fc
svml	fpsa_asa_linux_01_oracle_slobdata_fc_8	FlexPod-ASA-linux-cluster-fc	8	fc
10 entries were displayed.				

After LUNs are mapped to the igroup for the Oracle Linux host, perform a SCSI rescan operation in Oracle Linux as the root user.

[root@fpsa-asa-linux-01 ~]# rescan-scsi-bus.sh
Scanning SCSI subsystem for new devices
Scanning host 0 for all SCSI target IDs, all LUNs
Scanning for device 0 0 0 1 ...
OLD: Host: scsi0 Channel: 00 Id: 00 Lun: 01
Vendor: NETAPP Model: LUN C-Mode Rev: 9161
Type: Direct-Access ANSI SCSI revision: 05
Scanning for device 0 0 0 2 ...
...

Use sanlun tool to correlate the LUN UUID and LUN path for dm-mp device alias creation.

[root@fpsa-asa-linux-01 ~]# sanlun lun show -p   egrep "ONTAP Path Host Device"
...
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_8
Host Device: 3600a0980383234486724587338696e2d
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_1
Host Device: 3600a0980383234486724587338696e35
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_2

```

Host Device: 3600a0980383234486724587338696e36
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_3
Host Device: 3600a0980383234486724587338696e37
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_4
Host Device: 3600a0980383234486724587338696e38
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_6
Host Device: 3600a0980383234486724587338696e39
ONTAP Path: svml:fpsa_asa_linux_01_oracle_redolog_fc_1
Host Device: 3600a0980383234486724587338696e41
ONTAP Path: svml:fpsa_asa_linux_01_oracle_redolog_fc_2
Host Device: 3600a0980383234486724587338696e42
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_5
Host Device: 3600a098038323448723f5877434a5274
ONTAP Path: svml:fpsa_asa_linux_01_oracle_slobdata_fc_7
Host Device: 3600a098038323448723f5877434a5275

```

Create udev rules to map the FC devices to persistent device names.

```

[root@fpsa-asa-linux-01 ~]# cat /etc/udev/rules.d/76-fc-redolog.rules
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e41", SYMLINK+="fcredolog1",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e42", SYMLINK+="fcredolog2",
OWNER="grid", GROUP="oinstall", MODE="0660"

[root@fpsa-asa-linux-01 ~]# cat /etc/udev/rules.d/77-fc-slobdata.rules
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e35", SYMLINK+="fcslobdata1",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e36", SYMLINK+="fcslobdata2",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e37", SYMLINK+="fcslobdata3",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e38", SYMLINK+="fcslobdata4",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a098038323448723f5877434a5274", SYMLINK+="fcslobdata5",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e39", SYMLINK+="fcslobdata6",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a098038323448723f5877434a5275", SYMLINK+="fcslobdata7",
OWNER="grid", GROUP="oinstall", MODE="0660"
ACTION=="add|change", ENV{DM_NAME}=="3600a0980383234486724587338696e2d", SYMLINK+="fcslobdata8",
OWNER="grid", GROUP="oinstall", MODE="0660"

```

Reload udev rules for them to take effect and check for the persistent device name mappings.

```

[root@fpsa-asa-linux-01 ~]# udevadm control --reload-rules
[root@fpsa-asa-linux-01 ~]# udevadm trigger --type=devices --action=change

[root@fpsa-asa-linux-01 ~]# ls -l /dev/fc*
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcredolog1 -> dm-27
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcredolog2 -> dm-28
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata1 -> dm-19
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata2 -> dm-21
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata3 -> dm-20
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata4 -> dm-22
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata5 -> dm-23
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata6 -> dm-24
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata7 -> dm-25
lrwxrwxrwx 1 root root 5 Aug 29 15:59 /dev/fcslobdata8 -> dm-26

```

As grid user, use asmca tool to create two FC disk groups. Please refer to earlier iSCSI section for details. Afterwards, use the asmcmd to confirm the created disk groups.

```

ASMCMDB> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU      Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED    EXTERN    N           512           512      4096    4194304    409600    408848
0           408848           0           N      FCREDOLOG/
MOUNTED    EXTERN    N           512           512      4096    4194304    6553600    935664
0           935664           0           N      FCSLOBDATA/

```

MOUNTED	EXTERN	N	512	512	4096	4194304	204800	204464
0					Y	ISCSIOCR/VOTE/		
MOUNTED	EXTERN	N	512	512	4096	4194304	409600	408848
0					N	ISCSIREDOLOG/		
MOUNTED	EXTERN	N	512	512	4096	4194304	6553600	939036
0					N	ISCSISLOBDATA/		
MOUNTED	EXTERN	N	4096	4096	4096	4194304	409600	408848
0					N	NVMEREDOLOG/		
MOUNTED	EXTERN	N	4096	4096	4096	4194304	6553600	939176
0					N	NVMESLOBDATA/		

ASMCMD>

As oracle user, use dbca to create fcdB database and fcdB pluggable database. Refer to the procedure available in the iSCSI section for details.

Afterwards, you can use crsctl command to confirm that the created fcdB is running and then perform testing against the database, using the same procedures used for testing iSCSI protocol-based database.

### Confirm NVMe/FC related configurations in Linux host

1. Login to Linux host as root.
2. Start and check for NVMe/FC protocol support related module nvme\_fc is running along with other nvme related modules.

```
[root@fpsa-asa-linux-01 ~]# modprobe nvme
[root@fpsa-asa-linux-01 ~]# modprobe nvme-tcp

[root@fpsa-asa-linux-01 ~]# lsmod | grep nvme
nvme_tcp          57344  0
nvme              61440  0
nvme_fc          61440  1 fnic
nvme_fabrics      36864  2 nvme_tcp,nvme_fc
nvme_core        208896  598 nvme_tcp,nvme,nvme_fc,nvme_fabrics
nvme_common      24576  1 nvme_core
t10_pi           16384  2 sd_mod,nvme_core
```

3. Verify that in-kernel NVMe multipath is enabled.

```
[root@fpsa-asa-linux-01 ~]# cat /sys/module/nvme_core/parameters/multipath
Y
```

4. Check for installed nvme-cli package.

```
[root@fpsa-asa-linux-01 ~]# dnf list installed | grep nvme-cli
nvme-cli.x86_64          1.16-9.el8
@anaconda
```

5. Check for the host NQN in /etc/nvme/hostnqn file.

```
[root@fpsa-asa-linux-01 ~]# cat /etc/nvme/hostnqn
nqn.2014-08.org.nvmexpress:uuid:0000a0aa-0000-0100-aaa0-000000000004
```

6. Confirm dm-multipath configuration file /etc/multipath.conf excludes NVMe namespaces so the in-kernel multipathing is used for NVMe namespaces with enable\_foreign parameter set to NONE.

```
[root@fpsa-asa-linux-01 ~]# cat /etc/multipath.conf
defaults {
    find_multipaths yes
    user_friendly_names yes
    enable_foreign NONE
}

blacklist {
}
```

7. Restart multipathd to use the updated configuration if needed.

```
[root@fpsa-asa-linux-01 ~]# systemctl restart multipathd
```

8. Login to ONTAP storage as admin user.

## 9. Confirm NVMe/FC LIFs are up.

```
fpsa-a50-u0909:> net int show fc-nvme-lif*
(network interface show)
Vserver      Logical      Status      Network      Current      Current Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
svm1
      fc-nvme-lif-01a up/up 20:0b:d0:39:ea:c6:a7:94 fpsa-a50-u0909-01 1b true
      fc-nvme-lif-01b up/up 20:0c:d0:39:ea:c6:a7:94 fpsa-a50-u0909-01 1d true
      fc-nvme-lif-02a up/up 20:0d:d0:39:ea:c6:a7:94 fpsa-a50-u0909-02 1b true
      fc-nvme-lif-02b up/up 20:1c:d0:39:ea:c6:a7:94 fpsa-a50-u0909-02 1d true
4 entries were displayed.
```

## 10. Login to Linux host as admin user.

## 11. Show the discovered NVMe subsystem.

```
[admin@fpsa-asa-linux-01 ~]$ nvme list-subsys
nvme-subsys0 - NQN=nqn.1992-08.com.netapp:sn.57cd8838ea2911ef9608d039eac6a795:subsystem.FlexPod-
ASA-linux-cluster-nvme
\
+- nvme0 fc traddr=nn-0x2007d039eac6a794:pn-0x201cd039eac6a794 host_traddr=nn-
0x20000025b5150003:pn-0x20000025b509000f live
+- nvme1 fc traddr=nn-0x2007d039eac6a794:pn-0x200dd039eac6a794 host_traddr=nn-
0x20000025b5150003:pn-0x20000025b509000e live
+- nvme2 fc traddr=nn-0x2007d039eac6a794:pn-0x200cd039eac6a794 host_traddr=nn-
0x20000025b5150003:pn-0x20000025b509000f live
+- nvme3 fc traddr=nn-0x2007d039eac6a794:pn-0x200bd039eac6a794 host_traddr=nn-
0x20000025b5150003:pn-0x20000025b509000e live
```

## 12. Show the discovered NVMe devices.

```
[admin@fpsa-asa-linux-01 ~]$ sudo nvme netapp ontapdevices -o json
{
  "ONTAPdevices" : [
    {
      "Device" : "/dev/nvme0n1",
      "Vserver" : "svm1",
      "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_redolog_nvme_1/blocks",
      "NSID" : 1,
      "UUID" : "704f70ff-0d75-11f0-a707-d039eac6a795",
      "Size" : "214.75GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 52428800
    },
    {
      "Device" : "/dev/nvme0n10",
      "Vserver" : "svm1",
      "Namespace_Path" : "/vol/fpsa_asa_linux_01_oracle_slobdata_nvme_8/blocks",
      "NSID" : 10,
      "UUID" : "b58774fb-0d75-11f0-a707-d039eac6a795",
      "Size" : "858.99GB",
      "LBA_Data_Size" : 4096,
      "Namespace_Size" : 209715200
    },
    ...
  ]
}
```

## 13. Check for NVMe device persistent naming.

```
[admin@fpsa-asa-linux-01 ~]$ sudo ls -l /dev/nvmeredolog*
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeredolog1 -> nvme0n1
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeredolog2 -> nvme0n2
[admin@fpsa-asa-linux-01 ~]$ sudo ls -l /dev/nvmeslobdata*
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata1 -> nvme0n3
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata2 -> nvme0n4
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata3 -> nvme0n5
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata4 -> nvme0n6
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata5 -> nvme0n7
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata6 -> nvme0n8
lrwxrwxrwx 1 root root 7 Aug 18 16:37 /dev/nvmeslobdata7 -> nvme0n9
```



```
lrwxrwxrwx 1 root root 8 Aug 18 16:37 /dev/nvmeslobdata8 -> nvme0n10
```

## Start Oracle disk groups and database on NVMe/FC storage

Since we are accessing the same set of namespaces previously configured for NVMe/TCP protocol access but now with NVMe/FC, we can just mount the disk group and access the previously created database on them.

1. Login to Linux host as grid user.
2. Invoke asmcmd tool and list the configured disk groups.

```
[grid@fpsa-asa-linux-01 ~]$ asmcmd
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU  Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED  EXTERN  N      512      512      4096  4194304  204800  204468
0         204468      0      Y  ISCSIOCRVOTE/
MOUNTED  EXTERN  N      512      512      4096  4194304  409600  408848
0         408848      0      N  ISCSIREDOLOG/
MOUNTED  EXTERN  N      512      512      4096  4194304  6553600  939044
0         939044      0      N  ISCSISLOBDATA/
```

3. mount the two NVMe disk groups and then exit the asmcmd tool.

```
ASMCMD> mount NVMEREDOLOG
ASMCMD> mount NVMESLOBDATA
ASMCMD> lsdg
State      Type      Rebal  Sector  Logical_Sector  Block      AU  Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED  EXTERN  N      512      512      4096  4194304  204800  204468
0         204468      0      Y  ISCSIOCRVOTE/
MOUNTED  EXTERN  N      512      512      4096  4194304  409600  408848
0         408848      0      N  ISCSIREDOLOG/
MOUNTED  EXTERN  N      512      512      4096  4194304  6553600  939044
0         939044      0      N  ISCSISLOBDATA/
MOUNTED  EXTERN  N      4096     4096     4096  4194304  409600  408848
0         408848      0      N  NVMEREDOLOG/
MOUNTED  EXTERN  N      4096     4096     4096  4194304  6553600  967128
0         967128      0      N  NVMESLOBDATA/
```

4. Login to Linux host as oracle user.
5. Start the Oracle database nvmedb that was created on the NVMe storage.

```
[oracle@fpsa-asa-linux-01 ~]$ srvctl status database -db nvmedb
Instance NVMEDB1 is not running on node fpsa-asa-linux-01

[oracle@fpsa-asa-linux-01 ~]$ srvctl start database -db nvmedb

[oracle@fpsa-asa-linux-01 ~]$ srvctl status database -db nvmedb
Instance NVMEDB1 is running on node fpsa-asa-linux-01
```

**Note:** Now that the database is running on the NVMe storage using NVMe/FC protocol, access to the nvmedb database created previously can resume and you can follow the documented SLOB testing procedures to conduct testing.

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- FlexPod home page: <https://www.flexpod.com>
- FlexPod design guides: <https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>

- FlexPod datacenter with Oracle 21c RAC on Cisco UCS X-Series M7 and NetApp AFF900 with NVMe/FC: [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_oracle\\_xseries\\_m7.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_oracle_xseries_m7.html)
- FlexPod security hardening: [https://www.netapp.com/pdf.html?item=/media/99202-tr\\_4984\\_flexpod\\_security\\_hardening.pdf](https://www.netapp.com/pdf.html?item=/media/99202-tr_4984_flexpod_security_hardening.pdf)
- FlexPod datacenter zero trust framework: [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_zero\\_trust.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_zero_trust.html)
- FlexPod datacenter for Microsoft SQL Server 2022 and VMware vSphere 8.0: [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_sql2022\\_xseries.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_sql2022_xseries.html)
- Cisco UCS X-Series Direct data sheet: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-x-series-modular-system/ucs-x-series-direct-ds.html>
- Cisco UCS X-Series Direct Fabric Interconnect 9108 100G installation and service guide: [https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/x-series-direct/hw/s9108/install/ucs-x-series-direct-9108-100g.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/x-series-direct/hw/s9108/install/ucs-x-series-direct-9108-100g.html)
- Performance tuning for Cisco UCS M8 platforms with AMD EPYC 4<sup>th</sup> Gen and 5<sup>th</sup> Gen processors: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/ucs-c245-m8-rack-ser-4th-gen-amd-epyc-pro-wp.html>
- Cisco Intersight: <https://intersight.com>
- NetApp ASA r2 documentation: <https://docs.netapp.com/us-en/asa-r2/index.html>
- Installation and setup workflow for ASA r2 storage systems: <https://docs.netapp.com/us-en/asa-r2/install-setup/install-setup-workflow.html>
- ONTAP tools for VMware vSphere documentation: <https://docs.netapp.com/us-en/ontap-tools-vmware-vsphere-10/index.html>
- Oracle 21c install and upgrade: <https://docs.oracle.com/en/database/oracle/oracle-database/21/install-and-upgrade.html>
- Installing SQL server from the installation wizard (Setup): <https://learn.microsoft.com/en-us/sql/database-engine/install-windows/install-sql-server-from-the-installation-wizard-setup?view=sql-server-ver16>
- Microsoft SQL server download: <https://www.microsoft.com/en-us/sql-server/sql-server-downloads>
- Microsoft SQL server manager download: <https://learn.microsoft.com/en-us/ssms/download-sql-server-management-studio-ssms>
- Load testing SQL Server using HammerDB: <https://cloud.google.com/compute/docs/tutorials/load-testing-sql-server-hammerdb>
- HammerDB download: <https://www.hammerdb.com/download.html>
- SLOB Resources: <https://kevinclosson.net/slob/>
- SLOB 2 documentation: [https://kevinclosson.net/wp-content/uploads/2013/12/slob2\\_readme-2012-05-04.pdf](https://kevinclosson.net/wp-content/uploads/2013/12/slob2_readme-2012-05-04.pdf)
- SLOB deployment – a picture tutorial: <https://kevinclosson.net/2014/08/04/slob-deployment-a-picture-tutorial/>
- Broadcom KB 369525 - Host TPM Attestation Alarm present in the vSphere UI: <https://knowledge.broadcom.com/external/article?articleNumber=369525>
- Broadcom KB 316512 - The new host TPM endorsement key doesn't match the one stored in the DB: <https://knowledge.broadcom.com/external/article?articleNumber=316512>
- NetApp product documentation: <https://www.netapp.com/support-and-training/documentation/>
- NetApp ASA datasheet: <https://www.netapp.com/media/85736-ds-4254-asa.pdf>
- NetApp support site: <https://support.netapp.com>

- NetApp hardware universe: <https://hwu.netapp.com>
- NetApp interoperability matrix tool: <http://support.netapp.com/matrix>
- Cisco UCS hardware and software compatibility list: <https://ucshcltool.cloudapps.cisco.com/public/>
- Broadcom compatibility guide: <https://compatibilityguide.broadcom.com/>

## Acknowledgement

The author extends sincere appreciations for the collaborations and support of the following teams: Cisco and NetApp TME teams, Cisco and NetApp Interoperability teams, and NetApp engineering teams. Special thanks go to Bobby Oommen (NetApp), Roney Daniel (NetApp), John George (Cisco), Reese Lloyd (NetApp), Kamini Singh (NetApp), Abhinav Singh (NetApp), Suzie Morin (NetApp), and others who provided help, support, guidance, feedback, and assistance with the solution. Thank you!

## Version history

As an option, use the NetApp Table style to create a Version History table. Do not add a table number or caption.

Version	Date	Document version history
Version 1.0	May, 2025	Initial release.
Version 1.1	October, 2025	Added FC-based SAN configuration information in Appendices.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright information**

Copyright © 2025 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

### **Trademark information**

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