

FlexPod Datacenter with Microsoft Hyper-V Windows Server 2016 and Cisco ACI 3.0

Deployment Guide for FlexPod Datacenter with Microsoft Hyper-V
Windows Server 2016, Cisco ACI 3.0, and NetApp AFF A-Series

Last Updated: May 28, 2018



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)

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	12
Solution Overview.....	13
Introduction	13
Audience	13
Purpose of this Document.....	13
What's New?	13
Solution Design.....	14
Architecture.....	14
Physical Topology.....	15
Deployment Hardware and Software	17
Software Revisions	17
Configuration Guidelines.....	17
Physical Infrastructure.....	19
FlexPod Cabling	19
Infrastructure Servers Prerequisites	21
Active Directory DC/DNS.....	21
Microsoft System Center 2016	21
Network Switch Configuration.....	22
Physical Connectivity	22
Cisco Application Policy Infrastructure Controller (APIC) Verification	22
Cisco ACI Fabric Discovery.....	24
Initial ACI Fabric Setup Verification	26
Software Upgrade	26
Setting Up Out-of-Band Management IP Addresses for New Leaf and Switches.....	27
Verifying Time Zone and NTP Server	28
Verifying Domain Name Servers	29
Verifying BGP Route Reflectors.....	29
Set Up Fabric Access Policy Setup	31
Create Link Level Policies	31
Create CDP Policy	32
Create LLDP Interface Policies.....	33
Create Port Channel Policy	34
Create BPDU Filter/Guard Policies	36

Create Global VLAN Policy	37
Create Firewall Policy	38
Create Virtual Port Channels (vPCs).....	40
VPC - Management Switch	40
VPC - UCS Fabric Interconnects.....	43
VPC - NetApp AFF Cluster.....	47
Configuring Common Tenant for Management Access	52
Create VRFs	53
Create Bridge Domains.....	54
Create Application Profile	55
Create EPG.....	56
Create Security Filters in Tenant Common	65
Deploy FP-Foundation Tenant	68
Create Bridge Domain	69
Create Application Profile for IB-Management Access	70
Create Application Profile for Host Connectivity.....	75
Storage Configuration.....	81
NetApp All Flash FAS A300 Controllers	81
NetApp Hardware Universe	81
Controllers.....	81
Disk Shelves	81
NetApp ONTAP 9.1.....	82
Complete Configuration Worksheet	82
Configure ONTAP Nodes	82
Login to the Cluster	91
Zero All Spare Disks	91
Set Onboard Unified Target Adapter 2 Port Personality	92
Set Auto-Revert on Cluster Management	92
Set Up Management Broadcast Domain	93
Set Up Service Processor Network Interface	93
Create Aggregates	93
Verify Storage Failover.....	94
Disable Flow Control on 10GbE and 40GbE Ports	95
Disable Unused FCoE Capability on CNA Ports.....	95
Configure Network Time Protocol	95

Configure Simple Network Management Protocol.....	96
Configure AutoSupport.....	96
Enable Cisco Discovery Protocol.....	96
Create Jumbo Frame MTU Broadcast Domains in ONTAP.....	96
Create Interface Groups.....	97
Create VLANs.....	97
Create Storage Virtual Machine.....	97
Create the CIFS Service.....	98
Modify Storage Virtual Machine Options.....	99
Create Load-Sharing Mirrors of SVM Root Volume.....	99
Create Block Protocol Service(s).....	99
Configure HTTPS Access.....	100
Create SMB Export Policy.....	101
Create NetApp FlexVol Volumes.....	101
Create CIFS Shares.....	102
Create Gold Management Host Boot LUN.....	102
Create Witness and iSCSI Datastore LUNs.....	102
Schedule Deduplication.....	102
Create SAN LIFs.....	103
Create SMB LIFs.....	103
Add Infrastructure SVM Administrator.....	104
Server Configuration.....	105
Cisco UCS Base Configuration.....	105
Perform Initial Setup.....	105
Cisco UCS Setup.....	107
Log in to Cisco UCS Manager.....	107
Upgrade Cisco UCS Manager Software to Version 3.2(1d).....	107
Anonymous Reporting.....	107
Configure Cisco UCS Call Home.....	108
Add Block of IP Addresses for KVM Access.....	108
Synchronize Cisco UCS to NTP.....	109
Edit Policy to Automatically Discover Server Ports.....	109
Edit Chassis Discovery Policy.....	110
Verify Server and Enable Uplink Ports.....	111
Acknowledge Cisco UCS Chassis and FEX.....	112

Re-Acknowledge Any Inaccessible C-Series Servers	113
Create Uplink Port Channels to Cisco Nexus 9332 Switches	113
Create an IQN Pool for iSCSI Boot	114
Create iSCSI Boot IP Address Pools	116
Create MAC Address Pools	119
Create UUID Suffix Pool	121
Create Server Pool	122
Create VLANs	122
Modify Default Host Firmware Package	126
Set Jumbo Frames in Cisco UCS Fabric	127
Create Local Disk Configuration Policy (Optional)	128
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	130
Create Power Control Policy	131
Create Server Pool Qualification Policy (Optional)	132
Create Server BIOS Policy	133
Update the Default Maintenance Policy	135
Create vNIC Templates	136
Create LAN Connectivity Policy for iSCSI Boot	142
Create iSCSI Boot Policies	147
Create iSCSI Boot Service Profile Templates	149
Create Multipath Service Profile Template	157
Create Service Profiles	162
Storage Configuration - Boot LUNs	163
NetApp ONTAP Boot Storage Setup	163
Create igroups	163
Map LUNs to igroups	163
Microsoft Windows Server 2016 Hyper-V Deployment Procedure	164
Setting Up Microsoft Windows Server 2016	164
Install Windows Server 2016	165
Host Renaming and Join to Domain	166
Install NetApp Windows Unified Host Utilities	167
Set Up Multipathing and iSCSI	169
Install NetApp SnapDrive 7.1.4 for Windows	170
Configuring Access for SnapDrive for Windows	170
Downloading SnapDrive 7.1.4 for Windows	171

Installing SnapDrive for Windows.....	171
Configure Server for Cloning	180
Clone and Remap Server LUNs for Sysprep Image	182
Boot and Set Up Sysprep Clone	182
Clone and Remap Server LUNs for Production Images	182
Boot and Set Up Clones	183
Install Roles and Features Required for Hyper-V.....	186
Set Up Hyper-V Networking	187
Build System Center Virtual Machine Manager (SCVMM) Virtual Machine (VM).....	188
Deploying and Managing the Management Hyper-V Cluster Using System Center 2016 VMM	197
Settings	197
Create Run As accounts in VMM.....	197
Fabric – Servers - I	197
Create Host Group.....	197
Add Hosts to the Host Group	198
Creating APIC-Controlled Hyper-V Networking	201
Create Windows Failover Cluster	215
Build a Windows Server 2016 Virtual Machine for Cloning	226
NetApp SMI-S Provider Configuration	227
NetApp SMI-S Integration with VMM	229
Build Windows Active Directory Servers for ACI Fabric Core Services	233
Build Microsoft Systems Center Operations Manager (SCOM) Server VM.....	234
Cisco UCS Management Pack Suite Installation and Configuration.....	235
Cisco UCS Manager Integration with SCOM	235
About Cisco UCS Management Pack Suite	235
Installing Cisco UCS Monitoring Service	235
Adding a Firewall Exception for the Cisco UCS Monitoring Service.....	237
Installing the Cisco UCS Management Pack Suite	237
Adding Cisco UCS Domains to the Operations Manager	240
Cisco UCS Manager Monitoring Dashboards	244
Cisco UCS Manager Plug-in for SCVMM	249
Cisco UCS Manager Plug-in Installation.....	249
Cisco UCS Domain Registration:.....	250
Using the Cisco UCS SCVMM Plugin	251
Viewing the Server Details from the Hypervisor Host View.....	251

Viewing Registered UCS Domains	252
Viewing the UCS Blade Server Details	253
Viewing the UCS Rack-Mount Server Details:	254
Viewing the Service Profile Details	256
Viewing the Service Profile Template Details	258
Viewing the Host Firmware Package Details	259
NetApp FlexPod Management Tools Setup	261
OnCommand Unified Manager 7.2	261
NetApp SnapManager for Hyper-V	263
Downloading SnapManager for Hyper-V	263
Installing SnapManager for Hyper-V	264
NetApp OnCommand Plug-in for Microsoft	267
Downloading OnCommand Plug-in for Microsoft	267
Installing NetApp OnCommand Plug-In for Microsoft	268
Storage Configuration - Boot LUNs for Tenant Hyper-V Hosts	272
NetApp ONTAP Boot Storage Setup	272
Create igroups	272
Map LUNs to igroups	272
Sample Tenant Setup	273
Add Supernet Routes to Core-Services Devices	273
Adding the Supernet Route in a Windows VM or Host	273
ACI Shared Layer 3 Out Setup	273
Configuring the Nexus 7000s for ACI Connectivity (Sample)	274
Configuring ACI Shared Layer 3 Out	277
Lab Validation Tenant Configuration	295
Deploy ACI Application (MS-TNT-A) Tenant	296
Configure Tenant Storage	329
Create Tenant IPspace	330
Create Tenant Broadcast Domains in ONTAP	330
Create VLAN Interfaces	330
Create Storage Virtual Machine	331
Setup SVM Management Access	331
Create the CIFS Service	332
Modify Storage Virtual Machine Options	332
Create Load-Sharing Mirrors of SVM Root Volume	332

Create Block Protocol Service(s)	333
Configure HTTPS Access	333
Create SMB Export Policy	334
Create NetApp FlexVol Volumes	334
Schedule Deduplication	335
Create SAN LIFs	335
Create SMB LIFs	335
Add Quality of Service (QoS) Policy to Monitor Application Workload	336
Configure Cisco UCS for the Tenant	336
Add Tenant Host Management vNIC Template	336
Create Tenant LAN Connectivity Policy for iSCSI Boot	338
Create Tenant Service Profile Template	341
Add New Application-Specific Server Pool	341
Create New Service Profiles for Tenant Servers	342
Configure Storage SAN Boot for the Tenant	342
Hyper-V Boot LUN in Infra-MS-SVM for First Tenant Host	342
Clustered Data ONTAP iSCSI Boot Storage Setup	343
Map Boot LUN to igroup	343
Microsoft Hyper-V Server Deployment Procedure for Tenant Hosts	344
Setting Up Microsoft Hyper-V Server 2016	344
Host Renaming and Join to Domain	346
Install NetApp Windows Unified Host Utilities	347
Set Up Multipathing and iSCSI	349
Install SnapDrive for Windows	350
Configure Server for Cloning	358
Clone and Remap Server LUNs for Sysprep Image	358
Boot and Set Up Sysprep Clone	359
Clone and Remap Server LUNs for Production Image	359
Boot and Set Up Clones	359
Deploying and Managing the Tenant Hyper-V Cluster Using System Center 2016 VMM	361
Fabric - Servers - I	361
Create Host Groups	361
Add Hosts to the Host Group	362
Fabric - Networking - Install APIC Hyper-V Agent and Add Host to SCVMM Virtual Switch	363
Set Up Hyper-V Networking	365

Add TNT iSCSI Sessions to Hosts.....	367
Create Windows Failover Cluster.....	368
Add Tenant iSCSI Datastores (Optional).....	378
Build a Second Tenant (Optional).....	383
Appendix - FC Solution.....	384
Storage Configuration.....	385
Set Onboard Unified Target Adapter 2 Port Personality.....	385
Add FCP Storage Protocol to Infrastructure SVM.....	385
Create FCP Storage Protocol in Infrastructure SVM.....	386
Create FC LIFs.....	386
Server Configuration.....	386
Configure FC Unified Ports (UP) on UCS Fabric Interconnects.....	386
Place Cisco UCS Fabric Interconnects in Fiber Channel Switching Mode.....	387
Create Storage VSANs.....	388
Configure FC Storage Ports.....	390
Assign VSANs to FC Storage Ports.....	390
Create a WWNN Pool for FC Boot.....	391
Create WWPN Pools.....	393
Create vHBA Templates.....	396
Create SAN Connectivity Policy.....	398
Create Server Pool.....	400
Create LAN Connectivity Policy for FC Boot.....	400
Create Boot Policy (FC Boot).....	402
Create Boot Policy (FC Boot) With a Single Path for Windows Installation.....	405
Create Service Profile Templates.....	408
Create Service Profiles.....	415
Add More Servers to FlexPod Unit.....	416
Gather Necessary Information.....	416
Adding Direct Connected Tenant FC Storage.....	417
Create Storage Connection Policies.....	417
Map Storage Connection Policies vHBA Initiator Groups in SAN Connectivity Policy.....	418
Create igroups.....	419
Map Boot LUNs to igroups.....	419
FlexPod Backups.....	420
Cisco UCS Backup.....	420

About the Authors.....	422
Acknowledgements	422



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 3.2(1d), Cisco Application Centric Infrastructure (ACI) 3.0(1k), and Microsoft Hyper-V 2016. Cisco UCS Manager (UCSM) 3.2 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, including Cisco UCS B200M5 servers. FlexPod Datacenter with Cisco UCS unified software release 3.2(1d), 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, Cisco Application Policy Infrastructure Controller (APIC), and NetApp All Flash FAS (AFF).

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



FC storage traffic does not flow through the ACI Fabric and is not covered by the ACI policy 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 ACI solution. This document primarily focuses on deploying Microsoft Hyper-V 2016 Cluster on FlexPod Datacenter using iSCSI and SMB storage protocols. The Appendix section covers the delta changes on the configuration steps using FC 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 3.2(1d) unified software release, Cisco UCS B200-M5 servers, Cisco UCS B200-M4 servers, and Cisco UCS C220-M4 servers
- Support for Cisco ACI version 3.0(1k)
- Support for the latest release of NetApp ONTAP® 9.1
- SMB, iSCSI, and FC storage design
- Validation of Microsoft Hyper-V 2016

Solution Design

Architecture

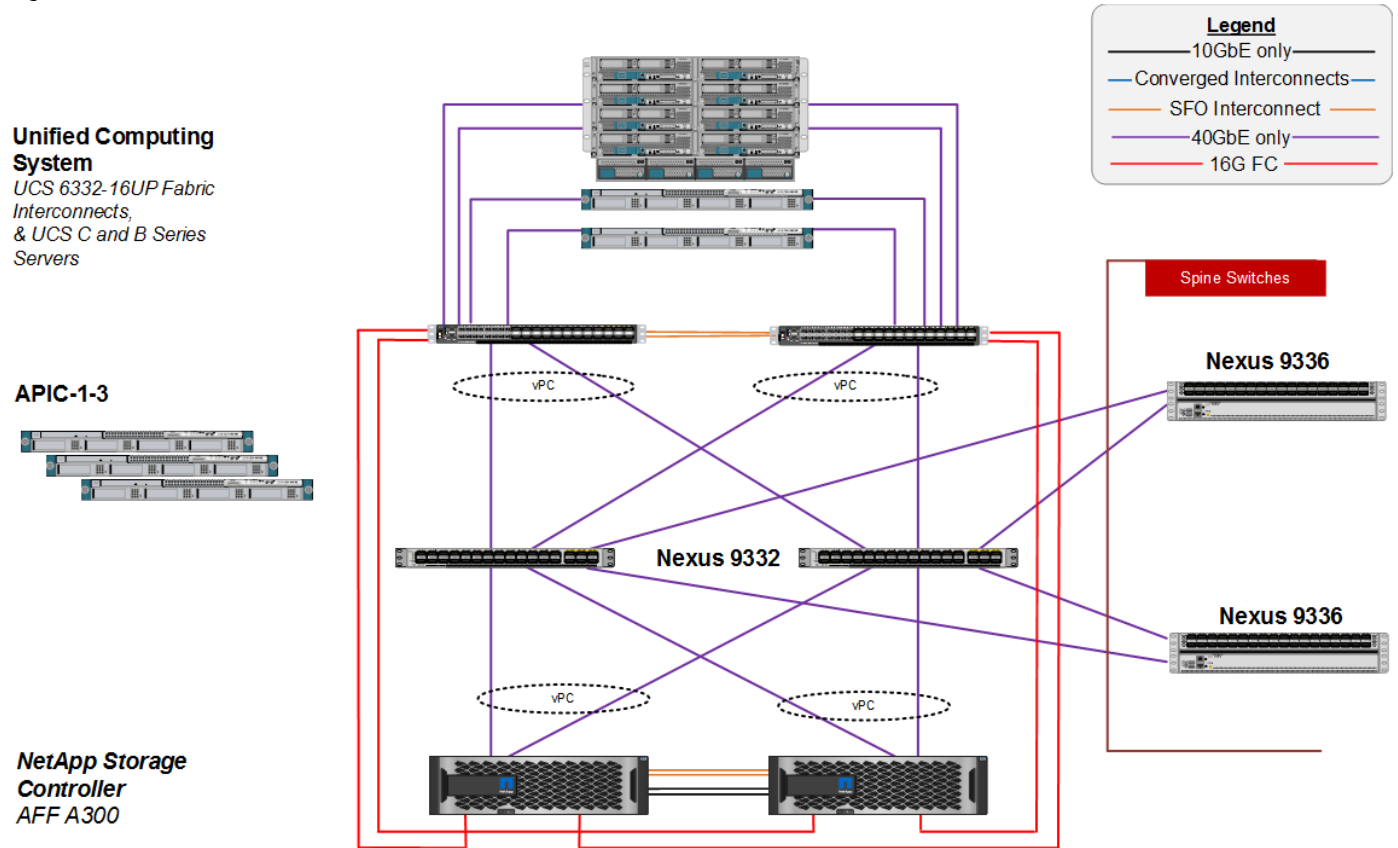
FlexPod architecture is highly modular, or pod-like. Although each customer's FlexPod unit might vary in its exact configuration, after a FlexPod unit is built, it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a FlexPod unit) 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 ACI® networking, Cisco Unified Computing System (Cisco UCS®), Microsoft Systems Center Operations Manager and Microsoft Systems Center 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 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. The Nexus 9336PQ switches shown serve as spine switches in the Cisco ACI Fabric Spine-Leaf Architecture, while the Nexus 9332PQ switches serve as 40GE leaf switches. The Cisco APICs shown attach to the ACI Fabric with 10GE connections. This attachment can be accomplished with either other leaf switches in the fabric with 10 GE ports, such as the Nexus 93180YC-EX, or with Cisco QSFP to SFP/SFP+ Adapter (QSA) modules in the Nexus 9332s. 10GE breakout cables are not supported when the 9332 is in ACI mode. This design has end-to-end 40 Gb Ethernet connections from Cisco UCS Blades, Cisco UCS C-Series rackmount servers, a pair of Cisco UCS Fabric Interconnects, Cisco Nexus 9000 switches, through to NetApp AFF A300. These 40 GE paths carry SMB, iSCSI, and Virtual Machine (VM) traffic that has Cisco ACI policy applied. This infrastructure option can be expanded by connecting 16G FC or 10G FCoE links between the Cisco UCS Fabric Interconnects and the NetApp AFF A300 as shown below, or introducing a pair of Cisco MDS switches between the Cisco UCS Fabric Interconnects and the NetApp AFF A300 to provide FC/FCoE block-level shared storage access. Note that FC/FCoE storage access does not have ACI policy applied. The FC configuration shown below is covered in the appendix of this document, but the FCoE and MDS options are also supported. 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:

- Three Cisco APICs
- Two Cisco Nexus 9336PQ fixed spine switches
- Two Cisco Nexus 9332PQ leaf switches
- Two Cisco UCS 6332-16UP fabric interconnects
- 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)



A 10GE-based design with Cisco UCS 6200 Fabric Interconnects is also supported, but not covered in this deployment Guide. All systems and fabric links feature redundancy and provide end-to-end high availability. For server virtualization, this 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 ca-

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 M5, B-200 M4, UCS C-220 M4 	<ul style="list-style-type: none"> 3.2(1d) - Infrastructure Bundle 3.2(1d) - Server Bundle 	Includes the Cisco UCS-IOM 2304 Cisco UCS Manager, Cisco UCS VIC 1340 and Cisco UCS VIC 1385
Network	Cisco APIC	3.0(1k)	
	Cisco Nexus 9000 ACI	n9000-13.0(1k)	
Storage	NetApp AFF A300	ONTAP 9.1P5	
Software	Cisco UCS Manager	3.2(1d)	
	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 01 or 02 or A and B. In this CVD we have used node01 and node02 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 i0a-<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. In this table VS indicates dynamically assigned VLANs from the APIC-Controlled Microsoft Virtual Switch.

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	3911
MS-IB-MGMT	VLAN for in-band management interfaces	118/218/318/418/VS
Native-VLAN	VLAN to which untagged frames are assigned	2
MS-Infra-SMB-VLAN	VLAN for SMB traffic	3053/3153/VS
MS-LVMN-VLAN	VLAN designated for the movement of VMs from one physical host to another.	906/VS
MS-Cluster-VLAN	VLAN for cluster connectivity	907/VS
MS-Infra-iSCSI-A	VLAN for iSCSI Boot on Fabric A	3013/3113
MS-Infra-iSCSI-B	VLAN for iSCSI Boot on Fabric B	3023/3123

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)	ACI-FP-AD1, ACI-FP-AD2

Virtual Machine Description	Host Name
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.



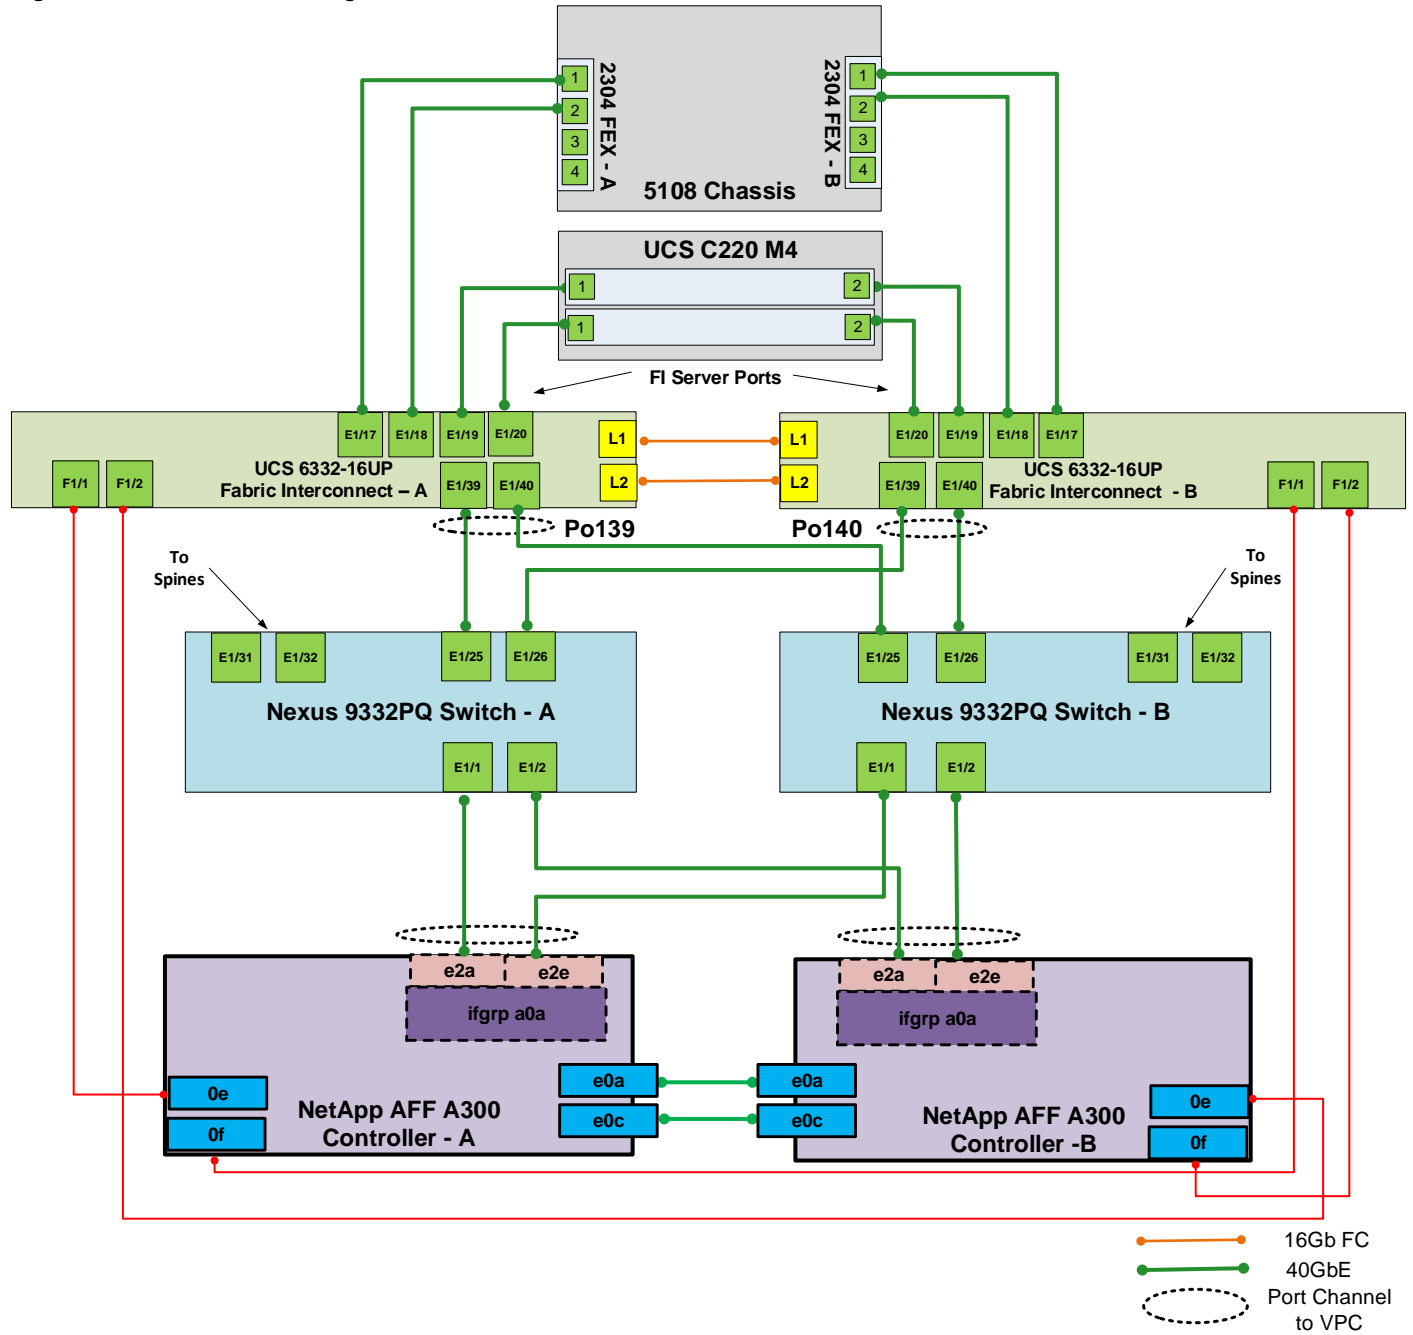
For any modifications of this prescribed architecture, consult the [NetApp Interoperability Matrix Tool \(IMT\)](#). Cisco HyperFlex documents need Cisco.com login credentials. Please login to access these documents.

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 2 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. Four 16Gb links connect directly to the NetApp AFF controllers from the Cisco UCS Fabric Interconnects. 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 2 FlexPod Cabling with Cisco UCS 6332-16UP Fabric Interconnect



Infrastructure Servers Prerequisites

Active Directory DC/DNS

Production environments at most customers' locations 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 controller 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. We will configure two additional AD/DNS servers connected to the Core Services End Point Group (EPG) in the ACI Fabric. These AD/DNS servers will be configured as additional Domain Controllers in the same domain as the prerequisite AD/DNS server.

Microsoft System Center 2016

This document details the steps to install Microsoft System Center Operations Manager (SCOM) and Virtual Machine Manager (SCVMM). The Microsoft guidelines to install SCOM and SCVMM 2016 can be found at:

- 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 ACI fabric for use in a FlexPod environment and is written where the FlexPod components are added to an existing Cisco ACI fabric in several new ACI tenants. Required fabric setup is verified, but previous configuration of the ACI fabric is assumed.



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 section FlexPod Cabling.

In ACI, both spine and leaf switches are configured using APIC, individual configuration of the switches is not required. Cisco APIC discovers the ACI infrastructure switches using LLDP and acts as the central control and management point for the entire configuration.

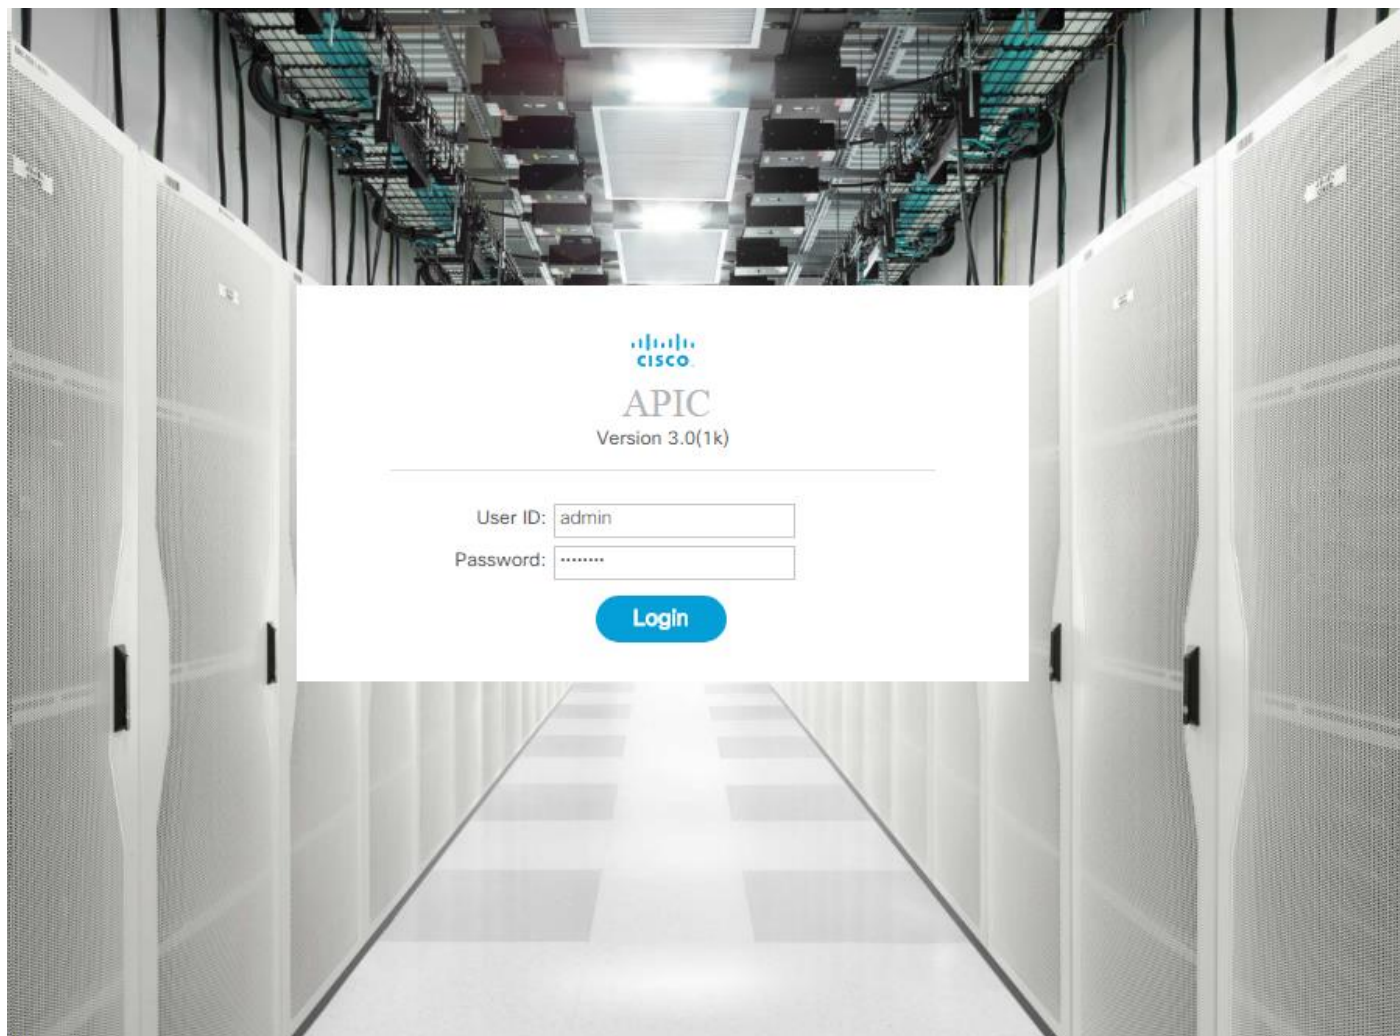
Cisco Application Policy Infrastructure Controller (APIC) Verification

This sub-section verifies the setup the Cisco APIC. Cisco recommends a cluster of at least 3 APICs controlling an ACI Fabric.

1. Log into the APIC GUI using a web browser, by browsing to the out of band IP address configured for APIC. Login with the admin user id and password.

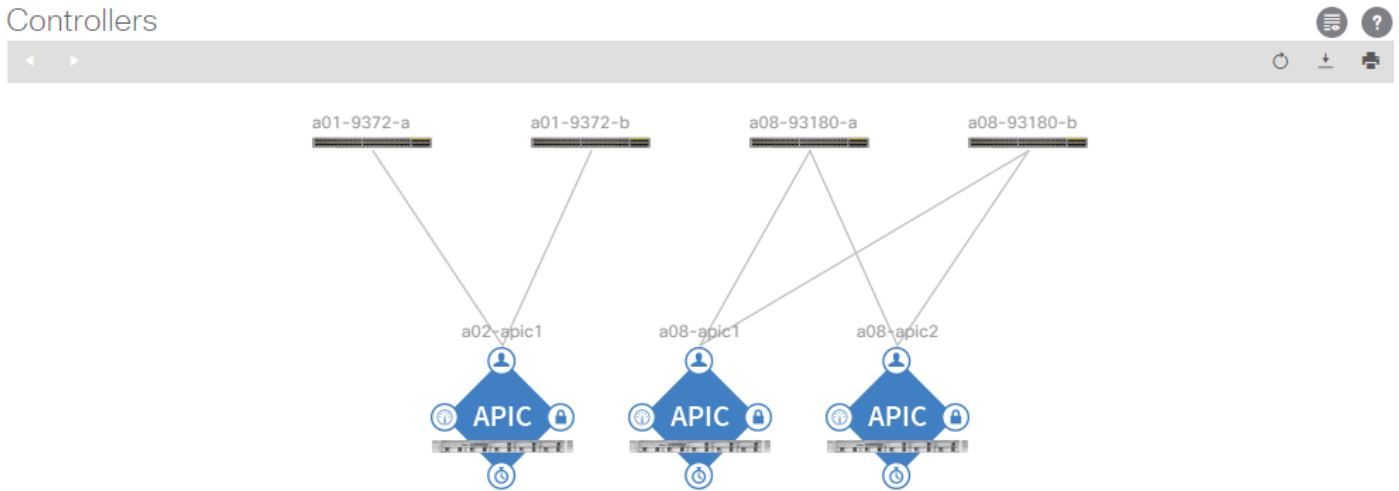


In this validation, Google Chrome was used as the web browser. It might take a few minutes before APIC GUI is available after the initial setup.



2. Take the appropriate action to close any warning or information screens.
3. At the top in the APIC home page, select the System tab followed by Controllers.
4. On the left, select the Controllers folder. Verify that at least 3 APICs are available and have redundant connections to the fabric.

Controllers



Cluster Health

ID	Name	IP	Admin State	Operational State	Health State
1	a08-apic2	10.0.0.1	In Service	Available	Fully Fit
2	a08-apic1	10.0.0.2	In Service	Available	Fully Fit
3	a02-apic1	10.0.0.3	In Service	Available	Fully Fit

Cisco ACI Fabric Discovery

This section details the steps for adding the two Nexus 9332PQ leaf switches to the Fabric. These switches are automatically discovered in the ACI Fabric and are manually assigned node IDs. To add Nexus 9332PQ leaf switches to the Fabric, complete the following steps:

1. At the top in the APIC home page, select the Fabric tab.
2. In the left pane, select and expand Fabric Membership.
3. The two 9332 Leaf Switches will be listed on the Fabric Membership page with Node ID 0 as shown:

Fabric Membership

Serial Number	Pod ID	Node ID	Node Name	Rack Name	Model	Role	IP	Supported Model	SSL Certificate	Status
FDO21131U...	1	104	a08-93180...		N9K-C931...	leaf	10.0.248.5...	True	yes	Active
FDO21131U...	1	103	a08-93180...		N9K-C931...	leaf	10.0.248.4...	True	yes	Active
SAL1913CJXR	1	102	a01-9372-b		N9K-C937...	leaf	10.0.248.1...	True	yes	Active
SAL1914CN42	1	101	a01-9372-a		N9K-C937...	leaf	10.0.152.1...	True	yes	Active
SAL2009ZQJ9	1	0			N9K-C933...	leaf	0.0.0.0	True	n/a	
SAL2009ZQNF	1	0			N9K-C933...	leaf	0.0.0.0	True	n/a	
SAL18391DXU	1	201	a02-9336-1		N9K-C933...	spine	10.0.56.94...	True	yes	Active
SAL18391DYH	1	202	a02-9336-2		N9K-C933...	spine	10.0.56.93...	True	yes	Active

- Connect to the two Nexus 9332 leaf switches using serial consoles and login in as admin with no password (press enter). Use show inventory to get the leaf's serial number.

```
(none)# show inventory
NAME: "Chassis",  DESCR: "Nexus C9332PQ Chassis"
PID: N9K-C9332PQ      ,  VID: V03      ,  SN: SAL2009ZQJ9

NAME: "Slot 1 ",  DESCR: "32x40G Supervisor  "
PID: N9K-C9332PQ      ,  VID: V03      ,  SN: SAL2009ZQJ9
```

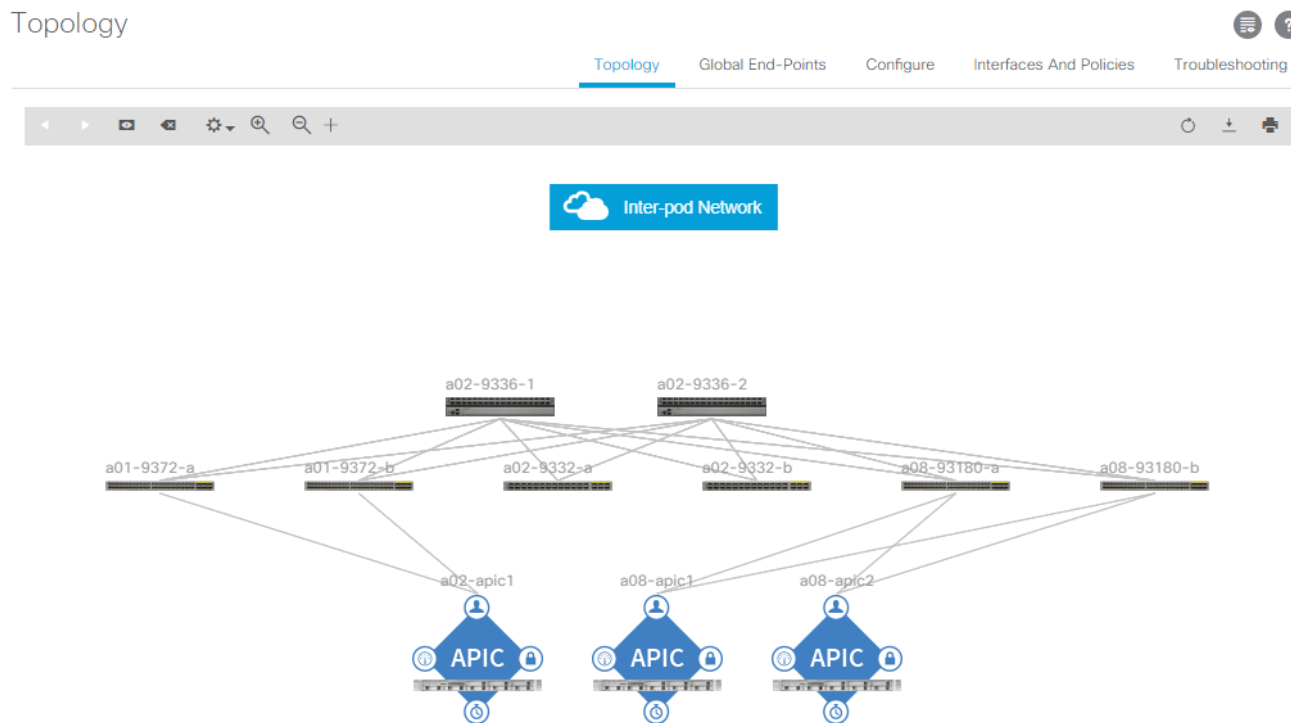
- Match the serial numbers from the leaf listing to determine the A and B switches under Fabric Membership.
- In the APIC GUI, under Fabric Membership, double-click the A leaf in the list. Enter a Node ID and a Node Name for the Leaf switch and click Update.


Fabric Membership

Serial Number	Pod ID	Node ID	Node Name	Rack Name	Model	Role	IP	Supported Model	SSL Certificate	Status
FDO21131U...	1	104	a08-93180...		N9K-C931...	leaf	10.0.248.5...	True	yes	Active
FDO21131U...	1	103	a08-93180...		N9K-C931...	leaf	10.0.248.4...	True	yes	Active
SAL1913CJXR	1	102	a01-9372-b		N9K-C937...	leaf	10.0.248.1...	True	yes	Active
SAL1914CN42	1	101	a01-9372-a		N9K-C937...	leaf	10.0.152.1...	True	yes	Active
SAL2009ZQJ9	1	105	a02-9332-a	select	N9K-C9332PQ	leaf	0.0.0.0	True	n/a	
SAL2009ZQNF	1	0			N9K-C933...	leaf	0.0.0.0	True	n/a	
SAL18391DXU	1	201	a02-9336-1		N9K-C933...	spine	10.0.56.94...	True	yes	Active
SAL18391DYH	1	202	a02-9336-2		N9K-C933...	spine	10.0.56.93...	True	yes	Active

- Repeat step 6 for the B leaf in the list.

- Click Topology in the left pane. The discovered ACI Fabric topology will appear. It may take a few minutes for the Nexus 9332 Leaf switches to appear and you will need to click the refresh button for the complete topology to appear.



 The topology shown in the screenshot above is the topology of the validation lab fabric containing 6 leaf switches, 2 spine switches, and 3 APICs. Notice that the APICs are connected to 10GE ports. Customer topology will vary depending on number and type of devices. Cisco recommends a cluster of at least 3 APICs in a production environment.

Initial ACI Fabric Setup Verification

This section details the steps for the initial setup of the Cisco ACI Fabric, where the software release is validated, out of band management IPs are assigned to the new leaves, NTP setup is verified, and the fabric BGP route reflectors are verified.

Software Upgrade

To upgrade the software, complete the following steps:

- In the APIC GUI, at the top select Admin > Firmware.
- This document was validated with ACI software release 3.0(1k). Select Fabric Node Firmware in the left pane under Firmware Management. All switches should show the same firmware release and the release version should be at minimum n9000-13.0(1k). The switch software version should also match the APIC version.

Fabric Node Firmware

Policy Faults History

Firmware Default Policy
Enforce Bootscript Version Validation:

All Nodes

Node id	Node name	Model	Current Firmware	Status	Role	Firmware Group	Maintenance Group
Current Firmware: n9000-13.0(1k) (8 Nodes)							
101	a01-937...	N9K-C9372PX	n9000-13.0(1k)	Upgraded successfully on 2017-08...	leaf	Odd-Switches	Odd-Switches
102	a01-937...	N9K-C9372PX	n9000-13.0(1k)	Upgraded successfully on 2017-08...	leaf	Even-Switches	Even-Switches
103	a08-931...	N9K-C93180Y...	n9000-13.0(1k)	Upgraded successfully on 2017-08...	leaf	Odd-Switches	Odd-Switches
104	a08-931...	N9K-C93180Y...	n9000-13.0(1k)	Upgraded successfully on 2017-08...	leaf	Even-Switches	Even-Switches
105	a02-933...	N9K-C9332PQ	n9000-13.0(1k)	Upgraded successfully on 2017-09...	leaf	Odd-Switches	Odd-Switches
106	a02-933...	N9K-C9332PQ	n9000-13.0(1k)	Upgraded successfully on 2017-09...	leaf	Even-Switches	Even-Switches
201	a02-933...	N9K-C9336PQ	n9000-13.0(1k)	Upgraded successfully on 2017-08...	spine	Odd-Switches	Odd-Switches
202	a02-933...	N9K-C9336PQ	n9000-13.0(1k)	Upgraded successfully on 2017-08...	spine	Even-Switches	Even-Switches

- Click Admin > Firmware > Controller Firmware. If all APICs are not at the same release at a minimum of 3.0(1k), follow the [Cisco APIC Controller and Switch Software Upgrade and Downgrade Guide](#) to upgrade both the APICs and switches to a minimum release of 3.0(1k) on APIC and 13.0(1k) on the switches.

Setting Up Out-of-Band Management IP Addresses for New Leaf and Switches

To set up out-of-band management IP addresses, complete the following steps:

- To add out-of-band management interfaces for all the switches in the ACI Fabric, select Tenants > mgmt.
- Expand Tenant mgmt on the left. Right-click Node Management Addresses and select Create Static Node Management Addresses.
- Enter the node number range for the new leaf switches (105-106 in this example).
- Select the checkbox for Out-of-Band Addresses.
- Select default for Out-of-Band Management EPG.
- Considering that the IPs will be applied in a consecutive range of two IPs, enter a starting IP address and netmask in the Out-Of-Band IPV4 Address field.
- Enter the out of band management gateway address in the Gateway field.
- Click SUBMIT, then click YES.
- On the left, expand Node Management Addresses and select Static Node Management Addresses. Verify the mapping of IPs to switching nodes.

Static Node Management Addresses

Node	Type	EPG	IPv4 Address	IPv4 Gateway	IPv6 Address	IPv6 Gateway
pod-1/node-105	Out-Of-Band	default	192.168.1.21/24	192.168.1.254	::	::
pod-1/node-106	Out-Of-Band	default	192.168.1.22/24	192.168.1.254	::	::
pod-1/node-101	Out-Of-Band	default	192.168.1.35/24	192.168.1.254	::	::
pod-1/node-102	Out-Of-Band	default	192.168.1.36/24	192.168.1.254	::	::
pod-1/node-103	Out-Of-Band	default	192.168.1.37/24	192.168.1.254	::	::
pod-1/node-104	Out-Of-Band	default	192.168.1.38/24	192.168.1.254	::	::
pod-1/node-201	Out-Of-Band	default	192.168.1.39/24	192.168.1.254	::	::
pod-1/node-202	Out-Of-Band	default	192.168.1.40/24	192.168.1.254	::	::

10. Direct out-of-band access to the switches is now available for SSH.

Verifying Time Zone and NTP Server

This procedure will allow customers to verify setup of an NTP server for synchronizing the fabric time. To verify the time zone and NTP server set up, complete the following steps:

1. To verify NTP setup in the fabric, select and expand Fabric > Fabric Policies > Pod Policies > Policies > Date and Time.
2. Select default. In the Datetime Format - default pane, verify the correct Time Zone is selected and that Offset State is enabled. Adjust as necessary and click Submit and Submit Changes.
3. On the left, select Policy default. Verify that at least one NTP Server is listed.

Date and Time Policy - Policy default

?

↻ ⬇ ✖

Properties

Name: default

Description:

Administrative State: disabled enabled

Authentication State: disabled enabled

NTP Servers:

Host Name/IP Address	Preferred	Minimum Polling Interval	Maximum Polling Interval	Management EPG
192.168.1.254	False	4	6	default (Out-of-Band)

🗑 +

4. If necessary, on the right use the + sign to add NTP servers accessible on the out of band management subnet. Enter an IP address accessible on the out of band management subnet and select the default (Out-of-Band) Management EPG. Click Submit to add the NTP server. Repeat this process to add all NTP servers.

Verifying Domain Name Servers

To verify optional DNS in the ACI fabric, complete the following steps:

1. Select and expand Fabric > Fabric Policies > Global Policies > DNS Profiles > default.
2. Verify the DNS Providers and DNS Domains.
3. If necessary, in the Management EPG drop-down, select the default (Out-of-Band) Management EPG. Use the + signs to the right of DNS Providers and DNS Domains to add DNS servers and the DNS domain name. Note that the DNS servers should be reachable from the out of band management subnet. Click SUBMIT to complete the DNS configuration.

DNS Profile - default

Policy History

Properties

Description: optional

Management EPG: default (Out-of-Band)

DNS Providers:

Address	Preferred
172.26.163.251	False

DNS Domains:

Name	Default	Description
flexpod.cisco.com	False	

Verifying BGP Route Reflectors

In this ACI deployment, both the spine switches should be set up as BGP route-reflectors to distribute the leaf routes throughout the fabric. To verify the BGP Route Reflector, complete the following steps:

1. Select and expand System > System Settings > BGP Route Reflector.
2. Verify that a unique Autonomous System Number has been selected for this ACI fabric. If necessary, use the + sign on the right to add the two spines to the list of Route Reflector Nodes. Click SUBMIT to complete configuring the BGP Route Reflector.

BGP Route Reflector Policy - BGP Route Reflector

Policy Faults History

🔄 ⚠️ 🔍 📄
🔄 ⬇️ 🗑️

Properties

Name: default

Description: optional

Autonomous System Number: 101

Route Reflector Nodes:

Node ID	Node Name	Description
201	a02-9336-1	
202	a02-9336-2	

External Route Reflector Nodes:

Node ID	Node Name	Description
No items have been found. Select Actions to create a new item.		

- To verify the BGP Route Reflector has been enabled, select and expand Fabric > Fabric Policies > Pod Policies > Policy Groups. Under Policy Groups make sure a policy group has been created and select it. The BGP Route Reflector Policy field should show “default.”

Pod Policy Group - pod1-policygrp

Policy Faults History

🔄 ⚠️ 🔍 📄
🔄 ⬇️ 🗑️

Properties

Name: pod1-policygrp

Description: optional

Date Time Policy: select a value

Resolved Date Time Policy: default

ISIS Policy: select a value

Resolved ISIS Policy: default

COOP Group Policy: select a value

Resolved COOP Group Policy: default

BGP Route Reflector Policy: default

Resolved BGP Route Reflector Policy: default

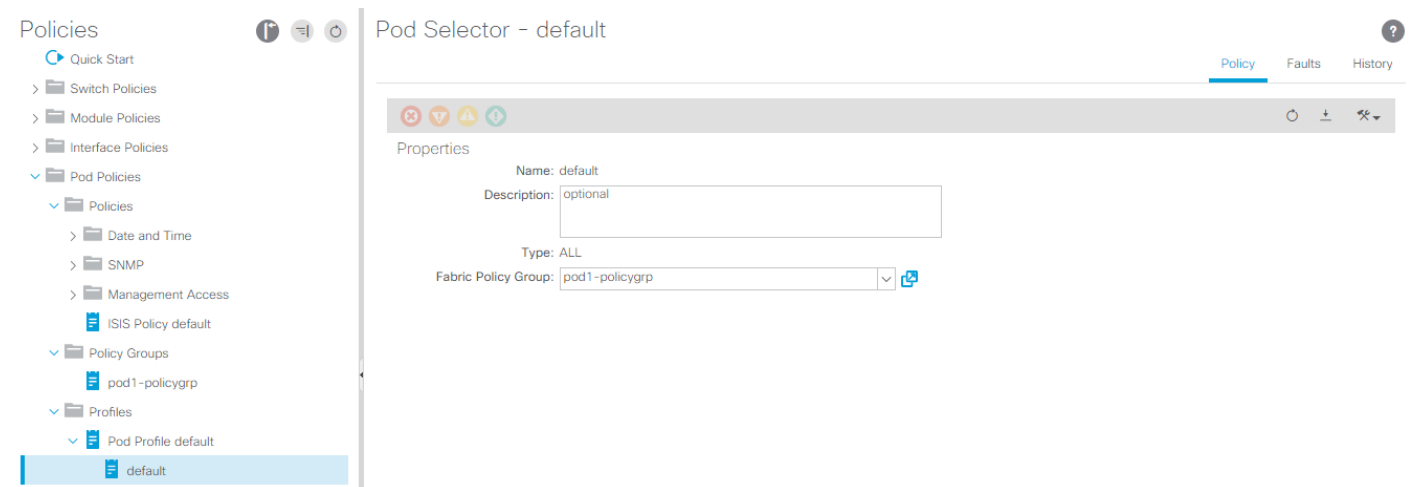
Management Access Policy: select a value

Resolved Management Access Policy: default

SNMP Policy: select a value

Resolved SNMP Policy: default

- If a Policy Group has not been created, on the left, right-click Policy Groups under Pod Policies and select Create Pod Policy Group. In the Create Pod Policy Group window, name the Policy Group pod1-policygrp. Select the default BGP Route Reflector Policy. Click SUBMIT to complete creating the Policy Group.
- On the left expand Profiles under Pod Policies and select Pod Profile default > default.
- Verify that the pod1-policygrp or the Fabric Policy Group identified above is selected. If the Fabric Policy Group is not selected, view the drop-down list to select it and click Submit.



Set Up Fabric Access Policy Setup

This section details the steps to create various access policies creating parameters for CDP, LLDP, LACP, etc. These policies are used during vPC and VM domain creation. In an existing fabric, these policies may already exist. The existing policies can be used if configured the same way as listed. To define fabric access policies, complete the following steps:

- Log into the APIC AGUI.
- In the APIC UI, select and expand Fabric > Access Policies > Interface Policies > Policies.

Create Link Level Policies

This procedure will create link level policies for setting up the 1Gbps, 10Gbps, and 40Gbps link speeds. To create the link level policies, complete the following steps:

- In the left pane, right-click Link Level and select Create Link Level Policy.
- Name the policy as 1Gbps-Auto and select the 1Gbps Speed.

Create Link Level Policy



Specify the Physical Interface Policy Identity

Name:

Description:

Alias:

Auto Negotiation:

Speed:

Link debounce interval (msec):

Forwarding Error Correction:

Cancel

Submit

3. Click Submit to complete creating the policy.
4. In the left pane, right-click Link Level and select Create Link Level Policy.
5. Name the policy 10Gbps-Auto and select the 10Gbps Speed.
6. Click Submit to complete creating the policy.
7. In the left pane, right-click Link Level and select Create Link Level Policy.
8. Name the policy 40Gbps-Auto and select the 40Gbps Speed.
9. Click Submit to complete creating the policy.

Create CDP Policy

This procedure creates policies to enable or disable CDP on a link. To create a CDP policy, complete the following steps:

1. In the left pane, right-click CDP interface and select Create CDP Interface Policy.

2. Name the policy as CDP-Enabled and enable the Admin State.

Create CDP Interface Policy



Specify the CDP Interface Policy Identity

Name:

Description:

Alias:

Admin State: Disabled Enabled

Cancel

Submit

3. Click Submit to complete creating the policy.
4. In the left pane, right-click the CDP Interface and select Create CDP Interface Policy.
5. Name the policy CDP-Disabled and disable the Admin State.
6. Click Submit to complete creating the policy.

Create LLDP Interface Policies

This procedure will create policies to enable or disable LLDP on a link. To create an LLDP Interface policy, complete the following steps:

1. In the left pane, right-click LLDP Interface and select Create LLDP Interface Policy.
2. Name the policy as LLDP-Enabled and enable both Transmit State and Receive State.

Create LLDP Interface Policy



Specify the LLDP Interface Policy Properties

Name:

Description:

Alias:

Receive State: Disabled Enabled

Transmit State: Disabled Enabled

3. Click Submit to complete creating the policy.
4. In the left, right-click the LLDP Interface and select Create LLDP Interface Policy.
5. Name the policy as LLDP-Disabled and disable both the Transmit State and Receive State.
6. Click Submit to complete creating the policy.

Create Port Channel Policy

This procedure will create policies to set LACP active mode configuration, LACP Mode On configuration and the MAC-Pinning mode configuration. To create Port Channel policy, complete the following steps:

1. In the left pane, right-click the Port Channel and select Create Port Channel Policy.
2. Name the policy as LACP-Active and select LACP Active for the Mode. Do not change any of the other values.

Create Port Channel Policy



Specify the Port Channel Policy

Name:

Description:

Alias:

Mode:

Control:

Minimum Number of Links:
Not Applicable for FEX PC/VPC

Maximum Number of Links:
Not Applicable for FEX PC/VPC

Cancel

Submit

3. Click Submit to complete creating the policy.
4. In the left pane, right-click Port Channel and select Create Port Channel Policy.
5. Name the policy as MAC-Pinning and select MAC Pinning-Physical-NIC-load for the Mode. Do not change any of the other values.

Create Port Channel Policy



Specify the Port Channel Policy

Name:

Description:

Alias:

Mode:

Minimum Number of Links:
Not Applicable for FEX PC/VPC

Maximum Number of Links:
Not Applicable for FEX PC/VPC

Cancel

Submit

6. Click Submit to complete creating the policy.
7. In the left pane, right-click Port Channel and select Create Port Channel Policy.

Create BPDU Filter/Guard Policies

This procedure will create policies to enable or disable BPDU filter and guard. To create a BPDU filter/Guard policy, complete the following steps:

1. In the left pane, right-click Spanning Tree Interface and select Create Spanning Tree Interface Policy.
2. Name the policy as BPDU-FG-Enabled and select both the BPDU filter and BPDU Guard Interface Controls.

Create Spanning Tree Interface Policy



Define the STP Interface Policy

Name:

Description:

Alias:

Interface controls: BPDU filter enabled
 BPDU Guard enabled

Cancel

Submit

3. Click Submit to complete creating the policy.
4. In the left pane, right-click Spanning Tree Interface and select Create Spanning Tree Interface Policy.
5. Name the policy as BPDU-FG-Disabled and make sure both the BPDU filter and BPDU Guard Interface Controls are cleared.
6. Click Submit to complete creating the policy.

Create Global VLAN Policy

To create policies to enable global scope for all the VLANs, complete the following steps:

1. In the left pane, right-click the L2 Interface and select Create L2 Interface Policy.
2. Name the policy as VLAN-Scope-Global and make sure Global scope is selected. Do not change any of the other values.

Create L2 Interface Policy



Define the L2 Interface Policy

Name:

Description:

QinQ:

Reflective Relay (802.1Qbg):

VLAN Scope:

Cancel

Submit

3. Click Submit to complete creating the policy.

Create Firewall Policy

To create policies to disable a firewall, complete the following steps:

1. In the left pane, right-click Firewall and select Create Firewall Policy.
2. Name the policy Firewall-Disabled and select Disabled for Mode. Do not change any of the other values.

Create Firewall Policy



Specify the Firewall Policy Properties

Name:

Description:

Mode: Disabled Enabled Learning

SysLog

Administrative State:

Included Flows:

Polling Interval (seconds):

Log Level:

Dest Group:

- Click Submit to complete creating the policy.

Create Virtual Port Channels (vPCs)


This section details the steps to setup vPCs for connectivity to the In-Band Management Network, Cisco UCS, and NetApp Storage.

VPC - Management Switch

To setup vPCs for connectivity to the existing In-Band Management Network, complete the following steps:



This deployment guide covers the configuration for a pre-existing Cisco Nexus management switch. You can adjust the management configuration depending on your connectivity setup. The In-Band Management Network provides connectivity of Management Virtual Machines and Hosts in the ACI fabric to existing services on the In-Band Management network outside of the ACI fabric. Layer 3 connectivity is assumed between the In-Band and Out-of-Band Management networks. This setup creates management networks that are physically isolated from tenant networks. In this validation, a 10GE vPC from two 10GE capable leaf switches in the fabric is connected to a port-channel on a Nexus 5K switch outside the fabric. Note that this vPC is not created on the Nexus 9332 leaves, but on existing leaves that have 10GE ports.



1. In the APIC GUI, at the top select Fabric > Access Policies > Quick Start.
2. In the right pane select Configure an interface, PC and VPC.
3. In the configuration window, configure a VPC domain between the leaf switches by clicking “+” under VPC Switch Pairs. If a VPC Domain already exists between the two switches being used for this vPC, skip to step 7.

VPC Switch Pairs



4. Enter a VPC Domain ID (1 in this example).
5. From the drop-down list, select Switch A and Switch B IDs to select the two leaf switches.

Select two switches to be paired for VPC.

Only switches with interfaces in the same VPC policy group can be paired together.

VPC Domain ID:  

Switch 1:  

Switch 2:  

- Click SAVE.
- Click the “+” under Configured Switch Interfaces.


Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
 			

- From the Switches drop-down list on the right, select both the leaf switches being used for this vPC.
- Leave the system generated Switch Profile Name in place.
- Click the **big green “+”** to configure switch interfaces.

Select Switches To Configure Interfaces: Quick Advanced

Switches:  Switch Profile Name:



- Configure various fields as shown in the figure below. In this screen shot, port 1/21 on both leaf switches is connected to a Nexus switch using 10Gbps links.

Select Switches To Configure Interfaces: Quick Advanced

Switches: Switch Profile Name:

Interface Type: Individual PC VPC

Interfaces: Interface Selector Name:
Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One


Link Level Policy: <input type="text" value="10Gbps-Auto"/>	CDP Policy: <input type="text" value="CDP-Enabled"/>
MCP Policy: <input type="text" value="select a value"/>	LLDP Policy: <input type="text" value="LLDP-Enabled"/>
STP Interface Policy: <input type="text" value="BPDU-FG-Disabled"/>	Monitoring Policy: <input type="text" value="select a value"/>
Storm Control Policy: <input type="text" value="select a value"/>	L2 Interface Policy: <input type="text" value="VLAN-Scope-Global"/>
Port Security Policy: <input type="text" value="select a value"/>	Egress Data Plane Policing Policy: <input type="text" value="select a value"/>
Ingress Data Plane Policing Policy: <input type="text" value="select a value"/>	IPv4 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Priority Flow Control Policy: <input type="text" value="select a value"/>	IPv6 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Slow Drain Policy: <input type="text" value="select a value"/>	Layer2-Switched (CE type) NetFlow Monitor Policy: <input type="text" value="select a value"/>
Fibre Channel Interface Policy: <input type="text" value="select a value"/>	
Port Channel Policy: <input type="text" value="LACP-Active"/>	

Attached Device Type:

Domain: Create One Choose One Domain Name:

VLAN: Create One Choose One VLAN Range:
Please use comma to separate VLANs.

12. Click Save.
13. Click Save again to finish the configuring switch interfaces.
14. Click Submit.

 To validate the configuration, log into the Nexus switch and verify the port-channel is up (`show port-channel summary`).

VPC – UCS Fabric Interconnects

Complete the following steps to setup vPCs for connectivity to the UCS Fabric Interconnects.

Figure 3 VLANs Configured for Cisco UCS






Table 4 VLANs for Cisco UCS Hosts

Name	VLAN
Native	<2>
MS-Core-Services	<318>
MS-IB-Mgmt	<418>
MS-LVMN	<906>
MS-Cluster	<907>
MS-Infra-SMB	<3153>
MS-Infra-iSCSI-A	<3113>
MS-Infra-iSCSI-B	<3123>

 MS-Core-Services and MS-IB-MGMT will be in the same bridge domain and subnet; they have to be in different VLANs because we are using the VLAN-Scope-Global L2 Interface Policy.

 MS-LVMN, MS-Cluster, and MS-Infra-SMB VLANs are configured in this section and will be in place if needed on the manually created Hyper-V Virtual Switch. The EPG should be configured; it is not necessary to configure the actual VLAN or UCS static port mapping, but you can configure these without any negative effects.

1. In the APIC GUI, select Fabric > Access Policies > Quick Start.
2. In the right pane, select Configure and interface, PC and VPC.
3. In the configuration window, configure a VPC domain between the 9332 leaf switches by clicking “+” under VPC Switch Pairs.

VPC Switch Pairs

VPC Domain Id	Switch 1	Switch 2

4. Enter a VPC Domain ID (10 in this example).
5. From the drop-down list, select 9332 Switch A and 9332 Switch B IDs to select the two leaf switches.

Select two switches to be paired for VPC.
Only switches with interfaces in the same VPC policy group can be paired together.


VPC Domain ID:   

Switch 1:  

Switch 2:  


6. Click Save.
7. Click the “+” under Configured Switch Interfaces.
8. Select the two Nexus 9332 switches under the Switches drop-down list.

Select Switches To Configure Interfaces: Quick Advanced

Switches:  Switch Profile Name:

 Click '+' to configure switch interfaces



9. Click  to add switch interfaces.
10. Configure various fields as shown in the figure below. In this screenshot, port 1/25 on both leaf switches is connected to UCS Fabric Interconnect A using 40Gbps links.

Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,...			
> 102			
> 103			
> 103,...			
> 104			
> 106,...			
	1/25	VPC	L2 (VLANs: 318,906,2,90...
	1/26	VPC	L2 (VLANs: 318,906,2,90...
	1/1	VPC	Bare Metal (VLANs: 218,3...
	1/2	VPC	Bare Metal (VLANs: 218,3...

VPC Switch Pairs

VPC Domain Id	Switch 1	Switch 2
1	102	101
2	103	104
10	105	106

Select Switches To Configure Interfaces: Quick Advanced

Switches: 105-106 Switch Profile Name: Switch105-106_Profile

Interface Type: Individual PC VPC

Interfaces: 1/25 Interface Selector Name: A02-6332-AJ-ports-25

Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One

Link Level Policy: 40Gbps-Auto	CDP Policy: CDP-Enabled
MCP Policy: select a value	LLDP Policy: LLDP-Enabled
STP Interface Policy: BPDU-FG-Enabled	Monitoring Policy: select a value
Storm Control Policy: select a value	L2 Interface Policy: VLAN-Scope-Global
Port Security Policy: select a value	Egress Data Plane Policing Policy: select a value
Ingress Data Plane Policing Policy: select a value	IPv4 NetFlow Monitor Policy: select a value
Priority Flow Control Policy: select a value	IPv6 NetFlow Monitor Policy: select a value
Slow Drain Policy: select a value	Layer2-Switched (CE type) NetFlow Monitor Policy: select a value
Fibre Channel Interface Policy: select a value	
Port Channel Policy: LACP-Active	

Attached Device Type: External Bridged Devices

Domain: Create One Choose One Domain Name: UCS

VLAN: Create One Choose One VLAN Range: 2,318,418,906,907,3153,3113,3123

Please use comma to separate VLANs.

Cancel
Save

11. Click Save.
12. Click Save again to finish the configuring switch interfaces.
13. Click Submit.
14. From the right pane, select Configure and interface, PC and VPC.
15. Select the switches configured in the last step under Configured Switch Interfaces.

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,102			
> 102			
> 103			
> 103,104			
> 104			
✓ 106,105	1/25	VPC	L2 (VLANs: 906,318,3...

16. Click to add switch interfaces.

17. Configure various fields as shown in the screenshot. In this screenshot, port 1/26 on both leaf switches is connected to UCS Fabric Interconnect B using 40Gbps links. Instead of creating a new domain, the External Bridge Domain created in the last step (UCS) is attached to the FI-B as shown below.

Select Switches To Configure Interfaces: Quick Advanced

Switches: Switch Profile Name:

Interface Type: Individual PC VPC

Interfaces: Interface Selector Name:

Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One

Link Level Policy: <input type="text" value="40Gbps-Auto"/>	CDP Policy: <input type="text" value="CDP-Enabled"/>
MCP Policy: <input type="text" value="select a value"/>	LLDP Policy: <input type="text" value="LLDP-Enabled"/>
STP Interface Policy: <input type="text" value="BPDU-FG-Enabled"/>	Monitoring Policy: <input type="text" value="select a value"/>
Storm Control Policy: <input type="text" value="select a value"/>	L2 Interface Policy: <input type="text" value="VLAN-Scope-Global"/>
Port Security Policy: <input type="text" value="select a value"/>	Egress Data Plane Policing Policy: <input type="text" value="select a value"/>
Ingress Data Plane Policing Policy: <input type="text" value="select a value"/>	IPv4 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Priority Flow Control Policy: <input type="text" value="select a value"/>	IPv6 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Slow Drain Policy: <input type="text" value="select a value"/>	Layer2-Switched (CE type) NetFlow Monitor Policy: <input type="text" value="select a value"/>
Fibre Channel Interface Policy: <input type="text" value="select a value"/>	
Port Channel Policy: <input type="text" value="LACP-Active"/>	

Attached Device Type:

Domain: Create One Choose One External Bridge Domain:

18. Click Save.
19. Click Save again to finish the configuring switch interfaces.
20. Click Submit.
21. Optional: Repeat this procedure to configure any additional UCS domains. For a uniform configuration, the External Bridge Domain (UCS) will be utilized for all the Fabric Interconnects.

VPC – NetApp AFF Cluster

Complete the following steps to setup vPCs for connectivity to the NetApp AFF storage controllers. The VLANs configured for NetApp are shown in the table below.



Since Global VLAN Scope is being used in this environment, unique VLAN IDs must be used for each different entry point into the ACI fabric. The VLAN IDs for the same named VLANs are different.




Table 5 VLANs for Storage

Name	VLAN
MS-IB-MGMT	<218>
MS-Infra-SMB	<3053>
MS-Infra-iSCSI-A	<3013>
MS-Infra-iSCSI-B	<3023>

1. In the APIC GUI, select Fabric > Access Policies > Quick Start.
2. In the right pane, select Configure and interface, PC and VPC.
3. Select the paired Nexus 9332 switches configured in the last step under Configured Switch Interfaces.

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,102			
> 102			
> 103			
> 103,104			
> 104			
✓ 106,105			
	1/25	VPC	L2 (VLANs: 318,906,3...
	1/26	VPC	L2 (VLANs: 318,906,3...

- Click to add switch interfaces.
- Configure various fields as shown in the screenshot below. In this screen shot, port 1/1 on both leaf switches is connected to Storage Controller 1 using 40Gbps links.

Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,...			
> 102			
> 103			
> 103,...			
> 104			
> 106,...			
1/25	VPC	L2 (VLANs: 318,906,2,90...	
1/26	VPC	L2 (VLANs: 318,906,2,90...	
1/1	VPC	Bare Metal (VLANs: 218,3...	
1/2	VPC	Bare Metal (VLANs: 218,3...	

VPC Switch Pairs

VPC Domain Id	Switch 1	Switch 2
1	102	101
2	103	104
10	105	106

Select Switches To Configure Interfaces: Quick Advanced

Switches: Switch Profile Name:

Interface Type: Individual PC VPC

Interfaces: Interface Selector Name:
Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One

Link Level Policy: <input type="text" value="40Gbps-Auto"/>	CDP Policy: <input type="text" value="CDP-Enabled"/>
MCP Policy: <input type="text" value="select a value"/>	LLDP Policy: <input type="text" value="LLDP-Disabled"/>
STP Interface Policy: <input type="text" value="BPDU-FG-Enabled"/>	Monitoring Policy: <input type="text" value="select a value"/>
Storm Control Policy: <input type="text" value="select a value"/>	L2 Interface Policy: <input type="text" value="VLAN-Scope-Global"/>
Port Security Policy: <input type="text" value="select a value"/>	
Ingress Data Plane Policing Policy: <input type="text" value="select a value"/>	Egress Data Plane Policing Policy: <input type="text" value="select a value"/>
Priority Flow Control Policy: <input type="text" value="select a value"/>	IPv4 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Slow Drain Policy: <input type="text" value="select a value"/>	IPv6 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Fibre Channel Interface Policy: <input type="text" value="select a value"/>	Layer2-Switched (CE type) NetFlow Monitor Policy: <input type="text" value="select a value"/>
Port Channel Policy: <input type="text" value="LACP-Active"/>	

Attached Device Type:

Domain: Create One Choose One Domain Name:

VLAN: Create One Choose One VLAN Range:
Please use comma to separate VLANs.

6. Click Save.
7. Click Save again to finish the configuring switch interfaces.
8. Click Submit.
9. From the right pane, select Configure and interface, PC and VPC.
10. Select the paired Nexus 9332 switches configured in the last step under Configured Switch Interfaces.

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,102			
> 102			
> 103			
> 103,104			
> 104			
✓ 106,105			
	1/25	VPC	L2 (VLANs: 318,906,3...
	1/26	VPC	L2 (VLANs: 318,906,3...
	1/1	VPC	Bare Metal (VLANs: 31...

11. Click to add switch interfaces.

12. Configure various fields as shown in the screenshot below. In this screenshot, port 1/2 on both leaf switches is connected to Storage Controller 2 using 40Gbps links. Instead of creating a new domain, the Bare Metal Domain created in the previous step (NetApp-AFF) is attached to the storage controller 2 as shown below.

Select Switches To Configure Interfaces: Quick Advanced

Switches: Switch Profile Name:

Interface Type: Individual PC VPC

Interfaces: Interface Selector Name:
Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One

Link Level Policy: <input type="text" value="40Gbps-Auto"/>	CDP Policy: <input type="text" value="CDP-Enabled"/>
MCP Policy: <input type="text" value="select a value"/>	LLDP Policy: <input type="text" value="LLDP-Disabled"/>
STP Interface Policy: <input type="text" value="BPDU-FG-Enabled"/>	Monitoring Policy: <input type="text" value="select a value"/>
Storm Control Policy: <input type="text" value="select a value"/>	L2 Interface Policy: <input type="text" value="VLAN-Scope-Global"/>
Port Security Policy: <input type="text" value="select a value"/>	Egress Data Plane Policing Policy: <input type="text" value="select a value"/>
Ingress Data Plane Policing Policy: <input type="text" value="select a value"/>	IPv4 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Priority Flow Control Policy: <input type="text" value="select a value"/>	IPv6 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Slow Drain Policy: <input type="text" value="select a value"/>	Layer2-Switched (CE type) NetFlow Monitor Policy: <input type="text" value="select a value"/>
Fibre Channel Interface Policy: <input type="text" value="select a value"/>	
Port Channel Policy: <input type="text" value="LACP-Active"/>	

Attached Device Type:

Domain: Create One Choose One Physical Domain:

13. Click Save.
14. Click Save again to finish the configuring switch interfaces.
15. Click Submit.
16. Optional: Repeat this procedure to configure any additional NetApp AFF storage controllers. For a uniform configuration, the Bare Metal Domain (NetApp-AFF) will be utilized for all the Storage Controllers.

Configuring Common Tenant for Management Access

This section details the steps to setup in-band management access in the Tenant common. This design will allow all the other tenant EPGs to access the common management segment for Core Services VMs such as AD/DNS.

1. In the APIC GUI, select Tenants > common.
2. In the left pane, expand Tenant common and Networking.

Create VRFs

To create VRFs, complete the following steps:

1. Right-click VRFs and select Create VRF.
2. Enter vrf-FP-Common-IB-MGMT as the name of the VRF.
3. Uncheck Create A Bridge Domain.
4. Click Finish.

Create VRF

STEP 1 > VRF

Specify Tenant VRF

Name:

Alias:

Description:

Policy Control Enforcement Preference: Enforced Unenforced

Policy Control Enforcement Direction: Egress Ingress

BD Enforcement Status:

End Point Retention Policy:
This policy only applies to remote L3 entries

Monitoring Policy:

DNS Labels:
enter names separated by comma

Route Tag Policy:

Create A Bridge Domain:

Configure BGP Policies:

Configure OSPF Policies:

Configure EIGRP Policies:

Create Bridge Domains

To create Bridge domains, complete the following steps:

1. In the APIC GUI, select Tenants > common.
2. In the left pane, expand Tenant common and Networking.

3. Right-click the Bridge Domain and select Create Bridge Domain.
4. Name the Bridge Domain as BD-FP-Common-IB-Mgmt
5. Select vrf-FP-Common-IB-MGMT from the VRF drop-down list.
6. Select Custom under Forwarding and enable the flooding as shown in the screenshot below.

Create Bridge Domain

STEP 1 > Main

1. Main

2. L3 Configurations

3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type: fc regular

VRF:

Forwarding:

L2 Unknown Unicast:

L3 Unknown Multicast Flooding:

Multi Destination Flooding:

ARP Flooding: Enabled

Clear Remote MAC Entries:

End Point Retention Policy:

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

7. Click Next.
8. Do not change any configuration on next screen (L3 Configurations). Select Next.
9. No changes are needed Advanced/Troubleshooting. Click FINISH.

Create Application Profile

To create an Application profile, complete the following steps:



When the APIC-Controlled Microsoft Virtual Switch is used later in this document, for port-group naming, the Tenant name, Application Profile name, and EPG name are concatenated together to form the port-group name. Since the port-group name must be less than 64 characters, short names are used here for Tenant, Application Profile, and EPG.

1. In the APIC GUI, select Tenants > common.
2. In the left pane, expand Tenant common and Application Profiles.

3. Right-click the Application Profiles and select Create Application Profiles.
4. Enter MS-IB-MGMT as the name of the application profile.

Create Application Profile



Specify Tenant Application Profile

Name:

Alias:

Description:

Tags: ▼
enter tags separated by comma

Monitoring Policy: ▼

EPGs

Name	Alias	BD	Domain	Static Path	Static Path VLAN	Provided Contract	Consumed Contract

5. Click Submit.

Create EPG

To create EPG, complete the following steps:

1. Expand the MS-IB-MGMT Application Profile and right-click the Application EPGs.
2. Select Create Application EPG.
3. Enter MS-Core-Services as the name of the EPG.
4. Select BD-FP-Common-IB-MGMT from the drop-down list for Bridge Domain.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description:

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

5. Click Finish.

Set Domains

To set Domains, complete the following steps:

1. Expand the newly create EPG and click Domains.
2. Right-click Domains and select Add L2 External Domain Association.
3. Select the FP-Mgmt-Sw as the L2 External Domain Profile.

Add L2 External Domain Association



Choose the L2 External domain to associate

L2 External Domain Profile:  

Cancel

Submit

4. Click Submit.
5. Right-click Domains and select Add L2 External Domain Association.
6. Select the UCS as the L2 External Domain Profile.
7. Click Submit.

Set Static Ports

To set Static Ports, complete the following steps:

1. In the left pane, right-click Static Ports.
2. Select Deploy Static EPG on PC, VPC, or Interface.
3. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for FP-Mgmt-Sw configured earlier.
4. Enter the IB-MGMT VLAN under Port Encap.
5. Change Deployment Immediacy to Immediate.
6. Set the Mode to Trunk.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Port Encap (or Secondary VLAN for Micro-Seg):
Integer Value

Deployment Immediacy: Immediate On Demand

Primary VLAN for Micro-Seg:
Integer Value

Mode: Trunk Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

Cancel

Submit

7. Click Submit.
8. In the left pane, right-click Static Ports.
9. Select Deploy Static EPG on PC, VPC, or Interface.
10. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for UCS Fabric Interconnect A configured earlier.
11. Enter the UCS Core-Services VLAN under Port Encap.



This VLAN should be a different VLAN than the one entered above for the Management Switch.

12. Change Deployment Immediacy to Immediate.
13. Set the Mode to Trunk.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel **Virtual Port Channel**

Path:

Port Encap (or Secondary VLAN for Micro-Seg): VLAN Integer Value

Deployment Immediacy: **Immediate** On Demand

Primary VLAN for Micro-Seg: VLAN Integer Value

Mode: **Trunk** Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

Cancel Submit

14. Click Submit.
15. In the left pane, right-click Static Ports.
16. Select Deploy Static EPG on PC, VPC, or Interface.
17. In the next screen, for the Path Type, select Virtual Port Channel and from the Path drop-down list, select the VPC for UCS Fabric Interconnect B configured earlier.
18. Enter the UCS MS-IB-MGMT VLAN under Port Encap.



This VLAN should be a different VLAN than the one entered above for the Management Switch.

19. Change Deployment Immediacy to Immediate.
20. Set the Mode to Trunk.
21. Click Submit.

Create EPG Subnet

A subnet gateway for this Core Services EPG provides Layer 3 connectivity to Tenant subnets. To create a EPG Subnet, complete the following steps:

1. In the left pane, right-click Subnets and select Create EPG Subnet.
2. In CIDR notation, enter an IP address and subnet mask to serve as the gateway within the ACI fabric for routing between the Core Services subnet and Tenant subnets. This IP should be different than the IB-MGMT subnet gateway. In this lab validation, 10.1.118.1/24 is the IB-MGMT subnet gateway and is configured externally to the ACI fabric. 10.1.118.254/24 will be used for the EPG subnet gateway. Set the Scope of the subnet to Shared between VRFs.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy: ▼

Cancel

Submit

3. Click Submit to create the Subnet.

Create Provided Contract

To create Provided Contract, complete the following steps:

1. In the left pane, right-click Contracts and select Add Provided Contract.
2. In the Add Provided Contract window, select Create Contract from the drop-down list.
3. Name the Contract FP-Allow-Common-Core-Services.
4. Set the scope to Global.
5. Click + to add a Subject to the Contract.



The following steps create a contract to allow all the traffic between various tenants and the common management segment. You are encouraged to limit the traffic by setting restrictive filters.

6. Name the subject Allow-All-Traffic.
7. Click + under Filter Chain to add a Filter.
8. From the drop-down Name list, select common/default.
9. In the Create Contract Subject window, click Update to add the Filter Chain to the Contract Subject.

Create Contract Subject



Specify Identity Of Subject

Name:

Alias:

Description:

Target DSCP:

Apply Both Directions:

Reverse Filter Ports:

Filter Chain

Filters			+
Name	Directives		
common/default	none		

L4-L7 SERVICE GRAPH

Service Graph:

PRIORITY

QoS:

Cancel

OK

10. Click OK to add the Contract Subject.



The Contract Subject Filter Chain can be modified later.

11. Click Submit to finish creating the Contract.

Create Contract



Specify Identity Of Contract

Name:

Alias:

Scope: ▼

QoS Class: ▼

Target DSCP: ▼

Description:

Tags: ▼
enter tags separated by comma

Subjects: 🗑️ +

Name	Description
Allow-All-Traffic	

12. Click Submit to finish adding a Provided Contract.

Add Provided Contract



Select a contract

Contract: 

QoS:

Contract Label:

Subject Label:

Cancel

Submit

Create Security Filters in Tenant Common

To create Security Filters for SMB/CIFS with NetApp Storage and for iSCSI, complete the following steps. This section can also be used to set up other filters necessary to your environment.

1. In the APIC GUI, at the top select Tenants > common.
2. On the left, expand Tenant common, Security Policies, and Filters.
3. Right-click Filters and select Create Filter.
4. Name the filter Allow-All.
5. Click the + sign to add an Entry to the Filter.
6. Name the Entry Allow-All and select EtherType IP.
7. Leave the IP Protocol set at Unspecified.
8. Click UPDATE to add the Entry.

Create Filter



Specify the Filter Identity

Name:

Alias:

Description: optional

Entries: 🗑️ +

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
							From	To	From	To	
Allow-All		IP		unspecified	False	False					

Cancel
Submit

9. Click SUBMIT to complete adding the Filter.
10. Right-click Filters and select Create Filter.
11. Name the filter NetApp-SMB.
12. Click the + sign to add an Entry to the Filter.
13. Name the Entry tcp-445 and select EtherType IP.
14. Select the tcp IP Protocol and enter 445 for From and To under the Destination Port / Range by back-spacing over Unspecified and entering the number.
15. Click UPDATE to add the Entry.
16. Click the + sign to add another Entry to the Filter.
17. Name the Entry udp-445 and select EtherType IP.
18. Select the tcp IP Protocol and enter 445 for From and To under the Destination Port / Range by back-spacing over Unspecified and entering the number.
19. Click UPDATE to add the Entry.

Create Filter



Specify the Filter Identity

Name: NetApp-SMB

Alias:

Description: optional

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
							From	To	From	To	
tcp-445		IP		tcp	False	False	unspecified	unspecified	445	445	Unspecified
udp-445		IP		udp	False	False	unspecified	unspecified	445	445	

Cancel

Submit

20. Click SUBMIT to complete adding the Filter.

21. Right-click Filters and select Create Filter.

22. Name the filter iSCSI.

23. Click the + sign to add an Entry to the Filter.

24. Name the Entry iSCSI and select EtherType IP.

25. Select the TCP IP Protocol and enter 3260 for From and To under the Destination Port / Range by back-spacing over Unspecified and entering the number.

26. Click UPDATE to add the Entry.

Create Filter



Specify the Filter Identity

Name: iSCSI

Alias:

Description: optional

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range		Destination Port / Range		TCP Session Rules
							From	To	From	To	
iSCSI		IP		tcp	False	False	unspecified	unspecified	3260	3260	Unspecified

Cancel

Submit

27. Click SUBMIT to complete adding the Filter.



By adding these Filters to Tenant common, they can be used from within any Tenant in the ACI Fabric

Deploy FP-Foundation Tenant

This section details the steps for creating the Foundation Tenant in the ACI Fabric. This tenant will host infrastructure connectivity for the compute (Microsoft Hyper-V on UCS nodes) and the storage environments. To deploy the FP-Foundation Tenant, complete the following steps:

1. In the APIC GUI, select Tenants > Add Tenant.
2. Name the Tenant as FP-Foundation.
3. For the VRF Name, enter FP-Foundation. Keep the check box “Take me to **this tenant when I click finish**” checked.

Create Tenant



Specify tenant details

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

GUID:

Provider	GUID	Account Name

Monitoring Policy:

Security Domains:

Name	Description

VRF Name:

Take me to this tenant when I click finish

Cancel

Submit

- Click Submit to finish creating the Tenant.

Create Bridge Domain

To create a Bridge Domain, complete the following steps:

- In the left pane, expand Tenant FP-Foundation and Networking.
- Right-click Bridge Domains and select Create Bridge Domain.
- Name the Bridge Domain BD-FP-Foundation-Internal.
- Select FP-Foundation from the VRF drop-down list.
- Select Custom under Forwarding and enable the flooding.

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type: fc regular

VRF:

Forwarding:

L2 Unknown Unicast:

L3 Unknown Multicast Flooding:

Multi Destination Flooding:

ARP Flooding: Enabled

Clear Remote MAC Entries:

End Point Retention Policy:
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

6. Click Next.
7. Do not change any configuration on the next screen (L3 Configurations). Select Next.
8. No changes are needed for Advanced/Troubleshooting. Click Finish to finish creating Bridge Domain.

Create Application Profile for IB-Management Access

To create an application profile for IB-Management Access, complete the following steps:

1. In the left pane, expand tenant FP-Foundation, right-click Application Profiles and select Create Application Profile.
2. Name the Application Profile as AP-IB-MGMT and click Submit to complete adding the Application Profile.

Create EPG for IB-MGMT Access

This EPG will be used for Hyper-V hosts and management virtual machines that are in the IB-MGMT subnet, but that do not provide ACI fabric Core Services. For example, AD server VMs could be placed in the Core Services EPG defined earlier to provide DNS services to tenants in the Fabric. The SCVMM VM can be placed in the IB-MGMT EPG; it will have access to the Core Services VMs, but will not be reachable from Tenant VMs.

To create EPG for IB-MGMT access, complete the following steps:

1. In the left pane, expand the Application Profiles and right-click the AP-IB-MGMT EPG and select Create Application EPG.
2. Name the EPG EPG-IB-MGMT.
3. From the Bridge Domain drop-down list, select Bridge Domain BD-FP-Common-IB-MGMT from Tenant common.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description:

Tags: enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

4. Click Finish to complete creating the EPG.
5. In the left menu, expand the newly created EPG, right-click Domains and select Add Physical Domain Association.
6. Select the NetApp-AFF Physical Domain Profile and click Submit.
7. In the left menu, right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.
8. Select the Virtual Port Channel Path Type, then for Path select the vPC for the first NetApp AFF storage controller.
9. For Port Encap leave VLAN selected and fill in the storage IB-MGMT VLAN ID.

10. Set the Deployment Immediacy to Immediate and click Submit.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Port Encap (or Secondary VLAN for Micro-Seg):
Integer Value

Deployment Immediacy: Immediate On Demand

Primary VLAN for Micro-Seg:
Integer Value

Mode: Trunk Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

11. Repeat steps 7-10 to add the Static Port mapping for the second NetApp AFF storage controller.
12. In the left menu, right-click Domains and select Add L2 External Domain Association.
13. Select the UCS L2 External Domain Profile and click Submit.
14. In the left menu, right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.
15. Select the Virtual Port Channel Path Type, then for Path select the vPC for the first UCS Fabric Interconnect.
16. For Port Encap leave VLAN selected and fill in the UCS IB-MGMT VLAN ID.

17. Set the Deployment Immediacy to Immediate and click Submit.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel **Virtual Port Channel**

Path:

Port Encap (or Secondary VLAN for Micro-Seg): VLAN Integer Value

Deployment Immediacy: **Immediate** On Demand

Primary VLAN for Micro-Seg: VLAN Integer Value

Mode: **Trunk** Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

Cancel Submit

18. Repeat steps 7-10 to add the Static Port mapping for the second UCS Fabric Interconnect.

19. In the left menu, right-click Contracts and select Add Consumed Contract.

20. From the drop-down list for the Contract, select FP-Allow-Common-Core-Services from Tenant common.

Add Consumed Contract



Select a contract

Contract:

QoS:

Contract Label:

Subject Label:

Cancel

Submit

21. Click Submit.

This EPG will be utilized to provide Hyper-V hosts as well as the VMs that do not provide Core Services access to the existing in-band management network.

Create Application Profile for Host Connectivity

To create an application profile for host connectivity, complete the following steps:

1. In the left pane, under the Tenant FP-Foundation, right-click Application Profiles and select Create Application Profile.
2. Name the Profile AP-Host-Connectivity and click Submit to complete adding the Application Profile.

The following EPGs and the corresponding mappings will be created under this application profile.



Refer to [Error! Reference source not found.](#) for the information required during the following configuration. Items marked by { } will need to be updated according to Table 6 . Note that since all storage interfaces on a single Interface Group on a NetApp AFFA300 share the same MAC address, that different bridge domains must be used for each storage EPG.

Table 6 EPGs and mappings for AP-Host-Connectivity

EPG Name	Bridge Domain	Domain	Static Port - Compute	Static Port - Storage
EPG-MS-LVMN	BD-FP-Foundation-Internal	L2 External: UCS	VPC for all UCS Fis VLAN 906	N/A
EPG-MS-Clust	BD-FP-Foundation-Internal	L2 External: UCS	VPC for all UCS Fis VLAN 907	N/A

EPG Name	Bridge Domain	Domain	Static Port - Compute	Static Port - Storage
EPG-Infra-iSCSI-A	BD-FP-Foundation-iSCSI-A	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS Fis VLAN 3113	VPC for all NetApp AFFs VLAN 3013
EPG-Infra-iSCSI-B	BD-FP-Foundation-iSCSI-B	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS Fis VLAN 3123	VPC for all NetApp AFFs VLAN 3023
EPG-Infra-SMB	BD-FP-Foundation-SMB	L2 External: UCS Physical: NetApp-AFF	VPC for all UCS Fis VLAN 3153	VPC for all NetApp AFFs VLAN 3053



The MS-LVMN, MS-Cluster and MS-Infra-SMB VLANs are configured in Cisco UCS here and will be in place if needed on the manually created Hyper-V Virtual Switch. The EPG should be configured; it is not necessary to configure the actual VLAN or UCS static port mapping, but you can configure these without any negative effects.

Create Bridge Domains and EPGs

To create bridge domains and EPGs, complete the following steps:

1. For each row in the table above, if the Bridge Domain does not already exist, in the left pane, expand Networking > Bridge Domains.
2. Right-click Bridge Domains and select Create Bridge Domain.
3. Name the Bridge Domain {BD-FP-Foundation-iSCSI-A}.
4. Select the FP-Foundation VRF.
5. Select Custom for Forwarding and setup forwarding as shown in the screenshot.

Create Bridge Domain



STEP 1 > Main

1. Main

2. L3 Configurations

3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type:

VRF:

Forwarding:

L2 Unknown Unicast:

L3 Unknown Multicast Flooding:

Multi Destination Flooding:

ARP Flooding: Enabled

Clear Remote MAC Entries:

End Point Retention Policy:

This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

Previous

Cancel

Next

6. Click Next.
7. Do not change any configuration on the next screen (L3 Configurations). Select Next.
8. No changes are needed for Advanced/Troubleshooting. Click Finish to finish creating Bridge Domain.
9. In the left pane, expand Application Profiles > AP-Host-Connectivity. Right-click on Application EPGs and select Create Application EPG.
10. Name the EPG {EPG-MS-LVMN}.
11. From the Bridge Domain drop-down list, select the Bridge Domain from the table.
12. Click Finish to complete creating the EPG.
13. In the left pane, expand the Application EPGs and EPG {EPG-LVMN}.
14. Right-click Domains and select Add L2 External Domain Association.
15. From the drop-down list, select the previously defined {UCS} L2 External Domain Profile.

Add L2 External Domain Association



Choose the L2 External domain to associate

L2 External Domain Profile: 

Cancel

Submit

16. Click Submit to complete the L2 External Domain Association.
17. Repeat the Domain Association steps (6-9) to add appropriate EPG specific domains from Table 7 .
18. Right-click Static Ports and select Deploy EPG on PC, VPC, or Interface.
19. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.
20. From the drop-down list, select the appropriate VPCs.
21. Enter VLAN from Error! Reference source not found. {906} for Port Encap.
22. Select Immediate for Deployment Immediacy and for Mode select Trunk.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel **Virtual Port Channel**

Path:

Port Encap (or Secondary VLAN for Micro-Seg): VLAN Integer Value

Deployment Immediacy: **Immediate** On Demand

Primary VLAN for Micro-Seg: VLAN Integer Value

Mode: **Trunk** Access (802.1P) Access (Untagged)

IGMP Snoop Static Group: 🗑️ +

Group Address	Source Address

Cancel Submit

23. Click Submit to complete adding the Static Path Mapping.

24. Repeat the above steps to add all the Static Path Mappings for the EPG listed in Table 6 .

The screenshot shows the APIC interface for Tenant FP-Foundation. The 'Static Ports' section is active, displaying a table of configurations for Node: Pod-1.

Path	Primary VLAN for Micro-Seg	Port Encap (or Secondary VLAN for Micro-Seg)	Deployment Immediacy	Mode
Pod-1/Node-105-106/A02-6332-A-...	vlan-3152	vlan-3152	Immediate	Trunk
Pod-1/Node-105-106/A02-6332-B-...	vlan-3152	vlan-3152	Immediate	Trunk
Pod-1/Node-105-106/A02-AFFA300...	vlan-3052	vlan-3052	Immediate	Trunk
Pod-1/Node-105-106/A02-AFFA300...	vlan-3052	vlan-3052	Immediate	Trunk

Table 7 EPGs and Subnets for AP-Host-Connectivity

EPG Name	Subnet
EPG-MS-LVMN	192.168.96.254/24
EPG-MS-Clust	192.168.97.254/24
EPG-Infra- iSCSI-A	192.168.12.254/24
EPG-Infra- iSCSI-B	192.168.22.254/24
EPG-Infra-SMB	192.168.53.254/24

25. On the left under the EPG, right-click Subnets and select Create EPG Subnet.

26. In the Create EPG Subnet window, enter the Subnet from Table 7 as the Default Gateway IP.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy: ▼

Cancel Submit

27. Click Submit to complete adding the subnet.

28. Repeat the above steps to complete adding the EPGs and subnets in Table 6 and Table 7 .

Storage Configuration



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 bootarg.tmgr.disable_pit_hp 1
```

For more information about the workaround, see: <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 All Flash FAS A300 Controllers

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.

To access the [HWU](#) application to view the System Configuration guides, complete the following steps:

1. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. 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 8 lists the information needed to configure two ONTAP nodes. Customize the cluster detail values with the information applicable to your deployment.

Table 8 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 `e0M` 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 02 configuration while the disks for node 01 are zeroing.

Configure Node 02

To configure node 02, 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 e0M 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.

```
<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

To set up a node, complete the following steps:

1. 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.
2. Follow the prompts to set up node 01:

```
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the 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:
```

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

Table 9 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 Detail	Cluster Detail Value
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 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.

4. 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

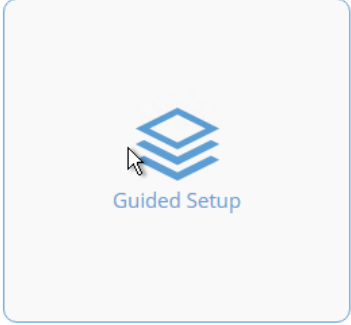
i For information related to setting up the cluster, [click here](#)

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


5. In the Cluster screen, complete the following steps:
 - a. Enter the cluster and node names.
 - b. Select the cluster configuration.
 - c. Enter and confirm the password.
 - d. 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.

6. Click Submit.

7. On the network page, complete the following sections:
 - a. Cluster Management
 - Enter the IP address, netmask, gateway and port details.
 - b. Node Management
 - Enter the node management IP addresses and port details for all the nodes.
 - c. Service Processor Management
 - Enter the IP addresses for all the nodes.
 - d. DNS Details
 - Enter the DNS domain names and server address.
 - e. NTP Details
 - Enter the primary and alternate NTP server.
8. Click Submit.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



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.

	IP Address	Netmask	Gateway (Optional)	Port
Cluster Management	<input type="text" value="192.168.156.60"/>	<input type="text" value="255.255.255.0"/>	<input type="text" value="192.168.156.1"/>	<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.			
	<input type="text" value="bb04-aff300-1"/>	<input type="text" value="192.168.156.61"/>	<input type="text" value="e0M"/>	<input type="text" value=""/>
	<input type="text" value="bb04-aff300-2"/>	<input type="text" value="192.168.156.62"/>	<input type="text" value="e0M"/>	<input type="text" value=""/>
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.			
	<input type="text" value="bb04-aff300-1"/>	<input type="text" value="192.168.156.58"/>		
	<input type="text" value="bb04-aff300-2"/>	<input type="text" value="192.168.156.59"/>		

DNS Details

DNS Domain Names

DNS Server IP Address

NTP Details

Primary NTP Server

Alternative NTP Server (Optional)

9. 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

10. Click Submit.
11. 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
```

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 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 `fc 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 `fc 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 -vservers <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, 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 <st-node01>:e2a,<st-node01>:e2e,<st-node02>:e2a,<st-node01>: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.

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



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

```
cluster ha show
```

4. Continue with step 5 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
```

5. 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:

1. 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
```

2. 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

If NTP was not configured during guided setup, it can be configured via the CLI as follows:

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 <ccyyymmddhhmm.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-<MS-IB-MGMT-VLAN> -mtu 9000
broadcast-domain create -broadcast-domain Infra_MS_SMB -mtu 9000
```

If using iSCSI, create two iSCSI broadcast domains with an MTU of 9000, via the following command:

```
broadcast-domain create -broadcast-domain iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain iSCSI-B -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-<MS-IB-MGMT-VLAN> -ports <st-node01>:a0a-<MS-IB-MGMT-VLAN>,<st-node02>:a0a-<MS-IB-MGMT-VLAN>
```

If using iSCSI, create two iSCSI VLANs and add them to the corresponding broadcast domains:

```
network port vlan create -node <st-node01> -vlan-name a0a-<iSCSI-A-VLAN>
network port vlan create -node <st-node02> -vlan-name a0a-<iSCSI-A-VLAN>
broadcast-domain add-ports -broadcast-domain iSCSI-A -ports <st-node01>:a0a-<iSCSI-A-VLAN>,<st-node02>:a0a-<iSCSI-A-VLAN>
network port vlan create -node <st-node01> -vlan-name a0a-<iSCSI-B-VLAN>
network port vlan create -node <st-node02> -vlan-name a0a-<iSCSI-B-VLAN>
broadcast-domain add-ports -broadcast-domain iSCSI-B -ports <st-node01>:a0a-<iSCSI-B-VLAN>,<st-node02>:a0a-<iSCSI-B-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
```

- Remove the unused data protocols (NFS and NDMP) from the SVM.

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

- Add the two data aggregates to the Infra-MS-SVM aggregate list.

```
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:

- 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.

- 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
```

- Depending on how your network is configured, you may need to add a default route for the SVM to reach the Active Directory domain controller via the in-band management network. Here is how that would look:

```
a02-affa300::> net route show
(network route show)
Vserver          Destination      Gateway          Metric
-----
a02-affa300
                0.0.0.0/0       192.168.1.254   20

a02-affa300::> net route create -vs Infra-MS-SVM -destination 0.0.0.0/0 -gateway 10.1.118.1
(network route create)

a02-affa300::> net route show
(network route show)
Vserver          Destination      Gateway          Metric
-----
Infra-MS-SVM
                0.0.0.0/0       10.1.118.1      20
a02-affa300
                0.0.0.0/0       192.168.1.254   20
2 entries were displayed.
```

```
a02-affa300::>
```

4. 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 Service(s)

If the deployment is using FCP, create the FCP service on each SVM using the following command. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM.

```

fcpx create -vserver Infra-MS-SVM

fcpx show

```

If the deployment is using iSCSI, create the iSCSI service on each SVM using the following command. This command also starts the iSCSI service and sets the IQN for the SVM.

```

iscsi create -vserver Infra-MS-SVM

iscsi show

```



The licenses for FCP and iSCSI must be installed before the services can be started. If the license(s) weren't installed during cluster setup, install them before this step.

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 previous 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 3 (<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 smb -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 iscsi_datastore_1 -aggregate aggr1_node01 -size 500GB -
state online -policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-MS-SVM -volume iscsi_datastore_2 -aggregate aggr1_node02 -size 500GB -
state online -policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-MS-SVM -volume witness -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 shares, run the following commands:

```
cifs share access-control create -vserver Infra-MS-SVM -share infra_share_1_Share -user-or-group
Flexpod\Administrator -user-group-type windows -permission Full_Control

cifs share access-control create -vserver Infra-MS-SVM -share infra_share_2_Share -user-or-group
Flexpod\Administrator -user-group-type windows -permission Full_Control

cifs share access-control create -vserver Infra-MS-SVM -share infra_share_1_Share -user-or-group
flexpod\<FC_Host_1>$ -user-group-type windows -permission Full_Control

cifs share access-control create -vserver Infra-MS-SVM -share infra_share_2_Share -user-or-group
flexpod\<FC_Host_1>$ -user-group-type windows -permission Full_Control

cifs share access-control create -vserver Infra-MS-SVM -share infra_share_1_Share -user-or-group
flexpod\<FC_Host_2>$ -user-group-type windows -permission Full_Control

cifs share access-control create -vserver Infra-MS-SVM -share infra_share_2_Share -user-or-group
flexpod\<FC_Host_2>$ -user-group-type windows -permission Full_Control
```

Create Gold Management Host Boot LUN

To create one boot LUN, run the following commands:

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun MGMT-Win2016-Gold -size 200GB -ostype
windows_2008 -space-reserve disabled
```

Create Witness and iSCSI Datastore LUNs

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 -lun witness -size 1GB -ostype windows_2008 -space-
reserve disabled

lun create -vserver Infra-MS-SVM -volume iscsi_datastore_1 -lun iscsi_datastore_1 -size 500GB -ostype
windows_2008 -space-reserve disabled

lun create -vserver Infra-MS-SVM -volume iscsi_datastore_2 -lun iscsi_datastore_2 -size 500GB -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
efficiency modify -vserver Infra-MS-SVM -volume iscsi_datastore_1 -schedule sun-sat@0
efficiency modify -vserver Infra-MS-SVM -volume iscsi_datastore_2 -schedule sun-sat@0
```

Create SAN LIFs

If using FCP, 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
```

If using iSCSI, 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-<iSCSI-A-VLAN> -address <iscsi_lif01a_ip> -netmask
<iscsi_lif01a_mask> -status-admin up

network interface create -vserver Infra-MS-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -
home-node <st-node01> -home-port a0a-<iSCSI-B-VLAN> -address <iscsi_lif01b_ip> -netmask
<iscsi_lif01b_mask> -status-admin up

network interface create -vserver Infra-MS-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -
home-node <st-node02> -home-port a0a-<iSCSI-A-VLAN> -address <iscsi_lif02a_ip> -netmask
<iscsi_lif02a_mask> -status-admin up

network interface create -vserver Infra-MS-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -
home-node <st-node02> -home-port a0a-<iSCSI-B-VLAN> -address <iscsi_lif02b_ip> -netmask
<iscsi_lif02b_mask> -status-admin up
```

Create SMB LIFs

To create SMB LIFs, 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
```



The two SMB LIF IPs need to be entered in the Domain DNS server with the CIFS Server name (infra-cifs) created above. Create two host records with the different IPs and same host name.

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 iSCSI boot to the NetApp AFF through the Cisco ACI Fabric. An alternative Fibre Channel (FC) boot delta configuration is covered in the appendix of this document. If FC boot is desired, execute the following procedure not executing iSCSI-related steps and then execute the procedure in the appendix.



The MS-LVMN, MS-Cluster and MS-Infra-SMB VLANs are configured here in the UCS and will be in place if needed on the manually created Hyper-V Virtual Switch. It is not necessary to configure the actual VLAN or add it to the vNIC interfaces, but you can configure these without any negative effects.

Perform Initial Setup

This section provides detailed steps to configure 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 `http://<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>`.

Cisco UCS Manager Initial Setup

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>
```

- Using a supported web browser, connect to `http://<ucsb-mgmt-ip>`, accept the security prompts, and click the **'Express Setup'** link under HTML.
- Under System setup, enter the Admin Password entered above and click Submit.
- Enter `<ucsb-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 Cisco 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 3.2(1d)

This document assumes the use of Cisco UCS 3.2(1d). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.2(1d), 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 Cisco UCS Manager. Configuring Call Home will accelerate resolution of support cases. To configure Call Home, complete the following steps:

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

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 the LAN icon 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 :	<input type="text" value="192.168.1.209"/>	Size :	<input type="text" value="16"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="192.168.1.254"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

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 the Admin icon 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 <global-ntp-ip> and click OK. Click OK on the confirmation.

Add NTP Server



NTP Server :



8. Add any other NTP servers as necessary.

Edit Policy to Automatically Discover Server Ports

If the UCS Port Auto-Discovery Policy is enabled, server ports will be discovered automatically. To enable the Port Auto-Discovery Policy, complete the following steps:

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

- Under Policies, select the Port Auto-Discovery Policy tab.
- Under Properties, set Auto Configure Server Port to Enabled.

Equipment

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

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

Actions

Use Global

Properties

Owner : **Local**

Auto Configure Server Port : Disabled Enabled

- Click Save Changes.
- Click OK.

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:

- In Cisco UCS Manager, click the Equipment icon on the left and select Equipment in the second list.
- 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, and the Multicast Hardware Hash setting appears, set the Multicast Hardware Hash setting to Enabled.

Equipment

- [Main Topology View](#)
[Fabric Interconnects](#)
[Servers](#)
[Thermal](#)
[Decommissioned](#)
[Firmware Management](#)
[Policies](#)
- [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.

Verify Server and Enable Uplink Ports

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

1. In Cisco UCS Manager, click the Equipment icon on the left.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand and select Ethernet Ports.
4. On the right, verify that the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers are configured as Server ports. If any Server ports are not configured correctly, right-click them, and select “Configure as Server Port.” Click Yes to confirm server ports and click OK.



In lab testing, for C220M4 servers with VIC 1385 PCIE cards, it has been necessary to manually configure Server ports.


5. Verify that the ports connected to the chassis, C-series servers and Cisco FEX are now configured as server ports.
6. Select the ports that are connected to the Cisco Nexus 9332 switches, right-click them, and select Configure as Uplink Port.




The last 6 ports (ALE) of the UCS 6332 and UCS 6332-16UP FIs require the use of active (optical) or AOC cables when connected to a Nexus 9332. It may also be necessary to remove and reinsert these cables on the switch end to get them to come up the first time.

7. Click Yes to confirm uplink ports and click OK.
8. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

- Expand and select Ethernet Ports.
- On the right, verify that the ports that are connected to the chassis, Cisco FEX, and direct connect UCS C-Series servers are configured as Server ports. If any Server ports are not configured correctly, right-click them, and select “Configure as Server Port.” Click Yes to confirm server ports and click OK.

 In lab testing, for C220M4 servers with VIC 1385 PCIE cards, it has been necessary to manually configure Server ports.

- 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 (ALE) of the UCS 6332 and UCS 6332-16UP FIs require the use of active (optical) or AOC cables when connected to a Nexus 9332. It may also be necessary to remove and reinsert these cables on the switch end to get them to come up the first time.

- 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 the Equipment icon on the left.
- Expand Chassis and select each chassis that is listed.
- Right-click each chassis and select Acknowledge Chassis.

Acknowledge Chassis



Are you sure you want to acknowledge Chassis 1 ?

This operation will rebuild the network connectivity between the Chassis and the Fabrics it is connected to. Currently there are 2 active links to Fabric A and there are 2 active links to Fabric B.

Yes

No

- Click Yes and then click OK to complete acknowledging the chassis.
- If Nexus 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.

Re-Acknowledge Any Inaccessible C-Series Servers

If any C-Series servers show an Inaccessible Status, complete the following steps:

1. In Cisco UCS Manager, click the Equipment icon on the left.
2. Under Equipment > Rack Mounts, expand Servers.
3. If any of the servers have a status of Inaccessible, right-click the server and select Server Maintenance. Select Re-Acknowledge and click OK. Click Yes and OK. The server should then be Discovered properly.

Create Uplink Port Channels to Cisco Nexus 9332 Switches

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

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



In this procedure, two port channels are created: one from fabric A to both Cisco Nexus 9332 switches and one from fabric B to both Cisco Nexus 9332 switches.

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 139 as the unique ID of the port channel.
6. Enter Po139-ACI 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. Expand Port Channels and select Port-Channel 139. Since the vPC has already been configured in the ACI fabric, this port channel should come up.
13. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
14. Right-click Port Channels.

15. Select Create Port Channel.
16. Enter 140 as the unique ID of the port channel.
17. Enter Po140-ACI as the name of the port channel.
18. Click Next.
19. Select the ports connected to the Nexus switches to be added to the port channel:
20. Click >> to add the ports to the port channel.
21. Click Finish to create the port channel.
22. Click OK.
23. Expand Port Channels and select Port-Channel 140. Since the vPC has already been configured in the ACI fabric, this port channel should come up.

Create an IQN Pool for iSCSI Boot

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

1. Select the SAN icon on the left.
2. Select Pools > root.
3. Right-click IQN Pools under the root organization.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-P001 for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter iqn.2010-11.com.flexpod for the Prefix
8. Select Sequential for Assignment Order.

1 Define Name and Description

2 Add IQN Blocks

Create IQN Suffix Pool

Name : IQN-Pool

Description :

Prefix : iqn.2010-11.com.flexpod

IQN Prefix must have the following format: **iqn.yyyy-mm.naming-authority**, where *naming-authority* is usually the reverse syntax of the Internet domain name of the naming authority.

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

9. Click Next.

10. Click Add.

11. Enter a name to identify the individual UCS host for the Suffix.

12. Enter 1 for the From field.

13. Specify a size of the IQN block sufficient to support the available server resources.

Create a Block of IQN Suffixes



Suffix :

From :

Size :



14. Click OK.

15. Click Finish and OK to complete creating the IQN pool.

Create iSCSI Boot IP Address Pools

To configure the necessary iSCSI IP Address pools for the Cisco UCS environment, complete the following steps:

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

1 Define Name and Description

2 Add IPv4 Blocks

3 Add IPv6 Blocks

Create IP Pool

Name : iSCSI-IP-Pool-A

Description :

Assignment Order : Default Sequential

< Prev Next > Finish Cancel

9. Click Next.

10. Click Add to add a Block of IPs to the pool.

11. Specify a starting IP address and subnet mask in the subnet for iSCSI boot on Fabric A. It is not necessary to specify the Default Gateway or DNS server addresses.

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

Create Block of IPv4 Addresses



From :	<input type="text" value="192.168.12.201"/>	Size :	<input type="text" value="16"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="0.0.0.0"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

13. Click OK.
14. Click Next.
15. Click Finish.
16. In the confirmation message, click OK.
17. Right-click IP Pools under the root organization.
18. Select Create IP Pool to create the IP pool.
19. Enter `iSCSI-IP-Pool-B` as the name of the second IP pool.
20. Optional: Enter a description for the IP pool.
21. Select Sequential for Assignment Order
22. Click Next.
23. Click Add to add a Block of IPs to the pool.
24. Specify a starting IP address and subnet mask in the subnet for iSCSI boot on Fabric B. It is not necessary to specify the Default Gateway or DNS server addresses.
25. Specify a size for the IP pool that is sufficient to support the available blade or server resources.

Create Block of IPv4 Addresses



From :	<input type="text" value="192.168.22.201"/>	Size :	<input type="text" value="16"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="0.0.0.0"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

26. Click OK.
27. Click Next.
28. Click Finish.
29. In the confirmation message, click OK.

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 the LAN icon 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 also embedded the cabinet number (A2) information giving us 00:25:B5:A2:0A: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 assuming that multiple vNICs can be configured on each server.

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

OK

Cancel

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, 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 also embedded the cabinet number (A2) information giving us 00:25:B5:A2:0A: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 the Servers icon on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter `UUID-Pool1` 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. Optionally, specify identifiers such as UCS location.
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 Hyper-V management 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 the Servers icon on the left.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter Hyper-V-MGMT-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 Hyper-V-MGMT-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 the LAN icon on the left.

 In this procedure, seven unique and 100 sequential VLANs are created. See Table 2

 Note that MS-LVMN, MS-Cluster and MS-Infra-SMB VLANs are configured here in the UCS and will be in place if needed on the manually created Hyper-V Virtual Switch. It is not necessary here to configure the actual VLAN or add it to the vNIC interfaces, but it does not hurt to configure these.

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 <2>.
8. Keep the Sharing Type as None.
9. Click OK and then click OK again.

Create VLANs



VLAN Name/Prefix :

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

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

VLAN IDs :

Sharing Type : None Primary Isolated Community

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-Core-Services` 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 UCS Core Services VLAN ID <318>.
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-IB-MGMT` as the name of the VLAN to be used for management traffic.
22. Keep the Common/Global option selected for the scope of the VLAN.
23. Enter the UCS In-Band management VLAN ID <418>.
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-Infra-SMB` as the name of the VLAN to be used for SMB File share.
29. Keep the Common/Global option selected for the scope of the VLAN.
30. Enter the UCS SMB File Share VLAN ID <3153>.
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-Infra-iSCSI-A` as the name of the VLAN to be used for UCS Fabric A iSCSI boot.
36. Keep the Common/Global option selected for the scope of the VLAN.

37. Enter the UCS Fabric A iSCSI boot VLAN ID <3112>.
38. Keep the Sharing Type as None.
39. Click OK, and then click OK again.
40. Right-click VLANs.
41. Select Create VLANs.
42. Enter `MS-Infra-iSCSI-B` as the name of the VLAN to be used for UCS Fabric B iSCSI boot.
43. Keep the Common/Global option selected for the scope of the VLAN.
44. Enter the UCS Fabric B iSCSI boot VLAN ID <3122>.
45. Keep the Sharing Type as None.
46. Click OK, and then click OK again.
47. Right-click VLANs.
48. Select Create VLANs.
49. Enter `MS-LVMN` as the name of the VLAN to be used for Live Migration.
50. Keep the Common/Global option selected for the scope of the VLAN.
51. Enter the Live Migration VLAN ID <906>.
52. Keep the Sharing Type as None.
53. Click OK, and then click OK again.
54. Right-click VLANs.
55. Select Create VLANs.
56. Enter `MS-Cluster` as the name of the VLAN to be used for Cluster communication network.
57. Keep the Common/Global option selected for the scope of the VLAN.
58. Enter the Cluster network VLAN ID <907>.
59. Keep the Sharing Type as None.
60. Click OK, and then click OK again.
61. Right-click VLANs.

62. Select Create VLANs.
63. Enter `ACI-system` as the name of the VLAN to be used for OpFlex communication to the Hyper-V Virtual Switch.
64. Keep the Common/Global option selected for the scope of the VLAN.
65. Enter the ACI System VLAN ID `<4093>`.



The ACI system VLAN ID can be determined by using ssh to connect to the APIC CLI and typing “`ifconfig | grep bond0`”. You should see a `bond0.xxxx` interface listed. The `xxxx` is the ACI system VLAN id.

66. Keep the Sharing Type as None.
67. Click OK, and then click OK again.
68. Right-click VLANs.
69. Select Create VLANs.
70. Enter `Virtual-Switch-Pool` as the prefix for this VLAN pool.
71. Keep the Common/Global option selected for the scope of the VLAN.
72. Enter a range of 100 VLANs for VLAN ID. `<1200-1299>` was used in this validation.
73. Keep the Sharing Type as None.
74. Click OK and then click OK again.

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 specify the UCS 3.2(1d) release for the Default firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers icon 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 3.2(1d) for both the Blade and Rack Packages.

Modify Package Versions



Blade Package :

Rack Package :

Service Pack :

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

Excluded Components:

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

- Click OK then click 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 the LAN icon 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 the Servers icon 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 :

Description :

Mode :

FlexFlash

FlexFlash State : Disable Enable

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

FlexFlash RAID Reporting State : Disable Enable

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 the LAN icon 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 the Servers icon 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.

- Click OK.

Create Power Control Policy



Name :

Description :

Fan Speed Policy :

Power Capping

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

No Cap cap

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

OK

Cancel

Create Server Pool Qualification Policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, 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.

- In Cisco UCS Manager, click the Servers icon on the left.
- 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 :	<input type="text" value="Xeon"/>	PID (RegEx) :	<input type="text" value="UCS-CPU-E52660E"/>
Min Number of Cores :	<input checked="" type="radio"/> Unspecified <input type="radio"/> select	Max Number of Cores :	<input checked="" type="radio"/> Unspecified <input type="radio"/> select
Min Number of Threads :	<input checked="" type="radio"/> Unspecified <input type="radio"/> select	Max Number of Threads :	<input checked="" type="radio"/> Unspecified <input type="radio"/> select
CPU Speed (MHz) :	<input checked="" type="radio"/> Unspecified <input type="radio"/> select	CPU Stepping :	<input checked="" type="radio"/> Unspecified <input type="radio"/> 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 the Servers icon on the left.
2. Select Policies > root.

3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `Virtual-Host` as the BIOS policy name.
6. Click OK.
7. Expand BIOS Policies and Select Virtual-Host.
8. On the right, change the Quiet Boot setting to Disabled.
9. Change CDN Control to Enabled.

Properties

Name : **Virtual-Host**

Description :

Owner : **Local**

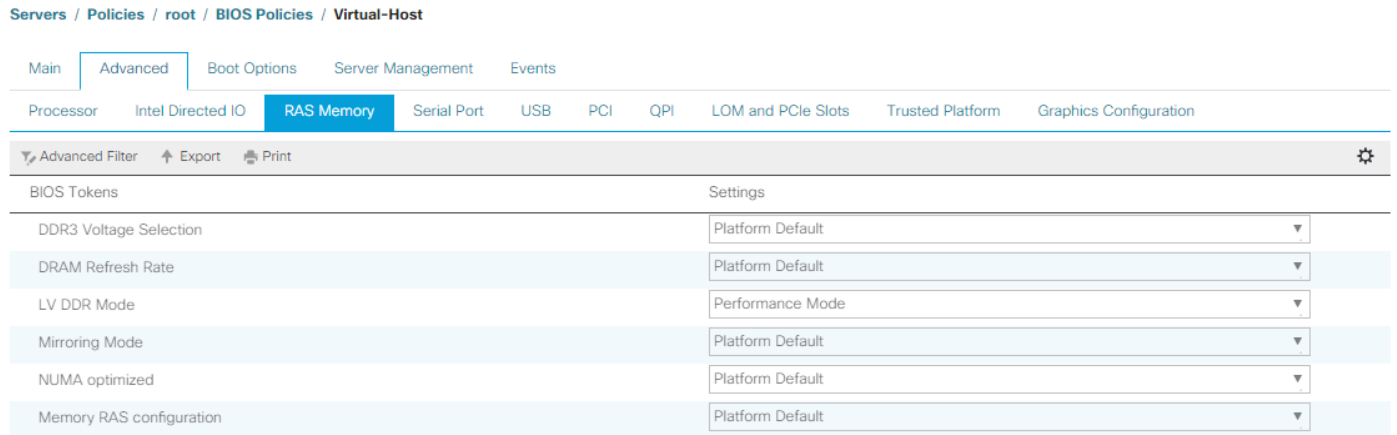
Reboot on BIOS Settings Change :

Advanced Filter Export Print ⚙️

BIOS Tokens	Settings
CDN Control	Enabled ▼
Front panel lockout	Platform Default ▼
POST error pause	Platform Default ▼
Quiet Boot	Disabled ▼
Resume on AC power loss	Platform Default ▼

10. Click Save Changes and OK.
11. Click the Advanced tab and then the Processor tab.
12. Set the following within the Processor tab:
 - a. Processor C State -> Disabled
 - b. Processor C1E -> Disabled
 - c. Processor C3 Report -> Disabled
 - d. Processor C7 Report -> Disabled
 - e. Energy Performance -> Performance
 - f. Frequency Floor Override -> Enabled
 - g. DRAM Clock Throttling -> Performance
13. Click Save Changes and OK.
14. Click the RAS Memory tab and select:

a. LV DDR Mode -> Performance-Mode



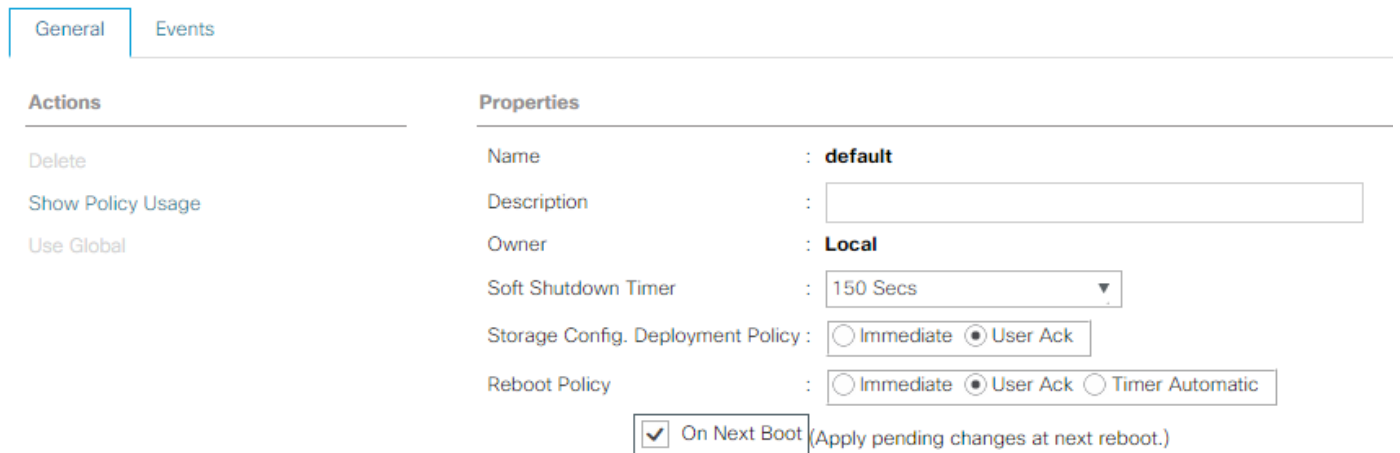
15. Click Save Changes and OK.

Update the Default Maintenance Policy

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

1. In Cisco UCS Manager, click the Servers icon 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 Policies / default



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 6 vNIC Templates will be created.

Create Infrastructure vNICs Templates

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Infra-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-Cluster, MS-Core-Services, MS-IB-MGMT, MS-Infra-SMB, and MS-MS-LVMN VLANs.
13. Set MS-IB-MGMT 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



Select	Name	Native VLAN
<input type="checkbox"/>	ACI-System	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Cluster	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Core-Services	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-IB-MGMT	<input checked="" type="radio"/>
<input type="checkbox"/>	MS-Infra-iSCSI-A	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool :

QoS Policy :

Network Control Policy :

Pin Group :

Stats Threshold Policy :

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy :

OK **Cancel**

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template Infra-Host-B:

1. Select the LAN icon on the left.
2. Select Policies > root.

3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter `Infra-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-Host-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 iSCSI Boot vNICs

To create iSCSI Boot vNICs, complete the following steps:

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Infra-iSCSI-A` as the vNIC template name.
6. Keep Fabric A selected.
7. Do not select the Enable Failover checkbox.
8. Select No Redundancy for Redundancy Type.
9. Under Target, make sure that only the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select the checkbox for `MS-Infra-iSCSI-A`.
12. Set `MS-Infra-iSCSI-A` as the native VLAN.
13. Select vNIC Name for the CDN Source.

14. For MTU, enter 9000.
15. In the MAC Pool list, select MAC-Pool-A.
16. In the Network Control Policy list, select Enable-CDP-LLDP.
17. Click OK to create the vNIC template.
18. Click OK.

Create the Infra-iSCSI-B template:

1. Select the LAN icon on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter `Infra-iSCSI-B` as the vNIC template name.
6. Select Fabric B.
7. Do not elect the Enable Failover checkbox.
8. Select No Redundancy for Redundancy Type.
9. Under Target, make sure that only the Adapter checkbox is selected.
10. Select Updating Template as the Template Type.
11. Under VLANs, select the checkbox for MS-Infra-iSCSI-B.
12. Set MS-Infra-iSCSI-B as the native VLAN.
13. Select vNIC Name for the CDN Source.
14. For MTU, enter 9000.
15. In the MAC Pool list, select MAC-Pool-B.
16. In the Network Control Policy list, select Enable-CDP-LLDP.
17. Click OK to create the vNIC template.
18. Click OK.

Create vNIC Templates for APIC-controlled Virtual Switch

To create vNIC templates for APIC-controlled virtual switch, complete the following steps:

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `APIC-MS-VS-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 `ACI-System` and all 100 Virtual-Switch-Pool VLANs.
13. Do not set a 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



VLANS | VLAN Groups

Advanced Filter | Export | Print |

Select	Name	Native VLAN
<input type="checkbox"/>	Virtual-Switch-Pool1294	<input type="radio"/>
<input checked="" type="checkbox"/>	Virtual-Switch-Pool1295	<input type="radio"/>
<input checked="" type="checkbox"/>	Virtual-Switch-Pool1296	<input type="radio"/>
<input checked="" type="checkbox"/>	Virtual-Switch-Pool1297	<input type="radio"/>
<input checked="" type="checkbox"/>	Virtual-Switch-Pool1298	<input type="radio"/>
<input checked="" type="checkbox"/>	Virtual-Switch-Pool1299	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

Connection Policies

Dynamic vNIC usNIC VMQ

Dynamic vNIC Connection Policy : ▼

OK **Cancel**

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template APIC-MS-VS-B:

1. Select the LAN icon on the left.
2. Select Policies > root.

3. Right-click vNIC Templates.
4. Select Create vNIC Template
5. Enter `APIC-MS-VS-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 `APIC-MS-VS-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 iSCSI Boot

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

1. In Cisco UCS Manager, click the LAN icon 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 `00-Infra-Host-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select `Infra-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 : Peer Name : vNIC Template : [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-Infra-Host-B as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select Infra-Host-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 another vNIC to the policy.
19. In the Create vNIC box, enter `02-Infra-iSCSI-A` as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select Infra-iSCSI-A.
22. In the Adapter Policy list, select Windows.
23. Click OK to add the vNIC to the policy.
24. Click the upper Add button to add another vNIC to the policy.
25. In the Create vNIC box, enter `03-Infra-iSCSI-B` as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select Infra-iSCSI-B.
28. In the Adapter Policy list, select Windows.
29. Click OK to add the vNIC to the policy.
30. Click the upper Add button to add another vNIC to the policy.
31. In the Create vNIC box, enter `04-APIC-MS-VS-A` as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select APIC-MS-VS-A.
34. In the Adapter Policy list, select Windows.
35. Click OK to add the vNIC to the policy.
36. Click the upper Add button to add another vNIC to the policy.
37. In the Create vNIC box, enter `05-APIC-MS-VS-B` as the name of the vNIC.
38. Select the Use vNIC Template checkbox.
39. In the vNIC Template list, select APIC-MS-VS-B.
40. In the Adapter Policy list, select Windows.
41. Click OK to add the vNIC to the policy.
42. Expand the Add iSCSI vNICs section.

43. Click the lower Add button to add an iSCSI boot vNIC to the policy.
44. In the Create iSCSI vNIC box, enter `iSCSI-Boot-A` as the name of the vNIC.
45. Select `02-Infra-iSCSI-A` for the Overlay vNIC.
46. Select the default iSCSI Adapter Policy.
47. `MS-Infra-iSCSI-A (native)` should be selected as the VLAN.
48. Do not select anything for MAC Address Assignment.

Create iSCSI vNIC



Name :

Overlay vNIC :

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

VLAN :

iSCSI MAC Address

MAC Address Assignment:

[Create MAC Pool](#)

49. Click OK to add the vNIC to the policy.
50. Click the lower Add button to add an iSCSI boot vNIC to the policy.
51. In the Create iSCSI vNIC box, enter `iSCSI-Boot-B` as the name of the vNIC.

- 52. Select 03-Infra-iSCSI-B for the Overlay vNIC.
- 53. Select the default iSCSI Adapter Policy.
- 54. MS-Infra-iSCSI-B (native) should be selected as the VLAN.
- 55. Do not select anything for MAC Address Assignment.
- 56. Click OK to add the vNIC to the policy.

Create LAN Connectivity Policy



Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 05-APIC-MS-VS-B	Derived	
vNIC 04-APIC-MS-VS-A	Derived	
vNIC 03-Infra-iSCSI-B	Derived	
vNIC 02-Infra-iSCSI-A	Derived	
vNIC 01-Infra-Host-B	Derived	
vNIC 00-Infra-Host-A	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-B	03-Infra-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-Boot-A	02-Infra-iSCSI-A	default	Derived

Add Delete Modify

OK
Cancel

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

Create iSCSI Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif02a and ficsi_lif02b).

Two boot policies are configured in this procedure. The first policy configures the primary target to be iscsi_lif01a with four SAN paths. The second policy only configures one iSCSI target with one SAN path for Windows installation.

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

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

Add iSCSI Boot



iSCSI vNIC :

OK

Cancel

11. Click OK.
12. From the iSCSI vNICs drop-down list, select Add iSCSI Boot.
13. Enter `iSCSI-Boot-B` in the iSCSI vNIC field.
14. Click OK.
15. Click OK, then click OK again to create the boot policy.
16. Right-click Boot Policies.
17. Select Create Boot Policy.
18. Enter `iSCSI-One-Path` as the name of the boot policy.
19. Optional: Enter a description for the boot policy.
20. Keep the Reboot on Boot Order Change option cleared.
21. Expand the Local Devices drop-down list and select Remote CD/DVD.
22. Expand the iSCSI vNICs drop-down list and select Add iSCSI Boot.

23. Enter `iSCSI-Boot-A` in the iSCSI vNIC field.
24. Click OK.
25. Click OK, then click OK again to create the boot policy.

Create iSCSI Boot Service Profile Templates

In this procedure, one service profile template for installation of Windows on 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 the Servers icon 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 `Install-Win-iSCSI-Host` as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. **Select the “Updating Template” option.**
7. Under UUID, select `UUID_Pool` as the UUID pool.

Create Service Profile Template ? 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.

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 iSCSI-Boot from the LAN Connectivity Policy drop-down list.
4. Select IQN-Pool from the Initiator Name Assignment drop-down list.

5. Click Next.

Configure SAN Connectivity Options

1. Select the No vHBAs option for the “How would you like to configure SAN connectivity?” field.

1 Identify Service Profile Template

2 Storage Provisioning

3 Networking

4 SAN Connectivity

5 Zoning

6 vNIC/vHBA Placement

7 vMedia Policy

8 Server Boot Order

9 Maintenance Policy

10 Server Assignment

11 Operational Policies

Create Service Profile Template

Optionally specify disk policies and SAN configuration information.

How would you like to configure SAN connectivity?

Simple Expert No vHBAs Use Connectivity Policy

This server associated with this service profile will not be connected to a storage area network.

< Prev Next > Finish Cancel

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-One-Path` for Boot Policy.
2. Under Boot Order, expand Boot Order and select the `iSCSI-Boot-A` row.
3. Select the Set iSCSI Boot Parameters button.

4. Select iSCSI-IP-Pool-A for the Initiator IP Address Policy.
5. Scroll to the bottom of the window and click Add.
6. Enter the IQN (Target Name) from the Infra-MS-SVM iSCSI Target Name. To get this IQN, ssh into the **storage cluster interface** and type **“iscsi show”**.
7. For IPv4 address, enter the IP address of iscsi_lif01a from the Infra-MS-SVM. To get this IP, ssh into the **storage cluster interface** and type **“network interface show -vserver Infra-MS-SVM”**.

Create iSCSI Static Target



iSCSI Target Name	:	<input type="text" value="iqn.1992-08.com.netapp:"/>	
Priority	:	<input type="text" value="1"/>	
Port	:	<input type="text" value="3260"/>	
Authentication Profile	:	<input type="text" value="<not set> ▼"/>	Create iSCSI Authentication Profile
IPv4 Address	:	<input type="text" value="192.168.12.61"/>	
LUN ID	:	<input type="text" value="0"/>	

8. Click OK to complete configuring the iSCSI target.

Set iSCSI Boot Parameters



Initiator Name

Initiator Name Assignment:

[Create IQN Suffix Pool](#)

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

Initiator Address

Initiator IP Address Policy:

IPv4 Address : **0.0.0.0**
 Subnet Mask : **255.255.255.0**
 Default Gateway : **0.0.0.0**
 Primary DNS : **0.0.0.0**
 Secondary DNS : **0.0.0.0**

[Create IP Pool](#)

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface
 iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPv4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.12.61	0

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

- Click OK to complete Setting the iSCSI Boot Parameters.



At this point, we are only putting in one iSCSI SAN path because the Windows installer does not support multipathing.

- Click Next.

Configure Maintenance Policy

To configure the Maintenance Policy, complete the following steps:

- Change the Maintenance Policy to default.

Create Service Profile Template ? ×

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

Maintenance Policy

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

Maintenance Policy: [Create Maintenance Policy](#)

Name : **default**
 Description :
 Soft Shutdown Timer : **150 Secs**
 Storage Config. Deployment Policy : **User Ack**
 Reboot Policy : **User Ack**

< Prev Next > **Finish** Cancel

- Click Next.

Configure Server Assignment

To configure server assignment, complete the following steps:

- In the Pool Assignment list, select `Hyper-V-MGMT-Pool1`.
- Select Down as the power state to be applied when the profile is associated with the server.
- Optional: select **“UCS-Broadwell”** for the Server Pool Qualification.

- Expand Firmware Management at the bottom of the page and select the default policy.

Create Service Profile Template ? X

Optionally specify a server pool for this service profile template.

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

Pool Assignment: [Create Server Pool](#)

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

Up Down

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

Server Pool Qualification : ▼

Restrict Migration :

⊖ Firmware Management (BIOS, Disk Controller, Adapter)

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

Host Firmware Package: ▼

[Create Host Firmware Package](#)

- Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

- In the BIOS Policy list, select **Virtual-Host**.
- 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

⊕ Graphics Card Policy

< Prev Next > **Finish** Cancel

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

Create Multipath Service Profile Template

To create a Service Profile Template with 4 iSCSI SAN paths to be used once multipathing software is installed in Windows, complete the following steps:

1. Select the Servers icon on the left.
2. Expand Service Profile Templates > root.
3. Right-click the newly-created Service Template Install-Win-iSCSI-Host and select Create a Clone.
4. Name the clone Hyper-V-iSCSI-Host.
5. Click OK then OK again to complete creating the clone.
6. Select the newly cloned Service Template Hyper-V-iSCSI-Host. On the right, select the Boot Order tab.
7. In the middle of the screen, select iSCSI-Boot-A Primary and click the Set iSCSI Boot Parameters button.

8. Scroll to the bottom of the window and select the configured iSCSI static target. Click Info at the bottom of the screen.
9. Change the IPv4 Address to the IP address of iscsi_lif02a in the Infra-MS-SVM. To get this address, ssh **into the storage cluster and type “network interface show -vserver Infra-MS-SVM”**.
10. Fully select, right-click and copy the iSCSI Target Name.
11. Click OK to complete modifying the iSCSI Target.
12. Click Add.
13. For the iSCSI Target Name, either paste in the IQN of the Infra-MS-SVM or retrieve it from the storage cluster. To get this IQN, ssh into the storage cluster and type **“iscsi show”**.
14. For the IPv4 Address, enter the IP address of iscsi_lif01a in the Infra-MS-SVM. To get this address, ssh **into the storage cluster and type “network interface show -vserver Infra-MS-SVM”**.
15. Click OK to complete Creating the iSCSI Static Target.

Set iSCSI Boot Parameters



Name : **iSCSI-Boot-A**

Authentication Profile : <not set> ▼

[Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment: <not set> ▼

[Create IQN Suffix Pool](#)

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

Initiator Address

Initiator IP Address Policy: iSCSI-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](#)

[Reset Initiator Address](#)

The IP address will be automatically assigned from the selected pool.

iSCSI Static Target Interface iSCSI Auto Target Interface

Name	Priority	Port	Authentication Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08.c...	1	3260		192.168.12.62	0
iqn.1992-08.c...	2	3260		192.168.12.61	0

OK

Cancel

16. Click OK then OK again to complete Setting the iSCSI Boot Parameters.
17. In the middle of the screen, select iSCSI-Boot-B Secondary and click the Set iSCSI Boot Parameters button.
18. For the Initiator IP Address Policy, select iSCSI-IP-Pool-B.
19. Scroll to the bottom of the window and click Add.
20. For the iSCSI Target Name, either paste in the IQN of the Infra-MS-SVM or retrieve it from the storage cluster. **To get this IQN, ssh into the storage cluster and type “iscsi show”.**
21. For the IPv4 Address, enter the IP address of iscsi_lif02b in the Infra-MS-SVM. To get this address, ssh **into the storage cluster and type “network interface show -vserver Infra-MS-SVM”.**
22. Fully select, right-click and copy the iSCSI Target Name.
23. Click OK to complete modifying the iSCSI Target.
24. Click Add.
25. For the iSCSI Target Name, either paste in the IQN of the Infra-MS-SVM or retrieve it from the storage cluster. **To get this IQN, ssh into the storage cluster and type “iscsi show”.**
26. For the IPv4 Address, enter the IP address of iscsi_lif01b in the Infra-MS-SVM. To get this address, ssh **into the storage cluster and type “network interface show -vserver Infra-MS-SVM”.**
27. Click OK to complete Creating the iSCSI Static Target.

Set iSCSI Boot Parameters



Initiator Name

Initiator Name Assignment: <not set> ▼

[Create IQN Suffix Pool](#)

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

Initiator Address

Initiator IP Address Policy: iSCSI-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 Pr...	iSCSI IPV4 Address	LUN Id
iqn.1992-08....	1	3260		192.168.22.62	0
iqn.1992-08....	2	3260		192.168.22.61	0

⊕ Add
🗑 Delete
ℹ Info

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

OK

Cancel

28. Click OK then OK again to complete Setting the iSCSI Boot Parameters.

Create Service Profiles

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

1. Connect to UCS Manager and click the Servers icon on the left.
2. Select Service Profile Templates > root > Service Template Install-Win-iSCSI-Host.
3. Right-click Install-Win-iSCSI-Host and select Create Service Profiles from Template.
4. Enter `Hyper-V-MGMT-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 ? ×

Naming Prefix :

Name Suffix Starting Number :

Number of Instances :

8. Click OK in the confirmation message.

Storage Configuration – Boot LUNs

NetApp ONTAP Boot Storage Setup

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-01 -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-01-iqn>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-02 -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-02-iqn>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-All -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-01-iqn>,<hyper-v-mgmt-02-iqn>
```

1. To get the management host IQNs, log in to Cisco UCS Manager and click the Servers icon on the left.
2. Select Servers > Service Profiles > root and the host Service Profile. The host IQN is listed under the iSCSI vNICs tab on the right.

Map LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun MGMT-Win2016-Gold -igroup Hyper-V-MGMT-01 -lun-id
0
lun map -vserver Infra-MS-SVM -volume witness -lun witness -igroup Hyper-V-MGMT-All -lun-id 1
lun map -vserver Infra-MS-SVM -volume iscsi_datastore_1 -lun iscsi_datastore_1 -igroup Hyper-V-MGMT-
All -lun-id 2
lun map -vserver Infra-MS-SVM -volume iscsi_datastore_2 -lun iscsi_datastore_2 -igroup Hyper-V-MGMT-
All -lun-id 3
```

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 a FlexPod 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 the Servers icon on the left.
7. Select Servers > Service Profiles > root > `Hyper-V-MGMT-Host-01`.
8. **On the right under the General tab, select the “>>” icon to the right of “KVM Console”.**
9. Follow the prompts to launch the Java KVM console.
10. From the KVM Console, under the Virtual Media tab, select Activate Virtual Devices. Follow the prompts and select Apply.
11. From the KVM Console, under the Virtual Media tab, select Map CD/DVD.
12. Click Browse.
13. Browse to the Windows Server 2016 installation ISO image file and click Open.
14. Map the image that you just added by selecting Map Device.
15. To boot the server, click the Boot Server icon above the KVM Console tab.

16. Click OK then OK again to boot the server.

Install Windows Server 2016


To install Windows Server 2016 to the first management host, complete the following steps:

1. On boot, the machine detects the presence of the Windows installation media.
2. After the installer has finished loading, select the relevant region information and click Next.
3. Click Install now.
4. Enter the Product Key and click Next.
5. Select Windows Server 2016 Datacenter (Desktop Experience) and click Next.



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

6. After reviewing the EULA, accept the license terms and click Next.
7. Select Custom: Install Windows only (advanced).
8. In the Windows Setup window, select Load driver.
9. Under Virtual Media, select the Windows Server 2016 item to unmap it. Click Yes to complete the un-mapping.
10. Under Virtual Media, select map CD/DVD.
11. Click Browse and browse to the ucs-bxxx-drivers-windows.3.2.1 iso. Select this iso and click Open. Click Map Device to map this iso.
12. In the Load driver window, click Browse.
13. Browse to the CD Drive and expand Network > Cisco > VIC > W2K16. Select x64 under W2K16. Click OK.
14. Back in the Windows Setup window, make sure Cisco VIC Ethernet Interface is selected and click Next.



15. If you are booting with FC, also load the VIC fNIC storage driver.
16. Two disk drives should now appear in the Windows Setup window. In the Virtual Media menu, unmap the ucs-bxxx-drivers-windows.3.2.1 iso and remap the Windows Server 2016 installation iso.
17. In the Windows Setup window, click Refresh. Make sure the 200 GB drive is selected and click Next.
18. When Windows is finished installing, enter an administrator password on the settings page and click Finish.
19. Under Virtual Media, unmap the Windows Server 2016 Installation iso.

Host Renaming and Join to Domain

To rename host and join to a domain, complete the following steps:

1. Login to the host and open PowerShell.
2. Rename the host.

```
Rename-Computer -NewName Win2016-Gold -restart
```

3. Set the MTU of the current management interface to 1500.

```
netsh interface ipv4 set subinterface 00-Infra-Host-A mtu=1500 store=persistent
```

4. Assign an IP address to the management interface.

```
Get-NetAdapter - determine the ifIndex of the 00-Infra-Host-A adapter
```

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

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

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

6. Add the host to Active Directory.

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

7. Set the timezone of the host to the appropriate timezone.

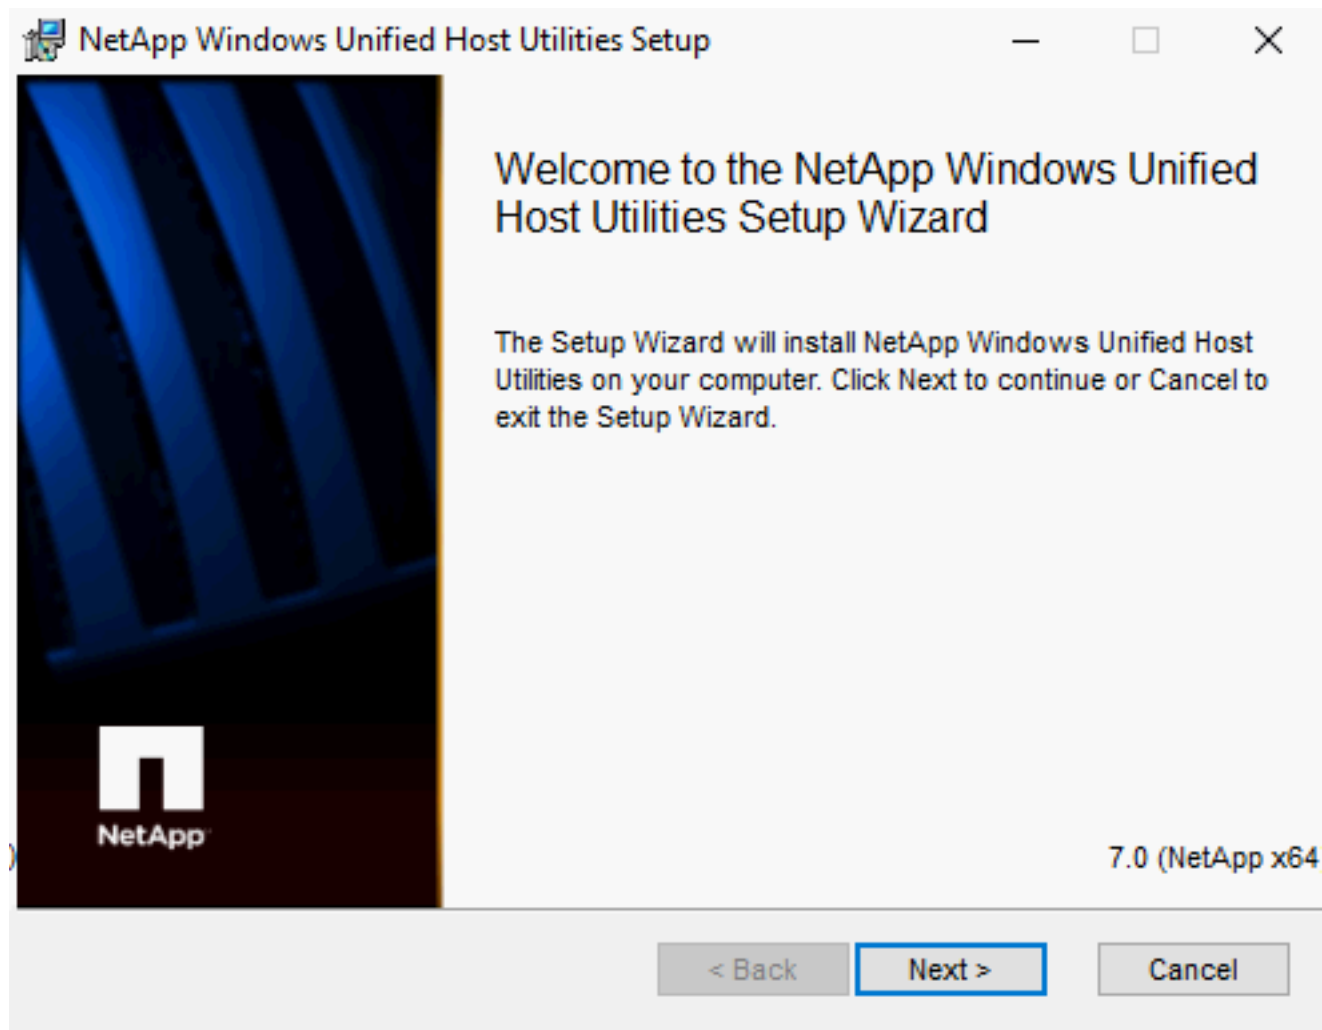
Install NetApp Windows Unified Host Utilities

After enabling the MPIO feature in Windows, download and install NetApp Windows Unified Host Utilities. To download and install the host utilities, complete the following steps:


1. Download the x64 version of the NetApp host utilities v7.0 for Windows from the link below:

https://mysupport.netapp.com/NOW/download/software/sanhost_win/7.0

2. Run the .msi file. The NetApp Windows Unified Host Utilities setup wizard is launched. Click OK on the hotfix warning and then click Next.



3. Click the checkbox to accept the license agreement and click Next.
4. Select "Yes, install support for Multipath IO" and click Next.



5. Accept the default destination folder and click Next.
6. Click Install, OK, and Finish to complete the installation of host utilities.
7. Click No to not restart the computer.
8. Shut down the server.

Set Up Multipathing and iSCSI

To set up multipathing and iSCSI, complete the following steps:

1. In Cisco UCS Manager, select the Hyper-V-MGMT-Host-01 Service Profile.
2. Under the General tab on the right, select Bind to a Template.
3. Select the Hyper-V-iSCSI-Host template which contains 4 SAN paths for multipath I/O. Click OK, then Yes and OK to complete binding to the new Service Profile Template.

4. Follow steps 1-3 to bind the Hyper-V-MGMT-Host-02 Service Profile to the Hyper-V-iSCSI-Host template.
5. When the host with Service Profile Hyper-V-MGMT-Host-01 returns to the Power Off state, go to the KVM console, and select Boot Server to boot the host. Click OK and OK again to proceed. When the server boots, 4 iSCSI SAN paths should be seen.
6. Log into the server as Administrator.

Install NetApp SnapDrive 7.1.4 for Windows

NetApp SnapDrive® 7.1.4 for Windows allows you to automate storage provisioning tasks and to manage data in physical or virtual Microsoft Windows hosts within SMB 3.0 environments.

To install SnapDrive 7.1.4 for Windows on Microsoft Hyper-V Hosts, complete the steps in this section. It is assumed that the prerequisites have been verified.

Configuring Access for SnapDrive for Windows

To configure access for SnapDrive 7.1.4 for Windows, complete the following steps:

1. Create a user account on the storage virtual machine by entering the following command:

```
security login create -vserver -user -authentication-method -application -role
```

The variables in this command represent the following values:

- `-vserver` is the name of the Vserver for the user to be created.
- `-user` is the SnapDrive user name.
- `-authentication-method` is the method used for authentication.
- `-application` is the communications application the user will use to access SnapDrive.
- `-role` is the user privileges.

For example,

To add a user called `snapdrive` to the `BUILTIN\Administrators` group on the storage system, run the following command:

```
security login create -vserver Infra-MS-SVM -user snapdrive -authentication-method password -
application http -role vsadmin

security login create -vserver Infra-MS-SVM -user snapdrive -authentication-method password -
application ontapi -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.

1. When prompted, enter a password, for the user account you are creating.

2. You will be prompted to enter this password twice during the SnapDrive installation, so make a note of it, including letter case.
3. Make sure that the user account you just created belongs to the local administrator's group on the storage virtual machine 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.

4. On the Active Directory domain, create a snapdrive user and add it to the Domain Admins user group.



Set up the domain snapdrive user account so that the password for the account never expires. For detailed instructions on how to create domain user accounts, see your Windows documentation.

Downloading SnapDrive 7.1.4 for Windows

To download SnapDrive 7.1.4 for Windows, complete the following steps:



Before you install SnapDrive for Windows, download the software package from the [NetApp Support site](#) (requires login credentials).

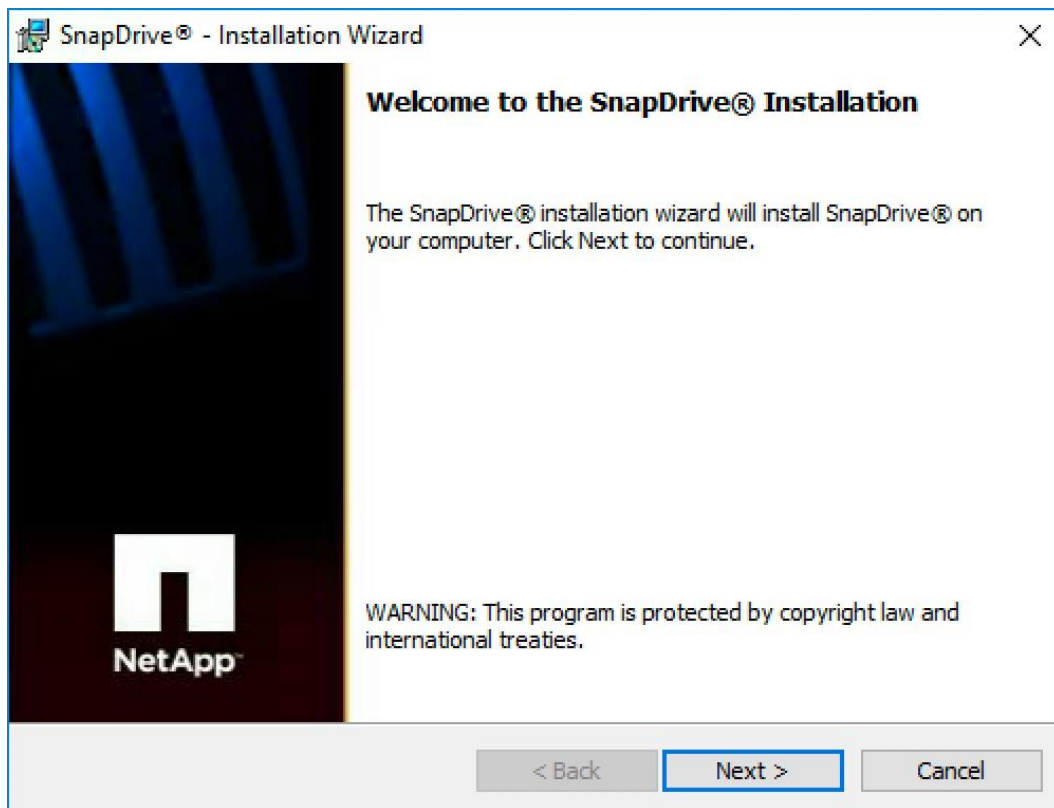
1. Log in to the [NetApp Support site](#).
2. Go to the Download Software page.
3. From the drop-down list, select Windows for the operating system on which you are installing SnapDrive, and click Go!
4. Click View & Download for the software version 7.1.4.
5. On the SnapDrive for Windows Description page, click Continue.
6. Review and accept the terms of the license agreement.
7. On the Download page, click the link for the installation file.
8. Save the SnapDrive installation 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

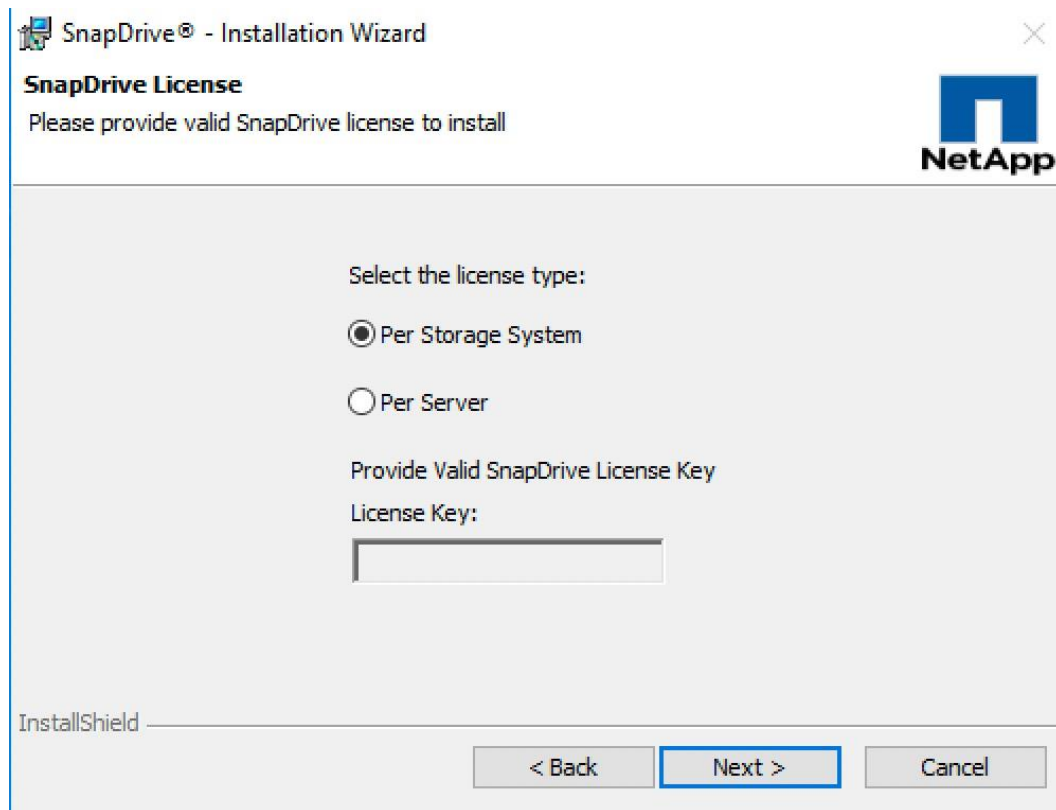
To install SnapDrive for Windows, complete the following steps:

1. Login to the Windows Host as a Domain Admin user

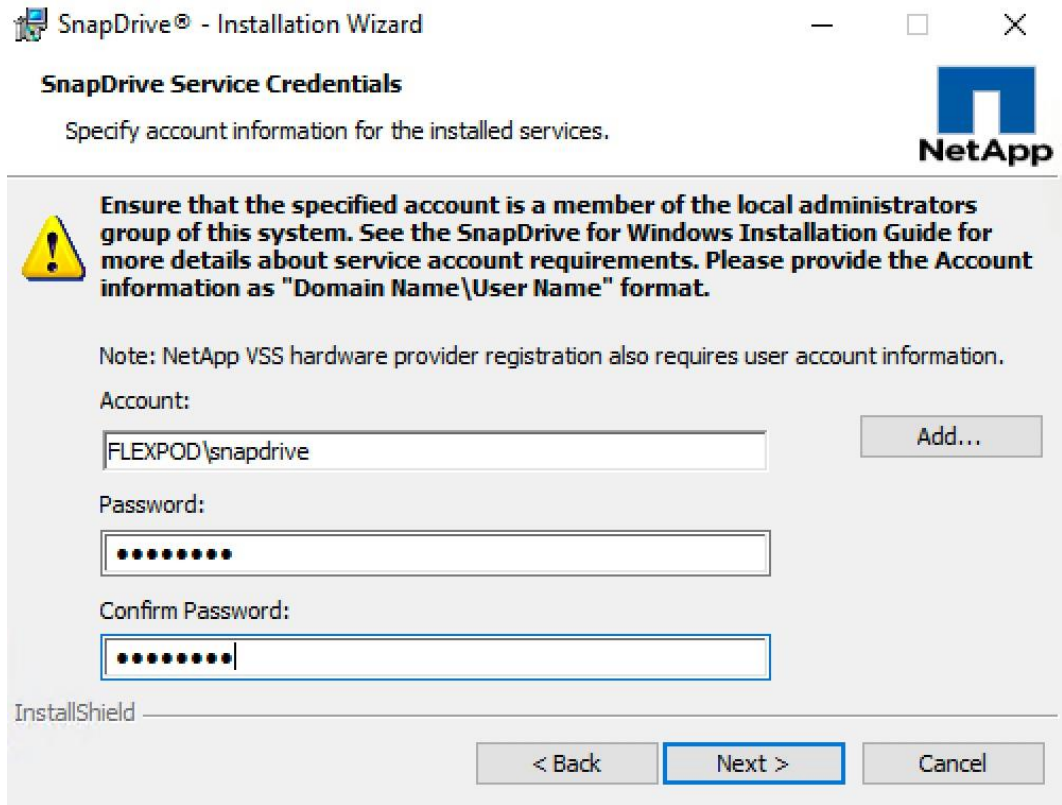
2. Make sure that the Microsoft .NET Framework 3.5 Feature is installed on the Windows Host.
3. Launch the SnapDrive for Windows installer as Administrator and then follow the instructions in the wizard.



4. On the SnapDrive License page, select the appropriate license type and click Next.



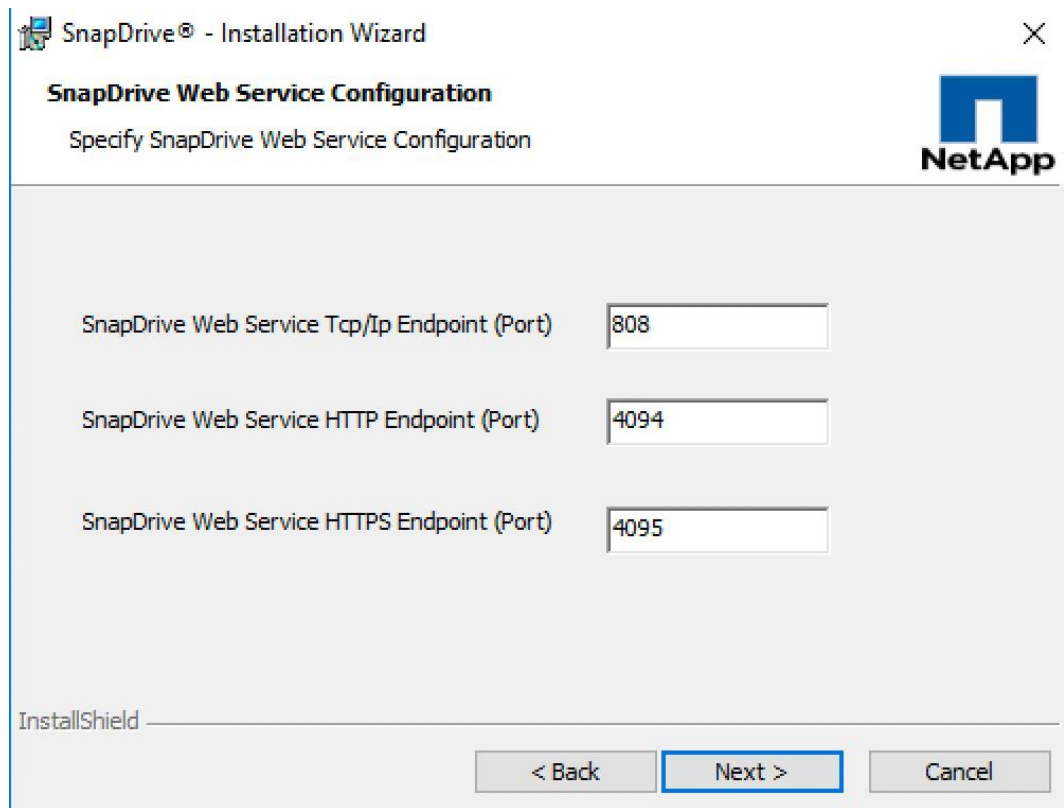
5. On the Customer Information page, enter the appropriate information and click Next.
6. On the Destination Folder page, enter the appropriate destination or accept the default. Click Next.
7. On the Set Firewall Rules page, select for Enable SnapDrive to Communicate Through the Windows Firewall option and click Next.
8. On the SnapDrive Service Credentials page, enter the account and password information of the account created earlier that is a member of the local administrators group.




9. On the SnapDrive Web Service Configuration page, accept the default port numbers and click Next.

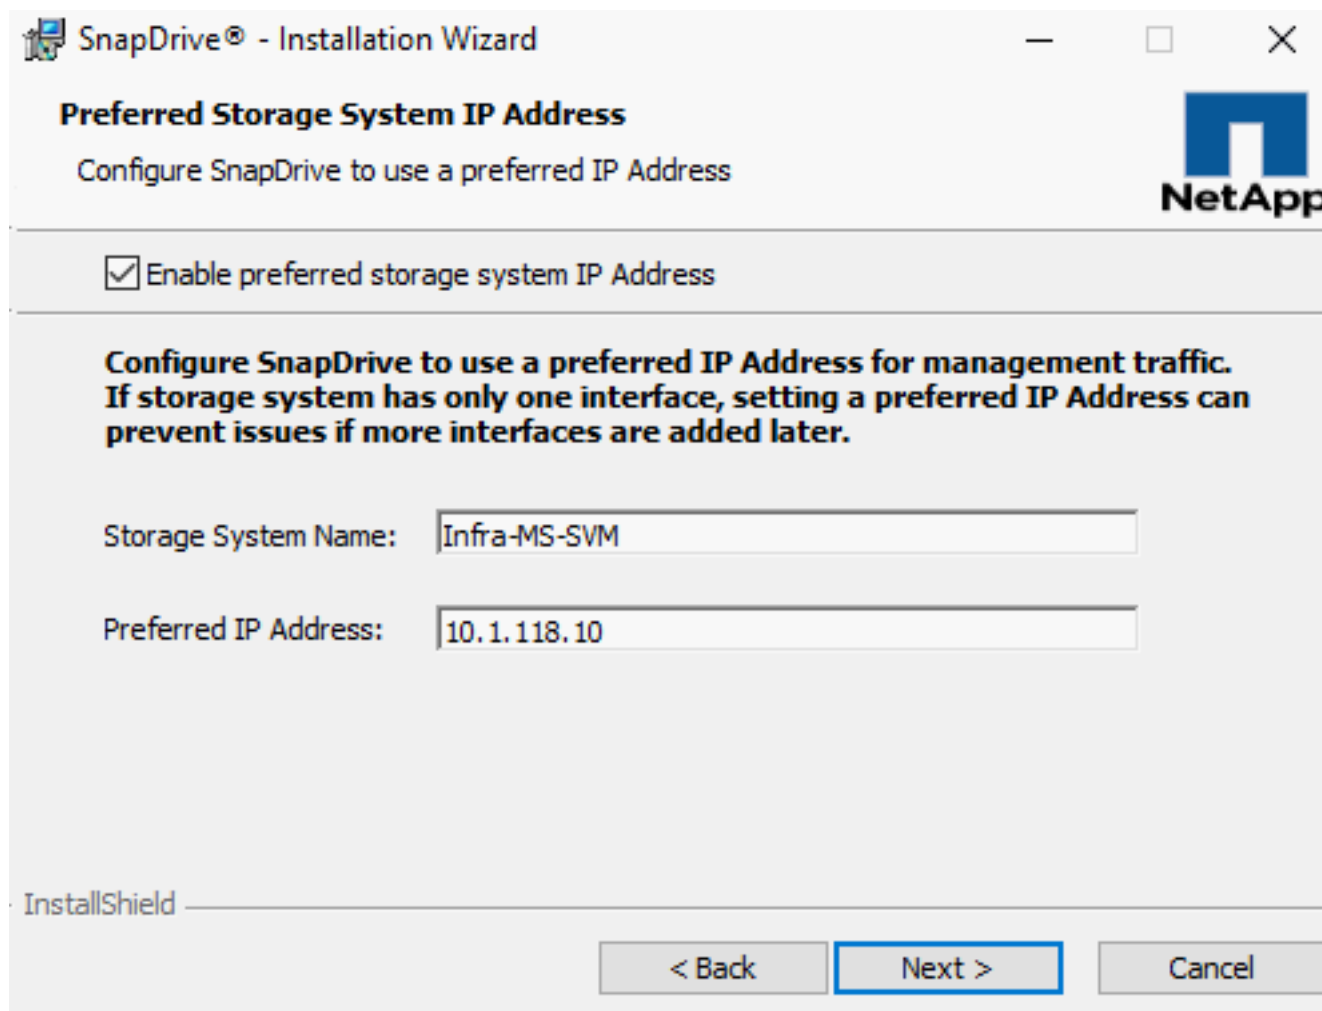


If you want to change the port numbers, you should also change the port numbers for the other Snap-Drive hosts.

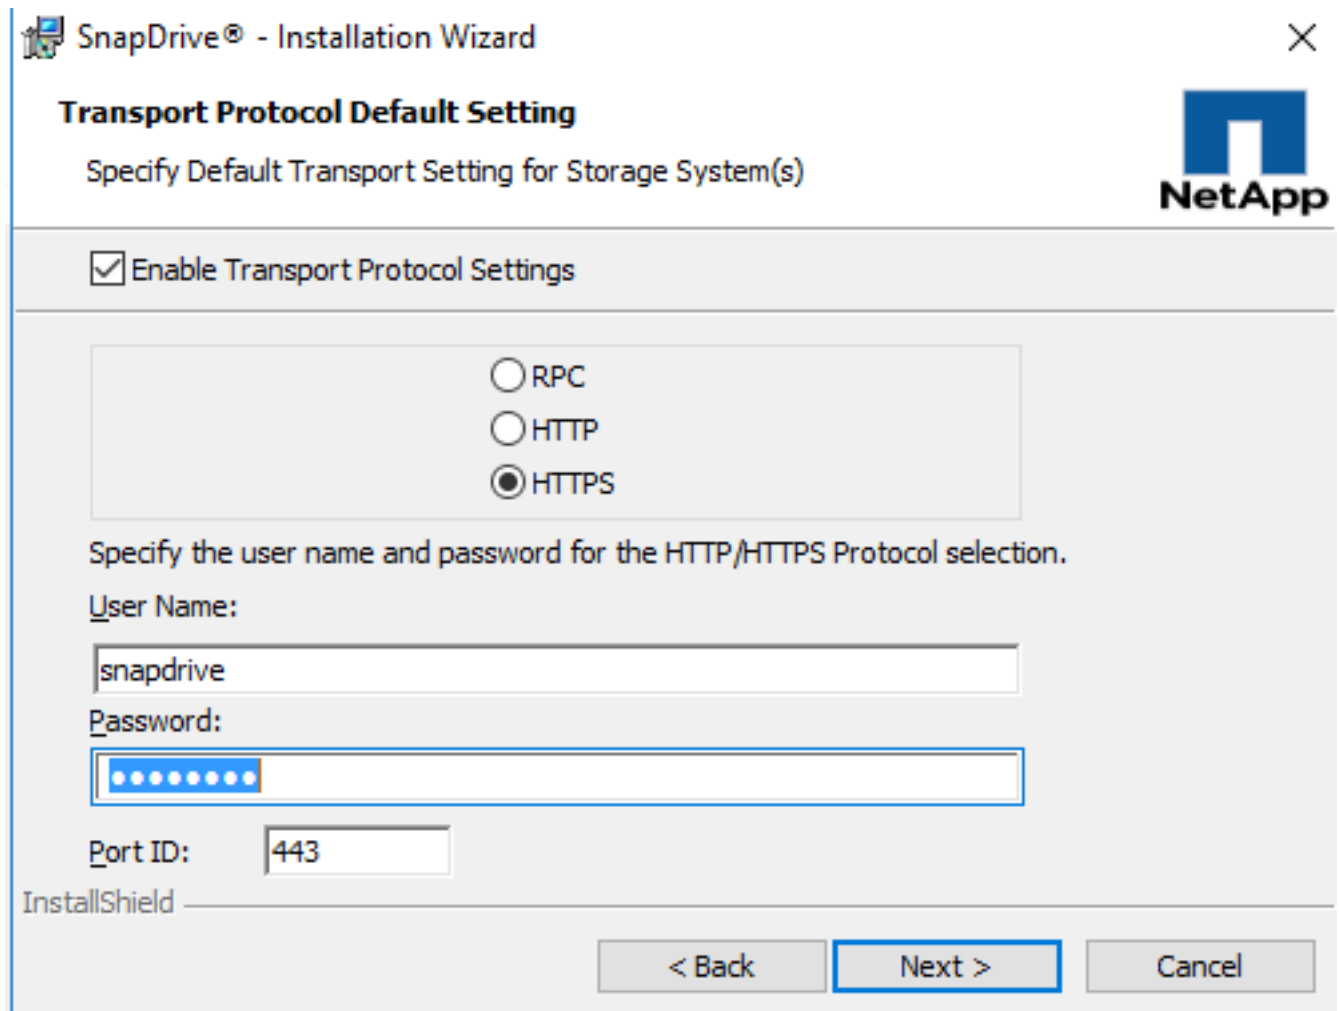


10. On the Preferred IP Address screen, identify the IP address you want to use to communicate with the storage system and click Next.


 You should configure the preferred IP address, because doing this improves performance and scalability.



11. On the Transport Protocol Default Setting page, enable the storage protocol settings and click Next. RPC is not supported in clustered Data ONTAP.

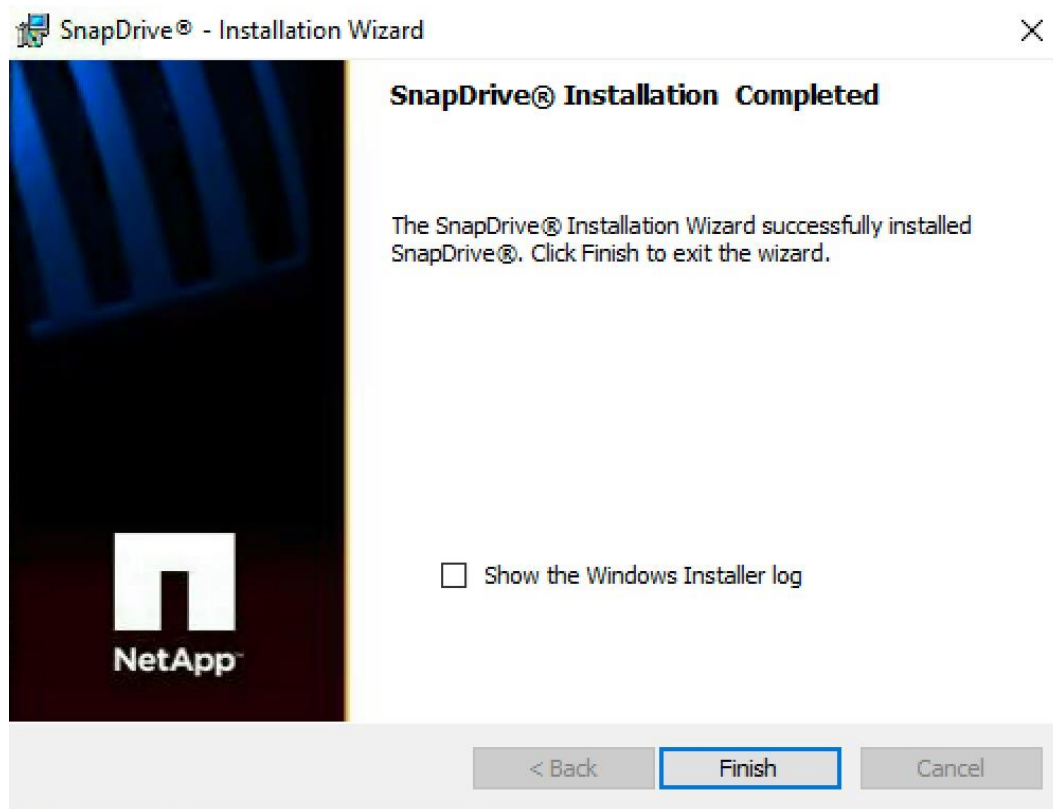


12. On the Unified Manager Configuration Screen, click Next.

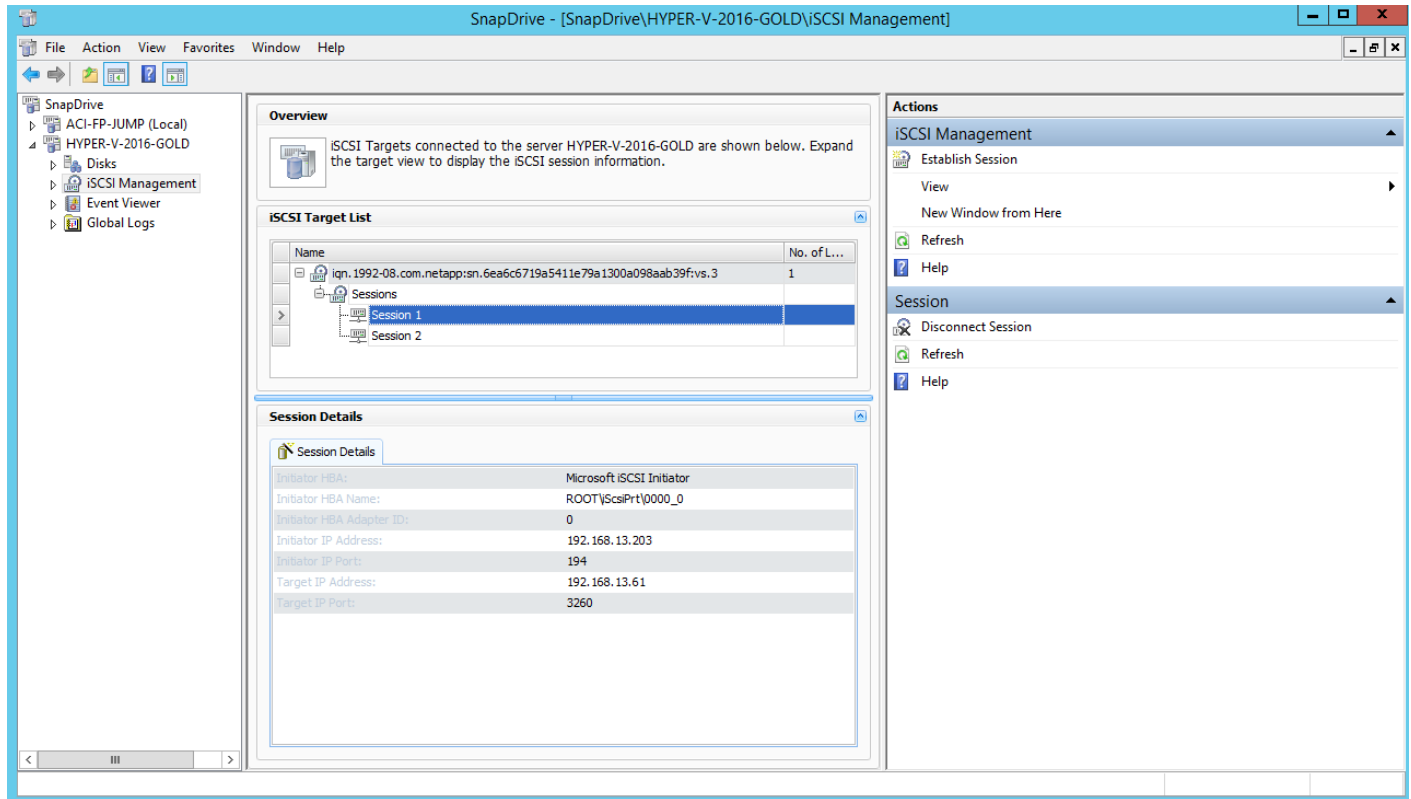
 OnCommand Unified Manager Core Package data protection capabilities are available only in 7- Mode environments.

13. On the Ready to Install page, click Install.


14. When the SnapDrive installation is complete, click Finish.



- 15. Launch the SnapDrive snap-in.
- 16. Select iSCSI Management and expand to show sessions. Two sessions will display.



17. Use “Establish Session” to add sessions for the two missing iSCSI LIFs. Four sessions will display.




18. Close SnapDrive.

Configure Server for Cloning

To configure the server for cloning, complete the following steps:

1. Under Windows Administrative tools, open the Computer Management tool.
2. On the left, select Disk Management. Click Cancel on the Disk 1 initialization window. Near the bottom middle of the window, right-click Disk 0 and select Properties.
3. Select the MPIO tab and verify the disk now has two Active/Optimized paths and two Active/Un-optimized paths. Click OK to close the Disk Device Properties window and close the Computer Management tool.



4. Install all available Windows Updates on the server.
5. Open Windows Powershell and enter `SCONFIG` to configure the server.
6. Select Remote Desktop by typing `7` and pressing `Enter`.
7. Enter `E` to enable Remote Desktop.
8. Enter `2` to allow any version of Remote Desktop. Click `OK` to acknowledge Remote Desktop enablement.
9. Enter `14` to shut down the server. Click `Yes` to complete shutdown.



The boot LUN cloning procedure used in this document will only work when the clones are applied to the same server hardware. If the clone source image was created on a Cisco UCS B200 M4, and you want to apply the image to a Cisco UCS C220 M4, you will need to follow the steps above to install Windows on the Cisco UCS C220 M4.

Clone and Remap Server LUNs for Sysprep Image

To clone and remap the server LUNs for the sysprep image, complete the following steps:

1. In the storage cluster interface, unmap the MGMT-Win2016-Gold LUN.

```
lun unmap -path /vol/HV_boot/MGMT-Win2016-Gold -igroup Hyper-V-MGMT-01
```

2. Make a clone of the MGMT-Win2016-Gold LUN for the Sysprep clone.

```
clone start -source-path /vol/HV_boot/MGMT-Win2016-Gold -destination-path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep
```

3. Map the Sysprep clone boot LUN to the first Hyper-V management host.

```
lun map -path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep -igroup Hyper-V-MGMT-01 -lun-id 0
```

Boot and Set Up Sysprep Clone

To boot and set up the sysprep clone, complete the following steps:

1. Back in the UCS KVM Console for Hyper-V-MGMT-Host-01, click Boot Server then OK two times to boot the Sysprep Clone LUN.
2. Once the server boots up, log in as the local machine Administrator.
3. Open the Windows Powershell prompt and enter `C:\Windows\System32\Sysprep\sysprep /generalize /oobe /shutdown` **to reset the machine's security id. The server will shut down.**

Clone and Remap Server LUNs for Production Images

To clone and remap server LUNs for production images, complete the following steps:

1. In the storage cluster interface, unmap the MGMT-Win2016-Gold-Sysprep LUN.

```
lun unmap -path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep -igroup Hyper-V-MGMT-01
```

2. Make two clones of the MGMT-Win2016-Gold-Sysprep LUN for the Hyper-V-MGMT hosts.

```
clone start -source-path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep -destination-path /vol/HV_boot/Hyper-V-MGMT-01
clone start -source-path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep -destination-path /vol/HV_boot/Hyper-V-MGMT-02
```

3. Map the Hyper-V-MGMT LUNs to the hosts.

```
lun map -path /vol/HV_boot/Hyper-V-MGMT-01 -igroup Hyper-V-MGMT-01 -lun-id 0
lun map -path /vol/HV_boot/Hyper-V-MGMT-02 -igroup Hyper-V-MGMT-02 -lun-id 0
```

Boot and Set Up Clones

To boot and set up clones, complete the following steps:

1. Back in the UCS KVM Console for Hyper-V-MGMT-Host-01, click Boot Server then OK two times to boot Hyper-V-MGMT-Host-01.
2. When the server boots up, select the appropriate Regional and Language information and click Next.
3. Enter the server Product Key and click Next.
4. Click Accept to Accept the License terms.
5. Log into the server as Administrator and open Powershell.
6. Rename the host.

```
Rename-Computer -NewName Hyper-V-MGMT-01 -Restart
```

7. The server will reboot. Return to Powershell. Assign an IP address to the management interface.

```
Get-NetAdapter - determine the ifIndex of the 00-Infra-Host-A adapter
```

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


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

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

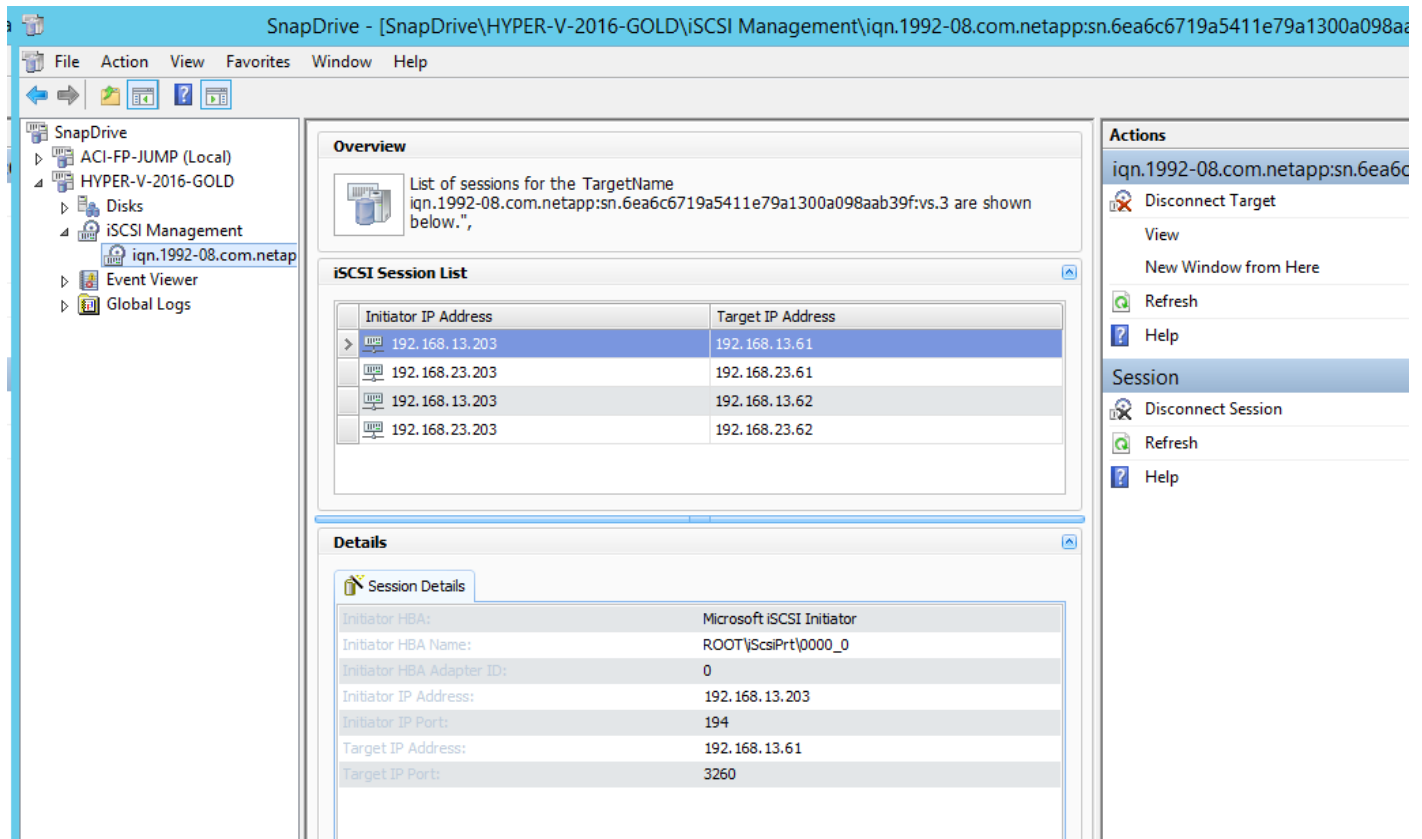
9. Add the host to Active Directory.

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

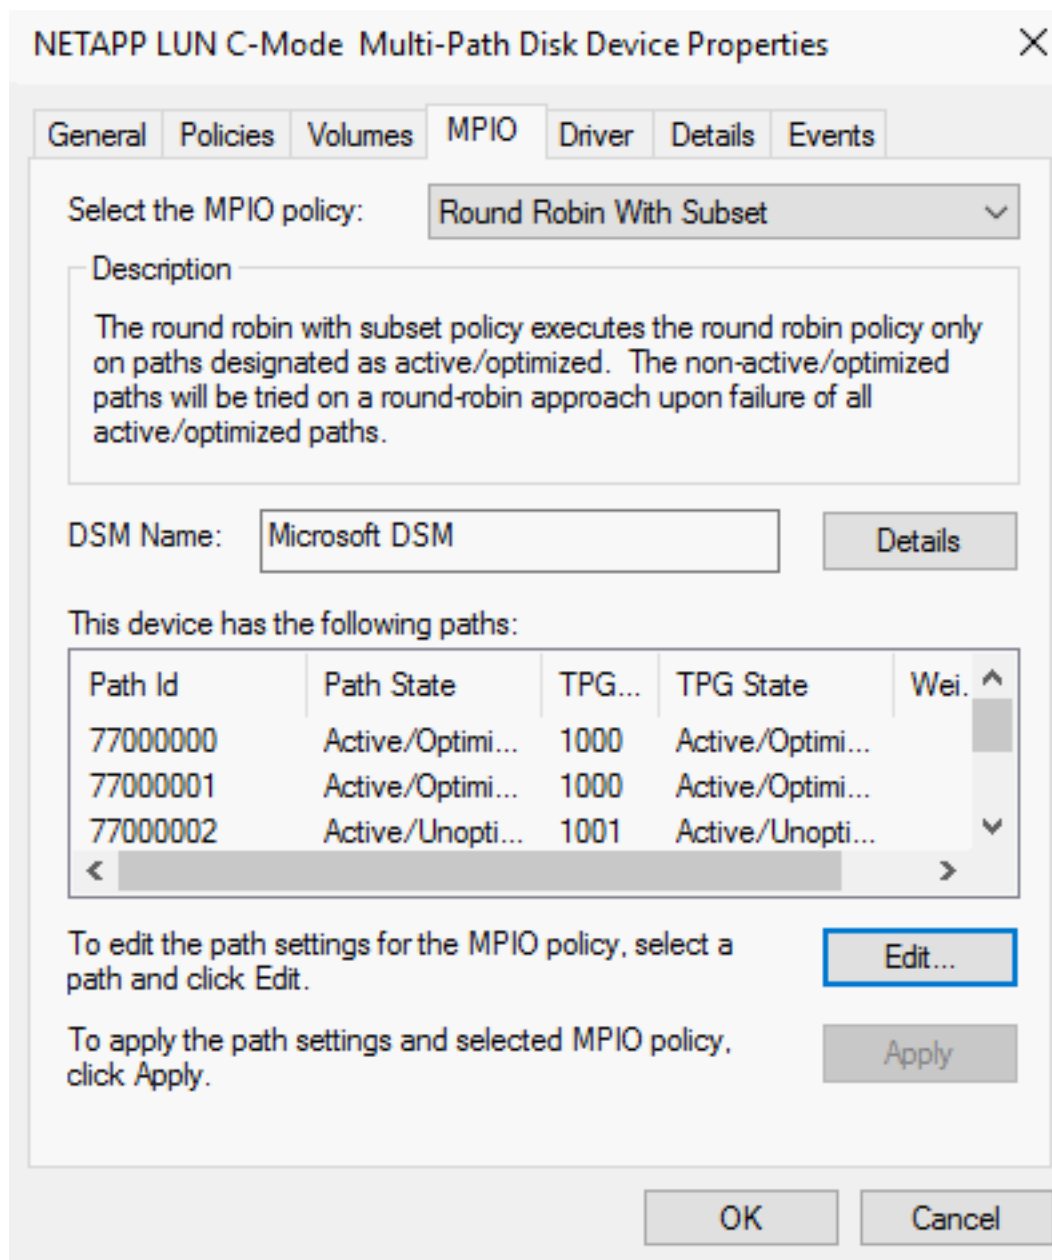
10. Enter the appropriate Domain Admin credentials and click OK. The server will reboot. Login to the server as a Domain Admin.
11. Launch the SnapDrive snap-in.
12. Select iSCSI Management and expand to show sessions. At least two sessions will display.



13. Use “Establish Session” to add sessions for any missing iSCSI LIFs. Four sessions will display.



14. Close SnapDrive.
15. Under Windows Administrative tools, open the Computer Management tool.
16. Select Disk Management.
17. Only on the first Hyper-V Management host, on the Initialize Disk window, make sure Disks 1-3 are selected and select the GPT radio button. Click OK.
18. Only on the first Hyper-V Management host, for Disk 1 (the 1GB disk), right-click the Unallocated area and select New Simple Volume. Click Next. Leave the size at default maximum and click Next. Select **“Do not assign a drive letter or drive path”** and click Next. **Change the Volume label to Quorum** and click Next. Click Finish.
19. Only on the first Hyper-V Management host, for Disk 2 (the first 500GB disk), right-click the Unallocated area and select New Simple Volume. Click Next. Leave the size at default maximum and click Next. Select **“Do not assign a drive letter or drive path”** and click Next. **Change the Volume label to iscsi_datastore_1** and click Next. Click Finish.
20. Only on the first Hyper-V Management host, for Disk 3 (the second 500GB disk), right-click the Unallocated area and select New Simple Volume. Click Next. Leave the size at default maximum and click Next. Select **“Do not assign a drive letter or drive path”** and click Next. **Change the Volume label to iscsi_datastore_2** and click Next. Click Finish.
21. Only on the second Hyper-V Management host, launch Disk Management. At the Initialize Disk window, click Cancel. Under the Action Menu, select Rescan Disks. Verify that Disks 1-3 now show a status of Online.
22. Near the bottom middle of the window, right-click Disk 0 and select Properties.
23. Select the MPIO tab and verify the disk now has two Active/Optimized paths and two Active/Un-optimized paths. Click OK to close the Disk Device Properties window and close the Computer Management tool.



24. Open a second UCS KVM console and repeat steps 1-23 for the Hyper-V-MGMT-02 host.

Install Roles and Features Required for Hyper-V

To install roles and features on both Hyper-V-MGMT Hosts, complete the following steps:

1. Log in with a Domain Administrator User ID.
2. Open Powershell with elevated rights (Run as Administrator) and add Hyper-V and Windows Failover Clustering by entering the following command:

```
Add-WindowsFeature Hyper-V, Failover-Clustering -IncludeManagementTools -Restart
```

- The servers will reboot two times. When the reboots have completed, log in with a Domain Administration User ID.

Set Up Hyper-V Networking

To set up networking on both Hyper-V-MGMT Hosts from a UCS KVM Console, complete the following steps:

- Open Powershell with elevated rights (Run as Administrator). Set the MTU of the 00-Infra-Host-A network interface back to 9000.

```
netsh interface ipv4 set subinterface 00-Infra-Host-A mtu=9000 store=persistent
netsh interface ipv4 show subinterface
```

- Run the Get-NetAdapter command to confirm the vNIC names.

```
PS C:\Windows\system32> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
04-APIC-MS-VS-A	Cisco VIC Ethernet Interface #3	15	Up	00-25-B5-A2-0A-02	40 Gbps
02-Infra-iSCSI-A	Cisco VIC Ethernet Interface	7	Up	00-25-B5-A2-0A-01	40 Gbps
00-Infra-Host-A	Cisco VIC Ethernet Interface #2	5	Up	00-25-B5-A2-0A-00	40 Gbps
05-APIC-MS-VS-B	Cisco VIC Ethernet Interface #6	6	Up	00-25-B5-A2-0B-02	40 Gbps
03-Infra-iSCSI-B	Cisco VIC Ethernet Interface #5	9	Up	00-25-B5-A2-0B-01	40 Gbps
01-Infra-Host-B	Cisco VIC Ethernet Interface #4	14	Up	00-25-B5-A2-0B-00	40 Gbps

- Create a NIC team using the 00-Infra-Host-A and 01-Infra-Host-B interfaces.

```
New-NetLbfoTeam -Name HV-Infra-Team -TeamMembers 00-Infra-Host-A, 01-Infra-Host-B -TeamingMode SwitchIndependent -LoadBalancing HyperVPort
```

- Press Y to confirm.

- Remove the management IP stack from the Teamed Interface.

```
Get-NetAdapter HV-Infra-Team | Set-NetAdapterBinding -ComponentID ms_tcpip* -Enabled $false
```

- Create a Hyper-V virtual switch for the management, storage, and VM traffic.

```
New-VMSwitch -Name HV-Infra-vSwitch -NetAdapterName HV-Infra-Team -AllowManagementOS $false
```

- Create Virtual NIC.

```
Add-VMNetworkAdapter -ManagementOS -Name MS-IB-MGMT -SwitchName HV-Infra-vSwitch
```

- Make sure MTU of MS-IB-MGMT virtual adapter is 1500.

```
netsh interface ipv4 show subinterface
```

- Set IP Address for MS-IB-MGMT host virtual NIC.

```
New-NetIPAddress -InterfaceAlias "vEthernet (MS-IB-MGMT)" -IPAddress <host-mgmt-ip> -DefaultGateway <mgmt-gateway> -PrefixLength <mgmt-net-prefix>
```



You will not add a VLAN to this interface since the IB-MGMT-VLAN <418> is the native VLAN for the UCS VIC vNIC interfaces used in the team for the virtual switch.

10. Disable DNS registration for all NICs

```
Set-DnsClient -InterfaceAlias * -Register $false
```

11. Turn registration back on and configure DNS for the Management NIC

```
Set-DnsClient -InterfaceAlias "vEthernet (MS-IB-MGMT)" -Register $true -  
ConnectionSpecificSuffix <dns-domain-name>
```

```
Set-DnsClientServerAddress -InterfaceAlias "vEthernet (MS-IB-MGMT)" -  
ServerAddresses <dns-server-ip>
```



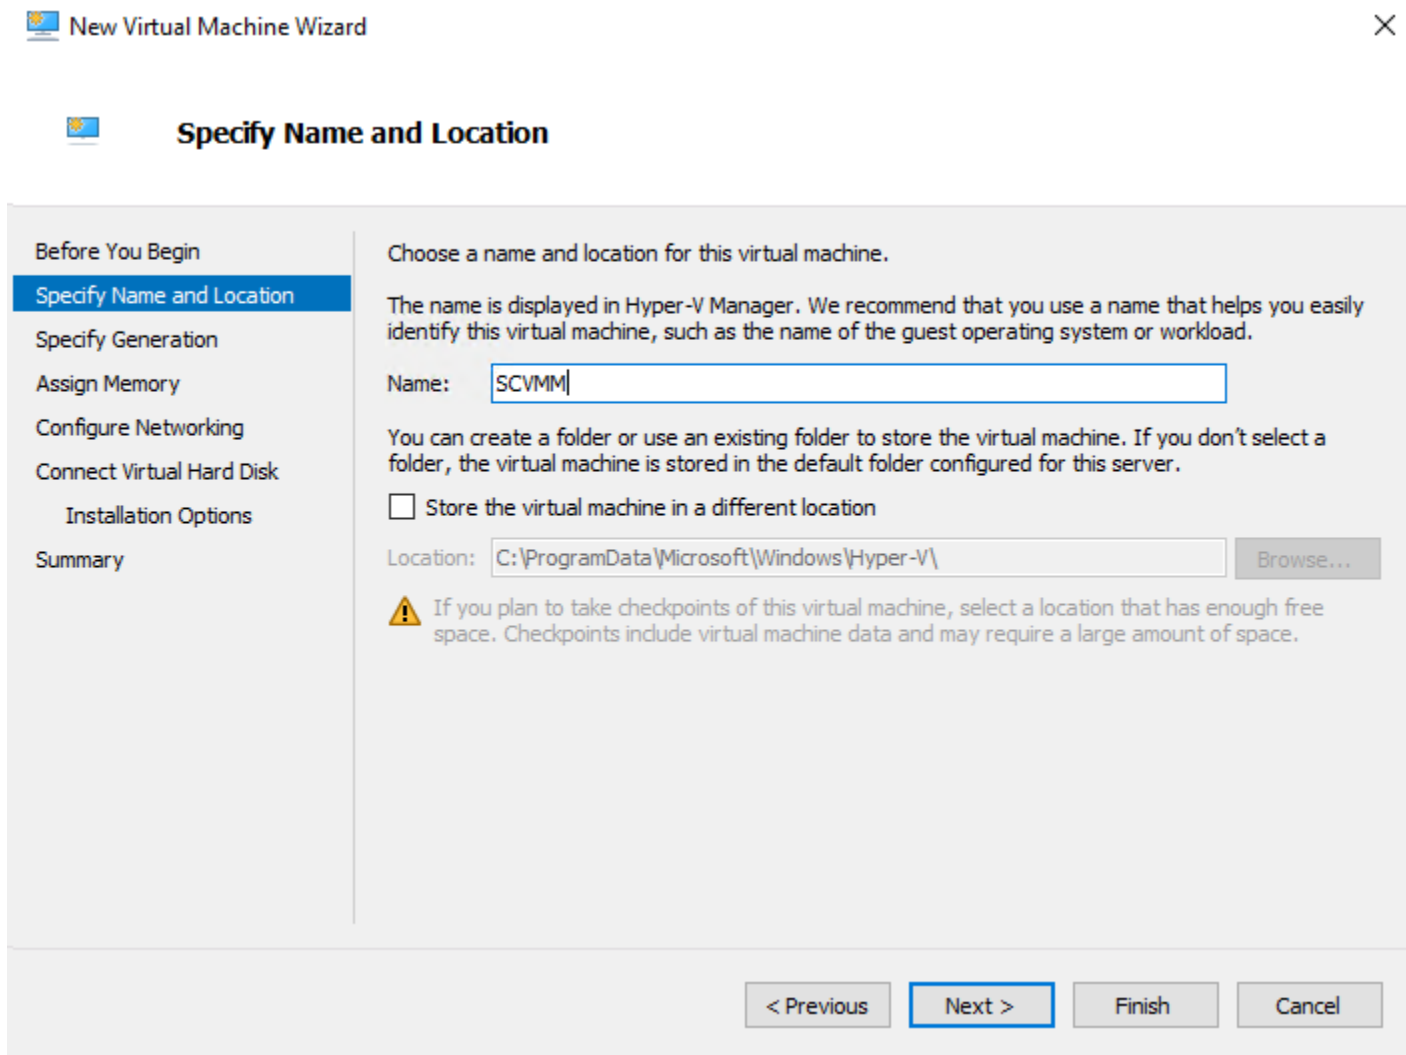
Now that the host networking setup is complete, a Microsoft Remote Desktop session can be used on both hosts going forward.

Build System Center Virtual Machine Manager (SCVMM) Virtual Machine (VM)

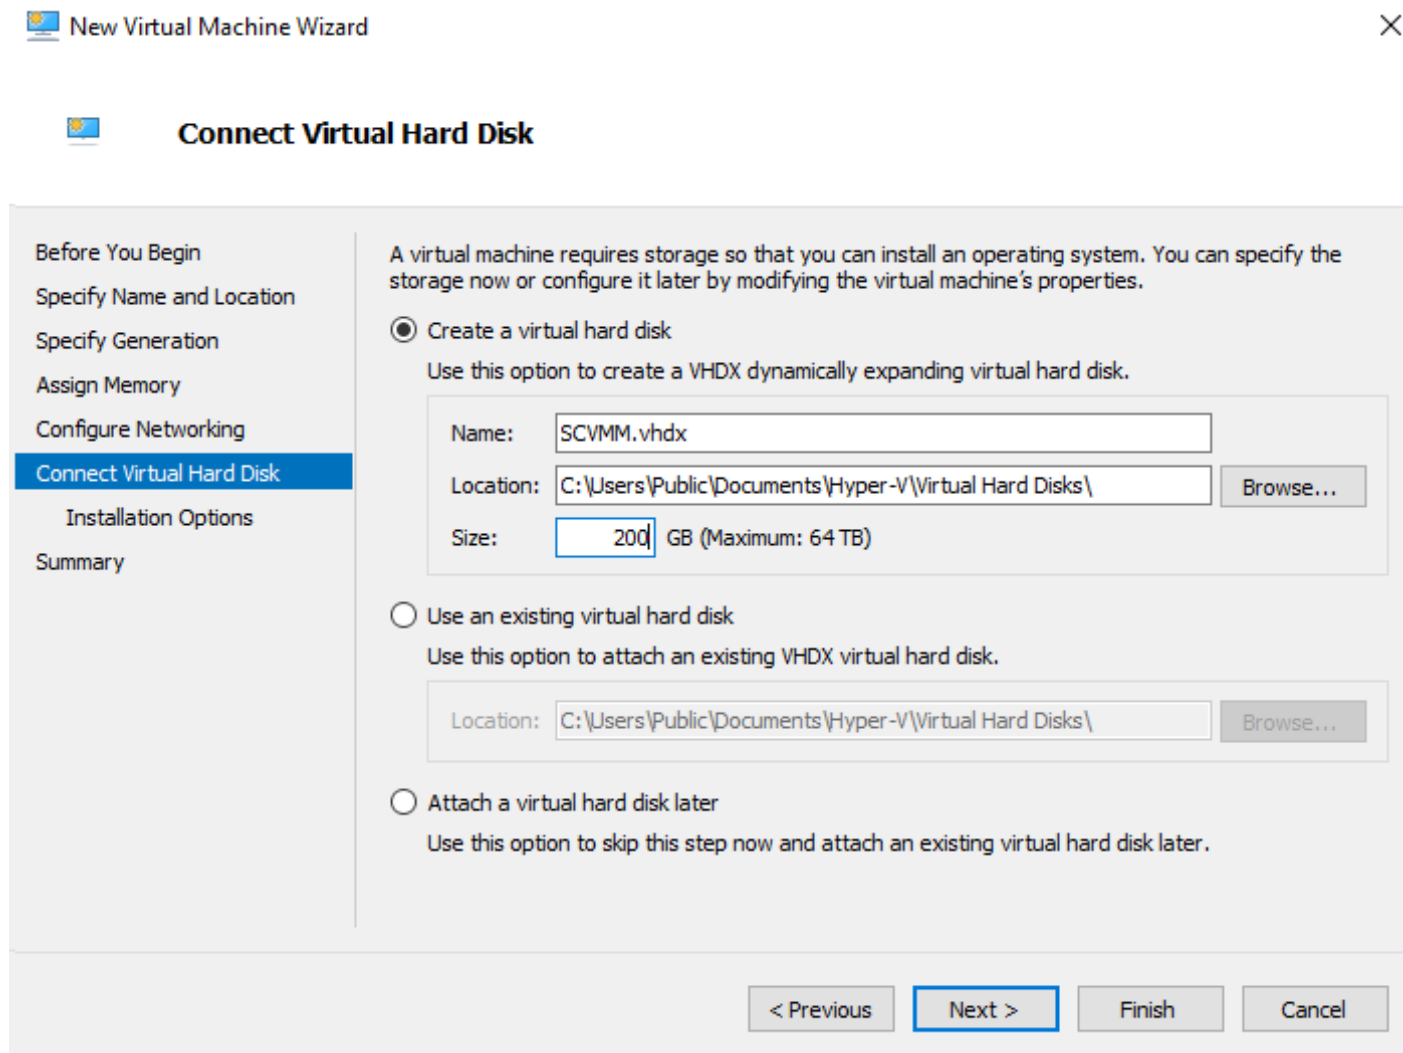
To build SCVMM virtual machine, complete the following steps:

First Hyper-V Management Host Only

1. Connect to the first Hyper-V Management Host with Windows Remote Desktop.
2. Copy the Windows Server 2016 Installation ISO and the SQL Server 2016 ISO to the host desktop.
3. From the Start Menu under Windows Administrative Tools, open Hyper-V Manager.
4. On the left, right-click the host and select New Virtual Machine.
5. In the New Virtual Machine Wizard, click Next.
6. Name the virtual machine SCVMM and leave the default location for the virtual machine selected. Click Next.



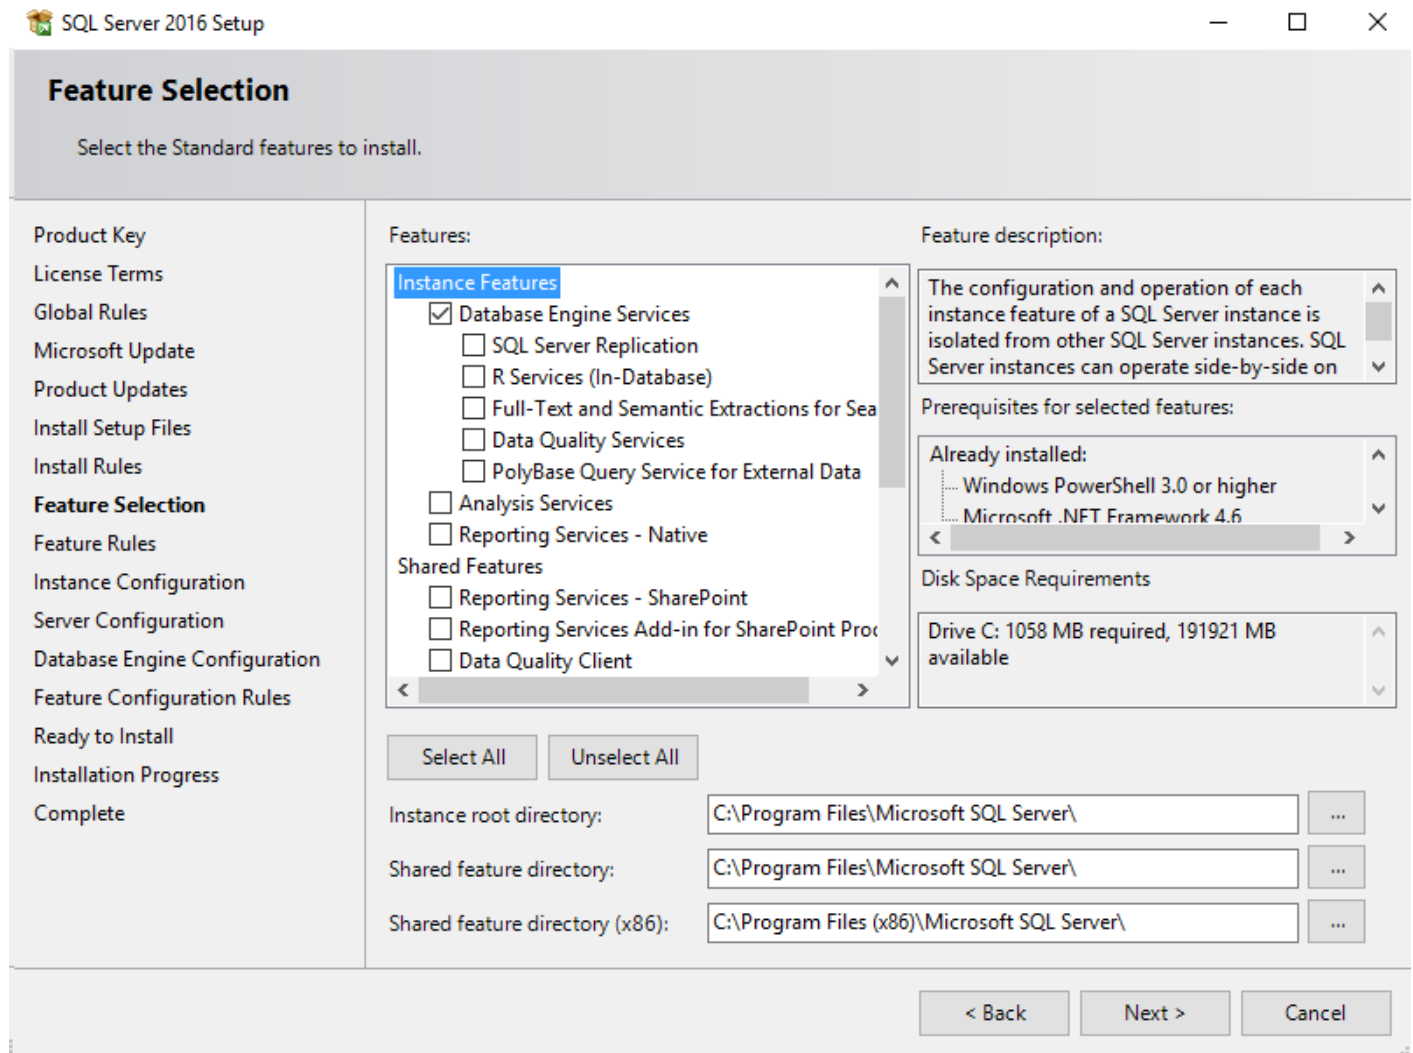
7. Select Generation 2 and click Next.
8. Enter 16384 for the Startup memory and select the checkbox for “Use Dynamic Memory for this virtual machine”. Click Next.
9. For the Networking Connection, select HV-Infra-vSwitch. Click Next.
10. Set the virtual hard disk size to 200GB and click Next.



11. Select “Install an operating system from a bootable image file”. Click Browse and browse to the Windows Server 2016 Installation ISO. Select the ISO and click Open. Click Next.
12. Verify all information and click Finish.
13. The SCVMM VM should now appear in Hyper-V Manager when the host is selected. Select the VM and on the right select Settings.
14. Under Memory, change the Minimum RAM to 4096 MB.
15. Under Processor change the Number of virtual processors to 8.
16. Since the UCS IB-MGMT VLAN is set as the default VLAN on the virtual switch network uplinks, it is not necessary to set a VLAN for this VM.
17. Click Apply then OK to complete changing the VM settings.
18. Right-click the SCVMM VM and click Connect.

19. In the SCVMM Virtual Machine Connection window, under the Action menu, select Start. Immediately press a key when you see “Press any key to boot from CD or DVD”. If the VM tries to boot from the network, use Action > Reset to reset the VM and immediately press a key when you see “Press any key to boot from CD or DVD”.
20. In the Windows Setup Window, select the appropriate language and regional format and click Next.
21. Click Install now.
22. In the Activate Windows window, click “I don’t have a product key”.
23. Select Windows Server 2016 Datacenter (Desktop Experience) and click Next.
24. Click to accept the license terms and click Next.
25. Select Custom: Install Windows only (advanced).
26. Select Drive 0 Unallocated Space and click Next.
27. Windows Installation will complete and the VM will reboot.
28. After reboot, set the Administrator password on the SCVMM VM and log into the VM as Administrator.
29. Set an IP address on the VM in the IB-MGMT subnet and join the machine to the AD Domain. On reboot, login as the local machine administrator.
30. Set the correct time zone in the VM and enable Remote Desktop.
31. Verify that the server has been Activated (this may require a Product Key to be entered).
32. In Server Manager under Local Server, turn Internet Explorer Enhanced Security Configuration Off for Administrators and click OK.
33. In the Virtual Machine Connection window, select Media > DVD Drive > Eject Windows Installation ISO.
34. Select Media > DVD Drive > Insert Disk.
35. Browse to the SQL 2016 Installation ISO and click Open.
36. Open Windows Explorer and browse to the DVD Drive where the SQL 2016 Installation ISO is mounted.
37. Double-click setup.
38. When the SQL Server Installation Center has loaded, click Installation on the left, then New Server stand-alone installation.
39. On the Product Key window, enter your product key and click Next.
40. On the License Terms screen, read the license terms, then click the checkbox next to “I accept the license terms”. Click Next.

41. On the Global Rules window, a system check runs.
42. On the Microsoft Update window, check the box to check for updates and click Next.
43. Click Next.
44. Setup files will be installed and an Install Rules check will be run. The Windows Firewall warning can be disregarded. Click Next.
45. On the Feature Selection window, select the checkbox next to “Database Engine Services” under Instance Features. Click Next.



46. A Features Rules check will run.
47. On the Instance Configuration window, click Next unless a different instance id is needed. If a different instance id is needed, change the Instance ID field and click Next.
48. At the Server Configuration Window, click Next.

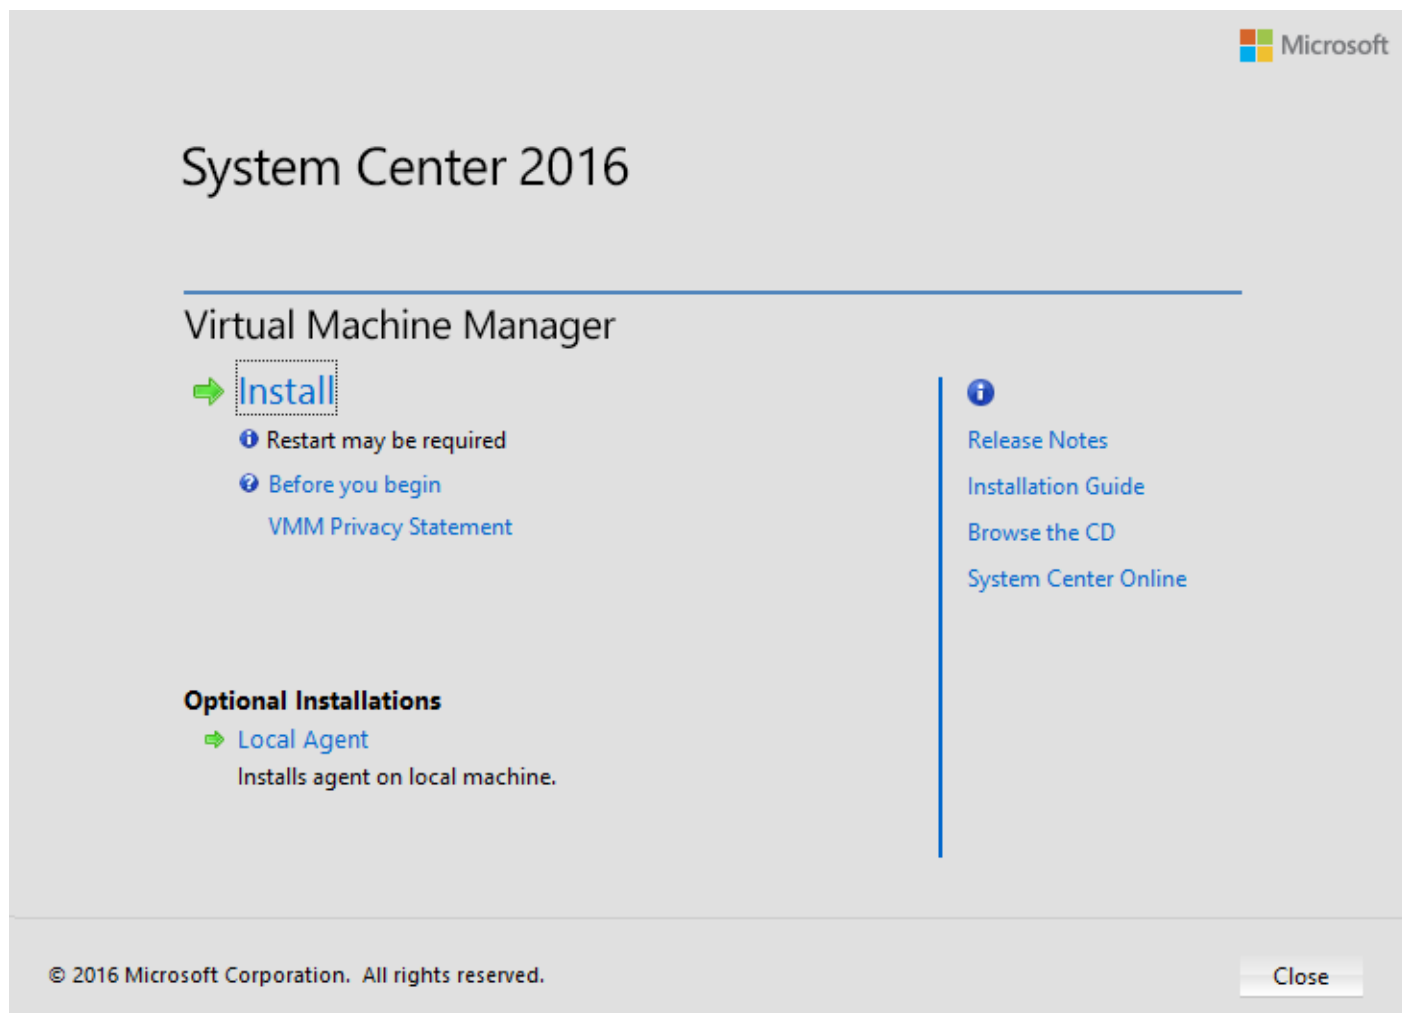
49. At the Database Engine Configuration window, click Add Current User to add the SCVMM local Administrator as a SQL Server Administrator. Click Next.
50. Click Install to install SQL Server 2016 Database Engine.
51. Once the installation has completed, click Close to close the Setup window.
52. Back at the SQL Server Installation Center window, click Install SQL Server Management Tools. A web browser will launch to download the tools. Click OK to use default Internet Explorer settings. Click to Download SQL Server Management Studio.
53. Click Run to install the SQL Server Management Tools. Click Install.
54. When the installation is complete, click Restart to restart the SCVMM VM.
55. Log in as the local Administrator.
56. Download and install the x64 version of Microsoft Command Line Utilities for SQL Server from <http://go.microsoft.com/fwlink/?LinkID=797863>.
57. Select the option to run SqlCmdLnUtils.msi.
58. At the Welcome window, click Next.
59. Accept the terms in the license agreement and click Next.
60. Click Install.
61. Click Finish.
62. Install the Windows Assessment and Deployment Kit (ADK) after downloading it from the following URL: <http://go.microsoft.com/fwlink/?LinkID=614942>. Download the Windows ADK for Windows 10, version 1607.



Even though the ADK download says it is for Windows 10, it also supports Windows Server 2016.

63. Click Run to install the ADK.
64. Click Next to install the ADK in the default location.
65. Respond to the Privacy prompt and click Next.
66. Click Accept to accept the license agreement.
67. Ensure that only Deployment Tools and Windows Pre-installation Environment are selected. Click Install.
68. When the installation completes, click Close to close the installation window.
69. Close the Web Browser Window.

70. In the Virtual Machine Connection window, use the Media menu to mount the SCVMM 2016 ISO to the DVD drive.
71. In the AD server, create an SCVMM user and place it in the Domain Admins group.
72. Log out and log back into the SCVMM VM as the SCVMM user just created.
73. Open Windows Explorer and browse to the DVD Drive where the SCVMM 2016 Installation ISO is mounted.
74. Double-click SC2016_SCVMM to open the SCVMM file extractor. Click Yes to allow the app to make changes.
75. Click Next at the Welcome window.
76. Click to accept the license agreement and click Next.
77. Change the location to “C:\System Center 2016 Virtual Machine Manager” and click Next.
78. Click Extract to extract the SCVMM files.
79. Click Finish to close the file extractor.
80. Open Windows Explorer and navigate to C:\System Center 2016 Virtual Machine Manager.
81. Double-click the setup application. Click Yes to allow the app to make changes.
82. Click Install to begin the installation.



83. Select both the VMM management server and VMM console features. Click Next.
84. Enter a Name, Organization, and the SCVMM Product key and click Next.
85. Click to accept the license agreement and click Next.
86. Click Next.
87. Click Next to install in the default location.
88. A hardware and software check will run. If a Pending Restart is necessary, restart the machine and return to this point.
89. At the Database configuration window, use the Browse button to browse AD for the local computer name. **This will populate the database Instance Name. Select the checkbox next to “Use the following credentials”.** Enter “servername\Administrator” as the User name and the local Administrator password. Click Next.

90. At the Configure service account and distributed key management window, click Select and select the domain SCVMM user created above. Enter the password for this user. Do not select the checkbox next to **“Store my keys in Active Directory”**. **Click Next.**
91. At the Port configuration window, click Next.
92. At the Library configuration window, leave the default settings and click Next.
93. Click Install.
94. When installation has completed, follow any instructions in the window and click Close and click Close again to close the installer.
95. If it is not already opened, open Virtual Machine Manager and Connect with the current Microsoft Windows session identity.
96. Unmap the SCVMM Installation ISO from the SCVMM VM.
97. In Server Manager, under Manage, select Add Roles and Features and under the feature Remote Server Administration Tools > Role Administration Tools, install the Hyper-V Management Tools feature.
98. Install all available Windows Updates on the SCVMM VM.

Deploying and Managing the Management Hyper-V Cluster Using System Center 2016 VMM

This section will focus only on configuring the Networking, Storage and Servers in VMM to deploy and manage Hyper-V failover clusters.



System Center 2016 VMM must be running in your environment.

Settings

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. To create a Run As account in VMM, complete the following steps:

1. Connect to the AD Domain and create an scvmmrunas account and place in the Domain Admins group.
2. Click Settings, and at the top of the window, click Create Run As Account.
3. In Create Run As Account specify name and optional description to identify the credentials in VMM.
4. In User name and Password specify the credentials. The credentials can be a valid Active Directory user or group account, or local credentials. The scvmmrunas account created in step 1 should be used here.
5. Clear Validate domain credentials if you don't need it, and click Finish to create the Run As account.

Fabric – Servers - I

This section details:

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


Create Host Group

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

To create a host group structure in Virtual Machine Manager (VMM) for the Hyper-V Management Cluster, complete the following steps:

1. To create a host group structure, open the Fabric workspace.
2. In the Fabric pane, expand Servers.

3. Right-click All Hosts, and then click Create Host Group. VMM creates a new host group that is named New host group, with the host group name highlighted.
4. Type Hyper-V Management for the name and then press Enter.



Add Hosts to the Host Group

When the Host Group is created, to add the Hyper-V hosts to Virtual Machine Manager, complete the following steps:

1. Open the Fabric workspace.
2. Select the just created 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.

Add Resource Wizard ×

Discovery Scope

Resource Location

Credentials

Discovery Scope

Target Resources

Host Settings

Summary

Specify the search scope for virtual machine host candidates

Search for computers by whole or partial names, FQDNs, and IP addresses. Alternatively, you may generate an Active Directory query to discover the desired computers.

Specify Windows Server computers by names
 Specify an Active Directory query to search for Windows Server computers

Enter the computer names of the hosts or host candidates that you want VMM to manage. Each computer name must be on a separate line.

Computer names:

Hyper-V-MGMT-01
 Hyper-V-MGMT-02

Skip AD verification

Examples: server1
 server1.contoso.com
 10.0.1.1
 2a01:110:1e:3:f8ffcf44:23

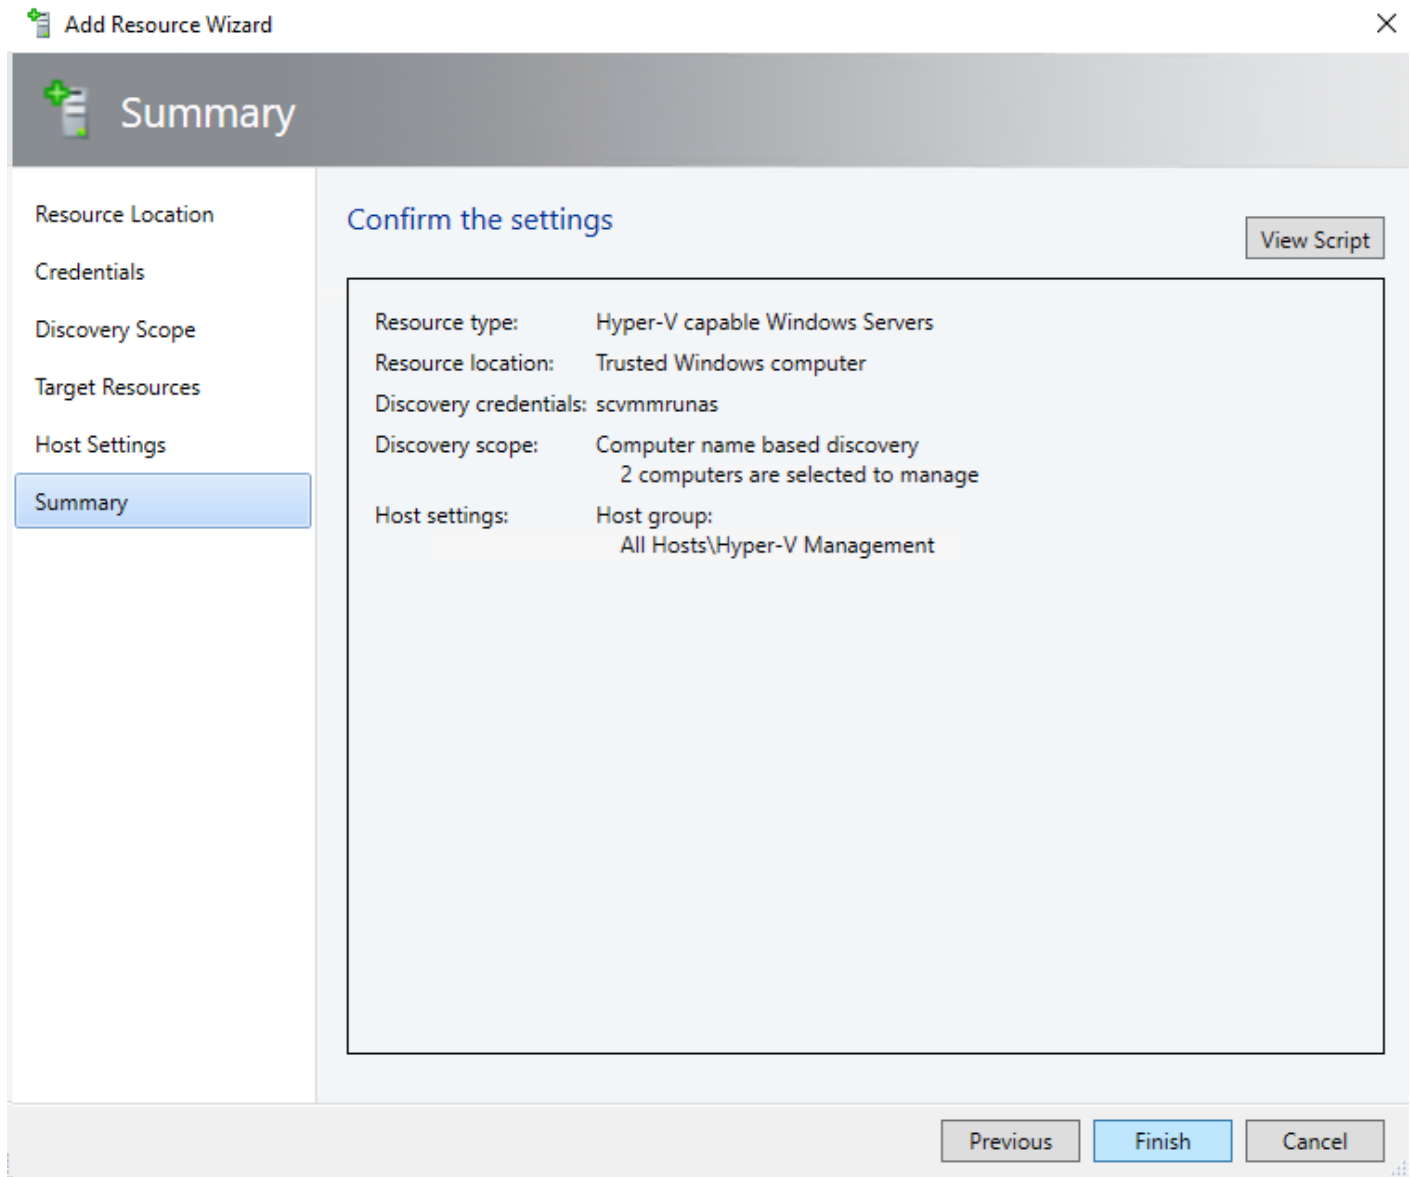
Previous
Next
Cancel

6. Under Target Resources, select the check boxes next to the two Hyper-V management hosts that need to be added. Click Next.




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, use the pulldown to select the Hyper-V Management Host Group.
8. On the Summary page, confirm the settings, and then click Finish.



9. A Jobs log window will open showing completion status. It may be necessary to reboot the two Hyper-V Management hosts. If the log indicates a reboot of the two hosts is required, on the left, select VMs and Services. In the expanded Host Group and Cluster, select the second Hyper-V-MGMT host. Right-click the host and select Restart. Shutdown the SCVMM VM. Then connect to either a console or RDP session on the first host and reboot it. Finally, once the first host has rebooted, used Hyper-V Manager to restart the SCVMM VM.



Creating APIC-Controlled Hyper-V Networking

To create ACI APIC-controlled Hyper-V Networking in SCVMM, complete the following steps. This networking can then be assigned to Hyper-V hosts.

SCVMM VM

1. Connect to the SCVMM VM with Windows Remote Desktop and login as the SCVMM Service User.
2. Open Virtual Machine Manager and create a private cloud in SCVMM by selecting the VMs and Services workspace and selecting Create Cloud at the top.
3. Name the cloud ACI-Cloud and click Next.
4. Under Resources, select the checkbox next to All Hosts and click Next.
5. Under Logical Networks, click Next.
6. Under Load Balancers, click Next.
7. Under VIP Templates, click Next.
8. Under Port Classifications, click Next.
9. Under Storage, click Next.
10. Under Library, browse to the MSSCVMMLibrary, click OK, and click Next.
11. Under Capacity, click Next.
12. Under Capability Profiles, select Hyper-V and click Next.
13. Under Replication Groups, click Next.
14. Under Summary, click Finish to complete creating the Cloud.
15. Using Internet Explorer, go to Cisco's Application Policy Infrastructure Controller (APIC) Website: <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>.
16. Select [All Downloads for this Product](#).
17. Select [APIC Software](#).
18. **Select Release 3.0(1k) and download the “MSFT Package for 3.0(1k) Release” to the Desktop.** This download will require login to cisco.com.
19. Close the web browser.
20. Right-click the **“aci-msft-pkg-3.0.1k.zip”** file and select Extract All.

21. Extract the files to the “aci-msft-pkg-3.0.1k” folder on the Desktop.
22. The extraction should open the “aci-msft-pkg-3.0.1k” folder to a window.
23. Right-click “APIC SCVMM Agent” and select Install.
24. Click Run.
25. At the Welcome window, click Next.
26. Check the box to accept the terms in the License Agreement and click Next.
27. Enter the credentials for the SCVMM Service Account and click Next.
28. Click Install to begin the installation. Click Yes to allow the app to make changes.
29. Click Finish to complete the installation.
30. Under the Start Menu, navigate to Windows PowerShell, and right-click Window PowerShell. Select More > Run as Administrator to open the Windows PowerShell Window. Run the following commands.

```
cd "C:\Program Files (x86)\ApicVMMService"

Import-Module .\ACIScvmnPSCmdlets.dll

Get-Command -Module ACIScvmnPSCmdlets
```

CommandType	Name	Version
Source		
-----	----	-----
----		--
Cmdlet	Get-ACIScvmnOpflexInfo	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	Get-ApicConnInfo	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	Get-ApicCredentials	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	New-ApicOpflexCert	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	Read-ApicOpflexCert	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	Set-ApicConnInfo	3.1.0.11
ACIScvmnPSCmdlets		
Cmdlet	Set-ApicCredentials	3.1.0.11
ACIScvmnPSCmdlets		

```
$pfpassword = ConvertTo-SecureString "MyPassword" -AsPlainText -Force
```

```

New-APICOpflexCert -ValidNotBefore 1/1/2017 -ValidNotAfter 1/1/2022 -Email
t0@domain.com -Country USA -State NC -Locality "RTP" -Organization MyOrg -
PfxPassword $pfxpassword

Successfully created:

C:\Program Files (x86)\ApicVMMService\OpflexAgent.pfx

Read-APICOpflexCert -PfxFile "C:\Program Files (x86)\ApicVMMService\OpflexAgen
t.pfx" -PfxPassword $pfxpassword

-----BEGIN CERTIFICATE-----

MIIDrDCCApSgAwIBAgIQFf+yYqe7haFKNunb2HSETDANBgkqhkiG9w0BAQ0FADBkMSEwHwYJKoZI
hvcNAQkBFhJqb2dlb3JnMkBJaXNjby5jb20xZDjAMBGNVBAoMBUNpc2NvMQswCQYDVQQIDAJQzEM
MAoGA1UEBhMDVnNBMRQwEgYDVQQDDAtPcGZsZXhBZ2VudDAeFw0xNzAxMDEwMDAwMDBaFw0yMjAx
MDEwMDAwMDBaMGQxITAfBgkqhkiG9w0BCQEWEmpvZ2VvcncyQGhpc2NvLmNvbTEOMAwGA1UECgwF
Q2lZ28xZCZAJBgNVBAGMAk5DMQwwCgYDVQQGEwNVU0ExFDASBgNVBAMMC09wZmxleEFnZW50MIIB
IjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEArg90/UoLzg5/xdIdgAieXOBROQBjDieHkung
/DK803q3rMBU8/sMN9JkxudWPNToufTElMahruLf2oNjXJNOJEvltPa9UnipMjUgn0f+EJU5rnS
egF0zRyKVDrfoghhoIiBaFgBZ5m+m2KcBaMqH39b3IXFLytRxqEBhs/WQkKH5eNbddHOph56jwN
jIimy+IvcQHpVhqvAR+drXU9fsArNRFFw+4Q+ZAVA4CoBAvSTjD7wMlHDwEG0aH7xP88+YqPoBhr
XhDTW6kJ9yfpBH6oYh9ZWnsJRmBykH2zd9ut4oX2GyhylBgLRCOFc5JbV9wzbNR0VoHiWm1fGgph
HwIDAQABo1owWDASBgNVHRMBAf8ECDAGAQH/AgEAMBMGA1UdJQQMMAoGCCsGAQUFBwMBMB0GA1Ud
DgQWBQBqB+cK2vOLBV3E1ht5DANDsyEevjAOBgNVHQ8BAf8EBAMCAQYwDQYJKoZIhvcNAQENBQAD
ggEBAKkJs1stbRBGulgbOWImaeLqdWWfsKGlqsL8ytaTfoW8rmLhMgPX8lR3HHwc06EMaQPQriB
ZDwnJACaaB4l/tBDUfbDobPhOo9hRtWa6tVWnu/TUiic+QACuSlegQGeB9voNiSXgCKizKvsMM2r
9nG0UOgZg7CicJJFeCofr+kIyQ6oypDFhA4M4Qu0MMeLKppqLUxsCfptf6RhTfPsw7zToExY/49A
w2xC7rgG8uK1lh+k06z5lL2kt4Wu2ij9s4UHphYknYsffp052c+hIv6mhXmk50o8HtPG4jJV3nY2
+lpk/tentloeAqgnUm+Qd8dMH4I6CDdiE1BBEVuul7U=

-----END CERTIFICATE-----

```

31. Using Chrome (this will need to be installed on the SCVMM VM), connect to the APIC GUI and log in as admin.
32. On the menu bar at the top, select Admin > AAA.
33. In the Navigation pane, choose Security Management > Local Users and click on admin.
34. Right-click admin and select Create X509 Certificate.

35. Name the certificate “OpflexAgent”. In the Data field copy and paste the output of the Read-OpflexAgent.ps1 powershell cmdlet. Click Submit.

Create X509 Certificate



Define a User Certificate

Name:

Data:

```
-----BEGIN CERTIFICATE-----
uS1egQGeB9voNiSXgCKizKvsMM2r
9nG0UOgZg7CicJJFeCofr+kIyQ6oypDFhA4M4Qu0MMeLKpp
qLUxsCfptf6RhTfPsw7zToExY/49A
w2xC7rgG8uK1lh+k06z5IL2kt4Wu2ij9s4UHpHYknYsffpO52c+
hlv6mhXmk50o8HtPG4jJV3nY2
+lpk/tentloeAqgnUm+Qd8dMH4l6CDdiE1BBEVuu17U=
-----END CERTIFICATE-----
```

Cancel

Submit

36. In the Properties pane under User Certificates, the certificate will now be displayed.
37. Back on the SCVMM VM, in the still open Powershell window, enter `mmc` and press Enter.
38. In Console Root, under File, select Add/Remove Snap-In.
39. In the Available Snap-Ins field, choose Certificates and click Add.
40. In the Certificates snap-in dialog box, choose the Computer Account radio button and click Next.
41. In the Select Computer dialog box, choose the Local computer radio button and click Finish.
42. Click OK to go back to the main MMC Console window.
43. In the MMC Console window, double-click Certificates (Local Computer) to expand its view.
44. Expand Certificates > Personal and right-click Certificates under Personal. Select All Tasks > Import.
45. In the Certificate Import Wizard, click Next.
46. Browse to the OpflexAgent file created earlier and click Next.



You will need to change the file selection type to Personal Information Exchange to see the OpflexAgent file.



Certificate Import Wizard

File to Import

Specify the file you want to import.

File name:

C:\Program Files (x86)\ApicVMMService\OpflexAgent.pfx

Browse...

Note: More than one certificate can be stored in a single file in the following formats:

Personal Information Exchange- PKCS #12 (.PFX,.P12)

Cryptographic Message Syntax Standard- PKCS #7 Certificates (.P7B)

Microsoft Serialized Certificate Store (.SST)

Next

Cancel

47. Enter the password provided when the certificate was created. Select the checkbox next to “Mark this key as exportable.” Leave the checkbox next to “Include all extended properties.” selected. Click Next.

48. The “Personal” certificate store should be selected. Click Next.

49. Click Finish and OK to complete importing the certificate.
50. Close the mmc console.
51. Back in the still open Windows PowerShell window, set and verify the APIC connection settings.

```
Set-APICConnInfo -ApicNameorIPAddress <APIC-IP> -CertificateSubjectName
OpflexAgent
```

```
Get-APICConnInfo
```

```
EndpointAddress      :
Username             :
Password             :
ApicAddresses        : 192.168.1.46
ConnectionStatus     : Connected
adminSettingsFlags   : 0
certificateSubjectName : OpflexAgent
ExtensionData        :
```

52. From the SCVMM VM Start Menu, open Remote Desktop Connection and enter the hostname or IP address of the first Hyper-V Management host. Select Show options. Under the Local Resources tab, make sure the Clipboard and Drives are selected under Local devices and resources. Click Connect and Connect again. The SCVMM Service Account should be selected. Enter the password for this account and click OK.
53. Open Windows Explorer and navigate to the C drive on the SCVMM VM. Navigate to Users > SCVMM Service Account Name > Desktop > aci-msft-pkg-3.0.1k.
54. Copy the APIC Hyper-V Agent.msi file to the Desktop. Right-click the APIC Hyper-V Agent.msi file copy and choose Install.
55. Click the checkbox to accept the terms in the License Agreement and click Install. Click Yes to allow the app to make changes.
56. Click Finish to complete the installation.
57. Back in Windows Explorer on **the C drive on the SCVMM VM**, navigate to **“Program Files (x86)\ApicVMMService”**. **Copy the OpflexAgent.pfx file to the local Desktop.**
58. Click the search icon to the right of the Start Menu icon and enter Run. Click to open the Run Desktop App.
59. In the box to the right of Open, enter mmc and click OK. Click Yes to allow the app to make changes.

60. In Console Root, under File, select Add/Remove Snap-In.
61. In the Available Snap-Ins field, choose Certificates and click Add.
62. In the Certificates snap-in dialog box, choose the Computer Account radio button and click Next.
63. In the Select Computer dialog box, choose the Local computer radio button and click Finish.
64. Click OK to go back to the main MMC Console window.
65. In the MMC Console window, double-click Certificates (Local Computer) to expand its view.
66. Expand Certificates > Personal and right-click Certificates under Personal. Select All Tasks > Import.
67. In the Certificate Import Wizard, click Next.
68. Browse to the OpflexAgent file on the Desktop and click Next.



You will need to change the file selection type to Personal Information Exchange to see the OpflexAgent file.

69. Enter the password provided when the certificate was created. Select the checkbox next to “Mark this key as exportable.” Leave the checkbox next to “Include all extended properties.” selected. Click Next.
70. The “Personal” certificate store should be selected. Click Next.
71. Click Finish and OK to complete importing the certificate.
72. Close the mmc console.
73. You can now delete the two files that were added on the Desktop.
74. Log out of the Remote Desktop Session.
75. Repeat steps 52-74 to install the APIC Hyper-V Agent and opflex certificate on the second Hyper-V Management host.
76. Back on the SCVMM Remote Desktop, using Chrome, connect to the APIC GUI as admin.
77. On the menu bar, choose Fabric > Access Policies.
78. In the Navigation pane, expand Global Policies > Attachable Access Entity Profiles (AEP). If the External L2 Bridged Domain for the 6332 UCS Fabric Interconnects was named UCS when vPCs were created, the AEP will be named UCS_AttEntityP. Select UCS_AttEntityP under Attachable Access Entity Profiles.
79. In the Properties pane, select the checkbox next to Enable Infrastructure VLAN and click Submit and then Submit Changes.

Attachable Access Entity Profile - UCS_AttEntityP

Policy Operational Faults History

Properties

Name: UCS_AttEntityP

Description: optional

Enable Infrastructure VLAN:

Domains (VMM, Physical or External) Associated to Interfaces:

name	State
UCS (L2)	formed

80. On the menu bar, choose VM Networking > Inventory.
81. In the Navigation pane, right-click the VM Provider Microsoft and choose Create SCVMM Domain.
82. In the Create SCVMM Domain window, in the Name field, enter a Domain Name (APIC-MS-vSwitch).
83. Use the Associated Attachable Entity Profile pulldown to select UCS_AttEntityP.
84. Use the VLAN Pool pulldown to select Create VLAN Pool.
85. In the Create VLAN Pool window, enter a name for the VLAN Pool (VP-APIC-MS-vSwitch).
86. **Leave Dynamic Allocation selected and click the “+” to the right of Encap Blocks to add a VLAN range.**
87. Enter the VLAN IDs for the start and end of the VLAN range that was entered in the UCS (1200-1299) and click OK.
88. Click Submit to complete creating the VLAN pool. Using the pulldown, select the VLAN Pool just created.
89. **Click the “+” to the right of SCVMM Controllers to add the SCVMM controller.**
90. In the Create SCVMM Controller window, put the SCVMM VM hostname in the Name field. In the Fully Qualified Domain Name (FQDN) field enter the fqdn of the SCVMM VM. In the SCVMM Cloud Name field, enter the SCVMM Cloud Name of the cloud created above (ACI-Cloud). Click OK.
91. **Use the Port Channel Mode pulldown to select “MAC Pinning-Physical-NIC-load”.**

Create SCVMM Domain



Specify SCVMM Domain users and controllers

Name:

Associated Attachable Entity Profile:

VLAN Pool:

Delimiter:

SCVMM Controllers:

Name	IP	Type	Stats Collection
aci-scvmm	aci-scvmm.flexpod.cis...	SCVMM	Disabled

Port Channel Mode:

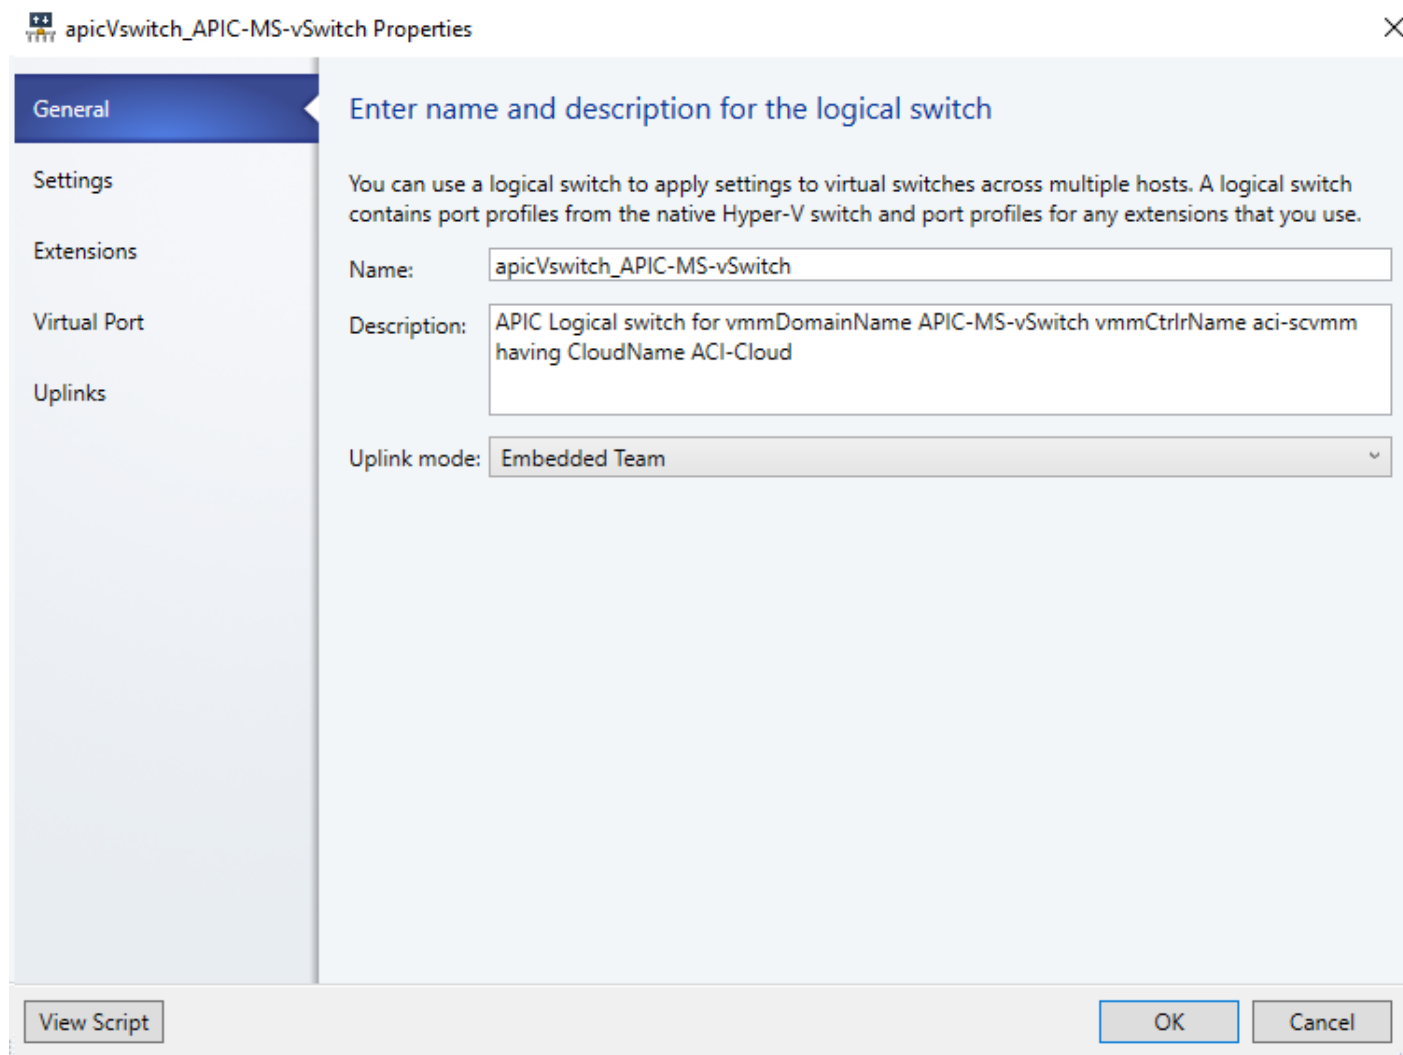
Cancel

Submit

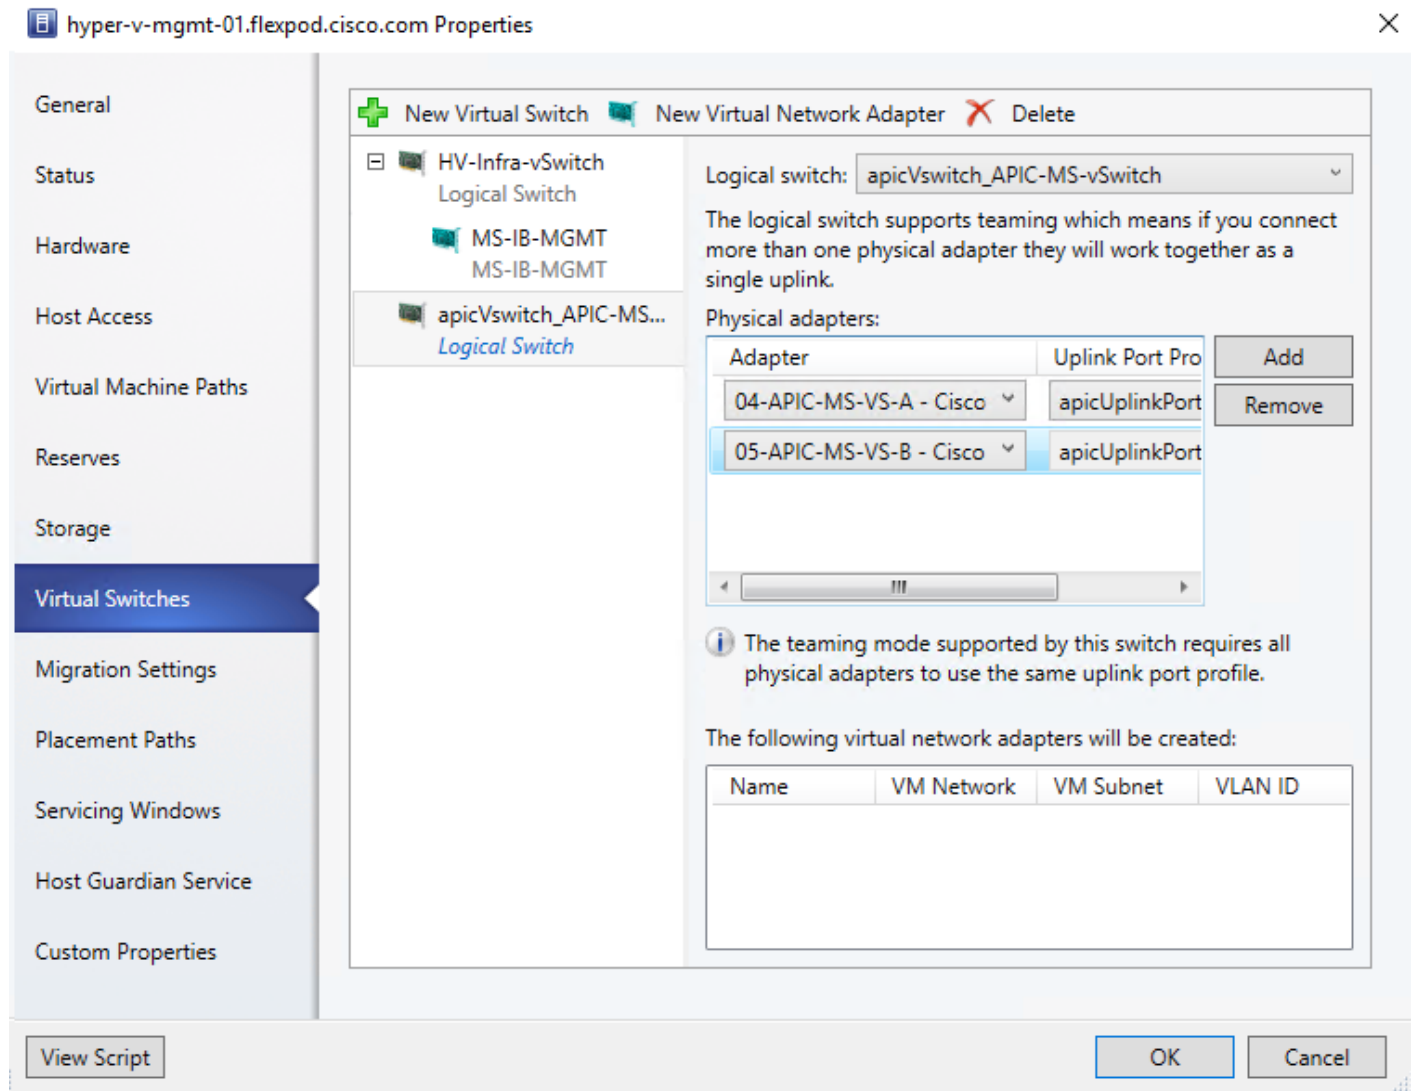
92. Click Submit.

93. Back in Virtual Machine Manager verify the APIC-VMM Integration by selecting the Fabric Workspace and selecting Networking > Logical Switches, Networking > Logical Networks, and Networking > Port Profiles. **Verify entries beginning with “apic”. Select the VMs and Services Workspace and select VM Networks. Again, verify entries beginning with “apic”.**

94. In Virtual Machine Manager, under the Fabric Workspace, select Networking > Logical Switches. In the center pane, under Logical Switches, right-click the newly-created Logical Switch from APIC and select Properties. Under General, change the Uplink mode to Embedded Team and click OK.







95. In Virtual Machine Manager, select the VMs and Services workspace and expand All Hosts > Hyper-V Management.
96. Right-click the first Hyper-V Management host and choose Properties.
97. **Select Virtual Switches on the left, then select “+ New Virtual Switch”. Choose New Logical Switch.**
98. The apicVswitch should be selected. Under Physical Adapters, select 04-APIC-MS-VS-A and add 05-APIC-MS-VS-B. The Uplink Port Profile should be populated automatically.



99. With the apicVswitch still selected, at the **top**, select **“New Virtual Network Adapter”**.
100. Name the virtual network adapter <hostname-vtep>. The ACI System VLAN should already be filled in.
101. Click OK and Yes to add the virtual switch to the host.
102. Repeat this process to add the APIC-controlled virtual switch to the second Hyper-V Management host.
103. Back in the Cisco ACI APIC GUI, select VM Networking > Inventory > Microsoft > Your APIC Virtual Switch. Set the vSwitch Policies as shown below and click Submit and Submit Changes.

vSwitch Policies

Port Channel Policy:	MAC-Pinning	▼	
LLDP Policy:	LLDP-Enabled	▼	
CDP Policy:	CDP-Enabled	▼	
STP Policy:	BPDU-FG-Enabled	▼	

104. Expand the vSwitch, Controllers, the SCVMM, and Hypervisors. Select each host and verify that the OPFLEX Status is Connected under General on the right.
105. Select Tenants > common > Application Profiles > MS-IB-MGMT > Application EPGs > MS-Core-Services. Expand MS-Core-Services. Right-click Domains (VMs and Bare-Metals) and select Add VMM Domain Association. Use the pulldown to select the Microsoft vSwitch and set the Deploy Immediacy and Resolution Immediacy to Immediate. Click Submit.

Add VMM Domain Association



Choose the VMM domain to associate

VMM Domain Profile:

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

VLAN Mode: Dynamic Static

Delimiter:

Cancel

Submit

106. Repeat the procedure in the previous step to associate the Infra-SMB-1, MS-Clust, MS-LVMN, and the IB-MGMT EPGs in the FP-Foundation tenant to the Microsoft vSwitch.
107. Back in Virtual Machine Manager under the VMs and Services workspace, select VM Networks. Ensure that the five networks associated in steps 105 and 106 are listed as VM Networks.



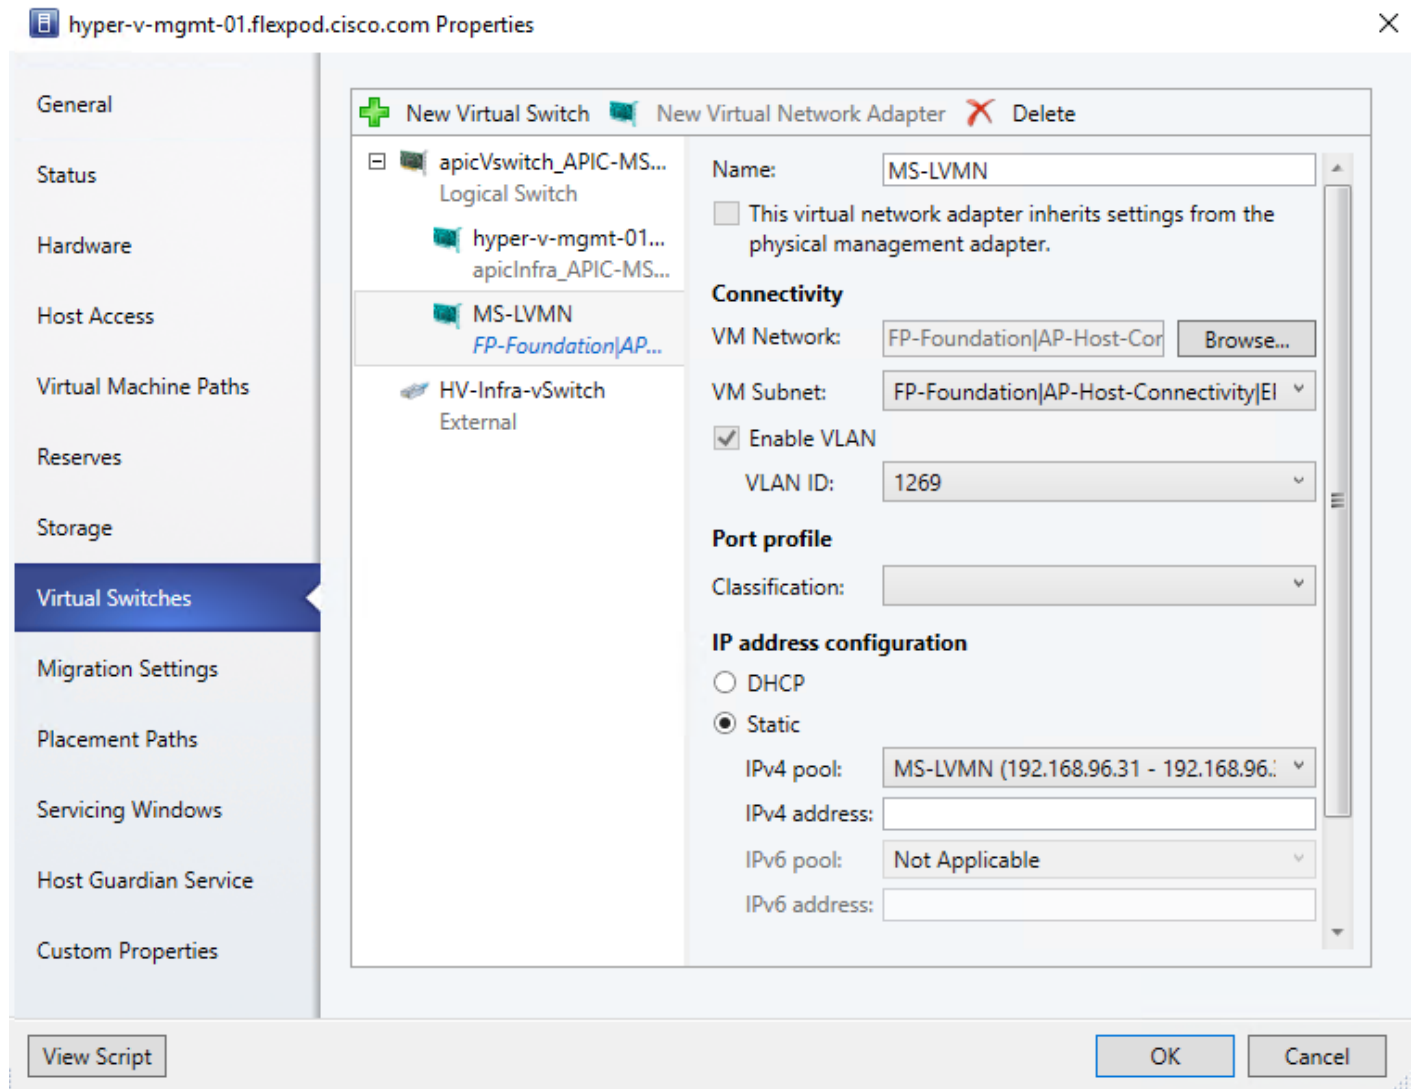
Any VMs created in Virtual Machine Manager will need to use VM Networks in the APIC-controlled vSwitch. To use networks from the manually built Hyper-V vSwitch on the two management hosts, VMs would need to be created on the host in either Failover Cluster Manager or Hyper-V Manager.

Table 10 EPGs and Subnets for Logical Network IP Pools

EPG Name	Pool Name	Subnet	Pool Start IP	Pool End IP
EPG-MS-LVMN	MS-LVMN	192.168.96.254/24	192.168.96.31	192.168.96.39
EPG-MS-Clust	MS-Clust	192.168.97.254/24	192.168.97.31	192.168.97.39

EPG Name	Pool Name	Subnet	Pool Start IP	Pool End IP
EPG-Infra-SMB	Infra-SMB	192.168.53.254/24	192.168.53.31	192.168.53.39

108. For each entry in the table above, in Virtual Machine Manager in the Fabric workspace, select Networking > Logical Networks.
109. In the center pane, right-click apicLogicalNetwork_APIC-MS-vSwitch and select Create IP Pool.
110. Name the Pool according to the Pool Name in the table and click Next.
111. Select the Network site according to the line in the table, ensure the IP subnet is correct and click Next.
112. Enter the starting and ending IP address from the table and click Next.
113. Click Next three times and click Finish to create IP Pool.
114. Repeat these steps for each row in the table.
115. In Virtual Machine Manager in the VMs and Services workspace, right click on the first Hyper-V Management host and select Properties.
116. **Select Virtual Switches. Make sure the apicVswitch is selected and select “New Virtual Network Adapter”.**
117. Give the Adapter the same name as the Pool Name in the table above.
118. Use the Browse button and select the appropriate VM Network and click OK.
119. Under IP address configuration, select Static and select the appropriate IP pool.



120. Click OK to complete creating the Virtual Network Adapter.
121. Repeat the above steps for all three rows in the table.
122. Repeat the above steps for the second Hyper-V Management host.
123. Using Remote Desktop, connect to both Hyper-V Management hosts and open Powershell as Administrator. Type `netsh interface ipv4 show subinterface` to see the MTU of each interface. Then, use the following example command to set the MTU to 9000 for the MS-LVMN, MS-Clust, and Infra-SMB interfaces:

```
netsh interface ipv4 set subinterface "vEthernet (MS-LVMN)" mtu=9000
store=persistent
```

Create Windows Failover Cluster

To create a Windows Failover Cluster, complete the following steps:



Be sure to create DNS records for the Cluster name. The IP address for cluster management should be on the IB-MGMT Subnet.

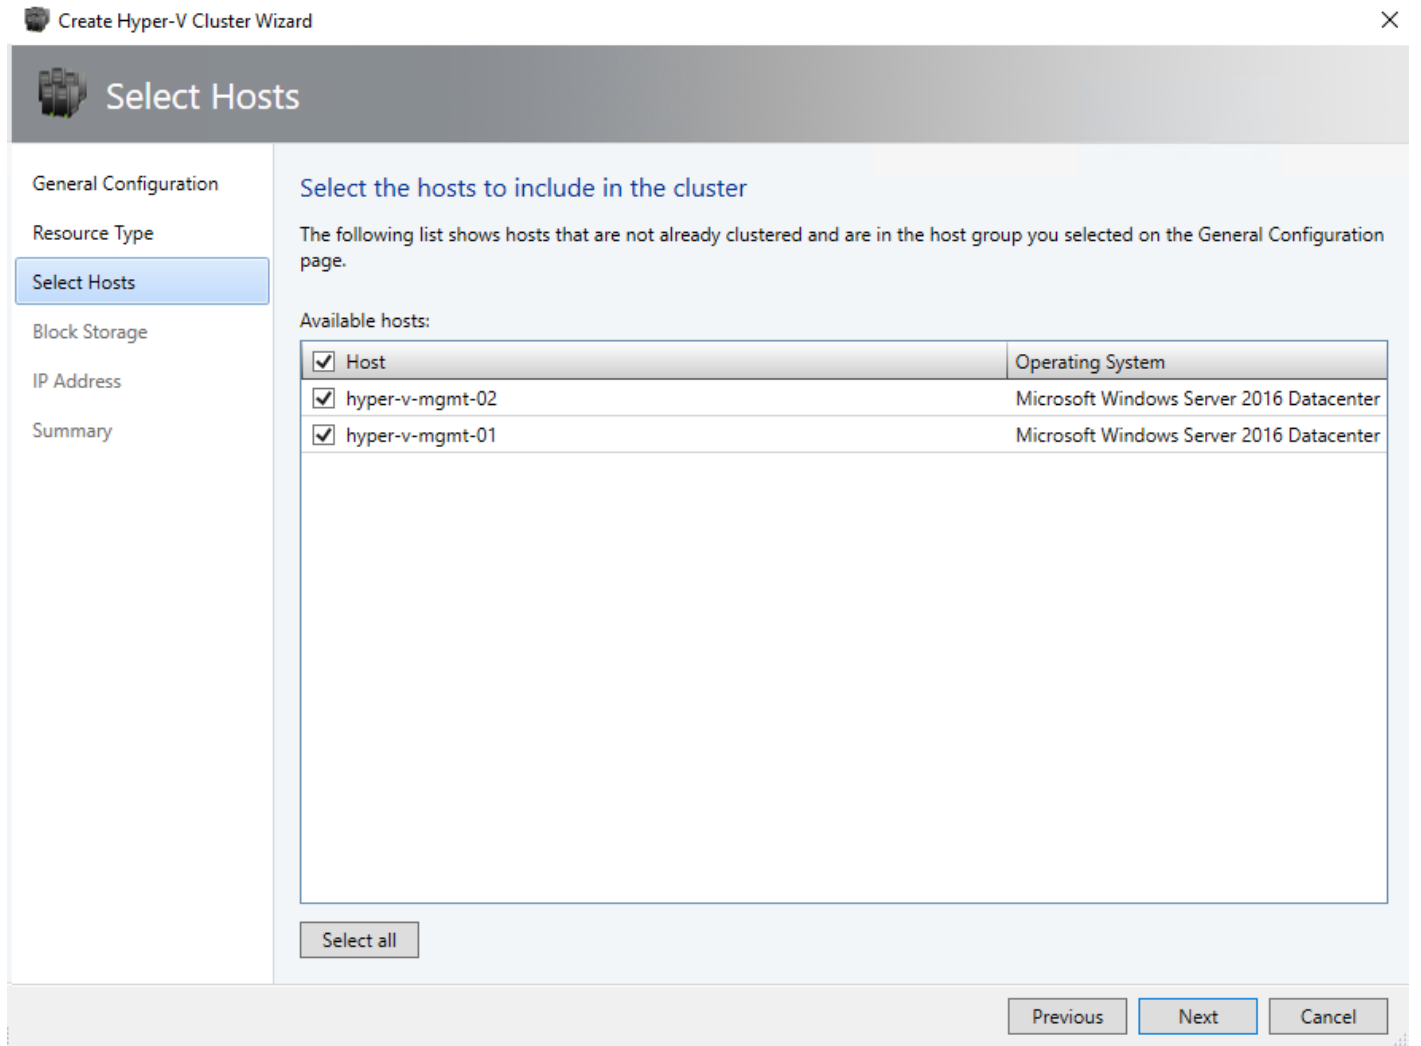
1. In the VMM console, click Fabric > Create > Hyper-V Cluster to open the Create Hyper-V Cluster wizard.
2. In General Configuration, specify a cluster name and choose the Hyper-V Management host group in which the existing Hyper-V hosts are located. Click Next.

3. In Resource Type, select the SCVMM Run As account that you'll use to create the cluster. Make sure **“Existing servers running a Windows Server operating system”** is selected and click Next.

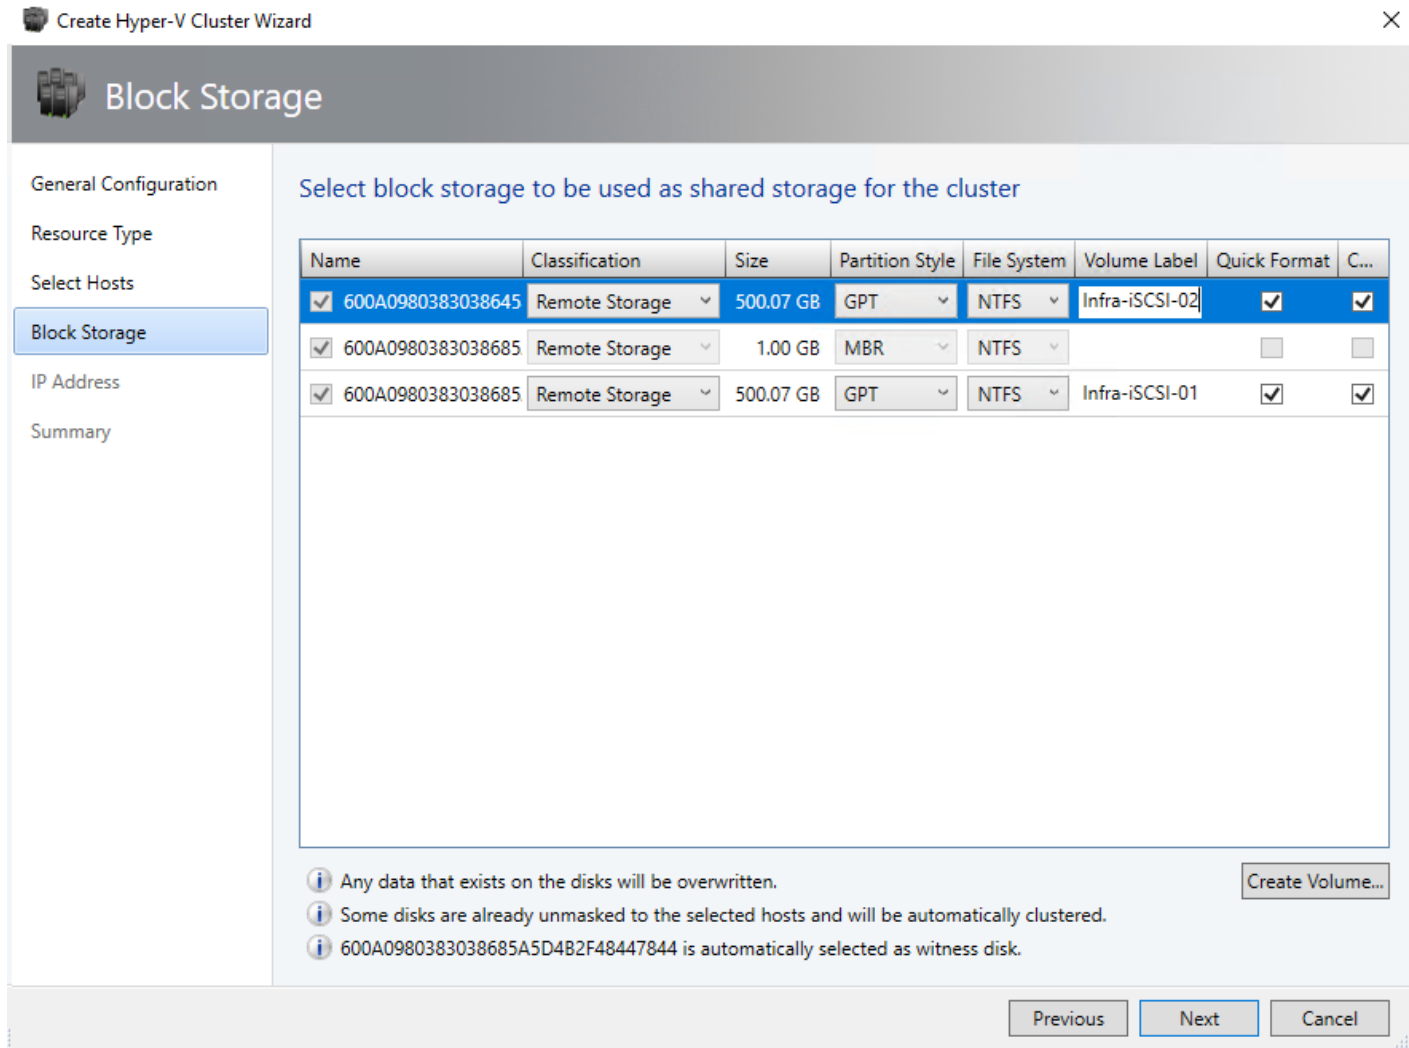


The accounts that you use must have administrative permissions on the servers that will become cluster nodes, and must belong to the same domain as the Hyper-V hosts that you want to cluster. Also, the account requires Create Computer objects permission in the container that is used for Computer accounts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

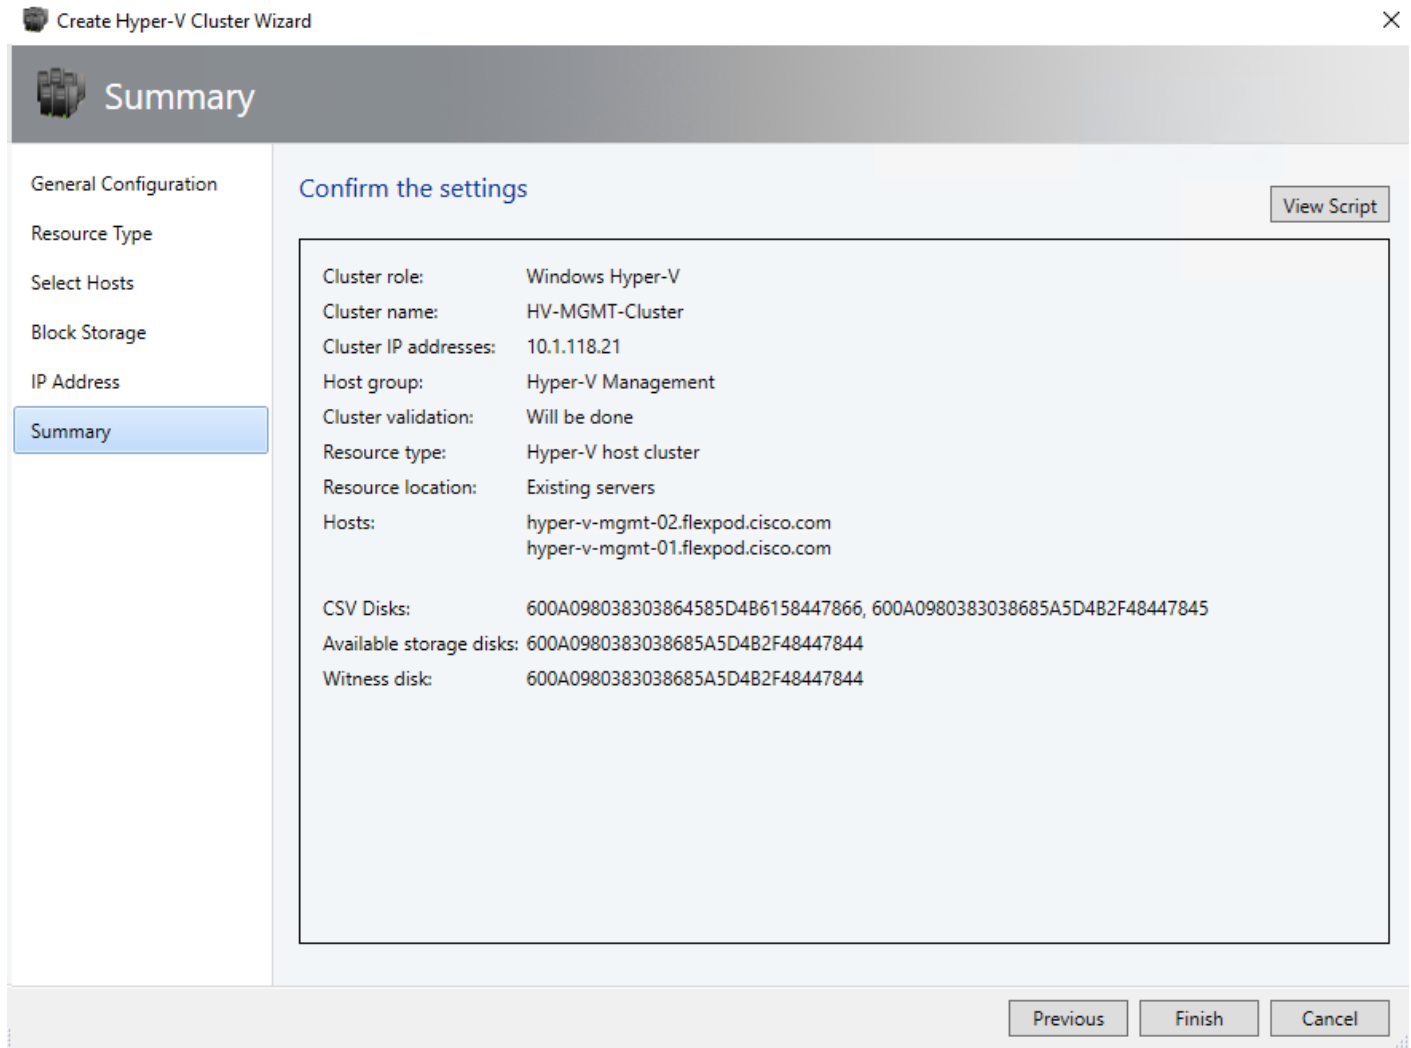
4. In Nodes, select the Hyper-V host servers that you want to include in the cluster. Click Next.



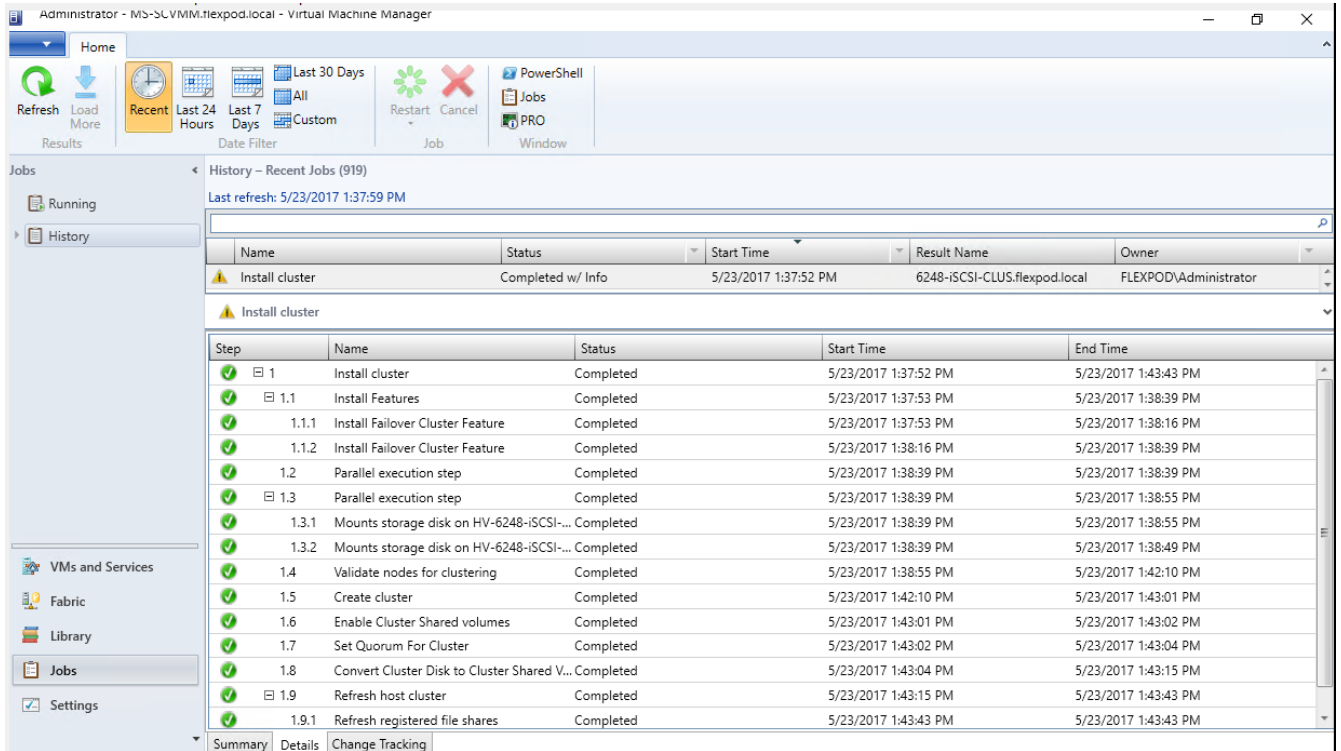
5. In Block Storage, for the two 500GB disks, select the GPT partition style, Quick Format, and CSV. For the 500 GB disk with the Name closest to the name for the 1 GB disk, make the Volume Label Infra-iSCSI-01. Make the Volume Label for the other 500GB disk Infra-iSCSI-02. Click Next.




- In IP address, select the IB-MGMT/Core-Services subnet and type in the IP address you want to use for the cluster. Click Next.
- In Summary, confirm the settings and then click Finish.

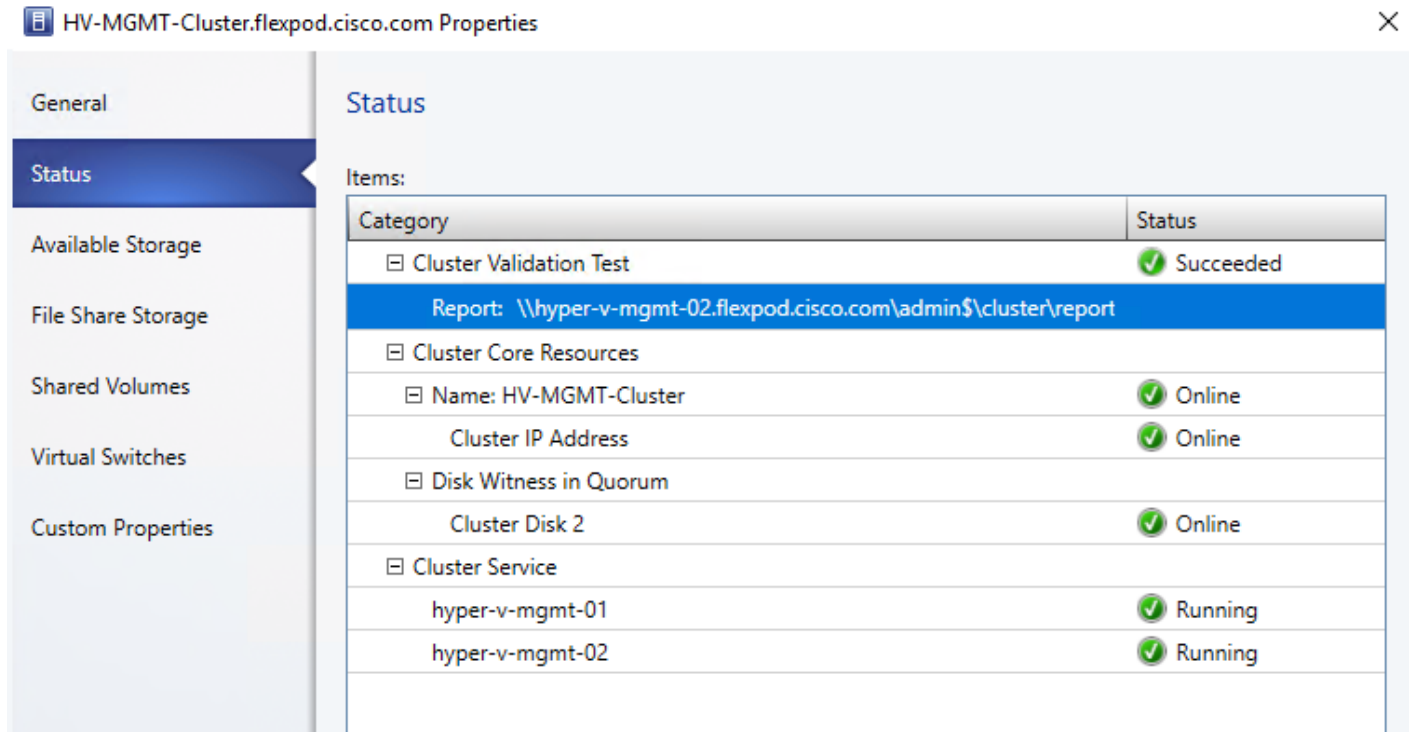


8. You can go to the jobs workspace and click on “Install Cluster” job to see the status of cluster installation. Fix and troubleshoot any errors or warnings and revalidate the cluster.



9. After the cluster is installed, a new cluster icon is seen after expanding the Servers > All Hosts > Hyper-V Management 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 11 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.

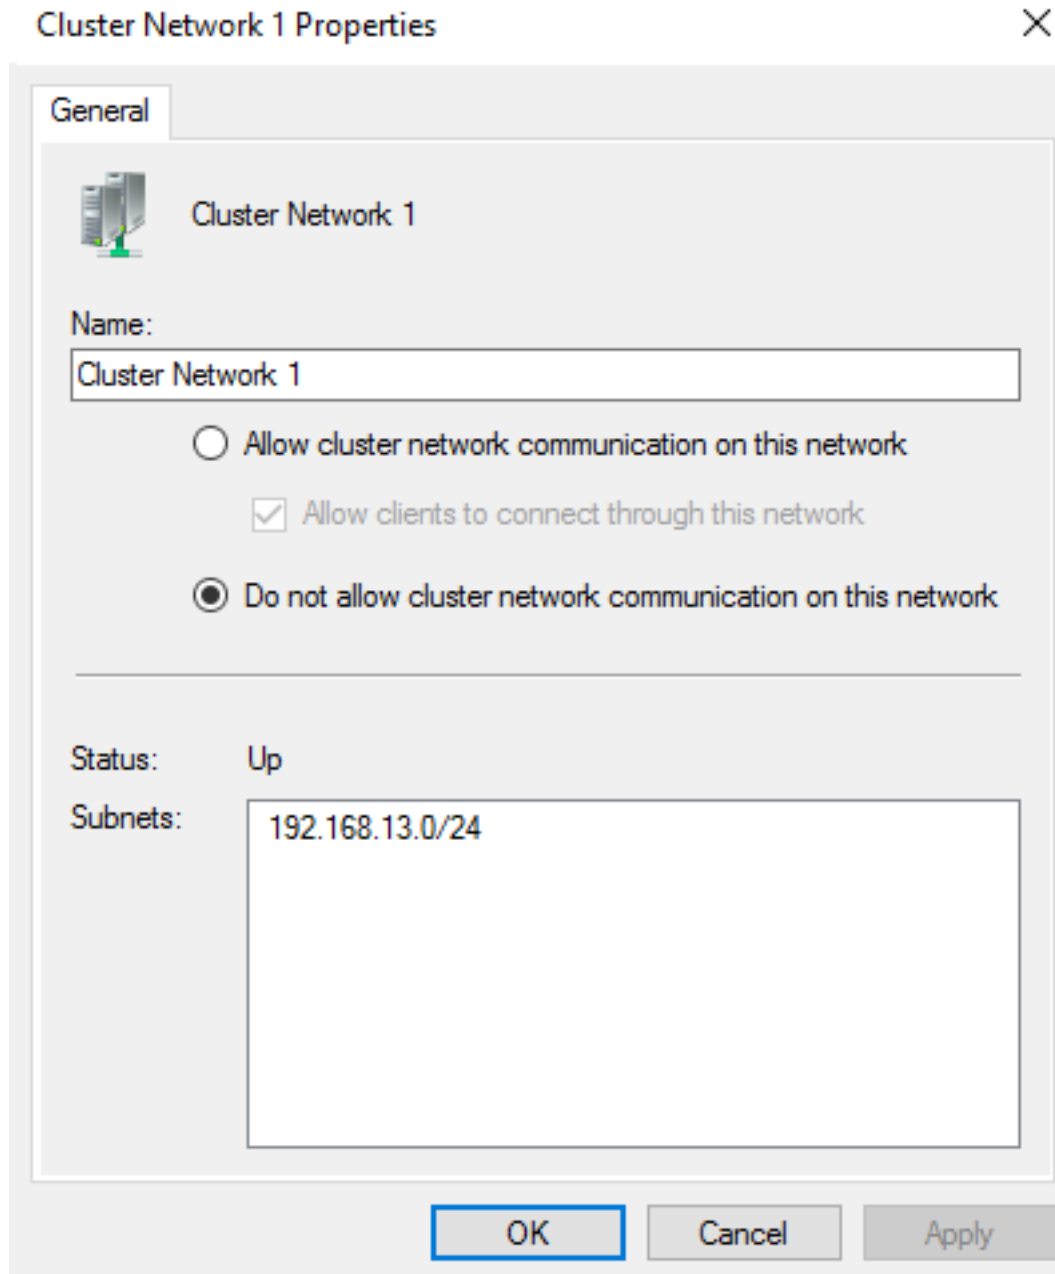
Network Type	Recommended Setting
Storage	Do not allow cluster network communication on this network

1. On the SCVMM VM, open Failover Cluster Manager and connect to the Failover Cluster just created. Click Networks in the navigation tree.

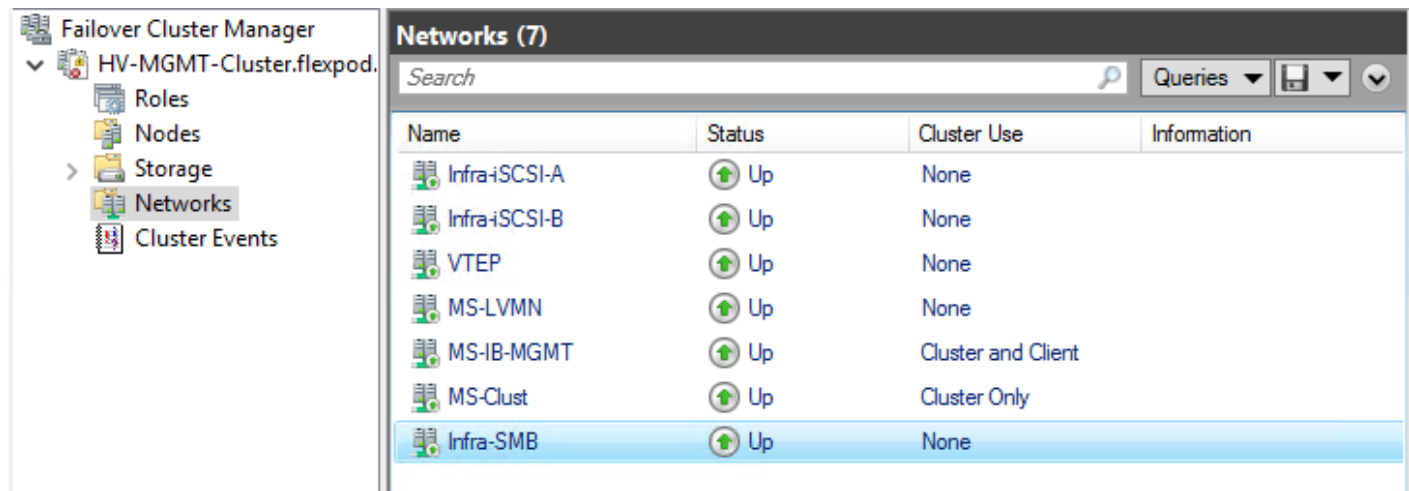


It may be necessary to install the Failover Clustering Tools Feature under Features > Remote Server Administration Tools > Feature Administration Tools in the Add Roles and Features Wizard to install Failover Cluster Manager.

2. In the Networks pane, right-click a network, and then click Properties.



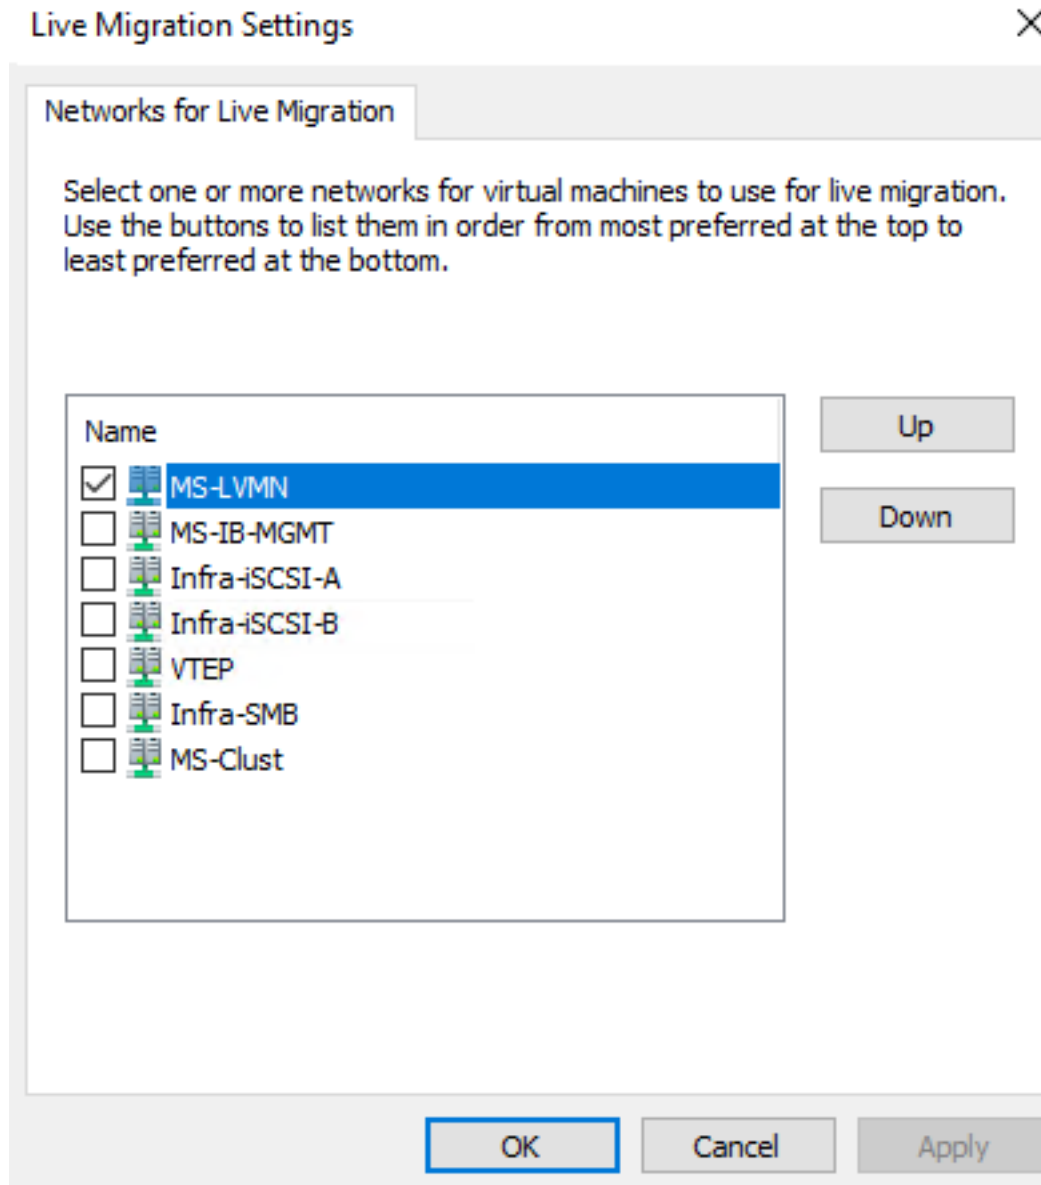
3. Using the subnet information reset the Name of the Cluster Network to the appropriate name, adjust the communication setting and click OK. Storage Networks, Live Migration, and the VTEP network should not allow cluster communication. It may be necessary to edit the network two times to get the name and Cluster Use setting correctly input.
4. Repeat step 3 to assign a descriptive name to all Cluster Networks.



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. Complete the following steps:

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



4. Click Apply and OK to save this setting.

Cluster Storage Settings

1. On the left, expand Storage and select Disks.
2. Right-click the Witness disk and select Properties.
3. Change the Disk Name to Witness and click OK.
4. Looking at the volume names for the Cluster Shared Volumes, rename them to match to storage volume names. Note, the mount pointy for each CSV.

Make SCVMM a Highly Available VM

1. In Virtual Machine Manager, in the VMs and Services workspace, select the first Hyper-V Management host and in the center pane, right-click the SCVMM VM and select Migrate Storage.
2. Browse to C:\ClusterStorage\Volume1. Select the radio button to Automatically place all VHDs with the configuration. Click Next.
3. Click Move. The migration will take several minutes.
4. Right-click the SCVMM VM again and select Properties. Select Hardware Configuration. Ensure that Availability is High.

Build a Windows Server 2016 Virtual Machine for Cloning

A Windows Server 2016 virtual machine can be built in Virtual Machine Manager and cloned to create other virtual machines for management and tenant functions.

1. In Virtual Machine Manager, in the Library workspace, at the top select Import Physical Resource. Click **“Add resource” and browse to the location of the Windows Server 2016 ISO.** Select the ISO and click Open. Back in the Import Library Resources window, click Browse, select the MSSCVMMLibrary, and click OK. Click Import. When the resource is successfully imported, close the Jobs window.
2. On the SCVMM VM use Windows Explorer to navigate to C:\ProgramData. Add the Everyone user to the **Sharing on the Virtual Machine Manager Library Files” folder with Read-only permissions.**
3. In Virtual Machine Manager, in the VMs and Services workspace, select the HV-MGMT-Cluster. Right-click the cluster and select Create Virtual Machine.
4. **Under Select Source, select “Create the new virtual machine with a blank virtual hard disk” and click Next.**
5. Under Identity, name the virtual machine **Win2016-DC-GUI”, select Generation 2,** and click Next.
6. Under Configure Hardware, select the Hyper-V Cloud Capability Profile, 2 Processors, 4096 MB Virtual machine memory, a new Dynamic 120 GB Virtual Hard Disk that contains the operating system, the **Windows Server 2016 Installation ISO connected to the Virtual DVD drive with “Share file instead of copying it” selected, Network Adapter 1 connected to EPG-IB-MGMT in the FP Foundation ACI tenant, and Availability set to “Make the virtual machine highly available. Click Next.**
7. **Under Select Destination, select “Place the virtual machine on a host” and make sure the “Hyper-V Management” Destination is selected. Click Next.**
8. Under Select Host, allow the placement process to run, accept the recommendation, and click Next.
9. Under Configure Settings, Virtual Machine Location, browse to C:\ClusterStorage\Volume1. Also under Machine Resources, browse to C:\ClusterStorage\Volume1 for the Destination path of the Virtual Hard Disk. Click Next.
10. Under Add Properties, select the Windows Server 2016 Datacenter Operating system and click Next.

11. Under Summary, click Create.
12. When the virtual machine is successfully created, close the Jobs window.
13. In Virtual Machine Manager, in the VMs and Services workspace, right-click the newly created Win2016-DC-GUI VM and select Power On. Then right-click the VM and select Connect or View > Connect via Console. Click the icon to send **Ctrl-Alt-Del to the VM and press Enter when you see “Press any key to boot from CD or DVD”**.
14. Install Windows Server 2016 Datacenter with Desktop Experience on the VM, assign it an IP address and hostname, do not join the VM to the Windows Domain, and install all Windows Updates on the VM. Then shut down the VM.
15. Right-click the Win2016-DC-GUI VM in Virtual Machine Manager and select Properties. Under Hardware Configuration, disconnect the Windows Server 2016 Installation ISO from the Virtual DVD drive and click OK.

NetApp SMI-S Provider Configuration

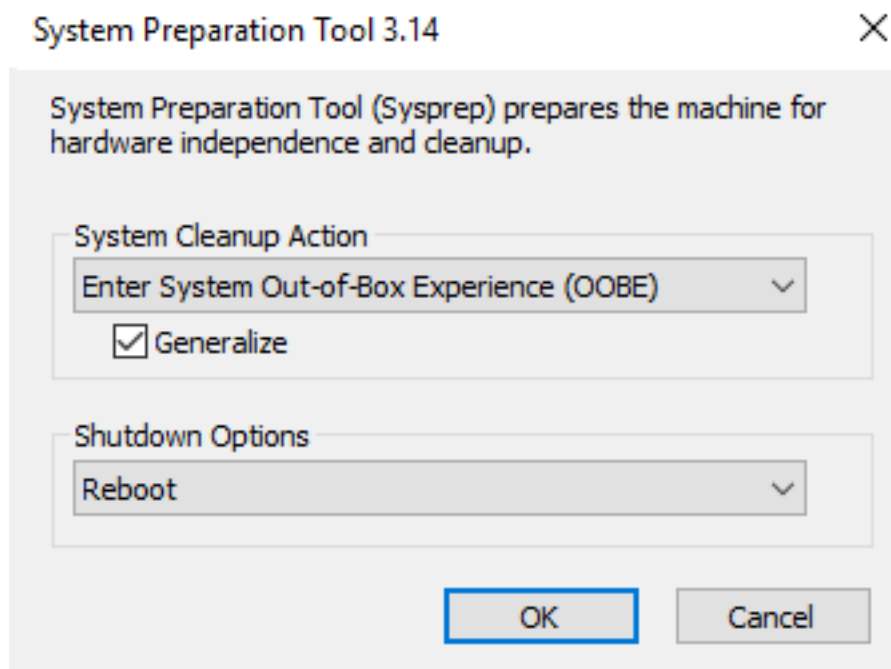


The NetApp SMI-S Provider can be downloaded from <http://mysupport.netapp.com>.

Install the NetApp SMI-S Provider to iSCSI Storage

1. In Virtual Machine Manager, right-click the Win2016-DC-GUI VM and select Create > Clone.
2. Under Identity, name the VM SMI-S-Provider and click Next.
3. Under Configure Hardware, if you will have a large Hyper-V environment, change the Memory to 8192 MB. Click Next.
4. **Under Select Destination, select “Place the virtual machine on a host” and select the Hyper-V Management Destination using the pulldown.** Click Next.
5. Under Select Host, allow the placement process to run, accept the recommendation, and click Next.
6. Under Select Path, Browse to and select C:\ClusterStorage\Volume1. Click Next.
7. Under Select Networks, select the MS-Core-Services EPG in Tenant common and click Next.
8. Under Add Properties, click Next.
9. Under Summary, click Create. When the VM has successfully been created, close the Jobs window.
10. In Virtual Machine Manager, right click the newly created SMI-S-Provider VM and select Properties. Verify that the Hardware Configuration is correct.
11. Right-click the SMI-S-Provider VM and select Power On. Then right-click the SMI-S-Provider VM and select Connect or View > Connect via Console. Click the icon to send Ctrl-Alt-Del to the VM and log into the VM as Administrator.

12. Using File Explorer, navigate to C:\Windows\System32\Sysprep. Double-click sysprep.exe.
13. Select options as shown below and click OK.



14. Respond to the prompts and log into the Windows VM as Administrator. Assign the VM an IP address and hostname, join the VM to the Windows Domain, and install all Windows Updates on the VM. This process will require at least one reboot. After reboot log back into the VM as Administrator.
15. Download the NetApp SMI-S Provider version 5.2.4 to the local desktop from <http://mysupport.netapp.com>. Save the file as smisprovider-5-2-4.msi.
16. Navigate to the directory that contains the NetApp SMI-S Provider software package. Double-click the package name.
17. Complete the steps in the setup wizard to complete the install.

Create the Local Administrator

1. Using the search icon, enter run and 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 Start Menu, navigate to NetApp SMI-S Provider.
2. Right click and select Run as Administrator. A command line prompt should open.

- 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.
```

- 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\flexadmin
Please enter your password: *****
Please re-enter your password: *****
User added successfully.
```

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

```
C:\Program Files (x86)\NetApp\smis\pegasus\bin>smis addsecure 10.1.118.10 vsadmin
Enter password: *****
Returned Path  ONTAP_FilerData.hostName="10.1.118.10",port=443
Successfully added 10.1.118.10
```

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.

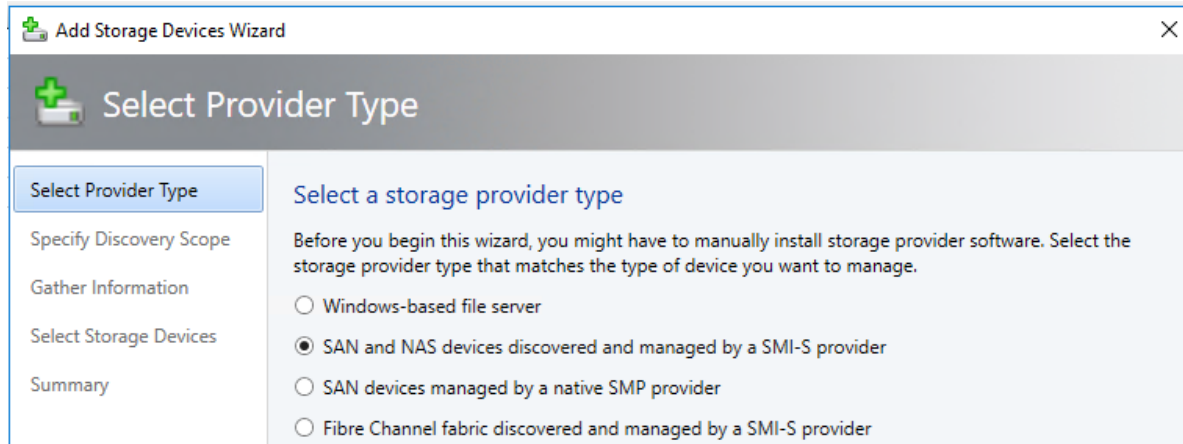
To add an SMI-S storage device, make sure that you have installed the SMI-S provider for the array on a server that the VMM management server can access over the network by IP address or by fully qualified domain name (FQDN).



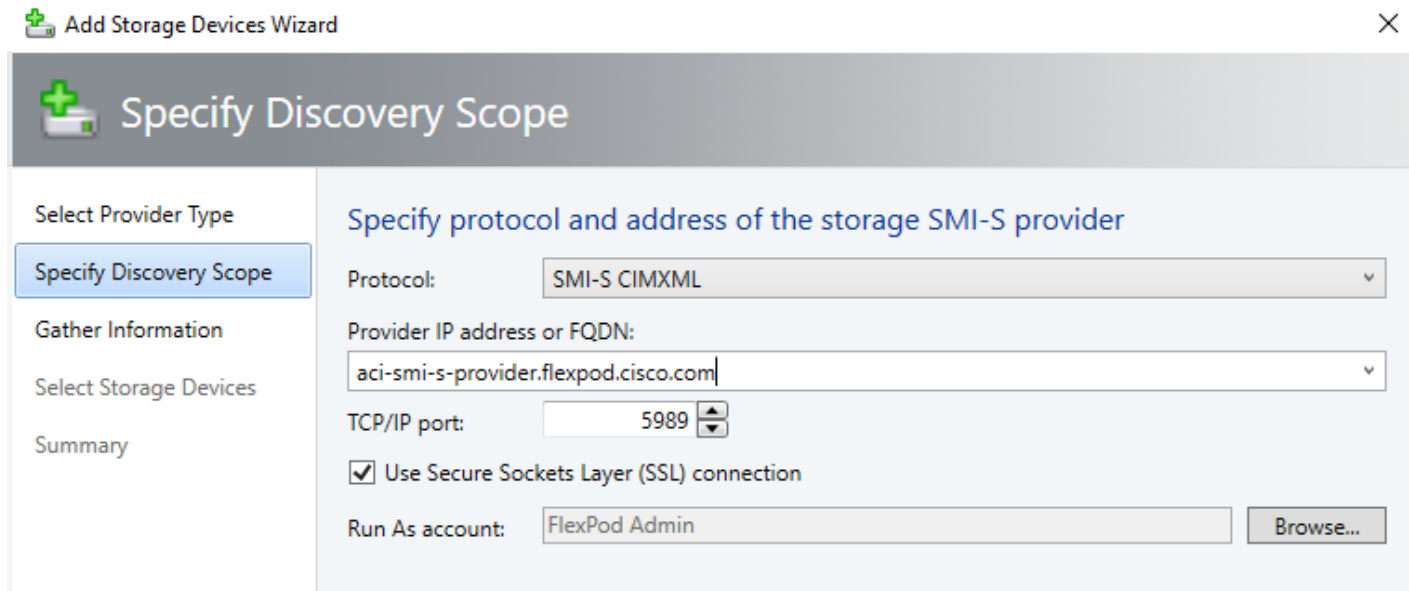
Do not install the SMI-S provider on the VMM management server. This configuration is not supported.

Add a storage device

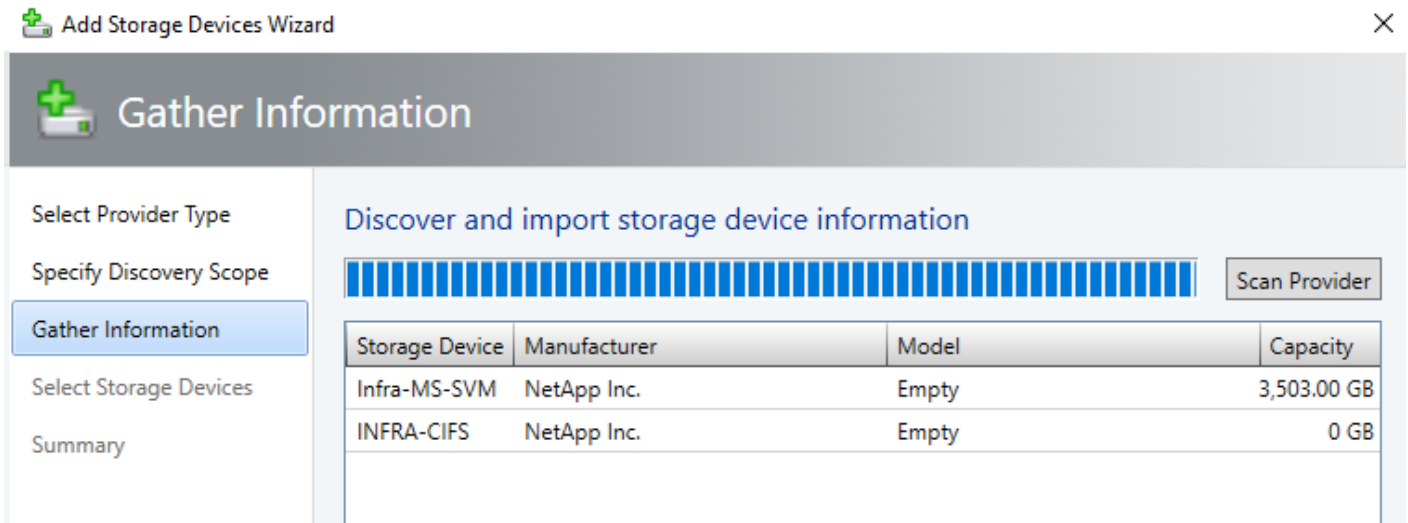
- In Virtual Machine Manager, click Fabric > Storage > Add Resources > Storage Devices.
- In Add Storage Devices Wizard > Select Provider Type, select to add a storage device with SMI-S. Click Next.



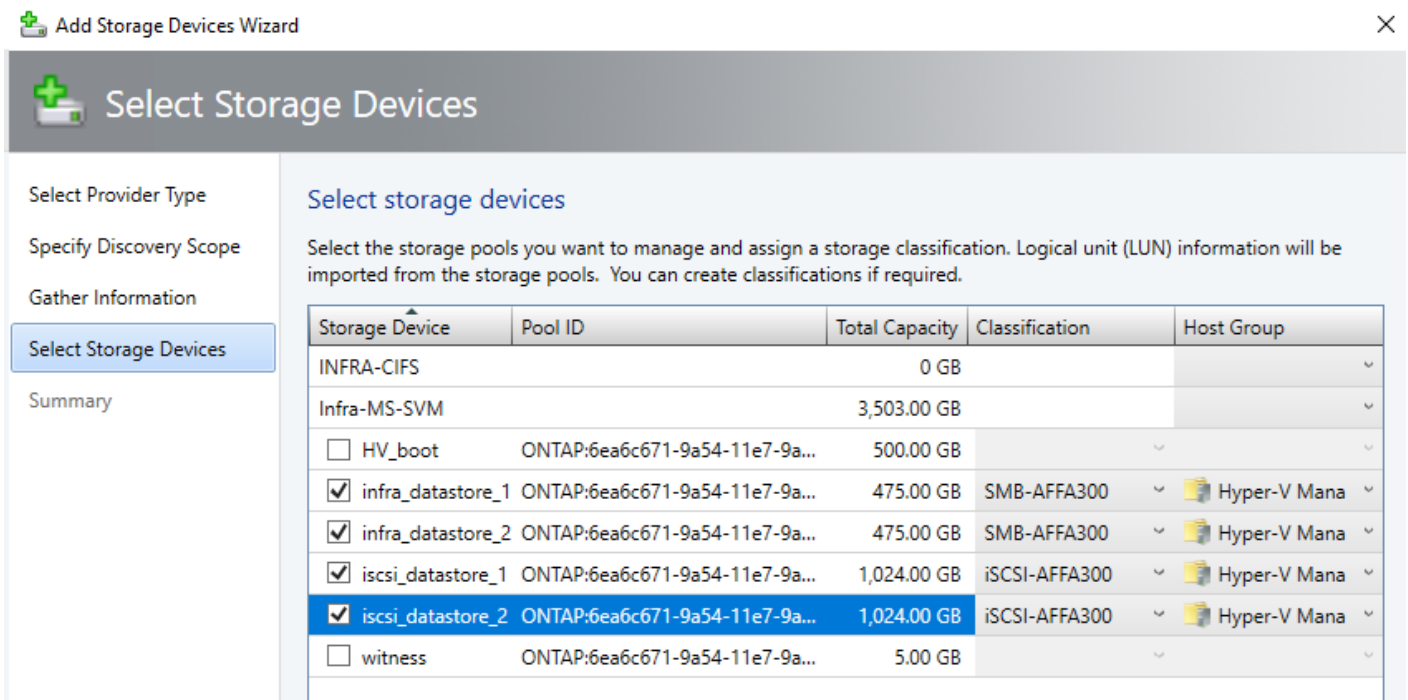
- In Specify Discovery Scope, select Protocol - SMI-S CIMXML, add the IP address/FQDN of the SMI-S Provider, and add the port used to connect to the provider on the remote server. You can enable SSL if you're using CIMXML. Then specify an account for connecting to the provider. You will need to create a Run As account for the flexadmin account added to the CIM server above. Click Next.



- In Gather Information, VMM automatically tries to discover and import the storage device information. You will need to import the security certificate.
- 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. Only select storage where VMs will be stored. Click Next.



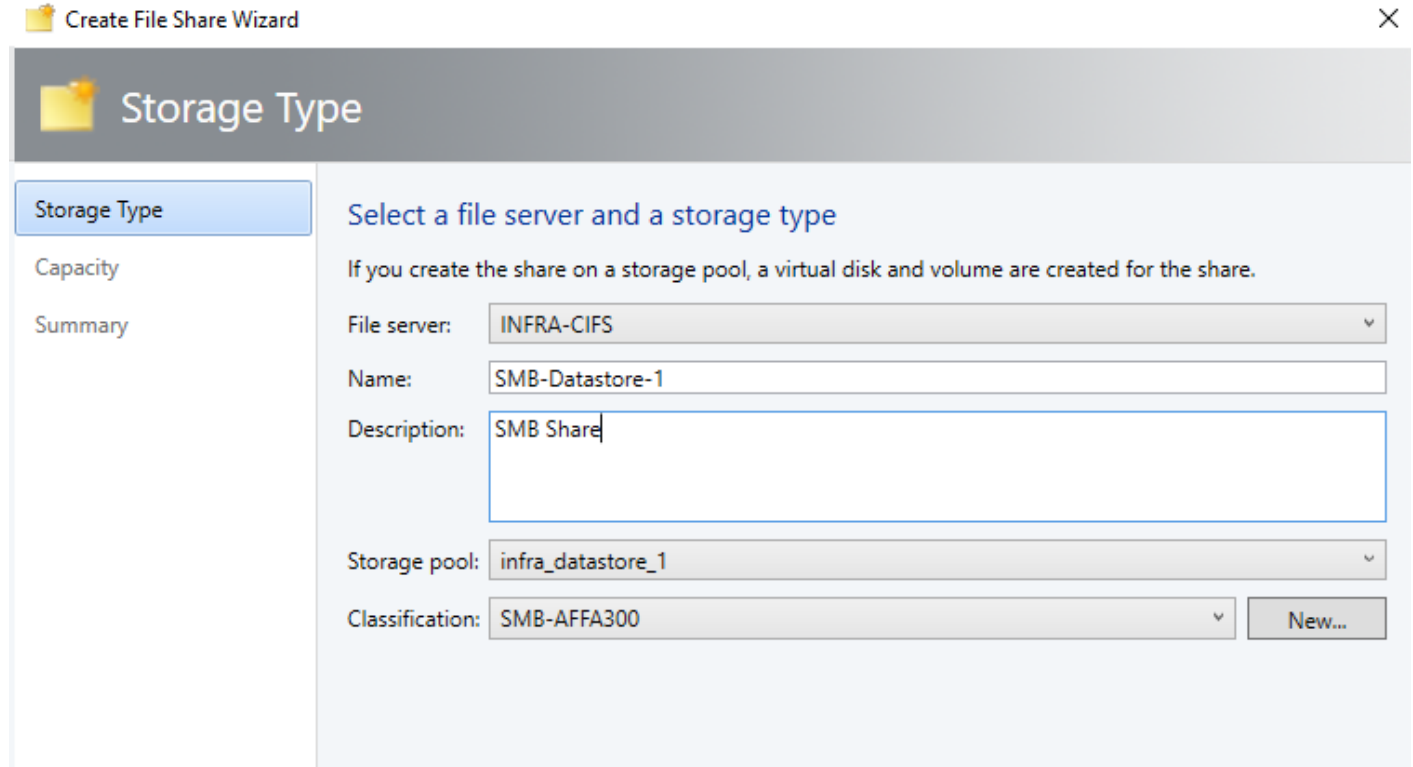
- On the Summary page, confirm the settings, and then click Finish. The Jobs dialog box appears. When status is Completed you can verify the storage in Fabric > Storage > Classifications and Pools.

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

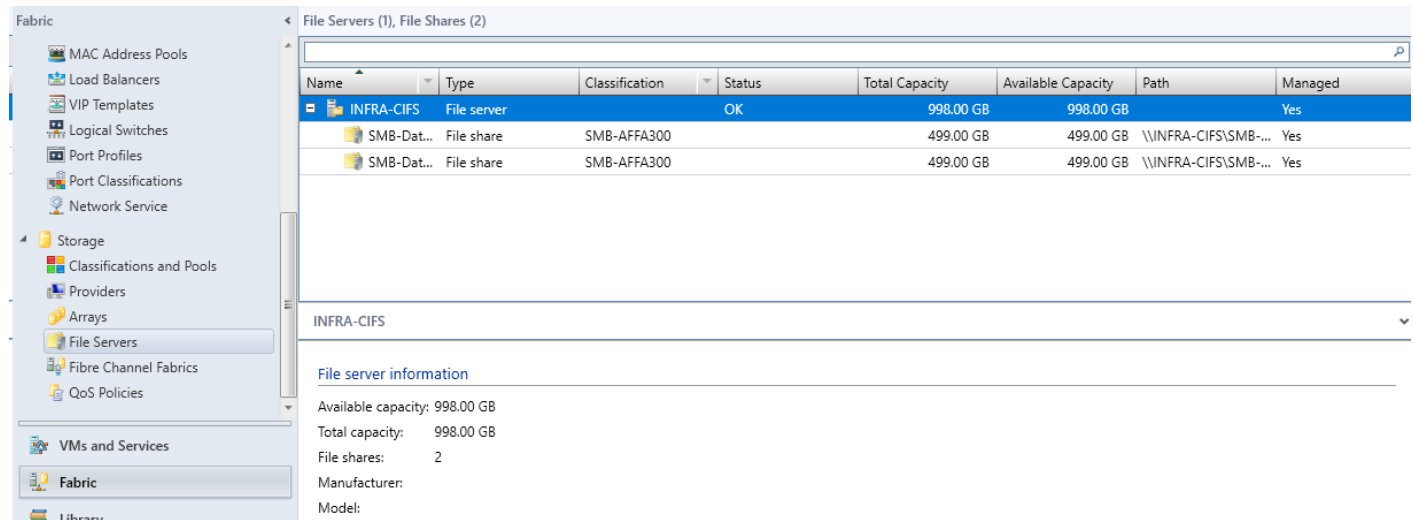
SMB file shares can be used by Hyper-V hosts as a shared storage to store virtual machine files. This section covers steps to create and assign SMB file shares to stand-alone Hyper-V servers and host cluster.

- To Add a storage device, refer to the steps covered in the above section.

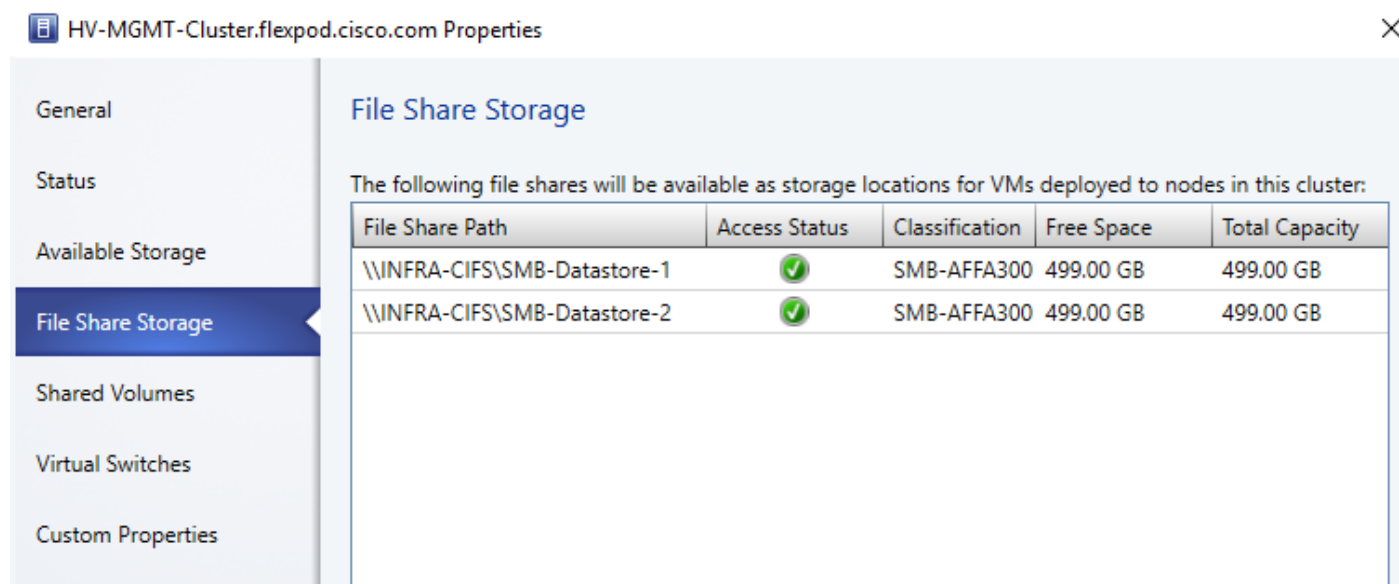
2. To create a file share, open Fabric workspace, expand Storage and click File Servers.
3. Select the File Server and click Create File Share and in the Create File Share wizard, enter a name for the share and select Storage Type, Pool and Classification. Click Next.



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 File Servers.



7. Repeat this process to add a file share for the infra_datastore_2 Storage pool.
8. Assign the file shares to the host cluster by navigating to Fabric > Servers > All Hosts > Hyper-V Management > HV-MGMT-Cluster.
9. Locate and right-click the cluster icon and click on Properties.
10. Click on File Share Storage and click Add.
11. From the drop-down list next to the File Share Path, select a share and click OK.
12. Repeat this step to select the other share. Click OK.



13. VMs can now be Storage Migrated into the SMB Datastores if desired.

Build Windows Active Directory Servers for ACI Fabric Core Services

Two Windows Server 2016 virtual machines will be cloned from the Win2016-DC-GUI VM and provisioned as Active Directory (AD) Domain Controllers in the existing AD Domain.

1. Create two high-availability clones of the Win2016-DC-GUI VM connected to the Core-Services EPG in ACI tenant common. Place one of these VMs in SMB-Datastore-1 on Host 1 and the other in SMB-Datastore-2 on Host 2.
2. Boot each clone and sysprep it. Then assign the VM an IP address and hostname. Do not join the VM to the AD domain.
3. Install Active Directory Domain Services on each VM and make it a Domain Controller and DNS server in the AD domain. Ensure the DNS server Forwarders are set correctly.
4. Add a persistent route to each VM to route to the tenant IP address space (172.18.0.0/16 in this validation) through the Core-Services EPG gateway address (10.1.118.254):

```
route ADD -p 172.18.0.0 MASK 255.255.0.0 10.1.118.254  
route print
```

5. Reset the DNS in all existing VMs, servers, and hardware components to point to the two just-created DNS servers.
6. The two new AD servers can be placed in a separate site if the original AD server is in a different subnet and will not be reachable from tenant subnets.

Build Microsoft Systems Center Operations Manager (SCOM) Server VM

To install SCOM VM, complete the following steps:

1. Create a high-availability clone of the Win2016-DC-GUI VM, 4 CPUs, 8 GB RAM, and connected to the IB-MGMT EPG in ACI tenant FP-Foundation. Place this VMs in either an iSCSI or SMB datastore.
2. Boot the clone and sysprep it. Then assign the VM an IP address and hostname. Join the VM to the AD domain.
3. Install and configure a single server deployment of SCOM 2016 on the VM according to <https://docs.microsoft.com/en-us/system-center/scom/deploy-overview>.

Cisco UCS Management Pack Suite Installation and Configuration

Cisco UCS Manager Integration with SCOM

About Cisco UCS Management Pack Suite

A 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 Centrals. It provides visibility of health and alerts through familiar and easy-to-use interface.

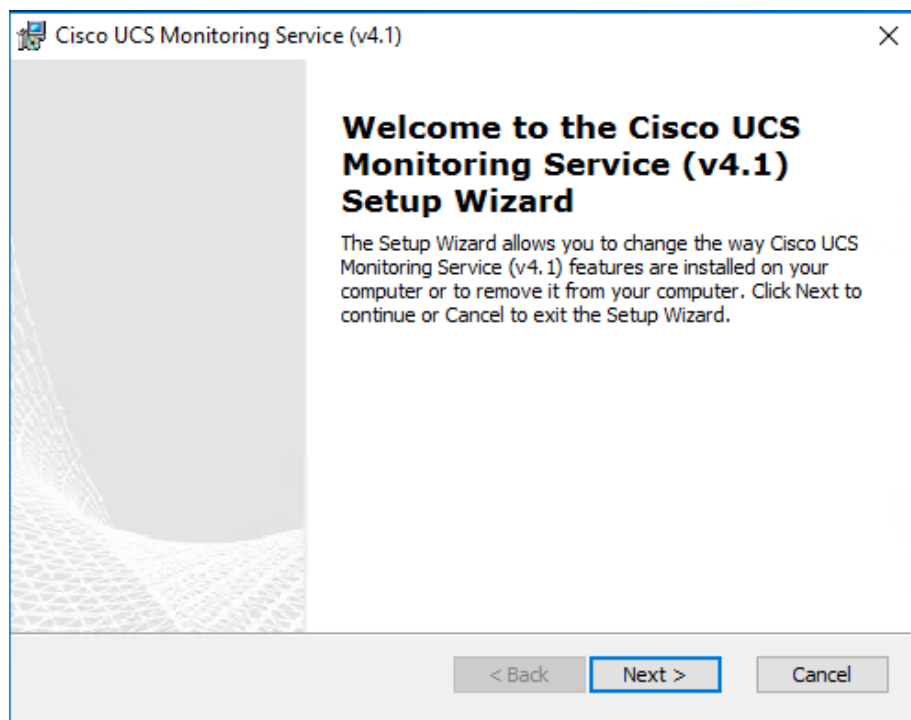
For more information, see:

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

To install the Cisco UCS monitoring service, complete the following steps:

1. On the Cisco.com download site for Cisco UCS Management Partner Ecosystem Software, download the Cisco UCS management pack suite 4.1(1) file and unzip the file into a folder.
2. Navigate to the folder in which the unzipped Cisco UCS Management Pack Suite is stored.
3. Select the monitoring service installer .msi file, and launch the installer.
4. In the Setup wizard, click Next.



5. In the License Agreement page, do the following:
 - a. Review and accept the EULA.
 - b. Click Next.
6. 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.
7. In the Select Installation Folder page, accept the default installation folder or click Browse to navigate to a different folder, and then click Next.
8. On the Ready to Install page, click Install to start the installation.
9. When the Cisco UCS monitoring service is successfully installed, the Installation Complete message appears.
10. 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

To add a firewall exception, complete the following steps:



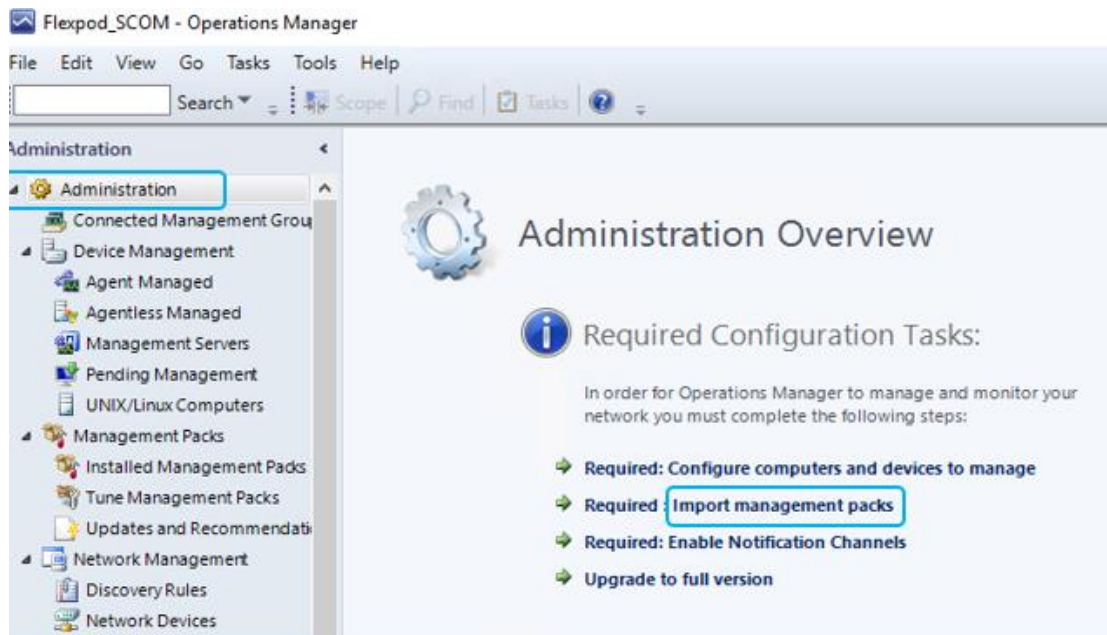
Before you monitor a Cisco UCS domain, enable the following inbound rules in the Windows Firewall with Advanced Security on the computer where you run the Cisco UCS Management Service.

1. File and Printer Sharing:
 - a. Echo-Request-ICMPv4-In
 - b. Echo-Request-ICMPv6-In
2. Remote Service Management (RPC)
3. Remote Service Management (RPC-EPMAP)

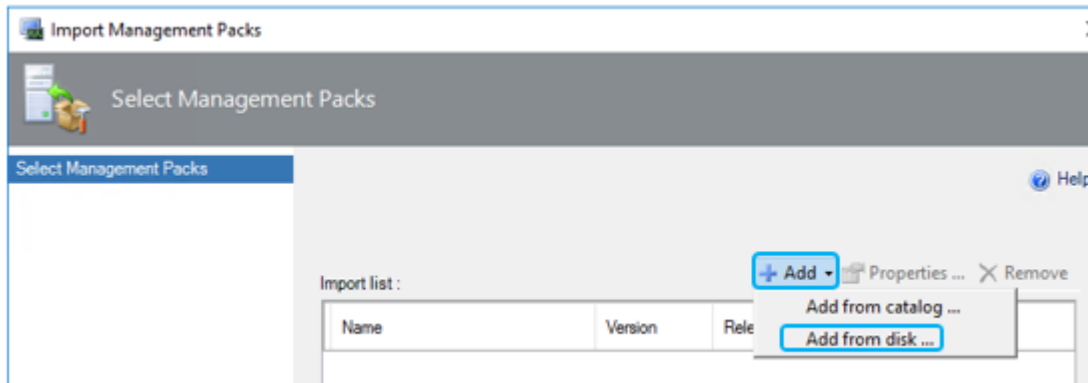
Installing the Cisco UCS Management Pack Suite

To install the Cisco UCS Management Pack Suite, complete the following steps:

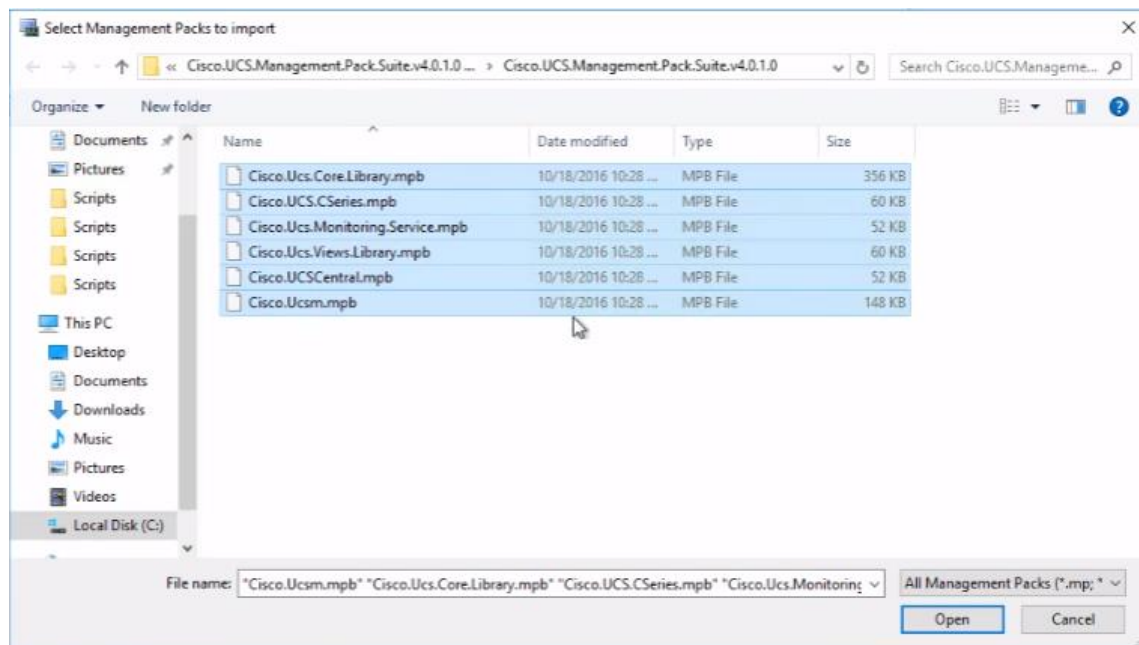
1. For importing Management Packs using Operations Manager console, you must have administrative privileges. For more information on the access privileges, see <https://technet.microsoft.com/en-in/library/hh212691.aspx>. On the Cisco.com download site for Cisco UCS Management Partner Ecosystem Software, 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.



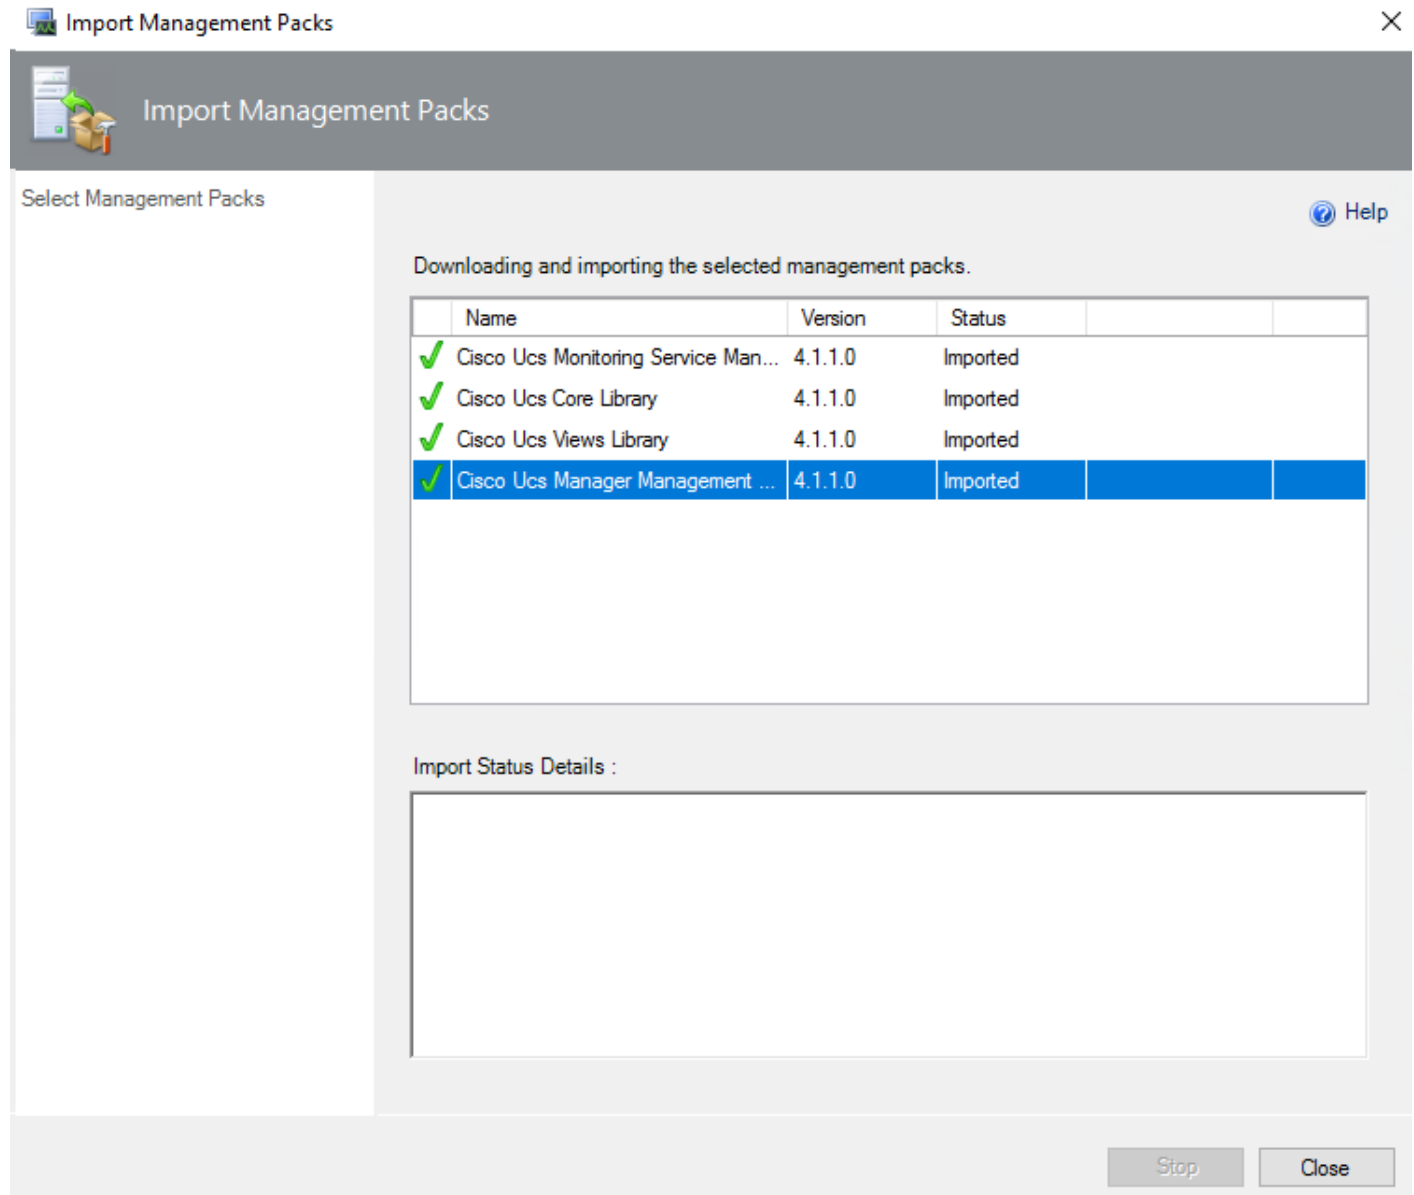
- On the Import Management Pack page, click Add and select Add from disk. An Online Catalog Connection dialog box appears.



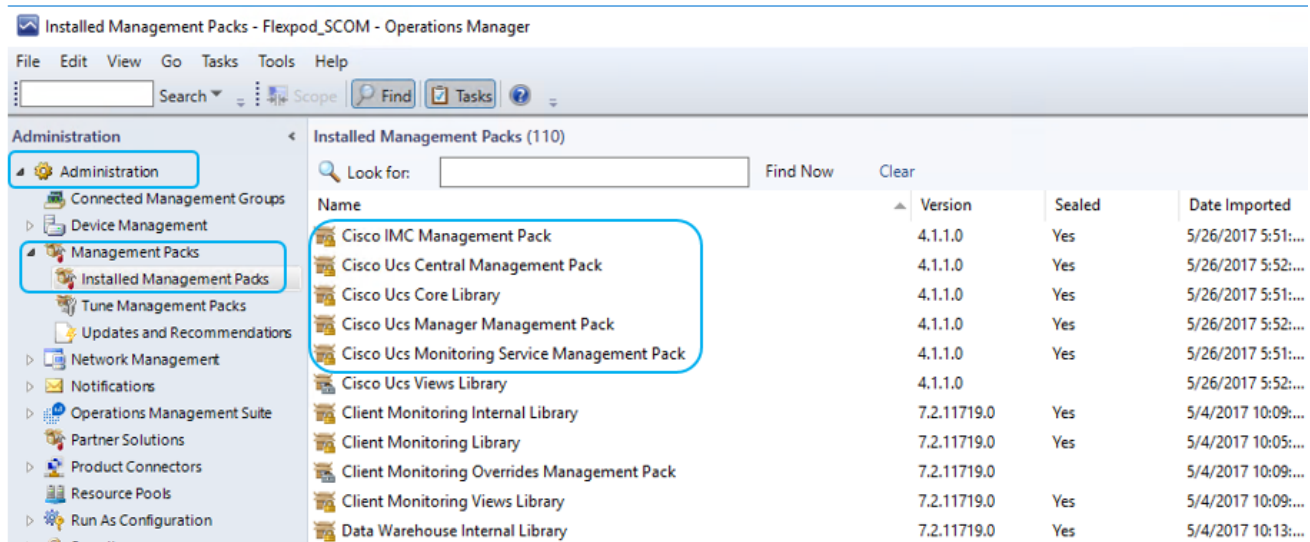
- Click No, if you do not want to search the management pack dependencies online.
- Navigate to the unzipped management pack suite files folder.
- From the list of files, select the mandatory files:
 - Cisco.Ucs.Views.Library.mpb
 - Cisco.Ucs.Core.Library.mpb
 - Cisco.Ucs.Monitoring.Service.mpb
- Other 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.



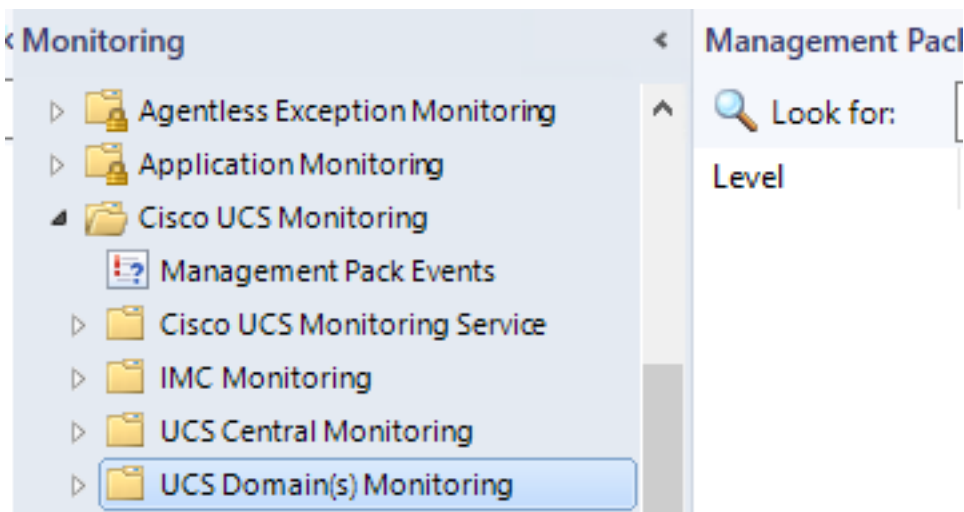
9. Click Open.
10. Click Install on the Import Management Packs page.
11. It may take a few minutes to import the files. When the import is complete, click Close.



12. Verify the installation by navigating to the Administration > Management Packs and click Installed Management Packs.



13. 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 Cisco UCS Domains to the Operations Manager

To add Cisco UCS domains on the servers, where either management pack is imported or the Cisco UCS Management Service is installed, complete the following steps:

1. Launch the Operations Manager console.
2. Navigate to Authoring > Management Pack Templates > 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 screenshot below:

Add Monitoring Wizard ✕

Specify IP Address, Port and Connection Mode Help

Monitoring Type

- General Information
- Instance Name
- Run As Account
- Summary

Cisco UCS Manager

Connection

IP Address* / Hostname:

Connection Mode: Secure Port Number:

Proxy Server

Enable Proxy Configuration

IP Address * / Hostname :

Port:

Enable Proxy Authentication

Username: Password:

* IPv4 Address or IPv6 Address
* IPv6 address should be enclosed in "[" and "]" brackets

Cisco UCS Monitoring Service

Machine Type:


Service Machine:

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 screenshot below and click Next.

Add Monitoring Wizard ✕

Cisco UCSM Instance Name

Monitoring Type
General Information
Instance Name
Run As Account
Summary

 Help

Enter UCS name and description

Name:
a02-6332

Description:

Configuration

Org Discovery Level 3 ▼ Show Unassociated Profiles
 Collect Performance Statistics

Management Pack

Create destination management pack:
a02-6332

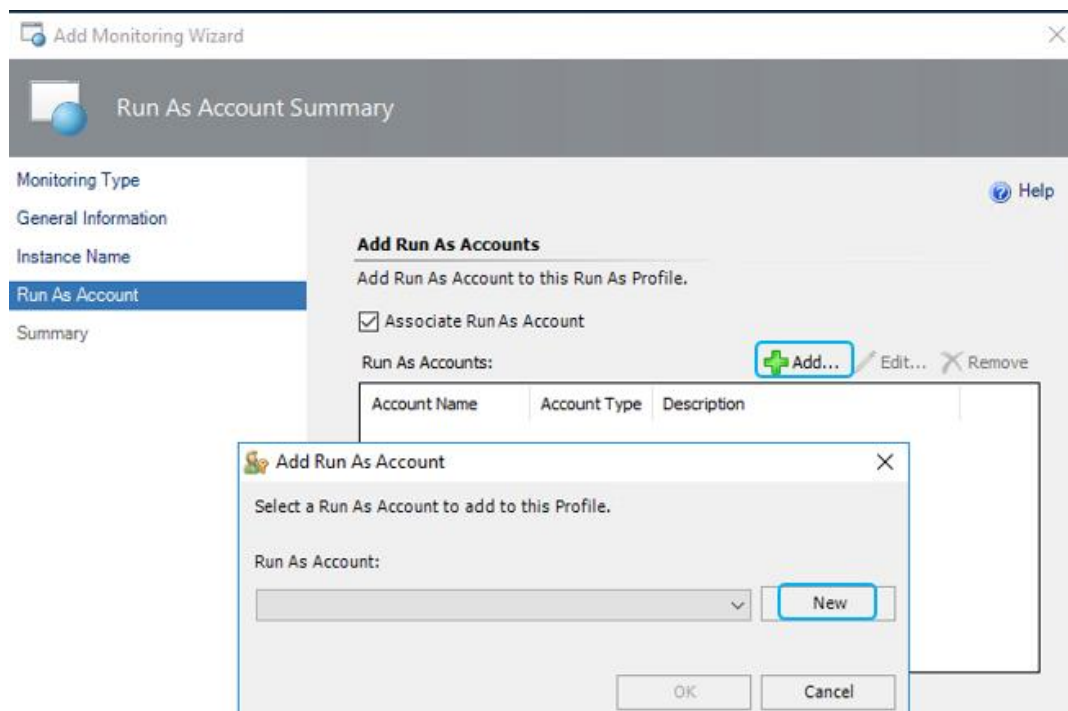
Use existing management pack or create new

<Select Management Pack> ▼ New...

< Previous **Next >** Create Cancel

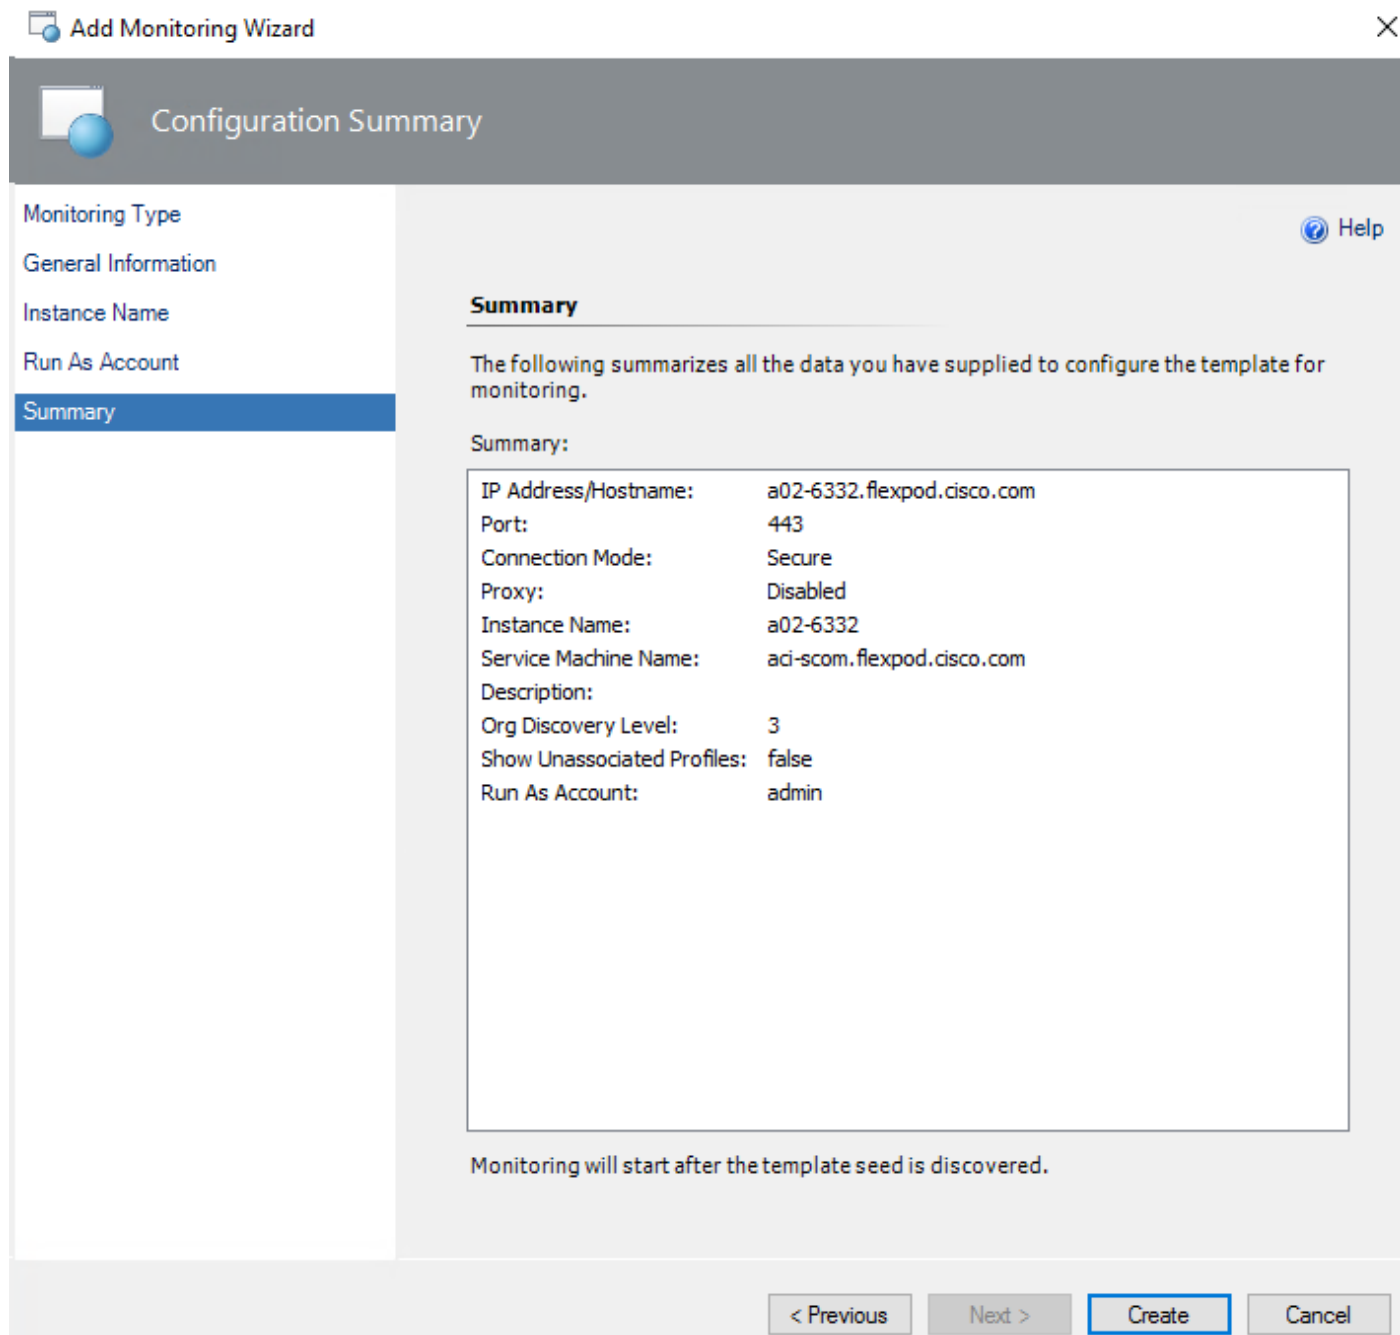
10. On the Run As Account tab, click Add

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

The screenshot displays the 'Ucs Domain Alert Dashboard - ACI-SCOM - Operations Manager' interface. The left sidebar contains a navigation tree with categories like Monitoring, Authoring, Reporting, Administration, and My Workspace. The main area is divided into three sections: Active Alerts (3), Acknowledged Alerts (0), and Cleared Alerts (1). The Active Alerts table shows a critical alert for 'F0831: membership-down' and two warning alerts. The Cleared Alerts table shows one critical alert for 'FabricInterconn...'. The Alert Details section provides information for the selected alert, including its source, full path name, alert rule, and a detailed description of the link-down event.

Icon	Source	Name	Resolution State
Severity: Critical (1)			
🚫	a02-6332	F0831: membership-down	New
Severity: Warning (2)			

Icon	Source	Name
🚫	FabricInterconn...	FabricInterconn...

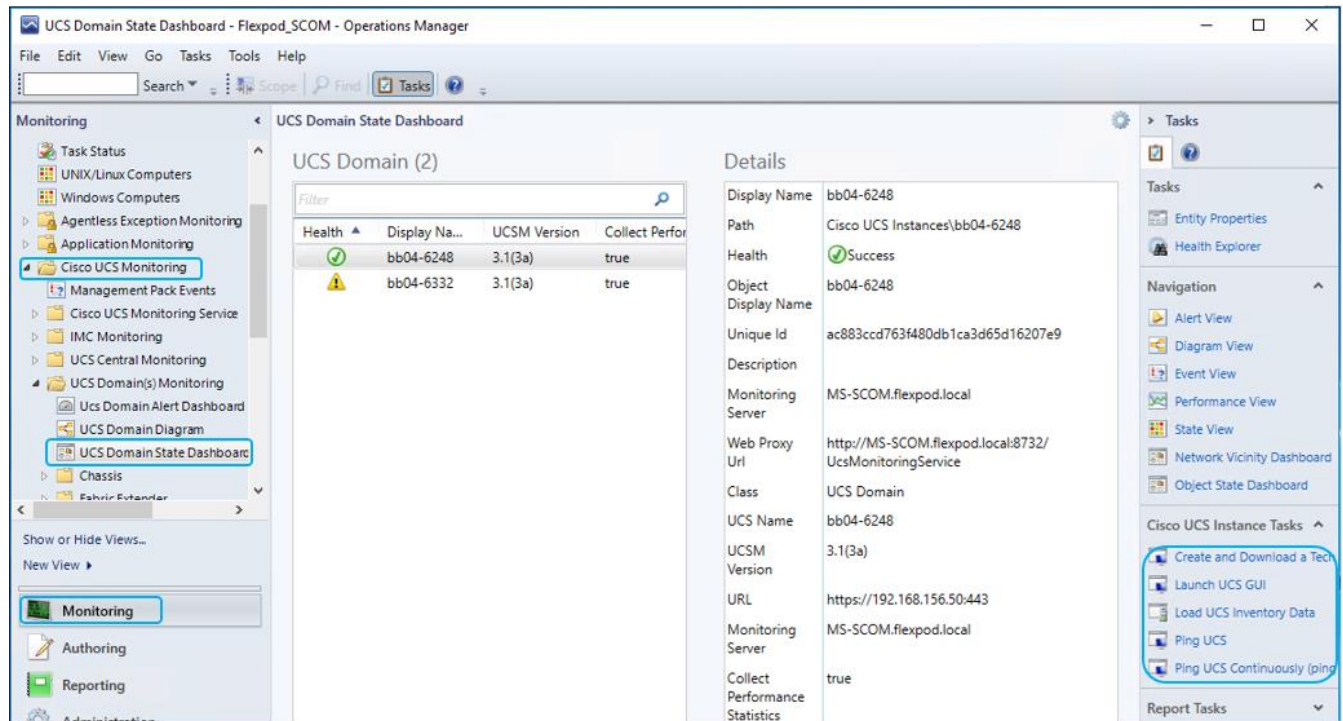
🚫 FabricInterconnect.F0276: link-down	Alert Description
Source: 🚫 FabricInterconnect B	[Instance Name:a02-6332; DN:sys/switch-B/slot-1/switch-ether/port-17/]
Full Path Name: Cisco UCS Instances\A02-6332\FabricInterconnect B	Description: ether port 1/17 on fabric interconnect B oper state: link-down, reason: Link failure or not-connected
Alert Rule: 🔵 Fault Rule : FabricInterconnect.F0276 (link-down)	

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

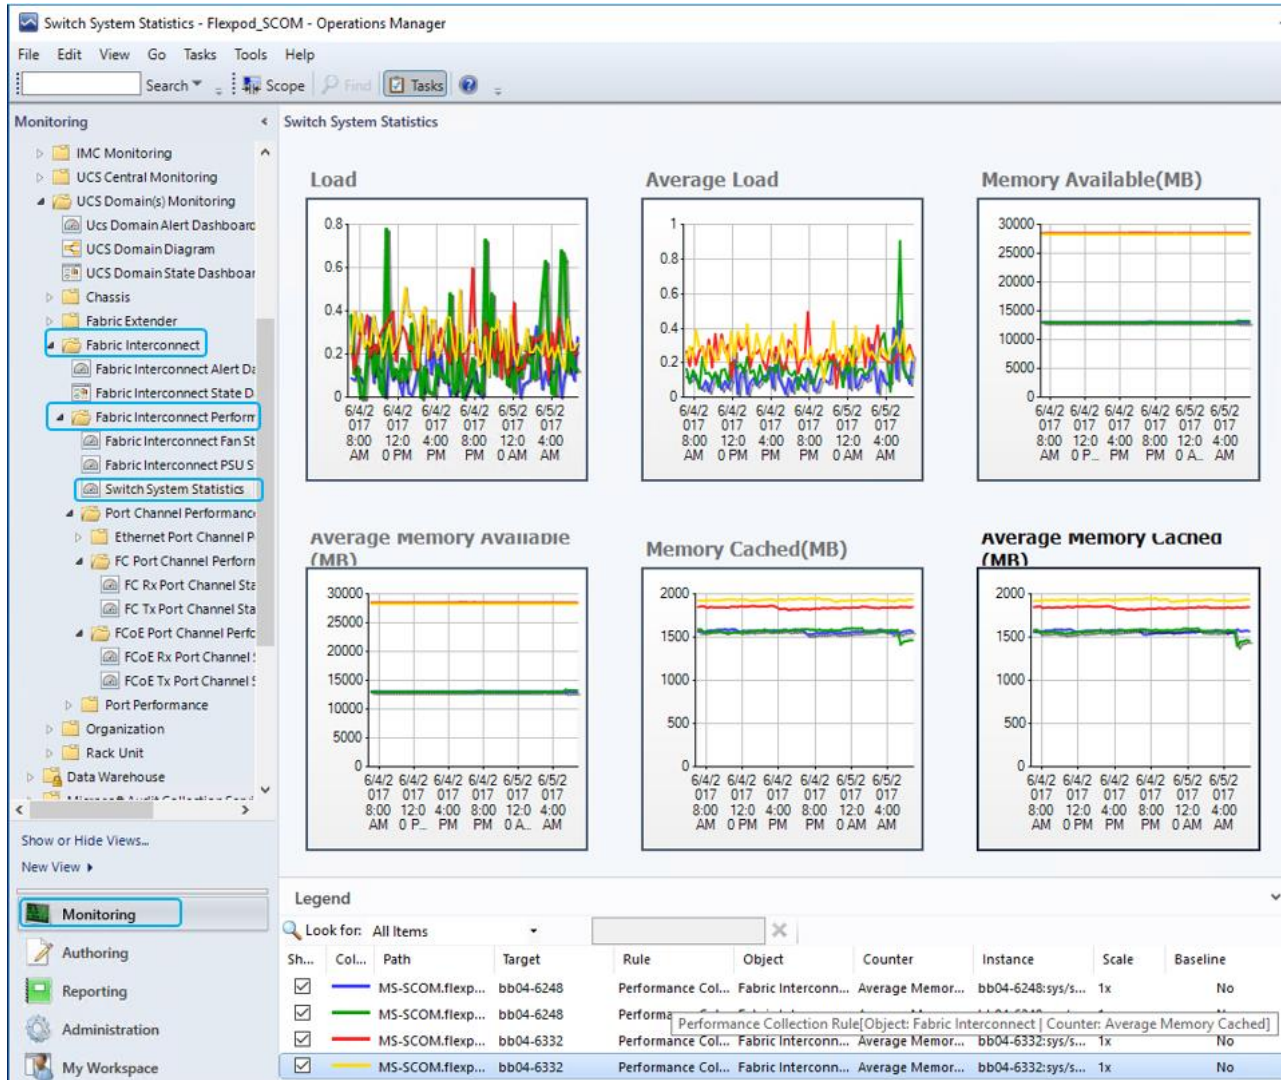
The screenshot displays the Cisco UCS Management Pack Suite interface. On the left is a navigation pane with a 'Monitoring' section expanded to show 'UCS Domain(s) Monitoring' and 'UCS Domain Diagram'. The main area shows a hierarchical tree diagram of the UCS domain. At the top is 'Cisco Ucs m.AllEqu...' with a red error icon. Below it is 'a02-6332' with a red error icon. This node branches into 'a02-6332 HW Invent...' (red error icon) and 'a02-6332 Logical In...' (green checkmark icon). 'a02-6332 HW Invent...' further branches into 'Chassis 1' (red error icon), 'Fabricinter connect A' (green checkmark icon), 'Fabricinter connect B' (green checkmark icon), and 'Rack Unit 1' (green checkmark icon). 'Chassis 1' contains 'Blade 2' (green checkmark), 'Blade 3' (green checkmark), 'Blade 4' (green checkmark), 'Blade 5' (red error icon), and 'Blade 6' (green checkmark). 'Rack Unit 1' contains 'IO Module 1' (green checkmark) and 'IO Module 2' (green checkmark). 'a02-6332 Logical In...' branches into 'Organization root' (green checkmark icon), which contains 'Hyper-V: MGMT-Ho...' (green checkmark icon) and another 'Hyper-V: MGMT-Ho...' (green checkmark icon). Below the diagram is a 'Detail View' section for 'Chassis properties of Chassis 1'.

Chassis properties of Chassis 1	
Display Name	Chassis 1
Unique Id	3f34fae0247547529d5f5b6fa96fb05
Description	
Monitoring Server	aci-scom.flexpod.cisco.com
Web Proxy Uri	http://aci-scom.flexpod.cisco.com:8732/UcsMonitoringService
Class	Chassis
Moniker	sys/chassis-1/
UCS Name/Host Name	a02-6332
Unique Moniker	NA
Model	N20-C6508
Revision	0
Serial Number	FOX1509H5TL
Vendor	Cisco Systems Inc

- 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



- You can view performance metrics for the various Cisco UCS components as shown in 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 Cisco UCS servers (blades or rack-mount servers) on which the host is running.

Cisco UCS Manager Plug-in Installation

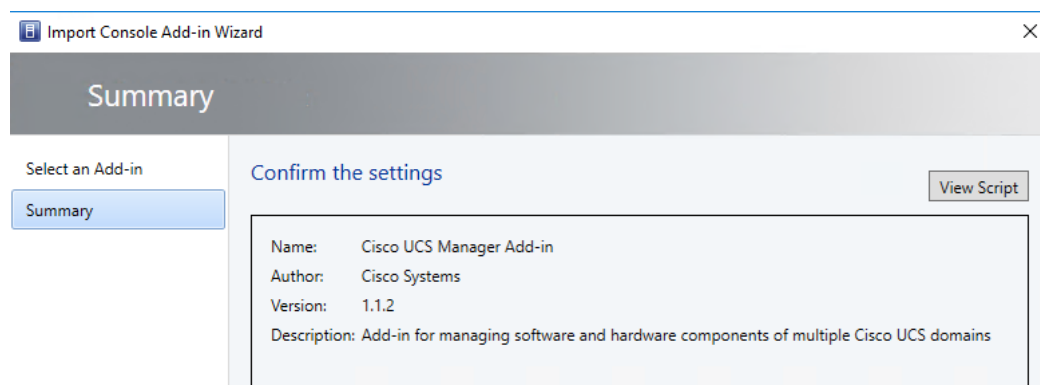
To install the Cisco UCS virtual machine manager add-in, complete the following steps:

1. From the Virtual Machine Manager console, 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.

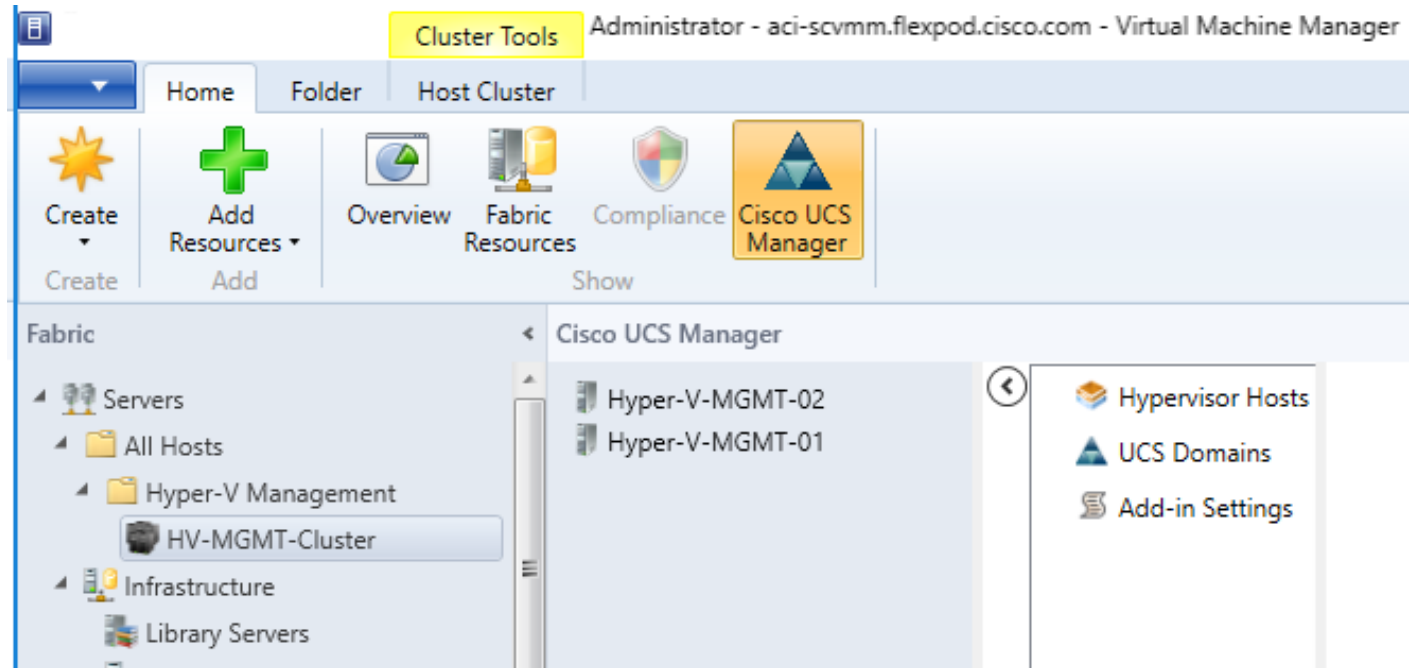


The add-in is made available as a zipped file that has to be imported into the virtual machine manager to install it.

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.

To register a UCS domain, complete the following steps:

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 UCS domain details in the dialog box.



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

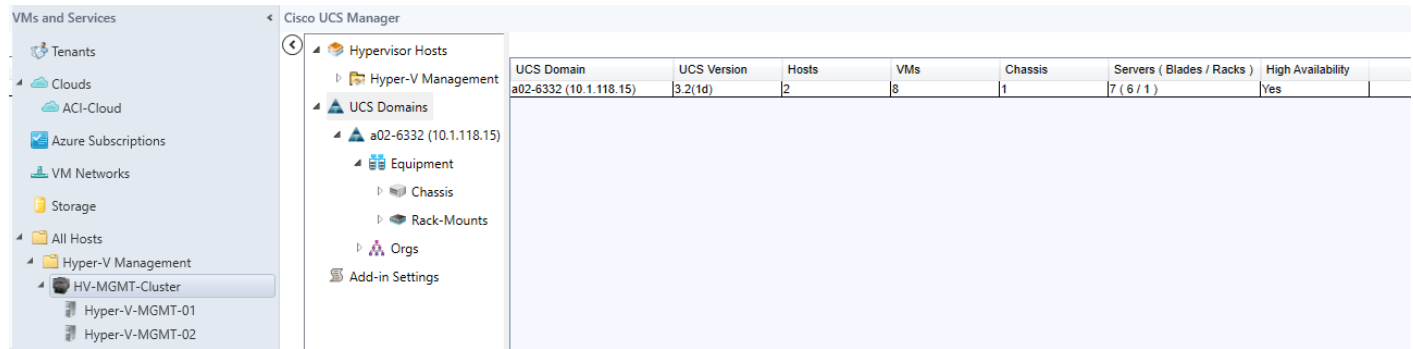
6. If necessary, click Proxy Settings. The Proxy Settings dialog box appears.
7. In the Proxy Settings dialog box, click Use Custom Proxy Settings radio button and enter the details.




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.



 You can also 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

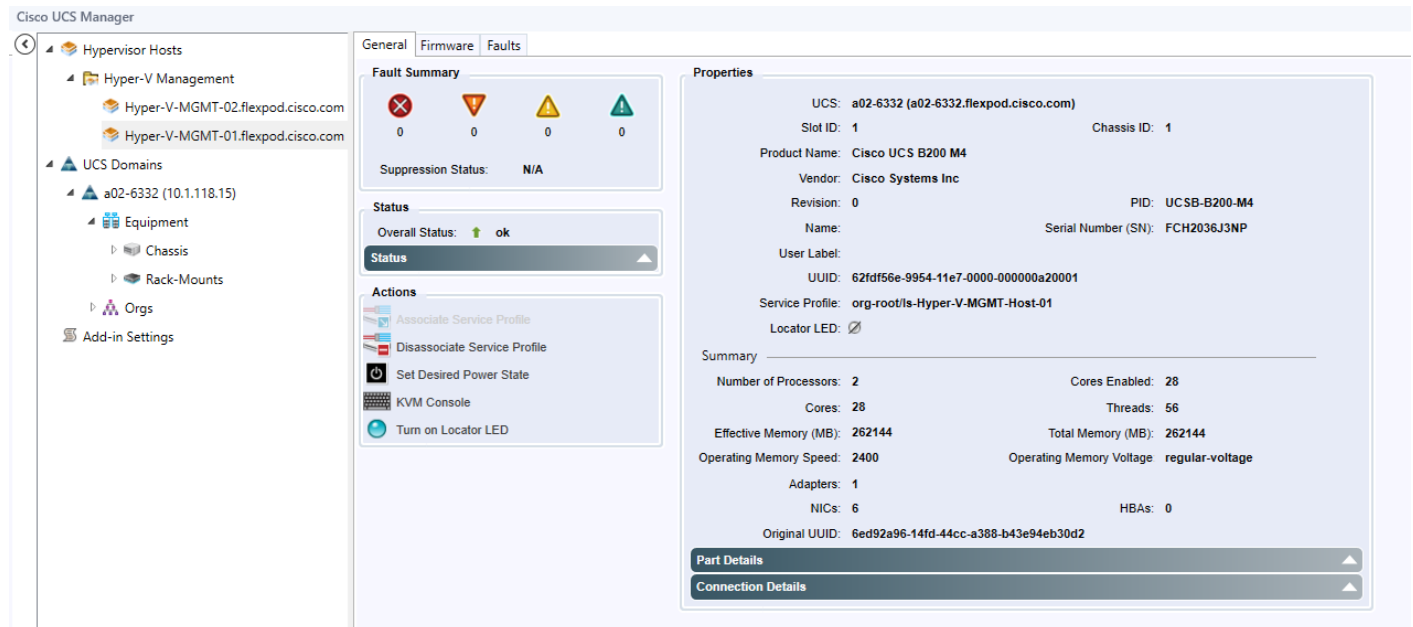
To view server details, complete the following steps:

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.

Name	Description
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.

1. On the right pane of the window, you can view the following information of the server on which the host is running:



Viewing Registered UCS Domains

To view registered UCS domains, complete the following steps:

1. On the toolbar, click Cisco UCS Manager.
2. 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.
3. (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.

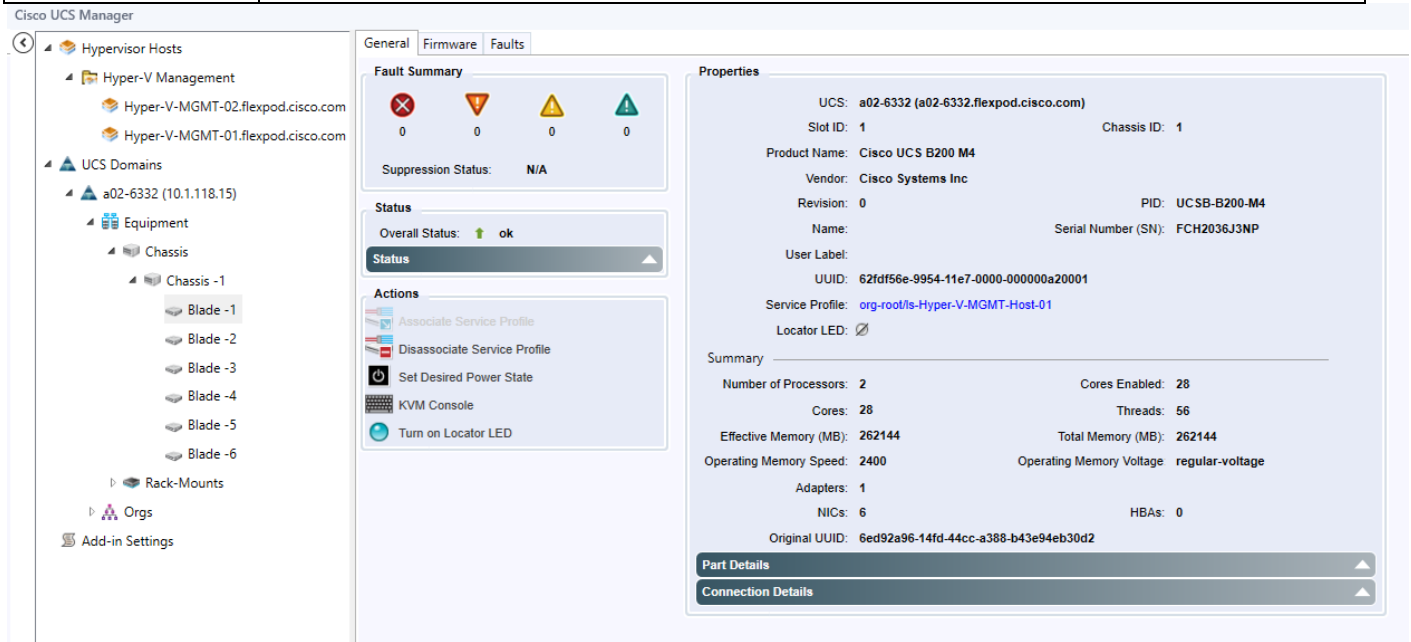
Viewing the UCS Blade Server Details

Using the add-in, you can view the server details, such as properties, faults information, and firmware details. To view server details, complete the following steps:

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. To view server details, complete the following steps:

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. To view the service profile details, complete the following steps:

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.

Name	Description
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.

The screenshot displays the Cisco UCS Manager interface. On the left is a navigation tree with the following structure:

- Hypervisor Hosts
 - Hyper-V Management
 - Hyper-V-MGMT-02.flexpod.cisco.com
 - Hyper-V-MGMT-01.flexpod.cisco.com
 - UCS Domains
 - a02-6332 (10.1.118.15)
 - Equipment
 - Chassis
 - Chassis -1
 - Blade -1
 - Blade -2
 - Blade -3
 - Blade -4
 - Blade -5
 - Blade -6
 - Rack-Mounts
 - Orgs
 - root
 - Service Profiles
 - Hyper-V-MGMT-Host-01
 - Hyper-V-MGMT-Host-02
 - Service Profile Templates
 - Host Firmware Packages
 - Sub-Organizations

The main content area is titled 'General' and contains:

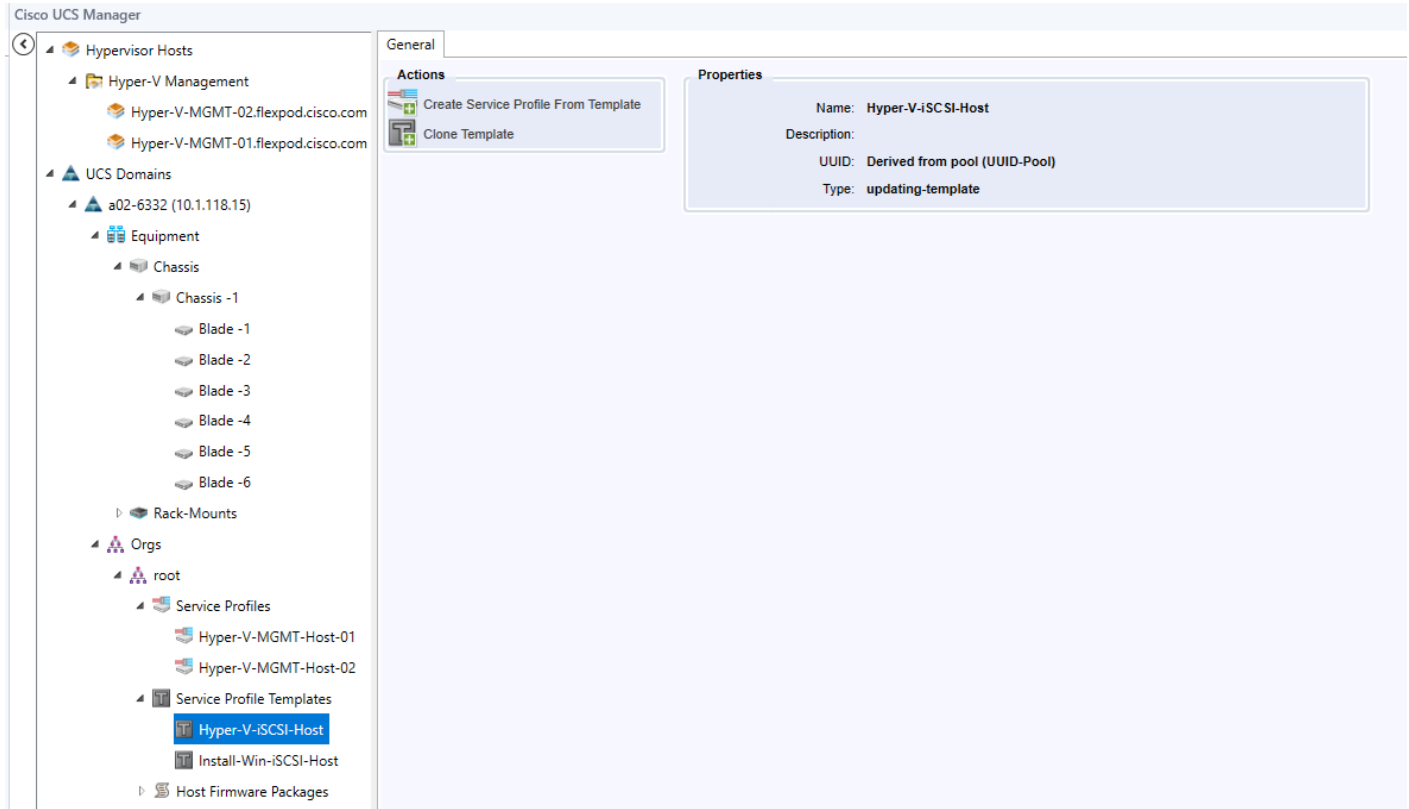
- Fault Summary:** Shows four status icons (red X, orange triangle, yellow triangle, green triangle) with a count of 0 for each. The suppression status is 'N/A'.
- Status:** Overall Status is 'ok'. A 'Status' dropdown menu is visible.
- Actions:** A list of available actions including: Set Desired Power State, KVM Console, Rename Service Profile, Clone Service Profile, Create Service Profile Template, Disassociate Service Profile, Change Host Firmware Package, and Change Service Profile Association.
- Properties:**
 - Name: Hyper-V-MGMT-Host-01
 - User Label:
 - Description:
 - UUID: 62fdf56e-9954-11e7-0000-000000a20001
 - UUID Pool: UUID-Pool
 - UUID Pool Instance: org-root/uuid-pool-UUID-Pool
 - Associated Server: sys/chassis-1/blade-1
 - Service Profile Template: Hyper-V-iSCSI-Host
 - Template Instance: org-root/ls-Hyper-V-iSCSI-Host
 - Assigned Server or Server Pool:** (Dropdown menu)
 - Server Pool: Hyper-V-MGMT-Pool
 - Server Pool Qualification: UCS-Broadwell
 - Restrict Migration: no

Viewing the Service Profile Template Details

Using the add-in you can view the service profile template details, such as properties, and faults information. To view the service profile template details, complete the following steps:

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. To view the host firmware packages details, complete the following steps:

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	Enables you to modify Blade package version and Rack package version

Versions	properties.
----------	-------------

NetApp FlexPod Management Tools Setup

OnCommand Unified Manager 7.2

1. Use the cloning method to deploy a Windows Server 2016 Virtual Machine for the OnCommand installation. Setup the VM with 12 GB RAM, 4 CPUs, a 200 GB hard drive, and connected to the IB-MGMT EPG in the FP-Foundation tenant. Boot the cloned VM, sysprep it, and assign it an IP Address and Hostname. Join the VM to the Windows Domain. It may be necessary to extend the C: system disk to get a 200 GB hard drive. Log into the VM as the FlexPod Administrator.
2. Download and review the [OnCommand Unified Manager 7.2 Installation and Setup Guide for Microsoft Windows](#)
3. Download OnCommand Unified Manager version 7.2P1 .exe file for Microsoft Windows from <http://mysupport.netapp.com> to the virtual machine.
4. Follow the instructions to install OnCommand after running the executable as Administrator.
5. After installation, log into OnCommand Unified Manager and fill in Setup Email and enable Autosupport.
6. To add your cluster for monitoring, click Add from the main dashboard and fill in your cluster details as seen below.

Add Cluster

Host Name or IP Address

a02-affa300

User Name

admin

Password

••••••••|



Protocol

HTTP

HTTPS

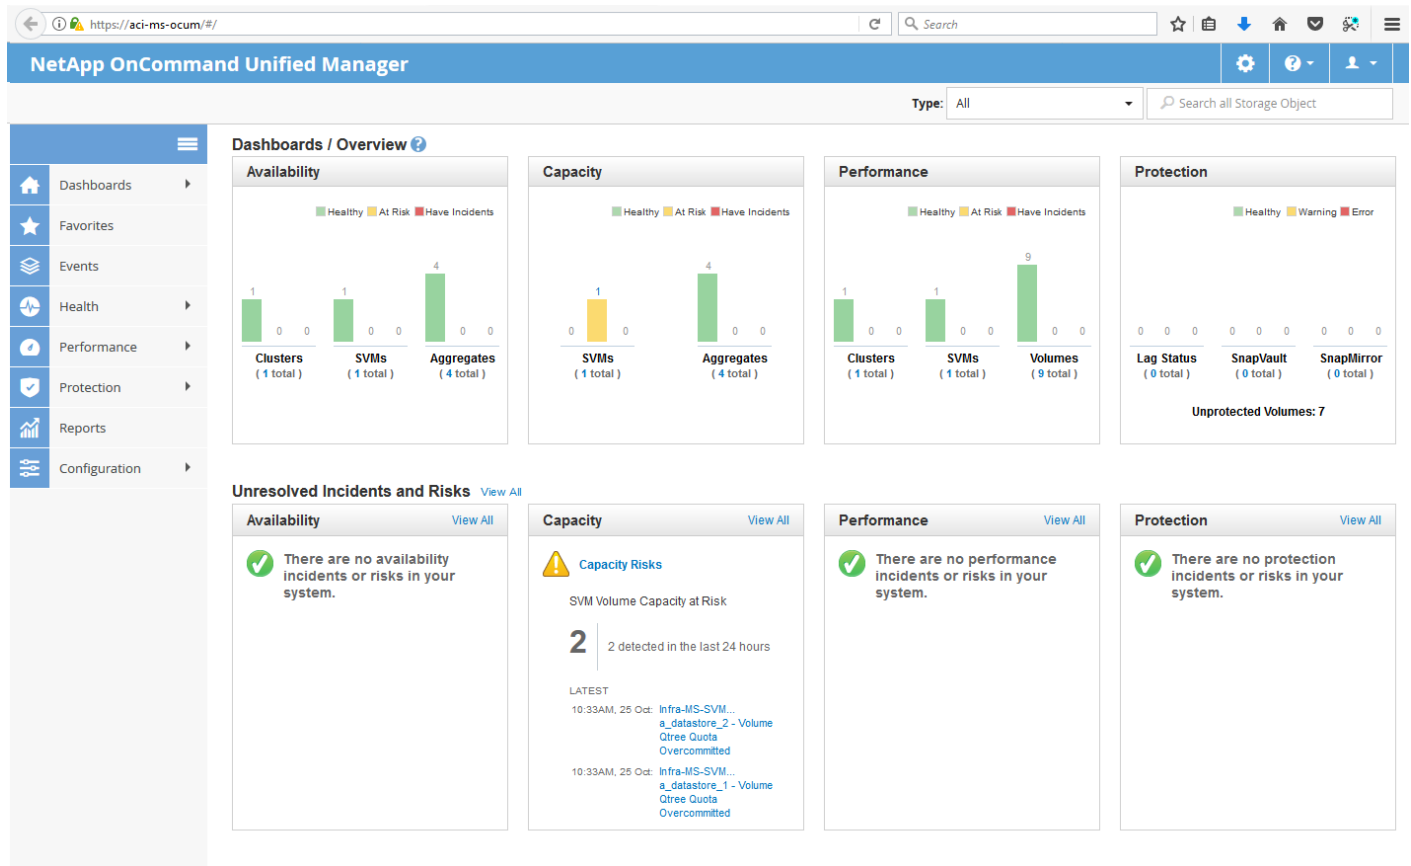
Port

443

Cancel

Submit

7. Add an Inbound Rule to Windows Firewall with Advanced Security to allow TCP port 443 (https) to allow external connectivity to the OnCommand Unified Manager web interface.
8. Storage information can then be viewed by logging into OnCommand Unified Manager at <https://OncommandUnifiedManagerHostname> as show below.



NetApp SnapManager for Hyper-V

NetApp 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 steps in the sections that follow to install SnapManager 2.1.2 for Hyper-V. NetApp SnapManager for Hyper-V will need to be installed on every Hyper-V host that has virtual machines that should be backed up.

Downloading SnapManager for Hyper-V

To download SnapManager for Hyper-V, complete the following steps:



Before you install SnapManager for Hyper-V, download the software package from the [NetApp Support site](#) (requires login credentials).



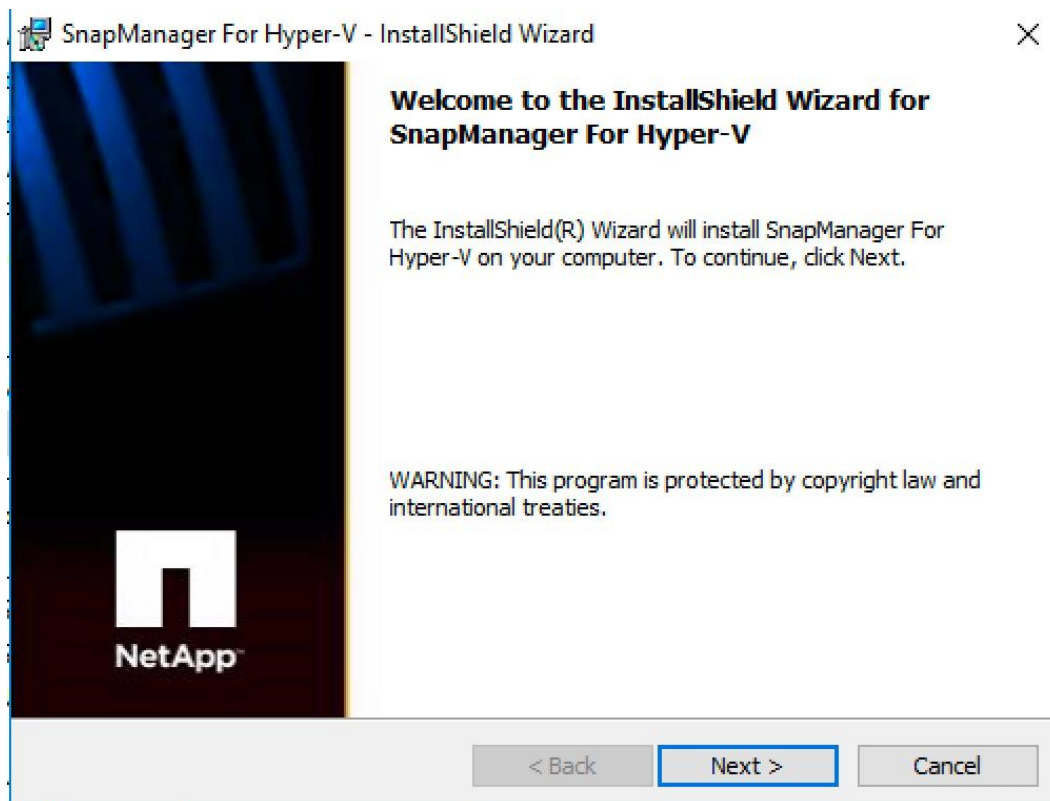
Before you begin you must have login credentials for the NetApp Support Site.

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 terms of 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.

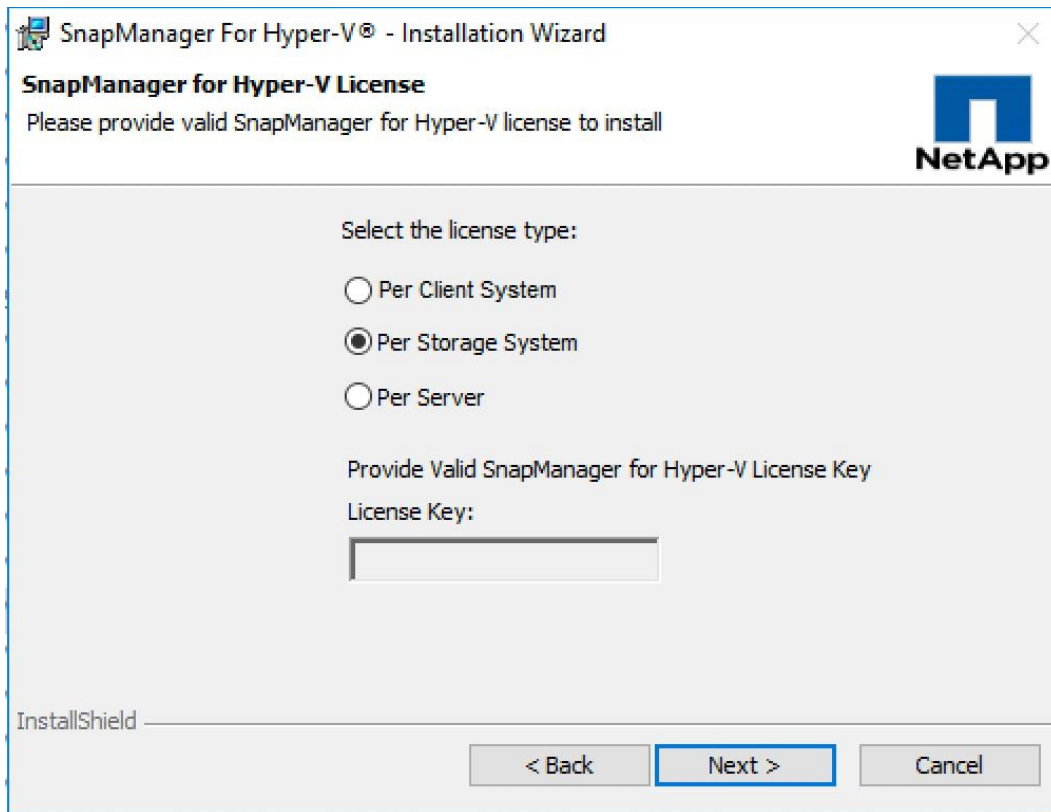
Installing SnapManager for Hyper-V

To install SnapManager for Hyper-V, complete the following steps:

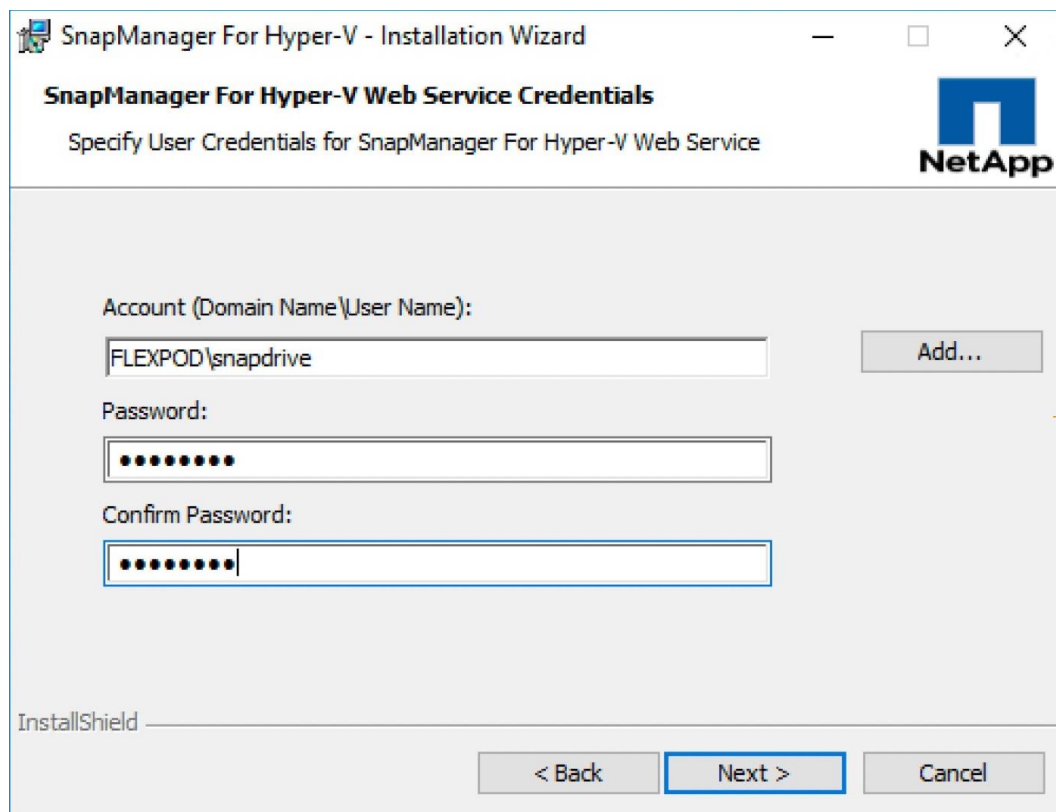
1. Launch the SnapManager for Hyper-V executable file and then follow the Installation wizard instructions.



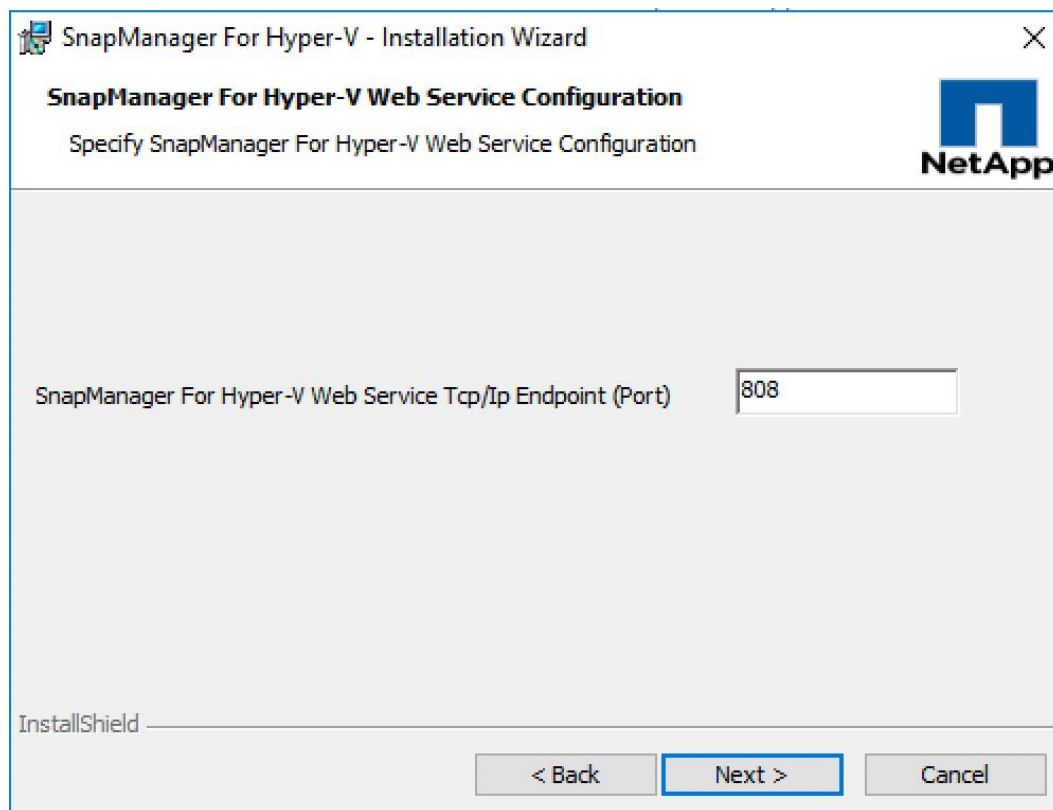
2. On the SnapManager for Hyper-V License page, select the appropriate license type and click Next.



3. Select the installation location and click Next.
4. On the SnapManager for Hyper-V Credentials page, enter the account and password information of an account that is a member of the local administrators. Click Next.



5. On the SnapManager for Hyper-V Web Service Configuration page, accept the default port number and click Next.



6. On the Ready to Install page, click Install.
7. Review the summary of your selections and click Finish.

NetApp OnCommand Plug-in for Microsoft

The NetApp OnCommand Plug-in for Microsoft is an enterprise-class storage monitoring and management application that integrates with SCOM. The plug-in enables administrators to monitor, manage, and report on their NetApp storage systems from within SCOM.

Downloading OnCommand Plug-in for Microsoft

To download OnCommand plug-in for Microsoft, complete the following steps:



Before you install OnCommand Plug-in for Microsoft, download the software package from the [NetApp Support site](#) (requires login credentials).



Before you begin you must have login credentials for the NetApp Support Site.

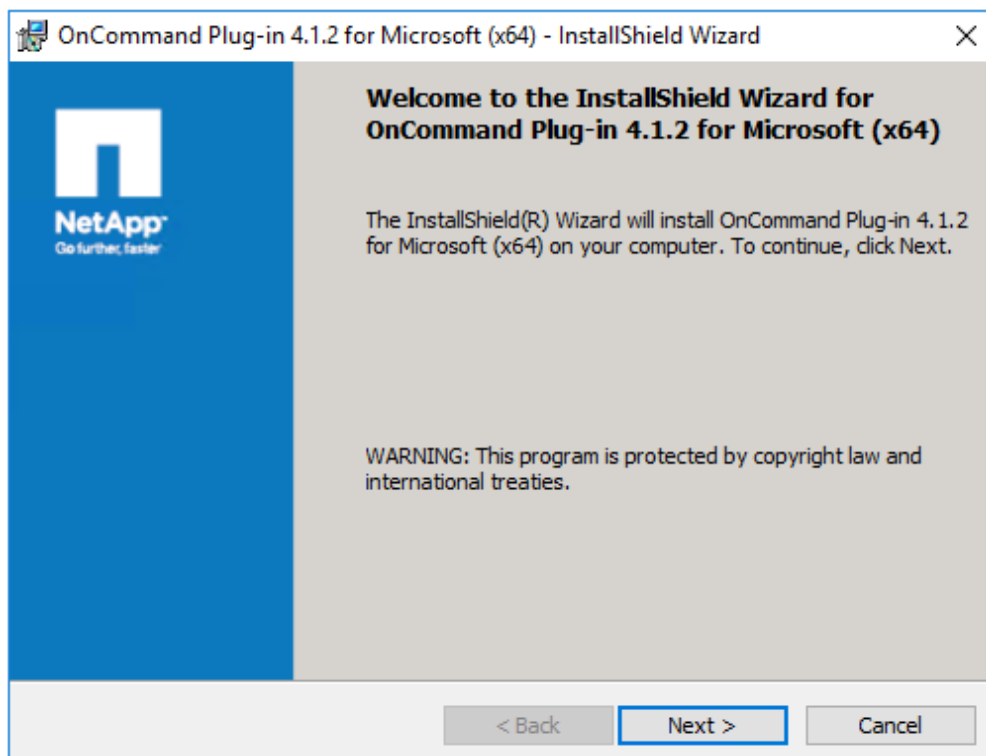
1. Log in to the [NetApp Support site](#).
2. Go to the Download Software page.

3. From the drop-down list, select the platform on which you are installing OnCommand Plug-in for Microsoft and click Go!.
4. Click View & Download for the software version (4.1.2) you want to install.
5. On the Description page, click Continue.
6. Review and accept the terms of the license agreement.
7. On the Download page, click the link for the installation file.
8. Save the installation file to a local or network directory accessible by the SCOM host.
9. Click Save File.
10. Verify the checksum to ensure that the software downloaded correctly.


Installing NetApp OnCommand Plug-In for Microsoft

To install NetApp OnCommand plug-in for Microsoft, complete the following steps:

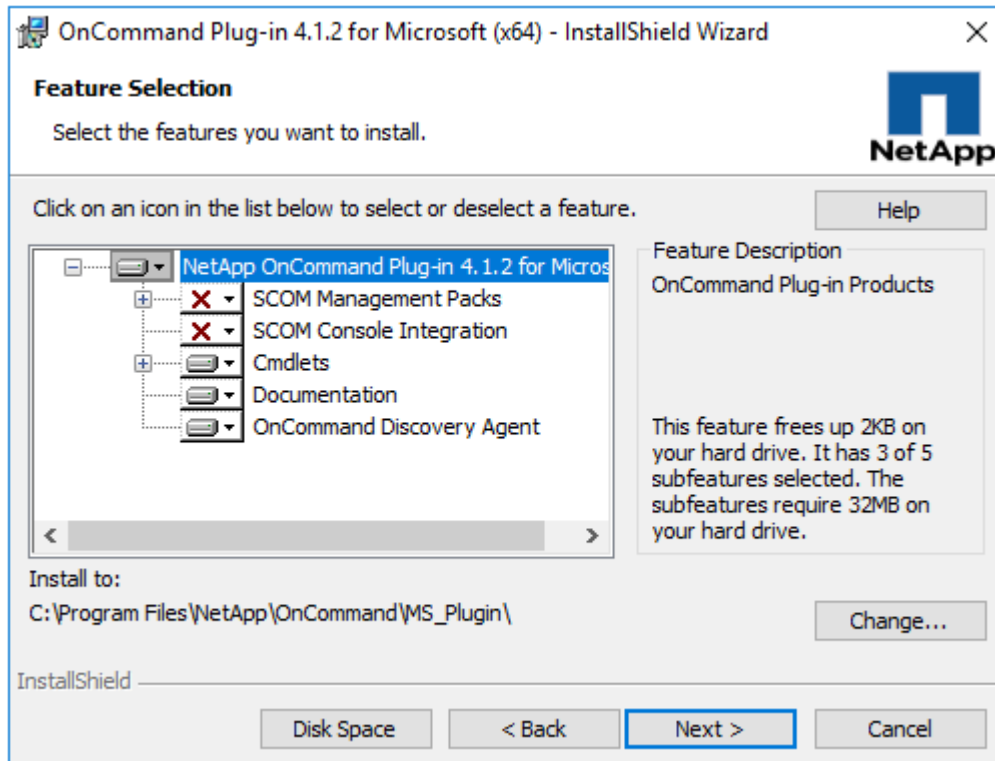
1. Launch the OnCommand Plug-in for Microsoft executable file and then follow the Installation wizard instructions.



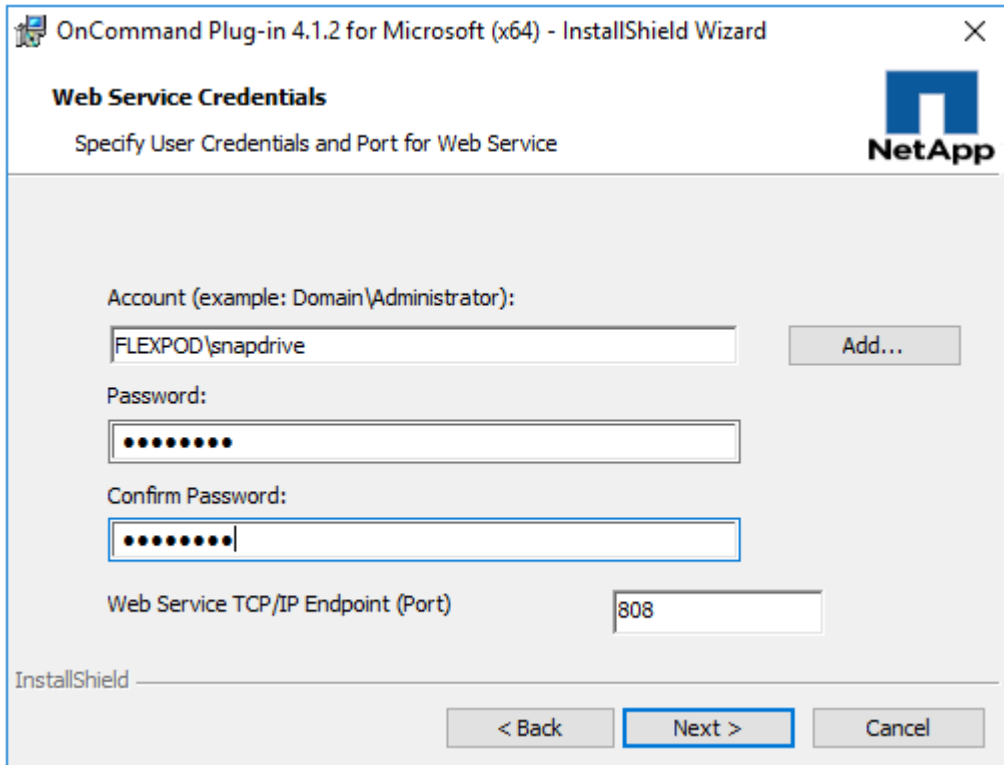
2. On the Welcome page, click Next.



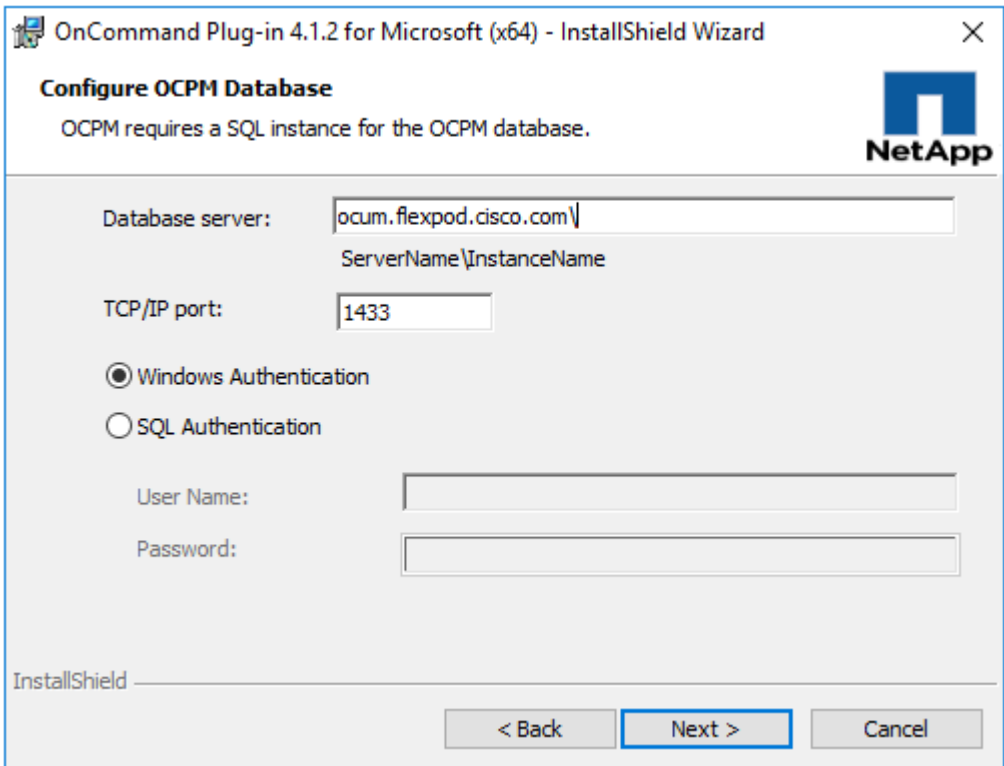
3. On the Destination page, select the destination and click Next.



4. On the Feature Selection page, select the desired features, and click Next.



5. On the Web Service Credentials page, enter the appropriate credentials and TCP port (the default port 808 should work fine), and click Next.



6. On the Configure OCPM Database page, enter the name, port, and authentication method for the OCUM server. If no instance is specified, the default OCPM database instance on the OCUM server will be used.
7. On the Ready to Install page, click Install.
8. When the installation completes click Finish.

Storage Configuration – Boot LUNs for Tenant Hyper-V Hosts

NetApp ONTAP Boot Storage Setup

Create igroups

To create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-01 -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-01-iqn>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-02 -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-02-iqn>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-All -protocol iscsi -ostype windows -
initiator <hyper-v-mgmt-01-iqn>,<hyper-v-mgmt-02-iqn>
```



To get the management host IQNs, log in to UCS Manager and click the Servers icon on the left. Then select Servers > Service Profiles > root and the host Service Profile. The host IQN is listed under the iSCSI vNICs tab on the right.

Map LUNs to igroups

To map LUNs to igroups, run the following commands:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun MGMT-Win2016-Gold -igroup Hyper-V-MGMT-01 -lun-id
0
lun map -vserver Infra-MS-SVM -volume witness -lun witness -igroup Hyper-V-MGMT-All -lun-id 1
lun map -vserver Infra-MS-SVM -volume iscsi_datastore_1 -lun iscsi_datastore_1 -igroup Hyper-V-MGMT-
All -lun-id 2
lun map -vserver Infra-MS-SVM -volume iscsi_datastore_2 -lun iscsi_datastore_2 -igroup Hyper-V-MGMT-
All -lun-id 3
```


Sample Tenant Setup

Add Supernet Routes to Core-Services Devices

In this FlexPod with Cisco ACI lab validation, a Core-Services subnet was setup in Tenant common to allow Tenant VMs to access Core Services such as DNS and Active Directory Authentication. Tenant VMs access the Core-Services devices over Layer 3 using their EPG subnet gateway. In this implementation, the Core-Services devices were setup connected by contract to the Bridged FP-Mgmt-Sw Network that had a default gateway outside of the ACI Fabric. Since the Core-Services devices use this default gateway that is outside of the ACI Fabric, persistent, static routes must be placed in the Core-Services devices to reach the Tenant VMs. To simplify this setup, all tenant VMs and devices connected to Core-Services had their IP subnets mapped from a range (172.18.0.0/16 in this deployment), allowing one supernet route to be put into each Core-Services device. This section describes the procedure for deploying these supernet routes to each type of Core-Services device.

Adding the Supernet Route in a Windows VM or Host

To add a persistent Supernet Route in a Windows VM (AD servers) or a Tenant Host, open a command prompt with Administrator privileges in Windows and type the following command:

```
route -p ADD 172.18.0.0 MASK 255.255.0.0 <core-services-EPG-gateway>
route print
```

ACI Shared Layer 3 Out Setup

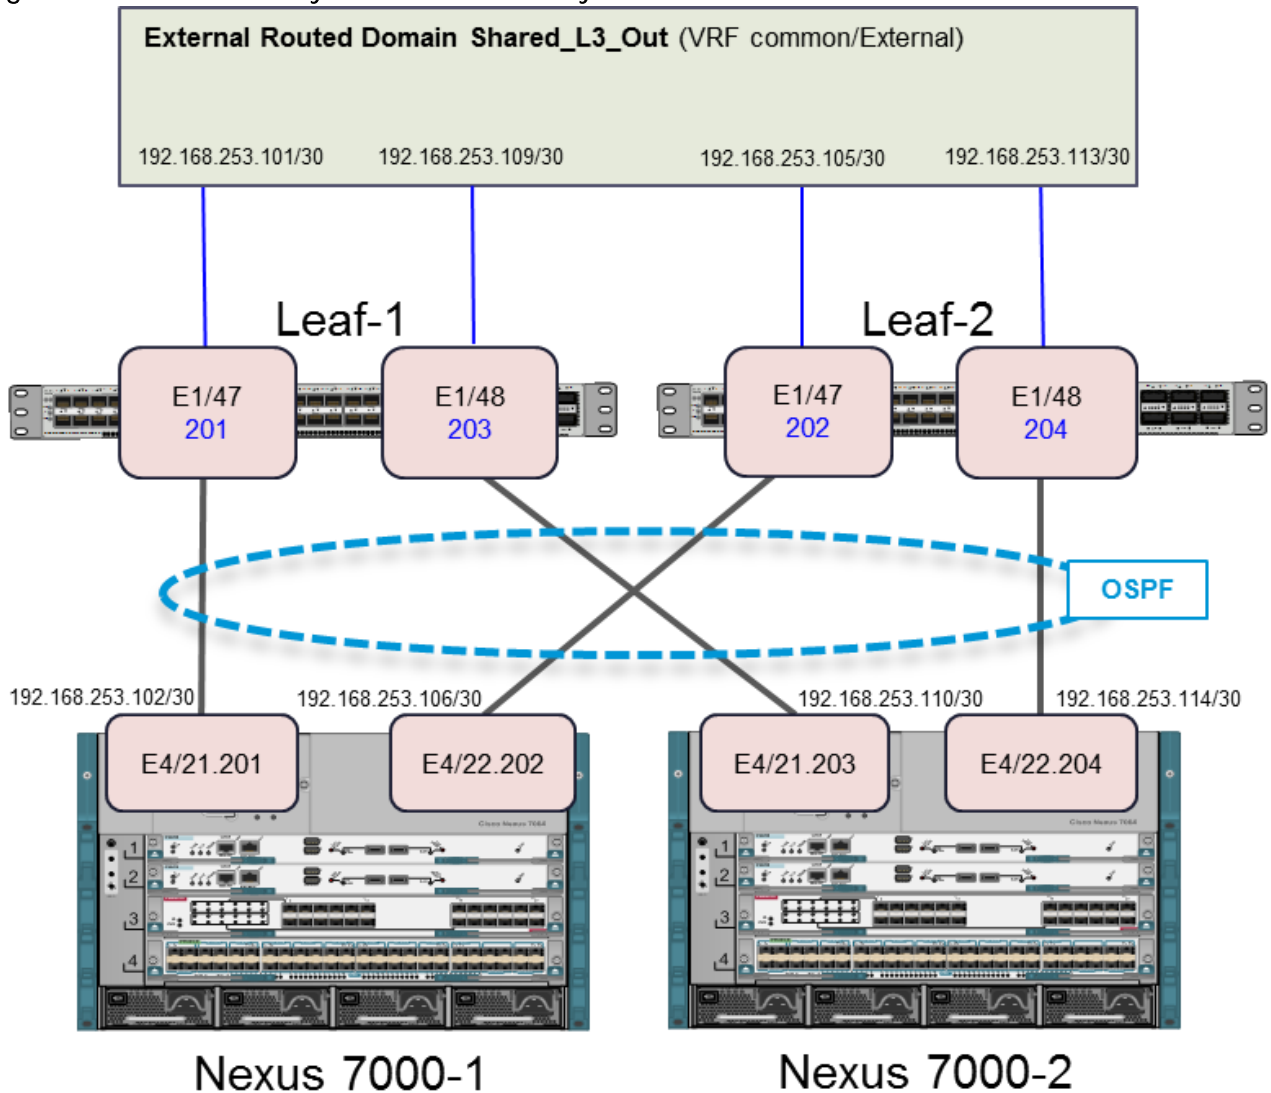
This section describes the procedure for deploying the ACI Shared Layer 3 Out. This external network is setup with a routing protocol and provides ACI tenants with a gateway to enter and leave the fabric.

This section provides a detailed procedure for setting up the Shared Layer 3 Out in Tenant common to existing Nexus 7000 core routers using sub-interfaces and VRF aware OSPF. Some highlights of this connectivity are:

- A new bridge domain and associated VRF is configured in Tenant common for external connectivity.
- **The shared Layer 3 Out created in Tenant common “provides” an external connectivity contract that can be “consumed” from any tenant.**
- Routes to tenant EPG subnets connected by contract are shared across VRFs with the Nexus 7000 core routers using OSPF.
- **The Nexus 7000s’ default gateway is shared with the ACI fabric using OSPF.**
- Each of the two Nexus 7000s is connected to each of the two Nexus 9000 leaf switches.
- Sub-interfaces are configured and used for external connectivity.
- The Nexus 7000s are configured to originate and send a default route to the Nexus 9000 leaf switches.

- This Shared Layer 3 Out was set up on a set of 10GE leaves that were part of the ACI fabric and not the 9332 leaves (which were also part of the fabric) used in this validation.

Figure 4 ACI Shared Layer 3 Out Connectivity Details



Configuring the Nexus 7000s for ACI Connectivity (Sample)

The following configuration is a sample from the virtual device contexts (VDCs) from two Nexus 7004s. Interfaces and a default route from the two Nexus 7000s also needs to be set up, but is not shown here because this would be set up according to customer security policy.

Nexus 7004-1 VDC

```
feature ospf
```

```
vlan 100
```

```
name OSPF-Peering

interface Vlan100
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.253.253/30
  no ipv6 redirects
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.0

interface Ethernet4/21
  no shutdown

interface Ethernet4/21.201
  encapsulation dot1q 201
  ip address 192.168.253.102/30
  ip ospf network point-to-point
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown

interface Ethernet4/22
  no shutdown

interface Ethernet4/22.202
  encapsulation dot1q 202
  ip address 192.168.253.106/30
  ip ospf cost 5
  ip ospf network point-to-point
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
```

```
no shutdown
```

```
interface loopback0  
  ip address 192.168.254.3/32  
  ip router ospf 10 area 0.0.0.0
```

```
router ospf 10  
  router-id 192.168.254.3  
  area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
```

Nexus 7004-2 VDC

```
feature ospf
```

```
vlan 100  
  name OSPF-Peering
```

```
interface Vlan100  
  no shutdown  
  mtu 9216  
  no ip redirects  
  ip address 192.168.253.254/30  
  no ipv6 redirects  
  ip ospf mtu-ignore  
  ip router ospf 10 area 0.0.0.0
```

```
interface Ethernet4/21  
  no shutdown
```

```
interface Ethernet4/21.203  
  encapsulation dot1q 203  
  ip address 192.168.253.110/30  
  ip ospf cost 21  
  ip ospf network point-to-point
```

```
ip ospf mtu-ignore
ip router ospf 10 area 0.0.0.10
no shutdown

interface Ethernet4/22
no shutdown

interface Ethernet4/22.204
encapsulation dot1q 204
ip address 192.168.253.114/30
ip ospf cost 30
ip ospf network point-to-point
ip ospf mtu-ignore
ip router ospf 10 area 0.0.0.10
no shutdown

interface loopback0
ip address 192.168.254.4/32
ip router ospf 10 area 0.0.0.0

router ospf 10
router-id 192.168.254.4
area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
```

Configuring ACI Shared Layer 3 Out

ACI GUI

1. At the top, select Fabric > Access Policies.
2. On the left, expand Physical and External Domains.
3. Right-click External Routed Domains and select Create Layer 3 Domain.
4. Name the Domain Shared-L3-Out.
5. Use the Associated Attachable Entity Profile pulldown to select Create Attachable Entity Profile.

6. Name the Profile AEP-Shared-L3-Out and click Next.

Create Attachable Access Entity Profile

STEP 1 > Profile

? ✕

1. Profile

2. Association To Interfaces

Specify the name, domains and infrastructure encaps

Name:

Description:

Enable Infrastructure VLAN:

EPG DEPLOYMENT (All Selected EPGs will be deployed on all the interfaces associated.)

Application EPGs	Encap	Primary Encap	Mode

Previous

Cancel

Next

7. Click Finish to continue without specifying interfaces.

8. Back in the Create Layer 3 Domain window, use the VLAN Pool pulldown to select Create VLAN Pool.

9. Name the VLAN Pool VP-Shared-L3-Out and select Static Allocation.

10. Click the + sign to add and Encap Block.

11. In the Create Ranges window, enter the From and To VLAN IDs for the Shared-L3-Out VLAN range (201-204). Select Static Allocation.

Create Ranges



Specify the Encap Block Range

Type: VLAN

Range: - - -
Integer Value Integer Value

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

- 12. Click OK to complete adding the VLAN range.
- 13. Click Submit to complete creating the VLAN Pool.

Create Layer 3 Domain



Specify the Layer 3 Domain

Name:

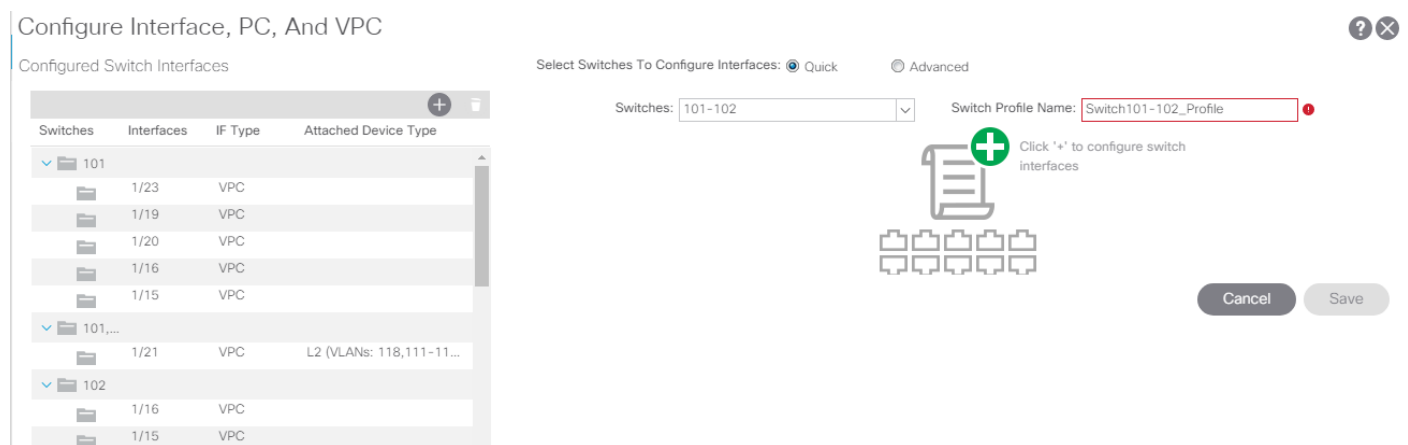
Associated Attachable Entity Profile:

VLAN Pool:

Security Domains:

Select	Name	Description
--------	------	-------------

14. Click Submit to complete creating the Layer 3 Domain.
15. At the top, select Fabric > Access Policies.
16. On the left, select Quick Start. Under Steps, select Configure an interface, PC, or VPC.
17. If an Interface, PC, or vPC has already been configured on the leaf pair being used here, select that switch pair in the list on the left and skip to step 19. Otherwise, in the center pane, click the green plus sign to select switches.
18. Using the Switches pull-down, select the two leaf switches connected to the Nexus 7000s and click away from the list to get the two switches filled in next to Switches. The Switch Profile Name will be automatically filled in.



19. Click the green plus sign to configure switch interfaces.
20. Next to interfaces, enter the 2-port identifiers for the ports connected to the Nexus 7000s and used for Shared-L3-Out. It is important to use the same two ports on each leaf. Fill in the policies, Attached Device Type, and External Route Domain as shown below.

Configure Interface, PC, And VPC

Configured Switch Interfaces

Switches	Interfaces	IF Type	Attached Device Type
> 101			
> 101,...			
	1/21	VPC	L2 (VLANs: 118,111-11...
> 102			
> 103			
> 103,...			
> 104			
> 106,...			
	1/25	VPC	L2 (VLANs: 318,906,2,9...
	1/26	VPC	L2 (VLANs: 318,906,2,9...
	1/1	VPC	Bare Metal (VLANs: 218,...
...	1/2	VPC	Bare Metal (VLANs: 218...

VPC Switch Pairs

VPC Domain Id	Switch 1	Switch 2
1	102	101
2	103	104
10	105	106

Select Switches To Configure Interfaces: Quick Advanced

Switches: Switch Profile Name:

Interface Type: Individual PC VPC

Interfaces: Interface Selector Name:
Select interfaces by typing, e.g. 1/17-18.

Interface Policy Group: Create One Choose One

Link Level Policy: <input type="text" value="10Gbps-Auto"/>	CDP Policy: <input type="text" value="CDP-Enabled"/>
MCP Policy: <input type="text" value="select a value"/>	LLDP Policy: <input type="text" value="LLDP-Enabled"/>
STP Interface Policy: <input type="text" value="BPDU-FG-Disabled"/>	Monitoring Policy: <input type="text" value="select a value"/>
Storm Control Policy: <input type="text" value="select a value"/>	L2 Interface Policy: <input type="text" value="VLAN-Scope-Global"/>
Port Security Policy: <input type="text" value="select a value"/>	Egress Data Plane Policing Policy: <input type="text" value="select a value"/>
Ingress Data Plane Policing Policy: <input type="text" value="select a value"/>	IPv4 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Priority Flow Control Policy: <input type="text" value="select a value"/>	IPv6 NetFlow Monitor Policy: <input type="text" value="select a value"/>
Slow Drain Policy: <input type="text" value="select a value"/>	Layer2-Switched (CE type) NetFlow Monitor Policy: <input type="text" value="select a value"/>

Fibre Channel Interface Policy:

Attached Device Type:

Domain: Create One Choose One External Route Domain:

21. On the lower right, click Save. Click Save again and click Submit.
22. At the top, select Tenants > common.
23. On the left, expand Tenant common and Networking.
24. Right-click VRFs and select create VRF.
25. Name the VRF common-External. Select default for both the End Point Retention Policy and Monitoring Policy.

Create VRF

STEP 1 > VRF

Specify Tenant VRF

Name:

Alias:

Description:

Policy Control Enforcement Preference: Enforced Unenforced

Policy Control Enforcement Direction: Egress Ingress

BD Enforcement Status:

End Point Retention Policy:
This policy only applies to remote L3 entries

Monitoring Policy:

DNS Labels:
enter names separated by comma

Route Tag Policy:

Create A Bridge Domain:

Configure BGP Policies:

Configure OSPF Policies:

Configure EIGRP Policies:



1. VRF **2. Bridge Domain**

26. Leave Create A Bridge Domain selected and click Next.

27. Name the Bridge Domain BD-common-External. Leave all other values unchanged.

Create VRF

STEP 2 > Bridge Domain

1. VRF

2. Bridge Domain

Specify Bridge Domain for the VRF

Name:	<input type="text" value="BD-common-External"/>	IGMP Snoop Policy:	<input type="text" value="select a value"/>
Alias:	<input type="text"/>	Monitoring Policy:	<input type="text" value="select a value"/>
Description:	<input type="text" value="optional"/>	ND policy:	<input type="text" value="select a value"/>
Type:	<input type="text" value="fc"/> <input checked="" type="text" value="regular"/>	ARP Flooding:	<input type="checkbox"/> Enabled
Forwarding:	<input type="text" value="Optimize"/>		
Endpoint Dataplane Learning:	<input checked="" type="checkbox"/>		
Limit IP Learning To Subnet:	<input checked="" type="checkbox"/>		
Config BD MAC Address:	<input checked="" type="checkbox"/>		
MAC Address:	<input type="text" value="00:22:BD:F8:19:FF"/>		

Previous

Cancel

Finish

28. Click Finish to complete creating the VRF.
29. On the left, right-click External Routed Networks and select Create Routed Outside.
30. Name the Routed Outside Shared-L3-Out.
31. Select the checkbox next to OSPF.
32. Enter 0.0.0.10 (configured in the Nexus 7000s) as the OSPF Area ID.
33. Using the VRF pulldown, select common/common-External.
34. Using the External Routed Domain pulldown, select Shared-L3-Out.
35. Click the + sign to the right of Nodes and Interfaces Protocol Profiles to add a Node Profile.
36. Name the Node Profile Nodes-101-102 for the Nexus 9000 Leaf Switches.

37. Click the + sign to the right of Nodes to add a Node.

38. In the select Node window, select Leaf switch 101.

39. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.101).

Select Node



Select Node and Configure Static Routes

Node ID: a01-93180-a (Node-101) ▼

Router ID: 192.168.254.101

Use Router ID as Loopback Address:

Loopback Addresses:

IP
192.168.254.101

Static Routes:

IP Address	Next Hop IP
------------	-------------

Cancel

OK

40. Click OK to complete selecting the Node.

41. Click the + sign to the right of Nodes to add a Node.

42. In the select Node window, select Leaf switch 102.

43. Provide a Router ID IP address that will also be used as the Loopback Address (192.168.254.102).

Create Routed Outside

1. Identity 2. External FPC Networks

Select Node

Select Node and Configure Static Routes

Node ID: a01-93180-b (Node-102)

Router ID: 192.168.254.102

Use Router ID as Loopback Address:

Loopback Addresses:

IP
192.168.254.102

Static Routes:

IP Address	Next Hop IP
------------	-------------

Cancel OK

44. Click OK to complete selecting the Node.

45. Click the + sign to the right of OSPF Interface Profiles to create an OSPF Interface Profile.

46. Name the profile OIP-Nodes-101-102.

Create Interface Profile



STEP 1 > Identity

1. Identity

2. Protocol Profiles

3. Interfaces

Specify the Interface Profile

Name:

Description:

ND policy: ▼

Egress Data Plane Policing Policy: ▼

Ingress Data Plane Policing Policy: ▼

NetFlow Monitor Policies: 🗑️ +

NetFlow IP Filter Type	NetFlow Monitor Policy
------------------------	------------------------

Config Protocol Profiles:

Previous

Cancel

Next

47. Click Next.

48. Using the OSPF Policy pulldown, select Create OSPF Interface Policy.

49. Name the policy To-7K.

50. Select the Point-to-Point Network Type.

51. Select the Advertise subnet and MTU ignore Interface Controls.

Create OSPF Interface Policy



Define OSPF Interface Policy

Name:

Description:

Network Type: Broadcast Point-to-point Unspecified

Priority:

Cost of Interface:

Interface Controls:

- Advertise subnet
- BFD
- MTU ignore
- Passive participation

Hello Interval (sec):

Dead Interval (sec):

Retransmit Interval (sec):

Transmit Delay (sec):

Cancel

Submit

52. Click SUBMIT to complete creating the policy.

Create Interface Profile



STEP 2 > Protocol Profiles

1. Identity

2. Protocol Profiles

3. Interfaces

Specify the Protocol Profiles

OSPF Profile

Authentication Type:

Authentication Key:

Confirm Key:

OSPF Policy:

BFD Interface Profile

Authentication Type:

BFD Interface Policy:

HSRP Interface Profile

Enable HSRP: HSRP version: HSRP Interface Policy:

HSRP Interface Groups:

Name	Group ID	IP	MAC	Group Name	Group Type	IP Obtain Mode

Previous

Cancel

Next

53. Click Next.

54. Select Routed Sub-Interface under Interfaces.

55. Click the + sign to the right of Routed Sub-Interfaces to add a routed sub-interface.

56. In the Select Routed Sub-Interface window, select the interface on Node 101 that is connected to Nexus 7000-1.

57. Enter VLAN 201 for Encap.

58. Enter the IPv4 Primary Address (192.168.253.101/30)

59. Leave the MTU set to inherit.

Select Routed Sub-Interface



Specify the Interface

Node:
Ex: topology/pod-1/node-17

Path:
Ex: Pod-1/Node-101/[Fex-110]/eth1/2

Description:

Encap:
Integer Value

IPv4 Primary / IPv6 Preferred Address:
address/mask

IPv4 Secondary / IPv6 Additional
 Addresses:
Address

MAC Address:

MTU (bytes):

Link-local Address:

Cancel

OK

60. Click OK to complete creating the routed sub-interface.

61. Repeat steps 54-59 to add the second Leaf 1 interface (VLAN 203, IP 192.168.253.109/30), the first Leaf 2 interface (VLAN 202, IP 192.168.253.105/30), and the second Leaf 2 interface (VLAN 204, IP 192.168.253.113/30).

Create Interface Profile



STEP 3 > Interfaces

1. Identity 2. Protocol Profiles **3. Interfaces**

Specify the Interfaces

Routed Interfaces SVI **Routed Sub-Interface**

Routed Sub-Interfaces			
Path	IP Address	MAC Address	MTU (bytes)
Pod-1/Node-101/eth1/47	192.168.253.101/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-101/eth1/48	192.168.253.109/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-102/eth1/47	192.168.253.105/30	00:22:BD:F8:19:FF	inherit
Pod-1/Node-102/eth1/48	192.168.253.113/30	00:22:BD:F8:19:FF	inherit

Previous Cancel **OK**

62. Click OK to complete creating the Node Interface Profile.

Create Node Profile



Specify the Node Profile

Name:

Description:

Target DSCP: ▼

Nodes: 🗑️ +

Node ID	Router ID	Static Routes	Loopback Address
topology/pod-1/...	192.168.254.101		
topology/pod-1/...	192.168.254.102		

OSPF Interface Profiles: 🗑️ +

Name	Description	Interfaces	OSPF Policy
OIP-Nodes-101-102		[eth1/47], [eth1/47], [eth1/48], [eth1/48]	To-7K

Cancel OK

63. Click OK to complete creating the Node Profile.

Create Routed Outside

STEP 1 > Identity

1. Identity | 2. External EPG Networks

Define the Routed Outside

Name:

Alias:

Description:

Tags:

PIM:

Route Control Enforcement: Import Export

Target DSCP:

VRF:

External Routed Domain:

Route Profile for Interleak:

Provider Label:

Consumer Label:

BGP EIGRP OSPF

OSPF Area ID:

OSPF Area Control: Send redistributed LSAs into NSSA area
 Originate summary LSA
 Suppress forwarding address in translated LSA

OSPF Area Type: NSSA area Regular area Stub area

OSPF Area Cost:

Route Control For Dampening:

Address Family Type	Route Dampening Policy

Nodes and Interfaces Protocol Profiles

Name	Description	DSCP	Nodes
Nodes-101-102		Unspecified	101, 102

Previous | Cancel | Next

64. Click Next.
65. Click the + sign under External EPG Networks to create and External EPG Network.
66. Name the External Network Default-Route.
67. Click the + sign to add a Subnet.
68. Enter 0.0.0.0/0 as the IP Address. Select the checkboxes for External Subnets for the External EPG, Shared Route Control Subnet, and Shared Security Import Subnet.

Create Subnet



Specify the Subnet

IP Address:
address/mask

- scope:
- Export Route Control Subnet
 - Import Route Control Subnet
 - External Subnets for the External EPG
 - Shared Route Control Subnet
 - Shared Security Import Subnet

OSPF Route Summarization Policy: ▼

- aggregate:
- Aggregate Export
 - Aggregate Import
 - Aggregate Shared Routes

Route Control Profile:

Name	Direction		

Cancel OK

69. Click OK to complete creating the subnet.

Create External Network



Define an External Network

Name:

Alias:

Tags:
enter tags separated by comma

QoS class:

Description:

Target DSCP:

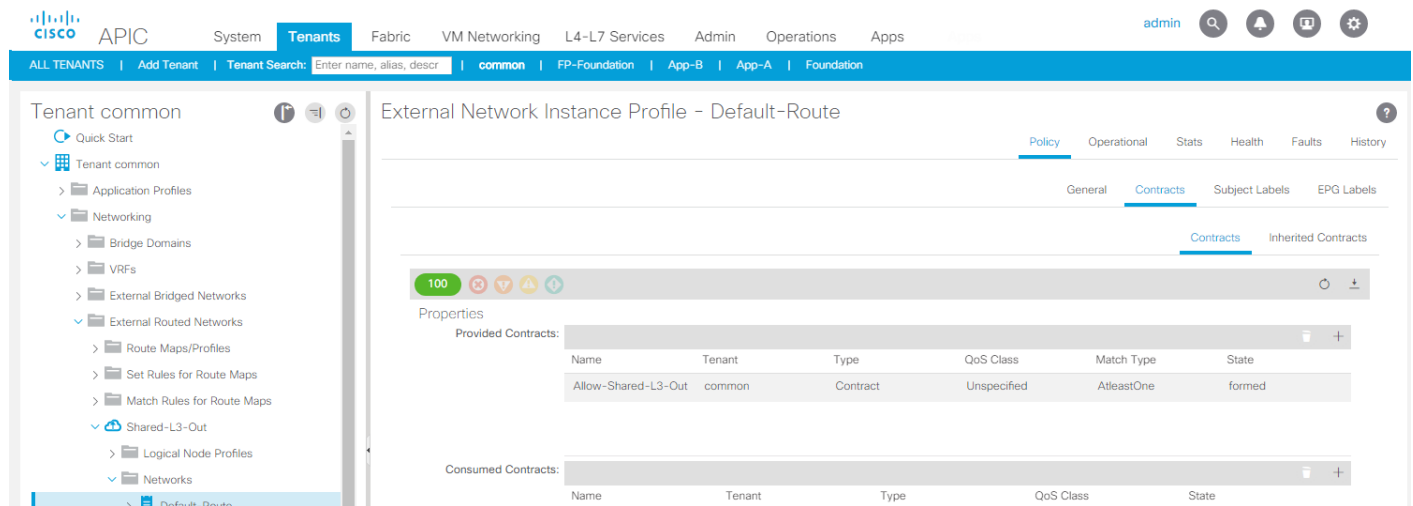
Preferred Group Member:


Subnet

IP Address	Scope	Aggregate	Route Control Profile	Route Summarization Policy
0.0.0.0/0	External Subnets for the Ex...	Shared Route Control Subn...	Shared Security Import Su...	

70. Click OK to complete creating the external network.
71. Click Finish to complete creating the Shared-L3-Out.
72. On the left, right-click Security Policies and select Create Contract.
73. Name the contract Allow-Shared-L3-Out.
74. Select the Global Scope to allow the contract to be consumed from all tenants.
75. Click the + sign to the right of Subjects to add a contract subject.
76. Name the subject Allow-All.
77. Click the + sign to the right of Filters to add a filter.
78. Use the pulldown to select the Allow-All filter from Tenant common.
79. Click Update.

80. Click OK to complete creating the contract subject.
81. Click Submit to complete creating the contract.
82. On the left, expand Tenant common, Networking, External Routed Networks, Shared-L3-Out, and Networks. Select Default-Route.
83. On the right, under Policy, select Contracts.
84. Click the + sign to the right of Provided Contracts to add a Provided Contract.
85. Select the common/Allow-Shared-L3-Out contract and click Update.



 Tenant EPGs can now consume the Allow-Shared-L3-Out contract and connect outside of the fabric. More restrictive contracts can be built and provided here for more restrictive access to the outside.

Lab Validation Tenant Configuration

The following table shows the VLANs, Subnets, and Bridge Domains for the sample App-A Tenant set up as part of this lab validation:

Table 12 Lab Validation Tenant MS-TNT-A Configuration

EPG	Storage VLAN	UCS VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3014	Virtual Switch	192.168.14.0/24 – L2	BD-iSCSI-A
iSCSI-B	3024	Virtual Switch	192.168.24.0/24 – L2	BD-iSCSI-B
SMB	3055	Virtual Switch	192.168.55.0/24 – L2	BD-SMB
SVM-MGMT	219	Virtual Switch	172.18.254.6/29	BD-Internal
Web	N/A	Virtual Switch	172.18.0.254/24	BD-Internal
App	N/A	Virtual Switch	172.18.1.254/24	BD-Internal
DB	N/A	Virtual Switch	172.18.2.254/24	BD-Internal

Deploy ACI Application (MS-TNT-A) Tenant

This section details the steps for creation of the MS-TNT-A Sample Tenant in the ACI Fabric. This tenant will host application connectivity between the compute (Hyper-V Server on UCS) and the storage (NetApp) environments. This tenant will also host the three application tiers of the sample three-tier application. A corresponding MS-TNT-A-SVM will be created on the NetApp storage to align with this tenant. To deploy the MS-TNT-A Tenant, complete the following steps:

1. In the APIC GUI, select Fabric > Access Policies.
2. On the left, expand Pools and VLAN.
3. Select the storage VLAN Pool created earlier (NetApp-AFF_vlans).
4. In the center pane, click the + sign to add an encapsulation block.
5. Enter <storage-MS-TNT-A-SMB-VLAN> for the From and To fields.
6. Select Static Allocation.

Create Ranges



Specify the Encap Block Range

Type: VLAN

Range: VLAN - VLAN
Integer Value Integer Value

Allocation Mode: Dynamic Allocation Inherit allocMode from parent Static Allocation

Cancel

Submit

7. Click Submit to add the Encap Block Range.
8. In the center pane, click the + sign to add another encapsulation block.
9. Enter <storage-MS-TNT-A-SVM-MGMT-VLAN> for the From and To fields.
10. Select Static Allocation.
11. Click Submit to add the Encap Block Range.
12. If iSCSI LUN access is being provided by the MS-TNT-A tenant, complete steps 13-29. Otherwise, continue at step 30.
13. In the center pane, click the + sign to add an encapsulation block.

14. Enter <storage-MS-TNT-A-iSCSI-A-VLAN> for the From and To fields.
15. Select Static Allocation.
16. Click SUBMIT to add the Encap Block Range.
17. In the center pane, click the + sign to add an encapsulation block.
18. Enter <storage-MS-TNT-A-iSCSI-B-VLAN> for the From and To fields.
19. Select Static Allocation.
20. Click SUBMIT to add the Encap Block Range.
21. At the top select Tenants > Add Tenant.
22. Name the Tenant MS-TNT-A. Select the default Monitoring Policy.
23. For the VRF Name, also enter MS-TNT-A. Leave the Take me to this tenant when I click finish checkbox checked.

Create Tenant



Specify tenant details

Name:

Alias:

Description: optional

Tags:
 enter tags separated by comma

GUID:

Provider	GUID	Account Name

Monitoring Policy:

Security Domains:

Name	Description

VRF Name:

Take me to this tenant when I click finish

Cancel

Submit

24. Click Submit to finish creating the Tenant.
25. On the left under Tenant MS-TNT-A, right-click Application Profiles and select Create Application Profile.
26. Name the Application Profile Host-Conn select the default Monitoring Policy, and click Submit to complete adding the Application Profile.
27. If you are using providing iSCSI LUN access from this tenant, complete steps 28-53. Otherwise, continue to step 54.
28. On the left, expand Application Profiles and Host-Conn.
29. Right-click Application EPGs and select Create Application EPG.
30. Name the EPG iSCSI-A. Leave Intra EPG Isolation Unenforced.

31. Use the Bridge Domain pulldown to select Create Bridge Domain.
32. Name the Bridge Domain BD-iSCSI-A.
33. Select the MS-TNT-A VRF.
34. Use the Forwarding pulldown to select Custom.
35. Select Flood for the L2 Unknown Unicast and default for the End Point Retention Policy and IGMP Snoop Policy.

Create Bridge Domain ? ✕

STEP 1 > Main 1. Main 2. L3 Configurations 3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type:

VRF: ↗

Forwarding: ▼

L2 Unknown Unicast: ▼

L3 Unknown Multicast Flooding: ▼

Multi Destination Flooding: ▼

ARP Flooding: Enabled

Clear Remote MAC Entries:

End Point Retention Policy: ↗
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy: ↗

Previous
Cancel
Next

36. At the bottom right, click Next.
37. Make sure Unicast Routing is Enabled and click Next.
38. Select the default Monitoring Policy and click Finish.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

Previous

Cancel

Finish

39. Select the default Monitoring Policy and click Finish to complete creating the EPG.

40. On the left, expand Application EPGs and EPG iSCSI-A. Right-click Domains and select Add Physical Domain Association.

41. Using the pulldown, select the NetApp-AFF Physical Domain Profile.

Add Physical Domain Association



Choose the Physical domain to associate

Physical Domain Profile:  

- 42. Click Submit to complete the Physical Domain Association.
- 43. Right-click Domains and select Add VMM Domain Association.
- 44. From the drop-down list, select APIC-MS-vSwitch. Set Deploy Immediacy to Immediate.

Add VMM Domain Association



Choose the VMM domain to associate

VMM Domain Profile:

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

VLAN Mode: Dynamic Static

Delimiter:

Cancel

Submit

45. Click Submit to complete the VMM Domain Association.



In this deployment for iSCSI, you are adding both the NetApp LIF endpoints and the Windows Host Interface endpoints in a single EPG. This method allows unrestricted communication within the EPG. You also had the choice to put the LIFs in one EPG and the Host Interfaces in a second EPG and connect them with a filtered contract. You will deploy SMB that way next.

46. Right-click Static-Ports and select Deploy Static EPG on PC, VPC, or Interface.

47. In the Deploy Static EPG on PC, VPC, Or Interface Window, select the Virtual Port Channel Path Type.

48. Using the Path pulldown, select the VPC for NetApp Storage Controller 01.

49. Enter <storage-MS-TNT-A-iSCSI-A-VLAN> for Port Encap.

50. Select the Immediate Deployment Immediacy and the Trunk Mode.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Port Encap (or Secondary VLAN for Micro-Seg):
Integer Value

Deployment Immediacy: Immediate On Demand

Primary VLAN for Micro-Seg:
Integer Value

Mode: Trunk Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

51. Click Submit to complete adding the Static Port Mapping.

52. Repeat steps 46-51 to add the Static Path Mapping for NetApp Storage Controller 02.

Path	Primary VLAN for Micro-Seg	Port Encap (or Secondary VLAN for Micro-Seg)	Deployment Immediacy	Mode
Node: Pod-1				
Pod-1/Node-105-106/A02-AFFA300-...	vlan-3014		Immediate	Trunk
Pod-1/Node-105-106/A02-AFFA300-...	vlan-3014		Immediate	Trunk

53. Repeat steps 28-52 to build the iSCSI-B EPG. Make sure to create a separate Bridge Domain for this EPG and use the MS-TNT-A iSCSI-B VLAN IDs.


Path	Primary VLAN For Micro-Seg	Port Encap (Or Secondary VLAN For Micro-Seg)	Deployment Immediacy	Mode
Node: Node-101-102				
Node-101-102/VPC-a01-6248-a	vlan-3121		Immediate	Trunk
Node-101-102/VPC-a01-6248-b	vlan-3121		Immediate	Trunk
Node-101-102/VPC-a01-aff8040-01	vlan-3021		Immediate	Trunk
Node-101-102/VPC-a01-aff8040-02	vlan-3021		Immediate	Trunk

54. On the left, under Tenant MS-TNT-A, right-click the Host-Conn Application Profile and select Create Application EPG.

55. Name the EPG `SMB-LIF` and leave Intra EPG Isolation set at Unenforced.

56. Use the Bridge Domain pulldown to select Create Bridge Domain.

57. Name the Bridge Domain `BD-SMB` and select the MS-TNT-A VRF.

 It is important to create a new Bridge Domain for each traffic VLAN coming from the NetApp Storage Controllers. All of the VLAN interfaces on a given NetApp Interface Group share the same MAC address, and separating to different bridge domains in the ACI Fabric allows all the traffic to be forwarded properly.

58. For Forwarding, select Custom and select Flood for L2 Unknown Unicast. Select default for the End Point Retention Policy and the IGMP Snoop Policy.

Create Bridge Domain



STEP 1 > Main

1. Main

2. L3 Configurations

3. Advanced/Troubleshooting

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Type:

VRF:

Forwarding:

L2 Unknown Unicast:

L3 Unknown Multicast Flooding:

Multi Destination Flooding:

ARP Flooding: Enabled

Clear Remote MAC Entries:

End Point Retention Policy:
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

Previous

Cancel

Next

59. At the bottom right, click Next.

60. Make sure Unicast Routing is enabled and click Next.

61. Select the default Monitoring Policy and click Finish.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description:

Tags: enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

62. Select the default Monitoring Policy and click Finish to complete creating the EPG.

63. On the left expand Host-Conn, Application EPGs, and EPG SMB-LIF.


64. Right-click Domains and select Add Physical Domain Association.

65. Select the NetApp-AFF Physical Domain Profile.

Add Physical Domain Association



Choose the Physical domain to associate

Physical Domain Profile:  

Cancel

Submit

-

66. Click Submit to complete adding the Physical Domain Association.
67. Right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.
68. Select the Virtual Port Channel Path Type.
69. Using the Path pulldown, select the VPC for NetApp Storage Controller 01.
70. For Port Encap, enter `vlan-<storage-MS-TNT-A-SMB-VLAN>`.
71. Select Immediate for Deployment Immediacy and Trunk for Mode.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Port Encap (or Secondary VLAN for Micro-Seg):
Integer Value

Deployment Immediacy: Immediate On Demand

Primary VLAN for Micro-Seg:
Integer Value

Mode: Trunk Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

Cancel

Submit

72. Click Submit to finish adding the EPG Static Binding.
73. Repeat steps 67-72 for the Static Port for NetApp Storage Controller 02.
74. On the left under EPG SMB-LIF, right-click Contracts and select Add Provided Contract.
75. In the Add Provided Contract window, use the Contract pulldown to select Create Contract.
76. Name the contract `Allow-SMB`. Leave the Scope set at VRF.
77. Click the + sign to add a Contract Subject.
78. Name the subject `Allow-SMB`.
79. Click the + sign to add a Filter to the Filter Chain.

80. Click the pulldown and select NetApp-SMB from Tenant common.

81. Click Update.

Create Contract Subject



Specify Identity Of Subject

Name:

Alias:

Description:

Target DSCP:

Apply Both Directions:

Reverse Filter Ports:

Filter Chain

Filters		
Name	Directives	
common/NetApp-SMB	none	

L4-L7 SERVICE GRAPH
Service Graph:

PRIORITY
QoS:

Cancel OK



Optionally, add ICMP to the filter chain to allow ping in this contract for troubleshooting purposes.

82. Click OK to complete the Contract Subject.

Create Contract



Specify Identity Of Contract

Name:

Alias:

Scope: ▼

QoS Class: ▼

Target DSCP: ▼

Description:

Tags: ▼
enter tags separated by comma

Subjects: 🗑️ +

Name	Description
Allow-SMB	

Cancel Submit

83. Click Submit to complete creating the Contract.

Add Provided Contract



Select a contract

Contract: 

QoS:

Contract Label:

Subject Label:

Cancel

Submit

84. Click Submit to complete Adding the Provided Contract.

85. Right-click Application EPGs under the Host-Conn Application Profile and select Create Application EPG.

86. Name the EPG `SMB-Host` and leave Intra EPG Isolation set at Unenforced.

87. Use the Bridge Domain drop-down list to select `BD-SMB`. Select the default Monitoring Policy.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

Application EPGs

88. Click Finish to complete creating the EPG.

89. On the left expand Host-Conn, Application EPGs, and EPG SMB-Host.

90. Under EPG SMB-Host, right-click Domains and select Add VMM Domain Association.


91. Select the APIC-MS-vSwitch VMM Domain Profile.

92. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy.

Add VMM Domain Association



Choose the VMM domain to associate

VMM Domain Profile: 

Deploy Immediacy: Immediate On Demand

Resolution Immediacy: Immediate On Demand Pre-provision

VLAN Mode: Dynamic Static

Delimiter:

Cancel

Submit

93. Click Submit to complete adding the VMM Domain Association.
94. On the left under EPG SMB-Host, right-click Contracts and select Add Consumed Contract.
95. In the Add Consumed Contract window, from the Contract drop-down list, select MS-TNT-A/Allow-SMB.

Add Consumed Contract



Select a contract

Contract:  

QoS: 

Contract Label:

Subject Label:

Cancel

Submit

96. Click Submit to complete adding the Consumed Contract.
97. On the left, under Tenant MS-TNT-A, right-click Application Profiles and select Create Application Profile.
98. Name the Profile `sVM-MGMT`, select the default Monitoring Policy, and click SUBMIT.
99. Right-click the SVM-MGMT Application Profile and select Create Application EPG.
100. Name the EPG `sVM-MGMT` and leave Intra EPG Isolation set at Unenforced.
101. Use the Bridge Domain pulldown to select create Bridge Domain.
102. Name the Bridge Domain `BD-Internal` and select the MS-TNT-A VRF.
103. Leave Forwarding set at optimize, select the default End Point Retention Policy and IGMP Snoop Policy.
104. Click Next.
105. Make sure Unicast Routing is enabled and click Next.
106. Select the default Monitoring Policy and click Finish to complete creating the Bridge Domain.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description:

Tags: enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

107. Select the default Monitoring Policy and click Finish to complete creating the EPG.
108. On the left expand SVM-MGMT, Application EPGs, and EPG SVM-MGMT.
109. Right-click Domains and select Add Physical Domain Association.
110. Select the NetApp-AFF Physical Domain Profile.

Add Physical Domain Association



Choose the Physical domain to associate

Physical Domain Profile:  

Cancel

Submit

111. Click Submit to complete adding the Physical Domain Association.
112. Right-click Static Ports and select Deploy Static EPG on PC, VPC, or Interface.
113. Select the Virtual Port Channel Path Type.
114. Using the Path pulldown, select the VPC for NetApp Storage Controller 01.
115. For Port Encap, enter <storage-MS-TNT-A-SVM-MGMT-VLAN>.
116. Select Immediate for Deployment Immediacy and Trunk for Mode.

Deploy Static EPG On PC, VPC, Or Interface



Select PC, VPC, or Interface

Path Type: Port Direct Port Channel Virtual Port Channel

Path:

Port Encap (or Secondary VLAN for Micro-Seg):
Integer Value

Deployment Immediacy: Immediate On Demand

Primary VLAN for Micro-Seg:
Integer Value

Mode: Trunk Access (802.1P) Access (Untagged)

IGMP Snoop Static Group:

Group Address	Source Address

Cancel

Submit

117. Click Submit to finish adding the EPG Static Binding.
118. Repeat steps 112-117 for the Static Port Mapping to NetApp Storage Controller 02.
119. On the left under EPG SVM-MGMT, right-click Contracts and select Add Provided Contract.
120. In the Add Provided Contract window, use the Contract pulldown to select Create Contract.
121. Name the contract `allow-svm-mgmt`. Leave the Scope set at VRF.
122. Click the + sign to add a Contract Subject.
123. Name the subject `allow-all`. Click the + sign to add a filter.
124. Use the pulldown to select the Allow-All filter from Tenant common. Click Update.

125. Click OK to complete adding the Contract Subject.

Create Contract



Specify Identity Of Contract

Name:

Alias:

Scope: ▼

QoS Class: ▼

Target DSCP: ▼

Description:

Tags: ▼
enter tags separated by comma

Subjects: 🗑️ +

Name	Description
Allow-All	

Cancel

Submit

126. Click Submit to complete creating the Contract.

127. Click Submit to complete adding the Provided Contract.

128. On the left under EPG SVM-MGMT, right-click Subnets and select Create EPG Subnet.

129. For the Default Gateway IP, enter the SVM gateway IP address and mask or the MS-TNT-A tenant.
130. Select only the Shared between VRFs scope.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy:

Cancel

Submit

131. Click Submit to complete adding the EPG subnet.
132. On the left under EPG SVM-MGMT, right-click Contracts and select Add Consumed Contract.
133. In the Add Consumed Contract window, use the Contract drop-down list to select FP-Allow-Common-Core-Services in Tenant common.

Add Consumed Contract



Select a contract

Contract: 

QoS:

Contract Label:

Subject Label:

Cancel

Submit

134. Click Submit to complete adding the Consumed Contract.
135. On the left, right-click Application Profiles and select Create Application Profile.
136. Name the Application Profile `3-Tier-App` and select the default Monitoring Policy.
137. Click Submit to complete creating the Application Profile.
138. Expand `3-Tier-App`, right-click Application EPGs under `3-Tier-App` and select Create Application EPG.
139. Name the EPG `web` and leave Intra EPG Isolation set at Unenforced.
140. Use the Bridge Domain pulldown to select `MS-TNT-A/BD-Internal`. Select the default Monitoring Policy.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

141. Click Finish to complete creating the EPG.
142. On the left expand 3-Tier-App, Application EPGs, and EPG Web.
143. Under EPG Web, right-click Domains and select Add VMM Domain Association.
144. Select the APIC-MS-vSwitch VMM Domain Profile.
145. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Leave VLAN Mode set at Dynamic.
146. Click Submit to complete adding the VMM Domain Association.
147. On the left under EPG Web, right-click Contracts and select Add Provided Contract.

148. In the Add Provided Contract window, use the Contract pulldown to select Create Contract.
149. Name the Contract Allow-Web-App. Select the Application Profile Scope.
150. Click the + sign to add a Contract Subject.
151. Name the subject Allow-All.
152. Click the + sign to add a Contract filter.
153. Use the pulldown to select the Allow-All filter from Tenant common. Click Update.
154. Click OK to complete creating the Contract Subject.

Create Contract



Specify Identity Of Contract

Name:

Alias:

Scope:

QoS Class:

Target DSCP:

Description:

Tags:
enter tags separated by comma

Subjects:

Name	Description
Allow-All	

- 155. Click Submit to complete creating the Contract.
- 156. Click Submit to complete adding the Provided Contract.
- 157. Right-click Contracts and select Add Consumed Contract.

158. In the Add Consumed Contract window, use the Contract pulldown to select the common/Allow-Shared-L3-Out contract.
159. Click Submit to complete adding the Consumed Contract.
160. Optionally, repeat steps 157-159 to add the common/FP-Allow-Common-Core-Services Consumed Contract.
161. On the left under EPG Web, right-click Subnets and select Create EPG Subnet.
162. For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.18.0.0/16) that was set up for assigning Tenant IP addresses.
163. For scope, select Advertised Externally and Shared between VRFs.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy: ▼

164. Click Submit to complete creating the EPG Subnet.
165. Right-click Application EPGs under 3-Tier-App and select Create Application EPG.
166. Name the EPG `App` and leave Intra EPG Isolation set at Unenforced.
167. Use the Bridge Domain pulldown to select MS-TNT-A/BD-Internal. Select the default Monitoring Policy.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description:

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

Application EPGs

168. Click Finish to complete creating the EPG.
169. On the left expand 3-Tier-App, Application EPGs, and EPG App.
170. Under EPG Web, right-click Domains and select Add VMM Domain Association.
171. Select the APIC-MS-vSwitch VMM Domain Profile.
172. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Select the Dynamic VLAN Mode.
173. Click Submit to compete adding the VMM Domain Association.
174. On the left under EPG App, right-click Contracts and select Add Provided Contract.

175. In the Add Provided Contract window, use the Contract pulldown to select Create Contract.
176. Name the Contract `Allow-App-DB`. Select the Application Profile Scope.
177. Click the + sign to add a Contract Subject.
178. Name the subject `Allow-All`.
179. Click the + sign to add a Contract filter.
180. Use the pulldown to select the Allow-All filter from Tenant common. Click Update.
181. Click OK to complete creating the Contract Subject.
182. Click Submit to complete creating the Contract.
183. Click Submit to complete adding the Provided Contract.
184. Right-click Contracts and select Add Consumed Contract.
185. In the Add Consumed Contract window, use the Contract pulldown to select the `MS-TNT-A/Allow-Web-App` contract.
186. Click Submit to complete adding the Consumed Contract.
187. Optionally, repeat steps 184-186 to add the `common/FP-Allow-Common-Core-Services` Consumed Contract.
188. On the left under EPG App, right-click Subnets and select Create EPG Subnet.
189. For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.18.0.0/16) that was set up for assigning Tenant IP addresses.
190. If this EPG was connected to Core-Services by contract, select only the Shared between VRFs scope. Otherwise, if the tenant SVM management interface will only be accessed from EPGs within the tenant, leave only the Private to VRF Scope selected.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy:

191. Click Submit to complete creating the EPG Subnet.
192. Right-click Application EPGs under 3-Tier-App and select Create Application EPG.
193. Name the EPG DB and leave Intra EPG Isolation set at Unenforced.
194. Use the Bridge Domain pulldown to select MS-TNT-A/BD-Internal. Select the default Monitoring Policy.

Create Application EPG



1. Identity

STEP 1 > Identity

Specify the EPG Identity

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Bridge Domain:

Monitoring Policy:

FHS Trust Control Policy:

Associate to VM Domain Profiles:

Statically Link with Leaves/Paths:

EPG Contract Master:

Application EPGs

195. Click Finish to complete creating the EPG.
196. On the left expand 3-Tier-App, Application EPGs, and EPG DB.
197. Under EPG DB, right-click Domains and select Add VMM Domain Association.
198. Select the APIC-MS-vSwitch VMM Domain Profile.
199. Select Immediate for both the Deploy Immediacy and the Resolution Immediacy. Select the Dynamic VLAN Mode.
200. Click Submit to compete adding the VMM Domain Association.
201. On the left under EPG DB, right-click Contracts and select Add Consumed Contract.

202. In the Add Consumed Contract window, use the Contract pulldown to select the MS-TNT-A/Allow-App-DB contract.
203. Click Submit to complete adding the Consumed Contract.
204. Repeat steps 201-203 to add the common/FP-Allow-Common-Core-Services (optional) and MS-TNT-A/Allow-SVM-MGMT Consumed Contracts.
205. On the left under EPG DB, right-click Subnets and select Create EPG Subnet.
206. For the Default Gateway IP, enter a gateway IP address and mask from a subnet in the Supernet (172.18.0.0/16) that was set up for assigning Tenant IP addresses.
207. If the FP-Allow-Common-Core-Services contract was consumed, select only the Shared between VRFs scope. Otherwise, select only the Private to VRF scope.

Create EPG Subnet



Specify the Subnet Identity

Default Gateway IP:
address/mask

Treat as virtual IP address:

Scope: Private to VRF
 Advertised Externally
 Shared between VRFs

Description:

Subnet Control: ND RA Prefix
 No Default SVI Gateway
 Querier IP

ND RA Prefix policy: ▼

Cancel

Submit

208. Click Submit to complete creating the EPG Subnet.

Configure Tenant Storage

This section describes the procedure for deploying a NetApp storage SVM for a tenant named MS-TNT-A. In this section, VLAN interface ports, a separate IPspace, the tenant SVM, storage protocols within the SVM,

tenant logical interfaces (LIFs), and tenant data volumes are deployed. All procedures in this section are completed using a SSH connection to the storage cluster CLI.

Create Tenant IPspace

To create the tenant IPspace, run the following commands:

```
ipspace create -ipspace MS-TNT-A
ipspace show
```

Create Tenant Broadcast Domains in ONTAP

To create data broadcast domains in the tenant IPspace, run the following commands. If you are not setting up access to iSCSI application data LUNs in this tenant, do not create the iSCSI broadcast domains.

```
broadcast-domain create -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-SMB -mtu 9000
broadcast-domain create -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-SVM-MGMT -mtu 1500
broadcast-domain create -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-iSCSI-A -mtu 9000
broadcast-domain create -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-iSCSI-B -mtu 9000
broadcast-domain show -ipspace MS-TNT-A
```

Create VLAN Interfaces

To create tenant-storage VLAN interfaces, complete the following steps:

1. Create SMB VLAN ports and add them to the data broadcast domain.

```
network port vlan create -node <node01> -vlan-name a0a-<storage-MS-TNT-A-smb-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-MS-TNT-A-smb-vlan-id>

broadcast-domain add-ports -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-SMB -ports <node01>:a0a-
<storage-MS-TNT-A-smb-vlan-id>, <node02>:a0a-<storage-MS-TNT-A-smb-vlan-id>
```

2. Create SVM management VLAN ports and add them to the data broadcast domain.

```
network port vlan create -node <node01> -vlan-name a0a-<storage-MS-TNT-A-svm-mgmt-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-MS-TNT-A-svm-mgmt-vlan-id>

broadcast-domain add-ports -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-SVM-MGMT -ports <node01>:a0a-
<storage-MS-TNT-A-svm-mgmt-vlan-id>, <node02>:a0a-<storage-MS-TNT-A-svm-mgmt-vlan-id>
```

3. Create tenant iSCSI VLAN ports and add them to the data broadcast domain. If you are not setting up access to iSCSI application data LUNs in this tenant, do not create the iSCSI VLAN ports.

```
network port vlan create -node <node01> -vlan-name a0a-<storage-MS-TNT-A-iscsi-A-vlan-id>
network port vlan create -node <node01> -vlan-name a0a-<storage-MS-TNT-A-iscsi-B-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-MS-TNT-A-iscsi-A-vlan-id>
network port vlan create -node <node02> -vlan-name a0a-<storage-MS-TNT-A-iscsi-B-vlan-id>

broadcast-domain add-ports -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-iSCSI-A -ports <node01>:a0a-
<storage-MS-TNT-A-iscsi-A-vlan-id>, <node02>:a0a-<storage-MS-TNT-A-iscsi-A-vlan-id>
broadcast-domain add-ports -ipspace MS-TNT-A -broadcast-domain MS-TNT-A-iSCSI-B -ports <node01>:a0a-
<storage-MS-TNT-A-iscsi-B-vlan-id>, <node02>:a0a-<storage-MS-TNT-A-iscsi-B-vlan-id>
broadcast-domain show -ipspace MS-TNT-A
```

Create Storage Virtual Machine

To create an infrastructure SVM, complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver MS-TNT-A-SVM -rootvolume rootvol -aggregate aggr1_node02 -rootvolume-
security-style ntfs -ip-space MS-TNT-A
```

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

```
vserver remove-protocols -vserver MS-TNT-A-SVM -protocols nfs,ndmp,fcv
```

3. Add the two data aggregates to the MS-TNT-A-SVM aggregate list.

```
vserver modify -vserver MS-TNT-A-SVM -aggr-list aggr1_node01,aggr1_node02
```

Setup SVM Management Access

1. Create the SVM Management interface.

```
network interface create -vserver MS-TNT-A-SVM -lif SVM-MGMT -role data -data-protocol none -home-
node <st-node-02> -home-port a0a-<MS-TNT-A-SVM-MGMT-VLAN> -address <svm-mgmt-ip> -netmask <svm-mgmt-
mask> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

2. Place an entry in your AD DNS server for the MS-TNT-A-SVM management interface.

3. Create SVM default route.

```
network route create -vserver MS-TNT-A-SVM -destination 0.0.0.0/0 -gateway <MS-TNT-A-SVM-Gateway>
network route show
(network route show)
Vserver          Destination      Gateway          Metric
-----
Infra-MS-SVM     0.0.0.0/0       10.1.118.1      20
MS-TNT-A-SVM     0.0.0.0/0       172.18.254.6    20
a02-affa300     0.0.0.0/0       192.168.1.254   20
3 entries were displayed.
```

4. Create snapdrive SVM user and unlock vsadmin SVM user.

```
security login create -user snapdrive -application http -authentication-method password -role vsadmin
-vserver MS-TNT-A-SVM

Please enter a password for user 'snapdrive':
Please enter it again:

security login create -user snapdrive -application ontapi -authentication-method password -role
vsadmin -vserver MS-TNT-A-SVM

security login password -username vsadmin -vserver MS-TNT-A-SVM

Enter a new password:
Enter it again:

security login unlock -username vsadmin -vserver MS-TNT-A-SVM
```

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 MS-TNT-A-SVM -domains <domain_name> -name-servers
<dns_server1_ip>,<dns_server2_ip>
```

The SVM Management EPG was connected to the Core Services EPG by contract, allowing access to the DNS servers. For more than one DNS server, separate the IPs by comma.

2. Create the CIFS service.

```
vserver cifs create -vserver MS-TNT-A-SVM -cifs-server MS-TNT-A-CIFS -domain <domain_name>
```

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 "<domain_name>" domain.

Enter the user name: <domain_admin_user>@<domain_name>

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 MS-TNT-A-SVM -is-referral-enabled true
set -privilege admin
```

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 MS-TNT-A-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver MS-TNT-A-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

2. Create the mirroring relationships.

```
snapmirror create -source-path MS-TNT-A-SVM:rootvol -destination-path MS-TNT-A-SVM:rootvol_m01 -type
LS -schedule 15min
```

```
snapmirror create -source-path MS-TNT-A-SVM:rootvol -destination-path MS-TNT-A-SVM:rootvol_m02 -type
LS -schedule 15min
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set -source-path MS-TNT-A-SVM:rootvol
snapmirror show
```

Create Block Protocol Service(s)

If the deployment is using FCP, create the FCP service on each SVM using the following command. This command also starts the FCP service and sets the worldwide name (WWN) for the SVM.

```
fcv create -vserver MS-TNT-A-SVM
fcv show
```

If the deployment is using iSCSI, create the iSCSI service on each SVM using the following command. This command also starts the iSCSI service and sets the IQN for the SVM.

```
iscsi create -vserver MS-TNT-A-SVM
iscsi show
```



The licenses for FCP and iSCSI must be installed before the services can be started. If the license(s) weren't installed during cluster setup, install them before this step.

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 MS-TNT-A-SVM -common-name MS-TNT-A-SVM -ca MS-TNT-A-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 previous command, use TAB completion to select and delete each default certificate.

- 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 MS-TNT-A-SVM
```

- To obtain the values for the parameters required in step 5 (<cert-ca> and <cert-serial>), run the security certificate show command.
- 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 <vserver-name> -server-enabled true -client-enabled false -ca <cert-ca> -serial <cert-serial> -common-name <cert-common-name>
```

- 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 MS-TN-A-SVM -policyname smb
export-policy rule create -vserver MS-TNT-A-SVM -policyname smb -clientmatch <domain_name> -rorule krb5i,krb5p -rwrule krb5i,krb5p
```

Create NetApp FlexVol Volumes

```
volume create -vserver MS-TNT-A-SVM -volume tnt_a_smb_datastore_1 -aggregate aggr1_node01 -size 500GB -state online -policy smb -security-style ntfs -junction-path /smb_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver MS-TNT-A-SVM -volume tnt_a_smb_datastore_2 -aggregate aggr1_node02 -size 500GB -state online -policy smb -security-style ntfs -junction-path /smb_datastore_2 -space-guarantee none -percent-snapshot-space 0
volume create -vserver MS-TNT-A-SVM -volume tnt_a_iscsi_datastore_1 -aggregate aggr1_node01 -size 1TB -state online -policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0
volume create -vserver MS-TNT-A-MS-SVM -volume tnt_a_iscsi_datastore_2 -aggregate aggr1_node02 -size 1TB -state online -policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0
volume create -vserver MS-TNT-A-SVM -volume witness -aggregate aggr1_node02 -size 5GB -state online -policy default -security-style ntfs -space-guarantee none -percent-snapshot-space 0
snapmirror update-ls-set -source-path MS-TNT-A-SVM:rootvol
```

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 `smb_datastore_1`, `ismb_datastore_2`, `iscsi_datastore_1`, and `iscsi_datastore_2`:

```
efficiency modify -vserver MS-TNT-A-SVM -volume smb_datastore_1 -schedule sun-sat@0
efficiency modify -vserver MS-TNT-A-SVM -volume smb_datastore_2 -schedule sun-sat@0
efficiency modify -vserver MS-TNT-A-SVM -volume iscsi_datastore_1 -schedule sun-sat@0
efficiency modify -vserver MS-TNT-A-SVM -volume iscsi_datastore_2 -schedule sun-sat@0
```

Create SAN LIFs

If using FCP, run the following commands to create four FC LIFs (two on each node):

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

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

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

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

If using iSCSI, run the following commands to create four iSCSI LIFs (two on each node):

```
network interface create -vserver MS-TNT-A-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -
home-node <st-node01> -home-port a0a-<iSCSI-A-VLAN> -address <iscsi_lif01a_ip> -netmask
<iscsi_lif01a_mask> -status-admin up

network interface create -vserver MS-TNT-A-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -
home-node <st-node01> -home-port a0a-<iSCSI-B-VLAN> -address <iscsi_lif01b_ip> -netmask
<iscsi_lif01b_mask> -status-admin up

network interface create -vserver MS-TNT-A-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -
home-node <st-node02> -home-port a0a-<iSCSI-A-VLAN> -address <iscsi_lif02a_ip> -netmask
<iscsi_lif02a_mask> -status-admin up

network interface create -vserver MS-TNT-A-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -
home-node <st-node02> -home-port a0a-<iSCSI-B-VLAN> -address <iscsi_lif02b_ip> -netmask
<iscsi_lif02b_mask> -status-admin up
```

Create SMB LIFs

To create SMB LIFs, run the following commands:

```
network interface create -vserver MS-TNT-A-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 MS-TNT-A-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
```



The two SMB LIF IPs need to be entered in the Domain DNS server with the CIFS Server name (ms-tnt-a-cifs) created above. Create two host records with the different IPs and same host name.

Add Quality of Service (QoS) Policy to Monitor Application Workload

To add a storage QoS policy to monitor both the IOPs and bandwidth delivered from the APP-A-SVM, complete the following steps:

1. Create the QoS policy-group to measure the SVM output without an upper limit.

```
qos policy-group create -policy-group MS-TNT-A -vserver MS-TNT-A-SVM -max-throughput INF
vserver modify -vserver MS-TNT-A-SVM -qos-policy-group MS-TNT-A
```

2. Monitor the QoS policy group output.

```
qos statistics performance show
```

Configure Cisco UCS for the Tenant

This section describes procedures for deploying UCS Servers for a tenant named MS-TNT-A. It is assumed in this FlexPod Deployment that a tenant is most likely an application or group of applications. Because of this assumption, it is assumed that a new set of Hyper-V servers will be setup for the tenant in a separate Hyper-V cluster. It is also assumed in this implementation that the new Hyper-V servers will be booted from the storage Infrastructure SVM, although server boot could be moved into the tenant SVM. In this section, required additions to Cisco UCS will be detailed, including adding a new tenant vNIC template and LAN connectivity policy, creating additional Service Profile Templates, and generating the new Service Profiles for the Hyper-V hosts for the tenant. All procedures in this section are completed using the Cisco UCS Manager HTML 5 User Interface.

Add Tenant Host Management vNIC Template

The tenant Hyper-V host management interfaces must be in the Core Services EPG in order for NetApp SnapDrive (installed on these hosts) to be able to reach the tenant Storage Virtual Machine (SVM). Complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Expand LAN > Policies > root > vNIC Templates.
3. Right-click vNIC Templates and select Create vNIC Template.
4. Name the template TNT-Core-MGMT-A.
5. Leave Fabric A selected and also select Enable Failover.
6. Leave Redundancy Type set to No Redundancy and Target set to Adapter.

7. Select the Updating Template Type.
8. Under VLANs select only the MS-Core-Services VLAN and make MS-Core-Services the Native VLAN.
9. Leave the MTU set at 1500 and select MAC-Pool-A and the Enable-CDP-LLDP Network Control Policy.

Create vNIC Template



Warning

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

Template Type : Initial Template Updating Template

VLANs
VLAN Groups

Advanced Filter Export Print ⚙️

Select	Name	Native VLAN
<input type="checkbox"/>	ACI-System	<input type="radio"/>
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	MS-Cluster	<input type="radio"/>
<input checked="" type="checkbox"/>	MS-Core-Services	<input checked="" type="radio"/>
<input type="checkbox"/>	MS-IB-MGMT	<input type="radio"/>
<input type="checkbox"/>	MS-Infra-iSCSI-A	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU :

MAC Pool : ▼

QoS Policy : ▼

Network Control Policy : ▼

OK
Cancel

10. Click OK twice to complete creating the template.

Create Tenant LAN Connectivity Policy for iSCSI Boot

To configure the necessary LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click the LAN icon on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter `TNT-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 `00-TNT-Core-MGMT` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select `TNT-Core-MGM-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 :

Peer Name :

vNIC Template :

Adapter Performance Profile

Adapter Policy :

[Create vNIC Template](#)

[Create Ethernet Adapter Policy](#)

12. Click the upper Add button to add another vNIC to the policy.
13. In the Create vNIC box, enter `01-Infra-iSCSI-A` as the name of the vNIC.
14. Select the Use vNIC Template checkbox.
15. In the vNIC Template list, select `Infra-iSCSI-A`.
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 another vNIC to the policy.
19. In the Create vNIC box, enter `02-Infra-iSCSI-B` as the name of the vNIC.
20. Select the Use vNIC Template checkbox.
21. In the vNIC Template list, select `Infra-iSCSI-B`.
22. In the Adapter Policy list, select `Windows`.
23. Click OK to add the vNIC to the policy.
24. Click the upper Add button to add another vNIC to the policy.
25. In the Create vNIC box, enter `03-APIC-MS-VS-A` as the name of the vNIC.
26. Select the Use vNIC Template checkbox.
27. In the vNIC Template list, select `APIC-MS-VS-A`.
28. In the Adapter Policy list, select `Windows`.
29. Click OK to add the vNIC to the policy.
30. Click the upper Add button to add another vNIC to the policy.
31. In the Create vNIC box, enter `04-APIC-MS-VS-B` as the name of the vNIC.
32. Select the Use vNIC Template checkbox.
33. In the vNIC Template list, select `APIC-MS-VS-B`.
34. In the Adapter Policy list, select `Windows`.
35. Click OK to add the vNIC to the policy.
36. Expand the Add iSCSI vNICs section.
37. Click the lower Add button to add an iSCSI boot vNIC to the policy.
38. In the Create iSCSI vNIC box, enter `iSCSI-Boot-A` as the name of the vNIC.
39. Select `01-Infra-iSCSI-A` for the Overlay vNIC.
40. Select the default iSCSI Adapter Policy.
41. `MS-Infra-iSCSI-A (native)` should be selected as the VLAN.

42. Do not select anything for MAC Address Assignment.

Create iSCSI vNIC



Name :

Overlay vNIC :

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

VLAN :

iSCSI MAC Address

MAC Address Assignment:

[Create MAC Pool](#)

43. Click OK to add the vNIC to the policy.

44. Click the lower Add button to add an iSCSI boot vNIC to the policy.

45. In the Create iSCSI vNIC box, enter `iSCSI-Boot-B` as the name of the vNIC.

46. Select `02-Infra-iSCSI-B` for the Overlay vNIC.

47. Select the default iSCSI Adapter Policy.

48. `MS-Infra-iSCSI-B (native)` should be selected as the VLAN.

49. Do not select anything for MAC Address Assignment.

50. Click OK to add the vNIC to the policy.

Create LAN Connectivity Policy



Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 04-APIC-MS-VS-B	Derived	
vNIC 03-APIC-MS-VS-A	Derived	
vNIC 02-Infra-iSCSI-B	Derived	
vNIC 01-Infra-iSCSI-A	Derived	
vNIC 00-TNT-Core-MGMT	Derived	

Delete Add Modify

⊖ Add iSCSI vNICs

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address
iSCSI vNIC iSCSI-Boot-B	02-Infra-iSCSI-B	default	Derived
iSCSI vNIC iSCSI-Boot-A	01-Infra-iSCSI-A	default	Derived

Add Delete Modify

OK

Cancel

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

Create Tenant Service Profile Template

It is recommended to create two new Service Profile Templates for the servers running the applications in the new tenant. These templates can be created by cloning the existing Service Profile Templates and modifying them by changing the LAN Connectivity Policy.

Add New Application-Specific Server Pool

Since new Service Profile Templates for the servers running the applications in the new tenant are being created, a new tenant-specific server pool can be created and mapped in the new Service Profile Templates. Create this pool and map it in the new Service Profile Templates.

Create New Service Profiles for Tenant Servers

Using the cloned and modified Tenant Service Profile Templates, create Service Profiles associated to servers for the new tenant.

Gather Necessary Information

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

Table 13 iSCSI LIFs for iSCSI IQN.

Vserver	iSCSI Target IQN
Infra-MS-SVM	
MS-TNT-A-SVM	



To gather the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface.

Table 14 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN
Hyper-V-TNT-A-Host-01	
Hyper-V-TNT-A-Host--02	



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

Configure Storage SAN Boot for the Tenant

This section details the steps for setting up SAN boot for the tenant ESXi Host servers.

Hyper-V Boot LUN in Infra-MS-SVM for First Tenant Host

1. From the cluster management node SSH connection, enter the following:

```
lun create -vserver Infra-MS-SVM -volume HV_boot -lun TNT-HV2016-Gold -size 200GB -ostype windows_2008 -space-reserve disabled
lun show
```

Clustered Data ONTAP iSCSI Boot Storage Setup

Create igroups

1. From the cluster management node SSH connection, enter the following:

```
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-TNT-A-Host-01 -protocol iscsi -ostype windows -
initiator <hyper-v-tnt-a-host-01-iqn>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-TNT-A-Host-02 -protocol iscsi -ostype windows -
initiator < hyper-v-tnt-a-host-02-iqn>
igroup create -vserver MS-TNT-A-SVM -igroup Hyper-V-TNT-A-Host-All -protocol iscsi -ostype windows -
initiator <hyper-v-tnt-a-host-01-iqn>,< hyper-v-tnt-a-host-02-iqn>
```



Use the values listed in Table 8 and Table 9 for the IQN information.



To view the igroups just created, type `igroup show`.

Map Boot LUN to igroup

1. From the storage cluster management SSH connection, enter the following:

```
lun map -vserver Infra-MS-SVM -volume HV_boot -lun TNT-HV2016-Gold -igroup Hyper-V-TNT-A-Host-01 -
lun-id 0

lun show -m
```

Microsoft Hyper-V Server Deployment Procedure for Tenant Hosts

Setting Up Microsoft Hyper-V Server 2016

This section provides detailed instructions for installing Microsoft Hyper-V Server 2016 for tenant application hosts in a FlexPod environment. After the procedures are completed, two booted Hyper-V Server 2016 hosts will be provisioned.

Several methods exist for installing Microsoft Hyper-V 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:

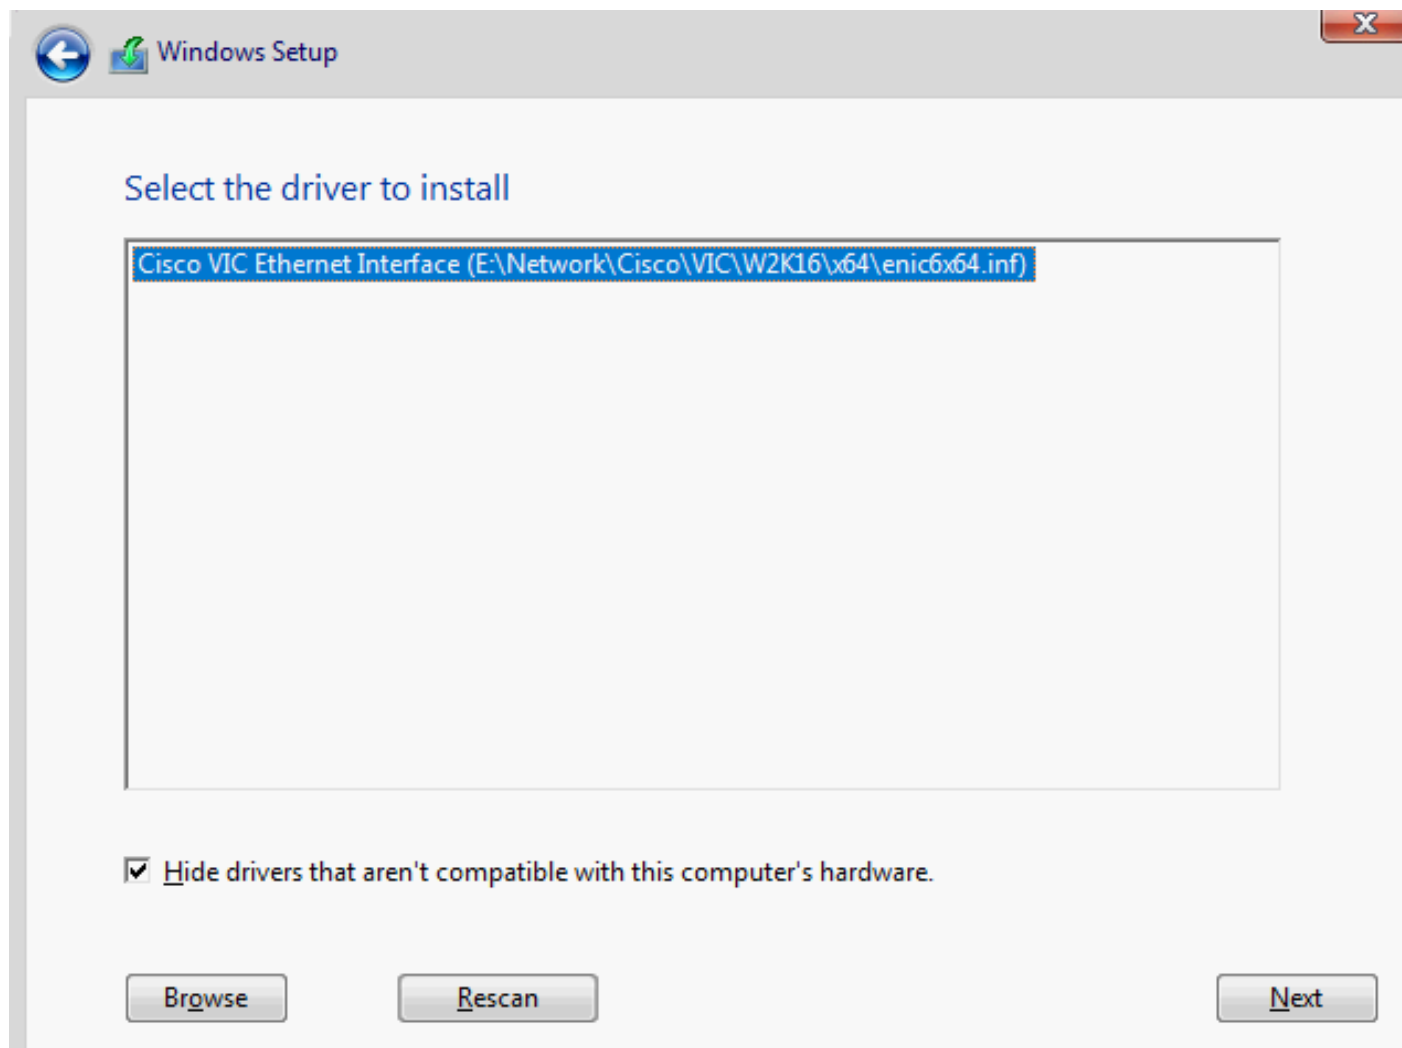
1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, 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 the Servers icon on the left.
7. Select Servers > Service Profiles > root > `Hyper-V-TNT-A-Host-01`.
8. **On the right under the General tab, select the “>>” icon to the right of “KVM Console”.**
9. Follow the prompts to launch the Java KVM console.
10. From the KVM Console, under the Virtual Media tab, select Activate Virtual Devices. Follow the prompts and select Apply.
11. From the KVM Console, under the Virtual Media tab, select Map CD/DVD.
12. Click Browse.
13. Browse to the Windows Hyper-V 2016 installation ISO image file and click Open.
14. Map the image that you just added by selecting Map Device.
15. To boot the server, click the Boot Server icon above the KVM Console tab.

16. Click OK then OK again to boot the server.

Install Hyper-V Server 2016

To install Hyper-V Server 2016 to the first management host, complete the following steps:

1. On boot, the machine detects the presence of the Windows installation media.
2. After the installer has finished loading, select the relevant region information and click Next.
3. Click Install now.
4. After reviewing the EULA, accept the license terms and click Next.
5. Select Custom: Install the newer version of Hyper-V Server only (advanced).
6. In the Windows Setup window, select Load driver.
7. Under Virtual Media, select the Hyper-V Server 2016 item to unmap it. Click Yes to complete the un-mapping.
8. Under Virtual Media, select map CD/DVD.
9. Click Browse and browse to the ucs-bxxx-drivers-windows.3.2.1 iso. Select this iso and click Open. Click Map Device to map this iso.
10. In the Load driver window, click Browse.
11. Browse to the CD Drive and expand Network > Cisco > VIC > W2K16. Select x64 under W2K16. Click OK.
12. Back in the Windows Setup window, make sure Cisco VIC Ethernet Interface is selected and click Next.



13. If booting with FC, also load the Cisco VIC fnic storage driver.
14. One disk drive should now appear in the Windows Setup window. In the Virtual Media menu, unmap the ucs-bxxx-drivers-windows.3.2.1 iso and remap the Windows Server 2016 installation iso.
15. In the Windows Setup window, click Refresh. Make sure the 200 GB drive is selected and click Next.
16. When Windows is finished installing, press Enter and enter and confirm an administrator password. Press Enter.
17. Under Virtual Media, unmap the Windows Server 2016 Installation iso.

Host Renaming and Join to Domain

1. In the Command Prompt window, type powershell to open Windows Powershell.
2. Type Get-NetAdapter to get the Cisco VIC Ethernet Interface # for the 00-TNT-Core-MGMT interface.

```

Administrator: C:\Windows\system32\cmd.exe - powershell
C:\Users\Administrator>powershell
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation. All rights reserved.

PS C:\Users\Administrator> Get-NetAdapter

Name                               InterfaceDescription          ifIndex Status      MacAddress          LinkSpeed
----                               -
03-APIC-MS-VS-A                   Cisco VIC Ethernet Interface #2 11 Up         00-25-B5-A2-0A-08   40 Gbps
01-Infra-iSCSI-A                   Cisco VIC Ethernet Interface     12 Up         00-25-B5-A2-0A-07   40 Gbps
04-APIC-MS-VS-B                   Cisco VIC Ethernet Interface #5 10 Up         00-25-B5-A2-0B-07   40 Gbps
02-Infra-iSCSI-B                   Cisco VIC Ethernet Interface #4  8 Up         00-25-B5-A2-0B-06   40 Gbps
00-TNT-Core-MGMT                   Cisco VIC Ethernet Interface #3  7 Up         00-25-B5-A2-0A-0F   40 Gbps

PS C:\Users\Administrator>
    
```

3. Use Server Configuration to set a Computer Name (Hyper-V-2016-Gold, a restart will be required), set Network Settings in the IB-MGMT/Core Services subnet for the 00-TNT-Core-MGMT interface, set the Timezone, enable Remote Desktop, and join the machine to the AD domain (another restart will be required).

Install NetApp Windows Unified Host Utilities

Download and install NetApp Windows Unified Host Utilities. This section provides the steps to download and install the host utilities. In this and further procedures where software is installed on Hyper-V Server 2016 (Server Core), it is best to have an SMB server available where software can be stored. The Hyper-V Server machine can then connect to this SMB server and load software from there.

1. Download the x64 version of the NetApp host utilities v7.0 for Windows from the link below and place in the SMB share:

https://mysupport.netapp.com/NOW/download/software/sanhost_win/7.0

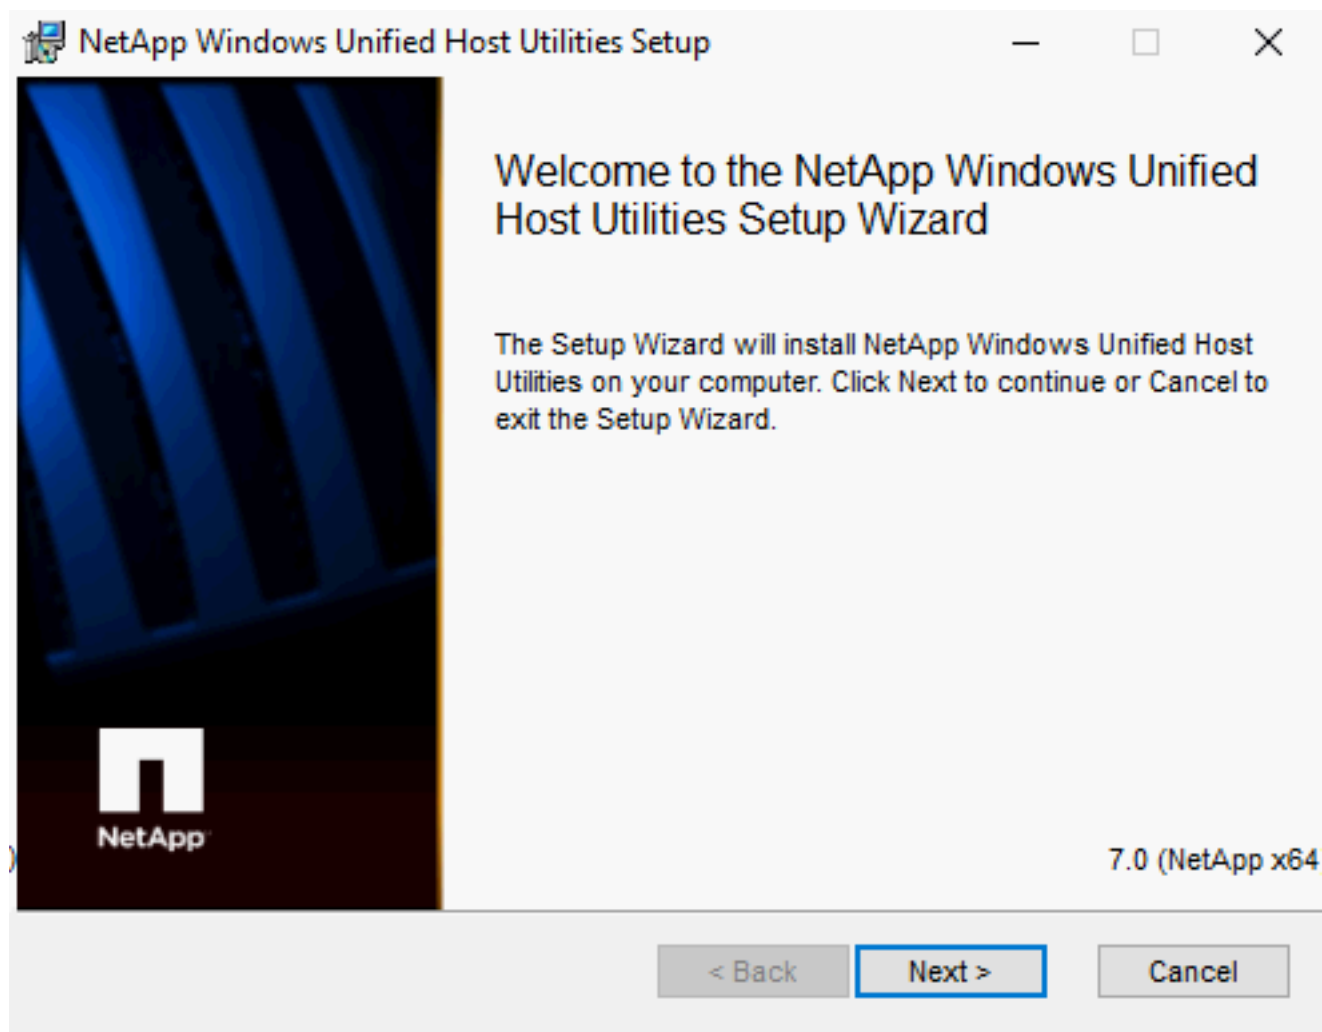
2. From a command prompt in the KVM console, use the “net use” command to connect to the SMB share.

```


net use S: \\<sharepath> <password> /USER:<domain\username>

net use
    
```

3. Run the NetApp Windows Unified Host Utilizes .msi file from the SMB share. The NetApp Windows Unified Host Utilities setup wizard is launched. Click OK on the hotfix warning and then click Next.



4. Click the checkbox to accept the license agreement and click Next.
5. Select “Yes, install support for Multipath IO” and click Next.



6. Accept the default destination folder and click Next.
7. Click Install, OK, and Finish to complete the installation of host utilities.
8. Click No to not restart the computer.
9. Shut down the server.

Set Up Multipathing and iSCSI

1. In Cisco UCS Manager, select the Hyper-V-TNT-A-Host-01 Service Profile.
2. Under the General tab on the right, select Bind to a Template.
3. Select the Hyper-V-TNT-iSCSI template which contains 4 SAN paths for multipath I/O. Click OK, then Yes and OK to complete binding to the new Service Profile Template.
4. Follow steps 1-3 to bind the Hyper-V-TNT-A-Host-02 Service Profile to the Hyper-V-TNT-iSCSI template.

5. When the host with Service Profile Hyper-V-TNT-A-Host-01 returns to the Power Off state, go to the KVM console, and select Boot Server to boot the host. Click OK and OK again to proceed. When the server boots, 4 iSCSI SAN paths should be seen.

6. Log into the server as Administrator and open Powershell.

7. Install the .NET 3.5 Windows feature.

```
Install-WindowsFeature Net-Framework-Core
```

8. Enable Server Remote Management on the Windows Firewall.

```
netsh firewall set service type=remoteadmin mode=enable  
netsh advfirewall firewall set rule group="remote administration" new enable=yes
```

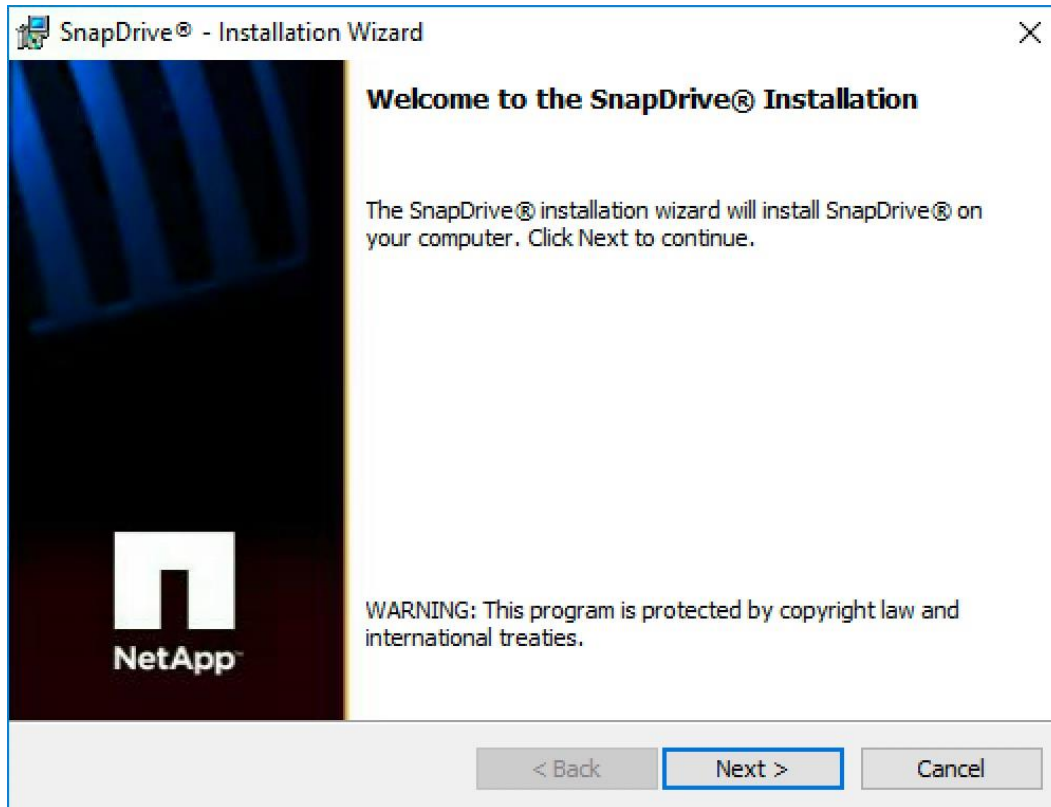
9. Start the Microsoft iSCSI Initiator Service and set automatic startup of the service.

```
Start-Service msiscsi  
Set-Service msiscsi -startuptype "automatic"
```

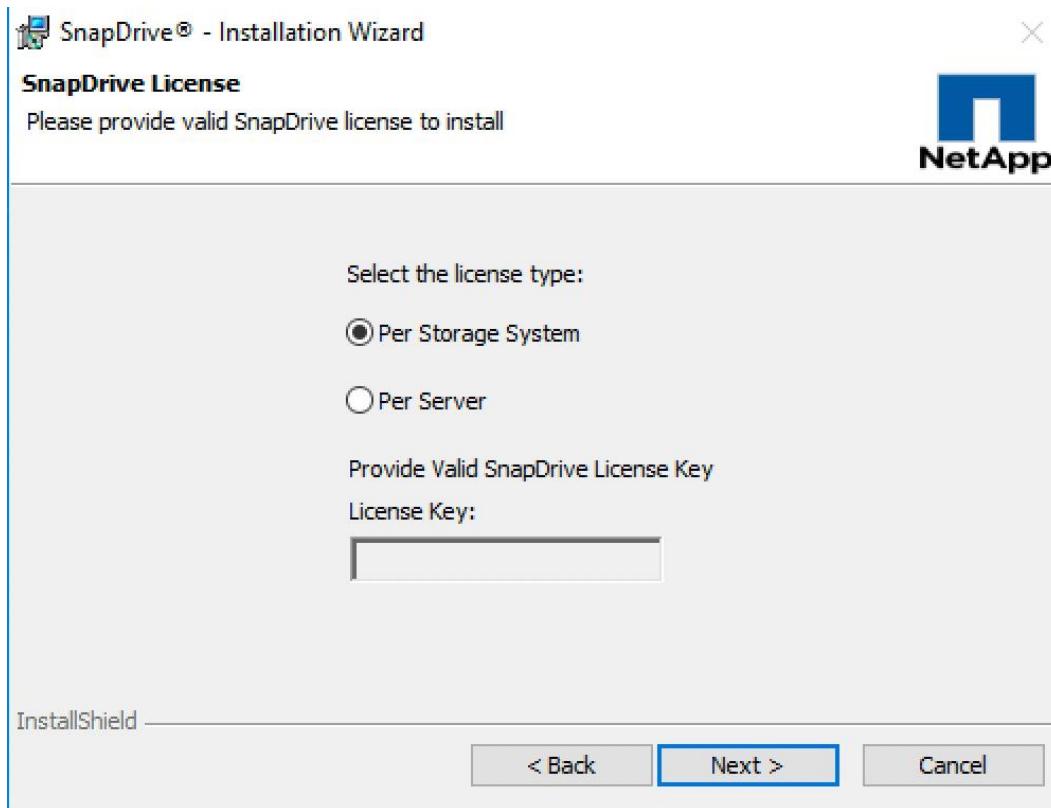
Install SnapDrive for Windows

1. Log out and log back in to Windows Hyper-V Server 2016 with a Domain Admin user account.

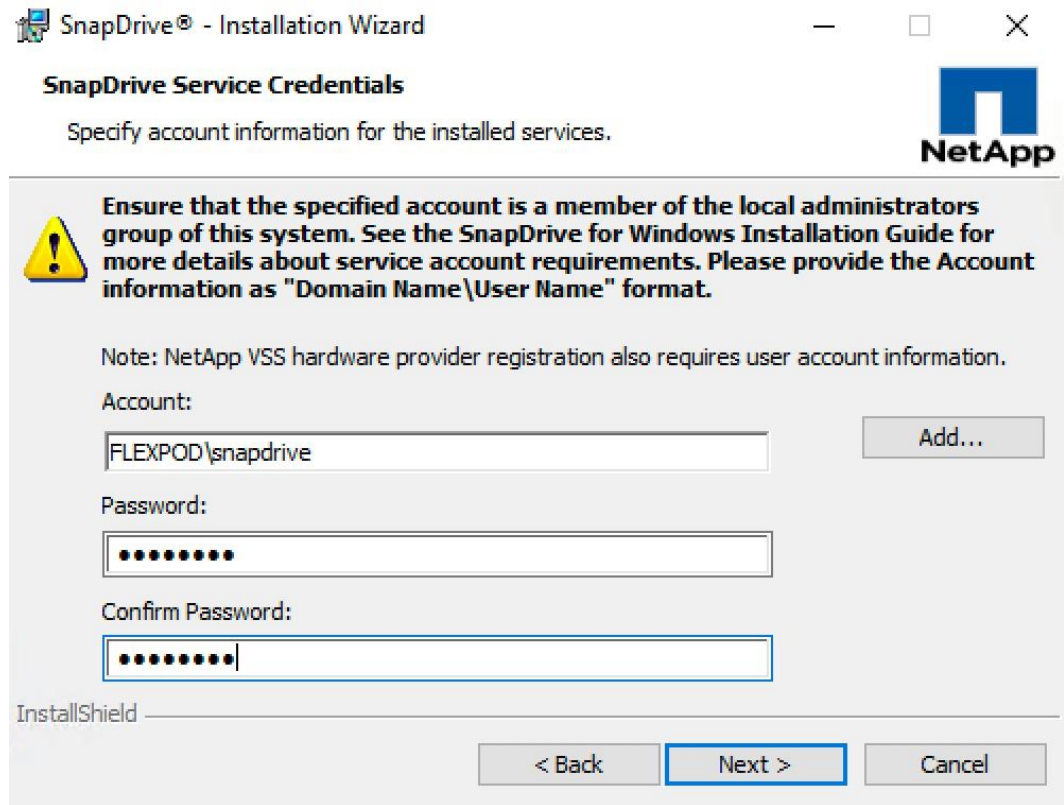
2. From the SMB share, launch the SnapDrive for Windows installer, and then follow the instructions in the wizard.




3. On the SnapDrive License page, select the appropriate license type and click Next.




4. On the Customer Information page, enter the appropriate information and click Next.
5. On the Destination Folder page, enter the appropriate destination or accept the default. Click Next.
6. On the Set Firewall Rules page, select for Enable SnapDrive to Communicate Through the Windows Firewall option and click Next.
7. On the SnapDrive Service Credentials page, enter the account and password information of the account created earlier that is a member of the local administrators group.



8. On the SnapDrive Web Service Configuration page, accept the default port numbers and click Next.

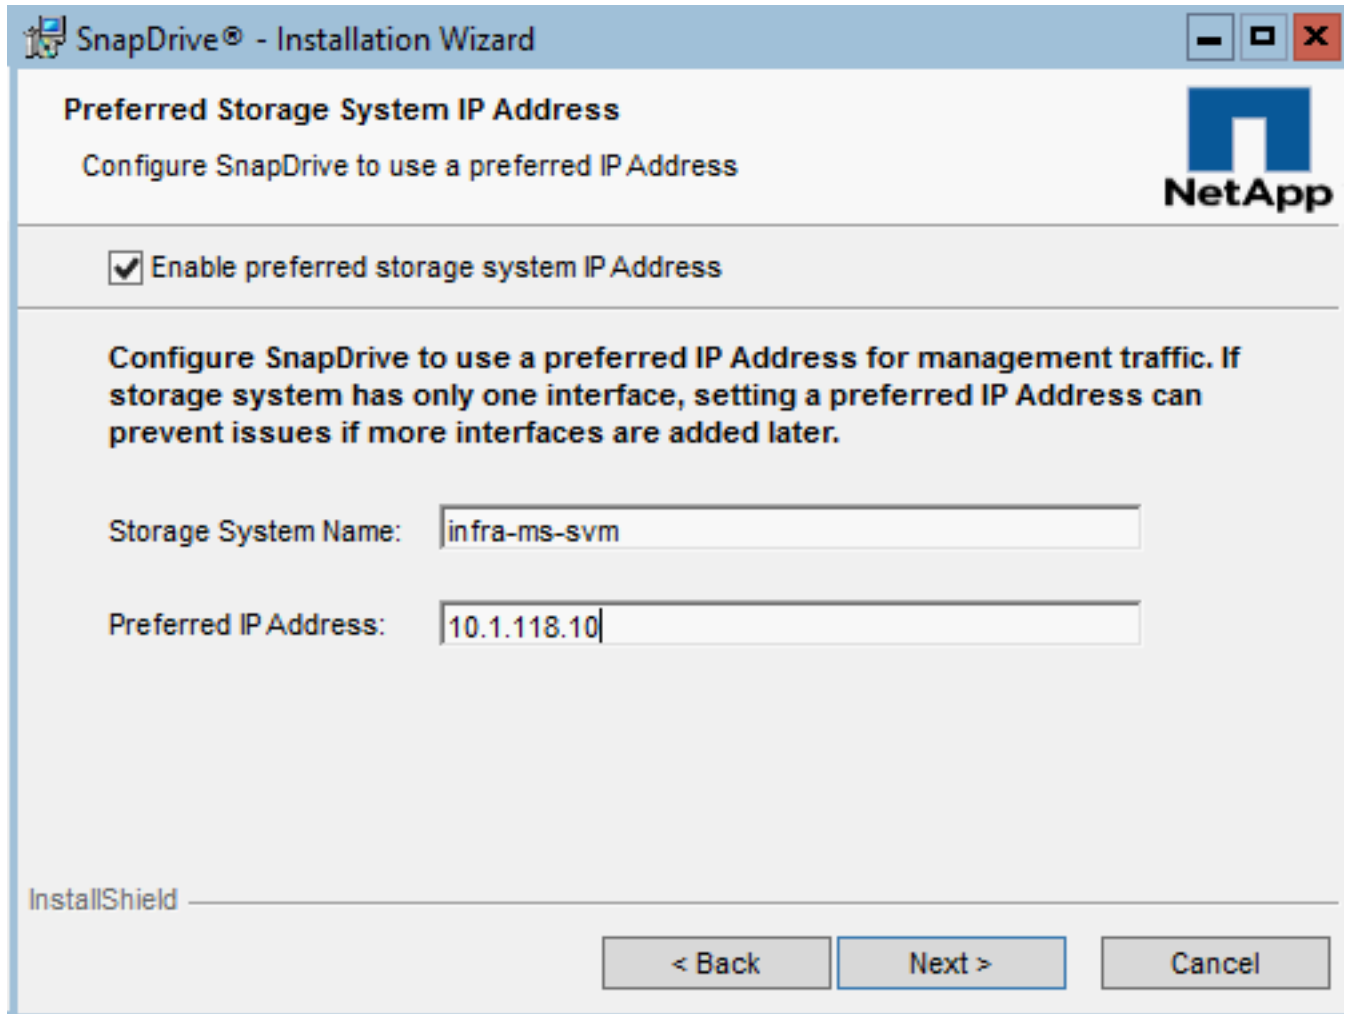
 If you want to change the port numbers, you should also change the port numbers for the other Snap-Drive hosts.



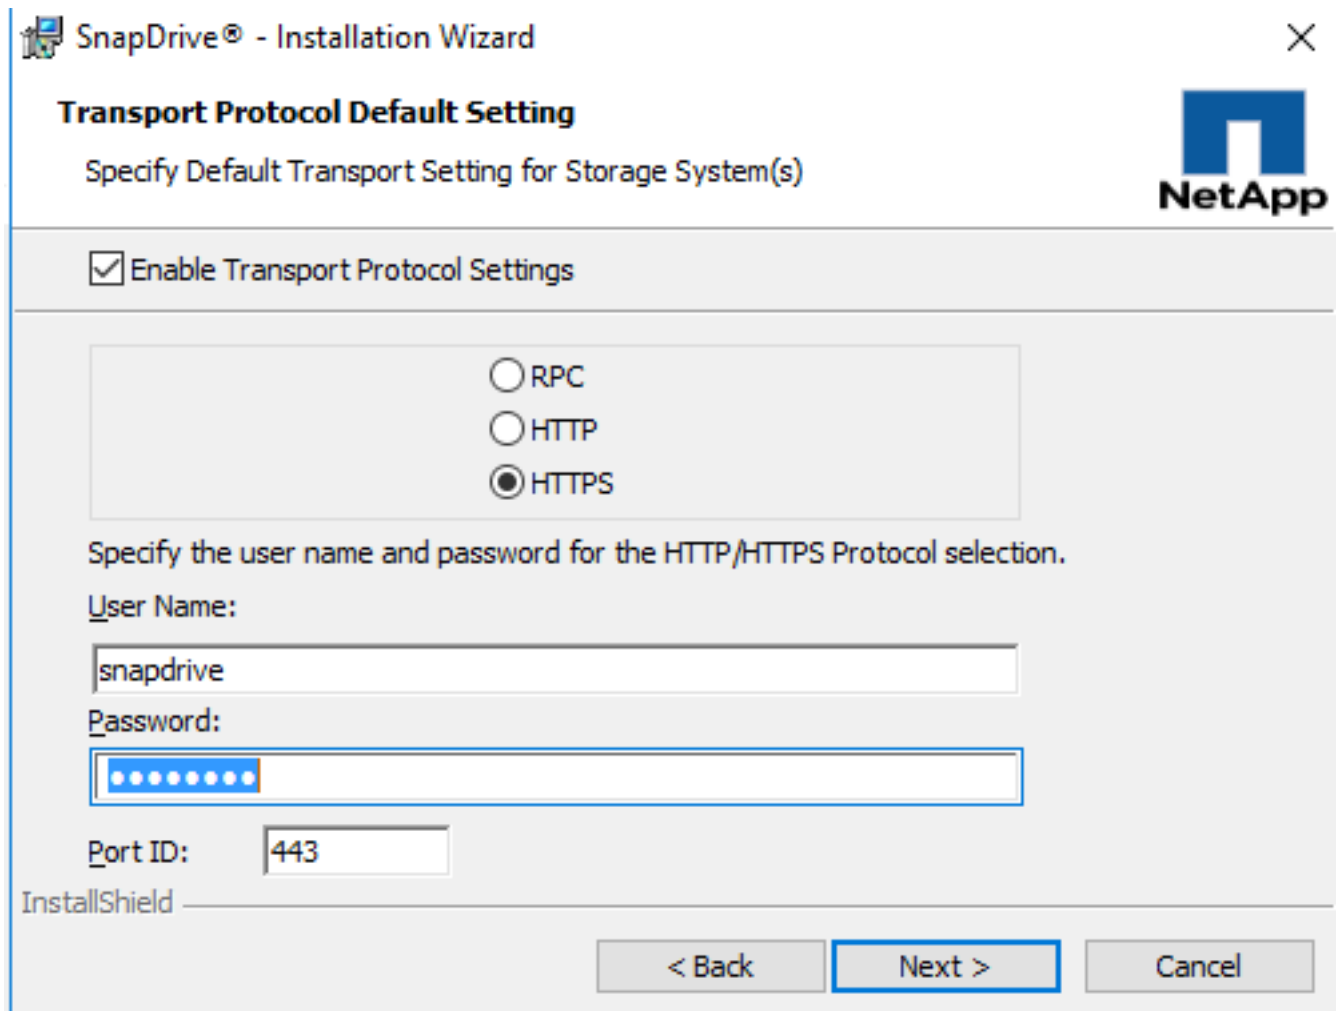
9. On the Preferred IP Address screen, identify the IP address you want to use to communicate with the storage system and click Next.




You should configure the preferred IP address, because doing this improves performance and scalability.



10. On the Transport Protocol Default Setting page, enable the storage protocol settings and click Next. RPC is not supported in clustered Data ONTAP.

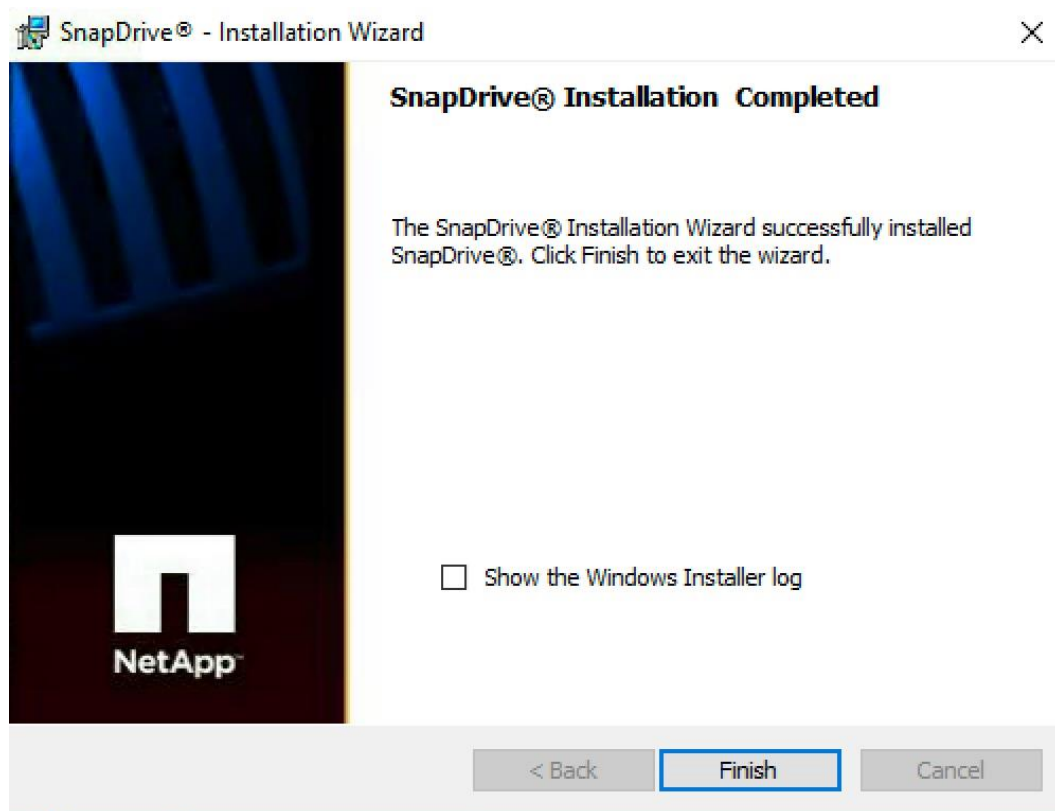


11. On the Unified Manager Configuration Screen, click Next.

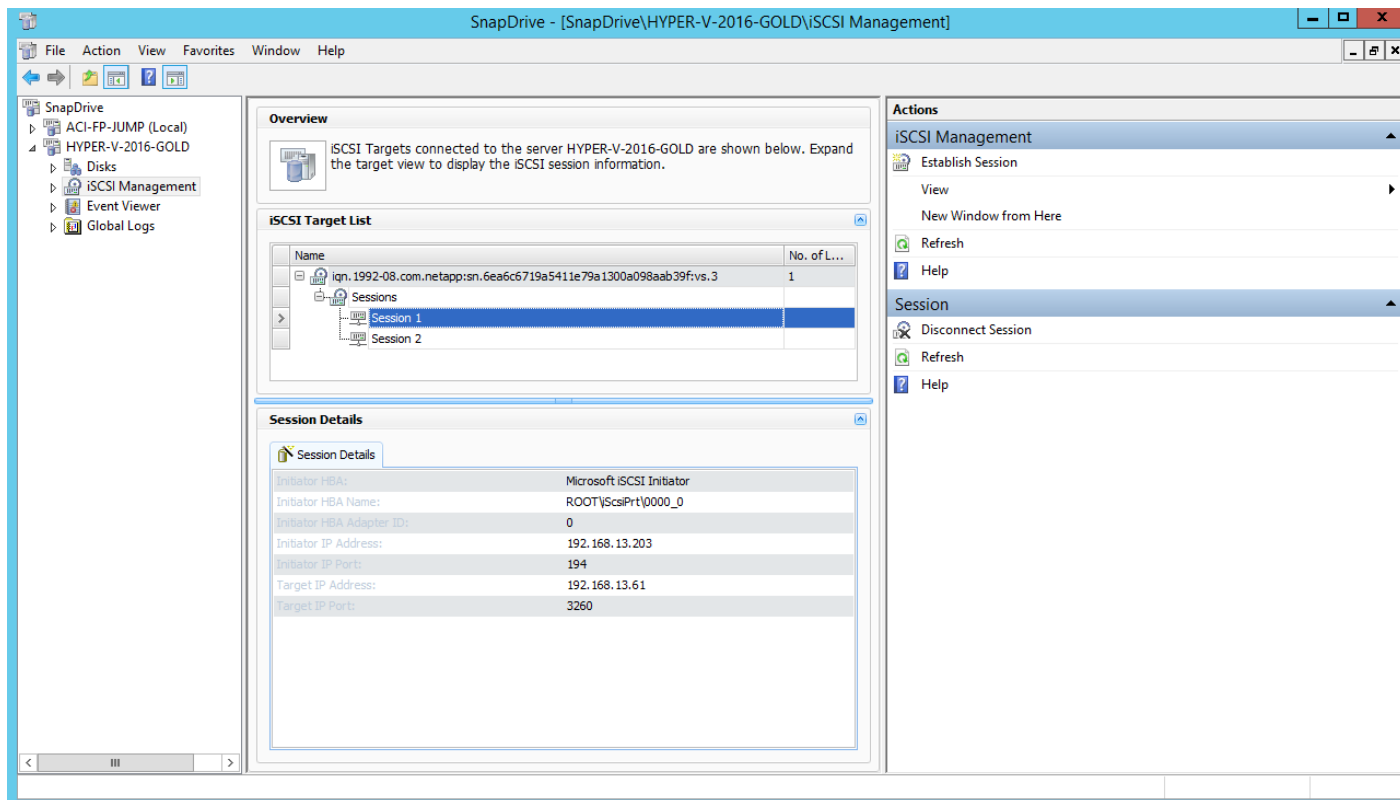
 OnCommand Unified Manager Core Package data protection capabilities are available only in 7- Mode environments.

12. On the Ready to Install page, click Install.

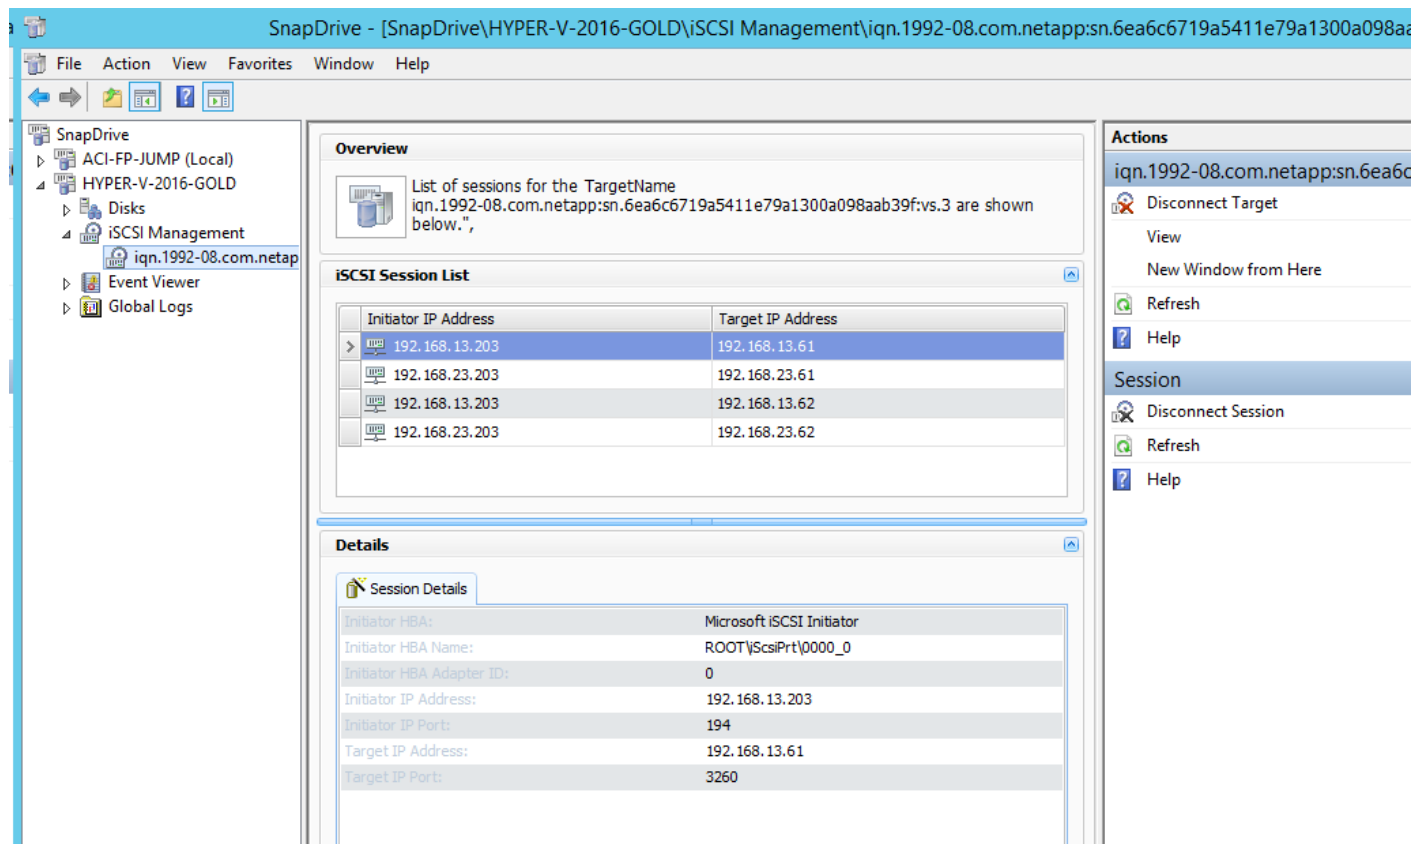
13. When the SnapDrive installation is complete, click Finish.



14. On another machine that has the Windows GUI and SnapDrive installed, launch the SnapDrive snap-in as the snapdrive domain user and add the Hyper-V-2016-Gold Server.
15. Select iSCSI Management and expand to show sessions. Two sessions should be shown.




16. Use “Establish Session” to add sessions for the two missing iSCSI LIFs. Four sessions should now be shown.



Configure Server for Cloning

To configure the server for cloning, complete the followings steps:

1. Install all available Windows Updates on the server.
2. Go to Server Configuration
3. Enter 13 to shut down the server. Click Yes to complete shutdown.

 The boot LUN cloning procedure used in this document will only work when the clones are applied to the same server hardware. If the clone source image was created on a Cisco UCS B200 M4, and you want to apply the image to a Cisco UCS C220 M4, you will need to follow the steps above to install Windows on the Cisco UCS C220 M4.

Clone and Remap Server LUNs for Sysprep Image

To clone and remap server LUNS for the sysprep image, complete the following steps:

1. In the storage cluster interface, unmap the TNT-HV2016-Gold LUN.

```
lun unmap -path /vol/HV_boot/TNT-HV2016-Gold -igroup Hyper-V-TNT-A-Host-01
```

2. Make a clone of the MGMT-Win2016-Gold LUN for the Sysprep clone.

```
clone start -vserver Infra-MS-SVM -source-path /vol/HV_boot/TNT-HV2016-Gold -destination-path /vol/HV_boot/TNT-HV2016-Gold-Sysprep
```

3. Map the Sysprep clone boot LUN to the first Hyper-V tenant host.

```
lun map -path /vol/HV_boot/TNT-HV2016-Gold-Sysprep -igroup Hyper-V-TNT-A-Host-01 -lun-id 0
```

Boot and Set Up Sysprep Clone

To boot and set up the sysprep clone, complete the following steps:

1. Back in the UCS KVM Console for Hyper-V-TNT-A-Host-01, click Boot Server then OK two times to boot the Sysprep Clone LUN.
2. Once the server boots up, log in as a Domain Admin.
3. In the Windows command prompt enter `C:\windows\System32\Sysprep\sysprep /generalize /oobe /shutdown` **to reset the machine's security id. The server will shut down.**

Clone and Remap Server LUNs for Production Image

To clone and remap the server LUNs for the production image, complete the following steps:

1. In the storage cluster interface, unmap the TNT-HV2016-Gold-Sysprep LUN.

```
lun unmap -vserver Infra-MS-SVM -path /vol/HV_boot/MGMT-Win2016-Gold-Sysprep -igroup Hyper-V-MGMT-01
```

2. Make two clones of the TNT-HV2016-Gold-Sysprep LUN for the Hyper-V tenant hosts.

```
clone start -vserver Infra-MS-SVM -source-path /vol/HV_boot/TNT-HV2016-Gold-Sysprep -destination-path /vol/HV_boot/Hyper-V-TNT-A-Host-01
clone start -vserver Infra-MS-SVM -source-path /vol/HV_boot/TNT-HV2016-Gold-Sysprep -destination-path /vol/HV_boot/Hyper-V-TNT-A-Host-02
```

3. Map the Hyper-V-MGMT LUNs to the hosts.

```
lun map -path /vol/HV_boot/Hyper-V-TNT-A-Host-01 -vserver Infra-MS-SVM -igroup Hyper-V-TNT-A-Host-01 -lun-id 0
lun map -path /vol/HV_boot/Hyper-V-TNT-A-Host-02 -vserver Infra-MS-SVM -igroup Hyper-V-TNT-A-Host-02 -lun-id 0
```

Boot and Set Up Clones

To boot and set up clones, complete the following steps:

1. Back in the UCS KVM Console for Hyper-V-TNT-A-Host-01, click Boot Server then OK two times to boot Hyper-V-TNT-A-Host-01.
2. Once the server boots up, login and in the Command Prompt window, type powershell to open Windows Powershell.
3. Type `Get-NetAdapter` to get the Cisco VIC Ethernet Interface # for the 00-TNT-Core-MGMT interface.

4. Use Server Configuration to set Network Settings in the IB-MGMT/Core Services subnet for the first network interface, set the Timezone, set a Computer Name (a restart will be required), and add the host to the Windows Domain (another restart will be required). Then Download and Install Updates.
5. In the command prompt window add a persistent route to the tenant supernet to allow access to the tenant SVM management interface.

```
route -p ADD 172.18.0.0 MASK 255.255.0.0 <core-services-gateway>
```

6. On another machine that has the Windows GUI and SnapDrive installed, launch the SnapDrive snap-in as the snapdrive domain user and add the Hyper-V-TNT-A-Host-01 Server. Look at the iSCSI Sessions and add any missing sessions to get to 4 unique sessions.
7. Open a second UCS KVM console and repeat steps 1-6 for the Hyper-V-TNT-A-Host-02 host.

Deploying and Managing the Tenant Hyper-V Cluster Using System Center 2016 VMM

This section will focus only on configuring the Networking, Storage and Servers in VMM to deploy and manage Hyper-V failover clusters.



You must have System Center 2016 VMM running in your environment.

Fabric – Servers – I

This section covers:

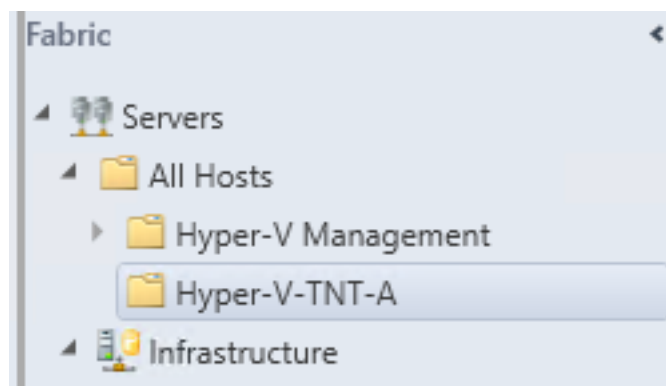
- Create Host Groups
- Add Windows Hosts to the Host Group

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.

To create a host group structure in Virtual Machine Manager (VMM) that aligns to your organizational needs, complete the following steps:

1. To create a host group structure, open the Fabric workspace.
2. In the Fabric pane, expand Servers, and then do either of the following:
3. Right-click All Hosts, and then click Create Host Group.
4. 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.
5. Type a new name and then press ENTER.
6. Repeat the steps in this procedure to create the rest of the host group structure.



Add Hosts to the Host Group


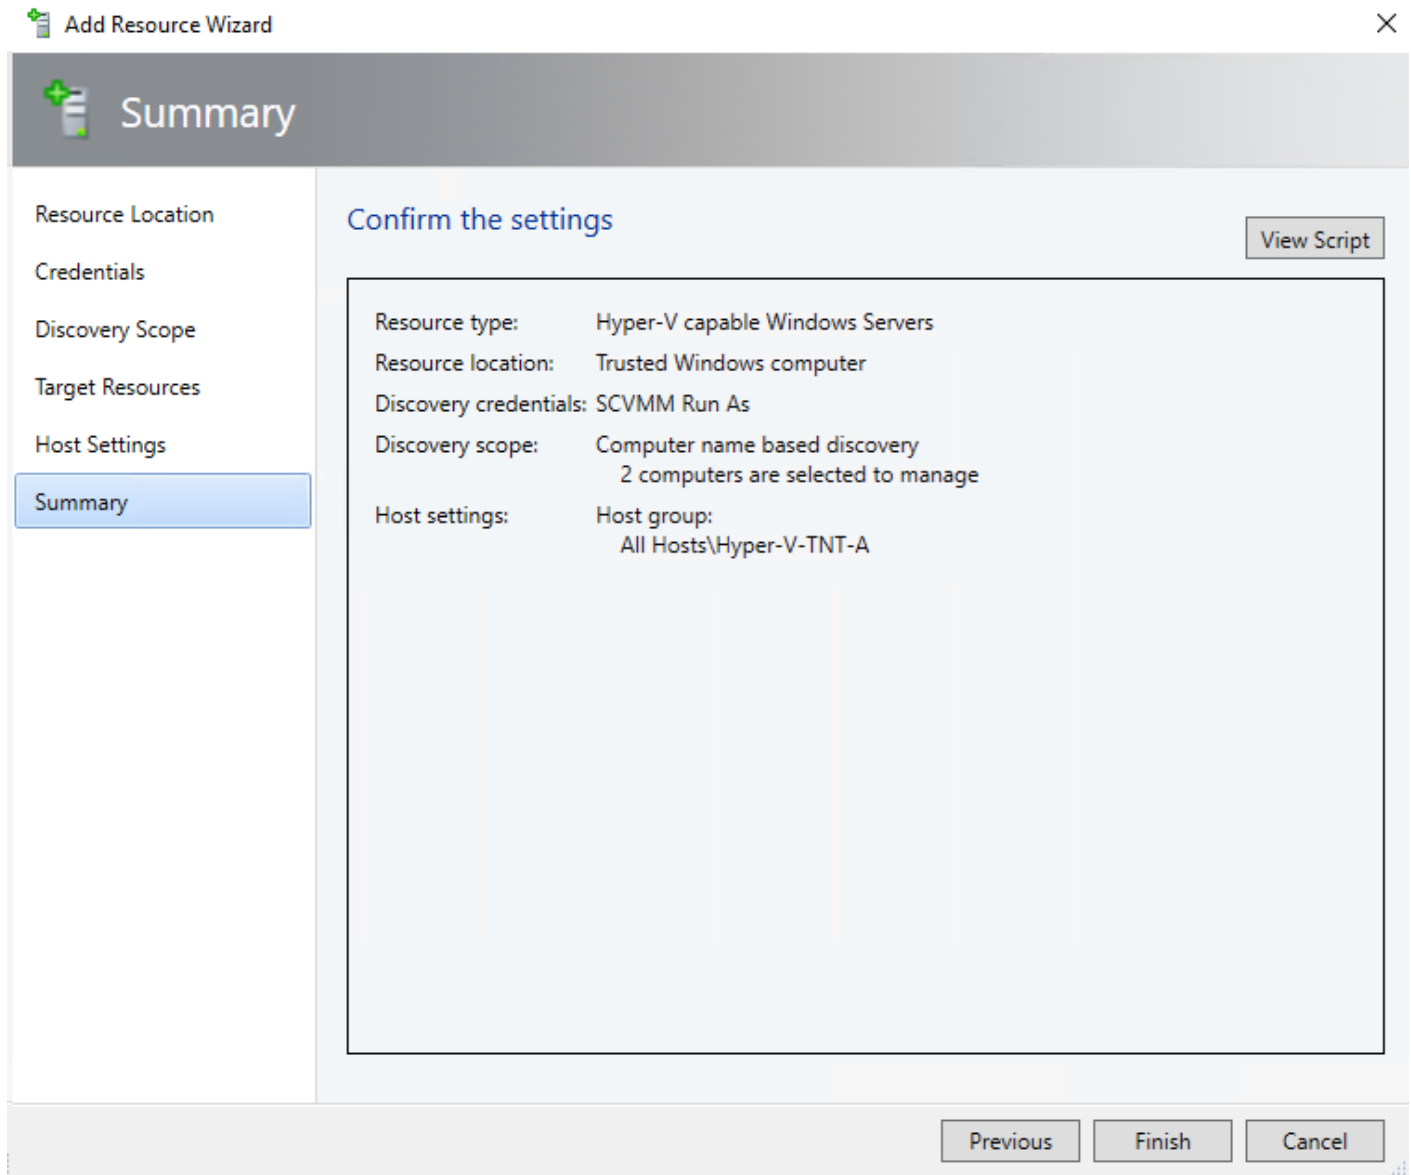
To add the Hyper-V hosts to Virtual Machine Manager, complete the following steps:

1. Open the Fabric workspace.
2. Select a host group, and On the Home tab, 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 SCVMM 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 need to be the part of the Hyper-V cluster. Click Next.



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. Click Next.
8. On the Summary page, confirm the settings, and then click Finish.



9. Restart both Tenant Hosts.

Fabric – Networking – Install APIC Hyper-V Agent and Add Host to SCVMM Virtual Switch

To install APIC Hyper-V and add a host to SCVMM virtual switch, complete the following steps:

1. Open Virtual Machine Manager and select the first Hyper-V Tenant host. Right-click the first Hyper-V Tenant host and select Start Maintenance Mode.
2. **Leave** "Place all running virtual machines into a saved state" **selected** and click Finish.
3. From the Start Menu, open Remote Desktop Connection and enter the hostname or IP address of the first Hyper-V Tenant host. Select Show options. Under the Local Resources tab, make sure the Clipboard and Drives are selected under Local devices and resources. Click Connect and Connect again. The SCVMM Service Account should be selected. Enter the password for this account and click OK.
4. Copy the APIC Hyper-V Agent.msi file from the SCVMM Service Account Desktop and APIC Certificate on the SCVMM VM to the local host.

```
cd c:\temp
```

```
copy "\\TSCLIENT\C\Users\scvmm\Desktop\aci-msft-pkg-3.0.1k\APIC Hyper-V Agent.msi" .\
```

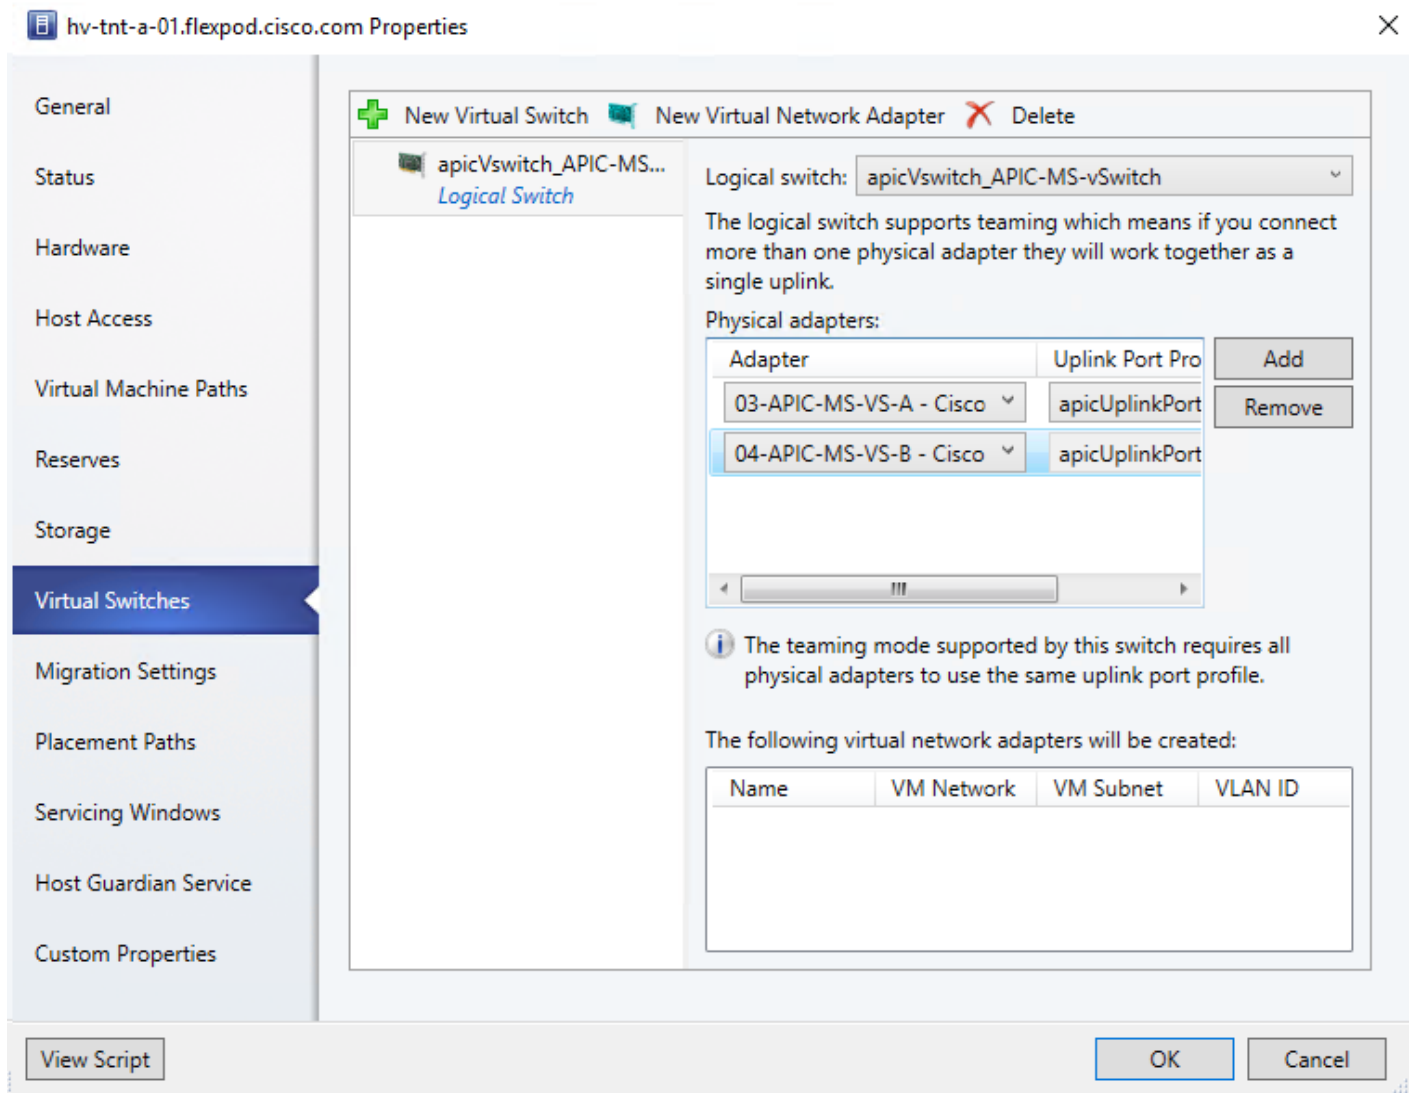
```
copy "\\TSCLIENT\C\Program Files (x86)\ApicVMMService\OpflexAgent.pfx" .\
```

5. Install the APIC Hyper-V Agent.

```
msiexec.exe /I "APIC Hyper-V Agent.msi" /log "C:\InstallLog.txt"
```

6. Click the checkbox to accept the terms in the License Agreement and click Install.
7. Click Finish to complete the installation.
8. Install the OpflexAgent certificate in the Personal Certificate Store.

```
certutil -importPFX My .\OpflexAgent.pfx
```
9. Enter the password provided when the certificate was created.
10. Log out of the Remote Desktop Session.
11. Open Virtual Machine Manager and Stop Maintenance Mode on the Hyper-V Management host.
12. Repeat steps 1-11 to install the APIC Hyper-V Agent and opflex certificate on the second Hyper-V Tenant host.
13. In Virtual Machine Manager, select the VMs and Services workspace and expand All Hosts > Hyper-V TNT-A.
14. Right-click the first Hyper-V Tenant host and choose Properties.
15. **Select Virtual Switches on the left, then select "+ New Virtual Switch". Choose New Logical Switch.**
16. The apicVswitch should be selected. Under Physical Adapters, select 03-APIC-MS-VS-A and add 04-APIC-MS-VS-B. The Uplink Port Profile should be populated automatically.



17. At the top, select “New Virtual Network Adapter”.
18. Name the virtual network adapter <hostname>-vtep. The ACI System VLAN should already be filled in.
19. Click OK and Yes to add the virtual switch to the host.
20. Repeat this process to add the APIC-controlled virtual switch to the second Hyper-V Tenant host.
21. Back in the Cisco ACI APIC GUI, select VM Networking > Inventory > Microsoft > Your APIC Virtual Switch. Expand the vSwitch, Controllers, the SCVMM, and Hypervisors. Select each host and verify that the OPFLEX Status is Connected under General on the right.

Set Up Hyper-V Networking

To set up Hyper-V networking, complete the following steps:

1. In virtual Machine Manager, open the VMs and Services workspace. In the list on the left, expand All Hosts > Hyper-V-TNT-A. Right-click on the first tenant host and select Properties.
2. **Select Virtual Switches, then select “New Virtual Network Adapter”. Name the adapter MS-Clust.** Click Browse to the right of VM Network.
3. In the Select a VM Network window, select EPG-MS-Clust in the FP-Foundation tenant. Click OK.
4. Back in the Properties window, the VLAN ID should be automatically filled in. Select the Static IP address configuration and select the MS-Clust IPv4 pool. Click OK and Yes to complete creating the virtual network adapter.
5. Repeat steps 1-4 to create the MS-LVMN adapter connected to EPG-MS-LVMN in the FP-Foundation tenant with the MS-LVMN IPv4 pool.
6. Right-click the first tenant host and select Properties.
7. **Select Virtual Switches, then select “New Virtual Network Adapter”. Name the adapter MS-TNT-A-iSCSI-A.** Click Browse to the right of VM Network.
8. In the Select a VM Network window, select EPG iSCSI-A in the MS-TNT-A tenant. Click OK.
9. Back in the Properties window, the VLAN ID should be automatically filled in. Click OK and Yes to complete creating the virtual network adapter.
10. Repeat steps 6-9 to create the MS-TNT-A-iSCSI-B adapter connected to EPG iSCSI-B in the MS-TNT-A tenant.
11. Repeat steps 6-9 to create the MS-TNT-A-SMB adapter connected to EPG SMB in the MS-TNT-A tenant.
12. Connect to either a Remote Desktop connection or a KVM console connection on the first Hyper-V tenant host as a Domain Admin. In the Windows command prompt, type powershell to open Powershell.
13. Configure Jumbo frames on select interfaces.

```
netsh interface ipv4 set subinterface "vEthernet (MS-Clust)" mtu=9000
store=persistent
```

```
netsh interface ipv4 set subinterface "vEthernet (MS-LVMN)" mtu=9000
store=persistent
```

```
netsh interface ipv4 set subinterface "vEthernet (MS-TNT-A-iSCSI-A)" mtu=9000
store=persistent
```

```
netsh interface ipv4 set subinterface "vEthernet (MS-TNT-A-iSCSI-B)" mtu=9000
store=persistent
```

```
netsh interface ipv4 set subinterface "vEthernet (MS-TNT-A-SMB)" mtu=9000
store=persistent
```

```
netsh interface ipv4 show subinterface
```

- Set IP Address for each host virtual NIC not set from an IPv4 pool.

```
New-NetIPAddress -InterfaceAlias "vEthernet (MS-TNT-A-iSCSI-A)" -IPAddress
<host-tnt-a-iscsi-a-ip> -PrefixLength <tnt-a-iscsi-a-net-prefix>
```

```
New-NetIPAddress -InterfaceAlias "vEthernet (MS-TNT-A-iSCSI-B)" -IPAddress
<host-tnt-a-iscsi-b-ip> -PrefixLength <tnt-a-iscsi-b-net-prefix>
```

```
New-NetIPAddress -InterfaceAlias "vEthernet (MS-TNT-A-SMB)" -IPAddress <host-
tnt-a-smb-ip> -PrefixLength <tnt-b-smb-net-prefix>
```

- Disable DNS registration for all NICs

```
Set-DnsClient -InterfaceAlias * -Register $false
```

- Turn registration back on and configure DNS for the Management NIC

```
Set-DnsClient -InterfaceAlias "00-TNT-Core-MGMT" -Register $true -
ConnectionSpecificSuffix <dns-domain-name>
```

- Restart the NetApp SnapDrive service on the host

```
net stop SWSvc
```

```
y
```

```
net start SWSvc
```

```
net start SnapDriveService
```

- Repeat steps 1-17 for the second Hyper-V tenant host.

Add TNT iSCSI Sessions to Hosts

To add TNT iSCSI sessions to hosts, complete the following steps:

- On another machine that has the Windows GUI and SnapDrive installed, launch the SnapDrive snap-in as the snapdrive domain user and add the Hyper-V-TNT-A-Host-01 server. Add the MS-TNT-A-SVM management interface as a storage system under Transport Protocol Settings. Add 4 iSCSI sessions to the 4 MS-TNT-A-SVM iSCSI LIFs.
- Repeat this process on the Hyper-V-TNT-A-Host-02 server.
- Using SnapDrive, select Disks under the first Hyper-V-TNT-A host.
- On the right, select Create Disk.
- In the Create Disk Wizard, click Next.
- Leave Dedicated selected and click Next.
- Enter ms-tnt-a-svm for the storage system name and click Add.
- In the list of volumes, select the witness volume. Enter witness for LUN name. Click Next.

9. Select **“Do not assign a Drive letter of Volume Mount Point”**. Select Yes to create a Thin-Provisioned LUN and Set the LUN Size to 1 GB. Click Next.
10. Click Next.
11. Select the Microsoft iSCSI Initiator and click Next.
12. Select Automatic and click Next.
13. Click Finish to create the witness LUN and format it.
14. On the storage cluster CLI, unmap the witness LUN and remap it to both tenant hosts.

```

lun show -m -vserver MS-TNT-A-SVM
Vserver      Path                                     Igroup      LUN ID      Protocol
-----
MS-TNT-A-SVM
              /vol/witness/witness                   viaRPC.iqn.2010-11.com.flexpod:a02-6332-host:3
                                                0 iscsi
lun unmap -vserver MS-TNT-A-SVM -path /vol/witness/witness -igroup viaRPC.iqn.2010-
11.com.flexpod:a02-6332-host:3
lun map -vserver MS-TNT-A-SVM -path /vol/witness/witness -igroup Hyper-V-TNT-A-Host-All -lun-id 1
    
```

Create Windows Failover Cluster

To create a Windows Failover Cluster, complete the following steps. Be sure to create DNS records for the Cluster name. The IP address for cluster management should be on the IB-MGMT Subnet.

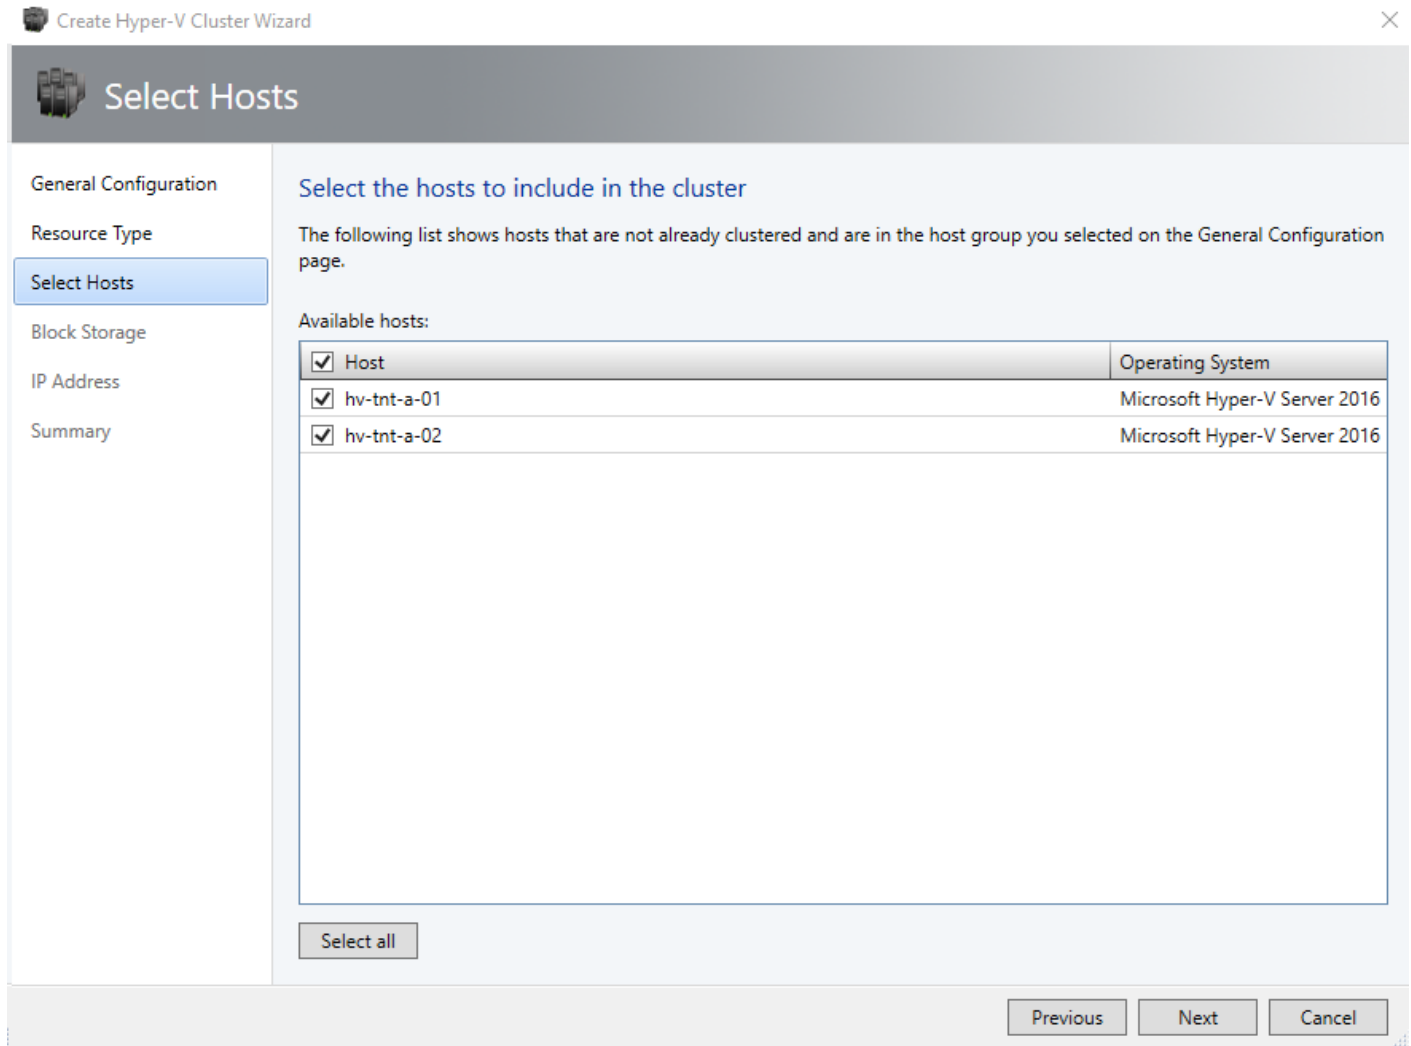
1. In the VMM console, click Fabric > Create > Hyper-V Cluster to open the Create Hyper-V Cluster wizard.
2. In General Configuration, specify a cluster name and choose the tenant host group in which the existing Hyper-V hosts are located. Click Next.

3. In Resource Type, select the SCVMM Run As account that you'll use to create the cluster. Make sure **“Existing servers running a Windows Server operating system”** is selected and click **Next**.

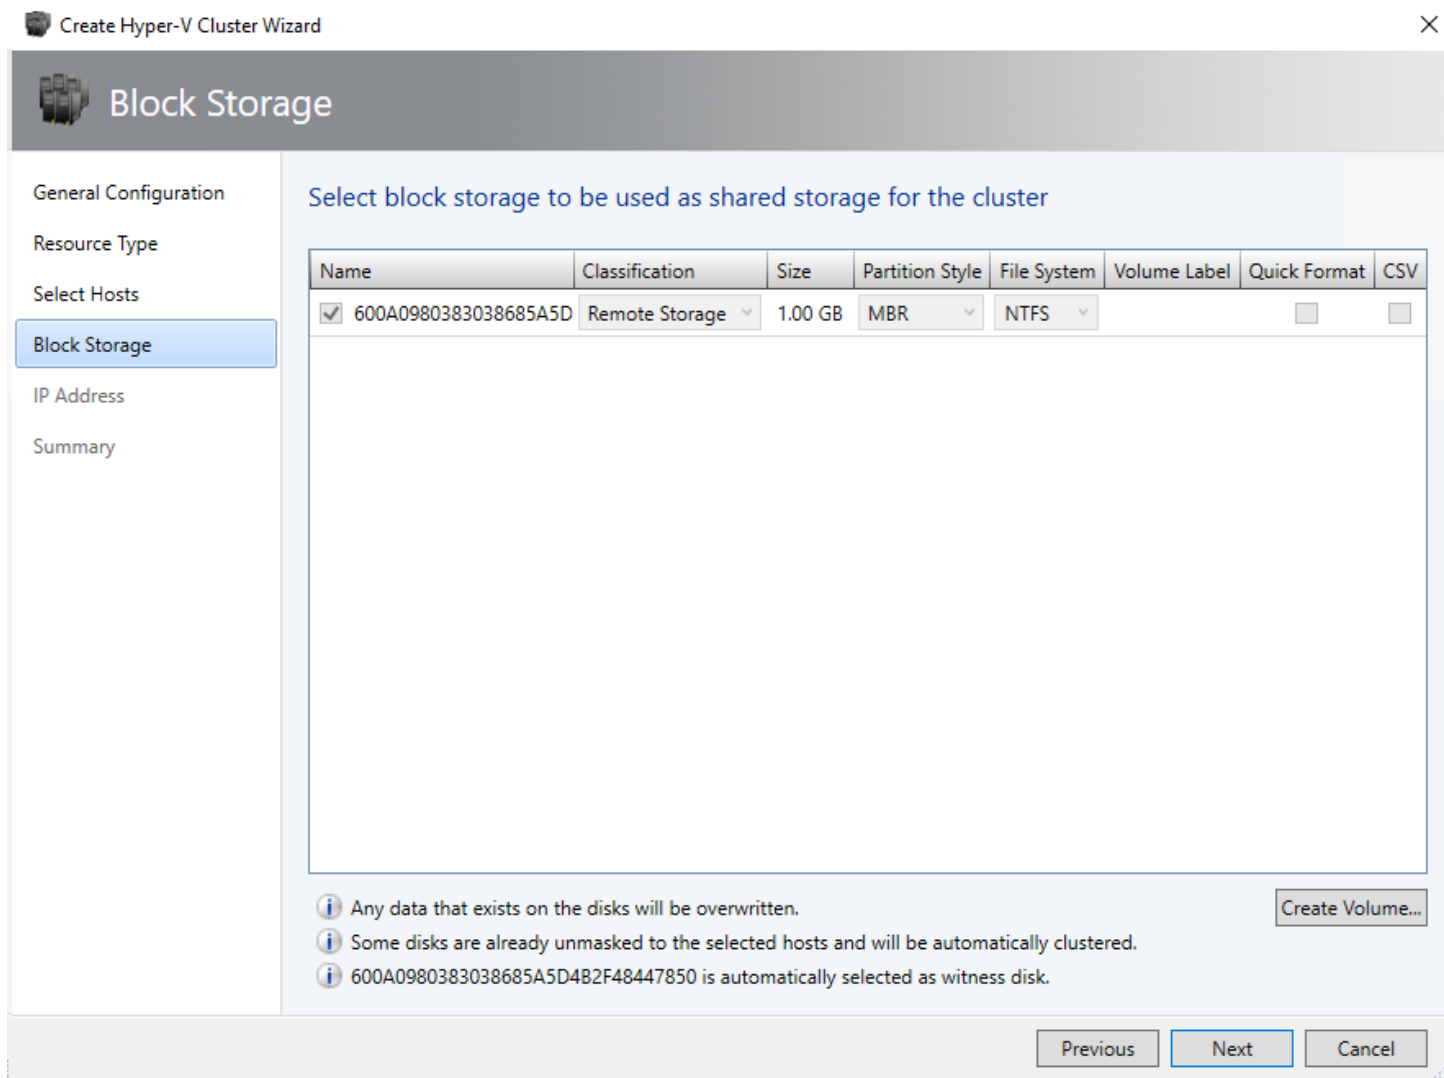


The accounts that you use must have administrative permissions on the servers that will become cluster nodes, and must belong to the same domain as the Hyper-V hosts that you want to cluster. Also, the account requires Create Computer objects permission in the container that is used for Computer accounts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

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

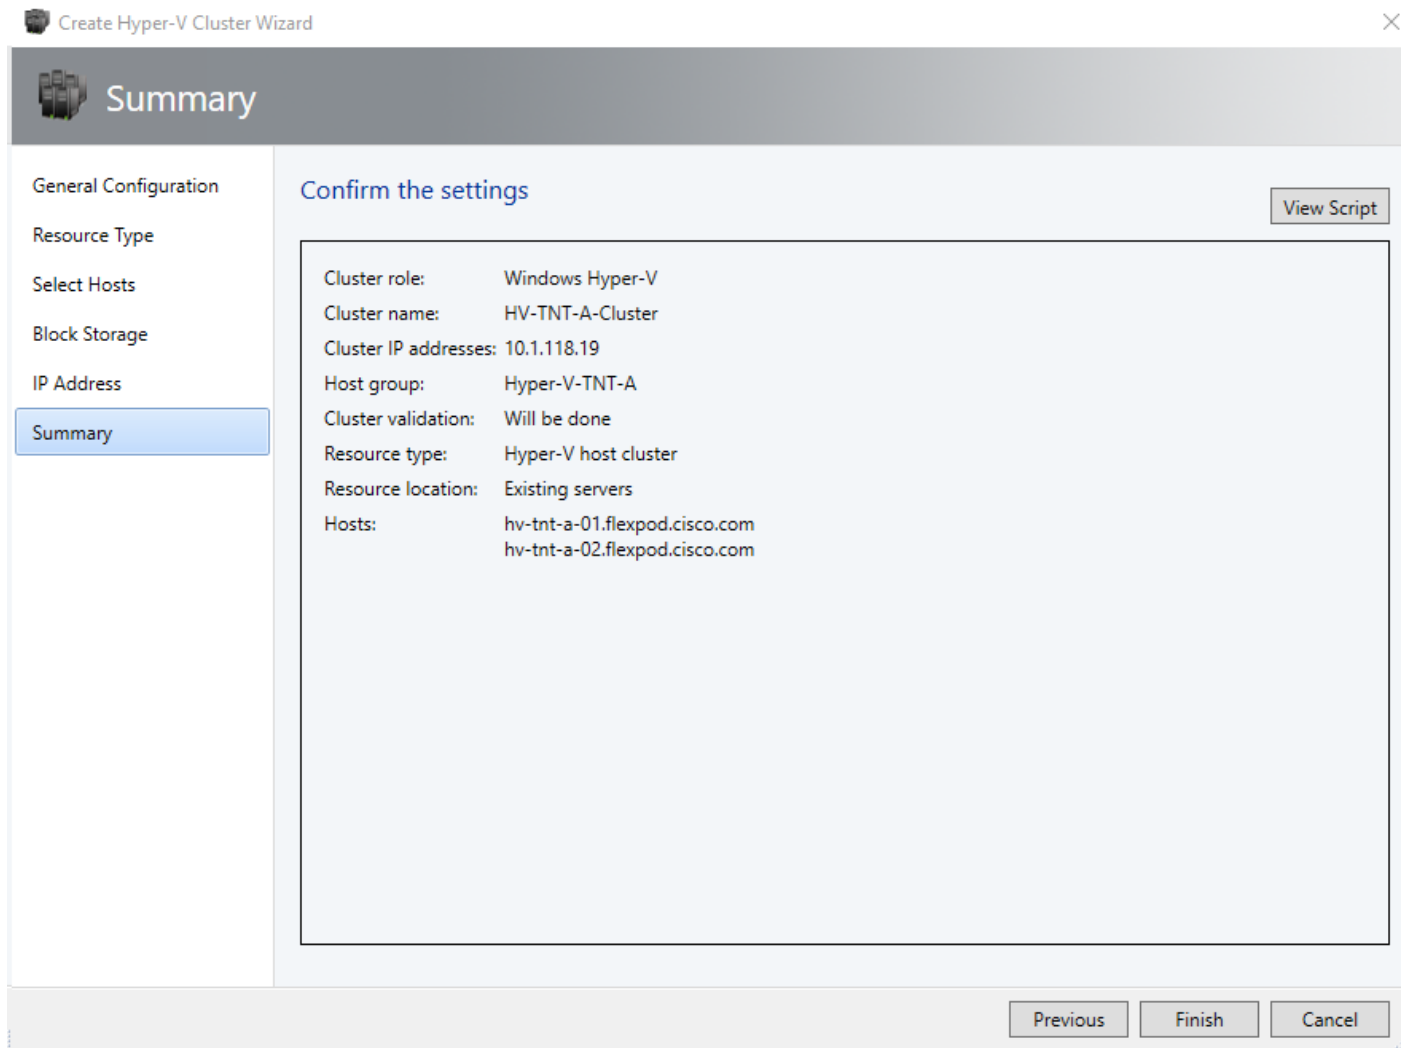


5. In Block Storage, click Next.

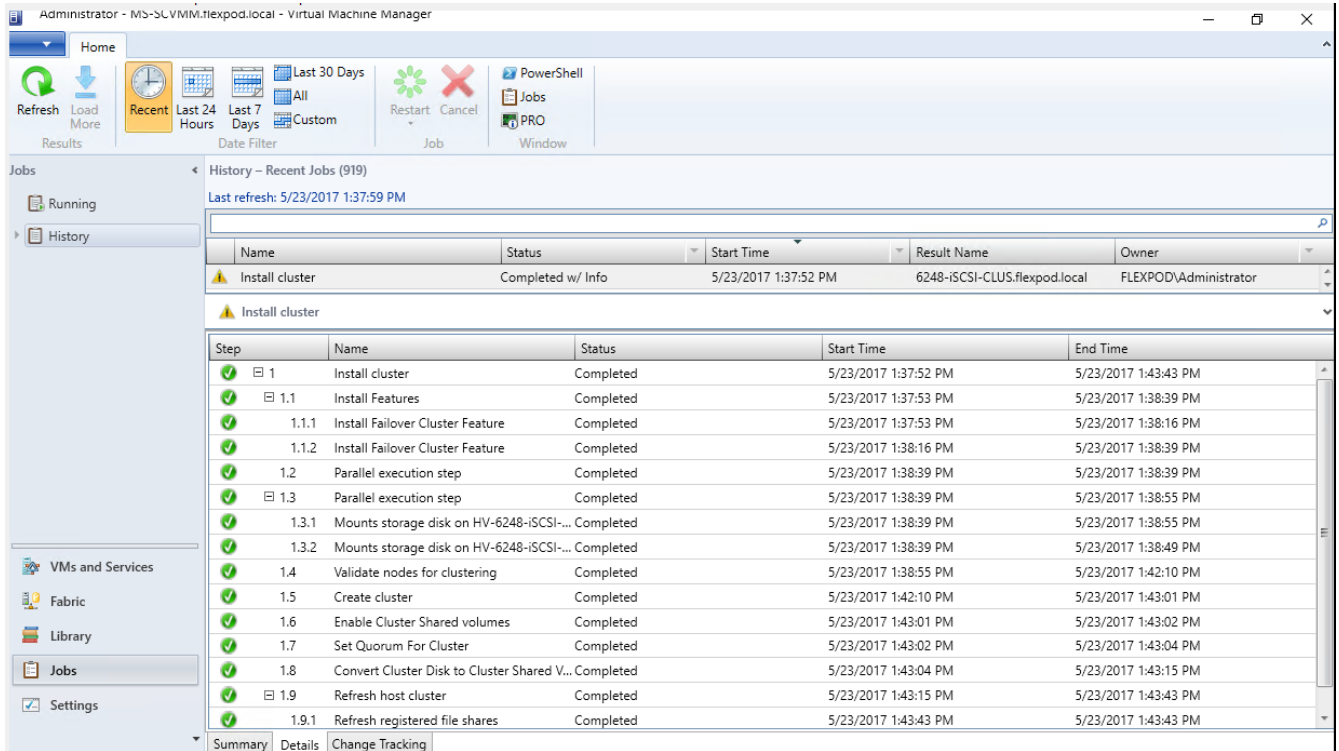


6. In IP address, type in the IP address you want to use for the cluster. Click Next.


7. In Summary, confirm the settings and then click Finish.

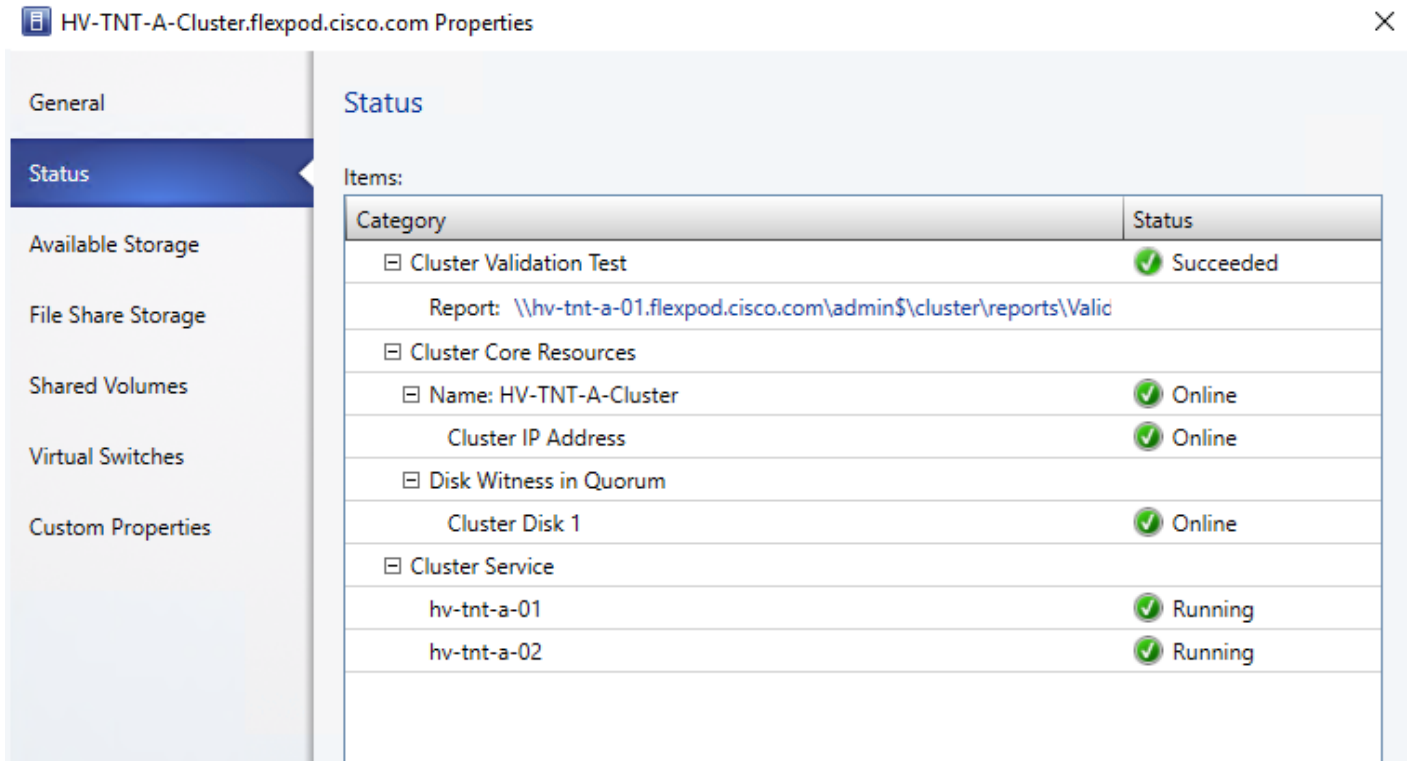


8. You can go to the jobs workspace and click on “Install Cluster” job to see the status of cluster installation. Fix and troubleshoot any errors or warnings and revalidate the cluster.



- 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 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 15 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.

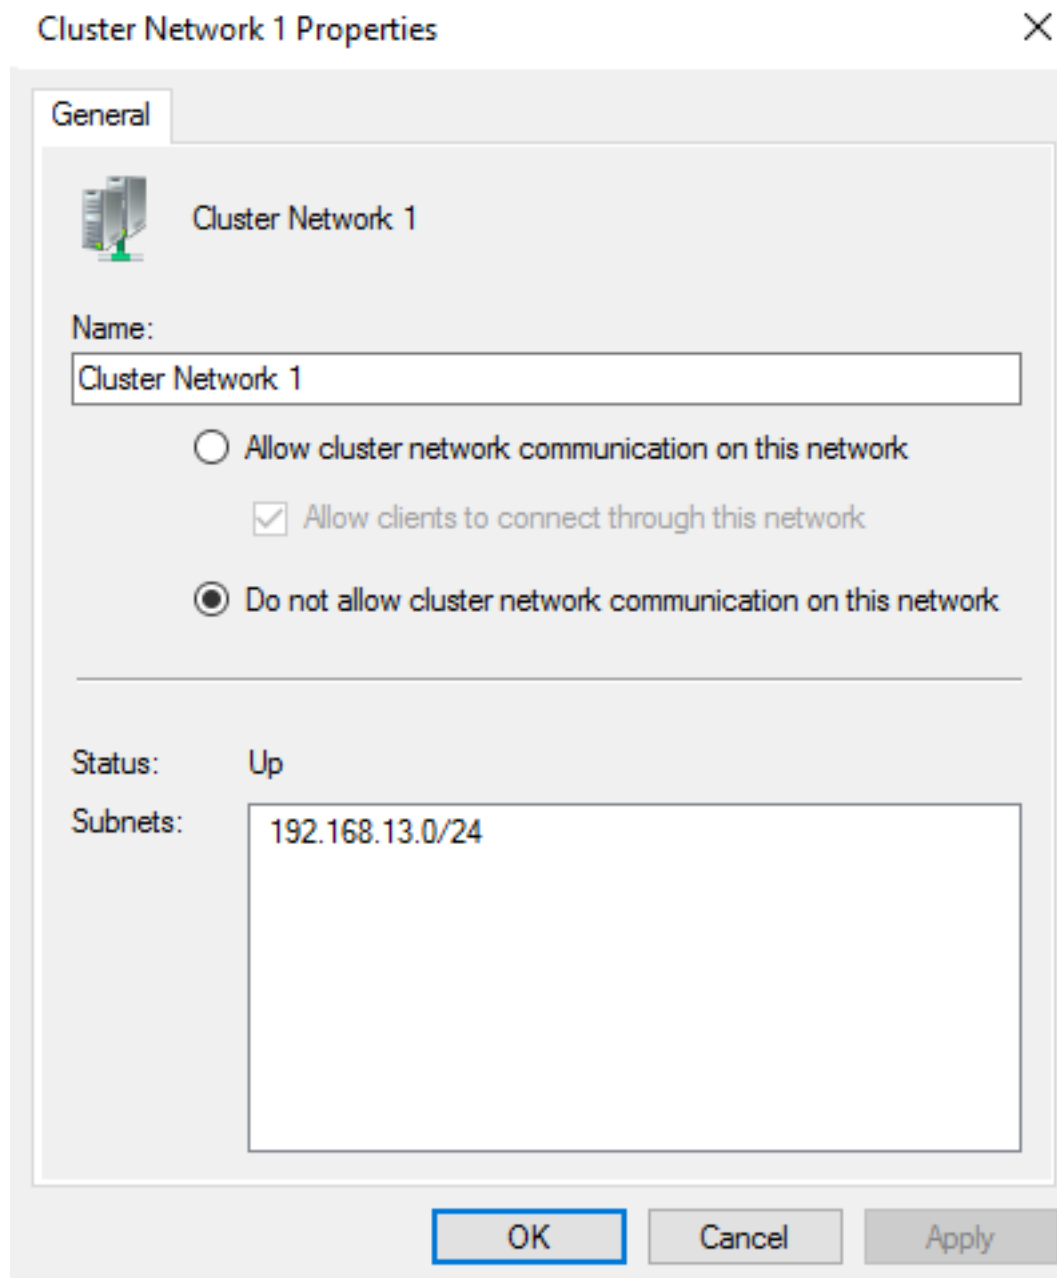
Network Type	Recommended Setting
Live migration	Allow cluster network communication on this network Note: Clear the Allow clients to connect through this network check box.
Storage	Do not allow cluster network communication on this network

1. Open Failover Cluster Manager and connect to the Failover Cluster just created. Click Networks in the navigation tree.

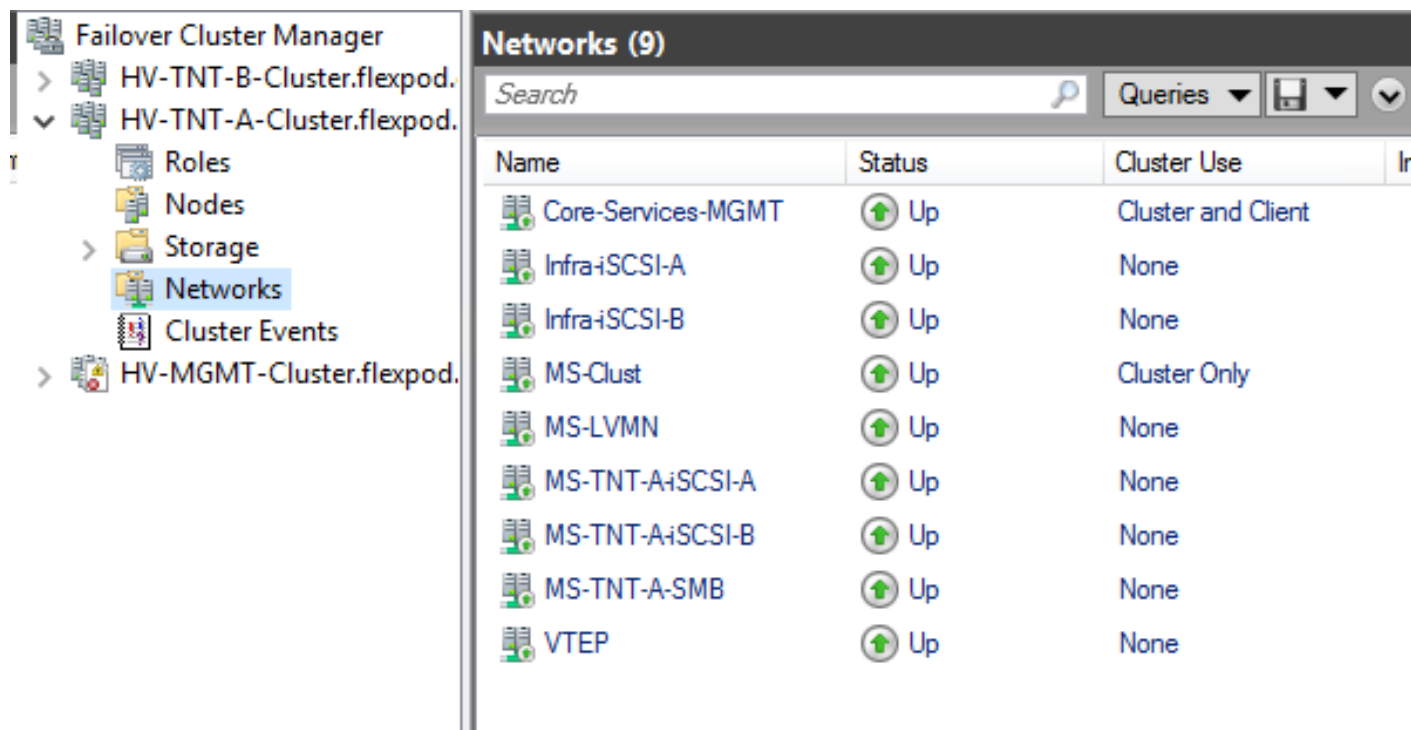


It may be necessary to install the Failover Clustering Tools Feature under Features > Remote Server Administration Tools > Feature Administration Tools in the Add Roles and Features Wizard to install Failover Cluster Manager.

2. In the Networks pane, right-click a network, and then click Properties.



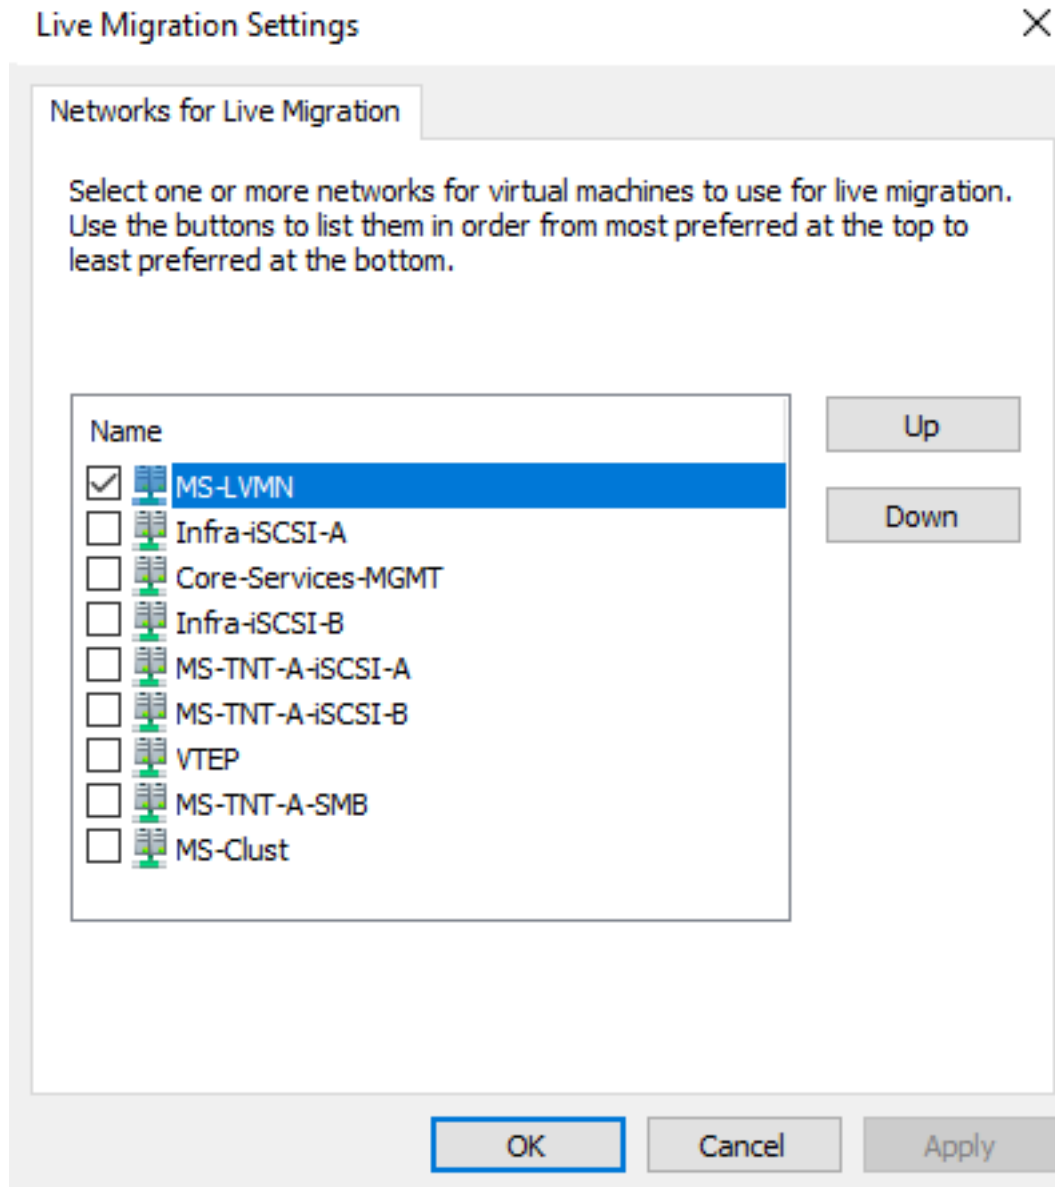
3. Using the subnet information reset the Name of the Cluster Network to the appropriate name, adjust the communication setting and click OK. Storage Networks, the VTEP network, and the LVMN network should not allow cluster communication.
4. Repeat step 3 to assign the descriptive name to all Cluster Networks.



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. Complete the following steps:

1. Open Failover Cluster Manager and connect to the tenant cluster.
2. In the navigation tree, right-click Networks, and then click Live Migration Settings.
3. Select only the Live Migration network (MS-LVMN).



4. Click Apply and OK to save this setting.

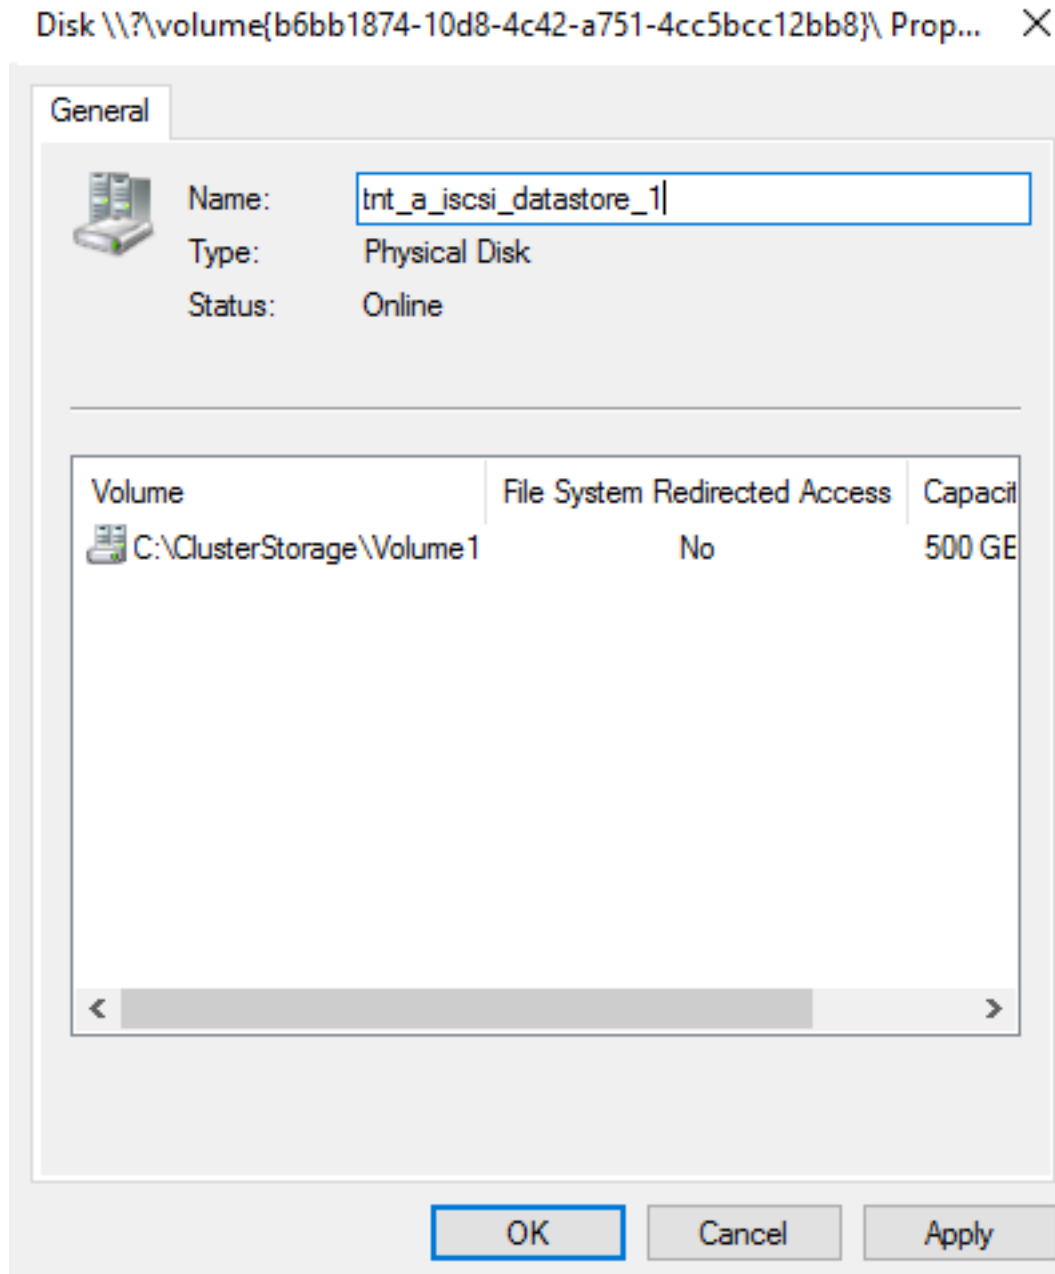
Cluster Storage Settings

1. On the left, expand Storage and select Disks.
2. Right-click the Witness disk and select Properties.
3. Change the Disk Name to Witness and click OK.
4. Note the Owner Node of the Disk.

Add Tenant iSCSI Datastores (Optional)

To add tenant iSCSI datastores, complete the following steps:

1. On another machine that has the Windows GUI and SnapDrive installed, launch the SnapDrive snap-in as the snapdrive domain user and add the Hyper-V-TNT-A-Host-01 server making sure to select the **checkbox next to “Add SnapDrive instances from all the Microsoft Cluster nodes”**. Hyper-V-TNT-A-Host-02 should appear as well.
2. Select Disks on the Witness Disk Owner Node and on the right select Create Disk.
3. In the Create Disk Wizard, click Next.
4. **Select “Shared (Microsoft Cluster Services only)” and click Next.**
5. Select the MS-TNT-A-SVM. In the list, select the tnt_a_iscsi_datastore_1 volume. Enter iscsi_datastore_1 for the LUN Name. Click Next.
6. Make sure both Tenant Hyper-V servers are selected and click Next.
7. **Select “Do not assign a Drive letter or Volume Mount Point.** Select Yes to Thin-Provision the LUN and enter 500 GB for the LUN Size. Click Next.
8. Click Next.
9. Select each Host Node and its corresponding Microsoft iSCSI Initiator. Click Next.
10. Select Automatic Initiator Group management and click Next.
11. **Select “Add to cluster shared volumes” and click Next.**
12. Click Finish to create the Cluster Shared Volume.
13. Repeat steps 2-12 to create iscsi_datastore_2 in the tnt_a_iscsi_datastore_2 volume.
14. Back in Failover Cluster Manager on the SCVMM VM, select the first Cluster Shared Volume, right-click and select Properties.
15. By referencing SnapDrive and looking at the Volume Name, determine the iSCSI datastore and adjust the volume name in Failover Cluster Manager to match. Click OK.



Configure the NetApp SMI-S Provider

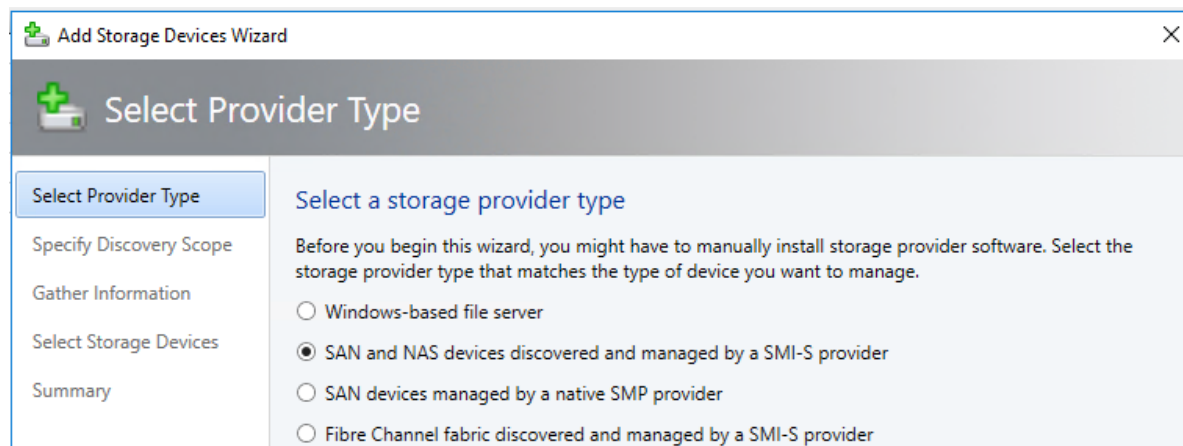
1. On a console interface to the SMI-S Provider VM, add the Tenant SVM to the SMI-S Provider using the following command:

```
C:\Program Files (x86)\NetApp\smis\pegasus\bin>smis addsecure 172.18.254.2 vsadmin
Enter password: *****
Returned Path  ONTAP_FilerData.hostName="172.18.254.2",port=443

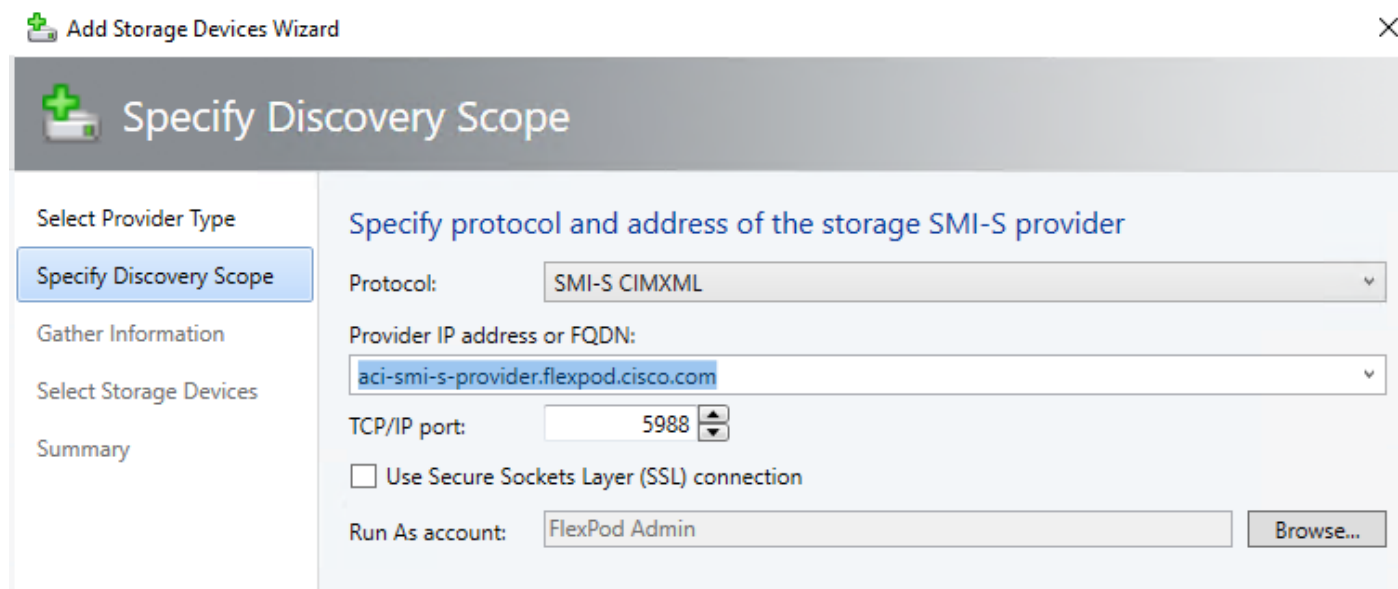
Successfully added 172.18.254.2
```

Add a Storage Device


1. In Virtual Machine Manager, 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. Click Next.



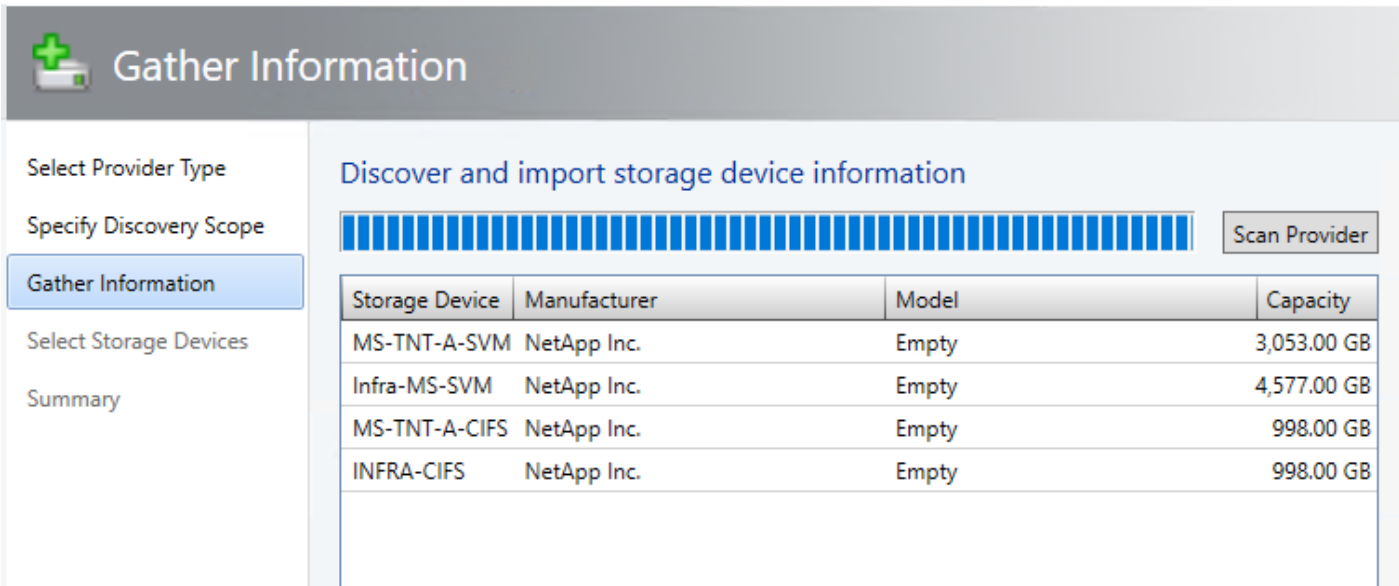
3. In Specify Discovery Scope, select Protocol - SMI-S CIMXML, add the IP address/FQDN of the SMI-S Provider, and add the port used to connect to the provider on the remote server. You can enable SSL if you're using CIMXML. Then specify an account for connecting to the provider. You will need to create a Run As account for the flexadmin account added to the CIM server above. Click Next.



4. In Gather Information, VMM automatically tries to discover and import the storage device information. You may need to import the security certificate.

 You may need to wait up to 30 minutes for the tenant SVM to appear in the Add Storage Devices Wizard.

- 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.



The screenshot shows the 'Add Storage Devices Wizard' window, specifically the 'Gather Information' step. The window title is 'Add Storage Devices Wizard' with a close button (X) in the top right corner. The main title is 'Gather Information'. On the left, there is a navigation pane with the following options: 'Select Provider Type', 'Specify Discovery Scope', 'Gather Information' (which is selected and highlighted in blue), 'Select Storage Devices', and 'Summary'. The main area is titled 'Discover and import storage device information'. It features a progress bar with 20 blue segments, a 'Scan Provider' button, and a table with the following data:

Storage Device	Manufacturer	Model	Capacity
MS-TNT-A-SVM	NetApp Inc.	Empty	3,053.00 GB
Infra-MS-SVM	NetApp Inc.	Empty	4,577.00 GB
MS-TNT-A-CIFS	NetApp Inc.	Empty	998.00 GB
INFRA-CIFS	NetApp Inc.	Empty	998.00 GB

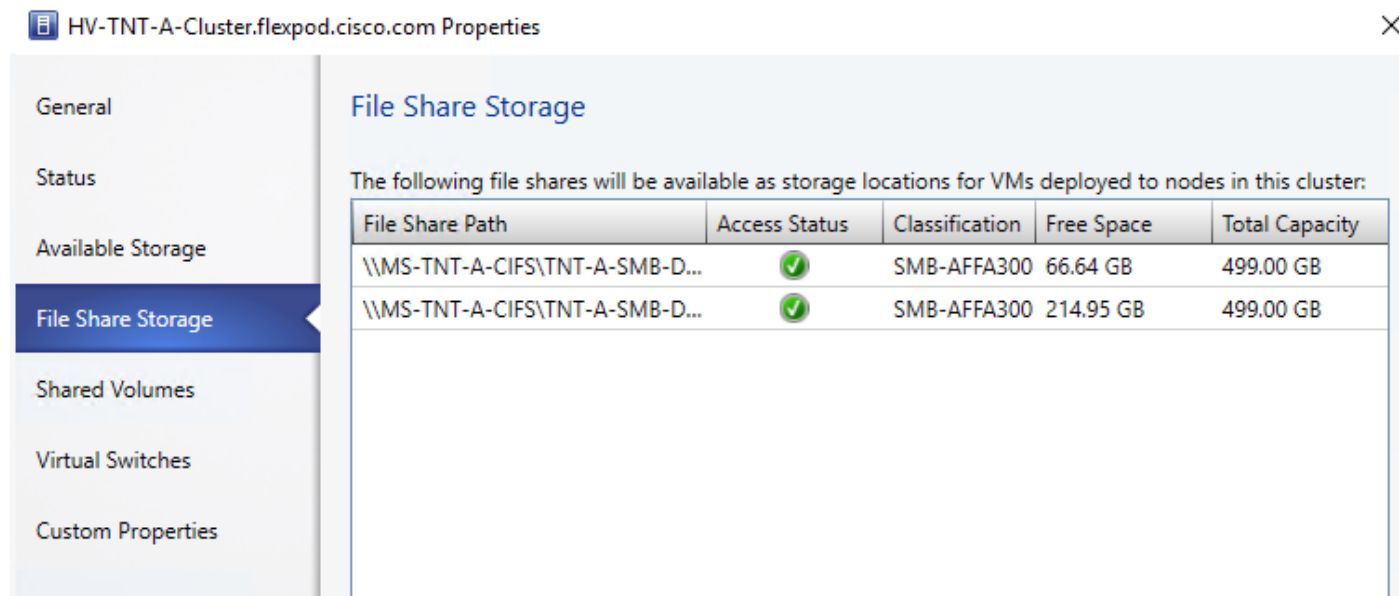
- In Select Storage Devices, specify a classification and host group from the drop-down list for each storage pool. Create storage classifications if none exists to group storage pools with similar characteristics. Only select storage where VMs will be stored. Click Next.
- On the Summary page, confirm the settings, and then click Finish. The Jobs dialog box appears. When status is Completed you can verify the storage in Fabric > Storage > Classifications and Pools.

Create and Assign SMB 3.0 File Shares to the Hyper-V Tenant Cluster

SMB file shares can be used by Hyper-V hosts as a shared storage to store virtual machine files. To create and assign SMB file shares to stand-alone Hyper-V servers and host cluster, complete the following steps:

- To Add a storage device, refer to the steps covered in the above section.
- To create a file share, open Fabric workspace, expand Storage and click File Servers.
- Select the File Server and click 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.
- In the Capacity page, enter a size and click Next.
- In the Summary page, confirm the setting and click Finish.
- Verify the file share created in the above steps by navigating to Fabric > Storage and click File Servers.
- Repeat this process to add a file share for the tnt_a_smb_datastore_2 Storage pool.

8. Assign the file shares to the host cluster by navigating to Fabric > Servers > All Hosts > Hyper-V Management > HV-TNT-A-Cluster.
9. Locate and right-click the cluster icon and click Properties.
10. Click File Share Storage and click Add.
11. From the drop-down list next to the File Share Path, select a share and click OK.
12. Repeat this step to select the other share. Click OK.



Build a Second Tenant (Optional)

In this lab validation a second tenant was built to demonstrate that multiple tenants could access and use the Shared-L3-Out and that tenants can be completely logically separated, but can also have overlapping IP address spaces. The second tenant built in this lab validation had the following characteristics and was built with the same contract structure as the TNT-A tenant:

Table 16 Lab Validation Tenant TNT-B Configuration

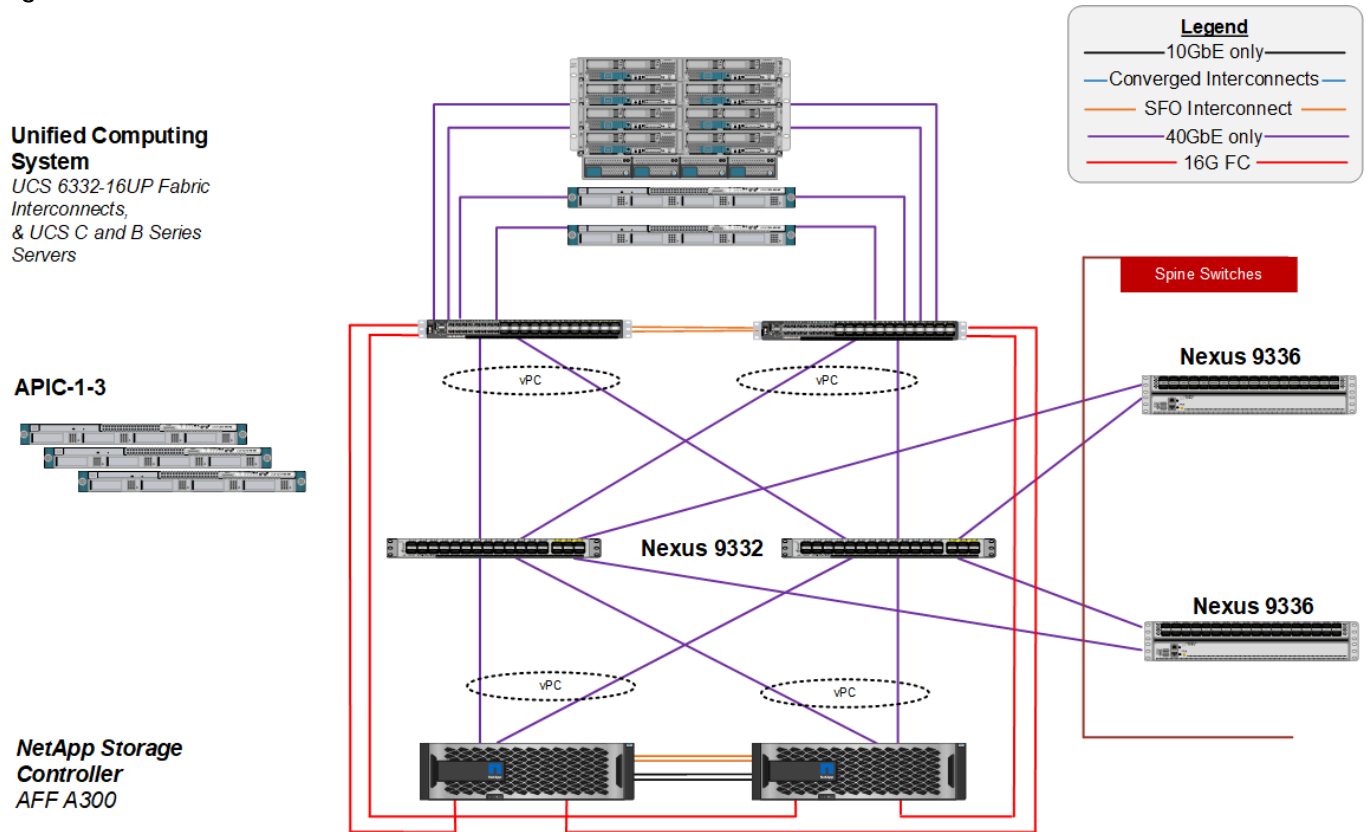
EPG	Storage VLAN	Subnet / Gateway	Bridge Domain
iSCSI-A	3015	192.168.14.0/24 – L2	BD-iSCSI-A
iSCSI-B	3025	192.168.24.0/24 – L2	BD-iSCSI-B
SMB-LIF	3056	192.168.55.0/24 – L2	BD-SMB
SMB-Host	N/A	192.168.55.0/24 – L2	BD-SMB
SVM-MGMT	220	172.18.254.10/29	BD-Internal
Web	N/A	172.18.3.254/24	BD-Internal
App	N/A	172.18.1.254/24	BD-Internal
DB	N/A	172.18.5.254/24	BD-Internal

Appendix - FC Solution

This section details the configuration steps for the Cisco UCS 6332-16UP Fabric Interconnects (FI) in a design that will support direct connectivity to NetApp AFF using 16 Gb/s Fibre Channel.

Figure 5 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 FC 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 16Gb FC connection between the Cisco UCS Fabric Interconnect and NetApp AFF A300. FC zoning is done in the Cisco UCS Fabric Interconnect. This infrastructure is deployed to provide FC-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 5 FlexPod with Cisco UCS 6332-16UP Fabric Interconnects and Cisco UCS Direct Connect SAN



This appendix will only detail the delta configuration required for the Storage, UCS, Windows Server, and Hyper-V setup. As mentioned in other sections of the document, any iSCSI steps can be skipped if iSCSI is not being implemented.

Storage Configuration

Set Onboard Unified Target Adapter 2 Port Personality

In order to use FC storage targets, FC 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	fc	target	-	-	online
<st-node01>	0h	fc	target	-	-	online
<st-node02>	0e	fc	target	-	-	online
<st-node02>	0f	fc	target	-	-	online
<st-node02>	0g	fc	target	-	-	online
<st-node02>	0h	fc	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 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 `fcport 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 `fcport adapter modify -node <home-node-of-the-port> -adapter <port-name> -state up`.

Add FCP Storage Protocol to Infrastructure SVM

Run the following command to add the FCP storage protocol to the Infrastructure SVM. It is assumed that an FCP license has been installed on each cluster node:

```
vserver add-protocols -vserver Infra-MS-SVM -protocols fcp
vserver show -fields allowed-protocols
```

Create FCP Storage Protocol in Infrastructure SVM

Run the following command to create the FCP storage protocol in the Infrastructure SVM. It is assumed that an FCP license has been installed on each cluster node:

```
fcv create -vserver Infra-MS-SVM
fcv show
```

Create FC LIFs

Run the following commands to create four FC LIFs (two on each node) in the Infrastructure SVM by using the previously configured FC 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

network interface show -vserver Infra-MS-SVM -lif fcp*
```

Server Configuration

This section details the delta steps in the UCS setup to provide FC-based boot and storage.

Configure FC Unified Ports (UP) on UCS Fabric Interconnects

In order to use Fiber Channel Storage Ports for storage directly connected to the Cisco UCS fabric interconnects, 10G Ethernet ports need to be converted to 16G FC ports..

To convert the first six ports of the UCS 6322-16 UP to FC, 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 Configure Unified Ports. Click Yes to proceed.
4. The 6322-16UP requires ports to be converted in groups of 6 ports from the left. To convert the first six ports, move the slider to the right until ports 1-6 are highlighted.

Configure Unified Ports



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	

- Click OK, then click Yes, then click OK to convert the ports. The Fabric Interconnect will reboot. Wait until the reboot is complete and the FI is back in the UCS domain cluster. This can be checked by logging into the FI's CLI and typing `show cluster state`. Wait until the "HA Ready" state appears.
- Repeat this process to convert the first six ports of FI B.

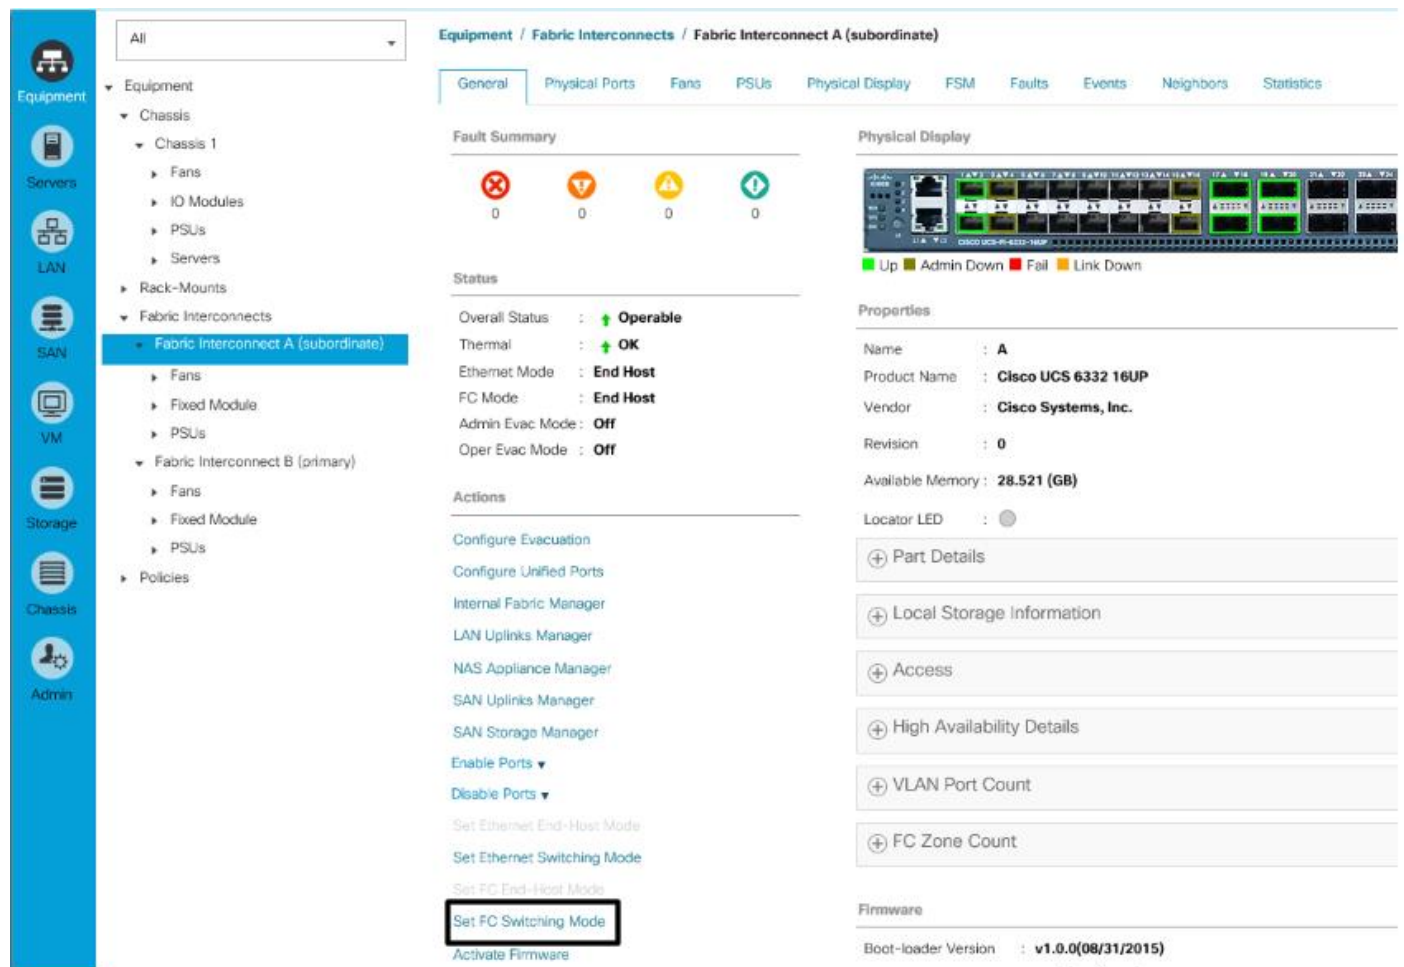
Place Cisco UCS Fabric Interconnects in Fiber Channel Switching Mode

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:

- 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.




4. Wait for both Fabric Interconnects to reboot by monitoring the console ports and log back into Cisco UCS Manager.

 It may be necessary to go to the Pending Activities window and select Reboot Fabric Interconnect to reboot the Primary Fabric Interconnect.

Create Storage VSANs

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 icon on the left.

 In this step, 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 ? 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

<p>You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.</p> <p>Enter the VSAN ID that maps to this VSAN.</p> <p>VSAN ID: <input type="text" value="101"/></p>	<p>A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.</p> <p>Enter the VLAN ID that maps to this VSAN.</p> <p>FCoE VLAN: <input type="text" value="101"/></p>
---	---

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.

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 :

16. Click OK, and then click OK again

Configure FC Storage Ports

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 (subordinate) > Fixed Module > FC Ports
3. Select the ports (1 and 2 for this document) that are connected to the NetApp array, right-click them, **and select "Configure as FC Storage Port"**
4. Click Yes to confirm and then click OK.
5. Verify that the ports connected to the NetApp array are now configured as FC Storage ports.
6. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module > FC Ports
7. Select the ports (1 and 2 for this document) that are connected to the NetApp array, right-click them, **and select "Configure as FC 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 FC Storage ports.

Assign VSANs to FC Storage Ports

To assign storage VSANs to FC 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 FC Interfaces.
4. Select the first FC Interface (1/1)
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 second FC Interface (1/2)
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 FC Interfaces.
11. Select the first FC Interface (1/1)
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 second FC Interface (1/2)
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 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 steps: '1 Define Name and Description' (highlighted in blue) and '2 Add WWNN Blocks'. The main content area is titled 'Create WWNN Pool' and contains three input fields: 'Name' with the value 'WWNN-POOL', 'Description' which is empty, and 'Assignment Order' with radio buttons for 'Default' and 'Sequential' (the latter is selected). At the bottom of the dialog 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 and click OK.



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 A2 to represent as identifying information for this being in the UCS 6332 in the 2nd cabinet of the A row. 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.

Create WWN Block



From : Size :

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

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



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 [?] [X]

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:A2:0A:00

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

Create WWN Block



From : Size :

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

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



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 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 :

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : [Create VSAN](#)

Template Type : Initial Template Updating Template

Max Data Field Size :

WWPN Pool : ▼

QoS Policy : ▼

Pin Group : ▼

Stats Threshold Policy : ▼

OK

Cancel

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-POOI-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



Name :

Use vHBA Template :

Redundancy Pair :

Peer Name :

vHBA Template :

[Create vHBA Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Fibre Channel Adapter Policy](#)

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



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
ℹ Modify

OK
Cancel

19. Click OK to create the SAN Connectivity Policy.
20. Click OK to confirm creation.

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 `FC-Boot-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 `FC-Boot-Pool` server pool.
9. Click Finish.
10. 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-Infra-Host-A` as the name of the vNIC.
8. Select the Use vNIC Template checkbox.
9. In the vNIC Template list, select Infra-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](#)

12. Click the upper Add button to add another vNIC to the policy.

13. In the Create vNIC box, enter 01-Infra-Host-B as the name of the vNIC.

14. Select the Use vNIC Template checkbox.

15. In the vNIC Template list, select Infra-Host-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-APIC-MS-VS-A as the name of the vNIC.

20. Select the Use vNIC Template checkbox.

21. In the vNIC Template list, select APIC-MS-VSA.

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 another vNIC to the policy.

25. In the Create vNIC box, enter 03-APIC-MS-VS-B as the name of the vNIC.

26. Select the Use vNIC Template checkbox.

27. In the vNIC Template list, select APIC-MS-VS-B.

28. In the Adapter Policy list, select Windows.

29. Click OK to add the vNIC to the policy.

Create LAN Connectivity Policy



Name :

Description :

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

Name	MAC Address	Native VLAN
vNIC 03-APIC-MS-VS-B	Derived	
vNIC 02-APIC-MS-VS-A	Derived	
vNIC 01-Infra-Host-B	Derived	
vNIC 00-Infra-Host-A	Derived	

Delete Add Modify

Add iSCSI vNICs

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



If creating an FC-booted Tenant host, replace vNICs 00 and 01 in this procedure with vNIC 00-TNT-Core-MGMT with the TNT-Core-MGMT-A vNIC Template.

Create Boot Policy (FC Boot)

This procedure applies to a Cisco UCS environment in which two FC logical interfaces (LIFs) are on cluster node 1 (fcp_lif01a and fcp_lif01b) and two FC LIFs are on cluster node 2 (fcp_lif02a and fcp_lif02b).

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 list and select Add Remote CD/DVD.

9. Expand the vHBAs drop-down list 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 list, select Add SAN Boot Target.

14. Keep 0 as the value for Boot Target LUN.

15. Enter the WWPN for fcp_lif01a.



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



Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary



17. Click OK to add the SAN boot target.
18. From the vHBA drop-down list, select Add SAN Boot Target.
19. Enter 0 as the value for Boot Target LUN.
20. Enter the WWPN for fcp_lif02a.
21. Click OK to add the SAN boot target.
22. From the vHBA drop-down list, 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 list, select Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for fcp_lif01b.

29. Select Primary for the SAN boot target type.
30. Click OK to add the SAN boot target.
31. From the vHBA drop-down list, select Add SAN Boot Target.
32. Keep 0 as the value for Boot Target LUN.
33. Enter the WWPN for fcp_lif02b.
34. Click OK to add the SAN boot target.

Create Boot Policy



Name : FC-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
- iSCSI vNICs
- CIMC Mounted vMedia
- EFI Shell

Boot Order

+ - Advanced Filter Export Print

N...	vNIC/vHBA/i...	Type	WWN	LU...	Slo...	Bo...	Bo...	Des...
	Fabric-A	Primary						
		Primary	20:01:00:A0:98:AA:50:8B	0				
		Secondary	20:03:00:A0:98:AA:50:8B	0				
	Fabric-B	Secondary						
		Primary	20:02:00:A0:98:AA:50:8B	0				
		Secondary	20:04:00:A0:98:AA:50:8B	0				

Move Up Move Down Delete

Set Uefi Boot Parameters

35. Click OK, then click OK again to create the boot policy.

Create Boot Policy (FC Boot) With a Single Path for Windows Installation

This procedure applies to a Cisco UCS environment in which two FC logical interfaces (LIFs) are on cluster node 1 (fcp_lif01a and fcp_lif01b) and two FC LIFs are on cluster node 2 (fcp_lif02a and fcp_lif02b).

To create a boot policy with a single boot path 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-One-Path` 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 list and select `Add Remote CD/DVD`.
9. Expand the vHBAs drop-down list and select `Add SAN Boot`.
10. Select the Primary for type field.
11. Enter `Fabric-A` in vHBA field.

Add SAN Boot ? X

vHBA :

Type : Primary Secondary Any

12. Click OK.
13. From the vHBA drop-down list, select `Add SAN Boot Target`.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for `fcp_lif01a`.



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



Boot Target LUN :

Boot Target WWPN :

Type : Primary Secondary

OK

Cancel

17. Click OK to add the SAN boot target.

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.

[Add SAN Boot](#)

[Add SAN Boot Target](#)

Boot Order

Name	vNI...	Type	WWN	LU...	Slot...	Boo...	Boo...	Des...
Remote CD/...	1							
San	2							
SAN Prim...		Fab...	Primary					
SAN T...			Primary	20:01:00:A0:98:AA:50:8B	0			

18. Click OK, then click OK again to create the boot policy.

Create Service Profile Templates

In this procedure, two service profile templates for Infrastructure Hyper-V hosts are 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-FC-Host` as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the "Updating Template" option.

- Under UUID Assignment, 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

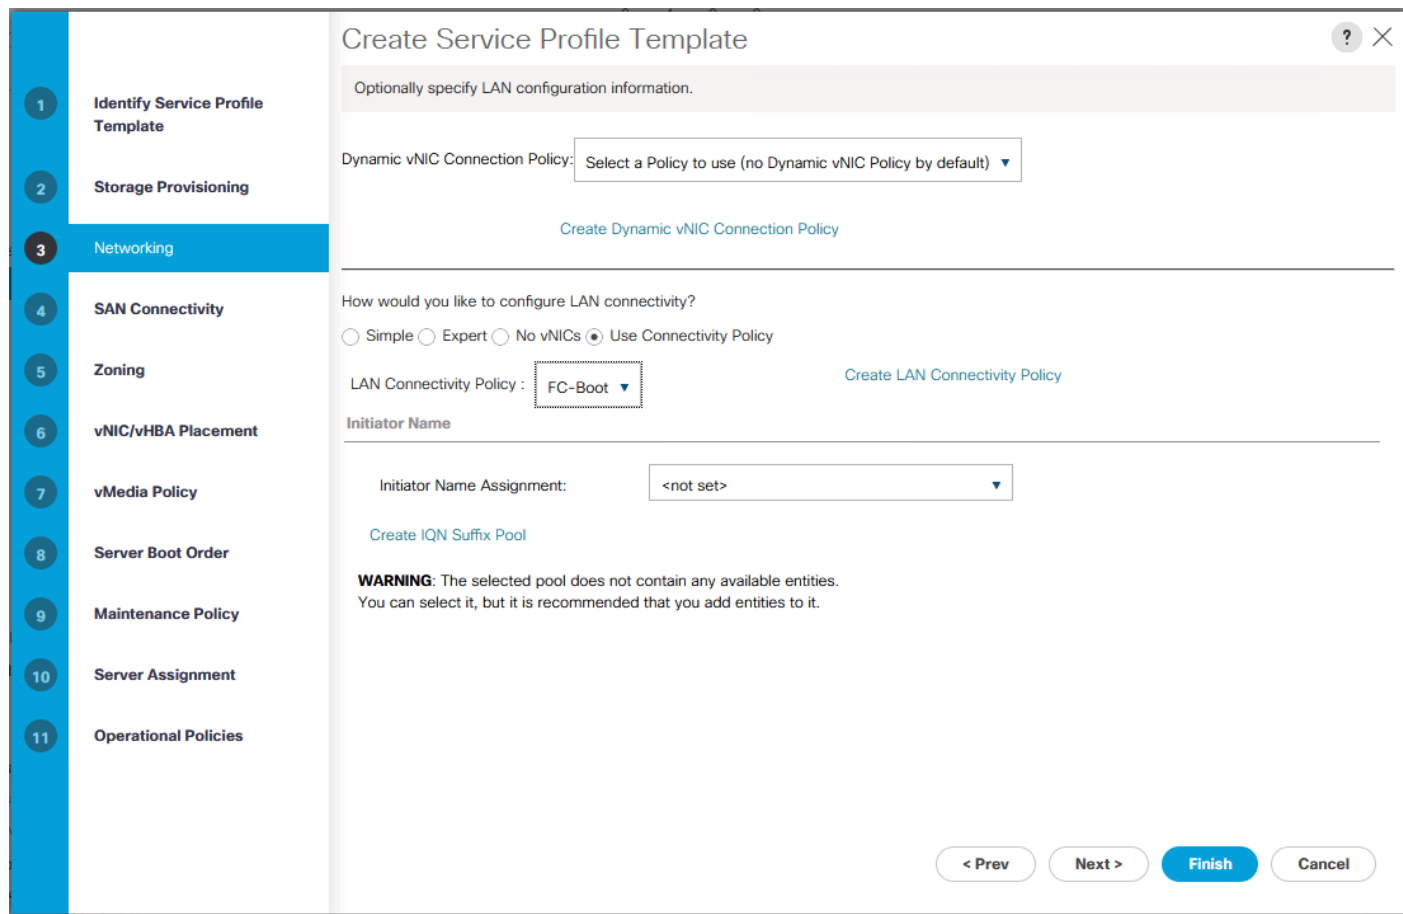
- Click Next.

Configure Storage Provisioning

- 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.
- Click Next.

Configure Networking Options

- Keep the setting at default for Dynamic vNIC Connection Policy.
- Select the **“Use Connectivity Policy”** option to configure the LAN connectivity.
- Select FC-Boot from the LAN Connectivity Policy drop-down.
- Leave Initiator Name Assignment at <not set>.

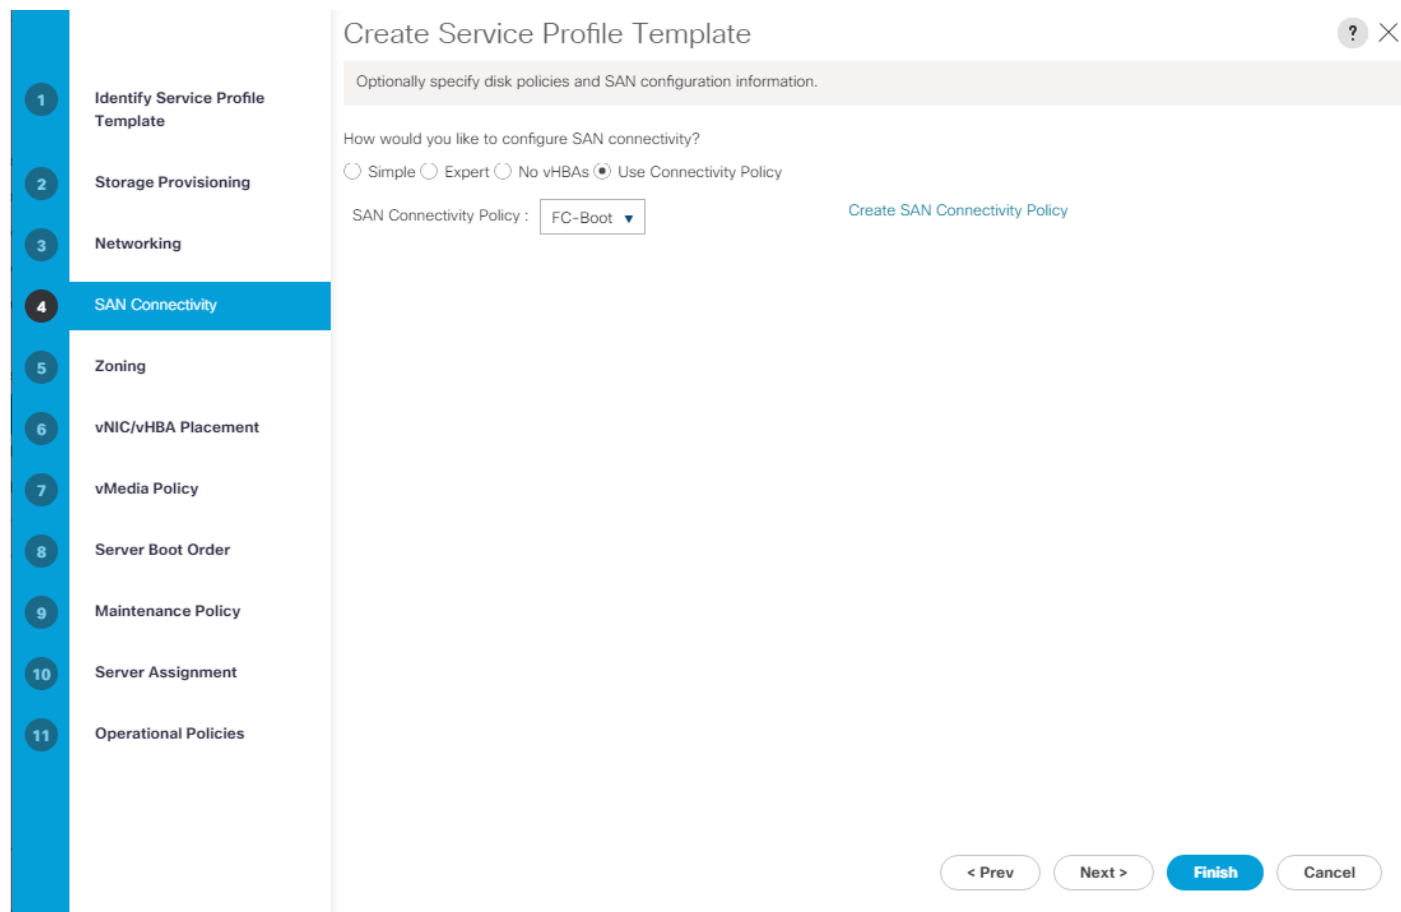


5. Click Next.

Configure Storage Options

To configure storage options, complete the following steps:

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 drop-down list.



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 FC-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 : **FC-Boot**
 Description :
 Reboot on Boot Order Change : **No**
 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/vHB...	Type	WWN	LUN Name	Slot Numb...	Boot Name	Boot Path	Description
Remot...	1								
▼ San	2								
▶ SA...		Fabric-A	Primary						
▶ SA...		Fabric-B	Secondary						

[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.

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 `FC-Boot-Pool1`.
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 **Virtual-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

⊕ Graphics Card Policy

< Prev Next > **Finish** Cancel

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

Create Windows Installation Service Profile Template

1. Right-click the just-created Hyper-V-FC-Host Service Profile Template and select Create a Clone.
2. Name the clone Install-Win-FC-Host and click OK.
3. On the left, select the Install-Win-FC-Host Service Profile Template. In the center pane, select the Boot Order tab. Click Modify Boot Policy. Use the pulldown to select the FC-One-Path Boot Policy and click OK, then OK again.

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 Install-Win-FC-Host.
3. Right-click Install-Win-FC-Host and select Create Service Profiles from Template.

4. Enter `Hyper-V-MGMT-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 17 and Table 18 .

Table 17 WWPNS from NetApp storage

SVM	Adapter	MDS Switch	Target: WWPNS
-----	---------	------------	---------------

SVM	Adapter	MDS Switch	Target: WWPN
Infra-MS-SVM	fcp_lif01a	Fabric A	<fcp_lif01a-wwpn>
	fcp_lif01b	Fabric B	<fcp_lif01b-wwpn>
	fcp_lif02a	Fabric A	<fcp_lif02a-wwpn>
	fcp_lif02b	Fabric B	<fcp_lif02b-wwpn>



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

Table 18 WWPNs for UCS Service Profiles

Cisco UCS Service Profile Name	MDS Switch	Initiator WWPN
Hyper-V-MGMT-Host-01	Fabric A	Hyper-V-MGMT-Host -01-wwpna
	Fabric B	Hyper-V-MGMT-Host -01-wwpnb
Hyper-V-MGMT-Host -02	Fabric A	Hyper-V-MGMT-Host -02-wwpna
	Fabric B	Hyper-V-MGMT-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 (fcp_lif01a) 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 (fcp_lif02a) 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 (fcp_lif01b) 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 (fcp_lif02b) 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 section, 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.
3. Create a duplicate of the FC-Boot SAN Connectivity Policy for the tenant. This policy will need to be mapped into the Tenant Service Profile Template.
4. Select the Tenant San Connectivity Policy.

5. In the center pane, select the vHBA Initiator Groups tab.
6. Click Add to add a vHBA Initiator Group.
7. Name the group Fabric A and select the Fabric A Initiator.
8. Use the pulldown to select the Tenant Fabric A Storage Connection Policy.
9. Click OK and OK to complete adding the Initiator Group.
10. Click Add to add a vHBA Initiator Group.
11. Name the group Fabric B and select the Fabric B Initiator.
12. Use the pulldown to select the Fabric B Storage Connection Policy.
13. Click OK and OK to complete adding the Initiator Group.

Create igroups

From the storage cluster CLI, to create igroups, run the following commands:

```
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-Host-01 -protocol fcp -ostype windows -
initiator <hyper-v-mgmt-host-01-wwpna>,<hyper-v-mgmt-host-01-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup Hyper-V-MGMT-Host-02 -protocol fcp -ostype windows -
initiator <hyper-v-mgmt-host-02-wwpna>,<hyper-v-mgmt-host-02-wwpnb>
igroup create -vserver Infra-MS-SVM -igroup VM-Host-Infra-All -protocol fcp -ostype windows -
initiator <hyper-v-mgmt-host-01-wwpna>,<hyper-v-mgmt-host-01-wwpnb>,<hyper-v-mgmt-host-02-
wwpna>,<hyper-v-mgmt-host-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 Hyper-V-MGMT-01 -igroup Hyper-V-MGMT-Host-01 -lun-
id 0

lun map -vserver Infra-MS-SVM -volume HV_boot -lun Hyper-V-MGMT-02 -igroup Hyper-V-MGMT-Host-02 -lun-
id 0

lun map -vserver Infra-MS-SVM -volume witness -lun witness -igroup VM-Host-Infra-All -lun-id 1
```



For additional storage related tasks, please see the storage configuration portion of this document.



FC storage in the tenant SVM can be mapped with NetApp SnapDrive.

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.

To create a backup using the Cisco UCS Manager GUI, complete the following steps:

1. Select Admin within the Navigation pane and select All.
2. Click 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]

The screenshot displays the UCS Manager interface for configuring backup policies. The left sidebar shows the navigation menu with categories like Equipment, Servers, LAN, SAN, VM, Storage, Chassis, and Admin. The main content area is titled 'All' and contains a 'Policy Backup & Export' tab. Under this tab, there are three sections: 'Full State Backup Policy', 'All Configuration Backup Policy', and 'Backup/Export Config Reminder'. Each policy section includes fields for Hostname, Protocol (FTP, TFTP, SCP, SFTP), User, Password, Remote File, Admin State (Disable/Enable), Schedule (Daily, Weekly, Bi Weekly), and Max Files. The 'Full State Backup Policy' is configured with Hostname 10.1.156.150, Protocol SCP, User root, Password masked, Remote File /var/www/html/bears/configs/ucs/6332.full, Admin State Enable, Schedule Bi Weekly, and Max Files 0. The 'All Configuration Backup Policy' is configured with Hostname 10.1.156.150, Protocol SCP, User root, Password masked, Remote File /var/www/html/bears/configs/ucs/6332.config, Admin State Enable, Schedule Daily, and Max Files 0. The 'Backup/Export Config Reminder' section has Admin State set to Disable. At the bottom right, there are 'Save Changes' and 'Reset Values' buttons.

4. Click Save Changes to create the Policy.

About the Authors

John George, Technical Marketing Engineer, Data Center Solutions Engineering, Cisco Systems, Inc.

John has over seven years of experience in designing, developing, validating, and supporting the FlexPod Converged Infrastructure. Before his roles with FlexPod, he supported and administered a large worldwide training network and VPN infrastructure. John holds a Master's degree in Computer Engineering from Clemson University.

Dave Derry, Technical Marketing Engineer, Infrastructure and Cloud Engineering, NetApp

Dave Derry is a Technical Marketing Engineer in the Converged Infrastructure Engineering team at NetApp. He has been with NetApp since 2012, serving in a variety of engineering roles. Prior to that, he was an engineer at Cisco Systems for over ten years, in a variety of development and solution test roles.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Haseeb Niazi, Cisco Systems, Inc.
- Sanjeev Naldurgkar, Cisco Systems, Inc.
- Rajendra Yogendra, Cisco Systems, Inc.
- Melissa Palmer, NetApp
- Lindsey Street, NetApp