



NetApp Verified Architecture

FlexPod Express with Microsoft Windows 2016 Hyper-V and FAS2600

NVA Deployment

Melissa Palmer, Lindsey Street, NetApp
May 2017 | NVA-1114-DEPLOY | Version 1.0

Reviewed by Cisco Systems, Inc.



TABLE OF CONTENTS

1	Program Summary	4
2	Solution Overview	4
2.1	FlexPod Converged Infrastructure Program.....	4
2.2	NetApp Verified Architecture Program	5
2.3	Solution Technology	5
2.4	Use Case Summary.....	6
3	Technology Requirements	6
3.1	Hardware Requirements	6
3.2	Software Requirements	7
4	FlexPod Express Cabling Information	8
5	Deployment Procedures	9
5.1	Cisco Nexus 31108 Deployment Procedure	10
5.2	NetApp FAS Storage Deployment Procedure: Part 1	17
5.3	Continuation of Node A Configuration and Cluster Configuration	21
5.4	Cisco UCS C-Series Rack Server Deployment Procedure	37
5.5	NetApp FAS Storage Deployment Procedure: Part 2	46
5.6	Microsoft Windows Server 2016 Deployment Procedure.....	47
5.7	Install Microsoft Windows Features	50
5.8	Configure Microsoft Windows	54
5.9	Create Windows Failover Cluster.....	58
5.10	Install and Configure Hyper-V Management Software	67
6	Conclusion	87
	About the Authors	87
	Acknowledgements	88
	Version History	88

LIST OF TABLES

Table 1)	Hardware requirements for the base configuration.....	7
Table 2)	Hardware for scaling the solution by using two hypervisor nodes.	7
Table 3)	Software requirements for the base FlexPod Express implementation.	7
Table 4)	Software requirements for a Microsoft Windows Server Hyper-V 2016 implementation.	7
Table 5)	Cabling information for Cisco Nexus switch 31108 A.	8
Table 6)	Cabling information for Cisco Nexus switch 31108 B.	9

Table 7) Cabling information for NetApp FAS2650 storage controller A.	9
Table 8) Cabling information for NetApp FAS2650 storage controller B.	9
Table 9) Required VLANs.....	10
Table 10) Hyper-V virtual machines created.	10
Table 11) ONTAP 9.1 installation and configuration information.	18
Table 12) Information required for iSCSI configuration.	35
Table 13) Information required for SVM administrator addition.	36
Table 14) Information required for CIMC configuration.	37
Table 15) Information required for iSCSI boot configuration.....	40
Table 16) Information required for configuring Hyper-V hosts.....	54

LIST OF FIGURES

Figure 1) FlexPod portfolio.	5
Figure 2) FlexPod Express with Microsoft Windows Server Hyper-V 2016 10GbE architecture.	6
Figure 3) Reference validation cabling.	8

1 Program Summary

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology with which they are familiar in their data center.

FlexPod® Express is a predesigned, best practice datacenter architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp® storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools that they are accustomed to. New FlexPod Express customers can easily adapt to managing a FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

2 Solution Overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure Program.

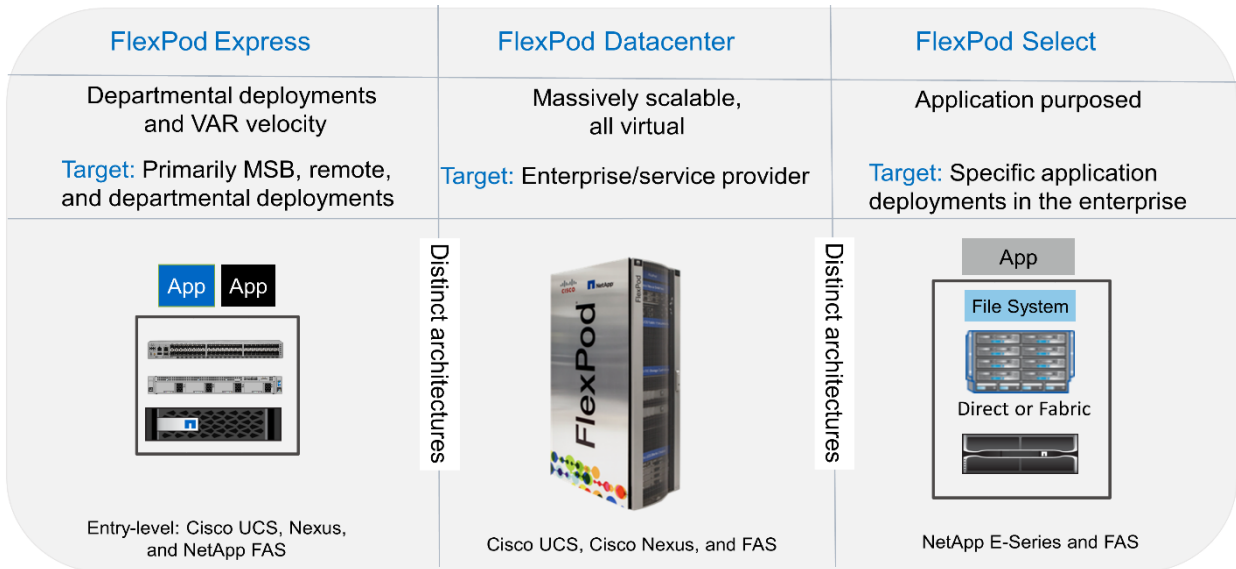
2.1 FlexPod Converged Infrastructure Program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

As depicted in Figure 1, the FlexPod program includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express** offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select** incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.

Figure 1) FlexPod portfolio.



2.2 NetApp Verified Architecture Program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. A NetApp Verified Architecture provides a NetApp solution architecture with the following qualities:


- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

This guide details the design of FlexPod Express with Microsoft Windows Server Hyper-V 2016. In addition, this design uses the all-new FAS2650 system, which runs NetApp ONTAP® 9.1; the Cisco Nexus 31108; and Cisco UCS C-Series C220 M4 servers as hypervisor nodes.

2.3 Solution Technology

This solution leverages the latest technologies from NetApp, Cisco, and Microsoft. This solution features the new NetApp FAS2650 running ONTAP 9.1, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M4 rack servers that run Microsoft Windows Server Hyper-V 2016. This validated solution uses 10-Gigabit Ethernet (10GbE) technology. Guidance is also provided on how to scale compute capacity by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.

Figure 2) FlexPod Express with Microsoft Windows Server Hyper-V 2016 10GbE architecture.



2.4 Use Case Summary

The FlexPod Express solution can be applied to several use cases, including the following:

- Remote offices or branch offices (ROBOs)
- Small and midsize businesses
- Environments that require a dedicated or cost-effective solution

FlexPod Express is best suited for virtualization and mixed workloads.

3 Technology Requirements

A FlexPod Express system requires a combination of hardware and software components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

3.1 Hardware Requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

Table 1 lists the hardware components that are required for all FlexPod Express configurations.

Table 1) Hardware requirements for the base configuration.

Hardware	Quantity
FAS2650 two-node cluster	1
Cisco C220 M4 server	2
Cisco Nexus 31108 switch	2
Cisco UCS virtual interface card (VIC) 1227 for the C220 M4 server	2

Table 2 lists the hardware that is required in addition to the base configuration for implementing 10GbE.

Table 2) Hardware for scaling the solution by using two hypervisor nodes.

Hardware	Quantity
Cisco UCS C220 M4 server	2
Cisco VIC 1227	2

3.2 Software Requirements

Table 3 lists the software components that are required to implement the architectures of the FlexPod Express solutions.

Table 3) Software requirements for the base FlexPod Express implementation.

Software	Version	Details
Cisco Integrated Management Controller (CIMC)	2.0 (13h)	For Cisco UCS C220 M4 rack servers
Cisco VIC Driver	4.0.0.2 Ethernet	For VIC 1227 interface cards
Cisco NX-OS	7.0(3)I5(2)	For Cisco Nexus 31108 switches
NetApp ONTAP	9.1	For FAS2650 controllers

Table 4 lists the software that is required for all Microsoft Windows Server Hyper-V 2016 implementations on FlexPod Express. You need ISO files for the Microsoft products listed in the table. The NetApp products can be downloaded from the [NetApp Support site](#).

Table 4) Software requirements for a Microsoft Windows Server Hyper-V 2016 implementation.

Software	Version
Microsoft Windows Server Hyper-V	2016
Microsoft SQL 2016	2016
Microsoft System Center Virtual Machine Manager	2016
NetApp Windows Unified Host Utilities	7.0
NetApp SMI-S Provider	5.2.4

4 FlexPod Express Cabling Information

The reference validation provided in this document is cabled as shown in Figure 3 and Table 5 through Table 8.

Figure 3) Reference validation cabling.




Table 5) Cabling information for Cisco Nexus switch 31108 A.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus switch 31108 A	Eth1/1	NetApp FAS2650 storage controller A	e0c
	Eth1/2	NetApp FAS2650 storage controller B	e0c
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM1
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM1
	Eth1/25	Cisco Nexus switch 31108 B	Eth1/25
	Eth1/26	Cisco Nexus switch 31108 B	Eth1/26
	Eth1/33	NetApp FAS2650 storage controller A	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server A	CIMC

Table 6) Cabling information for Cisco Nexus switch 31108 B.

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus switch 31108 B	Eth1/1	NetApp FAS2650 storage controller A	e0d
	Eth1/2	NetApp FAS2650 storage controller B	e0d
	Eth1/3	Cisco UCS C220 C-Series standalone server A	MLOM2
	Eth1/4	Cisco UCS C220 C-Series standalone server B	MLOM2
	Eth1/25	Cisco Nexus switch 31108 A	Eth1/25
	Eth1/26	Cisco Nexus switch 31108 A	Eth1/26
	Eth1/33	NetApp FAS2650 storage controller B	e0M
	Eth1/34	Cisco UCS C220 C-Series standalone server B	CIMC

Table 7) Cabling information for NetApp FAS2650 storage controller A.

Local Device	Local Port	Remote Device	Remote Port
NetApp FAS2650 storage controller A	e0a	NetApp FAS2650 storage controller B	e0a
	e0b	NetApp FAS2650 storage controller B	e0b
	e0c	Cisco Nexus switch 31108 A	Eth1/1
	e0d	Cisco Nexus switch 31108 B	Eth1/1
	e0M	Cisco Nexus switch 31108 A	Eth1/33

Table 8) Cabling information for NetApp FAS2650 storage controller B.

Local Device	Local Port	Remote Device	Remote Port
NetApp FAS2650 storage controller B	e0a	NetApp FAS2650 storage controller A	e0a
	e0b	NetApp FAS2650 storage controller A	e0b
	e0c	Cisco Nexus switch 31108 A	Eth1/2
	e0d	Cisco Nexus switch 31108 B	Eth1/2
	e0M	Cisco Nexus switch 31108 B	Eth1/33

5 Deployment Procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. Table 9 describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.

Note: If you use separate in-band and out-of-band management VLANs, you must create a layer 3 route between them. For this validation, a common management VLAN was used.

Table 9) Required VLANs.

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Management VLAN	VLAN for management interfaces	3437
Native VLAN	VLAN to which untagged frames are assigned	2
Cluster shared volumes (CSV) VLAN	VLAN for CSV	3438
Live migration VLAN	VLAN designated for the movement of virtual machines from one physical host to another	3441
Virtual machine traffic VLAN	VLAN for virtual machine application traffic	3442
iSCSI-A-VLAN	VLAN for iSCSI traffic on fabric A	3439
iSCSI-B-VLAN	VLAN for iSCSI traffic on fabric B	3440

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where xxxx is the purpose of the VLAN (such as iSCSI-A).

Table 10 lists the Hyper-V virtual machines created.

Table 10) Hyper-V virtual machines created.

Virtual Machine Description	Host Name
Microsoft System Center Virtual Machine Manager (SCVMM)	
NetApp SMI-S Provider	

5.1 Cisco Nexus 31108 Deployment Procedure

The following section details the Cisco Nexus 31108 switch configuration used in a FlexPod Express environment.

Initial Setup of Cisco Nexus 31108 Switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.

Note: This procedure assumes that you are using a Cisco Nexus 31108 running NX-OS software release 7.0(3)I5(2).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 31108 switches can be connected to an existing management network, or the

mgmt0 interfaces of the 31108 switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.

Note: In this deployment guide, the FlexPod Express Cisco Nexus 31108 switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108 switches, power on the switch and follow the on-screen prompts as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

```
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

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

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

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

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

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

Enter the switch name : 31108-B

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

  Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>

  Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>

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

  IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>

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

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

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

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

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

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

  NTP server IPv4 address : <<var_ntp_ip>>

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

Configure default switchport interface state (shut/noshut) [noshut]: <enter>

Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: <enter>
```

4. You then see a summary of your configuration, and you are asked if you would like to edit it. If your configuration is correct, enter n.

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

5. You are then asked if you would like to use this configuration and save it. If so, enter `y`.

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

6. Repeat this procedure for Cisco Nexus switch B.

Enable Advanced Features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options.

Note: The `interface-vlan` feature is required only if you use the back-to-back `mgmt0` option described throughout this document. This feature allows you to assign an IP address to the interface VLAN (switch virtual interface), which enables in-band management communication to the switch (such as through SSH).

1. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode using the command (`config t`) and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```

Note: The default PortChannel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the PortChannel. You can achieve better distribution across the members of the PortChannel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

2. From configuration mode (`config t`), enter the following commands to set the global PortChannel load-balancing configuration on Cisco Nexus switch A and switch B:

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

Perform Global Spanning-Tree Configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (`config t`), type the following commands to configure the default spanning-tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

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

Define VLANs

Before configuring individual ports with different VLANs, the layer 2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (`config t`), type the following commands to define and give descriptions to the layer 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<CSV_vlan_id>>
  name CSV-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<LiveMigration_vlan_id>>
  name LiveMigration-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure Access and Management Port Descriptions

As is the case with assigning names to the layer 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (`config t`) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus Switch A

```
int eth1/1
  description FAS2650-A e0c
int eth1/2
  description FAS2650-B e0c
int eth1/3
  description UCS-Server-A: VIC Port 1
int eth1/4
  description UCS-Server-B: VIC Port 1
int eth1/25
  description vPC peer-link 31108-B 1/25
int eth1/26
  description vPC peer-link 31108-B 1/26
int eth1/33
  description FAS2650-A e0M
int eth1/34
  description UCS Server A: CIMC
```

Cisco Nexus Switch B

```
int eth1/1
  description FAS2650-A e0d
int eth1/2
  description FAS2650-B e0d
int eth1/3
  description UCS-Server-A: VIC Port 2
int eth1/4
  description UCS-Server-B: VIC Port 2
int eth1/25
  description vPC peer-link 31108-A 1/25
int eth1/26
  description vPC peer-link 31108-A 1/26
int eth1/33
  description FAS2650-B e0M
int eth1/34
```

Configure Server and Storage Management Interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (`config t`), enter the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus Switch A

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Cisco Nexus Switch B

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Perform Virtual PortChannel Global Configuration

A virtual PortChannel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single PortChannel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer 2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a PortChannel across two upstream devices
- Eliminating Spanning-Tree Protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back `mgmt0` configuration, use the addresses defined on the interfaces and verify that they can communicate by using the `ping <<switch_A/B_mgmt0_ip_addr>>vrf management` command.

From configuration mode (`config t`), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus Switch A

```
vpc domain 1
  1-switch
  role priority 10
```

```

    peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source <<switch_A_mgmt0_ip_addr>> vrf
management
    peer-gateway
    auto-recovery
    ip arp synchronize

int eth1/25-26
    channel-group 10 mode active
int Po10
    description vPC peer-link
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<CSV_vlan_id>>,<<LiveMigration_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>>, <<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
    spanning-tree port type network
    vpc peer-link
    no shut
exit
copy run start

```

Cisco Nexus Switch B

```

vpc domain 1
    peer-switch
    role priority 20
    peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source <<switch_B_mgmt0_ip_addr>> vrf
management
    peer-gateway
    auto-recovery
    ip arp synchronize

int eth1/25-26
    channel-group 10 mode active
int Po10
    description vPC peer-link
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<CSV_vlan_id>>,<<LiveMigration_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>>, <<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
    spanning-tree port type network
    vpc peer-link
no shut
exit
copy run start

```

Configure Storage PortChannels

The NetApp storage controllers allow an active-active connection to the network using the Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach enables you to have active-active connections from the storage to separate physical switches. Each controller has two links to each of the switches. However, all four links are part of the same vPC and interface group (IFGRP).

From configuration mode (`config t`), run the following commands on each of the switches to configure the individual interfaces and the resulting PortChannel configuration for the ports connected to the NetApp FAS controller.

Run the following commands on switch A and switch B to configure the PortChannels for storage controller A.

```

int eth1/1
    channel-group 11 mode active
int Po11
    description vPC to Controller-A
    switchport

```

```

switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
spanning-tree port type edge trunk
mtu 9216
vpc 11
no shut

```

Run the following commands on switch A and switch B to configure the PortChannels for storage controller B.

```

int eth1/2
 channel-group 12 mode active
int Po12
 description vPC to Controller-B
 switchport
 switchport mode trunk
 switchport trunk native vlan <<native_vlan_id>>
 switchport trunk allowed vlan <<mgmt_vlan_id>>, <<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
 spanning-tree port type edge trunk
 mtu 9216
 vpc 12
 no shut
exit
copy run start

```

Note: Jumbo frames should be configured throughout the network to enable any applications and operating systems to transmit these larger frames without fragmentation. Both the endpoints and all the interfaces between the endpoints (layer 2 and layer 3) must support and be configured for jumbo frames to prevent performance problems caused by fragmenting frames.

Configure Server Connections

The Cisco UCS servers have a two-port virtual interface card, VIC1227, that is used for data traffic and booting the Windows operating system using iSCSI. These interfaces are configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

From configuration mode (`config t`), run the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B Configuration

```

int eth1/3-4
 switchport mode trunk
 switchport trunk native vlan <<native_vlan_id>>
 switchport trunk allowed vlan
<<iSCSI_A_vlan_id>>,<<CSV_vlan_id>>,<<LiveMigration_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_i
d>>
 spanning-tree port type edge trunk
 mtu9216
 no shut
exit
copy run start

```

Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B Configuration

```

int eth1/3-4
 switchport mode trunk
 switchport trunk native vlan <<native_vlan_id>>
 switchport trunk allowed vlan
<<iSCSI_B_vlan_id>>,<<CSV_vlan_id>>,<<LiveMigration_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_i
d>>
 spanning-tree port type edge trunk
 mtu 9216
 no shut

```



```
exit
copy run start
```

Note: Jumbo frames should be configured throughout the network to enable any applications and operating systems to transmit these larger frames without fragmentation. Both the endpoints and all the interfaces between the endpoints (layer 2 and layer 3) must support and be configured for jumbo frames to prevent performance problems by fragmenting frames.

Note: To scale the solution by adding additional Cisco UCS servers, run the previous commands with the switch ports that the newly added servers have been plugged into on switches A and B.

Uplink into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 31108 switches included in the FlexPod environment into the infrastructure. The uplinks may be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is completed.

5.2 NetApp FAS Storage Deployment Procedure: Part 1

This section describes the first part of the NetApp FAS storage deployment procedure.

NetApp Storage Controller FAS26xx Series Installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you want to use are supported with the version of ONTAP that you plan to install by using the [HWU application](#) at the [NetApp Support](#) site:

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

NetApp FAS26XX Series Prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections:

- Electrical requirements
- Supported power cords
- Onboard ports and cables

Storage Controllers

Follow the physical installation procedures for the controllers in the [FAS26xx documentation](#) available at the [NetApp Support](#) site.

NetApp ONTAP 9.1

Configuration Worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.1 Software Setup Guide](#) at the [NetApp Support](#) site.

Note: This system is set up in a two-node switchless cluster configuration.

Table 11) ONTAP 9.1 installation and configuration information.

Cluster Detail	Cluster Detail Value
Cluster node A IP address	<<var_nodeA_mgmt_ip>>
Cluster node A netmask	<<var_nodeA_mgmt_mask>>
Cluster node A gateway	<<var_nodeA_mgmt_gateway>>
Cluster node A name	<<var_nodeA>>
Cluster node B IP address	<<var_nodeB_mgmt_ip>>
Cluster node B netmask	<<var_nodeB_mgmt_mask>>
Cluster node B gateway	<<var_nodeB_mgmt_gateway>>
Cluster node B name	<<var_nodeB>>
ONTAP 9.1 URL	<<var_url_boot_software>>
Name for cluster	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster B gateway	<<var_clustermgmt_gateway>>
Cluster B netmask	<<var_clustermgmt_mask>>
Domain name	<<var_domain_name>>
DNS server IP (you can enter more than one)	<var_dns_server_ip
NTP server IP (you can enter more than one)	<<var_ntp_server_ip>>

Configure Node A

To configure node A, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

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

- To install new software, select option 7.

7

- Enter `y` to perform an upgrade.

y

- Select `e0M` for the network port you want to use for the download.

e0M

- Enter `y` to reboot now.

y

- Enter the IP address, netmask, and default gateway for `e0M` in their respective places.

<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>

- Enter the URL where the software can be found.

Note: This web server must be pingable.

<<var_url_boot_software>>

- Press Enter for the user name, indicating no user name.

- Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

y

- Enter `y` to reboot the node.

y

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

- Press Ctrl-C when you see this message:

Press Ctrl-C for Boot Menu

- Select option 4 for Clean Configuration and Initialize All Disks.

4

- Enter `y` to zero disks, reset config, and install a new file system.

y

- Enter `y` to erase all the data on the disks.

y

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

- While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

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

4. To install new software, select option 7.

```
7
```

5. Enter `y` to perform an upgrade.

```
y
```

6. Select e0M for the network port you want to use for the download.

```
e0M
```

7. Enter `y` to reboot now.

```
y
```

8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

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

Note: This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.

11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

12. Enter `y` to reboot the node.

```
y
```

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

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

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

14. Select option 4 for Clean Configuration and Initialize All Disks.

```
4
```

15. Enter `y` to zero disks, reset config, and install a new file system.

```
y
```

16. Enter `y` to erase all the data on the disks.

```
y
```

Note: The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

5.3 Continuation of Node A Configuration and Cluster Configuration

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

Note: The node and cluster setup procedure has changed slightly in ONTAP 9.1. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.

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

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

This system will send event messages and periodic reports to NetApp Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]:
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>
Enter the node management interface default gateway: <<var_nodeA_mgmt_gateway>>
A node management interface on port e0M with IP address <<var_nodeA_mgmt_ip>> has been created.

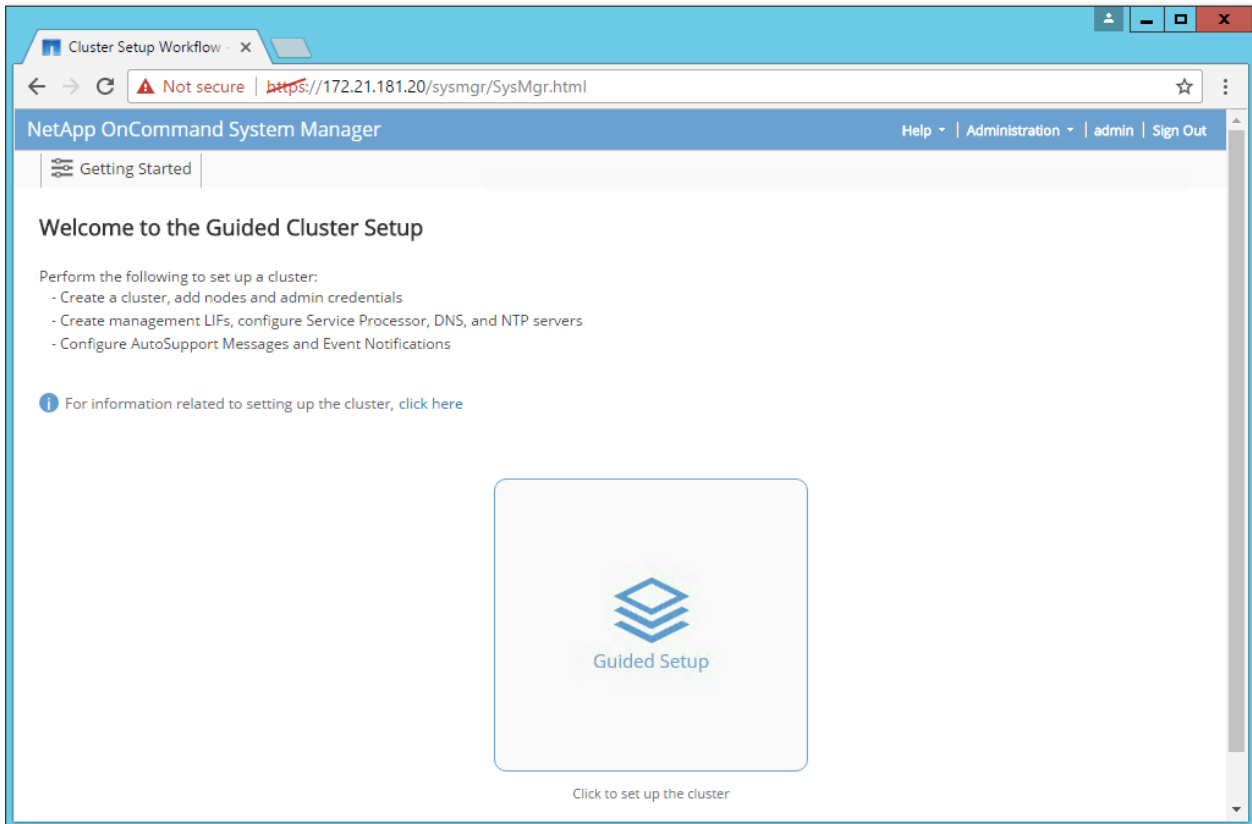
Use your web browser to complete cluster setup by accessing
https://<<var_nodeA_mgmt_ip>>

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

2. Navigate to the IP address of the node's management interface.

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


3. Click Guided Setup to configure the cluster.



4. Enter `<<var_clustertype>>` for the cluster name and `<<var_nodeA>>` and `<<var_nodeB>>` for each of the nodes you are configuring. Enter the password you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



Cluster Name

Nodes

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



Cluster Configuration: Switched Cluster Switchless Cluster

? Username admin

Password

Confirm Password


Cluster Base License (Optional)

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

Feature Licenses (Optional)

i Cluster Base License is mandatory to add Feature Licenses.


5. You can also enter feature licenses.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.



7. Next, you are guided through the process of configuring the network. The network screen looks like this.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



1 Network (Management)

IP Addresses (IPv4) required

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

IP Address Range

It is recommended not to manually modify the Cluster Management, Node Management, and Service Processor Management IP addresses. If you are enabling the IP Address Range, entering the Gateway address is mandatory. If you do not want to enter the Gateway address, ensure that the IP Address Range is disabled.

Starting address to Ending address Netmask Gateway

Apply Sequentially

Cluster Management IP Address Port

Node Management

FAS2650_A e0M

FAS2650_B e0M

Service Processor Management Default values have been detected for the Service Processor. Override the default values (Gateway is mandatory)

FAS2650_A

FAS2650_B

2 DNS Details

DNS Domain Names

DNS Server IP Address

3 NTP Details

Primary NTP Server


Alternative NTP Server (Optional)

8. Begin with the Network Management section.

? Network (Management)

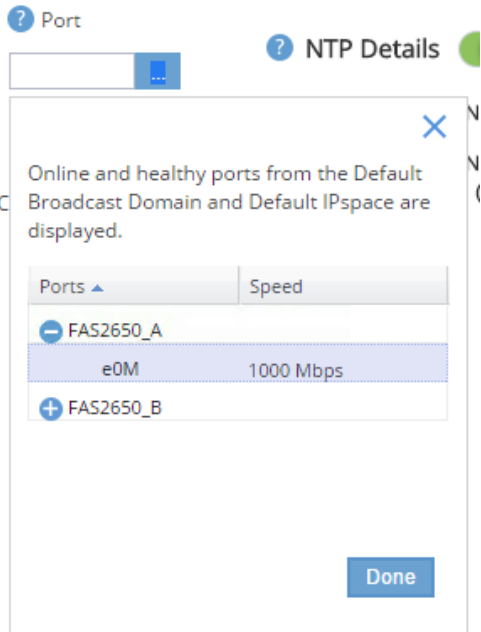
IP Addresses (IPv4) required

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

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

	IP Address	Netmask	Gateway (Optional)	? Port
Cluster Management	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text" value="..."/>
Node Management	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
FAS2650_A	<input type="text" value="172.21.181.20"/>	<input type="text" value="e0M"/>	<input type="text" value="..."/>	
FAS2650_B	<input type="text"/>	<input type="text" value="e0M"/>	<input type="text" value="..."/>	
Service Processor Management	Default values have been detected for the Service Processor.			
	<input type="checkbox"/> Override the default values (Gateway is mandatory)			
	<input checked="" type="checkbox"/> Retain Netmask and Gateway configuration of the Cluster Management.			
FAS2650_A	<input type="text" value="172.21.181.11"/>			
FAS2650_B	<input type="text" value="172.21.181.12"/>			

- Deselect IP Address Range.
- Enter `<<var_clustermgmt_ip>>` in the Cluster Management IP Address field, `<<var_clustermgmt_mask>>` in the Netmask field, and `<<var_clustermgmt_gateway>>` in the Gateway field. Use the ... selector in the port field to select e0M of node A.



c. The node management IP for node A is already populated. Enter <<var_nodeA_mgmt_ip>> for node B.

9. Configure the DNS details and NTP details.

? DNS Details

DNS Domain Names

DNS Server IP Address

? NTP Details

Primary NTP Server

Alternative NTP Server (Optional)

a. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

Note: You can enter multiple DNS server IP addresses.

b. Enter <<var_ntp_server_ip>> in the Primary NTP Server field.

Note: You can also enter an alternate NTP server.

10. Click Submit.

11. Now configure the support information.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:

<input checked="" type="checkbox"/>	Email	SMTP Mail Host <input type="text"/>	Email Addresses <input type="text" value="Separate email addresses with a comma..."/>
<input type="checkbox"/>	SNMP	SNMP Trap Host <input type="text"/>	
<input type="checkbox"/>	Syslog	Syslog Server <input type="text"/>	

Submit

- a. If your environment requires a proxy to access AutoSupport®, enter it next to the Proxy URL field.
- b. Enter the SMTP mail host and e-mail address for event notifications.

Note: You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

12. The following message indicates that the cluster configuration has completed. Click Manage Your Cluster to begin to configure the storage.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



[Click here to view the summary](#)

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

[Manage your cluster](#)

Continuation of Storage Cluster Configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero All Spare Disks

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

```
disk zerospares
```

Set Onboard UTA2 Ports Personality

1. Verify the current mode and the current type of the ports by using the `ucadmin show` command.

```
FAS2650::> ucadmin show
```

Node	Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
FAS2650_A	0c	fc	target	-	-	online
FAS2650_A	0d	fc	target	-	-	online
FAS2650_A	0e	fc	target	-	-	online
FAS2650_A	0f	fc	target	-	-	online
FAS2650_B	0c	fc	target	-	-	online
FAS2650_B	0d	fc	target	-	-	online
FAS2650_B	0e	fc	target	-	-	online
FAS2650_B	0f	fc	target	-	-	online

8 entries were displayed.

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If this is not the case, change the port personality by using the following command:

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

Note: The ports must be offline to run the previous command. To take a port offline, run the following command:

```
network fcp adapter modify -node <home node of the port> -adapter <port name> -state down
```

Note: If you change the port personality, you must reboot each node for the change to take effect.

Rename Management Logical Interfaces

To rename the management logical interfaces (LIFs), complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver FAS2650
```

2. Rename the cluster management LIF.

```
network interface rename -vserver FAS2650 -lif cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver FAS2650 -lif cluster_setup_node_mgmt_lif_FAS2650_B_1 -newname FAS2650-02_mgmt1
```

Set Auto-Revert on Cluster Management

Set the `auto-revert` parameter on the cluster management interface.

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

Set Up Service Processor Network Interface

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

```
system service-processor network modify -node <<var_nodeA>> -address-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>> -netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>
```

```
system service-processor network modify -node <<var_nodeB>> -address-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>> -netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```

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

Enable Storage Failover in ONTAP

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

1. Verify the status of storage failover.

```
storage failover show
```

Note: Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

Note: Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.

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

```
cluster ha show
```

- Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

- Enable HA mode only for the two-node cluster.

Note: Do not run this command for clusters with more than two nodes because it causes problems with failover.

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

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

```
storage failover hwassist show
```

Note: The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node <<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node <<var_nodeB>>
```

Create Jumbo Frame MTU Broadcast Domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove Data Ports from Default Broadcast Domain

The 10GbE data ports are used for iSCSI traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports <<var_nodeA>>:e0c,
<<var_nodeA>>:e0d, <<var_nodeA>>:e0e, <<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable Flow Control on UTA2 Ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following commands:

```
net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
```

```
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

Configure IFGRP LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Make sure the switch is properly configured.

From the cluster prompt, run the following commands.

```
ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d

ifgrp create -node << var_nodeB>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d
```

Configure Jumbo Frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (which usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```
FAS2650::> network port modify -node FAS2650_A -port a0a -mtu 9000

Warning: This command will cause a several second interruption of service on
this network port.
Do you want to continue? {y|n}: y

FAS2650::> network port modify -node FAS2650_B -port a0a -mtu 9000

Warning: This command will cause a several second interruption of service on
this network port.
Do you want to continue? {y|n}: y
```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create iSCSI VLAN ports and add them to the data broadcast domain.

```
network port vlan create -node <<var_nodeA>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_nodeA>>:a0a-
<<var_iscsi_vlan_A_id>>,<<var_nodeB>>:a0a-<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_nodeA>>:a0a-
<<var_iscsi_vlan_B_id>>,<<var_nodeB>>:a0a-<<var_iscsi_vlan_B_id>>
```

2. Create MGMT-VLAN ports.

```
network port vlan create -node <<var_nodeA>> -vlan-name a0a-<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-<<mgmt_vlan_id>>
```

Create Aggregates in ONTAP

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

To create new aggregates, run the following commands:

```
aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount <<var_num_disks>>
```

Note: Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Note: Start with five disks; you can add disks to an aggregate when additional storage is required.

Note: The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1_nodeA` is online.

Configure Time Zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```

Note: For example, in the eastern United States, the time zone is `America/New_York`. After you begin typing the time zone name, press the Tab key to see available options.

Configure Simple Network Management Protocol in ONTAP

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

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

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

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

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

```
snmp community add ro <<var_snmp_community>>
```

Note: Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select `des` as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

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

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a Storage Virtual Machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the Infra-SVM aggregate list.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc, nfs
```

Note: Commands are prefaced by `vserver` in the command line because storage virtual machines were previously called vservers.

Create iSCSI Service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM  
iscsi show
```

Create Load-Sharing Mirror of SVM Root Volume in ONTAP

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

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate aggr1_nodeA -size 1GB -type DP  
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate aggr1_nodeB -size 1GB -type DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m01 -type LS -schedule 15min  
snapmirror create -source-path Infra-SVM:rootvol -destination-path Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

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

Configure HTTPS Access in ONTAP

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

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

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

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority. To delete the default certificates, run the following commands:

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

```
security certificate delete [TAB] ...
Example: security certificate delete -vserver Infra-SVM -common-name Infra-SVM -ca Infra-SVM -
type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Use tab completion to aid in completing these commands.

```
security certificate create [TAB] ...
Example: security certificate create -common-name infra-svm.netapp.com -type server -size 2048 -
country US -state "North Carolina" -locality "RTP" -organization "NetApp" -unit "FlexPod" -email-
addr "abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-SVM
```

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.

6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use tab completion.

```
security ssl modify [TAB] ...
Example: security ssl modify -vserver Infra-SVM -server-enabled true -client-enabled false -ca
infra-svm.netapp.com -serial 55243646 -common-name infra-svm.netapp.com
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web service requests to be
interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http -vserver <<var_clustername>>
```

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

8. Revert to the admin privilege level and create a setup that allows SVM to be available through the web.

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

Create FlexVol Volumes in NetApp ONTAP

The following information is required to create a FlexVol® volume: the volume's name, size, and the aggregate on which it will exist. To create the NetApp FlexVol volumes, complete the following step:

1. Create a provisioning volume to store virtual machines, a witness volume for the Windows failover cluster, and a boot volume.

```

volume create -vserver Infra-SVM -volume hyperv_boot -aggregate aggr1_nodeA -size 500GB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume provisioning -aggregate aggr1_nodeA -size 500GB -state
online -policy default -space-guarantee none -percent-snapshot-space 0

volume create -vserver Infra-SVM -volume witness -aggregate aggr1_nodeA -size 10GB -state online
-policy default -space-guarantee none -percent-snapshot-space 0

```

Create LUNs in NetApp ONTAP

To create LUNs, complete the following steps:

1. Create two boot LUNs.

```

lun create -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-A -size 200GB -ostype
hyper_v -space-reserve disabled
lun create -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-B -size 200GB -ostype
hyper_v -space-reserve disabled

```

Note: When adding an additional Cisco UCS C-Series server, an additional boot LUN must be created.

2. Create provisioning and witness LUNs.

```

lun create -vserver Infra-SVM -volume provisioning -lun Provisioning_LUN -size 500GB -ostype
hyper_v -space-reserve disabled

lun create -vserver Infra-SVM -volume witness -lun witness_LUN -size 5GB -ostype hyper_v -space-
reserve disabled

```

Enable Deduplication in ONTAP

To enable deduplication on appropriate volumes, run the following commands:

```

volume efficiency on -vserver Infra-SVM -volume witness
volume efficiency on -vserver Infra-SVM -volume hyperv_boot
volume efficiency on -vserver Infra-SVM -volume provisioning

```

Note: To provision additional LUNs, follow the steps outlined earlier in the Create FlexVol Volumes in NetApp ONTAP, Create LUNs in NetApp ONTAP, and Enable Deduplication in NetApp ONTAP sections.

Create Boot LUNs in ONTAP

To create two boot LUNs, run the following commands:

```

lun create -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-A -size 15GB -ostype hyperv
-space-reserve disabled
lun create -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-B -size 15GB -ostype hyperv
-space-reserve disabled

```

Note: When adding an additional Cisco UCS C-Series server, an additional boot LUN must be created.

Create iSCSI LIFs in ONTAP

You need the information in Table 12 to complete this configuration step.

Table 12) Information required for iSCSI configuration.

Detail	Detail Value
Storage node A iSCSI LIF01A	<<var_nodeA_iscsi_lif01a_ip>>

Detail	Detail Value
Storage node A iSCSI LIF01A network mask	<<var_nodeA_iscsi_lif01a_mask>>
Storage node A iSCSI LIF01B	<<var_nodeA_iscsi_lif01b_ip>>
Storage node A iSCSI LIF01B network mask	<<var_nodeA_iscsi_lif01b_mask>>
Storage node B iSCSI LIF01A	<<var_nodeB_iscsi_lif01a_ip>>
Storage node B iSCSI LIF01A network mask	<<var_nodeB_iscsi_lif01a_mask>>
Storage node B iSCSI LIF01B	<<var_nodeB_iscsi_lif01b_ip>>
Storage node B iSCSI LIF01B network mask	<<var_nodeB_iscsi_lif01b_mask>>

1. Create four iSCSI LIFs, two on each node.

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -
home-node <<var_nodeA>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_nodeA_iscsi_lif01a_ip>> -netmask <<var_nodeA_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -
home-node <<var_nodeA>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
<<var_nodeA_iscsi_lif01b_ip>> -netmask <<var_nodeA_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -
home-node <<var_nodeB>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_nodeB_iscsi_lif01a_ip>> -netmask <<var_nodeB_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -
home-node <<var_nodeB>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
<<var_nodeB_iscsi_lif01b_ip>> -netmask <<var_nodeB_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface show
```

Add Infrastructure SVM Administrator

You need the information in Table 13 to complete these configuration steps.

Table 13) Information required for SVM administrator addition.

Detail	Detail Value
Vsmgmt IP	<<var_svm_mgmt_ip>>
Vsmgmt network mask	<<var_svm_mgmt_mask>>
Vsmgmt default gateway	<<var_svm_mgmt_gateway>>

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data -data-protocol none -home-node
<<var_nodeB>> -home-port e0M -address <<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -
status-admin up -failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-revert true
```

Note: The SVM management IP here should be in the same subnet as the storage cluster management IP.

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

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway <<var_svm_mgmt_gateway>>
network route show
```

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

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

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

5.4 Cisco UCS C-Series Rack Server Deployment Procedure

The following section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in the FlexPod Express configuration.

Perform Initial Cisco UCS C-Series Standalone Server Setup for Cisco Integrated Management Server

Complete these steps for the initial setup of the CIMC interface for Cisco UCS C-Series standalone servers.

You need the information in Table 14 to configure CIMC for each Cisco UCS C-Series standalone server.

Table 14) Information required for CIMC configuration.

Detail	Detail Value
CIMC IP address	<<cimc_ip>>
CIMC subnet mask	<<cimc_netmask
CIMC default gateway	<<cimc_gateway>>

Note: The CIMC version used in this validation is CIMC 2.0(13h).

All Servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.
2. Power on the server and press F8 when prompted to enter the CIMC configuration.



```
Press <F2> Setup, <F6> Boot Menu, <F7> Diagnostics, <F8>Cisco IMC Configuration,  
<F12> Network Boot
```

```
Bios Version : C220M4.2.0.13g.0.1113162259  
Platform ID : C220M4
```

```
Cisco IMC IPv4 Address : 10.61.186.43  
Cisco IMC MAC Address : 08:96:AD:AC:41:52
```

```
Processor(s) Intel(R) Xeon(R) CPU E5-2650 v4 @ 2.20GHz  
Total Memory = 128 GB Effective Memory = 128 GB  
Memory Operating Speed 2400 Mhz  
Entering CIMC Configuration Utility...
```

3. In the CIMC configuration utility, set the following options:

- Network interface card (NIC) mode:
 - Dedicated
- IP (Basic):
 - IPV4:
 - DHCP enabled:
 - CIMC IP: <<cimc_ip>>
 - Prefix/Subnet: <<cimc_netmask>>
 - Gateway: <<cimc_gateway>>
- VLAN (Advanced): Leave cleared to disable VLAN tagging.
 - NIC redundancy
 - None:

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode                               NIC redundancy
Dedicated:      [X]                    None:          [X]
Shared LDM:     [ ]                    Active-standby: [ ]
Cisco Card:
  Riser1:       [ ]                    Active-active:  [ ]
  Riser2:       [ ]                    VLAN (Advanced)
  MLom:         [ ]                    VLAN enabled:   [ ]
  Shared LDM Ext: [ ]                    VLAN ID:        1
                                           Priority:        0
IP (Basic)
IPV4:           [X]                    IPV6:          [ ]
DHCP enabled   [ ]
CIMC IP:       10.61.186.43
Prefix/Subnet: 255.255.255.0
Gateway:       10.61.186.1
Pref DNS Server: 0.0.0.0

*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings

```

4. Press F1 to see additional settings.
 - Common properties:
 - Host name: <<hyperv_host_name>>
 - Dynamic DNS: []
 - Factory defaults: Leave cleared.
 - Default user (basic):
 - Default password: <<admin_password>>
 - Reenter password: <<admin_password>>
 - Port properties: Use default values.
 - Port profiles: Leave cleared.

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
Common Properties
  Hostname:      c1e-c220m4-f0246
  Dynamic DNS:   [ ]
  DDNS Domain:
FactoryDefaults
  Factory Default: [ ]
Default User(Basic)
  Default password:
  Reenter password:
Port Properties
  Auto Negotiation: [X]
                Admin Mode      Operation Mode
  Speed [1000/100/10Mbps]:      Auto          1000
  Duplex mode [half/full]:      Auto          full
Port Profiles
  Reset: [ ]
  Name:
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>PreviousPageettings

```

5. Press F10 to save the CIMC interface configuration.
6. After the configuration is saved, press Esc to exit.

Configure Cisco UCS C-Series Servers iSCSI Boot

In this FlexPod Express configuration, the VIC1227 is used for iSCSI boot.

You need the information in Table 15 to configure each Hyper-V host.

Table 15) Information required for iSCSI boot configuration.

Detail	Detail Value
Hyper-V host initiator A name	<<var_ucs_initiator_name_A>>
Hyper-Vi host iSCSI-A boot IP	<<var_hyperv_host_iscsiA_ip>>
Hyper-V host iSCSI-A network mask	<<var_hyperv_host_iscsiA_mask>>
Hyper-V host iSCSI A default gateway	<<var_hyperv_host_iscsiA_gateway>>
Hyper-V host initiator B name	<<var_ucs_initiator_name_B>>
Hyper-V host iSCSI-B boot IP	<<var_hyperv_host_iscsiB_ip>>
Hyper-V host iSCSI-B network mask	<<var_hyperv_host_iscsiB_mask>>
Hyper-V host iSCSI-B gateway	<<var_hyperv_host_iscsiB_gateway>>
IP address iscsi_lif01a	
IP address iscsi_lif02a	

Detail	Detail Value
IP address <i>iscsi_lif01b</i>	
IP address <i>iscsi_lif02b</i>	
Infra_SVM IQN	

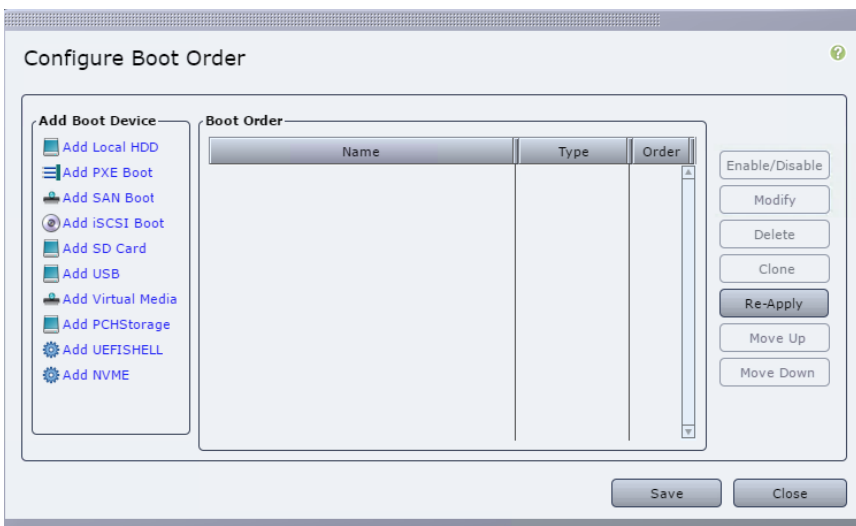
Note: Italicized font indicates variables that are unique for each Hyper-V host.

Boot Order Configuration

1. From the Cisco IMC interface browser window, click the Server tab and select BIOS.
2. Select Configure Boot Order and click OK.

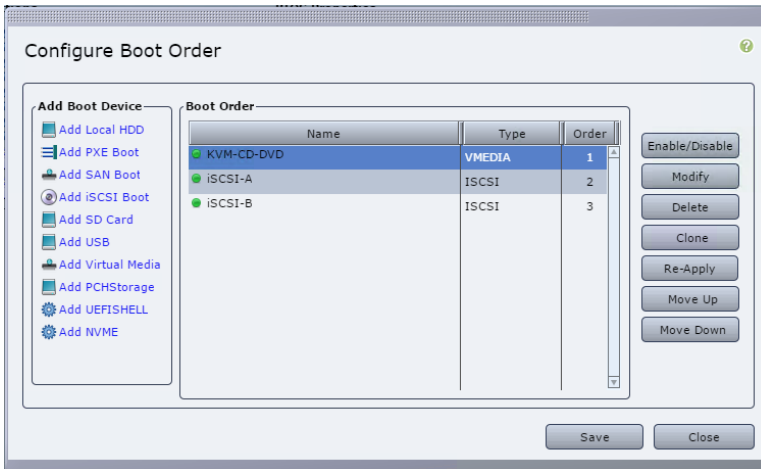


3. Click the following devices listed under Add Boot Device to configure them.



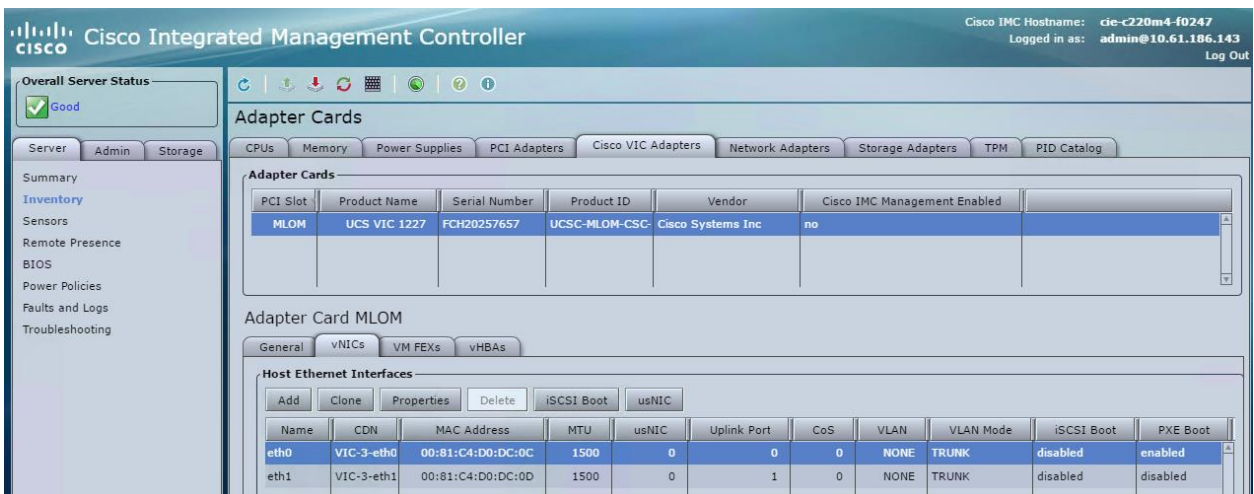
- Add Virtual Media:
 - Name: KVM-CD-DVD
 - Sub Type: KVM MAPPED DVD
 - State: Enabled
- Add iSCSI Boot:

- Name: iSCSI-A
 - State: Enabled
 - Order: 2
 - Slot: MLOM
 - Port: 0
4. Click Add Device:
 - Name: iSCSI-B
 - State: Enabled
 - Order: 3
 - Slot: MLOM
 - Port: 1
 5. Click Add Device.
 6. Click Save. Click Close.



Configure vNICs for Hyper-V

1. From the Cisco IMC interface browser window, click Inventory, then click Cisco VIC Adapters in the right pane.
2. Under Adapter Cards, select UCS VIC 1227, then select the vNICs underneath.



3. Select eth0 and click Properties.
4. Set the MTU to 9000. Click Save Changes.

vNIC Properties

General

Name: **eth0**

MTU: (1500 - 9000)

Uplink Port:

MAC Address: AUTO

Class of Service:

Trust Host CoS:

PCI Order: ANY (0 - 17)

Default VLAN: NONE (1 - 4094)

VLAN Mode:

Rate Limit: OFF (1 - 10000 Mbps)

Enable PXE Boot:

Channel Number: (1 - 1000) **N/A**

5. Repeat steps 3 and 4 for eth1.

Note: This procedure must be repeated for each initial Cisco UCS server node, as well as for each additional Cisco UCS server node that is added to the environment.

Configure Cisco VIC1227 for iSCSI Boot: Part 1

Initially one working path is configured between the server and the iSCSI LUN. The second path is configured in a later step.

These steps provide details for configuring the Cisco VIC1227 for iSCSI boot.

Create First iSCSI vNIC for VLAN A

1. Click Add to create a new vNIC.
2. In the Add vNIC window, complete the following settings:
 - Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <<var_iscsi_vlan_a>>
 - VLAN Mode: TRUNK
 - Enable PXE Boot: Select the checkbox

Add vNIC

General

Name:

MTU: (1500 - 9000)

Uplink Port:

MAC Address: AUTO

Class of Service:

Trust Host CoS:

PCI Order: ANY (0 - 17)

Default VLAN: NONE (1 - 4094)

VLAN Mode:

Rate Limit: OFF (1 - 10000 Mbps)

Enable PXE Boot:

Channel Number: (1 - 1000)

3. Click Add vNIC. Click OK.
4. Select the vNIC “iSCSI-vNIC-A” and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.

Adapter Card MLOM

General vNICs VM FEXs vHBAs

Host Ethernet Interfaces

Add Clone Properties Delete iSCSI Boot usNIC

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot
eth0	VIC-3-eth0	00:81:C4:D0:DC:0C	9000	0	0	0	NONE	TRUNK	disabled
eth1	VIC-3-eth1	00:81:C4:D0:DC:0D	9000	0	1	0	NONE	TRUNK	disabled
iSCSI-vNIC-A	VIC-3-iSCS	00:81:C4:D0:DC:10	9000	0	0	0	3439	TRUNK	disabled

5. From the iSCSI Boot Configuration window, enter the initiator details:
 - Name: <<var_ucsa_initiator_name_a>>
 - IP Address: <<var_hyperv_hostA_iscsiA_ip>>
 - Subnet Mask: <<var_hyperv_hostA_iscsiA_mask>>
 - Gateway: <<var_hyperv_iscsiA_gateway>>

iSCSI Boot Configuration

IP Version: **IPv4**

Initiator

Name: (0 - 223) chars

IP Address:

Subnet Mask:

Gateway:

Primary DNS:

Secondary DNS:

TCP Timeout: (0 - 255)

CHAP Name: (0 - 50) chars

CHAP Secret: (0 - 50) chars

Configure iSCSI Unconfigure iSCSI Reset Values Cancel

6. Enter the primary target details:
 - Name: IQN number of Infra-SVM
 - IP Address: IP address of `iscsi_lif01a`
 - Boot LUN: 0

Note: You can obtain the storage IQN number by using the `vserver iscsi show` command.

Be sure to record the IQN names for each vNIC. You need them for a later step.

The image shows a 'iSCSI Boot Configuration' dialog box. It has two sections: 'Primary Target' and 'Secondary Target'. Each section contains fields for Name, IP Address, TCP Port, Boot LUN, CHAP Name, and CHAP Secret. The Primary Target fields are: Name: eb800a098aa77ec:vs.3 (1 - 223) chars, IP Address: 172.21.183.33, TCP Port: 3260, Boot LUN: 0 (0 - 65535), CHAP Name: (0 - 50) chars, CHAP Secret: (0 - 50) chars. The Secondary Target fields are: Name: eb800a098aa77ec:vs.3 (1 - 223) chars, IP Address: 172.21.183.34, TCP Port: 3260, Boot LUN: 0 (0 - 65535). At the bottom, there are four buttons: 'Configure iSCSI', 'Unconfigure iSCSI', 'Reset Values', and 'Cancel'.

7. Click Configure iSCSI.

5.5 NetApp FAS Storage Deployment Procedure: Part 2

The following procedures continue the deployment of the NetApp FAS2600 storage system.

NetApp ONTAP SAN Boot Storage Setup

The following steps configure the NetApp FAS2600 storage system for iSCSI boot of the Microsoft Windows Server Hyper-V 2016 hosts.

Create iSCSI Igroups

To create igroups, complete the following step.

Note: You need the iSCSI initiator IQNs from the server configuration for this step.

1. From the cluster management node SSH connection, run the following commands.

```
igroup create -vserver Infra-SVM -igroup HV-Host-Infra-A -protocol iscsi -ostype hyper_v
-initiator <<var_vm_host_inofra_a_iSCSI-A_vNIC_IQN>>,<<var_vm_host_infra_a_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup HV-Host-Infra-B -protocol iscsi -ostype hyper_v
-initiator <<var_vm_host_infra_b_iSCSI-A_vNIC_IQN>>,<<var_vm_host_infra_b_iSCSI-B_vNIC_IQN>>
```

Note: This step must be completed when adding additional Cisco UCS C-Series servers. These igroups are used for all LUN access.

To view the igroups created in this step, run the `igroup show` command.

Map Boot LUNs to Igroups

To map boot LUNs to igroups, complete the following step:

1. From the cluster management SSH connection, run the following commands:

```
lun map -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-A -igroup HV-Host-Infra-A -lun-id 0
lun map -vserver Infra-SVM -volume hyperv_boot -lun HV-Host-Infra-B -igroup HV-Host-Infra-B -lun-id 0
```

Note: This step must be completed when adding additional Cisco UCS C-Series servers.

5.6 Microsoft Windows Server 2016 Deployment Procedure

This section provides detailed procedures for installing Windows Server 2016 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Several methods exist for installing Windows Server 2016 in such an environment. This procedure uses the virtual KVM console and virtual media features of the Cisco IMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.

Log In to Cisco IMC Interface for Cisco UCS C-Series Standalone Servers

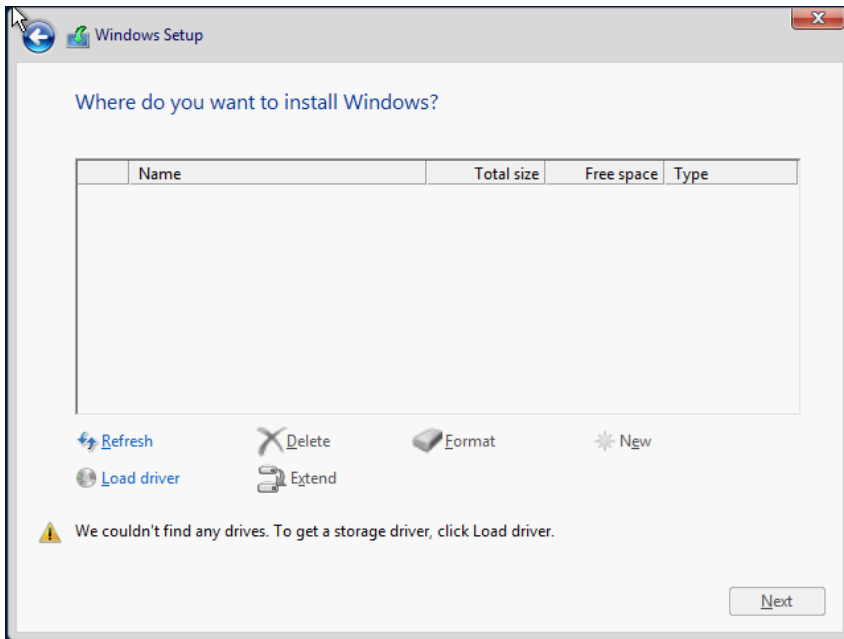
The following steps detail the method for logging in to the Cisco IMC interface for Cisco UCS C-Series standalone servers. You must log in to the Cisco IMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

All Hosts

1. Navigate to a web browser and enter the IP address for the Cisco IMC interface for the Cisco UCS C-Series.
This step launches the Cisco IMC GUI application.
2. Log in to the Cisco IMC GUI using the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.



5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.
 - Note:** You might first need to click Activate Virtual Devices. Select Accept This Session if prompted.
7. Browse to the Windows 2016 image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (cold boot). Click Yes.
The Windows Server 2016 files are loaded.
9. On the Windows setup screen, click Next.
10. Click Install Now. You see a message that setup is starting.
11. Enter the product key and click Next.
12. Select Windows Server 2016 Datacenter (Desktop Experience) and click Next.
13. Accept the license terms and click Next.
14. Select Custom: Install Windows only (advanced).
You do not see a boot drive. You must first load the driver to continue.



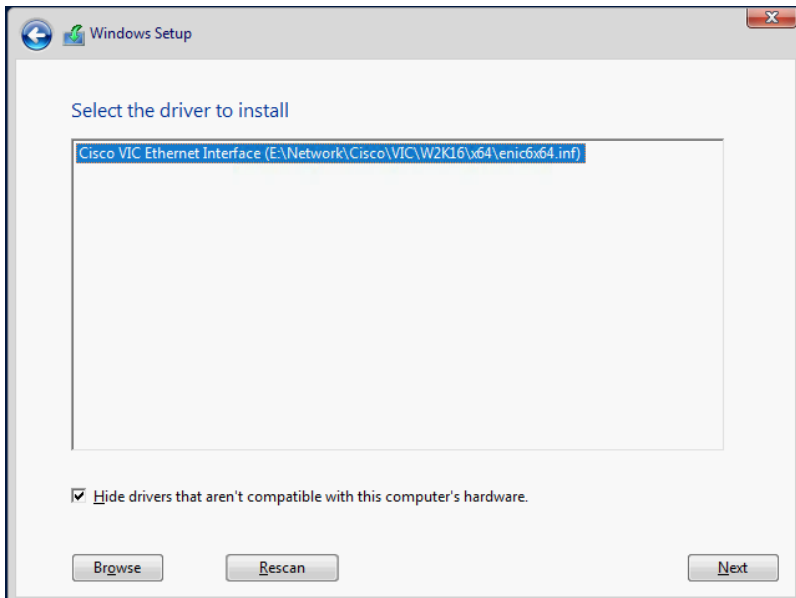
15. Unmap the Windows 2016 installation media and map the Windows driver ISO downloaded from the [Cisco UCS M4 Rack Server Software page](#).

Note: This configuration uses CIMC 2.0(13h) and the 2.0(13f) driver software version.

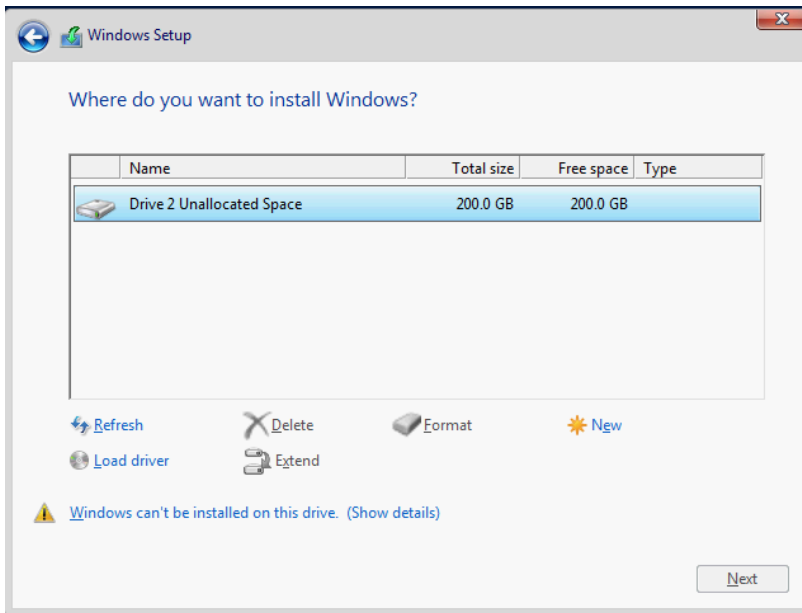
16. Click Load Driver to load the storage drivers to view the iSCSI boot LUN.

17. Browse to the driver file by navigating to Network > Cisco > VIC > W2K16 > x64 and then click OK.

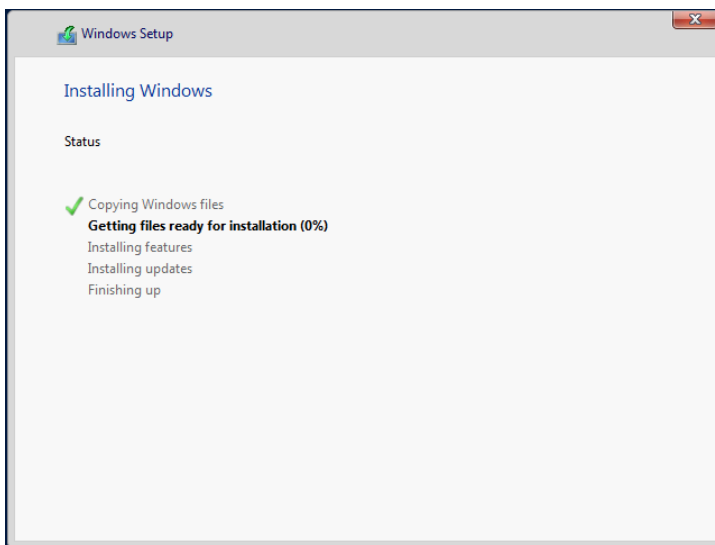
18. You now see the Cisco VIC Ethernet interface. Click Next to install the driver.



19. You now see the iSCSI LUN. Click Next.



20. Unmap the driver ISO and remap the Windows Server 2016 installation ISO to continue with the Windows installation.
21. Click Refresh.
22. Select the iSCSI LUN and click Next.
23. The Windows 2016 installation process now begins. This takes several minutes.



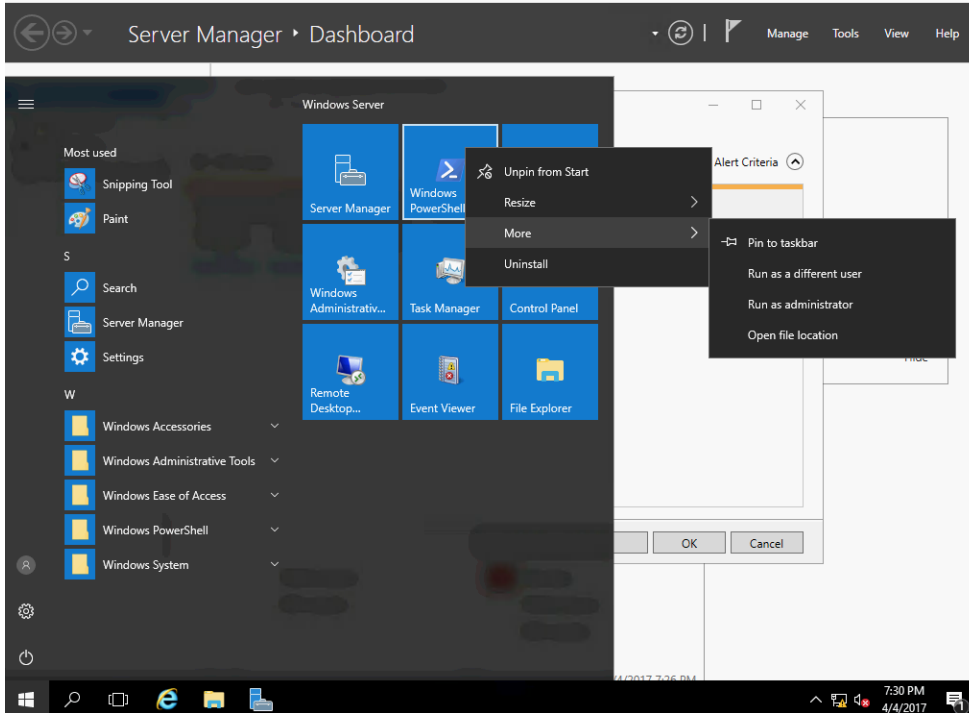
24. The server automatically reboots several times during the Windows 2016 installation process.
25. At the end of the installation, you are prompted for an administrator password.

5.7 Install Microsoft Windows Features

To install the required Microsoft Windows Server 2016 features, complete the following steps:

All Servers

1. Verify that the Windows 2016 installation media are mounted through the Cisco IMC KVM. If not, mount the ISO.
2. Use the Macros menu in the KVM to select Static Macros, then press Ctrl+Alt+Del.
3. Log in to Windows with the administrator password entered during installation.
4. Click the Start menu icon and right-click Windows PowerShell. Select Run as Administrator.



5. Type the following command at the PowerShell command prompt. This enables remote management.

```
SCONFIG
```

6. Select Configure Remote Management by typing 4 and pressing Enter.
7. Select Enable Remote Management by typing 1 and pressing Enter.
8. You see a message that you have successfully enabled remote management. Click OK.
9. Return to the main menu by typing 4 and pressing Enter.
10. Select Remote Desktop by typing 7 and pressing Enter.
11. Enter E to enable.
12. Enter 2 to allow any version of remote desktop. A message displays notifying you that remote desktop has been enabled. Click OK.

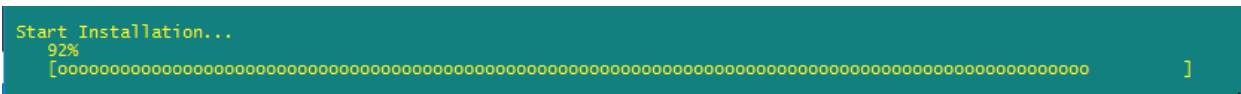
After this step, you are returned to the main menu.

13. Select Exit to Command Line by typing 15 and pressing Enter.
14. Add the Hyper-V, MPIO, and Windows failover clustering features by entering the following command:

```
Add-WindowsFeature Hyper-V, Failover-Clustering, Multipath-IO  
-IncludeManagementTools -Source F:\sources\sxs -Restart
```

Note: If the ISO image is not mounted to drive F:\, the source path needs to be changed to reflect the drive letter.

15. You can see the installation begin at the top of the PowerShell window. This takes several minutes.



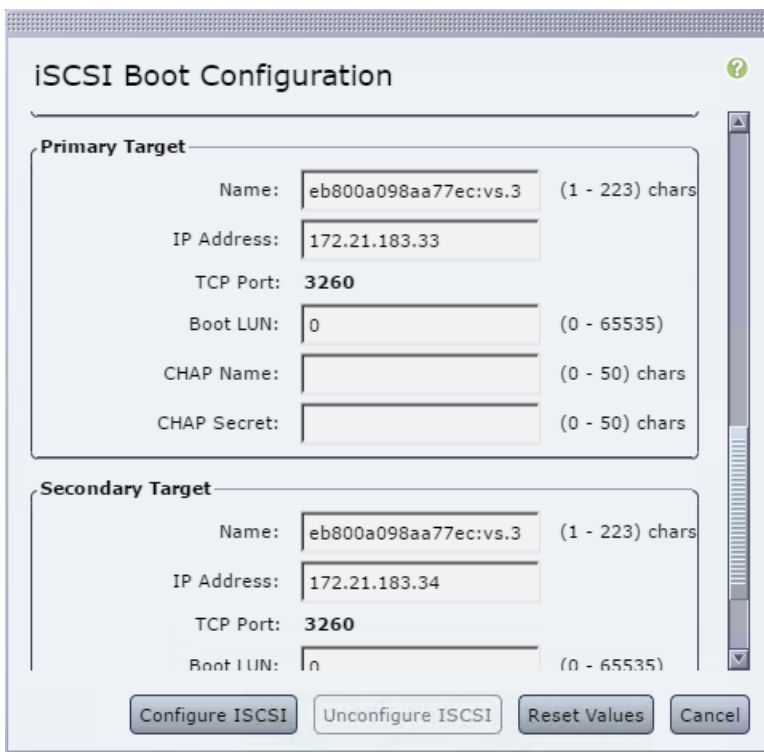
16. Unmap the Windows Server 2016 installation media from the Virtual Media tab when the installation has completed.

Configure Cisco VIC1227 Adapter for iSCSI Boot: Part 2

Now the remaining paths are configured between the server and the iSCSI LUN, since MPIO has been configured.

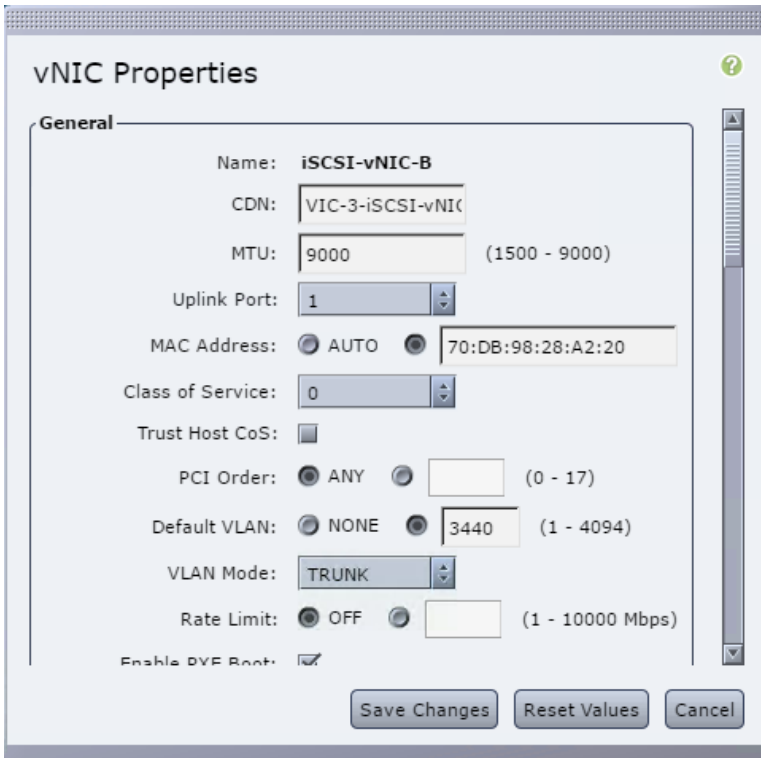
1. From the Cisco IMC interface browser window, click Inventory.
2. In the right pane, click Cisco VIC Adapters.
3. From the Adapter Cards section, select UCS VIC 1227.
4. From the Host Ethernet Interfaces section, select the vNICs tab.
5. Select iSCSI-vNIC-A and click iSCSI Boot, located on the top of the Host Ethernet Interfaces section.
6. Enter the secondary target details:
 - Name: IQN number of Infra-SVM
 - IP Address: IP address of iscsi_lif02a
 - Boot LUN: 0

Note: You can obtain the storage IQN number by using the `vserver iscsi show` command.



7. Click Configure iSCSI.

8. In the Host Ethernet Interfaces section, click Add to create a new vNIC.
9. In the Add vNIC window, complete the following settings:
 - Name: iSCSI-vNIC-B
 - MTU: 9000
 - Uplink Port: 1
 - Default VLAN: <<var_iscsi_vlan_b>>
 - VLAN Mode: TRUNK
 - Enable PXE Boot: check



10. Click OK.
11. Select the newly created vNIC iSCSI-vNIC-B and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.
12. From the iSCSI Boot Configuration window, enter the initiator details:
 - Name: <<var_ucsa_initiator_name_b>>
 - IP Address: <<var_hyperv_hostb_iscsib_ip>>
 - Subnet Mask: <<var_hyperv_hostb_iscsib_mask>>
 - Gateway: <<var_hyperv_hostb_iscsib_gateway>>
13. Enter the primary target details:
 - Name: IQN number of Infra-SVM
 - IP Address: IP address of iscsi_lif01b
 - Boot LUN: 0
14. Enter the secondary target details:
 - Name: IQN number of Infra-SVM

- IP Address: IP address of iscsi_lif02b
- Boot LUN: 0

Note: You can obtain the storage IQN number by using the `vserver iscsi show` command.

15. Click Configure iSCSI.

16. Reboot the Windows host.

Note: This process must be completed for Cisco UCS server B and any additional Cisco UCS C-Series servers added.

5.8 Configure Microsoft Windows

To configure the network for each Hyper-V host, complete the following steps. You need the following additional IP addresses for each Hyper-V host.

Table 16) Information required for configuring Hyper-V hosts.

Detail	Value
Hyper-V host name	
Hyper-V host management IP	
Hyper-V host management mask	
Hyper-V host management gateway	
Hyper-V host CSV	
Hyper-V host CSV mask	
Hyper-V host CSV gateway	
Hyper-V host live migration (LM) IP	
Hyper-V host live migration (LM) mask	
Hyper-V host live migration (LM) gateway	
Hyper-V host iSCSI-A data IP	
Hyper-V host iSCSI-A mask	
Hyper-V host iSCSI-A gateway	
Hyper-V host iSCSI-B data IP	
Hyper-V host iSCSI-B mask	
Hyper-V host iSCSI-B gateway	

All Servers

1. Log in to the server with the administrator account using the password you provided during installation.
2. Click the Start menu icon and right-click Windows PowerShell. Select Run as Administrator.
3. Find the 10GbE interfaces by running the `Get-NetAdapter` command.

```
PS C:\Users\Administrator> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
Ethernet 4	Cisco VIC Ethernet Interface #4	8	Up	00-81-C4-D0-DC-11	10 Gbps
Ethernet	Cisco VIC Ethernet Interface	10	Up	00-81-C4-D0-DC-10	10 Gbps
Ethernet 3	Cisco VIC Ethernet Interface #3	7	Up	00-81-C4-D0-DC-0D	10 Gbps
Ethernet 2	Cisco VIC Ethernet Interface #2	3	Up	00-81-C4-D0-DC-0C	10 Gbps


4. Configure jumbo frames on the physical interfaces.

```
Set-NetAdapterAdvancedProperty -Name Ethernet* -DisplayName "Jumbo Packet" -DisplayValue "9014 Bytes" -EA SilentlyContinue
Set-NetAdapterAdvancedProperty -Name Ethernet* -DisplayName "Jumbo Packet" -DisplayValue "9014" -EA SilentlyContinue
```

5. Log in to the Cisco IMC web interface, select Inventory from the left pane, and click Cisco VIC Adapters in the right pane.

6. Select the Cisco UCS VIC 1227 and click vNICs.

7. Note the MAC address of interfaces eth0 and eth1. These are the interfaces that are used for 10GbE data connectivity.



8. Rename the NICs in Windows in order to easily identify the NICs to be used for data.

9. Rename the Windows NIC corresponding to eth0:

- a. Find the Windows NIC that corresponds to eth0 by comparing the MAC addresses. In this case, it is Ethernet 2.
- b. Issue the following command, where Ethernet 2 is the Windows NIC corresponding to eth0:

```
Rename-NetAdapter -Name "Ethernet 2" -NewName eth0
```

10. Rename the Windows NIC corresponding to eth1:

- a. Find the Windows NIC that corresponds to eth1 by comparing the MAC addresses. In this case, it is Ethernet 3.
- b. Issue the following command, where Ethernet 3 is the Windows NIC corresponding to eth1:

```
Rename-NetAdapter -Name "Ethernet 3" -NewName eth1
```

11. Run the Get-NetAdapter command to confirm that the NIC names have been changed and correspond to their counterparts as listed in CIMC.

```
PS C:\Users\Administrator> Get-NetAdapter
```

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
Ethernet 4	Cisco VIC Ethernet Interface #4	8	Up	00-81-C4-D0-DC-11	10 Gbps
Ethernet	Cisco VIC Ethernet Interface	10	Up	00-81-C4-D0-DC-10	10 Gbps
eth1	Cisco VIC Ethernet Interface #3	7	Up	00-81-C4-D0-DC-0D	10 Gbps
eth0	Cisco VIC Ethernet Interface #2	3	Up	00-81-C4-D0-DC-0C	10 Gbps

12. Create a NIC team using the eth0 and eth1 10GbE interfaces.

```
New-NetLbfoTeam -Name 10GigTeam -TeamMembers eth0, eth1 -TeamingMode SwitchIndependent -LoadBalancing HyperVPort
```

13. Press Y to confirm.

14. Remove the IP stack from the TM NIC interface.

```
Get-NetAdapter 10GigTeam | set-NetAdapterBinding -ComponentID ms_tcpip* -Enabled $false
```

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

```
New-VMSwitch -Name VSwitch -NetAdapterName 10GigTeam -AllowManagementOS $false
```

16. Create VM NICs.

```
Add-VMNetworkAdapter -ManagementOS -Name Mgmt -SwitchName VSwitch
Add-VMNetworkAdapter -ManagementOS -Name Cluster -SwitchName VSwitch
Add-VMNetworkAdapter -ManagementOS -Name LM -SwitchName VSwitch
Add-VMNetworkAdapter -ManagementOS -Name iSCSI-A -SwitchName VSwitch
Add-VMNetworkAdapter -ManagementOS -Name iSCSI-B -SwitchName VSwitch

Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Mgmt -Access -AccessVlanId
<<mgmt_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName LM -Access -AccessVlanId
<<livemigraion_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Cluster -Access -AccessVlanId
<<cluster_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName iSCSI-A -Access -AccessVlanId
<<iscsi_b_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName iSCSI-B -Access -AccessVlanId
<<iscsi_a_vlan_id>>
```

17. Configure jumbo frames on the select interfaces.

```
Set-NetAdapterAdvancedProperty -Name *LM*,*Cluster*, *iSCSI-A*, *iSCSI-B* -DisplayName "Jumbo
Packet" -DisplayValue "9014 Bytes" -EA SilentlyContinue
Set-NetAdapterAdvancedProperty -Name *LM*,*Cluster*, *iSCSI-A*, *iSCSI-B* -DisplayName "Jumbo
Packet" -DisplayValue "9014" -EA SilentlyContinue
```

18. Set IP address information for each host NIC.

```
New-NetIPAddress -InterfaceAlias 'vEthernet (Mgmt)' -IPAddress <Mgmt_ipaddress> -DefaultGateway
<<Mgmt_gateway>> -PrefixLength <Mgmt_network_prefix>
New-NetIPAddress -InterfaceAlias 'vEthernet (Cluster)' -IPAddress <Cluster_ipaddress> -
PrefixLength <Cluster_prefix>
New-NetIPAddress -InterfaceAlias 'LM' -IPAddress <LM_ipaddress> -Prefix <LM_prefix>
New-NetIPAddress -InterfaceAlias 'iSCSI-A' -IPAddress <ISCSI-A_ipaddress> -PrefixLength <ISCSI-A
prefix>
New-NetIPAddress -InterfaceAlias 'iSCSI-A' -IPAddress <ISCSI-A_ipaddress> -PrefixLength <ISCSI-A
prefix>
```

19. Disable DNS registration for all NICs.

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

20. Turn registration back on and configure DNS for the mgmt NIC.

```
Set-DnsClient -InterfaceAlias 'vEthernet (Mgmt)' -Register $true -ConnectionSpecificSuffix
<dns_connection_suffix>
Set-DnsClientServerAddress -InterfaceAlias 'vEthernet (Mgmt)' -ServerAddresses <dns_server_ips>
```

21. Rename the host.

```
Rename-Computer <ServerName> -restart
```

22. Add the host to Active Directory.

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

Note: A dialog box prompts for a user name and password. After the password is entered, the server reboots.

23. Verify that DNS records have been created for the host name and management IP address of the Hyper-V host.

Install NetApp Windows iSCSI Host Utilities

The following section describes how to perform an unattended installation of the NetApp Windows Unified Host Utilities. For detailed information about the installation, see the [Windows Host Utilities 7.0 Installation and Setup Guide](#).

All Servers

1. Download Windows iSCSI Host Utilities from this location:
http://mysupport.netapp.com/NOW/download/software/sanhost_win/7.0/netapp_windows_host_utilities_7.0_x64.msi
2. Click the Start menu icon and right-click Windows PowerShell. Select Run as Administrator.
3. Unblock the downloaded file.

```
Unblock-file ~\Downloads\netapp_windows_host_utilities_7.0_x64.msi
```

4. Install the Host Utilities.

```
~\Downloads\netapp_windows_host_utilities_6.0.2_x64.msi /qn "MULTIPATHING=1"
```

Note: Press tab to populate the file path if desired. The system reboots during this process.

Configure Windows Host iSCSI Initiator

To configure the built-in Microsoft iSCSI initiator, complete the following steps.

All Servers

1. Click the Start menu icon and right-click Windows PowerShell. Select Run as Administrator.
2. Configure the iSCSI service to start automatically.

```
Set-Service -Name MSiSCSI -StartupType Automatic
```

3. Start the iSCSI service.

```
Start-Service -Name MSiSCSI
```

4. Configure the MPIO to claim any iSCSI device.

```
Enable-MSDSMAutomaticClaim -BusType iSCSI
```

5. Set the default load balance policy of all newly claimed devices to round robin.

```
Set-MSDSMGlobalDefaultLoadBalancePolicy -Policy LQD
```

6. Configure an iSCSI target for each controller.

```
New-IscsiTargetPortal InitiatorPortalAddress <iscsia_ipaddress> -TargetPortalAddress <<controllerA_iscsi_lif01a_ip>>  
New-IscsiTargetPortal -InitiatorPortalAddress <iscsib_ipaddress> -TargetPortalAddress <<controllerA_iscsi_lif01b_ip>>  
New-IscsiTargetPortal InitiatorPortalAddress <iscsia_ipaddress> -TargetPortalAddress <<controllerB_iscsi_lif02a_ip>>  
New-IscsiTargetPortal InitiatorPortalAddress <iscsib_ipaddress> -TargetPortalAddress <<controllerB_iscsi_lif02b_ip>>
```

7. Connect a session for each iSCSI network to each target.

```
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPortalAddress <iscsia_ipaddress> -TargetPortalAddress <<controllerA_iscsi_lif01a_ip>>  
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPortalAddress <iscsia_ipaddress> -TargetPortalAddress <<controllerB_iscsi_lif02a_ip>>  
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPortalAddress <iscsib_ipaddress> -TargetPortalAddress <<controllerA_iscsi_lif01b_ip>>  
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPortalAddress <iscsib_ipaddress> -TargetPortalAddress <<controllerB_iscsi_lif02b_ip>>
```

Note: Verify that you have completed these configuration steps for each Hyper-V host before continuing.


Note: These steps must be completed when adding additional Hyper-V hosts to FlexPod Express.

5.9 Create Windows Failover Cluster

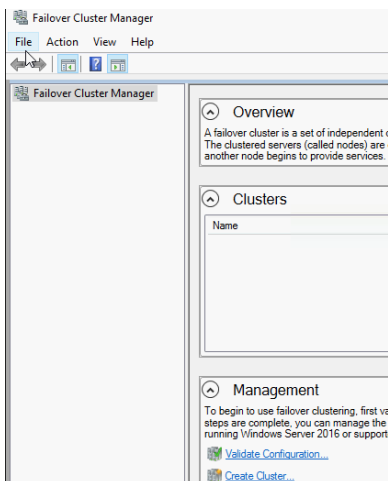
This process creates a Windows failover cluster. Be sure to create DNS records for the cluster name. The IP address for cluster management should be on the <<mgmt_vlan>>.

One Server Only

1. Launch Failover Cluster Manager. It can be found under Windows Administrative Tools on the Start menu.

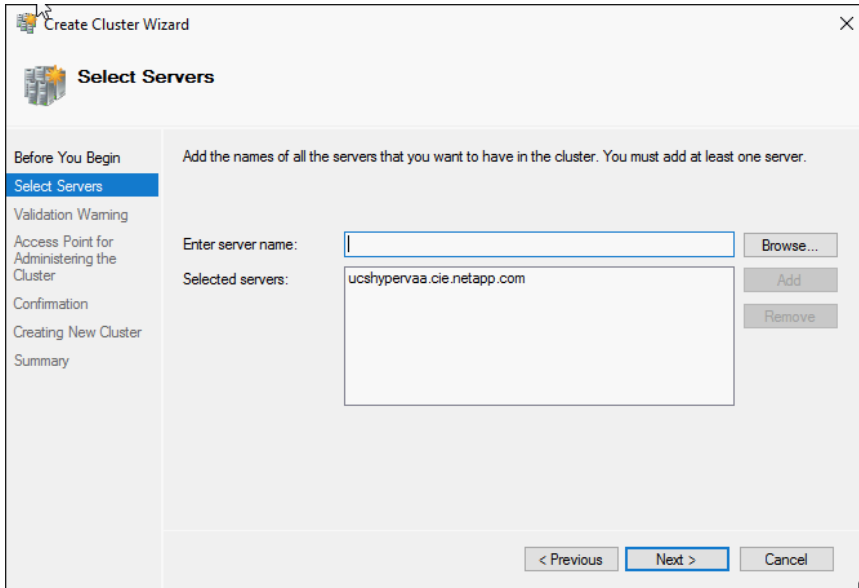


2. Select Create Cluster.

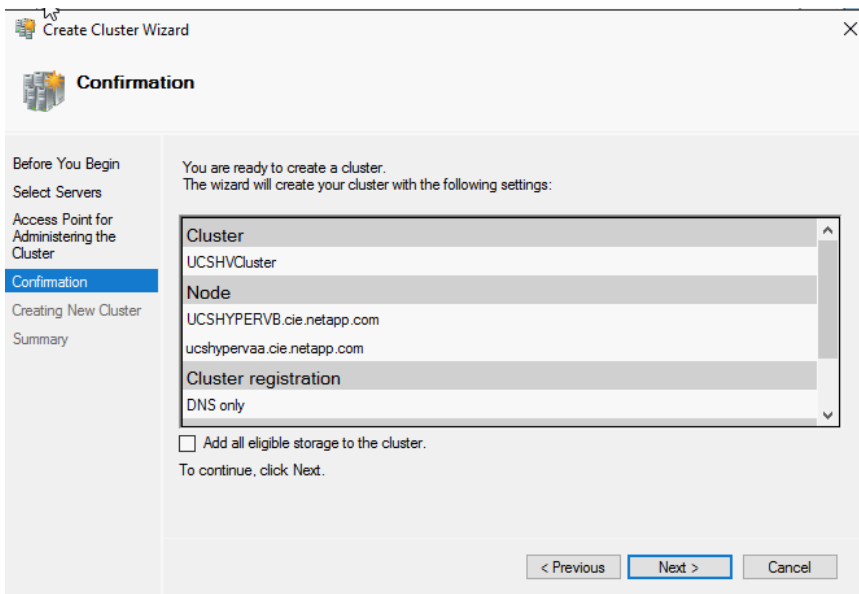


3. Click Next on the Before You Begin screen.

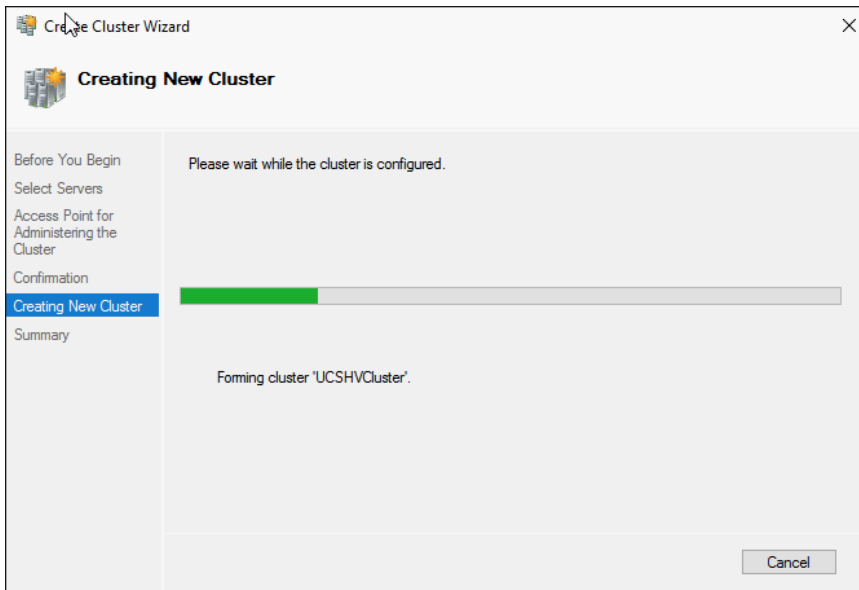
4. In the Select Servers screen, click Browse to browse for the Windows 2016 Hyper-V hosts in Active Directory. Add both Cisco UCS servers. Click Next.



5. On the Access Point for Administering the Cluster screen, enter the cluster name and the IP address to be used for the Hyper-V cluster management next to the appropriate network. Clear the networks you are not going to use. Click Next.
6. At the Confirmation screen, clear the checkbox next to Add all eligible storage to the cluster. Confirm your settings and click Next.



You see the Creating New Cluster dialog box. This takes several minutes.



7. You then see the Summary screen stating the cluster has been successfully created. Click Finish.

Note: If the cluster creation fails, see the Validate Cluster Configuration section, later in this document, to aid with troubleshooting.

Note: You may also add additional nodes to the cluster if additional Cisco UCS C-Series servers are added to the FlexPod Express configuration.

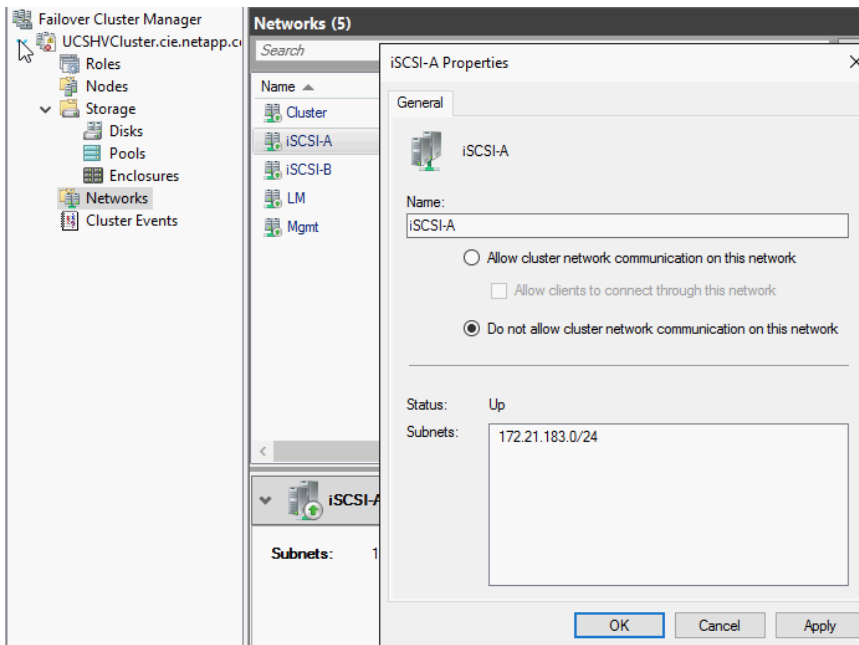
8. Rename the cluster networks according to their purpose by issuing the following commands:

```
Get-ClusterNetworkInterface | ? Name -like *Cluster* | Group Network| %{ (Get-ClusterNetwork
$_.Name).Name = 'Cluster'}
Get-ClusterNetworkInterface | ? Name -like *LM* | Group Network| %{ (Get-ClusterNetwork
$_.Name).Name = 'LM'}
Get-ClusterNetworkInterface | ? Name -like *iSCSI-A* | Group Network| %{ (Get-ClusterNetwork
$_.Name).Name = 'iSCSI-A'}
Get-ClusterNetworkInterface | ? Name -like *iSCSI-B* | Group Network| %{ (Get-ClusterNetwork
$_.Name).Name = 'iSCSI-B'}
Get-ClusterNetworkInterface | ? Name -like *Mgmt* | Group Network| %{ (Get-ClusterNetwork
$_.Name).Name = 'Mgmt'}
```

9. Set the cluster network.

```
(Get-ClusterNetwork -Name Cluster).Metric = 900
```

10. Right-click the iSCSI-A network under Networks in Failover Cluster Manager. Click the Do Not Allow Cluster Network Communication on this Network option. Click Apply. Click OK.

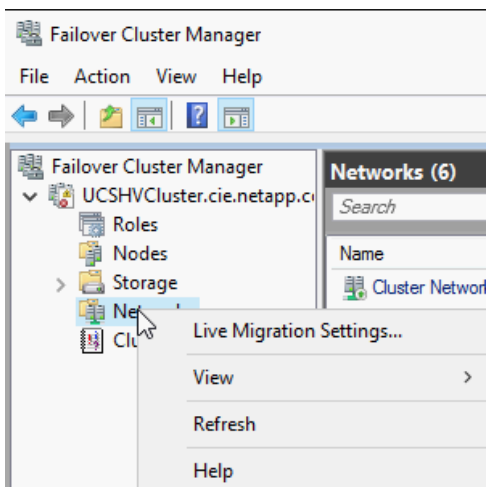


11. Repeat this process for the iSCSI-B network.

Configure Live Migration

To configure live migration, complete the following steps:

1. Launch Failover Cluster Manager.
2. Right-click Networks, then click Live Migration Settings.



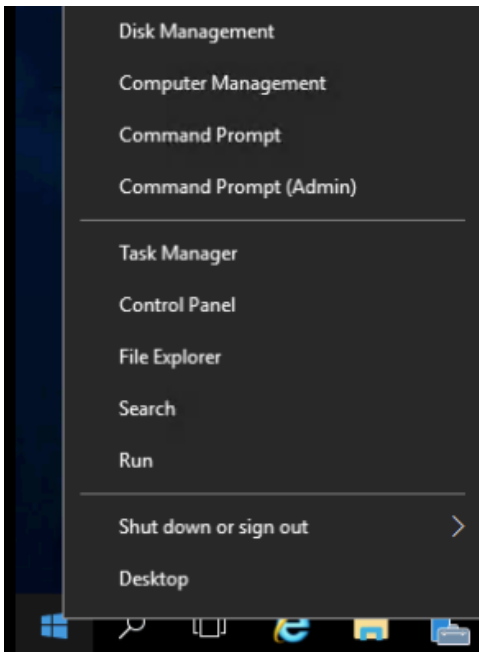
3. Verify that there is only one checkbox selected next to the network to be used for live migration, named LM. Click Apply, then click OK.

Mount Storage to Hyper-V Hosts


Complete the following steps to mount the storage to the Hyper-V hosts.

On First Hyper-V Host

1. Right-click the Start menu icon and select Disk Management.



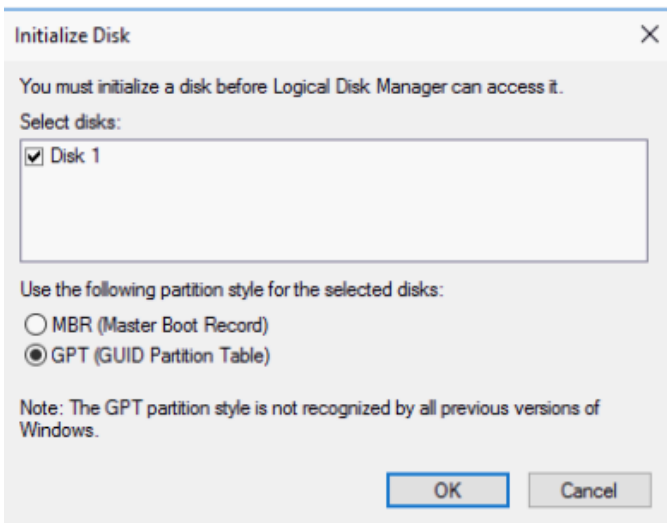
2. After launching disk management, you see a 500GB LUN that represents the provisioning LUN and a 5GB LUN that represents the witness LUN.



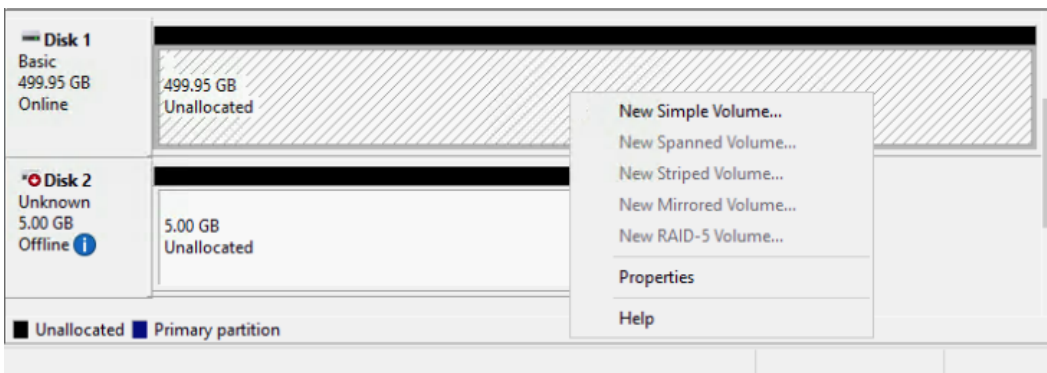
3. Right-click Disk 1 or the 500GB LUN and select Online.



- Right-click Disk 1 again and select Initialize Disk. Select GPT as the partition style in the Initialize Disk menu. Click OK.



- Right-click the unallocated space of disk 1 and select New Simple Volume.



The New Simple Volume wizard launches.

- Click Next on the Specify Volume Size screen.
- Select Do not assign a drive letter or drive path. Click Next.
- Leave the defaults on the Format Partition screen, but change the volume label to Provisioning. Click Next.
- Click Finish.
- Repeat steps 3 through 9 for disk 2. Assign Q: as the drive letter and give it the drive label Quorum.

Note: Complete this process for any additional LUNs you want to add.

On Additional Hyper-V Hosts

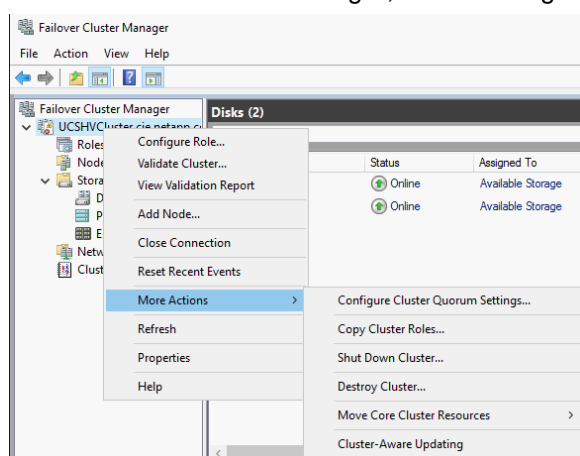
1. Right-click the Start menu icon and select Disk Management.
After launching disk management, you see a 500GB LUN that represents the provisioning LUN and a 5GB LUN that represents the witness LUN. Both LUNs are in the offline state.
2. Right-click the 500GB LUN. Select Online.
You now see the provisioning LUN online.
3. Right-click the 5GB LUN. Select Online.
You now see the quorum LUN online.

Note: These steps must be completed when additional Hyper-V hosts are added to the FlexPod Express configuration. Complete this process for any additional LUNs you want to add.

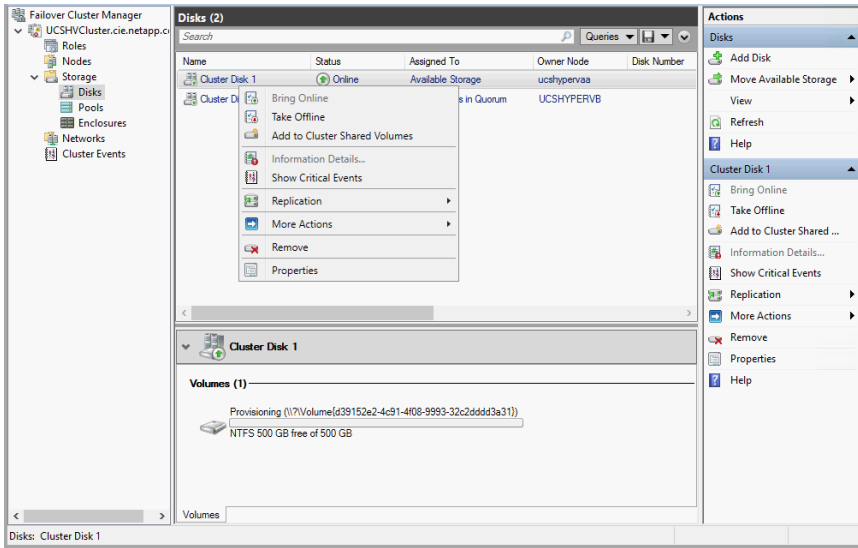
Configure Cluster Disks

Complete the following steps to assign a quorum disk to the cluster and to add the provisioning disk to the cluster.

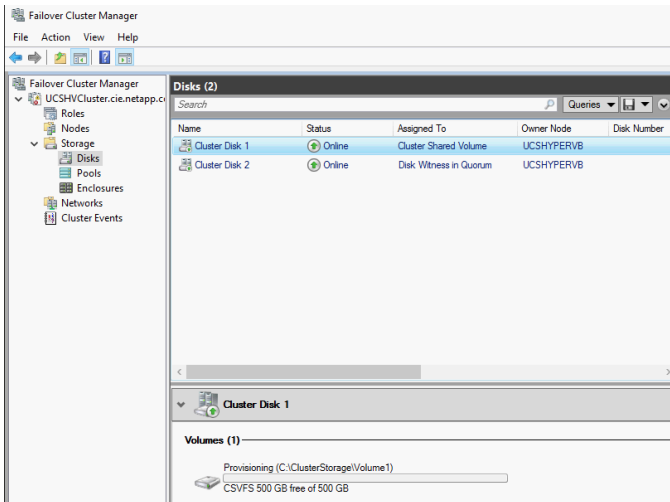
1. In Failover Cluster Manager, select Configure Cluster Quorum Settings.



2. Click Next on the Before You Begin screen.
3. Select the Select the quorum witness option on the Select Quorum Configuration Option screen. Click Next.
4. Select Configure a disk witness.
5. Select the 5GB quorum disk Q:.
6. Click Next on the Confirmation screen.
7. Click Finish.
8. Right-click the 500GB provisioning disk and select Add to Cluster Shared Volumes. Alternatively, you may click Add to Cluster Shared Volumes in the right navigation pane.



9. The disk is now added and assigned to a cluster shared volume.

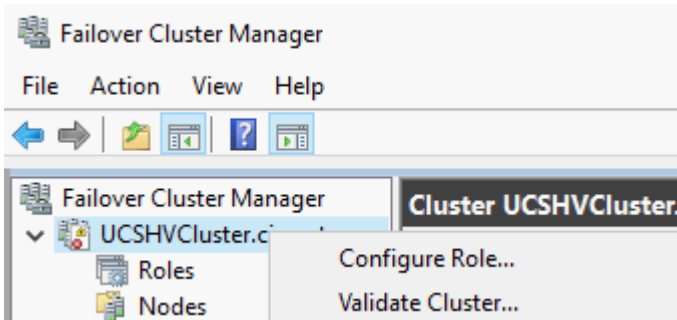


Note: The provisioning LUN is available at `C:\ClusterStorage\Volume1` on the Hyper-V hosts. Complete this step for any additional LUNs you want to add.

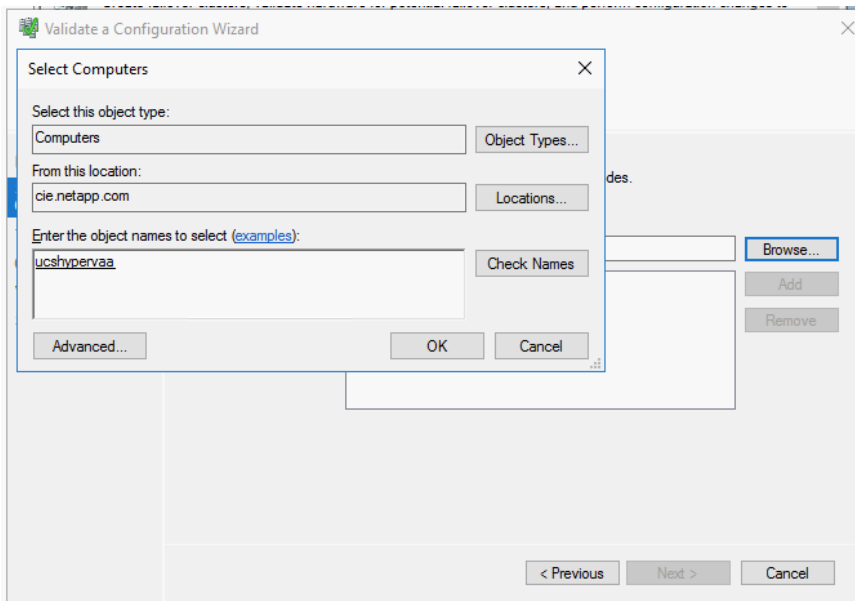
Validate Cluster Configuration

After creating a Windows failover cluster, NetApp recommends validating the configuration. This process can also be helpful in determining the cause of issues if the cluster cannot be created.

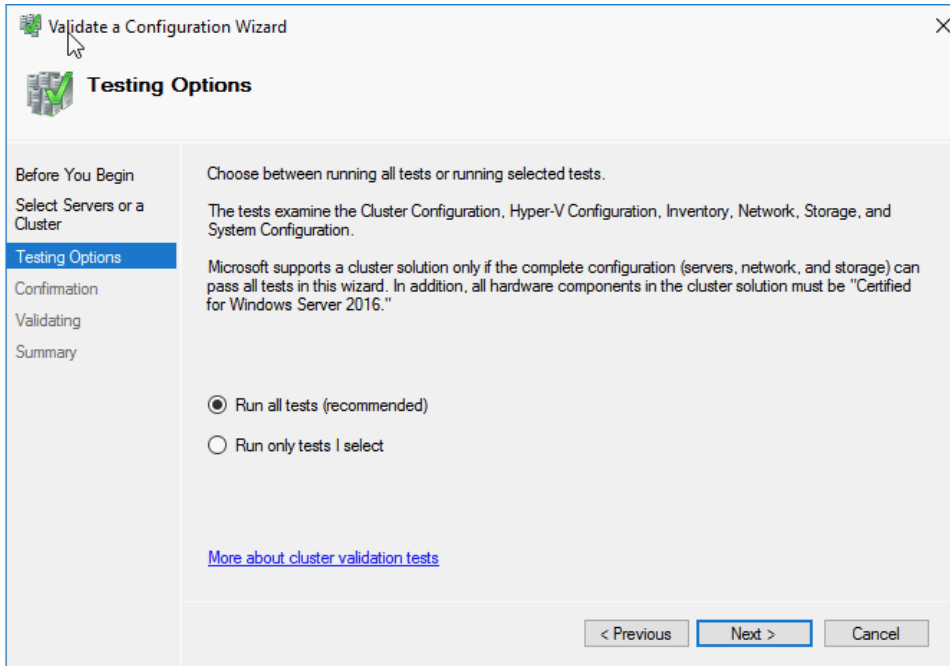
1. Launch Failover Cluster Manager.
2. Right-click the cluster name and select Validate Cluster.



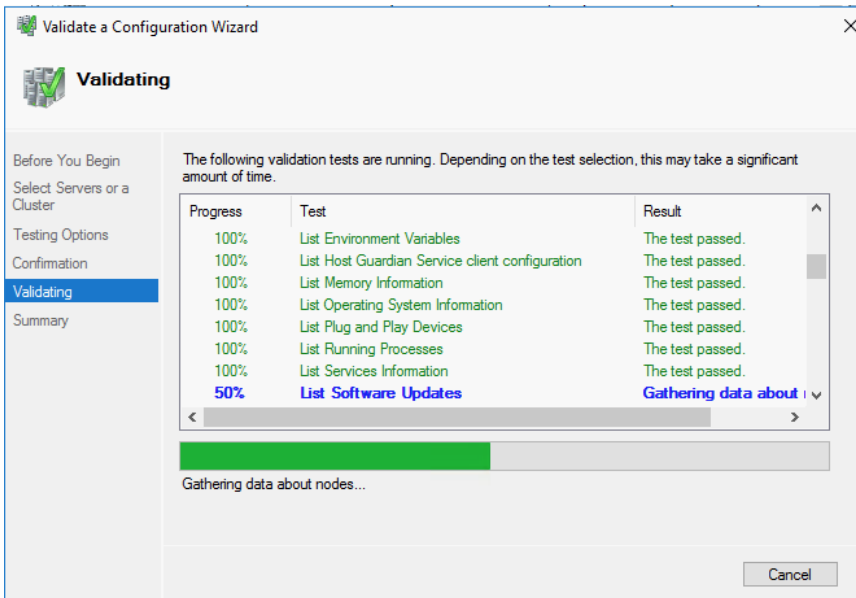
3. Click Next on the Before You Begin screen.
4. Browse AD for the nodes on which you would like to run the Validate Configuration wizard.



5. Click Next after you have added the desired nodes.
6. Select Run All Tests in the testing options screen. Click Next.



- Click Next to continue to the Confirmation screen.
The test is now performed. This takes several minutes.



- You then see the Summary screen. If any of the tests have failed, click View Report... to see details.
Note: All tests must succeed (with or without warnings) for the cluster solution to be supported by Microsoft.

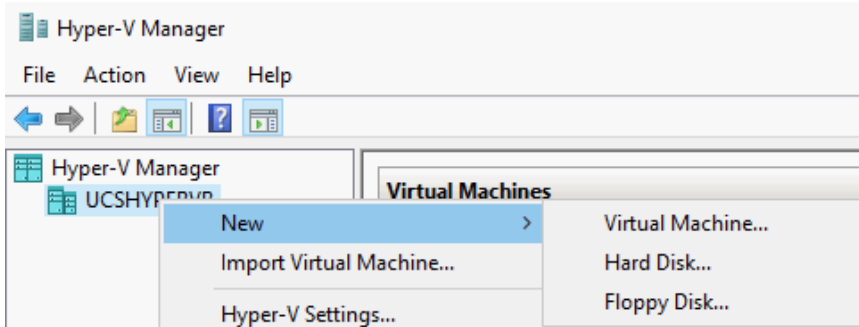
5.10 Install and Configure Hyper-V Management Software

The following steps illustrate the installation of System Center Virtual Machine Manager (referred to as SCVMM) and the NetApp SMI-S Provider.

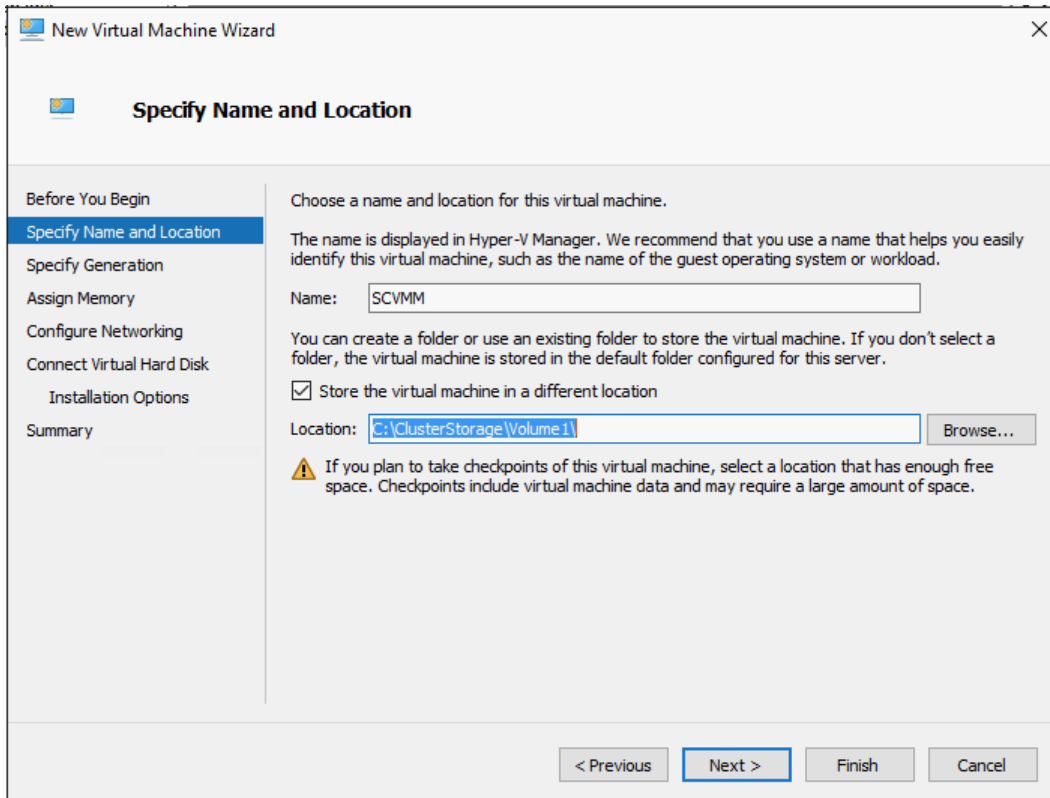
Create SCVMM and SMI-S Virtual Machines

Create the virtual machines for SCVMM and the SMI-S Provider. Complete the following steps on the Hyper-V host of your choice:

1. Launch Hyper-V Manager by clicking the Start menu, then Windows Administrative Tools, then Hyper-V Manager.
2. Right-click the host, select New, then select Virtual Machine.

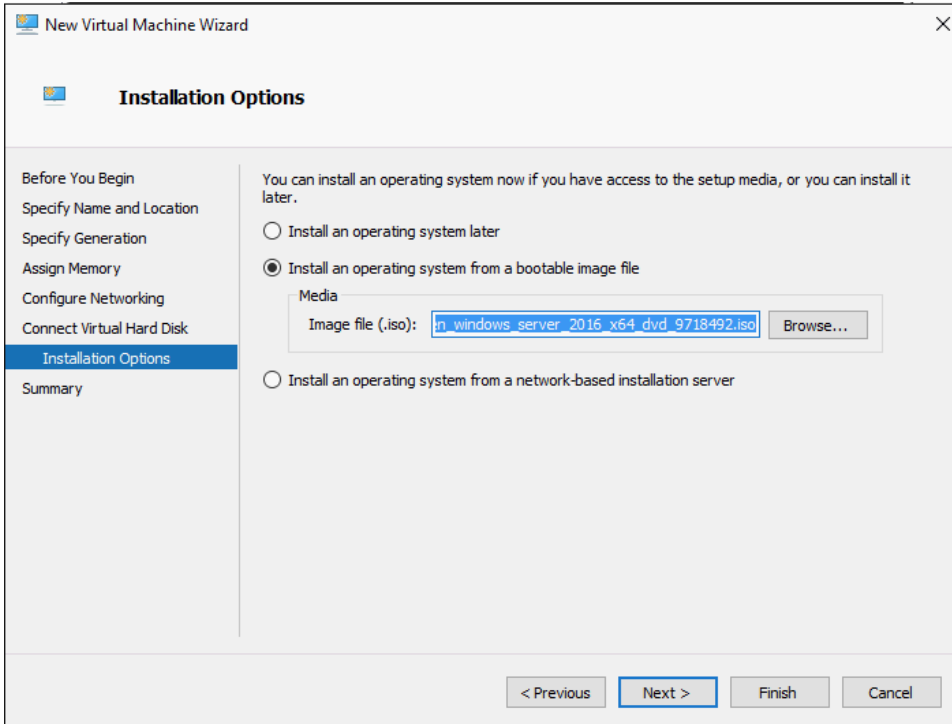


3. Click Next at the Before You Begin screen.
4. Enter a name for the virtual machine, such as SCVMM. Select the checkbox next to Store the virtual machine in a different location and navigate to the cluster shared volume you have created. Click Next.

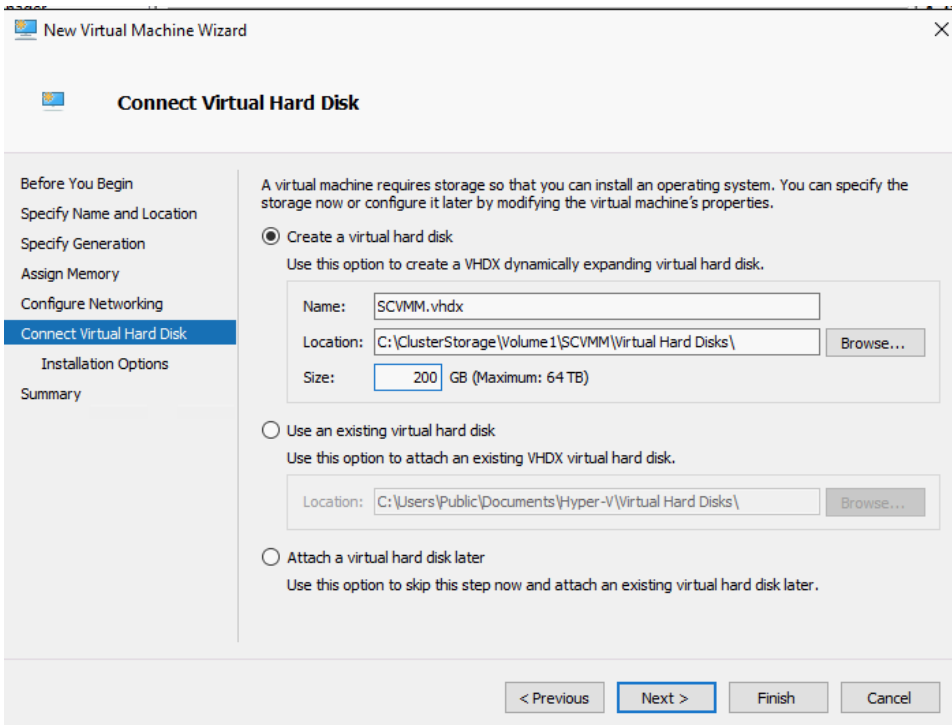


5. Select Generation 2 in the Specify Generation screen. Click Next.
6. On the Assign Memory screen, enter 16384MB as the startup memory for the virtual machine, then select the Use Dynamic Memory for this virtual machine checkbox. Click Next.

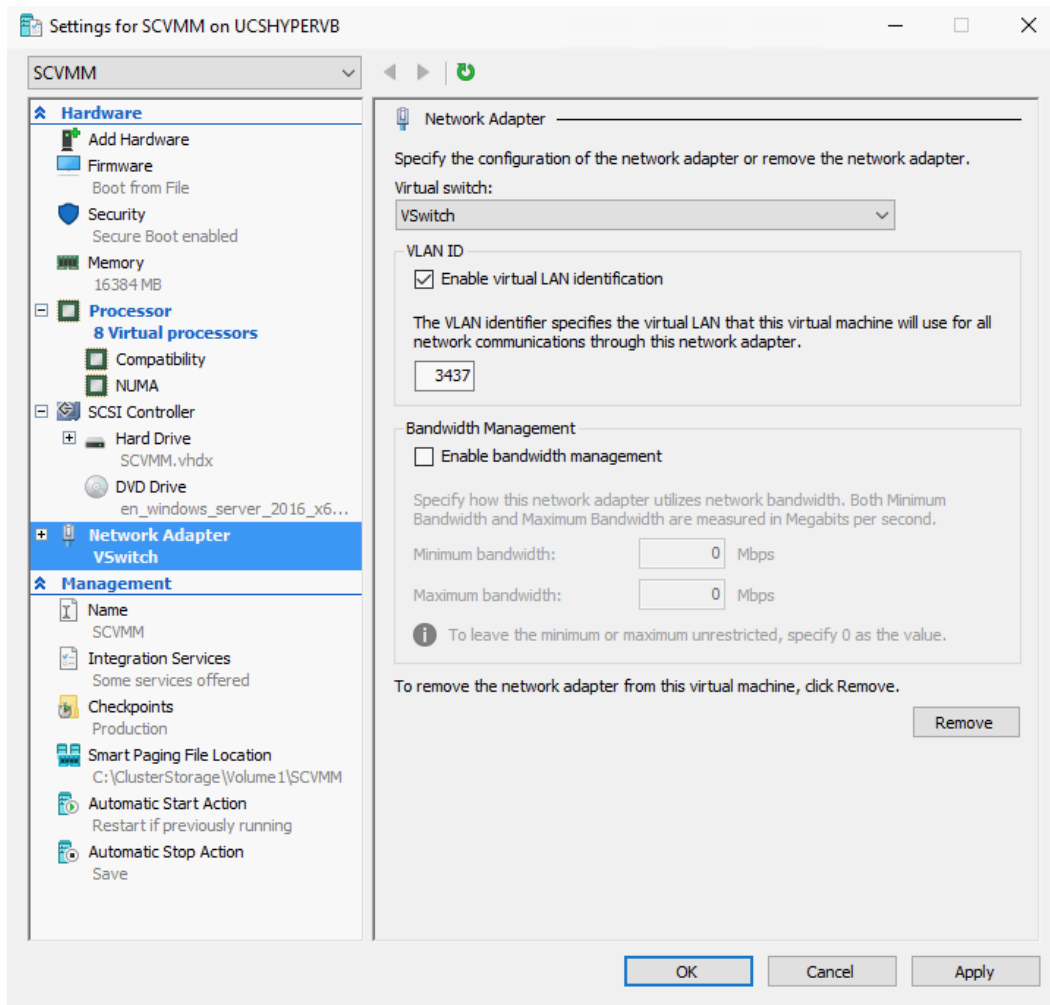
7. In the Configure Networking screen, select VSwitch in the Connection pull-down menu. Click Next.
8. In the Connect Virtual Hard Disk screen, verify that the Create a virtual hard disk option is checked. Change the size to 200GB.
9. In the Installation Options screen, select the Install an Operating System from a Bootable Image File option and browse to the location of the Windows 2016 ISO. Click Next.



10. View the Summary screen and click Finish.

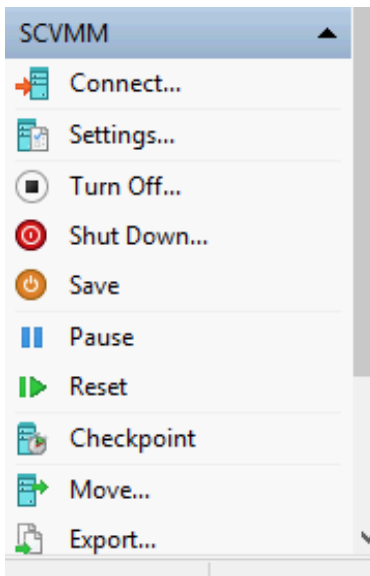


11. Right-click the virtual machine in Hyper-V Manager and select Settings:
 - Under Processors, change the number of virtual processors to 8.
 - Under Network Adapter, select the box that says Enable virtual LAN identification and enter the mgmt VLAN number.
 - Under Memory, change the minimum RAM to 4096MB.
 - Click Apply, then click OK.



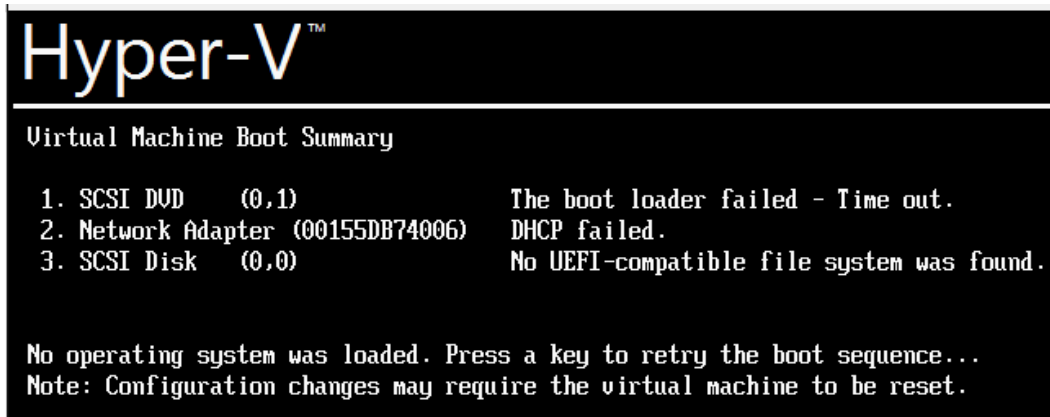
12. Start the virtual machine in Hyper-V Manager by right-clicking the virtual machine name and selecting Start.
13. Connect to the virtual machine by right-clicking the virtual machine name and clicking Connect.

Note: Alternatively, you may use the buttons in the right navigation pane after you have selected the virtual machine.



14. Install Windows 2016 on the local disk of the machine.

Note: If you have missed the “Press any key to boot from CD or DVD” message, you need to reboot the virtual machine. Press any key to retry the boot sequence. If you see the message in the following image, then you should immediately see the “Press any key to boot from CD or DVD” message. Press any key to boot from CD or DVD and begin the Windows installation.



15. Repeat this process to create a new virtual machine to host the SMI-S provider:

- Name the virtual machine.
- Specify Generation 2.
- Place the virtual machine in the same folder as the previous one.
- Assign 8192MB of RAM and select Use Dynamic Memory for this virtual machine.
- Select VSwitch as the connection.
- Create a new a 200GB virtual disk.
- Specify the Windows 2016 ISO.
- Click Finish.
- Change the settings of the virtual machine to 4 CPU, the minimum RAM to 2048MB, and specify the mgmt VLAN when enabling virtual LAN identification.

- Start the virtual machine.
- Install Windows 2016.

16. Assign each of the servers you have created an IP address on the mgmt network.

```
New-NetIPAddress -InterfaceAlias Ethernet -IPAddress <IPaddress> -PrefixLength <prefixlength> -
DefaultGateway <defaultgateway IP>
```

17. Add the DNS servers.

```
Set-DnsClientServerAddress -InterfaceAlias Ethernet -ServerAddresses <DNS server 1 IP>,<DNS
server 2 IP>
```

18. Rename each server with the desired computer name, then reboot the computer.

```
Rename-computer <servername> -restart
```

19. Join each server to the domain.

```
Add-Computer -DomainName <domain name> -restart
```

Note: The preceding customization steps can be done a number of ways. In this case, PowerShell was used. You may use another method if desired. Repeat this process for both the SCVMM and SMI-S virtual machines you have created.

20. Create DNS records for the newly created virtual machines.

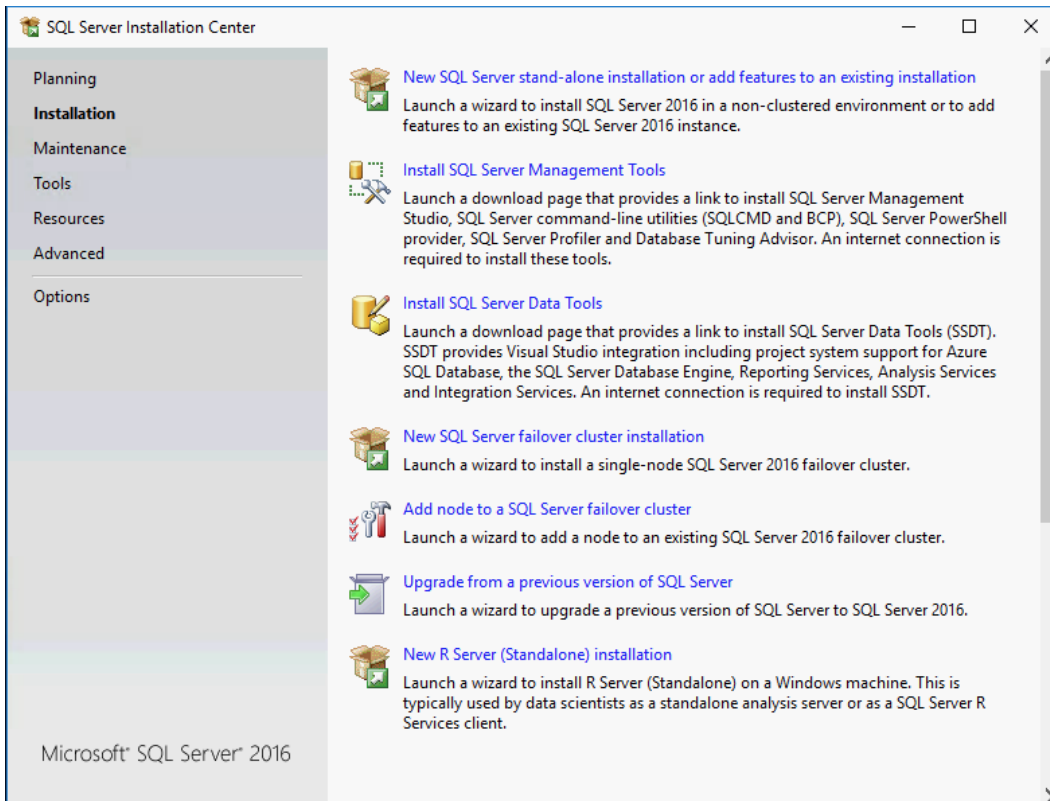
Prepare for SCVMM Installation by Installing SQL Server 2016

Microsoft SQL Server 2016 must be installed on the SCVMM virtual machine before the SCVMM installation. SQL is installed on the same virtual machine as SCVMM.

1. You need a domain service account to install SCVMM. It should be added to the local administrators group of the virtual machine.
2. Open a console to the SCVMM virtual machine.
3. Click Media, then click Insert disk in the console window.

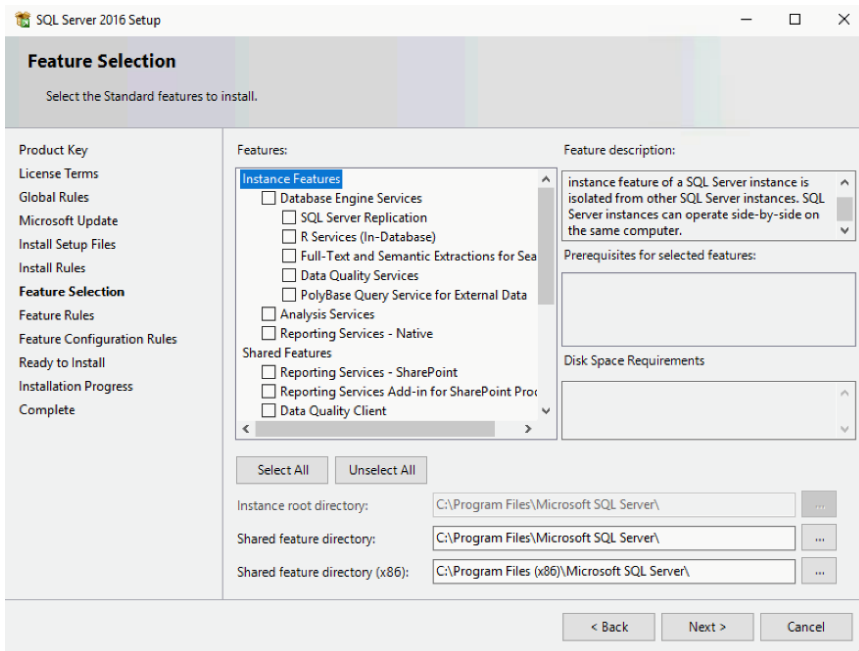


4. Browse to the DVD drive and click Setup.
5. Select the SQL 2016 installation ISO.
6. Double-click Setup.
7. When the SQL Server Installation Center has loaded, click Installation, then click the first option, New SQL Server stand-alone installation or add features to an existing installation.



The SQL Server 2016 setup launches.

8. On the Product Key screen, enter your product key. Click Next.
9. On the License Terms screen, read the license terms, then select the checkbox next to I accept the license terms. Click Next.
On the Global Rules screen, you see a system check run.
10. The next screen is the Microsoft Update screen. Select the box next to Use Microsoft Update to check for updates (recommended). Click Next.
You see the setup files install on the Install Setup screen.
11. You then see the Install Rules screen. Click Next.
12. On the Feature Selection screen, select the checkbox next to Database Engine Services under Instance Features.



There is a check at the Features Rules screen.

13. You then see the Instance Configuration screen. Click Next unless you require a different instance name. Click Next.
14. Click Next at the Server Configuration screen.
15. On the Database Engine Configuration screen, click Add Current User to add the administrative user. Click Next.

Note: You may add another user as needed, such as a domain account. If you use a domain account, it must be added to the sysadmin role. You also use this account during the SCVMM install process.

There is a check at the Features Configuration Rules screen.

16. Click Install at the Ready to Install screen.

You see the installation progress. The installation takes several minutes.

When the installation has completed, you see a window showing the setup operations have succeeded.

17. Click Close.

Install SQL Server Management Tools

To install the SQL Server Management Tools, complete the following steps:

1. From the Setup screen, click Install SQL Server Management Tools. This launches a web browser to download the tools.

Note: The tools can be found at <https://go.microsoft.com/fwlink/?LinkID=531355>.

2. Download the SQL Management Studio for Production Use, `SSMS-Setup-ENU.exe`.
3. Double-click the downloaded file to launch the installer.
4. Select Install to begin.



RELEASE 16.5.3

Microsoft SQL Server Management Studio

Welcome. Click "Install" to begin.

By clicking the "Install" button, I acknowledge that I accept the [License Terms](#) and [Privacy Statement](#).

SQL Server Management Studio transmits information about your installation experience, as well as other usage and performance data, to Microsoft to help improve the product. To learn more about SQL Server Management Studio data processing and privacy controls, see the privacy statement link above.

Install

Close

You see a message that the packages are loading, and then the packages install. This takes several minutes.

5. When installation has completed, a Setup Completed message displays. Click Close.

Install Windows Assessment and Deployment Kit

To install SCVMM, the Windows Assessment and Deployment Kit (ADK) must first be installed. If it has not been installed, SCVMM prompts for this installation during the deployment process.

1. Download the Windows ADK from the following URL: <http://go.microsoft.com/fwlink/?LinkID=614942>

Note: The ADK states it is for Windows 10. It is compatible with Windows 2016.

2. Launch the installer.
3. You are asked where you would like to install the ADK. Click Browse to specify a location or click Next to install it in the default location shown.
4. Read the Windows Kits Privacy message and select the appropriate response.
5. Click Accept to accept the license agreement.
6. At the Features screen, make sure that both Deployment Tools and Windows Preinstallation Environment (Windows PE) are selected. You may leave additional default features and add additional features. Click Install.

The features now install. This takes several minutes.

7. You see a welcome message after the installation has completed. Click Close.

Install SCVMM

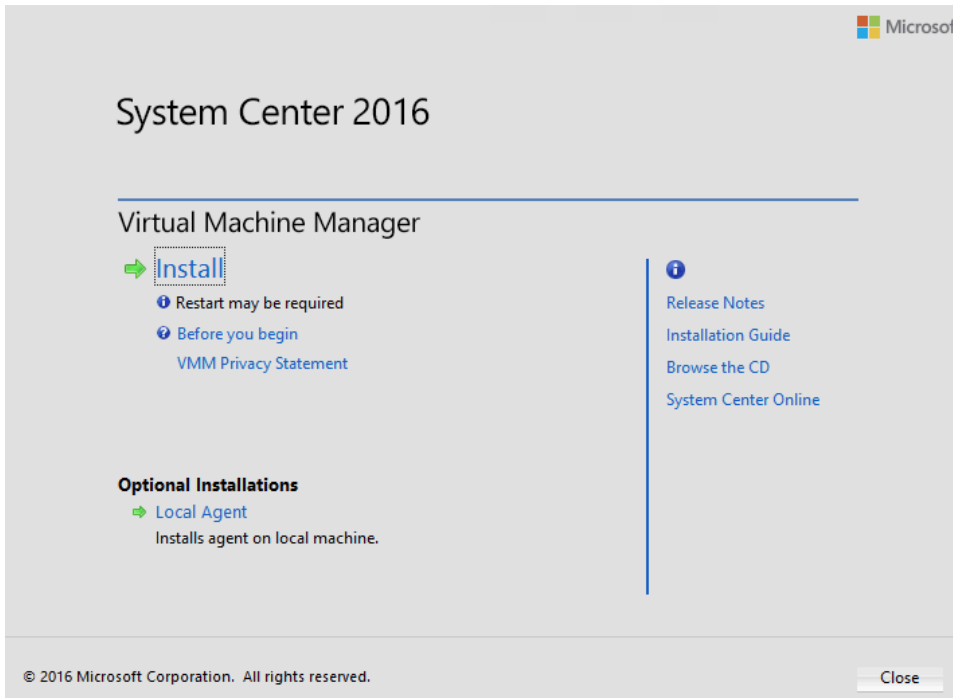
You must be logged in with the domain SCVMM service account to install SCVMM. The following steps show the installation of SCVMM. The files must be extracted before SCVMM can be installed.

1. Log in to the server with the domain SCVMM service account.
2. Mount the ISO file containing SCVMM to the virtual machine.

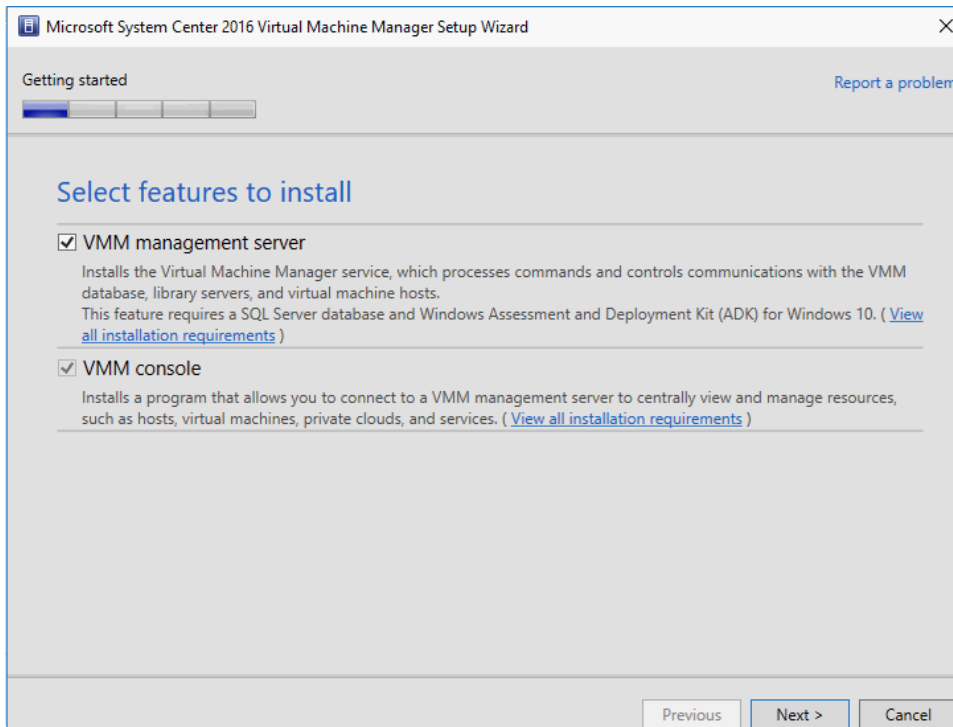
3. Launch the SCVMM Installer.
4. Click Next at the Welcome screen.



5. Select the radio button next to I accept the agreement in the License Agreement window. Click Next.
6. Click Next to extract the files to the location shown. Click Browse to extract the files to another directory, then click Next.
Note: You need this location in a later step.
7. Click Extract at the Ready to Extract screen.
The files now extract. This takes several moments.
8. Click Finish to exit the extractor.
9. Navigate to the directory to which the files were extracted.
10. Double-click `setup.exe`.
11. Click Install to begin the installation of SCVMM.
Note: You must be logged in under a domain account that has administrative credentials on the local computer to install SCVMM.



12. Select the VMM management server feature, which also selects the VMM console feature to install. Click Next.



13. Enter the information for the Name and Organization fields. Enter your Product Key and click Next.
14. Read the license agreement and select the checkbox next to I have read, understood, and agree with the terms of the license agreement. Click Next.

15. Click Next at the Diagnostic and Usage Data screen.
16. Click the On (Recommended) option on the Microsoft Update screen.

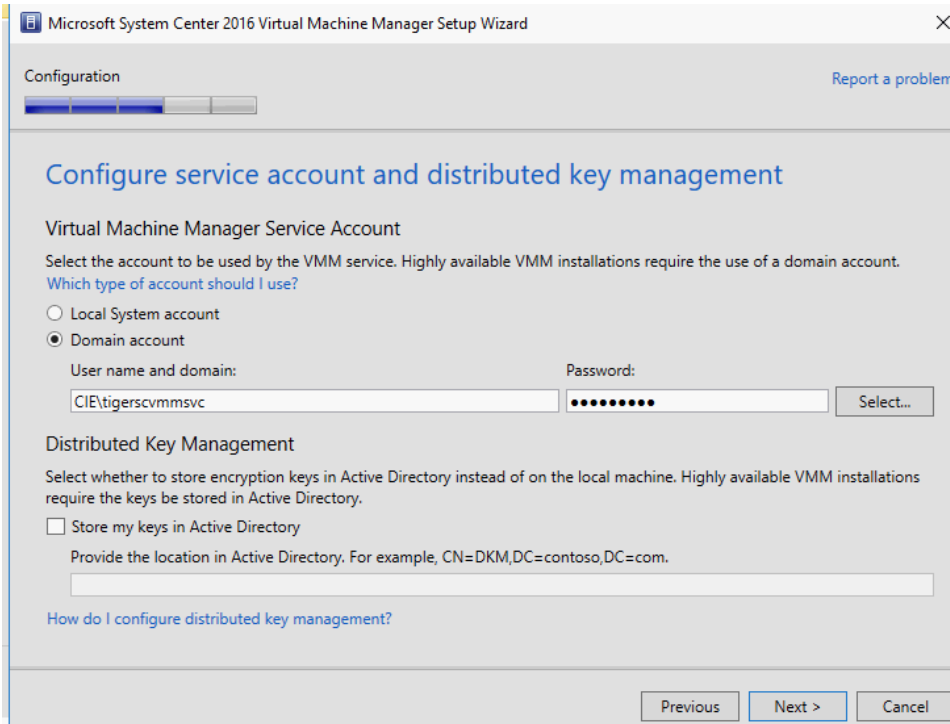
Note: If this violates your organizational policy, you may select Off.
17. Click Next at the installation Location screen, unless a different installation location is desired. If so, use the Browse button to select the installation location. Click Next.
18. The hardware and software check runs.
19. At the Database configuration screen, use the browse button to browse AD for the local computer name. This populates the instance name. Log in with the account you are using for administration that you determined during the SQL install.

The screenshot shows the 'Database configuration' step of the Microsoft System Center 2016 Virtual Machine Manager Setup Wizard. The window title is 'Microsoft System Center 2016 Virtual Machine Manager Setup Wizard'. The configuration progress bar shows the current step is 'Database configuration'. The main heading is 'Database configuration'. Below the heading, there is a prompt: 'Provide information about the database that you would like to use for your VMM management server.' The form contains the following fields and options:

- Server name:** A text box containing 'scvmm-tiger' and a 'Browse' button to its right.
- Port:** An empty text box.
- Use the following credentials
- User name:** An empty text box.
- Format:** A dropdown menu showing 'Domain\UserName'.
- Password:** An empty text box.
- Instance name:** A dropdown menu showing 'MSSQLSERVER'.
- Select an existing database or create a new database.**
- New database:** A text box containing 'VirtualManagerDB'.
- Existing database:** A dropdown menu.

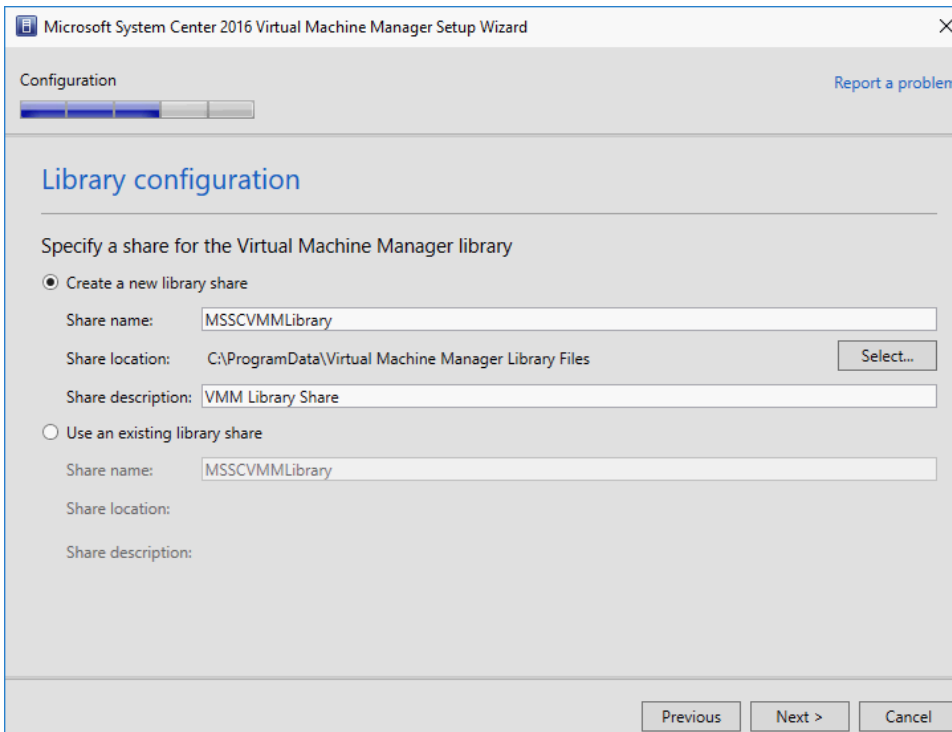
At the bottom of the window, there are three buttons: 'Previous', 'Next >', and 'Cancel'.

20. At the Configure service account and distributed key management screen, enter the name of the account you are using. Use the Select button to add an AD account. Use the Store my keys in Active Directory for a highly available SCVMM installation option checkbox (recommended but not required).



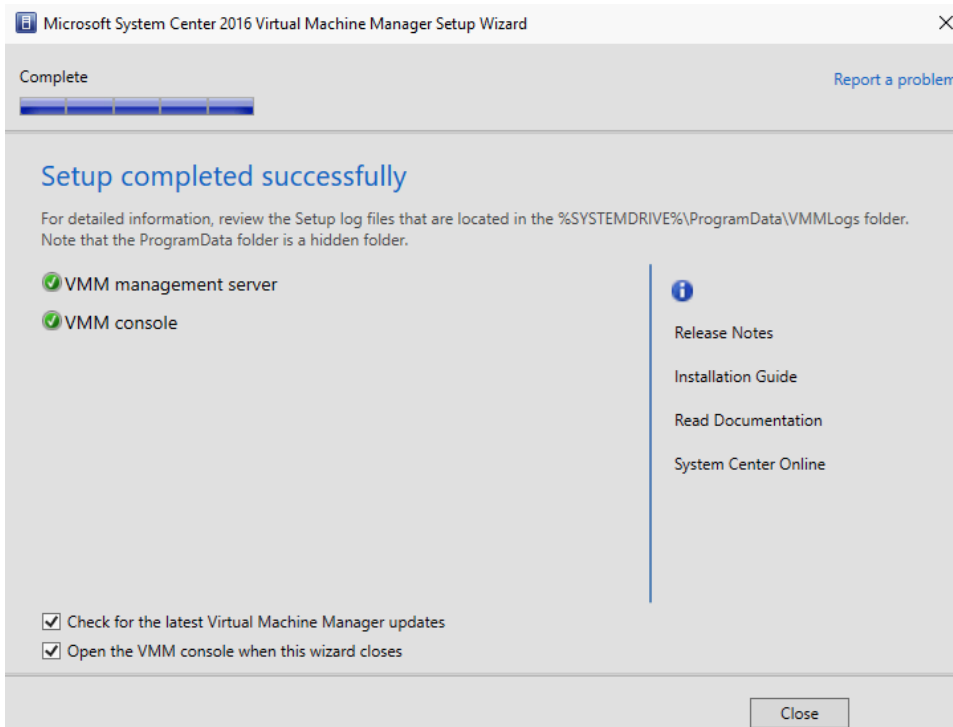
21. Click Next at the Port configuration screen.

22. Create a new library share. Click Next.

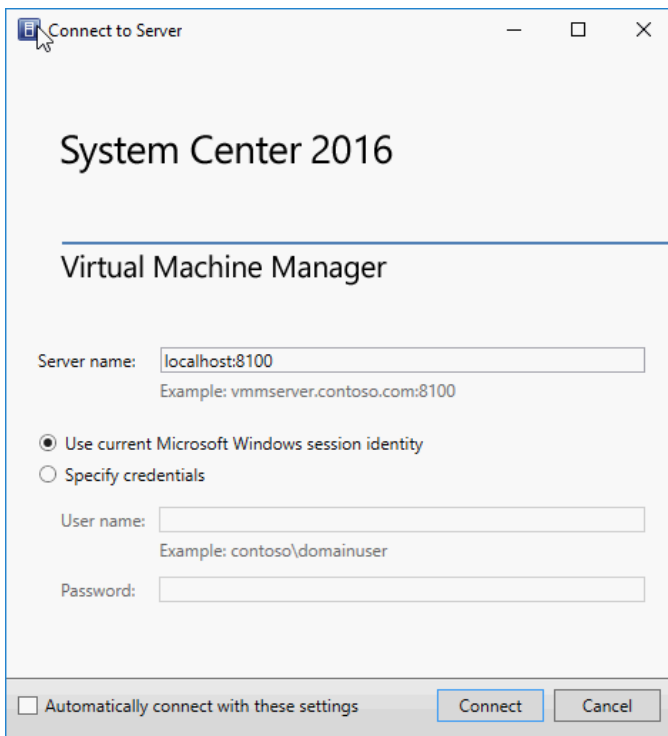


Note: To take full advantage of Hyper-V Offloaded Data Transfers (ODX), the library share should be a NetApp ONTAP LUN.

23. Review the installation summary and click Install. The installation takes several minutes.
24. You see a message that the installation has completed successfully. Click Close. The VMM console launches.



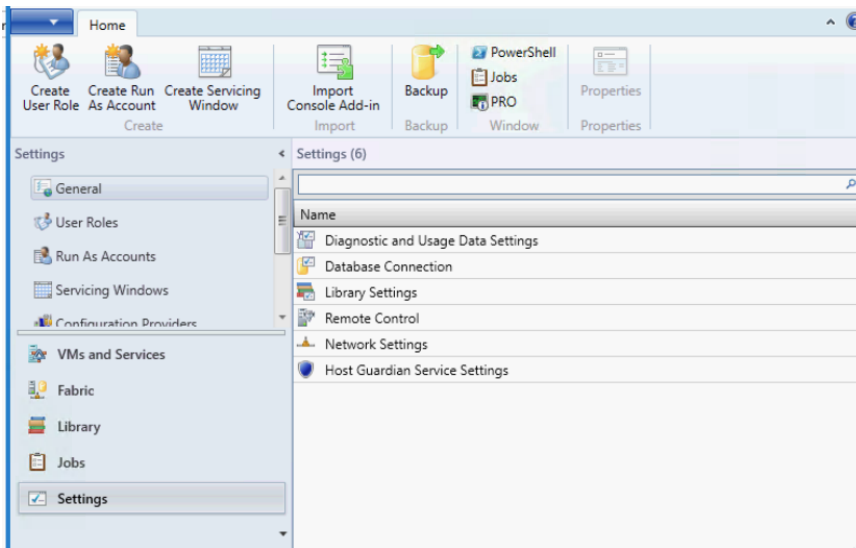
25. Connect to SCVMM using the VMM console.



Configure Run as Account for SCVMM

To configure a run as account for SCVMM, complete the following steps:

1. Click Settings in the VMM console, then click Create Run As Account.



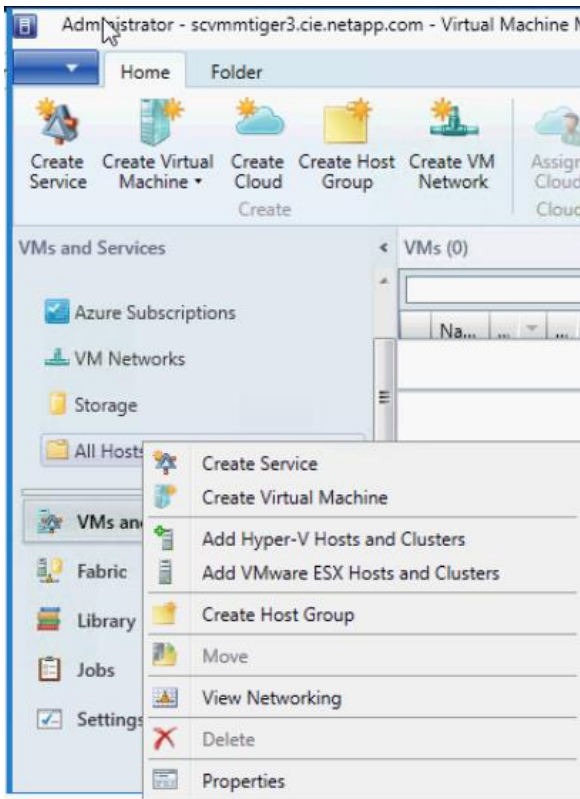
2. Provide an account name and description. Enter the domain and user name of the administrative account you would like to use. Click Finish.

Note: The administrative account must have administrative rights to the Hyper-V hosts. It cannot be the same account as the SCVMM service account.

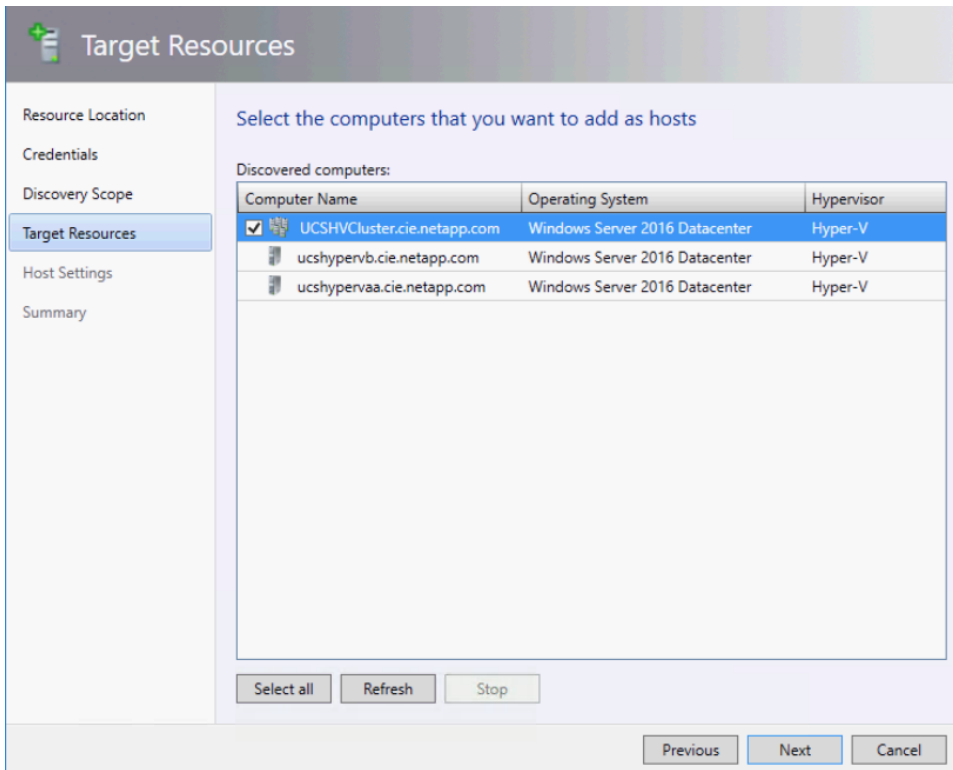
Add Hyper-V Cluster to SCVMM

To add the Hyper-V cluster to SCVMM, complete the following steps:

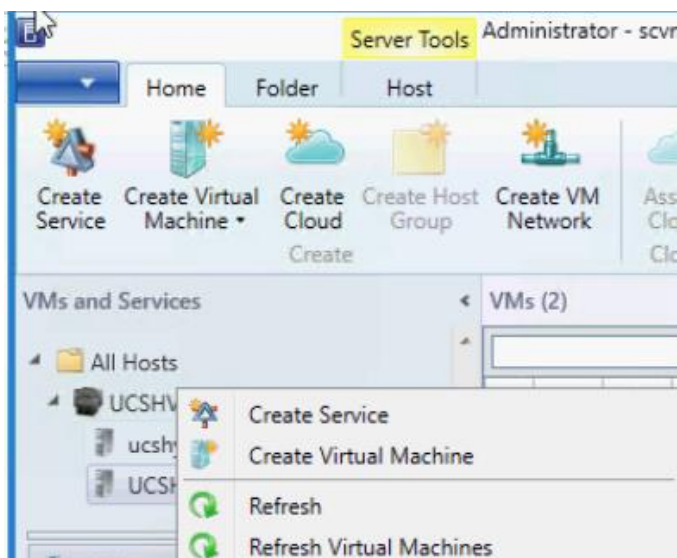
1. Right-click All Hosts and select Add Hyper-V hosts and Clusters.



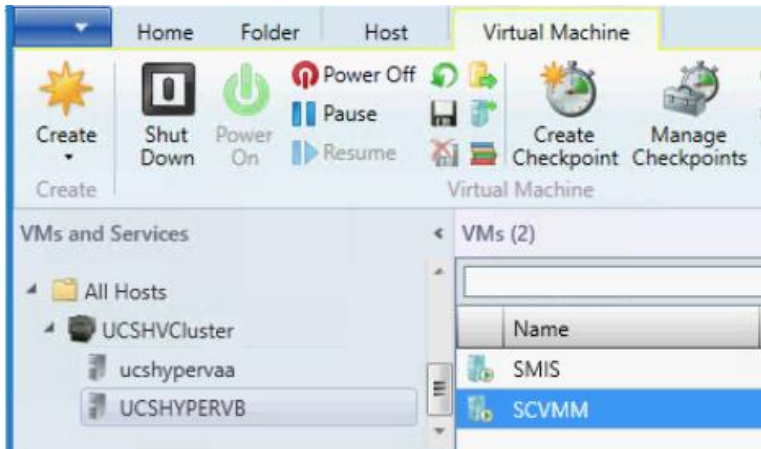
2. In the Resource Location screen, choose Windows Servers in a trusted Active Directory domain.
3. In the Credentials screen, select the run as account you have created. If you have not created a run as account, enter the credentials for the Hyper-V hosts.
4. In the Discovery Scope screen, enter the FQDN of the names of the Hyper-V hosts.
5. In the Target Resources screen, select the box next to the cluster name. Click Next.



6. Click Next at the Host Settings screen.
7. Click Finish at the summary screen.
You can now see the cluster hosts under the VMs and Services, and All Hosts.
8. Right-click the host on which you have built the SCVMM and SMI-S virtual machines on and select Refresh Virtual Machines.



You can now see the virtual machines in SCVMM.



Install SMI-S Provider

To install the SMI-S Provider on the virtual machine you created, complete the following steps:

1. Log in to the NetApp Support site located at <https://mysupport.netapp.com>.
2. Click Downloads, then Software. Scroll down to SMI-S Provider (formerly Data ONTAP SMI-S Agent). Select Windows from the pull-down menu and click Go.
3. Follow the instructions to download the SMI-S Provider 5.2.4.
4. Double-click the downloaded file to launch the installer.
5. Click Next at the Installer Welcome screen.
6. Click Install to install the SMI-S Provider.
You see the status of the installation. This takes several minutes.
7. Click Finish to exit the setup wizard.

Create SMI-S Local User Account

To create a local user account for SMI-S, complete the following steps:

Note: It must be placed in the local administrators group.

1. Open a command prompt.
2. Issue the following command to create the user:

```
Net user local-SMIS-User Password /Add /passwordchg:no /expires:never
```

Password is the password you want to create.

3. Add the account to the local administrators group.

```
Net localgroup add administrators local-SMIS-User /add
```

Configure SMI-S Provider

To configure the SMI-S Provider, complete the following steps:

1. From the Start menu, find the NetApp SMI-S Provider.
2. Click NetApp SMI-S Provider to launch it.
Note: You must be logged in with the user account with which you installed the SMIS-Provider. Commands are case sensitive.
3. Enable authentication with the following command:

```
cimconfig -p -s enableAuthentication=true
```

4. Restart the SMI-S Provider.

```
smis cimserver restart
```

5. Add the local user account you have created.

```
cimuser -a -u Local-SMIS-User -w Password
```

Password is the password you assigned to the user account.

6. Add the SVM.

```
smis addsecure <MgmtIPofSVM> vsadmin
```

7. You are prompted for the password. Enter the password, and the SVM is added.

Register SMI-S Provider with SCVMM

To register the SMI-S provider with SCVMM, complete the following steps:

1. Create a run as account with the SMI-S local user credentials.

Create Run As Account

Provide the details for this Run As account

Name: SMI-S

Description:

User name: Local-SMIS-User
Example: contoso\domainuser or localuser

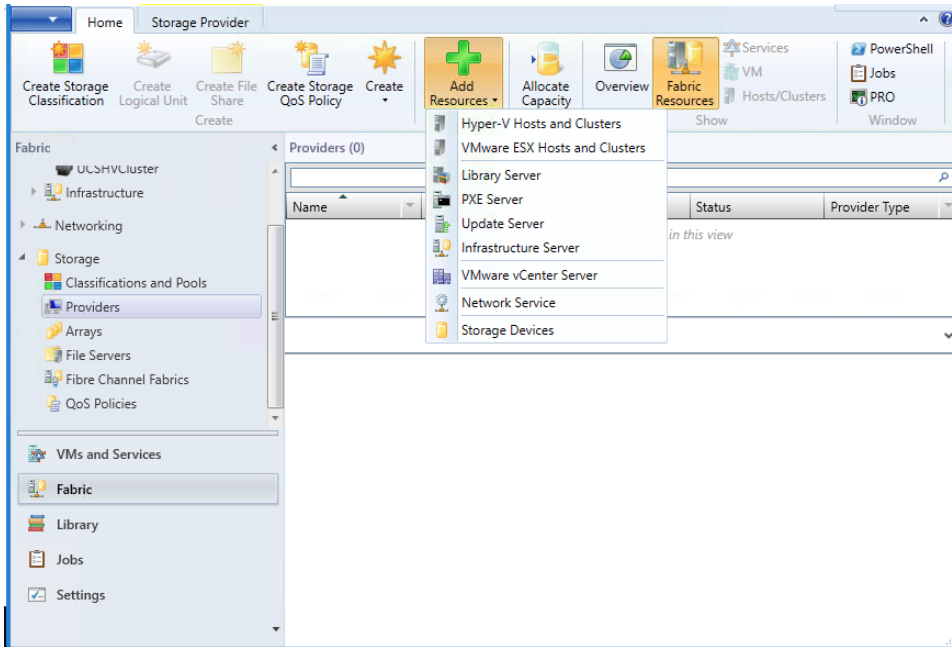
Password: ●●●●●●

Confirm password: ●●●●●●

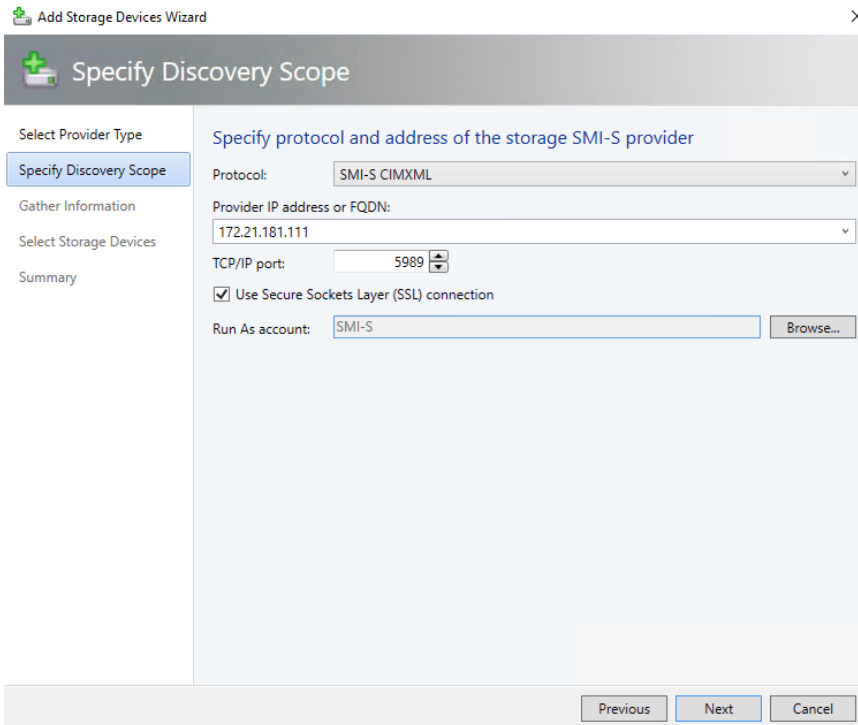
Validate domain credentials

View Script Finish Cancel

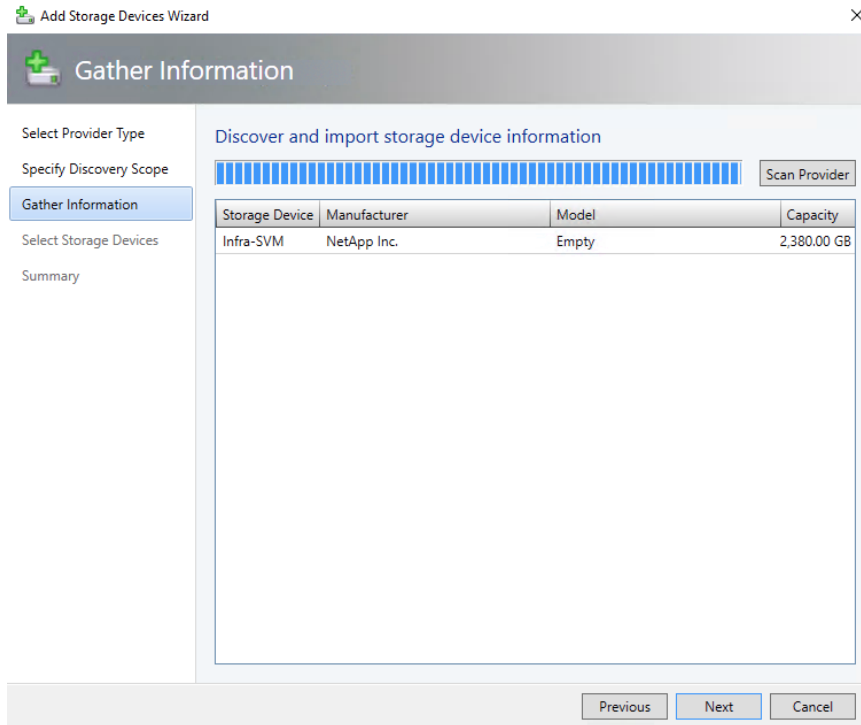
2. In the VMM console, click Fabric in the left navigation pane.
3. Expand Storage.
4. Click Providers.
5. Click Add Resources in the top navigation pane and select Storage Devices.




6. Select the SAN and NAS Devices Discovered and Managed by an SMI-S Provider option on the Storage Provider Type screen.
7. On the Specify Discovery Scope screen, verify that SMI-S CIMXML is selected as the protocol. Enter the IP address of the SMI-S Provider. Select the Use Secure Sockets Layer (SSL) connection checkbox. For run as account, select SMI-S. Click Next.



8. The Gather Information Screen is displayed. Information about the storage array is now being discovered. Click Import if prompted to import the root certificate from the storage array. Click Next after you see the storage array as having been discovered.



9. In the Select Storage Devices screen, click Create classification.
10. Enter HyperVStorage as the name.
11. Select the boxes next to the desired volumes. In this case, we are selecting the provisioning volume we created. Assign the HyperVStorage classification and select All Hosts as the host group. Click Next.



12. Click Finish.

6 Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of other components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

About the Authors

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a solutions architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started

her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has a Bachelor of Science degree in Computer Networking and a Masters of Science degree in Information Security from East Carolina University.

Melissa Palmer, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Melissa Palmer is a solutions architect in the NetApp Infrastructure and Cloud Engineering team. She is also VMware Certified Design Expert (VCDX) #236. Prior to joining the Infrastructure and Cloud Engineering team, Melissa was a systems engineer for NetApp and a VMware engineer for a number of enterprise environments. Melissa has her Bachelor of Engineering and Master of Engineering degrees from Stevens Institute of Technology.

Acknowledgements

The authors would like to acknowledge the following people for their support and contribution to this design:

- Dave Derry, NetApp
- John George, Cisco
- Chris O'Brien, Cisco
- Karthick Radhakrishnan, NetApp
- Glenn Sizemore, NetApp

Version History

Version	Date	Document Version History
Version 1.0	May 2017	Initial release.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE 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 NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE 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 AND USE OR RELIANCE UPON THIS DOCUMENT. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY NETAPP OR ITS PARTNERS.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-1114-DEPLOY-0517