



NetApp Verified Architecture

FlexPod Express with Microsoft Windows Server 2012 R2 Hyper-V: Large Configuration NVA Deployment Guide

Karthick Radhakrishnan, Arvind Ramakrishnan, Glenn Sizemore, NetApp
Chris O'Brien, Cisco
September 2015 | NVA-0018-DEPLOY | Version 1.0

TABLE OF CONTENTS

1	Solution Overview	4
1.1	Solution Technology	4
1.2	Use Case Summary	5
2	Technology Requirements	5
2.1	Hardware Requirements	5
2.2	Software Requirements	5
3	FlexPod Express Cabling Information.....	6
3.1	FlexPod Express Large Configuration	6
4	Deployment Procedures	8
4.1	Cisco Nexus Switch Deployment Procedure	9
4.2	NetApp FAS Storage Deployment Procedure (Part 1)	17
4.3	Cisco UCS C-Series Rack Server Deployment Procedure	36
4.4	NetApp FAS Storage Deployment Procedure (Part 2)	42
4.5	Microsoft Windows Server 2012 R2 Deployment Procedure	43
4.6	Update Display Adapter Driver.....	46
4.7	Install Microsoft Windows Features	46
4.8	Configure Microsoft Windows	50
4.9	Create Windows Failover Cluster.....	57
5	Microsoft System Center 2012 R2 Virtual Machine Manager	60
5.1	Build the SMI-S and SCVMM VMs.....	61
5.2	Configure SMI-S and SCVMM VMs	61
5.3	Add SMB Network Adapter to the SCVMM VM.....	62
5.4	Install Windows Server 2012 R2 on the VMs	62
5.5	Install the NetApp SMI-S Agent in the SMI-S Virtual Machine	63
5.6	Install System Center Virtual Machine Manager	64
6	Bill of Materials	77
7	Conclusion	79
	Acknowledgments	79
	References.....	79

LIST OF TABLES

Table 1)	FlexPod Express large configuration hardware requirements.	5
Table 2)	Software requirements.	5

Table 3) Cabling information for the FlexPod Express large configuration.	6
Table 4) Required VLANs.....	9
Table 5) Windows virtual machines created.	9
Table 6) FAS25XX controller series prerequisites.	18
Table 7) Clustered Data ONTAP software installation prerequisites.	18
Table 8) Cluster create in clustered Data ONTAP prerequisites.	23
Table 9) Cluster join in clustered Data ONTAP prerequisites.	24
Table 10) VM requirements.	60
Table 11) Large configuration components.	77

LIST OF FIGURES

Figure 1) Physical topology of FlexPod Express large configuration.	4
Figure 2) FlexPod Express large configuration cabling diagram.	6

1 Solution Overview

FlexPod[®] Express is a suitable platform for running a variety of virtualization hypervisors as well as bare-metal operating systems and enterprise workloads. FlexPod Express delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements. The large FlexPod Express configuration is a low-cost, standardized infrastructure solution developed to meet the needs of small and midsize businesses. Each configuration provides a standardized base platform capable of running a number of business-critical applications while providing scalability options to enable the infrastructure to grow with the demands of the business.

FlexPod Express:

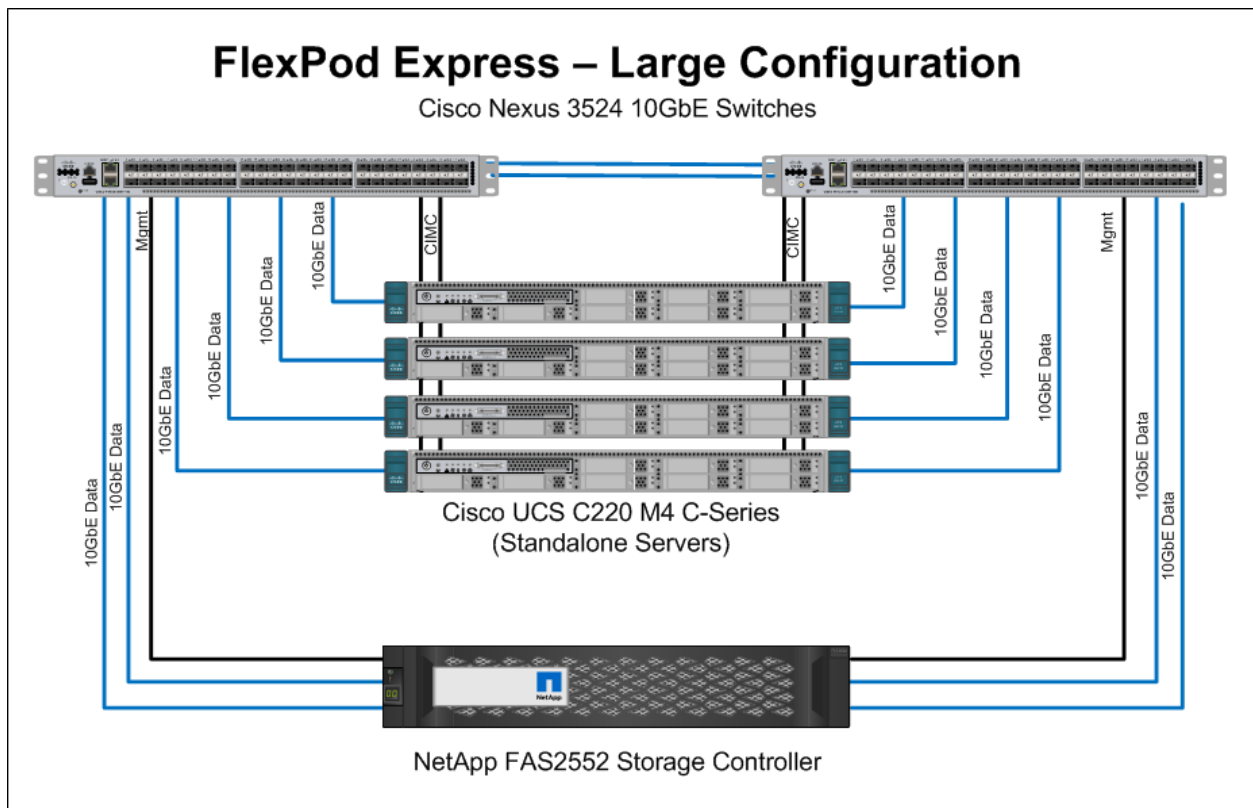
- Combines all application and data needs into one platform
- Is suitable for small to midsize organizations, remote and departmental deployments
- Provides easy infrastructure scaling
- Reduces cost and complexity

1.1 Solution Technology

The large FlexPod Express configuration uses Cisco UCS C-Series rack servers, Cisco Nexus switches (10GbE), and NetApp[®] FAS storage systems (the NetApp clustered Data ONTAP[®] operating system: switchless). This document describes the implementation of Microsoft Windows 2012 R2 on large FlexPod Express offerings. The configurations are based on best practices for each component in the solution architecture to enable a reliable enterprise-class infrastructure.

Figure 1 depicts the topology of the FlexPod Express large configuration.

Figure 1) Physical topology of FlexPod Express large configuration.



1.2 Use Case Summary

This document describes the deployment procedures and best practices to set up a FlexPod Express large configuration with Microsoft Windows Server 2012 R2 Hyper-V as the workload. The server operating system/hypervisor is Microsoft Hyper-V, and an instance of Microsoft System Center 2012 R2 Virtual Machine Manager is installed to manage the Hyper-V instances. The whole infrastructure is supported by NetApp FAS storage systems that serve data over storage area network (SAN) and network-attached storage (NAS) protocols.

2 Technology Requirements

This section details the hardware and software components required to implement the FlexPod Express large configuration.

2.1 Hardware Requirements

Table 1 lists the hardware components required to implement the FlexPod Express large configuration solution.

Table 1) FlexPod Express large configuration hardware requirements.

Layer	Hardware	Quantity
Compute	Cisco UCS C220 M4 rack servers (standalone)	4
Network	Cisco Nexus 3524 switches	2
Storage	NetApp FAS2552 (high-availability pair)	1
Disks	900GB, 10,000-rpm SAS with Advanced Drive Partitioning	24

2.2 Software Requirements

Table 2 lists the software components required to implement the FlexPod Express large configuration.

Table 2) Software requirements.

Layer	Component	Version or Release	Details
Compute	Cisco UCS C220 M4 rack servers	2.0(3j)	Cisco Integrated Management Controller (IMC) software
Network	Cisco Nexus 3524 switches	6.0(2)A6(1)	Cisco NX-OS software
Storage	NetApp FAS2552 high-availability storage	8.3	NetApp Data ONTAP software
Software	Microsoft Windows Server 2012 R2 Hyper-V	2012 R2	Virtualization hypervisor
	System Center Virtual Machine Manager	2012 R2	Virtualization management
	NetApp Data ONTAP SMI-S Agent	5.2	SMI-S Agent
	NetApp Windows Host Utilities Kit	6.0.2	NetApp Plug-in for Windows
	NetApp SnapDrive® for Windows	7.1.1	LUN provisioning and

Layer	Component	Version or Release	Details
			Snapshot management
	NetApp SnapManager® for Hyper-V	2.1	NetApp Plug-in for Hyper-V

3 FlexPod Express Cabling Information

3.1 FlexPod Express Large Configuration

Figure 2 provides a cabling diagram for the FlexPod Express large configuration. Table 3 provides cabling information.

Figure 2) FlexPod Express large configuration cabling diagram.




Table 3) Cabling information for the FlexPod Express large configuration.

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3524 Switch A	Eth1/1	NetApp FAS2552 Storage Controller 01	e0M	1
	Eth1/2	NetApp FAS2552 Storage Controller 02	e0a	2
	Eth1/3	Cisco UCS C220 C-Series Standalone Server 1	LOM1	3

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
	Eth1/4	Cisco UCS C220 C-Series Standalone Server 2	LOM1	4
	Eth1/5	Cisco UCS C220 C-Series Standalone Server 3	LOM1	5
	Eth1/6	Cisco UCS C220 C-Series Standalone Server 4	LOM1	6
	Eth1/7	NetApp FAS2552 Storage Controller 01	e0c	7
	Eth1/8	NetApp FAS2552 Storage Controller 02	e0c	8
	Eth1/9	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 1	9
	Eth1/10	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 1	10
	Eth1/11	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 1	11
	Eth1/12	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port 1	12
	Eth1/13	Cisco Nexus 3524 Switch B	Eth 1/13	13
	Eth1/14	Cisco Nexus 3524 Switch B	Eth 1/14	14

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
Cisco Nexus 3524 Switch B	Eth1/1	NetApp FAS2552 Storage Controller 02	e0M	15
	Eth1/2	NetApp FAS2552 Storage Controller 01	e0a	16
	Eth1/3	Cisco UCS C220 C-Series Standalone Server 1	LOM2	17
	Eth1/4	Cisco UCS C220 C-Series Standalone Server 2	LOM2	18
	Eth1/5	Cisco UCS C220 C-Series Standalone Server 3	LOM2	19
	Eth1/6	Cisco UCS C220 C-Series Standalone Server 4	LOM2	20
	Eth1/7	NetApp FAS2552 Storage Controller 01	e0d	21
	Eth1/8	NetApp FAS2552 Storage Controller 02	e0d	22

	Eth1/9	Cisco UCS C220 C-Series Standalone Server 1	MLOM Port 2	23
	Eth1/10	Cisco UCS C220 C-Series Standalone Server 2	MLOM Port 2	24
	Eth1/11	Cisco UCS C220 C-Series Standalone Server 3	MLOM Port 2	25
	Eth1/12	Cisco UCS C220 C-Series Standalone Server 4	MLOM Port 2	26
	Eth1/13	Cisco Nexus 3524 Switch A	Eth1/13	13
	Eth1/14	Cisco Nexus 3524 Switch A	Eth1/14	14

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2552 Storage Controller 01	e0e	NetApp FAS2552 Storage Controller 02	e0e	29
	e0f	NetApp FAS2552 Storage Controller 02	e0f	30
	ACP	NetApp FAS2552 Storage Controller 02	ACP	31
	SAS 0b	NetApp FAS2552 Storage Controller 01	SAS 0a	27
	SAS 0a	NetApp FAS2552 Storage Controller 01	SAS 0b	28

Local Device	Local Port	Remote Device	Remote Port	Cabling Code
NetApp FAS2552 Storage Controller 02	e0e	NetApp FAS2552 Storage Controller 01	e0e	29
	e0f	NetApp FAS2552 Storage Controller 01	e0f	30
	ACP	NetApp FAS2552 Storage Controller 01	ACP	31
	SAS 0b	NetApp FAS2552 Storage Controller 01	SAS 0a	27
	SAS 0a	NetApp FAS2552 Storage Controller 01	SAS 0b	28

4 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 01 or Component 02. For example, Controller 01 and Controller 02 identify the two NetApp storage controllers that are provisioned in this document, and Switch A and Switch B identify the pair of Cisco Nexus switches that are configured.

Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: Server-1, Server-2, and so on.

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

```
Controller01>vlan create vif0 <<ib_mgmt_vlan_id>>
```

This document is intended to enable you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes. Table 4 describes the VLANs necessary for deployment as outlined in this guide. This table can be completed based on the specific site variables and used in implementing 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 these VLANs. For this validation, a common management VLAN was used.

Table 4) Required VLANs.

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Native VLAN	VLAN to which untagged frames are assigned	2
Management VLAN	VLAN for management interfaces	3051
LiveMigration	VLAN designated for the movement of VMs from one physical host to another	3052
VM traffic	VLAN for virtual machine application traffic	3053
SMB	VLAN for SMB traffic	3054
Cluster	VLAN for cluster communication	3055
iSCSI-A-VLAN	VLAN for iSCSI traffic on Fabric A	3056
iSCSI-B-VLAN	VLAN for iSCSI traffic on Fabric B	3057

Table 5) Windows virtual machines created.

Virtual Machine Description	Host Name
System Center 2012 R2 Virtual Machine Manager	
NetApp SMI-S Agent	

4.1 Cisco Nexus Switch Deployment Procedure

A pair of Cisco Nexus switches that support 10GbE traffic is required to build the network backbone of this FlexPod Express infrastructure.

This document details the implementation of a FlexPod Express solution with the Cisco Nexus 3524 switches. However, these switches can also be replaced with the Cisco Nexus 9000 Series switches.

Cisco Nexus 9000 Series

The Cisco Nexus 9000 Series switches deliver proven high performance and density, low latency, and exceptional power efficiency in a broad range of compact form factors. Operating in Cisco NX-OS

software mode (standalone mode) or in Application Centric Infrastructure (ACI) mode, these switches are ideal for traditional or fully automated data center deployments.

The Cisco Nexus 9000 standalone mode FlexPod Express design consists of a single pair of Cisco Nexus 9000 top-of-rack switches. When leveraging the ACI mode, the Cisco Nexus 9500 and 9300 switches are deployed in a spine-leaf architecture.

ACI is a holistic architecture with centralized automation and policy-driven application profiles. ACI delivers software flexibility with the scalability of hardware performance. Key characteristics of ACI include:

- Simplified automation by an application-driven policy model
- Centralized visibility with real-time application health monitoring
- Open software flexibility for DevOps teams and ecosystem partner integration
- Scalable performance and multi-tenancy in hardware

The future of networking with Cisco ACI is to provide a network that is deployed, monitored, and managed in a fashion that supports DevOps and rapid application change.

Users will also be able to start with the Cisco Nexus 9000 switches in standalone mode and easily migrate to the ACI mode.

Cisco Nexus 3524 Switch Initial Setup

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, and it defines the control-plane policing policy.

The first major decision involves the configuration of the management network for the switches. For FlexPod Express, there are two main options for configuring the mgmt0 interfaces. The first involves configuring and cabling the mgmt0 interfaces into an existing out-of-band network. In this instance, when a management network already exists, all you need are valid IP addresses and the netmask configuration for this network and a connection from the mgmt0 interfaces to this network.

The other option, for installations without a dedicated management network, involves cabling together the mgmt0 interfaces of each Cisco Nexus 3524 switch in a back-to-back configuration. Any valid IP address and network mask can be configured on each mgmt0 interface as long as they are in the same network. Because they are configured back to back with no switch or other device in between, no default gateway configuration is needed, and they should be able to communicate with each other. This link cannot be used for external management access, such as SSH access, but it will be used for the virtual PortChannel (vPC) peer keepalive traffic. To enable SSH management access to the switch, configure the in-band interface VLAN IP address on a switched virtual interface (SVI), as discussed later in this document.

1. Power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both switches, substituting the appropriate values for the switch-specific information.

Cisco Nexus Switch A and Switch B

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

---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no): yes
Enter the password for "admin":<<admin_password>>
Confirm the password for "admin":<<admin_password>>

---- Basic System Configuration Dialog ----

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

of the system.

Please register Cisco Nexus 3500 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus devices must be registered to receive entitled support services.

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): yes
Create another login account (yes/no) [n]: Enter
Configure read-only SNMP community string (yes/no) [n]:Enter
Configure read-write SNMP community string (yes/no) [n]:Enter
Enter the switch name : <<switch_A/B_hostname>>
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:Enter
Mgmt0 IPv4 address : <<switch_A/B_mgmt0_ip_addr>>
Mgmt0 IPv4 netmask : <<switch_A/B_mgmt0_netmask>>
Configure the default gateway? (yes/no) [y]:Enter
```

Note: Do not configure the default gateway if the mgmt ports of the Cisco Nexus 3524 switches are connected back to back.

```
IPv4 address of the default gateway : <<switch_A/B_mgmt0_gateway_ip_addr>>
Enable the telnet service? (yes/no) [n]:Enter
Enable the ssh service? (yes/no) [y]:Enter
Type of ssh key you would like to generate (dsa/rsa) : rsa
Number of key bits <768-2048> : 1024
Configure the ntp server? (yes/no) [n]:Enter
Configure default interface layer (L3/L2) [L2]:Enter
Configure default switchport interface state (shut/noshut) [noshut]:Enter
Configure CoPP System Policy Profile ( default / 12 / 13 ) [default]:Enter
```

```
The following configuration will be applied:
switchname <<switch_A/B_hostname>>
interface mgmt0
ip address <<switch_A/B_mgmt0_ip_addr>> <<switch_A/B_mgmt0_netmask>>
no shutdown
exit
vrf context management
ip route 0.0.0.0/0 <<switch_A/B_mgmt0_gateway_ip_addr>>
exit
no telnet server enable
ssh key rsa 1024 force
ssh server enable
system default switchport
no system default switchport shutdown
policy-map type control-plane copp-system-policy ( default )
```

```
Would you like to edit the configuration? (yes/no) [n]:Enter
Use this configuration and save it? (yes/no) [y]:Enter
```

Upgrade Cisco NX-OS (Optional)

Perform any required software upgrades on the switch at this point in the configuration process. Download and install the latest available Cisco NX-OS software for the Cisco Nexus 3524 switch from the [Cisco software download site](#). There are multiple ways to transfer both the kickstart and system images for Cisco NX-OS to the switch. The most straightforward procedure uses the on-board USB port on the switch. Download the Cisco NX-OS kickstart and system files to a USB drive and plug the USB drive into the external USB port on the Cisco Nexus 3524 switch.

Note: This solution uses Cisco NX-OS software release 6.0(2)A6(1).

1. Copy the files to the local bootflash memory and update the switch by following the procedure shown here.

Cisco Nexus Switch A and Switch B

```
copy usb1:<<kickstart_image_file>> bootflash:
copy usb1:<<system_image_file>> bootflash:
install all kickstart bootflash:<<kickstart_image_file>> system bootflash:<<system_image_file>>
```

Note: The switches will install the updated Cisco NX-OS files and reboot.

Enable Advanced Features

Certain advanced features need to 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 an IP address to be assigned to the interface VLAN (SVI), which enables in-band management communication to the switch, such as through SSH.

Enter configuration mode using the command (`conf t`) and type the following commands to enable the appropriate features on each switch.

Cisco Nexus Switch A and Switch B

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

Perform Global PortChannel Configuration

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 PortChannels by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For that reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

From configuration mode (`conf t`) type the following commands to configure the global PortChannel load-balancing configuration on each switch.

Cisco Nexus Switch A and Switch B

```
port-channel load-balance ethernet source-dest-port
```

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 and 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 and edge, depending on the platform.

The recommended setting for bridge assurance is to consider all ports to be network ports by default. This setting forces the network administrator to review the configuration of each port and helps reveal the most common configuration errors, such as unidentified edge ports or a neighbor that does not have bridge assurance enabled. Also, it is safer to have spanning tree block too many ports than not enough, enabling 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 those cases, you might need to change the port type to make the ports active.

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 will shut 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 each switch.

Cisco Nexus Switch A and Switch B

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

Configure Jumbo Frames

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

From configuration mode (`config t`) type the following commands to enable jumbo frames on each switch.

Cisco Switch A and Switch B

```
policy-map type network-qos jumbo
  class type network-qos class-default
    mtu 9216
system qos
  service-policy type network-qos jumbo
exit
```

Define VLANs

Before configuring individual ports with different VLANs, those Layer 2 VLANs must be defined on the switch. It is also good practice to name the VLANs to help with troubleshooting in the future.

From configuration mode (`config t`) type the following commands to define and give descriptions to the Layer 2 VLANs.

Cisco Switch A and Switch B

```
vlan <<smb_vlan_id>>
  name SMB-VLAN
vlan <<livemigration_vlan_id>>
  name LiveMigration-VLAN
vlan <<cluster_vlan_id>>
  name Cluster-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<ib_mgmt_vlan_id>>
  name IB-MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
exit
```

Configure Access and Management Port Descriptions

As with the assignment of names to the Layer 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

For the small configuration, the descriptions for the management ports and data ports associated with Server-3 and Server-4 are not required because the FlexPod Express small configuration contains only two servers.

From configuration mode (`config t`) in each switch, type the following commands to set up the port descriptions.

FlexPod Express Large Configuration

Enter the following port descriptions for the FlexPod Express large configuration.

Cisco Nexus Switch A

```
int eth1/1
  description Controller-01:e0M
int eth1/2
  description Controller-02:e0a
int eth1/3
  description Server1:LOM1
int eth1/4
  description Server2:LOM1
int eth1/5
  description Server3:LOM1
int eth1/6
  description Server4:LOM1
int eth1/7
  description Controller-01:e0c
int eth1/8
  description Controller-02:e0c
int eth1/9
  description Server-1:VIC Port 1
int eth1/10
  description Server-2:VIC Port 1
int eth1/11
  description Server-3:VIC Port 1
int eth1/12
  description Server-4:VIC Port 1
int eth1/13
  description vPC peer-link NX3524-B:1/13
int eth1/14
  description vPC peer-link NX3524-B:1/14
```

Cisco Nexus Switch B

```
int eth1/1
  description Controller-02:e0M
int eth1/2
  description Controller-01:e0a
int eth1/3
  description Server1:LOM2
int eth1/4
  description Server2:LOM2
int eth1/5
  description Server3:LOM2
int eth1/6
  description Server4:LOM2
int eth1/7
  description Controller-01:e0d
int eth1/8
  description Controller-02:e0d
int eth1/9
  description Server-1:VIC Port 2
int eth1/10
  description Server-2:VIC Port 2
int eth1/11
  description Server-3:VIC Port 2
int eth1/12
  description Server-4:VIC Port 2
int eth1/13
  description vPC peer-link NX3524-A:1/13
int eth1/14
  description vPC peer-link NX3524-A:1/14
```

Configure Server and Storage Management Interfaces

The management interfaces for both the server and 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`) type the following commands to configure the port settings for the management interfaces of both the servers and storage.

Cisco Nexus Switch A

```
int eth1/1-6
  switchport access vlan <<ib_mgmt_vlan_id>>
  spanning-tree port type edge
exit
```

Cisco Nexus Switch B

```
int eth1/1-6
  switchport access vlan <<ib_mgmt_vlan_id>>
  spanning-tree port type edge
int eth1/3-6
  vpc orphan-port suspend
exit
```

Perform Virtual PortChannel Global Configuration

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, be sure to 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) type the following commands to configure the vPC global configuration for switch A.

Cisco Nexus Switch A

```
vpc domain 1
  peer-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 eth1/13-14
  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 <<smb_vlan_id>>,<<livemigration_vlan_id>>, <<cluster_vlan_id>>,
<<vmtraffic_vlan_id>>, <<ib_mgmt_vlan_id>>,<<oob_mgmt_vlan_id>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start
```

From configuration mode (config t) type the following commands to configure the vPC global configuration for switch B.

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 eth1/13-14
  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 <<smb_vlan_id>>,<<livemigration_vlan_id>>, <<cluster_vlan_id>>,
<<vmtraffic_vlan_id>>, <<ib_mgmt_vlan_id>>,<<oob_mgmt_vlan_id>>, <<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 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 completely separate physical switches. Each controller will have two links to each switch, but all four are part of the same vPC and interface group (IFGRP).

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

Cisco Nexus Switch A and Switch B and Controller-01 Configuration

```
int eth1/7
  channel-group 11 mode active
int Po11
  description vPC to Controller-01
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<smb_vlan_id>>,<<ib_mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  vpc 11
  no shut
```

Cisco Nexus Switch A and Switch B and Controller-02 Configuration

```
int eth1/8
  channel-group 12 mode active
int Po12
  description vPC to Controller-02
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<smb_vlan_id>>,<<ib_mgmt_vlan_id>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  vpc 12
  no shut
exit
copy run start
```

Configure Server Connections

The Cisco UCS servers have a two-port Cisco VIC1227 that will be used for data traffic and booting the ESXi operating system using iSCSI. These interfaces will be 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`) type the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A – Cisco UCS Server-1, Server-2, Server-3, and Server-4 Configuration

```
int eth1/9-12
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<smb_vlan_id>>, <<livemigration_vlan_id>>, <<cluster_vlan_id>>,
<<vmtraffic_vlan_id>>, <<ib_mgmt_vlan_id>>, <<iSCSI_A_vlan_id>>
  spanning-tree port type edge trunk
  no shut
```



```
exit
copy run start
```

Cisco Nexus Switch B – Cisco UCS Server-1, Server-2, Server-3, and Server-4 Configuration

```
int eth1/9-12
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<smb_vlan_id>>, <<livemigration_vlan_id>>, <<cluster_vlan_id>>,
<<vmtraffic_vlan_id>>, <<ib_mgmt_vlan_id>>, <<iSCSI_B_vlan_id>>
  spanning-tree port type edge trunk
  vpc orphan-port suspend
  no shut
exit
copy run start
```

Perform in-Band Management SVI Configuration

In-band management using SSH in the FlexPod Express environment is handled by an SVI. To configure the in-band management on each switch, you must configure an IP address on the interface VLAN and set up a default gateway.

From configuration mode (`config t`) type the following commands to configure the Layer 3 SVI for management purposes.

Cisco Nexus Switch A

```
int Vlan <<oob_mgmt_vlan_id>>
ip address <<outofband_mgmt_ip_address_A>>/<<outofband_mgmt_netmask>>
no shut
ip route 0.0.0.0/0 <<outofband_mgmt_gateway>>
```

Cisco Nexus Switch B

```
int Vlan <<oob_mgmt_vlan_id>>
ip address <<outofband_mgmt_ip_address_B>>/<<outofband_mgmt_netmask>>
no shut
ip route 0.0.0.0/0 <<outofband_mgmt_gateway>>
```

4.2 NetApp FAS Storage Deployment Procedure (Part 1)

This section describes the NetApp FAS storage deployment procedure.

Controller FAS25xx Series

NetApp Hardware Universe

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

1. Check the [NetApp Hardware Universe](#) at the [NetApp Support](#) site to make sure that the hardware and software components are supported with the version of Data ONTAP you plan to install.
2. Access the [Hardware Universe](#) Application to view the System Configuration guides. Click the Controllers tab to view compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.
3. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Table 6) FAS25XX controller series prerequisites.

Controller FAS255X Series Prerequisites
<p>To plan the physical location of the storage systems, refer to the following sections in the NetApp Hardware Universe:</p> <ul style="list-style-type: none">• Electrical requirements• Supported power cords• On-board ports and cables <p>Refer to the site requirements guide replacement tutorial to find NetApp FAS platform information using the Hardware Universe.</p>

Storage Controllers

Follow the physical installation procedures for the controllers in the [FAS25xx documentation](#) available at the NetApp Support site.

NetApp Clustered Data ONTAP 8.3

Complete the Configuration Worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [Clustered Data ONTAP 8.3 Software Setup Guide](#) at the [NetApp Support](#) site.

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

Table 7) Clustered Data ONTAP software installation prerequisites.

Cluster Detail	Cluster Detail Value
Cluster node 01 IP address	<<var_node01_mgmt_ip>>
Cluster node 01 netmask	<<var_node01_mgmt_mask>>
Cluster node 01 gateway	<<var_node01_mgmt_gateway>>
Cluster node 02 IP address	<<var_node02_mgmt_ip>>
Cluster node 02 netmask	<<var_node02_mgmt_mask>>
Cluster node 02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3 URL	<<var_url_boot_software>>

Node 01

To configure node 01, complete the following steps:

1. Connect to the storage system console port. You will see a Loader-A prompt. 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. Set boot monitor defaults.

```
Set-defaults
```

3. Allow the system to boot.

```
autoboot
```

4. Press Ctrl-C when prompted.

Note: If Data ONTAP 8.3 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3 is the version being booted, select option 8 and `yes` to reboot the node and go to step 14.

5. To install new software, select option 7.

```
7
```

6. Answer `yes` to perform an upgrade.

```
y
```

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

```
e0M
```

8. Select `yes` to reboot now.

```
y
```

9. After reboot, enter the IP address, network mask, and default gateway for e0M in their respective places.

```
<<var_node01_mgmt_ip>> <<var_node01_mgmt_mask>> <<var_node01_mgmt_gateway>>
```

10. Enter the URL where the software is located.

Note: This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

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

```
y
```

13. Enter `yes` 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.

14. When you see `Press Ctrl-C` for the boot menu:

```
Ctrl - C
```

15. Select option 5 to enter into maintenance mode.

```
5
```

16. Remove the disk ownership. Enter `Y` to remove the disk ownership and offline the existing volumes or aggregates.

```
disk remove_ownership
```

```
All disks owned by system ID 536902178 will have their ownership information removed.  
Do you wish to continue? y
```

```
Volumes must be taken offline. Are all impacted volumes offline(y/n)? y  
Removing the ownership of aggregate disks may lead to partition of aggregates between high-  
availability pair.
```

```
Do you want to continue(y/n)? y
```

17. Halt the node. The node will enter the Loader prompt.

```
halt
```

18. Start Data ONTAP.

```
autoboot
```

19. Press Ctrl-C for the boot menu:

```
Ctrl - C
```

20. Select option 4 for a clean configuration and initialize all disks.

```
4
```

21. Answer yes to Zero disks, reset config and install a new file system.

```
y
```

22. Enter yes to erase all the data on the disks.

```
y
```

Note: The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. After initialization is complete, the storage system reboots. You can continue to the node 02 configuration while the disks for node 01 zero.

Node 02

To configure node 02, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. 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. Set the boot monitor defaults.

```
set-defaults
```

3. Allow the system to boot.

```
autoboot
```

4. Press Ctrl-C when prompted.

```
Ctrl-C
```

Note: If Data ONTAP 8.3 is not the version of software being booted, continue with the following steps to install new software. If Data ONTAP 8.3 is the version being booted, select option 8 and *yes* to reboot the node and go to step 14.

5. To install new software first, select option 7.

```
7
```

6. Answer yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

8. Select yes to reboot now.

```
y
```

9. Enter the IP address, network mask, and default gateway for e0M in their respective places.

```
<<var_node02_mgmt_ip>> <<var_node02_mgmt_mask>> <<var_node02_mgmt_gateway>>
```

10. Enter the URL where the software is located.

Note: This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

12. Select yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

13. Select yes 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.

14. Press Ctrl-C for the boot menu:

```
Ctrl - C
```

15. Select option 5 to enter maintenance mode.

```
5
```

16. Remove the disk ownership. Enter Y to do so and offline the existing volumes or aggregates.

```
disk remove_ownership
```

```
All disks owned by system ID 536902178 will have their ownership information removed. Do you wish to continue? y
```

```
Volumes must be taken offline. Are all impacted volumes offline(y/n)? y
```

```
Removing the ownership of aggregate disks may lead to partition of aggregates between high-availability pair.
```

```
Do you want to continue(y/n)? y
```

17. Halt the node. The node will enter the Loader prompt.

```
halt
```

18. Start Data ONTAP.

```
autoboot
```

19. Press Ctrl-C for the boot menu:

```
Ctrl - C
```

20. Select option 4 for clean configuration and initialize all disks.

```
4
```

21. Answer yes to Zero disks, reset config and install a new file system.

```
y
```

22. Enter yes to erase all the data on the disks.

```
y
```

Note: The initialization and creation of the root volume can take 90 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Node Setup in Clustered Data ONTAP

From a console port program attached to the storage controller A (Node 01) console port, execute the node setup script. This script will come up when Data ONTAP 8.3 first boots on a node.

1. Follow the prompts below:

```
Welcome to node setup.

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

To accept a default or omit a question, do not enter a value.

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

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address: <<var_node01_mgmt_ip>>
Enter the node management interface netmask: <<var_node01_mgmt_mask>>
Enter the node management interface default gateway: <<var_node01_mgmt_gateway>>
A node management interface on port e0M with IP address <<var_node01_mgmt_ip>> has been created.

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility
from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.
```

2. Press Return and log in to the node using the admin user ID and no password to get a node command prompt.

```
::> storage failover modify -mode ha
Mode set to HA. Reboot node to activate HA.

::> system node reboot

Warning: Are you sure you want to reboot node "localhost"? {y|n}: y
```

3. After reboot, go through the node setup procedure with preassigned values.

```
Welcome to node setup.

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

To accept a default or omit a question, do not enter a value.

Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
```

```

Enter the node management interface default gateway [<<var_node01_mgmt_gateway>>]: Enter

This node has its management address assigned and is ready for cluster setup.

To complete cluster setup after all nodes are ready, download and run the System Setup utility
from the NetApp Support Site and use it to discover the configured nodes.

For System Setup, this node's management address is: <<var_node01_mgmt_ip>>.

Alternatively, you can use the "cluster setup" command to configure the cluster.

```

4. Log in to the node with the admin user and no password.
5. Repeat this entire procedure for node 2 of the storage cluster.

Cluster Create in Clustered Data ONTAP

Table 8) Cluster create in clustered Data ONTAP prerequisites.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Clustered Data ONTAP base license	<<var_cluster_base_license_key>>
Cluster management IP address	<<var_clustermgmt_ip>>
Cluster management netmask	<<var_clustermgmt_mask>>
Cluster management port	<<var_clustermgmt_port>>
Cluster management gateway	<<var_clustermgmt_gateway>>
Cluster node01 IP address	<<var_node01_mgmt_ip>>
Cluster node01 netmask	<<var_node01_mgmt_mask>>
Cluster node01 gateway	<<var_node01_mgmt_gateway>>

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01.

Using the console session to node 01, the Cluster Setup wizard is brought up by typing `cluster setup`.

```

cluster setup
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.
Do you want to create a new cluster or join an existing cluster? {create, join}:

```

Note: If a login prompt appears instead of the Cluster Setup wizard, start the wizard by using the factory default settings and then enter the `cluster setup` command.

To create a new cluster, complete the following steps:

1. Run the following command to create a new cluster:

```
create
```

2. Type `no` for the single-node cluster option.

Do you intend for this node to be used as a single node cluster? {yes, no} [no]: no

3. Type no for cluster network using network switches.

Will the cluster network be configured to use network switches? [yes]:no

4. The system defaults are displayed. Enter yes to use the system defaults. Use the following prompts to configure the cluster ports.

Existing cluster interface configuration found:

Port	MTU	IP	Netmask
e0d	9000	169.254.128.103	255.255.0.0
e0f	9000	169.254.52.249	255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:

5. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username "admin") password: <<var_password>>
Retype the password: <<var_password>>
Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>
Enter an additional license key []:<<var_cifs_license>>
```

Note: The cluster is created. This can take a minute or two.

Note: For this validated architecture NetApp recommends installing license keys for NetApp SnapRestore®, NetApp FlexClone®, and NetApp SnapManager Suite technologies. Additionally, install all required storage protocol licenses (CIFS, iSCSI). After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0a]: e0a
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway: <<var_clustermgmt_gateway>>
```

6. Enter the DNS domain name.

```
Enter the DNS domain names:<<var_dns_domain_name>>
Enter the name server IP addresses:<<var_nameserver_ip>>
```

Note: If you have more than one name server IP address, separate the IP addresses with a comma.

7. Set up the node.

```
Where is the controller located []:<<var_node_location>>
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
Enter the node management interface default gateway [<<var_node01_mgmt_gateway>>]: Enter
```

Note: The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet.

Cluster Join in Clustered Data ONTAP

Table 9) Cluster join in clustered Data ONTAP prerequisites.

Cluster Detail	Cluster Detail Value
Cluster name	<<var_clustername>>
Cluster management IP address	<<var_clustermgmt_ip>>

Cluster Detail	Cluster Detail Value
Cluster node02 IP address	<<var_node02_mgmt_ip>>
Cluster node02 netmask	<<var_node02_mgmt_mask>>
Cluster node02 gateway	<<var_node02_mgmt_gateway>>

The first node in the cluster performs the `cluster create` operation. All other nodes perform a `cluster join` operation. The first node in the cluster is considered node 01, and the node joining the cluster in this example is node 02.

To join the cluster, complete the following steps from the console session of node 02:

1. If prompted, enter `admin` in the login prompt.

```
admin
```

2. Start the Cluster Setup wizard by typing `cluster setup`.

```
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.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
```

Note: If a login prompt is displayed instead of the Cluster Setup wizard, start the wizard by logging in using the factory default settings and then enter the `cluster setup` command.

3. Run the following command to join a cluster:

```
join
```

4. Data ONTAP detects the existing cluster and agrees to join the same cluster. Follow these prompts to join the cluster:

```
Existing cluster interface configuration found:

Port      MTU      IP          Netmask
e0d       9000     169.254.144.37 255.255.0.0
e0f       9000     169.254.134.33 255.255.0.0

Do you want to use this configuration? {yes, no} [yes]:
```

5. The steps to join a cluster are displayed.

```
Enter the name of the cluster you would like to join [<<var_clustername>>]: Enter
```

Note: The node should find the cluster name. The cluster joining can take a few minutes.

6. Set up the node.

```
Enter the node management interface port [e0M]: e0M
Enter the node management interface IP address [<<var_node02_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node02_netmask>>]: Enter
Enter the node management interface default gateway [<<var_node02_gw>>]: Enter
```

Note: The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet.

Log in to the Cluster

Open an SSH connection using the cluster IP or host name and log in as the admin user with the password provided during setup.

Zero All Spare Disks

To zero all spare disks in the cluster, complete the following step:

1. Run the following command:

```
disk zerospares
```

Set on-Board UTA2 Ports Personality

2. Verify the Current Mode and Current Type of the ports by using the `ucadmin show` command.

```
iceel-stcl:> ucadmin show
Node      Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Admin Status
-----
iceel-stcl-01  0c      cna      target      -            -            online
iceel-stcl-01  0d      cna      target      -            -            online
iceel-stcl-01  0e      cna      target      -            -            online
iceel-stcl-01  0f      cna      target      -            -            online
iceel-stcl-02  0c      cna      target      -            -            online
iceel-stcl-02  0d      cna      target      -            -            online
iceel-stcl-02  0e      cna      target      -            -            online
iceel-stcl-02  0f      cna      target      -            -            online
8 entries were displayed.
```

3. 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.

Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, enter:

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

Setting Up Management Broadcast Domain

To set up the Default broadcast domain for management network interfaces, complete the following step:

The unused ports and data ports will need to be removed from the Default broadcast domain.

Port `e0b` is unused and `e0c` and `e0d` are used for iSCSI and SMB data traffic.

```
broadcast-domain remove-ports -broadcast-domain Default -ports <<var_node01>>:e0b,
<<var_node01>>:e0c, <<var_node01>>:e0d, <<var_node02>>:e0b, <<var_node02>>:e0c,
<<var_node02>>:e0d
```

Enable Cisco Discovery Protocol (CDP) in Clustered Data ONTAP

To enable CDP on the NetApp storage controllers, complete the following step:

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

1. Enable CDP on Data ONTAP.

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

Set Auto-Revert on Cluster Management

To set the `auto-revert` parameter on the cluster management interface, complete the following step:

1. Run the following command:

```
network interface modify -vserver <<var_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, complete the following step:

1. Run the following commands:

```
system service-processor network modify -node <<var_node01>> -address-family IPv4 -enable true -  
dhcp none -ip-address <<var_node01_sp_ip>> -netmask <<var_node01_sp_mask>> -gateway  
<<var_node01_sp_gateway>>
```

```
system service-processor network modify -node <<var_node02>> -address-family IPv4 -enable true -  
dhcp none -ip-address <<var_node02_sp_ip>> -netmask <<var_node02_sp_mask>> -gateway  
<<var_node02_sp_gateway>>
```

Note: The Service Processor IP addresses should be in the same subnet as the Node Management IP addresses.

Enable Storage Failover in Clustered Data 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 the nodes `<<var_node01>>` and `<<var_node02>>` must be capable of performing a takeover. Go to step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_node01>> -enabled true
```

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

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

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

```
cluster ha show
```

4. Go to step 6 if high availability is configured.
5. Enable HA mode only for the two-node cluster.

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

```
cluster ha modify -configured true
```

```
Do you want to continue? {y|n}: y
```

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

```
storage failover hwassist show
storage failover modify -hwassist-partner-ip <<var_node02_mgmt_ip>> -node <<var_node01>>
storage failover modify -hwassist-partner-ip <<var_node01_mgmt_ip>> -node <<var_node02>>
```

Create Jumbo Frame MTU Broadcast Domain in Clustered Data ONTAP

To create a data broadcast domain with an MTU of 9000, complete the following step:

1. Create broadcast domain on Data ONTAP.

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

Disable Flow Control on UTA2 Ports

A NetApp best practice is to disable flow control on all the UTA2 ports that are connected to external devices.

To disable flow control, run the following command:

```
net port modify -node <<controller01>> -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 <<controller01>> -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 <<controller01>> -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 <<controller01>> -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 <<controller02>> -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 <<controller02>> -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 <<controller02>> -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 <<controller02>> -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 Clustered Data ONTAP

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

1. From the cluster prompt, complete the following steps.

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0d
ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode multimode_lacp
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0d
```

Configure Jumbo Frames in Clustered Data ONTAP

1. To configure a clustered Data ONTAP network port to use jumbo frames (which usually have a maximum transmission unit [MTU] of 9,000 bytes), run the following command from the cluster shell:

```
nbice-fpel::> network port modify -node <<var_node01>> -port a0a -mtu 9000

Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y

nbice-fpel::> network port modify -node <<var_node02>> -port a0a -mtu 9000

Warning: Changing the network port settings will cause a several second interruption in carrier.
Do you want to continue? {y|n}: y
```

Create VLAN in Clustered Data ONTAP

1. Create SMB VLAN ports and add them to the Data Broadcast Domain.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<smb_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<smb_vlan_id>>

broadcast-domain add-ports -broadcast-domain Infra_SMB -ports <<var_node01>>:a0a-<<smb_vlan_id>>,
<<var_node02>>:a0a-<<smb_vlan_id>>
```

2. Create iSCSI VLAN ports and add them to the Data Broadcast Domain.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_node01>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<var_iscsi_vlan_B_id>>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports <<var_node01>>:a0a-
<<var_iscsi_vlan_A_id>>,<<var_node02>>:a0a-<<var_iscsi_vlan_A_id>>

broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports <<var_node01>>:a0a-
<<var_iscsi_vlan_B_id>>,<<var_node02>>:a0a-<<var_iscsi_vlan_B_id>>
```

3. Create IB-MGMT-VLAN ports.

```
network port vlan create -node <<var_node01>> -vlan-name a0a-<<ib_mgmt_vlan_id>>
network port vlan create -node <<var_node02>> -vlan-name a0a-<<ib_mgmt_vlan_id>>
```

Create Aggregates in Clustered Data ONTAP

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

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -node <<var_node01>> -diskcount <<var_num_disks>>
aggr create -aggregate aggr1_node02 -node <<var_node02>> -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 initially; 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 aggregate creation status. Do not proceed until `aggr1_node1` is online.

2. Disable NetApp Snapshot[®] copies for the data aggregate recently created.

```
node run <<var_node01>> aggr options aggr1_node01 nosnap on
```

```
node run <<var_node02>> aggr options aggr1_node02 nosnap on
```

3. Delete any existing Snapshot copies for the two data aggregates.

```
node run <<var_node01>> snap delete -A -a -f aggr1_node01  
node run <<var_node02>> snap delete -A -a -f aggr1_node02
```

4. Rename the root aggregate on node 01 to match the naming convention for this aggregate on node 02.

```
aggr show  
aggr rename -aggregate aggr0 -newname <<var_node01_rootaggrname>>
```

Configure NTP in Clustered Data ONTAP

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

1. To set the time zone for the cluster, run the following command:

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

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

2. To set the date for the cluster, run the following command:

```
date <ccyyymmddhhmm.ss>
```

Note: The format for the date is <[Century][Year][Month][Day][Hour][Minute].[Second]>; for example, 201505181453.17.

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

```
cluster time-service ntp server create -server <<var_global_ntp_server_ip>>
```

Configure SNMP in Clustered Data ONTAP

To configure 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, such as a DFM server or another fault management system.

```
snmp traphost add <<var_oncommand_server_fqdn>>
```

Configure SNMPv1 in Clustered Data ONTAP

To configure SNMPv1, complete the following step:

1. Set the shared secret plain-text password, which is 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 will remove them.

Configure SNMPv3 in Clustered Data ONTAP

SNMPv3 requires that a user be defined and configured for authentication. To configure SNMPv3, complete the following step:

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 Clustered Data ONTAP

The NetApp AutoSupport® tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, complete the following step:

1. 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 Storage Virtual Machine (Vserver)

To create an infrastructure Vserver, complete the following steps:

1. Run the `vserver create` command.

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

2. Add the data aggregate to the `Infra_Vserver` aggregate list for the NetApp Virtual Storage Console.

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

3. Select the Vserver data protocols to configure, leaving NFS, FCP, and iSCSI.

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

Create Load Sharing Mirror of Vserver Root Volume in Clustered Data ONTAP

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

```
volume create -vserver Infra-SVM -volume rootvol_m01 -aggregate aggr1_node01 -size 1GB -type DP
volume create -vserver Infra-SVM -volume rootvol_m02 -aggregate aggr1_node02 -size 1GB -type DP
```

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

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

3. Create the mirroring relationships.

```
snapmirror create -source-path //Infra-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.

```
snapmirror initialize-ls-set -source-path //Infra-SVM/rootvol
```

Create iSCSI Service in Clustered Data ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on each Vserver. This command also starts the iSCSI service and sets the iSCSI IQN for the Vserver.

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

Configure HTTPS Access in Clustered Data 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 Vserver shown, the certificate common name should match the DNS FQDN of the Vserver. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a Certificate Authority (CA). 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 Infra-SVM and the cluster Vserver. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...
Example: security certificate create -common-name infra-svm.ciscorobo.com -type server -size
2048 -country US -state "California" -locality "San Jose" -organization "Cisco" -unit "UCS" -
email-addr "abc@cisco.com" -expire-days 365 -protocol SSL -hash-function SHA256 -vserver Infra-
SVM
```

5. To obtain the values for the parameters that would be required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created by using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...
Example: security ssl modify -vserver clus -server-enabled true -client-enabled false -ca
clus.ciscorobo.com -serial 55243646 -common-name clus.ciscorobo.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>>
```

8. It is normal for some of these commands to return an error message stating that the entry does not exist.
9. Change back to normal admin privilege level and set up to allow Vserver logs to be available by web.

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


Create FlexVol Volume in Clustered Data ONTAP

To create a NetApp FlexVol® volume, complete the following step:

1. To create a FlexVol volume you need the volume's name and size and the aggregate on which it will exist. Create two datastore volumes and a server boot volume.

```
volume create -vsriver Infra-SVM -volume witness -aggregate aggr1_controller01 -size 100GB -state
online -policy default -space-guarantee none -percent-snapshot-space 0 -junction-path /quorum

volume create -vsriver Infra-SVM -volume sc_sql_db -aggregate aggr1_controller01 -size 1TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0 -junction-path /sc_sql_db

volume create -vsriver Infra-SVM -volume scvmm_pool0 -aggregate aggr1_controller02 -size 4TB -
state online -policy default -space-guarantee none -percent-snapshot-space 0 -junction-path
/scvmm_pool0

volume create -vsriver Infra-SVM -volume boot_luns -aggregate aggr1_node01 -size 1TB -state
online -policy default -space-guarantee none -percent-snapshot-space 0
```

Create LUNs in Clustered Data ONTAP

To create LUNs, complete the following step:

1. Create four boot LUNs.

```
lun create -vsriver Infra-SVM -volume boot_luns -lun VMHost-Infra-01 -size 200GB -ostype
windows_2008 -space-reserve disabled
lun create -vsriver Infra-SVM -volume boot_luns -lun VMHost-Infra-02 -size 200GB -ostype
windows_2008 -space-reserve disabled
lun create -vsriver Infra-SVM -volume boot_luns -lun VMHost-Infra-03 -size 200GB -ostype
windows_2008 -space-reserve disabled
lun create -vsriver Infra-SVM -volume boot_luns -lun VMHost-Infra-04 -size 200GB -ostype
windows_2008 -space-reserve disabled
```

Enable Deduplication in Clustered Data ONTAP

To enable deduplication on appropriate volumes, complete the following step:

1. Run the following commands:

```
volume efficiency on -vsriver Infra-SVM -volume scvmm_pool0
volume efficiency on -vsriver Infra-SVM -volume boot_luns
```

Create iSCSI LIFs in Clustered Data ONTAP

To create iSCSI LIFs, complete the following step:

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

```
network interface create -vsriver Infra-SVM -lif iscsi_lif01a -role data -data-protocol iscsi -
home-node <<var_node01>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_node01_iscsi_lif01a_ip>> -netmask <<var_node01_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vsriver Infra-SVM -lif iscsi_lif01b -role data -data-protocol iscsi -
home-node <<var_node01>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
<<var_node01_iscsi_lif01b_ip>> -netmask <<var_node01_iscsi_lif01b_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vsriver Infra-SVM -lif iscsi_lif02a -role data -data-protocol iscsi -
home-node <<var_node02>> -home-port a0a-<<var_iscsi_vlan_A_id>> -address
<<var_node02_iscsi_lif01a_ip>> -netmask <<var_node02_iscsi_lif01a_mask>> -status-admin up -
failover-policy disabled -firewall-policy data -auto-revert false

network interface create -vsriver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -
home-node <<var_node02>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address
```

```
<<var_node02_iscsi_lif01b_ip>> -netmask <<var_node02_iscsi_lif01b_mask>> -status-admin up -  
failover-policy disabled -firewall-policy data -auto-revert false  
  
network interface show
```

Create an SMB LIF in Clustered Data ONTAP

1. Create an SMB logical interface (LIF).

```
network interface create -vserver Infra-SVM -lif smb_lif01 -role data -data-protocol cifs -home-  
node <<var_node01>> -home-port a0a-<<var_smb_vlan_id>> -address <<var_node01_smb_lif_ip>> -  
netmask <<var_node01_smb_lif_mask>> -status-admin up -failover-policy broadcast-domain-wide -  
firewall-policy data -auto-revert true  
  
network interface create -vserver Infra-SVM -lif smb_lif02 -role data -data-protocol cifs -home-  
node <<var_node02>> -home-port a0a-<<var_smb_vlan_id>> -address <<var_node02_smb_lif_ip>> -  
netmask <<var_node02_smb_lif_mask>> -status-admin up -failover-policy broadcast-domain-wide -  
firewall-policy data -auto-revert true  
  
network interface show
```

Note: NetApp recommends creating a new LIF for each datastore.

Add Infrastructure Vserver Administrator

To add the infrastructure Vserver administrator and Vserver administration logical interface in the out-of-band management network, complete the following step:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif vsmsgmt -role data -data-protocol none -home-node  
<<var_node02>> -home-port e0a -address <<var_vserver_mgmt_ip>> -netmask  
<<var_vserver_mgmt_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-  
policy mgmt -auto-revert true
```

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

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

```
network routing-groups route create -vserver Infra-SVM -routing-group d<<var_out-band-network>> -  
destination 0.0.0.0/0 -gateway <<var_in-band_gateway>>  
  
network route show
```

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

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

Configure SMB in Clustered Data ONTAP

Run all commands to configure SMB on the Vserver.

1. Secure the default rule for the default export policy and create the FlexPod export policy.

```
vserver export-policy rule modify -vserver Infra-SVM -policyname default -ruleindex 1 -rorule  
never -rwrule never -superuser none
```

2. Create a new rule for the FlexPod export policy.

Note: For each of the Hyper-V, SCVMM, and SQL hosts, create a rule. Each host will have its own rule index. Your first Hyper-V host will have rule index 1, the second Hyper-V host will have rule index

2, and so on. Alternatively, you can assign the subnet by allocating the entire network by a single rule.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 1 -protocol cifs -clientmatch <<var_vmhost_host1_smb_ip>> -rorule sys -rwrule sys -superuser sys -allow-suid false
```

3. Assign the default export policy to the infrastructure Vserver root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```

4. Create the CIFS service and add it to Active Directory.

```
vserver cifs create -vserver Infra-SVM -cifs-server InfraSVM -domain <<var_dnsdomain>>
```

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

Enter the user name: adminXX

Enter the password: XXnetapp!

Create SMB Shares

1. Create a qtree in the SCVMM pool to house the infrastructure VMs and SCVMM library.

```
qtree create -volume scvmm_pool0 -qtree infrastructure -security-style ntfs -vserver Infra-SVM
qtree create -volume scvmm_pool0 -qtree vmmlibrary -security-style ntfs -vserver Infra-SVM
```

Note: These are colocated to enable ODX to rapidly deploy VMs utilizing FlexClone for files.

2. Create a qtree in the witness volume to house the file share witness.

```
qtree create -volume witness -qtree hyperv -security-style ntfs -vserver Infra-SVM
```

3. Create a qtree in the System Center SQL volume to house the SCVMM database.

```
qtree create -volume sc_sql_db -qtree scvmm -security-style ntfs -vserver Infra-SVM
```

4. Create the qtree quota policy for the infrastructure VM share and SCVMM library.

```
quota policy rule create -vserver Infra-SVM -policy-name default -volume scvmm_pool0 -type tree -
disk-limit 500g -target infrastructure
quota policy rule create -vserver Infra-SVM -policy-name default -volume scvmm_pool0 -type tree -
disk-limit 500g -target vmmlibrary
```

5. Create the qtree quota policy for the file share witness.

```
quota policy rule create -vserver Infra-SVM -policy-name default -volume witness -type tree -
disk-limit 5g -target hyperv
```

6. Create the qtree quota policy for the SCVMM database share.

```
quota policy rule create -vserver Infra-SVM -policy-name default -volume sc_sql_db -type tree -
disk-limit 500g -target scvmm
```

7. Create the SMB share to house the infrastructure VMs and SCVMM library.

```
share create -share-name infrastructure -path /scvmm_pool0/infrastructure -share-properties
browsable,continuously-available -vserver Infra-SVM
share create -share-name vmmlibrary -path /scvmm_pool0/vmmlibrary -share-properties
browsable,continuously-available -vserver Infra-SVM
```

8. Create the SMB share to house the file share witness.

```
share create -share-name hyperv-witness -path /quorum/hyperv -share-properties browsable -vserver
Infra-SVM
```

9. Create the SMB share to house the SCVMM database.

```
share create -share-name scvmddb -path /sc_sql_db/scvmm -share-properties browsable,continuously-  
available -vserver Infra-SVM
```

Add a Domain Name Service Record for the SMB LIFs

1. Open DNS Manager and navigate to the forward lookup zone for the domain. Right-click the forward lookup zone and select New Host (A or AAAA).
2. Enter the SVM CIF's host name and IP address. Click Add Host.
3. Click OK to acknowledge the DNS record creation.
4. Repeat for the second SMB LIF.
5. Click Done to close the New Host window.

infrasmv	Host (A)	192.168.172.30	static
infrasmv	Host (A)	192.168.172.31	static

4.3 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 large FlexPod Express configurations.

Performing Initial Cisco UCS C-Series Standalone Server Setup for Cisco IMC

These steps provide details for the initial setup of the Cisco IMC interface for Cisco UCS C-Series standalone servers.

All Servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM 1.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 Cisco IMC configuration.



3. In the Cisco IMC 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:      [ ]                    None:          [ ]
Shared LOM:     [X]                    Active-standby: [X]
Cisco Card:     [ ]                    Active-active:  [ ]
  Riser1:       [ ]                    VLAN (Advanced)
  Riser2:       [ ]                    VLAN enabled:   [ ]
  MLom:         [ ]                    VLAN ID:        1
Shared LOM Ext: [ ]                    Priority:       0
IP (Basic)
IPV4:           [X]                    IPV6:          [ ]
DHCP enabled    [ ]
CIMC IP:        192.168.50.18
Prefix/Subnet:  255.255.255.0
Gateway:        192.168.50.1
Pref DNS Server: 10.61.186.19

*****
<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: <<esxi_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:      icee1-ucs2-cimc
  Dynamic DNS:   [ ]
  DDNS Domain:
FactoryDefaults
  Factory Default:      [ ]
Default User(Basic)
  Default password:
  Reenter password:
Port Properties
  Auto Negotiation:      [ ]
  Speed[1000/100 Mbps]:  100
  Duplex mode[half/full]: full
Port Profiles
  Reset:                 [ ]
  Name:
-no_pp
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>PreviousPage

```

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

Note: Upgrade the Cisco C-Series rack-mount server software to the latest version. This document uses version 2.0(3j).

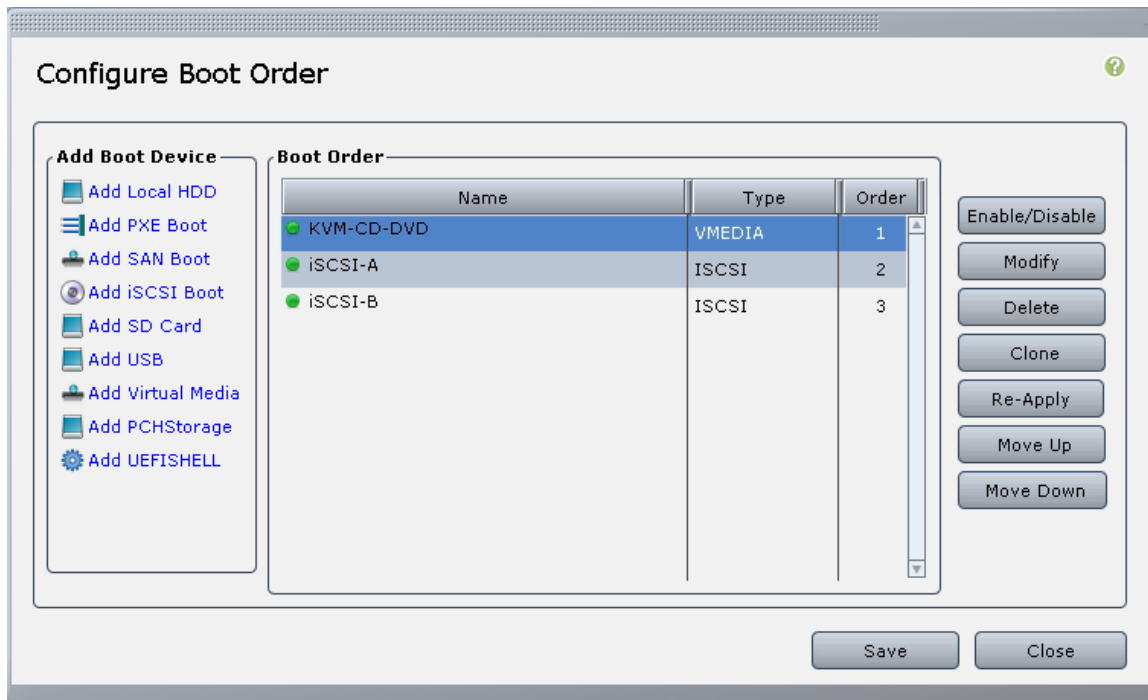
Configuring Cisco UCS C-Series Servers iSCSI Boot

In this document, VIC1227 is used for iSCSI boot.

Boot Order Configuration

1. From the Cisco IMC interface browser window (do not close the virtual KVM window), click the Server tab and choose BIOS.
2. Choose Configure Boot Order and click OK.
3. In the Boot Order section, remove all the entries and configure the following:
 - 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 and Close.
- 7. Click Save Changes.



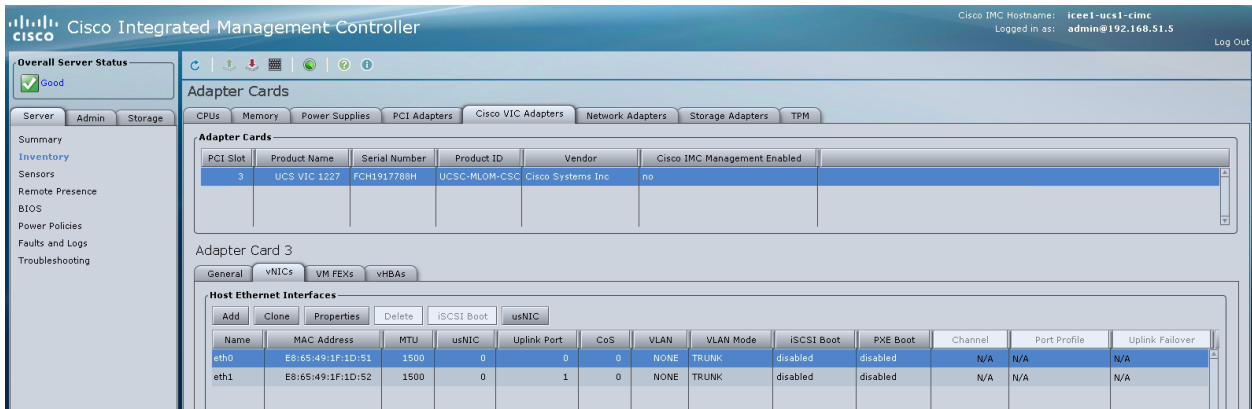
Configuring Cisco VIC1227 for iSCSI Boot: Part 1

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

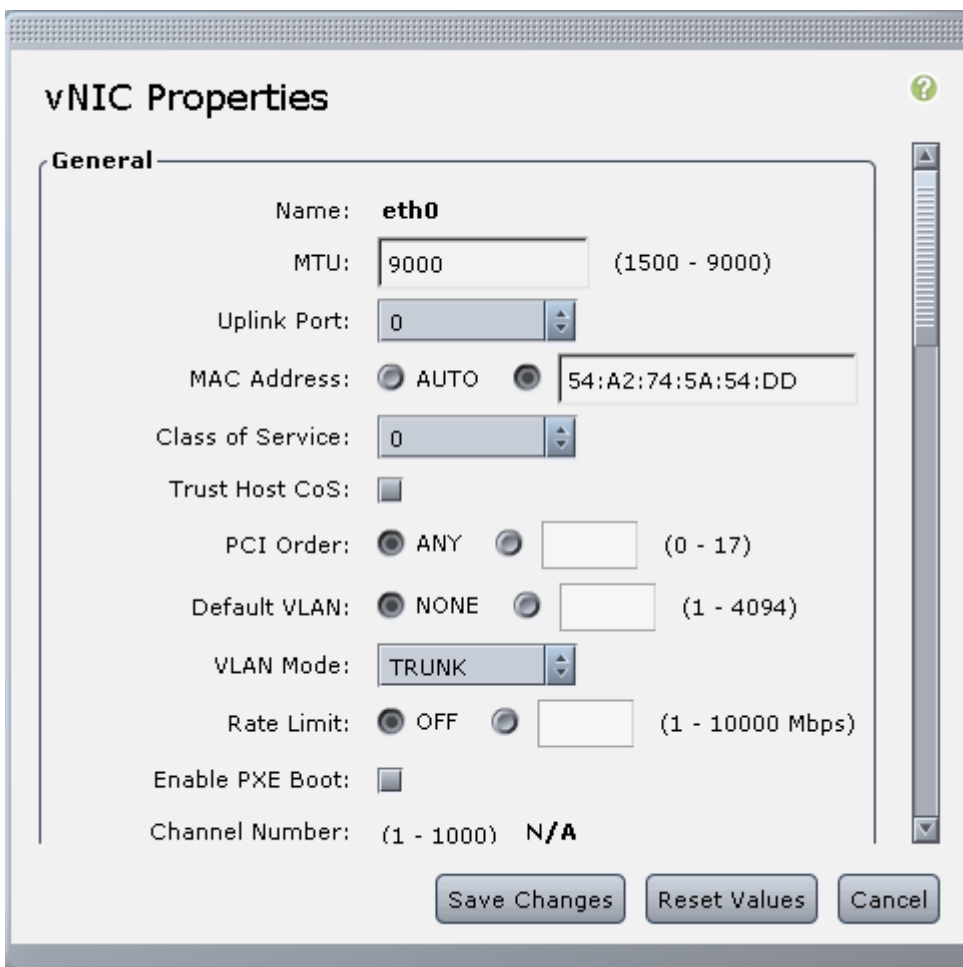
To start with, only one working path will be configured between the server and the iSCSI LUN because Windows cannot be installed on a LUN with multiple paths.

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

1. From the Cisco IMC interface browser window, click Inventory.
2. On 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 eth0 and click Properties.
6. Set the MTU to 9000.



1. Repeat steps 5 and 6 for eth1.
2. Click Add to create a new vNIC.
3. 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: check

vNIC Properties

General

Name: **iSCSI-VNIC-A**

MTU: (1500 - 9000)

Uplink Port:

MAC Address: ☐ AUTO ☒

Class of Service:

Trust Host CoS: ☐

PCI Order: ☒ ANY ☐ (0 - 17)

Default VLAN: ☒ NONE ☐ (1 - 4094)

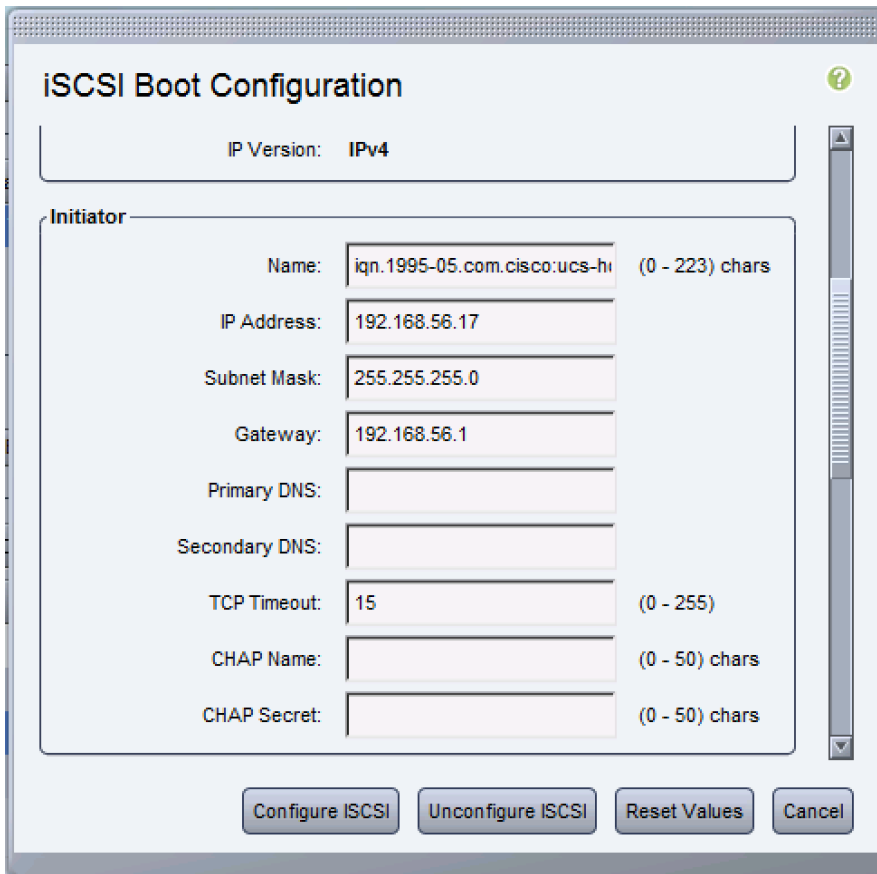
VLAN Mode:

Rate Limit: ☒ OFF ☐ (1 - 10000 Mbps)

Enable PXE Boot: ☒

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

4. Click Add vNIC. Click OK.
5. Select the newly created vNIC **iSCSI-VNIC-A** and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.
6. From the iSCSI Boot Configuration window, enter the Initiator details, as shown below.



The image shows a 'iSCSI Boot Configuration' dialog box. At the top, it says 'IP Version: IPv4'. Below this is a section titled 'Initiator' containing several input fields: 'Name' (with value 'iqn.1995-05.com.cisco:ucs-hi' and a '(0 - 223) chars' limit), 'IP Address' (with value '192.168.56.17'), 'Subnet Mask' (with value '255.255.255.0'), 'Gateway' (with value '192.168.56.1'), 'Primary DNS' (empty), 'Secondary DNS' (empty), 'TCP Timeout' (with value '15' and a '(0 - 255)' limit), 'CHAP Name' (empty and '(0 - 50) chars' limit), and 'CHAP Secret' (empty and '(0 - 50) chars' limit). At the bottom of the dialog are four buttons: 'Configure iSCSI', 'Unconfigure iSCSI', 'Reset Values', and 'Cancel'.

7. Enter the primary target details.
 - Name: IQN number of Infra-SVM.
 - IP Address: IP address of iscsi_lif01a
 - Boot LUN: 0
8. Click Configure iSCSI.

4.4 NetApp FAS Storage Deployment Procedure (Part 2)

Clustered Data ONTAP SAN Boot Storage Setup

Create iSCSI Igroups

To create igroups, complete the following step:

Note: Get the iSCSI-A IQN from the server configuration and use the proposed iSCSI-B IQN for the following steps.

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

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol iscsi -ostype windows -
initiator <<var_vm_host_infra_01_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_01_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol iscsi -ostype windows -
initiator <<var_vm_host_infra_02_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_02_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-03 -protocol iscsi -ostype windows -
initiator <<var_vm_host_infra_03_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_03_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-04 -protocol iscsi -ostype windows -
initiator <<var_vm_host_infra_04_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_04_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol iscsi -ostype windows -initiator
```

```
<<var_vm_host_infra_01_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_01_iSCSI-B_vNIC_IQN>>,  
<<var_vm_host_infra_02_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_02_iSCSI-B_vNIC_IQN>>,  
<<var_vm_host_infra_03_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_03_iSCSI-B_vNIC_IQN>>,  
<<var_vm_host_infra_04_iSCSI-A_vNIC_IQN>>, <<var_vm_host_infra_04_iSCSI-B_vNIC_IQN>>
```

Note: To view the five 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 boot_luns -lun VM-Host-Infra-01 -igroup VM-Host-Infra-01 -lun-  
id 0  
lun map -vserver Infra-SVM -volume boot_luns -lun VM-Host-Infra-02 -igroup VM-Host-Infra-02 -lun-  
id 0  
lun map -vserver Infra-SVM -volume boot_luns -lun VM-Host-Infra-03 -igroup VM-Host-Infra-03 -lun-  
id 0  
lun map -vserver Infra-SVM -volume boot_luns -lun VM-Host-Infra-04 -igroup VM-Host-Infra-04 -lun-  
id 0
```

4.5 Microsoft Windows Server 2012 R2 Deployment Procedure

This section provides detailed procedures for installing Windows Server 2012 R2 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 2012 R2 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.

Logging 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 installing the operating system through remote media.

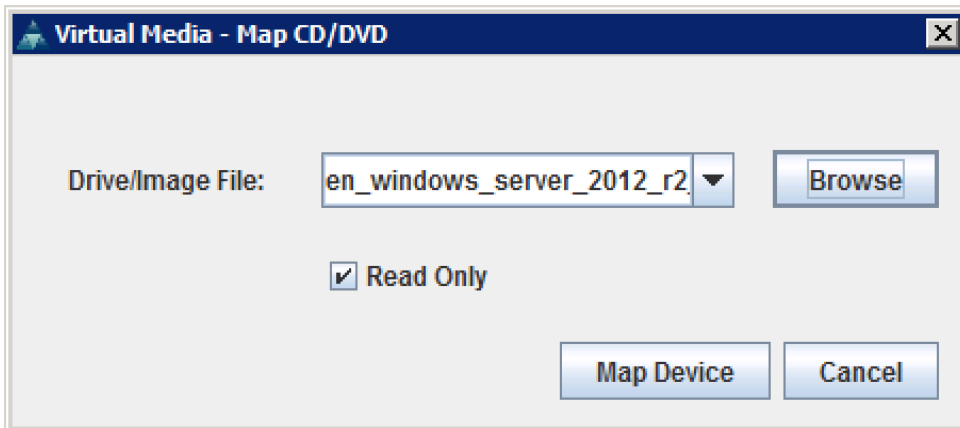
All Hosts

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. From 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.
7. Browse to the Windows Server 2012 R2 installer ISO image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (cold boot). Click Yes.

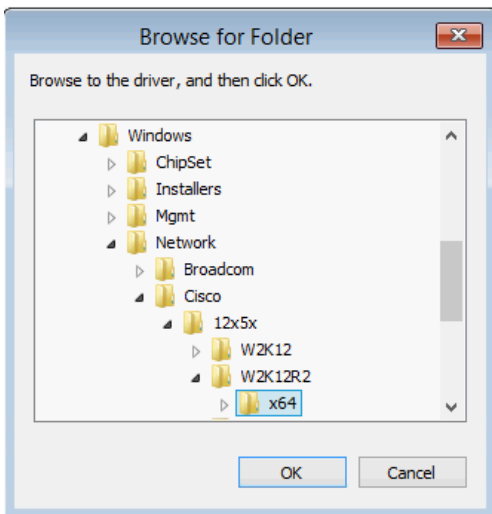
All Servers

1. Open a web browser and browse to the CIMC interface IP address.
2. Log in to the CIMC interface; the default user name is admin. Use the admin password <<admin_password>> set in the CIMC interface setup.
3. After successfully logging in, click the Server tab and then choose Summary. Select Launch KVM Console.
4. The virtual KVM window opens. Select Virtual Media at the top of the window.

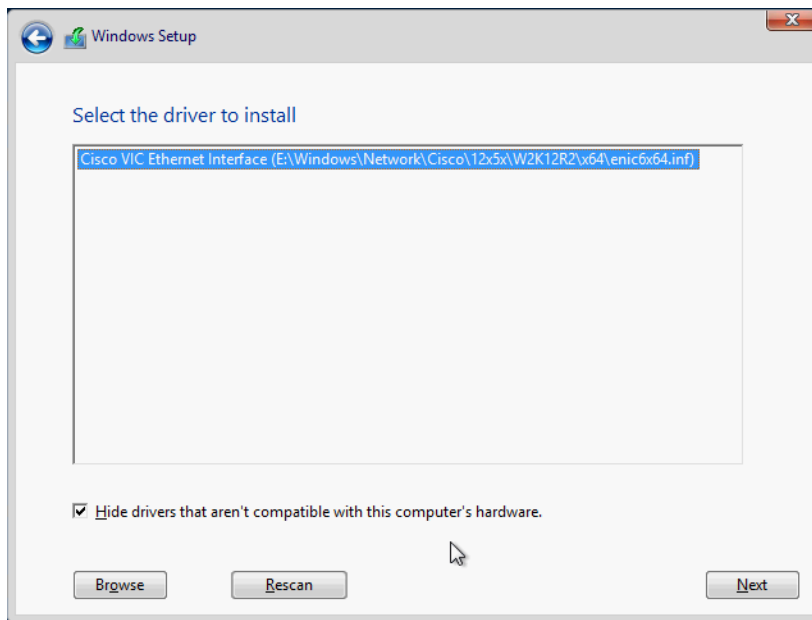
5. Click Activate Virtual Devices.
6. Click Map CD/DVD.
7. Browse to the location of the Windows Server 2012 R2 installation media and select it. Click Map Device.



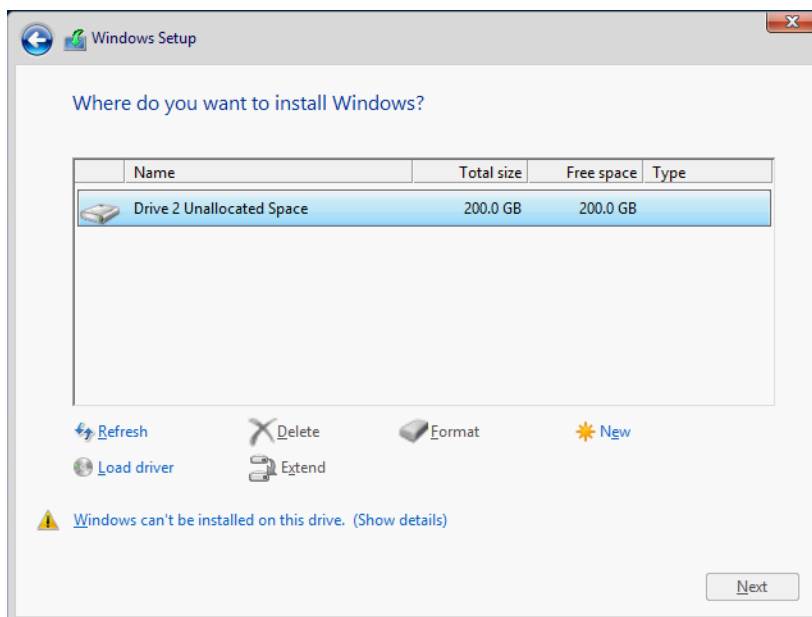
8. Click the Power tab and select Power Cycle System (cold boot). Click Yes.
The Windows Server 2012 R2 files are loaded.
9. On the Windows setup screen click Next.
10. Click Install Now.
11. Enter the product key and click Next.
12. Select Windows Server 2012 R2 Datacenter (server with a GUI) and click Next.
13. Accept the license terms and click Next.
14. Select Custom: Install Windows only (advanced).
15. Unmap the Windows 2012 R2 Installation media and map the Driver ISO (2.0.3c) downloaded from <https://software.cisco.com/download/release.html?mdfid=286281345&softwareid=283853158&release=2.0.3c&reind=AVAILABLE&rellifecycle=&reltype=latest>.
16. Click Load Driver to load the storage drivers to view the iSCSI Boot LUN.
17. Browse to the driver file by Navigating to Windows > Network > Cisco > 12x5x > W2K12R2 > x64 and then click OK.



18. Select Cisco VIC Ethernet Interface and click Next.



19. The iSCSI LUN should now be displayed.



20. Unmap the Driver ISO and remap the Windows Server 2012 R2 installation media to continue with the Windows installation.

21. Click Refresh.

22. Select the iSCSI LUN and click Next.

The server automatically reboots and performs an automated installation of Windows Server 2012 R2.

4.6 Update Display Adapter Driver

To update the driver for the display adapter, complete the following steps.

1. Download the Cisco UCS Server Configuration Utility Device Drivers Package for Windows 2012 R2 from <https://software.cisco.com/download/release.html?mdfid=286281345&flowid=71442&softwareid=283291009&os=Windows%202012r2%2064-bit&release=2.0%283b%29&relind=AVAILABLE&rellifecycle=&reltype=latest>.
2. From the Windows server navigate to Device Manager.
3. Expand the Display adapters. Right-click Microsoft Basic Display Adapter (Low Resolution) and select Update Driver Software.
4. Browse to the downloaded Cisco UCS Server Configuration Utility Device Drivers Package and select the folder.
5. Click Next and then click Close.

4.7 Install Microsoft Windows Features

To install the required Microsoft Windows Server 2012 R2 features, complete the following steps.

All Servers

1. From the Cisco IMC virtual KVM console, select the Virtual Media tab.
2. Click Map CD/DVD.
3. Browse to the Windows Server 2012 R2 installer ISO image file and click Map Device.
4. Log in to Windows with the administrator password entered during installation.
5. Launch the PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.
6. Enable remote management and remote desktop by using the SCONFIG application.

SCONFIG

7. Select Configure Remote Management by typing 4 and pressing Enter.
8. Select Enable Remote Management by typing 1 and pressing Enter.
9. Click OK.
10. Return to the main menu by typing 4 and pressing Enter.
11. Select Remote Desktop by typing 7 and pressing Enter.
12. Enter E to Enable.
13. Enter 2 to allow any version of remote desktop. After this step, you are returned to the main menu.
14. Click OK.
15. Select Exit to Command Line by typing 15 and pressing Enter.
16. Add the .NET 3.5, Hyper-V, MPIO, and clustering features by entering the following command:

```
Add-WindowsFeature Hyper-V, NET-Framework-Core, Failover-Clustering, Multipath-IO `
-IncludeManagementTools -Source E:\sources\sxs -Restart
```

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

17. Unmap the Windows Server 2012 R2 installation media from the Virtual Media tab.

Configuring Cisco VIC1227 Adapter for iSCSI Boot: Part 2

Now the remaining paths will be configured between the server and the iSCSI LUN because MPIO has been configured.

1. From the Cisco IMC interface browser window, click Inventory.
2. On 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.

The image shows a screenshot of the 'iSCSI Boot Configuration' window. It has a title bar with a question mark icon. The window is divided into two main sections: 'Primary Target' and 'Secondary Target'. Each section contains several input fields with labels and character counts. At the bottom, there are four buttons: 'Configure iSCSI', 'Unconfigure iSCSI', 'Reset Values', and 'Cancel'. A vertical scrollbar is on the right side of the window.

Field	Value	Character Count
Name	iqn.1992-08.com.netapp:sn.9!	(1 - 223) chars
IP Address	192.168.56.13	
TCP Port	3260	
Boot LUN	0	(0 - 65535)
CHAP Name		(0 - 50) chars
CHAP Secret		(0 - 50) chars

Field	Value	Character Count
Name	iqn.1992-08.com.netapp:sn.9!	(1 - 223) chars
IP Address	192.168.56.14	
TCP Port	3260	
Boot LUN	0	(0 - 65535)

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

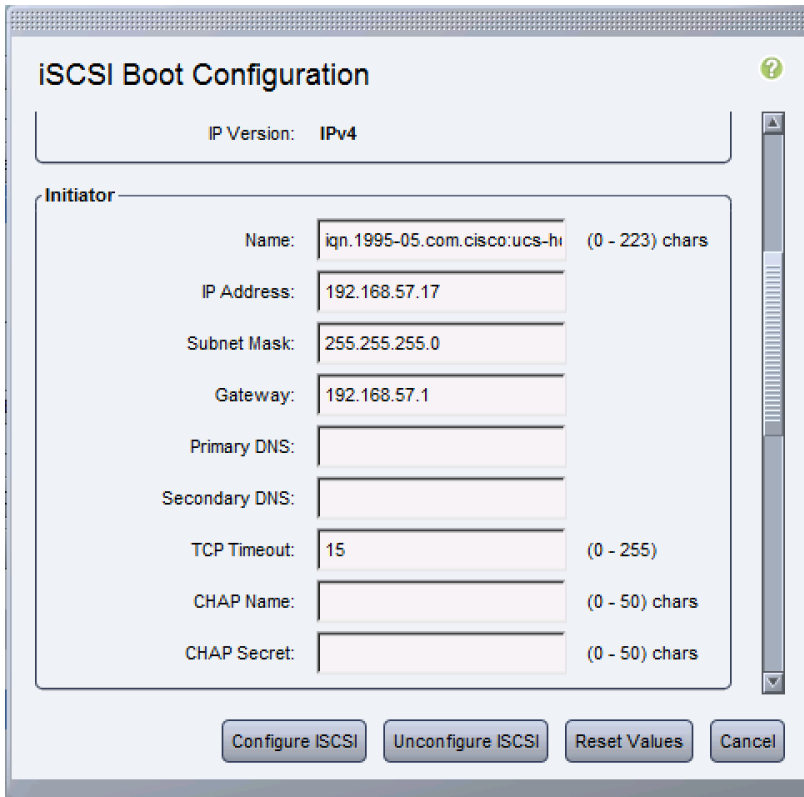
- MTU: 9000
- Uplink Port: 1
- Default VLAN: <<var_iscsi_vlan_b>>
- VLAN Mode: TRUNK
- Enable PXE Boot: check

The image shows a 'vNIC Properties' dialog box with a 'General' tab. The configuration is as follows:

- Name:** iSCSI-VNIC-B
- MTU:** 9000 (range: 1500 - 9000)
- Uplink Port:** 0
- MAC Address:** ☐ AUTO ☒ E8:65:49:1F:1D:5B
- Class of Service:** 0
- Trust Host CoS:** ☐
- PCI Order:** ☒ ANY ☐ (range: 0 - 17)
- Default VLAN:** ☒ NONE ☐ 3057 (range: 1 - 4094)
- VLAN Mode:** TRUNK
- Rate Limit:** ☒ OFF ☐ (range: 1 - 10000 Mbps)
- Enable PXE Boot:** ☒
- Channel Number:** (range: 1 - 1000) N/A

Buttons at the bottom: Save Changes, Reset Values, Cancel.

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, as shown below.



The image shows a 'iSCSI Boot Configuration' window. At the top, it says 'IP Version: IPv4'. Below this is a section titled 'Initiator'. Inside this section, there are several input fields: 'Name' (with a value 'iqn.1995-05.com.cisco:ucs-hi' and a note '(0 - 223) chars'), 'IP Address' (with '192.168.57.17'), 'Subnet Mask' (with '255.255.255.0'), 'Gateway' (with '192.168.57.1'), 'Primary DNS' (empty), 'Secondary DNS' (empty), 'TCP Timeout' (with '15' and a note '(0 - 255)'), 'CHAP Name' (empty and note '(0 - 50) chars'), and 'CHAP Secret' (empty and note '(0 - 50) chars'). At the bottom of the window are four buttons: 'Configure iSCSI', 'Unconfigure iSCSI', 'Reset Values', and 'Cancel'.

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.

iSCSI Boot Configuration

Primary Target

Name: (1 - 223) chars

IP Address:

TCP Port:

Boot LUN: (0 - 65535)

CHAP Name: (0 - 50) chars

CHAP Secret: (0 - 50) chars

Secondary Target

Name: (1 - 223) chars

IP Address:

TCP Port:

Boot LUN: (0 - 65535)

15. Click Configure iSCSI.
16. Reboot the Windows host.

4.8 Configure Microsoft Windows

To configure the network for each Hyper-V host, complete the following steps.

All Servers

1. Log in with the administrator password entered during installation.
2. Launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.
3. Find the 10GbE interfaces by running the `Get-NetAdapter` command.

Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
Ethernet 5	Cisco VIC Ethernet Interface #4	45	Up	E8-65-49-1F-1D-5B	10 Gbps
Ethernet 2	Cisco VIC Ethernet Interface #2	12	Up	E8-65-49-1F-1D-5A	10 Gbps
Ethernet	Cisco VIC Ethernet Interface	13	Up	E8-65-49-1F-1D-52	10 Gbps
Ethernet 3	Cisco VIC Ethernet Interface #3	14	Up	E8-65-49-1F-1D-51	10 Gbps
LOM Port 2	Intel(R) I350 Gigabit Network Connec...	15	Up	A0-EC-F9-CE-32-C5	1 Gbps
LOM Port 1	Intel(R) I350 Gigabit Network Conn...#2	16	Up	A0-EC-F9-CE-32-C4	1 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
```

- Log in to the Cisco IMC web interface, select Inventory from the left pane, and click Cisco VIC Adapters in the right pane.
- Select the Cisco UCS VIC 1227 and click vNICs
- Note the MAC Address of interfaces eth0 and eth1; these interfaces are not used for iSCSI boot.

Adapter Card 3

General vNICs VM FEXs vHBAs

Host Ethernet Interfaces

Add Clone Properties Delete iSCSI Boot usNIC

Name	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile
eth0	E8:65:49:1F:1D:51	9000	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A
eth1	E8:65:49:1F:1D:52	9000	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A

- Create a NIC team using the 10GbE interfaces that are not used for iSCSI boot.

```
New-NetLbfoTeam -Name TM1 -TeamMembers <10GBE_nic1>, <10GBE_nic2> -TeamingMode SwitchIndependent -LoadBalancing HyperVPort
```

Note: Using the example output from step 3 and step 7, the command would be:

```
New-NetLbfoTeam -Name TM1 -TeamMembers 'Ethernet','Ethernet 3' -TeamingMode SwitchIndependent -LoadBalancing HyperVPort
```

- Press Enter to confirm.
- In this document, LOM port 1 and LOM port are used for Cisco IMC. Disable the LOM adapters.

```
Get-NetAdapter *LOM* | Disable-NetAdapter
```

- Remove the IP stack from the TM NIC interface.

```
Get-NetAdapter TM1 | set-NetAdapterBinding -ComponentID ms_tcpip* -Enabled $false
```

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

```
New-VMSwitch -Name VMComm -NetAdapterName TM1 -AllowManagementOS $false
```

- Create VM NICs.

```
Add-VMNetworkAdapter -ManagementOS -Name Mgmt -SwitchName VMComm
Add-VMNetworkAdapter -ManagementOS -Name Cluster -SwitchName VMComm
Add-VMNetworkAdapter -ManagementOS -Name LM -SwitchName VMComm
Add-VMNetworkAdapter -ManagementOS -Name SMB -SwitchName VMComm

Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Mgmt -Access -AccessVlanId <<ib_mgmt_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName SMB -Access -AccessVlanId <<smb_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName LM -Access -AccessVlanId <<livemigraion_vlan_id>>
Set-VMNetworkAdapterVlan -ManagementOS -VMNetworkAdapterName Cluster -Access -AccessVlanId <<cluster_vlan_id>>
```

- Configure jumbo frames on the select interfaces.

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

- 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> -Prefix <Cluster_prefix>
```

```
New-NetIPAddress -InterfaceAlias 'vEthernet (LM)' -IPAddress <LM_ipaddress> -Prefix <LM_prefix>
New-NetIPAddress -InterfaceAlias 'vEthernet (SMB)' -IPAddress <SMB_ipaddress> -Prefix <SMB_prefix>
```

16. Disable DNS registration for all NICs.

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

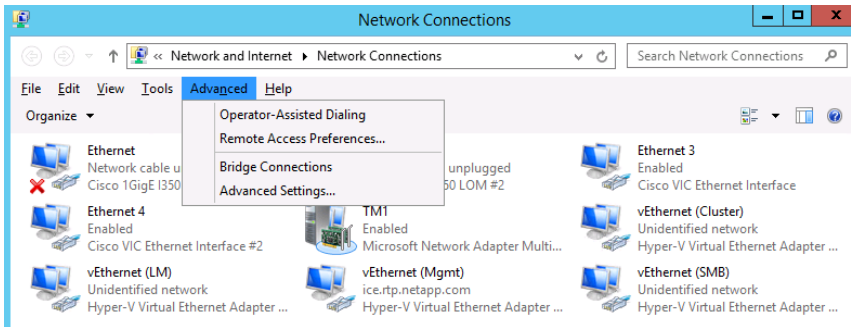
17. Turn registration back on and configure DNS for the Mgmt NIC.

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

18. From the CLI, enter `control netconnections` to open the Network Connections control panel.

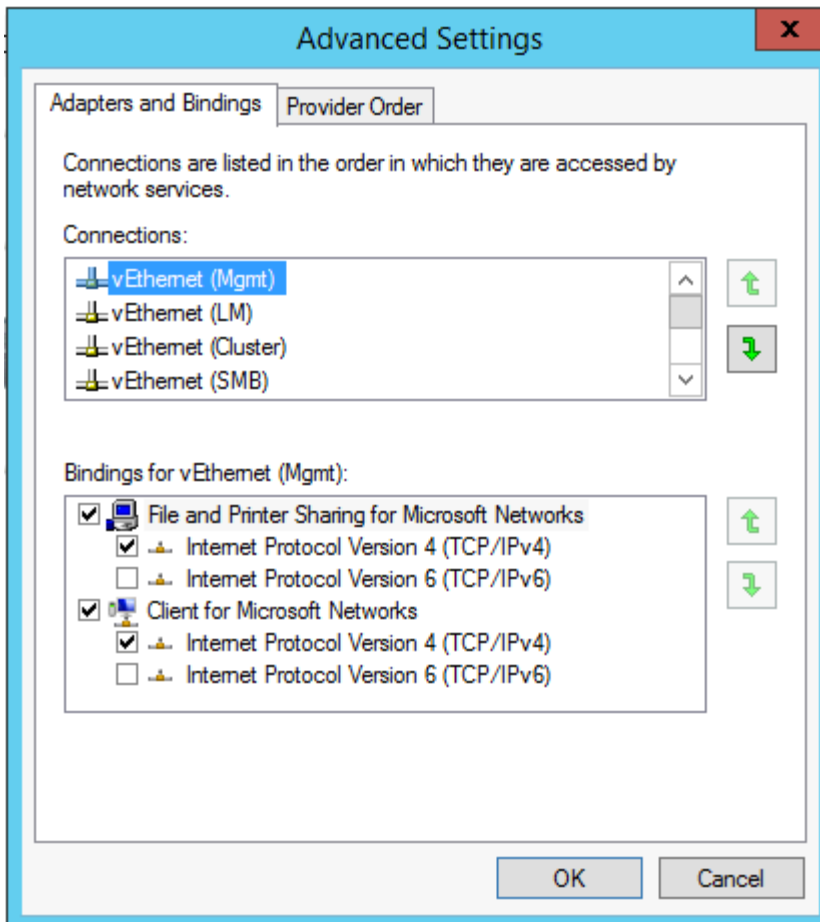
19. Press the Alt key to access the Advanced menu.

20. Click Advanced and select Advanced Settings.



21. Use the green arrows to modify the connection binding order as follows:

- a. vEthernet (Mgmt)
- b. vEthernet (LM)
- c. vEthernet (Cluster)
- d. vEthernet (SMB)



22. Rename the server and join the domain.

```
Rename-Computer <ServerName> -restart
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. After the server reboots, launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.

Installing NetApp Windows iSCSI Host Utilities

The following section describes how to perform an unattended installation of the NetApp Windows iSCSI Host Utilities. For detailed information about the installation, see the [Windows Host Utilities 6.0.2 Installation and Setup Guide](#).

All Servers

1. Download Windows iSCSI Host Utilities from http://mysupport.netapp.com/NOW/download/software/sanhost_win/6.0.2/netapp_windows_host_utilities_6.0.2_x64.msi.

2. Unblock the downloaded file.

```
Unblock-file ~\Downloads\netapp_windows_host_utilities_6.0.2_x64.msi
```

3. Install the Host Utilities.

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

Note: The system reboots during this process.

Configuring Windows Host iSCSI Initiator

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

All Servers

1. Launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting 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 -TargetPortalAddress <<controller01_iscsi_lif01a_ip>> -  
InitiatorPortalAddress <iscsia_ipaddress>  
New-IscsiTargetPortal -TargetPortalAddress <<controller01_iscsi_lif01b_ip>> -  
InitiatorPortalAddress <iscsib_ipaddress>  
New-IscsiTargetPortal -TargetPortalAddress <<controller02_iscsi_lif02a_ip>> -  
InitiatorPortalAddress <iscsia_ipaddress>  
New-IscsiTargetPortal -TargetPortalAddress <<controller02_iscsi_lif02b_ip>> -  
InitiatorPortalAddress <iscsib_ipaddress>
```

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

```
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPo  
rtalAddress <iscsia_ipaddress>  
Get-IscsiTarget | Connect-IscsiTarget -IsPersistent $true -IsMultipathEnabled $true -InitiatorPo  
rtalAddress <iscsib_ipaddress>
```

Install NetApp SnapDrive

The following section describes how to install NetApp SnapDrive for Windows. For detailed installation procedures, refer to the [Administration and Installation Guide](#).

1. In Active Directory, create a SnapDrive service account.

Note: This account requires no special delegation; the same account can be used for multiple hosts.

New Object - User

Create in: flexpod.test/Users

First name: SnapDrive Initials:

Last name:

Full name: SnapDrive

User logon name: SnapDrive @flexpod.test

User logon name (pre-Windows 2000): FLEXP0D\ SnapDrive

< Back Next > Cancel

2. Add the SnapDrive service account to the local administrators group in Windows.

Administrators Properties

General

Administrators

Description: Administrators have complete and unrestricted access to the computer/domain

Members:

- Administrator
- FLEXP0D\Domain Admins
- FLEXP0D\SnapDrive (SnapDrive@flexpod.test)

Add... Remove

Changes to a user's group membership are not effective until the next time the user logs on.

OK Cancel Apply Help

3. Download the SnapDrive installer v7.1.1 from the [NetApp Support site](#).
4. Launch the installer and click Next.

Note: If the installer prompts for a hotfix or patches to be downloaded, complete the activity and then proceed with SnapDrive installation.
5. Select the storage-based licensing method and click Next.

SnapDrive® - Installation Wizard

SnapDrive License
Please provide valid SnapDrive license to install

Select the license type:

☒ Per Storage System

☐ Per Server

Provide Valid SnapDrive License Key

License Key:

InstallShield

< Back Next > Cancel

6. Enter your user name and organization information and click Next.
7. Validate the installation path and click Next.
8. Select the Enable SnapDrive to Communicate Through the Windows Firewall checkbox and click Next.
9. Enter the information for the SnapDrive service account and click Next.
10. On the SnapDrive Web Service Configuration page, click Next.
11. Clear the Enable Preferred Storage System IP Address checkbox and click Next.
12. Clear the Enable Transport Protocol Settings checkbox and click Next.

SnapDrive - InstallShield Wizard

Transport Protocol Default Setting
Specify Default Transport Setting for Storage System(s)

☐ Enable Transport Protocol Settings

☒ RPC

☐ HTTP

☐ HTTPS

Specify the user name and password for the HTTP/HTTPS Protocol selection.

User Name:

Password:

Port ID:

InstallShield

< Back Next > Cancel

13. Leave the Enable Dataset Protection Integration option cleared and click Next.
14. Click Install.
15. Click Finish.
16. After the installation is finished, launch a new PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.

Note: A new prompt is required to register the `sdcli` executable.

17. Configure the SnapDrive preferred IP settings for each storage controller.

```
sdcli preferredIP set -f <<var_vserver_name>> -IP << var_vserver_mgmt_ip>>
```

18. Configure the SnapDrive transport protocol authentication configuration for each storage controller.

```
Set-SdStorageConnectionSetting -StorageSystem <<var_vserver_mgmt_ip>> -protocol https -credential vsadmin
```

19. In the pop-up windows, provide the password for the vsadmin account.

Install NetApp SnapManager for Hyper-V

This section describes how to install NetApp SnapManager for Hyper-V (SMHV). For detailed installation procedures, refer to the [Administration and Installation Guide](#).

1. In Active Directory, create an SMHV service account.

Note: This account requires no special delegation; the same account can be used for multiple hosts.

2. Add the SMHV service account to the local administrators group in Windows.
3. Download the SMHV installer v2.1 from the [NetApp Support](#) site.
4. Launch the installer and click Next.
5. Select the Storage-Based Licensing method and click Next.
6. Validate the installation path and click Next.
7. Enter the information for the SMHV service account and click Next.
8. On the SMHV Web Service Configuration page, click Next.
9. Click Install.
10. Click Finish.

4.9 Create Windows Failover Cluster

To create the Windows failover cluster, log in to any one of the Hyper-V hosts and complete the following steps:

1. Make sure Windows 2012 R2 is updated.
2. Create the Windows failover cluster.

```
New-Cluster -Name <cluster_name> -Node <hostnames> -NoStorage -StaticAddress <cluster_ip_address>
```

Note: You can add additional servers to the cluster at a later stage.

3. Name the cluster networks according to their usage. Use the following commands to name the networks:

```
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 *SMB* | Group Network| %{
    (Get-ClusterNetwork $_.Name).Name = 'SMB'}
```

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

4. Set cluster network role types to allow only client connections on the management network and to prevent cluster communication on the iSCSI network.

```
(Get-ClusterNetwork Cluster).role = 1  
(Get-ClusterNetwork LM).role = 1  
(Get-ClusterNetwork SMB).role = 0  
(Get-ClusterNetwork Mgmt).role = 3
```

Note: A setting of 0 prevents all cluster traffic, 1 allows cluster communication, and 3 allows both client and cluster communication.

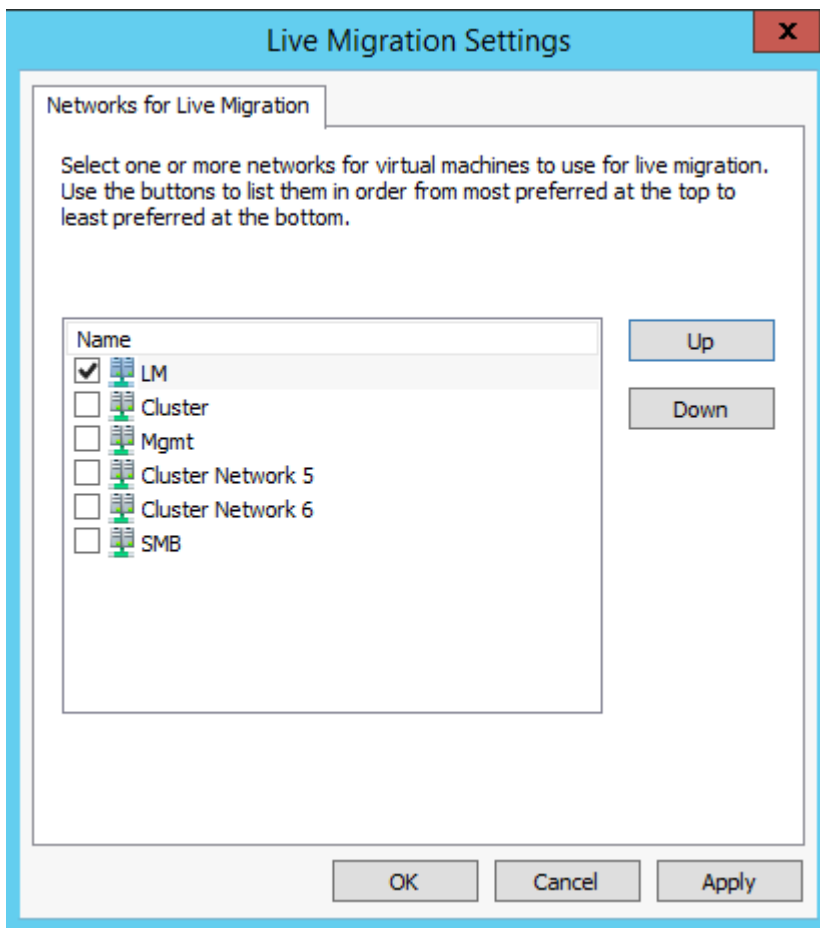
Configuring Live Migration

The preferred Live Migration network can be set either through the GUI or the CLI. The CLI requires modification to the registry; as such, NetApp recommends using the GUI.

Configure Through GUI

To configure the live migration network through the GUI, complete the following steps:

1. From Server Manager, select Tools and then select Failover Cluster Manager.
2. Expand the cluster tree on the left, right-click Networks, and select Live Migration Settings.
3. Make sure that the only option selected is the live migration network. Click OK.



Configure Through CLI

To configure the live migration network through the CLI, complete the following steps:

1. Configure the live migration network.

```
$ClusterNetwork = Get-ClusterNetwork
$includeIDs = $ClusterNetwork | Where-Object { $_.Name -eq 'LM'} |
    Select-Object -ExpandProperty ID
$excludeIDs = $ClusterNetwork | Where-Object { $_.Name -ne 'LM'} |
    Select-Object -ExpandProperty ID
$excludeIDs = $excludeIDs -join ';'

Set-ItemProperty -Path "HKLM:\Cluster\ResourceTypes\Virtual Machine\Parameters" `
    -Name MigrationExcludeNetworks -Value $excludeIDs
Set-ItemProperty -Path "HKLM:\Cluster\ResourceTypes\Virtual Machine\Parameters" `
    -Name MigrationNetworkOrder -Value $includeIDs
```

Configuring Constrained Delegation for Hyper-V Hosts

Although the hosts have the required permissions to access the SMB share, you might encounter access-denied errors when trying to remotely manage the hosts. To avoid these error messages, configure constrained delegation for the Hyper-V hosts by completing the following steps:

1. In Active Directory Users and Computers, browse to the Computer objects for each Hyper-V host.
2. Right-click the object and select Properties.
3. Select the Delegation tab.
 - a. Select Trust This Computer for Delegation to Specified Services Only.
 - b. Select Use Kerberos Only.
 - c. Click Add.
4. Click the Users or Computers button on the top of the Add Services dialog box.
5. Enter the name of the infrastructure SVM and click OK.
6. Select CIFS and click OK.
7. Click OK.
8. Repeat steps 1 through 7 for each Hyper-V host.

Changing Management Cluster to Use a File Share Witness

To change the management cluster to use the quorum disk, complete the following steps on one server only:

1. Open an SSH connection to the cluster IP or host name and log in using the admin user and password.
2. Remove the Everyone permission from the witness share.

```
cifs share access-control delete -share hyperv-witness -user-or-group Everyone -vserver infra_svm
```

3. Add permissions to the following accounts with NTFS full control permissions over the share:
 - Hyper-V Node 1
 - Hyper-V Node 2
 - Hyper-V Node 3
 - Hyper-V Node 4
 - Hyper-V Cluster Name Object (CNO)

```
share access-control create -share hyperv-witness -user-or-group <<var_domain>>\HyperV1$ -
permission full_Control -vserver infra_svm
```

```
share access-control create -share hyperv-witness -user-or-group <<var_domain>>\HyperV2$ -
permission full_Control -vserver infra_svm
share access-control create -share hyperv-witness -user-or-group <<var_domain>>\HyperV3$ -
permission full_Control -vserver infra_svm
share access-control create -share hyperv-witness -user-or-group <<var_domain>>\HyperV4$ -
permission full_Control -vserver infra_svm
share access-control create -share hyperv-witness -user-or-group <<var_domain>>\HyperV$ -
permission full_Control -vserver infra_svm
```

4. Launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.
5. Set the cluster to use the SMB share created previously.

```
Set-ClusterQuorum -FileShareWitness \\infrasm\hyperv-witness
```

Validating the Management Cluster

To validate the management cluster, complete the following steps on one server only:

1. Select the cluster in the navigation pane and click Validate Cluster.
2. Click Next.
3. Select Run all Tests (recommended) and click Next.
4. Click Next.
5. Select the shared disks on the cluster and click Next.
6. Confirm the selected options and click Next.
7. Review and correct any failures that are listed in the validation report.

Note: The Validation wizard reports the following warnings, which can be safely disregarded.

Successfully issued call to Persistent Reservation REGISTER using Invalid RESERVATION KEY 0xc, SERVICE ACTION RESERVATION KEY 0xd, for Test Disk 0 from node VMHost-Mgmt1.flexpod.local.

Test Disk 0 does not support SCSI-3 Persistent Reservations commands needed by clustered storage pools that use the Storage Spaces subsystem. Some storage devices require specific firmware versions or settings to function properly with failover clusters. Contact your storage administrator or storage vendor for help with configuring the storage to function properly with failover clusters that use Storage Spaces.

5 Microsoft System Center 2012 R2 Virtual Machine Manager

The procedures in the following subsections provide detailed instructions for installing System Center 2012 R2 Virtual Machine Manager in a FlexPod environment.

Table 10) VM requirements.

Role	Virtual CPU	RAM (GB)	Virtual Hard Disk (GB)
Virtual Machine Manager	4	8	60
SMI-S Agent	1	4	60

5.1 Build the SMI-S and SCVMM VMs

One Server Only

1. In Failover Cluster Manager, right-click Roles and select Virtual Machine. Select New Virtual Machine.
2. Select the host for the new virtual machine and click OK.
3. On the New Virtual Machine welcome screen, click Next.
4. Enter a name for the VM (for example, SCVMM), select the Store the virtual machine in a different location checkbox, and enter the UNC path for the infrastructure SMB share. Click Next.

Choose a name and location for this virtual machine.


The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.

Name:

You can create a folder or use an existing folder to store the virtual machine. If you don't select a folder, the virtual machine is stored in the default folder configured for this server.

☒ Store the virtual machine in a different location

Location:

 If you plan to take checkpoints of this virtual machine, select a location that has enough free space. Checkpoints include virtual machine data and may require a large amount of space.

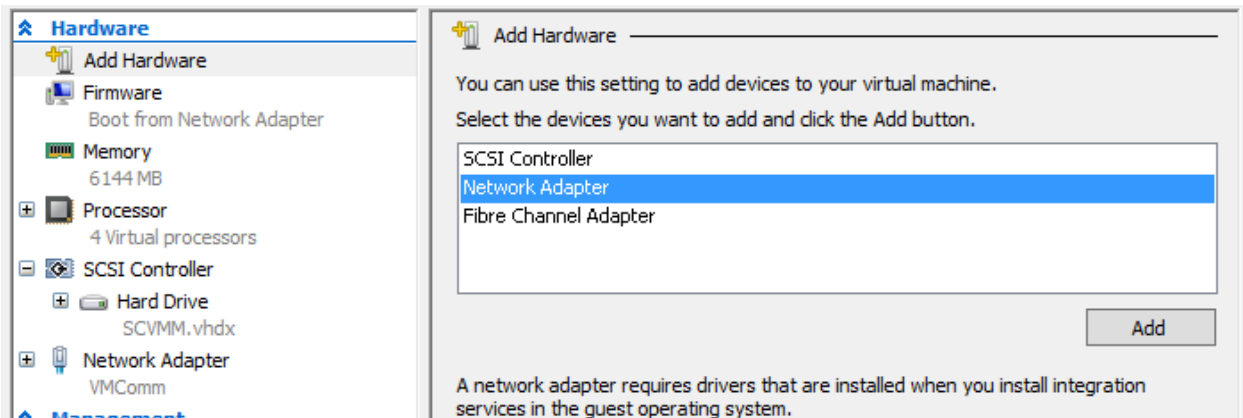
5. Select Generation 2 and click Next.
6. Enter the startup memory for the VM from Table 10, and select the Use Dynamic Memory for this virtual machine checkbox. Click Next.
7. Select the VMComm network and click Next.
8. Set the size for the new VHDX from Table 10, and click Next.
9. Select Install an operating system from a bootable image file, and provide the path to the Windows Server 2012 R2 ISO. Click Next.
10. Click Finish after reviewing the VM settings to be configured.
11. Click Finish in the High Availability Wizard Summary.
12. Repeat steps 1 through 11 for the remaining VMs.

5.2 Configure SMI-S and SCVMM VMs

1. In Failover Cluster Manager, select Roles and right-click the VM to be modified. Select Settings.
2. Select Memory and set the Dynamic Memory Maximum RAM to the Startup RAM.
3. Select CPU and set the CPU to the value outlined in Table 10.
4. Select Network Adapter, select Enable Virtual LAN Identification, and enter the <<ib_mgmt_vlan_id>>.
5. Select Automatic Start Action and select the Always Start This Virtual Machine Automatically option.
6. Select Automatic Stop Action and select the Shut Down the Guest Operating System option.
7. Click OK to save the modifications.
8. Repeat steps 1 through 7 for the remaining VMs.

5.3 Add SMB Network Adapter to the SCVMM VM

1. In the Microsoft Failover Cluster Manager, select Roles. Right-click the SCVMM VM and select Settings.
2. From the Add Hardware section, select Network Adapter and click Add.



3. Select the VMComm virtual switch, enable the Enable Virtual LAN Identification option, and enter the <<smb_vlan_id>>.
4. Click OK to save the modifications.

5.4 Install Windows Server 2012 R2 on the VMs

1. In Failover Cluster Manager, select Roles. Right-click the required VM and select Connect.
2. Click Start to power on the VM and boot into the Windows installer.
3. After the installer loads, enter the relevant region information and click Next.
4. Click Install Now.
5. Enter the product key and click Next.
6. Select Windows Server 2012 R2 Data Center (server with a GUI) and click Next.
7. Review and accept the license agreement and click Next.
8. Select Custom: Install Windows only (advanced).
9. Select Drive 0 and click Next.
10. After the installation is complete, enter a password for the administrator on the Settings page and click Finish.
11. Log in to the server console and launch a PowerShell prompt. Install .NET 3.5 by running the following command:

```
Add-WindowsFeature -Name NET-Framework-Core -Source D:\sources\sxs
```

Note: NetApp assumes that the Windows ISO is mapped to drive D. If the ISO is mounted on a different drive, update the source path accordingly.

12. Unmap the Windows ISO.
13. Configure the network adapter settings if using static IPs.

Note: The SCVMM virtual machine has two network adapters. Look for the MAC address of the adapters in the VM Settings menu and appropriately assign the IP addresses.

14. Install important and recommended Windows updates and reboot.
15. Rename the VM and add it to Active Directory.
16. Repeat steps 1 through 15 for the remaining VMs.

5.5 Install the NetApp SMI-S Agent in the SMI-S Virtual Machine

To install the NetApp SMI-S Agent, complete the following steps.

Prerequisites

You must meet the following environment prerequisites before proceeding.

Accounts

Verify that the following local account has been created:

User Name	Purpose	Permissions
SMIS-User	SMI-S access account	This account does not need any special delegation. Note: This is NOT a domain account. It must be a local account in Windows.

1. Verify that the following account is a member of the local administrator's group.
 - SMIS-User

Install the SMI-S Agent

To install the NetApp SMI-S agent, complete the following step.

1. Download the Data ONTAP SMI-S Agent installer from <http://mysupport.netapp.com/NOW/download/software/smis/Windows/5.2/smisagent-5-2.msi>.

Install the SMI-S Provider

Complete the following steps to install the NetApp SMI-S provider.

1. Right-click smisagent-5-2 and select Install.
2. On the Welcome to the Data ONTAP SMI-S Agent Setup Wizard page, click Next.
3. On the Ready to Install Data ONTAP SMI-S Agent page, click Install.
4. On the Completed the Data ONTAP SMI-S Agent Setup Wizard page, click Finish to complete the installation.

Configure the SMI-S Provider

To configure the NetApp SMI-S provider, complete the following steps.

1. Open the App screen, right-click Data ONTAP SMI-S Agent, and select Run as Administrator at the bottom of the screen.
2. Change the directory into the SMI-S program files.

```
cd %ProgramFiles(x86)%\ONTAP\smis\pegasus\bin
```

3. Add the SVM credentials to the SMI-S Agent.

```
Smis addsecure <VserverIpAddress> <VserverAdmin> <VserverAdminPassword>
```

4. Enable user authentication.

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

5. Restart the agent/cimserver.

```
Smis cimserver restart
```

6. Add the SMI-S Run As account to the SMIS configuration.

```
cimuser -a -u SMIS-User -w <password>
```

5.6 Install System Center Virtual Machine Manager

To install SCVMM in a minimal configuration, complete the following.

Prerequisites

You must meet the following environment prerequisites before proceeding.

Accounts

Verify that the following accounts have been created:

User Name	Purpose	Permissions
<DOMAIN>\SCVMM-SVC	Virtual Machine Manager Service Account	Requires full admin permissions on the SCVMM virtual machine and runs the Virtual Machine Manager service. Should also be added to the SQL-Admins group and as a sysadmin in all instances.
<DOMAIN>\SQL-SVC	SQL Server Service Account	Requires full admin permissions on the SCVMM virtual machine and runs the service account for all instances. This account must also be added to the SQL-Admins group and as a sysadmin in all instances.
<DOMAIN>\SnapDrive	SnapDrive for Windows	Needs to be an administrator on the SCVMM virtual machine.

Groups

Verify that the following security groups have been created:

Security Group Name	Group Scope	Members
<DOMAIN>\SCVMM-Admins	Global	SCVMM-SVC
<DOMAIN>\SQL-Admins	Global	All SQL Server administrators for the fabric management solution. The user account being used to install SQL must belong in this group.

7. Verify that the following accounts and/or groups are members of the Local Administrators group on the Virtual Machine Manager virtual machine:
- SnapDrive
 - Virtual Machine Manager Admins group
 - Virtual Machine Manager Service account
 - SQL Service account
 - SQL Admins group

Install the Windows Assessment and Deployment Kit

The SCVMM installation requires that the Windows Assessment and Deployment Kit (ADK) be installed on the Virtual Machine Manager management server. You can download the Windows ADK from <http://www.microsoft.com/en-us/download/details.aspx?id=39982>.

During installation, only the Deployment Tools and the Windows Preinstallation Environment features are selected. This installation also assumes that the VMM servers have Internet access. If that is not the case, an offline installation can be performed.

The following steps outline how to install the Windows ADK on the SCVMM management server.

1. From the Windows ADK installation media source, right-click `adksetup.exe` and select **Run as administrator**. If prompted by user account control, click **Yes** to allow the installation to make changes to the computer.
2. In the **Specify Location** page, accept the default folder location of `%ProgramFiles%\Windows Kits\8.1` and click **Next**.
3. In the **Join the Customer Experience Improvement Program (CEIP)** page, select the option to either participate or not participate in the CEIP by providing selected system information to Microsoft. Click **Next**.
4. In the **License Agreement** page, click **Accept**.
5. In the **Select the features you want to install** page, select the following checkboxes:
 - **Deployment Tools**
 - **Windows Preinstallation Environment (Windows PE)**
6. Make sure that all other option checkboxes are cleared. Click **Install** to begin the installation.
7. After installation is complete, clear the **Launch the Getting Started Guide** checkbox and click **Close** to exit the installation wizard.

Install the WSUS RSAT Tools

The Virtual Machine Manager installation requires the WSUS Administration Tools to be installed on the Virtual Machine Manager management server.

1. Launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting **Run as Administrator**.
2. Add Failover Cluster, Multipath-IO, and the WSUS console by entering the following command:

```
Add-WindowsFeature -Name UpdateServices-RSAT -IncludeManagementTools -Restart
```

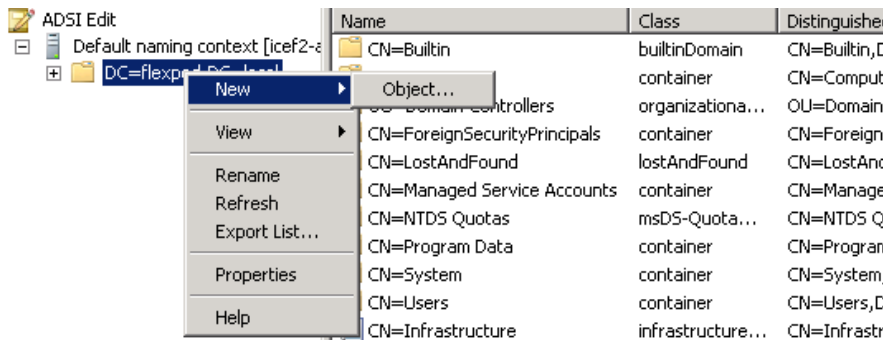
Create the Virtual Machine Manager Distributed Key Management Container in Active Directory Domain Services

The Virtual Machine Manager installation requires that an Active Directory container be created to house the distributed key information for Virtual Machine Manager.

Note: If Virtual Machine Manager will be deployed using an account with rights to create containers in AD DS, you can skip this step.

Perform the following steps to create an AD DS container to house the distributed key information. These instructions assume that a Windows Server 2012 R2 domain controller is in use; similar steps would be followed for other versions of Active Directory including Windows Server 2008 and Windows Server 2012.

1. Log in to a Domain Controller with a user that has Domain Admin privileges and run `adsiedit.msc`.
2. Right-click the **ADSI Edit** node and select **Connect to**.
3. In the **Connections Settings** page, from the **Connection Point** section, select the **Select a well-known Naming Context** option. Select **Default naming context** from the drop-down menu and click **OK**.
4. Expand **Domain Default naming context** [`<computer fully qualified domain name>`] and expand `<distinguished name of domain>`. Right-click the root node and select **New – Object**.

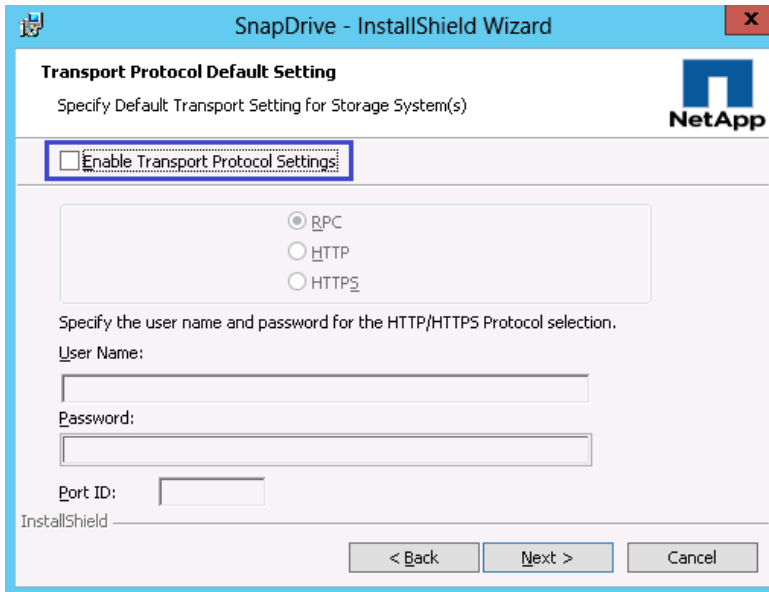


5. In the Create Object page, select Container and then click Next.
6. In the Value text box, type VMMDKM and then click Next.
7. Click Finish to create the container object.
8. Within ADSI Edit, right-click the new VMMDKM object and then click Properties.
9. In the VMMDKM Properties dialog box, click the Security tab. Click Add to add the VMM Service account and the VMM Admins group. Grant the security principles Full Control permissions.
10. Click OK three times and close ADSI Edit.

Install NetApp SnapDrive

The following section describes the installation of SnapDrive for Windows. For detailed information regarding the installation, see the [Administration and Installation Guide](#).

1. Download SnapDrive installer
http://mysupport.netapp.com/NOW/download/software/snapdrive_win/7.1.1/SnapDrive7.1.1_x64.exe.
Note: If the installer prompts you to download a hotfix or patches, complete the activity and then proceed with SnapDrive installation.
2. Launch the installer and click Next.
3. Select the Storage based Licensing method and click Next.
4. Enter your user name and organization information, and click Next.
5. Validate the installation path and click Next.
6. Select the Enable SnapDrive to communicate through the Windows Firewall checkbox and click Next.
7. Enter the Account information for the SnapDrive service account and click Next.
8. Click Next, through the SnapDrive Web Service configuration.
9. Clear the Enable Preferred storage system IP Address checkbox, and click Next.
10. Clear the Enable Transport Protocol Settings checkbox, and click Next.



11. Leave the Enable Unified Manager Configuration checkbox cleared and click Next.
12. Leave the Enable Hyper-V Server Pass-Through Disk checkbox cleared and click Next.
13. Click Install.
14. After the installation is finished, reboot the server to finish the installation.
15. Launch a PowerShell prompt by right-clicking the PowerShell icon in the taskbar and selecting Run as Administrator.

Note: A new prompt is required to register the sdcli executable.

16. Configure the SnapDrive preferred IP settings for each controller.

```
sdcli preferredIP set -f <<var_vserver_name> -IP << var_vserver_mgmt_ip>>
```

17. Configure the SnapDrive transport protocol authentication configuration for each controller.

```
Set-SdStorageConnectionSetting -StorageSystem <<var_vserver_mgmt_ip>> -protocol https -credential vsadmin
```

18. In the pop-up windows, provide the password for the vsadmin account.

Preparing SMB Shares

Prepare the SQL Server and SCVMM Library shares by assigning the required permissions:

1. Open an SSH connection to the cluster IP or host name and log in to the admin user with the password you provided earlier.
2. Remove the Everyone permission from the SQL share.

```
cifs share access-control delete -share scvmmdb -user-or-group Everyone -vserver infra_svm
```

3. Add permissions to the following accounts with NTFS full control permissions over the share:
 - SQL Admins Group

```
share access-control create -share scvmmdb -user-or-group <<var_domain>>\SQL-Admins -permission full_Control -vserver infra_svm
```

4. Remove the Everyone permission from the SCVMM Library share.

```
cifs share access-control delete -share vmmlibrary -user-or-group Everyone -vserver infra_svm
```

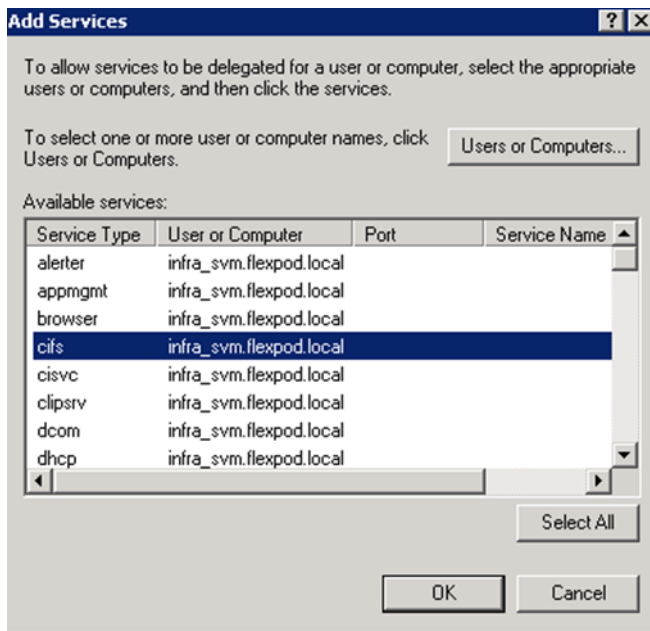
5. Add permissions to the following accounts with NTFS full control permissions over the share:
 - SCVMM SVC Account
 - SCVMM Machine Account

```
share access-control create -share vmmlibrary -user-or-group <<var_domain>>\SCVMM$ -permission full_Control -vserver infra_svm
share access-control create -share vmmlibrary -user-or-group <<var_domain>>\SCVMM-SVC -permission full_Control -vserver infra_svm
```

Configure Constrained Delegation

Although the hosts have the required permissions to access the SMB share, you might encounter access-denied errors when trying to remotely manage the database. To avoid these error messages, configure constrained delegation for the SCVMM host by completing the following steps:

1. In Active Directory Users and Computers, browse to the Computer objects for the SCVMM VM.
2. Right-click the object and select Properties.
3. Select the Delegation tab.
 - a. Select Trust This Computer for Delegation to the Specific Services Only.
 - b. Select Use Kerberos Only.
 - c. Click Add.
4. Click the Users or Computers button on the top of the Add Services page.
5. Enter the name of the infrastructure storage virtual machine (SVM) and click OK.
6. Select CIFS and click OK.

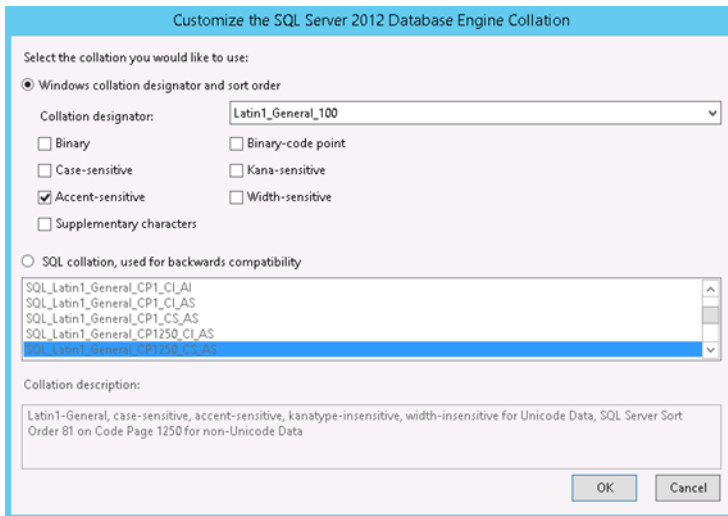


7. Click OK.

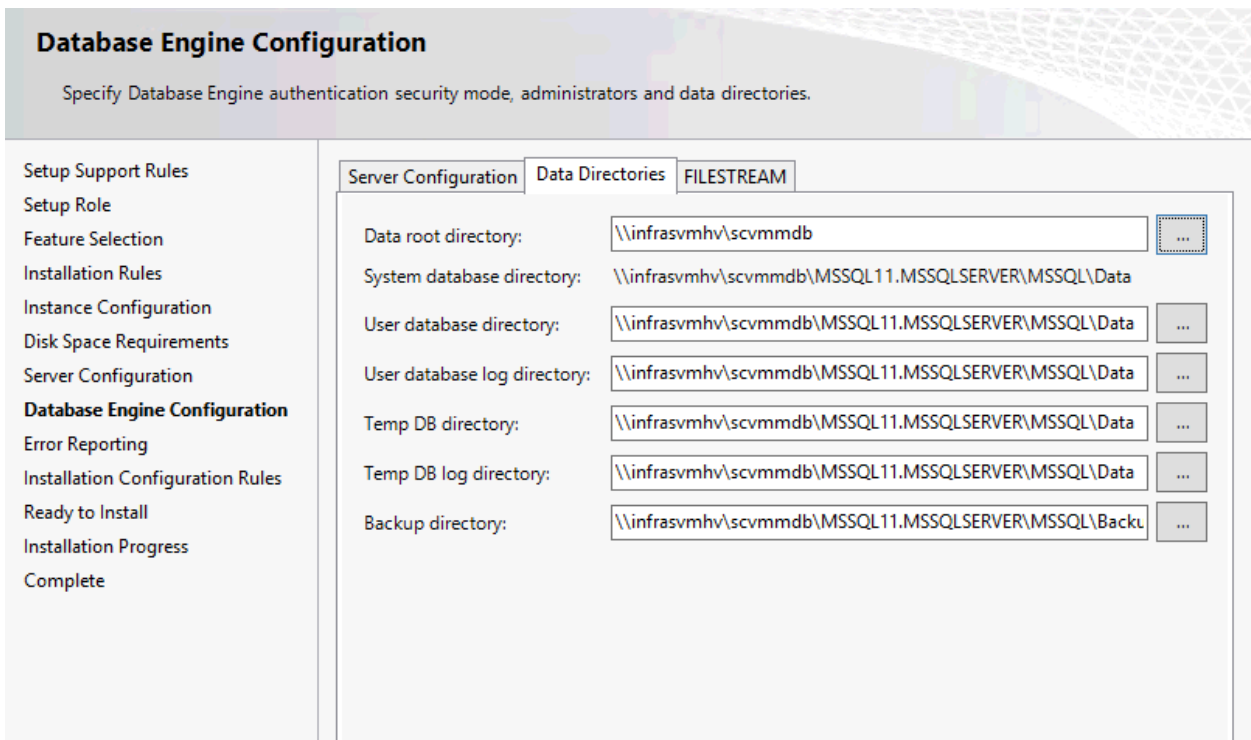
Install SQL Server 2012 SP2

Install SQL Server 2012 SP2 on the SCVMM server. For a deployment of more than 150 hosts, consider using a dedicated SQL cluster. To install SQL onto the SCVMM VM, complete the following steps.

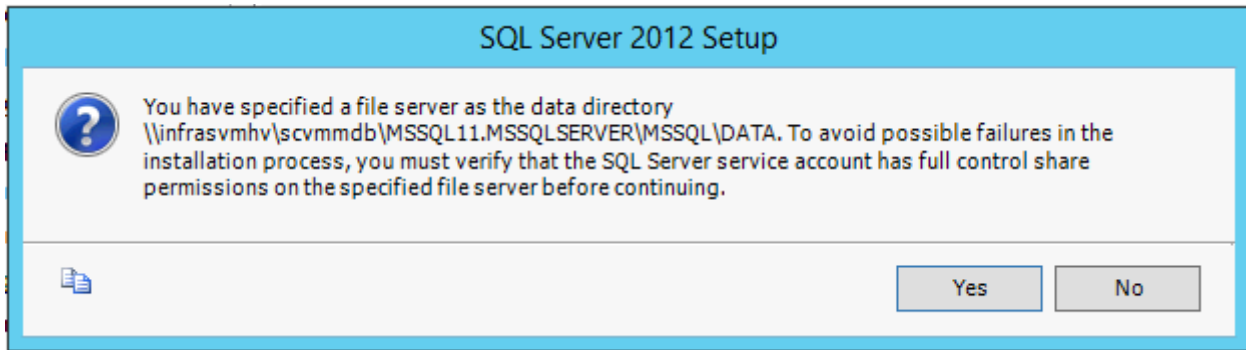
1. Log in to SCVMM VM as SQL Server service account. From the SQL Server 2012 SP2 installation media source, right-click setup.exe and select Run as administrator. The SQL Server Installation Center appears. Select the Installation menu option.
2. From the SQL Server Installation Center, click the New SQL Server standalone installation or add features to an existing installation link.
3. The SQL Server 2012 SP2 Setup wizard appears. In the Setup Support Rules page, verify that each rule shows a Passed status. If any rule requires attention, remediate the issue and rerun the validation check. Click OK.
4. In the Product Key page, select the Enter the product key option and enter the associated product key in the provided text box. Click Next.
Note: If you do not have a product key, select the Specify a free edition option and select Evaluation from the drop-down menu for a 180-day evaluation period.
5. Accept the license terms. Select or clear the Send feature usage data to Microsoft checkbox based on your organization's policies and click Next.
6. In the Product Updates page, select the Include SQL Server product updates checkbox and click Next.
7. In the Install Setup Files page, click Install and allow the support files to install.
8. In the Setup Support Rules page, verify that each rule shows a Passed status. If any rule requires attention, remediate the issue and rerun the validation check.
Note: Those common issues include MSDTC, MSCS, and Windows Firewall warnings. The use of MSDTC is not required for the System Center 2012 SP2 environment.
9. Click Next.
10. In the Setup Role page, select the SQL Server feature Installation option and click Next.
11. In the Feature Selection page, select the following:
 - Database Engine Services
 - Management Tools – Basic
 - Management Tools – Complete
12. In the Installation Rules page, click Next. The Show details and View detailed report can be viewed if required.
13. In the Disk Space Requirements page, verify that you have sufficient disk space and click.
14. In the Server Configuration page, select the Service Accounts tab. Specify the SQL Service account and associated password for the SQL Server Agent and SQL Server Database Engine services.
15. In the Server Configuration page, select the Collation tab. Click Customize.
In the Customize the SQL Server 2012 Database Engine Collation page, click Next.
 - a. Select the Windows collation designator and sort order option.
 - b. Select Latin1_General_100, and select the Accent-sensitive checkbox.
 - c. Click OK to set the collation to Latin1_General_100_CI_AS.



16. In the Database Engine Configuration page, select the Server Configuration tab. In the Authentication Mode section, select the Windows authentication mode option. In the Specify SQL Server administrators section, select the Add Current User button to add the current installation user. Click the Add button, and add the BUILTIN\Administrators, and any other groups who should have administrator access to the SQL Instance.
17. In the same Database Engine Configuration page, select the Data Directories tab. Under the root data directory field enter the UNC path to the SCVMM SQL share. Once complete, click Next.



18. A pop-up opens warning that if the SQL Service account does not have full control of the share, the installation will fail. Click Yes to acknowledge the warning.



19. In the Error Reporting page, select or clear the Send Windows and SQL Server Error Reports to Microsoft or your corporate report server checkbox based on your organization's policies. Click Next.
20. In the Installation Rules page, verify that each rule shows a Passed status. If any rule requires attention, remediate the issue and rerun the validation check. Click Next.
21. In the Ready to Install page, verify all of the settings that were entered during the setup process. Click Install to begin the installation of the SQL Server instance.
22. In the Installation Progress page, the installation progress is displayed.
23. Once complete, the Complete page appears. Click Close to complete the installation of this SQL Server database instance.
24. Verify the installation by inspecting the instances in Failover Cluster Manager and in SQL Server 2012 Management Studio before moving to the next step of installation.

Install System Center 2012 R2 Virtual Machine Manager

Perform the following procedure on one of the Virtual Machine Manager virtual machines.

1. From the System Center 2012 R2 Virtual Machine Manager installation media source, right-click setup.exe and select Run as administrator to begin setup. If prompted by user account control, click Yes to allow the installation to make changes to the computer.
2. The Virtual Machine Manager Installation wizard will begin. At the splash page, click Install to begin the Virtual Machine Manager Server installation.
3. In the Select features to install page, select the VMM management server installation checkbox. After selecting it, the VMM console installation checkbox will be selected by default. Click Next.
4. In the product registration information page, enter the following information and click Next.
 - Name. Specify the name of the primary user or responsible party within your organization.
 - Organization. Specify the name of the licensed organization.
 - Product key. Provide a valid product key for installation of Virtual Machine Manager. If no key is provided, Virtual Machine Manager will be installed in evaluation mode.
5. Accept the license agreement and click Next.
6. In the Join the Customer Experience Improvement Program (CEIP) page, select the option to either participate or not participate in the CEIP by providing selected system information to Microsoft. Click Next.
7. In the Microsoft Update page, select the option to either allow or not allow Virtual Machine Manager to use Microsoft Update to check for and perform Automatic Updates based on your organization's policies. Click Next.
8. In the Select installation location dialog, specify a location or accept the default location of %ProgramFiles%\Microsoft System Center 2012 R2\Virtual Machine Manager for the installation. Click Next.

Note: The setup wizard has a prerequisite checker built in. If for any reason a prerequisite is not met, the setup UI will notify you of the discrepancy. If the system passes the prerequisite check, no screen is displayed and the setup wizard proceeds to the Database configuration screen.

9. In the Database configuration page, enter the following information and click Next.
 - Server name. Specify the name of the SQL Server. Should be the local machine.
 - Port. Specify the TCP port used for the SQL Server. Leave blank if using a local instance.
 - Verify that the Use the following credentials checkbox is cleared. In the Instance name drop-down menu, select the Virtual Machine Manager database instance deployed previously (for example, MSSQLSERVER).
 - In the Select an existing database or create a new database option, select the New database option and accept the default database name of VirtualManagerDB.
10. In the Configure service account and distributed key management page, in the `Virtual Machine Manager Service account` section, select the Domain account option. Enter the following information and click Next.
 - User name and domain. Specify the Virtual Machine Manager service account identified in the previous section in the following format: <DOMAIN>\<USERNAME>.
 - Password. Specify the password for the Virtual Machine Manager service account identified previously.
 - In the Distributed Key Management section, select the Store my keys in Active Directory checkbox. In the provided text box, type the distinguished name (DN) location previously created within Active Directory: cn=VMMDKM,DC=domain,...

Configure service account and distributed key management

Virtual Machine Manager Service Account

Select the account to be used by the VMM service. Highly available VMM installations require the use of a domain account.
[Which type of account should I use?](#)

☐ Local System account
☒ Domain account

User name and domain: Password:

Distributed Key Management

Select whether to store encryption keys in Active Directory instead of on the local machine. Highly available VMM installations require the keys be stored in Active Directory.

☒ Store my keys in Active Directory

Provide the location in Active Directory. For example, CN=DKM,DC=contoso,DC=com.

[How do I configure distributed key management?](#)

11. In the Port configuration page, accept the default values in the provided text boxes and click Next:
 - Communication with the VMM console—default: 8100
 - Communication to agents on hosts and library servers—default: 5985
 - File transfers to agents on hosts and library servers—default: 443
 - Communication with Windows Deployment Services—default: 8102
 - Communication with Windows Preinstallation Environment (Windows PE) agents—default: 8101

- Communication with Windows PE agent for time synchronization—default: 8103
12. In the Library configuration page, click Next.
 13. The Installation summary page appears and displays the selections made during the installation wizard. Review the options selected and click Install to continue.
The wizard will display the progress while installing features.
 14. After the installation completes, the wizard indicates that the setup was successful. Click Close to complete the installation.
 15. Launch the Virtual Machine Manager console to verify that the installation occurred properly. Verify that the console launches and connects to the Virtual Machine Manager instance installed.

Create VMM Run as Account

1. From the Virtual Machine Manager console, click Settings in the left tree view and click Create Run As Account.
2. Name the account. Provide the Active Directory account name and password with administrator rights to all Hyper-V hosts and clusters.
3. Click OK to create the Run-As Account.

Register SMI-S in SCVMM

To register the NetApp SMI-S provider in SCVMM, complete the following steps.

1. In the Virtual Machine Manager console, navigate to the Fabric pane and expand the Storage node. Select the Providers subnode.
2. Click Add Resources and select Storage Devices from the drop-down menu.
3. In the Add Storage Devices Wizard, select Add a Storage device that is managed by a SMI-S provider, and click Next.
4. Select the SAN and NAS devices discovered and managed by the SMI-S provider and click Next.
5. On the Specify Discovery Scope page:
 - a. Select SMI-S CIMXML for the protocol.
 - b. Enter the IP or FQDN for the SMI-S provider.
 - c. Select the Use Secure Sockets Layer checkbox.
 - d. Click Browse and, in the resulting pop-up, select Create Run As Account.
 - Enter a display name.
 - Enter the user name (for example, SMI-S User).
 - Enter the password.
 - Click OK.

Provide the details for this Run As account

Name:

Description:

User name:
Example: contoso\domainuser or localuser

Password:

Confirm password:

☐ Validate domain credentials

- Click Next.

Specify protocol and address of the storage SMI-S provider

Protocol:

Provider IP address or FQDN:

TCP/IP port:

☒ Use Secure Sockets Layer (SSL) connection

Run As account:

- During the discovery phase a pop-up opens prompting you to import the SMI-S provider's certificate. Click Import.
- After the discovery is complete, the wizard shows all the storage controllers registered with the SMI-S provider. Click Next.
- On the Select Storage Devices page, click Create Classification. Enter a name for the storage pool.
- Check the scvmm_pool0 storage pool and assign a classification. Optionally, you can select and assign classifications for all the shares and pools.

Select storage pools to place under management and assign a classification

Logical unit information will be imported from the selected storage pools. The assigned classification describes the capabilities of the selected storage pools.

Storage Device	Pool ID	Classification	Total Capa...	A
Infra-SVM-HV			1,505.00 GB	1...
<input checked="" type="checkbox"/> hyperv-witness		Cluster	5.00 GB	5...
<input checked="" type="checkbox"/> infrastructure		Gold	500.00 GB	4...
<input checked="" type="checkbox"/> scvmmdb		SQL	500.00 GB	4...
<input checked="" type="checkbox"/> vmmlibrary		Gold	500.00 GB	5...
Infra-SVM-HV			6,244.00 GB	4...
<input type="checkbox"/> boot_luns	ONTAP:956ec188-3431-11e5-...		1,024.00 GB	9...
<input type="checkbox"/> sc_sql_db	ONTAP:956ec188-3431-11e5-...		1,024.00 GB	1...
<input checked="" type="checkbox"/> scvmm_pool0	ONTAP:956ec188-3431-11e5-...	SCVMM Pool	4,096.00 GB	2...
<input type="checkbox"/> witness	ONTAP:956ec188-3431-11e5-...		100.00 GB	9...

- Click Next and Finish to end the wizard.

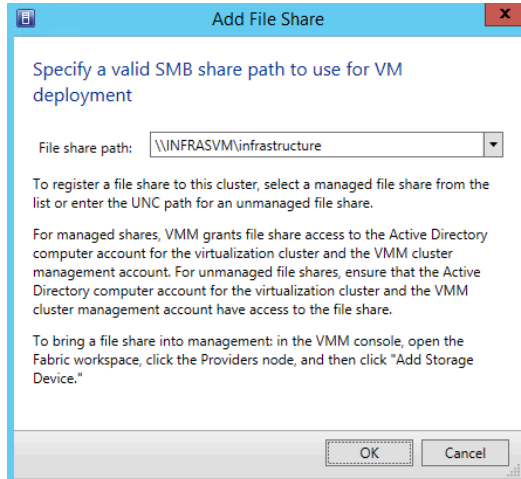
Add Fabric Management Resources Virtual Machine Manager

- Click Fabric in the left tree view and right-click All Hosts under the Servers section. Select Create Host Group and provide a name for the host group.
- Select Fabric and All Hosts. Click Add Resources, Hyper-V Hosts and Clusters.
- In the Indicate the Windows computer location window, select Windows server computers in a trusted Active Directory domain. Click Next.
- Select Use an Existing Run As account and click Browse.
- Select the previously created account and click OK.
- Click Next.
- Enter the cluster name and click Next.
- Click Select All and click Next.
- Select the Host Group created earlier and click Next.
- Click Finish.
- Verify job completion.
- Verify that the hosts were added.

Register File Share to Management Cluster

If the infrastructure VMs were provisioned onto an SMB share, complete the following steps to register the file share to the management cluster.

1. Click Fabric in the left tree view.
2. Expand Servers, All Hosts, and Management Fabric.
3. Right-click the management cluster and select Properties.
4. Select File Share Storage and click Add.
5. From the drop-down menu select the infrastructure share and click OK.



6. Click OK to register the file share.

Assign Permissions

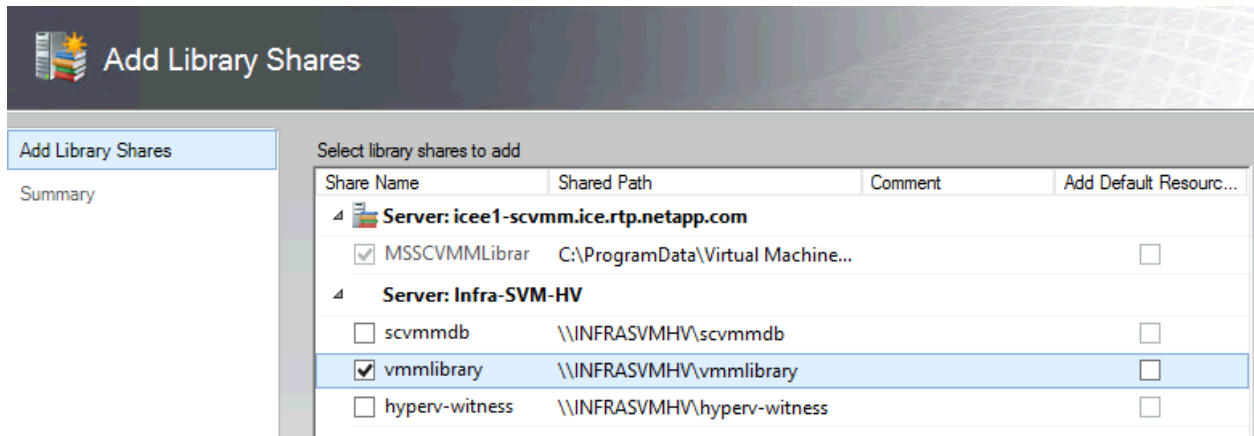
1. Open an SSH connection to the NetApp storage cluster IP or host name and log in to the admin user with the password you provided previously.
2. Add NTFS full control permissions for the following accounts over the share:
 - VMM service account
 - VMM admins group
 - VMM computer accounts

```
share access-control create -share vmmlibrary -user-or-group <<var_domain>>\SCVMM01$ -permission full_Control -vserver infra_svm
share access-control create -share vmmlibrary -user-or-group <<var_domain>>\SCVMM-Admins -permission full_Control -vserver infra_svm
share access-control create -share vmmlibrary -user-or-group <<var_domain>>\VMM-SVC -permission full_Control -vserver infra_svm
```

Add Library Share to SCVMM Library Server

1. From the SCVMM management console, select Library.
2. Browse to Library Servers, right-click the SCVMM library, and select Properties.
3. In the SCVMM Properties window, find the Library Management Credentials entry and click Browse.
4. Select the Action account created in the section “Create VMM Run as Account.” Click OK to select the account.
5. Click OK to save the account selection.
6. Right-click the SCVMM library and select Add Library Shares.

7. In the Add Library Server wizard, select the vmmlibrary share configured previously.



8. Click Next. Review the Summary page and click Add Library Shares.

6 Bill of Materials

This section details the hardware and software components used in validating the large FlexPod Express configuration.

Table 11) Large configuration components.

Part Number	Product Description	Quantity Required
Cisco Components		
Network Switches		
N3K-C3524P-10G	Cisco Nexus 3524 24 10G Ports	2
N2200-PAC-400W	N2K/N3K AC Power Supply Std airflow (port side exhaust)	4
CAB-C13-C14-AC	Power cord C13 to C14 (recessed receptacle) 10A	4
N3548-24P-LIC	Cisco Nexus 3524 Factory Installed 24 port license	2
N3K-C3064-ACC-KIT	Cisco Nexus 3064PQ Accessory Kit	2
N3548-BAS1K9	Cisco Nexus 3500 Base License	2
NXA-FAN-30CFM-F	Cisco Nexus 2K/3K Single Fan forward airflow (port side exhaust)	8
N3KUK9-602A1.1D	NX-OS Release 6.0(2)A1(1d)	2
CON-SNT-3524P10G	Cisco SMARTNET 8X5XNBD Nexus 3524 24 10G Ports	2
Cisco UCS Compute		
UCSC-C220-M4L	UCS C220 M4 LFF w/o CPU mem HD PCIe PSU rail kit	4
UCS-CPU-E52640D	2.60 GHz E5-2640 v3/90W 8C/20MB Cache/DDR4 1866MHz	8
UCS-MR-1X162RU-A	16GB DDR4-2133-MHz RDIMM/PC4-17000/dual rank/x4/1.2v	32
UCSC-MLOM-CSC-02	Cisco UCS VIC1227 VIC MLOM - Dual Port 10Gb SFP+	4

CAB-C13-C14-AC	Power cord C13 to C14 (recessed receptacle) 10A	8
UCSC-PSU1-770W	770W AC Hot-Plug Power Supply for 1U C-Series Rack Server	8
UCSC-BBLKD-L	3.5-inch HDD Blanking Panel	16
UCSC-HS-C220M4	Heat sink for UCS C220 M4 rack servers	8
UCSC-RAILB-M4	Ball Bearing Rail Kit for C220 M4 and C240 M4 rack servers	4
C1UCS-OPT-OUT	Cisco ONE Data Center Compute Opt Out Option	4
CON-SNT-C220M4L	SMARTNET 8X5XNBD UCS C220 M4 LFF w/o CPU mem HD	4
NetApp Components		
FAS2552A-001-R6	FAS2552 High Availability System	2
X80101A-R6-C	Bezel,FAS2552,R6,-C	1
FAS2552-213-R6-C	FAS2552,24x900GB,10K,-C	1
X1558A-R6-C	Power Cable,In-Cabinet,48-IN,C13-C14,-C	2
SVC-FLEXPOD-SYSTEMS	Systems Used in FlexPod Solution, Attach PN	1
X6560-R6-C	Cable,Ethernet,0.5m RJ45 CAT6,-C	1
X1983-3-R6	Cable,Twinax CU,SFP+,3M,X1962/X1963/X1968	4
X6557-EN-R6-C	Cbl,SAS Cntrl-Shelf/Shelf-Shelf/HA,0.5m,EN,-C	2
X6566B-2-R6	Cable,Direct Attach CU SFP+ 10G,2M	2
DOC-2552-C	Documents,2552,-C	1
X5526A-R6-C	Rackmount Kit,4-Post,Universal,-C,R6	1
OS-ONTAP-CAP2-1P-C	OS Enable,Per-0.1TB,ONTAP,Perf-Stor,1P,-C	216
SWITCHLESS	2-Node Switchless Cluster	1
SW-2-2552A-SMGR-C	SW-2,SnapManager Suite,2552A,-C	2
SW-2-2552A-SRESTORE-C	SW-2,SnapRestore,2552A,-C	2
SW-2-2552A-FLEXCLN-C	SW-2,FlexClone,2552A,-C	2
SW-2-2552A-ISCSI-C	SW-2,iSCSI,2552A,-C	2
SW-ONTAP8.2.2-CLM	SW,Data ONTAP8.2.2,Cluster-Mode	2
SW-2-2552A-CIFS-C	SW-2,CIFS,2552A,-C	2
SW-2-2552A-NFS-C	SW-2,NFS,2552A,-C	2
SVC-A2-IN-NBR-E	HW Support,Standard2 Replace,Inst,NBD,e	1
SW-SSP-A2-IN-NBR-E	SW Subs,Standard2 Replace,Inst,NBD,e	1

NetApp Components		
SVC-INST-A2-IN1-NBR-E	Initial Install, Standard2 Replace, Inst, NBD, e	1
CS-OS-SUPPORT-ONTAP	OS Support Entitlement, ONTAP	1
SES-SYSTEM	SupportEdge Standard, Premium or equivalent service from an authorized support services partner ¹	1

Note: The 1Gb management connection to the Cisco Nexus 3524 requires GLC-Ts.

¹SupportEdge Premium is required for cooperative support.

7 Conclusion

FlexPod Express is the optimal shared infrastructure foundation on which to deploy a variety of IT workloads. Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and applications. One common use case is to deploy Microsoft Windows Hyper-V as the virtualization solution, as described in this document. The flexibility and scalability of FlexPod also enable customers to start out with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements.

Acknowledgments

The authors thank John George at NetApp for his contribution to this report.

References

This report references the following documents and resources:

- NetApp FAS2500 storage
<http://www.netapp.com/in/products/storage-systems/fas2500/>
- Cisco UCS C-Series Rack Servers
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html>
- Cisco UCS Virtual Interface Card 1227
<http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1227/index.html>
- Microsoft Windows Hyper-V
<http://www.microsoft.com/en-us/server-cloud/solutions/virtualization.aspx>
- NetApp, Cisco UCS, and Microsoft Windows
<http://support.netapp.com/matrix>

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 © 1994–2015 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.

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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

NVA-0018-DEPLOY-0915