

FlashStack Virtual Server Infrastructure with Cisco UCS 4.2(1) in UCS Managed Mode, VMware vSphere 7.0 U2, and Purity//FA 6.1

Deployment Guide for FlashStack with VMware vSphere 7.0 U2, Cisco UCS M6 Servers with 3rd Generation Intel Xeon Scalable Processors, and Pure Storage FlashArray//X R3 Series

Published: December 2021



In partnership with:





About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_P3)

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2021 Cisco Systems, Inc. All rights reserved.

Contents

Executive Summary	6
Solution Overview.....	7
Deployment Hardware and Software	9
Network Switch Configuration.....	26
Storage Configuration.....	35
Cisco UCS Configuration	44
SAN Switch Configuration.....	133
FlashStack Cisco MDS Switch Configuration.....	137
Storage Configuration - Boot LUNs.....	143
VMware vSphere 7.0 U2 Setup.....	154
FlashStack Management Tools Setup	200
Appendix	274
FlashStack Automated Deployment with Ansible	330
About the Authors.....	348
Feedback.....	349

Executive Summary

Cisco Validated Designs (CVDs) consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.

This document details the design in the FlashStack Virtual Server Infrastructure Design Guide for VMware vSphere 7.0 U2, which describes a validated Converged Infrastructure (CI) jointly developed by Cisco and Pure Storage. The solution covers the deployment of a predesigned, best-practice data center architecture with VMware vSphere built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® 9000 family of switches, Cisco MDS 9000 family of Fibre Channel switches and Pure Storage FlashArray//X R3 all flash array supporting either iSCSI or Fibre Channel storage access.

In addition to that, this FlashStack solution is also delivered as Infrastructure as Code (IaC) to eliminate error-prone manual tasks, allowing quicker and more consistent solution deployments. Cisco Inter-sight cloud platform delivers monitoring, orchestration, workload optimization and lifecycle management capabilities for the FlashStack solution.

When deployed, the architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Server Infrastructure (VSI).

Solution Overview

Introduction

Currently, the industry trend is for pre-engineered solutions which standardize the data center infrastructure, offering the business operational efficiencies, agility, and scale to address cloud, bi-modal IT, and their business. Their challenge is complexity, diverse application support, efficiency, and risk; all these are met by FlashStack with:

- Reduced complexity, automatable infrastructure and easily deployed resources
- Robust components capable of supporting high performance and high bandwidth virtualized applications
- Efficiency through optimization of network bandwidth and in-line storage compression with deduplication
- Risk reduction at each level of the design with resiliency built into each touch point
- Cloud based monitoring, management, and support of your physical and virtual infrastructure

Cisco and Pure Storage have partnered to deliver this Cisco Validated Design, which uses best of breed storage, server, and network components to serve as the foundation for virtualized workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

In this document we will describe a reference architecture detailing a Virtual Server Infrastructure composed of Cisco Nexus switches, Cisco UCS Compute, Cisco MDS Multilayer Fabric Switches, and a Pure Storage FlashArray//X50 R3 delivering VMware vSphere 7.0 U2 hypervisor environment.

Audience

The intended audience of this document includes but is not limited to data scientists, IT architects, sales engineers, field consultants, professional services, IT managers, partner engineering, DevOps, and Site Reliability Engineers (SREs) and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step-by-step configuration and implementation guide along with automated deployment guidance for the FlashStack, implemented with either FC or iSCSI, centered around the Cisco UCS 6454 Fabric Interconnect and the Pure Storage FlashArray//X50 R3, delivering a Virtual Server Infrastructure on Cisco UCS B200 M6 Blade Servers running VMware vSphere 7.0 U2.

What's New in this Release?

This version of the FlashStack VSI Design introduces the Cisco UCS M6 Servers featuring the 3rd Gen Intel Xeon Scalable processors. The design incorporates options for 25 iSCSI as well as 32Gb Fibre Channel protocols, both delivered with new design options and features. Highlights for this design include:

-
- Support for Cisco UCS B200 M6 blade servers with 3rd Gen Intel Xeon Scalable Family processors and 3200 MHz memory
 - Support for Intel Optane Persistent Memory (PMem)
 - Support for the Cisco UCS Manager 4.2
 - Support for Pure Storage FlashArray//X50 R3 with Purity version 6.1.6
 - Support for NVMe over Fibre Channel (FC-NVMe) Datastores
 - Support for VMware vSphere 7.0 U2
 - Fully automated solution deployment covering FlashStack infrastructure and vSphere virtualization
 - Support for Cisco Intersight Software as a Service (SaaS) Management
 - Support for Cisco Data Center Network Manager (DCNM)-SAN Version 11.5(1)
 - Unified Extensible Firmware Interface (UEFI) Secure Boot of VMware ESXi 7.0 Update 2
 - Trusted Platform Module (TPM) 2.0 Attestation of UEFI Secure Boot of VMware ESXi 7.0 Update 2

Deployment Hardware and Software

Architecture

FlashStack with Cisco UCS M6 servers and vSphere 7.0 U2 delivers a Virtual Server Infrastructure that is redundant, using the best practices of Cisco and Pure Storage. The solution includes VMware vSphere 7.0 U2 hypervisor installed on the Cisco UCS M6 compute nodes configured for stateless compute design using boot from SAN. Pure Storage FlashArray//X50 R3 provides the storage infrastructure required for setting up the VMware environment. Cisco UCS manager is utilized to configure and manage the UCS infrastructure with Cisco Intersight providing lifecycle management capabilities. The solution requirements and design details are described in this section.


Physical Topology

FlashStack with Cisco UCS M6 servers supports both IP and Fibre Channel (FC) based storage access design. For the IP based solution, iSCSI configuration on Cisco UCS and Pure Storage FlashArray is utilized to setup storage access including boot from SAN for the compute node. For the FC designs, Pure Storage FlashArray and Cisco UCS are connected through Cisco MDS 9132T switches and storage access utilizes the FC network.

IP-based Storage Access

The physical topology for the IP based FlashStack is shown in [Figure 1](#).

Figure 1. FlashStack - physical topology for IP connectivity



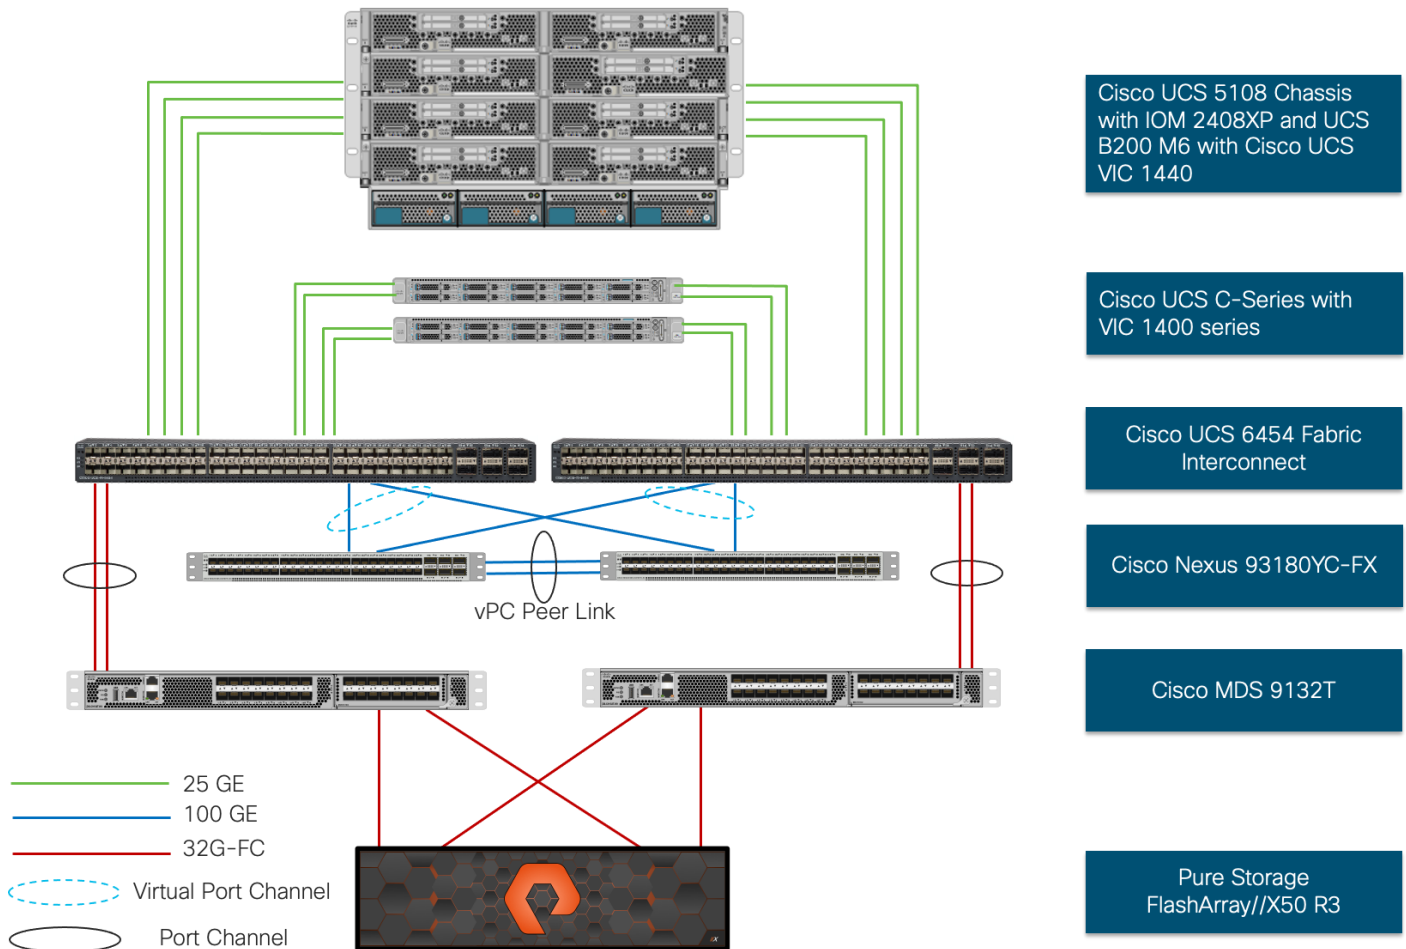
To validate the IP based storage access in a FlashStack configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Modular Chassis connects to fabric interconnects using the Cisco 2408XP IOM within modules hosted within the chassis, where four 25 Gigabit Ethernet ports are used on each IOM to connect to appropriate FI. Depending on customer workload requirements, for additional bandwidth all eight ports can be used to connect IOM to FI.
- Cisco UCS B200 M6 servers contain fourth-generation Cisco 1440 virtual interface cards.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX Switches in a virtual port channel (vPC) configuration.
- The Pure Storage FlashArray//50 R3 connects to the Cisco Nexus 93180YC-FX Switches using four 25 GE ports.
- VMware 7.0 U2 ESXi software is installed on Cisco UCS B200 M6 servers to validate the infrastructure.

FC-based Storage Access

[Figure 2](#) illustrates the FlashStack physical topology for FC connectivity.

Figure 2. FlashStack- physical topology for FC connectivity



To validate the FC based storage access in a FlashStack configuration, the components are set up as follows:

- Cisco UCS 6454 Fabric Interconnects provide the chassis and network connectivity.
- The Cisco UCS 5108 Modular Chassis connects to fabric interconnects using the Cisco 2408XP IOM within modules hosted within the chassis, where four 25 Gigabit Ethernet ports are used on each IOM to connect to the appropriate FI.
- Cisco UCS B200 M6 servers contain fourth-generation Cisco 1440 virtual interface cards.
- Cisco Nexus 93180YC-FX Switches in Cisco NX-OS mode provide the switching fabric.
- Cisco UCS 6454 Fabric Interconnect 100 Gigabit Ethernet uplink ports connect to Cisco Nexus 93180YC-FX3 Switches in a vPC configuration.

- Cisco UCS 6454 Fabric Interconnects are connected to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections configured as a port channel for SAN connectivity.
- The Pure Storage FlashArray//X 50 R3 connects to the Cisco MDS 9132T switches using 32-Gbps Fibre Channel connections for SAN connectivity.
- VMware 7.0 U2 ESXi software is installed on Cisco UCS B200 M6 servers to validate the infrastructure.

Software Revisions

[Table 1](#) lists the software revisions for this solution. The software versions for hardware and virtual components used in this solution. Each of these versions have been certified within interoperability matrixes supported by Cisco, Pure Storage, and VMware. For more current supported version information, consult the following sources:

- Cisco UCS Hardware and Software Interoperability Tool: <http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>
- Pure Storage Interoperability (note, this interoperability list will require a support login from Pure): https://support.purestorage.com/FlashArray/Getting_Started/Compatibility_Matrix
- Pure Storage FlashStack Compatibility Matrix (note, this interoperability list will require a support login from Pure): https://support.purestorage.com/FlashStack/Product_Information/FlashStack_Compatibility_Matrix
- VMware Compatibility Guide: <http://www.vmware.com/resources/compatibility/search.php>
- Additionally, it is also strongly suggested to align FlashStack deployments with the recommended release for the Cisco Nexus 9000 switches used in the architecture:
- Nexus: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html
- MDS: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html

Table 1. Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6454, Cisco UCS M6 Servers with 3rd Generation Intel Xeon Scalable Processors	4.2(1f)	Includes the Cisco UCS Manager and Cisco UCS VIC 1440
Network	Cisco Nexus 93180YC-FX NX-OS	9.3(7a)	Software version

Layer	Device	Image	Comments
	Cisco MDS 9132T	8.5(1a)	Software version
Storage	Pure Storage FlashArray//X50 R3	6.1.6	Software version
Software	Cisco UCS Manager	4.2(1f)	Software version
	Cisco Data Center Network Manager (SAN)	11.5(1)	Software version
	VMware vSphere	7.0 U2	Software version
	VMware ESXi nfnic FC Driver	5.0.0.15	Software version
	VMware ESXi nenic Ethernet Driver	1.0.35.0	Software version
	Pure Storage Plugin	4.5.0	Software version
	VASA Provider	3.5	Software version
Management	Cisco Intersight	N/A	

Configuration Guidelines

This document details the step-by-step configuration of a fully redundant and highly available Virtual Server Infrastructure built on Cisco and Pure Storage components. References are made to which component is being configured with each step, either 01 or 02 or A and B. For example, controller-1 and controller-2 are used to identify the two controllers within the Pure Storage FlashArray//X that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02 to represent Fibre Channel booted infrastructure and production hosts deployed to the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in each step, <<text>> appears as part of the command structure. The following is an example of a configuration step for both Cisco Nexus switches:

```
BB08-93180YC-FX-A (config)# ntp server <<var_oob_ntp>> use-vrf management
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. [Table 2](#) lists the VLANs necessary for deployment as outlined in this guide, and [Table 3](#) lists the external dependencies necessary for deployment as outlined in this guide.

Table 2. Necessary VLANs

VLAN ID	Name	Usage
2	Native-VLAN	Use VLAN 2 as Native VLAN instead of default VLAN (1)
15	OOB-MGMT-VLAN	Out-of-Band Management VLAN to connect the management ports for various devices
115	IB-MGMT-VLAN	In Band Management VLAN utilized for all in-band management connectivity for example, ESXi hosts, VM management, and so on.
1101	VM-Traffic-VLAN	VM data traffic VLAN.
1130	vMotion-VLAN	VMware vMotion traffic.
901*	iSCSI-A-VLAN	iSCSI-A path for supporting boot-from-san for both Cisco UCS B-Series and Cisco UCS C-Series servers
902*	iSCSI-B-VLAN	iSCSI-B path for supporting boot-from-san for both Cisco UCS B-Series and Cisco UCS C-Series servers

[Table 3](#) lists the VMs necessary for deployment as outlined in this document.

Table 3. Virtual Machines

Virtual Machine Description	Host Name	IP Address
vCenter Server		
Cisco Data Center Network Manager (DCNM)		
Cisco Intersight Assist		

Table 4. Configuration Variables

Variable Name	Variable Description	Customer Variable Name
<<var_nexus_A_hostname>>	Cisco Nexus switch A Host name (Example: BB08-91380YX-FX-A)	
<<var_nexus_A_mgmt_ip>>	Out-of-band management IP for Cisco Nexus switch A (Example: 10.1.164.61)	
<<var_oob_mgmt_mask>>	Out-of-band network mask (Example: 255.255.255.0)	
<<var_oob_gateway>>	Out-of-band network gateway (Example: 10.1.164.254)	
<<var_oob_ntp>>	Out-of-band management network NTP Server (Example: 10.1.164.254)	

Variable Name	Variable Description	Customer Variable Name
<<var_nexus_B_hostname>>	Cisco Nexus switch B Host name (Example: BB08-91380YX-FX-B)	
<<var_nexus_B_mgmt_ip>>	Out-of-band management IP for Nexus switch B (Example: 10.1.164.62)	
<<var_flasharray_hostname>>	Array Hostname set during setup (Example: BB08-FlashArrayR3)	
<<var_flasharray_vip>>	Virtual IP that will answer for the active management controller (Example: 10.2.164.100)	
<<var_contoller-1_mgmt_ip>>	Out-of-band management IP for FlashArray controller-1 (Example:10.2.164.101)	
<<var_contoller-1_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)	
<<var_contoller-1_mgmt_gateway>>	Out-of-band management network default gateway (Example: 10.2.164.254)	
<<var_contoller-2_mgmt_ip>>	Out-of-band management IP for FlashArray controller-2 (Example:10.2.164.102)	
<<var_contoller-2_mgmt_mask>>	Out-of-band management network netmask (Example: 255.255.255.0)	
<<var_ contoller-2_mgmt_gateway>>	Out-of-band management network default gateway (Example: 10.2.165.254)	
<<var_password>>	Administrative password (Example: FI@shSt4x)	
<<var_dns_domain_name>>	DNS domain name (Example: flashstack.cisco.com)	
<<var_nameserver_ip>>	DNS server IP(s) (Example: 10.1.164.125)	
<<var_smtp_ip>>	Email Relay Server IP Address or FQDN (Example: smtp.flashstack.cisco.com)	
<<var_smtp_domain_name>>	Email Domain Name (Example: flashstack.cisco.com)	
<<var_timezone>>	FlashStack time zone (Example: America/New_York)	
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID (Example: 15)	
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID (Example: 215)	

Variable Name	Variable Description	Customer Variable Name
<<var_ib_mgmt_vlan_netmask_length>>	Length of IB-MGMT-VLAN Netmask (Example: /24)	
<<var_ib_gateway_ip>>	In-band management network VLAN ID (Example: 10.2.164.254)	
<<var_vmotion_vlan_id>>	vMotion network VLAN ID (Example: 1130)	
<<var_vmotion_vlan_netmask_length>>	Length of vMotion VLAN Netmask (Example: /24)	
<<var_native_vlan_id>>	Native network VLAN ID (Example: 2)	
<<var_app_vlan_id>>	Example Application network VLAN ID (Example: 1101)	
<<var_snmp_contact>>	Administrator e-mail address (Example: admin@flashstack.cisco.com)	
<<var_snmp_location>>	Cluster location string (Example: RTP9-BB08)	
<<var_mds_A_mgmt_ip>>	Cisco MDS Management IP address (Example: 10.1.164.63)	
<<var_mds_A_hostname>>	Cisco MDS hostname (Example: BB08-MDS-9132T-A)	
<<var_mds_B_mgmt_ip>>	Cisco MDS Management IP address (Example: 10.1.164.64)	
<<var_mds_B_hostname>>	Cisco MDS hostname (Example: BB08-MDS-9132T-B)	
<<var_vsan_a_id>>	VSAN used for the A Fabric between the FlashArray/MDS/FI (Example: 100)	
<<var_vsan_b_id>>	VSAN used for the B Fabric between the FlashArray/MDS/FI (Example: 200)	
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name (Example: BB08-FI-6454)	
<<var_ucs_a_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address (Example: 10.1.164.51)	
<<var_ucs_mgmt_vip>>	Cisco UCS fabric interconnect (FI) Cluster out-of-band management IP address (Example: 10.1.164.50)	
<<var_ucs_b_mgmt_ip>>	Cisco UCS fabric interconnect (FI) Cluster out-of-band management IP address (Example: 10.1.164.52)	

Variable Name	Variable Description	Customer Variable Name
<<var_vm_host_fc_01_ip>>	VMware ESXi host 01 in-band management IP (Example:10.1.164.111)	
<<var_vm_host_fc_vmotion_01_ip>>	VMware ESXi host 01 vMotion IP (Example: 192.168.130.101)	
<<var_vm_host_fc_02_ip>>	VMware ESXi host 02 in-band management IP (Example:10.1.164.112)	
<<var_vm_host_fc_vmotion_02_ip>>	VMware ESXi host 02 vMotion IP (Example: 192.168.130.102)	
<<var_vmotion_subnet_mask>>	vMotion subnet mask (Example: 255.255.255.0)	
<<var_vcenter_server_ip>>	IP address of the vCenter Server (Example: 10.1.164.110)	

Physical Infrastructure

FlashStack Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlashStack environment. To simplify cabling requirements, a cabling diagram was used.

The cabling diagram in this section contains details for the prescribed and supported configuration of the Pure FlashArray//X R3 running Purity 6.1.6.


This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps.



Be sure to use the cabling directions in this section as a guide.

[Figure 3](#) details the cable connections used in the validation lab for FlashStack topology based on the Cisco UCS 6454 fabric interconnect. Four 32Gb uplinks connect as port-channels to each Cisco UCS Fabric Interconnect from the MDS switches, and a total of eight 32Gb links connect the MDS switches to the Pure FlashArray//X R3 controllers, four of these have been used for scsi-fc and the other four to support nvme-fc. Also, 25Gb links connect the Cisco UCS Fabric Interconnects to the Cisco Nexus Switches and the Pure FlashArray//X R3 controllers to the Cisco Nexus Switches. Additional 1Gb management connections will be needed for an out-of-band network switch that sits apart from the FlashStack infrastructure. Each Cisco UCS fabric interconnect and Cisco Nexus switch is connected to the out-of-band network switch, and each FlashArray controller has a connection to the out-of-band network switch. Layer 3 network connectivity is required between the Out-of-Band (OOB) and In-Band (IB) Management Subnets.

Figure 3. FlashStack Cabling with Cisco UCS 6454 Fabric Interconnect



Although this diagram includes the Cisco C4200 chassis and Cisco UCS C-220 servers, this document describes the configuration of only the Cisco UCS 5108 chassis with the Cisco UCS B-Series M6 servers with Intel 3rd Generation Intel Xeon Scalable Processors. For configuration of Cisco UCS AMD-based servers, please see the [FlashStack Datacenter with VMware vSphere 7.0 and Pure FlashArray//X R3](#) CVD.



Cisco UCS Fabric Interconnect's to the Cisco Nexus 93180YC-FX switches connectivity can be done using the 100Gbe or 25Gbe ports based on the bandwidth requirements, this document includes the usage of 25Gbe ports with aggregate bandwidth of 200Gbe per port channel from the Cisco UCS FI to the Cisco Nexus switches.



* iSCSI connectivity is not required if iSCSI storage access is not being implemented.

Table 5. Cisco Nexus 93180YC-FX-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco Nexus 93180YC-FX-A	Eth 1/1	25Gbe	Cisco UCS 6454-A	Eth 1/33
	Eth 1/2	25Gbe	Cisco UCS 6454-A	Eth 1/34
	Eth 1/3	25Gbe	Cisco UCS 6454-A	Eth 1/35
	Eth 1/4	25Gbe	Cisco UCS 6454-A	Eth 1/36
	Eth 1/5	25Gbe	Cisco UCS 6454-B	Eth 1/33
	Eth 1/6	25Gbe	Cisco UCS 6454-B	Eth 1/34
	Eth 1/7	25Gbe	Cisco UCS 6454-B	Eth 1/35
	Eth 1/8	25Gbe	Cisco UCS 6454-B	Eth 1/36
	Eth 1/53	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/53
	Eth 1/54	100Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/54
	Eth 1/9	10Gbe or 25 Gbe	Upstream Network Switch	Any
	Mgmt0	Gbe	Gbe Management Switch	Any

Local Device	Local Port	Connection	Remote Device	Remote port
	Eth 1/37 *	25Gbe	FlashArray//X50 R3 Controller 1	CT0.ETH4
	Eth 1/38 *	25Gbe	FlashArray//X50 R3 Controller 2	CT1.ETH4

Table 6. Cisco Nexus 93180YC-FX-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco Nexus 93180YC-FX-B	Eth 1/1	25Gbe	Cisco UCS 6454-A	Eth 1/37
	Eth 1/2	25Gbe	Cisco UCS 6454-A	Eth 1/38
	Eth 1/3	25Gbe	Cisco UCS 6454-A	Eth 1/39
	Eth 1/4	25Gbe	Cisco UCS 6454-A	Eth 1/40
	Eth 1/5	25Gbe	Cisco UCS 6454-B	Eth 1/37
	Eth 1/6	25Gbe	Cisco UCS 6454-B	Eth 1/38
	Eth 1/7	25Gbe	Cisco UCS 6454-B	Eth 1/39
	Eth 1/8	25Gbe	Cisco UCS 6454-B	Eth 1/40
	Eth 1/9	10Gbe or 25 Gbe	Upstream Network Switch	Any
	Mgmt0	Gbe	Gbe Management Switch	Any
	Eth 1/37 *	25Gbe	FlashArray//X50 R3 Controller 1	CT0.ETH5
	Eth 1/38 *	25Gbe	FlashArray//X50 R3 Controller 2	CT1.ETH5

Table 7. Cisco UCS-6545-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
--------------	------------	------------	---------------	-------------

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco UCS-6454-A	Eth 1/33	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/1
	Eth 1/34	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/2
	Eth 1/35	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/3
	Eth 1/36	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/4
	Eth 1/37	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/1
	Eth 1/38	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/2
	Eth 1/39	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/3
	Eth 1/40	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/4
	Eth 1/17	25Gbe	Cisco UCS Chassis 1 2408 FEX A	IOM 1/1
	Eth 1/18	25Gbe	Cisco UCS Chassis 1 2408 FEX A	IOM 1/2
	Eth 1/19	25Gbe	Cisco UCS Chassis 1 2408 FEX A	IOM 1/3
	Eth 1/20	25Gbe	Cisco UCS Chassis 1 2408 FEX A	IOM 1/4
	FC1/1	32G FC	Cisco MDS 9132T-A	FC1/1
	FC1/2	32G FC	Cisco MDS 9132T-A	FC1/2
	FC1/3	32G FC	Cisco MDS 9132T-A	FC1/3
	FC1/4	32G FC	Cisco MDS 9132T-A	FC1/4
	Mgmt0	Gbe	Gbe Management Switch	Any

Table 8. Cisco UCS-6545-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
--------------	------------	------------	---------------	-------------

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco UCS-6454-B	Eth 1/33	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/5
	Eth 1/34	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/6
	Eth 1/35	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/7
	Eth 1/36	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/8
	Eth 1/37	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/5
	Eth 1/38	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/6
	Eth 1/39	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/7
	Eth 1/40	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/8
	Eth 1/17	25Gbe	Cisco UCS Chassis 1 2408 FEX B	IOM 1/1
	Eth 1/18	25Gbe	Cisco UCS Chassis 1 2408 FEX B	IOM 1/2
	Eth 1/19	25Gbe	Cisco UCS Chassis 1 2408 FEX B	IOM 1/3
	Eth 1/20	25Gbe	Cisco UCS Chassis 1 2408 FEX B	IOM 1/4
	FC1/1	32G FC	Cisco MDS 9132T-B	FC1/1
	FC1/2	32G FC	Cisco MDS 9132T-B	FC1/2
	FC1/3	32G FC	Cisco MDS 9132T-B	FC1/3
	FC1/4	32G FC	Cisco MDS 9132T-B	FC1/4
	Mgmt0	Gbe	Gbe Management Switch	Any

Table 9. Cisco MDS-9132T-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco MDS-9132T-A	FC1/5	32Gb FC	Cisco UCS 6454-A	FC1/1
	FC1/6	32Gb FC	Cisco UCS 6454-A	FC1/2

Local Device	Local Port	Connection	Remote Device	Remote port
	FC 1/7	32Gb FC	Cisco UCS 6454-A	FC1/3
	FC 1/8	32Gb FC	Cisco UCS 6454-A	FC1/4
	FC1/1	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC0 (scsi-fc)
	FC1/2	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC0 (scsi-fc)
	FC1/3	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC1 (nvme-fc)
	FC1/4	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC1 (nvme-fc)
	Mgmt0	Gbe	Gbe Management Switch	Any



This design uses SCSI-FCP for boot and datastore storage access and Port numbers 0 and 2 on each Pure FlashArray Controller have been used for the fibre channel connectivity, the ports 1 and 3 are used for FC-NVMe datastore access. All the four ports can be used for SCSI-FCP or FC-NVMe as needed but each port can only function as an SCSI-FCP or FC-NVMe port.

Table 10. Cisco MDS-9132T-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
Cisco MDS-9132T-B	FC1/5	32Gb FC	Cisco UCS 6454-B	FC1/1
	FC1/6	32Gb FC	Cisco UCS 6454-B	FC1/2
	FC 1/7	32Gb FC	Cisco UCS 6454-B	FC1/3
	FC 1/8	32Gb FC	Cisco UCS 6454-B	FC1/4
	FC1/1	32Gb FC	FlashArray//X50 R3 Controller 0	CT0.FC2 (scsi-fc)
	FC1/2	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC2 (scsi-fc)
	FC1/3	32Gb FC	FlashArray//X50 R3 Controller	CT0.FC3 (nvme-fc)

Local Device	Local Port	Connection	Remote Device	Remote port
			0	
	FC1/4	32Gb FC	FlashArray//X50 R3 Controller 1	CT1.FC3 (nvme-fc)
	Mgmt0	Gbe	Gbe Management Switch	Any



This design uses SCSI-FCP for boot and datastore storage access and Port numbers 0 and 2 on each Pure FlashArray Controller have been used for the fibre channel connectivity, the ports 1 and 3 are used for FC-NVMe datastore access. All the four ports can be used for SCSI-FCP or FC-NVMe as needed but each port can only function as an SCSI-FCP or FC-NVMe port.

Table 11. Pure Storage FlashArray//X50 R3 Controller 1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
FlashArray//X50 R3 Controller 1	CT0.FC0 (scsi-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/1
	CT0.FC2 (scsi-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/1
	CT0.FC1 (nvme-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/3
	CT0.FC3 (nvme-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/3
	CT0.ETH4 *	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/37
	CT0.ETH5 *	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/37



* Required only if iSCSI storage access is implemented.

Table 12. Pure Storage FlashArray//X50 R3 Controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote port
FlashArray//X50 R3 Controller 2	CT1.FC0 (scsi-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/2
	CT1.FC2 (scsi-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/2

Local Device	Local Port	Connection	Remote Device	Remote port
	CT1.FC1 (nvme-fc)	32Gb FC	Cisco MDS 9132T-A	FC 1/4
	CT1.FC3 (nvme-fc)	32Gb FC	Cisco MDS 9132T-B	FC 1/4
	CT1.ETH4 *	25Gbe	Cisco Nexus 93180YC-FX-A	Eth 1/38
	CT1.ETH5 *	25Gbe	Cisco Nexus 93180YC-FX-B	Eth 1/38




* Required only if iSCSI storage access is implemented.

Network Switch Configuration

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 93180YC-FX switches running NX-OS 9.3(7a). Configuring on a differing model of Cisco Nexus 9000 series switches should be comparable but may differ slightly with model and changes in NX-OS release. The Cisco Nexus 93180YC-FX switch and the NX-OS 9.3(7a) release were used in validating this FlashStack solution, so the steps will reflect this model and release.

Figure 4. Network Configuration workflow



Physical Connectivity

Physical cabling should be completed by following the diagram and table references in section [FlashStack Cabling](#).

FlashStack Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus 93180YC-FX switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 9000 9.3(7a), the Cisco suggested Nexus switch release at the time of this validation.



The following procedure includes the setup of NTP distribution on both the mgmt0 port and the in-band management VLAN. The interface-vlan feature and ntp commands are used to set this up. This procedure also assumes that the default VRF is used to route the in-band management VLAN.

Cisco Nexus A

To set up the initial configuration for the Cisco Nexus A switch on <nexus-A-hostname>, follow these steps:

1. Configure the switch.



On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)

----- System Admin Account Setup -----

Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-A-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-A-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-A-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-A-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: n
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

Cisco Nexus B

To set up the initial configuration for the Cisco Nexus B switch on <nexus-B-hostname>, follow these steps:

1. Configure the switch.



On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning.

```
Abort Power On Auto Provisioning [yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning] (yes/skip/no) [no]: yes
Disabling POAP.....Disabling POAP
poap: Rolling back, please wait... (This may take 5-15 minutes)
```

```
---- System Admin Account Setup ----
```

```
Do you want to enforce secure password standard (yes/no) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>
Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]: Enter
Configure read-write SNMP community string (yes/no) [n]: Enter
Enter the switch name: <nexus-B-hostname>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter
Mgmt0 IPv4 address: <nexus-B-mgmt0-ip>
Mgmt0 IPv4 netmask: <nexus-B-mgmt0-netmask>
Configure the default gateway? (yes/no) [y]: Enter
IPv4 address of the default gateway: <nexus-B-mgmt0-gw>
Configure advanced IP options? (yes/no) [n]: Enter
Enable the telnet service? (yes/no) [n]: Enter
Enable the ssh service? (yes/no) [y]: Enter
Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter
Number of rsa key bits <1024-2048> [1024]: Enter
Configure the ntp server? (yes/no) [n]: Enter
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Enter basic FC configurations (yes/no) [n]: Enter
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter
```

2. Review the configuration summary before enabling the configuration.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

FlashStack Cisco Nexus Switch Configuration

Enable Features

Cisco Nexus A and Cisco Nexus B

To enable the appropriate features on the Cisco Nexus switches, follow these steps:

1. Log in as admin.
2. Run the following commands:

```
config t
feature udld
feature interface-vlan
feature lacp
feature vpc
feature lldp
feature nxapi
```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

To set global configurations, follow this step on both switches:

3. Run the following commands to set global configurations:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
system default switchport
system default switchport shutdown
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
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>
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```



It is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3\(x\)](#). Sample clock commands for the United States Eastern timezone are:

```
clock timezone EST -5 0
```

```
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

Create VLANs

Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
vlan <oob-mgmt-vlan-id>
name OOB-MGMT
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-Vlan
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus A

1. From the global configuration mode, run the following commands:

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

Cisco Nexus B

1. From the global configuration mode, run the following commands:

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

Add Individual Port Descriptions for Troubleshooting and Enable UDLD for Cisco UCS Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, follow these steps:



In this step and in the following sections, configure the Cisco UCS 6454 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

1. From the global configuration mode, run the following commands:

```
interface Eth1/1
description <ucs-clustername>-A:1/33
udld enable
interface Eth1/2
description <ucs-clustername>-A:1/34
udld enable
interface Eth1/3
description <ucs-clustername>-A:1/35
udld enable
interface Eth1/4
description <ucs-clustername>-A:1/36
udld enable
interface Eth1/5
description <ucs-clustername>-B:1/33
udld enable
interface Eth1/6
description <ucs-clustername>-B:1/34
udld enable
interface Eth1/7
description <ucs-clustername>-B:1/35
udld enable
interface Eth1/8
description <ucs-clustername>-B:1/36
udld enable
```



For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected. If you have fibre optic connections, do not enter the `udld enable` command.

```
interface Ethernet1/53
description Peer Link <<nexus-B-hostname>>:Eth1/53
interface Ethernet1/54
description Peer Link <<nexus-B-hostname>>:Eth1/54
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B and to enable aggressive UDLD on copper interfaces connected to Cisco UCS systems, follow this step:

1. From the global configuration mode, run the following commands:

```
interface Eth1/1
description <ucs-clustertype>-A:1/37
udld enable
interface Eth1/2
description <ucs-clustertype>-A:1/38
udld enable
interface Eth1/3
description <ucs-clustertype>-A:1/39
udld enable
interface Eth1/4
description <ucs-clustertype>-A:1/40
udld enable
interface Eth1/5
description <ucs-clustertype>-B:1/37
udld enable
interface Eth1/6
description <ucs-clustertype>-B:1/38
udld enable
interface Eth1/7
description <ucs-clustertype>-B:1/39
udld enable
interface Eth1/8
description <ucs-clustertype>-B:1/40
udld enable
```



For fibre optic connections to Cisco UCS systems (AOC or SFP-based), entering `udld enable` will result in a message stating that this command is not applicable to fiber ports. This message is expected.

```
interface Ethernet1/53
description Peer Link <<nexus-A-hostname>>:Eth1/53
interface Ethernet1/54
description Peer Link <<nexus-A-hostname>>:Eth1/54
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

To create the necessary port channels between devices, follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/53-54
channel-group 10 mode active
no shutdown
interface Po121
description <ucs-clustername>-A
interface Eth1/1-4
channel-group 121 mode active
no shutdown
interface Po123
description <ucs-clustername>-B
interface Eth1/5-8
channel-group 123 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type network
speed 100000
duplex full
state enabled

interface Po121
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled

interface Po123
switchport mode trunk
switchport trunk native vlan <native-vlan-id>
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <vmotion-vlan-id>, <vm-traffic-vlan-id>, <oob-mgmt-vlan-id>
spanning-tree port type edge trunk
mtu 9216
state enabled
exit
```



```
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, follow this step:

1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, follow this step:

1. From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
interface Po10
vpc peer-link
interface Po121
vpc 121
interface Po123
vpc 123
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlashStack environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlashStack environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing

environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Switch Testing Commands

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



Some of these commands need to run after further configuration of the FlashStack components are complete to see complete results.

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

Storage Configuration

Pure Storage FlashArray//X50 R3 Initial Configuration

FlashArray Initial Configuration

The following information should be gathered to enable the installation and configuration of the FlashArray. An official representative of Pure Storage will help rack and configure the new installation of the FlashArray.

Array Settings	Variable Name
Array Name (Hostname for Pure Array):	<<var_flasharray_hostname>>
Virtual IP Address for Management:	<<var_flasharray_vip>>
Physical IP Address for Management on Controller 0 (CT0):	<<var_contoller-1_mgmt_ip >>
Physical IP Address for Management on Controller 1 (CT1):	<<var_contoller-2_mgmt_ip>>
Netmask:	<<var_contoller-1_mgmt_mask>>
Gateway IP Address:	<<var_contoller-1_mgmt_gateway>>
DNS Server IP Address(es):	<<var_nameserver_ip>>
DNS Domain Suffix: (Optional)	<<var_dns_domain_name>>
NTP Server IP Address or FQDN:	<<var_oob_ntp>>
Email Relay Server (SMTP Gateway IP address or FQDN): (Optional)	<<var_smtp_ip>>
Email Domain Name:	<<var_smtp_domain_name>>
Alert Email Recipients Address(es): (Optional)	
HTTP Proxy Server ad Port (For Pure1): (Optional)	
Time Zone:	<<var_timezone>>

When the FlashArray has completed initial configuration, it is important to configure the Cloud Assist phone-home connection to provide the best pro-active support experience possible. Furthermore, this will enable the analytics functionalities provided by Pure1.

Add an Alert Recipient

The Alerts sub-view is used to manage the list of addresses to which Purity delivers alert notifications, and the attributes of alert message delivery. You can designate up to 19 alert recipients. The Alert Recipients section displays a list of email addresses that are designated to receive Purity alert messages. Up to 20 alert recipients can be designated.



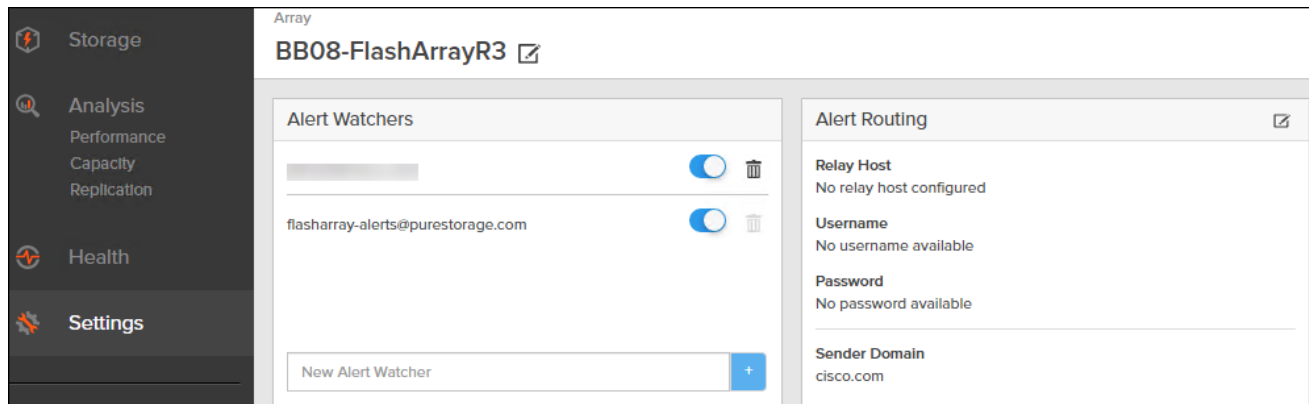
The list includes the built-in flasharray-alerts@purestorage.com address, which cannot be deleted.

The email address that Purity uses to send alert messages includes the sender domain name and is comprised of the following components:

<Array_Name>-<Controller_Name>@<Sender_Domain_Name>.com

To add an alert recipient, follow these steps:

1. Select Settings.
2. In the Alert Watchers section, enter the email address of the alert recipient and click the + icon.



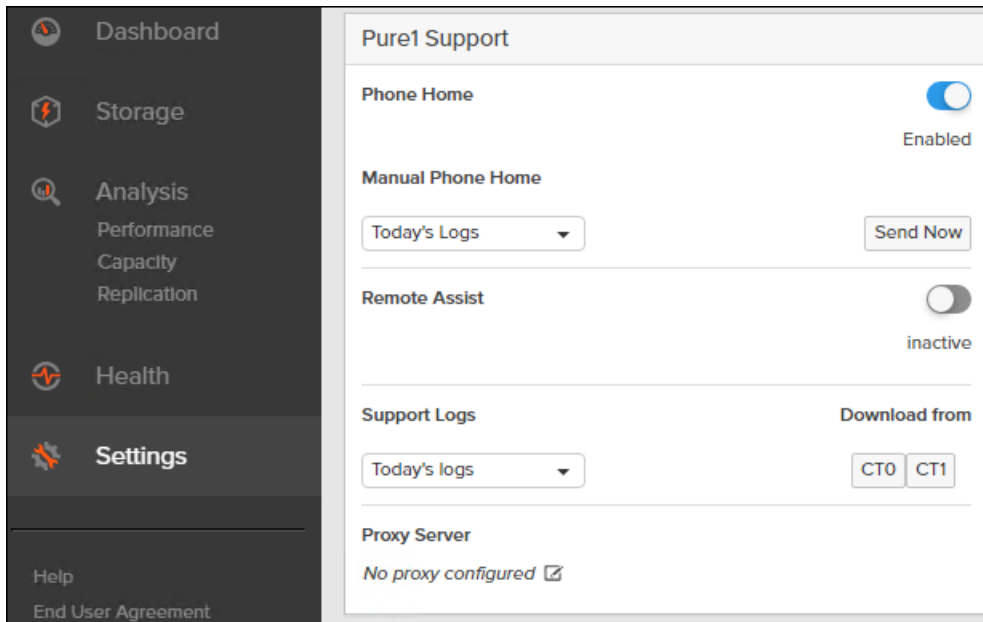
The Relay Host section displays the hostname or IP address of an SMTP relay host, if one is configured for the array. If you specify a relay host, Purity routes the email messages via the relay (mail forwarding) address rather than sending them directly to the alert recipient addresses.

In the Sender Domain section, the sender domain determines how Purity logs are parsed and treated by Pure Storage Support and Escalations. By default, the sender domain is set to the domain name please-configure.me.

It is crucial that you set the sender domain to the correct domain name. If the array is not a Pure Storage test array, set the sender domain to the actual customer domain name. For example, mycompany.com.

Configure Pure1 Support

The Pure1 Support section manages settings for Phone Home, Remote Assist, and Support Logs.



- The phone home facility provides a secure direct link between the array and the Pure Storage Technical Support web site. The link is used to transmit log contents and alert messages to the Pure Storage Support team so that when diagnosis or remedial action is required, complete recent history about array performance and significant events is available. By default, the phone home facility is enabled. If the phone home facility is enabled to send information automatically, Purity transmits log and alert information directly to Pure Storage Support via a secure network connection. Log contents are transmitted hourly and stored at the support web site, enabling detection of array performance and error rate trends. Alerts are reported immediately when they occur so that timely action can be taken.
- Phone home logs can also be sent to Pure Storage Technical support on demand, with options including Today's Logs, Yesterday's Logs, or All Log History.


The Remote Assist section displays the remote assist status as "Connected" or "Disconnected". By default, remote assist is disconnected. A connected remote assist status means that a remote assist session has been opened, allowing Pure Storage Support to connect to the array. Disconnect the remote assist session to close the session.

- The Support Logs section allows you to download the Purity log contents of the specified controller to the current administrative workstation. Purity continuously logs a variety of array activities, including performance summaries, hardware and operating status reports, and administrative actions.

Configure DNS Server IP Addresses

To configure the DNS server IP addresses, follow these steps:

1. Select Settings > Network.
2. In the DNS section, hover over the domain name and click the pencil icon. The Edit DNS dialog box appears.



3. Complete the following fields:
 - a. Domain: Specify the domain suffix to be appended by the array when doing DNS lookups.
 - b. NS#: Specify up to three DNS server IP addresses for Purity to use to resolve hostnames to IP addresses. Enter one IP address in each DNS# field. Purity queries the DNS servers in the order that the IP addresses are listed.
4. Click Save.

Directory Service

The Directory Service manages the integration of FlashArray with an existing directory service. When the Directory Service sub-view is configured and enabled, the FlashArray leverages a directory service to perform user account and permission level searches. Configuring directory services is OPTIONAL.

The screenshot displays the FlashArray web interface. The top navigation bar includes 'System', 'Network', 'Users', and 'Software', with 'Users' selected. The left sidebar contains 'Dashboard', 'Storage', 'Analysis', 'Health', and 'Settings'. The main content area is divided into two sections: 'Users' and 'Directory Service'.

Users Table:

Name	Role	Type	Public Key	API Token	Lockout Remaining
pureuser	array_admin	local		****	-

Directory Service Configuration:

Configuration

Enabled	False
URIs	-
Base DN	-
Bind User	-
Bind Password	-
User Login Attribute	-
User Object Class	-
Check Peer	False
CA Certificate	- Edit

Roles

Name	Group	Group Base
array_admin		
ops_admin		
readonly		
storage_admin		

The FlashArray is delivered with a single local user, named pureuser, with array-wide (Array Admin) permissions.

To support multiple FlashArray users, integrate the array with a directory service, such as Microsoft Active Directory or OpenLDAP.

Role-based access control is achieved by configuring groups in the directory that correspond to the following permission groups (roles) on the array:

- Read Only Group. Read Only users have read-only privilege to run commands that convey the state of the array. Read Only users cannot alter the state of the array.
- Storage Admin Group. Storage Admin users have all the privileges of Read Only users, plus the ability to run commands related to storage operations, such as administering volumes, hosts, and host groups. Storage Admin users cannot perform operations that deal with global and system configurations.
- Array Admin Group. Array Admin users have all the privileges of Storage Admin users, plus the ability to perform array-wide changes. In other words, Array Admin users can perform all FlashArray operations.

To configure the Directory Service, follow these steps:

1. Select Settings > Access > Users.

2. Select the  icon in the Directory Services panel:

- Enabled: Select the check box to leverage the directory service to perform user account and permission level searches.
- URI: Enter the comma-separated list of up to 30 URIs of the directory servers. The URI must include a URL scheme (ldap, or ldaps for LDAP over SSL), the hostname, and the domain. You can optionally specify a port. For example, ldap://ad.company.com configures the directory service with the hostname "ad" in the domain "company.com" while specifying the unencrypted LDAP protocol.
- Base DN: Enter the base distinguished name (DN) of the directory service. The Base DN is built from the domain and should consist only of domain components (DCs). For example, for ldap://ad.storage.company.com, the Base DN would be: "DC=storage,DC=company,DC=com"
- Bind User: Username used to bind to and query the directory. For Active Directory, enter the username - often referred to as sAMAccountName or User Logon Name - of the account that is used to perform directory lookups. The username cannot contain the characters " [] : ; | = + * ? < > / \ and cannot exceed 20 characters in length. For OpenLDAP, enter the full DN of the user. For example, " CN=John,OU=Users,DC=example,DC=com" .
- Bind Password: Enter the password for the bind user account.
- Group Base: Enter the organizational unit (OU) to the configured groups in the directory tree. The Group Base consists of OUs that, when combined with the base DN attribute and the configured group CNs, complete the full Distinguished Name of each groups. The group base should specify " OU=" for each OU and multiple OUs should be separated by commas. The order of OUs should get larger in scope from left to right. In the following example, SANManagers contains the sub-organizational unit PureGroups: " OU=PureGroups,OU=SANManagers" .
- Array Admin Group: Common Name (CN) of the directory service group containing administrators with full privileges to manage the FlashArray. Array Admin Group administrators have the same privileges as pureuser. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, " pureadmins,OU=PureStorage" , where pureadmins is the common name of the directory service group.
- Storage Admin Group: Common Name (CN) of the configured directory service group containing administrators with storage related privileges on the FlashArray. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify the OU. For example, " pureusers,OU=PureStorage" , where pureusers is the common name of the directory service group.
- Read Only Group: Common Name (CN) of the configured directory service group containing users with read-only privileges on the FlashArray. The name should be the Common Name of the group without the " CN=" specifier. If the configured groups are not in the same OU, also specify

the OU. For example, " purereadonly,OU=PureStorage" , where purereadonly is the common name of the directory service group.

- Check Peer: Select the check box to validate the authenticity of the directory servers using the CA Certificate. If you enable Check Peer, you must provide a CA Certificate.
- CA Certificate: Enter the certificate of the issuing certificate authority. Only one certificate can be configured at a time, so the same certificate authority should be the issuer of all directory server certificates. The certificate must be PEM formatted (Base64 encoded) and include the " -----BEGIN CERTIFICATE-----" and " -----END CERTIFICATE-----" lines. The certificate cannot exceed 3000 characters in total length.

3. Click Save.

4. Click Test to test the configuration settings. The LDAP Test Results pop-up window appears. Green squares represent successful checks. Red squares represent failed checks.

SSL Certificate

Self-Signed Certificate

Purity creates a self-signed certificate and private key when you start the system for the first time. The SSL Certificate sub-view allows you to view and change certificate attributes, create a new self-signed certificate, construct certificate signing requests, import certificates and private keys, and export certificates.

Creating a self-signed certificate replaces the current certificate. When you create a self-signed certificate, include any attribute changes, specify the validity period of the new certificate, and optionally generate a new private key.

SSL Certificate	
Status	self-signed
Key Size	2048
Issued To	-
Issued By	-
Valid From	2020-07-15 10:15:04
Valid To	2030-07-13 09:15:04
State/Province	-
Locality	-
Organization	Pure Storage, Inc.
Organizational Unit	Pure Storage, Inc.
Email	-

When you create the self-signed certificate, you can generate a private key and specify a different key size. If you do not generate a private key, the new certificate uses the existing key.

You can change the validity period of the new self-signed certificate. By default, self-signed certificates are valid for 3650 days

CA-Signed Certificate

Certificate authorities (CA) are third party entities outside the organization that issue certificates. To obtain a CA certificate, you must first construct a certificate signing request (CSR) on the array.

Construct Certificate Signing Request ×

Country	Two-letter ISO country code
State/Province	State, province, country or region
Locality	Full city name
Organization	Pure Storage, Inc.
Organization Unit	Pure Storage, Inc.
Common Name	FQDN or management IP address of the server
Email	Email address

The CSR represents a block of encrypted data specific to your organization. You can change the certificate attributes when you construct the CSR; otherwise, Purity will reuse the attributes of the current certificate (self-signed or imported) to construct the new one. Note that the certificate attribute changes will only be visible after you import the signed certificate from the CA.

Send the CSR to a certificate authority for signing. The certificate authority returns the SSL certificate for you to import. Verify that the signed certificate is PEM formatted (Base64 encoded), includes the "-----BEGIN CERTIFICATE-----" and "-----END CERTIFICATE-----" lines, and does not exceed 3000 characters in total length. When you import the certificate, also import the intermediate certificate if it is not bundled with the CA certificate.

Import Certificate

Certificate	<input type="button" value="Choose File"/> No file chosen
Private Key	<input type="button" value="Choose File"/> No file chosen
Intermediate Certificate (optional)	<input type="button" value="Choose File"/> No file chosen
Key Passphrase (optional)	<input type="text"/>

If the certificate is signed with the CSR that was constructed on the current array and you did not change the private key, you do not need to import the key. However, if the CSR was not constructed on the current array or if the private key has changed since you constructed the CSR, you must import the private key. If the private key is encrypted, also specify the passphrase.

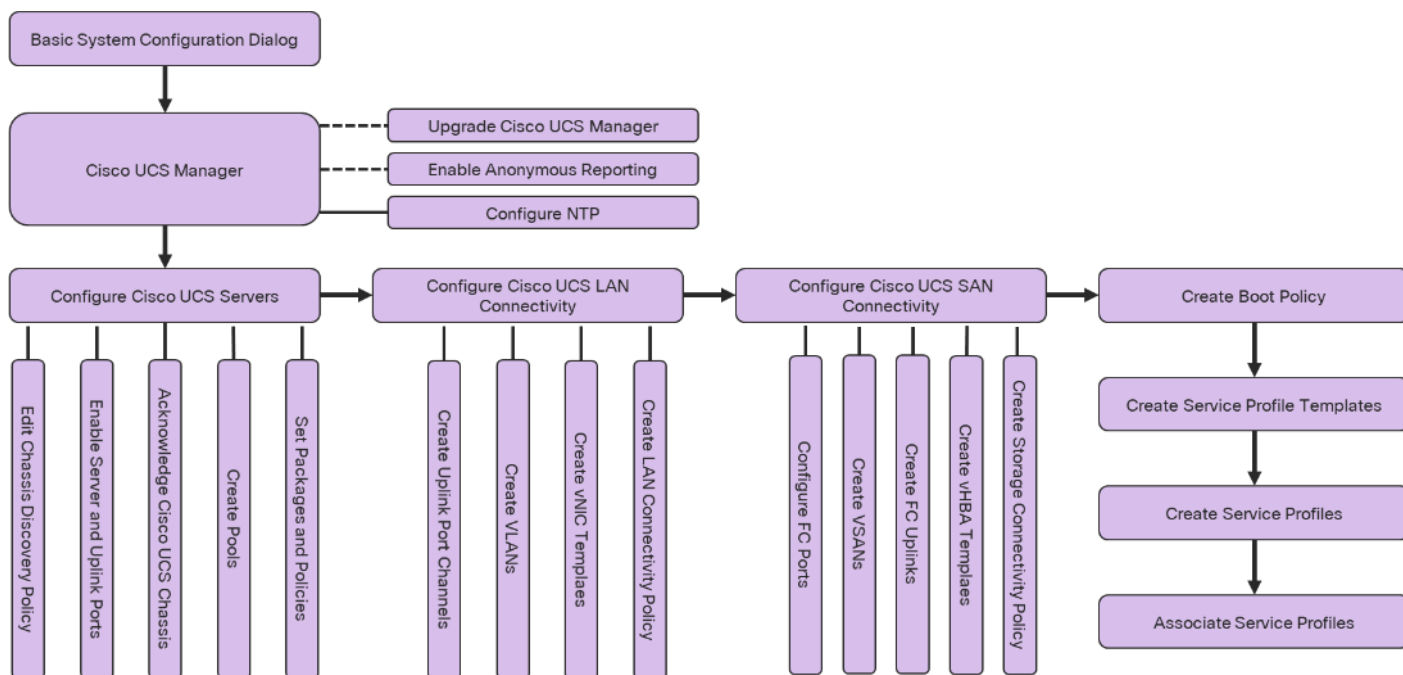


If FC-NVMe is being implemented, the FC ports personality on the FlashArray need to be converted to nvme-fc from the default scsi-fc. In this design we have used two scsi-fc and two nvme-fc ports to support both SCSI and NVMe over Fibre Channel. The ports can be converted to nvme-fc with the help off Pure support.

Cisco UCS Configuration

The following procedures describe how to configure the Cisco UCS domain for use in a base FlashStack environment. This procedure assumes the use of Cisco UCS Fabric Interconnects running 4.2(1f). Configuring a differing model of Cisco UCS Fabric Interconnects should be comparable but may differ slightly with model and changes in Cisco UCS Manager (UCSM) release. The Cisco USC 6454 Fabric Interconnects and Cisco UCS Manager 4.2(1f) release were used in validation of this FlashStack solution, so the steps will reflect this model and release.

Figure 5. Cisco UCS Configuration Workflow



Physical Connectivity

Physical cabling should be completed by following the diagram and table references in section [FlashStack Cabling](#).

Cisco UCS Base Configuration

This FlashStack deployment explains the configuration steps for the Cisco UCS 6454 Fabric Interconnects (FI) in a design that will support Fibre Channel SAN boot.



If setting up a system with iSCSI boot, the sections with (FCP) in the heading can be skipped and then complete the [Cisco UCS iSCSI Configuration](#) section in the Appendix.

Perform Initial Setup of Cisco UCS 6454 Fabric Interconnects for FlashStack Environments

This section provides the detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlashStack environment. These steps are necessary to provision the Cisco UCS B-Series and C-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlashStack environment in ucsd managed mode, follow these steps:

1. Connect to the console port on the first Cisco UCS fabric interconnect.

```
Enter the configuration method. (console/gui) ? console
Enter the management mode. (ucsm/intersight)? ucsm
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect in "ucsm" managed mode. Continue? (y/n): y
Enforce strong password? (y/n) [y]: Enter
Enter the password for "admin": <password>
Confirm the password for "admin": <password>

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: y
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 : <ucsa-mgmt-gateway>
Cluster IPv4 address : <ucs-cluster-ip>
Configure the DNS Server IP address? (yes/no) [n]: y
    DNS IP address : <dns-server-1-ip>
Configure the default domain name? (yes/no) [n]: y
    Default domain name : <ad-dns-domain-name>
Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

2. Wait for the login prompt for Cisco UCS Fabric Interconnect A before proceeding to the next section.

Cisco UCS Fabric Interconnect B

To configure the Cisco UCS for use in a FlashStack environment, follow these steps:

1. Connect to the console port on the second Cisco UCS fabric interconnect.

```
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: <password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <ucsa-mgmt-ip>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <ucsa-mgmt-mask>
Cluster IPv4 address      : <ucs-cluster-ip>

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

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

Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect is compatible with the local fabric interconnect. Continuing with the installer...

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

2. Wait for the login prompt for UCS Fabric Interconnect B before proceeding to the next section.

Cisco UCS Setup

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, follow these steps:

1. Open a web browser and navigate to the Cisco UCS fabric interconnect cluster address.



You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to open.

2. Click the Launch UCS Manager link to launch Cisco UCS Manager.

3. If prompted to accept security certificates, accept as necessary.

4. When prompted, enter admin as the username and enter the administrative password.

5. Click Login to log into Cisco UCS Manager.

Anonymous Reporting

To enable anonymous reporting, follow this step:

1. In the Anonymous Reporting window, choose whether to send anonymous data to Cisco for improving future products. If you choose Yes, enter the IP address of your SMTP Server. Click OK.

Anonymous Reporting

Cisco Systems, Inc. will be collecting feature configuration and usage statistics which will be sent to Cisco Smart Call Home server anonymously. This data helps us prioritize the features and improvements that will most benefit our customers.

If you decide to enable this feature in future, you can do so from the " Anonymous Reporting" in the Call Home settings under the Admin tab.

[View Sample Data](#)

Do you authorize the disclosure of this information to Cisco Smart CallHome?

Yes No

SMTP Server

Host (IP Address or Hostname):

Port:

Don't show this message again.

OK

Cancel

Upgrade Cisco UCS Manager Software to Version 4.2(1f)

This document assumes the use of Cisco UCS 4.2(1f). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.2(1f), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Cisco Intersight can also be used to upgrade the Cisco UCS Infrastructure (Cisco UCS Manager, Cisco UCS Fabric Interconnects, and Cisco UCS Fabric Extenders) to version 4.2(1f). Before the upgrade can be done from Cisco Intersight, the UCS cluster will need to be claimed into Intersight. Please see the Cisco Intersight section of this document for the Cisco Intersight-based upgrade procedure, please see https://intersight.com/help/features#firmware_upgrade for more detailed procedure. This upgrade does require interacting with Cisco UCS Manager to reboot the Primary Fabric Interconnect when upgrading. Because the Cisco UCS servers are not yet connected to the Cisco UCS Infrastructure, the servers will not be upgraded using Cisco Intersight. However, the Cisco UCS B and C-Series 4.2(1f) bundles need to be manually downloaded to the Cisco UCS system.

Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Choose All > Communication Management > Call Home.

-
3. Change the State to On.
 4. Fill in all the fields according to your management preferences and click Save Changes and then click OK to complete configuring Call Home.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Cisco Nexus switches, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand All > Time Zone Management.
3. Choose Timezone.
4. In the Properties pane, choose the appropriate time zone in the Timezone menu.
5. Click Save Changes and then click OK.
6. Click Add NTP Server.
7. Enter <nexus-A-mgmt0-ip> and click OK. Click OK on the confirmation.



Add NTP Server ? ×

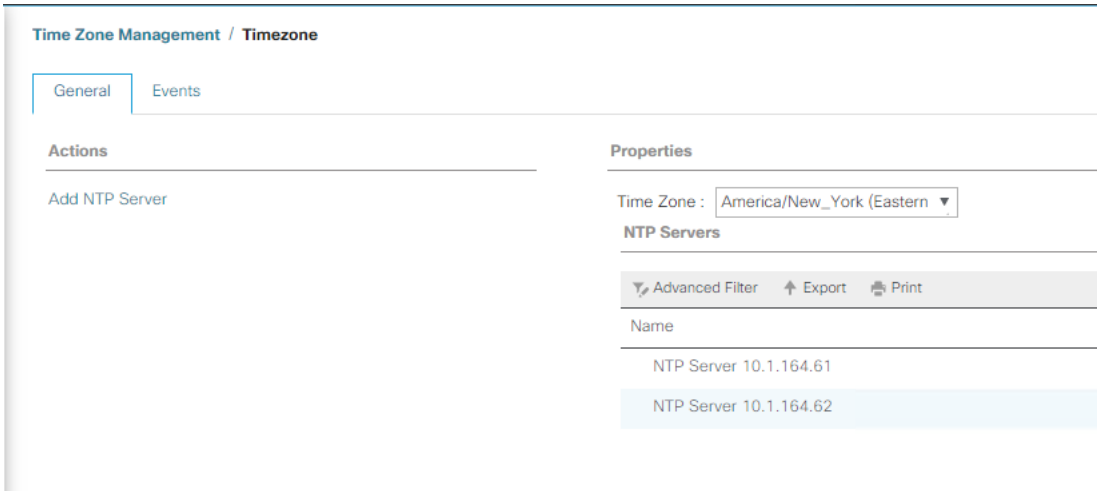
NTP Server :

OK Cancel



We used the Cisco Nexus switch mgmt0 interface IP here because it is in the same L2 domain as the UCS mgmt0 IPs. We could also use the Nexus NTP IPs, but that traffic would then have to pass through an L3 router.

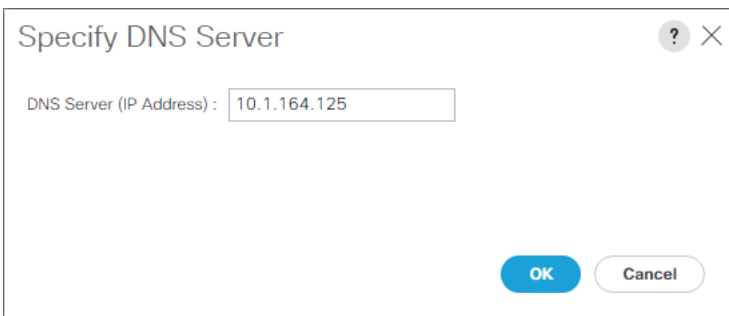
8. Click Add NTP Server.
9. Enter <nexus-B-mgmt0-ip> and click OK, then click OK again.



Add Additional DNS Server(s)

To add one or more additional DNS servers to the UCS environment, follow these steps:

1. In Cisco UCS Manager, click Admin.
2. Expand All > Communications Management.
3. Choose DNS Management.
4. In the Properties pane, choose Specify DNS Server.
5. Enter the IP address of the additional DNS server.



6. Click OK and then click OK again.
7. Repeat steps 1-6 for additional DNS servers.

Add an Additional Administrative User

To add an additional locally authenticated administrative user (flashadmin) to the Cisco UCS environment in case issues arise with the admin user, follow these steps:

1. In Cisco UCS Manager, click Admin.

2. Expand User Management > User Services > Locally Authenticated Users.
3. Right-click Locally Authenticated Users and choose Create User.
4. In the Create User fields it is only necessary to fill in the Login ID, Password, and Confirm Password fields. Fill in the Create User fields according to your local security policy.
5. Leave the Account Status field set to Active.
6. Set Account Expires according to your local security policy.
7. Under Roles, choose admin.
8. Leave Password Required selected for the SSH Type field.

Create User

Login ID : flashadmiin

First Name : FlashStack

Last Name : Administrator

Email :

Phone :

Password : *****

Confirm Password : *****

Account Status : Active Inactive

Account Expires :

Roles

- aaa
- admin
- facility-manager
- network
- operations
- read-only
- server-compute
- server-equipment
- server-profile
- server-security
- storage

Locales

OK Cancel

9. Click OK and then click OK again to complete adding the user.

Configure Unified Ports (FCP)

Fibre Channel port configurations differ between the Cisco UCS 6454, 6332-16UP and the 6248UP fabric interconnects. All fabric interconnects have a slider mechanism within the Cisco UCS Manager GUI interface, but the fibre channel port selection options for the Cisco UCS 6454 are from the first 16 ports starting from the first port and configured in increments of 4 ports from the left. For the Cisco

UCS 6332-16UP the port selection options are from the first 16 ports starting from the first port, and configured in increments of the first 6, 12, or all 16 of the unified ports. With the Cisco UCS 6248UP, the port selection options will start from the right of the 32 fixed ports, or the right of the 16 ports of the expansion module, going down in contiguous increments of 2. The remainder of this section shows configuration of the Cisco UCS 6454. Modify as necessary for the Cisco UCS 6332-16UP or Cisco UCS 6248UP.

To enable the fibre channel ports, follow these steps for the Cisco UCS 6454:

1. In Cisco UCS Manager, click Equipment.
2. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate).
3. Choose Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to choose either 4, 8, 12, or 16 ports to be set as FC Uplinks.

Configure Unified Ports



6. Click OK, then click Yes, then click OK to continue.
7. Choose Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
8. Choose Configure Unified Ports.
9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to choose either 4 or 8 ports to be set as FC Uplinks.
11. Click OK, then click Yes, then OK to continue.
12. Wait for both Fabric Interconnects to reboot.
13. Log back into Cisco UCS Manager.

Edit Chassis Discovery Policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS C-Series connectivity. To modify the chassis discovery policy, follow these steps:

1. In Cisco UCS Manager, click Equipment and choose the Policies tab.
2. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.



If varying numbers of links between chassis and the Fabric Interconnects will be used, set Action to 2 Link, the minimum recommended number of links for a FlashStack.

3. On the Cisco UCS 6454 Fabric Interconnects, the Link Grouping Preference is automatically set to Port Channel and is greyed out. On a Cisco UCS 6300 Series or Cisco UCS 6200 Series Fabric Interconnect, set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G.

Equipment

Main Topology View Fabric Interconnects Servers Thermal Decommissioned Firmware Management Policies Faults Diagnostics

Global Policies Autoconfig Policies Server Inheritance Policies Server Discovery Policies SEL Policy Power Groups Port Auto-Discovery Policy Security

Chassis/FEX Discovery Policy

Action : 2 Link

Link Grouping Preference : None Port Channel

Equipment

Main Topology View Fabric Interconnects Servers Thermal Decommissioned Firmware Management Policies Faults Diagnostics

Global Policies Autoconfig Policies Server Inheritance Policies Server Discovery Policies SEL Policy Power Groups Port Auto-Discovery Policy Security

Chassis/FEX Discovery Policy

Action : 4 Link

Link Grouping Preference : None Port Channel

Warning: Chassis should be re-acked to apply the link aggregation preference change on the fabric interconnect, as this change may cause the IOM to lose connectivity due to fabric port-channel being re-configured.

4. If any changes have been made, click Save Changes, and then click OK.

Enable Port Auto-Discovery Policy

Setting the port auto-discovery policy enables automatic discovery of Cisco UCS B-Series chassis server ports. If you plan to attach B-Series servers to this UCS Domain, enable this policy. To modify the port auto-discovery policy, follow these steps:

1. In Cisco UCS Manager, click Equipment, choose All > Equipment in the Navigation Pane, and choose the Policies tab.

2. Under Port Auto-Discovery Policy, set Auto Configure Server Port to Enabled.

Equipment

The screenshot shows the Cisco UCS Manager interface. At the top, there is a navigation bar with tabs: Main Topology View, Fabric Interconnects, Servers, Thermal, Decommissioned, Firmware Management, Policies (highlighted), Faults, and Diagnostics. Below this is a sub-navigation bar with tabs: Global Policies, Autoconfig Policies, Server Inheritance Policies, Server Discovery Policies, SEL Policy, Power Groups, Port Auto-Discovery Policy (highlighted), and Security. The main content area is divided into three sections: Actions, Properties, and a form. The Actions section contains a link for 'Use Global'. The Properties section shows 'Owner : Local'. The form contains a label 'Auto Configure Server Port:' followed by two radio buttons: 'Disabled' (unselected) and 'Enabled' (selected).

Save Changes

Reset Values

3. Click Save Changes and then click OK.

Enable Server and Uplink Ports

To enable and verify server and uplink ports, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand and choose Ethernet Ports.

-
4. Verify that all ports connected to any Cisco UCS 5108 chassis are configured as Server ports and have a status of Up.
 5. If any ports connected to UCS are missing, choose the ports and right-click them, and choose Configure as Server Port.
 6. Click Yes to confirm server ports and click OK.
 7. Verify that the ports are now configured as server ports.
 8. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
 9. Click Yes to confirm uplink ports and click OK.
 10. Choose Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
 11. Expand and choose Ethernet Ports.
 12. Verify that all ports connected to Cisco UCS chassis and rack mounts are configured as Server ports and have a status of Up.
 13. If any ports are missing, choose the ports and right-click them, and choose Configure as Server Port.
 14. Click Yes to confirm server ports and click OK.
 15. Verify that the ports are now configured as server ports.
 16. Choose the ports that are connected to the Cisco Nexus switches, right-click them, and choose Configure as Uplink Port.
 17. Click Yes to confirm the uplink ports and click OK.

Enable Info Policy for Neighbor Discovery

Enabling the info policy enables Fabric Interconnect neighbor information to be displayed. To modify the info policy, follow these steps:

1. In Cisco UCS Manager, click Equipment, choose All > Equipment in the Navigation Pane, and choose the Policies tab on the right.
2. Under Global Policies, scroll down to Info Policy and choose Enabled for Action.

Info Policy

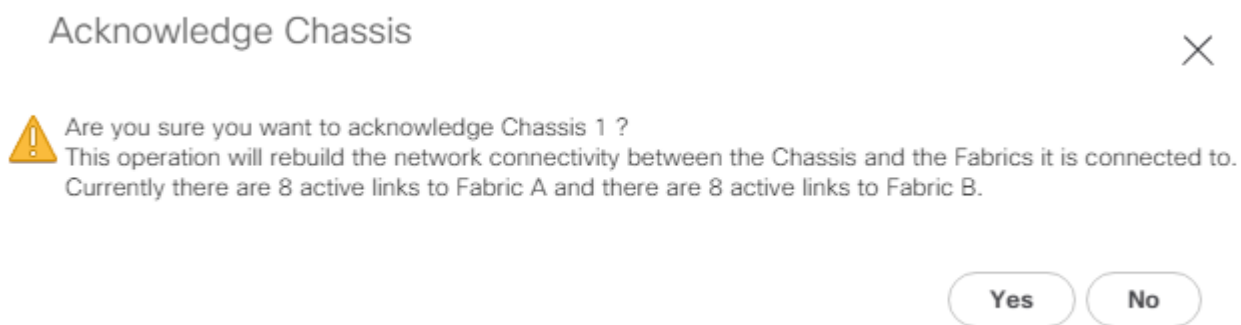
Action : Disabled Enabled

3. Click Save Changes and then click OK.
4. Under Equipment, choose Fabric Interconnect A or B. On the right, choose the Neighbors tab. CDP information is shown under the LAN tab and LLDP information is shown under the LLDP tab.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge any Cisco UCS chassis and any external FEX modules, follow these steps:

1. In Cisco UCS Manager, click Equipment.
2. Expand Chassis and choose each chassis that is listed.
3. Right-click each chassis and choose Acknowledge Chassis.



4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus FEXes are part of the configuration, expand Rack Mounts and FEX.
6. Right-click each FEX that is listed and choose Acknowledge FEX.
7. Click Yes and then click OK to complete acknowledging the FEX.

Create an organization

To this point in the Cisco UCS deployment, all items have been deployed at the root level in Cisco UCS Manager. To allow this Cisco UCS to be shared among different projects, Cisco UCS Organizations can be created. In this validation, the organization for this FlashStack deployment is FlashStack. To create an organization for this FlashStack deployment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. In the Navigation Pane, expand Servers > Service Profiles.
3. Right-click root under Service Profiles and choose Create Organization.

4. Provide a name for the Organization to indicate this FlashStack deployment and optionally provide a Description.

Create Organization ? ×

Name :

Description :

5. Click OK then click OK again to complete creating the organization.

Create a WWNN Pool for FC Boot (FCP)

In this FlashStack implementation, a WWNN pool is created at the root organization level to avoid WWNN address pool overlaps. If your deployment plan calls for different WWNN ranges in different Cisco UCS organizations, place the WWNN pool at the organizational level.

To configure the necessary WWNN pool for the Cisco UCS environment, follow these steps on Cisco UCS Manager:

1. Choose SAN.
2. Choose Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Choose Create WWNN Pool to create the WWNN pool.
5. Enter WWNN-Pool for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Choose Sequential for Assignment Order.

1 Define Name and Description

2 Add WWN Blocks

Create WWNN Pool

Name : WWNN-Pool

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.

9. Click Add.

10. Modify the From field as necessary for the UCS Environment



Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the Cisco UCS domain. Within the From field in our example, the sixth and seventh octets were changed from 00:00 to A4:00.



When there are multiple UCS domains sitting in adjacency, it is important that these blocks; the WWNN, WWPN, and MAC, hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources. In this example, with the WWNN block modification, a maximum of 32 addresses are available.

Create WWN Block



From : Size :

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

OK

Cancel

12. Click OK.

13. Click Finish and click OK to complete creating the WWNN pool.

Create WWPN Pools (FCP)

In this FlashStack implementation, WWPN address pools are created at the root organization level to avoid WWPN address pool overlaps. If your deployment plan calls for different WWPN address ranges in different UCS organizations, place the WWPN pools at the organizational level.

To configure the necessary WWPN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Choose Pools > root.



In this procedure, two WWPN pools are created, one for each switching fabric.

3. Right-click WWPN Pools under the root organization.
4. Choose Create WWPN Pool to create the WWPN pool.
5. Enter WWPN-Pool-A as the name of the WWPN pool.
6. Optional: Enter a description for the WWPN pool.
7. Choose Sequential for Assignment Order.

1 Define Name and Description

2 Add WWN Blocks

Create WWPN Pool

Name : WWPN-Pool-A

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.

9. Click Add.

10. Specify a starting WWPN.



For the FlashStack solution, the recommendation is to place **A** in the next-to-last octet of the starting WWPN to identify all the WWPNs as fabric A addresses. We used a WWPN block starting with `20:00:00:25:B5:A4:0A:00`

11. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created. In this example, with the WWPN block modification, a maximum of 32 addresses are available.

Create WWN Block



From : Size :

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:

20:00:00:25:b5:xx:xx:xx

OK

Cancel

12. Click OK.
13. Click Finish.
14. In the confirmation message, click OK.
15. Right-click WWPN Pools under the root organization.
16. Choose Create WWPN Pool to create the WWPN pool.
17. Enter WWPN-Pool-B as the name of the WWPN pool.
18. Optional: Enter a description for the WWPN pool.
19. Choose Sequential for Assignment Order.
20. Click Next.
21. Click Add.
22. Specify a starting WWPN.



For the FlashStack solution, the recommendation is to place B in the next-to-last octet of the starting WWPN to identify all the WWPNs as fabric B addresses. We used a WWPN block starting with 20:00:00:25:B5:A4:0B:00.

23. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources remembering that servers could have multiple vHBAs and unassociated service profiles could be created. In this example, with the WWPN block modification, a maximum of 32 addresses are available.

24. Click OK.

25. Click Finish.

26. In the confirmation message, click OK.

Create VSANs (FCP)

To configure the necessary virtual storage area networks (VSANs) for the FlashStack-VSI Organization in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.



In this procedure, two VSANs are created, one for each SAN switching fabric.

2. Choose SAN > SAN Cloud.

3. Right-click VSANs.

4. Choose Create VSAN.

5. Enter FlashStack-Fabric-A as the name of the VSAN to be used for Fabric A.

6. Leave FC Zoning set at Disabled.

7. Choose Fabric A.

8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

Create VSAN



Name :

FC Zoning Settings

FC Zoning : Disabled Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VSAN ID that maps to this VSAN.

Enter the VLAN ID that maps to this VSAN.

VSAN ID :

FCoE VLAN :

OK

Cancel

9. Click OK and then click OK again.
10. Under SAN Cloud, right-click VSANs.
11. Choose Create VSAN.
12. Enter FlashStack-Fabric-B as the name of the VSAN to be used for Fabric B.
13. Leave FC Zoning set at Disabled.
14. Choose Fabric B.
15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric B. It is recommended use the same ID for both parameters and to use something other than 1.
16. Click OK and then click OK again.

Enable FC Uplink VSAN Trunking (FCP)

To enable VSAN trunking on the FC Uplinks in the Cisco UCS environment, follow these steps:




Enabling VSAN trunking is optional. It is important that the Cisco MDS VSAN trunking configuration match the configuration set in Cisco UCS Manager.

-
1. In Cisco UCS Manager, click SAN.
 2. Expand SAN > SAN Cloud.
 3. Choose Fabric A and in the actions pane under General tab, choose Enable FC Uplink Trunking.
 4. Click Yes on the Confirmation and Warning.
 5. Click OK.
 6. Choose Fabric B and in the actions pane under General tab, choose Enable FC Uplink Trunking.
 7. Click Yes on the Confirmation and Warning.
 8. Click OK.

Create FC Uplink Port Channels (FCP)

To create the FC Uplink Port Channels and assign the appropriate VSANs to them for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Choose SAN > SAN Cloud.
3. Expand Fabric A and choose FC Port Channels.
4. Right-click FC Port Channels and choose Create FC Port Channel.
5. Set a unique ID for the port channel and provide a unique name for the port channel.
6. Click Next.
7. Choose the appropriate Port Channel Admin Speed.
8. Choose the ports connected to Cisco MDS 9132T A and use >> to add them to the port channel.

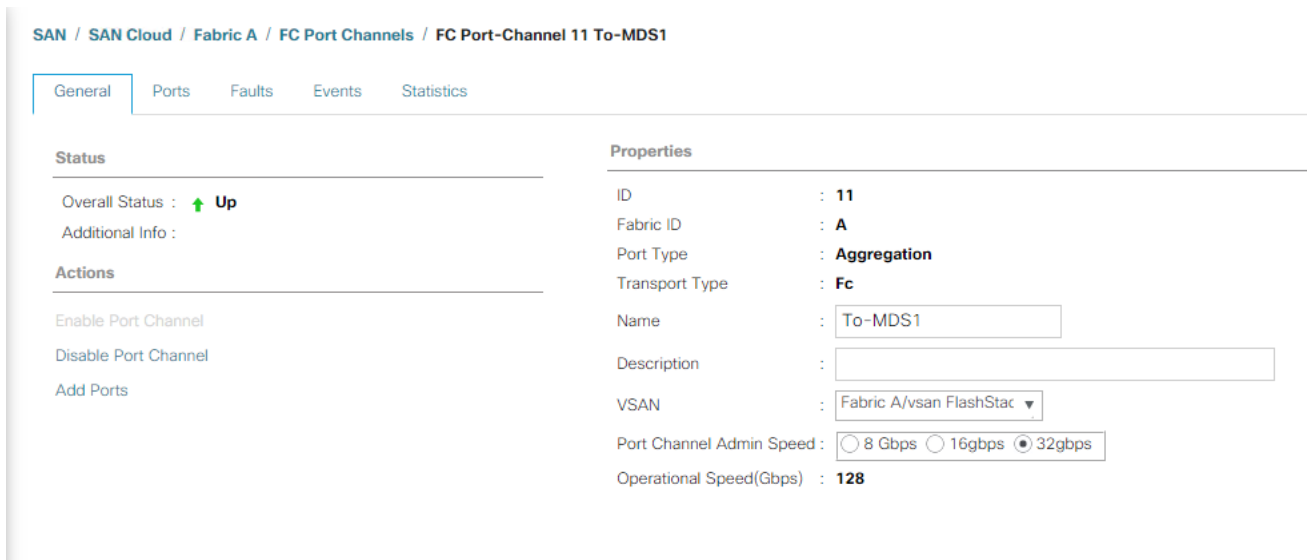


9. Click Finish to complete creating the port channel.

10. Click OK on the confirmation.

11. Under FC Port-Channels, choose the newly created port channel.

12. From the drop-down list to choose FlashStack-Fabric-A.



13. Click Save Changes to assign the VSAN.

-
14. Click OK.
 15. On the left under FC Port Channels, expand the newly created FC Port-Channel. Under the port-channel choose the first FC Interface. Enter a User Label to indicate the connectivity on the MDS 9132T switch, such as <mds-A-hostname>:fc1/5. Click Save Changes and then click OK. Repeat this process for the other FC Interface.
 16. Expand Fabric B and choose FC Port Channels.
 17. Right-click FC Port Channels and choose Create FC Port Channel.
 18. Set a unique ID for the port channel and provide a unique name for the port channel.
 19. Click Next.
 20. Choose the ports connected to Cisco MDS 9132T B and use >> to add them to the port channel.
 21. Click Finish to complete creating the port channel.
 22. Click OK on the confirmation.
 23. Under FC Port-Channels, choose the newly created port channel.
 24. In the right pane, use the drop-down to choose FlashStack-Fabric-B.
 25. Click Save Changes to assign the VSAN.
 26. Click OK.
 27. On the left under FC Port Channels, expand the newly created FC Port-Channel. Under the FC Port-Channel choose the first FC Interface. Enter a User Label to indicate the connectivity on the MDS 9132T switch, such as <mds-B-hostname>:fc1/5. Click Save Changes and then click OK. Repeat this process for the other FC Interface.

Disable Unused FC Uplink Ports (FCP) - Optional

When Unified Ports were configured earlier in this procedure, on the Cisco UCS 6454 FI and the Cisco UCS 6332-16UP FI, FC ports were configured in groups. Because of this group configuration, some FC ports may be unused and need to be disabled to prevent alerts.

To disable the unused FC ports 5 and 6 for example on the Cisco UCS 6454 FIs, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. In the Navigation Pane, expand SAN > SAN Cloud > Fabric A > Uplink FC Interfaces.
3. Right-click FC Interface 1/5 and choose Disable Interface.

-
4. Click Yes and then click OK to complete disabling FC Interface 1/5.
 5. Repeat this process to disable FC Interface 1/6.
 6. In the Navigation Pane, expand SAN > SAN Cloud > Fabric B > Uplink FC Interfaces.
 7. Right-click FC Interface 1/1 and choose Disable Interface.
 8. Click Yes and then click OK to complete disabling FC Interface 1/5.
 9. Repeat step 1-8 to disable FC Interface 1/6.

Create vHBA Templates (FCP)

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment within the FlashStack-VSI Organization, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Expand Policies > root > Sub-Organizations > FlashStack-VSI.
3. Right-click vHBA Templates under the FlashStack-VSI Organization.
4. Choose Create vHBA Template.
5. Enter vHBA-A as the vHBA template name.
6. Keep Fabric A selected.
7. Leave Redundancy Type set to No Redundancy.
8. Choose FlashStack-Fabric-A.
9. Leave Initial Template as the Template Type.
10. Choose WWPN-Pool-A as the WWPN Pool.

Create vHBA Template



Name : vHBA-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : FlashStack-Fabric-A [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-Pool-A(29/32)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

OK

Cancel

11. Click OK to create the vHBA template.
12. Click OK.
13. Right-click vHBA Templates under the FlashStack-VSI Organization.
14. Choose Create vHBA Template.
15. Enter vHBA-B as the vHBA template name.
16. Choose B as the Fabric ID.
17. Leave Redundancy Type set to No Redundancy.
18. Choose FlashStack-Fabric-B.
19. Leave Initial Template as the Template Type.
20. Choose WWPN-Pool-B as the WWPN Pool.
21. Click OK to create the vHBA template.
22. Click OK.

Create SAN Connectivity Policy (FCP)

To configure the necessary Infrastructure SAN Connectivity Policy within the FlashStack-VSI Organization, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Choose SAN > Policies > root > Sub-Organizations > FlashStack-VSI.
3. Right-click SAN Connectivity Policies under the FlashStack-VSI Organization.
4. Choose Create SAN Connectivity Policy.
5. Enter FC-Boot as the name of the policy.
6. Choose the previously created WWNN-Pool for the WWNN Assignment.
7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter FCP-Fabric-A as the name of the vHBA.
9. Choose the Use vHBA Template checkbox.
10. In the vHBA Template list, choose vHBA-A.
11. In the Adapter Policy list, choose VMWare.

Create vHBA



Name : FC-Fabric-A

Use vHBA Template :

Redundancy Pair :

Peer Name :

vHBA Template : vHBA-A ▼

Create vHBA Template

Adapter Performance Profile

Adapter Policy : VMWare ▼

Create Fibre Channel Adapter Policy

OK

Cancel

12. Click OK.

13. Click the Add button at the bottom to add a second vHBA.

14. In the Create vHBA dialog box, enter FCP-Fabric-B as the name of the vHBA.

15. Choose the Use vHBA Template checkbox.

16. In the vHBA Template list, choose vHBA-B.

17. In the Adapter Policy list, choose VMWare.

18. Click OK.

19. If configuring FC-NVMe in this FlashStack, click the Add button at the bottom to add an FC-NVMe vHBA.



Skip creating the FC-NVMe initiators if FC-NVMe storage connectivity is not required.

20. In the Create vHBA dialog box, enter FC-NVMe-Fabric-A as the name of the vHBA.

21. Choose the Use vHBA Template checkbox.

22. In the vHBA Template list, choose vHBA-A.

23. In the Adapter Policy list, choose FCNVMeInitiator.

Create vHBA



Name :

Use vHBA Template :

Redundancy Pair :

Peer Name :

vHBA Template :

[Create vHBA Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Fibre Channel Adapter Policy](#)

OK

Cancel

24. Click OK.

25. Click the Add button at the bottom to add a second FC-NVMe vHBA.

26. In the Create vHBA dialog box, enter FC-NVMe-Fabric-B as the name of the vHBA.

27. Choose the Use vHBA Template checkbox.

28. In the vHBA Template list, choose vHBA-B.

29. In the Adapter Policy list, choose FCNVMInitiator.

30. Click OK.

Create SAN Connectivity Policy



Name :

Description :

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

World Wide Node Name

WWNN Assignment:

[Create WWNN Pool](#)

The WWNN will be assigned from the selected pool.
The available/total WWNNs are displayed after the pool name.

Name	WWPN
▶ vHBA FC-NVMe-Fabric-B	Derived
▶ vHBA FC-NVMe-Fabric-A	Derived
▶ vHBA FCP-Fabric-B	Derived
▶ vHBA FCP-Fabric-A	Derived

Delete Add Modify

OK

Cancel

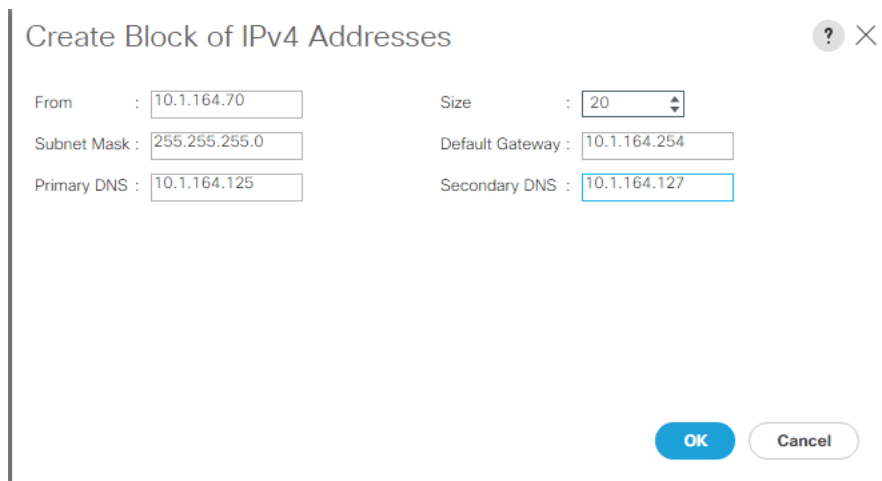
31. Click OK to create the SAN Connectivity Policy.

32. Click OK to confirm creation.

Add Block of IP Addresses for KVM Access

To create a block of IP addresses for in band server Keyboard, Video, Mouse (KVM) access in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and choose Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information. Optionally, enter the Primary and Secondary DNS server addresses.



Create Block of IPv4 Addresses

From : 10.1.164.70 Size : 20

Subnet Mask : 255.255.255.0 Default Gateway : 10.1.164.254

Primary DNS : 10.1.164.125 Secondary DNS : 10.1.164.127

OK Cancel

5. Click OK to create the block.
6. Click OK in the confirmation message.

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.



In this procedure, two port channels are created: one from fabric A to both Cisco Nexus switches and one from fabric B to both Cisco Nexus switches.

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels under Fabric A.

4. Choose Create Port Channel.
5. Enter 145 as the unique ID of the port channel.
6. Enter Po145-Nexus as the name of the port channel.
7. Click Next.
8. Choose the uplink ports connected to the Nexus switches to be added to the port channel.
9. Click >> to add the ports to the port channel.

Create Port Channel

1 Set Port Channel Name

2 Add Ports

Ports			
Slot ID	Aggr. Po...	Port	MAC
No data available			

>>

<<

Ports in the port channel			
Slot ID	Aggr. Po...	Port	MAC
1	0	33	00:3A:9...
1	0	34	00:3A:9...
1	0	35	00:3A:9...
1	0	36	00:3A:9...
1	0	37	00:3A:9...
1	0	38	00:3A:9...
1	0	39	00:3A:9...
1	0	40	00:3A:9...

< Prev Next > **Finish** Cancel

10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, choose Port-Channel 145. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 145 Po145-Nexus

General Ports Faults Events Statistics

Status

Overall Status : ↑ **Up**
 Additional Info : **none**

Actions

Enable Port Channel
 Disable Port Channel
 Add Ports

Properties

ID : **145**
 Fabric ID : **A**
 Port Type : **Aggregation**
 Transport Type : **Ether**
 Name :
 Description :
 Flow Control Policy :
 LACP Policy :
 Note: Changing LACP policy may flap the port-channel if the suspend-individual value changes!
 Admin Speed : 1 Gbps 10 Gbps 40 Gbps 25 Gbps 100 Gbps Auto
 Operational Speed(Gbps) : **200**

13. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.

14. Right-click Port Channels under Fabric B.

15. Choose Create Port Channel.

16. Enter 146 as the unique ID of the port channel.

17. Enter Po146-Nexus as the name of the port channel.

18. Click Next.

19. Choose the ports connected to the Nexus switches to be added to the port channel:

20. Click >> to add the ports to the port channel.

21. Click Finish to create the port channel.

22. Click OK.

23. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, choose Port-Channel 146. Ensure Auto is selected for the Admin Speed. After a few minutes, verify that the Overall Status is Up, and the Operational Speed is correct.

24. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 145. Under Port-Channel 145, choose Eth Interface 1/45. In the center pane under Properties, enter a User Label to indicate the port connectivity, such as <nexus-a-hostname>:Eth1/1. Click Save Changes and then click OK.

25. Repeat steps 1-24 for the remaining seven uplink ports.

Add UDLD to Uplink Port Channels

To configure the unidirectional link detection (UDLD) on the Uplink Port Channels to the Cisco Nexus switches for fibre optic connections, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > LAN Cloud > UDLD Link Policy.
3. Right-click UDLD Link Policy and choose Create UDLD Link Policy.
4. Name the Policy UDLD-Normal and choose Enabled for the Admin State and Normal for the Mode.

Create UDLD Link Policy



Name :

Admin State : Enabled Disabled

Mode : Normal Aggressive



5. Click OK, then click OK again to complete creating the policy.
6. Expand Policies > LAN Cloud > Link Profile.
7. Right-click Link Profile and choose Create Link Profile.
8. Name the Profile UDLD-Normal and choose the UDLD-Normal Link Policy created above.

Create Link Profile



Name :

UDLD Link Policy :

- Click OK, then click OK again to complete creating the profile.
- In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 145. Choose the first Eth Interface under Port-Channel 145. From the drop-down list, choose the UDLD-Normal Link Profile created above, click Save Changes and then click OK. Repeat this process for each Eth Interface under Port-Channel 145 and for each Eth Interface under Port-Channel 146 on Fabric B.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 145 Po145-Nexus / Eth Interface 1/33 (N9K-A:Eth1/1)

General | Faults | Events

Actions	Properties
Delete	ID : 33
Enable Interface	Slot ID : 1
Disable Interface	Fabric ID : A
	Transport Type : Ether
	Port : <i>sys/switch-A/slot-1/switch-ether/port-33</i>
	Membership : Up
	Link Profile : <input type="text" value="UDLD-Normal"/>
	User Label : <input type="text" value="N9K-A:Eth1/1"/>

Set Jumbo Frames in Cisco UCS Fabric

Jumbo Frames are used in FlashStack for the iSCSI storage protocols. The normal best practice in FlashStack has been to set the MTU of the Best Effort QoS System Class in Cisco UCS Manager to 9216 for Jumbo Frames. In the Cisco UCS 6454 Fabric Interconnect with Cisco UCS Manager version 4.0 software the MTU for the Best Effort QoS System Class is fixed at normal and cannot be changed. With this setting of normal in the 6454, Jumbo Frames can pass through the Cisco UCS fabric without being dropped. In Cisco UCS Manager version 4.1 and 4.2, the MTU for the Best Effort QoS System Class is again modifiable.

To configure jumbo frames in the Cisco UCS fabric, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes.
6. Click OK.

LAN / LAN Cloud

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input type="checkbox"/>	5	<input type="checkbox"/>	10	N/A	normal	<input type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	50	9216	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	50	fc	N/A

[Configure Slow Drain Timers](#)

[Configure WD timers](#)



Only the Fibre Channel and Best Effort QoS System Classes are enabled in this FlashStack implementation. The Cisco UCS and Cisco Nexus switches are intentionally configured this way so that all IP traffic within the FlashStack will be treated as Best Effort. Enabling the other QoS

System Classes without having a comprehensive, end-to-end QoS setup in place can cause difficulty in troubleshoot issues.

Create VLANs

To configure the necessary virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.



In this procedure, five unique VLANs are created. See [Table2](#).

2. Expand LAN > LAN Cloud.
3. Right-click VLANs.
4. Choose Create VLANs.
5. Enter Native-VLAN as the name of the VLAN to be used as the native VLAN.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.

Create VLANs



VLAN Name/Prefix :

Multicast Policy Name : [Create Multicast Policy](#)

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. " 2009-2019", " 29,35,40-45", " 23", " 23,34-45")

VLAN IDs :

Sharing Type : None Primary Isolated Community

Check Overlap

OK

Cancel

10. Expand the list of VLANs in the navigation pane, right-click the newly created Native-VLAN and choose Set as Native VLAN.
11. Click Yes and then click OK.
12. Right-click VLANs.
13. Choose Create VLANs
14. Enter IB-MGMT as the name of the VLAN to be used for management traffic.



Modify these VLAN names as necessary for your environment.

15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.

17. Keep the Sharing Type as None.
18. Click OK, and then click OK again.
19. Right-click VLANs.
20. Choose Create VLANs.
21. Enter vMotion-VLAN as the name of the VLAN to be used for vMotion.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the vMotion VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK and then click OK again.
26. Choose Create VLANs.
27. Enter VM-Traffic-VLAN as the name of the VLAN to be used for VM Traffic.
28. Keep the Common/Global option selected for the scope of the VLAN.
29. Enter the VM-Traffic VLAN ID.
30. Keep the Sharing Type as None.
31. Click OK and then click OK again.

LAN / LAN Cloud / VLANs

VLANs

Advanced Filter Export Print

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN vMotion-VLAN (1130)	1130	Lan	Ether	No	None
VLAN VM-Traffic-VLAN (1101)	1101	Lan	Ether	No	None
VLAN Native-Vlan (2)	2	Lan	Ether	Yes	None
VLAN IB-MGMT-VLAN (115)	115	Lan	Ether	No	None
VLAN default (1)	1	Lan	Ether	No	None

Add Delete Info

Create MAC Address Pools

In this FlashStack implementation, MAC address pools are created at the root organization level to avoid MAC address pool overlaps. If your deployment plan calls for different MAC address ranges in different UCS organizations, place the MAC pools at the organizational level.

To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root.



In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Choose Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Choose Sequential as the option for Assignment Order.
8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



For the FlashStack solution, the recommendation is to place A in the next-to-last octet of the starting MAC address to identify all the MAC addresses as fabric A addresses. In our example, we have used 00:25:B5:91:1A:00 as our first MAC address.

11. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. In this example, with the MAC block modification, a maximum of 256 addresses are available.

Create a Block of MAC Addresses ? ×

First MAC Address : Size :

To ensure uniqueness of MACs in the LAN fabric, you are strongly encouraged to use the following MAC prefix:
00:25:B5:xx:xx:xx

12. Click OK.

-
13. Click Finish.
 14. In the confirmation message, click OK.
 15. Right-click MAC Pools under the root organization.
 16. Choose Create MAC Pool to create the MAC address pool.
 17. Enter MAC-Pool-B as the name of the MAC pool.
 18. Optional: Enter a description for the MAC pool.
 19. Choose Sequential as the option for Assignment Order.
 20. Click Next.
 21. Click Add.
 22. Specify a starting MAC address.



For the FlashStack solution, it is recommended to place B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, our example uses 00:25:B5:91:1B:00 as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources remembering that a server may contain multiple vNICs and that multiple unassociated Service Profiles can be created. In this example, with the MAC block modification, a maximum of 256 addresses are available.
24. Click OK.
25. Click Finish.
26. In the confirmation message, click OK.

Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)

To create a network control policy that enables CDP and LLDP on server virtual network controller (vNIC) ports, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root.
3. Right-click Network Control Policies.

4. Choose Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, choose the Enabled option.
7. For LLDP, scroll down and choose Enabled for both Transmit and Receive.

Create Network Control Policy



CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

OK

Cancel

8. Click OK to create the network control policy.
9. Click OK.

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates within the FlashStack-VSI Organization, follow these steps. A total of 4 vNIC Templates will be created. Two of the vNIC templates (vSwitch0-A and vSwitch0-B) will be created for vNICs to connect to VMware ESXi vSwitch0. vSwitch0 will have port groups for the IB-MGMT, vMotion, and VM-Traffic VLANs. The third and fourth vNIC templates (vDS0-A and vDS0-B) will be created for vNICs to connect to the VMware Virtual Distributed Switch (vDS0). The vDS will have port groups for the vMotion and VM-Traffic VLANs. The vMotion VLAN is being placed on both vSwitch0 and vDS0 so that the vMotion VMkernel port can initially be created on vSwitch0 then migrated to the vDS to allow QoS marking of vMotion packets to occur within the vDS if QoS policies need to be applied to vMotion in the future. Any tenant or application VLANs can be placed on the vDS in the future.

Create Infrastructure vNIC Templates

To create the infrastructure vNIC templates, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand Policies > root > Sub-Organizations > FlashStack-VSI.
3. Under the FlashStack-VSI Organization, right-click vNIC Templates.
4. Choose Create vNIC Template.
5. Enter vSwitch0-A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Choose Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Choose Updating Template as the Template Type.
12. Under VLANs, choose the checkboxes for IB-MGMT-VLAN, vMotion-VLAN, and Native-VLAN VLANs.
13. Set IB-MGMT-VLAN as the native VLAN.
14. Choose vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, choose MAC-Pool-A.
17. In the Network Control Policy list, choose Enable-CDP-LLDP.

Create vNIC Template

Advanced Filter Export Print

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	default	<input type="radio"/>	1
<input checked="" type="checkbox"/>	IB-MGMT-VLAN	<input checked="" type="radio"/>	115
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>	901
<input type="checkbox"/>	Native-Vlan	<input type="radio"/>	2
<input checked="" type="checkbox"/>	OOB-MGMT-VLAN	<input type="radio"/>	15
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>	1101

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

usNIC Connection Policy : ▼

OK Cancel

18. Click OK to create the vNIC template.

19. Click OK.

20. Under the FlashStack-VSI organization, right-click vNIC Templates.

21. Choose Create vNIC Template.

22. Enter vSwitch0-B as the vNIC template name.

23. Choose Fabric B.

24. Do not select the Enable Failover checkbox.

25. Set Redundancy Type to Secondary Template.

26. Choose vSwitch0-A for the Peer Redundancy Template.

27. In the MAC Pool list, choose MAC-Pool-B.



The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

28. Click OK to create the vNIC template.

29. Click OK.

30. Under the FlashStack-VSI Organization, right-click vNIC Templates.

31. Choose Create vNIC Template.

32. Enter vDS0-A as the vNIC template name.

33. Keep Fabric A selected.

34. Do not select the Enable Failover checkbox.

35. Choose Primary Template for Redundancy Type.

36. Leave the Peer Redundancy Template set to <not set>.

37. Under Target, make sure that only the Adapter checkbox is selected.

38. Choose Updating Template as the Template Type.

39. Under VLANs, choose the checkboxes for vMotion-VLAN-VLAN, and Native-VLAN VLANs.

40. Set IB-MGMT-VLAN as the native VLAN.

41. Choose vNIC Name for the CDN Source.

42. For MTU, enter 9000.

43. In the MAC Pool list, choose MAC-Pool-A.

44. In the Network Control Policy list, choose Enable-CDP-LLDP.

Create vNIC Template

Advanced Filter Export Print

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	IB-MGMT-VLAN	<input type="radio"/>	115
<input type="checkbox"/>	iSCSI-A-VLAN	<input type="radio"/>	901
<input checked="" type="checkbox"/>	Native-Vlan	<input checked="" type="radio"/>	2
<input type="checkbox"/>	OOB-MGMT-VLAN	<input type="radio"/>	15
<input checked="" type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>	1101
<input checked="" type="checkbox"/>	vMotion-VLAN	<input type="radio"/>	1130

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

usNIC Connection Policy : ▼

OK Cancel

45. Click OK to create the vNIC template.

46. Click OK.

47. Under the FlashStack-VSI organization, right-click vNIC Templates.

48. Choose Create vNIC Template

49. Enter vDS0-B as the vNIC template name.

50. Choose Fabric B.

51. Do not select the Enable Failover checkbox.

52. Set Redundancy Type to Secondary Template.

53. Choose vDS0-A for the Peer Redundancy Template.

54. In the MAC Pool list, choose MAC-Pool-B.



The MAC Pool is all that needs to be selected for the Secondary Template, all other values will either be propagated from the Primary Template or set at default values.

55. Click OK to create the vNIC template.

56. Click OK.

Create High Traffic VMware Adapter Policy

To create the optional VMware-High-Traffic Ethernet Adapter policy to provide higher vNIC performance, follow these steps:



This Ethernet Adapter policy can be attached to vNICs when creating the LAN Connectivity policy for vNICs that have large amounts of traffic on multiple flows or TCP sessions. This policy provides more hardware receive queues handled by multiple CPUs to the vNIC.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Right-click Adapter Policies and choose Create Ethernet Adapter Policy.
4. Name the policy VMware-HighTrf.
5. Expand Resources and set the values as shown below.

Resources

Pooled	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Transmit Queues	:	<input type="text" value="8"/> [1-1000]
Ring Size	:	<input type="text" value="256"/> [64-4096]
<hr/>		
Receive Queues	:	<input type="text" value="8"/> [1-1000]
Ring Size	:	<input type="text" value="512"/> [64-4096]
<hr/>		
Completion Queues	:	<input type="text" value="16"/> [1-2000]
Interrupts	:	<input type="text" value="18"/> [1-1024]



In this policy, Receive Queues can be set to 1-16. Completion Queues = Transmit Queues + Receive Queues. Interrupts = Completion Queues + 2. For more information, see [Cisco UCS Manager Network Management Guide, Release 4.1, Network-Related Policies](#).



Although previous versions of this document set the Ring Sizes for the Transmit and Receive Queues to 4096, [Tuning Guidelines for Cisco UCS Virtual Interface Cards](#) states that the sizes should be increased only if packet drops are observed on the vNIC interfaces.

6. Expand Options and choose Enabled for Receive Side Scaling (RSS).

Options

Transmit Checksum Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled	
Receive Checksum Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled	
TCP Segmentation Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled	
TCP Large Receive Offload	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled	
Receive Side Scaling (RSS)	:	<input type="radio"/> Disabled <input checked="" type="radio"/> Enabled	
Accelerated Receive Flow Steering	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Network Virtualization using Generic Routing Encapsulation	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Virtual Extensible LAN	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
GENEVE	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
AzureStack-Host QoS	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Fallback Timeout (Seconds)	:	<input type="text" value="5"/>	[0-600]
Interrupt Mode	:	<input checked="" type="radio"/> MSI X <input type="radio"/> MSI <input type="radio"/> IN Tx	
Interrupt Coalescing Type	:	<input checked="" type="radio"/> Min <input type="radio"/> Idle	
Interrupt Timer (us)	:	<input type="text" value="125"/>	[0-65535]
RoCE	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Advance Filter	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	
Interrupt Scaling	:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled	

7. Click OK, then click OK again to complete creating the Ethernet Adapter Policy.

Create LAN Connectivity Policy for FC Boot (FCP)

To configure the necessary Infrastructure LAN Connectivity Policy within the FlashStack-VSI Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > FlashStack-VSI.
3. Under the FlashStack-VSI Organization, right-click LAN Connectivity Policies.
4. Choose Create LAN Connectivity Policy.
5. Enter FCP-Boot as the name of the policy.
6. Click OK then OK again to add the policy.

7. In the menu on the left under LAN > Policies > root > Sub-Organizations > FlashStack-VSI > LAN Connectivity Policies, choose FC-Boot.
8. Click the Add button to add a vNIC.
9. In the Create vNIC dialog box, enter 00-vSwitch0-A as the name of the vNIC.
10. Choose the Use vNIC Template checkbox.
11. In the vNIC Template list, choose vSwitch0-A.
12. In the Adapter Policy list, choose VMWare.

Create vNIC



Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

13. Click OK to add this vNIC to the policy.

-
14. Click Save Changes and then click OK.
 15. Click Add to add another vNIC to the policy.
 16. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
 17. Check the box for the Use vNIC Template.
 18. In the vNIC Template list, choose vSwitch0-B.
 19. In the Adapter Policy list, choose VMWare.
 20. Click OK to add the vNIC to the policy.
 21. Click Save Changes and then click OK.
 22. Click Add to add another vNIC to the policy.
 23. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.
 24. Choose the Use vNIC Template checkbox.
 25. In the vNIC Template list, choose vDS0-A.
 26. In the Adapter Policy list, choose VMWare-HighTrf.



The VMware Adapter Policy can also be selected for this vNIC.

27. Click OK to add this vNIC to the policy.
28. Click Save Changes and then click OK.
29. Click Add to add another vNIC to the policy.
30. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
31. Choose the Use vNIC Template checkbox.
32. In the vNIC Template list, choose vDS0-B.
33. In the Adapter Policy list, choose VMWare-HighTrf.



Choose the same Adapter Policy that was selected for 02-Infra-vDS-A.

34. Click OK to add this vNIC to the policy.

35. Click Save Changes and then click OK.

General Events

Actions

Delete

Show Policy Usage

Use Global

Name : **FCP-Boot**

Description :

Owner : **Local**

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address
▶ vNIC 00-vSwitch0-A	Derived
▶ vNIC 01-vSwitch0-B	Derived
▶ vNIC 02-VDS-A	Derived
▶ vNIC 03-VDS-B	Derived

Delete Add Modify

+ Add iSCSI vNICs

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment in the FlashStack-VSI Organization, follow these steps:



Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click Servers.
2. Expand Pools > root > Sub-Organizations > FlashStack-VSI.
3. Right-click Server Pools under the FlashStack-VSI Organization.
4. Choose Create Server Pool.
5. Enter Infra-Pool as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Choose three (or more) servers to be used for the VMware management cluster and click >> to add them to the Infra-Pool server pool.



Although the VMware minimum host cluster size is two, in most use cases three servers are recommended.

9. Click Finish.

10. Click OK.

Create UUID Suffix Pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.

2. Expand Pools > root.

3. Right-click UUID Suffix Pools.

4. Choose Create UUID Suffix Pool.

5. Enter UUID-Pool as the name of the UUID suffix pool.

6. Optional: Enter a description for the UUID suffix pool.

7. Keep the prefix at the derived option.

8. Choose Sequential for the Assignment Order.

9. Click Next.

10. Click Add to add a block of UUIDs.

11. Keep the From field at the default setting.

12. Specify a size for the UUID block that is sufficient to support the available blade or server resources and the number of Service Profiles that will be created.

13. Click OK.

14. Click Finish.

15. Click OK.

Modify Default Host Firmware Package

Firmware management policies allow the administrator to choose the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To modify the default firmware management policy in the Cisco UCS environment, follow these steps:

1. Choose version 4.2(1f) for both the Blade and Rack Packages.
2. In Cisco UCS Manager, click Servers.
3. Expand Policies > root.
4. Expand Host Firmware Packages.
5. Choose default.
6. In the Actions pane, choose Modify Package Versions.

Modify Package Versions ×

Blade Package :

Rack Package :

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

<input type="checkbox"/>	Adapter
<input type="checkbox"/>	BIOS
<input type="checkbox"/>	Board Controller
<input type="checkbox"/>	CIMC
<input type="checkbox"/>	FC Adapters
<input type="checkbox"/>	Flex Flash Controller
<input type="checkbox"/>	GPUs
<input type="checkbox"/>	HBA Option ROM
<input type="checkbox"/>	Host NIC
<input type="checkbox"/>	Host NIC Option ROM
<input checked="" type="checkbox"/>	Local Disk
<input type="checkbox"/>	NVME Mswitch Firmware
<input type="checkbox"/>	PSU
<input type="checkbox"/>	Pci Switch Firmware

7. Click OK, then click OK again to modify the host firmware package.

Create Local Disk Configuration Policy (Optional)

A local disk configuration specifying no local disks for the Cisco UCS environment can be used to ensure that servers with no local disks are used for SAN Boot.



This policy should not be used on servers that contain local disks.

To create a local disk configuration policy, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Right-click Local Disk Config Policies.
4. Choose Create Local Disk Configuration Policy.
5. Enter IgnoreDisk as the local disk configuration policy name.
6. Change the mode to No Local Storage.

Create Local Disk Configuration Policy ? ×

Name :

Description :

Mode :

FlexFlash

FlexFlash State : Disable Enable

If **FlexFlash State** is disabled, SD cards will become unavailable immediately.
Please ensure SD cards are not in use before disabling the FlexFlash State.

FlexFlash RAID Reporting State : Disable Enable

FlexFlash Removable State : Yes No No Change

If **FlexFlash Removable State** is changed, SD cards will become unavailable temporarily.
Please ensure SD cards are not in use before changing the FlexFlash Removable State.

7. Click OK to create the local disk configuration policy.

8. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Right-click Power Control Policies.
4. Choose Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.

Create Power Control Policy



Name :

Description :

Fan Speed Policy :

Power Capping

If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping.

No Cap cap

Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority.

OK

Cancel

7. Click OK to create the power control policy.
8. Click OK.

Create Server Pool Qualification Policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, follow these steps:



This example creates a policy for Cisco UCS M6 servers for a server pool.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Choose Create Server Pool Policy Qualification.
5. Name the policy UCSB-B200M6.
6. Choose Create Server PID Qualifications.
7. Choose UCSB-B200-M6 from the PID drop-down list.

Create Server PID Qualifications ? ×

PID :

8. Click OK
9. Optionally, choose additional qualifications to refine server selection parameters for the server pool.
10. Click OK to create the policy then OK for the confirmation.

Update the Default Maintenance Policy

To update the default Maintenance Policy to either require user acknowledgement before server boot when service profiles change or to make the changes on the next server reboot, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root.
3. Choose Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.

5. Choose “On Next Boot” to delegate maintenance windows to server administrators.

Servers / Policies / root / Maintenance Policies / default

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : default

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs

Storage Config. Deployment Policy : Immediate User Ack

Reboot Policy : Immediate User Ack Timer Automatic

On Next Boot (Apply pending changes at next reboot.)

6. Click Save Changes.
7. Click OK to accept the changes.

Create vMedia Policy for VMware ESXi 7.0 U2 ISO Install Boot

The vMedia policy uses a HTTP web server, which is used for hosting VMware software. The vMedia Policy created will map the [Cisco Custom ISO for UCS 4.1.3a](#) to the Cisco UCS server in order to boot the ESXi installation. To create this policy, follow these steps:



The Cisco Custom ISO for UCS 4.1.3a should also be used for Cisco UCS software release 4.2(1f) and VMware vSphere 7.0 U2.

1. In Cisco UCS Manager, choose Servers.

-
2. Expand Policies > root.
 3. Right-click vMedia Policies.
 4. Choose Create vMedia Policy.
 5. Name the policy ESXi-7U2-CC-HTTP.
 6. Enter optional description in the Description field.
 7. Click Add to add a vMedia Mount.
 8. Name the mount ESXi-7U2-CC-HTTP.
 9. Choose the CDD Device Type.
 10. Choose the HTTP Protocol.
 11. Enter the IP Address of the web server.



To avoid any DNS lookup issues, enter the IP of the web server instead of the hostname.

12. Enter VMWare_ESXi_7.0.2_17867351_Custom_Cisco_4.1.3_a.iso as the Remote File name.



This VMware ESXi 7.0 U2 Cisco Custom ISO can be downloaded from VMware Downloads.



If a working vCenter 7.0 U2 installation is already in your environment, a FlashStack custom ISO for installing ESXi 7.0 U2 with all the necessary drivers for this FlashStack deployment can be created. See section [Create a FlashStack ESXi Custom ISO using VMware vCenter](#) in the Appendix for a procedure to build this custom ISO.

13. Enter the web server path to the ISO file in the Remote Path field.

Create vMedia Mount



Name	:	<input type="text" value="ESXi-7U2-CC-HTTP"/>
Description	:	<input type="text"/>
Device Type	:	<input checked="" type="radio"/> CDD <input type="radio"/> HDD
Protocol	:	<input type="radio"/> NFS <input type="radio"/> CIFS <input checked="" type="radio"/> HTTP <input type="radio"/> HTTPS
Hostname/IP Address	:	<input type="text" value="10.1.164.127"/>
Image Name Variable	:	<input checked="" type="radio"/> None <input type="radio"/> Service Profile Name
Remote File	:	<input type="text" value="VMware_ESXi_7.0.2_17867351_Custom_Cisco_4.1"/>
Remote Path	:	<input type="text" value="software/vSphere-7-Update-2"/>
Username	:	<input type="text"/>
Password	:	<input type="text"/>
Remap on Eject	:	<input type="checkbox"/>

OK

Cancel

14. Click OK to create the vMedia Mount.

15. Click OK then click OK again to complete creating the vMedia Policy.



For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot the host will boot into the ESXi installer since the SAN mounted disk is empty. After ESXi is installed, the vMedia will not be referenced if the boot disk is accessible.

Create Server BIOS Policies

To create a server BIOS policy for VMware ESXi hosts within the root organization, follow these steps:



Since the design supports Cisco UCS M5, UCS M6 and the Cisco UCS C4200 chassis with C125 AMD servers, three BIOS policies will be created as part of this procedure, one for Cisco AMD servers, one for Cisco UCS Intel M6 and one for Cisco UCS Intel M5 servers, the respective policy needs to be applied in the service profile depending on the server platform used. Cisco UCS Intel M6 policy will be used as default in this design.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Policies.
3. Right-click BIOS Policies.
4. Choose Create BIOS Policy.

5. Enter Intel-M6-Virt as the BIOS policy name.

Create BIOS Policy ? ×

Name :

Description :

Reboot on BIOS Settings Change :

6. Click OK, then click OK again to create the BIOS Policy.

7. Under the root Organization, expand BIOS Policies and choose the newly created BIOS Policy. Set the following within the Main tab of the Policy:

a. CDN Control > Enabled

b. Quiet Boot > Disabled

Servers / Policies / root / BIOS Policies / Intel-M6-Virt

Main | Advanced | Boot Options | Server Management | Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : **Intel-M6-Virt**
Description :
Owner : **Local**
Reboot on BIOS Settings Change :

Advanced Filter | Export | Print

BIOS Setting	Value
PCIe Slots CDN Control	Platform Default
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

Main | Advanced | Boot Options | Server Management | Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : **Intel-M6-Virt**
Description :
Owner : **Local**
Reboot on BIOS Settings Change :

Advanced Filter | Export | Print

BIOS Setting	Value
PCIe Slots CDN Control	Platform Default
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

8. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Scroll down and set the following within the Processor tab:
- Enhanced CPU Performance > Auto
 - Energy Efficient Turbo > Enabled
 - Sub NUMA Clustering > Enabled
 - Processor C1E > Enabled
 - Processor C6 Report > Enabled
 - UPI Prefetch > Enabled
 - LLC Prefetch > Disabled
 - UPI Link Speed > Auto
 - XPT Prefetch > Enabled
 - Patrol Scrub > Disabled
 - UPI Link Enablement > 1
 - UPI Power Management > Enabled
 - Workload Configuration > Balanced



BIOS Setting	Value
Boot Performance mode	Platform Default
CPU Performance	Platform Default
Configurable TDP Level	Platform Default
Core Multi Processing	Platform Default
Enhanced CPU Performance	Auto
DCPMM Firmware Downgrade	Platform Default
DRAM Clock Throttling	Platform Default
Direct Cache Access	Platform Default
Energy Performance Tuning	Platform Default
Enhanced Intel SpeedStep Tech	Platform Default
Execute Disable Bit	Platform Default
Frequency Floor Override	Platform Default
Intel HyperThreading Tech	Platform Default
Energy Efficient Turbo	Enabled
Intel Turbo Boost Tech	Platform Default
Intel Virtualization Technology	Platform Default
Intel Dynamic Speed Select	Platform Default
Intel Speed Select	Platform Default
Channel Interleaving	Platform Default
IMC Inteleave	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Sub NUMA Clustering	Enabled
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default



BIOS Setting	Value
Processor C State	Platform Default
Processor C1E	Enabled
Processor C3 Report	Platform Default
Processor C6 Report	Enabled
Processor C7 Report	Platform Default
Processor CMC1	Platform Default
Power Technology	Platform Default
Energy Performance	Platform Default
Processor EPP Enable	Platform Default
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Enabled
LLC Prefetch	Disabled
UPI Link Speed	Auto
XPT Prefetch	Enabled
Burst and Postponed Refresh	Platform Default

Main **Advanced** Boot Options Server Management Events

Processor Intel Directed IO RAS Memory Serial Port USB PCI QPI LOM and PCIe Slots Trusted Platform Graphics Configuration

Advanced Filter Export Print

BIOS Setting	Value
Multikey Total Memory Encryption (MK-TME)	Platform Default
SW Guard Extensions (SGX)	Platform Default
Total Memory Encryption (TME)	Platform Default
Select Owner EPOCH input type	Platform Default
Operation Mode	Platform Default
SEV	Platform Default
SMEE	Platform Default
SProcessor Epoch 0	Platform Default [0-ffffffff] [Step Value: 1]
SProcessor Epoch 1	Platform Default [0-ffffffff] [Step Value: 1]
SGX Factory Reset	Platform Default
SGX PUBKEY HASH0	Platform Default [0-ffffffff] [Step Value: 1]
SGX PUBKEY HASH1	Platform Default [0-ffffffff] [Step Value: 1]
SGX PUBKEY HASH2	Platform Default [0-ffffffff] [Step Value: 1]
SGX PUBKEY HASH3	Platform Default [0-ffffffff] [Step Value: 1]
SGX Write Enable	Platform Default
SGX Pkg info In-Band Access	Platform Default
SGX QoS	Platform Default
SMT Mode	Platform Default
SVM Mode	Platform Default
TSME	Platform Default
SGX Auto MP Registration Agent	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Disabled
UPI Link Enablement	1
UPI Power Management	Enabled
Uncore Frequency Scaling	Platform Default
Workload Configuration	Balanced

9. Click Save changes.

10. Click the Advanced tab, leaving the RAS Memory tab selected within the Advanced tab. Scroll down and set the following within the Processor tab:

- a. Panic and High Watermark > High
- b. Memory Refresh Rate > 1x
- c. Partial Cache Line Sparing > Disabled
- d. Memory Thermal Throttling Mode > Disabled
- e. NVM Performance Setting > Balanced Profile

11. Click Save Changes.

12. Click OK.

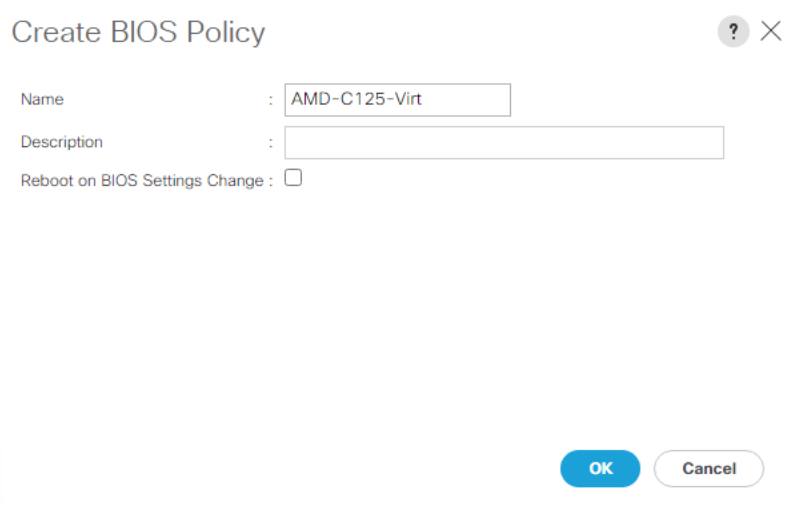
13. In Cisco UCS Manager, click Servers.

14. Expand Policies > root > Policies.

15. Right-click BIOS Policies.

16. Choose Create BIOS Policy.

17. Enter AMD-C125-Virt as the BIOS policy name.



Create BIOS Policy ? ×

Name : AMD-C125-Virt

Description :

Reboot on BIOS Settings Change :

OK Cancel

18. Click OK, then click OK again to create the BIOS Policy.

19. Under the root Organization, expand BIOS Policies and choose the newly created BIOS Policy. Set the following within the Main tab of the Policy:

- a. CDN Control > Enabled
- b. Quiet Boot > Disabled

Main | **Advanced** | Boot Options | Server Management | Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : **AMD-C125-Virt**
Description :
Owner : **Local**
Reboot on BIOS Settings Change :

Advanced Filter | Export | Print

BIOS Setting	Value
PCIe Slots CDN Control	Platform Default
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

20. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Scroll down and set the following within the Processor tab:

a. Determinism slider > Power

Main | **Advanced** | Boot Options | Server Management | Events

Processor | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform | Graphics Configuration

Advanced Filter | Export | Print

BIOS Setting	Value
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
UPI Link Speed	Platform Default
XPT Prefetch	Platform Default
Burst and Postponed Refresh	Platform Default
Core Performance Boost	Platform Default
Downcore control	Platform Default
Global C-state Control	Platform Default
L1 Stream HW Prefetcher	Platform Default
L2 Stream HW Prefetcher	Platform Default
Determinism Slider	Power
IOMMU	Platform Default
Bank Group Swap	Platform Default

21. Click Save Changes.

22. Click OK.

23. In Cisco UCS Manager, click Servers.

24. Expand Policies > root > Policies.

25. Right-click BIOS Policies.

26. Choose Create BIOS Policy.

27. Enter Intel-M5-Virt as the BIOS policy name.

Create BIOS Policy ? ×

Name :

Description :

Reboot on BIOS Settings Change :

28. Click OK, then click OK again to create the BIOS Policy.

29. Under the root Organization, expand BIOS Policies and choose the newly created BIOS Policy. Set the following within the Main tab of the Policy:

- a. CDN Control > Enabled
- b. Quiet Boot > Disabled

Main | Advanced | Boot Options | Server Management | Events

Actions

Delete
Show Policy Usage
Use Global

Properties

Name : **Intel-M5-Virt**
Description :
Owner : **Local**
Reboot on BIOS Settings Change :

Advanced Filter | Export | Print |

BIOS Setting	Value
PCIe Slots CDN Control	Platform Default
CDN Control	Enabled
Front panel lockout	Platform Default
POST error pause	Platform Default
Quiet Boot	Disabled
Resume on AC power loss	Platform Default

30. Click the Advanced tab, leaving the Processor tab selected within the Advanced tab. Scroll down and set the following within the Processor tab:

- a. Processor C State > Disabled
- b. Processor C1E > disabled
- c. Processor C3 Report > disabled
- d. Processor C6 Report > disabled
- e. Processor C7 Report > disabled
- f. Power Technology > Custom

Main | **Advanced** | Boot Options | Server Management | Events

Processor | Intel Directed IO | RAS Memory | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform | Graphics Configuration

Advanced Filter | Export | Print

BIOS Setting	Value
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Disabled
Processor C7 Report	Disabled
Processor CMC1	Platform Default
Power Technology	Custom
Energy Performance	Platform Default

31. Click the RAS Memory tab and select:

a. NVM Performance setting > Balanced Profile

Servers / Policies / root / BIOS Policies / Intel-M5-Virt

Main | **Advanced** | Boot Options | Server Management | Events

Processor | Intel Directed IO | **RAS Memory** | Serial Port | USB | PCI | QPI | LOM and PCIe Slots | Trusted Platform | Graphics Configuration

Advanced Filter | Export | Print

BIOS Setting	Value
DRAM Refresh Rate	Platform Default
eADR Support	Platform Default
LLC Dead Line	Platform Default
LV DDR Mode	Platform Default
Memory Refresh Rate	Platform Default
Memory Thermal Throttling Mode	Platform Default
Memory Bandwidth Boost	Platform Default
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
NVM Performance Setting	Balanced Profile
Panic and High Watermark	Platform Default

32. Click Save Changes.

33. Click OK.



For more information, see:

[Performance Tuning Guide for Cisco UCS M6 Servers](#)

[Performance Tuning Guide for Cisco UCS M5 Servers](#)

[Performance Tuning for Cisco UCS C125 Rack Server Nodes with AMD Processors.](#)

Create Persistent Memory Policies (Optional)

To create a persistent memory policy for VMware ESXi hosts within the root organization, follow these steps:



Two persistent memory policies will be created as part of this procedure, one for the Intel Optane Persistent Memory modules to be used in memory mode for 100% volatile memory, the other for the modules to be used in App direct mode for 100% persistent memory.

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Policies.
3. Right-click Persistent memory Policy.
4. Choose Create Persistent Memory Policy.
5. Enter App-Direct-Mode as the policy name.
6. To create a goal, click the Add button in the Goals area of the Create Persistent Memory Policy dialog box.
7. Click OK, the default mode is App Direct mode with Memory Mode (%) value 0.

Create Persistent Memory Policy



Properties

Name :

Description :

General

Security

Goals

Advanced Filter Export Print

Socket Id	Memory Mode (%)	Persistent Memory Type
All Sockets	0	App Direct

Add Delete Modify

Configure Namespace


Advanced Filter Export Print

Name	Socket Id	Socket Local DIMM ...	Mode	Capacity (GiB)
No data available				

OK

Cancel

- Click OK again.
- Right-click Persistent memory Policy.
- Choose Create Persistent Memory Policy.
- Enter Memory-Mode as the policy name.
- To create a goal, click the Add button in the Goals area of the Create Persistent Memory Policy dialog box.
- In the Create Goal dialog box, enter Memory Mode (%) value as 100



14. Click OK and click OK again.

Create Persistent Memory Policy ? X

Properties

Name :

Description :

General Security

Goals

Advanced Filter Export Print ⚙

Socket Id	Memory Mode (%)	Persistent Memory Type
All Sockets	0	App Direct

+ Add Delete Modify

Configure Namespace

Advanced Filter Export Print ⚙

Name	Socket Id	Socket Local DIMM ...	Mode	Capacity (GiB)
No data available				

OK Cancel

15. Click OK again.

Create FC Boot Policy (FCP)

This procedure applies to a Cisco UCS environment in which two Fibre Channel interfaces used are on Pure FlashArray controller 0 (CT0.FC0 and CT0.FC2) and two on controller 1 (CT1.FC0 and CT1.FC2). Additional FC ports can be used as needed for more bandwidth. With FC0 and FC2 being used for FCP protocol, the other two ports can be used for FC-NVMe if required.

Collect the WWPNs from each controller on the Pure Storage FlashArray that are visible from the Network tab under the Settings section of the FlashArray Web GUI.

Settings					
System Network Access Software					
Fibre Channel ^ 18 of 8					
Name	Enabled	WWN	Speed	Services	
CT0.FC0	true	52:4A:93:77:DE:D7:21:00	32 Gb/s	scsi-fc	<input checked="" type="checkbox"/>
CT0.FC1	true	52:4A:93:77:DE:D7:21:01	32 Gb/s	nvme-fc	<input checked="" type="checkbox"/>
CT0.FC2	true	52:4A:93:77:DE:D7:21:02	32 Gb/s	scsi-fc	<input checked="" type="checkbox"/>
CT0.FC3	true	52:4A:93:77:DE:D7:21:03	32 Gb/s	nvme-fc	<input checked="" type="checkbox"/>
CT1.FC0	true	52:4A:93:77:DE:D7:21:10	32 Gb/s	scsi-fc	<input checked="" type="checkbox"/>
CT1.FC1	true	52:4A:93:77:DE:D7:21:11	32 Gb/s	nvme-fc	<input checked="" type="checkbox"/>
CT1.FC2	true	52:4A:93:77:DE:D7:21:12	32 Gb/s	scsi-fc	<input checked="" type="checkbox"/>
CT1.FC3	true	52:4A:93:77:DE:D7:21:13	32 Gb/s	nvme-fc	<input checked="" type="checkbox"/>

As an alternative to the GUI, connect to the FlashArray//X via ssh using the pureuser account and find the WWNs using the pureport list command:

```
pureuser@BB08-FlashArrayR3> pureport list
Name      WWN      Portal      IQN      NON      Failover
CT0.ETH4  -        -           -        -        -
CT0.ETH5  -        192.168.102.146:3260  iqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT0.FC0   52:4A:93:77:DE:D7:21:00  -           -        -        -
CT0.FC1   52:4A:93:77:DE:D7:21:01  -           -        nqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT0.FC2   52:4A:93:77:DE:D7:21:02  -           -        nqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT0.FC3   52:4A:93:77:DE:D7:21:03  -           -        nqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT1.ETH4  -        192.168.102.147:3260  iqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT1.ETH5  -        192.168.101.147:3260  iqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT1.FC0   52:4A:93:77:DE:D7:21:10  -           -        -        -
CT1.FC1   52:4A:93:77:DE:D7:21:11  -           -        nqn.2010-06.com.purestorage:flasharray.779962553908b056  -
CT1.FC2   52:4A:93:77:DE:D7:21:12  -           -        -        -
CT1.FC3   52:4A:93:77:DE:D7:21:13  -           -        nqn.2010-06.com.purestorage:flasharray.779962553908b056  -
```

Find the FC0 ports for each controller from within the System view and record the values to be used for Primary and Secondary Targets. In the example lab environment, these appear as the first ports on the right side of each controller shown.

FlashArray Controller	FC Port	Primary or Secondary path	WWPN
FlashArray//X Controller 0	CT0.FC0	Primary	52:4A:93:77:DE:D7:21:00
FlashArray//X Controller 1	CT1.FC0	Secondary	52:4A:93:77:DE:D7:21:10

Within the same System view, find the FC2 ports for each controller and record the values to be used for Primary and Secondary Targets. In the example lab environment, these appear as the second ports on the right side of each controller shown.

FlashArray Controller	FC Port	Primary or Secondary path	WWPN
FlashArray//X Controller 0	CT0.FC2	Primary	52:4A:93:77:DE:D7:21:02
FlashArray//X Controller 1	CT1.FC2	Secondary	52:4A:93:77:DE:D7:21:12



One boot policy is configured in this procedure.

To create a boot policy within the FlashStack-VSI organization, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > FlashStack-VSI.
3. Under the FlashStack-VSI Organization, right-click Boot Policies.
4. Choose Create Boot Policy.
5. Enter Boot-FCP as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not select the Reboot on Boot Order Change checkbox.
8. Choose the Uefi Boot Mode.
9. Choose the Boot Security checkbox.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descripti...
------	-------	------------	------	-----------	-----	------------	------------	-----------	--------------



UEFI Secure Boot can be used to boot VMware ESXi 7.0 U2 with or without a TPM 2.0 module in the UCS server.

10. Expand Local Devices and choose Add Remote CD/DVD.
11. Expand vHBAs and choose Add SAN Boot.

12. Choose Primary for the Type field.

13. Enter FCP-Fabric-A in the vHBA field.

Add SAN Boot ? ×

vHBA :

Type : Primary Secondary Any

14. Click OK.

15. From vHBAs, choose Add SAN Boot Target.

16. Keep 1 as the value for Boot Target LUN.

17. Enter the WWPN for CT0.FC0.

18. Choose Primary for the SAN boot target type.

Add SAN Boot Target ? ×

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

19. Click OK to add the SAN boot target.

20. From vHBAs, choose Add SAN Boot Target.

21. Enter 1 as the value for Boot Target LUN.

22. Enter the WWPN for CT1.FC0.

-
23. Click OK to add the SAN boot target.
 24. From vHBAs, choose Add SAN Boot.
 25. In the Add SAN Boot dialog box, enter FCP-Fabric-B in the vHBA box.
 26. The SAN boot type should automatically be set to Secondary.
 27. Click OK.
 28. From vHBAs, choose Add SAN Boot Target.
 29. Keep 1 as the value for Boot Target LUN.
 30. Enter the WWPN for CT0.FC2.
 31. Choose Primary for the SAN boot target type.
 32. Click OK to add the SAN boot target.
 33. From vHBAs, choose Add SAN Boot Target.
 34. Keep 1 as the value for Boot Target LUN.
 35. Enter the WWPN for CT1.FC2.
 36. Click OK to add the SAN boot target.
 37. Expand CIMC Mounted Media and choose Add CIMC Mounted CD/DVD.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

- Local Devices
- CIMC Mounted vMedia
 - Add CIMC Mounted CD/DVD
 - Add CIMC Mounted HDD
- vNICs
- vHBAs
 - Add SAN Boot
 - Add SAN Boot Target
- iSCSI vNICs
- EFI Shell

Boot Order

+ - Advanced Filter ↑ Export Print ⚙									
Name	Order	vNIC/vH...	Type	LUN Na...	WWN	Slot Nu...	Boot Na...	Boot Path	Descript...
Rem...	1								
▼ San	2								
▶ S...		Fabric-A	Primary						
▶ S...		Fabric-B	Second...						
CIM...	3								

↑ Move Up ↓ Move Down 🗑 Delete

Set UEFI Boot Parameters

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Local Devices

CIMC Mounted vMedia

vNICs

vHBAs

Add SAN Boot

Add SAN Boot Target

iSCSI vNICs

EFI Shell

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vHBA/iSCSI ...	Type	LUN N...	WWN	Slot Nu...	Boot N...	Boot P...	Descrip...
Re...	1								
San	2								
▶ S...		FCP-Fabric-A	Primary						
▶ S...		FCP-Fabric-B	Secon...						
CIM...	3								

Move Up Move Down Delete

Set Uefi Boot Parameters

OK Cancel

38. Expand San > SAN Primary and select SAN Target Primary. Select Set Uefi Boot Parameters.



For Cisco UCS B200 M6 and M5, and Cisco UCS C220 M6 and M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for Cisco UCS M4 and earlier servers, VMware ESXi 7.0 and above will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

39. Fill in the Set Uefi Boot Parameters exactly as shown below:

Set Uefi Boot Parameters



Uefi Boot Parameters

Boot Loader Name	:	<input type="text" value="BOOTX64.EFI"/>
Boot Loader Path	:	<input "="" type="text" value="\EFI\BOOT\"/>
Boot Loader Description	:	<input type="text"/>



40. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target and click OK for the confirmation.
41. Repeat this process to set Uefi Boot Parameters for each of the 4 SAN Boot Targets.
42. Click OK, then click OK again to create the boot policy.

Create Service Profile Template (FCP)

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot within the FlashStack-VSI Organization. To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > FlashStack-VSI.
3. Right-click the FlashStack-VSI Organization.
4. Choose Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-FCP as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Choose the Updating Template option.
7. Under UUID, choose UUID_Pool as the UUID pool.

Create Service Profile Template ? X

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-FlashStack-VSI**

The template will be created in the following organization. Its name must be unique within this organization.
Type : Initial Template Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.
UUID

UUID Assignment:

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

8. Click Next.

Configure Storage Provisioning

To configure storage provisioning, follow these steps:

1. If you have servers with no physical disks in the UCS chassis, click on the Local Disk Configuration Policy tab and choose the IgnoreDisk Local Storage Policy. Otherwise, choose the default Local Storage Policy.
2. Click Next.

Configure Networking

To configure networking, follow these steps:

1. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.
2. Choose FC-Boot from the LAN Connectivity Policy drop-down list.
3. Leave the Initiator Name Assignment at <not set>.

Create Service Profile Template ? X

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: ▼

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

LAN Connectivity Policy : ▼ [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: ▼

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

4. Click Next.

Configure SAN Connectivity

To configure SAN connectivity, follow these steps:

1. Choose the Use Connectivity Policy option for the “How would you like to configure SAN connectivity?” field.
2. Choose the FC-Boot option from the SAN Connectivity Policy drop-down list.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

Simple Expert No vHBAs Use Connectivity Policy

SAN Connectivity Policy : [Create SAN Connectivity Policy](#)

< Prev Next > **Finish** Cancel

3. Click Next.

Configure Zoning

To configure zoning, follow this step:

1. Set no zoning options and click Next.



Set no zoning options here since the fabric interconnects are in end host (NPV) mode and zoning is being done in the upstream SAN switch.

Configure vNIC/HBA Placement

To configure vNIC/HBA placement, follow these steps:

1. In the Select Placement list, retain the placement policy as Let System Perform Placement.
2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

1. Do not select a vMedia Policy.
2. Click Next.

Configure Server Boot Order

To configure the server boot order, follow these steps:

1. Choose Boot-FCP-A for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **Boot-FCP**

Description :

Reboot on Boot Order Change : **No**

Enforce vNIC/vHBA/iSCSI Name : **Yes**

Boot Mode : **Uefi**

Boot Security : **Yes**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHB...	Type	LUN Name	WWN	Slot Numb...	Boot Name	Boot Path	Description
Remot...	1								
San	2								
CIMC ...	3								

[Create ISCSI vNIC](#) [Set ISCSI Boot Parameters](#) [Set Uefi Boot Parameters](#)

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

2. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

⊖ Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: [Create Maintenance Policy](#)

Name : **default**

Description :

Soft Shutdown Timer : **150 Secs**

Storage Config. Deployment Policy : **User Ack**

Reboot Policy : **User Ack**

< Prev Next > **Finish** Cancel

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, choose Infra-Pool.
2. Choose Down as the power state to be applied when the profile is associated with the server.
3. Optional: Choose “UCSB-B200-M6” for the Server Pool Qualification to choose only UCS M6 servers in the pool.
4. Expand Firmware Management and choose the default Host Firmware Package.

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment:

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification :

Restrict Migration :

+ Firmware Management (BIOS, Disk Controller, Adapter)

< Prev Next > Finish Cancel

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, choose Intel-M6-Virt.
2. Expand Power Control Policy Configuration and choose No-Power-Cap in the Power Control Policy list.

Create Service Profile Template ? X

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy :

External IPMI/Redfish Management Configuration

Management IP Address

Monitoring Configuration (Thresholds)

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : [Create Power Control Policy](#)

Scrub Policy

KVM Management Policy

Graphics Card Policy

3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create vMedia-Enabled Service Profile Template

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click VM-Host-Infra-FCP and choose Create a Clone.
4. Name the clone VM-Host-Infra-FCP-vM.
5. Click OK then click OK again to create the Service Profile Template clone.

-
6. Choose the newly created VM-Host-Infra-FCP-VM and choose the vMedia Policy tab.
 7. Click Modify vMedia Policy.
 8. Choose the ESXi-7U2-CC-HTTP vMedia Policy and click OK.
 9. Click OK to confirm.

Create Intel Optane Memory Mode Service Profile Template (Optional)

To create a service profile template with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click VM-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-MM-Host-Infra-FCP.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-MM-Host-Infra-FCP and choose the Policies tab.
7. Expand Persistent Memory Policy and use the pulldown to select the Memory-Mode Policy.
8. Click save Changes.
9. Click OK to confirm.

Create vMedia-Enabled Intel Optane Memory Mode Service Profile Template (Optional)

To create a service profile template with vMedia enabled for servers with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click Intel-MM-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-MM-Host-Infra-FCP-VM.
5. Click OK then click OK again to create the Service Profile Template clone.

-
6. Choose the newly created Intel-MM-Host-Infra-FCP-vM and choose the vMedia Policy tab.
 7. Click Modify vMedia Policy.
 8. Choose the ESXi-7U2-CC-HTTP vMedia Policy and click OK.
 9. Click OK to confirm.

Create Intel Optane App Direct Mode Service Profile Template (Optional)

To create a service profile template with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click VM-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-AD-Host-Infra-FCP.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-AD-Host-Infra-FCP and choose the Policies tab.
7. Expand Persistent Memory Policy and use the pulldown to select the Memory-Mode Policy.
8. Click save Changes.
9. Click OK to confirm.

Create vMedia-Enabled Intel Optane App Direct Mode Service Profile Template (Optional)

To create a service profile template with vMedia enabled for servers with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

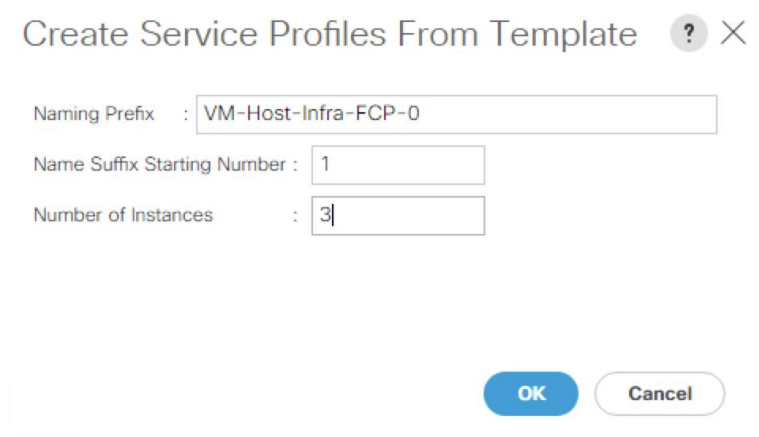
1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click Intel-AD-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-AD-Host-Infra-FCP-vM.
5. Click OK then click OK again to create the Service Profile Template clone.

6. Choose the newly created Intel-AD-Host-Infra-FCP-VM and choose the vMedia Policy tab.
7. Click Modify vMedia Policy.
8. Choose the ESXi-7U2-CC-HTTP vMedia Policy and click OK.
9. Click OK to confirm.

Create Service Profiles

To create service profiles from the service profile template within the FlashStack-VSI Organization, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack > Service Template VM-Host-Infra-FCP-VM.
3. Right-click VM-Host-Infra-FCP-VM and choose Create Service Profiles from Template.
4. Enter VM-Host-Infra-FCP-0 as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 3 as the “Number of Instances.”



Create Service Profiles From Template ? X

Naming Prefix : VM-Host-Infra-FCP-0

Name Suffix Starting Number : 1

Number of Instances : 3

OK Cancel

7. Click OK to create the service profiles.
8. Click OK in the confirmation message.
9. When VMware ESXi 7.0 U2 has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-FCP Service Profile Template to remove the vMedia Mapping from the host.

Add More Servers to FlashStack Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlashStack unit. All pools and policies created at the organizational level will need to be recreated within other organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlashStack deployment, specific information must be gathered from each Cisco UCS server and from the Pure FlashArray controllers.

Table 13. WWPNs from Pure FlashArray//X R3 Storage

FlashArray	Adapter	MDS Switch	Target: WWPN
BB08-FlashArray//X-R3	CT0.FC0	Fabric A	<CT0.FC0-wwpn>
	CT0.FC2	Fabric B	<CT0.FC2-wwpn>
	CT1.FC0	Fabric A	<CT1.FC0-wwpn>
	CT1.FC2	Fabric B	<CT1.FC2-wwpn>

Table 14. WWPNs for Cisco UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
VM-Host-Infra-FCP-01	Fabric A	<vm-host-infra-fcp-01-wwpna>
	Fabric B	<vm-host-infra-fcp-01-wwpnb>
VM-Host-Infra-FCP-02	Fabric A	<vm-host-infra-fcp-02-wwpna>
	Fabric B	<vm-host-infra-fcp-02-wwpnb>
VM-Host-Infra-FCP-03	Fabric A	<vm-host-infra-fcp-03-wwpna>
	Fabric B	<vm-host-infra-fcp-03-wwpnb>




To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root > Sub-Organizations > Organization. Expand each service profile and then expand vHBAs. Select each vHBA. The WWPN is shown under Properties on the right.

SAN Switch Configuration

This section explains how to configure the Cisco MDS 9000s for use in a FlashStack environment.



Follow the steps precisely because failure to do so could result in an improper configuration.



If directly connecting storage to the Cisco UCS fabric interconnects, skip this section.

Physical Connectivity

Follow the physical connectivity guidelines for FlashStack as explained in section [FlashStack Cabling](#).

FlashStack Cisco MDS Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.5(1a).

Cisco MDS 9132T A

To set up the initial configuration for the Cisco MDS A switch, <mds-A-hostname>, follow these steps:

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

```
----- System Admin Account Setup -----

Do you want to enforce secure password standard (yes/no) [y]: Enter

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

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name : <mds-A-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address : <mds-A-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-A-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway : <mds-A-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: Enter

Configure clock? (yes/no) [n]: Enter

Configure timezone? (yes/no) [n]: Enter

Configure summertime? (yes/no) [n]: Enter

Configure the ntp server? (yes/no) [n]: yes

NTP server IPv4 address : <nexus-A-mgmt0-ip>

Configure default switchport interface state (shut/noshut) [shut]: Enter

Configure default switchport trunk mode (on/off/auto) [on]: auto

Configure default switchport port mode F (yes/no) [n]: yes

Configure default zone policy (permit/deny) [deny]: Enter

Enable full zoneset distribution? (yes/no) [n]: Enter
```

```
Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
```

```
Use this configuration and save it? (yes/no) [y]: Enter
```

Cisco MDS 9132T B

To set up the initial configuration for the Cisco MDS B switch, <mds-B-hostname>, follow these steps:

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start and attempt to enter Power on Auto Provisioning. Enter y to get to the System Admin Account Setup.

1. Configure the switch using the command line.

```
----- System Admin Account Setup -----

Do you want to enforce secure password standard (yes/no) [y]: Enter

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

Would you like to enter the basic configuration dialog (yes/no): yes

Create another login account (yes/no) [n]: Enter

Configure read-only SNMP community string (yes/no) [n]: Enter

Configure read-write SNMP community string (yes/no) [n]: Enter

Enter the switch name : <mds-B-hostname>

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter

Mgmt0 IPv4 address : <mds-B-mgmt0-ip>

Mgmt0 IPv4 netmask : <mds-B-mgmt0-netmask>

Configure the default gateway? (yes/no) [y]: Enter

IPv4 address of the default gateway : <mds-B-mgmt0-gw>

Configure advanced IP options? (yes/no) [n]: Enter

Enable the ssh service? (yes/no) [y]: Enter

Type of ssh key you would like to generate (dsa/rsa) [rsa]: Enter

Number of rsa key bits <1024-2048> [1024]: Enter

Enable the telnet service? (yes/no) [n]: Enter

Configure congestion/no_credit drop for fc interfaces? (yes/no) [y]: Enter

Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]: Enter

Enter milliseconds in multiples of 10 for congestion-drop for logical-type edge
in range (<200-500>/default), where default is 500. [d]: Enter
```

```
Enable the http-server? (yes/no) [y]: Enter
Configure clock? (yes/no) [n]: Enter
Configure timezone? (yes/no) [n]: Enter
Configure summertime? (yes/no) [n]: Enter
Configure the ntp server? (yes/no) [n]: yes
NTP server IPv4 address : <nexus-A-mgmt0-ip>
Configure default switchport interface state (shut/noshut) [shut]: Enter
Configure default switchport trunk mode (on/off/auto) [on]: auto
Configure default switchport port mode F (yes/no) [n]: yes
Configure default zone policy (permit/deny) [deny]: Enter
Enable full zoneset distribution? (yes/no) [n]: Enter
Configure default zone mode (basic/enhanced) [basic]: Enter
```

2. Review the configuration.

```
Would you like to edit the configuration? (yes/no) [n]: Enter
Use this configuration and save it? (yes/no) [y]: Enter
```

FlashStack Cisco MDS Switch Configuration

Enable Licenses

Cisco MDS 9132T A and Cisco MDS 9132T B

To enable the correct features on the Cisco MDS switches, follow these steps:

1. Log in as admin.
2. Run the following commands:

```
configure terminal
feature npiv
feature fport-channel-trunk
```

Add Second NTP Server and Local Time Configuration

Cisco MDS 9132T A and Cisco MDS 9132T B

To configure the second NTP server and add local time configuration, follow this step:

1. From the global configuration mode, run the following command:

```
ntp server <nexus-B-mgmt0-ip>
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>
```



It is important to configure the local time so that logging time alignment, any backup schedules, and SAN Analytics forwarding are correct. For more information on configuring the timezone and daylight savings time or summer time, please see [Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 8.x](#). Sample clock commands for the United States Eastern time-zone are:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

Configure Individual Ports

Cisco MDS 9132T A

To configure individual ports and port-channels for switch A, follow this step:

1. From the global configuration mode, run the following commands:

```
interface fc1/1
switchport description BB08-X50R3-ct0fc0
switchport speed 32000
switchport trunk mode off
no shutdown
```

```
exit

interface fc1/2
switchport description BB08-X50R3-ct1fc0
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/1
switchport description BB08-X50R3-ct0fc1
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description BB08-X50R3-ct1fc1
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description BB08-6454-A:fc1/1
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/6
switchport description BB08-6454-A:fc1/2
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/7
switchport description BB08-6454-A:fc1/3
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/8
switchport description BB08-6454-A:fc1/4
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface port-channel15
switchport mode F
switchport trunk allowed vsan 100
switchport description BB08-6454-A
switchport speed 32000
switchport rate-mode dedicated
switchport trunk mode auto
```

If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-a-id>” for interface port-channel15. Note also that the default setting of switchport trunk mode auto is being used for the port channel.

Cisco MDS 9132T B

To configure individual ports and port-channels for switch B, follow these steps:

1. From the global configuration mode, run the following commands:

```
interface fc1/1
switchport description BB08-X50R3-ct0fc2
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/2
switchport description BB08-X50R3-ct1fc2
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/3
switchport description BB08-X50R3-ct0fc3
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/4
switchport description BB08-X50R3-ct1fc3
switchport speed 32000
switchport trunk mode off
no shutdown
exit

interface fc1/5
switchport description BB08-6454-B:fc1/1
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/6
switchport description BB08-6454-B:fc1/2
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/7
switchport description BB08-6454-B:fc1/3
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown

interface fc1/8
switchport description BB08-6454-B:fc1/4
switchport trunk mode auto
port-license acquire
channel-group 15 force
no shutdown
```

If VSAN trunking is not being used between the Cisco UCS Fabric Interconnects and the MDS switches, do not enter “switchport trunk allowed vsan <vsan-b-id>” for interface port-channel15. Note also that the default setting of switchport trunk mode auto is being used for the port channel.

Create VSANs

Cisco MDS 9132T A

To create the necessary VSANs for fabric A and add ports to them, follow this step:

1. From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fcl/1
vsan <vsan-a-id> interface fcl/2
vsan <vsan-a-id> interface fcl/3
vsan <vsan-a-id> interface fcl/4
vsan <vsan-a-id> interface port-channel15
exit
```

Cisco MDS 9132T B

To create the necessary VSANs for fabric B and add ports to them, follow these steps:

1. From the global configuration mode, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fcl/1
vsan <vsan-b-id> interface fcl/2
vsan <vsan-b-id> interface fcl/3
vsan <vsan-b-id> interface fcl/4
vsan <vsan-b-id> interface port-channel15
exit
```

At this point, it may be necessary to go into Cisco UCS Manager and disable and enable the FC port-channel interfaces to get the port-channels to come up.

Create Device Aliases

Cisco MDS 9132T A

To create device aliases for Fabric A that will be used to create zones, follow this step:

1. From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
```

```
device-alias name FlashArray-CT0FC0 pwnn 52:4a:93:77:de:d7:21:00
device-alias name FlashArray-CT1FC0 pwnn 52:4a:93:77:de:d7:21:10
device-alias name FlashArray-CT0FC1 pwnn 52:4a:93:77:de:d7:21:01
device-alias name FlashArray-CT1FC1 pwnn 52:4a:93:77:de:d7:21:11
device-alias name VM-Host-Infra-FCP-01-A pwnn 20:00:00:25:b5:a4:0a:00
device-alias name VM-Host-Infra-FCP-02-A pwnn 20:00:00:25:b5:a4:0a:01
device-alias name VM-Host-Infra-FCP-03-A pwnn 20:00:00:25:b5:a4:0a:02
device-alias name VM-Host-Infra-FC-NVMe-01-A pwnn 20:00:00:25:b5:a4:0a:03
device-alias name VM-Host-Infra-FC-NVMe-02-A pwnn 20:00:00:25:b5:a4:0a:04
device-alias name VM-Host-Infra-FC-NVMe-03-A pwnn 20:00:00:25:b5:a4:0a:05
device-alias commit
```

Cisco MDS 9132T B

To create device aliases for Fabric B that will be used to create zones, follow this step:

1. From the global configuration mode, run the following commands:

```
device-alias mode enhanced
device-alias database
device-alias name FlashArray-CT0FC2 pwnn 52:4a:93:77:de:d7:21:02
device-alias name FlashArray-CT1FC2 pwnn 52:4a:93:77:de:d7:21:12
device-alias name FlashArray-CT0FC2 pwnn 52:4a:93:77:de:d7:21:03
device-alias name FlashArray-CT1FC2 pwnn 52:4a:93:77:de:d7:21:13
device-alias name VM-Host-Infra-FCP-01-B pwnn 20:00:00:25:b5:a4:0b:00
device-alias name VM-Host-Infra-FCP-02-B pwnn 20:00:00:25:b5:a4:0b:01
device-alias name VM-Host-Infra-FCP-03-B pwnn 20:00:00:25:b5:a4:0b:02
device-alias name VM-Host-Infra-FC-NVMe-01-B pwnn 20:00:00:25:b5:a4:0b:03
device-alias name VM-Host-Infra-FC-NVMe-02-B pwnn 20:00:00:25:b5:a4:0b:04
device-alias name VM-Host-Infra-FC-NVMe-03-B pwnn 20:00:00:25:b5:a4:0b:05
device-alias commit
```

Create Zones and Zoneset

Cisco MDS 9132T A

To create the required zones and zoneset on Fabric A, run the following commands:

```
configure terminal
zone name Infra-VSI-Fabric-A vsan <vsan-a-id>
member device-alias FlashArray-CT0FC0 target
member device-alias FlashArray-CT1FC0 target
member device-alias Infra-Host-FCP-01-A init
member device-alias Infra-Host-FCP-02-A init
member device-alias Infra-Host-FCP-03-A init
exit
zone name Infra-VSI-NVMe-Fabric-A vsan <vsan-a-id>
member device-alias FlashArray-CT0FC1 target
member device-alias FlashArray-CT1FC1 target
member device-alias Infra-Host-FC-NVMe-01-A init
member device-alias Infra-Host-FC-NVMe-02-A init
member device-alias Infra-Host-FC-NVMe-03-A init
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Infra-VSI-Fabric-A
member Infra-VSI-NVMe-Fabric-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
show zoneset active
copy r s
```



Since Smart Zoning is enabled, a single zone for each storage protocol (FCP and FC-NVMe) is created with all host boot initiators and boot targets for the FlashArray//X R3 instead of creating a separate zone for each host with the host initiator and boot targets. If a new host is added, its boot initiator can simply be added to the single zone in each MDS switch and then the zoneset reactivated. If another FlashArray is added to the FlashStack with FC targets, a new zone can be added for that FlashArray.

Cisco MDS 9132T B

To create the required zones and zoneset on Fabric B, run the following commands:

```
configure terminal

zone name Infra-VSI-Fabric-B vsan <vsan-b-id>
member device-alias FlashArray-CT0FC2 target
member device-alias FlashArray-CT1FC2 target
member device-alias Infra-Host-FCP-01-B init
member device-alias Infra-Host-FCP-02-B init
member device-alias Infra-Host-FCP-03-B init
exit
zone name Infra-VSI-NVMe-Fabric-B vsan <vsan-b-id>
member device-alias FlashArray-CT0FC3 target
member device-alias FlashArray-CT1FC3 target
member device-alias Infra-Host-FC-NVMe-01-B init
member device-alias Infra-Host-FC-NVMe-02-B init
member device-alias Infra-Host-FC-NVMe-03-B init
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Infra-VSI-Fabric-B
member Infra-VSI-NVMe-Fabric-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active
copy r s
```

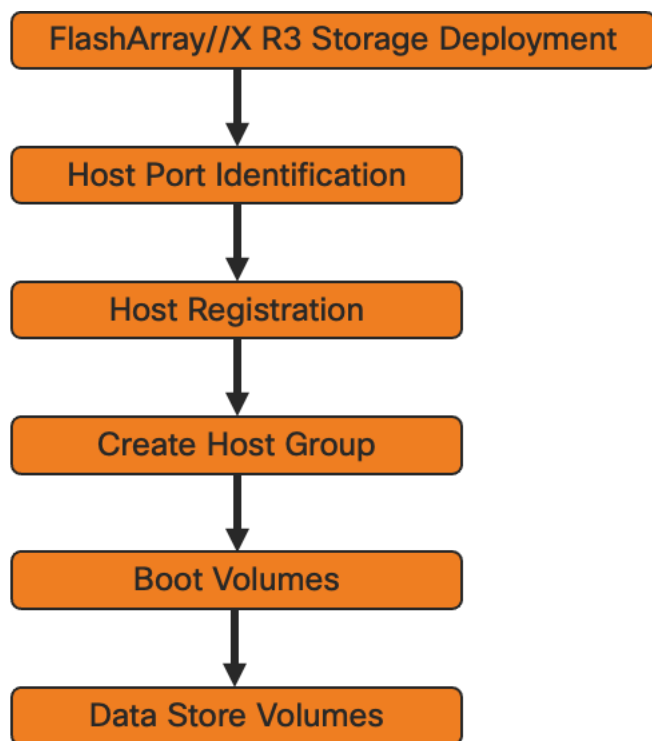
Storage Configuration – Boot LUNs

FlashArray Storage Deployment

The Pure Storage FlashArray//X is accessible to the FlashStack, but no storage has been deployed at this point. The storage to be deployed will include:

- ESXi FC Boot LUNs
- VMFS Datastores
- FC-NVMe Data stores

The FC Boot LUNs will need to be setup from the Pure Storage Web Portal, and the VMFS datastores can be provisioned from the Pure Storage Web Portal or can be directly provisioned from the vSphere Web Client after the Pure Storage vSphere Web Client Plugin has later been registered with the vCenter.



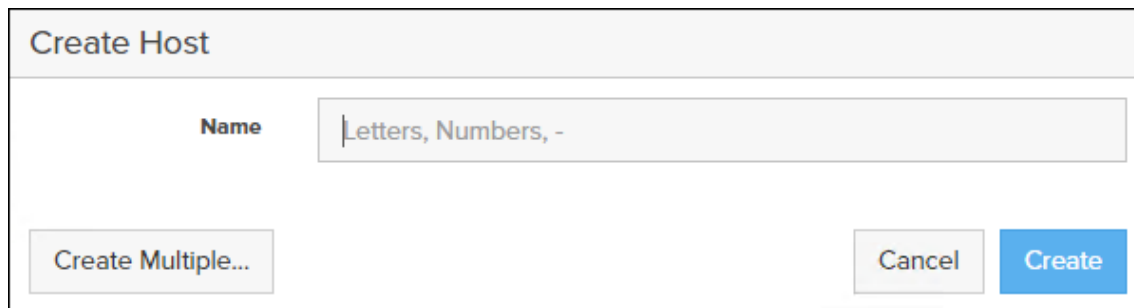
Host Port Identification

FC Boot LUNs will be mapped by the FlashArray//X using the assigned Initiator PWWN to the provisioned service profiles. This information can be found within the service profile located within the UCSM Servers > Service Profile > root > Sub-Organizations > FlashStack-VSI > Profiles:

Host Registration

To register the Host, follow these steps in the Pure Storage Web Portal:

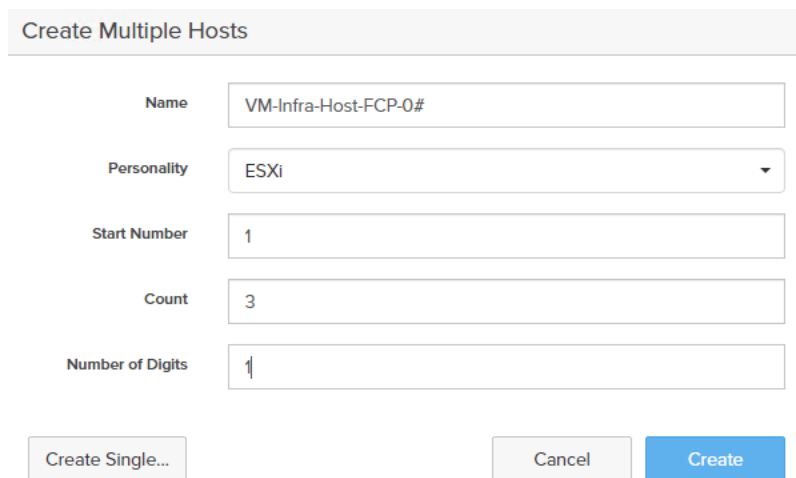
1. Select Storage > Hosts.
2. Select the + icon in the Hosts Panel.
3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.



Create Host

Name

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a “#” appearing in the name where an iterating number will appear:



Create Multiple Hosts

Name

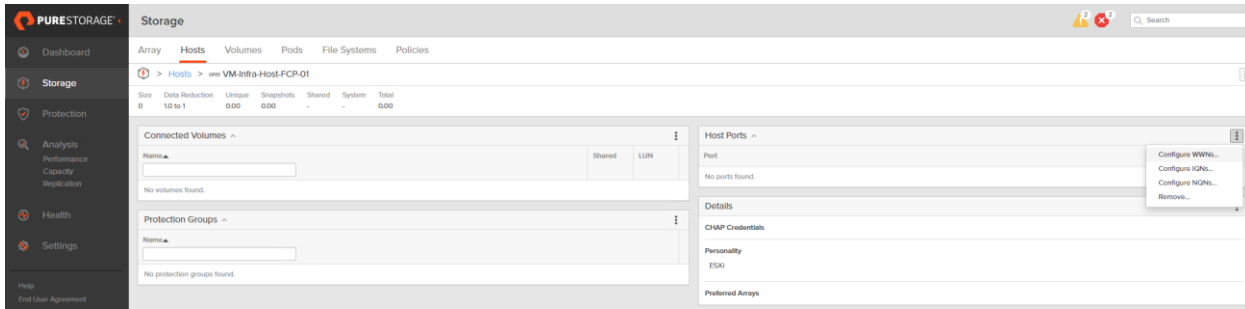
Personality

Start Number


Count

Number of Digits



5. Click Create to add the hosts.
6. For each host created, select the host.
7. In the Host view, select 'Configure WWNs...' from the Host Ports menu.



8. A pop-up will appear for Configure Fibre Channel WWNs <host being configured>. Within this pop-up, select the appropriate Existing WWNs from the discovered list.



9. Or you may enter the WWN manually by Selecting the +.



10. After entering the PWWN/WWPN, click Add to add the Host Ports.

11. Repeat steps 1-10 for each host created.

Create Host Group

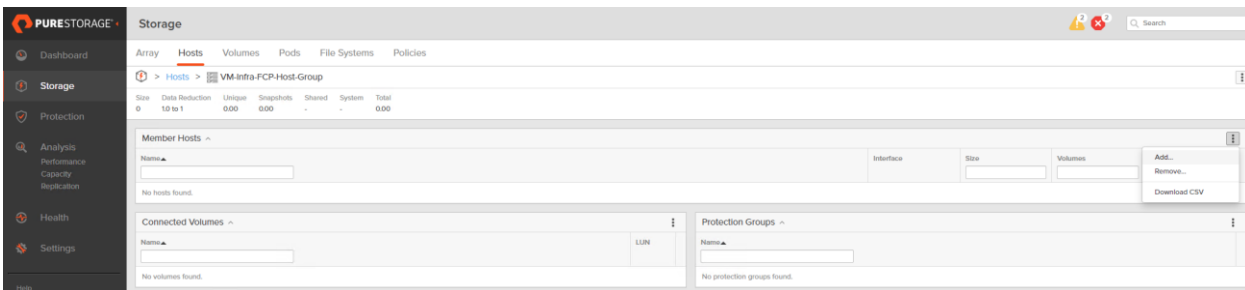
Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.
2. Select the + icon in the Host Groups Panel.
3. A pop-up will appear to create a host group on the FlashArray.

Create Host Group

Name

4. Provide a name for the group and click Create.
5. Select the group in the Host Groups Panel.
6. In the Host Group view, select 'Add...' from the Member Hosts menu.



7. Select the host to be part of the host group.

Add Hosts to Host Group

Existing Hosts

1-3 of 3

VM-Infra-Host-FCP-01

VM-Infra-Host-FCP-02

VM-Infra-Host-FCP-03

Selected Hosts

3 selected Clear all

VM-Infra-Host-FCP-01 ×

VM-Infra-Host-FCP-02 ×

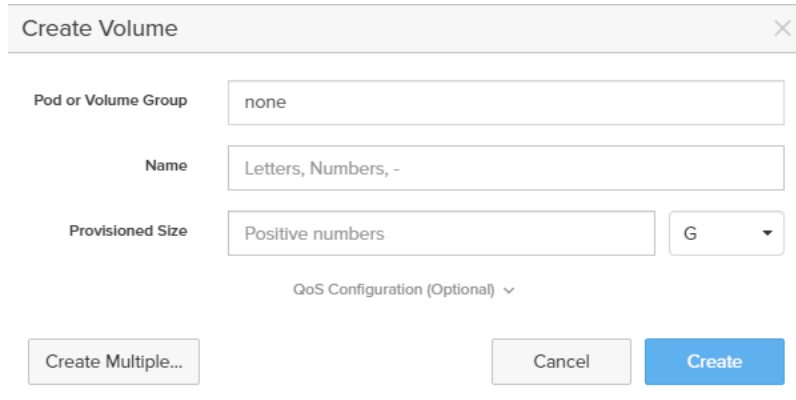
VM-Infra-Host-FCP-03 ×

8. Click Add.

Private Boot Volumes for each ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

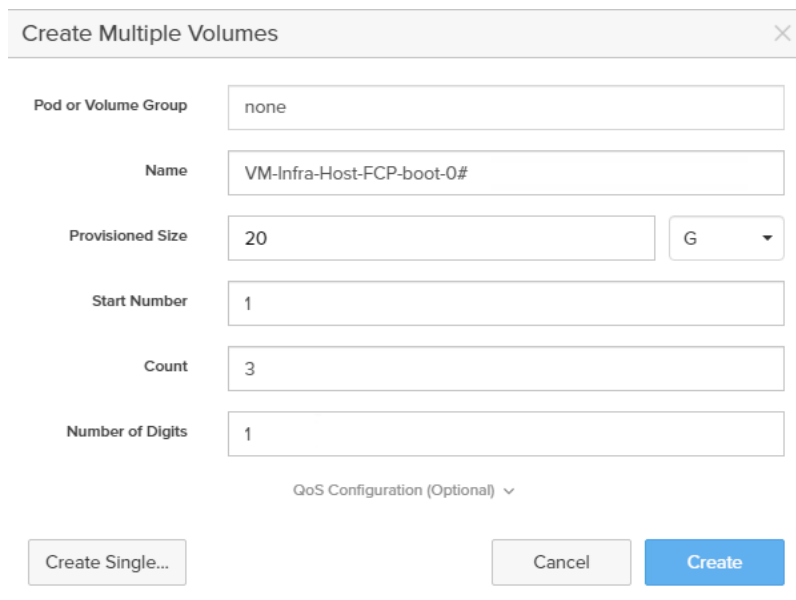
1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.



The 'Create Volume' dialog box contains the following fields and controls:

- Pod or Volume Group:** none
- Name:** Letters, Numbers, -
- Provisioned Size:** Positive numbers
- Unit:** G
- QoS Configuration (Optional):** (Optional) ▾
- Buttons:** Create Multiple..., Cancel, Create

4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Starting Number, Count, and Number of Digits, with a “#” appearing in the name where an iterating number will appear.

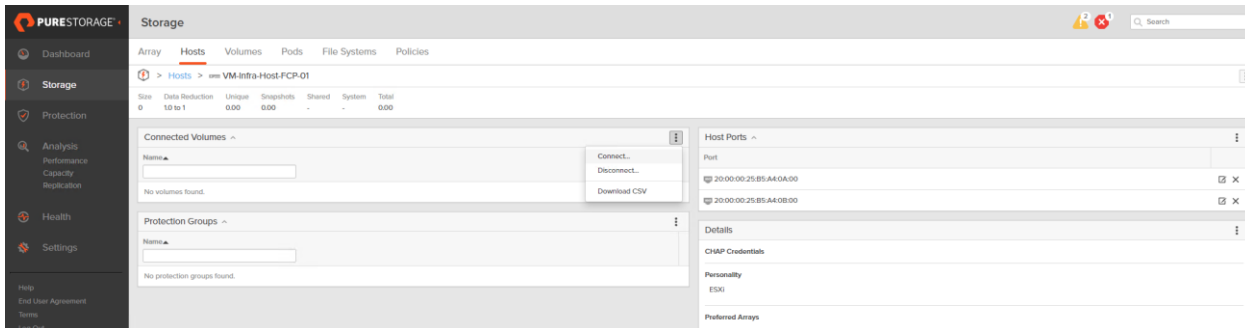


The 'Create Multiple Volumes' dialog box contains the following fields and controls:


- Pod or Volume Group:** none
- Name:** VM-Infra-Host-FCP-boot-0#
- Provisioned Size:** 20
- Unit:** G
- Start Number:** 1
- Count:** 3
- Number of Digits:** 1
- QoS Configuration (Optional):** (Optional) ▾
- Buttons:** Create Single..., Cancel, Create

5. Click Create to provision the volumes to be used as FC boot LUNs.

- Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon pull-down within the Connected Volumes tab within that host.



- From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.



 LUN ID 1 should be used for the boot.

- Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Create Infra and Swap Datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

- Select Storage > Volumes.

2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

The 'Create Volume' dialog box contains the following fields and controls:

- Pod or Volume Group:** A text input field with the value 'none'.
- Name:** A text input field with the value 'Letters, Numbers, -'.
- Provisioned Size:** A text input field with the value 'Positive numbers' and a dropdown menu set to 'G'.
- QoS Configuration (Optional):** A dropdown menu.
- Buttons:** 'Create Multiple...', 'Cancel', and 'Create'.

4. Fill in the Name and Provisioned Size.
5. Click Create to provision the volumes to be used as Infra datastore LUN.
6. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon drop-down within the Connected Volumes tab within that host group.

The screenshot shows the Pure Storage Storage console interface. The main content area displays the 'Hosts' section for the host group 'VM-Infra-Host-PCP-01'. A table shows storage metrics:

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	10 to 1	0.00	0.00	-	-	0.00

Below the table, there are sections for 'Connected Volumes' and 'Protection Groups', both showing 'No volumes found' and 'No protection groups found' respectively. A gear icon dropdown menu is open over the 'Connected Volumes' section, showing options: 'Connect...', 'Disconnect...', and 'Download CSV'. To the right, the 'Host Ports' section shows two ports with their respective IP addresses and MAC addresses, and a 'Details' section for 'CHAP Credentials' and 'Personality' (ESXi).

7. Within the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

Connect Volumes to Host

Existing Volumes 1-1 of 1

- vm-infra-boot-fcp
- VM-Infra-Boot-FCP-boot-01

Selected Volumes 1 selected Clear all

- VM-Infra-Boot-FCP-boot-01 x

LUN 1

Cancel Connect

8. Select the Infra datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.
9. Select Storage > Volumes.
10. Select the + icon in the Volumes Panel.
11. A pop-up will appear to create a volume on the FlashArray.

Create Volume

Pod or Volume Group none


Name Letters, Numbers, -

Provisioned Size Positive numbers G

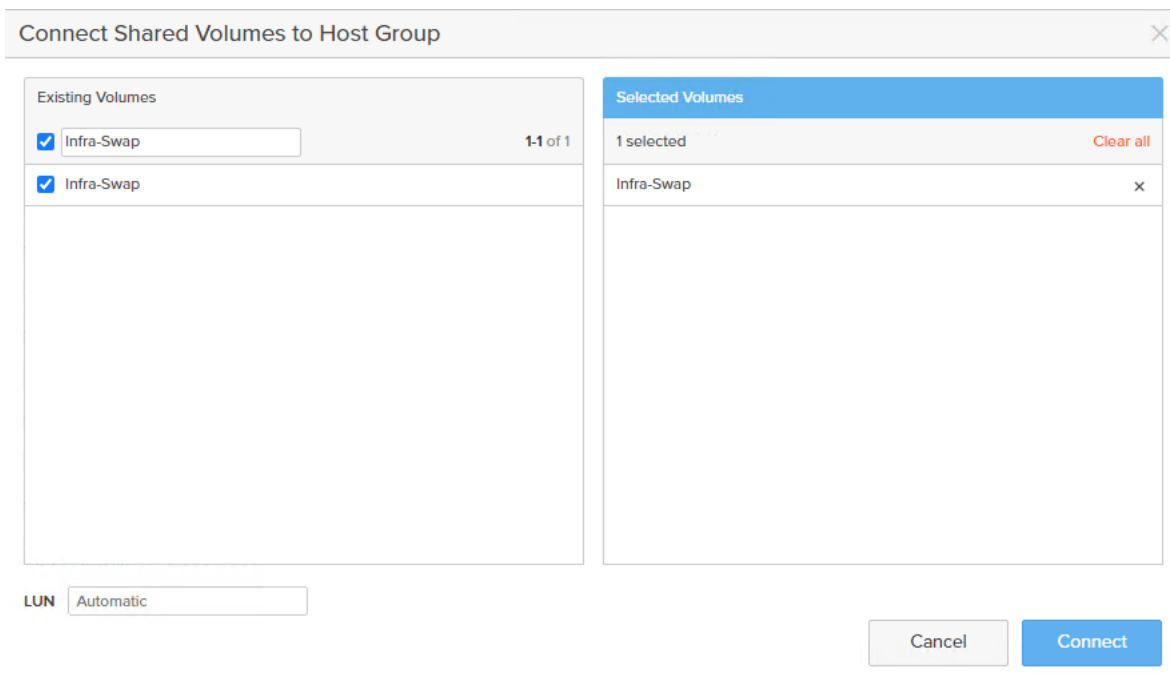
QoS Configuration (Optional) v

Create Multiple... Cancel Create

12. Fill in the Name and Provisioned Size.
13. Click Create to provision the volumes to be used as Swap datastore LUN.
14. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon drop-down within the Connected Volumes tab within that host group.



15. Within the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.



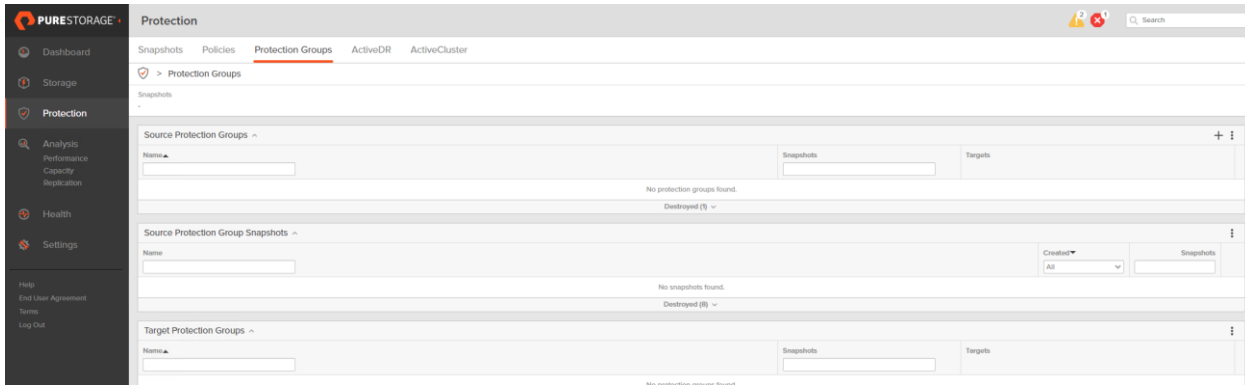
16. Select the Swap datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.

Configure Storage Policy Based Management

vSphere can communicate to the array via VASA provider to find out what features it supports and allow the vSphere administrator to assign, change, or remove functionality on a VVol on demand and via policies. Below is an example of how to configure a Protection group that will provide hourly snapshots that will be retained for 1 day, with 4 snapshots per day retained for 7 days. These policies should be configured based on application snapshot need.

To configure Storage Policy Based Management, follow these steps

1. From the Pure Storage Web Portal, Select Protection > Protection Groups > select the + icon in the Source Protection Groups.



2. Enter a name.

The screenshot shows a 'Create Protection Group' dialog box. It has a title bar with a close button (X). The dialog contains two input fields: 'Pod' with the value 'none' and 'Name' with the value 'Platinum'. At the bottom right, there are two buttons: 'Cancel' and 'Create'.

3. Select the protection group.

4. Edit the Snapshot Schedule based on your operational requirements.

The screenshot shows an 'Edit Snapshot Schedule' dialog box. It has a title bar with a close button (X). The dialog contains a toggle switch labeled 'Enabled' which is turned on. Below the toggle, there are three input fields: 'Create a snapshot on source every' with the value '1', a dropdown menu set to 'hours', and 'at' with a dropdown menu set to '-'. Below that, there are two input fields: 'Retain all snapshots on source for' with the value '1' and a dropdown menu set to 'days'. Below that, there are three input fields: 'then retain' with the value '4', 'snapshots per day for' with the value '7', and 'more days'. At the bottom right, there are two buttons: 'Cancel' and 'Save'.

5. Click Save.

VMware vSphere 7.0 U2 Setup

VMware ESXi 7.0 U2

This section provides detailed instructions for installing VMware ESXi 7.0 U2 in a FlashStack environment. After the procedures are completed, three booted ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in keyboard, video, mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot logical unit numbers (LUNs).

Download ESXi 7.0 U2 from VMware

If the VMware ESXi ISO has not already been downloaded, follow these steps:

1. Click this link: [Cisco Custom ISO for UCS 4.1.3a](#).



You will need a user id and password on vmware.com to download this software.



The Cisco Custom ISO for UCS 4.1.3a should also be used for Cisco UCS software release 4.2(1f) and VMware vSphere 7.0 U2.

2. Download the .iso file.

Log into Cisco UCS 6454 Fabric Interconnect

Cisco UCS Manager

The Cisco UCS IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log into the Cisco UCS environment, follow these steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. To log into Cisco UCS Manager, click Login.
6. From the main menu, click Servers.

-
7. Choose Servers > Service Profiles > root > Sub-Organizations > FlashStack-VSI Organization > VM-Host-Infra-FCP-01.
 8. In the Actions pane, click KVM Console.
 9. Follow the prompts to launch the HTML5 KVM console.
 10. Choose Servers > Service Profiles > root > Sub-Organizations > FlashStack-VSI Organization > VM-Host-Infra-FCP-02.
 11. In the Actions pane, click KVM Console.
 12. Follow the prompts to launch the HTML5 KVM console.
 13. Choose Servers > Service Profiles > root > Sub-Organizations > FlashStack-VSI Organization > VM-Host-Infra-FCP-03.
 14. In the Actions pane, click KVM Console.
 15. Follow the prompts to launch the HTML5 KVM console.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03



Skip this section if you're using vMedia policies; the ISO file will already be connected to KVM.

To prepare the server for the OS installation, follow these steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Choose Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and choose Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM Console tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To install VMware ESXi to the bootable LUN of the hosts, follow these steps on each host:

1. Boot the server by selecting Boot Server in the KVM and click OK, then click OK again.
2. On boot, the machine detects the presence of the ESXi installation media and loads the ESXi installer.



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. Then the ESXi installer should load properly.

3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.



It may be necessary to map function keys as User Defined Macros under the Macros menu in the UCS KVM console.

5. Choose the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Choose the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, press Enter to reboot the server.



The ESXi installation image will be automatically unmapped in the KVM when Enter is pressed.

10. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. The following section details how to add a management network for the VMware hosts.

ESXi Host VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To configure each ESXi host with access to the management network, follow these steps:

1. After the server has finished rebooting, in the UCS KVM console, press F2 to customize VMware ESXi.
2. Log in as root, enter the corresponding password, and press Enter to log in.
3. Use the down arrow key to choose Troubleshooting Options and press Enter.
4. Choose Enable ESXi Shell and press Enter.
5. Choose Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Choose the Configure Management Network option and press Enter.
8. Choose Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field. If the numbers do not match, note the mapping of vmnic ports to vNIC ports for later use.
10. Using the spacebar, choose vmnic1.

```
Network Adapters
Select the adapters for this host's default management network
connection. Use two or more adapters for fault-tolerance and
load-balancing.

Device Name  Hardware Label (MAC Address)  Status
[X] vmnic0   00-vSwitch0-A (...:91:1a:00)  Connected (...)
[X] vmnic1   01-vSwitch0-B (...:91:1b:00)  Connected (...)
[ ] vmnic2   02-VDS-A (00:25:b5:91:1a:01)  Connected
[ ] vmnic3   03-VDS-B (00:25:b5:91:1b:01)  Connected

<D> View Details  <Space> Toggle Selected          <Enter> OK  <Esc> Cancel
```



In lab testing, examples have been seen where the vmnic and device ordering do not match. If this is the case, use the Consistent Device Naming (CDN) to note which vmnics are mapped to which vNICs and adjust the upcoming procedure accordingly.

-
11. Press Enter.
 12. Choose the VLAN (Optional) option and press Enter.
 13. Enter the <ib-mgmt-vlan-id> and press Enter.
 14. Choose IPv4 Configuration and press Enter.
 15. Choose the “Set static IPv4 address and network configuration” option by using the arrow keys and space bar.
 16. Move to the IPv4 Address field and enter the IP address for managing the ESXi host.
 17. Move to the Subnet Mask field and enter the subnet mask for the ESXi host.
 18. Move to the Default Gateway field and enter the default gateway for the ESXi host.
 19. Press Enter to accept the changes to the IP configuration.
 20. Choose the IPv6 Configuration option and press Enter.
 21. Using the spacebar, choose Disable IPv6 (restart required) and press Enter.
 22. Choose the DNS Configuration option and press Enter.



Since the IP address is assigned manually, the DNS information must also be entered manually.

23. Using the spacebar, choose “Use the following DNS server addresses and hostname:”
24. Move to the Primary DNS Server field and enter the IP address of the primary DNS server.
25. Optional: Move to the Alternate DNS Server field and enter the IP address of the secondary DNS server.
26. Move to the Hostname field and enter the fully qualified domain name (FQDN) for the ESXi host.
27. Press Enter to accept the changes to the DNS configuration.
28. Press Esc to exit the Configure Management Network submenu.
29. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi Host VMkernel Port vmk0 MAC Address (Optional)

ESXi VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port it is placed on. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will exist because vmk0 will retain 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, follow these steps:

1. From the ESXi console menu main screen, type Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of Static Macros.
2. Log in as root.
3. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk0. vmk0 should be a part of the "Management Network" port group. Note the IP address and netmask of vmk0.
4. To remove vmk0, type `esxcfg-vmknic -d "Management Network"`.
5. To re-add vmk0 with a random MAC address, type `esxcfg-vmknic -a -i <vmk0-ip> -n <vmk0-netmask> "Management Network"`.
6. Verify vmk0 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.
7. Tag vmk0 as the management interface by typing `esxcli network ip interface tag add -i vmk0 -t Management`.
8. When vmk0 was re-added, if a message popped up saying vmk1 was marked as the management interface, type `esxcli network ip interface tag remove -i vmk1 -t Management`.
9. If this VMware ESXi host is iSCSI booted, the vmk1, iScsiBootPG-A interface's MAC address can also be reset to a random, VMware-assigned MAC address.
 - a. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk1. vmk1 should be a part of the "iScsiBootPG-A" port group and should have a MAC address from the UCS MAC Pool. Note the IP address and netmask of vmk1.
 - b. To remove vmk1, type `esxcfg-vmknic -d "iScsiBootPG-A"`.
 - c. To re-add vmk1 with a random MAC address, type `esxcfg-vmknic -a -i <vmk1-ip> -n <vmk1-netmask> -m 9000 "iScsiBootPG-A"`.
 - d. Verify vmk1 has been re-added with a random MAC address by typing `esxcfg-vmknic -l`.
 - e. Type `exit` to log out of the command line interface.
10. Type Ctrl-Alt-F2 to return to the ESXi console menu interface.

Install VMware and Cisco VIC Drivers for the ESXi Host

Download the offline bundle for the UCS Tools Component and VMware VIC Driver to the Management workstation:

[UCS Tools Component for ESXi 7.0 1.2.1](#) (ucs-tool-esxi_1.2.1-1OEM.zip)

[VMware ESXi 7.0 nfnic 5.0.0.15 Driver for Cisco VIC Adapters](#) (Cisco-nfnic_5.0.0.15-1OEM.700.1.0.15843807_18697950.zip)

[nfnic Driver version 1.0.35.0 \(nfnic driver is included with the Cisco ESXi installation ISO\).](#)

This document is using the driver versions shown above along with Cisco VIC nfnic version 1.0.35.0 and nfnic version 5.0.0.15 along with VMware vSphere version 7.0 U2, Cisco UCS version 4.2(1f), and the Pure Purity version 6.1.6. These were the versions validated and supported at the time this document was published. This document can be used as a guide for configuring future versions of software. Consult the Cisco UCS Hardware Compatibility List and the Pure Interoperability Matrix Tool to determine supported combinations

ESXi Hosts VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03

To install UCS Tools on the ESXi host ESXi VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02, and VM-Host-Infra-FCP-03, follow these steps:



The latest nfnic driver is already included with the ESXi install ISO and is not required to be updated if the [Cisco Custom ISO for UCS 4.1.3a is used.](#)

1. Using an SCP program such as WinSCP, copy the two offline bundles referenced above to the /tmp directory on each ESXi host.
2. Using a ssh tool such as PuTTY, ssh to each VMware ESXi host. Log in as root with the root password.
3. Type cd /tmp.
4. Run the following commands on each host:

```
esxcli software component apply -d /tmp/Cisco-nfnic_5.0.0.15-1OEM.700.1.0.15843807_18697950.zip
esxcli software component apply -d /tmp/ucs-tool-esxi_1.2.1-1OEM.zip
reboot
```

5. After reboot, log back into each host and run the following commands and ensure the correct version is installed:


```
esxcli software vib list | grep nenic
esxcli software component list | grep nfnic
esxcli software component list | grep ucs
```

Log into the First VMware ESXi Host by Using VMware Host Client

ESXi Host VM-Host-Infra-FCP-01

To log into the VM-Host-Infra-FCP-01 ESXi host by using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-FCP-01 management IP address.
2. Enter root for the User name.
3. Enter the root password.
4. Click Login to connect.
5. Decide whether to join the VMware Customer Experience Improvement Program and click OK.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-FCP-01

To set up the VMkernel ports and the virtual switches on the first ESXi host, follow these steps:



In this procedure, you're only setting up the first ESXi host. The second and third hosts will be added to vCenter and setup from the vCenter HTML5 Interface.

1. From the Host Client Navigator, choose Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the vSwitch0 line.
4. Choose Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, choose vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.

-
9. Click Save.
 10. Choose Networking, then choose the Port groups tab.
 11. In the center pane, right-click VM Network and choose Edit settings.
 12. Name the port group IB-MGMT Network and enter <ib-mgmt-vlan-id> in the VLAN ID field.
 13. Click Save to finalize the edits for the IB-MGMT Network.
 14. Click Add port group.
 15. Name the port group OOB-MGMT Network and enter the <OOB-MGMT-vlan-id> for the VLAN ID.
 16. Click Add to finalize the edits for the OOB-MGMT port group.
 17. At the top, choose the VMkernel NICs tab.
 18. Click VMkernel NICs tab.
 19. Click Add VMkernel NIC.
 20. For New port group, enter VMkernel-vMotion.
 21. For Virtual switch, choose vSwitch0.
 22. Enter <vmotion-vlan-id> for the VLAN ID.
 23. Change the MTU to 9000.
 24. Choose Static IPv4 settings and expand IPv4 settings.
 25. Enter the ESXi host vMotion IP address and netmask.
 26. Choose the vMotion stack for TCP/IP stack.
 27. Click Create.
 28. Choose the Virtual Switches tab, then vSwitch0. The properties for vSwitch0 VMkernel NICs should be like the following example:

[Add uplink](#) | [Edit settings](#) | [Refresh](#) | [Actions](#)

vSwitch0
 Type: Standard vSwitch
 Port groups: 3
 Uplinks: 2

vSwitch Details	
MTU	1500
Ports	5374 (5365 available)
Link discovery	Listen / Cisco discovery protocol (CDP)
Attached VMs	0 (0 active)
Beacon interval	1

NIC teaming policy	
Notify switches	Yes
Policy	Route based on originating port ID
Reverse policy	Yes
Failback	Yes

Security policy	
Allow promiscuous mode	No
Allow forged transmits	No
Allow MAC changes	No

Shaping policy	
Enabled	No

vSwitch topology

29. Choose Networking and the VMkernel NICs tab to confirm configured virtual adapter. The adapter listed should be like the following example:

VM-Host-Infra-FCP-01.flashstack.com - Networking

[Port groups](#) | [Virtual switches](#) | [Physical NICs](#) | **VMkernel NICs** | [TCP/IP stacks](#) | [Firewall rules](#)

[Add VMkernel NIC](#) | [Edit settings](#) | [Refresh](#) | [Actions](#)

Name	Portgroup	TCP/IP stack	Services	IPv4 address	IPv6 addresses
vmk0	Management Network	Default TCP/IP stack	Management	10.1.164.111	fe80::225:b5ff:fe91:1a00/64

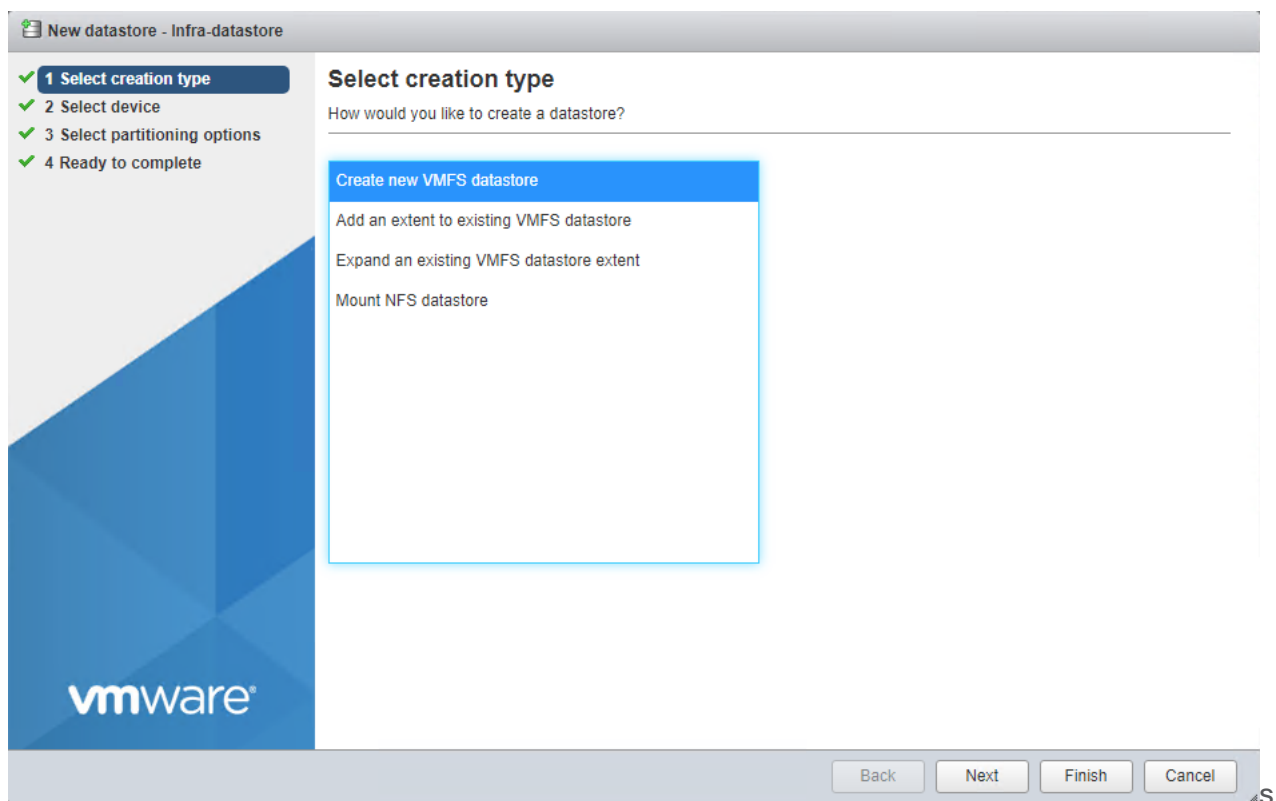
1 items

Mount Required Datastores

ESXi Host VM-Host-Infra-FCP-01

To mount the required datastores, follow these steps on the first ESXi host:

1. From the Host Client, choose Storage.
2. In the center pane, choose the Datastores tab.
3. Click New datastore to add a new datastore.
4. In the New datastore popup, choose Create new VMFS datastore and click Next.



5. Input Infra-Datastore1 for the datastore name.
6. Select the Pure LUN that will be used for the data store.
7. Click Next.

New datastore - Infra-Data Store1

- ✓ 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
Infra-DataStore1

The following devices are unclaimed and can be used to create a new VMFS datastore


Name	Type	Capacity	Free space
Local ATA Disk (t10.ATA____Micron_5100_MTFDD...)	Disk (SSD)	223.57 GB	223.57 GB
PURE Fibre Channel Disk (naa.624a9370b6c770713...)	Disk (SSD)	1,024 GB	1,024 GB

2 items

vmware

Back Next Finish Cancel

8. Click Next.




9. Click Finish. The datastore should now appear in the datastore list.

10. In the center pane, choose the Datastores tab.

11. Click New datastore to add a new datastore.

12. In the New datastore popup, choose Create new VMFS datastore and click Next.



13. Input Infra-Swap for the datastore name.

14. Select the Pure LUN that will be used for the data store.

15. Click Next.

16. Click Next again.

17. Click Finish. The datastore should now appear in the datastore list.

Name	Drive Type	Capacity	Provisioned	Free	Type	Thin provisioning	Access
datastore1	SSD	95.5 GB	2 GB	93.5 GB	VMFS6	Supported	Single
Infra-DataStore1	SSD	1,023.75 GB	1.42 GB	1,022.33 GB	VMFS6	Supported	Single
Infra-Swap	SSD	499.75 GB	1.41 GB	498.34 GB	VMFS6	Supported	Single

Configure NTP on First ESXi Host

ESXi Host VM-Host-Infra-FCP-01

To configure Network Time Protocol (NTP) on the first ESXi host, follow these steps:

1. From the Host Client, choose Manage.
2. In the center pane, choose System > Time & date.

3. Click Edit NTP settings.
4. Make sure “Manually configure the date and time on this host and enter the approximate date and time.
5. Select Use Network Time Protocol (enable NTP client).
6. Use the drop-down list to choose Start and stop with host.
7. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

Specify how the date and time of this host should be set.

Manually configure the date and time on this host

09/03/2021 11:52 AM

Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.164.61, 10.1.164.62

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

Save Cancel

8. Click Save to save the configuration changes.



It currently is not possible to start NTP from the ESXi Host Client. NTP will be started from vCenter. The NTP server time may initially vary slightly from the host time.

Configure ESXi Host Swap

ESXi Host VM-Host-Infra-FCP-01

To configure host swap on the first ESXi host, follow these steps on the host:

1. From the Host Client, choose Manage.
2. In the center pane, choose System > Swap.
3. Click Edit settings.
4. From the drop-down list choose Infra-Swap. Leave all other settings unchanged.

Edit swap configuration	
Enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No
Datastore	Infra-Swap
Local swap enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No
Host cache enabled	<input checked="" type="radio"/> Yes <input type="radio"/> No

Save Cancel

5. Click Save to save the configuration changes.

Configure Host Power Policy

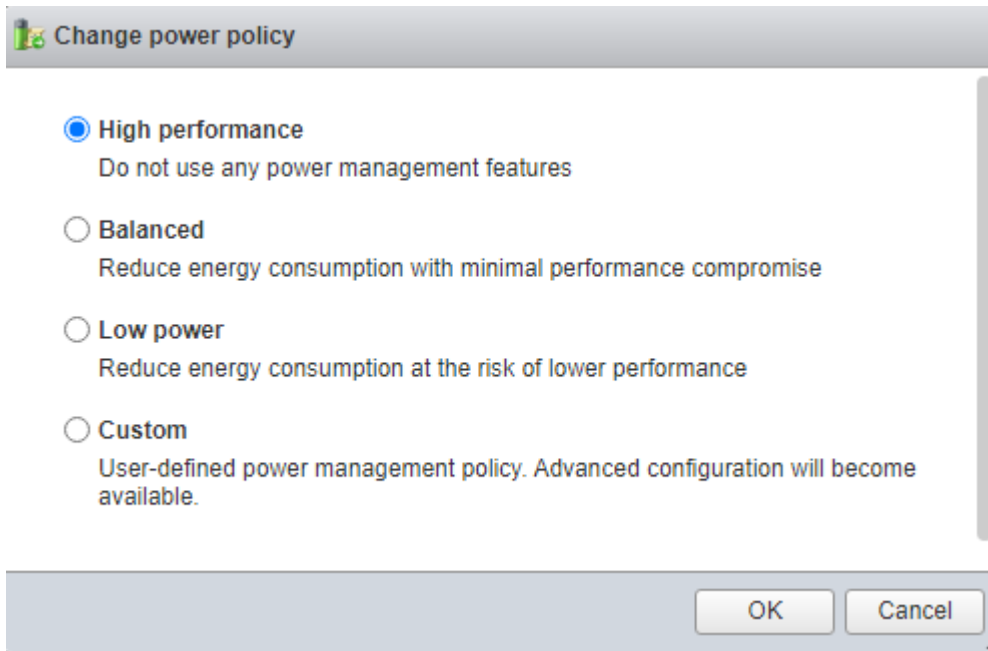
ESXi Host VM-Host-Infra-FCP-01

To configure the host power policy on the first ESXi host, follow these steps on the host:



Implementation of this policy is recommended in [Performance Tuning for Cisco UCS M6 Servers](#) for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. From the Host Client, choose Manage.
2. In the center pane, choose Hardware > Power Management.
3. Choose Change policy.
4. Choose High performance and click OK.



If you are implementing iSCSI boot, execute the VMware ESXi setup scripts in the [iSCSI Addition](#) appendix.

VMware vCenter 7.0 U2B (Optional)

The procedures in the following subsections provide detailed instructions for installing the VMware vCenter 7.0U2B Server Appliance in a FlashStack environment. After the procedures are completed, a VMware vCenter Server will be configured.

Build the VMware vCenter Server Appliance

The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, follow these steps:

1. Locate and copy the VMware-VCSA-all-7.0.2-17958471.iso file to the desktop of the management workstation. This ISO is for the VMware vSphere 7.0 U2 vCenter Server Appliance.



It is important to use at minimum VMware vCenter release 7.0U2 to ensure access to all needed features.

2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012 and above).
3. 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.
4. Click Install to start the vCenter Server Appliance deployment wizard.

-
5. Click NEXT in the Introduction section.
 6. Read and accept the license agreement and click NEXT.
 7. In the “vCenter Server deployment target” window, enter the host name or IP address of the first ESXi host, User name (root) and Password. Click NEXT.
 8. Click YES to accept the certificate.
 9. Enter the Appliance VM name and password details in the “Set up vCenter Server VM” section. Click NEXT.
 10. In the “Select deployment size” section, choose the Deployment size and Storage size. For example, choose “Small” and “Default”. Click NEXT.
 11. Choose Infra-DataStore1 for storage. Click NEXT.
 12. In the “Network Settings” section, configure the below settings:
 - a. Choose a Network: IB-MGMT Network.



It is important that the vCenter VM stay on the IB-MGMT Network on vSwitch0 and that it isn't moved to a vDS. If vCenter is moved to a vDS and the virtual environment is completely shut down and then brought back up, and it is attempted to bring up vCenter on a different host than the one it was running on before the shutdown, vCenter will not have a functional network connection. 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 to be brought up always occurs correctly without requiring vCenter to already be up and running.

- b. IP version: IPV4
 - c. IP assignment: static
 - d. FQDN: <vcenter-fqdn>
 - e. IP address: <vcenter-ip>
 - f. Subnet mask or prefix length: <vcenter-subnet-mask>
 - g. Default gateway: <vcenter-gateway>
 - h. DNS Servers: <dns-server1>,<dns-server2>
13. Click NEXT.
 14. Review all values and click FINISH to complete the installation.



The vCenter Server appliance installation will take a few minutes to complete.

15. Click CONTINUE to proceed with stage 2 configuration.
16. Click NEXT.
17. In the vCenter Server configuration window, configure these settings:
 - a. Time Synchronization Mode: Synchronize time with NTP servers.
 - b. NTP Servers: <nexus-a-ntp-ip>,<nexus-b-ntp-ip>
 - c. SSH access: Enabled.
18. Click NEXT.
19. Complete the SSO configuration as shown below, or according to your organization's security policies:
20. Click NEXT.
21. Decide whether to join VMware's Customer Experience Improvement Program (CEIP).
22. Click NEXT.
23. Review the configuration and click FINISH.
24. Click OK.



vCenter Server setup will take a few minutes to complete.


25. Click CLOSE. Eject or unmount the VCSA installer ISO.

Adjust vCenter CPU Settings


If 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 B and C-Series servers are normally 2-socket servers. In this validation, the Small deployment size was selected and vCenter was setup for a 4-socket server. This setup will cause issues in the VMware ESXi cluster Admission Control. To resolve the Admission Control issue, follow these steps:

1. Open a web browser on the management workstation and navigate to the VM-Host-Infra-FCP-01 management IP address.
2. Enter root for the user name.
3. Enter the root password.

4. Click Login to connect.
5. On the left, choose Virtual Machines.
6. In the center pane, right-click the vCenter VM and choose Edit settings.
7. In the Edit settings window, expand CPU and check the value of Sockets.



8. If the number of Sockets does not match your server configuration, it will need to be adjusted. Click Cancel.
9. If the number of Sockets needs to be adjusted:
 - a. Right-click the vCenter VM and choose Guest OS > Shut down. Click Yes on the confirmation.
 - b. Once vCenter is shut down, right-click the vCenter VM and choose Edit settings.
 - c. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value equal to your server configuration (normally 2).



- d. Click Save.
- e. Right-click the vCenter VM and choose Power > Power on. Wait approximately 10 minutes for vCenter to come up.

Setup VMware vCenter Server

To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to <https://<vcenter-ip-address>:5480>. You will need to navigate security screens.
2. Log into the VMware vCenter Server Management interface as root with the root password set in the vCenter installation.
3. In the menu on the left, choose Time.
4. Choose EDIT to the right of Time zone.
5. Choose the appropriate Time zone and click SAVE.
6. In the menu choose Administration.
7. According to your Security Policy, adjust the settings for the root user and password.
8. In the menu on the left choose Update.
9. Follow the prompts to STAGE AND INSTALL any available vCenter updates. In this validation, vCenter version 7.0.2.00200 was installed.
10. In the upper right-hand corner of the screen, choose root > Logout to logout of the Appliance Management interface.
11. Using a web browser, navigate to <https://<vcenter-fqdn>>. You will need to navigate security screens.



With VMware vCenter 7.0, the use of the vCenter FQDN is required.

12. Choose LAUNCH VSPHERE CLIENT (HTML5).



Although the previous versions of this document used the FLEX vSphere Web Client, the VMware vSphere HTML5 Client is the only option starting with vSphere 7 and will be used going forward.

13. Log in using the Single Sign-On username (`administrator@vsphere.local`) and password created during the vCenter installation. Dismiss the Licensing warning currently.
14. In the center pane, choose ACTIONS > New Datacenter.
15. Type “FlashStack-DC” in the Datacenter name field.

New Datacenter



Name

FlashStack-DC

Location:

 vcenter1.flashstack.com

CANCEL

OK

16. Click OK.

17. Expand the vCenter on the left.

18. Right-click the datacenter FlashStack-DC in the list in the left pane. Choose New Cluster.

19. Name the cluster FlashStack-Management.

20. Turn on DRS and vSphere HA. Do not turn on vSAN.

New Cluster

- 1 Basics
- 2 Review

Basics

Name	FlashStack-Management
Location	Datacenter-FC
vSphere DRS	<input checked="" type="checkbox"/>
vSphere HA	<input checked="" type="checkbox"/>
vSAN	<input type="checkbox"/>

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

Manage all hosts in the cluster with a single image

CANCEL **NEXT**

21. Click OK to create the new cluster.

22. Right-click “FlashStack-Management” and choose Settings.

23. Choose Configuration > General in the list located on the left and choose EDIT located on the right of General.

24. Choose Datastore specified by host and click OK.


Edit Cluster Settings | FlashStack-Management X

Virtual machine directory

Store the swap files in the same directory as the virtual machine.

Datastore specified by host

Store the swap files in the datastore specified by the host to be used for swap files. If not possible, store the swap files in the same directory as the virtual machine.

 Using a datastore that is not visible to both hosts during vMotion might affect the vMotion performance for the affected virtual machines.

CANCEL

OK

25. Right-click “FlashStack-Management” and click Add Hosts.
26. In the IP address or FQDN field, enter either the IP address or the FQDN of the first VMware ESXi host. Enter root as the Username and the root password. Click NEXT.
27. In the Security Alert window, choose the host and click OK.
28. Verify the Host summary information and click NEXT.
29. Ignore warnings about the host being moved to Maintenance Mode and click FINISH to complete adding the host to the cluster.
30. The added ESXi host will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.
31. In the list, right-click the added ESXi host and choose Settings.
32. In the center pane under Virtual Machines, choose Swap File location.
33. On the right, click EDIT.
34. Choose the Infra-Swap datastore and click OK.
35. In the list under System, choose Time Configuration.
36. Click EDIT to the right of Manual Time Configuration. Set the time and date to the correct local time and click OK.
37. Click EDIT to the right of Network Time Protocol.
38. In the Edit Network Time Protocol window, select Enable and then select Start NTP Service. Ensure the other fields are filled in correctly and click OK.

Edit Network Time Protocol | 10.1.164.117



Enable ⓘ

NTP Servers	10.1.164.61, 10.1.164.62
Separate servers with commas, e.g. 10.31.21.2, fe00::2800	
NTP Service Status:	Stopped <input checked="" type="checkbox"/> Start NTP Service
NTP Service Startup Policy:	Start and stop manually

CANCEL

OK

39. In the list under Hardware, choose Overview. Scroll to the bottom and ensure the Power Management Active policy is High Performance. If the Power Management Active policy is not High Performance, to the right of Power Management, choose EDIT POWER POLICY. Choose High performance and click OK.
40. In the list under Storage, choose Storage Devices. Make sure the Pure Fibre Channel Disk LUN 1 or Pure iSCSI Disk LUN 1 is selected.
41. Choose the Paths tab.
42. Ensure that 4 paths appear, which should have the status Active (I/O).

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS HDD DISK MARK AS LOCAL MARK AS PERENNIALY RESERVED

Name	LUN	Type	Capacity	Datastore	Operational
<input checked="" type="checkbox"/> PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd4000141af)	1	disk	20.00 GB	Not Consumed	Attached
<input type="checkbox"/> Local USB Direct-Access (mpx.vmhba32:CO:T0:L2)	2	disk	0.00 B	Not Consumed	Attached

1 - 20 of 29 items | 1 / 2

Properties Paths Partition Details

ENABLE DISABLE

Runtime Name	Status	Target	Name	Preferred
<input type="radio"/> vmhba1:CO:T22:L1	Active (I/O)	52.4a:93:77:de:d7:21:10 52...	vmhba1:CO:T22:L1	
<input type="radio"/> vmhba1:CO:T21:L1	Active (I/O)	52.4a:93:77:de:d7:21:00 52...	vmhba1:CO:T21:L1	
<input type="radio"/> vmhba0:CO:T26:L1	Active (I/O)	52.4a:93:77:de:d7:21:12 52...	vmhba0:CO:T26:L1	
<input type="radio"/> vmhba0:CO:T25:L1	Active (I/O)	52.4a:93:77:de:d7:21:02 52...	vmhba0:CO:T25:L1	

Add AD User Authentication to vCenter (Optional)

If an AD Infrastructure is set up in this FlashStack environment, you can set up in AD and authenticate from vCenter.

To add an AD user authentication to the vCenter, follow these steps:

1. In the AD Infrastructure, using the Active Directory Users and Computers tool, setup a Domain Administrator user with a user name such as flashadmin (FlashStack Admin).
2. Connect to <https://<vcenter-ip>> and choose LAUNCH VSPHERE CLIENT (HTML5).
3. Log in as Administrator@vsphere.local (or the SSO user set up in vCenter installation) with the corresponding password.
4. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
5. In the center pane, under Configuration, choose the Identity Provider tab.
6. In the list under Type, select Active Directory Domain.
7. Choose JOIN AD.
8. Fill in the AD domain name, the Administrator user, and the domain Administrator password. Do not fill in an Organizational unit. Click JOIN.
9. Click Acknowledge.
10. In the list on the left under Deployment, choose System Configuration. Choose the radio button to choose the vCenter, then choose REBOOT NODE.
11. Input a reboot reason and click OK. The reboot will take approximately 10 minutes for full vCenter initialization.
12. Log back into the vCenter vSphere HTML5 Client as Administrator@vsphere.local.
13. Under Menu, choose Administration. In the list on the left, under Single Sign On, choose Configuration.
14. In the center pane, under Configuration, choose the Identity Provider tab. Under Type, select Identity Sources. Click ADD.
15. Make sure your Active Directory (Integrated Windows Authentication) is selected, your Windows Domain name is listed, and Use machine account is selected. Click ADD.
16. In the list select the Active Directory (Integrated Windows Authentication) Identity source type. If desired, select SET AS DEFAULT and click OK.
17. On the left under Access Control, choose Global Permissions.
18. In the center pane, click the + sign to add a Global Permission.

19. In the Add Permission window, choose your AD domain for the Domain.

20. On the User/Group line, enter either the FlashStack Admin username or the Domain Admins group. Leave the Role set to Administrator. Choose the Propagate to children checkbox.



The FlashStack Admin user was created in the Domain Admins group. The selection here depends on whether the FlashStack Admin user will be the only user used in this FlashStack or you would like to add other users later. By selecting the Domain Admins group, any user placed in that group in the AD domain will be able to login to vCenter as an Administrator.

21. Click OK to add the selected User or Group. The user or group should now appear in the Global Permissions list with the Administrator role.

22. Log out and log back into the vCenter HTML5 Client as the FlashStack Admin user. You will need to add the domain name to the user, for example, flashadmin@domain.

FlashStack VMware vSphere Distributed Switch (vDS)

This section provides detailed procedures for installing the VMware vDS in vCenter and on the first FlashStack ESXi Management Host.

In the Cisco UCS setup section of this document two sets of vNICs were setup. The vmnic ports associated with the vDS0-A and B vNICs will be placed on the VMware vDS in this procedure. The vMotion VMkernel port(s) will be placed on the vDS.

A vMotion, and a VM-Traffic port group will be added to the vDS. Any additional VLAN-based port groups added to the vDS would need to have the corresponding VLANs added to the Cisco UCS LAN cloud, to the Cisco UCS vDS0-A and B vNIC templates, and to the Cisco Nexus 9K switches and vPC peer-link interfaces on the switches.

In this document, the infrastructure ESXi management VMkernel ports, the In-Band management interfaces including the vCenter management interface are left on vSwitch0 to facilitate bringing the virtual environment back up in the event it needs to be completely shut down. The vMotion VMkernel ports are moved to the vDS to allow QoS marking of vMotion to be done at the VLAN level in the vDS if vMotion needs to have QoS policies applied in the future. The vMotion port group is also pinned to Cisco UCS fabric B. Pinning should be done in a vDS to ensure consistency across all ESXi hosts.

Configure the VMware vDS in vCenter for the VMware vSphere Web Client

To configure the vDS, follow these steps:

1. After logging into the VMware vSphere HTML5 Client, choose Networking under Menu.
2. Right-click the FlashStack-DC datacenter and choose Distributed Switch > New Distributed Switch.
3. Give the Distributed Switch a descriptive name (vDS0) and click NEXT.

4. Make sure version 7.0.2 – ESXi 7.0.2 and later is selected and click NEXT.
5. 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.
6. Review the information and click FINISH to complete creating the vDS.

The screenshot shows the 'New Distributed Switch' wizard in the vSphere interface. The left sidebar contains a progress indicator with four steps: 1 Name and location, 2 Select version, 3 Configure settings, and 4 Ready to complete (highlighted). The main area is titled 'Ready to complete' and includes a close button (X). Below the title is a summary of the selected settings:

Name	vDS0
Version	7.0.2
Number of uplinks	2
Network I/O Control	Enabled
Default port group	VM-Traffic

Below the settings is a section for 'Suggested next actions' with two options: 'New Distributed Port Group' and 'Add and Manage Hosts'. A note at the bottom states: 'These actions will be available in the Actions menu of the new distributed switch.' At the bottom right, there are three buttons: 'CANCEL', 'BACK', and 'FINISH' (highlighted in green).

7. Expand the FlashStack-DC datacenter and the newly created vDS. Choose the newly created vDS.
8. Right-click the VM-Traffic port group and choose Edit Settings.
9. Choose VLAN.
10. Choose VLAN for VLAN type and enter the VM-Traffic VLAN ID. Click OK.
11. Right-click the vDS and choose Settings > Edit Settings.
12. In the Edit Settings window, choose Advanced.

13. Change the MTU to 9000. The Discovery Protocol can optionally be changed to Link Layer Discovery Protocol and the Operation to Both. Click OK.

Distributed Switch - Edit Settings vDS0 ×

General **Advanced** Uplinks

MTU (Bytes)

Multicast filtering mode

Discovery protocol

Type

Operation

Administrator contact

Name

Other details

CANCEL

OK

14. For the vMotion port group, right-click the vDS, choose Distributed Port Group, and choose New Distributed Port Group.

15. Enter VMkernel-vMotion as the name and click NEXT.
16. Set the VLAN type to VLAN, enter the VLAN ID used for vMotion, click the Customize default policies configuration check box, and click NEXT.
17. Leave the Security options set to Reject and click NEXT.
18. Leave the Ingress and Egress traffic shaping options as Disabled and click NEXT.
19. Choose Uplink 1 from the list of Active uplinks and click the move down tab twice to place Uplink 1 in the list of Standby uplinks. This will pin all vMotion traffic to Cisco UCS Fabric Interconnect B except when a failure occurs.

The screenshot shows the configuration window for a 'New Distributed Port Group'. On the left is a navigation pane with steps 1 through 8. Step 5, 'Teaming and failover', is selected and highlighted in dark blue. The main area is titled 'Teaming and failover' and contains the following settings:

- Load balancing:** Route based on originating virtual port (dropdown)
- Network failure detection:** Link status only (dropdown)
- Notify switches:** Yes (dropdown)
- Failback:** Yes (dropdown)
- Failover order:** Includes a help icon and buttons for MOVE UP, MOVE DOWN, SELECT ALL, and DESELECT ALL.
- Active uplinks:** Uplink 2 (checkbox is unchecked)
- Standby uplinks:** Uplink 1 (checkbox is checked)
- Unused uplinks:** (empty list)

At the bottom right of the window are three buttons: CANCEL, BACK, and NEXT. The NEXT button is highlighted in dark blue.

20. Click NEXT.
21. Leave NetFlow disabled and click NEXT.
22. Leave Block all ports set as No and click NEXT.
23. Confirm the options and click FINISH to create the port group.
24. Right-click the vDS and choose Add and Manage Hosts.

25. Make sure Add hosts is selected and click NEXT.

26. Click the + sign to add New hosts. Choose the FlashStack ESXi hosts and click OK. Click NEXT.

27. Choose vmnic2 and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 and click Assign uplink. Choose Uplink 2 and click OK. If more than one host is being connected to the vDS, use the Apply this uplink assignment to the rest of the hosts checkbox.



It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

vDS0 - Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Manage physical adapters**
- 4 Manage VMkernel adapt...
- 5 Migrate VM networking
- 6 Ready to complete

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Assign uplink Unassign adapter View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
vm-host-infra-fcp-01.flashstack.c...			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vm-host-infra-fcp-02.flashstack...			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			

CANCEL BACK NEXT

28. Click NEXT.

29. Do not migrate any VMkernel ports and click NEXT.

30. Do not migrate any virtual machine networking ports. Click NEXT.

31. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port(s) to the ESXi Host

ESXi Host VM-Host-Infra-FCP-01, VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To add the vMotion VMkernel Port to the ESXi host(s) on the VMware vDS, follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.
2. In the center pane, click the Configure tab.
3. In the list under Networking, choose VMkernel adapters.
4. Choose Add Networking to Add host networking.
5. Make sure VMkernel Network Adapter is selected and click NEXT.
6. Choose BROWSE to the right of Select an existing network.
7. Choose vMotion on the vDS and click OK.
8. Click NEXT.
9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.

vm-host-infra-fcp-01.flashstack.com - 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.

VMkernel port settings

Network label

IP settings

MTU

TCP/IP stack

Available services

Enabled services

- vMotion
- Provisioning
- Fault Tolerance logging
- Management
- vSphere Replication
- vSphere Replication NFC
- vSAN
- vSphere Backup NFC

CANCEL

BACK

NEXT

10. Choose Use static IPv4 settings and input the host's vMotion IPv4 address and Subnet mask.

11. Click NEXT.

vm-host-infra-fcp-01.flashstack.com - 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 settings selections before finishing the wizard.

Distributed port group	VMkernel-vmotion
Distributed switch	vDS0
vMotion	Enabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled
vSphere Backup NFC	Disabled

NIC settings

MTU	9000
TCP/IP stack	Default

IPv4 settings

IPv4 address	192.168.30.111 (static)
Subnet mask	255.255.255.0

CANCEL BACK FINISH

12. Review the parameters and click FINISH to add the vMotion VMkernel port.

Add and Configure a VMware ESXi Host in vCenter

This section details the steps to add and configure an ESXi host in vCenter. This section assumes the host has had VMware ESXi 7.0 U2 installed, the management IP address set, the nfnic driver updated and the Cisco UCS Tool installed. This procedure is initially being run on the second and third ESXi management hosts but can be run on any added ESXi host.

Add the ESXi Hosts to vCenter

To add the ESXi host(s) to vCenter, follow these steps:

1. From the Home screen in the VMware vCenter HTML5 Interface, choose Menu > Hosts and Clusters.
2. Right-click the “FlashStack-Management” cluster and click Add Hosts.
3. In the IP address or FQDN field, enter either the IP address or the FQDN name of the configured VMware ESXi host. Also enter the user id (root) and associated password. If more than one host is being added, add the corresponding host information, optionally selecting “Use the same credentials for all hosts”. Click NEXT.
4. Choose all hosts being added and click OK to accept the certificate(s).

5. Review the host details and click NEXT to continue.

6. Review the configuration parameters and click FINISH to add the host(s).

The screenshot shows the 'Add hosts' wizard in vCenter. The left pane is titled 'Add hosts' and contains a progress list with three steps: '1 Add hosts', '2 Host Summary', and '3 Ready to Complete'. The '3 Ready to Complete' step is highlighted with a dark background. The right pane is titled 'Review and finish' and contains the text: '2 new hosts will be connected to vCenter Server and moved to this cluster:' followed by the host names 'VM-Host-Infra-FCP-02' and 'VM-Host-Infra-FCP-03'. At the bottom right of the wizard, there are three buttons: 'CANCEL', 'BACK', and 'FINISH'. The 'FINISH' button is highlighted in green.

The added ESXi host(s) will be placed in Maintenance Mode and will have Warnings that the ESXi Shell and SSH have been enabled. These warnings can be suppressed.

Set Up VMkernel Ports and Virtual Switch for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To set up the VMkernel ports and the virtual switches on the ESXi host, follow these steps:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.
2. In the center pane, choose the Configure tab.
3. In the list, choose Virtual switches under Networking.
4. Expand Standard Switch: vSwitch0.
5. Choose EDIT to Edit settings.

6. Change the MTU to 9000.
7. Choose Teaming and failover located on the left.
8. In the Failover order section, use the arrow icons to move the vmnics until both are Active adapters.

vSwitch0 - Edit Settings

Properties

Load balancing: Route based on originating virtual port

Security

Network failure detection: Link status only

Traffic shaping

Teaming and failover

Notify switches: Yes

Failback: Yes

Failover order

↑ ↓

Active adapters
vmnic1
vmnic0

Standby adapters

Unused adapters

All Properties CDP LLDP

Adapter

Name: Cisco Systems Inc Cisco VIC Ethernet NI
vmnic1

Location: PCI 0000:69:00.1

Driver: nenic

Status

Status: Connected

Actual speed, Duplex: 40 Gbit/s, Full Duplex

Configured speed, Duplex: 40 Gbit/s, Full Duplex

Networks: 10.1.164.1-10.1.164.31 (VLAN115)

SR-IOV

Status: Not supported

Select active and standby adapters. During a failover, standby adapters activate in the order specified above.

9. Click OK.
10. In the center pane, to the right of VM Network click ... > Remove to remove the port group. Click YES on the confirmation.
11. Click ADD NETWORKING to add a new VM port group.
12. Choose Virtual Machine Port Group for a Standard Switch and click NEXT.
13. Ensure vSwitch0 is shown for Select an existing standard switch and click NEXT.
14. Name the port group "IB-MGMT Network" and leave the VLAN ID field set to None (0). Click NEXT.



In the Cisco UCS section of this document, the IB-MGMT VLAN was set as the native VLAN for the vSwitch0 vNIC templates, allowing DHCP to be used on ESXi vmk0 without putting in a VLAN ID for this port. Since this port group is in the same VLAN, the port group's VLAN ID should also be set to 0.

vm-host-infra-fcp-02.flashstack.com - Add Networking

✓ 1 Select connection type
✓ 2 Select target device
3 Connection settings
4 Ready to complete

Connection settings
Use network labels to identify migration-compatible connections common to two or more hosts.

Network label	IB-MGMT Network
VLAN ID	None (0) ▼

CANCEL BACK NEXT

15. Click FINISH to complete adding the IB-MGMT Network VM port group.
16. Click ADD NETWORKING to add a new VM port group.
17. Choose Virtual Machine Port Group for a Standard Switch and click NEXT.
18. Ensure vSwitch0 is shown for Select an existing standard switch and click NEXT.
19. Name the port group "OOB-MGMT Network" and input <OOB-MGMT-vlan-id> for the VLAN ID field. Click NEXT.

vm-host-infra-fcp-02.flashstack.com - Add Networking

- ✓ 1 Select connection type
- ✓ 2 Select target device
- 3 Connection settings**
- 4 Ready to complete

Connection settings

Use network labels to identify migration-compatible connections common to two or more hosts.

Network label	OOB-MGMT Network
VLAN ID	15

CANCEL

BACK

NEXT

20. Click FINISH to complete adding the OOB-MGMT Network VM port group.

21. Under Networking, choose Virtual switches. Expand vSwitch0. The properties for vSwitch0 should be like the following example:

Standard Switch: vSwitch0 | ADD NETWORKING | EDIT | MANAGE PHYSICAL ADAPTERS | ...

The screenshot displays the configuration for a Standard Switch named vSwitch0. On the left, three network port groups are listed:

- IB-MGMT Network**: VLAN ID: --, Virtual Machines (0)
- Management Network**: VLAN ID: --, VMkernel Ports (1) including vmk0 : 10.1164.112
- OOB-MGMT Network**: VLAN ID: 15, Virtual Machines (0)

On the right, the **Physical Adapters** section is expanded, showing two adapters:

- vmnic0 40000 Full
- vmnic1 40000 Full

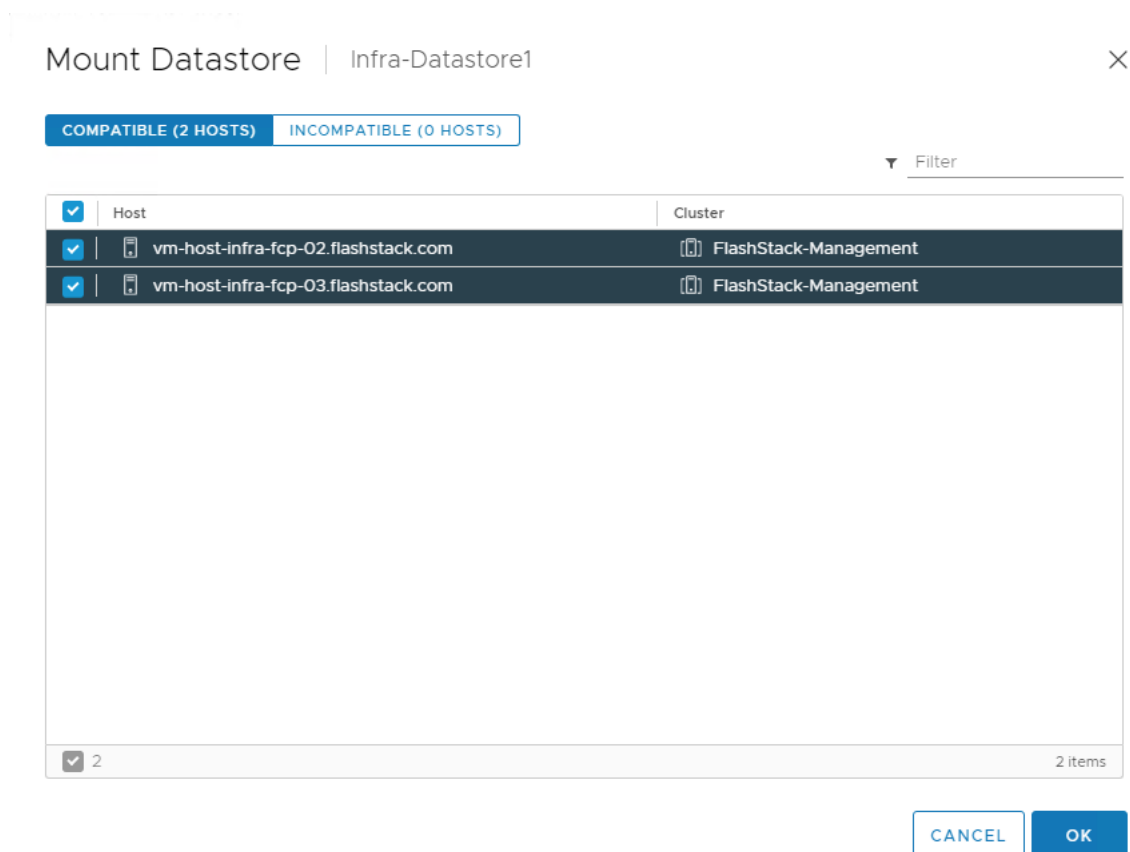
A central vertical bar indicates the connection status of each port group to the physical adapters, with green squares indicating active connections.

22. Repeat steps 1-21 for all hosts being added.

Mount Required Datastores for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To mount the required datastores, follow these steps on the ESXi host(s):

1. From the vCenter Home screen, choose Menu > Storage.
2. Expand FlashStack-DC.
3. Located on the left, right-click Infra-DataStore1 and choose Mount Datastore to Additional Hosts.
4. Choose the ESXi host(s) and click OK.



5. Repeat steps 1-4 to mount the Infra-Swap datastore to the ESXi host(s).
6. Choose Infra-DataStore1. In the center pane, choose Hosts. Verify the ESXi host(s) now has the datastore mounted. Repeat this process to also verify that Infra-Swap is also mounted.

Configure NTP on ESXi Host for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To configure Network Time Protocol (NTP) on the ESXi host(s), follow these steps:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.
2. In the center pane, choose the Configure tab.
3. In the list under System, choose Time Configuration.
4. To the right of Manual Time Configuration, click EDIT.
5. Set the correct local time and click OK.
6. To the right of Network Time Protocol, click EDIT.
7. Choose the Enable checkbox.
8. Enter the two Nexus switch NTP IP addresses in the NTP servers box separated by a comma.
9. Click the Start NTP Service checkbox.
10. Use the drop-down list to choose Start and stop with host.

Edit Network Time Protocol | vm-infra-esxi-01.flashstack.com

Enable ⓘ

NTP Servers: 10.1164.61,10.1164.62

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

NTP Service Status: Stopped

Start NTP Service

NTP Service Startup Policy: Start and stop with host

CANCEL OK

11. Click OK to save the configuration changes.
12. Verify that NTP service is now enabled and running, and the clock is now set to approximately the correct time.

Configure ESXi Host Swap for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To configure host swap on the ESXi host(s), follow these steps on the host:

-
1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.
 2. In the center pane, choose the Configure tab.
 3. In the list under System, choose System Swap.
 4. Located on the right, click EDIT.
 5. Choose Can use datastore and use the drop-down list to choose infra_swap. Leave all other settings unchanged.
 6. Click OK to save the configuration changes.
 7. In the list under Virtual Machines, choose Swap File Location.
 8. Located on the right, click EDIT.
 9. Choose Infra-Swap and click OK.

Change ESXi Power Management Policy for Cisco UCS M6 Hosts for the ESXi Host VM-Host-Infra-02 and VM-Host-Infra-03

To change the ESXi power management policy for the Cisco UCS M6 hosts, follow these steps:



Implementation of this policy is recommended in Performance Tuning for Cisco UCS M6 Server with Intel 3rd Gen Processors for maximum performance. If your organization has specific power policies, please set this policy accordingly.

1. In the list under Hardware, choose Overview. Scroll to the bottom and to the right of Power Management, choose EDIT POWER POLICY.
2. Choose High performance and click OK.

Edit Power Policy Settings

vm-host-infra-fcp... X

High performance

Do not use any power management features

Balanced

Reduce energy consumption with minimal performance compromise

Low power

Reduce energy consumption at the risk of lower performance

Custom

User-defined power management policy

CANCEL

OK

Check ESXi Host Fibre Channel Pathing for the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

For the fibre channel SAN-booted ESXi hosts, to ensure that the host(s) boot disk contains all required fibre channel paths, follow these steps:

1. In the list under Storage, choose Storage Devices. Make sure the Pure FlashArray Fibre Channel Disk is selected.
2. Choose the Paths tab.
3. Ensure that 4 fibre channel paths appear, all four should have the status Active (I/O).

Storage Devices

REFRESH ATTACH DETACH RENAME TURN ON LED TURN OFF LED ERASE PARTITIONS MARK AS HDD DISK MARK AS LOCAL MARK AS PERENNIALY RESERVED

<input type="checkbox"/>	Name	LUN	Type	Capacity	Datastore	Operational Status
<input type="checkbox"/>	Local ATA Disk (t10.ATA_____CISCO_VD_____0468bd9cb866001000000000...)	0	disk	223.51 GB	Not Consumed	Attached
<input type="checkbox"/>	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd400011fa9)	3	disk	1.00 TB	ESXi-Swap	Attached
<input checked="" type="checkbox"/>	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd4000141a2)	238	disk	1.00 TB	Infra-Data...	Attached
<input type="checkbox"/>	PURE Fibre Channel Disk (naa.624a9370b6c770713cae4dd400011fa9)	353	disk	1.00 TB	Infra-Data...	Attached

1 EXPORT 1 - 20 of 28 items

Properties Paths Partition Details

ENABLE DISABLE

<input type="radio"/>	Runtime Name	Status	Target	Name	Preferred
<input type="radio"/>	vmhba1:CO:T20:L238	Active (I/O)	52:4a:93:77:de:d7:21:00 52...	vmhba1:CO:T20:L238	
<input type="radio"/>	vmhba0:CO:T23:L238	Active (I/O)	52:4a:93:77:de:d7:21:02 52...	vmhba0:CO:T23:L238	
<input type="radio"/>	vmhba1:CO:T21:L238	Active (I/O)	52:4a:93:77:de:d7:21:10 52...	vmhba1:CO:T21:L238	
<input type="radio"/>	vmhba0:CO:T24:L238	Active (I/O)	52:4a:93:77:de:d7:21:12 52...	vmhba0:CO:T24:L238	

Add the ESXi Host(s) to the VMware Virtual Distributed Switch to the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

Follow this procedure if there are hosts to be added to vDS, skip if already added earlier. To add the ESXi host(s) to the VMware vDS, follow these steps on the host:

1. After logging into the VMware vSphere HTML5 Client, choose Networking under Menu.
2. Right-click the vDS (vDS0) and click Add and Manage Hosts.
3. Make sure Add hosts is selected and click NEXT.
4. Click the green + sign to add New hosts. Choose the configured FlashStack Management host(s) and click OK. Click NEXT.
5. Choose vmnic2 on each host and click Assign uplink. Choose Uplink 1 and click OK. Choose vmnic3 on each host and click Assign uplink. Choose Uplink 2 and click OK. If more than one host is being connected to the vDS, use the Apply this uplink assignment to the rest of the hosts checkbox.



It is important to assign the uplinks as shown below. This allows the port groups to be pinned to the appropriate Cisco UCS fabric.

vDS0 - Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Manage physical adapters**
- 4 Manage VMkernel adapt...
- 5 Migrate VM networking
- 6 Ready to complete

Manage physical adapters

Add or remove physical network adapters to this distributed switch.

Assign uplink Unassign adapter View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
vm-host-infra-fcp-02.flashstack...			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vm-host-infra-fcp-03.flashstack...			
On this switch			
vmnic2 (Assigned)	--	Uplink 1	vDS0-DVUplinks-...
vmnic3 (Assigned)	--	Uplink 2	vDS0-DVUplinks-...
On other switches/unclaimed			

CANCEL

BACK

NEXT

6. Click NEXT.
7. Do not migrate any VMkernel ports and click NEXT.
8. Do not migrate any VM ports and click NEXT.
9. Click FINISH to complete adding the ESXi host(s) to the vDS.

Add the vMotion VMkernel Port(s) to the ESXi Host to the ESXi Host VM-Host-Infra-FCP-02 and VM-Host-Infra-FCP-03

To add the vMotion VMkernel Port to the ESXi host(s) on the VMware vDS, follow these steps on the host:

1. In the vCenter HTML5 Interface, under Hosts and Clusters choose the ESXi host.
2. In the center pane, click the Configure tab.
3. In the list under Networking, choose VMkernel adapters.
4. Choose Add Networking to Add host networking.
5. Make sure VMkernel Network Adapter is selected and click NEXT.

-
6. Choose BROWSE to the right of Select an existing network.
 7. Choose vMotion on the vDS and click OK.
 8. Click NEXT.
 9. Make sure the Network label is vMotion with the vDS in parenthesis. From the drop-down list, select Custom for MTU and make sure the MTU is set to 9000. Choose the vMotion TCP/IP stack and click NEXT.
 10. Choose Use static IPv4 settings and input the host's vMotion IPv4 address and Subnet mask.
 11. Click NEXT.
 12. Review the parameters and click FINISH to add the vMotion VMkernel port.
 13. If this is an iSCSI-booted host, execute the instructions in the Appendix for an iSCSI-booted host being added in vCenter.
 14. Exit Maintenance Mode on each ESXi host in Maintenance Mode.

VMware ESXi 7.0 U2 TPM Attestation

If your Cisco UCS servers have Trusted Platform Module (TPM) 2.0 modules installed, the TPM can provide assurance that ESXi has booted with UEFI Secure Boot enabled and using only digitally signed code. In the Cisco UCS section of this document, UEFI secure boot was enabled in the boot 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. Follow these steps:

1. If your Cisco UCS servers have TPM 2.0 modules installed, TPM Attestation can be verified in the vSphere HTML5 Client. To get to the HTML5 client from the Web Client, click "Launch vSphere Client (HTML5) in the upper center portion of the Web Client window.
2. From the Hosts and Clusters window in the vSphere Client, click the FlashStack-Management cluster. In the center pane, click Monitor > Security. The Attestation status will appear as shown below, where 2 of the 3 hosts have TPM 2.0 modules installed:

FlashStack-Management | ACTIONS

Summary Monitor Configure Permissions Hosts VMs Datastores Networks Updates

Overview
Advanced
Tasks and Events
Tasks
Events
vSphere DRS
Recommendations
Faults
History
VM DRS Score
CPU Utilization
Memory Utilization
Network Utilization
vSphere HA
Summary
Heartbeat

Security

Name ↑	Attestation	Last verified	Attested by	TPM version	TXT	Message
vm-host-infra-fcp-01.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	
vm-host-infra-fcp-02.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	
vm-host-infra-fcp-03.flashstack.com	Passed	09/13/2021, 9:...	vCenter Server	2.0	N/A	



It may be necessary to disconnect and reconnect a host from vCenter to get it to pass attestation the first time. Also, in this example, only the second host had a TPM module installed.

FlashStack Management Tools Setup

Cisco Data Center Network Manager (DCNM)-SAN

Cisco DCNM-SAN can be used to monitor, configure, and analyze Cisco fibre channel fabrics. Cisco DCNM-SAN is deployed as a virtual appliance from an OVA and is managed through a web browser. SAN Analytics can be added to provide insights into your fabric by allowing you to monitor, analyze, identify, and troubleshoot performance issues.

Prerequisites

The following prerequisites need to be configured:

1. Licensing. Cisco DCNM-SAN includes a 60-day server-based trial license that can be used to monitor and configure Cisco MDS Fibre Channel switches and monitor Cisco Nexus switches. Both DCNM server-based and switch-based licenses can be purchased. Additionally, SAN Insights and SAN Analytics requires an additional switch-based license on each switch. Cisco MDS 32Gbps Fibre Channel switches provide a 120-day grace period to trial SAN Analytics.



If using the Cisco Nexus 93180YC-FX for SAN switching, it does not support SAN Analytics.

2. Passwords. Cisco DCNM-SAN passwords should adhere to the following password requirements:
 - f. It must be at least eight characters long and contain at least one alphabet and one numeral.
 - g. It can contain a combination of alphabets, numerals, and special characters.
 - h. Do not use any of these special characters in the DCNM password for all platforms: <SPACE> " & \$ % ' ^ = < > ; : ` \ | / , . *
3. DCNM SNMPv3 user on switches. Each switch (both Cisco MDS and Nexus) needs an SNMPv3 user added for DCNM to use to query and configure the switch. On each switch, enter the following command in configure terminal mode (in the example, the userid is snmpuser):
snmp-server user snmpadmin network-admin auth sha <password> priv aes-128 <privacy-password>
4. On Cisco MDS switches, type show run. If snmpadmin passphrase lifetime 0 is present, enter
username snmpadmin passphrase lifetime 99999 warntime 14 gracetime 3



It is important to use auth type sha and privacy auth aes-128 for both the switch and UCS snmpadmin users.

5. DCNM SNMPv3 user in UCSM. A SNMPv3 user needs to be added to UCSM to allow DCNM to query the LAN side of the fabric interconnects. In Cisco UCS Manager, click Admin. Navigate to All > Communication Management > Communication Services. Under SNMP, click Enabled, click Save

Changes, and then click OK. Under SNMP Users, click Add. Enter the user name and enter and confirm the Password and Privacy Password.

Create SNMP User



Name	:	<input type="text" value="snmpadmin"/>
Auth Type	:	SHA
Use AES-128	:	Yes
Password	:	<input type="password" value="*****"/>
Confirm Password	:	<input type="password" value="*****"/>
Privacy Password	:	<input type="password" value="*****"/>
Confirm Privacy Password	:	<input type="password" value="*****"/>



6. Click OK and then click OK again to complete adding the user.

Deploy the Cisco DCNM-SAN OVA

To deploy the Cisco DCNM-SAN OVA, follow these steps:

1. Download the Cisco DCNM 11.5.1 Open Virtual Appliance for VMware from [https://software.cisco.com/download/home/281722751/type/282088134/release/11.5\(1\)](https://software.cisco.com/download/home/281722751/type/282088134/release/11.5(1)). Extract dcnm-va.11.5.1.ova from the ZIP file.
2. In the VMware vCenter HTML5 interface, click Menu > Hosts and Clusters.
3. Right-click the FlashStack-Management cluster and select Deploy OVF Template.
4. Choose Local file then click UPLOAD FILES. Navigate to choose dcnm-va.11.5.1.ova and click Open. Click NEXT.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

Local file

dcnm-va.11.5.1.ova

5. Name the virtual machine and choose the FlashStack-DC datacenter. Click NEXT.
6. Choose the FlashStack-Management cluster and click NEXT.
7. Review the details and click NEXT.
8. Scroll through and accept the license agreements. Click NEXT.
9. Choose the appropriate deployment configuration size and click NEXT.



If using the SAN Insights and SAN Analytics feature, it is recommended to use the Huge size.

Deploy OVF Template

- Select an OVF template
- Select a name and folder
- Select a compute resource
- Review details
- License agreements
- 6 Configuration**
- Select storage
- Select networks
- Customize template
- Ready to complete

Configuration

Select a deployment configuration

<input type="radio"/> Large (Production)	Description Use this deployment option to configure a huge version of appliance with 32vCPUs and 128GB RAM. This is recommended when using SAN Insights feature.
<input type="radio"/> Small (Lab/PoC)	
<input checked="" type="radio"/> Huge	
<input type="radio"/> Compute	
<input type="radio"/> ComputeHuge	

5 Items

CANCEL BACK NEXT

10. Choose Infra-DataStore1 and the Thin Provision virtual disk format. Click NEXT.

11. Choose IB-MGMT Network for all three Source Networks. Click NEXT.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 License agreements
- 6 Configuration
- 7 Select storage
- 8 Select networks**
- 9 Customize template
- 10 Ready to complete

Select networks ✕

Select a destination network for each source network.

Source Network	Destination Network
dcnm-mgmt	IB-Mgmt ▼
enhanced-fabric-mgmt	IB-Mgmt ▼
enhanced-fabric-inband	IB-Mgmt ▼

3 items

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL
BACK
NEXT

12. Fill-in the management IP address, subnet mask, and gateway. Set the Extra Disk Size according to how many Cisco MDS switches you will be monitoring with this DCNM. If you are only monitoring the two Cisco MDS switches in this FlashStack deployment, set this field to 32. Click NEXT.

13. Review the settings and click FINISH to deploy the OVA.

Deploy OVF Template

- 1 Select an OVF template
- 2 Select a name and folder
- 3 Select a compute resource
- 4 Review details
- 5 License agreements
- 6 Configuration
- 7 Select storage
- 8 Select networks
- 9 Customize template
- 10 Ready to complete

Ready to complete

Click Finish to start creation.

Name	DCNM
Template name	dcnm
Download size	5.3 GB
Size on disk	Unknown
Folder	FlashStack-DC
Resource	FlashStack-Management
Storage mapping	1
All disks	Datastore: Infra-DataStore1; Format: Thick provision lazy zeroed
Network mapping	3
dcnm-mgmt	IB-Mgmt
enhanced-fabric-mgmt	IB-Mgmt
enhanced-fabric-inband	IB-Mgmt
IP allocation settings	
IP protocol	IPV4
IP allocation	Static - Manual
Properties	1.IP Address = 10.1.164.41 2.Subnet Mask = 255.255.255.0

CANCEL
BACK
FINISH

14. After deployment is complete, right-click the newly deployed DCNM VM and click Edit Settings. Expand CPU and adjust the Cores per Socket setting until the number of Sockets is set to match the number of CPUs in the UCS servers used in this deployment. The following example shows 2 sockets.

Edit Settings

dcnm



Virtual Hardware


VM Options

ADD NEW DEVICE

▼ CPU *	32 ▼	
Cores per Socket	16 ▼ Sockets: 2	
CPU Hot Plug	<input type="checkbox"/> Enable CPU Hot Add	
Reservation	0 <input type="text"/> MHz ▼	
Limit	Unlimited <input type="text"/> MHz ▼	
Shares	Normal ▼ 32000	
CPUID Mask	Expose the NX/XD flag to guest ▼ Advanced...	
Hardware virtualization	<input type="checkbox"/> Expose hardware assisted virtualization to the guest OS	
Performance Counters	<input type="checkbox"/> Enable virtualized CPU performance counters	

CANCEL

OK

15. Click OK to complete the change.
16. Right-click the newly deployed DCNM VM and click Open Remote Console. Once the console is up, click  to power on the VM. Once the VM has powered up, point a web browser to the URL displayed on the console.
17. Navigate the security prompts and click Get started.
18. Make sure Fresh installation - Standalone is selected and click Continue.
19. Choose SAN only for the Installation mode and leave Cisco Systems, Inc. for the OEM vendor and click Next.
20. Enter and repeat the administrator and database passwords and click Next.
21. Enter the DCNM FQDN, a comma-separated list of DNS servers, a comma-separated list of NTP servers, and select the appropriate time zone. Click Next.

-
22. The Management Network settings should be filled in. For Out-of-Band Network, a different IP address in the same subnet as the management address should be used. Only input the IPV4 address with prefix. Do not put in the Gateway IPv4 Address. Scroll down and click Next.
 23. Leave Internal Application Services Network set at the default setting and click Next.
 24. Review the Summary details and click Start installation.
 25. When the Installation status is complete, click Continue.
 26. In the vCenter HTML5 client under Hosts and Clusters, choose the DCNM VM and click the Summary Tab. If an alert is present that states “A newer version of VMware Tools is available for this virtual machine.”, click Upgrade VMware Tools. Choose Automatic Upgrade and click UPGRADE. Wait for the VMware Tools upgrade to complete.

Configure DCNM-SAN

To configure the DCNM-SAN, follow these steps:




When the DCNM installation is complete, the browser should redirect to the DCNM management URL.

1. Log in as admin with the password entered above.
2. On the message that appears, choose Do not show this message again and click No.



If you have purchased DCNM server-based or switch-based licenses, follow the instructions that came with the licenses to install them. A new DCNM installation also has a 60-day trial license.

3. In the menu on the left, click Inventory > Discovery > LAN Switches.
4. Click  to add LAN switches. In the Add LAN Devices window, enter the mgmt0 IP address of Nexus switch A in the Seed Switch box. Enter the snmpadmin user name and password set up in the Prerequisites section above. Set Auth-Privacy to SHA_AES. Click Next.

Add LAN Devices

Discovery Type: Hops from seed switch Switch list

Seed Switch:

Max Hops from Seed:

User Name:

Password:

Auth-Privacy:

Add Switches To Group:

Scan Time:

- LAN switch discovery will take a few minutes. In the LAN Discovery list that appears, the two Nexus switches and two Fabric Interconnects that are part of this FlashStack should appear with a status of “manageable”. Using the checkboxes on the left, choose the two Nexus switches and two Fabric Interconnects that are part of this FlashStack. Click Add.
- After a few minutes (click the Refresh icon in the upper right-hand corner), the two Nexus switches and two Fabric Interconnects that are part of this FlashStack will appear with detailed information. The SSH warning under SNMP Status can be ignored since only SNMP can be used to monitor Fabric Interconnects.


Data Center Network Manager

Inventory / Discovery / LAN Switches

	Switch	IP Address	Serial No	Managed	SNMP Status	Role	Last Updated Time	Group	User
1	<input type="checkbox"/> BB08-91380YX-FX-01	10.1.164.61	FDO24240CU3	true	ok		2021-09-05 19:24:49	Default_LAN	admin
2	<input type="checkbox"/> BB08-91380YX-FX-02	10.1.164.62	FDO24240CTN	true	ok		2021-09-05 19:24:49	Default_LAN	admin

- In the menu on the left, click Inventory > Discovery > SAN Switches.



- Click  to add a switching fabric.
- Enter either the IP address or hostname of the first Cisco MDS 9132T switch. Leave Use SNMPv3/SSH selected. Set Auth-Privacy to SHA_AES. Enter the snmpadmin user name and password set up in the Prerequisites section above. Click Options>>. Enter the UCS admin user name and password. Click Add.



If the Cisco Nexus 93180YC-FX switches are being used for SAN switching, substitute them here for MDS 9132Ts. They will need to be added again under SAN switches since LAN and SAN switching are handled separately in DCNM.

Add Fabric

Fabric Seed Switch:

SNMP: Use SNMPv3/SSH

Auth-Privacy: ▼

User Name:

Password:

Limit Discovery by VSAN

Enable NPV Discovery in All Fabrics

10. Repeat steps 1–9 to add the second Cisco MDS 9132T and Fabric Interconnect.

The two SAN fabrics should now appear in the Inventory.

The screenshot shows the Data Center Network Manager interface. The breadcrumb navigation is "Inventory / Discovery / SAN Switches". Below the navigation are buttons for "+", "X", "Move", "Rediscover", and "Purge". A table displays the discovered SAN switches:

Name	SeedSwitch	Status	SNMPv3/SSH
Fabric_BB08-MDS-9132T-B	10.1.164.64	managedContinuously	true
Fabric_BB08-MDS-9132T-A	10.1.164.63	managedContinuously	true

11. Choose Inventory > Discovery > Virtual Machine Manager.



12. Click to add the vCenter.

13. In the Add VCenter window, enter the IP address of the vCenter VCSA. Enter the administrator@vsphere.local user name and password. Click Add.

14. The vCenter should now appear in the inventory.


15. Choose Administration > Performance Setup > LAN Collections.


16. Choose the Default_LAN group and all information you would like to collect. Click Apply. Click Yes to restart the Performance Collector.

For all selected licensed LAN Switches collect: Trunks Access Errors & Discards Temperature Sensor

Apply

▼ Default_LAN

 BB08-91380YX-FX-01

 BB08-91380YX-FX-02

17. Choose Administration > Performance Setup > SAN Collections.

18. Choose both fabrics. Choose all information you would like to collect and click Apply. Click Yes to restart the Performance Collector.

Apply

		Name	ISL/NPV Links	Hosts	Storage	FC Flows
1	<input type="checkbox"/>	Fabric_AA12-FS-9132T-2				
2	<input type="checkbox"/>	Fabric_AA12-FS-Prod-UCS645...				
3	<input checked="" type="checkbox"/>	Fabric_BB08-MDS-9132T-A	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
4	<input checked="" type="checkbox"/>	Fabric_BB08-MDS-9132T-B	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

19. Choose Configure > SAN > Device Alias. Since device-alias mode enhanced was configured in the Cisco MDS 9132T switches, Device Aliases can be created and deleted from DCNM and pushed to the MDS switches.

20. Choose Configure > SAN > Zoning. Just as Device Aliases can be created and deleted from DCNM, zones can be created, deleted, and modified in DCNM and pushed to the MDS switches. Remember to enable Smart Zoning and to Zone by Device Alias.

You can now explore all of the different options and information provided by DCNM SAN. See [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5\(1\)](#).

Configure SAN Insights in DCNM SAN

The SAN Insights feature enables you to configure, monitor, and view the flow analytics in fabrics. Cisco DCNM enables you to visually see health-related indicators in the interface so that you can quickly identify issues in fabrics. Also, the health indicators enable you to understand the problems in fabrics. The SAN Insights feature also provides more comprehensive end-to-end flow-based data from host to LUN.

- Ensure that the time configurations set above, including daylight savings settings are consistent across the MDS switches and Cisco DCNM.

- SAN Insights requires installation of a switch-based SAN Analytics license on each switch. To trial the feature, each switch includes a one-time 120-day grace period for SAN Analytics from the time the feature is first enabled.
- SAN Insights supports current Fibre Channel Protocol (SCSI) and NVMe over Fibre Channel (NVMe).
- SAN Insights works by enabling SAN Analytics and Telemetry Streaming on each switch. The switches then stream the SAN Analytics data to DCNM, which collects, correlates, and displays statistics. All configurations can be done from DCNM.
- Only Cisco MDS switches support SAN Analytics. Nexus 93180YC-FX switches do not support SAN Analytics.
- For more information on SAN Insights, see the SAN Insights sections: [Cisco DCNM SAN Management for OVA and ISO Deployments Configuration Guide, Release 11.5\(1\)](#).
- For more information on SAN Analytics, see: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/san_analytics/cisco-mds9000-san-analytics-telemetry-streaming-config-guide-8x.html.

To configure SAN Insights in DCNM SAN, follow these steps:

1. Click Configure > SAN > SAN Insights. Click Continue.
2. Choose Fabric A. Click Continue.
3. Choose the Fabric A Cisco MDS switch. Under Install Query click None and from the drop-down list click Storage. Under Subscriptions, choose SCSI & NVMe. Optionally, under Receiver, choose the second IP address in the In-Band Management subnet configured for DCNM. Click Save, then click Continue.

2. Select Switches

Choose the switch(es) on which SAN Insights is to be configured in Fabric_aa13-9132t-a

DCNM server time: 10:06:10.494 EDT Tuesday August 11 2020

Selected 1 / Total 1

Disable Analytics...		Show Quick Filter						
<input type="checkbox"/>	Switch	Model	Release	Licensed	Switch Time	Subscriptions	Install Query	Receiver
<input checked="" type="checkbox"/>	aa13-9132t-a	DS-C9132T-K9	8.4(1a)	Yes	10:06:12.790 EDT Tue Aug 11 2020	SCSI	Storage	10.1.156.210

4. Review the information and click Continue.
5. Expand the switch and then the module. Under Enable / Disable SCSI Telemetry, click the left icon to enable telemetry on the ports connected to the FlashArray//X R3 Click Continue.

4. Select Interfaces

Choose the switch interfaces that will generate analytics data within Fabric_BB08-MDS-9132T-A

Total Top Level Rows 1

Switch	Module	Interface	Connected To	Type	Analytics Status	Enable / Disable SCSI Telemetry	Enable / Disable NVMe Telemetry
▼ BB08-MDS-9132...	1 module(s)	4 interface(s)		Storage			
▼	DS-C9132T-K9-S...	4 interface(s)					
		fc1/1	FlashArray-CT0FC0	both	disabled	<input checked="" type="checkbox"/> <input type="checkbox"/> pending enable	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/2	52.4a:93:77:de:d7:21:01	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/3	FlashArray-CT1FC0	both	disabled	<input checked="" type="checkbox"/> <input type="checkbox"/> pending enable	<input type="checkbox"/> <input checked="" type="checkbox"/>
		fc1/4	52.4a:93:77:de:d7:21:11	Storage	disabled	<input type="checkbox"/> <input checked="" type="checkbox"/>	<input type="checkbox"/> <input checked="" type="checkbox"/>

6. Review the information and click Commit to push the configuration to the Cisco MDS switch.
7. Ensure that the two operations were successful and click Close.
8. Repeat steps 1-7 to install SAN Analytics and Telemetry on the Fabric B switch.
9. After approximately two hours, you can view SAN Analytics data under the Dashboard and Monitor.

Cisco Intersight

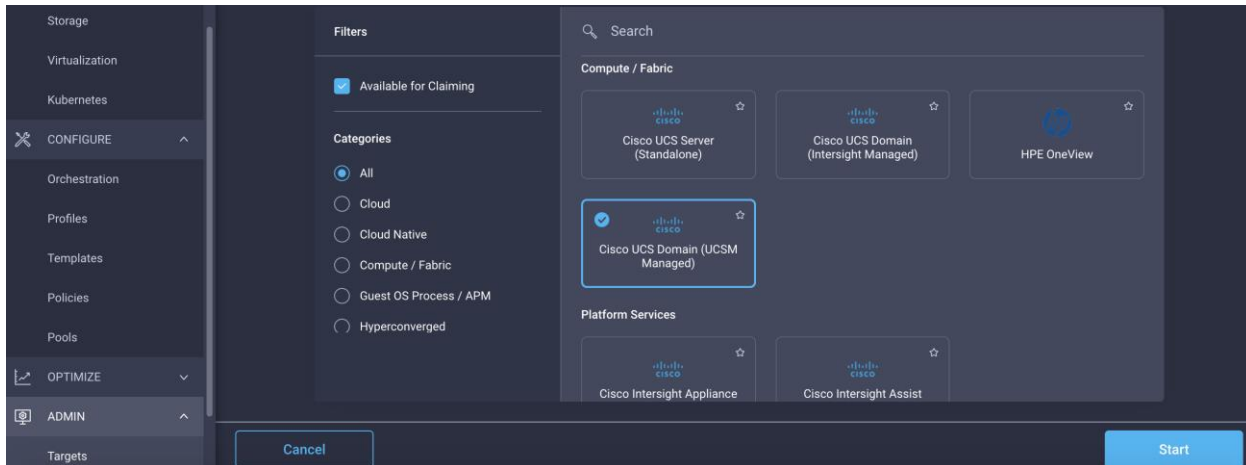
Cisco Intersight™ is a management platform delivered as a service with embedded analytics for your Cisco and 3rd party IT infrastructure. This platform offers an intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in more advanced ways than the prior generations of tools. Cisco Intersight provides an integrated and intuitive management experience for resources in the traditional data center and at the edge. With flexible deployment options to address complex security needs, getting started with Intersight is quick and easy.

Cisco Intersight offers flexible deployment either as Software as a Service (SaaS) on Intersight.com or running on your premises as Cisco Intersight Virtual Appliance. The virtual appliance provides the benefits of Cisco Intersight while allowing more flexibility for those with additional data locality and security requirements. The remainder of this section details Intersight deployment as SaaS on Intersight.com. To learn more about the virtual appliance, see the [Cisco Intersight Virtual Appliance Getting Started Guide](#).

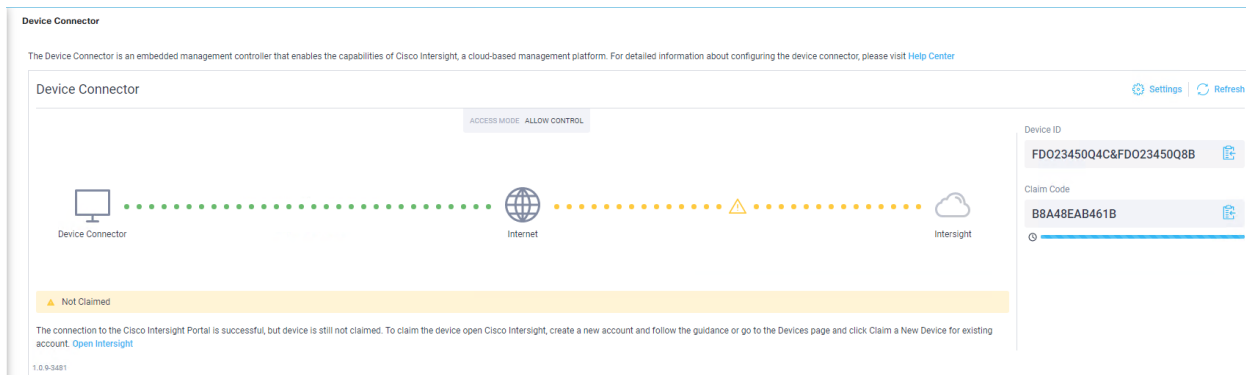
To configure Cisco Intersight, follow these steps:

1. If you do not already have a Cisco Intersight account, to claim your Cisco UCS system into a new account on Cisco Intersight, connect to <https://intersight.com>. If you have an existing Intersight account, connect to <https://intersight.com> and sign in with your Cisco ID, select the appropriate account, and skip to step 6.
2. Click Create an account.

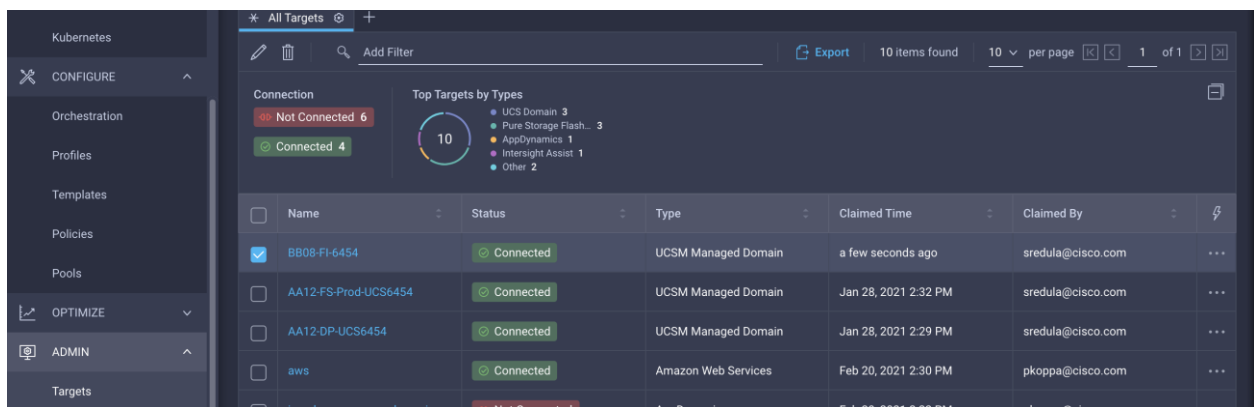
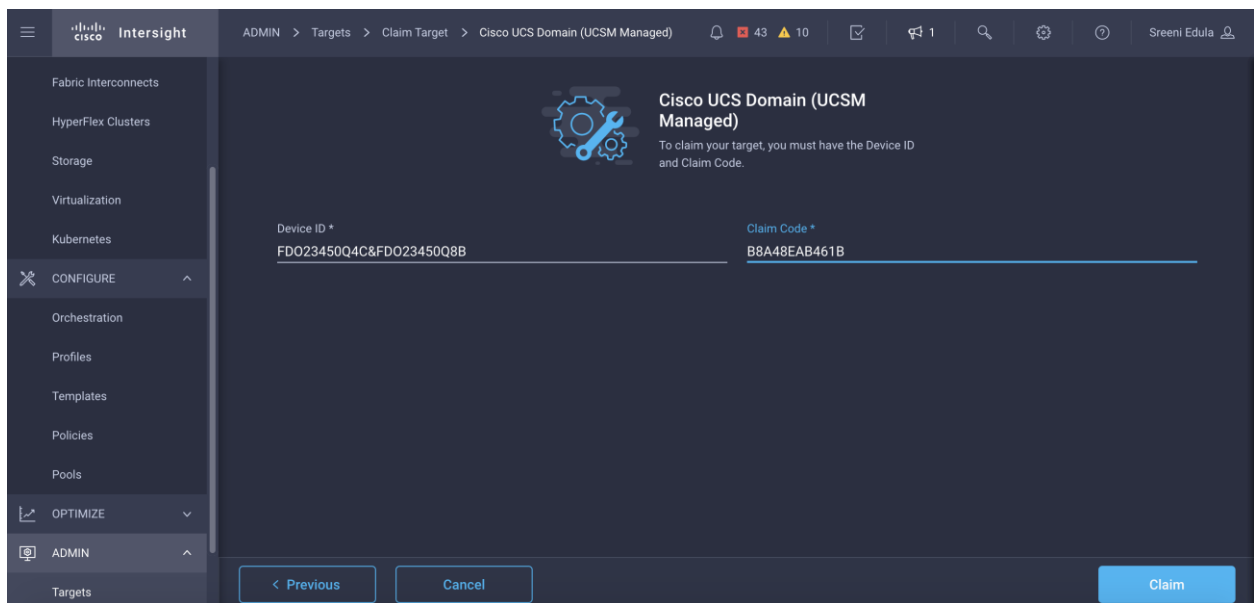
3. Sign in with your Cisco ID.
4. Read, scroll through, and accept the End User License Agreement and click Next.
5. Enter an Account Name and click Create.
6. Choose ADMIN > Targets. Click Claim a New Target. Select Cisco UCS Domain (UCSM Managed) and click Start. Fill in the Device ID and Claim Code and click Claim.




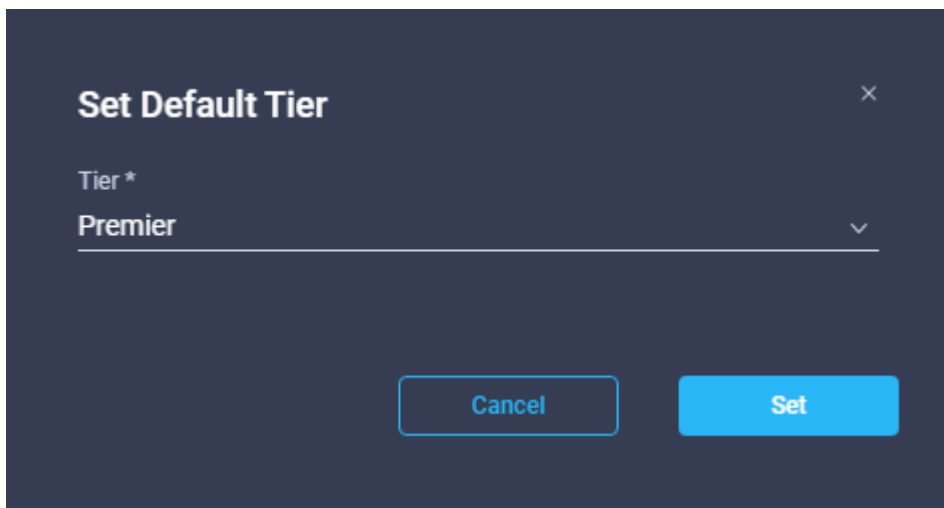
7. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right.






8. To claim your Cisco UCS system into an existing Intersight account, log into the account at <https://Intersight.com>. Choose Administration > Devices. Click Claim a New Device. Under Direct Claim, fill in the Device ID and Claim Code. The Device ID and Claim Code can be obtained by connecting to Cisco UCS Manager and selecting Admin > All > Device Connector. The Device ID and Claim Code are on the right.




9. From the Cisco Intersight window, click  and then click Licensing. If this is a new account, all servers connected to the UCS Domain will appear under the Base license Tier. If you have purchased Cisco Intersight licenses and have them in your Cisco Smart Account, click Register and follow the prompts to register this Cisco Intersight account to your Cisco Smart Account. Cisco Intersight also offers a one-time 90-day trial of Premier licensing for new accounts. Click Start Trial and then Start to begin this evaluation. The remainder of this section will assume Premier licensing.
10. From the Licensing Window, click Set Default Tier. From the drop-down list choose Premier for Tier and click Set.



11. To set all Cisco UCS Servers to Premier licensing, click Servers. Click  to the left of the Name heading to choose all servers. Click  above the headings and click Set License Tier. From the drop-down list choose Premier for the Tier and click Set License Tier.

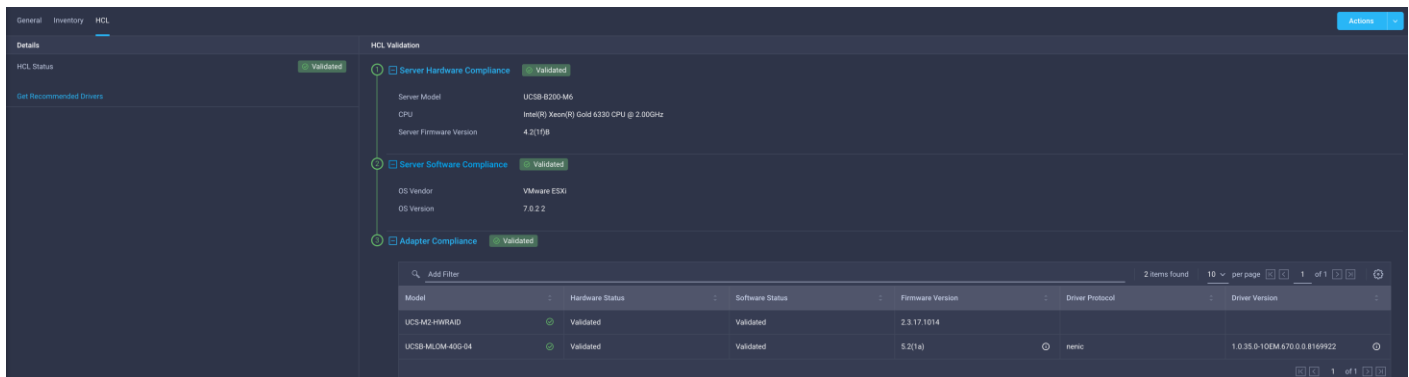


12. Click Refresh to refresh the Intersight window with Premier, Advantage, and Essentials features added.

13. Click  in the Intersight window and click Take a Site Tour. Follow the prompts for a tour of Cisco Intersight.

14. The Essentials tier of Cisco Intersight includes a Cisco driver check against the Cisco Hardware Compatibility List (HCL). In the Servers list, choose one of the servers in your VMware FlashStack-Management cluster by clicking the server name. Review the detailed General and Inventory infor-

mation for the server. Click the HCL tab. Review the server information, the version of VMware ESXi, and the Cisco VIC driver versions.



- Using the Intersight Assist personality of the Cisco Intersight Virtual Appliance VMware vCenter, it can be monitored (Advantage Licensing Tier) and configured (Premier Licensing Tier). To install Intersight Assist from an Open Virtual Appliance (OVA) in your VMware FlashStack-Management Cluster, first download the latest release of the OVA from <https://software.cisco.com/download/home/286319499/type/286323047/release/1.0.9-230>.
- Refer to https://www.cisco.com/c/en/us/td/docs/unified_computing/Intersight/cisco-intersight-assist-getting-started-guide/m-installing-cisco-intersight-assist.html and set up the DNS entries for the Intersight Assist hostname as specified under Before you begin.
- From Hosts and Clusters in the VMware vCenter HTML5 client, right-click the FlashStack-Management cluster and click Deploy OVF Template.
- Specify a URL or browse to the intersight-virtual-appliance-1.0.9-230.ova file. Click NEXT.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

URL

<http> | <https://remoteserver-address/filetoinstall.ovf> | .ova

Local file

UPLOAD FILES

intersight-virtual-appliance-1.0.9-148.ova

CANCEL

BACK

NEXT

19. Name the Intersight Assist VM and choose the location. Click NEXT.

20. Choose the FlashStack-Management cluster and click NEXT.

21. Review details and click NEXT.

22. Choose a deployment configuration (Tiny recommended) and click NEXT.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Configuration**
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Configuration

Select a deployment configuration

<input type="radio"/> Small(16 vCPU, 32 Gi RAM)	Description Deployment size supports Intersight Assist only.
<input type="radio"/> Medium(24 vCPU, 64 Gi RAM)	
<input checked="" type="radio"/> Tiny(8 vCPU, 16 Gi RAM)	
3 Items	

CANCEL

BACK

NEXT

23. Choose Infra-DataStore1 for storage and choose the Thin Provision virtual disk format. Click NEXT.

24. Choose IB-MGMT Network for the VM Network. Click NEXT.

25. Fill in all values to customize the template. Click NEXT.

26. Review the deployment information and click FINISH to deploy the appliance.

27. Once the OVA deployment is complete, right-click the Intersight Assist VM and click Edit Settings.

28. Expand CPU and adjust the Cores per Socket so that 2 Sockets are shown. Click OK.

Edit Settings | nx-intersight-assist




Virtual Hardware | VM Options

ADD NEW DEVICE

▼ CPU	8	▼	
Cores per Socket	4	▼	Sockets: 2
CPU Hot Plug	<input checked="" type="checkbox"/>	Enable CPU Hot Add	
Reservation	0	▼ MHz ▼	
Limit	Unlimited	▼ MHz ▼	
Shares	Normal	▼	8000
CPUID Mask	Expose the NX/XD flag to guest	▼	Advanced...
Hardware virtualization	<input type="checkbox"/>	Expose hardware assisted virtualization to the guest OS	
Performance Counters	<input type="checkbox"/>	Enable virtualized CPU performance counters	
CPU/MMU Virtualization	Automatic	▼	
> Memory	16	▼ GB ▼	
> Hard disks	8 total 500 GB		
> SCSI controller 0	LSI Logic SAS		

CANCEL OK

29. Right-click the Intersight Assist VM and choose Open Remote Console.

30. Click  to power on the VM.

31. When you see the login prompt, close the Remote Console and connect to <https://intersight-assist-fqdn>.



It may take a few minutes for <https://intersight-assist-fqdn> to respond.

32. Navigate the security prompts and select Intersight Assist. Click Proceed.

What would you like to Install ?

Intersight Connected Virtual Appliance ⓘ

Intersight Private Virtual Appliance ⓘ

Intersight Assist ⓘ

 Recover from backup

Proceed

33. From Cisco Intersight, click ADMIN > Devices. Click Claim a New Device. Copy and paste the Device ID and Claim Code shown in the Intersight Assist web interface to the Cisco Intersight Device Claim Direct Claim window. In Cisco Intersight, click Claim.

34. In the Intersight Assist web interface, click Continue.

35. The Intersight Assist software will now be downloaded and installed into the Intersight Assist VM. This can take up to an hour to complete.



The Intersight Assist VM will reboot during the software download process. It will be necessary to refresh the Web Browser after the reboot is complete to follow the status of the download process.

36. When the software download is complete, navigate the security prompts and an Intersight Assist login screen will appear. Log into Intersight Assist with the admin user and the password supplied in the OVA installation. Check the Intersight Assist status and log out of Intersight Assist.

37. To claim the vCenter, from Cisco Intersight, click ADMIN > Targets. Click Claim a New Target. In the Select Target Type window, select VMware vCenter under Hypervisor and click Start. In the VMware vCenter window, make sure the Intersight Assist is correctly selected, fill in the vCenter information, and click Claim.



VMware vCenter

To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

Intersight Assist *

flash-assist.flashstack.com

Hostname/IP Address *

vcenter.flashstack.com

Port

0 - 65535

Username *

administrator@vsphere.com

Password *

.....

Secure

Datastore Browsing Enabled

38. After a few minutes, the VMware vCenter will appear in the Devices list. It also can be viewed by clicking Intersight Assist in the Devices list.

39. Detailed information obtained from the vCenter can now be viewed by clicking Virtualization from the menu.

OPERATE > Virtualization > Datacenters

Datacenters Clusters Hosts Virtual Machines Datastores Datastore Clusters



Add Filter



Name

Datastores

Networks



FlashStack_DC

26

OPERATE > Virtualization > Hosts

Datcenters Clusters Hosts Virtual Machines Datastores Datastore Clusters

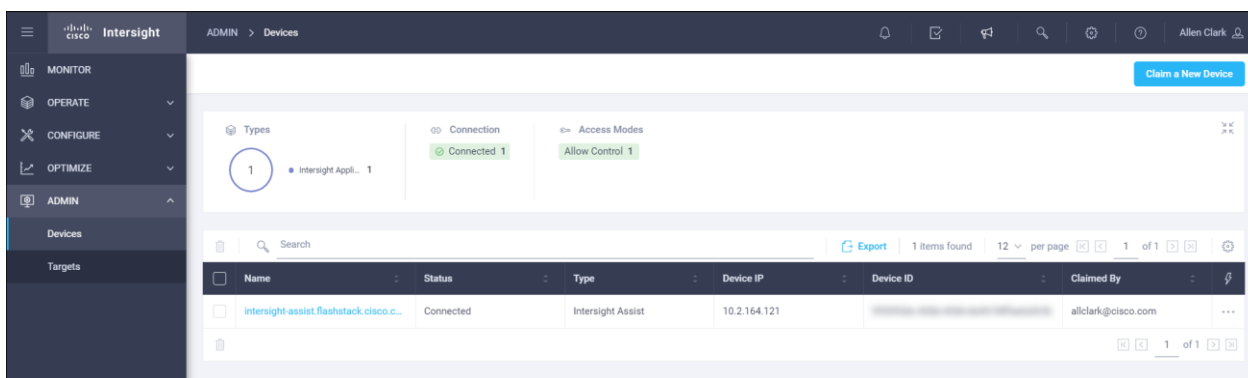
🔍 Add Filter

<input type="checkbox"/>	Name	Datacenter	Cluster	CPU
<input type="checkbox"/>	🟢 vm-infra-esxi-03.flashstack.com	FlashStack_DC	FlashStack-VSI	
<input type="checkbox"/>	🟢 vm-infra-esxi-04.flashstack.com	FlashStack_DC	FlashStack-VSI	
<input type="checkbox"/>	🟢 vm-infra-esxi-02.flashstack.com	FlashStack_DC	FlashStack-VSI	
<input type="checkbox"/>	🟢 vm-infra-esxi-01.flashstack.com	FlashStack_DC	FlashStack-VSI	


Claim FlashArray//X in Cisco Intersight

Claiming a Pure Storage FlashArray also requires the use of an Intersight Assist virtual machine. Deploy an Intersight assist appliance using the above described procedure if one doesn't exist. To claim FlashArray//X in Cisco Intersight, follow these steps:

1. Open a browser to Cisco Intersight, <https://intersight.com> and log in to your Intersight account.
2. Select Admin > Devices.




3. Click Claim a New Device and select Claim Though Intersight Assist.
4. Set Type to Pure Storage FlashArray.



5. Click Start.

6. Enter FlashArray Hostname/ IP address and credentials.



7. Click Claim.

<input type="checkbox"/>	Name	Status	Type	Claimed Time
<input type="checkbox"/>	dcnm.flashstack.local	Connected	Cisco DCNM	3 hours ago
<input type="checkbox"/>	iwo-demo.saas.appdynamics...	Connected	AppDynamics	Feb 20, 2021 2:22 PM
<input type="checkbox"/>	10.2.164.45	Connected	Pure Storage FlashArray	Feb 20, 2021 2:19 PM
<input type="checkbox"/>	10.2.164.110	Connected	VMware vCenter	Feb 20, 2021 2:10 PM
<input type="checkbox"/>	10.2.164.100	Connected	Pure Storage FlashArray	Aug 25, 2021 11:13 AM

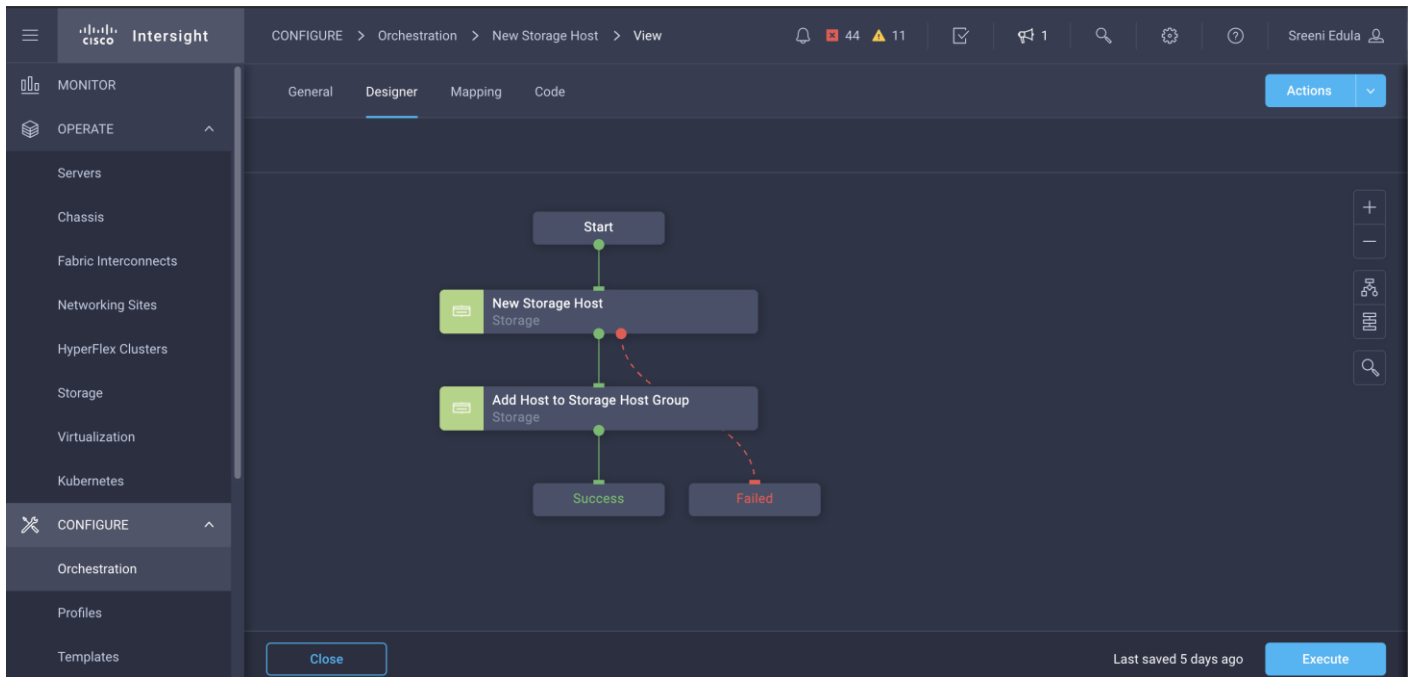
FC Host Registration using Cisco Intersight

To register the FC host using Cisco Intersight, follow these steps:

1. Selection Configure > Orchestration.
2. Select New Storage Host .

<input type="checkbox"/>	Display Name	Descrip...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update	
<input type="checkbox"/>	New Virtual Machine	Create a new...	Yes	1	0		Valid	Sep 2, 2021
<input type="checkbox"/>	New Storage Virtual Machine	Create a stor...	Yes	1	0		Valid	Sep 2, 2021
<input type="checkbox"/>	New Storage Interface	Create a stor...	Yes	1	0		Valid	Sep 2, 2021
<input type="checkbox"/>	New Storage Host Group	Create a new...	Yes	2	0		Valid	Sep 2, 2021
<input checked="" type="checkbox"/>	New Storage Host	Create a new...	Yes	4	0		Valid	Sep 2, 2021
<input type="checkbox"/>	New NAS Datastore	Create a NFS...	Yes	1	0		Valid	Sep 2, 2021
<input type="checkbox"/>	New Storage Export Policy	Create a stor...	Yes	1	0		Valid	Sep 2, 2021
<input type="checkbox"/>	Operating System Install	Workflow to i...	Yes	3	0		Valid	Sep 2, 2021
<input type="checkbox"/>	Deploy Infrastructure Kubernetes Cluster	Deploy a Kub...	Yes	1	0		Valid	Aug 26, 2021...	...

3. Select Execute.



4. Select the appropriate Organization (default by default).
5. Select the appropriate Pure Storage device.
6. Enter the name of the Host name and WWNs for host VM-Host-Infra-FCP-01.

Enter Workflow Input - New Storage Host ✕

Organization *

FlashStack-BB ▼ ⓘ

Workflow Instance Name

New Storage Host ⓘ

Storage Device * ⓘ

Selected Storage Device BB08-FlashArrayR3



Host Group ⓘ

Select Host Group

Host *

VM-Host-Infra-FCP-01 ⓘ

Name
dep

WWNs

20:00:00:25:B5:A4:0A:00 ⓘ



WWNs

20:00:00:25:B5:A4:0B:00 ⓘ



Cancel

Execute

7. Select Execute.

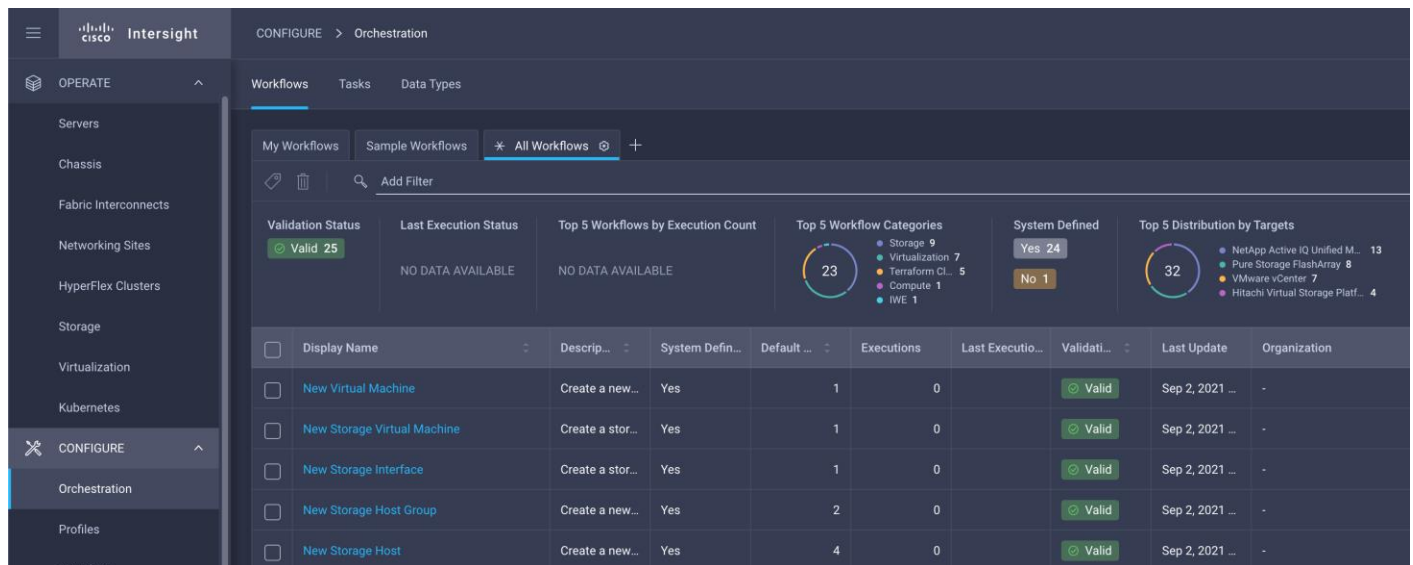
8. Repeat Steps 2-7 for all host.

Create FC Host Group using Cisco Intersight

To create a FC host group using Cisco Intersight, follow these steps:

1. Selection Configure > Orchestration.

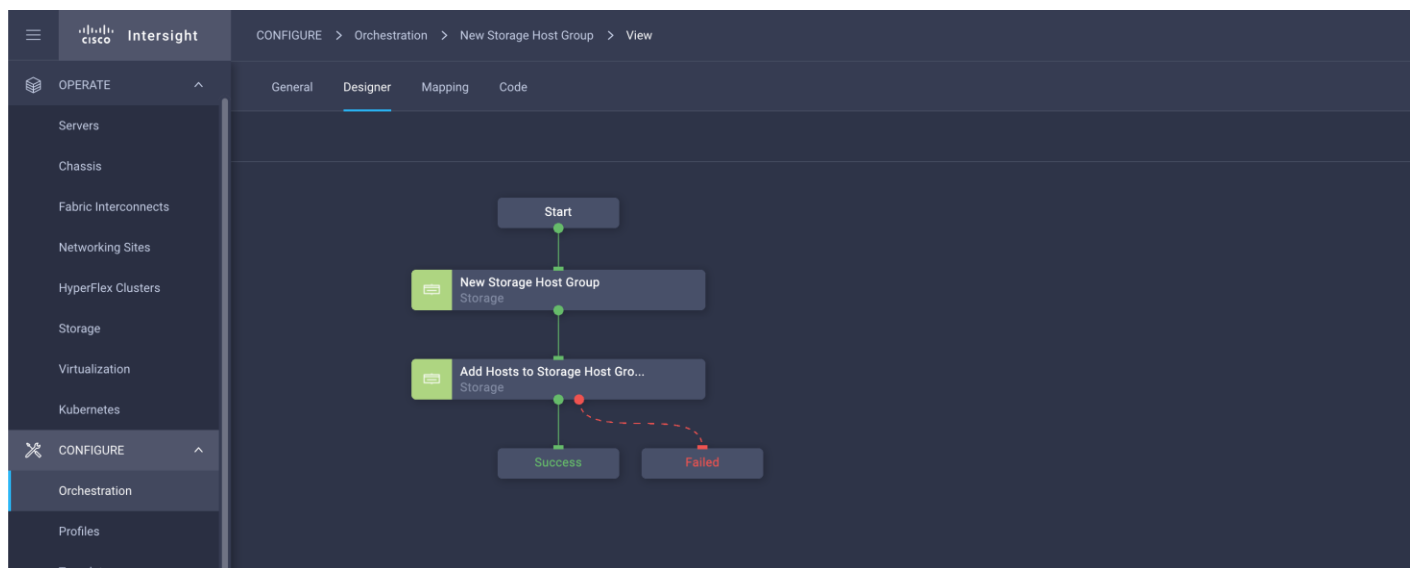
2. Select New Storage Host Group.



The screenshot shows the Cisco Intersight interface in the 'CONFIGURE > Orchestration' section. The 'Workflows' tab is active, displaying a list of workflows under 'All Workflows'. The workflow 'New Storage Host Group' is highlighted, showing it is 'Valid' and has been executed 2 times. The interface also includes summary cards for validation status, execution counts, and distribution by targets.

Display Name	Descr...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update	Organization
<input type="checkbox"/> New Virtual Machine	Create a new...	Yes	1	0		Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Virtual Machine	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Interface	Create a stor...	Yes	1	0		Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Host Group	Create a new...	Yes	2	0		Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Host	Create a new...	Yes	4	0		Valid	Sep 2, 2021 ...	-

3. Select Execute.



The screenshot shows the Cisco Intersight interface in the 'CONFIGURE > Orchestration > New Storage Host Group > View' section. The 'Designer' tab is active, displaying a workflow diagram. The workflow starts with a 'Start' node, followed by a 'New Storage Host Group' node, and then an 'Add Hosts to Storage Host Gro...' node. The workflow ends with 'Success' and 'Failed' nodes.

```
graph TD; Start[Start] --> NewStorageHostGroup[New Storage Host Group]; NewStorageHostGroup --> AddHosts[Add Hosts to Storage Host Gro...]; AddHosts --> Success[Success]; AddHosts --> Failed[Failed];
```

-
4. Select the appropriate Organization (default by default).
 5. Select the appropriate Pure Storage device.
 6. Enter the name of the Host Group and of the Hosts created during Host Registration. VM-Infra-Host-FCP-01, VM-Infra-Host-FC-02 and VM-Infra-Host-FCP-03 are the hosts used in this deployment.

Enter Workflow Input - New Storage Host Group ✕

Organization *
FlashStack-BB ▼ ⓘ

Workflow Instance Name
New Storage Host Group ⓘ

Storage Device * ⓘ
Selected Storage Device BB08-FlashArrayR3 ✎ | ✕

Host Group * ⓘ
VM-Infra-Host-Group ⓘ

Hosts

VM-Infra-Host-FCP-01 ⓘ 🗑️

Hosts

VM-Infra-Host-FCP-02 ⓘ 🗑️

Hosts

VM-Infra-Host-FCP-03 ⓘ 🗑️ +

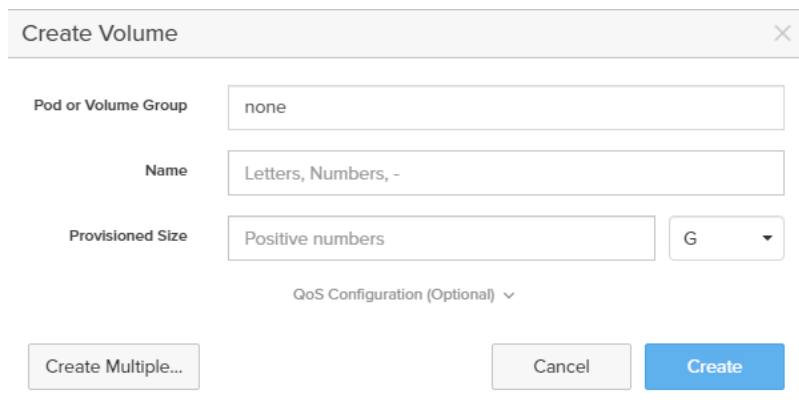
Cancel Execute

7. Select Execute

Private FC Boot Volumes for each ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

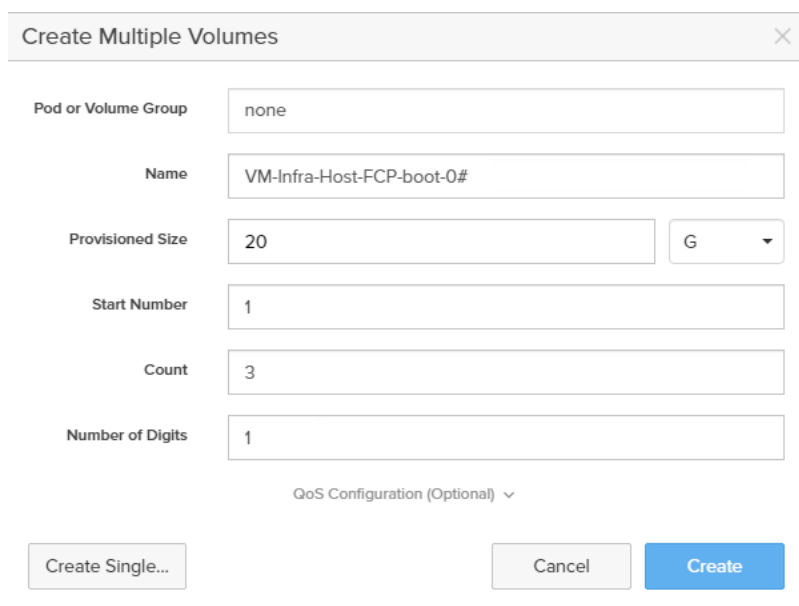
1. Select Storage > Volumes
2. Select the + icon in the Volumes Panel
3. A pop-up will appear to create a volume on the FlashArray.



The 'Create Volume' dialog box contains the following fields and controls:

- Pod or Volume Group:** A text input field containing the value 'none'.
- Name:** A text input field containing the value 'Letters, Numbers, -'.
- Provisioned Size:** A text input field containing 'Positive numbers' and a dropdown menu set to 'G'.
- QoS Configuration (Optional):** A dropdown menu.
- Buttons:** 'Create Multiple...' (disabled), 'Cancel', and 'Create'.


4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Starting Number, Count, and Number of Digits, with a “#” appearing in the name where an iterating number will appear.



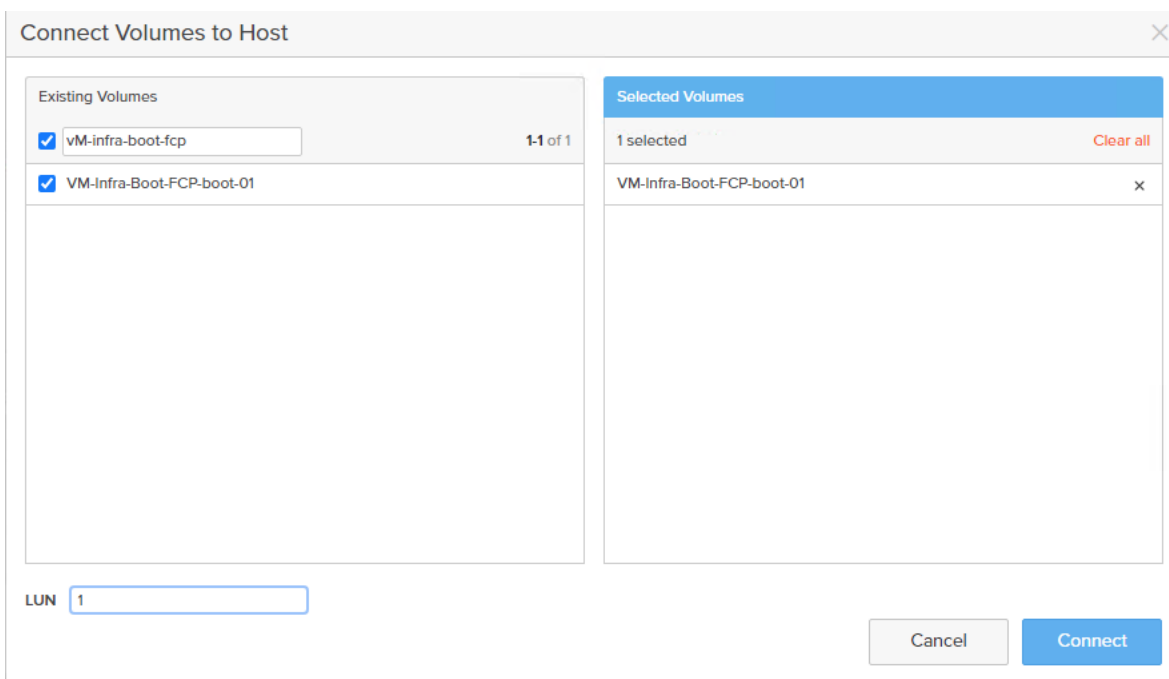
The 'Create Multiple Volumes' dialog box contains the following fields and controls:

- Pod or Volume Group:** A text input field containing the value 'none'.
- Name:** A text input field containing the value 'VM-Infra-Host-FCP-boot-0#'.
- Provisioned Size:** A text input field containing '20' and a dropdown menu set to 'G'.
- Start Number:** A text input field containing the value '1'.
- Count:** A text input field containing the value '3'.
- Number of Digits:** A text input field containing the value '1'.
- QoS Configuration (Optional):** A dropdown menu.
- Buttons:** 'Create Single...' (disabled), 'Cancel', and 'Create'.

5. Click Create to provision the volumes to be used as FC boot LUNs.
6. Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon pull-down within the Connected Volumes tab within that host.



7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.



 LUN ID 1 should be used for the boot

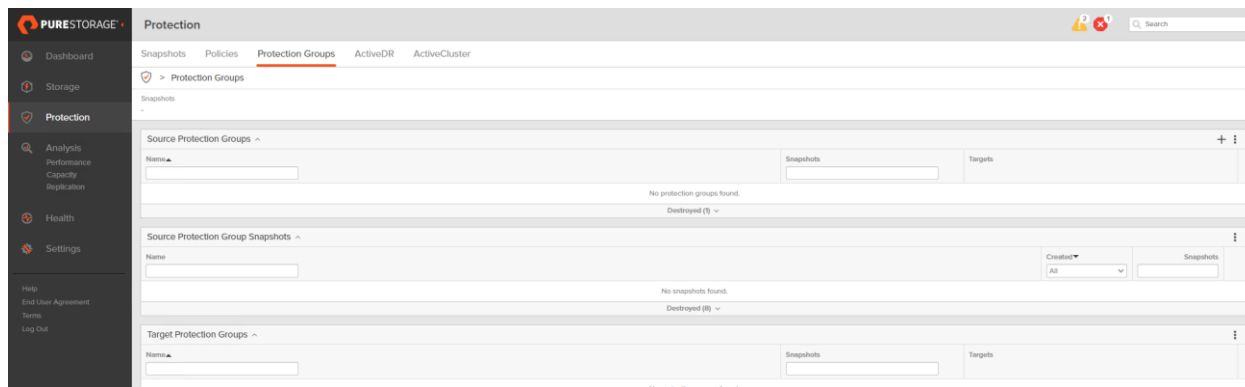
8. Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Configure Storage Policy Based Management

VMware vSphere can communicate to the array via VASA provider to find out what features it supports and allow the vSphere administrator to assign, change, or remove functionality on a VVol on demand and via policies. Below is an example of how to configure a Protection group that will provide hourly snapshots that will be retained for 1 day, with 4 snapshots per day retained for 7 days. These policies should be configured based on application snapshot need.

To configure Storage Policy Based Management, follow these steps:

1. In the Pure Storage Web Portal, select Protection > Protection Groups > select the + icon in the Source Protection Groups.



2. Enter a name.

Create Protection Group

Pod

Name

3. Select the protection group.

4. Edit the Snapshot Schedule based on your operational requirements.

Edit Snapshot Schedule

Enabled

Create a snapshot on source every hours at

Retain all snapshots on source for days

then retain snapshots per day for more days

5. Click Save.

iSCSI Host Registration using Cisco Intersight

To register the iSCSI Host using Cisco Intersight, follow these steps:

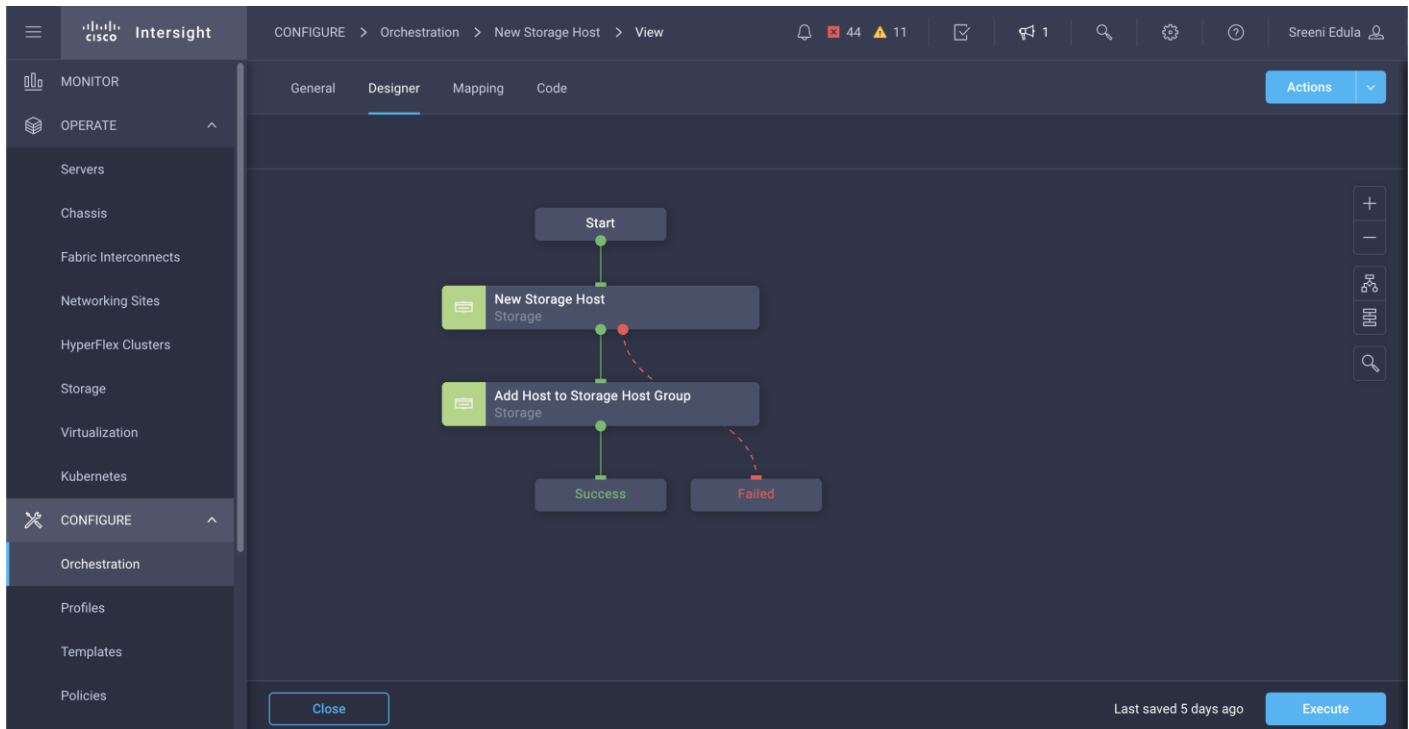
1. Selection Configure > Orchestration.

2. Select New Storage Host.

The screenshot shows the Cisco Intersight interface in the 'CONFIGURE > Orchestration' section. The left sidebar is expanded to the 'CONFIGURE' section, with 'Storage' selected. The main content area displays a table of tasks under the 'Workflows' tab. The table has columns for 'Display Name', 'Description', 'System Defined', 'Default Value', 'Executions', and 'Last Run'. A 'Valid 25' status indicator is visible at the top left of the table area. A circular progress indicator shows 23% completion. A legend on the right indicates 5 Terraform Clusters, 1 Compute, and 1 IWE.

<input type="checkbox"/>	Display Name	Description	System Defined	Default Value	Executions	Last Run
<input type="checkbox"/>	New VMFS Datastore	Create a stor...	Yes	5	0	
<input type="checkbox"/>	New Virtual Machine	Create a new...	Yes	1	0	
<input type="checkbox"/>	New Storage Virtual Machine	Create a stor...	Yes	1	0	
<input type="checkbox"/>	New Storage Interface	Create a stor...	Yes	1	0	
<input type="checkbox"/>	New Storage Host Group	Create a new...	Yes	2	0	
<input type="checkbox"/>	New Storage Host	Create a new...	Yes	4	0	
<input type="checkbox"/>	New NAS Datastore	Create a NFS...	Yes	1	0	

3. Select Execute.



4. Select the appropriate Organization (default by default).
5. Select the appropriate Pure Storage device.
6. Enter the name of the Host name and IQN for host VM-Host-Infra-iSCSI-01.

Enter Workflow Input - New Storage Host ✕

Organization *
FlashStack-BB ▼ ⓘ

Workflow Instance Name
New Storage Host ⓘ

Storage Device * ⓘ
Selected Storage Device BB08-FlashArrayR3 ✎ | ✕

Host Group ⓘ
[Select Host Group](#)

Host *
VM-Host-Infra-iSCSI-01 ⓘ

WWNs ⓘ +

IQNs ⓘ +
iqn.2010-11.com.flashstack:infra-ucs-host:1 ⓘ

Cancel Execute

7. Select Execute.

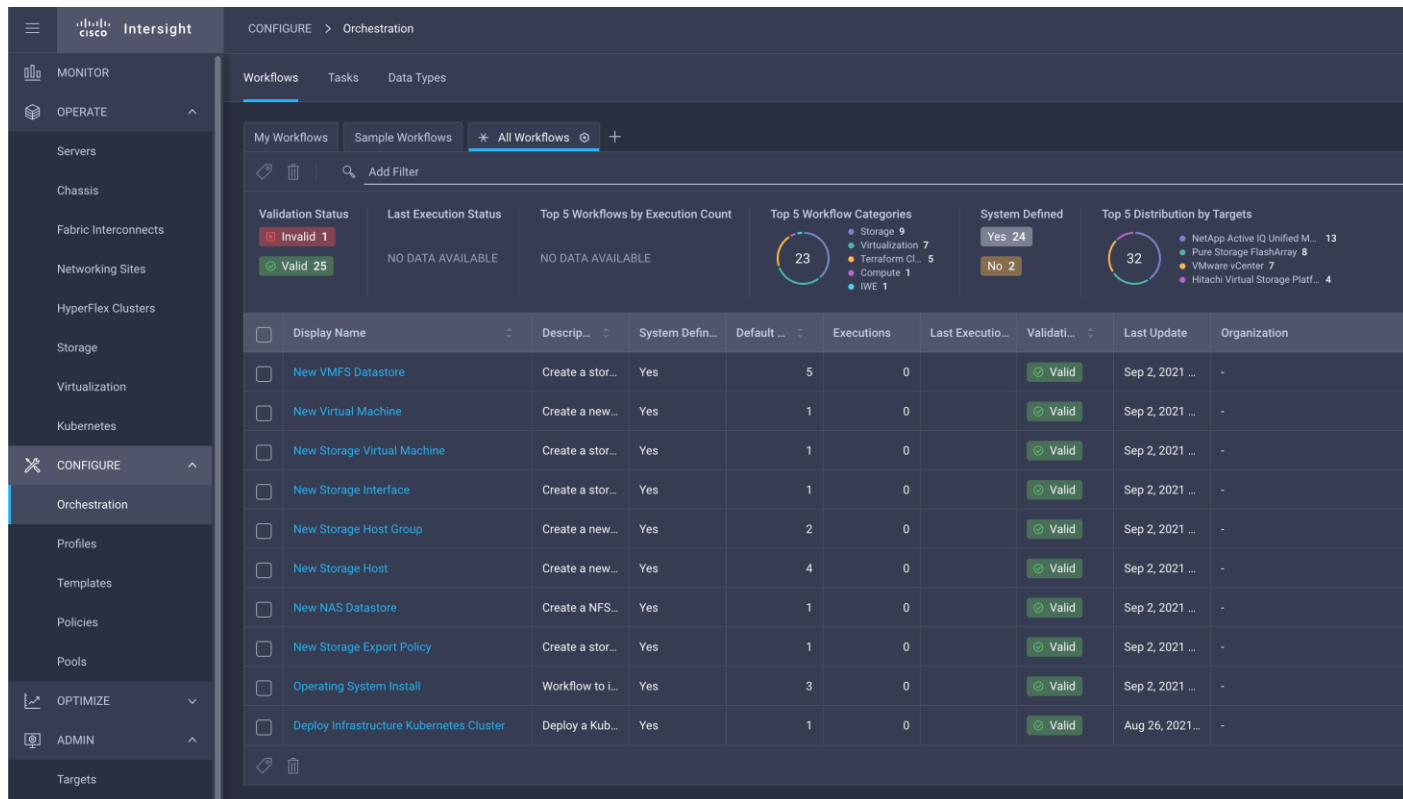
8. Repeat Steps 2-7 for all host.

Create Host Group using Cisco Intersight

To create a Host group using Cisco Intersight, follow these steps:

1. Selection Configure > Orchestration.

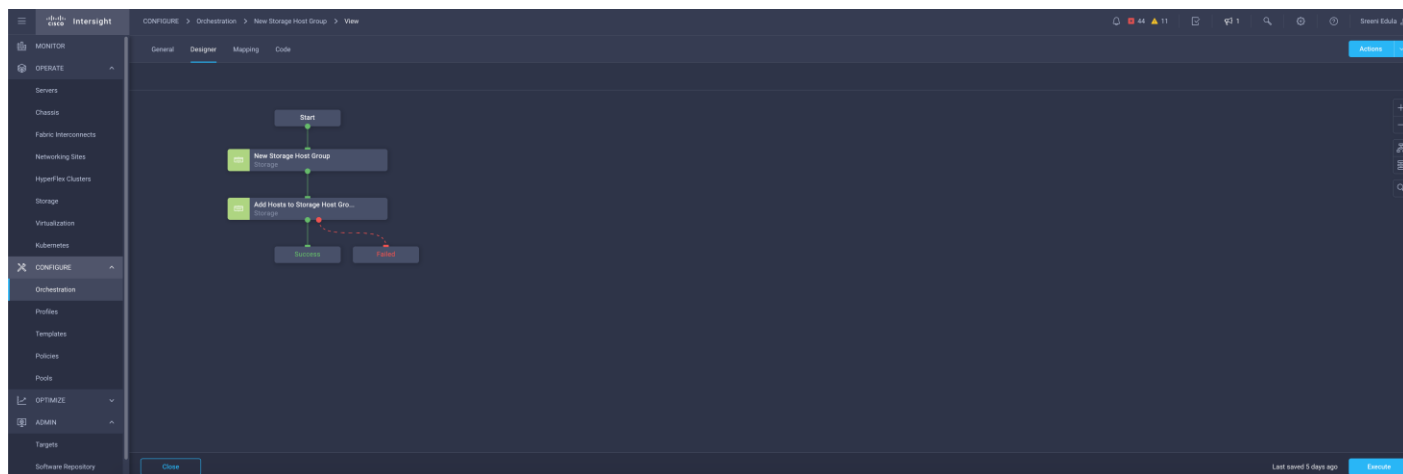
2. Select New Storage Host Group.



The screenshot shows the Cisco Intersight interface for the Orchestration section. The left sidebar is expanded to 'CONFIGURE > Orchestration'. The main content area displays a list of workflows under the 'All Workflows' tab. The workflows are listed in a table with columns for Display Name, Description, System Defined, Default, Executions, Last Execution, Validation Status, Last Update, and Organization. The 'New Storage Host Group' workflow is highlighted in blue.

Display Name	Descr...	System Defin...	Default ...	Executions	Last Executio...	Validati...	Last Update	Organization
<input type="checkbox"/> New VMFS Datastore	Create a stor...	Yes		5	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Virtual Machine	Create a new...	Yes		1	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Virtual Machine	Create a stor...	Yes		1	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Interface	Create a stor...	Yes		1	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Host Group	Create a new...	Yes		2	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Host	Create a new...	Yes		4	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New NAS Datastore	Create a NFS...	Yes		1	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> New Storage Export Policy	Create a stor...	Yes		1	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> Operating System Install	Workflow to i...	Yes		3	0	Valid	Sep 2, 2021 ...	-
<input type="checkbox"/> Deploy Infrastructure Kubernetes Cluster	Deploy a Kub...	Yes		1	0	Valid	Aug 26, 2021...	-

3. Select Execute.



The screenshot shows the Cisco Intersight interface for the Orchestration section, specifically the workflow designer for the 'New Storage Host Group' workflow. The workflow is displayed in a flowchart format with the following steps: Start, New Storage Host Group (Storage), Add Hosts to Storage Host Gro... (Storage), Success, and Failure. The 'New Storage Host Group' and 'Add Hosts to Storage Host Gro...' steps are highlighted in green. The 'Execute' button is visible at the bottom right.

4. Select the appropriate Organization (default by default).

5. Select the appropriate Pure Storage device.

6. Enter the name of the Host Group and of the Hosts created during Host Registration. VM-Infra-Host-iSCSI-01, VM-Infra-Host-iSCSI-02 and VM-Infra-Host-iSCSI-03 are the hosts used in this deployment.

Enter Workflow Input - New Storage Host Group ✕

Organization *
FlashStack-BB ▼ ⓘ

Workflow Instance Name
New Storage Host Group ⓘ

Storage Device * ⓘ
Selected Storage Device BB08-FlashArrayR3 ✎ | ✕

Host Group *
VM-Infra-iSCSI-Host-Group ⓘ

Hosts

VM-Infra-Host-iSCSI-01 ⓘ 🗑️

Hosts

VM-Infra-Host-iSCSI-02 ⓘ 🗑️

Hosts

VM-Infra-Host-iSCSI-03 ⓘ 🗑️ +

Cancel Execute

7. Select Execute.

Private Boot Volumes for each iSCSI ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

Create Volume

Pod or Volume Group: none

Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

QoS Configuration (Optional) v

Create Multiple... Cancel Create

4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Starting Number, Count, and Number of Digits, with a “#” appearing in the name where an iterating number will appear.

Create Multiple Volumes

Pod or Volume Group: none

Name: VM-Infra-Boot-iSCSI-0#

Provisioned Size: 20 G

Start Number: 1


Count: 3

Number of Digits: 1

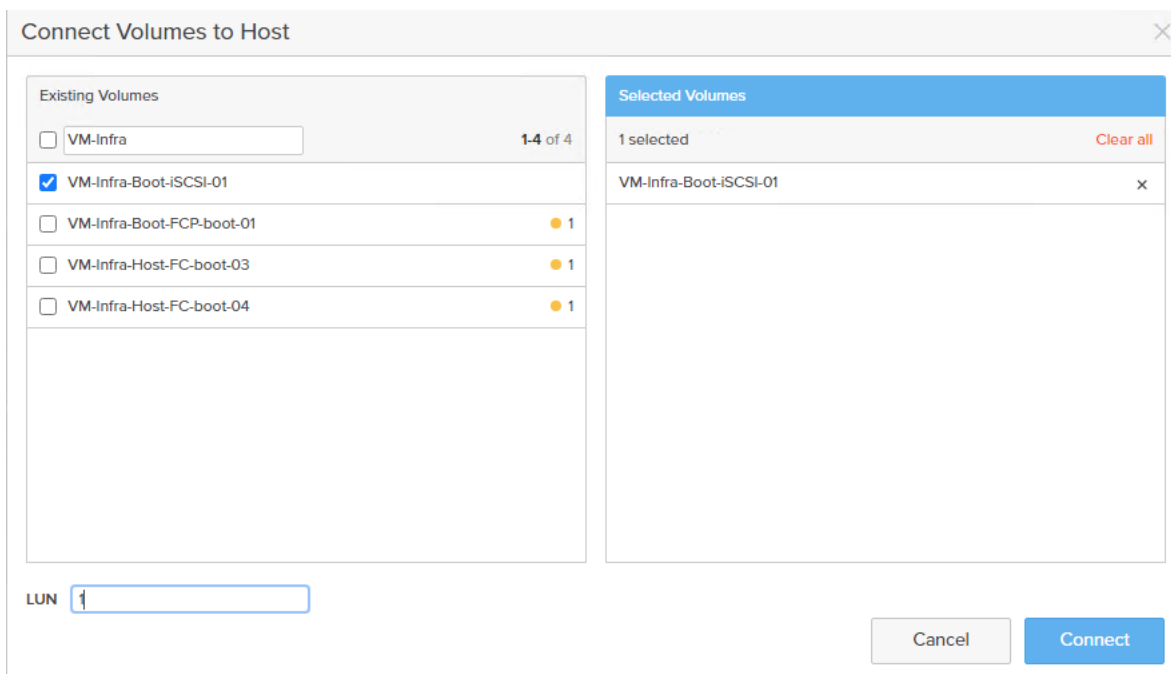
QoS Configuration (Optional) v

Create Single... Cancel Create

5. Click Create to provision the volumes to be used as iSCSI boot LUNs.
6. Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon pull-down within the Connected Volumes tab within that host.



7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.



LUN ID 1 should be used for the boot.

8. Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Cisco Infrastructure Firmware Upgrade (Fabric Interconnects) using Cisco Intersight

To upgrade Cisco UCS Fabric Interconnects using Cisco Intersight, follow these steps in Intersight SaaS Portal:

1. From the left navigation pane, click Fabric Interconnects, select a Fabric Interconnect, and perform an Upgrade Firmware action on it.

The screenshot shows the UCS Manager interface with a table of servers. The table has columns for Name, Health, Contract Status, M..., M..., E..., UCS D..., Total, Ports Used, and Ports Available. Two servers are selected for a firmware upgrade:

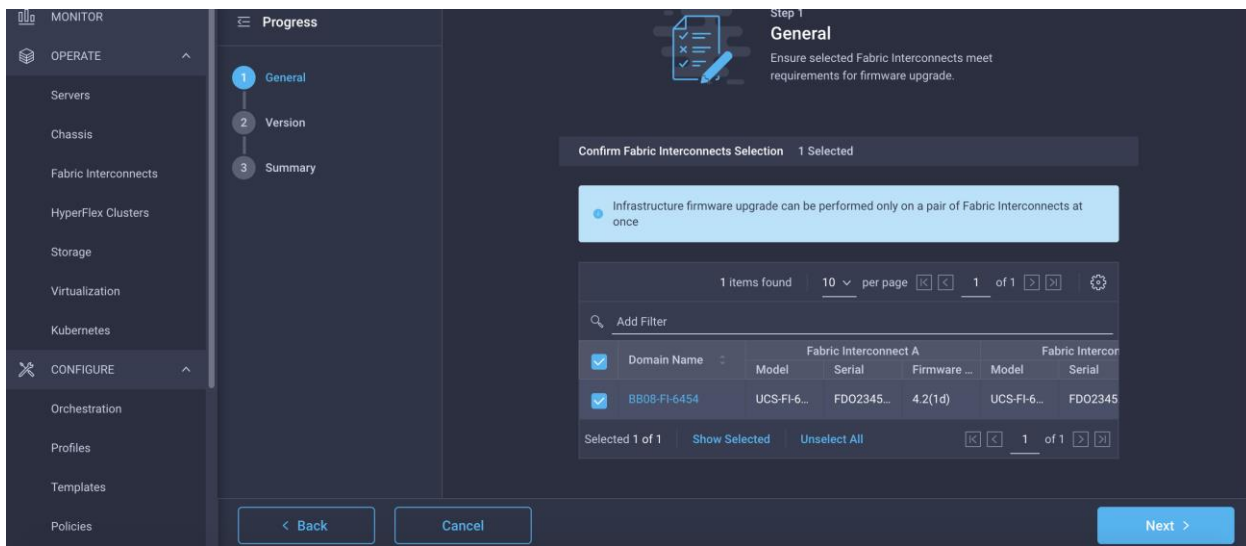
Name	Health	Contract Status	M...	M...	E...	UCS D...	Total	Ports Used	Ports Available
AA12-DP-UCS6454 FI-A	Critical	Not Covered	192.1...	UCS-FI...	0	54	16	38	
AA12-DP-UCS6454 FI-B	Critical	Not Covered	192.1...	UCS-FI...	0	54	16	38	
AA12-FS-Prod-UCS6454 FI-B	Critical	Not Covered	10.2.1...	UCS-FI...	0	54	16	38	
AA12-FS-Prod-UCS6454 FI-A	Critical	Not Covered	10.2.1...	UCS-FI...	0	54	16	38	
BB08-FI-6454 FI-A	Critical	Not Covered	10.1.1...	UCS-FI...	0	54	16	38	
BB08-FI-6454 FI-B	Critical	Not Covered	10.1.1...	UCS-FI...	0	54	16	38	

At the bottom of the table, it says "Selected 2 of 6".

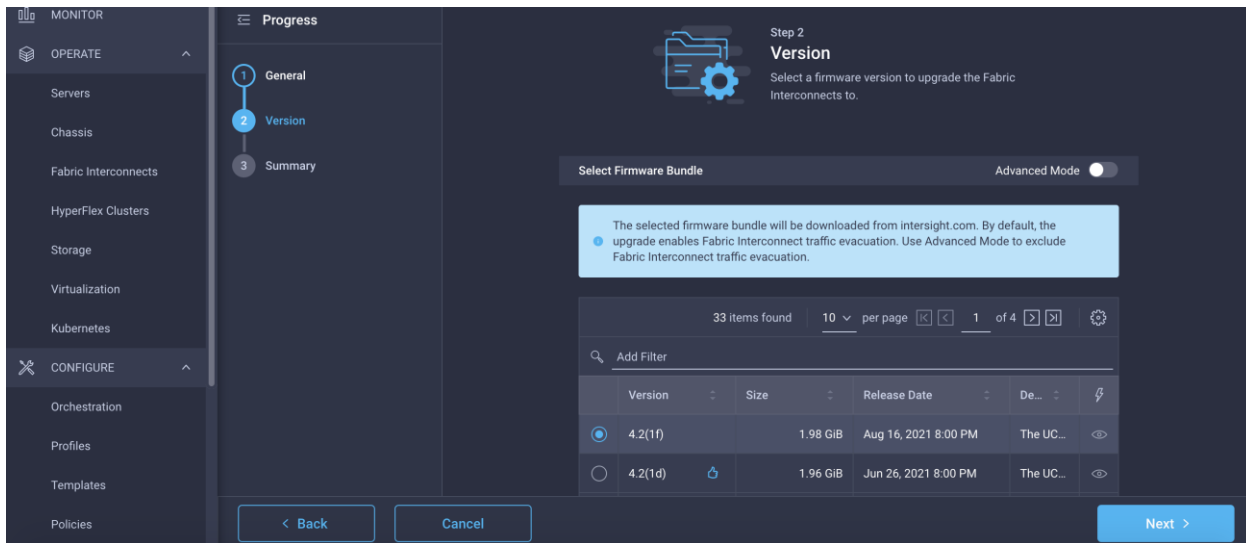
2. On the Upgrade Firmware page, click Start.

The screenshot shows the "Summary & Firmware Upgrade" dialog box. The title is "Summary & Firmware Upgrade" and the subtitle is "Confirm configuration and initiate the upgrade." There is a "Start" button at the bottom right and a "Cancel" button at the bottom left. There is also a link "About Firmware Upgrade" and a checkbox "Do not show this page again".

3. On the General page, confirm selection of the switch Domain and click Next.



- On the Version page, select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded, and click Next.



- On the Summary screen, verify the summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded, and click Upgrade.

MONITOR

OPERATE

Servers

Chassis

Fabric Interconnects

HyperFlex Clusters

Storage

Virtualization

Kubernetes

CONFIGURE

Orchestration

Profiles

Templates

Policies

Progress

1 General

2 Version

3 Summary

Selected firmware bundle will be downloaded to the Fabric Interconnects and upgraded. Click on Requests to monitor the progress of the firmware upgrade.

Firmware

Version 4.2(1f) Size 1.98 GiB

Fabric Interconnects to be Upgraded

1 Items found 10 per page 1 of 1

Add Filter

Domain Name	Fabric Interconnect A			Fabric Interconnect B		
	Model	Serial	Firmware...	Model	Serial	Firmw...
BB08-FI-6454	UCS-FI-6...	FDO234...	4.2(...)	UCS-FI-6...	FDO234...	4.2(...)

Back Cancel Upgrade

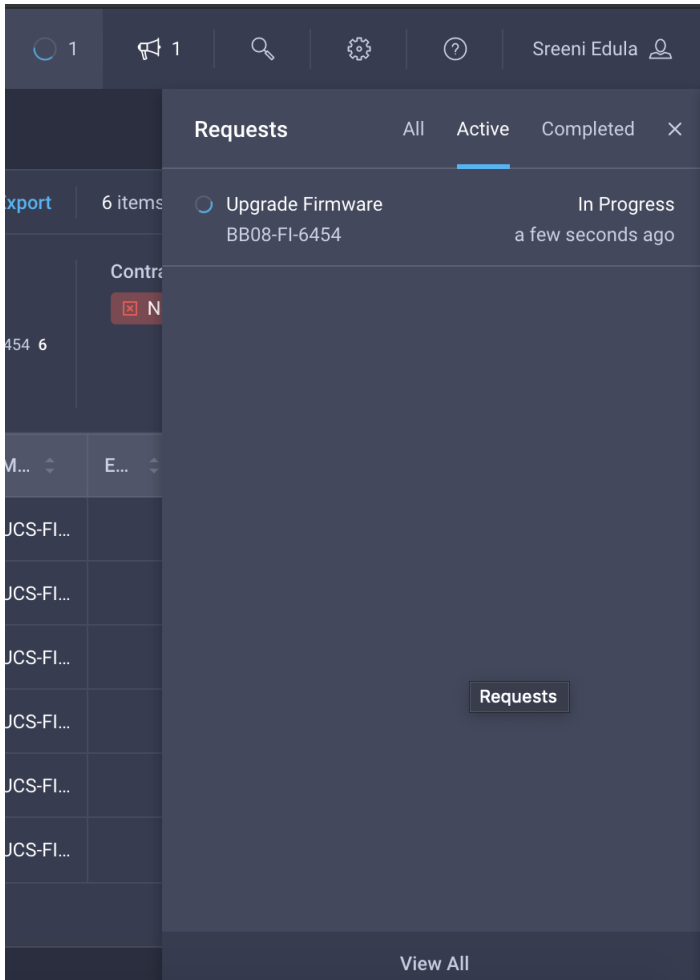
6. Confirm the upgrade request.

Upgrade Firmware

Firmware will be installed on the selected Fabric Interconnects. Are you sure you want to upgrade firmware?

Cancel Upgrade

The firmware upgrade workflow begins.



7. You can check the status of the upgrade workflow in the Execution Flow pane. Acknowledge any messages in the Execution Flow pane and click Continue to proceed with the upgrade.

Requests > Upgrade Firmware

45 12 1 1

Sreeni Edula

Details	Execution Flow
<p>Status: Action Required</p> <p>Name: Upgrade Firmware</p> <p>ID: 6124ff69696f6e2d32ece6ce</p> <p>Target Type: Fabric Interconnect</p> <p>Target Name: BB08-FI-6454 FI-A, BB08-FI-6454 FI-B</p> <p>Source Type: Firmware Upgrade</p> <p>Source Name: BB08-FI-6454</p> <p>Initiator: sredula@cisco.com</p> <p>Start Time: Aug 24, 2021 10:17 AM</p> <p>End Time: -</p> <p>Duration: 1 h 44 m 30 s</p> <p>Organizations: -</p>	<p>Progress: 75%</p> <p>Wait for user acknowledgement on primary Fabric Interconnect.</p> <p>Firmware upgrade for subordinate Fabric Interconnect is complete. Ensure Fabric Interconnects meets requirements to continue upgrade. Please acknowledge to continue with primary Fabric Interconnect upgrade. Learn more at Help Center.</p> <p>Continue</p> <p>Wait for peer Fabric Interconnect activation to complete. Aug 24, 2021 11:17 AM Waiting for User acknowledgement</p> <p>Activate peer Fabric Interconnect. Aug 24, 2021 10:39 AM</p> <p>Wait for image download to complete in UCS Manager. Aug 24, 2021 10:39 AM ucs-6400-k9-bundle-infra.4.2.1f.A.bin downloaded.</p> <p>Initiate image download to UCS Manager. Aug 24, 2021 10:36 AM</p>

8. Click Continue.

Acknowledge Primary Fabric Interconnect Upgrade

Firmware upgrade on subordinate Fabric Interconnect is complete. Please make sure the Fabric Interconnect is in the required state to proceed further with the firmware upgrade process.

Are you sure you want to upgrade?

Cancel Continue

9. Verify if the upgrade is successful.

Requests > Upgrade Firmware		Execution Flow	
Status	Success	Wait for infra upgrade to complete.	
Name	Upgrade Firmware	Wait for user acknowledgement.	
ID	6124ff69696f6e2d32ece6ce	Wait for user acknowledgement on primary Fabric Interconnect.	
Target Type	Fabric Interconnect	Wait for peer Fabric Interconnect activation to complete.	Waiting for User acknowledgement
Target Name	BB08-FI-6454 FI-A BB08-FI-6454 FI-B	Activate peer Fabric Interconnect.	
Source Type	Firmware Upgrade	Wait for image download to complete in UCS Manager.	ucs-6400-k9-bundle-infra.4.2.1f.A.bin downloaded.
Source Name	BB08-FI-6454	Initiate image download to UCS Manager.	
Initiator	sredula@cisco.com	Wait for image download to complete in Fabric Interconnect.	
Start Time	Aug 24, 2021 10:17 AM		
End Time	Aug 24, 2021 12:28 PM		
Duration	2 h 11 m 22 s		
Organizations	-		

Cisco UCS Server Upgrades

To upgrade the Cisco UCS Servers using Intersight, follow these steps in Intersight SaaS Portal:



Only servers in associated state can be upgraded.



Servers associated with server profiles bound to updating templates cannot be upgraded.



Servers associated with global server profiles cannot be upgraded.

1. From the left navigation pane, click Servers, select a server, and perform an Upgrade Firmware action on it.

OPERATE > Servers

Health: 19 (Critical 3, Warning 5, Healthy 11)

Power: On 13, Off 6

HCL Status: Incomplete 15, Not Listed 4

Models: 19 (B200 M6 10, B200 M5 6, UCS S3260-M5... 1, Other 2)

Contract Status: Not Covered 19

Profile Status: 2 (Not Assign...)

Name	Health	Contract Status	M...	M...	⌚	M...	UCS D...	⋮
AA12-DP-UCS6454-1-1	Healthy	Not Covered	192.1...	UCS-S...	92.8	384.0	AA12-...	⋮
BB08-FI-6454-1-6	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-8	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-1	Healthy	Not Covered	10.1.1...	UCSB-...	73.6	384.0	BB08-...	⋮
BB08-FI-6454-1-7	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-4	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-5	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-3	Healthy	Not Covered	10.1.1...	UCSB-...	128.0	512.0	BB08-...	⋮
BB08-FI-6454-1-2	Healthy	Not Covered	10.1.1...	UCSB-...	84.0	384.0	BB08-...	⋮

Context Menu:

- Upgrade Firmware
- Launch vKVM
- Launch UCS Manager
- Open TAC Case
- Set License Tier

2. On the Upgrade Firmware page, click Start.

Servers > Upgrade Firmware

Upgrade Firmware

Version

Select a firmware version to upgrade the servers to.

⏪ ⏩

⬢ ⬤

[About Firmware Upgrade](#)

Do not show this page again

Cancel Start >

3. On the General page, confirm selection of the server and click Next.

OPERATE > Servers

45 11

Sreeni Edula

Health: 19 (Critical 3, Warning 5, Healthy 11)

Power: On 13, Off 6

HCL Status: Not Listed 4, Incomplete 15

Models: 19 (B200 M6 10, B200 M5 6, UCS S3260-M5... 1, Other 2)

Contract Status: Not Covered 19

Profile Status: 2 (Not Assig...)

Name	Health	Contract Status	M...	M...	M...	M...	UCS D...
AA12-FS-Prod-UCS6454-1-8	Warning	Not Covered	10.2.1...	UCSB...	128.0	128.0	AA12-...
BB08-FI-6454-1-6	Healthy	Not Covered	10.1.1...	UCSB...	128.0	512.0	BB08-...
BB08-FI-6454-1-8	Healthy	Not Covered	10.1.1...	UCSB...	128.0	5120.0	BB08-...

Servers > Upgrade Firmware

45 11

Sreeni Edula

Progress

- General
- Version
- Summary

Step 1
General
Ensure selected servers meet requirements for firmware upgrade.

Confirm Servers Selection 1 Selected

1 items found 10 per page 1 of 1

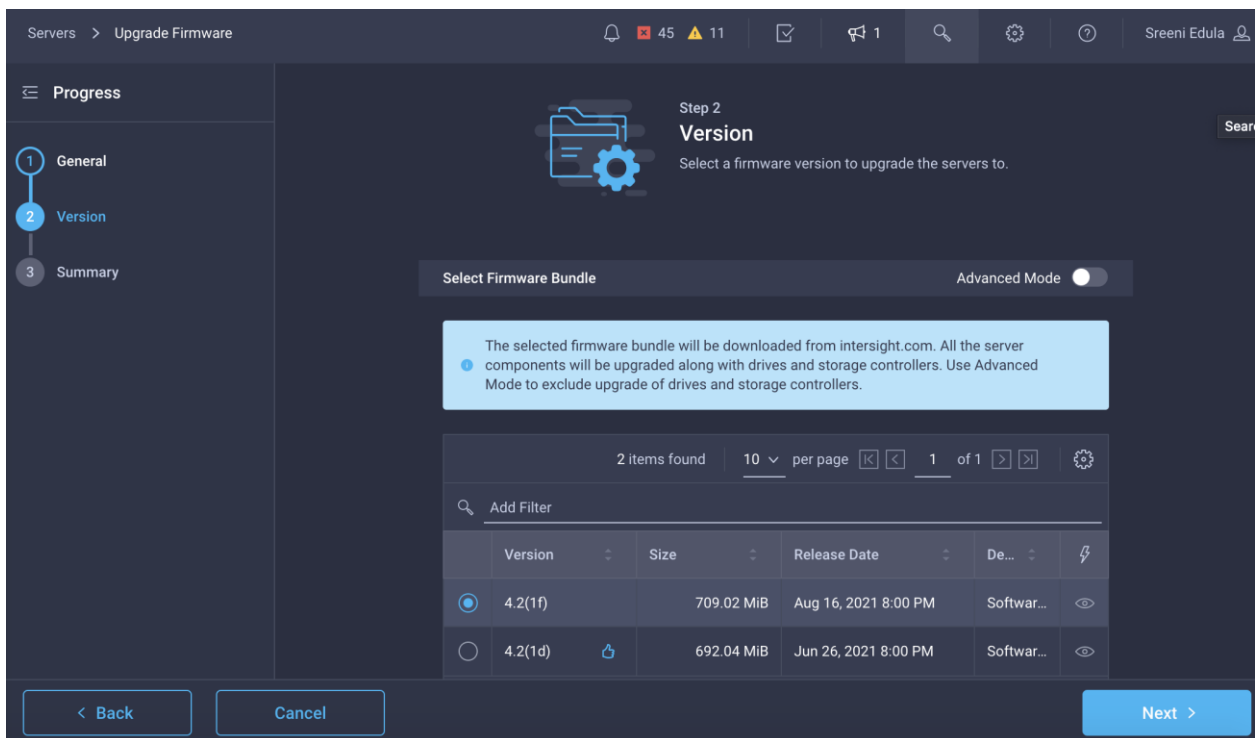
Add Filter

Name	User La...	Model	Firmwar...	UCS Domain
BB08-FI-6454...		UCSB-B200...	4.2(1a)	BB08-FI-6454

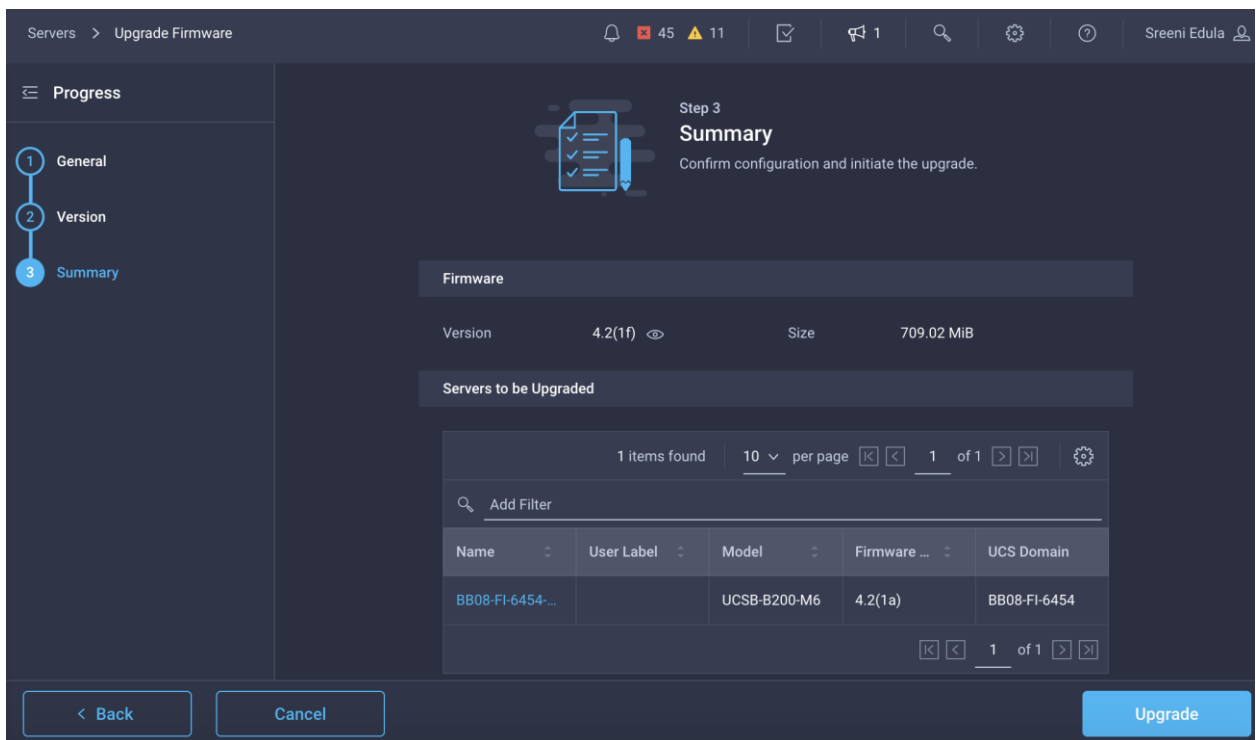
Selected 1 of 1 Show Selected Unselect All 1 of 1

< Back Cancel Next >

- On the Version page, select the fabric firmware bundle to which the Fabric Interconnects need to be upgraded, and click Next.



- On the Summary screen, verify the summary of the selected switches, the firmware version running on them, and the firmware version to which they will be upgraded, and click Upgrade.



- Confirm the upgrade request and monitor the process for successful upgrade.

Upgrade Firmware

Validate if server reboot is required

Skip validation



You will be able to see the impact and cancel or continue the upgrade once the image is uploaded to Fabric Interconnect.

Cancel

Upgrade


Pure Storage vSphere Client Plugin

The Pure Storage Plugin for the vSphere Client provides the ability to VMware users to have insight into and control of their Pure Storage FlashArray environment while directly logged into the vSphere Client. The Pure Storage plugin extends the vSphere Client interface to include environmental statistics and objects that underpin the VMware objects in use and to provision new resources as needed.

The Pure Storage vSphere Client Plugin will be accessible through the vSphere Client after registration through the Pure Storage Web Portal.

To access the Pure Storage vSphere Client Plugin, follow these steps:

1. Go to Settings > Software.
2. Select the edit icon in the vSphere Plugin panel.



3. Enter the vCenter information in the pop-up window and click Save.

Edit vSphere Plugin Configuration

vCenter Host

Administrator User

Administrator Password

4. After the discovery completes, click Install.

vSphere Plugin

vCenter Host vcenter1.flashstack.com

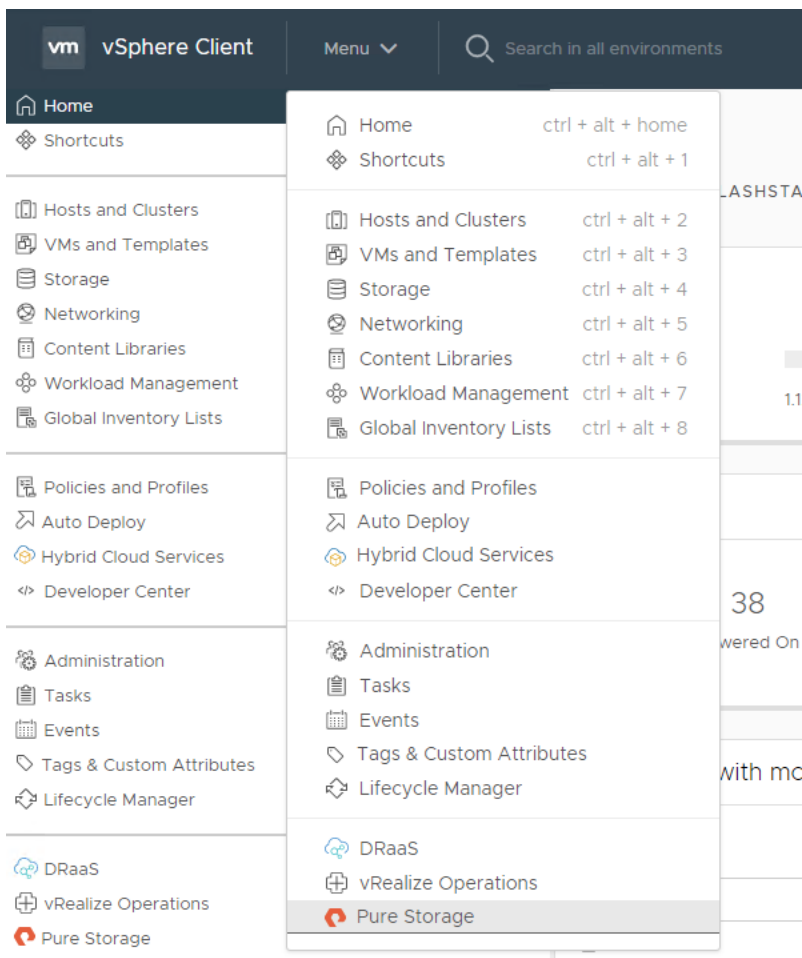
Administrator User administrator@vsphere.local

Administrator Password ****

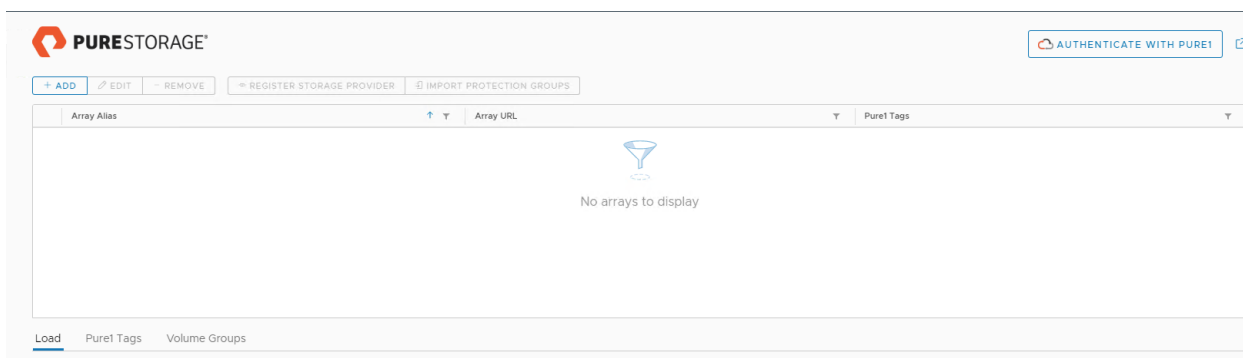
Version on vCenter -

Available Version 4.5.0

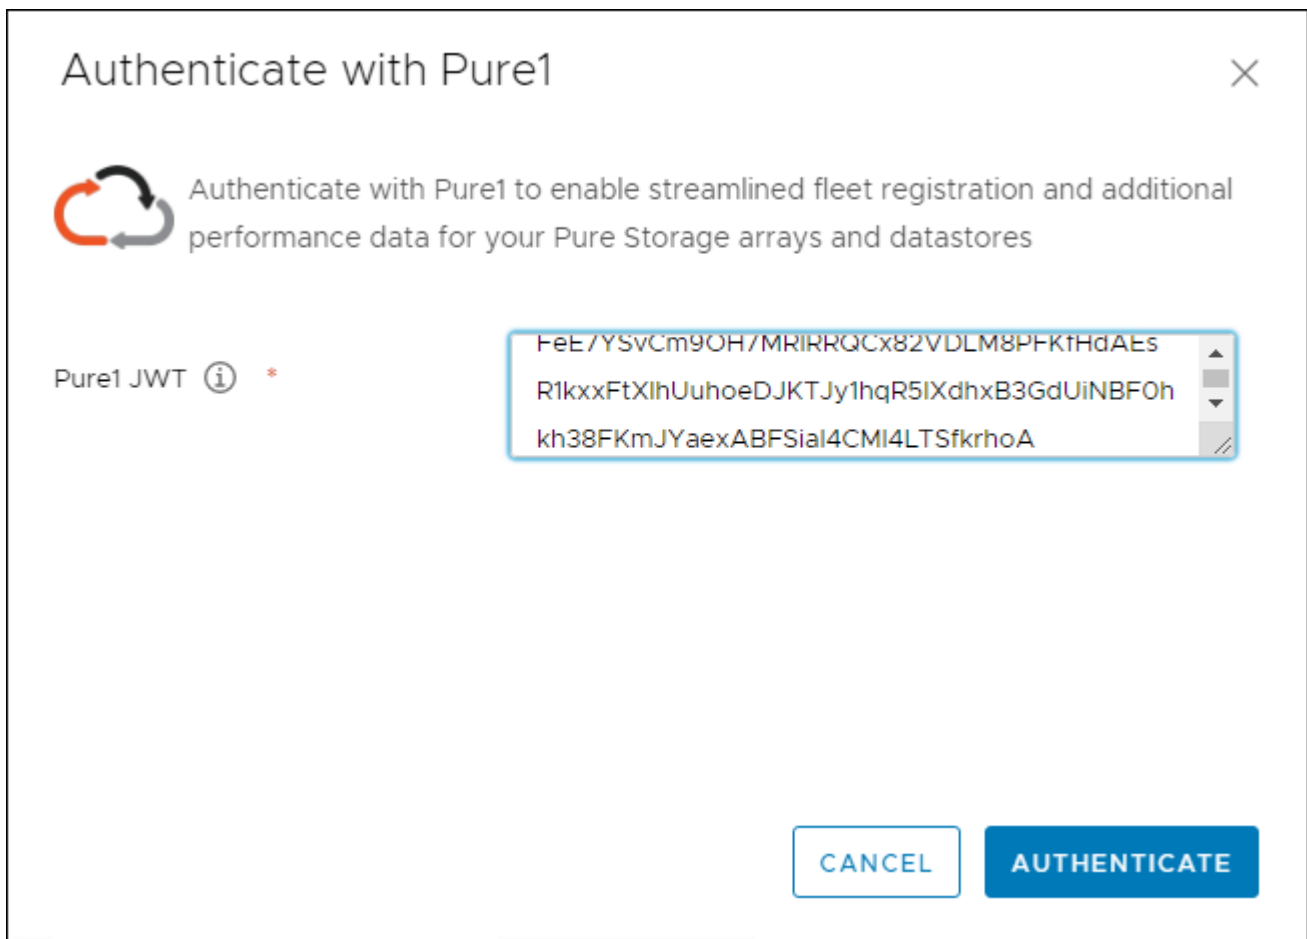
5. In vCenter, select Pure Storage from the Menu.



6. Select Authenticate with Pure1.



7. Input your Pure1 JWT (link).



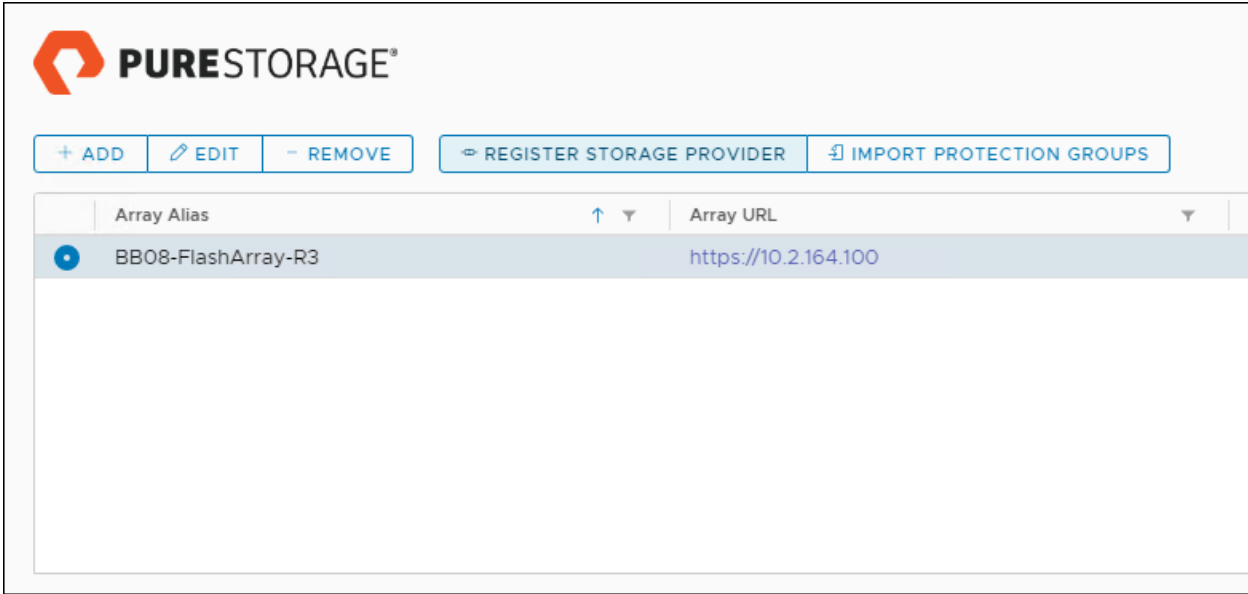
8. Select Authenticate.

9. Select Add.

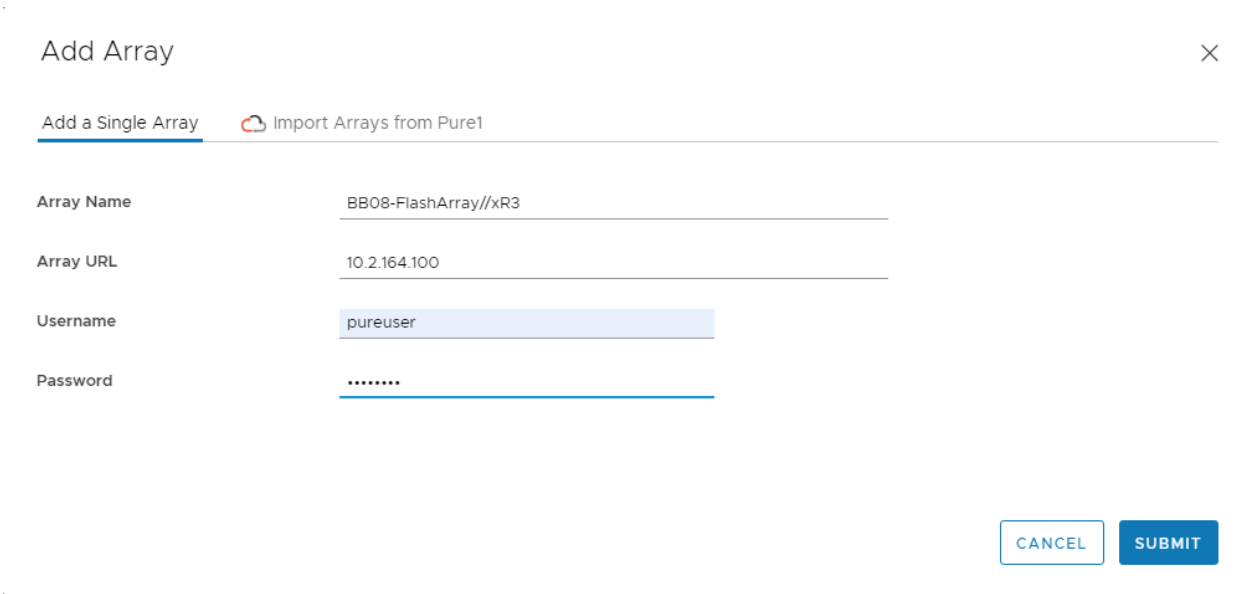
10. Select Import Arrays from Pure1 and input the Username and Password.

11. Select Import Arrays from Pure1 and input the Username and Password.

12. Select Done.



13. Alternatively, provide array details in the Add a Single Array tab to add the Array manually.



14. Select the newly added array.

15. Select Register Storage Provider.

+ ADD

EDIT

- REMOVE

REGISTER STORAGE PROVIDER

IMPORT PROTECTION GROUPS

Array Alias

↑ ↓

Array URL



BB08-FlashArray//xR3

https://10.2.164.100

16. Enter Username and Password.

Register Storage Provider ×

i Registering the storage provider requires a valid username and password.

Username *


Password *

17. Select Register.

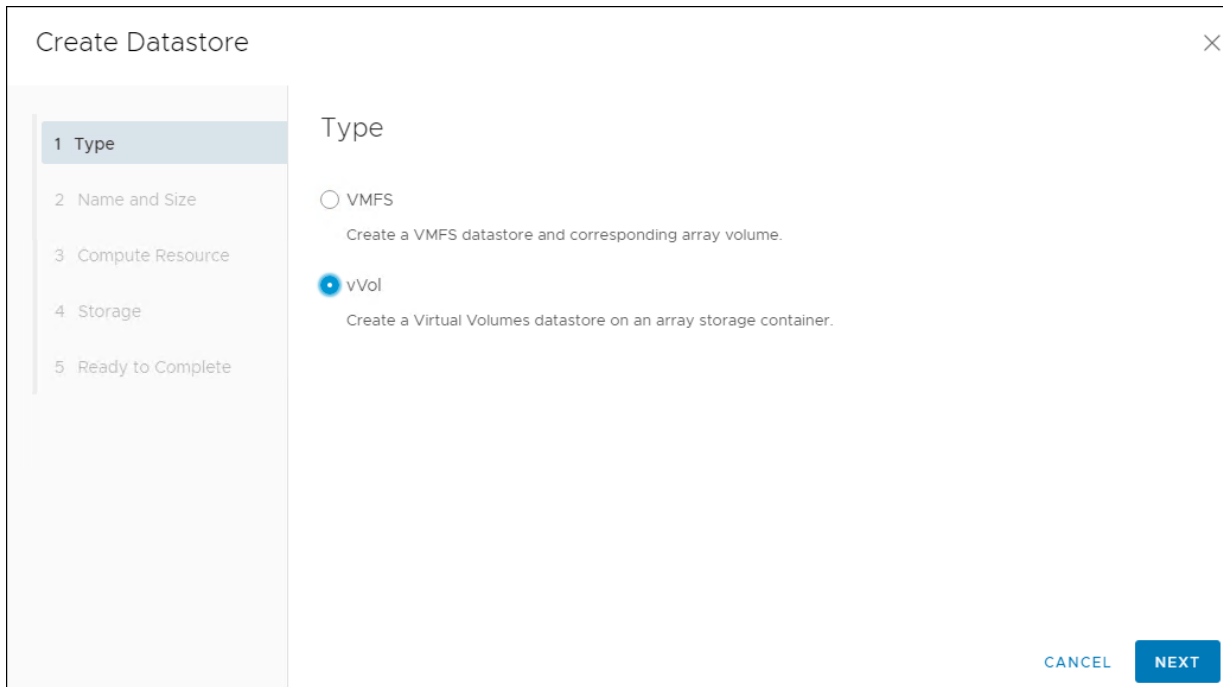
Create VMDS Datastore using Pure vSphere Plugin

To create VMDS datastore using the Pure vSphere plugin, follow these steps:

1. In vCenter, Select Host and Clusters.
2. Right-click the FlashStack Cluster and Select Pure Storage > Create Datastore.



3. Select VMFS.



4. Click Next.
5. Keep VMFS 6 selected.
6. Click Next.
7. Enter a Datastore Name and Datastore Size.

Create Datastore



1 Type

2 VMFS Version

3 Name and Size

4 Compute Resource

5 Storage

6 Protection Groups

7 Volume Group & QoS

8 Ready to Complete

Name and Size

Datastore Name:

Datastore Size: **GB**

CANCEL

BACK

NEXT

8. Select Next.

9. Select the Cluster under Compute Resources.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource**
- 4 Storage
- 5 Ready to Complete

Compute Resource

Compute Resource
<input checked="" type="radio"/> FlashStack-FC
<input type="radio"/> vm-host-infra-fc-01.flashstack.com
<input type="radio"/> vm-host-infra-fc-02.flashstack.com
<input type="radio"/> vm-host-infra-fc-03.flashstack.com

1 - 4 of 4 clusters/hosts

CANCEL BACK NEXT

10. Click Next.

11. Select the Registered FlashArray.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource
- 4 Storage**
- 5 Ready to Complete

Storage

Array
<input checked="" type="radio"/> BB08-FlashArray//xR3

1 - 1 of 1

CANCEL BACK NEXT

12. Optionally, add to the protection group created earlier and click Next.

Create Datastore ×

- 1 Type
- 2 VMFS Version
- 3 Name and Size
- 4 Compute Resource
- 5 Storage
- 6 Protection Groups**
- 7 Volume Group & QoS
- 8 Ready to Complete

Protection Groups

Add to Protection Group(s) ↑ ▼

Platinum (local snapshot every 1 hour, no remote replication)

1 - 1 of 1 protection groups

CANCEL BACK NEXT

13. Click Next on the Volume Group & QoS page.

- 1 Type
- 2 VMFS Version
- 3 Name and Size
- 4 Compute Resource
- 5 Storage
- 6 Protection Groups
- 7 Volume Group & QoS**
- 8 Ready to Complete

Volume Group & QoS

Bandwidth Limit (optional) _____ MB/s ▾

IOPS Limit (optional) _____ K ▾

Volume Group	Bandwidth Limit	IOPS Limit
None	-	-

1 - 1 of 1

CANCEL BACK NEXT

14. Review the information and select Finish.

Create Datastore



- 1 Type
- 2 VMFS Version
- 3 Name and Size
- 4 Compute Resource
- 5 Storage
- 6 Protection Groups
- 7 Volume Group & QoS
- 8 Ready to Complete**

Ready to Complete

Datastore Name:	FS-DS
Type:	VMFS
VMFS Version:	VMFS 6
Datastore Size:	200 GB
Compute Resource:	FlashStack-FC
Array:	BB08-FlashArray//XR3
Pod:	None
Volume Bandwidth Limit:	-
Volume IOPS Limit:	-
Volume Group:	None
Protection Groups:	None

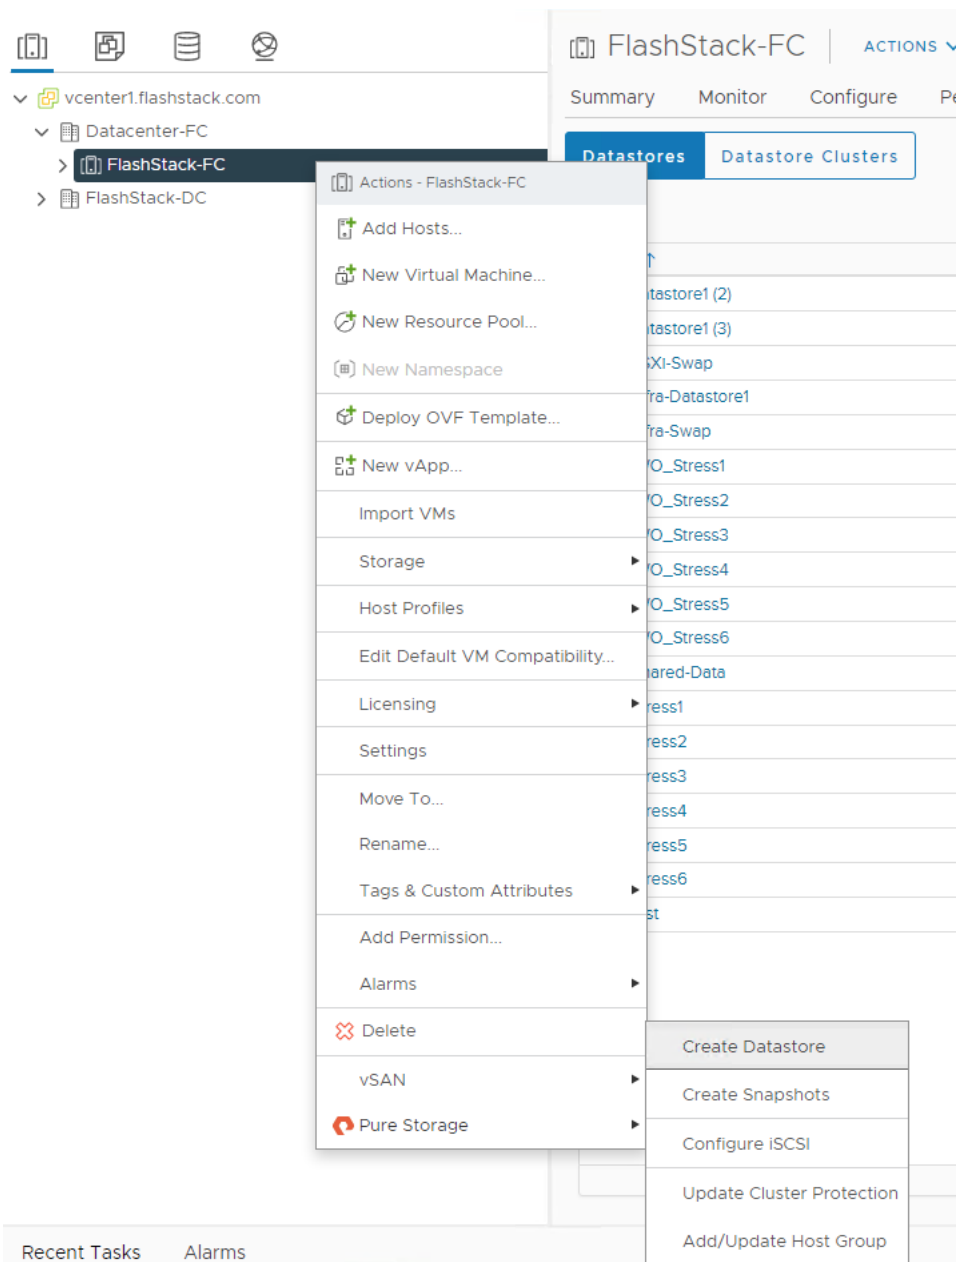
CANCEL

BACK

FINISH

Create vVol Datastore

1. In vCenter, Select Host and Clusters.
2. Right-click the FlashStack Cluster and Select Pure Storage > Create Datastore.



3. Select vVol.

The screenshot shows a 'Create Datastore' dialog box with a sidebar on the left containing five steps: 1 Type, 2 Name and Size, 3 Compute Resource, 4 Storage, and 5 Ready to Complete. The main area is titled 'Type' and contains two radio button options: 'VMFS' (unselected) and 'vVol' (selected). Below 'VMFS' is the text 'Create a VMFS datastore and corresponding array volume.' Below 'vVol' is the text 'Create a Virtual Volumes datastore on an array storage container.' At the bottom right, there are 'CANCEL' and 'NEXT' buttons.

4. Select Next.

5. Enter a Datastore Name.

The screenshot shows the 'Create Datastore' dialog box at the 'Name and Size' step. The sidebar on the left has '2 Name and Size' highlighted. The main area is titled 'Name and Size' and contains a 'Datastore Name:' label with a red asterisk, followed by a text input field containing 'FlashStack-VSI-vVol'. Below the input field is the text 'FlashArray Virtual Volume Datastores are automatically created using the maximum size.' At the bottom right, there are 'CANCEL', 'BACK', and 'NEXT' buttons.

6. Select Next.

7. Select the Cluster under Compute Resources.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource**
- 4 Storage
- 5 Ready to Complete

Compute Resource

Compute Resource		▼
<input checked="" type="radio"/>	FlashStack-FC	
<input type="radio"/>	vm-host-infra-fc-01.flashstack.com	
<input type="radio"/>	vm-host-infra-fc-02.flashstack.com	
<input type="radio"/>	vm-host-infra-fc-03.flashstack.com	

1 - 4 of 4 clusters/hosts

CANCEL BACK NEXT

8. Select Next.

9. Select the Registered FlashArray.

Create Datastore ×

- 1 Type
- 2 Name and Size
- 3 Compute Resource
- 4 Storage**
- 5 Ready to Complete

Storage

Array		▼
<input checked="" type="radio"/>	BB08-FlashArray//xR3	

1 - 1 of 1

CANCEL BACK NEXT

10. Select Next.

Create Datastore ✕

- 1 Type
- 2 Name and Size
- 3 Compute Resource
- 4 Storage
- 5 Ready to Complete**

Ready to Complete

Datstore Name:	FlashStack-VSI-vVOL
Type:	vVol
Compute Resource:	FlashStack-FC
Array:	BB08-FlashArray//XR3
Pod:	None
Storage Provider:	✔ 2 / 2
Storage Container:	✔ vvol container
Protocol Endpoint Verified:	✔ Yes

[CANCEL](#) [BACK](#) [FINISH](#)

11. Review the information and select Finish.

Configure NVMe over FC on ESXi Host

To configure the NVMe over FC on ESXi host, follow these steps:

1. Login to vCenter and on the ESXi host verify the storage adapter information, there will be four adapters listed, two among them are the FC-NVMe initiators.
2. Once you click on one, you will see more information appear in the details panel:

Name	Subsystem NQN	Transport Type	FUSE Support	Model	Firmv Versi
<input type="radio"/> nqn.2010-06.com.purestorage:flasharray.71a3f052a63267d7#vmhba4#524a9375f2e3d511524a93	nqn.2010-06.com.purestorage:flasharray.71a3f052a63267d7	fibreChannel	true	Pure Storage FlashArray	99.9

- If the zoning is complete at this point no additional steps are required.
- The next step is to create the host and host group objects on the FlashArray. In NVMe-oF, initiators use something called an **NVMe Qualified Name (NQN)**.



The initiator has one and so does the target (the FlashArray). With NVMe-oF/FC, NQNs do not **replace** FC WWNs—they both exist.



The WWN of each side is what is advertised on the FC layer to enable physical connectivity and zoning. The NQN is what enables the NVMe layer to communicate to the correct endpoints on the FC fabric. You can look at it in a similar way as networking in IP (MAC addresses and IPs).

- For each ESXi host, you need to create a host object on the FlashArray, then add the NQN to it. So where do you get the NQN? However, not from the vSphere Client. For now, you need to use `esxcli`.
- So, SSH back into the ESXi host and run:

```
esxcli nvme info get
```

- Copy the NQN.
- Log into the FlashArray.

Host Registration

For Host registration, follow these steps in the Pure Storage Web Portal:

- Select Storage > Hosts.
- Select the + icon in the Hosts Panel.

3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.

Create Host

Name

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a “#” appearing in the name where an iterating number will appear:

Create Multiple Hosts

Name


Personality

Start Number


Count

Number of Digits

5. Click Create to add the hosts.
6. For each host created, select the host.
7. In the Host view, select ‘Configure NQNs...’ from the Host Ports menu.



8. A pop-up will appear for Configure NVMe-oF NQNs for <Host> Within this pop-up, enter the appropriate NQN of this specific host.

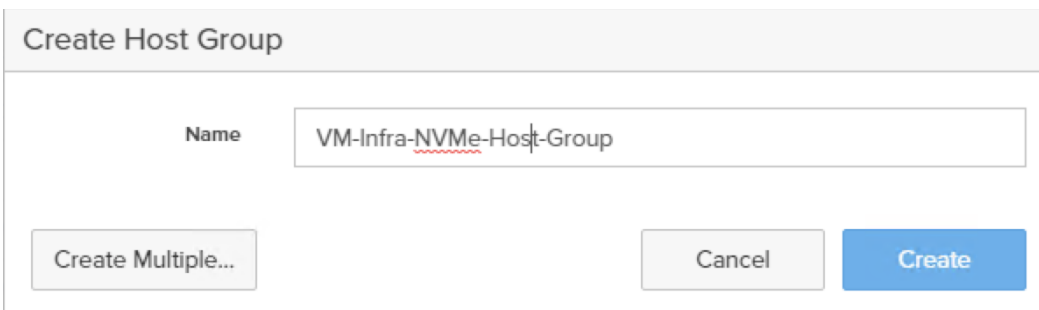


9. Click Add.
10. Repeat steps 1-9 for each host created.

Create NVMe Host Group

Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.
2. Select the + icon in the Host Groups Panel.
3. A pop-up will appear to create a host group on the FlashArray.



- Provide a name for the group and click Create.
- Select the group in the Host Groups Panel.
- In the Host Group view, select 'Add...' from the Member Hosts menu.

The screenshot shows the 'Storage' interface with the 'Hosts' tab selected. The breadcrumb path is 'Hosts > VM-Infra-NVMe-Host-Group'. A table at the top shows storage metrics: Size (0), Data Reduction (1.0 to 1), Unique (0.00), Snapshots (0.00), Shared (-), System (-), and Total (0.00). Below this are three sections: 'Member Hosts', 'Connected Volumes', and 'Protection Groups'. The 'Member Hosts' section has a search bar and a dropdown menu open with options: 'Add...', 'Remove...', and 'Download CSV'. The other two sections are currently empty, showing 'No hosts found.', 'No volumes found.', and 'No protection groups found.' respectively.

- Select the host to be part of the host group.

The screenshot shows the 'Add Hosts to Host Group' dialog box. It is divided into two main panels: 'Existing Hosts' and 'Selected Hosts'. In the 'Existing Hosts' panel, there are four items, each with a checked checkbox: 'NVMe', 'VM-Host-Infra-NVMe-01', 'VM-Host-Infra-NVMe-02', and 'VM-Host-Infra-NVMe-03'. The 'NVMe' item is highlighted. In the 'Selected Hosts' panel, there are three items, each with an 'x' icon for removal: 'VM-Host-Infra-NVMe-01', 'VM-Host-Infra-NVMe-02', and 'VM-Host-Infra-NVMe-03'. At the bottom right, there are 'Cancel' and 'Add' buttons.

- Click Add.

Create NVMe datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

Create Multiple Volumes ✕

Pod or Volume Group	<input type="text" value="none"/>
Name	<input type="text" value="VM-Infra-NVMe-DS#"/>
Provisioned Size	<input type="text" value="1"/> <input type="text" value="T"/>
Start Number	<input type="text" value="1"/>
Count	<input type="text" value="2"/>
Number of Digits	<input type="text" value="1"/>
QoS Configuration (Optional) ▾	

4. Fill in the Name and Provisioned Size.
5. Click Create to provision the volumes to be used as Infra datastore LUN.
6. Go back to the Hosts section under the Storage tab. Click ESXi cluster NVMe host group created earlier and select the gear icon pull-down within the Connected Volumes tab within that host group.

Storage 🚨 1 🔍 Search

Array **Hosts** Volumes Pods File Systems Policies

> Hosts > VM-Infra-NVMe-Host-Group

Size	Data Reduction	Unique	Snapshots	Shared	System	Total
0	1.0 to 1	0.00	0.00	-	-	0.00

Member Hosts ^ 1-3 of 3

Name	Interface	Size	Volumes	Reduction
VM-Host-Infra-NVMe-01	NVMe-oF	0	0.00	1.0 to 1
VM-Host-Infra-NVMe-02		0	0.00	1.0 to 1
VM-Host-Infra-NVMe-03		0	0.00	1.0 to 1

Connected Volumes ^

Name

No volumes found.

Protection Groups ^

Name

No protection groups found.

7. Within the drop-down of the gear icon, select Connect Volumes, and a pop-up will appear.

Connect Shared Volumes to Host Group ✕

Existing Volumes 1-3 of 3

NVMe

VM-Infra-NVMe-DS1

VM-Infra-NVMe-DS2

Infra_DS-NVME ● 1

Selected Volumes Clear all

2 selected

VM-Infra-NVMe-DS1 ✕

VM-Infra-NVMe-DS2 ✕

LUN

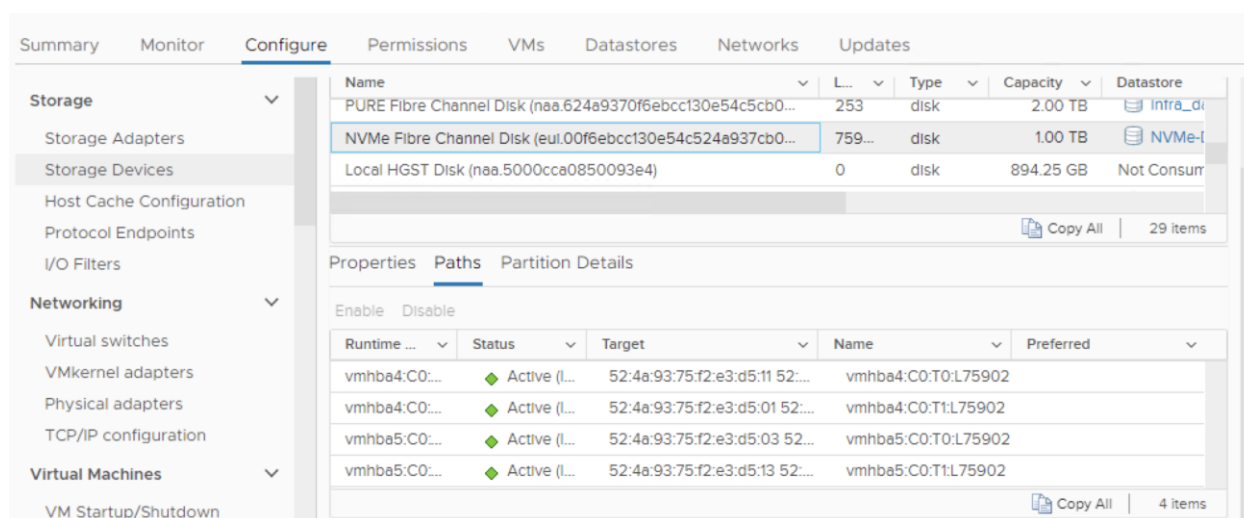
Cancel
Connect

8. Select the Infra datastore NVMe volumes that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.

ESXi Host NVMe over FC Datastore Configuration

To configure the ESXi host NVMe over FC datastore, follow these steps:

1. The remaining steps in the VMware vSphere Client are manual steps that should be completed whether Ansible configuration or manual configuration is being done. Verify that the NVMe Fibre Channel Disk is mounted on each ESXi host. Under Hosts and Clusters select the ESXi host. In the center pane, select Configure > Storage > Storage Devices. The NVMe Fibre Channel Disk should be listed under Storage Devices. Select the NVMe Fibre Channel Disk, then select Paths underneath. Verify 4 paths have a status of Active (I/O). Repeat this for all 3 hosts.



Runtime ...	Status	Target	Name	Preferred
vmhba4:C0:...	Active (I/O)	52:4a:93:75:f2:e3:d5:11 52:...	vmhba4:C0:T0:L75902	
vmhba4:C0:...	Active (I/O)	52:4a:93:75:f2:e3:d5:01 52:...	vmhba4:C0:T1:L75902	
vmhba5:C0:...	Active (I/O)	52:4a:93:75:f2:e3:d5:03 52:...	vmhba5:C0:T0:L75902	
vmhba5:C0:...	Active (I/O)	52:4a:93:75:f2:e3:d5:13 52:...	vmhba5:C0:T1:L75902	

2. For any of the three hosts, right-click the host under Hosts and Clusters and select Storage > New Datastore. Leave VMFS selected and click NEXT.
3. Name the datastore and select the NVMe Fibre Channel Disk. Click NEXT.
4. Leave VMFS 6 selected and click NEXT.
5. Leave all Partition configuration values at the default values and click NEXT.
6. Review the information and click FINISH.
7. Select Storage and select the just-created NVMe datastore. In the center pane, select Hosts. Ensure all three hosts have the datastore mounted.

ESXi Host Multipathing Configuration

To configure the ESXi host multipathing, follow these steps:

1. From the vCenter management GUI.
2. Go to Hosts and Clusters view.

3. Select a Host.
4. Click the Configure tab.
5. Select Storage Devices.
6. Select an NVMe device.
7. Click Edit Multipathing.

Edit Multipathing Policies | eui.00f6ebcc130e54c524a937cb0001287f ×

Path selection policy LB-Latency ▾

Latency evaluation time ⓘ 180000 ▴ ▾
The value must be between 10000 and 300000

Sampling I/Os per path ⓘ 16
The value must be between 16 and 160

CANCEL SAVE

Appendix

FlashStack iSCSI Addition

Cisco Nexus Switch Configuration

This section is a delta section for adding infrastructure iSCSI to the Cisco Nexus switches. This section should be executed after the Cisco Nexus Switch Configuration section in the main document is completed.

Create Infrastructure iSCSI VLANs on Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
config t
vlan <infra-iscsi-a-vlan-id>
name Infra-iSCSI-A-VLAN
vlan <infra-iscsi-b-vlan-id>
name Infra-iSCSI-B-VLAN
exit
```

Add iSCSI Individual Port Descriptions for Troubleshooting and Enable UDLD for Pure iSCSI Interfaces

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A connected to Cisco Pure FlashArray//X R3, follow this step, follow this step:

1. From the global configuration mode, run the following commands:

```
config t
interface Ethernet1/37
description <<var_flasharray_hostname>>-CT0.ETH4
interface Ethernet1/38
description <<var_flasharray_hostname>>-CT1.ETH4
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B connected to Cisco Pure FlashArray//X R3, follow this step:

1. From the global configuration mode, run the following commands:

```
config t
interface Ethernet1/37
description <<var_flasharray_hostname>>-CT0.ETH5
interface Ethernet1/38
description <<var_flasharray_hostname>>-CT1.ETH5
```

Configure iSCSI interfaces for Cisco Nexus 93180YC-FX-A

To configure iSCSI interfaces for this deployment, run the following commands on Cisco Nexus 93180YC-FX-A:

```
config t
interface Ethernet1/37
switchport
switchport access valn <<var-iscsi-a-vlan-id>>
mtu 9216
no negotiate auto
no shut
interface Ethernet1/38
switchport
switchport access valn <<var-iscsi-a-vlan-id>>
mtu 9216
no negotiate auto
no shut
```

Configure iSCSI interfaces for Cisco Nexus 93180YC-FX-B

To configure iSCSI interfaces for this deployment, run the following commands on Cisco Nexus 93180YC-FX-B:

```
config t
interface Ethernet1/37
switchport
switchport access valn <<var-iscsi-b-vlan-id>>
mtu 9216
no negotiate auto
no shut
interface Ethernet1/38
switchport
switchport access valn <<var-iscsi-b-vlan-id>>
mtu 9216
no negotiate auto
no shut
```

Add Infrastructure iSCSI VLANs to Port-Channels on Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), follow this step on both switches:

1. From the global configuration mode, run the following commands:

```
interface Po10
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
interface Po121
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
interface Po123
switchport trunk allowed vlan add <infra-iscsi-a-vlan-id>,<infra-iscsi-b-vlan-id>
exit
copy run start
```

FlashArray //X R3 iSCSI Interface Configuration

The iSCSI traffic will be carried on two VLANs, A (901) and B (902) that are configured in our example with the following values:

Table 15.iSCSI A FlashArray//X50 R3 Interface Configuration Settings

FlashArray Controller	iSCSI Port	IP Address	Subnet Mask
FlashArray//X R3 Controller 0	CT0.ETH4	192.168.101.146	255.255.255.0
FlashArray//X R3 Controller 1	CT0.ETH4	192.168.101.147	255.255.255.0

Table 16.iSCSI B FlashArray//X50 R3 Interface Configuration Settings

FlashArray Controller	iSCSI Port	IP Address	Subnet Mask
FlashArray//X R3 Controller 0	CT0.ETH5	192.168.102.146	255.255.255.0
FlashArray//X R3 Controller 1	CT0.ETH5	192.168.102.147	255.255.255.0

To configure iSCSI interfaces for environments deploying iSCSI boot LUNs and/or datastores, follow these steps from Pure FlashArray Web Portal:.

1. Select Settings > Network
2. Click the Edit Icon for interface CT0.eth4.
3. Select Enable and add the IP information from above tables and set the MTU to 9000.

The screenshot shows a window titled "Edit Network Interface" with a close button (X) in the top right corner. The window contains several configuration fields:

- Name:** ct0.eth4
- Enabled:** A toggle switch is turned on (blue).
- Address:** 192.168.101.146
- Netmask:** 255.255.255.0
- Gateway:** 192.168.101.254
- MAC:** 24:a9:37:0d:df:b3
- MTU:** 9000
- Service(s):** iscsi

At the bottom right of the window, there are two buttons: "Cancel" (light gray) and "Save" (blue).

4. Click Save.
5. Repeat steps 1-4 for CT0.eth5, CT1.eth4, and CT1.eth5.

Cisco UCS iSCSI Configuration

The following subsections can be completed to add infrastructure iSCSI to the Cisco UCS. These subsections can be completed in place of the subsections in the Cisco UCS Configuration section of this document labeled (FCP), or they can be completed in addition to the FCP sections to have the option of FCP or iSCSI boot.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click SAN.
2. Expand Pools > root.
3. Right-click IQN Pools.
4. Choose Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.

7. Enter iqn.2010-11.com.flashStack as the prefix.

8. Choose Sequential for Assignment Order.

9. Click Next.

10. Click Add.

11. Enter ucs-host as the suffix.




If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.

12. Enter 1 in the From field.

13. Specify the size of the IQN block sufficient to support the available server resources.

Create a Block of IQN Suffixes



Suffix :

From :

Size :

OK

Cancel

14. Click OK.

15. Click Finish and then click OK to complete creating the IQN pool.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools for iSCSI boot for the Cisco UCS environment, follow these steps:



The IP Pools for iSCSI Boot are created here in the root organization. If servers will be booted from UCS tenant organizations, consider creating the IP Pools for iSCSI Boot in the tenant organization.

1. In Cisco UCS Manager, click LAN.
2. Expand Pools > root.
3. Right-click IP Pools.
4. Choose Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Choose Sequential for the assignment order.
8. Click Next.
9. Click Add to add a block of IP addresses.
10. In the From field, enter the beginning of the range to assign as iSCSI boot IP addresses on Fabric A.
11. Set the size to enough addresses to accommodate the servers.
12. Enter the appropriate Subnet Mask.
13. Click OK.
14. Click Next.
15. Click Finish and then click OK to complete creating the Fabric A iSCSI IP Pool.
16. Right-click IP Pools.
17. Choose Create IP Pool.
18. Enter iSCSI-IP-Pool-B as the name of IP pool.
19. Optional: Enter a description for the IP pool.
20. Choose Sequential for the assignment order.
21. Click Next.

22. Click Add to add a block of IP addresses.

23. In the From field, enter the beginning of the range to assign as iSCSI IP addresses on Fabric B.

24. Set the size to enough addresses to accommodate the servers.

25. Enter the appropriate Subnet Mask.

26. Click OK.

27. Click Next.

28. Click Finish and then click OK to complete creating the Fabric B iSCSI IP Pool.

Create iSCSI VLANs

To configure the necessary iSCSI virtual local area networks (VLANs) for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > LAN Cloud.
3. Right-click VLANs.
4. Choose Create VLANs.
5. Enter iSCSI-A-VLAN as the name of the VLAN to be used for iSCSI-A.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the iSCSI-A VLAN ID.
8. Keep the Sharing Type as None.

Create VLANs



VLAN Name/Prefix :

Multicast Policy Name : [Create Multicast Policy](#)

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. " 2009-2019", " 29,35,40-45", " 23", " 23,34-45")

VLAN IDs :

Sharing Type : None Primary Isolated Community

Check Overlap

OK

Cancel

9. Click OK and then click OK again.

10. Right-click VLANs.

11. Choose Create VLANs.

12. Enter iSCSI-B-VLAN as the name of the VLAN to be used for iSCSI-B.

13. Keep the Common/Global option selected for the scope of the VLAN.

14. Enter the iSCSI-B VLAN ID.

15. Keep the Sharing Type as None.

16. Click OK and then click OK again.

Create iSCSI vNIC Templates

To create iSCSI virtual network interface card (vNIC) templates for the Cisco UCS environment within the FlashStack-VSI Organization, follow these steps:

1. Choose LAN.

2. Expand Policies > root > Sub-Organizations > FlashStack-VSI Organization.

-
3. Right-click vNIC Templates under the FlashStack-VSI Organization.
 4. Choose Create vNIC Template.
 5. Enter iSCSI-A as the vNIC template name.
 6. Choose Fabric A. Do not choose the Enable Failover checkbox.
 7. Leave Redundancy Type set at No Redundancy.
 8. Under Target, make sure that only the Adapter checkbox is selected.
 9. Choose Updating Template for Template Type.
 10. Under VLANs, choose only iSCSI-A-VLAN.
 11. Choose iSCSI-A-VLAN as the native VLAN.
 12. Leave vNIC Name set for the CDN Source.
 13. Under MTU, enter 9000.
 14. From the MAC Pool list, choose MAC-Pool-A.
 15. From the Network Control Policy list, choose Enable-CDP-LLDP.

Create vNIC Template ? X

VLANs | VLAN Groups

Advanced Filter | Export | Print ⚙️

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	default	<input type="radio"/>	1
<input type="checkbox"/>	IB-MGMT-VLAN	<input type="radio"/>	115
<input checked="" type="checkbox"/>	iSCSI-A-VLAN	<input checked="" type="radio"/>	901
<input type="checkbox"/>	Native-Vlan	<input type="radio"/>	2
<input type="checkbox"/>	oob-mgmt	<input type="radio"/>	15
<input type="checkbox"/>	VM-Traffic-VLAN	<input type="radio"/>	1101

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

16. Click OK to complete creating the vNIC template.

17. Click OK.

18. Right-click vNIC Templates.

19. Choose Create vNIC Template.

-
20. Enter iSCSI-B as the vNIC template name.
 21. Choose Fabric B. Do not choose the Enable Failover checkbox.
 22. Leave Redundancy Type set at No Redundancy.
 23. Under Target, make sure that only the Adapter checkbox is selected.
 24. Choose Updating Template for Template Type.
 25. Under VLANs, choose only iSCSI-B-VLAN.
 26. Choose iSCSI-B-VLAN as the native VLAN.
 27. Leave vNIC Name set for the CDN Source.
 28. Under MTU, enter 9000.
 29. From the MAC Pool list, choose MAC-Pool-B.
 30. From the Network Control Policy list, choose Enable-CDP-LLDP.
 31. Click OK to complete creating the vNIC template.
 32. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

To configure the necessary Infrastructure LAN Connectivity Policy within the FlashStack-VSI Organization, follow these steps:

1. In Cisco UCS Manager, click LAN.
2. Expand LAN > Policies > root > Sub-Organizations > FlashStack-VSI Organization.
3. Right-click LAN Connectivity Policies under the FlashStack-VSI Organization.
4. Choose Create LAN Connectivity Policy.
5. Enter iSCSI-Boot as the name of the policy.
6. Click OK then OK again to create the policy.
7. On the left under LAN > Policies > root > Sub-Organizations > FlashStack-VSI Organization > LAN Connectivity Policies, choose iSCSI-Boot.
8. Click Add to add a vNIC.
9. In the Create vNIC dialog box, enter 00-vSwitch0-A as the name of the vNIC.

-
10. Choose the Use vNIC Template checkbox.
 11. In the vNIC Template list, choose vSwitch0-A.
 12. In the Adapter Policy list, choose VMWare.
 13. Click OK to add this vNIC to the policy.
 14. Click Save Changes and then click OK.
 15. Click the Add button to add another vNIC to the policy.
 16. In the Create vNIC box, enter 01-vSwitch0-B as the name of the vNIC.
 17. Choose the Use vNIC Template checkbox.
 18. In the vNIC Template list, choose vSwitch0-B.
 19. In the Adapter Policy list, choose VMWare.
 20. Click OK to add the vNIC to the policy.
 21. Click Save Changes and then click OK.
 22. Click the Add button to add a vNIC.
 23. In the Create vNIC dialog box, enter 02-vDS0-A as the name of the vNIC.
 24. Choose the Use vNIC Template checkbox.
 25. In the vNIC Template list, choose vDS0-A.
 26. In the Adapter Policy list, choose VMWare-HighTrf.
 27. Click OK to add this vNIC to the policy.
 28. Click Save Changes and then click OK.
 29. Click the Add button to add another vNIC to the policy.
 30. In the Create vNIC box, enter 03-vDS0-B as the name of the vNIC.
 31. Choose the Use vNIC Template checkbox.
 32. In the vNIC Template list, choose vDS0-B.
 33. In the Adapter Policy list, choose VMWare-HighTrf.

-
34. Click OK to add the vNIC to the policy.
 35. Click Save Changes and then click OK.
 36. Click the Add button to add a vNIC.
 37. In the Create vNIC dialog box, enter 04-iSCSI-A as the name of the vNIC.
 38. Choose the Use vNIC Template checkbox.
 39. In the vNIC Template list, choose iSCSI-A.
 40. In the Adapter Policy list, choose VMWare.
 41. Click OK to add this vNIC to the policy.
 42. Click Save Changes and then click OK.
 43. Click Add to add a vNIC to the policy.
 44. In the Create vNIC dialog box, enter 05-iSCSI-B as the name of the vNIC.
 45. Choose the Use vNIC Template checkbox.
 46. In the vNIC Template list, choose iSCSI-B.
 47. In the Adapter Policy list, choose VMWare.
 48. Click OK to add this vNIC to the policy.
 49. Click Save Changes and then click OK.
 50. Expand Add iSCSI vNICs.
 51. Choose Add in the Add iSCSI vNICs section.
 52. Set the name to iSCSI-Boot-A.
 53. Choose 04-iSCSI-A as the Overlay vNIC.
 54. Set the iSCSI Adapter Policy to default.
 55. Leave the VLAN set to Infra-iSCSI-A (native).
 56. Leave the MAC Address set to None.
 57. Click OK.

58. Click Save Changes and then click OK.
59. Choose Add in the Add iSCSI vNICs section.
60. Set the name to iSCSI-Boot-B.
61. Choose 05-iSCSI-B as the Overlay vNIC.
62. Set the iSCSI Adapter Policy to default.
63. Leave the VLAN set to Infra-iSCSI-B (native).
64. Leave the MAC Address set to None.
65. Click OK.
66. Click Save Changes and then click OK.

General
Events

Actions

Delete

Show Policy Usage

Use Global

Name : **iSCSI-Boot**

Description :

Owner : **Local**

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Native VLAN
▶ vNIC 00-vSwitch0-A	Derived	
▶ vNIC 01-vSwitch0-B	Derived	
▶ vNIC 02-VDS-A	Derived	
▶ vNIC 03-VDS-B	Derived	
▶ vNIC 04-iSCSI-A	Derived	
▶ vNIC 05-iSCSI-B	Derived	

🗑️ Delete
➕ Add
⚙️ Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-A	04-iSCSI-A	default	Derived
iSCSI vNIC iSCSI-Boot-B	05-iSCSI-B	default	Derived

Create iSCSI Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI interfaces on FlashArray//X R3 controller 1 (ct0.eth4 and ct0.eth5) and two iSCSI interfaces on FlashArray//X R3 controller 2 (ct1.eth4 and ct1.eth5). Also, it is assumed that the ports numbered 4 are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the ports numbered 5 are connected to Fabric B (Cisco UCS Fabric Interconnect B).



One boot policy is configured in this procedure. The policy configures the primary target to be `iscsi-lif01a`.

To create a boot policy for the Cisco UCS environment within the FlashStack-VSI Organization, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Policies > root > Sub-Organizations > FlashStack-VSI Organization.
3. Right-click Boot Policies under the FlashStack-VSI Organization.
4. Choose Create Boot Policy.
5. Enter Boot-iSCSI as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Do not choose the Reboot on Boot Order Change checkbox.
8. Choose the Uefi Boot Mode.
9. Check the checkbox for Boot Security.
10. Expand the Local Devices drop-down list and click Add Remote CD/DVD.
11. Expand the iSCSI vNICs drop-down list and click Add iSCSI Boot.
12. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-A.
13. Click OK.
14. Choose Add iSCSI Boot.
15. In the Add iSCSI Boot dialog box, enter iSCSI-Boot-B.
16. Click OK.
17. Expand CIMC Mounted Media and select Add CIMC Mounted CD/DVD.

Create Boot Policy



Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

Boot Security :

WARNINGS:

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

- Local Devices
- CIMC Mounted vMedia
- vNICs
- vHBAs
- iSCSI vNICs
-
- EFI Shell

Boot Order

+ - 🔍 Advanced Filter ↑ Export 🖨 Print ⚙

Name	O...	vNIC/vHBA/iSCSI vNIC	Type	LUN ...	WWN	Slot ...	Boot...	Boot...	Des...
Remote CD/DVD	1								
▼ iSCSI	2								
iSCSI		iSCSI-Boot-A	Prim...						
iSCSI		iSCSI-Boot-B	Sec...						
CIMC Mounted C...	3								

↑ Move Up ↓ Move Down 🗑 Delete

Set Uefi Boot Parameters

18. Expand iSCSI and select iSCSI-Boot-A. Select Set Uefi Boot Parameters.



For Cisco UCS B200 M5 and Cisco UCS C220 M5 servers it is not necessary to set the Uefi Boot Parameters. These servers will boot properly with or without these parameters set. However, for Cisco UCS M4 and earlier servers, VMware ESXi 7.0 will not boot with Uefi Secure Boot unless these parameters are set exactly as shown.

19. Fill-in the Set Uefi Boot Parameters exactly as shown in the following screenshot:

Set Uefi Boot Parameters



Uefi Boot Parameters

Boot Loader Name	:	<input type="text" value="BOOTX64.EFI"/>
Boot Loader Path	:	<input "="" type="text" value="\EFI\BOOT\"/>
Boot Loader Description	:	<input type="text"/>



20. Click OK to complete setting the Uefi Boot Parameters for the SAN Boot Target and click OK for the confirmation.

21. Repeat steps 1-20 to set Uefi Boot Parameters for each of the 2 iSCSI Boot Targets.

22. Click OK then click OK again to create the policy.

Create iSCSI Boot Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts within the FlashStack-VSI Organization is created for Fabric A boot.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click Servers.
2. Expand Service Profile Templates > root > Sub-Organizations > FlashStack-VSI Organization.
3. Right-click the FlashStack-VSI Organization.
4. Choose Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter AMD-VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Choose the Updating Template option.
7. Under UUID Assignment, choose UUID_Pool.

Create Service Profile Template ? X

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to this template and enter a description.

Name :

The template will be created in the following organization. Its name must be unique within this organization.
Where : **org-root/org-FlashStack-VSI**

The template will be created in the following organization. Its name must be unique within this organization.
Type : Initial Template Updating Template

Specify how the UUID will be assigned to the server associated with the service generated by this template.
UUID

UUID Assignment:

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

8. Click Next.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

1. If you have servers with no physical disks, click on the Local Disk Configuration Policy tab and choose the ignoreDisk Local Storage Policy. Otherwise, choose the default Local Storage Policy.
2. Click Next.

Configure Networking Options

To configure the network options, follow these steps:

1. Choose the “Use Connectivity Policy” option to configure the LAN connectivity.
2. Choose iSCSI-Boot from the LAN Connectivity Policy drop-down list.
3. Choose IQN_Pool in Initiator Name Assignment.

Create Service Profile Template ? ×

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: ▼

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple
 Expert
 No vNICs
 Use Connectivity Policy

LAN Connectivity Policy : ▼ [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: ▼

Initiator Name :

[Create IQN Suffix Pool](#)

The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

< Prev Next > **Finish** Cancel

4. Click Next.

Configure Storage Options

To configure the storage options, follow these steps:

1. Choose No vHBAs for the “How would you like to configure SAN connectivity?” field.
2. Click Next.

Configure Zoning Options

To configure the zoning options, follow this step:

1. Make no changes and click Next.

Configure vNIC/HBA Placement

To configure the vNIC/HBA placement, follow these steps:

1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement”.

2. Click Next.

Configure vMedia Policy

To configure the vMedia policy, follow these steps:

1. Do not select a vMedia Policy.
2. Click Next.

Configure Server Boot Order

To configure the server boot orders, follow these steps:

1. Choose Boot-iSCSI for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Create Boot Policy

Name : **Boot-iSCSI**

Description :

Reboot on Boot Order Change : **No**

Enforce vNIC/vHBA/iSCSI Name : **Yes**

Boot Mode : **Uefi**

Boot Security : **Yes**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

+ - Advanced Filter Export Print

Name	Order	vNIC/vHB...	Type	LUN Name	WWN	Slot Numb...	Boot Name	Boot Path	Description
Remot...	1								
▶ iSCSI	2								
CIMC ...	3								

Create iSCSI vNIC Set iSCSI Boot Parameters Set Uefi Boot Parameters

< Prev Next > Finish Cancel

2. In the Boot order, expand iSCSI and choose iSCSI-Boot.

3. Click Set iSCSI Boot Parameters.

4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have in-dependently created one appropriate to your environment.

5. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI-IP-Pool-A as the “Initiator IP address Policy.”
7. Choose iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI Target Name.
10. Enter the IP address of ct0.eth4 for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

11. Click OK to add the iSCSI static target.
12. Click Add.
13. Enter the iSCSI Target Name.
14. Enter the IP address of ct0.eth5 for the IPv4 Address field.
15. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



Name : **iSCSI-Boot-A**

Authentication Profile : <not set> ▼

[Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment: <not set> ▼

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI-Pool-A(47/50) ▼

IPv4 Address : **0.0.0.0**
Subnet Mask : **255.255.255.0**
Default Gateway : **0.0.0.0**
Primary DNS : **0.0.0.0**
Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

[Reset Initiator Address](#)

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPv4 Addre...	LUN Id
iqn.2010-06.c...	1	3260		192.168.101.146	1
iqn.2010-06.c...	2	3260		192.168.101.147	1

OK

Cancel

-
16. Click OK to complete setting the iSCSI Boot Parameters.
 17. In the Boot order, choose iSCSI-Boot-B.
 18. Click Set iSCSI Boot Parameters.
 19. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have in-dependently created one appropriate to your environment.
 20. Leave the “Initiator Name Assignment” dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.
 21. Set iSCSI-IP-Pool-B as the “Initiator IP address Policy”.
 22. Choose the iSCSI Static Target Interface option.
 23. Click Add.
 24. Enter the iSCSI Target Name.
 25. Enter the IP address of ct0.eth5 for the IPv4 Address field.
 26. Click OK to add the iSCSI static target.
 27. Click Add.
 28. Enter the iSCSI Target Name.
 29. Enter the IP address of ct1.eth5 for the IPv4 Address field.
 30. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters



Name : **iSCSI-Boot-B**

Authentication Profile : <not set> ▼

[Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment: <not set> ▼

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI-Pool-B(47/50) ▼

IPv4 Address : **0.0.0.0**

Subnet Mask : **255.255.255.0**

Default Gateway : **0.0.0.0**

Primary DNS : **0.0.0.0**

Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

[Reset Initiator Address](#)

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Addre...	LUN Id
iqn.2010-06.c...	1	3260		192.168.102.146	1
iqn.2010-06.c...	2	3260		192.168.102.147	1

OK

Cancel

31. Click OK to complete setting the iSCSI Boot Parameters.

32. Click Next.

Configure Maintenance Policy

To configure the maintenance policy, follow these steps:

1. Change the Maintenance Policy to default.

Create Service Profile Template ? ×

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

⊖ Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: [Create Maintenance Policy](#)

Name	: default
Description	:
Soft Shutdown Timer	: 150 Secs
Storage Config. Deployment Policy	: User Ack
Reboot Policy	: User Ack

< Prev Next > **Finish** Cancel

2. Click Next.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, choose Infra-Pool.
2. Choose Down as the power state to be applied when the profile is associated with the server.
3. Optional: Choose “UCSB-B200-M6” for the Server Pool Qualification to choose only UCS M6 servers in the pool.

4. Expand Firmware Management and choose the default Host Firmware Package.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment:

Select the power state to be applied when this profile is associated with the server.

Up Down

The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification :

Restrict Migration :

Firmware Management (BIOS, Disk Controller, Adapter)

5. Click Next.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, choose Intel-M6-Virt.
2. Expand Power Control Policy Configuration and choose No-Power-Cap in the Power Control Policy list.

Create Service Profile Template ? ×

Optionally specify information that affects how the system operates.

BIOS Configuration

If you want to override the default BIOS settings, select a BIOS policy that will be associated with this service profile

BIOS Policy :

External IPMI/Redfish Management Configuration

Management IP Address

Monitoring Configuration (Thresholds)

Power Control Policy Configuration

Power control policy determines power allocation for a server in a given power group.

Power Control Policy : [Create Power Control Policy](#)

Scrub Policy

KVM Management Policy

Graphics Card Policy

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create vMedia-Enabled Service Profile Template

To create a service profile template with vMedia enabled, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI Organization > Service Template VM-Host-Infra-iSCSI.
3. Right-click VM-Host-Infra-iSCSI and click Create a Clone.
4. Name the clone VM-Host-Infra-iSCSI-vM and click OK then click OK again to create the clone.
5. Choose the newly created VM-Host-Infra-iSCSI-vM and choose the vMedia Policy tab.

-
6. Click Modify vMedia Policy.
 7. Choose the VM-Host-Infra-iSCSI vMedia Policy and click OK.
 8. Click OK to confirm.

Create Intel Optane Memory Mode Service Profile Template (Optional)

To create a service profile template with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click VM-Host-Infra-iSCSI and choose Create a Clone.
4. Name the clone Intel-MM-Host-Infra-iSCSI.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-MM-Host-Infra-iSCSI and choose the Policies tab.
7. Expand Persistent Memory Policy and use the pulldown to select the Memory-Mode Policy.
8. Click save Changes.
9. Click OK to confirm.

Create vMedia-Enabled Intel Optane Memory Mode Service Profile Template (Optional)

To create a service profile template with vMedia enabled for servers with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click Intel-MM-Host-Infra-iSCSI and choose Create a Clone.
4. Name the clone Intel-MM-Host-Infra-iSCSI-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-MM-Host-Infra-iSCSI-vM and choose the vMedia Policy tab.

-
7. Click Modify vMedia Policy.
 8. Choose the ESXi-7U2-CC-HTTP vMedia Policy and click OK.
 9. Click OK to confirm.

Create Intel Optane App Direct Mode Service Profile Template (Optional)

To create a service profile template with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click VM-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-AD-Host-Infra-FCP.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-AD-Host-Infra-iSCSI and choose the Policies tab.
7. Expand Persistent Memory Policy and use the pulldown to select the Memory-Mode Policy.
8. Click save Changes.
9. Click OK to confirm.

Create vMedia-Enabled Intel Optane App Direct Mode Service Profile Template (Optional)

To create a service profile template with vMedia enabled for servers with Intel Optane DC PMEM installed and Memory Mode enabled, follow these steps:

1. Connect to UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI > Service Template VM-Host-Infra-FCP.
3. Right-click Intel-AD-Host-Infra-FCP and choose Create a Clone.
4. Name the clone Intel-AD-Host-Infra-iSCSI-vM.
5. Click OK then click OK again to create the Service Profile Template clone.
6. Choose the newly created Intel-AD-Host-Infra-iSCSI-vM and choose the vMedia Policy tab.

-
7. Click Modify vMedia Policy.
 8. Choose the ESXi-7U2-CC-HTTP vMedia Policy and click OK.
 9. Click OK to confirm.

Create Service Profiles

To create service profiles from the service profile template, follow these steps:

1. Connect to Cisco UCS Manager and click Servers.
2. Choose Service Profile Templates > root > Sub-Organizations > FlashStack-VSI Organization > Service Template VM-Host-Infra-iSCSI-vM.
3. Right-click VM-Host-Infra-iSCSI-vM and choose Create Service Profiles from Template.
4. For Naming Prefix, enter Infra-ESXi-iSCSI-0.
5. For Name Suffix Starting Number, enter 1.
6. For Number of Instances, enter 3.

Create Service Profiles From Template ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

OK

Cancel

7. Click OK to create the service profiles.
8. Click OK in the confirmation message.



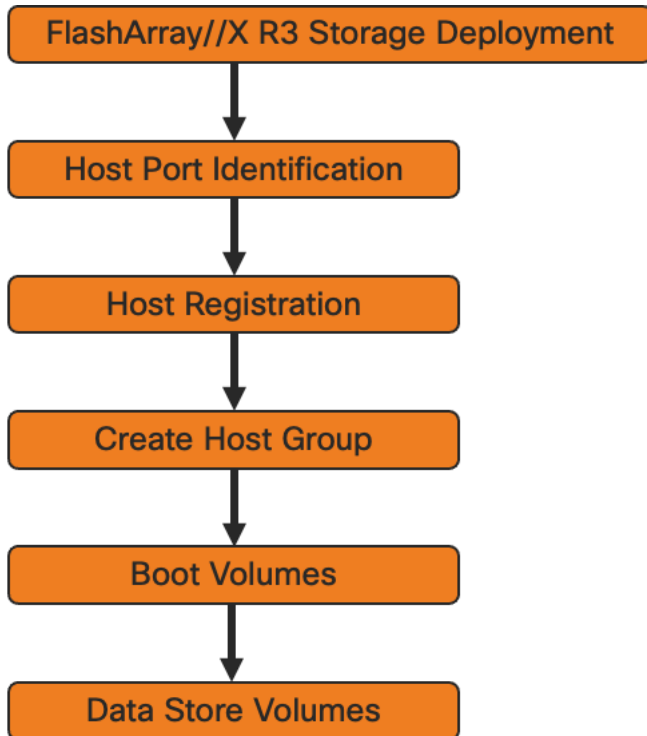
When VMware ESXi 7.0 has been installed on the hosts, the host Service Profiles can be bound to the AMD-VM-Host-Infra-iSCSI-A Service Profile Template to remove the vMedia Mapping from the host.

FlashArray Storage Deployment

The Pure Storage FlashArray//X is accessible to the FlashStack, but no storage has been deployed at this point. The storage to be deployed will include:

- ESXi iSCSI Boot LUNs
- VMFS Datastores
- vVOL Data Stores

The iSCSI Boot LUNs will need to be setup from the Pure Storage Web Portal, and the VMFS datastores can be directly provisioned from the vSphere Web Client after the Pure Storage vSphere Web Client Plugin has later been registered with the vCenter.



Host Port Identification

iSCSI Boot LUNs will be mapped by the FlashArray//X R3 using the assigned Initiator IQN to the provisioned service profiles. This information can be found within the service profile located within the iSCSI vNIC tab:

Servers / Service Profiles / root / Service Profile VM-Host...

< General Storage Network **iSCSI vNICs** vMedia Policy Boot Order Virtual Machines FC Zones Policies Server Details CIMC Sessions FSM VIF Paths >

Actions

Change Initiator Name
Reset Initiator Name

Service Profile Initiator Name

IQN Pool Name : **IQN-Pool**
Initiator Name : **iqn.1992-08.cisco.com:ucs-host:1**

No Configuration Change of vNICs/vHBAs/iSCSI vNICs is allowed due to connectivity policy.

iSCSI vNICs

+ - Advanced Filter Export Print

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-A-vNIC	06-iSCSI-A	default	Derived
iSCSI vNIC iSCSI-B-vNIC	07-iSCSI-B	default	Derived

Host Registration

For Host registration, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.
2. Select the + icon in the Hosts Panel.
3. After clicking the Create Host (+) option, a pop-up will appear to create an individual host entry on the FlashArray.

Create Host

Name

4. To create more than one host entry, click the Create Multiple... option, filling in the Name, Start Number, Count, Personality as ESXi and Number of Digits, with a “#” appearing in the name where an iterating number will appear:

Create Multiple Hosts

Name:


Personality:

Start Number:

Count:

Number of Digits:

5. Click Create to add the hosts.
6. For each host created, select the host.
7. In the Host view, select 'Configure IQNs...' from the Host Ports menu.

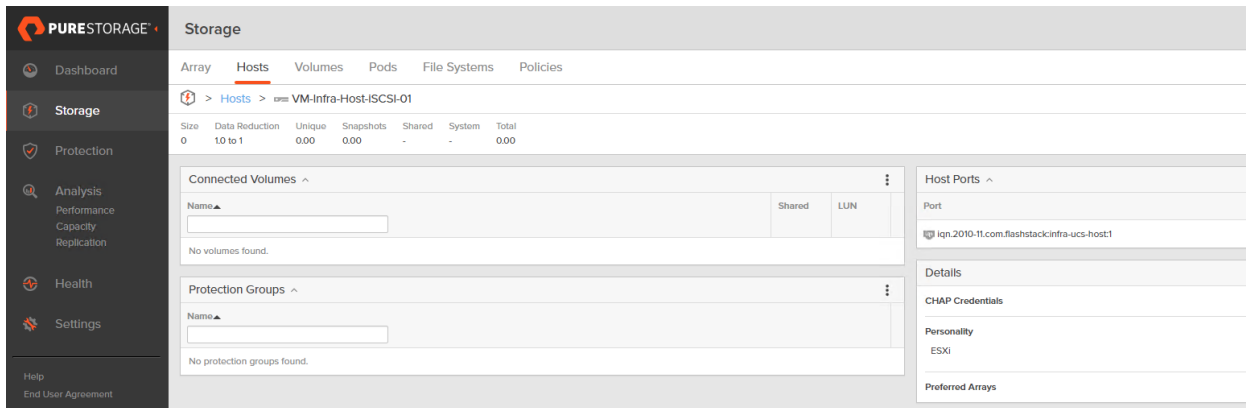


8. A pop-up will appear for Configure iSCSI IQNs for Host <host being configured>. Within this pop-up, enter the IQN Initiator Name found within the service profile for the host being configured:

Configure iSCSI IQNs for 'VM-Infra-Host-iSCSI-01'

Port IQNs:

9. After entering the IQN, click Add to add the Host Ports.




10. Repeat steps 1-9 for each host created.

Create Host Group

Host Groups allow the Administrator to map Volumes to a group of hosts at once with the same LUN ID. To create a Host Group, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Hosts.
2. Select the + icon in the Host Groups Panel.
3. A pop-up will appear to create a host group on the FlashArray.



4. Provide a name for the group and click Create.
5. Select the group in the Host Groups Panel.
6. In the Host Group view, select 'Add...' from the Member Hosts menu.



7. Select the host to be part of the host group.

Add Hosts to Host Group
✕

Existing Hosts	Selected Hosts
<input type="checkbox"/> <input style="width: 100%;" type="text"/> 1-4 of 4	<div style="display: flex; justify-content: space-between;"> 3 selected Clear all </div>
<input type="checkbox"/> iSCSI-Test1	VM-Infra-Host-iSCSI-01 ✕
<input checked="" type="checkbox"/> VM-Infra-Host-iSCSI-01	VM-Infra-Host-iSCSI-02 ✕
<input checked="" type="checkbox"/> VM-Infra-Host-iSCSI-02	VM-Infra-Host-iSCSI-03 ✕
<input checked="" type="checkbox"/> VM-Infra-Host-iSCSI-03	

Cancel
Add

8. Click Add.

Private Boot Volumes for each ESXi Host

To create private boot volumes for each ESXi Host, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

Create Volume
✕

Pod or Volume Group

Name

Provisioned Size

G ▼

QoS Configuration (Optional) ▼

Create Multiple...
Cancel
Create

4. To create more than one volume, click the Create Multiple... option, filling in the Name, Provisioned Size, Starting Number, Count, and Number of Digits, with a “#” appearing in the name where an iterating number will appear.

Create Multiple Volumes ✕

Pod or Volume Group

Name

Provisioned Size G ▼

Start Number

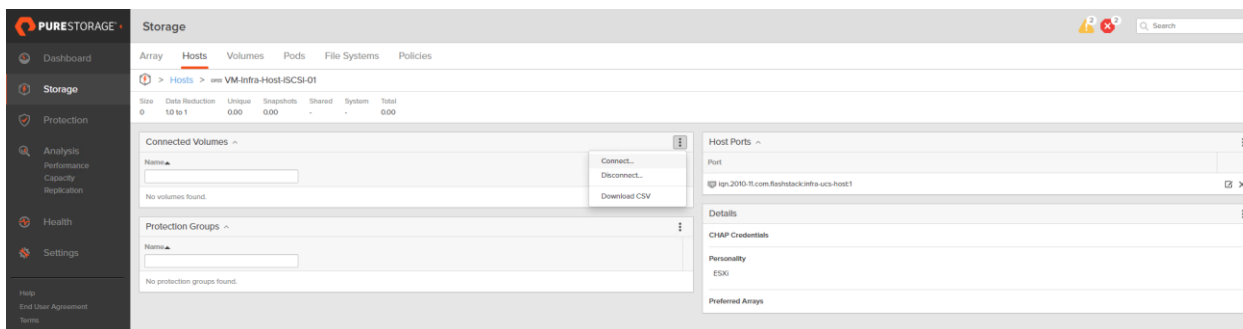
Count

Number of Digits

QoS Configuration (Optional) ▼

Create Single...
Cancel
Create

5. Click Create to provision the volumes to be used as iSCSI boot LUNs.
6. Go back to the Hosts section under the Storage tab. Click one of the hosts and select the gear icon pull-down within the Connected Volumes tab within that host.



7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.

Connect Volumes to Host

Existing Volumes 1-4 of 4

- VM-Infra
- VM-Infra-Boot-iSCSI-01
- VM-Infra-Boot-FCP-boot-01
- VM-Infra-Host-FC-boot-03
- VM-Infra-Host-FC-boot-04

Selected Volumes 1 selected Clear all

- VM-Infra-Boot-iSCSI-01

LUN 1

Cancel Connect



LUN ID 1 should be used for the boot .

8. Select the volume that has been provisioned for the host, set the LUN ID for the volume, click the + next to the volume, and select Confirm to proceed. Repeat the steps for connecting volumes for each of the host/volume pairs configured.

Create Infra and Swap Datastores

To create datastore volumes for the ESXi Cluster, follow these steps in the Pure Storage Web Portal:

1. Select Storage > Volumes.
2. Select the + icon in the Volumes Panel.
3. A pop-up will appear to create a volume on the FlashArray.

Create Volume

Pod or Volume Group: none

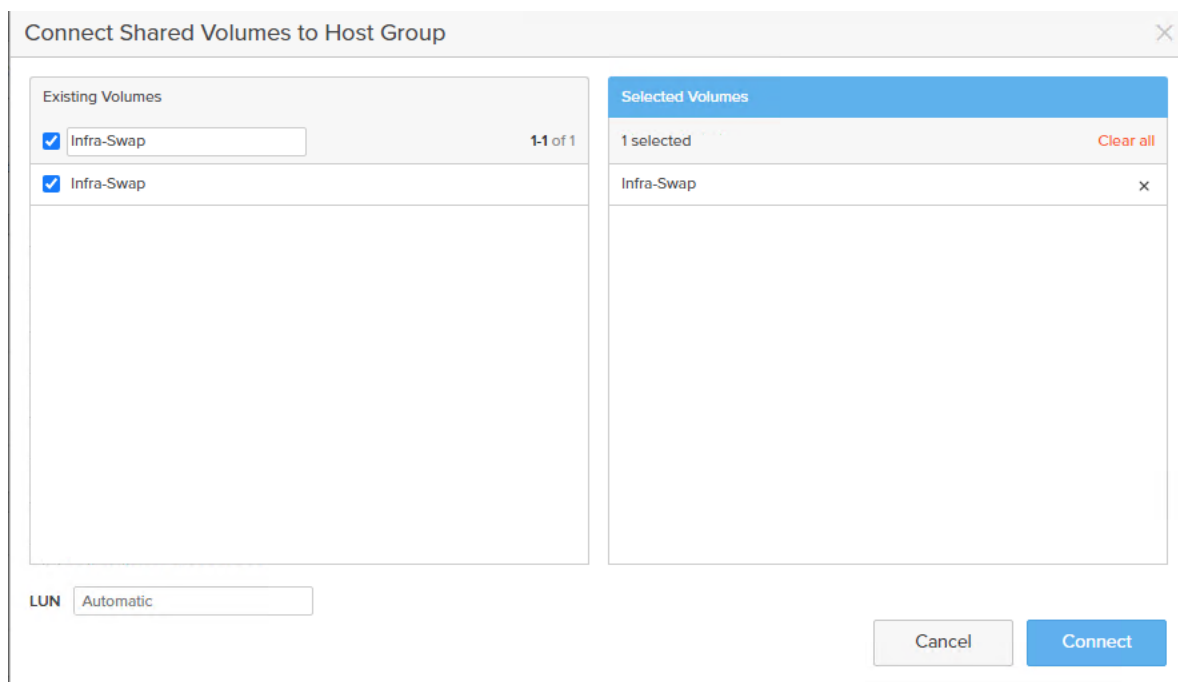
Name: Letters, Numbers, -

Provisioned Size: Positive numbers G

QoS Configuration (Optional) v

Create Multiple... Cancel Create

4. Fill in the Name and Provisioned Size.
5. Click Create to provision the volumes to be used as Infra datastore LUN.
6. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon pull-down within the Connected Volumes tab within that host group.
7. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.
8. Select the Infra datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.
9. Select Storage > Volumes.
10. Select the + icon in the Volumes Panel.
11. A pop-up will appear to create a volume on the FlashArray.
12. Fill in the Name and Provisioned Size.
13. Click Create to provision the volumes to be used as Swap datastore LUN.
14. Go back to the Hosts section under the Storage tab. Click ESXi cluster host group created earlier and select the gear icon pull-down within the Connected Volumes tab within that host group.
15. From the drop-down list of the gear icon, select Connect Volumes, and a pop-up will appear.



16. Select the Swap datastore volume that has been provisioned for the host group, leave the LUN ID for the volume to Automatic, click Connect.


VMware vSphere Configuration


Set Up VMkernel Ports and Virtual Switch on ESXi Host VM-Host-Infra-iSCSI-01


To add the iSCSI networking configuration on the first ESXi host, follow the steps at the end of section [Set Up VMkernel Ports and Virtual Switch](#). In this section, a single iSCSI Boot vSwitch is configured with two uplinks, one to UCS fabric A and the other to fabric B. The first VMkernel port will be mapped only to the fabric A uplink and the second one will be mapped to the fabric B uplink.

To setup VMkernel ports and virtual switches on ESXi hosts on VM-Host-Infra-iSCSI-01, follow these steps:


1. From the Host Client Navigator, click Networking.
2. In the center pane, choose the Virtual switches tab.
3. Highlight the iScsiBootvSwitch line.
4. Choose Edit settings.
5. Change the MTU to 9000.




 Edit standard virtual switch - iScsiBootvSwitch

 Add uplink

MTU	<input type="text" value="9000"/>
Uplink 1	<input type="text" value="vmnic4 - Up, 40000 mbps"/> 
▶ Link discovery	Click to expand
▶ Security	Click to expand
▶ NIC teaming	Click to expand
▶ Traffic shaping	Click to expand

6. Click Save to save the changes to iScsiBootvSwitch.
7. Click vmk1 entry.
8. Click Edit Settings.
9. From Port properties update the MTU value to 9000.

 Edit settings - vmk1

Port group	<input type="text" value="iScsiBootPG"/> 
MTU	<input type="text" value="9000"/>
IP version	<input type="text" value="IPv4 and IPv6"/> 
▶ IPv4 settings	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
▶ IPv6 settings	Click to expand
TCP/IP stack	<input type="text" value="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

10. Click the IPv4 Settings.
11. Change the IPv4 settings from the Cisco UCS Manager iSCSI-A-Pool assigned IP to one that is not in the IP block.

Edit settings - vmk1

Port group	iScsiBootPG
MTU	9000
IP version	IPv4 and IPv6
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	192.168.101.63
Subnet mask	255.255.255.0
▶ IPv6 settings	Click to expand
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

Save Cancel

12. Click OK to apply the changes.

Configure iSCSI B vSwitch and VMkernel

To configure the iSCSI vSwitch and VMkernel, follow these steps:

1. From the Host Client Navigator, click Networking.
2. In the center pane, choose the Virtual switches tab.
3. Click add standard virtual switch.
4. Name the switch iScsiBootvSwitch-B.
5. Change the MTU to 9000.
6. From the drop-down list select vmnic5 for Uplink 1.

Add standard virtual switch - iScsiBootvSwitch-B

Add uplink

vSwitch Name	<input type="text" value="iScsiBootvSwitch-B"/>
MTU	<input type="text" value="9000"/>
Uplink 1	<input type="text" value="vmnic5 - Up, 40000 mbps"/> ✕
▸ Link discovery	Click to expand
▸ Security	Click to expand

7. Choose Add to add iScsiBootvSwitch-B.
8. In the center pane, choose the VMkernel NICs tab.
9. Choose Add VMkernel NIC.
10. For New port group, enter iScsiBootPG-B.
11. For Virtual switch, use the pull-down to choose vSwitch1.
12. Change the MTU to 9000.
13. For IPv4 settings, choose Static.
14. Expand IPv4 Settings and enter a unique IP address in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B.

Port group	New port group
New port group	iScsiBootPG-B
Virtual switch	iScsiBootvSwitch
VLAN ID	0
MTU	9000
IP version	IPv4 only
▼ IPv4 settings	
Configuration	<input type="radio"/> DHCP <input checked="" type="radio"/> Static
Address	192.168.102.63
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

Create Cancel

15. Click Create to complete creating the VMkernel NIC.

16. In the center pane, choose the Port groups tab.

17. Highlight the iScsiBootPG line.

18. Choose Edit settings.

19. Change the Name to iScsiBootPG-A.

20. Click Save to complete editing the port group name.

21. Click Storage, then in the center pane choose the Adapters tab.

22. Click Software iSCSI to configure software iSCSI for the host.

23. In the Configure iSCSI window, under Dynamic targets, click Add dynamic target.

24. Choose to add address and enter the IP address of ct0.eth4 from Pure FlashArray//X R3. Press Return.

25. Repeat above steps to add the IP addresses for ct0.eth5, ct1.eth4 and ct1.eth5.

26. Click Save configuration.

27. Click Software iSCSI to configure software iSCSI for the host.

28. Verify that four static targets and four dynamic targets are listed for the host.

The screenshot shows the 'Configure iSCSI - vmhba64' window. It has several sections:

- iSCSI enabled:** Radio buttons for 'Disabled' and 'Enabled' (selected).
- Name & alias:** 'iqn.2010-11.com.flashstack.infra-ucs-host:3 (iscsi_vmk)'
- CHAP authentication:** 'Do not use CHAP' (dropdown)
- Mutual CHAP authentication:** 'Do not use CHAP' (dropdown)
- Advanced settings:** 'Click to expand'
- Network port bindings:** Includes 'Add port binding' and 'Remove port binding' buttons. Below are dropdowns for 'VMkernel NIC', 'Port group', and 'IPv4 address'. A message says 'No port bindings'.
- Static targets:** Includes 'Add static target', 'Remove static target', and 'Edit settings' buttons. A search bar is present. Below is a table with 4 rows and 3 columns: Target, Address, and Port.
- Dynamic targets:** Includes 'Add dynamic target', 'Remove dynamic target', and 'Edit settings' buttons. A search bar is present. Below is a table with 4 rows and 2 columns: Address and Port.

At the bottom right, there are 'Save configuration' and 'Cancel' buttons.

Target	Address	Port
iqn.2010-06.com.purestorage.flasharray.779962553908b056	192.168.101.146	3260
iqn.2010-06.com.purestorage.flasharray.779962553908b056	192.168.102.147	3260
iqn.2010-06.com.purestorage.flasharray.779962553908b056	192.168.101.147	3260
iqn.2010-06.com.purestorage.flasharray.779962553908b056	192.168.102.146	3260

Address	Port
192.168.101.146	3260
192.168.102.146	3260
192.168.101.147	3260
192.168.102.147	3260

29. Click Cancel to close the window.



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.

Add iSCSI Configuration to a VMware ESXi Host Added in vCenter

This section details the steps to add iSCSI configuration to an ESXi host added and configured in vCenter. This section assumes the host has been added to vCenter and the basic networking completed, and the time configuration and swap files added.


To add an iSCSI configuration to an ESXi host, follow these steps:

1. In the vSphere HTML5 Client, under Hosts and Clusters, choose the ESXi host.

-
2. In the center pane, click Configure. In the list under Networking, select Virtual switches.
 3. In the center pane, expand iScsiBootvSwitch. Click EDIT to edit settings for the vSwitch.
 4. Change the MTU to 9000 and click OK.
 5. Choose ... > Edit Settings to the right of iScsiBootPG. Change the Network label to iScsiBootPG-A and click OK.
 6. Choose ... > Edit Settings to the right of the VMkernel Port IP address. Change the MTU to 9000.
 7. Click IPv4 settings on the left. Change the IP address to a unique IP address in the Infra-iSCSI-A subnet but outside of the Cisco UCS iSCSI-IP-Pool-A.



It is recommended to enter a unique IP address for this VMkernel port to avoid any issues related to IP Pool reassignments.

8. Click OK.
9. In the upper right-hand corner, choose ADD NETWORKING to add another vSwitch.
10. Make sure VMkernel Network Adapter is selected and click NEXT.
11. Choose New standard switch and change the MTU to 9000. Click NEXT.
12. Choose  to add an adapter. Make sure vmnic5 is highlighted and click OK. vmnic5 should now be under Active adapters. Click NEXT.
13. Enter iScsiBootPG-B for the Network label, leave VLAN ID set to None (0), choose Custom - 9000 for MTU, and click NEXT.
14. Choose Use static IPv4 settings. Enter a unique IP address and netmask in the Infra-iSCSI-B subnet but outside of the Cisco UCS iSCSI-IP-Pool-B. Click NEXT.
15. Click FINISH to complete creating the vSwitch and the VMkernel port.
16. In the list under Storage, choose Storage Adapters.
17. Choose the iSCSI Software Adapter and below, choose the Dynamic Discovery tab.
18. Click Add.
19. Enter the IP address of the pure FlashArray storage controller's ct0.eth4 and click OK.
20. Repeat this process to add the IPs for ct0.eth5, ct1.eth4 and ct1.eth5.

-
21. Under Storage Adapters, click Rescan Adapter to rescan the iSCSI Software Adapter.
 22. Under Static Discovery, four static targets should now be listed.
 23. Under Paths, four paths should now be listed with two of the paths having the “Active (I/O)” Status.

Create a FlashStack ESXi Custom ISO using VMware vCenter

In this validation document, the [Cisco Custom ISO for UCS 4.1.3a](#) was used to install VMware ESXi. After this installation the Cisco UCS Tools and the Cisco VIC nfnic drivers had to be updated during the FlashStack deployment. vCenter 7.0 U2 or later can be used to produce a FlashStack custom ISO containing the updated UCS Tools and VIC drivers. This ISO can be used to install VMware ESXi 7.0 U2 without having to do any additional driver updates. To create the FlashStack ESXi custom ISO, follow these steps:



The Cisco Custom ISO for UCS 4.1.3a should also be used for Cisco UCS software release 4.1(2b) and VMware vSphere 7.0 U2.

1. Download the [Cisco Custom Offline Bundle](#) for UCS 4.1.3a. This file (VMware_ESXi_7.0.2_17867351_Custom_Cisco_4.1.3_a_Bundle.zip) can be used to produce the FlashStack ESXi 7.0 U2 CD ISO.
2. Download the following listed .zip files:
 - [UCS Tools Component for ESXi 7.0 1.2.1](#) (ucs-tool-esxi_1.1.5-1OEM.zip)
 - [VMware ESXi 7.0 nfnic 5.0.0.15 Driver for Cisco VIC Adapters](#) (Cisco-nfnic_5.0.0.15-1OEM.700.1.0.15843807_18697950.zip)
 - [nfnic Driver version 1.0.35.0](#) (Already part of install ISO, but including as a reference for updating to a newer nfnic version when available)
3. Log into the VMware vCenter HTML5 Client as administrator@vsphere.local.
4. Under Menu, choose Auto Deploy.
5. If you receive the message “Auto Deploy and Image Builder are disabled in this vCenter”, click ENABLE IMAGE BUILDER.



Auto Deploy and Image Builder are disabled in this vCenter.

To access full-featured auto deploy, enable both Image Builder and Auto Deploy.

To manage software depots only, enable Image Builder.

ENABLE AUTO DEPLOY AND IMAGE BUILDER

ENABLE IMAGE BUILDER

6. Click IMPORT to upload a software depot.
7. Name the depot Cisco Custom ESXi 7.0 for UCS 4.1(3a). Click BROWSE. Browse to the local location of the VMware-ESXi-7.0.0-16324942-Custom-Cisco-4.1.3a-Bundle.zip file downloaded above, highlight it, and click Open.

Import Software Depot ×

Name *

File * [BROWSE](#)

8. Click UPLOAD to upload the software depot.
9. Repeat steps 6-8 to add software depots for ucs-tool-esxi_2.1.5-1OEM, nfnic-5.0.0.15 and nenic-1.0.35.0.
10. Click NEW to add a custom software depot.
11. Choose Custom depot and name the custom depot FlashStack-ESXi-7.0U2.

Add Software Depot



Online depot

Name:

URL:

Custom depot

Name: *

FlashStack-ESXi-7.0U2

CANCEL

ADD

12. Click ADD to add the custom software depot.

13. From the drop-down list, choose the Cisco Custom ESXi-7.0 for UCS 4.1(3a) (ZIP) software depot. Make sure the Image Profiles tab is selected and then click the radio button to select the Cisco-UCS-Custom-ESXi-7-1632492_4.1.3-a image profile. Click CLONE to clone the image profile.

14. Name the clone FlashStack-ESXi-7.0U2. For Description, enter "Cisco Custom ISO ESXi 7.0U2 for UCS 4.1(3a) with nfnic-5.0.0.15, nenic-1.0.35.0 and ucs-tool-2.1.5". Choose FlashStack-ESXi-7.0U2 for Software depot.

Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

Name and details

Name * FlashStack-ESXi-7.0U2

Vendor * Cisco Systems, Inc.

Description

Cisco Custom ISO ESXi 7.0U2 for UCS 4.1(3a) with nfnic-5.0.0.12, nenic-1.0.35.0 and ucs-tool-2.1.5

Software depot * FlashStack-ESXi-7.0U2 ⓘ

CANCEL NEXT

15. Click NEXT.

16. Under Available software packages, check nfnic_5.0.0.15-1OEM.700.1.0.15843807 and uncheck 4.0.0.65-1OEM.670.0.0.8169922, make sure 1.0.35.0-1OEM.670.0.0.8169922 is checked and 1.0.33.0-1vmw.702.0.0.17867351 is unchecked. Uncheck 1.1.6-1OEM and check 1.2.1-1OEM. Leave the remaining selections unchanged.

Clone Image Profile

1 Name and details

2 Select software packages

3 Ready to complete

Select software packages



Acceptance level

Partner supported ▼

<input type="checkbox"/>	Name ▼	Version ▼	Acceptance Level ▼	Vendor
<input checked="" type="checkbox"/>	lsuv2-smartpqiv2...	1.0.0-6vmw.702.0.0.17867351	VMware certified	VMwar
<input checked="" type="checkbox"/>	mtip32xx-native	3.9.8-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	native-misc-drive...	7.0.2-0.0.17867351	VMware certified	VMwar
<input checked="" type="checkbox"/>	ne1000	0.8.4-11vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	nenic	1.0.35.0-1OEM.670.0.0.8169922	VMware certified	Cisco
<input type="checkbox"/>	nenic	1.0.33.0-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	nenic-ens	1.0.4.0-1OEM.700.1.0.15843807	VMware certified	Cisco
<input checked="" type="checkbox"/>	nfnic	5.0.0.12-1OEM.700.1.0.15843807	VMware certified	Cisco
<input type="checkbox"/>	nfnic	4.0.0.63-1vmw.702.0.0.17867351	VMware certified	VMW
<input type="checkbox"/>	nfnic	5.0.0.11-1OEM.700.1.0.15843807	VMware certified	Cisco
<input type="checkbox"/>	nfnic	4.0.0.65-1OEM.670.0.0.8169922	VMware certified	Cisco
<input checked="" type="checkbox"/>	nhpsa	70.0051.0.100-2vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	nmlx4-core	3.19.16.8-2vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	nmlx4-en	3.19.16.8-2vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	nmlx4-rdma	3.19.16.8-2vmw.702.0.0.17867351	VMware certified	VMW

80 selected of 96 items

CANCEL

BACK

NEXT

Clone Image Profile

- 1 Name and details
- 2 Select software packages
- 3 Ready to complete

Select software packages ✕

Acceptance level Partner supported ▼

<input type="checkbox"/>	Name	Version	Acceptance Level	Vendor
<input checked="" type="checkbox"/>	sfvmk	2.4.0.2010-4vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	smartpqi	70.4000.0.100-6vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	tools-light	11.2.5.17337674-17867351	VMware certified	VMware
<input type="checkbox"/>	ucs-tool-esxi	1.1.6-1OEM	Partner supported	CIS
<input checked="" type="checkbox"/>	ucs-tool-esxi	1.2.1-1OEM	Partner supported	CIS
<input checked="" type="checkbox"/>	vdfs	7.0.2-0.0.17867351	VMware certified	VMware
<input checked="" type="checkbox"/>	vmkata	0.1-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	vmkfcoc	1.0.0.2-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	vmkusb	0.1-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	vmw-ahci	2.0.9-1vmw.702.0.0.17867351	VMware certified	VMW
<input checked="" type="checkbox"/>	vmware-esx-esx...	1.2.0.42-1vmw.702.0.0.17867351	VMware certified	VMware
<input checked="" type="checkbox"/>	vsan	7.0.2-0.0.17867351	VMware certified	VMware
<input checked="" type="checkbox"/>	vsanhealth	7.0.2-0.0.17867351	VMware certified	VMware

80 selected of 95 items

CANCEL
BACK
NEXT

17. Click NEXT.

18. Click FINISH.

19. From the Software Depot drop-down list, choose the FlashStack-ESXi-7.0U2 (Custom) software depot. Under Image Profiles choose the FlashStack-ESXi-7.0U2 image profile. Click EXPORT to export an image profile. The ISO should be highlighted.

Software Depots
Deploy Rules
Deployed Hosts
Discovered Hosts
Script Bundles
Configure

Software Depot
FlashStack-ESXi-7.0U2 (Custom) ▼
REMOVE
NEW
IMPORT

Image Profiles
Software Packages

NEW IMAGE PROFILE
VIEW SOFTWARE PACKAGES
DELETE
...

Name	Acceptance Level	Vendor	Last Modified	Description	# Software Packages	Download Image Profiles
<input checked="" type="checkbox"/> FlashStack-ESXi-7.0U2	Partner supported	Cisco Systems, Inc.	09/11/2021, 07:13 PM	Cisco Custom ISO ESXi ...	80	EXPORT

20. Click OK to generate a bootable ESXi installable image.

Export Image Profile | FlashStack-ESXi-7.0U2 ×

Generate an image profile and download it from the "Download Image Profiles" column of the selected image profile.

- ISO - Generate a bootable ISO image from the image profile.
 - Do not include an installer on the ISO.
- ZIP - Generate a ZIP archive containing the software packages in the image profile.
- Skip acceptance level checking.

CANCEL

OK

21. Once the Image profile export completes, click **DOWNLOAD** to download the ISO.
22. Once downloaded, you can rename the ISO to a more descriptive name.
23. Optionally, generate the ZIP archive to generate an offline bundle for the FlashStack image using ... > Export.

FlashStack Backups

Cisco UCS Backup

Automated backup of the UCS domain is important for recovery of the UCS Domain from issues ranging catastrophic failure to human error. There is a native backup solution within Cisco UCS that allows local or remote backup using FTP/TFTP/SCP/SFTP as options.

Backups created can be a binary file containing the Full State, which can be used for a restore to the original or a replacement pair of fabric interconnects. Alternately create the XML configuration file consisting of All configurations, just System configurations, or just Logical configurations of the UCS Domain. For scheduled backups, options will be Full State or All Configuration, backup of just the System or Logical configurations can be manually initiated.

To configure the backup, using the Cisco UCS Manager GUI, follow these steps:

1. Choose Admin within the Navigation pane and choose All.
2. Click the Policy Backup & Export tab within All.
3. For a Full State Backup, All Configuration Backup, or both, specify the following:
 - a. Hostname: <IP or FQDN of host that will receive the backup>
 - b. Protocol: [FTP/TFTP/SCP/SFTP]

- c. User: <account on host to authenticate>
- d. Password: <password for account on host>
- e. Remote File: <full path and filename prefix for backup file>



Admin State must be Enabled to fill in the Remote File field.

- f. Admin State: <choose Enable to activate the schedule on save, Disable to disable schedule on Save>
- g. Schedule: [Daily/Weekly/Bi Weekly]

All

General Policy Backup & Export

Full State Backup Policy

Hostname :

Protocol : FTP TFTP SCP SFTP

User :

Password :

Remote File :

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : 0

Description : Database Backup Policy

All Configuration Backup Policy

Hostname : fs-ftp.flashstack.com

Protocol : FTP TFTP SCP SFTP

User : admin

Password :

Remote File : /var/www/html/software/configs/bb08-6454/bb08-

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : 0

Description : Configuration Export Policy

Backup/Export Config Reminder

Admin State : Disable Enable

Remind me after(Days) : 30

4. Click Save Changes to create the Policy.

Cisco Nexus and MDS Backups

The configuration of the Cisco Nexus 9000 and Cisco MDS 9132T switches can be backed up manually at any time with the copy command, but automated backups can be put in place with the NX-OS feature scheduler. An example of setting up automated configuration backups of one of the FlashStack 93180YC-FX switches is shown below:

```
conf t
feature scheduler
scheduler logfile size 1024
scheduler job name backup-cfg
copy running-config tftp://<server-ip>/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
exit
scheduler schedule name daily
job name backup-cfg
time daily 2:00
end
```



On the Cisco MDS 9132T, remove “vrf management” from the copy command.

Show the job that has been setup:

```
show scheduler job
Job Name: backup-cfg
-----
copy running-config tftp://10.1.164.150/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
=====

show scheduler schedule
Schedule Name      : daily
-----
User Name         : admin
Schedule Type     : Run every day at 2 Hrs 0 Mins
Last Execution Time : Yet to be executed
-----
Job Name          Last Execution Status
-----
backup-cfg       -NA-
=====
```

The documentation for the feature scheduler can be found here:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01010.html

VMware VCSA Backup

Basic scheduled backup of the vCenter Server Appliance is available within the native capabilities of the VCSA. To create a scheduled backup, follow these steps:

1. Connect to the VCSA Console at <https://<VCSA IP>:5480> as root.
2. Click Backup in the list to open up the Backup Appliance Dialogue.

3. To the right of Backup Schedule, click CONFIGURE.

4. Specify the following:

- a. The Backup location with the protocol to use [FTPS, HTTPS, SFTP, FTP, NFS, SMB, HTTP]
- b. The User name and password.
- c. The Number of backups to retain.

Create Backup Schedule

Backup location ⓘ	<input type="text" value="http://10.164.127/var/www/html/software/"/>	
Backup server credentials	User name	<input type="text" value="root"/>
	Password	<input type="password" value="....."/>
Schedule ⓘ	Daily ▾ 11 : 59 P.M. Etc/UTC	
Encrypt backup (optional)	Encryption Password	<input type="password"/>
	Confirm Password	<input type="password"/>
DB Health Check ⓘ	<input checked="" type="checkbox"/> Enabled	
Number of backups to retain	<input checked="" type="radio"/> Retain all backups	
	<input type="radio"/> Retain last <input type="text" value="0"/> backups	
Data	<input checked="" type="checkbox"/> Stats, Events, and Tasks	80 MB
	<input checked="" type="checkbox"/> Inventory and configuration	198 MB
	Total size (compressed) 278 MB	
<input type="button" value="CANCEL"/> <input type="button" value="CREATE"/>		

5. Click CREATE.

Backup Schedule		EDIT	DISABLE	DELETE
▼ Status	Enabled			
Schedule	Daily , 11:59 P.M. Etc/UTC			
Backup Location	http://10.164.127/var/www/html/Software			
Backup data	<ul style="list-style-type: none">Stats, Events, and TasksInventory and configuration			
Number of backups to retain	Retain all backups			

6. The Backup Schedule should now show a Status of Enabled.

-
7. Restoration can be initiated with the backed-up files using the Restore function of the VCSA 7.0 U2 Installer.

FlashStack Automated Deployment with Ansible


If using the published Ansible playbooks to configure the FlashStack infrastructure, complete this section of the document.

Ansible Automation Workflow and Solution Deployment

This FlashStack with vSphere 7.0 U2 and Cisco UCS M6 solution uses a management workstation (control machine) to run Ansible playbooks to configure Cisco Nexus, Cisco UCS, Pure Storage and Install VMware Cluster.

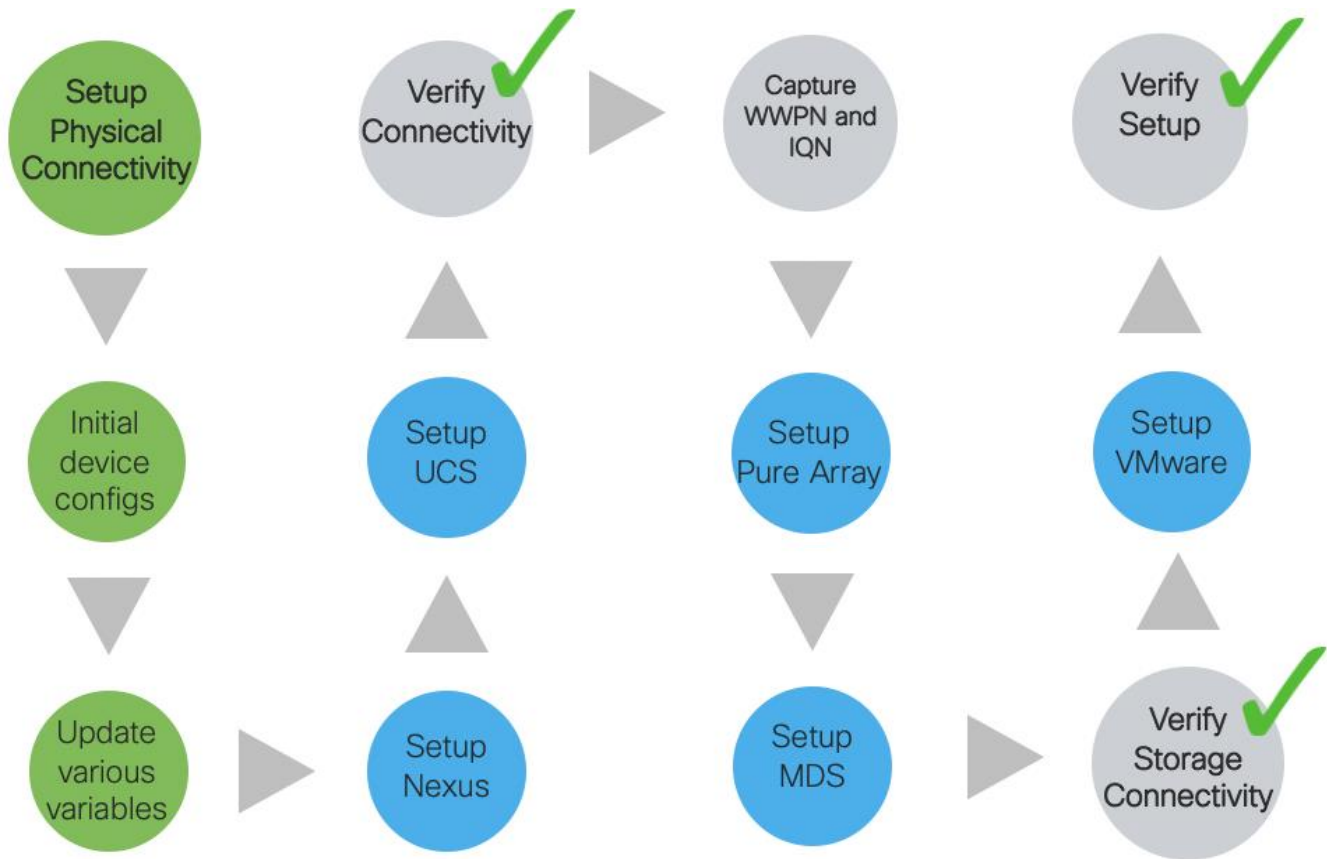
[Figure 6](#) illustrates the FlashStack with vSphere 7.0 U2 and Cisco UCS solution implementation workflow which is explained in the following sections. The FlashStack Ansible based automation is depicted in the following [Figure 7](#).

Figure 6. High-level FlashStack Automation



The FlashStack Automated deployment workflow is depicted in [Figure 7](#).

Figure 7. FlashStack Automated Deployment Workflow



Prerequisites

Setting up the solution begins with a management workstation that has access to the internet and has a working installation of Ansible. The management workstation runs a variant of Linux or MacOS for ease of use with these command-line-based tools. Instructions for installing the workstation are not included in this document, but the basic installation and configuration of Ansible is explained. The following is a list of prerequisites:

- [Getting Started with Red Hat Ansible](#)

To use the Ansible playbooks demonstrated in this document, the management workstation must also have a working installation of Git and access to the Cisco DevNet public GitHub repository. The Ansible playbooks used in this document are cloned from the public repositories, located at the following links:

- Cisco DevNet: <https://developer.cisco.com/codeexchange/github/repo/ucs-compute-solutions/FlashStack-laC-UCSM6>
- GitHub repository for FlashStack infrastructure setup: <https://github.com/ucs-compute-solutions/FlashStack-laC-UCSM6.git>

- The Cisco Nexus Switches, Pure Storage and Cisco UCS must be physically racked, cabled, powered, and configured with the management IP addresses before the Ansible-based installation procedure can begin as shown in the cabling diagram (Figure 6.). If necessary, upgrade the Nexus Switches to release 9.3(7) and the UCS System to 4.2(1f) with the default firmware packages for both blades and rack servers set to 4.2(1f).
- Before running each Ansible Playbook to setup the Network, Storage, UCS and VMware, various variables must be updated based on the customers environment and specific implementation with values such as the VLANs, pools & ports on Cisco UCS, IP addresses for iSCSI interfaces and values needed for the ESXi installation and configuration.



Day 2 Configuration tasks such as adding datastores or ESXi servers have been performed manually or with Cisco Intersight Cloud Orchestrator (ICO) and the information has been provided in the respective sections of this document.

Prepare Management Workstation (Control Machine)

In this section, the installation steps are performed on the CentOS management host to prepare the host for solution deployment to support the automation of Cisco UCS, Cisco Nexus, Pure Storage and VMware installation using Ansible Playbooks.

To prepare the management workstation, follow these steps:

1. Install the EPEL repository on the management host.

```
[root@FSV-Automation ~]# yum install epel-release
```

2. Install Ansible engine.

```
[root@FSV-Automation ~]# yum install ansible
```

3. Verify the Ansible version to make sure it's at least release 2.9.

```
[root@FS-Automation tasks]# ansible --version
ansible 2.10.7
  config file = None
  configured module search path = ['/root/.ansible/plugins/modules',
  '/usr/share/ansible/plugins/modules']
  ansible python module location = /usr/local/lib/python3.6/site-packages/ansible
  executable location = /usr/local/bin/ansible
  python version = 3.6.8 (default, Aug 24 2020, 17:57:11) [GCC 8.3.1 20191121 (Red
  Hat 8.3.1-5)]
```

4. Install **pip** the package installer for Python.

```
[root@FSV-Automation ~]# yum install python-pip
```

5. Install the Cisco UCS SDK.

```
[root@FSV-Automation ~]# pip3 install ucsm sdk
```

6. Install the **paramiko** package for Cisco Nexus automation.

```
[root@FSV-Automation ~]# pip3 install paramiko
```

7. SSH into each of the Cisco Nexus and Cisco MDS switches using Ansible so that the SSH keys are cached.

```
[root@FSV-Automation ~]# ssh admin@10.1.164.61
The authenticity of host '10.1.164.61 (10.1.164.61)' can't be established.
RSA key fingerprint is SHA256:mtomJluZVkcITgSLhVygocSnojlyPPDPmcJLQX2dfu4.
RSA key fingerprint is MD5:b4:e3:86:97:99:58:df:0d:5d:20:b2:5b:d5:69:aa:23.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.1.164.61' (RSA) to the list of known hosts.
User Access Verification
Password:
```

8. Install the Pure Storage SDK.

```
[root@FSV-Automation ~]# pip3 install purestorage
```

9. Install ansible-galaxy collections for Cisco UCS, Cisco Nexus/MDS switches and Pure Storage Array as follows:

```
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.nxos
[root@FSV-Automation ~]# ansible-galaxy collection install cisco.ucs
[root@FSV-Automation ~]# ansible-galaxy collection install purestorage.flasharray
```



We validated the Ansible automation with both python 2.7.5 and python 3.6 as the python interpreter for Ansible.

Clone GitHub Collection

You will use GitHub repos from two public locations; the first step in the process is to clone the GitHub collection named FlashStack-IaC-UCSM6 (<https://github.com/ucs-compute-solutions/FlashStack-IaC-UCSM6>) to the new empty folders on the management workstation. Cloning the collections creates a local copy, which is then used to run the playbooks that have been created for this solution. To clone the GitHub collection, follow these steps:

1. From the management workstation, create a new folder for the project. The GitHub collection will be cloned in a new folder inside this one, named ucsm6.
2. Open a command-line or console interface on the management workstation and change directories to the new folder just created.
3. Change directories to the new folder named ucsm6.
4. Clone the GitHub collection using the following command:

```
git clone https://github.com/ucs-compute-solutions/FlashStack-IaC-UCSM6.git
```

5. Change directories to the folder named **FlashStack-IaC-UCSM6**.

FlashStack Deployment using Playbooks

The following sections explain the installation and configuration of all the infrastructure layers with in FlashStack. The Ansible Playbook tree structure is shown below with the directory structure and various roles and tasks:




```
├── configure_mds_interfaces.yml
├── configure_mds_ntp.yml
├── configure_mds_vsans.yml
├── configure_mds_zoneset.yml
├── configure_mds_zones.yml
├── main.yml
├── save_mds_config.yml
├── NEXUSconfig
│   ├── defaults
│   │   └── main.yml
│   └── tasks
│       ├── configure_default_gw.yml
│       ├── configure_nxos_features.yml
│       ├── configure_nxos_global_settings.yml
│       ├── configure_nxos_ntp.yml
│       ├── configure_nxos_vlans.yml
│       ├── configure_nxos_vpc.yml
│       ├── initiate_nxos_config_backup.yml
│       ├── main.yml
│       ├── save_nxos_config.yml
│       └── set_nxos_interfaces.yml
├── PUREconfig
│   ├── meta
│   │   └── main.yml
│   ├── tasks
│   │   ├── ConfigPure.yml
│   │   ├── main.yml
│   │   └── SetupPure.yml
│   └── vars
│       ├── main.yml
│       └── main.yml.true
├── UCSadmin
│   ├── defaults
│   │   └── main.yml
│   └── tasks
│       ├── add_ucs_alternate_user.yml
│       ├── configure_dns.yml
│       ├── configure_ntp_server.yml
│       ├── configure_ucs_timezone.yml
│       ├── create_ucs_org.yml
│       └── main.yml
├── UCSequipment
│   ├── defaults
│   │   └── main.yml
│   └── tasks
│       ├── configure_ucs_auto_discovery.yml
│       ├── configure_ucs_chassis_discovery.yml
│       ├── configure_ucs_info_policy.yml
│       ├── configure_ucs_server_port.yml
│       ├── configure_ucs_udld_policy.yml
│       ├── configure_ucs_uplink_PC.yml
│       ├── disable_unused_FC_ports.yml
│       └── main.yml
├── UCSlan
│   ├── defaults
│   │   └── main.yml
│   └── tasks
│       ├── configure_ucs_adapter_policy.yml
│       ├── configure_ucs_fc_LCP.yml
│       ├── configure_ucs_iscsi_ip_pools.yml
│       ├── configure_ucs_iscsi_LCP.yml
│       ├── configure_ucs_iscsi_vnic_templates.yml
│       ├── configure_ucs_mac_pools.yml
│       ├── configure_ucs_mgmt_ip_pool.yml
│       ├── configure_ucs_NCP.yml
│       ├── configure_ucs_system_qos.yml
│       ├── configure_ucs_vlans.yml
│       ├── configure_ucs_vnic_templates.yml
│       └── main.yml
```



The following information must be modified based on your environment and more information needs to be modified specific to each device automation. This is explained later in the document.

- inventory - contains the variables such as device IP addresses and authentication details:
 - group_vars/all.yml - contains the VLAN ids required for the solution deployment, update this file based on your environment.



FlashStack Network Configuration

Before the Ansible Nexus switch setup playbook can be run, the Nexus switches must be brought up with a management IP address. The following procedures describe this basic configuration of the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 9000 9.3(7), the Cisco suggested Nexus switch release at the time of this validation.



Make sure the FlashStack cabling and initial configuration has been completed on the Cisco Nexus switches. The Nexus automation includes the VPC connectivity between the Cisco UCS FI's and the Nexus 93180YC-FC switches using 25G ports, but 100G ports can be leveraged to reduce the number of cables and when used the variable parameters must be changed accordingly.

The following information has to be modified based on your specific environment, before running the Nexus Automation Playbook:


1. Add Nexus switch ssh keys to `/root/.ssh/known_hosts`. Adjust `known_hosts` as necessary if errors occur.

```
ssh admin@<nexus-A-mgmt0-ip>
exit
ssh admin@<nexus-B-mgmt0-ip>
exit
```


2. Edit the following variable files to ensure proper Nexus variables are entered:

- `ucsm6/FlashStack-laC-UCSM6/inventory`
- `ucsm6/FlashStack-laC-UCSM6/group_vars/all.yml`
- `ucsm6/FlashStack-laC-UCSM6/host_vars/n9kA.yml`
- `ucsm6/FlashStack-laC-UCSM6/host_vars/n9kB.yml`
- `ucsm6/FlashStack-laC-UCSM6/roles/NEXUSconfig/defaults/main.yml`

3. Switch Interface details in the following files if using different ports.



4. vPC domain id, Port Channel details and Nexus features in the following files if using different port channel ids or features.



5. From /root/ucsm6/FlashStack-IaC-UCSM6, run the Setup_Nexus.yml Ansible playbook.

```
ansible-playbook ./Setup_Nexus.yml -i inventory
```

6. Once the Ansible playbook has been run on both switches, it is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summertime, please see Cisco Nexus 9000 Series NX-OS Fundamentals Configuration Guide, Release 9.3(x). Sample clock commands for the United States Eastern timezone are:

```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

7. ssh into each switch and execute the following commands.

```
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>
```

8. Login into the Nexus switches and verify the configuration has been completed as desired before proceeding with the next section to configure Pure Storage and Cisco UCS.


FlashStack initial Storage Configuration



Skip this section if the initial configuration of FlashArray is performed by a Pure Implementation engineer.

To configure the FlashStack storage, follow these steps:

1. Update the following information as required based on your environment before running the MDS and UCS Automation Playbook.
2. There are three variables defined in the `group_vars/all.yml` file as follows, comment out the lines based on what configuration is required:
 - `initial_fa_config`: "yes" - required to perform the initial configuration of FlashArray
 - `configure_iscsi`: "yes" - required to configure the iSCSI ports on the FlashArray
 - `configure_fc`: "yes" - comment this line during initial configuration of FlashArray, it needs to be enabled or disabled when configuring the storage on FlashArray at a later point in time.
3. Change directory to `"/root/ucsm6/FlashStack-IaC-UCSM6/roles/PUREconfig/vars"` on your management host.
4. Following details need to be updated in the **main.yml** file:



Change the values in the above-mentioned files with caution, only change the information that is required. All the other files can be left to defaults, modify them only if you want to go with a different naming convention or if you do not have the identical hardware discussed in this design.

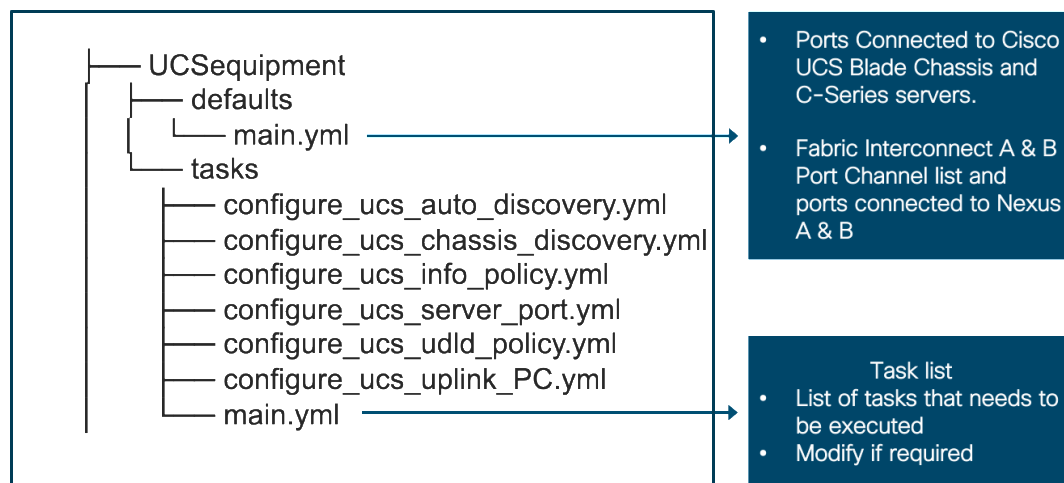
5. When the information has been updated in the respective files, run the UCS Ansible playbook to setup the initial configuration of FlashStack, this can be skipped if this is already completed :

```
[root@FSV-Automation FlashStack-IaC-UCSM6]# ansible-playbook ./Setup_Pure.yml
```

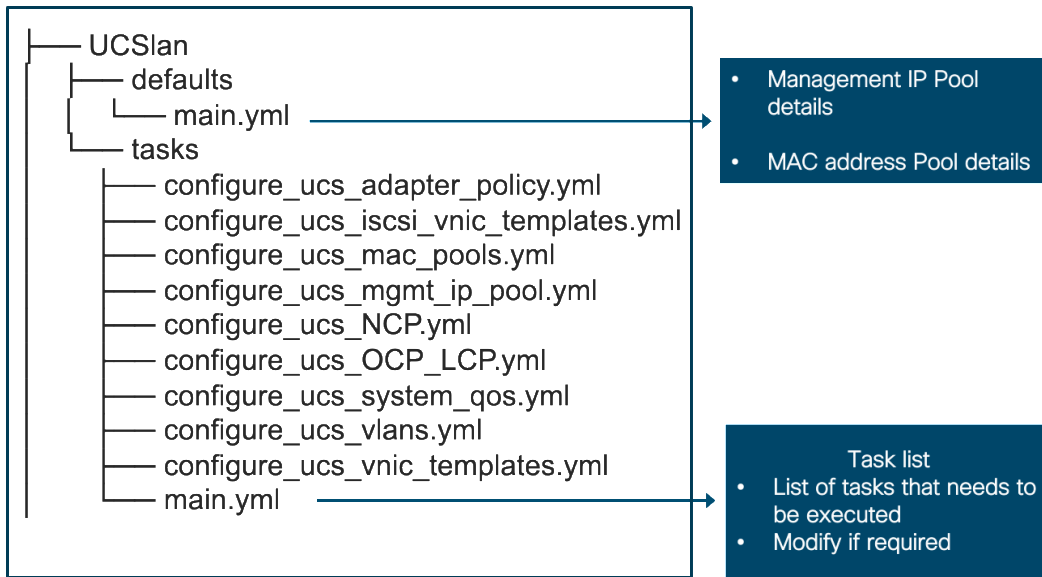
FlashStack UCS Compute Configuration

To configure the FlashStack UCS compute, follow these steps. Update the following information as required based on your environment before running the UCS Automation Playbook.

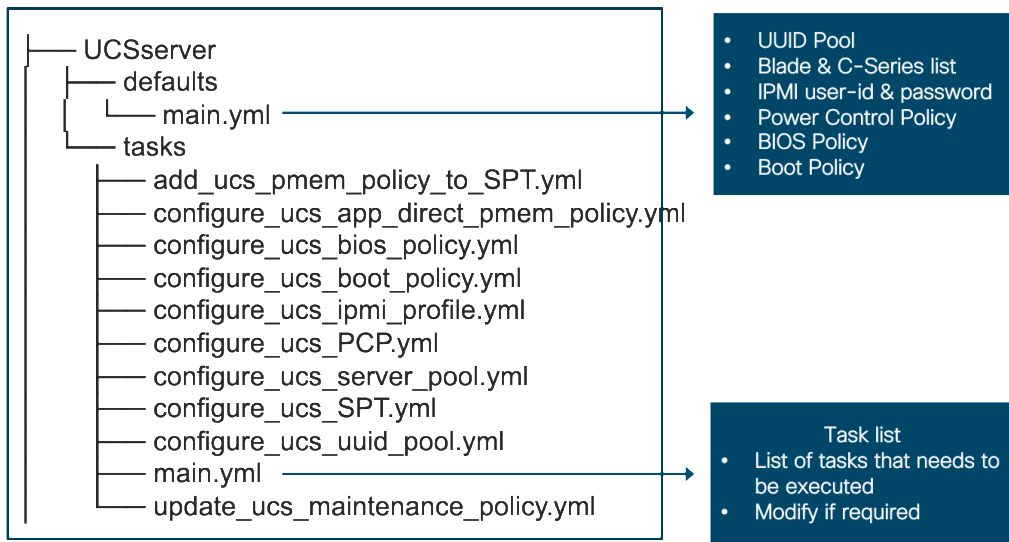
1. The following procedure can be used to configure the Cisco UCS from the Ansible management workstation.
2. Edit the following variable files to ensure proper Nexus variables are entered:
 - ucs6/FlashStack-IaC-UCSM6/inventory
 - ucs6/FlashStack-IaC-UCSM6/group_vars/all.yml
 - ucs6/FlashStack-IaC-UCSM6/group_vars/ucs.yml
 - ucs6/FlashStack-IaC-UCSM6/roles/UCSequipment/defaults/main.yml
 - ucs6/FlashStack-IaC-UCSM6/roles/UCSadmin/defaults/main.yml
 - ucs6/FlashStack-IaC-UCSM6/roles/UCSln/defaults/main.yml
 - ucs6/FlashStack-IaC-UCSM6/roles/UCSsan/defaults/main.yml
 - ucs6/FlashStack-IaC-UCSM6/UCSserver/defaults/main.yml
3. The port details and tasks to be included for **UCSequipment** configuration role if different from the defaults.



4. Management and MAC address pool details for **UCSln** configuration role.



5. UUID pool, UCS servers list and IPMI details for **UCSServer** configuration role.



Change the values in the mentioned files with caution; only change the information that is required. All the other files can be left to defaults, modify them only if you want to go with a different naming convention or if you do not have the identical hardware discussed in this design.

6. Once the information has been updated in the respective files, run the UCS Ansible playbook:

```
[root@FSV-Automation FlashStack-IaC-UCSM6]# ansible-playbook ./Setup_UCS.yml
```

7. Login into Cisco UCS Manager and verify the configuration has been completed as desired.

-
- The cloning process used in **Error! Reference source not found.** below can be used to create other Service Profile templates that can be modified to accommodate additional features such as Intel Datacenter Persistent Memory (DCPMem) in Memory or App-Direct Mode.

Create Service Profiles

To create service profiles from the service profile template within the FlashStack-VSI Organization, follow these steps:



Use the appropriate service profile template based on the storage protocol being used to setup FlashStack, which is either iSCSI or FC. Below procedure is using a FCP template as an example.

- Connect to UCS Manager and click Servers.
- Choose Service Profile Templates > root > Sub-Organizations > FlashStack > Service Template VM-Host-Infra-FCP-vM.
- Right-click VM-Host-Infra-FCP-vM and choose Create Service Profiles from Template.
- Enter VM-Host-Infra-FCP-0 as the service profile prefix.
- Enter 1 as “Name Suffix Starting Number.”
- Enter 3 as the “Number of Instances.”

Create Service Profiles From Template ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

OK

Cancel

- Click OK to create the service profiles.
- Click OK in the confirmation message.
- When VMware ESXi 7.0 U2 has been installed on the hosts, the host Service Profiles can be bound to the VM-Host-Infra-FCP Service Profile Template to remove the vMedia Mapping from the host.

FlashStack Cisco MDS Ansible Switch Configuration

The following procedure can be used to configure the Cisco MDS switches from the management workstation.

To configure the MDS Switches, follow these steps:



This section can be skipped if iSCSI is being setup.

1. Add MDS switch ssh keys to `/root/.ssh/known_hosts`. Adjust `known_hosts` as necessary if errors occur.

```
ssh admin@<mds-A-mgmt0-ip>
exit
ssh admin@<mds-B-mgmt0-ip>
exit
```

2. Edit the following variable files to ensure proper MDS variables are entered:
 - `ucsm6/FlashStack-laC-UCSM6/inventory`
 - `ucsm6/FlashStack-laC-UCSM6/group_vars/all.yml`
 - `ucsm6/FlashStack-laC-UCSM6/host_vars/mdsA.yml`
 - `ucsm6/FlashStack-laC-UCSM6/host_vars/mdsB.yml`
 - `ucsm6/FlashStack-laC-UCSM6/roles/MDSconfig/defaults/main.yml`
3. There are two variables defined in the `group_vars/all.yml` file:
4. `initial_fa_config`: "yes" - required to perform the initial configuration of FlashArray
5. `configure_iscsi`: "yes" - required to configure the iSCSI ports on the FlashArray
6. `configure_fc`: "yes" - comment this line during initial configuration of FlashArray, it needs to be enabled or disabled when configuring the storage on FlashArray at a later point in time.
7. From `/root/ucsm6/FlashStack-laC-UCSM6`, run the `Setup_MDS.yml` Ansible playbook.

```
ansible-playbook ./Setup_MDS.yml -i inventory
```

8. Once the Ansible playbook has been run and configured both switches, it is important to configure the local time so that logging time alignment and any backup schedules are correct. For more information on configuring the timezone and daylight savings time or summertime, please see [Cisco MDS 9000 Series Fundamentals Configuration Guide, Release 8.x](#). Sample clock commands for the United States Eastern timezone are:


```
clock timezone EST -5 0
clock summer-time EDT 2 Sunday March 02:00 1 Sunday November 02:00 60
```

```
ssh into each switch and execute the following commands
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>
```

FlashStack Storage Configuration

To configure the FlashStack storage, follow these steps. Update the following information as required based on your environment before running the UCS Automation Playbook.

1. Change directory to “/root/ucsm6/FlashStack-IaC-UCSM6/roles/PUREconfig/vars” on your management host.
2. Following details need to be updated in the **main.yaml** file:



Change the values in the above-mentioned files with caution, only change the information that is required. All the other files can be left to defaults, modify them only if you want to go with a different naming convention or if you do not have the identical hardware discussed in this design.

There are three variables defined in the `group_vars/all.yml` file as follows, comment out the lines based on what configuration is required:

- `configure_fc`: “yes” - required to configure scsi-fc setup on the MDS.
 - `configure_fc-nvme`: “yes” - uncomment this variable if nvme-fc configuration is also required.
3. When the information has been updated in the respective files, run the UCS Ansible playbook:

```
[root@FSV-Automation FlashStack-IaC-UCSM6]# ansible-playbook ./Setup_Pure.yml
```

VMware vSphere 7.0 U2 Installation and Configuration

The following procedure can be used to configure the three VMware ESXi hosts from the management workstation.

4. Edit the following variable files to ensure proper Nexus variables are entered:

- ucsm6/FlashStack-1aC-UCSM6/inventory
- ucsm6/FlashStack-1aC-UCSM6/group_vars/all.yml
- ucsm6/FlashStack-1aC-UCSM6/roles/ESXihosts/defaults/main.yml
- ucsm6/FlashStack-1aC-UCSM6/roles/ESXilscsi/defaults/main.yml (If using iSCSI boot)

5. From /root/ucsm6/FlashStack-1aC-UCSM6, run the Setup_ESXi.yml Ansible playbook.

```
ansible-playbook ./Setup_ESXi.yml -i inventory
```

To complete the FC-NVMe configuration on the ESXi hosts, follow these steps:

The remaining steps in the VMware vSphere Client are manual steps that should be completed whether an Ansible configuration or manual configuration is being done.

1. Verify that the NVMe Fibre Channel Disk is mounted on each ESXi host. Under Hosts and Clusters select the ESXi host.
2. In the center pane, select Configure > Storage > Storage Devices. The NVMe Fibre Channel Disk should be listed under Storage Devices.
3. Select the NVMe Fibre Channel Disk, then select Paths underneath. Verify 4 paths have a status of Active (I/O).
4. Repeat steps 1-3 for all 3 hosts.

Name	L...	Type	Capacity	Datastore
PURE Fibre Channel Disk (naa.624a9370f6ebcc130e54c5cb0...	253	disk	2.00 TB	Intra-di...
NVMe Fibre Channel Disk (eu1.00f6ebcc130e54c524a937cb0...	759...	disk	1.00 TB	NVMe-t...
Local HGST Disk (naa.5000cca0850093e4)	0	disk	894.25 GB	Not Consum...

Runtime ...	Status	Target	Name	Preferred
vmhba4:C0:...	Active (L...	52:4a:93:75:f2:e3:d5:11 52:...	vmhba4:C0:T0:L75902	
vmhba4:C0:...	Active (L...	52:4a:93:75:f2:e3:d5:01 52:...	vmhba4:C0:T1:L75902	
vmhba5:C0:...	Active (L...	52:4a:93:75:f2:e3:d5:03 52:...	vmhba5:C0:T0:L75902	
vmhba5:C0:...	Active (L...	52:4a:93:75:f2:e3:d5:13 52:...	vmhba5:C0:T1:L75902	

5. For any of the three hosts, right-click the host under Hosts and Clusters and select Storage > New Datastore. Leave VMFS selected and click NEXT.
6. Name the datastore and select the NVMe Fibre Channel Disk. Click NEXT.
7. Leave VMFS 6 selected and click NEXT.

8. Leave all Partition configuration values at the default values and click NEXT.
9. Review the information and click FINISH.
10. Select Storage and select the just-created NVMe datastore. In the center pane, select Hosts. Ensure all three hosts have the datastore mounted.

ESXi Host Multipathing Configuration

To configure the ESXi Host multipathing, follow these steps:

1. From the vCenter management GUI.
2. Go to Hosts and Clusters view.
3. Select a Host.
4. Click on the Configure tab.
5. Select Storage Devices.
6. Select an NVMe device.
7. Click Edit Multipathing.

Edit Multipathing Policies | eui.00f6ebcc130e54c524a937cb0001287f ×

Path selection policy LB-Latency ▼

Latency evaluation time 180000 ⓘ
The value must be between 10000 and 300000

Sampling I/Os per path 16 ⓘ
The value must be between 16 and 160

CANCEL SAVE

vCenter and Final ESXi Ansible Setup

The following procedure can be used to complete the configuration of the VMware vCenter and the three management ESXi hosts.

1. Edit the following variable files to ensure proper variables are entered:

- ucsm6/FlashStack-1aC-UCSM6/inventory
- ucsm6/FlashStack-1aC-UCSM6/group_vars/all.yml
- ucsm6/FlashStack-1aC-UCSM6/roles/ESXlpostvC/defaults/main.yml

2. From /root/ucsm6/FlashStack-1aC-UCSM6, run the Setup_vCenter.yml Ansible playbook.

```
ansible-playbook ./Setup_vCenter.yml -i inventory
```

About the Authors

Sreenivasa Edula, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Sreeni is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Engineering team focusing on converged and hyper-converged infrastructure solutions, prior to that he worked as a Solutions Architect at EMC Corporation. He has experience in Information Systems with expertise across Cisco Data Center technology portfolio, including DC architecture design, virtualization, compute, network, storage, and cloud computing.

Joe Houghes, Senior Solutions Architect, Pure Storage, Inc.

Joe is a Senior Solutions Architect in the Portfolio Solutions team within Pure Storage, focused on solutions on the FlashStack platform along with automation and integration. He has experience from over 15 years in Information Technology across various customer/vendor organizations with architecture and operations expertise covering compute, networking, storage, virtualization, business continuity and disaster recovery, along with cloud computing technologies, plus automation and integration across many applications and vendor platforms.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- John George, Technical Marketing Engineer, Cisco Systems, Inc.
- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Craig Waters, Technical Director, Pure Storage, Inc.
- Simon Dodsley, Principal Field Solutions Architect, Pure Storage, Inc.

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://community.cisco.com/t5/Network-Design/Cisco-Community) at <https://cs.co/en-cvds>.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)