

FlexPod Datacenter with Microsoft Hyper-V Windows Server 2016

Deployment Guide for FlexPod Datacenter with Microsoft Hyper-V Windows Server 2016

Last Updated: August 28, 2017



About the Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

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

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

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

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

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

© 2017 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary.....	11
Solution Overview	12
Introduction.....	12
Audience	12
Purpose of this Document.....	12
What's New?.....	12
Solution Design	13
Architecture.....	13
Physical Topology.....	14
Deployment Hardware and Software.....	18
Software Revisions	18
Configuration Guidelines	18
Physical Infrastructure	19
FlexPod Cabling	19
Infrastructure Servers Prerequisites	23
Active Directory DC/DNS	23
Microsoft System Center 2016	23
Network Switch Configuration.....	24
Physical Connectivity.....	24
FlexPod Cisco Nexus Base.....	24
Set Up Initial Configuration	24
FlexPod Cisco Nexus Switch Configuration.....	26
Enable Licenses.....	26
Set Global Configurations.....	27
Create VLANs.....	27
Add NTP Distribution Interface.....	28
Add Individual Port Descriptions for Troubleshooting	28
Create Port Channels.....	29
Configure Port Channel Parameters	30
Configure Virtual Port Channels	31
Uplink into Existing Network Infrastructure	33
Storage Configuration	34
NetApp All Flash FAS A300 Controllers	34
NetApp Hardware Universe.....	34
Controllers	34
Disk Shelves.....	34

NetApp ONTAP 9.1.....	34
Complete Configuration Worksheet	34
Configure ONTAP Nodes.....	35
Login to the Cluster.....	44
Zero All Spare Disks	44
Set Onboard Unified Target Adapter 2 Port Personality.....	44
Set Auto-Revert on Cluster Management	45
Set Up Management Broadcast Domain.....	45
Set Up Service Processor Network Interface.....	45
Create Aggregates	46
Verify Storage Failover	46
Disable Flow Control on 10GbE and 40GbE ports.....	47
Disable Unused FCoE Capability on CNA Ports	47
Configure Network Time Protocol.....	48
Configure Simple Network Management Protocol.....	48
Configure AutoSupport	48
Enable Cisco Discovery Protocol.....	49
Create Jumbo Frame MTU Broadcast Domains in ONTAP	49
Create Interface Groups	49
Create VLANs.....	49
Create Storage Virtual Machine.....	49
Create the CIFS Service	50
Modify Storage Virtual Machine Options	50
Create Load-Sharing Mirrors of SVM Root Volume	51
Create Block Protocol FC Service.....	51
Configure HTTPS Access	51
Create SMB Export Policy.....	52
Create NetApp FlexVol Volumes.....	52
Create CIFS Shares	53
Create Boot LUNs	53
Create Witness LUN	53
Schedule Deduplication.....	53
Create FC LIFs	53
Create SMB LIF	54
Add Infrastructure SVM Administrator	54
Server Configuration	55
Cisco UCS Base Configuration.....	55
Perform Initial Setup	55

Cisco UCS Setup.....	57
Log in to Cisco UCS Manager.....	57
Upgrade Cisco UCS Manager Software to Version 3.1(3a)	57
Anonymous Reporting	57
Configure Cisco UCS Call Home.....	58
Configure Unified Ports	58
Add Block of IP Addresses for KVM Access	60
Synchronize Cisco UCS to NTP	60
Edit Chassis Discovery Policy.....	61
Enable Server and Uplink Ports.....	62
Acknowledge Cisco UCS Chassis and FEX.....	63
Create Uplink Port Channels to Cisco Nexus Switches.....	63
Create a WWNN Pool for FC Boot.....	64
Create WWPN Pools.....	66
Create VSAN.....	69
Create FC Uplink Port Channels.....	71
Create vHBA Templates	73
Create SAN Connectivity Policy.....	75
Create MAC Address Pools	77
Create UUID Suffix Pool	79
Create Server Pool	79
Create VLANs.....	80
Modify Default Host Firmware Package.....	83
Set Jumbo Frames in Cisco UCS Fabric	84
Create Local Disk Configuration Policy (Optional).....	85
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	86
Create Power Control Policy.....	87
Create Server Pool Qualification Policy (Optional)	88
Create Server BIOS Policy	89
Update the Default Maintenance Policy.....	92
Create vNIC Templates.....	93
Create LAN Connectivity Policy for FC Boot.....	96
Create FC Boot Policy.....	98
Create Service Profile Templates.....	101
Create Service Profiles	108
Add More Servers to FlexPod Unit.....	109
Gather Necessary Information.....	109
FlexPod Cisco MDS Switch Configuration.....	116

Enable Licenses	116
Configure Individual Ports	116
Create VSANs	120
Create Device Aliases	121
Create Zones.....	121
Storage Configuration – Boot LUNs.....	123
NetApp ONTAP Boot Storage Setup.....	123
Create igroups.....	123
Map Boot LUNs to igroups	123
Microsoft Windows Server 2016 Hyper-V Deployment Procedure	124
Setting Up Microsoft Windows Server 2016.....	124
Install Chipset and Windows eNIC Drivers.....	127
Install Windows Roles and Features	129
Install NetApp Host Utilities	131
Host Renaming and Join to Domain.....	133
Storage Configuration – Boot LUNs (Continued).....	134
NetApp ONTAP Boot Storage Setup.....	134
Deploying and Managing Hyper-V Clusters using System Center 2016 VMM.....	135
Settings	135
Configuring Network Settings	135
Create Run As accounts in VMM	136
Fabric – Servers - I.....	136
Create Host Groups	136
Add Hosts to the Host Group.....	137
Fabric – Networking.....	139
Creating Logical Networks, Sites and IP Pools	139
Create VM Networks	143
Create Uplink Port Profiles and Hyper-V Port Profiles	145
Create Logical Switch using SET	147
Fabric - Storage	151
NetApp SMI-S Provider Configuration	151
NetApp SMI-S Integration with VMM	152
Fabric – Servers - II.....	160
Configure Network on Host – Applying Logical Switch.....	160
Deploy Hyper-V Cluster	162
Cisco UCS Management Pack Suite Installation and Configuration	169
Cisco UCS Manager Integration with SCOM	169
About Cisco UCS Management Pack Suite.....	169

Installing Cisco UCS Monitoring Service.....	169
Adding a Firewall Exception for the Cisco UCS Monitoring Service.....	171
Installing the Cisco UCS Management Pack Suite	171
Adding a Cisco UCS Domains to the Operations Manager.....	174
Cisco UCS Manager Monitoring Dashboards.....	177
Cisco UCS Manager Plug-in for SCVMM.....	182
Cisco UCS Manager Plug-in Installation.....	182
Cisco UCS Domain Registration:.....	183
Using the Cisco UCS SCVMM Plugin.....	185
Viewing the Server Details from the Hypervisor Host View	185
Viewing Registered UCS Domains	186
Viewing the UCS Blade Server Details	187
Viewing the UCS Rack-Mount Server Details:	188
Viewing the Service Profile Details	189
Viewing the Service Profile Template Details.....	191
Viewing the Host Firmware Package Details.....	192
FlexPod Management Tools Setup	193
OnCommand Unified Manager 7.2	193
NetApp SnapDrive 7.1.3	195
Configuring access for SnapDrive for Windows.....	195
Downloading SnapDrive for Windows	196
Installing SnapDrive for Windows	197
NetApp SnapManager for Hyper-V	203
Downloading SnapManager for Hyper-V	203
Installing SnapManager for Hyper-V.....	204
Sample Tenant Provisioning	209
Provisioning a Sample Application Tenant	209
Appendix - ISCSI 10/40GbE Solution	211
Network Switch Configuration.....	213
Physical Connectivity	214
Add NTP Distribution Interface.....	216
Add Individual Port Descriptions for Troubleshooting	216
Create Port Channels.....	218
Configure Port Channel Parameters	219
Configure Virtual Port Channels	220
Uplink into Existing Network Infrastructure	221
Storage Configuration	221
Create Jumbo Frame MTU Broadcast Domains in ONTAP	221

Create VLANs.....	222
Create iSCSI LIFs	222
Server Configuration.....	222
Perform Initial Setup of Cisco UCS 6332-16UP Fabric Interconnects for FlexPod Environments	222
Create Uplink Port Channels to Cisco Nexus Switches.....	222
Create IQN Pools for iSCSI Boot	225
Create IP Pools for iSCSI Boot.....	226
Create UUID Suffix Pool	227
Create Server Pool	228
Create VLANs.....	228
Modify Default Host Firmware Package.....	231
Set Jumbo Frames in Cisco UCS Fabric	232
Create Local Disk Configuration Policy (Optional).....	233
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	234
Create Power Control Policy.....	235
Create Server Pool Qualification Policy (Optional)	236
Create Server BIOS Policy	237
Update the Default Maintenance Policy.....	240
Create vNIC Templates.....	241
Create iSCSI Boot Policy.....	248
Create Service Profile Templates.....	250
Create Service Profiles	261
Add More Servers to FlexPod Unit.....	262
Gather Necessary Information.....	262
Storage Configuration – Boot LUNs and Igroups	263
Create Boot LUNs	263
Create Witness and Data LUN	263
Create igroups.....	263
Map Boot LUNs to igroups	263
Microsoft Windows Server 2016 Hyper-V Deployment Procedure	264
Setup the Microsoft Windows 2016 install	264
Install Windows Server 2016.....	264
Install Chipset and Windows eNIC Drivers.....	264
Install Windows Roles and Features	264
Install NetApp Host Utilities	264
Configure Microsoft iSCSI Initiator	265
Host Renaming and Join to Domain.....	268
Deploying and Managing Hyper-V Clusters using System Center 2016 VMM.....	269

Appendix - FCoE Solution	270
Network Switch Configuration	272
Physical Connectivity	273
FlexPod Cisco Nexus Base	274
FlexPod Cisco Nexus Switch Configuration	274
Storage Configuration	275
Set Onboard Unified Target Adapter 2 Port Personality	275
Create FC LIFs	275
Create Boot LUNs	276
Create Witness and Data LUN	276
Create igroups	276
Map Boot LUNs to igroups	276
Server Configuration	277
Perform Initial Setup of Cisco UCS 6332-16UP and 6248UP Fabric Interconnects for FlexPod Environments	277
Cisco UCS Direct Storage Connect Setup	279
Log in to Cisco UCS Manager	279
Upgrade Cisco UCS Manager Software to Version 3.1(3a)	279
Anonymous Reporting	279
Configure Cisco UCS Call Home	280
Place UCS Fabric Interconnects in Fiber Channel Switching Mode	280
Add Block of IP Addresses for KVM Access	281
Synchronize Cisco UCS to NTP	282
Edit Chassis Discovery Policy	283
Enable Server and Uplink Ports	284
Acknowledge Cisco UCS Chassis and FEX	285
Create Uplink Port Channels to Cisco Nexus Switches	285
Create a WWNN Pool for FC Boot	286
Create WWPN Pools	288
Create Storage VSAN	291
Configure FCoE Storage Port	292
Assign VSANs to FCoE Storage Ports	293
Create vHBA Templates	294
Create SAN Connectivity Policy	296
Create MAC Address Pools	298
Create UUID Suffix Pool	300
Create Server Pool	300
Create VLANs	301
Modify Default Host Firmware Package	304

Set Jumbo Frames in Cisco UCS Fabric	305
Create Local Disk Configuration Policy (Optional).....	306
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	307
Create Power Control Policy.....	308
Create Server Pool Qualification Policy (Optional)	309
Create Server BIOS Policy	310
Update the Default Maintenance Policy.....	313
Create vNIC Templates.....	314
Create LAN Connectivity Policy for FC Boot.....	317
Create Boot Policy (FCoE Boot)	319
Create Service Profile Templates	323
Create Service Profiles	330
Add More Servers to FlexPod Unit.....	331
Gather Necessary Information.....	331
Adding Direct Connected Tenant FC Storage	332
Create Storage Connection Policies	332
Map Storage Connection Policies vHBA Initiator Groups in SAN Connectivity Policy	333
Microsoft Windows Server 2016 Hyper-V Deployment Procedure	334
Setup the Microsoft Windows 2016 install	334
Install Windows Server 2016	334
Install Chipset and Windows eNIC Drivers.....	334
Install Windows Roles and Features	334
Install NetApp Host Utilities	334
Host Renaming and Join to Domain.....	334
Deploying and Managing Hyper-V Clusters using System Center 2016 VMM	334
FlexPod Backups	335
Cisco UCS Backup.....	335
Cisco Nexus Backups.....	336
Breakout Interface Configuration in the Nexus 9332PQ Switches	338
About the Authors	339
Acknowledgements	339



Executive Summary

Cisco Validated Designs include 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 customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

This document describes the Cisco and NetApp® FlexPod Datacenter with Cisco UCS Manager unified software release 3.1(3a) and Microsoft Hyper-V 2016. Cisco UCS Manager (UCSM) 3.1 provides consolidated support for all the current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release 3.1(3a), and Microsoft Hyper-V 2016 is a predesigned, best-practice data center architecture built on Cisco Unified Computing System (UCS), Cisco Nexus® 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp All Flash FAS (AFF).

This document primarily focuses on deploying Microsoft Hyper-V 2016 Cluster on FlexPod Datacenter using Fibre Channel and SMB storage protocols. The Appendix section covers the delta changes on the configuration steps using iSCSI and FCoE storage protocol for the same deployment model.

Solution Overview

Introduction

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed storage, server and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

Audience

The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step by step configuration and implementation guidelines for the FlexPod Datacenter with Cisco UCS Fabric Interconnects, NetApp AFF, and Cisco Nexus 9000 solution. This document primarily focuses on deploying Microsoft Hyper-V 2016 Cluster on FlexPod Datacenter using Fibre Channel and SMB storage protocols. The Appendix section covers the delta changes on the configuration steps using iSCSI and FCoE storage protocol for the same deployment model.

What's New?

The following design elements distinguish this version of FlexPod from previous FlexPod models:

- Support for the Cisco UCS 3.1(3a) unified software release, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for the latest release of NetApp ONTAP® 9.1
- SMB, Fibre channel, FCoE and iSCSI storage design
- Validation of Microsoft Hyper-V 2016

Solution Design

Architecture


FlexPod architecture is highly modular, or pod-like. Although each customer's FlexPod unit might vary in its exact configuration, after a FlexPod unit is built, it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a FlexPod unit) and scaling out (adding additional FlexPod units). Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for all virtualization solutions. FlexPod validated with Microsoft Hyper-V 2016 includes NetApp All Flash FAS storage, Cisco Nexus® networking, Cisco Unified Computing System (Cisco UCS®), Microsoft System Center Operation Manager and Microsoft Virtual Machine Manager in a single package. The design is flexible enough that the networking, computing, and storage can fit in a single data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

The reference architectures detailed in this document highlight the resiliency, cost benefit, and ease of deployment across multiple storage protocols. A storage system capable of serving multiple protocols across a single interface allows for the customer choice and investment protection because it truly is a wire-once architecture.

Figure 1 shows the Microsoft Hyper-V built on FlexPod components and its physical cabling with the Cisco UCS 6332-16UP Fabric Interconnects. This design has end-to-end 40 Gb Ethernet connections from Cisco UCS 5108 Blade Chassis, Cisco UCS C-Series rackmount servers, a pair of Cisco UCS Fabric Interconnects, Cisco Nexus 9000 switches, through NetApp AFF A300. This infrastructure option can be expanded by introducing a pair of Cisco MDS switches between the UCS Fabric Interconnects and the NetApp AFF A300 to provide the FC-booted hosts with a file-level shared storage access. The reference architecture reinforces the "wire-once" strategy, because the additional storage can be introduced into the existing architecture without a need for re-cabling from the hosts to the Cisco UCS Fabric Interconnects.

Physical Topology

Figure 1 FlexPod with Cisco UCS 6332-16UP Fabric Interconnects




The reference 40Gb based hardware configuration includes:

- Two Cisco Nexus 9332PQ switches

- Two Cisco UCS 6332-16UP fabric interconnects
- Two Cisco MDS 9148S multilayer fabric switches
- One chassis of Cisco UCS blade servers
- Two Cisco UCS C220M4 rack servers
- One NetApp AFF A300 (HA pair) running ONTAP with disk shelves and solid state drives (SSD)

Figure 2 shows the Microsoft Hyper-V built on FlexPod components and its physical connections with the Cisco UCS 6248UP Fabric Interconnects. This design is identical to the 6332-16UP based topology, but has 10Gb Ethernet connecting through a pair of Cisco Nexus 93180YC-EX switches to access SMB file share to the AFF A300. Alternately, the same Nexus 9332PQ switch can be used as the UCS 6332-16UP with a QSFP breakout cable and port configuration setting on the 9332PQ switch.

Figure 2 FlexPod with Cisco UCS 6248UP Fabric Interconnects



The reference hardware configuration includes:

- Two Cisco Nexus 93180YC-EX switches
- Two Cisco UCS 6248UP fabric interconnects
- Two Cisco MDS 9148S multilayer fabric switches
- One NetApp AFF A300 (HA pair) running ONTAP with Disk shelves and Solid State Drives (SSD)

All systems and fabric links feature redundancy and provide end-to-end high availability. For server virtualization, the deployment includes Microsoft Hyper-V 2016. Although this is the base design, each of the components can be scaled

flexibly to support specific business requirements. For example, more (or different) blades and chassis could be deployed to increase compute capacity, additional disk shelves could be deployed to improve I/O capacity and throughput, or special hardware or software features could be added to introduce new features.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	<ul style="list-style-type: none"> Cisco UCS Fabric Interconnects 6200 and 6300 Series. UCS B-200 M4, UCS C-220 M4 	<ul style="list-style-type: none"> 3.1(3a) - Infrastructure Bundle 3.1(2f) – Server Bundle 	Includes the Cisco UCS-IOM 2304 Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
Network	Cisco Nexus 9000 NX-OS	7.0(3)I4(5)	
Storage	NetApp AFF A300	ONTAP 9.1	
	Cisco MDS 9148S	7.3(0)D1(1)	
Software	Cisco UCS Manager	3.1(3a)	
	Microsoft System Center Virtual Machine Manager	2016 (version: 4.0.2051.0)	
	Microsoft Hyper-V	2016	
	Microsoft System Center Operation Manager	2016 (version: 7.2.11878.0)	

Configuration Guidelines

This document provides details on configuring a fully redundant, highly available reference model for a FlexPod unit with NetApp ONTAP storage. Therefore, reference is made to the component being configured with each step, as either o1 or o2 or A and B. In this CVD we have used nodeo1 and nodeo2 to identify the two NetApp storage controllers provisioned in this deployment model. Similarly, Cisco Nexus A and Cisco Nexus B refer to the pair of Cisco Nexus switches configured. Likewise the Cisco UCS Fabric Interconnects are also configured in the same way. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these examples are identified as: Hyper-V-Host-01, Hyper-V-Host-02 to represent infrastructure hosts deployed to each of the fabric interconnects in this document. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the network port vlan create command:

Usage:

```
network port vlan create ?
```

```
[-node] <nodename>      Node
```

```
{ [-vlan-name] }{<netport>|<ifgrp>} VLAN Name
```

| -port {<netport>|<ifgrp>} Associated Network Port
 [-vlan-id] <integer> } Network Switch VLAN Identifier

Example:

network port vlan -node <node01> -vlan-name ioa-<vlan id>

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

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out-of-Band-Mgmt	VLAN for out-of-band management interfaces	13
MS-IB-MGMT	VLAN for in-band management interfaces	904
Native-VLAN	VLAN to which untagged frames are assigned	2
MS-SMB-1-VLAN & MS-SMB-2-VLAN	VLAN for SMB traffic	3052/3053
MS-LVMN-VLAN	VLAN designated for the movement of VMs from one physical host to another.	906
MS-Cluster-VLAN	VLAN for cluster connectivity	907
MS-Tenant-VM-VLAN	VLAN for Production VM Interfaces	908

Table 3 lists the VMs necessary for deployment as outlined in this document.

Table 3 Virtual Machines

Virtual Machine Description	Host Name
Active Directory (AD)	MS-AD
Microsoft System Center Virtual Machine Manager	MS-SCVMM
Microsoft System Center Operation Manager	MS-SCOM

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain details for the prescribed and supported configuration of the NetApp AFF A300 running NetApp ONTAP® 9.1.



For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool](#) (IMT). Cisco HyperFlex documents need Cisco.com login credentials. Please login to access these documents.

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps. Make sure to use the cabling directions in this section as a guide.

The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 3 details the cable connections used in the validation lab for the 40Gb end-to-end with Fibre Channel topology based on the Cisco UCS 6332-16UP Fabric Interconnect. Two 16Gb uplinks connect as port-channels to each of the FIs from the Cisco MDS switches, and a total of four 16Gb links connect to the NetApp AFF controllers from the MDS switches. An additional 1Gb management connection is required for an out-of-band network switch apart from the FlexPod infrastructure. Cisco UCS fabric interconnects and Cisco Nexus switches are connected to the out-of-band network switch, and each NetApp AFF controller has two connections to the out-of-band network switch.

Figure 3 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect

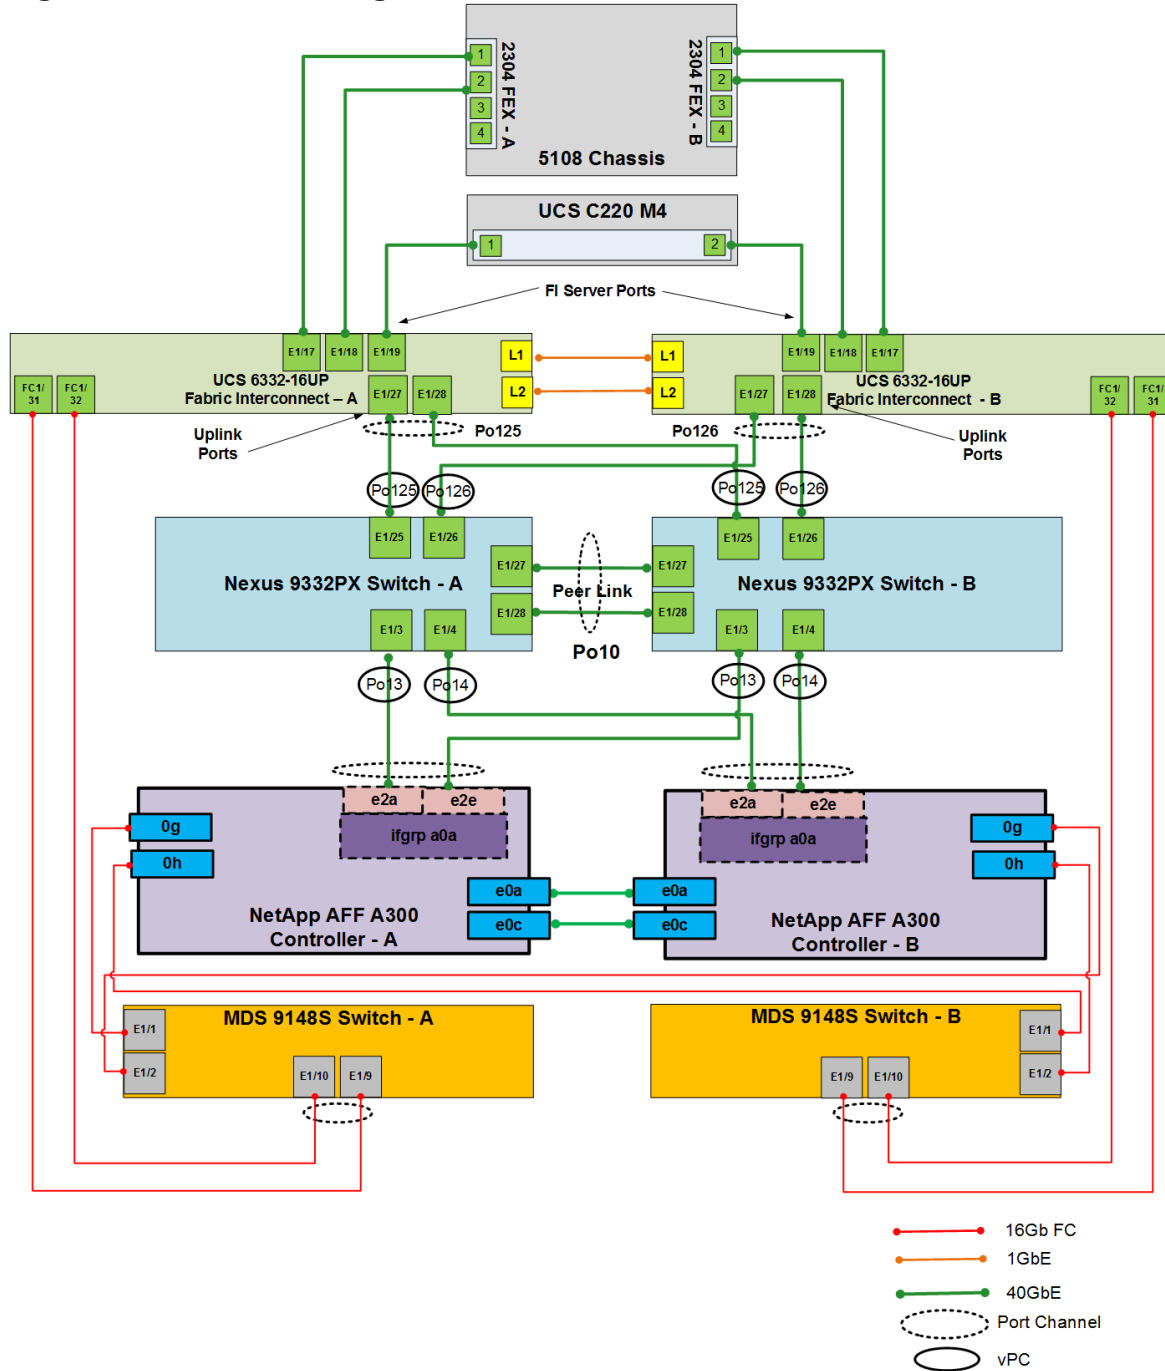
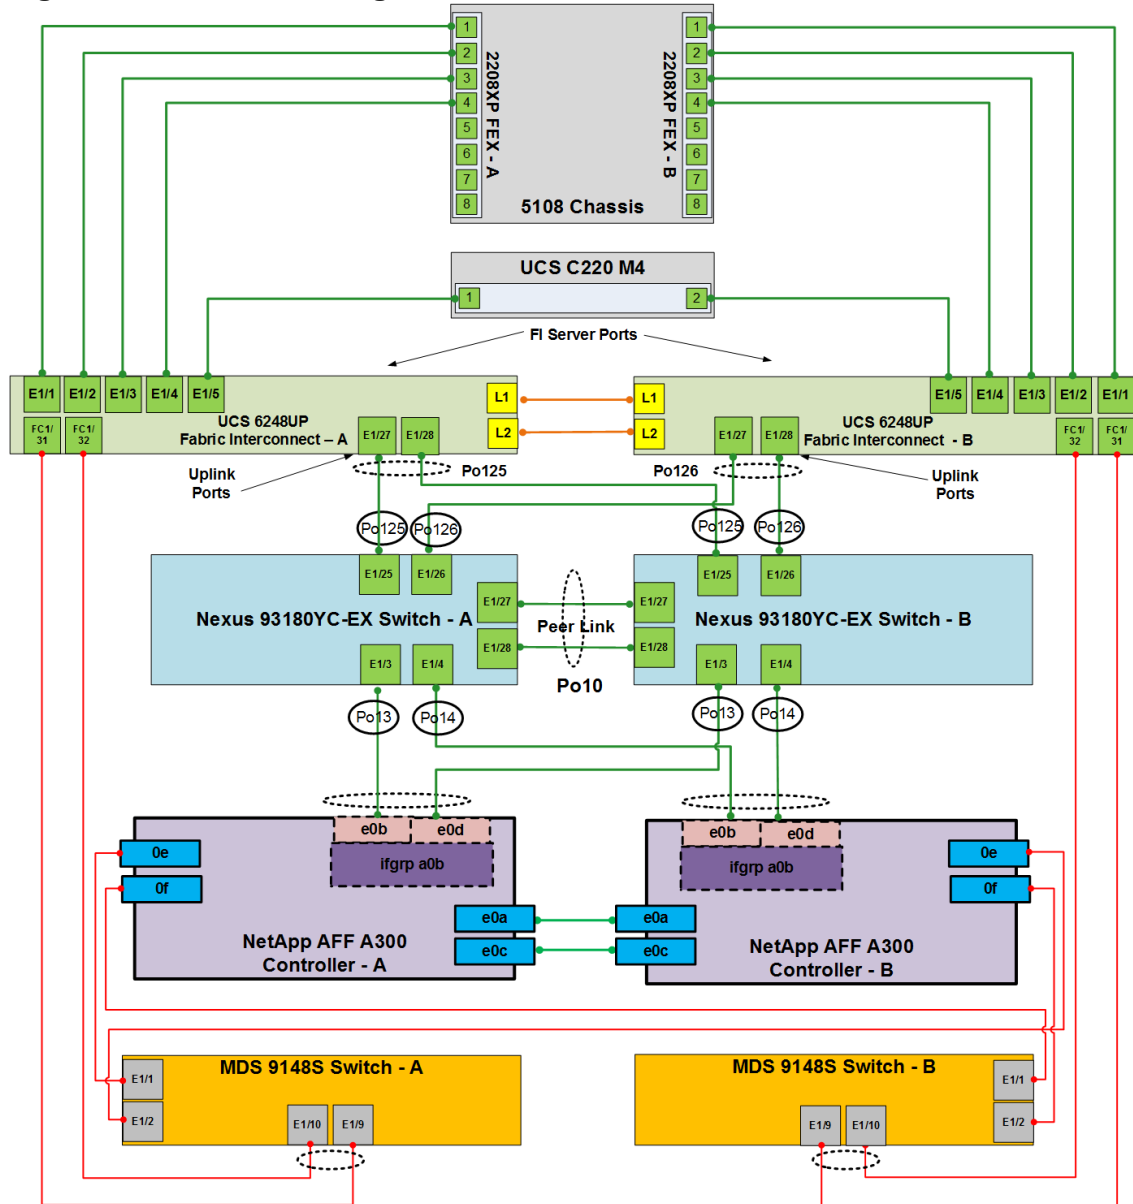


Figure 4 details the cabling connections used in the alternate 10Gb end-to-end topology based on the Cisco UCS 6248UP Fabric Interconnect using Cisco MDS switches for 8Gb Fibre Channel links. As with the 40Gb topology, an out-of-band connection is also required. Cisco UCS fabric interconnect and Cisco Nexus connects to the out-of-band network switch, and each NetApp AFF controller will have two connections to the out-of-band network switch.

Figure 4 FlexPod Cabling with Cisco UCS 6248UP Fabric Interconnect



- 8Gb FC
- 1GbE
- 10GbE
- Port Channel
- vPC

Infrastructure Servers Prerequisites

Active Directory DC/DNS

Production environments at most customer's location might have an active directory and DNS infrastructure configured; the FlexPod with Microsoft Windows Server 2016 Hyper-V deployment model does not require an additional domain controller to be setup. The optional domain controllers is omitted from the configuration in this case or used as a resource domain. In this document we have used an existing AD domain controller and an AD integrated DNS server role running on the same server, which is available in our lab environment.

Microsoft System Center 2016

This document does not cover the steps to install Microsoft System Center Operations Manager (SCOM) and Virtual Machine Manager (SCVMM). Follow Microsoft guidelines to install SCOM and SCVMM 2016:

- SCOM: <https://docs.microsoft.com/en-us/system-center/scom/deploy-overview>
- SCVMM: <https://docs.microsoft.com/en-us/system-center/vmm/install-console>

Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in the section "FlexPod Cabling".

FlexPod Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I4(5), and is valid for both the Cisco Nexus 9332PQ switches deployed with the 40Gb end-to-end topology, and the Cisco Nexus 9318oYC-EX switches used in the 10Gb based topology.



The following procedure includes the setup of NTP distribution on 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.

Set Up Initial Configuration

Cisco Nexus A

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

1. Configure the switch.



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

Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes

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 (mgmt) management configuration? (yes/no) [y]: Enter

Mgmt IPv4 address: <nexus-A-mgmt-ip>

Mgmt IPv4 netmask: <nexus-A-mgmt-netmask>

- Configure the default gateway? (yes/no) [y]: Enter
 - IPv4 address of the default gateway: <nexus-A-mgmt-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) [L3]: L2
 - Configure default switchport interface state (shut/noshut) [shut]: Enter
 - 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.
- Use this configuration and save it? (yes/no) [y]: Enter

Cisco Nexus B

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

1. Configure the switch.



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

- Abort Power on Auto Provisioning and continue with normal setup? (yes/no) [n]: yes
- 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 (mgmt) management configuration? (yes/no) [y]: Enter

Mgmtmto IPv4 address: <nexus-B-mgmtmto-ip>

Mgmtmto IPv4 netmask: <nexus-B-mgmtmto-netmask>

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

IPv4 address of the default gateway: <nexus-B-mgmtmto-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) [L3]: L2

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

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.

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

FlexPod Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus A and Cisco Nexus B

To license the Cisco Nexus switches, complete the following steps:

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

```
config t
```

```
feature interface-vlan
```

```
feature lacp
```

```
feature vpc
```

```
feature lldp
```

```
feature nxapi
```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

To set global configurations, complete the following step on both switches.

Run the following commands to set global configurations:

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

Create VLANs

Cisco Nexus A and Cisco Nexus B

To create the necessary virtual local area networks (VLANs), complete the following step on both the switches.

From the global configuration mode, run the following commands:

```
vlan <ms-ib-mgmt-vlan-id>
name MS-IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan < ms-lvmn-vlan-id>
name MS-LVMN-VLAN
vlan <ms-tenant-vm-vlan-id>
name MS-Tenant-VM-VLAN
vlan <ms-cluster-vlan-id>
name MS-Cluster-VLAN
vlan <ms-SMB-1-vlan-id>
name MS-SMB-1-VLAN
vlan <ms-SMB-2-vlan-id>
name MS-SMB-2-VLAN
exit
```

Add NTP Distribution Interface

Cisco Nexus A

From the global configuration mode, run the following commands:

```
ntp source <switch-a-ntp-ip>
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
```

Cisco Nexus B

From the global configuration mode, run the following commands:

```
ntp source <switch-b-ntp-ip>
interface Vlan<ib-mgmt-vlan-id>
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>
no shutdown
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, complete the following step:



In this step and in the later sections, configure the AFF nodename <st-node> and Cisco UCS 6332-16UP or UCS 6248UP fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/3
description <st-node>-1:e2a
interface Eth1/4
description <st-node>-2:e2a
interface Eth1/25
description <ucs-clustername>-a:1/27
interface Eth1/26
description <ucs-clustername>-b:1/27
interface Eth1/27
```

```
description <nexus-hostname>-b:1/27
interface Eth1/28
description <nexus-hostname>-b:1/28
exit
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B, complete the following step:

From the global configuration mode, run the following commands:

```
interface Eth1/3
description <st-node>-1:e2e
interface Eth1/4
description <st-node>-2:e2e
interface Eth1/25
description <ucs-clustname>-a:1/28
interface Eth1/26
description <ucs-clustname>-b:1/28
interface Eth1/27
description <nexus-hostname>-a:1/27
interface Eth1/28
description <nexus-hostname>-a:1/28
exit
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

To create necessary port channels between the devices, complete the following step on both the switches.

From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/27-28
channel-group 10 mode active
no shutdown
interface Po13
description <st-node>-1
```

```
interface Eth1/3
channel-group 13 mode active
no shutdown
interface Po14
description <st-node>-2
interface Eth1/4
channel-group 14 mode active
no shutdown
interface Po125
description <ucs-clustername>-a
interface Eth1/25
channel-group 125 mode active
no shutdown
interface Po126
description <ucs-clustername>-b
interface Eth1/26
channel-group 126 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, complete the following step on both the switches.

From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-SMB-1-id>, <infra-SMB-2-id>, <LiveMigration-vlan-id>,
<vm-traffic-vlan-id>, <infra-ClusterComm-id>,
spanning-tree port type network
interface Po13
```

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-SMB-1-id>, <infra-SMB-2-id>
spanning-tree port type edge trunk
mtu 9216
interface Po14
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-SMB-1-id>, <infra-SMB-2-id>
spanning-tree port type edge trunk
mtu 9216
interface Po125
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-SMB-1-id>, <infra-SMB-2-id>, <LiveMigration-vlan-id>,
<vm-traffic-vlan-id>, <infra-ClusterComm-id>
spanning-tree port type edge trunk
mtu 9216
interface Po126
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <infra-SMB-1-id>, <infra-SMB-2-id>, <LiveMigration-vlan-id>,
<vm-traffic-vlan-id>, <infra-ClusterComm-id>
spanning-tree port type edge trunk
mtu 9216
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, complete the following step.

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt-ip> source <nexus-A-mgmt-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po13
vpc 13
interface Po14
vpc 14
interface Po125
vpc 125
interface Po126
vpc 126
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, complete the following step.

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt-ip> source <nexus-B-mgmt-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
```



```
interface Po13  
vpc 13  
interface Po14  
vpc 14  
interface Po125  
vpc 125  
interface Po126  
vpc 126  
exit  
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to provide uplink connectivity to the FlexPod environment. If a Cisco Nexus environment is present, we recommend using vPCs with the Cisco Nexus switches included in the FlexPod environment. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after completing the configuration.

Storage Configuration

NetApp All Flash FAS A300 Controllers



Pursuant to best practices, NetApp recommends the following command on the LOADER prompt of the NetApp controllers to assist with LUN stability during copy operations. To access the LOADER prompt, connect to the controller via serial console port or Service Processor connection and press Ctrl-C to halt the boot process when prompted: `setenv boot-arg.tmgr.disable_pit_hp 1`



For more information about the workaround, please see the NetApp public report. Note that a NetApp login is required to view the report: <http://nt-ap.com/2w6myr4>



For more information about Windows Offloaded Data Transfers see: [https://technet.microsoft.com/en-us/library/hh831628\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/hh831628(v=ws.11).aspx)

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities. Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install by using the [HWU application](#) at the [NetApp Support](#) site.

Access the [HWU](#) application to view the System Configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers found in the [AFF A300 Series product documentation](#) at the [NetApp Support](#) site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of [disk shelves](#) that are supported by the AFF A300 is available at the [NetApp Support](#) site.

For SAS disk shelves with NetApp storage controllers, refer to the [SAS Disk Shelves Universal SAS and ACP Cabling Guide](#) for proper cabling guidelines.

NetApp ONTAP 9.1

Complete Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the [ONTAP 9.1 Software Setup Guide](#). You must have access to the [NetApp Support](#) site to open the cluster setup worksheet.

Configure ONTAP Nodes

Before running the setup script, review the configuration worksheets in the [ONTAP 9.1 Software Setup Guide](#) to learn about configuring ONTAP. Table 4 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 4 ONTAP software installation prerequisites

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Data ONTAP 9.1 URL	<url-boot-software>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e00M for the network port you want to use for the download.
7. Enter y to reboot now.

8. Enter the IP address, netmask, and default gateway for e0M.

```
<node01-mgmt-ip> <node01-mgmt-mask> <node01-mgmt-gateway>
```

9. Enter the URL where the software can be found.



This web server must be reachable.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.



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

13. Press **Ctrl-C** when the following message displays:

```
Press Ctrl-C for Boot Menu
```

14. Select option `4` for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with node `o2` configuration while the disks for node `o1` are zeroing.

Configure Node `o2`

To configure node `o2`, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press **Ctrl-C** to exit the autoboot loop when the following message displays:

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

2. Allow the system to boot up.

```
autoboot
```

3. Press **Ctrl-C** when prompted.



If ONTAP 9.1 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.1 is the version being booted, select option 8 and y to reboot the node. Then continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select eoM for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for eoM.

```
<node02-mgmt-ip> <node02-mgmt-mask> <node02-mgmt-gateway>
```

9. Enter the URL where the software can be found.



This web server must be reachable.

```
<url-boot-software>
```

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.



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

13. Press Ctrl-C when you see this message:

```
Press Ctrl-C for Boot Menu
```

14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with node 02 configuration while the disks for node 01 are zeroing.

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when ONTAP 9.1 boots on the node for the first time.

1. Follow the prompts to set up node 01:

```

Welcome to node setup.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
    Any changes you made before quitting will be saved.

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

This system will send event messages and weekly reports to NetApp Technical Support.
To disable this feature, enter "autosupport modify -support disable" within 24 hours.
Enabling AutoSupport can significantly speed problem determination and resolution should a problem occur on
your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

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

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

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

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

Table 5 Cluster create in ONTAP prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	<clustername>
ONTAP base license	<cluster-base-license-key>
Cluster management IP address	<clustermgmt-ip>
Cluster management netmask	<clustermgmt-mask>
Cluster management gateway	<clustermgmt-gateway>
Cluster node 01 IP address	<node01-mgmt-ip>
Cluster node 01 netmask	<node01-mgmt-mask>
Cluster node 01 gateway	<node01-mgmt-gateway>
Cluster node 02 IP address	<node02-mgmt-ip>

Cluster Detail	Cluster Detail Value
Cluster node 02 netmask	<node02-mgmt-mask>
Cluster node 02 gateway	<node02-mgmt-gateway>
Node 01 service processor IP address	<node01-SP-ip>
Node 02 service processor IP address	<node02-SP-ip>
DNS domain name	<dns-domain-name>
DNS server IP address	<dns-ip>
NTP server IP address	<ntp-ip>



Cluster setup can also be performed with the command line interface. This document describes the cluster setup using the NetApp System Manager guided setup.

3. Click Guided Setup on the welcome screen.

Cluster Setup Workflow - X

Not Secure | <https://192.168.156.61/sysmgr/SysMgr.html>

NetApp OnCommand System Manager

Getting Started

Language English (English)

Welcome to the Guided Cluster Setup

Perform the following to set up a cluster:

- Create a cluster, add nodes and admin credentials
- Create management LIFs, configure Service Processor, DNS, and NTP servers
- Configure AutoSupport Messages and Event Notifications


i For information related to setting up the cluster, [click here](#)

Template File

Browse to select a .csv file...


i To download the template, click [file.csv](#) or [file.xlsx](#)

Important: You can download the template in ".csv" or ".xlsx" format. However, you can upload only those templates that are in ".csv" format.



Click to set up the cluster


4. In the Cluster screen, complete the following steps:
 - Enter the cluster and node names.
 - Select the cluster configuration.
 - Enter and confirm the password.
 - (Optional) Enter the cluster base and feature licenses.



Cluster Name

Nodes

Not sure all nodes have been discovered? [Refresh](#)



Cluster Configuration: Switched Cluster Switchless Cluster

Ensure that the hardware connectivity is set up for the two-node switchless cluster.

Username

Password

Confirm Password

Cluster Base License (Optional)

For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional)

Cluster Base License is mandatory to add Feature Licenses.



The nodes are discovered automatically, if they are not discovered, click the Refresh link. By default, the cluster interfaces are created on all new storage controllers shipped from the factory. If all the nodes are not discovered, then configure the cluster using the command line. Cluster license and feature licenses can also be installed after completing the cluster creation.

5. Click Submit.


6. On the network page, complete the following sections:

- Cluster Management
 - Enter the IP address, netmask, gateway and port details.
- Node Management
 - Enter the node management IP addresses and port details for all the nodes.
- Service Processor Management
 - Enter the IP addresses for all the nodes.
- DNS Details
 - Enter the DNS domain names and server address.
- NTP Details
 - Enter the primary and alternate NTP server.

7. Click Submit.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



2 Network (Management)

IP Addresses (IPv4) required Enter 1 Cluster Management, 1 Node Management, and 2 Service Processor IP Addresses. You can override the Service Processor IP Address.

IP Address Range You must enter the default network details manually.

Cluster Management	IP Address: <input type="text" value="192.168.156.60"/>	Netmask: <input type="text" value="255.255.255.0"/>	Gateway (Optional): <input type="text" value="192.168.156.1"/>	Port: <input type="text" value="e0c"/>
⚠ Ensure that the cluster management LIF is reachable or a Gateway is configured for the same subnet in which the cluster management LIF is present.				
Node Management	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
bb04-aff300-1	<input type="text" value="192.168.156.61"/>	<input type="text" value="e0M"/>		
bb04-aff300-2	<input type="text" value="192.168.156.62"/>	<input type="text" value="e0M"/>		
Service Processor Management	Default values have been detected for the Service Processor.			
<input type="checkbox"/> Override the default values (Gateway is mandatory)				
<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.				
bb04-aff300-1	<input type="text" value="192.168.156.58"/>			
bb04-aff300-2	<input type="text" value="192.168.156.59"/>			

3 DNS Details

DNS Domain Names:

DNS Server IP Address:

4 NTP Details

Primary NTP Server:

Alternative NTP Server (Optional):

8. On the Support page, configure the AutoSupport and Event Notifications sections.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:

<input checked="" type="checkbox"/>	Email	SMTP Mail Host <input type="text" value="testvikings.smtp.cisco.com"/>	Email Addresses <input type="text" value="adminvikings@cisco.com"/>
<input type="checkbox"/>	SNMP	SNMP Trap Host <input type="text"/>	
<input type="checkbox"/>	Syslog	Syslog Server <input type="text"/>	

Submit

9. Click Submit.
10. On the Summary page, review the configuration details.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



[Click here to view the summary](#)

The next step will be to configure your aggregates, SVM and Storage Objects. Click the button below to start provisioning your storage.

[Manage your cluster](#)



The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document, we assume that it is on the same subnet.

Login to the Cluster

To log in to the cluster, complete the following steps:

1. Open an SSH connection to either the cluster IP or host name.
2. Log in to the admin user with the password you provided earlier.

Zero All Spare Disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares.
```



Advanced Data Partitioning creates a root partition and two data partitions on each SSD drive in an All Flash FAS configuration. Disk autoassign should have assigned one data partition to each node in an HA pair.



If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the `disk option modify` command. Spare partitions can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard Unified Target Adapter 2 Port Personality

To set the personality of the onboard unified target adapter 2 (UTA2), complete the following steps:

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin show` command.

```
ucadmin show
Current Current Pending Pending Admin
```

Node	Adapter	Mode	Type	Mode	Type	Status
<st-node01>	0e	fc	target	-	-	online
<st-node01>	0f	fc	target	-	-	online
<st-node01>	0g	cna	target	-	-	online
<st-node01>	0h	cna	target	-	-	online
<st-node02>	0e	fc	target	-	-	online
<st-node02>	0f	fc	target	-	-	online
<st-node02>	0g	cna	target	-	-	online
<st-node02>	0h	cna	target	-	-	online

8 entries were displayed.

- Verify that the Current Mode and Current Type properties for all ports are set properly. Set the ports used for FC connectivity to mode `fc`. The port type for all protocols should be set to `target`. Change the port personality by running the following command:

```
ucadmin modify -node <home-node-of-the-port> -adapter <port-name> -mode fc -type target.
```



The ports must be offline to run this command. To take an adapter offline, run the `fcip adapter modify -node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, 0e and 0f).



After conversion, a reboot is required. After reboot, bring the ports online by running `fcip adapter modify -node <home-node-of-the-port> -adapter <port-name> -state up`.

Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, run the following command:



A storage virtual machine (SVM) is referred to as a Vserver (or vservers) in the GUI and CLI.

Run the following command:

```
network interface modify -vserver <clustername> -lif cluster_mgmt -auto-revert true
```

Set Up Management Broadcast Domain

By default, all network ports are included in the default broadcast domain. Network ports used for data services (for example, e0d, e2a, and e2e) should be removed from the default broadcast domain, leaving just the management network ports (e0c and e0M). To perform this task, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports bb04-affa300-1:e0d,bb04-affa300-1:e0g,bb04-
affa300-1:e0h,bb04-affa300-1:e2a,bb04-affa300-1:e2e,bb04-affa300-2:e0d,bb04-affa300-2:e0g,bb04-affa300-
2:e0h,bb04-affa300-2:e2a,bb04-affa300-2:e2e
broadcast-domain show
```

Set Up Service Processor Network Interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

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

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



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

Create Aggregates

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it should contain.

To create new aggregates, run the following commands:

```
aggr create -aggregate aggr1_node01 -node <st-node01> -diskcount <num-disks>
aggr create -aggregate aggr1_node02 -node <st-node02> -diskcount <num-disks>
```



You should have the minimum number of hot spare disks for hot spare disk partitions recommended for your aggregate. For all-flash aggregates, you should have a minimum of one hot spare disk or disk partition. For non-flash homogenous aggregates, you should have a minimum of two hot spare disks or disk partitions.



For Flash Pool aggregates, you should have a minimum of two hot spare disks or disk partitions for each disk type. Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDs, you might want to create an aggregate with all but one remaining disk (spare) assigned to the controller.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both `aggr1_node1` and `aggr1_node2` are online.

(Optional) Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02. The aggregate is automatically renamed if system-guided setup is used.

```
aggr show
aggr rename -aggregate aggr0 -newname <node01-rootaggrname>
```

Verify Storage Failover

To confirm that storage failover is enabled, run the following commands for a failover pair:

1. Verify the status of the storage failover.

```
storage failover show
```



Both `<st-node01>` and `<st-node02>` must be able to perform a takeover. Continue with step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <st-node01> -enabled true
```



Enabling failover on one node enables it for both nodes.

- Verify the HA status for a two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

- Continue with step 6 if high availability is configured.



Only enable HA mode for two-node clusters. Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true
Do you want to continue? {y|n}: y
```

- Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <node02-mgmt-ip> -node <st-node01>
storage failover modify -hwassist-partner-ip <node01-mgmt-ip> -node <st-node02>
```

Disable Flow Control on 10GbE and 40GbE ports

NetApp recommends disabling flow control on all the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

- Run the following commands to configure node 01:

```
network port modify -node <st-node01> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

- Run the following commands to configure node 02:

```
network port modify -node <st-node02> -port e0a,e0b,e0e,e0f,e0g,e0h,e2a,e2e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
network port show -fields flowcontrol-admin
```

Disable Unused FCoE Capability on CNA Ports

If the UTA2 port is set to CNA mode and is only expected to handle Ethernet data traffic (for example CIFS), then the unused FCoE capability of the port should be disabled by setting the corresponding FCP adapter to state down with the `fc adapter modify` command. Here are some examples:

```
fc adapter modify -node <st-node01> -adapter 0g -status-admin down
fc adapter modify -node <st-node01> -adapter 0h -status-admin down
fc adapter modify -node <st-node02> -adapter 0g -status-admin down
fc adapter modify -node <st-node02> -adapter 0h -status-admin down
fc adapter show -fields status-admin
```

Configure Network Time Protocol

To configure time synchronization on the cluster, complete the following steps:

1. Set the time zone for the cluster.

```
timezone <timezone>
```



For example, in the eastern United States, the time zone is `America/New_York`.

2. Set the date for the cluster.

```
date <ccyymmddhhmm.ss>
```



The format for the date is `<[Century][Year][Month][Day][Hour][Minute].[Second]>` (for example, `201703231549.30`).

3. Configure the Network Time Protocol (NTP) servers for the cluster.

```
cluster time-service ntp server create -server <switch-a-ntp-ip>
cluster time-service ntp server create -server <switch-b-ntp-ip>
```

Configure Simple Network Management Protocol

To configure the Simple Network Management Protocol (SNMP), complete the following steps:

1. Configure basic SNMP information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <snmp-contact>
snmp location "<snmp-location>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts, such as a DFM server or another fault management system.

```
snmp traphost add <oncommand-um-server-fqdn>
```

Configure SNMPv1 Access

To configure SNMPv1 access, set the shared, secret plain-text password (called a community).

```
snmp community add ro <snmp-community>
```

Configure AutoSupport

NetApp AutoSupport® sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <mailhost> -transport https -support enable
-noteto <storage-admin-email>
```


Enable Cisco Discovery Protocol

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

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



To be effective, CDP must also be enabled on directly connected networking equipment such as switches and routers.

Create Jumbo Frame MTU Broadcast Domains in ONTAP

To create a data broadcast domain with an MTU of 9000 for SMB and management on ONTAP, run the following command:

```
broadcast-domain create -broadcast-domain IB-MGMT-904 -mtu 9000
broadcast-domain create -broadcast-domain Infra_MS_SMB -mtu 9000
```

Create Interface Groups

To create LACP interface groups for the 10GbE data interfaces, run the following commands:

```
ifgrp create -node <st-node01> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node01> -ifgrp a0a -port e2e

ifgrp create -node <st-node02> -ifgrp a0a -distr-func port -mode multimode_lacp
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2a
ifgrp add-port -node <st-node02> -ifgrp a0a -port e2e

ifgrp show
```

Create VLANs

To create SMB VLAN, create SMB VLAN ports and add them to the SMB broadcast domain:

```
network port modify -node <st-node01> -port a0a -mtu 9000
network port modify -node <st-node02> -port a0a -mtu 9000

network port vlan create -node <st-node01> -vlan-name a0a-<infra-smb-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-smb-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_MS_SMB -ports <st-node01>:a0a-<infra-smb-vlan-id>, <st-
node02>:a0a-<infra-smb-vlan-id>
```

To create In-Band-Management VLAN and add them to the management broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<MS-IB-MGMT-VLAN>
network port vlan create -node <st-node02> -vlan-name a0a-<MS-IB-MGMT-VLAN>

broadcast-domain add-ports -broadcast-domain IB-MGMT-904 -ports <st-node01>:a0a-<MS-IB-MGMT-VLAN>, <st-
node02>:a0a-<MS-IB-MGMT-VLAN>
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-MS-SVM -rootvolume ms_rootvol -aggregate aggr1_node01 -rootvolume-security-style ntfs
```

2. Remove the unused data protocols (NFS, iSCSI, and NDMP) from the SVM.

```
vserver remove-protocols -vserver Infra-MS-SVM -protocols nfs,ndmp,iscsi
```

3. Add the two data aggregates to the Infra-MS-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-MS-SVM -aggr-list aggr1_node01,aggr1_node02
```

Create the CIFS Service

You can enable and configure CIFS servers on storage virtual machines (SVMs) with NetApp FlexVol® volumes to let SMB clients access files on your cluster. Each data SVM in the cluster can be bound to exactly one Active Directory domain. However, the data SVMs do not need to be bound to the same domain. Each data SVM can be bound to a unique Active Directory domain.

Before configuring the CIFS service on your SVM, the DNS must be configured. To do so, complete the following steps:

1. Configure the DNS for your SVM.

```
dns create -vserver Infra-Hyper-V -domains <<domain_name>> -name-servers <<dns_server_ip>>
```

The node management network interfaces should be able to route to the Active Directory domain controller to which you want to join the CIFS server. Alternatively, a data network interface must exist on the SVM that can route to the Active Directory domain controller.

2. Create a network interface on the in-band VLAN.

```
network interface create -vserver Infra-MS-SVM -lif <<svm_mgmt_lif_name>> -role data -data-protocol none -home-node <<st-node-01>> -home-port a0a-<MS-IB-MGMT-VLAN> -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

3. Create the CIFS service.

```
vserver cifs create -vserver Infra-MS-SVM -cifs-server Infra-CIFS -domain flexpod.local
```

In order to create an Active Directory machine account for the CIFS server, you must supply the name and password of a Windows account with sufficient privileges to add computers to the "CN=Computers" container within the "FLEXPOD.LOCAL" domain.

Enter the user name: Administrator@flexpod.local

Enter the password:

Modify Storage Virtual Machine Options

NetApp ONTAP can use automatic node referrals to increase SMB client performance on SVMs with FlexVol volumes. This feature allows the SVM to automatically redirect a client request to a network interface on the node where the FlexVol volume resides.

To enable automatic node referrals on your SVM, run the following command:

```
set -privilege advanced
vserver cifs options modify -vserver Infra-MS-SVM -is-referral-enabled true
```

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra-MS-SVM -volume ms_rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-MS-SVM -volume ms_rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-MS-SVM:ms_rootvol -destination-path Infra-MS-SVM:ms_rootvol_m01 -type LS
-schedule 15min
snapmirror create -source-path Infra-MS-SVM:ms_rootvol -destination-path Infra-MS-SVM:ms_rootvol_m02 -type LS
-schedule 15min
```

4. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path Infra-MS-SVM:ms_rootvol
snapmirror show
```

Create Block Protocol FC Service

To create the FCP service on each SVM, run the following command. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM.fcp create -vserver Infra-MS-SVM

```
fcp show
```



If an FC license is not installed during the cluster configuration, you must install the FC service license.

Configure HTTPS Access

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate and obtain parameters (for example, <serial-number>) by running the following command:

```
security certificate show
```

For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. Delete the two default certificates and replace them with either self-signed certificates or certificates from a certificate authority (CA). To delete the default certificates, run the following commands:

```
security certificate delete -vserver Infra-MS-SVM -common-name Infra-MS-SVM -ca Infra-MS-SVM -type server -
serial <serial-number>
```



Deleting expired certificates before creating new certificates is a best practice. Run the security certificate delete command to delete the expired certificates. In the following command, use TAB completion to select and delete each default certificate.

3. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the Infra-MS-SVM and the cluster SVM. Use TAB completion to aid in the completion of these commands.

```
security certificate create -common-name <cert-common-name> -type server -size 2048 -country <cert-country>
-state <cert-state> -locality <cert-locality> -organization <cert-org> -unit <cert-unit> -email-addr <cert-
email> -expire-days <cert-days> -protocol SSL -hash-function SHA256 -vserver Infra-MS-SVM
```

4. To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the security certificate show command.
5. Enable each certificate that was just created by using the --server-enabled true and --client-enabled false parameters. Use TAB completion to aid in the completion of these commands.

```
security ssl modify -vserver <clustername> -server-enabled true -client-enabled false -ca <cert-ca> -serial
<cert-serial> -common-name <cert-common-name>
```

6. Disable HTTP cluster management access.

```
system services firewall policy delete -policy mgmt -service http -vserver <clustername>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

7. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled true
```

Create SMB Export Policy

Optionally, you can use export policies to restrict SMB access to files and folders on SMB volumes. You can use export policies in combination with share level and file level permissions to determine effective access rights.

To create an export policy that limits access to devices in the domain, run the following command:

```
export-policy create -vserver Infra-MS-SVM -policyname smb

export-policy rule create -vserver Infra-MS-SVM -policyname cifs -clientmatch flexpod.local -rorule
krb5i,krb5p -rwrule krb5i,krb5p
```

Create NetApp FlexVol Volumes

```
volume create -vserver Infra-MS-SVM -volume infra_datastore_1 -aggregate aggr1_node01 -size 500GB -state
online -policy smb -security-style ntfs -junction-path /infra_datastore_1 -space-guarantee none -percent-
snapshot-space 5

volume create -vserver Infra-MS-SVM -volume infra_datastore_2 -aggregate aggr1_node02 -size 500GB -state
online -policy smb -security-style ntfs -junction-path /infra_datastore_2 -space-guarantee none -percent-
snapshot-space 5

volume create -vserver Infra-MS-SVM -volume witness_FC_6332 -aggregate aggr1_node01 -size 5GB -state online -
policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-MS-SVM -volume HV_boot -aggregate aggr1_node01 -size 500GB -state online -policy
default -security-style ntfs -space-guarantee none -percent-snapshot-space 0

snapmirror update-ls-set -source-path Infra-MS-SVM:ms_rootvol
```

Create CIFS Shares

A CIFS share is a named access point in a volume that enables CIFS clients to view, browse, and manipulate files on a file server.

```
cifs share create -vserver Infra-MS-SVM -share-name infra_share_1_Share -path /infra_datastore_1 -share-properties oplocks,browsable,continuously-available,showsnapshot
```

```
cifs share create -vserver Infra-MS-SVM -share-name infra_share_2_Share -path /infra_datastore_2 -share-properties oplocks,browsable,continuously-available,showsnapshot
```

Configuring share permissions by creating access control lists (ACLs) for SMB shares enables you to control the level of access to a share for users and groups.

To configure the Administrators and Hyper-V hosts to access to the CIFS Share, run the following commands:

```
cifs share access-control create -vserver Infra-MS-SVM -share 6332_Share -user-or-group Flexpod\Administrator -user-group-type windows -permission Full_Control
```

```
cifs share access-control create -vserver Infra-MS-SVM -share 6332_Share -user-or-group flexpod\

```

```
cifs share access-control create -vserver Infra-MS-SVM -share 6332_Share -user-or-group flexpod\

```

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -size 200GB -ostype windows_2008 -space-reserve disabled
```

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -size 200GB -ostype windows_2008 -space-reserve disabled
```

Create Witness LUN

A witness LUN is required in a Hyper-V cluster. To create the witness LUN, run the following command:

```
lun create -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -size 1GB -ostype windows_2008 -space-reserve disabled
```

Schedule Deduplication

On NetApp All Flash FAS systems, deduplication is enabled by default. To schedule deduplication, complete the following steps:

1. After the volumes are created, assign a once-a-day deduplication schedule to HV_boot, infra_datastore_1 and infra_datastore_2:

```
efficiency modify -vserver Infra-MS-SVM -volume HV_boot -schedule sun-sat@0
efficiency modify -vserver Infra-MS-SVM -volume infra_datastore_1 -schedule sun-sat@0
efficiency modify -vserver Infra-MS-SVM -volume infra_datastore_2 -schedule sun-sat@0
```

Create FC LIFs

Run the following commands to create four FC LIFs (two on each node):

```
network interface create -vserver Infra-MS-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <st-node01> -home-port 0e -status-admin up
```

```
network interface create -vserver Infra-MS-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <st-node01> -home-port 0f -status-admin up

network interface create -vserver Infra-MS-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node <st-node02> -home-port 0e -status-admin up

network interface create -vserver Infra-MS-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node <st-node02> -home-port 0f -status-admin up
```

Create SMB LIF

To create an SMB LIF, run the following commands:

```
network interface create -vserver Infra-MS-SVM -lif smb_lif01 -role data -data-protocol cifs -home-node <st-node01> -home-port a0a-<infra-smb-vlan-id> -address <node01-smb_lif01-ip> -netmask <node01-smb_lif01-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-MS-SVM -lif smb_lif02 -role data -data-protocol cifs -home-node <st-node02> -home-port a0a-<infra-smb-vlan-id> -address <node02-smb_lif02-ip> -netmask <node02-smb_lif02-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface show
```

Add Infrastructure SVM Administrator

To add an infrastructure SVM administrator and an SVM administration LIF in the out-of-band management network, complete the following steps:



If the network interface created during the Create the CIFS Service step was created on the out-of-band network, skip to step 2.

1. Create a network interface.

```
network interface create -vserver Infra-MS-SVM -lif svm-mgmt -role data -data-protocol none -home-node <st-node02> -home-port e0c -address <svm-mgmt-ip> -netmask <svm-mgmt-mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```



The SVM management IP in this step should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-MS-SVM -destination 0.0.0.0/0 -gateway <svm-mgmt-gateway>

network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-MS-SVM
Enter a new password: <password>
Enter it again: <password>

security login unlock -username vsadmin -vserver Infra-MS-SVM
```



A cluster serves data through at least one and possibly several SVMs. We have just described the creation of a single SVM. If you would like to configure your environment with multiple SVMs, this is a good time to create them.

Server Configuration

Cisco UCS Base Configuration

This FlexPod deployment will show configuration steps for the Cisco UCS 6332-16UP Fabric Interconnects (FI) in a design that will support Fibre Channel to the NetApp AFF through the Cisco Nexus.

Perform Initial Setup

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

Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

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

```
Enter the configuration method: gui
```

```
Physical switch Mgmt0 IP address: <ucsa-mgmt-ip>
```

```
Physical switch Mgmt0 IPv4 netmask: <ucsa-mgmt-mask>
```

```
IPv4 address of the default gateway: <ucsa-mgmt-gateway>
```

2. Using a supported web browser, connect to <ucsa-mgmt-ip>, accept the security prompts, and click the 'Express Setup' link under HTML.
3. Select Initial Setup and click Submit.
4. Select Enable clustering, Fabric A, and IPv4.
5. Fill in the Virtual IP Address with the UCS cluster IP.
6. Completely fill in the System setup section. For system name, use the overall UCS system name. For the Mgmt IP Address, use <ucsa-mgmt-ip>.

Basic Settings

Cluster and Fabric setup

Enable clustering
 Standalone mode
 Synchronize

Fabric Setup: Fabric A Fabric B

IPv4
 IPv6

Virtual IP Address: . . .

System setup

Enforce strong password?: Yes No

System name:

Admin Password: Confirm Admin password:

Mgmt IP Address: . . . Mgmt IP Netmask: . . .

Default Gateway: . . .

DNS Server IP: . . . Domain Name :

UCS Central managed environment

UCS Central IP: . . . Shared Secret:

7. Click Submit.

Cisco UCS Fabric Interconnect B

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

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

Enter the configuration method: `gui`

Physical switch Mgmt0 IP address: `<ucsb-mgmt-ip>`

Physical switch Mgmt0 IPv4 netmask: `<ucsb-mgmt-mask>`

IPv4 address of the default gateway: `<ucsb-mgmt-gateway>`

2. Using a supported web browser, connect to <ucs-b-mgmt-ip>, accept the security prompts, and click the 'Express Setup' link under HTML.
3. Under System setup, enter the Admin Password entered above and click Submit.
4. Enter <ucs-b-mgmt-ip> for the Mgmt IP Address and click Submit.

Cisco UCS Setup

Log in to Cisco UCS Manager

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

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



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

2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

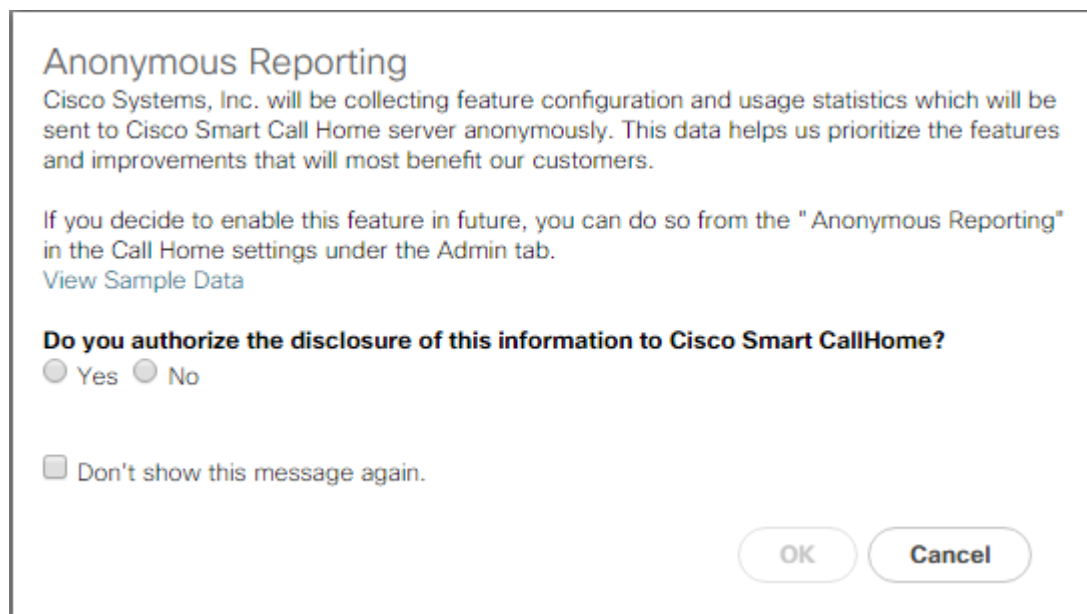
Upgrade Cisco UCS Manager Software to Version 3.1(3a)

This document assumes the use of Cisco UCS 3.1(3a). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(3a), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following 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.



Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in UCSM. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, complete the following steps:

1. In Cisco UCS Manager, click Admin on the 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.

Configure Unified Ports


Fiber Channel port configurations differ slightly between the 6332-16UP and the 6248UP Fabric Interconnects. Both Fabric Interconnects have a slider mechanism within the UCSM GUI interface, but the fiber channel port selection options for the 6332-16UP are from the first 16 ports starting from the first port on the left, 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.

To enable the fiber channel ports, complete the following steps for the 6332-16UP:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. Select Configure Unified Ports.
4. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.

5. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.

Configure Unified Ports ? X



Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Fibre Channel ports (Purple), while the ports to the right are Ethernet ports (Blue).

Port	Transport	If Role or Port Channel Membership	Desired If Role
Port 1	ether	Unconfigured	FC Uplink
Port 2	ether	Unconfigured	FC Uplink
Port 3	ether	Unconfigured	FC Uplink
Port 4	ether	Unconfigured	FC Uplink
Port 5	ether	Unconfigured	FC Uplink
Port 6	ether	Unconfigured	FC Uplink
Port 7	ether	Unconfigured	
Port 8	ether	Unconfigured	
Port 9	ether	Unconfigured	
Port 10	ether	Unconfigured	
Port 11	ether	Unconfigured	
Port 12	ether	Unconfigured	
Port 13	ether	Unconfigured	
Port 14	ether	Unconfigured	
Port 15	ether	Unconfigured	
Port 16	ether	Unconfigured	

OK
Cancel

6. Click OK, then Yes, then OK to continue
7. Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary)
8. Select Configure Unified Ports.
9. Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
10. Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 6, 12, or 16 ports to be set as FC Uplinks.
11. Click OK, then Yes, then OK to continue.
12. Wait for both the Fabric Interconnects to reboot.

13. Log back into UCS Manager.



This process will be similar for the UCS 6248UP Fabric Interconnect, but will be in increments of two unified ports that can be converted to FC uplinks, and will slide from the right to the left instead of the left to the right process used with the UCS 6332-16UP Fabric Interconnects.

Add Block of IP Addresses for KVM Access

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

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

Create Block of IPv4 Addresses

From : 192.168.94.241 Size : 8

Subnet Mask : 255.255.255.0 Default Gateway : 192.168.94.254

Primary DNS : 0.0.0.0 Secondary DNS : 0.0.0.0

OK Cancel

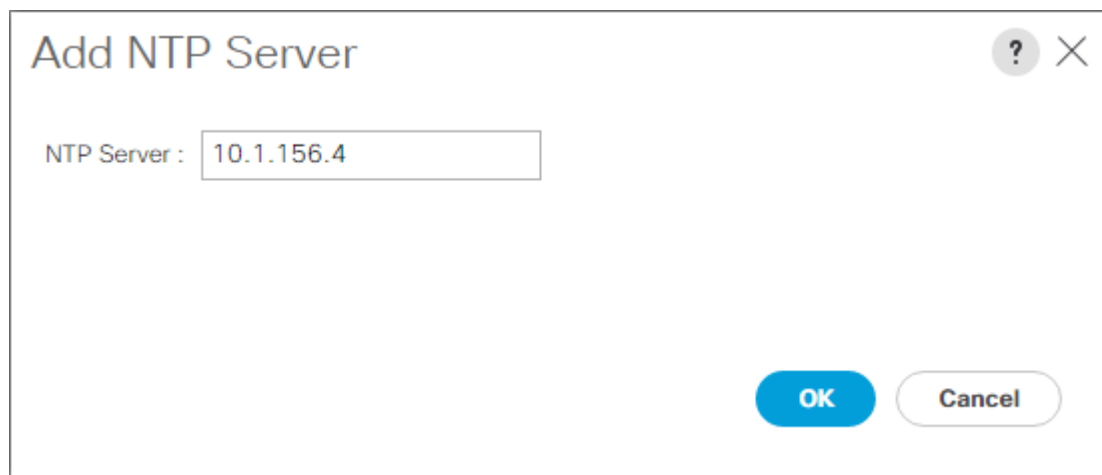
5. Click OK to create the block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

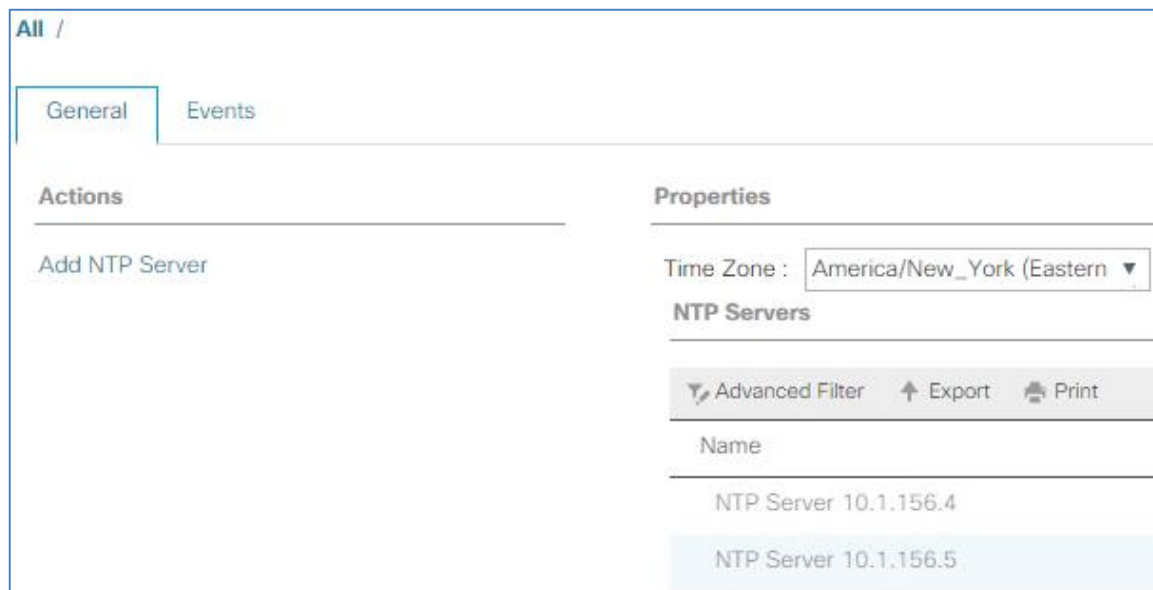
To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Expand All > Time Zone Management.
3. Select Timezone.
4. In the Properties pane, select the appropriate time zone in the Timezone menu.

5. Click Save Changes, and then click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> and click OK. Click OK on the confirmation.



8. Click Add NTP Server.
9. Enter <switch-b-ntp-ip> and click OK. Click OK on the confirmation.



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, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, click the Policies tab.

- Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
- Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.

The screenshot shows the 'Equipment' section of the Cisco UCS Manager interface. The 'Policies' tab is selected, and the 'Global Policies' sub-tab is active. The 'Chassis/FEX Discovery Policy' configuration is shown with the following settings:

- Action: 2 Link
- Link Grouping Preference: Port Channel (selected)
- Backplane Speed Preference: 40G (selected)

- Click Save Changes.
- Click OK.

Enable Server and Uplink Ports

To enable server and uplink ports, complete the following steps:

- In Cisco UCS Manager, click Equipment on the left.
- Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
- Expand Ethernet Ports.
- Select the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select "Configure as Server Port."
- Click Yes to confirm server ports and click OK.
- Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
- Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.



The last 6 ports of the UCS 6332 and UCS 6332-16UP FIs will only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only.

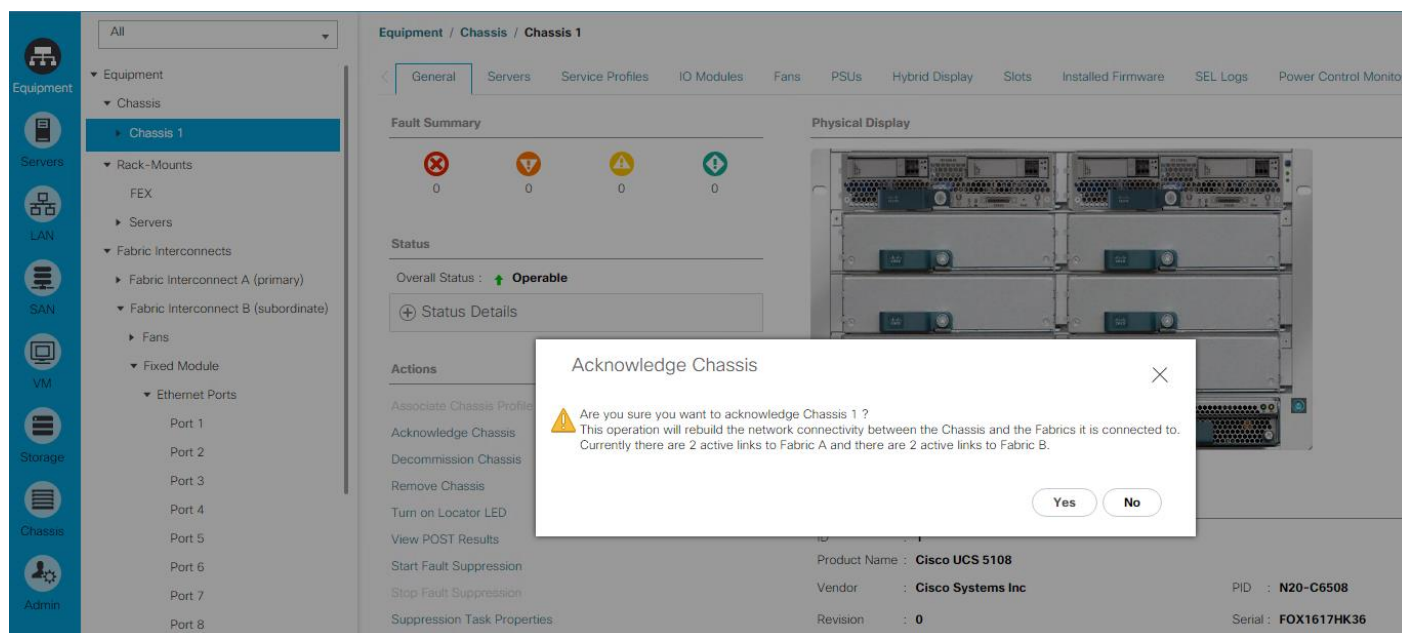
- Click Yes to confirm uplink ports and click OK.
- Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
- Expand Ethernet Ports.

11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
14. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external 2232 FEX modules, complete the following steps:

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



4. Click Yes and then click OK to complete acknowledging the chassis.
5. If Nexus 2232 FEX 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.

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.



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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 125 as the unique ID of the port channel.
6. Enter `vPC-125-Nexus` as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 126 as the unique ID of the port channel.
16. Enter `vPC-126-Nexus` as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create a WWNN Pool for FC Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

1. Select SAN on the left.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.

4. Select Create WWNN Pool to create the WWNN pool.
5. Enter `WWNN-POOL` for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.
7. Select **Sequential** for Assignment Order.

The screenshot shows a 'Create WWNN Pool' dialog box. On the left, a vertical sidebar contains two numbered steps: '1 Define Name and Description' (highlighted in blue) and '2 Add WWN Blocks'. The main content area of the dialog has the following fields and controls:

- Name**: A text input field containing 'WWNN-POOL'.
- Description**: An empty text input field.
- Assignment Order**: A radio button group with 'Default' (unselected) and 'Sequential' (selected).

At the bottom of the dialog, there are four buttons: '< Prev' (disabled), 'Next >' (active), 'Finish' (disabled), and 'Cancel'.

8. Click Next.
9. Click Add.
10. Modify the From field as necessary for the UCS Environment

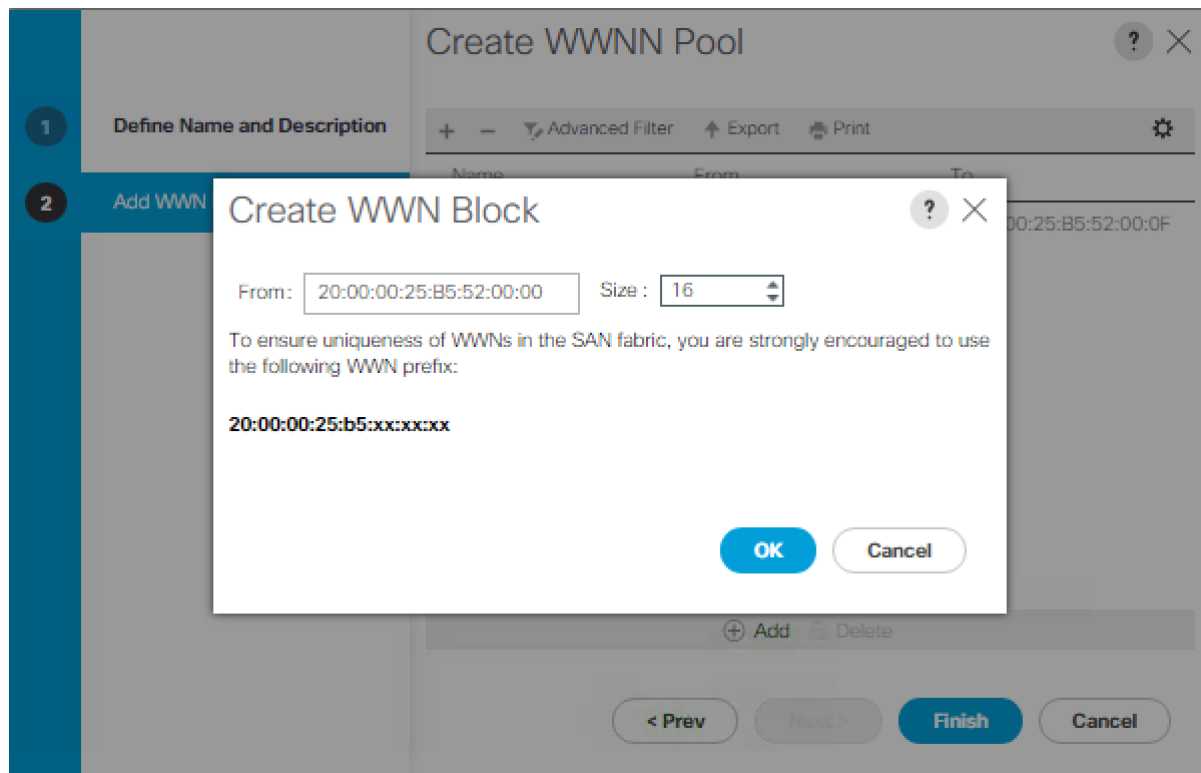


Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to 52 to represent as identifying information for this being in the UCS 6332 in the 4th cabinet.



Also, when having multiple UCS domains sitting in adjacency, it is important that these blocks, the WWNN, WWPN, and MAC hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources.



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

Create WWPN Pools

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.
4. Right-click WWPN Pools under the root organization.
5. Select Create WWPN Pool to create the WWPN pool.
6. Enter `WWPN-POOL-A` as the name of the WWPN pool.
7. Optional: Enter a description for the WWPN pool.
8. Select **Sequential** for Assignment Order

Create WWPN Pool

1 Define Name and Description

2 Add WWN Blocks

Name : WWPN-POOL-A

Description :

Assignment Order: Default Sequential

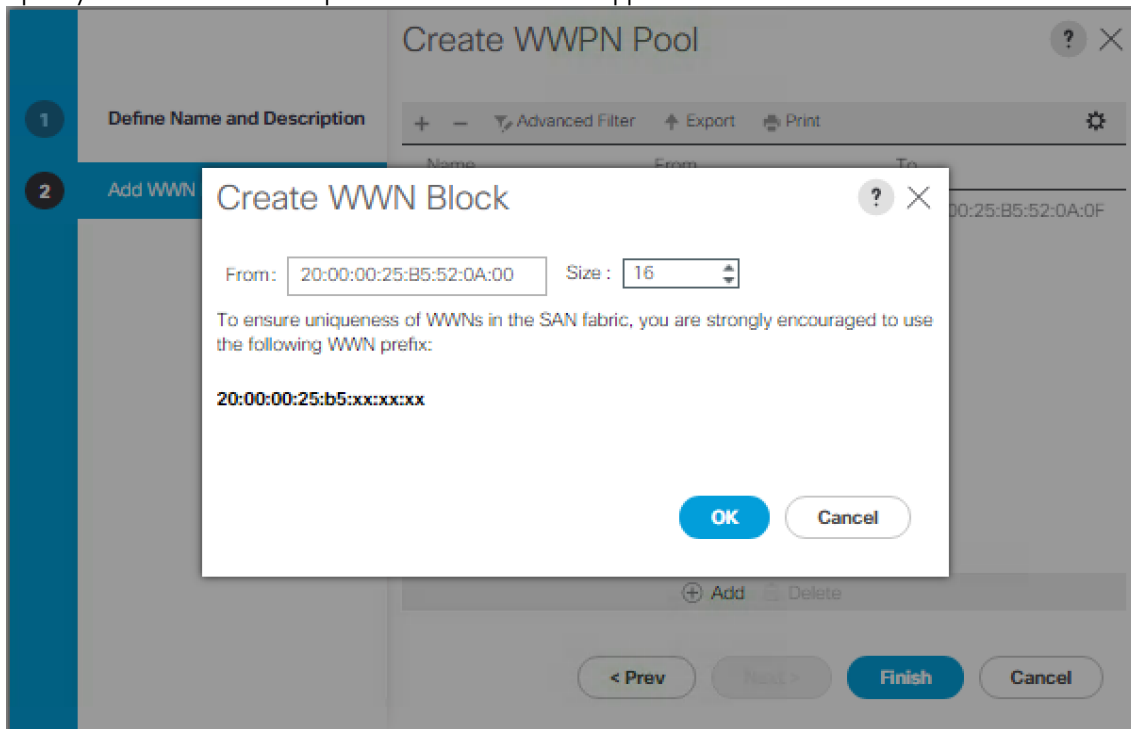
< Prev Next > Finish Cancel

9. Click Next.
10. Click Add.
11. Specify a starting WWPN



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with 20:00:00:25:B5:52:0A:00

12. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources.



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



For the FlexPod solution, the recommendation is to place `0B` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:52:0B:00`.

24. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources.

25. Click OK.
26. Click Finish.
27. In the confirmation message, click OK

Create VSAN

1. To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, complete the following steps:
2. In Cisco UCS Manager, click the SAN on the left.



In this procedure, two VSANs are created.

3. Select SAN > SAN Cloud.
4. Right-click VSANs.
5. Select Create VSAN.
6. Enter `VSAN-A` as the name of the VSAN to be used for Fabric A
7. Leave FC Zoning set at Disabled.
8. Select Fabric A.
9. Enter a unique VSAN ID and a corresponding FCoE VLAN ID that matches the configuration in the MDS switch for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

Create VSAN ? ×

Name :

FC Zoning Settings

FC Zoning : Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A. A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VSAN ID that maps to this VSAN. Enter the VLAN ID that maps to this VSAN.

VSAN ID : FCoE VLAN :

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

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 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 :

17. Click OK, and then click OK again

Create FC Uplink Port Channels

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

1. In Cisco UCS Manager, click SAN on the left.
2. Select SAN > SAN Cloud.
3. Expand Fabric A and select FC Port Channels.
4. Right-click FC Port Channels and select Create FC Port Channel.
5. Set a unique ID for the port channel and provide a unique name for the port channel.
6. Click Next.
7. Select the ports connected to Cisco MDS A and use >> to add them to the port channel

Create FC Port Channel

Port Channel Admin Speed : Auto

Ports		
Port	Slot ID	WWPN
3	1	20:03:8C:60...
4	1	20:04:8C:60...
5	1	20:05:8C:60...
6	1	20:06:8C:60...

Ports in the port channel		
Port	Slot ID	WWPN
1	1	20:01:8C:60...
2	1	20:02:8C:60...

Slot ID:
WWPN:

Slot ID:
WWPN:

< Prev Next > **Finish** Cancel

8. Click Finish to complete creating the port channel.
9. Click OK the confirmation.
10. Under FC Port-Channels, select the newly created port channel.
11. In the right pane, use the pulldown to select VSAN-A

SAN / SAN Cloud / Fabric A / FC Port Channels /

General Ports Faults Events Statistics

Status

Overall Status : ▼ **Failed**

Additional Info : **No operational members**

< III >

Actions

Enable Port Channel

Disable Port Channel

Add Ports

< III >

Properties

ID : **101**

Fabric ID : **A**

Port Type : **Aggregation**

Transport Type : **Fc**

Name :

Description :

VSAN :

Port Channel Admin Speed :

Operational Speed(Gbps) : **0**

12. Click Save Changes to assign the VSAN.
13. Click OK.
14. Expand Fabric B and select FC Port Channels.
15. Right-click FC Port Channels and select Create FC Port Channel.
16. Set a unique ID for the port channel and provide a unique name for the port channel.
17. Click Next.
18. Select the ports connected to Cisco MDS B and use >> to add them to the port channel.
19. Click Finish to complete creating the port channel.
20. Click OK on the confirmation.
21. Under FC Port-Channels, select the newly created port channel.
22. In the right pane, use the pulldown to select VSAN-B.
23. Click Save Changes to assign the VSAN.
24. Click OK.

Create vHBA Templates

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.

2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter vHBA-Template-A as the vHBA template name.
6. Keep Fabric A selected.
7. Leave Redundancy Type set to No Redundancy.
8. Select VSAN-A.
9. Leave Initial Template as the Template Type.
10. Select WWPN-POOL-A as the WWPN Pool.
11. Click OK to create the vHBA template.
12. Click OK

Create vHBA Template [?] [X]

Name : vHBA-Template-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

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

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-POOL-A(16/16)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

OK **Cancel**

13. Right-click vHBA Templates.
14. Select Create vHBA Template.

15. Enter `vHBA-Template-B` as the vHBA template name.
16. Leave Redundancy Type set to No Redundancy.
17. Select Fabric B as the Fabric ID.
18. Select VSAN-B.
19. Leave Initial Template as the Template Type.
20. Select WWPN-POOL-B as the WWPN Pool.
21. Click OK to create the vHBA template.
22. Click OK.

Create SAN Connectivity Policy

To configure the necessary Infrastructure SAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter `FC-BOOT` as the name of the policy.
6. Select the previously created WWNN-POOL for the WWNN Assignment.
7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter FABRIC-A as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select vHBA-Template-A.
11. In the Adapter Policy list, select WindowsBoot

Create vHBA ? X

Name : FABRIC-A

Use vHBA Template :

Redundancy Pair :

vHBA Template : vHBA-Template-A ▼

Peer Name :

Create vHBA Template

Adapter Performance Profile

Adapter Policy: WindowsBoot ▼

Create Fibre Channel Adapter Policy

OK Cancel

12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.
14. In the Create vHBA dialog box, enter FABRIC-B as the name of the vHBA.
15. Select the Use vHBA Template checkbox.
16. In the vHBA Template list, select vHBA-Template-B.
17. In the Adapter Policy list, select WindowsBoot.
18. Click OK

Create SAN Connectivity Policy ? X

Name :

Description :

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

World Wide Node Name

WWNN Assignment:

[Create WWNN Pool](#)

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

Name	WWPN
▶ vHBA FABRIC-B	Derived
▶ vHBA FABRIC-A	Derived

🗑 Delete ➕ Add ⚙ Modify

OK
Cancel

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
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 **Sequential** as the option for Assignment Order.

8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the of also embedding the UCS domain number information giving us 00:25:B5:52:0A:00 as our first MAC address.

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

Create a Block of MAC Addresses

First MAC Address : Size :

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

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



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the UCS domain number information giving us 00:25:B5:52:0B:00 as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
24. Click OK.
25. Click Finish.
26. 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, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-POOL` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select **Sequential** for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `MS-Server-Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the Hyper-V management cluster and click >> to add them to the `MS-Server-Pool` server pool.
9. Click Finish.
10. Click OK.

Create VLANs

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

1. In Cisco UCS Manager, click LAN on the left.



In this procedure, five unique VLANs are created. See Table 1.

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

Create VLANs ? X

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

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select `Set as Native VLAN`.
11. Click `Yes`, and then click `OK`.
12. Right-click `VLANs`.
13. Select `Create VLANs`
14. Enter `MS-IB-MGMT` as the name of the VLAN to be used for management traffic.
15. Keep the `Common/Global` option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.
17. Keep the `Sharing Type` as `None`.
18. Click `OK`, and then click `OK` again.
19. Right-click `VLANs`.
20. Select `Create VLANs`.

21. Enter `MS-SMB-1` as the name of the VLAN to be used for SMB File share.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the SMB File Share VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter `MS-SMB-2` as the name of the VLAN to be used for 2nd SMB File share.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the Infrastructure SMB File Share VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK, and then click OK again
33. Right-click VLANs.
34. Select Create VLANs.
35. Enter `MS-LVMN` as the name of the VLAN to be used for Live Migration.
36. Keep the Common/Global option selected for the scope of the VLAN.
37. Enter the Live Migration VLAN ID.
38. Keep the Sharing Type as None.
39. Click OK, and then click OK again.
40. Select Create VLANs.
41. Enter `MS-Cluster` as the name of the VLAN to be used for Cluster communication network.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the Cluster network VLAN ID.
44. Keep the Sharing Type as None.
45. Click OK, and then click OK again.
46. Select Create VLANs.
47. Enter `MS-Tenant-VM` as the name of the VLAN to be used for VM Traffic.

48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

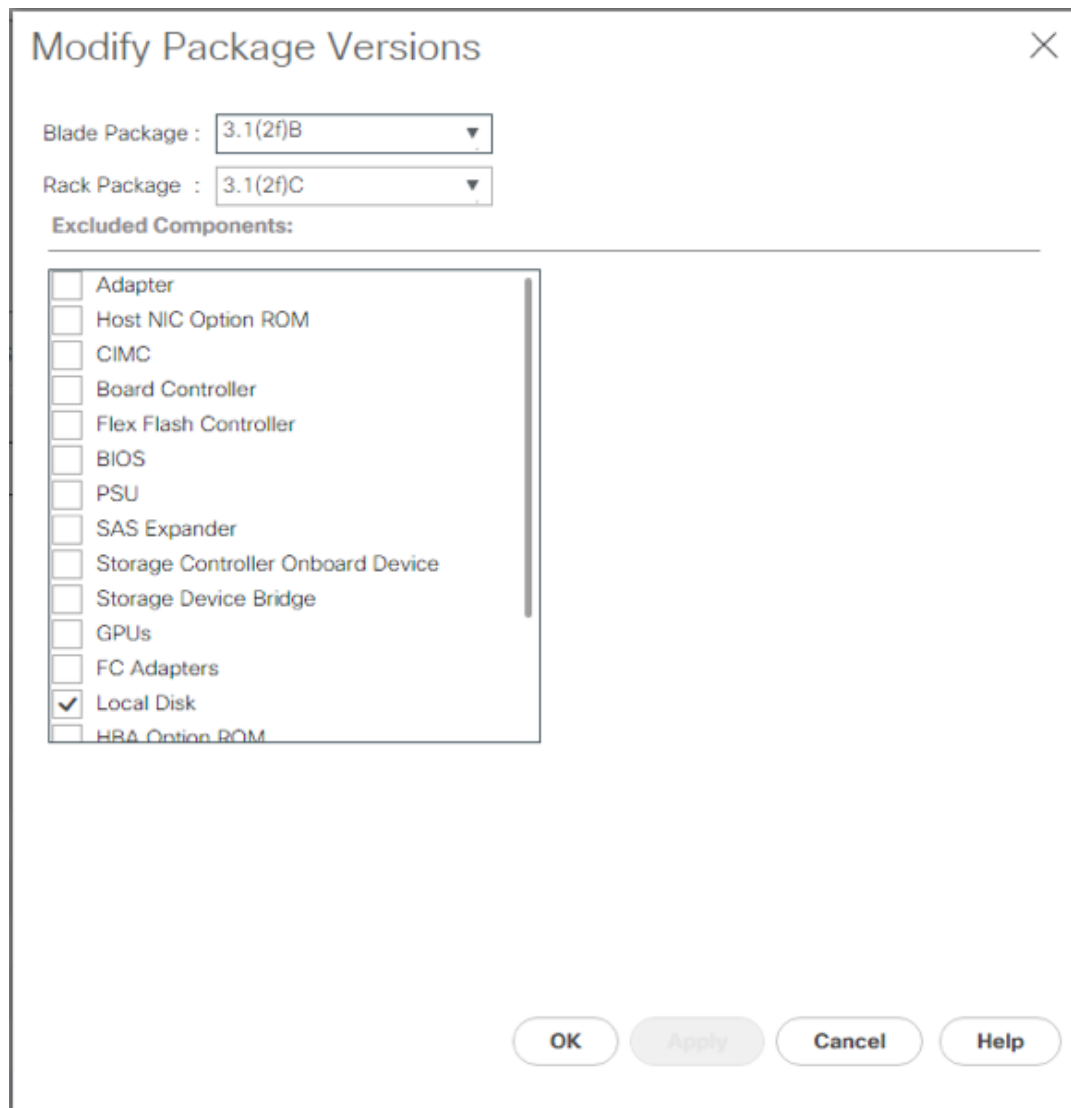
Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Na...	Multicast Policy N...
VLAN default (1)	1	Lan	Ether	No	None		
VLAN Native-VLAN (2)	2	Lan	Ether	Yes	None		
VLAN MS-IB-MGMT (90...	904	Lan	Ether	No	None		
VLAN MS-CSV (905)	905	Lan	Ether	No	None		
VLAN MS-LVMN (906)	906	Lan	Ether	No	None		
VLAN MS-Cluster (907)	907	Lan	Ether	No	None		
VLAN MS-Tenant-VM (...)	908	Lan	Ether	No	None		
VLAN MS-SMB-1 (3052)	3052	Lan	Ether	No	None		
VLAN MS-SMB-2 (3053)	3053	Lan	Ether	No	None		

Modify Default 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 properties.

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(2f) for both the Blade and Rack Packages.



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

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
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.
6. Click OK

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

Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.



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

To create a local disk configuration policy, complete the following steps:

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

Create Local Disk Configuration Policy ? X

Name : SAN-Boot

Description :

Mode : No Local Storage ▼

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

OK Cancel

8. Click OK.

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

To create a network control policy that enables CDP and LLDP on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable-CDP-LLDP` as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

Create Network Control Policy ? X

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

9. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter `No-Power-Cap` as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Create Power Control Policy ? X

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.

Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for Cisco UCS B-Series and Cisco UCS C-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-Broadwell.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter UCS-CPU-E52660E as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.

Create CPU/Cores Qualifications

Processor Architecture : PID (RegEx) :

Min Number of Cores : Unspecified select Max Number of Cores : Unspecified select

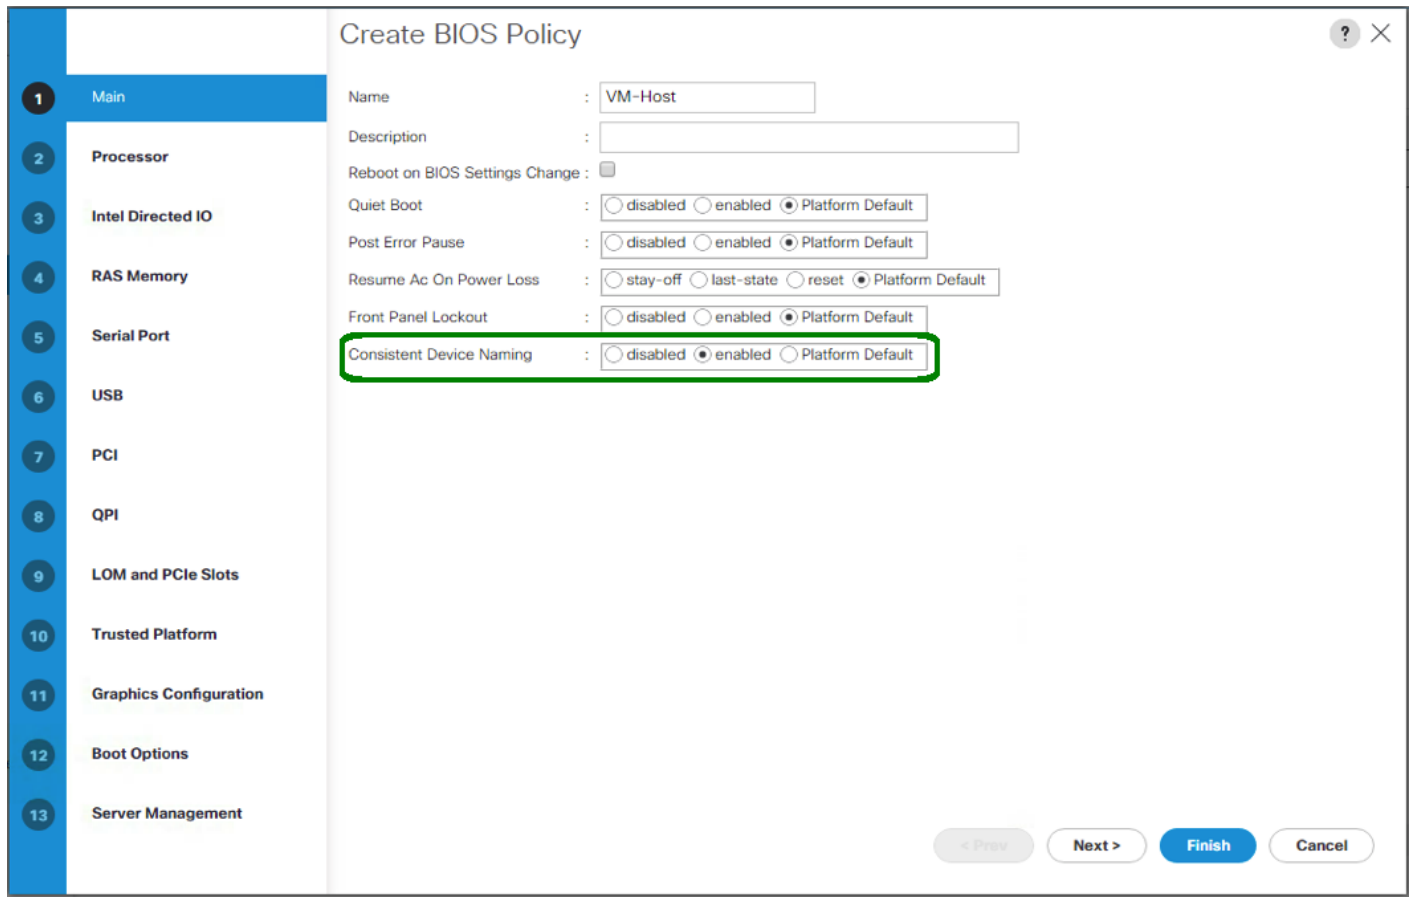
Min Number of Threads : Unspecified select Max Number of Threads : Unspecified select

CPU Speed (MHz) : Unspecified select CPU Stepping : Unspecified select


Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `MS-Host` as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



8. Click on the Processor tab on the left.
9. Set the following within the Processor tab
10. Processor C State -> disabled
11. Processor C1E -> disabled
12. Processor C3 Report -> disabled
13. Processor C7 Report -> disabled

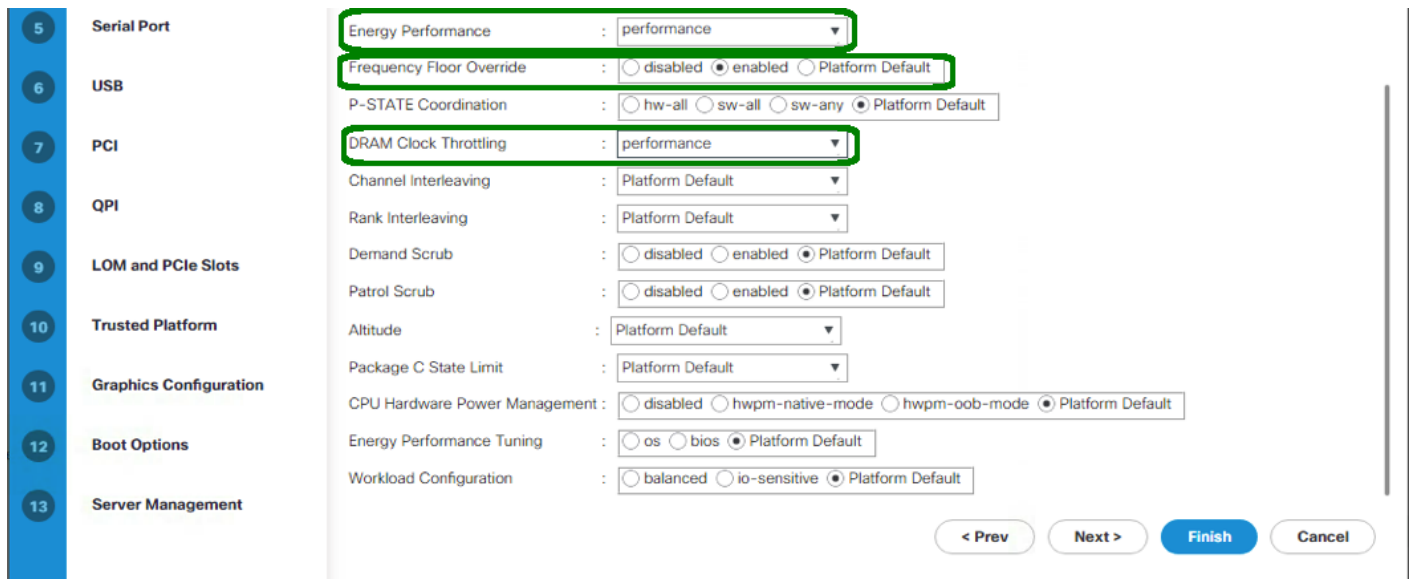


14. Scroll down to the remaining Processor options, and select:

15. Energy Performance -> performance

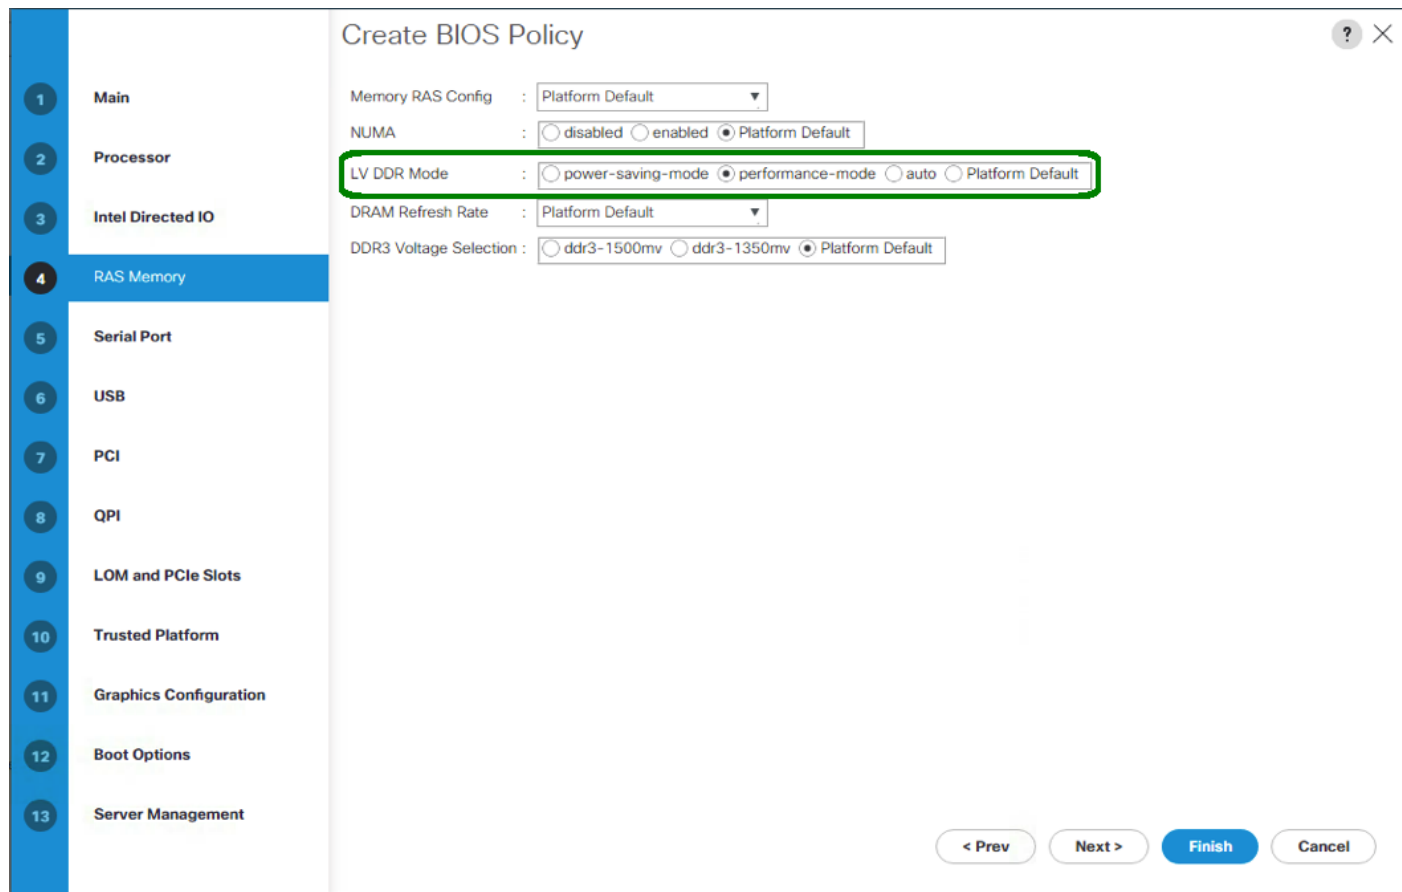
16. Frequency Floor Override -> enabled

17. DRAM Clock Throttling -> performance



18. Click on the RAS Memory option, and select:

19. LV DDR Mode -> performance-mode



20. Click Finish to create the BIOS policy.

21. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select "On Next Boot" to delegate maintenance windows to server administrators.

General	Events
Actions Delete Show Policy Usage Use Global	Properties Name : default Description : <input type="text"/> Owner : Local Soft Shutdown Timer : <input type="text" value="150 Secs"/> Reboot Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic <input checked="" type="checkbox"/> On Next Boot (Apply pending changes at next reboot.)

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

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 2 vNIC Templates will be created.

Create Infrastructure vNICs

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Host-A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for MS-IB-MGMT, MS-Cluster, MS-CSV, MS-SMB1, MS-SMB2, and MS-Tenant-VM VLANs.

13. Set Native-VLAN as the native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-POOL-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



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

Advanced Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	MS-Cluster	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-CSV	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-B	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-LVMN	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-SMB-1	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-SMB-2	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Tenant-VM	<input type="radio"/>

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 :

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template Infra-B:

1. Select LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter `Host-B` as the vNIC template name.
6. Select Fabric B.
7. Do not elect the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select Infra-A for the Peer Redundancy Template.
10. In the MAC Pool list, select `MAC-POOL-B`. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Click OK to create the vNIC template.
12. Click OK.

Create LAN Connectivity Policy for FC Boot

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following 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.
5. Enter `FC-Boot` as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter `00-Host-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Host-A.
10. In the Adapter Policy list, select Windows.
11. Click OK to add this vNIC to the policy.

Create vNIC

Name :

Use vNIC Template:

Redundancy Pair:

vNIC Template:

Peer Name:

Create vNIC Template

Adapter Performance Profile

Adapter Policy :

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-Host-B` as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select Host-B.
16. In the Adapter Policy list, select Windows.
17. Click OK to add the vNIC to the policy.

Name : **FC-Boot**

Description:

Owner : **Local**

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

Name	MAC Address	Native VLAN
▶ vNIC 00-Host-A	Derived	
▶ vNIC 01-Host-B	Derived	

18. Click OK, then OK again to create the LAN Connectivity Policy.

Create FC Boot Policy

This procedure applies to a Cisco UCS environment in which two Fibre Channel logical interfaces (LIFs) are on cluster node 1 (fcp_lifo1a and fcp_lifo1b) and two Fibre Channel LIFs are on cluster node 2 (fcp_lifo2a and fcp_lifo2b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).

One boot policy is configured in this procedure. The policy configures the primary target to be fcp_lifo1a.

To create a boot policy for the Cisco UCS environment, complete the following steps:

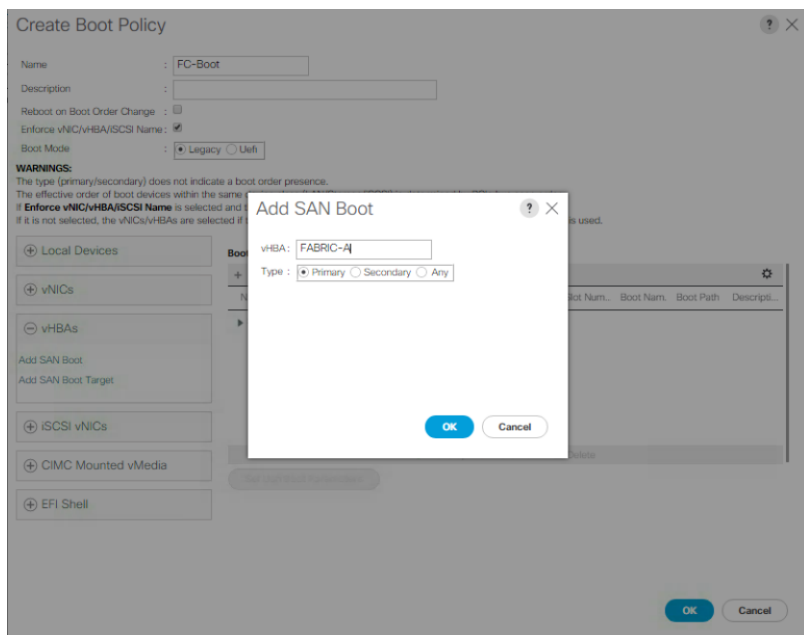
1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `FC-Boot` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.



Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Local CD/DVD`.
9. Expand the vHBAs drop-down menu and select `Add SAN Boot`.

10. Select the Primary for type field.
11. Enter FABRIC-A in vHBA field.

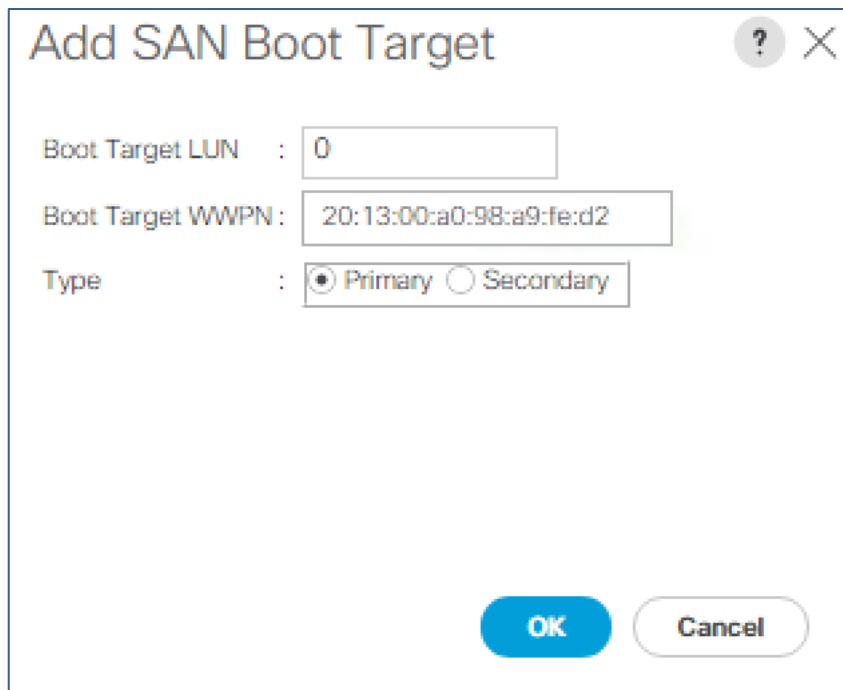


12. Click OK.
13. From the vHBA drop-down menu, select Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for fcp_lifo1a



To obtain this information, log in to the storage cluster and run the network interface show command

16. Select Primary for the SAN boot target type.



The image shows a dialog box titled "Add SAN Boot Target" with a question mark icon and a close button (X) in the top right corner. The dialog contains three input fields: "Boot Target LUN" with the value "0", "Boot Target WWPN" with the value "20:13:00:a0:98:a9:fe:d2", and "Type" with radio buttons for "Primary" (selected) and "Secondary". At the bottom right, there are two buttons: "OK" (blue) and "Cancel" (white with blue border).

17. Click OK to add the SAN boot target.
18. From the vHBA drop-down menu, select Add SAN Boot Target.
19. Enter 0 as the value for Boot Target LUN.
20. Enter the WWPN for fcp_lifo2a.
21. Click OK to add the SAN boot target.
22. From the vHBA drop-down menu, select Add SAN Boot.
23. In the Add SAN Boot dialog box, enter FABRIC-B in the vHBA box.
24. The SAN boot type should automatically be set to Secondary.
25. Click OK to add the SAN boot.
26. From the vHBA drop-down menu, select Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for fcp_lifo1b.
29. Select Primary for the SAN boot target type.
30. Click OK to add the SAN boot target.
31. From the vHBA drop-down menu, select Add SAN Boot Target.
32. Keep 0 as the value for Boot Target LUN.

33. Enter the WWPN for fcp_lifo2b.
34. Click OK to add the SAN boot target. Click OK, then click OK again to create the boot policy.

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

- 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
 - Add Local CD/DVD
 - Add Remote CD/DVD
- Add Floppy
 - Add Local Floppy
 - Add Remote Floppy
- Add Remote Virtual Drive

Boot Order

Name	O...	vNIC/vHBA/IS...	Type	WWN	LUN ...	Slot ...	Boot ...	Boot ...	Desc...
CD/DVD	1								
▼ San	2								
▶ SAN Primary		FABRIC-A	Primary						
▶ SAN Secondary		FABRIC-B	Secondary						

Create Service Profile Templates

In this procedure, one service profile template for Infrastructure Hyper-V hosts is created for Fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
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 `Hyper-V-Host-FC` 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 ? 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.

< Prev Next > **Finish** Cancel

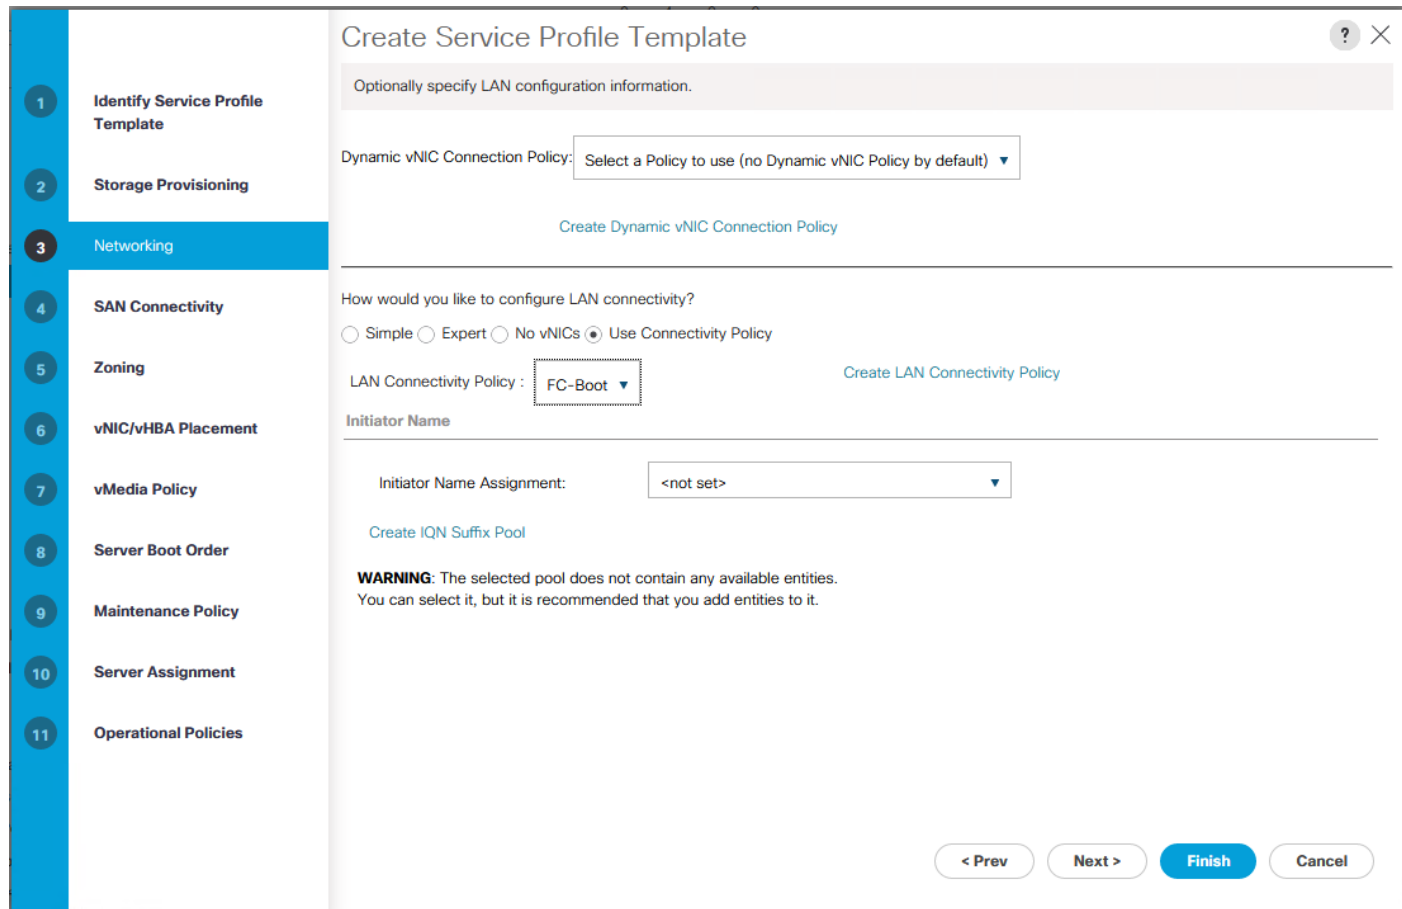
8. Click Next.

Configure Storage Provisioning

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

Configure Networking Options

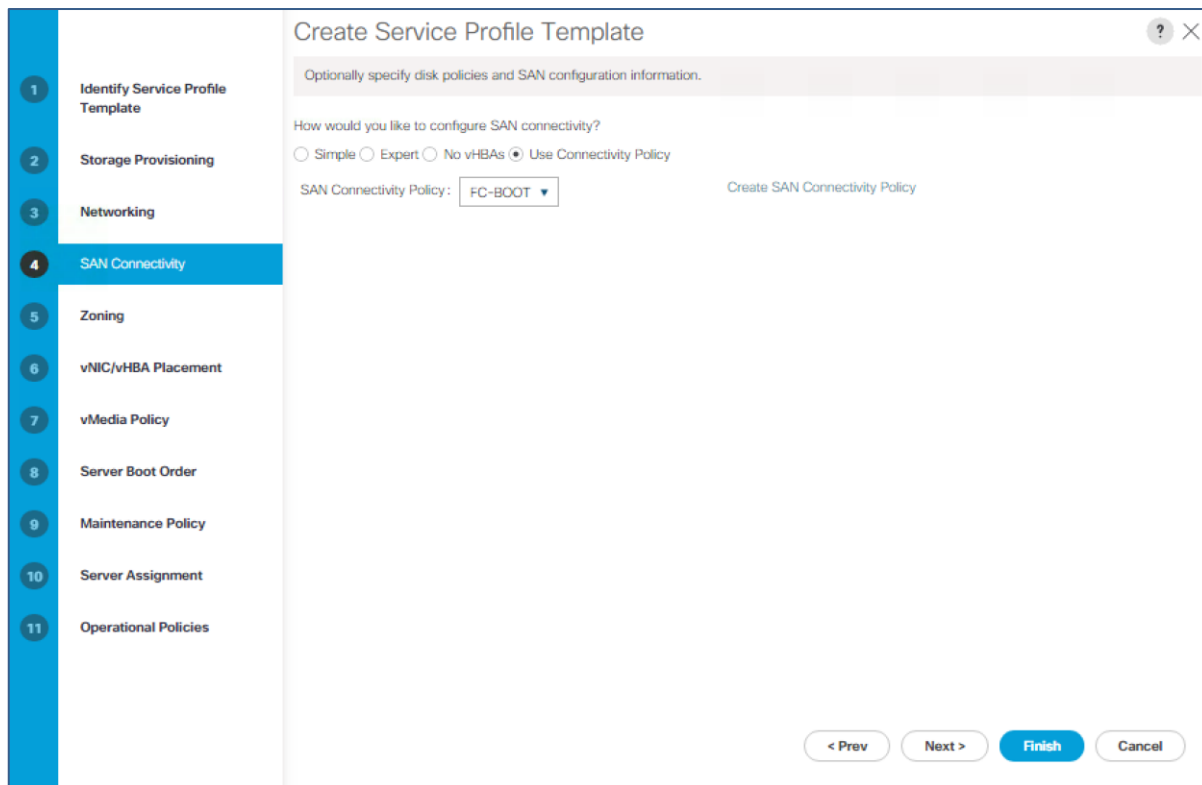
1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the "Use Connectivity Policy" option to configure the LAN connectivity.
3. Select FC-Boot from the LAN Connectivity Policy pull-down.
4. Leave Initiator Name Assignment at <not set>.



5. Click Next.

Configure Storage Options

1. Select the Use Connectivity Policy option for the "How would you like to configure SAN connectivity?" field.
2. Select the FC-BOOT option from the SAN Connectivity Policy pull-down.



3. Click Next.

Configure Zoning Options

1. 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 FC-Boot for Boot Policy.

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 the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **FC-Boot**

Description :

Reboot on Boot Order Change : **No**

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

Boot Mode : **Legacy**

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/SCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/SCSI Name** is selected and the vNIC/vHBA/SCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA...	Type	WWN	LUN Name	Slot Number	Boot Name	Boot Path	Description
CD/DVD	1								
San	2								

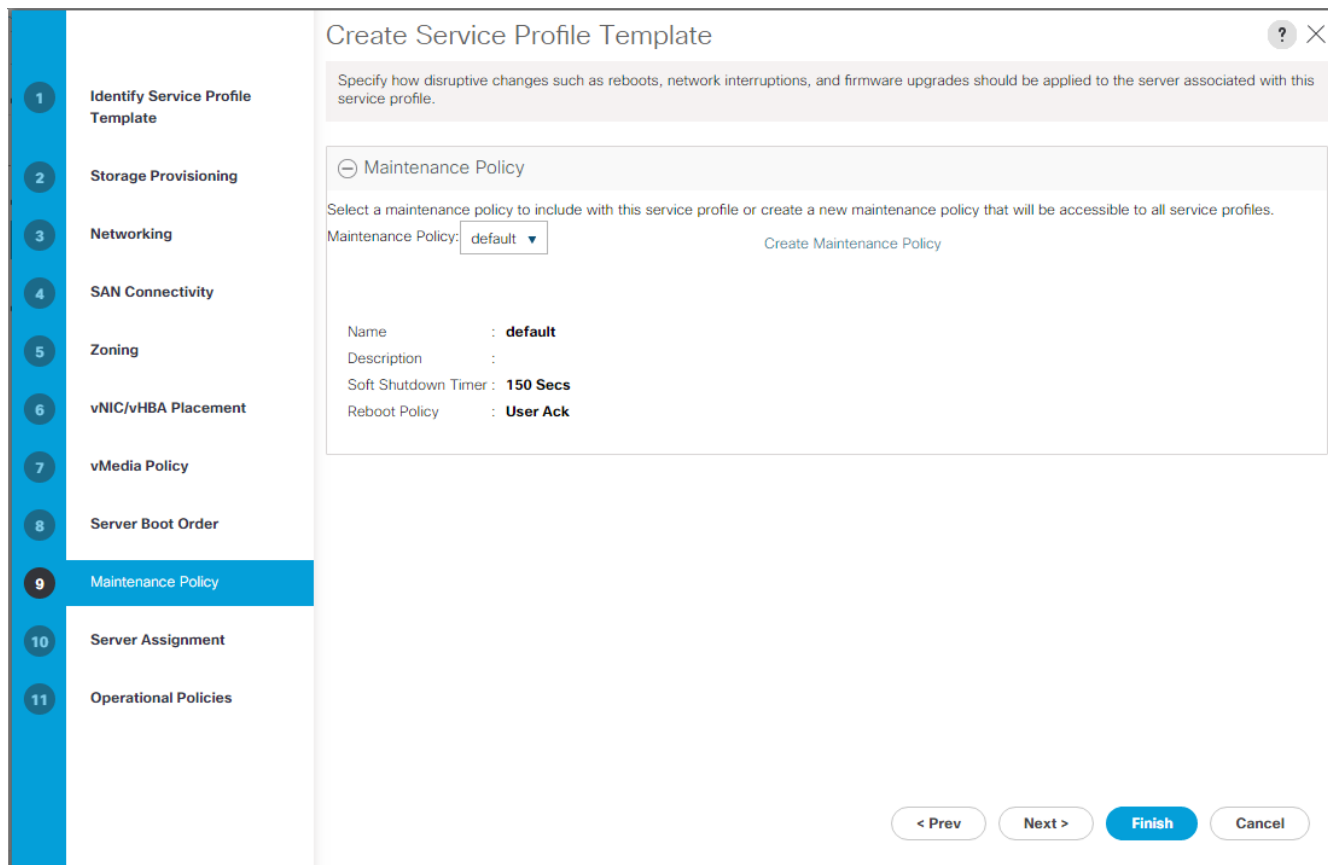
[Create iSCSI vNIC](#) [Set iSCSI Boot Parameters](#) [Set UEFI Boot Parameters](#)

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

2. Click Next.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.



2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select MS-Server-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select "UCS-Broadwell" for the Server Pool Qualification.
4. Expand Firmware Management at the bottom of the page and select the default policy

5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template ? ×

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

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

BIOS Policy:

+ External IPMI 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

< Prev Next > **Finish** Cancel

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

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template Hyper-V-Host-FC.
3. Right-click Hyper-V-Host-FC and select Create Service Profiles from Template.
4. Enter `Hyper-V-Host-0` as the service profile prefix.
5. Enter 1 as "Name Suffix Starting Number."
6. Enter 2 as the "Number of Instances."
7. Click OK to create the service profiles.

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :



- Click OK in the confirmation message.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 6 and Table 7 below.

Table 6 WWPNS from NetApp storage

SVM	Adapter	MDS Switch	Target: WWPNS
Infra-MS-SVM	fcplifo1a	Fabric A	<fcplifo1a-wwpn>
	fcplifo1b	Fabric B	<fcplifo1b-wwpn>
	fcplifo2a	Fabric A	<fcplifo2a-wwpn>
	fcplifo2b	Fabric B	<fcplifo2b-wwpn>



To obtain the FC WWPNS, run the network interface show command on the storage cluster management interface.

Table 7 WWPNS for UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPNS
--------------------------------	------------	-----------------

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
Hyper-V-Host-01	Fabric A	Hyper-V-Host -01-wwpna
	Fabric B	Hyper-V-Host -01-wwpnb
Hyper-V-Host -02	Fabric A	Hyper-V-Host -02-wwpna
	Fabric B	Hyper-V-Host -02-wwpnb



To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "Storage" tab, then "vHBAs" tab on the right. The WWPNs are displayed in the table at the bottom of the page.

SAN Switch Configuration

This section provides a detailed procedure for configuring the Cisco MDS 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

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

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod as covered in the section "FlexPod Cabling."

FlexPod Cisco MDS Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod environment. This procedure assumes the use of the Cisco MDS 9148s with NX-OS

Set Up Initial Configuration

Cisco MDS 9148S A

To set up the initial configuration for the Cisco MDS A switch, <mds-A-hostname>, complete the following 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 using the command line.

```
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> Enter
```

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

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

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

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

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

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 port mode F

in range (<100-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: 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 : <switch-a-ntp-ip>

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

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

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]: yes

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

2. Review the configuration.

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

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

Cisco MDS 9148S B

To set up the initial configuration for the Cisco MDS B switch, <mds-B-hostname>, complete the following 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 using the command line.

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

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

Server Configuration

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> Enter

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

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

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

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

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

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 port mode F
in range (<100-500>/default), where default is 500. [d]: Enter

Enable the http-server? (yes/no) [y]: 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 : <<var_global_ntp_server_ip>>

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

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

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]: yes

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

2. Review the configuration.

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

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

FlexPod Cisco MDS Switch Configuration

Enable Licenses

Cisco MDS 9148S A and Cisco MDS 9148S B

To enable the correct features on the Cisco MDS switches, complete the following steps:

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

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

Configure Individual Ports

Cisco MDS 9148S A

To configure individual ports and port-channels for switch A, complete the following step:



In this step and in further sections, configure the <ucs-6248-clustername> and <ucs-6332-clustername> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface fc1/1  
switchport description <st-node01>:0e  
switchport trunk mode off  
port-license acquire  
no shut  
  
interface fc1/2  
switchport description <st-node02>:0e  
switchport trunk mode off  
port-license acquire  
no shutdown  
exit
```

Server Configuration

```
interface fc1/9
switchport description <ucs-6248-clustername>-a:1/31
port-license acquire
channel-group 110
no shutdown
exit
```

```
interface fc1/10
switchport description <ucs-6248-clustername>-b:1/31
port-license acquire
channel-group 110
no shutdown
exit
```

```
interface fc1/11
switchport description <ucs-6332-clustername>-a:1/1
port-license acquire
channel-group 112
no shutdown
exit
```

```
interface fc1/12
switchport description <ucs-6332-clustername>-b:1/1
port-license acquire
channel-group 112
no shutdown
exit
```

```
interface port-channel110
channel mode active
switchport mode F
switchport trunk allowed vsan <vsan-a-id>
```

```
switchport description <ucs-6248-clustername>
switchport rate-mode dedicated

interface port-channel112
channel mode active
switchport mode F
switchport trunk allowed vsan <vsan-a-id>
switchport description <ucs-6332-clustername>
switchport rate-mode dedicated
```

Cisco MDS 9148S B

To configure individual ports and port-channels for switch B, complete the following step:

From the global configuration mode, run the following commands:

```
interface fc1/1
switchport description <st-node01>:0f
switchport trunk mode off
port-license acquire
no shut

interface fc1/2
switchport description <st-node02>:0f
switchport trunk mode off
port-license acquire
no shutdown
exit

interface fc1/9
switchport description <ucs-6248-clustername>-a:1/32
port-license acquire
channel-group 111
no shutdown
exit
```

Server Configuration

```
interface fc1/10
switchport description <ucs-6248-clustername>-a:1/32
port-license acquire
channel-group 111
no shutdown
exit
```

```
interface fc1/11
switchport description <ucs-6332-clustername>-a:1/2
port-license acquire
channel-group 113
no shutdown
exit
```

```
interface fc1/12
switchport description <ucs-6332-clustername>-a:1/2
port-license acquire
channel-group 113
no shutdown
exit
```

```
interface port-channel111
channel mode active
switchport mode F
switchport trunk allowed vsan <vsan-b-id>
switchport description <ucs-6248-clustername>
switchport rate-mode dedicated
```

```
interface port-channel113
channel mode active
switchport mode F
switchport trunk allowed vsan <vsan-b-id>
```

```
switchport description <ucs-6332-clustername>  
switchport rate-mode dedicated
```

Create VSANs

Cisco MDS 9148S A

To create the necessary VSANs for fabric A and add ports to them, complete the following 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 fcl/1  
vsan <vsan-a-id> interface fcl/2  
vsan <vsan-a-id> interface port-channel110  
vsan <vsan-a-id> interface port-channel112
```

Cisco MDS 9148S B

To create the necessary VSANs for fabric A and add ports to them, complete the following 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 fcl/1  
vsan <vsan-b-id> interface fcl/2  
vsan <vsan-b-id> interface port-channel111  
vsan <vsan-b-id> interface port-channel113
```


Create Device Aliases

Cisco MDS 9148S A

To create device aliases for Fabric A that will be used to create zones, complete the following steps:

From the global configuration mode, run the following commands:

```
configure terminal
device-alias database
device-alias name Infra-MS-SVM-fcp_lif01a pwwn <fcp_lif01a-wwpn>
device-alias name Infra-MS-SVM-fcp_lif02a pwwn <fcp_lif02a-wwpn>
device-alias name Hyper-V-Host-01-A pwwn <vm-host-infra-01-wwpna>
device-alias name Hyper-V-Host-02-A pwwn <vm-host-infra-02-wwpna>
device-alias commit
```

Cisco MDS 9148S B

To create device aliases for Fabric B that will be used to create zones, complete the following steps:

From the global configuration mode, run the following commands:

```
configure terminal
device-alias database
device-alias name Infra-MS-SVM-fcp_lif01b pwwn <fcp_lif01b-wwpn>
device-alias name Infra-MS-SVM-fcp_lif02b pwwn <fcp_lif02b-wwpn>
device-alias name Hyper-V-Host-01-B pwwn <vm-host-infra-01-wwpnb>
device-alias name Hyper-V-Host-02-B pwwn <vm-host-infra-02-wwpnb>
device-alias commit
```

Create Zones

Cisco MDS 9148S A

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

```
configure terminal
zone name Hyper-V-Host-01-A vsan <vsan-a-id>
member device-alias Hyper-V-Host-01-A init
member device-alias Infra-MS-SVM-fcp_lif01a target
member device-alias Infra-MS-SVM-fcp_lif02a target
exit
zone name Hyper-V-Host-02-A vsan <vsan-a-id>
```

```
member device-alias Hyper-V-Host-02-A init
member device-alias Infra-MS-SVM-fcp_lif01a target
member device-alias Infra-MS-SVM-fcp_lif02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member Hyper-V-Host-01-A
member Hyper-V-Host-02-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
exit
show zoneset active vsan <vsan-a-id>
```

Cisco MDS 9148S B

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

```
configure terminal
zone name Hyper-V-Host-01-B vsan <vsan-b-id>
member device-alias Hyper-V-Host-01-B init
member device-alias Infra-MS-SVM-fcp_lif01b target
member device-alias Infra-MS-SVM-fcp_lif02b target
exit
zone name Hyper-V-Host-02-B vsan <vsan-b-id>
member device-alias Hyper-V-Host-02-B init
member device-alias Infra-MS-SVM-fcp_lif01b target
member device-alias Infra-MS-SVM-fcp_lif02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member Hyper-V-Host-01-B
member Hyper-V-Host-02-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active vsan <vsan-b-id>
```

Storage Configuration – Boot LUNs

NetApp ONTAP Boot Storage Setup



Disable network interface - A target network interface is disabled prior to Windows 2016 installation because Windows 2016 does not support multipathing. After multipathing is installed and enabled, the interface will be enabled.

To disable one network interface on each node, run the following commands:

```
network interface modify -vserver Infra-MS-SVM -lif fcp_lif01b -home-node bb04-affa300-1 -status-admin down
network interface modify -vserver Infra-MS-SVM -lif fcp_lif02b -home-node bb04-affa300-2 -status-admin down
```

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-01 -protocol fcp -ostype windows -initiator <vm-
host-infra-01-wwpna>,<vm-host-infra-01-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-02 -protocol fcp -ostype windows -initiator <vm-
host-infra-02-wwpna>,<vm-host-infra-02-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-All -protocol fcp -ostype windows -initiator <vm-
host-infra-01-wwpna>,<vm-host-infra-01-wwpnb>,<vm-host-infra-02-wwpna>,<vm-host-infra-02-wwpnb>
```

Map Boot LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-id 0
lun map -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -igroup VM-Host-Infra-All -lun-id
1
```

Microsoft Windows Server 2016 Hyper-V Deployment Procedure

Setting Up Microsoft Windows Server 2016

This section provides detailed instructions for installing Microsoft Windows Server 2016 in an environment. After the procedures are completed, two booted Windows Server 2016 hosts will be provisioned.

Several methods exist for installing Microsoft Windows Server 2016. 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).

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

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click Servers on the left.
7. Select Servers > Service Profiles > root > Hyper-V-Host-01.
8. Right-click Hyper-V-Host-01 and select KVM Console.
9. Follow the prompts to launch the Java-based KVM console.
10. Select Servers > Service Profiles > root > Hyper-V-Host-02.
11. Right-click Hyper-V-Host-02. and select KVM Console.
12. Follow the prompts to launch the Java-based KVM console.
13. From the virtual KVM Console, select the Virtual Media tab.
14. Select Add Image in the right pane.
15. Browse to the Windows Server 2016 installation ISO image file and click Open.
16. Map the image that you just added by selecting Mapped.
17. To boot the server, select the KVM tab.
18. Select Power On Server in the KVM interface Summary tab, and then click OK.

Install Windows Server 2016

1. The following steps describe the installation of Windows Server 2016 to each host:
2. All Hosts
3. On boot, the machine detects the presence of the Windows installation media.
4. After the installer has finished loading, Enter the relevant region information and click Next.
5. Click Install now.
6. Enter the Product Key and click Next.
7. Select Windows Server 2016 Datacenter (Server with a GUI) and click Next.



You may optionally remove the GUI after the Hyper-V cluster is operational.

8. After reviewing the EULA, accept the license terms and click Next.
9. Select Custom: Install Windows only (advanced).
10. Select Custom (advanced) installation.
11. In the Virtual Media Session manager uncheck the Mapped checkbox for the Windows ISO and select yes to confirm.
12. Click Add Image.
13. Browse to the Cisco fNIC driver ISO, click Open.
14. Check the Mapped checkbox next to the Cisco fNIC Driver ISO. Download the latest driver iso image from the cisco.com site.

Browse for driver software on your computer

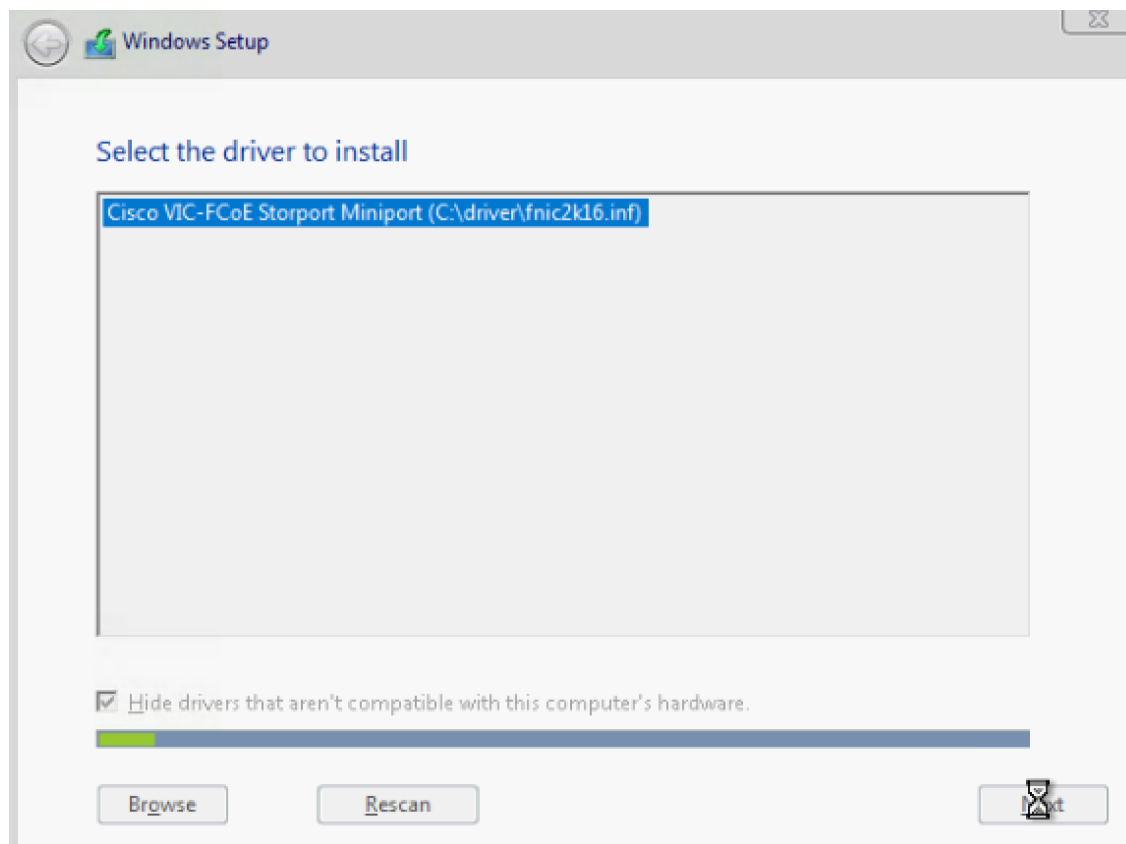
Search for driver software in this location:

F:\Network\Cisco\VIC\W2K16\x64

Browse...

Include subfolders

15. Back in the KVM Console, click Load Driver and then, click OK.
16. The Cisco VIC FCoE Storport Miniport driver should auto detected; Click Next.



17. You should see a LUN listed in the drive selection screen.




Only a single LUN instance should be displayed. Multiple instance of the same LUN indicates that there are multiple paths to the installation LUN. Verify that the SAN zoning is correct and restart the installation.



The message "Windows Can't be installed on this drive" appears because the Windows installation ISO image is not mapped at this time.



The Cisco eNIC driver can be loaded at this point in the same way as the fNIC driver. Loading the eNIC driver at this time bypasses the need to load the eNIC driver in the section titled "Installing Windows eNIC Driver".



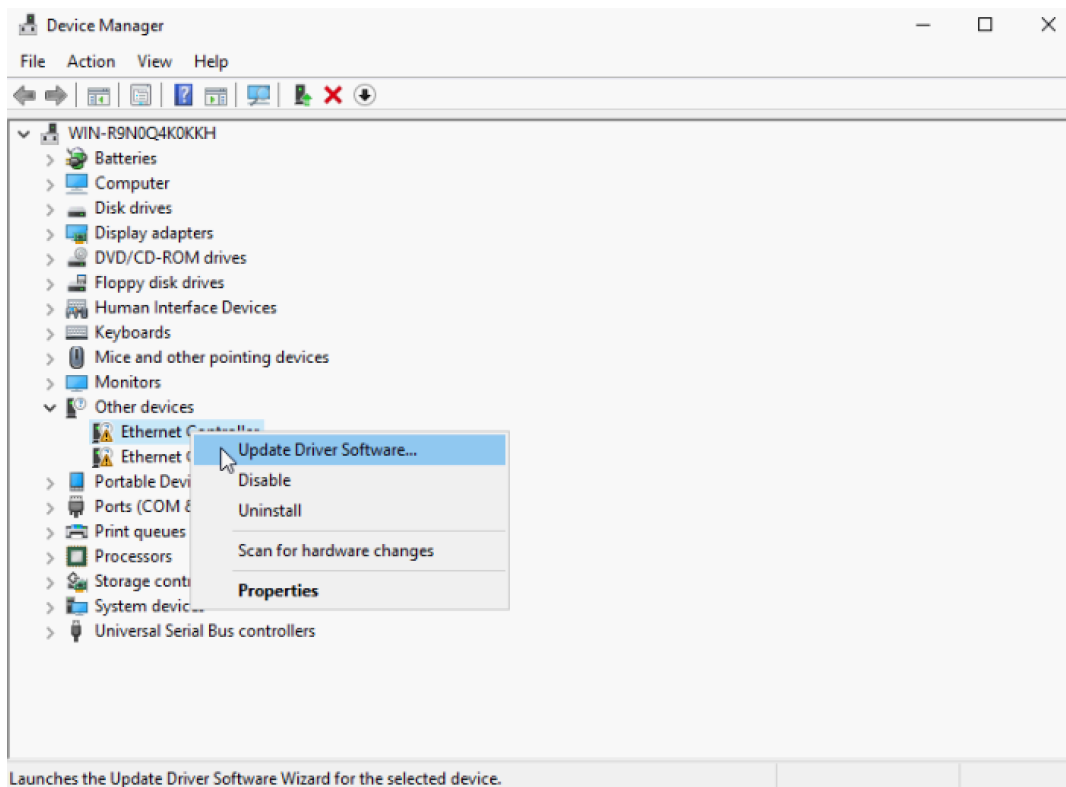
18. Select the LUN, and click Next to continue with the install.
19. When Windows is finished installing, enter an administrator password on the settings page and click Finish.

Install Chipset and Windows eNIC Drivers

This section provides detailed information on installing the Intel chipset drivers and the Cisco VIC enic drivers.

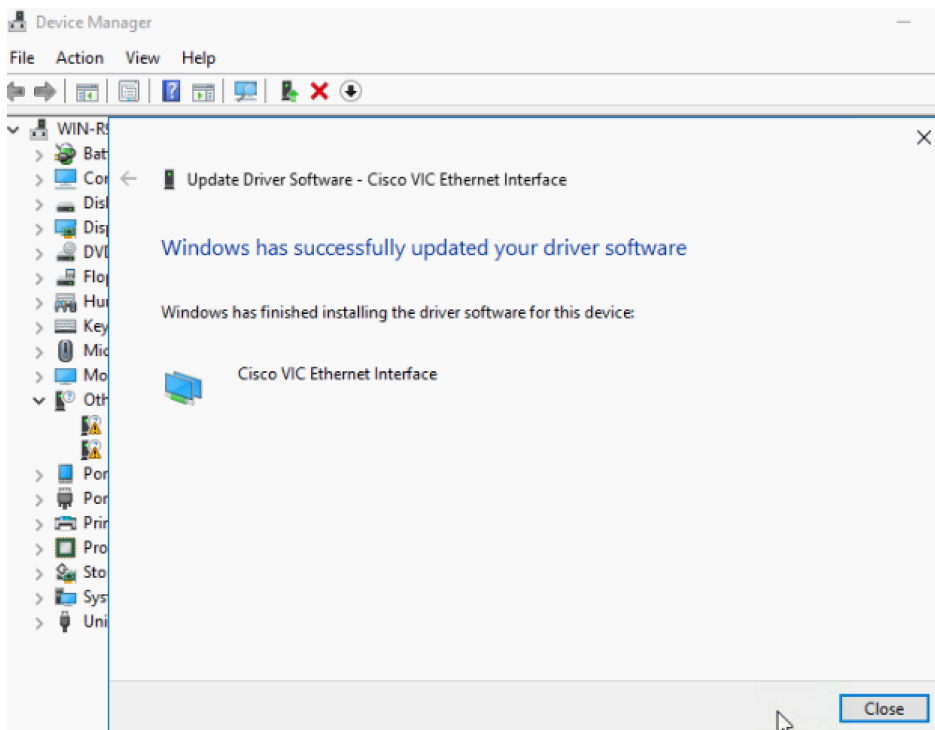
All Hosts

1. In the Virtual Media Session manager, uncheck the Mapped checkbox for the Windows ISO.
2. Click Add Image
3. Browse to the Cisco UCS driver ISO, click Open
4. Check the Mapped checkbox for the Cisco UCS driver ISO.
5. Browse to the CD ROM > Chipset > Intel > <Server Model> W2K16 > x64
6. Double click on Setup Chipset to install the chipset driver and reboot the system
7. In the KVM console, open Server Manager, and select Tools > Computer Management
8. In Computer Manager, select System Tools > Device Manager > Other devices.
9. Right-click one of the Ethernet Controller, and select Update Driver Software.

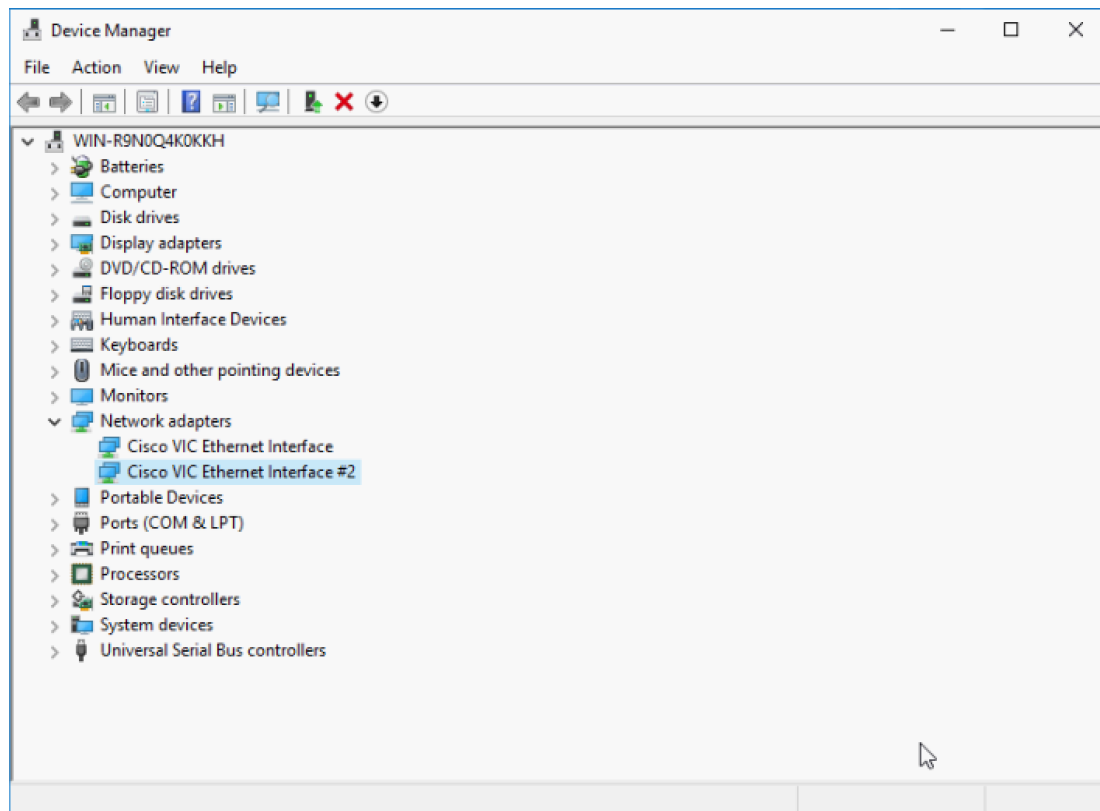


10. Click Browse, and select CDROM drive, click OK.

11. Click Next > Close.




12. Right-click the next Ethernet Controller and select Update Driver Software.
13. Click Search automatically for update driver software.
14. Click Close.
15. Repeat these steps for the remaining Ethernet Controllers.
16. All Cisco VIC Ethernet devices will appear under Network Adapters.



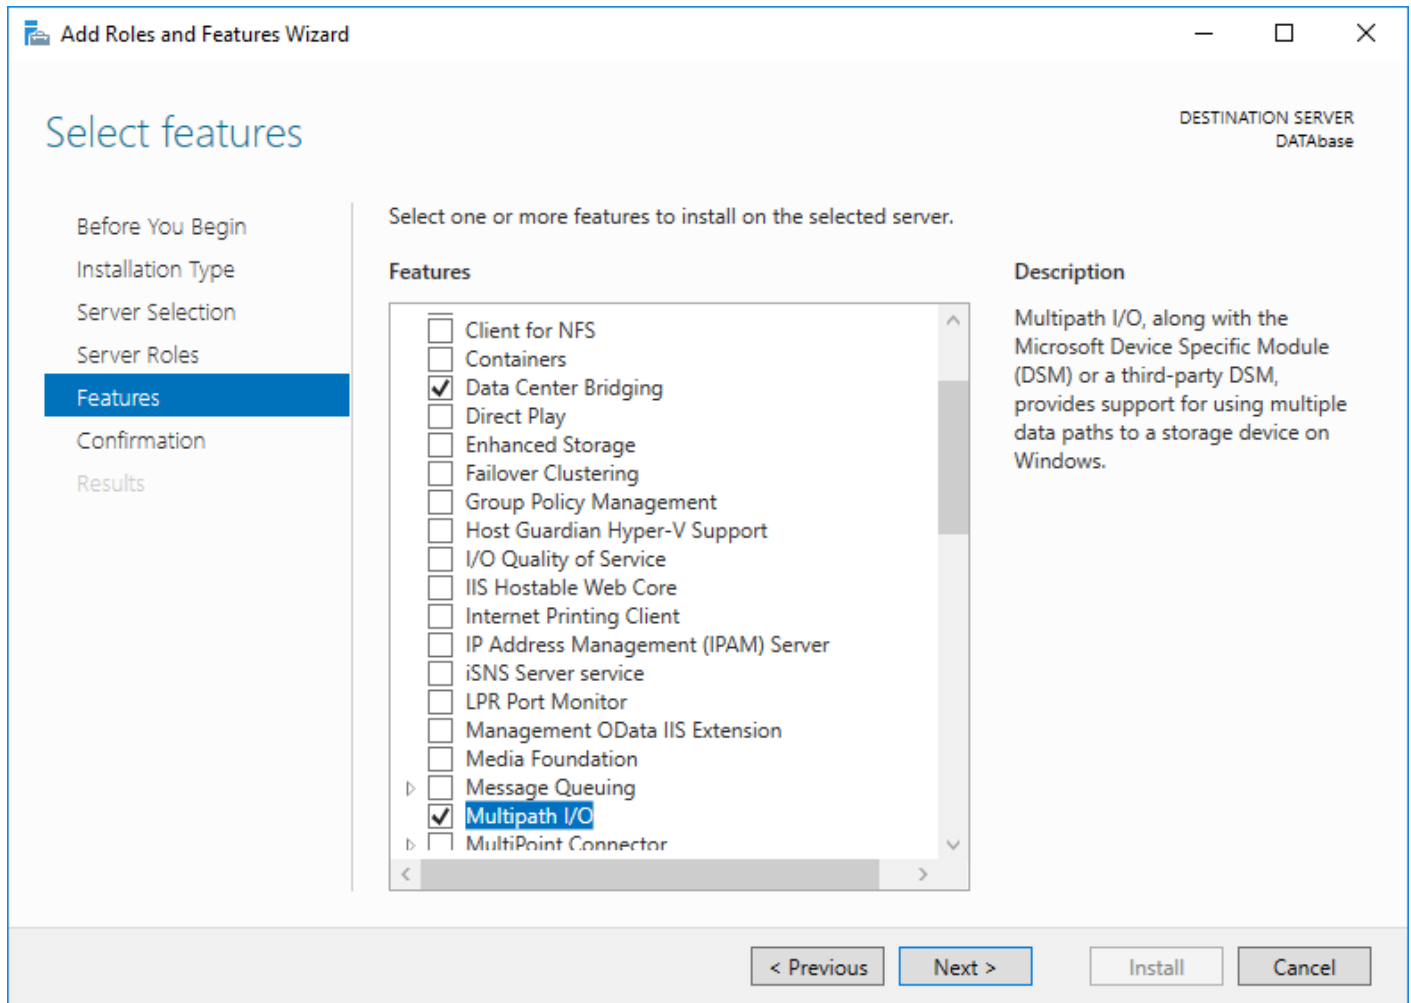
Install Windows Roles and Features

This section provides the steps to enable the MPIO feature from the Server Manager.

1. In Server Manager, select Manage > Add Roles and Features to launch the wizard.



2. Click Next in the 'Before you begin' section of the wizard to continue
3. In the 'Installation Type' section, select 'Role-based or feature-based installation' and click Next.
4. In the 'Server selection' section, select the server.
5. Click Next in the 'Server Roles' section without making any selection.
6. Select Multipath I/O and Data Center Bridging in the 'Features' section and click Next
7. Click on Install in the confirmation page to install the feature.



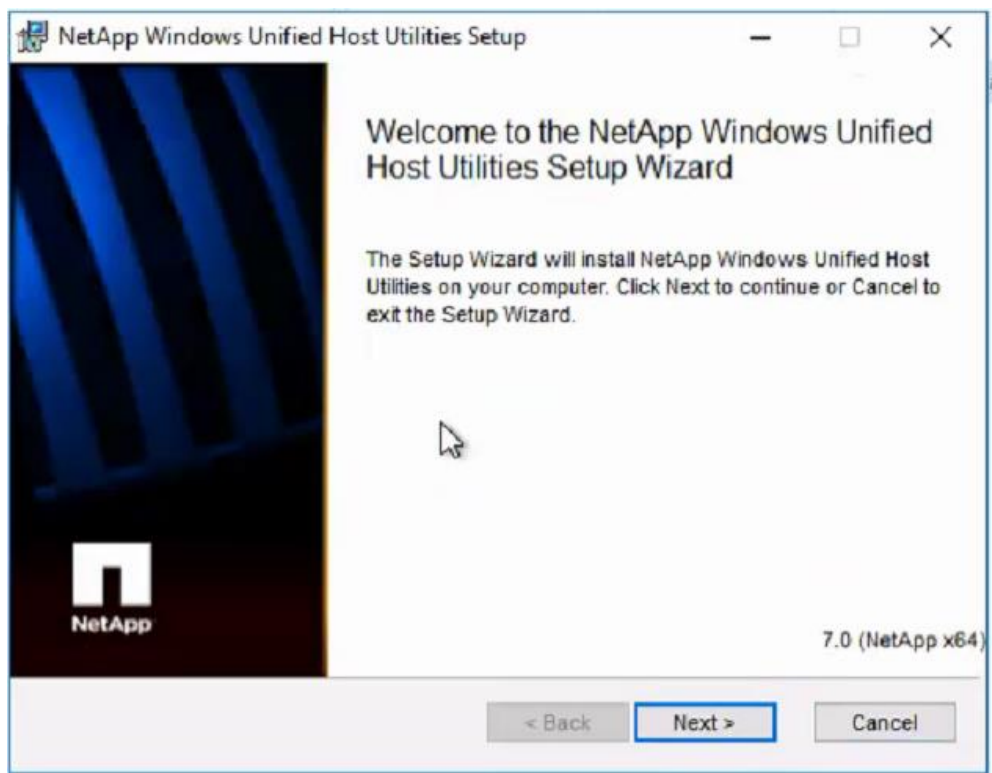
Install NetApp Host Utilities

After enabling the MPIO feature in Windows, download and install NetApp Windows Unified Host Utilities. This section provides the steps to download and install the host utilities.


1. Download the NetApp host utilities v7.0 for Windows from the link below:

<https://mysupport.netapp.com/documentation/productlibrary/index.html?productID=61343>

2. Unzip the file and run the executable file. The NetApp Windows Unified Host Utilities setup wizard is launched and Click Next.



3. Select "Yes, install support for Multipath IO" and click Next.



4. Accept the default destination folder and click Next.

5. Click Ok and Next to finish the installation of host utilities.

Host Renaming and Join to Domain

Login to the host and open PowerShell and enter the below commands to

1. Rename the host.

```
Rename-Computer -NewName <hostname> -restart
```

2. Assign an IP address to the management interface.

```
new-netipaddress -interfaceindex <UInt32> -ipaddress <string> -prefixlength <Byte> -DefaultGateway <string>
```

3. Assign DNS server IP address to the above management interface

```
Set-DnsClientServerAddress -InterfaceIndex <UInt32[]> -ServerAddresses <String[]>
```

4. Add the host to Active Directory.

```
Add-Computer -DomainName <domain_name> -Restart
```

Storage Configuration – Boot LUNs (Continued)

NetApp ONTAP Boot Storage Setup



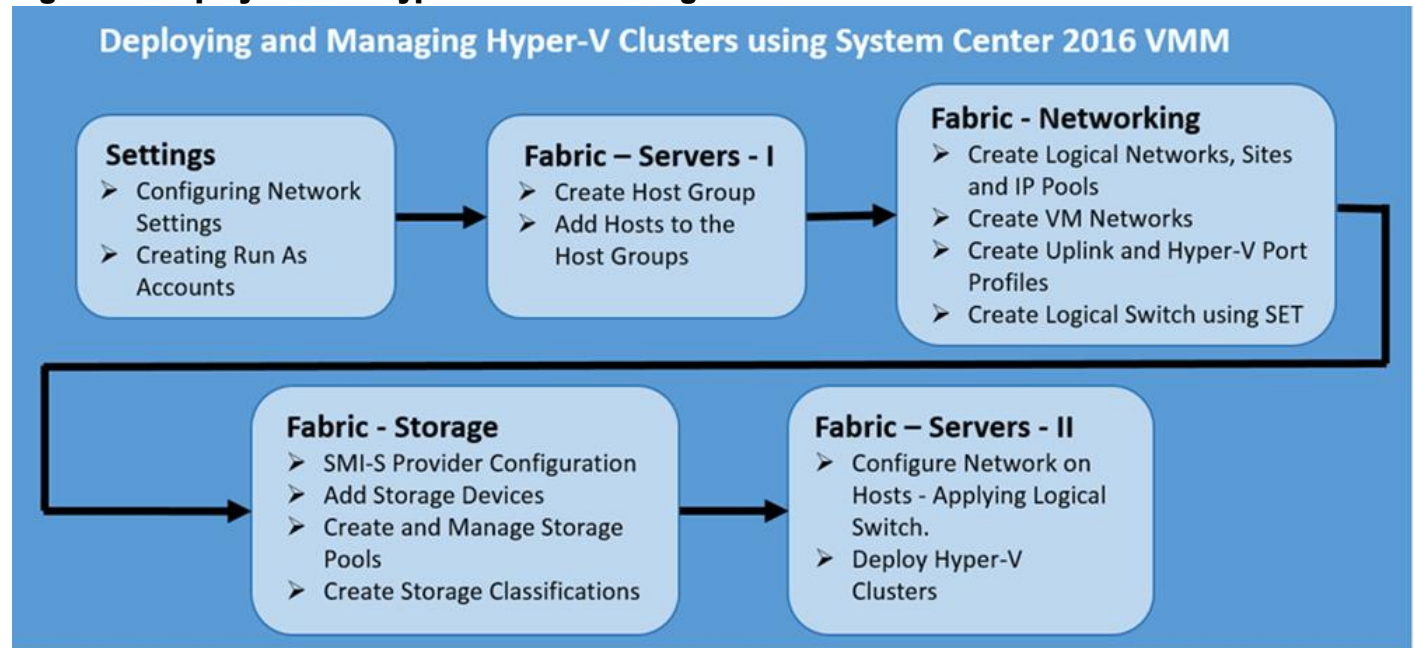
Enable All Network Interfaces - now that multipathing has been installed and enabled in Windows 2016, all network interfaces can be enabled.

```
network interface modify -vserver Infra-MS-SVM -lif fcp_lif01b -home-node bb04-affa300-1 -status-admin up
network interface modify -vserver Infra-MS-SVM -lif fcp_lif02b -home-node bb04-affa300-2 -status-admin up
```

Deploying and Managing Hyper-V Clusters using System Center 2016 VMM

For this part of the section, we assume that System Center 2016 VMM is up-and-running in your environment. This section will focus only on configuring the Networking, Storage and Servers in VMM to deploy and manage Hyper-V failover clusters. Figure 5 provides a high-level view of the steps that will be covered in detail in the following sections.

Figure 5 Deployment of Hyper-V Cluster Using SCVMM



Settings

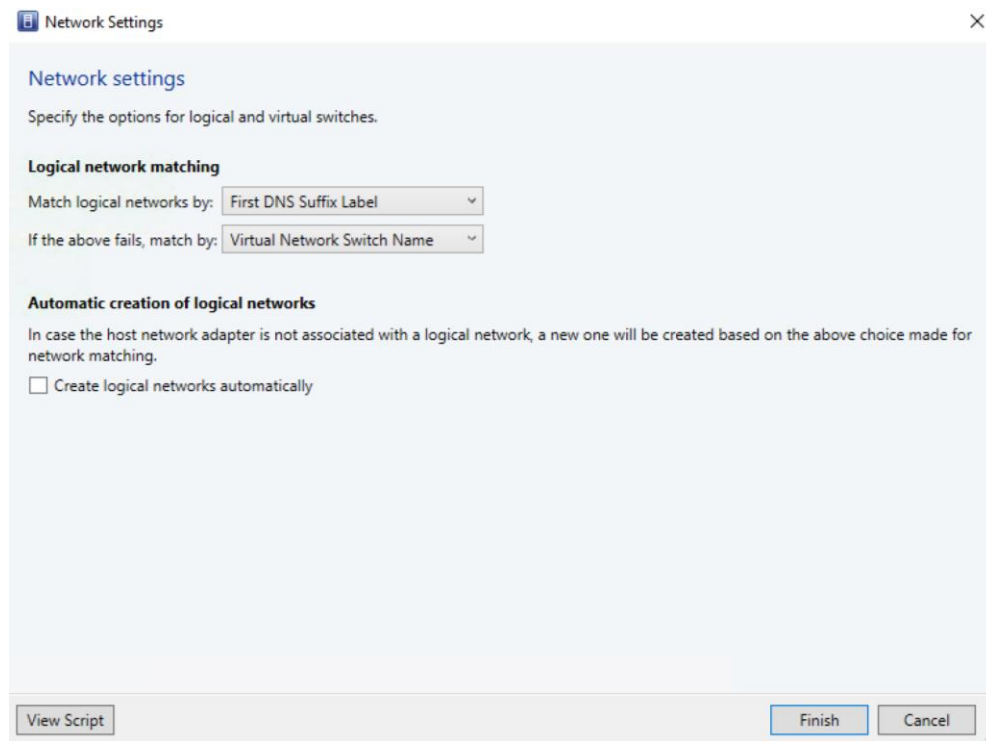
Configuring Network Settings

By default, VMM creates logical networks automatically. When you provision a host in the VMM fabric and there's no VMM logical network associated with a physical network adapter on that host, VMM automatically creates a logical network and associates it with an adapter. Follow the below procedure to disable automatic logical network creation.

1. Open Virtual Machine Manager.
2. Open the Settings workspace.
3. Select the General navigation node.
4. Double-click Network Settings in the details pane.
5. In the Network Settings dialog box, uncheck the Create Logical Networks Automatically option and click OK.



Notice also in this dialog box that you can change the logical network matching behavior to a scheme that may better suit your naming conventions and design.



Create Run As accounts in VMM

A Run As account is a container for a set of stored credentials. In VMM a Run As account can be provided for any process that requires credentials. Administrators and Delegated Administrators can create Run As accounts. For this deployment, a Run As account should be created for adding Hyper-V hosts and integrating NetApp SMI-S provider.

1. Click Settings, and in Create click Create Run As Account.
2. In Create Run As Account specify name and optional description to identify the credentials in VMM.
3. In User name and Password specify the credentials. The credentials can be a valid Active Directory user or group account, or local credentials.
4. Clear Validate domain credentials if you don't need it, and click OK to create the Run As account.

Fabric – Servers - I

This section covers:

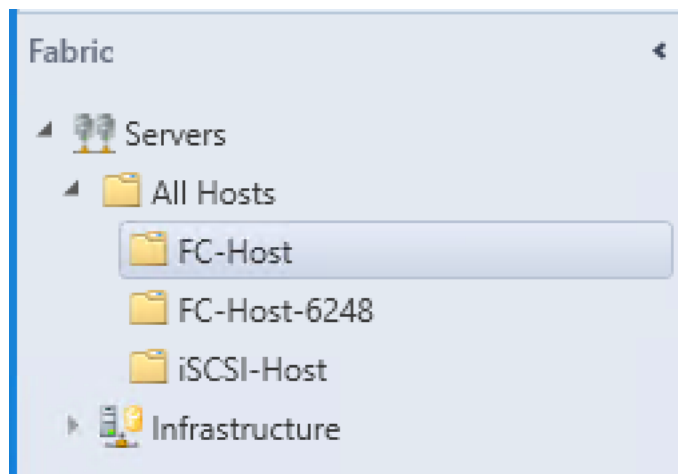
- Create Host Groups
- Add Windows Hosts to the Host Group

Create Host Groups

You can use host groups to group virtual machine hosts in meaningful ways, often based on physical site location and resource allocation.

Follow the below procedures to create a host group structure in Virtual Machine Manager (VMM) that aligns to your organizational needs.

1. To create a host group structure.
2. Open the Fabric workspace.
3. In the Fabric pane, expand Servers, and then do either of the following:
4. Right-click All Hosts, and then click Create Host Group.
5. Click All Hosts. On the Folder tab, in the Create group, click Create Host Group. VMM creates a new host group that is named New host group, with the host group name highlighted.
6. Type a new name, and then press ENTER.
7. Repeat the steps in this procedure to create the rest of the host group structure.



Add Hosts to the Host Group

Once the virtual switch is created, you can add the Hyper-V hosts to Virtual Machine Manager:

1. Open the Fabric workspace.
2. Select a host group, and On the Home tab, in the Add group, click Add Resources, and then click Hyper-V Hosts and Clusters. The Add Resource Wizard starts.
3. On the Resource location page, click Windows Server computers in a trusted Active Directory domain, and then click Next.
4. On Credentials page, select Use an Run As account, click Browse and add the Run as account created earlier. Click Next.
5. On Discovery scope, select Specify Windows Server computers by names and enter the Computer names. Click Next.
6. Under Target Resources, select the check box next to the computer names that needs to be the part of the Hyper-V cluster.



If the Hyper-V role is not enabled on a selected server, you receive a message that VMM will install the Hyper-V role and restart the server. Click OK to continue.

7. On the Host settings page, In the Host group list, click the host group to which you want to assign the host or host cluster.
8. On the Summary page, confirm the settings, and then click Finish.

Add Resource Wizard X

+ **Summary**

Resource Location

Credentials

Discovery Scope

Target Resources

Host Settings

Summary

Confirm the settings

View Script

Resource type: Hyper-V capable Windows Servers

Resource location: Trusted Windows computer

Discovery credentials: flexpod\administrator

Discovery scope: Computer name based discovery
2 computers are selected to manage

Host settings: Host group:
All Hosts\FC-Host

Previous
Finish
Cancel

Administrator - MS-SCVMM.flexpod.local - Virtual Machine Manager

+ Home Folder

Create
 Add Resources
 Overview
 Fabric Resources
 Compliance
 Cisco UCS Manager
 Scan
 Remediate
 Compliance Properties
 Update Agent
 Reassociate Agent
 Connect via RDP
 PowerShell
 Jobs
 PRO

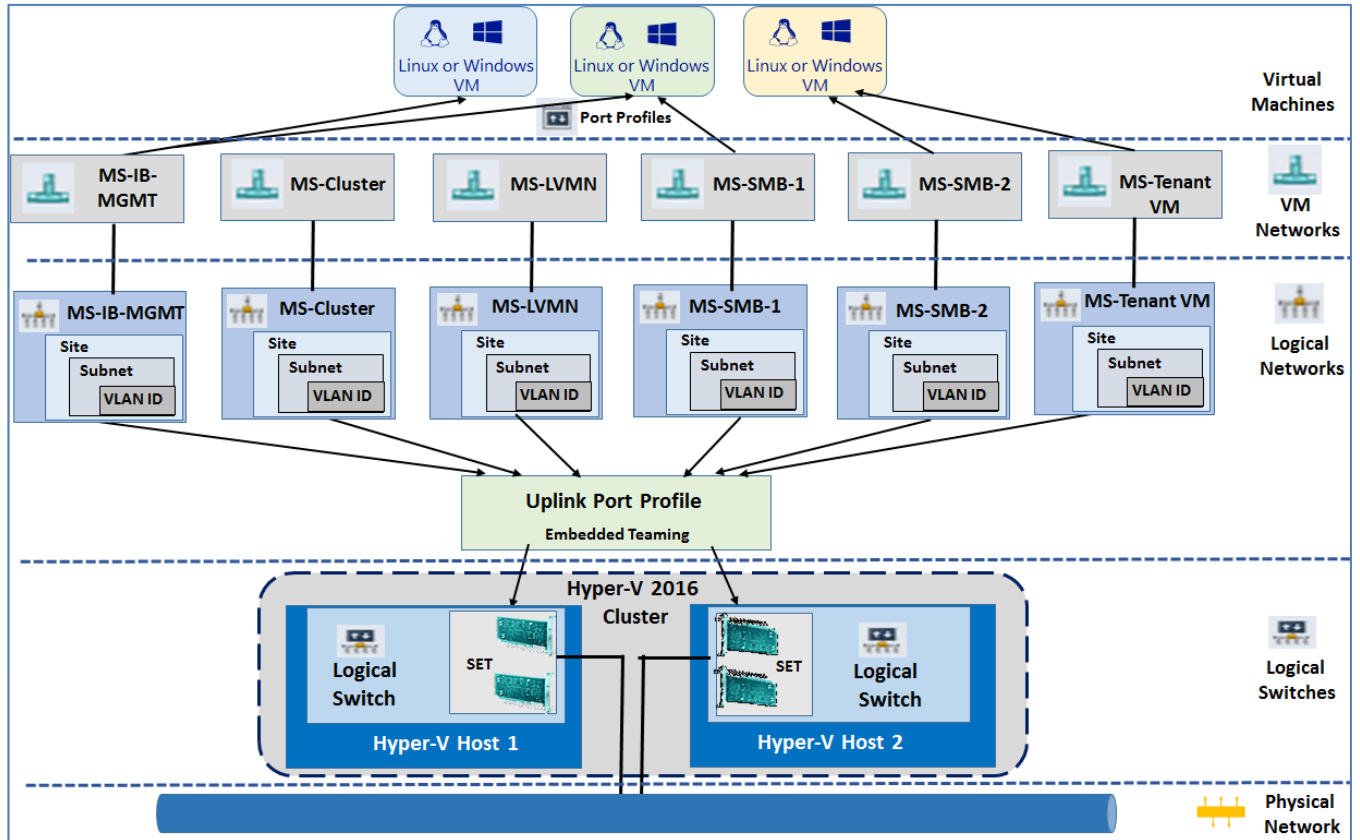
Fabric

- Servers
- All Hosts
 - FC-Host
 - FC-Host-6248

Hosts (2)

Name	Host...	Role	Job S...	CPU Aver...	Available Me...	Operating System
hv-fc-01.flexpod.local	OK	Host	Completed	4 %	242.92 GB	Microsoft Windows Server 2016 Datacenter Evaluation
hv-fc-02.flexpod.local	OK	Host	Completed	0 %	230.89 GB	Microsoft Windows Server 2016 Datacenter Evaluation

Fabric – Networking

Figure 6 SCVMM Logical Network

The above figure shows the logical representation of the network that will be configured in this section using the System Center 2016 VMM and applied later to configure the network settings of Hyper-V hosts before deploying the failover cluster. For this document, we are going to use and deploy "Switch Embedded Teaming (SET); a new feature released in Windows server 2016. SET is a new teaming solution integrated with the Hyper-V switch.

The topics that will be covered in this Networking section are:

- Create Logical Networks, Sites and IP Pools
- Create VM Networks
- Create Uplink and Hyper-V Port Profiles
- Create Logical Switch using SET

Creating Logical Networks, Sites and IP Pools

In this environment, we have six networks available that we will model as logical networks. However, they are all separate VLANs on the same physical network that will be controlled by setting the VLAN ID on the virtual network adapter. The physical ports on the switch have been configured to allow all of the various VLANs that can be configured (similar to a trunk port):

- **MS-IB-MGMT:** This logical network will be used for management traffic and has its own IP subnet and VLAN.

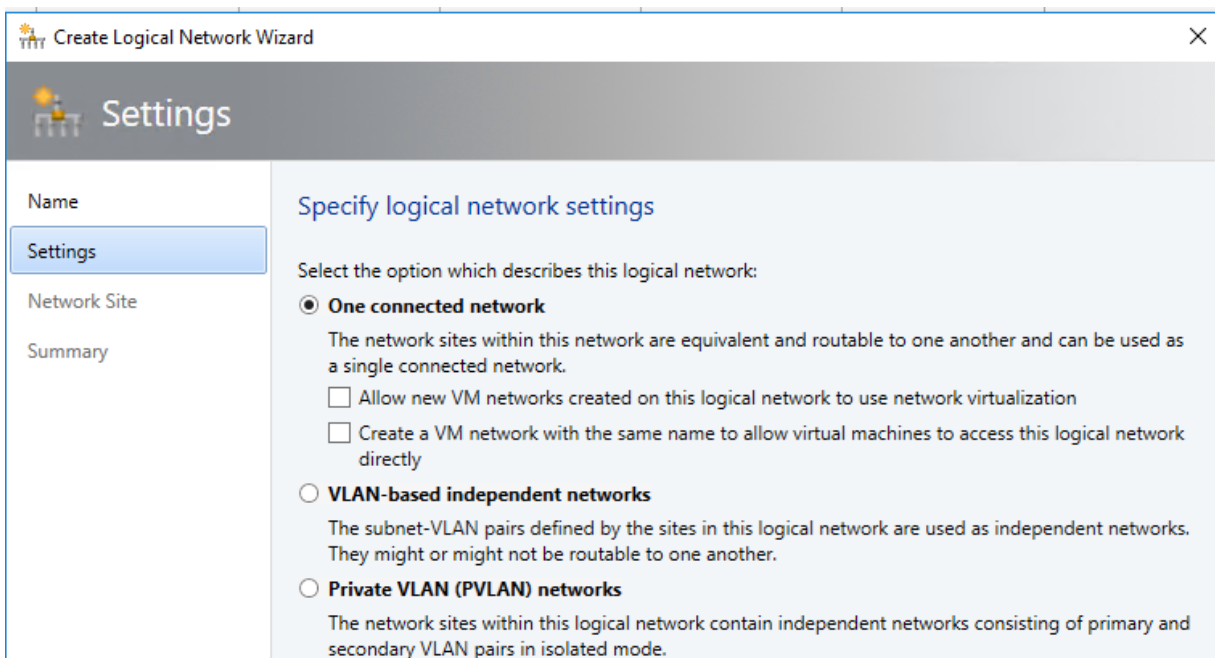
- **MS-Cluster:** This network will be used for Microsoft Hyper-V cluster communication and will have its own IP subnet and VLAN.
- **MS-LVMN:** This network will be used for Live Migration traffic and will have its own IP subnet and VLAN
- **MS-SMB-1 and MS-SMB-2:** This network will be used for SMB file share access/traffic and has its own IP subnet and VLAN.
- **MS-Tenant-VM (optional):** This network will be used for all the VM traffic.

Perform the following steps to create logical networks and sites:

1. Open Virtual Machine Manager Console.
2. Open the Fabric workspace.
3. Select the Networking > Logical Networks navigation node.
4. Click the Create Logical Network button, which launches the Create Logical Network Wizard.
5. Enter a name and description for the logical network and click Next.
6. In the Settings tab, select VLAN-based independent networks.



You can select a type of network. It can be a connected network that allows multiple sites to communicate with each other and use network virtualization, a VLAN-based independent network, or a PVLAN-based network.

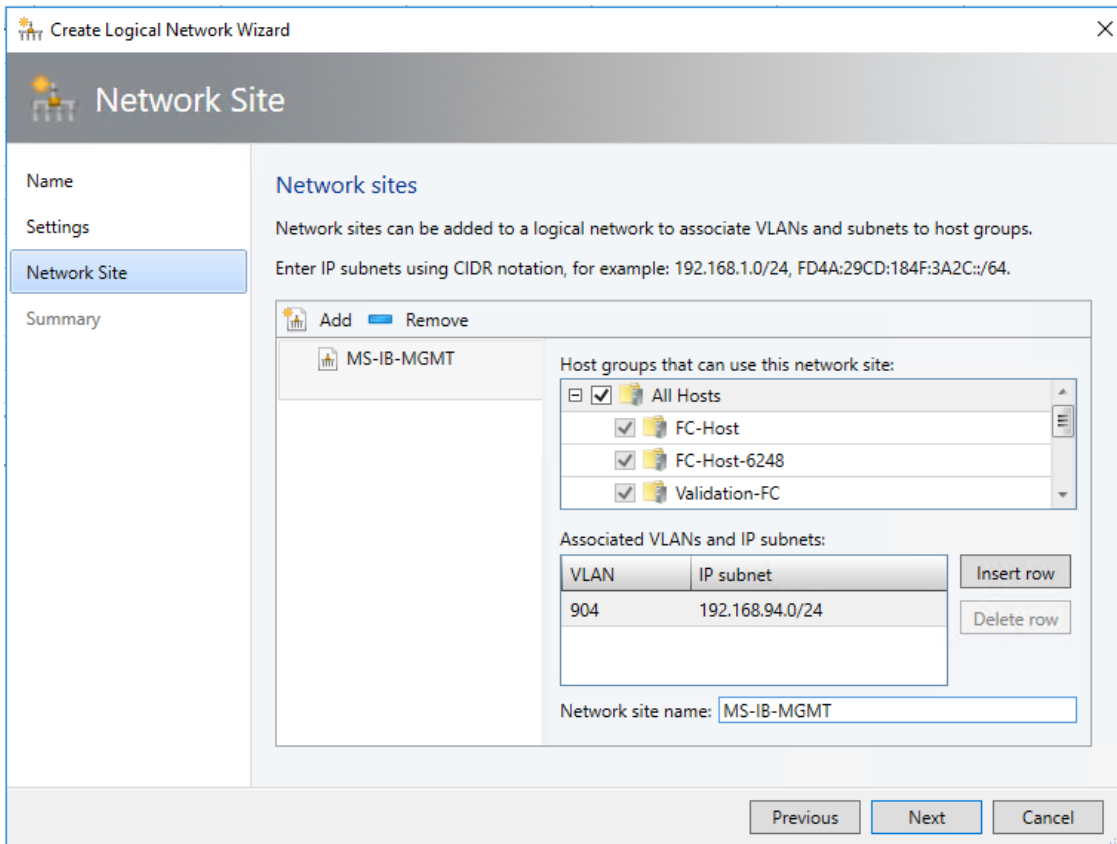


7. Select the sites and Host Group, where you want to apply the Management VLAN. Click the Add button to add a site and then click Insert Row to add VLAN/IP details for the site.

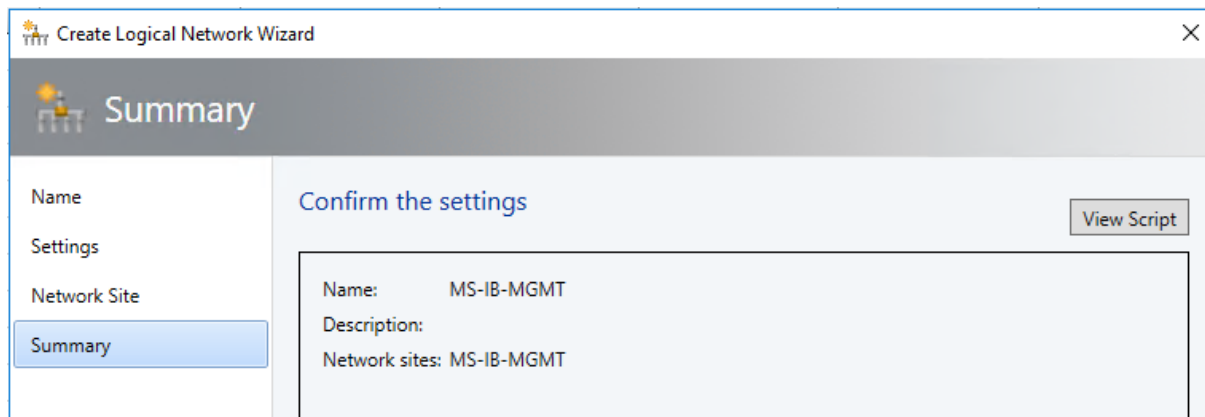


If IP space is configured by corporate DHCP servers, leave the IP subnet blank, which tells SCVMM to just configure the VM for DHCP. If the network does not use VLANs, set the VLAN ID to 0; this tells SCVMM that VLANs

are not to be configured. By default, sites are given the name <Logical Network>_<number>, but you should re-name this to something more useful.



8. The Summary screen is displayed. It includes a View Script button that when clicked shows the PowerShell code that can be used to automate the creation. This can be useful when you are creating many logical networks, or more likely, many sites.
9. Click Finish to create the logical network.



10. Follow the above steps to create all the Logical Networks for the environment. The figure below shows the all the logical networks created for this document.

Name	Network Compliance	Subnet	Begin Address	End Address
MS-Cluster	Fully compliant			
MS-IB-MGMT	Fully compliant			
MS-iSCSI-A	Fully compliant			
MS-iSCSI-B	Fully compliant			
MS-LVMN	Fully compliant			
MS-Tenant-VM	Fully compliant			
MS-SMB-1	Fully compliant			
MS-SMB-2	Fully compliant			

MS-Cluster

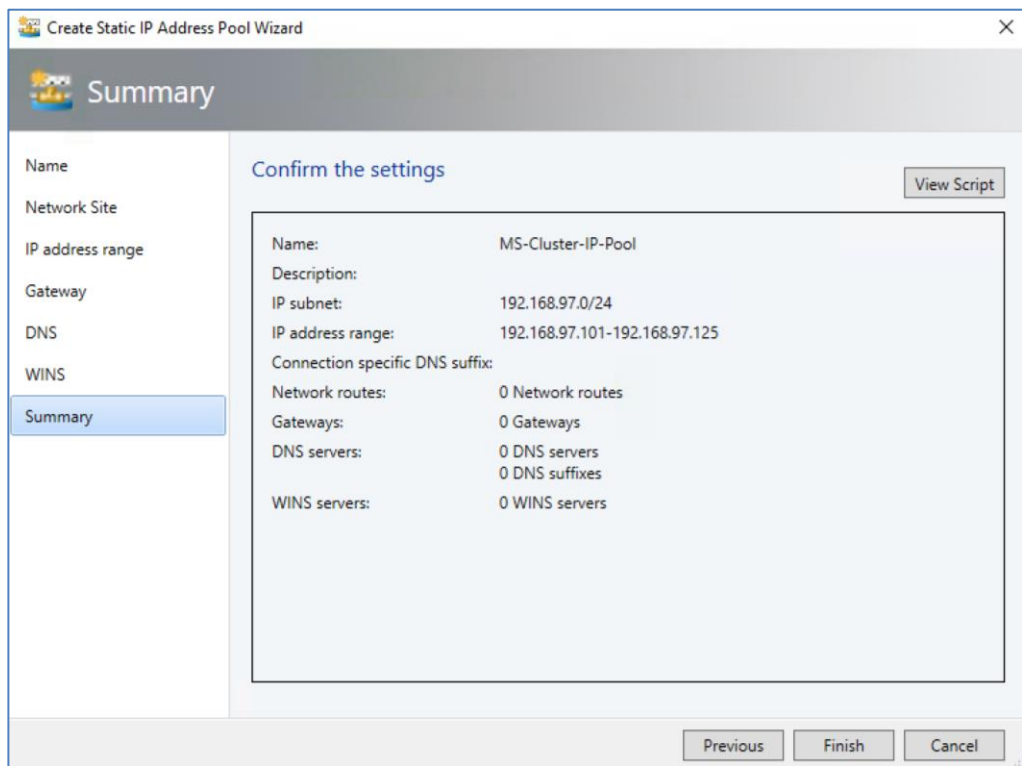
[Logical network information](#)

Description: Logical Network for Cluster Communication

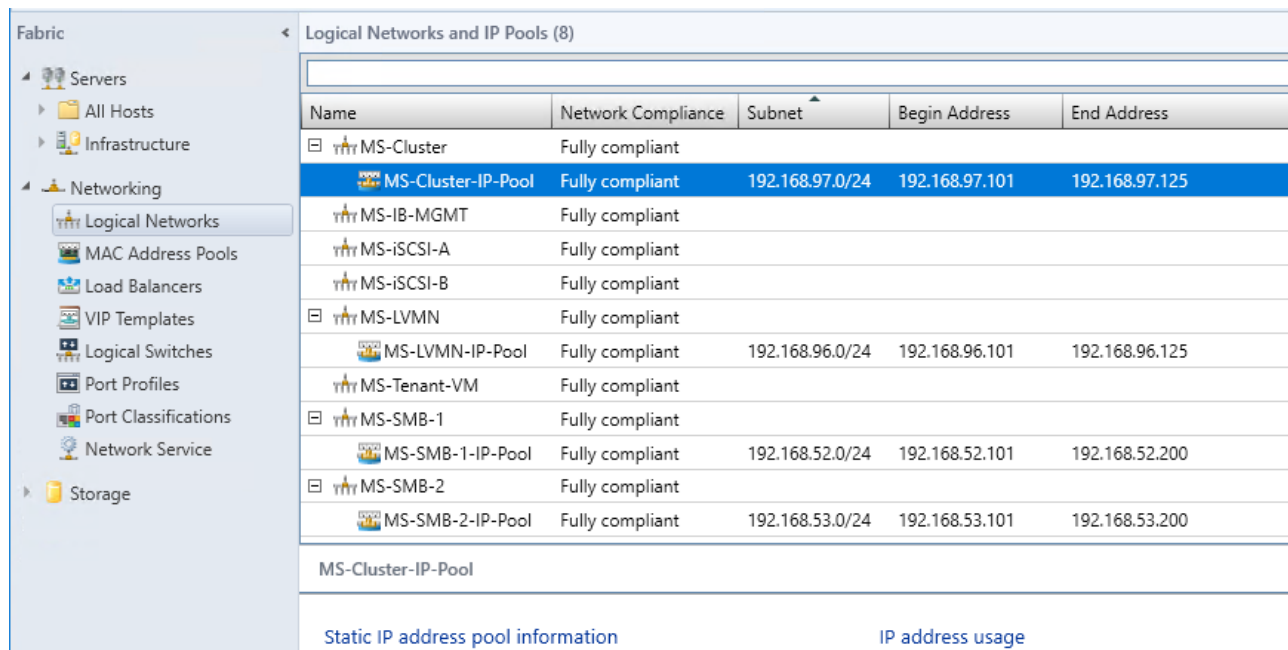
Perform the following steps to create a static IP address pool for a logical network in Virtual Machine Manager (VMM). With static IP address pools,

IP address management for the virtual environment is brought within the scope of the VMM administrator.

1. From the Fabric workspace, Click the Create IP Pool button, or right-click the logical network and select the Create IP Pool context menu action.
2. Enter a name and description. From the drop-down list, select the logical network for the IP pool.
3. The next screen, allows you to use an existing network site or create a new one. Choose to use an existing one and then click Next.
4. The next screen, allows you to use an existing network site or create a new one. Choose to use an existing one and then click Next.
5. The IP Address Range page allows configuration of the IP address range that SCVMM will manage and allocate to resources such as virtual machines and load balancers. Within the range, you can configure specific addresses to be reserved for other purposes or for use by load-balancer virtual IPs (VIPs) that SCVMM can allocate.
6. Click the Insert button, and enter the gateway IP address. Then click Next.
7. Configure the DNS servers, DNS suffix, and additional DNS suffixes to append, and then click Next.
8. Enter the WINS server details if used, and click Next.
9. On the Summary screen, confirm the configuration, click the View Script button to see the PowerShell that will be used, and then click Finish to create the IP pool.



10. Create IP Pools for all the Logical Networks as shown the figure below.

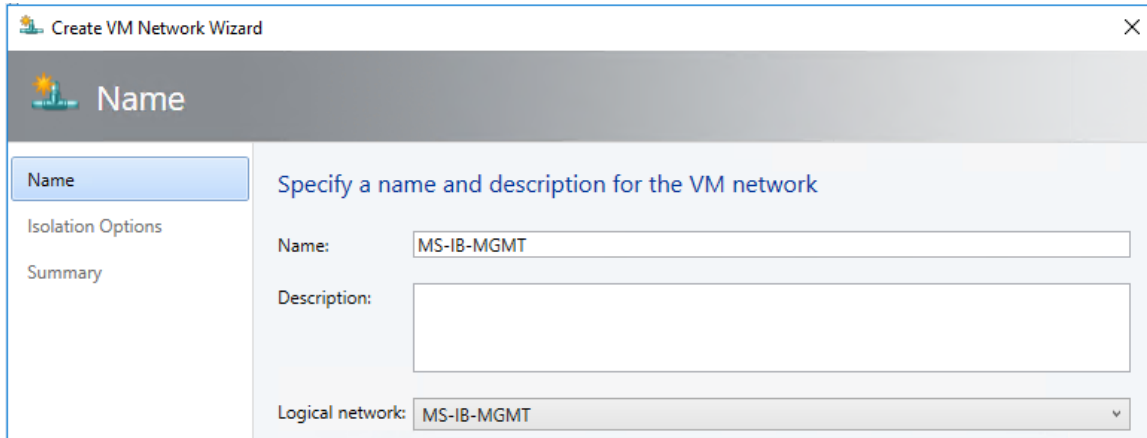


Create VM Networks

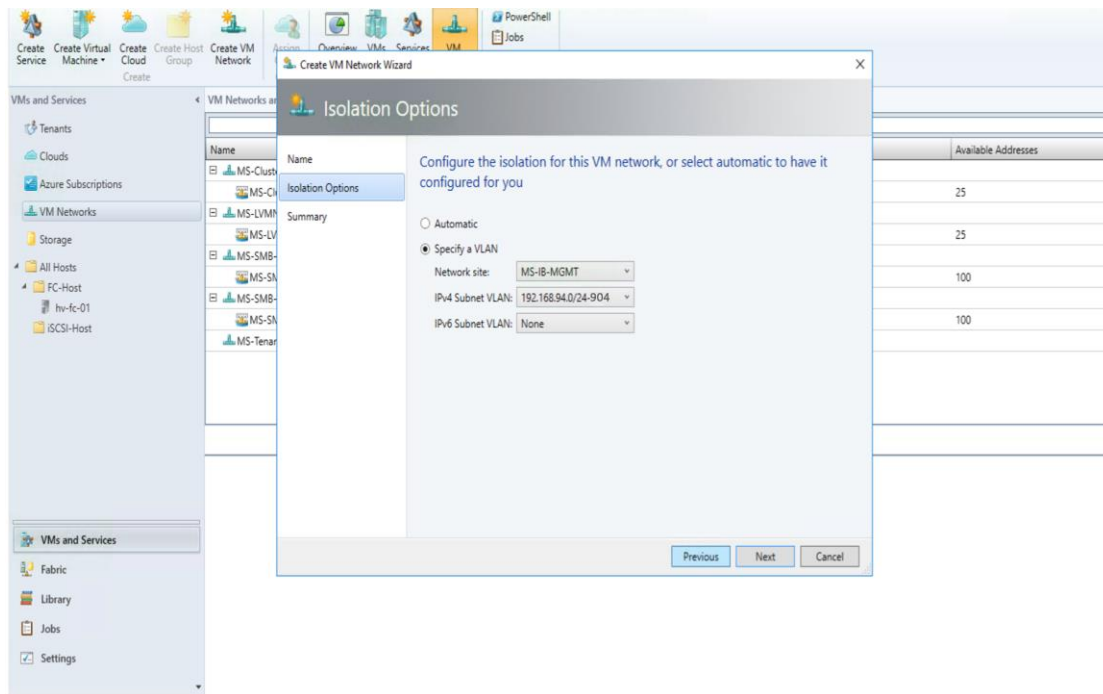
With logical networks created, the next step is to create the VM networks to which virtual machines can be connected.

1. Open Virtual Machine Manager.
2. Open the VMs and Services workspace.

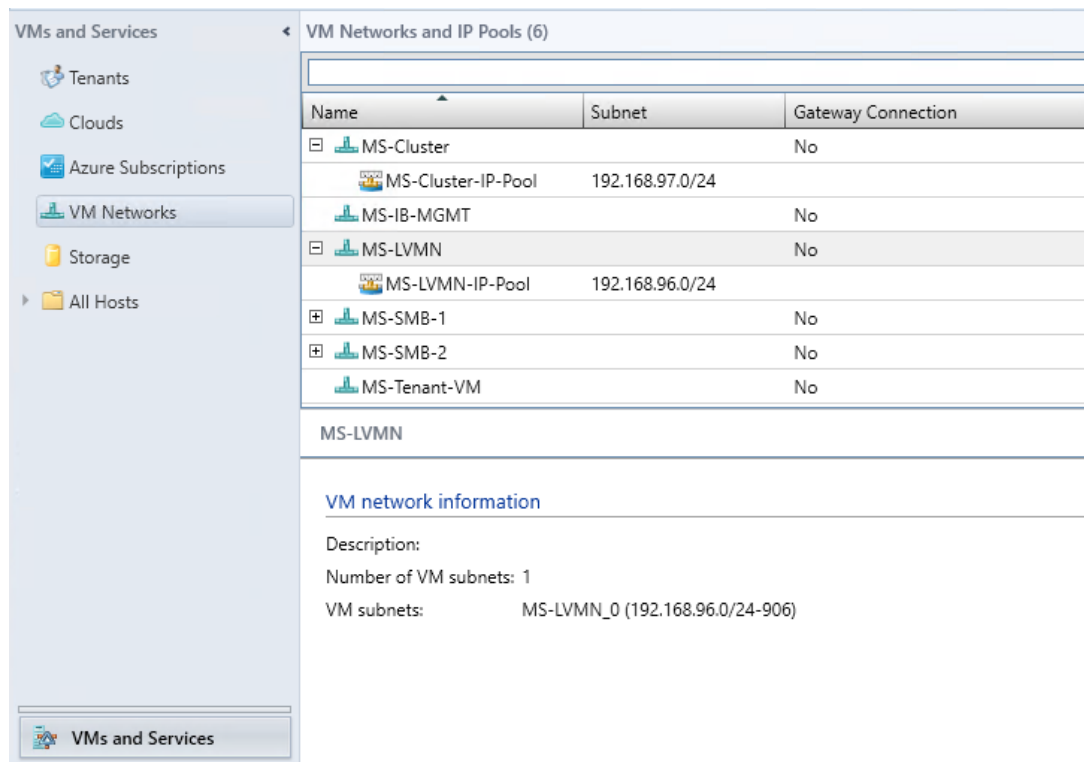
3. Select the VM Networks navigation node.
4. Click the Create VM Network button.
5. Enter a name and description for the VM network, select the logical network, and click Next.



6. In the Isolation Options screen, select Specify a VLAN and select Network Site, IPv4 Subnet VLAN and click Next



7. Click Finish to complete the VM network creation process.
8. Repeat the above steps to create all the required VM Networks.

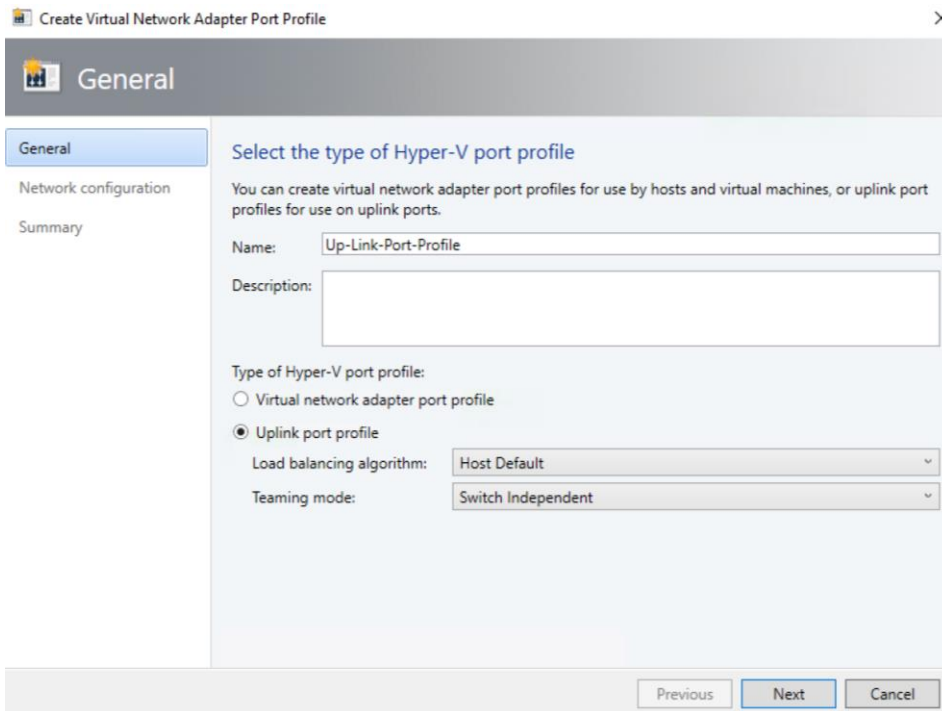


Create Uplink Port Profiles and Hyper-V Port Profiles

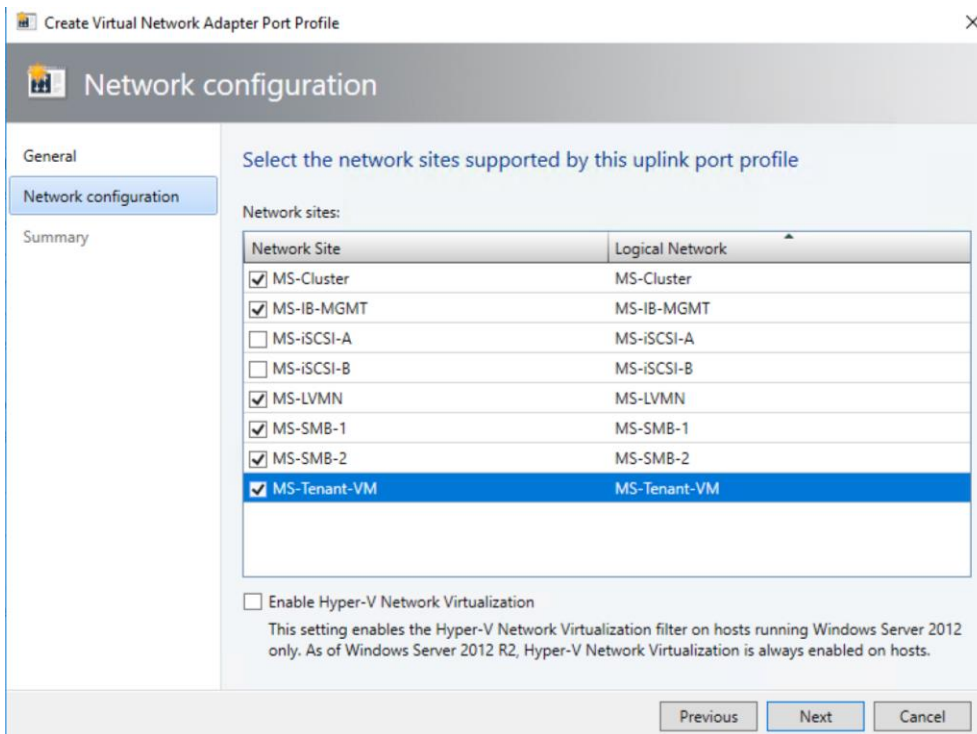
Create uplink port profile

Uplink port profiles define the load balancing algorithm for an adapter, and specify how to team multiple network adapters on a host that use the same uplink port profile. This profile is used in conjunction with the logical network that you associated with the adapter.

1. Open Virtual Machine Manager.
2. Open the Fabric workspace.
3. Select the Networking > Port Profiles navigation node.
4. Click the Create button drop-down, and select Hyper-V Port Profile.
5. Enter a name and description for the new port profile, Select the Uplink Port Profile radio button. Leave the Load balancing algorithm to Host Default and set Teaming Mode to Switch Independent and click Next.



6. Select the network sites (which are part of your logical networks) that can be connected to via this uplink port profile. Click Next.



7. Click Finish to complete the creation of the uplink port profile.

Create Logical Switch using SET

A logical switch brings virtual switch extensions, port profiles, and port classifications together so that you can configure each network adapter with the settings you need, and have consistent settings on network adapters across multiple hosts.

This section covers the steps to create logical switch using embedded team as the uplink mode. Windows Server 2016 introduces Switch Embedded Teaming (SET) which, as the name suggests, teams multiple adapters directly in the VM Switch instead of creating a separate NIC team by using the Load Balancing and Failover (LBFO) functionality. SET has the benefit of enabling mixed use of adapters with the VM Switch and utilizing RDMA.

The logical switch will bring all of the components together. Follow these steps to create the Logical Switch:

1. Open Virtual Machine Manager
2. Click Fabric tab > Networking > Logical Switches > Create Logical Switch.
3. In Create Logical Switch Wizard > Getting Started, review the information, Click Next.
4. Enter a name, description for the new logical switch and select Uplink Mode as Embedded Team to deploy the switch with SET-based teaming and click Next

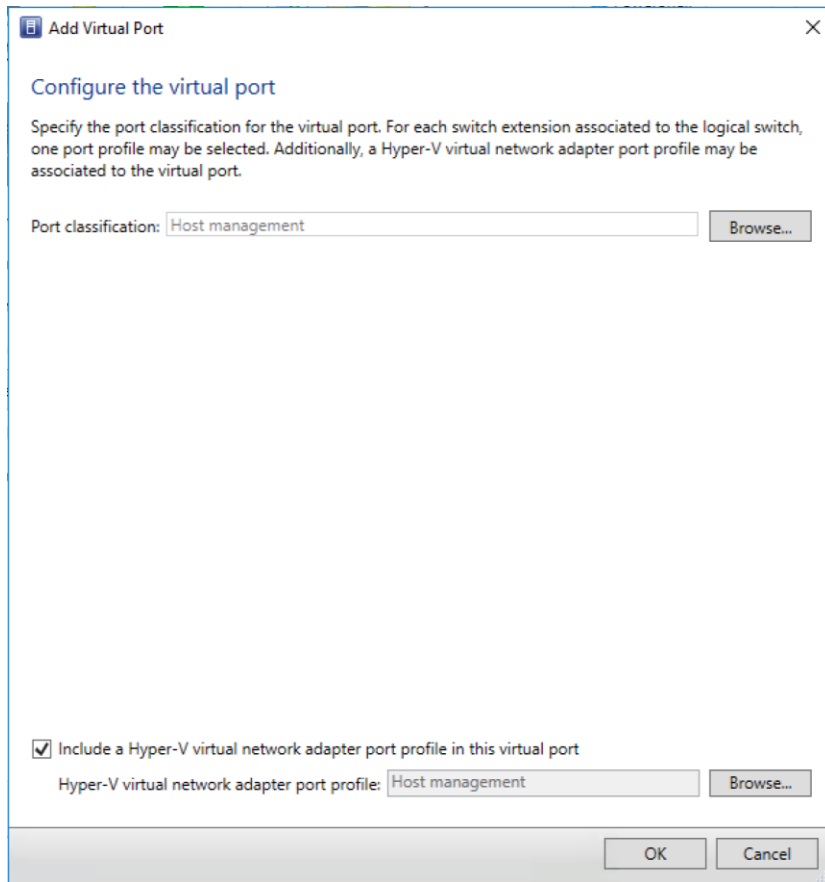
The screenshot shows the 'Create Logical Switch Wizard' dialog box with the 'General' tab selected. The dialog has a sidebar on the left with the following options: Getting Started, General (selected), Settings, Extensions, Virtual Port, Uplinks, and Summary. The main area is titled 'Enter name and description for the logical switch' and contains the following text: 'You can use a logical switch to apply settings to virtual switches across multiple hosts. A logical switch contains port profiles from the native Hyper-V switch and port profiles for any extensions that you use.' Below this text are three input fields: 'Name:' with the value 'Cluster_Logical_Switch', 'Description:' (empty), and 'Uplink mode:' with a dropdown menu showing 'Embedded Team'. At the bottom of the dialog are three buttons: 'Previous', 'Next', and 'Cancel'.

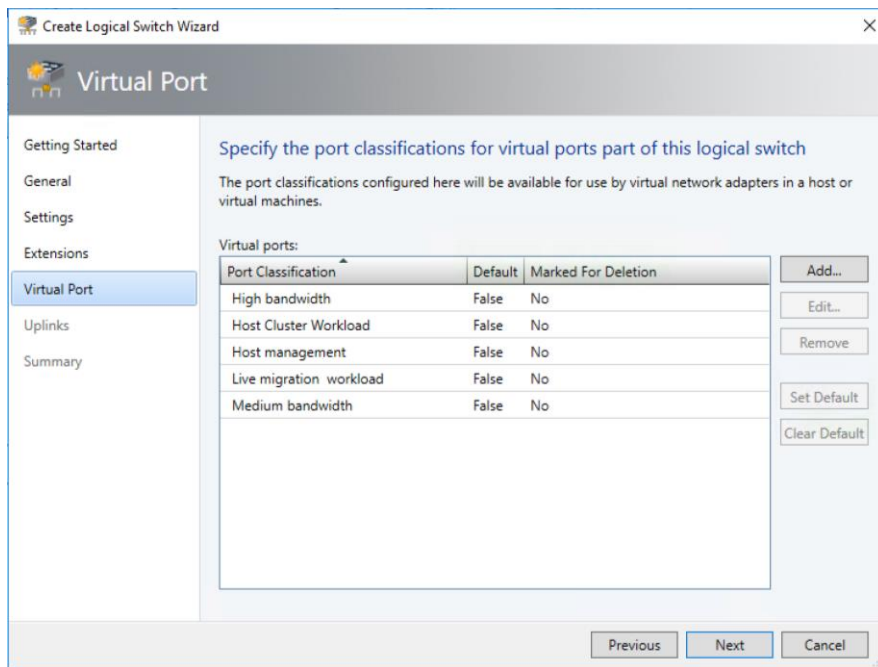
5. Select the minimum bandwidth mode as **Weight**, which quantifies minimum bandwidth for workloads, click Next
6. In Extensions selection window, leave the default and click Next.



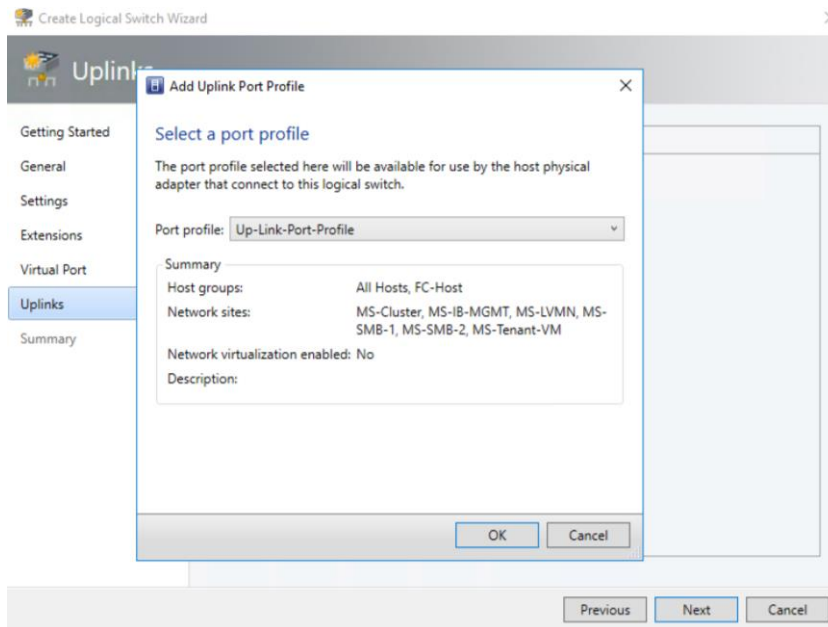
The list of installed virtual switch extensions is displayed and can be selected for deployment as part of the logical switch usage. This can be changed in the future if required.

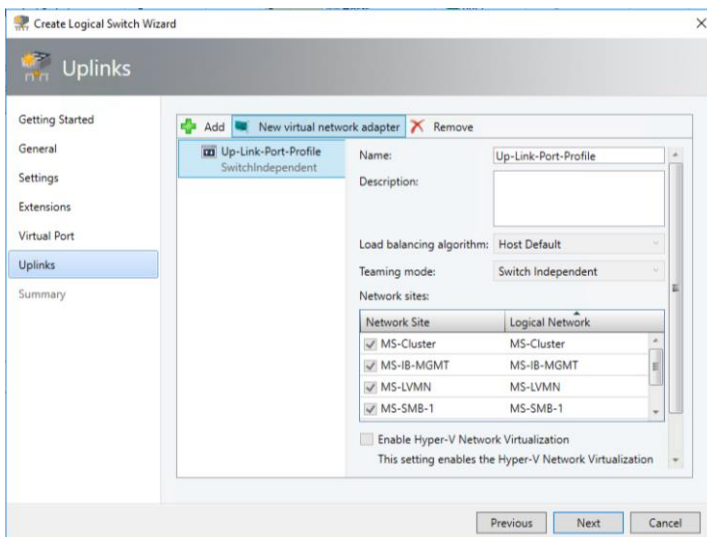
7. In Virtual Port window, Click the Add button, and in the dialog box that appears, click the Browse button to select the port classification. Then select the "Include a virtual network adapter port profile in this virtual port" check box and select the virtual port profile that corresponds. For example, if you select the high-bandwidth port classification, then most likely you would select the High Bandwidth Adapter virtual port profile object. Click OK. Repeat to add classifications. Select the classification that you would like to be the default, and click the Set Default button. Click Next.





8. In the Uplinks window, Click the Add button and then select Existing Uplink Port Profile - Up-Link-Port-Profile.

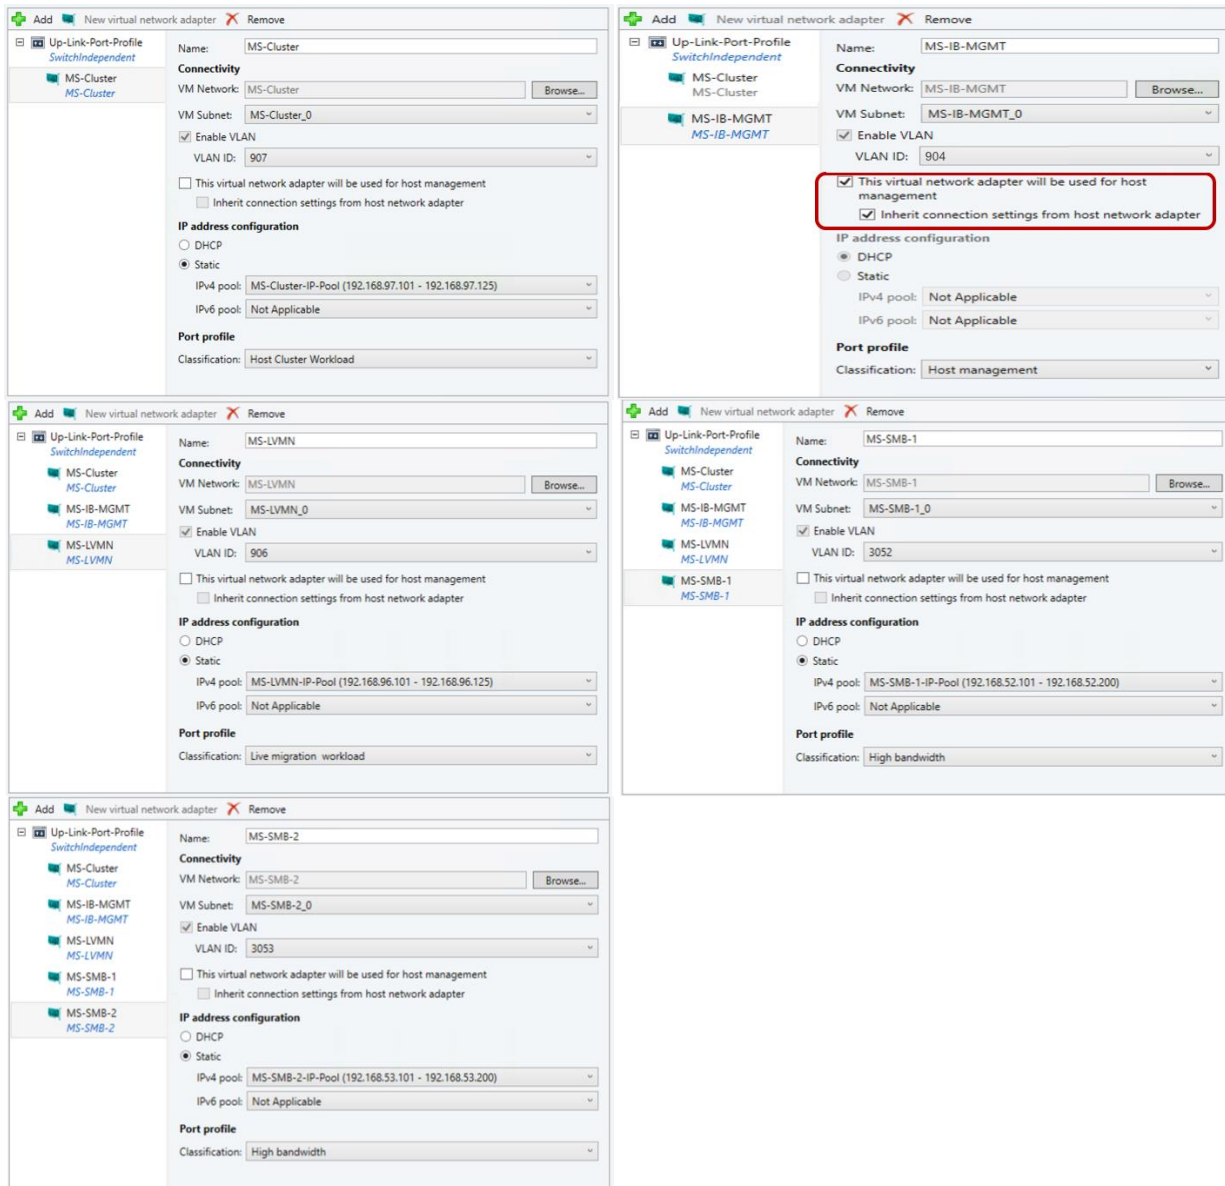




- Highlight Up link port profile, and click new virtual network adapter to add a virtual network adapter, click Browse to add the VM Networks and enter the name to match the VM Network under Connectivity. Under IP address configuration, select Static, select the IP Pool for the VM Network and Port Profile for the virtual adapter. Add all the virtual network adapters needed for your infrastructure and Click Next.



Only the MS-IB-MGMT virtual network adapter will have check box enabled for "This virtual network adapter will be used for host management" and "Inherit connection settings from host network adapter". This ensures continued connectivity for the host.



10. Click Finish to create the logical switch.

Fabric - Storage

NetApp SMI-S Provider Configuration



The NetApp SMI-S Provider can be downloaded from <http://mysupport.netapp.com>

Install the NetApp SMI-S Provider

1. Log into the Windows VM as Administrator.
2. Navigate to the directory that contains the NetApp SMI-S Provider software package. Double-click the package name.

3. Complete the steps in the setup wizard to complete the install.

Create the Local Administrator

1. Enter Win-R to open the Run application.
2. Open the Local Users and Groups window by entering `lusrmgr.msc` and pressing Enter.
3. Add a user named SMIS-User as a local Administrator

Configure the NetApp SMI-S Provider

1. In the Programs menu, navigate to NetApp SMI-S Provider.
2. Right click and select Run as Administrator. A command line prompt should open.
3. Run the command `smis cimserver status` to ensure the NetApp SMI-S Provider is running

```
C:\Program Files (x86)\NetApp\smis\pegasus\bin>smis cimserver status  
NetApp SMI-S Provider is running.
```

4. Add a user to the CIM server by running the following command:



The added user should be a valid domain administrator on your domain.

```
C:\Program Files (x86)\NetApp\smis\pegasus\bin>cimuser -a -u flexpod\SMIS  
Please enter your password: *****  
Please re-enter your password: *****  
User added successfully.
```

5. Add the SVM to the SMI-S Provider using the following command:

```
C:\Program Files (x86)\NetApp\smis\pegasus\bin>smis addsecure 192.168.94.80 vsadmin  
Enter password: *****  
Returned Path ONTAP_FilerData.hostName="192.168.94.80",port=443  
Successfully added 192.168.94.80
```

NetApp SMI-S Integration with VMM

To add a remote storage device in Virtual Machine Manager (VMM), you can add and discover external storage arrays that are managed by Storage Management Initiative – Specification (SMI-S) or Store Management Provider (SMP) providers.

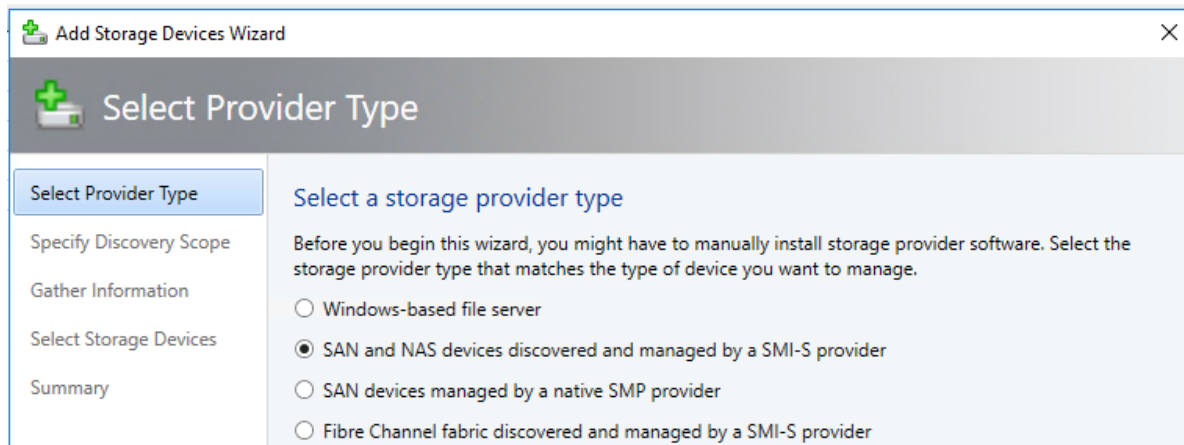
To add an SMI-S storage device, ensure that you have installed the SMI-S provider for the array on a server that the VMM management server can access over the network by IP address or by fully qualified domain name (FQDN).



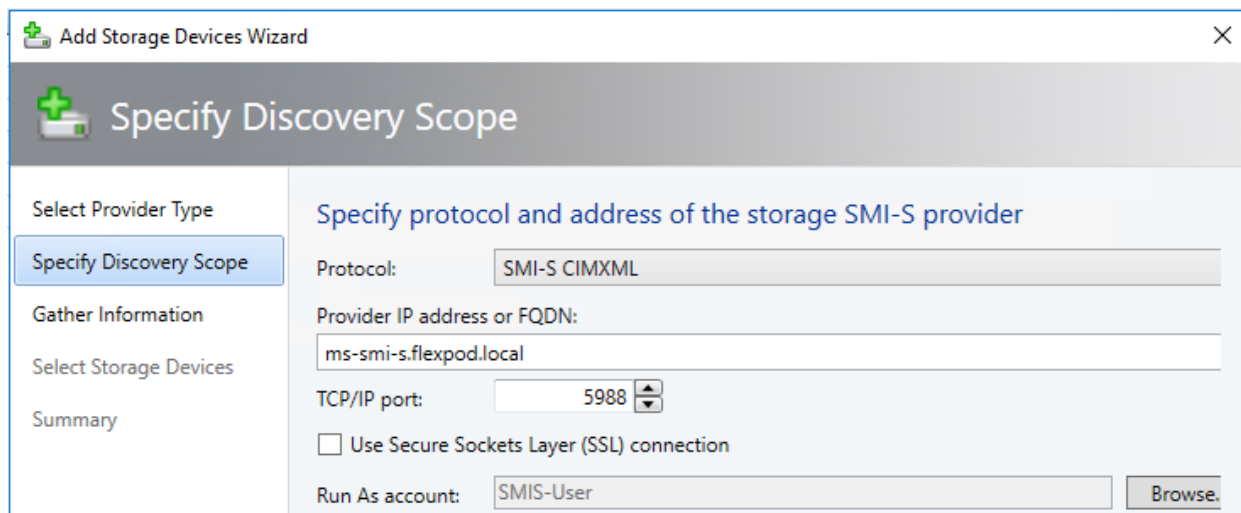
Do not install the SMI-S provider on the VMM management server. This configuration is not supported.

Add a storage device

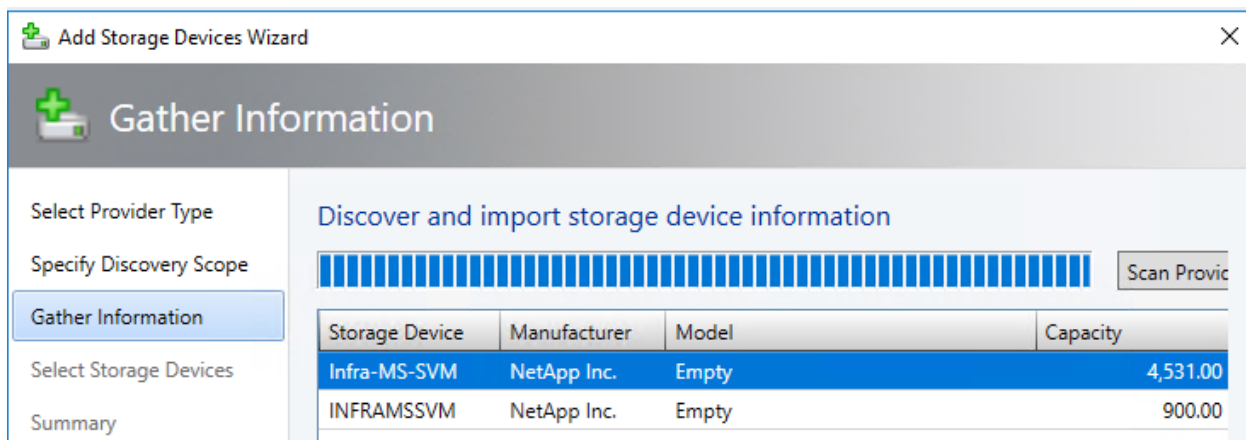
1. Click Fabric > Storage > Add Resources > Storage Devices.
2. In Add Storage Devices Wizard > Select Provider Type, select to add a storage device with SMI-S.



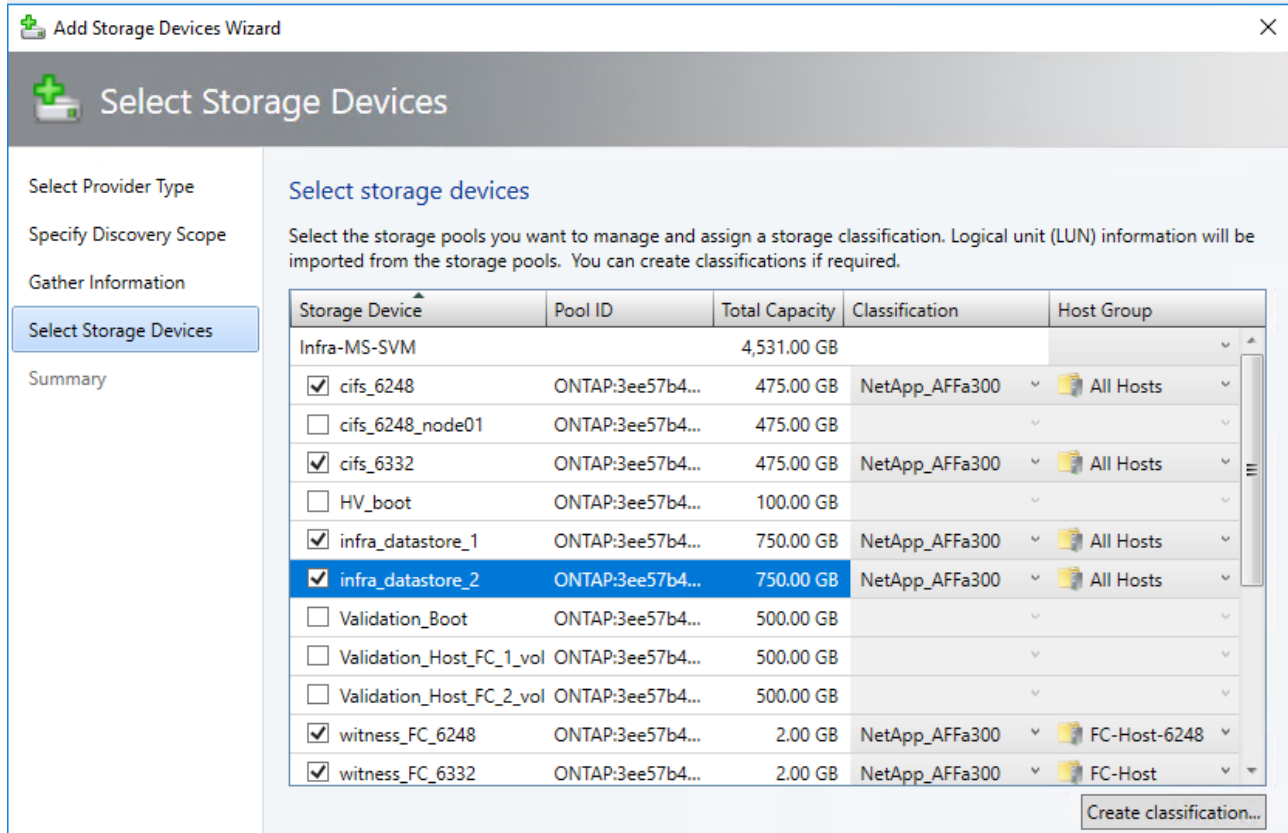
3. In Specify Discovery Scope, select Protocol - SMI-S CIMXML, add the IP address/FQDN, and add the port used to connect to the provider on the remote server. You can enable SSL if you're using CIMXML. Then specify an account for connecting to the provider.



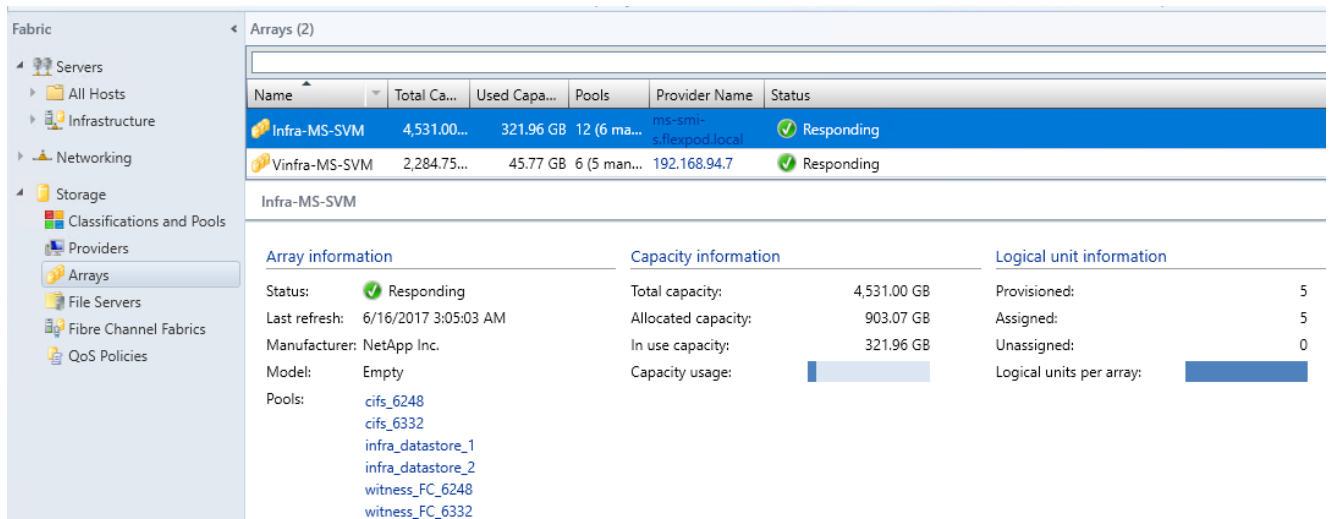
4. In Gather Information, VMM automatically tries to discover and import the storage device information.
5. If the discovery process succeeds, the discovered storage arrays, storage pools, manufacturer, model, and capacity are listed as shown in the below figure. When the process finishes, click Next.



- In Select Storage Devices, specify a classification and host group from the drop-down list for each storage pool. Create storage classifications if none exists to group storage pools with similar characteristics.



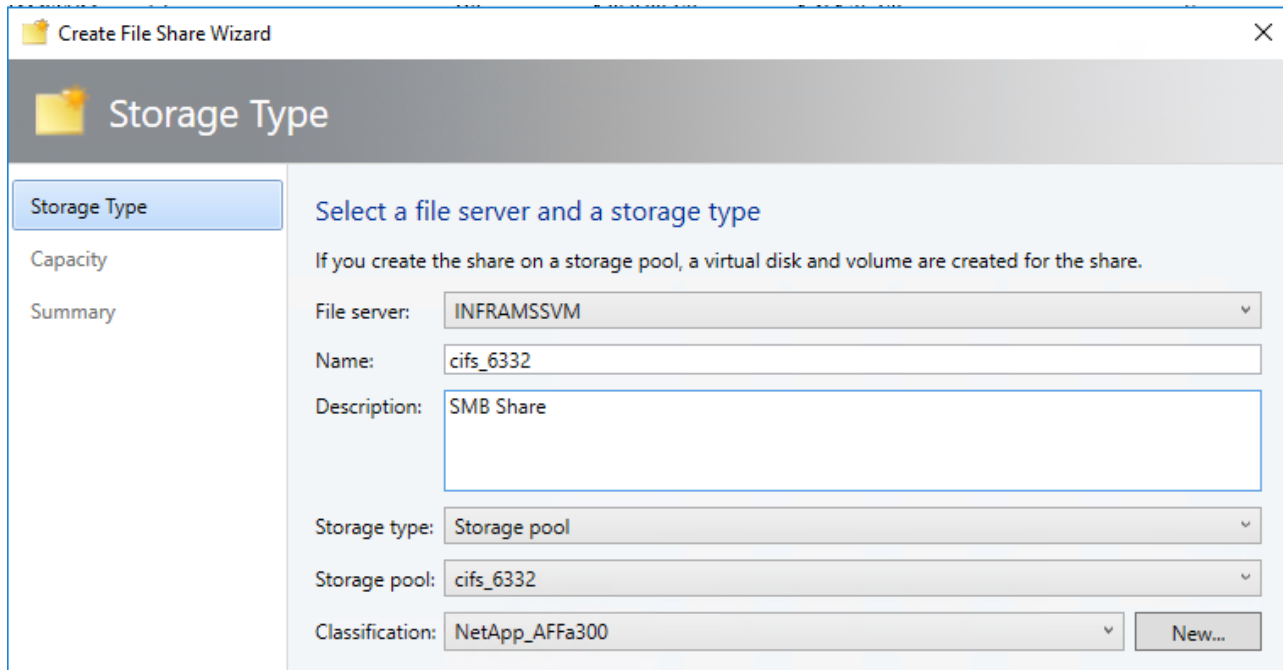
- On the Summary page, confirm the settings, and then click Finish. The Jobs dialog box appears. When status is Completed you can verify the storage in Fabric > Storage.



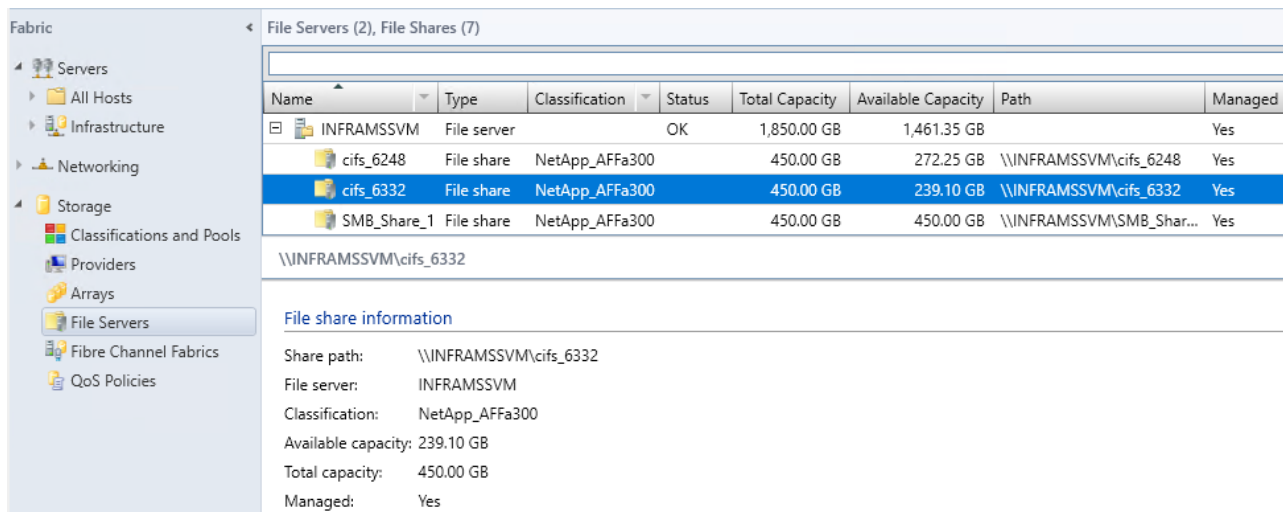
Create and Assign SMB 3.0 file shares to the Hyper-V host clusters

SMB file shares can be used by Hyper-V hosts as a shared storage to store virtual machine files. This section covers steps to create and assign SMB file shares to stand-alone Hyper-V servers and host cluster.

1. To add a storage device, refer to the steps covered in the above section.
2. To create a file share, open Fabric workspace, expand Storage and click on File Servers.
3. Select the File Server and click on Create File Share and in the Create File Share wizard, enter a name for the share and select Storage Type, Pool and Classification. Click Next.

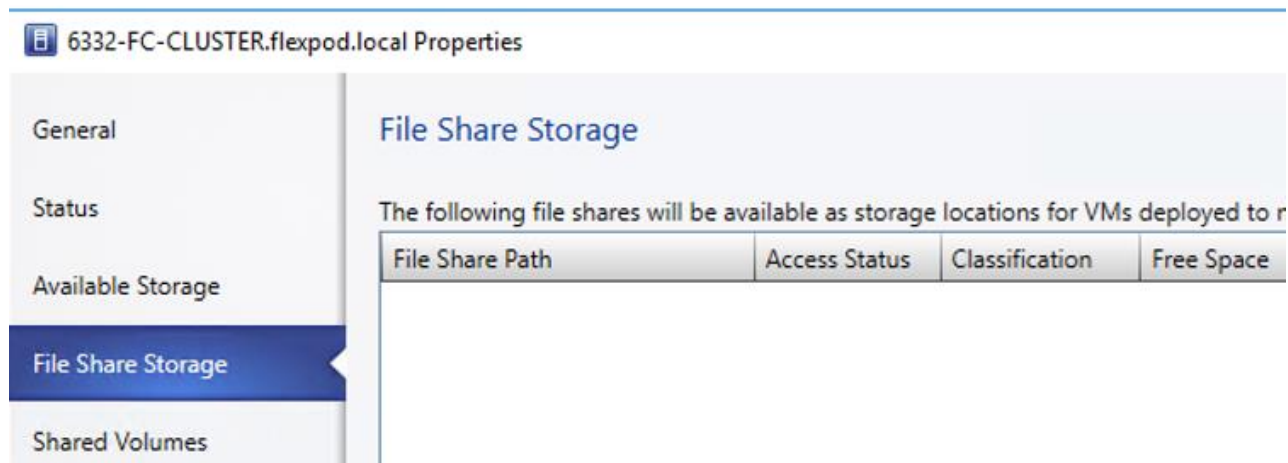


4. In the Capacity page, enter a size and click Next.
5. In the Summary page, confirm the setting and click Finish.
6. Verify the file share created in the above steps by navigating to Fabric > Storage and click on File Servers.

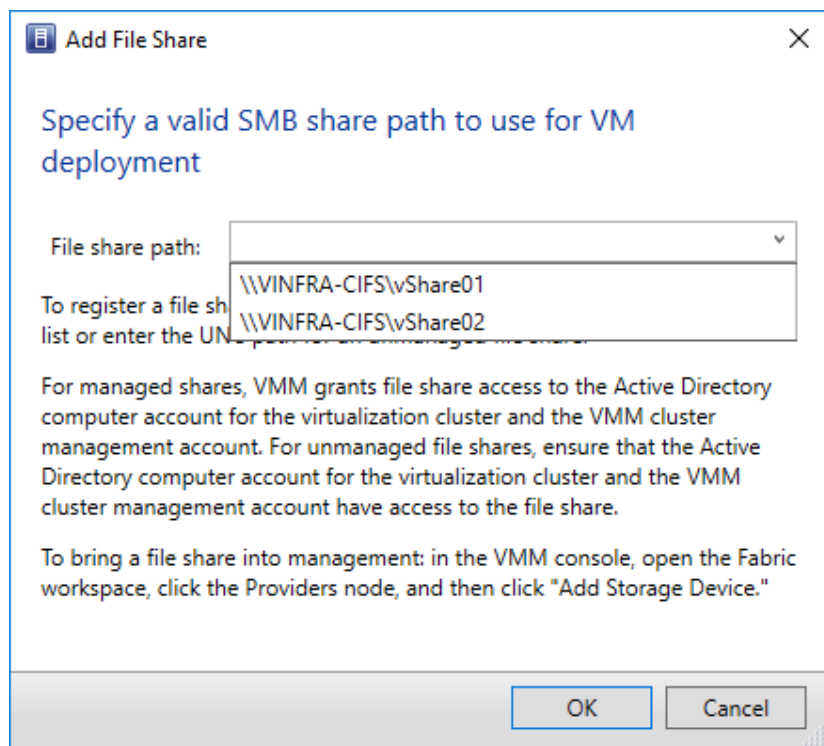


7. Assign the file share to the host cluster by navigating to Fabric > Servers > All Hosts.
8. Locate and right-click on the cluster icon and click on Properties.

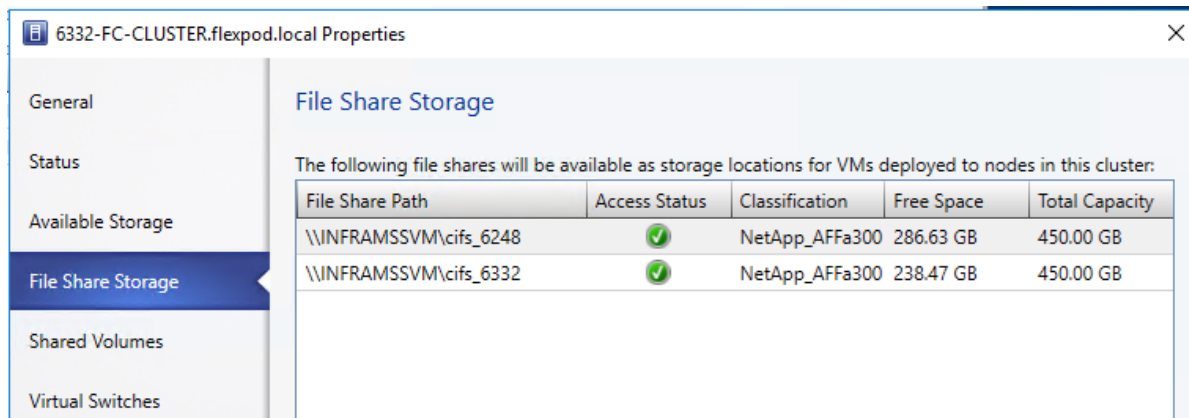
- Click on File Share Storage and click Add.



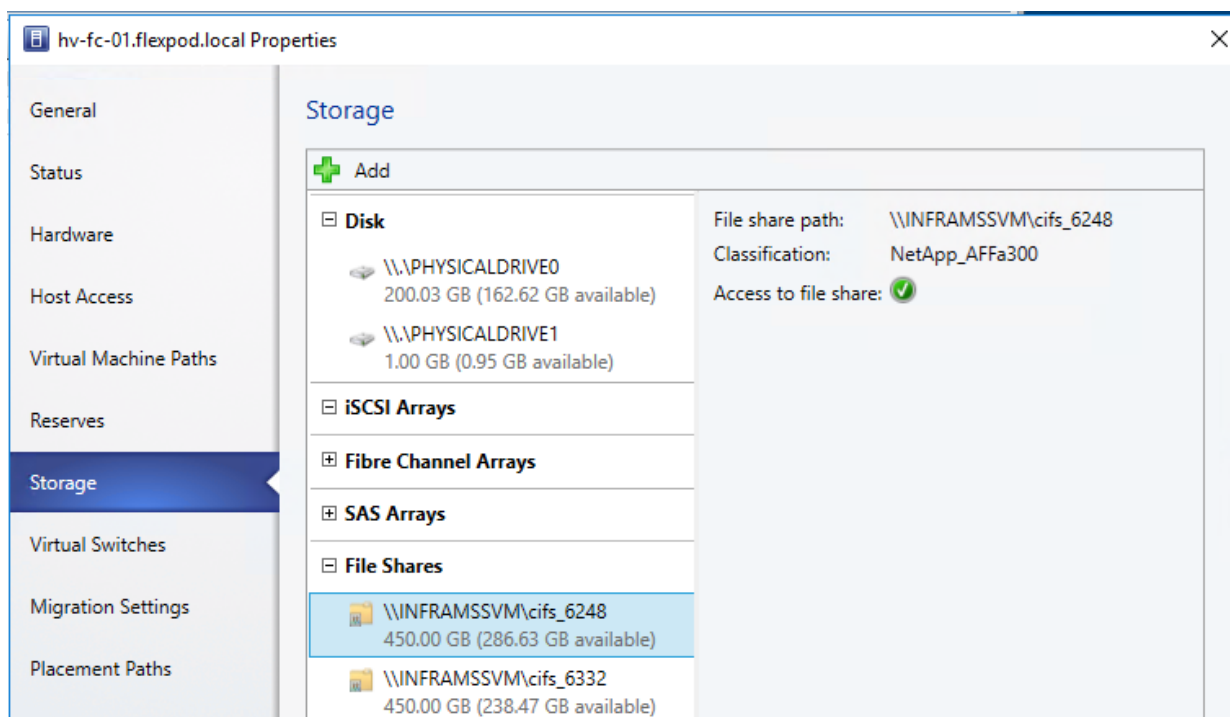
- From the drop-down menu next to the File Share Path, select a share.



- Repeat this step to select other shares.



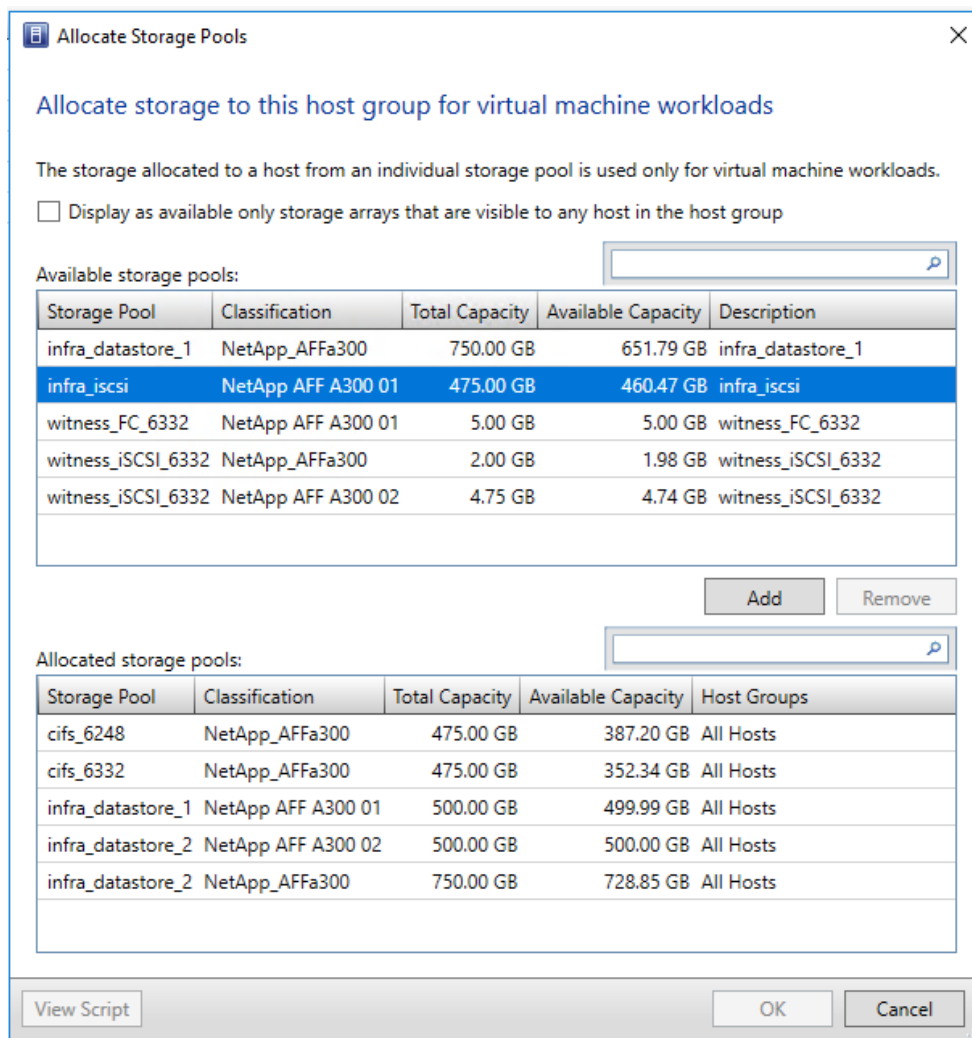
12. Verify the allocation by checking each cluster node properties for the file share allocation under the storage as shown in the figure below.



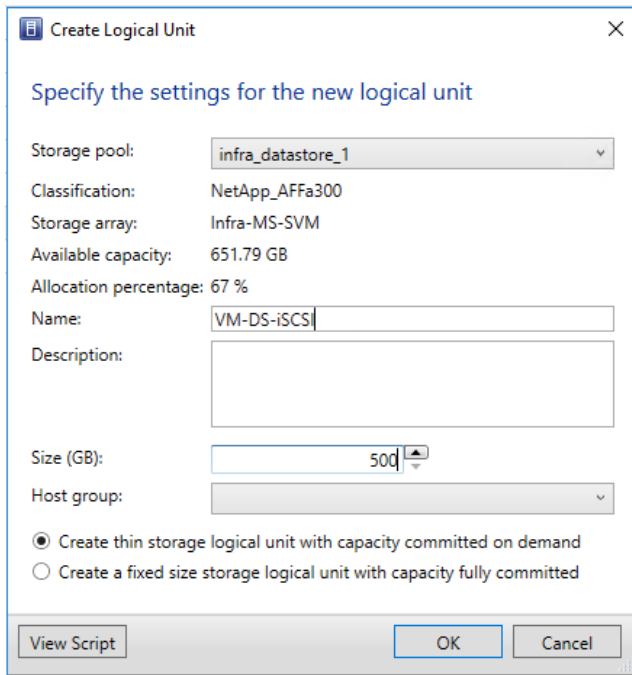
Allocate a storage pool to a host group

For this document purpose, the necessary SAN LUNs required for deploying Windows Hyper-V clusters were already created on the array before the SMI-S integration with VMM section. During the integration process of SMI-S with VMM, these storage pools were classified and associated with the appropriate host groups. The steps in the sections below show how to create storage pools and LUNs from the VMM console after the SMI-S integration as an example.

1. Click Fabric > Storage > Allocate Capacity, and click the host group.
2. The total and available storage capacity information is displayed for the host group. The storage capacity information includes the total and available capacity for both local and remote storage, and total and available allocated storage. Click Allocate Storage Pools.
3. Click a storage pool > Add.



4. Create a LUN in VMM In the SCVMM console, Click Fabric > Storage > Create Logical Unit.
5. Specify the storage pool, a name and description for the LUN, and the size. Click OK to create the LUN.

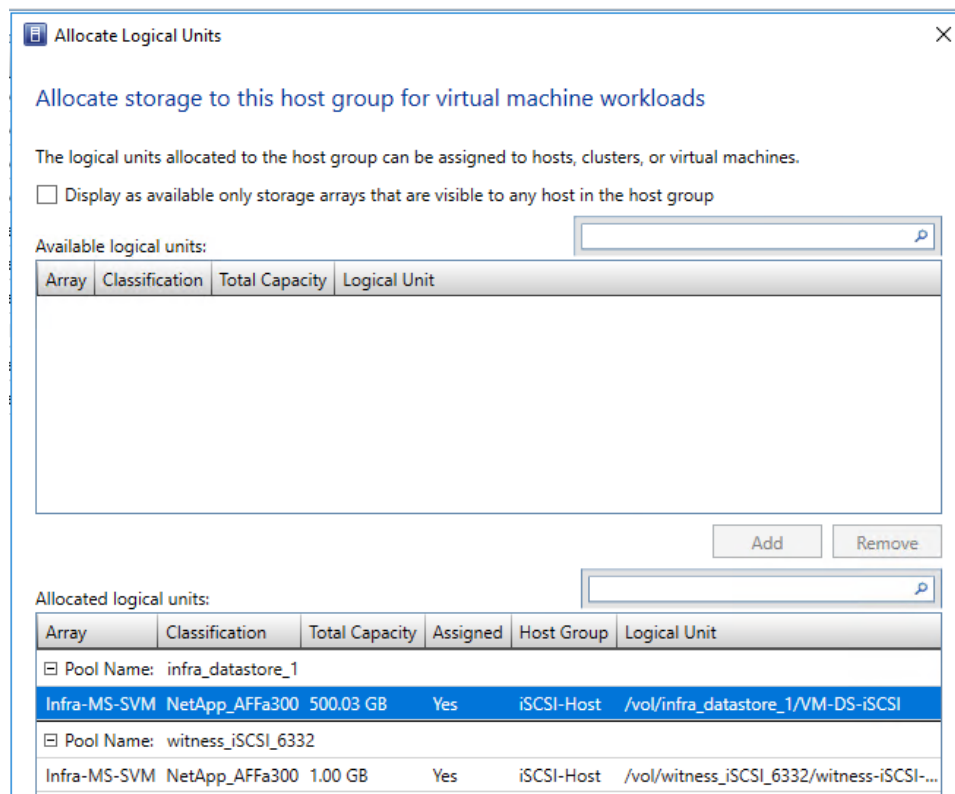


6. Verify that the LUN was created in **Fabric Resources > Classifications, Storage Pools, and Logical Units**.

Name	Type	Size	Available...	A	Description	Provisioning Type
\\192.168.52.61\smb_ds_2	Classification	0 GB	0 GB			
NetApp AFF A300 01	Classification	980.00 GB	979.99 GB			
NetApp AFF A300 02	Classification	504.75 GB	504.75 GB			
NetApp_AFFa300	Classification	2,452.00 GB	2,122.16 GB		SMB Share-1	
cifs_6248	Storage pool	475.00 GB	387.20 GB		cifs_6248	
cifs_6332	Storage pool	475.00 GB	352.34 GB		cifs_6332	
infra_datastore_1	Storage pool	750.00 GB	651.79 GB		infra_datastore_1	
/vol/infra_datastore_1/VM-DS-iSCSI	Logical unit	500.03 GB			... /vol/infra_datastore_1/VM-DS-iSCSI	Thin

Allocate a LUN to a host group

1. Click Fabric > Storage > Allocate Capacity > Allocate Storage Capacity and click the host group.
2. Click Allocate Logical Units, select a unit > Add.



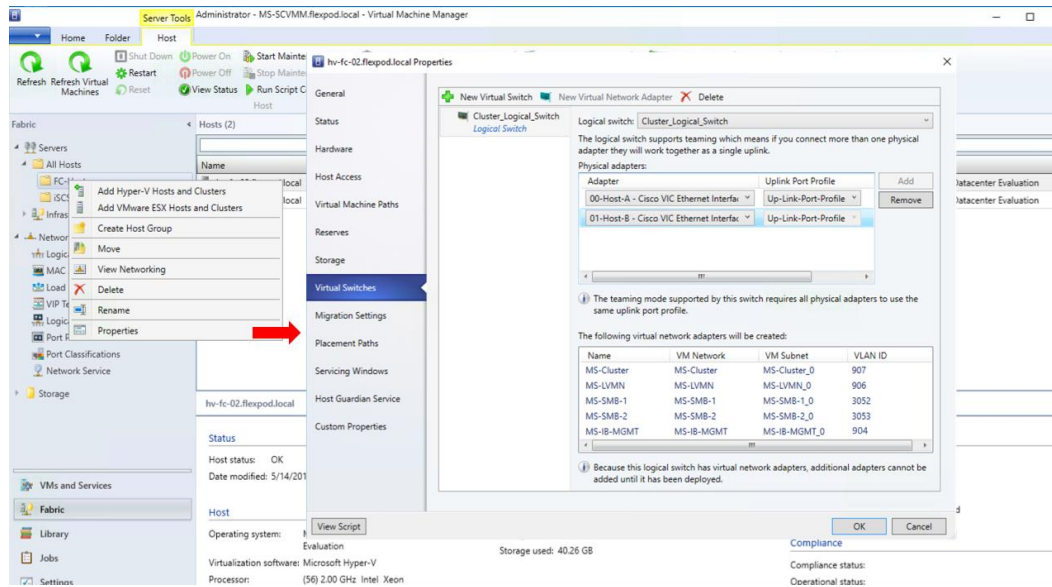
Fabric – Servers - II

Once the networking and storage configuration is complete and associated to the host groups, the next step is to deploy Hyper-V failover cluster.


- Configure Network on Hosts by Applying Logical Switch
- Deploy Hyper-V Failover Cluster

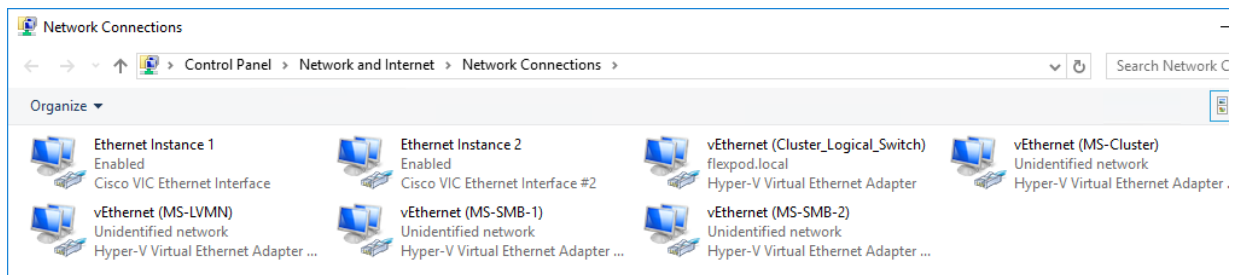
Configure Network on Host – Applying Logical Switch

1. Click Fabric > Storage > All Hosts > host group > Hosts > Host > Properties > Virtual Switches.
2. Open Fabric > Servers > click New Virtual Switch.
3. Select the logical switch you created. Under Adapter, select both the physical adapter to apply to the logical switch.



4. In the Uplink Port Profile list, select the uplink port profile Up-Link-Port-Profile and click Ok.
5. Repeat the above steps to configure the Logical Switch on all the hosts in the Host Group.
6. After applying the logical switch, you can check that the network adapter settings and verify whether they're in compliance with the switch:
 - a. Click Fabric > Networking > Logical Switches > Home > Show > Hosts.
 - b. Login to the host and verify the Network Adapters under Network Connections.

 In Logical Switch Information for Hosts verify the settings. Fully compliant indicates that the host settings are compliant with the logical switch. Partially compliant indicates some issues. Check the reasons in Compliance errors. Non-compliant indicates that none of the IP subnets and VLANs defined for the logical network are assigned to the physical adapter. Click the switch > Remediate to fix this.

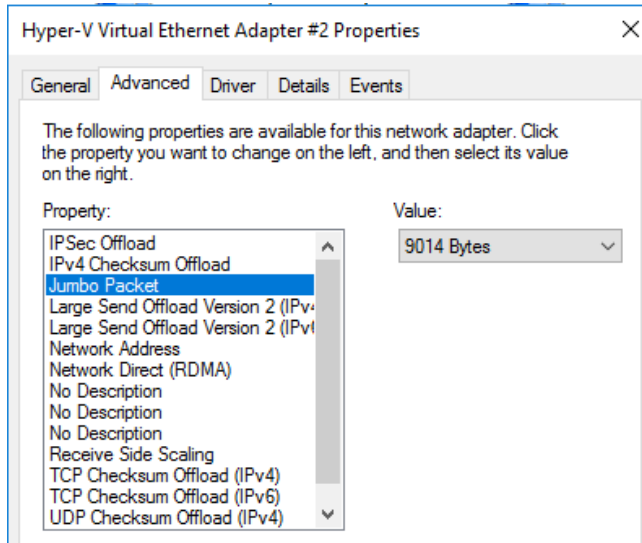


Enable Jumbo Frames

SET Team virtual switch does not require jumbo frame settings, however, the jumbo frames need to be enabled on the virtual network adapters of the Windows OS. Set the Jumbo Frames for the following vEthernet/virtual adapters.

- MS-LVMN
- MS-SMB-1
- MS-SMB-2

- MS-Cluster
1. Login to the Windows Operating System, under Network Connections, right click on the virtual adapters, select Properties.
 2. Check the Advanced Properties on the NIC in windows and set the Jumbo Packet value to 9014 Bytes



3. Verify and validate the jumbo packet settings as shown in the below commands

```
PS C:\Users\administrator.FLEXP0D> Get-NetAdapterAdvancedProperty -DisplayName "Jumbo Packet" | ft -AutoSize
```

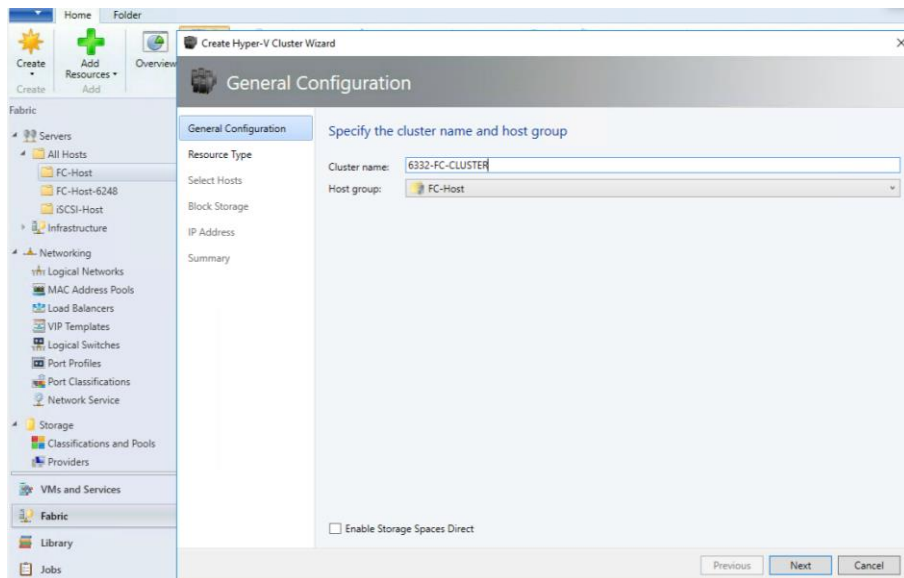
Name	DisplayName	DisplayValue	RegistryKeyword	RegistryValue
vEthernet (MS-SMB-1)	Jumbo Packet	9014 Bytes	*JumboPacket	{9014}
vEthernet (MS-Cluster)	Jumbo Packet	9014 Bytes	*JumboPacket	{9014}
vEthernet (MS-SMB-2)	Jumbo Packet	9014 Bytes	*JumboPacket	{9014}
vEthernet (MS-LVMN)	Jumbo Packet	9014 Bytes	*JumboPacket	{9014}
vEthernet (Cluster_Logical_Switch)	Jumbo Packet	Disabled	*JumboPacket	{1514}

```
PS C:\Users\administrator.FLEXP0D> ping 192.168.96.103 -f -l 8972
```

```
Pinging 192.168.96.103 with 8972 bytes of data:
Reply from 192.168.96.103: bytes=8972 time<1ms TTL=128
Reply from 192.168.96.103: bytes=8972 time<1ms TTL=128
Reply from 192.168.96.103: bytes=8972 time<1ms TTL=128
Reply from 192.168.96.103: bytes=8972 time<1ms TTL=128
```

Deploy Hyper-V Cluster

1. In the VMM console, click Fabric > Create > Hyper-V Cluster to open the Create Hyper-V Cluster wizard.
2. In General, specify a cluster name and choose the host group in which the existing Hyper-V hosts are located.

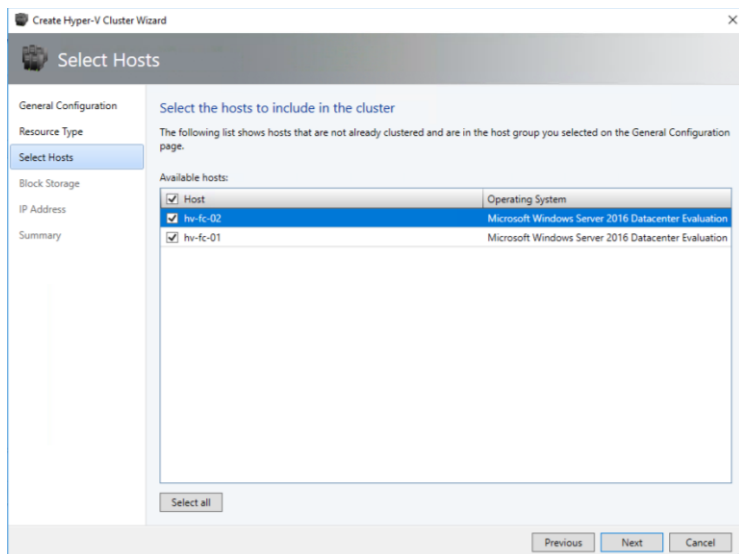


- In Resource Type, select the Run As account that you'll use to create the cluster.

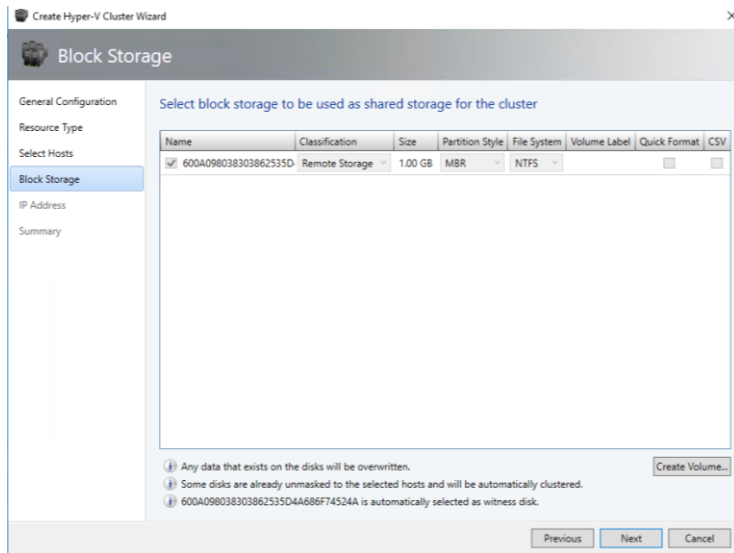


The accounts that you use must have administrative permissions on the servers that will become cluster nodes, and must belong to the same domain as the Hyper-V hosts that you want to cluster. Also, the account requires Create Computer objects permission in the container that is used for Computer accounts in the domain. Ensure that the option Existing Windows servers is selected.

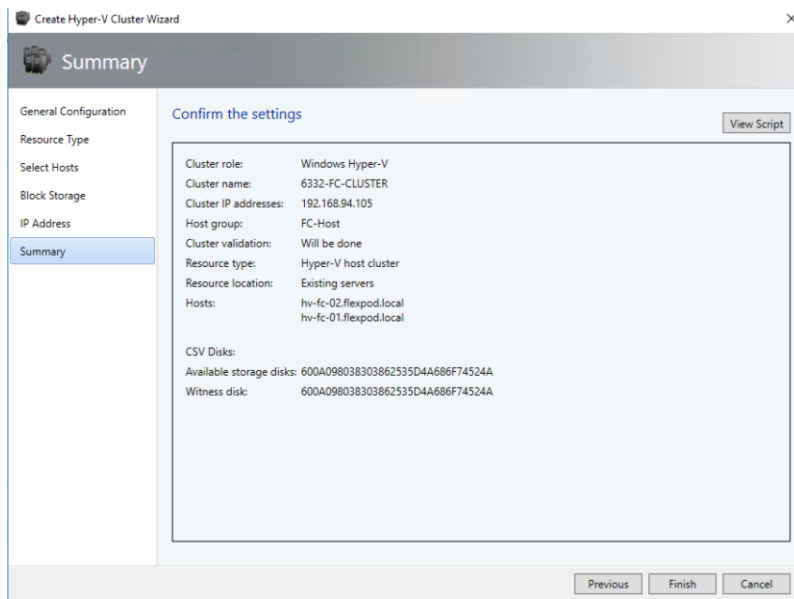
- In Nodes, select the Hyper-V host servers that you want to include in the cluster.



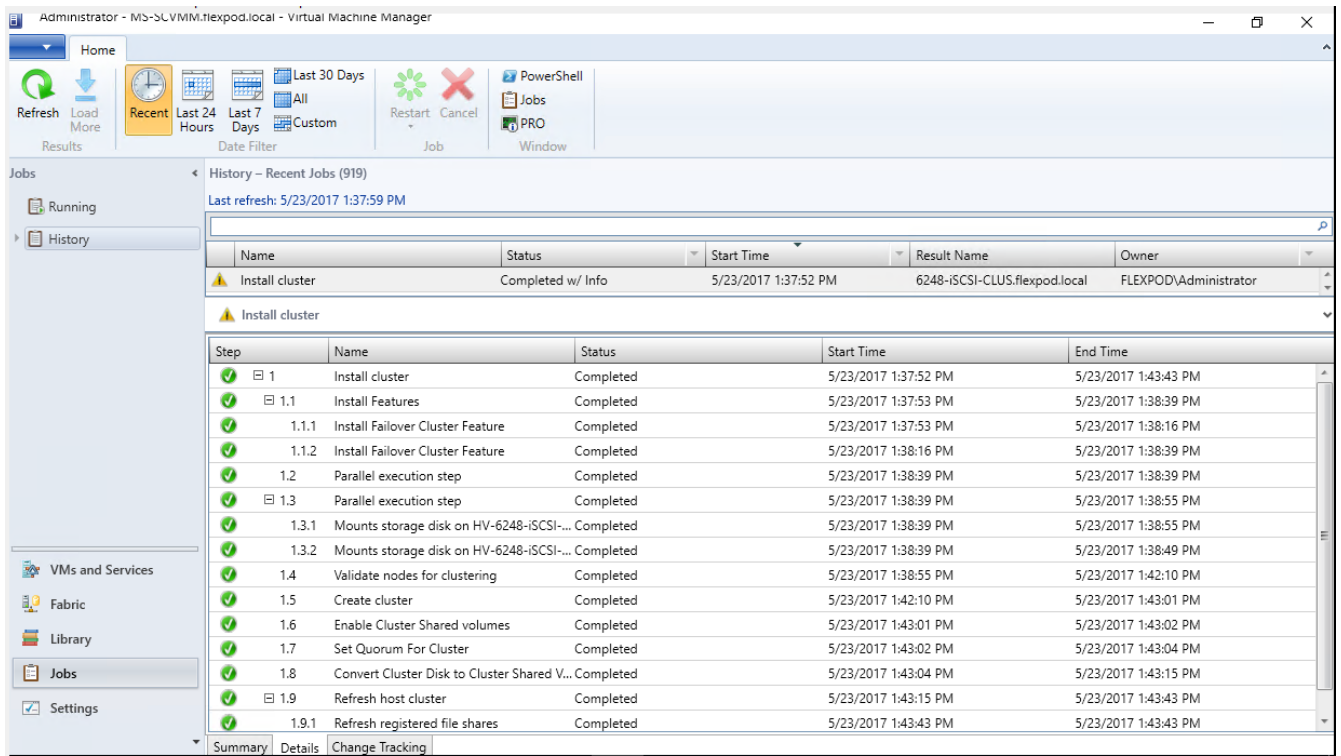
- In Block Storage, select the data disks you want the cluster to use as witness disk.




6. In IP address, type in the IP address you want to use for the cluster.
7. In Summary, confirm the settings and then click Finish.

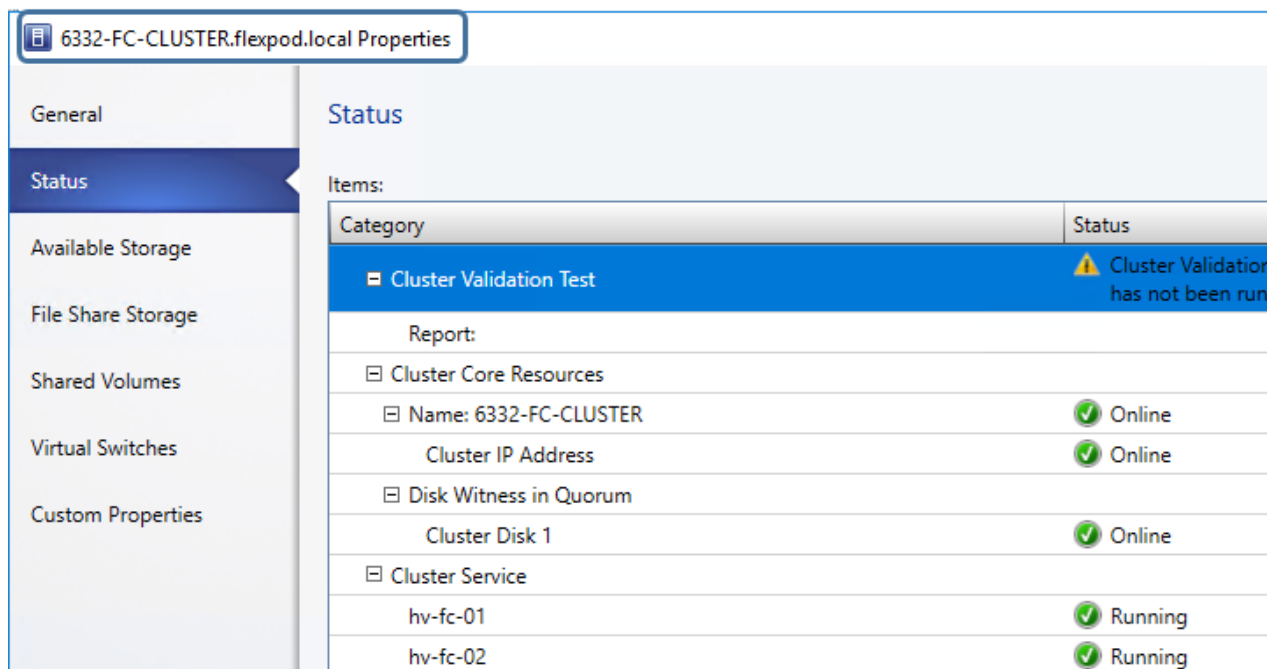


8. You can go to the jobs workspace and click on "Install Cluster" job to see the status of cluster installation. Fix and troubleshoot any errors or warnings and revalidate the cluster.



- After the cluster is installed, a new cluster icon is seen after expanding the Servers>All Hosts>FC-Host host group in the fabric workspace. Right-click on the cluster and click on properties to view the status and other information about the cluster.





Hyper-V Cluster Communication Network Configuration

A failover cluster can use any network that allows cluster network communication for cluster monitoring, state communication, and for CSV-related communication.

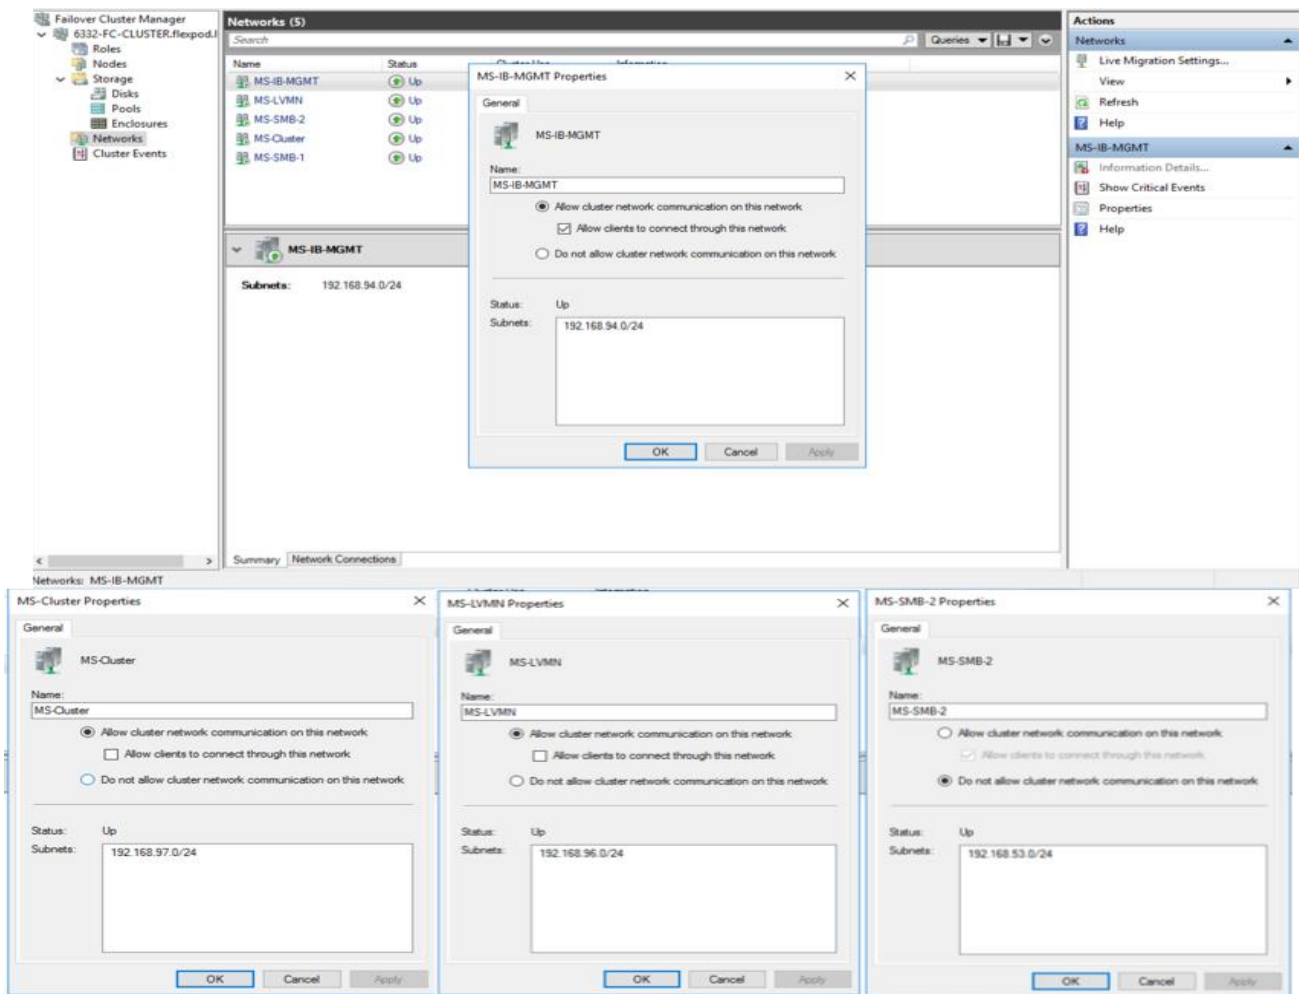
The following table shows the recommended settings for each type of network traffic.

To configure a network to allow or not to allow cluster network communication, you can use Failover Cluster Manager or Windows PowerShell.

Table 8 Recommended Settings for Network Traffic

Network Type	Recommended Setting
Management	Both of the following: - Allow cluster network communication on this network - Allow clients to connect through this network
Cluster	Allow cluster network communication on this network Note: Clear the Allow clients to connect through this network check box.
Live migration	Allow cluster network communication on this network Note: Clear the Allow clients to connect through this network check box.
Storage	Do not allow cluster network communication on this network

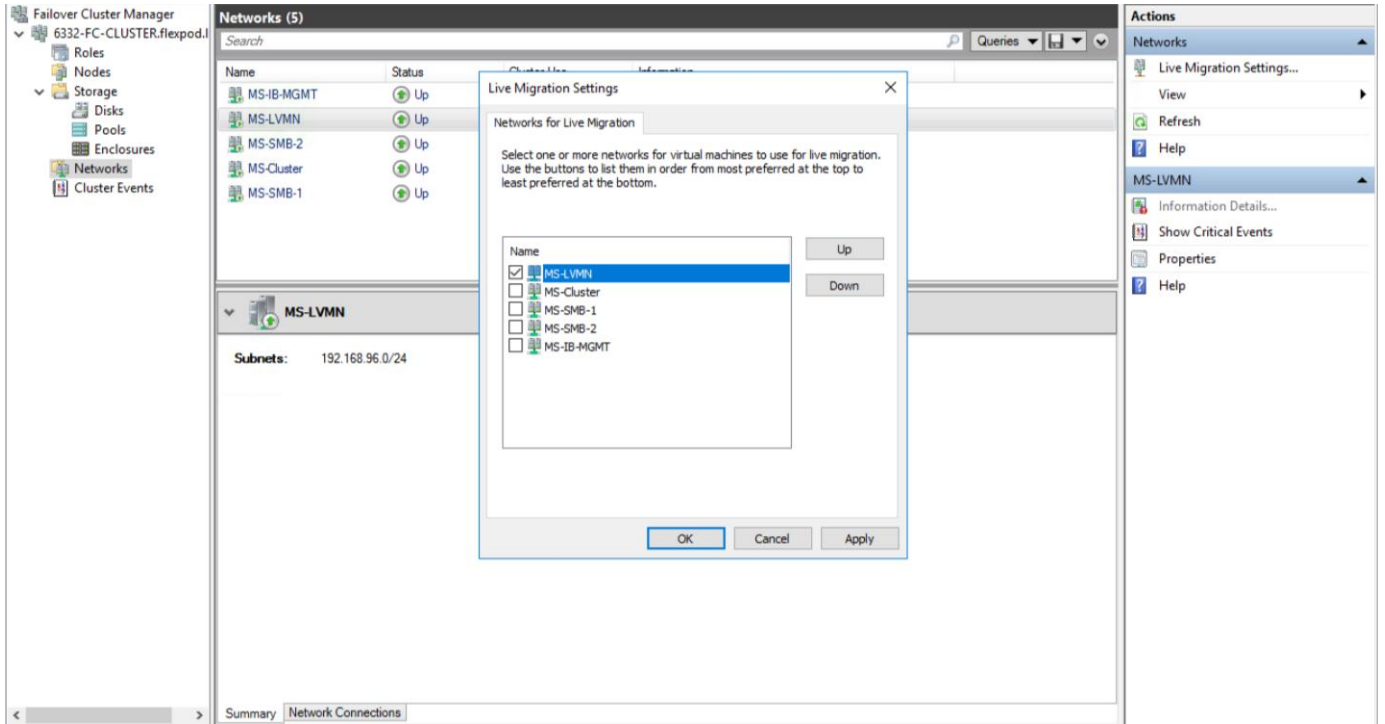
1. Open Failover Cluster Manager, click Networks in the navigation tree.
2. In the Networks pane, right-click a network, and then click Properties.



Live Migration Network Settings

By default, live migration traffic uses the cluster network topology to discover available networks and to establish priority. However, you can manually configure live migration preferences to isolate live migration traffic to only the networks that you define.

1. Open Failover Cluster Manager.
2. In the navigation tree, right-click **Networks**, and then click **Live Migration Settings**.
3. Select the Live Migration network.



Cisco UCS Management Pack Suite Installation and Configuration

Cisco UCS Manager Integration with SCOM

About Cisco UCS Management Pack Suite

Management Pack is a definition file with predefined monitoring settings. It enables you to monitor a specific service or application in Operations Manager. These predefined settings include discovery information which allows Operations Manager to automatically detect and start the monitoring services and applications. It also has a knowledge base which contains error details, troubleshooting information, alerts, and reports which helps to resolve the problems detected in the environment.

The Cisco UCS Manager Management Pack provides visibility to the health, performance, and availability of a Cisco UCS domain through a single, familiar, and easy-to-use interface. The management pack contains rules to monitor chassis, blade servers, rack servers, and service profiles across multiple Cisco UCS domains.

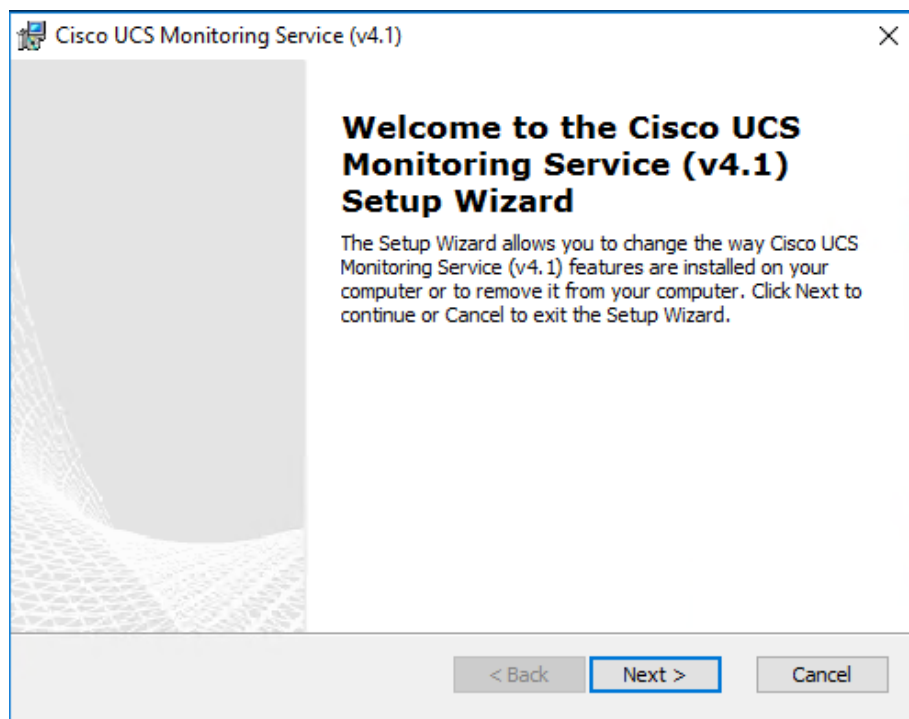
The Cisco UCS Central Management Pack has rules to monitor global service profiles and organizations across multiple Cisco UCS Central. It provides visibility of health and alerts through familiar and easy-to-use interface.

For more information, see:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/msft_tools/installation_guide/SCOM/b_Management_Pack_Installation_Guide.html

Installing Cisco UCS Monitoring Service

1. Navigate to the folder in which the unzipped Cisco UCS Management Pack Suite is stored.
2. Select the monitoring service installer .msi file, and launch the installer.
3. In the Setup wizard, click Next.



4. In the License Agreement page, do the following:
 - a. Review and accept the EULA.
 - b. Click Next.
5. In the Product Registration page, complete the following:
 - a. Enter a username.
 - b. **Optional:** Enter the name of your organization. The username is required, but the organization name is optional.
 - c. Click Next.
6. In the Select Installation Folder page, accept the default installation folder or click Browse to navigate to a different folder, and then click Next.
7. On the Ready to Install page, click Install to start the installation.
8. Once the Cisco UCS monitoring service is successfully installed, the Installation Complete message appears.
9. Click Finish.



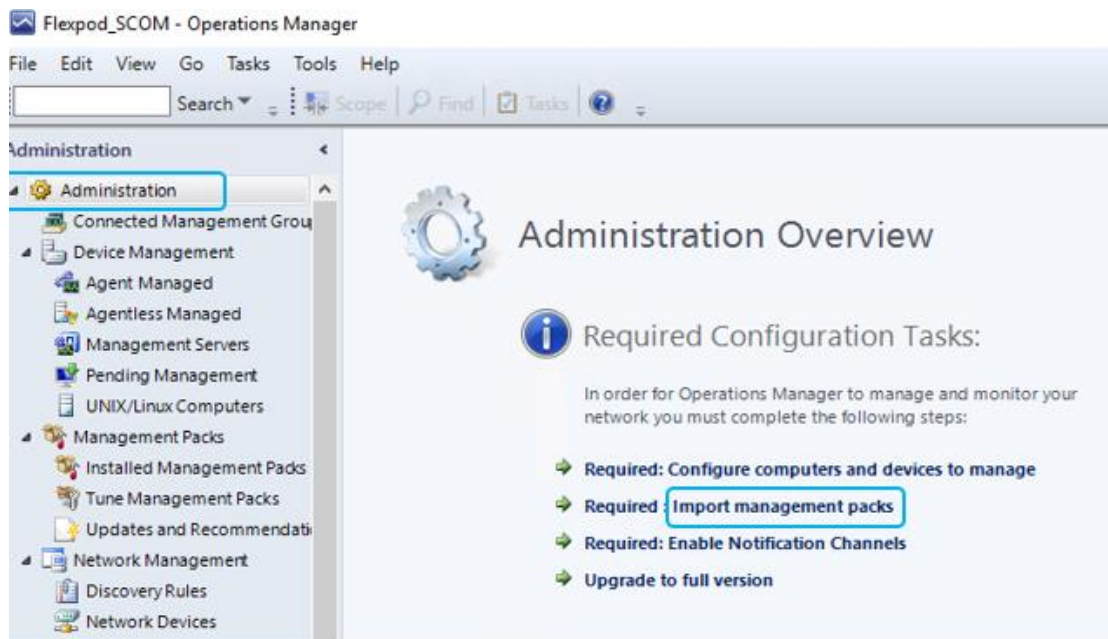
The same installation procedure is followed to install the monitoring service on agent managed computers and gateway servers.

Adding a Firewall Exception for the Cisco UCS Monitoring Service

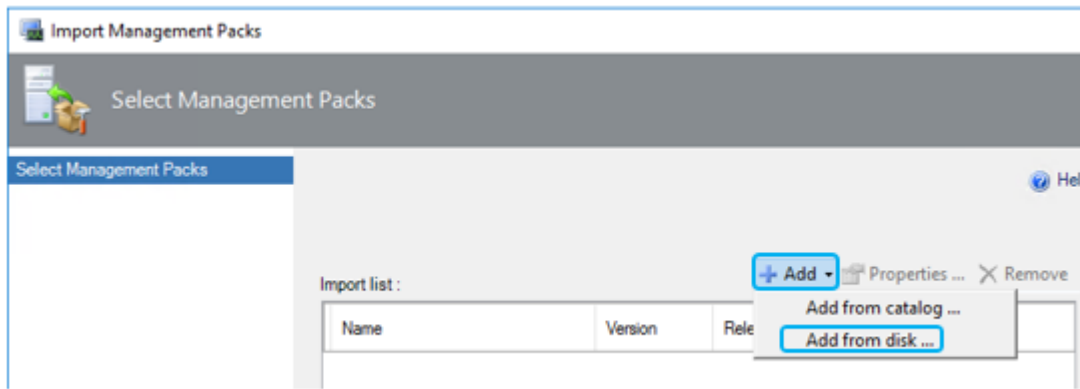
1. Before you monitor a Cisco UCS domain, enable the following inbound rules in the Windows Firewall with Advanced Security on the computer where you run the Cisco UCS Management Service.
2. File and Printer Sharing:
 - a. Echo-Request—ICMPv4-In
 - b. Echo-Request—ICMPv6-In
3. Remote Service Management (RPC)
4. Remote Service Management (RPC-EPMAP)

Installing the Cisco UCS Management Pack Suite

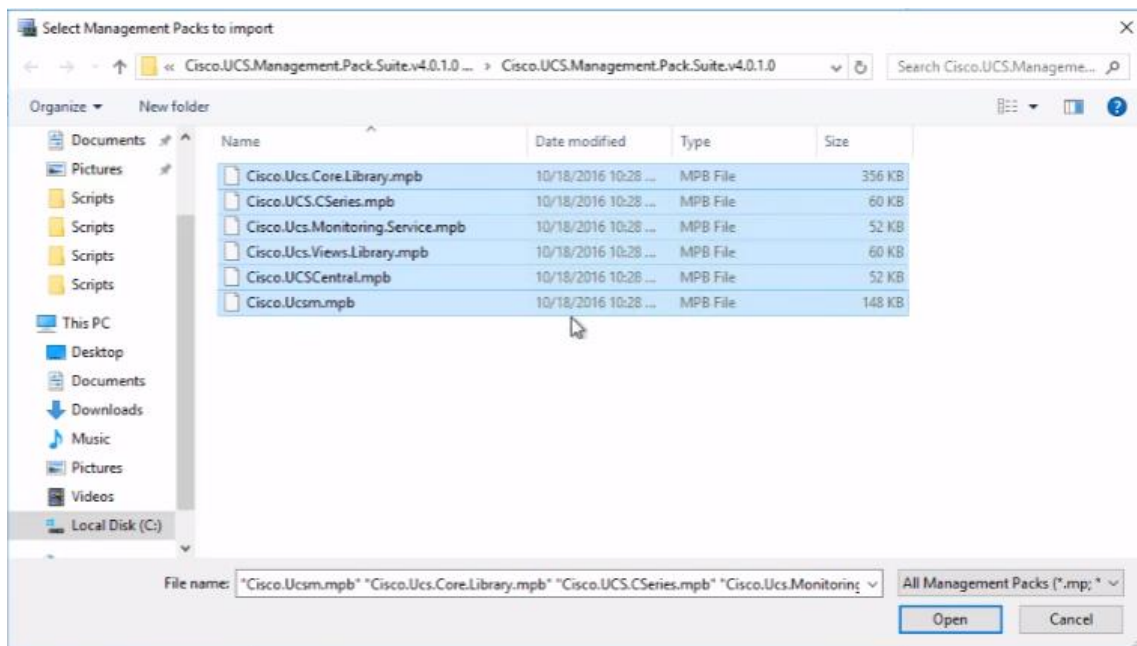
1. For importing Management Packs using Operations Manager console, you must have administrative privileges. For more information on the access privileges, see <https://technet.microsoft.com/en-in/library/hh212691.aspx>. On the Cisco.com download site for Cisco UCS Management Partner Ecosystem Software, download the Cisco UCS management pack suite file and unzip the file into a folder.
2. Launch Operations Manager console.
3. Navigate to the Administration > Management Packs > Import Management Packs tab.



4. On the Import Management Pack page, click Add and select Add from the disk. An Online Catalog Connection dialog box appears.

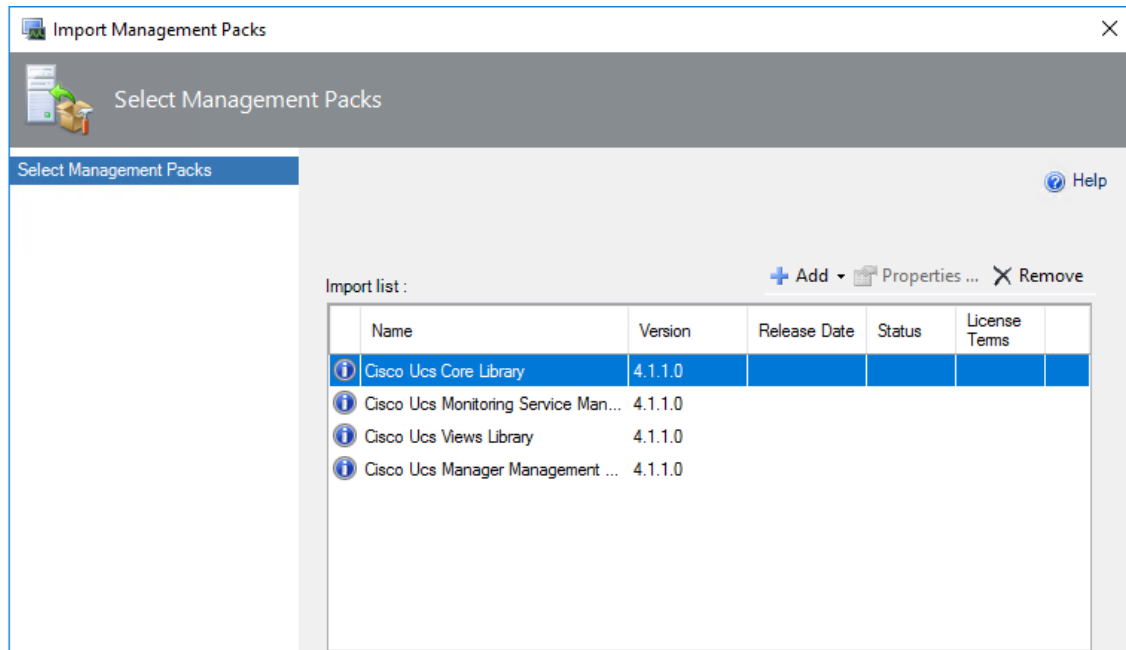


5. Click No, if you do not want to search the management pack dependencies online.
6. Navigate to the unzipped management pack suite files folder.
7. From the list of files, select the mandatory files:
 - Cisco.Ucs.Views.Library.mpb
 - Cisco.Ucs.Core.Library.mpb
 - Cisco.Ucs.Monitoring.Service.mpb
8. Other management pack files can be imported based on your machine requirements. For example, select *Cisco.Ucsm.mpb* for UCS Manager, *Cisco.UCS.CSeries.mpb* for Cisco IMC, and *Cisco.UCSCentral.mpb* for UCS Central.

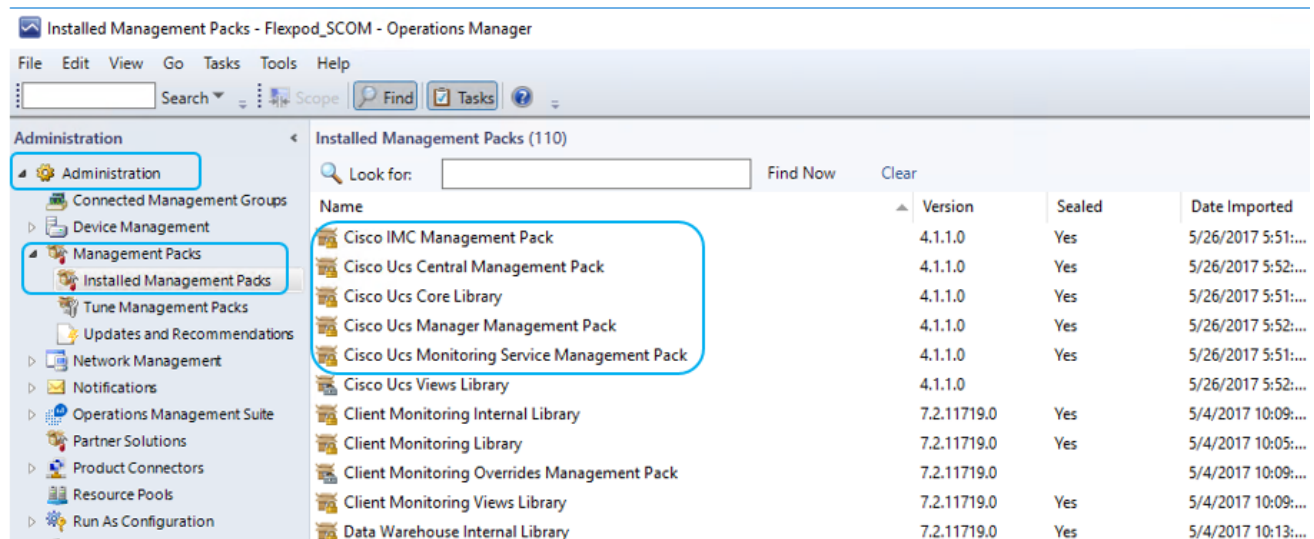


9. Click Open.
10. Click Install on the Import Management Packs page.

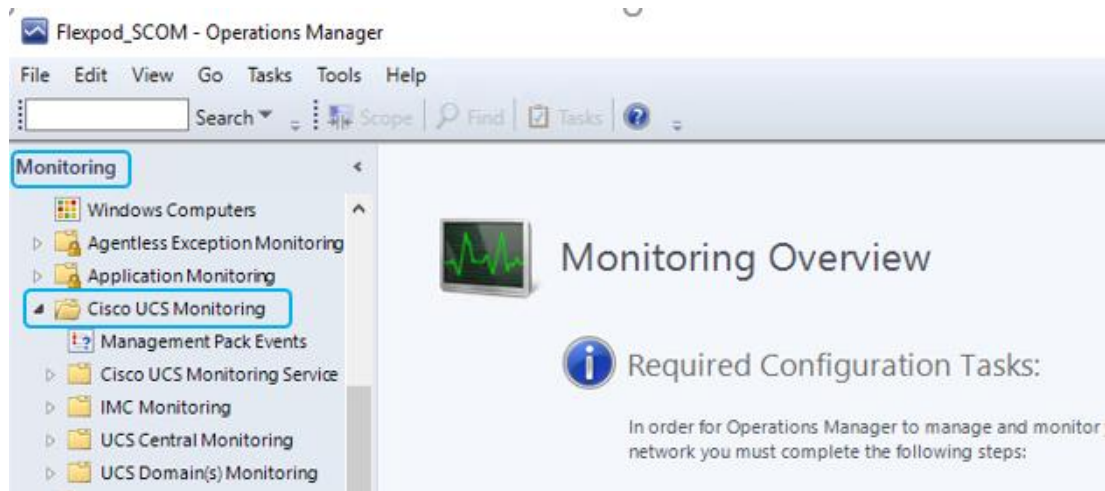
It may take few minutes to import the files.



- Verify the installation by navigating to the Administration > Management Packs and click on Installed Management Packs



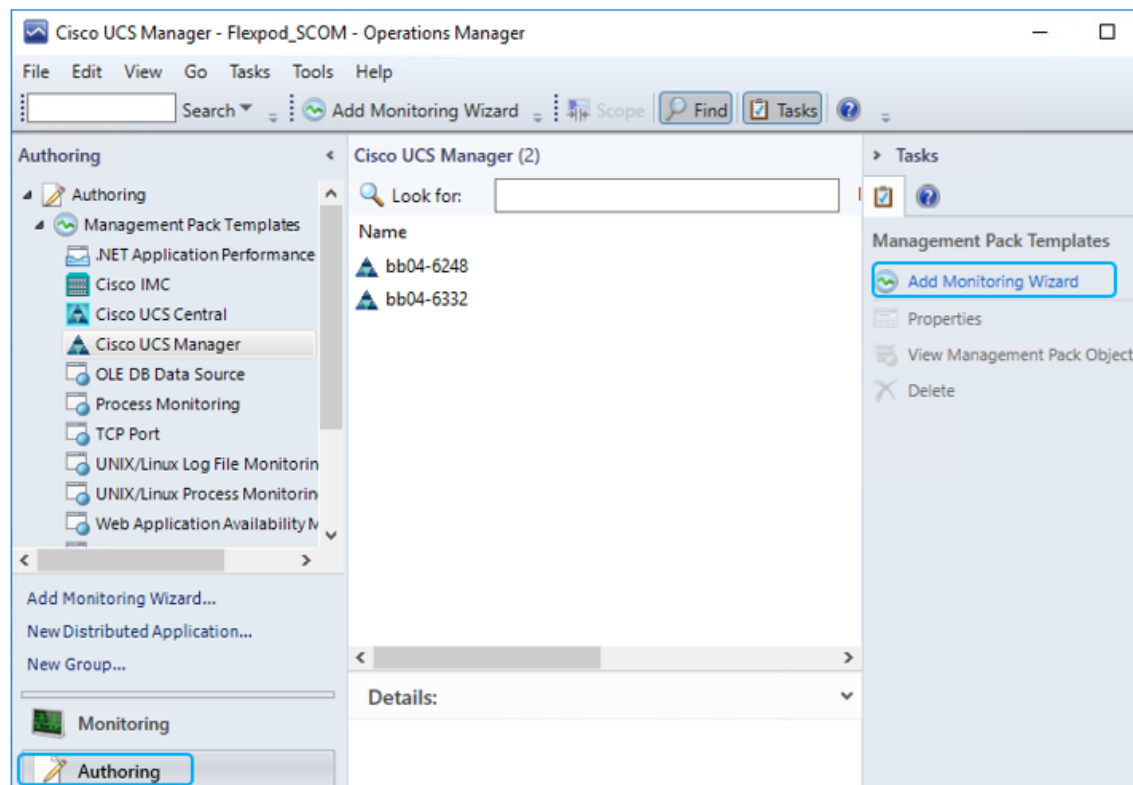
- In the Monitoring pane, a Cisco UCS folder is also created. When the folder is expanded, it lists the Cisco UCS Monitoring Service, IMC, UCS Central and UCS Domain monitoring folders.



Adding a Cisco UCS Domains to the Operations Manager

You can add Cisco UCS domains on the servers, where either management pack is imported or the Cisco UCS Management Service is installed.

1. Launch the Operations Manager console.
2. Navigate to Authoring > Cisco UCS Manager.
3. From the Tasks pane, click Add Monitoring Wizard.



4. On the Monitoring Type tab, click Cisco UCS Manager.

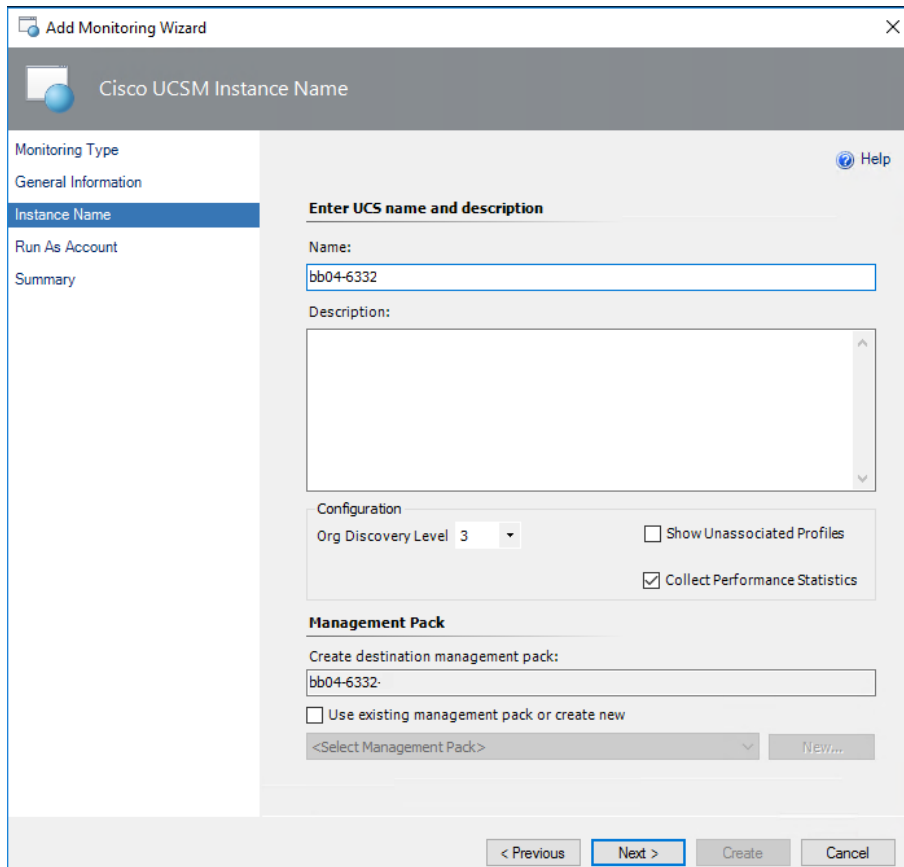
5. Click Next.
6. On the General Information tab, review and complete the following as shown in the below figure:

The screenshot shows the 'Add Monitoring Wizard' dialog box with the 'Specify IP Address, Port and Connection Mode' tab selected. The 'Monitoring Type' sidebar on the left has 'General Information' selected. The main content area is titled 'Cisco UCS Manager' and contains the following fields:

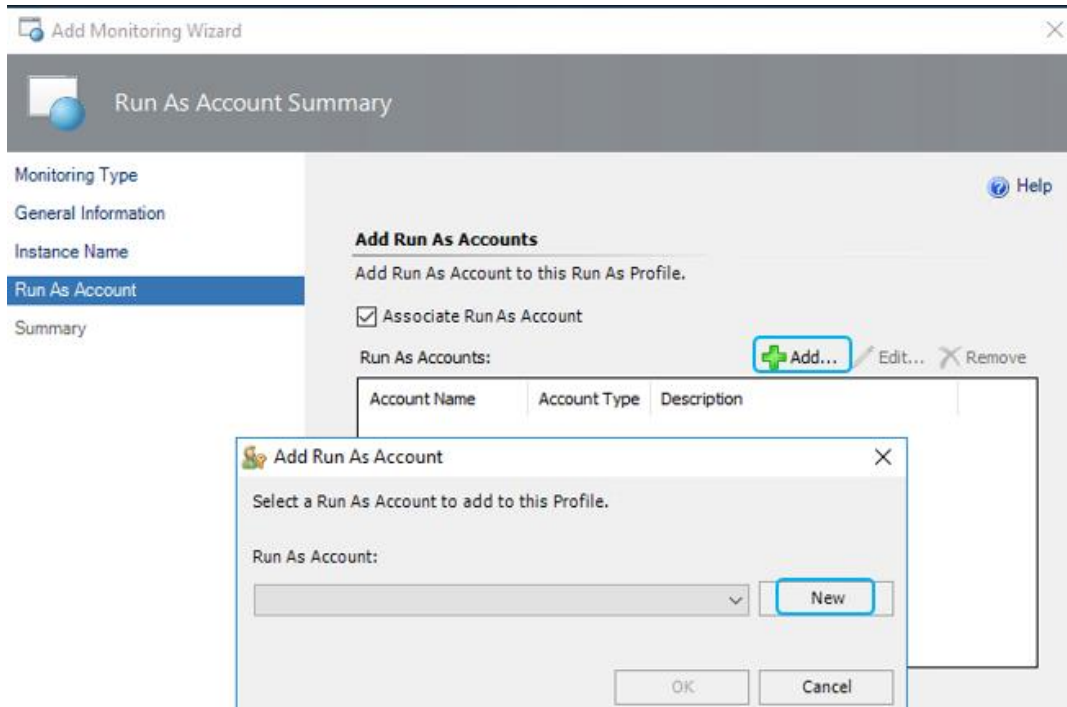
- Connection:**
 - IP Address* / Hostname: 192.168.156.12
 - Connection Mode: Secure
 - Port Number: 443
- Proxy Server:**
 - Enable Proxy Configuration
 - IP Address * / Hostname : []
 - Port: []
 - Enable Proxy Authentication
 - Username: [] Password: []
- Cisco UCS Monitoring Service:**
 - Machine Type: Management Server
 - Service Machine: MS-SCOM.flexpod.local

At the bottom right, there is a 'Test Connection' button.

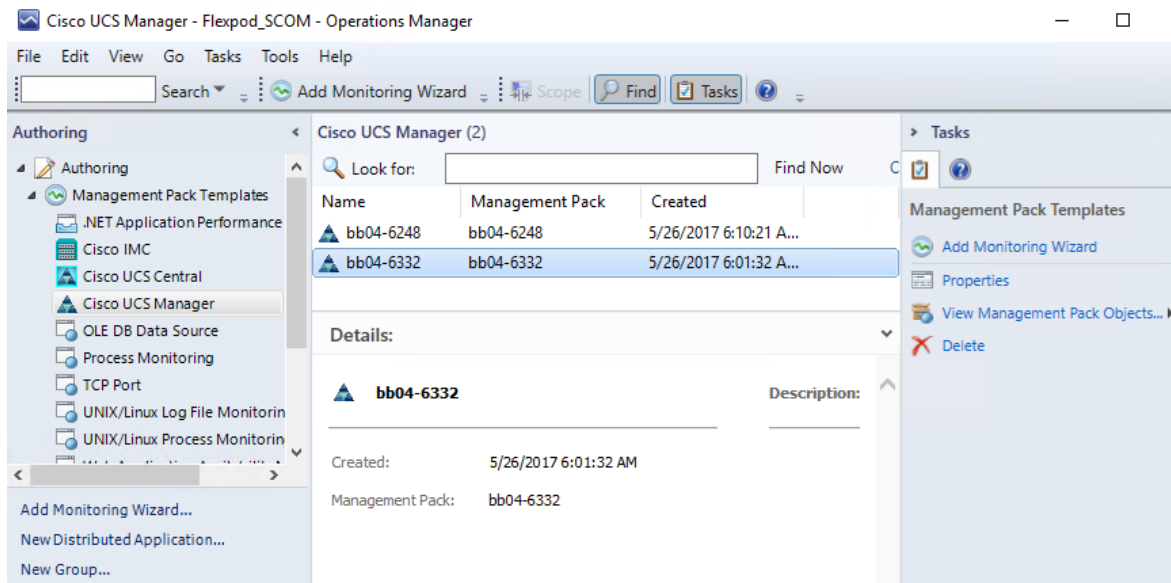
7. To check Operations Manager connectivity to UCS Manager, click Test Connection.
8. In the Authentication dialog box, enter the username and password, and click OK and Click Next.
9. On the Instance Name tab, complete the following as shown in the figure below and Click Next.



10. On the Run As Account tab, click Add
11. If you want to associate a new run-as account to the UCS domain instance, click New.



12. Click Next.
13. On the Summary tab, review the configuration summary, and click Create. The template for monitoring the UCS domain is created.



Cisco UCS Manager Monitoring Dashboards

The UCS Domain(s) Monitoring folder contains the following views:

- Ucs Domain Alert Dashboard—Displays all alerts generated in the UCS domain. The alerts are further categorized into the following views:
 - Active Alerts
 - Acknowledge Alerts
 - Cleared Alerts

Ucs Domain Alert Dashboard - Flexpod_SCOM - Operations Manager

File Edit View Go Tasks Tools Help

Search Overrides Scope Find Tasks

Monitoring

- Windows Computers
- Agentless Exception Monitoring
- Application Monitoring
- Cisco UCS Monitoring
 - Management Pack Events
 - Cisco UCS Monitoring Service
 - IMC Monitoring
 - UCS Central Monitoring
 - UCS Domain(s) Monitoring
 - Ucs Domain Alert Dashboard**
 - UCS Domain Diagram
 - UCS Domain State Dashboard
 - Chassis
 - Fabric Extender
 - Fabric Interconnect
 - Organization
 - Rack Unit
- Data Warehouse
- Microsoft Audit Collection Servi
- Microsoft System Center Virtual
- Microsoft Windows Client
- Microsoft Windows Server
- Network Monitoring
- Operations Management Suite
- Operations Manager
- Synthetic Transaction

Ucs Domain Alert Dashboard

Active Alerts (67)

Icon	Source	Name	Resolution State	Created	Age
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,
✖	Blade 4	Server.F0283: link-down	New	6/3/2017 7:47:11 PM	1 Day,

Acknowledged Alerts

Icon	Source	Name

Cleared Alerts (205)

Severity: Critical (183)

Icon	Source	Name
✖	IO Module 1	IOModule.F0481: equipment-p
✖	IO Module 2	IOModule.F0481: equipment-p
✖	FabricInterconn...	FabricInterconnect.F0276: link-
✖	FabricInterconn...	FabricInterconnect.F0276: link-
✖	FabricInterconn...	FabricInterconnect.F0276: link-
✖	FabricInterconn...	FabricInterconnect.F0276: link-
✖	FabricInterconn...	FabricInterconnect.F0276: link-

Alert Details

✖ Server.F0283: link-down

Source: Blade 4

Full Path Name: Cisco UCS Instances\bb04-6248\Chassis 2 \Blade 4

Alert Rule: Fault Rule : Server.F0283 (link-down)

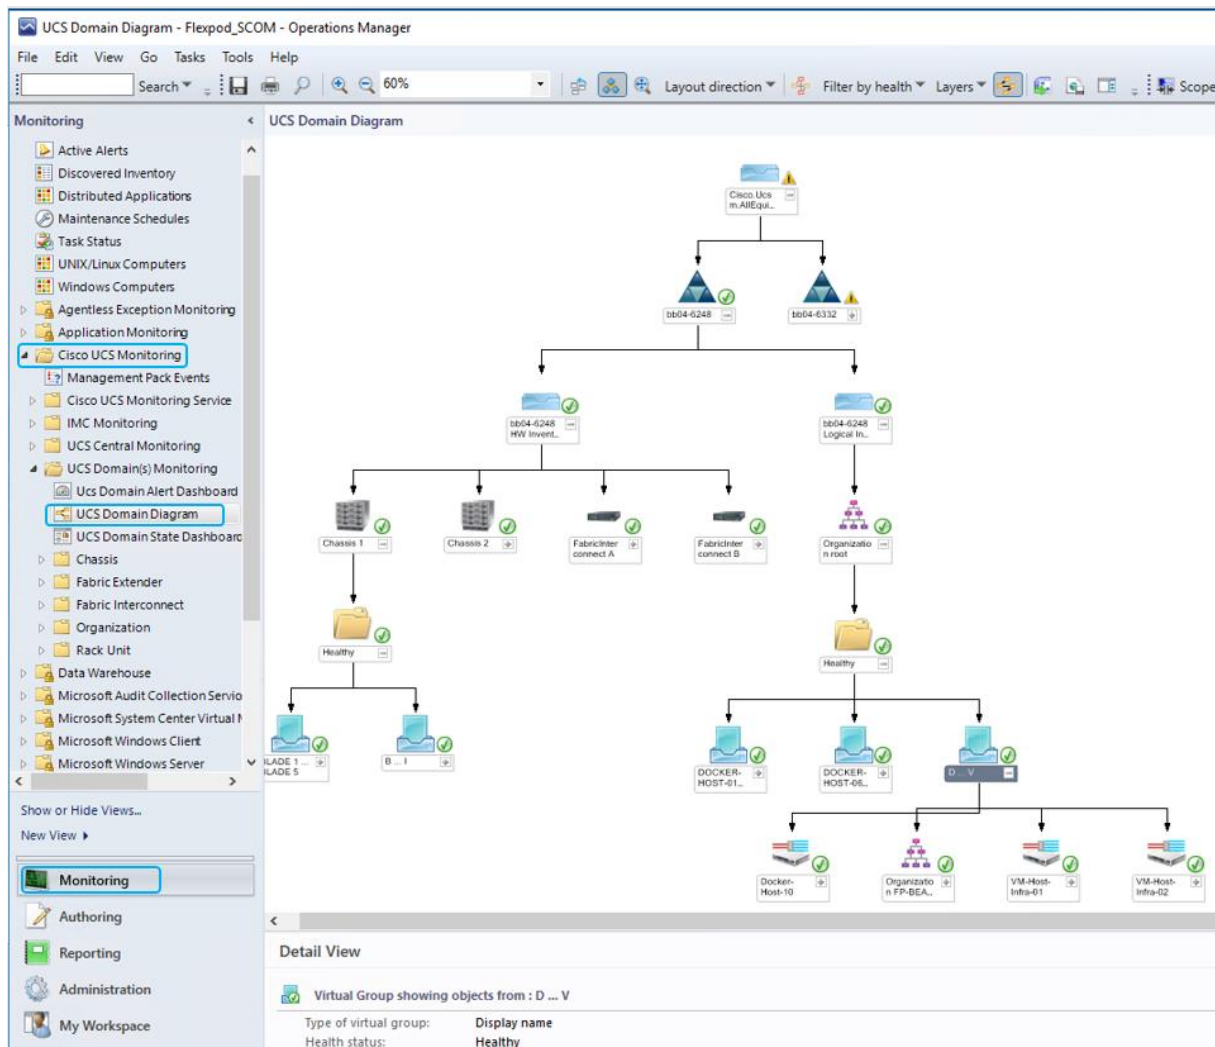
Created: 6/3/2017 7:47:19 PM

Alert Description

[Instance Name:bb04-6248; DN:sys/chassis-2/blade-4/fabric-B/path-1/vc-750/]

Description: fc VIF 750 on server 2 / 4 of switch B down, reason: None

- UCS Domain Diagram—Displays a graphical view of the relationship between different Cisco UCS Domain(s) components for all Instances.



- UCS Domain State Dashboard—Displays the list of domains added and its health state and other inventory information.
- When you select a UCS domain from the State dashboard, you can perform the tasks listed in the following sections.
 - Generating Cisco UCS Domain Technical Support Bundle
 - Launching UCS GUI
 - Loading the UCS Inventory Data
 - Ping UCS
 - Ping UCS Continuously
 - Physical and Logical Inventory
 - Launching KVM Console
 - Alert Operations

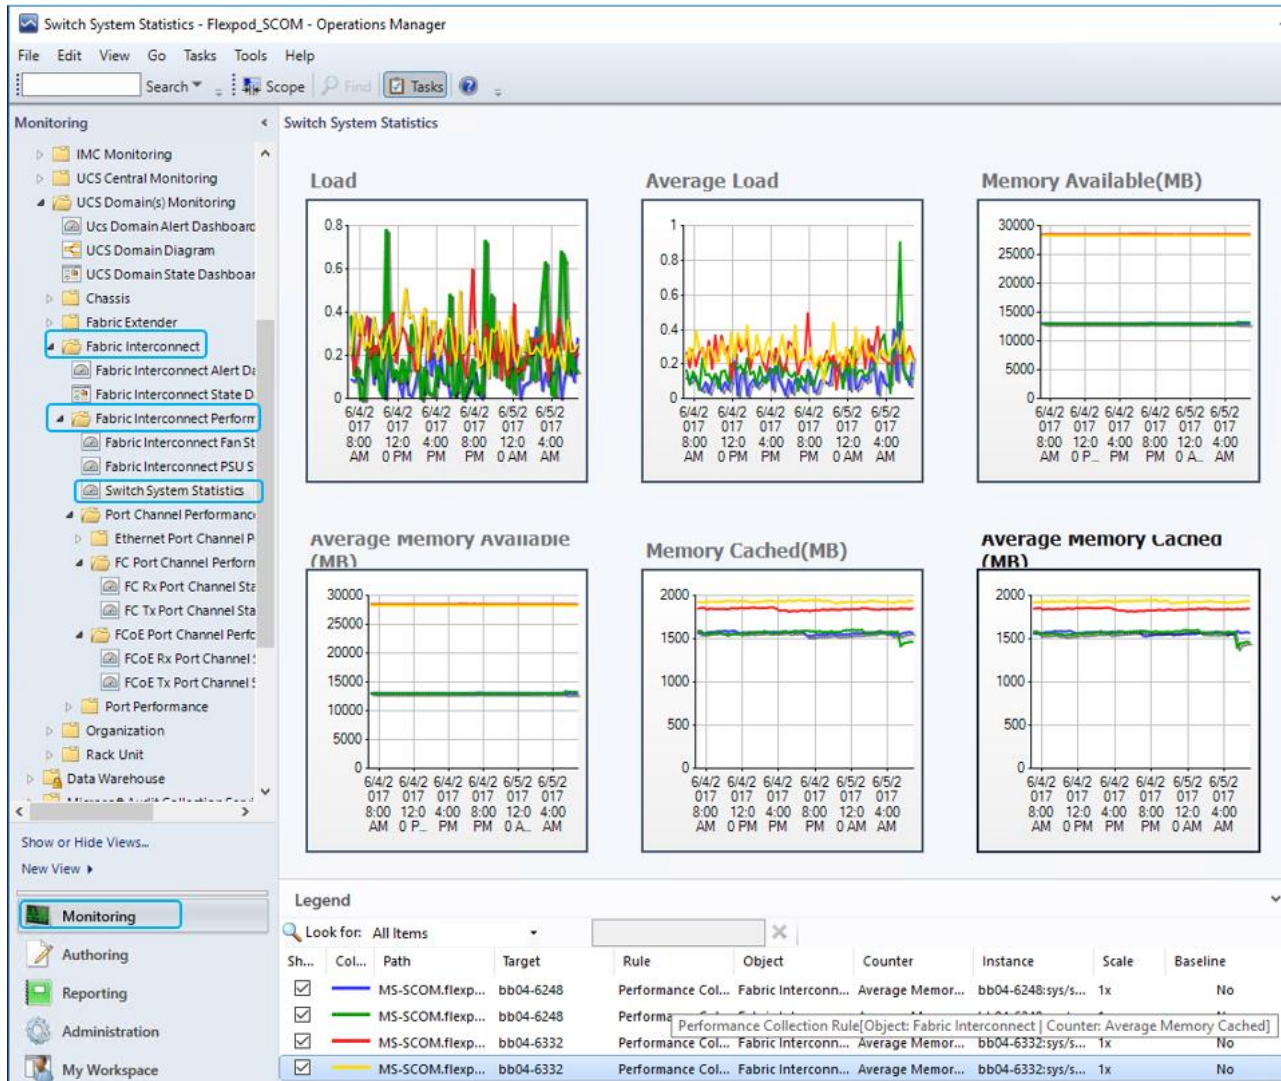
The screenshot displays the 'UCS Domain State Dashboard' in the 'Flexpod_SCOM - Operations Manager' application. The interface is divided into several sections:

- Monitoring (Left Sidebar):** A tree view showing various monitoring categories. 'Cisco UCS Monitoring' and 'UCS Domain State Dashboard' are highlighted with blue boxes.
- UCS Domain (2) Table:** A table listing UCS domains with columns for Health, Display Name, UCSM Version, and Collect Performance.

Health	Display Na...	UCSM Version	Collect Perfor
Success	bb04-6248	3.1(3a)	true
Warning	bb04-6332	3.1(3a)	true
- Details (Right Pane):** A detailed view of the selected domain 'bb04-6248'.

Display Name	bb04-6248
Path	Cisco UCS Instances\bb04-6248
Health	Success
Object Display Name	bb04-6248
Unique Id	ac883ccd763f480db1ca3d65d16207e9
Description	
Monitoring Server	MS-SCOM.flexpod.local
Web Proxy Url	http://MS-SCOM.flexpod.local:8732/UcsMonitoringService
Class	UCS Domain
UCS Name	bb04-6248
UCSM Version	3.1(3a)
URL	https://192.168.156.50:443
Monitoring Server	MS-SCOM.flexpod.local
Collect Performance Statistics	true
- Tasks (Far Right):** A list of tasks for the selected domain. 'Create and Download a Tech...' is highlighted with a blue box.

- You can view performance metrics for the various Cisco UCS components as shown the below figure.



Cisco UCS Manager Plug-in for SCVMM

Using the Cisco UCS Manager add-in you can view the details such as properties, faults information, and firmware details of the servers (blades or rack-mount servers) on which the host is running.

Cisco UCS Manager Plug-in Installation

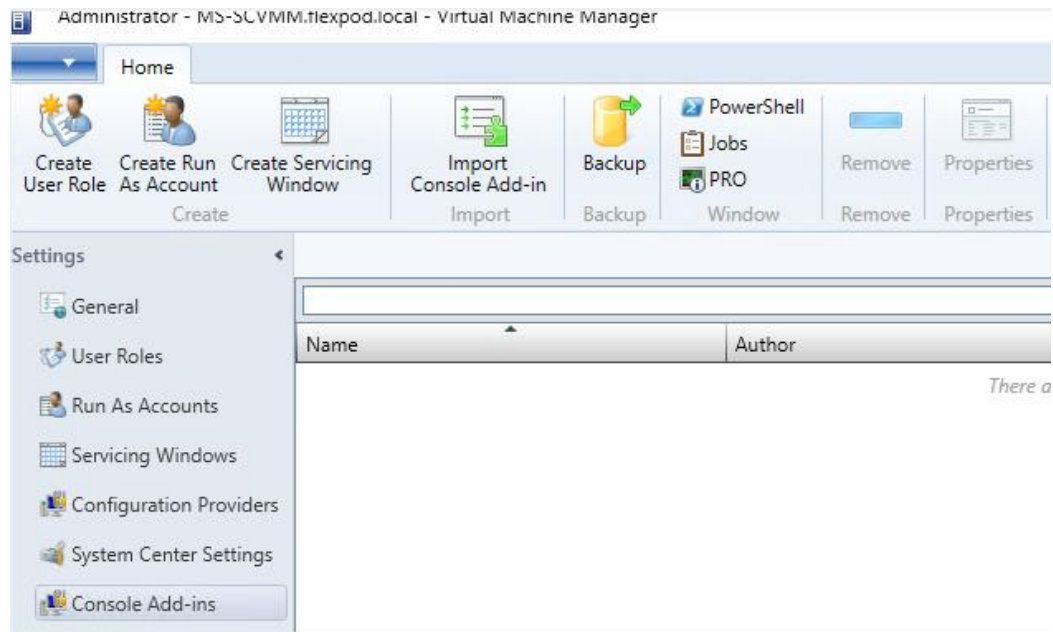
Perform the following steps to install the Cisco UCS virtual machine manager add-in:

1. Open <https://software.cisco.com/download/type.html?mdfid=286282669&flowid=72562>
2. Click Unified Computing System (UCS) Microsoft System Center Virtual Machine Manager to view the list of available versions for download (CiscoUCS-Scvmm-1.1.2.zip).
3. Download and save the zipped folder.

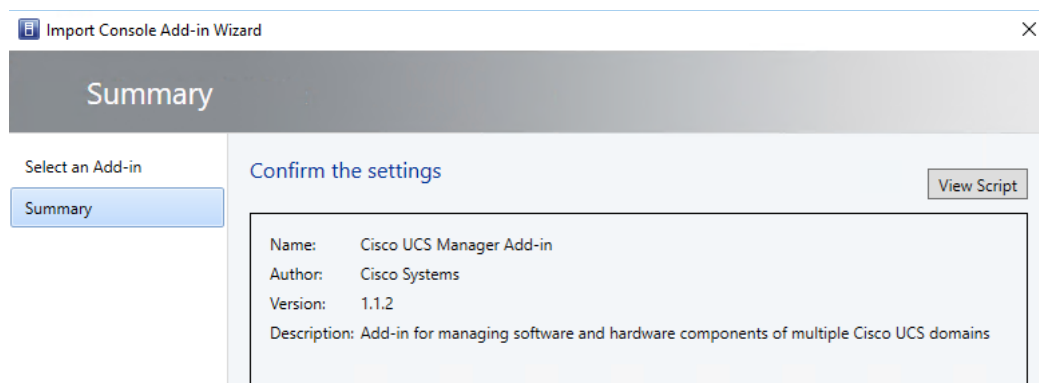


The add-in is made available as a zipped file that has to be imported into the virtual machine manager to install it.

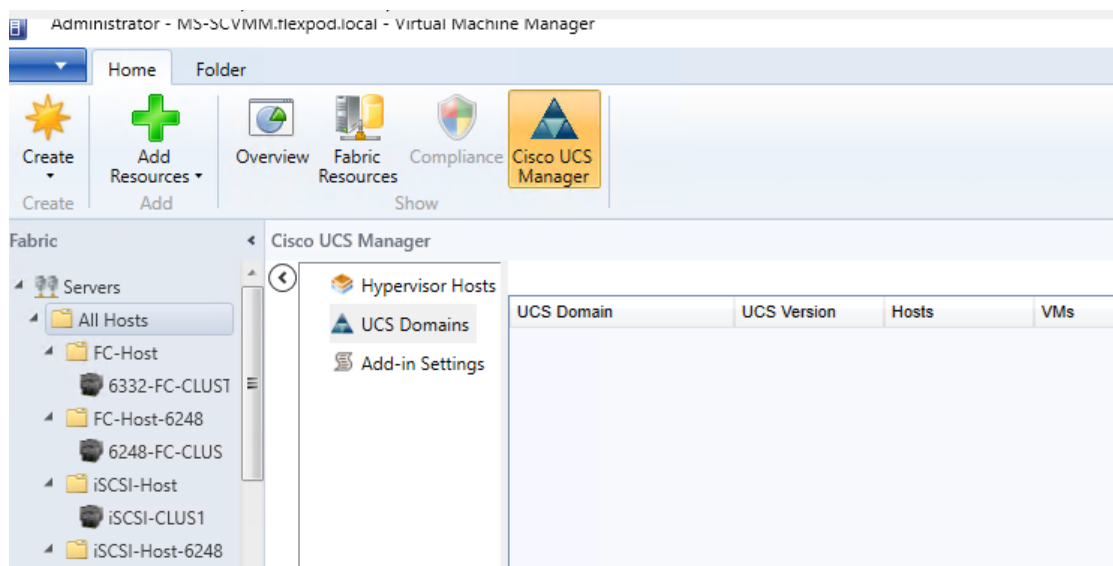
4. Open an instance of the Virtual Machine Manager console.
5. In the Navigation pane, click Settings.
6. In the toolbar, click Import Console Add-in. The Import Console Add-in wizard appears.



7. Click Browse and navigate to the location where the zipped file is saved.
8. Select the zip file and click Open. Click Next. Click Finish.



9. The add-in is installed and a new icon called Cisco UCS Manager appears in the toolbar.

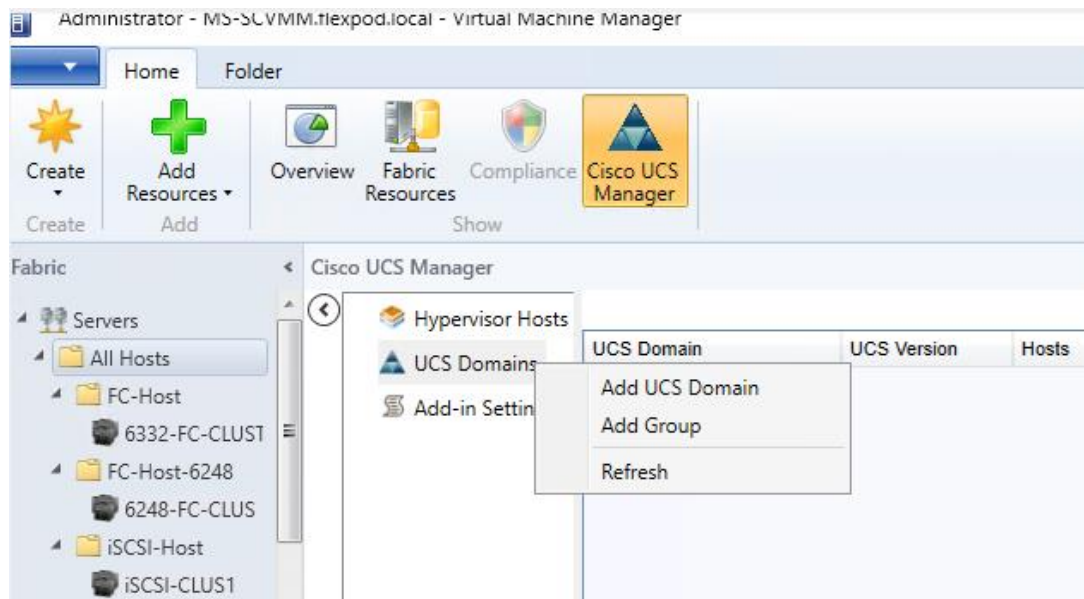


Cisco UCS Domain Registration:

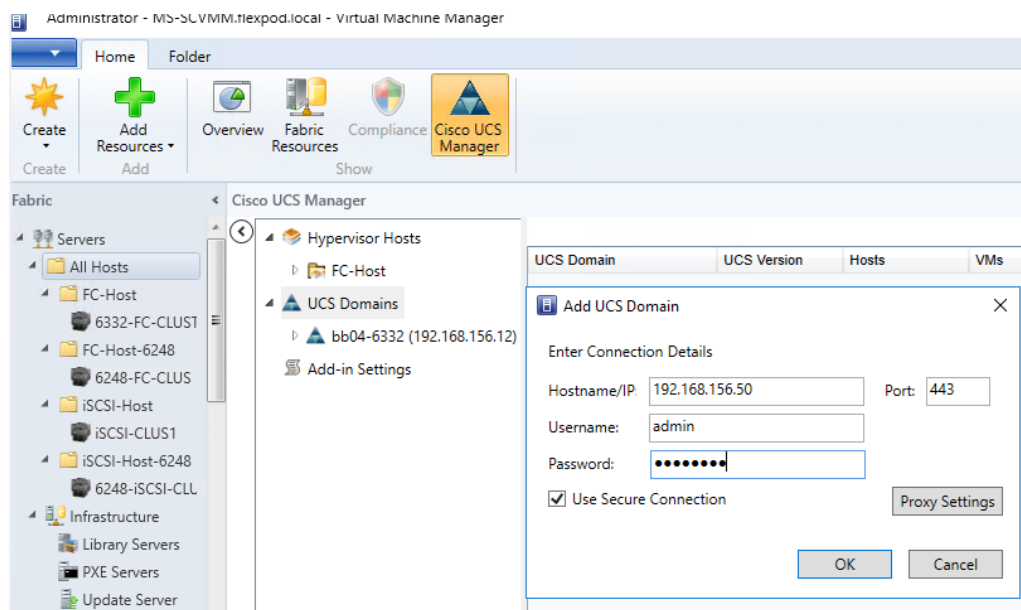
You can register domains using any access privileges. Depending on the privileges available to the user with which UCS domain is registered, some or all actions may be disabled.

Perform the following steps to register a UCS domain:

1. On the toolbar, click Cisco UCS Manager.
2. Right-click on UCS Domains.
3. Click Add UCS Domain.
4. The Add UCS Domain dialog box appears.



5. Enter the following details in the dialog box:



If required, you can edit the UCS domain details at a later time.

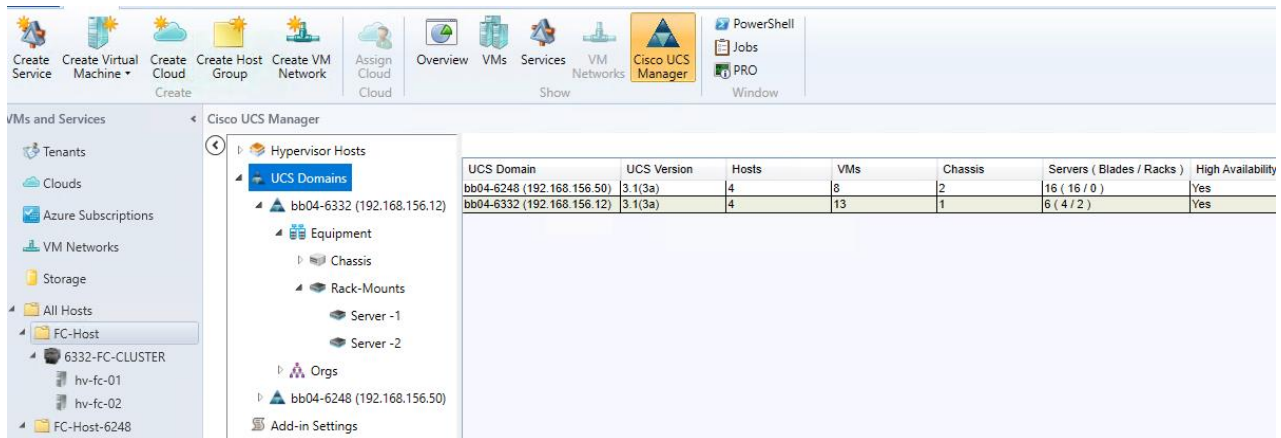
6. Click Proxy Settings. Proxy Settings dialog box appears.
7. In the Proxy Settings dialog box, click Use Custom Proxy Settings radio button and enter the following details:



If required, you can edit the proxy settings at a later time.

8. Click OK.

The registered UCS domain appears under the UCS domains node. Upon adding a UCS domain, the Hyper-Visor hosts running on the newly added UCS domain appear under the Hyper-Visor Hosts node.



UCS Domain	UCS Version	Hosts	VMs	Chassis	Servers (Blades / Racks)	High Availability
bb04-6248 (192.168.156.50)	3.1(3a)	4	8	2	16 (16 / 0)	Yes
bb04-6332 (192.168.156.12)	3.1(3a)	4	13	1	6 (4 / 2)	Yes



You also can add UCS domains within groups. If you want to add a UCS domain within a group, right-click on the group and follow steps 3 through step 7 in the preceding procedure.

Using the Cisco UCS SCVMM Plugin

Viewing the Server Details from the Hypervisor Host View

The following procedure provides steps to view server details:

1. On the toolbar, click Cisco UCS Manager.
2. In the Hypervisors node, select the Hypervisor host which is associated with the server.

Name	Description
General tab	
Fault summary	Displays the number of faults categorized based on fault severity. You can click on the severity fault icons in this section to view the fault details.
Properties	Displays the properties of the server such as, server ID, UUID, serial number, associated service profiles and so on. If a service profile is associated with the server, a link to the location of the service profile is provided. Clicking on the link displays the properties of the associated service profile.
Status	Indicates the status of the tasks running on the host.
Actions area	
Associate Service Profile	Enables you to associate a service profile to the server.
Disassociate Service Profile	Enables you to disassociate a service profile from the server.

Set Desired Power State	Provides options to set the power state of a service profile.
KVM Console	Enables you to launch the KVM console.
Turn on Locator LED	Enables you to either turn on or turn off the locator LED depending on the current state.
Firmware tab	Provides the firmware details such as BIOS, CIMC, adaptors and storage device part IDs, and the firmware versions. If there are any changes to the firmware details on the server, those changes will reflect here.
Faults tab	Displays the faults' details specific to the server, such as properties, severity, fault codes and IDs, description, affected objects, and so on. Provides options to filter the faults based on severity, and option under the Actions area to acknowledge the fault on UCS.

- On the right pane of the window, you can view the following information of the server on which the host is running:

The screenshot displays the Cisco UCS Manager interface. The left-hand navigation pane shows a tree structure under 'Hypervisor Hosts' and 'UCS Domains'. The right-hand pane is divided into two main sections: 'Fault Summary' and 'Properties'. The 'Fault Summary' section shows four categories of faults (Critical, Major, Minor, Warning) with zero counts each, and a suppression status of 'N/A'. The 'Properties' section provides detailed information about the UCS domain, including its name (bb04-6332), slot ID (1), chassis ID (1), product name (Cisco UCS B200 M4), vendor (Cisco Systems Inc), revision (0), and various hardware specifications like 28 cores, 262144 MB of memory, and 2 HBAs.

Viewing Registered UCS Domains

- On the toolbar, click Cisco UCS Manager.
- Click UCS Domains. The list of registered UCS domains and consolidated UCS information for each domain, such as the name and version, number of associated hosts, VMs and servers appear on the right pane of the window as shown in the above figure.
- (Optional) You can view the details in the grid view or the card view by clicking View option on the right-top corner and choosing the appropriate option.

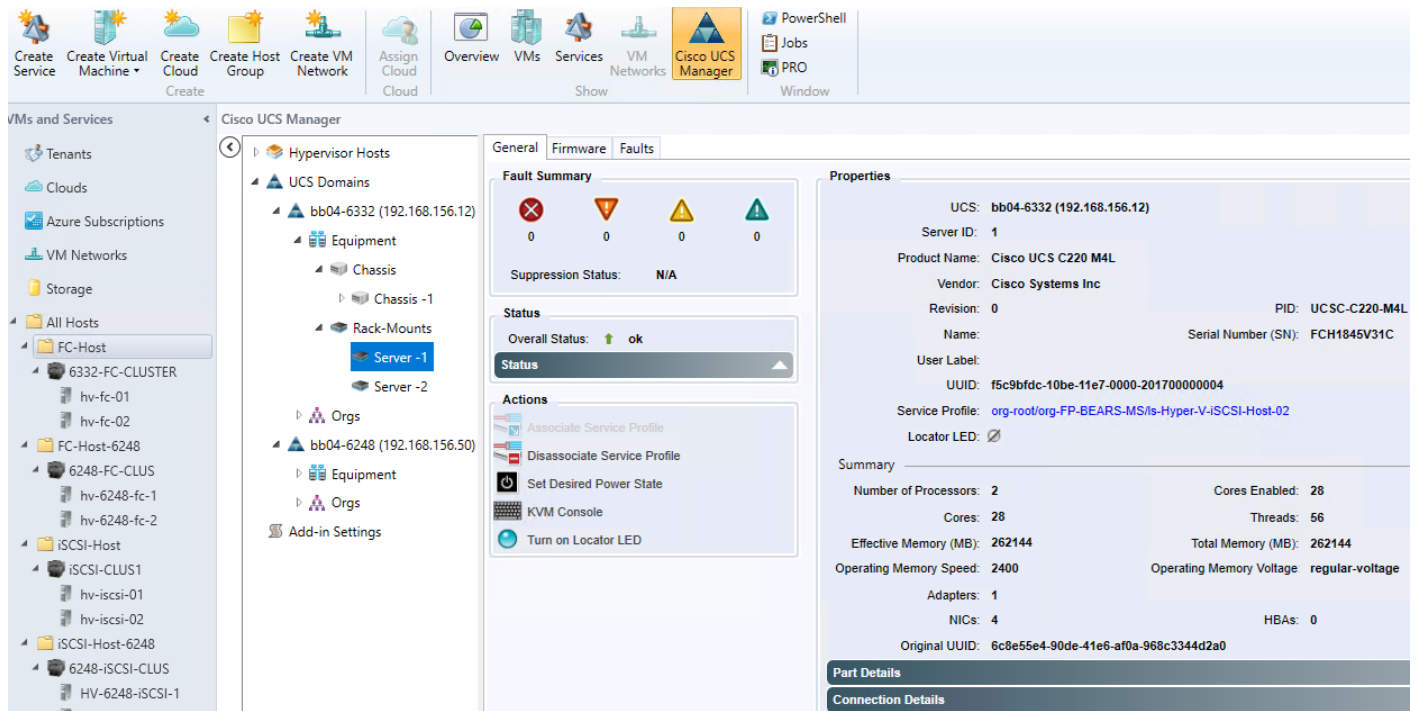
Viewing the UCS Blade Server Details

Using the add-in you can view the server details, such as properties, faults information, and firmware details. The following procedure provides steps to view server details:

1. On the toolbar, click Cisco UCS Manager.
2. Under the UCS Domains node, expand the UCS domain.
3. Expand Equipment > Chassis. A list of chassis appears.
4. Choose a chassis.
5. The list of blade servers on the chassis appears under the chassis on the left pane. You can also view the list of blade servers on the right pane on the window.
6. Select the blade for which you want to view the details.
7. The properties of the blade appear on the right pane of the window. You can view the following server information as shown in the table below.

Name	Description
General tab	
Fault summary	Displays the number of faults categorized based on fault severity. You can click on the severity fault icons in this section to view the fault details.
Properties	Displays the properties of the server such as, chassis ID, UUID, serial number, associated service profiles and so on. If a service profile is associated with the blade, a link to the location of the service profile is provided. Clicking on the link displays the properties of the associated service profile.
Status	Indicates the status of the server.
Actions area	
Set Desired Power State	Provides options to set the power state of a service profile.
KVM Console	Enables you to launch the KVM console.
Rename Service Profile	Enables you to rename a service profile.
Associate Service Profiles	Enables you to associate a service profile to the server.
Turn on Locator LED	Enables you to either turn on or turn off the locator LED depending on the current state.
Disassociate Service Profile	Enables you to disassociate a service profile from the server.
Firmware tab	Provides the firmware details such as BIOS, CIMC, adaptors and storage device part IDs and the firmware versions. If there are any changes to the firmware details on the server, those

	changes will reflect here.
Faults tab	Displays the faults' details specific to the server such as properties, severity, fault codes and IDs, description, affected objects, and so on. Provides options to filter the faults based on severity, and option under the Actions area to acknowledge the fault on UCS.



Viewing the UCS Rack-Mount Server Details:

Using the add-in you can view the details such as properties, faults information, and firmware details of the servers on which the host is running. The following procedure provides steps to view server details:

1. On the toolbar, click Cisco UCS Manager.
2. In the UCS Domains node, expand the UCS domain.
3. Expand Equipment > Rack-Mounts.
4. The list of registered UCS rack-mount servers appears.
5. Choose the server for which you want to view the details.
6. The properties of the rack-mount server appear on the right pane of the window. You can view the following server information:

Name	Description
General tab	

Fault summary	Displays the number of faults categorized based on fault severity. You can click on the severity fault icons in this section to view the fault details.
Properties	Displays the properties of the server such as, server ID, UUID, serial number, associated service profiles and so on. If a service profile is associated with the server, a link to the location of the service profile is provided. Clicking on the link displays the properties of the associated service profile.
Status	Indicates the status of the server.
Actions area	
Set Desired Power State	Provides options to set the power state of a service profile.
KVM Console	Enables you to launch the KVM console.
Rename Service Profile	Enables you to rename a service profile.
Associate Service Profiles	Enables you to associate a service profile to the server.
Turn on Locator LED	Enables you to either turn on or turn off the locator LED depending on the current state.
Disassociate Service Profile	Enables you to disassociate a service profile from the server.
Firmware tab	Provides the firmware details such as BIOS, Cisco IMC, adaptors and storage device part IDs and the firmware versions. If there are any changes to the firmware details on the server, those changes will reflect here.
Faults tab	Displays the faults' details specific to the server such as properties, severity, fault codes and IDs, description, affected objects, and so on. Provides options to filter the faults based on severity, and option under the Actions area to acknowledge the fault on UCS.

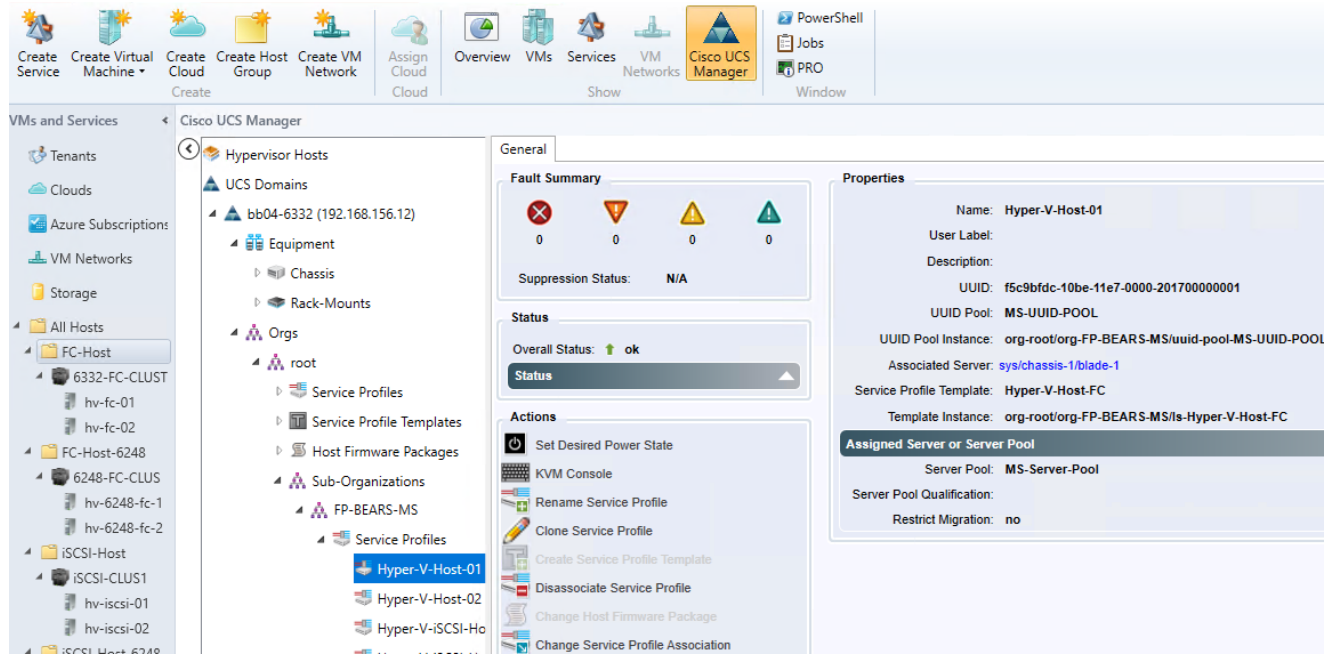
Viewing the Service Profile Details

Using the add-in you can view the service profile details, such as properties, and faults information. The following procedure provides steps to view the service profile details:

1. On the toolbar, click Cisco UCS Manager.
2. In the UCS Domains node, expand the UCS domain.
3. Expand Orgs > root.
4. Choose Service Profiles.
5. The list of service profiles and associated information appear on the right pane of the window. The server column lists the links to the servers that the service profile is associated with. Click the link to view the details of the server.
6. Click the service profile for which you want to view the details.

7. The service profile details appear on the right pane of the window. You can view the following service profile information:

Name	Description
General tab	
Fault summary	Displays the number of faults and the severity of the faults.
Properties	Displays the properties of the service profile such as, name, associated server, service profile template used and so on.
Status	Indicates the status of the service profile.
Actions area	
Set Desired Power State	Provides options to set the power state of the server.
KVM Console	Enables you to launch the KVM console.
Rename Service Profile	Enables you to rename a service profile.
Create a Clone	Enables you to create a clone of the service profile by inheriting the attributes of the service profile.
Disassociate Service Profile	Enables you to disassociate the service profile from the server.
Change Host Firmware Package	Enables you to change the host firmware association.
Change Service Profile Association	Enables you to upgrade the host firmware on the servers.

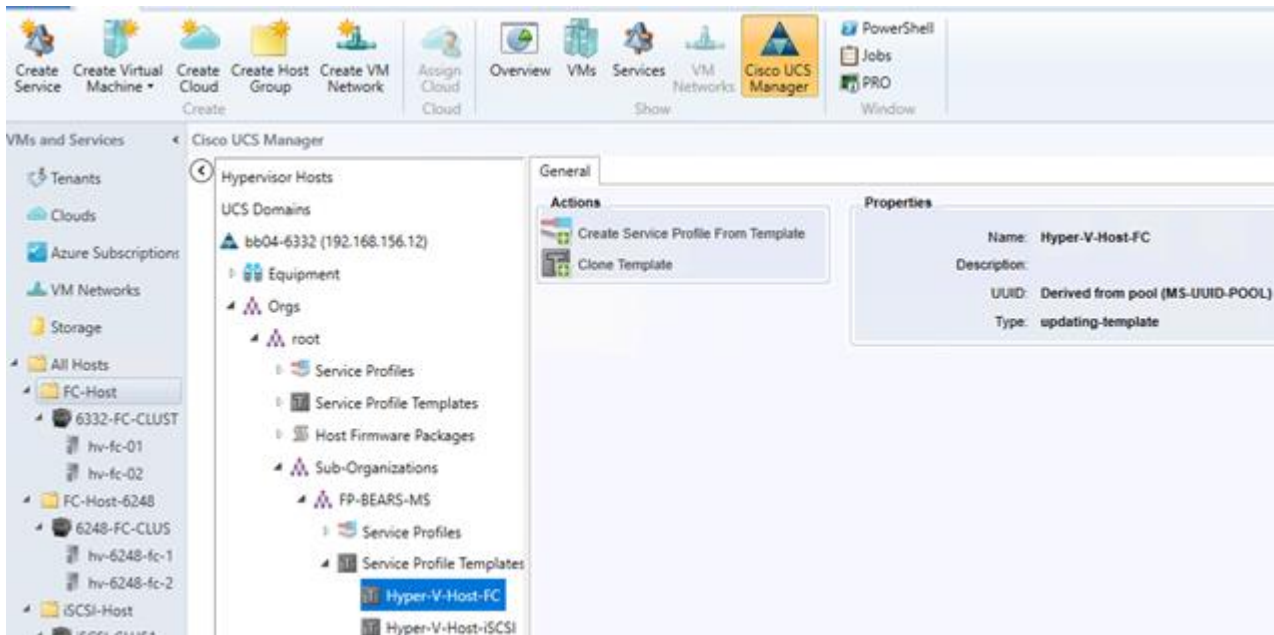


Viewing the Service Profile Template Details

Using the add-in you can view the service profile template details, such as properties, and faults information. The following procedure provides steps to view the service profile template details:

1. On the toolbar, click Cisco UCS Manager.
2. In the UCS Domains node, expand the UCS domain.
3. Expand Orgs > root.
4. Expand Service Profile Templates and select the service profile template for which you want to view the details.
5. You can view the following service profile template information on the right pane of the window:

Name	Description
General tab	
Properties area	Displays the properties of the service profile template, such as name, type and so on.
Actions area	
Create Service Profile from Templates	Enables you to use the template to create a service profile.
Create a Clone	Enables you to create a clone of the service profile template by inheriting the attributes of the



Viewing the Host Firmware Package Details

Using the add-in you can view the host firmware packages properties. The following procedure provides steps to view the host firmware packages details:

1. On the toolbar, click Cisco UCS Manager.
2. In the UCS Domains node, expand the UCS domain.
3. Expand Orgs > root.
4. Expand Host Firmware Packages and select the host firmware package for which you want to view the details.

You can view the following host firmware package information on the right pane of the window:

Name	Description
General tab	
Properties area	Displays the properties of the host firmware package, such as name, description, ownership information, package version and so on.
Actions area	
Modify Package Versions	Enables you to modify Blade package version and Rack package version properties.

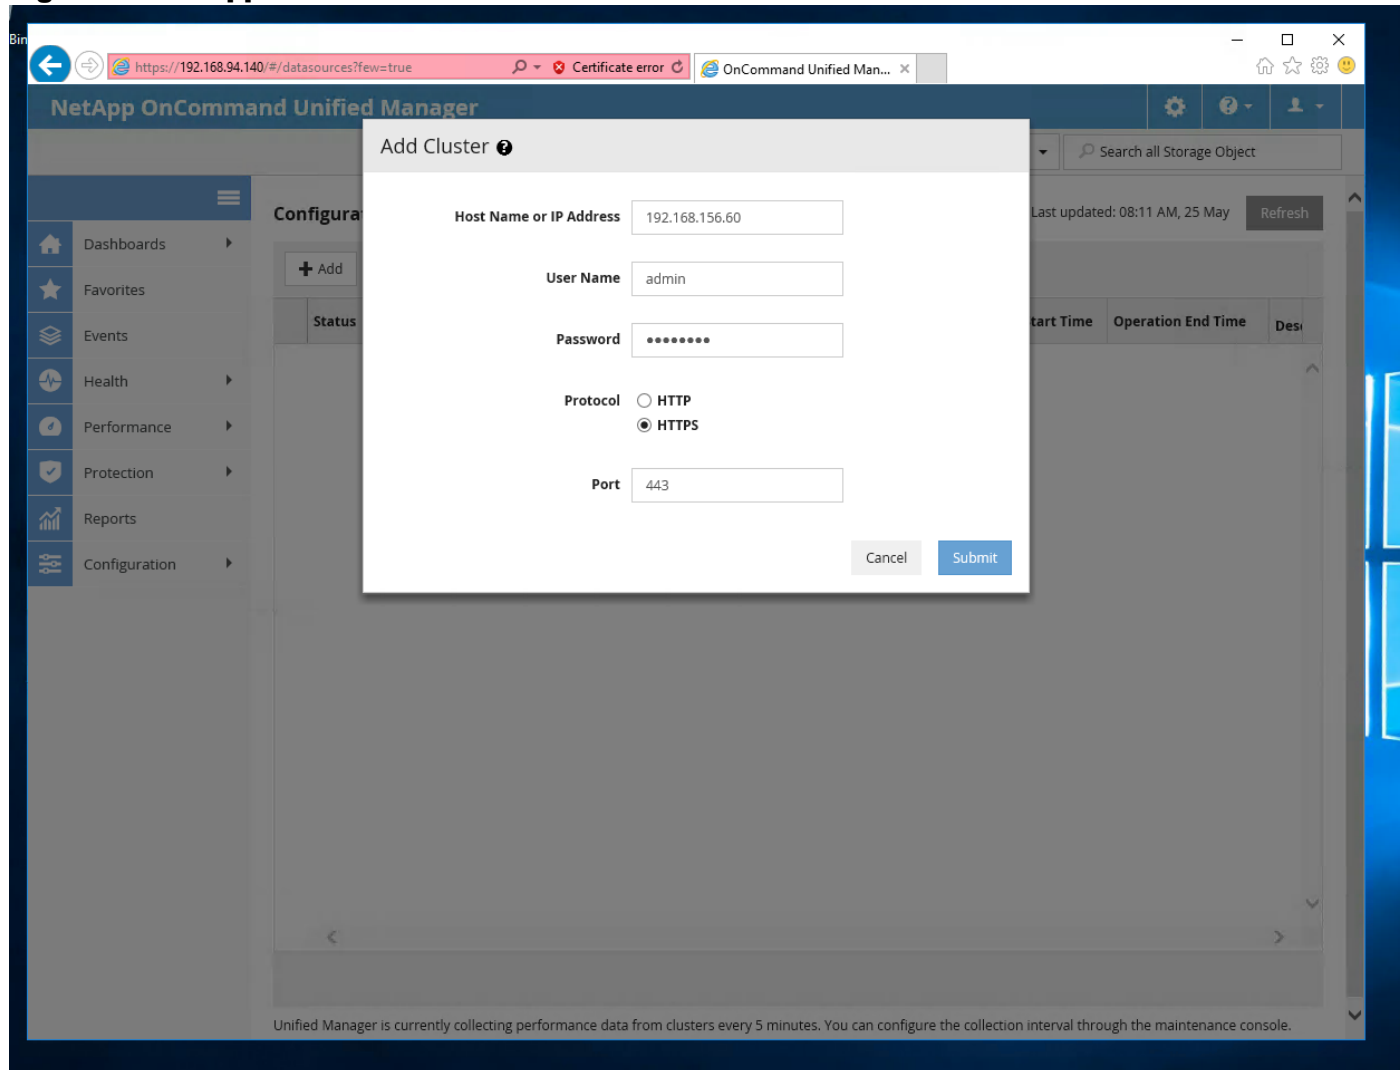
FlexPod Management Tools Setup

OnCommand Unified Manager 7.2

1. Deploy a Windows Server 2016 Virtual Machine for the OnCommand installation.
2. Download and review the [OnCommand Unified Manager 7.2 Installation and Setup Guide for Microsoft Windows](#)
3. Download OnCommand Unified Manager version 7.2 from <http://mysupport.netapp.com> to the virtual machine.
4. Follow the instructions to install OnCommand after double-clicking on the executable.
5. After installation, fill in Setup Email, Time Settings and enable Autosupport
6. To add your cluster for monitoring, Click on Add from the main dashboard as seen in figure 7 and fill in your cluster details as seen in figure 8.

Figure 7 NetApp OnCommand Dashboard

The screenshot shows the NetApp OnCommand Unified Manager interface. The browser address bar displays `https://192.168.94.140/#/datasources?few=true` with a 'Certificate error' warning. The page title is 'NetApp OnCommand Unified Manager'. A search bar at the top right contains the text 'Search all Storage Object'. The left sidebar contains a navigation menu with the following items: Dashboards, Favorites, Events, Health, Performance, Protection, Reports, and Configuration. The main content area is titled 'Configuration / Cluster Data Sources' and includes a 'Last updated: 08:11 AM, 25 May' timestamp and a 'Refresh' button. Below the title are buttons for '+ Add', 'Edit', 'Remove', and 'Rediscover'. A table with the following columns is visible: Status, Name, Host, Protocol, Port, User Name, Operation, Operation State, Operation Start Time, Operation End Time, and Description. The table is currently empty. At the bottom of the page, a message states: 'Unified Manager is currently collecting performance data from clusters every 5 minutes. You can configure the collection interval through the maintenance console.'

Figure 8 NetApp OnCommand - Add a cluster

NetApp SnapDrive 7.1.3

SnapDrive 7.1.3 for Windows enables you to automate storage provisioning tasks and to manage data in physical or virtual Microsoft Windows hosts, in SMB 3.0 environments.

You can use the following steps to install SnapDrive 7.1.3 for Windows on Microsoft Hyper-V Hosts. These steps assume that prerequisites have been verified.

Configuring access for SnapDrive for Windows

You can use the following steps to configure access for SnapDrive 7.1.3 for Windows:

1. Create a user account on the storage system by entering the following command: `security login create -vserver -user -authentication-method -application -role`.
The variables represent the following values:
 - Vserver is the name of the Vserver for the user to be created.
 - User is the name of the SnapDrive user.

- authentication-method is the method used for authentication
- application is the option you use to specify how the user will access.
- role is the privileges this user will have.

For example,

The following command adds a user called "snapdrive" to the BUILTIN\Administrators group on the storage system:

```
security login create -vserver Infra-MS-SVM -user snapdrive -authentication
method password -application ssh -role vsadmin
```



You must provide this user name later in this procedure. Therefore, make a note of the user name, including the letter case (lowercase or uppercase) of each character in the user name.

2. Enter a password, when prompted to do so, for the user account you are creating.
3. You are prompted to enter the password twice. You are required to provide this password later, so make a note of it, including letter case.
4. Check to ensure that the user account you just created belongs to the local administrator's group on the storage system by entering the following command:

```
security login show
```

For additional information, see the section about creating local groups on the storage system in the Data ONTAP File Access and Protocols Management Guide for 7-Mode.

5. On each Windows host that needs access to the storage system, create a local user account with administrative rights on the host, using the same user name and password that you specified in Step 1 and Step 2.



Set up the local user account so that the password for the account never expires. For detailed instructions on how to create local user accounts, see your Windows documentation.

Downloading SnapDrive for Windows

Before installing SnapDrive for Windows, you must download the software package from the NetApp Support Site.

Before you begin you need the NetApp Support Site login credentials .

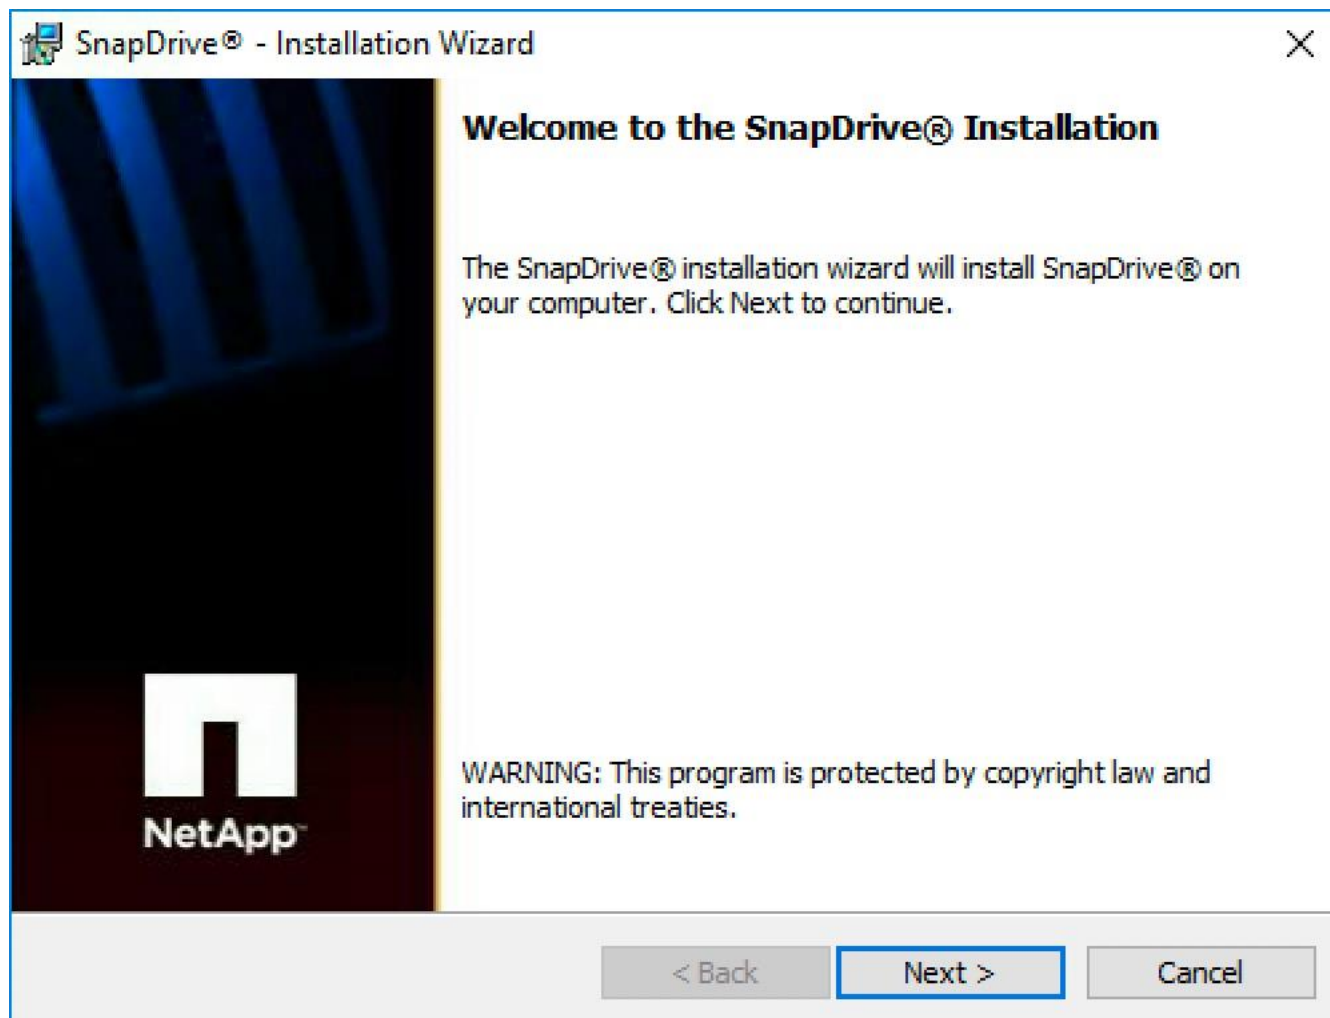
Steps

1. Log in to the NetApp Support Site.
2. Go to the Download Software page.
3. From the drop-down list, select the operating system on which you are installing SnapDrive and click Go!.
4. Click View & Download for the software version you want to install.
5. On the SnapDrive for Windows Description page, click Continue.
6. Review and accept the license agreement.

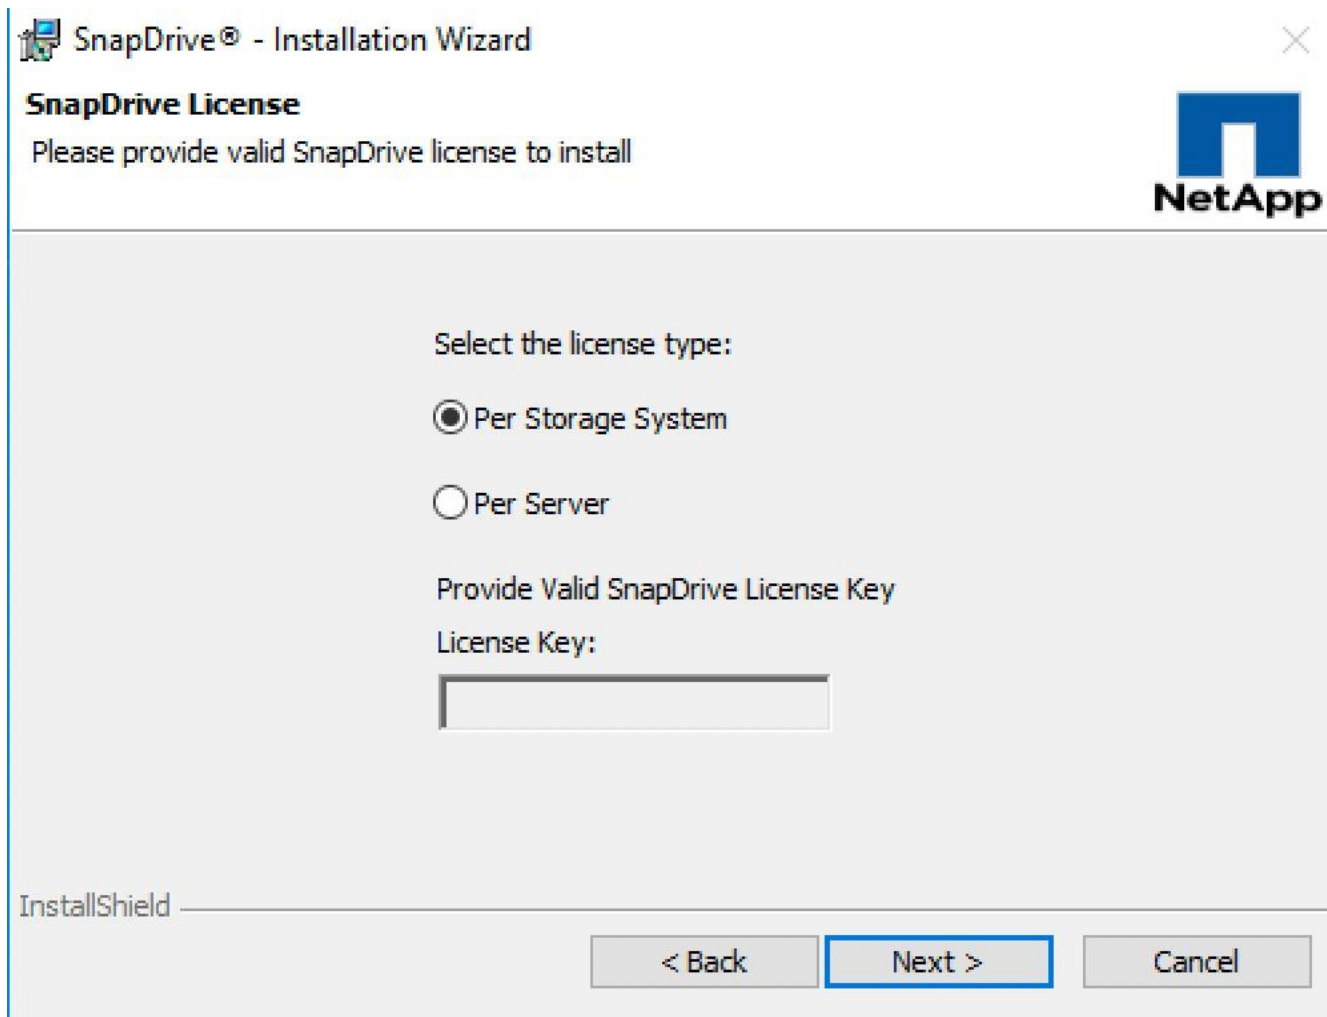
7. On the Download page, click the link for the installation file.
8. Save the SnapManager for Hyper-V file to a local or network directory.
9. Click Save File.
10. Verify the checksum to ensure that the software downloaded correctly.

Installing SnapDrive for Windows

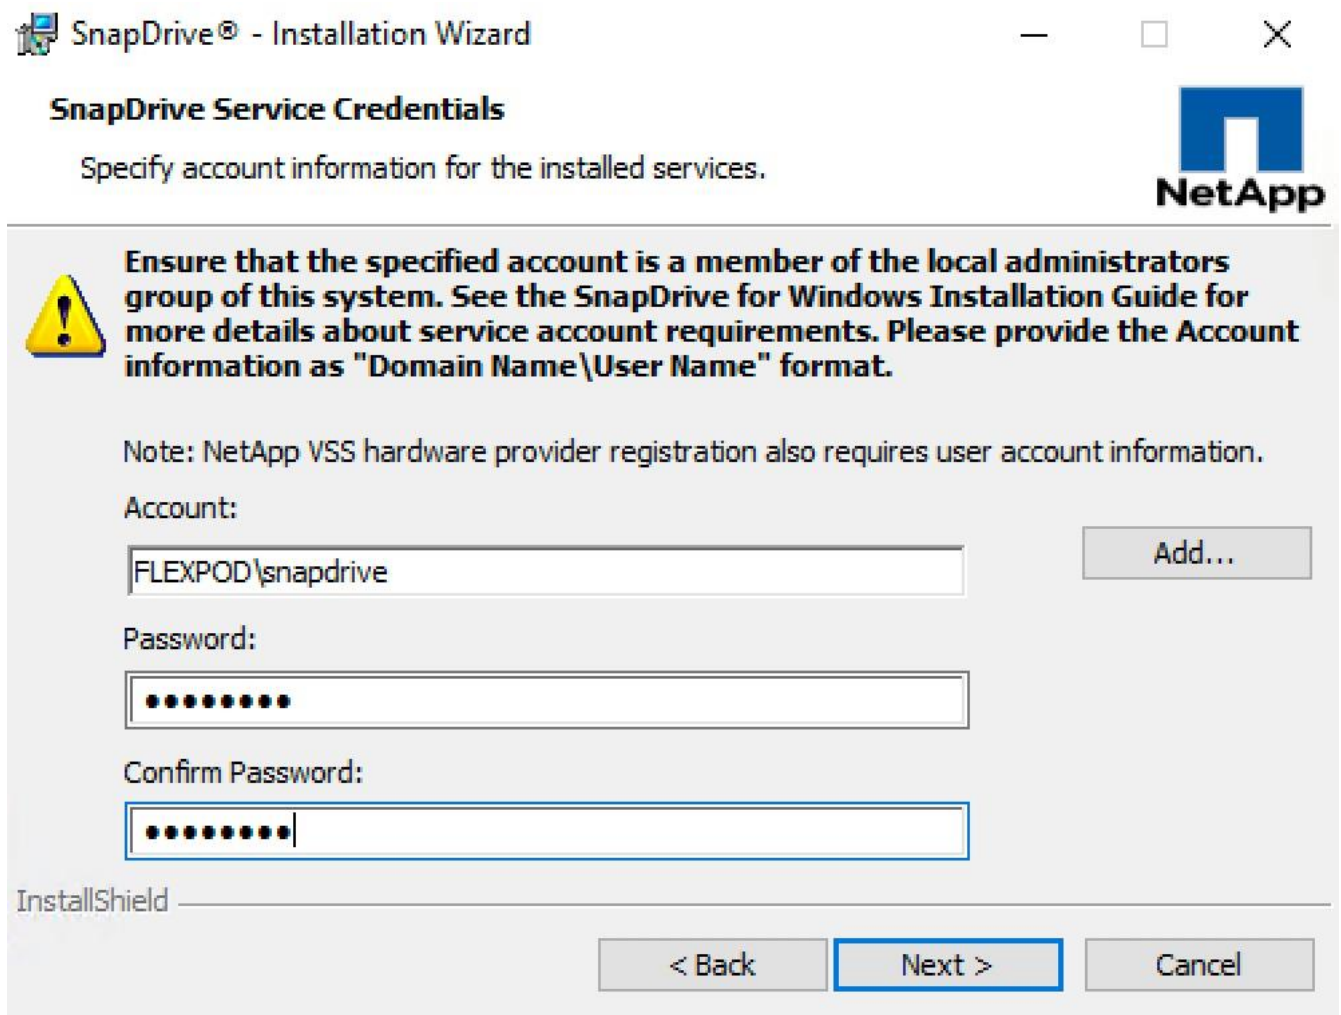
1. Launch the SnapDrive for Windows installer, and then follow the Installation wizard instructions.



2. In the SnapDrive License screen, select the appropriate license type.



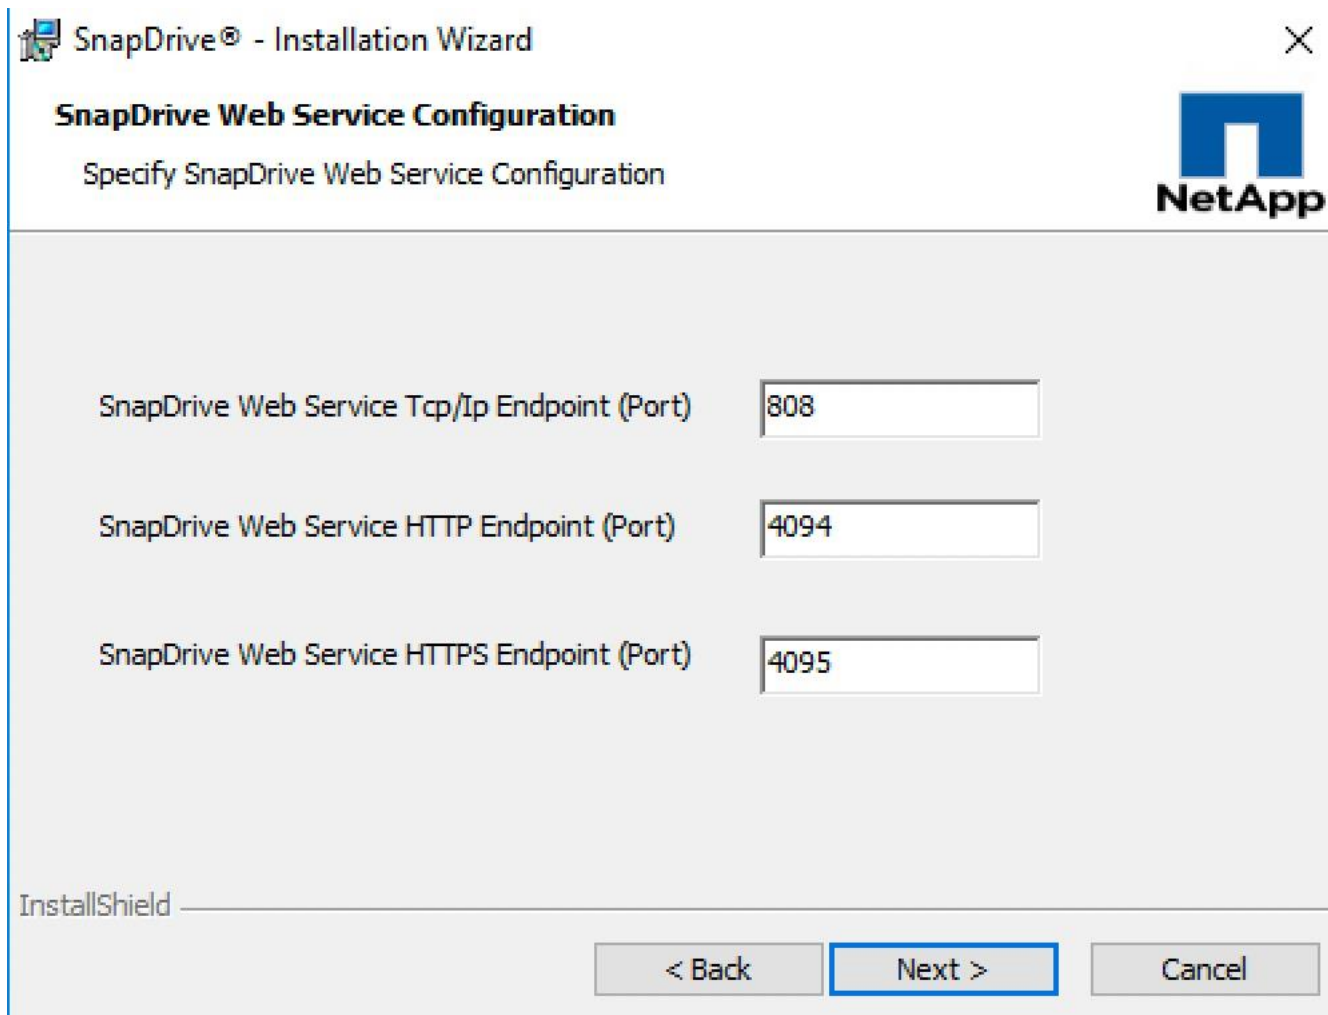
3. In the Customer Information screen, enter the appropriate information.
4. In the Destination Folder screen, enter the appropriate destination or accept default.
5. In the SnapDrive Credentials screen, enter the account and password information of an account that is a member of the local administrators



The image shows a Windows installation wizard window titled "SnapDrive® - Installation Wizard". The window has a title bar with standard minimize, maximize, and close buttons. Below the title bar, the text "SnapDrive Service Credentials" is displayed in a bold font, followed by the instruction "Specify account information for the installed services." The NetApp logo is located in the top right corner. A yellow warning triangle icon is on the left side of the main content area. The main text reads: "Ensure that the specified account is a member of the local administrators group of this system. See the SnapDrive for Windows Installation Guide for more details about service account requirements. Please provide the Account information as 'Domain Name\User Name' format." Below this, a note states: "Note: NetApp VSS hardware provider registration also requires user account information." There are three input fields: "Account:" with the text "FLEXPOD\snapdrive" and an "Add..." button to its right; "Password:" with a masked field of ten dots; and "Confirm Password:" with a masked field of ten dots. At the bottom left, the text "InstallShield" is visible. At the bottom right, there are three buttons: "< Back", "Next >" (which is highlighted with a blue border), and "Cancel".

6. In the SnapDrive Web Service Configuration screen, accept the default port numbers.

If you want to change the port numbers, you should also change the port numbers for other SnapDrive hosts.




7. In the Preferred IP Address screen, identify the IP address you want to use to communicate with the storage system.

You should configure the preferred IP address, because doing this improves performance and scalability.

SnapDrive® - Installation Wizard

Preferred Storage System IP Address

Configure SnapDrive to use a preferred IP Address



Enable preferred storage system IP Address

Configure SnapDrive to use a preferred IP Address for management traffic. If storage system has only one interface, setting a preferred IP Address can prevent issues if more interfaces are added later.

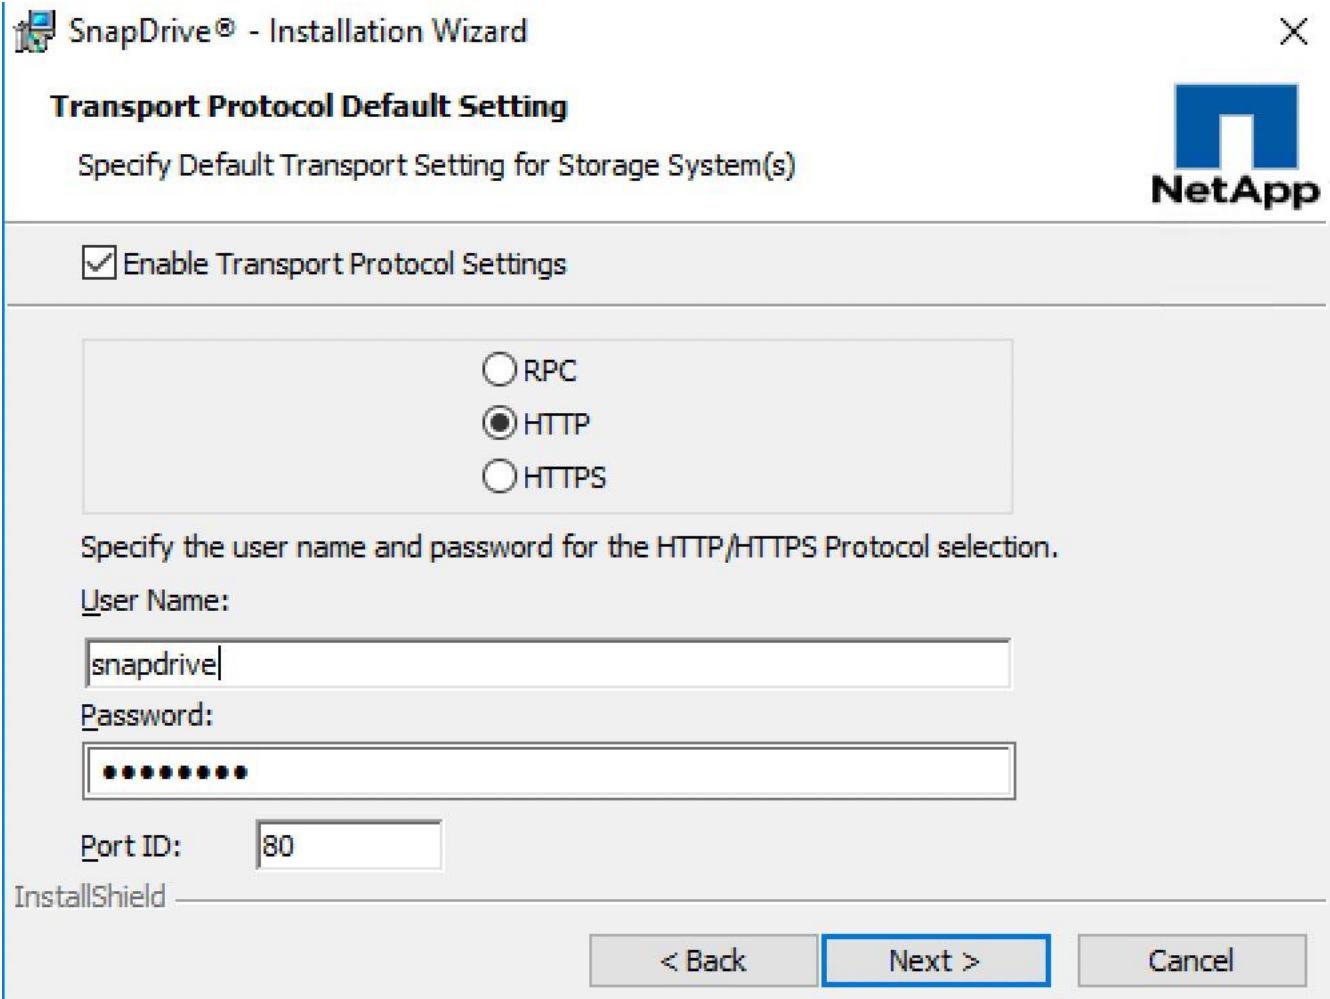
Storage System Name:

Preferred IP Address:

InstallShield

< Back Next > Cancel

8. In the Transport Protocol Default Setting screen, enable the storage protocol settings. RPC is not supported in clustered Data ONTAP.



The image shows a Windows-style dialog box titled "SnapDrive® - Installation Wizard" with a close button (X) in the top right corner. The main heading is "Transport Protocol Default Setting" with a subtitle "Specify Default Transport Setting for Storage System(s)". The NetApp logo is in the top right. A checkbox labeled "Enable Transport Protocol Settings" is checked. Below this, there are three radio button options: "RPC", "HTTP" (which is selected), and "HTTPS". A text instruction says "Specify the user name and password for the HTTP/HTTPS Protocol selection." There are three input fields: "User Name:" containing "snapdrive", "Password:" containing ten dots, and "Port ID:" containing "80". At the bottom left is the "InstallShield" logo. At the bottom right are three buttons: "< Back", "Next >" (highlighted with a blue border), and "Cancel".

SnapDrive® - Installation Wizard

Transport Protocol Default Setting
Specify Default Transport Setting for Storage System(s)

Enable Transport Protocol Settings

RPC
 HTTP
 HTTPS

Specify the user name and password for the HTTP/HTTPS Protocol selection.

User Name:
snapdrive

Password:
●●●●●●●●●●

Port ID: 80

InstallShield

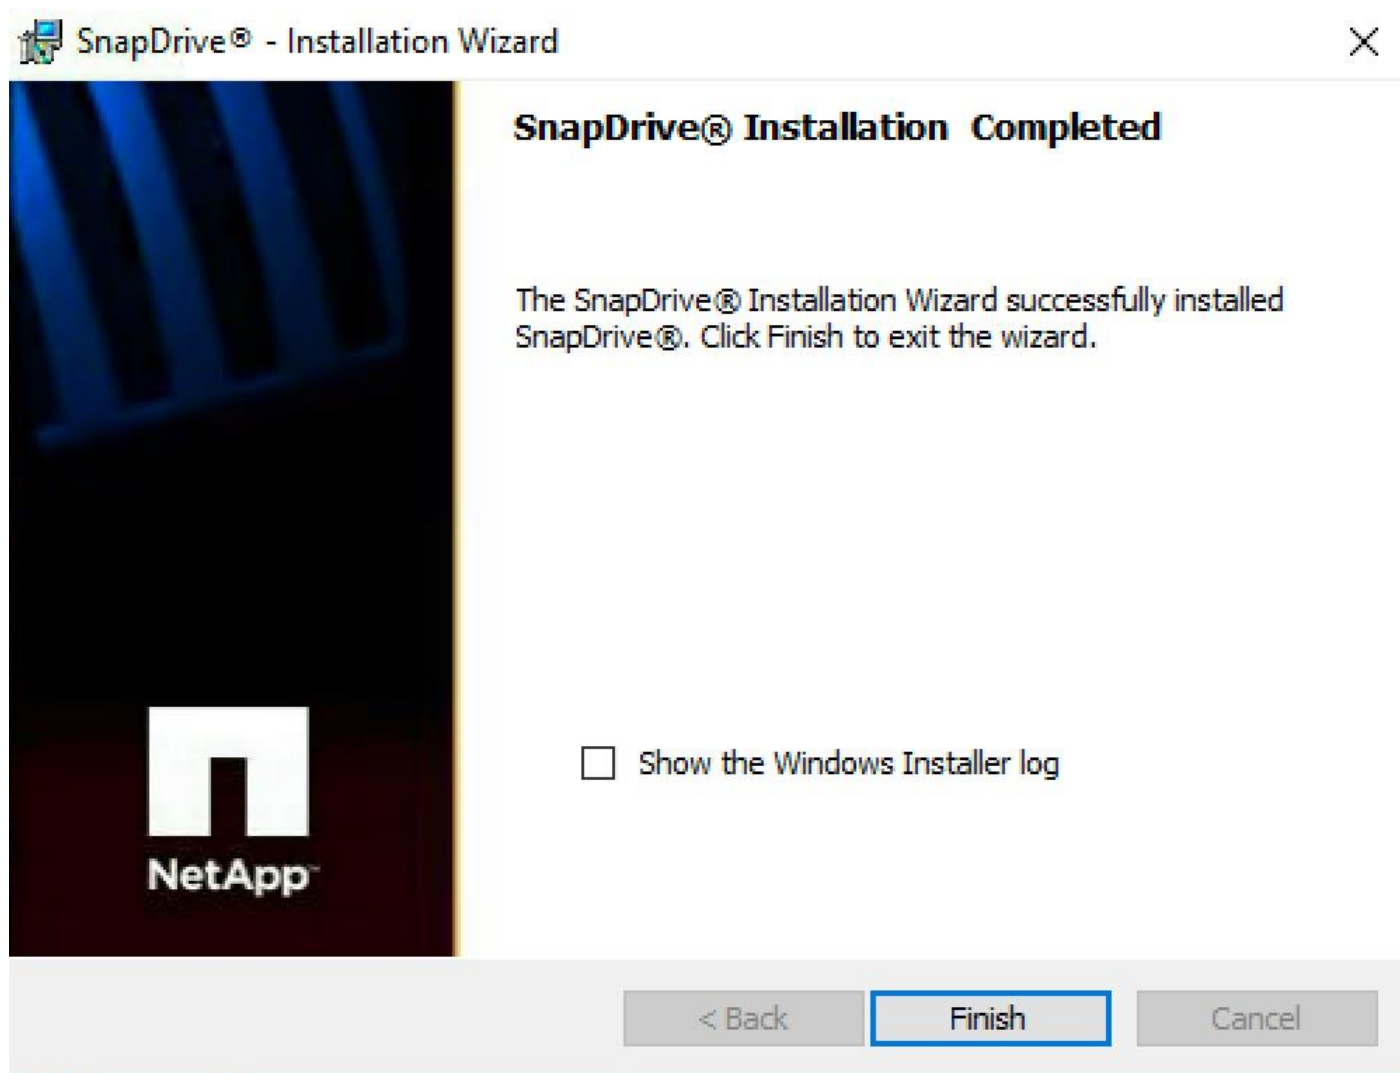
< Back Next > Cancel

9. In the Unified Manager Configuration Screen, click next.



OnCommand Unified Manager Core Package data protection capabilities are available only in 7- Mode environments.

10. In the Ready to Install screen, click Install.
11. When you have completed the Installation wizard instructions, click Finish.



NetApp SnapManager for Hyper-V

SnapManager for Hyper-V provides a solution for data protection and recovery for Microsoft® Hyper-V virtual machines (VMs) running on ONTAP® software. You can perform application-consistent and crash-consistent dataset backups according to protection policies set by your backup administrator. You can also restore VMs from these backups. Reporting features enable you to monitor the status of and get detailed information about your backup and restore jobs.

You can use the following steps to install SnapManager 2.1.2 for Hyper-V.

Downloading SnapManager for Hyper-V

Before installing SnapManager for Hyper-V, you must download the software package from the NetApp Support Site.

Before you begin you must have login credentials for the NetApp Support Site.

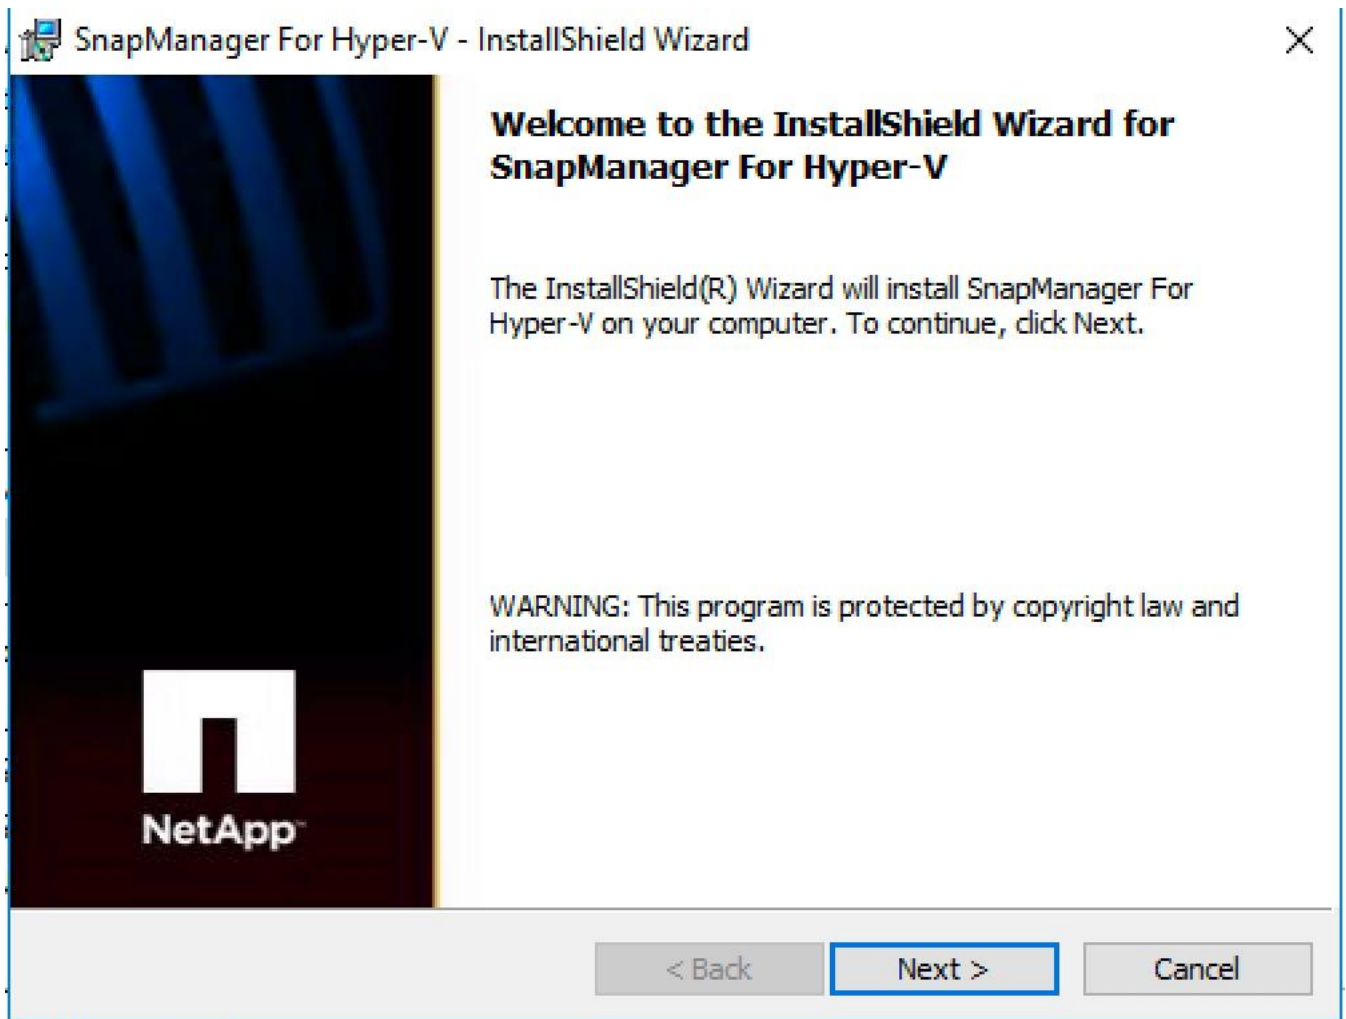
Steps

1. Log in to the NetApp Support Site.

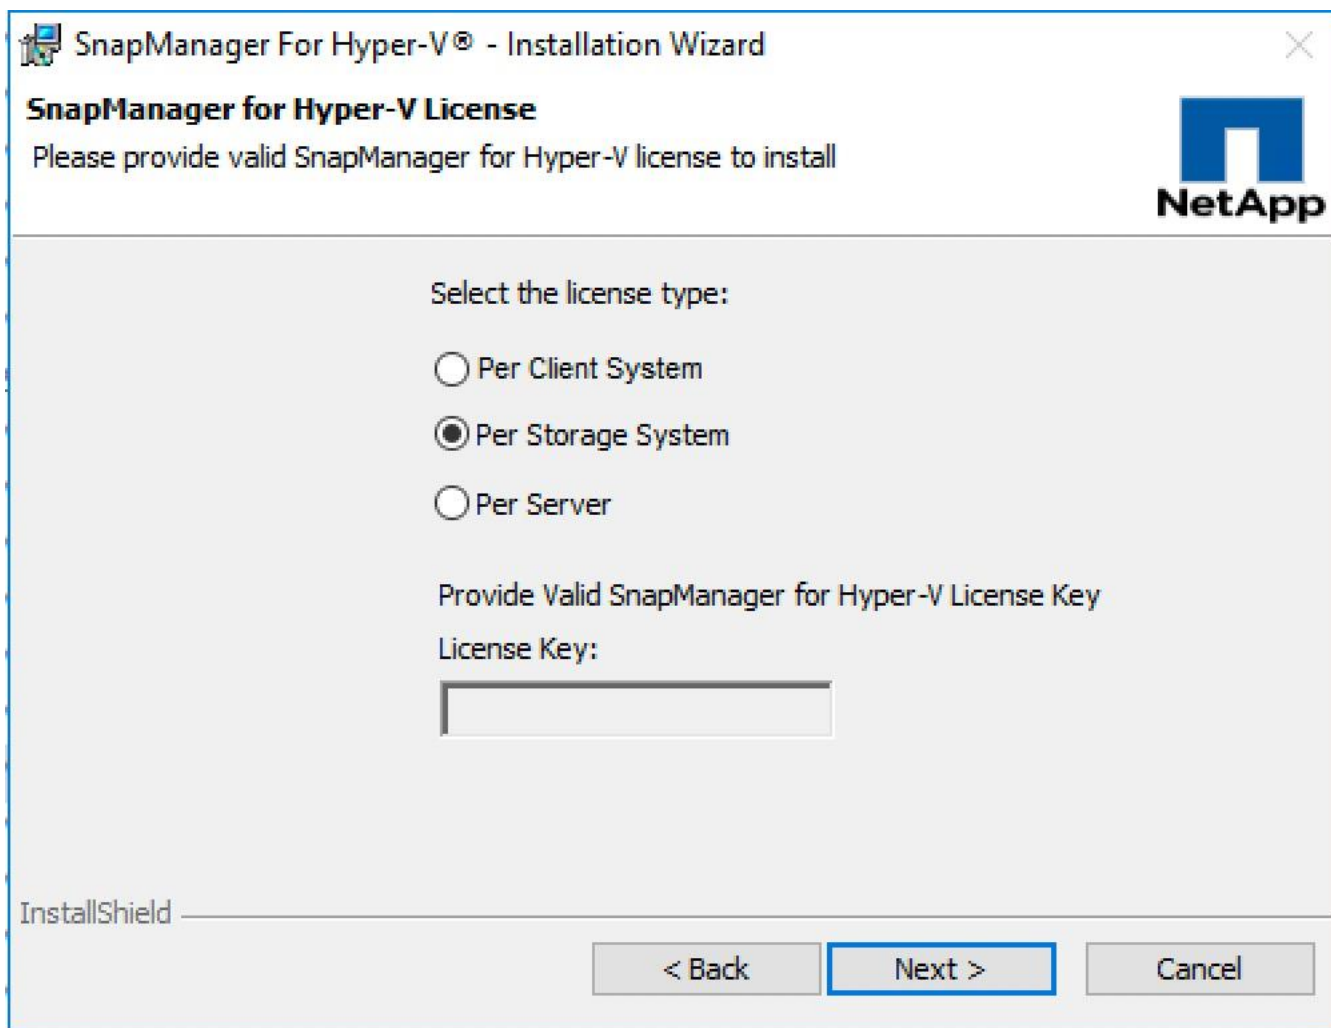
2. Go to the Download Software page.
3. From the drop-down list, select the operating system on which you are installing SnapManager for Hyper-V and click Go!.
4. Click View & Download for the software version you want to install.
5. On the Description page, click Continue.
6. Review and accept the license agreement.
7. On the Download page, click the link for the installation file.
8. Save the SnapManager for Hyper-V file to a local or network directory.
9. Click Save File.
10. Verify the checksum to ensure that the software downloaded correctly.

Installing SnapManager for Hyper-V

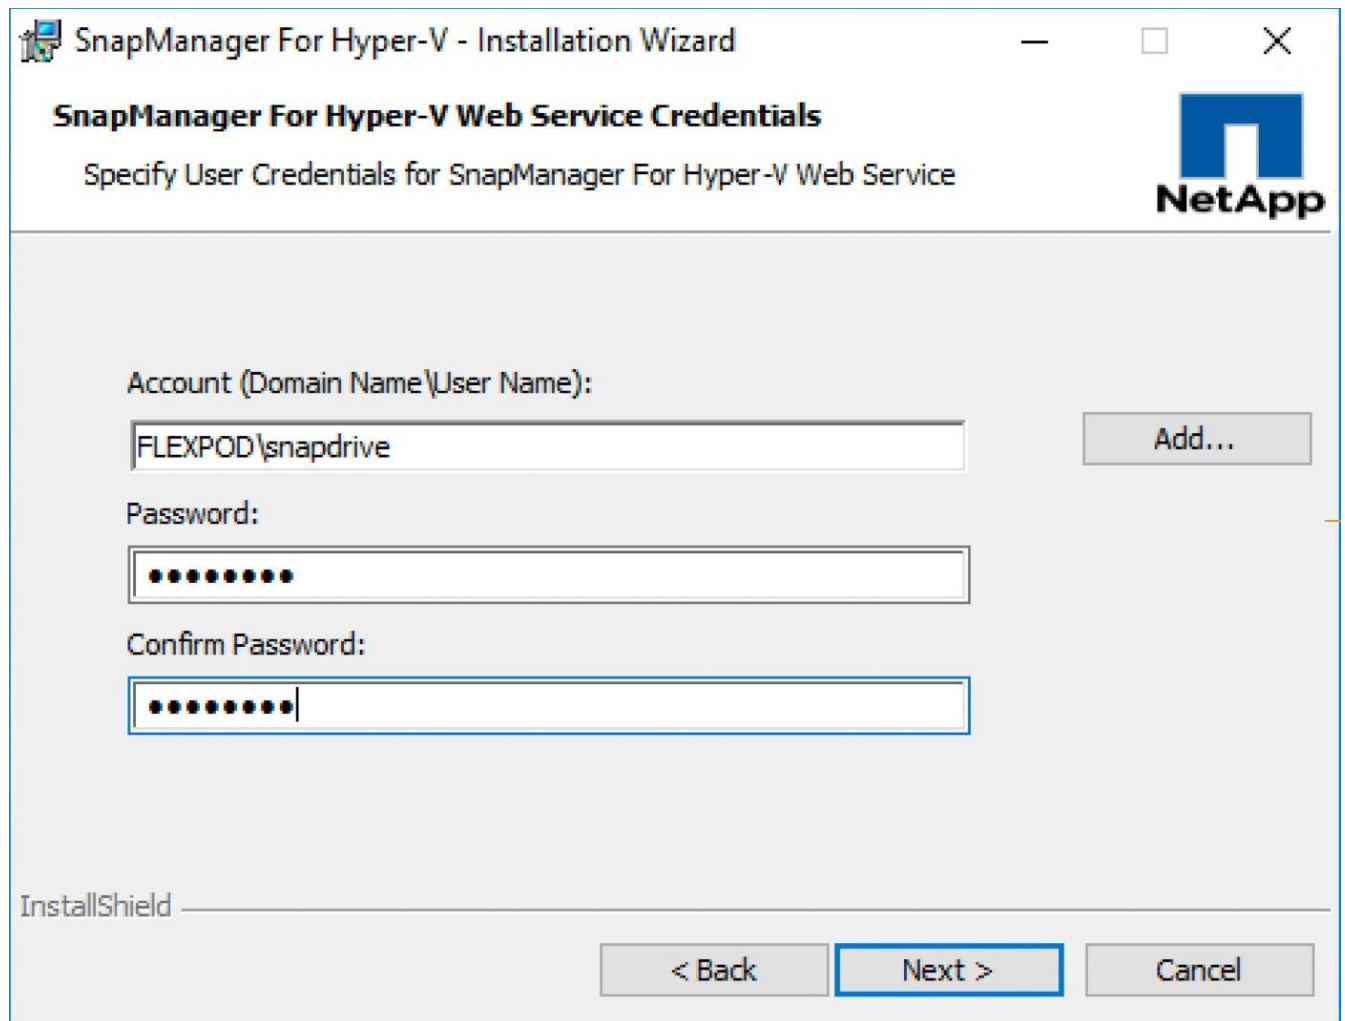
1. Launch the SnapManager for Hyper-V executable file, and then follow the Installation wizard instructions.



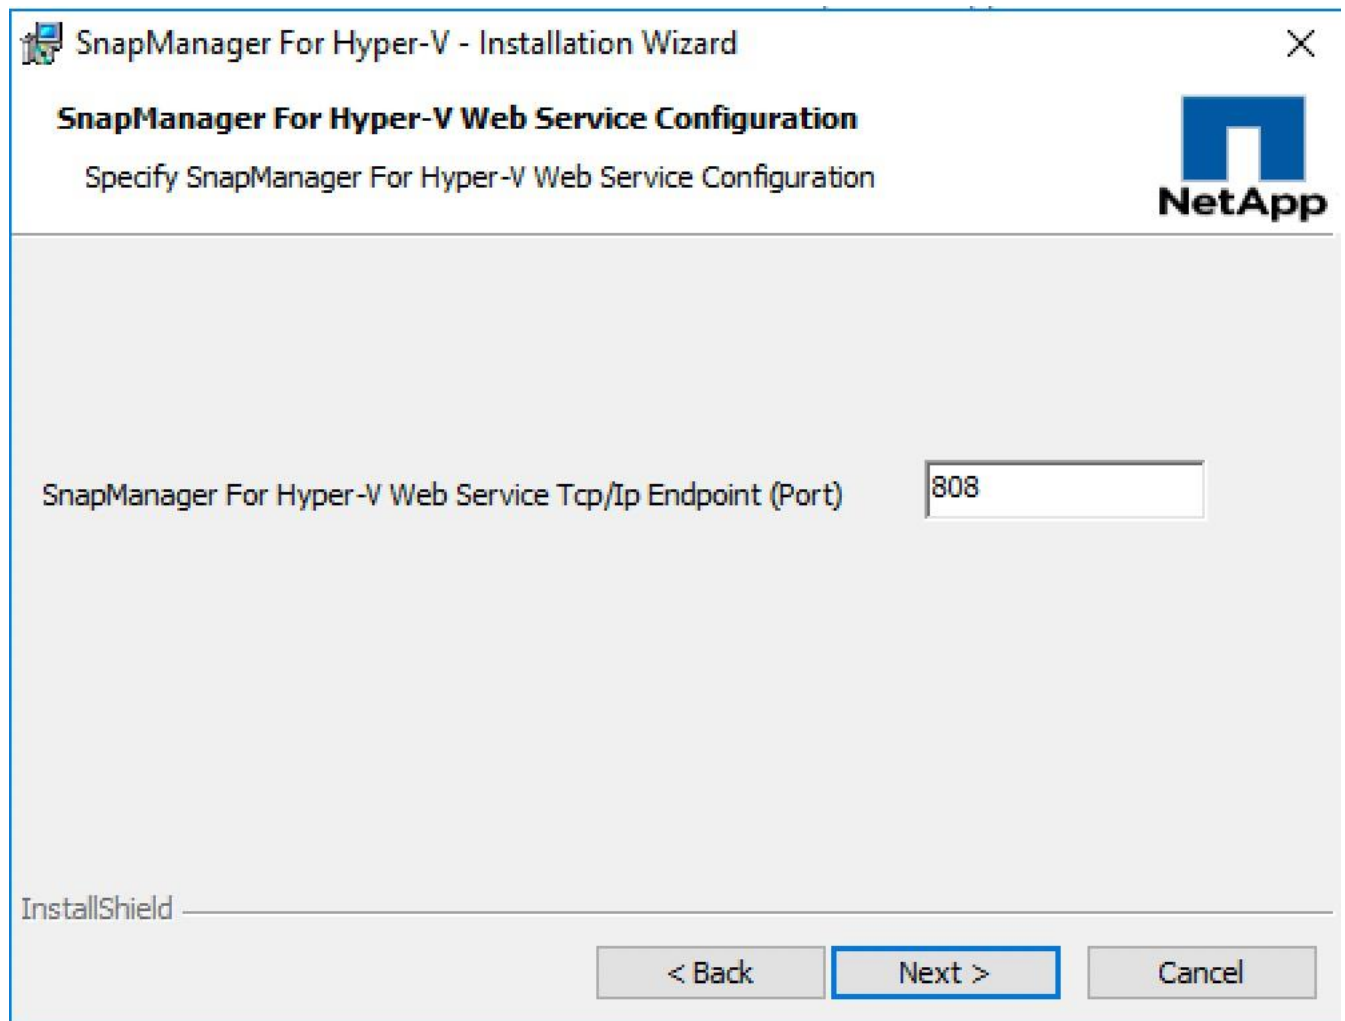
2. In the SnapManager for Hyper-V License screen, select the appropriate license type.



3. Select the installation location and click Next.
4. In the SnapManager for Hyper-V Credentials screen, enter the account and password information of an account that is a member of the local administrators.



5. In the SnapDrive Web Service Configuration screen, accept the default port numbers.



6. Click Install on the Ready to Install page.
7. Review the summary of your selections and click Finish.

Sample Tenant Provisioning

Provisioning a Sample Application Tenant

This section describes a sample procedure for provisioning an application tenant. The procedure here refers back to previous sections of this document and can be used as a guide and modified as needed when provisioning an application tenant.

1. Plan your application tenant and determine what storage protocols will be provided in the tenant. In the architecture covered in this document, fiber channel, iSCSI, FCoE and CIFS/SMB can be provided to the tenant. Also, plan what network VLANs the tenant will use.
2. In the Nexus switches, declare all added VLANs and configure the VM VLAN as an allowed VLAN on the UCS port channels and the vPC peer link. Also, Layer 3 with HSRP or VRRP can be configured in the Nexus switches to provide this VLAN access to the outside. Layer 3 setup is not covered in this document, but is covered in the Nexus 9000 documentation. Configure the storage VLANs on the UCS and storage port channels, and on the vPC peer link. The VM VLAN can also be added to the storage port channels in order to configure the tenant SVM management interface on this VLAN.
3. In the storage cluster:
 - a. Create a broadcast domain with MTU 1500 for the tenant SVM management interface. Create a broadcast domain with MTU 9000 for each tenant storage protocol except fiber channel.
 - b. Create VLAN interface ports on the node interface group on each node for tenant SVM management (VM VLAN) and for the VLAN for each storage protocol except fiber channel. Add these VLAN ports to the appropriate broadcast domains.
 - c. Create the tenant SVM and follow all procedures in that section.
 - d. Create Load-Sharing Mirrors for the tenant SVM.
 - e. Create the FC service for the tenant SVM if fiber channel is being deployed in this tenant.
 - f. Optionally, create a self-signed security certificate for the tenant SVM.
 - g. Configure CIFS for the tenant SVM.
 - h. Create a volume and CIFS share in the tenant SVM.
 - i. Create a once-a-day deduplication schedule on the volume.
 - j. If fiber channel is being deployed in this tenant, configure four FCP LIFs in the tenant SVM on the same fiber channel ports as in the Infrastructure SVM.
 - k. Create SMB LIFs in the tenant SVM on each storage node.
 - l. Create a boot LUN in the HV_boot volume in the Infra-MS-SVM for each tenant Hyper-V host.
 - m. Add the tenant SVM Administrator, SVM management LIF on the SVM management VLAN port, and default route for the SVM.
4. In the UCS one method of tenant setup is to dedicate a Hyper-V cluster and set of UCS servers to each tenant. Service profiles will be generated for at least two tenant Hyper-V hosts. These hosts can boot from LUNs from the HV_boot volume in the Infra-MS-SVM, but will also have access to FC storage in the tenant SVM.
 - a. Create a Server Pool for the tenant Hyper-V host servers.
 - b. Create all tenant VLANs in the LAN Cloud.

- c. Add the tenant VLANs to the DVS vNIC templates.

Generate service profiles from the service profile template with the vMedia policy for the tenant Hyper-V hosts. Remember to bind these service profiles to the service profile template without the vMedia policy after Windows Server 2016 installation.

5. In the Cisco MDS 9148S switches:
 - a. Create device aliases for the tenant Hyper-V host vHBAs and the FC LIFs in the tenant storage SVM.
 - b. Create zones for the tenant Hyper-V hosts with fiber channel targets from both the storage Infra-MS-SVM and the tenant SVM.
 - c. Add these zones to the Fabric zoneset and activate the zoneset.
6. In the storage cluster:
 - a. Create igroups for the tenant Hyper-V hosts in both the Infra-MS-SVM and tenant SVM. Also, create an igroup in the tenant SVM that includes the WWPNs for all tenant Hyper-V hosts to support shared storage from the tenant SVM.
 - b. In Infra-MS-SVM, map the boot LUNs created earlier to the tenant Hyper-V hosts.
7. Install and configure Windows Server 2016 on all tenant host servers. It is not necessary to map infra_datastore_1.
8. In System Center VMM, create a cluster for the tenant Hyper-V hosts. Add the hosts to the cluster.
9. Using the VMM console, add the tenant hosts to the Host groups and apply the logical switch to the hosts.
10. After applying the logical switch to the Hyper-V host from the VMM, verify, check and enable the VM adapters created on the hosts for MTU, RSS and RDMA are set correctly. Mount the tenant CIFS datastore on the tenant cluster if one was created.

Appendix - iSCSI 10/40GbE Solution

This FlexPod deployment will show configuration steps for both the Cisco UCS 6332-16UP and Cisco UCS 6248UP Fabric Interconnects (FI) in a design that will support Ethernet connectivity to the NetApp AFF through the Cisco Nexus 9000.

Configuration steps will be referenced for both fabric interconnects and will be called out by the specific model where steps have differed.

This section contains the UCS deployment for end-to-end 10/40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and C-Series rackmounts and the Cisco UCS Fabric Interconnect, between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A300. This infrastructure is deployed to provide iSCSI-booted hosts with block-level access to shared storage.

Figure 9 shows the Microsoft Hyper-V 2016 built on FlexPod components and the network connections for a configuration with the Cisco UCS 6332-16UP Fabric Interconnects. This design has end-to-end 40 Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and C-Series rackmounts and the Cisco UCS Fabric Interconnect, between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000, and between Cisco Nexus 9000 and NetApp AFF A300.

Figure 9 Nexus 9000 based FlexPod – Cisco UCS 6332 Fabric Interconnects and 40GbE End to End – iSCSI

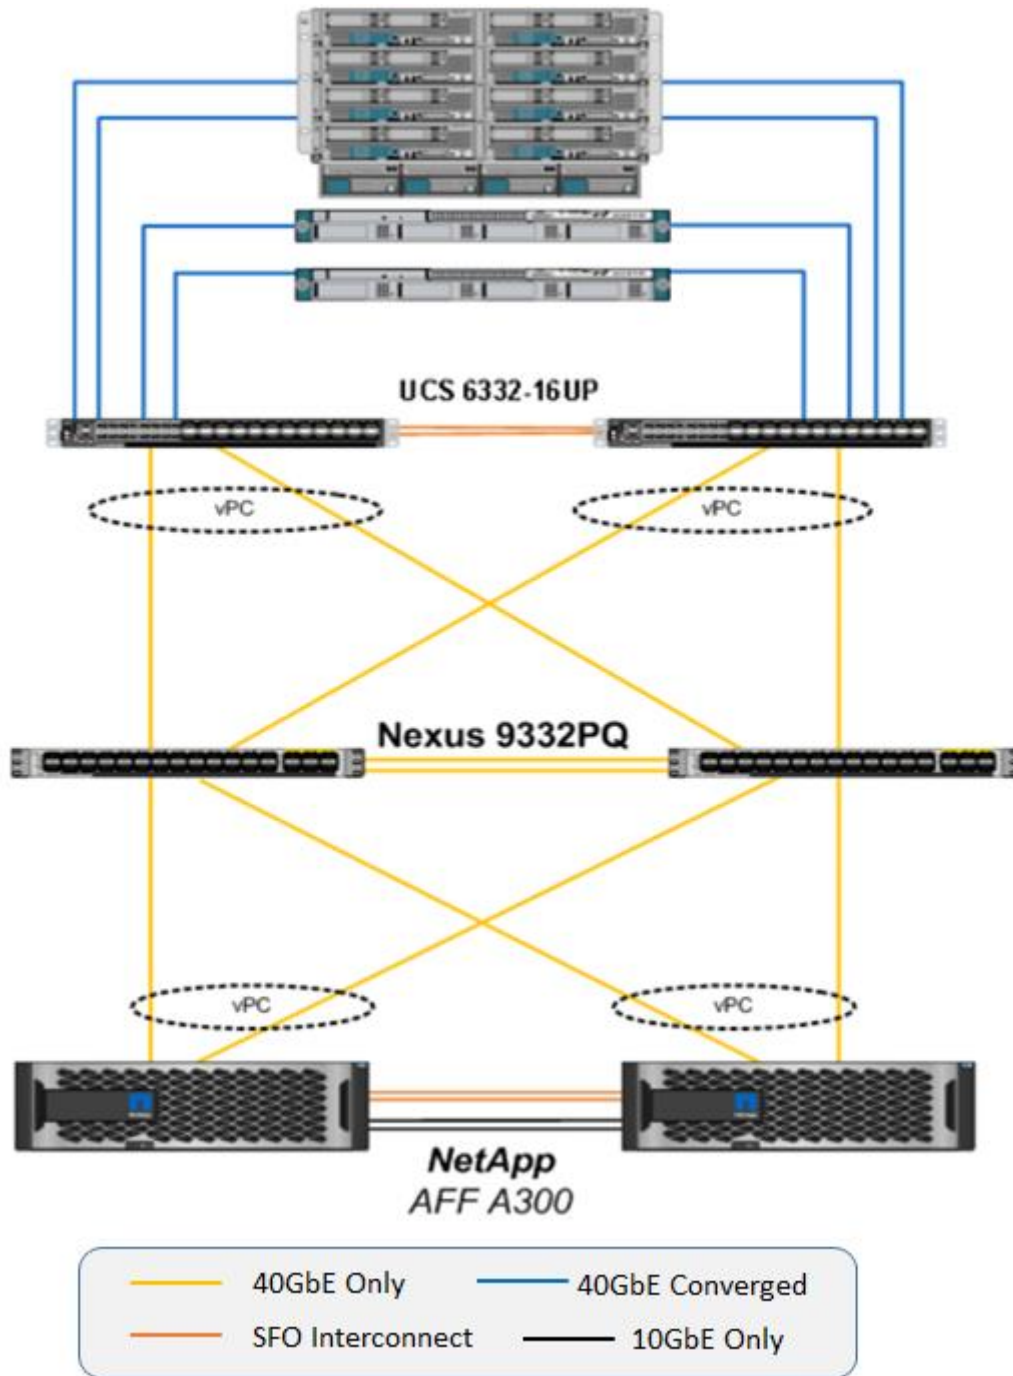



Figure 10 shows the Hyper-V built on FlexPod components and the network connections for a configuration with the Cisco UCS 6248UP Fabric Interconnects. This design is identical to the 6332-16UP based topology, but has 10 Gb Ethernet connecting through a pair of Cisco Nexus 9318oYC-EX switches to access iSCSI access to the AFF A300. Alternately, the Cisco Nexus 9332PQ switch can be used with the Cisco UCS 6248UP with a QSFP breakout cable and port configuration setting on the 9332PQ switch.

Figure 10 Nexus 9000 based FlexPod – Cisco UCS 6248UP Fabric Interconnects and 10GbE End to End – iSCSI



Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Figure 11 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect

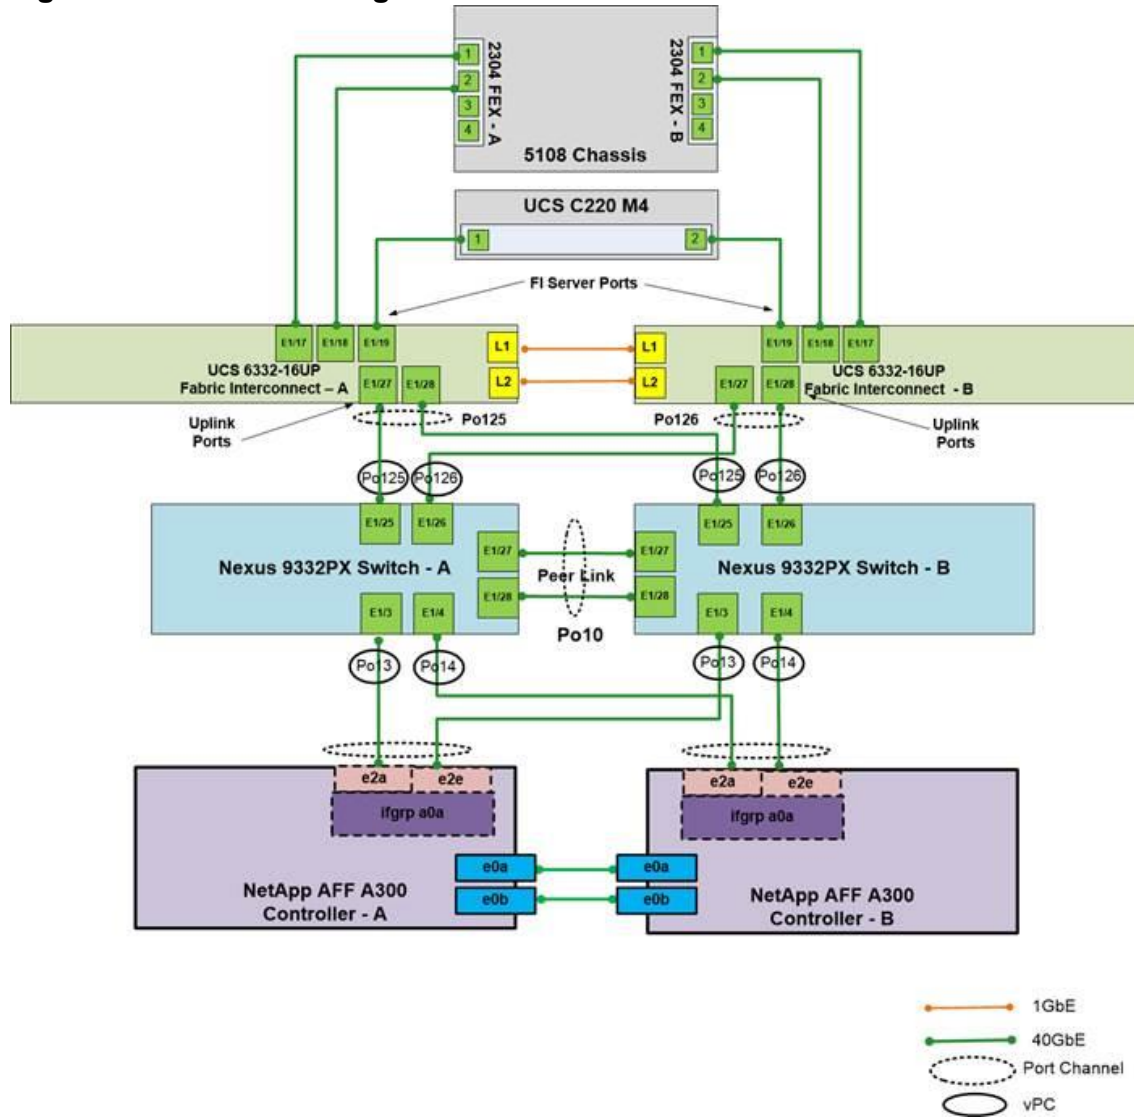
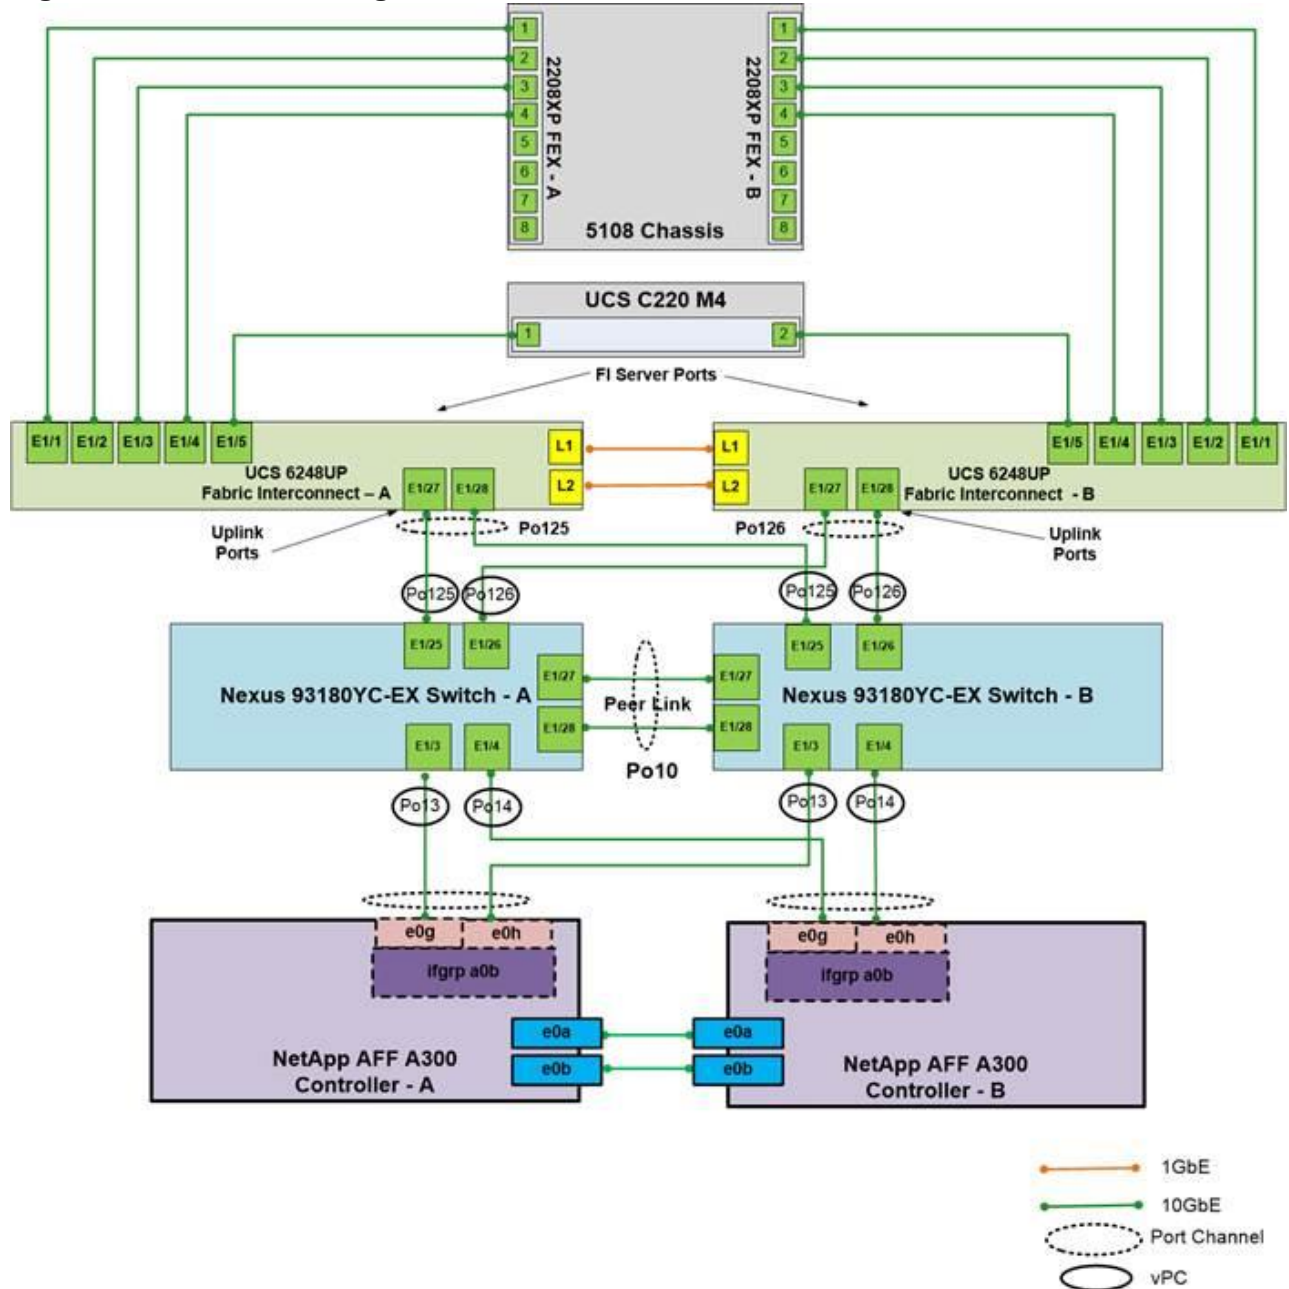


Figure 12 FlexPod Cabling with Cisco UCS 6248UP Fabric Interconnect



Create VLANs

To create the necessary virtual local area networks (VLANs), complete the following step on both switches.

From the global configuration mode, run the following commands:

```

vlan <ms-ib-mgmt-vlan-id>
name MS-IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
    
```

```
vlan <ms-lvmn-vlan-id>  
name MS-LVMN-VLAN  
vlan <ms-tenant-vm-vlan-id>  
name MS-Tenant-VM-VLAN  
vlan <ms-Cluster-vlan-id>  
name MS-Cluster-VLAN  
vlan <ms--iSCSI-A-vlan-id>  
name MS-iSCSI-A-VLAN  
vlan <ms-iSCSI-B-vlan-id>  
name MS-iSCSI-B-VLAN  
exit
```

Add NTP Distribution Interface

Cisco Nexus A

From the global configuration mode, run the following commands:

```
ntp source <switch-a-ntp-ip>  
interface Vlan<ib-mgmt-vlan-id>  
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>  
no shutdown  
exit
```

Cisco Nexus B

From the global configuration mode, run the following commands:

```
ntp source <switch-b-ntp-ip>  
interface Vlan<ib-mgmt-vlan-id>  
ip address <switch-b-ntp-ip>/<ib-mgmt-vlan-netmask-length>  
no shutdown  
exit
```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus A

To add individual port descriptions for troubleshooting activity and verification for switch A, complete the following step:



In this step and in the later sections, configure the AFF nodename <st-node> and Cisco UCS 6332-16UP or Cisco UCS 6248UP fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/3
description <st-node>-1:e2a
interface Eth1/4
description <st-node>-2:e2a
interface Eth1/25
description <ucs-clustername>-a:1/27
interface Eth1/26
description <ucs-clustername>-b:1/27
interface Eth1/27
description <nexus-hostname>-b:1/27
interface Eth1/28
description <nexus-hostname>-b:1/28
exit
```

Cisco Nexus B

To add individual port descriptions for troubleshooting activity and verification for switch B, complete the following step:

From the global configuration mode, run the following commands:

```
interface Eth1/3
description <st-node>-1:e2e
interface Eth1/4
description <st-node>-2:e2e
interface Eth1/25
description <ucs-clustername>-a:1/28
interface Eth1/26
description <ucs-clustername>-b:1/28
interface Eth1/27
description <nexus-hostname>-a:1/27
interface Eth1/28
description <nexus-hostname>-a:1/28
```

```
exit
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

To create the necessary port channels between devices, complete the following step on both switches:

From the global configuration mode, run the following commands:

```
interface Po10
description vPC peer-link
interface Eth1/27-28
channel-group 10 mode active
no shutdown
interface Po13
description <st-node>-1
interface Eth1/3
channel-group 13 mode active
no shutdown
interface Po14
description <st-node>-2
interface Eth1/4
channel-group 14 mode active
no shutdown
interface Po125
description <ucs-clustername>-a
interface Eth1/25
channel-group 125 mode active
no shutdown
interface Po126
description <ucs-clustername>-b
interface Eth1/26
channel-group 126 mode active
no shutdown
exit
```

```
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

To configure port channel parameters, complete the following step on both switches:

From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ms-ib-mgmt-vlan-id>, < ms-lvmnvlan-id>, <ms-tenant-vm-vlan-id>, <ms-iSCSI-A-
vlan-id>, <ms-iSCSI-B-vlan-id>, < ms-cluster-vlan-id>
spanning-tree port type network
interface Po13
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan < ms-iSCSI-A-vlan-id >, < ms-iSCSI-B-vlan-id >
spanning-tree port type edge trunk
mtu 9216
interface Po14
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ms-iSCSI-A-vlan-id>, <ms-iSCSI-B-vlan-id>
spanning-tree port type edge trunk
mtu 9216
interface Po125
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ms-ib-mgmt-vlan-id>, < ms-lvmn-vlan-id>, <ms-tenant-vm-vlan-id>, <ms-iSCSI-A-
vlan-id>, <ms-iSCSI-B-vlan-id>
spanning-tree port type edge trunk
mtu 9216
interface Po126
```

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ms-ib-mgmt-vlan-id>, <ms-lvmn-vlan-id>, <ms-tenant-vm-vlan-id>, <ms-iSCSI-A-
vlan-id>, <ms-iSCSI-B-vlan-id>,<ms-cluster-vlan-id>
spanning-tree port type edge trunk
mtu 9216
exit
copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, complete the following step:

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 10
peer-keepalive destination <nexus-B-mgmt-ip> source <nexus-A-mgmt-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po13
vpc 13
interface Po14
vpc 14
interface Po125
vpc 125
interface Po126
vpc 126
exit
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, complete the following step:

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt-ip> source <nexus-B-mgmt-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po13
vpc 13
interface Po14
vpc 14
interface Po125
vpc 125
interface Po126
vpc 126
exit
copy run start
```

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, we recommend using vPCs to uplink the Cisco Nexus switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is completed.

Storage Configuration

Create Jumbo Frame MTU Broadcast Domains in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands to create a broadcast domain for SMB and iSCSI on ONTAP:

```
broadcast-domain create -broadcast-domain Infra_MS_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_MS_iSCSI-B -mtu 9000
```

Create VLANs

To create iSCSI VLAN, create iSCSI VLAN ports and add them to the iSCSI broadcast domain:

```
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-A-vlan-id>
network port vlan create -node <st-node01> -vlan-name a0a-<infra-iscsi-B-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-A-vlan-id>
network port vlan create -node <st-node02> -vlan-name a0a-<infra-iscsi-B-vlan-id>

broadcast-domain add-ports -broadcast-domain Infra_MS_iSCSI-A -ports <st-node01>:a0a-<infra-iscsi-A-vlan-id>,
<st-node02>:a0a-<infra-iscsi-A-vlan-id>
broadcast-domain add-ports -broadcast-domain Infra_MS_iSCSI-B -ports <st-node01>:a0a-<infra-iscsi-B-vlan-id>,
<st-node02>:a0a-<infra-iscsi-B-vlan-id>
```

Run the following command to create the iSCSI service on each SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM.

```
iscsi create -vserver Infra-MS-SVM
iscsi show
```

Create iSCSI LIFs

Run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver Infra-MS-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -home-node
<st-node01> -home-port a0a-<infra-iscsi-a-vlan-id> -address <var_node01_iscsi_lif01a_ip> -netmask
<var_node01_iscsi_lif01a_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false

network interface create -vserver Infra-MS-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -home-node
<st-node01> -home-port a0a-<infra-iscsi-b-vlan-id> -address <var_node01_iscsi_lif01b_ip> -netmask
<var_node01_iscsi_lif01b_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false

network interface create -vserver Infra-MS-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -home-node
<st-node02> -home-port a0a-<infra-iscsi-a-vlan-id> -address <var_node02_iscsi_lif02a_ip> -netmask
<var_node02_iscsi_lif02a_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false

network interface create -vserver Infra-MS-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-node
<st-node02> -home-port a0a-<infra-iscsi-b-vlan-id> -address <var_node02_iscsi_lif02b_ip> -netmask
<var_node02_iscsi_lif02b_mask> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false

network interface show
```

Server Configuration

Perform Initial Setup of Cisco UCS 6332-16UP Fabric Interconnects for FlexPod Environments

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

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.



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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 125 as the unique ID of the port channel.
6. Enter `vPC-125-Nexus` as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 126 as the unique ID of the port channel.
16. Enter `vPC-126-Nexus` as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.
22. Create MAC Address Pools
23. To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:
24. 1 In Cisco UCS Manager, click LAN on the left.
25. Select Pools > root.



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

26. Right-click MAC Pools under the root organization.
27. Select Create MAC Pool to create the MAC address pool.
28. Enter `MAC-POOL-A` as the name of the MAC pool.
29. Optional: Enter a description for the MAC pool.
30. Select **Sequential** as the option for Assignment Order.
31. Click Next.
32. Click Add.
33. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the same and also have embedded the UCS domain number information to give us `00:25:B5:52:0A:00` as our first MAC address.

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

Create a Block of MAC Addresses

First MAC Address : Size :

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

35. Click OK.
36. Click Finish.
37. In the confirmation message, click OK.
38. Right-click MAC Pools under the root organization.
39. Select Create MAC Pool to create the MAC address pool.

40. Enter `MAC-POOL-B` as the name of the MAC pool.
41. Optional: Enter a description for the MAC pool.
42. Select **Sequential** as the option for Assignment Order.
43. Click Next.
44. Click Add.
45. Specify a starting MAC address.



For the FlexPod solution, it is recommended to place `0B` in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example and have embedded the UCS domain number information to give us `00:25:B5:52:0B:00` as our first MAC address.

46. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
47. Click OK.
48. Click Finish.
49. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

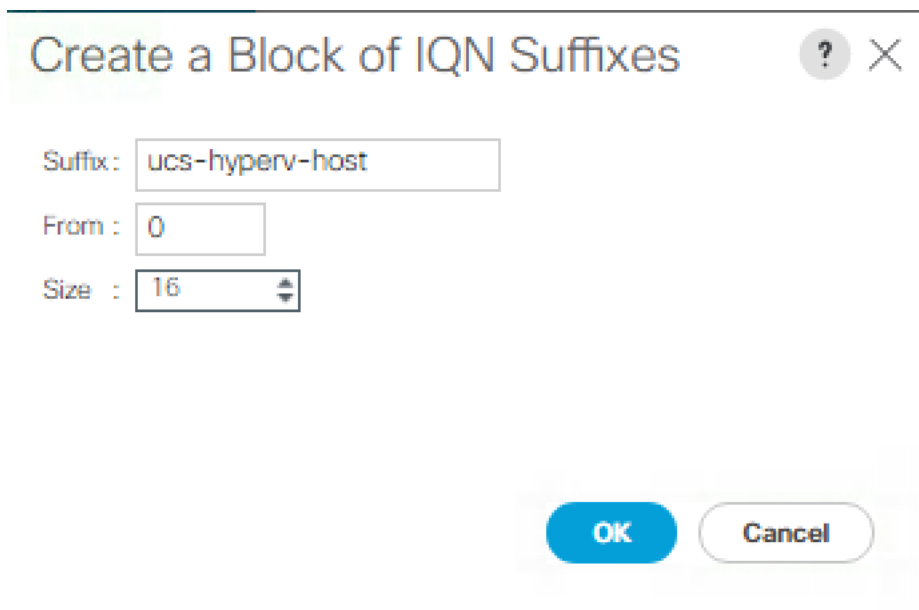
To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. Right click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter `IQN-Pool` for the name of the IQN pool
6. Optional: Enter a description for the IQN pool
7. Enter `iqn.1992-08.com.cisco` as the prefix.
8. Select **Sequential** for Assignment Order
9. Click Next.
10. Click Add.
11. Enter `ucs-host` as the suffix.



If multiple Cisco UCS domains are being used, a more specific IQN suffix may need to be used.

12. Enter 1 in the From field.
13. Specify the size of the IQN block sufficient to support the available server resources.
14. Click OK.



15. Click Finish.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.
3. Right-click IP Pools.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select **Sequential** for the assignment order.
8. Click Next.
9. Click Add to add a block of IP address.
10. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
11. Set the size to enough addresses to accommodate the servers.
12. Click OK.

13. Click Next.
14. Click Finish.
15. Right-click IP Pools.
16. Select Create IP Pool.
17. Enter iSCSI-IP-Pool-B as the name of IP pool.
18. Optional: Enter a description for the IP pool.
19. Select **Sequential** for the assignment order.
20. Click Next.
21. Click Add to add a block of IP address.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers.
24. Click OK.
25. Click Next.
26. Click Finish.

Create UUID Suffix Pool

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-POOL` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select **Sequential** for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.

12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `MS-Server-Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the Hyper-V management cluster and click >> to add them to the `MS-Server-Pool` server pool.
9. Click Finish.
10. Click OK.

Create VLANs

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

1. In Cisco UCS Manager, click LAN on the left.



In this procedure, five unique VLANs are created. See Table 1.

2. Select LAN > LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter `Native-VLAN` as the name of the VLAN to be used as the native VLAN.

6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the native VLAN ID.
8. Keep the Sharing Type as None.
9. Click OK, and then click OK again.

Create VLANs ? X

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

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs
14. Enter `MS-IB-MGMT` as the name of the VLAN to be used for management traffic.
15. Keep the Common/Global option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.

17. Keep the Sharing Type as None.
18. Click OK, and then click OK again.
19. Right-click VLANs.
20. Select Create VLANs.
21. Enter `MS-iSCSI-A-VLAN` as the name of the VLAN to be used for iSCSI-A.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the iSCSI-A VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter `MS-iSCSI-B-VLAN` as the name of the VLAN to be used for iSCSI-B
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the iSCSI-B VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK, and then click OK again
33. Right-click VLANs.
34. Select Create VLANs.
35. Enter `MS-LVMN` as the name of the VLAN to be used for Live Migration.
36. Keep the Common/Global option selected for the scope of the VLAN.
37. Enter the Live Migration VLAN ID.
38. Keep the Sharing Type as None.
39. Click OK, and then click OK again.
40. Select Create VLANs.
41. Enter `MS-Cluster` as the name of the VLAN to be used for Cluster Communication VLAN.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the Cluster network VLAN ID.

44. Keep the Sharing Type as None.
45. Click OK, and then click OK again.
46. Select Create VLANs.
47. Enter MS-Tenant-VM as the name of the VLAN to be used for VM Traffic.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

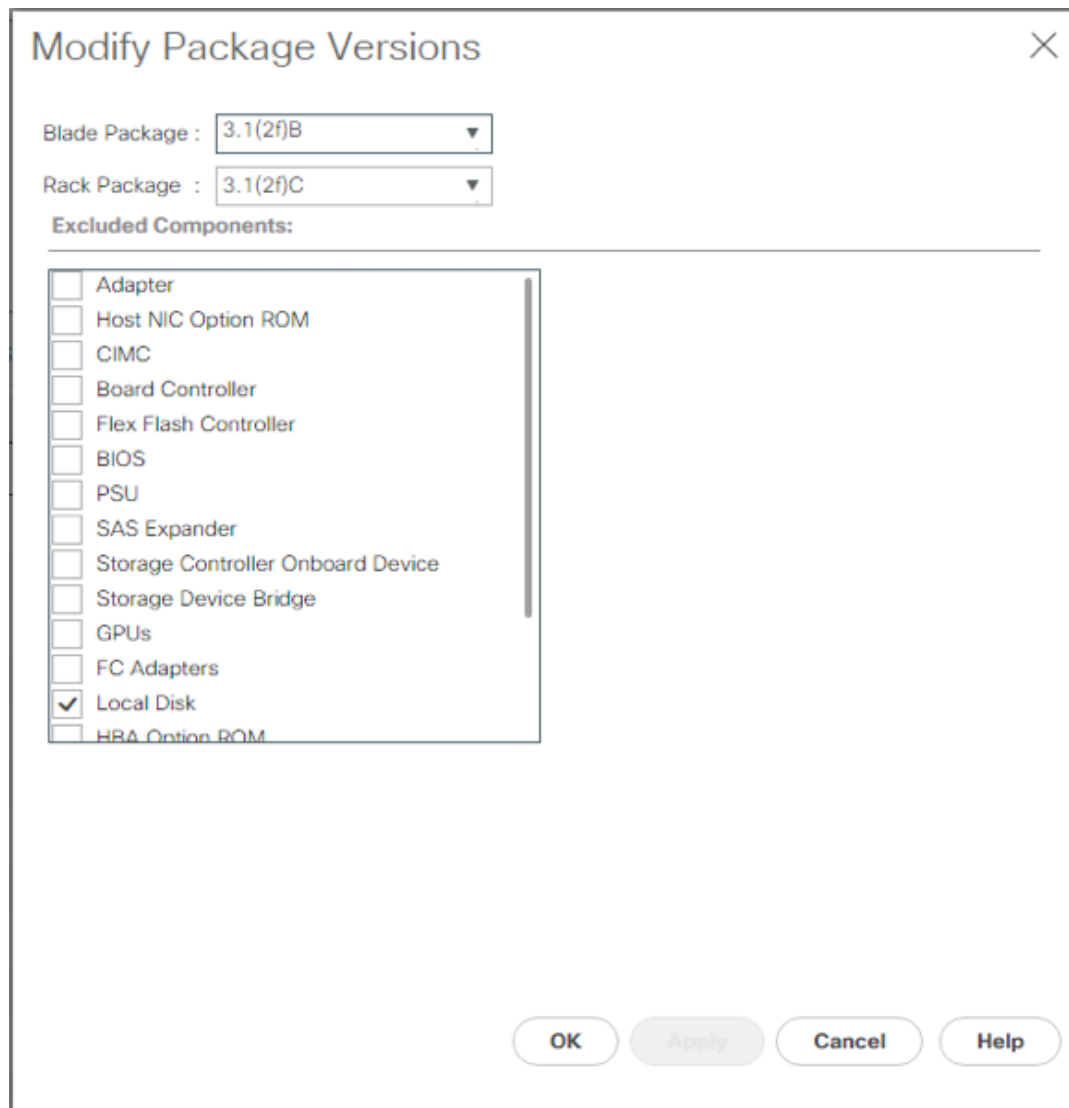
Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Na...	Multicast Policy N...
VLAN default (1)	1	Lan	Ether	No	None		
VLAN Native-VLAN (2)	2	Lan	Ether	Yes	None		
VLAN MS-IB-MGMT (90...	904	Lan	Ether	No	None		
VLAN MS-CSV (905)	905	Lan	Ether	No	None		
VLAN MS-LVMN (906)	906	Lan	Ether	No	None		
VLAN MS-Cluster (907)	907	Lan	Ether	No	None		
VLAN MS-Tenant-VM (...)	908	Lan	Ether	No	None		
VLAN MS-ISCSI-A (301...	3012	Lan	Ether	No	None		
VLAN MS-ISCSI-B (302...	3022	Lan	Ether	No	None		

Modify Default 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 properties.

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(2f) for both the Blade and Rack Packages.



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

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
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.
6. Click OK

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

Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.



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

To create a local disk configuration policy, complete the following steps:

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

Create Local Disk Configuration Policy ? X

Name : SAN-Boot

Description :

Mode : No Local Storage ▼

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

OK Cancel

8. Click OK.

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

To create a network control policy that enables CDP and LLDP on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable-CDP-LLDP` as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

Create Network Control Policy ? X

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

9. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter `No-Power-Cap` as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Create Power Control Policy ? X

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.

Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for Cisco UCS B-Series and Cisco UCS C-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-Broadwell.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter UCS-CPU-E52660E as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.

Create CPU/Cores Qualifications

Processor Architecture : PID (RegEx) :

Min Number of Cores : Unspecified select Max Number of Cores : Unspecified select

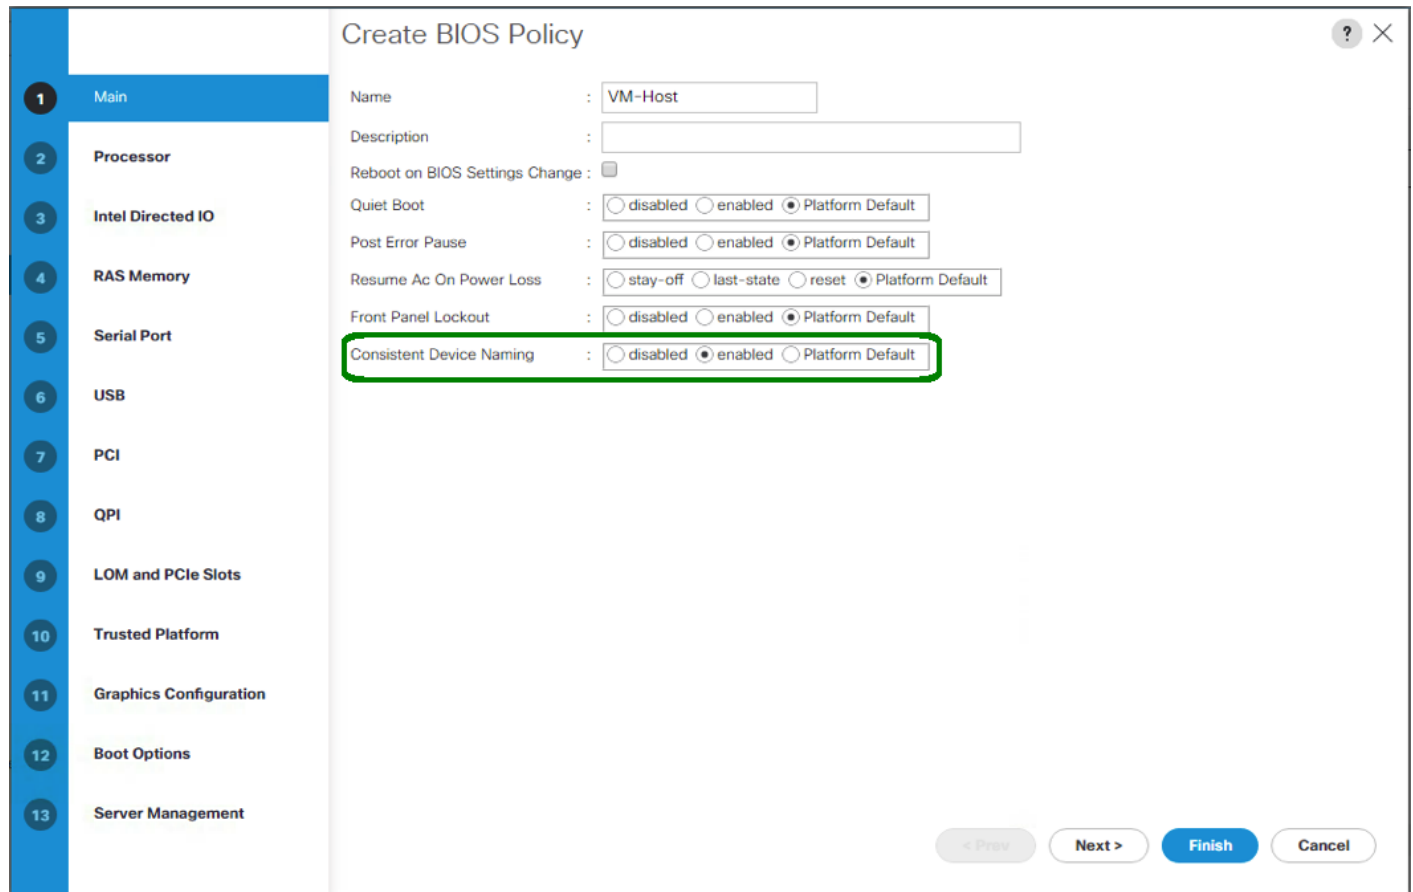
Min Number of Threads : Unspecified select Max Number of Threads : Unspecified select

CPU Speed (MHz) : Unspecified select CPU Stepping : Unspecified select

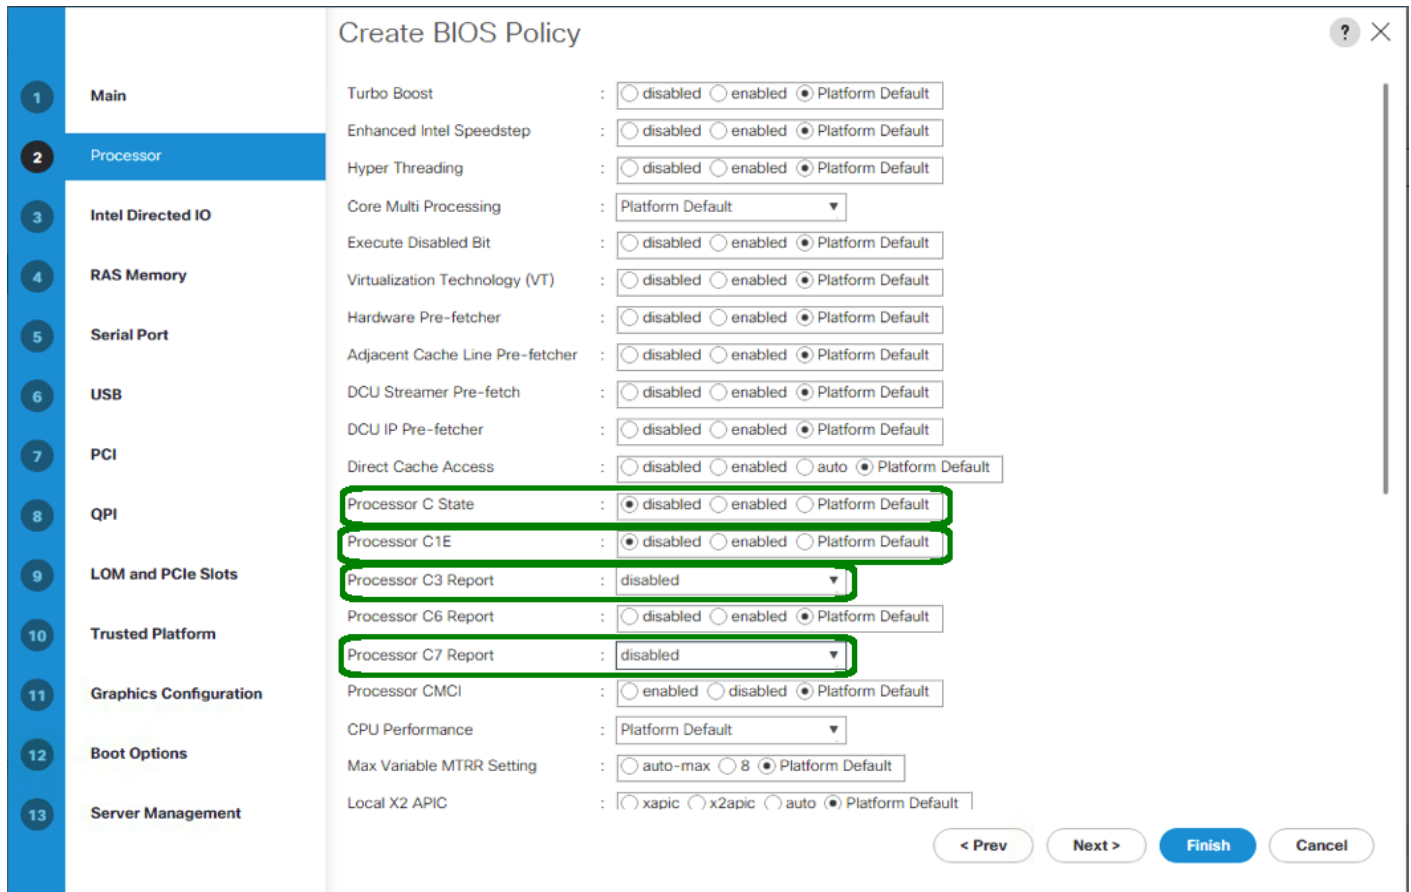
Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `MS-Host` as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



8. Click on the Processor tab on the left.
9. Set the following within the Processor tab
10. Processor C State -> disabled
11. Processor C1E -> disabled
12. Processor C3 Report -> disabled
13. Processor C7 Report -> disabled




14. Scroll down to the remaining Processor options, and select:

15. Energy Performance -> performance

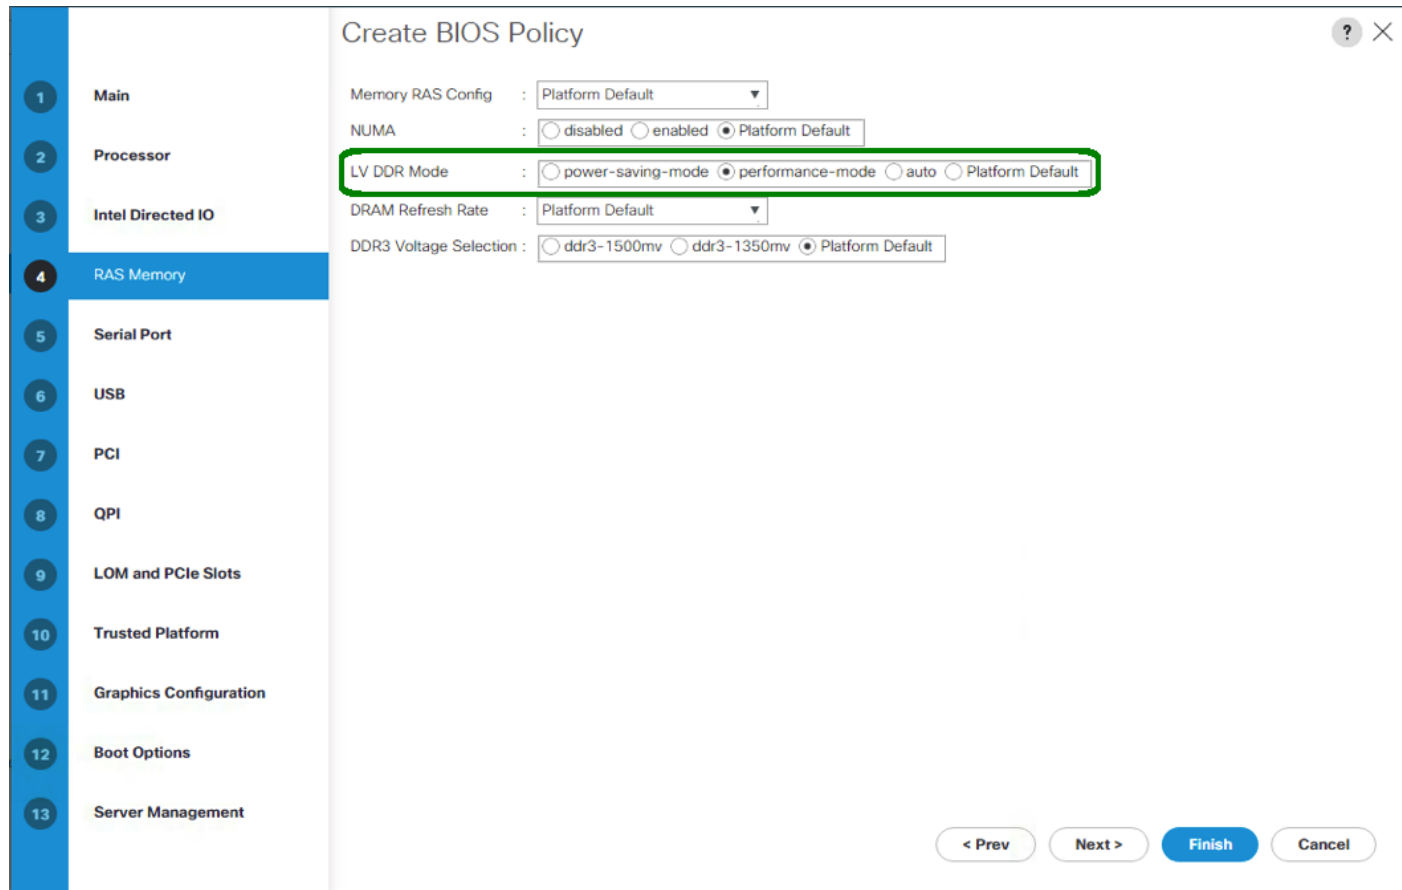
16. Frequency Floor Override -> enabled

17. DRAM Clock Throttling -> performance



18. Click on the RAS Memory option, and select:

19. LV DDR Mode -> performance-mode



20. Click Finish to create the BIOS policy.

21. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select "On Next Boot" to delegate maintenance windows to server administrators.

Servers / Policies / root / Maintenance Poli... / default

General	Events
Actions Delete Show Policy Usage Use Global	Properties Name : default Description : <input type="text"/> Owner : Local Soft Shutdown Timer : <input type="text" value="150 Secs"/> Reboot Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic <input checked="" type="checkbox"/> On Next Boot (Apply pending changes at next reboot.)

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

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 4 vNIC Templates will be created.

Create Infrastructure vNICs

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter vNIC-Template-A as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for MS-IB-MGMT, MS-Cluster, MS-CSV, and MS-Tenant-VM VLANs.
13. Set Native-VLAN as the native VLAN.

14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-POOL-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



Name : vNIC-Template-A

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

Advanced Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	MS-Tenant-VM	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-B	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-A	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-IB-MGMT	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-CSV	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Cluster	<input type="radio"/>

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC-POOL-A(42/48)

QoS Policy : <not set>

Network Control Policy : Enable-CDP-LLDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : <not set>

OK

Cancel

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

Create Secondary Redundancy Template Infra-B

1. Select LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter vNIC-Template-B as the vNIC template name.
6. Select Fabric B.
7. Select the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select Infra-A for the Peer Redundancy Template.
10. In the MAC Pool list, select MAC-POOL-B. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Click OK to create the vNIC template.
12. Click OK.
13. Create iSCSI vNICs
14. Select LAN on the left.
15. Select Policies > root.
16. Right-click vNIC Templates.
17. Select Create vNIC Template.
18. Enter iSCSI-Template-A as the vNIC template name.
19. Select Fabric A. Do not select the Enable Failover checkbox.
20. Leave Redundancy Type set at No Redundancy.
21. Under Target, make sure that only the Adapter checkbox is selected.
22. Select Updating Template for Template Type.
23. Under VLANs, select only MS-iSCSI-A-VLAN.
24. Select MS-iSCSI-A-VLAN as the native VLAN.

25. Leave vNIC Name set for the CDN Source.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select MAC-Pool-A.
28. From the Network Control Policy list, select Enable-CDP-LLDP.
29. Click OK to complete creating the vNIC template.
30. Click OK.
31. Select LAN on the left.
32. Select Policies > root.
33. Right-click vNIC Templates.
34. Select Create vNIC Template.
35. Enter iSCSI-Template-B as the vNIC template name.
36. Select Fabric B. Do not select the Enable Failover checkbox.
37. Leave Redundancy Type set at No Redundancy.
38. Under Target, make sure that only the Adapter checkbox is selected.
39. Select Updating Template for Template Type.
40. Under VLANs, select only MS-iSCSI-B-VLAN.
41. Select MS-iSCSI-B-VLAN as the native VLAN.
42. Leave vNIC Name set for the CDN Source.
43. Under MTU, enter 9000.
44. From the MAC Pool list, select MAC-Pool-B.
45. From the Network Control Policy list, select Enable-CDP-LLDP.
46. Click OK to complete creating the vNIC template.
47. Click OK.

Create LAN Connectivity Policy for iSCSI Boot

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following 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.
5. Enter iSCSI-BOOT 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-Infra-A as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select vNIC-Template-A.
10. In the Adapter Policy list, select Windows.
11. Click OK to add this vNIC to the policy.

The screenshot shows a 'Create vNIC' dialog box with the following configuration:

- Name: 00-Infra-A
- Use vNIC Template:
- Redundancy Pair:
- vNIC Template: vNIC-Template-A
- Adapter Policy: Windows

Buttons: OK, Cancel

12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter 01-Infra-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.

15. In the vNIC Template list, select vNIC-Template-B.
16. In the Adapter Policy list, select Windows.
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-iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select iSCSI-Template-A.
22. In the Adapter Policy list, select Windows.
23. Click OK to add this vNIC to the policy.
24. Click the upper Add button to add a vNIC to the policy.
25. In the Create vNIC dialog box, enter 03-iSCSI-B as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select iSCSI-Template-B.
28. In the Adapter Policy list, select Windows.
29. Click OK to add this vNIC to the policy.
30. Expand the Add iSCSI vNICs.
31. Select Add in the Add iSCSI vNICs section.
32. Set the name to iSCSI-A-vNIC.
33. Select the 02-iSCSI-A as Overlay vNIC.
34. Set the VLAN to iSCSI-A-VLAN (native).
35. Set the iSCSI Adapter Policy to default
36. Leave the MAC Address set to None.
37. Click OK.
38. Select Add in the Add iSCSI vNICs section.
39. Set the name to iSCSI-B-vNIC.
40. Select the 03-iSCSI-B as Overlay vNIC.
41. Set the VLAN to iSCSI-B-VLAN.

42. Set the VLAN to iSCSI-B-VLAN (native).
43. Set the iSCSI Adapter Policy to default.
44. Leave the MAC Address set to None.

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 03-iSCSI-B	Derived	
vNIC 02-iSCSI-A	Derived	
vNIC 01-Infra-B	Derived	
vNIC 00-Infra-A	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-B-vNIC	03-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-A-vNIC	02-iSCSI-A	default	Derived

Add Delete Modify

OK Cancel

45. Click OK, then click OK again to create the LAN Connectivity Policy.

Create iSCSI Boot Policy

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi_lifo1a and iscsi_lifo1b) and two iSCSI LIFs are on cluster node 2 (iscsi_lifo2a and iscsi_lifo2b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).



One boot policy is configured in this procedure. The policy configures the primary target to be iscsi_lifo1a.

To create a boot policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.

3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter iSCSI-Boot-Fab-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.



Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select Local CD/DVD.
9. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
10. In the Add iSCSI Boot dialog box, enter iSCSI-A-vNIC.
11. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter iSCSI-B-vNIC.
14. 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
- vNICs
- vHBAs
- iSCSI vNICs
- Add iSCSI Boot
- CIMC Mounted vMedia
- EFI Shell

Boot Order									
Name	Or...	vNIC/vHBA/iS...	Type	WWN	LUN ...	Slot N..	Boot ...	Boot ...	Descr...
Local CD/DVD	1								
iSCSI	2								
iSCSI		iSCSI-A-vNIC	Primary						
iSCSI		iSCSI-B-vNIC	Secondary						

[See iSCSI Boot Parameters](#)

15. Click OK to create the policy.

Create Service Profile Templates

In this procedure, one service profile template for Infrastructure Hyper-V hosts is created for Fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
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 `Hyper-V-Host-iSCSI` 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/org-FP-BEARS-MS**

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.

< Previous Next > **Finish** Cancel

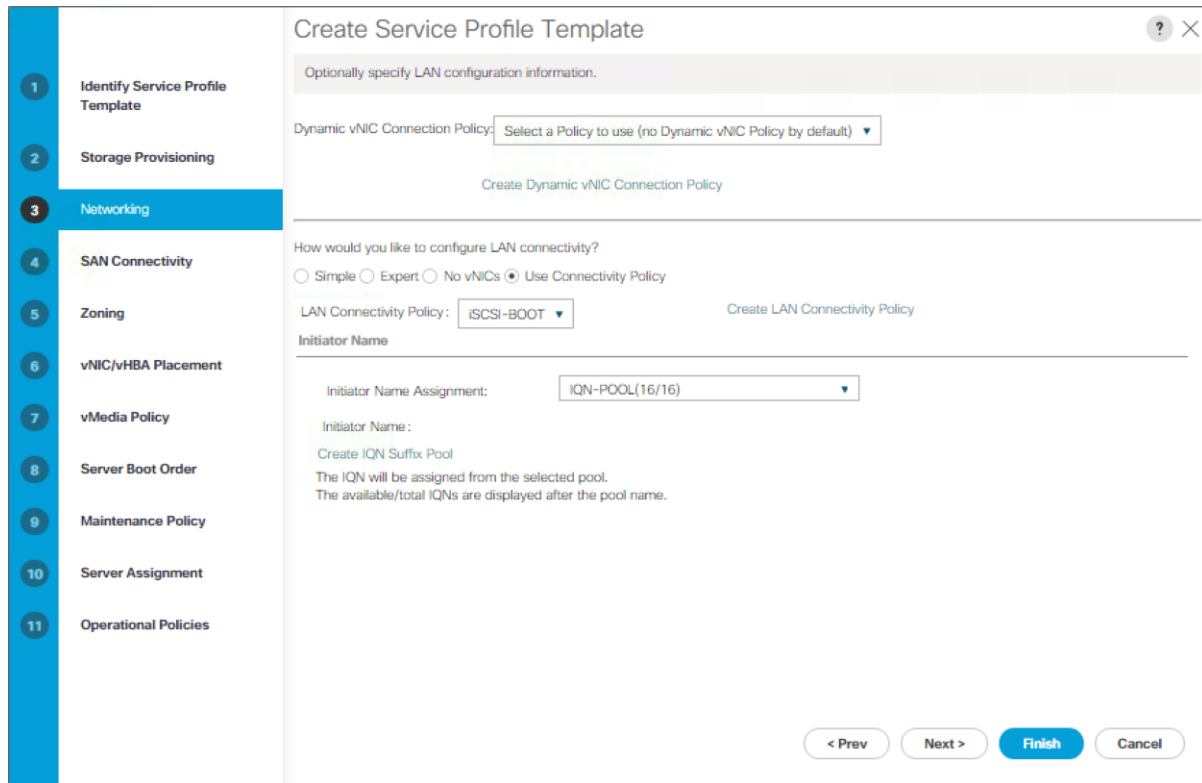
8. Click Next.

Configure Storage Provisioning

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

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 iSCSI-BOOT from the LAN Connectivity Policy pull-down.
4. Select IQN_POOL in Initiator Name Assignment.



5. Click Next.

Configure Storage Options

1. Select No vHBAs for the "How would you like to configure SAN connectivity?" field.
2. Click Next.

Configure Zoning Options

1. Click Next.

Configure vNIC/HBA Placement

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

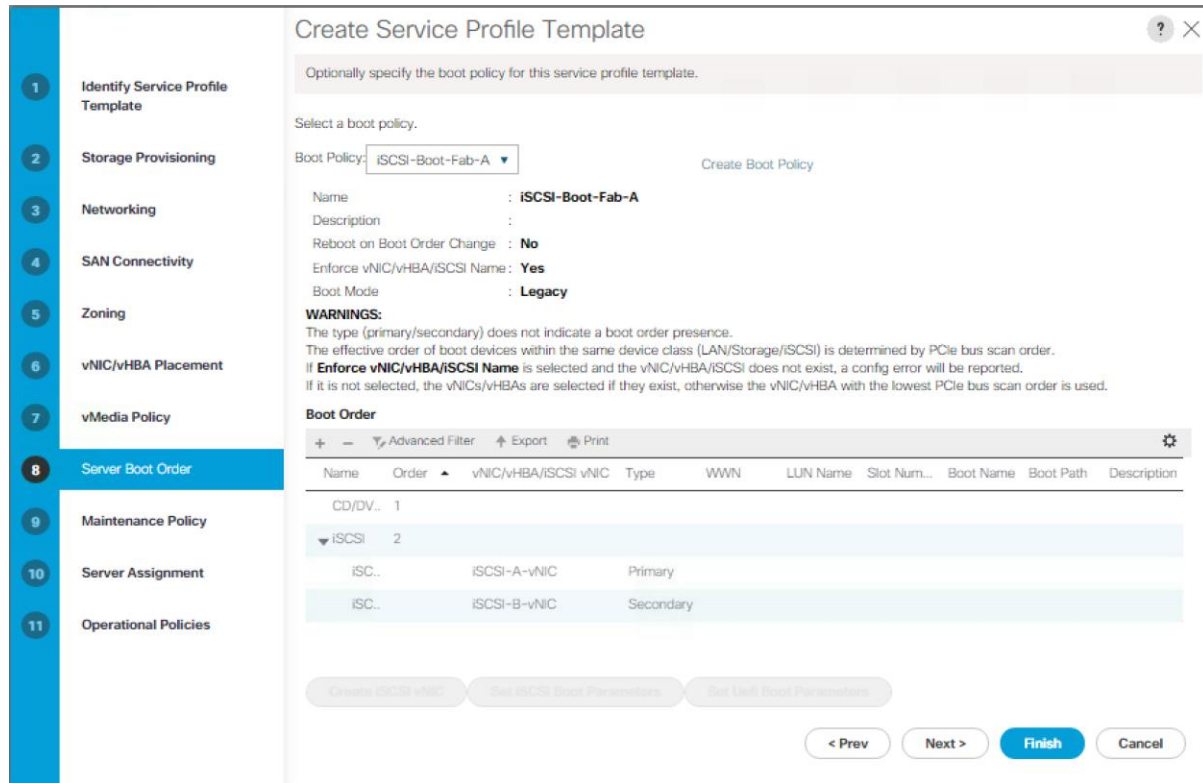
Configure vMedia Policy

1. Do not select a vMedia Policy.

2. Click Next.

Configure Server Boot Order

1. Select iSCSI-Boot-Fab-A for Boot Policy.



2. In the Boot 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_IP_POOL_A as the "Initiator IP address Policy".
7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-MS-SVM, login into storage cluster management interface and run "iscsi show" command".

```
bb04-affa300::> iscsi show
Vserver      Target      Target      Status
Name         Name        Alias       Admin
-----
Infra-MS-SVM
              iqn.1992-08.com.netapp:sn.3ee57b44298c11e7a85c00a098a9fec2:vs.5
                          Infra-MS-SVM
                          up
```

10. Enter the IP address of iscsi_lif_01a for the IPv4 Address field.

Create iSCSI Static Target



iSCSI Target Name :
 Priority : **1**
 Port :
 Authentication Profile: [Create iSCSI Authentication Profile](#)
 IPv4 Address :
 LUN ID :

OK

Cancel

11. Click OK to add the iSCSI static target.
12. Click Add.
13. Enter the iSCSI Target Name.
14. Enter the IP address of iscsi_lif_02a for the IPv4 Address field.

Create iSCSI Static Target ? ×

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.

Set iSCSI Boot Parameters

Name : **iSCSI-A-vNIC**

Authentication Profile: <not set> ▼

Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set> ▼

Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-POOL-A(16/16) ▼

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 Pro..	iSCSI IPv4 Adresse...	LUN Id
iqn.1992-08.c..	1	3260		192.168.12.61	0
iqn.1992-08.c..	2	3260		192.168.12.62	0

OK

Cancel

16. In the Boot order, select iSCSI-B-vNIC.
17. Click Set iSCSI Boot Parameters button.
18. In the Set iSCSI Boot Parameters pop-up, leave Authentication Profile to <not set> unless you have independently created one appropriate to your environment
19. Leave the "Initiator Name Assignment" dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps
20. Set iSCSI_IP_POOL_B as the "Initiator IP address Policy".
21. Select iSCSI Static Target Interface option.

22. Click Add.
23. Enter the iSCSI Target Name. To get the iSCSI target name of Infra-MS-SVM, login into storage cluster management interface and run "iscsi show" command".

```
bb04-affa300:~> iscsi show
Vserver      Target      Target      Status
  Name       Alias
-----
Infra-MS-SVM
iqn.1992-08.com.netapp:sn.3ee57b44298c11e7a85c00a098a9fec2:vs.5
                               Infra-MS-SVM      up
```

24. Enter the IP address of iscsi_lif_01b for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :

Priority :

Port :

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

IPv4 Address :

LUN ID :

25. Click OK to add the iSCSI static target.
26. Click Add.
27. Enter the iSCSI Target Name.
28. Enter the IP address of iscsi_lif_02b for the IPv4 Address field.

Create iSCSI Static Target ? X

iSCSI Target Name :
Priority : **2**
Port :
Authentication Profile : [Create iSCSI Authentication Profile](#)
IPv4 Address :
LUN ID :

29. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters ? X

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set>

Create IQN Suffix Pool

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

Initiator Address

Initiator IP Address Policy: iSCSI-IP-POOL-B(16/16)

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 Pro..	iSCSI IPv4 Addre...	LUN Id
iqn.1992-08.c..	1	3260		192.168.22.61	0
iqn.1992-08.c..	2	3260		192.168.22.62	0

OK
Cancel

30. Click Next.

Configure Maintenance Policy

1. Change the Maintenance Policy to default.

Create Service Profile Template ? X

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
Reboot Policy	: User Ack

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select MS-Server-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select "UCS-Broadwell" for the Server Pool Qualification.
4. Expand Firmware Management at the bottom of the page and select the default policy

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: MS-Server-Pool 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: <not set>

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: MS-HostFirmware Create Host Firmware Package

< Prev Next > Finish Cancel

5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template ? X

Optionally specify information that affects how the system operates.

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

+ External IPMI 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

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

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. Connect to UCS Manager, click Servers on the left pane.
5. Select Service Profile Templates > root > Service Template Hyper-V-Host-iSCSI.
6. Right-click Hyper-V-Host-iSCSI and select Create Service Profiles from Template.
7. Enter Hyper-V-iSCSI-Host-o as the service profile prefix.
8. Enter 1 as "Name Suffix Starting Number."
9. Enter 2 as the "Number of Instances."
10. Click OK to create the service profiles.

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

11. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 9 and Table 10.

Table 9 iSCSI LIFs for iSCSI IQN

SVM	Target: IQN
Infra-MS-SVM	



To obtain the iSCSI IQN, run `iscsi show` command on the storage cluster management interface.

Table 10 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
Hyper-V-iSCSI-Host-01		< Hyper-V-iSCSI-Host-01-iqn>
Hyper-V-iSCSI-Host-02		< Hyper-V-iSCSI-Host-02-iqn>



To obtain the iSCSI vNIC IQN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "iSCSI vNICs" tab on the right. The "Initiator Name" is displayed at the top of the page under the "Service Profile Initiator Name".

Storage Configuration – Boot LUNs and Igroups

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -size 200GB -ostype windows_2008 -space-reserve disabled
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -size 200GB -ostype windows_2008 -space-reserve disabled
```

Create Witness and Data LUN

A witness LUN is required in a Hyper-V cluster. To create the witness and Data LUN, run the following command:

```
lun create -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -size 1GB -ostype windows_2008 -space-reserve disabled
lun create -vserver Infra-MS-SVM -volume infra_datastore_1 -lun data_FC_6332 -size 250GB -ostype windows_2008 -space-reserve disabled
```

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-01 -protocol iscsi -ostype windows -initiator <vm-host-infra-01-iqn>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-02 -protocol iscsi -ostype windows -initiator <vm-host-infra-02-iqn>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-All -protocol iscsi -ostype windows -initiator <vm-host-infra-01-iqn>,<vm-host-infra-02-iqn>
```

Map Boot LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-id 0
lun map -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -igroup VM-Host-Infra-All -lun-id 1
lun map -vserver Infra-MS-SVM -volume infra_datastore_1 -lun data_FC_6332 -igroup VM-Host-Infra-All -lun-id 2
```

Microsoft Windows Server 2016 Hyper-V Deployment Procedure

Setup the Microsoft Windows 2016 install

Most of the sub-sections of Microsoft Windows 2016 install are similar to the steps covered in the main document; hence the delta changes required for iSCSI deployment are captured below.

Install Windows Server 2016

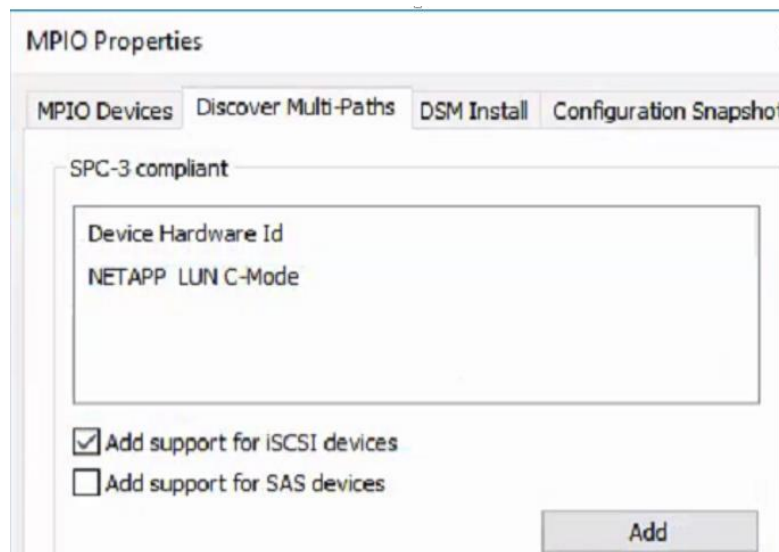
To complete this section, refer to the corresponding section of the main document and execute all the steps.

Install Chipset and Windows eNIC Drivers

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Install Windows Roles and Features

1. Complete the installation steps from the main section and enable the roles and features as mentioned in the main document.
2. For configuring MPIO, Open Server Manager and click Tools > MPIO.
3. Click Discover Multi-Paths tab.
4. Select the check box next to "Add Support for iSCSI Devices" and also select the "NetApp LUN C-Mode" under the Hardware ID. Click Add and reboot the server.

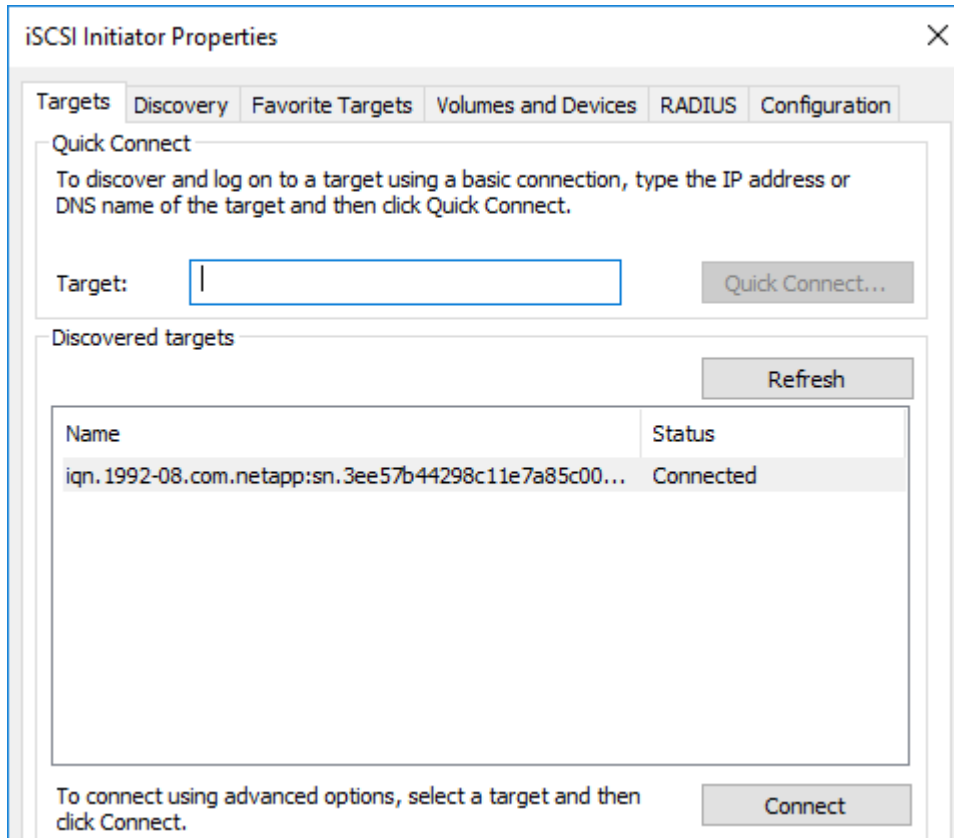


Install NetApp Host Utilities

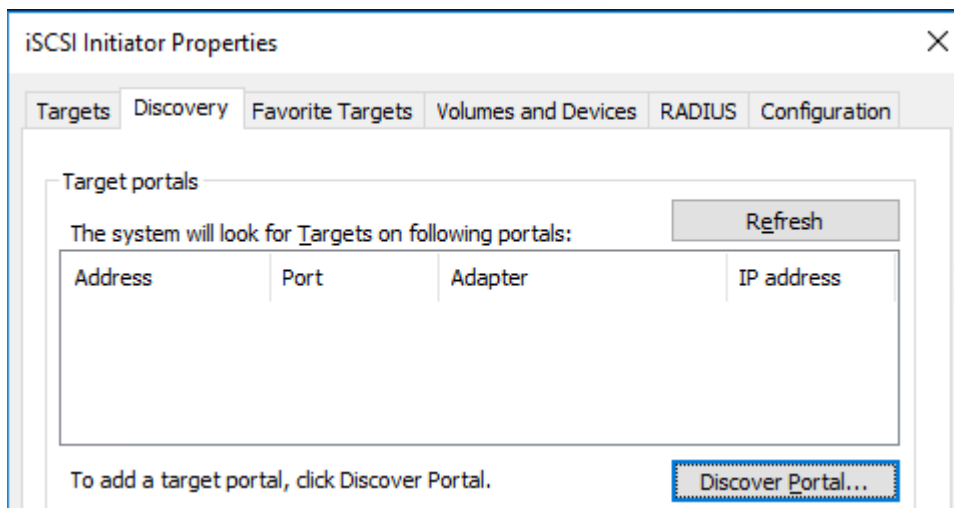
To complete this section, refer to corresponding section of the main document and execute all the steps.

Configure Microsoft iSCSI Initiator

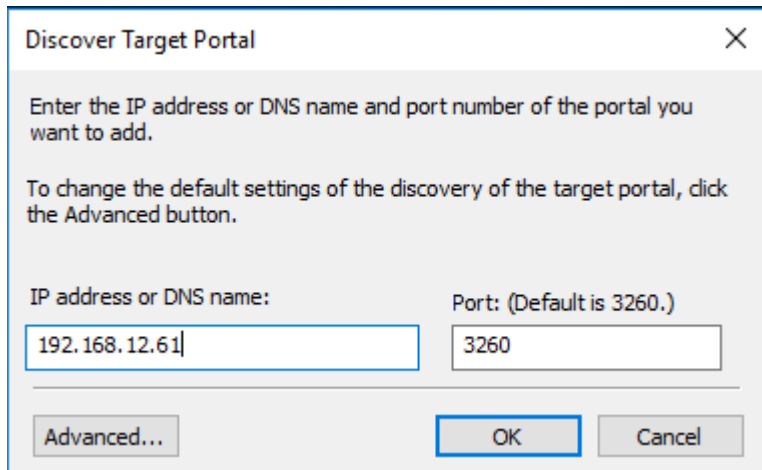
1. Open Server Manager, click Tools and then click iSCSI Initiator and click Yes to start the service.
2. Click on the Targets tab and you will see a discovered target with status as connected. This entry is because of the iSCSI SAN boot.



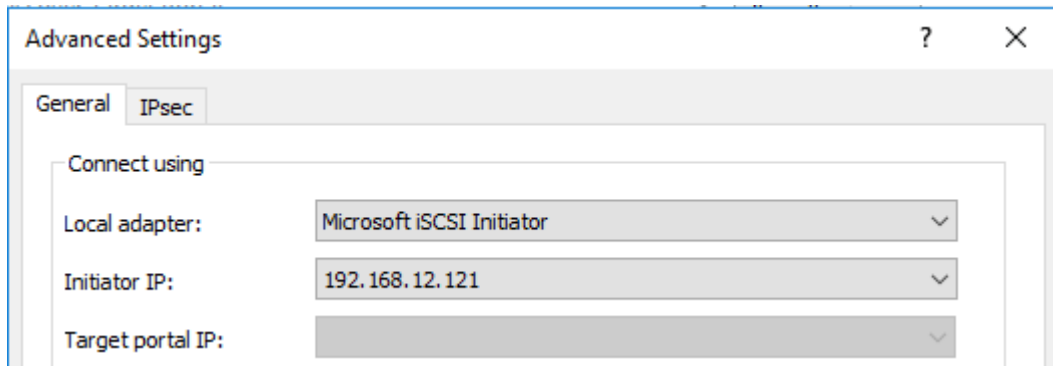
3. In the iSCSI Initiator properties, click the Discovery tab and then click Discover portal.



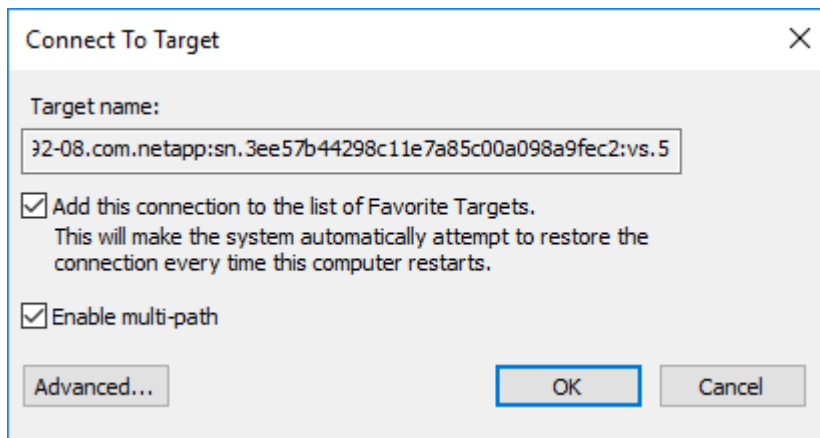
4. In the Discover Target Portal window, enter the IP address or DNS name of the iSCSI target and click Advanced.



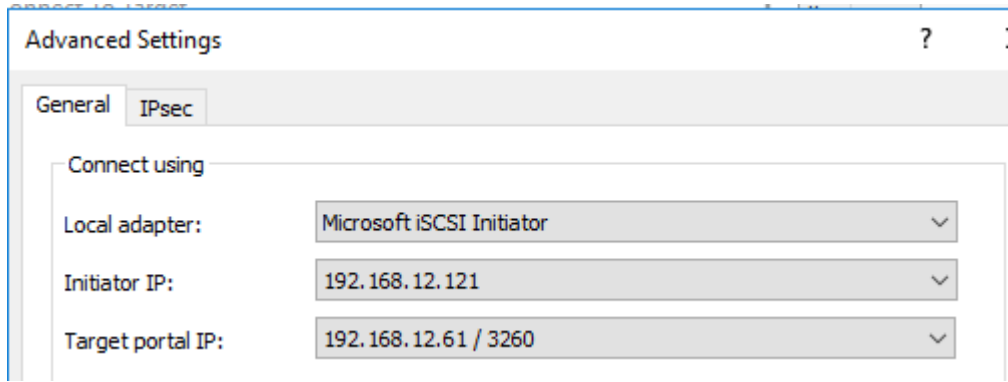
- 5. In the Advanced setting, select Microsoft iSCSI Initiator as local adapter and an iSCSI IP address for the Initiator IP as shown in the below figure.



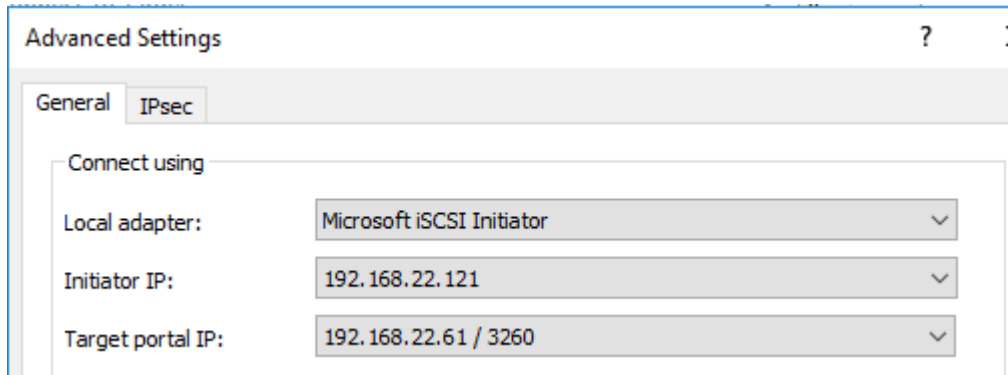
- 6. Click on Targets tab and select the discovered target and click on Connect.
- 7. In the Connect To Target, select Enable multi-path and click on Advanced.



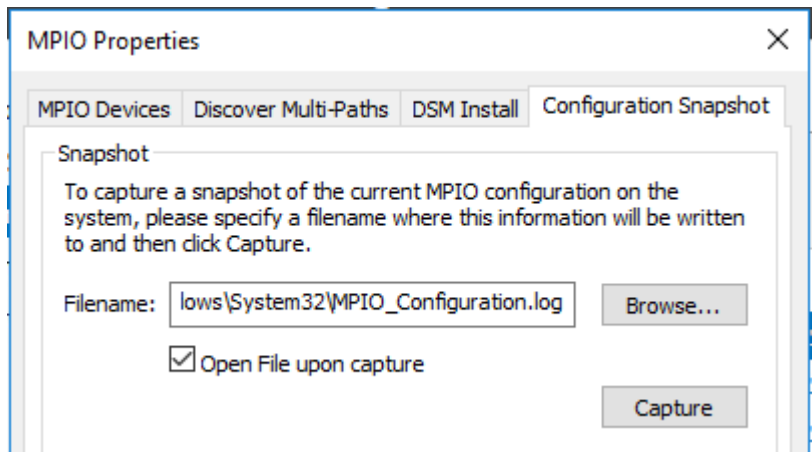
- 8. In the Advanced settings, configure the Initiator IP and the Target portal IP for the iSCSI-A path as shown in the figure below.



9. Repeat steps 6 to 8 and configure as shown in the figure below for the iSCSI-B path.



10. Verify the MPIO configuration by navigating to the Server Manager > Tools > MPIO > Configuration Snapshot. Select "Open File upon capture" and click Capture.



11. The MPIO configuration log file opens showing the number of paths and the load balance policies used to manage the LUNs/disks.

```

MPIO_Configuration - Notepad
File Edit Format View Help
MPIO Storage Snapshot on Sunday, 23 July 2017, at 23:06:41.587

Registered DSMs: 1
=====
+-----+-----+----+----+----+----+-----+
| DSM Name          | Version          | PRP | RC | RI | PVP | PVE |
+-----+-----+----+----+----+----+-----+
| Microsoft DSM     | 010.0000.14393.0000 | 0130|0006|0001|030|False|
+-----+-----+----+----+----+----+-----+

Microsoft DSM
=====
|
MPIO Disk1: 04 Paths, Round Robin with Subset, Implicit Only
SN: 600A0980383038625A244A6C3873524D
Supported Load Balance Policies: FOO RRWS LQD WP LB

Path ID          State          SCSI Address      Weight
-----
0000000077010003 Active/Unoptimized 001|000|003|001  0
TPG_State: Active/Unoptimized, TPG_Id: 1000, TP_Id: 7
Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

0000000077010002 Active/Unoptimized 001|000|002|001  0
TPG_State: Active/Unoptimized, TPG_Id: 1000, TP_Id: 6
Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

0000000077010001 Active/Optimized   001|000|001|001  0
TPG_State: Active/Optimized , TPG_Id: 1001, TP_Id: 9
Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

0000000077010000 Active/Optimized   001|000|000|001  0
TPG_State: Active/Optimized , TPG_Id: 1001, TP_Id: 8
Adapter: Microsoft iSCSI Initiator... (B|D|F: 000|000|000)
Controller: 46616B65436F6E74726F6C6C6572 (State: Active)

MPIO Disk0: 04 Paths, Round Robin with Subset, Implicit Only
SN: 600A098038303862535D4A686F74524D
Supported Load Balance Policies: FOO RRWS LQD WP LB
    
```

Host Renaming and Join to Domain

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Deploying and Managing Hyper-V Clusters using System Center 2016 VMM

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Appendix - FCoE Solution

This section of FlexPod deployment will show configuration steps for both the Cisco UCS 6332-16UP and Cisco UCS 6248UP Fabric Interconnects (FI) in a design that will support direct connectivity to NetApp AFF using Fibre Channel over Ethernet (FCoE).

Configuration steps will be referenced for both fabric interconnects and will be called out by the specific model where steps have differed.

Figure 13 shows the Microsoft Hyper-V 2016 built on FlexPod components and the network connections for a configuration with the Cisco UCS 6332-16UP Fabric Interconnects with storage FCoE connections directly connected to the fabric interconnect. This design has 40Gb Ethernet connections between the Cisco UCS 5108 Blade Chassis and C-Series rackmounts and the Cisco UCS Fabric Interconnect, and between the Cisco UCS Fabric Interconnect and Cisco Nexus 9000. This design also has a 10Gb FCoE connection between Cisco UCS Fabric Interconnect and NetApp AFF A300. FC zoning is done in the Cisco UCS Fabric Interconnect. This infrastructure is deployed to provide FCoE-booted hosts with file-level and block-level access to shared storage with use cases that do not require the Cisco MDS SAN connectivity or scale.

Figure 13 FlexPod with Cisco UCS 6332-16UP Fabric Interconnects and Cisco UCS Direct Connect SAN

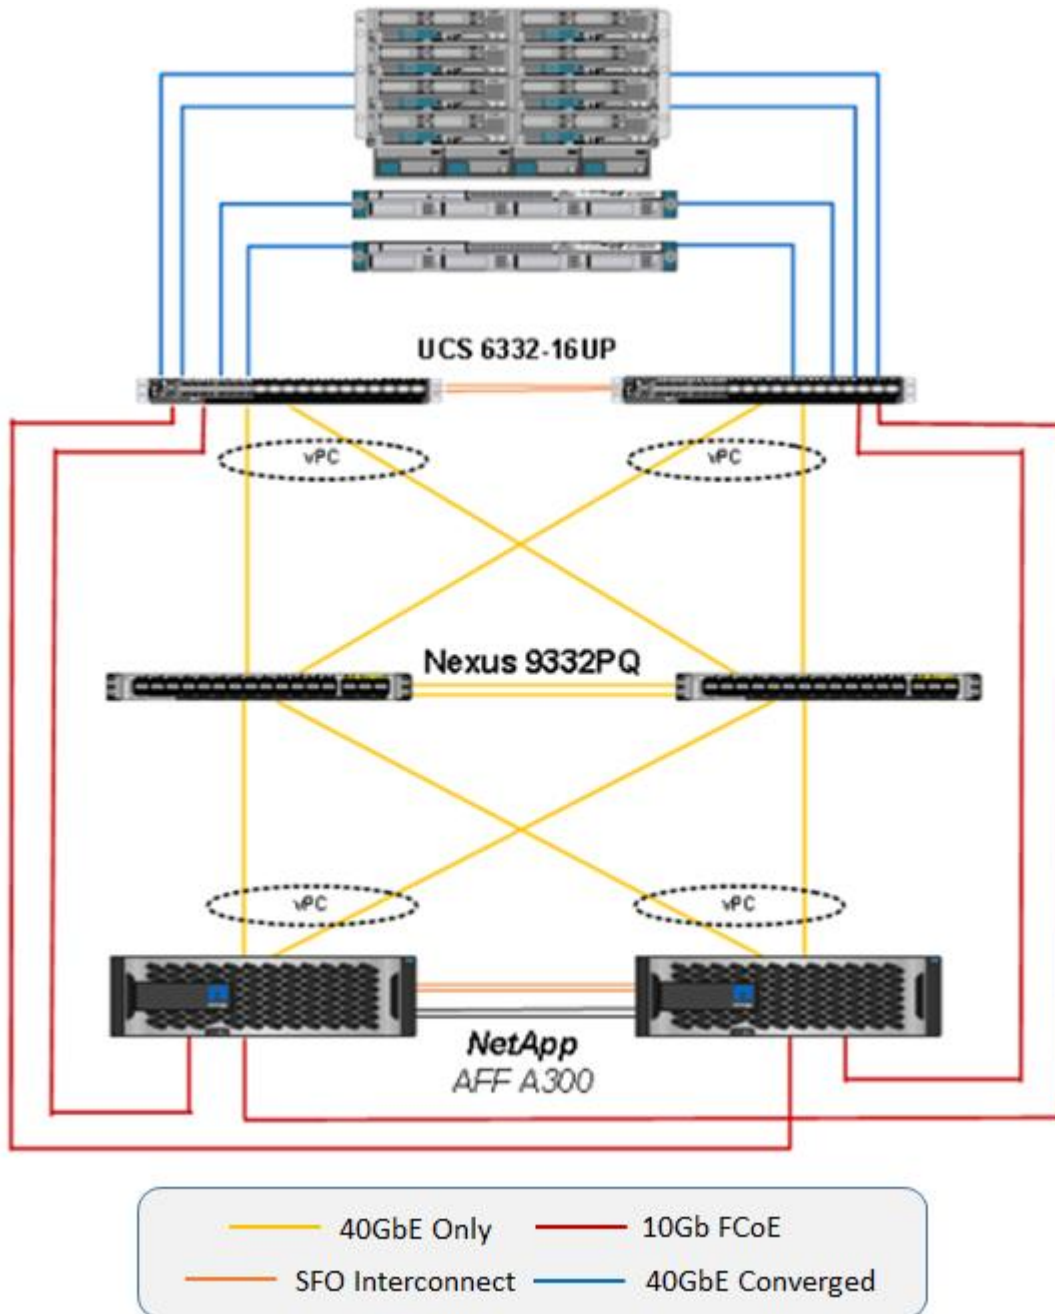



Figure 14 shows the Hyper-V built on FlexPod components and the network connections for a configuration with the Cisco UCS 6248UP Fabric Interconnects with storage FC connections directly connected to the fabric interconnect. This design has 10Gb Ethernet connections throughout the architecture. This design also has 10Gb FCoE connections between the Cisco UCS Fabric Interconnect and the NetApp AFF family of storage controllers. This infrastructure is also deployed to provide FCoE-booted hosts with file-level and block-level access to shared storage with use cases that do not require the Cisco MDS SAN connectivity or scale.

Figure 14 FlexPod with Cisco UCS 6248UP Fabric Interconnects and Cisco UCS Direct Connect SAN



Network Switch Configuration

This section provides a detailed procedure for configuring the Cisco Nexus 9000s for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Physical Connectivity

Figure 15 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect for Direct Connect FCoE SAN

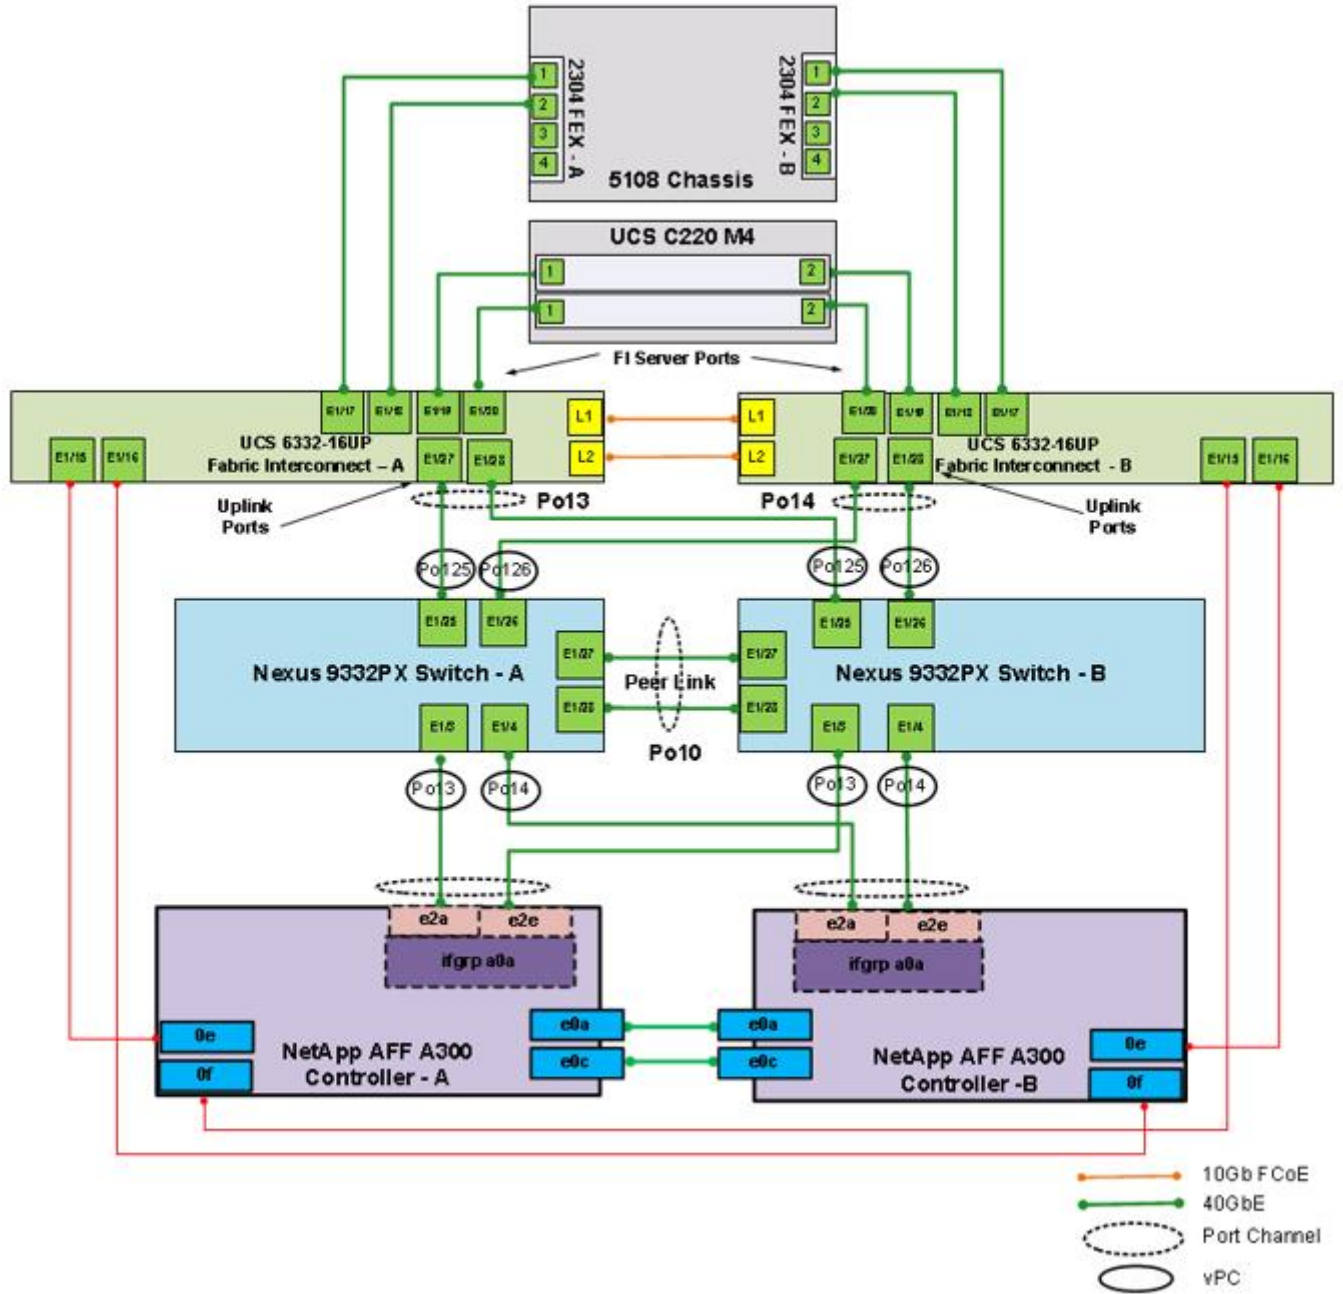
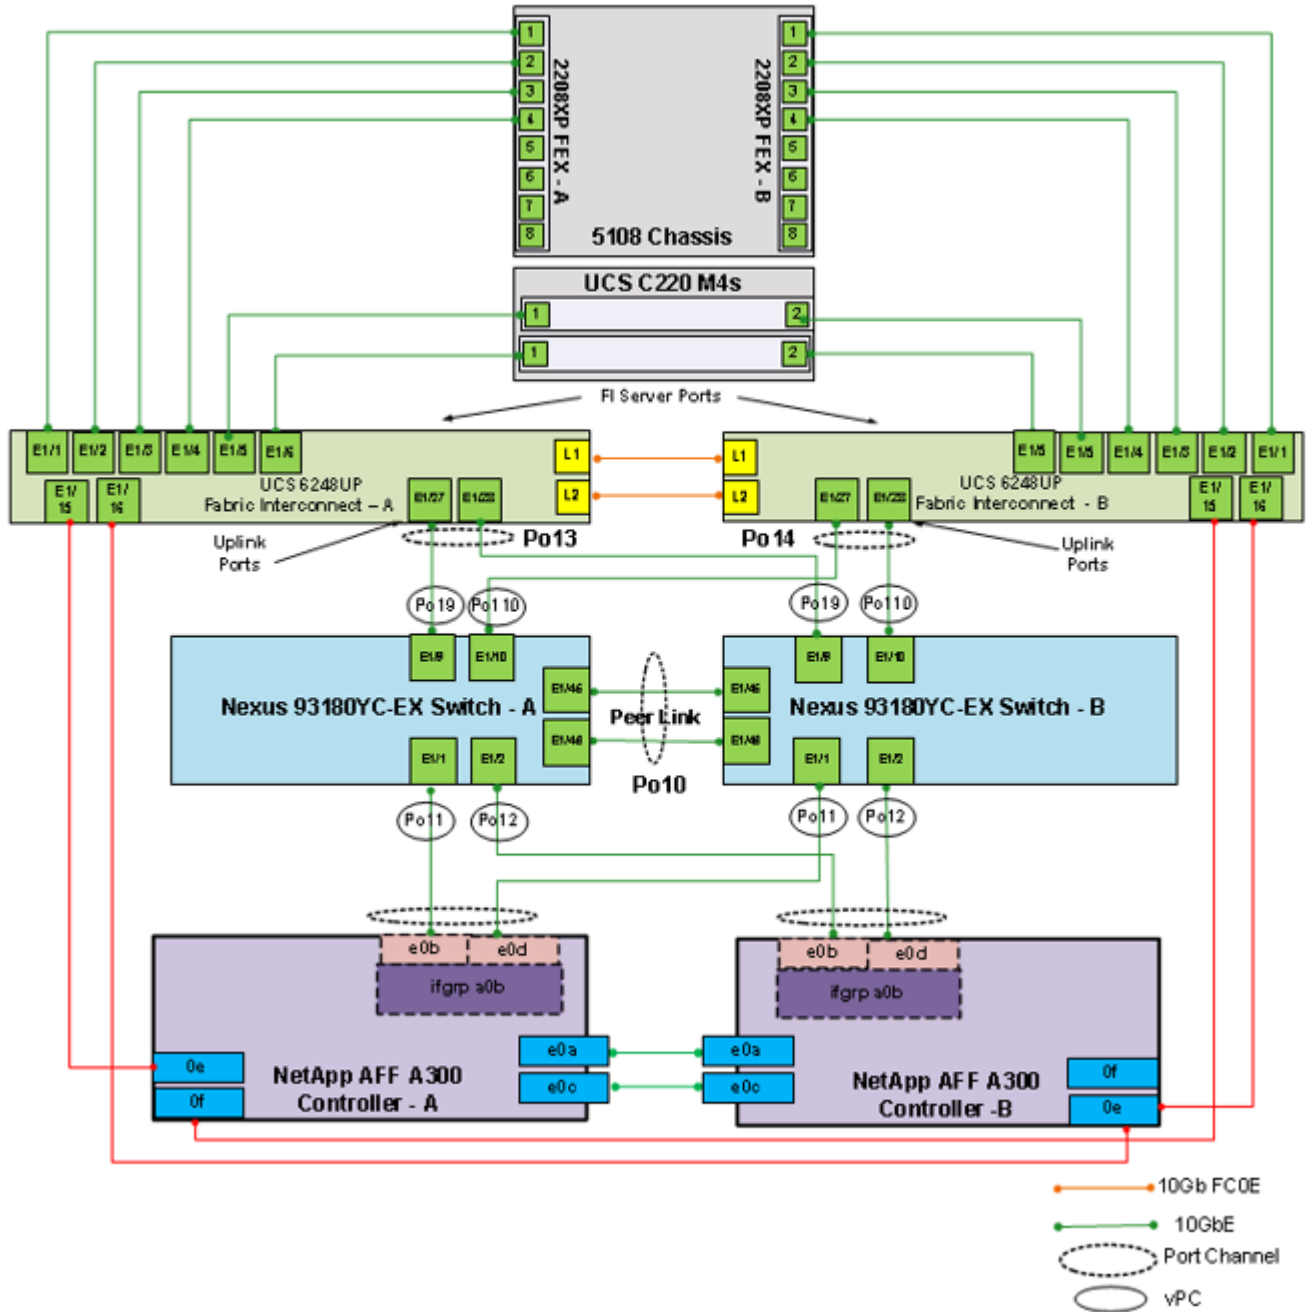


Figure 16 FlexPod Cabling with Cisco UCS 6248UP Fabric Interconnect for Direct Connect FCoE SAN



FlexPod Cisco Nexus Base

To complete this section, refer to the corresponding section of the main document and execute all the steps.

FlexPod Cisco Nexus Switch Configuration

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Storage Configuration

Set Onboard Unified Target Adapter 2 Port Personality

In order to use FCoE storage targets, CNA ports must be configured on the storage. To set the personality of the onboard unified target adapter 2 (UTA2), complete the following steps for both controllers in an HA pair:

1. Verify the Current Mode and Current Type properties of the ports by running the `ucadmin show` command.

```
ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
<st-node01>	0e	fc	target	-	-	online
<st-node01>	0f	fc	target	-	-	online
<st-node01>	0g	cna	target	-	-	online
<st-node01>	0h	cna	target	-	-	online
<st-node02>	0e	fc	target	-	-	online
<st-node02>	0f	fc	target	-	-	online
<st-node02>	0g	cna	target	-	-	online
<st-node02>	0h	cna	target	-	-	online

8 entries were displayed.

2. Verify that the Current Mode and Current Type properties for all ports are set properly. Set the ports used for FCoE connectivity to mode `cna`. The port type for all protocols should be set to `target`. Change the port personality by running the following command:

```
ucadmin modify -node <home-node-of-the-port> -adapter <port-name> -mode cna -type target.
```



The ports must be offline to run this command. To take an adapter offline, run the `fcplib adapter modify -node <home-node-of-the-port> -adapter <port-name> -state down` command. Ports must be converted in pairs (for example, 0e and 0f).



After conversion, a reboot is required. After reboot, bring the ports online by running `fcplib adapter modify -node <home-node-of-the-port> -adapter <port-name> -state up`.

Create FC LIFs

Run the following commands to create four FC LIFs (two on each node) by using the previously configured FCoE ports.:

```
network interface create -vserver Infra-MS-SVM -lif fcp_lif01a -role data -data-protocol fcp -home-node <st-node01> -home-port 0e -status-admin up

network interface create -vserver Infra-MS-SVM -lif fcp_lif01b -role data -data-protocol fcp -home-node <st-node01> -home-port 0f -status-admin up

network interface create -vserver Infra-MS-SVM -lif fcp_lif02a -role data -data-protocol fcp -home-node <st-node02> -home-port 0e -status-admin up

network interface create -vserver Infra-MS-SVM -lif fcp_lif02b -role data -data-protocol fcp -home-node <st-node02> -home-port 0f -status-admin up
```



From storage controller's perspective, LIF's are created according to the native protocol to be used. Therefore, FC LIFs are created using FCoE ports.



For additional storage related tasks, please see the storage configuration portion of this document.

Create Boot LUNs

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -size 200GB -ostype windows_2008 -space-reserve disabled
lun create -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -size 200GB -ostype windows_2008 -space-reserve disabled
```

Create Witness and Data LUN

A witness LUN is required in a Hyper-V cluster. To create the witness and Data LUN, run the following command:

```
lun create -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -size 1GB -ostype windows_2008 -space-reserve disabled
lun create -vserver Infra-MS-SVM -volume infra_datastore_1 -lun data_FC_6332 -size 250GB -ostype windows_2008 -space-reserve disabled
```

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-01 -protocol fcp -ostype windows -initiator <vm-host-infra-01-wwpna>,<vm-host-infra-01-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-02 -protocol fcp -ostype windows -initiator <vm-host-infra-02-wwpna>,<vm-host-infra-02-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-All -protocol fcp -ostype windows -initiator <vm-host-infra-01-wwpna>,<vm-host-infra-01-wwpnb>,<vm-host-infra-02-wwpna>,<vm-host-infra-02-wwpnb>
```

Map Boot LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-id 0
lun map -vserver Infra-MS-SVM -volume HV_boot -lun VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-id 0
lun map -vserver Infra-MS-SVM -volume witness_FC_6332 -lun witness_FC_6332 -igroup VM-Host-Infra-All -lun-id 1
lun map -vserver Infra-MS-SVM -volume infra_datastore_1 -lun data_FC_6332 -igroup VM-Host-Infra-All -lun-id 2
```



For additional storage related tasks, please see the storage configuration portion of this document.

Server Configuration

Perform Initial Setup of Cisco UCS 6332-16UP and 6248UP Fabric Interconnects for FlexPod Environments

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

Cisco UCS Fabric Interconnect A

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

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

```
Enter the configuration method: gui
```

```
Physical switch Mgmt0 IP address: <ucsa-mgmt-ip>
```

```
Physical switch Mgmt0 IPv4 netmask: <ucsa-mgmt-mask>
```

```
IPv4 address of the default gateway: <ucsa-mgmt-gateway>
```

2. Using a supported web browser, connect to <ucsa-mgmt-ip>, accept the security prompts, and click the 'Express Setup' link under HTML.
3. Select Initial Setup and click Submit.
4. Select Enable clustering, Fabric A, and IPv4.
5. Fill in the Virtual IP Address with the UCS cluster IP.
6. Completely fill in the System setup section. For system name, use the overall UCS system name. For the Mgmt IP Address, use <ucsa-mgmt-ip>.

Basic Settings

Cluster and Fabric setup

Enable clustering
 Standalone mode
 Synchronize

Fabric Setup: Fabric A Fabric B

IPv4
 IPv6

Virtual IP Address: . . .

System setup

Enforce strong password?: Yes No

System name:

Admin Password: Confirm Admin password:

Mgmt IP Address: . . . Mgmt IP Netmask: . . .

Default Gateway: . . .

DNS Server IP: . . . Domain Name :

UCS Central managed environment

UCS Central IP: . . . Shared Secret:

- Click Submit.

Cisco UCS Fabric Interconnect B

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

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

Enter the configuration method: `gui`

Physical switch Mgmt0 IP address: `<ucsb-mgmt-ip>`

Physical switch Mgmt0 IPv4 netmask: `<ucsb-mgmt-mask>`

IPv4 address of the default gateway: `<ucsb-mgmt-gateway>`

2. Using a supported web browser, connect to <ucs-b-mgmt-ip>, accept the security prompts, and click the 'Express Setup' link under HTML.
3. Under System setup, enter the Admin Password entered above and click Submit.
4. Enter <ucs-b-mgmt-ip> for the Mgmt IP Address and click Submit.

Cisco UCS Direct Storage Connect Setup

Log in to Cisco UCS Manager



The following steps are the same between the UCS 6332-16UP and the UCS 6248UP Fabric Interconnects unless otherwise noted

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

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



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

2. Click the Launch UCS Manager link under HTML to launch Cisco UCS Manager.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 3.1(3a)

This document assumes the use of Cisco UCS 3.1(3a). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.1(3a), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following 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.

Configure Cisco UCS Call Home

It is highly recommended by Cisco to configure Call Home in UCSM. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, complete the following steps:

1. In Cisco UCS Manager, click Admin on the 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.

Place UCS Fabric Interconnects in Fiber Channel Switching Mode

In order to use Fiber Channel Storage Ports for storage directly connected to the Cisco UCS fabric interconnects, the fabric interconnects must be changed from fiber channel end host mode to fiber channel switching mode.

To place the fabric interconnects in fiber channel switching mode, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. In the center pane, select set FC Switching Mode. Click Yes and OK for the confirmation message.

The screenshot displays the Cisco UCS Manager interface for configuring Fabric Interconnect A (subordinate). The left sidebar shows the navigation menu with 'Fabric Interconnect A (subordinate)' selected. The main content area is divided into several sections:

- Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate)**: Breadcrumb trail.
- General**: Tab selected.
- Fault Summary**: Shows four status indicators (red X, orange triangle, yellow triangle, green circle) with a count of 0 for each.
- Status**:
 - Overall Status: **Operable** (green up arrow)
 - Thermal: **OK** (green up arrow)
 - Ethernet Mode: **End Host**
 - FC Mode: **End Host**
 - Admin Evac Mode: **Off**
 - Oper Evac Mode: **Off**
- Actions**:
 - Configure Evacuation
 - Configure Unified Ports
 - Internal Fabric Manager
 - LAN Uplinks Manager
 - NAS Appliance Manager
 - SAN Uplinks Manager
 - SAN Storage Manager
 - Enable Ports ▼
 - Disable Ports ▼
 - Set Ethernet End-Host Mode
 - Set Ethernet Switching Mode
 - Set FC End-Host Mode
 - Set FC Switching Mode** (highlighted with a red box)
 - Activate Firmware
- Physical Display**: Shows a visual representation of the hardware components.
- Properties**:
 - Name: **A**
 - Product Name: **Cisco UCS 6332 16UP**
 - Vendor: **Cisco Systems, Inc.**
 - Revision: **0**
 - Available Memory: **28.521 (GB)**
 - Locator LED: **Off**
- Firmware**:
 - Boot-loader Version: **v1.0.0(08/31/2015)**

4. Wait for both Fabric Interconnects to reboot by monitoring the console ports and log back into Cisco UCS Manager.

Add Block of IP Addresses for KVM Access

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

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

From : 192.168.94.241

Subnet Mask : 255.255.255.0

Primary DNS : 0.0.0.0

Size : 8

Default Gateway : 192.168.94.254

Secondary DNS : 0.0.0.0

OK Cancel

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

Synchronize Cisco UCS to NTP

To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, complete the following steps:

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

8. Click Add NTP Server.
9. Enter <switch-b-ntp-ip> and click OK. Click OK on the confirmation.

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, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.

Equipment

[Main Topology View](#)
 [Fabric Interconnects](#)
 [Servers](#)
 [Thermal](#)
 [Decommissioned](#)
 [Firmware Management](#)
Policies
[Faults](#)

Global Policies
[Autoconfig Policies](#)
[Server Inheritance Policies](#)
[Server Discovery Policies](#)
[SEL Policy](#)
[Power Groups](#)

Chassis/FEX Discovery Policy

Action :

Link Grouping Preference : None Port Channel

Backplane Speed Preference : 40G 4x10G

5. Click Save Changes.
6. Click OK.

Enable Server and Uplink Ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers, right-click them, and select "Configure as Server Port."
5. Click Yes to confirm server ports and click OK.
6. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
7. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.



The last 6 ports of the UCS 6332 and UCS 6332-16UP FIs will only work with optical based QSFP transceivers and AOC cables, so they can be better utilized as uplinks to upstream resources that might be optical only.

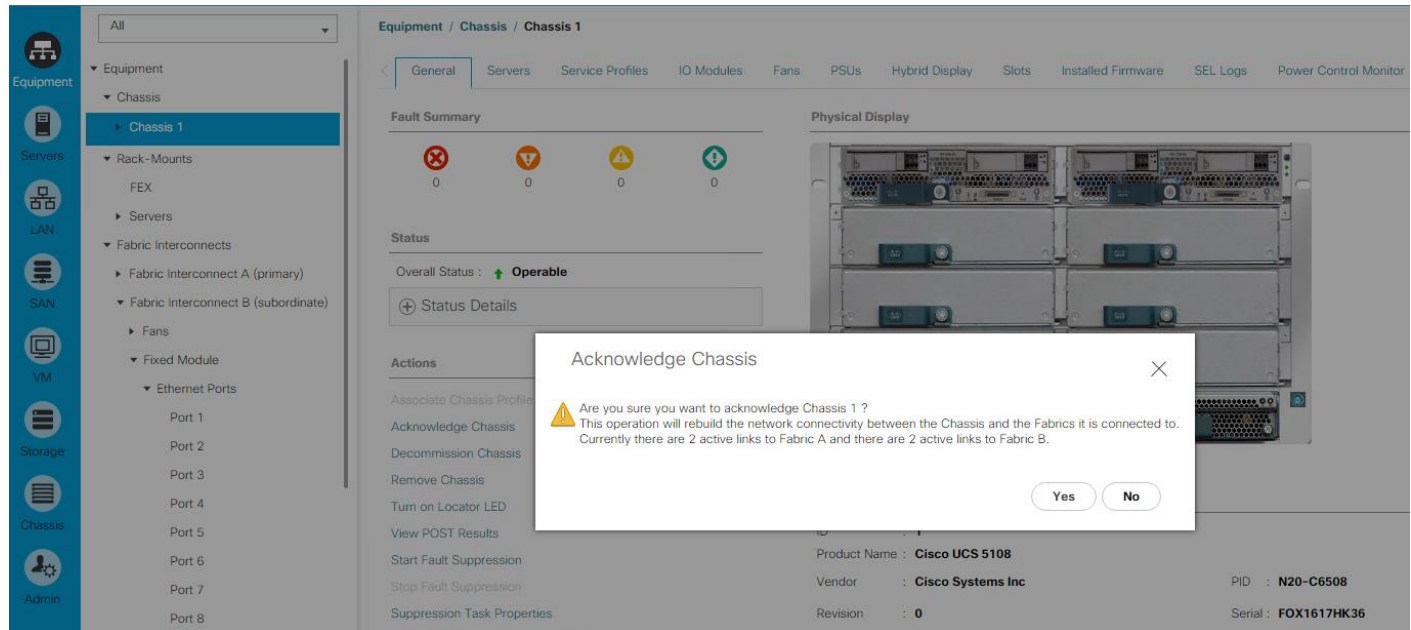
8. Click Yes to confirm uplink ports and click OK.
9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
10. Expand Ethernet Ports.
11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.
12. Click Yes to confirm server ports and click OK.
13. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

- Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and any external 2232 FEX modules, complete the following steps:

- In Cisco UCS Manager, click Equipment on the left.
- 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 2232 FEX are part of the configuration, expand Rack Mounts and FEX.
- Right-click each FEX that is listed and select Acknowledge FEX.
- Click Yes and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

To configure the necessary port channels out of the Cisco UCS environment, complete the following steps:

- In Cisco UCS Manager, click LAN on the left.



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

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

4. Select Create Port Channel.
5. Enter 125 as the unique ID of the port channel.
6. Enter `vPC-125-Nexus` as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 126 as the unique ID of the port channel.
16. Enter `vPC-126-Nexus` as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create a WWNN Pool for FC Boot

To configure the necessary WWNN pool for the Cisco UCS environment, complete the following steps on Cisco UCS Manager.

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

Create WWNN Pool

1 Define Name and Description

2 Add WWN Blocks

Name : WWNN-POOL

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

8. Click Next.
9. Click Add.
10. Modify the From field as necessary for the UCS Environment

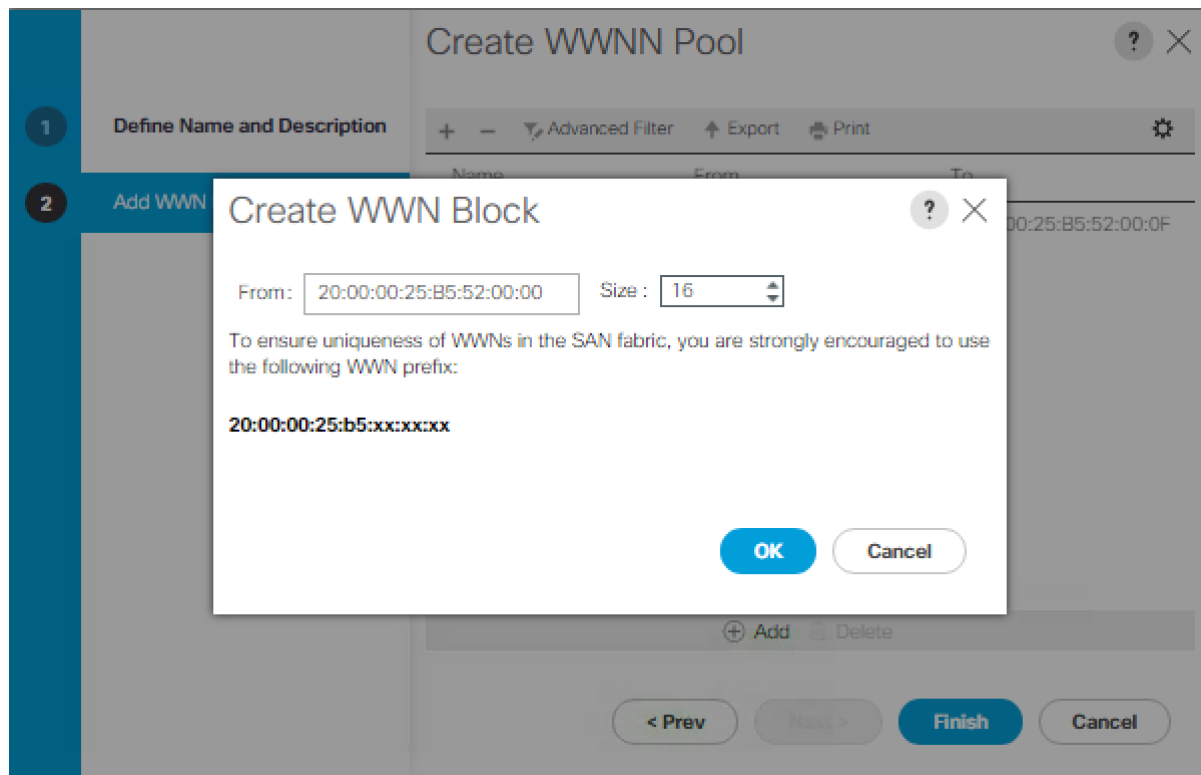


Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to 52 to represent as identifying information for this being in the UCS 6332 in the 4th cabinet



Also, when having multiple UCS domains sitting in adjacency, it is important that these blocks, the WWNN, WWPN, and MAC hold differing values between each set.

11. Specify a size of the WWNN block sufficient to support the available server resources.



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

Create WWPN Pools

To configure the necessary WWPN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. In this procedure, two WWPN pools are created, one for each switching fabric.
4. Right-click WWPN Pools under the root organization.
5. Select Create WWPN Pool to create the WWPN pool.
6. Enter `WWPN-POOL-A` as the name of the WWPN pool.
7. Optional: Enter a description for the WWPN pool.
8. Select **Sequential** for Assignment Order

Create WWPN Pool

1 Define Name and Description

2 Add WWN Blocks

Name : WWPN-POOL-A

Description :

Assignment Order: Default Sequential

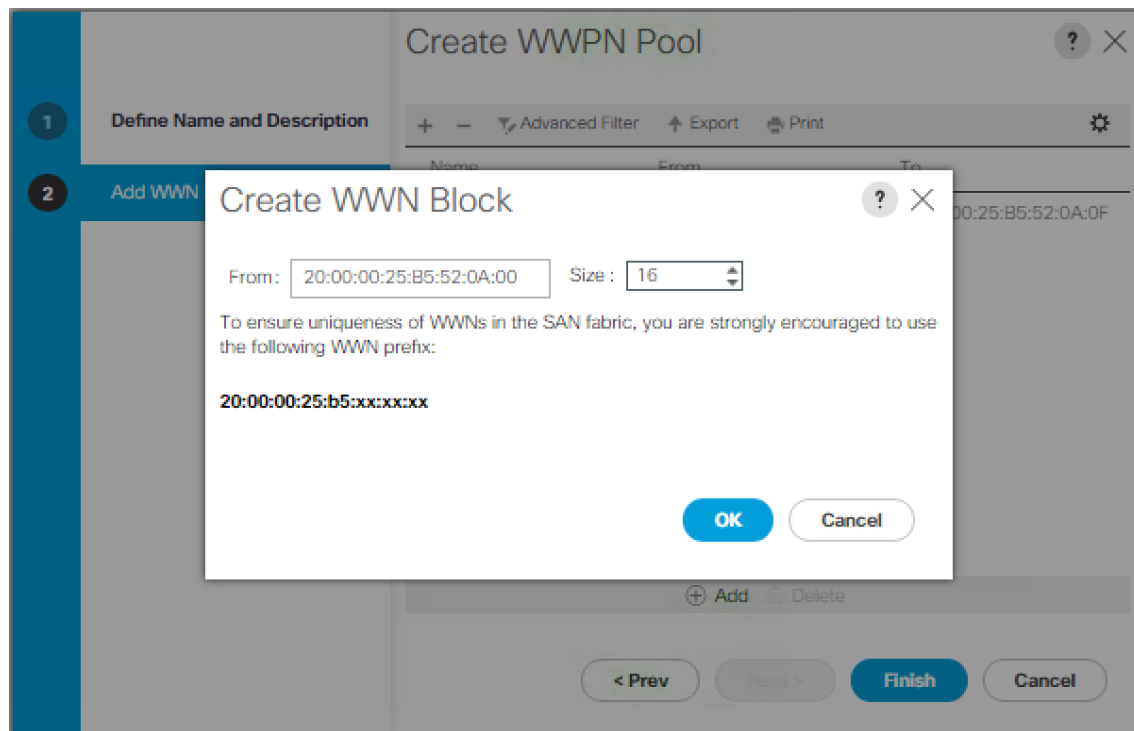
< Prev Next > Finish Cancel

9. Click Next.
10. Click Add.
11. Specify a starting WWPN



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with 20:00:00:25:B5:52:0A:00

12. Specify a size for the WWPN pool that is sufficient to support the available blade or server resources.



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



For the FlexPod solution, the recommendation is to place 0B in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses. Merging this with the pattern we used for the WWNN we see a WWPN block starting with `20:00:00:25:B5:52:0B:00`.

24. Specify a size for the WWPN address pool that is sufficient to support the available blade or server resources.
25. Click OK.

- 26. Click Finish.
- 27. In the confirmation message, click OK

Create Storage VSAN

To configure the necessary virtual storage area networks (VSANs) for the Cisco UCS environment, complete the following steps:

- 1. In Cisco UCS Manager, click the SAN on the left.



In this procedure, two VSANs are created.

- 2. Select SAN > Storage Cloud.
- 3. Right-click VSANs.
- 4. Select Create Storage VSAN.
- 5. Enter VSAN-A as the name of the VSAN to be used for Fabric A
- 6. Set FC Zoning to Enabled.
- 7. Select Fabric A.
- 8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID for Fabric A. It is recommended to use the same ID for both parameters and to use something other than 1.

Create Storage VSAN

Name:

FC Zoning Settings

FC Zoning: Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.	A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VSAN ID that maps to this VSAN.	Enter the VLAN ID that maps to this VSAN.
VSAN ID: <input type="text" value="101"/>	FCoE VLAN: <input type="text" value="101"/>

- 9. Click OK and then click OK again.
- 10. Under Storage Cloud, right-click VSANs.
- 11. Select Create Storage VSAN.
- 12. Enter VSAN-B as the name of the VSAN to be used for Fabric B.

13. Leave FC Zoning set at Disabled.
14. Select Fabric B.
15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID for Fabric B. It is recommended use the same ID for both parameters and to use something other than 1.

Create Storage VSAN ? X

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 B that maps to a VSAN ID that exists only in fabric B. A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.

Enter the VSAN ID that maps to this VSAN. Enter the VLAN ID that maps to this VSAN.

VSAN ID : FCoE VLAN :

16. Click OK, and then click OK again

Configure FCoE Storage Port

To configure the necessary FCoE Storage port for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module > Ethernet Ports
3. Select the ports (15 and 16 for this document) that are connected to the NetApp array, right-click them, and select "Configure as FCoE Storage Port"

Equipment / Fabric Interconnects / Fabric Interconnect A (primary) / Fixed Module / Ethernet Ports

Ethernet Ports

Slot	Aggr. Port ID	Port ID	MAC	If Role
1	0	7	8C:60:4F:BD:62:32	Unconfigured
1	0	8	8C:60:4F:BD:62:33	Unconfigured
1	0	9	8C:60:4F:BD:62:34	Unconfigured
1	0	10	8C:60:4F:BD:62:35	Unconfigured
1	0	11	8C:60:4F:BD:62:36	Unconfigured
1	0	12	8C:60:4F:BD:62:37	Unconfigured
1	0	13	8C:60:4F:BD:62:38	Server
1	0	14	8C:60:4F:BD:62:39	Unconfigured
1	0	15	8C:60:4F:BD:62:3A	Unconfigured
1	0	16	8C:60:4F:BD:62:3B	Unconfigured
1	0	17	8C:60:4F:BD:62:3C	Server
1	0	18	8C:60:4F:BD:62:40	Server
1	0	19	8C:60:4F:BD:62:44	Server
1	0	20	8C:60:4F:BD:62:48	Server
1	0	21	8C:60:4F:BD:62:4C	Unconfigured
1	0	22	8C:60:4F:BD:62:50	Unconfigured

Context menu options for port 16:

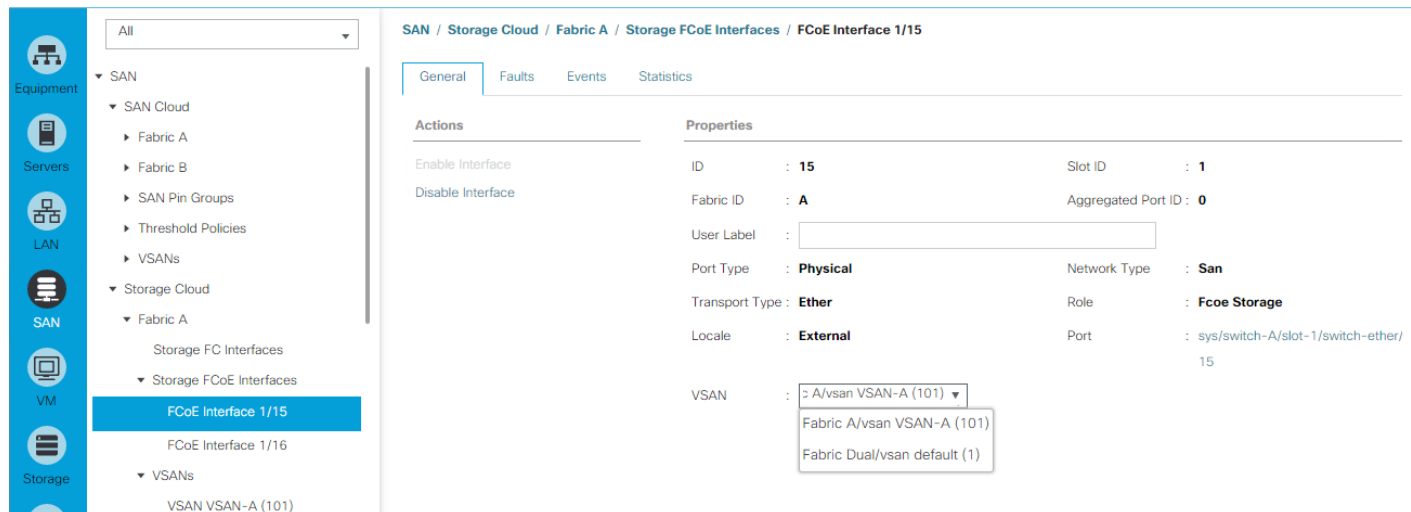
- Enable
- Disable
- Configure as Server Port
- Configure as Uplink Port
- Configure as FCoE Uplink Port
- Configure as FCoE Storage Port**
- Configure as Appliance Port

4. Click Yes to confirm and then click OK.
5. Verify that the ports connected to the NetApp array are now configured as FCoE Storage.
6. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module > Ethernet Ports
7. Select the ports (15 and 16 for this document) that are connected to the NetApp array, right-click them, and select "Configure as FCoE Storage Port"
8. Click Yes to confirm and then click OK.
9. Verify that the ports connected to the NetApp array are now configured as FCoE Storage.

Assign VSANs to FCoE Storage Ports

To assign storage VSANs to FCoE Storage Ports, complete the following steps:

1. In Cisco UCS Manager, Click SAN on the left.
2. Select SAN > Storage Cloud.
3. Expand Fabric A and Storage FCoE Interfaces.
4. Select the first FCoE Interface (1/15)
5. For User Label, enter the storage controller name and port. Click Save Changes and OK.
6. Use the pulldown to select VSAN VSAN-A (101). Click Save Changes and OK.



7. Select the first FCoE Interface (1/16)
8. For User Label, enter the storage controller name and port. Click Save Changes and OK.
9. Use the pulldown to select VSAN VSAN-A (101). Click Save Changes and OK.
10. Expand Fabric B and Storage FCoE Interfaces.
11. Select the first FCoE Interface (1/15)
12. For User Label, enter the storage controller name and port. Click Save Changes and OK.
13. Use the pulldown to select VSAN VSAN-B (102). Click Save Changes and OK.
14. Select the first FCoE Interface (1/16)
15. For User Label, enter the storage controller name and port. Click Save Changes and OK.
16. Use the pulldown to select VSAN VSAN-B (102). Click Save Changes and OK.

Create vHBA Templates

To create the necessary virtual host bus adapter (vHBA) templates for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter vHBA-Template-A as the vHBA template name.
6. Keep Fabric A selected.

7. Leave Redundancy Type set to No Redundancy.
8. Select VSAN-A.
9. Leave Initial Template as the Template Type.
10. Select WWPN-POOL-A as the WWPN Pool.
11. Click OK to create the vHBA template.
12. Click OK

Create vHBA Template

Name : vHBA-Template-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : VSAN-A

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

WWPN Pool : WWPN-POOL-A(16/16)

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

13. Right-click vHBA Templates.
14. Select Create vHBA Template.
15. Enter vHBA-Template-B as the vHBA template name.
16. Leave Redundancy Type set to No Redundancy.
17. Select Fabric B as the Fabric ID.
18. Select VSAN-B.
19. Leave Initial Template as the Template Type.

20. Select WWPN-POOL-B as the WWPN Pool.
21. Click OK to create the vHBA template.
22. Click OK.

Create SAN Connectivity Policy

To configure the necessary Infrastructure SAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select SAN > Policies > root.
3. Right-click SAN Connectivity Policies.
4. Select Create SAN Connectivity Policy.
5. Enter `FC-BOOT` as the name of the policy.
6. Select the previously created WWNN-POOL for the WWNN Assignment.
7. Click the Add button at the bottom to add a vHBA.
8. In the Create vHBA dialog box, enter FABRIC-A as the name of the vHBA.
9. Select the Use vHBA Template checkbox.
10. In the vHBA Template list, select vHBA-Template-A.
11. In the Adapter Policy list, select WindowsBoot.

Create vHBA [?] [X]

Name : FABRIC-A

Use vHBA Template :

Redundancy Pair :

vHBA Template : vHBA-Template-A ▼

Adapter Policy : WindowsBoot ▼

Peer Name : []

Create vHBA Template

Create Fibre Channel Adapter Policy

OK Cancel

12. Click OK.
13. Click the Add button at the bottom to add a second vHBA.
14. In the Create vHBA dialog box, enter FABRIC-B as the name of the vHBA.
15. Select the Use vHBA Template checkbox.
16. In the vHBA Template list, select vHBA-Template-B.
17. In the Adapter Policy list, select WindowsBoot.
18. Click OK.

Create SAN Connectivity Policy ? X

Name :

Description :

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

World Wide Node Name

WWNN Assignment:

[Create WWNN Pool](#)

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

Name	WWPN
▶ vHBA FABRIC-B	Derived
▶ vHBA FABRIC-A	Derived

🗑️ Delete ➕ Add ⚙️ Modify

OK
Cancel

19. Click OK to create the SAN Connectivity Policy.

20. Click OK to confirm creation.

Create MAC Address Pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
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 **Sequential** as the option for Assignment Order.

8. Click Next.
9. Click Add.
10. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the of also embedding the UCS domain number information giving us 00:25:B5:52:0A:00 as our first MAC address.

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

Create a Block of MAC Addresses

First MAC Address : Size :

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

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



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the UCS domain number information giving us 00:25:B5:52:0B:00 as our first MAC address.

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
24. Click OK.
25. Click Finish.
26. 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, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-POOL` as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select **Sequential** for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources.
13. Click OK.
14. Click Finish.
15. Click OK.

Create Server Pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `MS-Server-Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the Hyper-V management cluster and click >> to add them to the `MS-Server-Pool` server pool.
9. Click Finish.
10. Click OK.

Create VLANs

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

1. In Cisco UCS Manager, click LAN on the left.



In this procedure, five unique VLANs are created. See Table 1.

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

Create VLANs ? X

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

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select `Set as Native VLAN`.
11. Click `Yes`, and then click `OK`.
12. Right-click `VLANs`.
13. Select `Create VLANs`
14. Enter `MS-IB-MGMT` as the name of the VLAN to be used for management traffic.
15. Keep the `Common/Global` option selected for the scope of the VLAN.
16. Enter the In-Band management VLAN ID.
17. Keep the `Sharing Type` as `None`.
18. Click `OK`, and then click `OK` again.
19. Right-click `VLANs`.
20. Select `Create VLANs`.

21. Enter `MS-SMB-1` as the name of the VLAN to be used for SMB File share.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the SMB File Share VLAN ID.
24. Keep the Sharing Type as None.
25. Click OK, and then click OK again.
26. Right-click VLANs.
27. Select Create VLANs.
28. Enter `MS-SMB-2` as the name of the VLAN to be used for 2nd SMB File share.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the Infrastructure SMB File Share VLAN ID.
31. Keep the Sharing Type as None.
32. Click OK, and then click OK again
33. Right-click VLANs.
34. Select Create VLANs.
35. Enter `MS-LVMN` as the name of the VLAN to be used for Live Migration.
36. Keep the Common/Global option selected for the scope of the VLAN.
37. Enter the Live Migration VLAN ID.
38. Keep the Sharing Type as None.
39. Click OK, and then click OK again.
40. Select Create VLANs.
41. Enter `MS-Cluster` as the name of the VLAN to be used for Cluster communication network.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the Cluster network VLAN ID.
44. Keep the Sharing Type as None.
45. Click OK, and then click OK again.
46. Select Create VLANs.
47. Enter `MS-Tenant-VM` as the name of the VLAN to be used for VM Traffic.

48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VM-Traffic VLAN ID.
50. Keep the Sharing Type as None.
51. Click OK, and then click OK again.

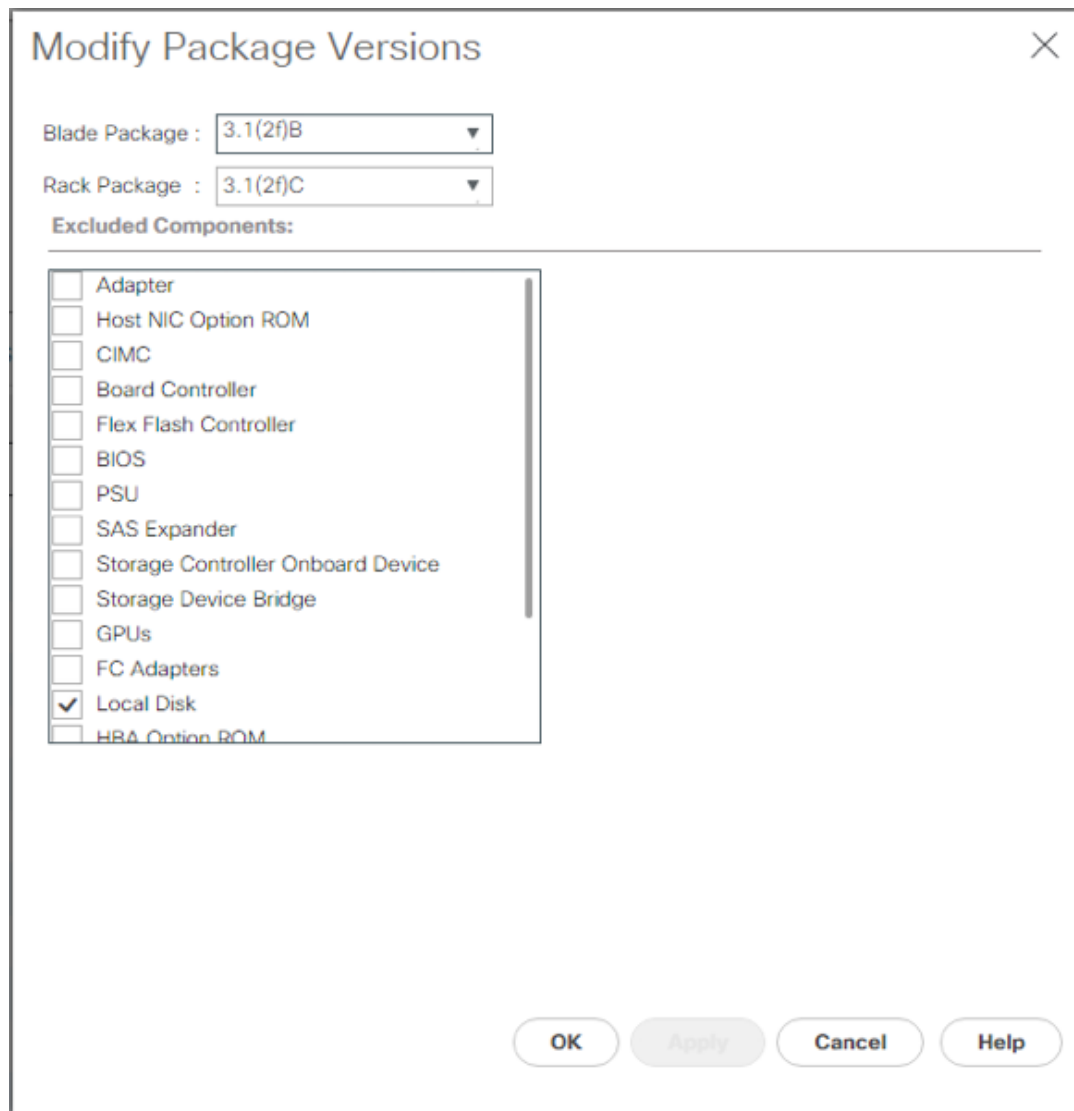
Name	ID	Type	Transport	Native	VLAN Sharing	Primary VLAN Na...	Multicast Policy N...
VLAN default (1)	1	Lan	Ether	No	None		
VLAN Native-VLAN (2)	2	Lan	Ether	Yes	None		
VLAN MS-IB-MGMT (90...	904	Lan	Ether	No	None		
VLAN MS-CSV (905)	905	Lan	Ether	No	None		
VLAN MS-LVMN (906)	906	Lan	Ether	No	None		
VLAN MS-Cluster (907)	907	Lan	Ether	No	None		
VLAN MS-Tenant-VM (...)	908	Lan	Ether	No	None		
VLAN MS-SMB-1 (3052)	3052	Lan	Ether	No	None		
VLAN MS-SMB-2 (3053)	3053	Lan	Ether	No	None		

Modify Default 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 properties.

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.1(2f) for both the Blade and Rack Packages.



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

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
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.
6. Click OK

LAN / LAN Cloud / QoS System Class

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

Create Local Disk Configuration Policy (Optional)

A local disk configuration for the Cisco UCS environment is necessary if the servers in the environment do not have a local disk.



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

To create a local disk configuration policy, complete the following steps:

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

Create Local Disk Configuration Policy

Name : SAN-Boot

Description :

Mode : No Local Storage

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

OK Cancel

8. Click OK.

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

To create a network control policy that enables CDP and LLDP on virtual network ports, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter `Enable-CDP-LLDP` as the policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

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**

9. Click OK.

Create Power Control Policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter `No-Power-Cap` as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy.
8. Click OK.

Create Power Control Policy ? X

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.

Create Server Pool Qualification Policy (Optional)

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



This example creates a policy for Cisco UCS B-Series and Cisco UCS C-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Name the policy UCS-Broadwell.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter UCS-CPU-E52660E as the PID.
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy then OK for the confirmation.

Create CPU/Cores Qualifications

Processor Architecture : PID (RegEx) :

Min Number of Cores : Unspecified select Max Number of Cores : Unspecified select

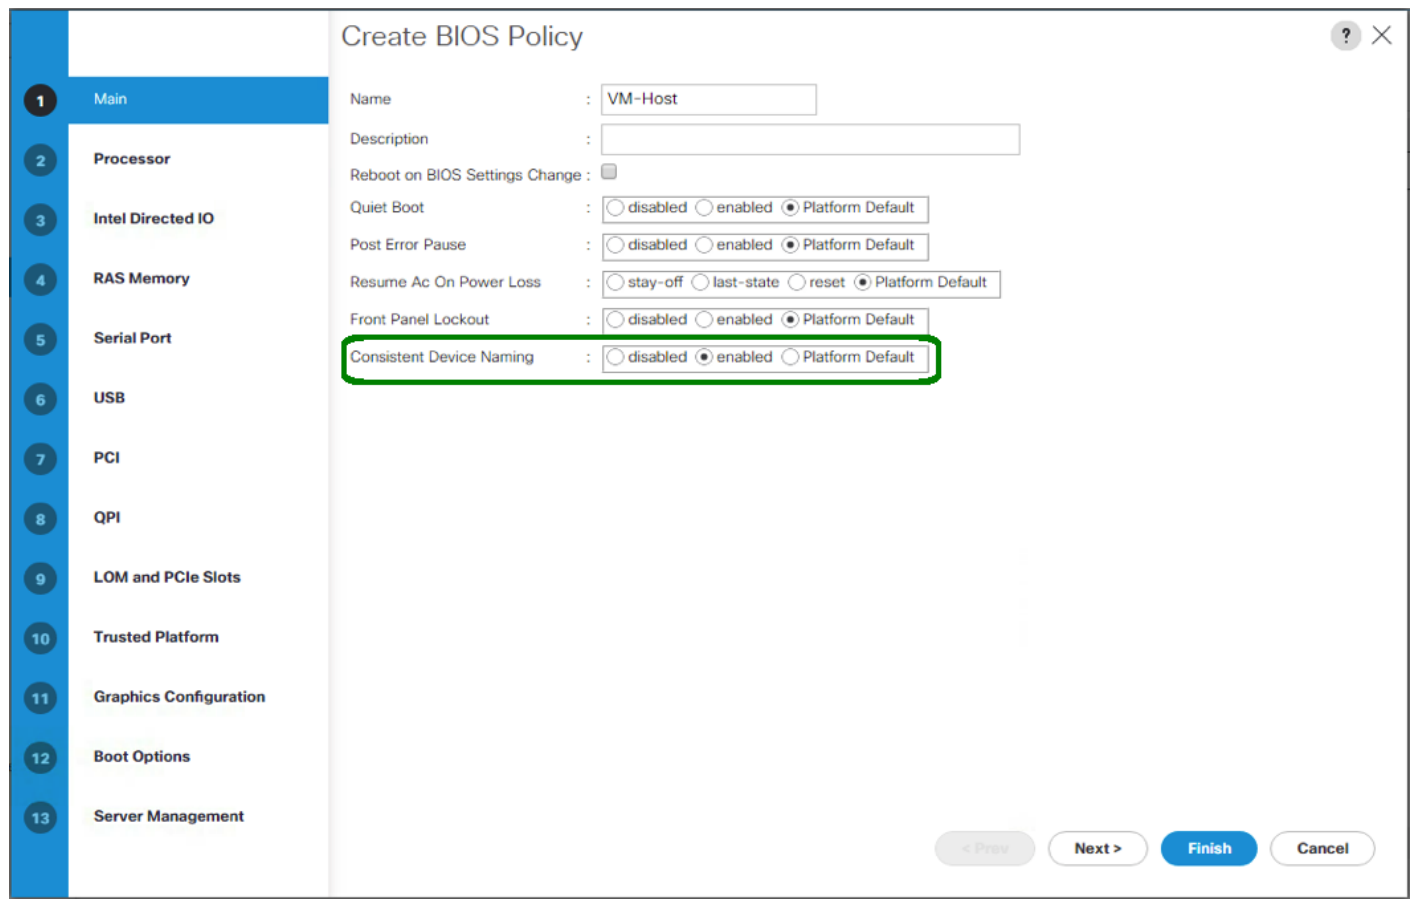
Min Number of Threads : Unspecified select Max Number of Threads : Unspecified select

CPU Speed (MHz) : Unspecified select CPU Stepping : Unspecified select

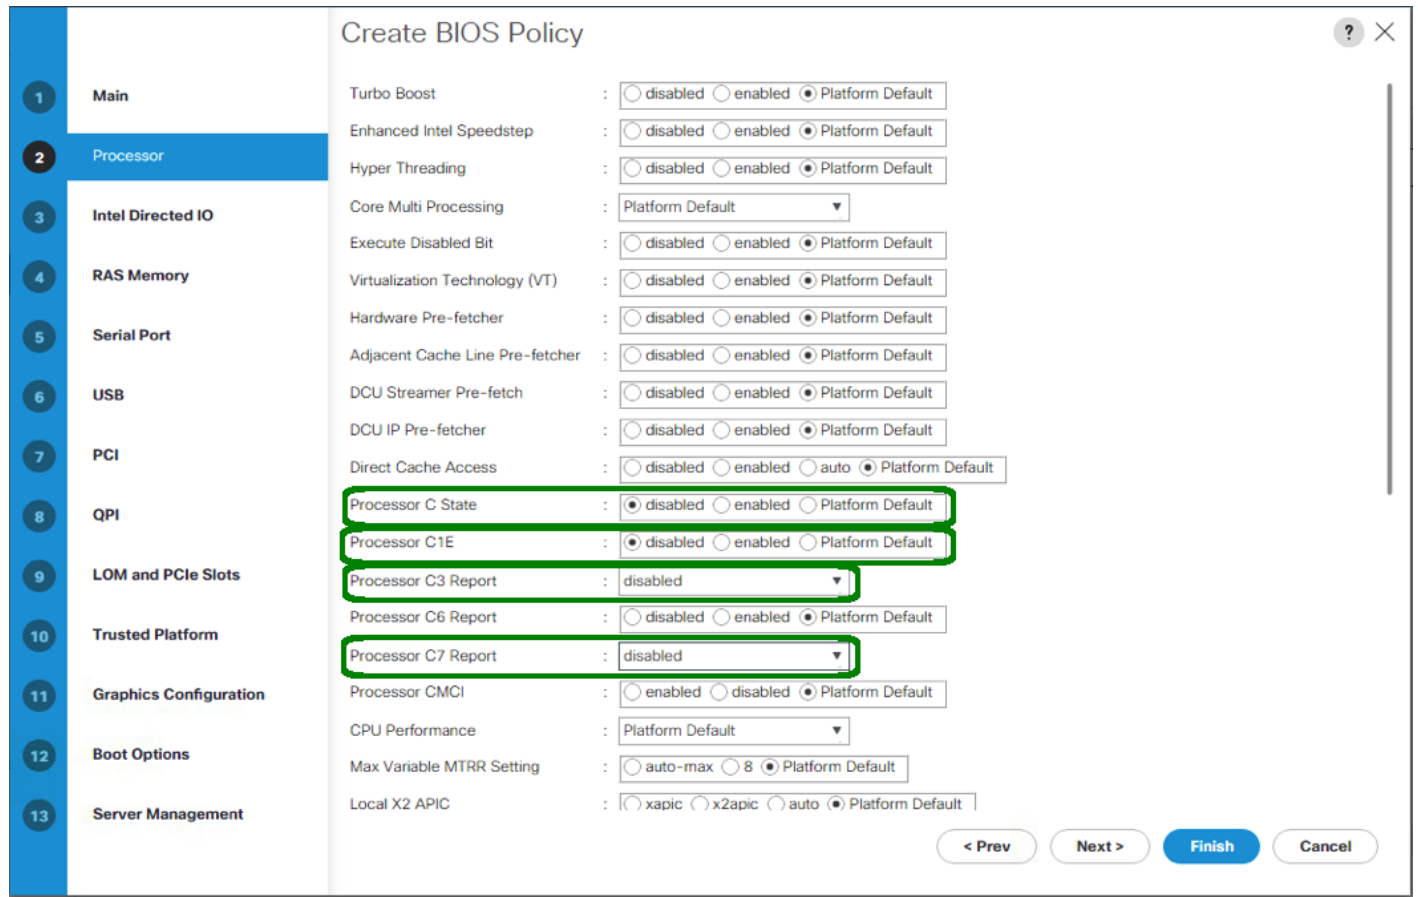
Create Server BIOS Policy

To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `MS-Host` as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



8. Click on the Processor tab on the left.
9. Set the following within the Processor tab
10. Processor C State -> disabled
11. Processor C1E -> disabled
12. Processor C3 Report -> disabled
13. Processor C7 Report -> disabled

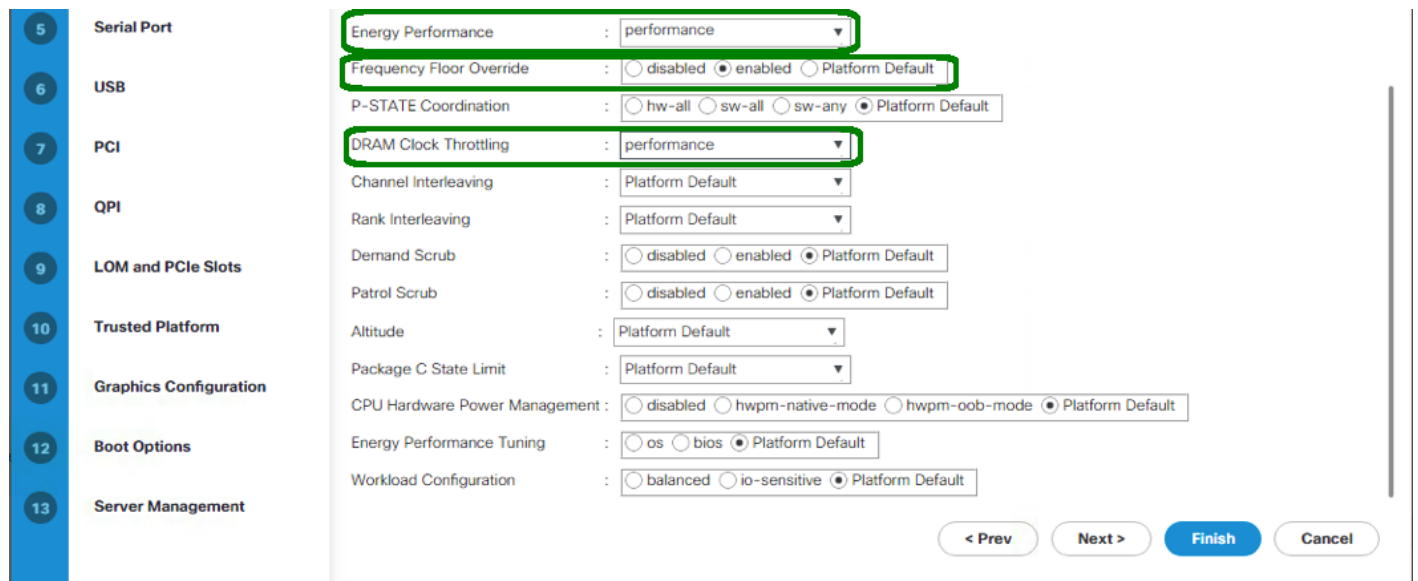


14. Scroll down to the remaining Processor options, and select:

15. Energy Performance -> performance

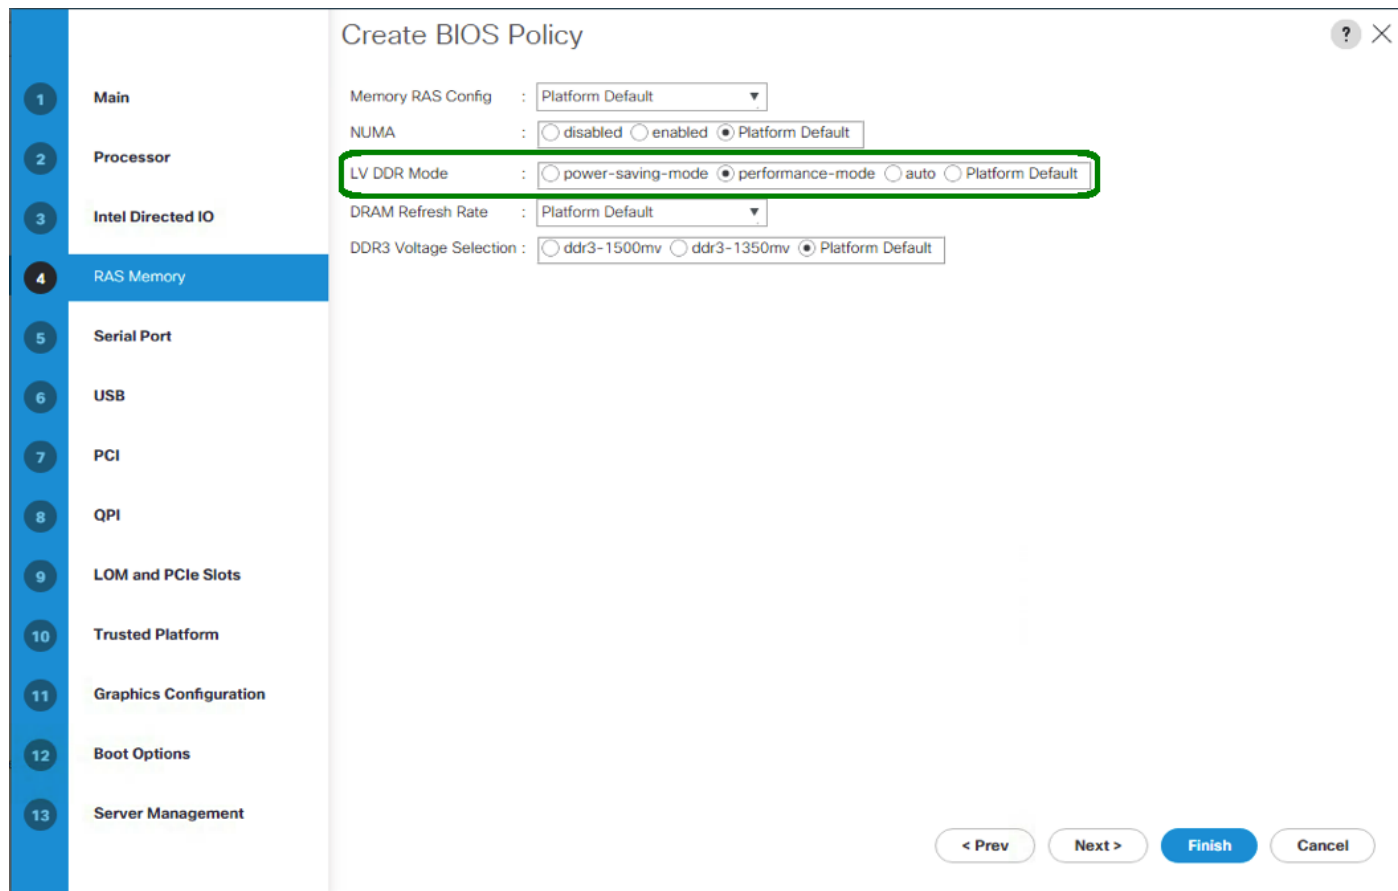
16. Frequency Floor Override -> enabled

17. DRAM Clock Throttling -> performance



18. Click on the RAS Memory option, and select:

19. LV DDR Mode -> performance-mode



20. Click Finish to create the BIOS policy.

21. Click OK.

Update the Default Maintenance Policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select "On Next Boot" to delegate maintenance windows to server administrators.

Servers / Policies / root / Maintenance Poli... / default

General	Events
Actions Delete Show Policy Usage Use Global	Properties Name : default Description : <input type="text"/> Owner : Local Soft Shutdown Timer : <input type="text" value="150 Secs"/> Reboot Policy : <input type="radio"/> Immediate <input checked="" type="radio"/> User Ack <input type="radio"/> Timer Automatic <input checked="" type="checkbox"/> On Next Boot (Apply pending changes at next reboot.)

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

Create vNIC Templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps. A total of 2 vNIC Templates will be created.

Create Infrastructure vNICs

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Host-A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select Primary Template for Redundancy Type.
9. Leave the Peer Redundancy Template set to <not set>.
10. Under Target, make sure that only the Adapter checkbox is selected.
11. Select Updating Template as the Template Type.
12. Under VLANs, select the checkboxes for MS-IB-MGMT, MS-Cluster, MS-CSV, MS-SMB1, MS-SMB2, and MS-Tenant-VM VLANs.

13. Set Native-VLAN as the native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-POOL-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template



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

Advanced Filter Export Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	MS-Cluster	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-CSV	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-A	<input type="radio"/>
<input type="checkbox"/>	MS-iSCSI-B	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-LVMN	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-SMB-1	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-SMB-2	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Tenant-VM	<input type="radio"/>

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 :

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template Infra-B:

1. Select LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter `Host-B` as the vNIC template name.
6. Select Fabric B.
7. Do not elect the Enable Failover checkbox.
8. Set Redundancy Type to Secondary Template.
9. Select Infra-A for the Peer Redundancy Template.
10. In the MAC Pool list, select `MAC-POOL-B`. The MAC Pool is all that needs to be selected for the Secondary Template.
11. Click OK to create the vNIC template.
12. Click OK.

Create LAN Connectivity Policy for FC Boot

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following 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.
5. Enter `FC-Boot` as the name of the policy.
6. Click the upper Add button to add a vNIC.
7. In the Create vNIC dialog box, enter `00-Host-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Host-A.
10. In the Adapter Policy list, select Windows.
11. Click OK to add this vNIC to the policy.

Create vNIC

Name : 00-Host-A

Use vNIC Template :

Redundancy Pair :

vNIC Template : Host-A

Peer Name :

Create vNIC Template

Adapter Performance Profile

Adapter Policy : Windows

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-Host-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select Host-B.
16. In the Adapter Policy list, select Windows.
17. Click OK to add the vNIC to the policy.

Name : **FC-Boot**

Description:

Owner : **Local**

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

Name	MAC Address	Native VLAN
▶ vNIC 00-Host-A	Derived	
▶ vNIC 01-Host-B	Derived	

Delete + Add Modify

+ Add iSCSI vNICs

18. Click OK, then OK again to create the LAN Connectivity Policy.

Create Boot Policy (FCoE Boot)

This procedure applies to a Cisco UCS environment in which two FCoE logical interfaces (LIFs) are on cluster node 1 (fcp_lifo3a_6332 and fcp_lifo3b_6332) and two FCoE LIFs are on cluster node 2 (fcp_lifo4a_6332 and fcp_lifo4b_6332).

To create a boot policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `FCoE-Boot` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.



Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Local CD/DVD`.
9. Expand the vHBAs drop-down menu and select `Add SAN Boot`.

Create Boot Policy

Name : FCoE-Boot

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

vNICs

vHBAs

Add SAN Boot

Add SAN Boot Target

Boot Order

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

10. Select the Primary for type field.
11. Enter FABRIC-A in vHBA field.

Add SAN Boot

vHBA : Fabric-A

Type : Primary Secondary Any

OK Cancel

12. Click OK.
13. From the vHBA drop-down menu, select Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.

15. Enter the WWPN for fcp_lifo3a_6332



To obtain this information, log in to the storage cluster and run the network interface show command

16. Select Primary for the SAN boot target type.

Add SAN Boot Target ? X

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK Cancel

17. Click OK to add the SAN boot target.
18. From the vHBA drop-down menu, select Add SAN Boot Target.
19. Enter 0 as the value for Boot Target LUN.
20. Enter the WWPN for fcp_lifo4a_6332.

Add SAN Boot Target ? X

Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

21. Click OK to add the SAN boot target.
22. From the vHBA drop-down menu, select Add SAN Boot.
23. In the Add SAN Boot dialog box, enter FABRIC-B in the vHBA box.
24. The SAN boot type should automatically be set to Secondary.
25. Click OK to add the SAN boot.
26. From the vHBA drop-down menu, select Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for fcp_lifo3b_6332.
29. Select Primary for the SAN boot target type.
30. Click OK to add the SAN boot target.
31. From the vHBA drop-down menu, select Add SAN Boot Target.
32. Keep 0 as the value for Boot Target LUN.
33. Enter the WWPN for fcp_lifo4b_6332.
34. Click OK to add the SAN boot target.

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.

Boot Order

+ - Advanced Filter Export Print ⚙

Name	vNIC/vH...	Type	WWN	LUN...	S...	B...	B...	D...
▼ SAN Primary	Fabric-A	Primary						
SAN Target Prim...		Primary	20:1C:00:A0:98:A9:FE:D2	0				
SAN Target Seco...		Second...	20:1E:00:A0:98:A9:FE:D2	0				
▼ SAN Secondary	Fabric-B	Second...						
SAN Target Prim...		Primary	20:1D:00:A0:98:A9:FE:D2	0				
SAN Target Seco...		Second...	20:1F:00:A0:98:A9:FE:D2	0				

↑ Move Up ↓ Move Down 🗑 Delete

35. Click OK, then click OK again to create the boot policy.

Create Service Profile Templates

In this procedure, one service profile template for Infrastructure Hyper-V hosts is created for Fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
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 `Hyper-V-Host-FC` 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 ? 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.

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

Name :

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

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

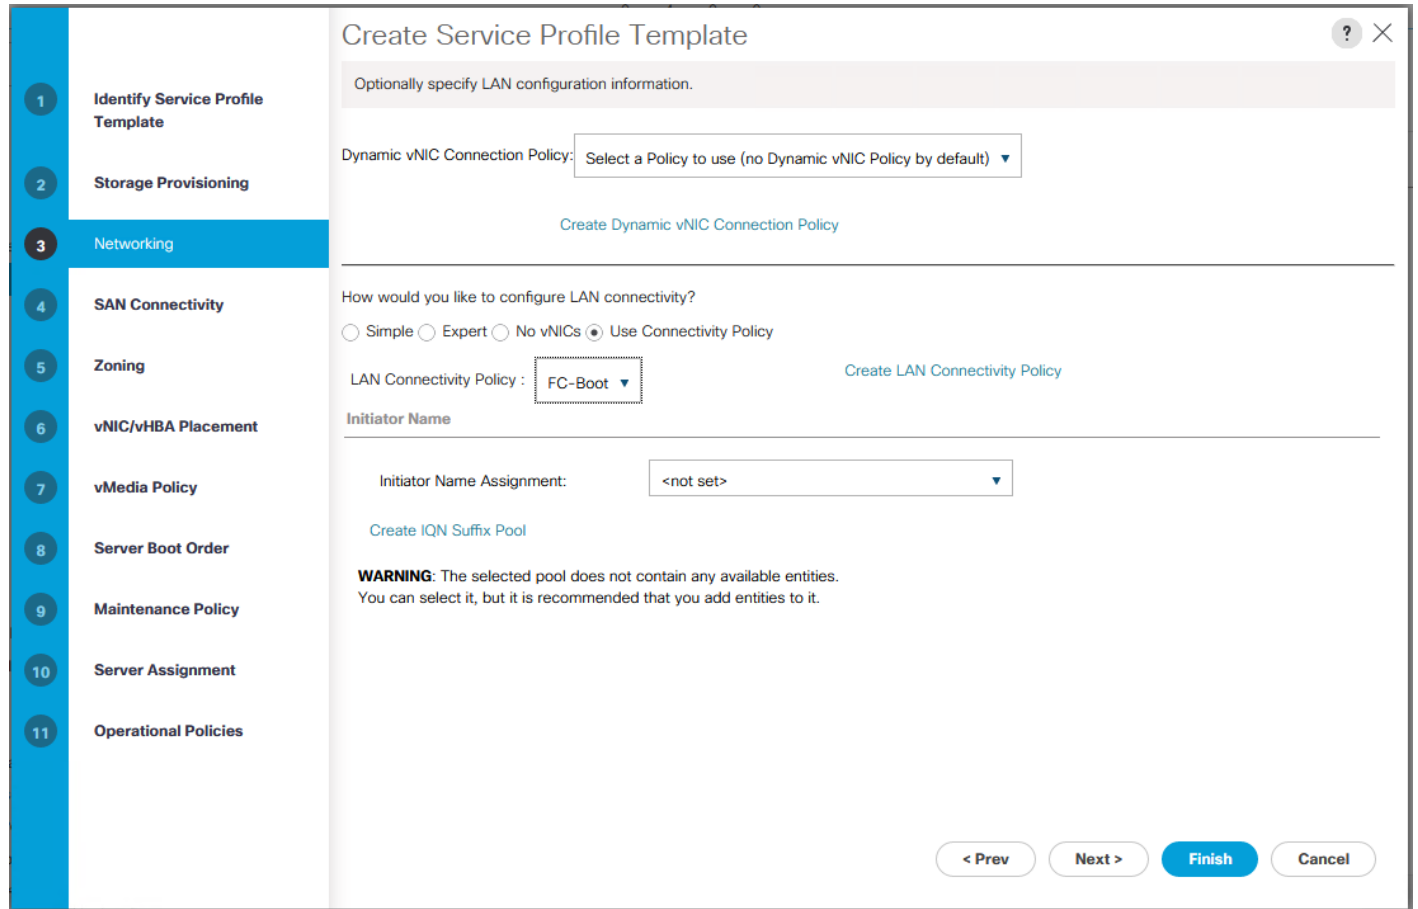
8. Click Next.

Configure Storage Provisioning

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

Configure Networking Options

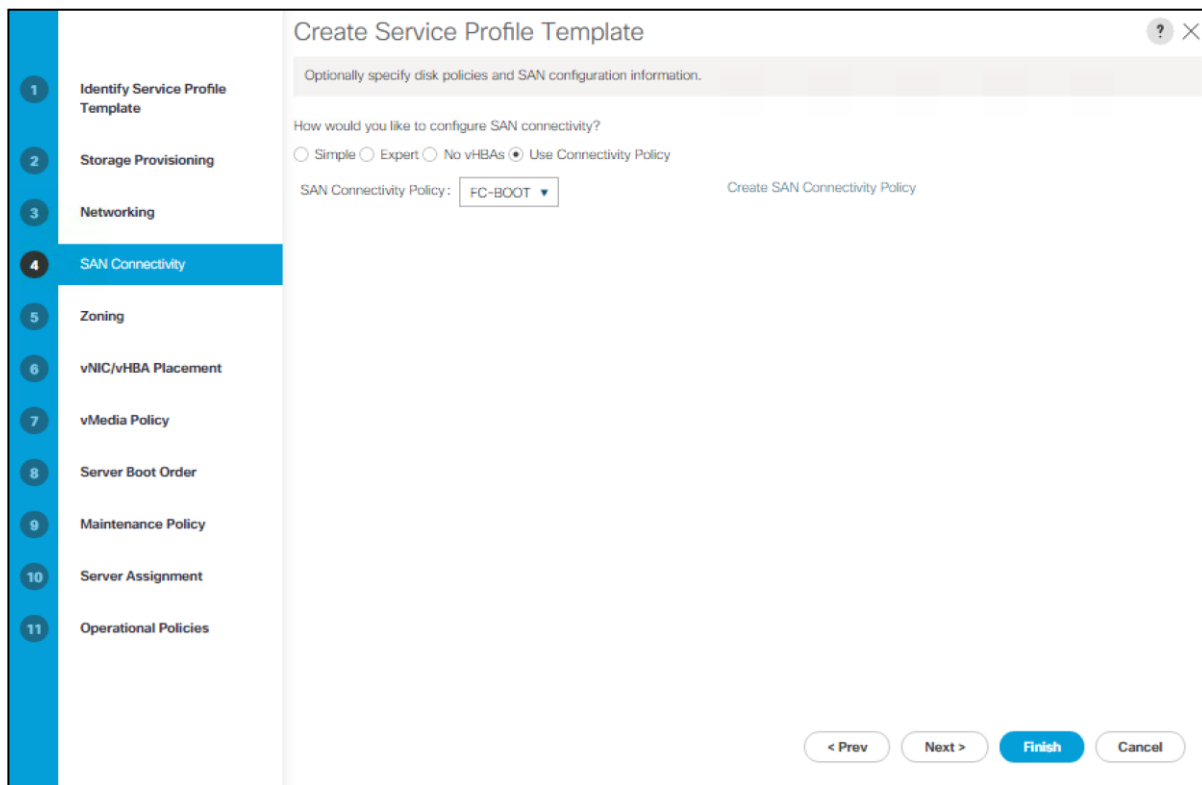
1. Keep the setting at default for Dynamic vNIC Connection Policy.
2. Select the "Use Connectivity Policy" option to configure the LAN connectivity.
3. Select FC-Boot from the LAN Connectivity Policy pull-down.
4. Leave Initiator Name Assignment at <not set>.



5. Click Next.

Configure Storage Options

1. Select the Use Connectivity Policy option for the "How would you like to configure SAN connectivity?" field.
2. Select the FC-BOOT option from the SAN Connectivity Policy pull-down.



3. Click Next.

Configure Zoning Options

1. Set no Zoning options 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 FCoE-Boot for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: [Create Boot Policy](#)

Name : **FCoE-Boot**
 Description :
 Reboot on Boot Order Change : **Yes**
 Enforce vNIC/vHBA/iSCSI Name : **Yes**
 Boot Mode : **Legacy**

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

Boot Order

Name	Order	vNIC/vHBA..	Type	WWN	LUN Name	Slot Numb...	Boot Name	Boot Path	Description
Local C...	1								
▼ San	2								
▶ SAN..		FABRIC-A	Primary						
▶ SAN..		FABRIC-B	Secondary						

2. Click Next.

Configure Maintenance Policy

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
Reboot Policy	: User Ack

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

2. Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

1. In the Pool Assignment list, select `MS-Server-Pool`.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Optional: select "UCS-Broadwell" for the Server Pool Qualification.
4. Expand Firmware Management at the bottom of the page and select the default policy

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify a server pool for this service profile template.

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

Pool Assignment: [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](#)

5. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template ? ×

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

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

BIOS Policy:

+ External IPMI 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

< Prev Next > **Finish** Cancel

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

Create Service Profiles

To create service profiles from the service profile template, complete the following steps:

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template Hyper-V-Host-FC.
3. Right-click Hyper-V-Host-FC and select Create Service Profiles from Template.
4. Enter `Hyper-V-Host-0` as the service profile prefix.
5. Enter 1 as "Name Suffix Starting Number."
6. Enter 2 as the "Number of Instances."
7. Click OK to create the service profiles.

Create Service Profiles From Template ? X

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

8. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

Additional server pools, service profile templates, and service profiles can be created in the respective organizations to add more servers to the FlexPod unit. All other pools and policies are at the root level and can be shared among the organizations.

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure server in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS server and from the NetApp controllers. Insert the required information into Table 6 and Table 7 below.

Table 1 WWPNs from NetApp storage

SVM	Adapter	MDS Switch	Target: WWPN
Infra-MS-SVM	fcp_lifo3a_6332	Fabric A	<fcp_lifo3a_6332-wwpn>
	fcp_lifo3b_6332	Fabric B	<fcp_lifo3b_6332-wwpn>
	fcp_lifo4a_6332	Fabric A	<fcp_lifo4a_6332-wwpn>
	fcp_lifo4b_6332	Fabric B	<fcp_lifo4b_6332-wwpn>



To obtain the FC WWPNs, run the `network interface show` command on the storage cluster management interface.

Table 2 WWPNs for UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
--------------------------------	------------	----------------

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
Hyper-V-Host-01	Fabric A	Hyper-V-Host -01-wwpna
	Fabric B	Hyper-V-Host -01-wwpnb
Hyper-V-Host -02	Fabric A	Hyper-V-Host -02-wwpna
	Fabric B	Hyper-V-Host -02-wwpnb



To obtain the FC vHBA WWPN information in Cisco UCS Manager GUI, go to Servers > Service Profiles > root. Click each service profile and then click the "Storage" tab, then "vHBAs" tab on the right. The WWPNs are displayed in the table at the bottom of the page.

Adding Direct Connected Tenant FC Storage

To add FC storage from an additional storage SVM, two storage connection policies, one for each fabric must be added in UCS Manager and attached to vHBA Initiator Groups in the SAN Connectivity Policy. These steps were not shown in the initial deployment above because it is not necessary to zone boot targets. Boot targets are automatically zoned in the fabric interconnect when zoning is enabled on the fabric VSAN. To add direct connected tenant FC storage from a tenant SVM, complete the following steps:

Create Storage Connection Policies

In this procedure, one storage connection policy is created for each fabric.

To create the storage connection policies, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Right-click SAN > Policies > root > Storage Connection Policies and select Create Storage Connection Policy.
3. Name the policy to indicate a tenant on Fabric A.
4. Select the Single Initiator Multiple Targets Zoning Type.
5. Click Add to add a target.
6. Enter the WWPN of the first fabric A FC LIF in the tenant SVM connected to fabric interconnect A. Select Path A and VSAN VSAN-A. Click OK.
7. Click Add to add a target.
8. Enter the WWPN of the second fabric A FC LIF in the tenant SVM connected to fabric interconnect A. Select Path A and VSAN VSAN-A. Click OK.
9. Click OK then OK again to complete adding the Storage Connection Policy.
10. Right-click SAN > Policies > root > Storage Connection Policies and select Create Storage Connection Policy.
11. Name the policy to indicate a tenant on Fabric B.
12. Select the Single Initiator Multiple Targets Zoning Type.

13. Click Add to add a target.
14. Enter the WWPN of the first fabric B FC LIF in the tenant SVM connected to fabric interconnect B. Select Path B and VSAN VSAN-B. Click OK.
15. Click Add to add a target.
16. Enter the WWPN of the second fabric B FC LIF in the tenant SVM connected to fabric interconnect B. Select Path B and VSAN VSAN-B. Click OK.
17. Click OK then OK again to complete adding the Storage Connection Policy.

Map Storage Connection Policies vHBA Initiator Groups in SAN Connectivity Policy

In this procedure, storage connection policies are mapped to vHBA initiator groups for each fabric.

To create the storage connection policy mappings, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select SAN > Policies > root > SAN Connectivity Policies > FC-Boot.
3. In the center pane, select the vHBA Initiator Groups tab.
4. Click Add to add a vHBA Initiator Group.
5. Name the group Fabric A and select the Fabric A Initiator.
6. Use the pulldown to select the Fabric A Storage Connection Policy.
7. Click OK and OK to complete adding the Initiator Group.
8. Click Add to add a vHBA Initiator Group.
9. Name the group Fabric B and select the Fabric B Initiator.
10. Use the pulldown to select the Fabric B Storage Connection Policy.
11. Click OK and OK to complete adding the Initiator Group.

Microsoft Windows Server 2016 Hyper-V Deployment Procedure

Setup the Microsoft Windows 2016 install

The Microsoft Windows 2016 install sub-sections listed below and their installation procedures follow the installation guidelines covered in the main document.

Install Windows Server 2016

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Install Chipset and Windows eNIC Drivers

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Install Windows Roles and Features

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Install NetApp Host Utilities

To complete this section, refer to corresponding section of the main document and execute all the steps.

Host Renaming and Join to Domain

To complete this section, refer to the corresponding section of the main document and execute all the steps.

Deploying and Managing Hyper-V Clusters using System Center 2016 VMM

To complete this section, refer to the corresponding section of the main document and execute all the steps.

FlexPod Backups

Cisco UCS Backup

Automated backup of the UCS domain is important for recovery of the Cisco UCS Domain from issues ranging catastrophic failure to human error. There is a native backup solution within UCS that allows local or remote backup using FTP/TFTP/SCP/SFTP as options.

Created backups can be a binary file containing the Full State, which can be used for a restore to the original or a replacement pair of Cisco UCS fabric interconnects. Alternately this XML configuration file consists of All configurations, just System configurations, or just Logical configurations of the UCS Domain. For scheduled backups, the available options are Full State or All Configuration, backup of just the System or Logical configurations can be manually initiated.

Specification of the backup can be done by following these steps within the UCSM GUI:

1. Select Admin within the Navigation pane and select All.
2. Click on the Policy Backup and 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]

4. Click Save Changes to create the Policy.

Cisco Nexus Backups

The configuration of the Cisco Nexus 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 an automated configuration backup for one of the FlexPod 9332PQ switches is shown below:

```
bb04-9332-a# conf t
```

Enter configuration commands, one per line. End with CNTL/Z.

```
bb04-9332-a(config)# feature scheduler
```

```
bb04-9332-a(config)# scheduler logfile size 1024
```

```
bb04-9332-a(config)# scheduler job name backup-cfg
```

```
bb04-9332-a(config-job)# copy running-config
```

```
tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management
```

```
bb04-9332-a(config-job)# exit
```



```
bb04-9332-a(config)# scheduler schedule name daily
bb04-9332-a(config-schedule)# job name backup-cfg
bb04-9332-a(config-schedule)# time daily 2:00
bb04-9332-a(config-schedule)# end
```

Show the job that has been setup:

```
bb04-9332-a# sh scheduler job
```

```
Job Name: backup-cfg
```

```
-----
```

```
copy running-config tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP)
vrf management
```

```
=====
```

```
bb04-9332-a# show scheduler schedule
```

```
Schedule Name      : daily
```

```
-----
```

```
User Name          : admin
```

```
Schedule Type      : Run every day at 2 Hrs 0 Mins
```

```
Last Execution Time : Sun Apr 9 02:00:00 2017
```

```
Last Completion Time: Sun Apr 9 02:00:01 2017
```

```
Execution count    : 3
```

```
-----
```

```
Job Name           Last Execution Status
```

```
-----
```

```
backup-cfg        Success (0)
```

```
=====
```

Full documentation for the feature scheduler can be found at:

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

Breakout Interface Configuration in the Nexus 9332PQ Switches

The 40Gb end to end FlexPod design in this document uses a pair of Nexus 9332PQ which is built with 40 Gbps Quad Small Form Factor Pluggable Plus (QSFP+) type on all ports. If there is a need to directly support a 10Gb Small Form Pluggable Plus (SFP+), this can be configured within the switch, and connected to the 10Gb SFP+ device using a supported QSFP+ Breakout Cable.

Configuration of the QSFP+ ports will use the interface breakout command as shown in this example to turn the 40Gb interface Ethernet 1/1 into 4x10Gb interfaces:

```
bb04-9332-a(config)# show running-config interface Ethernet1/1

interface Ethernet1/1

no switchport

bb04-9332-a(config)# interface breakout module 1 port 5 map 10g-4x

bb04-9332-a(config)# show running-config interface Ethernet1/1/1-4

interface Ethernet1/1/1

interface Ethernet1/1/2

interface Ethernet1/1/3

interface Ethernet1/1/4
```

Breakout configurations that are no longer needed can be reverted with the no interface breakout command:

```
bb04-9332-a(config)# no interface breakout module 1 port 1 map 10g-4x
```

About the Authors

Rajendra Yogendra, Technical Marketing Engineer, Data Center Solutions Engineering, Cisco Systems Inc.

Rajendra has over 9 years of experience in IT Infrastructure, Server Virtualization, and Cloud Computing. His current role includes building Infrastructure solutions, software defined storage solutions, and performance benchmarking on Cisco UCS platforms.

Aaron Kirk, Technical Marketing Engineer, Converged Infrastructure Engineering, NetApp Inc.

Aaron Kirk is a Technical Marketing Engineer in the NetApp Converged Infrastructure Engineering team. He focuses on producing validated reference architectures that promote the benefits of end-to-end data center solutions and cloud environments. Aaron has been at NetApp since 2010, previously working as an engineer on the MetroCluster product in the Data Protection Group. Aaron holds a Bachelor of Science in Computer Science and a Masters of Business Administration from North Carolina State University. [Acknowledgements](#)

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Suhas Rao
- Sanjeev Naldurgkar
- John George
- Chris Reno
- Karthick Radhakrishnan