



VersaStack with VMware vSphere 6.7, Cisco UCS 4th Generation Fabric, and IBM FS9100 NVMe-accelerated Storage

Deployment Guide for VersaStack with VMware vSphere 6.7 U2, Cisco UCS Manager 4.0(4), and IBM Spectrum Virtualize 8.2.1

Last Updated: October 8, 2019



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 NON-INFRINGEMENT 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.

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. (o8ogR)

© 2019 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary.....	8
Solution Overview	9
Introduction.....	9
Audience	9
Purpose of this Document.....	9
Solution Design	11
Architecture.....	11
Physical Topology	11
Software Revisions	12
Configuration Guidelines	13
Physical Infrastructure	14
Cisco UCS Connectivity to Nexus Switches	14
IBM FS9100 Connectivity to Nexus Switches	16
Cisco UCS connectivity to SAN Fabric.....	17
Network Configuration.....	19
Physical Connectivity	19
Cisco Nexus 9000 Initial Configuration Setup	19
Cisco Nexus A.....	19
Cisco Nexus B.....	20
Enable Appropriate Cisco Nexus 9000 Features and Settings.....	21
Enable Licenses.....	21
Set Global Configurations.....	22
Setup NTP (optional)	22
Cisco Nexus 9000 A and Cisco Nexus 9000 B.....	22
Add NTP Distribution Interface.....	22
Create VLANs for VersaStack IP Traffic	23
Cisco Nexus 9000 A and Cisco Nexus 9000 B.....	23
Configure Virtual Port Channel Domain.....	23
Cisco Nexus 9000 A	23
Cisco Nexus 9000 B	24
Configure Network Interfaces for the vPC Peer Links.....	25
Cisco Nexus 9000 A	25
Cisco Nexus 9000 B	26
Configure Network Interfaces to Cisco UCS Fabric Interconnects	27
Cisco Nexus 9000 A	27
Cisco Nexus 9000 B	29
Enable UDLD for Cisco UCS Interfaces	31
Cisco Nexus A and Cisco Nexus B.....	31
Configure Network Interfaces Connected to IBM FS9100 iSCSI Ports (iSCSI Deployment)	31
Management Uplink into Existing Network Infrastructure	34

Cisco Nexus 9000 A and B using Port Channel Example	34
Switch Testing Commands.....	35
Cisco MDS 9132T Configuration (FC Deployment).....	35
Physical Connectivity	35
VersaStack Cisco MDS Base Configuration	35
Cisco MDS 9132T Initial Configuration Setup	38
Add Second NTP server.....	38
Configure Individual Ports	38
Create VSANs	42
Initial Storage Configuration.....	43
IBM FlashSystem 9100	43
IBM Service Support Representative (SSR) Configuration	45
Customer Configuration Setup Tasks via the GUI.....	53
System Dashboard, and Post-Initialization Setup Tasks.....	62
Create Storage Pools and Allocate Storage.....	64
IBM FS9100 iSCSI Configuration (iSCSI Deployment)	70
Modify Interface MTU	74
Cisco UCS Server Configuration.....	75
Cisco UCS Initial Configuration.....	75
Cisco UCS 6454 A	75
Cisco UCS 6454 B	76
Cisco UCS Setup	76
Log into Cisco UCS Manager	76
Upgrade Cisco UCS Manager Software to Version 4.0(4c).....	76
Anonymous Reporting	76
Configure Cisco UCS Call Home	77
Add a Block of Management IP Addresses for KVM Access.....	78
Synchronize Cisco UCS to NTP.....	78
Add Additional DNS Server(s)	79
Add an Additional Administrator User.....	79
Enable Port Auto-Discovery Policy	80
Enable Info Policy for Neighbor Discovery.....	81
Edit Chassis Discovery Policy	81
Enable Server and Uplink Ports	82
Acknowledge Cisco UCS Chassis and FEX.....	83
Enable Fibre Channel Ports (FC Deployment)	84
Create VSAN for the Fibre Channel Interfaces.....	84
Create Port Channels for the Fibre Channel Interfaces	86
Create Port Channels for Ethernet Uplinks	88
Add UDLD to Uplink Port Channels	89
Create MAC Address Pools.....	91
Create UUID Suffix Pool.....	93

Create Server Pool	94
Create a WWNN Address Pool for FC based Storage Access	95
Create a WWPN Address Pools for FC Based Storage Access	96
Create IQN Pools for iSCSI Boot and LUN Access (iSCSI Deployment)	98
Create IP Pools for iSCSI Boot and LUN Access (iSCSI Deployment)	99
Create VLANs	101
Create Host Firmware Package	102
Set Jumbo Frames in Cisco UCS Fabric	103
Create Local Disk Configuration Policy	104
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	105
Create Power Control Policy	106
Create Server Pool Qualification Policy (Optional)	107
Create Server BIOS Policy	108
Update Default Maintenance Policy	111
Create vNIC/vHBA Placement Policy	112
Create vNIC Templates	113
Create Management vNICs	114
Create vMotion vNICs	118
Create Application vNICs	122
Create iSCSI vNICs	126
Create LAN Connectivity Policy	130
Add iSCSI vNICs in LAN Policy (iSCSI Deployment)	139
Create vHBA Templates for FC Connectivity (FC Deployment)	142
Create FC SAN Connectivity Policies (FC Deployment)	144
Create iSCSI Boot Policy (iSCSI Deployment)	146
Create FC Boot Policies (FC Deployment)	147
Create iSCSI Boot Service Profile Template (iSCSI Deployment)	155
Configure Storage Provisioning	155
Configure Networking Options	156
Configure Storage Options	157
Configure Zoning Options	157
Configure vNIC/HBA Placement	157
Configure vMedia Policy	157
Configure Server Boot Order	157
Configure Maintenance Policy	164
Configure Server Assignment	165
Configure Operational Policies	166
Create iSCSI Boot Service Profiles (iSCSI Deployment)	167
Create FC Boot Service Profile Template (FC Deployment)	168
Configure Storage Provisioning	169
Configure Networking Options	170
Configure SAN Connectivity	171

Configure Zoning	172
Configure vNIC/HBA Placement	172
Configure vMedia Policy	173
Configure Server Boot Order	173
Configure Maintenance Policy	174
Configure Server Assignment	175
Configure Operational Policies	176
Create FC Boot Service Profiles (FC Deployment)	177
Backup the Cisco UCS Manager Configuration	179
Add Servers	179
Gather Necessary WWPN Information (FC Deployment)	179
Gather Necessary IQN Information (iSCSI Deployment)	180
IBM FS9100 iSCSI Storage Configuration (iSCSI Deployment)	182
Create Volumes on the Storage System	182
Create Host Cluster & Host objects	185
Add Hosts to Host Cluster	188
Map Volumes to Hosts and Host Cluster	190
IBM FS9100 Fibre Channel Storage Configuration (FC Deployment)	195
Create Device Aliases and SAN Zoning	195
Cisco MDS - A Switch	195
Cisco MDS - B Switch	198
IBM FS9100 FC Configuration	200
Create Volumes on the Storage System	201
Create Host Cluster & Host Objects	203
Map Volumes to Hosts and Host Cluster	207
VMware vSphere Setup for Cisco UCS Host Environment	211
VMware ESXi 6.7 U2	211
Log into Cisco UCS Manager	211
Install ESXi on the UCS Servers	211
Set Up Management Networking for ESXi Hosts	213
VMware vSphere Configuration	216
Log into VMware ESXi Hosts Using VMware vSphere Client	216
Install VMware Drivers for the Cisco Virtual Interface Card (VIC)	216
Mount Required Datastores	217
Configure NTP on ESXi Hosts	220
Move VM Swap File Location	221
Deploy VMware vCenter Appliance 6.7 (Optional)	222
Adjust vCenter CPU Settings	229
Set Up VMware vCenter Server	230
Setup Data Center, Cluster, DRS and HA for ESXi Nodes	231
Add the VMware ESXi Hosts Using the VMware vSphere Web Client	232
ESXi Dump Collector Setup for iSCSI Hosts (iSCSI Configuration Only)	237

Configure ESXi Networking.....	238
Update Management vSwitcho Configuration	238
Create vSwitch1 for vMotion	241
Configure iSCSI Adapters (iSCSI Deployment Only)	252
Create a VMware vDS for Application and Production Networks.....	262
Add the ESXi Hosts to the vDS	267
References	281
Products and Solutions	281
Interoperability Matrixes.....	282
Appendix.....	283
VersaStack Configuration Backups.....	283
Cisco UCS Backup	283
Cisco Nexus and MDS Backups	285
VMware VCSA Backup	286
About the Authors.....	288

Executive Summary

Cisco Validated Designs (CVDs) deliver 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 the customers and to guide them from design to deployment.

Customers looking to deploy applications using shared data center infrastructure face a number of challenges. A recurrent infrastructure challenge is to achieve the levels of IT agility and efficiency that can effectively meet the company business objectives. Addressing these challenges requires having an optimal solution with the following key characteristics:

- **Availability:** Help ensure applications and services availability at all times with no single point of failure
- **Flexibility:** Ability to support new services without requiring underlying infrastructure modifications
- **Efficiency:** Facilitate efficient operation of the infrastructure through re-usable policies
- **Manageability:** Ease of deployment and ongoing management to minimize operating costs
- **Scalability:** Ability to expand and grow with significant investment protection
- **Compatibility:** Minimize risk by ensuring compatibility of integrated components
 - **Extensibility:** Extensible platform with support for various management applications and configuration tools

Cisco and IBM have partnered to deliver a series of VersaStack solutions that enable strategic data center platforms with the above characteristics. VersaStack solution delivers an integrated architecture that incorporates compute, storage and network design best practices thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance and support that can be used in various stages (planning, designing and implementation) of a deployment.

The VersaStack solution, described in this CVD, delivers a Converged Infrastructure platform (CI) specifically designed for high-performance Virtual Server Infrastructure (VSI), which is a validated solution jointly developed by Cisco and IBM. In this deployment, IBM® FlashSystem 9100 combines the performance of flash and Non-Volatile Memory Express (NVMe) with the reliability and innovation of IBM FlashCore technology and the rich features of IBM Spectrum Virtualize. With the addition of Cisco UCS M5 servers including 2nd gen Intel Xeon Scalable processors and Cisco UCS 6400 series Fabric Interconnects, the solution provides superior compute performance and network throughput with 10/25/40/100GbE support for ethernet using Nexus 9000 series switches in the LAN and 32G support for fibre channel connectivity with the Cisco MDS 9000 portfolio of switches in the SAN.

The design showcases:

- Cisco UCS 6400 Series Fabric Interconnects (FI)
- Cisco UCS 5108 Blade Server chassis
- Cisco Unified Computing System (Cisco UCS) servers with 2nd gen Intel Xeon scalable processors
- Cisco Nexus 9336C-FX2 Switches running NX-OS mode
- Cisco MDS 9132T Fabric Switches
- IBM FlashSystem 9100 NVMe-accelerated Storage
- VMware vSphere 6.7 Update 2

Solution Overview

Introduction

VersaStack solution is a pre-designed, integrated and validated architecture for the data center that combines Cisco UCS servers, Cisco Nexus family of switches, Cisco MDS fabric switches, IBM Storage offerings into a single, flexible architecture. VersaStack is designed for high availability, with no single points of failure, while maintaining cost-effectiveness and flexibility in design to support a wide variety of workloads.

VersaStack designs can support different hypervisor options, bare metal servers and can also be sized and optimized based on customer workload requirements. The VersaStack design discussed in this document has been validated for resiliency (under fair load) and fault tolerance during system upgrades, component failures, and partial loss of power scenarios.

This document discusses the design of the high-performance VersaStack with flash and NVMe based solution. The solution is a pre-designed, best-practice data center architecture with VMware vSphere built on the Cisco Unified Computing System (Cisco UCS). The solution architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Server Infrastructure (VSI).

Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, architects, partner engineering, 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 step-by-step configuration and implementation guidelines for setting up VersaStack. The following design elements distinguish this version of VersaStack from previous models:

- Support for UCS 6454 Fabric Interconnects
- Support for VIC 1400 series adapter cards on UCS M5 servers
- Support for Cisco UCS C480 M5 ML Servers, not validated in this document
- Support for the Second Generation Intel® Xeon® Scalable processor (Cascade Lake) refresh and Intel® Optane™ Data Center persistent memory modules on UCS Intel-based M5 servers
- Improved memory RAS features on M5 servers
- IBM FlashSystem 9100 release 8.2.1.6
- Support for the Cisco UCS release 4.0(4c)
- Validation of 25GbE IP-based storage design with Nexus NX-OS switches supporting iSCSI based storage access
- Validation of VMware vSphere 6.7 U2
- 100 Gigabit per second Ethernet Connectivity
- 32 Gigabit per second Fibre Channel Connectivity

The design that will be implemented is discussed in the VersaStack with VMware vSphere 6.7 Design Guide found at: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/versastack_vmw67_ibmfs9100_design.html

For more information on the complete portfolio of VersaStack solutions, please refer to the VersaStack guides:

<http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html>

Solution Design

Architecture

This VersaStack design aligns with the converged infrastructure configurations and best practices as identified in the previous VersaStack releases. The solution focuses on integration of IBM Flash System 9100 in to VersaStack architecture with Cisco UCS 4th Generation and support for VMware vSphere 6.7 U2.

The system includes hardware and software compatibility support between all components and aligns to the configuration best practices for each of these components. All core hardware components and software releases are listed and supported on the following lists:

http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html

and IBM Interoperability Matrix:

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

The system supports high availability at network, compute and storage layers such that no single point of failure exists in the design. The system utilizes 10/25/40/100 Gbps Ethernet jumbo-frame based connectivity combined with port aggregation technologies such as virtual port-channels (VPC) for non-blocking LAN traffic forwarding. A dual SAN 32Gbps environment provides redundant storage access from compute devices to the storage controllers.

Physical Topology

The VersaStack infrastructure satisfies the high-availability design requirements and is physically redundant across the network, compute and storage stacks. Figure 1 provides a high-level topology of the system connectivity.

To provide the compute to storage system connectivity, this design guides highlights two different storage connectivity options:

- Option 1: iSCSI based storage access through Cisco Nexus Fabric
- Option 2: FC based storage access through Cisco MDS Fabric

This VersaStack design utilizes Cisco UCS platform with Cisco UCS B200 M5 half-width blades and UCS C220 M5 servers connected and managed through Cisco UCS 6454 Fabric Interconnects and the integrated Cisco UCS manager. These high-performance servers are configured as stateless compute nodes where ESXi 6.7 U2 hypervisor is loaded using SAN (iSCSI and FC) boot. The boot disks to store ESXi hypervisor image and configuration along with the block based datastores to host application Virtual Machines (VMs) are provisioned on the IBM FS9100 storage array.

This design has following physical connectivity between the components of VersaStack:

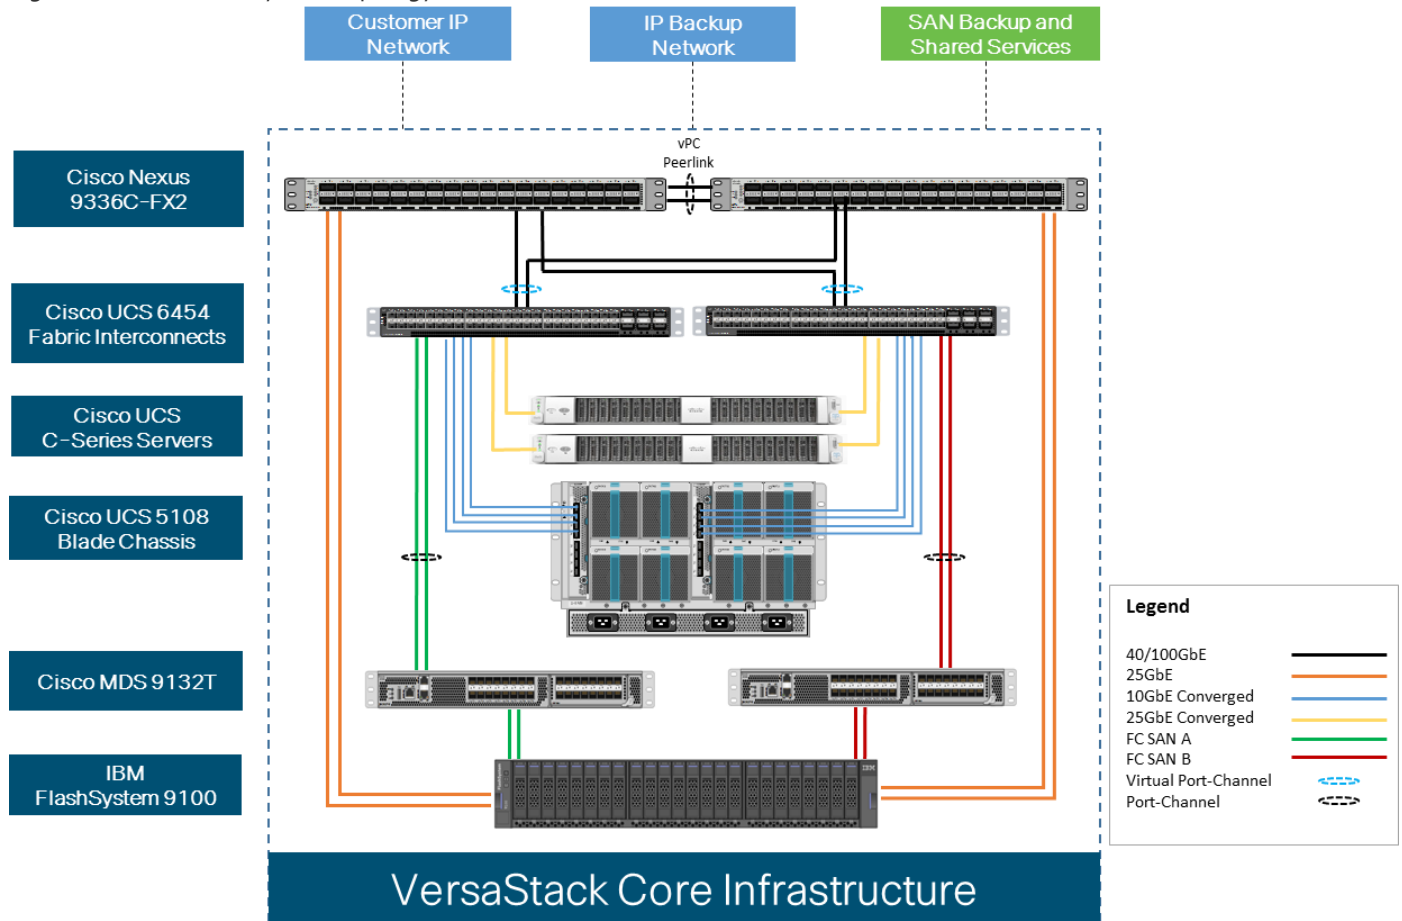
- 4 X 10 Gb Ethernet connections port-channelled between the Cisco UCS 5108 Blade Chassis and the Cisco UCS Fabric Interconnects
- 25 Gb Ethernet connections port-channelled between the Cisco UCS C-Series rackmounts and the Cisco UCS Fabric Interconnects
- 100 Gb Ethernet connections port-channelled between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000s
- 32 Gb Fibre Channel connections port-channelled between the Cisco UCS Fabric Interconnect and Cisco MDS 9132T
- 16 Gb Fibre Channel connections between the Cisco MDS 9132T and IBM FS9100 storage array for fibre channel block storage access

- 25 Gb Ethernet connections between the Cisco Nexus 9000s and IBM FS9100 storage array for iSCSI block storage access



Any supported connectivity to existing customer IP and SAN Networks from the VersaStack core infrastructure is allowed.

Figure 1 VersaStack Physical Topology



This document guides customers through the low-level steps for deploying the base architecture. These procedures explain everything from physical cabling to network, compute, and storage device configurations.

For detailed information about the VersaStack design, see:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/versastack_vmw67_ibmfsg100_design.html

Software Revisions

Table 1 lists the hardware and software versions used for the solution validation.

It is important to note that Cisco, IBM, and VMware have interoperability matrices that should be referenced to determine support for any specific implementation of VersaStack. See the following links for more information:

- [IBM System Storage Interoperation Center](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)

- [VMware Compatibility Guide](#)

Table 1 Hardware and Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6400 Series, Cisco UCS B-200 M5, Cisco UCS C-220 M5	4.0(4C)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, Cisco UCS VIC 1440 and Cisco UCS VIC 1457
	Cisco nenic Driver	1.0.29.0	Ethernet driver for Cisco VIC
	Cisco fnic Driver	4.0.0.40	FCoE driver for Cisco VIC
Network	Cisco Nexus Switches	7.0(3)I7(6)	NXOS
	Cisco MDS 9132T	8.4(1)	FC switch firmware version
Storage	IBM FlashSystem 9100	8.2.1.6	Software version
Software	VMware vSphere ESXi	6.7 update 2	Software version
	VMware vCenter	6.7 update 2	Software version

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available VersaStack configuration. Therefore, appropriate references are provided to indicate the component being configured at each step, such as o1 and o2 or A and B. For example, the Cisco UCS fabric interconnects are identified as FI-A or FI-B. This document is intended to enable customers and partners to fully configure the customer environment and during this process, various steps may require the use of customer-specific naming conventions, IP addresses, and VLAN schemes, as well as appropriate MAC addresses.



This document details network (Nexus and MDS), compute (Cisco UCS), virtualization (VMware) and related IBM FS9100 storage configurations (host to storage system connectivity).

Table 2 lists various VLANs, VSANs and subnets used to setup VersaStack infrastructure to provide connectivity between core elements of the design.

Table 2 VersaStack Infrastructure Configuration

VLAN Name	VLAN	Subnet	Usage
IB-MGMT	11	192.168.160.0/22	Management VLAN to access and manage the servers
iSCSI-A	3161	10.29.161.0/24	iSCSI-A path for booting both B Series and C Series servers and datastore access

VLAN Name	VLAN	Subnet	Usage
iSCSI-B	3162	10.29.162.0/24	iSCSI-B path for booting both B Series and C Series servers and datastore access
vMotion	3173	10.29.173.0/24	VMware vMotion traffic
Native-2	2	N/A	VLAN 2 used as Native VLAN instead of default VLAN (1)
VM Network	3174	10.29.174.0/24	VLAN to carry data traffic for both VM and bare-metal Servers
VSAN-A	101	N/A	Fabric A VSAN for FC Storage access
VSAN-B	102	N/A	Fabric B VSAN for FC Storage access

Physical Infrastructure

The information in this section is provided as a reference for cabling the equipment in VersaStack environment. To simplify the documentation, the architecture shown in Figure 2 is broken down into network, compute and storage related physical connectivity details.

This document assumes that the 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.



Customers can choose interfaces and ports of their liking but failure to follow the exact connectivity shown in figures below will result in changes to the deployment procedures since specific port information is used in various configuration steps



The Nexus 9336C-FX2 switches used in this design support 10/25/40/100 Gbps on all the ports. The switch supports breakout interfaces, each 100Gbps port on the switch can be split in to 4 X 25Gbps interfaces. The QSFP breakout cable has been leveraged to connect 25Gbps iSCSI ethernet ports on the FS9100 storage array to the 100Gbps QSFP port on the switch end. With this connectivity, IBM SFP transceiver on the FS9100 are not required.

Cisco UCS Connectivity to Nexus Switches

For physical connectivity details of Cisco UCS to the Cisco Nexus switches, refer to Figure 2.

Figure 2 Cisco UCS Connectivity to the Nexus Switches

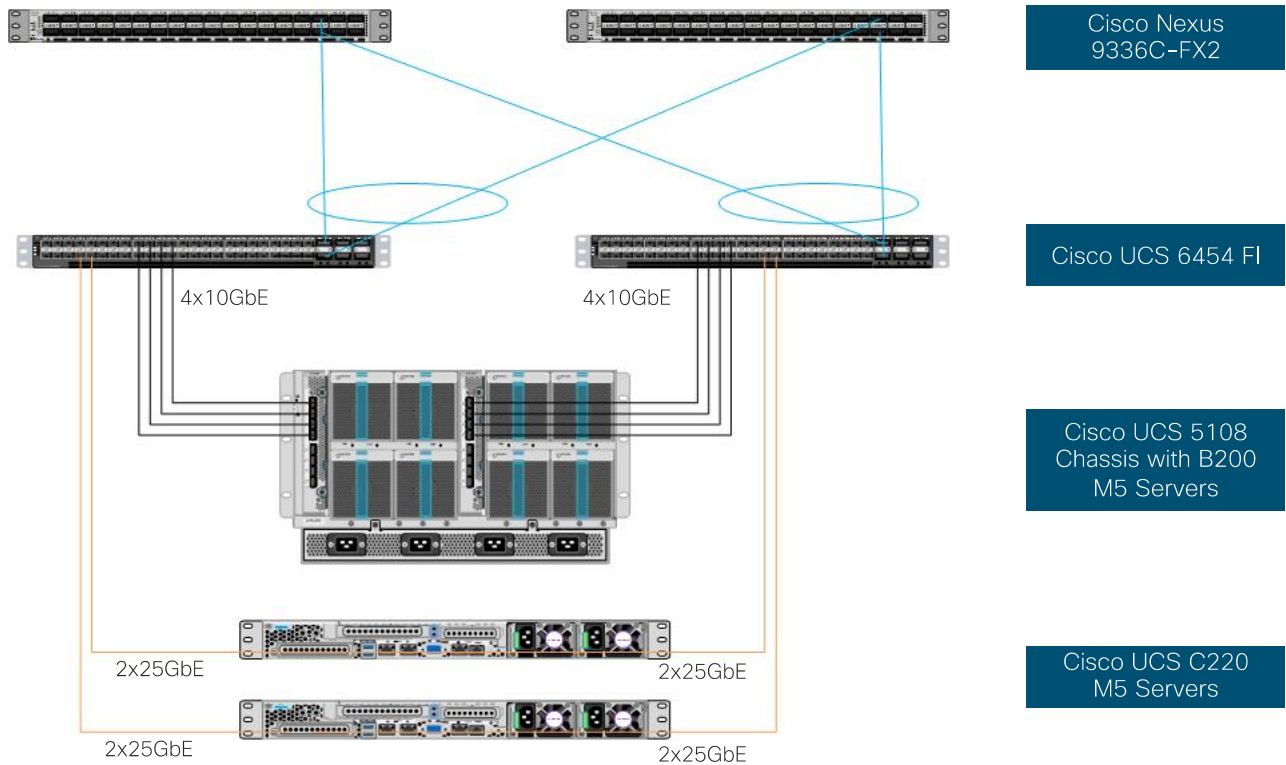


Table 3 Cisco UCS Connectivity to Nexus Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/17	10GbE	Cisco UCS Chassis FEX A	IOM 1/1
Cisco UCS Fabric Interconnect A	Eth1/18	10GbE	Cisco UCS Chassis FEX A	IOM 1/2
Cisco UCS Fabric Interconnect A	Eth1/19	10GbE	Cisco UCS Chassis FEX A	IOM 1/3
Cisco UCS Fabric Interconnect A	Eth1/20	10GbE	Cisco UCS Chassis FEX A	IOM 1/4
Cisco UCS Fabric Interconnect A	Eth1/53	100GbE	Cisco Nexus 9336C-FX2 A	Eth1/31
Cisco UCS Fabric Interconnect A	Eth1/54	100GbE	Cisco Nexus 9336C-FX2 B	Eth1/31
Cisco UCS Fabric Interconnect B	Eth1/17	10GbE	Cisco UCS Chassis FEX B	IOM 1/1
Cisco UCS Fabric Interconnect B	Eth1/18	10GbE	Cisco UCS Chassis FEX B	IOM 1/2
Cisco UCS Fabric Interconnect B	Eth1/19	10GbE	Cisco UCS Chassis FEX B	IOM 1/3
Cisco UCS Fabric Interconnect B	Eth1/20	10GbE	Cisco UCS Chassis FEX B	IOM 1/4
Cisco UCS Fabric Interconnect B	Eth1/53	100GbE	Cisco Nexus 9336C-FX2 A	Eth1/32
Cisco UCS Fabric Interconnect B	Eth1/54	100GbE	Cisco Nexus 9336C-FX2 B	Eth1/32

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C220 M5 Server 1	VIC Port 1, Port 3	25GbE	Cisco UCS Fabric Interconnect A	Eth1/21, Eth1/22
Cisco UCS C220 M5 Server 1	VIC Port 2, Port 4	25GbE	Cisco UCS Fabric Interconnect B	Eth1/21, Eth1/22
Cisco UCS C220 M5 Server 2	VIC Port 1, Port 3	25GbE	Cisco UCS Fabric Interconnect A	Eth1/23, Eth1/24
Cisco UCS C220 M5 Server 2	VIC Port 2, Port 4	25GbE	Cisco UCS Fabric Interconnect B	Eth1/23, Eth1/24
Cisco Nexus 9336C-FX2 A	Eth1/33	100GbE	Cisco Nexus 9336C-FX2 B	Eth1/33
Cisco Nexus 9336C-FX2 A	Eth1/34	100GbE	Cisco Nexus 9336C-FX2 B	Eth1/34

IBM FS9100 Connectivity to Nexus Switches

For physical connectivity details of IBM FS9100 node canisters to the Cisco Nexus Switches, refer to Table 3 . This deployment shows connectivity for a pair of IBM FS9100 node canisters. Additional nodes can be connected to open ports on Nexus switches as needed.

Figure 3 IBM FS9100 Connectivity to Nexus 9k Switches




Table 4 IBM FS9100 Connectivity to the Nexus Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
IBM FS9100 node 1	Port 5	25GbE	Cisco Nexus 9336C-FX2 A	Eth1/11/1*
IBM FS9100 node 1	Port 6	25GbE	Cisco Nexus 9336C-FX2 B	Eth1/11/1*
IBM FS9100 node 2	Port 5	25GbE	Cisco Nexus 9336C-FX2 A	Eth1/11/2*
IBM FS9100 node 2	Port 6	25GbE	Cisco Nexus 9336C-FX2 B	Eth1/11/2*



* Breakout interfaces with one 100G QSFP port on the Nexus 9336C-FX2 is split in to 4 X 25Gbps SFP interfaces connected to the IBM FS9100. Cisco QSFP100G-4SFP25G breakout cable has been leveraged for this connectivity.

Cisco UCS connectivity to SAN Fabric

For physical connectivity details of Cisco UCS to an MDS 9132T based redundant SAN fabric, refer to Figure 4.

Figure 4 Cisco UCS Connectivity to Cisco MDS Switches




Table 5 Cisco UCS Connectivity to Cisco MDS Switches

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	FC1/1	32Gbps	Cisco MDS 9132T A	FC1/1
Cisco UCS Fabric Interconnect A	FC1/2	32Gbps	Cisco MDS 9132T A	FC1/2
Cisco UCS Fabric Interconnect A	FC1/3	32Gbps	Cisco MDS 9132T A	FC1/3
Cisco UCS Fabric Interconnect A	FC1/4	32Gbps	Cisco MDS 9132T A	FC1/4
Cisco UCS Fabric Interconnect B	FC1/1	32Gbps	Cisco MDS 9132T B	FC1/1
Cisco UCS Fabric Interconnect B	FC1/2	32Gbps	Cisco MDS 9132T B	FC1/2
Cisco UCS Fabric Interconnect B	FC1/3	32Gbps	Cisco MDS 9132T B	FC1/3
Cisco UCS Fabric Interconnect B	FC1/4	32Gbps	Cisco MDS 9132T B	FC1/4

Figure 5 illustrates FC connectivity for IBM FS9100 storage array. Additional nodes can be connected and configured by following the same design guidelines.

Figure 5 IBM FS9100 and Storage System FC Connectivity




Table 6 IBM FS9100 Storage System FC Connectivity

Local Device	Local Ports	Connection	Remote Device	Remote Port
IBM FS9100 Node Canister 1	Port 1	16Gbps	Cisco MDS 9132T A	FC1/5
IBM FS9100 Node Canister 1	Port 2	16Gbps	Cisco MDS 9132T B	FC1/5
IBM FS9100 Node Canister 1	Port 3	16Gbps	Cisco MDS 9132T A	FC1/6
IBM FS9100 Node Canister 1	Port 4	16Gbps	Cisco MDS 9132T B	FC1/6
IBM FS9100 Node Canister 2	Port 1	16Gbps	Cisco MDS 9132T A	FC1/7
IBM FS9100 Node Canister 2	Port 2	16Gbps	Cisco MDS 9132T B	FC1/7
IBM FS9100 Node Canister 2	Port 3	16Gbps	Cisco MDS 9132T A	FC1/8
IBM FS9100 Node Canister 2	Port 4	16Gbps	Cisco MDS 9132T B	FC1/8

Network Configuration

The procedures in this section describe how to configure the Cisco Nexus switches for use in a base VersaStack environment. This procedure assumes the use of Nexus 9336C-FX2 switches running 7.0(3)I7(6) code. Configuration on a differing model of Nexus 9000 series switch should be comparable but may differ slightly with model and changes in NX-OS release. The Cisco Nexus 9336C-FX2 switch and NX-OS 7.0(3)I7(6) release were used in validation of this VersaStack solution, so steps will reflect this model and release.



Connectivity between the Nexus switches and IBM FS9100 for iSCSI access depends on the Nexus 9000 switch model used within the architecture. If any supported Nexus switch with 25Gbps capable SFP ports is used, breakout cable is not required and ports from the switch to IBM FS9100 can be connected directly using the SFP transceivers on both sides.



With Cisco Nexus 9000 release 7.0(3)I7(6), autonegotiation (40G/100G) is not supported on ports 1-6 and 33-36 on the Cisco Nexus 9336C-FX2 switch. If these ports are used for connectivity, port speed and duplex should be hard set at both ends of the connection.

Physical Connectivity

Physical cabling should be completed by following the diagram and table references in the previous sections.

Cisco Nexus 9000 Initial Configuration Setup

The steps provided in this section details for the initial Cisco Nexus 9336C-FX2 Switch setup. In this case, we are connected using a Cisco Terminal Server that is connected via the console port on the switch.



Cisco Nexus A

To set up the initial configuration for the first Cisco Nexus switch <nexus-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.

1. Configure the switch:

```
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]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

Cisco Nexus B

To set up the initial configuration for the second Cisco Nexus switch, 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.

1. Configure the switch:

```

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]: y
NTP server IPv4 address: <global-ntp-server-ip>
Configure default interface layer (L3/L2) [L2]: Enter
Configure default switchport interface state (shut/noshut) [noshut]: shut
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter
Would you like to edit the configuration? (yes/no) [n]: Enter

```

2. Review the configuration summary before enabling the configuration.

Enable Appropriate Cisco Nexus 9000 Features and Settings

Enable Licenses

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To license the Cisco Nexus switches, follow these steps:

1. Log in a admin.

2. Run the following commands:

```
config terminal
feature udd
feature lacp
feature vpc
```

Set Global Configurations

Cisco Nexus 9000 A and Cisco Nexus 9000 B

To set global configurations, follow these steps on both switches:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

Setup NTP (optional)

The following procedure can be used to optionally enable the NTP service on the Nexus switches. The procedure includes the setup of NTP distribution on both the mgmt 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 9000 A and Cisco Nexus 9000 B

1. Run the following commands

```
feature interface-vlan
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
```

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
```

Create VLANs for VersaStack IP Traffic

Cisco Nexus 9000 A and Cisco Nexus 9000 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 <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
vlan <iSCSI-A_VLAN_id>
name iSCSI-A-VLAN
vlan <iSCSI-B_VLAN id>
name iSCSI-B-VLAN
exit
copy run start
```

Configure Virtual Port Channel Domain

Cisco Nexus 9000 A

To configure vPC domain for switch A, follow these steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 10
```

2. Make the Nexus 9000A the primary vPC peer by defining a low priority value:

```
role priority 10
```

3. Use the management interfaces on the supervisors of the Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <Mgmt. IP address for Switch B> source <Mgmt. IP address for Switch A>
```

4. Enable the following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
ip arp synchronize
auto-recovery
copy run start
```

Cisco Nexus 9000 B

To configure the vPC domain for switch B, follow these steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 10
```

2. Make the Nexus 9000A the primary vPC peer by defining a low priority value:

```
role priority 20
```

3. Use the management interfaces on the supervisors of the Nexus 9000s to establish a keepalive link:

```
peer-keepalive destination <Mgmt. IP address for Switch A> source <Mgmt. IP address for Switch B>
```

4. Enable the following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
ip arp synchronize
auto-recovery
copy run start
```


Configure Network Interfaces for the vPC Peer Links

To configure the network interfaces for the vPC Peer links, follow these steps:

Cisco Nexus 9000 A

1. Define a port description for the interfaces connecting to vPC Peer <nexus-B-hostname>.

```
interface Eth1/33
description VPC Peer <Nexus-B Switch Name>:1/33
interface Eth1/34
description VPC Peer <Nexus-B Switch Name>:1/34
```

2. Apply a port channel to both vPC Peer links and bring up the interfaces.

```
interface Eth1/33,Eth1/34
channel-group 10 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <nexus_B_hostname>.

```
interface Po10
description vPC peer-link
```

4. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLAN.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

5. Make the port channel and associated interfaces spanning tree network ports.

```
spanning-tree port type network
```

6. Set port speed and duplex.

```
speed 100000
duplex full
no negotiate
```

7. Make this port-channel the VPC peer link and bring it up.

```
vpc peer-link
no shutdown
copy run start
```

Cisco Nexus 9000 B

1. Define a port description for the interfaces connecting to VPC Peer <nexus_A_hostname>.

```
interface Eth1/33
description VPC Peer <Nexus-A Switch Name>:1/33
interface Eth1/34
description VPC Peer <Nexus-A Switch Name>:1/34
```

2. Apply a port channel to both VPC Peer links and bring up the interfaces.

```
interface Eth1/33,Eth1/34
channel-group 10 mode active
no shutdown
```

3. Define a description for the port-channel connecting to <nexus_A_hostname>.

```
interface Po10
description vPC peer-link
```

4. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLAN.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

5. Make the port channel and associated interfaces spanning tree network ports.

```
spanning-tree port type network
```

6. Set port speed and duplex.

```
speed 100000
duplex full
no negotiate
```

7. Make this port-channel the VPC peer link and bring it up.

```
vpc peer-link
no shutdown
copy run start
```

Configure Network Interfaces to Cisco UCS Fabric Interconnects

To configure the network interfaces for the Cisco UCS Fabric Interconnects, follow these steps:

Cisco Nexus 9000 A

1. Define a description for the port-channel connecting to <UCS Cluster Name>-A.

```
interface Po13
description <UCS Cluster Name>-A
```

2. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion, iSCSI and the native VLANs.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

3. Set port speed and duplex.

```
speed 100000
duplex full
no negotiate auto
```

4. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

5. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

6. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

7. Define a port description for the interface connecting to <UCS Cluster Name>-A.

```
interface Eth1/31
description <UCS Cluster Name>-A:53
```

8. Apply it to a port channel and bring up the interface.

```
channel-group 13 force mode active
no shutdown
```

9. Define a description for the port-channel connecting to <UCS Cluster Name>-B.

```
interface Po14
description <UCS Cluster Name>-B
```

10. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

11. Set port speed.

```
speed 100000
duplex full
no negotiate auto
```

12. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

13. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

14. Make this a VPC port-channel and bring it up.

```
vpc 14
no shutdown
```

15. Define a port description for the interface connecting to <UCS Cluster Name>-B.

```
interface Eth1/32
description <UCS Cluster Name>-B:1/53
```

16. Apply it to a port channel and bring up the interface.

```
channel-group 14 force mode active
no shutdown
copy run start
```

Cisco Nexus 9000 B

1. Define a description for the port-channel connecting to <UCS Cluster Name>-B.

```
interface Po13
description <UCS Cluster Name>-A
```

2. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

3. Set port speed.

```
speed 100000
duplex full
no negotiate auto
```

4. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

5. Set the MTU to 9216 to support jumbo frames.

```
mtu 9216
```

6. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

7. Define a port description for the interface connecting to <UCS Cluster Name>-B.

```
interface Eth1/31
description <UCS Cluster Name>-A:1/54
```

8. Apply it to a port channel and bring up the interface.

```
channel-group 13 force mode active
no shutdown
```

9. Define a description for the port-channel connecting to <UCS Cluster Name>-A.

```
interface Po14
description <UCS Cluster Name>-B
```

10. Make the port-channel a switchport and configure a trunk to allow in-band management, VM traffic, vMotion and the native VLANs.

```
switchport
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>, <iSCSI-A
VLAN id>, <iSCSI-B VLAN id>
```

11. Set port speed.

```
speed 100000
duplex full
no negotiate auto
```

12. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
```

13. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
```

14. Make this a VPC port-channel and bring it up.

```
vpc 14
no shutdown
```

15. Define a port description for the interface connecting to <UCS Cluster Name>-A.

```
interface Eth1/32
description <UCS Cluster Name>-B:1/54
```

16. Apply it to a port channel and bring up the interface.

```
channel-group 14 force mode active
no shutdown
copy run start
```

Enable UDLD for Cisco UCS Interfaces

Enable aggressive unidirectional link detection (UDLD) on interfaces connected to Cisco UCS.

Cisco Nexus A and Cisco Nexus B

1. From the global configuration mode, run either of the following commands depending on the connectivity:

```
interface Eth1/x

#For Copper cable or twinnax connections use the following command
udld aggressive

#For fibre optic connections use the following command
udld enable
```

Configure Network Interfaces Connected to IBM FS9100 iSCSI Ports (iSCSI Deployment)



This configuration step can be skipped if the UCS environment does not need access to storage using iSCSI.

To configure the network interfaces for IBM FS9100 iSCSI ports, follow these steps:

Cisco Nexus 9000 A & B

The 100Gbps design in this document uses a pair of Nexus 9336C-FX2 switches built with all ports being capable of the 100Gbps Quad Small Form Factor Pluggable Plus (QSFP+) type. The IBM FS9100 has 25Gbps SFP+ ports for iSCSI connectivity. The 100Gbps QSFP+ ports on the Nexus 9336C-FX2 switches in this design have been connected to the IBM FS9100 iSCSI ethernet SFP+ ports using a QSFP+ Breakout Cable.

Configuration of the QSFP+ ports will use the interface breakout command as shown in this example to turn the 100G interface Ethernet 1/11 into 4x25G interfaces on both the Nexus switches:

```
show running-config interface Ethernet1/11

interface Ethernet1/11
  no switchport

interface breakout module 1 port 11 map 25g-4x

show running-config interface Ethernet1/11/1-4
```

```
interface Ethernet1/11/1
interface Ethernet1/11/2
interface Ethernet1/11/3
interface Ethernet1/11/4
```



Connectivity between the Nexus switches and IBM FS9100 for iSCSI access depends on the Nexus 9000 switch model used within the architecture. If any supported Nexus switch with 25Gbps capable SFP ports is used, breakout cable is not required and ports from the switch to IBM FS9100 can be connected directly using the SFP transceivers on both sides.

Cisco Nexus 9000 A

1. Define a description for the Ethernet port connecting to <FS9100 Node1, P5>.

```
interface Ethernet1/11/1
description <FS9100-Node1-iSCSI-P5>
```

2. Make the Interface access port and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-A VLAN id>
```

3. Make the interface spanning normal.

```
spanning-tree port type edge
```

4. Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

5. Define a description for the Ethernet port connecting to <FS9100 Node2, P5>.

```
interface Ethernet1/11/2
description <FS9100-Node2-iSCSI-P5>
```

6. Make the Interface a access port and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-A VLAN id>
```

7. Make the interface spanning normal.

```
spanning-tree port type edge
```


- Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

Cisco Nexus 9000 B

- Define a description for the Ethernet port connecting to <FS9100 Node1, P6>.

```
interface Ethernet1/11/1
description <FS9100-Node1-iSCSI-P6>
```

- Make the Interface access port and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-B VLAN id>
```

- Make the interface spanning normal.

```
spanning-tree port type edge
```

- Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
copy run start
```

- Define a description for the Ethernet port connecting to <FS9100 Node1, P6>.

```
interface Ethernet1/11/2
description <FS9100-Node2-iSCSI-P6>
```

- Make the Interface a access port and configure the switchport access VLAN.

```
switchport mode access
switchport access vlan <iSCSI-B VLAN id>
```

- Make the interface spanning normal.

```
spanning-tree port type edge
```

- Set the MTU to be 9216 to support jumbo frames.

```
mtu 9216
no shutdown
```

```
copy run start
```

Management Uplink into Existing Network Infrastructure

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

Cisco Nexus 9000 A and B using Port Channel Example

To enable management access across the IP switching environment leveraging port channel in config mode run the following commands:

1. Define a description for the port-channel connecting to management switch.

```
interface po6
description IB-MGMT
```

2. Configure the port as an access VLAN carrying the InBand management VLAN traffic.

```
switchport
switchport mode access
switchport access vlan <IB-MGMT VLAN id>
```

3. Make the port channel and associated interfaces normal spanning tree ports.

```
spanning-tree port type normal
```

4. Make this a VPC port-channel and bring it up.

```
vpc 6
no shutdown
```

5. Define a port description for the interface connecting to the management plane.

```
interface Eth1/30
description IB-MGMT-SWITCH_uplink
```

6. Apply it to a port channel and bring up the interface.

```
channel-group 6 force mode active
no shutdown
```

7. Save the running configuration to start-up in both Nexus 9000s and run commands to look at port and port channel.

```
Copy run start
sh int eth1/30 br
sh port-channel summary
```

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 VersaStack 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 udld neighbors
show run int
show int
```

Cisco MDS 9132T Configuration (FC Deployment)

This section explains how to configure the Cisco MDS 9000s for use in a VersaStack 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 or if only iSCSI storage access is required, skip this section.

Physical Connectivity

Follow the physical connectivity guidelines for VersaStack as explained in the section Physical Infrastructure.

VersaStack Cisco MDS Base Configuration

The following procedures describe how to configure the Cisco MDS switches for use in a base VersaStack environment. This procedure assumes you are using the Cisco MDS 9132T with NX-OS 8.3(1).

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.

Cisco MDS 9132T A

1. Configure the switch using 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 summary before enabling the configuration.

Cisco MDS 9132T B

1. Configure the switch using 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 summary before enabling the configuration.

Cisco MDS 9132T Initial Configuration Setup

To perform Cisco MDS Initial configuration, follow these steps:

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. Login as admin.
2. Run the following commands:

```
Configure terminal
feature npiv
feature fport-channel-trunk
feature lldp
device-alias mode enhanced
device-alias commit
```

Add Second NTP server

Cisco MDS 9132T A and Cisco MDS 9132T B

To configure the second NTP server, follow this step:

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

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

Configure Individual Ports

To configure Cisco MDS individual ports used for Cisco UCS and IBM FS9100 connectivity, follow these steps:

Cisco MDS 9132T A

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

From the global configuration mode, run the following commands:

```
interface fcl/1
```

```
switchport description <ucs-clustername>-a:1/1
channel-group 1 force
no shutdown
exit

interface fcl/2
switchport description <ucs-clustername>-a:1/2
channel-group 1 force
no shutdown
exit

interface fcl/3
switchport description <ucs-clustername>-a:1/3
channel-group 1 force
no shutdown
exit

interface fcl/4
switchport description <ucs-clustername>-a:1/4
channel-group 1 force
no shutdown
exit

interface fcl/5
switchport description <FS9100-Node1-FC1>
switchport speed 16000
no shutdown
exit

interface fcl/6
switchport description <FS9100-Node1-FC3>
switchport speed 16000
no shutdown
exit

interface fcl/7
switchport description <FS9100-Node2-FC1>
```

```
switchport speed 16000
no shutdown
exit

interface fcl/8
switchport description <FS9100-Node2-FC3>
switchport speed 16000
no shutdown
exit

interface port-channel1
channel mode active
switchport description <ucs-clustername>-a
no shutdown
exit
```

Cisco MDS 9132T B

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

From the global configuration mode, run the following commands:

```
interface fcl/1
switchport description <ucs-clustername>-b:1/1
channel-group 2 force
no shutdown
exit

interface fcl/2
switchport description <ucs-clustername>-b:1/2
channel-group 2 force
no shutdown
exit

interface fcl/3
switchport description <ucs-clustername>-b:1/3
channel-group 2 force
no shutdown
```



```
exit

interface fcl/4
switchport description <ucs-clustername>-b:1/4
channel-group 2 force
no shutdown
exit

interface fcl/5
switchport description <FS9100-Node1-FC2>
switchport speed 16000
no shutdown
exit

interface fcl/6
switchport description <FS9100-Node1-FC4>
switchport speed 16000
no shutdown
exit

interface fcl/7
switchport description <FS9100-Node2-FC2>
switchport speed 16000
no shutdown
exit

interface fcl/8
switchport description <FS9100-Node2-FC4>
switchport speed 16000
no shutdown
exit

interface port-channel2
channel mode active
switchport description <ucs-clustername>-b
no shutdown
```

```
exit
```

Create VSANs

Cisco MDS 9132T A

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

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 fc1/5
vsan <vsan-a-id> interface fc1/6
vsan <vsan-a-id> interface fc1/7
vsan <vsan-a-id> interface fc1/8
vsan <vsan-a-id> interface port-channel1
exit
```

Cisco MDS 9132T B

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

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 fc1/5
vsan <vsan-b-id> interface fc1/6
vsan <vsan-b-id> interface fc1/7
vsan <vsan-b-id> interface fc1/8
vsan <vsan-b-id> interface port-channel2
exit
```

Initial Storage Configuration


IBM FlashSystem 9100



FlashSystem 9100 systems have specific connection requirements. Care must be taken to note the orientation of each node canister in the control enclosure.

The FlashSystem 9100 control enclosure contains two node canisters. A label on the control enclosure identifies each node canister and power supply unit (PSU). As Figure 6 shows, node canister 1 is on top and node canister 2 is on the bottom. Because the node canisters are inverted, the location of the ports and the port numbering are oriented differently on each node canister. It is important to remember this orientation when installing adapters and cables.

Figure 6 Orientation of the Node Canisters and PSUs



For example, Figure 7 shows the top node canister. On this canister, the PCIe slot and port numbering goes from right to left. PCIe adapter slot 1 contains a 4-port 16 Gbps Fibre Channel adapter, PCIe slot 2 contains a 2-port 25 Gbps iWARP Ethernet adapter, and PCIe slot 3 contains a 4-port 12 Gbps SAS adapter. The onboard Ethernet and USB ports are also shown.

Figure 7 Orientation of Ports on Node Canister 1





Figure 8 shows the bottom node canister. This node canister has the same type and number of adapters installed. However, on the bottom canister, the PCI slot and port numbering goes from left to right.

Figure 8 Orientation of Ports on Node Canister 2



Four 10 Gb Ethernet ports on each node canister provide system management connections and iSCSI host connectivity. A separate technician port provides access to initialization and service assistant functions. Table 7 describes each port.

Table 7 Summary of Onboard Ethernet Ports

On board Ethernet Port	Speed	Function
1	10 Gbps	Management IP, Service IP, Host I/O
2	10 Gbps	Secondary Management IP, Host I/O
3	10 Gbps	Host I/O
4	10 Gbps	Host I/O
T	1 Gbps	Technician Port - DHCP/DNS for direct attach service management

The following connections are required for FlashSystem 9100 control enclosures:

- Each control enclosure requires two Ethernet cables to connect it to an Ethernet switch. One cable connects to port 1 of the top node canister, and the other cable connects to port 1 of the bottom node canister. For 10 Gbps ports, the minimum link speed is 1 Gbps. Both Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6) are supported.
- To ensure system failover operations, Ethernet port 1 on each node canister must be connected to the same set of subnets. If used, Ethernet port 2 on each node canister must also be connected to the same set of subnets. However, the subnets for Ethernet port 1 do not have to be the same as Ethernet port 2.
- If you have more than one control enclosure in your system, the control enclosures communicate through their Fibre Channel ports.
- Each FlashSystem 9100 node canister also has three PCIe interface slots to support optional host interface adapters. The host interface adapters can be supported in any of the interface slots. Table 8 provides an overview of the host interface adapters.
- The 2-port SAS host interface adapter supports expansion enclosures. In total, FlashSystem 9100 control enclosures can have up to 20 chain-linked expansion enclosures, 10 per port.

Table 8 Summary of Supported Host Interface Adapters

Protocol	Feature	Ports	FRU part number	Quantity supported
16 Gbs Fibre Channel	AHB3	4	01YM333	0-3
25 Gbs Ethernet (RoCE)	AHB6	2	01YM283	0-3
25 Gbs Ethernet (iWARP)	AHB7	2	01YM285	0-3
12 Gb SAS Expansion	AHBA	4, but only 2 are active for SAS expansion chains.	01YM338	0-1



Each node canister within the control enclosure (I/O group) must be configured with the same host interface adapters.

Each node canister has four onboard 10 Gbps Ethernet ports. A node canister can also support up to three 2-port 25 Gbps Ethernet host interface adapters.

Table 9 lists the fabric types that can be used for communicating between hosts, nodes, and RAID storage systems. These fabric types can be used at the same time.

Table 9 Communications types

Communications type	Host to node	Node to storage system	Node to node
Fibre Channel SAN	Yes	Yes	Yes
iSCSI 10 Gbps Ethernet 25 Gbps Ethernet	Yes	Yes	No
iSER 25 Gbps Ethernet	Yes	No	No

The feature codes for the 16 Gbps Fibre Channel adapter, 25Gbps iWarp adapter, and the 25Gbps RoCE adapter each include standard SFP transceivers for each adapter. In this design the 25Gbps RoCE adapter has been leveraged for iSCSI connectivity and the ports are connected to the Cisco Nexus 9336C-FX2 switches using breakout cables, SFP transceivers are not required with this connectivity.

The 2-port 25 GB Ethernet adapter for iWARP and the 2-port 25GB Ethernet adapter for RDMA over Converged Ethernet (RoCE) both support iSER host attachment. However, RoCE and iWARP are not cross-compatible; therefore, it is important to use the adapter that matches the iSER implementation on your SAN if iSER is planned to be implemented in the future.



This document implements traditional iSCSI, iSER based iSCSI implementation can be configured with the support of iSER on Cisco VIC 1400 series when available with the future releases of Cisco UCS software.

IBM Service Support Representative (SSR) Configuration

To install the FlashSystem 9100 hardware, an IBM SSR must complete the following tasks:



You must complete the planning tasks and provide completed worksheets to the IBM SSR before they can proceed with installing and initializing your system.

- An IBM SSR unpacks and installs the AF7/AF8 control enclosures and any optional SAS expansion enclosures in the rack.
- Referring to the worksheets that you completed, the IBM SSR completes the cabling.



If the IBM SSR is aware of your intent to add the FlashSystem 9100 to an existing system, the IBM SSR installs the FlashSystem 9100 control enclosure for you but does not initialize a system on it. If you are planning on adding a FlashSystem 9100 control enclosure to an existing Storwize® V7000 system, inform the IBM SSR of this intention. In these cases, the IBM SSR installs the FlashSystem 9100 control enclosure for you, but does not initialize a system on it, because the existing system is already initialized.

After the hardware is installed, an IBM SSR connects a workstation to an AF7/AF8 control enclosure technician port and completes the following tasks:

- Configuring the system with a name, and management and service IP addresses.
- Logging in to the control enclosure using the management GUI and completing the system setup wizard using information from the customer-supplied worksheets.

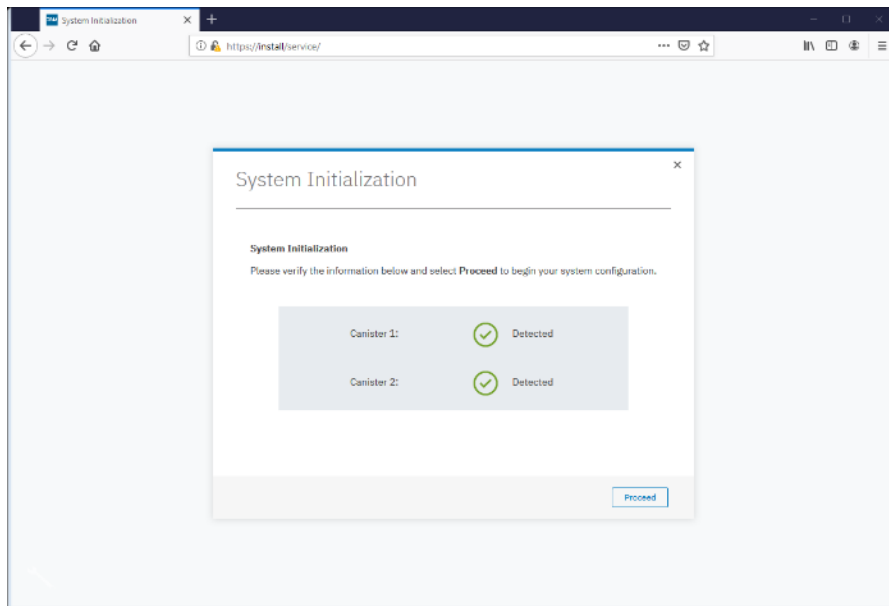
The SSR configuration steps are documented below.

Initialize the System

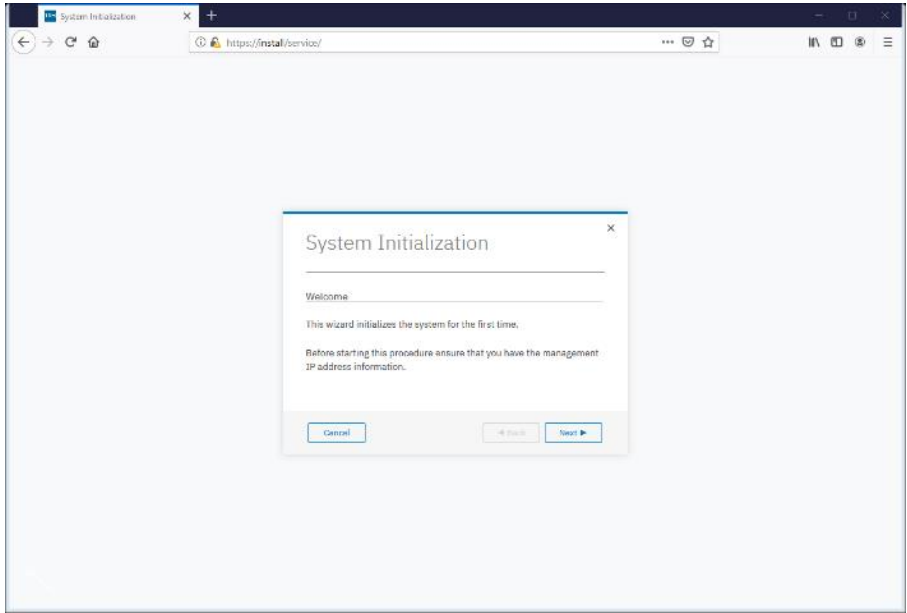
The initial configuration requires a workstation be locally attached to the Ethernet port labelled "T" on the Upper node canister in the FS9100 enclosure. "T" refers to Tech Port and will allocate an IP address to the connected workstation using DHCP and will redirect any DNS queries to the System Initialization page. This page shows the status of each node canister in the enclosure and will guide you through the initialization process.

To initialize the system, follow these steps:

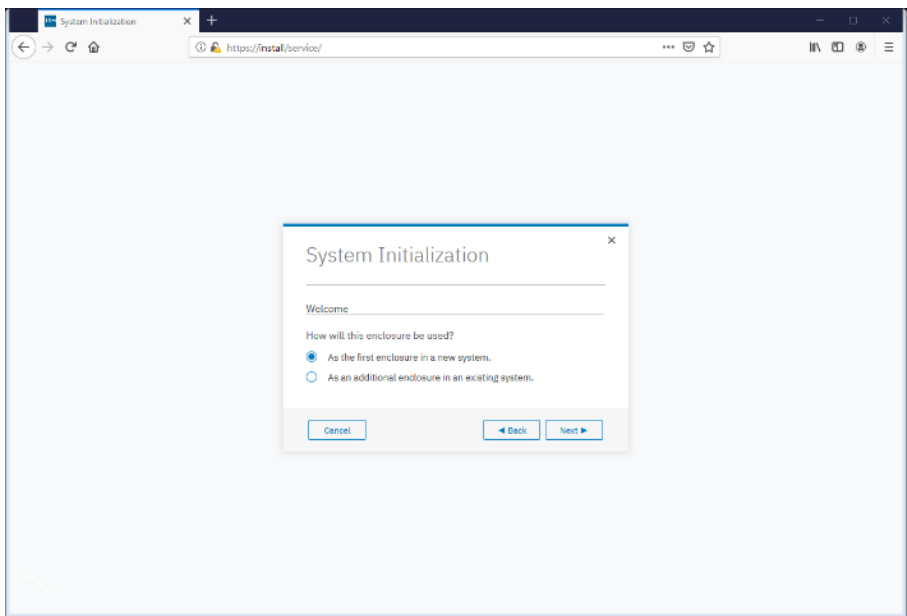
1. Ensure both node canisters have been detected and click **Proceed**.



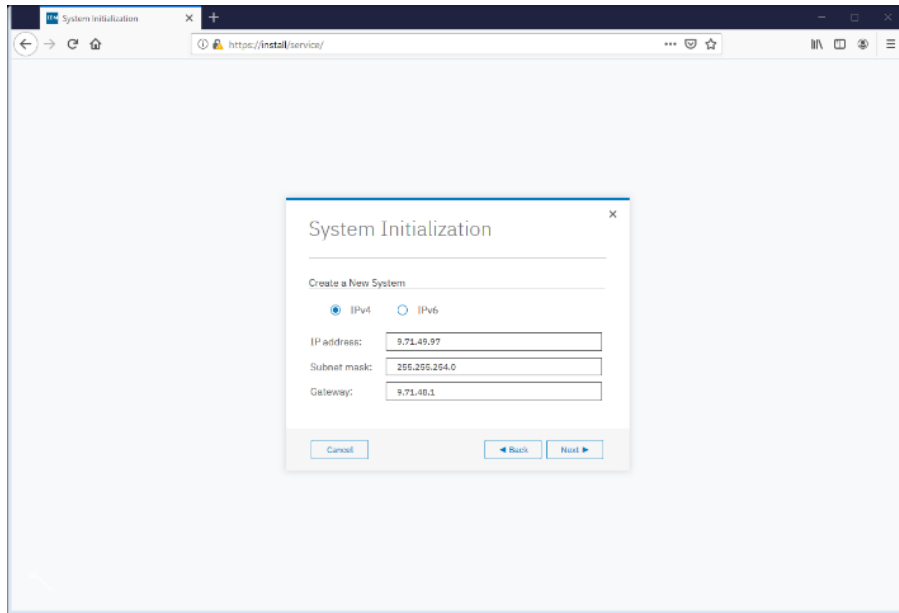
2. Click **Next** through the Welcome screen.



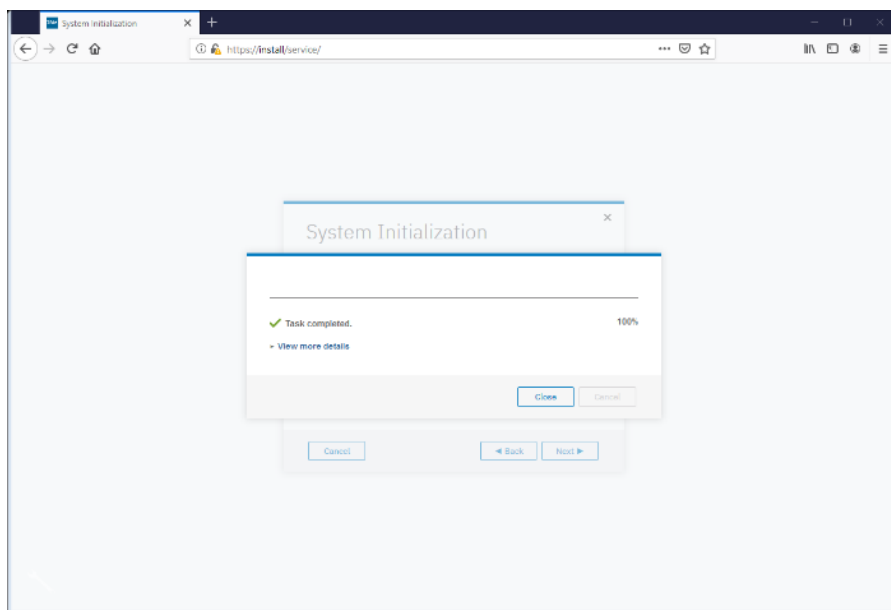
3. Select the option to define the enclosure as the first in a new system




4. Enter the network details for the management interface for the new system. This IP address is sometimes referred to as the Management IP, or Cluster IP and will be used to manage the FS9100 system via the web interface or CLI via SSH.



5. Acknowledge the Task Completion message.



6. The initial configuration steps are now complete, and the system will now restart the Web Server.



Prepare FS9100 for Customer Environments

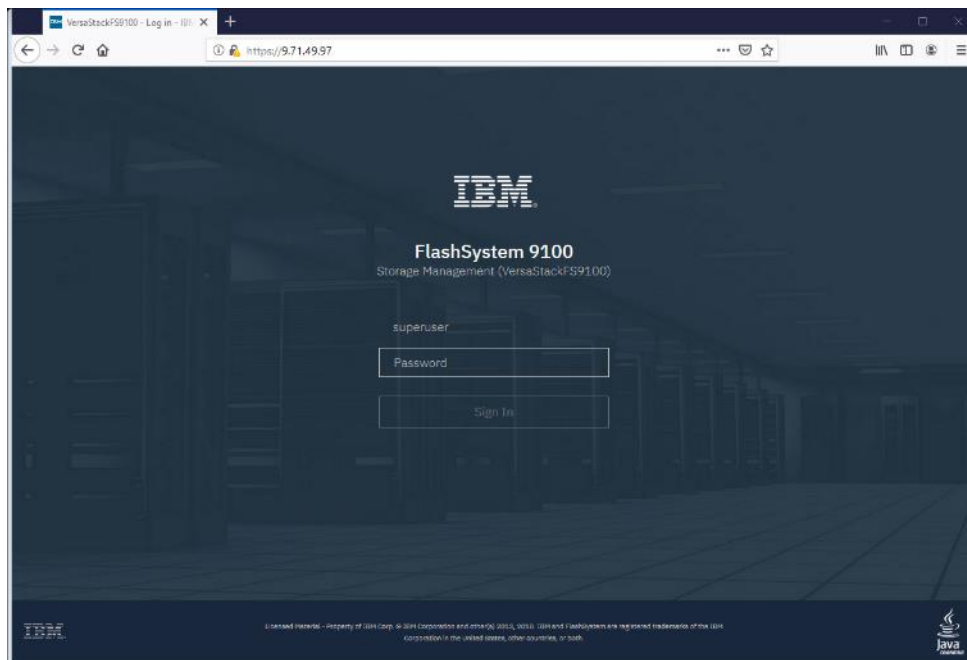
Now the Management IP is enabled, all future configuration steps are made with this interface.

To prepare the FS9100 for customer environments, follow these steps:


1. Log in using the default credentials:

Username: `superuser`

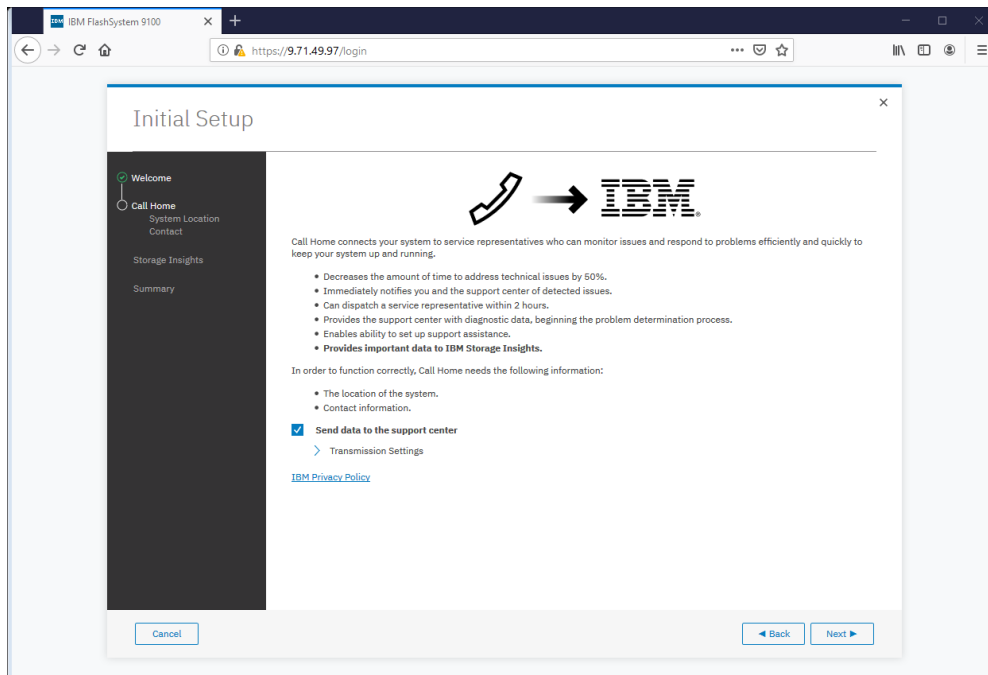
Password: `passw0rd`



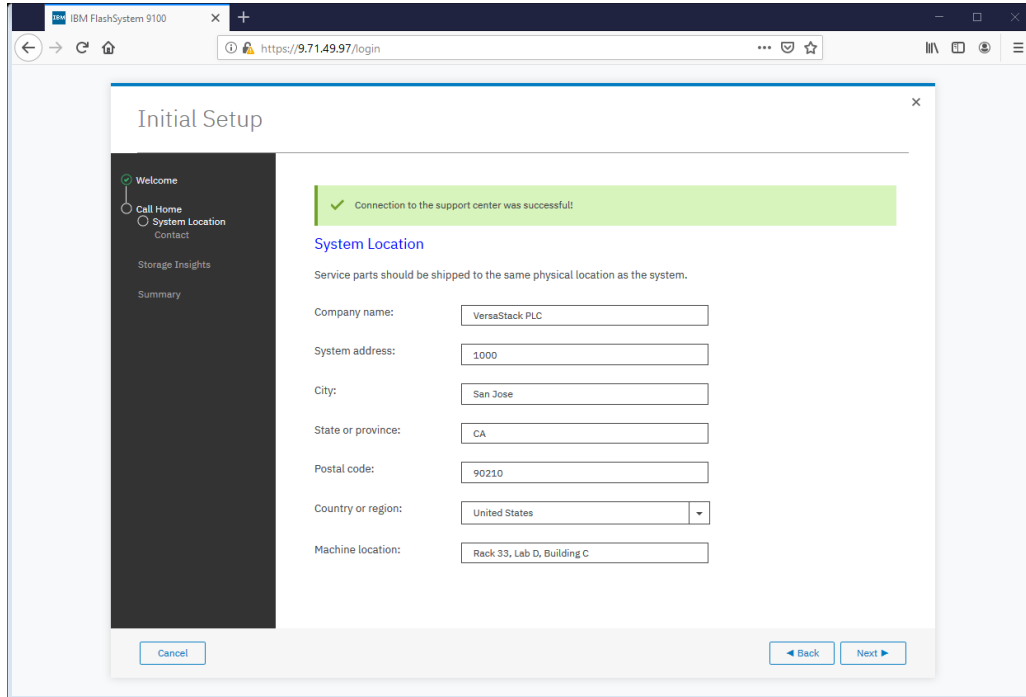
2. Click **Next** to proceed through the configuration wizard.



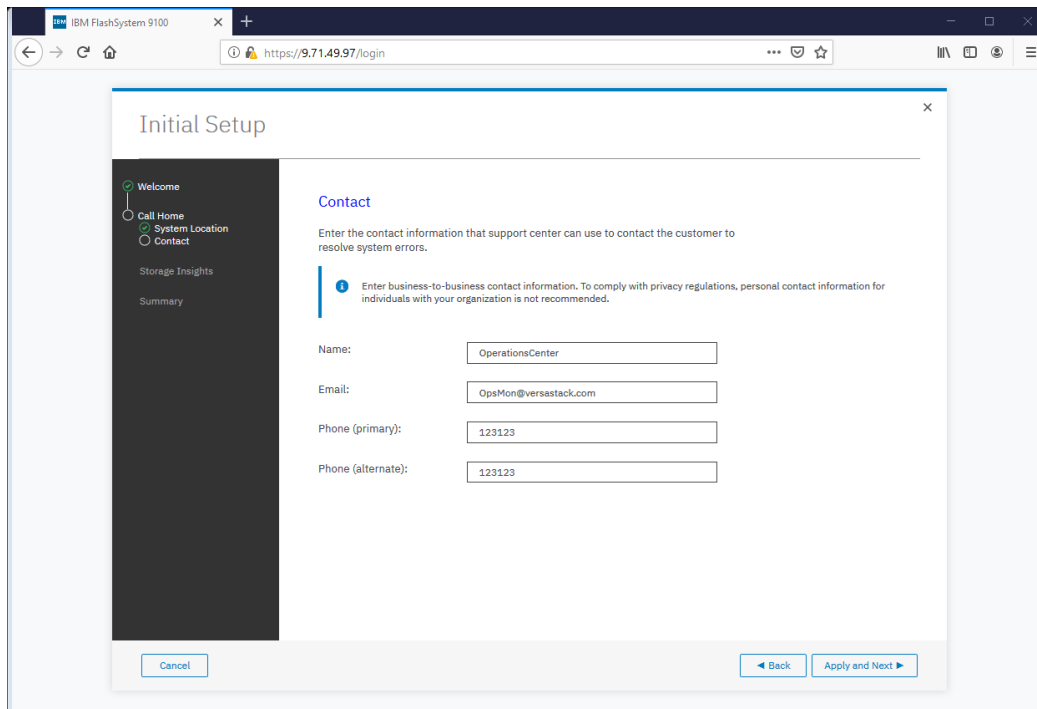
3. For optimal configuration, check the box to enable the Call Home feature.



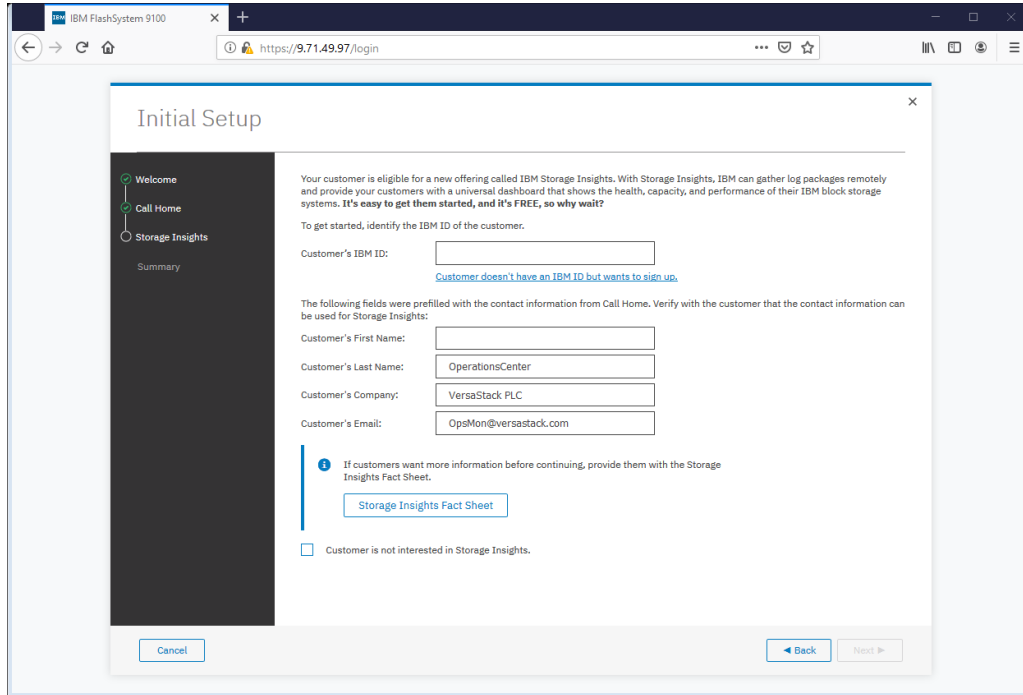
4. Detail the System Location.



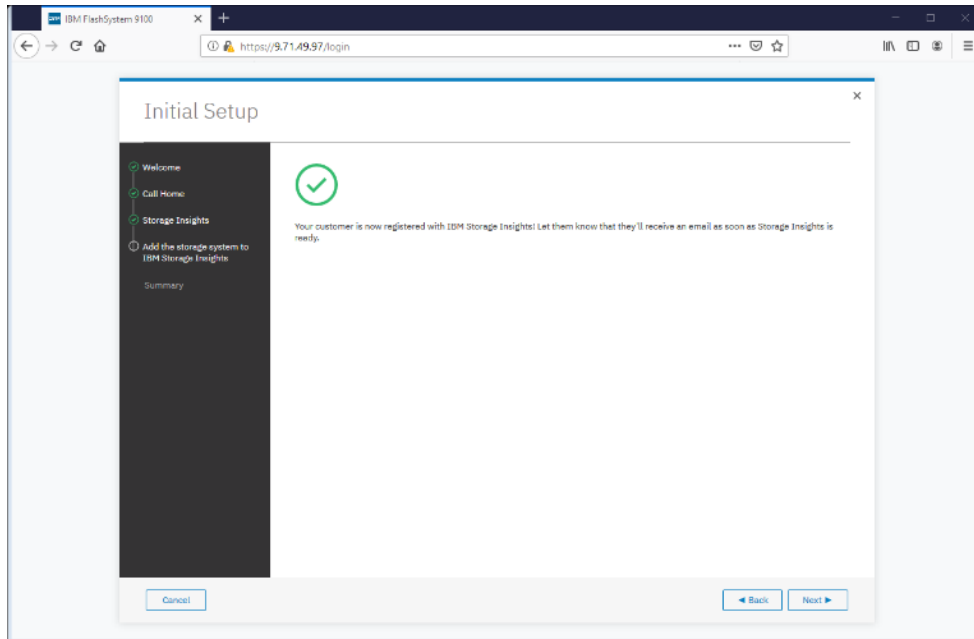
5. Specify the contact details.



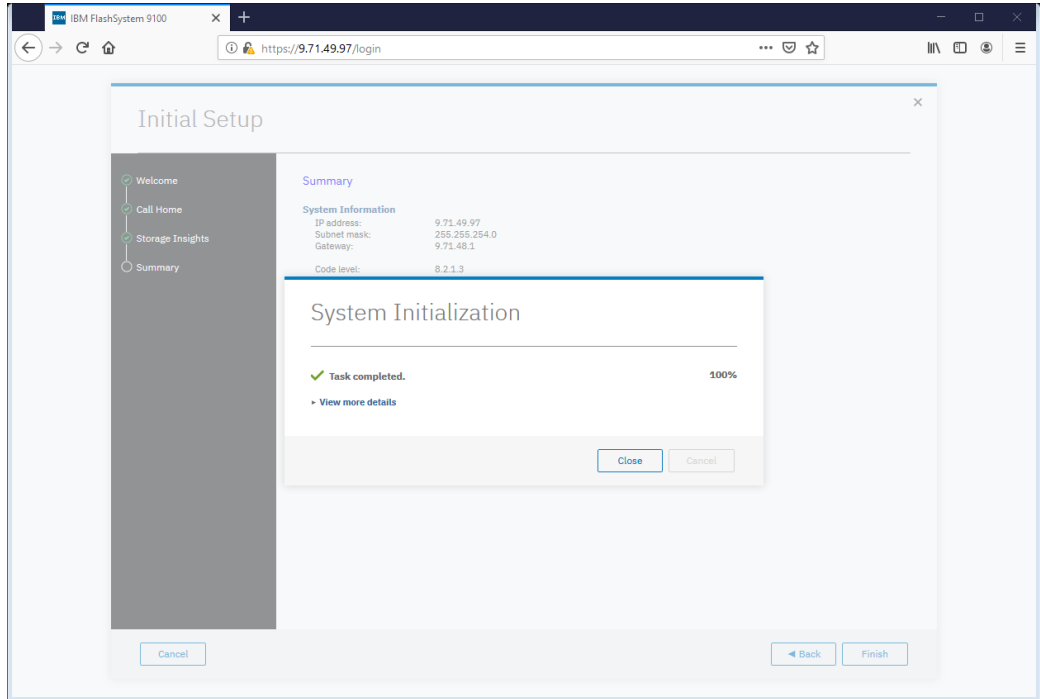
6. Specify the customer's IBM ID and contact details.



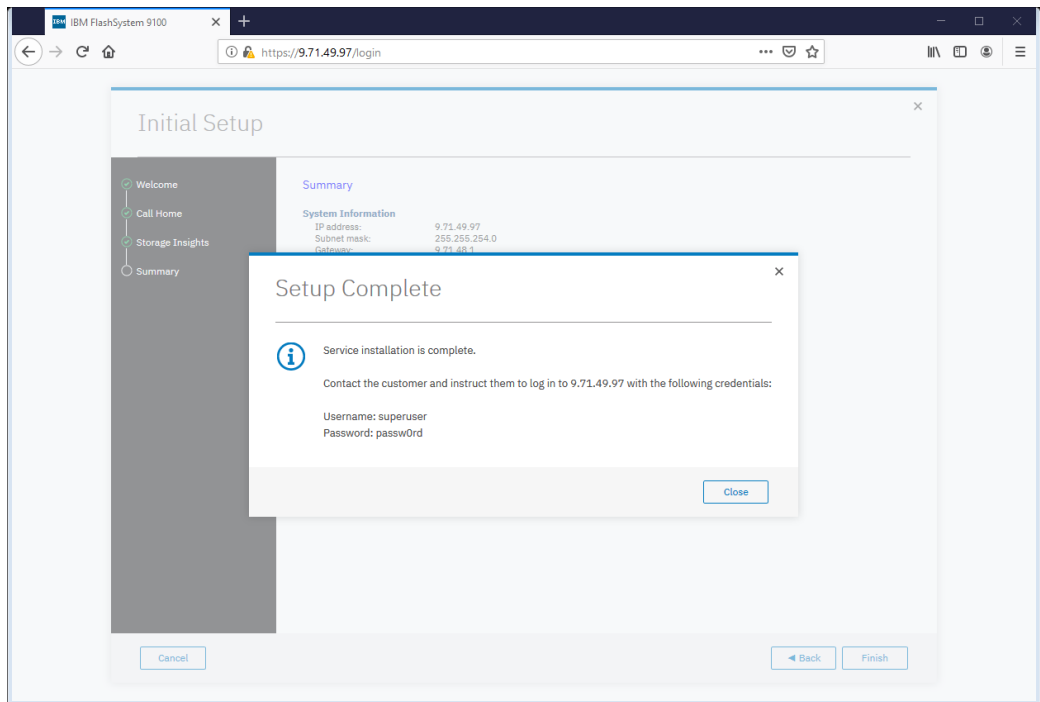
7. Click Next to finalize the IBM Storage Insights registration.



8. Review the Initial Setup summary and click **Finish**.



9. Click **Close** to complete the Service Initialization.



Customer Configuration Setup Tasks via the GUI

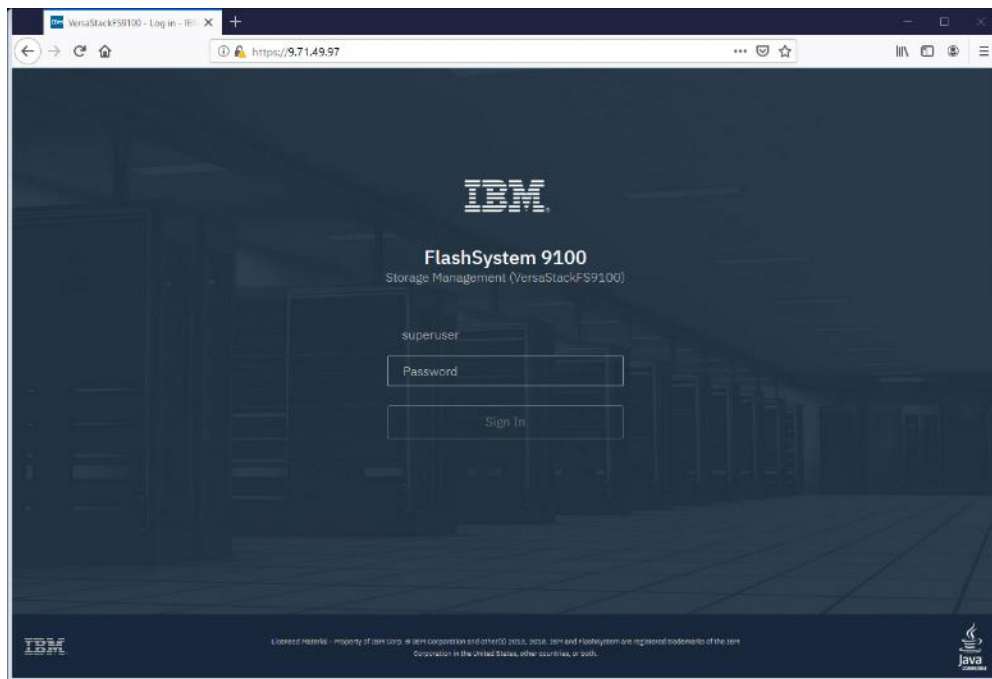
After completing the initial tasks above, launch the management GUI and continue configuring the IBM FlashSystem 9100.

To configure the customer's tasks, follow these steps:




Following e-Learning module introduces the IBM FlashSystem 9100 management interface and provides an overview of the system setup tasks, including configuring the system, migrating and configuring storage, creating hosts, creating and mapping volumes, and configuring email notifications: [Getting Started](#)

1. Log into the management GUI using the cluster IP address configured above.

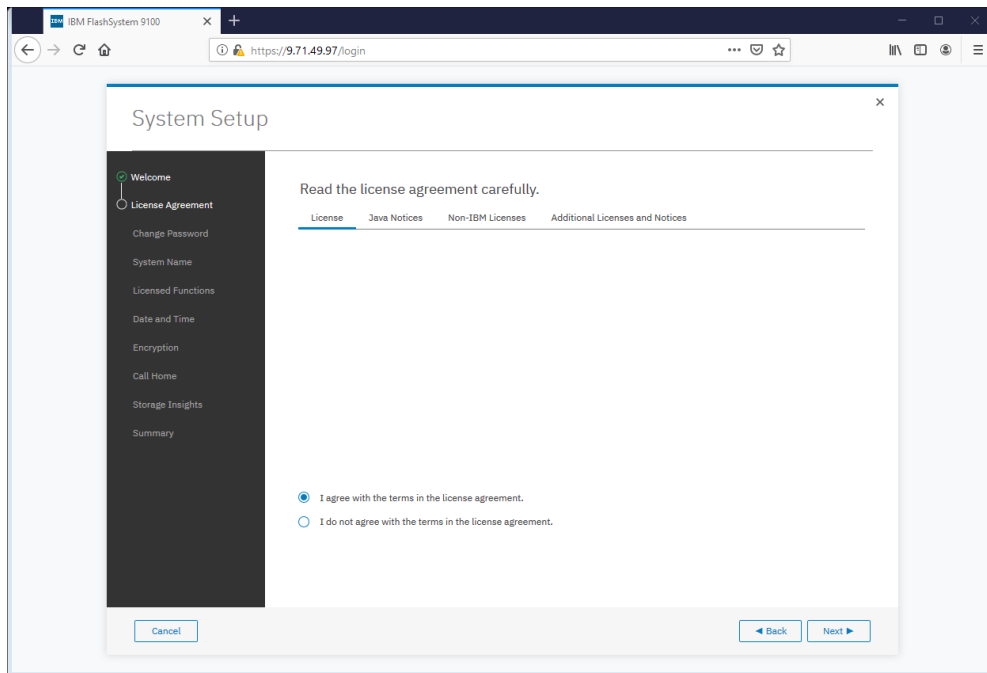


2. Log in using the default credentials:

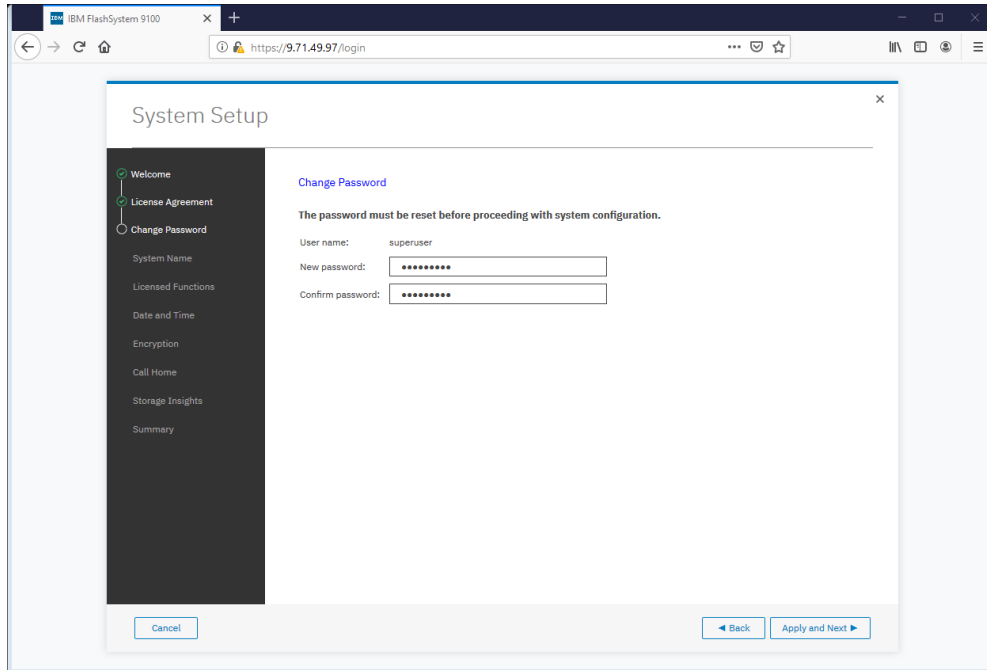
Username: `superuser`
Password: `passw0rd`
3. Click **Next** to skip the Welcome message.



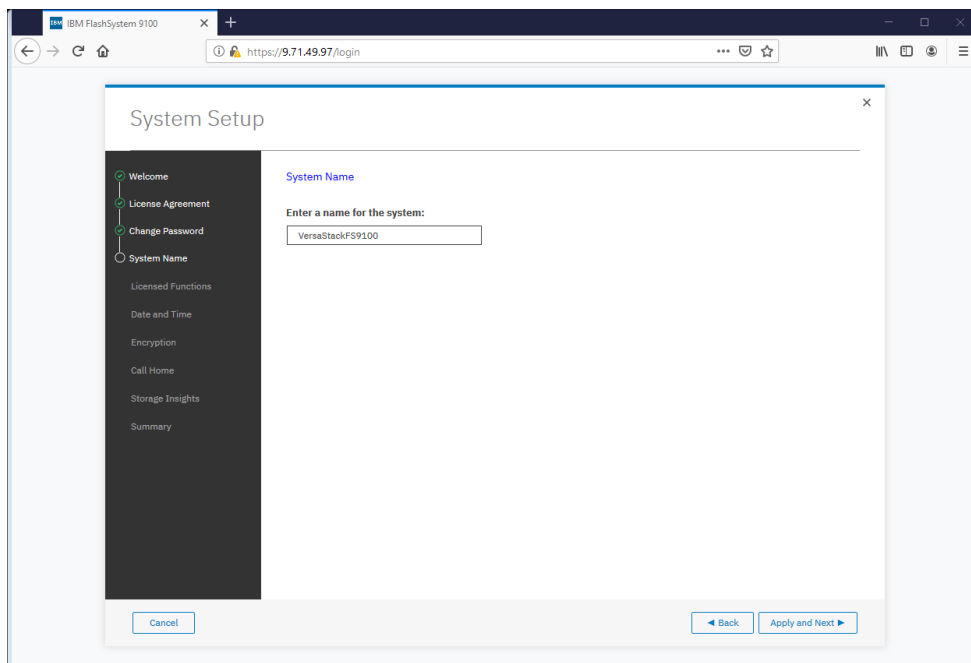
4. Read and accept the license agreement. Click **Accept**.



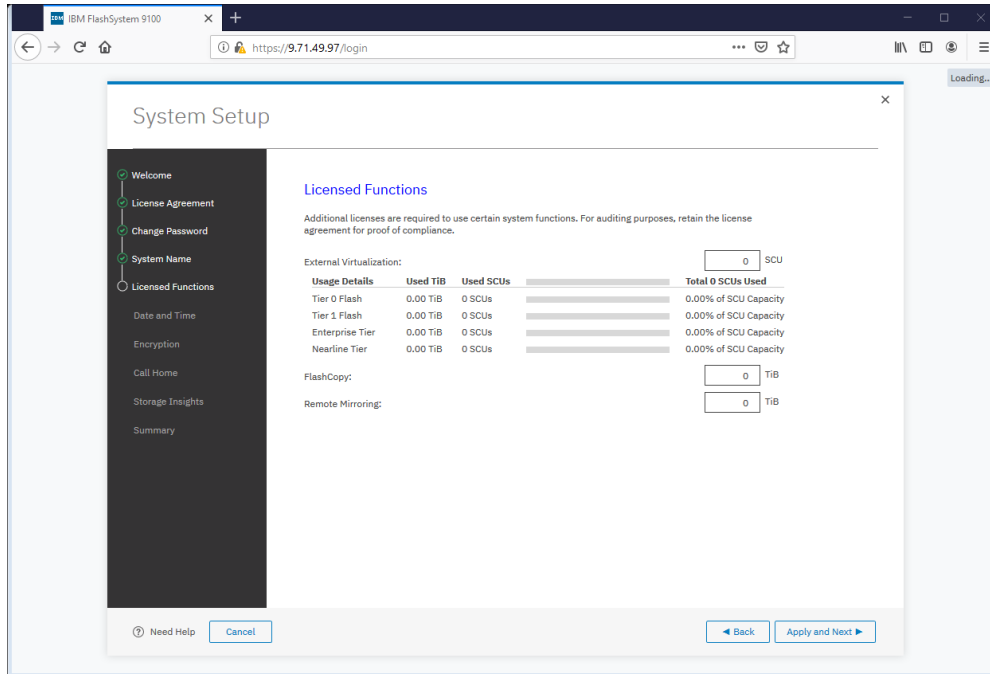
5. Define new credentials for the superuser user account.



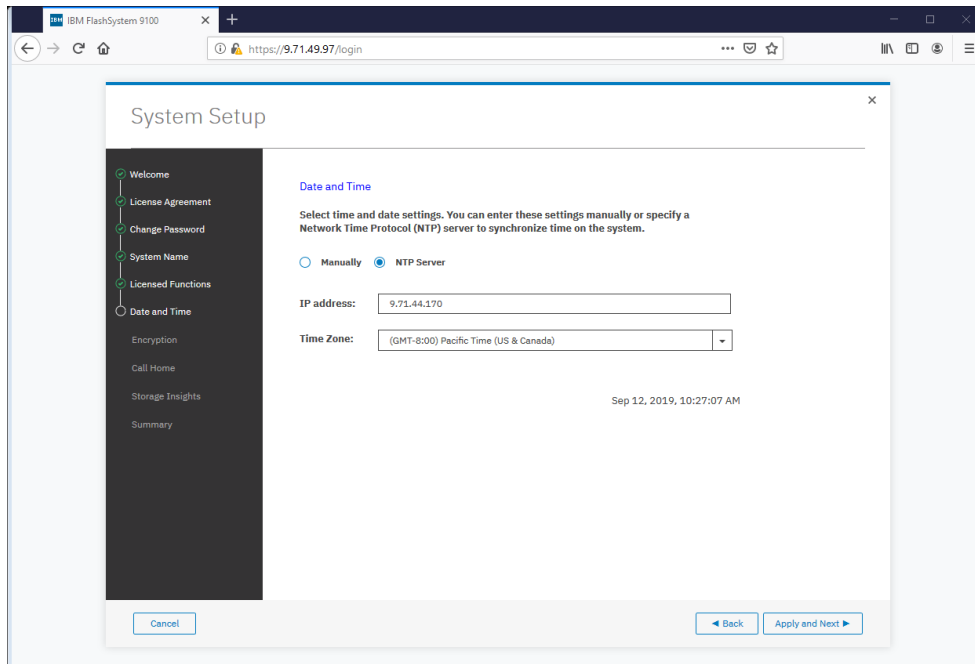
6. Enter the System Name and click **Apply and Next** to proceed.



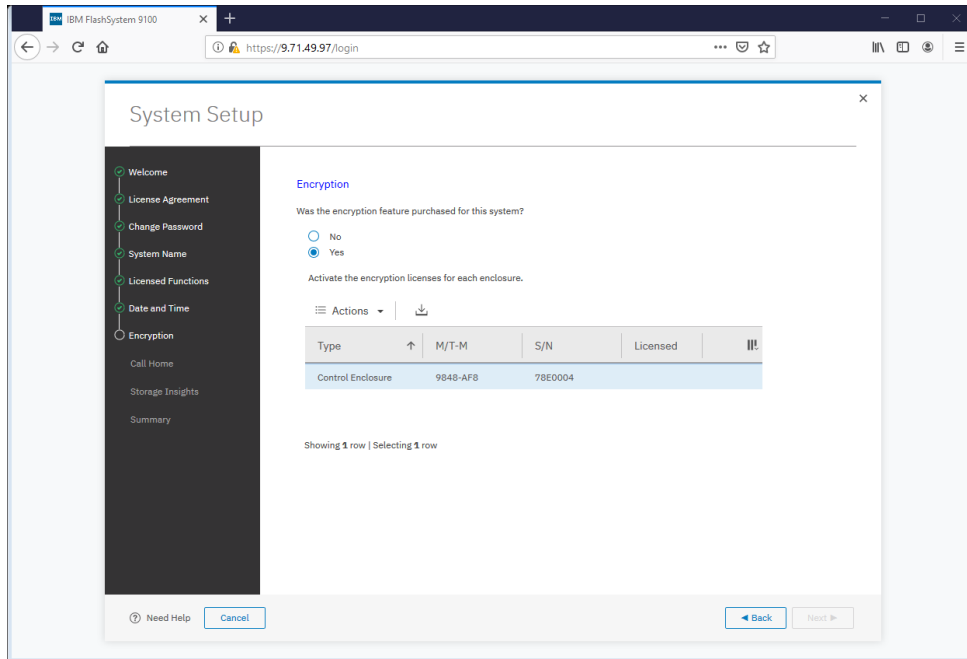
7. Enter the license details that was purchased for FlashCopy, Remote Mirroring, Easy Tier, and External Virtualization. Click **Apply and Next** to proceed.



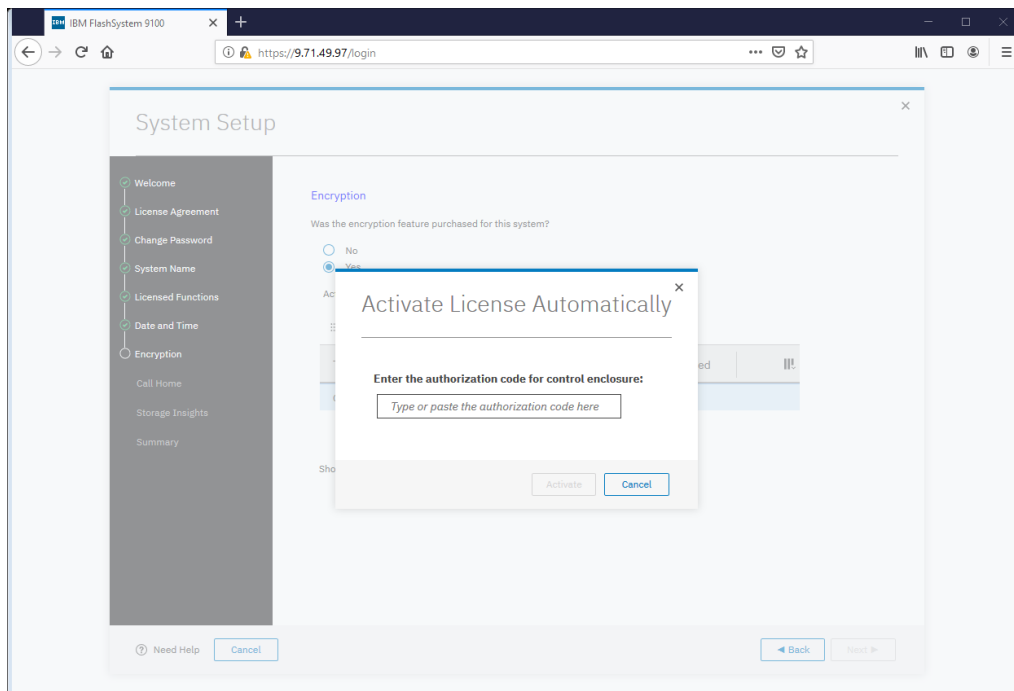
8. Configure the date and time settings, inputting NTP server details if available. Click **Apply and Next** to proceed.

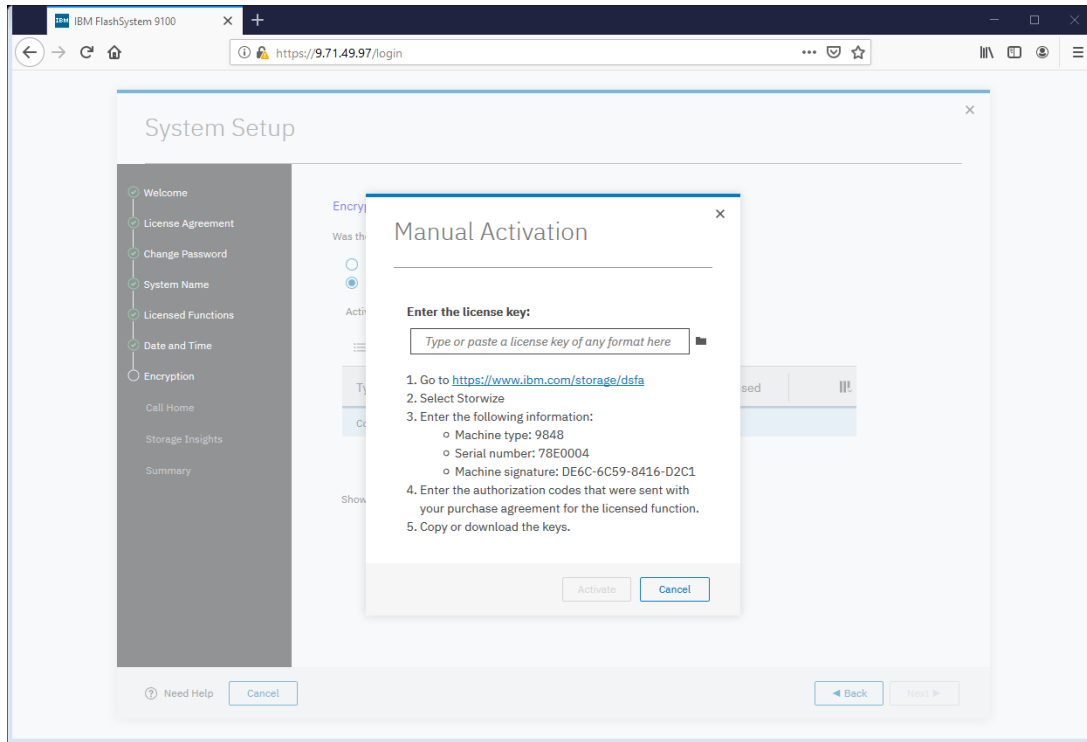


9. Enable the Encryption feature (or leave it disabled). Click **Next** to proceed.



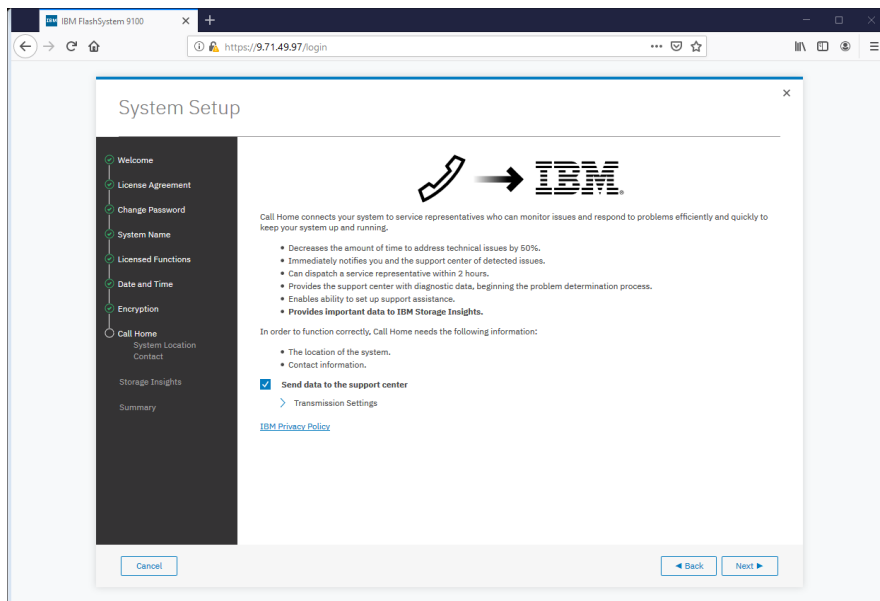
10. If using the encryption, select either Manual or Automatic activation and enter the authorization code or license key accordingly.







It is highly recommended to configure email event notifications which will automatically notify IBM support centers when problems occur.


11. Enter the complete company name and address and then click **Next**.



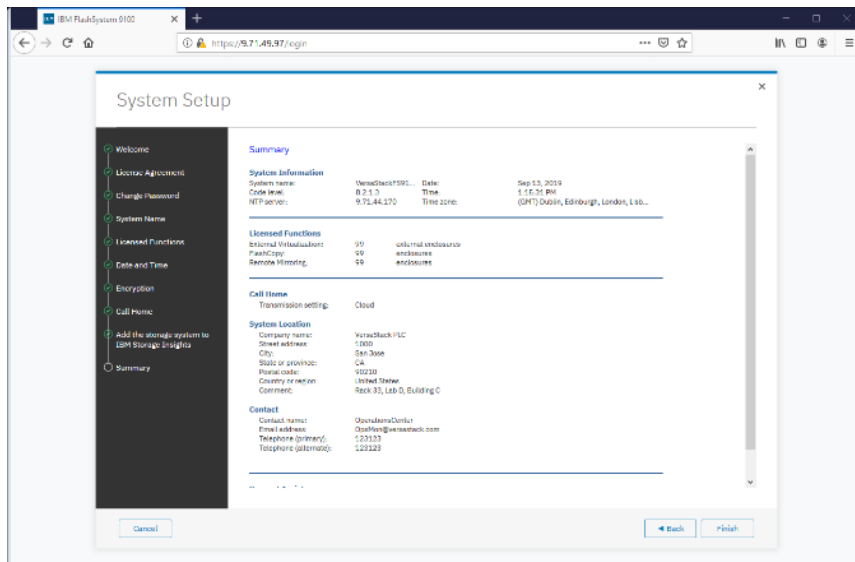
12. Enter the contact person for the support center calls. Click **Apply** and **Next**.



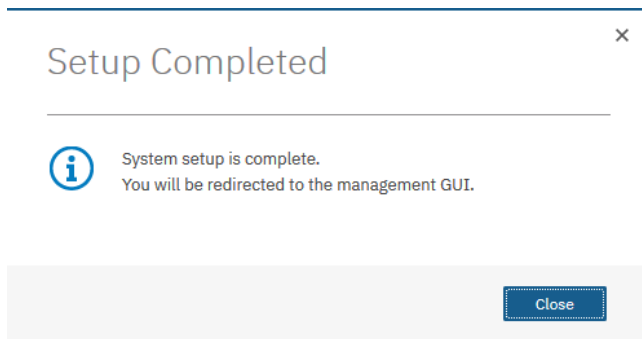
IBM Storage Insights is required to enable performance/health monitoring required by remote IBM Support representatives when assisting with any support issues.



13. Review the final summary page and click **Finish** to complete the System Setup wizard.



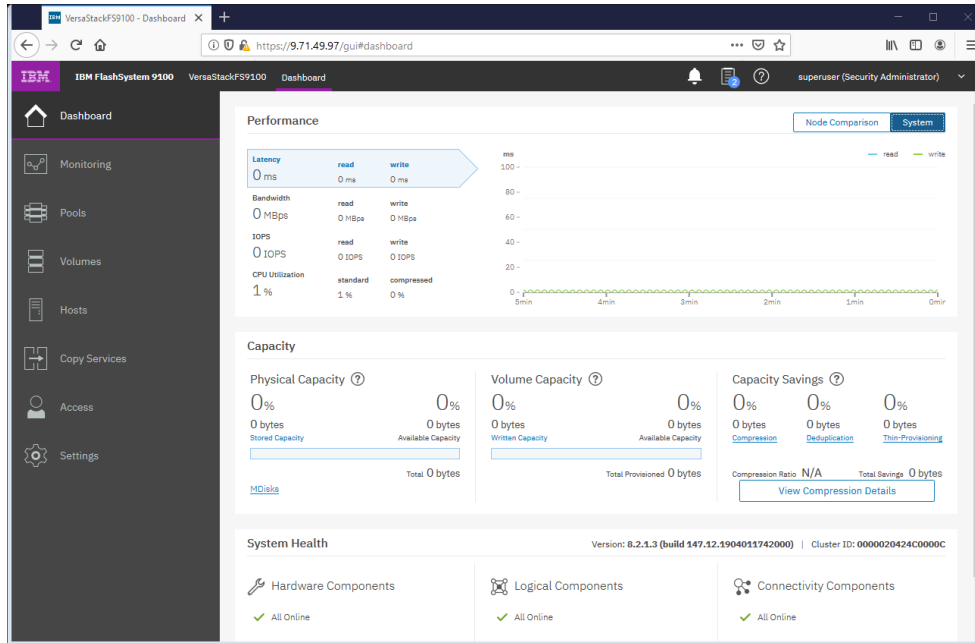
14. Setup Completed. Click **Close**.



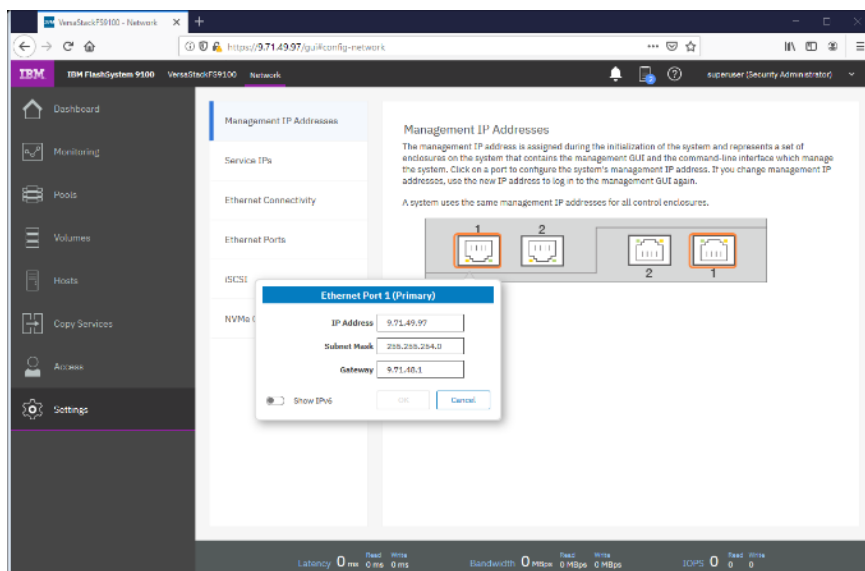
System Dashboard, and Post-Initialization Setup Tasks

To configure the necessary post-initialization setup tasks, follow these steps:

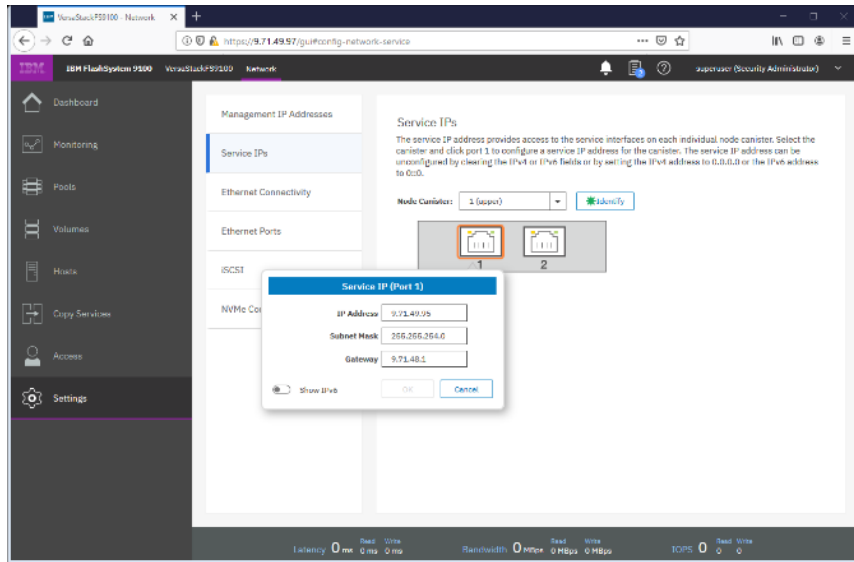
1. The System view of IBM FS9100 is now available, as shown below.



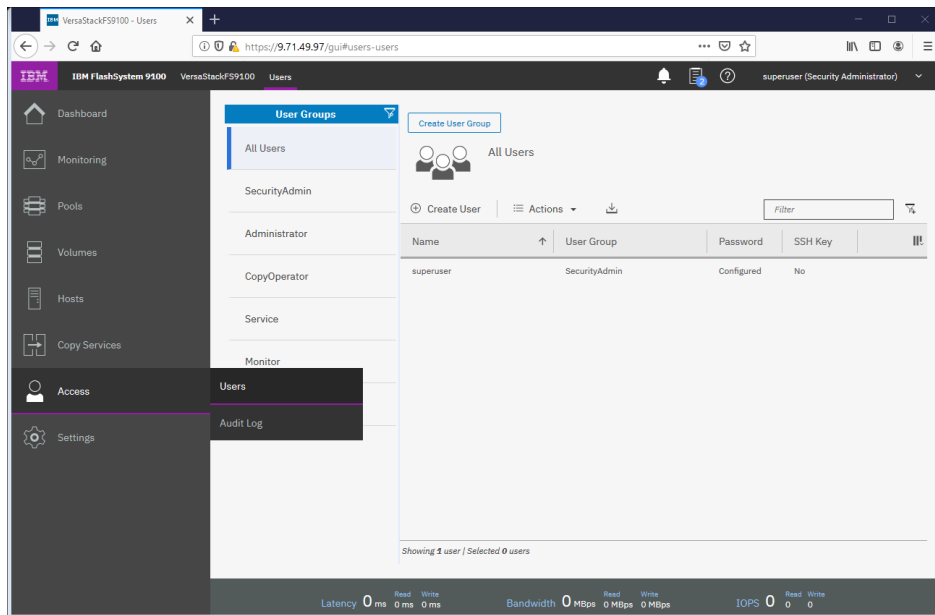
2. In the left side menu, hover over each of the icons on the Navigation Dock to become familiar with the options.
3. Verify the configured Management IP Address (Cluster IP) and configuring Service Assistant IP addresses for each node canister in the system.
4. On the Network screen, highlight the Management IP Addresses section. Then click the number 1 interface on the left-hand side to bring up the Ethernet port IP menu. If required, change the IP address if necessary and click OK.



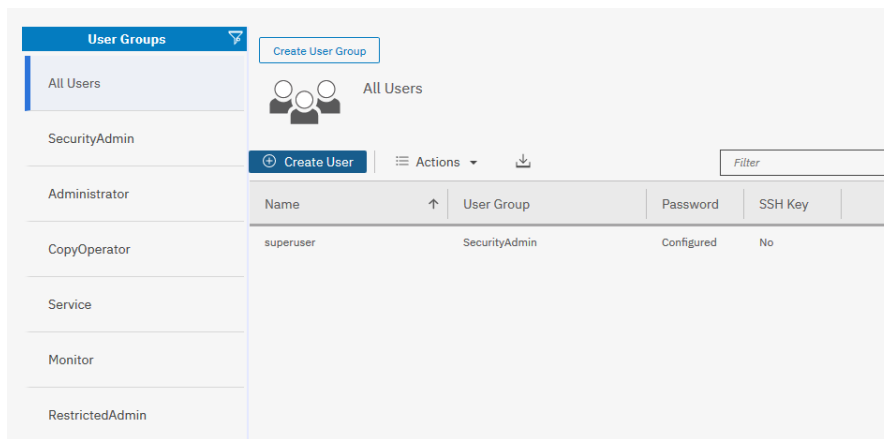
- While still on the Network screen, select 1) Service IP Addresses from the list on the left and select each Node Canister Upper/Lower in turn, and change the IP address for port 1, click OK.



- Repeat this process for port 1 on the other Node Canisters.



- Click the Access icon from the Navigation Dock on the left and select Users to access the Users screen.



8. Select Create User.

Create User

Name: VSAdmin

Authentication Mode: Local Remote

User Group: SecurityAdmin

Local Credentials

Users must have a password, an SSH public key, or both.

Password: [dots] Verify password: [dots]

SSH Public Key: No file selected.

- Enter a new name for an alternative admin account. Leave the `SecurityAdmin` default as the User Group, and input the new password, then click Create. Optionally, an SSH Public Key generated on a Unix server through the command `ssh-keygen -t rsa` can be copied to a public key file and associated with this user through the **Choose File** button.



Consider using Remote Authentication (via LDAP) to centrally manage authentication, roles, and responsibilities. For more information on Remote Authentication, refer to Redbook: [Implementing the IBM System Storage SAN Volume Controller with IBM Spectrum Virtualize V8.2.1](#).

Create Storage Pools and Allocate Storage

Typically, the NVMe drives within the FlashSystem 9100 enclosure are grouped together into a Distributed RAID array (sometimes referred to as a Managed Disk or mdisk) and are added to a storage resource called a Storage Pool (sometimes referred to as Managed Disk Group or mdiskgrp). Volumes are then created within this storage pool and presented to the

host(s) within the UCS chassis. Data from a UCS host is striped across multiple drives for performance, efficiency and redundancy.

Data Reduction Pools, SCSI UNMAP, and Data Deduplication

If enabling Data reduction on the pool during creation, the pool will be created as a Data Reduction Pool (DRP). Data Reduction Pools are a new type of storage pool, implementing techniques such as thin-provisioning, compression, and deduplication to reduce the amount of physical capacity required to store data.

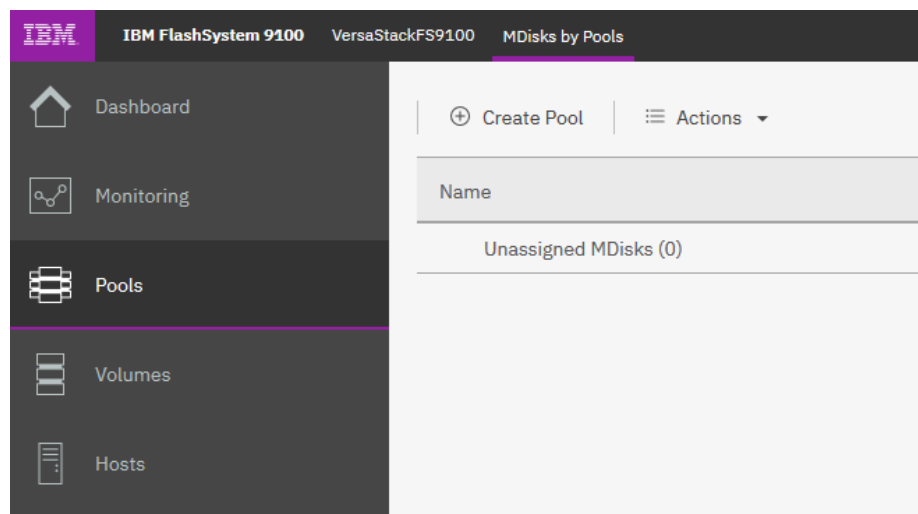
When using modern operating systems that support SCSI UNMAP, the storage pool also enables the automatic de-allocation and reclaim capacity occupied by deleted data and, for the first time, enable this reclaimed capacity to be reused by other volumes in the pool.

Data deduplication is one of the methods of reducing storage needs by eliminating redundant copies of data. Existing or new data is categorized into chunks that are examined for redundancy. If duplicate chunks are detected, then pointers are shifted to reference a single copy of the chunk, and the duplicate data sets are then released.

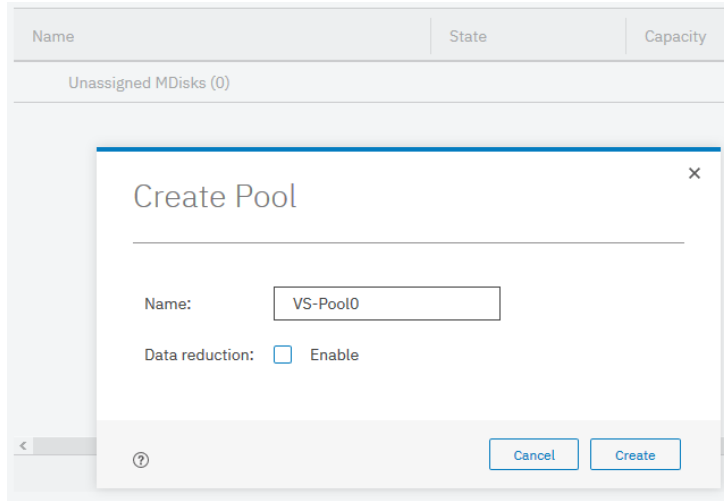
Deduplication has several benefits, such as storing more data per physical storage system, saving energy by using fewer disk drives, and decreasing the amount of data that must be sent across a network to another storage for backup replication and for disaster recovery.

However, these data savings come at a cost. There is a performance overhead when using DRPs when compared to traditional storage pools. And a percentage of the capacity of a DRP is reserved for system usage. For more information on Data Reduction Pools and techniques, refer to the Redbook publication: [Implementing the IBM System Storage SAN Volume Controller](#).


1. Select **Pools** from the Navigation Dock and select **MDisk by Pools**.



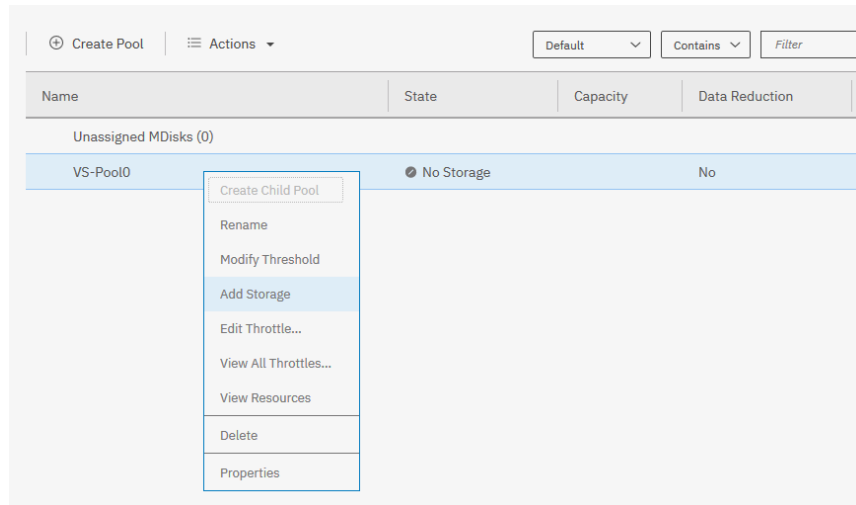
2. Click **Create Pool** and enter the name of the new storage pool. Click **Create**.



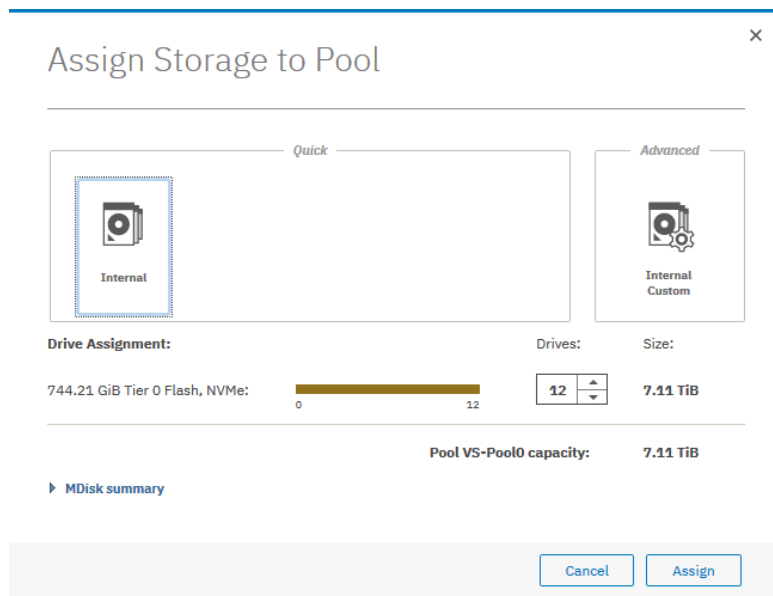
3. Identify the available drives along the bottom of the window




4. Right-click the new Pool and select **Add Storage**.



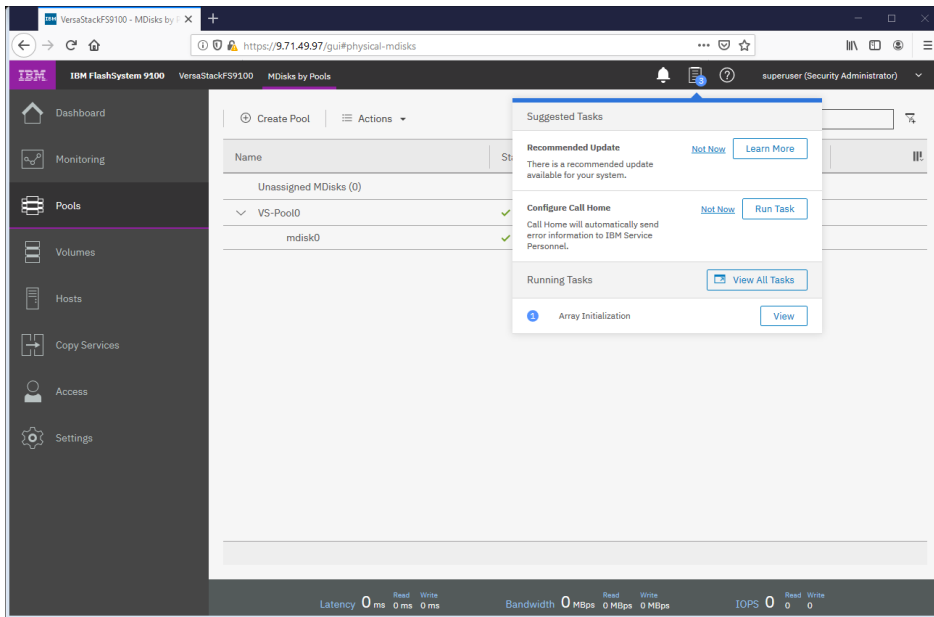
- 5. Select **Internal** to utilize drives within the enclosure, rather than from externally virtualized storage controllers.




- 6. The Managed Disk (`mdisk`) has now been created and allocated to the storage pool.



7. Reference the Running Tasks window to monitor the array initialization.



During the initialization, the array performance will be sub-optimal. Where possible, wait for the array initialization to complete before running resource intensive workloads.

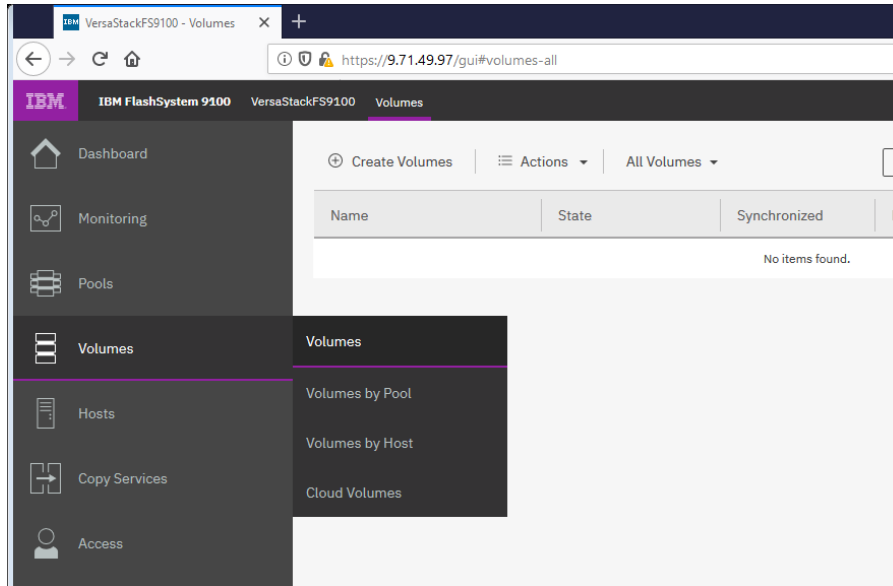


8. Select **Internal**, review the drive assignments and then select **Assign**.

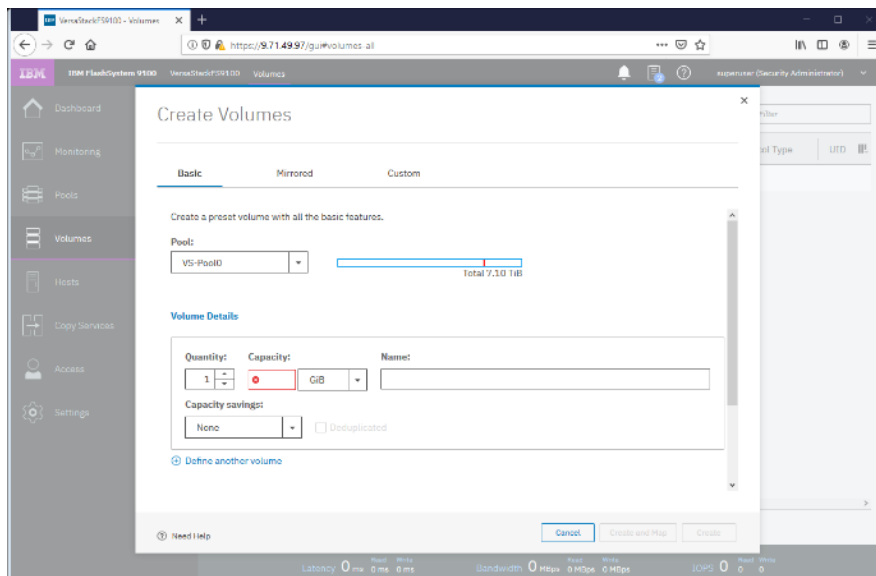


Depending on customer configuration, select Internal Custom to manually create tired storage pools grouping together disk by capabilities. In this deployment, Flash and Enterprise class disk are utilized for Silver pool and Near-line disks are utilized for Bronze storage pool.

9. Validate the pools are online and have the relevant storage assigned.
10. Select Volumes from the Navigation Dock and then select **Volumes**.



11. Click Create Volumes.



12. Define the volume characteristics, paying attention to any capacity saving, and/or high availability requirements, and specify a friendly name. Click **Create**.
13. Validate the created volumes.



Creating volumes will be explained in following sections of this document

IBM FS9100 iSCSI Configuration (iSCSI Deployment)



Cisco UCS configuration requires information about the iSCSI IQNs on IBM FS9100. Therefore, as part of the initial storage configuration, iSCSI ports are configured on IBM FS9100



This configuration step can be skipped if the UCS environment does not need to access storage environment using iSCSI.

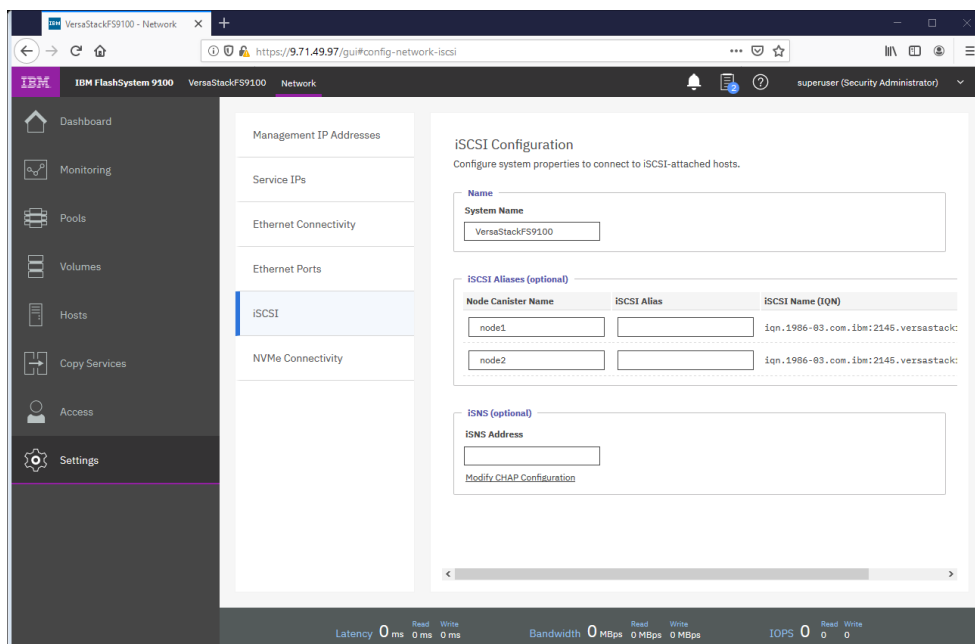
Two 25G ports from each of the IBM FS9100 node canisters are connected to each of Nexus 9336C-FX2 switches. These ports are configured as shown in Table 10 .

Table 10 IBM FS9100 iSCSI Interface Configuration

System	Port	Path	VLAN	IP address
Node canister 1	5	iSCSI-A	3161	10.29.161.249/24
Node canister 1	6	iSCSI-B	3162	10.29.162.249/24
Node canister 2	5	iSCSI-A	3161	10.29.161.250/24
Node canister 2	6	iSCSI-B	3162	10.29.162.250/24

To configure the IBM FS9100 system for iSCSI storage access, follow these steps:

1. Log into the IBM Management Interface GUI and navigate to **Settings > Network**.
2. Click the **iSCSI** icon and enter the system and node names as shown:

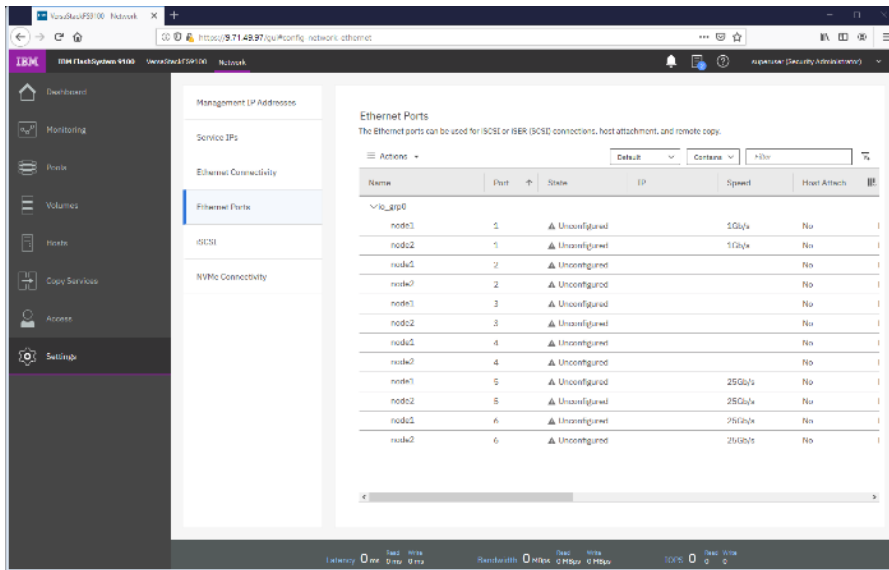


- Note the resulting iSCSI Name (IQN) in the Table 11 to be used later in the configuration procedure

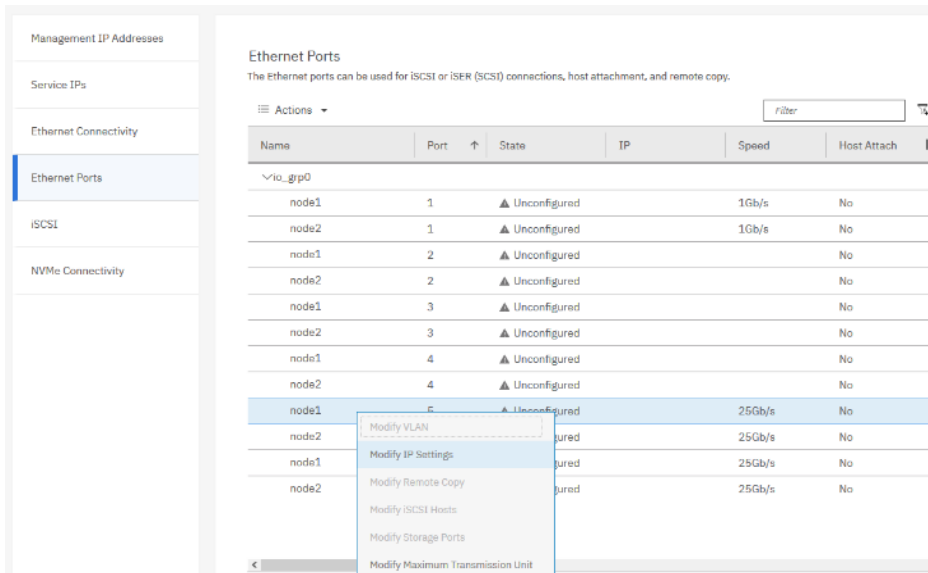
Table 11 IBM FS9100 IQN

Node	Example iSCSI name (IQN)
Node 1	iqn.1986-03.com.ibm:2145.versastack-fs9100.node1
Node 2	iqn.1986-03.com.ibm:2145.versastack-fs9100.node2

- Click the **Ethernet Ports** icon



- Click **Actions** and choose **Modify iSCSI Hosts**.



- Make sure IPv4 iSCSI hosts field is set to enable – if not, change the setting to Enabled and click **Modify**.
- If already set, click **Cancel** to close the configuration box.
- For each of the four ports listed in Table 10 Table 10 , repeat steps 1-7 .

9. Right-click the appropriate port and choose **Modify IP Settings**.
10. Enter the IP address, Subnet Mask and Gateway information in [Table 10](#) .

Modify Port 5 of Node node1 ✕

IPv4 address:

Subnet mask:

Gateway:

▶ **IPv6**

11. Click **Modify**.
12. Right-click the newly updated port and choose **Modify VLAN**.

Ethernet Ports
The Ethernet ports can be used for iSCSI or ISER (SCSI) connections, host attachment, and remote copy.

≡ Actions ▾ 🔍

Name	Port	↑	State	IP	Speed	Host Attach	!!!
▼ io_grp0							
node1	1		▲ Unconfigured		1Gb/s	No	
node2	1		▲ Unconfigured		1Gb/s	No	
node1	2		▲ Unconfigured			No	
node2	2		▲ Unconfigured			No	
node1	3		▲ Unconfigured			No	
node2	3		▲ Unconfigured			No	
node1	4		▲ Unconfigured			No	
node2	4		▲ Unconfigured			No	
node1	5		▲ Unconfigured	10.29.161.249	25Gb/s	Yes	
node2	5		▲ Unconfigured	10.29.161.250	25Gb/s	Yes	
node1	6		▲ Unconfigured	10.29.162.249	25Gb/s	Yes	
node2	6		▲ Unconfigured	10.29.162.250	25Gb/s	Yes	

- Modify VLAN
- Modify IP Settings
- Modify Remote Copy
- Modify iSCSI Hosts
- Modify Storage Ports
- Modify Maximum Transmission Unit

13. Check the box to **Enable** VLAN.

Modify VLAN for port 5 on Node 1 ✕

VLAN: Enable

VLAN tag:

Apply change to the failover port too ?

[2 ports affected](#)

? Need Help

Cancel

Modify

14. Enter the appropriate VLAN from Table 2 .



This is only needed if the VLAN is not set as native VLAN in the UCS, do not enable VLAN if the iSCSI VLAN is set as native VLAN.

15. Keep the Apply change to the failover port too check box checked.

16. Click **Modify**.

17. Repeat the steps for all for iSCSI ports listed in Table 10 .

18. Verify all ports are configured as shown below. The output below shows configuration for two FS9100 node canisters.

Management IP Addresses		Ethernet Ports																																																																																				
Service IPs		The Ethernet ports can be used for iSCSI or iSER (SCSI) connections, host attachment, and remote copy.																																																																																				
Ethernet Connectivity		Actions		Filter																																																																																		
Ethernet Ports		Name	Port	State	IP	Speed	Host Attach																																																																															
iSCSI		<div style="display: flex; align-items: center;"> ▼ io_grp0 <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th>Name</th> <th>Port</th> <th>State</th> <th>IP</th> <th>Speed</th> <th>Host Attach</th> </tr> </thead> <tbody> <tr> <td>node1</td> <td>1</td> <td>▲ Unconfigured</td> <td></td> <td>1Gb/s</td> <td>No</td> </tr> <tr> <td>node2</td> <td>1</td> <td>▲ Unconfigured</td> <td></td> <td>1Gb/s</td> <td>No</td> </tr> <tr> <td>node1</td> <td>2</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node2</td> <td>2</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node1</td> <td>3</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node2</td> <td>3</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node1</td> <td>4</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node2</td> <td>4</td> <td>▲ Unconfigured</td> <td></td> <td></td> <td>No</td> </tr> <tr> <td>node1</td> <td>5</td> <td>✓ Configured</td> <td>10.29.161.249</td> <td>25Gb/s</td> <td>Yes</td> </tr> <tr> <td>node2</td> <td>5</td> <td>✓ Configured</td> <td>10.29.161.250</td> <td>25Gb/s</td> <td>Yes</td> </tr> <tr> <td>node1</td> <td>6</td> <td>✓ Configured</td> <td>10.29.162.249</td> <td>25Gb/s</td> <td>Yes</td> </tr> <tr> <td>node2</td> <td>6</td> <td>✓ Configured</td> <td>10.29.162.250</td> <td>25Gb/s</td> <td>Yes</td> </tr> </tbody> </table> </div>							Name	Port	State	IP	Speed	Host Attach	node1	1	▲ Unconfigured		1Gb/s	No	node2	1	▲ Unconfigured		1Gb/s	No	node1	2	▲ Unconfigured			No	node2	2	▲ Unconfigured			No	node1	3	▲ Unconfigured			No	node2	3	▲ Unconfigured			No	node1	4	▲ Unconfigured			No	node2	4	▲ Unconfigured			No	node1	5	✓ Configured	10.29.161.249	25Gb/s	Yes	node2	5	✓ Configured	10.29.161.250	25Gb/s	Yes	node1	6	✓ Configured	10.29.162.249	25Gb/s	Yes	node2	6	✓ Configured	10.29.162.250	25Gb/s	Yes
Name	Port	State	IP	Speed	Host Attach																																																																																	
node1	1	▲ Unconfigured		1Gb/s	No																																																																																	
node2	1	▲ Unconfigured		1Gb/s	No																																																																																	
node1	2	▲ Unconfigured			No																																																																																	
node2	2	▲ Unconfigured			No																																																																																	
node1	3	▲ Unconfigured			No																																																																																	
node2	3	▲ Unconfigured			No																																																																																	
node1	4	▲ Unconfigured			No																																																																																	
node2	4	▲ Unconfigured			No																																																																																	
node1	5	✓ Configured	10.29.161.249	25Gb/s	Yes																																																																																	
node2	5	✓ Configured	10.29.161.250	25Gb/s	Yes																																																																																	
node1	6	✓ Configured	10.29.162.249	25Gb/s	Yes																																																																																	
node2	6	✓ Configured	10.29.162.250	25Gb/s	Yes																																																																																	
NVMe Connectivity																																																																																						

Modify Interface MTU

Use the `cfgportip` CLI command to set Jumbo Frames (MTU 9000). The default value of port MTU is 1500. An MTU of 9000 (jumbo frames) provides improved CPU utilization and increased efficiency by reducing the overhead and increasing the size of the payload.

To modify the interface MTU, follow these steps:

1. The MTU configuration can be verified using the command:

```
FS9100info lspportip <port number> | grep mtu
```

2. SSH to the IBM FS9100 management IP address and use following CLI command to set the MTU for ports 5 and 6 in the FS9100 in iogrp 0:

```
FS9100task cfgportip -mtu 9000 -iogrp 0 5  
FS9100task cfgportip -mtu 9000 -iogrp 0 6
```

This completes the initial configuration of the IBM systems. The next section explains the Cisco UCS configuration.

Cisco UCS Server Configuration

This section explains the Cisco UCS setup for VersaStack infrastructure. This section includes setup for both iSCSI as well as FC SAN boot and storage access.



If a customer environment does require implementing some of the storage protocols explained in this deployment guide, the relevant configuration sections can be skipped.

Cisco UCS Initial Configuration

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

Cisco UCS 6454 A

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

1. Connect to the console port on the first Cisco UCS 6454 fabric interconnect.

```

Enter the configuration method. (console/gui) ? console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? setup
You have chosen to setup a new Fabric interconnect? Continue? (y/n): y
Enforce strong password? (y/n) [y]: y
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]: yes
Which switch fabric (A/B)[]: A
Enter the system name: <Name of the System>
Physical Switch Mgmt0 IP address: <Mgmt. IP address for Fabric A>
Physical Switch Mgmt0 IPv4 netmask: <Mgmt. IP Subnet Mask>
IPv4 address of the default gateway: <Default GW for the Mgmt. IP >
Cluster IPv4 address: <Cluster Mgmt. IP address>
Configure the DNS Server IP address? (yes/no) [n]: y
DNS IP address: <DNS IP address>
Configure the default domain name? (yes/no) [n]: y
Default domain name: <DNS Domain Name>
Join centralized management environment (UCS Central)? (yes/no) [n]: n
Apply and save configuration (select no if you want to re-enter)? (yes/no): yes

```

2. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS 6454 B

To configure the second Cisco UCS Fabric Interconnect for use in a VersaStack environment, follow these steps:

1. Connect to the console port on the second Cisco UCS 6454 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 for the peer Fabric interconnect: <Admin Password>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <Address provided in last step>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <Mask provided in last step>
Cluster IPv4 address          : <Cluster IP provided in last step>
Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address
Physical switch Mgmt0 IP address: < Mgmt. IP address for Fabric B>
Apply and save the configuration (select no if you want to re-enter)?
(yes/no): yes

```

2. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS Setup

Log into Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6454 fabric interconnect cluster address.
2. Click the **Launch UCS Manager link** to launch the Cisco UCS Manager User Interface.
3. When prompted, enter admin as the username and enter the administrative password.
4. Click **Login** to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 4.0(4c)

This document assumes the use of Cisco UCS 4.0(4c). To upgrade the Cisco UCS Manager software and the UCS 6454 Fabric Interconnect software to version 4.0(4c), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To enable anonymous reporting, follow this step:

1. In the Anonymous Reporting window, select whether to send anonymous data to Cisco for improving future products. If you select **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

Don't show this message again.

OK

Cancel

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 the **Admin** tab in the navigation pane on left.
2. Select 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 **OK** to complete configuring Call Home.

The screenshot shows the Cisco UCS Manager interface for configuring Call Home. The left navigation pane is open to 'Communication Management > Call Home'. The main content area is titled 'Communication Management / Call Home' and has several tabs: 'General', 'Profiles', 'Call Home Policies', 'System Inventory', 'Anonymous Reporting', 'Events', and 'FSM'. The 'General' tab is active, and the 'Admin' section is expanded. The 'State' is set to 'On', 'Switch Priority' is 'Debugging', and 'Throttling' is 'On'. Below this are sections for 'Contact Information', 'Ids', 'Email Addresses', and 'SMTP Server', each with input fields for various details.

Communication Management / Call Home

General Profiles Call Home Policies System Inventory Anonymous Reporting Events FSM

Admin

State : Off On

Switch Priority : Debugging

Throttling : Off On

States

Contact Information

Contact :

Phone :

Email :

Address :

Ids

Customer ID :

Contract ID :

Site ID :

Email Addresses

From :

Reply To :

SMTP Server

Host: (IP Address or Hostname) :

Port : 25

Add a Block of Management IP Addresses for KVM Access

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

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
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, the number of IP addresses required, and the subnet and gateway information. Click **OK**.

Create Block of IPv4 Addresses



From :	<input type="text" value="192.168.163.181"/>	Size :	<input type="text" value="20"/>
Subnet Mask :	<input type="text" value="255.255.252.0"/>	Default Gateway :	<input type="text" value="192.168.160.1"/>
Primary DNS :	<input type="text" value="192.168.163.50"/>	Secondary DNS :	<input type="text" value="192.168.163.51"/>

OK Cancel



This block of IP addresses should be in the out of band management subnet.

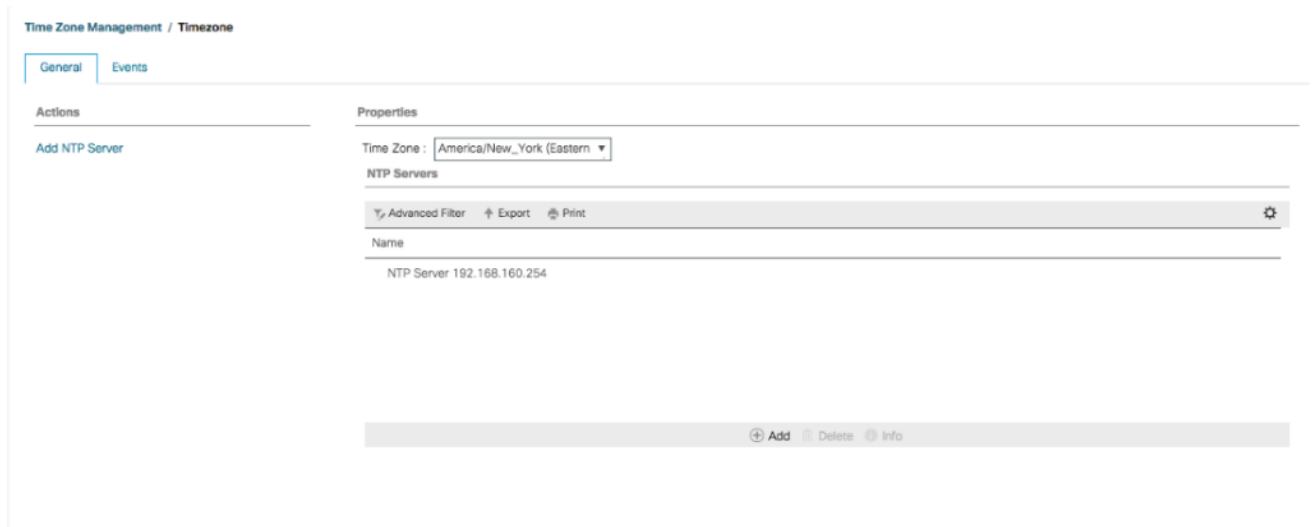
5. Click **OK**.
6. Click **OK** in the confirmation message.

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP server, follow these steps:

1. In Cisco UCS Manager, click the **Admin** tab in the navigation pane.
2. Select All > Timezone Management > Timezone.
3. In the Properties pane, select the appropriate time zone in the Timezone menu.
4. Click **Save Changes**, and then click **OK**.
5. Click Add NTP Server.

6. Enter <NTP Server IP Address> and click **OK**.
7. Click **OK**.



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. Select DNS Management.
4. In the Properties pane, select Specify DNS Server.
5. Enter the IP address of the additional DNS server.

Specify DNS Server ? ×

DNS Server (IP Address) :

6. Click **OK** and then click **OK** again. Repeat this process for any additional DNS servers.

Add an Additional Administrator User

To add an additional locally authenticated Administrative user (flexadmin) 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 select **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, select *admin*.
8. Leave Password Required selected for the *SSH Type* field.

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

Enable Port Auto-Discovery Policy

To enable the port auto-discovery policy, follow these steps:

1. Setting the port auto-discovery policy enables automatic discovery of Cisco UCS B-Series chassis server ports.
2. In Cisco UCS Manager, click Equipment, select All > Equipment in the Navigation Pane, and select the Policies tab on the right.
3. Under Port Auto-Discovery Policy, set Auto Configure Server Port to **Enabled**.

The screenshot shows the Cisco UCS Manager interface. At the top, the 'Equipment' tab is selected. Below it, a navigation bar includes 'Main Topology View', 'Fabric Interconnects', 'Servers', 'Thermal', 'Decommissioned', 'Firmware Management', 'Policies', 'Faults', and 'Diagnostics'. Under the 'Policies' tab, there are sub-tabs: 'Global Policies', 'Autoconfig Policies', 'Server Inheritance Policies', 'Server Discovery Policies', 'SEL Policy', 'Power Groups', 'Port Auto-Discovery Policy' (which is highlighted), and 'Security'. Below the navigation, there are sections for 'Actions' and 'Properties'. The 'Properties' section shows 'Owner : Local' and 'Auto Configure Server Port : Disabled Enabled'.

4. Click **Save Changes** and then **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, select All > Equipment in the Navigation Pane, and select the Policies tab on the right.
2. Under Global Policies, scroll down to Info Policy and select `Enabled` for Action.

Info Policy

Action : Disabled Enabled

3. Click **Save Changes** and then **OK**.
4. Under Equipment, select Fabric Interconnect A (primary). On the right, select the Neighbors tab. CDP information is shown under the LAN tab and LLDP information is shown under the LLDP tab.

Edit Chassis Discovery Policy

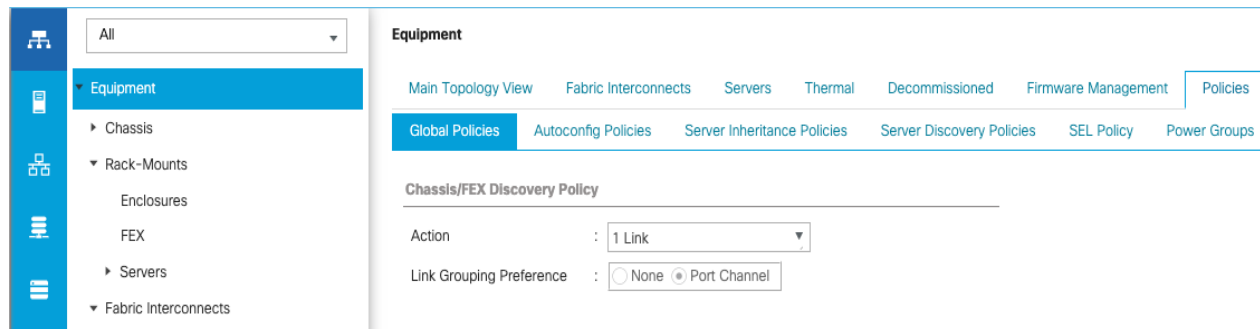
Setting the discovery policy simplifies the addition of B-Series Cisco UCS chassis and of additional fabric extenders for further C-Series connectivity. To modify the chassis discovery policy, follow these steps:

1. In Cisco UCS Manager, click the **Equipment** tab in the navigation pane and select Equipment from the list in the left pane.
2. In the right pane, click the **Policies** tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum the number of uplink ports that are cabled between any chassis IOM or fabric extender (FEX) and the fabric interconnects.



If varying numbers of links between chassis and the Fabric Interconnects will be used, leave Action set to 1 Link.

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

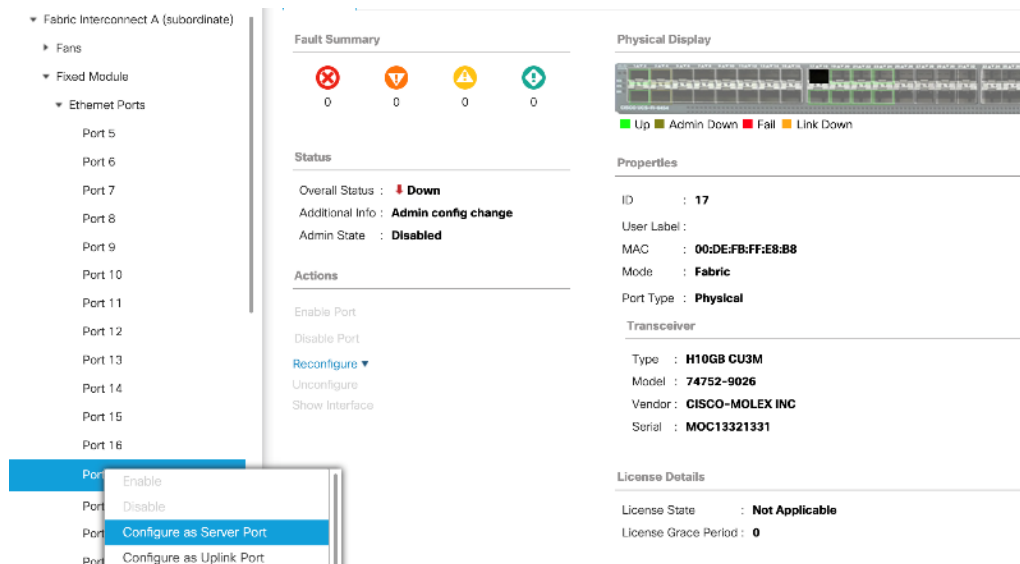


- If any changes have been made, Click **Save Changes**.
- Click **OK**.

Enable Server and Uplink Ports

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

- In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
- Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
- Expand Fixed Module.
- Expand and select **Ethernet Ports**.
- Select the ports that are connected to the Cisco UCS 5108 chassis and UCS C-Series servers, one by one, right-click and select **Configure as Server Port**.

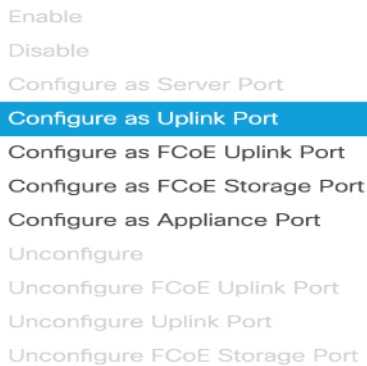


- Click **Yes** to confirm server ports and click **OK**.

- Verify that the ports connected to the UCS 5108 chassis and C-series servers are now configured as Server ports by selecting **Fabric Interconnect A** in the left and **Physical Ports** tab in the right pane.

Slot	Aggr. Port ID	Port ID	MAC	If Role	If Type	Overall Status	Admin State
1	0	17	00:DE:FB:FF:E8:B8	Server	Physical	Up	Enabled
1	0	18	00:DE:FB:FF:E8:B9	Server	Physical	Up	Enabled
1	0	19	00:DE:FB:FF:E8:BA	Server	Physical	Up	Enabled
1	0	20	00:DE:FB:FF:E8:BB	Server	Physical	Up	Enabled
1	0	21	00:DE:FB:FF:E8:BC	Server	Physical	Up	Enabled
1	0	22	00:DE:FB:FF:E8:BD	Server	Physical	Up	Enabled
1	0	23	00:DE:FB:FF:E8:BE	Server	Physical	Up	Enabled
1	0	24	00:DE:FB:FF:E8:BF	Server	Physical	Up	Enabled

- Select the ports that are connected to the Cisco Nexus 9336C-FX2 switches, one by one, right-click and select **Configure as Uplink Port**.



- Click **Yes** to confirm uplink ports and click **OK**.
- Verify that the uplink ports are now configured as Network ports by selecting **Fabric Interconnect A** in the left and **Physical Ports** tab in the right pane.
- Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
- Repeat above steps to configure server and uplink ports on Fabric Interconnect B.

Acknowledge Cisco UCS Chassis and FEX

When the UCS FI ports are configured as server ports, UCS chassis is automatically discovered and may need to be acknowledged. To acknowledge all Cisco UCS chassis, follow these steps:

- In Cisco UCS Manager, click the **Equipment** tab in the navigation pane.
- Expand **Chassis** and select each chassis that is listed.
- Right-click each chassis and select **Acknowledge Chassis**.
- Click **Yes** and then click **OK** to complete acknowledging the chassis.
- If Nexus FEXes are part of the configuration, expand Rack Mounts and FEX.

6. Right-click each FEX that is listed and select Acknowledge FEX.
7. Click **Yes** and then click **OK** to complete acknowledging the FEX.

Enable Fibre Channel Ports (FC Deployment)



The FC port and uplink configurations can be skipped if the UCS environment does not need access to IBM storage using fibre channel.

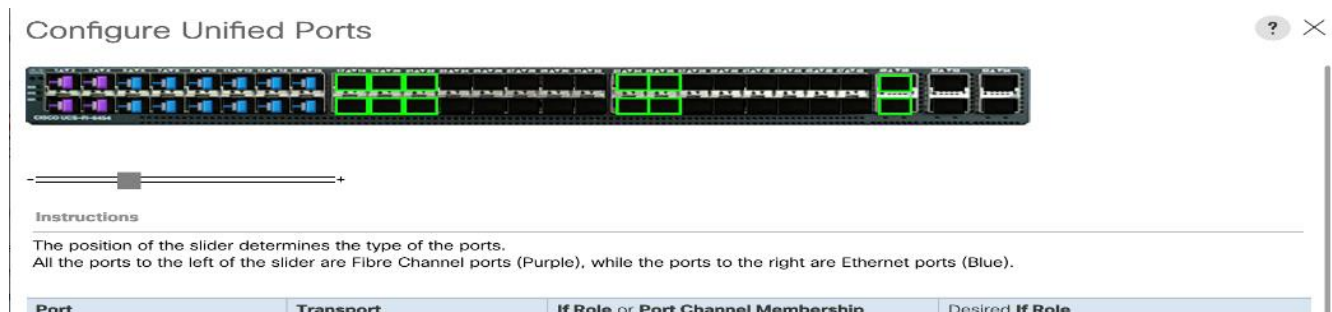
Fibre Channel port configurations differ between the 6454, 6332-16UP and the 6248UP Fabric Interconnects. All Fabric Interconnects have a slider mechanism within the Cisco UCS Manager GUI interface, but the fibre channel port selection options for the 6454 are from the first 8 ports starting from the first port and configured in increments of 4 ports from the left. For the 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 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 6454. Modify as necessary for the 6332-16UP or 6248UP.

To enable FC uplink ports, follow these steps.



This step requires a reboot. To avoid an unnecessary switchover, configure the subordinate Fabric Interconnect first.

1. In the **Equipment** tab, select the **Fabric Interconnect B** (subordinate FI in this example), and in the **Actions** pane, select **Configure Unified Ports**, and click **Yes** on the splash screen.
2. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 4 or 8 ports to be set as FC Uplinks.
3. Slide the lever to change the ports 1-4 to Fiber Channel. Click **Finish** followed by **Yes** to the reboot message. Click **OK**.



4. When the subordinate has completed reboot, repeat the procedure to configure FC ports on primary Fabric Interconnect. As before, the Fabric Interconnect will reboot after the configuration is complete.

Create VSAN for the Fibre Channel Interfaces

To configure the necessary virtual storage area networks (VSANs) for FC uplinks for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Expand the **SAN > SAN Cloud** and select **Fabric A**.

3. Right-click VSANs and choose **Create VSAN**.
4. Enter `VSAN-A` as the name of the VSAN for fabric A.
5. Keep the Disabled option selected for FC Zoning.
6. Click the **Fabric A** radio button.
7. Enter `101` as the VSAN ID for Fabric A.
8. Enter `101` as the FCoE VLAN ID for fabric A. Click **OK** twice.

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.

Enter the VSAN ID that maps to this VSAN.

VSAN ID :

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN :

9. In the **SAN** tab, expand **SAN > SAN Cloud > Fabric-B**.
10. Right-click VSANs and choose **Create VSAN**.
11. Enter `VSAN-B` as the name of the VSAN for fabric B.
12. Keep the Disabled option selected for FC Zoning.
13. Click the **Fabric B** radio button.
14. Enter `102` as the VSAN ID for Fabric B. Enter `102` as the FCoE VLAN ID for Fabric B. Click **OK** twice.

Create VSAN

Name :

FC Zoning Settings

FC Zoning : Disabled EnabledDo **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 B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID :

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN :

OK

Cancel

Create Port Channels for the Fibre Channel Interfaces

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

Fabric-A

1. In the navigation pane, under **SAN > SAN Cloud**, expand the Fabric A tree.
2. Right-click FC Port Channels and choose **Create Port Channel**.
3. Enter 1 for the port channel ID and P01 for the port channel name.
4. Click **Next** then choose ports 1-4 and click >> to add the ports to the port channel. Click **Finish**.

1 Set FC Port Channel Name

2 Add Ports

Create FC Port Channel

Port Channel Admin Speed : 4 Gbps 8 Gbps 16gbps 32gbps

Ports		
Port	Slot ID	WWPN
1	1	20:01:00:DE...
2	1	20:02:00:DE...
3	1	20:03:00:DE...
4	1	20:04:00:DE...

>>
<<

Ports in the port channel		
Port	Slot ID	WWPN
No data available		

Slot ID: 1
WWPN: 20:01:00:DE:FB:FD:40

Slot ID:
WWPN:

< Prev Next > Finish Cancel

5. Click **OK**.

6. Select FC Port-Channel 1 from the menu in the left pane and from the VSAN drop-down field, select VSAN 101 in the right pane.

SAN / SAN Cloud / Fabric A / FC Port Channels / FC Port-Channel 1 PC1

General Ports Faults Events Statistics

Status

Overall Status : ↑ **Up**

Additional Info :

Actions

Enable Port Channel

Disable Port Channel

Add Ports

Properties

ID : 1

Fabric ID : A

Port Type : **Aggregation**

Transport Type : **Fc**

Name : PC1

Description :

VSAN : A/vsan VSAN-A (101) ▼

Port Channel Admin Speed : Fabric A/vsan VSAN-A (101) 32gbps

Operational Speed(Gbps) : Fabric Dual/vsan default (1)

7. Click **Save Changes** and then click **OK**.

Fabric-B

1. Click the **SAN** tab. In the navigation pane, under **SAN > SAN Cloud**, expand the Fabric B.
2. Right-click FC Port Channels and choose **Create Port Channel**.

3. Enter 2 for the port channel ID and Po2 for the port channel name. Click **Next**.
4. Choose ports 1-4 and click >> to add the ports to the port channel.
5. Click **Finish**, and then click **OK**.
6. Select FC Port-Channel 2 from the menu in the left pane and from the VSAN drop-down list, select VSAN 102 in the right pane.
7. Click **Save Changes** and then click **OK**.



To initialize a quick sync of the connections to the MDS switch, right-click the recently created port channels, disable port channel and then re-enable the port channel.

Create Port Channels for Ethernet Uplinks

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



In this procedure, two port channels are created one from each Fabric Interconnect (A and B) to both the Cisco Nexus 9336C-FX2 switches.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Under **LAN > LAN Cloud**, expand the Fabric A tree.
3. Right-click Port Channels and choose Create Port Channel.
4. Enter 13 as the unique ID of the port channel.
5. Enter Po13 as the name of the port channel and click **Next**.

1 Set Port Channel Name

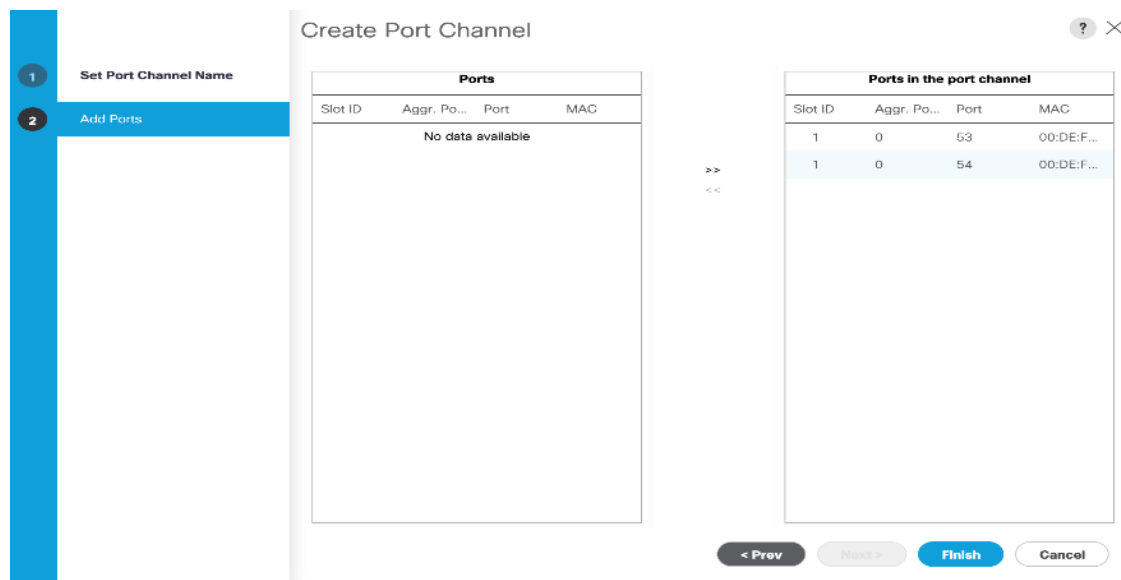
2 Add Ports

Create Port Channel

ID :

Name :

6. Select the network uplink ports to be added to the port channel.
7. Click >> to add the ports to the port channel (53 and 54 in this design).



8. Click **Finish** to create the port channel and then click **OK**.
9. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, select `Port-Channel 13`. Select 100 Gbps for the Admin Speed.
10. Click Save Changes and **OK**. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.
11. In the navigation pane, under **LAN > LAN Cloud**, expand the Fabric B tree.
12. Right-click Port Channels and choose **Create Port Channel**.
13. Enter 14 as the unique ID of the port channel.
14. Enter `Po14` as the name of the port channel and click **Next**.
15. Select the network uplink ports (53 and 54 in this design) to be added to the port channel.
16. Click >> to add the ports to the port channel.
17. Click **Finish** to create the port channel and click **OK**.
18. In the navigation pane, under LAN > LAN Cloud > Fabric B > Port Channels, select `Port-Channel 14`. Select 100 Gbps for the Admin Speed.
19. Click **Save Changes** and **OK**. After a few minutes, verify that the Overall Status is Up and the Operational Speed is correct.

Add UDLD to Uplink Port Channels

To configure the unidirectional link detection (UDLD) on the Uplink Port Channels to the Nexus switches, 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 select **Create UDLD Link Policy**.
4. If the uplink cables to the Nexus switches are copper cables, name the Policy UDLD-Aggressive and select Enabled for the Admin State and Aggressive for the Mode. If the uplink cables to the Nexus switches are fibre optic cables, name the Policy UDLD-Normal Aggressive and select Enabled for the Admin State and Aggressive for the Mode. In the validation lab configuration, UDLD-Aggressive was created.

Create UDLD Link Policy
? ×

Name :

Admin State : Enabled Disabled

Mode : Normal Aggressive

OK
Cancel

5. Click **OK**, then click **OK** again to complete creating the policy.



It is important that the Nexus switch port UDLD configurations (Aggressive or Normal) match the UCS port UDLD configurations.

6. Expand Policies > LAN Cloud > Link Profile.
7. Right-click Link Profile and select **Create Link Profile**.
8. Give the Link Profile the same name as the UDLD Link Policy above and select the UDLD Link Policy created above.

Create Link Profile
? ×

Name :

UDLD Link Policy :

OK
Cancel

9. Click **OK**, then click **OK** again to complete creating the profile.

10. In the navigation pane, under LAN > LAN Cloud > Fabric A > Port Channels, expand Port-Channel 13.
11. Select the first Eth Interface under Port-Channel 13. From the drop-down list, select the Link Profile created above, click **Save Changes** and **OK**. Repeat this process for each Eth Interface under Port-Channel 13 and for each Eth Interface under Port-Channel 14 on Fabric B.

LAN / LAN Cloud / Fabric A / Port Channels / Port-Channel 13 vPC13 / Eth Interface 1/53

General		Faults	Events
Actions		Properties	
Delete	Enable Interface	Disable Interface	
ID	Slot ID	Fabric ID	Transport Type
: 53	: 1	: A	: Ether
Port	Membership	Link Profile	User Label
: sys/switch-A/slot-1/switch-ether/port-53	: Up	: UDLD-Aggressive ▼	: <input type="text"/>



To see that UDLD is set up correctly, log into each Nexus switch and type `show udld neighbors`.

Create MAC Address Pools

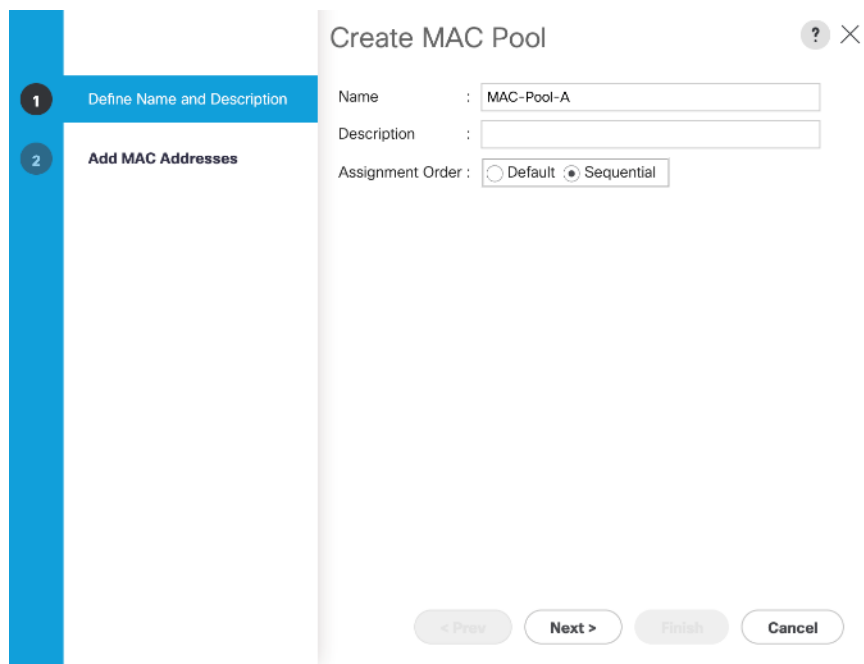
To configure the necessary MAC address pools for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select **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. Select **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. Select the option Sequential for the Assignment Order field and click **Next**.

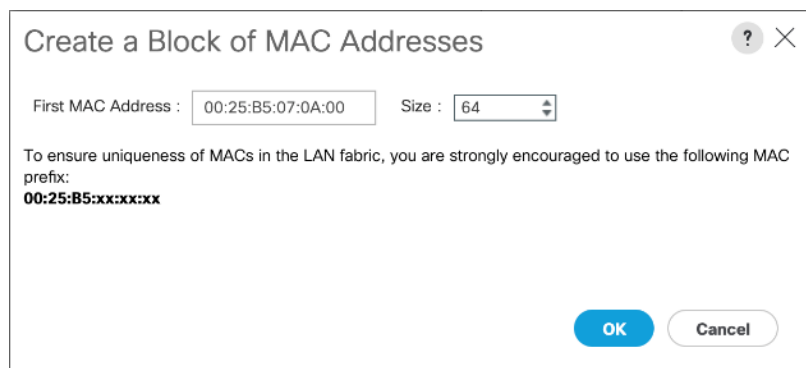


8. Click **Add**.
9. Specify a starting MAC address.



It is recommended to place oA in the second last octet of the starting MAC address to identify all of the MAC addresses as Fabric A addresses. It is also recommended to not change the first three octets of the MAC address.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or rack server resources. Remember that multiple Cisco VIC vNICs will be created on each server and each vNIC will be assigned a MAC address.



11. Click **OK** and then click **Finish**.
12. In the confirmation message, click **OK**.
13. Right-click MAC Pools under the root organization.
14. Select **Create MAC Pool** to create the MAC address pool.
15. Enter `MAC-Pool-B` as the name of the MAC pool.

16. **Optional:** Enter a description for the MAC pool.
17. Select the Sequential Assignment Order and click **Next**.
18. Click **Add**.
19. Specify a starting MAC address.



It is recommended to place 0B in the second last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. It is also recommended to not change the first three octets of the MAC address.

20. Specify a size for the MAC address pool that is sufficient to support the available blade or rack server resources.

21. Click **OK** and then click **Finish**.
22. In the confirmation message, 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 the **Servers** tab in the navigation pane.
2. Select **Pools > root**.
3. Right-click UUID Suffix Pools and choose Create UUID Suffix Pool.
4. Enter UUID-Pool as the name of the UUID suffix pool.
5. **Optional:** Enter a description for the UUID suffix pool.
6. Keep the prefix at the derived option.
7. Change the Assignment Order to Sequential.
8. Click **Next**.
9. Click **Add** to add a block of UUIDs.
10. Keep the **From** field at the default setting.

- Specify a size for the UUID block that is sufficient to support the available blade or rack server resources.

Create a Block of UUID Suffixes ? X

From : Size :

OK
Cancel

- Click **OK**. Click **Finish** and then click **OK**.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, follow these steps:

- In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Select **Pools > root**.
- Right-click Server Pools and choose **Create Server Pool**.
- Enter `Infra-Server-Pool` as the name of the server pool.
- Optional:** Enter a description for the server pool.
- Click **Next**.

1
Set Name and Description

2
Add Servers

Servers

Use...	PID	
1	UCSB-B200-M4	↓ F 2
2	UCSB-B200-M4	↓ F 2
3	UCSB-B200-M4	↓ F 2
8	UCSB-B200-M5	↑ F 1
1	UCSC-C220-M5SN	U V
2	UCSC-C220-M5SN	U V

>>
<<

Pooled Servers

Use...	PID	
4	UCSB-B200-M5	U F 1
5	UCSB-B200-M5	U F 3
6	UCSB-B200-M5	U F 3
7	UCSB-B200-M5	U F 3

Model:
 Serial Number:
 Vendor:

Model:
 Serial Number:
 Vendor:

< Prev
Next >
Finish
Cancel

94

7. Select at least two (or more) servers to be used for the setting up the VMware environment and click >> to add them to the Infra-Server-Pool server pool.
8. Click **Finish** and click **OK**.



If Cisco UCS C-Series servers are leveraged in the design, create storage pool by selecting the appropriate server models intended to be used.

Create a WWNN Address Pool for FC based Storage Access



This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

For FC boot as well as access to FC LUNs, create a World Wide Node Name (WWNN) pool by following these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Select **Pools > root**.
3. Right-click WWNN Pools under the root organization and choose **Create WWNN Pool** to create the WWNN address pool.
4. Enter `WWNN-Pool` as the name of the WWNN pool.
5. **Optional:** Enter a description for the WWNN pool.
6. Select the Sequential Assignment Order and click **Next**.
7. Click **Add**.
8. Specify a starting WWNN address.
9. Specify a size for the WWNN address pool that is sufficient to support the available blade or rack server resources. Each server will receive one WWNN.



Modifications of the WWNN block, as well as the WWPn and MAC Addresses, can convey identifying information for the Cisco UCS domain.



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. Modify the values accordingly to make them unique.

Create WWN Block ? X

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

10. Click **OK** and click **Finish**.
11. In the confirmation message, click **OK**.

Create a WWPN Address Pools for FC Based Storage Access



This configuration step can be skipped if the UCS environment does not need access to storage environment using FC.

If you are providing FC boot or access to FC LUNs, create a World Wide Port Name (WWPN) pool for each SAN switching fabric by completing the following steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Select **Pools > root**.
3. Right-click WWPN Pools under the root organization and choose **Create WWPN Pool** to create the first WWPN address pool.
4. Enter `WWPN-Pool-A` as the name of the WWPN pool.
5. **Optional:** Enter a description for the WWPN pool.
6. Select the Sequential Assignment Order and click **Next**.
7. Click **Add**.
8. Specify a starting WWPN address.



It is recommended to place `oA` in the second last octet of the starting WWPN address to identify all of the WWPN addresses as Fabric A addresses.

9. Specify a size for the WWPN address pool that is sufficient to support the available blade or rack server resources. Each server's Fabric A vHBA will receive one WWPN from this pool.

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

10. Click **OK** and click **Finish**.
11. In the confirmation message, click **OK**.
12. Right-click WWPN Pools under the root organization and choose **Create WWPN Pool** to create the second WWPN address pool.
13. Enter `WWPN-POOL-B` as the name of the WWPN pool.
14. **Optional:** Enter a description for the WWPN pool.
15. Select the Sequential Assignment Order and click **Next**.
16. Click **Add**.
17. Specify a starting WWPN address.



It is recommended to place `0B` in the second last octet of the starting WWPN address to identify all of the WWPN addresses as Fabric B addresses.

18. Specify a size for the WWPN address pool that is sufficient to support the available blade or rack server resources. Each server's Fabric B vHBA will receive one WWPN from this pool.

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

19. Click **OK** and click **Finish**.

- In the confirmation message, click **OK**.

Create IQN Pools for iSCSI Boot and LUN Access (iSCSI Deployment)



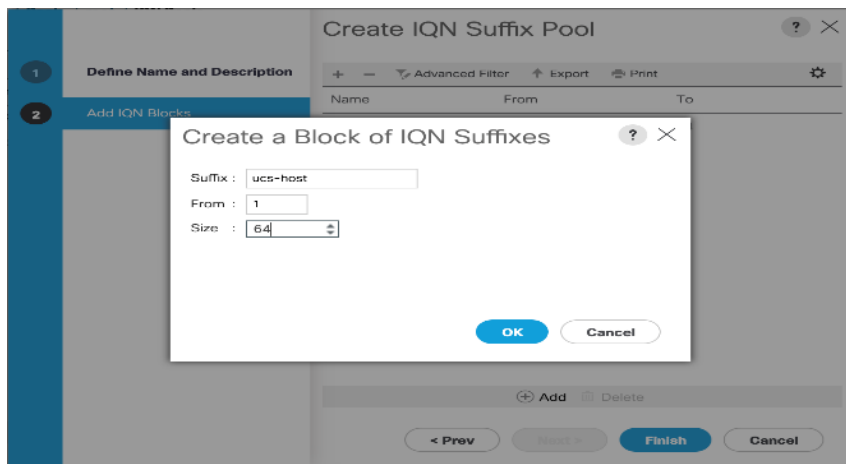
This configuration step can be skipped if the UCS environment does not need access to storage using iSCSI.

To enable iSCSI boot and provide access to iSCSI LUNs, configure the necessary IQN pools in the Cisco UCS Manager by completing the following steps:

- In the UCS Manager, select the **SAN** tab.
- Select **Pools > root**.
- Right-click IQN Pools under the root organization and choose **Create IQN Suffix Pool** to create the IQN pool.
- Enter `Infra-IQN-Pool` for the name of the IQN pool.
- Optional:** Enter a description for the IQN pool.
- Enter `iqn.1992-08.com.cisco` as the prefix
- Select the option Sequential for Assignment Order field. Click **Next**.

- Click **Add**.
- Enter an identifier with `ucs-host` as the suffix. Optionally a rack number or any other identifier can be added to the suffix to make the IQN unique within a DC.
- Enter 1 in the From field.

11. Specify a size of the IQN block sufficient to support the available server resources. Each server will receive one IQN.
12. Click **OK**.



13. Click **Finish**. In the message box that displays, click **OK**.

Create IP Pools for iSCSI Boot and LUN Access (iSCSI Deployment)



This configuration step can be skipped if the UCS environment does not need access to storage using iSCSI.

For enabling iSCSI storage access, these steps provide details for configuring the necessary IP pools in the Cisco UCS Manager:

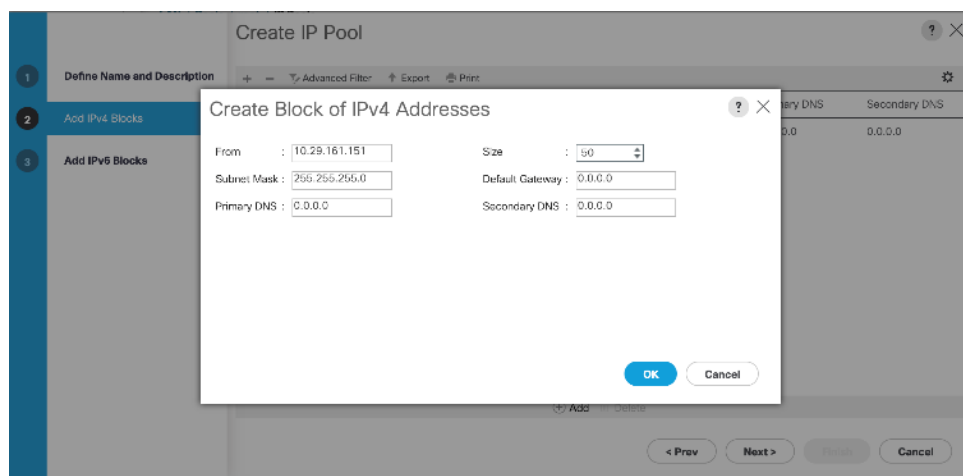


Two IP pools are created, one for each switching fabric.

1. In Cisco UCS Manager, select the **LAN** tab.
2. Select **Pools > root**.
3. Right-click IP Pools under the root organization and choose **Create IP Pool** to create the IP pool.
4. Enter `iSCSI-initiator-A` for the name of the IP pool.
5. **Optional:** Enter a description of the IP pool.
6. Select the option Sequential for the Assignment Order field. Click **Next**.

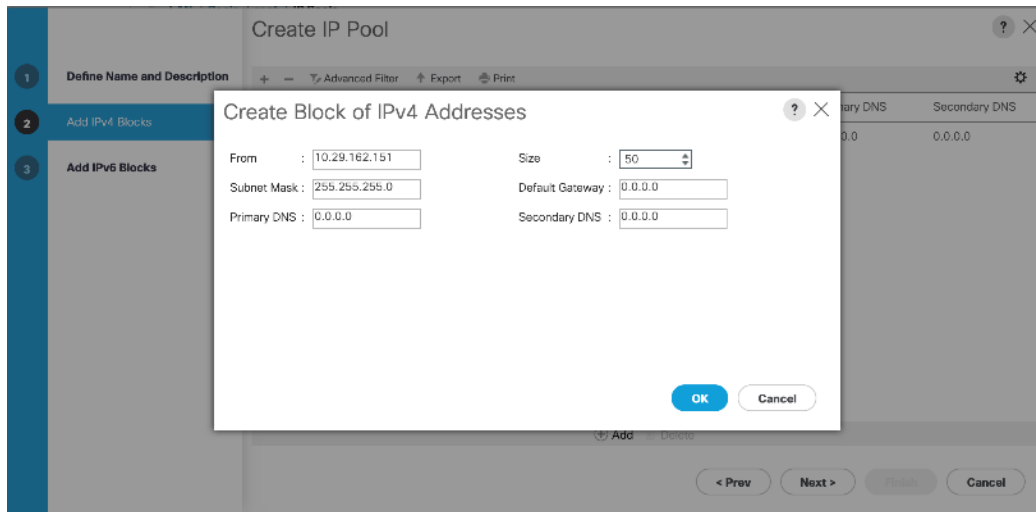


7. Click **Add**.
8. In the From field, enter the beginning of the range to assign an iSCSI IP addresses. These addresses are covered in Table 2 .
9. Enter the Subnet Mask.
10. Set the size with sufficient address range to accommodate the servers. Click **OK**.



11. Click **Next** and then click **Finish**.
12. Click **OK** in the confirmation message.
13. Right-click IP Pools under the root organization and choose **Create IP Pool** to create the IP pool.
14. Enter `iSCSI-initiator-B` for the name of the IP pool.
15. **Optional:** Enter a description of the IP pool.
16. Select the Sequential option for the Assignment Order field. Click **Next**.
17. Click **Add**.
18. In the From field, enter the beginning of the range to assign an iSCSI IP addresses. These addresses are covered in Table 2 .

19. Enter the Subnet Mask.
20. Set the size with sufficient address range to accommodate the servers. Click **OK**.



21. Click **Next** and then click **Finish**.
22. Click **OK** in the confirmation message.

Create VLANs

To configure the necessary VLANs in the Cisco UCS Manager, follow these steps for all the VLANs listed in Table 12 :

Table 12 VLANs on Cisco UCS

VLAN Name	VLAN
IB-Mgmt	11
iSCSI-A*	3161
iSCSI-B*	3162
Out of Band Mgmt	3171
VM Traffic	3174
vMotion	3173
Native-2	2



* iSCSI-A and iSCSI-B VLANs are required for iSCSI deployments only.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select LAN > LAN Cloud.
3. Right-click **VLANs** and choose **Create VLANs**.
4. Enter name from the VLAN Name column.

5. Keep the Common/Global option selected for the scope of the VLAN.
6. Enter the VLAN ID associated with the name.
7. Keep the Sharing Type as None.
8. 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

9. Click **Yes** and then click **OK** twice.
10. Repeat these steps for all the VLANs in Table 12 .

Create Host Firmware Package

Firmware management policies allow the administrator to select 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.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages and choose **Create Host Firmware Package**.
4. Enter `Infra-FW-Pack` as the name of the host firmware package.
5. Keep the Host Firmware Package as Simple.
6. Select the version 4.o(4c) for both the Blade and Rack Packages.

7. Click **OK** to create the host firmware package.
8. Click **OK**.

Create Host Firmware Package [?] [X]

Name :

Description :

How would you like to configure the Host Firmware Package?

Simple Advanced

Blade Package :

Rack Package :

Service Pack :

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

Excluded Components:

- Adapter
- BIOS
- Board Controller
- CIMC
- FC Adapters
- Flex Flash Controller
- GPUs
- HBA Option ROM
- Host NIC
- Host NIC Option ROM
- Local Disk

OK **Cancel**

Set Jumbo Frames in Cisco UCS Fabric

Jumbo Frames are used in VersaStack for the iSCSI storage protocols. The normal best practice in VersaStack 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 the MTU for the Best Effort QoS System Class is fixed at normal and cannot be changed. Testing has shown that even with this setting of normal in the 6454, Jumbo Frames can pass through the Cisco UCS fabric without being dropped. The screenshot below is from Cisco UCS Manager on a 6454 Fabric Interconnect, where the MTU for the Best Effort class is not configurable. To configure jumbo frames in the Cisco UCS fabric in a 6300 or 6200 series Fabric Interconnect, follow these steps:

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

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select 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** in the bottom of the window.

LAN / LAN Cloud / QoS System Class

General Events FSM

Actions Properties

Use Global Owner: Local

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

6. Click **OK**.

Create Local Disk Configuration Policy

When using an external storage system for OS boot, a local disk configuration for the Cisco UCS environment is necessary because the servers in the environment will not contain a local disk.



This policy should not be applied to the servers that contain local disks.

To create a local disk configuration policy for no local disks, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click Local Disk Config Policies and choose **Create Local Disk Configuration Policy**.
4. Enter `SAN-Boot` as the local disk configuration policy name.
5. Change the mode to **No Local Storage**.
6. Click **OK** to create the local disk configuration policy.

Create Local Disk Configuration Policy ? X

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** again.

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

To create a network control policy that enables Link Layer Discovery Protocol (LLDP) on virtual network ports, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click Network Control Policies and choose Create Network Control Policy.
4. Enter `Enable-CDP-LLDP` as the policy name.
5. For CDP, select Enabled option.
6. For LLDP, scroll down and select Enabled for both Transit 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

7. Click **OK** to create the network control 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 the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click Power Control Policies and choose Create Power Control Policy.
4. Enter `No-Power-Cap` as the power control policy name.
5. Change the power capping setting to No Cap.
6. Click **OK** to create the power control policy.
7. Click **OK**.

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

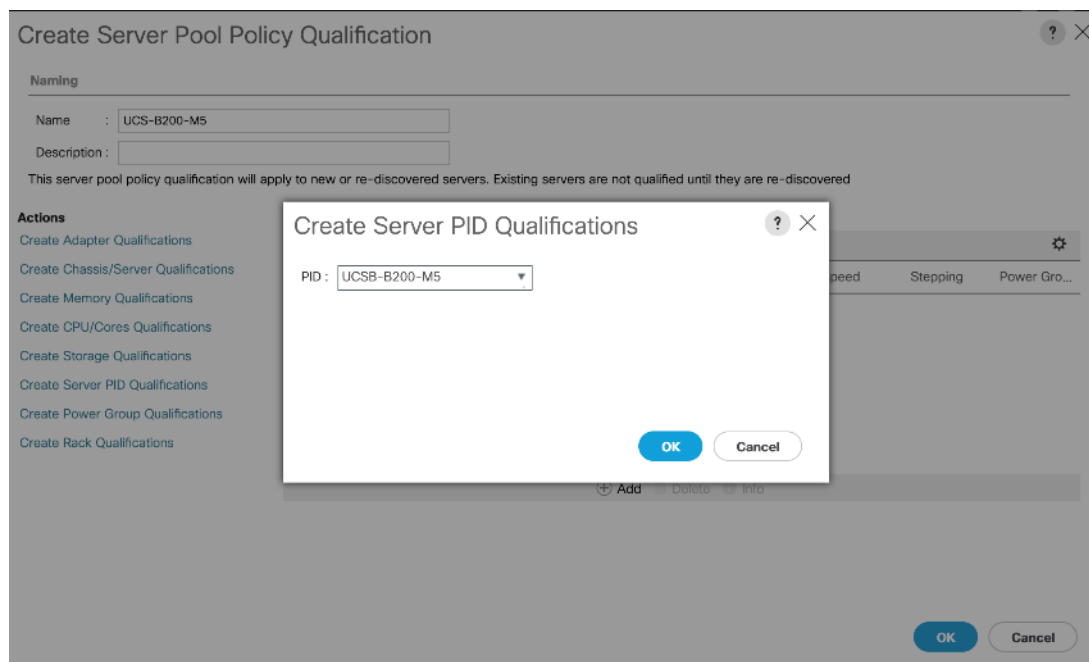
Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for selecting a Cisco UCS B200-M5 server.

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click **Server Pool Policy Qualifications** and choose Create Server Pool Policy Qualification.
4. Enter UCSB-B200-M5 as the name for the policy.
5. Choose Create Server PID Qualifications.
6. Select UCSB-B200-M5 as the PID.



7. Click **OK**.
8. Click **OK** to create the server pool policy qualification.



The server pool qualification policy name and the PID values varies if the UCS C-Series or other B-Series server models are used, select appropriate values based on the server model being used.

Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click **BIOS Policies** and choose Create BIOS Policy.
4. Enter `Infra-Host-BIOS` as the BIOS policy name.

Create BIOS Policy



Name :

Description :

Reboot on BIOS Settings Change :

OK **Cancel**

5. Click **OK**, then **OK** again to create the BIOS Policy.
6. Select the newly created BIOS Policy.
7. Set the following within the Main tab of the Policy:
 - a. CDN Control -> Enabled
 - b. Quiet Boot -> Disabled

Servers / Policies / root / BIOS Policies / Infra-Host-BIOS

Main | Advanced | Boot Options | Server Management | Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **Infra-Host-BIOS**

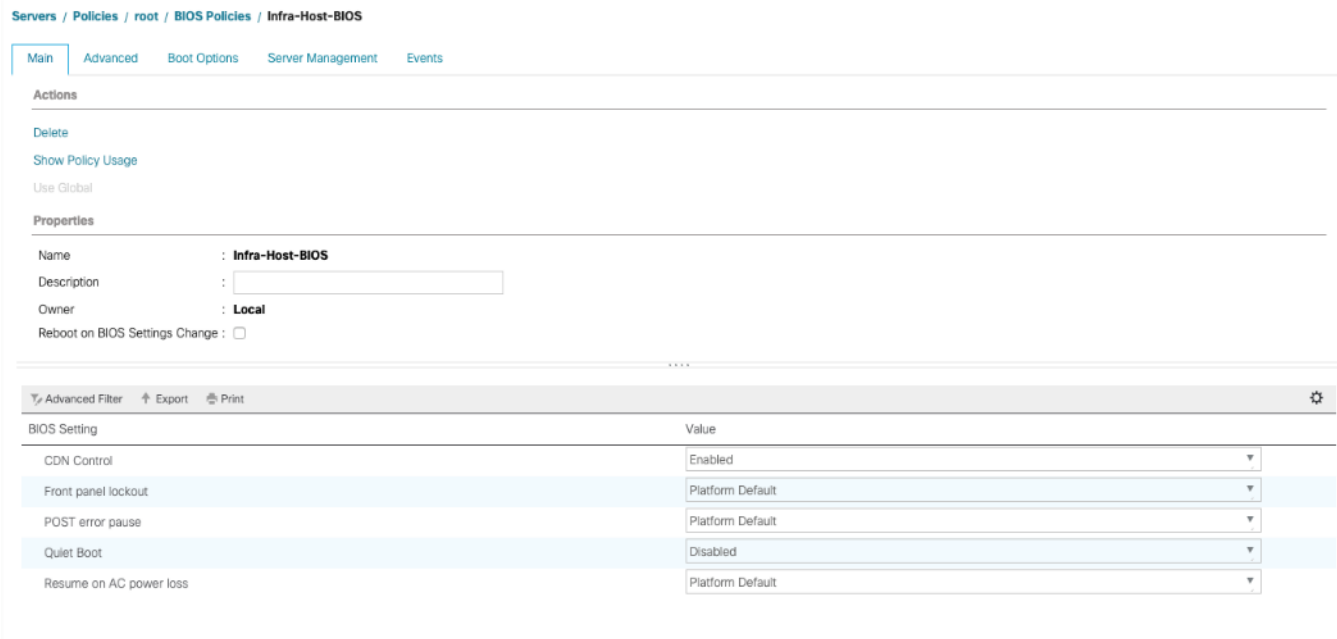
Description :

Owner : **Local**

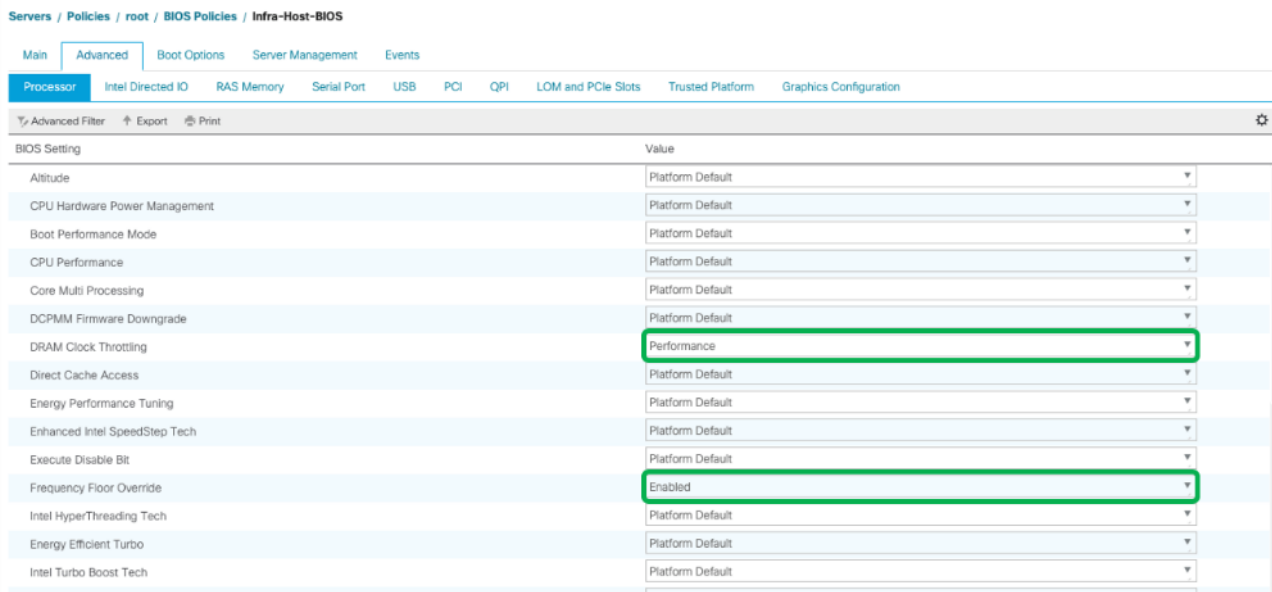
Reboot on BIOS Settings Change :

Advanced Filter | Export | Print

BIOS Setting	Value
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. Set the following within the Processor tab:
 - a. DRAM Clock Throttling -> Performance
 - b. Frequency Floor Override -> Enabled



9. Scroll down to the remaining Processor options and select:
 - a. Processor C State -> Disabled
 - b. Processor C1E -> Disabled

- c. Processor C3 Report -> Disabled
- d. Processor C7 Report -> Disabled
- e. Energy Performance -> Performance

BIOS Setting	Value
Package C State Limit	Platform Default
Autonomous Core C-state	Platform Default
Processor C State	Disabled
Processor C1E	Disabled
Processor C3 Report	Disabled
Processor C6 Report	Platform Default
Processor C7 Report	Disabled
Processor CMC1	Platform Default
Power Technology	Platform Default
Energy Performance	Performance
ProcessorEppProfile	Platform Default
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
IPI Prefetch	Platform Default

- 10. Click the RAS Memory tab, and select:
 - a. LV DDR Mode -> Performance Mode

BIOS Setting	Value
DDR3 Voltage Selection	Platform Default
DRAM Refresh Rate	Platform Default
LV DDR Mode	Performance Mode
Mirroring Mode	Platform Default
NUMA optimized	Platform Default
Memory RAS configuration	Platform Default

- 11. Click **Save Changes** to modify the BIOS policy.
- 12. Click **OK**.

Update Default Maintenance Policy

To update the default Maintenance Policy, follow these steps:

- 1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- 2. Select Policies > root and then select Maintenance Policies > default.

3. Change the Reboot Policy to User Ack.
4. Check the box to enable On Next Boot
5. Click Save Changes.
6. Click **OK** to accept the change.

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 Timer Automatic

Reboot Policy : Immediate User Ack Timer Automatic

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

Create vNIC/vHBA Placement Policy

To create a vNIC/vHBA placement policy for the infrastructure hosts, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies and choose Create Placement Policy.
4. Enter `Infra-Policy` as the name of the placement policy.
5. Click **1** and select Assigned Only.
6. Click **OK** and then click **OK** again.

Create Placement Policy ? X

Name :

Virtual Slot Mapping Scheme : Round Robin Linear Ordered

Virtual Slot	Selection Preference	Transport
1	All	ethernet,fc
2	All	ethernet,fc
3	All	ethernet,fc
4	All	ethernet,fc

Create vNIC Templates

Eight different vNIC Templates are covered in Table 13 . Not all vNICs need to be created in all deployments. The vNICs templates covered below are for iSCSI vNICs, infrastructure (management, vMotion etc.) vNICs, and data vNICs (VM traffic) for VMware VDS. Refer to Usage column in Table 13 to see if a vNIC is needed for a particular ESXi host.

Table 13 NIC Templates and Associated VLANs

Name	Fabric ID	VLANs	Native VLAN	MAC Pool	Usage
vNIC_Mgmt_A	A	IB-Mgmt, Native-2	Native-2	MAC-Pool-A	All ESXi Hosts
vNIC_Mgmt_B	B	IB-Mgmt, Native-2	Native-2	MAC-Pool-B	All ESXi Hosts
vNIC_vMotion_A	A	vMotion	vMotion	MAC-Pool-A	All ESXi Hosts
vNIC_vMotion_B	A	vMotion	vMotion	MAC-Pool-B	All ESXi Hosts
vNIC_VM_A	A	VM Network	Native-2	MAC-Pool-A	All ESXi Hosts
vNIC_VM_B	A	VM Network	Native-2	MAC-Pool-B	All ESXi Hosts

Name	Fabric ID	VLANs	Native VLAN	MAC Pool	Usage
vNIC_iSCSI_A	A	iSCSI-A	iSCSI-A	MAC-Pool-A	iSCSI hosts only
vNIC_iSCSI_B	B	iSCSI-B	iSCSI-B	MAC-Pool-B	iSCSI hosts only

Create Management vNICs

For the vNIC_Mgmt_A Template, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_Mgmt_A as the vNIC template name.
6. Keep Fabric A selected.
7. Optional: select the Enable Failover checkbox.



Selecting Failover can improve link failover time by handling it at the hardware level and can guard against any potential for NIC failure not being detected by the virtual switch.

8. Select Primary Template for the Redundancy Type.
9. Leave Peer Redundancy Template as **<not set>**



Redundancy Type and specification of Redundancy Template are configuration options to later allow changes to the Primary Template to automatically adjust onto the Secondary Template.

10. Under Target, make sure that the VM checkbox is not selected.
11. Select Updating Template as the Template Type.

Create vNIC Template



Name : vNIC_Mgmt_A

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : <not set>

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

OK Cancel

12. Under VLANs, select the checkboxes for IB-Mgmt and Native-VLAN VLANs.
13. Set Native-VLAN as the native VLAN.
14. Leave vNIC Name selected for the CDN Source.
15. Leave 1500 for the MTU.
16. In the MAC Pool list, select MAC_Pool_A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Select	Name	Native VLAN	VLAN ID
<input checked="" type="checkbox"/>	IB-Mgmt	<input type="radio"/>	11
<input type="checkbox"/>	ISCSI-A	<input type="radio"/>	3161
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>	3162
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>	2
<input type="checkbox"/>	VM-Net1	<input type="radio"/>	3174
<input type="checkbox"/>	vMotion	<input type="radio"/>	3173

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 1500

MAC Pool : MAC-Pool-A(52/64) ▼

QoS Policy : <not set> ▼

Network Control Policy : Enable-CDP-LLDP ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

Connection Policies

OK Cancel

18. Click **OK** to create the vNIC template.

19. Click **OK**.

For the vNIC_Mgmt_B Template, follow these steps:

1. In the navigation pane, select the LAN tab.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter vNIC_Mgmt_B as the vNIC template name.
6. Select Fabric B.
7. Select Secondary Template for Redundancy Type.
8. For the Peer Redundancy Template drop-down, select vNIC_Mgmt_A.



With Peer Redundancy Template selected, Failover specification, Template Type, VLANs, CDN Source, MTU, and Network Control Policy are all pulled from the Primary Template.

9. Under Target, make sure the VM checkbox is not selected.

Create vNIC Template

Name : vNIC_Mgmt_B

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : vNIC_Mgmt_A

Target

Adapter VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>

OK Cancel

10. In the MAC Pool list, select MAC_Pool_B.

Create vNIC Template

VLANs | VLAN Groups

Advanced Filter | |

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	OOB-MGMT	<input type="radio"/>

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

Dynamic vNIC Connection Policy :

11. Click **OK** to create the vNIC template.

12. Click **OK**.

Create vMotion vNICs

For the vNIC_vMotion_A Template, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_vMotion_A as the vNIC template name.
6. Keep Fabric A selected.
7. Optional: select the Enable Failover checkbox.
8. Select Primary Template for the Redundancy Type.

9. Leave Peer Redundancy Template as <not set>
10. Under Target, make sure that the VM checkbox is not selected.
11. Select Updating Template as the Template Type.

Create vNIC Template ? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template :

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print ⚙

Select	Name	Native VLAN
<input type="checkbox"/>	ISCSI-A	
<input type="checkbox"/>	ISCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	OOB-MGMT	<input type="radio"/>

12. Under VLANs, select the checkboxes vMotion as the only VLAN.
13. Set vMotion as the native VLAN.
14. For MTU, enter 9000.
15. In the MAC Pool list, select MAC_Pool_A.
16. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



<input type="checkbox"/>	ISCSI-B	<input type="radio"/>	3162
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>	2
<input type="checkbox"/>	VM-Net1	<input type="radio"/>	3174
<input checked="" type="checkbox"/>	vMotion	<input checked="" type="radio"/>	3173

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**

17. Click **OK** to create the vNIC template.

18. Click **OK**.

For the vNIC_vMotion_B Template, follow these steps:

1. In the navigation pane, select the LAN tab.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_vMotion_B as the vNIC template name.
6. Select Fabric B.
7. Select Secondary Template for Redundancy Type.
8. For the Peer Redundancy Template drop-down, select vNIC_vMotion_A.



With Peer Redundancy Template selected, MAC Pool will be the main configuration option left for this vNIC template.

9. Under Target, make sure the VM checkbox is not selected.

Create vNIC Template



Name : vNIC_vMotion_B

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template : <not set> ▼

<not set>

Domain Policies

vNIC_Mgmt_A

vNIC_vMotion_A

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
 If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print | ⚙️

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>

10. the MAC Pool list, select MAC_Pool_B.

Create vNIC Template ? X

VLANs | VLAN Groups

Advanced Filter | Export | Print ⚙️

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	ISCSI-A	<input type="radio"/>
<input type="checkbox"/>	ISCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	QDR-MGMT	<input type="radio"/>

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

Dynamic vNIC Connection Policy :

11. Click **OK** to create the vNIC template.

12. Click **OK**.

Create Application vNICs

To create the vNIC_VM_A Template, follow these steps:

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_VM_A as the vNIC template name.
6. Keep Fabric A selected.
7. Optional: select the Enable Failover checkbox.
8. Select Primary Template for the Redundancy Type.
9. Leave Peer Redundancy Template as <not set>.
10. Under Target, make sure that the VM checkbox is not selected.

11. Select Updating Template as the Template Type.
12. Set default as the native VLAN.

Create vNIC Template ? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template :

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	ISCSI-A	
<input type="checkbox"/>	ISCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	OOB-MGMT	<input type="radio"/>

13. Under VLANs, select the checkboxes for any application or production VLANs that should be delivered to the ESXi hosts.
14. For MTU, enter 9000.
15. In the MAC Pool list, select MAC_Pool_A.
16. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	ib-mgmt		11
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>	3161
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>	3162
<input checked="" type="checkbox"/>	Native-VLAN	<input checked="" type="radio"/>	2
<input checked="" type="checkbox"/>	VM-Net1	<input type="radio"/>	3174
<input type="checkbox"/>	vMotion	<input type="radio"/>	3173

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

17. Click **OK** to create the vNIC template.

18. Click **OK**.

To create the vNIC_VM_B Templates, follow these steps:

1. In the navigation pane, select the LAN tab.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter vNIC_VM_B as the vNIC template name.
6. Select Fabric B.
7. Select Secondary Template for Redundancy Type.
8. From the Peer Redundancy Template drop-down list, select vNIC_VM_A.



With Peer Redundancy Template selected, MAC Pool will be the main configuration option left for this vNIC template.

9. Under Target, make sure the VM checkbox is not selected.

Create vNIC Template



Name : vNIC_VM_B
 Description :
 Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template
 Peer Redundancy Template : vNIC_VM_A

Target

Adapter VM

<not set>

Domain Policies

vNIC_Mgmt_A

vNIC_VM_A

vNIC_vMotion_A

Warning

If **VM** is selected, a port profile by the same name will be created.
 If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	ISCSI-A	<input type="radio"/>
<input type="checkbox"/>	ISCSI-B	<input type="radio"/>

10. In the MAC Pool list, select MAC_Pool1_B.

Create vNIC Template



Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	QDR-MGMT	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 1500

MAC Pool : <not set>

QoS Policy : <not set>

Network Control Policy : Domain Pools

Pin Group : MAC-Pool-A(60/64)
MAC-Pool-B(60/64)
default(0/0)

OK Cancel

11. Click **OK** to create the vNIC template.

12. Click **OK**.

Create iSCSI vNICs



The configuration steps to create iSCSI vNICs can be skipped if the UCS environment does not need to access storage infrastructure using iSCSI.

1. In Cisco UCS Manager, click the **LAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_iSCSI_A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Keep the No Redundancy options selected for the Redundancy Type.

9. Under Target, make sure that the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select iSCSI-A-VLAN as the only VLAN and set it as the Native VLAN.

Create vNIC Template ? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter

VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	OOB-MGMT	<input type="radio"/>
<input checked="" type="checkbox"/>	iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>
<input type="checkbox"/>	OOB-MGMT	<input type="radio"/>

12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_A.
14. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Advanced Filter Export Print

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	ib-mgmt		11
<input checked="" type="checkbox"/>	iSCSI-A	<input checked="" type="radio"/>	3161
<input type="checkbox"/>	iSCSI-B	<input type="radio"/>	3162
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>	2
<input type="checkbox"/>	VM-Net1	<input type="radio"/>	3174
<input type="checkbox"/>	vMotion	<input type="radio"/>	3173

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

15. Click **OK** to create the vNIC template.

16. Click **OK**.

To create the vNIC_iSCSI_B Template, follow these steps:

1. In the navigation pane, select the **LAN** tab.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC_iSCSI_B as the vNIC template name.
6. Keep Fabric B selected.
7. Do not select the Enable Failover checkbox.
8. Keep the No Redundancy options selected for the Redundancy Type.
9. Under Target, make sure that the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select iSCSI-B-VLAN as the only VLAN and set it as the Native VLAN.

Create vNIC Template

Name : vNIC_ISCSI_B

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Target

Adapter VM

Warning

If **VM** is selected, a port profile by the same name will be created.
If a port profile of the same name exists, and updating template is selected, it will be overwritten

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	ISCSI-A	<input type="radio"/>
<input checked="" type="checkbox"/>	ISCSI-B	<input checked="" type="radio"/>
<input type="checkbox"/>	Native	<input type="radio"/>

OK Cancel

12. For MTU, enter 9000.
13. In the MAC Pool list, select MAC_Pool_B.
14. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Advanced Filter Export Print

Select	Name	Native VLAN	VLAN ID
<input type="checkbox"/>	IB-Mgmt		11
<input type="checkbox"/>	iSCSI-A	<input type="radio"/>	3161
<input checked="" type="checkbox"/>	iSCSI-B	<input checked="" type="radio"/>	3162
<input type="checkbox"/>	Native-VLAN	<input type="radio"/>	2
<input type="checkbox"/>	VM-Net1	<input type="radio"/>	3174
<input type="checkbox"/>	vMotion	<input type="radio"/>	3173

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

15. Click **OK** to create the vNIC template.

16. Click **OK**.

Create LAN Connectivity Policy

To configure the necessary Infrastructure LAN Connectivity Policy, follow these steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.



Use Infra-LAN-Pol as the name if hosts boot from FC only.

5. Enter `iSCSI-LAN-Policy` as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the **Create vNIC** dialog box, enter `oo-Mgmt-A` as the name of the vNIC.



The numeric prefix of "00-" and subsequent increments on the later vNICs are used in the vNIC naming to force the device ordering through Consistent Device Naming (CDN). Without this, some operating systems might not respect the device ordering that is set within Cisco UCS.

8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select 00-Mgmt-A.
10. In the Adapter Policy list, select VMWare.
11. Click **OK** to add this vNIC to the policy.

Create vNIC ? X

Name : 00-Mgmt-A

Use vNIC Template :

Redundancy Pair : Peer Name :

vNIC Template : vNIC_Mgmt_A Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare Create Ethernet Adapter Policy

OK Cancel

12. Click the upper Add button to add another vNIC to the policy.
13. In the **Create vNIC** box, enter 01-Mgmt-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select 01-Mgmt-B.
16. In the Adapter Policy list, select VMWare.

Create vNIC ? X

Name : 01-Mgmt-B

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template : vNIC_Mgmt_B

Create vNIC Template

Adapter Performance Profile

Adapter Policy : VMWare

Create Ethernet Adapter Policy

OK Cancel

17. Click **OK** to add the vNIC to the policy.
18. Click the upper Add button to add a vNIC.
19. In the **Create vNIC** dialog box, enter 02-vMotion-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select vNIC_vMotion_A.
22. In the Adapter Policy list, select VMWare.
23. Click **OK** to add this vNIC to the policy.

Create vNIC ? X

Name :

Use vNIC Template :

Redundancy Pair :

Peer Name :

vNIC Template : [Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy : [Create Ethernet Adapter Policy](#)

24. Click the upper Add button to add a vNIC to the policy.
25. In the **Create vNIC** dialog box, enter 03-vMotion-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select vNIC_vMotion_B.
28. In the Adapter Policy list, select VMWare.

Create vNIC

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

29. Click **OK** to add this vNIC to the policy.
30. Click the upper Add button to add a vNIC.
31. In the **Create vNIC** dialog box, enter `04-VM-A` as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select `vNIC_VM_A`.
34. In the Adapter Policy list, select `VMWare`.
35. Click **OK** to add this vNIC to the policy.

Create vNIC ? X

Name : 04-VM-A

Use vNIC Template :

Redundancy Pair :

vNIC Template : vNIC_VM_A

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy : VMWare

[Create Ethernet Adapter Policy](#)

OK Cancel

36. Click the upper Add button to add a vNIC to the policy.
37. In the **Create vNIC** dialog box, enter 05-VM-B as the name of the vNIC.
38. Select the Use vNIC Template checkbox.
39. In the vNIC Template list, select vNIC_VM_B.
40. In the Adapter Policy list, select VMWare.

Create vNIC ? X

Name : 05-VM_B

Use vNIC Template :

Redundancy Pair :

vNIC Template : vNIC_VM_B

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy : VMWare

[Create Ethernet Adapter Policy](#)

OK Cancel

41. Click **OK** to add this vNIC to the policy.
42. Click the upper Add button to add a vNIC.



The following iSCSI vNICs can be skipped if hosts need FC storage access only.

43. In the **Create vNIC** dialog box, enter 06-iSCSI-A as the name of the vNIC.
44. Select the Use vNIC Template checkbox.
45. In the vNIC Template list, select iSCSI-Template-A.
46. In the Adapter Policy list, select VMWare.

Create vNIC ? X

Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

47. Click **OK** to add this vNIC to the policy.
48. Click the upper Add button to add a vNIC to the policy.
49. In the Create vNIC dialog box, enter `07-iSCSI-B` as the name of the vNIC.
50. Select the Use vNIC Template checkbox.
51. In the vNIC Template list, select `iSCSI-Template-B`.
52. In the Adapter Policy list, select `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](#)

53. Click **OK** to add this vNIC to the policy.

Create LAN Connectivity Policy

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

Name	MAC Address	Native VLAN
vNIC 07-iSCSI-B	Derived	
vNIC 06-iSCSI-A	Derived	
vNIC 05-VM-B	Derived	
vNIC 04-VM-A	Derived	
vNIC 03-vMotion-B	Derived	
vNIC 02-vMotion-A	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
No data available			

Add Delete Modify

OK

Cancel

Add iSCSI vNICs in LAN Policy (iSCSI Deployment)



This configuration step can be skipped if the UCS environment does not need to access storage environment using iSCSI.

Follow these steps only if you are using iSCSI SAN access:

1. Verify the iSCSI base vNICs are already added as part of vNIC implementation.
2. Expand the **Add iSCSI vNICs** section to add the iSCSI boot vNICs.
3. Select **Add** in the Add iSCSI vNICs section.
4. Set the name to `iSCSI-A-vNIC`.
5. Select the `06-iSCSI-A` as Overlay vNIC.
6. Set the VLAN to `iSCSI-A (native) VLAN`.
7. Set the iSCSI Adapter Policy to default
8. Leave the MAC Address set to None.

Create iSCSI vNIC ? ×

Name :

Overlay vNIC :

iSCSI Adapter Policy : [Create iSCSI Adapter Policy](#)

VLAN :

iSCSI MAC Address

MAC Address Assignment:

[Create MAC Pool](#)

9. Click **OK**.
10. Select **Add** in the Add iSCSI vNICs section.
11. Set the name to `iSCSI-B-vNIC`.
12. Select the `07-iSCSI-B` as Overlay vNIC.
13. Set the VLAN to `iSCSI-B (native)` VLAN.
14. Set the iSCSI Adapter Policy to default.
15. Leave the MAC Address set to None.

Create iSCSI vNIC ? ×

Name :

Overlay vNIC :

iSCSI Adapter Policy : [Create iSCSI Adapter Policy](#)

VLAN :

iSCSI MAC Address

MAC Address Assignment: [Create MAC Pool](#)

16. Click **OK** then click **OK** again to create the LAN Connectivity Policy.

Create LAN Connectivity Policy

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

Name	MAC Address	Native VLAN
vNIC 07-ISCASI-B	Derived	
vNIC 06-ISCASI-A	Derived	
vNIC 05-VM_B	Derived	
vNIC 04-VM-A	Derived	
vNIC 03-vMotion-B	Derived	
vNIC 02-vMotion-A	Derived	

Delete + Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
ISCASI vNIC ISCASI-B-vNIC	07-ISCASI-B		Derived
ISCASI vNIC ISCASI-A-vNIC	06-ISCASI-A		Derived

+ Add Delete Modify

OK

Cancel

Create vHBA Templates for FC Connectivity (FC Deployment)



This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

To create virtual host bus adapter (vHBA) templates for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Select Policies > root.
3. Right-click **vHBA Templates** and choose Create vHBA Template.
4. Enter `Infra-vHBA-A` as the vHBA template name.
5. Click the radio button to select `Fabric A`.
6. In the Select VSAN list, Choose `VSAN-A`.
7. In the WWPN Pool list, Choose `WWPN-Pool-A`.

Create vHBA Template



Name :

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size :

WWPN Pool :

QoS Policy :

Pin Group :

Stats Threshold Policy :

OK

Cancel

8. Click **OK** to create the vHBA template.
9. Click **OK**.
10. Right-click **vHBA Templates** again and choose Create vHBA Template.
11. Enter `Infra-vHBA-B` as the vHBA template name.
12. Click the radio button to select Fabric `B`.
13. In the Select VSAN list, Choose `VSAN-B`.
14. In the WWPN Pool, Choose `WWPN-Pool-B`.

Create vHBA Template



Name :

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN :

Template Type : Initial Template Updating Template

Max Data Field Size :

WWPN Pool :

QoS Policy :

Pin Group :

Stats Threshold Policy :

15. Click **OK** to create the vHBA template.
16. Click **OK**.

Create FC SAN Connectivity Policies (FC Deployment)



This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

A SAN connectivity policy defines the vHBAs that will be created as part of a service profile deployment.

To configure the necessary FC SAN Connectivity Policies, follow these steps:

1. In Cisco UCS Manager, click the **SAN** tab in the navigation pane.
2. Select SAN > Policies > root.
3. Right-click **SAN Connectivity Policies** and choose Create SAN Connectivity Policy.
4. Enter `Infra-FC-pol` as the name of the policy.
5. Select WWNN-Pool from the drop-down list under World Wide Node Name.

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
No data available	

OK

Cancel

- Click **Add**. You might have to scroll down the screen to see the Add link.
- Under Create vHBA, enter `vHBA-A` in the Name field.
- Check the check box Use vHBA Template.
- From the vHBA Template drop-down list, select `Infra-vHBA-A`.
- From the Adapter Policy drop-down list, select `VMWare`.

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](#)

- Click **OK**.
- Click **Add**.

13. Under Create vHBA, enter vHBA-B in the Name field.
14. Check the check box next to Use vHBA Template.
15. From the vHBA Template drop-down list, select Infra-vHBA-B.
16. From the Adapter Policy drop-down list, select VMWare.

Create vHBA vHBA

Name : vHBA-B

Use vHBA Template :

Redundancy Pair : Peer Name :

vHBA Template : Infra-vHBA-B

Adapter Performance Profile

Adapter Policy : VMWare

17. Click **OK**.
18. Click **OK** again to accept creating the SAN connectivity policy.

Create iSCSI Boot Policy (iSCSI Deployment)



This configuration step can be skipped if the UCS environment does not need to access storage environment using iSCSI.

This procedure applies to a Cisco UCS environment in which iSCSI interface on Controller A is chosen as the primary target.

To create boot the policy for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies and choose Create Boot Policy.
4. Enter `Boot-iSCSI-X-A` as the name of the boot policy.
5. **Optional:** Enter a description for the boot policy.
6. Keep the Reboot on Boot Order Change option cleared.
7. Expand the Local Devices drop-down list and select Add Remote CD/DVD.
8. Expand the iSCSI vNICs section and select Add iSCSI Boot.
9. In the Add iSCSI Boot dialog box, enter `iSCSI-A-vNIC`.

10. Click **OK**.
11. Select Add iSCSI Boot.
12. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
13. Click **OK**.

Create Boot Policy

Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/iSCSI Name :

Boot Mode : Legacy Uefi

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

Add Local Disk

- Add Local LUN
- Add Local JBOD
- Add SD Card
- Add Internal USB
- Add External USB
- Add Embedded Local LUN
- Add Embedded Local Disk

Add CD/DVD

Boot Order

Name	Ord...	vNIC/v...	Type	LUN N...	WWN	Slot N...	Boot N...	Boot P...	Descri...
Remote CD/DVD									
	1								
iSCSI									
	2								
iSCSI		iSCSI-A	Primary						
iSCSI		iSCSI-B	Secon...						

↑ Move Up
 ↓ Move Down
 🗑️ Delete

14. Click **OK** then **OK** again to save the boot policy.

Create FC Boot Policies (FC Deployment)



This configuration step can be skipped if the UCS environment does not need to access storage environment using FC.

This procedure applies to a Cisco UCS environment in which two FC interfaces are used on each of the FS9100 node canisters for host connectivity. This procedure captures a single boot policy which defines Fabric-A as the primary fabric. Customer can choose to create a second boot policy which can use Fabric-B as primary fabric to spread the boot-from-san traffic load on both the nodes in case of disaster recovery.

WWPN information from the IBM FS9100 node canisters is required to complete this section. This information can be found by logging into the IBM FS9100 management address using SSH and issuing the commands as captured below. The information can be recorded in Table 14 .



Since NPIV feature is enabled on the IBM FS9100 systems, the WWPN permitted for host communication can be different from the physical WWPN. Refer to the example below.

1. Verify the node_id of the FS9100 node canisters using the following command (node1 through node 2 in this example):

lsportfc

```
IBM_FlashSystem:VersaStack-FS9100:superuser>lsportfc
id fc_io_port_id port_id type port_speed node_id node_name WWPN nportid status attachment cluster_use adapter_location adapter_port_id
0 1 1 fc 16Gb 1 node1 5005076810110516 720020 active switch local_partner 1 1
1 2 2 fc 16Gb 1 node1 5005076810120516 C60040 active switch local_partner 1 2
2 3 3 fc 16Gb 1 node1 5005076810130516 720040 active switch local_partner 1 3
3 4 4 fc 16Gb 1 node1 5005076810140516 C60060 active switch local_partner 1 4
16 1 1 fc 16Gb 2 node2 500507681011050D 720060 active switch local_partner 1 1
17 2 2 fc 16Gb 2 node2 500507681012050D C60080 active switch local_partner 1 2
18 3 3 fc 16Gb 2 node2 500507681013050D 720080 active switch local_partner 1 3
19 4 4 fc 16Gb 2 node2 500507681014050D C600A0 active switch local_partner 1 4
```

2. Use the following command to record the WWPN corresponding to ports connected to the SAN fabric:

lstargetportfc -filtervalue host_io_permitted=yes

```
IBM_FlashSystem:VersaStack-FS9100:superuser>lstargetportfc -filtervalue host_io_permitted=yes
id WWPN WWPN port_id owning_node_id current_node_id nportid host_io_permitted virtualized protocol
2 5005076810150516 5005076810000516 1 1 1 720021 yes yes scsi
3 5005076810190516 5005076810000516 1 1 1 720022 yes yes nvme
5 5005076810160516 5005076810000516 2 1 1 C60041 yes yes scsi
6 50050768101A0516 5005076810000516 2 1 1 C60042 yes yes nvme
8 5005076810170516 5005076810000516 3 1 1 720041 yes yes scsi
9 50050768101B0516 5005076810000516 3 1 1 720042 yes yes nvme
11 5005076810180516 5005076810000516 4 1 1 C60061 yes yes scsi
12 50050768101C0516 5005076810000516 4 1 1 C60062 yes yes nvme
50 500507681015050D 500507681000050D 1 2 2 720061 yes yes scsi
51 500507681019050D 500507681000050D 1 2 2 720062 yes yes nvme
53 500507681016050D 500507681000050D 2 2 2 C60081 yes yes scsi
54 50050768101A050D 500507681000050D 2 2 2 C60082 yes yes nvme
56 500507681017050D 500507681000050D 3 2 2 720081 yes yes scsi
57 50050768101B050D 500507681000050D 3 2 2 720082 yes yes nvme
59 500507681018050D 500507681000050D 4 2 2 C600A1 yes yes scsi
60 50050768101C050D 500507681000050D 4 2 2 C600A2 yes yes nvme
```

Table 14 IBM FS9100 – WWPN Information

Node	Port ID	WWPN	Variable	Fabric
FS9100 node canister 1	1		WWPN-FS9100-Node1-FC1	A
FS9100 node canister 1	3		WWPN-FS9100-Node1-FC3	A
FS9100 node canister 1	2		WWPN-FS9100-Node1-FC2	B
FS9100 node canister 1	4		WWPN-FS9100-Node1-FC4	B
FS9100 node canister 2	1		WWPN-FS9100-Node2-FC1	A
FS9100 node canister 2	3		WWPN-FS9100-Node2-FC3	A
FS9100 node canister 2	2		WWPN-FS9100-Node2-FC2	B
FS9100 node canister 2	4		WWPN-FS9100-Node2-FC4	B

To create boot policies for the Cisco UCS environment, follow these steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Choose Policies > root.
3. Right-click Boot Policies and choose Create Boot Policy.
4. Enter `Boot-Fabric-A` as the name of the boot policy.
5. **Optional:** Enter a description for the boot policy.
6. Keep the Reboot on the Boot Order Change check box unchecked.
7. Expand the Local Devices drop-down list and Choose **Add Remote CD/DVD**.
8. Expand the vHBAs drop-down list and Choose **Add SAN Boot**.

Create Boot Policy ? X

Name :

Description :

Reboot on Boot Order Change :

Enforce vNIC/vHBA/ISCSI Name :

Boot Mode : Legacy Uefi

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

Boot Order									
Name	Order	vNIC/v...	Type	LUN N...	WWN	Slot N...	Boot N...	Boot P...	Descr...
Remote CD/DVD	1								

9. Make sure to select Primary radio button as the Type.
10. Enter `vHBA-A` in the vHBA field.
11. Click **OK** to add the SAN boot initiator.

Add SAN Boot ? X

vHBA : vHBA-A

Type : Primary Secondary Any

OK Cancel

- 12. From the vHBA drop-down list, choose Add SAN Boot Target.
- 13. Keep 0 as the value for Boot Target LUN.
- 14. Enter the WWPN <WWPN-FS9100-Node1-FC1> from Table 14 .
- 15. Keep the Primary radio button selected as the SAN boot target type.
- 16. Click **OK** to add the SAN boot target.

Add SAN Boot Target ? X

Boot Target LUN : 0

Boot Target WWPN : 50:05:07:68:10:15:05:16

Type : Primary Secondary

OK Cancel

- 17. From the vHBA drop-down menu, choose **Add SAN Boot Target**.
- 18. Keep 0 as the value for Boot Target LUN.

- 19. Enter the WWPN <WWPN-FS9100-Node2-FC1> from Table 14 .
- 20. Click **OK** to add the SAN boot target.

Add SAN Boot Target ? X

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK **Cancel**

- 21. From the vHBA drop-down list, choose **Add SAN Boot**.
- 22. In the Add SAN Boot dialog box, enter vHBA-B in the vHBA box.
- 23. The SAN boot type should automatically be set to Secondary.
- 24. Click **OK** to add the SAN boot initiator.

Add SAN Boot ? ×

vHBA :

Type : Primary Secondary Any

25. From the vHBA drop-down list, choose Add SAN Boot Target.
26. Keep 0 as the value for Boot Target LUN.
27. Enter the WWPN <WWPN-FS9100-Node1-FC2> from Table 14 .
28. Keep Primary as the SAN boot target type.
29. Click **OK** to add the SAN boot target.

Add SAN Boot Target ? ×

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

30. From the vHBA drop-down list, choose **Add SAN Boot Target**.
31. Keep 0 as the value for Boot Target LUN.
32. Enter the WWPN <WWPN-FS9100-Node2-FC2> from Table 14 .
33. Click **OK** to add the SAN boot target.

Add SAN Boot Target



Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK

Cancel

34. Click **OK**, and then click **OK** again to create the boot policy.
35. Verify that your SAN boot configuration looks similar to the screenshot below.

Boot Order

Name	vNIC/vHBA/iSCS...	Type	WWN
▼ SAN Primary	vHBA-A	Primary	
SAN Target Primary		Primary	50:05:07:68:10:15:05:16
SAN Target Secondary		Secondary	50:05:07:68:10:15:05:0D
▼ SAN Secondary	vHBA-B	Secondary	
SAN Target Primary		Primary	50:05:07:68:10:16:05:16
SAN Target Secondary		Secondary	50:05:07:68:10:16:05:0D

↑ Move Up
↓ Move Down
🗑 Delete

Set Uefi Boot Parameters

Create iSCSI Boot Service Profile Template (iSCSI Deployment)

Service profile template configuration for the iSCSI-based SAN access is covered in this section.



This section can be skipped if iSCSI boot is not implemented in the customer environment.

To create the service profile template, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click **root**.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter `Infra-ESXi-iSCSI-Host` as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the “Updating Template” option.
7. Under UUID, select `UUID_Pool` as the UUID pool.

Create Service Profile Template

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**

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.

< Prev Next > **Finish** Cancel

8. Click **Next**.

Configure Storage Provisioning

To configure the storage provisioning, follow these steps:

1. If you have servers with no physical disks, click the Local Disk Configuration Policy tab and select the SAN-Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.

Create Service Profile Template

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Specific Storage Profile Storage Profile Policy **Local Disk Configuration Policy**

Local Storage:

[Create Local Disk Configuration Policy](#)

Mode : **No Local Storage**

Protect Configuration : **Yes**

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : **Disable**

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**

FlexFlash Removable State : **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.

< Prev Next > **Finish** Cancel

2. Click **Next**.

Configure Networking Options

To configure the network options, follow these steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the "Use Connectivity Policy" option to configure the LAN connectivity.
3. Select `iSCSI-LAN-Policy` from the LAN Connectivity Policy drop-down list.
4. Select `IQN_Pool` in Initiator Name Assignment.

5. Click **Next**.

Configure Storage Options

1. Select the **No vHBA** option for the “How would you like to configure SAN connectivity?” field.
2. Click **Next**.

Configure Zoning Options

1. Leave Zoning configuration unspecified and click **Next**.

Configure vNIC/HBA Placement

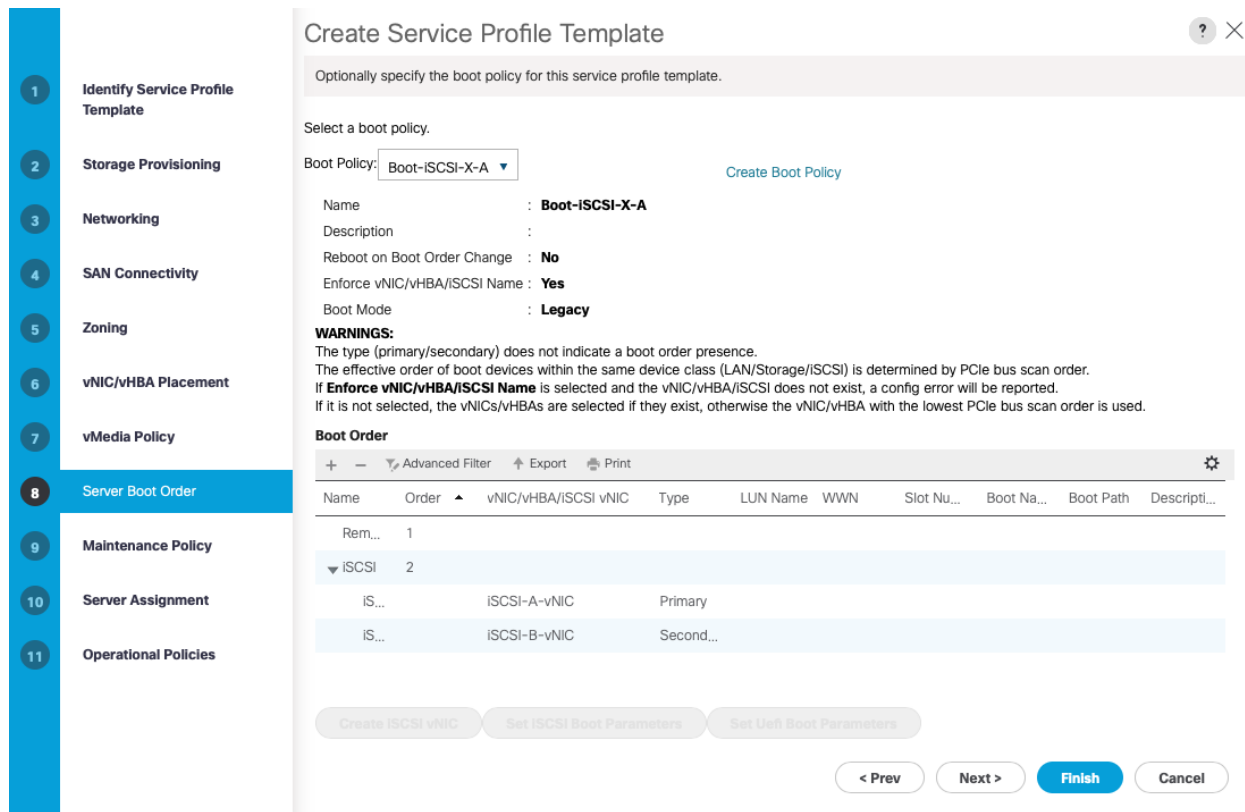
1. In the “Select Placement” list, leave the placement policy as “Let System Perform Placement.”
2. Click **Next**.

Configure vMedia Policy

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

Configure Server Boot Order

1. Select **Boot-iSCSI-X-A** for Boot Policy.



2. In the **Boor order**, select `iSCSI-A-vNIC`.
3. Click Set iSCSI Boot Parameters button.
4. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to `<not set>` unless you have independently 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-initiator-A` as the "Initiator IP address Policy."
7. Select iSCSI Static Target Interface option.
8. Click **Add**.
9. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 1 (IQN) from Table 11
10. Enter the IP address of Node 1 iSCSI-A interface from Table 10 .

Create iSCSI Static Target ? ×

iSCSI Target Name :

Priority : **1**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

11. Click **OK** to add the iSCSI Static Target.
12. Keep the iSCSI Static Target Interface option selected and click **Add**.
13. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 2 (IQN) from Table 11 .
14. Enter the IP address of Node 2 iSCSI-A interface from Table 10 .

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority : **2**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

15. Click **OK** to add the iSCSI Static Target.
16. Verify both the targets on iSCSI Path A as shown below:

Set iSCSI Boot Parameters



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-initiator-A(50/50) ▾

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](#)

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.1986-03...	1	3260		10.29.161.249	0
iqn.1986-03...	2	3260		10.29.161.250	0

[+](#) Add [-](#) Delete [i](#) Info

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

OK **Cancel**

17. Click **OK** to set the iSCSI-A-vNIC iSCSI Boot Parameters.
18. In the Boot order, select iSCSI-B-vNIC.
19. Click Set iSCSI Boot Parameters button.

20. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to **<not set>** unless you have independently created one appropriate to your environment.
21. Leave the "Initiator Name Assignment" dialog box **<not set>** to use the single Service Profile Initiator Name defined in the previous steps.
22. Set `iSCSI-initiator-B` as the "Initiator IP address Policy".
23. Select iSCSI Static Target Interface option.
24. Click **Add**.
25. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 1 (IQN) from Table 11 .
26. Enter the IP address of Node 1 iSCSI-B interface from Table 10 .

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority : **1**

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

27. Click **OK** to add the iSCSI Static Target.
28. Keep the iSCSI Static Target Interface option selected and click Add.
29. In the Create iSCSI Static Target dialog box, add the iSCSI target node name for Node 2 (IQN) from Table 11 .
30. Enter the IP address of Node 2 iSCSI-B interface from Table 10 .

Create iSCSI Static Target



iSCSI Target Name :

Priority : **2**

Port :

Authentication Profile :

[Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

OK **Cancel**

31. Click **OK** to add the iSCSI Static Target.

Set iSCSI Boot Parameters



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-initiator-B(50/50) ▼

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](#)

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.1986-03...	1	3260		10.29.162.249	0
iqn.1986-03...	2	3260		10.29.162.250	0

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

OK

Cancel

32. Click **OK** to set the iSCSI-B-vNIC iSCSI Boot Parameters.
33. Click **Next** to continue to the next section.

Configure Maintenance Policy

To configure the maintenance policy, follow these step:

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	: Immediate

< Prev Next > **Finish** Cancel

2. Click **Next**.

Configure Server Assignment

To configure server assignment, follow these steps:

1. In the Pool Assignment list, select `Infra-Server-Pool`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Optional: Select "UCS-B200M5" for the Server Pool Qualification.



Firmware Management at the bottom of the page can be left alone as it will use default from the Host Firmware 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

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: [Create Server Pool](#)

Select the power state to be applied when 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 pool will use.

Server Pool Qualification :

Restrict Migration :

Firmware Management (Server, Adapter)

- Domain Policies
- UCSB-B200-M5**
- all-chassis

5. Click **Next**.

Configure Operational Policies

To configure the operational policies, follow these steps:

1. In the BIOS Policy list, select **Infra-Host-BIOS**.
2. Expand Power Control Policy Configuration and select **No-Power-Cap** in the Power Control Policy 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 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

BIOS Policy :

External IPMI 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 Contr](#)

Scrub Policy

KVM Management

Graphics Card Policy

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

Create iSCSI Boot Service Profiles (iSCSI Deployment)

To create service profiles from the service profile template, follow these steps:

1. Connect to the UCS 6454 Fabric Interconnect UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root > Service Template `Infra-ESXi-iSCSI-Host`.
3. Right-click `Infra-ESXi-iSCSI-Host` and select Create Service Profiles from Template.
4. Enter `Infra-ESXi-Host-iSCSI-Host-` for iSCSI deployment as the service profile prefix
5. Enter 1 as the Name Suffix Starting Number.
6. Enter the Number of servers to be deploy in the Number of Instances field.

- Click **OK** to create the service profile.

Create Service Profiles From Template ? ✕

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

- Click **OK** in the confirmation message to provision four VersaStack Service Profiles.



Adjust the number of Service Profile instances based on the actual customer deployment with intended number of ESXi servers needed.

Create FC Boot Service Profile Template (FC Deployment)

In this procedure, a service profile template is created to use FC Fabric A as primary boot path.



This section can be skipped if FC boot is not implemented in the customer environment.

To create service profile templates, follow these steps:

- In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
- Choose Service Profile Templates > root.
- Right-click root and choose **Create Service Profile Template**. This opens the Create Service Profile Template wizard.
- Enter `Infra-ESXi-Host` as the name of the service profile template.
- Select the Updating Template option.
- Under UUID, select `UUID-Pool` as the UUID pool.
- Click **Next**.

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**

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.

Configure Storage Provisioning

1. Select the Local Disk Configuration Policy tab.
2. Select the SAN-Boot Local Storage Policy. This policy usage requires servers with no local HDDs.
3. Click **Next**.

Create Service Profile Template ? X

Optionally specify or create a Storage Profile, and select a local disk configuration policy.

Specific Storage Profile Storage Profile Policy **Local Disk Configuration Policy**

Local Storage: ▼

Create Local Disk Configuration Policy

Mode : **No Local Storage**

Protect Configuration : **Yes**

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server. In that case, a configuration error will be raised when a new service profile is associated with that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State : **Disable**

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**

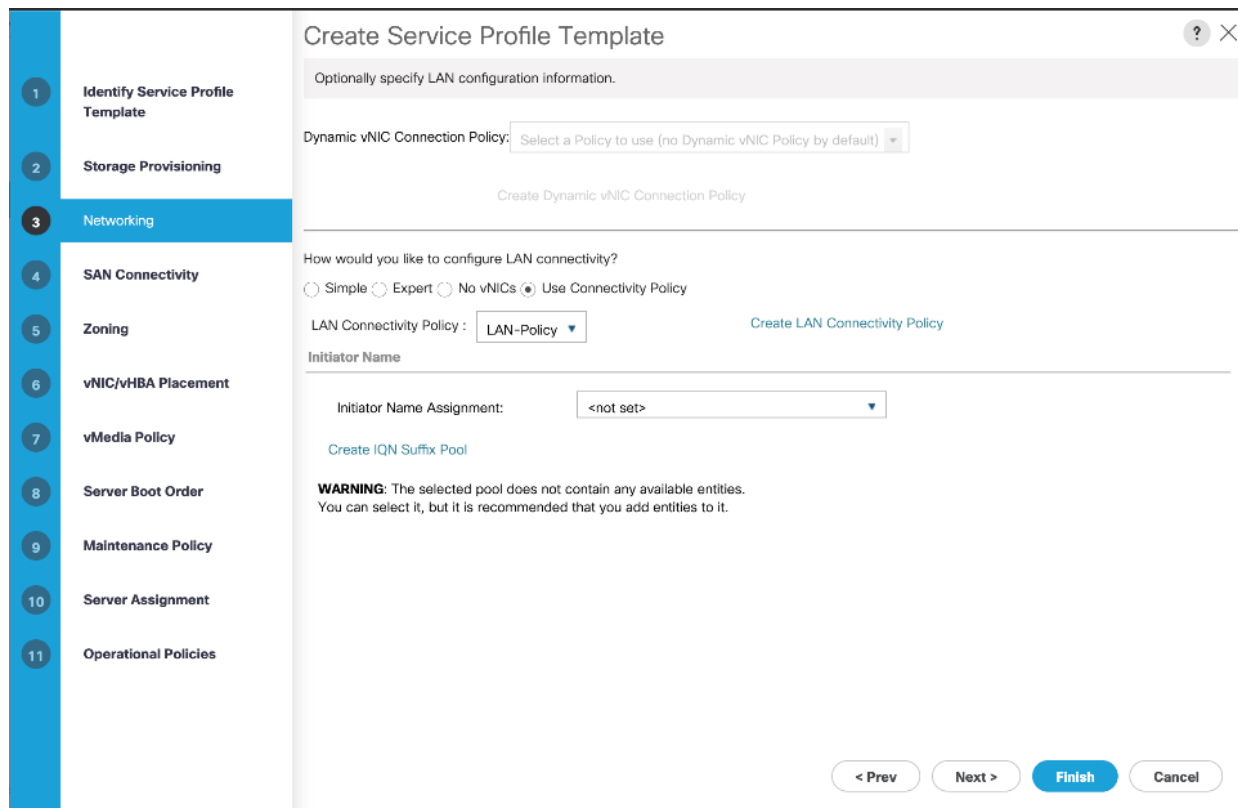
FlexFlash Removable State : **No Change**

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.

< Prev Next > **Finish** Cancel

Configure Networking Options

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Use Connectivity Policy option to configure the LAN connectivity.
3. Select the LAN-Policy as the LAN Connectivity Policy.
4. Click **Next**.



Configure SAN Connectivity

1. Select the Use Connectivity Policy option to configure the SAN connectivity.
2. Select the Infra-FC-Policy as the SAN Connectivity Policy.
3. Click **Next**.

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](#)

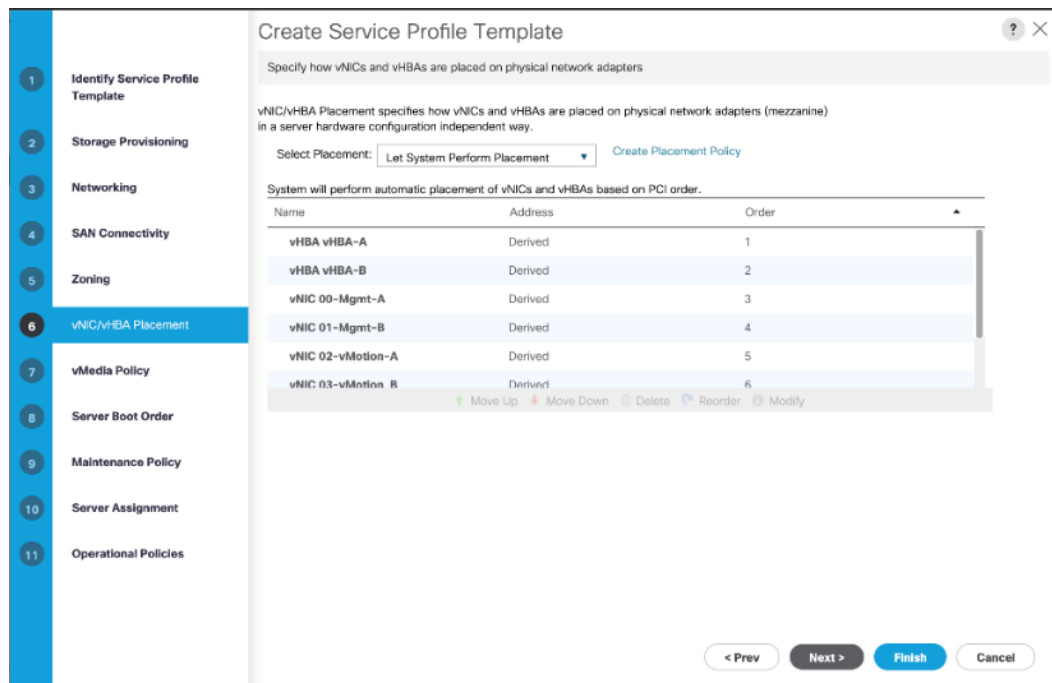
< Prev Next > **Finish** Cancel

Configure Zoning

1. It is not necessary to configure any Zoning options.
2. Click **Next**.

Configure vNIC/HBA Placement

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



Configure vMedia Policy

1. There is no need to set a vMedia Policy.
2. Click **Next**.

Configure Server Boot Order

1. Select **Boot-Fabric-A** as the Boot Policy
2. Verify all the boot devices are listed correctly
3. Click **Next**.

The screenshot shows the 'Create Service Profile Template' wizard. The left sidebar lists 11 steps, with '8 Server Boot Order' selected. The main panel shows the 'Boot Policy' configuration for 'Boot-Fabric-A'. The 'Boot Policy' dropdown is set to 'Boot-Fabric-A'. The configuration includes: Name: Boot-Fabric-A, Description: (empty), Reboot on Boot Order Change: No, Enforce vNIC/vHBA/ISCSI Name: Yes, and Boot Mode: Legacy. A 'WARNINGS' section explains that the type (primary/secondary) does not indicate a boot order presence and that the effective order is determined by PCIe bus scan order. Below this is a 'Boot Order' table with columns: Name, Order, vNIC/vHB..., Type, LUN Name, WWN, Slot Numb..., Boot Name, Boot Path, and Description. The table shows 'Remot...' at order 1 and 'San' at order 2. Under 'San', there are two entries: 'SA...' with 'Fabric-A' and 'Primary' type, and 'SA...' with 'Fabric-B' and 'Secondary' type. At the bottom, there are buttons for 'Create iSCSI vNIC', 'Set iSCSI Boot Parameters', 'Set UEFI Boot Parameters', '< Prev', 'Next >', 'Finish', and 'Cancel'.

Configure Maintenance Policy

1. Choose the default Maintenance Policy.
2. Click **Next**.

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	: Immediate

< Prev Next > **Finish** Cancel

Configure Server Assignment

1. For the Pool Assignment field, select `Infra-Server-Pool`.
2. **Optional:** Select a Server Pool Qualification policy.
3. Select the option Up for the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management and select `Infra-FW-Pack` from the Host Firmware list.
5. Click **Next**.

Create Service Profile Template ? X

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: [Create Server Pool](#)

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)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: [Create Host Firmware Package](#)

< Prev Next > **Finish** Cancel

Configure Operational Policies

1. For the BIOS Policy field, select `Infra-Host-BIOS`.
2. Expand Power Control Policy Configuration and select `No-Power-Cap` for the Power Control Policy field.

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 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 FC Boot Service Profiles (FC Deployment)

To create service profiles from the service profile template, follow these steps:

1. In Cisco UCS Manager, click the **Servers** tab in the navigation pane.
2. Choose Service Profile Templates > root > Service Template Infra-ESXi-Host (Infra-ESXi-iSCSI-Host for iSCSI Deployment).
3. Right-click and choose Create Service Profiles from Template.

The screenshot shows the Cisco UCS Manager interface. A context menu is open over the 'Service Template Infra-ESXi-Host' item in the left-hand tree view. The menu option 'Create Service Profiles From Template' is selected. The background shows the configuration page for this template, with the 'Network' tab active. On the right, a table titled 'vNIC/vHBA Placement Policy' is visible, showing four virtual slots, each with a selection preference of 'All'.

Virtual Slot	Selection Preference
1	All
2	All
3	All
4	All

4. Enter `Infra-ESXi-Host-` as the service profile prefix.
5. Enter 1 as the Name Suffix Starting Number.
6. Enter the Number of servers to be deploy in the Number of Instances field.
7. Click **OK** to create the service profile.

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

8. Click **OK** in the confirmation message.
9. Verify that the service profiles are successfully created and automatically associated with the servers from the pool.



Adjust the number of Service Profile instances based on the actual servers required for customer deployment.

Backup the Cisco UCS Manager Configuration

It is recommended to backup the Cisco UCS Configuration. For additional information, go to:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Management/4-o/b_Cisco_UCS_Admin_Mgmt_Guide_4-o/b_Cisco_UCS_Admin_Mgmt_Guide_4-o_chapter_01.html



Refer to the [Appendix](#) for example backup procedures

Add Servers

Additional server pools, service profile templates, and service profiles can be created under root or in organizations under the root. All the policies at the root level can be shared among the organizations. Any new physical blades can be added to the existing or new server pools and associated with the existing or new service profile templates.

Gather Necessary WWPN Information (FC Deployment)

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will be assigned certain unique configuration parameters. To proceed with the SAN configuration, this deployment specific information must be gathered from each Cisco UCS blade. Follow these steps:

1. To gather the vHBA WWPN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand **Servers > Service Profiles > root**. Select each service profile and expand to see the vHBAs.
2. Click **vHBAs** to see the WWPNs for both HBAs.

Name	WWPN	Desired Order	Actual Order	Fabric ID	Desired Placement	Actual Placement	Admin Host Port	Actual Host Port
vHBA vHBA-A	20:00:00:25:85:00:0A:02	1	7	A	Any	1	ANY	NONE
vHBA vHBA-B	20:00:00:25:85:00:0B:02	2	8	B	Any	1	ANY	NONE

3. Record the WWPN information that is displayed for both the Fabric A vHBA and the Fabric B vHBA for each service profile into the WWPN variable in Table 15 . Add or remove rows from the table depending on the number of ESXi hosts.

Table 15 Cisco UCS WWPN Information

Host	vHBA		Value
Infra-ESXi-Host-1	Fabric-A	WWPN-Infra-ESXi-Host-1-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-1-B	20:00:00:25:b5:

Infra-ESXi-Host-2	Fabric-A	WWPN-Infra-ESXi-Host-2-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-2-B	20:00:00:25:b5:
Infra-ESXi-Host-3	Fabric-A	WWPN-Infra-ESXi-Host-3-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-3-B	20:00:00:25:b5:
Infra-ESXi-Host-4	Fabric-A	WWPN-Infra-ESXi-Host-4-A	20:00:00:25:b5:
	Fabric-B	WWPN-Infra-ESXi-Host-4-B	20:00:00:25:b5:

Gather Necessary IQN Information (iSCSI Deployment)

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will be assigned certain unique configuration parameters. To proceed with the SAN configuration, this deployment specific information must be gathered from each Cisco UCS blade. Follow these steps:

1. To gather the vNIC IQN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the **Servers** tab. Expand Servers > Service Profiles > root.
2. Click each service profile and then click the "iSCSI vNICs" tab on the right. Note "Initiator Name" displayed at the top of the page under "Service Profile Initiator Name."

Actions

[Change Initiator Name](#)

[Reset Initiator Name](#)

Service Profile Initiator Name

IQN Pool Name : **Infra-IQN-Pool**

Initiator Name : **iqn.1992-08.com.cisco:ucs-host:3**

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		Derived
iSCSI vNIC iSCSI-B-vNIC	07-iSCSI-B		Derived

Table 16 Cisco UCS iSCSI IQNs

Cisco UCS Service Profile Name	iSCSI IQN
Infra-ESXi-iSCSI-Host-01	iqn.1992-08.com.cisco:ucs-host____
Infra-ESXi-iSCSI-Host-02	iqn.1992-08.com.cisco:ucs-host____
Infra-ESXi-iSCSI-Host-03	iqn.1992-08.com.cisco:ucs-host____

Cisco UCS Service Profile Name	iSCSI IQN
Infra-ESXi—iSCSI-Host-04	iqn.1992-08.com.cisco:ucs-host____

IBM FS9100 iSCSI Storage Configuration (iSCSI Deployment)



This configuration step can be skipped if the UCS environment does not need access to storage using iSCSI.

As part of IBM FS9100 storage configuration, follow these steps:

1. Create ESXi boot Volumes (Boot LUNs for all the ESXi hosts)
2. Create Share Storage Volumes (for hosting VMs)
3. Map Volumes to Hosts

Table 17 List of Volumes for iSCSI on IBM FS9100*

Volume Name	Capacity (GB)	Purpose	Mapping
Infra-ESXi-iSCSI-Host-01	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-01
Infra-ESXi-iSCSI-Host-02	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-02
Infra-ESXi-iSCSI-Host-03	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-03
Infra-ESXi-iSCSI-Host-04	10	Boot LUN for the Host	Infra-ESXi-iSCSI-Host-04
Infra-iSCSI-datastore-1	2000**	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-iSCSI-Host-01 to Infra-ESXi-iSCSI-Host-04
Infra-iSCSI-datastore-2	2000**	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-iSCSI-Host-01 to Infra-ESXi-iSCSI-Host-04
Infra-iSCSI-swap	500**	Shared volume to host VMware VM swap directory	All ESXi hosts: Infra-ESXi-iSCSI-Host-01 to Infra-ESXi-iSCSI-Host-04

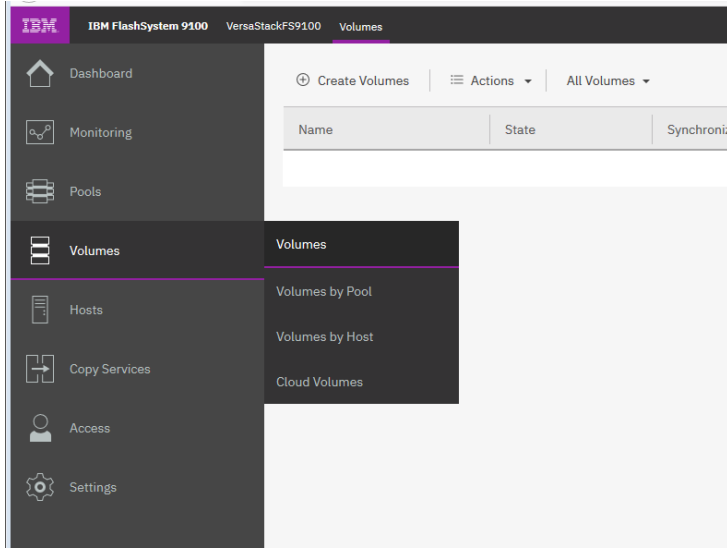
* Customers should adjust the names and values used for server and volumes names based on their deployment

** The volume size can be adjusted based on customer requirements

Create Volumes on the Storage System

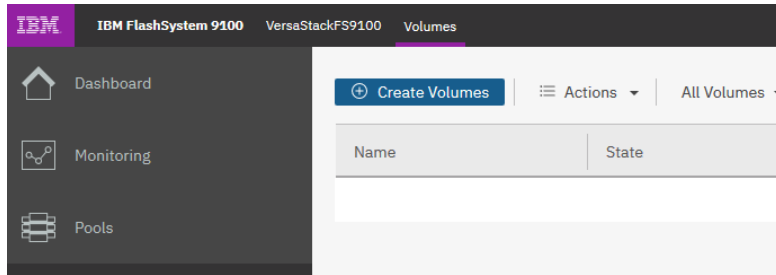
To create volumes on the storage system, follow these steps:

1. Log into the IBM FS9100 GUI and select the **Volumes** icon on the left screen and select **Volumes**



You will repeat the following steps to create and map the volumes shown in Table 17 .

2. Click **Create Volumes** as shown below.



3. Click **Basic** and then select the pool (VS-Pool10 in this example) from the drop-down list.
4. When creating single volumes, input quantity 1 and the capacity and name from Table 17 . Select *Thin-provisioned* for Capacity savings and enter the Name of the volume. Select I/O group io_grpo.
5. When creating multiple volumes in bulk enter the quantity required and review the Name field. The number value will be appended to the specified volume name.




IBM FS9100 and Spectrum Virtualize is optimized for environments with more than 30 volumes. Consider distributing Virtual Machines over multiple VMFS datastores for optimal performance.

Create Volumes

Basic
Mirrored
Custom

Create a preset volume with all the basic features.

Pool:
 VS-Pool0


Total 7.10 TiB

Volume Details


Quantity:	Capacity:	Name:	
4	10 GiB	<input style="width: 100%;" type="text" value="Infra-ESXi-iSCSI-Host-0"/>	1 - 4

Capacity savings:

Deduplicated

[+ Define another volume](#)

I/O group:
 Automatic

 **Summary**

6. Click **Create**.



During the volume creation, expand **view more details** to monitor the CLI commands utilized to create each volume. All commands run against the system by either the GUI or CLI will be stored in the Audit log, along with the associated user account and timestamp.

7. Repeat steps 1-6 to create all the required volumes and verify all the volumes have successfully been created as shown in the sample output below.

Name	State	Synchronized	Pool	Protocol Type	UID
Infra-ESXi-iSCSI-Host-01	Online		VS-Pool0		60050768109300003000000000000000...
Infra-ESXi-iSCSI-Host-02	Online		VS-Pool0		60050768109300003000000000000000...
Infra-ESXi-iSCSI-Host-03	Online		VS-Pool0		60050768109300003000000000000000...
Infra-ESXi-iSCSI-Host-04	Online		VS-Pool0		60050768109300003000000000000000...
Infra-iSCSI-datastore-1	Online		VS-Pool0		60050768109300003000000000000000...
Infra-iSCSI-swap	Online		VS-Pool0		60050768109300003000000000000000...

Create Host Cluster & Host objects

Host Cluster Shared & Private mappings

In traditional hypervisor environments such as VMware vSphere, each physical host requires access to the same shared datastores (or LUNs) in order to facilitate features such as vMotion, High Availability, Fault Tolerance. It is important for all ESXi hosts within a vSphere cluster to have identical access to LUNs presented from the FS9100.

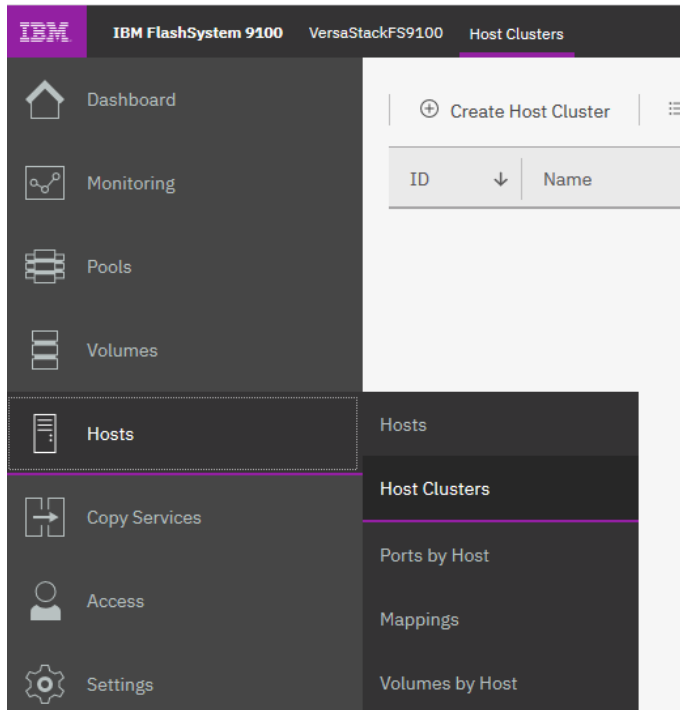
The Host Clusters feature in IBM Spectrum Virtualize products introduces a way to simplify administration when mapping volumes to host environments that require shared storage.



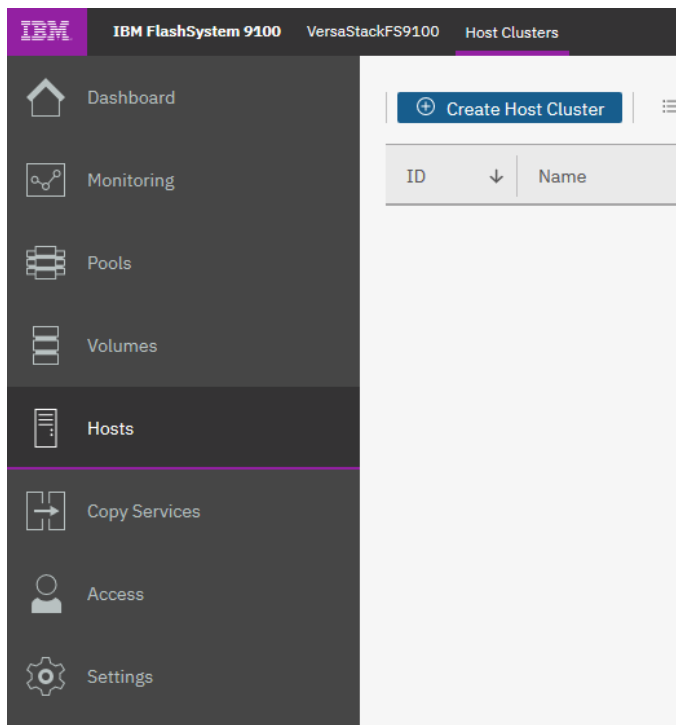
It is recommended that a Host Cluster object be created for each vSphere Cluster visible in vCenter, and any ESXi hosts within the vSphere cluster be defined as individual host objects within FS9100. This ensures that volume access is consistent across all members of the cluster and any hosts that are subsequently added to the Host Cluster will inherit the same LUN mappings.

To create host clusters and objects, follow these steps:

1. Click Hosts then click Host Clusters.



2. Click **Create Host Cluster**.



3. Give the Host Cluster a friendly name.

Create Host Cluster ✕

Name:

Optional: Select hosts to assign to a new host cluster. Any current volume mappings become the shared mappings for all the hosts in the host cluster.

i It is recommended that all hosts in a host cluster have access to the same I/O Groups.

↓ ↕

Name	Status	Host Type	Host Mappings	P III
------	--------	-----------	---------------	-------

No items found.

< >

[? Need Help](#)

[Cancel](#)

[◀ Back](#)

[Next ▶](#)

- Review the summary and click **Make Host Cluster**.

Create Host Cluster: Summary x

An empty host cluster **VS-UCS01** will be created.

Cancel

◀ Back

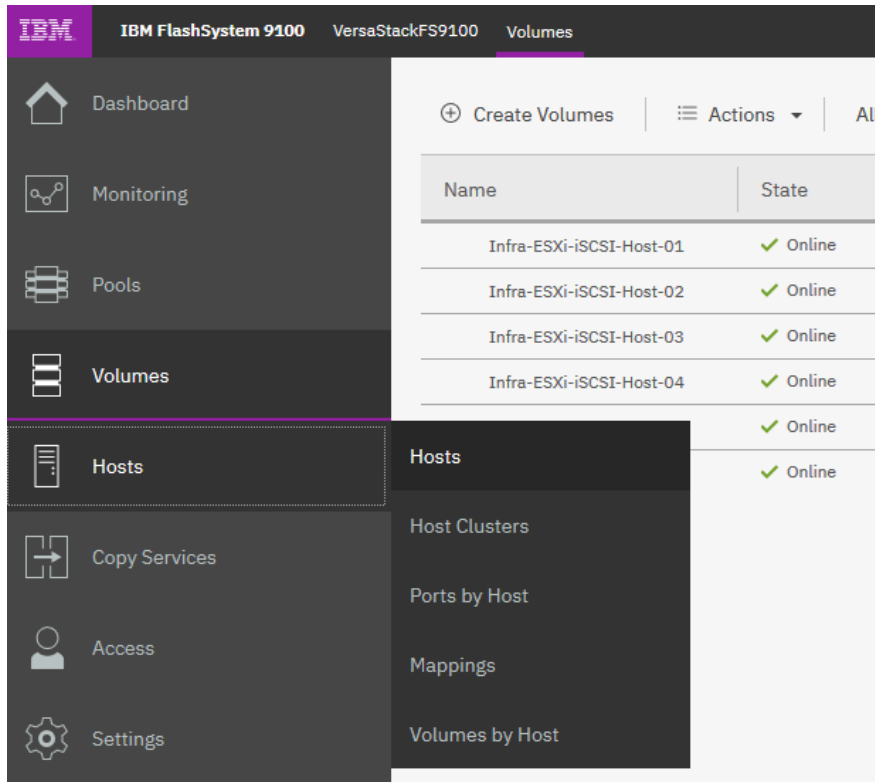
Make Host Cluster

Add Hosts to Host Cluster

Create iSCSI Host Definitions

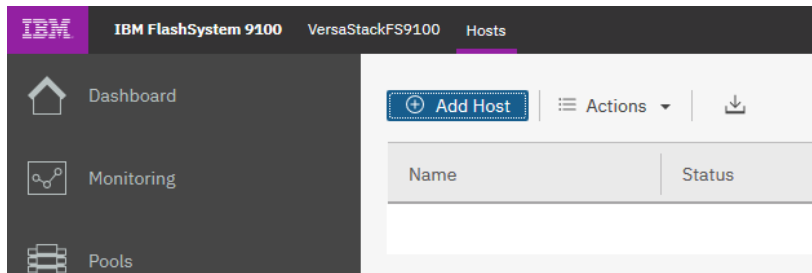
To create iSCSI host definitions, follow these steps:

1. Click **Hosts** and then **Hosts** from the navigation menu.



2. For each ESXi host (Table 16), follow these steps on the IBM FS9100 system:

- a. Click Add Host.



- b. Select iSCSI or iSER (SCSI) Host. Add the name of the host to match the ESXi service profile name from Table 17 . Type the IQN corresponding to the ESXi host from Table 16 and select the Host Cluster that we created in the previous step.

Add Host ✕

Required Fields

Name:

Host connections:

Host IQN: ⊕ ⊖

Optional Fields

CHAP authentication:

CHAP secret:

CHAP username:

Host type:

I/O groups:

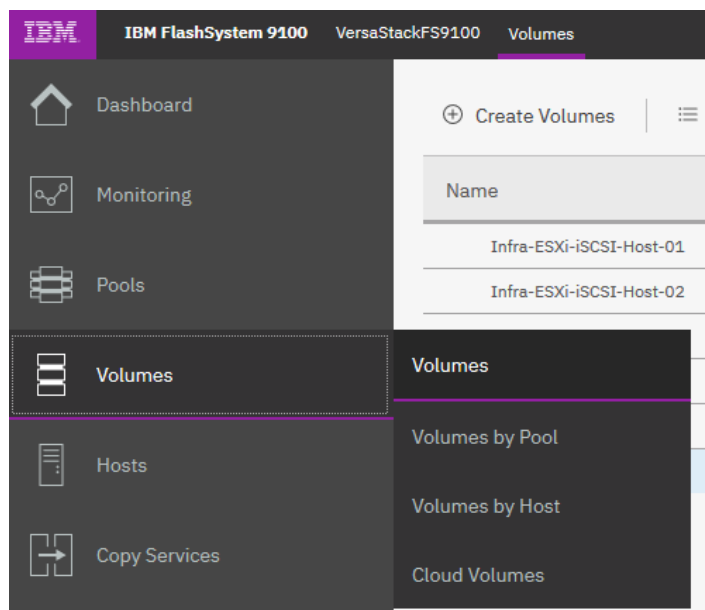
Host cluster:

3. Click **Add**.

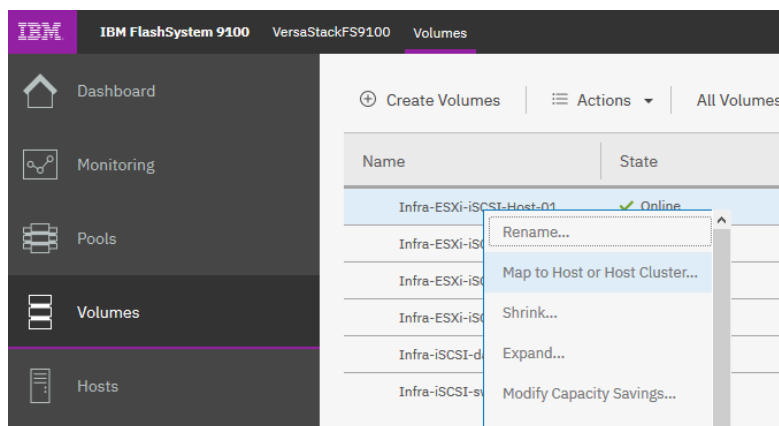
Map Volumes to Hosts and Host Cluster

To map volumes to hosts and clusters, follow these steps:

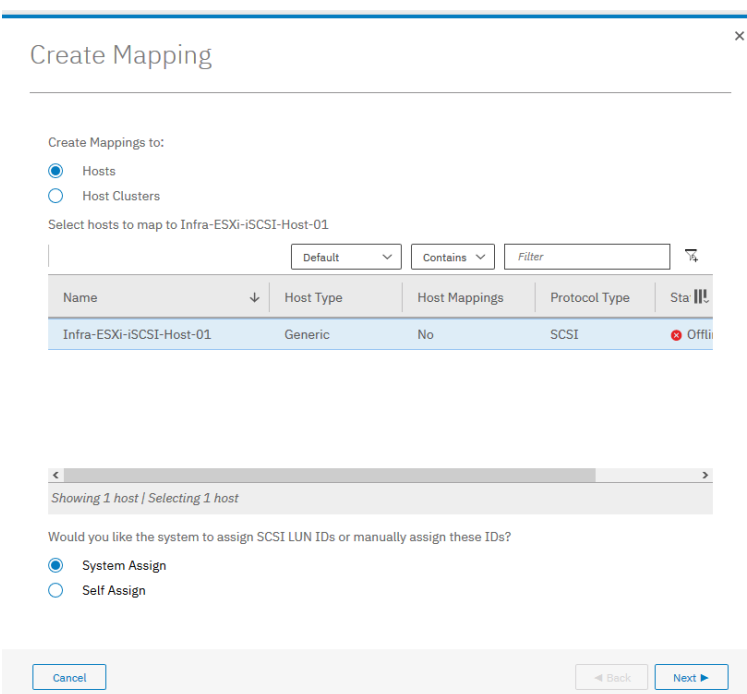
1. Now the Host Cluster and Host objects have been created, we need to map each LUN to the hosts.
2. Click **Volumes**.



3. Right-click the Boot LUN for each ESXi host in turn and choose **Map to Host**.



4. Select the **Hosts** radio button, and select the corresponding Host in the list and click **Next**



5. Click **Map Volumes** and when the process is complete, click **Close**.

Map Volumes to Infra-ESXi-iSCSI-Host-01: Summary x

The following volumes will be mapped to Infra-ESXi-iSCSI-Host-01:

Name	SCSI ID	Caching I/O Group ID	New Mapping
Infra-ESXi-iSCSI-Host-01	0	0	New

[Cancel](#)

[← Back](#)

[Map Volumes](#)

- Repeat steps 1-5 to map a Boot volume for each ESXi host in the cluster.
- When mapping shared volumes from Table 17 i.e. for shared VMFS datastores, right click on the volume in question (or select multiple volumes if mapping multiple LUNs) and select **Map to Host or Host Cluster**.

The screenshot displays the IBM FlashSystem 9100 Volumes management interface. The left sidebar contains navigation options: Dashboard, Monitoring, Pools, Volumes (selected), Hosts, Copy Services, Access, and Settings. The main content area shows a table of volumes with columns for Name, State, and Sync. The volumes listed are:

Name	State	Sync
Infra-ESXi-iSCSI-Host-01	✓ Online	
Infra-ESXi-iSCSI-Host-02	✓ Online	
Infra-ESXi-iSCSI-Host-03	✓ Online	
Infra-ESXi-iSCSI-Host-04	✓ Online	
Infra-iSCSI-datastore-1	✓ Online	
Infra-iSCSI-datastore-2	✓ Online	
Infra-iSCSI-datastore-3	✓ Online	
Infra-iSCSI-datastore-4	✓ Online	
Infra-iSCSI-datastore-5	✓ Online	
Infra-iSCSI-datastore-6	✓ Online	
Infra-iSCSI-swap		

A context menu is open over the 'Infra-iSCSI-datastore-5' volume, showing the following options:

- Rename...
- Map to Host or Host Cluster...
- Shrink...
- Expand...
- Modify Capacity Savings...

- Select the **Host Clusters** radio button.

Create Mapping ×

Create Mappings to:

Hosts

Host Clusters

Select host clusters to map to 6 volumes

Default Filter 1/4

Name	Status	Host Count	Mappings Count	
VS-UCS01	✘ Offline	1	0	1

< >

Showing 1 host cluster | Selecting 1 host cluster

Would you like the system to assign SCSI LUN IDs or manually assign these IDs?

System Assign

Self Assign

9. Review the summary and Click **Map Volumes** to confirm.

Map Volumes to VS-UCS01: Summary ×

The following volumes will be mapped to VS-UCS01:

Name	SCSI ID	Caching I/O Group ID	New Mapping	
Infra-iSCSI-datastore-1	1	0	New	
Infra-iSCSI-datastore-2	2	0	New	
Infra-iSCSI-datastore-3	3	0	New	
Infra-iSCSI-datastore-4	4	0	New	
Infra-iSCSI-datastore-5	5	0	New	
Infra-iSCSI-datastore-6	6	0	New	

[Cancel](#)

[← Back](#)

[Map Volumes](#)

- Any Shared host cluster mappings will be automatically inherited by any future ESXi hosts which are defined as members of the host cluster.

IBM FS9100 Fibre Channel Storage Configuration (FC Deployment)



This configuration step can be skipped if the UCS environment does not need access to storage using fibre channel.

As part of IBM FS9100 Fibre Channel storage configuration, follow these steps:

1. Setup Zoning on Cisco MDS switches
2. Setup Volumes on IBM FS9100
3. Map Volumes to Hosts

Create Device Aliases and SAN Zoning

The following steps will configure zoning for the WWPNs for the UCS hosts and the IBM FS9100 node canisters. WWPN information collected from the previous steps will be used in this section. Multiple zones will be created for servers in VSAN 101 on Switch A and VSAN 102 on Switch B.



The configuration below assumes 4 UCS services profiles have been deployed in this example. Customers can adjust the configuration according to their deployment size.

Cisco MDS - A Switch

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

The WWPNs recorded in Table 15 will be used in the next step. Replace the variables with actual WWPN values.

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

```
device-alias database
device-alias name Infra-ESXi-Host-01 pwwn <WWPN-Infra-ESXi-Host-1-A>
device-alias name Infra-ESXi-Host-02 pwwn <WWPN-Infra-ESXi-Host-2-A>
device-alias name Infra-ESXi-Host-03 pwwn <WWPN-Infra-ESXi-Host-3-A>
device-alias name Infra-ESXi-Host-04 pwwn <WWPN-Infra-ESXi-Host-4-A>
device-alias name FS9100-Node1-FC1 pwwn <WWPN-FS9100-Node1-FC1>
device-alias name FS9100-Node1-FC3 pwwn <WWPN-FS9100-Node1-FC3>
device-alias name FS9100-Node2-FC1 pwwn <WWPN-FS9100-Node2-FC1>
device-alias name FS9100-Node2-FC3 pwwn <WWPN-FS9100-Node2-FC3>
device-alias commit
```

2. Create the zones and add device-alias members for the 4 blades.

```
zone name Infra-ESXi-Host-01 vsan 101
member device-alias Infra-ESXi-Host-01
```

```

member device-alias FS9100-Node1-FC1
member device-alias FS9100-Node1-FC3
member device-alias FS9100-Node2-FC1
member device-alias FS9100-Node2-FC3
!
zone name Infra-ESXi-Host-02 vsan 101
member device-alias Infra-ESXi-Host-02
member device-alias FS9100-Node1-FC1
member device-alias FS9100-Node1-FC3
member device-alias FS9100-Node2-FC1
member device-alias FS9100-Node2-FC3
!
zone name Infra-ESXi-Host-03 vsan 101
member device-alias Infra-ESXi-Host-03
member device-alias FS9100-Node1-FC1
member device-alias FS9100-Node1-FC3
member device-alias FS9100-Node2-FC1
member device-alias FS9100-Node2-FC3
!
zone name Infra-ESXi-Host-04 vsan 101
member device-alias Infra-ESXi-Host-04
member device-alias FS9100-Node1-FC1
member device-alias FS9100-Node1-FC3
member device-alias FS9100-Node2-FC1
member device-alias FS9100-Node2-FC3
!

```

3. Add zones to zoneset.

```

zoneset name versastackzoneset vsan 101
    member Infra-ESXi-Host-01
    member Infra-ESXi-Host-02
    member Infra-ESXi-Host-03
    member Infra-ESXi-Host-04

```

4. Activate the zoneset.

```

zoneset activate name versastackzoneset vsan 101

```



Validate all the HBA's are logged into the MDS switch. The FS9100 nodes and the Cisco servers should be powered on. To start the Cisco servers from Cisco UCS Manager, select the server tab, then click Servers>Service>Profiles>root, and right-click service profile then select boot server.

5. Validate that all the powered-on system's HBAs are logged into the switch through the show zoneset command.

```
show zoneset active
```

```
VersaStack-MDS-A# sh zoneset active
zoneset name versastackzoneset vsan 101

  zone name Infra-ESXi-Host-01 vsan 101
    pwnn 20:00:00:25:b5:00:0a:00 [Infra-ESXi-Host-01]
    * fcid 0x720021 [pwnn 50:05:07:68:10:15:05:16] [VS-FS9100-Node1-FC1]
    * fcid 0x720041 [pwnn 50:05:07:68:10:17:05:16] [VS-FS9100-Node1-FC3]
    * fcid 0x720061 [pwnn 50:05:07:68:10:15:05:0d] [VS-FS9100-Node2-FC1]
    * fcid 0x720081 [pwnn 50:05:07:68:10:17:05:0d] [VS-FS9100-Node2-FC3]

  zone name Infra-ESXi-Host-02 vsan 101
    pwnn 20:00:00:25:b5:00:0a:01 [Infra-ESXi-Host-02]
    * fcid 0x720021 [pwnn 50:05:07:68:10:15:05:16] [VS-FS9100-Node1-FC1]
    * fcid 0x720041 [pwnn 50:05:07:68:10:17:05:16] [VS-FS9100-Node1-FC3]
    * fcid 0x720061 [pwnn 50:05:07:68:10:15:05:0d] [VS-FS9100-Node2-FC1]
    * fcid 0x720081 [pwnn 50:05:07:68:10:17:05:0d] [VS-FS9100-Node2-FC3]

  zone name Infra-ESXi-Host-03 vsan 101
    * fcid 0x7200e4 [pwnn 20:00:00:25:b5:00:0a:02] [Infra-ESXi-Host-03]
    * fcid 0x720021 [pwnn 50:05:07:68:10:15:05:16] [VS-FS9100-Node1-FC1]
    * fcid 0x720041 [pwnn 50:05:07:68:10:17:05:16] [VS-FS9100-Node1-FC3]
    * fcid 0x720061 [pwnn 50:05:07:68:10:15:05:0d] [VS-FS9100-Node2-FC1]
    * fcid 0x720081 [pwnn 50:05:07:68:10:17:05:0d] [VS-FS9100-Node2-FC3]

  zone name Infra-ESXi-Host-04 vsan 101
    * fcid 0x7200e1 [pwnn 20:00:00:25:b5:00:0a:03] [Infra-ESXi-Host-04]
    * fcid 0x720021 [pwnn 50:05:07:68:10:15:05:16] [VS-FS9100-Node1-FC1]
    * fcid 0x720041 [pwnn 50:05:07:68:10:17:05:16] [VS-FS9100-Node1-FC3]
    * fcid 0x720061 [pwnn 50:05:07:68:10:15:05:0d] [VS-FS9100-Node2-FC1]
    * fcid 0x720081 [pwnn 50:05:07:68:10:17:05:0d] [VS-FS9100-Node2-FC3]
```

6. Save the configuration.

```
copy run start
```

Cisco MDS - B Switch

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

The WWPNs recorded in Table 15 will be used in the next step. Replace the variables with actual WWPN values.

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

```
device-alias database
device-alias name Infra-ESXi-Host-01 pwwn <WWPN-Infra-ESXi-Host-1-B>
device-alias name Infra-ESXi-Host-02 pwwn <WWPN-Infra-ESXi-Host-2-B>
device-alias name Infra-ESXi-Host-03 pwwn <WWPN-Infra-ESXi-Host-3-B>
device-alias name Infra-ESXi-Host-04 pwwn <WWPN-Infra-ESXi-Host-4-B>
device-alias name FS9100-Node1-FC2-NPIV pwwn <WWPN-FS9100-Node1-FC2>
device-alias name FS9100-Node1-FC4-NPIV pwwn <WWPN-FS9100-Node1-FC4>
device-alias name FS9100-Node2-FC2-NPIV pwwn <WWPN-FS9100-Node2-FC2>
device-alias name FS9100-Node2-FC4-NPIV pwwn <WWPN-FS9100-Node2-FC4>
device-alias commit
```

2. Create the zones and add device-alias members for the 4 blades.

```
zone name Infra-ESXi-Host-01 vsan 102
member device-alias Infra-ESXi-Host-01
member device-alias FS9100-Node1-FC2
member device-alias FS9100-Node1-FC4
member device-alias FS9100-Node2-FC2
member device-alias FS9100-Node2-FC4
!
zone name Infra-ESXi-Host-02 vsan 102
member device-alias Infra-ESXi-Host-02
member device-alias FS9100-Node1-FC2-NPIV
member device-alias FS9100-Node1-FC4-NPIV
member device-alias FS9100-Node2-FC2-NPIV
member device-alias FS9100-Node2-FC4-NPIV
!
zone name Infra-ESXi-Host-03 vsan 102
member device-alias Infra-ESXi-Host-03
```

```

member device-alias FS9100-Node1-FC2-NPIV
member device-alias FS9100-Node1-FC4-NPIV
member device-alias FS9100-Node2-FC2-NPIV
member device-alias FS9100-Node2-FC4-NPIV
!
zone name Infra-ESXi-Host-04 vsan 102
member device-alias Infra-ESXi-Host-04
member device-alias FS9100-Node1-FC2-NPIV
member device-alias FS9100-Node1-FC4-NPIV
member device-alias FS9100-Node2-FC2-NPIV
member device-alias FS9100-Node2-FC4-NPIV
!

```

3. Add zones to zoneset.

```

zoneset name versastackzoneset vsan 102
    member Infra-ESXi-Host-01
    member Infra-ESXi-Host-02
    member Infra-ESXi-Host-03
    member Infra-ESXi-Host-04

```

4. Activate the zoneset.

```
zoneset activate name versastackzoneset vsan 102
```



Validate all the HBA's are logged into the MDS switch. The FS9100 nodes and the Cisco servers should be powered on. To start the Cisco servers from Cisco UCS Manager, select the server tab, then click Servers→Service→Profiles→root, and right-click service profile then select boot server.

5. Validate that all the powered-on system's HBAs are logged into the switch through the show zoneset command.

```
show zoneset active
```

```

VersaStack-MDS-B# sh zoneset active
zoneset name versastackzoneset vsan 102
  zone name Infra-ESXi-Host-01 vsan 102
    pwnn 20:00:00:25:b5:00:0b:00 [Infra-ESXi-Host-01]
    * fcid 0xc60041 [pwnn 50:05:07:68:10:16:05:16] [VS-FS9100-Node1-FC2]
    * fcid 0xc60061 [pwnn 50:05:07:68:10:18:05:16] [VS-FS9100-Node1-FC4]
    * fcid 0xc60081 [pwnn 50:05:07:68:10:16:05:0d] [VS-FS9100-Node2-FC2]
    * fcid 0xc600a1 [pwnn 50:05:07:68:10:18:05:0d] [VS-FS9100-Node2-FC4]

```

```

zone name Infra-ESXi-Host-02 vsan 102
  pwnn 20:00:00:25:b5:00:0b:01 [Infra-ESXi-Host-02]
* fcid 0xc60041 [pwnn 50:05:07:68:10:16:05:16] [VS-FS9100-Node1-FC2]
* fcid 0xc60061 [pwnn 50:05:07:68:10:18:05:16] [VS-FS9100-Node1-FC4]
* fcid 0xc60081 [pwnn 50:05:07:68:10:16:05:0d] [VS-FS9100-Node2-FC2]
* fcid 0xc600a1 [pwnn 50:05:07:68:10:18:05:0d] [VS-FS9100-Node2-FC4]

zone name Infra-ESXi-Host-03 vsan 102
* fcid 0xc60103 [pwnn 20:00:00:25:b5:00:0b:02] [Infra-ESXi-Host-03]
* fcid 0xc60041 [pwnn 50:05:07:68:10:16:05:16] [VS-FS9100-Node1-FC2]
* fcid 0xc60061 [pwnn 50:05:07:68:10:18:05:16] [VS-FS9100-Node1-FC4]
* fcid 0xc60081 [pwnn 50:05:07:68:10:16:05:0d] [VS-FS9100-Node2-FC2]
* fcid 0xc600a1 [pwnn 50:05:07:68:10:18:05:0d] [VS-FS9100-Node2-FC4]

zone name Infra-ESXi-Host-04 vsan 102
* fcid 0xc60104 [pwnn 20:00:00:25:b5:00:0b:03] [Infra-ESXi-Host-04]
* fcid 0xc60041 [pwnn 50:05:07:68:10:16:05:16] [VS-FS9100-Node1-FC2]
* fcid 0xc60061 [pwnn 50:05:07:68:10:18:05:16] [VS-FS9100-Node1-FC4]
* fcid 0xc60081 [pwnn 50:05:07:68:10:16:05:0d] [VS-FS9100-Node2-FC2]
* fcid 0xc600a1 [pwnn 50:05:07:68:10:18:05:0d] [VS-FS9100-Node2-FC4]
    
```

6. Save the configuration.

```
copy run start
```

IBM FS9100 FC Configuration

As part of IBM FS9100 FC configuration, follow these steps:

1. Create ESXi boot Volumes (Boot LUNs for all the ESXi hosts).
2. Create Share Storage Volumes (for hosting VMs).
3. Map Volumes to Hosts.



In this deployment example, there are four ESXi hosts. The volumes listed in Table 18 will be created in this process.

Table 18 List of FC Volumes on IBM FS9100*

Volume Name	Capacity (GB)	Purpose	Mapping
-------------	---------------	---------	---------

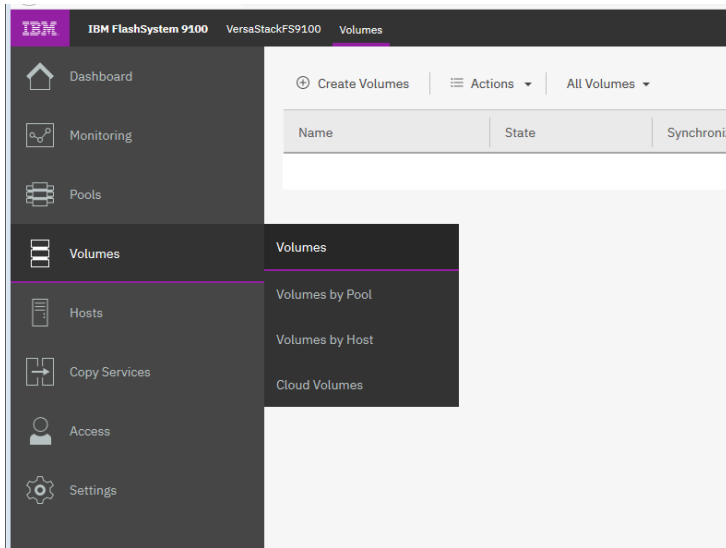
Volume Name	Capacity (GB)	Purpose	Mapping
Infra-ESXi-Host-01	10	Boot LUN for the Host	Infra-ESXi-Host-01
Infra-ESXi-Host-02	10	Boot LUN for the Host	Infra-ESXi-Host-02
Infra-ESXi-Host-03	10	Boot LUN for the Host	Infra-ESXi-Host-03
Infra-ESXi-Host-04	10	Boot LUN for the Host	Infra-ESXi-Host-04
Infra-datastore-1	2000*	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-Host-01 to Infra-ESXi-Host-04
Infra-datastore-2	2000*	Shared volume to host VMs	All ESXi hosts: Infra-ESXi-Host-01 to Infra-ESXi-Host-04
Infra-swap	500*	Shared volume to host VMware VM swap directory	All ESXi hosts: Infra-ESXi-Host-01 to Infra-ESXi-Host-04

* Customers should adjust the names and values based on their environment.

Create Volumes on the Storage System

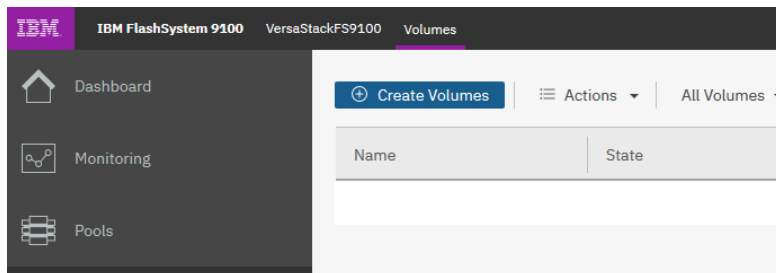
To create volumes on the storage system, follow these steps:

1. Log into the IBM FS9100 GUI and select the **Volumes** icon on the left screen and select **Volumes**.




You will repeat the following steps to create and map the volumes shown in [Table 18](#) .

2. Click **Create Volumes** as shown below.



3. Click **Basic** and the select the pool (VS-Pool0 in this example) from the drop-down list.
4. When creating single volumes, input quantity 1 and the capacity and name from Table 18 .Select Thin-provisioned for Capacity savings and enter the Name of the volume. Select I/O group io_grp0.
5. When creating multiple volumes in bulk enter the quantity required and review the Name field. The number value will be appended to the specified volume name.

 IBM FS9100 and Spectrum Virtualize is optimized for environments with more than 30 volumes. Consider distributing Virtual Machines over multiple VMFS datastores for optimal performance.

Create Volumes ×

Basic
Mirrored
Custom

Create a preset volume with all the basic features.

Pool:

VS-Pool0 ▾

Effective Capacity:

Total 6.98 TiB

Volume Details

Quantity:	Capacity:	Name:	
4 ▴ ▾	10 GiB ▾	Infra-ESXi-Host-0	1 - 4

Capacity savings:

Thin-provisioned ▾

 Deduplicated


?

Cancel

Create and Map

Create

6. Click **Create**.

 During the volume creation, expand **view more details** to monitor the CLI commands utilized to create each volume. All commands run against the system by either the GUI or CLI will be stored in the Audit log, along with the associated user account and timestamp.

7. Repeat steps 1-6 above to create all the required volumes and verify all the volumes have successfully been created as shown in the sample output below.

Name	State	Synchronized	Pool	Protocol Type	LUN	Host Mappings	Capa
Infra-ESX-Host-01	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra-ESX-Host-02	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra-ESX-Host-03	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra-ESX-Host-04	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra-ESX-iSCSI-Host-01	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra-ESX-iSCSI-Host-02	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra_Datastore1	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra_Datastore2	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	
Infra_Swap	Online		VS-Pool0	SCSI	630507681080028B00000000000000...	Yes	

Create Host Cluster & Host Objects

Host Cluster Shared & Private Mappings

In traditional hypervisor environments such as VMware vSphere, each physical host requires access to the same shared datastores (or LUNs) in order to facilitate features such as vMotion, High Availability, Fault Tolerance. It is important for all ESXi hosts within a vSphere cluster to have identical access to LUNs presented from the FS9100.

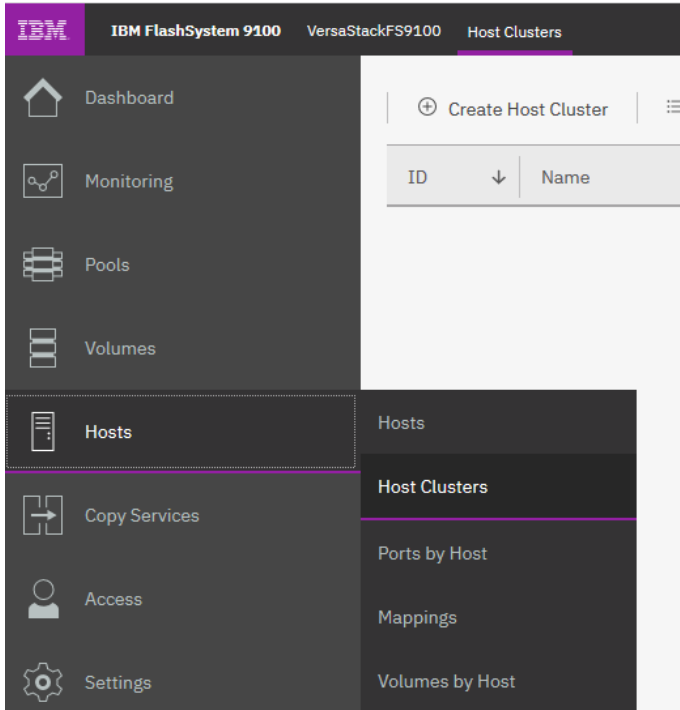
The Host Clusters feature in IBM Spectrum Virtualize products introduces a way to simplify administration when mapping volumes to host environments that require shared storage.




It is recommended that a Host Cluster object be created for each vSphere Cluster visible in vCenter, and any ESXi hosts within the vSphere cluster be defined as individual host objects within FS9100. This ensures that volume access is consistent across all members of the cluster and any hosts that are subsequently added to the Host Cluster will inherit the same LUN mappings.

To create host clusters and host objects, follow these steps:

1. Click Hosts then click Host Clusters.



2. Click Create Host Cluster.



3. Give the Host Cluster a friendly name.

Create Host Cluster ×

Name:

Optional: Select hosts to assign to a new host cluster. Any current volume mappings become the shared mappings for all the hosts in the host cluster.

i It is recommended that all hosts in a host cluster have access to the same I/O Groups.

↓
Default ▾
Contains ▾

↕

Name	↓	Status	Host Type	Host Mappings	P III
------	---	--------	-----------	---------------	-------

No items found.

<
>

Showing 0 hosts | Selecting 0 hosts

Need Help

◀ Back

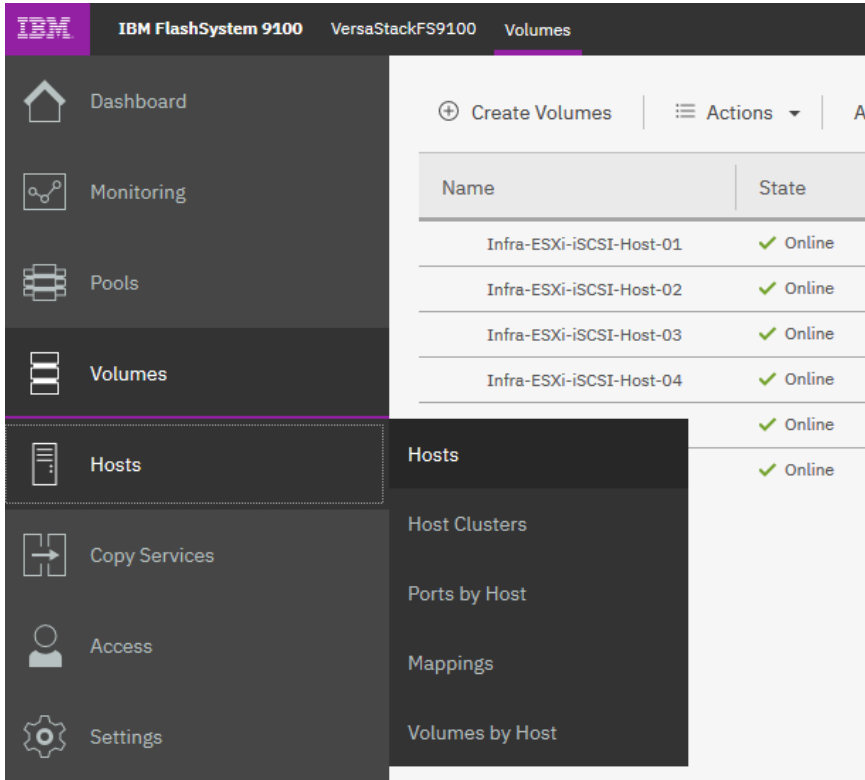
- Review the summary and click **Make Host Cluster**

Create Host Cluster: Summary ×

An empty host cluster **VS-UCS01** will be created.

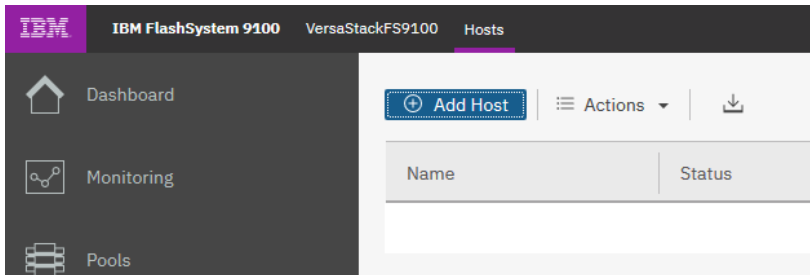
◀ Back

- Click **Hosts** and then **Hosts** from the navigation menu.



For each ESXi host (Table 18), follow these steps on the IBM FS9100 system:

6. Click Add Host.



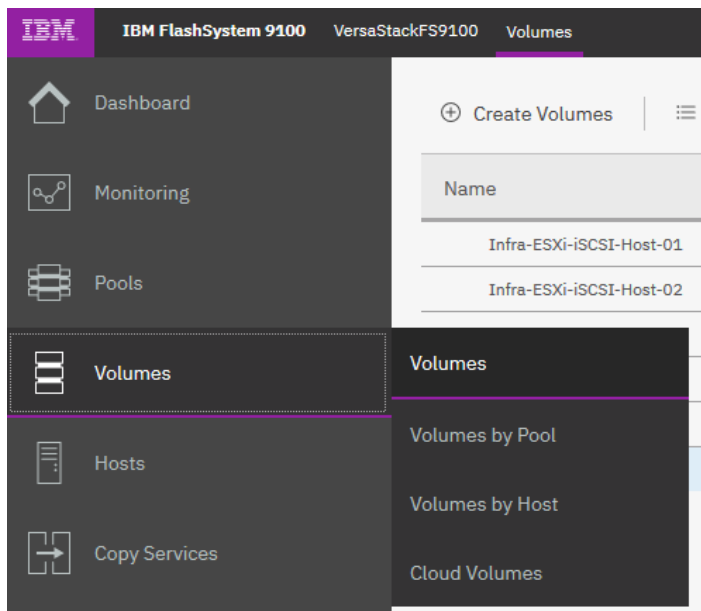
7. Select Fibre Channel (SCSI) . Add the name of the host to match the ESXi service profile name from Table 18 .
8. From the drop-down list, select both (Fabric A and B) WWPNs corresponding to the host in Table 15 .
9. Select the Host Cluster that we created in the previous step.

10. Click **Add**.

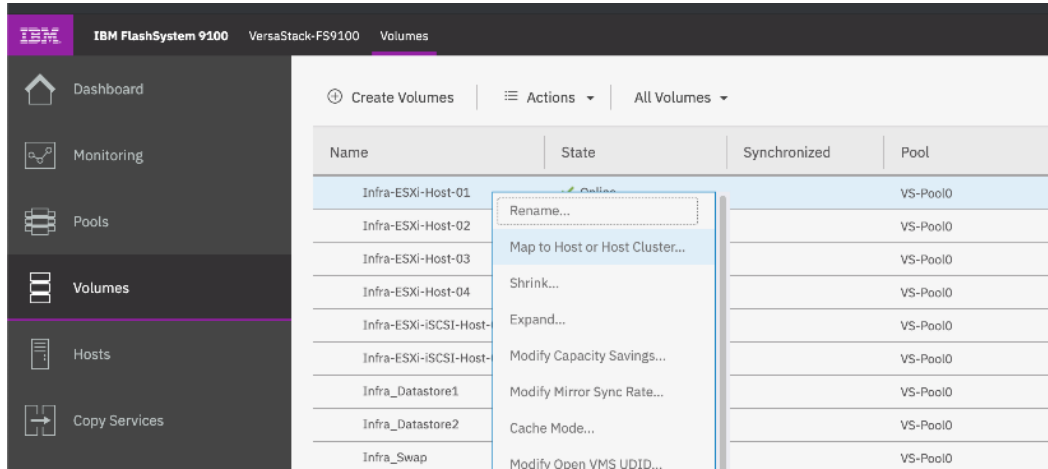
Map Volumes to Hosts and Host Cluster

To map volumes to hosts and host clusters, follow these steps:

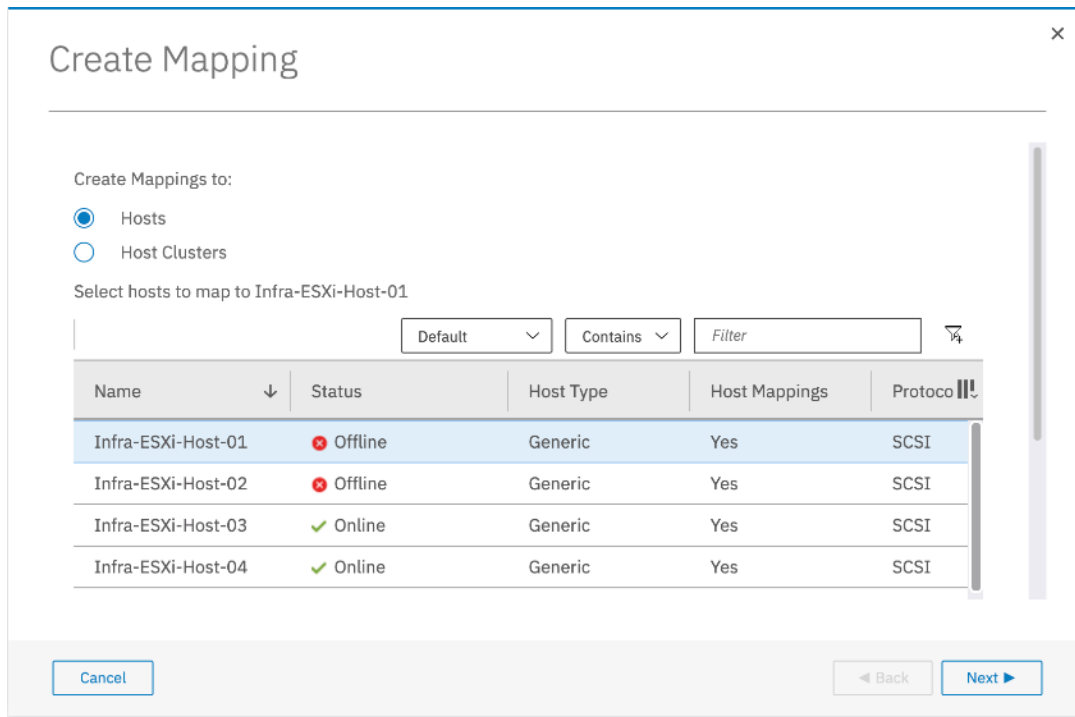
1. Click **Volumes**.



2. Right-click the Boot LUN for each ESXi host in turn and choose **Map to Host**.



3. Select the **Hosts** radio button and select the corresponding Host in the list and click **Next**.



4. Click **Map Volumes** and when the process is complete, click **Close**.

Map Volumes to Infra-ESXi-Host-01: Summary ×

The following volumes will be mapped to Infra-ESXi-Host-01:

Name	SCSI ID	Caching I/O Group ID	New Mapping
Infra-ESXi-Host-01	0	0	New

Cancel

◀ Back

Map Volumes

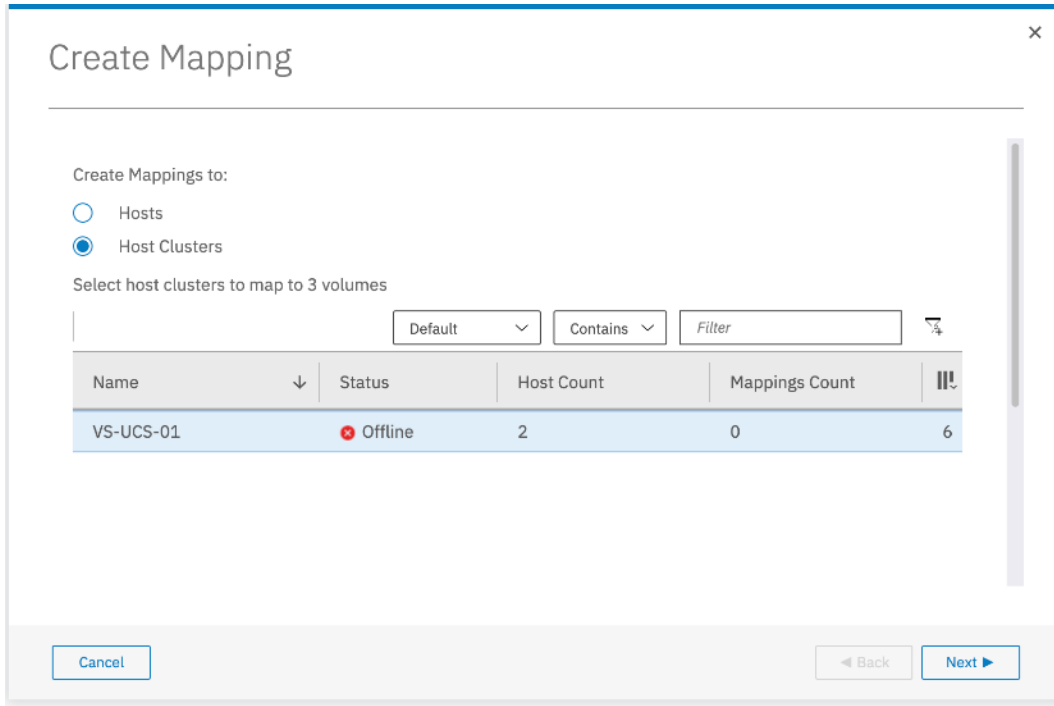
- Repeat above steps to map a Boot volume for each ESXi host in the cluster.
- When mapping shared volumes from Table 18 i.e. for shared VMFS datastores, right click on the volume in question (or select multiple volumes if mapping multiple LUNs) and select **Map to Host or Host Cluster**.

The screenshot shows a storage management interface with a table of volumes. The table has columns for Name, State, Synchronized, Pool, and Protocol Type. The 'Infra_Datastore1' row is selected, and a context menu is open over it. The menu options include: Rename..., Map to Host or Host Cluster..., Shrink..., Expand..., Modify Capacity Savings..., Modify Mirror Sync Rate..., Cache Mode..., Modify Open VMS UDID..., Remove Private Mappings..., View Mapped Hosts..., View Member MDisks..., Modify I/O Group..., and Cloud Volumes. The 'Map to Host or Host Cluster...' option is highlighted.

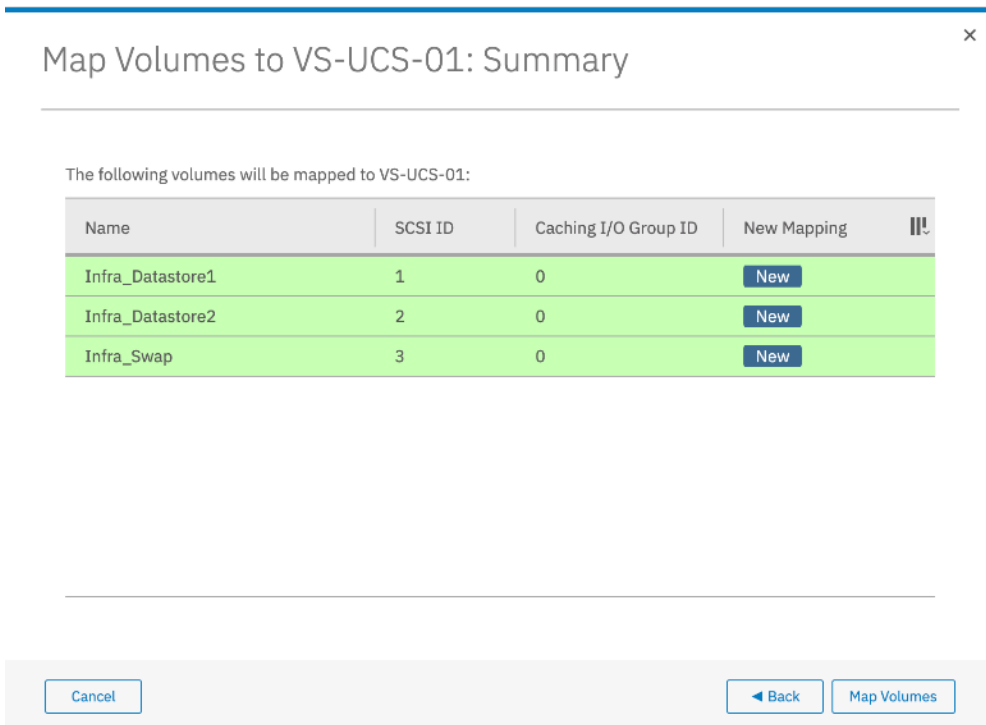
Name	State	Synchronized	Pool	Protocol Type
Infra_Datastore1	✓ Online			SCSI
Infra_Datastore2	✓ Online			SCSI
Infra_Swap	✓ Online			SCSI
RAW1	✓ Online			
RAW2	✓ Online			
RAW3	✓ Online			
RAW4	✓ Online			
RAW5	✓ Online			
RAW6	✓ Online			
RAW7	✓ Online			
RAW8	✓ Online			
RAW9	✓ Online			SCSI
RAW10	✓ Online			SCSI

Showing 42 volumes | Selecting 3 volumes (4.49 TiB)

- Select the **Host Clusters** radio button.



8. Review the summary and click **Map Volumes** to confirm.



9. Any Shared host cluster mappings will be automatically inherited by any future ESXi hosts which are defined as members of the host cluster.

VMware vSphere Setup for Cisco UCS Host Environment

VMware ESXi 6.7 U2

This section provides detailed instructions for installing VMware ESXi 6.7 U2 in the VersaStack UCS environment. After the procedures are completed, multiple ESXi hosts will be provisioned to host customer workloads.

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).

Log into Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log into the 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. Under HTML, click the **Launch UCS Manager link**.
3. When prompted, enter `admin` as the user name and enter the administrative password.
4. To log in to Cisco UCS Manager, click **Login**.
5. From the main menu, click the **Servers** tab.
6. Select Servers > Service Profiles > root > Infra-ESXi-Host-01.



For iSCSI setup, the name of the profile will be `Infra-ESXi-iSCSI-Host-0.1`

7. Right-click `Infra-ESXi-Host-01` and select KVM Console.
8. If prompted to accept an Unencrypted KVM session, accept as necessary.
9. Open KVM connection to all the hosts by right-clicking the Service Profile and launching the KVM console
10. Boot each server by selecting Boot Server and clicking **OK**. Click **OK** again.

Install ESXi on the UCS Servers

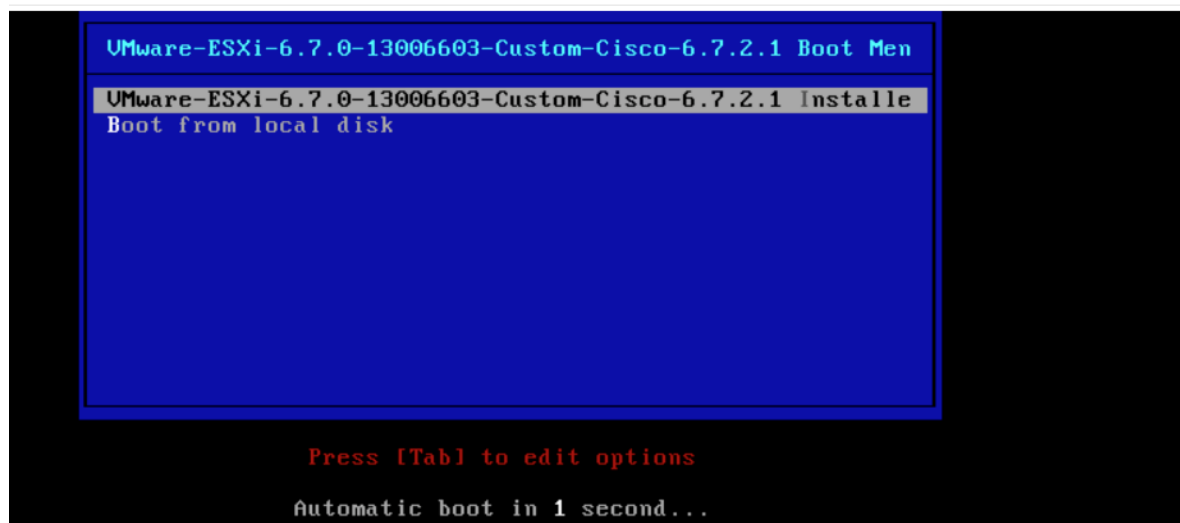
To install VMware ESXi to the boot LUN of the hosts, follow these steps on each host. The Cisco customer VMware ESXi image can be downloaded from:

<https://my.vmware.com/web/vmware/details?downloadGroup=OEM-ESXi67U2-CISCO&productId=742>

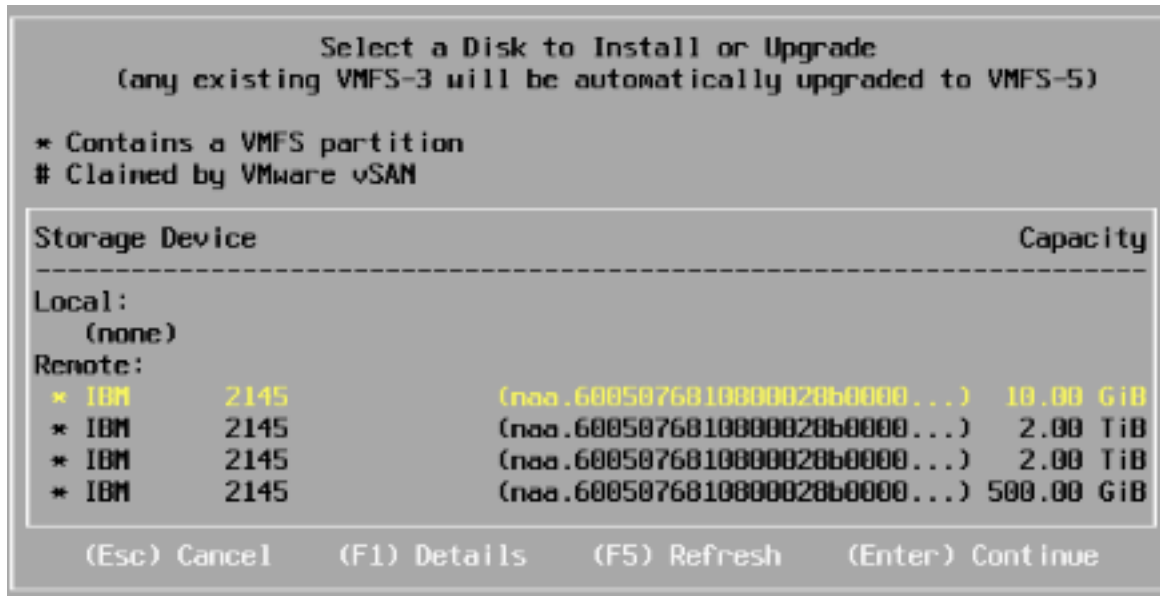


VMware ESXi will be installed on two Cisco UCS servers as part of the deployment covered in the following sections. The number of ESXi servers can vary based on the customer specific deployment.

1. In the KVM windows, click Virtual Media in the upper right of the screen.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select **Map CD/DVD**.
5. Browse to the ESXi installer ISO image file and click **Open**.
6. Click Map Device.
7. Click the **KVM** tab to monitor the server boot.
8. Reset the server by clicking Reset button. Click **OK**.
9. Select Power Cycle on the next window and click **OK** and **OK** again.
10. On reboot, the machine detects the presence of the boot LUNs (sample output below).
11. From the ESXi Boot Menu, select the ESXi installer.



12. After the installer has finished loading, press **Enter** to continue with the installation.
13. Read and accept the end-user license agreement (EULA). Press **F11** to accept and continue.
14. Select the LUN that was previously set up and discovered as the installation disk for ESXi and press **Enter** to continue with the installation.



15. Select the appropriate keyboard layout and press **Enter**.
16. Enter and confirm the root password and press **Enter**.
17. The installer issues a warning that the selected disk will be repartitioned. Press **F11** to continue with the installation.
18. After the installation is complete, press **Enter** to reboot the server.
19. Repeat the ESXi installation process for all the Service Profiles.



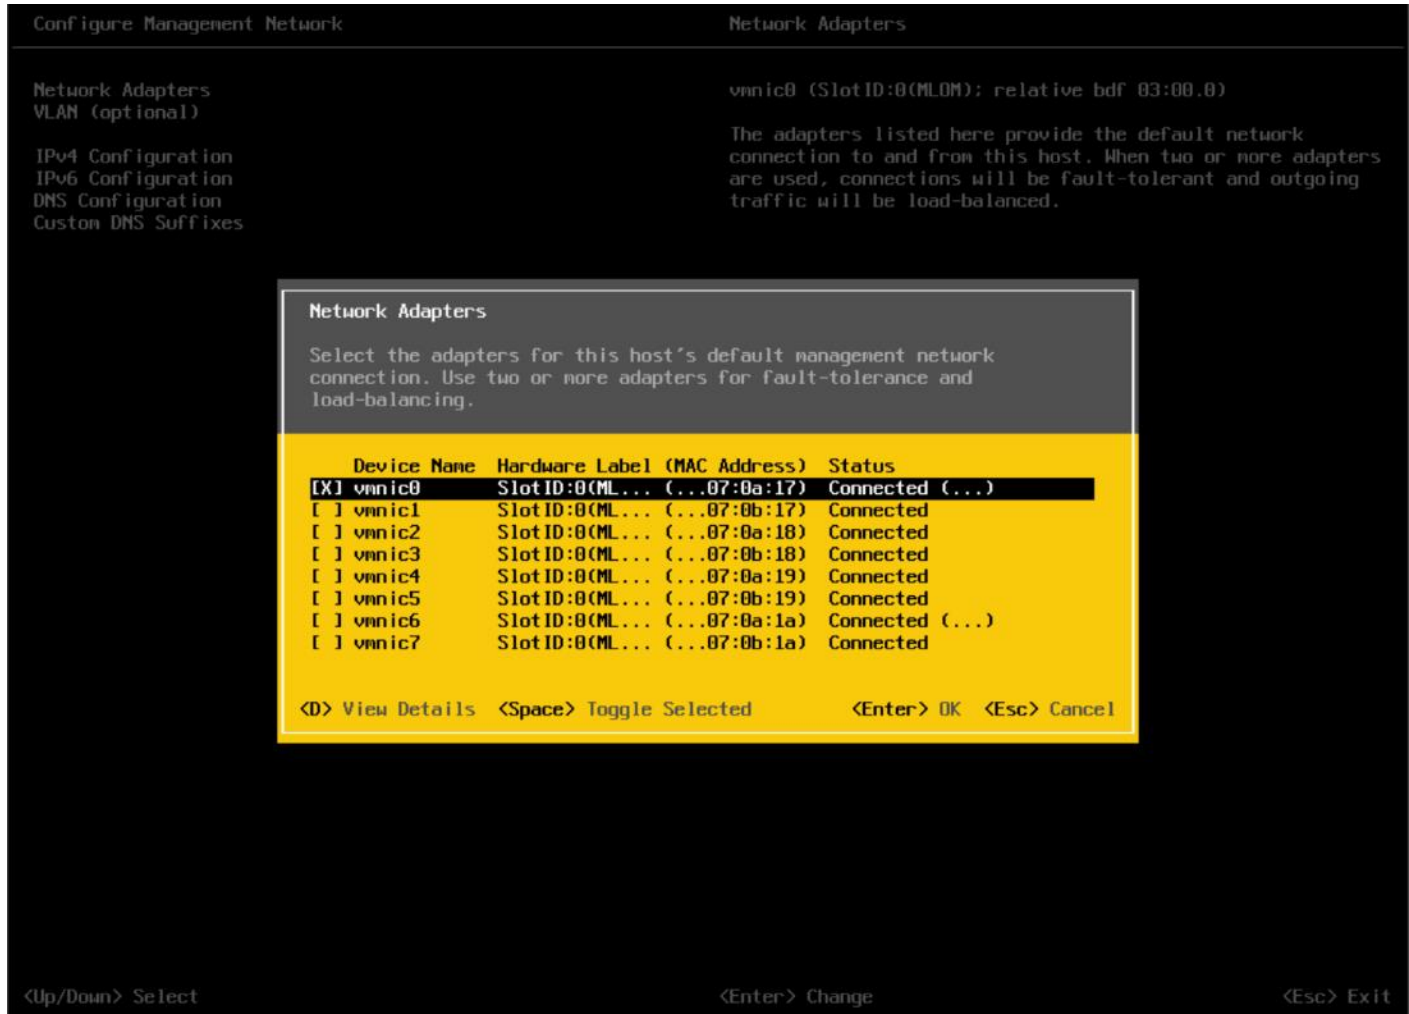
In this deployment, we used two UCS server blades for the VMware vSphere deployment.

Set Up Management Networking for ESXi Hosts

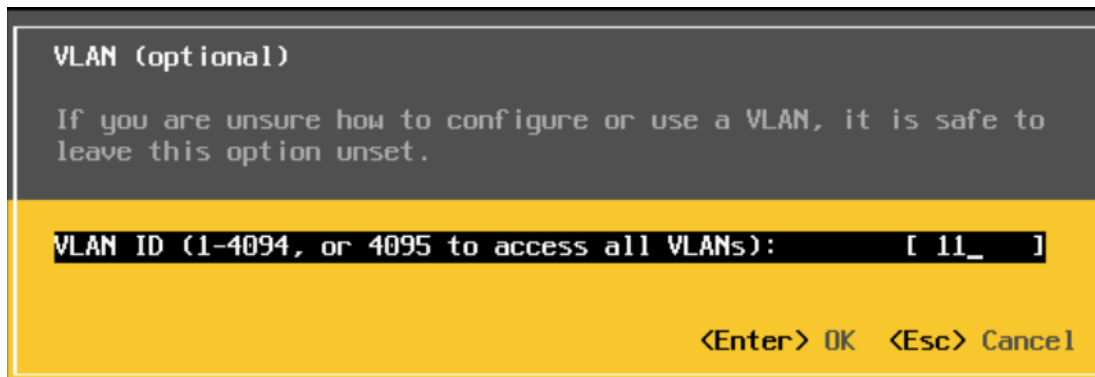
Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, follow these steps on each ESXi host.

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

1. After the server has finished post-installation rebooting, press **F2** to customize the system.
2. Log in as root, enter the password chosen during the initial setup, and press **Enter** to log in.
3. Select the Configure Management Network option and press **Enter**.
4. Select Network Adapters
5. Select `vmnic0` (if it is not already selected) by pressing the Space Bar.

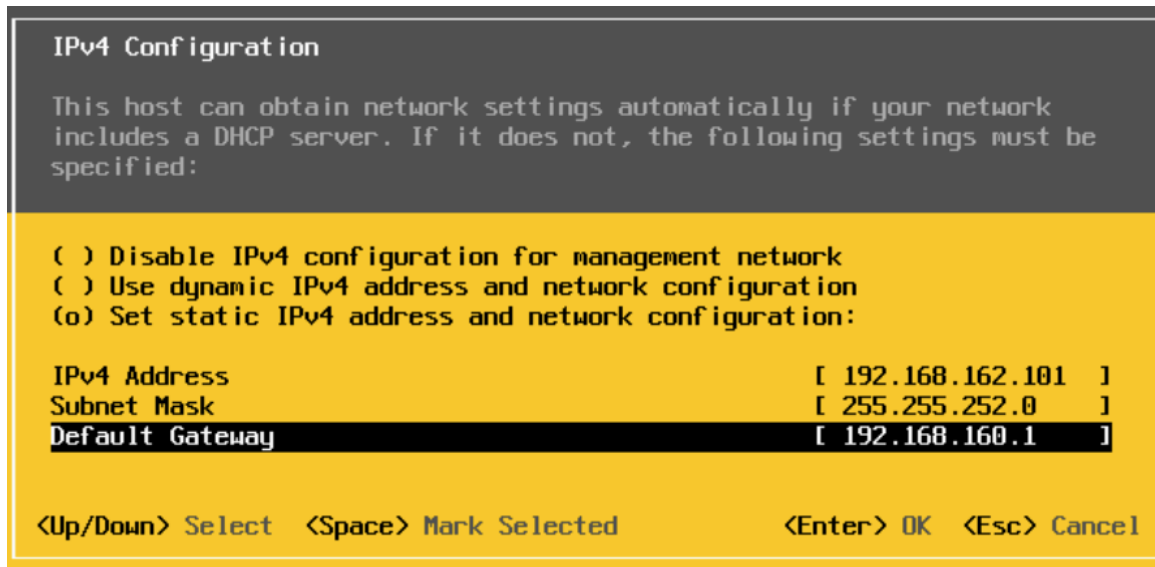


6. Press **Enter** to save and exit the Network Adapters window.
7. Select the VLAN (Optional) and press Enter.
8. Enter the <IB Mgmt VLAN> (11) and press Enter.



9. Select IPv4 Configuration and press **Enter**.

10. Select the Set Static IP Address and Network Configuration option by using the Space Bar.
11. Enter the IP address for managing the ESXi host.
12. Enter the subnet mask for the management network of the ESXi host.
13. Enter the default gateway for the ESXi host.

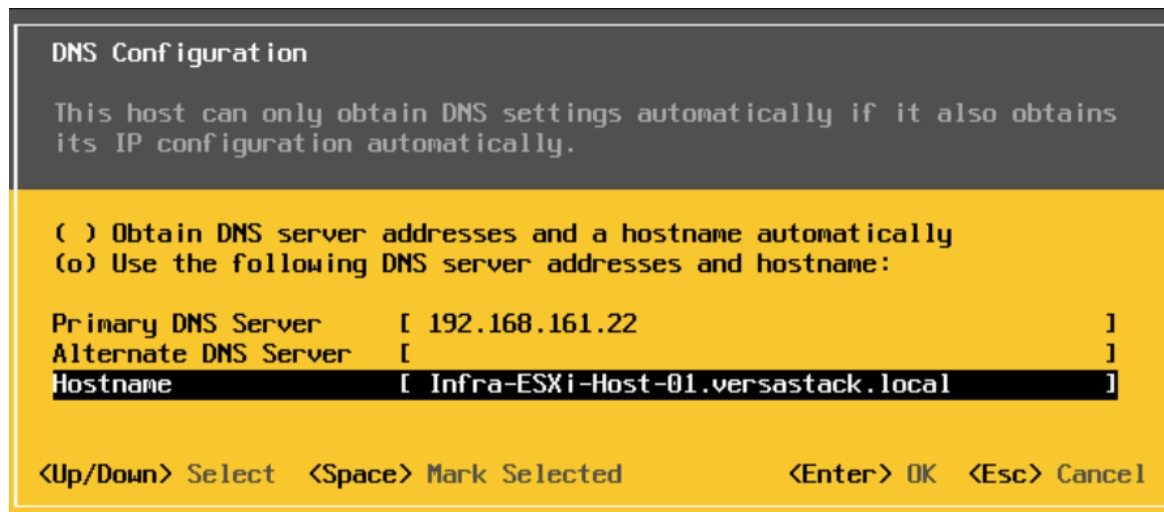


14. Press **Enter** to accept the changes to the IP configuration.
15. Select the IPv6 Configuration option and press **Enter**.
16. Using the Space Bar, select Disable IPv6 (restart required) and press **Enter**.
17. Select the DNS Configuration option and press **Enter**.



Because the IP address is assigned manually, the DNS information must also be entered manually.

18. Enter the IP address of the primary DNS server.
19. **Optional:** Enter the IP address of the secondary DNS server.
20. Enter the fully qualified domain name (FQDN) for the ESXi host.



21. Press **Enter** to accept the changes to the DNS configuration.
22. Press **Esc** to exit the Configure Management Network submenu.
23. Press **Y** to confirm the changes and reboot the host.
24. Repeat this procedure for all the ESXi hosts in the setup.

VMware vSphere Configuration

The vSphere configuration covered in this section is common to all the ESXi servers. In the procedure below, three shared datastores, two for hosting the VMs and another to host the VM swap files, will be mounted to all the ESXi servers. Customers can adjust the number and size of the shared datastores based on their particular deployments.

Log into VMware ESXi Hosts Using VMware vSphere Client

To log into the ESXi host using the VMware Host Client, follow these steps:

1. Open a web browser on the management workstation and navigate to the management IP address of the host.
2. Click Open the VMware Host Client.
3. Enter `root` for the user name.
4. Enter the root password configured during the installation process.
5. Click **Login** to connect.
6. Decide whether to join the VMware Customer Experience Improvement Program and click OK.
7. Repeat this process to log into all the ESXi hosts.

Install VMware Drivers for the Cisco Virtual Interface Card (VIC)

For the most recent versions, please refer to [Cisco UCS HW and SW Availability Interoperability Matrix](#). If a more recent driver is made available that is appropriate for VMware vSphere 6.7 U2, download and install the latest drivers.

To install VMware VIC Drivers on the ESXi hosts using esxcli, follow these steps:

1. Download and extract the following VIC Drivers to the Management workstation:

NFNIC Driver version 4.0.0.40:

<https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI67-CISCO-NFNIC-40040&productId=742>

NENIC Driver version 1.0.29.0:

<https://my.vmware.com/web/vmware/details?downloadGroup=DT-ESXI67-CISCO-NENIC-10290&productId=742>

To install VIC Drivers on ALL the ESXi hosts, follow these steps:

1. From each Host Client, select **Storage**.
2. Right-click `datastore1` and select **Browse**.
3. In the Datastore browser, click **Upload**.
4. Navigate to the saved location for the downloaded VIC drivers and select `VMW-ESX-6.7.0-nenic-1.0.29.0-offline_bundle-12897497.zip`.
5. In the Datastore browser, click **Upload**.
6. Navigate to the saved location for the downloaded VIC drivers and select `VMW-ESX-6.7.0-nfnic-4.0.0.40-offline_bundle-14303978.zip`.
7. Click **Open** to upload the file to `datastore1`.
8. Make sure the file has been uploaded to both ESXi hosts.
9. Place each host into Maintenance mode if it isn't already.
10. Connect to each ESXi host through ssh from a shell connection or putty terminal.
11. Login as root with the root password.
12. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-nenic-1.0.29.0-offline_bundle-12897497.zip

esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-nfnic-4.0.0.40-offline_bundle-14303978.zip

reboot
```

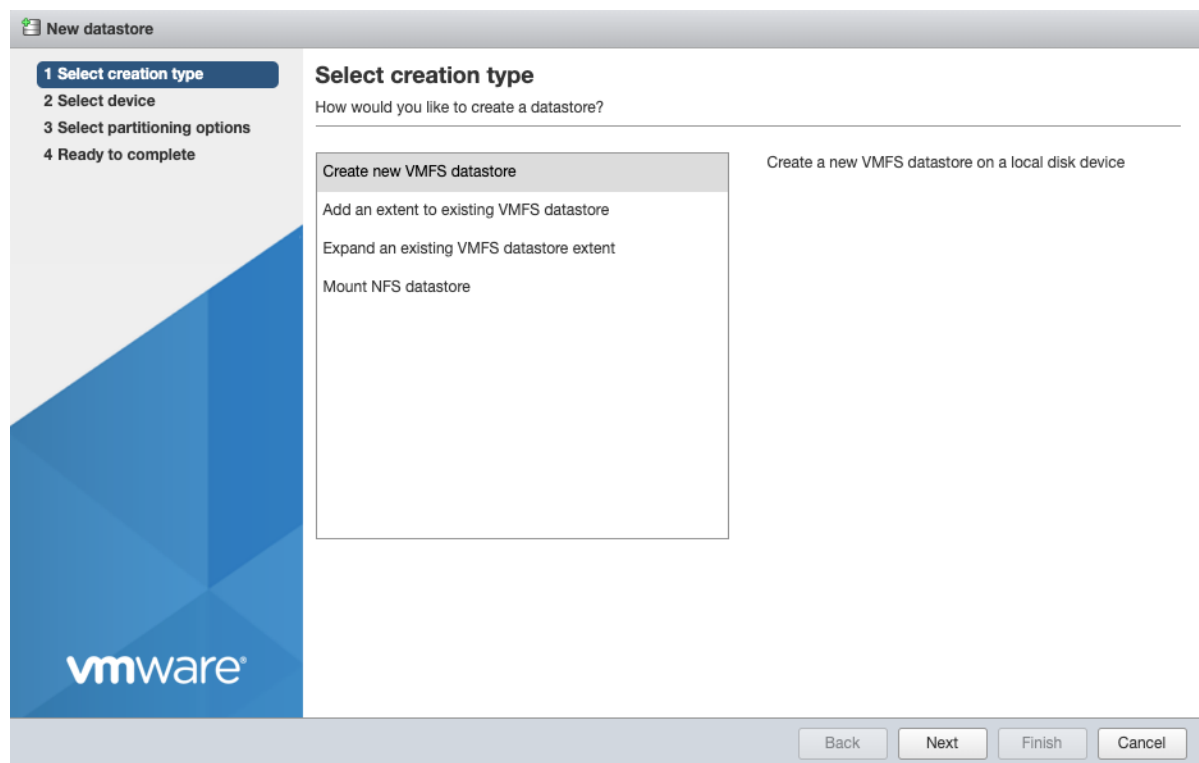
13. Log into the Host Client on each host once reboot is complete and exit Maintenance Mode.

Mount Required Datastores

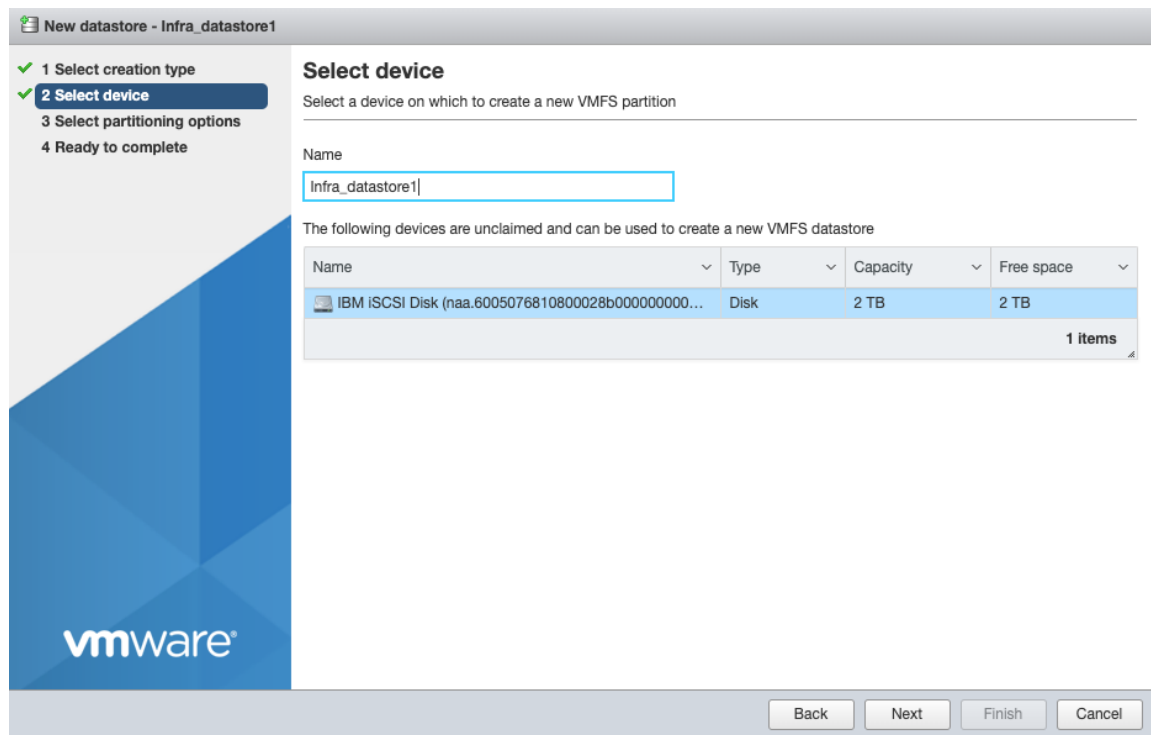
To mount the required datastores, follow these steps on each ESXi host:

1. From the Host Client, select **Storage**.

2. In the center pane, select the **Datastores** tab.
3. In the center pane, select New Datastore to add a new datastore.
4. In the New datastore popup, select Create new VMFS datastore.
5. Click **Next**.



6. Enter `Infra_datastore1` as the datastore name.
7. Verifying by using the size of the datastore LUN, select the LUN configured for VM hosting and click **Next**.



8. Accept default VMFS setting and Use full disk option to retain maximum available space.
9. Click **Next**
10. Verify the details and Click **Finish**.
11. In the center pane, select the Datastores tab.
12. In the center pane, select New Datastore to add a new datastore.
13. In the New datastore popup, select Create new VMFS datastore.
14. Click **Next**.
15. Enter `Infra_datastore2` as the datastore name.
16. Verifying by using the size of the datastore LUN, select the LUN configured for VM hosting and click **Next**.
17. Accept default VMFS setting and Use full disk option to retain maximum available space.
18. Click **Next**
19. Verify the details and Click **Finish**.
20. In the center pane, select the Datastores tab.
21. In the center pane, select New Datastore to add a new datastore.
22. In the New datastore popup, select Create new VMFS datastore.

23. Click **Next**.
24. Enter `Infra_swap` as the datastore name.
25. Verifying by using the size of the datastore LUN, select the LUN configured for VM hosting and click **Next**.
26. Accept default VMFS setting and Use full disk option to retain maximum available space.
27. Click **Next**
28. Verify the details and Click **Finish**.
29. The storage configuration should look similar to figure shown below.
30. Repeat these steps on all the ESXi hosts.

The screenshot shows the 'Storage' configuration page in vSphere. At the top, there are tabs for 'Datastores', 'Adapters', 'Devices', and 'Persistent Memory'. Below the tabs, there are several action buttons: 'New datastore', 'Increase capacity', 'Register a VM', 'Datastore browser', 'Refresh', and 'Actions'. The main area contains a table with the following data:

Name	Drive Type	Capacity	Provisioned	Free	Type
datastore1	Non-SSD	2.5 GB	1.41 GB	1.09 GB	VMFS6
Infra_datastore1	Non-SSD	2 TB	815.34 GB	1.2 TB	VMFS6
Infra_datastore2	Non-SSD	2 TB	1.01 TB	1,015.83 GB	VMFS6
Infra_swap	Non-SSD	496.75 GB	76.13 GB	423.62 GB	VMFS6

Configure NTP on ESXi Hosts

To configure NTP on the ESXi hosts, follow these steps on each host:

1. From the Host Client, select **Manage**.
2. In the center pane, select **Time & date**.
3. Click Edit settings.
4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the drop-down to select Start and stop with host.
6. Enter the NTP addresses in the NTP servers box separated by a comma, Nexus switch addresses can be entered if NTP service is configured on the switches.

Edit time configuration

Specify how the date and time of this host should be set.

Manually configure the date and time on this host

09/14/2019 10:46 AM

Use Network Time Protocol (enable NTP client)

NTP service startup policy	Start and stop manually ▼
NTP servers	<div style="border: 1px solid #ccc; padding: 2px;">192.168.162.128, 192.168.162.129</div> <p style="font-size: 0.8em; margin-top: 5px;">Separate servers with commas, e.g. 10.31.21.2, fe00::2800</p>

7. Click **Save** to save the configuration changes.
8. Select Actions > NTP service > Start.
9. Verify that NTP service is now running and the clock is now set to approximately the correct time.



The NTP server time may vary slightly from the host time.

Move VM Swap File Location

To move the VM swap file location, follow these steps on each ESXi host:

1. From the Host Client, select **Manage**.
2. In the center pane, select **Swap**.
3. Click Edit settings.
4. Use the drop-down list to select `Infra_swap`. Leave all other settings unchanged.

Property	Value
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

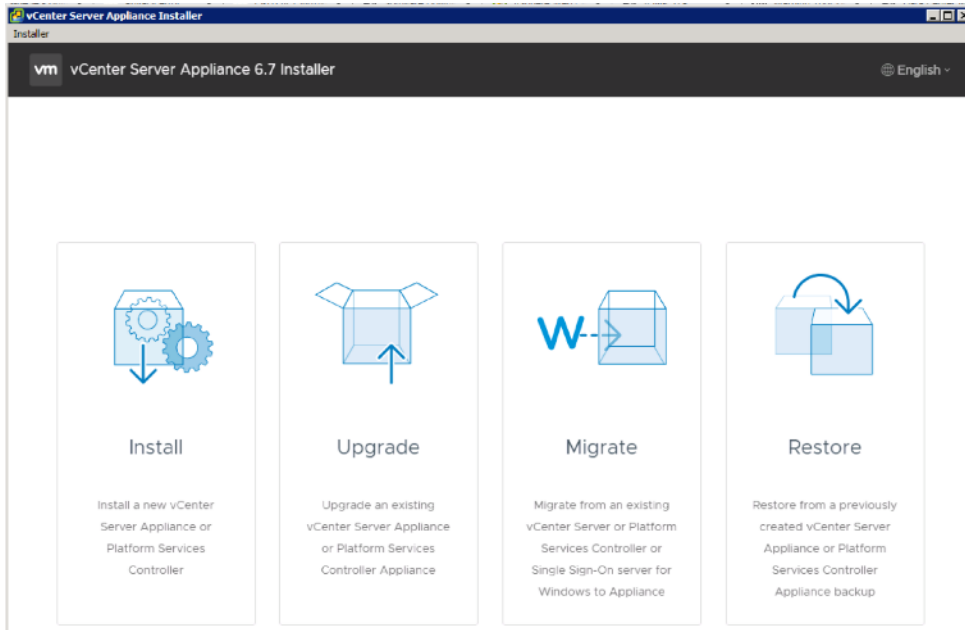
Save Cancel

5. Click **Save** to save the configuration changes.

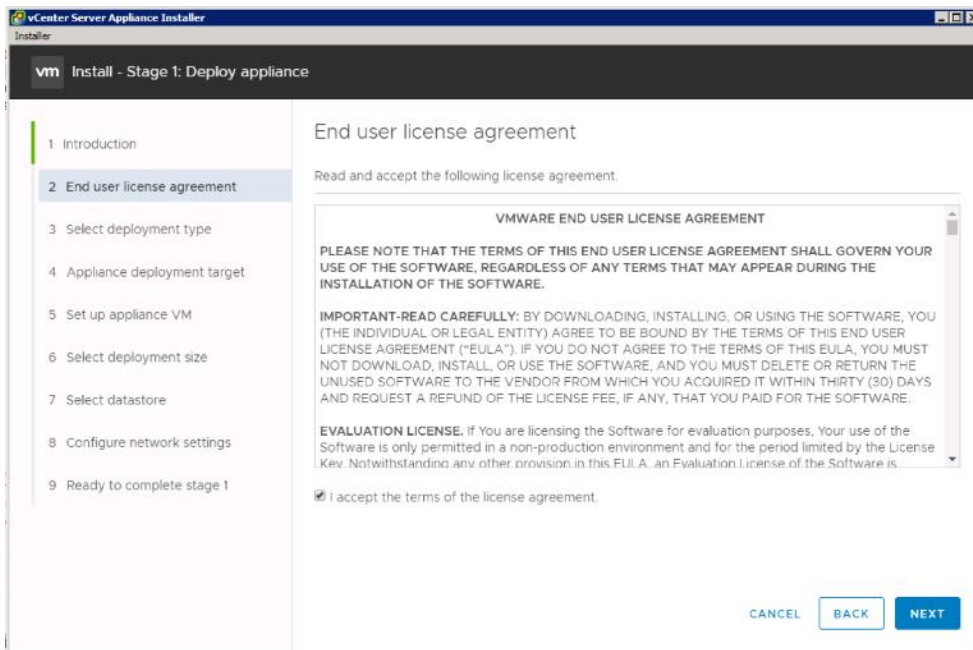
Deploy VMware vCenter Appliance 6.7 (Optional)

The VCSA deployment consists of 2 stages: install and configuration. To build the VMware vCenter virtual machine, follow these steps:

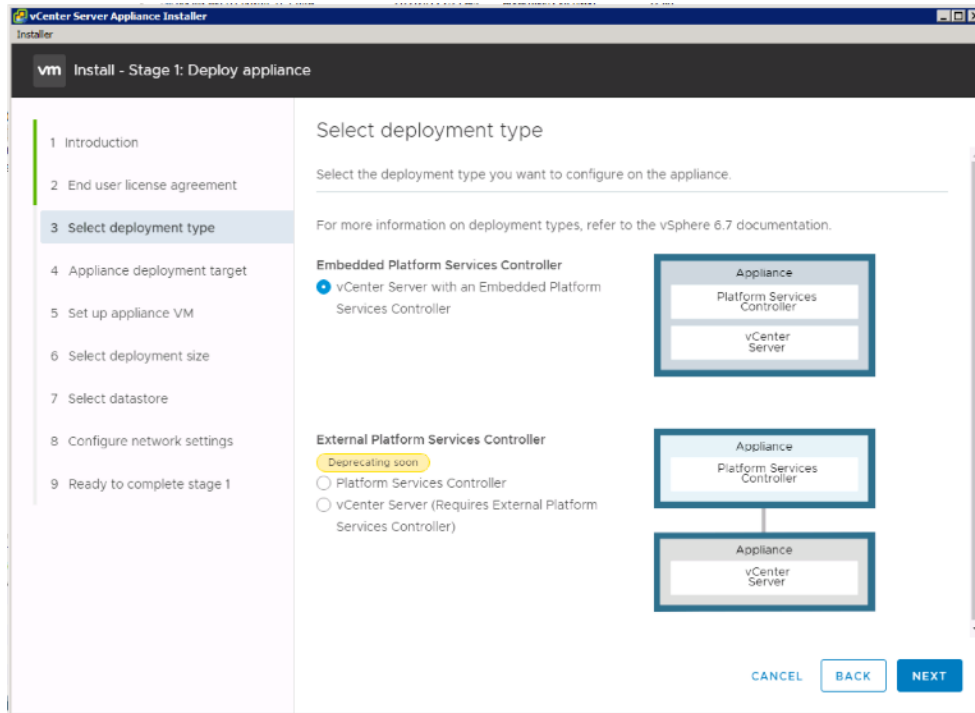
1. Download the VCSA ISO from VMware at <https://my.vmware.com/group/vmware/details?downloadGroup=VC67U2C&productId=742&rPId=35624>
2. Using ISO mounting software, mount the ISO image as a disk on the management workstation. (For example, with the Mount command in Windows Server 2012).
3. In the mounted disk directory, navigate to the vcsa-ui-installer > win32 directory and double-click installer.exe. The vCenter Server Appliance Installer wizard appears.



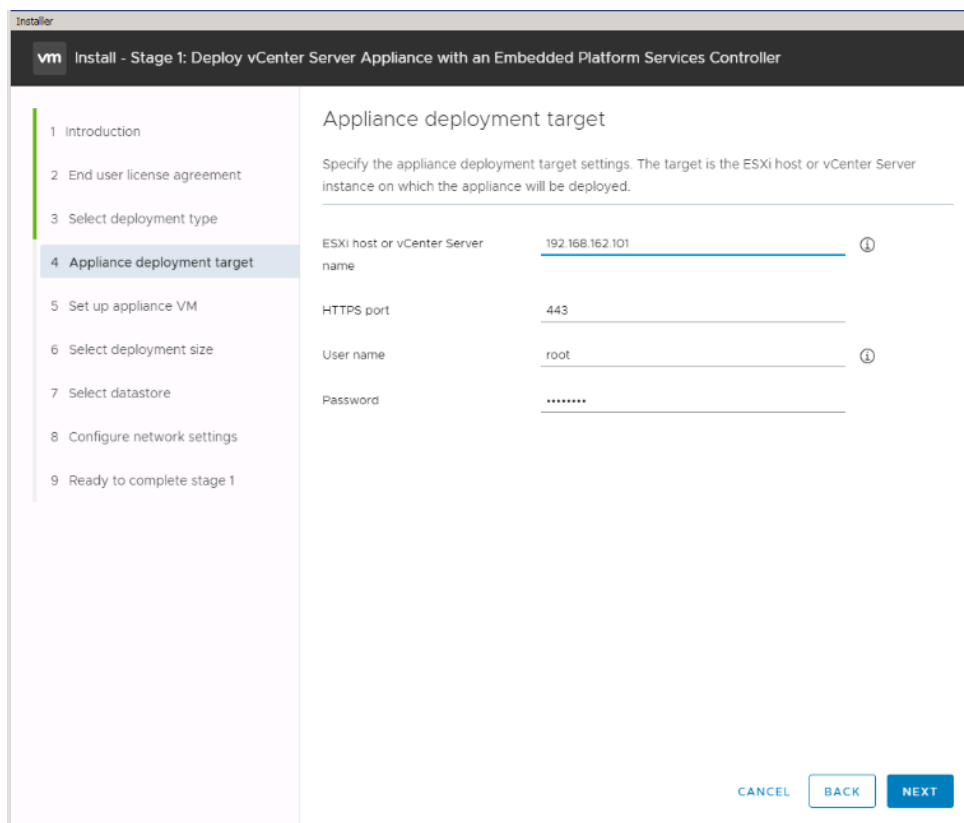
4. Click **Install** to start the vCenter Server Appliance deployment wizard.
5. Click **Next** in the Introduction section.
6. Read and accept the license agreement and click **Next**.



7. In the "Select deployment type" section, select vCenter Server with an Embedded Platform Services Controller and click **Next**.

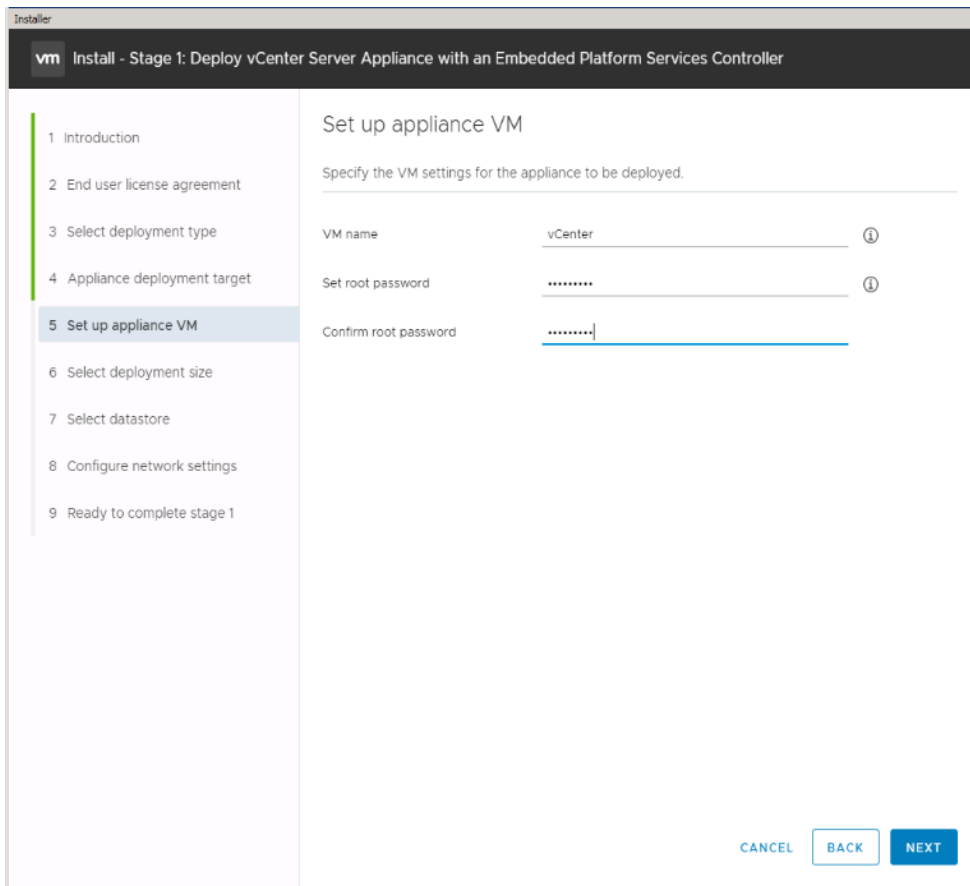


8. In the “Appliance deployment target”, enter the ESXi host name or IP address for the first configured VSI host, User name and Password. Click **Next**.

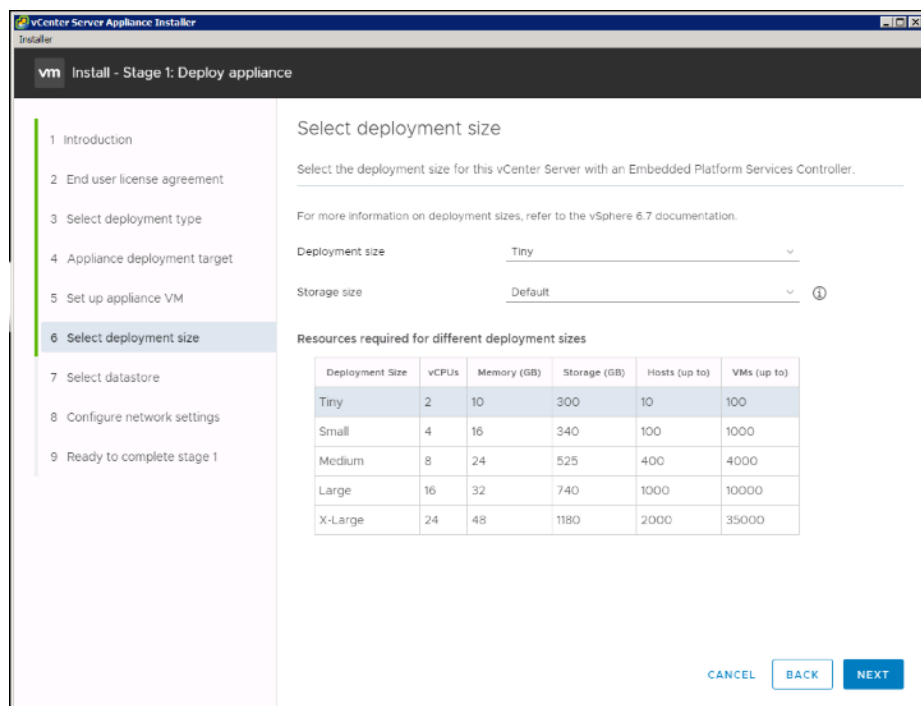


9. Click **Yes** to accept the certificate.

10. Enter the Appliance name and password details in the "Set up appliance VM" section. Click **Next**.

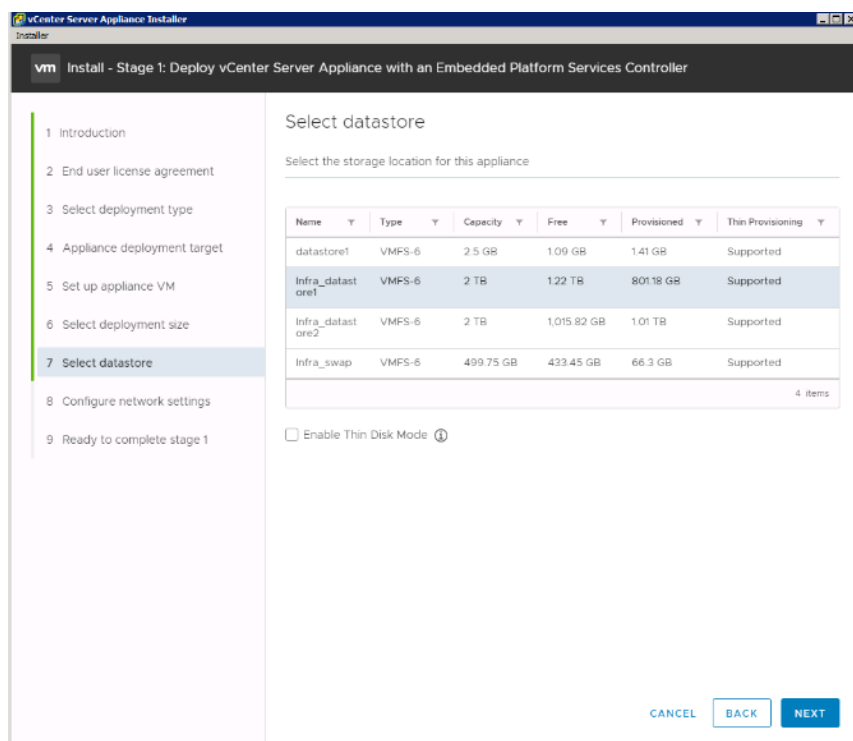


11. In the "Select deployment size" section, Select the deployment size and Storage size. For example, "Tiny" Deployment size was selected in this CVD.



12. Click **Next**.

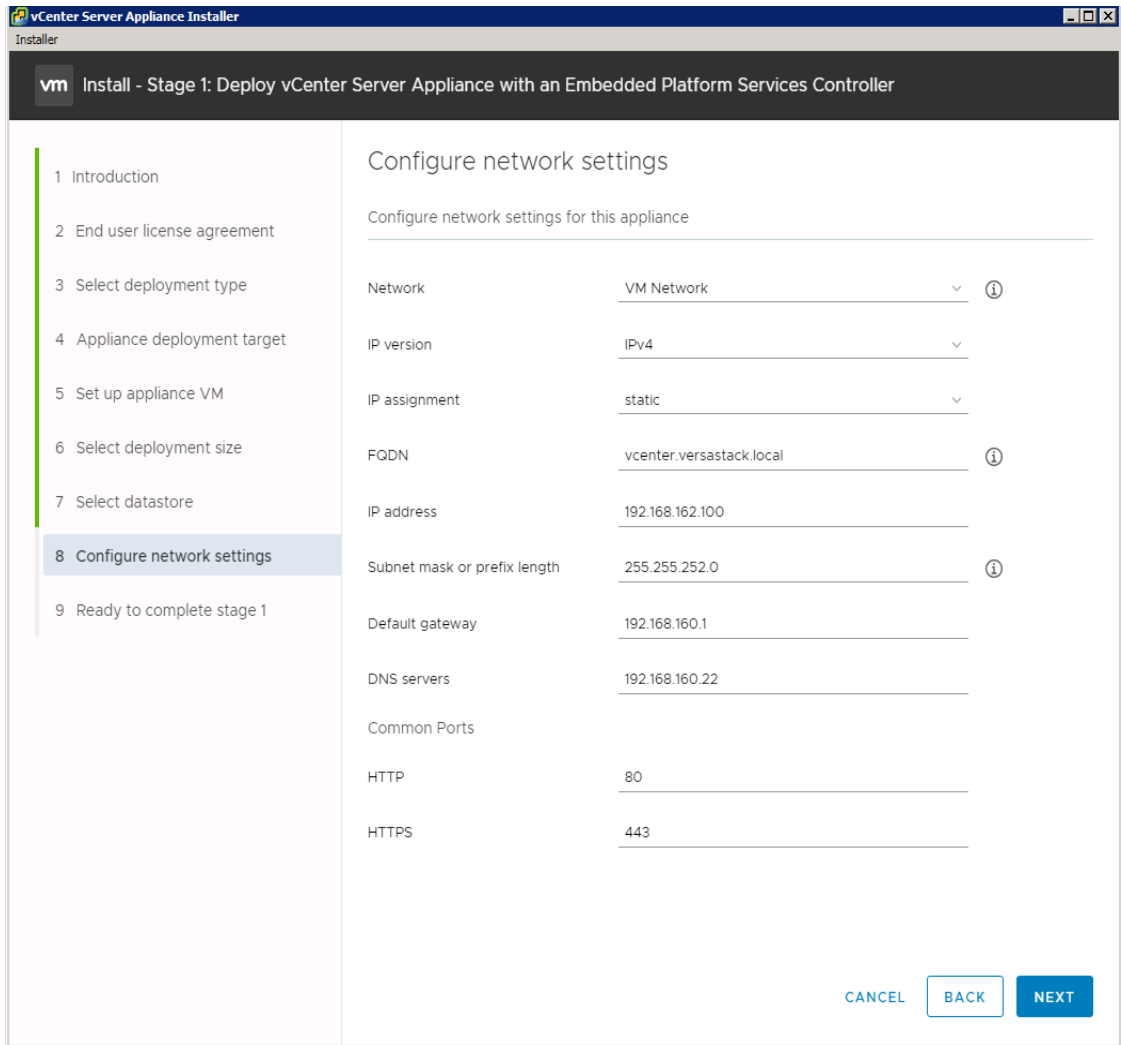
13. Select preferred datastore e.g. the "Infra_datastore1" was created previously.



14. Click **Next**.

15. In the "Network Settings" section, configure the following settings:

- a. Choose a Network: VM Network
- b. IP version: IPV4
- c. IP assignment: static
- d. System name: <vcenter-fqdn> (optional)
- e. IP address: <vcenter-ip>
- f. Subnet mask or prefix length: <vcenter-subnet-mask>
- g. Default gateway: <vcenter-gateway>
- h. DNS Servers: <dns-server>



- 16. Click **Next**.
- 17. Review all values and click **Finish** to complete the installation.
- 18. The vCenter appliance installation will take a few minutes to complete.
- 19. Click **Continue** to proceed with stage 2 configuration.

20. Click **Next**.

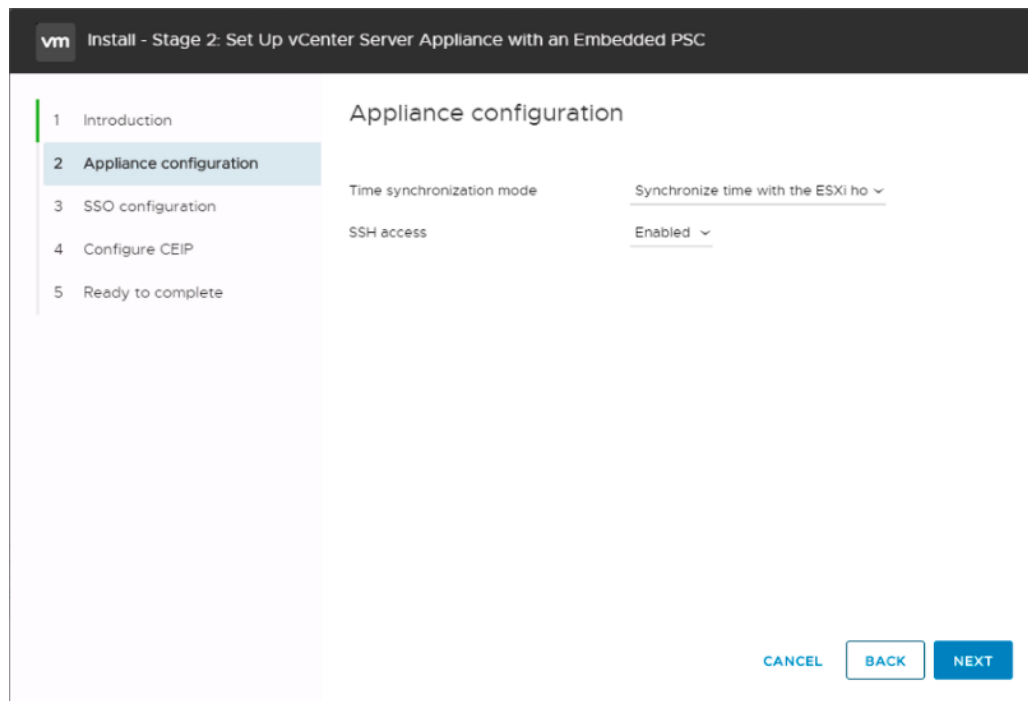
21. In the Appliance Configuration, configure the below settings:

- a. Time Synchronization Mode: Synchronize time with the ESXi host.



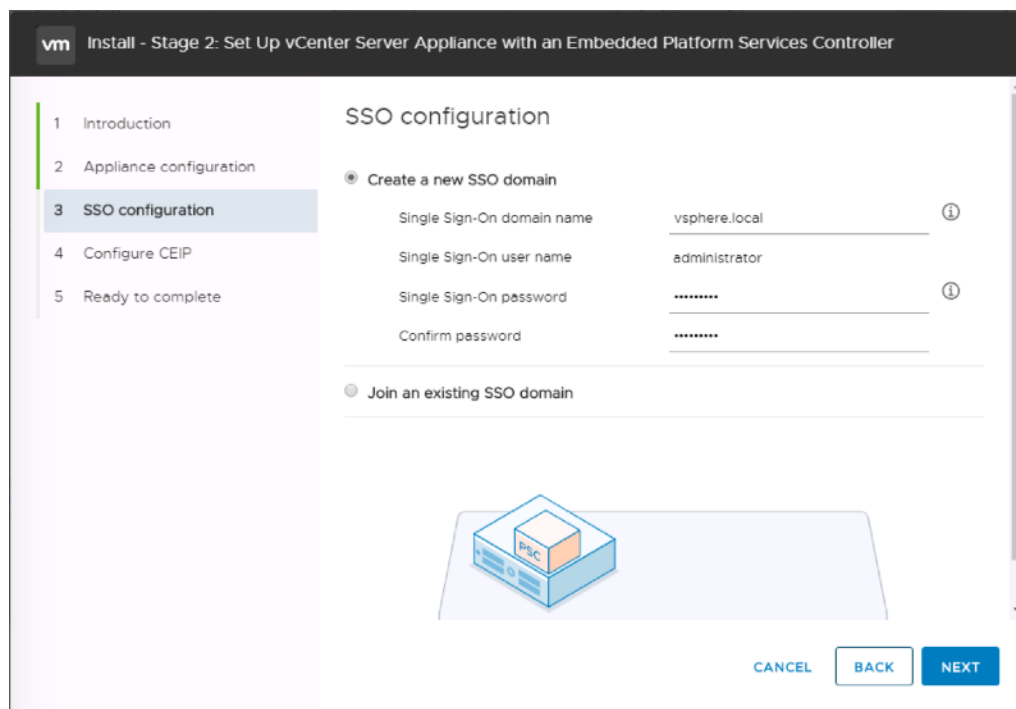
Since the ESXi host has been configured to synchronize the time with an NTP server, vCenter time can be synced to ESXi host. Customer can choose a different time synchronization setting.

- b. SSH access: Enabled.



22. Click Next.

23. Complete the SSO configuration as shown below.



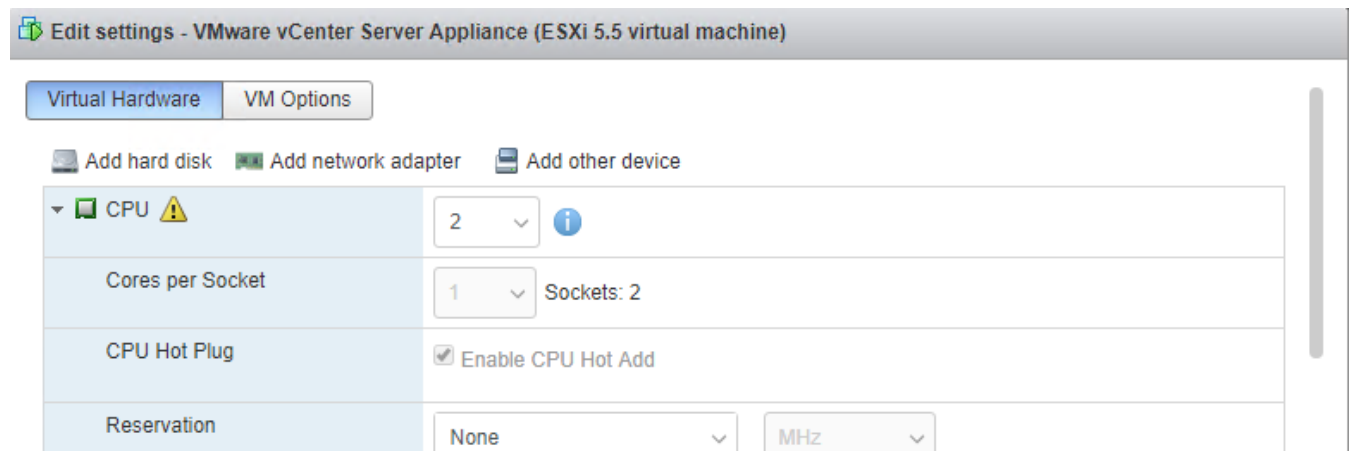
24. Click **Next**.
25. If preferred, select Join the VMware's Customer Experience Improvement Program (CEIP).
26. Click **Next**.
27. Review the configuration and click Finish.
28. Click **OK**.
29. Make note of the access URL shown in the completion screen.
30. Click **Close**.

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 UCS server CPU hardware configuration. Cisco UCS B200 and C220 servers are 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 `Infra-esxi-host-01` management IP address.
2. Click **Open** the VMware Host Client.
3. Enter `root` for the user name.
4. Enter the root password.
5. Click **Login** to connect.

6. In the center pane, right-click the vCenter VM and select **Edit settings**.
7. In the Edit settings window, expand CPU and check the value of Sockets is not greater than 2.



8. If the number of Sockets is greater than 2, it will need to be adjusted. Click **Cancel**.
9. If the number of Sockets needs to be adjusted:
10. Right-click the vCenter VM and select Guest OS > Shut down. Click **Yes** on the confirmation.
11. Once vCenter is shut down, right-click the vCenter VM and select Edit settings.
12. In the Edit settings window, expand CPU and change the Cores per Socket value to make the Sockets value 2.
13. Click **Save**.
14. Right-click the vCenter VM and select Power > Power on. Wait approximately 10 minutes for vCenter to come up.

Set Up VMware vCenter Server

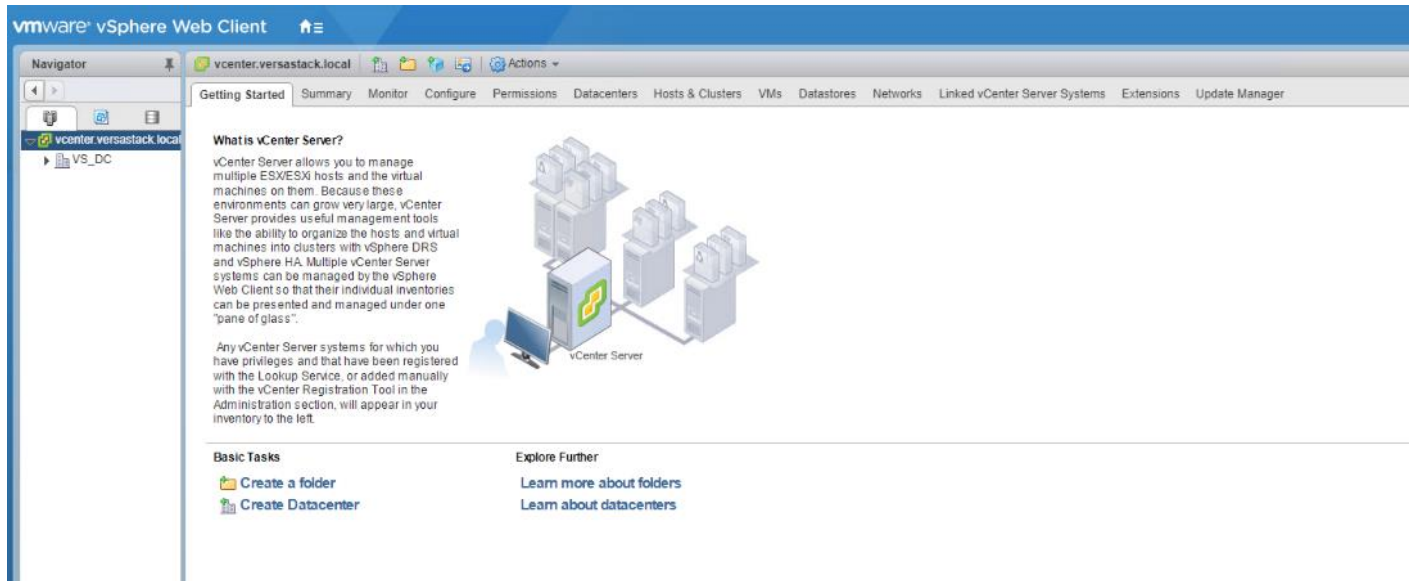
To setup the VMware vCenter Server, follow these steps:

1. Using a web browser, navigate to <https://<vcenter-ip>/vsphere-client>.



The VMware vSphere HTML5 Client is fully featured in vSphere 6.7U2 and can be used for setting up the vCenter if preferred by the customer. However, the Web Client is used in this document.

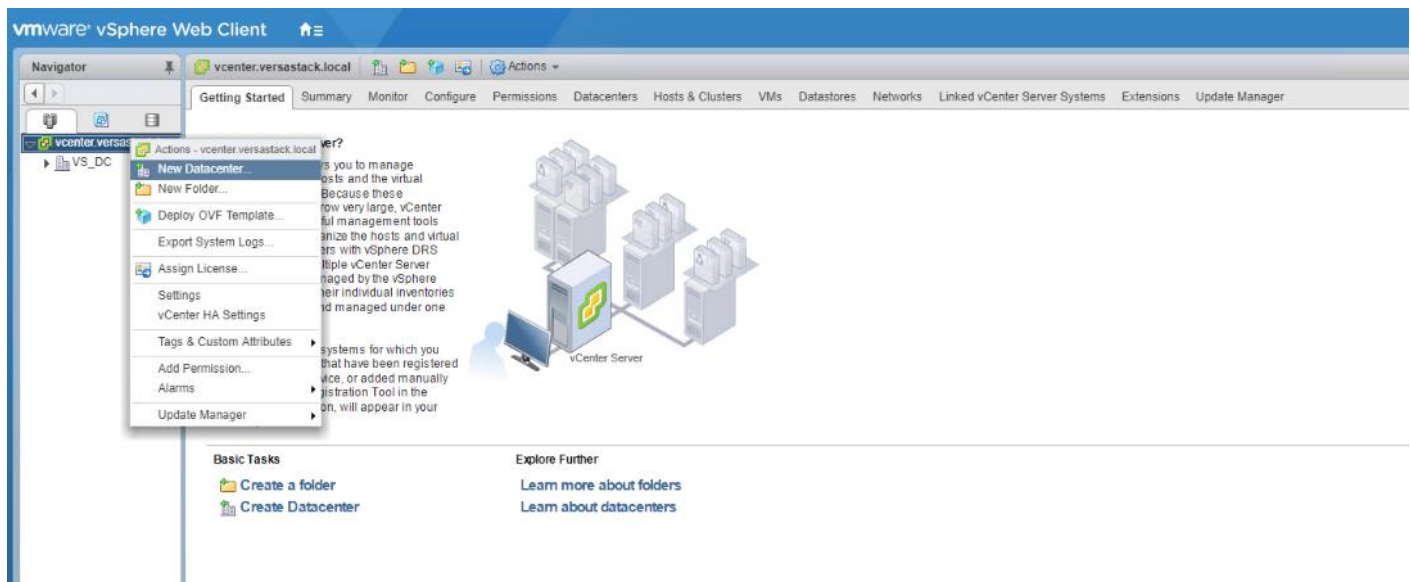
2. If the link is available, click Download Enhanced Authentication Plugin. Install the same by double-clicking the downloaded file.
3. Log in using the Single Sign-On username (administrator@vsphere.local) and password created during the vCenter installation.



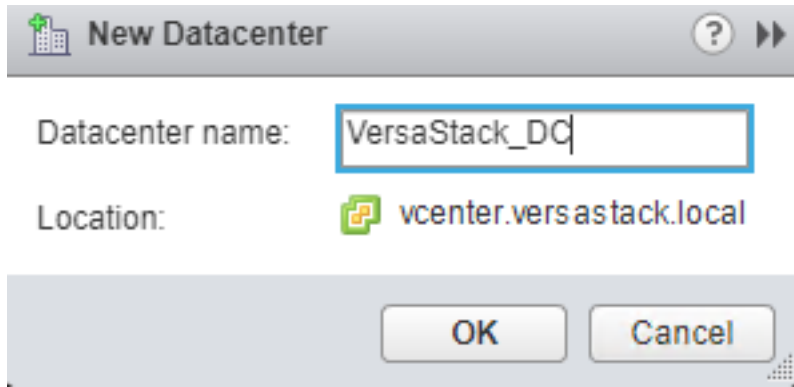
Setup Data Center, Cluster, DRS and HA for ESXi Nodes

If a new data center is needed for the VersaStack, follow these steps on the vCenter:

1. Connect to the vSphere Web Client and click **Hosts and Clusters** from the left side Navigator window or the **Hosts and Clusters** icon from the Home center window
2. From Hosts and Clusters:
3. Right-click the **vCenter** icon and from the drop-down list select **New Datacenter**.



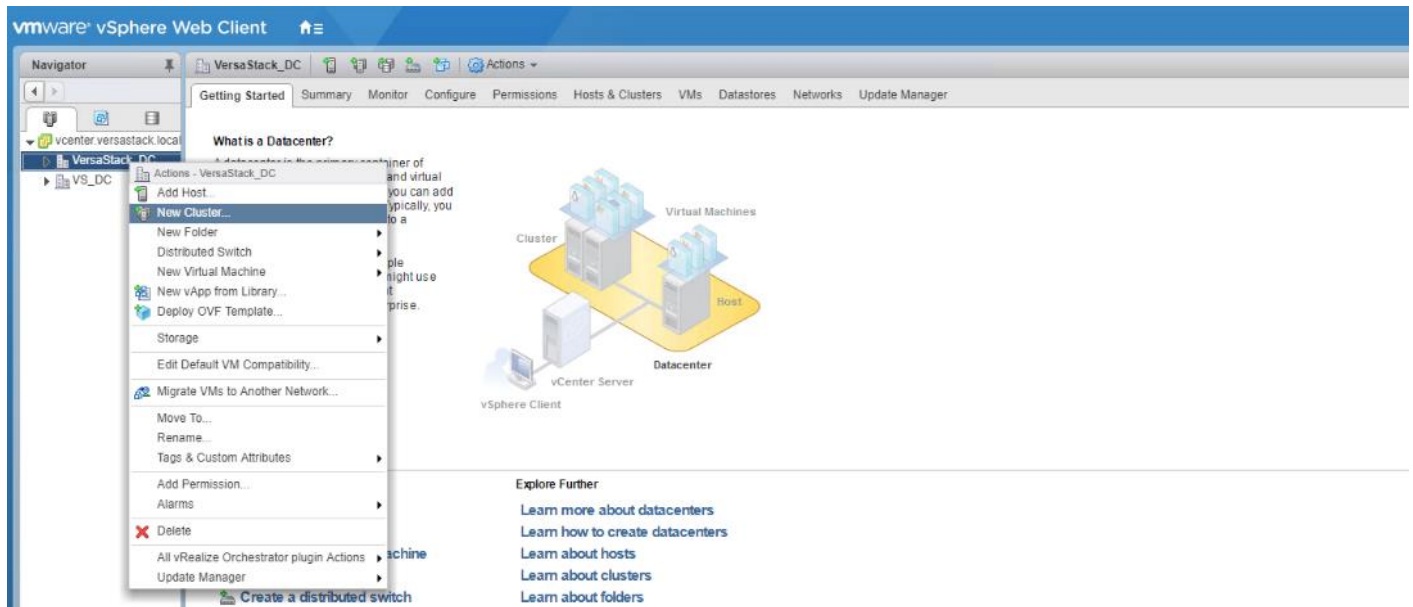
4. From the New Datacenter pop-up dialogue enter in a Datacenter name and click **OK**.



Add the VMware ESXi Hosts Using the VMware vSphere Web Client

To add the VMware ESXi Hosts using the VMware vSphere Web Client, follow these steps:

1. From the Hosts and Clusters tab, right-click the new or existing Datacenter within the Navigation window, and from the drop-down list select **New Cluster**.



2. Enter a name for the new cluster, select the DRS and HA checkmark boxes, leaving all other options with defaults.

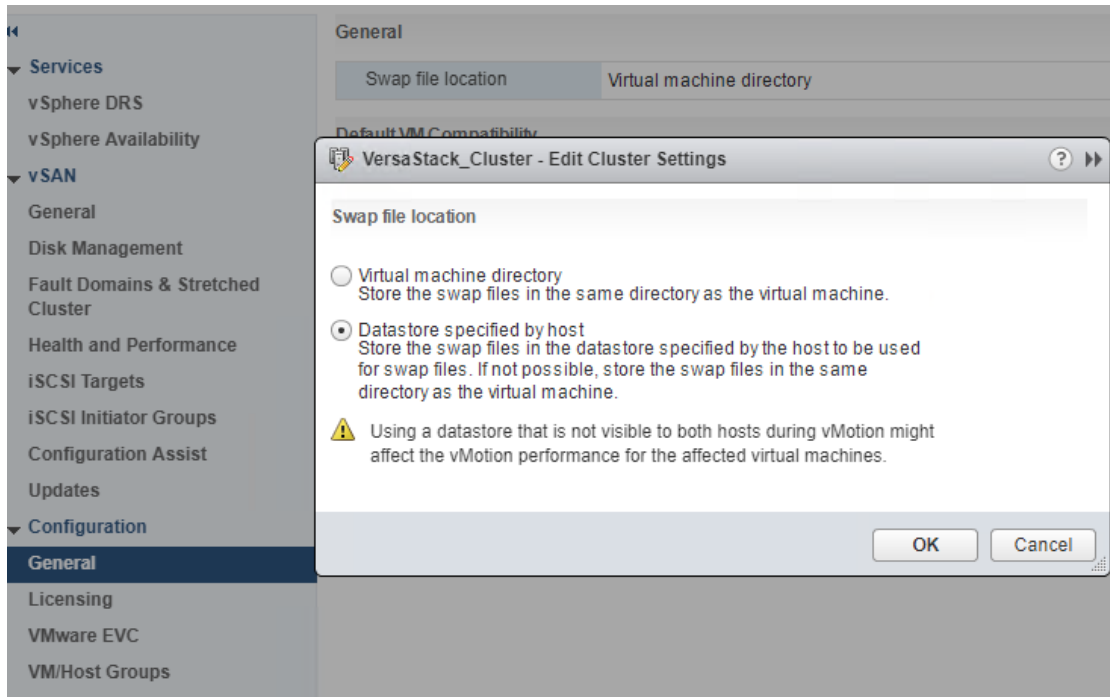
Name	VersaStack_Cluster
Location	VersaStack_DC
DRS	<input checked="" type="checkbox"/> Turn ON
Automation Level	Fully automated
Migration Threshold	Conservative ———— Aggressive
vSphere HA	<input checked="" type="checkbox"/> Turn ON
Host Monitoring	<input checked="" type="checkbox"/> Enable host monitoring
Admission Control	<input checked="" type="checkbox"/> Enable admission control
VM Monitoring	
VM Monitoring Status	Disabled Overrides for individual VMs can be set from the VM Overrides page from Manage Settings area.
Monitoring Sensitivity	Low ———— High
EVC	Disable
vSAN	<input type="checkbox"/> Turn ON

OK Cancel

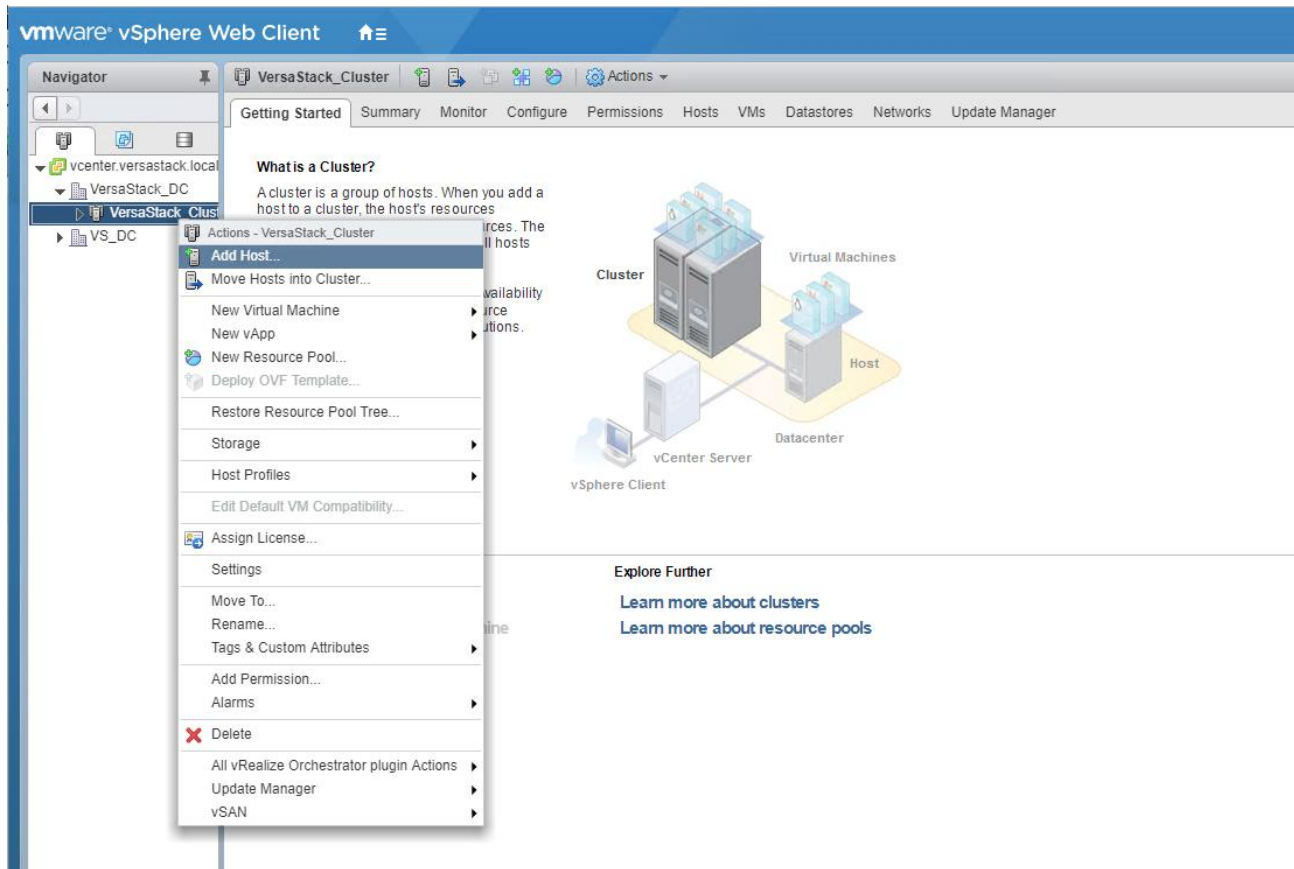


If mixing Cisco UCS B or C-Series M2, M3 or M4 servers within a vCenter cluster, it is necessary to enable VMware Enhanced vMotion Compatibility (EVC) mode. For more information about setting up EVC mode, refer to [Enhanced vMotion Compatibility \(EVC\) Processor Support](#).

3. Click **OK** to create the cluster.
4. Expand "VersaStack_DC".
5. Right-click "VersaStack_Cluster" and select Settings.
6. Select Configuration > General in the list and select **Edit** to the right of General.
7. Select Datastore specified by host and click **OK**.



8. Right-click the newly created cluster and from the drop-down list select the **Add Host**.



9. Enter the IP or FQDN of the first ESXi host and click **Next**.

Add Host

1 **Name and location**

2 Connection settings

3 Host summary

4 Resource pool

5 Ready to complete

Enter the name or IP address of the host to add to vCenter Server.

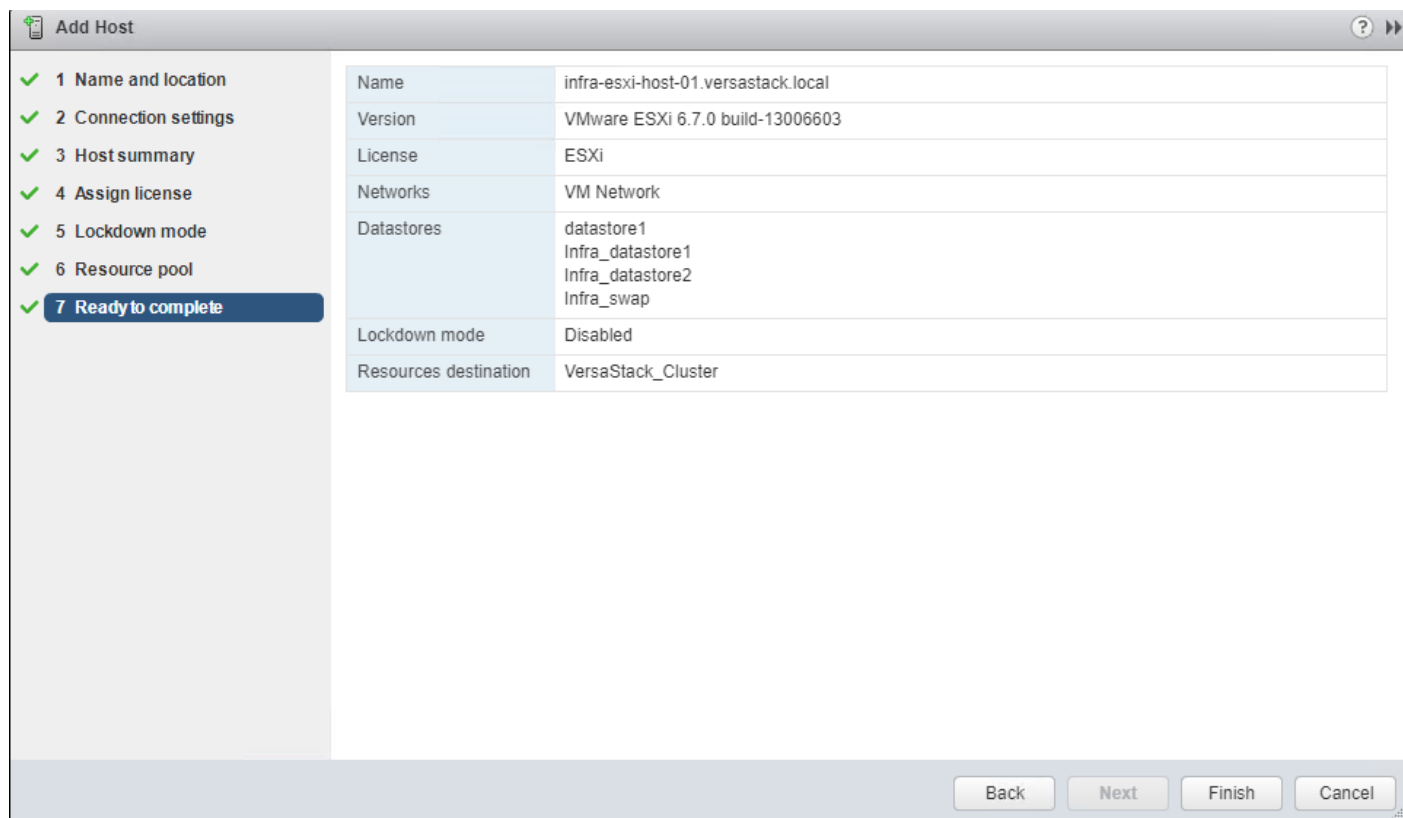
Host name or IP address:

Location:

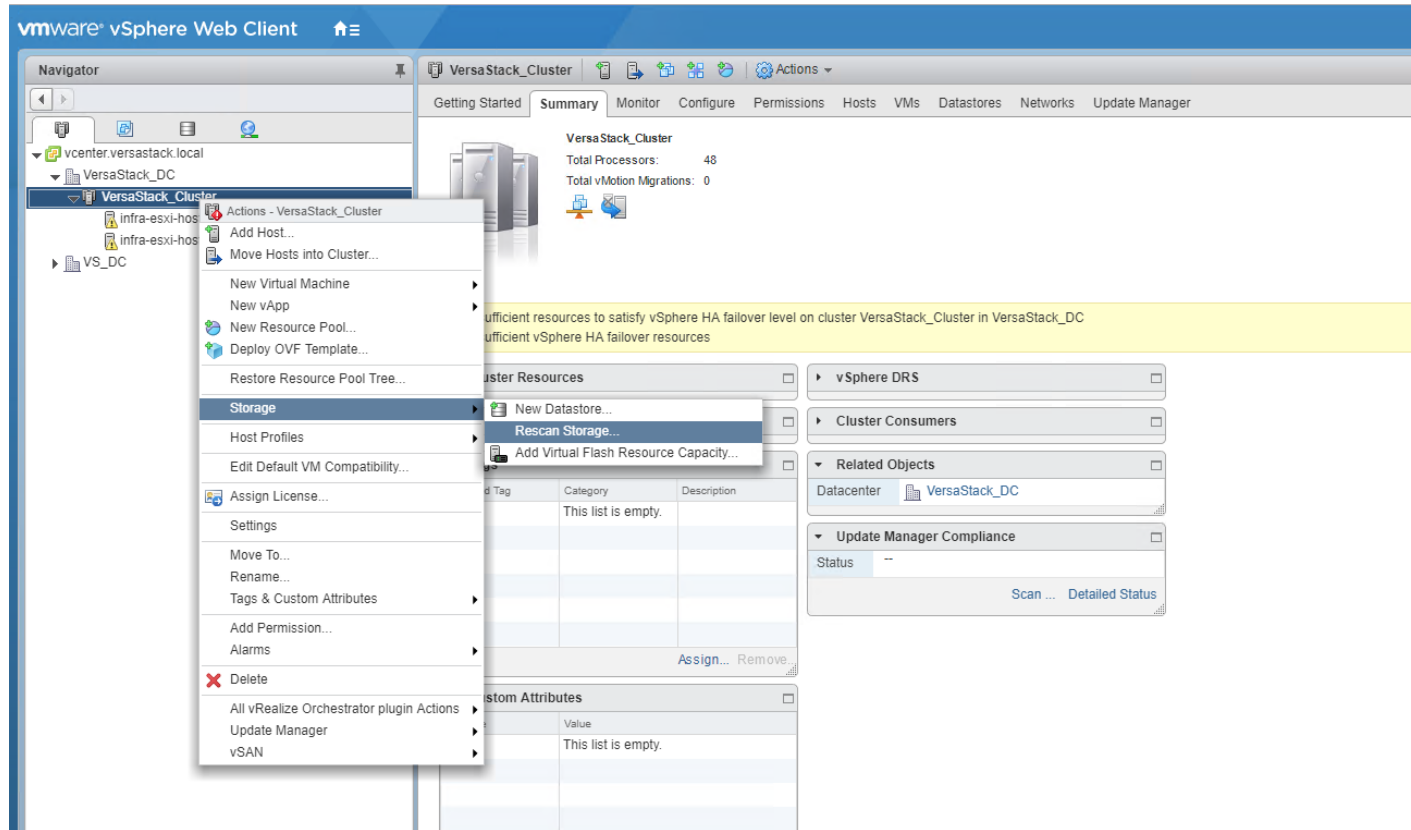
Type:

Back Next Finish Cancel

10. Enter root for the User Name, provide the password set during initial setup and click Next.
11. Click **Yes** in the Security Alert pop-up to confirm the host's certificate.
12. Click **Next** past the Host summary dialogue.
13. Provide a license by clicking the green + icon under the License title, select an existing license, or skip past the Assign license dialogue by clicking **Next**.
14. Leave lockdown mode Disabled within the Lockdown mode dialogue window and click **Next**.
15. Skip past the Resource pool dialogue by clicking **Next**.
16. Confirm the Summary dialogue and add the ESXi host to the cluster by clicking **Next**.



17. Repeat steps 1-16 for each ESXi host to be added to the cluster.
18. In vSphere, in the left pane right-click the newly created cluster, and under Storage click **Rescan Storage**.



19. Click **OK** on the Rescan Storage popup window.

ESXi Dump Collector Setup for iSCSI Hosts (iSCSI Configuration Only)

ESXi hosts booted with iSCSI need to be configured with ESXi dump collection. The Dump Collector functionality is supported by the vCenter but is not enabled by default on the vCenter Appliance.



Make sure the account used to login is Administrator@vsphere.local (or a system admin account).

To setup the ESXi dump collector for iSCSI hosts, follow these steps:

1. In the vSphere web client, select **Home**.
2. In the center pane, click **System Configuration**.
3. In the left-hand pane, select Services and select **VMware vSphere ESXi Dump Collector**.
4. In the Actions menu, choose **Start**.
5. In the Actions menu, click **Edit Startup Type**.
6. Select **Automatic**.
7. Click **OK**.

8. Select Home > Hosts and Clusters.
9. Expand the Data Center and Cluster.
10. For each ESXi host, right-click the host and select **Settings**. Scroll down and select **Security Profile**. Scroll down to Services and select **Edit**. Select **SSH** and click **Start**. Click **OK**.
11. SSH to each ESXi hosts and use root for the user id and the associated password to log into the system. Type the following commands to enable dump collection:

```
[root@Infra-ESXi-Host-01:~] esxcli system coredump network set --interface-name vmk0 --server-ipv4
192.168.162.100 --server-port 6500

[root@Infra-ESXi-Host-01:~] esxcli system coredump network set --enable true

[root@Infra-ESXi-Host-01:~] esxcli system coredump network check

Verified the configured netdump server is running
```

12. **Optional:** Turn off SSH on the host servers.

Configure ESXi Networking

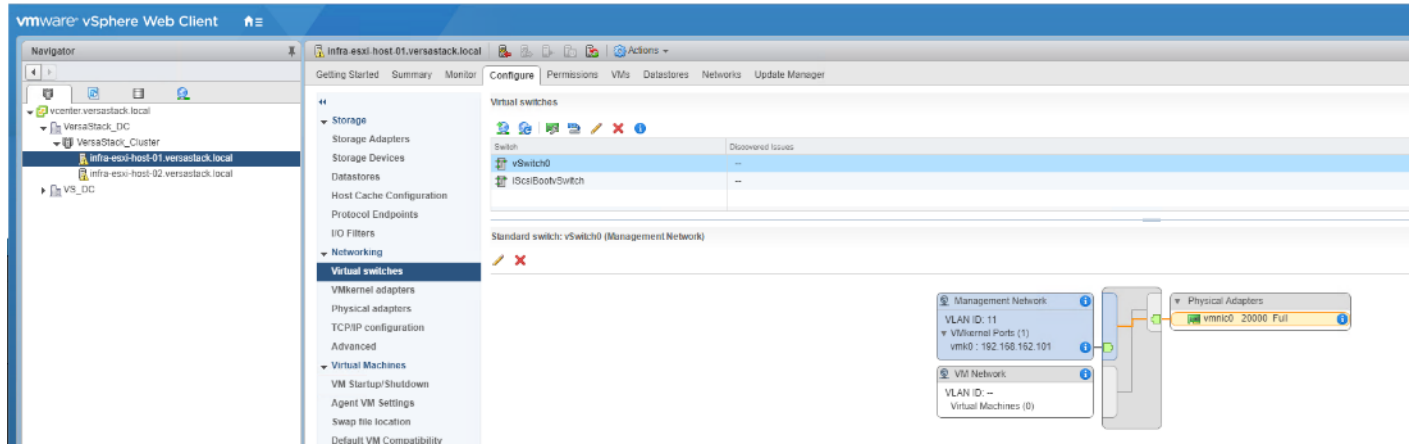
This section covers the virtual switch (vSwitch) setup for Management, vMotion and iSCSI storage traffic and vSphere Distributed Switch (vDS) for application traffic.

Update Management vSwitcho Configuration

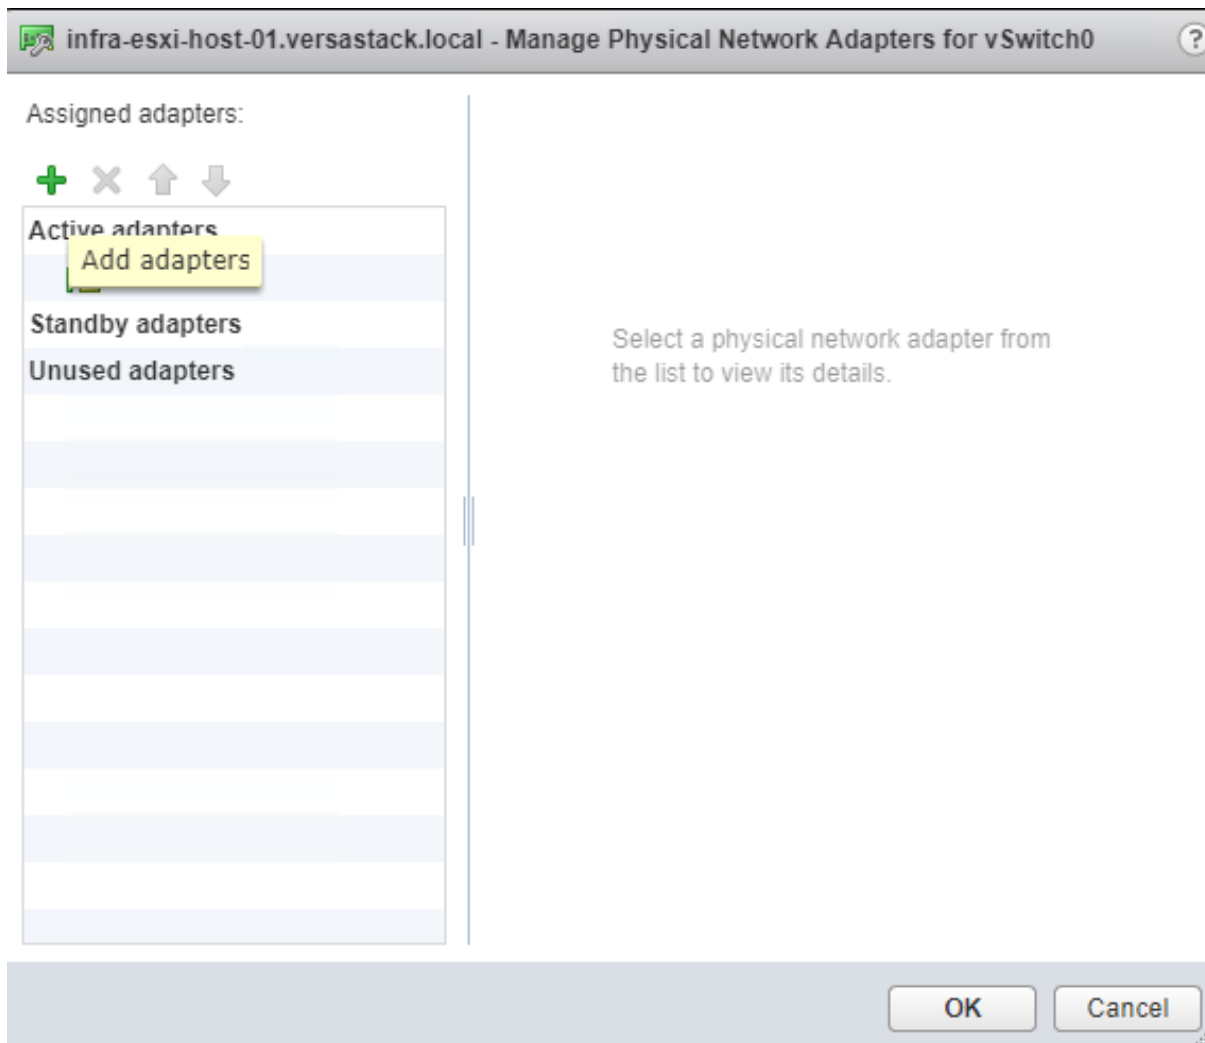
This design uses a separate vSwitch for management using two uplink vNICs with Active/Passive Failover for the port group and routing based on originating port ID for load balancing at the vSwitch level. Traffic from each vNIC take different paths across fabric to enable redundancy and load balancing.

To update the management vSwitcho configuration, follow these steps:

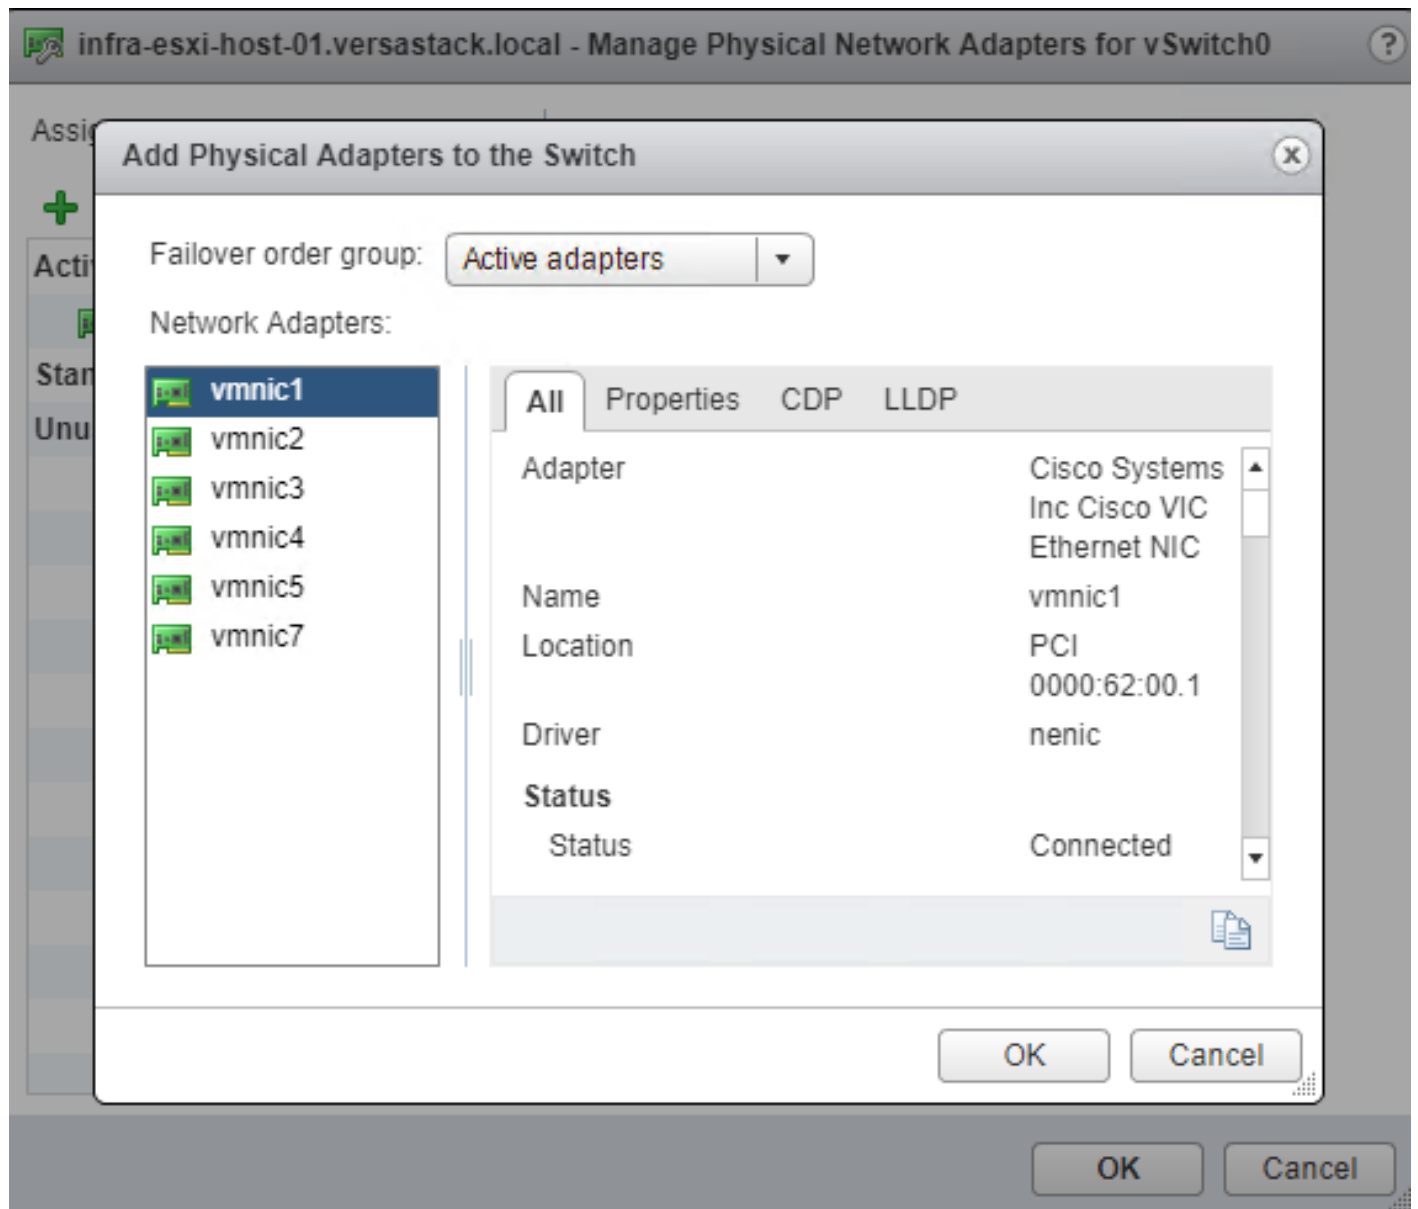
1. Using a web browser, browse to vCenter's IP address. Login to vCenter. From vSphere Web Client, navigate to the datacenter and cluster where the host resides.
2. Select the ESXi host `infra-esxi-host-01`. On the right windowpane, click the **Configure** Tab. Navigate to Networking > Virtual switches and select vSwitcho from the Virtual Switches list. Click the **Manage physical adapter** (third icon) to open the Manage Physical Network Adapters for vSwitcho window.



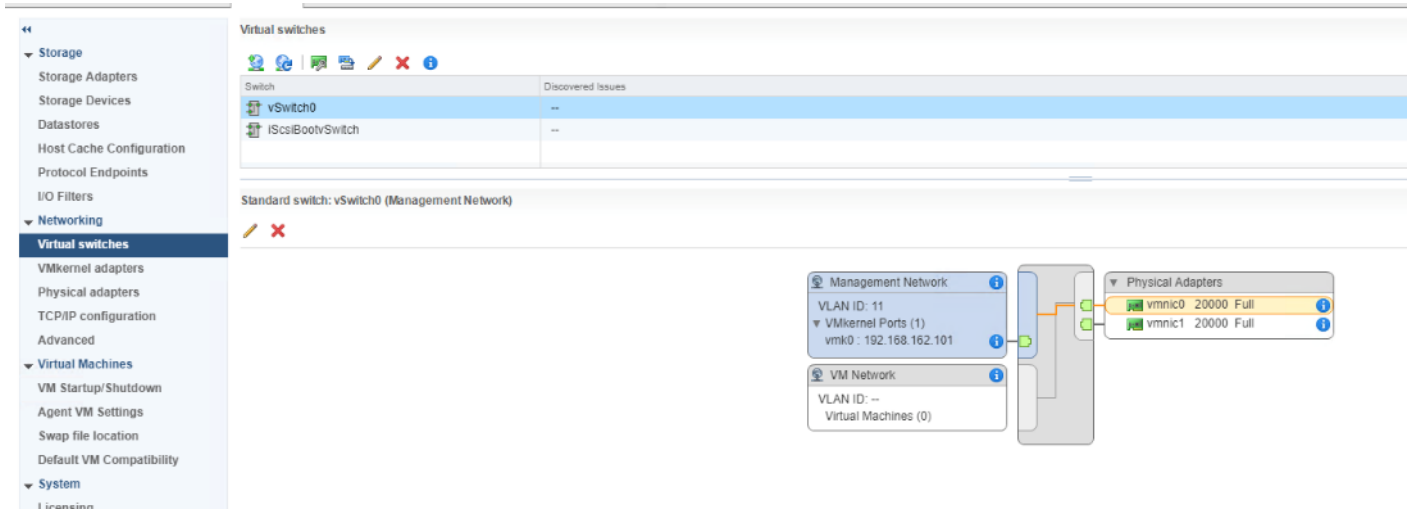
3. Click the **green [+]** to add a second Adapter. Select an unused vmnic from the list of Network Adapters.



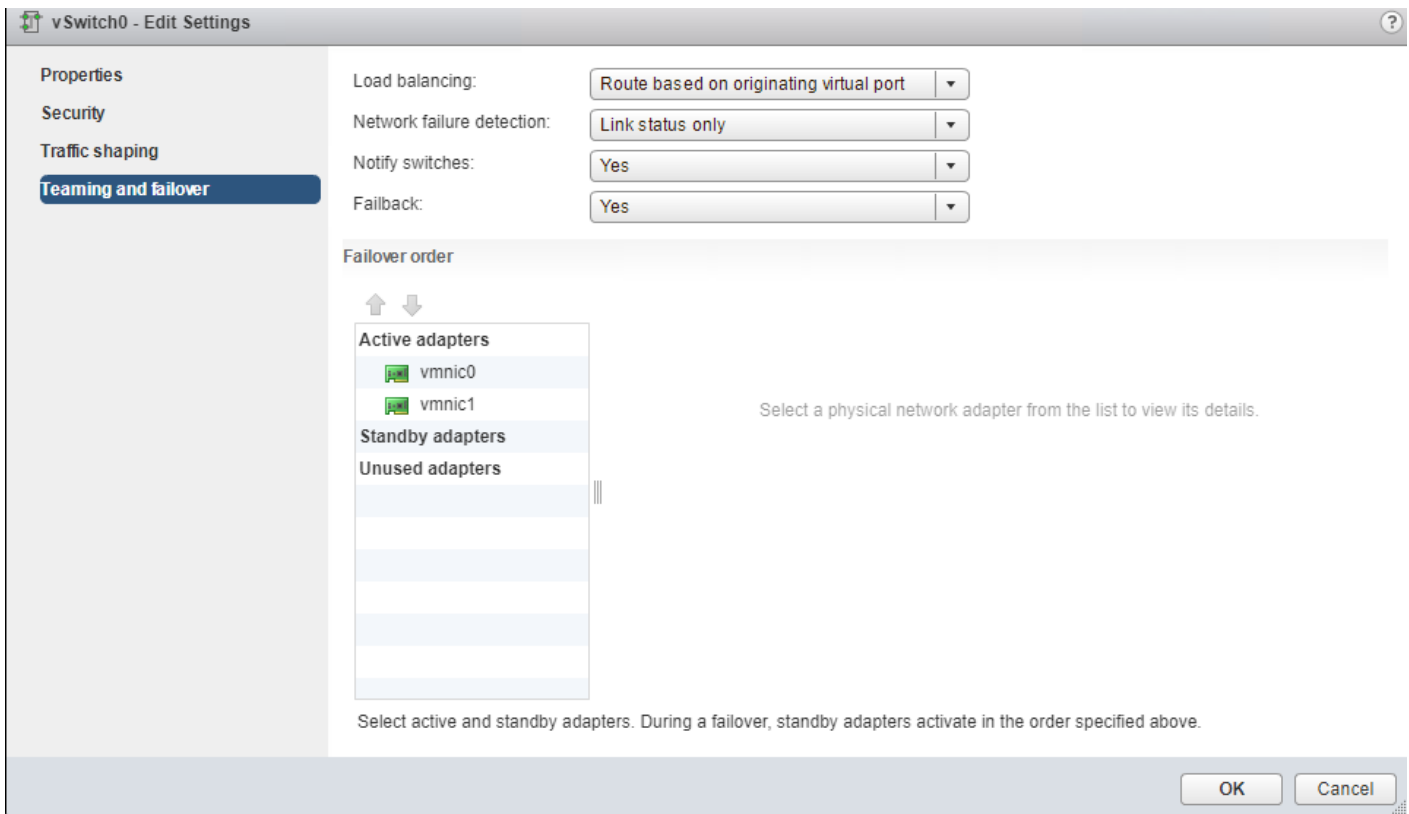
4. Select `vmnic1` and click **OK** to add the vmnic as a second Active adapter to vSwitch0.



5. Click **OK** to commit the change.



6. While the vSwitch0 of the host still selected, click the **Edit Settings** icon (5th icon) to open the Edit settings window.
7. Under Teaming and failover, verify the load balancing and failover configuration. Both vmnics should be listed as up-links under Active adapters.

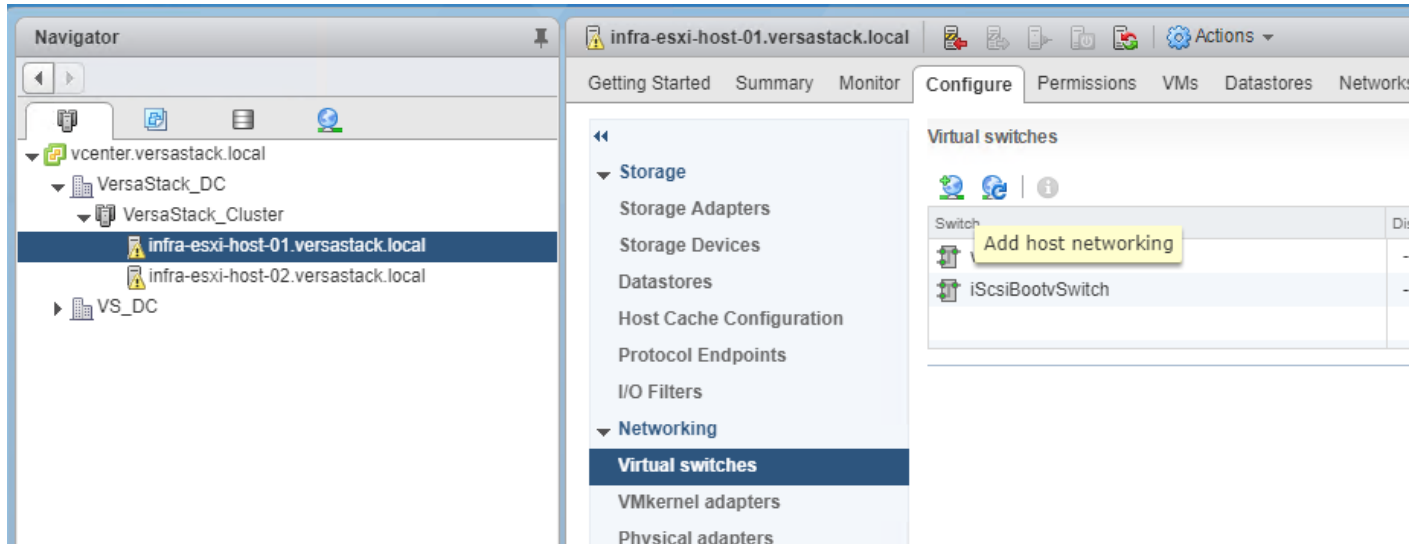


8. Repeat steps 1-7 for all the ESXi hosts in the Cluster.

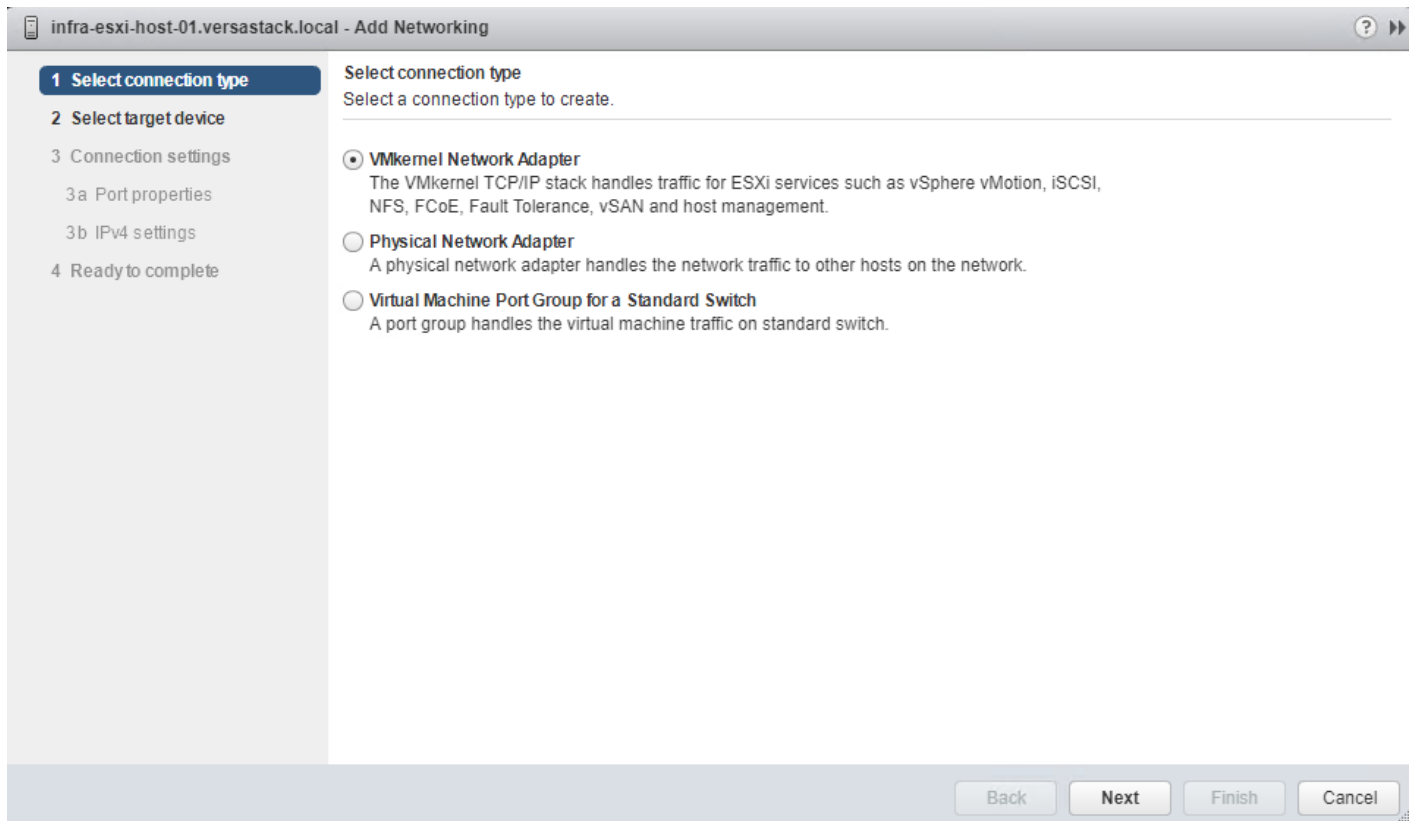
Create vSwitch1 for vMotion

This design uses a separate vSwitch for vMotion with two uplink vNICs. To create and setup vSwitch1 for vMotion, follow these steps:

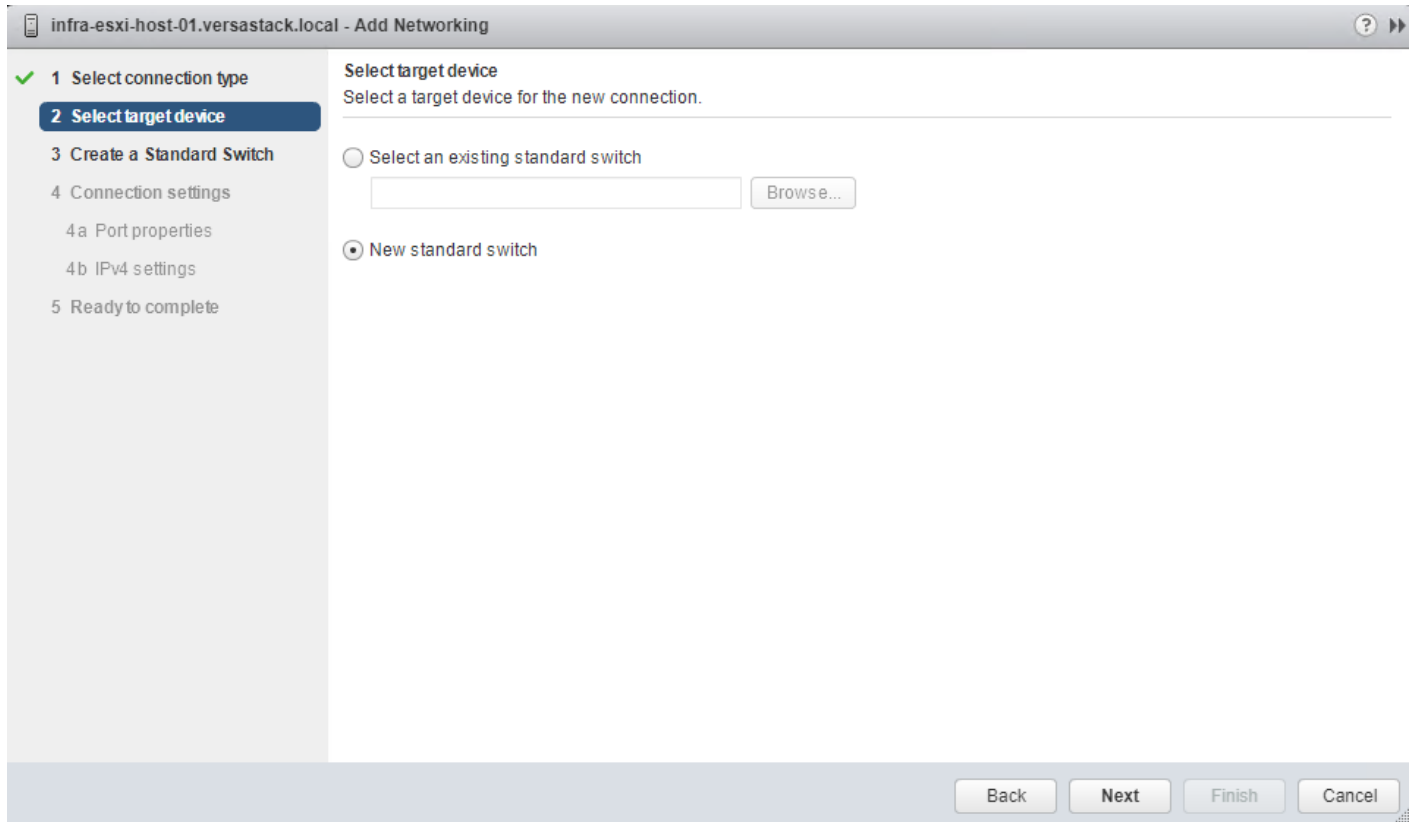
1. Using a web browser, browse to vCenter's IP address. Login to vCenter. From vSphere Web Client, navigate to the datacenter and cluster where the host resides
2. Select the host `infra-esxi-host-01`. On the right windowpane, select the Configure Tab. Navigate to Networking > Virtual Switches. Click the **Add host networking** icon (1st icon).



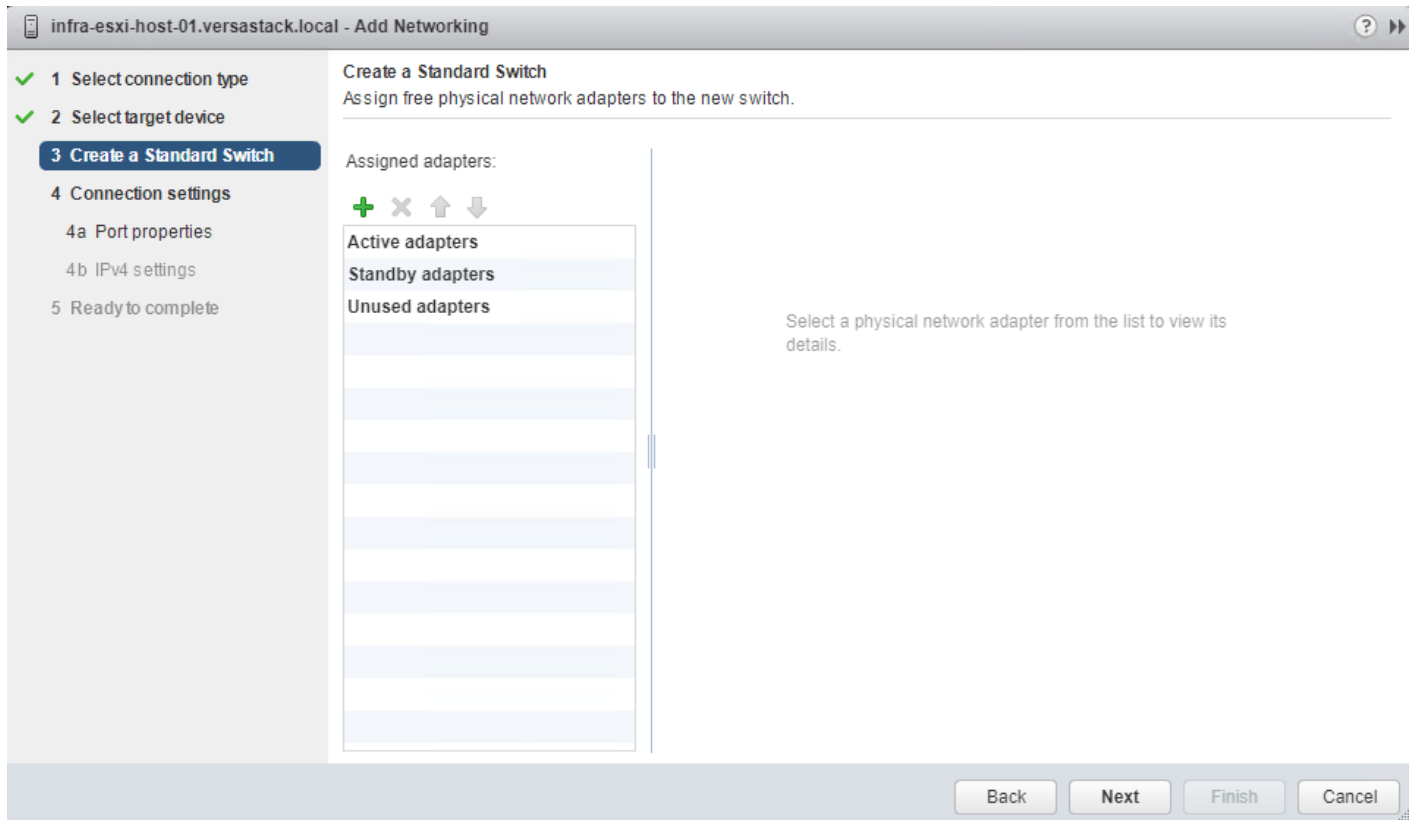
3. Leave VMkernel Network Adapter selected within Select connection type of the Add Networking pop-up window that is generated and click **Next**.



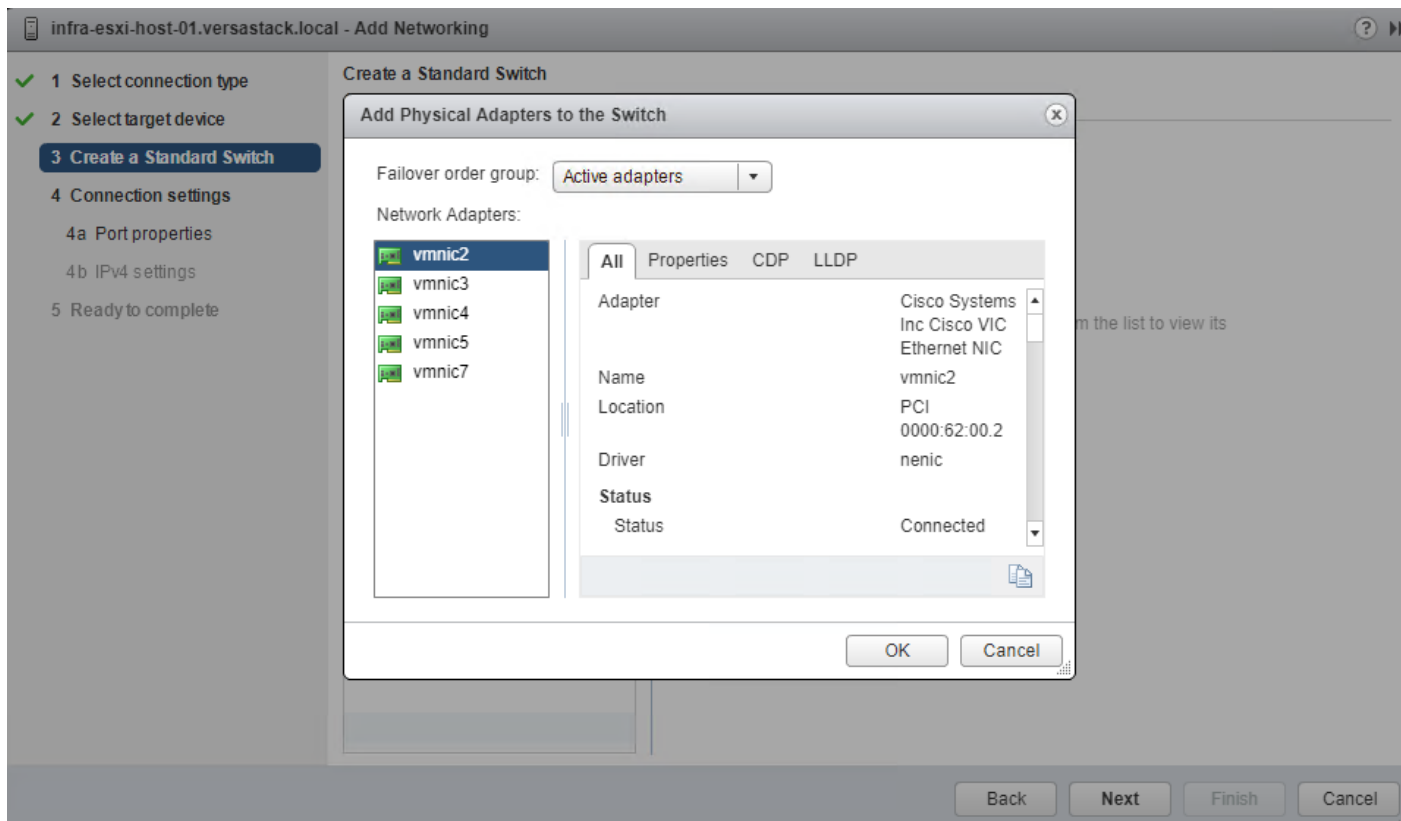
4. Select New Standard switch and click **Next**.



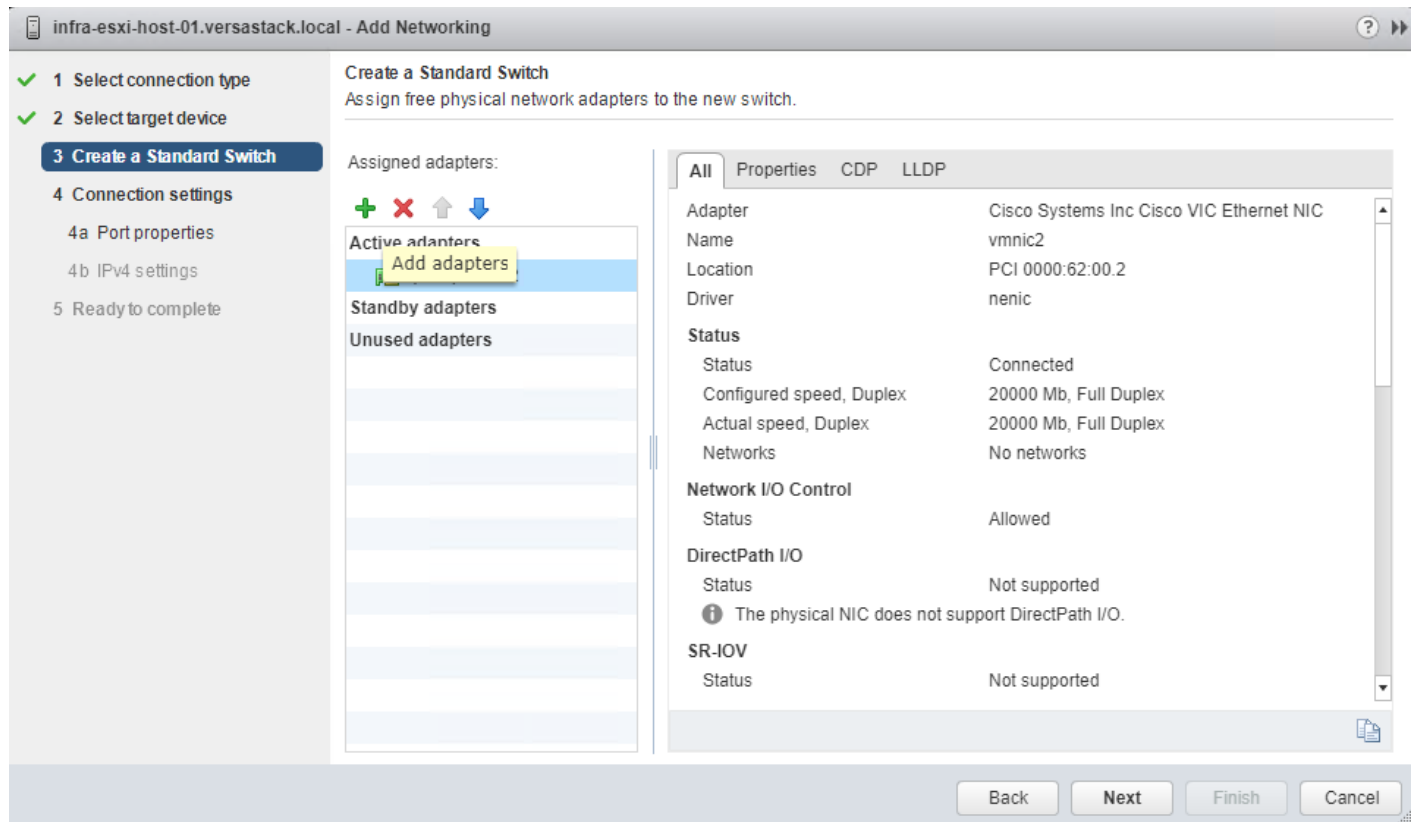
5. Within Select target device, click the New standard switch option, and click **Next**.




6. Within the Create Standard Switch dialogue press the green + icon below Assigned adapters.



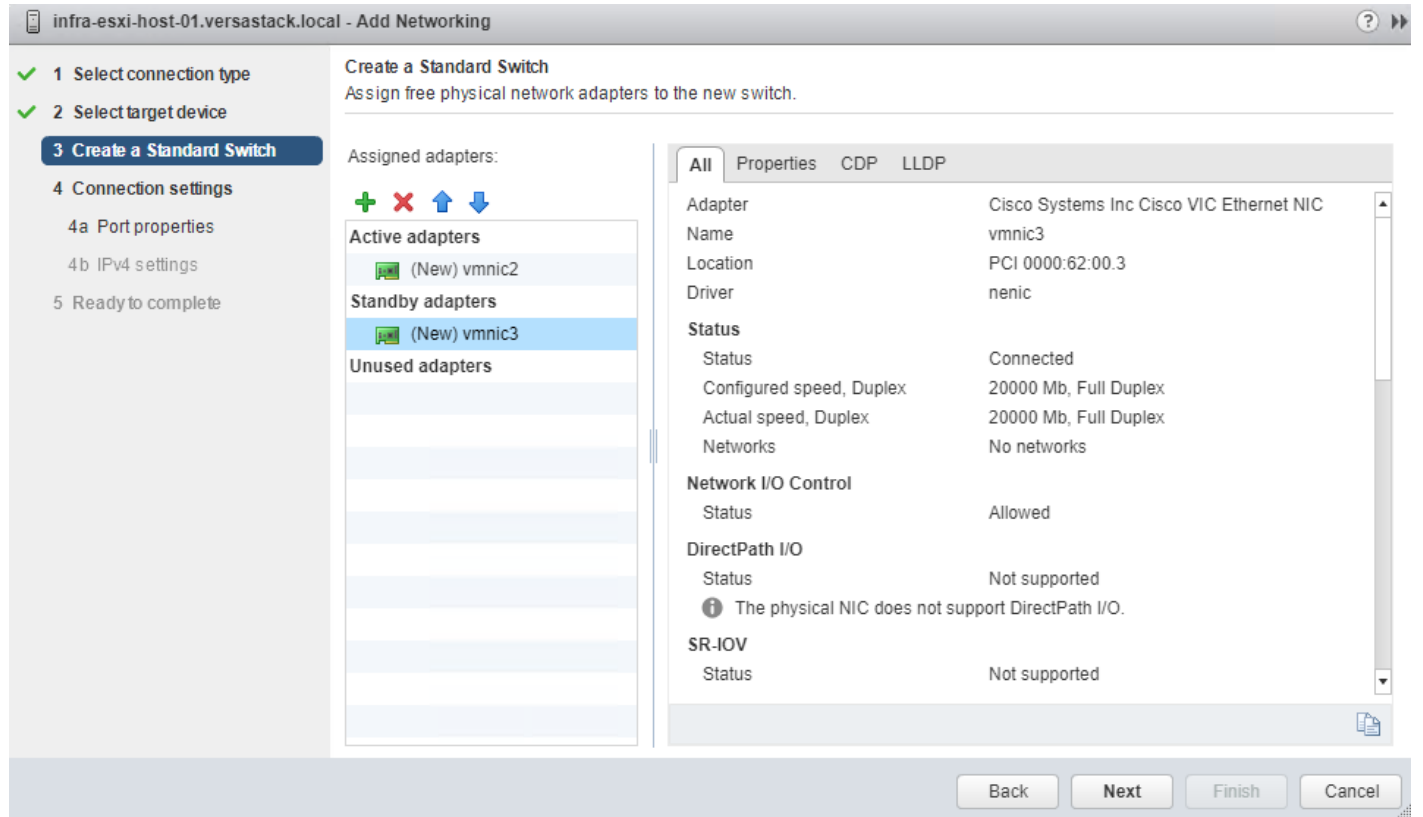
7. Select vmnic2 within the Network Adapters and click **OK**.



8. While still in the Create a Standard Switch dialogue, click the green + icon one more time.



9. Select vmnic3, and from the Failover order group drop-down list, select Standby adapters. Click **OK**.



10. Click **Next**.

infra-esxi-host-01.versastack.local - Add Networking

1 Select connection type
2 Select target device
3 Create a Standard Switch
4 Connection settings
4a Port properties
4b IPv4 settings
5 Ready to complete

Port properties
Specify VMkernel port settings.

VMkernel port settings

Network label: VMkernel vMotion
VLAN ID: 3173
IP settings: IPv4
TCP/IP stack: Default

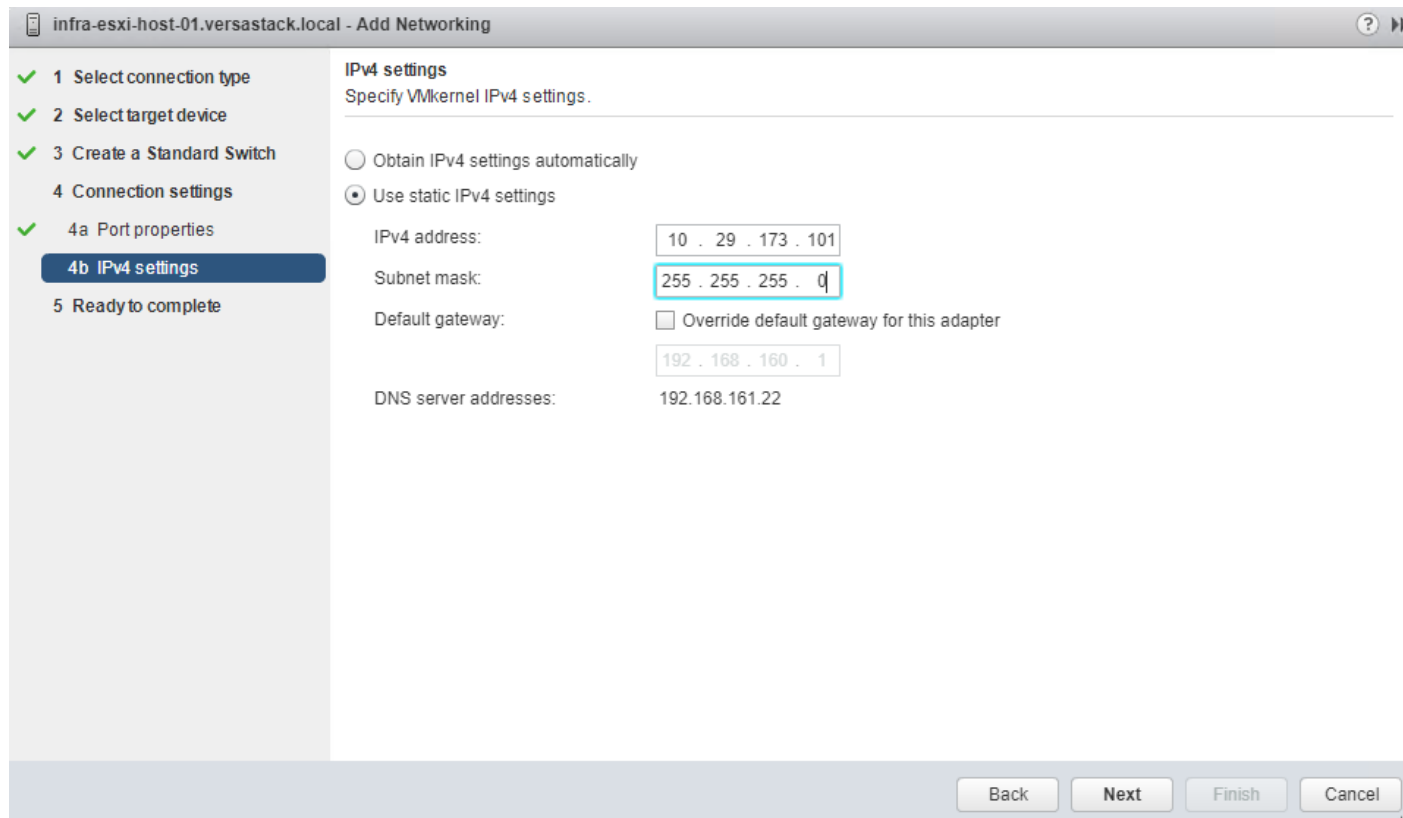
Available services

Enabled services:

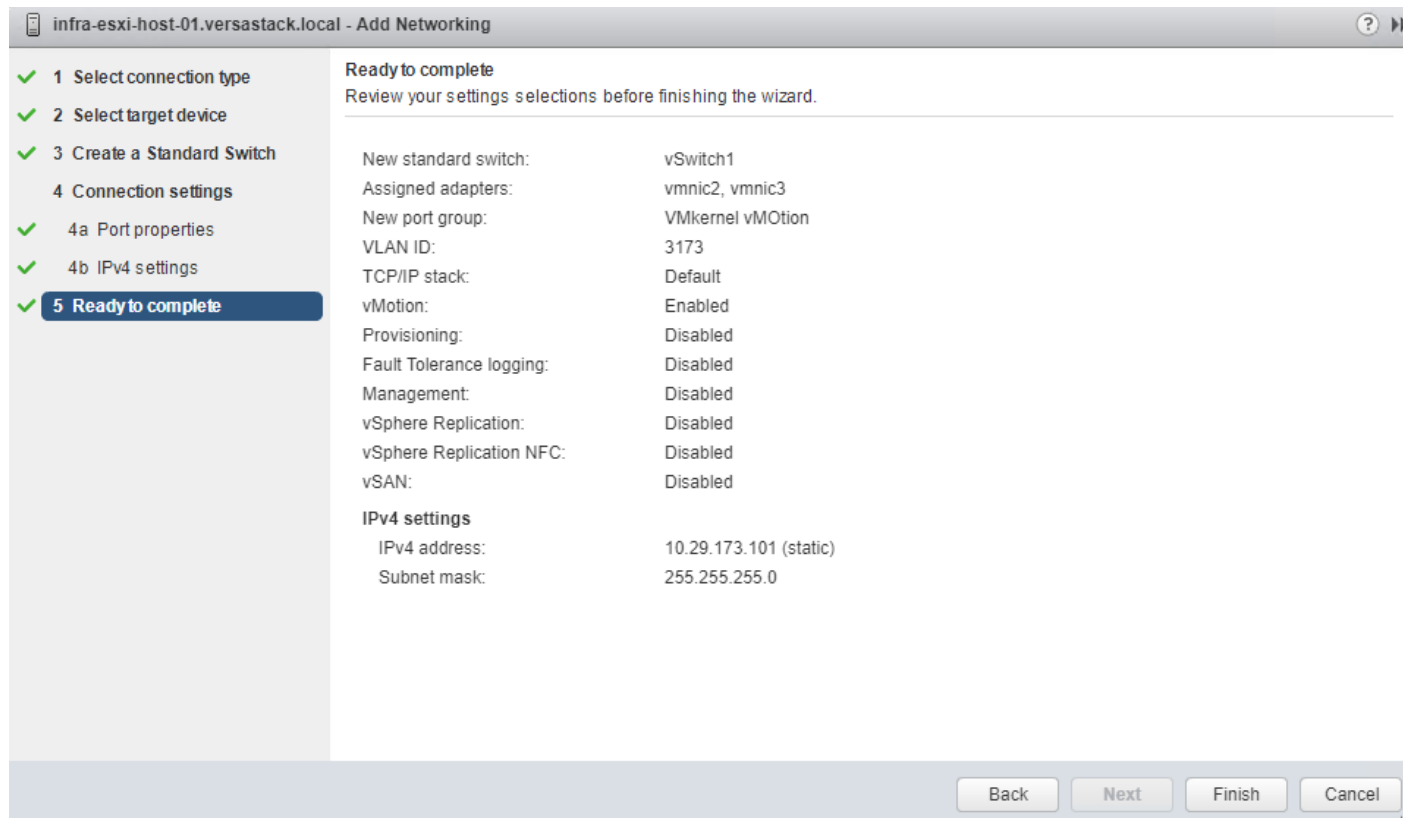
- vMotion
- Provisioning
- Fault Tolerance logging
- Management
- vSphere Replication
- vSphere Replication NFC
- vSAN

Back Next Finish Cancel

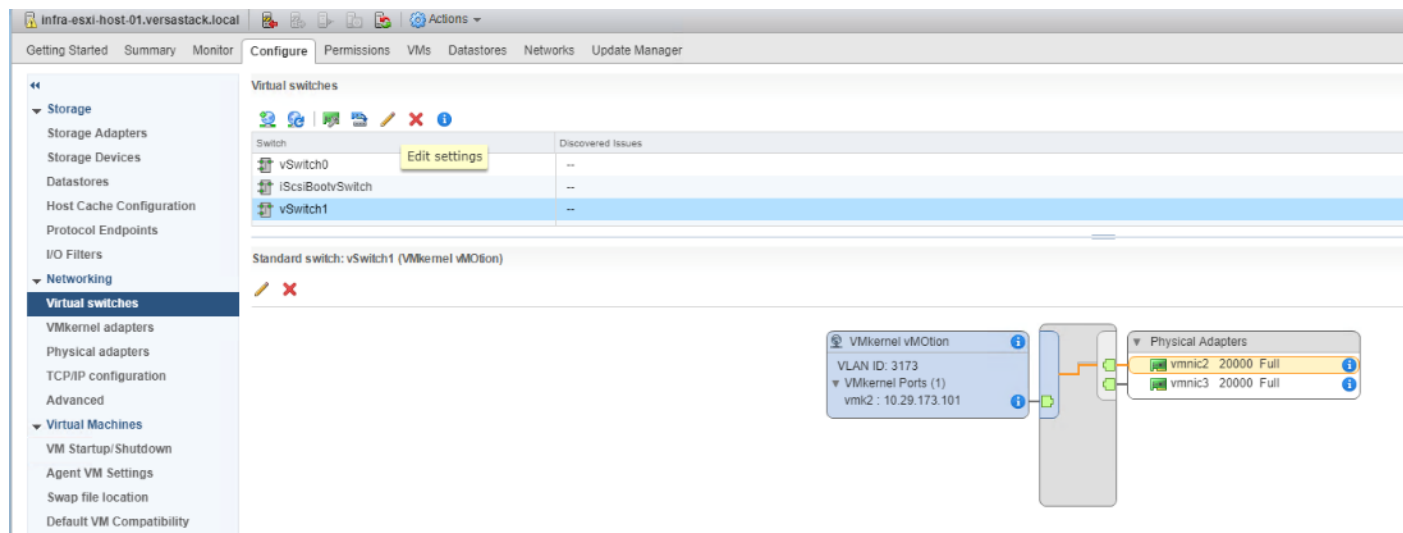
11. Within Port properties under Connection settings, set the Network label to be VMkernel vMotion, set the VLAN ID to the value for <vMotion VLAN id>, and checkmark vMotion traffic under Available services. Click **Next**.



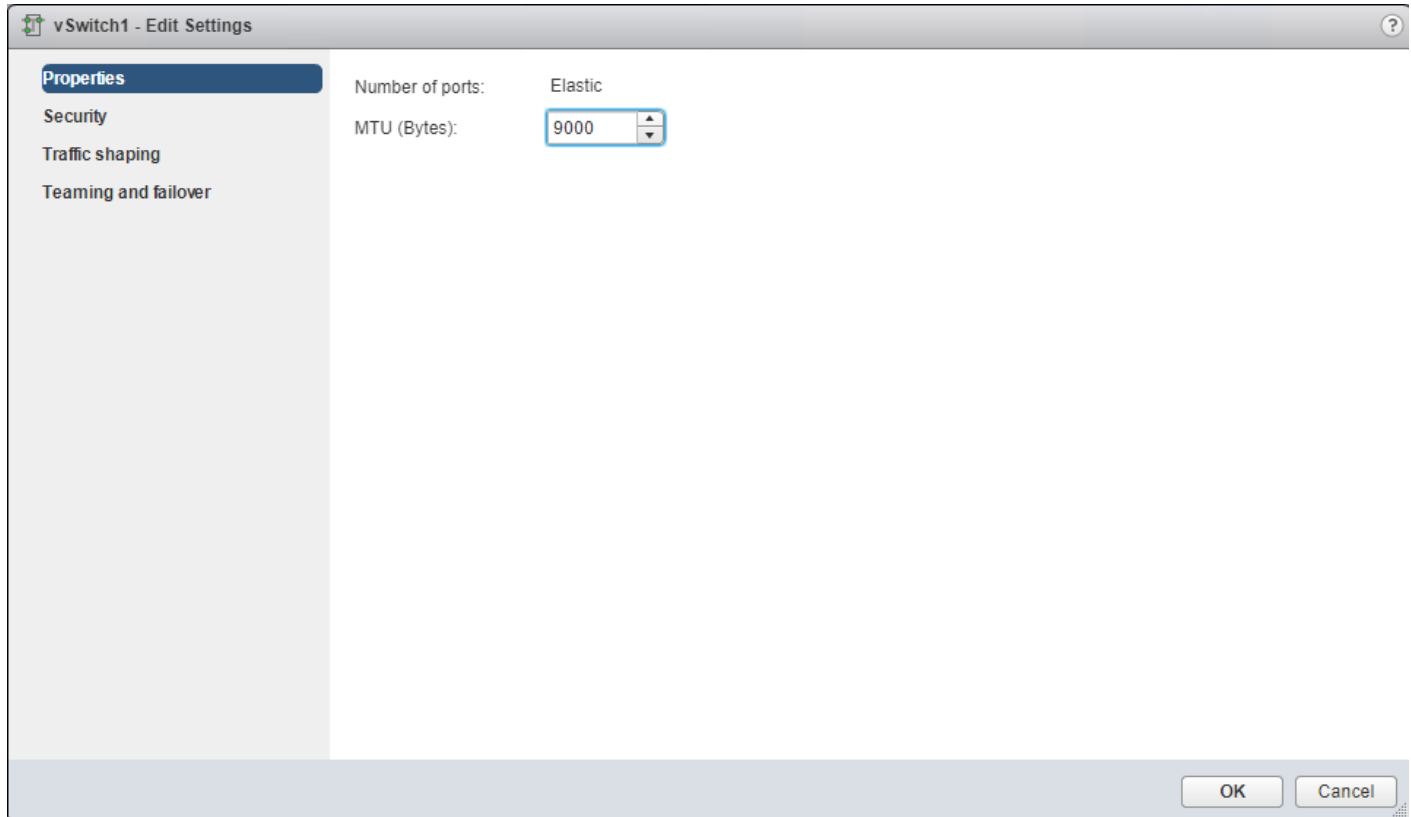
12. Enter <vMotion IP Address> in the field for IPv4 address, and <vMotion Subnet Mask> for the Subnet mask. Click **Next**.



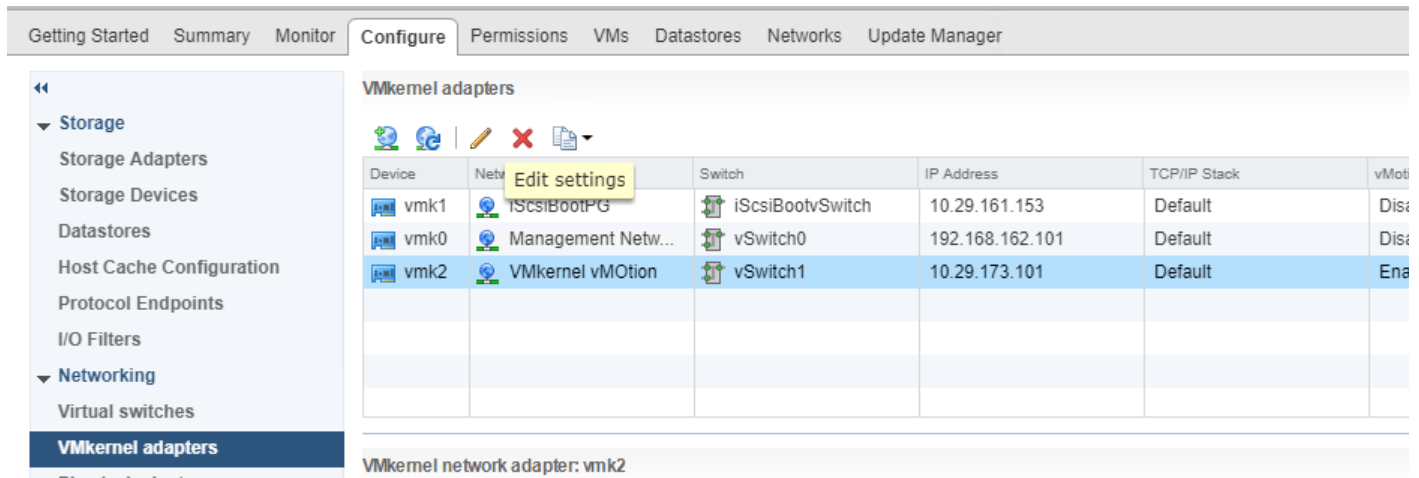
13. Confirm the values shown on the Ready to complete summary page and click **Finish** to create the vSwitch and VMkernel for vMotion.
14. Within the **Configure** tab for the host, under Networking -> Virtual switches, make sure that vSwitch1 is selected, and click on the pencil icon under the Virtual Switches title to edit the vSwitch properties to adjust the MTU for the vMotion vSwitch.




15. Enter 9000 in the Properties dialogue for the vSwitch1 – Edit Settings pop-up that appears. Click **OK** to apply the change.



- Click the VMkernel adapters under **Networking** for the host, and with the VMkernel for vMotion (vmk2) selected, click the pencil icon to edit the VMkernel settings.



- Click the NIC settings in the vmk2 – Edit Settings pop-up window that appears and enter 9000 for the MTU value to use for the VMkernel. Click **OK** to apply the change.



18. Repeat steps 1-17 for each host being added to the cluster, changing the vMotion VMkernel IP to an appropriate unique value for each host.

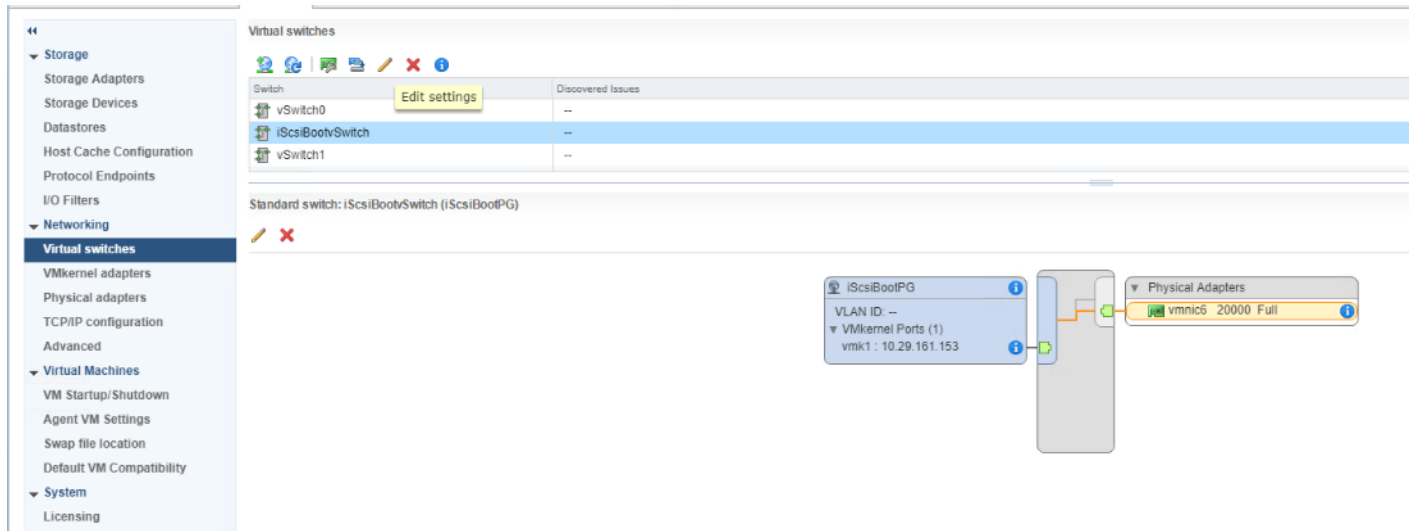
Configure iSCSI Adapters (iSCSI Deployment Only)

The base ESXi installation will set up one vmkernel adapter for the iSCSI boot, with a generated vSwitch named `iScsiBootvSwitch`. vSwitch changes will be needed, as well as the creation of a second vmkernel adapter used for the B side iSCSI boot. To make the vSwitch changes and create the vmkernel adapter, follow these steps for each host:

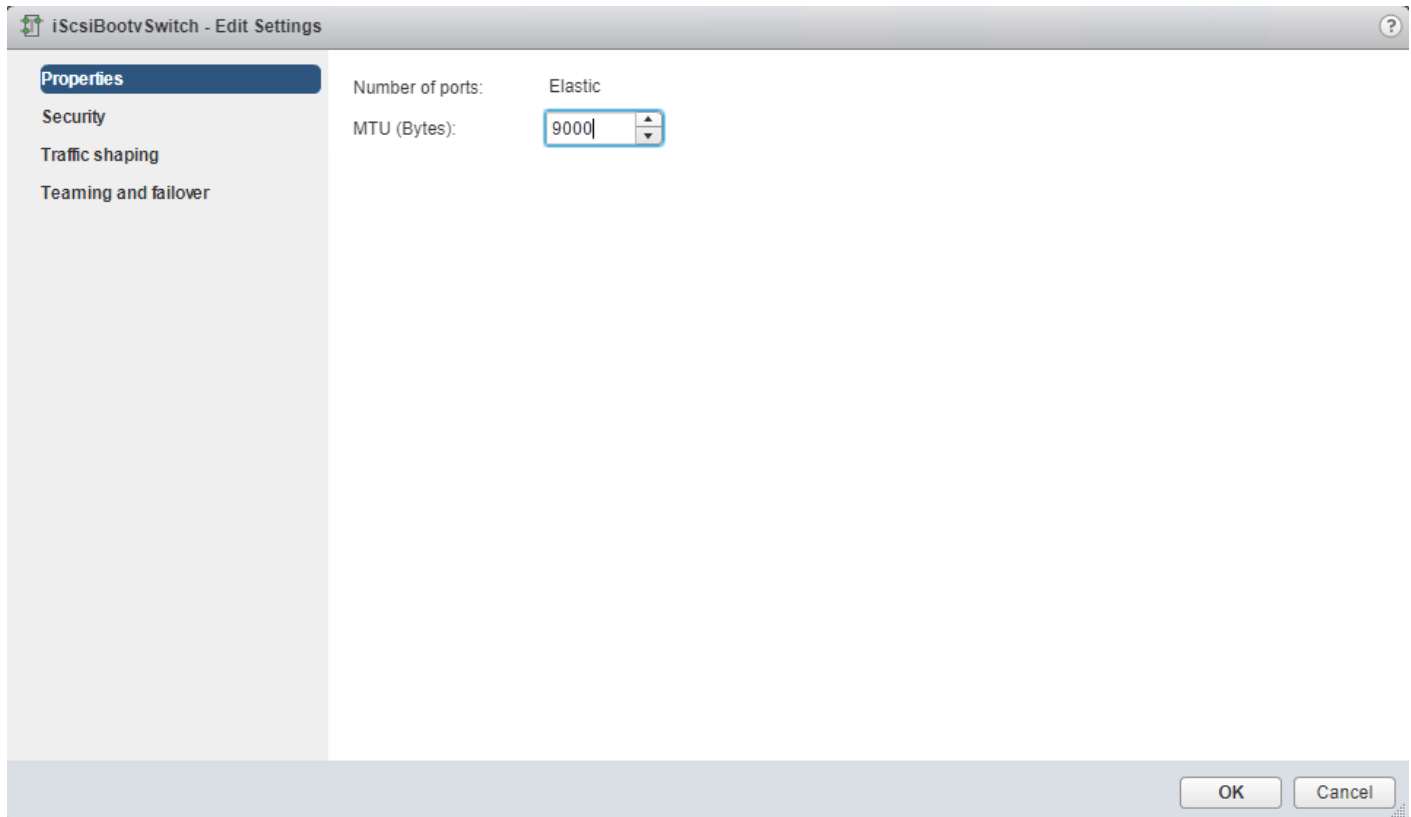
Adjust iSCSI A vSwitch MTU

To adjust iSCSI A vSwitch MTU, follow these steps:

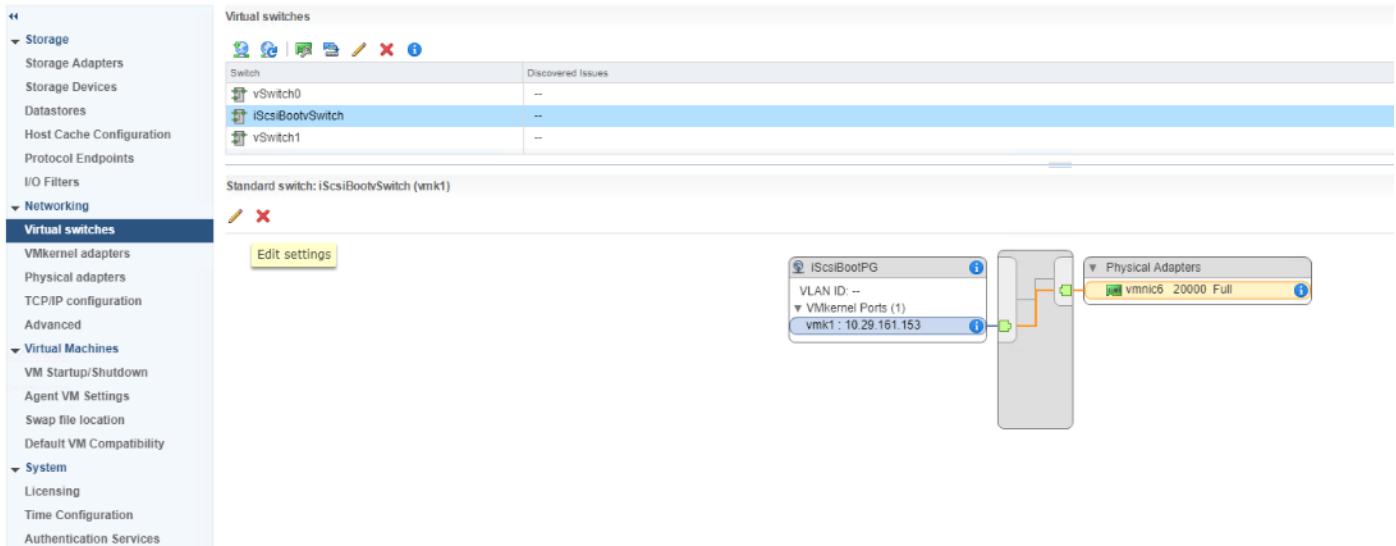
1. From the vSphere Web Client, select the installed iSCSI host, click the **Configure** tab, and select the Virtual switches section from the Networking section on the left.
2. Select the `iScsiBootvSwitch` and click the pencil icon to open up Edit settings for the vSwitch.




3. Within the Properties section, change the MTU from 1500 to 9000 and click **OK** to save the changes.



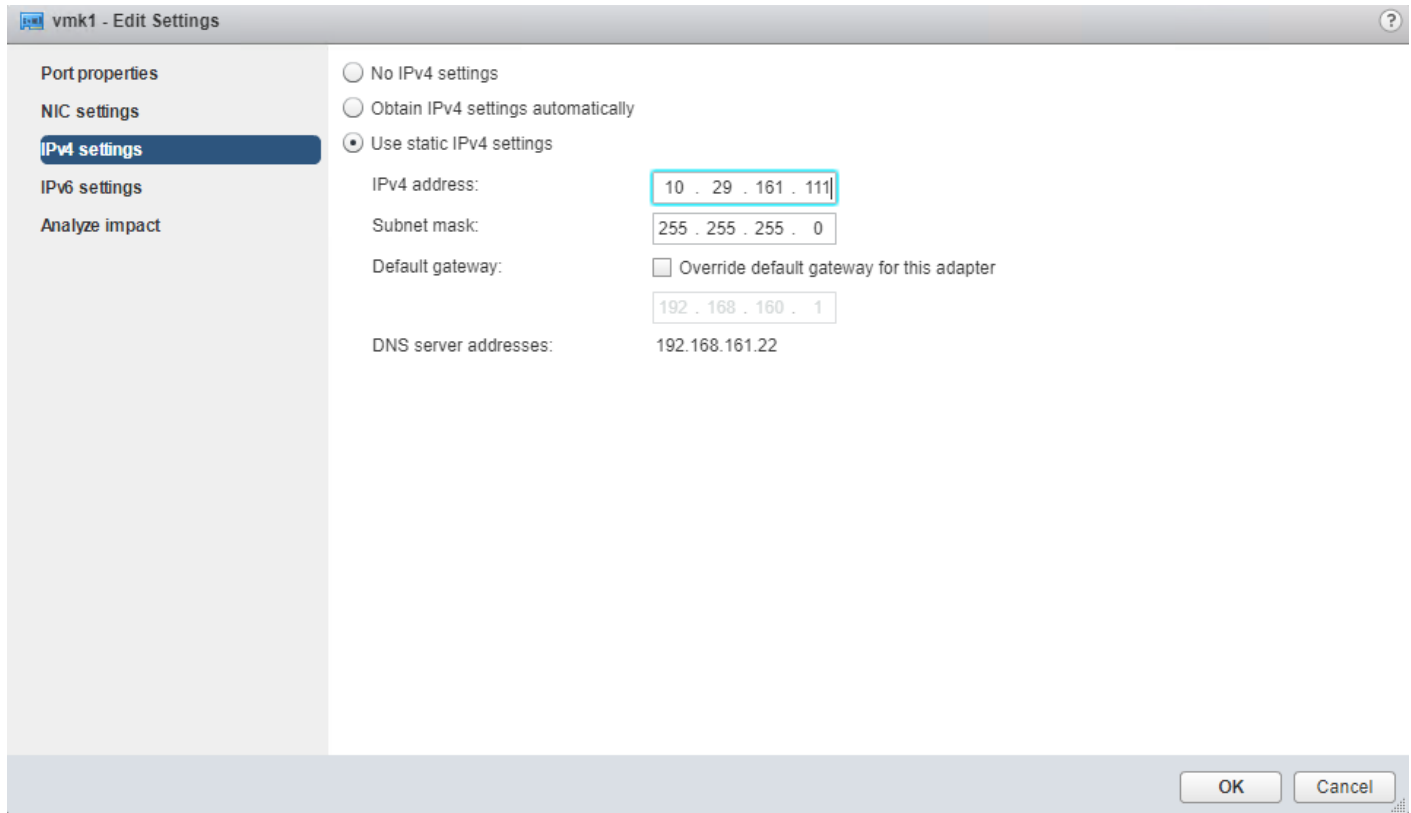
4. Click the vmk1 entry within the iScsiBootPG and select the pencil icon on the left to edit the settings of the vmkernel adapter.



5. Select NIC settings on the left side of the Edit Settings window and adjust the MTU from 1500 to 9000.



6. Click the IPv4 settings for vmk1 and change the IPv4 settings from the Cisco UCS Manager iSCSI-initiator-A assigned IP to one that is not in the IP block.

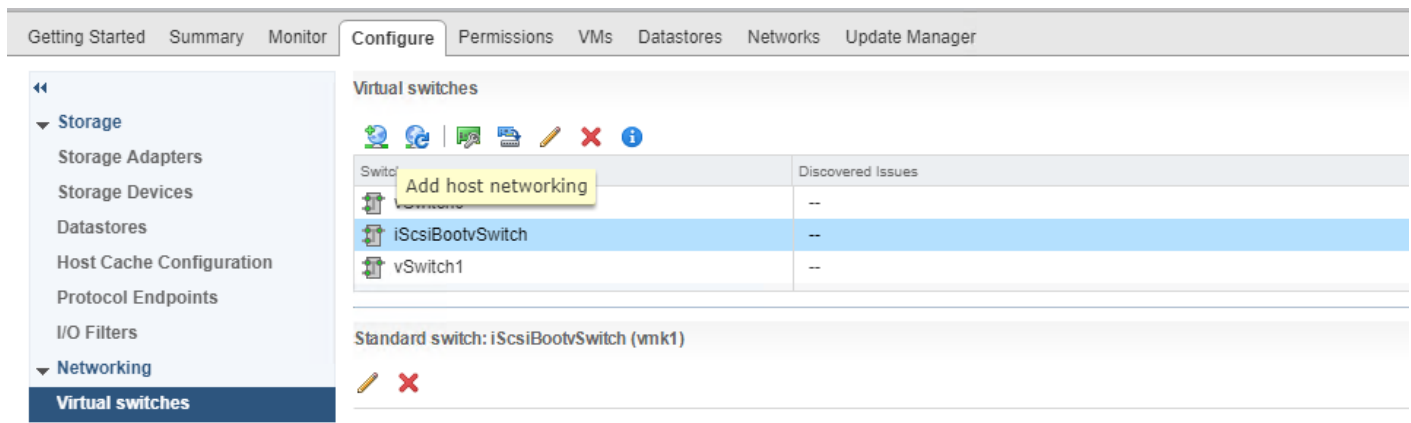


7. Click **OK** to apply the changes to the vmkernel adapter.

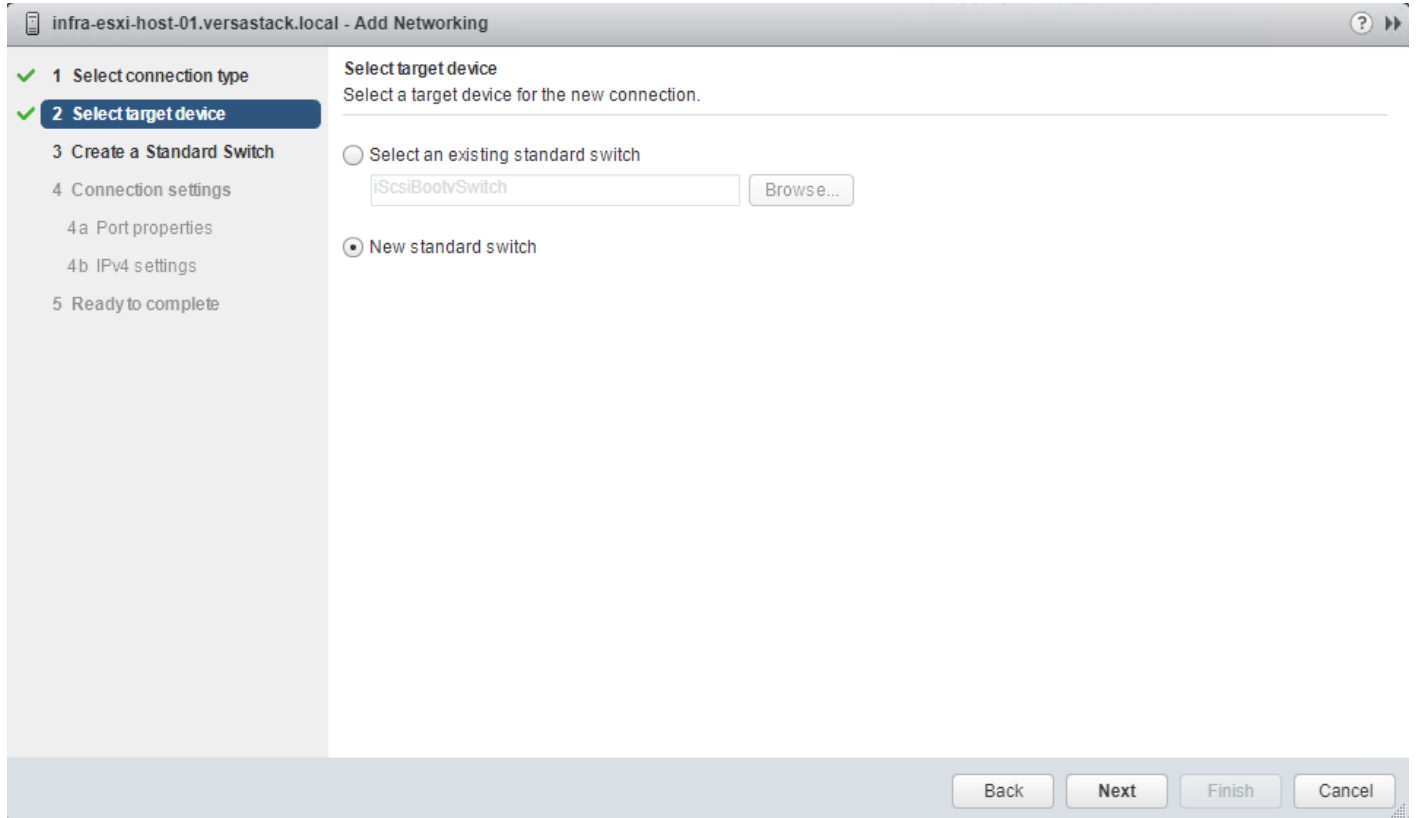
Create iSCSI B vSwitch and vmkernel Adapter

To create the iSCSI B vSwitch and vmkernel adapter, follow these steps:

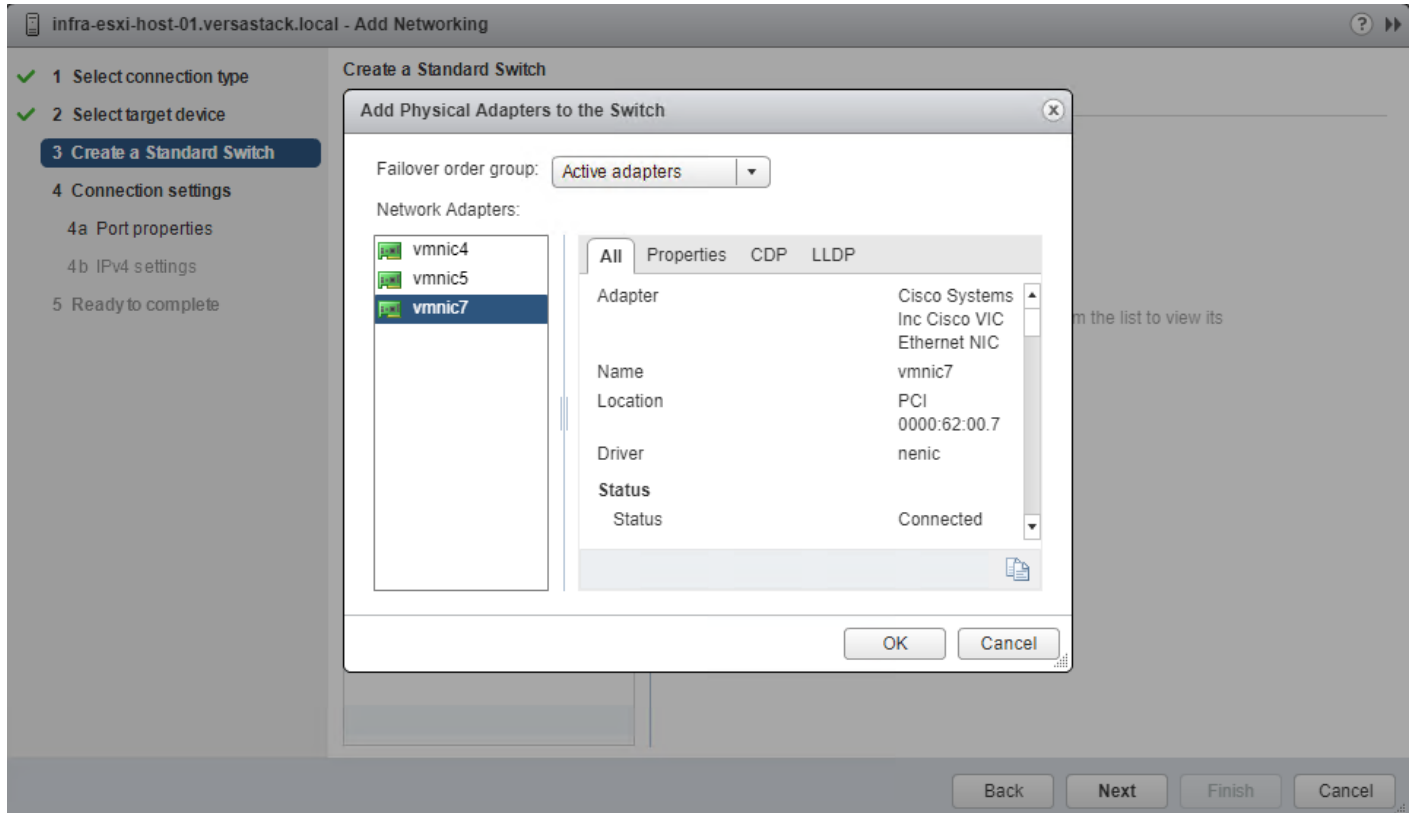
1. Click the Add host networking icon under Virtual switches.



2. Leave VMkernel Network Adapter selected and click **Next**.
3. Change the Select target device option to New standard switch and click **Next**.



4. Click the green plus icon under Assigned adapters and select vmnic7 from the listed adapters in the resulting window.



5. Click **OK** to add the vmnic to the vSwitch and click Next.
6. **(Optional)** Enter a relevant name for the Network label.

infra-esxi-host-01.versastack.local - Add Networking

1 Select connection type
2 Select target device
3 Create a Standard Switch
4 Connection settings
 4a Port properties
 4b IPv4 settings
5 Ready to complete

Port properties
Specify VMkernel port settings.

VMkernel port settings

Network label: VMkernel-iSCSI-B

VLAN ID: None (0)

IP settings: IPv4

TCP/IP stack: Default

Available services

Enabled services:

- vMotion
- Provisioning
- Fault Tolerance logging
- Management
- vSphere Replication
- vSphere Replication NFC
- vSAN

Back Next Finish Cancel


7. Click **Next**.
8. Change the option for IPv4 settings to Use static IPv4 settings and enter valid IP and subnet mask information that is outside of the iSCSI-initiator-B pool.

9. Click Next and click Finish in the resulting Summary window.

Setup iSCSI Multipathing

To setup the iSCSI multipathing on the ESXi hosts, follow these steps:

1. From the vSphere Web Client, select the host and select the Configure tab within the host view.
2. Select Storage Adapters from within the Storage section and vmhba64 under the iSCSI Software Adapter listing.
3. Select the Targets tab under the Adapter Details.



4. With Dynamic Discovery selected, click **Add**.
5. Enter the first iSCSI interface IP address for IBM FS9100 Node 1 storage from Table 10 and click **OK**.

vmhba64 - Add Send Target Server ?

iSCSI Server:






Port:

Authentication Settings

Inherit settings from parent

6. Click **OK** and repeat the previous step to add all IP addresses for all the FS9100 nodes.

Storage Adapters

Adapter	Type	Status	Identifier	Targets	Devices	Paths
Lewisburg SATA AHCI Controller						
vmhba0	Block SCSI	Unknown		0	0	0
iSCSI Software Adapter						
vmhba64	iSCSI	Online	iqn.1992-08.com.cisco:ucs-host:3	1	4	4

Due to recent configuration changes, a rescan of this storage adapter is recommended.

Adapter Details






Properties Devices Paths **Targets** Network Port Binding Advanced Options

iSCSI server

- 10.29.161.249:3260
- 10.29.161.250:3260
- 10.29.162.249:3260
- 10.29.162.250:3260

- Rescan the storage adapters with the third icon at the top of the page.
- Click **OK** on the pop-up window.

Storage Adapters

Adapter	Type	Status	Identifier	Targets
Lewisburg SATA AH				
vmhba0				0
iSCSI Software Adapter				
vmhba64	iSCSI	Online	iqn.1992-08.com.cisco:ucs-host:3	1

Rescans the host's storage adapter to discover newly added storage devices.

- Observed Paths should now be four times what it previously was.

Adapter Details

Properties Devices **Paths** Targets Network Port Binding Advanced Options

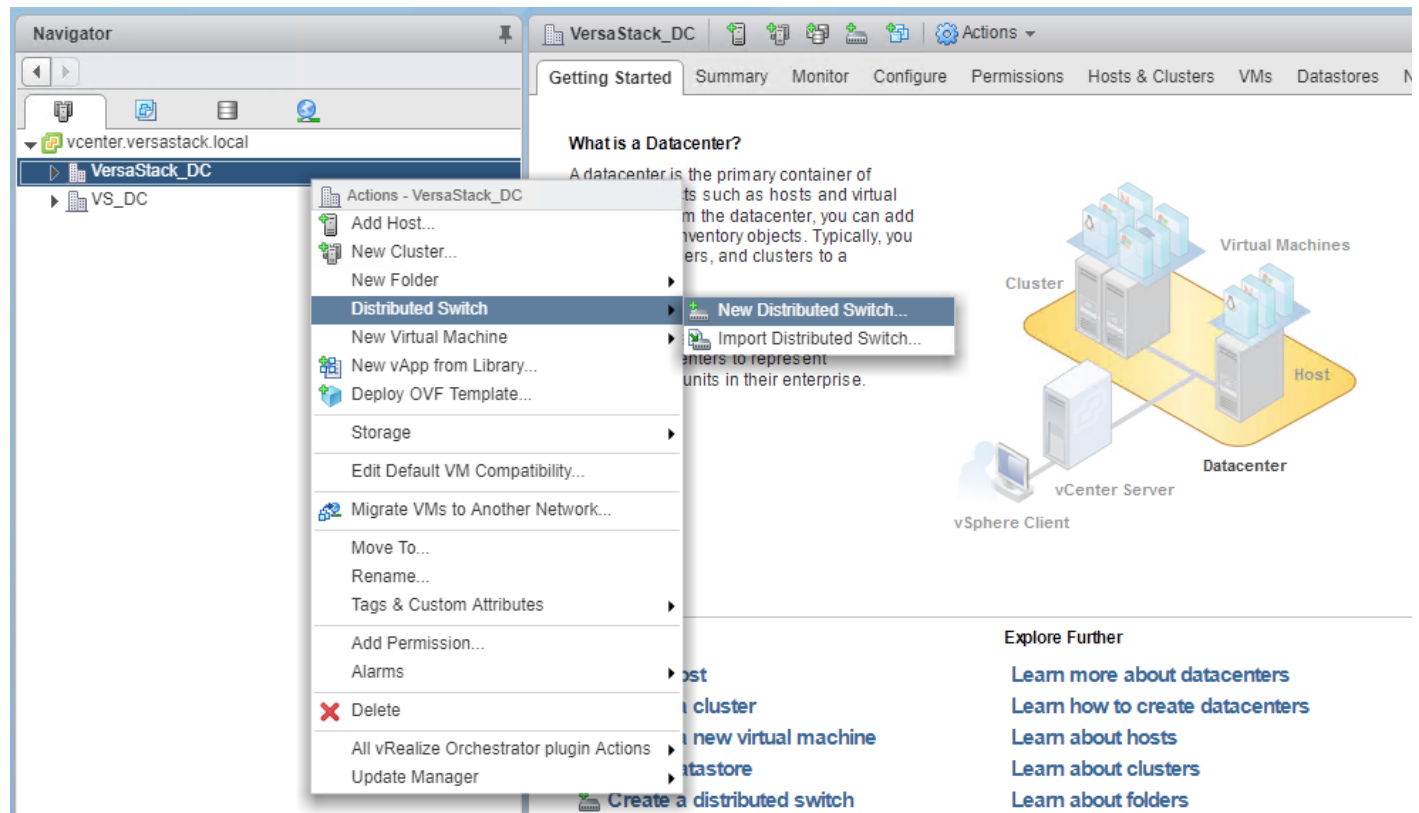
Enable Disable

Runtime Name	Target	LUN	Status
vmhba64:C0:T0:L0	iqn.1986-03.com.ibm:2145.versa...	0	Active (I/O)
vmhba64:C0:T0:L1	iqn.1986-03.com.ibm:2145.versa...	1	Active (I/O)
vmhba64:C0:T0:L2	iqn.1986-03.com.ibm:2145.versa...	2	Active
vmhba64:C0:T0:L3	iqn.1986-03.com.ibm:2145.versa...	3	Active (I/O)
vmhba64:C1:T1:L0	iqn.1986-03.com.ibm:2145.versa...	0	Active
vmhba64:C1:T1:L1	iqn.1986-03.com.ibm:2145.versa...	1	Active
vmhba64:C1:T1:L2	iqn.1986-03.com.ibm:2145.versa...	2	Active (I/O)
vmhba64:C1:T1:L3	iqn.1986-03.com.ibm:2145.versa...	3	Active
vmhba64:C1:T0:L0	iqn.1986-03.com.ibm:2145.versa...	0	Active (I/O)
vmhba64:C1:T0:L1	iqn.1986-03.com.ibm:2145.versa...	1	Active (I/O)
vmhba64:C1:T0:L2	iqn.1986-03.com.ibm:2145.versa...	2	Active
vmhba64:C1:T0:L3	iqn.1986-03.com.ibm:2145.versa...	3	Active (I/O)

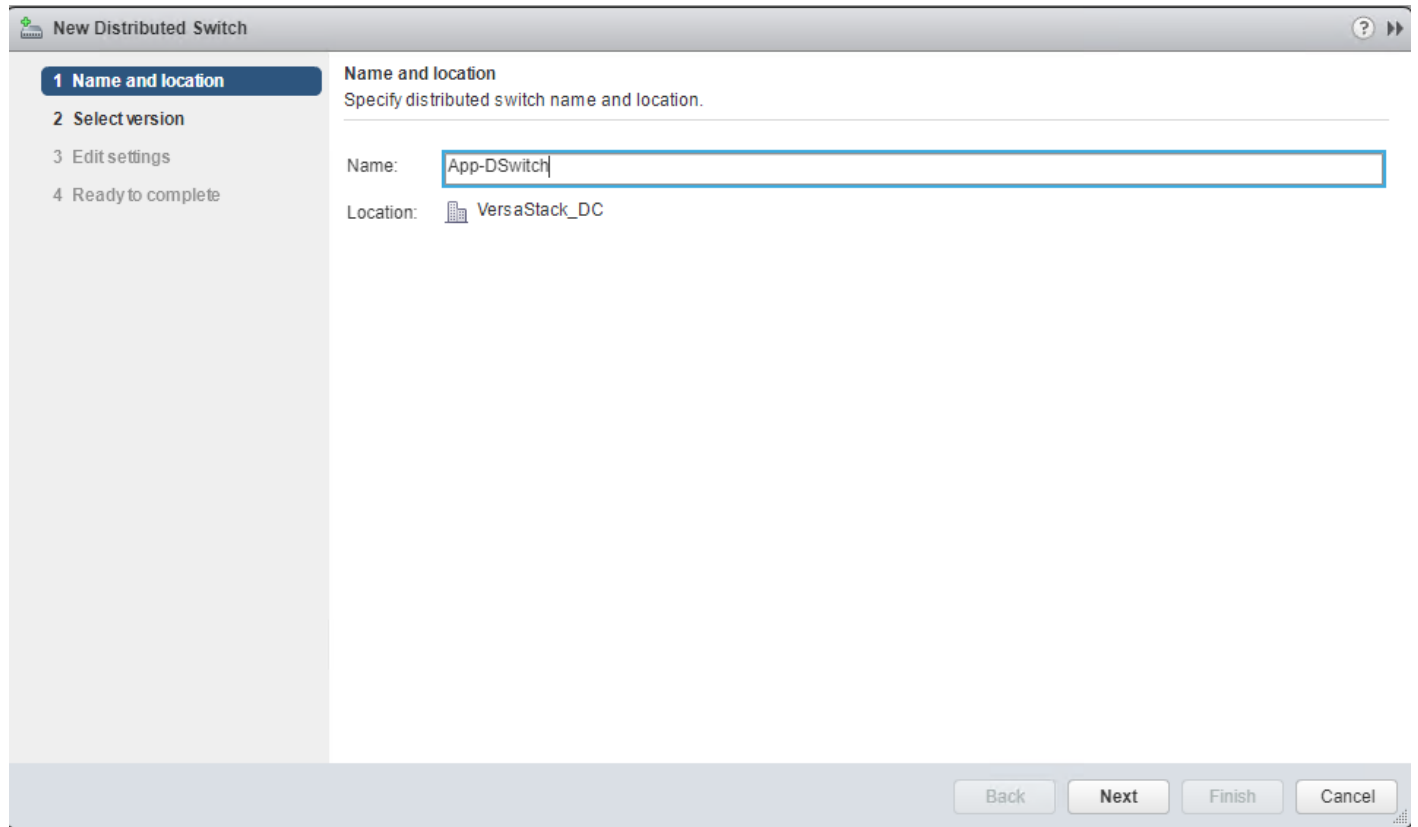
Create a VMware vDS for Application and Production Networks

Production networks will be configured on a VMware vDS to allow additional configuration, as well as consistency between hosts. To configure the VMware vDS, click the right-most icon within the Navigation window, and follow these steps:

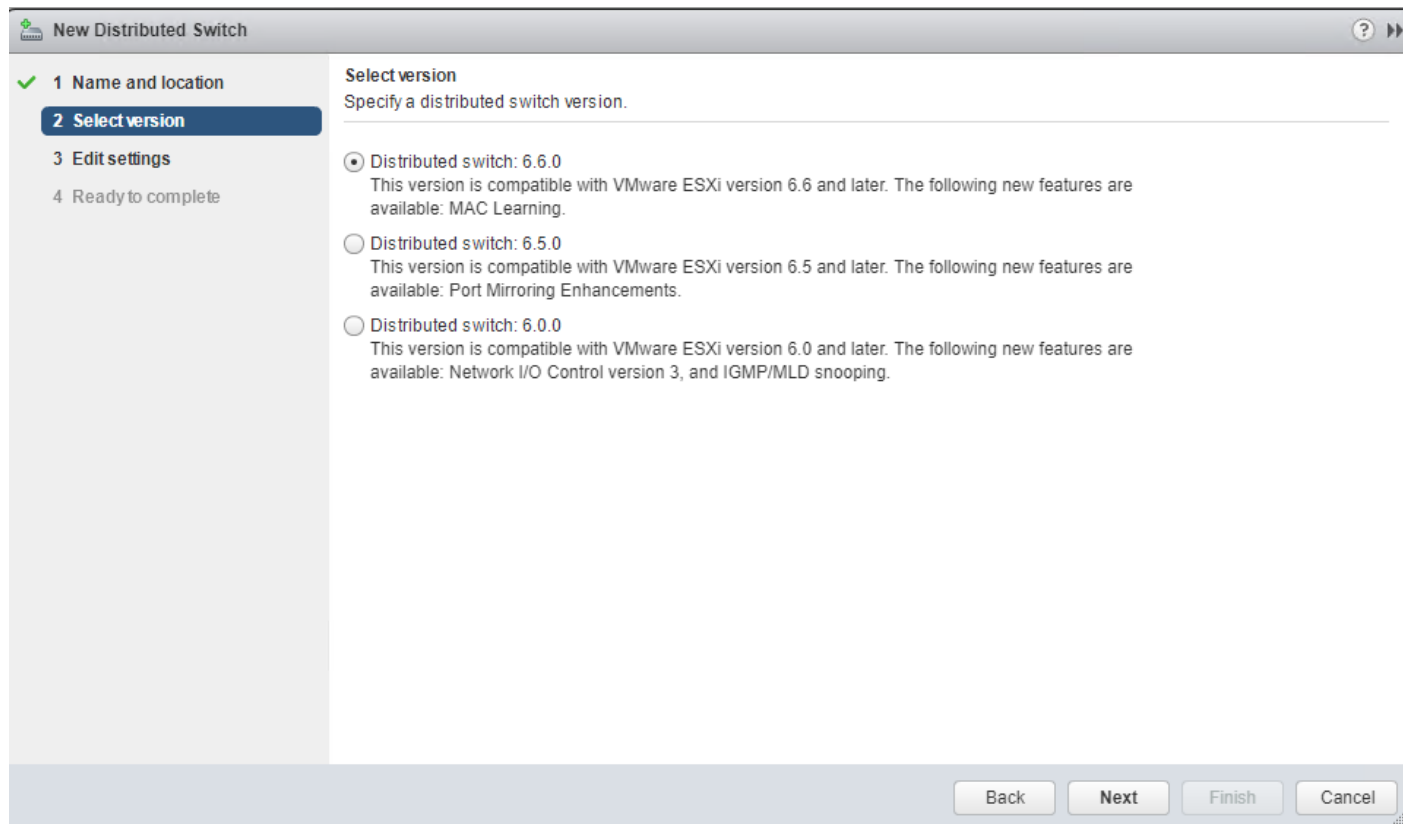
1. Right-click the Datacenter (VersaStack_DC in the example screenshot below), select from the drop-down list **Distributed Switch** > **New Distributed Switch**.




2. Provide a relevant name for the Name field and click **Next**.



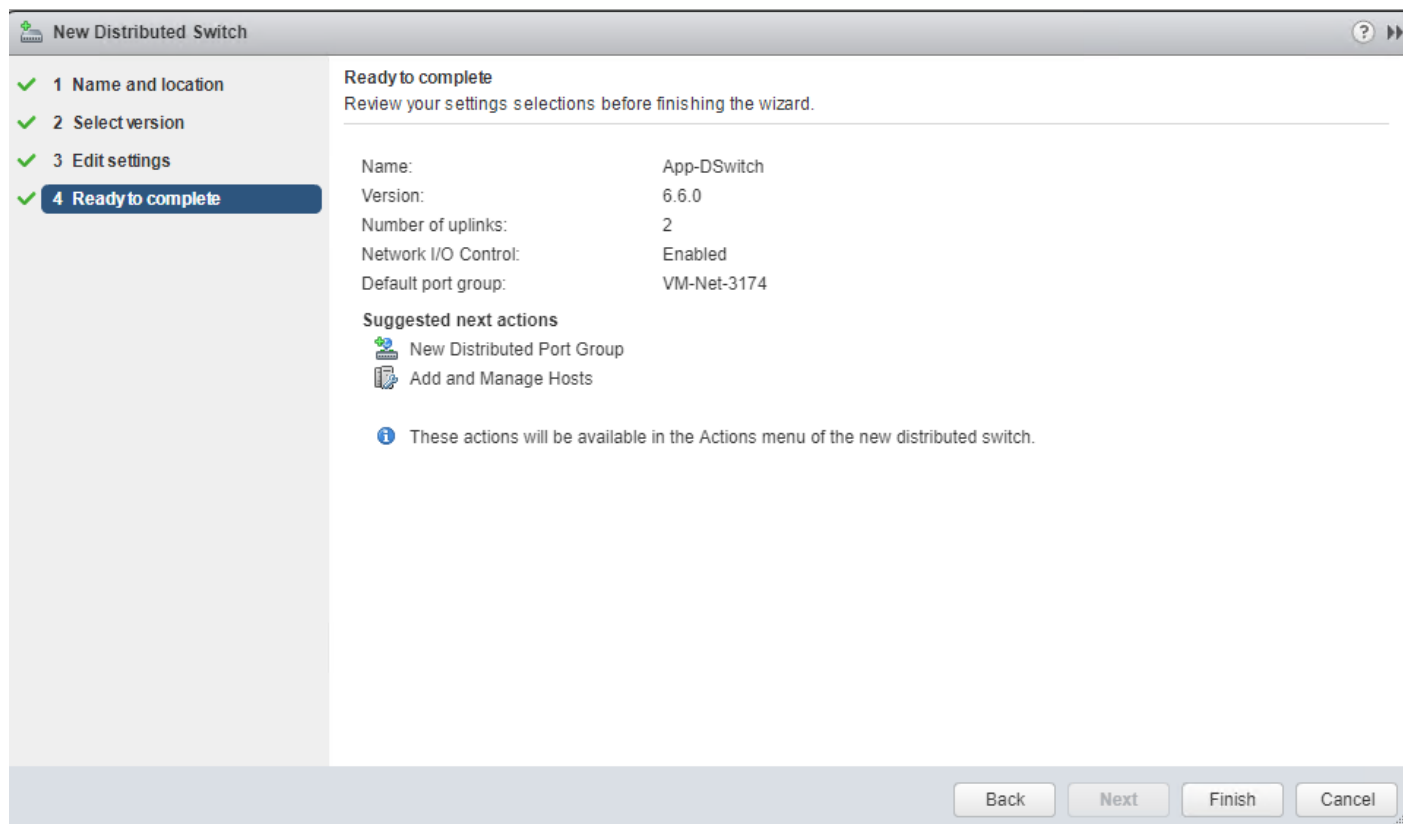
3. Leave the version selected as Distributed switch: 6.6.o and click **Next**.




4. Change the Number of uplinks from 4 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-Net-3174 for the name of the default Port group to be created. Click **Next**.



5. Review the summary in the Ready to complete page and click **Finish** to create the vDS.



6. Right-click the newly created `App-DSwitch` vDS by selecting the Networking sub-tab of Hosts and Clusters of the Navigator window, and select **Settings > Edit Settings...**
7. Click the Advanced option for the Edit Settings window and change the MTU from 1500 to 9000.
8. Click **OK** to save the changes.
9. Right-click the `VM-Net-3174` Distributed Port Group, and select **Edit Settings...**
10. Click **VLAN**, changing VLAN type from None to VLAN, and enter in the appropriate VLAN number for the first application network.



The application Distributed Port Groups will not need to adjust their NIC Teaming as they will be Active/Active within the two vNICs uplinks associated to the App-DSwitch, using the default VMware Route based on originating virtual port load balancing algorithm.

11. Click **OK** to save the changes.
12. Right-Click the `App-DSwitch`, selecting `Distributed Port Group > New Distributed Port Group...` for any additional application networks to be created, setting the appropriate VLAN for each new Distributed Port Group.

Add the ESXi Hosts to the vDS

With the vDS and the distributed port groups created within the vDS in place, the ESXi hosts will be added to the vDS.

To add the ESXi Hosts to the vDS, follow these steps:

1. Within the Networking sub-tab of Hosts and Clusters of the Navigator window, right-click the vDS and select **Add and Manage Hosts...**

What is a Distributed Switch?

A distributed switch acts as a single virtual switch across all associated hosts. This allows virtual machines to maintain consistent network configuration as they migrate across hosts.

Distributed virtual networking configuration consists of three parts. The first part takes place at the datacenter level, where distributed switches are created, and hosts and distributed port groups are added to distributed switches. The second part takes place at the host level, where host ports and networking services are associated with distributed switches either through individual host networking configuration or using host profiles. The third part takes place at the virtual machine level, where virtual machine NICs are connected to distributed port groups either through individual virtual machine NIC configuration or by migrating virtual machine networking from the distributed switch itself.

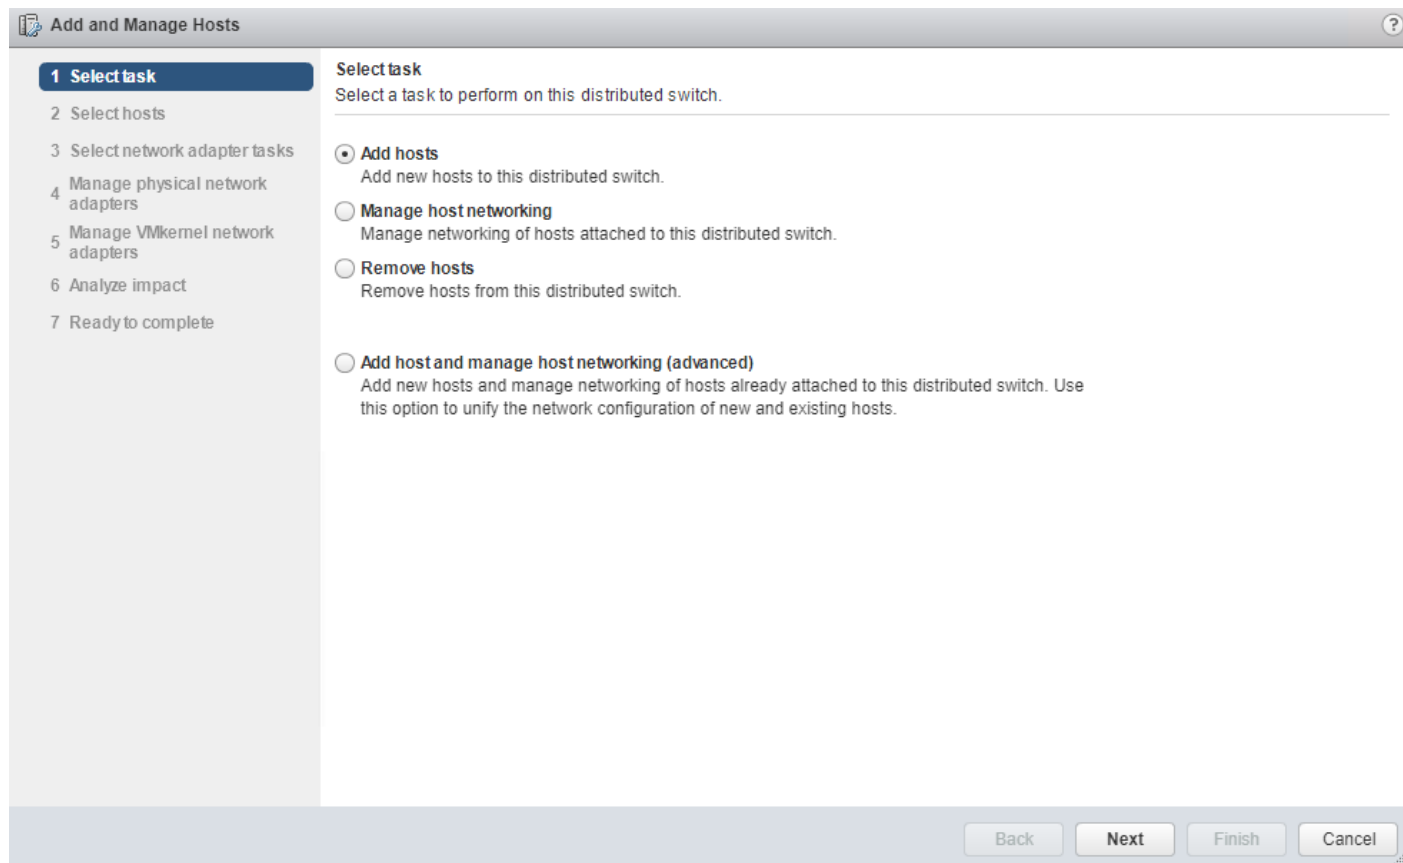
Basic Tasks

- [Add and manage hosts](#)
- [Manage this distributed switch](#)
- [Create a new port group](#)

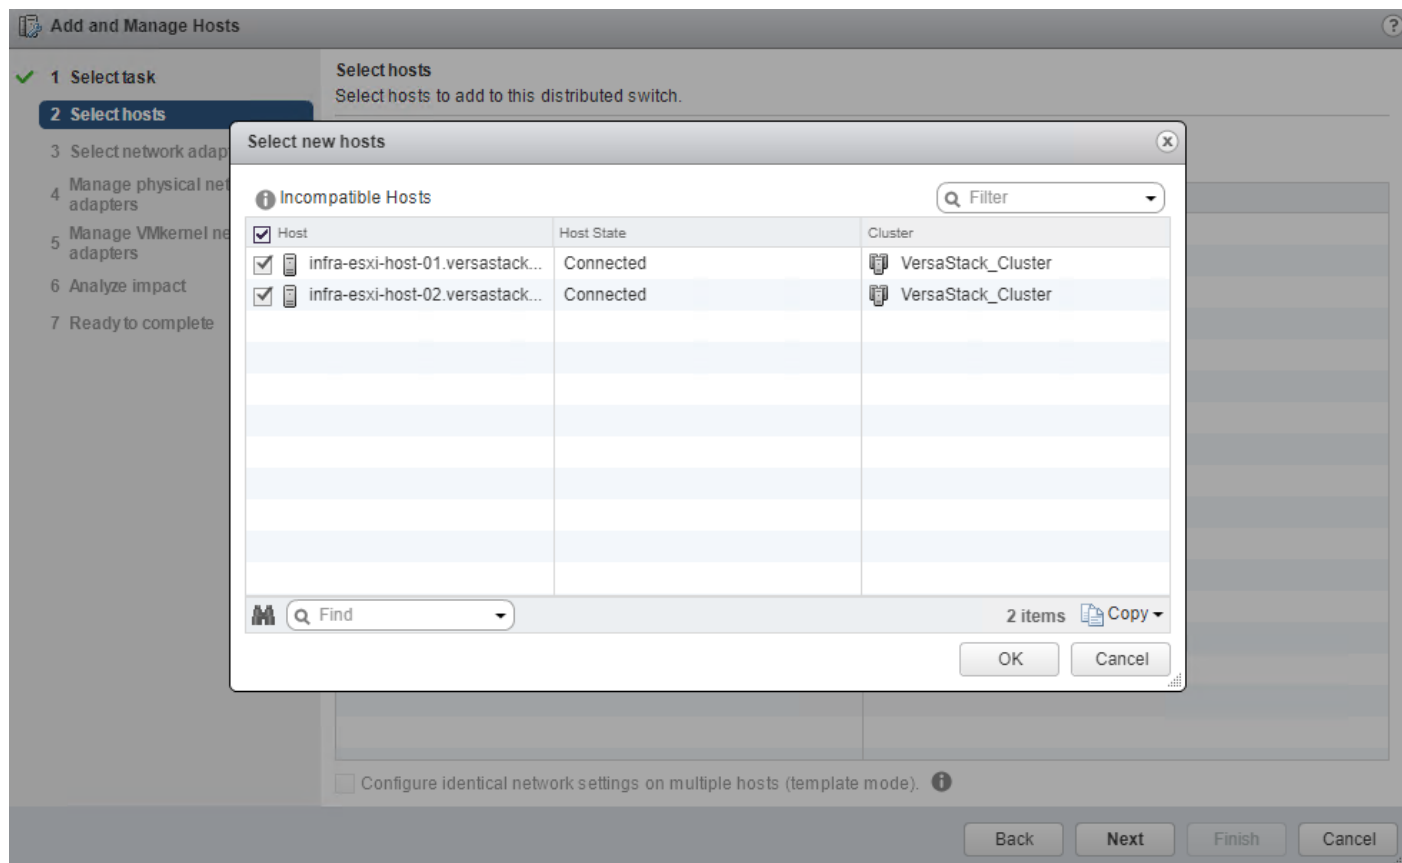
Explore Further

- [Learn more about distributed switches](#)
- [Learn how to set up a network with distributed switch](#)

2. Leave Add hosts selected and click **Next**.



3. Click the green + icon next to New hosts...



5. Click **Next**.
6. At the bottom of the dialog box, select **Configure identical networking settings on multiple hosts** and click **Next**.

Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Select template host**
- 4 Select network adapter tasks
- 5 Manage physical network adapters (template mode)
- 6 Manage VMkernel network adapters (template mode)
- 7 Analyze impact
- 8 Ready to complete

Select template host
 Select a template host to apply its network configuration on this switch to the other hosts.

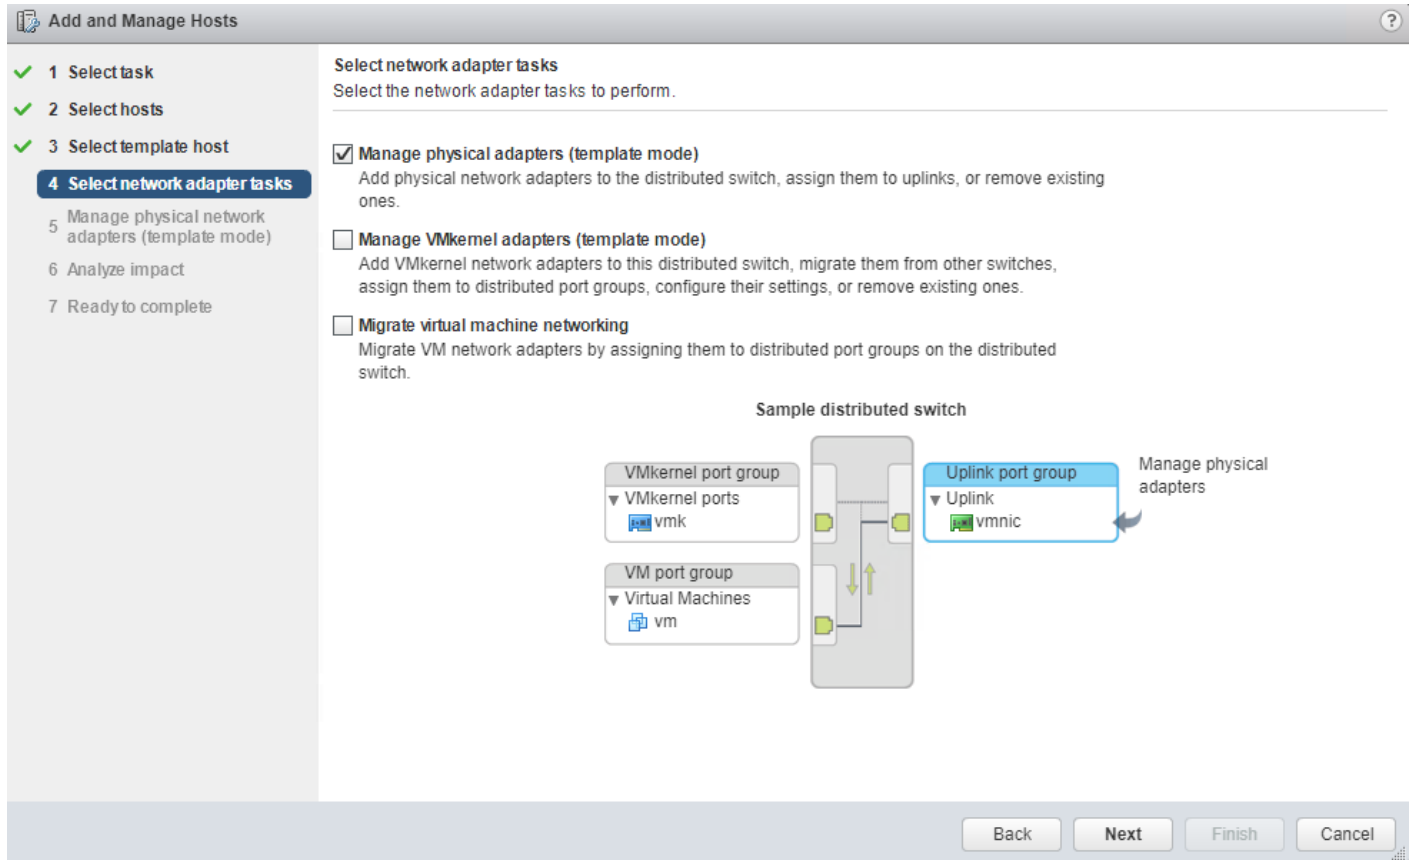
Host	Physical Adapters - On This Switch / All	VMkernel Adapters - On This Switch / All
<input checked="" type="radio"/> infra-esxi-host-01.versa...	0 / 8	0 / 4
<input type="radio"/> infra-esxi-host-02.versast...	0 / 8	0 / 4

Services (infra-esxi-host-01.versastack.local)

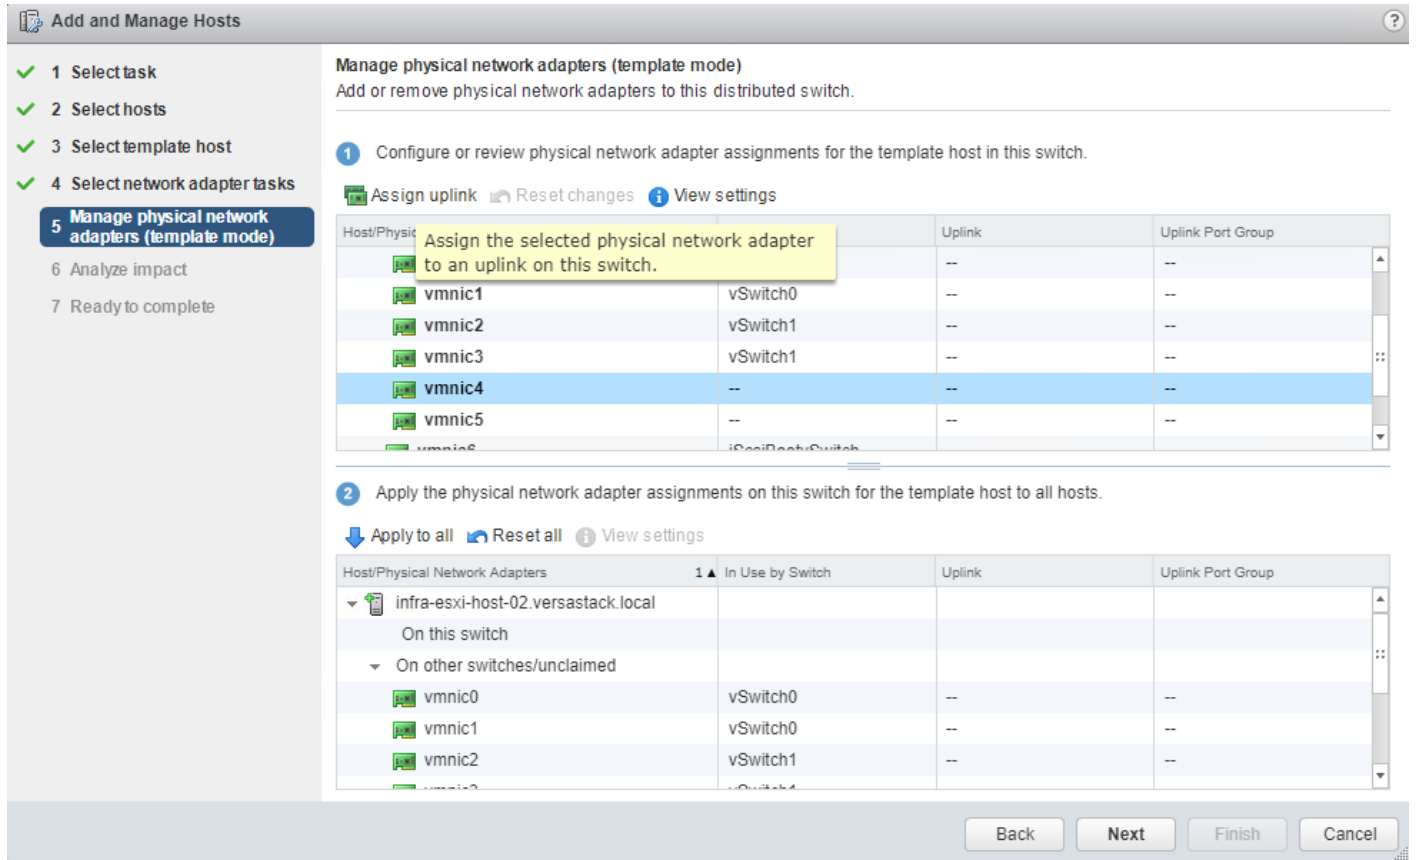
- Fault Tolerance logging: --
- Management: vmk0
- Provisioning: --
- vSphere Replication: --
- vSphere Replication NFC: --
- vMotion: vmk?

Back Next Finish Cancel

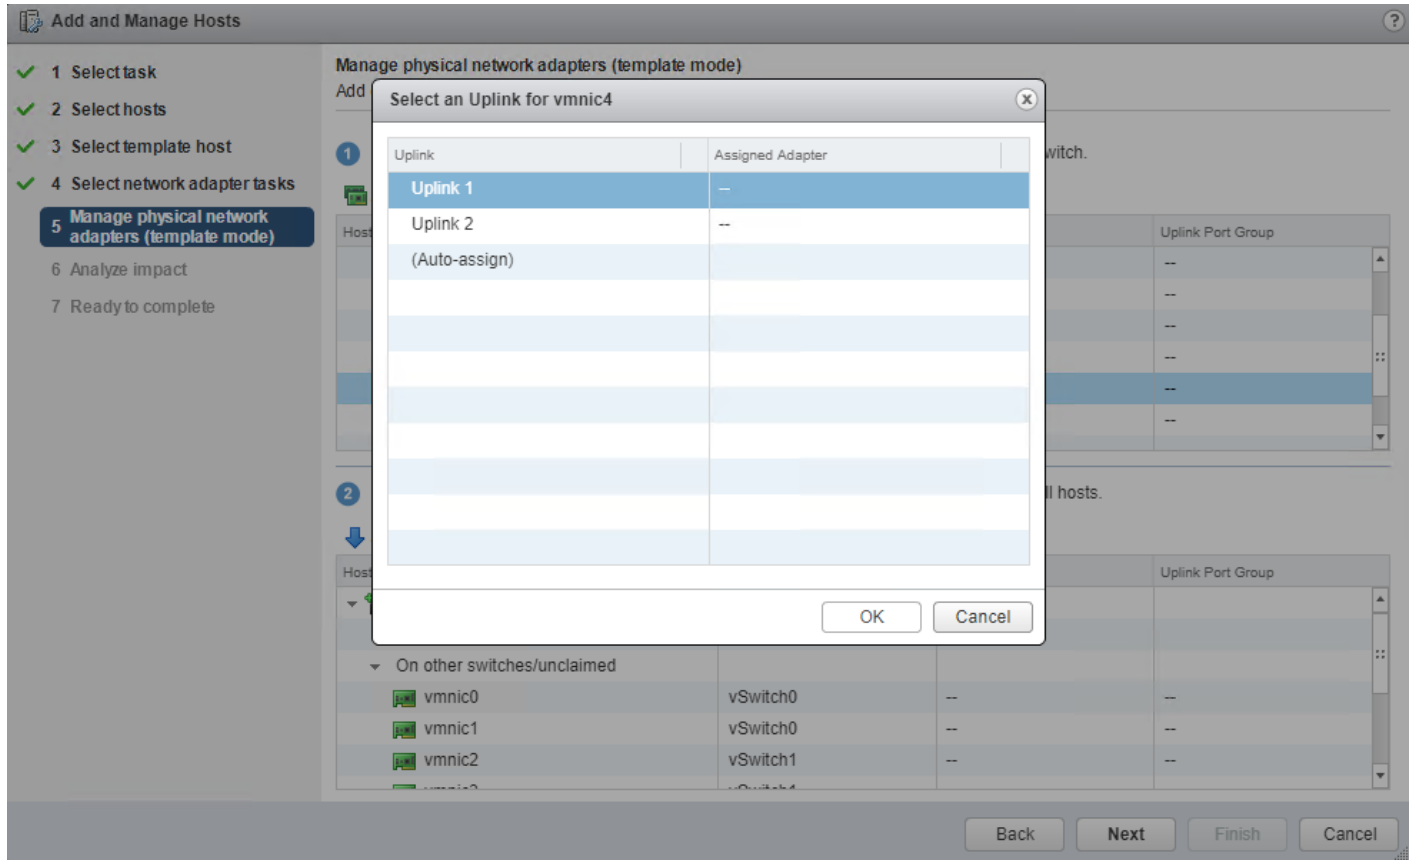
8. Unselect Manage VMkernel adapters (template mode) if it is selected and click **Next**.



9. For each vmnic (vmnic4 and vmnic5) to be assigned from the Host/Physical Network Adapters column, select the vmnic and click the **Assign uplink**.



10. Assign the first to **Uplink 1** and assign the second to **Uplink 2**.



11. Repeat steps 1-10 until all vmnics have been assigned.
12. Click **Next**.
13. Click **Apply to all** to create the same configuration on the other host.

Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- ✓ 3 Select template host
- ✓ 4 Select network adapter tasks
- 5 Manage physical network adapters (template mode)**
- 6 Analyze impact
- 7 Ready to complete

Manage physical network adapters (template mode)
Add or remove physical network adapters to this distributed switch.

1 Configure or review physical network adapter assignments for the template host in this switch.

Assign uplink Reset changes View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	vSwitch1	--	--
vmnic3	vSwitch1	--	--
vmnic6	iScsiBootvSwitch	--	--
vmnic7	vSwitch2	--	--

2 Apply the physical network adapter assignments on this switch for the template host to all hosts.

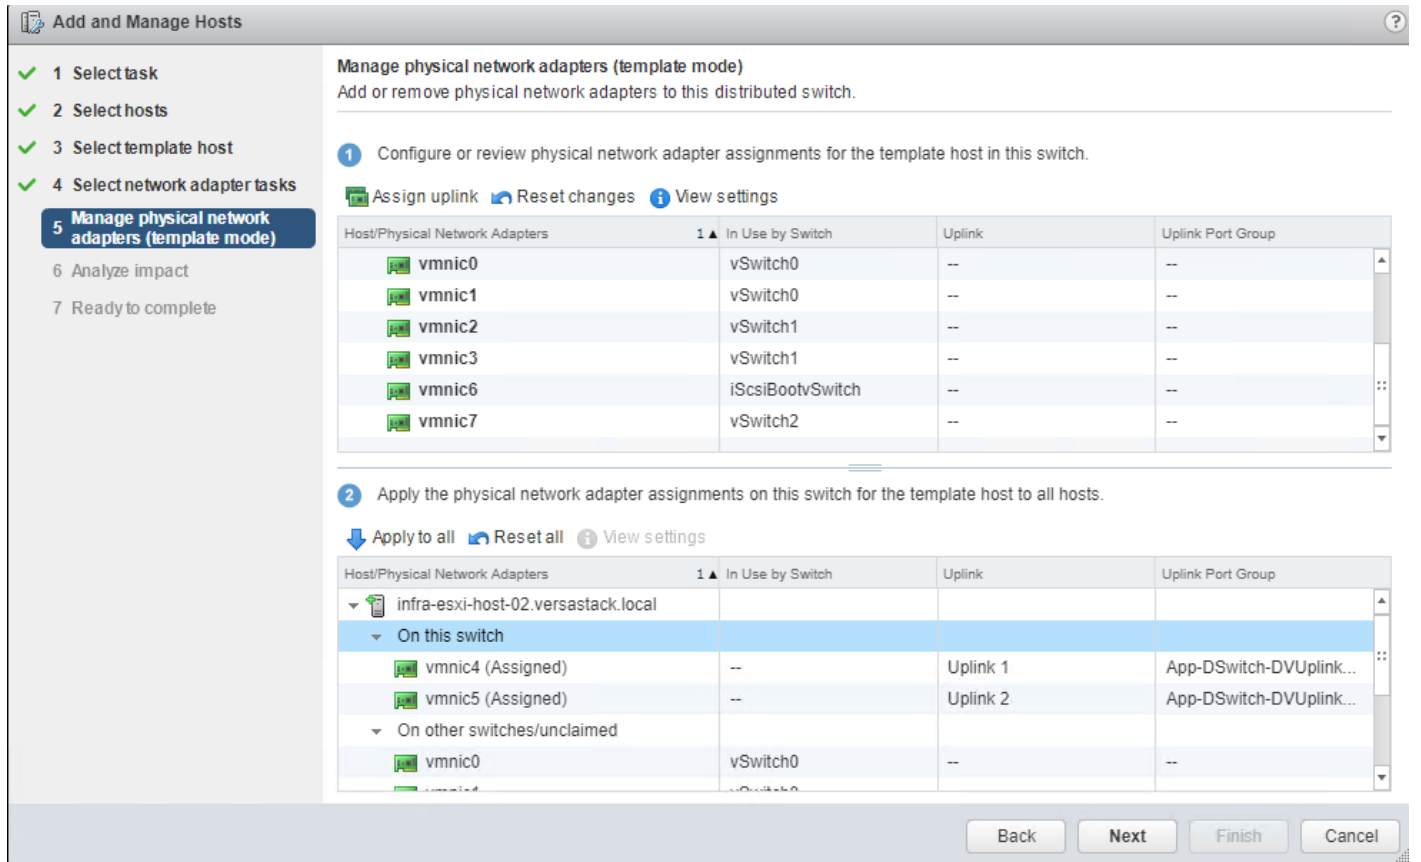
Apply to all Reset all View settings

Host/Physical Network Adapters	In Use by Switch	Uplink	Uplink Port Group
infra-es: On this switch.			
On other switches/unclaimed			
vmnic0	vSwitch0	--	--
vmnic1	vSwitch0	--	--
vmnic2	vSwitch1	--	--

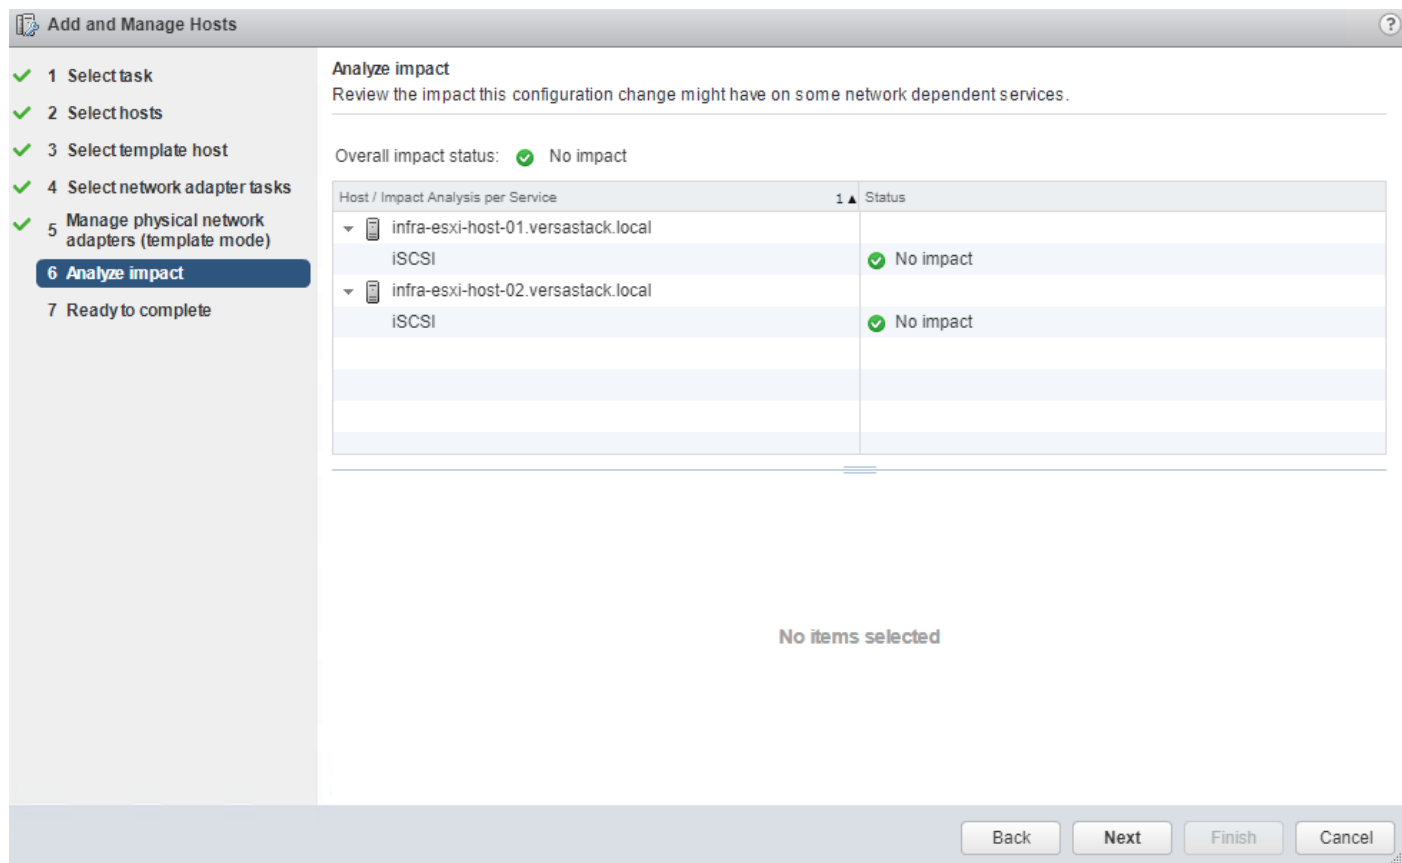
Apply to all hosts the configuration of physical network adapters on the template host for this switch.

Back Next Finish Cancel

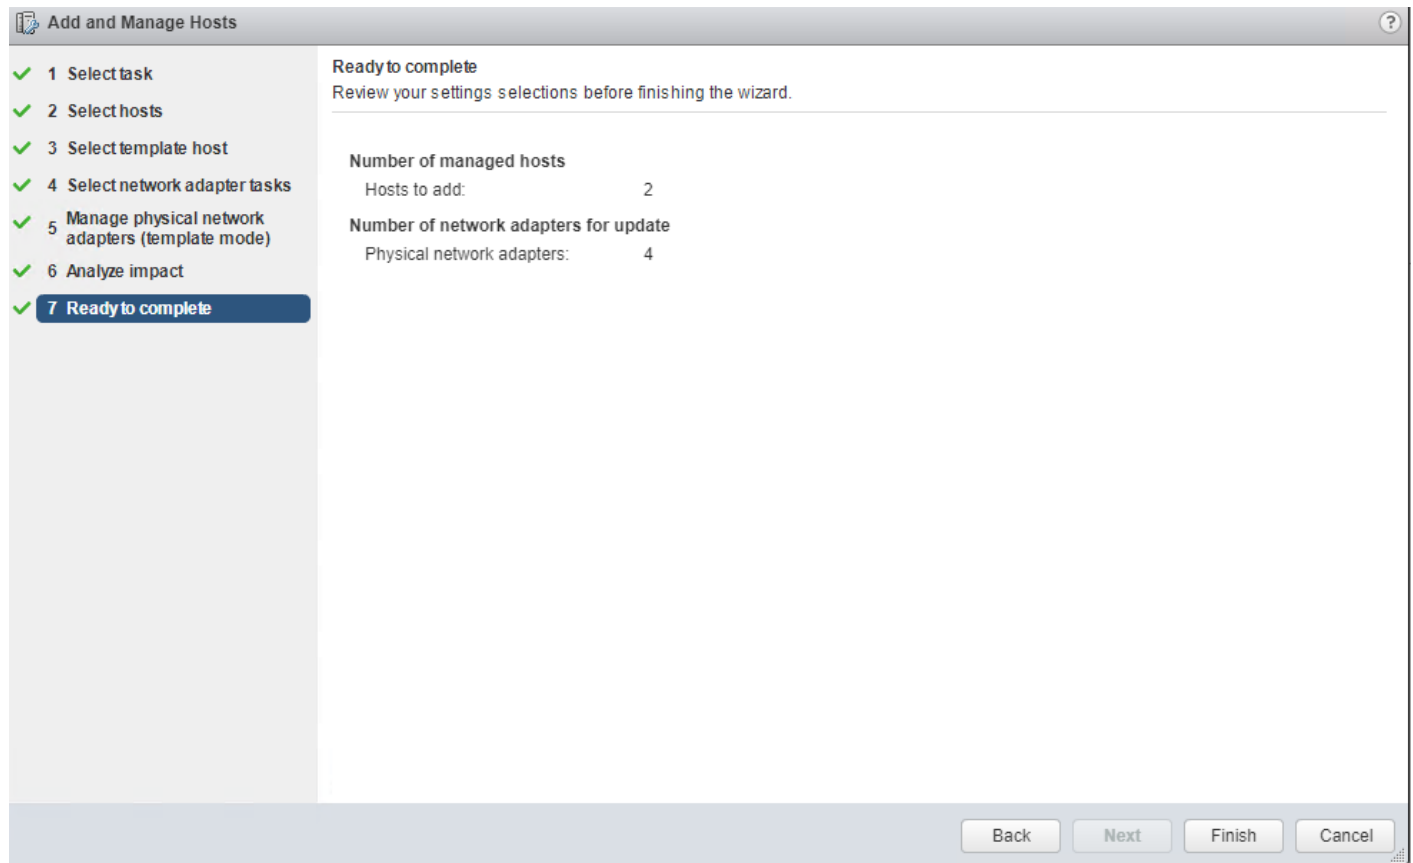
14. Verify the configuration has been applied to the second host and Click **Next**.



15. Proceed past the Analyze impact screen if no issues appear.



16. Review the Ready to complete summary and click **Finish** to add the hosts to the vDS.



References

Products and Solutions

Cisco Unified Computing System:

<http://www.cisco.com/en/US/products/ps10265/index.html>

Cisco UCS 6400 Series Fabric Interconnects:

<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-6400-series-fabric-interconnects/tsdproducts-support-series-home.html>

Cisco UCS 5100 Series Blade Server Chassis:

<http://www.cisco.com/en/US/products/ps10279/index.html>

Cisco UCS B-Series Blade Servers:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

Cisco UCS C-Series Rack Servers:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html>

Cisco UCS Adapters:

http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

<http://www.cisco.com/en/US/products/ps10281/index.html>

Cisco Intersight:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/intersight/index.html>

Cisco Nexus 9000 Series Switches:

<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html>

Cisco Application Centric Infrastructure:

<http://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html>

Cisco Data Center Network Manager:

<https://www.cisco.com/c/en/us/products/cloud-systems-management/prime-data-center-networkmanager/index.html>

Cisco UCS Director:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-director/index.html>

VMware vCenter Server:

<http://www.vmware.com/products/vcenter-server/overview.html>

VMware vSphere:

<https://www.vmware.com/products/vsphere>

IBM FlashSystem 9100:

<https://www.ibm.com/us-en/marketplace/flashsystem-9100>

Interoperability Matrixes

Cisco UCS Hardware Compatibility Matrix:

<https://ucshcltool.cloudapps.cisco.com/public/>

VMware and Cisco Unified Computing System:

<http://www.vmware.com/resources/compatibility>

IBM System Storage Interoperation Center:

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

Appendix

VersaStack Configuration Backups

Cisco UCS Backup

Automated backup of the Cisco UCS domain is important for recovery of the Cisco UCS Domain from issues ranging from 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 and is detailed below.

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 this XML configuration file consisting of All configurations, just System configurations, or just Logical configurations of the Cisco 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 schedule the backup, follow these steps within the Cisco UCS Manager GUI:

1. Select Admin within the Navigation pane and select 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>
 - f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on save>
 - g. Schedule: [Daily/Weekly/Bi Weekly]

General

Policy Backup & Export

Full State Backup Policy

Hostname : 192.168.160.99

Protocol : FTP TFTP SCP SFTP

User : root

Password :

Remote File : /var/www/html/vs/configs/ucs/6454.full

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : 0

Description : Database Backup Policy

All Configuration Backup Policy

Hostname : 192.168.160.99

Protocol : FTP TFTP SCP SFTP

User : root

Password :

Remote File : /var/www/html/vs/configs/ucs/6454.config

Admin State : Disable Enable

Schedule : Daily Weekly Bi Weekly

Max Files : 0

Description : Configuration Export Policy

Backup/Export Config Reminder

Admin State : Disable Enable

4. Click Save Changes to create the Policy.

Cisco Nexus and MDS Backups

The configuration of the Cisco Nexus 9000 and MDS 9000 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 VersaStack 9336C-FX2 switches is shown below:

```
AA10-9336-FX2-A# conf t
Enter configuration commands, one per line. End with CNTL/Z.
AA10-9336-FX2-A(config)# feature scheduler
AA10-9336-FX2-A(config)# scheduler logfile size 1024
AA10-9336-FX2-A(config)# scheduler job name backup-cfg
AA10-9336-FX2-A(config-job)# copy running-config tftp://192.168.160.242/9336/$(SWITCHNAMW)-
cfg.$(TIMESTAMP) vrf management
AA10-9336-FX2-A(config-job)# exit
AA10-9336-FX2-A(config)# scheduler schedule name daily
AA10-9336-FX2-A(config-schedule)# job name backup-cfg
AA10-9336-FX2-A(config-schedule)# time daily 2:00
AA10-9336-FX2-A(config-schedule)# end
```

Show the job that has been setup:

```
AA10-9336-FX2-A# show scheduler job
Job Name: backup-cfg
-----
copy running-config tftp://192.168.160.242/9336/$(SWITCHNAMW)-cfg.$(TIMESTAMP) vrf management
=====
AA10-9336-FX2-A# 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-
```

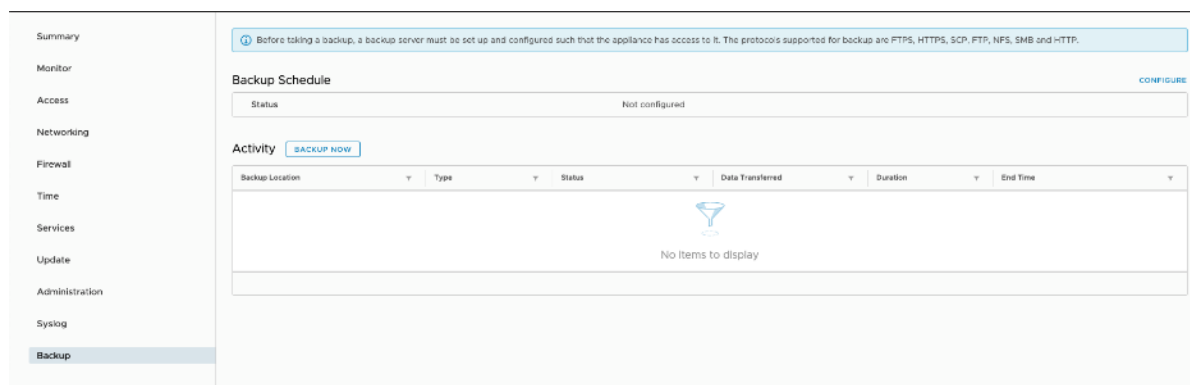
For detailed information about the scheduler, go to:

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 backup of the vCenter Server Appliance is also available within the native capabilities of the VCSA, though within the default solution this is manually initiated for each backup operation. To create a backup, follow these steps:

1. Connect to the VCSA Console at <https://<VCSA IP>:5480>
2. Click Backup In the left side menu.



3. Click Configure to open up the Backup Appliance Dialogue.
4. Fill in all the fields based on your requirement.

Create Backup Schedule

Backup location ⓘ

Backup server credentials

User name

Password

Schedule ⓘ :

Encrypt backup (optional)

Encryption Password

Confirm Password

Number of backups to retain

Retain all backups

Retain last _____ backups

Data

<input checked="" type="checkbox"/> Inventory and configuration	616 MB
<input checked="" type="checkbox"/> Stats, Events, and Tasks	54 MB

5. Review and click **CREATE** to create the backup schedule.

About the Authors

Sreenivasa Edula, Technical Marketing Engineer, UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Sreeni is a Technical Marketing Engineer in the 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.

Warren Hawkins, Virtualization Test Specialist for IBM Spectrum Virtualize, IBM

Working as part of the development organization within IBM Storage, Warren Hawkins is also a speaker and published author detailing best practices for integrating IBM Storage offerings into virtualized infrastructures. Warren has a background in supporting Windows and VMware environments working in second-line and third-line support in both public and private sector organizations. Since joining IBM in 2013, Warren has played a crucial part in customer engagements and, using his field experience, has established himself as the Test Lead for the IBM Spectrum Virtualize™ product family, focusing on clustered host environments.