

FlashStack Mini with Pure Storage FlashArray//M10 for 1250 Citrix XenDesktop Users

Cisco UCS B200 M5 Blade Servers with Pure Storage FlashArray//M10 on Citrix XenDesktop 7.15 and Hyper-V 2016

Updated: March 2, 2020

Published: June 18, 2018



Partnered with:

About the Cisco Validated Design Program

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

© 2020 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	8
Solution Overview.....	9
Introduction	9
Audience	9
Purpose of this Document.....	9
What's New?	9
Solution Design.....	10
Architecture.....	10
Physical Topology.....	11
Deployment Hardware and Software	12
Software Revisions	12
Configuration Guidelines.....	12
Physical Infrastructure.....	13
FlashStack Cabling	13
Infrastructure Servers Prerequisites	15
Active Directory DC/DNS	15
Microsoft System Center 2016	15
Network Switch Configuration.....	16
Physical Connectivity	16
FlashStack Cisco Nexus Base	16
Set Up Initial Configuration	16
FlashStack Cisco Nexus Switch Configuration	18
Enable Licenses.....	18
Set Global Configurations	19
Create VLANs.....	19
Add NTP Distribution Interface.....	20
Add Individual Port Descriptions for Troubleshooting.....	20
Create Port Channels.....	22
Configure Port Channel Parameters	23
Configure Virtual Port Channels	24
Uplink into Existing Network Infrastructure	26
Storage Configuration.....	27
Pure Storage All Flash //M10 Controllers	27
Controllers.....	27

Disk Shelves	27
Server Configuration	28
Cisco UCS Base Configuration.....	28
Perform Initial Setup	28
Cisco UCS Setup	30
Log in to Cisco UCS Manager	30
Upgrade Cisco UCS Manager Software to Version 3.2(2b)	30
Anonymous Reporting	30
Configure Cisco UCS Call Home.....	31
Configure Unified Ports.....	31
Add Block of IP Addresses for KVM Access	33
Synchronize Cisco UCS to NTP.....	33
Edit Chassis Discovery Policy	35
Enable Uplink Ports.....	35
Acknowledge Cisco UCS Chassis and FEX	36
Create Uplink Port Channels to Cisco Nexus Switches	36
Create a WWNN Pool for FC Boot.....	37
Create WWPN Pools	39
Configuring an FC Storage Port	42
Before You Begin.....	42
Create VSANs and Enable FC Zoning	44
Create vHBA Templates	46
Create MAC Address Pools	48
Create UUID Suffix Pool.....	50
Create Server Pool	50
Create VLANs	51
Modify Default Host Firmware Package	54
Set Jumbo Frames in Cisco UCS Fabric.....	55
Create Local Disk Configuration Policy (Optional)	56
Create Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP)	57
Create Power Control Policy.....	58
Update the Default Maintenance Policy.....	61
Create vNIC Templates.....	62
Create FC Boot Policy.....	65
Create Service Profile Templates.....	68

Create Service Profiles	75
Add More Servers to FlashStack Unit.....	76
Storage Configuration – Boot LUNs.....	77
Pure Boot Storage Setup	77
Create Hosts	77
Create Volumes and Attach to Hosts on the Boot	80
Microsoft Windows Server 2016 Hyper-V Deployment Procedure	82
Setting Up Microsoft Windows Server 2016	82
Install Chipset and Windows eNIC Drivers	85
Install Windows Roles and Features.....	87
Configuring MPIO for Pure Flash Array//M10.....	89
Multipath-IO Timers.....	91
Updating MPIO Timers using Windows PowerShell.....	91
Update MPIO Timer Values.....	92
Setup using Windows Disk Management	93
Host Renaming and Join to Domain	94
Storage Configuration – Boot LUNs.....	95
Pure Boot Storage Setup	95
Deploying and Managing Hyper-V Clusters using System Center 2016 VMM	96
Settings	96
Configuring Network Settings	96
Create Run As Account in VMM	97
Fabric – Servers - I	97
Create Host Groups.....	98
Add Hosts to the Host Group.....	98
Fabric – Networking	100
Creating Logical Networks, Sites, and IP Pools	101
Create VM Networks.....	106
Create Uplink Port Profiles and Hyper-V Port Profiles.....	108
Create Logical Switch using SET	109
Fabric – Storage	113
Pure Storage SMI-S Provider Configuration.....	113
Fabric – Servers - II	119
Configure Network on Host – Applying Logical Switch.....	119
Deploy Hyper-V Cluster.....	122

Deploy Volumes to the Hyper-V Cluster	129
Cisco UCS Management Pack Suite Installation and Configuration.....	135
Cisco UCS Manager Integration with SCOM	135
About Cisco UCS Management Pack Suite	135
Installing Cisco UCS Monitoring Service	135
Adding a Firewall Exception for the Cisco UCS Monitoring Service.....	137
Installing the Cisco UCS Management Pack Suite.....	137
Adding a Cisco UCS Domains to the Operations Manager	140
Cisco UCS Manager Monitoring Dashboards	143
Cisco UCS Manager Plug-in for SCVMM	148
Cisco UCS Manager Plug-in Installation.....	148
Cisco UCS Domain Registration.....	149
Using the Cisco UCS SCVMM Plugin	151
Viewing the Server Details from the Hypervisor Host View.....	151
Viewing Registered UCS Domains	153
Viewing the UCS Blade Server Details	153
Viewing the Service Profile Details.....	154
Viewing the Service Profile Template Details	156
Viewing the Host Firmware Package Details	157
Building the Virtual Machines and Environment for Workload Testing.....	158
Software Infrastructure Configuration	158
Preparing the Master Targets.....	159
Installing and Configuring XenDesktop and XenApp	160
Prerequisites	160
Install XenDesktop Delivery Controller, Citrix Licensing, and StoreFront.....	160
Installing Citrix Licenses	164
Configure the XenDesktop Site.....	166
Additional XenDesktop Controller Configuration	169
Add the Second Delivery Controller to the XenDesktop Site	171
Create Host Connections with Citrix Studio	173
Configuring StoreFront	175
Create Machine Catalogs.....	178
Create Delivery Groups.....	180
Citrix XenDesktop Policies and Profile Management	187
Configure Citrix XenDesktop Policies.....	187

Configuring User Profile Management	188
245 User Single Server Testing on Cisco UCS B200 M5 Server.....	201
1250 User Full Scale Testing on 6-node Hyper-V Cluster.....	204
Pure Storage FlashArray//M10 Test Results for 1250 Persistent Windows 10 x64 MCS Desktops.....	207
Scalability Considerations and Guidelines	209
Cisco UCS System Scalability.....	209
Scalability of Citrix XenDesktop 7.15 Configuration	209
FlashStack Backups.....	211
Cisco UCS Backup.....	211
Cisco Nexus Backups	212
Pure Storage Snapshots for Array Protection.....	214
References	217
Cisco UCS B-Series Servers.....	217
Cisco UCS Manager Configuration Guides.....	217
Citrix References	217
Microsoft References.....	217
Login VSI Documentation.....	217
Pure Storage Reference Documents	217
About the Authors.....	219
Acknowledgments	219



Executive Summary

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

This CVD describes the Cisco and Pure Storage® FlashStack Datacenter with Cisco UCS Manager unified software release 3.2(2b) and Microsoft Hyper-V 2016. Cisco UCS Manager (UCSM) 3.1 provides consolidated support for all the current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlashStack Datacenter with Cisco UCS unified software release 3.2(2b), and Microsoft Hyper-V 2016 is a predesigned, best-practice data center architecture built on Cisco Unified Computing System (UCS), Cisco Nexus® 9000 family of switches, MDS 9000 multilayer fabric switches, and Pure Storage.

This document primarily focuses on deploying Microsoft Hyper-V 2016 Cluster on FlashStack Datacenter using Fibre Channel and SMB storage protocols.

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 Pure Storage have partnered to deliver FlashStack, 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 FlashStack Datacenter with Cisco UCS Fabric Interconnects, Pure Storage /M10, and Cisco Nexus 9000 solution. This document primarily focuses on deploying Microsoft Hyper-V 2016 Cluster on FlashStack Datacenter using Fibre Channel protocols.

What's New?

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

- Support for the Cisco UCS 3.2(2b) unified software release and Cisco UCS B200-M5 servers
- Support for the latest release of Pure Storage PURITY4.10.4 fibre channel storage design
- Validation of Microsoft Hyper-V 2016

Solution Design

Architecture


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

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

Figure 1 illustrates the Microsoft Hyper-V built on FlashStack components and its physical cabling with the Cisco UCS 6324 Fabric Interconnects. This design has end-to-end 8 Gb FC connections from Cisco UCS 5108 Blade Chassis to Pure Storage //M10. This infrastructure is directly connected without the need for upstream FC switching. Networking is uplinked to a pair of Cisco 9000 switches for 40GB of network bandwidth.

Physical Topology

Figure 1 FlashStack with Cisco UCS 6324 Fabric Interconnects



The reference 40Gb based hardware configuration includes:

- Two Cisco 93180YC-EX switches
- Two Cisco UCS 6324 fabric interconnects
- One chassis of Cisco UCS blade servers
- One Pure Storage //M10 (HA pair) running PURITY with disk shelves and solid state drives (SSD)

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 6324 Series. UCS B-200 M5 	<ul style="list-style-type: none"> 3.2(2b) - Infrastructure Bundle 3.2(2b) - Server Bundle 	UCS VIC 1340
Network	Cisco Nexus 9000 NX-OS	7.0(3)I4(5)	
Storage	Pure Storage //M10	PURITY 4.10.4	
Software	Cisco UCS Manager	3.2(2b)	
	Microsoft System Center Virtual Machine Manager	2016 (version: 4.0.2051.0)	
	Microsoft Hyper-V	2016	
	Citrix XenDesktop	7.15 LTSR	

Configuration Guidelines

This document provides details on configuring a fully redundant, highly available reference model for a FlashStack unit with Pure Storage PURITY operating environment. 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 Pure Storage //M10 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.

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

necessary for deployment as outlined in this guide. Table 2 describes the VLANs necessary for deployment as outlined in this guide.

Table 2 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out-of-Band-Mgmt	VLAN for out-of-band management interfaces	132
MS-IB-MGMT	VLAN for in-band management interfaces	20
Native-VLAN	VLAN to which untagged frames are assigned	1
MS-LVMN-VLAN	VLAN designated for the movement of VMs from one physical host to another.	29
MS-Cluster-VLAN	VLAN for cluster connectivity	28
MS-Tenant-VM-VLAN	VLAN for Production VM Interfaces	21

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

Table 3 Virtual Machines

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

Physical Infrastructure

FlashStack Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlashStack environment. To simplify cabling requirements, 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 Pure Storage //M10 running Pure Storage PURITY 4.10.4. Future upgrade releases of Purity will not require re-cabling unless a new feature outside of the scope of this document requires it.



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.

Infrastructure Servers Prerequisites

Active Directory DC/DNS

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

Microsoft System Center 2016



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

Follow the Microsoft guidelines to install SCOM and SCVMM 2016:

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

Network Switch Configuration

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

Physical Connectivity

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

FlashStack Cisco Nexus Base

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlashStack environment. This procedure assumes the use of Cisco Nexus 9000 7.0(3)I4(5) and is valid for both the Cisco 93180YC-EX switches deployed with the 40Gb end-to-end topology, and the Cisco Nexus 93180YC-EX switches used in the 10Gb based topology.



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

Set Up Initial Configuration

Cisco Nexus A

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

1. Configure the switch.



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

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

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

Enter the password for " admin" : <password>

Confirm the password for " admin" : <password>

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

NTP server IPv4 address: <global-ntp-server-ip>

Configure default interface layer (L3/L2) [L3]: L2

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

Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]: Enter

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

2. Review the configuration summary before enabling the configuration.

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

Cisco Nexus B

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

1. Configure the switch.



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

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

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

Enter the password for " admin" : <password>

Confirm the password for " admin" : <password>

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

2. Review the configuration summary before enabling the configuration.

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

FlashStack Cisco Nexus Switch Configuration

Enable Licenses

Cisco Nexus A and Cisco Nexus B

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

1. Log in as admin.

2. Run the following commands:

```
config t
```

```
feature interface-vlan
```

```
feature lacp
```

```
feature vpc
```

```
feature lldp
```

```
feature nxapi
```

Set Global Configurations

Cisco Nexus A and Cisco Nexus B

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

1. Run the following commands to set global configurations:

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

```
spanning-tree port type edge bpduguard default
```

```
spanning-tree port type edge bpdufilter default
```

```
port-channel load-balance src-dst l4port
```

```
ntp server <global-ntp-server-ip> use-vrf management
```

```
ntp master 3
```

```
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
```

```
copy run start
```

Create VLANs

Cisco Nexus A and Cisco Nexus B

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

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

```
vlan <ms-ib-mgmt-vlan-id>
```

```
name MS-IB-MGMT-VLAN
```

```
vlan <native-vlan-id>
```

```
name Native-VLAN
```

```

vlan < ms-lvmn-vlan-id>
name MS-LVMN-VLAN
vlan <ms-tenant-vm-vlan-id>
name MS-Tenant-VM-VLAN
vlan <ms-cluster-vlan-id>
name MS-Cluster-VLAN
exit

```

Add NTP Distribution Interface

Cisco Nexus A

From the global configuration mode, run the following commands:

```

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

```

Cisco Nexus B

From the global configuration mode, run the following commands:

```

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

```

Add Individual Port Descriptions for Troubleshooting

Cisco Nexus A

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



In this step and in the later sections, configure the Pure //M10 nodename <st-node> and Cisco UCS 6324 fabric interconnect clustername <ucs-clustername> interfaces as appropriate to your deployment.

From the global configuration mode, run the following commands:

```
interface Eth1/3
description <st-node>-1:e2a
interface Eth1/4
description <st-node>-2:e2a
interface Eth1/25
description <ucs-clustername>-a:1/27
interface Eth1/26
description <ucs-clustername>-b:1/27
interface Eth1/27
description <nexus-hostname>-b:1/27
interface Eth1/28
description <nexus-hostname>-b:1/28
exit
```

Cisco Nexus B

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

From the global configuration mode, run the following commands:

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

```
description <nexus-hostname>-a:1/28  
exit
```

Create Port Channels

Cisco Nexus A and Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
interface Po10  
description vPC peer-link  
interface Eth1/27-28  
channel-group 10 mode active  
no shutdown  
interface Po13  
description <st-node>-1  
interface Eth1/3  
channel-group 13 mode active  
no shutdown  
interface Po14  
description <st-node>-2  
interface Eth1/4  
channel-group 14 mode active  
no shutdown  
interface Po125  
description <ucs-clustername>-a  
interface Eth1/25  
channel-group 125 mode active  
no shutdown  
interface Po126  
description <ucs-clustername>-b
```

```
interface Eth1/26
channel-group 126 mode active
no shutdown
exit
copy run start
```

Configure Port Channel Parameters

Cisco Nexus A and Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
interface Po10
switchport mode trunk
switchport trunk native vlan 1
switchport trunk allowed vlan <ib-mgmt-vlan-id>, <LiveMigration-vlan-id>, <vm-traffic-vlan-id>,
<infra-ClusterComm-id>,
spanning-tree port type network
interface Po13
switchport mode trunk
switchport trunk native vlan 1
switchport trunk allowed vlan <ib-mgmt-vlan-id>,
spanning-tree port type edge trunk
mtu 9216
interface Po14
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan <ib-mgmt-vlan-id>,
spanning-tree port type edge trunk
mtu 9216
interface Po125
switchport mode trunk
```

```
switchport trunk native vlan 2

switchport trunk allowed vlan <ib-mgmt-vlan-id>, <LiveMigration-vlan-id>, <vm-traffic-vlan-id>,
<infra-ClusterComm-id>

spanning-tree port type edge trunk

mtu 9216

interface Po126

switchport mode trunk

switchport trunk native vlan 2

switchport trunk allowed vlan <ib-mgmt-vlan-id>, <LiveMigration-vlan-id>, <vm-traffic-vlan-id>,
<infra-ClusterComm-id>

spanning-tree port type edge trunk

mtu 9216

exit

copy run start
```

Configure Virtual Port Channels

Cisco Nexus A

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

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>

role priority 10

peer-keepalive destination <nexus-B-mgmt0-ip> source <nexus-A-mgmt0-ip>

peer-switch

peer-gateway

auto-recovery

delay restore 150

interface Po10

vpc peer-link

interface Po13

vpc 13
```



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

Cisco Nexus B

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

From the global configuration mode, run the following commands:

```
vpc domain <nexus-vpc-domain-id>
role priority 20
peer-keepalive destination <nexus-A-mgmt0-ip> source <nexus-B-mgmt0-ip>
peer-switch
peer-gateway
auto-recovery
delay restore 150
interface Po10
vpc peer-link
interface Po13
vpc 13
interface Po14
vpc 14
interface Po125
vpc 125
interface Po126
vpc 126
exit
```

copy run start

Uplink into Existing Network Infrastructure

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

Storage Configuration

Pure Storage All Flash //M10 Controllers

Controllers



Pure Storage installation engineers or certified partners will conduct the physical installation and initial configuration of the array, up until the management UI is available on the network.

The implementation engineer will complete the following steps:

1. Rack the array.
2. Install power cables on redundant PDUs. (Controller A to a separate PDU than Controller B).
3. Connect 1GB Ethernet cables to the management interfaces on both controllers. (preferably to separate switches for redundancy).
4. A Pure Storage implementation engineer will console into the array locally and configure the IP addresses on the management interfaces, including a VIP, enable ODX, connect the array to Pure1 manage for remote support connectivity and monitoring (if desired) and if necessary, upgrade to a later required version of the Purity Operating Environment.
5. At this point, the web UI will be available to begin volume creation and host initiator management.

Disk Shelves

The Pure Storage //M10 array is targeted for smaller customer deployments and sites at low cost. As such it is either available in 5TB raw or 10TB raw capacity configurations, both of which fit in the 3U chassis. Customers who need further capacity or performance scaling can upgrade completely non-disruptively to the //M20 controller or above without any data migration or I/O interruption all within the same chassis. More dense capacity packs can be installed in the front of the chassis or via a 2U expansion shelf and 100% trade-in credit is given for both controllers as well as capacity. The complete list of controller and capacity configuration options that are supported by the FlashArray //M series is available [here](#).

Server Configuration

Cisco UCS Base Configuration

This FlashStack deployment will show configuration steps for the Cisco UCS 6324 Fabric Interconnects (FI) in a design that will support Fibre Channel to the Pure Storage through Cisco UCS.

Perform Initial Setup

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlashStack 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 FlashStack environment, complete the following steps:

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

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

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

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

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

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

Basic Settings

Cluster and Fabric setup

- Enable clustering
- Standalone mode
- Synchronize

Fabric Setup: Fabric A Fabric B

 IPv4

 IPv6

Virtual IP Address: . . .

System setup

- Enforce strong password?:** Yes No
- System name:**
- Admin Password:** **Confirm Admin password:**
- Mgmt IP Address:** . . . **Mgmt IP Netmask:** . . .
- Default Gateway:** . . .
- DNS Server IP:** . . . **Domain Name :**

UCS Central managed environment

- UCS Central IP:** . . . **Shared Secret:**

- 1.
7. Click Submit.

Cisco UCS Fabric Interconnect B

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

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

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

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

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

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

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

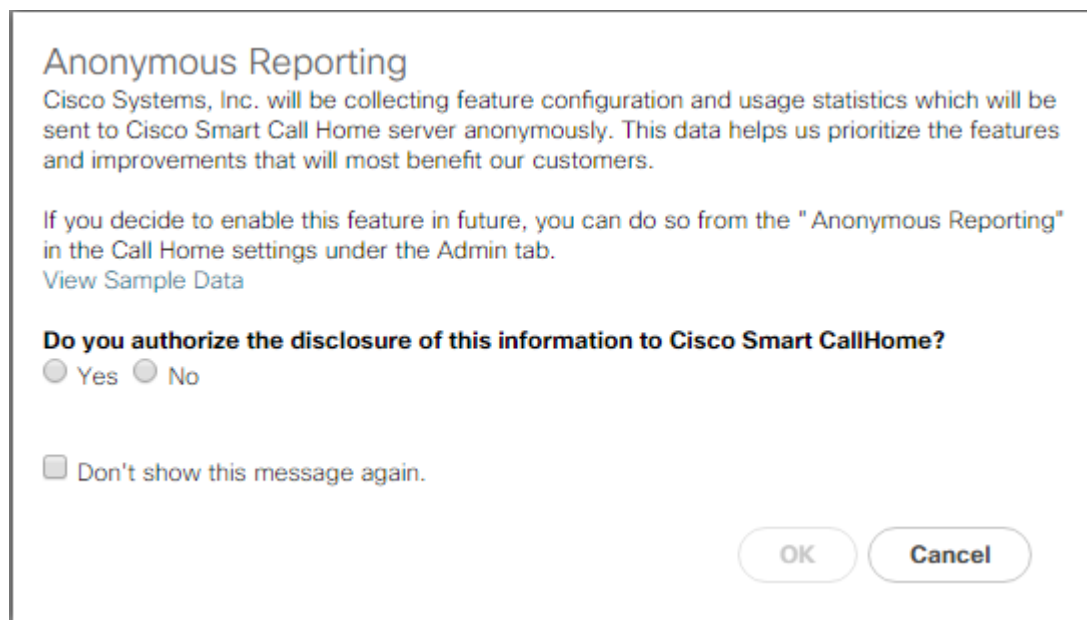
Upgrade Cisco UCS Manager Software to Version 3.2(2b)

This document assumes the use of Cisco UCS 3.2(2b). To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 3.2(2b), refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Anonymous Reporting

To create anonymous reporting, complete the following steps:

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.
2. Click OK.



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 Admin on the left.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Configure Unified Ports


Fiber Channel port configurations differ slightly between the 6324 and the 6248UP Fabric Interconnects. Both Fabric Interconnects have a slider mechanism within the Cisco UCS Manager GUI interface, but the fiber channel port selection options for the 6324 are from the first 4 ports starting from the first port on the left and configured in increments of the first 2 of the unified ports. With the 6248UP, the port selection options will start from the right of the 32 fixed ports, or the right of the 16 ports of the expansion module, going down in contiguous increments of 2.

To enable the fiber channel ports, complete the following steps for the 6324:

1. In Cisco UCS Manager, click Equipment on the left.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary).
3. Select Configure Unified Ports.

- Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
- Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select 2 ports to be set as FC Uplinks.

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
FC Port 1	fc	FC Storage	
FC Port 2	fc	FC Storage	
Port 3	ether	Ethernet Uplink Port Channel Member	
Port 4	ether	Ethernet Uplink Port Channel Member	

■ Up
 ■ Admin Down
 ■ Fail
 ■ Link Down



- Click OK, then click Yes, then click OK to continue.
- Select Equipment > Fabric Interconnects > Fabric Interconnect B (primary).
- Select Configure Unified Ports.
- Click Yes on the pop-up window warning that changes to the fixed module will require a reboot of the fabric interconnect and changes to the expansion module will require a reboot of that module.
- Within the Configured Fixed Ports pop-up window move the gray slider bar from the left to the right to select either 2, or 4 ports to be set as FC Uplinks.

11. Click OK, then click Yes, then click OK to continue.
12. Wait for both the Fabric Interconnects to reboot.
13. Log back into UCS Manager.

Add Block of IP Addresses for KVM Access

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

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



The screenshot shows a dialog box titled "Create Block of IPv4 Addresses". It has a question mark icon and a close button (X) in the top right corner. The dialog contains the following fields:

From :	<input type="text" value="10.29.132.124"/>	Size :	<input type="text" value="8"/>
Subnet Mask :	<input type="text" value="255.255.255.0"/>	Default Gateway :	<input type="text" value="10.29.132.1"/>
Primary DNS :	<input type="text" value="0.0.0.0"/>	Secondary DNS :	<input type="text" value="0.0.0.0"/>

At the bottom right, there are two buttons: "OK" (highlighted in blue) and "Cancel".

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

Synchronize Cisco UCS to NTP

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

1. In Cisco UCS Manager, click Admin on the left.
2. Expand All > Time Zone Management.
3. Select Timezone.

4. In the Properties pane, select the appropriate time zone in the Timezone menu.
5. Click Save Changes, and then click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> and click OK. Click OK on the confirmation.

Add NTP Server



NTP Server :



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

The screenshot shows the 'Timezone Management' interface. The breadcrumb path is 'All / Time Zone Management / Timezone'. There are two tabs: 'General' (selected) and 'Events'. On the left, under 'Actions', there is a link for 'Add NTP Server'. On the right, the 'Properties' pane shows 'Time Zone : America/Los_Angeles (Pacif)' with a dropdown arrow. Below this is a table titled 'NTP Servers' with columns for 'Name'. The table contains two entries: 'NTP Server 10.29.164.131' and 'NTP Server 10.29.164.132'. Above the table are icons for 'Advanced Filter', 'Export', and 'Print'. At the bottom right of the table are icons for '+ Add', 'Delete', and 'Info'.

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, click the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel. If Backplane Speed Preference appears, leave it set at 40G. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.

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

- Action: 2 Link
- Link Grouping Preference: None Port Channel
- Multicast Hardware Hash: Disabled Enabled

5. Click Save Changes.
6. Click OK.

Enable Uplink Ports

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

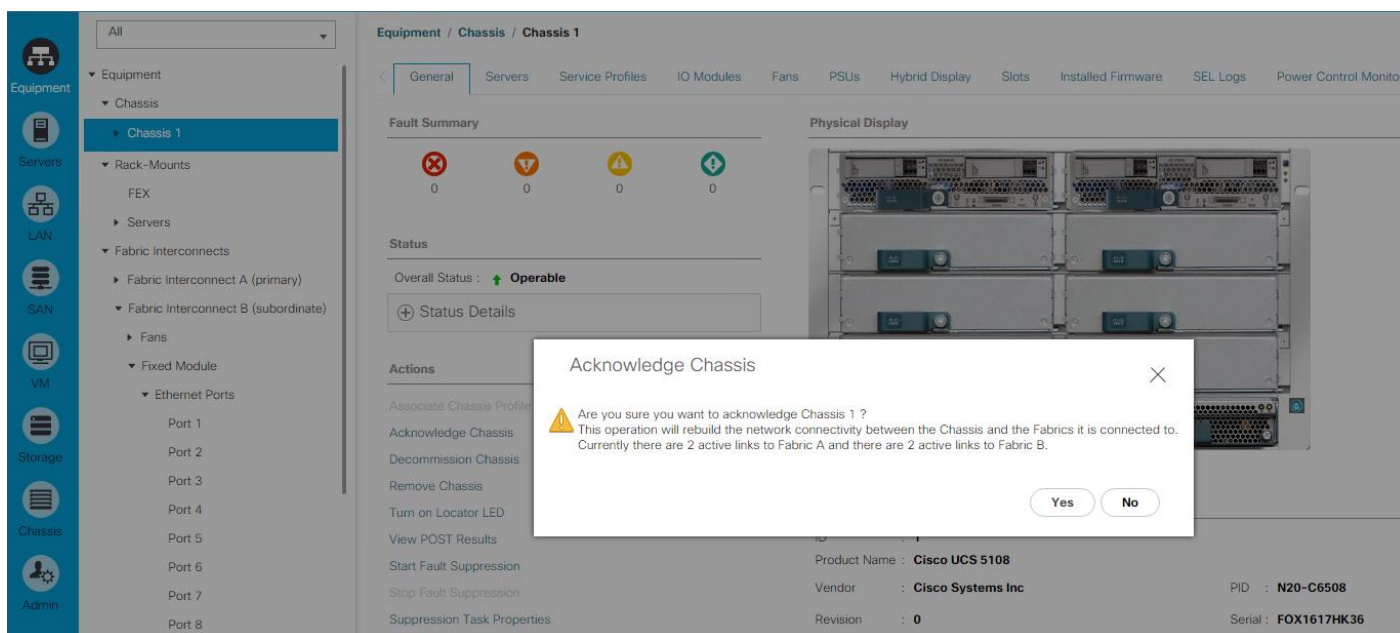
1. In Cisco UCS Manager, click Equipment on the left.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
5. Click Yes to confirm uplink ports and click OK.
6. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.
7. Expand Ethernet Ports.

8. Select the ports that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.
9. Click Yes to confirm the uplink ports and click OK.

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis, complete the following steps:

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



4. Click Yes and then click OK to complete acknowledging the chassis.

Create Uplink Port Channels to Cisco Nexus Switches

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

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



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

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.

5. Enter 11 as the unique ID of the port channel.
6. Enter Port-Channel11 as the name of the port channel.
7. Click Next.
8. Select the ports connected to the Nexus switches to be added to the port channel:
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel.
11. Click OK.
12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 12 as the unique ID of the port channel.
16. Enter Port-Channel12 as the name of the port channel.
17. Click Next.
18. Select the ports connected to the Nexus switches to be added to the port channel:
19. Click >> to add the ports to the port channel.
20. Click Finish to create the port channel.
21. Click OK.

Create a WWNN Pool for FC Boot

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

1. Select SAN on the left.
2. Select Pools > root.
3. Right-click WWNN Pools under the root organization.
4. Select Create WWNN Pool to create the WWNN pool.
5. Enter WWNN-POOL for the name of the WWNN pool.
6. Optional: Enter a description for the WWNN pool.

7. Select Sequential for Assignment Order.

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

- Name:** WWNN-POOL
- Description:** (empty text box)
- Assignment Order:** Default Sequential

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

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

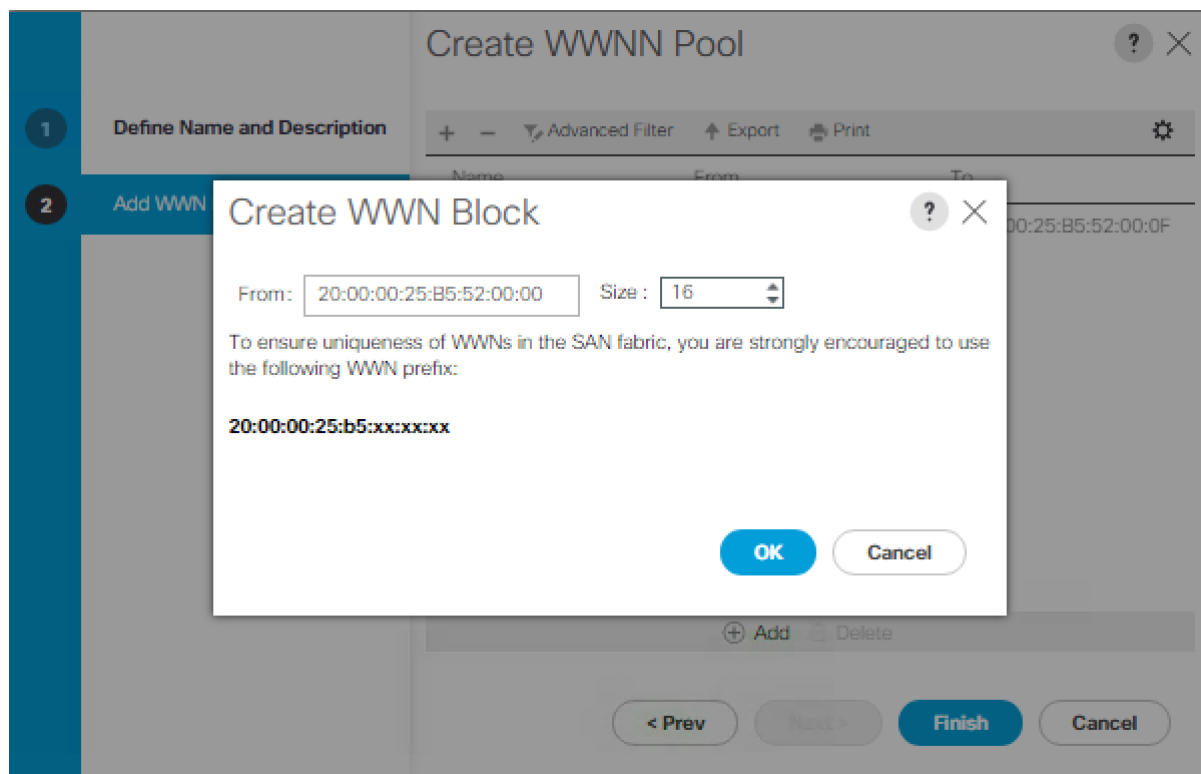


Modifications of the WWNN block, as well as the WWPN and MAC Addresses, can convey identifying information for the UCS domain. Within the From field in our example, the 6th octet was changed from 00 to 52 to represent as identifying information for this being in the UCS 6324 in the 4th cabinet.



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

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



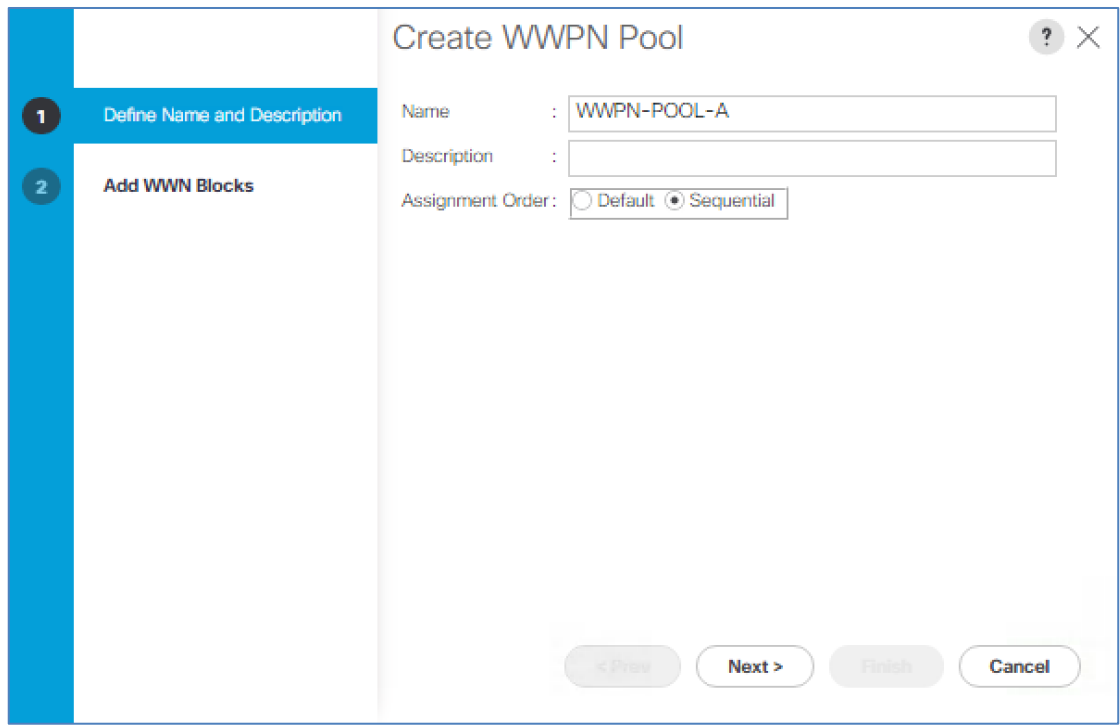
12. Click OK.

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

Create WWPN Pools

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

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



9. Click Next.

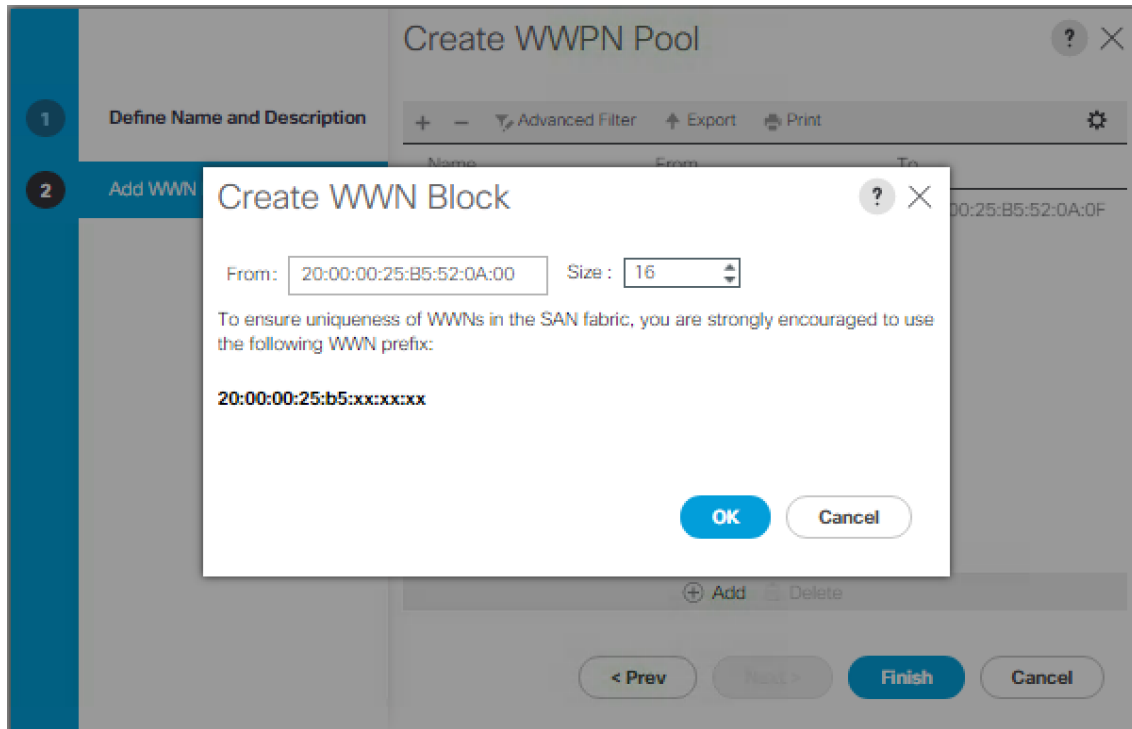
10. Click Add.

11. Specify a starting WWPN.



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

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



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



For the FlashStack solution, the recommendation is to place B 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

Configuring an FC Storage Port

This task describes only one method of configuring FC storage ports. You can also configure FCoE storage ports from the General tab for the port.

Before You Begin

The Fibre Channel switching mode must be set to Switching for these ports to be valid. The storage ports cannot function in end-host mode.

To use the Switching mode, complete the following steps:

1. In the Navigation pane, click Equipment.
2. Expand Equipment > Fabric Interconnects > Fabric_Interconnect_Name.
3. Click one or more of the ports under the FC Ports node.
4. Right-click the selected port or ports and choose Configure as FC Storage Port.
5. If a confirmation dialog box displays, click Yes.
6. Click OK.

Equipment / Fabric Interconnects / Fabric Interconnect A(primary)... / Fixed Module / FC Ports / FC Port 1

General | Faults | Events | FSM | Statistics

Fault Summary

0 1 0 0


Status

Overall Status : **Failed**
Additional Info : **Offline**
Admin State : **Enabled**

Actions

Enable Port
Disable Port
Configure as Uplink Port
Configure as FC Storage Port
Show Interface

Physical Display



Up Admin Down Fail Link Down

Properties

ID	: 1	Slot ID	: 1
User Label	:	Mode	: E
WWPN	: 20:01:F8:C2:88:D9:A5:80	Negotiated Speed	: Indeterminate
Port Type	: Physical		
VSAN	: Fabric []		

Transceiver

Type	: Sfp
Model	: FTLF8528P2BCV-CS
Vendor	: CISCO-FINISAR
Serial	: FNS15300AUX

License Details

License State	: License OK
License Grace Period	: 0

General | Faults | Events | FSM | Statistics

Fault Summary

✘
0

⚠
0

✔
0

Status

Overall Status : ✔ **Up**

Additional Info:

Admin State : **Enabled**

Actions

[Enable Port](#)


[Disable Port](#)

[Configure as Uplink Port](#)

[Configure as FC Storage Port](#)

[Show Interface](#)

Physical Display



■ Up
 ■ Admin Down
 ■ Fail
 ■ Link Down

Properties

ID : 1	Slot ID : 1
User Label :	
WWPN : 20:01:F8:C2:88:D9:A5:80	Mode : F
Port Type : Physical	Negotiated Speed : 8 Gbps
VSAN : Fabric L	

Transceiver

Type : **Sfp**

Model : **FTLF8528P2BCV-CS**

Vendor : **CISCO-FINISAR**

Serial : **FNS15300AUX**

License Details

License State : **License OK**

License Grace Period : **0**

Create VSANs and Enable FC Zoning

To create the necessary VSAN's for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Expand Storage Cloud.
3. Expand Fabric A.
4. Expand VSANs.
5. Click the + Add button to Add our VSAN for fabric A.

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 A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: ←

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN:

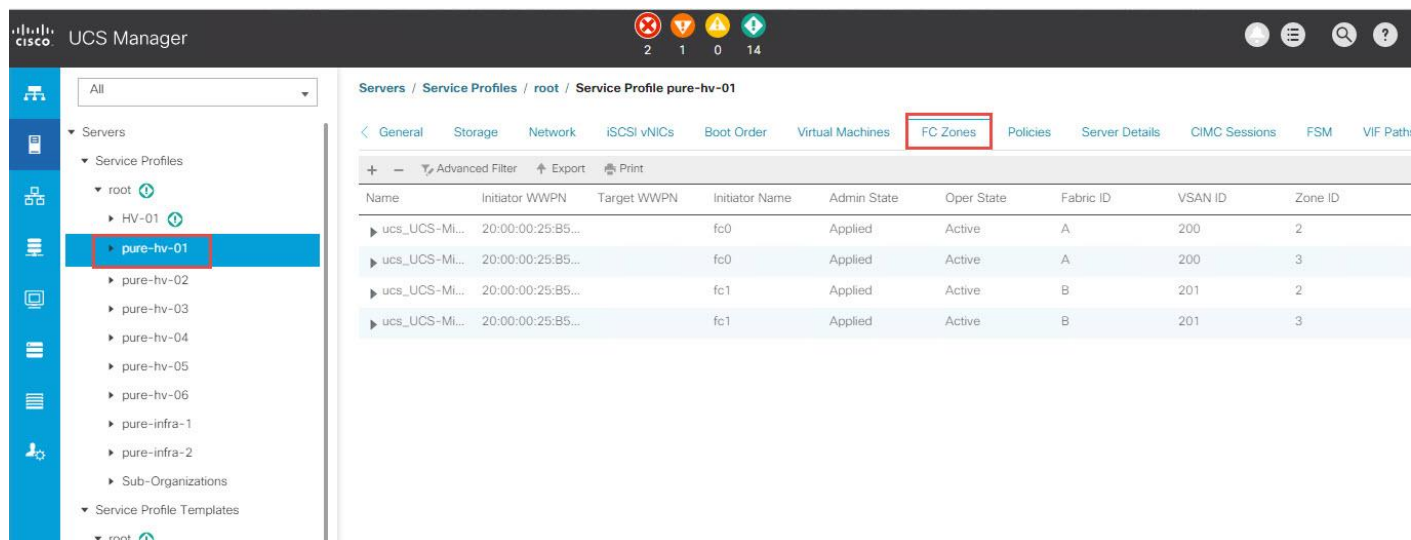
6. For 'FC Zoning' select 'Enabled'.

7. Repeat the steps for Fabric B and ensure a unique VSAN ID from Fabric A.



In this study we used VSAN 200 for A and 201 for B.

When this step and the SAN Boot policy are completed, FC Zones will be created when service profiles are deployed.



Create vHBA Templates

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

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

Create vHBA Template

Name : vHBA-Template-A

Description :

Fabric ID : A B

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Select VSAN : VSAN-A

Template Type : Initial Template Updating Template

Max Data Field Size : 2048

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

QoS Policy : <not set>

Pin Group : <not set>

Stats Threshold Policy : default

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

Create MAC Address Pools

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

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.



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

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



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

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

Create a Block of MAC Addresses ? X

First MAC Address : Size :

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

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



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

23. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources.
24. Click OK.

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

Create UUID Suffix Pool

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

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

Create Server Pool

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



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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.

4. Select Create Server Pool.
5. Enter `MS-Server-Pool1` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the Hyper-V management cluster and click >> to add them to the `MS-Server-Pool1` server pool.
9. Click Finish.
10. Click OK.

Create VLANs

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

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



In this procedure, five unique VLANs are created.

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

Create VLANs ? X

VLAN Name/Prefix :

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

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

VLAN IDs :

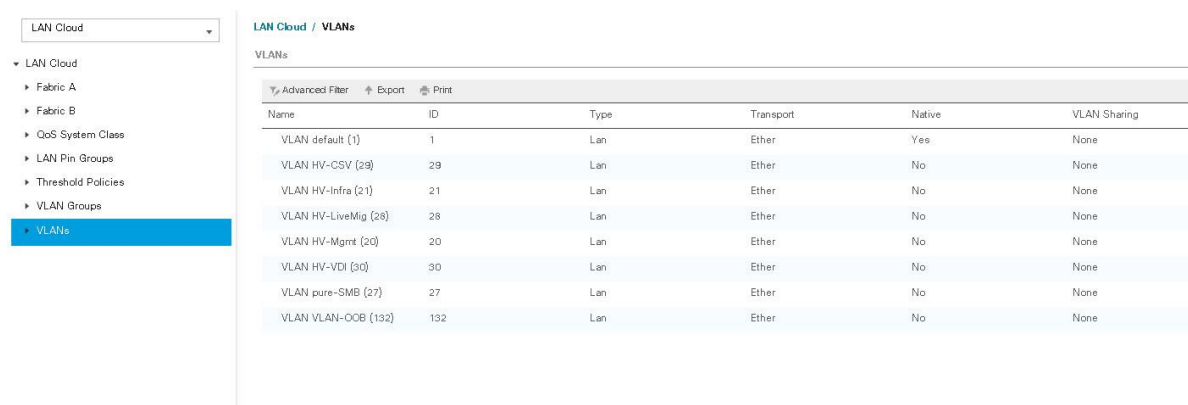
Sharing Type : None Primary Isolated Community

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

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

45. Keep the Sharing Type as None.

46. Click OK, and then click OK again.



LAN Cloud / VLANs

VLANs

Name	ID	Type	Transport	Native	VLAN Sharing
VLAN default (1)	1	Lan	Ether	Yes	None
VLAN HV-CSV (29)	29	Lan	Ether	No	None
VLAN HV-Infra (21)	21	Lan	Ether	No	None
VLAN HV-LiveMig (28)	28	Lan	Ether	No	None
VLAN HV-Mgmt (20)	20	Lan	Ether	No	None
VLAN HV-VDI (30)	30	Lan	Ether	No	None
VLAN pure-SMB (27)	27	Lan	Ether	No	None
VLAN VLAN-OCB (132)	132	Lan	Ether	No	None

Modify Default Host Firmware Package

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

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 3.2(2b) for both the Blade Package.

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

7. 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 LAN on the left.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.
5. Click Save Changes in the bottom of the window.

6. Click OK.

LAN Cloud / QoS System Class

General Events FSM

Actions Properties

Use Global Owner: Local

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

Create Local Disk Configuration Policy (Optional)

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



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

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

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

Create Local Disk Configuration Policy ? X

Name : SAN-Boot

Description :

Mode : No Local Storage

FlexFlash

FlexFlash State : Disable Enable

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

FlexFlash RAID Reporting State: Disable Enable

OK Cancel

8. Click OK.

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

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

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP as the policy name.
6. For CDP, select the Enabled option.

7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy.

Create Network Control Policy

CDP : Disabled Enabled

MAC Register Mode : Only Native Vlan All Host Vlans

Action on Uplink Fail : Link Down Warning

MAC Security

Forge : Allow Deny

LLDP

Transmit : Disabled Enabled

Receive : Disabled Enabled

OK **Cancel**

9. Click OK.

Create Power Control Policy

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

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

Create Power Control Policy ? ×

Name :

Description :

Fan Speed Policy :

Power Capping

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

No Cap cap

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

Create Server 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 tab in the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter B200-M5-BIOS as the BIOS policy name.
6. Configure the remaining BIOS policies as follows and click Finish.

Servers / Policies / root / BIOS Policies / B200-M5

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

Actions

Delete

Show Policy Usage

Use Global

Properties

Name : **B200-M5**

Description :

Owner : **Local**

Reboot on BIOS Settings Change :

Advanced Filter | Export | Print

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

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

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

Advanced Filter | Export | Print

BIOS Setting	Value
Altitude	Platform Default
CPU Hardware Power Management	Platform Default
Boot Performance Mode	Max Performance
CPU Performance	Enterprise
Core Multi Processing	All
DRAM Clock Throttling	Performance
Direct Cache Access	Enabled
Energy Performance Tuning	BIOS
Enhanced Intel SpeedStep Tech	Enabled
Execute Disable Bit	Enabled
Frequency Floor Override	Platform Default
Intel HyperThreading Tech	Enabled
Intel Turbo Boost Tech	Enabled
Intel Virtualization Technology	Enabled
Channel Interleaving	Platform Default
IMC Inteleave	Platform Default
Memory Interleaving	Platform Default
Rank Interleaving	Platform Default
Sub NUMA Clustering	Platform Default
Local X2 Apic	Platform Default
Max Variable MTRR Setting	Platform Default
P STATE Coordination	Platform Default
Package C State Limit	Platform Default
Processor C State	Disabled
Processor C1E	Disabled

+ Add | - Delete | i Info

Processor C3 Report	Disabled
Processor C6 Report	Enabled
Processor C7 Report	Disabled
Processor CMC1	Platform Default
Power Technology	Performance
Energy Performance	Performance
Adjacent Cache Line Prefetcher	Platform Default
DCU IP Prefetcher	Platform Default
DCU Streamer Prefetch	Platform Default
Hardware Prefetcher	Platform Default
UPI Prefetch	Platform Default
LLC Prefetch	Platform Default
XPT Prefetch	Platform Default
Demand Scrub	Platform Default
Patrol Scrub	Platform Default
Workload Configuration	Platform Default

+ Add - Delete i Info

Servers / Policies / root / BIOS Policies / B200-M5

Main **Advanced** Boot Options Server Management Events

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

Advanced Filter Export Print

BIOS Setting	Value
Intel VTD ATS support	Platform Default
Intel VTD coherency support	Platform Default
Intel VT for directed IO	Enabled
Intel VTD interrupt Remapping	Platform Default
Intel VTD pass through DMA support	Platform Default

7. Click Finish.

Update the Default Maintenance Policy

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

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

Servers / Policies / root / Maintenance Poli... / default

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

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

Create vNIC Templates

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

Create Infrastructure vNICs

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter HV_SET_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 HV-MGMT, HV-CSV, HV-LiveMig, HV-VDI, and HV-Infra VLANs.
13. Set Native-VLAN as the native VLAN.
14. Select vNIC Name for the CDN Source.
15. For MTU, enter 9000.
16. In the MAC Pool list, select MAC-POOL-A.
17. In the Network Control Policy list, select Enable-CDP-LLDP.

Create vNIC Template ? X

Name :

Description :

Fabric ID : Fabric A Fabric B Enable Failover

Redundancy

Redundancy Type : No Redundancy Primary Template Secondary Template

Peer Redundancy Template :

Target

Adapter

VM

Warning

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

Template Type : Initial Template Updating Template

VLANs | VLAN Groups

Advanced Filter | Export | Print

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	HV-LiveMig	<input type="radio"/>
<input checked="" type="checkbox"/>	HV-Mgmt	<input type="radio"/>
<input type="checkbox"/>	HV-VDI	<input type="radio"/>

Create VLAN

CDN Source : vNIC Name User Defined

MTU : 9000

MAC Pool : MAC_Pool_A(55/64) ▼

QoS Policy : <not set> ▼

Network Control Policy : CDP_Enabled ▼

Pin Group : <not set> ▼

Stats Threshold Policy : default ▼

Connection Policies

 Dynamic vNIC usNIC VMO

Dynamic vNIC Connection Policy : <not set> ▼

OK

Cancel

18. Click OK to create the vNIC template.

19. Click OK.

Create the secondary redundancy template Infra-B:

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

Create FC Boot Policy

This procedure applies to a Cisco UCS environment in which two Fibre Channel logical interfaces are on cluster node 1 (CT0.FC0 and FC1) and two Fibre Channel connections are on cluster node 2 (CT1.FC0 and FC1). Also, it is assumed that the CT0 are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the CT1 are connected to Fabric B (Cisco UCS Fabric Interconnect B).

One boot policy is configured in this procedure. The policy configures the primary target to be CT0-FC0.

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 `Pure_HV_Boot` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.



Do not select the Reboot on Boot Order Change checkbox.

7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Local CD/DVD`.
9. Expand the vHBAs drop-down menu and select Add SAN Boot.
10. Select the Primary for type field.
11. Enter `fc0` in vHBA field.

Add SAN Boot



vHBA :

Type : Primary Secondary Any

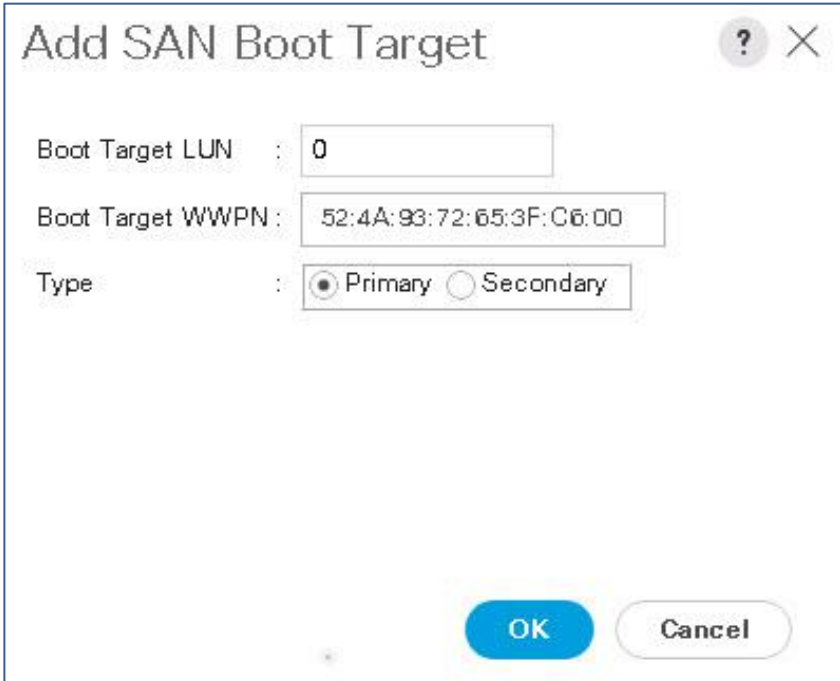
12. Click OK.
13. From the vHBA drop-down menu, select Add SAN Boot Target.
14. Keep 0 as the value for Boot Target LUN.
15. Enter the WWPN for CT0-FC0.



To obtain this information, log in to the Purity Interface and select System > Connections.

PORT	NAME	SPEED	FAILOVER	PORT	NAME	SPEED
CT0.FC0	52:4A:93:72:65:3F:C6:00	8 Gb/s		CT1.FC0	52:4A:93:72:65:3F:C6:10	8 Gb/s
CT0.FC1	52:4A:93:72:65:3F:C6:01	8 Gb/s		CT1.FC1	52:4A:93:72:65:3F:C6:11	8 Gb/s

16. Select Primary for the SAN boot target type.



The image shows a dialog box titled "Add SAN Boot Target". It contains three input fields: "Boot Target LUN" with the value "0", "Boot Target WWPN" with the value "52:4A:93:72:65:3F:C6:00", and "Type" with radio buttons for "Primary" (selected) and "Secondary". At the bottom right, there are "OK" and "Cancel" buttons.

17. Click OK to add the SAN boot target.
18. From the vHBA drop-down menu, select Add SAN Boot Target.
19. Enter 0 as the value for Boot Target LUN.
20. Enter the WWPN for CT0.FC1.
21. Click OK to add the SAN boot target.
22. From the vHBA drop-down menu, select Add SAN Boot.
23. In the Add SAN Boot dialog box, enter fc1 in the vHBA box.
24. The SAN boot type should automatically be set to Secondary.
25. Click OK to add the SAN boot.
26. From the vHBA drop-down menu, select Add SAN Boot Target.
27. Keep 0 as the value for Boot Target LUN.
28. Enter the WWPN for CT1.FC0.
29. Select Primary for the SAN boot target type.
30. Click OK to add the SAN boot target.
31. From the vHBA drop-down menu, select Add SAN Boot Target.

32. Keep 0 as the value for Boot Target LUN.

33. Enter the WWPN for CT1.FC1.

34. Click OK to add the SAN boot target. Click OK, then click OK again to create the boot policy.

The screenshot shows the configuration page for a Boot Policy named 'Pure_HV_Boot'. The interface is divided into 'Actions' and 'Properties' sections. The 'Properties' section includes fields for Name, Description, Owner, Reboot on Boot Order Change, Enforce vNIC/vHBA/iSCSI Name, and Boot Mode. The 'Boot Mode' is set to 'Legacy'. Below the properties is a 'Warning' section with text explaining boot order precedence. On the left, there are buttons for adding various device types: Local Devices, CIMC Mounted vMedia, vNICs, vHBAs, iSCSI vNICs, and EFI Shell. The main area displays a 'Boot Order' table with columns for Name, Order, vNIC/vHBA/iSCSI..., Type, WWN, LUN Name, and Slot. The table shows a hierarchy starting with CD/DVD (Order 1) and San (Order 2). Under San, there are SAN Primary and SAN Secondary entries. The SAN Primary entry has two sub-entries: SAN Target Pr... (Primary) and SAN Target S... (Secondary). The SAN Secondary entry has one sub-entry: SAN Secondary (Secondary). At the bottom of the table, there are 'Move Up', 'Move Down', and 'Delete' icons.

Create Service Profile Templates

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

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

1. In Cisco UCS Manager, click Servers on the left.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter pure-HV-FC as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the "Updating Template" option.

- Under UUID, select UUID_Pool as the UUID pool.

Create Service Profile Template ? X

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

Name :

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

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

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

UUID Assignment:

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

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

< Prev Next > **Finish** Cancel

- 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 default setting for Dynamic vNIC Connection Policy.
- Select the “Expert” option to configure the LAN connectivity.
- Click “Add” to add a vNIC.
- Label the vNIC ‘eth0’ and ‘eth1’ for the vNIC interfaces.
- Check the box for “Use vNIC Template” and use the vNIC templates created earlier.

Create vNIC



Name :

Use vNIC Template :

Redundancy Pair :

vNIC Template :

Peer Name :

[Create vNIC Template](#)

Adapter Performance Profile

Adapter Policy :

[Create Ethernet Adapter Policy](#)

OK

Cancel

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

Simple Expert No vNICs Use Connectivity Policy

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

Name	MAC Address	Fabric ID	Native VLAN
vNIC eth1	Derived	derived	
vNIC eth0	Derived	derived	

Delete Add Modify

< Prev Next > Finish Cancel

6. Click Next.

Configure Storage Options

1. Select the Expert option for the “How would you like to configure SAN connectivity?” field.
2. Click Add to add the vHBA Templates.

The screenshot shows a 'Create vHBA' dialog box with the following configuration:

- Name: fc0
- Use vHBA Template:
- Redundancy Pair:
- vHBA Template: pure-vHBA-A
- Adapter Policy: Windows
- Buttons: OK, Cancel

3. Click OK.

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 Pure-HV_Boot for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Pure_HV_Boot** [Create Boot Policy](#)

Name : **Pure_HV_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/v...	Type	WWN	LUN N...	Slot N...	Boot N...	Boot P...	Descri...
CD/DVD	1								
San	2								
SAN Primary		fc0	Primary						
SAN Target Primary			Primary	52:4A:...	1				
SAN Target Secondary			Secon...	52:4A:...	1				
SAN Secondary		fc1	Secon...						

[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

1. Keep Defaults.
2. Click Next.

Configure Operational Policies

To configure the operational policies, complete the following steps:

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

Create Service Profile Template ? X

Optionally specify information that affects how the system operates.

⊖ BIOS Configuration

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

BIOS Policy: MS-Host ▼

+ 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: No-Power-Cap ▼ [Create Power Control Policy](#)

+ Scrub Policy

+ KVM Management Policy

< Prev Next > **Finish** Cancel

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

Create Service Profiles

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

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template Pure_HV_FC.
3. Right-click Hyper-V-Host-FC and select Create Service Profiles from Template.
4. Enter pure-hv-0 as the service profile prefix.
5. Enter 1 as “Name Suffix Starting Number.”
6. Enter 6 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.

Add More Servers to FlashStack Unit

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

Storage Configuration – Boot LUNs

Pure Boot Storage Setup



Disable 3 of 4 FC interfaces on the Cisco Fabric Interconnects prior to installing Windows 2016 since Windows 2016 does not support multipathing. After multipathing is installed and enabled, enable all interfaces on the Fabric Interconnects and the Pure Storage array.

To disable FC interface on each node, complete the following steps:

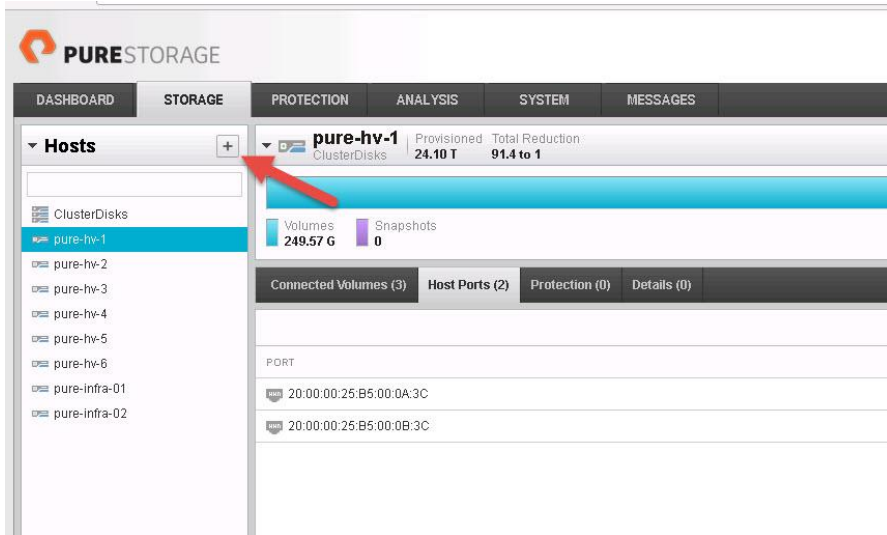
1. Navigate to FC Port 1.
2. Click the General tab.
3. Right-click in the Physical Display pane.
4. Click Disable.

The screenshot displays the configuration page for a specific FC port. The breadcrumb trail at the top reads: Equipment / Fabric Interconnects / Fabric Interconnect A(primary)... / Fixed Module / FC Ports / FC Port 1. The 'General' tab is selected. On the left, the 'Fault Summary' shows zero faults across four categories. The 'Status' section indicates the overall status is 'Up' and the admin state is 'Enabled'. The 'Actions' menu includes options like 'Enable Port', 'Disable Port', and 'Configure as Uplink Port'. The 'Physical Display' pane shows a visual representation of the port with a context menu open, highlighting the 'Disable' option. The 'Properties' section provides technical details: ID: 1, Slot ID: 1, User Label, WWPN: 20:01:fb:c2:86:d9:a5:80, Port Type: Physical, Mode: F, Negotiated Speed: 8 Gbps, and VSAN: Fabric Dual/vsan default. The transceiver is identified as an SFP with model FTLF8528P2BCV-CS and vendor CISCO-FINISAR.

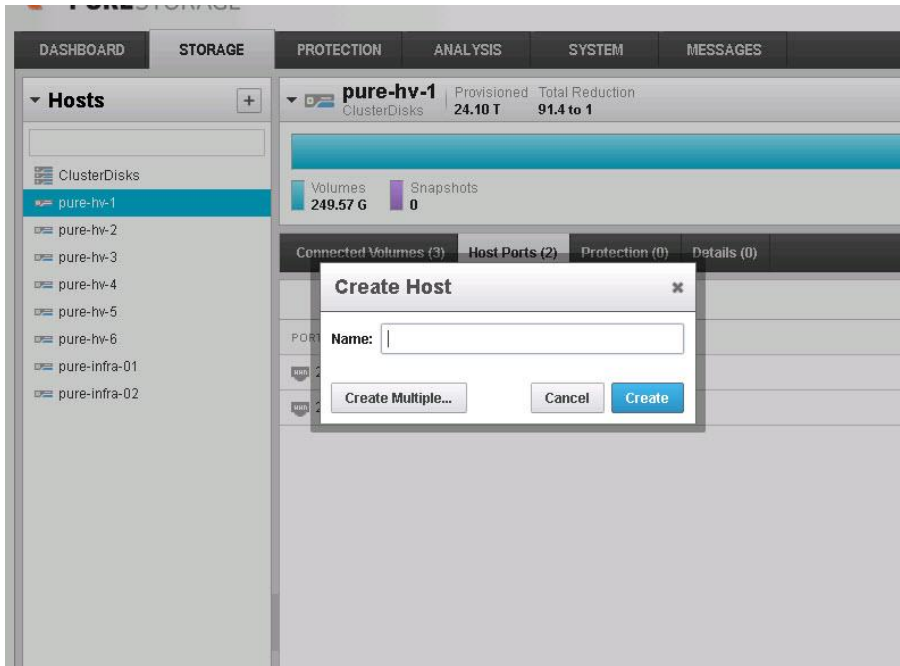
Create Hosts

To create Hosts in the Pure Storage, complete the following steps:

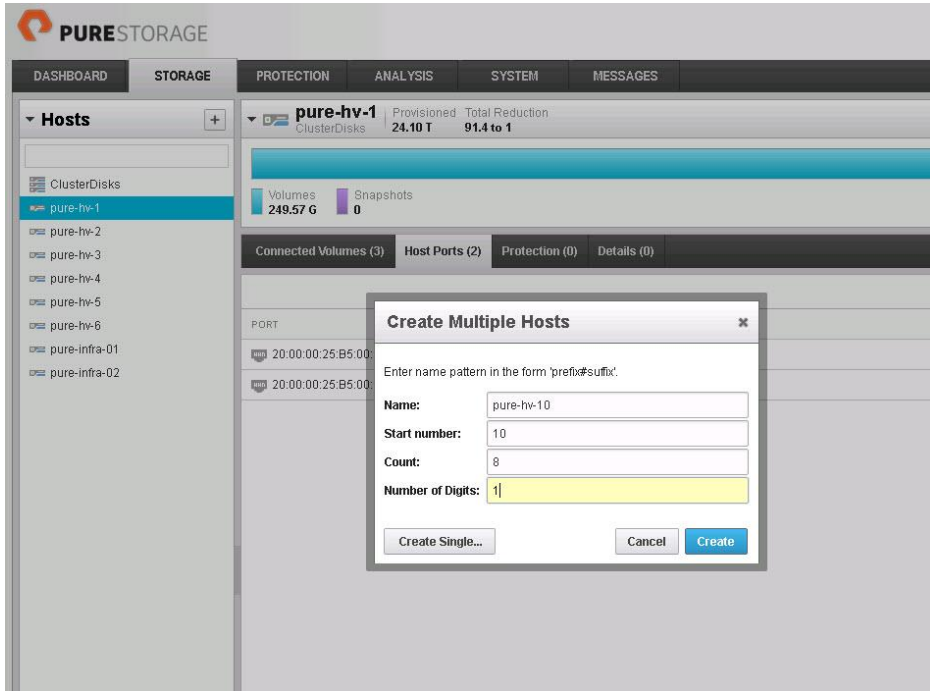
1. Select STORAGE on the Pure Menu.
2. In the Left Navigation Pane, **Select the '+' button next to Hosts.**



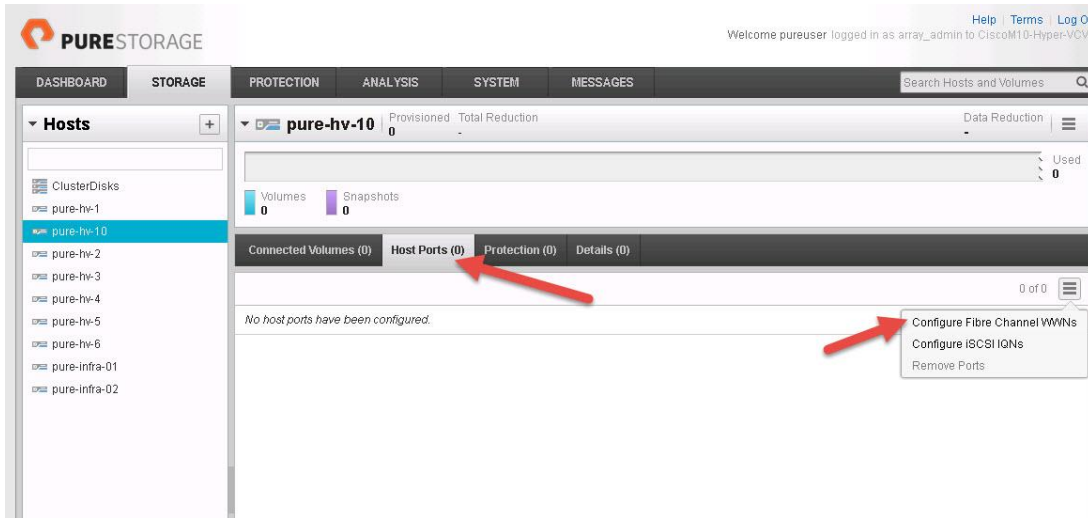
3. Click Create Host.




4. You can create a single host or a whole group of Hosts by selecting 'Create Multiple..'



5. When the host is created, click the 'Host Ports' tab, then click the menu button to configure Fibre Channel WWPNS.



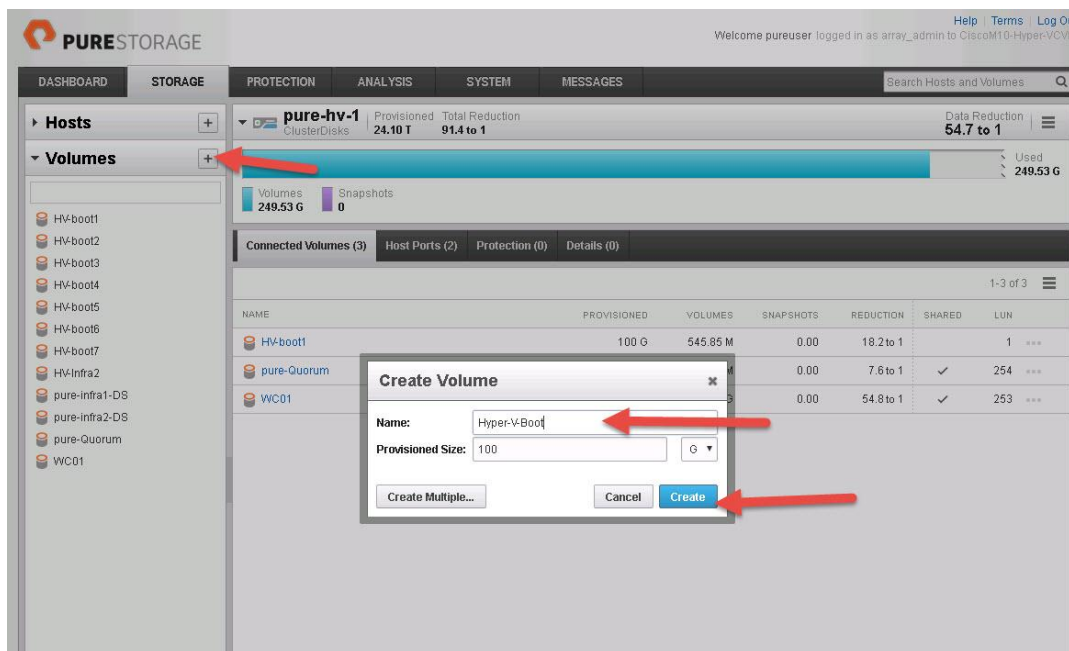
6. The WWPNS can either be entered manually or the Pure Array will discover vHBAs attached to it.



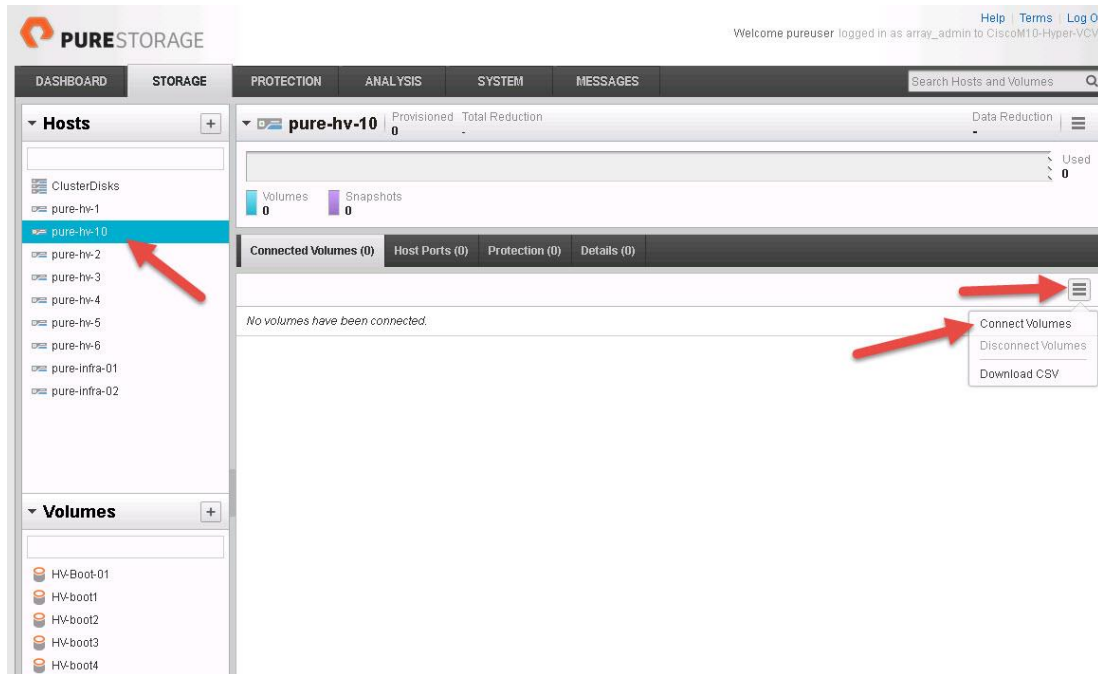
Create Volumes and Attach to Hosts on the Boot

To create Volumes in Pure Storage, complete the following steps:

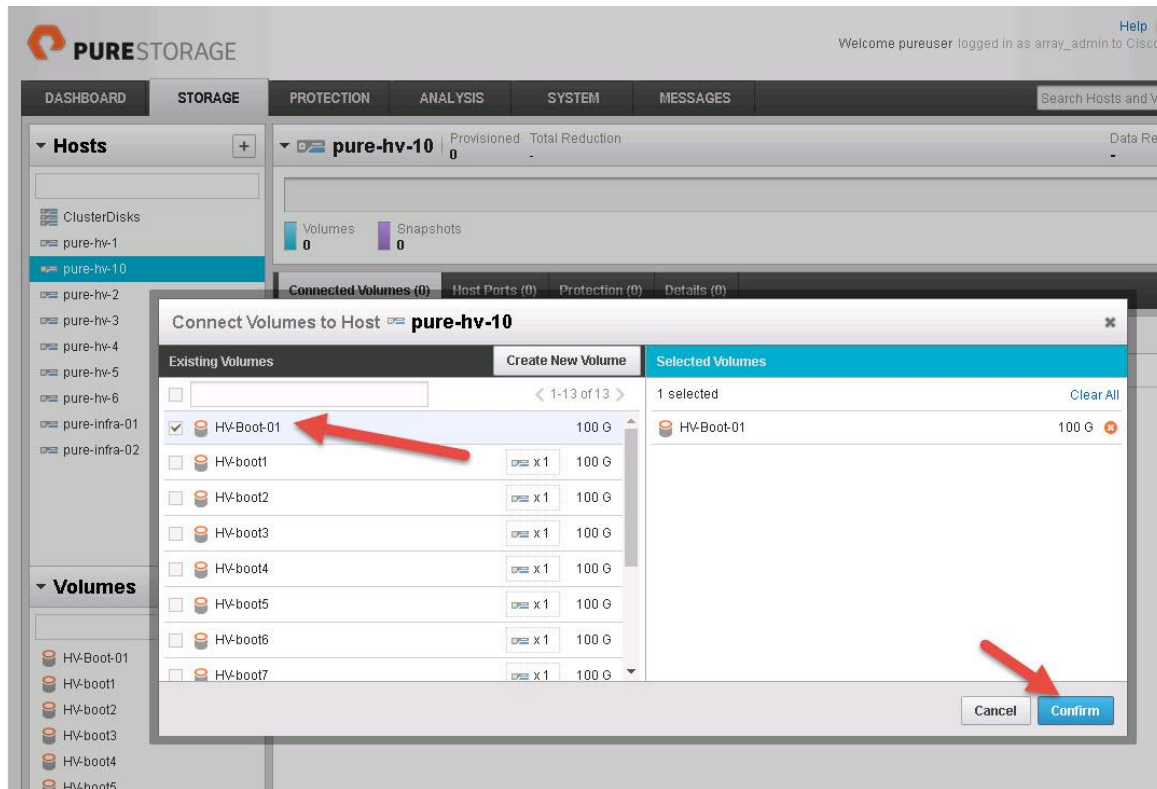
1. Select 'STORAGE' on the Pure Menu.
2. In the Left Navigation Pane, Select the '+' button next to 'Volumes'.



3. Enter a volume name and desired size. In this design, we used 100GB for Windows Boot Volume size. In this step there is also an option to create multiple Volumes at once.
4. When the Boot Volumes are created, return to the 'Hosts' section in the left navigation pane and use the menu button to connect the volumes to the Hosts created earlier.



5. Select the proper volume to their respective hosts (In this study, host 'pure-hv-1' is mapped to volume 'HV-boot-1' and so on).



Microsoft Windows Server 2016 Hyper-V Deployment Procedure

Setting Up Microsoft Windows Server 2016

This section provides detailed instructions for installing Microsoft Windows Server 2016 in an environment. After the procedures are completed, 8 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 Cisco UCS environment to run the IP KVM.

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

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

17. To boot the server, select the KVM tab.
18. Select Power On Server in the KVM interface Summary tab, and then click OK.

Install Windows Server 2016

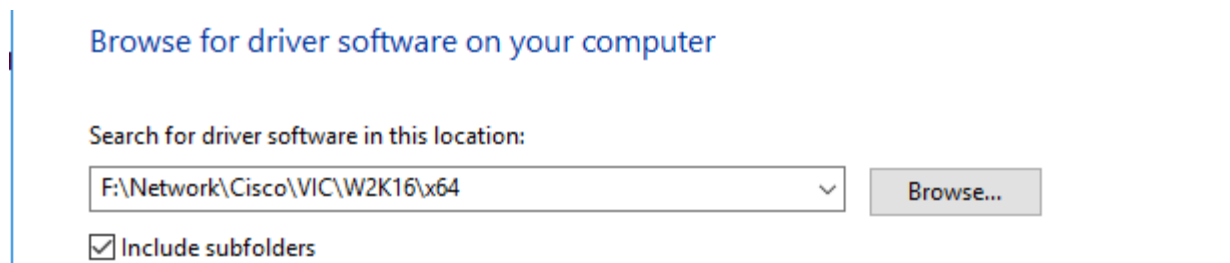
To install Windows Server 2016 to each host, complete the following steps:

1. On boot, the machine detects the presence of the Windows installation media.
2. After the installer has finished loading, Enter 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 (Server with a GUI) and click Next.

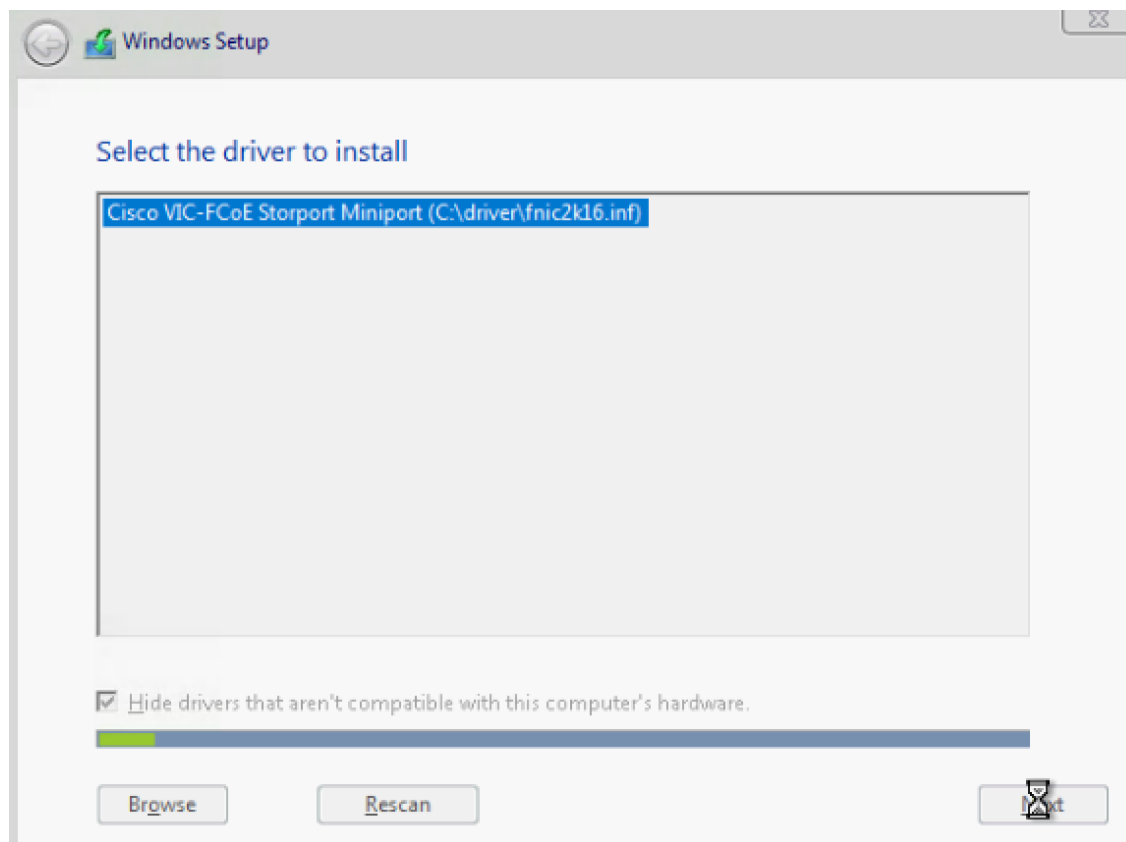


You may optionally remove the GUI after the Hyper-V cluster is operational.


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





13. Back in the KVM Console, click Load Driver and then click OK.
14. The Cisco VIC FCoE Storport Miniport driver is auto detected, click Next.

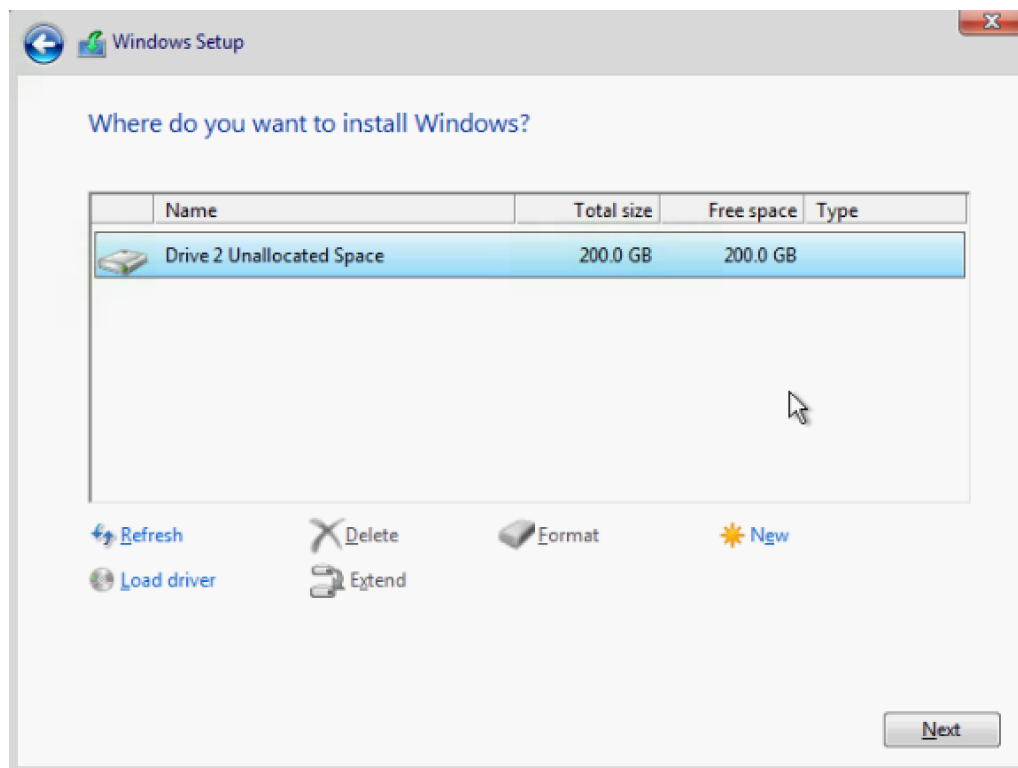


15. You should see a LUN listed in the drive selection screen.

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

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

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



16. Select the LUN and click Next to continue with the install.

17. When Windows is finished installing, enter an administrator password on the settings page and click Finish.

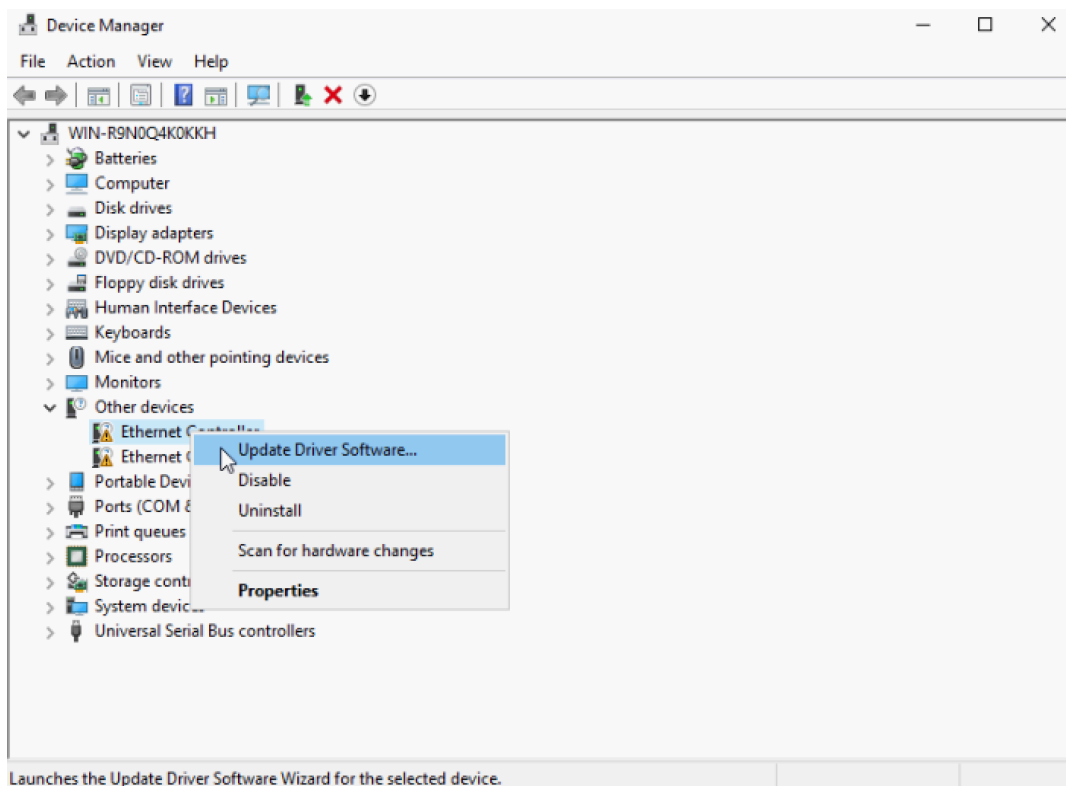
Install Chipset and Windows eNIC Drivers

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

All Hosts

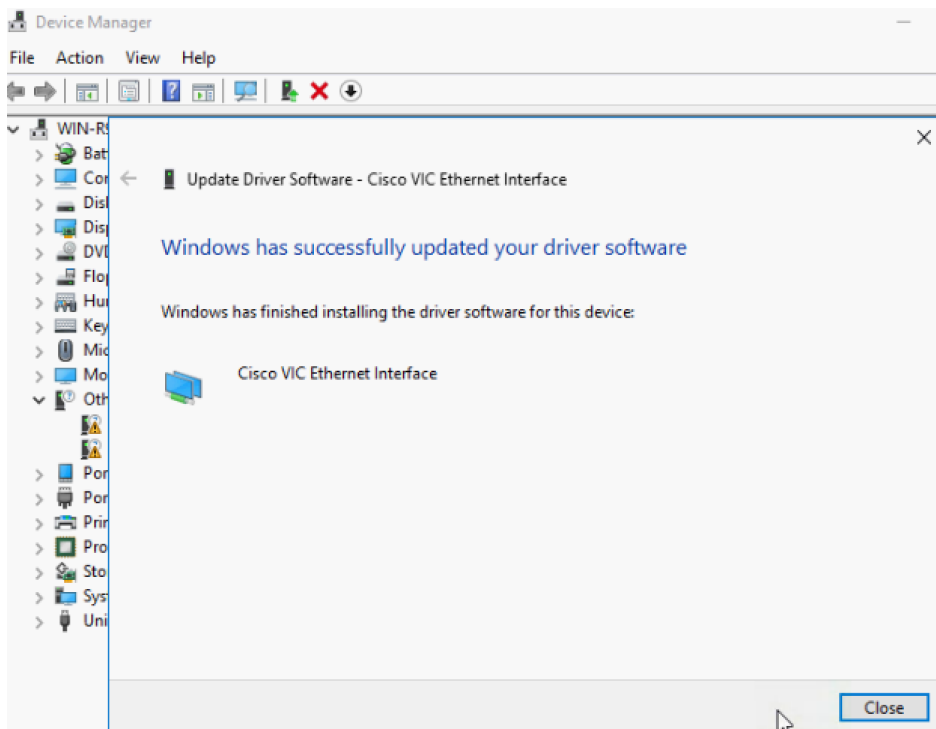
1. In the Virtual Media Session manager, uncheck the Mapped checkbox for the Windows ISO.
2. Click Add Image.
3. Browse to the Cisco UCS driver ISO, click Open.
4. Check the Mapped checkbox for the Cisco UCS driver ISO.
5. Browse to the CD ROM > Chipset > Intel > <Server Model> W2K16 > x64
6. Double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

9. Right-click one of the Ethernet Controller, and select Update Driver Software.

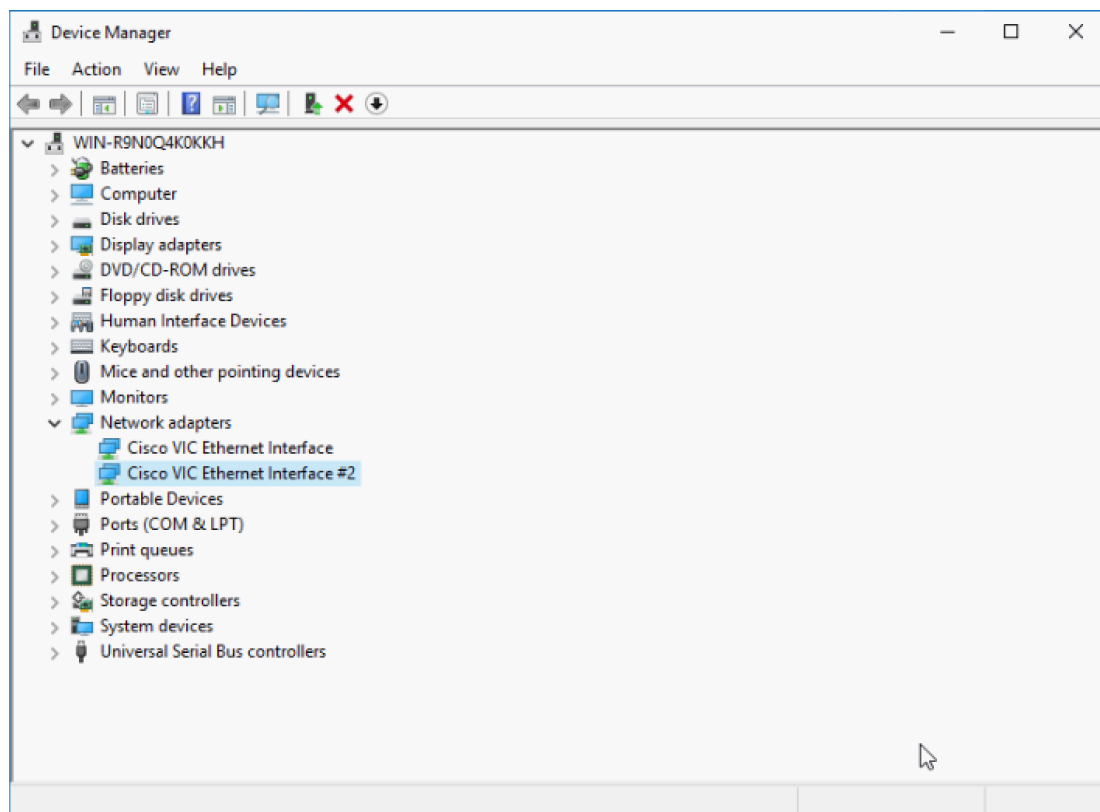


10. Click Browse, and select CDROM drive, click OK.

11. Click Next > Close.



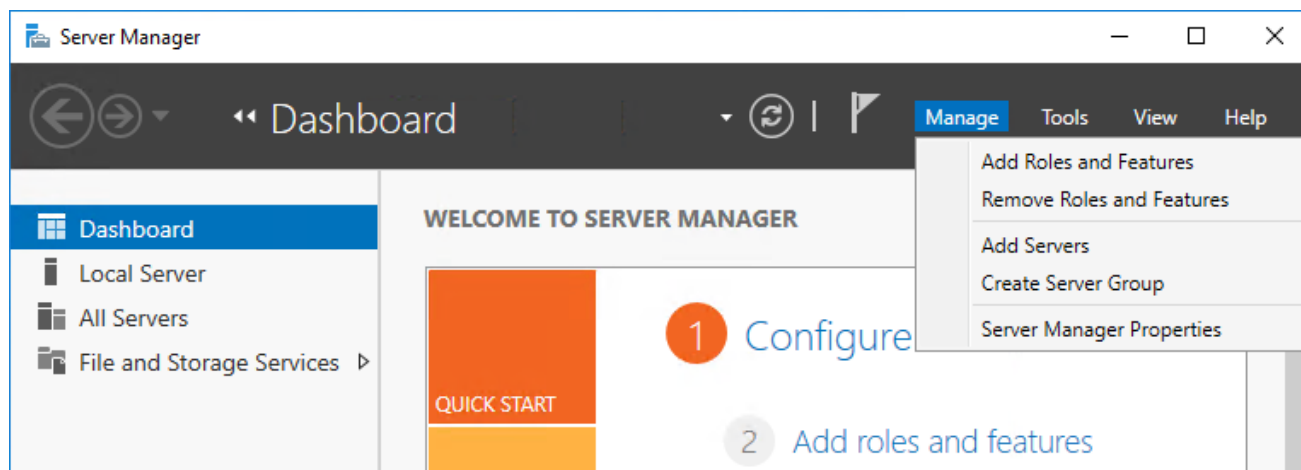
12. Right-click the next Ethernet Controller and select Update Driver Software.
13. Click Search automatically for update driver software.
14. Click Close.
15. Repeat these steps for the remaining Ethernet Controllers.
16. All Cisco VIC Ethernet devices will appear under Network Adapters.



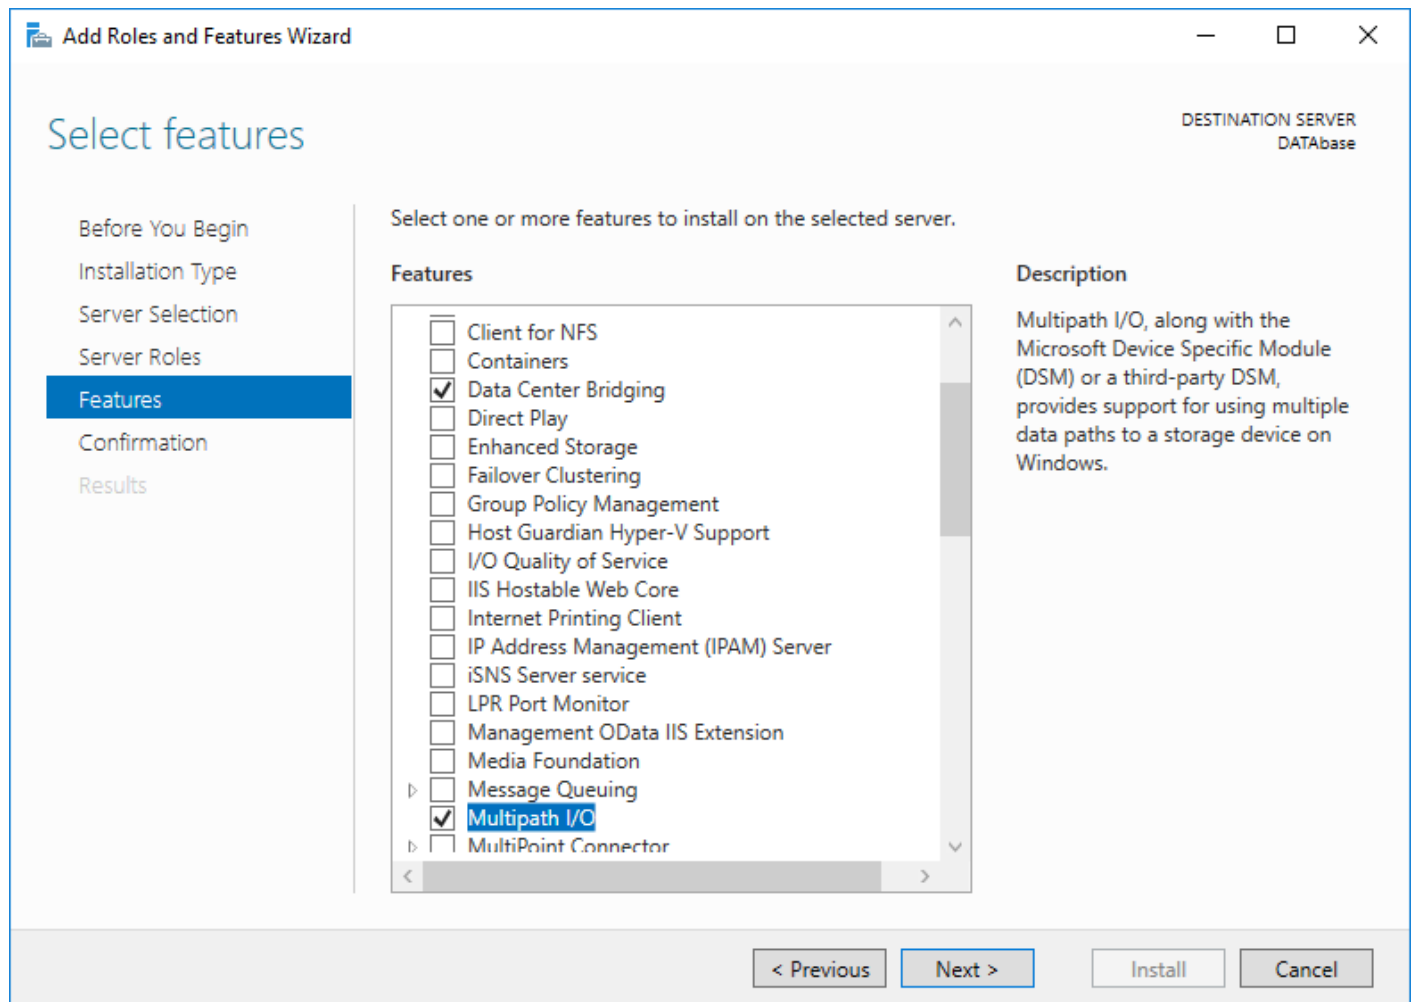
Install Windows Roles and Features

To enable the MPIO feature from the Server Manager, complete the following steps:

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



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



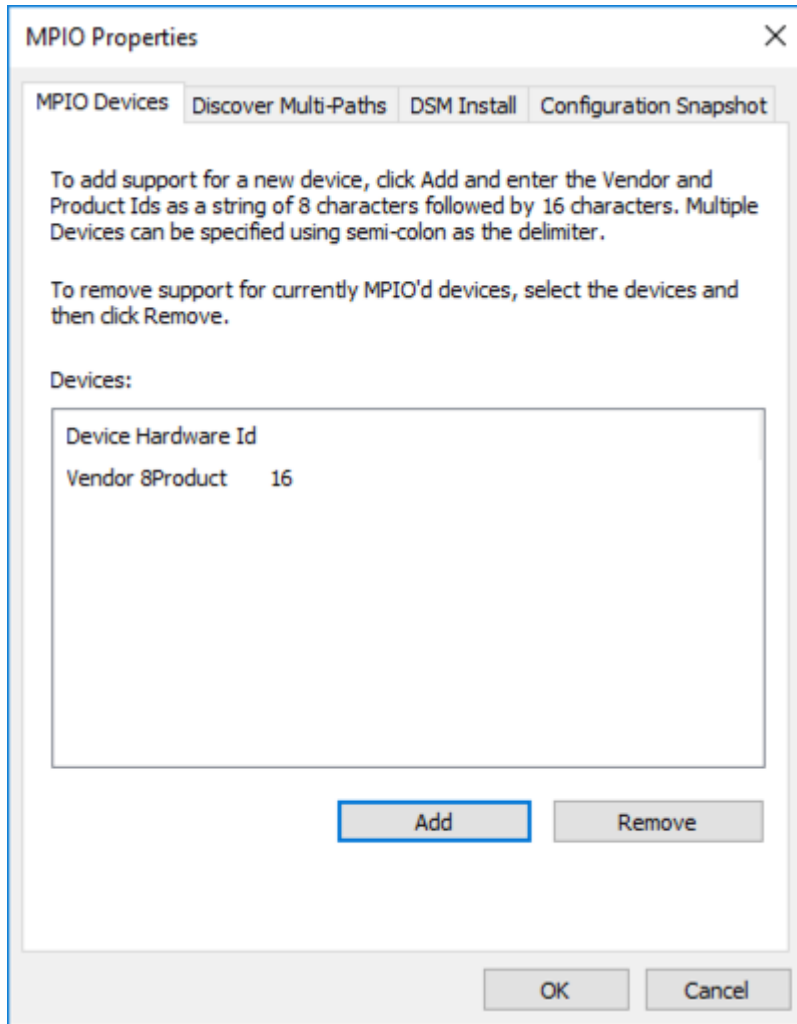
Configuring MPIO for Pure Flash Array//M10

To configure MPIO for Pure Flash Array//M10, complete the following steps:

1. Using the Control Panel double-click the MPIO applet to open up the MPIO Properties dialog box.



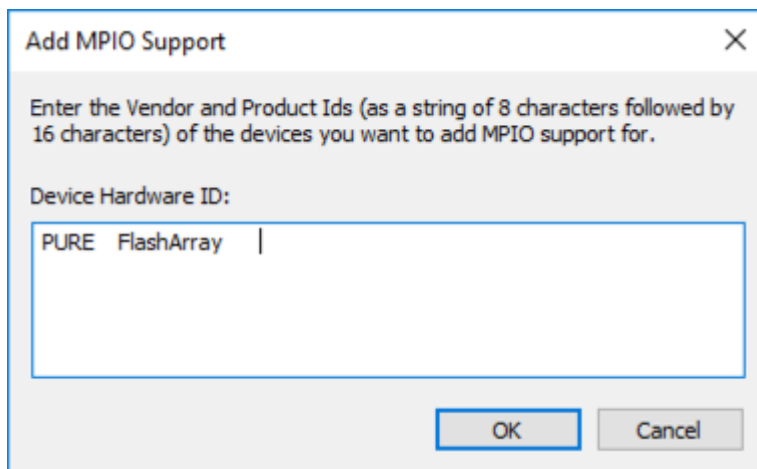
The MPIO applet can also be started from **Start > Run mpioctl**.



2. Click Add to show the Add MPIO Support dialog box.



Pay close attention to the instructions in the dialog box for string formatting. There should be 4 extra spaces after PURE and 6 extra spaces after FlashArray.



Multipath-IO Timers

There are four MPIO Timer values that are the recommended for use with a Pure Storage FlashArray for optimal performance. See [MPIO Timers](#) for full details.

Setting	Default	FlashArray	Definition
CustomPathRecovery	0	1	<p>Specifies whether MPIO performs custom path recovery.</p> <p>Type is boolean and must be filled with either 0 (disable) or 1 (enable). By default, it is disabled.</p>
NewPathRecoveryInterval	40	20	<p>Specifies a custom path recovery time, in seconds. This is the length of time before the server attempts path recovery. The default value is 40.</p> <p>Note: CustomPathRecovery parameter has to be enabled for this value to be used.</p>
PDORemovePeriod	20	30	<p>Specifies a physical device object (PDO) removal period, in seconds. This period is the length of time the server waits after all paths to a PDO have failed before it removes the PDO. The default value is 20.</p>
NewDiskTimeout	60	60	<p>Specifies the disk timeout value, in seconds. This value is the length of time the server waits before it marks the I/O request as timed out.</p> <p>Note: Microsoft documentation has an error and states the default DiskTimeoutValue is 120. On a newly installed Windows Server using Get-MPIOSetting shows the default value is 60. This is an error in Microsoft's documentation and this value should not be changed.</p>

Updating MPIO Timers using Windows PowerShell

Using Windows PowerShell is the preferred method for setting the [MPIO Timer](#) values when managing Windows Server 2012, 2012 R2 or 2016 using [Set-MPIOSetting](#) cmdlet which is part of the MPIO module.

To update MPIO Timers using Windows PowerShell, complete the following steps:



If you are managing a Windows Server 2008 R2 or 2008 R2 Service Pack 1, please refer to the later information in this section.

1. Start a Windows PowerShell session and run the following:

```
Windows Server 2016
```

```
Retrieving Current MPIO Timer Values
```

This will return the current MPIO Timer values. On a newly installed Windows Server all of the default settings will be set as shown below:

```
01 PS C:\> Get-MPIOSetting
02
03 PathVerificationState      : Disabled
04 PathVerificationPeriod    : 30
05 PDORemovePeriod           : 20
06 RetryCount                 : 3
07 RetryInterval              : 1
08 UseCustomPathRecoveryTime : Disabled
09 CustomPathRecoveryTime    : 40
10 DiskTimeoutValue          : 60
```

Update MPIO Timer Values

The following PowerShell shows running `Set-MPIOSetting` four different times with new parameter values. This was done to show each new timer value for clarity. The same can be accomplished with a single line of PowerShell using each of the parameters, this alternative is shown as well.

```
1 Set-MPIOSetting -NewPathRecoveryInterval 20
2 Set-MPIOSetting -CustomPathRecovery Enabled
3 Set-MPIOSetting -NewPDORemovePeriod 30
```

4 Set-MPIOSetting -NewDiskTimeout 60

5

6 OR

7

8 Set-MPIOSetting -NewPathRecoveryInterval 20 -CustomPathRecovery Enabled -
NewPDORemovePeriod 30 -NewDiskTimeout 60

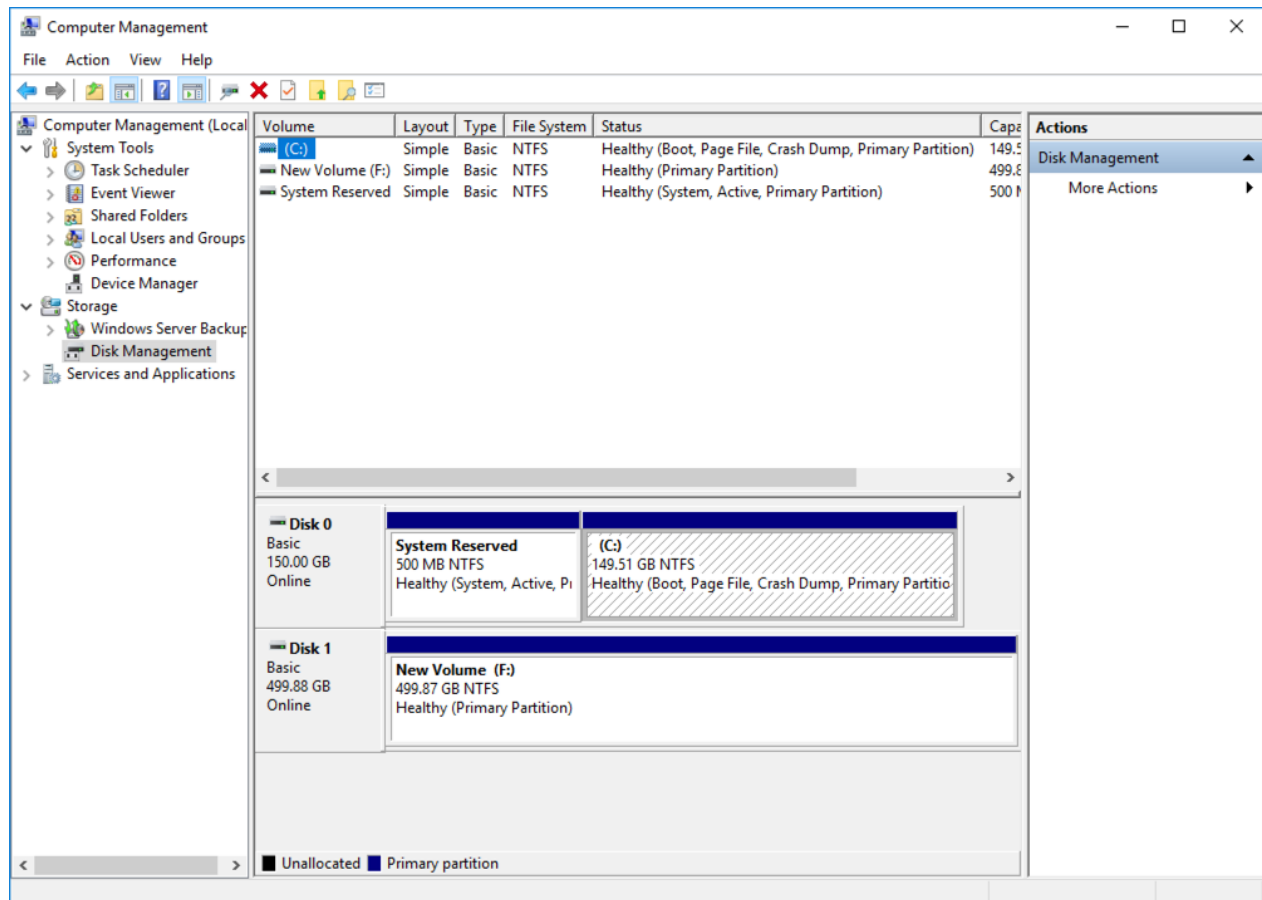


Setup the MPIO Policy (Least Queue Depth (LQD) is the Pure Storage recommended best practice).

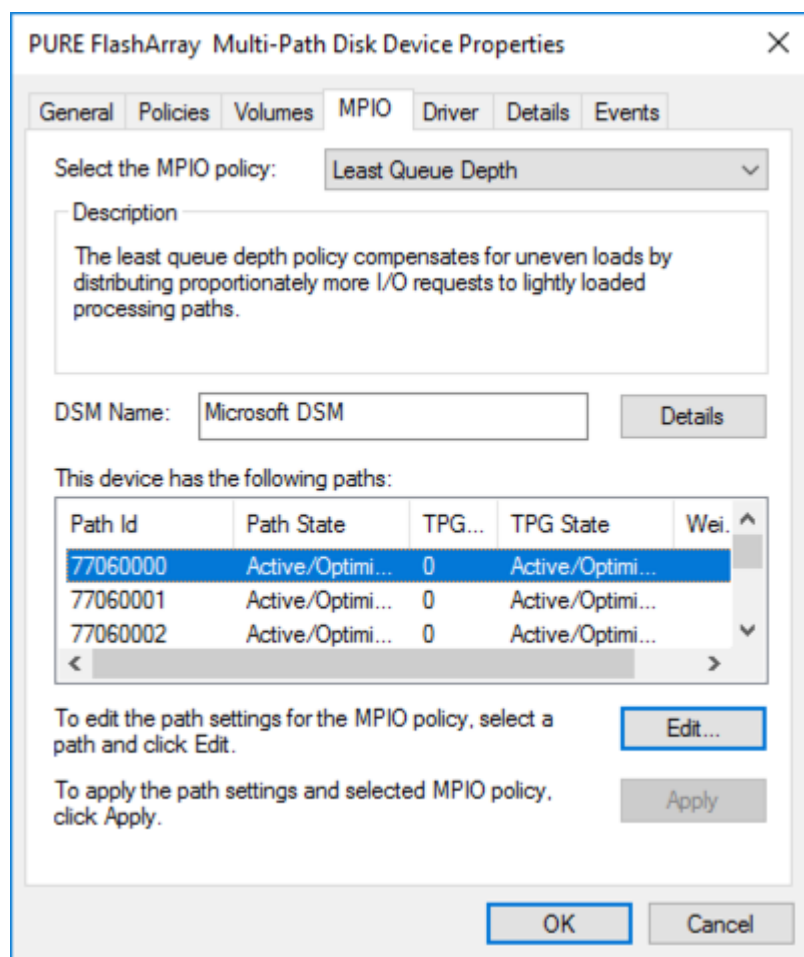
Setup using Windows Disk Management

To setup using Windows Disk Management, complete the following steps:

1. Open Computer Management using the **Tools** menu in **Server Manager**. Select **Disk Management** from the left hand node tree. This will display the currently connected Logical Units (LUNs) to the Windows Server host.



- Right-click **Disk 1** (or whatever Disk # represents the newly connected LUN) and select **Properties** from the menu. This will open the properties dialog for the **PURE FlashArray Multi-Path Disk Device**. Each of the device paths has a **Path State** of **Active/Optimized**.



Host Renaming and Join to Domain

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

- Login to the host and open PowerShell.
- Rename the host:

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

- Assign an IP address to the management interface:

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

- Assign DNS server IP address to the above management interface:

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

5. Add the host to Active Directory:

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

Storage Configuration – Boot LUNs

Pure Boot Storage Setup



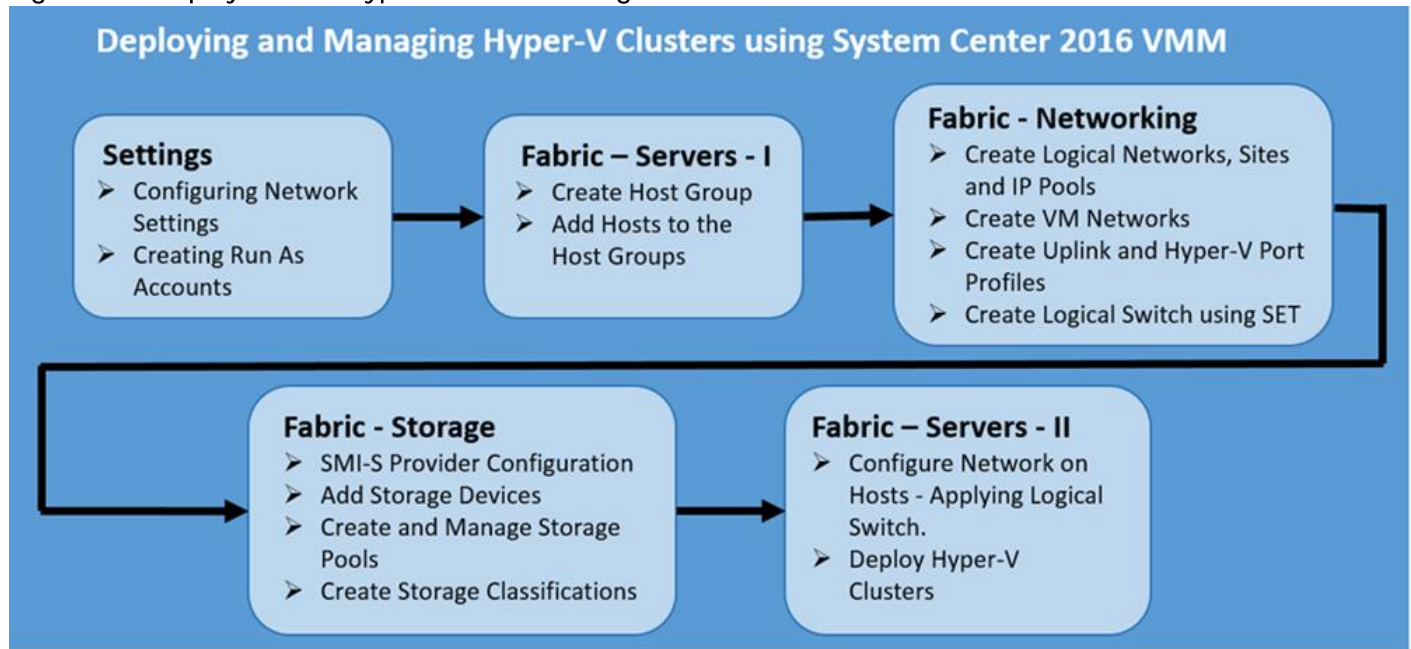
Enable All FC Interfaces on the Fabric Interconnects - now that multipathing has been installed and enabled in Windows 2016, all interfaces can be enabled on the Cisco Fabric Interconnects.

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

In this section, we assume that System Center 2016 VMM is up-and-running in your environment. If not, here is a reference to get that up and running: [Setup SCVMM 2016](#). This section focuses only on configuring the Networking, Storage, and Servers in VMM to deploy and manage Hyper-V failover clusters.

Figure 2 provides a high-level view of the steps that will be covered in detail in the following sections.

Figure 2 Deployment of Hyper-V Cluster Using SCVMM



Settings

Configuring Network Settings

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

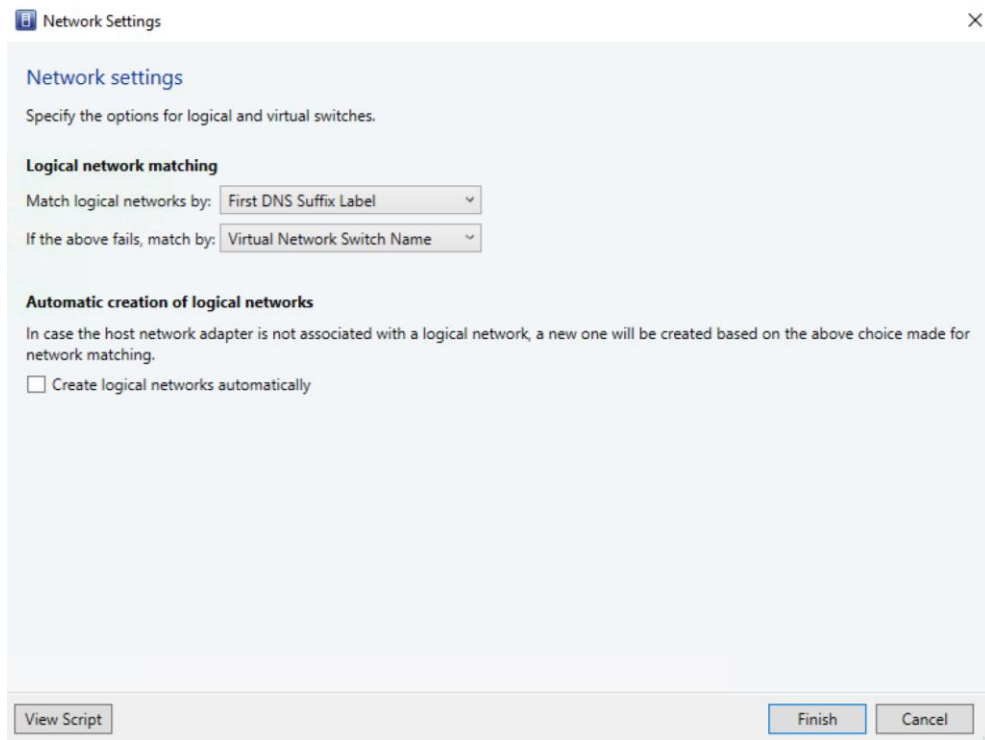
To disable automatic logical network creation, complete the following steps:

1. Open Virtual Machine Manager.
2. Open the Settings workspace.
3. Select the General navigation node.
4. Double-click Network Settings in the details pane.

5. In the Network Settings dialog box, uncheck the Create Logical Networks Automatically option and click OK.



You can change the logical network matching behavior to a scheme that may better suit your naming conventions and design.



Create Run As Account 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 Pure Storage SMI-S provider.

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

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

Fabric – Servers - I

This section covers the following:

- 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. Open the Fabric workspace.
2. In the Fabric pane, expand Servers, and then do one of the following:
 - a. Right-click All Hosts and then click Create Host Group.
 - b. 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.
3. Type a new name, and then press ENTER.
4. Repeat the steps in this procedure to create the rest of the host group structure.



Add Hosts to the Host Group

When the virtual switch is created, you can add the Hyper-V hosts to Virtual Machine Manager; complete the following steps:

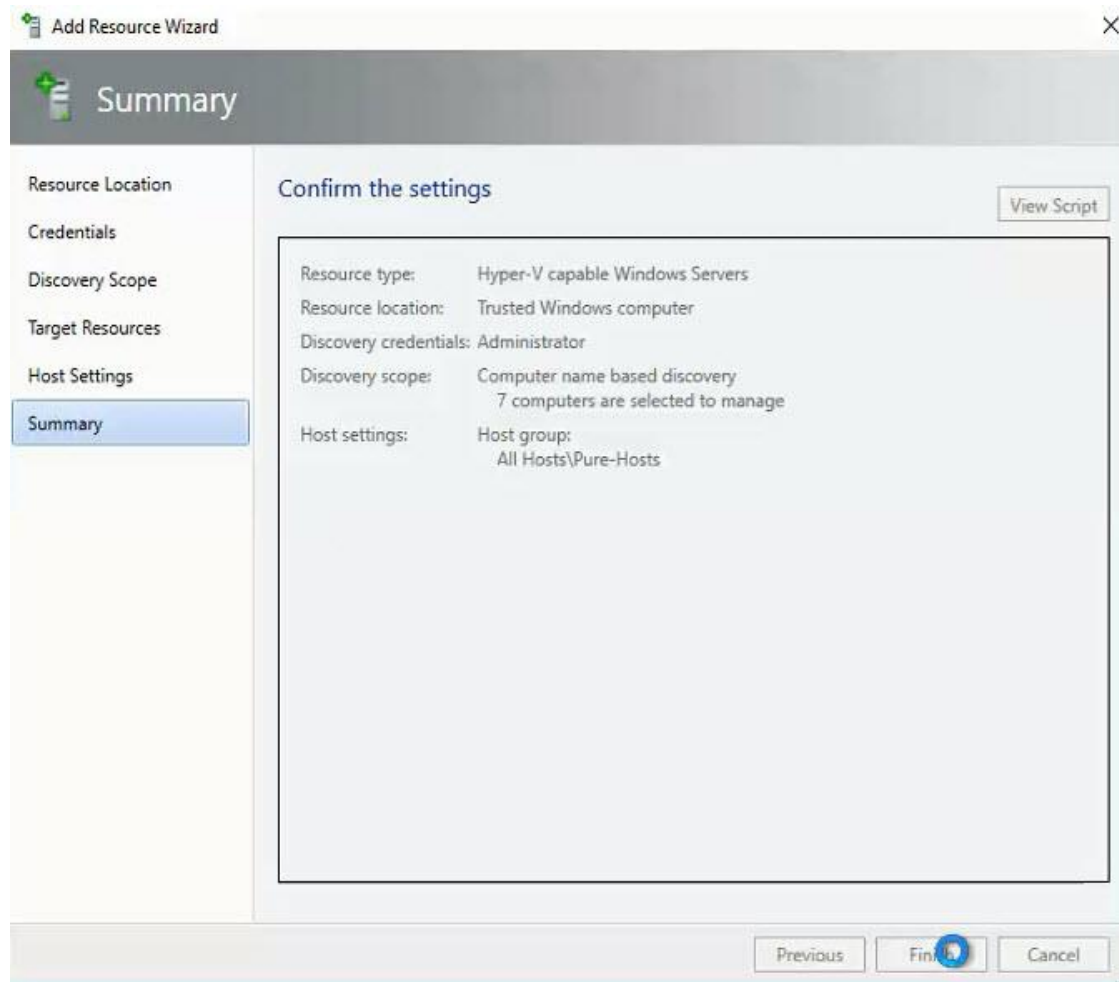
1. Open the Fabric workspace.
2. Select a host group, and On the Home tab, in the Add group, click Add Resources, and then click Hyper-V Hosts and Clusters. The Add Resource Wizard starts.
3. On the Resource location page, click Windows Server computers in a trusted Active Directory domain, and then click Next.

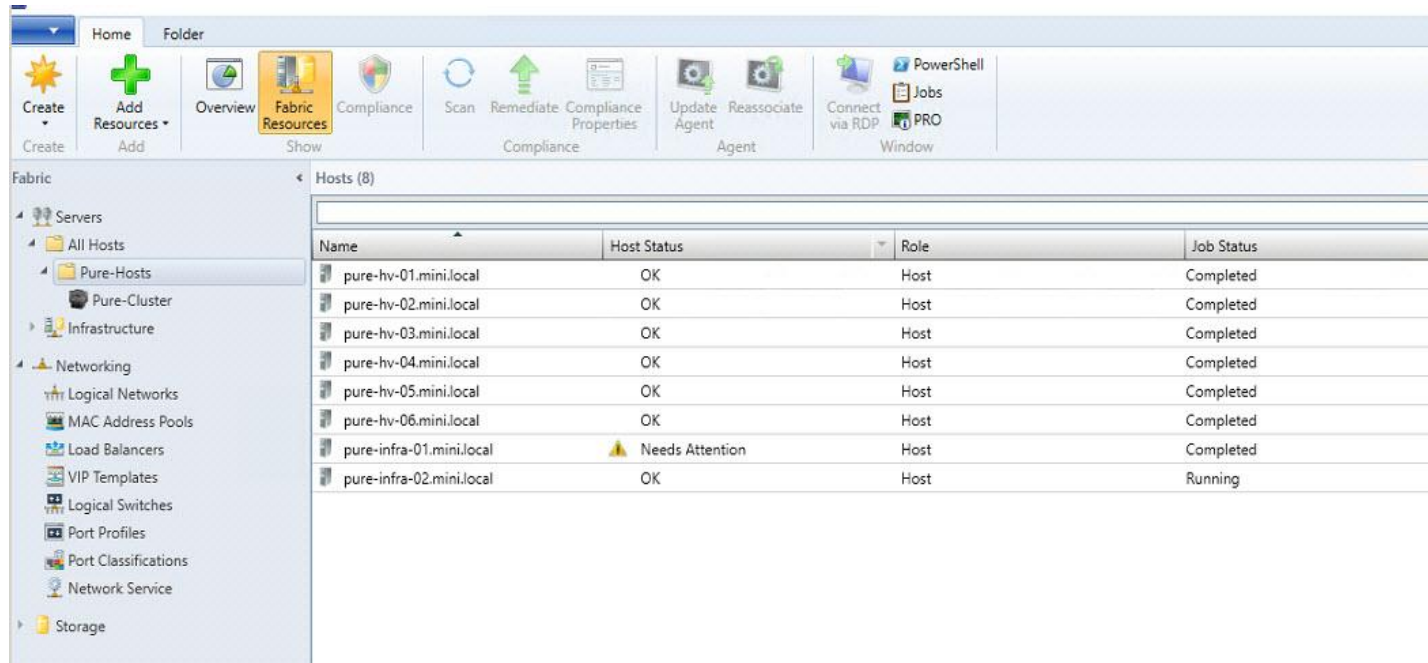
4. On Credentials page, select Use an Run As account, click Browse and add the Run as account created earlier. Click Next.
5. On Discovery scope, select Specify Windows Server computers by names and enter the Computer names. Click Next.
6. Under Target Resources, select the check box next to the computer names that needs to be the part of the Hyper-V cluster.



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

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





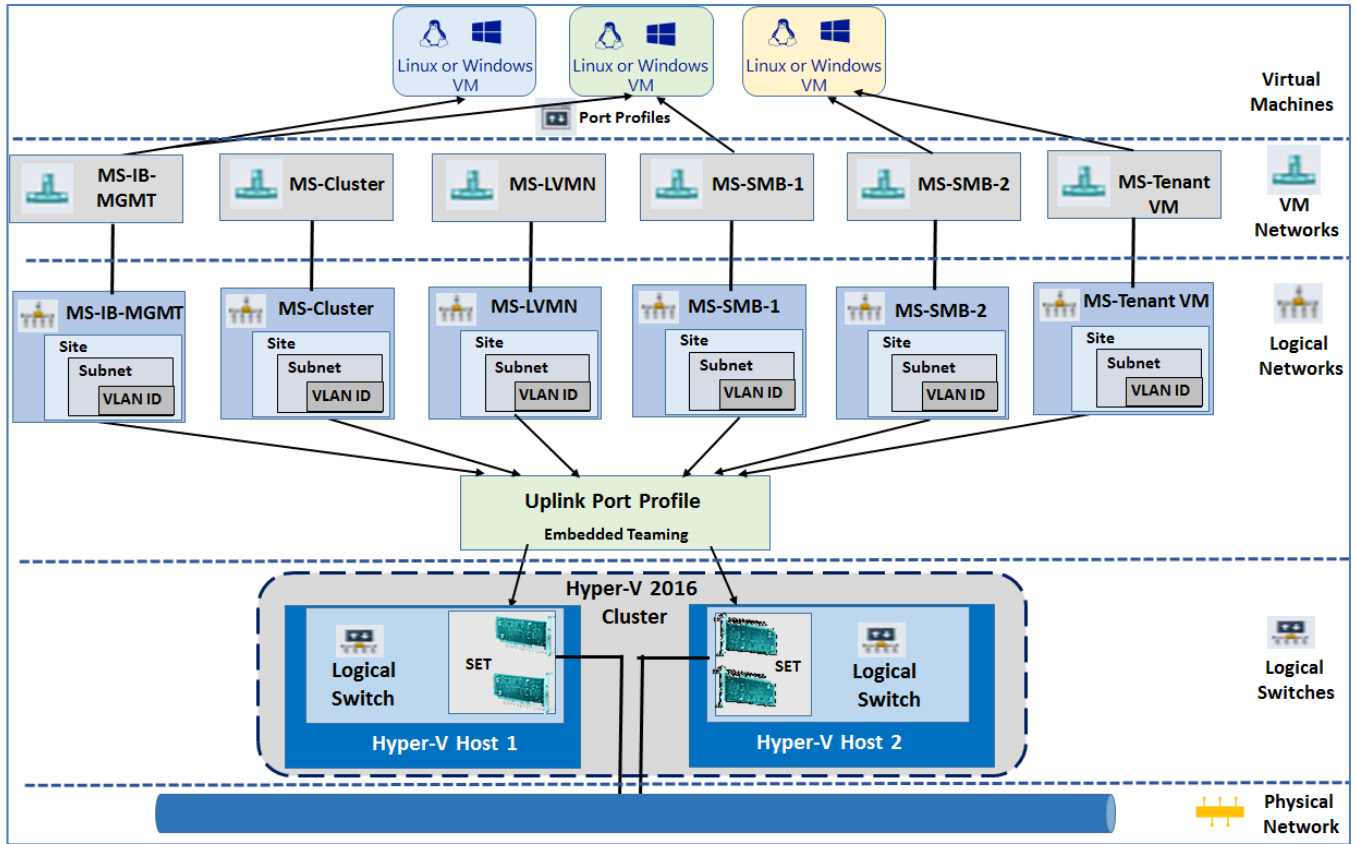
Fabric – Networking

Figure 3 shows the logical representation of the network that will be configured in this section using the System Center 2016 VMM and applied later to configure the network settings of Hyper-V hosts before deploying the failover cluster.



For this solution, you are going to use and deploy “Switch Embedded Teaming (SET)”; a new feature released in Windows server 2016. SET is a new teaming solution integrated with the Hyper-V switch.

Figure 3 SCVMM Logical Network



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

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

Creating Logical Networks, Sites, and IP Pools

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

- HV-MGMT: This logical network will be used for management traffic and has its own IP subnet and VLAN.
- HV-CSV: This network will be used for Microsoft Hyper-V cluster communication and will have its own IP subnet and VLAN.
- HV-LiveMig: This network will be used for Live Migration traffic and will have its own IP subnet and VLAN

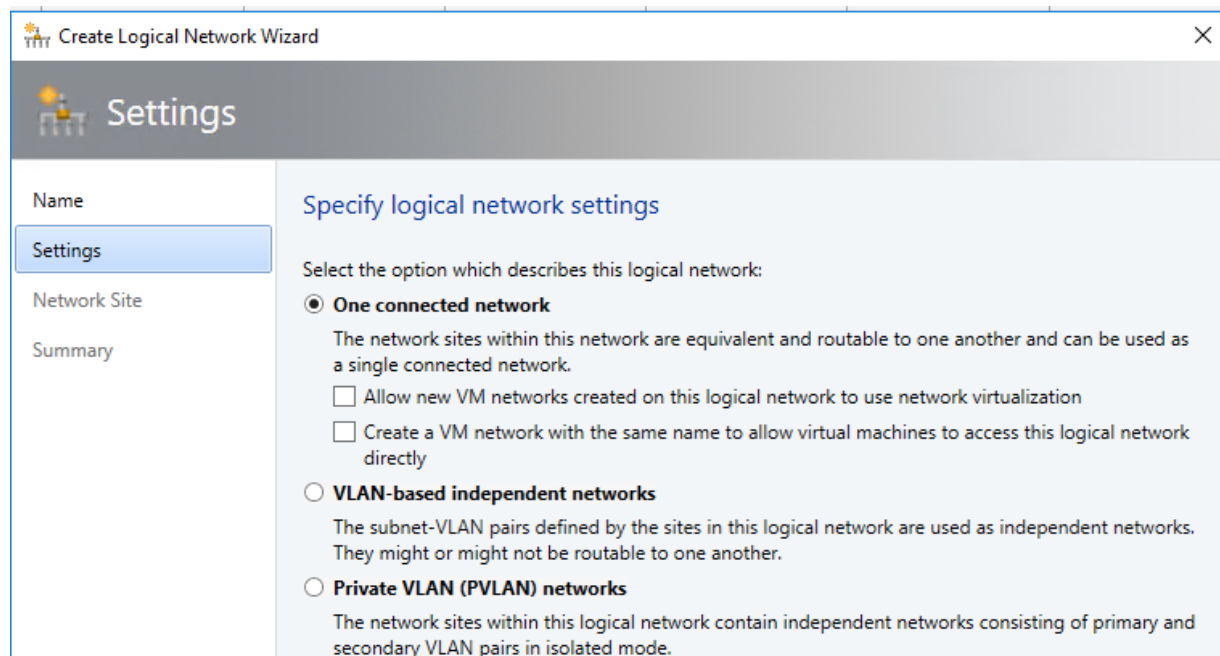
- HV-SMB: This network will be used for SMB file share access/traffic and has its own IP subnet and VLAN.
- HV-VDI: This network will be used for all the VM traffic.

To create logical networks and sites, complete the following steps:

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



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

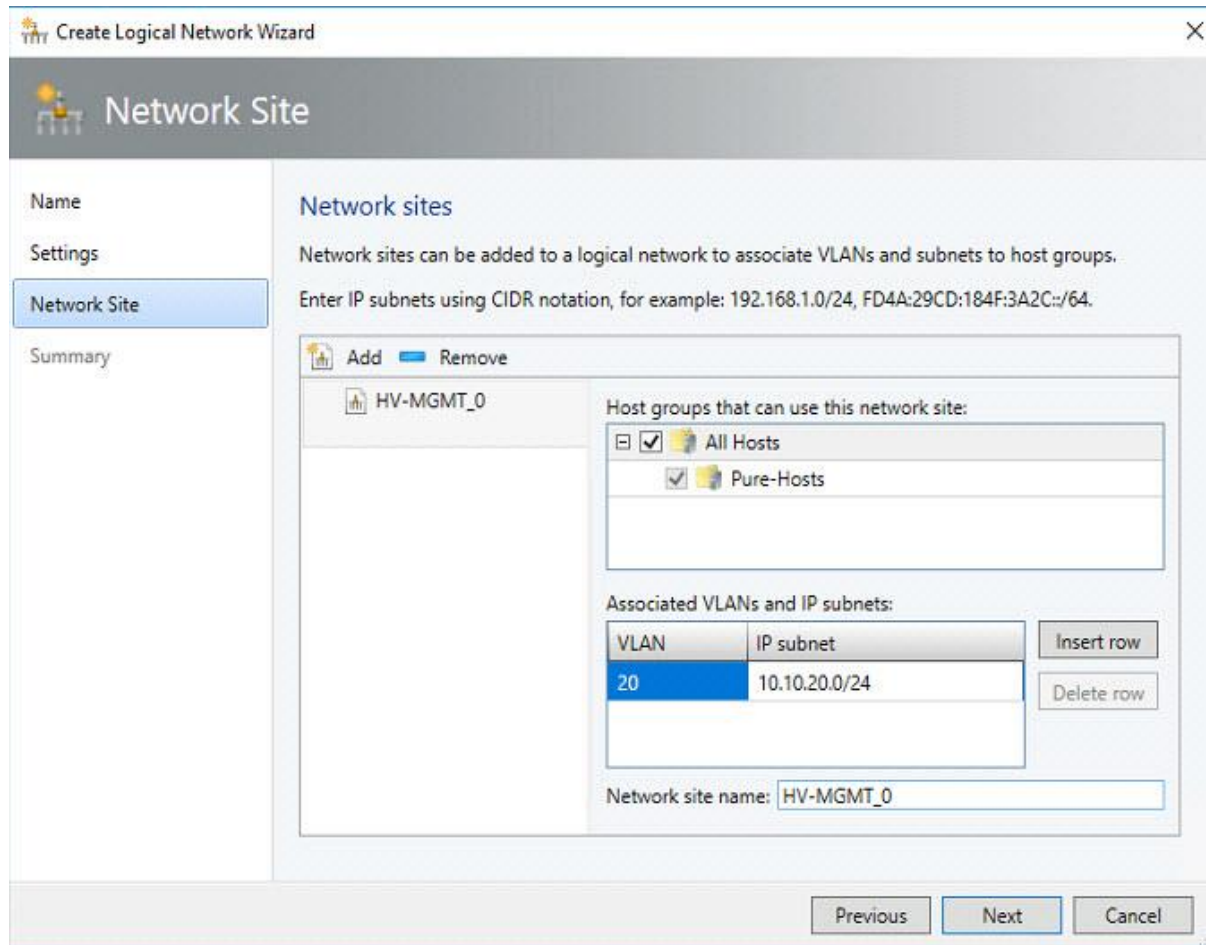


7. Select the sites and Host Group, where you want to apply the Management VLAN. Click the Add button to add a site and then click Insert Row to add VLAN/IP details for the site.

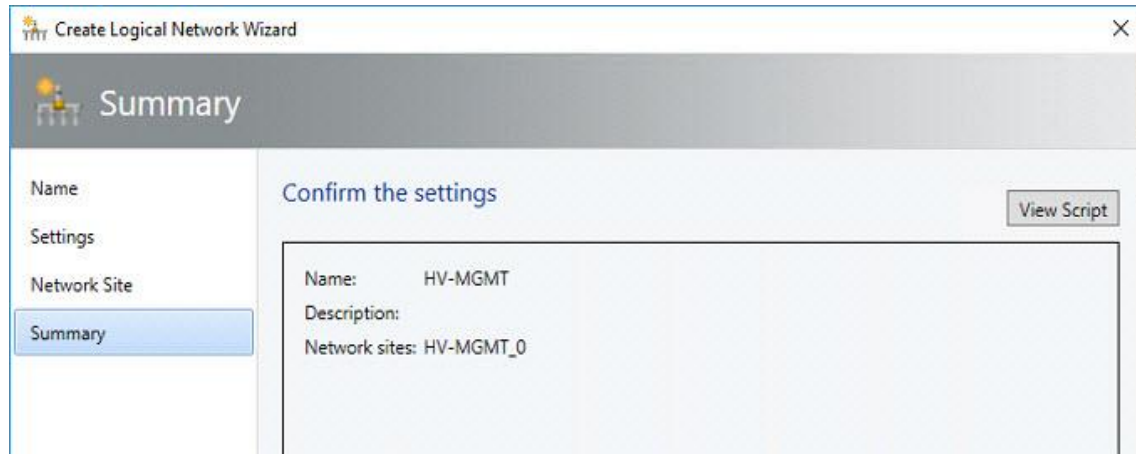


If IP space is configured by corporate DHCP servers, leave the IP subnet blank, which tells SCVMM to just configure the VM for DHCP. If the network does not use VLANs, set the VLAN ID to 0; this

tells SCVMM that VLANs are not to be configured. By default, sites are given the name <Logical Network>_<number>, but you should rename this to something more useful.



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



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

Name	Network Compliance	Subnet	Begin Address	End Address
Cluster	Fully compliant			
Cluster_Pool	Fully compliant	10.10.29.0/24	10.10.29.101	10.10.29.154
Infrastructure	Fully compliant			
LiveMig	Fully compliant			
LivMig_Pool	Fully compliant	10.10.28.0/24	10.10.28.101	10.10.28.154
SMB	Fully compliant			
SMB	Fully compliant	10.10.27.0/24	10.10.27.101	10.10.27.154
VDI	Fully compliant			
VDI2	Fully compliant			

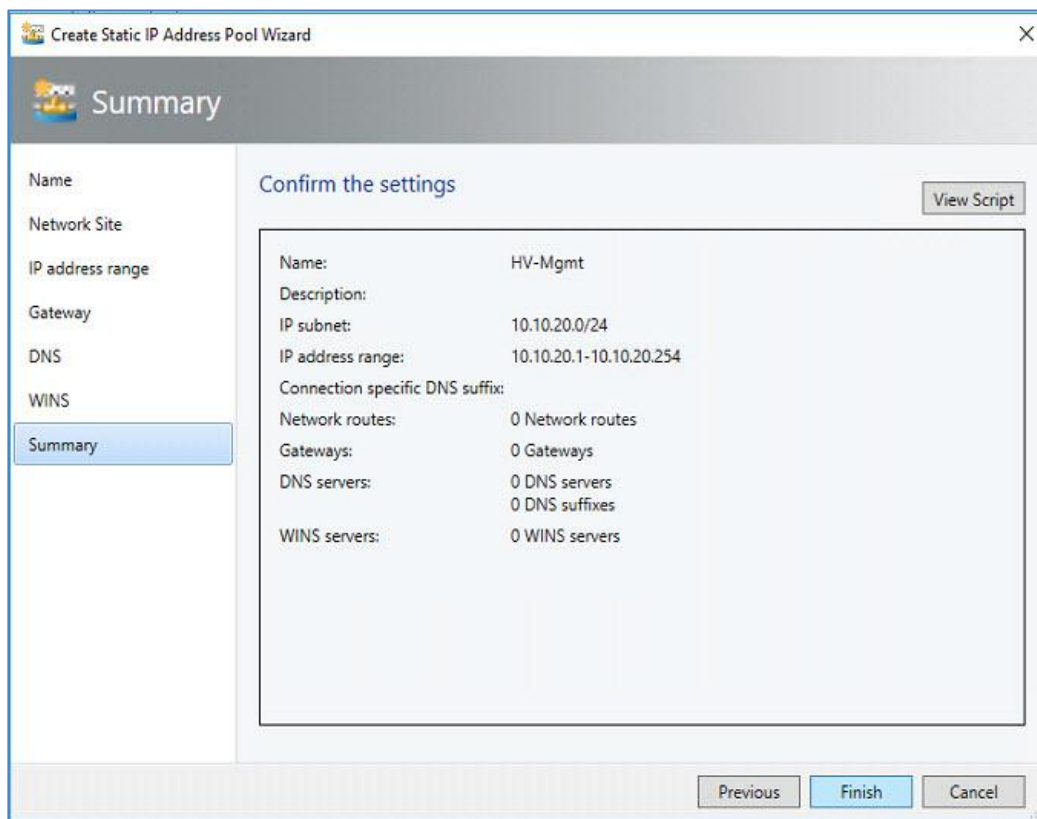
To create a static IP address pool for a logical network in Virtual Machine Manager (VMM), complete the following steps:



With static IP address pools, IP address management for the virtual environment is brought within the scope of the VMM administrator.

1. From the Fabric workspace, Click the Create IP Pool button, or right-click the logical network and select the Create IP Pool context menu action.
2. Enter a name and description. From the drop-down list, select the logical network for the IP pool.
3. The next screen, allows you to use an existing network site or create a new one. Choose to use an existing one and then click Next.
4. The next screen, allows you to use an existing network site or create a new one. Choose to use an existing one and then click Next.

5. The IP Address Range page allows configuration of the IP address range that SCVMM will manage and allocate to resources such as virtual machines and load balancers. Within the range, you can configure specific addresses to be reserved for other purposes or for use by load-balancer virtual IPs (VIPs) that SCVMM can allocate.
6. Click the Insert button, and enter the gateway IP address, then click Next.
7. Configure the DNS servers, DNS suffix, and additional DNS suffixes to append, and then click Next.
8. Enter the WINS server details if used and click Next.
9. On the Summary screen, confirm the configuration, click the View Script button to see the PowerShell that will be used, and then click Finish to create the IP pool.



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

Name	Network Compliance	Subnet	Begin Address	End Address
Cluster	Fully compliant			
Cluster_Pool	Fully compliant	10.10.29.0/24	10.10.29.101	10.10.29.154
Infrastructure	Fully compliant			
LiveMig	Fully compliant			
LivMig_Pool	Fully compliant	10.10.28.0/24	10.10.28.101	10.10.28.154
SMB	Fully compliant			
SMB	Fully compliant	10.10.27.0/24	10.10.27.101	10.10.27.154
VDI	Fully compliant			
VDI2	Fully compliant			

Create VM Networks

To create the VM networks to which virtual machines can be connected, complete the following steps:

1. Open Virtual Machine Manager.
2. Open the VMs and Services workspace.
3. Select the VM Networks navigation node.
4. Click the Create VM Network button.
5. Enter a name and description for the VM network, select the logical network, and click Next.

Create VM Network Wizard

Name

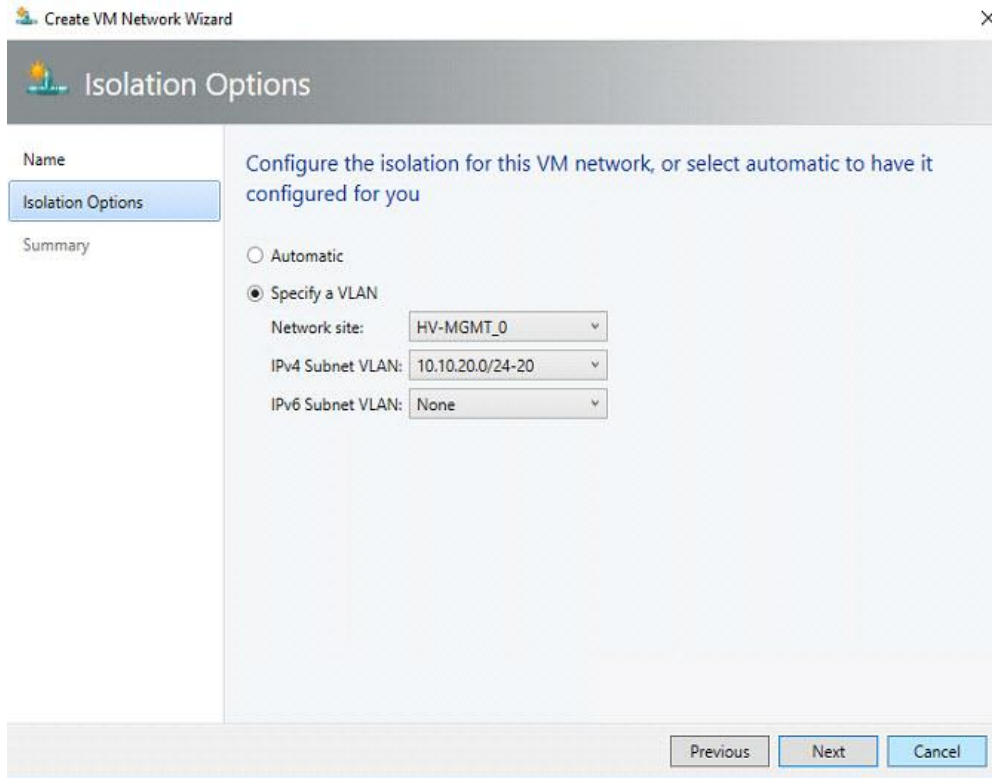
Specify a name and description for the VM network

Name:

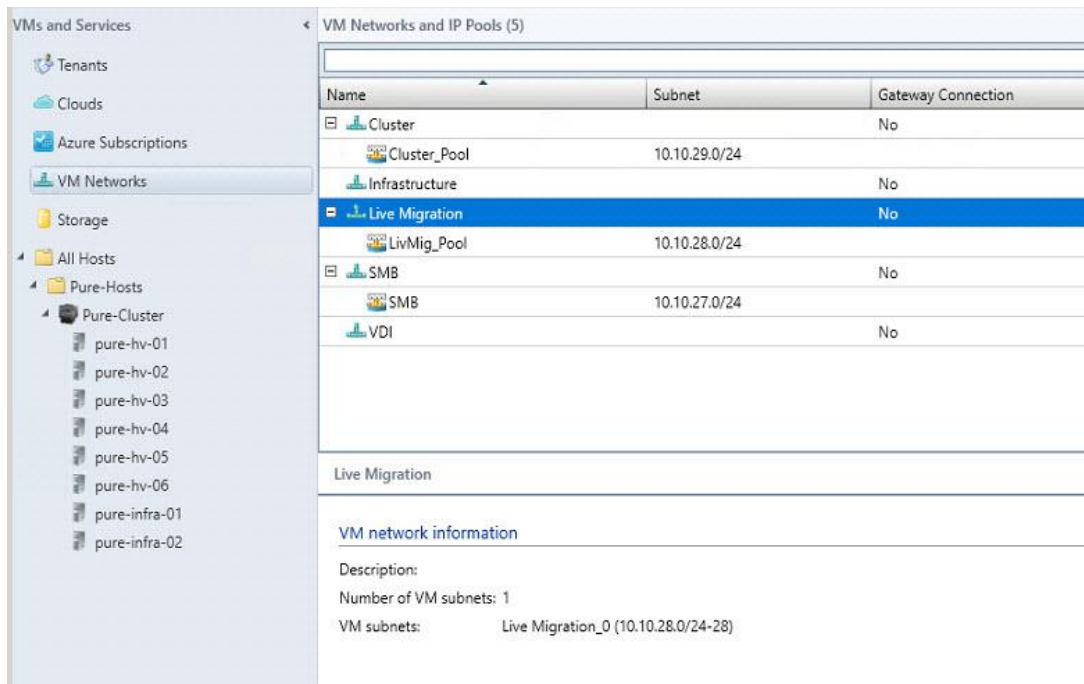
Description:

Logical network:

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



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



Create Uplink Port Profiles and Hyper-V Port Profiles

Create Uplink Port Profile

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

To create an uplink port profile, complete the following steps:

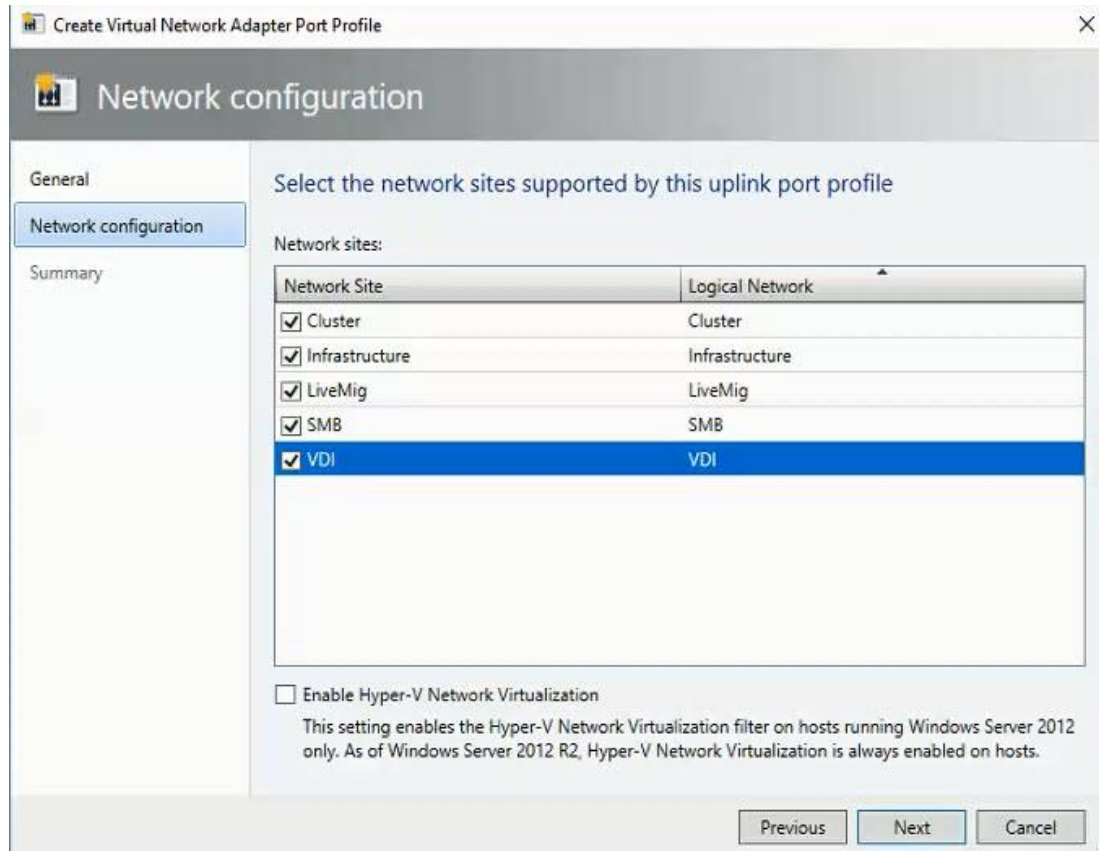
1. Open Virtual Machine Manager.
2. Open the Fabric workspace.
3. Select the Networking > Port Profiles navigation node.
4. Click the Create button drop-down and select Hyper-V Port Profile.
5. Enter a name and description for the new port profile, Select the Uplink Port Profile radio button. Leave the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

The screenshot shows the 'Create Virtual Network Adapter Port Profile' dialog box. The title bar reads 'Create Virtual Network Adapter Port Profile'. The main window has a 'General' tab selected. On the left, there is a navigation pane with 'General' selected, and 'Network configuration' and 'Summary' are visible below it. The main area is titled 'Select the type of Hyper-V port profile' and contains the following fields and options:

- Name:** UpLink-Port-Profile
- Description:** (Empty text box)
- Type of Hyper-V port profile:**
 - Virtual network adapter port profile
 - Uplink port profile
- Load balancing algorithm:** Host Default
- Teaming mode:** Switch Independent

At the bottom of the dialog, there are three buttons: 'Previous', 'Next', and 'Cancel'.

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



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

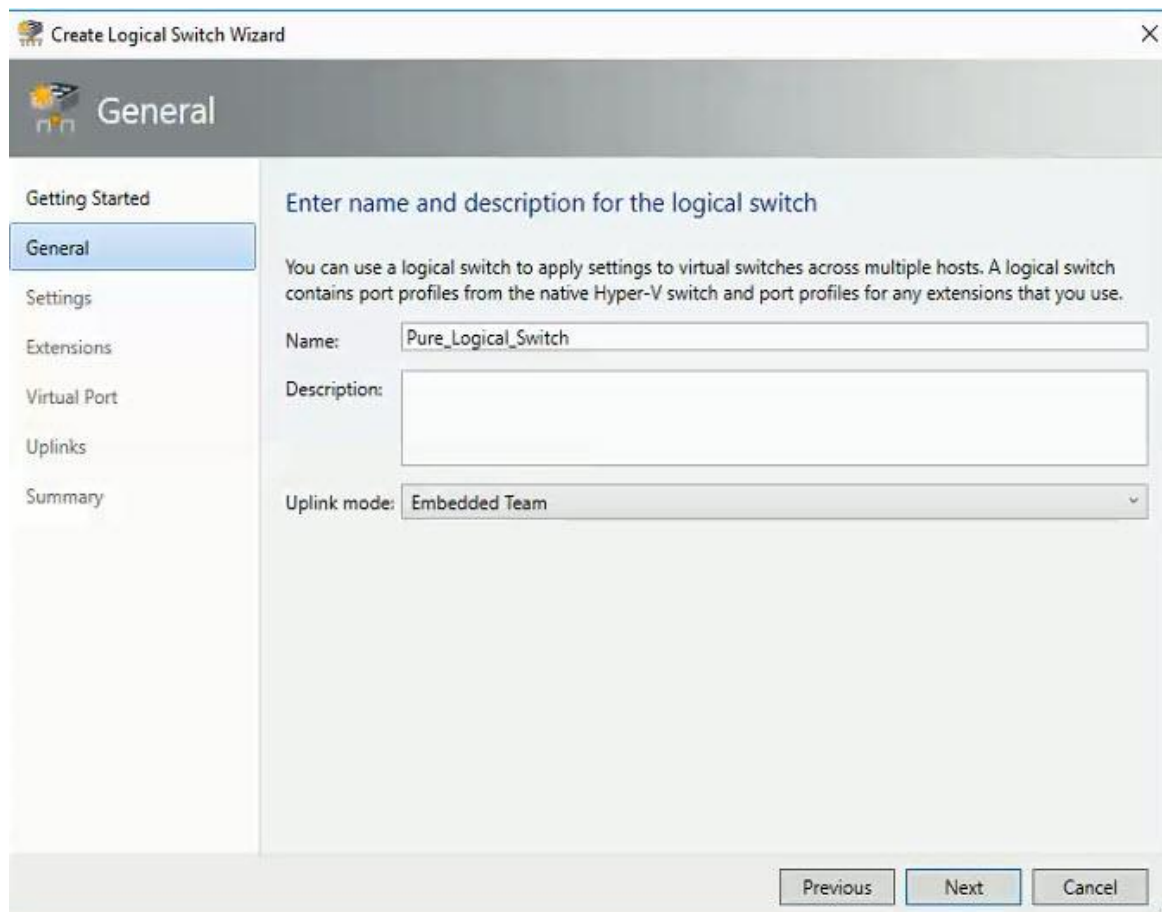
Create Logical Switch using SET

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

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

The logical switch will bring all of the components together. To create the Logical Switch, complete the following steps:

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

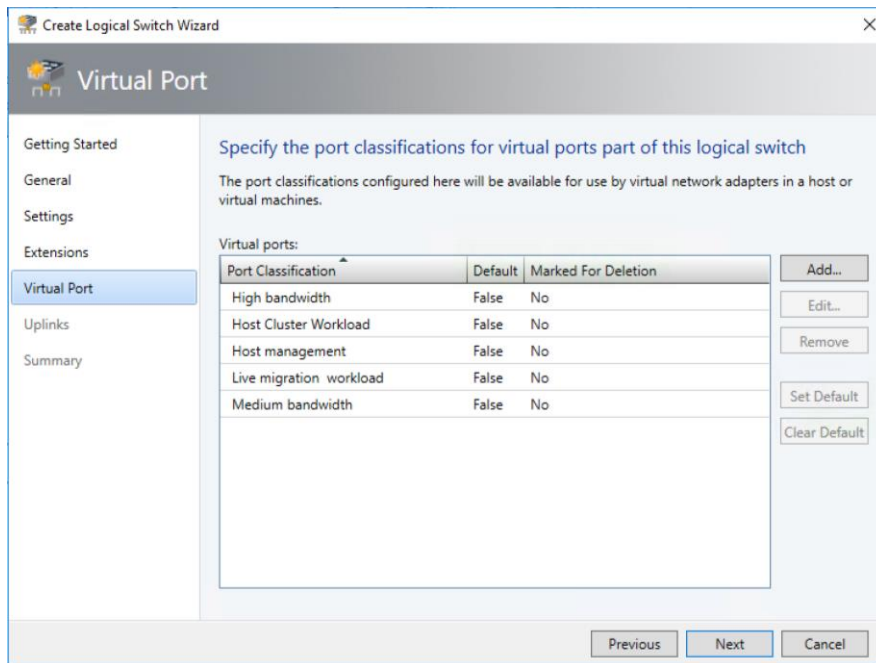
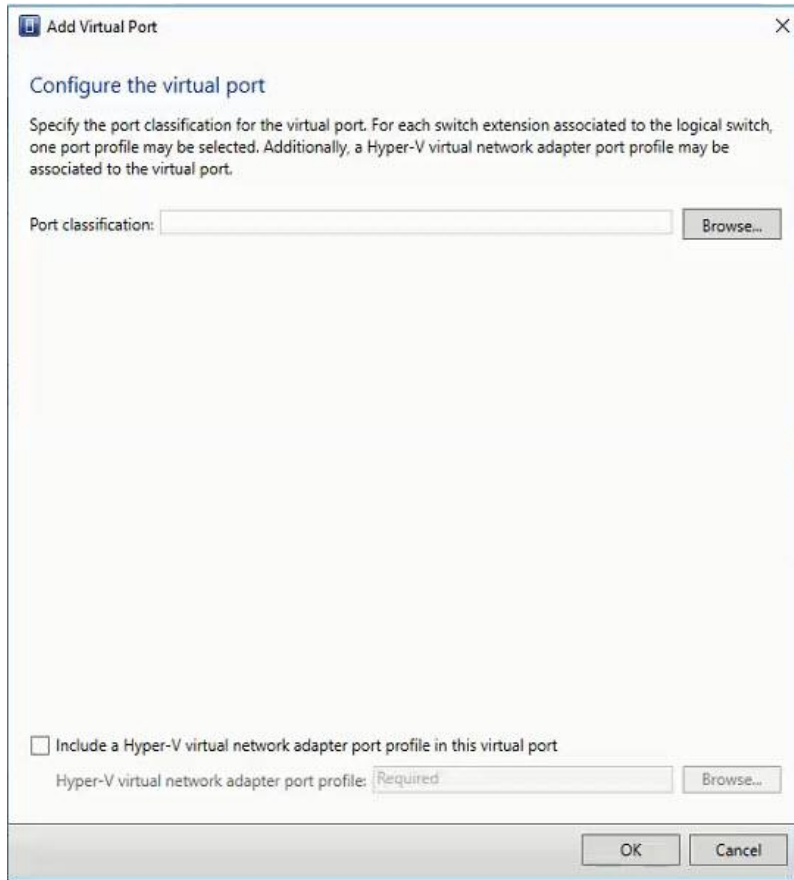


5. Select the minimum bandwidth mode as Weight, which quantifies minimum bandwidth for workloads, click Next
6. In Extensions selection window, leave the default and click Next.

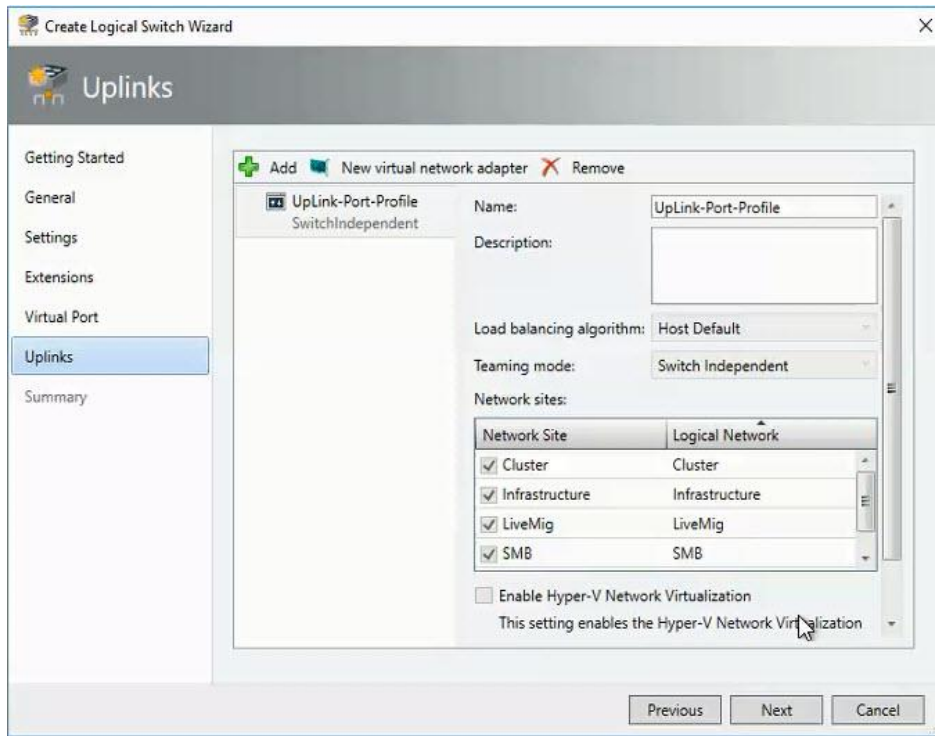
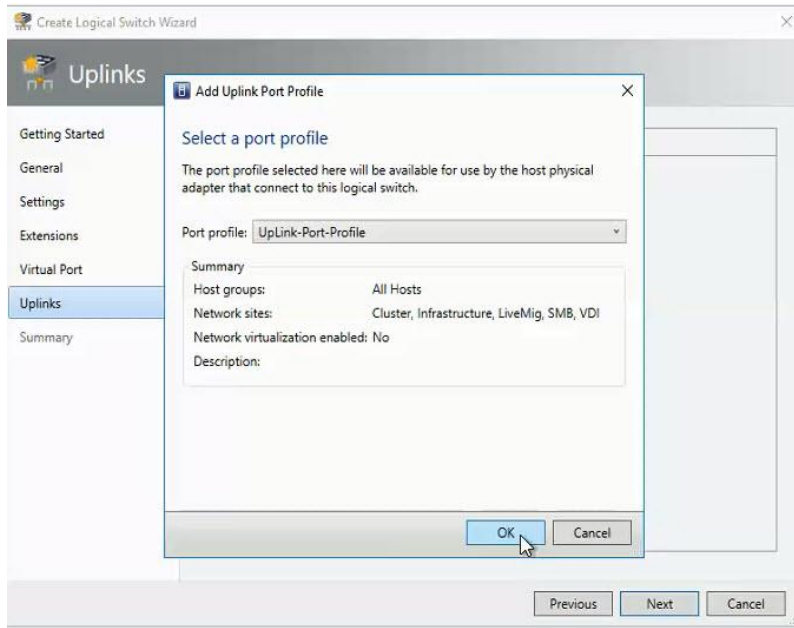


The list of installed virtual switch extensions is displayed and can be selected for deployment as part of the logical switch usage. This can be changed in the future if required.

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



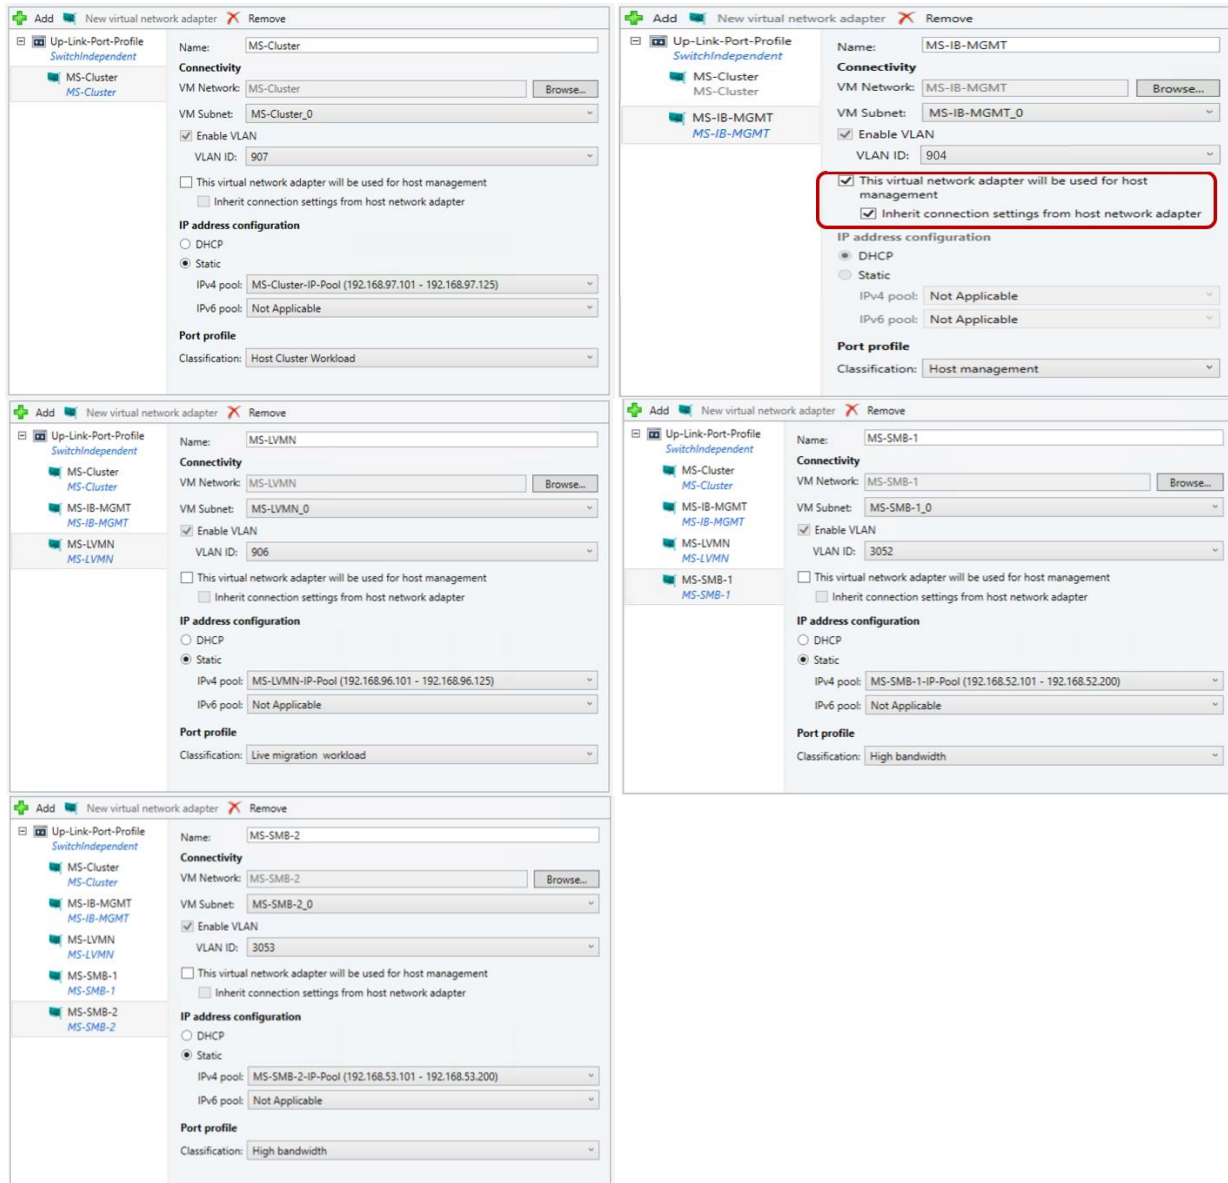
8. In the Uplinks window, click Add and then select Existing Uplink Port Profile - Up-Link-Port-Profile.



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



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



10. Click Finish to create the logical switch.

Fabric - Storage

Pure Storage SMI-S Provider Configuration

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



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

To add a storage device, complete the following steps:

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

The screenshot shows the 'Add Storage Devices Wizard' dialog box with the 'Select Provider Type' step selected. The title bar reads 'Add Storage Devices Wizard'. The main heading is 'Select Provider Type'. On the left, a navigation pane lists 'Select Provider Type', 'Specify Discovery Scope', 'Gather Information', 'Select Storage Devices', and 'Summary'. The main area is titled 'Select a storage provider type' and contains the following text: 'Before you begin this wizard, you might have to manually install storage provider software. Select the storage provider type that matches the type of device you want to manage.' Below this text are three radio button options:

- Windows-based file server
- SAN and NAS devices discovered and managed by a SMI-S provider
- SAN devices managed by a native SMP provider
- Fibre Channel fabric discovered and managed by a SMI-S provider

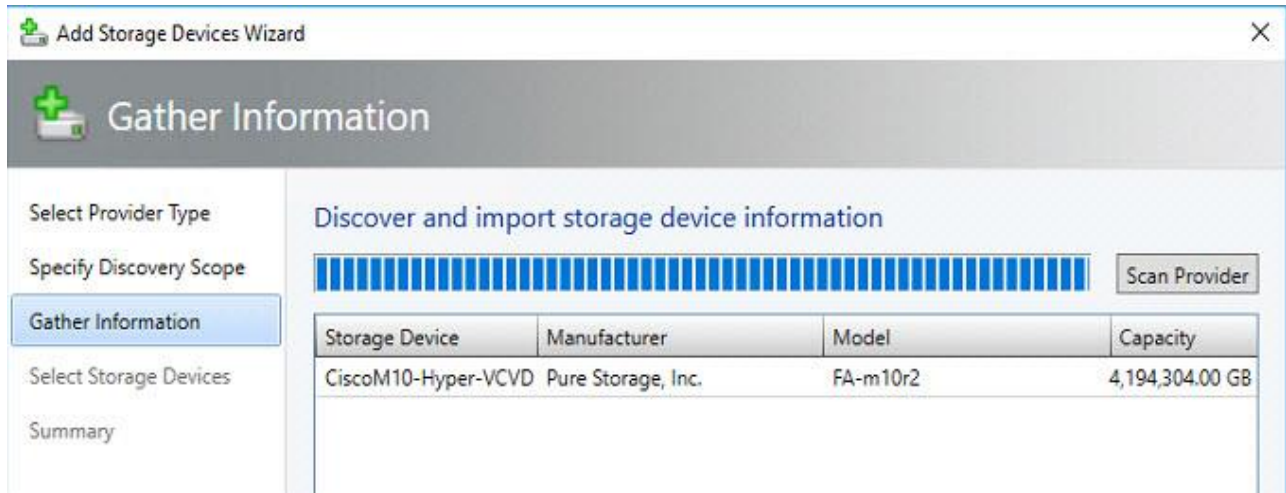
3. In Specify Discovery Scope, select Protocol - SMI-S CIMXML, add the IP address/FQDN, and add the port used to connect to the provider on the remote server. You can enable SSL if you're using CIMXML. Then specify an account for connecting to the provider.

The screenshot shows the 'Add Storage Devices Wizard' dialog box with the 'Specify Discovery Scope' step selected. The title bar reads 'Add Storage Devices Wizard'. The main heading is 'Specify Discovery Scope'. On the left, a navigation pane lists 'Select Provider Type', 'Specify Discovery Scope', 'Gather Information', 'Select Storage Devices', and 'Summary'. The main area is titled 'Specify protocol and address of the storage SMI-S provider' and contains the following fields:

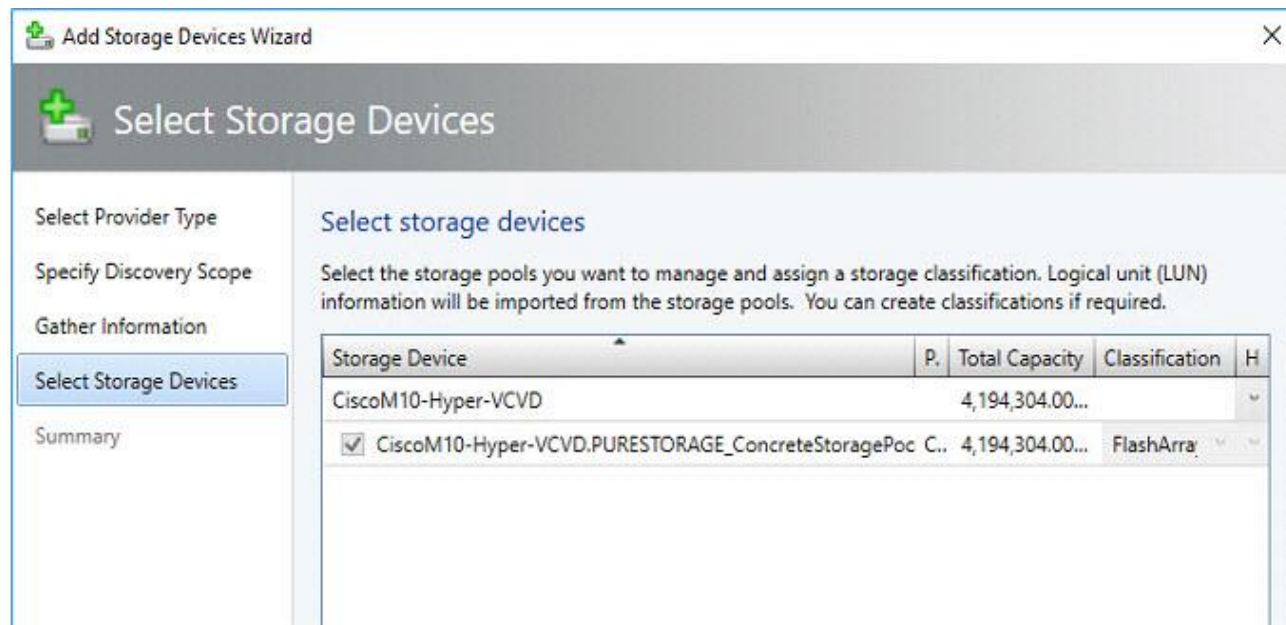
- Protocol: A dropdown menu with 'SMI-S CIMXML' selected.
- Provider IP address or FQDN: A text box containing '10.10.20.200'.
- TCP/IP port: A spin box with '5989'.
- Use Secure Sockets Layer (SSL) connection
- Run As account: A text box containing 'pureuser' and a 'Browse...' button.

4. In Gather Information, VMM automatically tries to discover and import the storage device information.

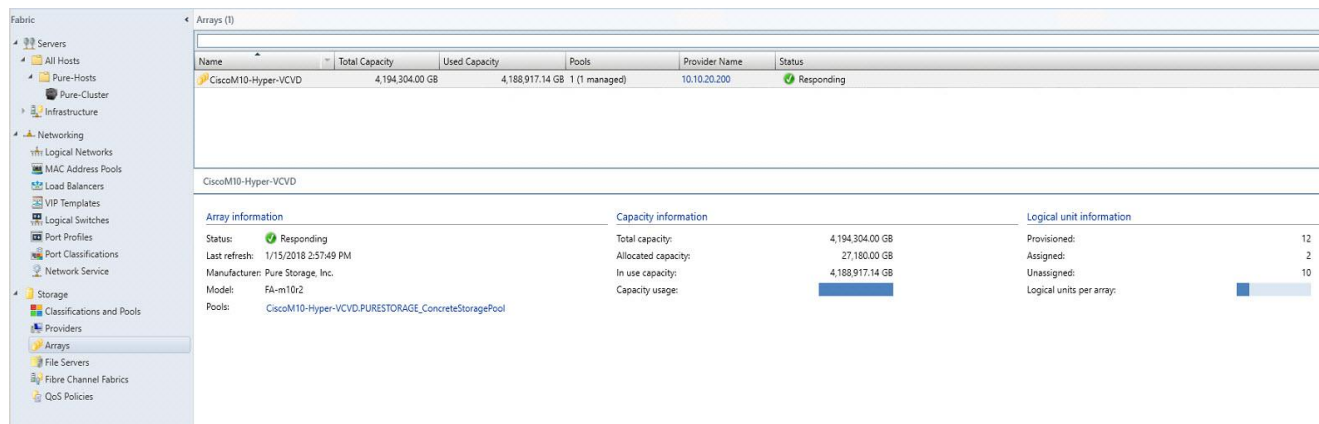
- If the discovery process succeeds, the discovered storage arrays, storage pools, manufacturer, model, and capacity are listed as shown in the below figure. When the process finishes, click Next.



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



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



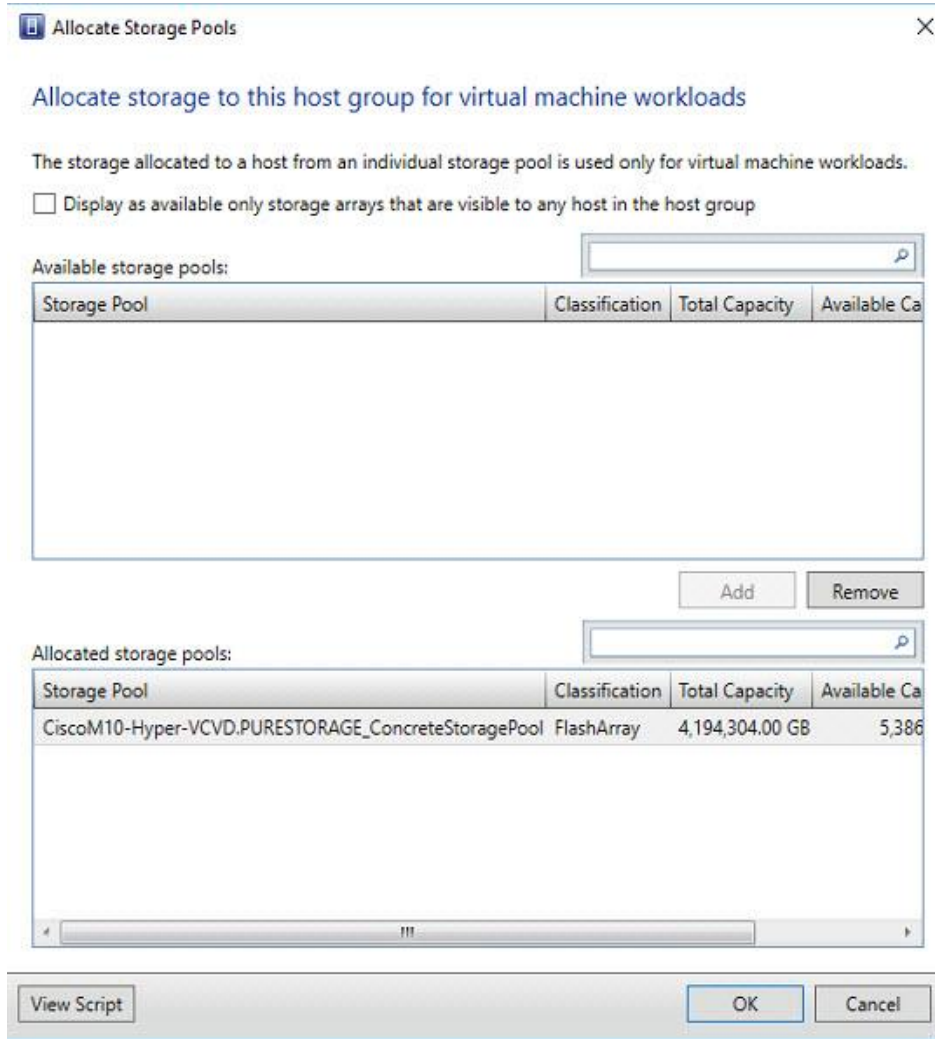
Allocate a Storage Pool to a Host Group

For this solution, the necessary SAN LUNs required for deploying Windows Hyper-V clusters were already created on the array before the SMI-S integration with VMM section. During the integration process of SMI-S with VMM, these storage pools were classified and associated with the appropriate host groups.

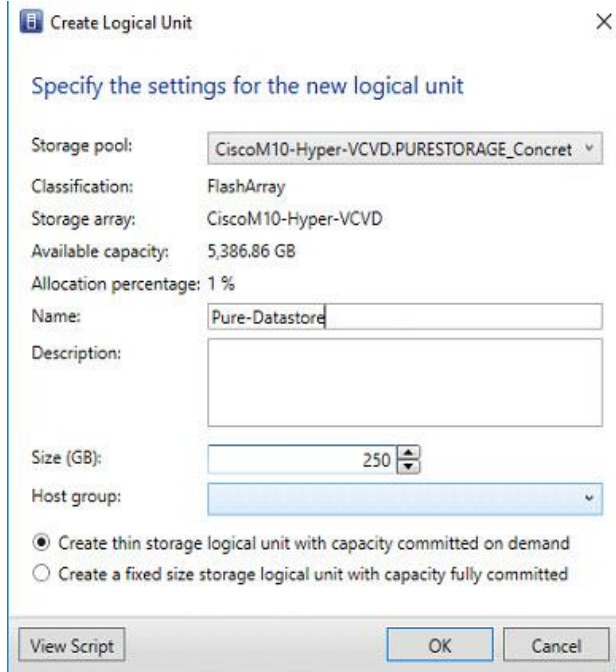
The steps in the following sections show how to create storage pools and LUNs from the VMM console after the SMI-S integration as an example.

To allocate a storage pool to a host group, complete the following steps:

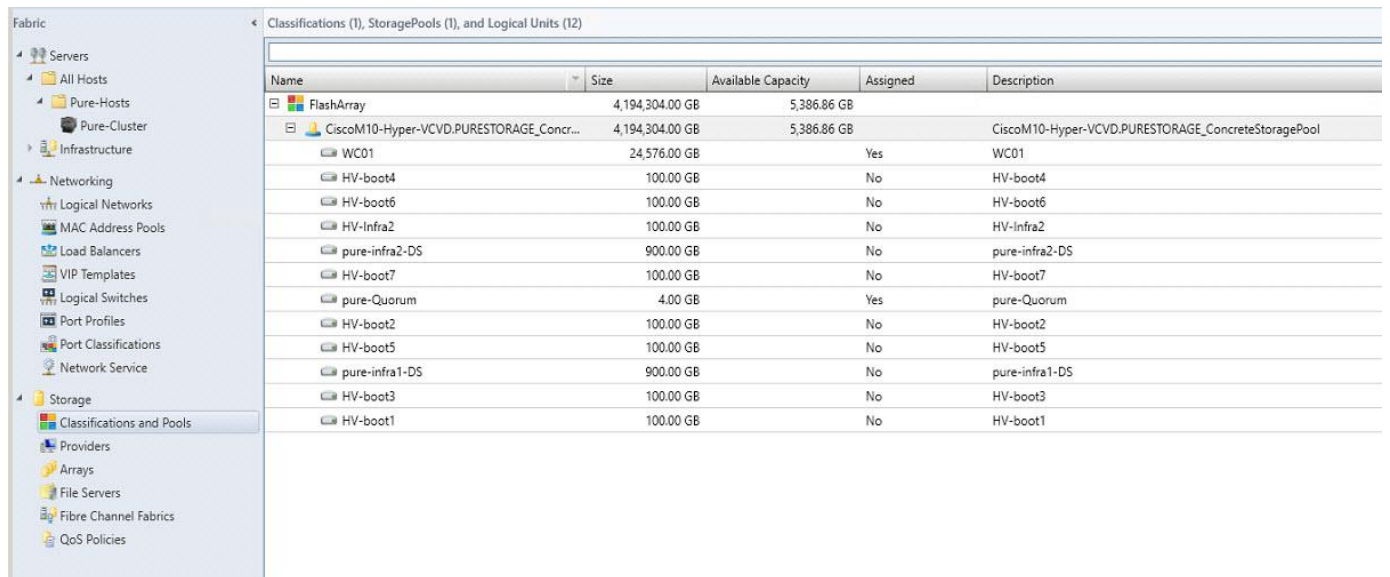
1. Click Fabric > Storage > Allocate Capacity and click the host group.
2. The total and available storage capacity information is displayed for the host group. The storage capacity information includes the total and available capacity for both local and remote storage, and total and available allocated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



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



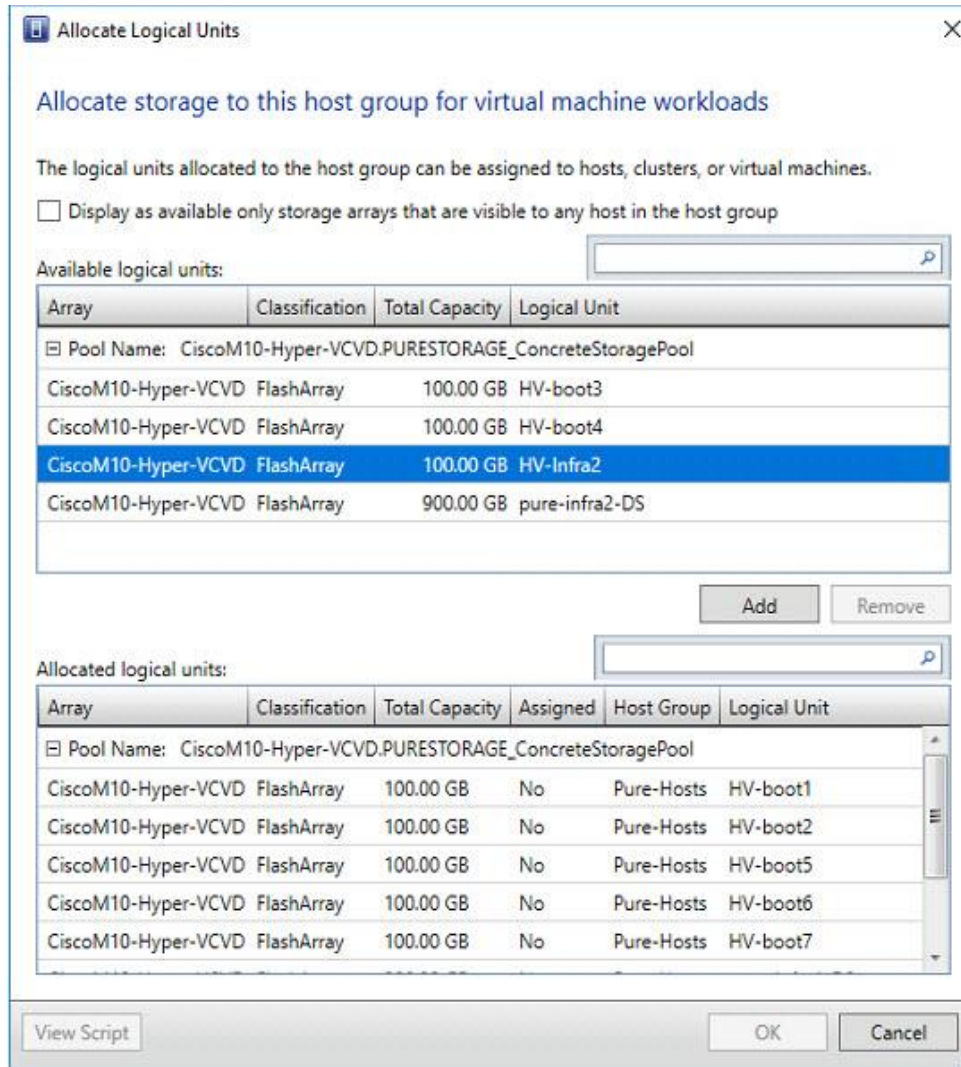
6. Verify that the LUN was created in Fabric Resources > Classifications, Storage Pools, and Logical Units.



Allocate a LUN to a Host Group

To allocate a LUN to a host group, complete the following steps:

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



Fabric – Servers – II

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

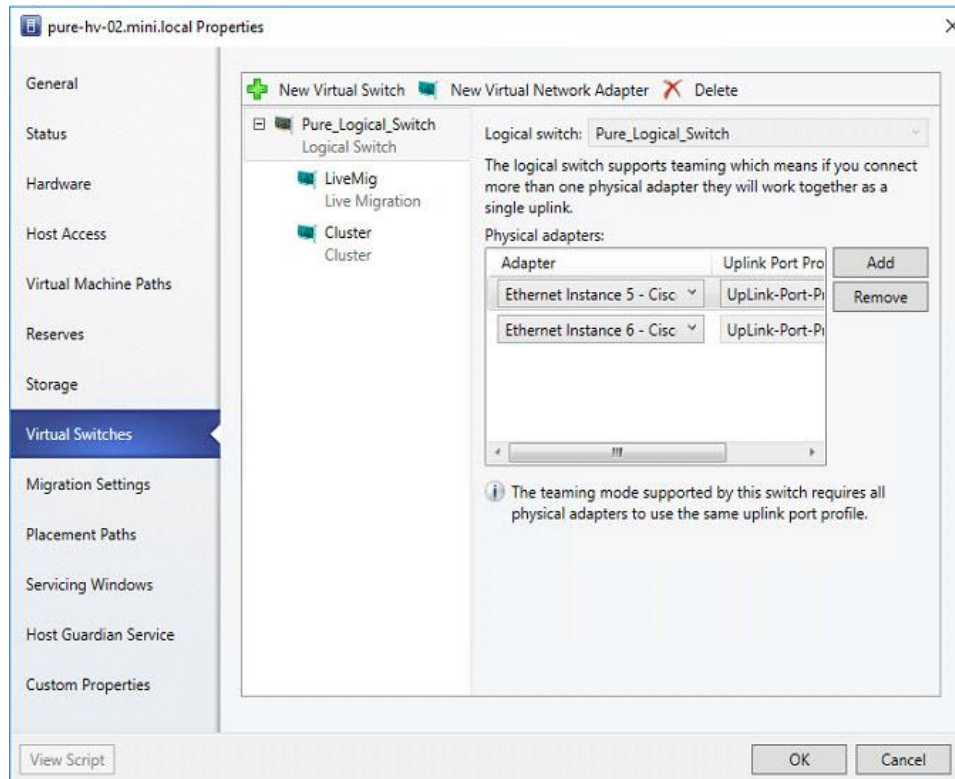
- Configure Network on Hosts by Applying Logical Switch
- Deploy Hyper-V Failover Cluster

Configure Network on Host – Applying Logical Switch

To apply the logical switch, complete the following steps:

1. Click Fabric > Storage > All Hosts > host group > Hosts > Host > Properties > Virtual Switches.
2. Open Fabric > Servers > click New Virtual Switch.

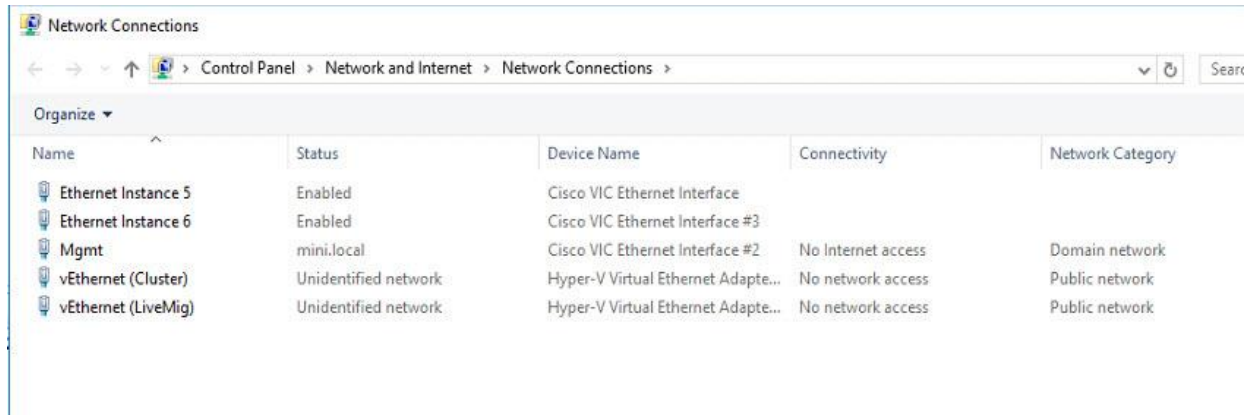
3. Select the logical switch you created. Under Adapter, select both the physical adapter to apply to the logical switch.



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



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



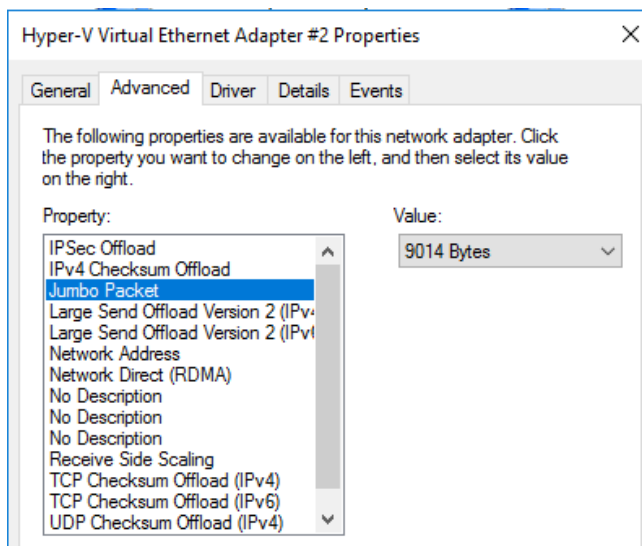
Enable Jumbo Frames

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

- HV-LiveMig
- HV-VDI
- HV-CSV

To enable jumbo frames, complete the following steps:

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



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

```

PS C:\Users\Administrator> Get-NetAdapterAdvancedProperty -DisplayName "jumbo Packet" | ft -AutoSize
Name                DisplayName  DisplayValue RegistryKeyword RegistryValue
-----
vEthernet (Cluster) Jumbo Packet 9014 Bytes   *JumboPacket  {9014}
vEthernet (LiveMig) Jumbo Packet 9014 Bytes   *JumboPacket  {9014}

PS C:\Users\Administrator> ping 10.10.20.81 -l 8972

Pinging 10.10.20.81 with 8972 bytes of data:
Reply from 10.10.20.81: bytes=8972 time<1ms TTL=128
Reply from 10.10.20.81: bytes=8972 time<1ms TTL=128
Reply from 10.10.20.81: bytes=8972 time<1ms TTL=128
Reply from 10.10.20.81: bytes=8972 time<1ms TTL=128

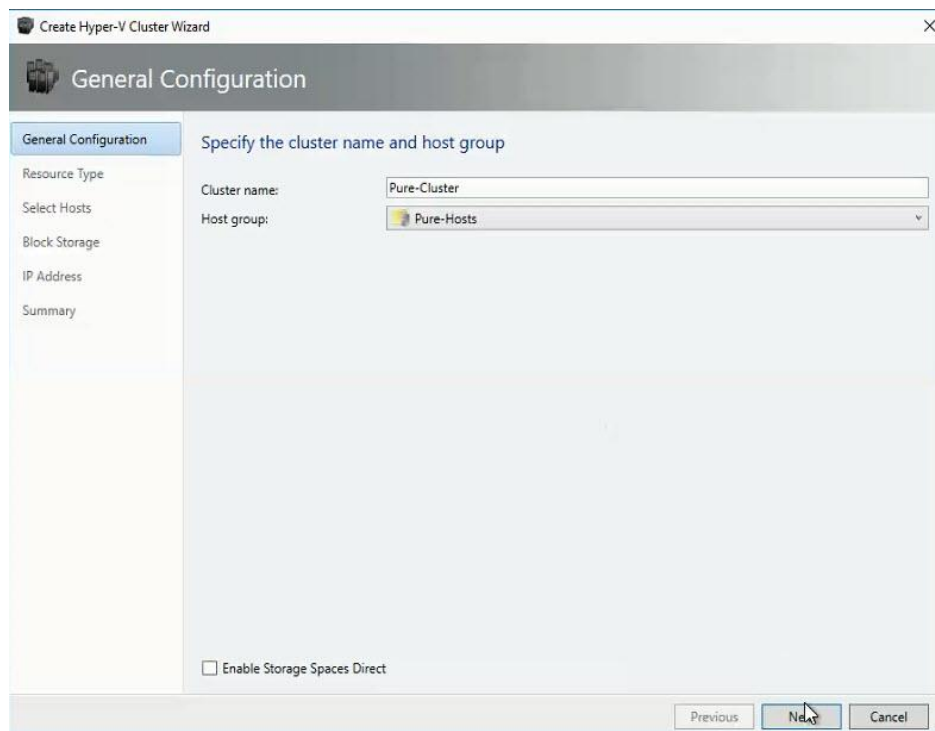
Ping statistics for 10.10.20.81:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
PS C:\Users\Administrator>

```

Deploy Hyper-V Cluster

To deploy the Hyper-V cluster, complete the following steps:

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

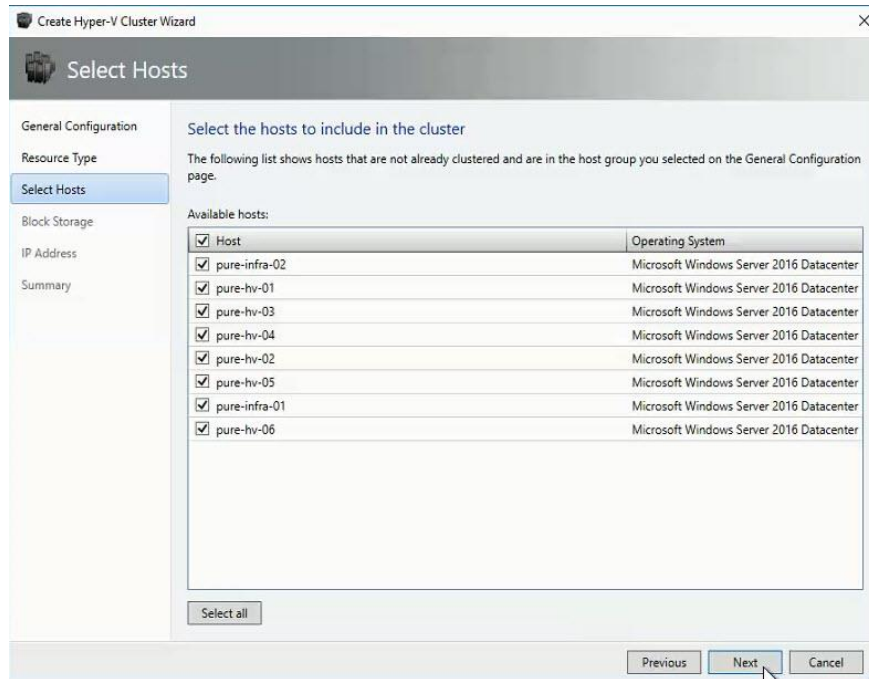


3. In Resource Type, select the Run As account that you will use to create the cluster.

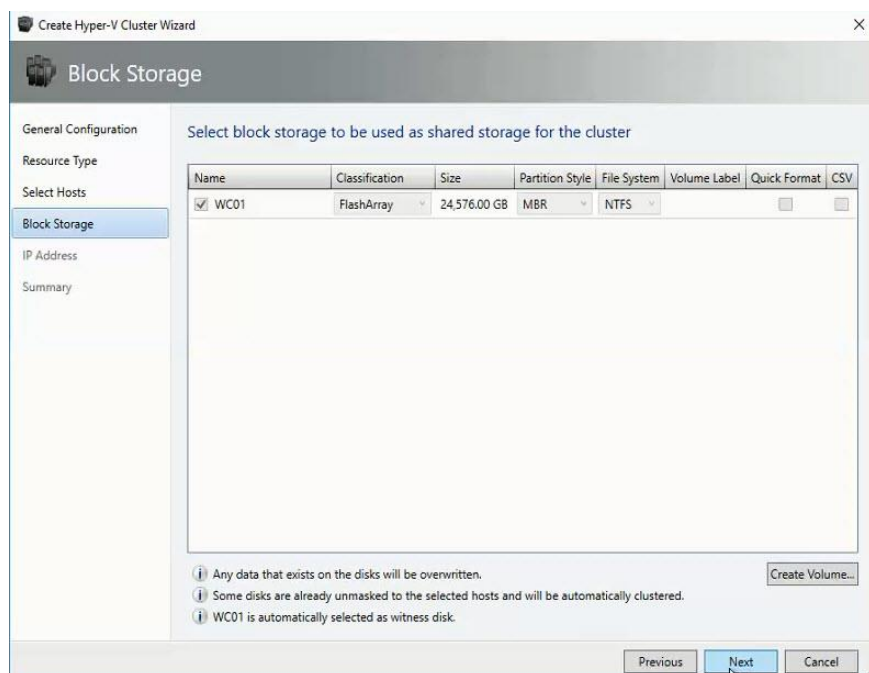


The accounts that you use must have administrative permissions on the servers that will become cluster nodes and must belong to the same domain as the Hyper-V hosts that you want to cluster. Also, the account requires Create Computer objects permission in the container that is used for Computer accounts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

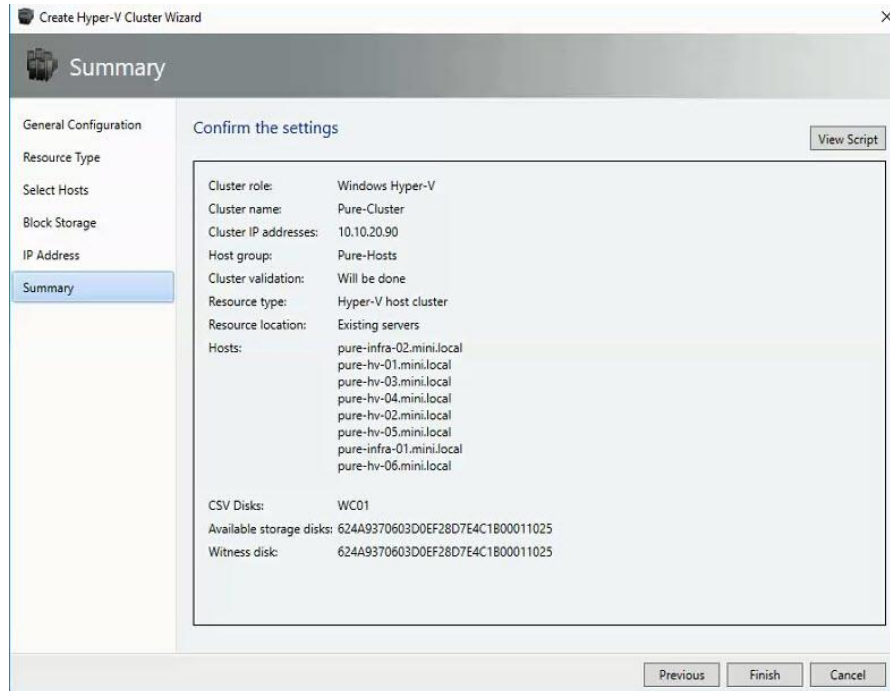
- In Nodes, select the Hyper-V host servers that you want to include in the cluster.



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



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



The screenshot below shows the PowerShell Script for creating the cluster:

```
# Get Host Group 'Pure-Hosts'
$HostGroup = Get-SCVMHostGroup -ID "cb6c4184-385e-4053-a007-40f7b7579dff"
# Get RunAs Account 'Administrator'
$AdminRunAsAccount = Get-SCRunAsAccount -ID "5694ea18-b1a0-4b45-bc58-612c9374edd0"

# Get Host 'pure-infra-02.mini.local, pure-hv-01.mini.local, pure-hv-03.mini.local, pure-hv-04.mini.local, pure-hv-02.mini.local, pure-hv-05.mini.local, pure-infra-01.mini.local, pure-hv-06.mini.local'
$VMHosts = @()
$VMHosts += Get-SCVMHost -ID "c1526b84-6750-4800-b159-809321bfa2ed"
$VMHosts += Get-SCVMHost -ID "1a5dd5f5-4177-4993-a07d-db3a7902496c"
$VMHosts += Get-SCVMHost -ID "72e9f8fd-693b-4118-8e22-d3dde47d444b"
$VMHosts += Get-SCVMHost -ID "f3c7c32f-c3a3-4728-b2c6-19920e9ec925"
$VMHosts += Get-SCVMHost -ID "c88851bb-4c24-4e97-b306-216775db9bc5"
$VMHosts += Get-SCVMHost -ID "0c6f8032-a6cd-4e44-afbe-ed9f25f5fa71"
$VMHosts += Get-SCVMHost -ID "42f18adb-8eb1-451d-82bf-8f2fef6a109e"
$VMHosts += Get-SCVMHost -ID "2e3ad0a7-39aa-44d6-810f-d8fd7c9dabc7"

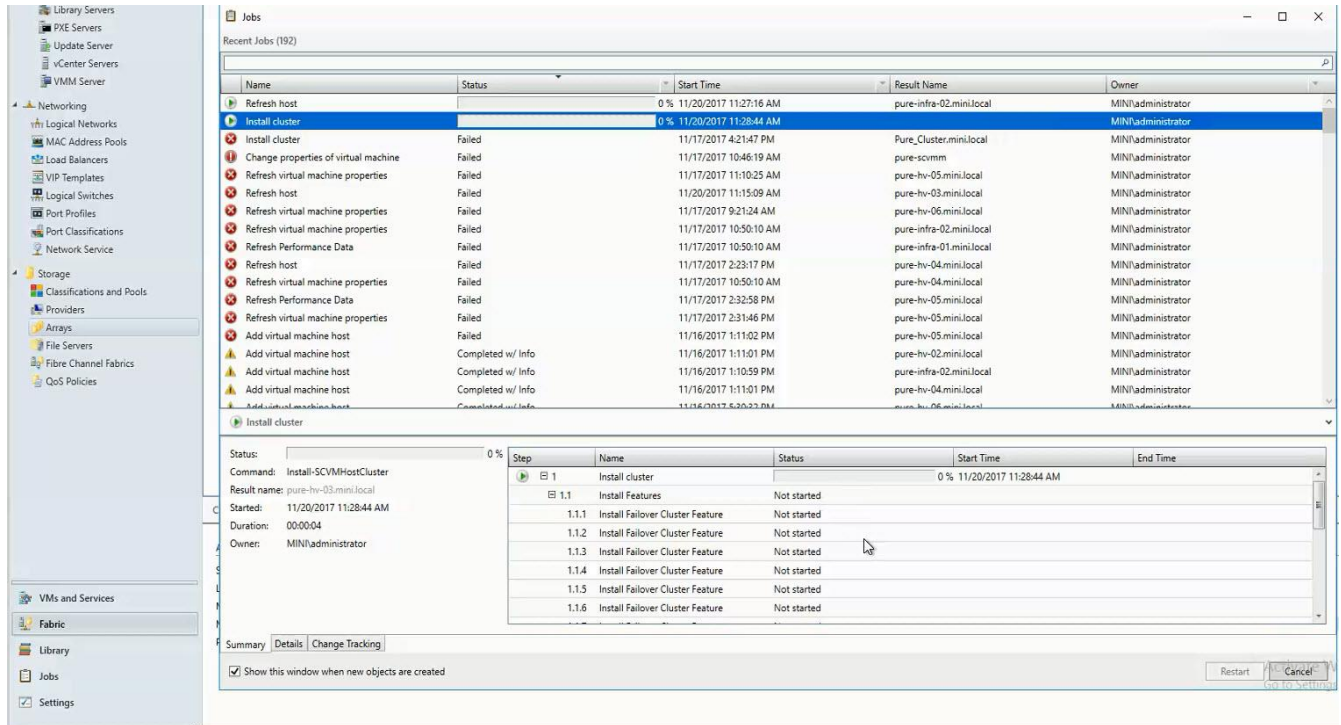
# Get Host 'pure-infra-02.mini.local'
$VMHost = Get-SCVMHost -ID "c1526b84-6750-4800-b159-809321bfa2ed"
$HostDisk = Get-SCStorageDisk -VMHost $VMHost | where { $_.SMLunId -eq "624A9370603D0EF28D7E4C1B00011025" }
Mount-SCStorageDisk -StorageDisk $HostDisk -MasterBootRecord -VolumeLabel "" -FullFormat -JobGroup "7f6927ae-5578-4e59-ba54-726dfade828f"

$HostDisk = Get-SCStorageDisk -VMHost $VMHost | where { $_.SMLunId -eq "624A9370603D0EF28D7E4C1B00011024" }
Mount-SCStorageDisk -StorageDisk $HostDisk -GuidPartitionTable -VolumeLabel "WC01" -QuickFormat -JobGroup "7f6927ae-5578-4e59-ba54-726dfade828f" -CreateClusterSharedVolume

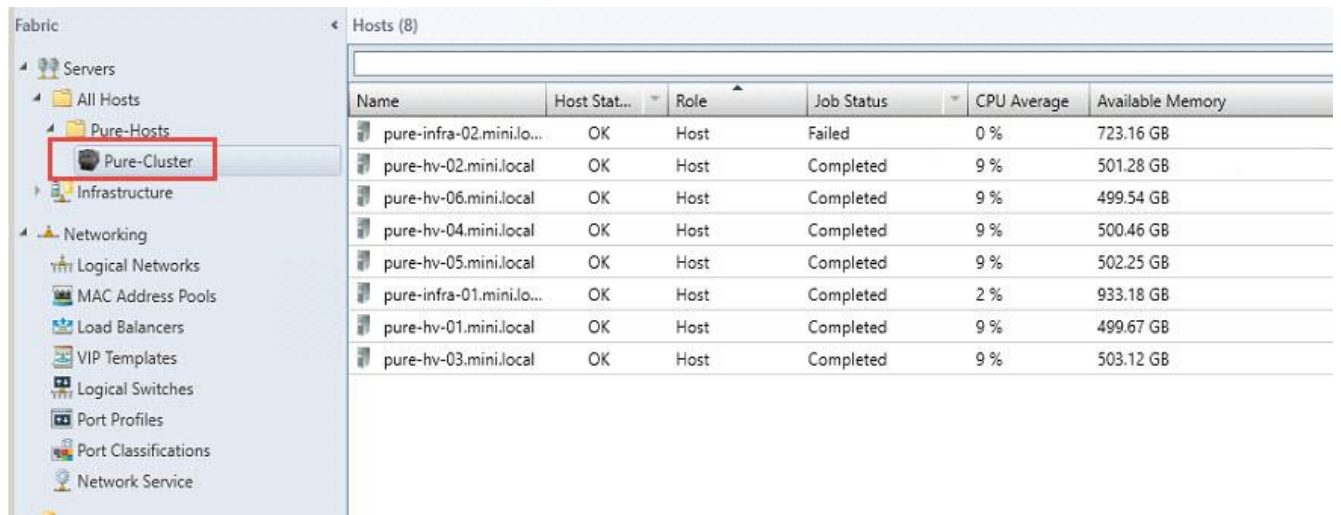
$WitnessDisk = Get-SCStorageDisk -Name "\\.\PHYSICALDRIVE6" | where { $_.ID -eq "0e1d7688-a131-49bc-8956-4777ab85c94a" }
$StaticIPAddress = @("10.10.20.90")

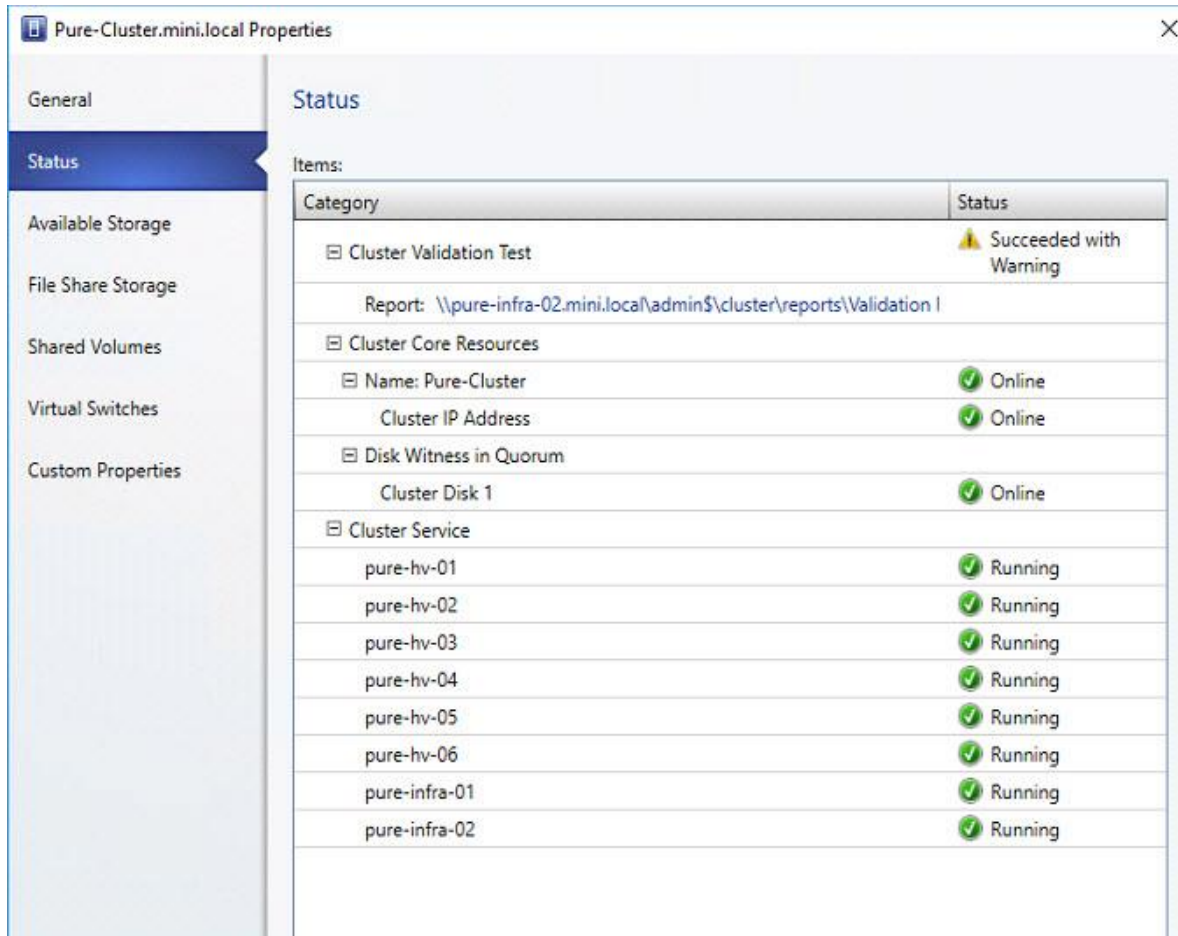
Install-SCVMHostCluster -ClusterName "Pure-Cluster" -JobGroup "7f6927ae-5578-4e59-ba54-726dfade828f" -RunAsynchronously -Credential $AdminRunAsAccount -VMHost $VMHosts -SetQuorumDisk $WitnessDisk -ClusterIPAddress $StaticIPAddress
```

8. You can go to the jobs workspace and click "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>Pure-Cluster host group in the fabric workspace. Right-click 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.

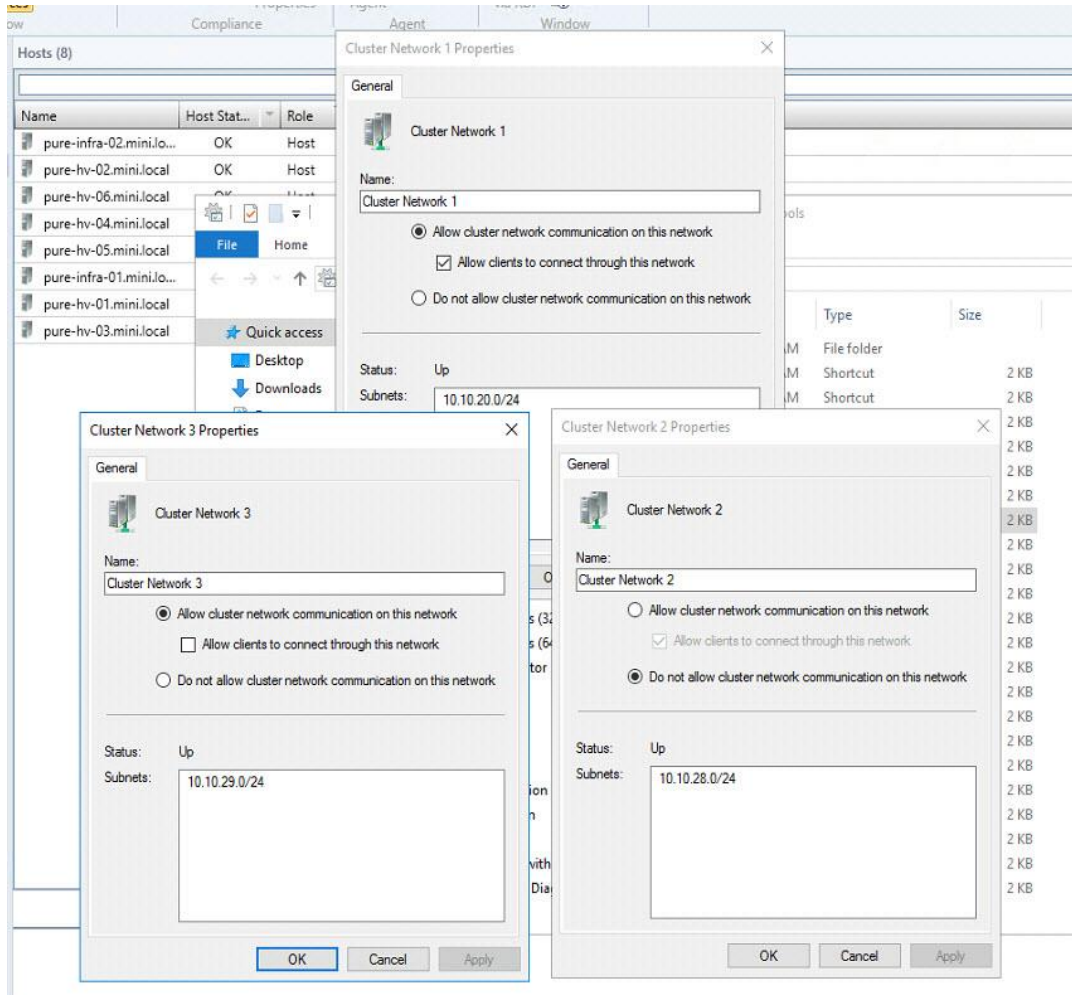
Table 4 Recommended Settings for Network Traffic

Network Type	Recommended Setting
Management	Both of the following: - Allow cluster network communication on this network - Allow clients to connect through this network

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

To configure a network to allow or not to allow cluster network communication, complete the following steps:

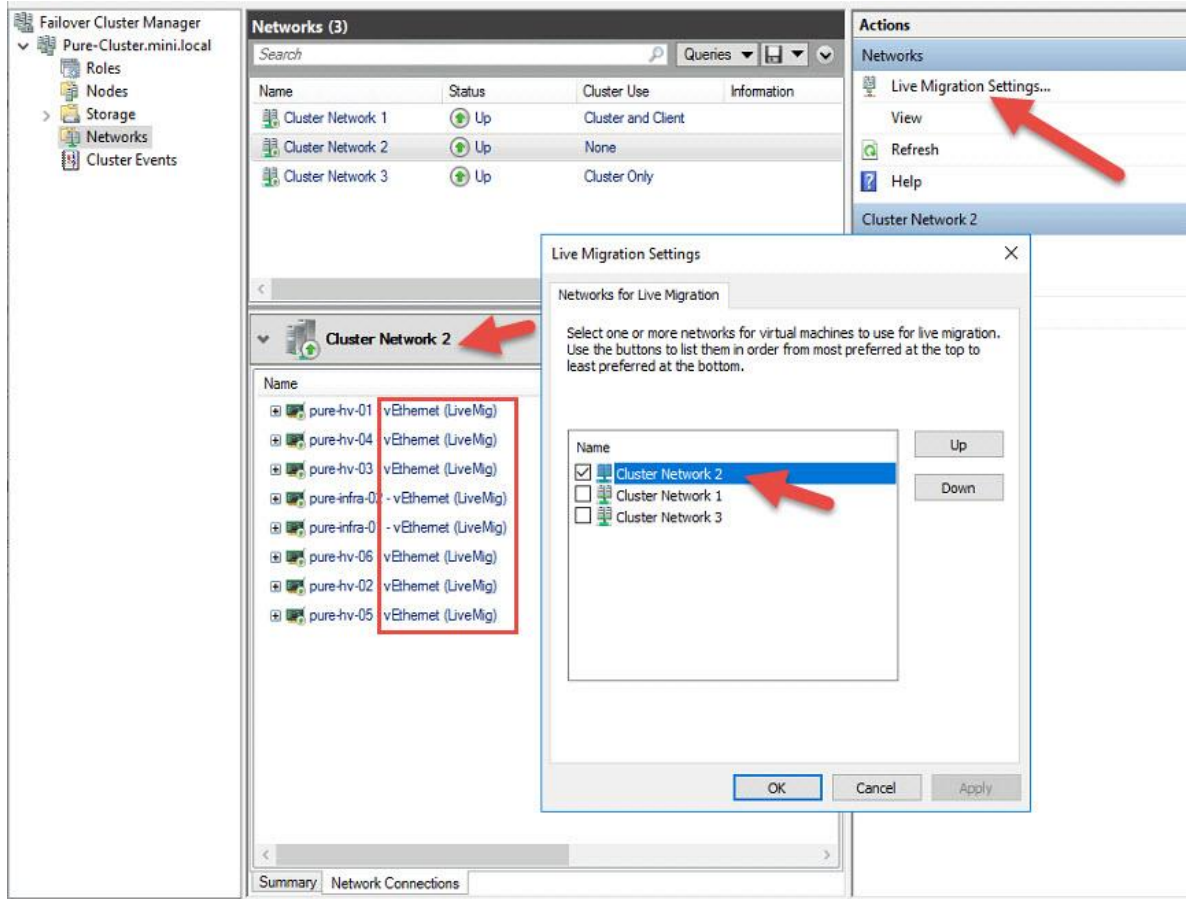
1. Open Failover Cluster Manager, click Networks in the navigation tree.
2. In the Networks pane right-click a network and then click Properties.



Live Migration Network Settings

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

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

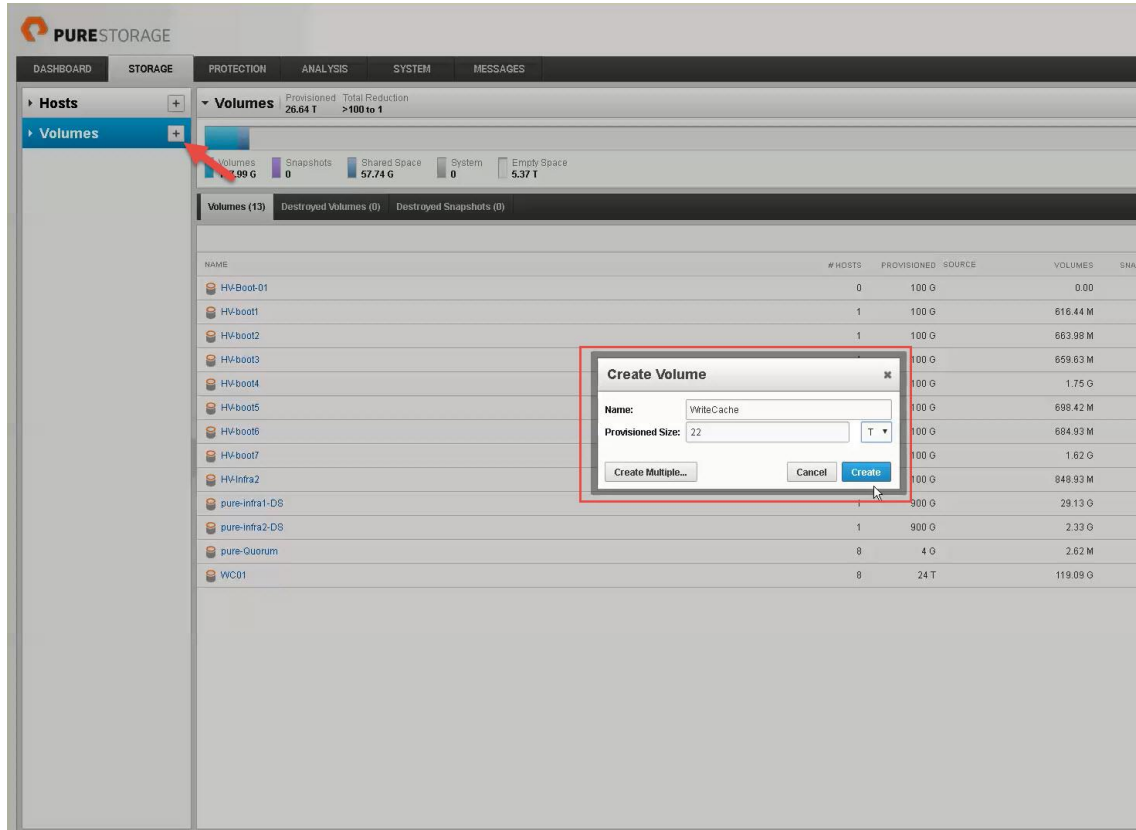


Deploy Volumes to the Hyper-V Cluster

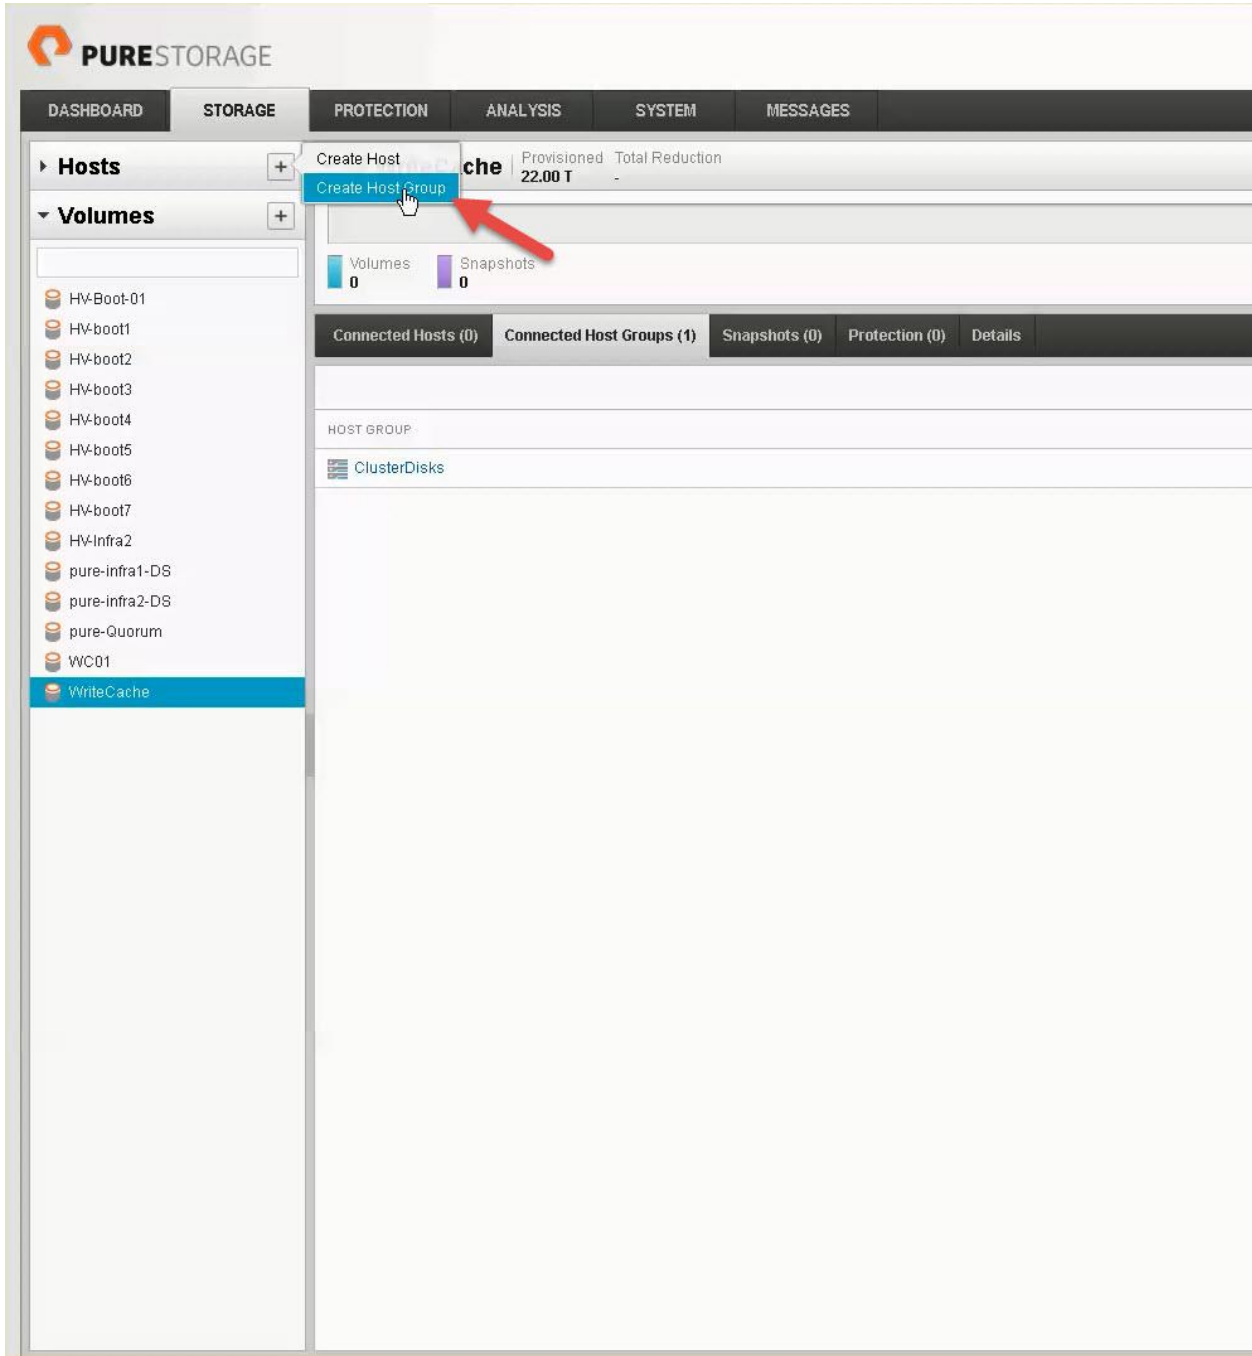
To deploy the Hyper-V cluster volumes, we used the Pure Array console with their Host Groups technology. **With Pure’s Host Group technology, we can present single volumes to groups of host. In this case, the cluster hosts.**

To achieve this task, complete the following steps:

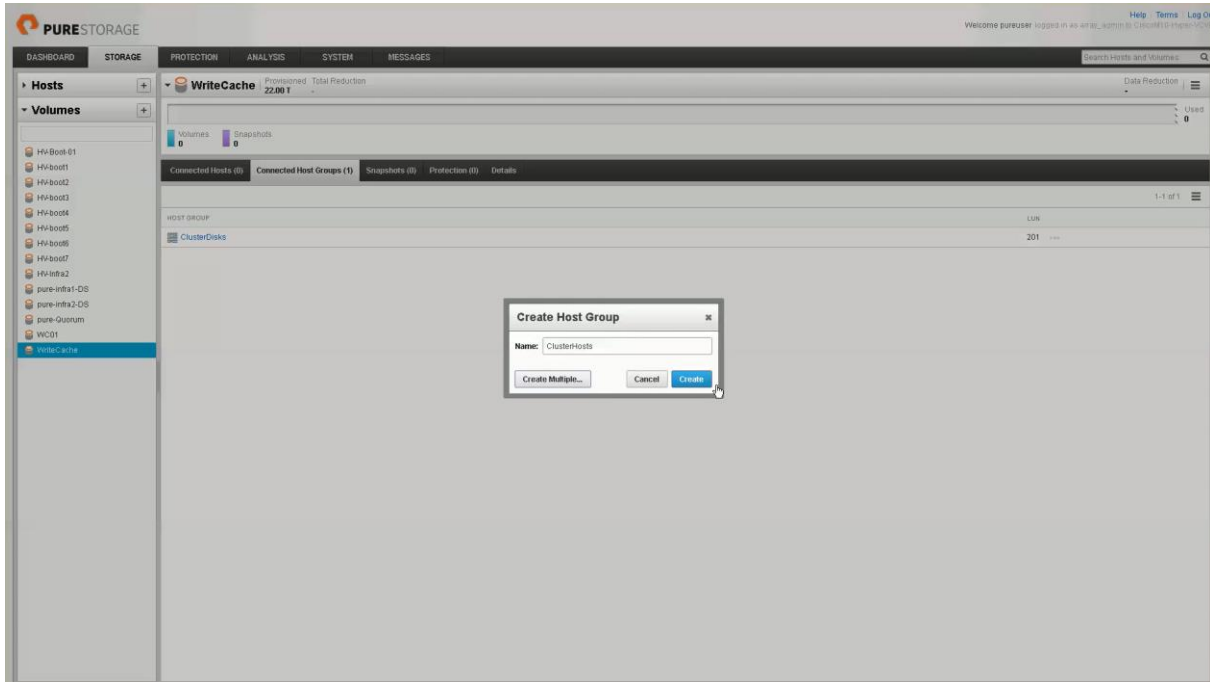
1. In the Pure Console, select ‘Volumes’ and the + sign to create the cluster volume.




2. Once the volume is created, it needs to be presented to the cluster hosts as shared storage.
3. In the Pure console, select 'Hosts' and the + sign to create a group of Hosts to present the cluster storage to.



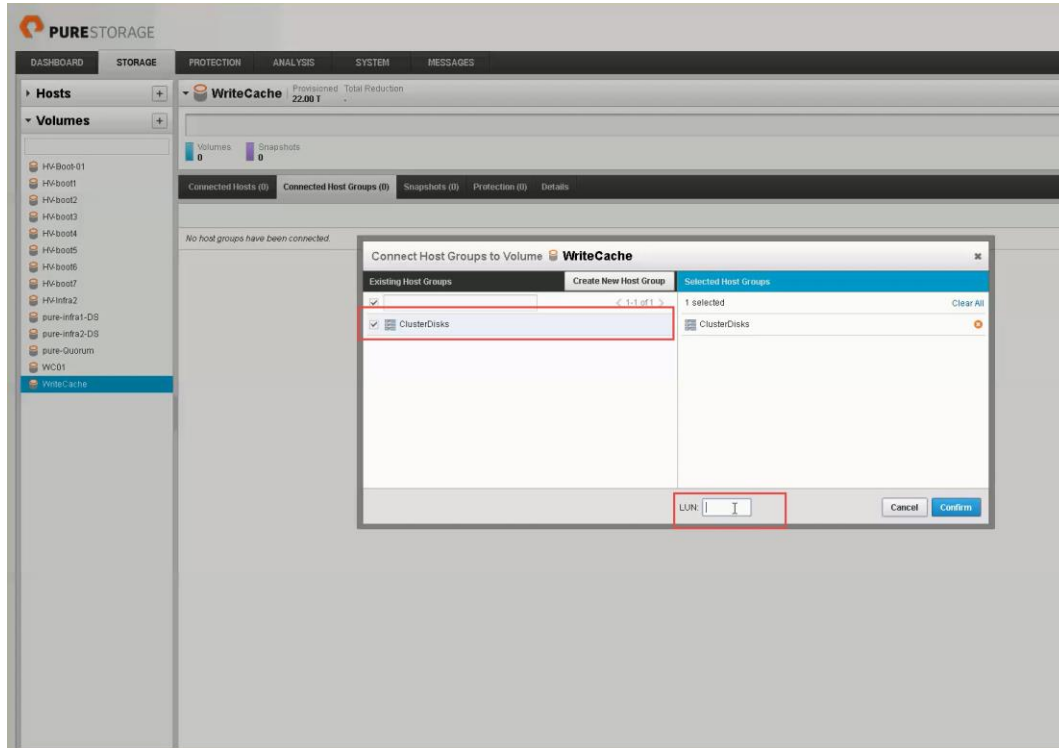
4. Name the Host Group and add the hosts previously created in the Boot LUNs step.




5. Once the Volumes and Host Groups are created, they need to be connected using the Pure Console.
6. In the Pure Console, navigate to Hosts -> click the Menu to the far right to select 'Connect Host Groups'.

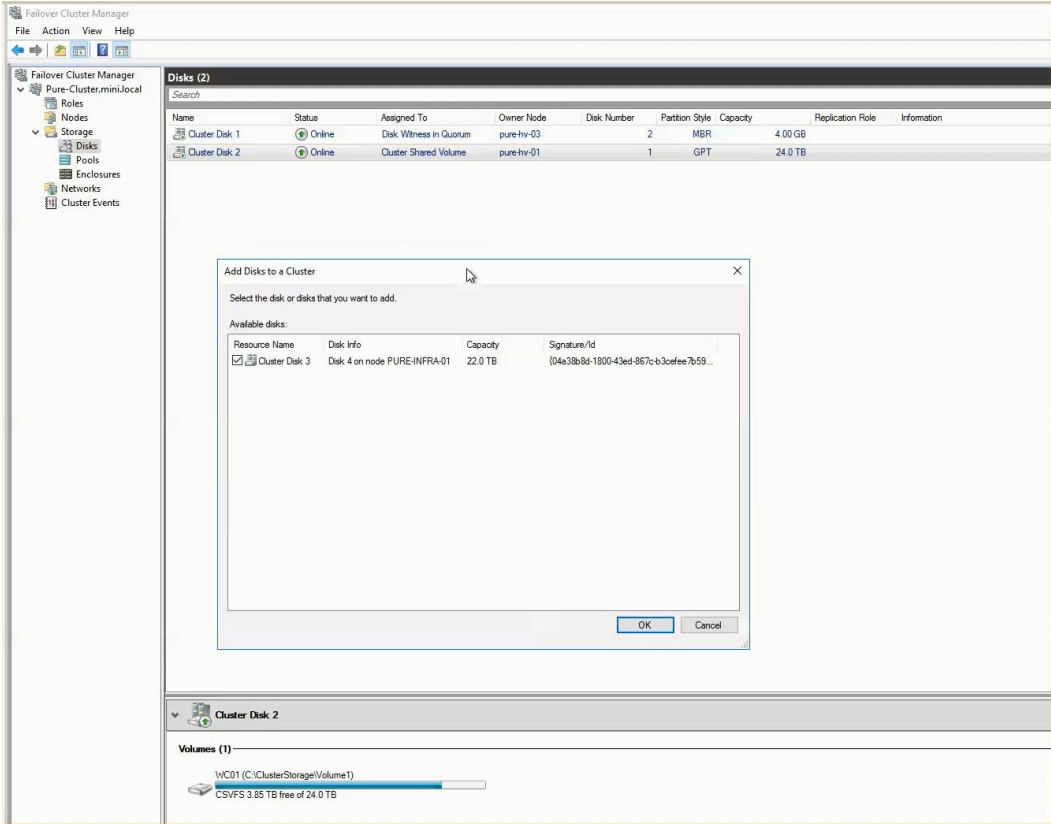


7. Select the Check box next to the Host Groups just created and assign a LUN ID for the volumes.



- Once the cluster disks are created, they can be added to the Microsoft cluster and used in Cluster Shared Volumes.





Cisco UCS Management Pack Suite Installation and Configuration

Cisco UCS Manager Integration with SCOM

About Cisco UCS Management Pack Suite

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

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

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


For more information, see:

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

Installing Cisco UCS Monitoring Service

To install Cisco UCS Monitoring Service, complete the following steps:

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



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



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

Adding a Firewall Exception for the Cisco UCS Monitoring Service

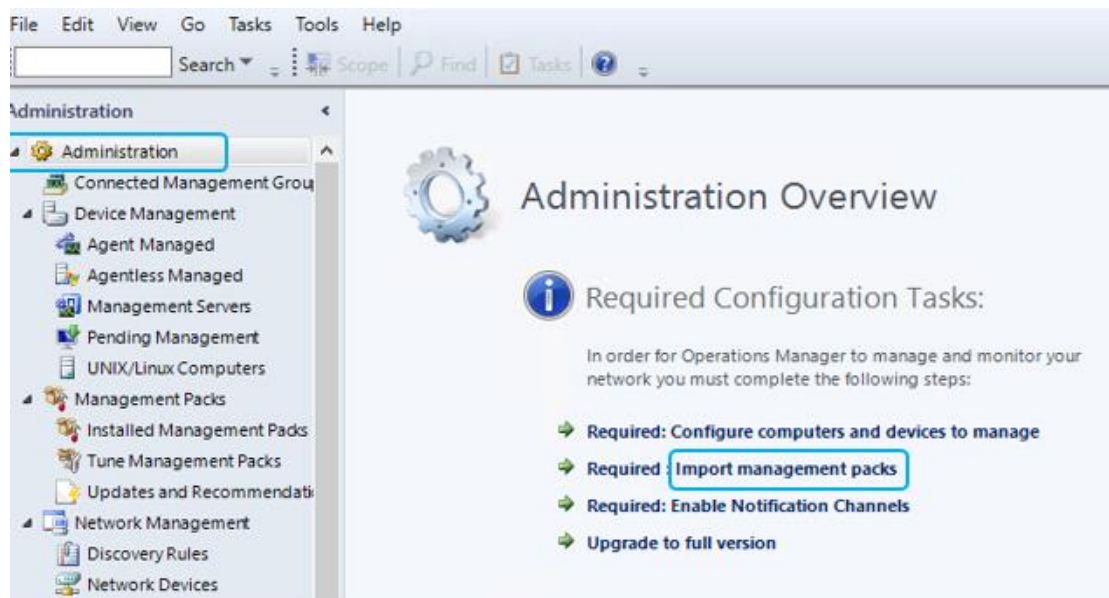
To add a firewall exception, complete the following steps:

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

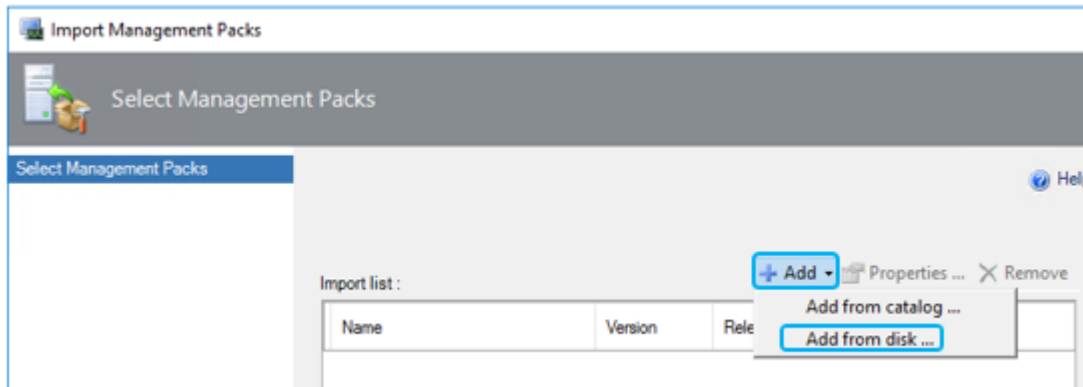
Installing the Cisco UCS Management Pack Suite

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

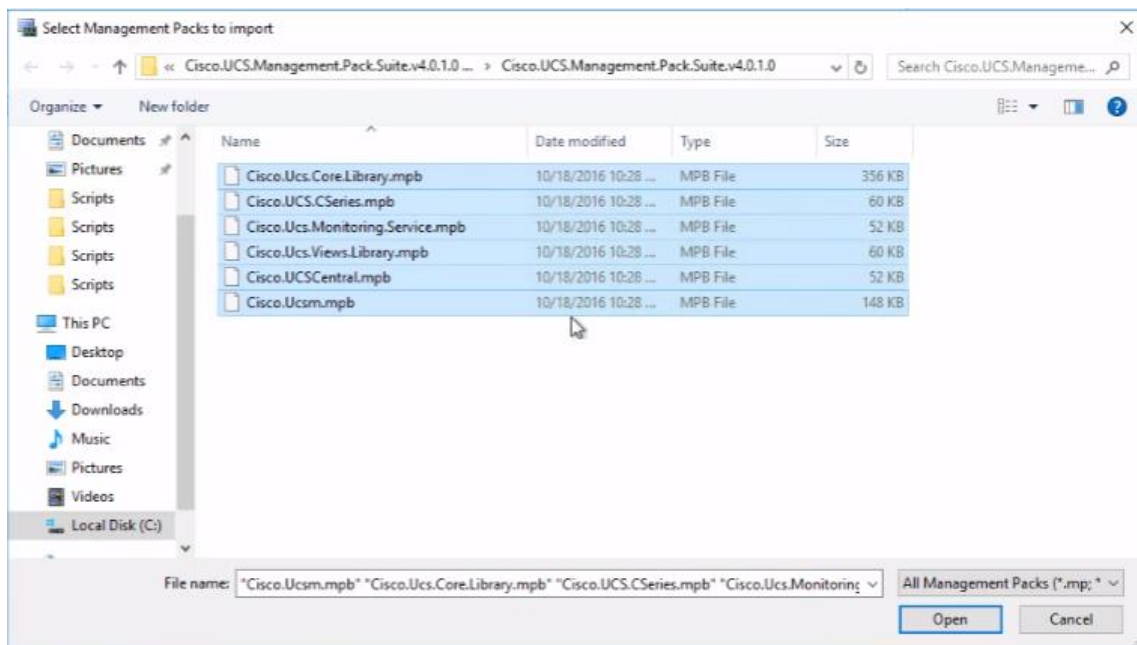
1. For importing Management Packs using Operations Manager console, you must have administrative privileges. For more information on the access privileges, see <https://technet.microsoft.com/en-in/library/hh212691.aspx>. On the Cisco.com download site for Cisco UCS Management Partner Ecosystem Software,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



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

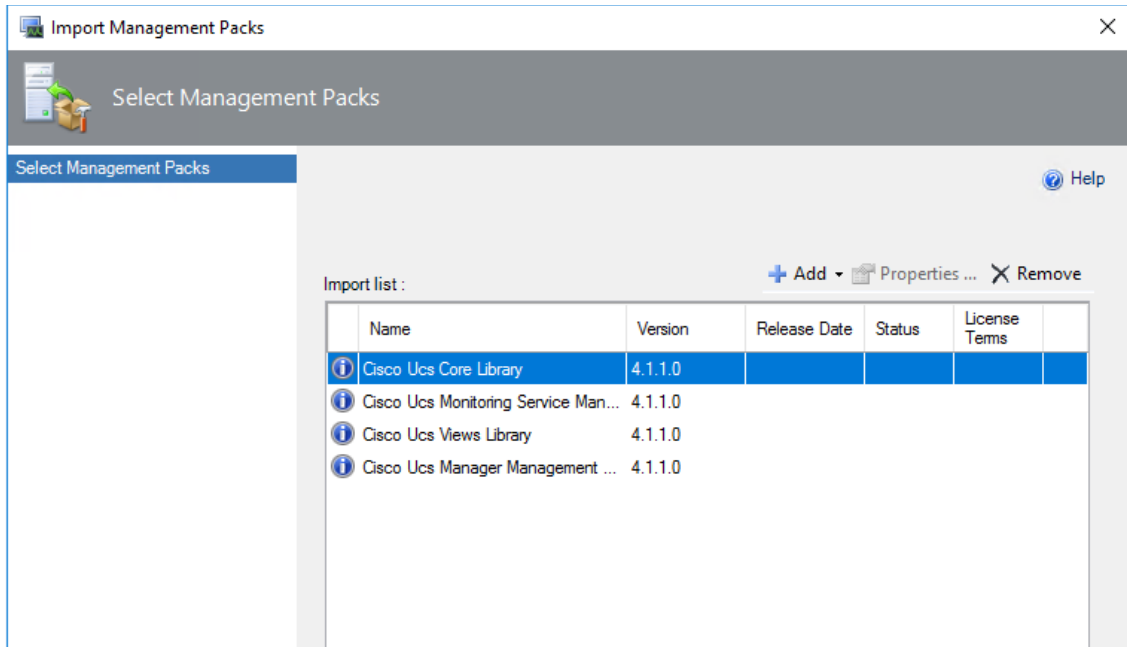


5. Click No, if you do not want to search the management pack dependencies online.
6. Navigate to the unzipped management pack suite files folder.
7. From the list of files, select the mandatory files:
 - Cisco.Ucs.Views.Library.mpb
 - Cisco.Ucs.Core.Library.mpb
 - Cisco.Ucs.Monitoring.Service.mpb
8.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

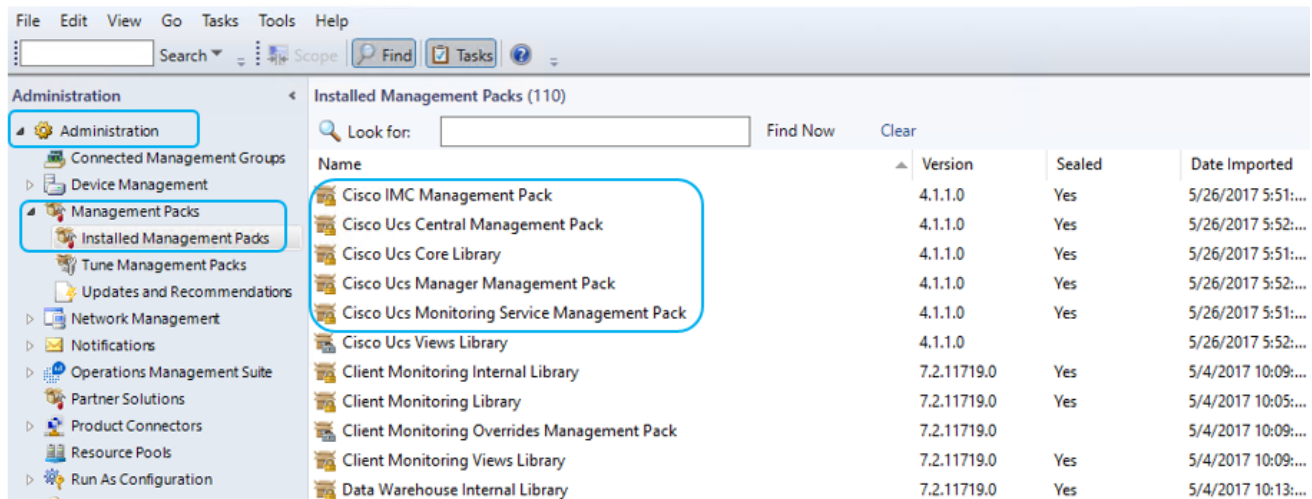


9. Click Open.
10. Click Install on the Import Management Packs page.

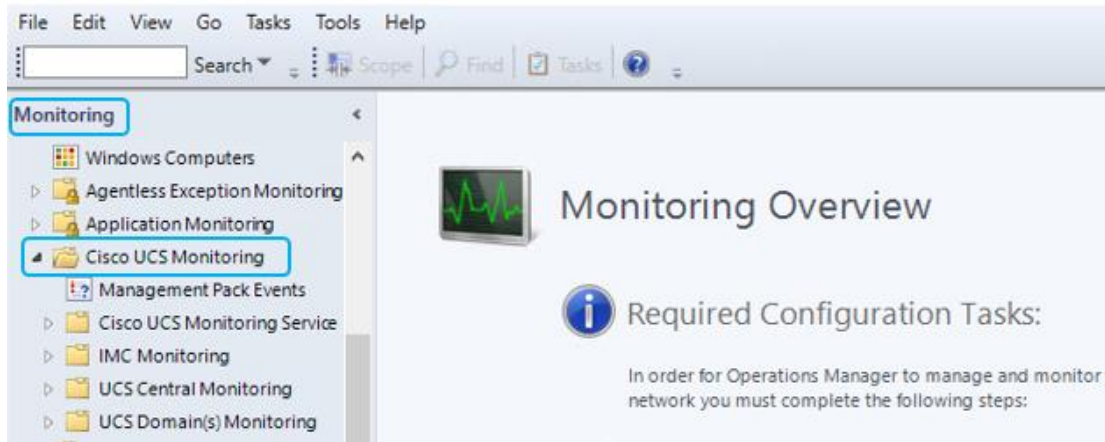
It may take few minutes to import the files.



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



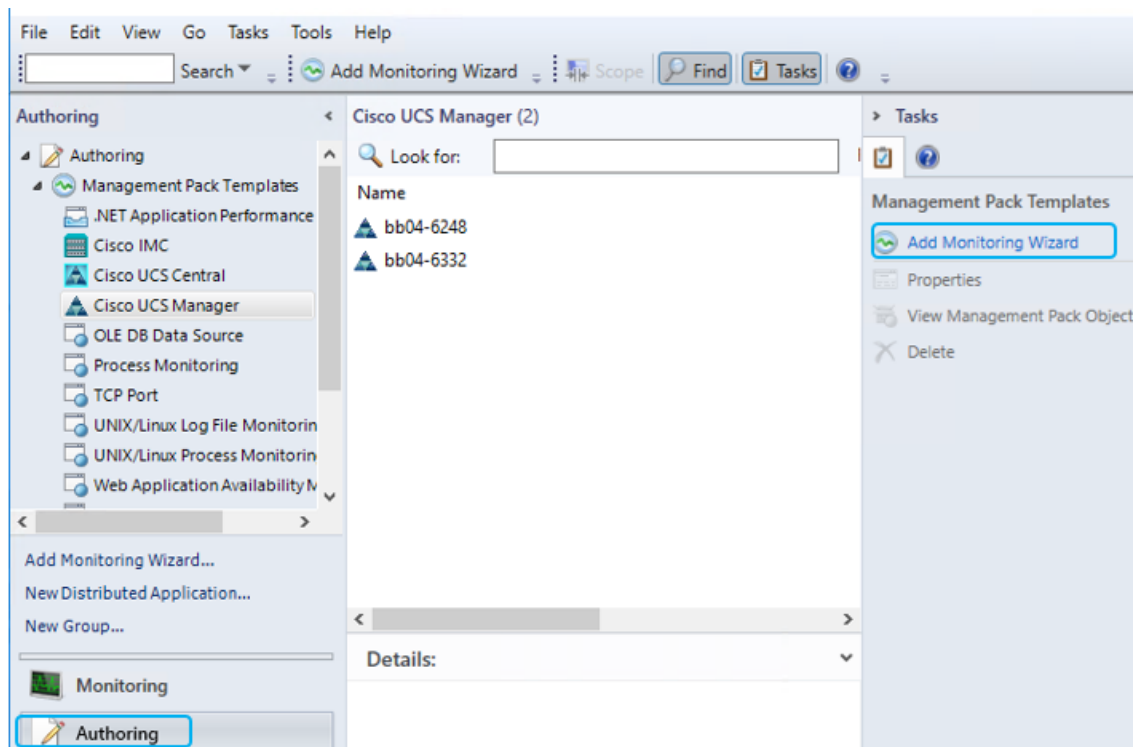
12. In the Monitoring pane, a Cisco UCS folder is also created. When the folder is expanded, it lists the Cisco UCS Monitoring Service, IMC, UCS Central and UCS Domain monitoring folders.



Adding a Cisco UCS Domains to the Operations Manager

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

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



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

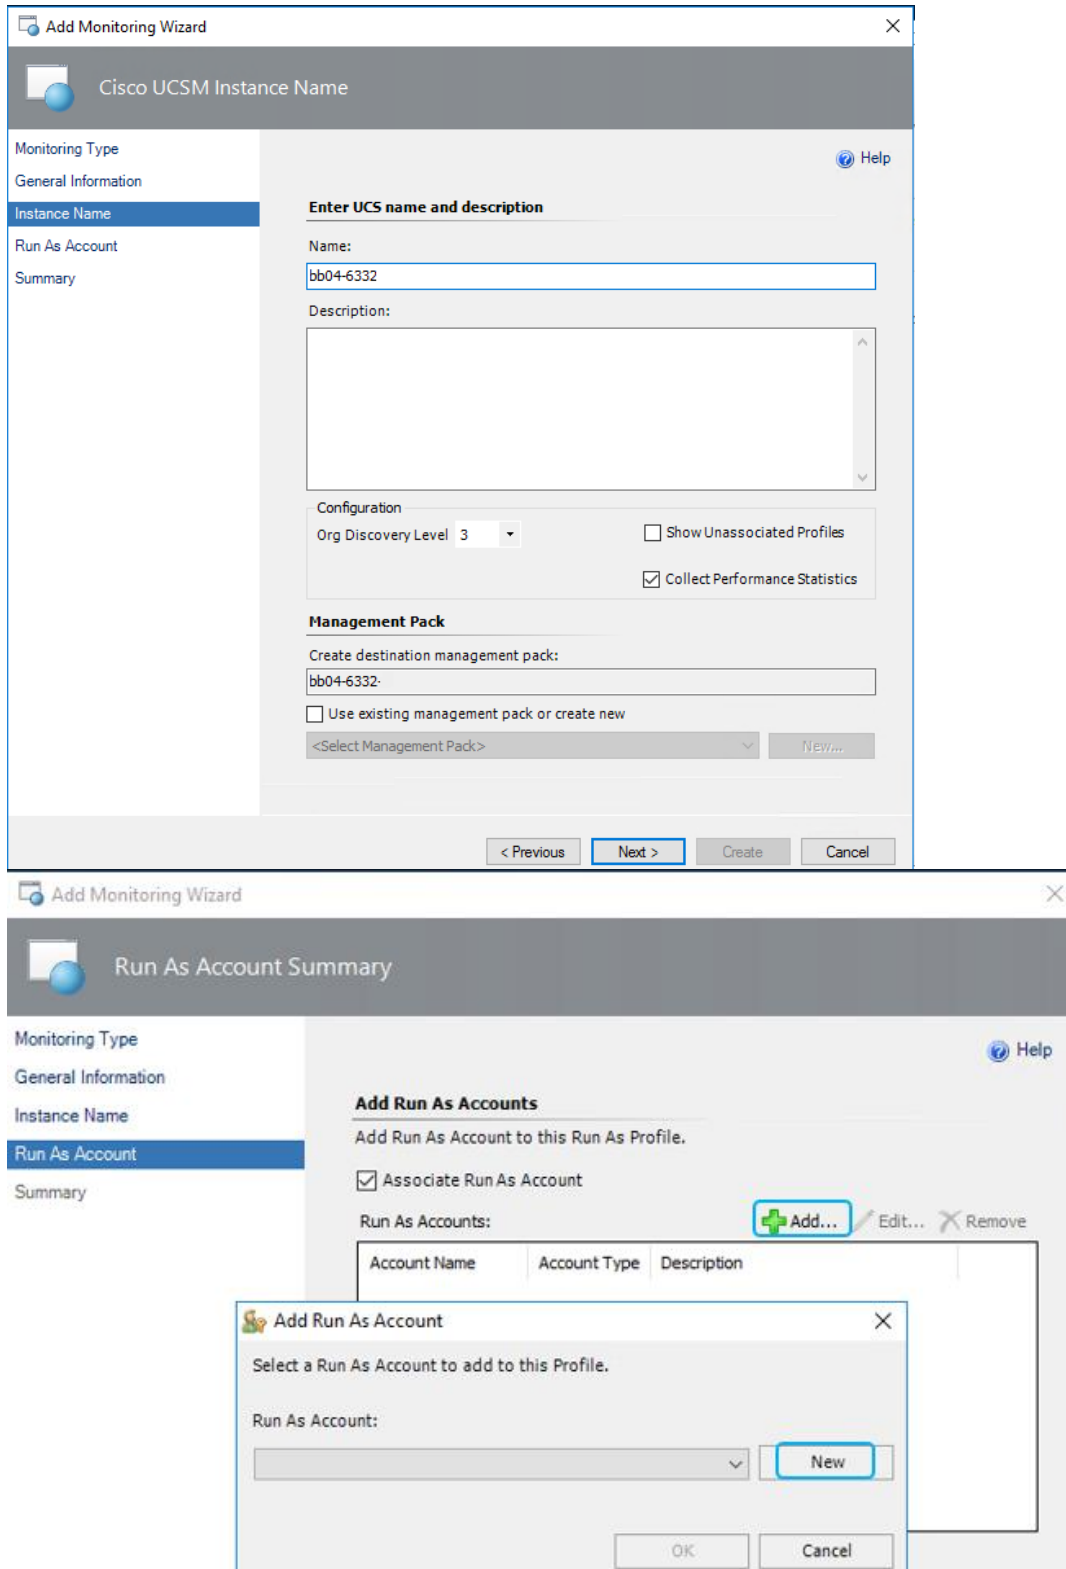
6. On the General Information tab, review and complete the following as shown in the figure below:

The screenshot shows the 'Add Monitoring Wizard' dialog box with the 'General Information' tab selected. The title bar reads 'Add Monitoring Wizard' and 'Specify IP Address, Port and Connection Mode'. The left sidebar lists 'Monitoring Type', 'General Information', 'Instance Name', 'Run As Account', and 'Summary'. The main area is titled 'Cisco UCS Manager' and contains the following fields:

- Connection:** IP Address* / Hostname: 192.168.156.12
- Connection Mode:** Secure
- Port Number:** 443
- Proxy Server:**
 - Enable Proxy Configuration
 - IP Address * / Hostname : []
 - Port: []
 - Enable Proxy Authentication
 - Username: [] Password: []
- Notes:**
 - * IPv4 Address or IPv6 Address
 - * IPv6 address should be enclosed in "[" and "]" brackets
- Cisco UCS Monitoring Service:**
 - Machine Type: Management Server
 - Service Machine: MS-SCOM. []

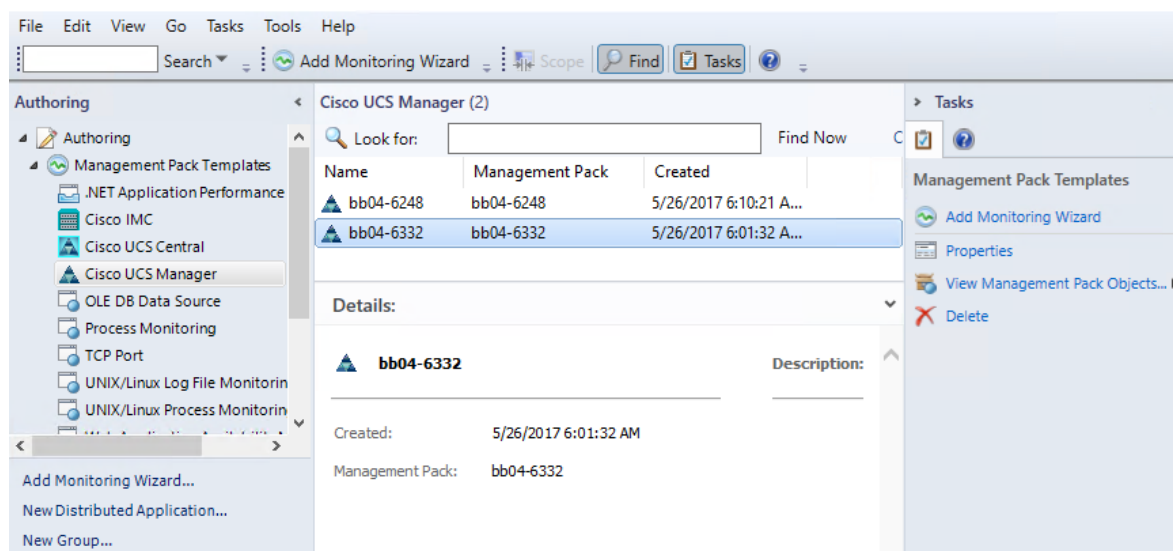
At the bottom right, there is a 'Test Connection' button.

7. To check Operations Manager connectivity to UCS Manager, click Test Connection.
8. In the Authentication dialog box, enter the username and password, and click OK and click Next.
9. On the Instance Name tab, complete the following as shown in the figure below and click Next.
10. On the Run As Account tab, click Add.
11. If you want to associate a new run-as account to the UCS domain instance, click New.



12. Click Next.

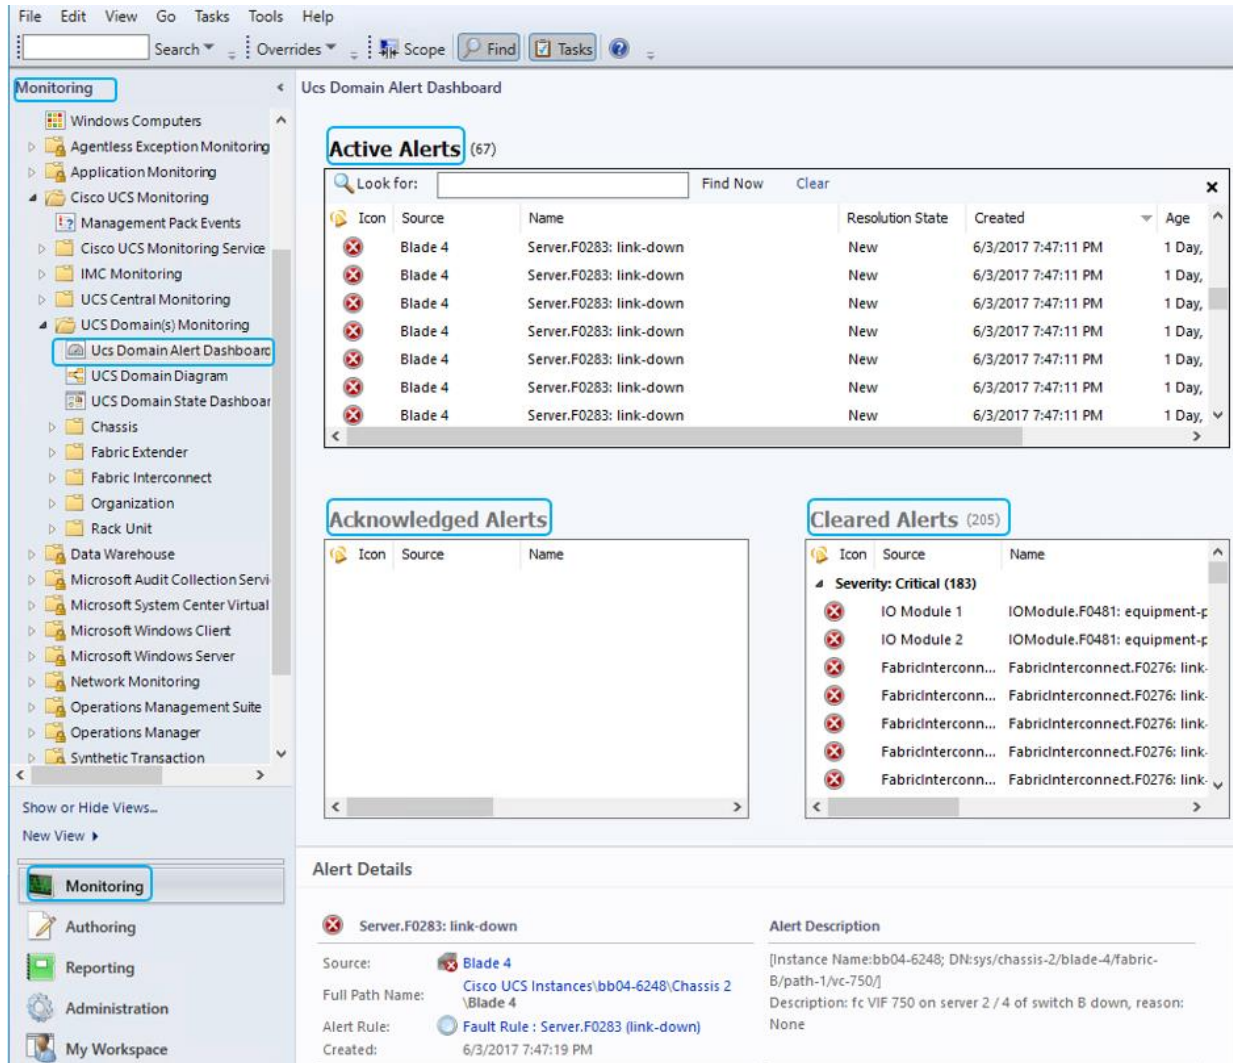
13. On the Summary tab, review the configuration summary, and click Create. The template for monitoring the UCS domain is created.



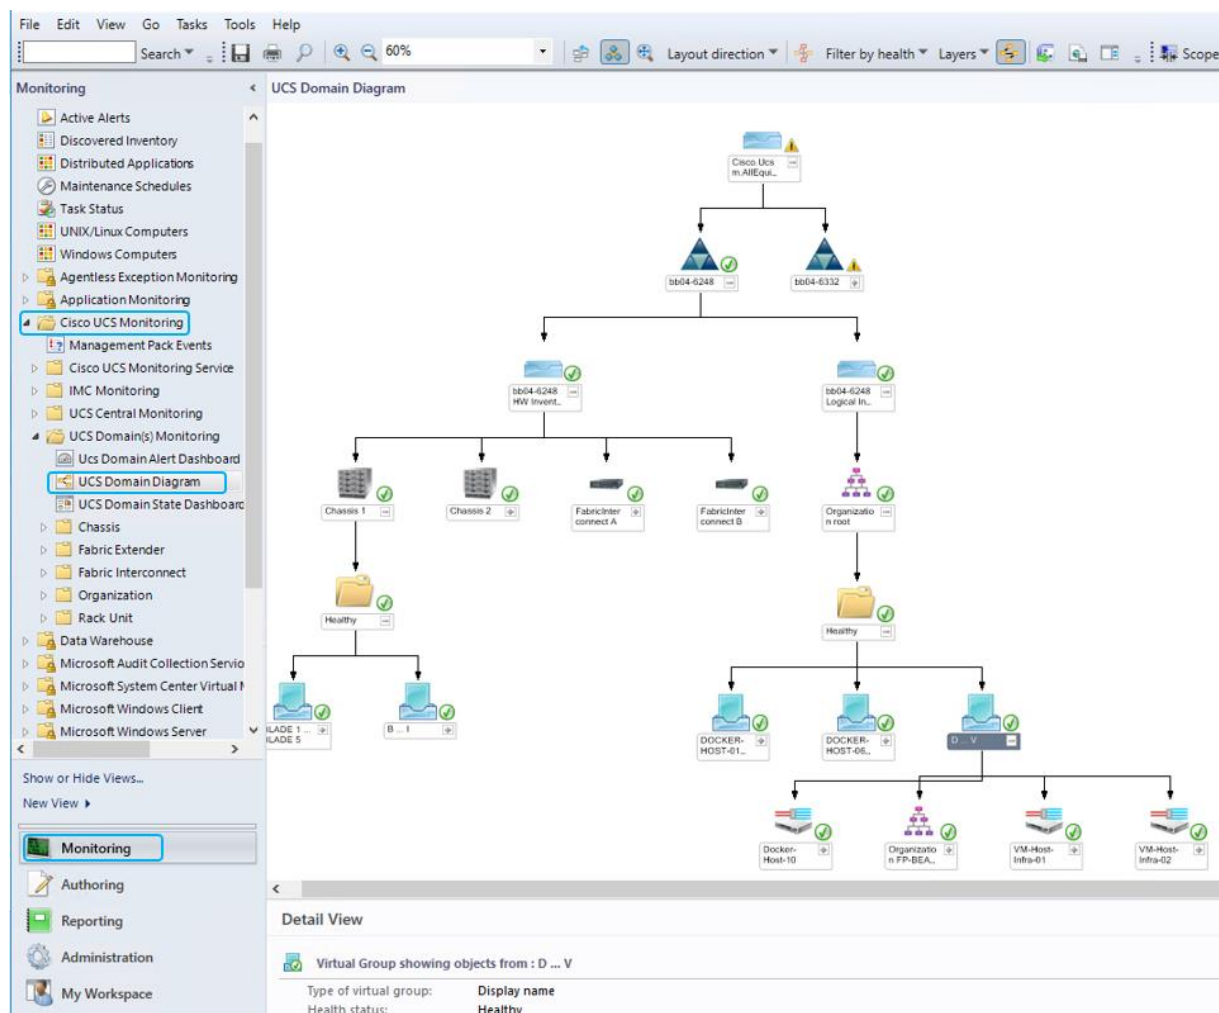
Cisco UCS Manager Monitoring Dashboards

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

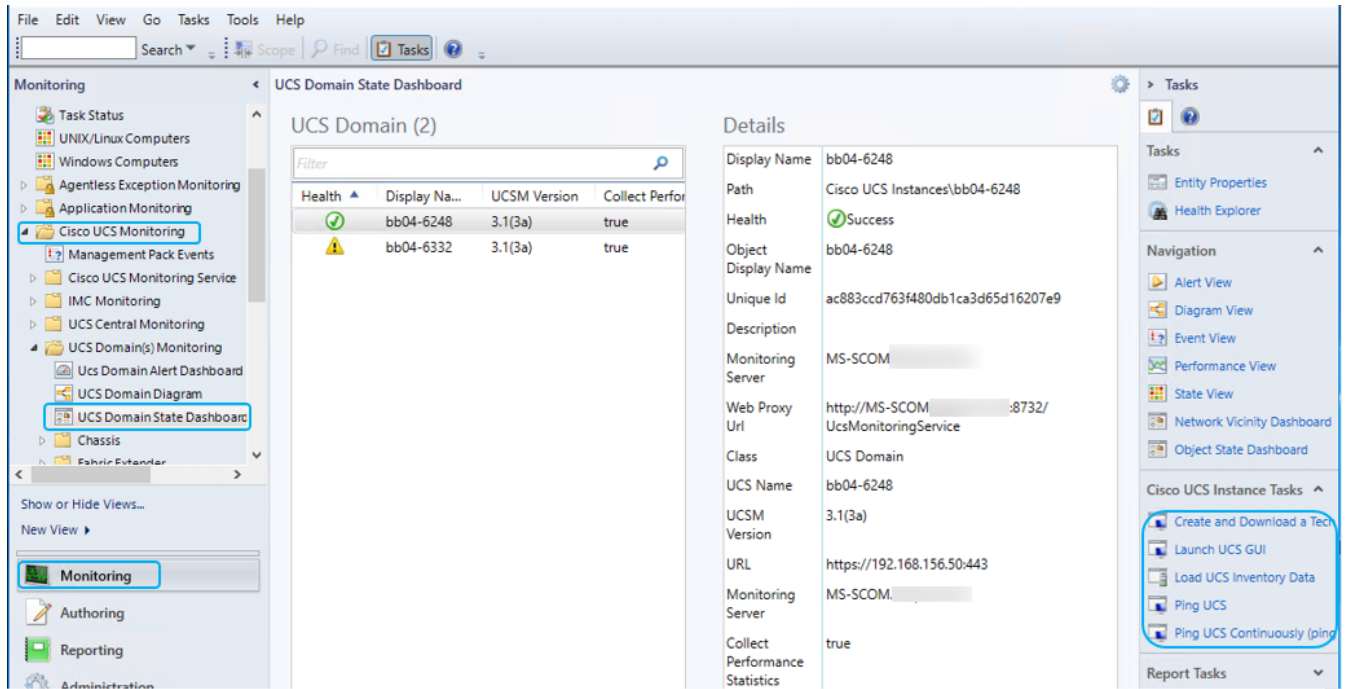
- UCS Domain Alert Dashboard—Displays all alerts generated in the UCS domain. The alerts are further categorized into the following views:
 - Active Alerts
 - Acknowledge Alerts
 - Cleared Alerts



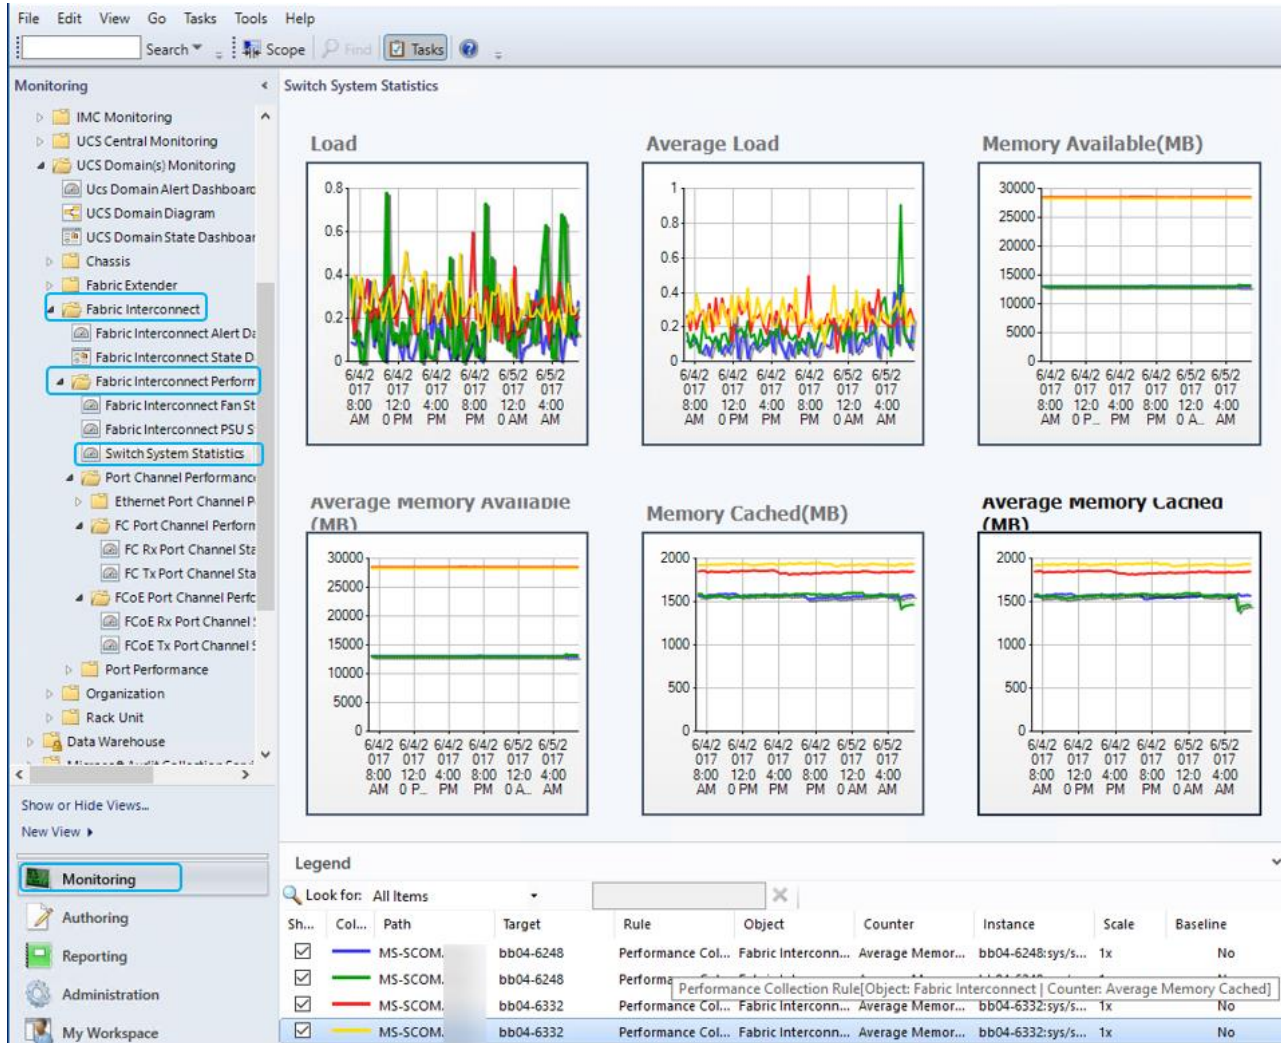
- UCS Domain Diagram—Displays a graphical view of the relationship between different Cisco UCS Domain(s) components for all Instances.



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



You can view performance metrics for the various Cisco UCS components as shown the screenshot below:



Cisco UCS Manager Plug-in for SCVMM

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

Cisco UCS Manager Plug-in Installation

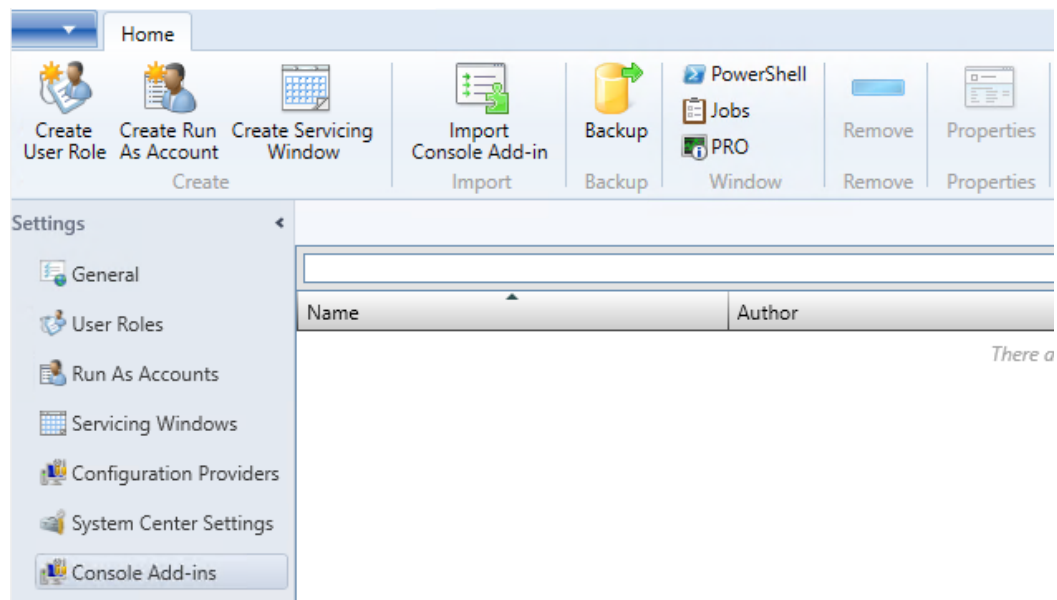
To install the Cisco UCS virtual machine manager add-in, complete the following steps:

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

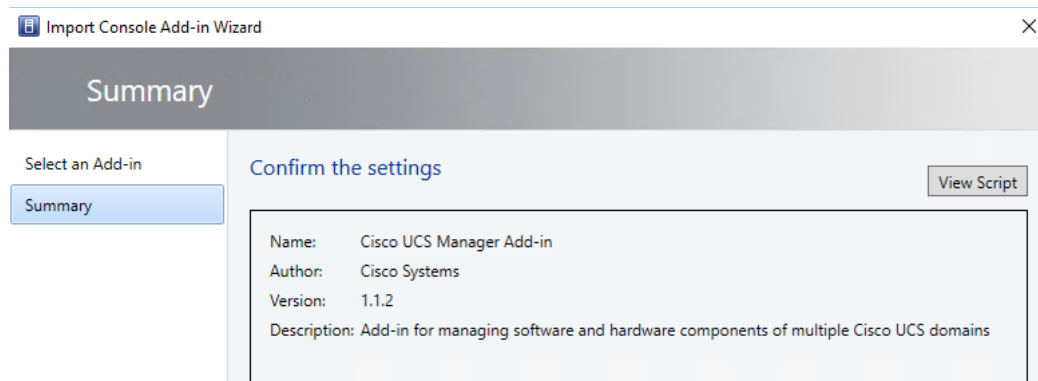


The add-in is made available as a zipped file that has to be imported into the virtual machine manager to install it.

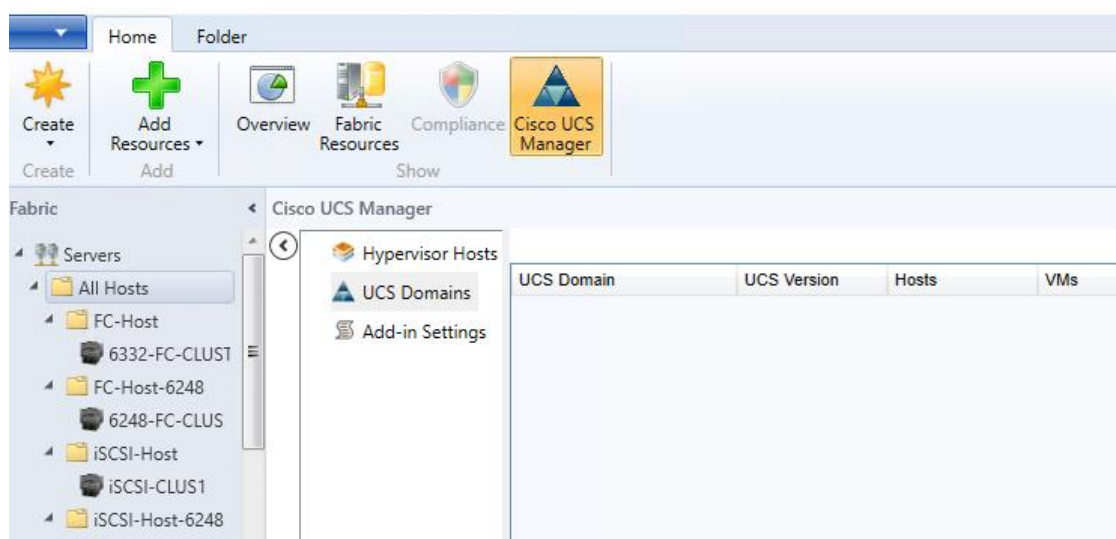
4. Open an instance of the Virtual Machine Manager console.
5. In the Navigation pane, click Settings.
6. In the toolbar, click Import Console Add-in. The Import Console Add-in wizard appears.



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



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



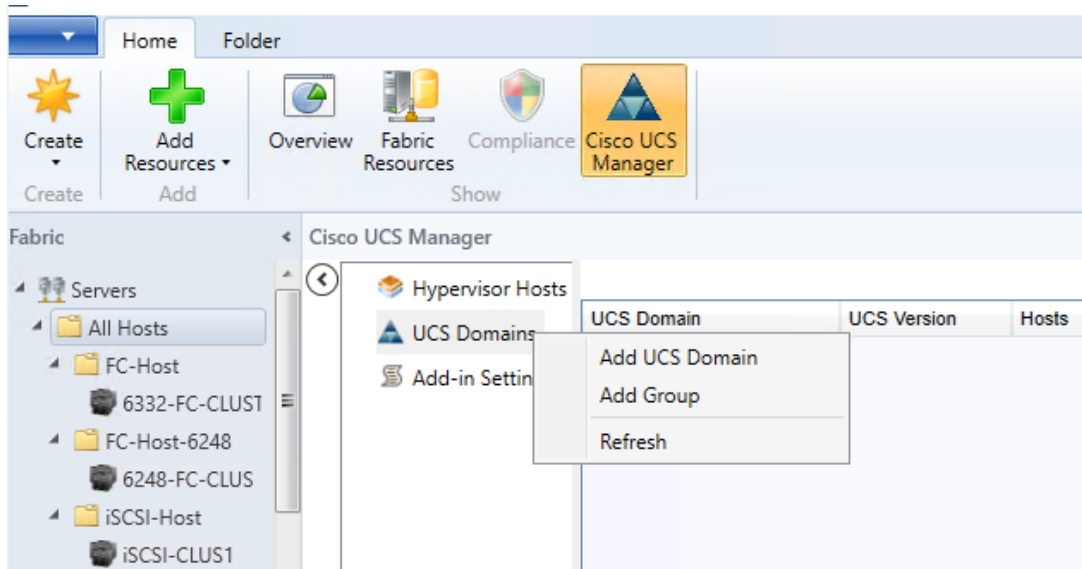
Cisco UCS Domain Registration

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.

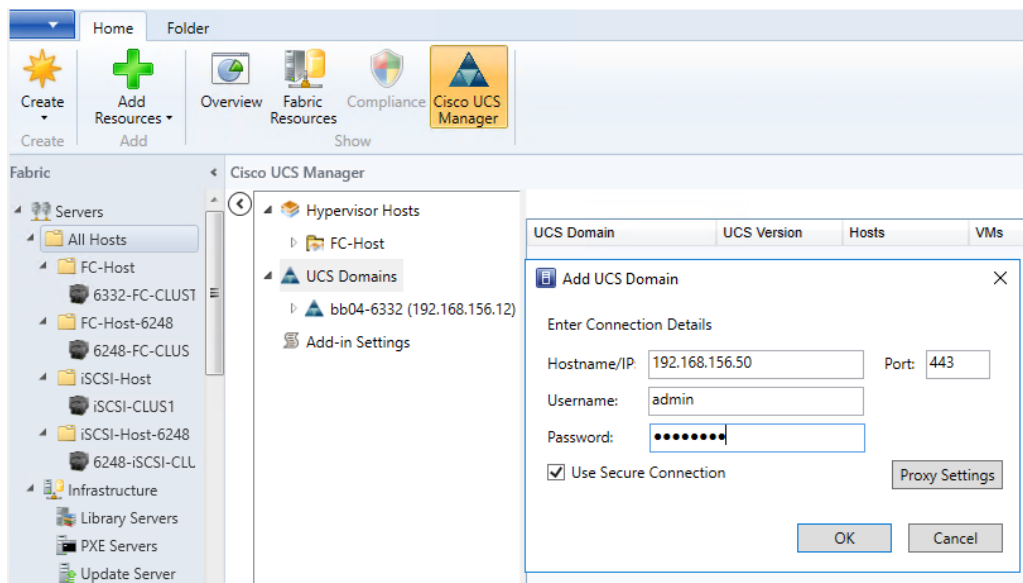
To register a UCS domain, complete the following steps:

1. On the toolbar, click Cisco UCS Manager.
2. Right-click UCS Domains.
3. Click Add UCS Domain.

The Add UCS Domain dialog box appears.



4. Enter the following details in the dialog box:



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

5. Click Proxy Settings. Proxy Settings dialog box appears.

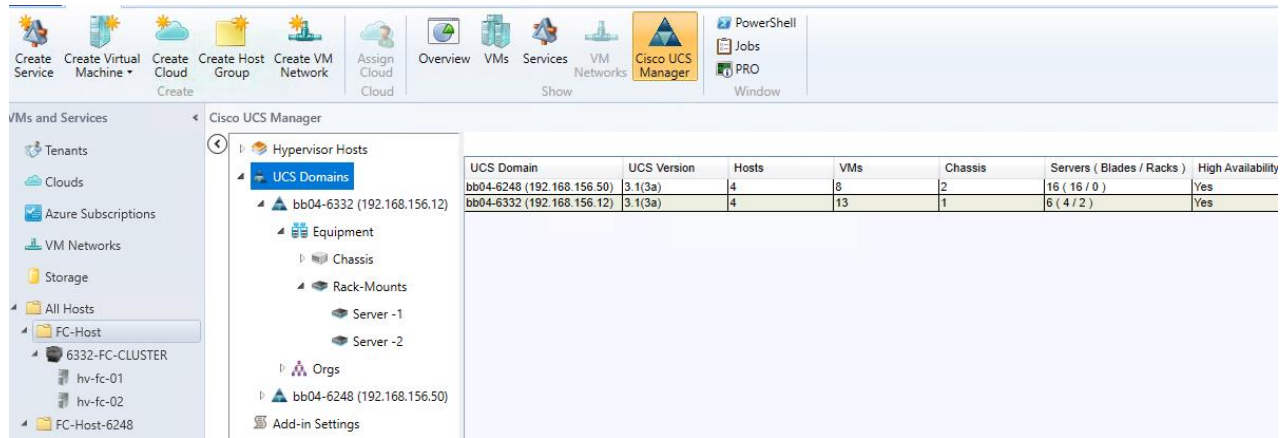
6. In the Proxy Settings dialog box, click Use Custom Proxy Settings radio button and enter the following details:



If required, you can edit the proxy settings at a later time.

7. 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 also can add UCS domains within groups. If you want to add a UCS domain within a group, right-click on the group and follow steps 3 through step 7 in the preceding procedure.

Using the Cisco UCS SCVMM Plugin

Viewing the Server Details from the Hypervisor Host View

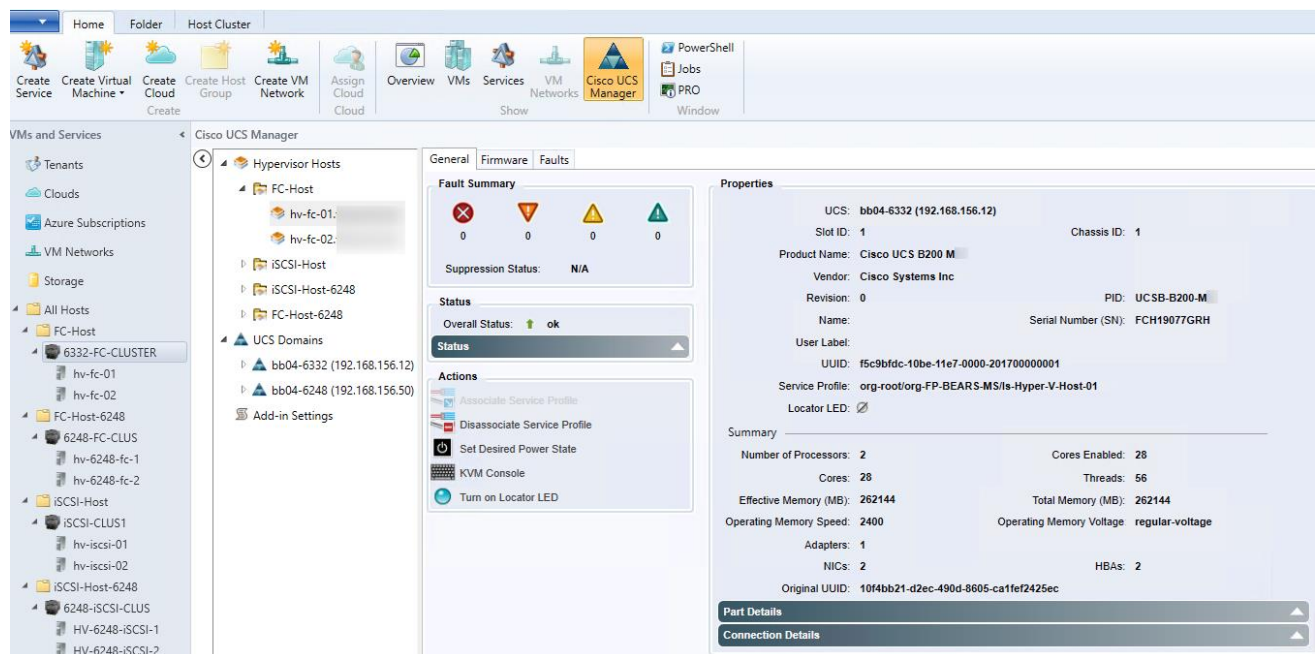
To view server details, complete the following steps:

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	

Name	Description
Associate Service Profile	Enables you to associate a service profile to the server.
Disassociate Service Profile	Enables you to disassociate a service profile from the server.
Set Desired Power State	Provides options to set the power state of a service profile.
KVM Console	Enables you to launch the KVM console.
Turn on Locator LED	Enables you to either turn on or turn off the locator LED depending on the current state.
Firmware tab	Provides the firmware details such as BIOS, CIMC, adaptors and storage device part IDs, and the firmware versions. If there are any changes to the firmware details on the server, those changes will reflect here.
Faults tab	Displays the faults' details specific to the server, such as properties, severity, fault codes and IDs, description, affected objects, and so on. Provides options to filter the faults based on severity, and option under the Actions area to acknowledge the fault on UCS.

3. 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.

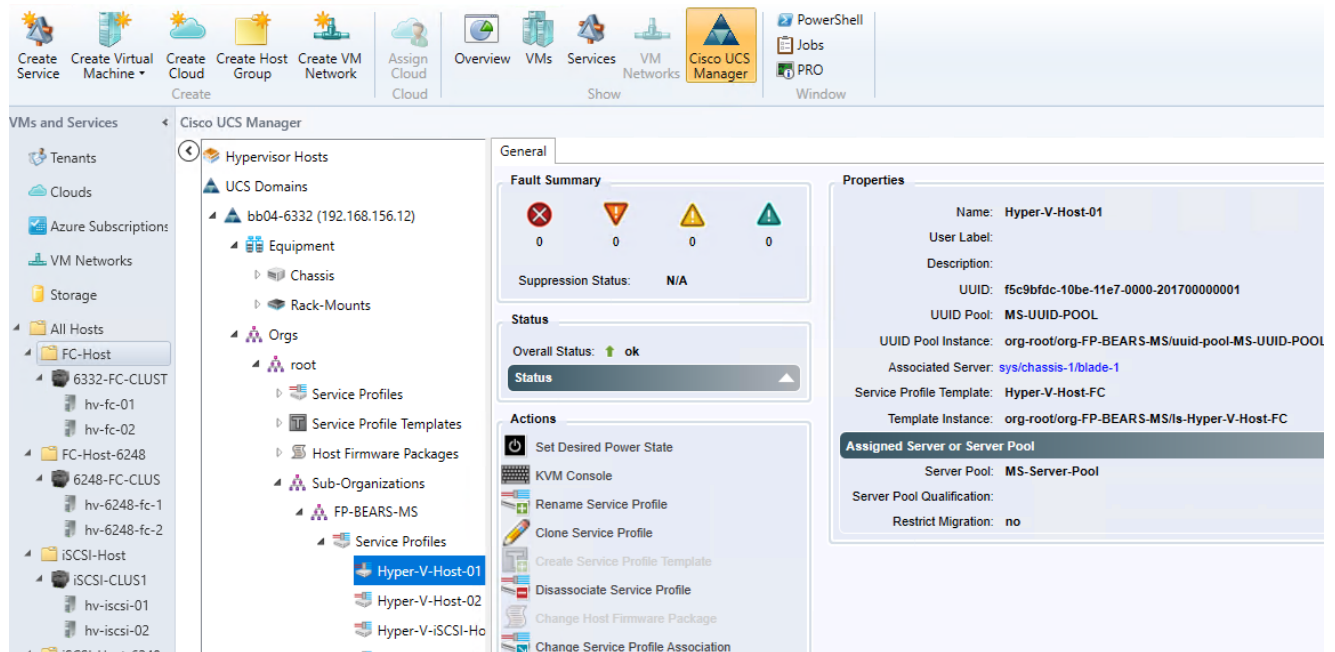
Viewing the Service Profile Details

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.
Rename Service Profile	Enables you to rename a service profile.
Create a Clone	Enables you to create a clone of the service profile by inheriting the attributes of the service profile.
Disassociate Service Profile	Enables you to disassociate the service profile from the server.
Change Host Firmware Package	Enables you to change the host firmware association.
Change Service Profile Association	Enables you to upgrade the host firmware on the servers.

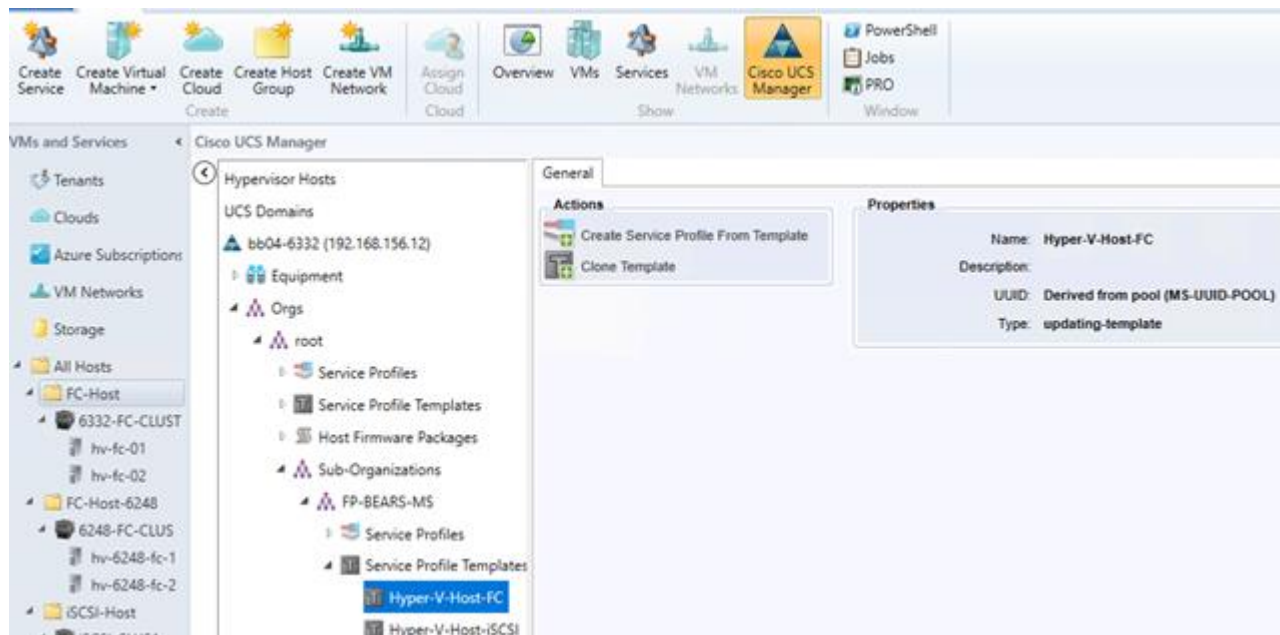


Viewing the Service Profile Template Details

Using the add-in, you can view the service profile template details, such as properties, and faults information. 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 Versions	Enables you to modify Blade package version and Rack package version properties.

Building the Virtual Machines and Environment for Workload Testing

Software Infrastructure Configuration

This section details how to configure the software infrastructure components that comprise this solution.

Install and configure the infrastructure virtual machines by following the process provided in the table below:

Configuration	Citrix XenDesktop Controllers Virtual Machines	Citrix Machine Creation Services Virtual Machines
Operating system	Microsoft Windows Server 2016	Microsoft Windows 10
Virtual CPU amount	4	2
Memory amount	22 GB	2 GB
Network	HV-Infra	HV-VDI
Disk-1 (OS) size and location	40 GB Infra-DS volume	40 GB Infra-DS volume
Configuration	Microsoft Active Directory DCs Virtual Machines	Microsoft System Center Virtual Machine Manager Virtual Machine
Operating system	Microsoft Windows Server 2016	Microsoft Windows Server 2016
Virtual CPU amount	4	8
Memory amount	4 GB	24 GB
Network	HV-Infra	HV-Infra
Disk size and location	40 GB Infra-DS volume	250 GB Infra-DS volume
Configuration	Microsoft SQL Server	

	Virtual Machine	
Operating system	Microsoft Windows Server 2016 Microsoft SQL Server 2016 SP1	
Virtual CPU amount	4	
Memory amount	4 GB	
Network	VMXNET3 HV-Infra	
Disk-1 (OS) size and location	40 GB Infra-DS volume	
Disk-2 size and location	100 GB Infra-DS volume SQL Logs	
Disk-3 size and location	150 GB Infra-DS volume SQL Databases	

Preparing the Master Targets

This section provides guidance around creating the golden (or master) images for the environment. VMs for the master targets must first be installed with the software components needed to build the golden images. For this CVD, the images contain the basics needed to run the Login VSI workload.

To prepare the master VMs for the Hosted Virtual Desktops (HVDs) installing the OS, installing the Virtual Delivery Agents (VDAs), and installing application software.

The master target HVD(VDI) was configured as listed in the table below:

Configuration	VDI Virtual Machines	
Operating system	Microsoft Windows 10 64-bit	
Virtual CPU amount	2	
Memory amount	2.0 GB (reserved)	

Configuration	VDI Virtual Machines	
Network	HV-VDI	
Additional software used for testing	Microsoft Office 2016 Login VSI 4.1.25 (Knowledge Worker Workload)	

Installing and Configuring XenDesktop and XenApp

This section details the installation of the core components of the XenDesktop/XenApp 7.15 system. This CVD installs two XenDesktop Delivery Controllers to support persistent virtual desktops (VDI).

Prerequisites

Citrix requires the install of the Microsoft System Center Virtual Machine Manager software to be installed on all XenDesktop Delivery Controllers

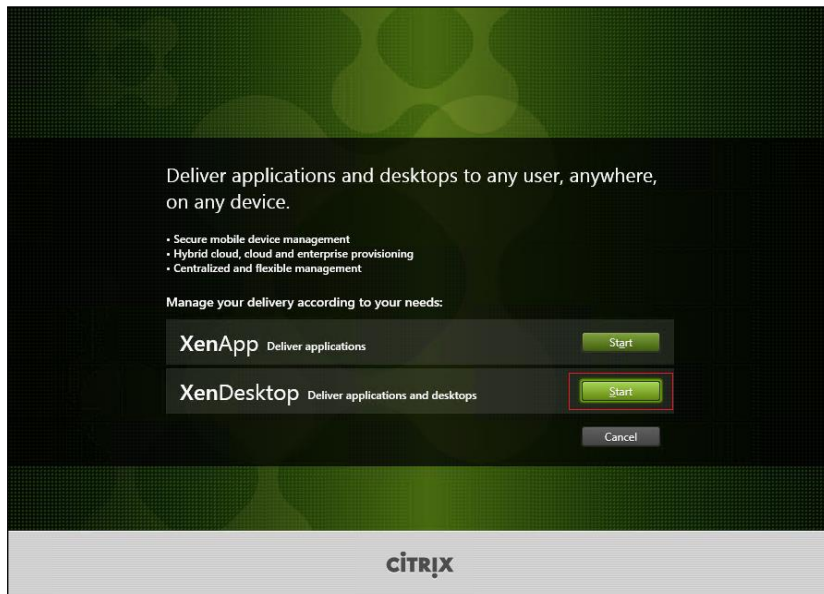
To install the SCVMM software, complete the following steps:

1. Insert the SCVMM 2016 ISO into the XenDesktop Delivery Controller VMs.
2. When prompted, un-check the VMM Server option, leaving only the VMM Console option checked.
3. Accept defaults through install.
4. Click Next and then click Finish.
5. Perform the above steps on all Delivery Controllers.

Install XenDesktop Delivery Controller, Citrix Licensing, and StoreFront

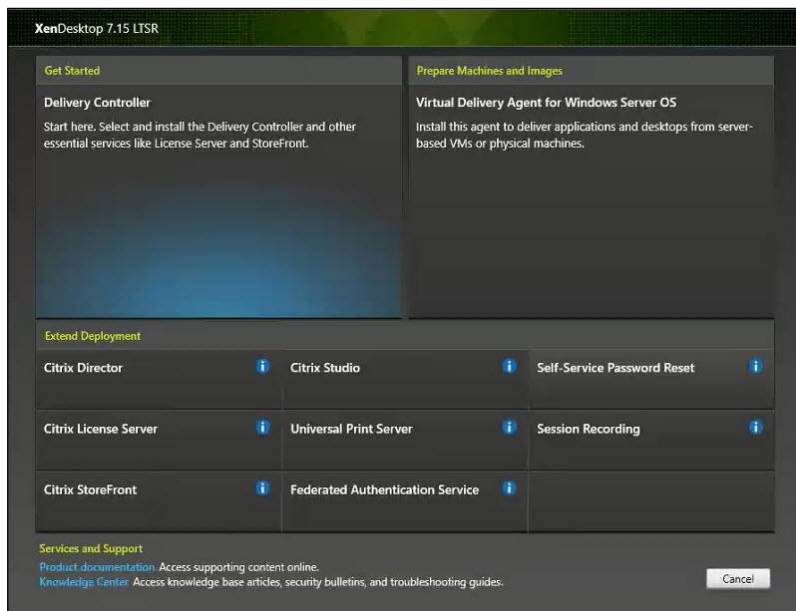
The process of installing the XenDesktop Delivery Controller also installs other key XenDesktop software components, including Studio, which is used to create and manage infrastructure components, and Director, which is used to monitor performance and troubleshoot problems.

1. To begin the installation, connect to the first XenDesktop server and launch the installer from the Citrix XenDesktop 7.15 ISO.
2. Click Start.



The installation wizard presents a menu with three subsections.

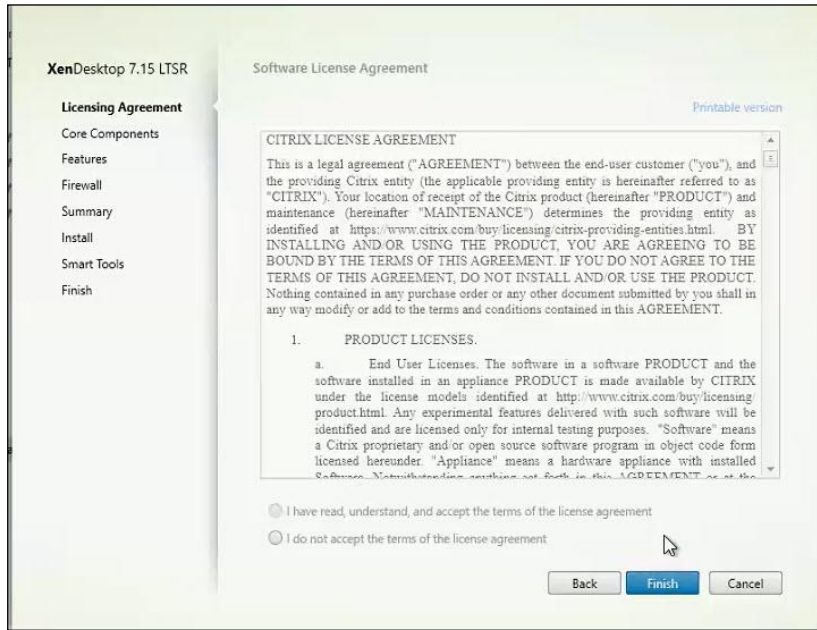
3. Click “Get Started - Delivery Controller.”



4. Read the Citrix License Agreement.

5. If acceptable, indicate your acceptance of the license by selecting the “I have read, understand, and accept the terms of the license agreement” radio button.

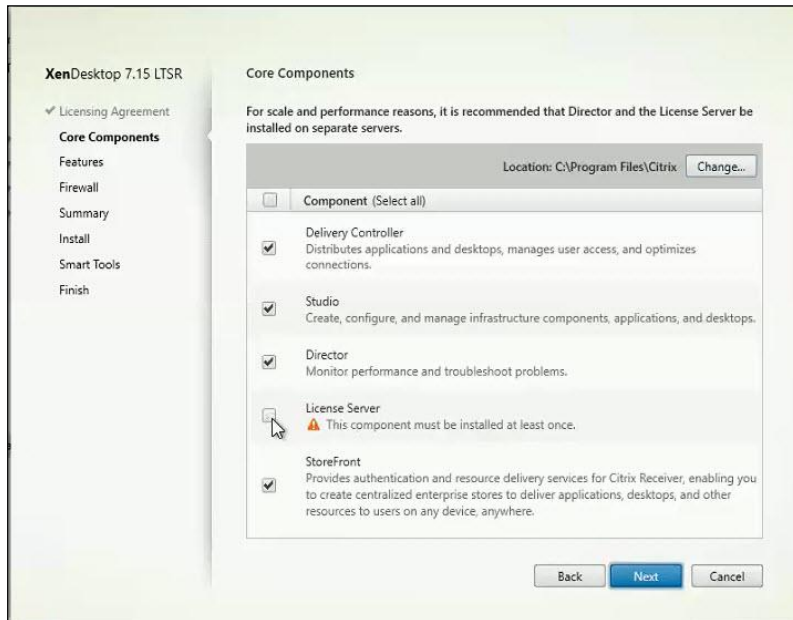
6. Click Next.



7. Select the components to be installed on the first Delivery Controller Server:

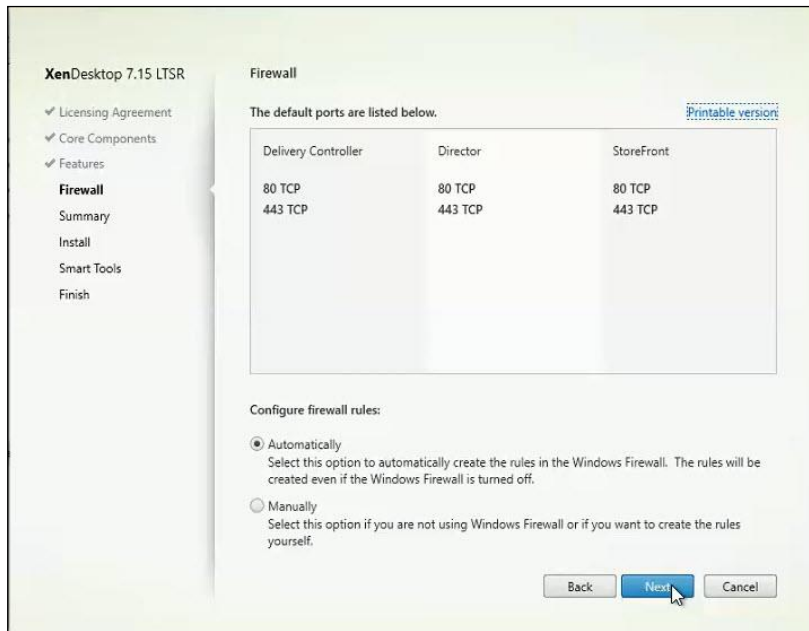
- a. Delivery Controller
- b. Studio
- c. License Server
- d. StoreFront

8. Click Next.

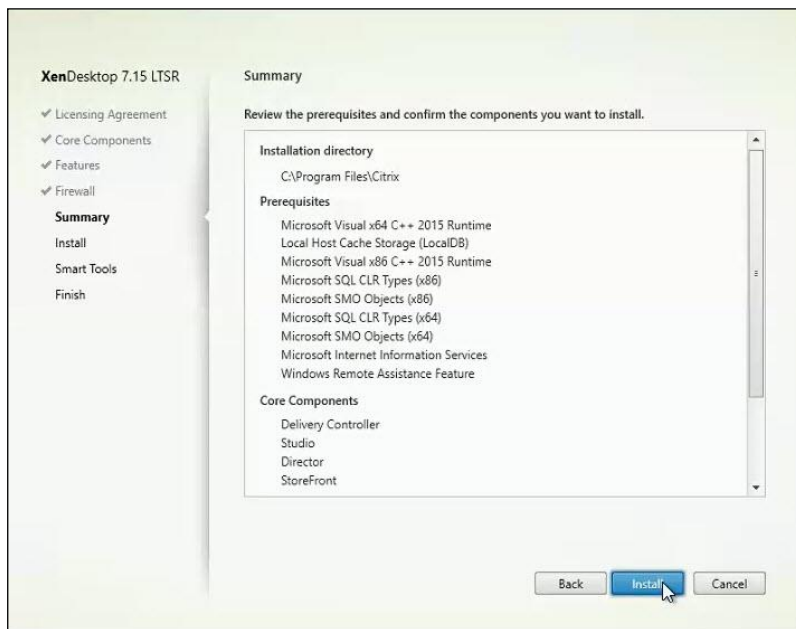


Dedicated StoreFront servers should be implemented for large scale deployments.

9. Since a SQL Server will be used to Store the Database, leave “Install Microsoft SQL Server 2012 SP1 Express” unchecked.
10. Click Next.
11. Select the default ports and automatically configured firewall rules.
12. Click Next.

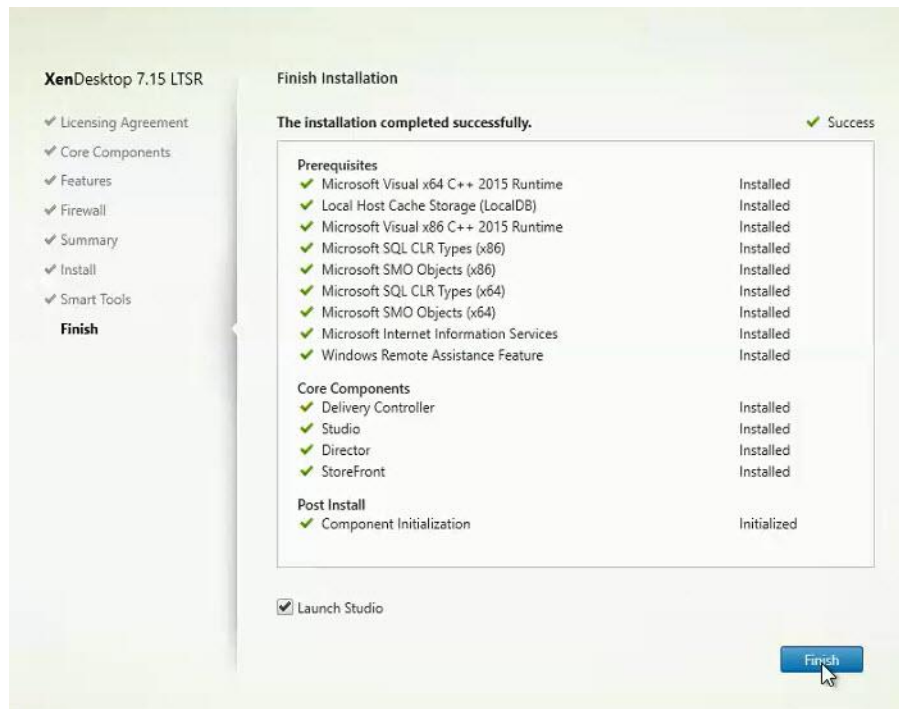
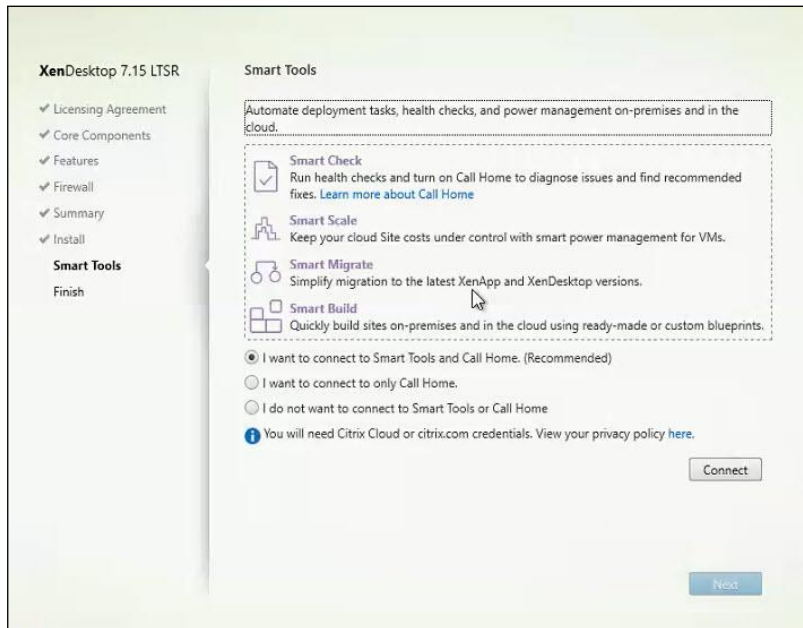


13. Click Install to begin the installation.



14. (Optional) Click the Call Home participation.

15. Click Finish.

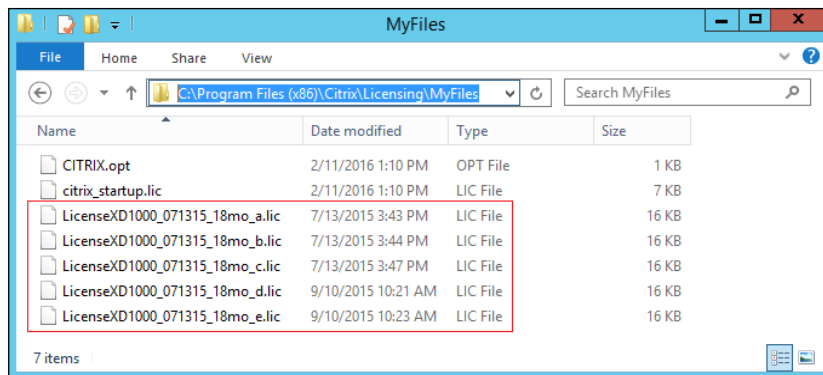


16. (Optional) Check Launch Studio to launch Citrix Studio Console.

Installing Citrix Licenses

To install the Citrix Licenses, complete the following steps:

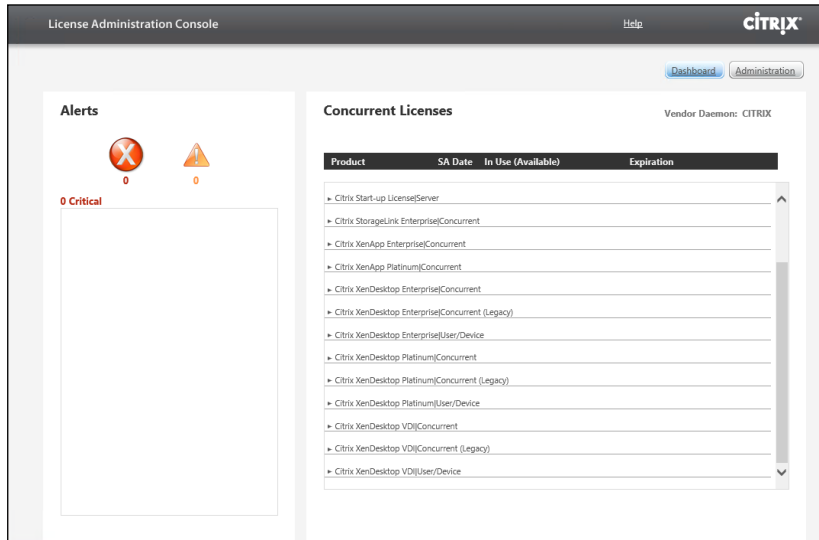
1. Copy the license files to the default location (C:\Program Files (x86)\Citrix\Licensing\ MyFiles) on the li-
cense server.



2. Restart the server or Citrix licensing services so that the licenses are activated.
3. Run the application Citrix License Administration Console.



4. Confirm that the license files have been read and enabled correctly.



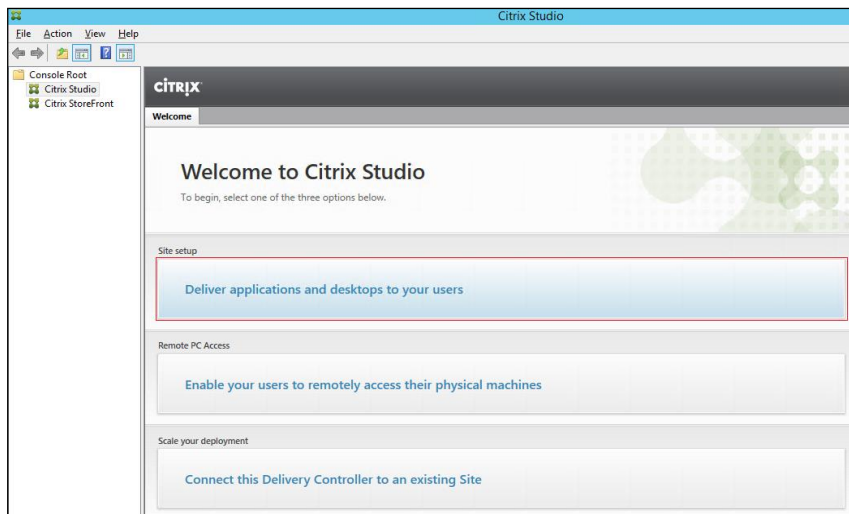
Configure the XenDesktop Site

Citrix Studio is a management console that allows you to create and manage infrastructure and resources to deliver desktops and applications. Replacing Desktop Studio from earlier releases, it provides wizards to set up your environment, create workloads to host applications and desktops, and assign applications and desktops to users.

Citrix Studio launches automatically after the XenDesktop Delivery Controller installation, or if necessary, it can be launched manually. Studio is used to create a Site, which is the core XenDesktop 7.15 environment consisting of the Delivery Controller and the Database.

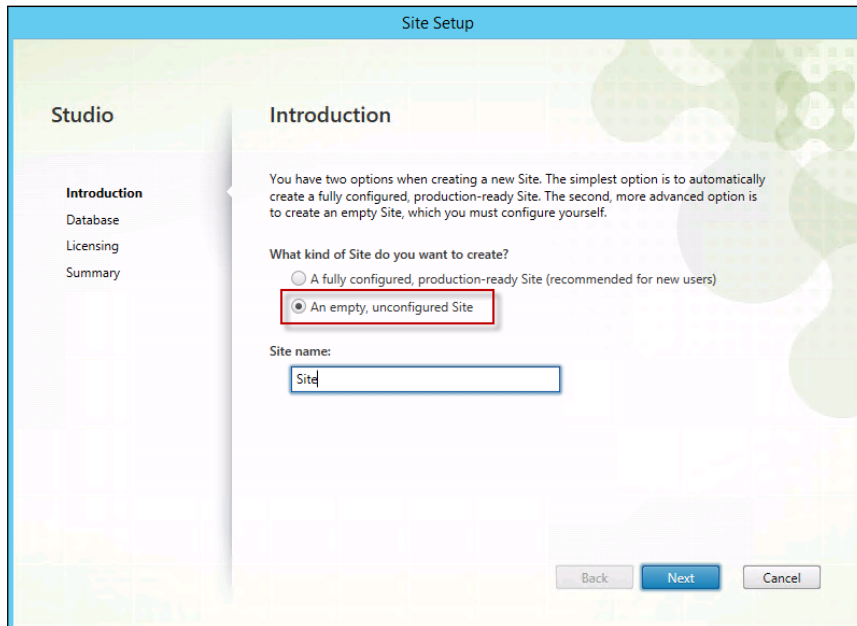
To configure XenDesktop, complete the following steps:

1. From Citrix Studio, click Deliver applications and desktops to your users.

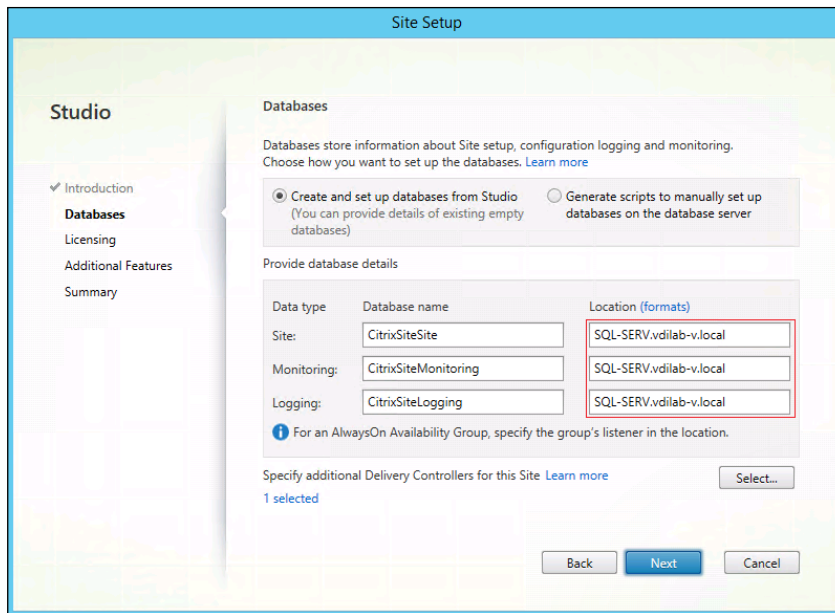


2. Select the “An empty, unconfigured Site” radio button.
3. Enter a site name.

4. Click Next




5. Provide the Database Server Locations for each data type and click Next.



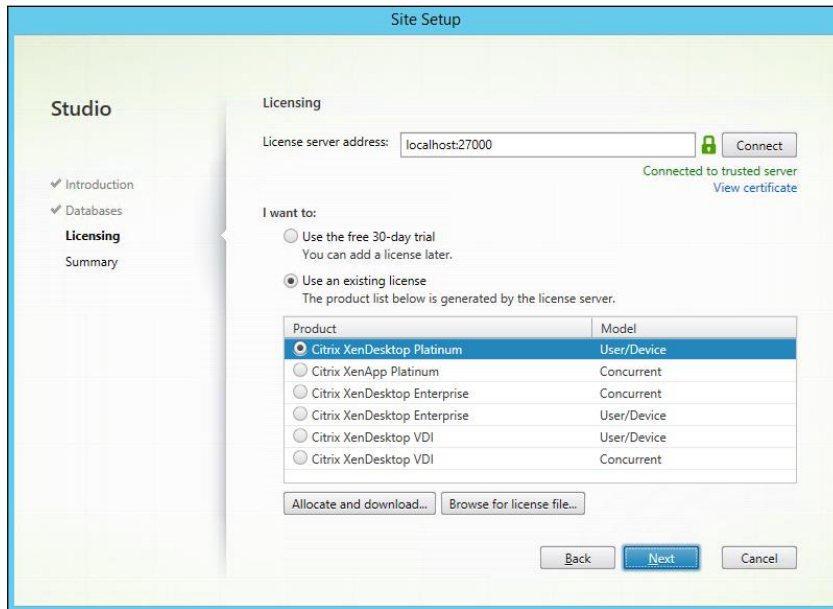
6. Provide the FQDN of the license server.

7. Click Connect to validate and retrieve any licenses from the server.

 If no licenses are available, you can use the 30-day free trial or activate a license file.

8. Select the appropriate product edition using the license radio button.

9. Click Next.

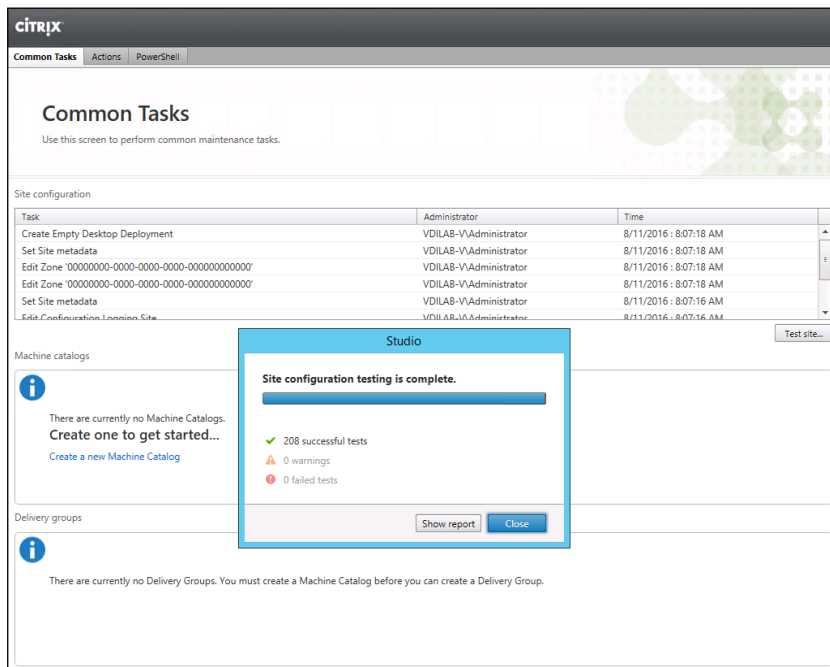


10. Click Finish to complete initial setup.



High availability will be available for the databases when added to the SQL AlwaysOn Availability Group.

11. Click Test site to determine the site creation success.

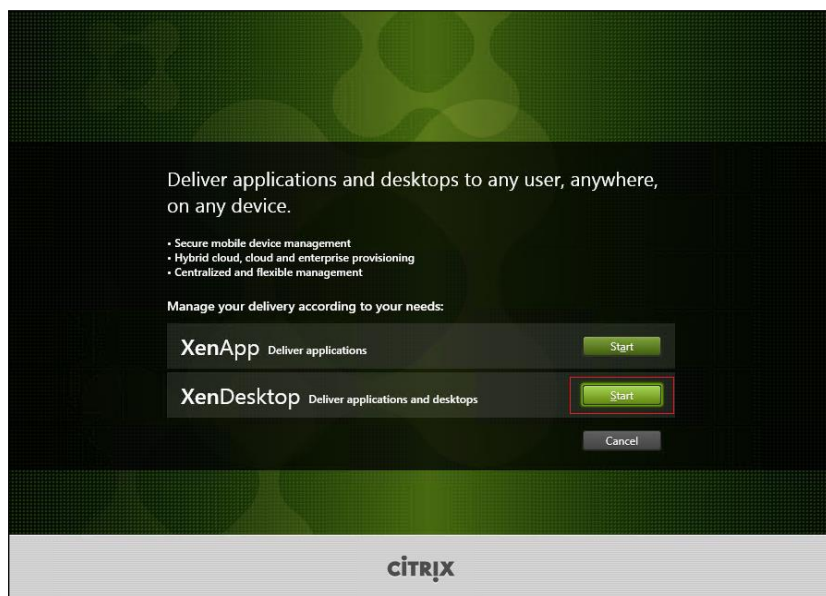


Additional XenDesktop Controller Configuration

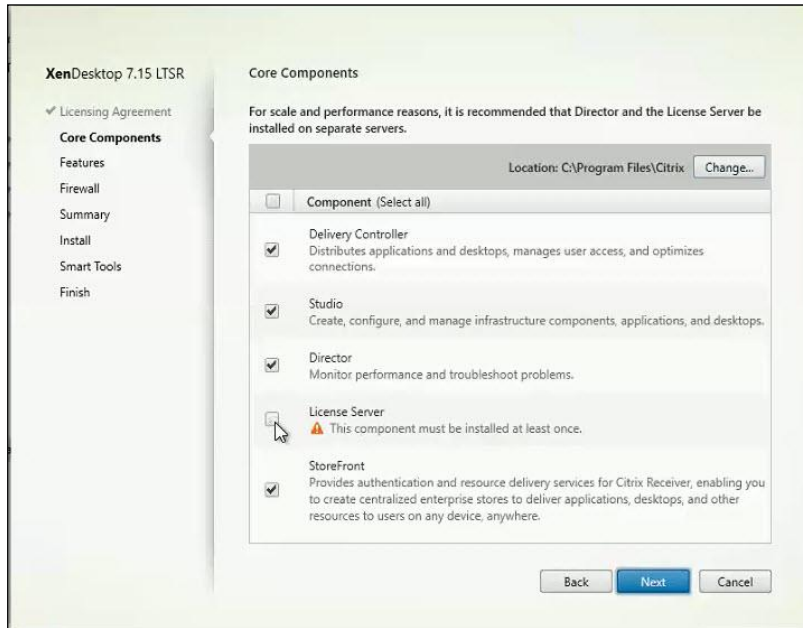
After the first controller is completely configured and the Site is operational, you can add additional controllers. In this CVD, we created two Delivery Controllers.

To configure additional XenDesktop controllers, complete the following steps:

1. To begin the installation of the second Delivery Controller, connect to the second XenDesktop server and launch the installer from the Citrix XenDesktop 7.15 ISO.
2. Click Start.
3. Click Delivery Controller.



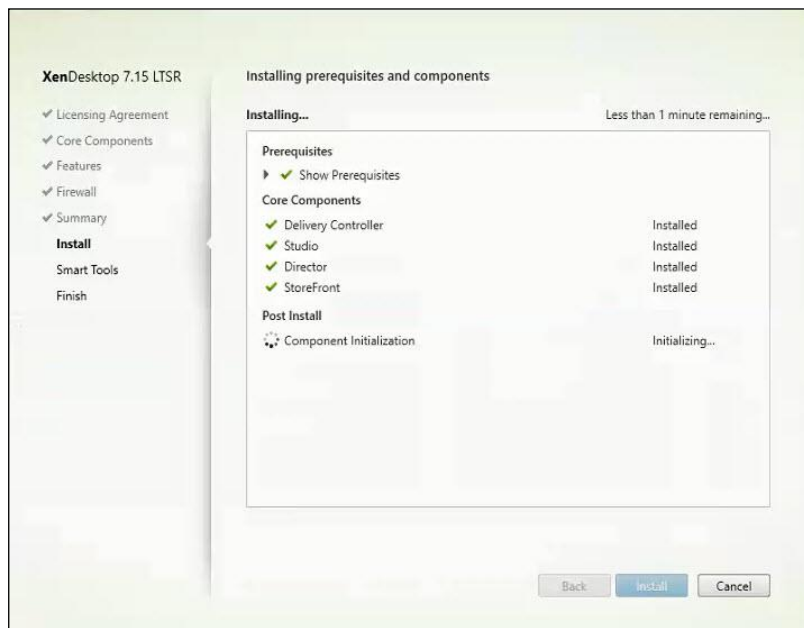
4. Select the components to be installed:
5. Delivery Controller:
 - a. Studio
 - b. Director
 - c. StoreFront (This solution uses two dedicated StoreFront servers)
6. Click Next.



7. Repeat the same steps used to install the first Delivery Controller.
8. Review the Summary configuration.
9. Click Install.



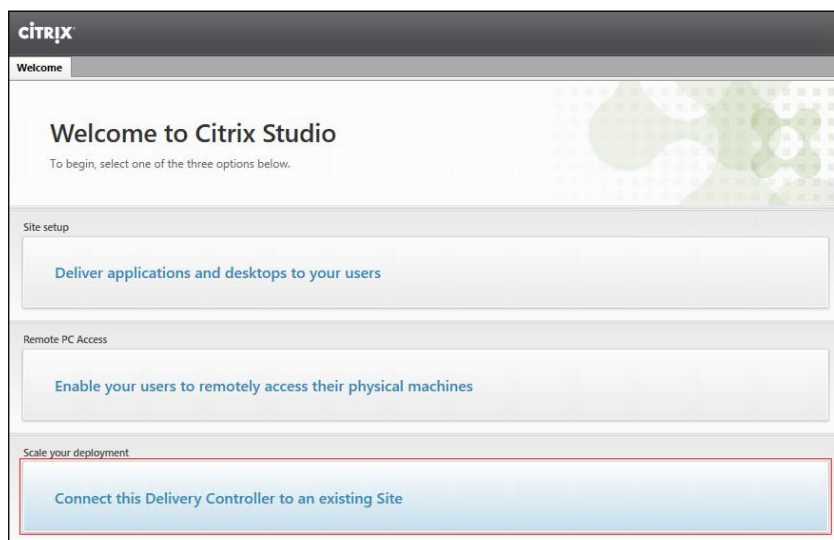
10. Confirm all selected components were successfully installed.
11. Verify the Launch Studio checkbox is checked.
12. Click Finish.



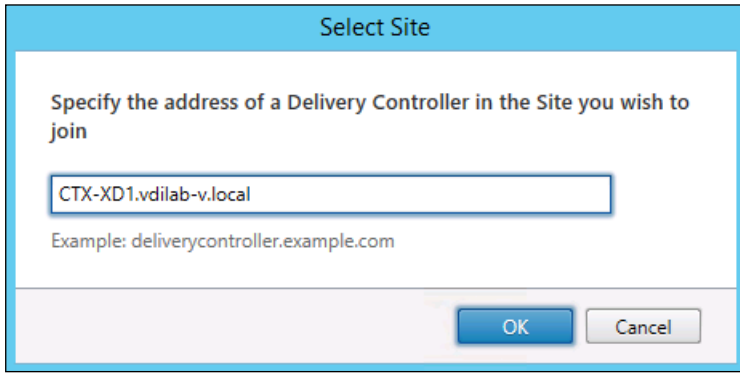
Add the Second Delivery Controller to the XenDesktop Site

To add the second Delivery Controller to the XenDesktop Site, complete the following steps:

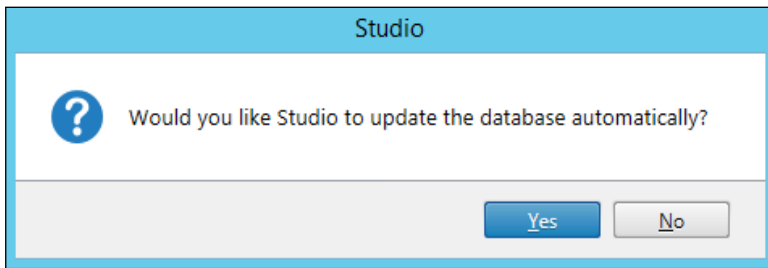
1. Click Connect this Delivery Controller to an existing Site.



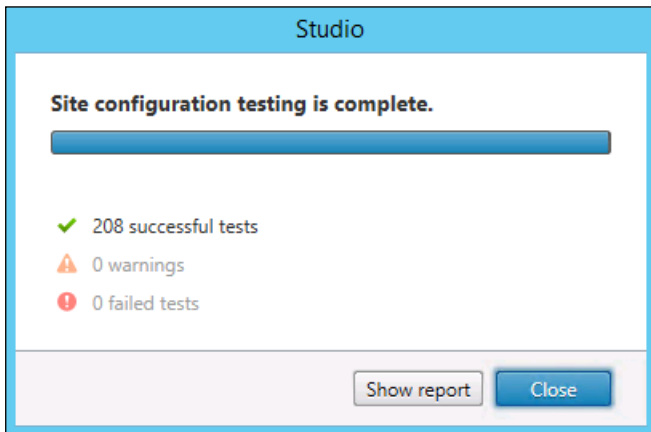
2. Enter the FQDN of the first delivery controller.
3. Click OK.

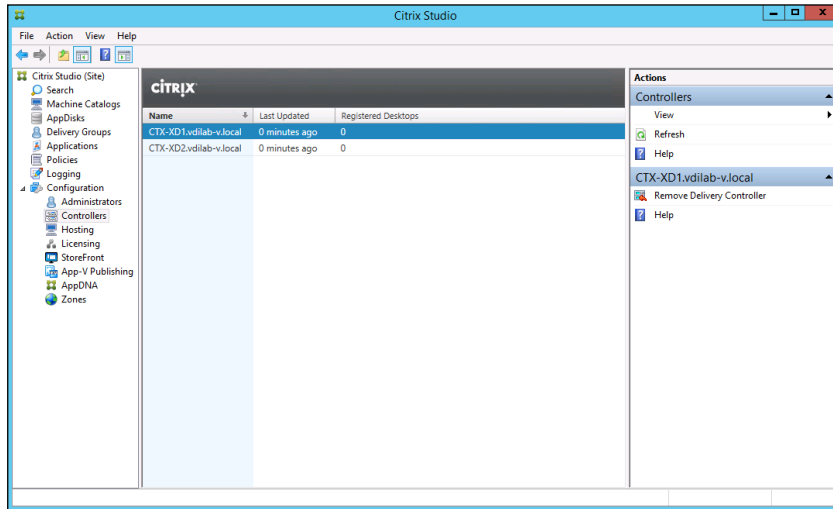


4. Click Yes to allow the database to be updated with this controller's information automatically.



5. When complete, test the site configuration and verify the Delivery Controller has been added to the list of Controllers.





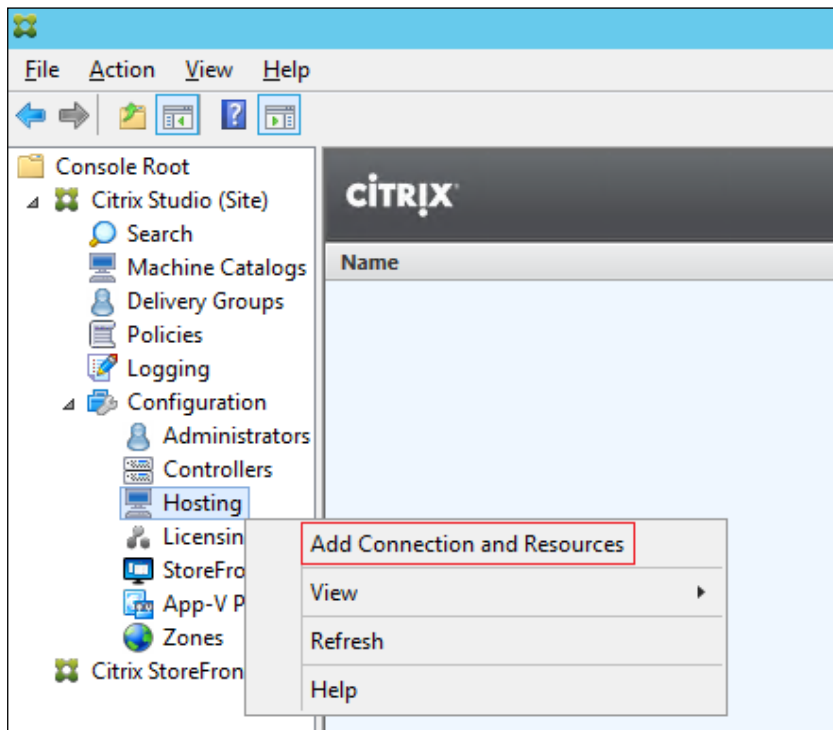
Create Host Connections with Citrix Studio

Citrix Studio provides wizards to guide the process of setting up an environment and creating desktops. To set up a host connection for a cluster of VMs for the VDI desktops, complete the following steps:

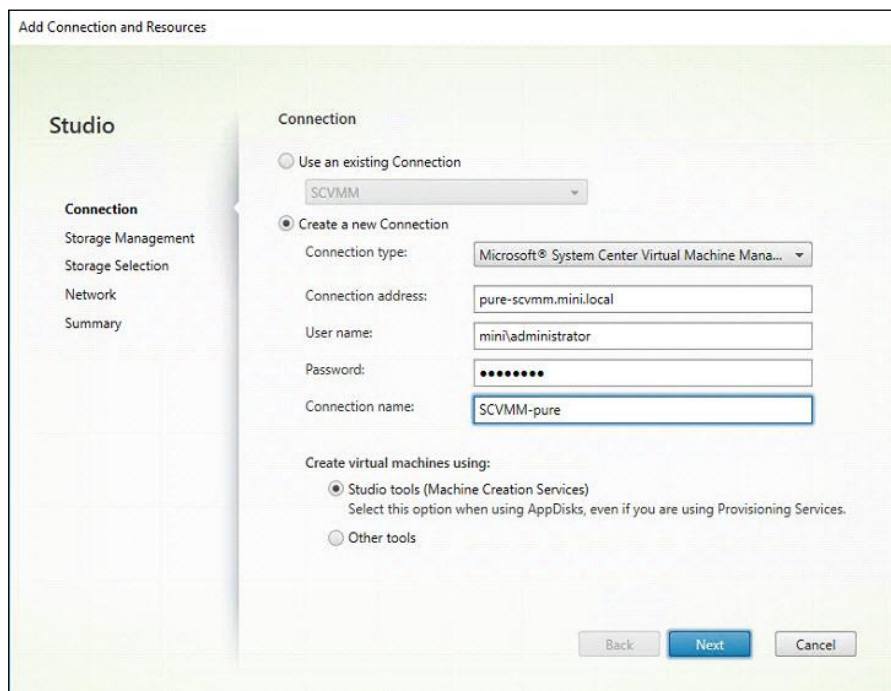


The instructions below outline the procedure to add a host connection and resources for HSD and VDI desktops.

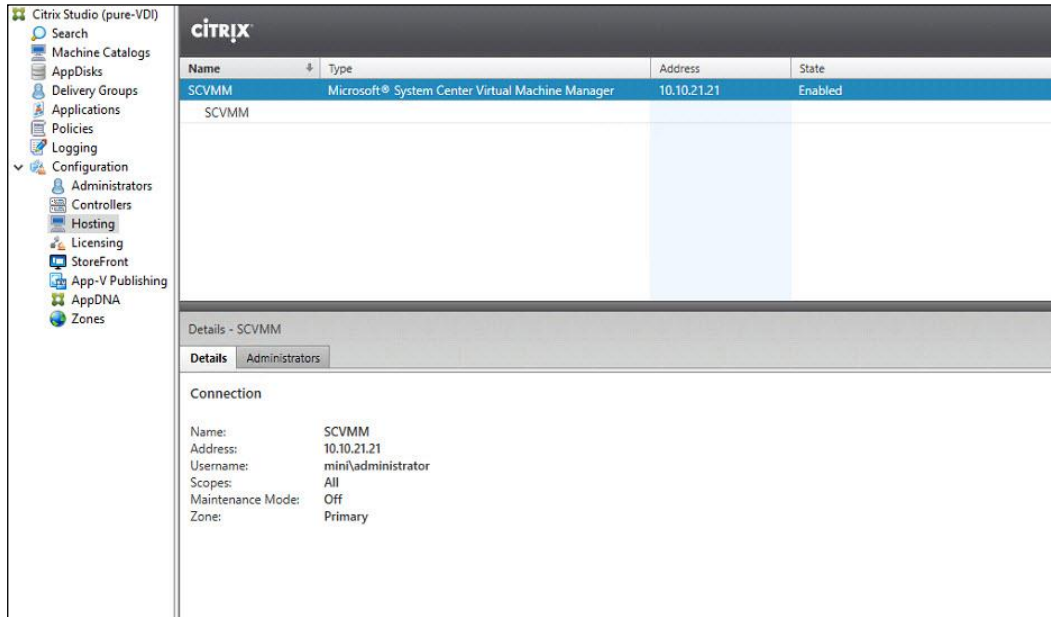
1. Connect to the XenDesktop server and launch Citrix Studio.
2. From the Configuration menu, right-click Hosting and select Add Connection and Resources.



3. Select the Host Type of Microsoft System Center Virtual Machine Manager.
4. Enter the FQDN of the SCVMM server.
5. Enter the username (in domain\username format) for the domain admin account.
6. Provide the password for the domain admin account.
7. Provide a connection name.
8. Select the Studio tools radio button since MCS will be used.
9. Click Next.



10. Review the Summary.
11. Click Finish.

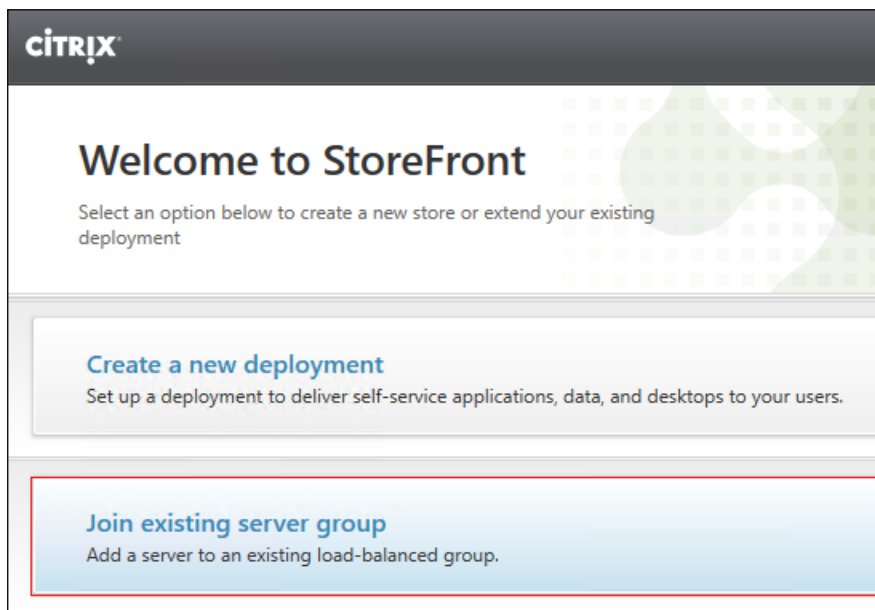


Configuring StoreFront

Citrix StoreFront stores aggregate desktops and applications from XenDesktop sites, making resources readily available to users. In this CVD, StoreFront is installed on the Delivery Controllers virtual machine as part of the initial Delivery Controller installation. Most of the StoreFront configuration is automatically done as part of the installer. To finalize the StoreFront configuration log into the second Delivery Controller and launch the StoreFront Console.

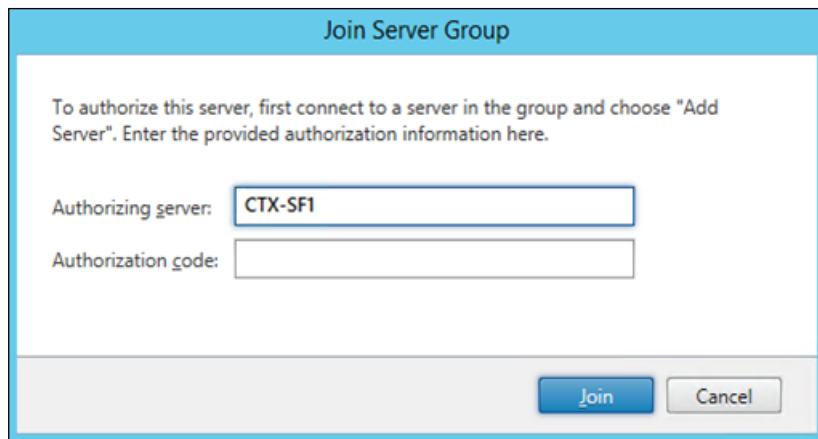
To configure StoreFront, complete the following steps:

1. From the StoreFront Console on the second server select **“Join existing server group.”**

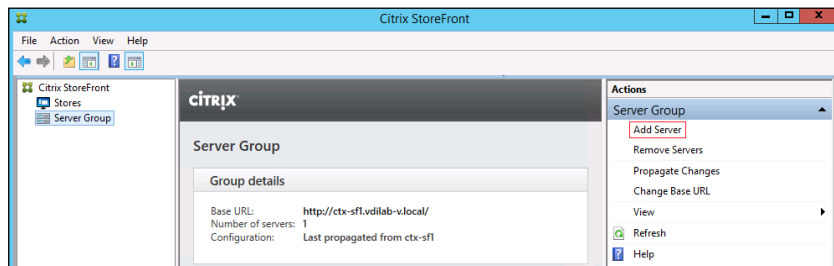


2. In the Join Server Group dialog, enter the name of the first Storefront server.

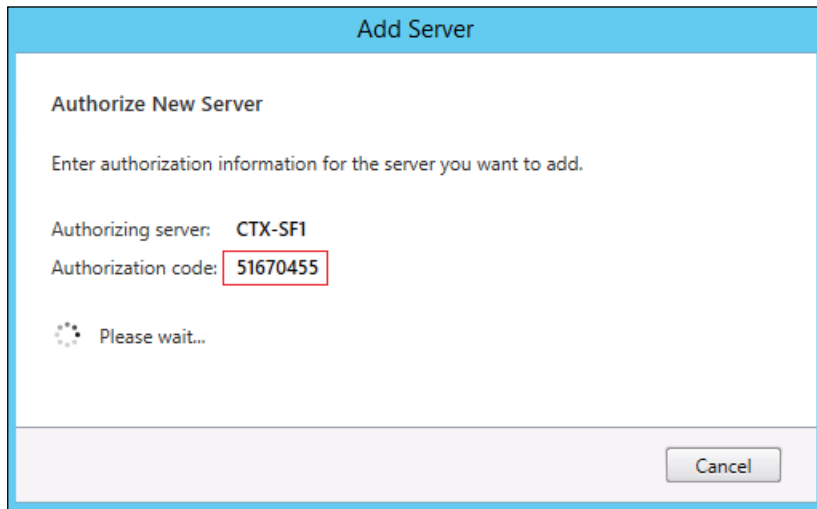
3. Before the additional StoreFront server can join the server group, you must connect to the first Storefront server, add the second server, and obtain the required authorization information.



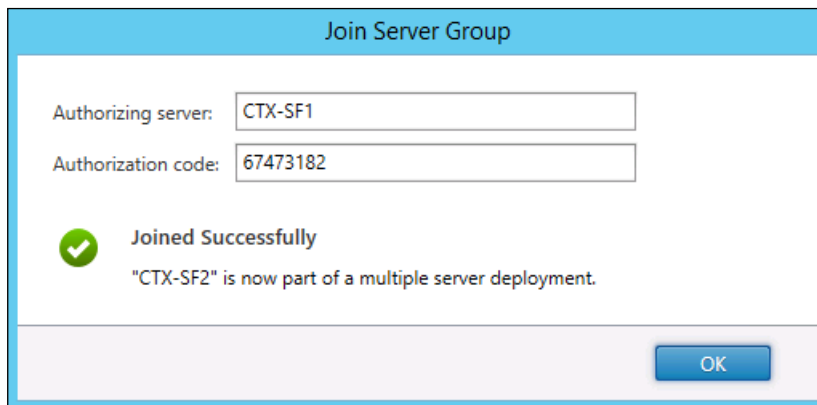
4. Connect to the first StoreFront server.
5. Using the StoreFront menu on the left, you can scroll through the StoreFront management options.
6. Select Server Group from the menu.
7. To add the second server and generate the authorization information that allows the additional StoreFront server to join the server group, select Add Server.



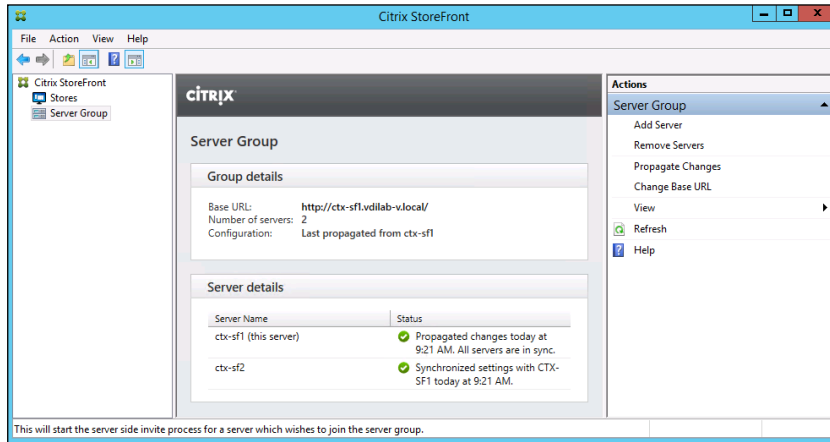
8. Copy the Authorization code from the Add Server dialog.



9. Connect to the second Storefront server and paste the Authorization code into the Join Server Group dialog.
10. Click Join.
11. A message appears when the second server has joined successfully.
12. Click OK.



13. The Server Group now lists both StoreFront servers in the group.

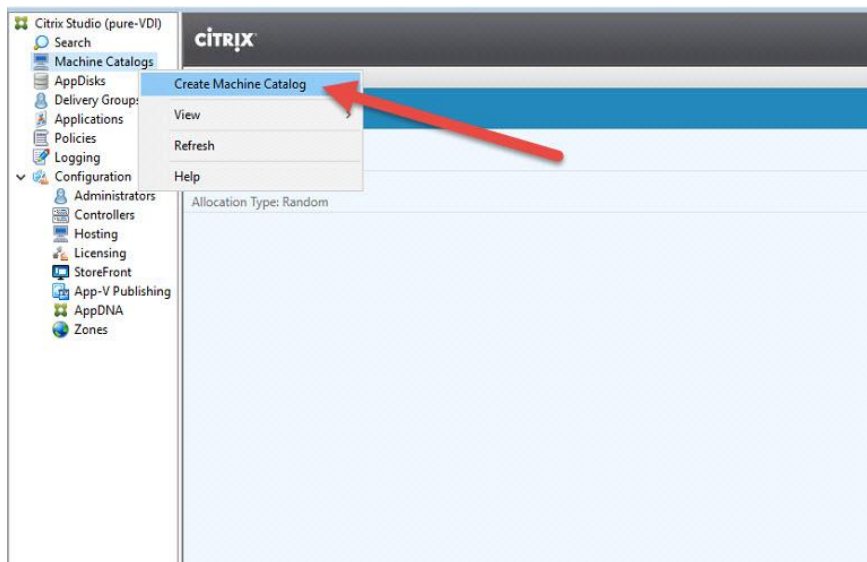


Create Machine Catalogs

Machine Catalogs are collections of machines that organize desktops and applications. With Machine Catalogs, you can organize machines based on what type or OS they are.

To create Machine Catalogs, complete the following steps:

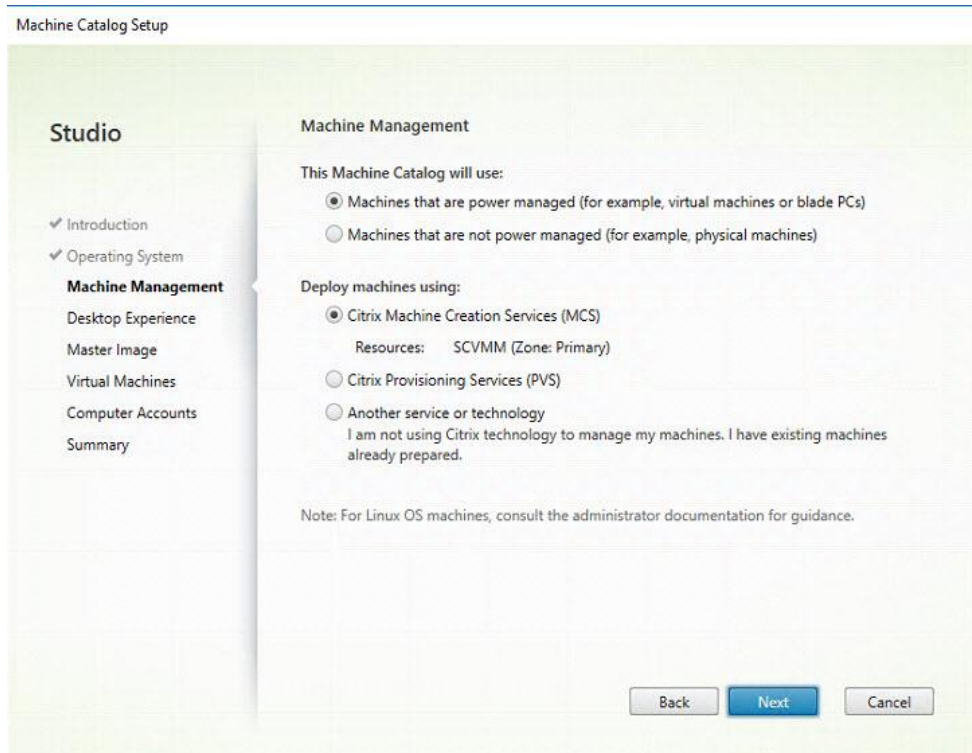
1. Connect to a XenDesktop server and launch Citrix Studio.
2. Right-click Site in Studio and choose Create Machine Catalogs.



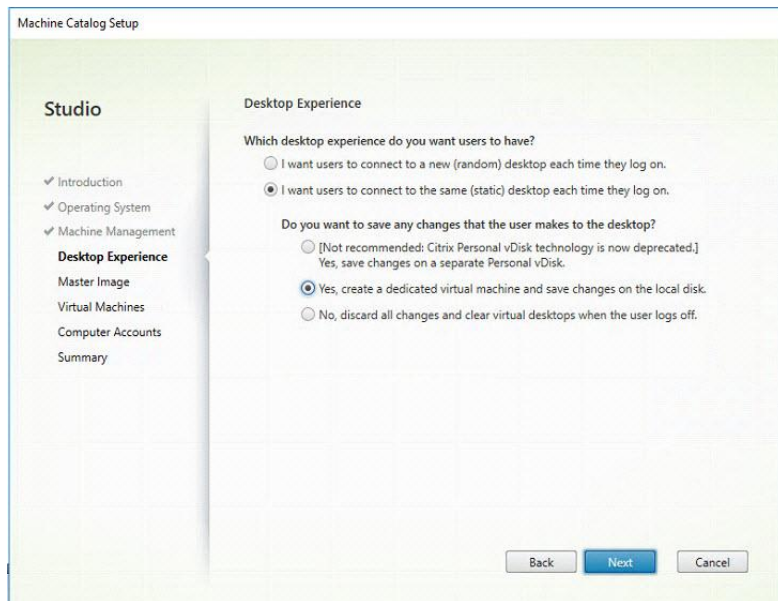
3. Select 'Desktop OS'.



You will be deploying MCS machines using the Host connection or Zone you created earlier.

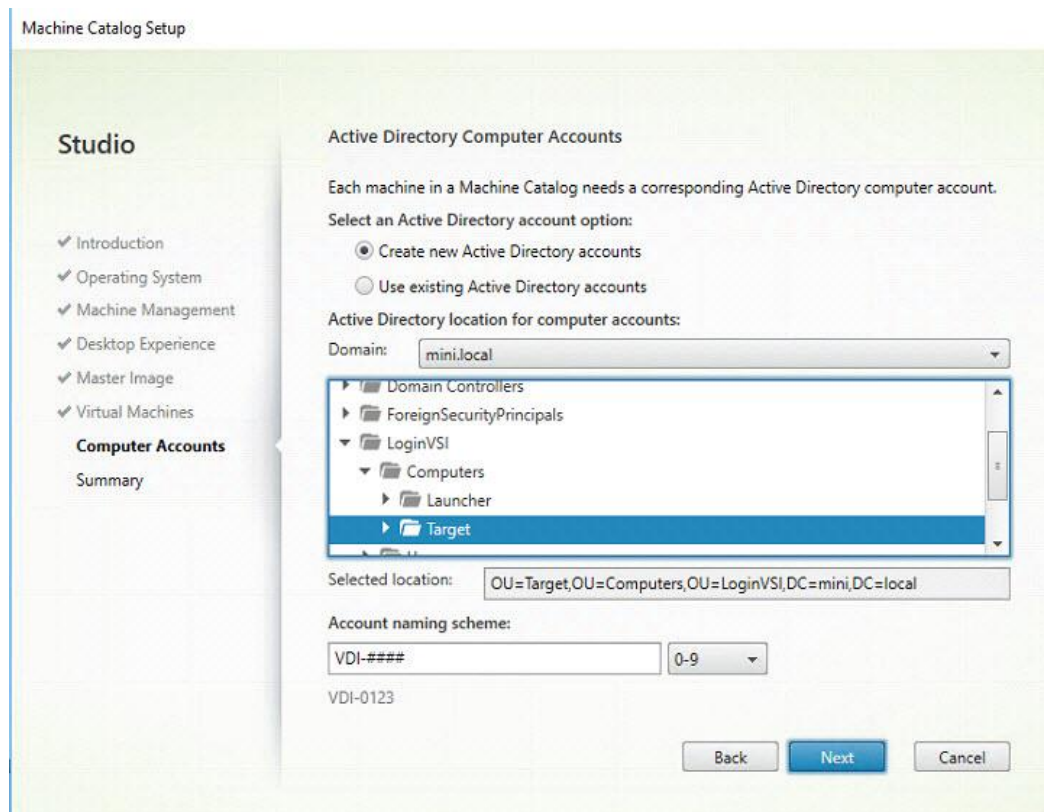


4. Select the option for a 'Static' desktop and a 'Dedicated' machine.



5. Select the VM for the Master Image click Next.

6. Navigate to the OU you would like the VMs to be placed and provide a naming convention.



7. Give the Catalog a name and click Finish.

Create Delivery Groups

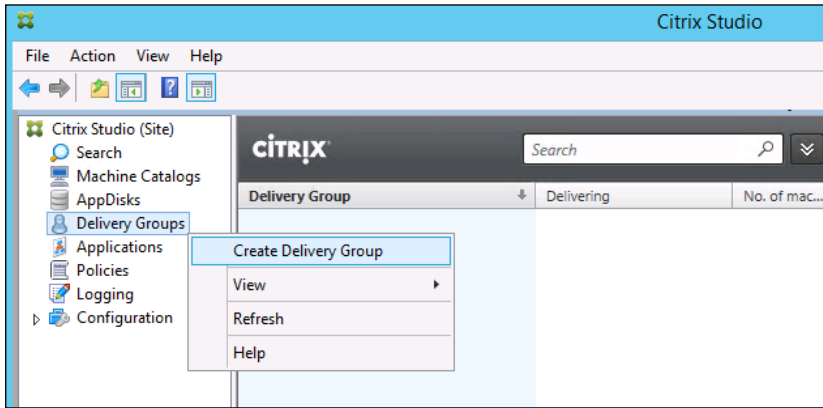
Delivery Groups are collections of machines that control access to desktops and applications. With Delivery Groups, you can specify which users and groups can access which desktops and applications.

To create delivery groups, complete the following steps:

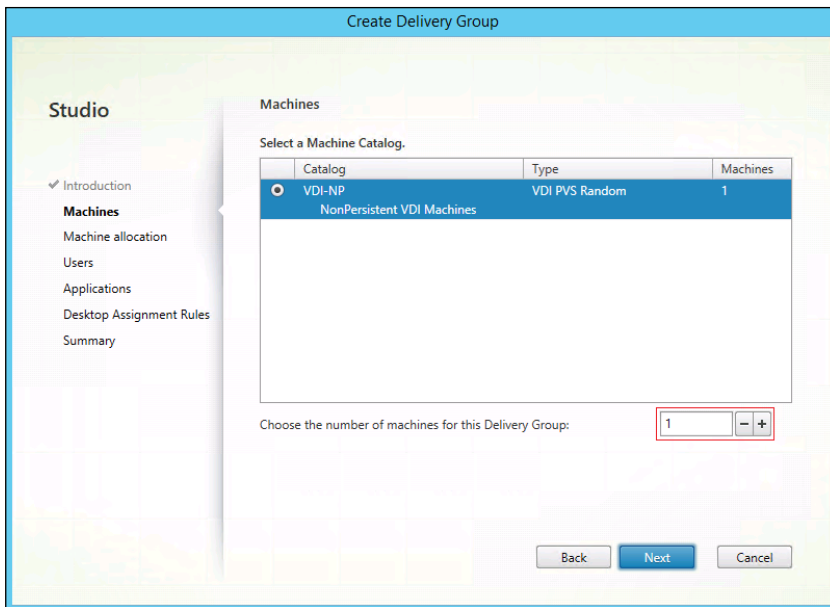


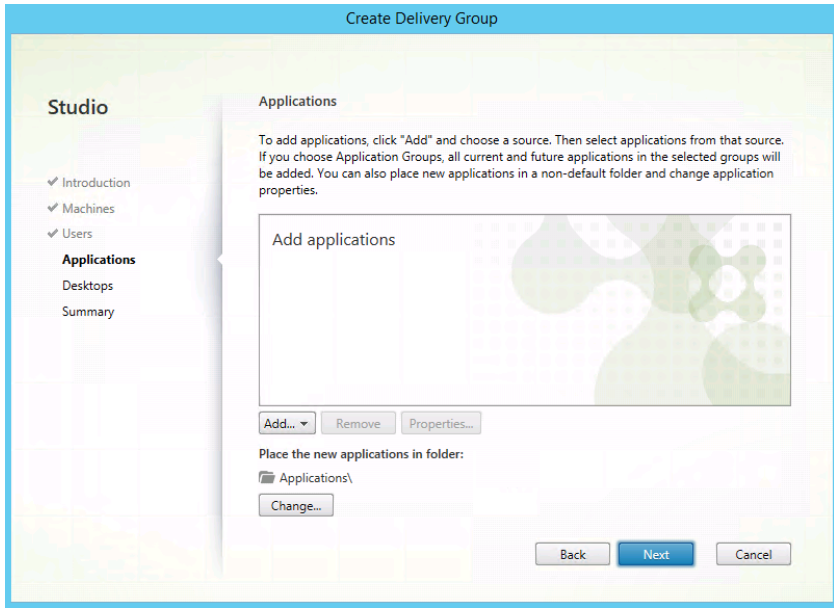
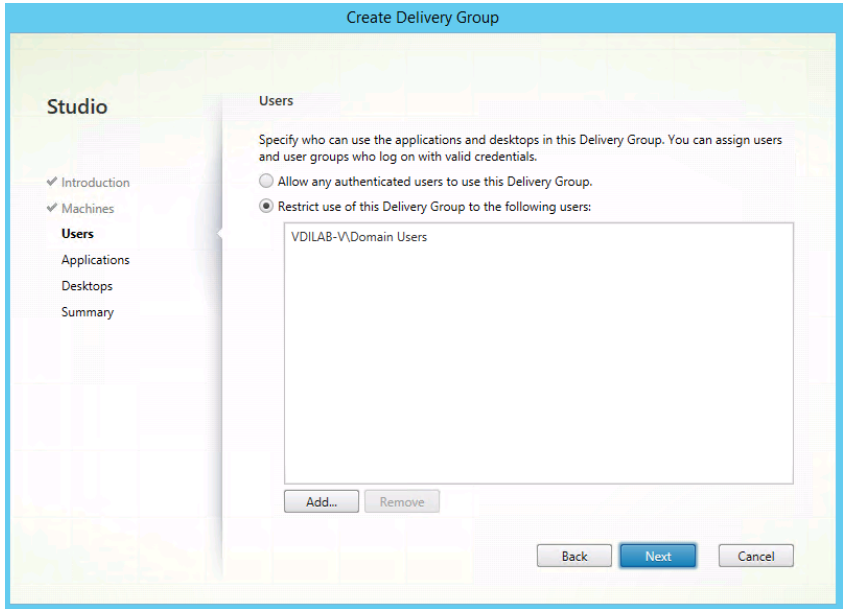
The instructions below outline the procedure to create a Delivery Group for VDI desktops. When you have completed these steps, repeat the procedure to a Delivery Group for RDS desktops.

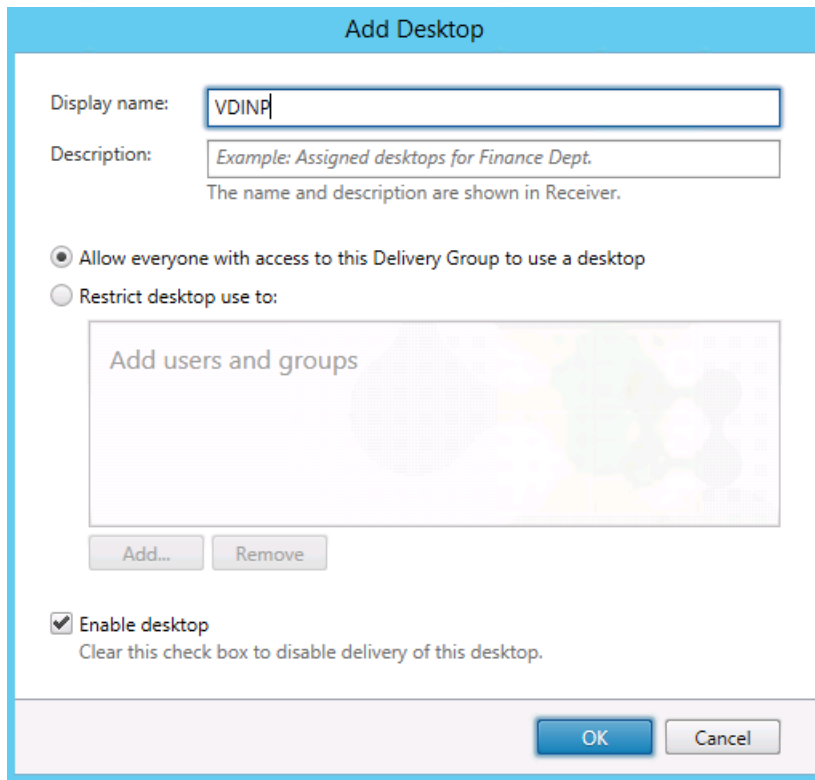
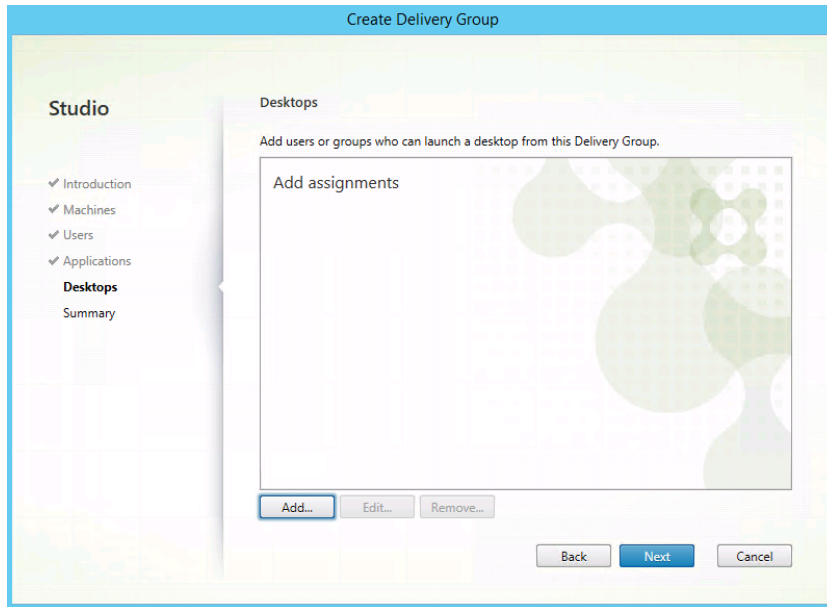
1. Connect to a XenDesktop server and launch Citrix Studio.
2. Choose Create Delivery Group from the drop-down menu.

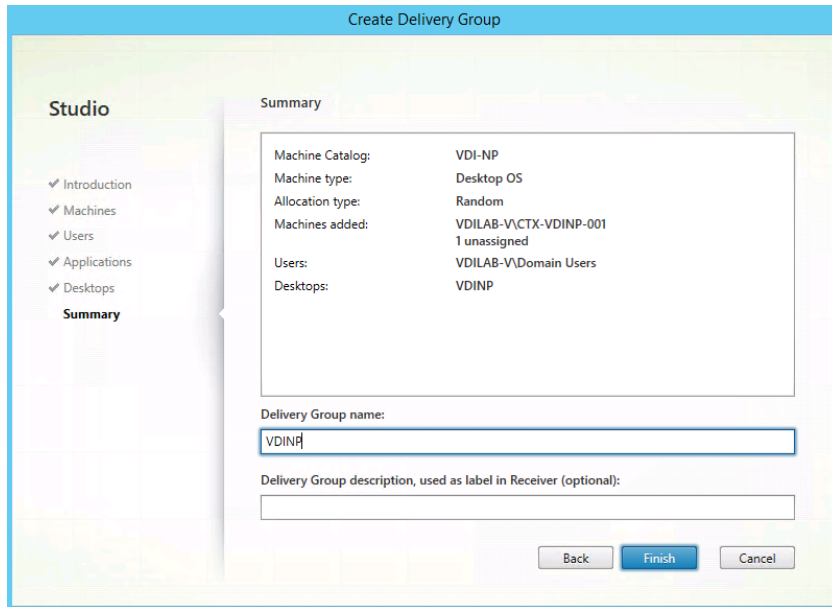


3. Specify the Machine Catalog and increment the number of machines to add.
4. Click Next.

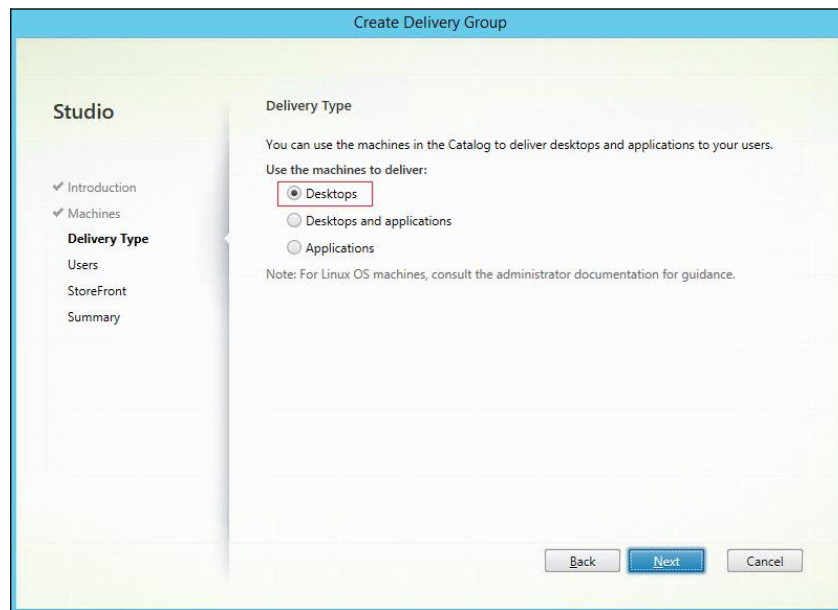




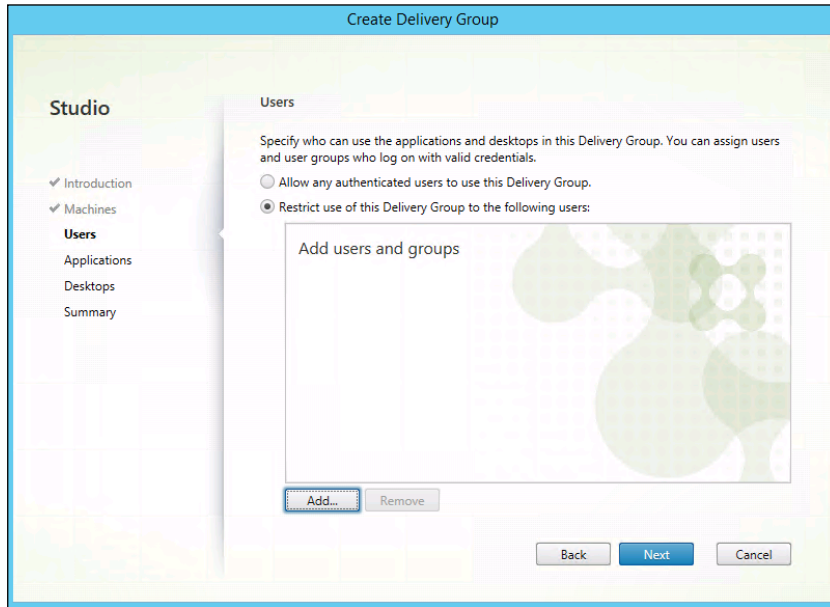




5. Specify what the machines in the catalog will deliver: Desktops, Desktops and Applications, or Applications.
6. Select Desktops.
7. Click Next.

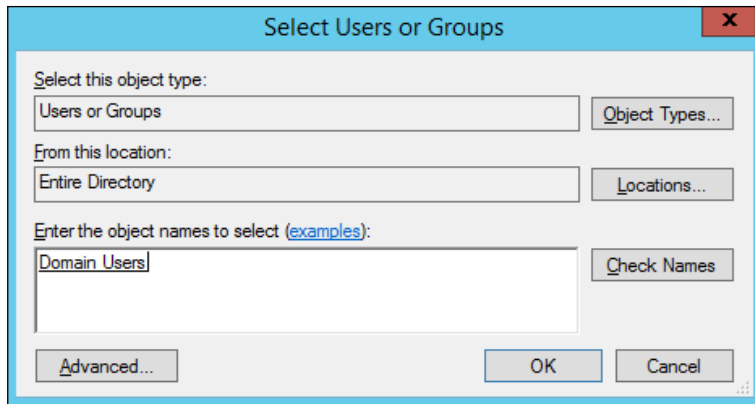


8. To make the Delivery Group accessible, you must add users, click Add.



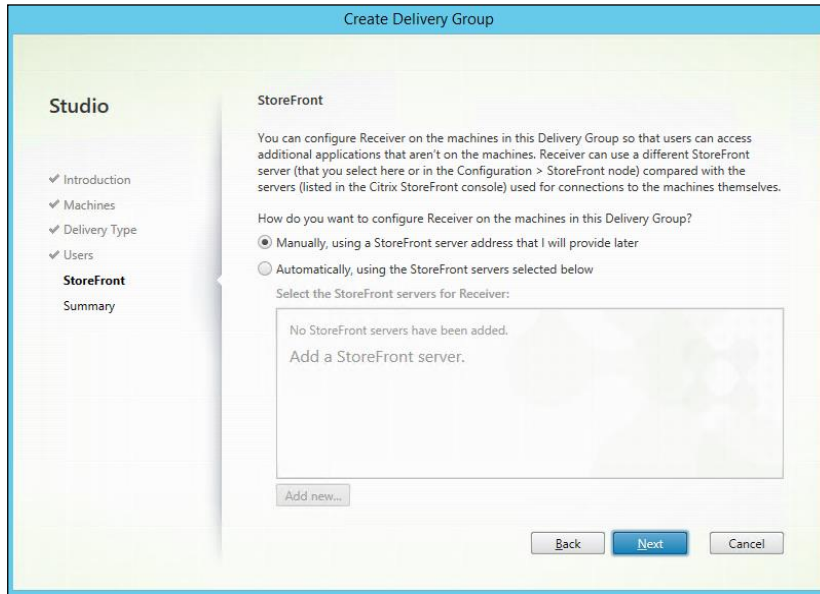
9. In the Select Users or Groups dialog, add users or groups.

10. Click OK. When users have been added, click Next.



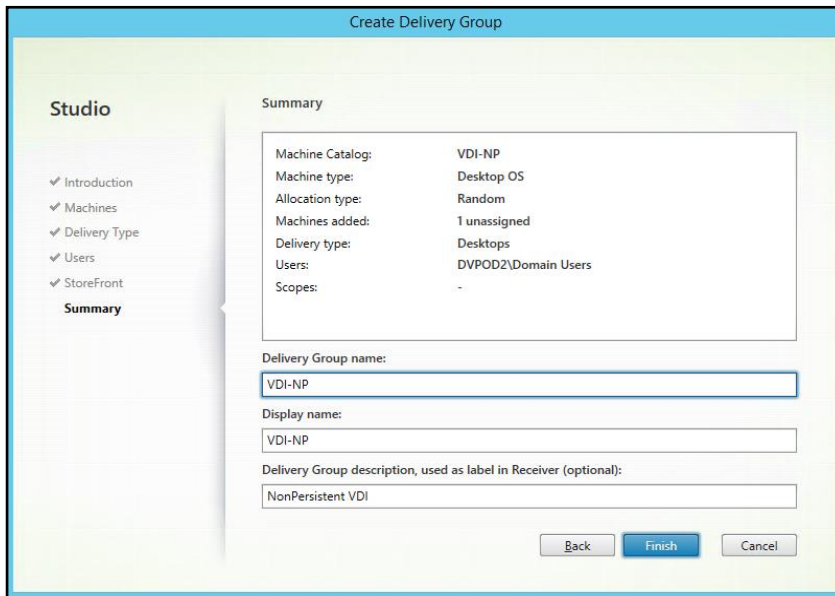
11. Enter the StoreFront configuration for how Receiver will be installed on the machines in this Delivery Group. Click “Manually, using a StoreFront server address that I will provide later.”

12. Click Next.



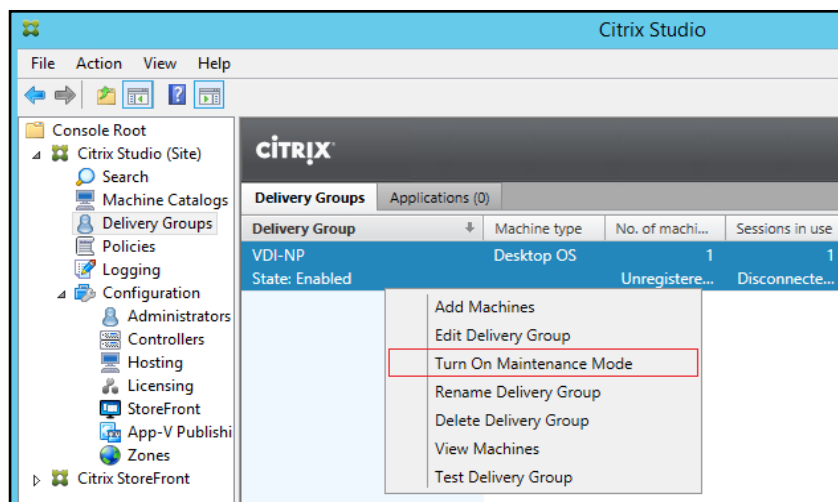
13. On the Summary dialog, review the configuration. Enter a Delivery Group name and a Display name (for example, VDI or RDS).

14. Click Finish.



15. Citrix Studio lists the created Delivery Groups and the type, number of machines created, sessions, and applications for each group in the Delivery Groups tab.

16. On the drop-down menu, select “Turn on Maintenance Mode.”



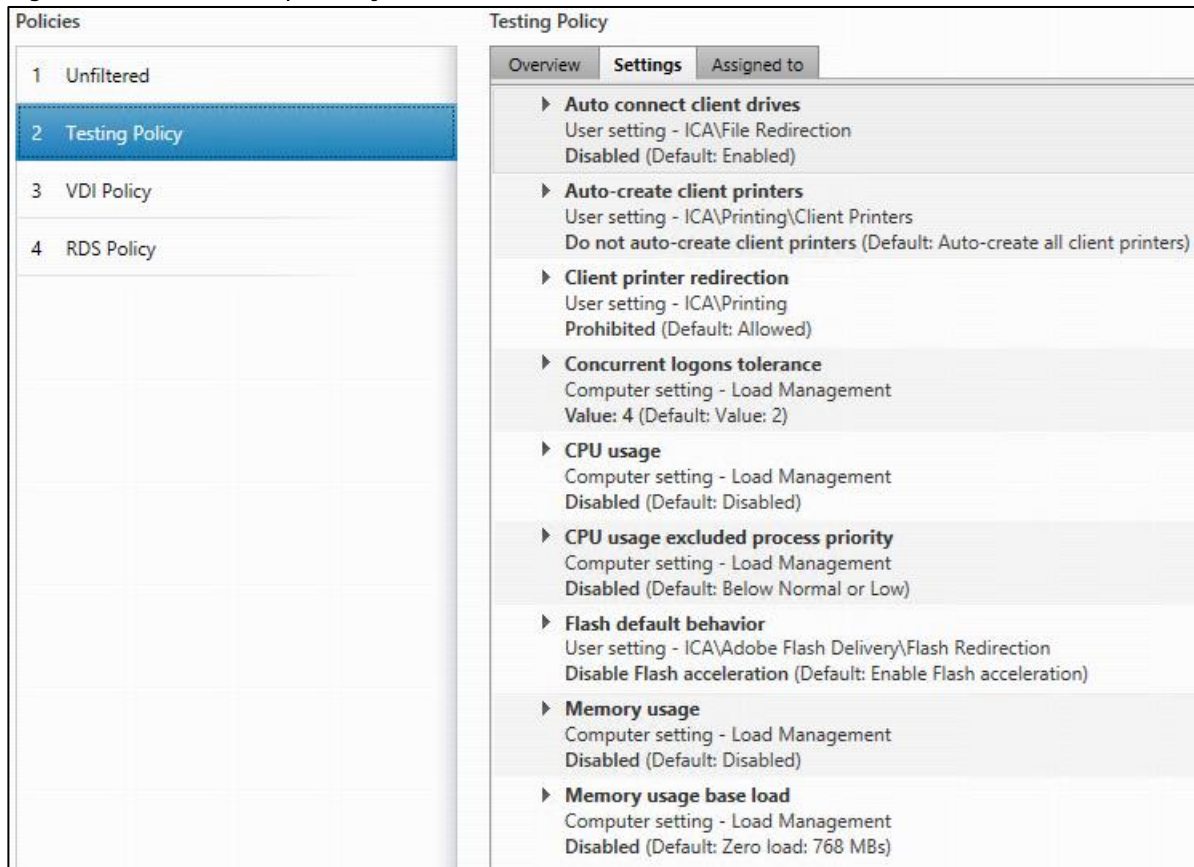
Citrix XenDesktop Policies and Profile Management

Policies and profiles allow the Citrix XenDesktop environment to be easily and efficiently customized.

Configure Citrix XenDesktop Policies

Citrix XenDesktop policies control user access and session environments, and are the most efficient method of controlling connection, security, and bandwidth settings. You can create policies for specific groups of users, devices, or connection types with each policy. Policies can contain multiple settings and are typically defined through Citrix Studio. (The Windows Group Policy Management Console can also be used if the network environment includes Microsoft Active Directory and permissions are set for managing Group Policy Objects). The figure below shows policies for Login VSI testing in this CVD.

Figure 4 XenDesktop Policy



Configuring User Profile Management

Profile management provides an easy, reliable, and high-performance way to manage user personalization settings in virtualized or physical Windows environments. It requires minimal infrastructure and administration and provides users with fast logons and logoffs. A Windows user profile is a collection of folders, files, registry settings, and configuration settings that define the environment for a user who logs on with a particular user account. These settings may be customizable by the user, depending on the administrative configuration. Examples of settings that can be customized are:

- Desktop settings such as wallpaper and screen saver
- Shortcuts and Start menu setting
- Internet Explorer Favorites and Home Page
- Microsoft Outlook signature
- Printers

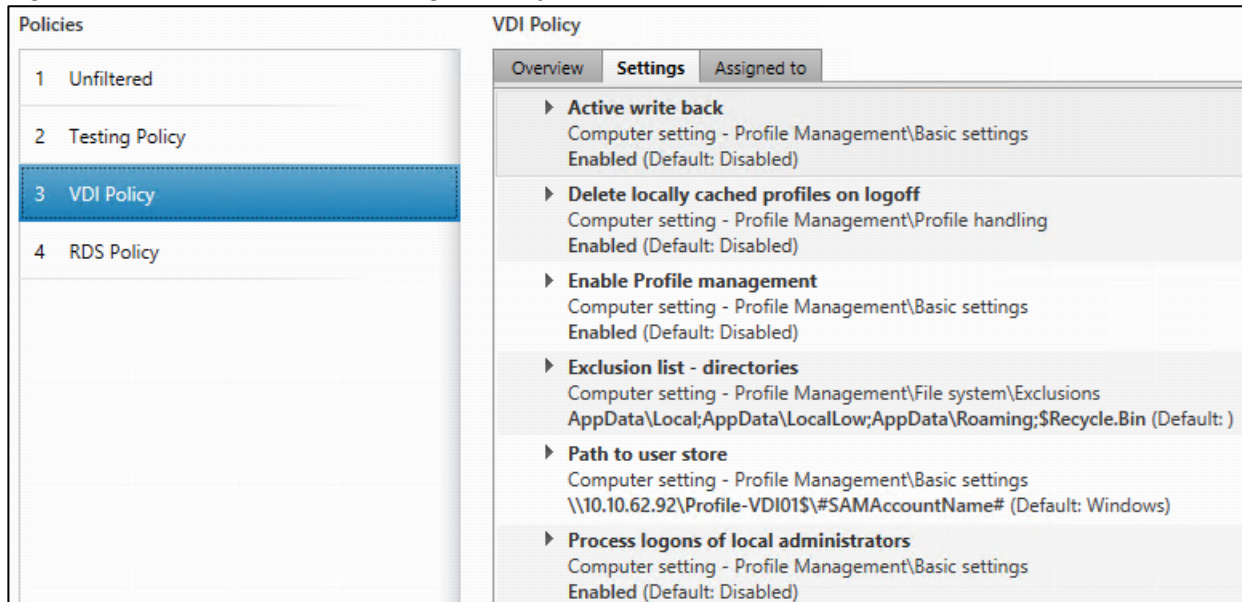
Some user settings and data can be redirected by means of folder redirection. However, if folder redirection is not used these settings are stored within the user profile.

The first stage in planning a profile management deployment is to decide on a set of policy settings that together form a suitable configuration for your environment and users. The automatic configuration feature

simplifies some of this decision-making for XenDesktop deployments. Screenshots of the User Profile Management interfaces that establish policies for this CVD's RDS and VDI users (for testing purposes) are shown below. Basic profile management policy settings are documented here:

<https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-15-ltsr.html>

Figure 5 VDI User Profile Manager Policy



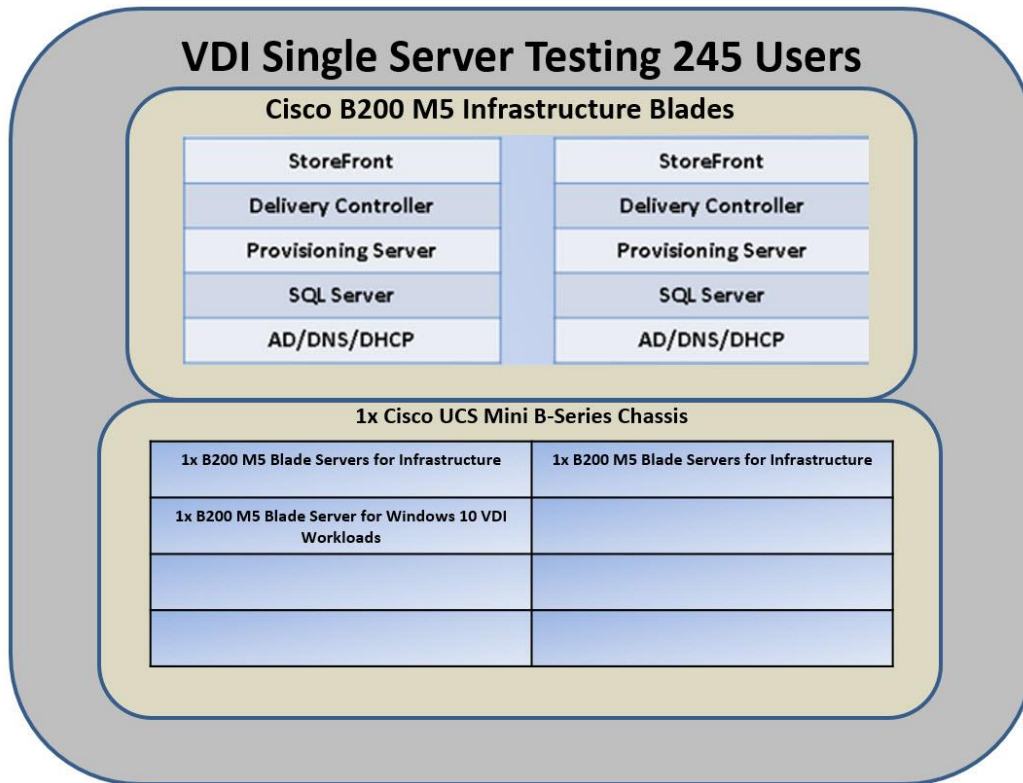
Test Setup and Configurations

In this project, we tested a single Cisco UCS Mini running eight (8) Cisco UCS B200 M5 Blade Servers in a single Cisco UCS domain. This solution is tested to illustrate linear scalability for each workload studied.

The configuration of this smaller footprint solution is to be able to achieve scalability from a very small footprint of 1 workload server and 2 Infrastructure servers that can run up to 245 users in the chassis. We then can expand all the way to 6 workload servers and 2 Infrastructure servers that can house up to 1250 users.

This solution allows customers that need a small footprint solution for smaller offices, branches, edge computing or disaster recovery sites, to name a few.

- With a single workload host and 2 Infrastructure servers, we achieve N+1 by creating a 3 node cluster and allowing all 245 user desktops and all required Infrastructure to failover and handle the workloads (i.e. if the workload blade fails, the two infrastructure blades can easily handle the Windows 10 desktops along with Infrastructure VMs).



- This solution will allow a customer to scale from a single workload server to a fully populated UCS Mini Chassis with 6 workload servers and 2 Infrastructure servers. This can grow a solution from 200+ users to approximately 1250 Users with N+1 capability utilizing Infrastructure servers as failover options.


VDI Full Chassis Testing, 1250 Users

Cisco B200 M5 Infrastructure Blades

StoreFront	StoreFront
Delivery Controller	Delivery Controller
Provisioning Server	Provisioning Server
SQL Server	SQL Server
AD/DNS/DHCP	AD/DNS/DHCP

1x Cisco UCS Mini B-Series Chassis

1x B200 M5 Blade Servers for Infrastructure	1x B200 M5 Blade Servers for Infrastructure
1x B200 M5 Blade Server for Windows 10 VDI Workloads	1x B200 M5 Blade Server for Windows 10 VDI Workloads
1x B200 M5 Blade Server for Windows 10 VDI Workloads	1x B200 M5 Blade Server for Windows 10 VDI Workloads
1x B200 M5 Blade Server for Windows 10 VDI Workloads	1x B200 M5 Blade Server for Windows 10 VDI Workloads



Hardware Components:

- 2 x Cisco UCS 6324 Fabric Interconnects
- 2 x Cisco Nexus 9372PX Access Switches
- 8 x Cisco UCS B200 M5 Blade Servers (2x Intel Scalable Processor 6140 Gold at 2.3 GHZ with 768 GB of memory per server [64 GB x 12 DIMM at 2666 MHz]).
- 1 x Pure Storage //M10 Flash Array

Software components:

- Cisco UCS firmware 3.2(2b)

- Microsoft Hyper-V Server 2016
- Citrix XenDesktop 7.15
- Citrix User Profile Management
- Citrix NetScaler VPX NS11.1 52.13.nc
- Microsoft SQL Server 2016
- Microsoft Windows 10
- Microsoft Windows 2016
- Microsoft Office 2016
- Login VSI 4.1.25

Testing Methodology and Success Criteria

All validation testing was conducted on-site within the Cisco labs in San Jose, California.

The testing results focused on the entire process of the virtual desktop lifecycle by capturing metrics during the desktop boot-up, user logon and virtual desktop acquisition (also referred to as ramp-up,) user workload execution (also referred to as steady state), and user logoff for the Hosted Shared Desktop Session under test.

Test metrics were gathered from the virtual desktop, storage, and load generation software to assess the overall success of an individual test cycle. Each test cycle was not considered passing unless all of the planned test users completed the ramp-up and steady state phases (described below) and unless all metrics were within the permissible thresholds as noted as success criteria.

Three successfully completed test cycles were conducted for each hardware configuration and results were found to be relatively consistent from one test to the next.

You can obtain additional information and a free test license from <http://www.loginvsi.com>.

Testing Procedure

The following protocol was used for each test cycle in this study to insure consistent results.

Pre-Test Setup for Testing

All machines were shut down utilizing the Citrix XenDesktop 7.15 LTSR Administrator Console.

All Launchers for the test were shut down. They were then restarted in groups of 10 each minute until the **required number of launchers was running with the Login VSI Agent at a “waiting for test to start” state.**

Test Run Protocol

To simulate severe, real-world environments, Cisco requires the log-on and start-work sequence, known as Ramp Up, to complete in 48 minutes. Additionally, we require all sessions started, whether 60 single server users or 4000 full scale test users to become active within two minutes after the last session is launched.

In addition, Cisco requires that the Login VSI Benchmark method is used for all single server and scale testing. This assures that our tests represent real-world scenarios. For each of the three consecutive runs on single server tests, the same process was followed. Complete the following steps:

1. Time 0:00:00 Start esxtop Logging on the following systems:

- – Infrastructure and VDI Host Blades used in test run
- – All Infrastructure VMs used in test run (AD, SQL, XenDesktop brokers, image mgmt., etc.)

Time 0:00:10 Start Storage Partner Performance Logging on Storage System.

Time 0:05: Boot VDI Machines using Citrix XenDesktop 7.15 Administrator Console.

Time 0:06 First machines boot.

Time 0:35 Single Server or Scale target number of VDI Desktops registered on XD.



No more than 60 Minutes of rest time is allowed after the last desktop is registered and available on Citrix XenDesktop 7.15 Administrator Console dashboard. Typically, a 20-30 minute rest period for Windows 10 desktops and 10 minutes for RDS VMs is sufficient.

Time 1:35 Start Login VSI 4.1.25 Knowledge Worker Benchmark Mode Test, setting auto-logoff time at 900 seconds, with Single Server or Scale target number of desktop VMs utilizing sufficient number of Launchers (at 20-25 sessions/Launcher).

Time 2:23 Single Server or Scale target number of desktop VMs desktops launched (48 minute benchmark launch rate).

Time 2:25 All launched sessions must become active.



All sessions launched must become active for a valid test run within this window.

Time 2:40 Login VSI Test Ends (based on Auto Logoff 900 Second period designated above).

Time 2:55 All active sessions logged off.

All sessions launched and active must be logged off for a valid test run. The Citrix XenDesktop 7.15 Administrator Dashboard must show that all desktops have been returned to the registered/available state as evidence of this condition being met.

Time 2:57 All logging terminated; Test complete.

Time 3:15 Copy all log files off to archive; Set virtual desktops to maintenance mode through broker; Shut-down all Windows 10 machines.

Time 3:30 Reboot all hypervisors.

Time 3:45 Ready for new test sequence.

Success Criteria

Our “pass” criteria for this testing is as follows: Cisco will run tests at a session count levels that effectively utilize the server capacity measured by CPU, memory, storage and network utilization. We use Login VSI version 4.1.25 to launch Knowledge Worker workload sessions. The number of launched sessions must equal active sessions within two minutes of the last session launched in a test as observed on the VSI Management console.

The Citrix XenDesktop Studio will be monitored throughout the steady state to make sure of the following:

All running sessions report In Use throughout the steady state

No sessions move to unregistered, unavailable or available state at any time during steady state

Within 20 minutes of the end of the test, all sessions on all launchers must have logged out automatically and **the Login VSI Agent must have shut down. Cisco’s tolerance** for Stuck Sessions is 0.5% (half of one percent.) If the Stuck Session count exceeds that value, we identify it as a test failure condition.

Cisco requires three consecutive runs with results within +/-1% variability to pass the Cisco Validated Design performance criteria. For white papers written by partners, two consecutive runs within +/-1% variability are accepted. (All test data from partner run testing must be supplied along with proposed white paper.)

We will publish Cisco Validated Designs with our recommended workload following the process above and will note that we did not reach a VSImax dynamic in our testing.

The purpose of this testing is to provide the data needed to validate Citrix XenDesktop 7.153 VDI with Citrix XenDesktop 7.15 MCS provisioning using Microsoft Windows 10 sessions on Cisco UCS B200 M5 servers.

The information contained in this section provides data points that a customer may reference in designing their own implementations. These validation results are an example of what is possible under the specific environment conditions outlined here and do not represent the full characterization of Citrix and VMware products.

Four test sequences, each containing three consecutive test runs generating the same result, were performed to establish system performance and linear scalability.

VSImax 4.1.x Description

The philosophy behind Login VSI is different to conventional benchmarks. In general, most system benchmarks are steady state benchmarks. These benchmarks execute one or multiple processes, and the measured execution time is the outcome of the test. Simply put: the faster the execution time or the bigger the throughput, the faster the system is according to the benchmark.

Login VSI is different in approach. Login VSI is not primarily designed to be a steady state benchmark (however, if needed, Login VSI can act like one). Login VSI was designed to perform benchmarks for SBC or

VDI workloads through system saturation. Login VSI loads the system with simulated user workloads using well-known desktop applications like Microsoft Office, Internet Explorer and Adobe PDF reader. By gradually increasing a number of simulated users, the system will eventually be saturated. Once the system is saturated, the response time of the applications will increase significantly. This latency in application response times shows a clear indication whether the system is (close to being) overloaded. As a result, by nearly overloading a system it is possible to find out what its true maximum user capacity is.

After a test is performed, the response times can be analyzed to calculate the maximum active session/desktop capacity. Within Login VSI this is calculated as VSImax. When the system is coming closer to its saturation point, response times will rise. When reviewing the average response time, it will be clear the response times escalate at saturation point.

This VSImax is the “Virtual Session Index (VSI)”. With Virtual Desktop Infrastructure (VDI) and Terminal Services (RDS) workloads this is valid and useful information. This index simplifies comparisons and makes it possible to understand the true impact of configuration changes on hypervisor host or guest level.

Server-Side Response Time Measurements

It is important to understand why specific Login VSI design choices have been made. An important design choice is to execute the workload directly on the target system within the session instead of using remote sessions. The scripts simulating the workloads are performed by an engine that executes workload scripts on every target system and are initiated at logon **within the simulated user’s desktop session context**.

An alternative to the Login VSI method would be to generate user actions client side through the remoting protocol. These methods are always specific to a product and vendor dependent. More importantly, some protocols simply do not have a method to script user actions client side.

For Login VSI the choice has been made to execute the scripts completely server side. This is the only practical and platform independent solutions, for a benchmark like Login VSI.

Calculating VSImax v4.1.x

The simulated desktop workload is scripted in a 48-minute loop when a simulated Login VSI user is logged on, performing generic Office worker activities. After the loop is finished it will restart automatically. Within each loop the response times of sixteen specific operations are measured in a regular interval: sixteen times in within each loop. The response times of these five operations are used to determine VSImax.

The five operations from which the response times are measured are:

- Notepad File Open (NFO)

Loading and initiating VSINotepad.exe and opening the openfile dialog. This operation is handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-**user’s** point of view.

- Notepad Start Load (NSLD)

Loading and initiating VSINotepad.exe and opening a file. This operation is also handled by the OS and by the VSINotepad.exe itself through execution. This operation seems almost instant from an end-**user’s point of view**.

- Zip High Compression (ZHC)

This action copy's a random file and compresses it (with 7zip) with high compression enabled. The compression will very briefly spike CPU and disk IO.

- Zip Low Compression (ZLC)

This action copy's a random file and compresses it (with 7zip) with low compression enabled. The compression will very briefly disk IO and creates some load on the CPU.

- CPU

Calculates a large array of random data and spikes the CPU for a short period of time.

These measured operations within Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations become consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. If such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive.

Figure 6 Sample of a VSI Max Response Time Graph, Representing a Normal Test





Figure 7 Sample of a VSI Test Response Time Graph with a Clear Performance Issue



When the test is finished, VSI_{max} can be calculated. When the system is not saturated, and it could complete the full test without exceeding the average response time latency threshold, VSI_{max} is not reached and a number of sessions ran successfully.

The response times are very different per measurement type, for instance Zip with compression can be around 2800 ms, while the Zip action without compression can only take 75ms. This response time of these actions are weighted before they are added to the total. This ensures that each activity has an equal impact on the total response time.

In comparison to previous VSImax models, this weighting much better represent system performance. All actions have very similar weight in the VSImax total. The following weighting of the response times are applied.

The following actions are part of the VSImax v4.1 calculation and are weighted as follows (US notation):

- Notepad File Open (NFO): 0.75
- Notepad Start Load (NSLD): 0.2
- Zip High Compression (ZHC): 0.125
- Zip Low Compression (ZLC): 0.2
- CPU: 0.75

This weighting is applied on the baseline and normal Login VSI response times.

With the introduction of Login VSI 4.1 we also created a new method to calculate the base phase of an environment. With the new workloads (Taskworker, Powerworker, etc.) enabling 'base phase' for a more reliable baseline has become obsolete. The calculation is explained below. In total 15 lowest VSI response time samples are taken from the entire test, the lowest 2 samples are removed and the 13 remaining samples are averaged. The result is the Baseline. The calculation is as follows:

- Take the lowest 15 samples of the complete test
- From those 15 samples remove the lowest 2
- Average the 13 results that are left is the baseline

The VSImax average response time in Login VSI 4.1.x is calculated on a number of active users that are logged on the system.

Always a 5 Login VSI response time samples are averaged + 40% of the amount of “active” sessions. For example, if the active sessions is 60, then latest 5 + 24 (=40% of 60) = 31 response time measurement are used for the average calculation.

To remove noise (accidental spikes) from the calculation, the top 5% and bottom 5% of the VSI response time samples are removed from the average calculation, with a minimum of 1 top and 1 bottom sample. As a result, with 60 active users, the last 31 VSI response time sample are taken. From those 31 samples the top 2 samples are removed and lowest 2 results are removed (5% of 31 = 1.55, rounded to 2). At 60 users the average is then calculated over the 27 remaining results.

VSImax v4.1.x is reached when the VSIbase + a 1000 ms latency threshold is not reached by the average VSI response time result. Depending on the tested system, VSImax response time can grow 2 - 3x the baseline average. In end-user computing, a 3x increase in response time in comparison to the baseline is typically regarded as the maximum performance degradation to be considered acceptable.

In VSImax v4.1.x this latency threshold is fixed to 1000ms, this allows better and fairer comparisons between two different systems, especially when they have different baseline results. Ultimately, in VSImax v4.1.x, the performance of the system is not decided by the total average response time, but by the latency it has under load. For all systems, this is now 1000ms (weighted).

The threshold for the total response time is: average weighted baseline response time + 1000ms.

When the system has a weighted baseline response time average of 1500ms, the maximum average response time may not be greater than 2500ms (1500+1000). If the average baseline is 3000 the maximum average response time may not be greater than 4000ms (3000+1000).

When the threshold is not exceeded by the average VSI response time during the test, VSImax is not hit and a number of sessions ran successfully. This approach is fundamentally different in comparison to previous VSImax methods, as it was always required to saturate the system beyond VSImax threshold.

Lastly, VSImax v4.1.x is now always reported with the average baseline VSI response time result. For **example: "The VSImax v4.1 was 125 with a baseline of 1526ms"**. This helps considerably in the comparison of systems and gives a more complete understanding of the system. The baseline performance helps to understand the best performance the system can give to an individual user. VSImax indicates what the total user capacity is for the system. These two are not automatically connected and related:

When a server with a very fast dual core CPU, running at 3.6 GHZ, is compared to a 10 core CPU, running at 2,26 GHZ, the dual core machine will give an individual user better performance than the 10 core machine. This is indicated by the baseline VSI response time. The lower this score is, the better performance an individual user can expect.

However, the server with the slower 10 core CPU will easily have a larger capacity than the faster dual core system. This is indicated by VSImax v4.1.x, and the higher VSImax is, the larger overall user capacity can be expected.

With Login VSI 4.1.x a new VSImax method is introduced: VSImax v4.1. This methodology gives much better insight into system performance and scales to extremely large systems.


Test Results

Boot Storms

A key performance metric for desktop virtualization environments is the ability to boot the virtual machines quickly and efficiently to minimize user wait time for their desktop.

As part of Cisco's virtual desktop test protocol, we shut down each virtual machine at the conclusion of a benchmark test. When we run a new test, we cold boot all 1250 desktops and measure the time it takes for the 1250th virtual machine to register as available in the Citrix XenDesktop Studio.

The Cisco UCS B200 M5 can accomplish this task in 30 minutes as shown in the following chart:



Recommended Maximum Workload and Configuration Guidelines

Eight Cisco UCS B200 M5 Blade Servers

For the Citrix XenDesktop 7.15 Hosted Virtual Desktop use case, the recommended maximum workload was determined based on both Login VSI Knowledge Worker workload end user experience measures, and B200 M5 server operating parameters.

This recommended maximum workload approach allows you to determine the server N+1 fault tolerance load the server can successfully support in the event of a server outage for maintenance or upgrade.

Our recommendation is that the Login VSI Average Response and VSI Index Average should not exceed the Baseline plus 2000 milliseconds to ensure that end user experience is outstanding. Additionally, during steady state, the processor utilization should average no more than 90-95%.



Memory should never be oversubscribed for Desktop Virtualization workloads.

Test Phase	Description
Boot	Start all RDS and/or VDI virtual machines at the same time.
Login	The Login VSI phase of test is where sessions are launched and start executing the workload over a 48 minutes duration.
Steady state	The steady state phase is where all users are logged in and performing various workload tasks such as using Microsoft Office, Web browsing, PDF printing, playing videos, and compressing files.
Logoff	Sessions finish executing the Login VSI workload and logoff.



The recommended maximum workload for a Cisco B200 M5 node with Intel 6140 Gold processors and 768GB of RAM for Windows 10 persistent Hosted Virtual Desktop users is 250 sessions with Office 2016 virtual desktops respectively.

245 User Single Server Testing on Cisco UCS B200 M5 Server

This section shows the key performance metrics captured on a single Cisco UCS workload blades and two Infrastructure blades (three Cisco UCS B200 M5), running VDI VMs. The Single Server testing with 245 users comprised of: 245 VDI (MCS. Full copy).

Test result highlights include:

0.665 second baseline response time (sub-second)

0.919 second average response time with 1250 desktop sessions running (sub-second)

Average CPU utilization of 45 percent during steady state

Average of 283 GB of RAM used out of 768 GB available

9700Mbps peak network utilization per host.

Average Read Latency 0.6ms/Max Read Latency 2.2ms

Average Write Latency 3.6ms/Max Write Latency 11.3ms

130000 peak I/O operations per second (IOPS) per cluster at steady state

2700MBps peak throughput per cluster at steady state

74% Deduplication savings

45% Compression savings

Total of 86% storage space savings

Figure 8 LoginVSI Analyzer Chart for 245 Users Test




Figure 9 CPU Utilization for Single Server Testing




Figure 10 LoginVSI Analyzer Chart for Three (3) Consecutive 245 User Tests Running Knowledge Workload on 1 Cisco UCS B200 M5





Figure 11 Sample Hyper-V Host Memory Usage in Mbytes Running 1250 User Test on 6 Node
 Cisco UCS B200 M5 Memory Utilization (\\memory\\available mbytes)
 Windows 10, Office 2016 245 Users



1250 User Full Scale Testing on 6-node Hyper-V Cluster

This section shows the key performance metrics captured on the Cisco UCS Hyper-V Cluster and Pure Storage Array, configured with six workload blades and two Infrastructure blades (eight Cisco UCS B200 M4), running VDI VMs. The full-scale testing with 1250 users comprised of: 1250 VDI (MCS. Full copy).

Test result highlights include:

- 0.665 second baseline response time (sub-second)
- 0.919 second average response time with 1250 desktop sessions running (sub-second)
- Average CPU utilization of 45 percent during steady state
- Average of 283 GB of RAM used out of 768 GB available
- 9700Mbps peak network utilization per host.
- Average Read Latency 0.6ms/Max Read Latency 2.2ms
- Average Write Latency 3.6ms/Max Write Latency 11.3ms
- 130000 peak I/O operations per second (IOPS) per cluster at steady state
- 2700MBps peak throughput per cluster at steady state
- 74% Deduplication savings

- 45% Compression savings
- Total of 86% storage space savings

Figure 12 LoginVSI Analyzer Chart for 1250 Users Test




Figure 13 LoginVSI Analyzer Chart for Three (3) Consecutive 1250 User Tests Running Knowledge Workload on 6 Node UCS Cluster




Figure 14 Sample Hyper-V CPU Core Utilization Running 1250 User Test on 6 Nodes





Figure 15 Sample Hyper-V Host Memory Usage in Mbytes running 1250 User Test on 6 Node



Pure Storage FlashArray//M10 Test Results for 1250 Persistent Windows 10 x64 MCS Desktops

The cluster-level simulation was to run 1250 persistent Windows 10 x64 desktops against the same FlashArray//M10. All Login VSI parameters were kept consistent with the Single Server test with the only change being to use 1250 desktops created via XenDesktop Machine Creation Services. As can be seen in the below storage metrics, the Pure Storage FlashArray//M10 was clearly able to handle this workload and continued to provide sub-millisecond latency for another impressive Login VSI result.

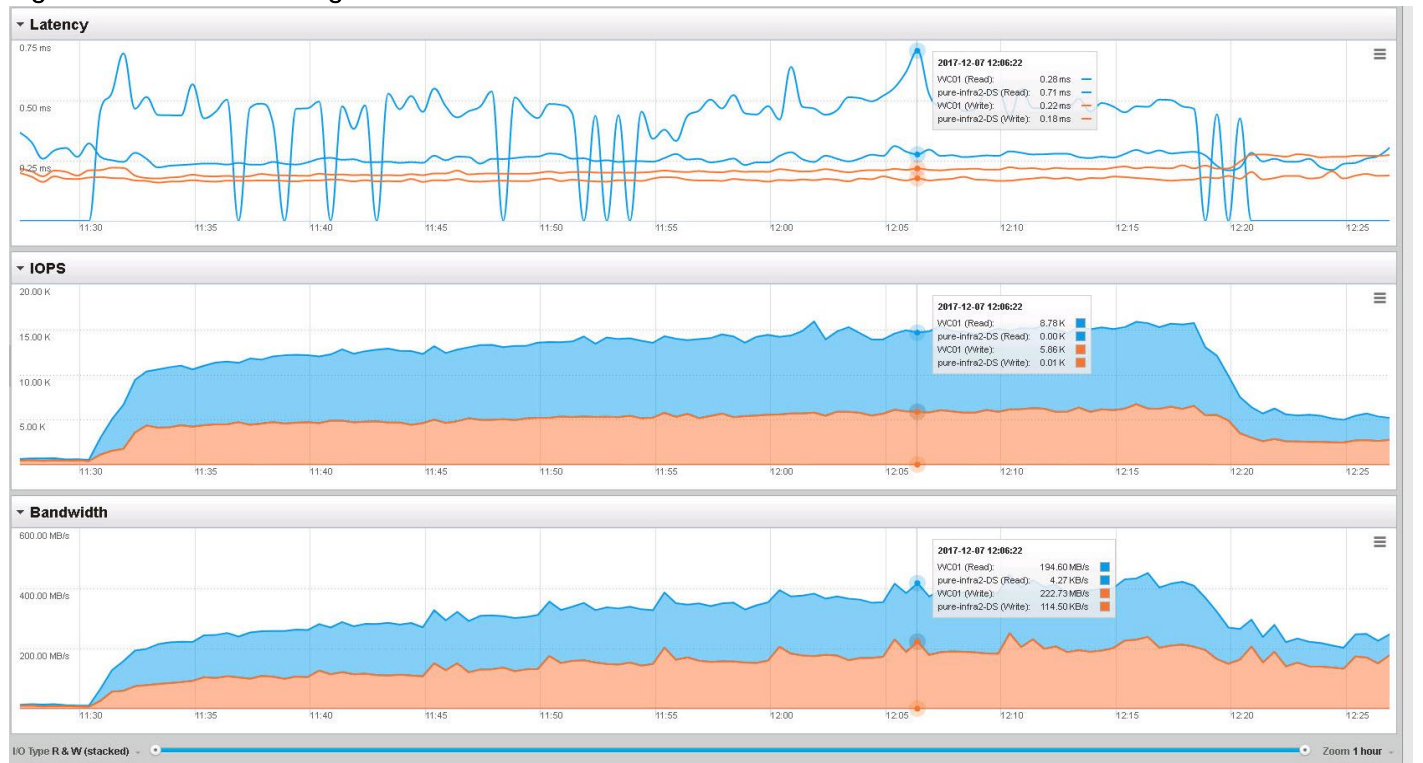
Another item worth noting from the GUI screenshot shows that latency was consistently sub-millisecond throughout all phases of the Login VSI simulation despite driving hundreds of megabytes of sustained write bandwidth.

1. During the boot storms the Pure Array handled more than 33 thousand IOPS with more than 1 GB of throughput all while maintaining a sub millisecond latency, Figure 16.
2. During the scale user tests of 1250 users, the Pure Array handled more than 15 thousand IOPS with more than 400 MBs of throughput. All while maintaining a sub millisecond latency, Figure 17.

Figure 16 Boot Storms



Figure 17 1250 User LoginVSI Test



Scalability Considerations and Guidelines

There are many factors to consider when you begin to scale beyond 1250 Users, one chassis 6 VDI workload host server configuration, which this reference architecture has successfully tested. In this section we give guidance to scale beyond the 1250 user system.

Cisco UCS System Scalability

As our results indicate, we have proven linear scalability in the Cisco UCS Reference Architecture as tested:

- Cisco UCS Manager Software supports up to 20 Cisco UCS chassis within a single Cisco UCS domain with Cisco UCS 6248UP Fabric Interconnect. A single UCS domain can grow to 160 blades for an enterprise deployment.
- Cisco UCS Central, the manager of managers, extends UCS domains and vastly increases the reach of the Cisco UCS system. Simplify daily operations by centrally managing and automating routine tasks and expediting problem resolution. Our powerful platform eliminates disparate management environments. Use it to support up to 10,000 Cisco UCS servers (blade, rack, hyperconverged and Mini) and manage multiple Cisco UCS instances or domains across globally-distributed locations.
- As scale grows, the value of the combined UCS fabric, Nexus physical switches and Nexus virtual switches increases dramatically to define the Quality of Services required to deliver excellent end user experience 100 percent of the time.
- To accommodate the Cisco Nexus 9000 upstream connectivity in the way we describe in the network configuration section, two Ethernet uplinks are needed to be configured on the Cisco UCS 6324 Fabric Interconnect.

The backend storage has to be scaled accordingly, based on the IOP considerations as described in the Pure Storage scaling section. Please refer the Pure Storage web site for scalability guidelines.

Scalability of Citrix XenDesktop 7.15 Configuration

XenDesktop environments can scale to large numbers. When implementing Citrix XenDesktop, consider the following in scaling the number of hosted shared and hosted virtual desktops:

- Types of storage in your environment
- Types of desktops that will be deployed
- Data protection requirements

These and other various aspects of scalability considerations are **described in greater detail in “XenDesktop - Modular Reference Architecture” document and should be a part of any XenDesktop design.**

When designing and deploying this CVD environment, best practices were followed including the following:

- Citrix recommends using N+1 schema for virtualization host servers to accommodate resiliency. In all Reference Architectures (such as this CVD), this recommendation is applied to all host servers.
- All Provisioning Server Network Adapters are configured to have a static IP and management.

We used the XenDesktop Machine Creation Services deployment wizard. The wizard does an excellent job of creating the desktops automatically and it's possible to run multiple instances of the wizard, provided the deployed desktops are placed in different catalogs and have different naming conventions.

FlashStack 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 Cisco 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 perform a backup, complete the following steps:

1. From Cisco UCS Manager, select Admin within the Navigation pane and select All.
2. Click on the Policy Backup and Export tab within All.
3. For a Full State Backup, All Configuration Backup, or both, specify the following:
 - a. Hostname : <IP or FQDN of host that will receive the backup>
 - b. Protocol: [FTP/TFTP/SCP/SFTP]
 - c. User: <account on host to authenticate>
 - d. Password: <password for account on host>
 - e. Remote File: <full path and filename prefix for backup file>
 - f. Admin State: <select Enable to activate the schedule on save, Disable to disable schedule on save>
 - g. Schedule: [Daily/Weekly/Bi Weekly]

4. Click Save Changes to create the Policy.

Cisco Nexus Backups

The configuration of the Cisco Nexus 9000 switches can be backed up manually at any time with the copy command, but automated backups can be put in place with the NX-OS feature scheduler. An example of setting up an automated configuration backup for one of the FlashStack 93180YC-EX switches is shown below:

```
bb04-9332-a# conf t
```

Enter configuration commands, one per line. End with CNTL/Z.

```
bb04-9332-a(config)# feature scheduler
```

```
bb04-9332-a(config)# scheduler logfile size 1024
```

```
bb04-9332-a(config)# scheduler job name backup-cfg
```

```
bb04-9332-a(config-job)# copy running-config
tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP) vrf management

bb04-9332-a(config-job)# exit

bb04-9332-a(config)# scheduler schedule name daily

bb04-9332-a(config-schedule)# job name backup-cfg

bb04-9332-a(config-schedule)# time daily 2:00

bb04-9332-a(config-schedule)# end
```

Show the job that has been setup:

```
bb04-9332-a# sh scheduler job

Job Name: backup-cfg

-----

copy running-config tftp://192.168.156.155/9332/$(SWITCHNAME)-cfg.$(TIMESTAMP)
vrf management

=====
```

```
bb04-9332-a# show scheduler schedule

Schedule Name      : daily
-----
User Name          : admin
Schedule Type      : Run every day at 2 Hrs 0 Mins
Last Execution Time : Sun Apr  9 02:00:00 2017
Last Completion Time: Sun Apr  9 02:00:01 2017
Execution count    : 3
```

```
-----

Job Name          Last Execution Status
-----
backup-cfg        Success (0)

=====
```

For detailed information about the feature scheduler can be found at:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/system_management/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x/b_Cisco_Nexus_9000_Series_NX-OS_System_Management_Configuration_Guide_7x_chapter_01010.html

Pure Storage Snapshots for Array Protection

A critical factor to the success of any VDI project is having a robust backup and disaster recovery policy and plan in place at the storage layer. Having the ability to revert an entire datastore to an earlier version if a mistake or catastrophic event occurs can save hundreds of hours and protect proprietary user data from loss or corruption. Pure Storage includes snapshots and both synchronous as well as asynchronous array-based replication as a feature of the Purity Operating Environment at no additional cost.

Pure Storage snapshots are immutable, take up zero space upon initial creation (future snapshots are comprised of metadata changes over time) and can be instantly recovered. This flexibility and efficiency enables administrators to design a backup policy logically based upon their specific workloads and risk tolerance rather than being artificially constrained due to complexity or free storage array space limitations. It is highly recommended to assign a frequent (at least once a day) snapshot policy to datastores containing core infrastructure VMs as well as persistent desktops with user data. We also recommend placing non-persistent desktops on a separate datastore with no snapshot policy since those VMs are destroyed and recreated upon user logoff and contain no unique data.

Similar to datastore creation, setting up a snapshot schedule is also accomplished easily and intuitively from within the Pure Storage GUI.

To setup a snapshot schedule, complete the following steps:

1. Click the Protection tab.



2. Click the '+' sign on the Source Groups menu to the left:



3. Give the Protection Group a recognizable name:

Create Protection Group ✕

Name:

4. Using the above Snapshot schedule for the Infrastructure datastore, assign the following values to create the snapshot schedule and retention policy:

Infra-Datastore

Snapshots
0 GB

Schedules (none) | Targets (0) | Members (0) | Snapshots (0)

Snapshot Schedule

Create a snapshot on source every days at

Retain all snapshots on source for days

then retain snapshots per day for more days

Replication Schedule

Replicate a snapshot to targets every hours at

except between and

Retain all snapshots on targets for days

then retain snapshots per day for more days

5. When the Snapshot and retention policy has been defined, add the Datastore to the policy by clicking the Members tab of the Source Group:

Infra-Datastore

Snapshots
0 GB

Schedules (snapshot) | Targets (0) | Members (0) | Snapshots (0) 0 of 0 ⚙


No members have been added

Add Hosts

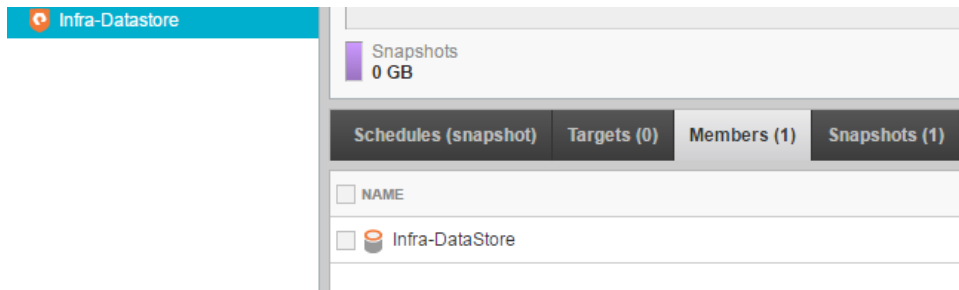
Add Host Groups

Add Volumes

6. Click the Volume(s) you wish to include in the Snapshot policy to move them to the right. Click Confirm when completed.



After clicking Confirm, the Datastore should be listed as a Snapshot policy member:



Repeat the above steps for any datastores containing user data and Hyper-V Boot LUNs datastores. It is recommended to include all Hyper-V boot LUNs in a single Snapshot policy for simplicity.

References

This section provides links to additional information for each partner's solution component of this document.

Cisco UCS B-Series Servers

- <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/datasheet-c78-739296.html>
- <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>

Cisco UCS Manager Configuration Guides

- <http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html>

Citrix References

- <https://docs.citrix.com/en-us/xenapp-and-xendesktop/7-15-ltsr.html>
- <https://www.citrix.com/go/jmp/upm.html>
- <https://www.citrix.com/virtualization/hdx/>
- <https://www.citrix.com/products/xenapp-xendesktop/hdx-3d-pro.html>

Microsoft References

- <https://docs.microsoft.com/en-us/system-center/vmm/install?view=sc-vmm-1711>
- <https://social.technet.microsoft.com/wiki/contents/articles/37890.windows-server-2016-installation.aspx>
- <https://docs.microsoft.com/en-us/windows-server/virtualization/hyper-v/get-started/install-the-hyper-v-role-on-windows-server>

Login VSI Documentation

- https://www.loginvsi.com/documentation/Main_Page
- https://www.loginvsi.com/documentation/Start_your_first_test

Pure Storage Reference Documents

- Pure Storage FlashArray//m Datasheet

- <https://www.purestorage.com/products/flash-array-m.html>
- Pure Storage FlashStack Converged Infrastructure Solutions
 - <http://www.purestorage.com/solutions/infrastructure/flashstack.html>
- Pure Storage Best Practices Guide for Hyper-V 2016
 - https://support.purestorage.com/Solutions/Microsoft_Platform_Guide/aaa_Quick_Setup_Steps/Step_01
 - https://support.purestorage.com/Solutions/Microsoft_Platform_Guide/aaa_Quick_Setup_Steps/Step_02
 - https://support.purestorage.com/Solutions/Microsoft_Platform_Guide/aaa_Quick_Setup_Steps/Step_03
 - https://support.purestorage.com/Solutions/Microsoft_Platform_Guide/Hyper-V_Role/aa_Design_guide_for_Microsoft_Applications_with_Hyper-V_on_FlashStack_Mini_CI

About the Authors

Jeff Nichols, Technical Marketing Engineer, Desktop Virtualization and Graphics Solutions, Cisco Systems, Inc.

Jeff Nichols is a Cisco Unified Computing System architect, focusing on Virtual Desktop and Application solutions with extensive experience with VMware ESX/ESXi, XenDesktop, XenApp and Microsoft Remote Desktop Services. He has expert product knowledge in application, desktop and server virtualization across all three major hypervisor platforms and supporting infrastructures including but not limited to Windows Active Directory and Group Policies, User Profiles, DNS, DHCP and major storage platforms.

Acknowledgments

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, the authors would like to thank:

- Kyle Grossmiller, Solutions Architect, Customer Solutions Group, Pure Storage, Inc.