FlexPod Data Center with Microsoft Hyper-V Windows Server 2012 with 7-Mode

Deployment Guide for FlexPod with Microsoft Hyper-V Windows Server 2012 with Data ONTAP 8.1.2 Operating in 7-Mode

Last Updated: November 22, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Mike Mankovsky, Technical Leader, Cisco Systems

Mike Mankovsky is a Cisco Unified Computing System architect, focusing on Microsoft solutions with extensive experience in Hyper-V, storage systems, and Microsoft Exchange Server. He has expert product knowledge in Microsoft Windows storage technologies and data protection technologies.

Glenn Sizemore, Technical Marketing Engineer, NetApp

Glenn Sizemore is a Technical Marketing Engineer in the Microsoft Solutions Group at NetApp, where he specializes in Cloud and Automation. Since joining NetApp, Glenn has delivered a variety of Microsoft based solutions ranging from general best practice guidance to co-authoring the NetApp Hyper-V Cloud Fast Track with Cisco reference architecture.

About Cisco Validated Design (CVD) Program

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

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

FlexPod with Microsoft Hyper-V Windows Server 2012 Deployment Guide

Reference Architecture

FlexPod architecture is highly modular, or pod-like. Although each customer's FlexPod unit might vary in its exact configuration, after a FlexPod unit is built, it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a FlexPod unit) and scaling out (adding additional FlexPod units).

Specifically, FlexPod is a defined set of hardware and software that serves as an integrated foundation for all virtualization solutions. FlexPod validated with Microsoft Server 2012 Hyper-V includes NetApp® FAS3200 Series storage, Cisco Nexus® 5500 Series network switches, the Cisco Unified Computing Systems™ (Cisco UCS™) platforms, and Microsoft virtualization software in a single package. The computing and storage can fit in one data center rack with networking residing in a separate rack or deployed according to a customer's data center design. Due to port density, the networking components can accommodate multiple configurations of this kind.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Figure 1 Architecture overview

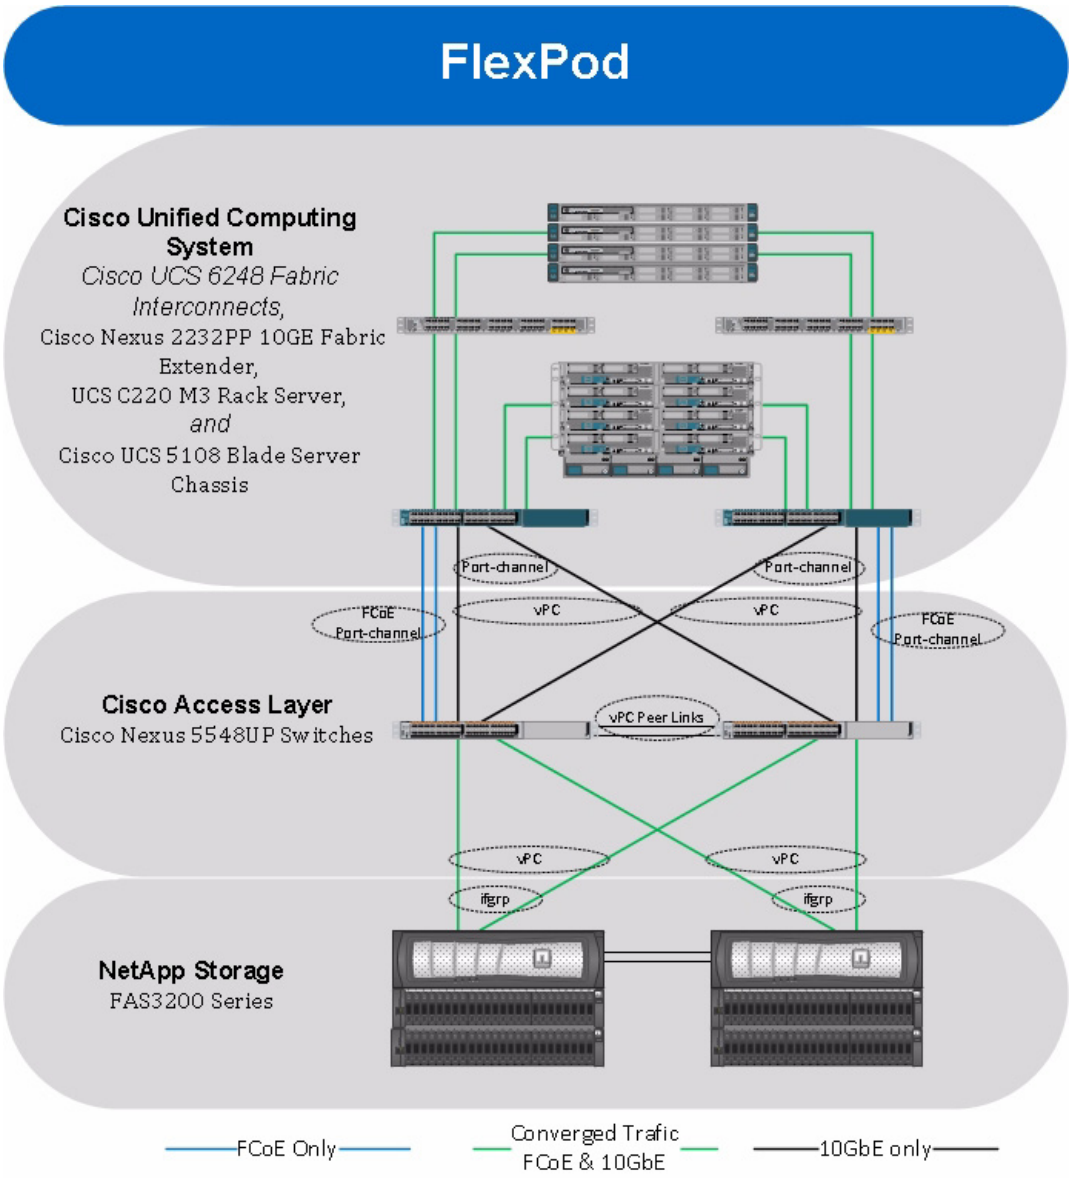

The reference configuration shown in Figure 1 includes:

- Two Cisco Nexus 5548UP Switches
- Two Cisco UCS 6248UP Fabric Interconnects
- Two Cisco Nexus 2232PP Fabric Extenders
- One chassis of Cisco UCS blades with two fabric extenders per chassis
- Four Cisco USC C220M3 Servers
- One FAS3240A (HA pair)

Storage is provided by a NetApp FAS3240A with accompanying disk shelves. All systems and fabric links feature redundancy and provide end-to-end high availability. For server virtualization, the deployment includes Hyper-V. Although this is the base design, each of the components can be scaled flexibly to support specific business requirements. For example, more (or different) blades and chassis could be deployed to increase compute capacity, additional disk shelves could be deployed to improve I/O capacity and throughput, or special hardware or software features could be added to introduce new features.

For more information on FlexPod for Windows Server 2012 Hyper-V design choices and deployment best practices, see FlexPod for Windows Server 2012 Hyper-V Design Guide:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/whitepaper__c07-727095.html

Note This is a sample bill of materials (BoM) only. This solution is certified for use with any configuration that meets the FlexPod Technical Specification rather than for a specific model. FlexPod and Fast Track programs allow customers to choose from within a model family to make sure that each FlexPod for Microsoft Windows Server 2012 Hyper-V solution meets the customers' requirements.

The remainder of this document guides you through the low-level steps for deploying the base architecture, as shown in Figure 1. This includes everything from physical cabling, to compute and storage configuration, to configuring virtualization with Hyper-V.

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration. Therefore, references are made as to which component is being configured with each step, whether it is A or B. For example, Controller A and Controller B, are used to identify the two NetApp storage controllers that are provisioned with this document, while Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. The Cisco UCS Fabric Interconnects are similarly configured. Additionally, this document details steps for provisioning multiple Cisco UCS hosts and these are identified sequentially as VM-Host-Infra-01 and VM-Host-Infra-02, and so on. Finally, to indicate that the reader should include information about their environment in a given step, **<management VLAN ID>** appears as part of the command structure. See the following commands show VLAN creation:

```
controller A> vlan create
```

Usage:

```
vlan create [-g {on|off}] <ifname> <vlanid_list>
vlan add <ifname> <vlanid_list>
vlan delete -q <ifname> [<vlanid_list>]
vlan modify -g {on|off} <ifname>
vlan stat <ifname> [<vlanid_list>]
```

Example:

```
controller A> vlan create vif0 <management VLAN ID>
```

This document is intended to allow readers to fully configure the their environment. In this process, various steps require the reader to insert customer specific naming conventions, IP addresses and VLAN schemes as well as to record appropriate WWPN, WWNN, or MAC addresses. Table 2 provides the list of VLANs necessary for deployment as outlined in this guide.

Note In this document the VM-Data VLAN is used for virtual machine management interfaces.

The VM-Mgmt VLAN is used for management interfaces of the Microsoft Hyper-V hosts. A Layer-3 route must exist between the VM-Mgmt and VM-Data VLANS.

Table 1 Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in This Document
Mgmt	VLAN for management interfaces	805
Native	VLAN to which untagged frames are assigned	2
CSV	VLAN for cluster shared volume	801
iSCSI-A	VLAN for iSCSI traffic for fabric A	802
iSCSI-B	VLAN for iSCSI traffic for fabric B	807
Live Migration	VLAN designated for the movement of VMs from one physical host to another.	803
VM Cluster Comm	VLAN for cluster connectivity	806
Public	VLAn for application data	804

Deployment

This document details the necessary steps to deploy base infrastructure components as well for provisioning Microsoft Hyper-V as the foundation for virtualized workloads. At the end of these deployment steps, you will be prepared to provision applications on top of a Microsoft Hyper-V virtualized infrastructure. The outlined procedure includes:

- Initial NetApp Controller configuration
- Initial Cisco UCS configuration
- Initial Cisco Nexus configuration
- Creation of necessary VLANs and VSANs for management, basic functionality, and virtualized infrastructure specific to the Microsoft
- Creation of necessary vPCs to provide HA among devices
- Creation of necessary service profile pools such as World Wide Port Name (WWPN), World Wide Node Name (WWNN), MAC, server, and so on.
- Creation of necessary service profile policies such as adapter policy, boot policy, and so on.
- Creation of two service profile templates from the created pools and policies, one each for fabric A and B
- Provisioning of four servers from the created service profiles in preparation for OS installation
- Initial configuration of the infrastructure components residing on the NetApp Controller

- Installation of Microsoft Windows Server 2012 Datacenter Edition
- Enabling Microsoft Hyper-V Role
- Configuring FM-FEX and SR-IOV adapters

The FlexPod Validated with Microsoft Private Cloud architecture is flexible; therefore, the configuration detailed provided in this section might vary for customer implementations depending on specific requirements. Although customer implementations might deviate from these details; the best practices, features, and configurations listed in this section can still be used as a reference for building a customized FlexPod, validated with Microsoft Private Cloud architecture.

Cabling Information

The following information is provided as a reference for cabling the physical equipment in a FlexPod environment. The tables include both local and remote device and port locations in order to simplify cabling requirements.

Table 2, Table 3, Table 4, Table 5, and Table 6 contain details for the prescribed and supported configuration of the NetApp FAS3240 running Data ONTAP 8.0.2. This configuration leverages a dual-port 10 Gigabit Ethernet adapter as well as the native FC target ports and the onboard SAS ports for disk shelf connectivity. For any modifications of this prescribed architecture, consult the currently available NetApp Interoperability Matrix Tool (IMT) at:

http://now.netapp.com/matrix

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site.

Be sure to follow the cable directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned.

It is possible to order a FAS3240A system in a different configuration from what is prescribed in the tables in this section. Before starting, be sure the configuration matches what is described in the tables and diagrams in this section

Figure 2 shows a FlexPod cabling diagram. The labels indicate connections to end points rather than port numbers on the physical device. For example, connection 1 is an FCoE target port connected from NetApp controller A to Nexus 5548 A. SAS connections 23, 24, 25, and 26 as well as ACP connections 27 and 28 should be connected to the NetApp storage controller and disk shelves according to best practices for the specific storage controller and disk shelf quantity.

Figure 2 Flexpod Cabling Diagram

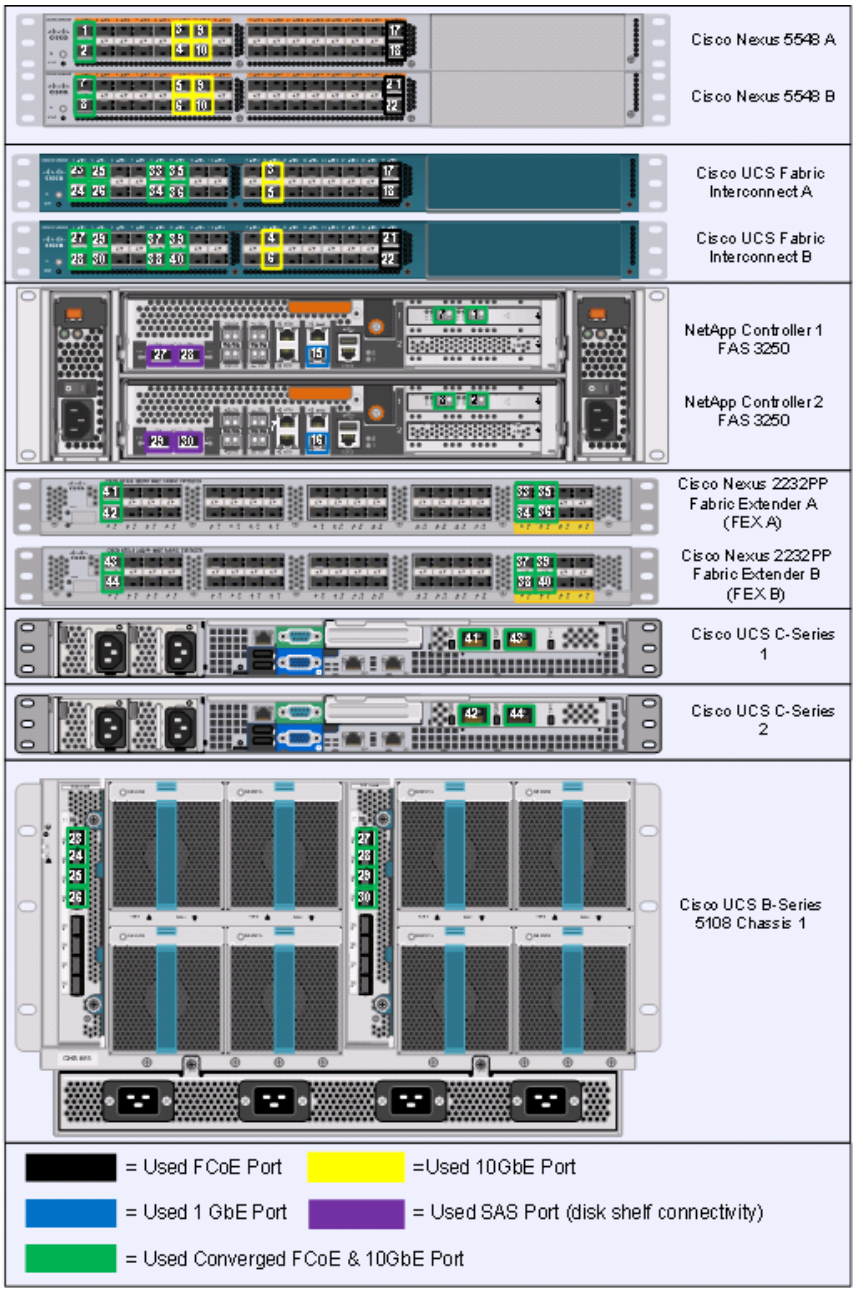

Table 2 Cisco Nexus 5548 A Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 A	Eth1/1	FCoE/10GbE	NetApp Controller A	2a
	Eth1/2	FCoE/10GbE	NetApp Controller B	2a
	Eth1/13	10GbE	Cisco Nexus 5548 B	Eth1/13
	Eth1/14	10GbE	Cisco Nexus 5548 B	Eth1/14
	Eth1/11	10GbE	Cisco UCS Fabric Interconnect A	Eth1/19
	Eth1/12	10GbE	Cisco UCS Fabric Interconnect B	Eth1/19
	Eth1/31	FCoE	Cisco UCS Fabric Interconnect A	Eth1/31
	Eth1/32	FCoE	Cisco UCS Fabric Interconnect A	Eth1/32
	MGMT0	100MbE	100MbE Management Switch	Any

Table 3 Cisco Nexus 5548 B Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 B	Eth1/1	FCoE/10GbE	NetApp Controller A	2b
	Eth1/2	FCoE/10GbE	NetApp Controller B	2b
	Eth1/13	10GbE	Cisco Nexus 5548 A	Eth1/13
	Eth1/14	10GbE	Cisco Nexus 5548 A	Eth1/14
	Eth1/11	10GbE	Cisco UCS Fabric Interconnect A	Eth1/20
	Eth1/12	10GbE	Cisco UCS Fabric Interconnect B	Eth1/20
	Eth1/31	FCoE	Cisco UCS Fabric Interconnect B	Eth1/31
	Eth1/32	FCoE	Cisco UCS Fabric Interconnect B	Eth1/32
	MGMT0	100MbE	100MbE Management Switch	Any

Table 4 NetApp Controller A Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp Controller A	e0M	100MbE	100MbE Management Switch	Any
	e0P	1GbE	SAS Shelves	ACP Port
	2a	FCoE/10GbE	Cisco Nexus 5548 A	Eth 1/1
	2b	FCoE/10GbE	Cisco Nexus 5548 B	Eth 1/1

Table 5 NetApp Controller B Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp Controller B	e0M	100MbE	100MbE Management Switch	Any
	e0P	1GbE	SAS Shelves	ACP Port
	2a	FCoE/10GbE	Cisco Nexus 5548 A	Eth 1/1
	2b	FCoE/10GbE	Cisco Nexus 5548 B	Eth 1/1

Table 6 Cisco UCS Fabric Interconnect A Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/11
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/11
	Eth1/1	FCoE/10GbE	Chassis 1 FEX 2208 A	Port 1
	Eth1/2	FCoE/10GbE	Chassis 1 FEX 2208 A	Port 2
	Eth1/3	FCoE/10GbE	Chassis 1 FEX 2208 A	Port 3
	Eth1/4	FCoE/10GbE	Chassis 1 FEX 2208 A	Port 4
	Eth1/5	FCoE/10GbE	Chassis 2 FEX 2208 A (if required)	Port 1
	Eth1/6	FCoE/10GbE	Chassis 2 FEX 2208 A (if required)	Port 2
	Eth1/7	FCoE/10GbE	Chassis 2 FEX 2208 A (if required)	Port 3
	Eth1/8	FCoE/10GbE	Chassis 2 FEX 2208 A (if required)	Port 4
	Eth1/9	FCoE/10GbE	FEX 2232PP A (if required)	Port 1
	Eth1/10	FCoE/10GbE	FEX 2232PP A (if required)	Port 2
	Eth1/11	FCoE/10GbE	FEX 2232PP A (if required)	Port 3
	Eth1/12	FCoE/10GbE	FEX 2232PP A (if required)	Port 4
	Eth1/31	FCoE	Cisco Nexus 5548 A	Eth1/31
	Eth1/32	FCoE	Cisco Nexus 5548 A	Eth1/32
	MGMT0	1GbE	1GbE Management Switch	Any
	L1	1GbE	Cisco UCS Fabric Interconnect B	L1
	L2	1GbE	Cisco UCS Fabric Interconnect B	L2

Table 7 Cisco UCS Fabric Interconnect B Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/19	10GbE	Cisco Nexus 5548 A	Eth1/12
	Eth1/20	10GbE	Cisco Nexus 5548 B	Eth1/12
	Eth1/1	FCoE/10GbE	Chassis 1 FEX 2208 B	Port 1
	Eth1/2	FCoE/10GbE	Chassis 1 FEX 2208 B	Port 2
	Eth1/3	FCoE/10GbE	Chassis 1 FEX 2208 B	Port 3
	Eth1/4	FCoE/10GbE	Chassis 1 FEX 2208 B	Port 4
	Eth1/5	FCoE/10GbE	Chassis 2 FEX 2208 B (if required)	Port 1
	Eth1/6	FCoE/10GbE	Chassis 2 FEX 2208 B (if required)	Port 2
	Eth1/7	FCoE/10GbE	Chassis 2 FEX 2208 B (if required)	Port 3
	Eth1/8	FCoE/10GbE	Chassis 2 FEX 2208 B (if required)	Port 4
	Eth1/9	FCoE/10GbE	FEX 2232PP B (if required)	Port 1
	Eth1/10	FCoE/10GbE	FEX 2232PP B (if required)	Port 2
	Eth1/11	FCoE/10GbE	FEX 2232PP B (if required)	Port 3
	Eth1/12	FCoE/10GbE	FEX 2232PP B (if required)	Port 4
	Eth1/31	FCoE	Cisco Nexus 5548 B	Eth1/31
	Eth1/32	FCoE	Cisco Nexus 5548 B	Eth1/32
	MGMT0	1GbE	1GbE Management Switch	Any
	L1	1GbE	Cisco UCS Fabric Interconnect A	L1
	L2	1GbE	Cisco UCS Fabric Interconnect A	L2

Table 8 Cisco Nexus 2232PP Fabric Extender A Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX A	Eth1/1	FCoE/10GbE	Cisco UCS C-Series 1	Port 1
	Eth1/2	FCoE/10GbE	Cisco UCS C-Series 2	Port 1

Table 9 Cisco Nexus 2232PP Fabric Extender B Ethernet cabling information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX B	Eth1/1	FCoE/10GbE	Cisco UCS C-Series 1	Port 2
	Eth1/2	FCoE/10GbE	Cisco UCS C-Series 2	Port 2

Nexus 5548UP Deployment Procedure

The following section provides a detailed procedure for configuring the Cisco Nexus 5548 switches for use in a FlexPod environment. Follow these steps precisely because failure to do so could result in an improper configuration.

Note The configuration steps detailed in this section provides guidance for configuring the Nexus 5548UP running release 5.2(1)N1(3). This configuration also leverages the native VLAN on the trunk ports to discard untagged packets, by setting the native VLAN on the Port Channel, but not including this VLAN in the allowed VLANs on the Port Channel.

Initial Cisco Nexus 5548UP Switch Configuration

These steps provide details for the initial Cisco Nexus 5548 Switch setup.

Nexus 5548 A

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start.

1. Enter yes to enforce secure password standards.
2. Enter the password for the admin user.
3. Enter the password a second time to commit the password.
4. Enter yes to enter the basic configuration dialog.
5. Create another login account (yes/no) [n]: Enter.
6. Configure read-only SNMP community string (yes/no) [n]: Enter.
7. Configure read-write SNMP community string (yes/no) [n]: Enter.
8. Enter the switch name: <Nexus A Switch name> Enter.
9. Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter.
10. Mgmt0 IPv4 address: <Nexus A mgmt0 IP> Enter.
11. Mgmt0 IPv4 netmask: <Nexus A mgmt0 netmask> Enter.
12. Configure the default gateway? (yes/no) [y]: Enter.
13. IPv4 address of the default gateway: <Nexus A mgmt0 gateway> Enter.
14. Enable the telnet service? (yes/no) [n]: Enter.
15. Enable the ssh service? (yes/no) [y]: Enter.
16. Type of ssh key you would like to generate (dsa/rsa):rsa.

17. Number of key bits <768–2048>:1024 Enter.
18. Configure the ntp server? (yes/no) [y]: Enter.
19. NTP server IPv4 address: <NTP Server IP> Enter.
20. Enter basic FC configurations (yes/no) [n]: Enter.
21. Would you like to edit the configuration? (yes/no) [n]: Enter.

Note Be sure to review the configuration summary before enabling it.

22. Use this configuration and save it? (yes/no) [y]: Enter.
23. Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus A.
24. Log in as user `admin` with the password previously entered.

Nexus 5548 B

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start.

1. Enter yes to enforce secure password standards.
2. Enter the password for the admin user.
3. Enter the password a second time to commit the password.
4. Enter yes to enter the basic configuration dialog.
5. Create another login account (yes/no) [n]: Enter.
6. Configure read-only SNMP community string (yes/no) [n]: Enter.
7. Configure read-write SNMP community string (yes/no) [n]: Enter.
8. Enter the switch name: <Nexus B Switch name> Enter.
9. Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter.
10. Mgmt0 IPv4 address: <Nexus B mgmt0 IP> Enter.
11. Mgmt0 IPv4 netmask: <Nexus B mgmt0 netmask> Enter.
12. Configure the default gateway? (yes/no) [y]: Enter.
13. IPv4 address of the default gateway: <Nexus B mgmt0 gateway> Enter.
14. Enable the telnet service? (yes/no) [n]: Enter.
15. Enable the ssh service? (yes/no) [y]: Enter.
16. Type of ssh key you would like to generate (dsa/rsa):rsa.
17. Number of key bits <768–2048>:1024 Enter.
18. Configure the ntp server? (yes/no) [y]: Enter.
19. NTP server IPv4 address: <NTP Server IP> Enter.
20. Enter basic FC configurations (yes/no) [n]: Enter.
21. Would you like to edit the configuration? (yes/no) [n]: Enter.

Note Be sure to review the configuration summary before enabling it.

22. Use this configuration and save it? (yes/no) [y]: Enter.
23. Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus A.
24. Log in as user admin with the password previously entered.

Enable Appropriate Cisco Nexus Features

These steps provide details for enabling the appropriate Cisco Nexus features.

Nexus A and Nexus B

1. Type config t to enter the global configuration mode.
2. Type feature lacp.
3. Type feature fcoe.
4. Type feature npiv.
5. Type feature vpc.
6. Type feature fport-channel-trunk.

Set Global Configurations

These steps provide details for setting global configurations.

Nexus A and Nexus B

1. From the global configuration mode, type spanning-tree port type network default to make sure that, by default, the ports are considered as network ports in regards to spanning-tree.
2. Type spanning-tree port type edge bpduguard default to enable bpduguard on all edge ports by default.
3. Type spanning-tree port type edge bpdufilter default to enable bpdufilter on all edge ports by default.
4. Type ip access-list classify_Silver.
5. Type 10 permit ip <iSCSI-A net address> any

Note Where the variable is the network address of the iSCSI-A VLAN in CIDR notation (i.e. 192.168.102.0/24).

6. Type 20 permit ip any <iSCSI-A net address>.
7. Type 30 permit ip <iSCSI-B net address> any.
8. Type 40 permit ip any <iSCSI-B net address>.
9. Type exit.
10. Type class-map type qos match-all class-gold.
11. Type match cos 4.
12. Type exit.

13. Type class-map type qos match-all class-silver.
14. Type match cos 2.
15. Type match access-group name classify_Silver.
16. Type exit.
17. Type class-map type queuing class-gold.
18. Type match qos-group 3.
19. Type exit.
20. Type class-map type queuing class-silver.
21. Type match qos-group 4.
22. Type exit.
23. Type policy-map type qos system_qos_policy.
24. Type class class-gold.
25. Type set qos-group 3.
26. Type class class-silver.
27. Type set qos-group 4.
28. Type class class-fcoe.
29. Type set qos-group 1.
30. Type exit.
31. Type exit.
32. Type policy-map type queuing system_q_in_policy.
33. Type class.
34. Type queuing class-fcoe.
35. Type bandwidth percent 20.
36. Type class type queuing class-gold.
37. Type bandwidth percent 33.
38. Type class type queuing class-silver.
39. Type bandwidth percent 29.
40. Type class type queuing class-default.
41. Type bandwidth percent 18.
42. Type exit.
43. Type exit.
44. Type policy-map type queuing system_q_out_policy.
45. Type class type queuing class-fcoe.
46. Type bandwidth percent 20.
47. Type class type queuing class-gold.
48. Type bandwidth percent 33.
49. Type class type queuing class-silver.
50. Type bandwidth percent 29.

51. Type class type queuing class-default.
52. Type bandwidth percent 18.
53. Type exit.
54. Type exit.
55. Type class-map type network-qos class-gold.
56. Type match qos-group 3.
57. Type exit.
58. Type class-map type network-qos class-silver.
59. Type match qos-group 4.
60. Type exit.
61. Type policy-map type network-qos system_nq_policy.
62. Type class type network-qos class-gold.
63. Type set cos 4.
64. Type mtu 9000.
65. Type class type network-qos class-fcoe.
66. Type pause no-drop.
67. Type mtu 2158.
68. Type class type network-qos class-silver.
69. Type set cos 2.
70. Type mtu 9000.
71. Type class type network-qos class-default.
72. Type mtu 9000.
73. Type exit.
74. Type system qos.
75. Type service-policy type qos input system_qos_policy.
76. Type service-policy type queuing input system_q_in_policy.
77. Type service-policy type queuing output system_q_out_policy.
78. Type service-policy type network-qos system_nq_policy.
79. Type exit.
80. Type copy run start.

Create Necessary VLANs

These steps provide details for creating the necessary VLANs.

Nexus A

1. Type vlan <<Fabric_A_FCoE_VLAN ID>>.
2. Type name FCoE_Fabric_A.

3. Type exit.

Nexus B

1. Type vlan <<Fabric_B_FCoE_VLAN ID>>.
2. Type name FCoE_Fabric_B.
3. Type exit.

Nexus A and Nexus B

1. Type vlan <<Native VLAN ID>>.
2. Type name Native-VLAN.
3. Type exit.
4. Type vlan <<CSV VLAN ID>>.
5. Type name CSV-VLAN.
6. Type exit.
7. Type vlan <<Live Migration VLAN ID>>.
8. Type name Live-Migration-VLAN.
9. Type exit.
10. Type vlan <<iSCSI A VLAN ID>>.
11. Type name iSCSI-A-VLAN.
12. Type exit.
13. Type vlan <<iSCSI B VLAN ID>>.
14. Type name iSCSI-B-VLAN.
15. Type exit.
16. Type vlan <<MGMT VLAN ID>>.
17. Type name Mgmt-VLAN.
18. Type exit.
19. Type vlan <<VM Data VLAN ID>>.
20. Type name VM-Public-VLAN.
21. Type exit.
22. Type vlan <<VM Cluster Comm VLAN ID>>.
23. Type name VM-Cluster-Comm-VLAN.
24. Type exit.

Add Individual Port Descriptions for Troubleshooting

These steps provide details for adding individual port descriptions for troubleshooting activity and verification.

Nexus 5548 A

1. From the global configuration mode,

2. Type interface Eth1/1.
3. Type description <Controller A:e2a>.
4. Type exit.
5. Type interface Eth1/2.
6. Type description <Controller B:e2a>.
7. Type exit.
8. Type interface Eth1/3.
9. Type description <UCSM A:Eth1/19>.
10. Type exit.
11. Type interface Eth1/4.
12. Type description <UCSM B:Eth1/19>.
13. Type exit.
14. Type interface Eth1/5.
15. Type description <Nexus B:Eth1/5>.
16. Type exit.
17. Type interface Eth1/6.
18. Type description <Nexus B:Eth1/6>.
19. Type exit.

Nexus 5548 B

1. From the global configuration mode,
2. Type interface Eth1/1.
3. Type description <Controller A:e2b>.
4. Type exit.
5. Type interface Eth1/2.
6. Type description <Controller B:e2b>.
7. Type exit.
8. Type interface Eth1/3.
9. Type description <UCSM A:Eth1/20>.
10. Type exit.
11. Type interface Eth1/4.
12. Type description <UCSM B:Eth1/20>.
13. Type exit.
14. Type interface Eth1/5.
15. Type description <Nexus A:Eth1/5>.
16. Type exit.
17. Type interface Eth1/6.
18. Type description <Nexus A:Eth1/6>.

19. Type exit.

Create Necessary Port Channels

These steps provide details for creating the necessary Port Channels between devices.

Nexus 5548 A

1. From the global configuration mode,
2. Type interface Po10.
3. Type description vPC peer-link.
4. Type exit.
5. Type interface Eth1/5-6.
6. Type channel-group 10 mode active.
7. Type no shutdown.
8. Type exit.
9. Type interface Po11.
10. Type description <Controller A>.
11. Type exit.
12. Type interface Eth1/1.
13. Type channel-group 11 mode active.
14. Type no shutdown.
15. Type exit.
16. Type interface Po12.
17. Type description <Controller B>.
18. Type exit.
19. Type interface Eth1/2.
20. Type channel-group 12 mode active.
21. Type no shutdown.
22. Type exit.
23. Type interface Po13.
24. Type description <UCSM A>.
25. Type exit.
26. Type interface Eth1/3.
27. Type channel-group 13 mode active.
28. Type no shutdown.
29. Type exit.
30. Type interface Po14.
31. Type description <UCSM B>.

32. Type exit.
33. Type interface Eth1/4.
34. Type channel-group 14 mode active.
35. Type no shutdown.
36. Type exit.
37. Type interface eth1/31.
38. Type switchport description <UCSM A:eth1/31>.
39. Type exit.
40. Type interface eth1/32.
41. Type switchport description <UCSM A:eth1/32>.
42. Type exit.
43. Type interface Eth1/31-32.
44. Type channel-group 15 mode active.
45. Type no shutdown.
46. Type copy run start.

Nexus 5548 B

1. From the global configuration mode, type `interface Po10`.
2. Type description vPC peer-link.
3. Type exit.
4. Type interface Eth1/5-6.
5. Type channel-group 10 mode active.
6. Type no shutdown.
7. Type exit.
8. Type interface Po11.
9. Type description <Controller A>.
10. Type exit.
11. Type interface Eth1/1.
12. Type channel-group 11 mode active.
13. Type no shutdown.
14. Type exit.
15. Type interface Po12.
16. Type description <Controller B>.
17. Type exit.
18. Type interface Eth1/2.
19. Type channel-group 12 mode active.
20. Type no shutdown.
21. Type exit.

22. Type interface Po13.
23. Type description <UCSM A>.
24. Type exit.
25. Type interface Eth1/3.
26. Type channel-group 13 mode active.
27. Type no shutdown.
28. Type exit.
29. Type interface Po14.
30. Type description <UCSM B>.
31. Type exit.
32. Type interface Eth1/4.
33. Type channel-group 14 mode active.
34. Type no shutdown
35. Type exit.
36. Type interface eth1/31.
37. Type switchport description <UCSM B:eth1/31>.
38. Type exit.
39. Type interface eth1/32.
40. Type switchport description <UCSM B:eth1/32>.
41. Type exit.
42. Type interface eth1/31-32.
43. Type channel-group 16 mode active.
44. Type no shutdown.
45. Type copy run start.

Add Port Channel Configurations

These steps provide details for adding Port Channel configurations.

Nexus 5548 A

1. From the global configuration mode,
2. Type interface Po10.
3. Type switchport mode trunk.
4. Type switchport trunk native vlan <<Native VLAN ID>>.
5. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<VM Cluster Comm VLAN ID>> <<Fabric A FCoE VLAN ID>>.
6. Type spanning-tree port type network.
7. Type no shutdown.

8. Type exit.
9. Type interface Po11.
10. Type switchport mode trunk.
11. Type switchport trunk native vlan <<Native VLAN ID>>.
12. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Fabric A FCoE VLAN ID>>.
13. Type spanning-tree port type edge trunk.
14. Type no shut.
15. Type exit.
16. Type interface Po12.
17. Type switchport mode trunk.
18. Type switchport trunk native vlan <<Native VLAN ID>>.
19. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Fabric A FCoE VLAN ID>>.
20. Type spanning-tree port type edge trunk.
21. Type no shut.
22. Type exit.
23. Type interface Po13.
24. Type switchport mode trunk.
25. Type switchport trunk native vlan <Native VLAN ID>.
26. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<VM Cluster Comm VLAN ID>> <<Fabric A FCoE VLAN ID>>.
27. Type spanning-tree port type edge trunk.
28. Type no shut.
29. Type exit.
30. Type interface Po14
31. Type switchport mode trunk
32. Type switchport trunk native vlan <Native VLAN ID>.
33. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<VM Cluster Comm VLAN ID>> <<Fabric A FCoE VLAN ID>>.
34. Type spanning-tree port type edge trunk.
35. Type no shutdown.
36. Type exit.
37. Type interface Po15.
38. Type switchport mode trunk.
39. Type switchport trunk allowed vlan <Fabric A FCoE VLAN ID>
40. Type no shutdown

41. Type exit.
42. Type copy run start.

Nexus 5548 B

1. From the global configuration mode,
2. Type interface Po10.
3. Type switchport mode trunk.
4. Type switchport trunk native vlan <<Native VLAN ID>>.
5. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<Fabric B FCoE VLAN ID>>.
6. Type spanning-tree port type network.
7. Type no shutdown.
8. Type exit.
9. Type interface Po11.
10. Type switchport mode trunk.
11. Type switchport trunk native vlan <<Native VLAN ID>>.
12. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Fabric B FCoE VLAN ID>>.
13. Type spanning-tree port type edge trunk.
14. Type no shut.
15. Type exit.
16. Type interface Po12.
17. Type switchport mode trunk.
18. Type switchport trunk native vlan <<Native VLAN ID>>.
19. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Fabric B FCoE VLAN ID>>.
20. Type spanning-tree port type edge trunk.
21. Type no shut.
22. Type exit.
23. Type interface Po13.
24. Type switchport mode trunk.
25. Type switchport trunk native vlan <Native VLAN ID>.
26. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<Fabric B FCoE VLAN ID>>.Type spanning-tree port type edge trunk.
27. Type no shut.
28. Type exit.
29. Type interface Po14.
30. Type switchport mode trunk.

31. Type switchport trunk native vlan <Native VLAN ID>.
32. Type switchport trunk allowed vlan <<MGMT VLAN ID>>, <<CSV VLAN ID>>, <<iSCSI A VLAN ID>>, <<iSCSI B VLAN ID>>, <<Live Migration VLAN ID>>, <<VM Data VLAN ID>>, <<Fabric B FCoE VLAN ID>>.Type spanning-tree port type edge trunk.
33. Type no shut.
34. Type exit.
35. Type interface Po16.
36. Type switchport mode trunk.
37. Type switchport trunk allowed vlan <Fabric B FCoE VLAN ID>
38. Type no shutdown.
39. Type exit.
40. Type copy run start.

Configure Virtual Port Channels

These steps provide details for configuring virtual Port Channels (vPCs).

Nexus 5548 A

1. From the global configuration mode,
2. Type vpc domain <Nexus vPC domain ID>.
3. Type role priority 10.
4. Type peer-keepalive destination <Nexus B mgmt0 IP> source <Nexus A mgmt0 IP>.
5. Type exit.
6. Type interface Po10.
7. Type vpc peer-link.
8. Type exit.
9. Type interface Po11.
10. Type vpc 11.
11. Type exit.
12. Type interface Po12.
13. Type vpc 12.
14. Type exit.
15. Type interface Po13.
16. Type vpc 13.
17. Type exit.
18. Type interface Po14.
19. Type vpc 14.
20. Type exit.
21. Type copy run start.

Nexus 5548 B

1. From the global configuration mode, type vpc domain <Nexus vPC domain ID>.
2. Type role priority 20.
3. Type peer-keepalive destination <Nexus A mgmt0 IP> source <Nexus B mgmt0 IP>.
4. Type exit.
5. Type interface Po10.
6. Type vpc peer-link.
7. Type exit.
8. Type interface Po11.
9. Type vpc 11.
10. Type exit.
11. Type interface Po12.
12. Type vpc 12.
13. Type exit
14. Type interface Po13.
15. Type vpc 13.
16. Type exit.
17. Type interface Po14.
18. Type vpc 14.
19. Type exit.
20. Type copy run start

Configure FCoE Fabric

These steps provide details for configuring Fiber Channel over Ethernet Fabric.

Nexus 5548 A

1. Type interface vfc11.
2. Type bind interface po11.
3. Type no shutdown.
4. Type exit.
5. Type interface vfc12.
6. Type bind interface po12.
7. Type no shutdown.
8. Type exit.
9. Type interface vfc15.
10. Type bind interface po15.
11. Type no shutdown.
12. Type exit.

13. Type vsan database.
14. Type vsan <VSAN A ID>
15. Type fcoe vsan <VSAN A ID>.
16. Type vsan <VSAN A ID> interface vfc11.
17. Type vsan <VSAN A ID> interface vfc12.
18. Type exit.
19. Type vlan <<Fabric_A_FCoE_VLAN ID>>
20. Type fcoe vsan <VSAN A ID>.
21. Type exit.
22. Type copy run start

Nexus 5548 B

1. Type interface vfc11.
2. Type bind interface po11.
3. Type no shutdown.
4. Type exit.
5. Type interface vfc12.
6. Type bind interface po12.
7. Type no shutdown
8. Type exit.
9. Type interface vfc16
10. Type bind interface po16
11. Type no shutdown
12. Type exit.
13. Type vsan database.
14. Type vsan <VSAN B ID>
15. Type vsan <VSAN B ID> name Fabric_B.
16. Type vsan <VSAN B ID> interface vfc11.
17. Type vsan <VSAN B ID> interface vfc12.
18. Type exit.
19. Type vlan <<Fabric_B_FCoE_VLAN ID>>
20. Type fcoe vsan <VSAN B ID>.
21. Type exit.
22. Type copy run start

Link into Existing Network Infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 5548 switches included in the FlexPod environment into the infrastructure. The previously described procedures can be used to create an uplink vPC to the existing environment.

NetApp FAS3240A Deployment Procedure - Part 1

Complete the Configuration Worksheet

Before running the setup script, complete the Configuration worksheet from the product manual.

	How to Access the Configuration Worksheet Configuration Guide	Comments
Configuration Worksheet	https://library.netapp.com/ecm/ecm_get_file/ECMM1249829	Requires access to the NetApp Support site.

Assign Controller Disk Ownership and initialize storage

These steps provide details for assigning disk ownership and disk initialization and verification.

Note Typical best practices should be followed when determining the number of disks to assign to each controller head. You may choose to assign a disproportionate number of disks to a given storage controller in an HA pair, depending on the intended workload.

In this reference architecture, half the total number of disks in the environment is assigned to one controller and the remainder to its partner.

Detail	Detail Value
Controller A MGMT IP	
Controller A netmask	
Controller A gateway	
URL of the Data ONTAP boot software	
Controller B MGMT IP	
Controller B netmask	
Controller B gateway	

Controller A

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

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

2. If the system is at the LOADER prompt, enter the following command to boot Data ONTAP:

```
autoboot
```

3. During system boot, press Ctrl – C when prompted for the Boot Menu:

```
Press Ctrl-C for Boot Menu…
```

Note If 8.1.2 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.2 is the version being booted, then proceed with step 14, maintenance mode boot.

4. To install new software first select option 7.

```
7
```

5. Type y indicating yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

7. Type y indicating yes to reboot now.

```
y
```

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

```
<<var_controller1_e0m_ip>>
<<var_controller1_mask>>>
<<var_controller1_mgmt_gateway>>.
```

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

Note This Web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

11. Type y indicating yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

12. Type y indicating yes to reboot the node.

```
y
```

13. When you see "`Press Ctrl-C for Boot Menu`", press:

```
Ctrl-C
```

14. To enter Maintenance mode boot, select option 5.

```
5
```

15. When you see the question "`Continue to Boot?`" type `yes`.

```
y
```

16. To verify the HA status of your environment, enter:

```
ha-config show
```

Note If either component is not in HA mode, use the **`ha-config modify`** command to put the components in HA mode.

17. To see how many disks are unowned, enter:

```
disk show -a
```

Note No disks should be owned in this list.

18. Assign disks.

```
disk assign –n <<var_#_of_disks>>
```

Note This reference architecture allocates half the disks to each controller. However, workload design could dictate different percentages.

19. Reboot the controller.

```
halt
```

20. At the LOADER-A prompt, enter:

```
autoboot
```

21. Press Ctrl – C for Boot Menu when prompted.

```
Ctrl-C
```

22. Select option 4 for Clean configuration and initialize all disks.

```
4
```

23. Enter y indicating yes to zero disks, reset config, and install a new file system.

```
y
```

24. Type y indicating yes to erase all the data on the disks.

```
y
```

Note The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots. You can continue with the Controller B configuration while the disks for Controller A are zeroing.

Controller B

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

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

2. If the system is at the LOADER prompt, enter the following command to boot Data ONTAP:

```
autoboot
```

3. During system boot, press Ctrl – C when prompted for the Boot Menu:

```
Press Ctrl-C for Boot Menu…
```

Note If 8.1.2 is not the version of software being booted, proceed with the steps below to install new software. If 8.1.2 is the version being booted, then proceed with step 14, maintenance mode boot.

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

```
7
```

5. Type y indicating yes to perform a nondisruptive upgrade.

```
y
```

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

```
e0M
```

7. Type y indicating yes to reboot now.

```
y
```

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

```
<<var_controller2_e0m_ip>>
<<var_controller2_mask>>
<<var_controller2_mgmt_gateway>>.
```

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

Note This Web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

11. Type y indicating yes to set the newly installed software as the default to be used for subsequent reboots.

```
y
```

12. Type y indicating yes to reboot the node.

```
y
```

13. When you see "`Press Ctrl-C for Boot Menu`", press:

```
Ctrl-C
```

14. To enter Maintenance mode boot, select option 5:

```
5
```

15. If you see the question "`Continue to Boot?`" type `yes`.

```
y
```

16. To verify the HA status of your environment, enter:

```
ha-config show
```

Note If either component is not in HA mode, use the **`ha-config modify`** command to put the components in HA mode.

17. To see how many disks are unowned, enter:

```
disk show -a
```

Note The remaining disks should be shown.

18. Assign disks by entering:

```
disk assign -n <<var_#_of_disks>>
```

Note This reference architecture allocates half the disks to each controller. However, workload design could dictate different percentages.

19. Reboot the controller.

```
halt
```

20. At the LOADER prompt, enter:

```
autoboot
```

21. Press Ctrl – C for Boot Menu when prompted.

```
Ctrl-C
```

22. Select option 4 for a Clean configuration and initialize all disks.

```
4
```

23. Type y indicating yes to zero disks, reset config, and install a new file system.

```
y
```

24. Type y indicating yes to erase all the data on the disks.

```
y
```

Note The initialization and creation of the root volume can take 75 minutes or more to complete, depending on the number of disks attached. When initialization is complete, the storage system reboots.

Run the Setup Process

When Data ONTAP is installed on your new storage system, the following files are not populated:

- **/etc/rc**
- **/etc/exports**
- **/etc/hosts**
- **/etc/hosts.equiv**

Controller A

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
2. Enter the following information:

```
Please enter the new hostname []:<<var_controller1>>
Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter
Please enter the IP address for Network Interface e0a []: Enter
```

Note Press Enter to accept the blank IP address.

```
Should interface e0a take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e0b []:Enter
Should interface e0b take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e1a []:Enter
Should interface e1a take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e1b []:Enter
Should interface e1b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for Network Interface e0M []:
<<var_controller1_e0m_ip>>
Please enter the netmaskfor the Network Interface e0M [255.255.255.0]:
<<var_controller1_mask>>

Should interface e0M take over a partner IP address during failover? [n]: y
```

```
Please enter the IPv4 address or interface name to be taken over by e0M []: e0M
Please enter flow control for e0M {none, receive, send, full} [full]: Enter
```

3. Enter the following information:

```
Please enter the name or IP address of the IPv4 default gateway:
<<var_controller1_mgmt_gateway>>

The administration host is given root access to the storage system's / etc files
for system administration. To allow /etc root access to all NFS clients enter
RETURN below.
Please enter the name or IP address for administrative host: <<var_adminhost_ip>>

Please enter timezone [GTM]: <<var_timezone>>
```

Note Example time zone: America/New_York.

```
Where is the filer located? <<var_location>>
Enter the root directory for HTTP files [home/http]: Enter
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:
```

Note Optionally enter up to three name server IP addresses.

```
Do you want to run NIS client? [n]: Enter
Press the Return key to continue through AutoSupport message
would you like to configure SP LAN interface [y]: Enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_sp_ip>>
Please enter the netmask for the SP []: <<var_sp_mask>>
Please enter the IP address for the SP gateway: <<var_sp_gateway>>
Please enter the name or IP address of the mail host [mailhost]: <<var_mailhost>>
Please enter the IP address for <<var_mailhost>> []: <<var_mailhost_ip>>
New password: <<var_ password>>
Retype new password <<var_ password>>
```

4. Enter the admin password to log in to Controller A.

Controller B

1. Enter the configuration values the first time you power on the new system. The configuration values populate these files and configure the installed functionality of the system.
2. Enter the following information:

```
Please enter the new hostname []: <<var_controller2>>
Do you want to enable IPv6? [n]: Enter

Do you want to configure interface groups? [n]: Enter
Please enter the IP address for Network Interface e0a []: Enter
```

Note Press Enter to accept the blank IP address.

```
Should interface e0a take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e0b []:Enter
Should interface e0b take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e1a []:Enter
```

```
Should interface e1a take over a partner IP address during failover? [n]: Enter
Please enter the IP address for the Network Interface e1b []:Enter
Should interface e1b take over a partner IP address during failover? [n]: Enter

Please enter the IP address for Network Interface e0M []:
<<var_controller2_e0m_ip>>
Please enter the netmaskfor the Network Interface e0M [255.255.255.0]:
<<var_controller2_mask>>

Should interface e0M take over a partner IP address during failover? [n]: y
Please enter the IPv4 address or interface name to be taken over by e0M []: e0M
Please enter flow control for e0M {none, receive, send, full} [full]: Enter
```

3. Enter the following information:

```
Please enter the name or IP address of the IPv4 default gateway:
<<var_controller2_mgmt_gateway>>

The administration host is given root access to the storage system's / etc files
for system administration. To allow /etc root access to all NFS clients enter
RETURN below.
Please enter the name or IP address for administrative host: <<var_adminhost_ip>>

Please enter timezone [GTM]: <<var_timezone>>
```

Note Example time zone: America/New York.

```
Where is the filer located? <<var_location>>
Enter the root directory for HTTP files [home/http]: Enter
Do you want to run DNS resolver? [n]: y
Please enter DNS domain name []: <<var_dns_domain_name>>
Please enter the IP address for first nameserver []: <<var_nameserver_ip>>
Do you want another nameserver? [n]:
```

Note Optionally enter up to three name server IP addresses.

```
Do you want to run NIS client? [n]: Enter
Press the Return key to continue through AutoSupport message
would you like to configure SP LAN interface [y]: Enter
Would you like to enable DHCP on the SP LAN interface [y]: n
Please enter the IP address for the SP: <<var_sp_ip>>
Please enter the netmask for the SP []: <<var_sp_mask>>
Please enter the IP address for the SP gateway: <<var_sp_gateway>>
Please enter the name or IP address of the mail host [mailhost]: <<var_mailhost>>
Please enter the IP address for <<var_mailhost>> []: <<var_mailhost_ip>>
New password: <<var_admin_passwd>>
Retype new password <<var_admin_passwd>>
```

4. Enter the admin password to log in to Controller B.

Upgrade the Service Processor on Each Node to the Latest Release

With Data ONTAP 8.1.2, you must upgrade to the latest Service Processor (SP) firmware to take advantage of the latest updates available for the remote management device.

1. Using a web browser, go to http://support.netapp.com/NOW/cgi-bin/fw.
2. Navigate to the Service Process Image for installation from the Data ONTAP prompt page for your storage platform.
3. Proceed to the Download page for the latest release of the SP Firmware for your storage platform.

4. Follow the instructions on this page; update the SPs on both controllers. You will need to download the .zip file to a Web server that is reachable from the management interfaces of the controllers.

64-Bit Aggregates

Note A 64-bit aggregate containing the root volume is created during the Data ONTAP setup process. To create additional 64-bit aggregates, determine the aggregate name, the node on which it can be created, and how many disks it should contain. Calculate the RAID group size to allow for roughly balanced (same size) RAID groups from 12 through 20 disks (for SAS disks) within the aggregate. For example, if 52 disks are assigned to the aggregate then, select a RAID group size of 18. A RAID group size of 18 would yield two 18-disk RAID groups and one 16-disk RAID group. Remember that the default RAID group size is 16 disks, and that the larger the RAID group size, the longer the disk rebuild time in case of a failure.

Controller A

1. Execute the following command to create a new aggregate:

```
aggr create aggr1 -B 64 -r <<var_raidsize>> <<var_num_disks>>
```

Note Leave 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.

Controller B

1. Execute the following command to create a new aggregate:

```
aggr create aggr1 -B 64 -r <<var_raidsize>> <<var_num_disks>>
```

Note Leave 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.

Flash Cache

Controller A and Controller B

1. Execute the following commands to enable Flash Cache:

```
options flexscale.enable on
options flexscale.lopri_blocks off
options flexscale.normal_data_blocks on
```

Note For directions on how to configure Flash Cache in metadata mode or low-priority data caching mode, refer to TR-3832: Flash Cache and PAM Best Practices Guide. Before customizing the settings, determine whether the custom settings are required or whether the default settings are sufficient.

IFGRP LACP

Since this type of interface group requires two or more Ethernet interfaces and a switch that supports Link Aggregation Control Protocol (LACP), make sure that the switch is configured properly.

Controller A and Controller B

1. Run the following command on the command line and also add it to the **/etc/rc** file, so it is activated upon boot:

```
ifgrp create lacp ifgrp0 -b port e2a e2b
wrfile -a /etc/rc "ifgrp create lacp ifgrp0 -b ip e1a e1b"
```

Note All interfaces must be in down status before being added to an interface group.

VLAN

Controller A and Controller B

1. Follow the steps below to create a VLAN interface for iSCSI data traffic.

```
vlan create ifgrp0 <<var_iscsi_a_vlan_id>>, <<var_iscsi_b_vlan_id>>
wrfile -a /etc/rc "vlan create ifgrp0 <<var_iscsi_a_vlan_id>>,
<<var_iscsi_b_vlan_id>>"
```

IP Config

Controller A and Controller B

1. Run the following commands on the command line.

```
ifconfig ifgrp0-<<var_iscsi_a_vlan_id>> <<var_iscsi_a_ip>> netmask
<<var_iscsi_a_mask>> mtusize 9000 partner ifgrp0-<<var_iscsi_a_vlan_id>>
ifconfig ifgrp0-<<var_iscsi_b_vlan_id>> <<var_iscsi_b_ip>> netmask
<<var_iscsi_b_mask>> mtusize 9000 partner ifgrp0-<<var_iscsi_b_vlan_id>>
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_iscsi_a_vlan_id>> <<var_iscsi_a_ip>>
netmask <<var_iscsi_a_mask>> mtusize 9000 partner ifgrp0-<<var_iscsi_a_vlan_id>>"
wrfile -a /etc/rc "ifconfig ifgrp0-<<var_iscsi_b_vlan_id>> <<var_iscsi_b_ip>>
netmask <<var_iscsi_b_mask>> mtusize 9000 partner ifgrp0-<<var_iscsi_b_vlan_id>>"
```

Storage Controller Active-Active Configuration

Controller A and Controller B

To enable two storage controllers to an active-active configuration, complete the following steps:

1. Enter the cluster license on both nodes.

```
license add <<var_cf_license>>
```

2. Reboot both the storage controllers.

```
reboot
```

3. Log back in to both the controllers.

Controller A

1. Enable failover on Controller A, if it is not enabled already.

```
cf enable
```

NTP

The following commands configure and enable time synchronization on the storage controller. You must have either a publicly available IP address or your company's standard NTP server name or IP address.

Controller A and Controller B

1. Run the following commands to configure and enable the NTP server:

```
date <<var_date>>
```

2. Enter the current date in the format of [[[[CC]yy]mm]dd]hhmm[.ss]].

   For example, date 201208311436; which means the date is set to August 31st 2012 at 14:36.

```
options timed.servers <<var_global_ntp_server_ip>>
options timed.enable on
```

Joining a Windows Domain (optional)

The following commands should be used to allow the NetApp controllers to join the existing Domain.

Controller A and Controller B

Add the controller to the domain by running CIFS setup.

```
CIFS setup
Do you want to make the system visible via WINS? [N] n
Choose (2) NTFS-only filer [2] 2
Enter the password for the root user []: <<var_root_password>>
Reenter the password: <<var_root_password>>
Would you like the change this name? [n]: Enter
Choose (1) Active Directory Domain Authentication : 1 Enter
Configure the DNS Resolver Service ?[y]: y
What is the filers DSN Domain name? []: <<var_dnsdomain>>
What the IPv4 Addresses of your Authoritative DNS servers? []:
<<var_ip_DNSserver>>
What is the name of the Active Directory Domain Controller? : <<var_dnsdomain>>
Would you like to configure time services? [y]: [Enter]
Enter the time server host []:<<var_dnsdomain>>
Enter the ame of the windows user [administrator@<<var_fas3240_dnsdomain>>]
[Enter]
Password for <<var_domainAccountUsed>>: <<password>>
Choose (1) Create the filers machine account in the "computers" container: 1
Do you want to configure a <<var_ntap_hostname>>/administrator account [Y]:
[Enter]
Password for the <<var_ntap_hostname>>/Administrator <<var_password>> [Enter]
Would you like to specify a user or group that can administer CIFS [n]: [Enter]
```

iSCSI

Controller A and Controller B

1. Add a license for iSCSI.

```
license add <<var_nfs_license>>
```

2. Start iSCSI

```
iscsi start
```

FCP

Controller A and Controller B

1. License FCP.

```
license add <<var_fc_license>>
```

2. Start the FCP service.

```
fcp start
```

3. Record the WWPN or FC port name for later use.

```
fcp show adapters
```

4. If using FC instead of FCoE between storage and the network, if necessary execute the following commands to make ports 0c and 0d target ports.

```
fcadmin config
```

5. Make an FC port into a target.

Note Only FC ports that are configured as targets can be used to connect to initiator hosts on the SAN.

For example, make a port called **<<var_fctarget01>>** into a target port by running the following command:

```
fcadmin config -t target <<var_fctarget01>>
```

Note If an initiator port is made into a target port, a reboot is required. NetApp recommends rebooting after completing the entire configuration because other configuration steps might also require a reboot.

Data ONTAP SecureAdmin

Secure API access to the storage controller must be configured.

Controller A

1. Execute the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup ssl
SSL Setup has already been done before. Do you want to proceed? [no] y
Country Name (2 letter code) [US]: <<var_country_code>>
State or Province Name (full name) [California]: <<var_state>>
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
Organization Name (company) [Your Company]: <<var_org>>
Organization Unit Name (division): <<var_unit>>
Common Name (fully qualified domain name) [<<var_controller1_fqdn>>]: Enter
Administrator email: <<var_admin_email>>
Days until expires [5475] : Enter
Key length (bits) [512] : <<var_key_length>>
```

Note NetApp recommends your key length to be 1024.

After the initialization, the CSR is available in the file:

```
/etc/keymgr/csr/secureadmin_tmp.pem.
```

2. Configure and enable SSL and HTTPS for API access using the following options.

```
options httpd.access none
options httpd.admin.enable off
options httpd.admin.ssl.enable on
options ssl.enable on
```

Controller B

1. Execute the following as a one-time command to generate the certificates used by the Web services for the API.

```
secureadmin setup ssl
SSL Setup has already been done before. Do you want to proceed? [no] y
Country Name (2 letter code) [US]: <<var_country_code>>
State or Province Name (full name) [California]: <<var_state>>
Locality Name (city, town, etc.) [Santa Clara]: <<var_city>>
Organization Name (company) [Your Company]: <<var_org>>
Organization Unit Name (division): <<var_unit>>
Common Name (fully qualified domain name) [<<var_controller2_fqdn>>]: Enter
Administrator email: <<var_admin_email>>
Days until expires [5475] : Enter
Key length (bits) [512] : <<var_key_length>>
```

Note NetApp recommends your key length to be 1024.

After the initialization, the CSR is available in the file

```
/etc/keymgr/csr/secureadmin_tmp.pem.
```

2. Configure and enable SSL and HTTPS for API access using the following options.

```
options httpd.access none
options httpd.admin.enable off
options httpd.admin.ssl.enable on
options ssl.enable on
```

Secure Shell

SSH must be configured and enabled.

Controller A and Controller B

1. Execute the following one-time command to generate host keys.

```
secureadmin disable ssh
secureadmin setup -f -q ssh 768 512 1024
```

2. Use the following options to configure and enable SSH.

```
options ssh.idle.timeout 60
options autologout.telnet.timeout 5
```

SNMP

Controller A and Controller B

1. Run the following commands to configure SNMP basics, such as the local and contact information. When polled, this information displays as the sysLocation and sysContact variables in SNMP.

```
snmp contact "<<var_admin_email>>"
```

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

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

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

SNMPv1

Controller A and Controller B

1. Set the shared secret plain-text password, which is called a community.

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

Note Use the **delete all** command with caution. If community strings are used for other monitoring products, the **delete all** command will remove them.

SNMPv3

SNMPv3 requires a user to be defined and configured for authentication.

Controller A and Controller B

1. Create a user called **snmpv3user**.

```
useradmin role add snmp_requests -a login-snmp
useradmin group add snmp_managers -r snmp_requests
useradmin user add snmpv3user -g snmp_managers
New Password: <<var_ password>>
Retype new password: <<var_ password>>
```

AutoSupport HTTPS

AutoSupport™ sends support summary information to NetApp through HTTPS.

Controller A and Controller B

1. Execute the following command to configure AutoSupport.

```
options autosupport.noteto <<var_admin_email>>
```

Security Best Practices

Note Apply the following commands according to local security policies.

Controller A and Controller B

1. Run the following commands to enhance security on the storage controller:

```
options rsh.access none
```

```
options webdav.enable off
options security.passwd.rules.maximum 14
options security.passwd.rules.minimum.symbol 1
options security.passwd.lockout.numtries 6
options autologout.console.timeout 5
```

Install Remaining Required Licenses and Enable MultiStore

Controller A and Controller B

1. Install the following licenses to enable SnapRestore® and FlexClone®.

```
license add <<var_snaprestore_license>>
license add <<var_flex_clone_license>>
options licensed_feature.multistore.enable on
```

Enable NDMP

Run the following commands to enable NDMP.

Controller A and Controller B

```
options ndmpd.enable on
```

Add Infrastructure Volumes

Controller A

1. Create a FlexVol® volume in Aggr1 to host the UCS boot LUNs, and cluster quorum.

```
vol create ucs_boot -s none aggr1 500g
vol create hyperv_quorum -s none aggr1 10g
```

2. Configure volume dedupe.

```
sis config -s auto /vol/ucs_boot
sis config -s auto /vol/hyperv_quorum
sis on /vol/ucs_boot
sis on /vol/hyperv_quorum
sis start -s /vol/ucs_boot
sis start -s /vol/hyperv_quorum
```

Controller B

1. Create a 500GB FlexVol volume in Aggr1 to host the management infrastructure virtual machines.

```
vol create ucs_boot -s none aggr1 500g
vol create fabric_mgmt_csv -s none aggr1 5t
```

2. Configure dedupe.

```
sis config -s auto /vol/fabric_mgmt_csv
sis config -s auto /vol/ucs_boot
sis on /vol/ucs_boot
sis on /vol/fabric_mgmt_csv
sis start -s /vol/ucs_boot
sis start -s /vol/fabric_mgmt_csv
```

Install SnapManager licenses

Controller A and Controller B

1. Add a license for SnapManager® for Hyper-V.

```
license add <<var_snapmanager_hyperv_license>>
```

2. Add a license for SnapDrive® for Windows.

```
license add <<var_snapdrive_windows_license>>
```

Cisco Unified Computing System Deployment Procedure

The following section provides a detailed procedure for configuring the Cisco Unified Computing System for use in a FlexPod environment. These steps should be followed precisely because a failure to do so could result in an improper configuration.

Perform Initial Setup of the Cisco UCS 6248 Fabric Interconnects

These steps provide details for initial setup of the Cisco UCS 6248 Fabric Interconnects.

Cisco UCS 6248 A

1. Connect to the console port on the first Cisco UCS 6248 Fabric Interconnect.
2. At the prompt to enter the configuration method, enter `console` to continue.
3. If asked to either do a new setup or restore from backup, enter `setup` to continue.
4. Enter `y` to continue to set up a new fabric interconnect.
5. Enter `y` to enforce strong passwords.
6. Enter the password for the admin user.
7. Enter the same password again to confirm the password for the admin user.
8. When asked if this fabric interconnect is part of a cluster, answer `y` to continue.
9. Enter `A` for the switch fabric.
10. Enter the cluster name for the system name.
11. Enter the Mgmt0 IPv4 address.
12. Enter the Mgmt0 IPv4 netmask.
13. Enter the IPv4 address of the default gateway.
14. Enter the cluster IPv4 address.
15. To configure DNS, answer `y`.
16. Enter the DNS IPv4 address.
17. Answer `y` to set up the default domain name.
18. Enter the default domain name.
19. Review the settings that were printed to the console, and if they are correct, answer `yes` to save the configuration.
20. Wait for the login prompt to make sure the configuration has been saved.

Cisco UCS 6248 B

1. Connect to the console port on the second Cisco UCS 6248 Fabric Interconnect.
2. When prompted to enter the configuration method, enter console to continue.
3. The installer detects the presence of the partner fabric interconnect and adds this fabric interconnect to the cluster. Enter `y` to continue the installation.
4. Enter the admin password for the first fabric interconnect.
5. Enter the Mgmt0 IPv4 address.
6. Answer yes to save the configuration.
7. Wait for the login prompt to confirm that the configuration has been saved.

Log into Cisco UCS Manager

These steps provide details for logging into the Cisco UCS environment.

1. Open a Web browser and navigate to the Cisco UCS 6248 Fabric Interconnect cluster address.
2. Select the Launch link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` for the username and enter the administrative password and click Login to log in to the Cisco UCS Manager software.

Add a Block of IP Addresses for KVM Access

These steps provide details for creating a block of KVM ip addresses for server access in the Cisco UCS environment.

1. Select the Admin tab at the top of the left window.
2. Select **All** > **Communication Management**.
3. Right-click Management IP Pool.
4. Select Create Block of IP Addresses.
5. Enter the starting IP address of the block and number of IPs needed as well as the subnet and gateway information.
6. Click **OK** to create the IP block.
7. Click **OK** in the message box.

Synchronize Cisco UCS to NTP

These steps provide details for synchronizing the Cisco UCS environment to the NTP server.

1. Select the Admin tab at the top of the left window.
2. Select **All** > **Timezone Management**.
3. Right-click Timezone Management.
4. In the right pane, select the appropriate timezone in the Timezone drop-down menu.
5. Click **Save Changes** and then **OK**.
6. Click **Add NTP Server**.

7. Input the NTP server IP and click **OK**.

Configure Unified Ports

These steps provide details for modifying an unconfigured Ethernet port into a FC uplink port ports in the Cisco UCS environment.

Note Modification of the unified ports leads to a reboot of the fabric interconnect in question. This reboot can take up to 10 minutes.

1. Navigate to the Equipment tab in the left pane.
2. Select Fabric Interconnect A.
3. In the right pane, select the General tab.
4. Select Configure Unified Ports.
5. Click **Yes** to launch the wizard.
6. Use the slider tool and move one position to the left to configure the last two ports (31 and 32) as FC uplink ports.
7. Ports 31 and 32 now have the “B” indicator indicating their reconfiguration as FC uplink ports.
8. Click **Finish**.
9. Click **OK**.
10. The Cisco UCS Manger GUI will close as the primary fabric interconnect reboots.
11. Upon successful reboot, open a Web browser and navigate to the Cisco UCS 6248 Fabric Interconnect cluster address.
12. When prompted, enter admin for the username and enter the administrative password and click **Login** to log in to the Cisco UCS Manager software.
13. Navigate to the Equipment tab in the left pane.
14. Select Fabric Interconnect B.
15. In the right pane, click the General tab.
16. Select Configure Unified Ports.
17. Click **Yes** to launch the wizard.
18. Use the slider tool and move one position to the left to configure the last two ports (31 and 32) as FC uplink ports.
19. Ports 31 and 32 now have the “B” indicator indicating their reconfiguration as FC uplink ports.
20. Click **Finish**.
21. Click **OK**.

Chassis Discovery Policy

These steps provide details for modifying the chassis discovery policy as the base architecture includes two uplinks from each fabric extender installed in the Cisco UCS chassis.

1. Navigate to the Equipment tab in the left pane.

2. In the right pane, click the Policies tab.
3. Under Global Policies, change the Chassis Discovery Policy to 4-link or set it to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Keep Link Grouping Preference set to None
5. Click **Save Changes**.

Enable Server and Uplink Ports

These steps provide details for enabling Fibre Channel, server and uplinks ports.

1. Select the Equipment tab on the top left of the window.
2. Select **Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module**.
3. Expand the Ethernet Ports object.
4. Select the ports that are connected to the chassis or to the Cisco 2232 FEX (four per FEX), right-click them, and select Configure as Server Port.
5. Click **Yes** to confirm the server ports, and then click **OK**.
6. The ports connected to the chassis or to the Cisco 2232 FEX are now configured as server ports.

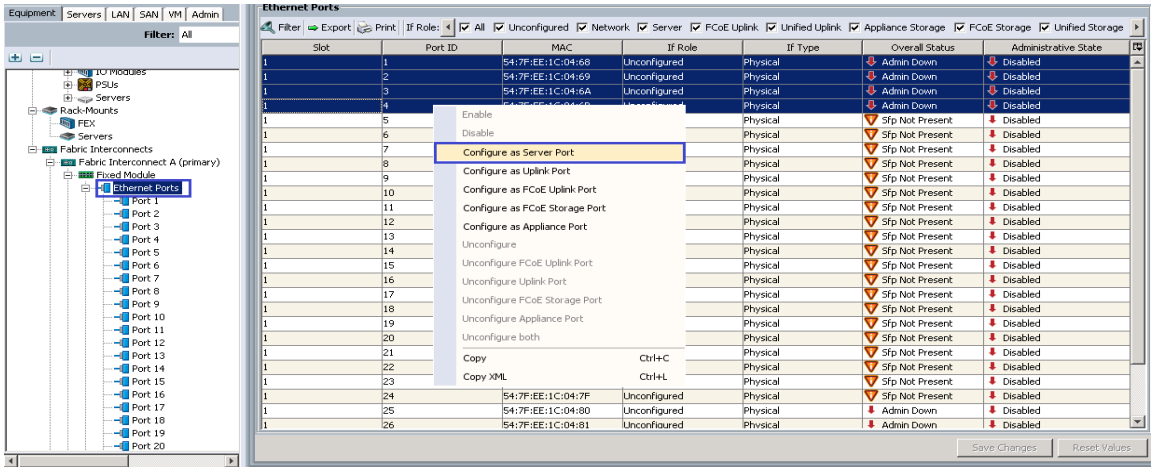

7. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
8. Select ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and select Configure as Uplink Port.
9. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
10. Select **Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module**.
11. Expand the Ethernet Ports object.
12. Select ports the number of ports that are connected to the Cisco UCS chassis (4 per chassis), right-click them, and select Configure as Server Port.
13. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.

14. Select ports 19 and 20 that are connected to the Cisco Nexus 5548 switches, right-click them, and select Configure as Uplink Port.
15. A prompt displays asking if this is what you want to do. Click **Yes**, then **OK** to continue.
16. At the prompt, click **Yes** to confirm the uplink ports, and then click **OK**.
17. If using the 2208 or 2204 FEX or the external 2232 FEX, navigate to each device by selecting **Equipment** > **Chassis** > **Rack-Mounts** > **FEX** > **<FEX #>**.
18. Select the Connectivity Policy tab in the right pane and change the administrative state of each fabric to Port Channel.
19. Click **Save Changes**, click **Yes**, and then click **OK**.

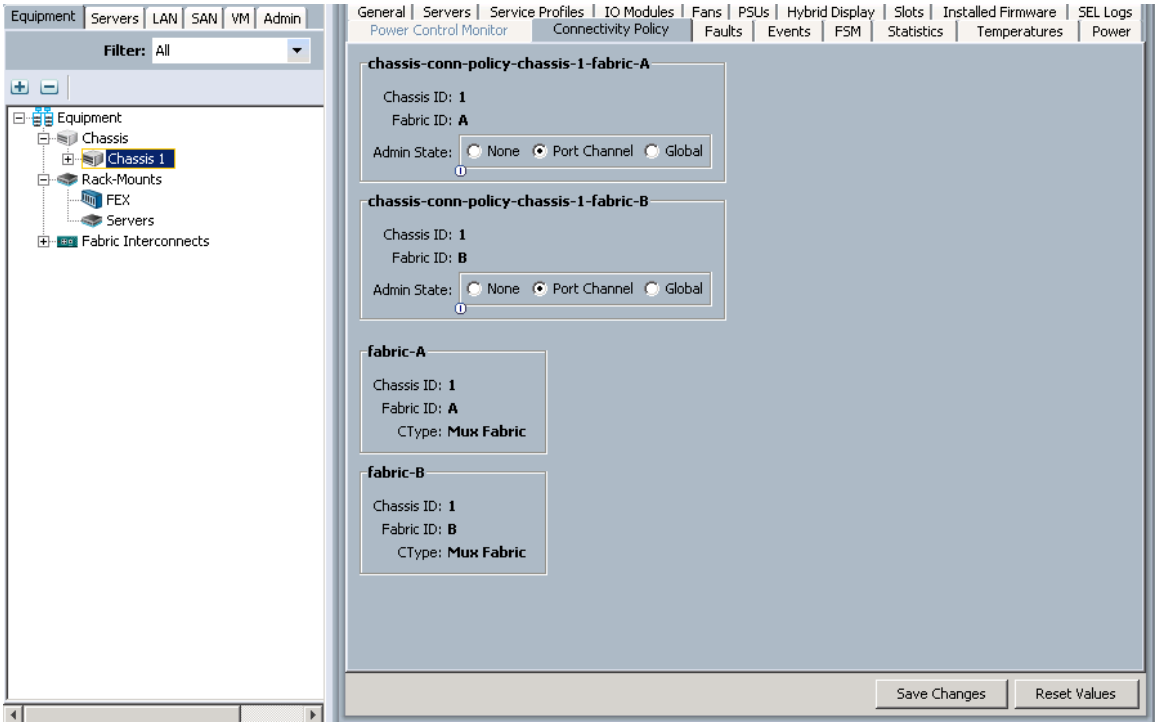

Acknowledge the Cisco UCS Chassis

The connected chassis needs to be acknowledged before it can be managed by Cisco UCS Manager.

1. Select Chassis 1 in the left pane.
2. Click **Acknowledge Chassis**.

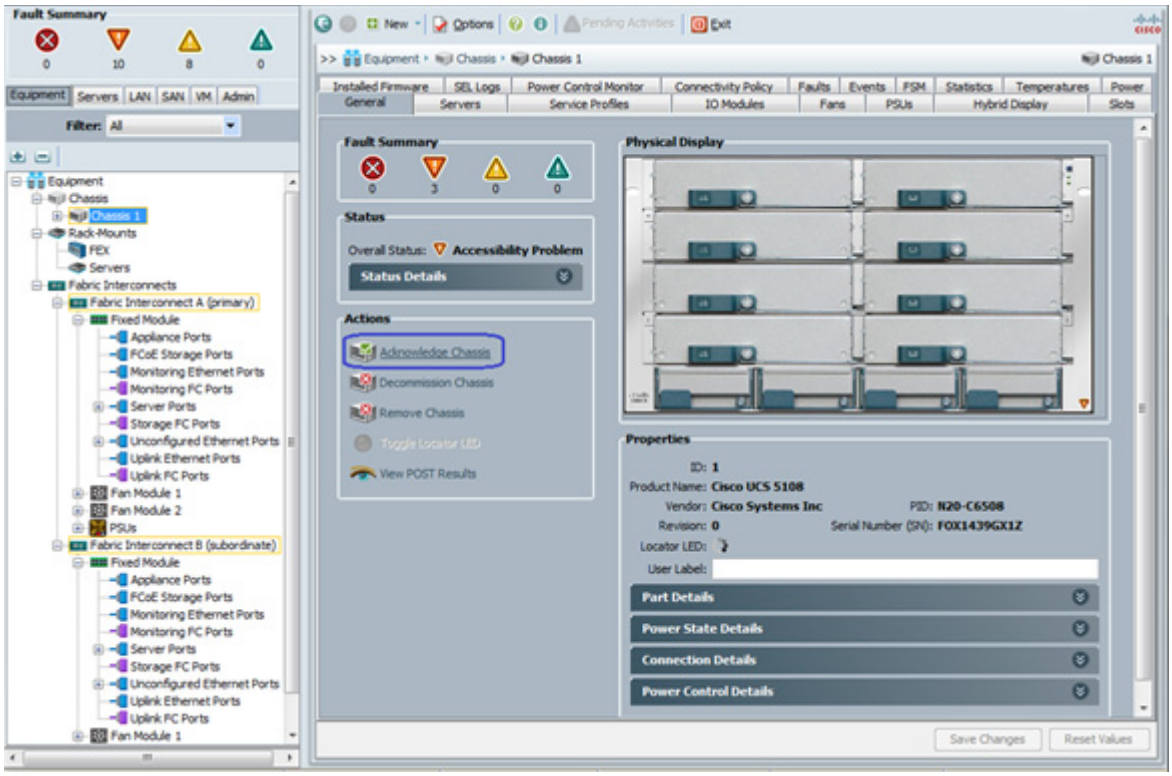

Create Uplink Port Channels to the Cisco Nexus 5548 Switches

These steps provide details for configuring the necessary Port Channels out of the Cisco UCS environment.

1. Select the LAN tab on the left of the window.

Note Two Port Channels are created, one from fabric A to both Cisco Nexus 5548 switches and one from fabric B to both Cisco Nexus 5548 switches.

2. Under LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the Port Channel.
6. Enter vPC-13-N5548 as the name of the Port Channel.
7. Click **Next**.

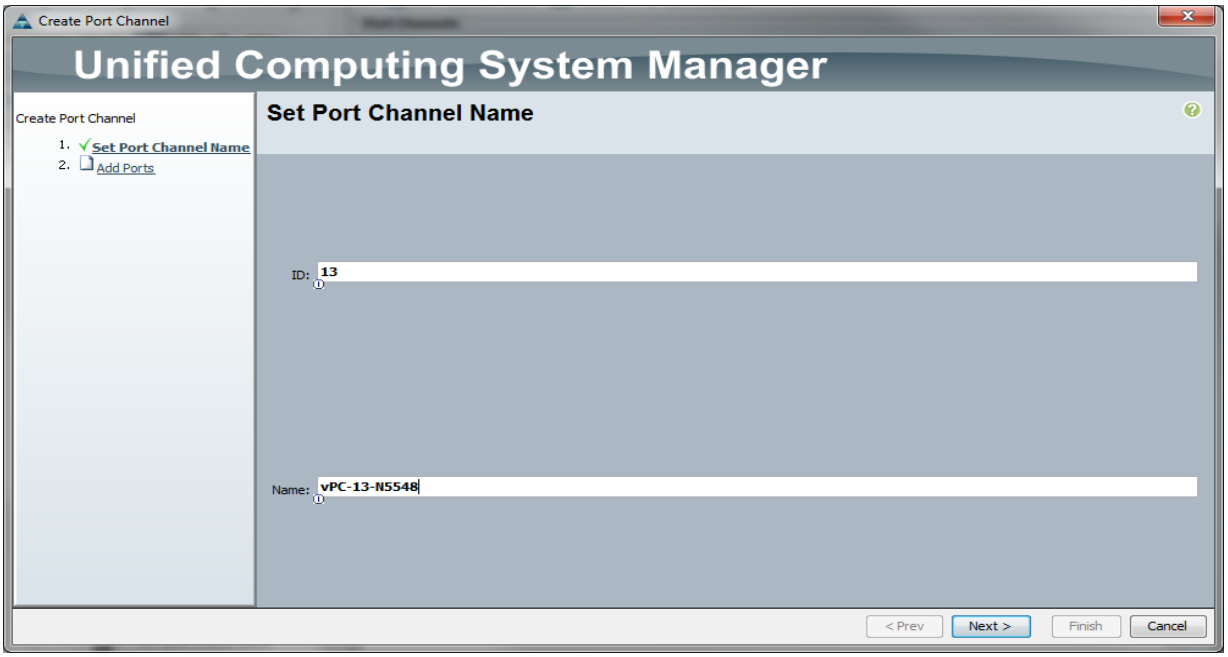

8. Select the port with slot ID: 1 and port: 19 and also the port with slot ID: 1 and port 20 to be added to the Port Channel.
9. Click **>>** to add the ports to the Port Channel.
10. Click **Finish** to create the Port Channel.
11. Check the Show navigator for Port-Channel 13 (Fabric A) checkbox.
12. Click **OK** to continue.
13. Under Actions, click **Enable Port Channel**.
14. In the pop-up box, click **Yes**, then **OK** to enable.

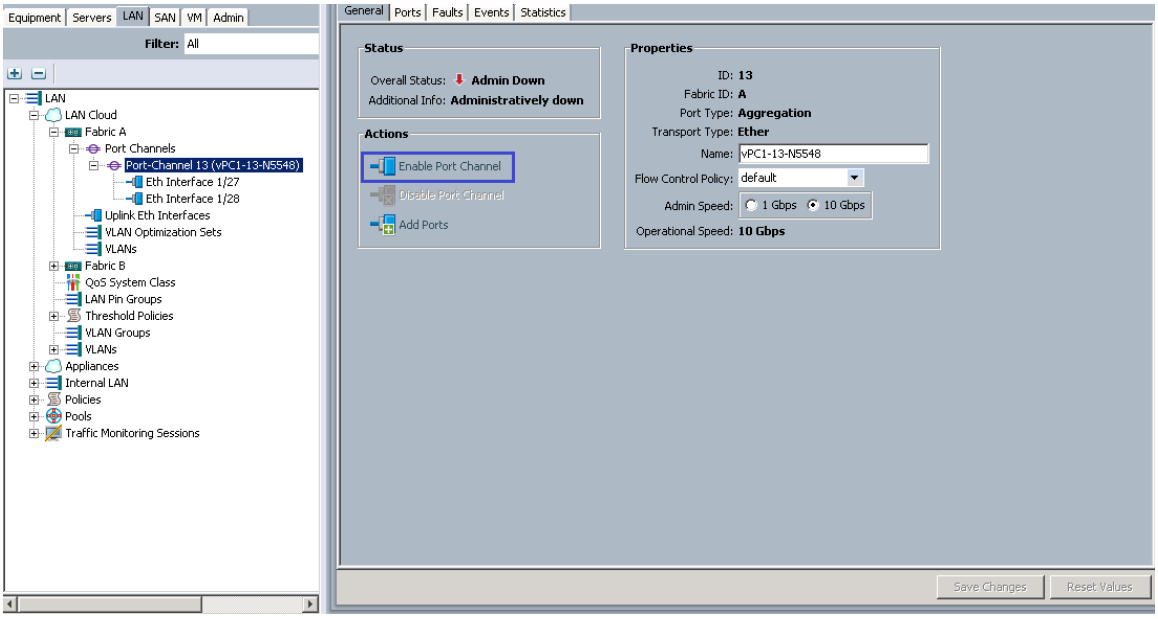

15. Wait until the overall status of the Port Channel is up.
16. Click **OK** to close the Navigator.
17. Under LAN Cloud, expand the Fabric B tree.
18. Right-click Port Channels.
19. Select Create Port Channel.
20. Enter 14 as the unique ID of the Port Channel.
21. Enter vPC-14-N5548 as the name of the Port Channel.
22. Click **Next**.
23. Select the port with slot ID: 1 and port: 19 and also the port with slot ID: 1 and port 20 to be added to the Port Channel.
24. Click **>>** to add the ports to the Port Channel.
25. Click **Finish** to create the Port Channel.
26. Check the Show navigator for Port-Channel 14 (Fabric B) checkbox.
27. Click **OK** to continue.
28. Under Actions, select Enable Port Channel.
29. In the pop-up box, click **Yes**, then **OK** to enable.
30. Wait until the overall status of the Port Channel is up
31. Click **OK** to close the Navigator.

Create an Organization

These steps provide details for configuring an organization in the Cisco UCS environment. Organizations are used as a means to organize and restrict access to various groups within the IT organization, thereby enabling multi-tenancy of the compute resources. This document does not assume the use of Organizations, however the necessary steps are included below.

1. From the New… menu at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Enter a description for the organization (optional).
4. Click **OK**.
5. In the message box that displays, click **OK**.

Create a MAC Address Pool

These steps provide details for configuring the necessary MAC address pool for the Cisco UCS environment.

1. Select the LAN tab on the left of the window. Select **Pools** > **Sub Organizations**.

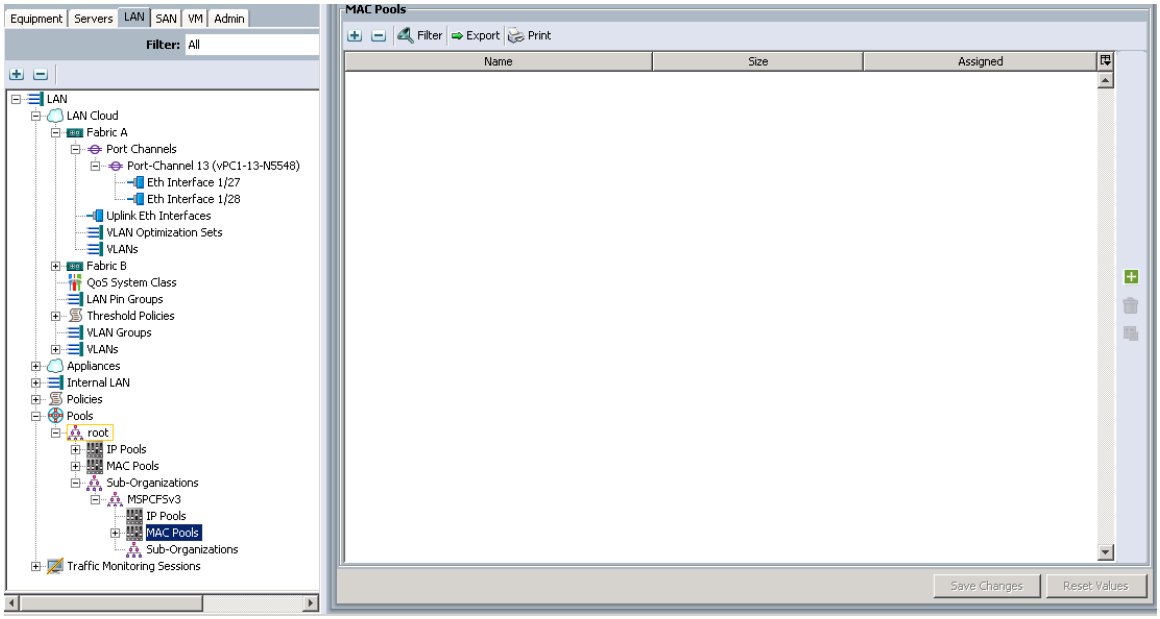

2. Right-click MAC Pools under the organization previously created.
3. Select Create MAC Pool to create the MAC address pool.
4. Enter MAC_Pool for the name of the MAC pool.
5. (Optional) Enter a description of the MAC pool.
6. Select Default assignment order.tesr
7. Click **Next**.
8. Click **Add**.
9. Specify a starting MAC address.
10. Specify a size of the MAC address pool sufficient to support the available blade resources.

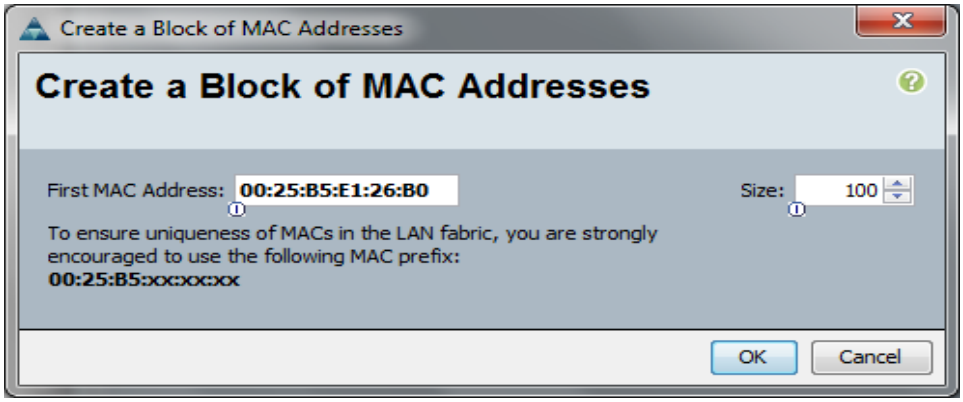

11. Click **OK**.
12. Click **Finish**.
13. In the message box that displays, click **OK**.

Create WWNN Pools

These steps provide details for configuring the necessary WWNN pools for the Cisco UCS environment.

1. Select the SAN tab at the top left of the window.
2. Select **Pools** > **root**.
3. Right-click WWNN Pools
4. Select Create WWNN Pool.
5. Enter WWNN_Pool as the name of the WWNN pool.
6. (Optional) Add a description for the WWNN pool.
7. Click **Next** to continue.
8. Click **Add** to add a block of WWNNs.

Note The default is appropriate for most configurations, modify if necessary.

9. Specify a size of the WWNN block sufficient to support the available blade resources.

10. Click **OK** to proceed.
11. Click **Finish** to proceed.
12. Click **OK** to finish.

Create WWPN Pools

These steps provide details for configuring the necessary WWPN pools for the Cisco UCS environment.

1. Select the SAN tab at the top left of the window. Select **Pools** > **root**.
2. Two WWPN pools are created, one for fabric A and one for fabric B.
3. Right-click WWPN Pools.

4. Select Create WWPN Pool.
5. Enter WWPN_Pool_A as the name for the WWPN pool for fabric A.
6. (Optional). Give the WWPN pool a description.
7. Click **Next**.
8. Click **Add** to add a block of WWPNs.
9. Enter the starting WWPN in the block for fabric A.
10. Specify a size of the WWPN block sufficient to support the available blade resources.

11. Click **OK**.
12. Click **Finish** to create the WWPN pool.
13. Click **OK**.
14. Right-click WWPN Pools
15. Select Create WWPN Pool.
16. Enter WWPN_Pool_B as the name for the WWPN pool for fabric B.
17. (Optional) Give the WWPN pool a description.
18. Click **Next**.
19. Click **Add** to add a block of WWPNs.
20. Enter the starting WWPN in the block for fabric B.
21. Specify a size of the WWPN block sufficient to support the available blade resources.
22. Click **OK**.
23. Click **Finish**.
24. Click **OK** to finish.

Create UUID Suffix Pools

These steps provide details for configuring the necessary UUID suffix pools for the Cisco UCS environment.

1. Select the Servers tab on the top left of the window. Select **Pools** > **root**.
2. Right-click UUID Suffix Pools.

3. Select Create UUID Suffix Pool.
4. Name the UUID suffix pool UUID_Pool.
5. (Optional) Give the UUID suffix pool a description.
6. Leave the prefix at the derived option.
7. Click **Next** to continue.
8. Click **Add** to add a block of UUIDs
9. The From field is fine at the default setting.
10. Specify a size of the UUID block sufficient to support the available blade resources.

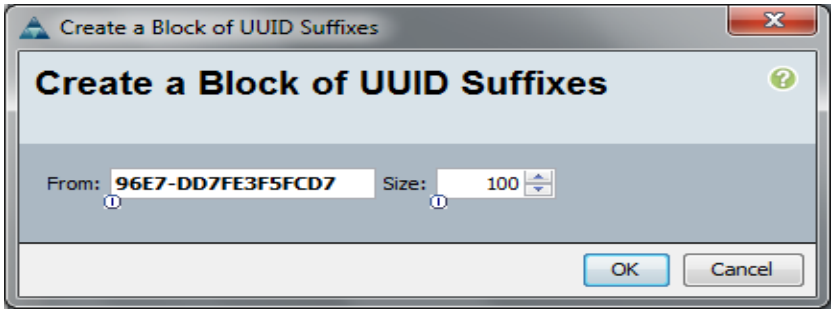

11. Click **OK**.
12. Click **Finish** to proceed.
13. Click **OK** to finish.

Create Server Pools

These steps provide details for configuring the necessary UUID suffix pools for the Cisco UCS environment.

1. Select the Servers tab at the top left of the window. Select **Pools** > **root**.
2. Right-click Server Pools.
3. Select Create Server Pool.
4. Name the server pool Infra_Pool.
5. (Optional) Give the server pool a description.
6. Click **Next** to continue to add servers.
7. Select two server to be used for the infrastructure cluster and Click **>>** to add them to the pool.
8. Click **Finish**.
9. Select **OK** to finish.

Create VLANs

These steps provide details for configuring the necessary VLANs for the Cisco UCS environment.

1. Select the LAN tab on the left of the window.

Note Eight VLANs are created.

2. Select LAN Cloud.
3. Right-click VLANs.
4. Select Create VLANs.
5. Enter MGMT-VLAN as the name of the VLAN to be used for management traffic.
6. Keep the Common/Global option selected for the scope of the VLAN.
7. Enter the VLAN ID for the management VLAN. Keep the sharing type as none.
8. Click **OK**.

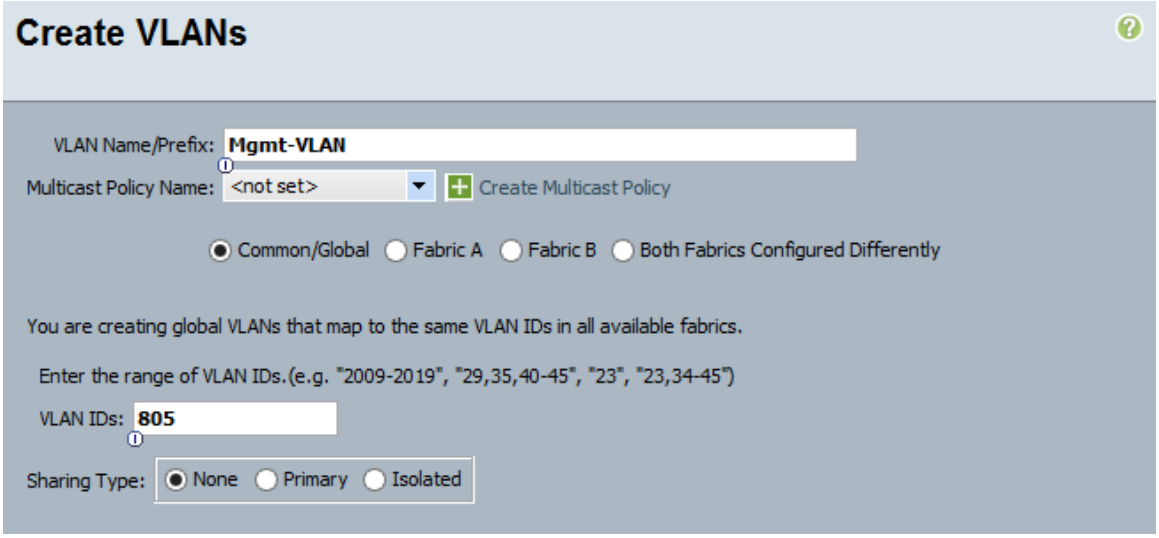

9. Right-click VLANs.
10. Select Create VLANs.
11. Enter CSV-VLAN as the name of the VLAN to be used for the CSV VLAN.
12. Keep the Common/Global option selected for the scope of the VLAN.
13. Enter the VLAN ID for the CSV VLAN.
14. Click **OK**.

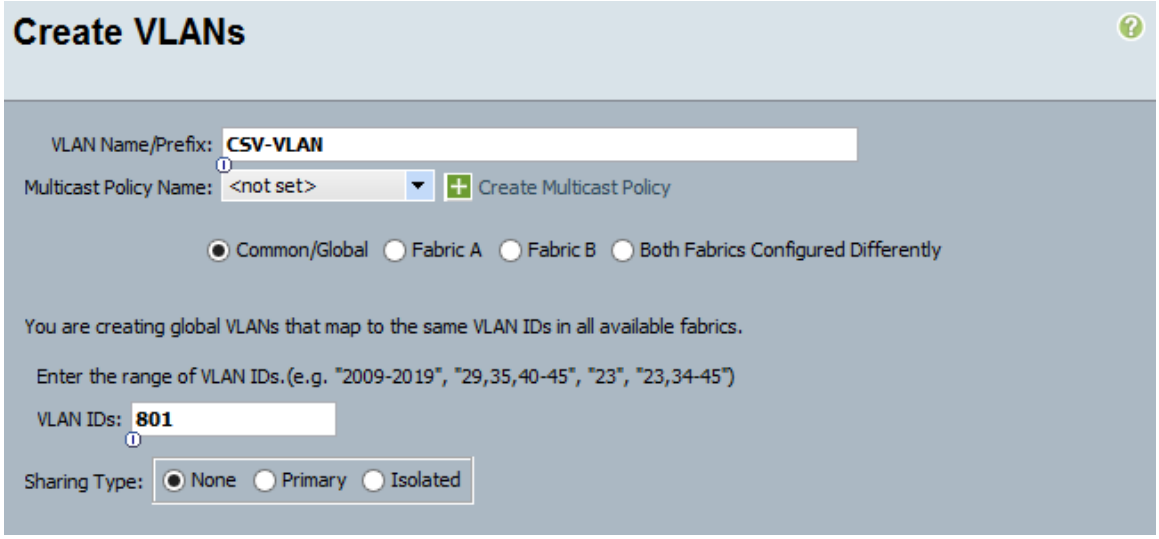

15. Right-click VLANs.
16. Select Create VLANs.
17. Enter iSCSI-A-VLAN as the name of the VLAN to be used for the first iSCSI VLAN.
18. Keep the Common/Global option selected for the scope of the VLAN.
19. Enter the VLAN ID for the first iSCSI VLAN.
20. Click **OK**.

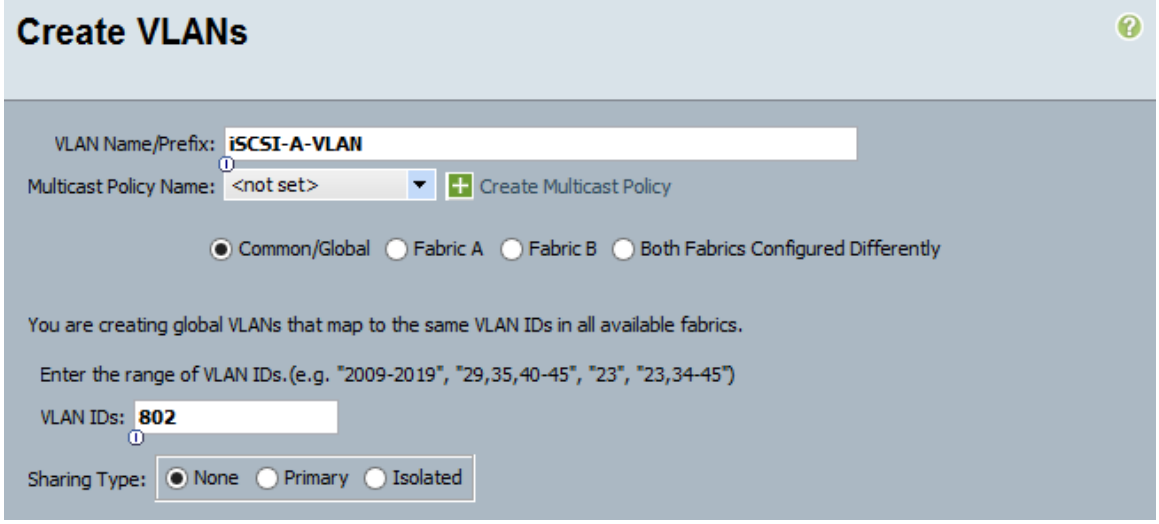

21. Right-click VLANs.
22. Select Create VLANs.
23. Enter iSCSI-VLAN-B as the name of the VLAN to be used for the second iSCSI VLAN.
24. Keep the Common/Global option selected for the scope of the VLAN.
25. Enter the VLAN ID for the second iSCSI VLAN.
26. Click **OK**.

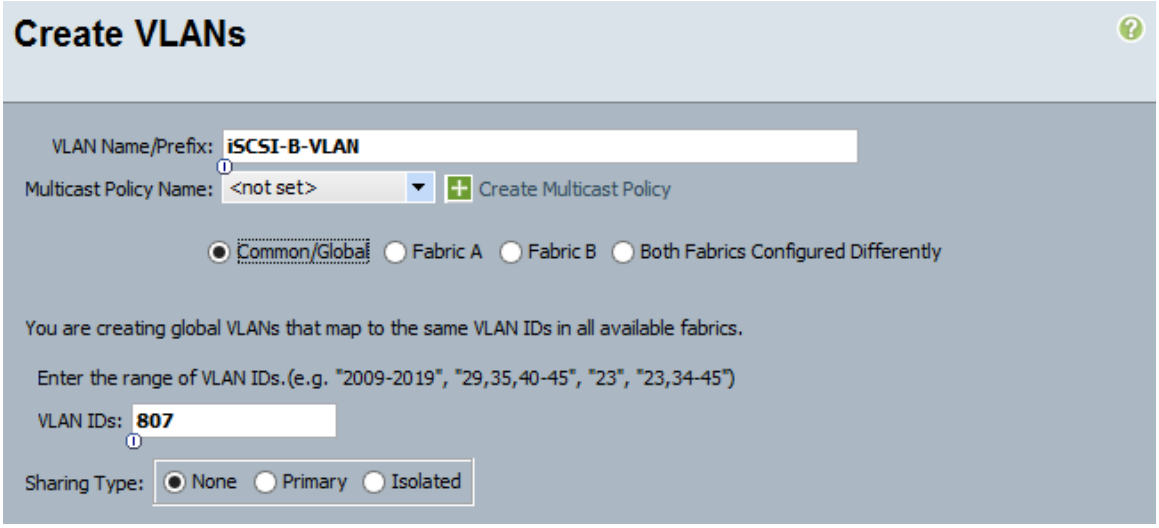

27. Right-click VLANs.
28. Select Create VLANs.
29. Enter Live Migration-VLAN as the name of the VLAN to be used for the live migration VLAN.
30. Keep the Common/Global option selected for the scope of the VLAN.
31. Enter the VLAN ID for the live migration VLAN.
32. Click **OK**.

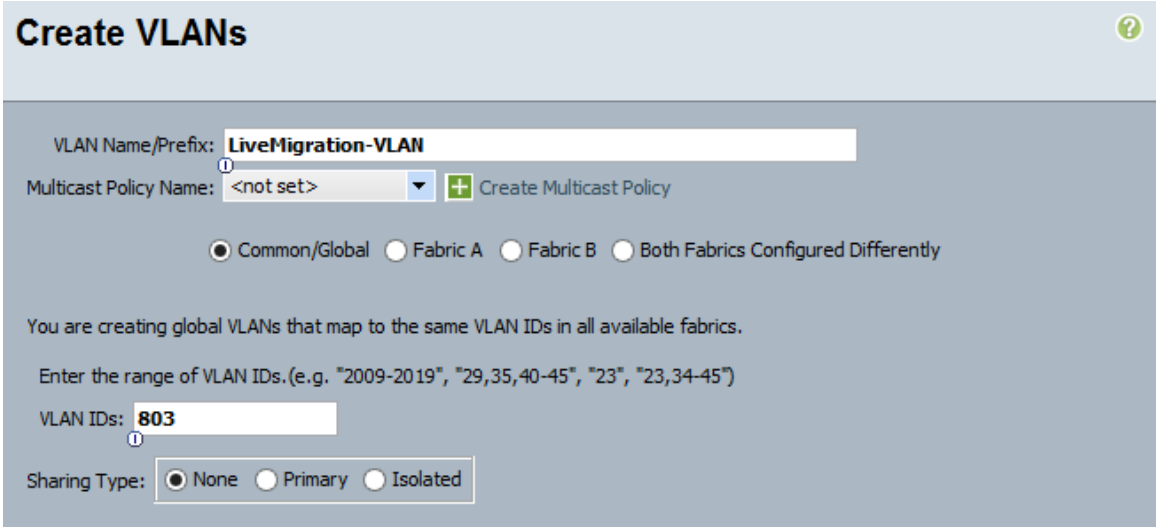

33. Right-click VLANs
34. Select Create VLANs.
35. Enter VM-Cluster-Comm-VLAN as the name of the VLAN to be used for the VM Cluster VLAN.
36. Keep the Common/Global option selected for the scope of the VLAN.
37. Enter the VLAN ID for the VM Cluster VLAN.
38. Click **OK**.

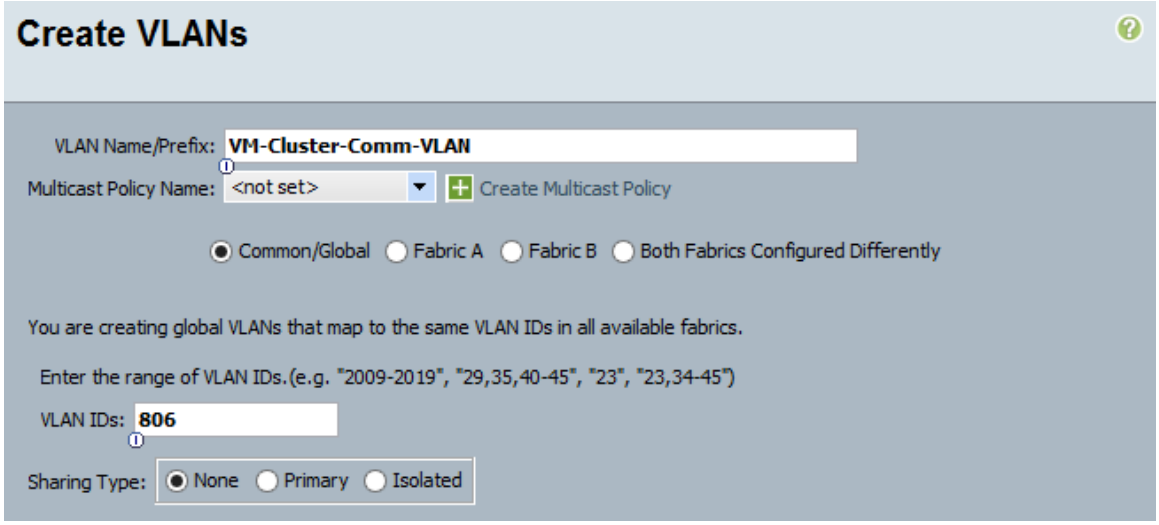

39. Right-click VLANs.
40. Select Create VLANs.
41. Enter VM-Public-VLAN as the name of the VLAN to be used for the VM data VLAN.
42. Keep the Common/Global option selected for the scope of the VLAN.
43. Enter the VLAN ID for the VM data VLAN.
44. Click **OK**.

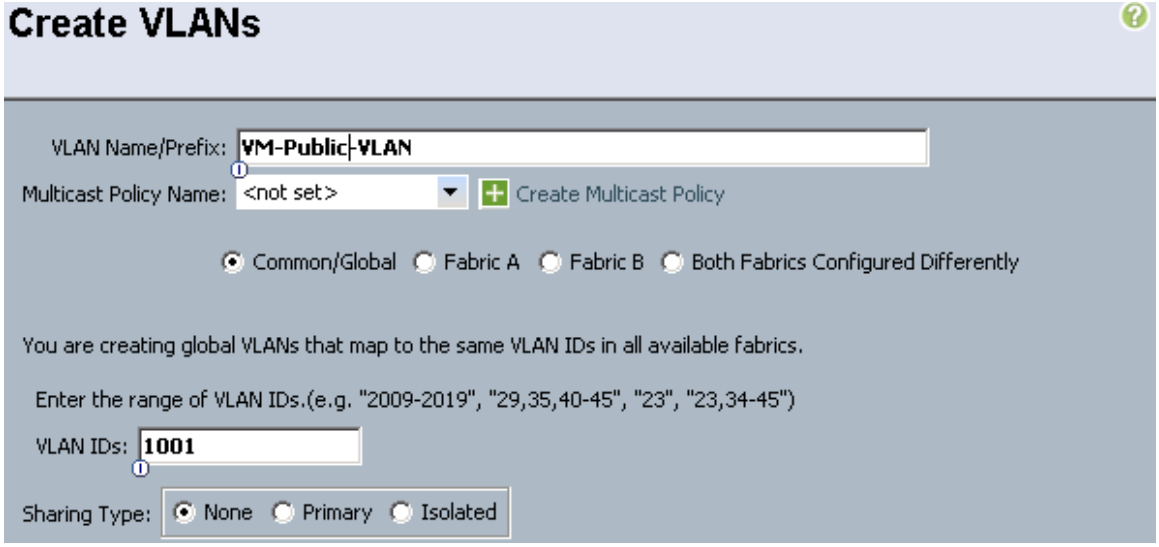

45. Right-click VLANs.
46. Select Create VLANs.
47. Enter Native-VLAN as the name of the VLAN to be used for the Native VLAN.
48. Keep the Common/Global option selected for the scope of the VLAN.
49. Enter the VLAN ID for the Native VLAN.
50. Click **OK**.

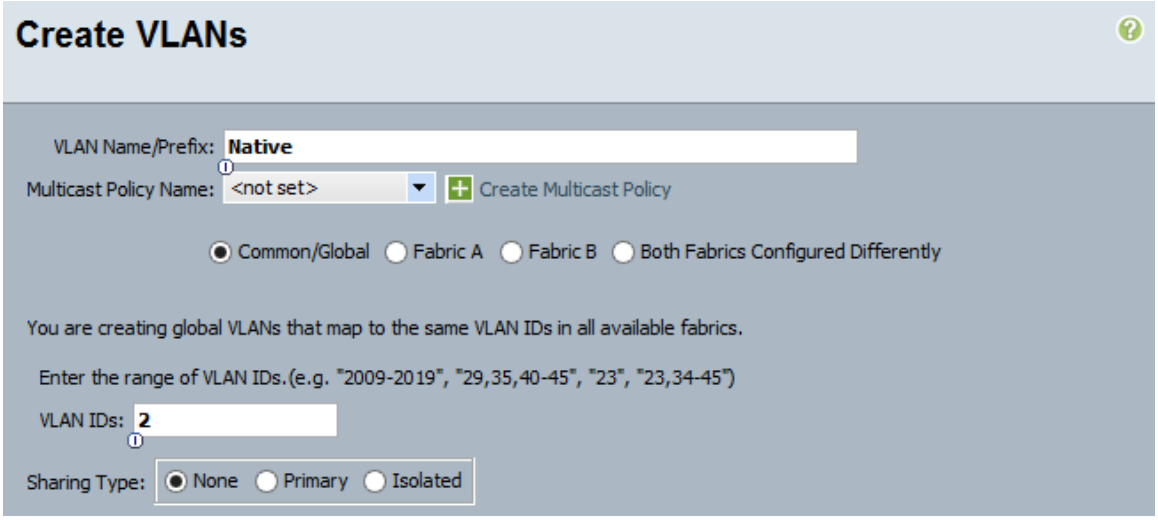

51. In the list of VLANs in the left pane, right-click the newly created Native-VLAN and select Set as Native VLAN.
52. Click **Yes** and **OK**.

Create VSANs and FCoE Port Channels

These steps provide details for configuring the necessary VSANs and FCoE Port Channels for the Cisco UCS environment.

1. Select the SAN tab at the top left of the window.
2. Expand the SAN Cloud tree.
3. Right-click VSANs
4. Select Create VSAN.
5. Enter VSAN_A as the VSAN name for fabric A.
6. Keep the Disabled option selected for the Default Zoning
7. Select Fabric A.
8. Enter the VSAN ID for fabric A.
9. Enter the FCoE VLAN ID for fabric A.
10. Click **OK** and then **OK** to create the VSAN.

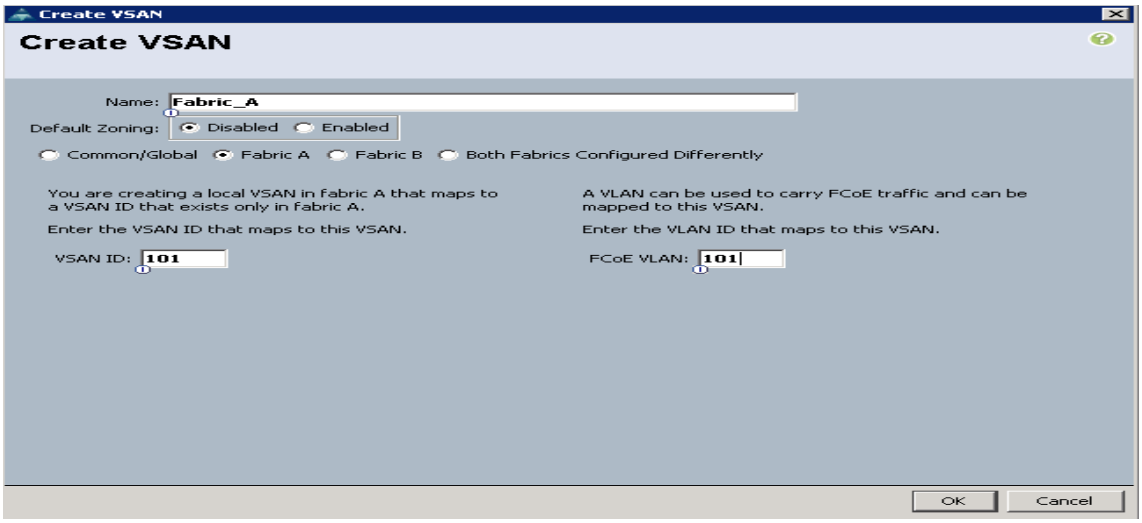

11. Right-click VSANs.
12. Select Create VSAN.
13. Enter VSAN_B as the VSAN name for fabric B.
14. Keep the Disabled option selected for the Default Zoning
15. Select Fabric B.
16. Enter the VSAN ID for fabric B.
17. Enter the FCoE VLAN ID for fabric B.
18. Click **OK** and then **OK** to create the VSAN.

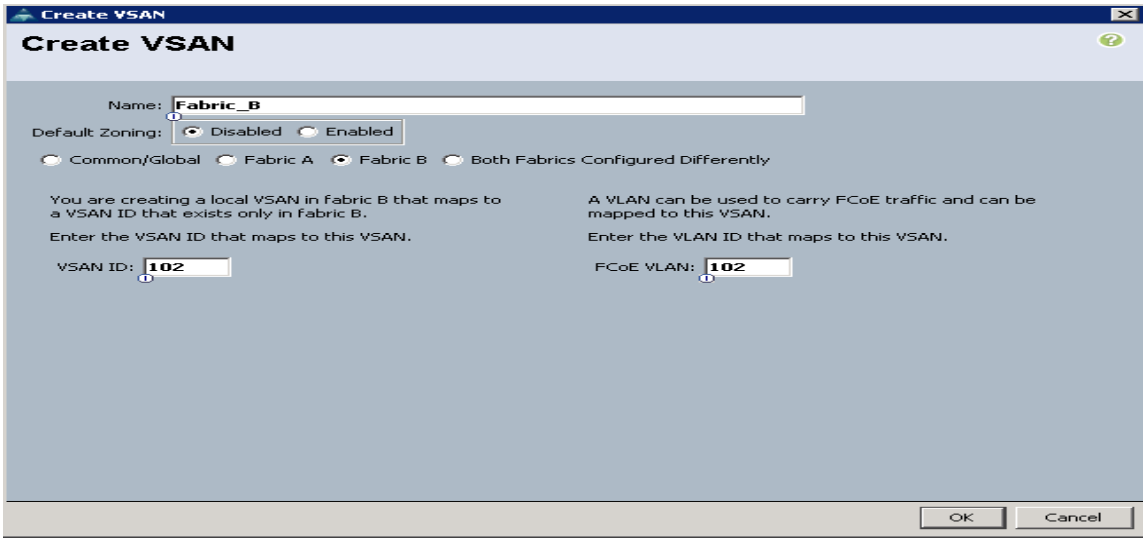

19. Under SAN Cloud, expand the Fabric A tree.
20. Right-click FCoE Port Channels
21. Select Create FCoE Port Channel.
22. Click **Yes** and then enter 101 for the Port Channel ID and FCoE_PC_Fabric-A for the Port Channel name.
23. Click **Next**.

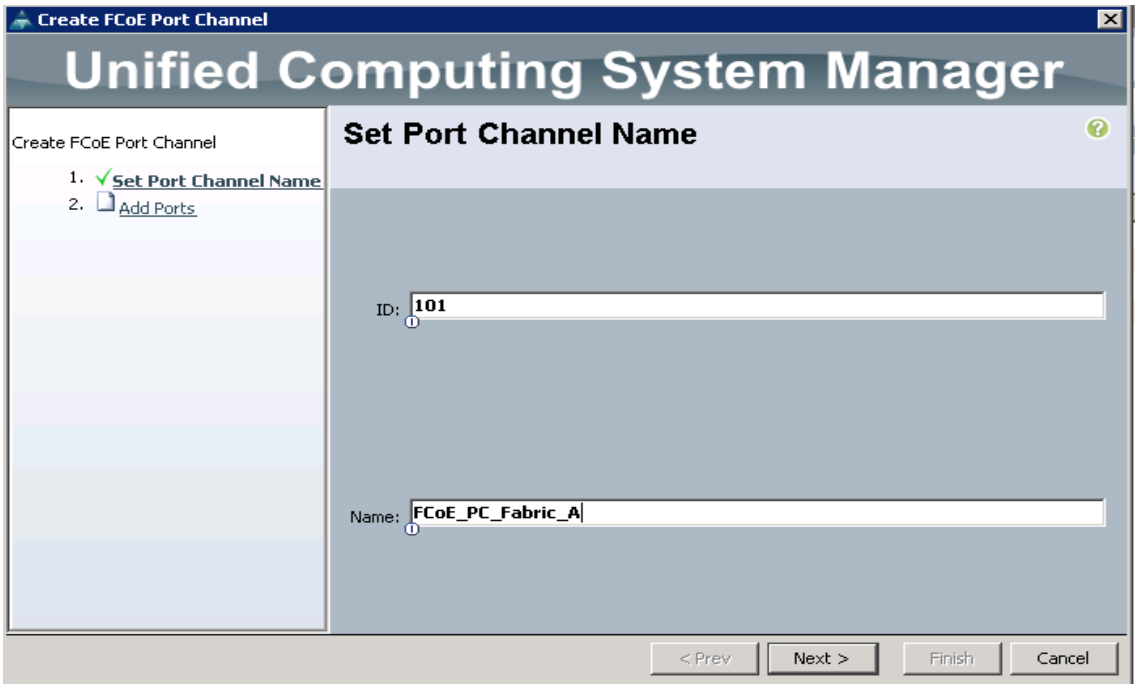

24. Select ports 31 and 32 and click **>>** to add the ports to the Port Channel.
25. Click **Finish**.

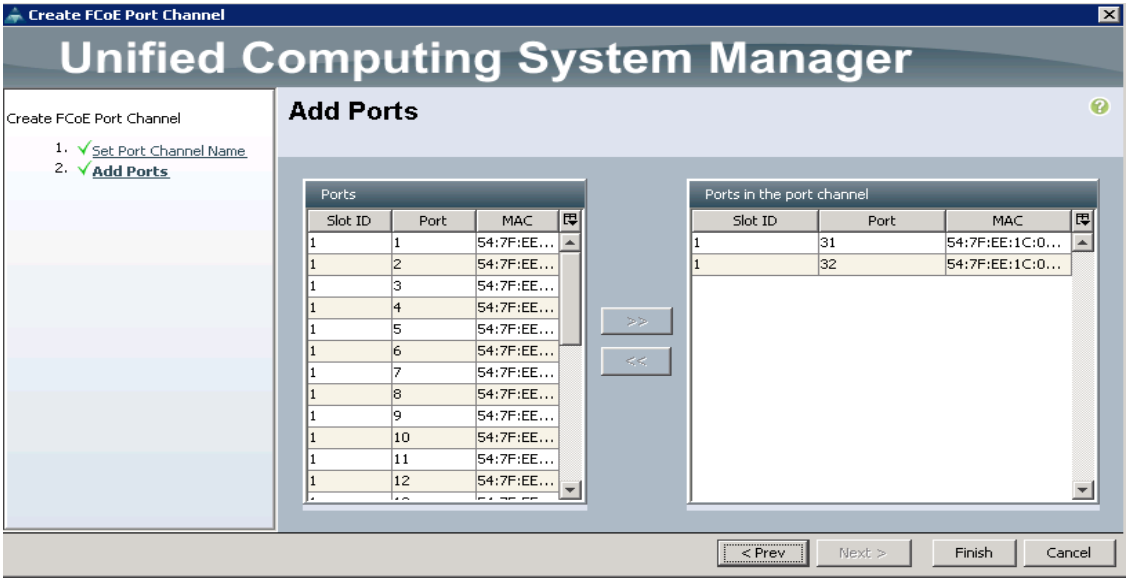

26. Check the Show navigator for FCoE Port-Channel 101 (Fabric A) checkbox.
27. Click **OK** to complete creating the FCoE Port Channel.
28. In the VSAN pull-down under Properties select the vsan VSAN_A for fabric A.
29. Click **Apply**, then click **OK**.
30. Under Actions, click **Enable Port Channel**.
31. Click **Yes** and then **OK** to enable the Port Channel. This action also enables the two FCoE ports in the Port Channel.

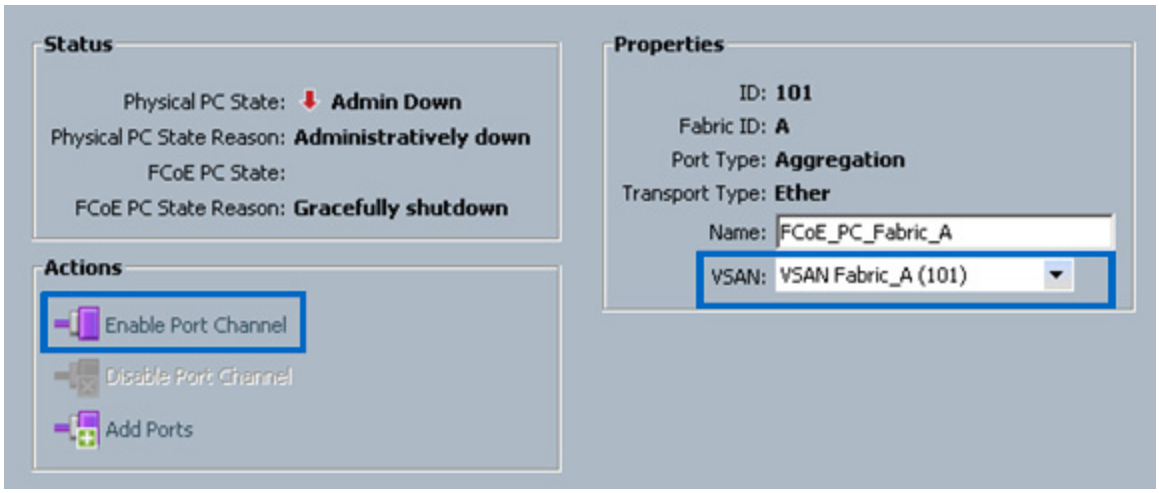

32. Click **OK** to Close the Navigator.

Note The FCoE Port Channel may take a few seconds to come up. The operational speed will be displayed when the link speed is negotiated. This may take approximately 30 seconds.

If the Overall State results in an error condition and does not clear after 30 seconds the FC uplink ports on the Nexus 5548UP will need to shut down and brought back up in order to establish the link.

33. Under SAN Cloud, expand the Fabric B tree.
34. Right-click FCoE Port Channels
35. Select Create FcoE Port Channel.
36. Click **Yes**, and then enter 102 for the Port Channel ID and FCoE_PC_Fabric_B for the Port Channel name.
37. Click **Next**.

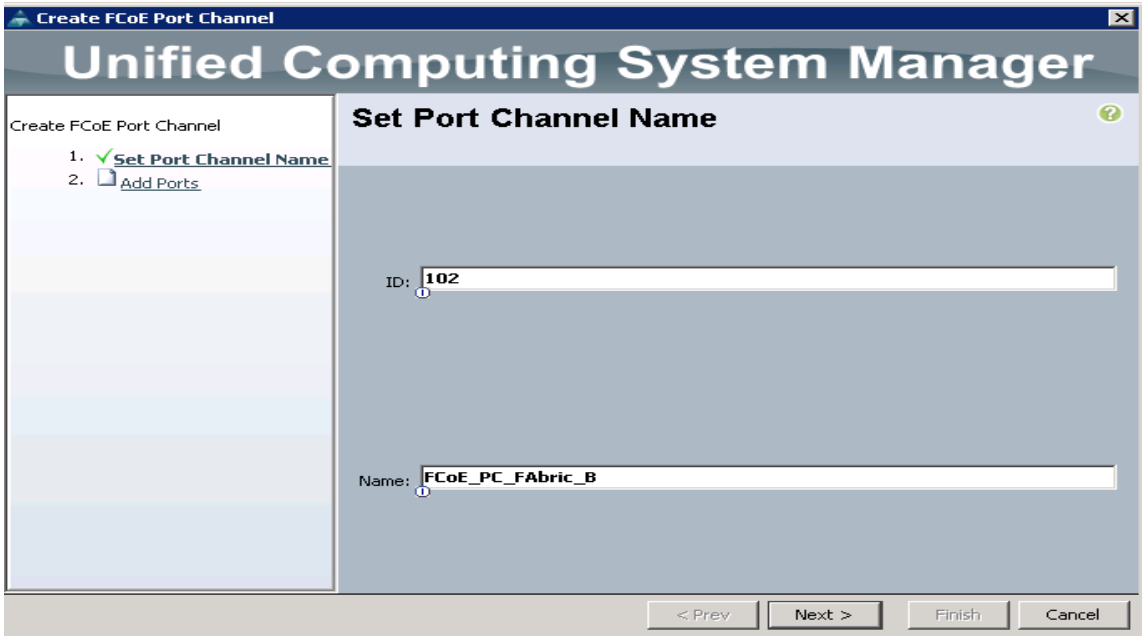

38. Select ports 31 and 32 and click **>>** to add the ports to the Port Channel.
39. Click **Finish**.

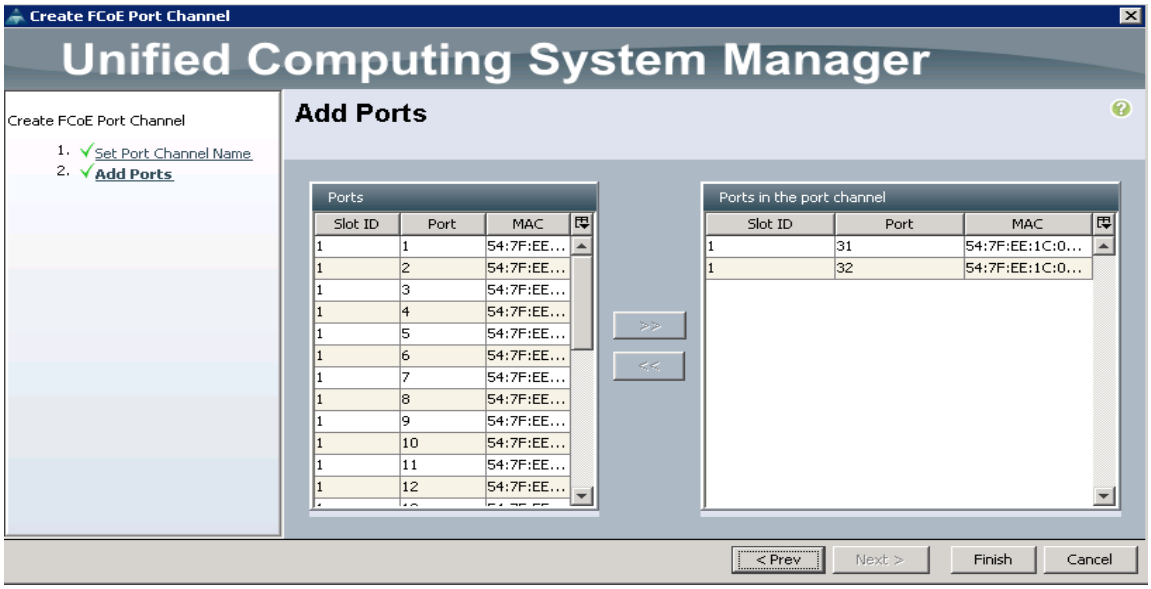

40. Select the Show navigator for FCoE Port-Channel 102 (Fabric B) checkbox.
41. Click **OK** to complete creating the Port Channel.
42. In the VSAN pull-down under Properties, select VSAN_B for fabric B.
43. Click **Apply**, then click **OK**.
44. Under Actions, click **Enable Port Channel**.
45. Click **Yes**, then **OK** to enable the Port Channel. This action also enables the two FC ports in the Port Channel.

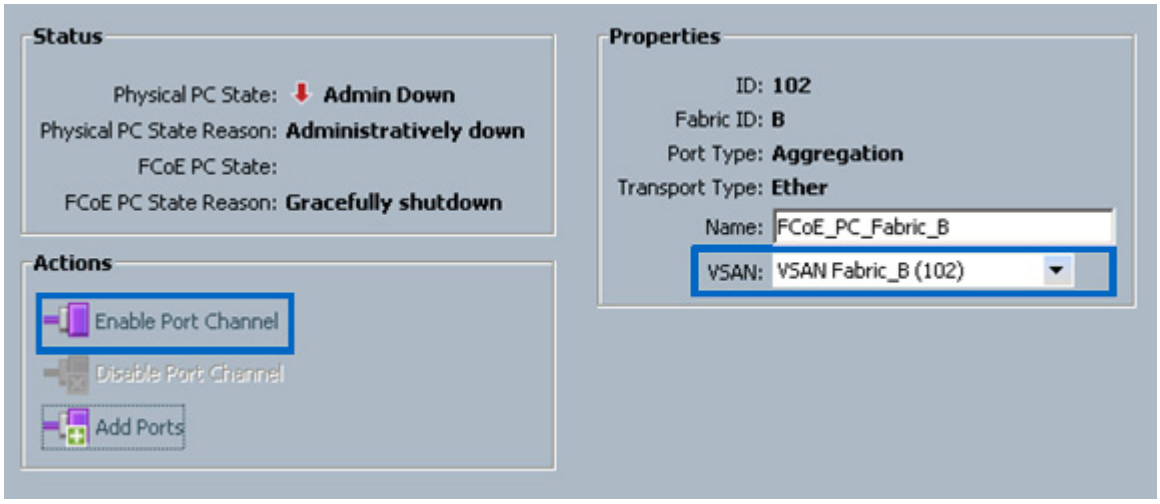

46. Click **OK** to Close the Navigator.

Note

- The FC Port Channel may take a few seconds to come up. The operational speed will be displayed when the link speed is negotiated. This may take approximately 30 seconds.
- If the Overall State results in an error condition and does not clear after 30 seconds the FC uplink ports on the Nexus 5548UP will need to shut down and brought back up in order to establish the link.

Create a FC Adapter Policy for NetApp Storage Arrays

These steps provide details for a FC adapter policy for NetApp storage arrays.

1. Select to the SAN tab at the top of the left window.
2. Go to **SAN** > **Policies** > **root**.
3. Right-click Fibre Channel Adapter Policies and click Create New Fibre Channel Adapter Policy.
4. Use Windows-NetApp as the name of the Fibre Channel Adapter Policy.
5. The default values are appropriate for most configurable items. Expand the Options dropdown. and set the Link Down Timeout (MS) option to 5000.
6. Click **OK** to complete creating the FC adapter policy.
7. Click **OK**.

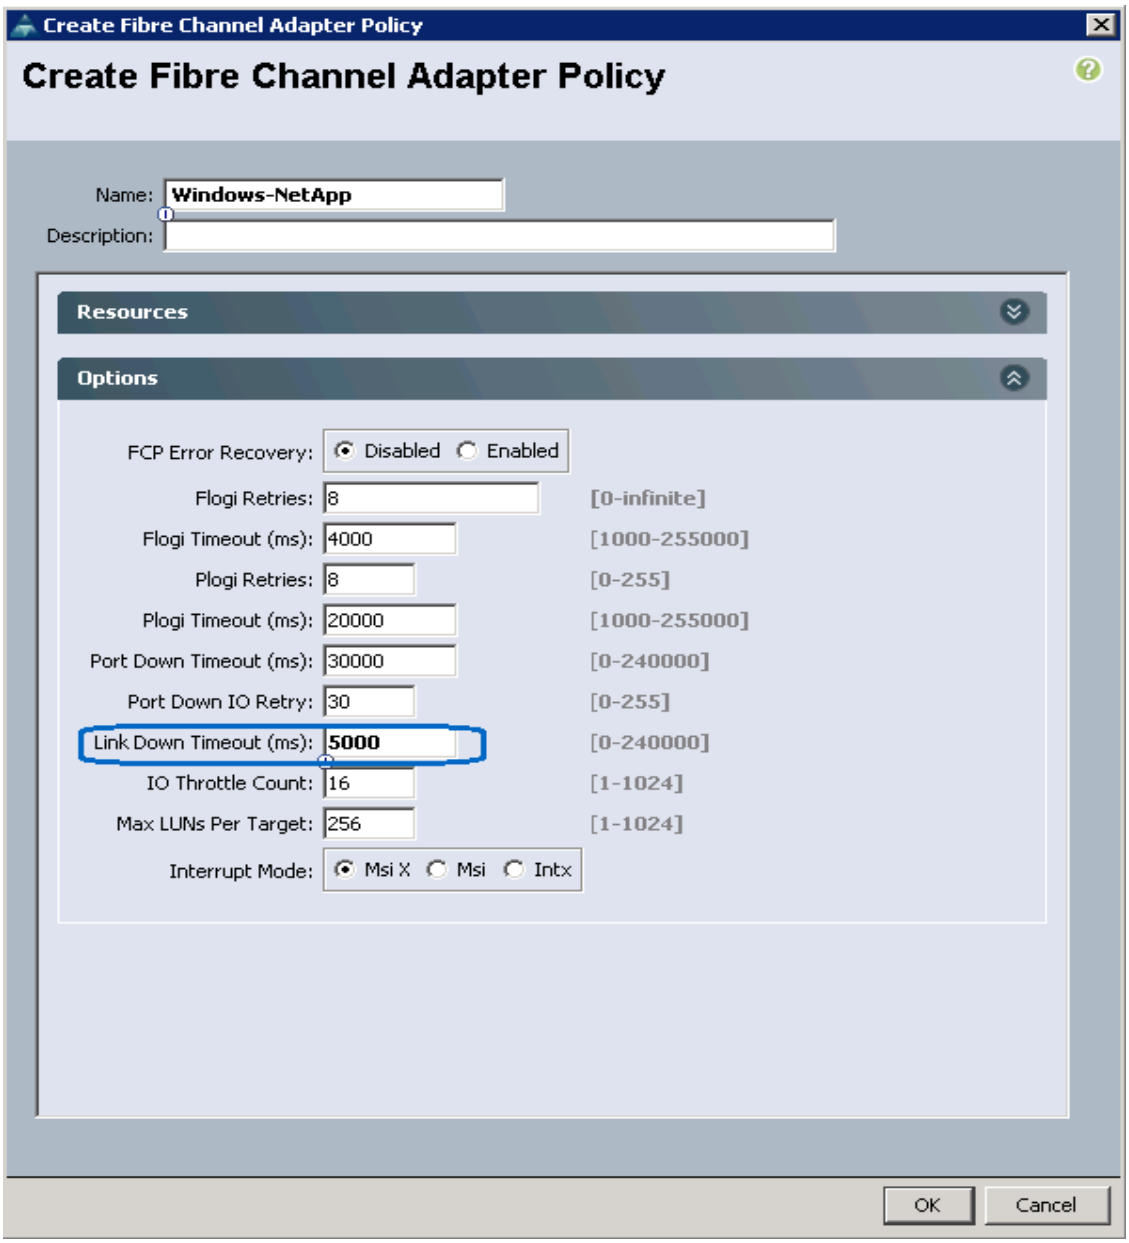

Create a Firmware Management Package

These steps provide details for a firmware management policy for n the Cisco UCS environment.

1. Select the Servers tab at the top left of the window.
2. Select **Policies** > **root**.
3. Right-click Management Firmware Packages
4. Select create Management Firmware Package.
5. Enter VM-Host-Infra as the management firmware package name.

6. Select the appropriate packages and versions of the Server Management Firmware For servers that you have.
7. Click **OK** to complete creating the management firmware package.
8. Click **OK**.

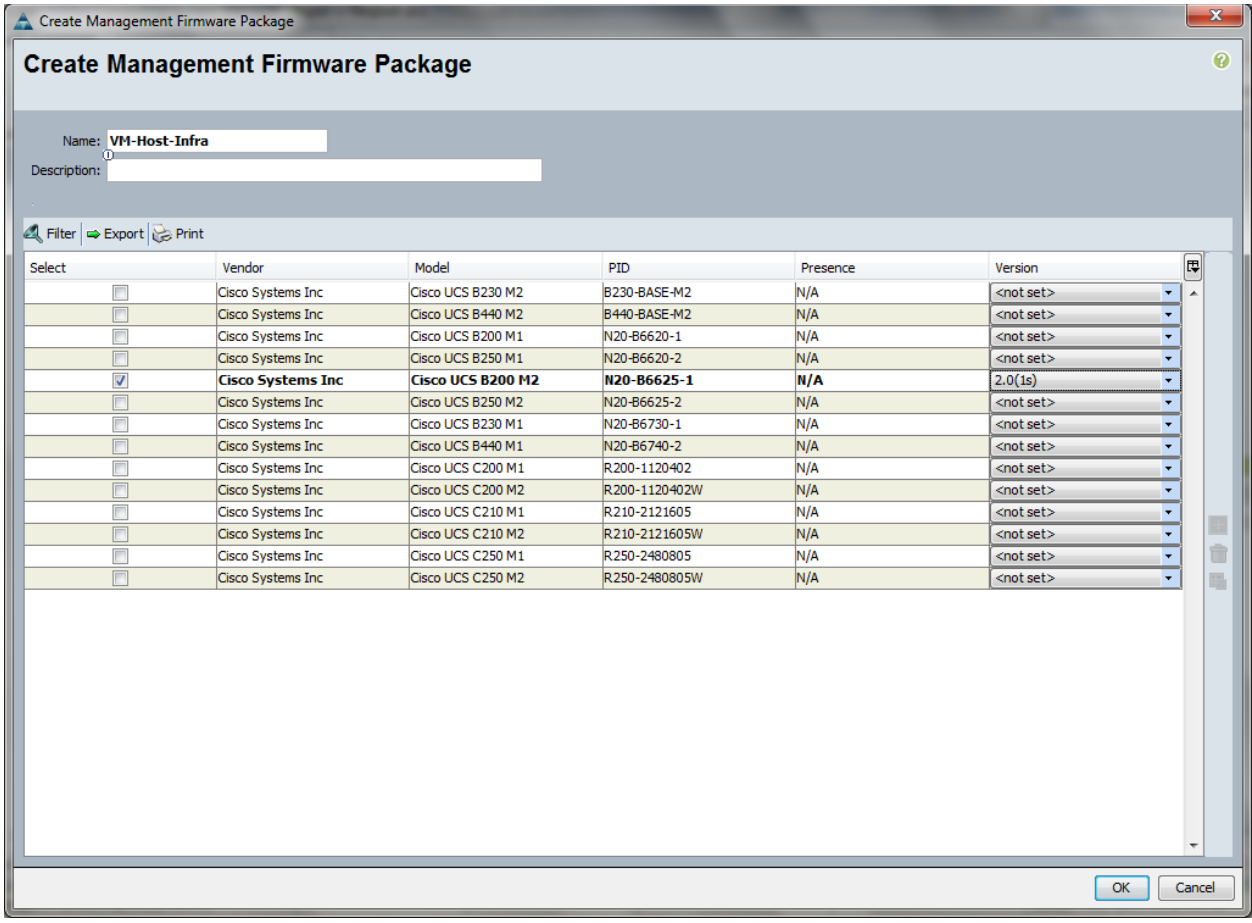

Create Firmware Package Policy

These steps provide details for creating a firmware management policy for a given server configuration in the Cisco UCS environment. Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These often include adapter, BIOS, board controller, FC adapters, HBA option ROM, and storage controller properties.

1. Select the Servers tab at the top left of the window.
2. Select **Policies** > **root**.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter the name of the host firmware package for the corresponding server configuration.
6. Navigate the tabs of the Create Host Firmware Package Navigator and select the appropriate packages and versions for the server configuration.
7. Click **OK** to complete creating the host firmware package.

8. Click **OK**.

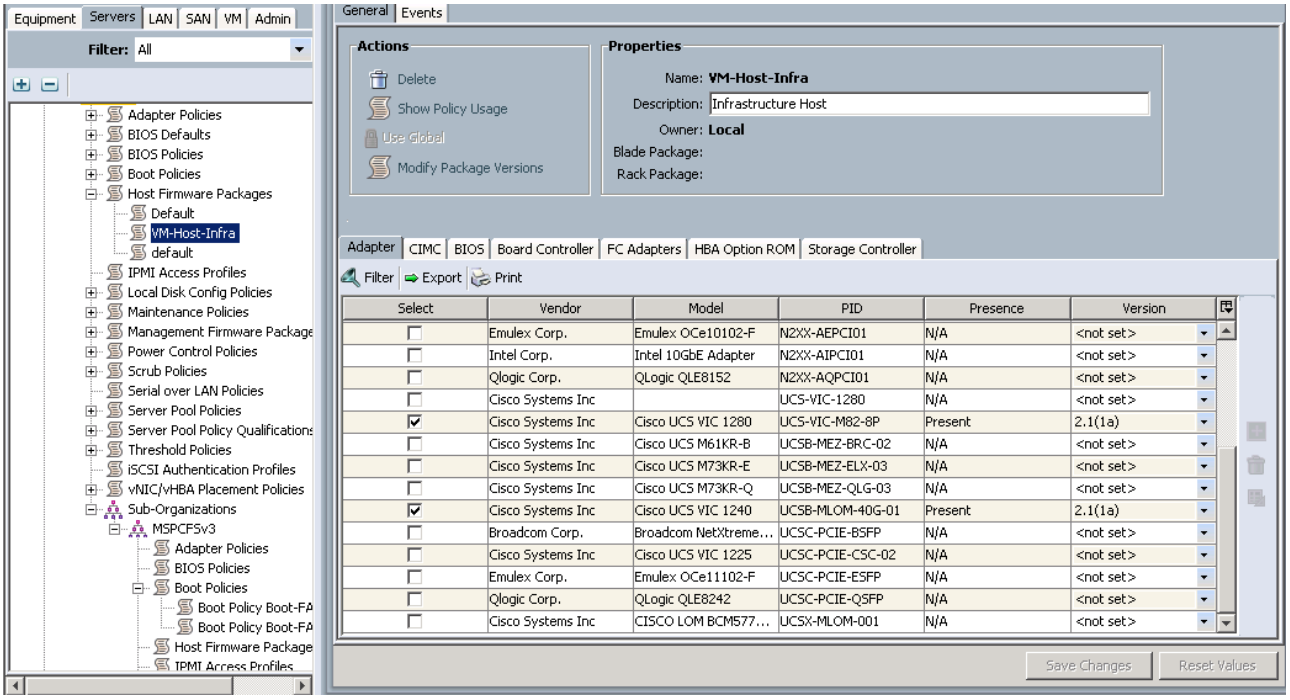
Equipment
Servers
LAN
SAN
VM
Admin
Filter: All
Adapter Policies
BIOS Defaults
BIOS Policies
Boot Policies
Host Firmware Packages
Default
VM-Host-Infra
default
IPMI Access Profiles
Local Disk Config Policies
Maintenance Policies
Management Firmware Package
Power Control Policies
Scrub Policies
Serial over LAN Policies
Server Pool Policies
Server Pool Policy Qualifications
Threshold Policies
iSCSI Authentication Profiles
vNIC/vHBA Placement Policies
Sub-Organizations
MSPCFSv3
Adapter Policies
BIOS Policies
Boot Policies
Boot Policy Boot-FA
Boot Policy Boot-FA
Host Firmware Package
IPMI Access Profiles
General
Events
Actions
Delete
Show Policy Usage
Use Global
Modify Package Versions
Properties
Name: VM-Host-Infra
Description: Infrastructure Host
Owner: Local
Blade Package:
Rack Package:
Adapter
CIMC
BIOS
Board Controller
FC Adapters
HBA Option ROM
Storage Controller
Filter
Export
Print
Select
Vendor
Model
PID
Presence
Version
Emulex Corp. Emulex OCe10102-F N2XX-AEPCI01 N/A <not set>
Intel Corp. Intel 10GbE Adapter N2XX-AIPCI01 N/A <not set>
Qlogic Corp. QLogic QLE8152 N2XX-AQPCI01 N/A <not set>
Cisco Systems Inc UCS-VIC-1280 N/A <not set>
Cisco Systems Inc Cisco UCS VIC 1280 UCS-VIC-M82-8P Present 2.1(1a)
Cisco Systems Inc Cisco UCS M61KR-B UCSB-MEZ-BRC-02 N/A <not set>
Cisco Systems Inc Cisco UCS M73KR-E UCSB-MEZ-ELX-03 N/A <not set>
Cisco Systems Inc Cisco UCS M73KR-Q UCSB-MEZ-QLG-03 N/A <not set>
Cisco Systems Inc Cisco UCS VIC 1240 UCSB-MLOM-40G-01 Present 2.1(1a)
Broadcom Corp. Broadcom NetXtreme... UCSC-PCIE-BSFP N/A <not set>
Cisco Systems Inc Cisco UCS VIC 1225 UCSC-PCIE-CSC-02 N/A <not set>
Emulex Corp. Emulex OCe11102-F UCSC-PCIE-ESFP N/A <not set>
Qlogic Corp. QLogic QLE8242 UCSC-PCIE-QSFP N/A <not set>
Cisco Systems Inc CISCO LOM BCM577... UCSX-MLOM-001 N/A <not set>
Save Changes
Reset Values

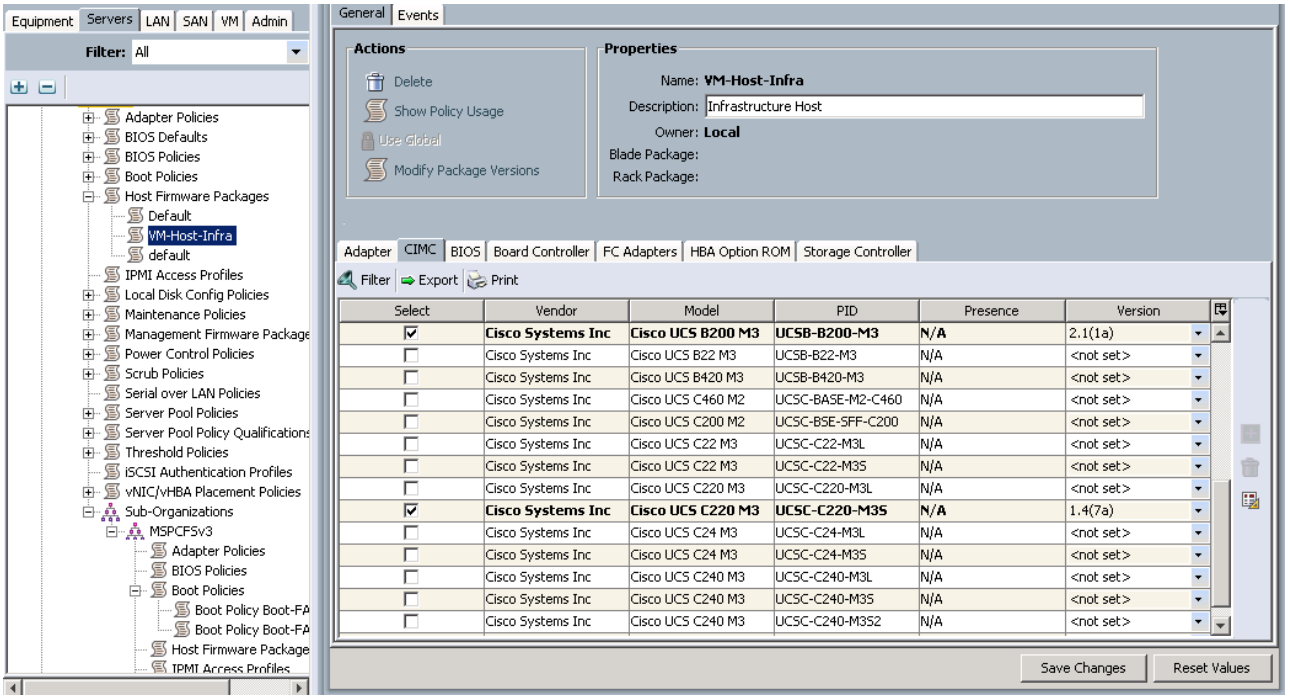
Equipment
Servers
LAN
SAN
VM
Admin
Filter: All
Adapter Policies
BIOS Defaults
BIOS Policies
Boot Policies
Host Firmware Packages
Default
VM-Host-Infra
default
IPMI Access Profiles
Local Disk Config Policies
Maintenance Policies
Management Firmware Package
Power Control Policies
Scrub Policies
Serial over LAN Policies
Server Pool Policies
Server Pool Policy Qualifications
Threshold Policies
iSCSI Authentication Profiles
vNIC/vHBA Placement Policies
Sub-Organizations
MSPCFSv3
Adapter Policies
BIOS Policies
Boot Policies
Boot Policy Boot-FA
Boot Policy Boot-FA
Host Firmware Package
IPMI Access Profiles
General
Events
Actions
Delete
Show Policy Usage
Use Global
Modify Package Versions
Properties
Name: VM-Host-Infra
Description: Infrastructure Host
Owner: Local
Blade Package:
Rack Package:
Adapter
CIMC
BIOS
Board Controller
FC Adapters
HBA Option ROM
Storage Controller
Filter
Export
Print
Select
Vendor
Model
PID
Presence
Version
Cisco Systems Inc Cisco UCS B200 M3 UCSB-B200-M3 N/A 2.1(1a)
Cisco Systems Inc Cisco UCS B22 M3 UCSB-B22-M3 N/A <not set>
Cisco Systems Inc Cisco UCS B420 M3 UCSB-B420-M3 N/A <not set>
Cisco Systems Inc Cisco UCS C460 M2 UCSC-BASE-M2-C460 N/A <not set>
Cisco Systems Inc Cisco UCS C200 M2 UCSC-BSE-SFF-C200 N/A <not set>
Cisco Systems Inc Cisco UCS C22 M3 UCSC-C22-M3L N/A <not set>
Cisco Systems Inc Cisco UCS C22 M3 UCSC-C22-M3S N/A <not set>
Cisco Systems Inc Cisco UCS C220 M3 UCSC-C220-M3L N/A <not set>
Cisco Systems Inc Cisco UCS C220 M3 UCSC-C220-M3S N/A 1.4(7a)
Cisco Systems Inc Cisco UCS C24 M3 UCSC-C24-M3L N/A <not set>
Cisco Systems Inc Cisco UCS C24 M3 UCSC-C24-M3S N/A <not set>
Cisco Systems Inc Cisco UCS C240 M3 UCSC-C240-M3L N/A <not set>
Cisco Systems Inc Cisco UCS C240 M3 UCSC-C240-M3S N/A <not set>
Cisco Systems Inc Cisco UCS C240 M3 UCSC-C240-M3S2 N/A <not set>
Save Changes
Reset Values

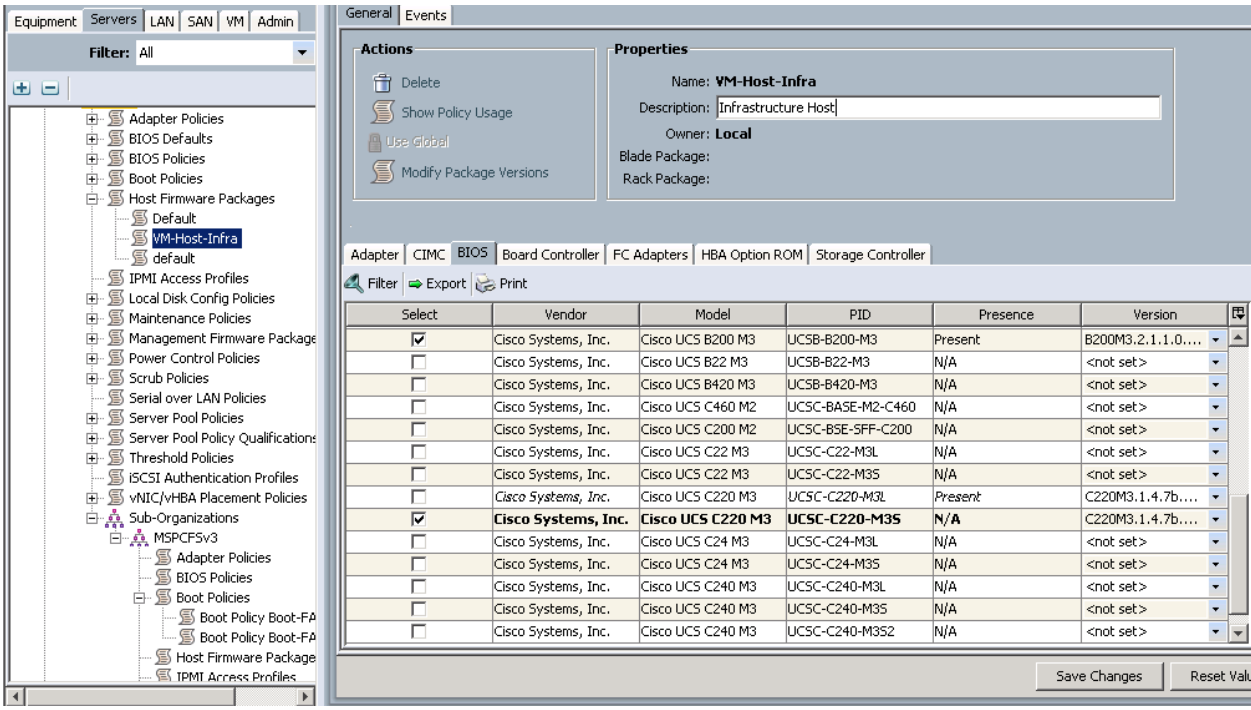

Set Jumbo Frames and Enable Quality of Service in Cisco UCS Fabric

These steps provide details for setting Jumbo frames and enabling the quality of server in the Cisco UCS Fabric.

1. Select the LAN tab at the top left of the window.
2. Go to **LAN Cloud** > **QoS System Class**.
3. In the right pane, select the General tab
4. On the Gold and Silver Priority, and Best Efforts row, type 9000 in the MTU boxes.
5. Click **Save Changes** in the bottom right corner.
6. Click **OK** to continue.

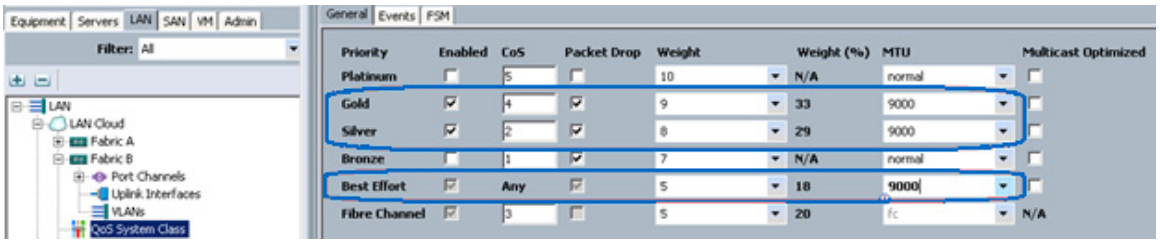

7. Select the LAN tab on the left of the window.
8. Go to **LAN** > **Policies** > **root**.
9. Right-click QoS Policies.
10. Select Create QoS Policy.
11. Enter LiveMigration as the QoS Policy name.

12. Change the Priority to Gold. Leave Burst (Bytes) set to 10240. Leave Rate (Kbps) set to line-rate. Leave Host Control set to None.
13. Click **OK** in the bottom right corner.

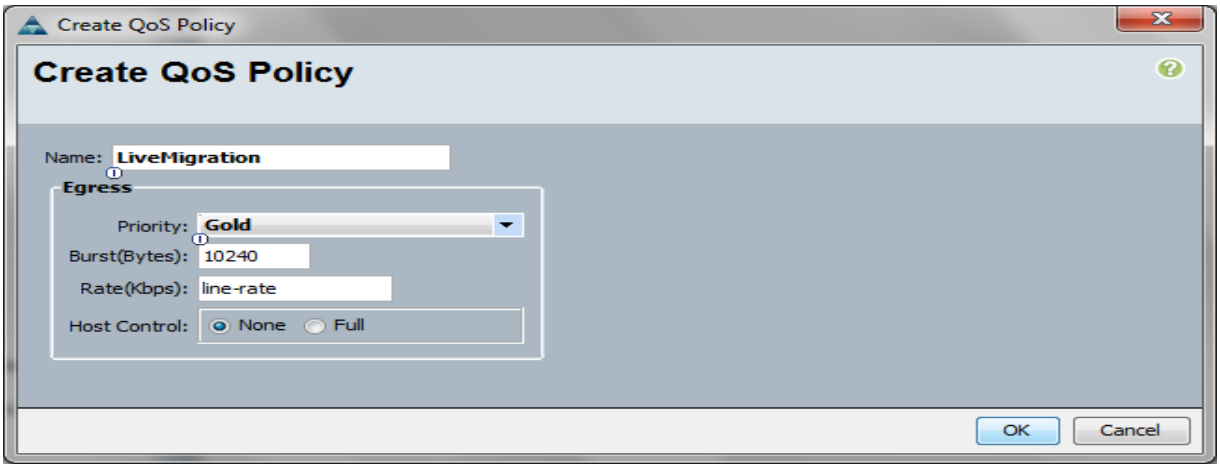

14. Right-click QoS Policies.
15. Select Create QoS Policy.
16. Enter CSV as the QoS Policy name.
17. Change the Priority to Gold. Leave Burst (Bytes) set to 10240. Leave Rate (Kbps) set to line-rate. Leave Host Control set to None.
18. Click **OK** in the bottom right corner.

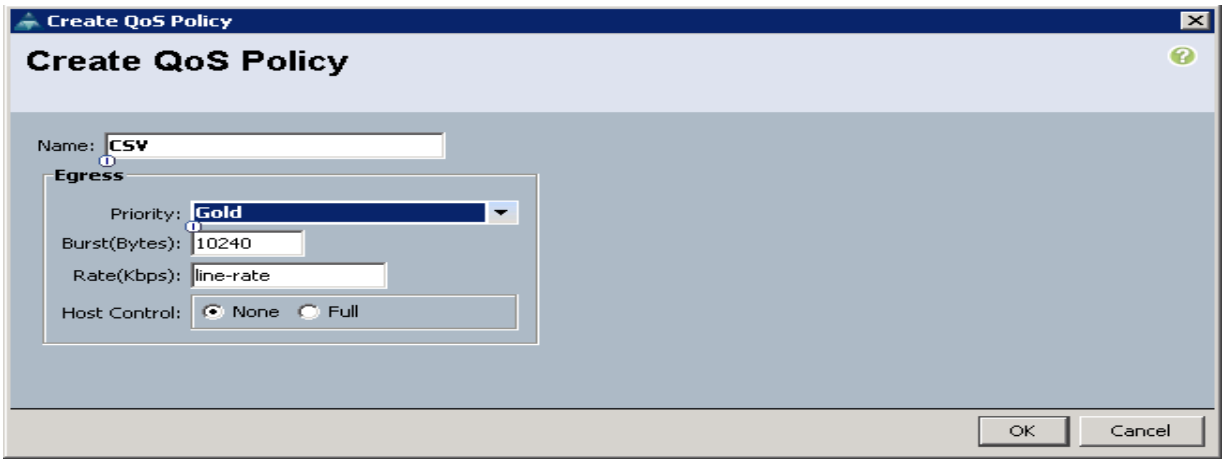

19. Right-click QoS Policies.
20. Select Create QoS Policy.
21. Enter iSCSI as the QoS Policy name.
22. Change the Priority to Silver. Leave Burst (Bytes) set to 10240. Leave Rate (Kbps) set to line-rate. Leave Host Control set to None.
23. Click **OK** in the bottom right corner.

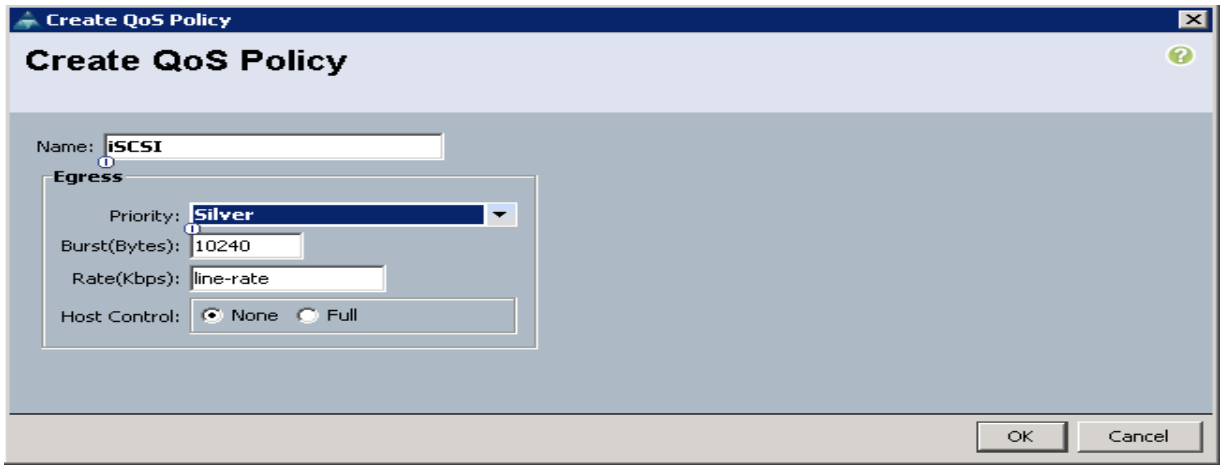

Create a Power Control Policy

These steps provide details for creating a Power Control Policy for the Cisco UCS environment.

1. Select the Servers tab at the top left of the window.
2. Go to **Policies** > **root**.
3. Right-click Power Controller Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the Power Capping to No Cap.
7. Click **OK** to complete creating the host firmware package.
8. Click **OK**.

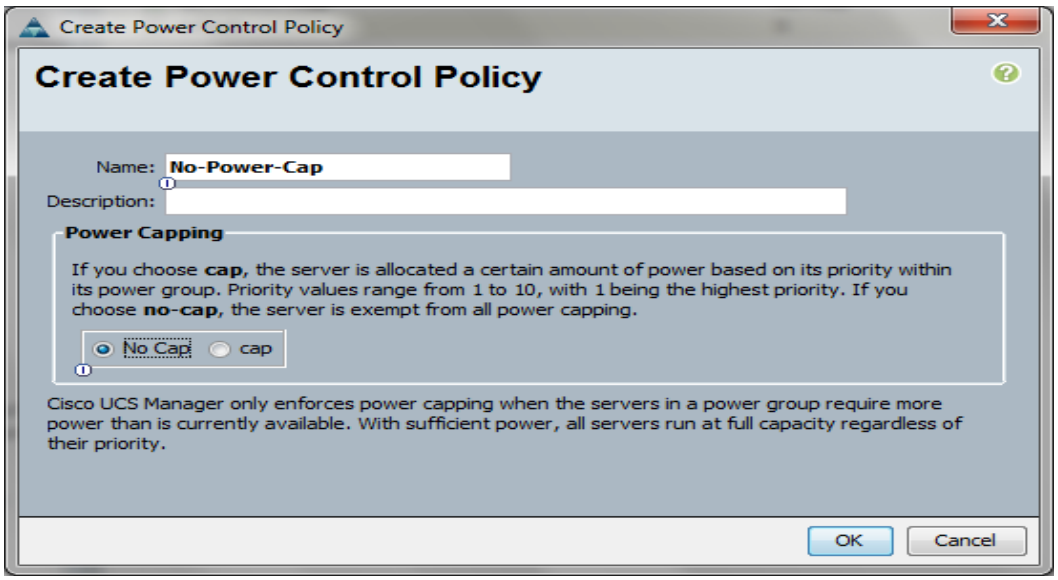

Create a Local Disk Configuration Policy

These steps provide details for creating a local disk configuration for the Cisco UCS environment, which is necessary if the servers in question do not have a local disk.

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

1. Select the Servers tab on the left of the window.
2. Go to **Policies** > **root**.
3. Right-click Local Disk Config Policies.
4. Select Create Local Disk Configuration Policy.
5. Enter SAN Boot as the local disk configuration policy name.
6. Change the Mode to No Local Storage. Uncheck the Protect Configuration checkbox.
7. Click **OK** to complete creating the host firmware package.
8. Click **OK**.

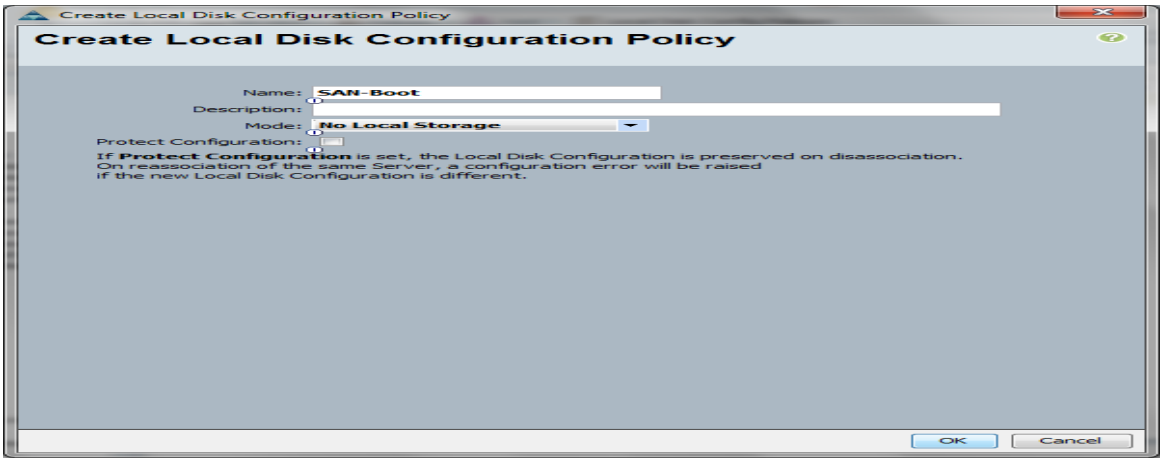

Create a Maintenance Policy

These steps provide details for creating a maintenance policy. The maintenance policy controls the timing of a server reboot after an update has been made that requires the server to reboot prior to the update taking affect.

1. Select the Servers tab on the left of the window.
2. Go to **Policies** > **root or sub-organization**
3. Right-click Maintenance Policy and select Create Maintenance Policy.
4. Name the policy User_Acknowledge
5. Select the User Ack option.
6. Click **OK** to create the policy.

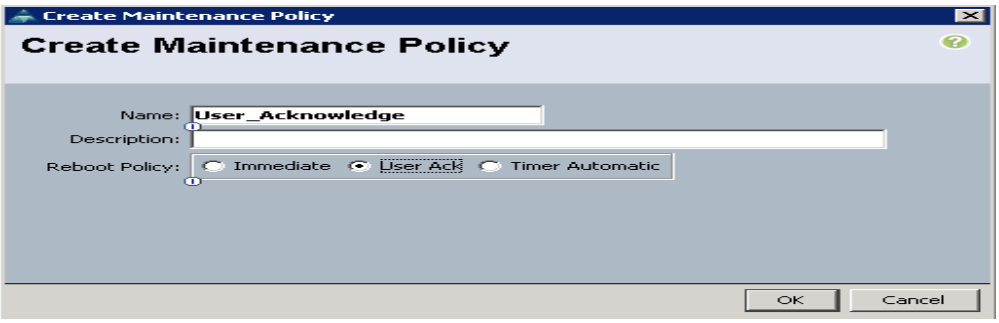

Create a Server Pool Qualification Policy

These steps provide details for creating a server pool qualification policy for the Cisco UCS environment.

1. Select the Servers tab on the left of the window
2. Go to **Policies** > **root**.
3. Right-click Server Pool Qualification Policies.
4. Select Create Server Pool Policy Qualification.
5. Select Server Model Qualifications.
6. Enter B200 M3 or C220 M3 as the Model (RegEx).
7. Click **OK** to complete creating the host firmware package.
8. Click **OK**.

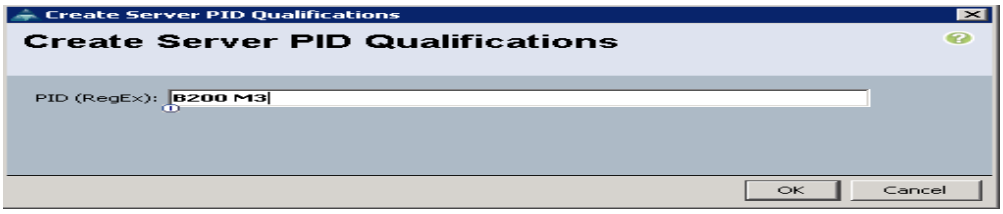

Create a Server BIOS Policy

These steps provide details for creating a server BIOS policy for the Cisco UCS environment.

1. Select the Servers tab on the left of the window.
2. Go to **Policies** > **root**.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host-Infra as the BIOS policy name.
6. Make the following changes to support VM-FEX (SR-IOV):

Property	Setting
Quiet Boot	Disabled
Virtual Technology (VT)	Enabled
VT Fort Direct IO	Enabled
Interrupt Remap	Enabled
Coherency Support	Disabled
ATS Support	Enabled
Pass Through DMA Support	Enabled

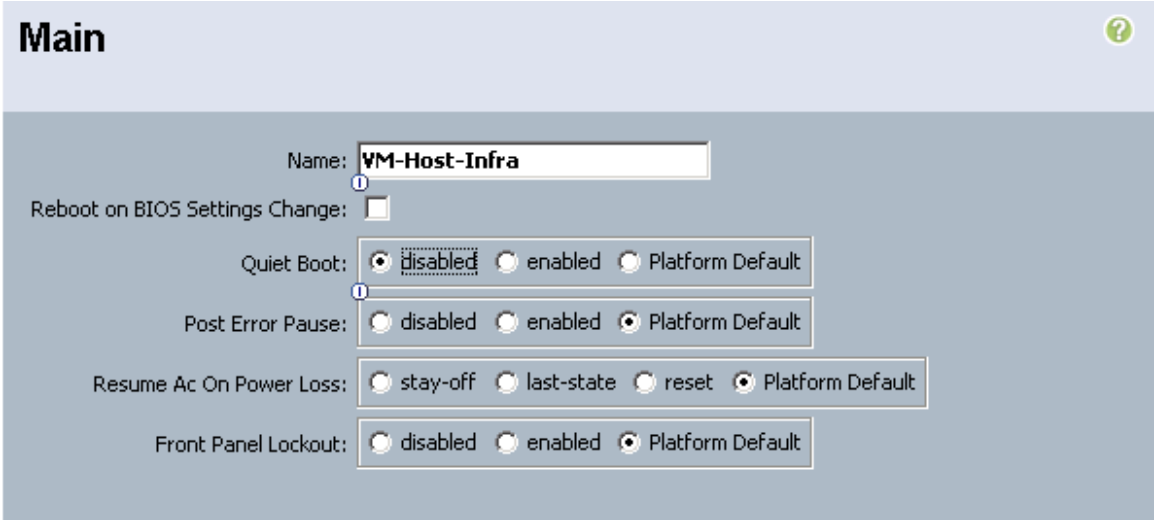

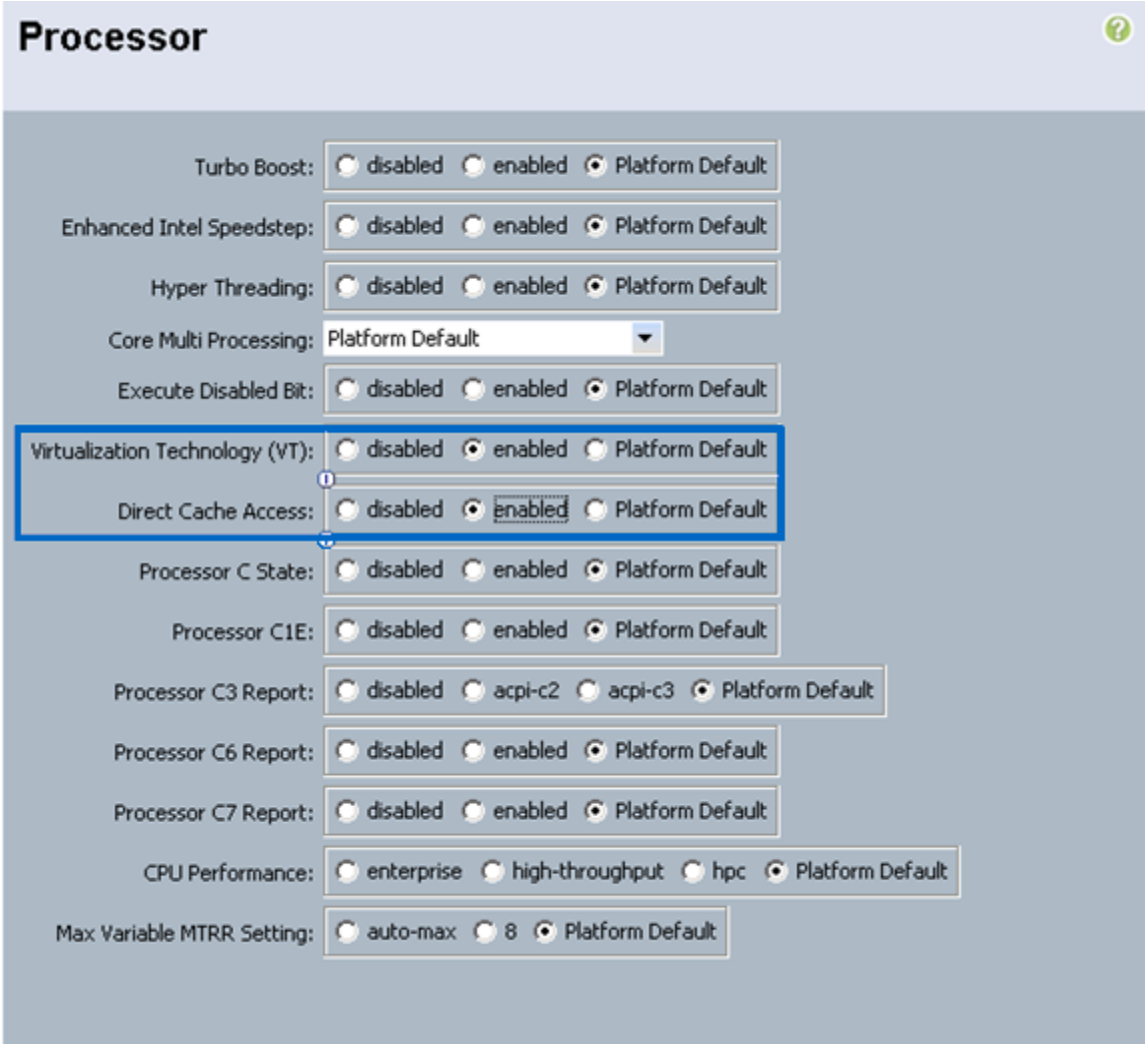

7. Click **Finish** to complete creating the BIOS policy.

8. Click **OK**.

Create Dynamic vNIC Connection Policy for VM-FEX (SR-IOV)

These steps provide details for creating the vNIC Connection Policy for use with VM-FEX (SR-IOV).

1. Select the LAN tab on the left of the window.
2. Go to **Policies** > **root**.
3. Right-click Dynamic vNIC Connection Policy and select create.
4. Enter the name VF-iSCSI-A
5. Enter 20 for the Dynamic vNIC value.

Note The number of Dynamic vNICs may need to be reduced depending on the adapter version and the number of uplinks between the IOM or FEX and the fabric interconnect. The service profile will fail to associate with the blade or rack server if there are too many Dynamic vNICs specified.

6. Select Windows adapter policy from the dropdown box.
7. Select Protected Perf A protection policy.
8. Click **OK** to create the dynamic vNIC policy for the iSCS-A virtual function.

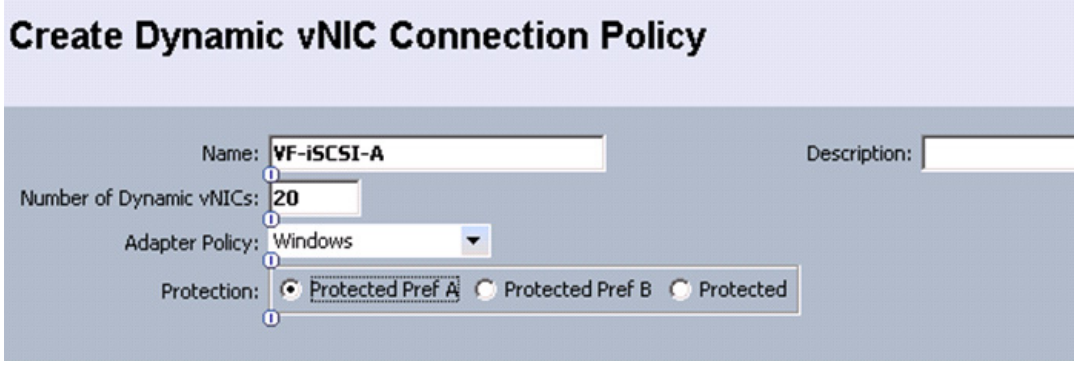

9. Right-click Dynamic vNIC Connection Policy and select create.
10. Enter the name VF-iSCSI-B
11. Enter 20 for the Dynamic vNIC value.

Note The number of Dynamic vNICs may need to be reduced depending on the adapter version and the number of uplinks between the IOM or FEX and the fabric interconnect. The service profile will fail to associate with the blade or rack server if there are too many Dynamic vNICs specified.

12. Select Windows adapter policy from the dropdown box.
13. Select Protected Perf B protection policy.
14. Click **OK** to create the dynamic vNIC policy for the iSCS-A virtual function.

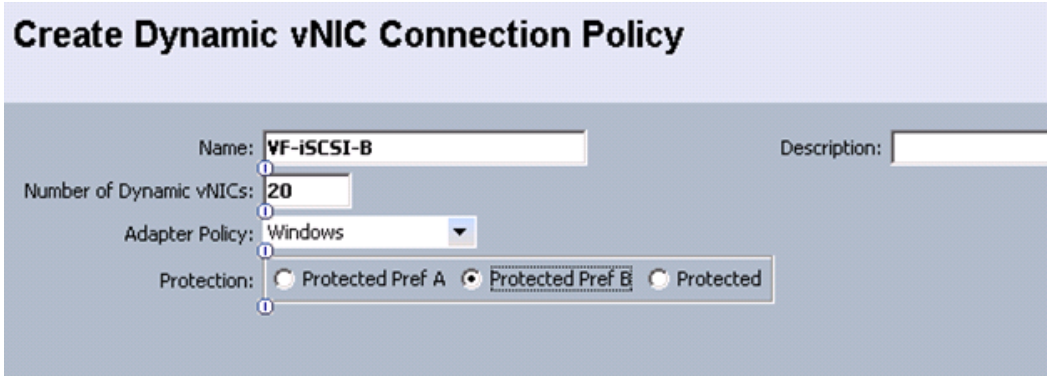

Create vNIC/HBA Placement Policy for Virtual Machine Infrastructure Hosts

1. Right-click vNIC/HBA Placement policy and select create.
2. Enter the name VM-Host-Infra.
3. Select 1 and select Assign Only.
4. Click **OK**.

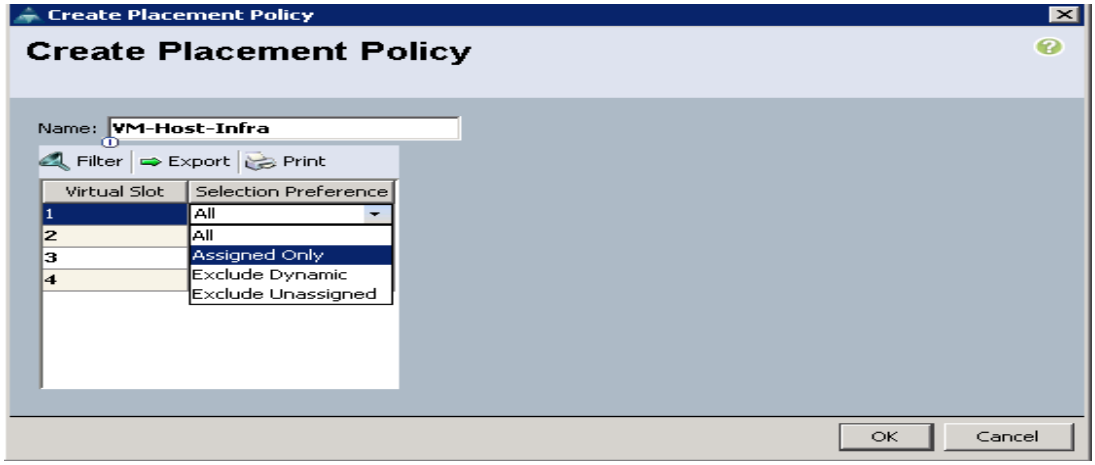

Create a vNIC Template

These steps provide details for creating multiple vNIC templates for the Cisco UCS environment.

1. Select the LAN tab on the left of the window.
2. Go to **Policies** > **root**.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter CSV as the vNIC template name.

6. Keep the **Fabric A** radio button selected. Check the Enable Failover checkbox. Under target, uncheck the VM checkbox. Select **Updating Template** radio button as the Template Type. Under VLANs, select CSV VLAN and set as Native VLAN. Under MTU, enter 9000. keep MAC Pool at default. Select QOS Policy as CSV.
7. Click **OK** to complete creating the vNIC template.
8. Click **OK**.

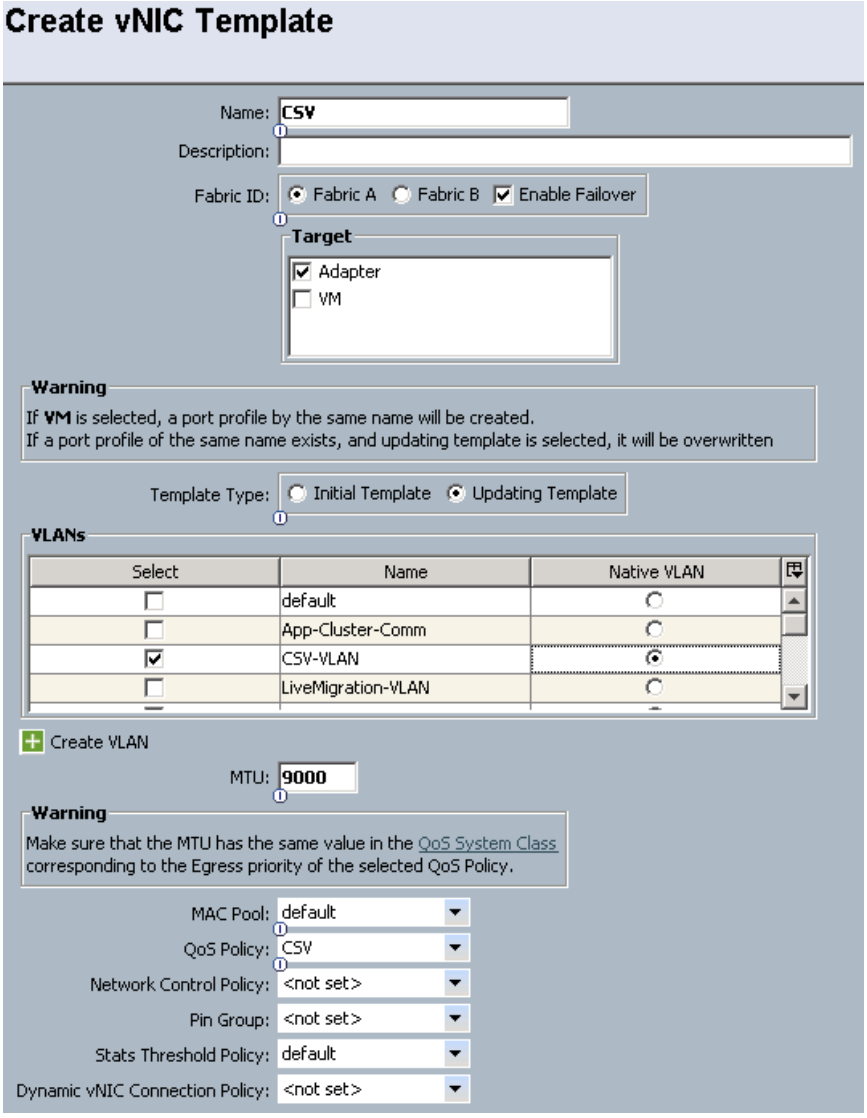

9. Select the LAN tab on the left of the window.
10. Go to **Policies** > **root**.
11. Right-click vNIC Templates.
12. Select Create vNIC Template.
13. Enter LiveMigration as the vNIC template name.

14. Select the **Fabric B** radio button. Check the Enable Failover box. Under target, uncheck the VM checkbox. Select **Updating Template** radio button for the Template Type. Under VLANs, select Live-Migration-VLAN and set as Native VLAN. For MTU, enter 9000. Keep MAC Pool as Default. For QoS Policy, select Live-Migration.
15. Click **OK** to complete creating the vNIC template.
16. Click **OK**.

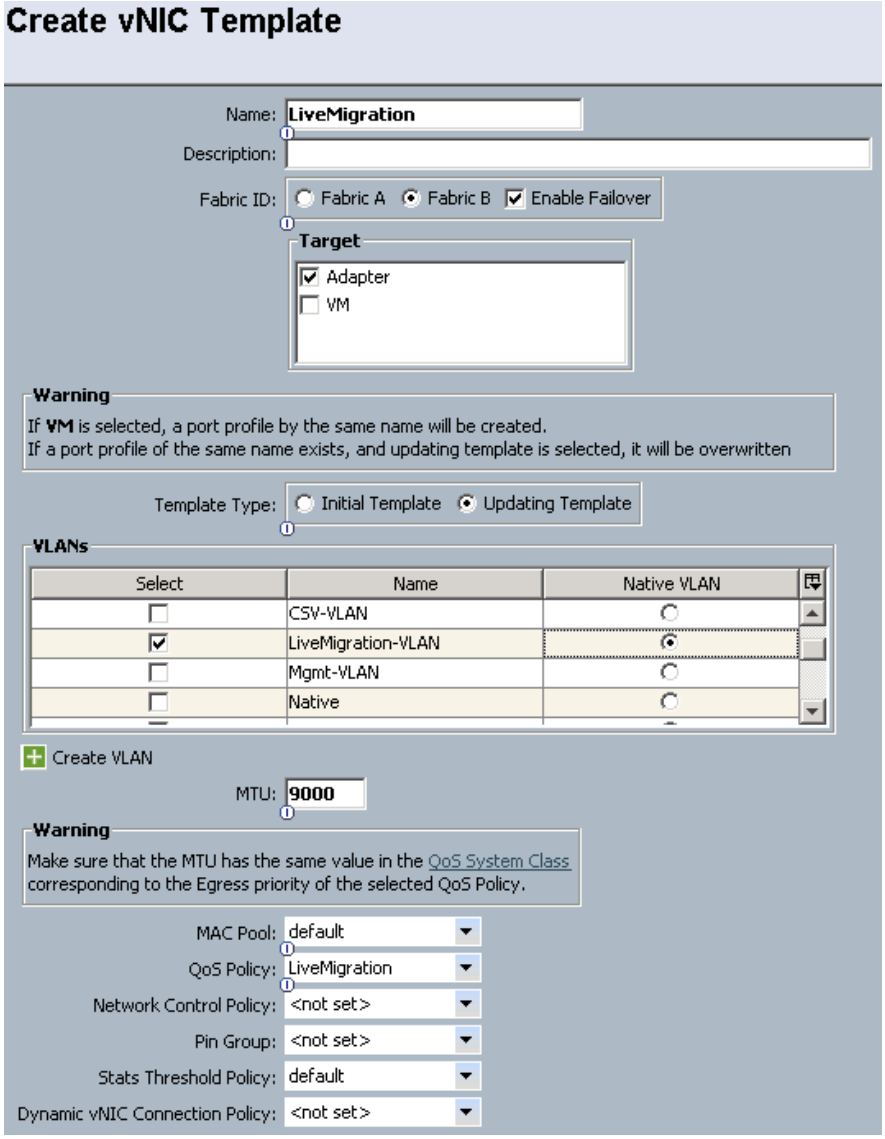

17. Select the LAN tab on the left of the window.
18. Go to **Policies** > **root**.
19. Right-click vNIC Templates.
20. Select Create vNIC Template.
21. Enter Mgmt as the vNIC template name.
22. Select the **Fabric A** radio button. Check the Enable Failover checkbox. Under target, uncheck the VM checkbox. Select **Updating Template** radio button for the Template Type. Under VLANs, select MGMT-VLAN. Set as Native VLAN. Select Default for MAC Pool.

23. Click **OK** to complete creating the vNIC template.
24. Click **OK**.

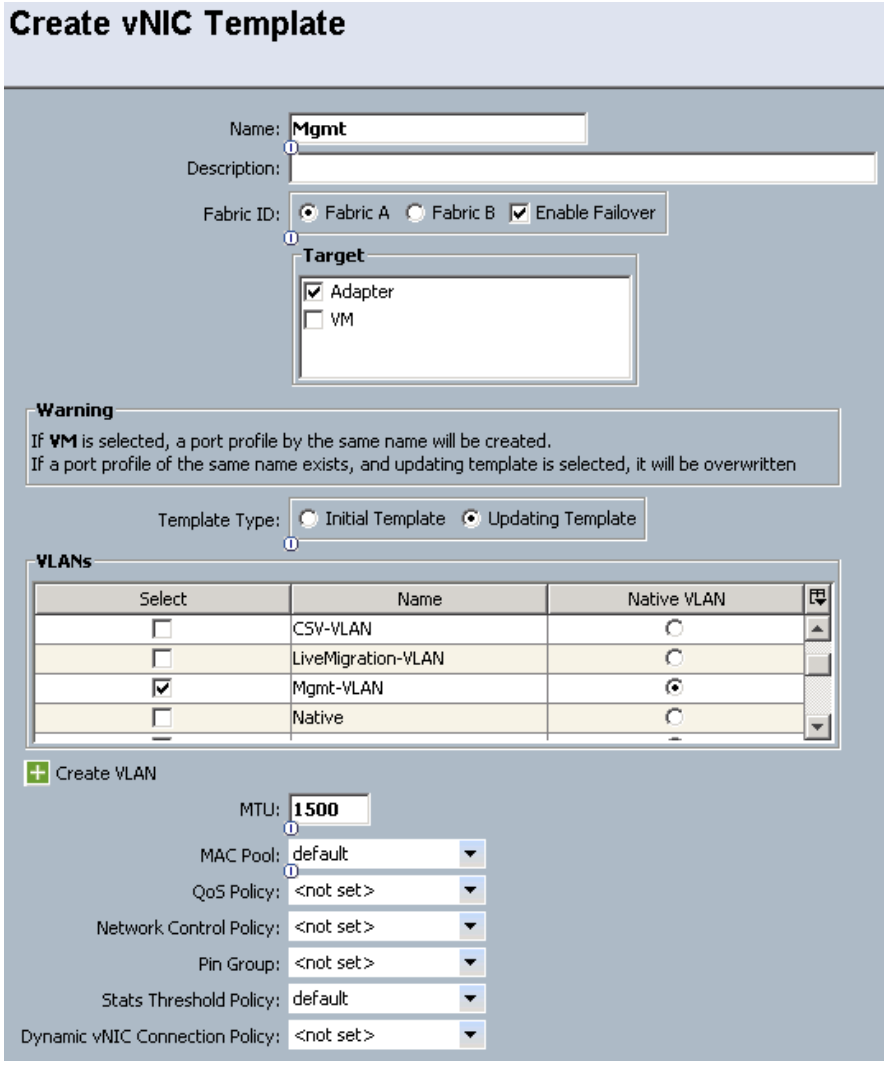

25. Select the LAN tab on the left of the window.
26. Go to **Policies** > **root**.
27. Right-click vNIC Templates.
28. Select Create vNIC Template.
29. Enter VM-Cluster-Comm as the vNIC template name.
30. Select the **Fabric B** radio button. Check the Enable Failover checkbox. Under target, uncheck the VM checkbox. Select **Updating Template** radio button for the Template Type. Under VLANs, select VM-Cluster-Comm. Do not set a Native VLAN. For MTU, enter 1500. Select Default for MAC Pool.
31. Click **OK** to complete creating the vNIC template.
32. Click **OK**.

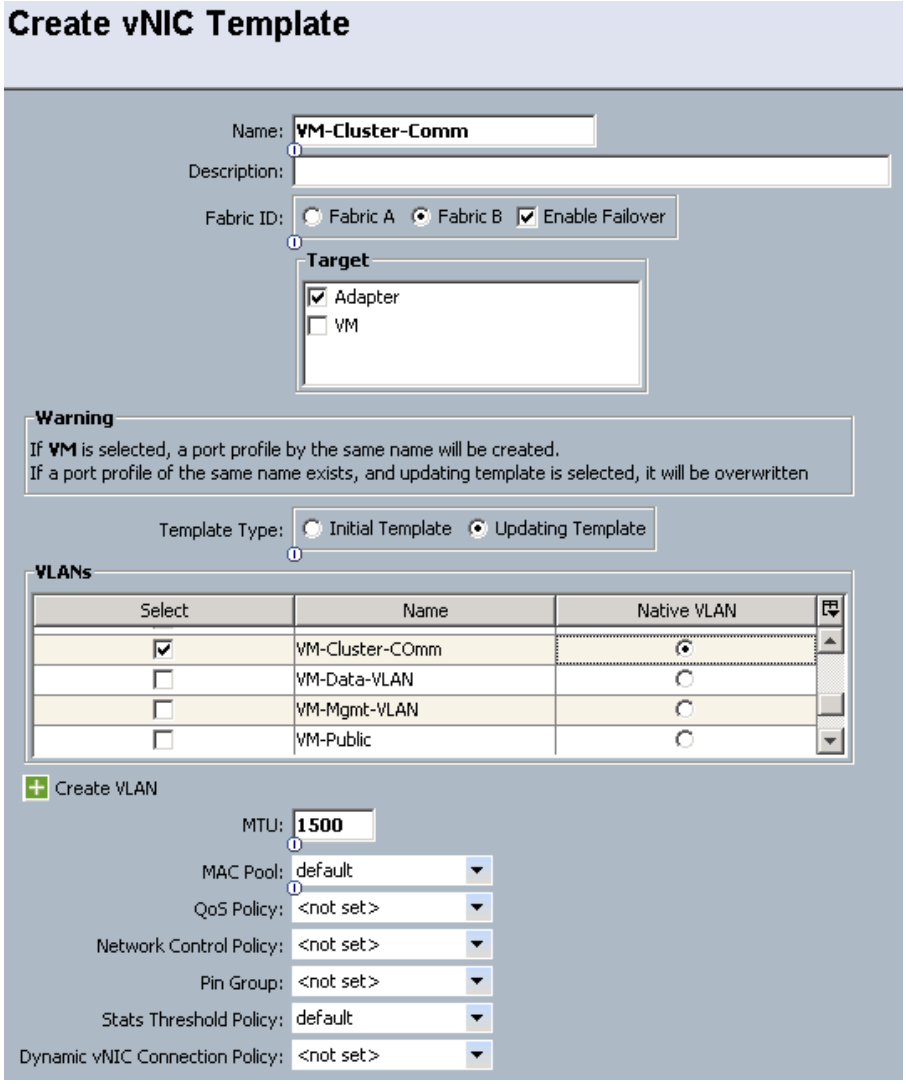

33. Select the LAN tab on the left of the window.
34. Go to **Policies** > **root**.
35. Right-click vNIC Templates.
36. Select Create vNIC Template.
37. Enter VM-Data as the vNIC template name.
38. Select the **Fabric A** radio button. Check the Enable Failover checkbox. Under target, uncheck the VM checkbox. Select **Updating Template** radio button for the Template Type. Under VLANs, select VM-Public. Do not set a Native VLAN. Select Default for MAC Pool.
39. Click **OK** to complete creating the vNIC template.
40. Click **OK**.

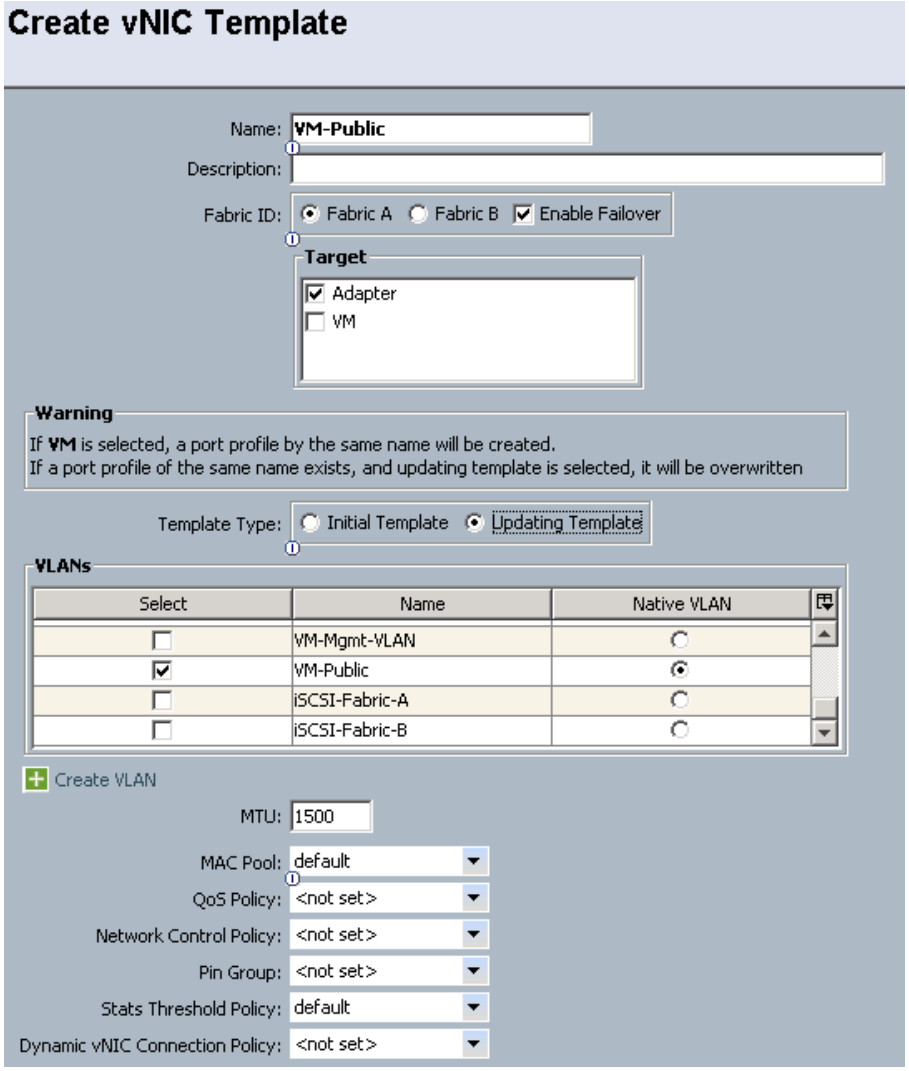

41. Select the LAN tab on the left of the window.
42. Go to **Policies** > **root**.
43. Right-click vNIC Templates.
44. Select Create vNIC Template.
45. Enter PF-iSCSI-A as the vNIC template name.
46. Select the **Fabric A** radio button. Uncheck the Enable Failover checkbox. Under target, check Adapter and VM checkboxes. Select **Updating Template** radio button as the Template Type. Under VLANs, select iSCSI-VLAN-A and set as Native VLAN. Under MTU, enter 9000. Under MAC Pool, select MAC_Pool. Under QoS Policy, select iSCSI. Under Dynamic vNIC Connection Policy, select VF-iSCSI-A.
47. Click **OK** to complete creating the vNIC template.
48. Click **OK**.

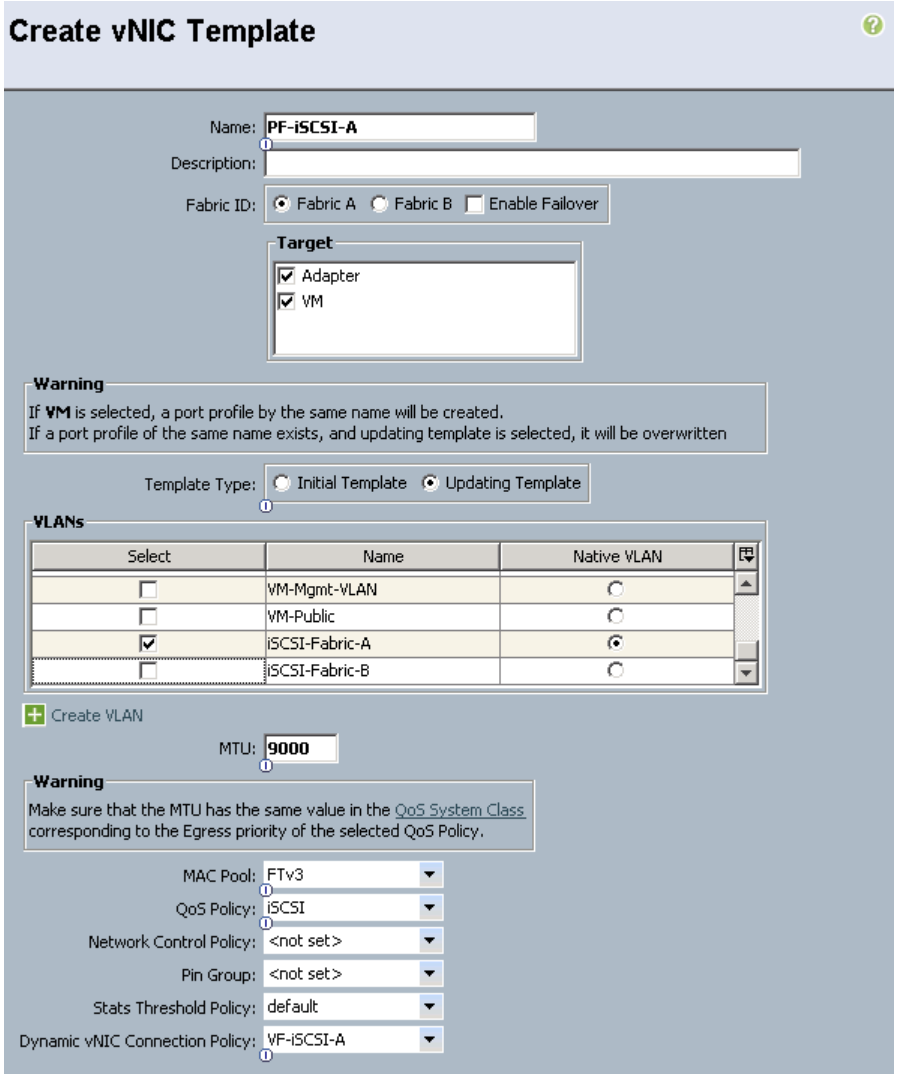

Note A Port profile is automatically created when creating a vNIC template and checking the VM check box

49. Select the LAN tab on the left of the window.
50. Go to **Policies > root or sub-organization**.
51. Right-click vNIC Templates.
52. Select Create vNIC Template.
53. Enter PF-iSCSI-B as the vNIC template name.
54. Select the **Fabric B** radio button. Uncheck the Enable Failover checkbox. Under target, check Adapter and VM checkboxes. Select **Updating Template** radio button for the Template Type. Under VLANs, select iSCSI-VLAN-B and set as Native VLAN. Under MTU, enter 9000. Under MAC Pool, select MAC_Pool. Under QoS Policy, select iSCSI. Under Dynamic vNIC Connection Policy, select VF-iSCSI-B. Click **OK** to complete creating the vNIC template.
55. Click **OK**.

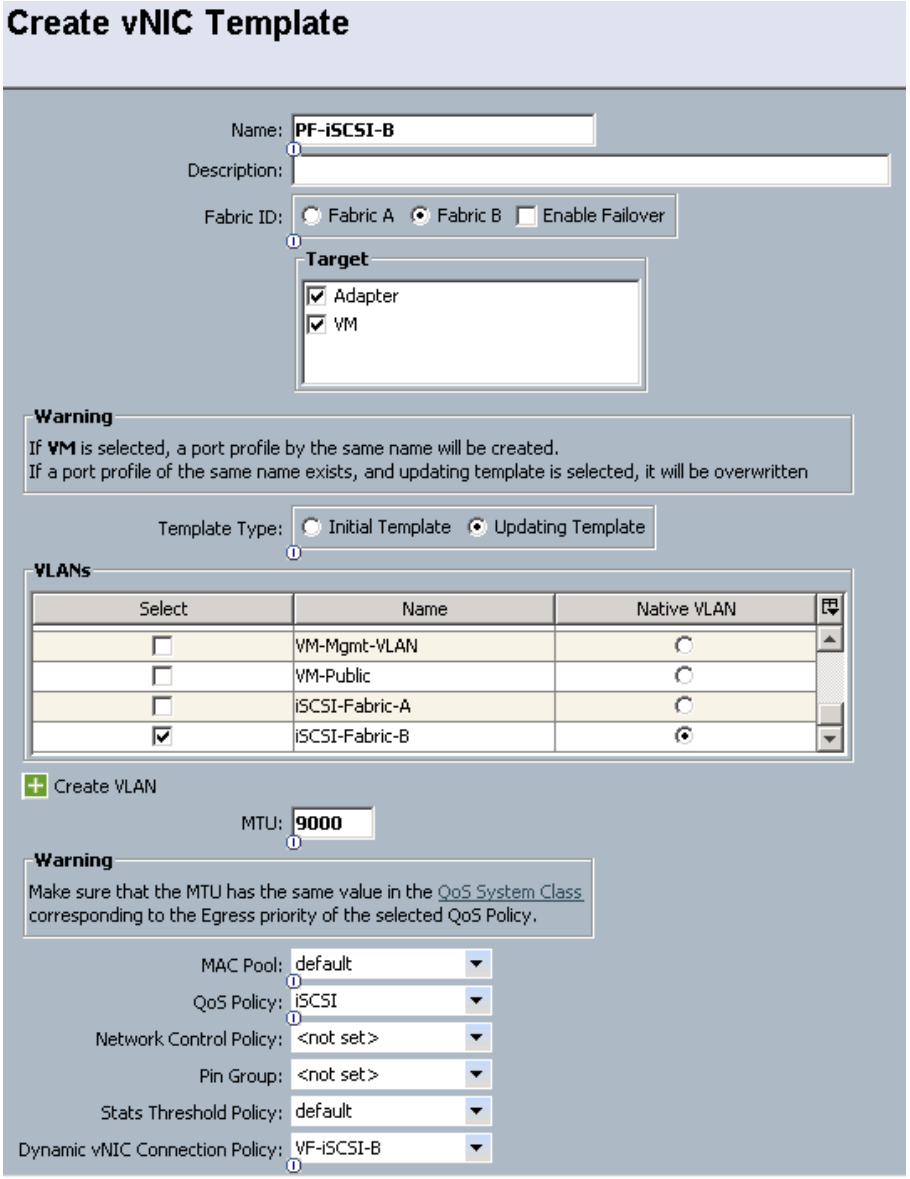

Note A Port profile is automatically created when creating a vNIC template and checking the VM check box

Create vHBA Templates for Fabric A and B

These steps provide details for creating multiple vHBA templates for the Cisco UCS environment.

1. Select the VSAN tab on the left of the window.
2. Go to **Policies** > **root**.
3. Right-click vHBA Templates.
4. Select Create vNIC Template.
5. Enter Fabric-A as the vHBA template name.

6. Select the Fabric A radio button. Under Select VSAN, select VSAN_A. Under WWN Pool, select default.
7. Click **OK** to complete creating the vHBA template.
8. Click **OK**.

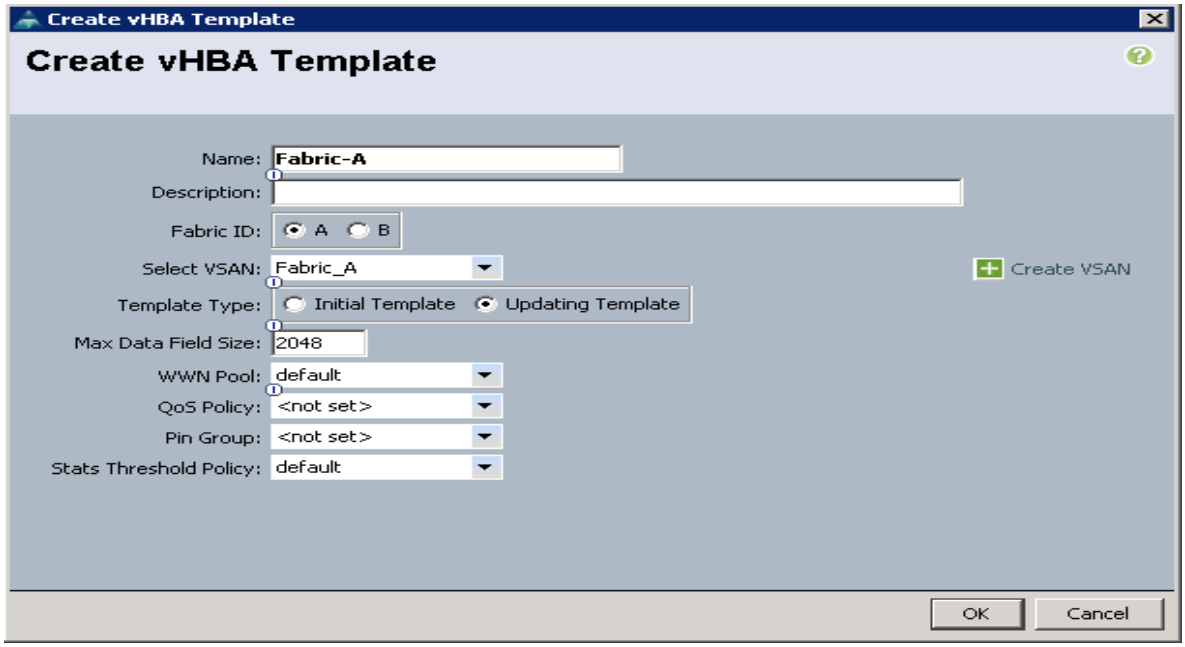

9. Select the VSAN tab on the left of the window.
10. Go to **Policies** > **root**.
11. Right-click vHBA Templates.
12. Select Create vHBA Template.
13. Enter Fabric-B as the vHBA template name.
14. Select the Fabric B radio button. Under Select VSAN, select VSAN_B. Under WWN Pool, select default.
15. Click **OK** to complete creating the vHBA template.
16. Click **OK**.

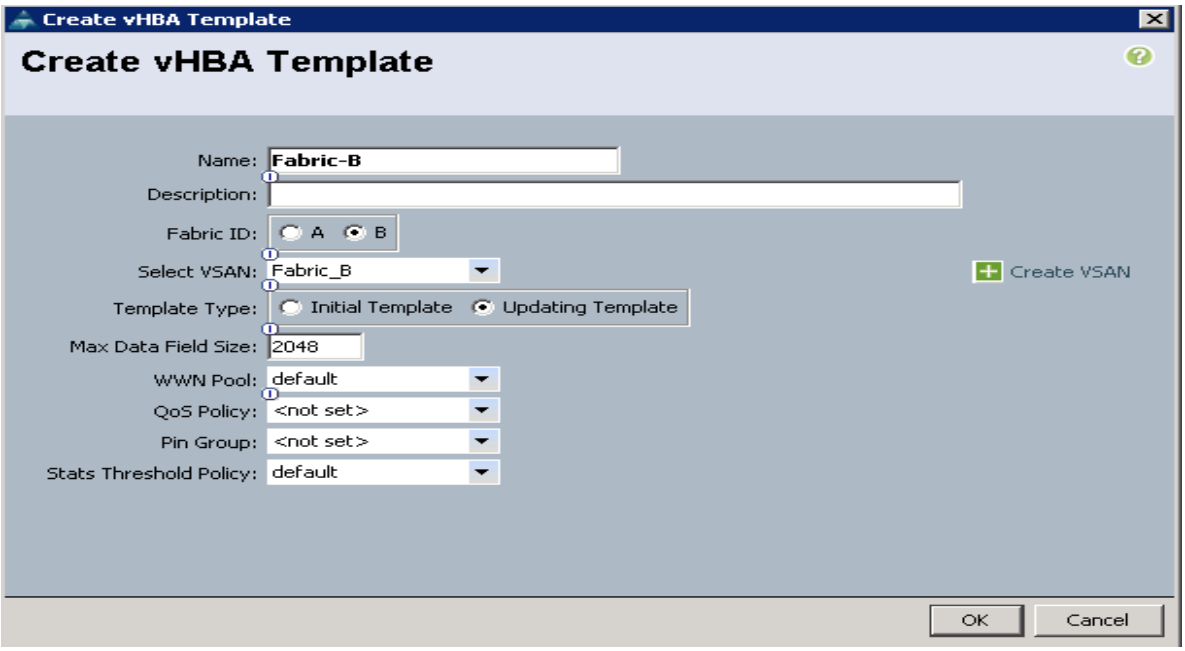

Create Boot Policies

These steps provide details for creating boot policies for the Cisco UCS environment. These directions apply to an environment in which each storage Controller Ba port is connected to fabric A and each storage Controller Bb port is connected to fabric B. In these steps, 2 boot policies will be configured. The first policy will configure the primary target to be controller A port 2a and the second boot policy primary target will be controller B port 2b.

1. Select the Servers tab at the top left of the window.
2. Go to **Policies** > **root**.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Name the boot policy Boot-FAS01-A.
6. (Optional) Give the boot policy a description.
7. Leave Reboot on Boot Order Change and Enforce vNIC/vHBA Name unchecked.
8. Expand the Local Devices drop-down menu and select Add CD-ROM.
9. Expand the vHBAs drop-down menu and select Add SAN Boot.
10. Enter Fabric-A in the vHBA field in the Add SAN Boot window that displays.
11. Make sure that Primary is selected as the type.
12. Click **OK** to add the SAN boot initiator.

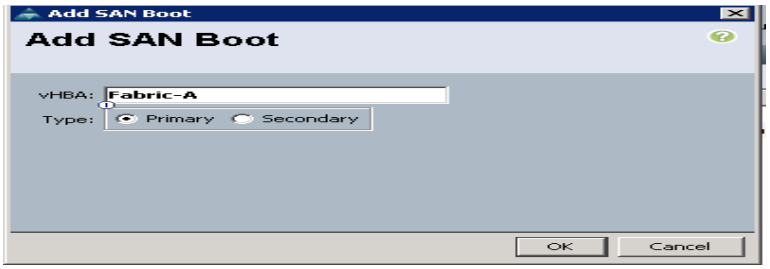

13. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
14. Enter the WWPN for the primary FC adapter interface 2a of controller A. To obtain this information, log in to controller A and run the fcp show adapters command.
15. Be sure to use the FC portname for 2a and not the FC node name.
16. Keep the type as Primary.
17. Click **OK** to add the SAN boot target.

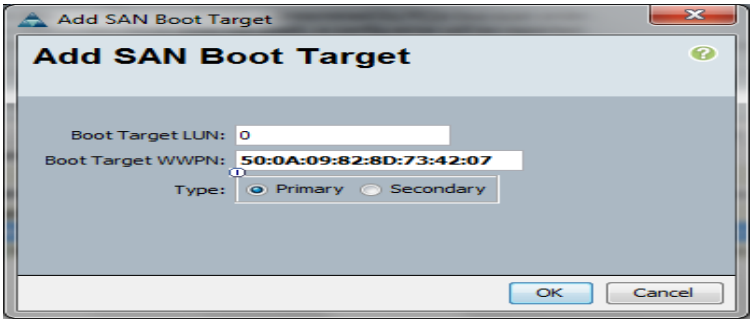

18. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
19. Enter the WWPN for the primary FC adapter interface 2a of controller B. To obtain this information, log in to the controller B and run the fcp show adapters command.
20. Be sure to use the FC portname for port 2a and not the FC node name.
21. Click **OK** to add the SAN boot target.

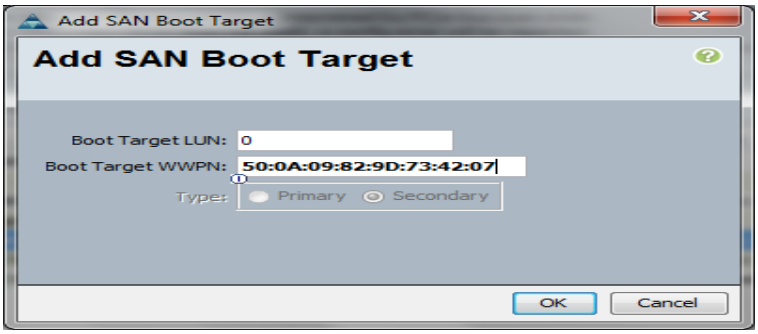

22. Select Add SAN Boot under the vHBA drop-down menu.
23. Enter Fabric-B in the vHBA field in the Add SAN Boot window that displays.
24. The type should automatically be set to Secondary and it should be grayed out.
25. Click **OK** to add the SAN boot target.

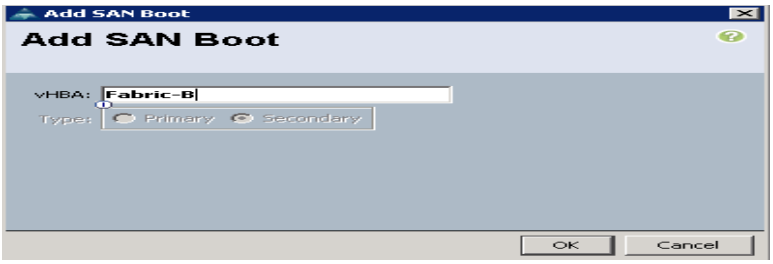

26. Select Add SAN Boot Target under the vHBA drop-down menu.
27. The Add SAN Boot Target window displays. Keep the value for Boot Target LUN as 0.
28. Enter the WWPN for the primary FC adapter interface 2b of the controller B. To obtain this information, log in to controller B and run the fcp show adapters command.
29. Be sure to use the FC portname for port 2b and not the FC node name.
30. Keep the type as Primary.
31. Click **OK** to add the SAN boot target.

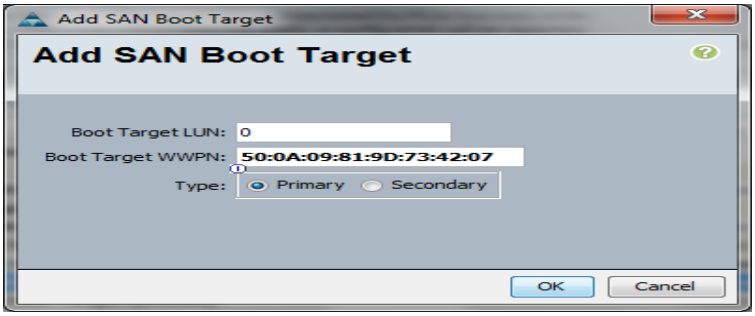

32. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
33. Enter the WWPN for the primary FCoE adapter interface 2b of controller A. To obtain this information, log in to controller A and run the fcp show adapters command.
34. Be sure to use the FC portname for port 2b and not the FC node name.
35. Click **OK** to add the SAN boot target.

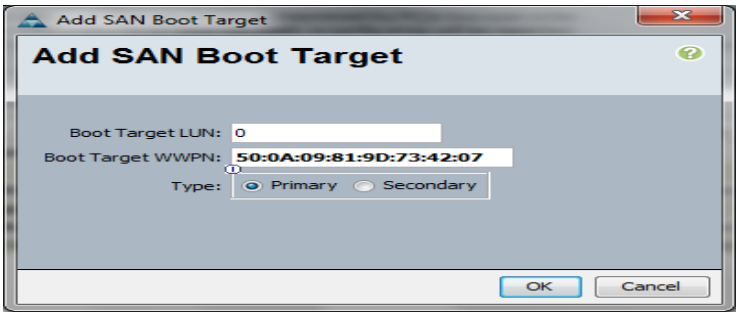

36. Right-click Boot Policies again.
37. Select Create Boot Policy.
38. Name the boot policy Boot-FAS01-B.
39. (Optional) Give the boot policy a description.
40. Leave Reboot on Boot Order Change and Enforce vNIC/vHBA Name unchecked.

41. Expand the Local Devices drop-down menu and select Add CD-ROM.
42. Click the vHBA drop-down menu and select Add SAN Boot.
43. Enter Fabric-B in the vHBA field in the Add SAN Boot window that displays.
44. Make sure that Primary is selected as the type.
45. Click **OK** to add the SAN boot target.

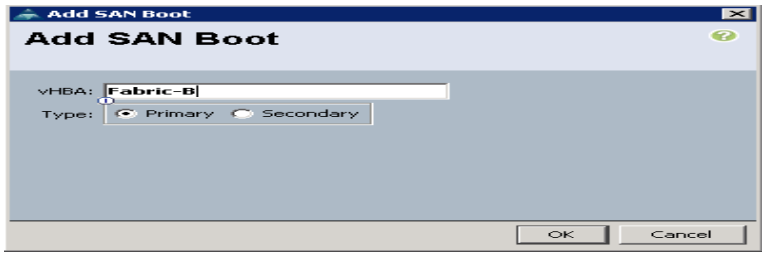

46. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
47. Enter the WWPN for the primary FCoE adapter interface 2b of controller B. To obtain this information, log in to controller B and run the fcp show adapters command.
48. Be sure to use the FC portname for port 2b and not the FC node name.
49. Keep the type as Primary.
50. Click **OK** to add the SAN boot target.

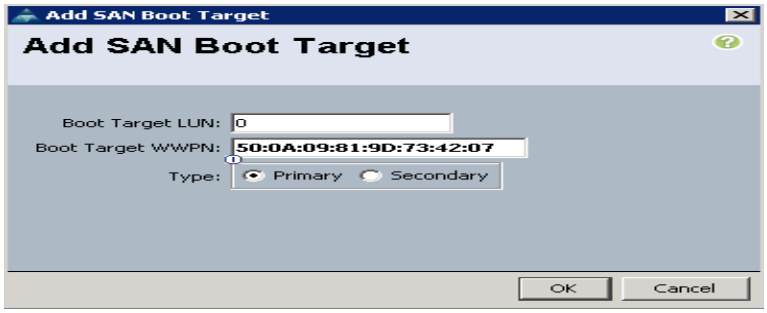

51. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
52. Enter the WWPN for the primary FC adapter interface 2b of controller A. To obtain this information, log in to controller A and run the fcp show adapters command.
53. Be sure to use the FC portname for port 2b and not the FC node name.
54. Click **OK** to add the SAN boot target.

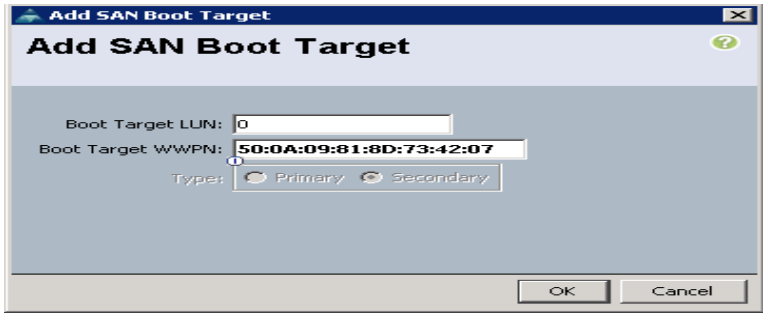

55. Select Add SAN Boot under the vHBA drop-down menu.
56. Enter Fabric-A in the vHBA field in the Add SAN Boot window that displays.
57. The type should automatically be set to Secondary and it should be grayed out.
58. Click **OK** to add the SAN boot target.

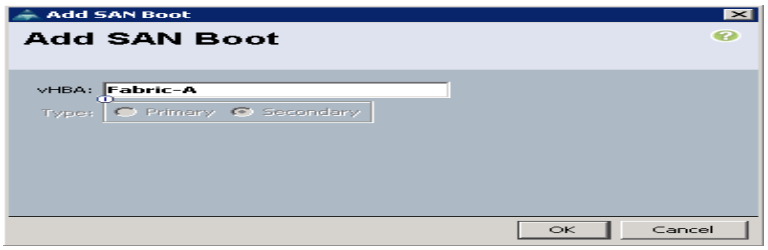

59. Select Add SAN Boot Target under the vHBA drop-down menu.
60. The Add SAN Boot Target window displays. Keep the value for Boot Target LUN as 0.
61. Enter the WWPN for the primary FC adapter interface 2a of controller A. To obtain this information, log in to controller A and run the fcp show adapters command.
62. Be sure to use the FC portname for port 2a and not the FC node name.
63. Keep the type as Primary.
64. Click **OK** to add the SAN boot target.

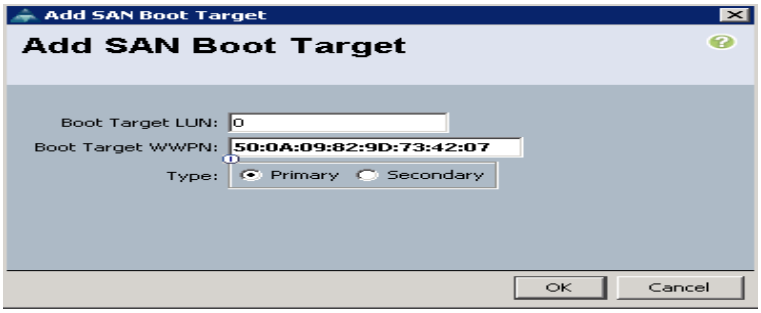

65. Under the vHBA drop-down menu, select Add SAN Boot Target. Keep the value for Boot Target LUN as 0.
66. Enter the WWPN for the primary FCoE adapter interface 2a of controller B. To obtain this information, log in to controller B and run the fcp show adapters command.
67. Be sure to use the FC portname for port 2a and not the FC node name.
68. Click **OK** to add the SAN boot target.

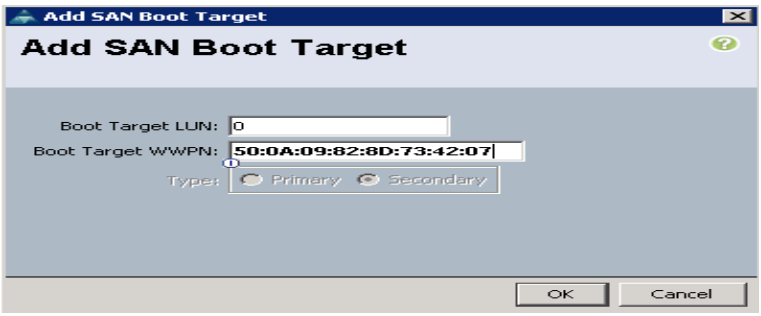

69. Click **OK** to create the boot policy in the Create Boot Policy pop-up window.

Create Service Profile Templates

This section details the creation of two service profile templates: one for fabric A and one for fabric B. The first profile is created and then cloned and modified for the second host.

1. Select the Servers tab at the top left of the window.
2. Go to **Service Profile Templates** > **root or sub-organization**.
3. Right-click root or sub-organization.
4. Select Create Service Profile Template.
5. The Create Service Profile Template window appears.

Identify the Service Profile Template

These steps detail configuration info for the Identify the Service Profile Template Section.

1. Name the service profile template VM-Host-Infra-Fabric-A. This service profile template is configured to boot from controller a on fabric A.
2. Select **Updating Template** radio button.
3. In the UUID section, select UUID_Pool as the UUID pool.
4. Click **Next** to continue to the next section.

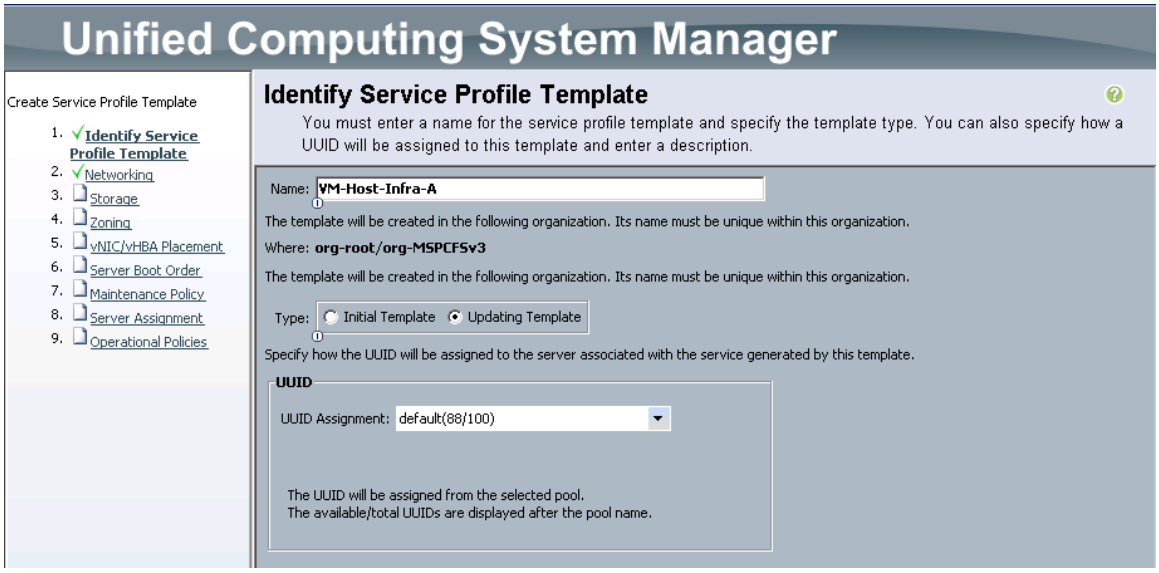

Networking Section

Leave the Dynamic vNIC Connection Policy field at the default.

1. Select **Expert** radio button for How would you like to configure LAN connectivity? option.

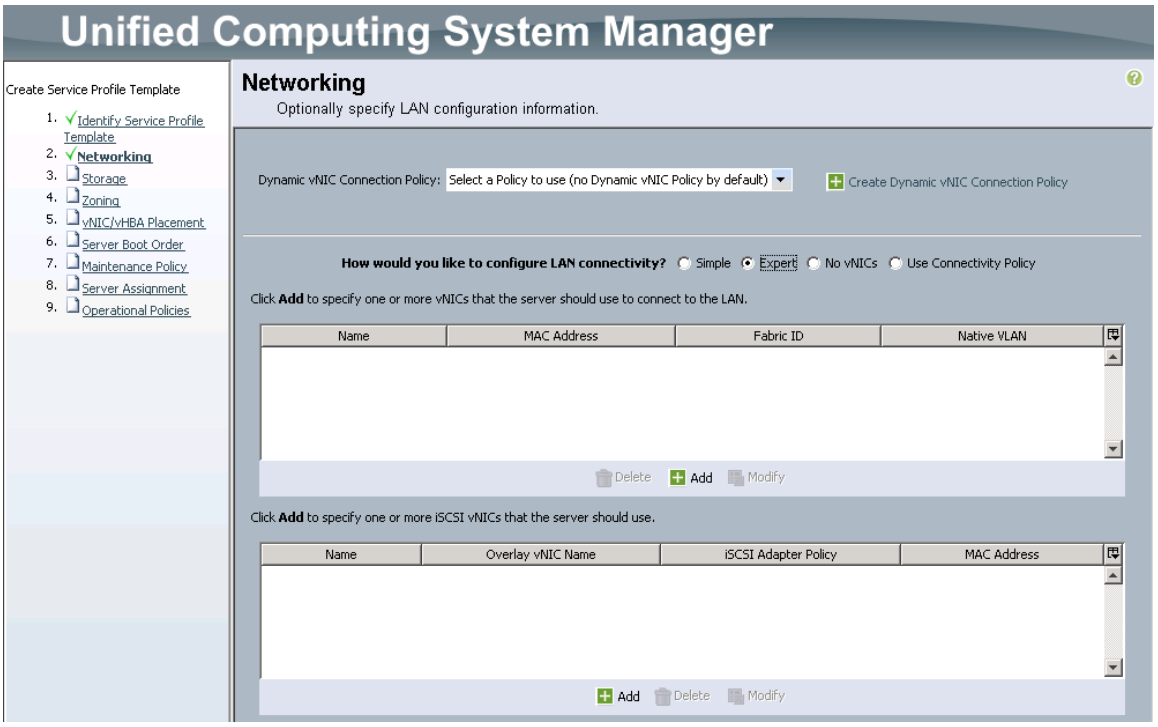

2. Click **Add** to add a vNIC to the template.
3. The Create vNIC window displays. Name the vNIC CSV.
4. Check the Use vNIC Template checkbox.
5. Select CSV for the vNIC Template field.
6. Select Windows in the Adapter Policy field.
7. Click **OK** to add the vNIC to the template.

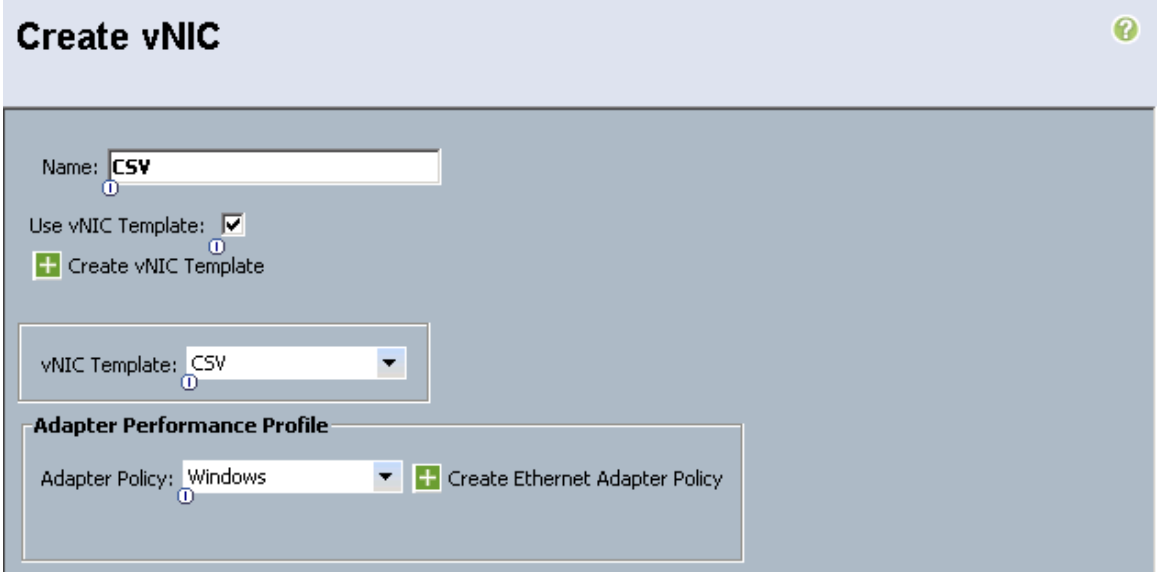

8. Click **Add** to add a vNIC to the template.

9. The Create vNIC window displays. Name the vNIC LiveMigration.
10. Check the Use LAN Connectivity Template checkbox.
11. Select LiveMigration for the vNIC Template field.
12. Select Windows in the Adapter Policy field.
13. Click **OK** to add the vNIC to the template.

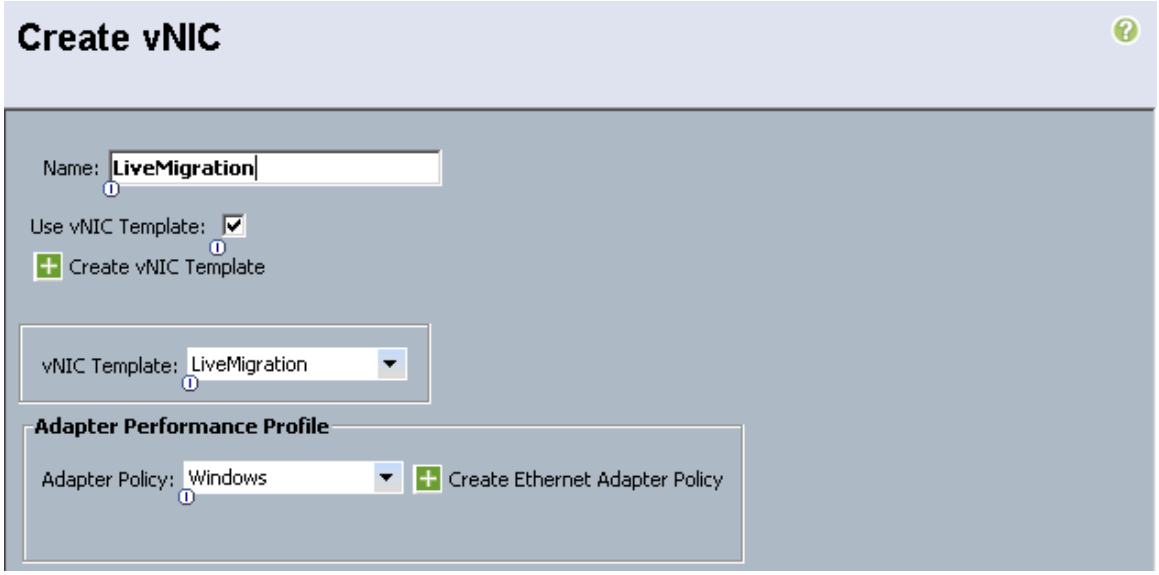

14. Click **Add** to add a vNIC to the template.
15. The Create vNIC window displays. Name the vNIC Mgmt.
16. Check the Use LAN Connectivity Template checkbox.
17. Select Mgmt for the vNIC Template field.
18. Select Windows in the Adapter Policy field.
19. Click **OK** to add the vNIC to the template.

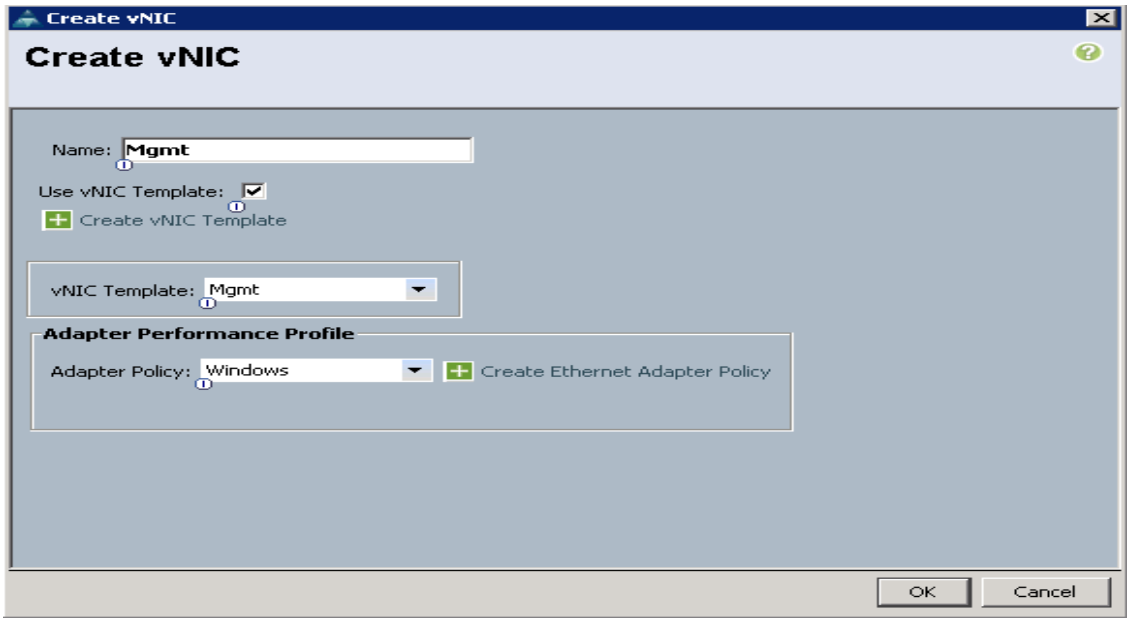

20. Click **Add** to add a vNIC to the template.
21. The Create vNIC window displays. Name the vNIC VM-Cluster-Comm.
22. Check the Use LAN Connectivity Template checkbox.
23. Select VM-Cluster-Comm for the vNIC Template field.
24. Select Windows in the Adapter Policy field.
25. Click **OK** to add the vNIC to the template.

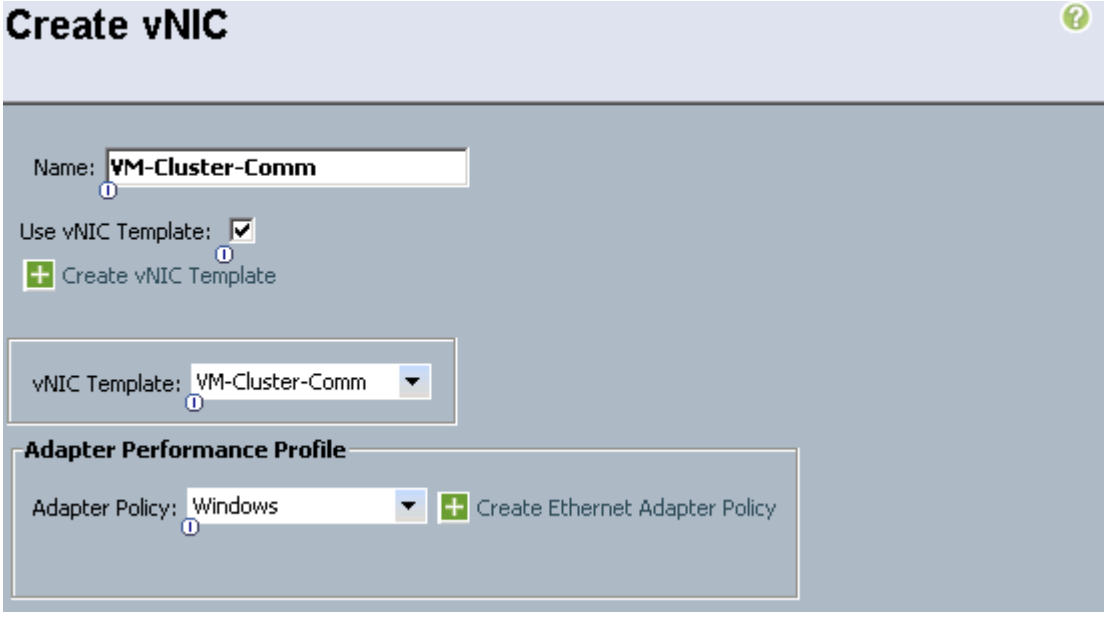

26. Click **Add** to add a vNIC to the template.
27. The Create vNIC window displays. Name the vNIC VM-Public.

28. Check the Use LAN Connectivity Template checkbox.
29. Select VM-Public for the vNIC Template field.
30. Select Windows in the Adapter Policy field.
31. Click **OK** to add the vNIC to the template.

32. Click **Add** to add a vNIC to the template.
33. The Create vNIC window displays. Name the vNIC PF-iSCSI-A.
34. Check the Use LAN Connectivity Template checkbox.
35. Select PF-iSCSI-A for the vNIC Template field.
36. Select SRIOV in the Adapter Policy field.
37. Click **OK** to add the vNIC to the template.

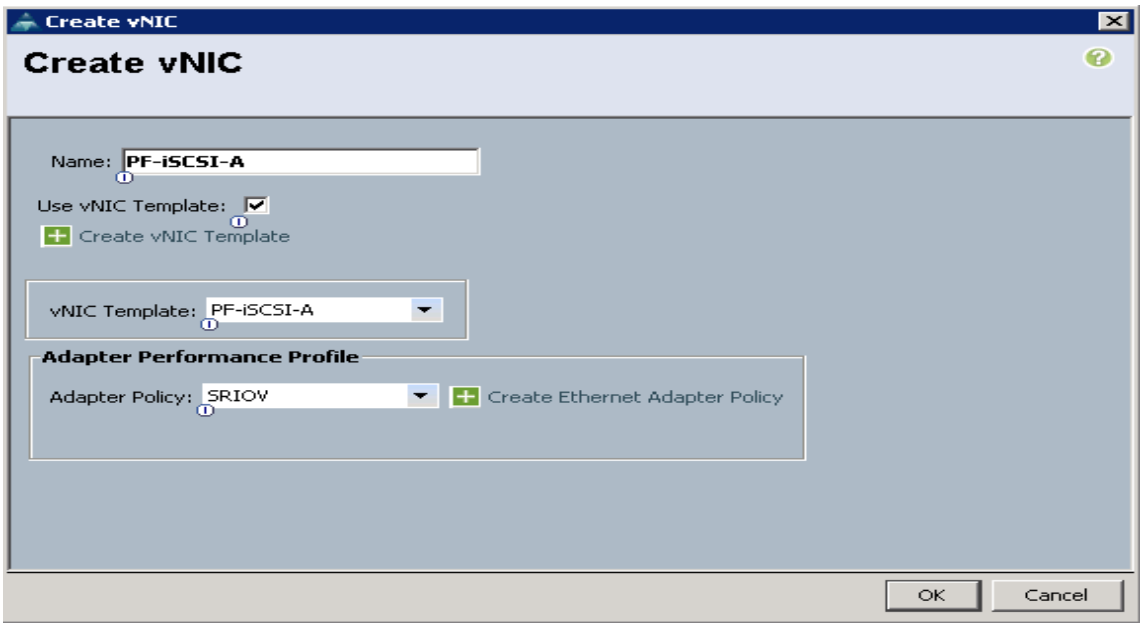

38. Click **Add** to add a vNIC to the template.
39. Check the Use LAN Connectivity Template checkbox.
40. Select PF-iSCSI-B for the vNIC Template field.
41. Select SRIOV in the Adapter Policy field.
42. Click **OK** to add the vNIC to the template.

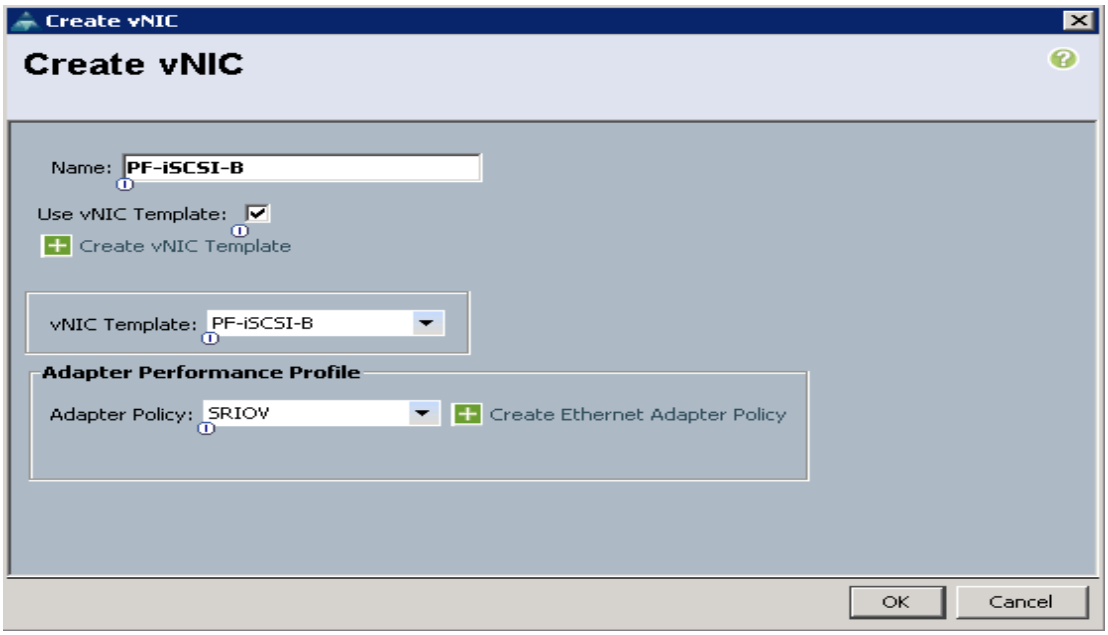

43. Verify: Review the table to make sure that all of the vNICs were created.
44. Click **Next** to continue to the next section.

Storage Section

1. Select Default for the Local Storage field.
2. Select the appropriate local storage policy if the server in question does not have local disk.
3. Select SAN-Boot for the local disk configuration policy.
4. Select the **Expert** radio button for How would you like to configure SAN connectivity option.
5. In the WWNN Assignment field, select WWNN_Pool.
6. Click **Add,** at the bottom of the window to add vHBAs to the template.
7. The Create vHBA window displays. Name the vHBA Fabric-A.
8. Check the Use SAN Connectivity Template checkbox.
9. Select Fabric-A in the vHBA Template field.
10. Select Windows-NetApp in the Adapter Policy field.
11. Click **OK** to add the vHBA to the template.

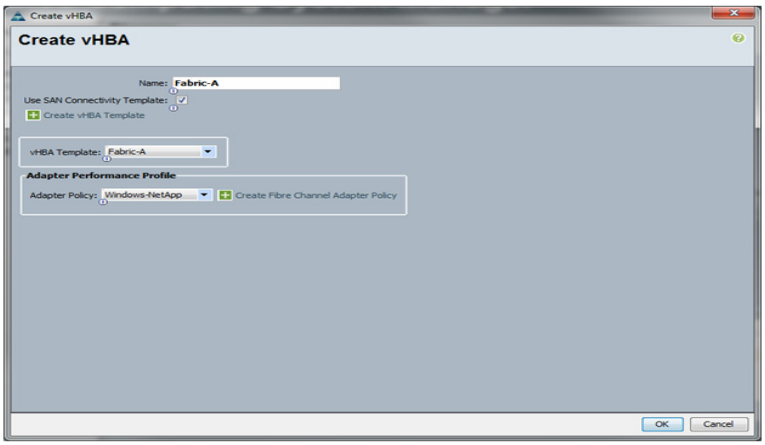

12. Click **Add,** at the bottom of the window to add vHBAs to the template.
13. The Create vHBA window displays. Name the vHBA Fabric-B.
14. Check the Use SAN Connectivity Template checkbox.
15. Select Fabric-B in the vHBA Template field.
16. Select Windows-NetApp in the Adapter Policy field.
17. Click **OK** to add the vHBA to the template.

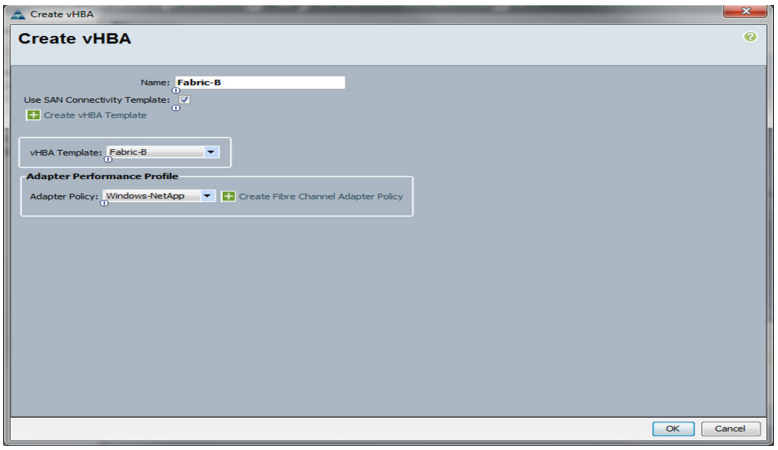

18. Verify – Review the table to make sure that both of the vHBAs were created.

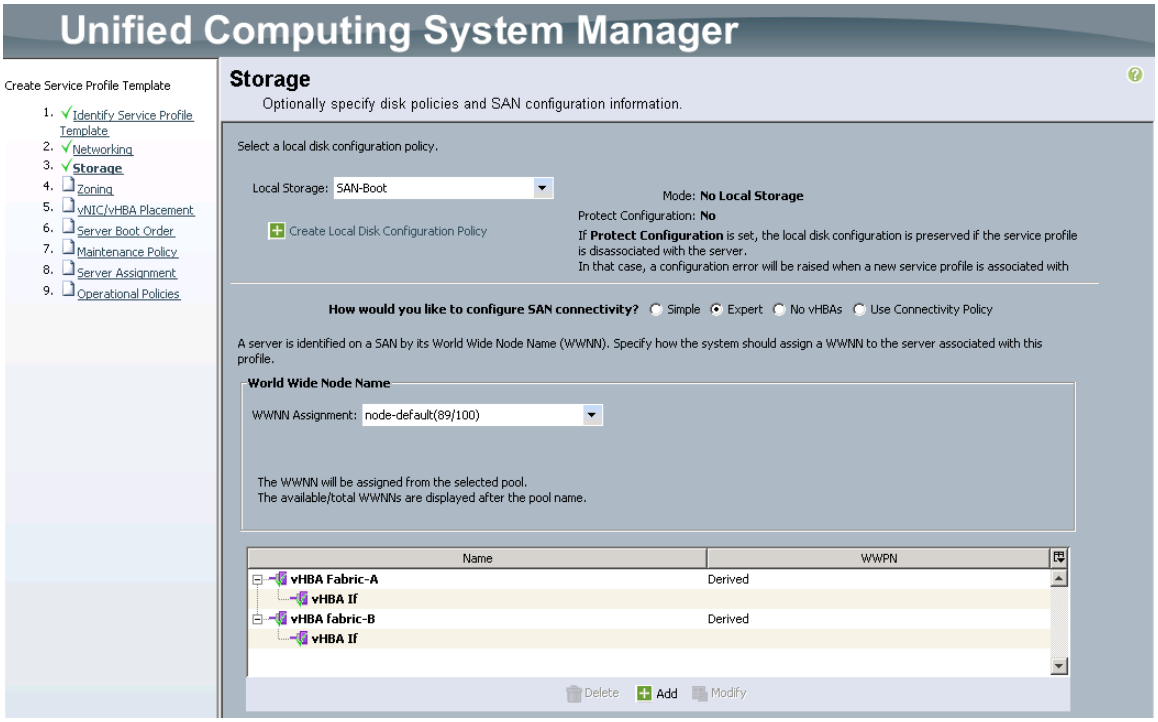

19. Click **Next** to continue to the next section.

Zoning Section

Note Zoning configuration in this section is not required because the fabric interconnects are in End-Host mode and zoning is configured on the Cisco Nexus 5548 switches.

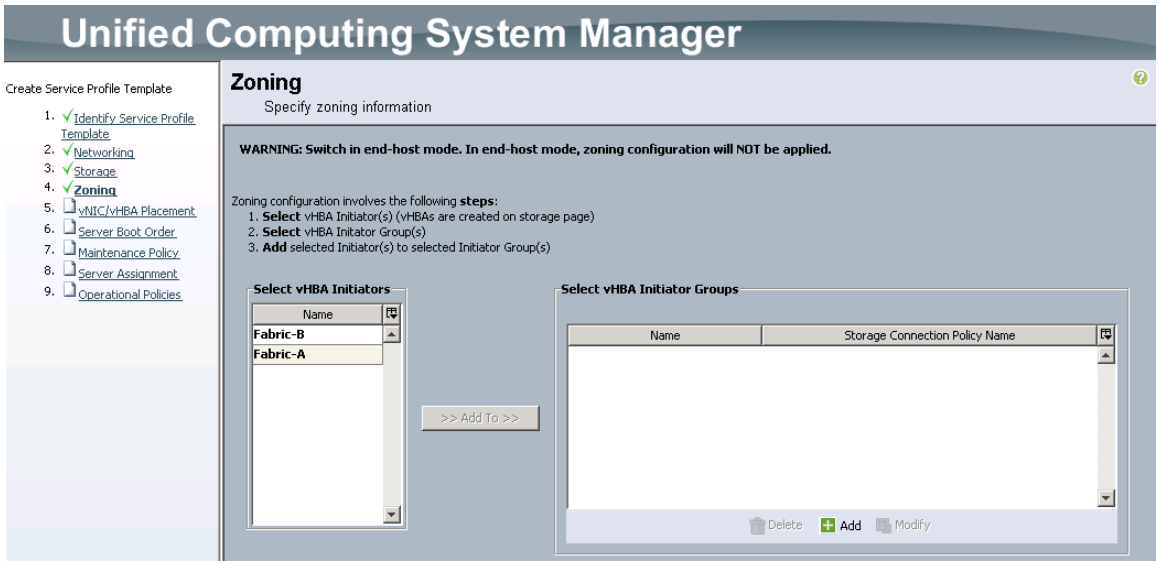

1. Click **Next** to continue the next section.

vNIC/vHBA Placement Section

Select the VM-Host-Infra Placement Policy in the Select Placement field.

1. Select vCon1 assign the vNICs in the following order:
   - Mgmt
   - CSV
   - LiveMigration
   - VM-Public
   - VM-Cluster-Comm
   - PF-iSCSI-A
   - PF-iSCSI-B
2. Click the vHBA tab and add the vHBAs in the following order:
   - Fabric-A
   - Fabric-B
3. Verify: Review the table to make sure that all of the vHBAs and vNICs were created. The order of the vNICs and vHBAs is not important.
4. Click **Next** to continue to the next section.

Server Boot Order Section

1. Select Boot-FAS01-A in the Boot Policy field.
2. Verify: Review the table to make sure that all of the boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.
3. Click **Next** to continue to the next section.

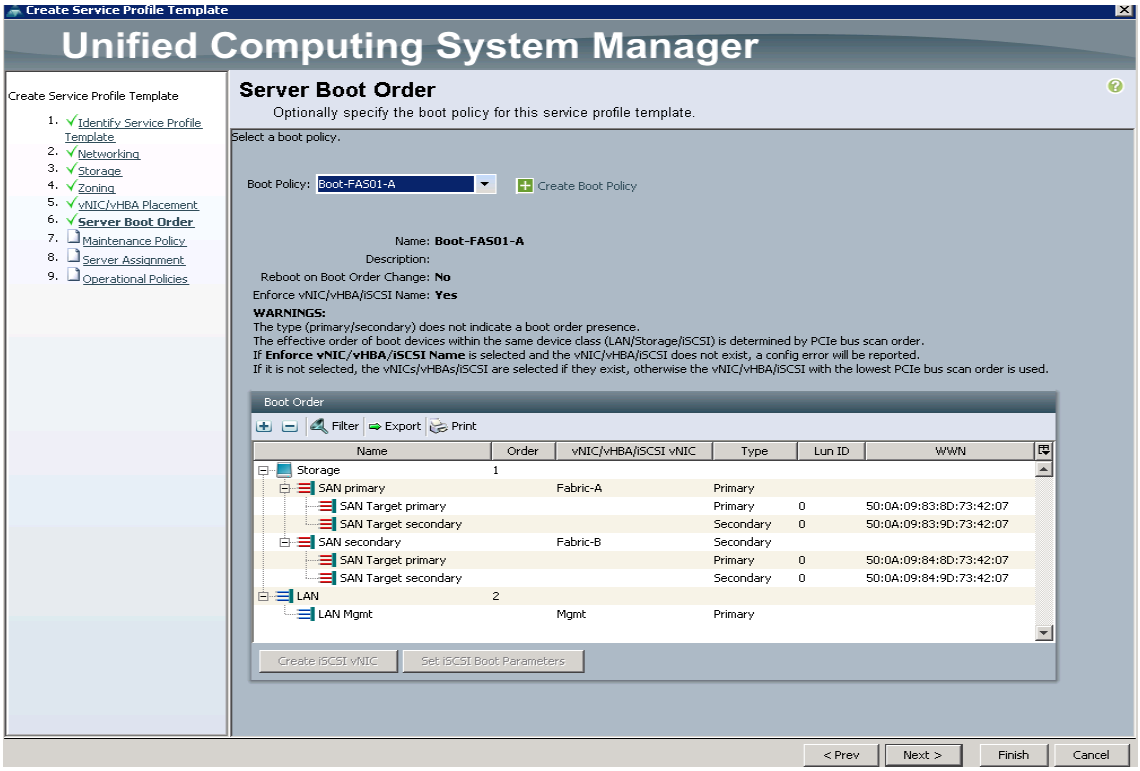

Maintenance Policy Section

1. Select the previously created policy User_Acknowledge.
2. Click **Next** to continue to the next section.

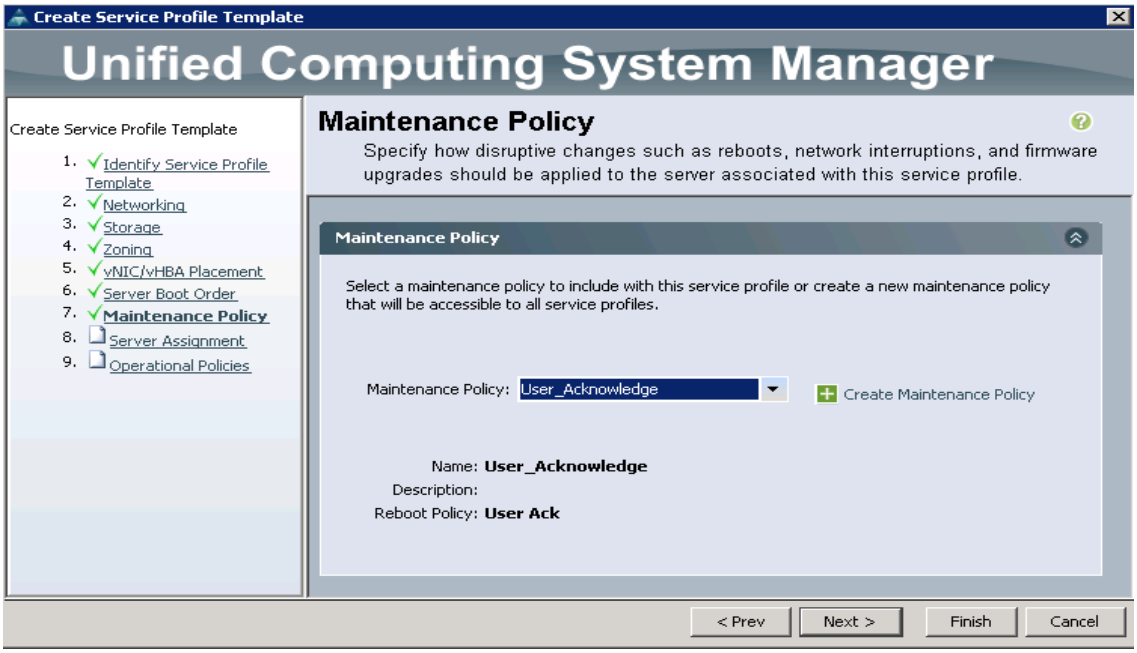

Server Assignment Section

1. Select VM-Host-Infra in the Pool Assignment field.
2. Select VM-Host-Infra for the Server Pool Qualification field.
3. Select Up for the power state.
4. Select VM-Host-Infra in the Host Firmware field.
5. Click **Next** to continue to the next section.

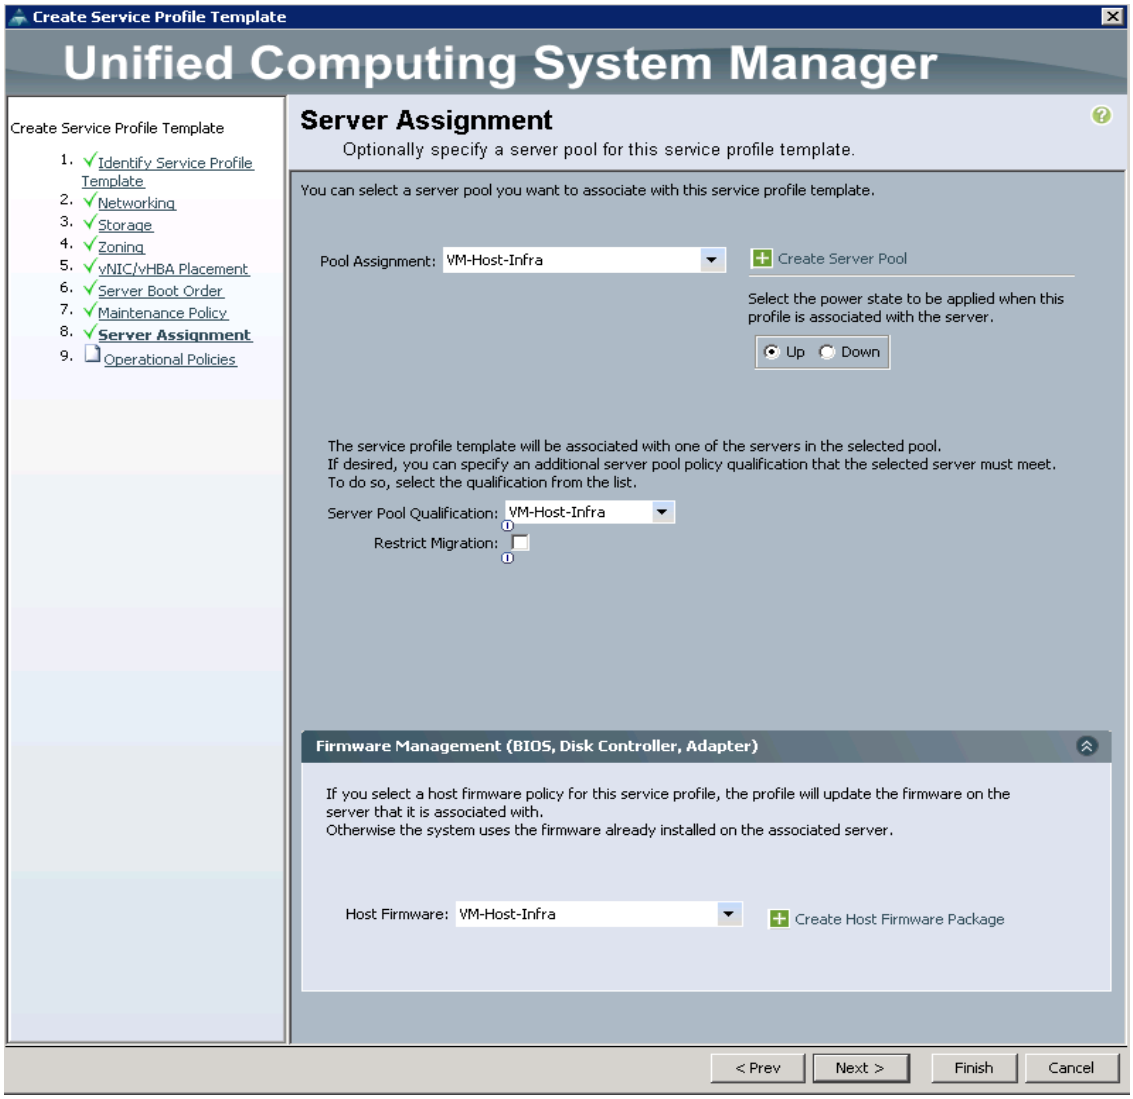

Operational Policies Section

1. Select VM-Host-Infra in the BIOS Policy field.
2. Expand Power Control Policy Configuration.
3. Select No-Power-Cap in the Power Control Policy field.
4. Click **Finish** to create the Service Profile template.

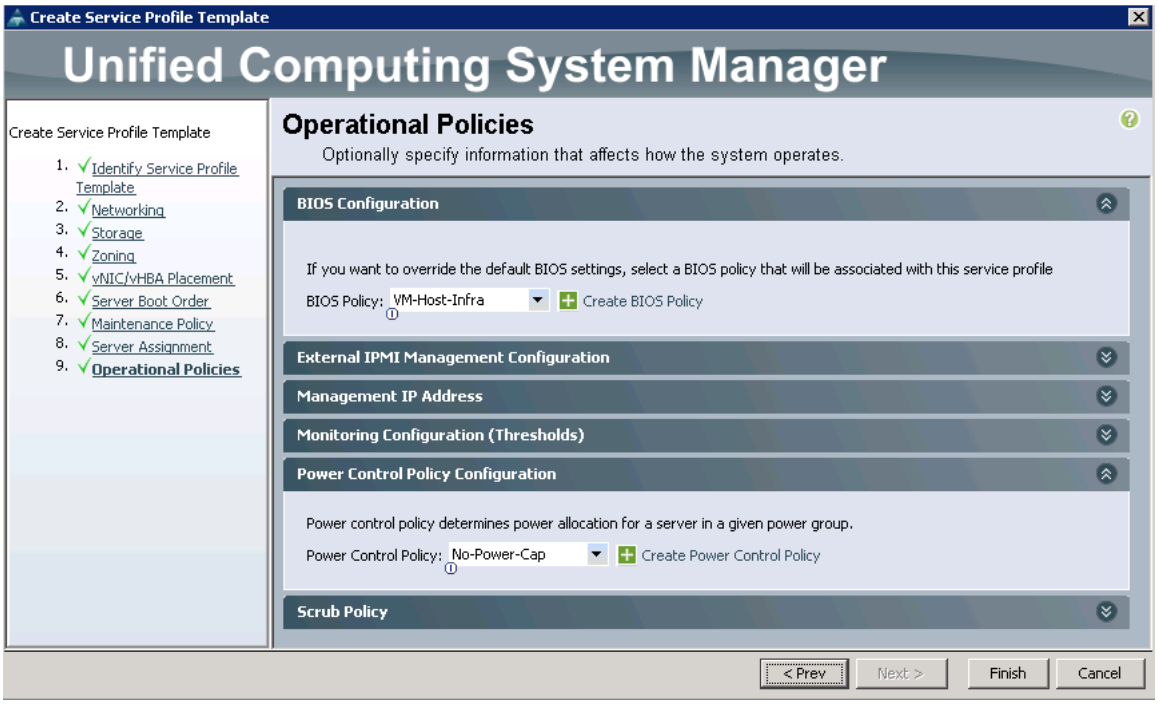

5. Click **OK** in the pop-up window to proceed.

Create the Fabric-B Template.

1. Select the Servers tab at the top left of the window.
2. Go to **Service Profile Templates > root or a sub-organization**.
3. Select the previously created VM-Host-Infra-Fabric-A template
4. Click **Create a Clone**.
5. Enter VM-Host-Infra-Fabric-B in the Clone Name field and click **OK**.

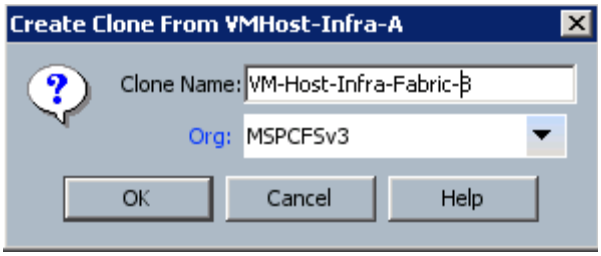

6. Select the newly created service profile template and select the Boot Order tab.
7. Click **Modify Boot Policy**.

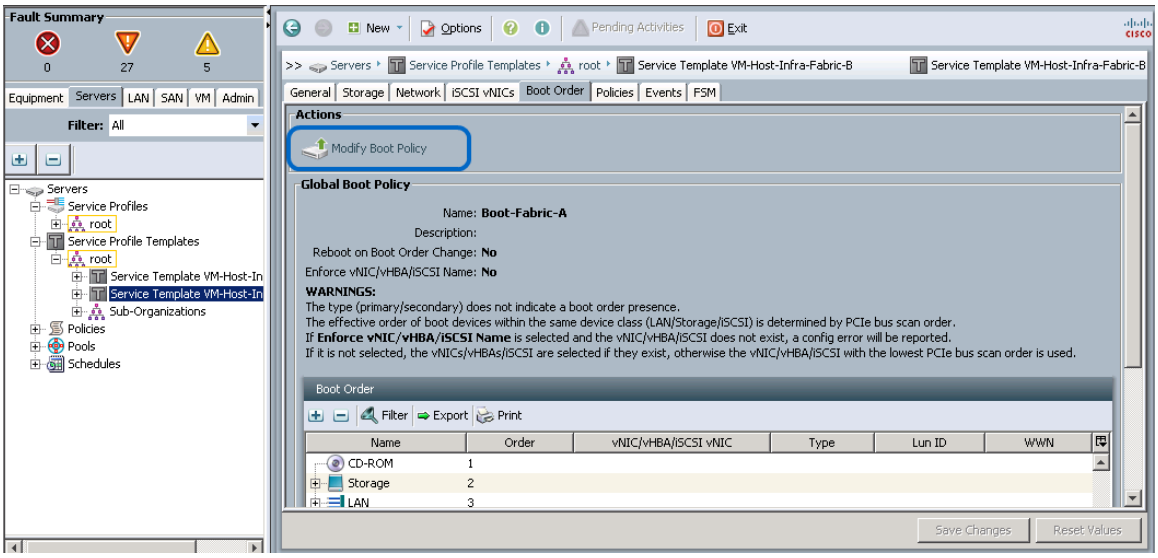

8. Select Boot-FAS01-B Boot Policy and click **OK**.

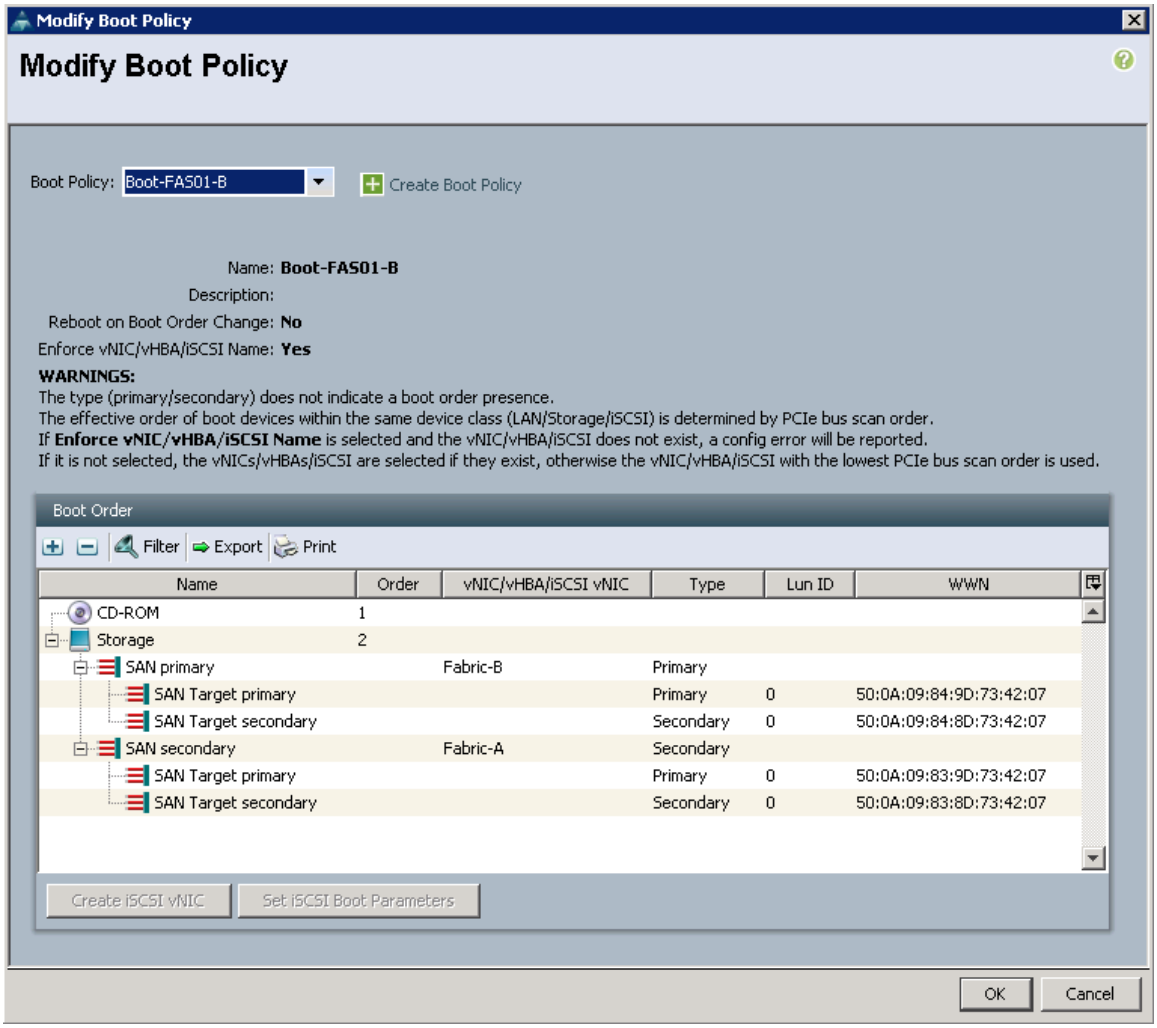

9. Select the Network tab and click **Modify vNIC/HBA Placement Policy**.

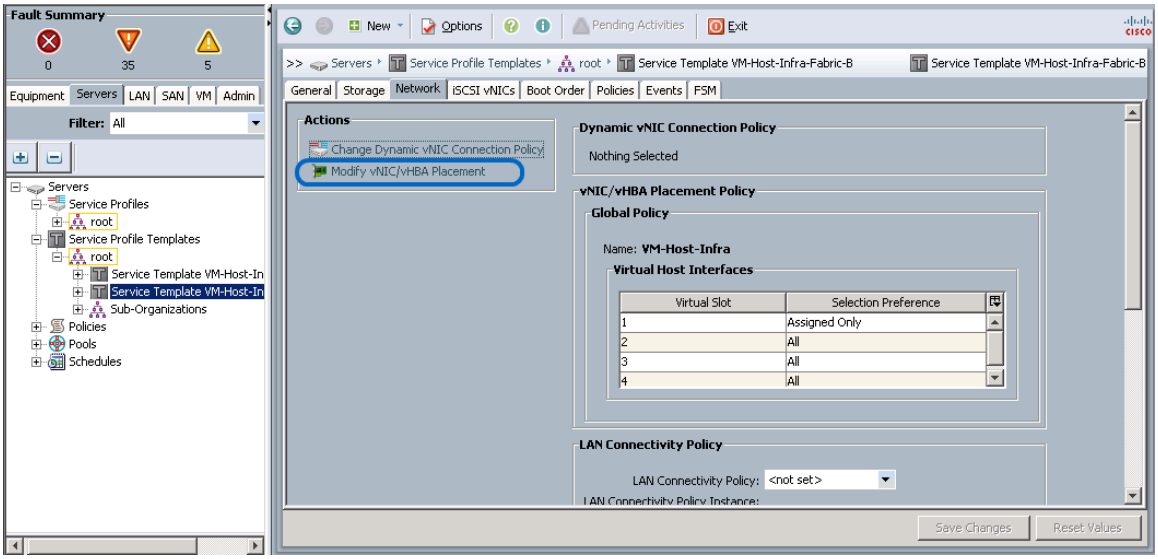

10. Move vHBA Fabric-B ahead of vHBA Fabric-A in the placement order and click **OK**.

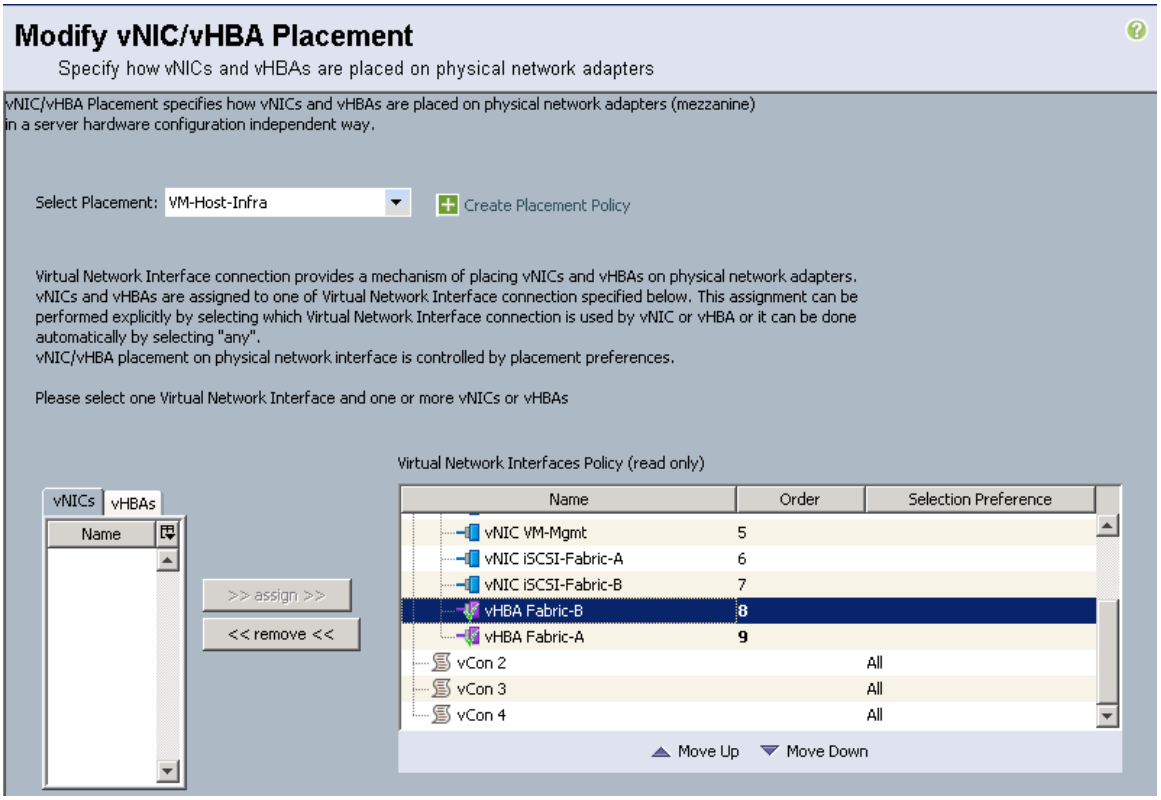

Create VM-FEX Port Profiles and Virtual Switch Clusters

These steps provide details for verifying that the port profiles and creating virtual Switch Clusters which will be used by the SR-IOV physical and virtual functions.

1. Select the VM tab at the top left of the window.
2. Note that the PF-iSCSI-A and PF-iSCSI-B port profiles have automatically been created because the VM option was selected in the vNIC Template creation step.

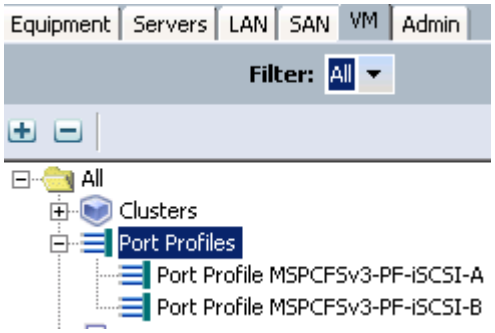

3. Right-click the cluster object and select Create Cluster
4. Enter the iSCSI-A fabric cluster name.
5. Open a PowerShell window and type the following command to generate a GUID.
   [system.guid]::NewGuid()
6. Copy the GUID generated in the previous step and past it Cluster ID text box.

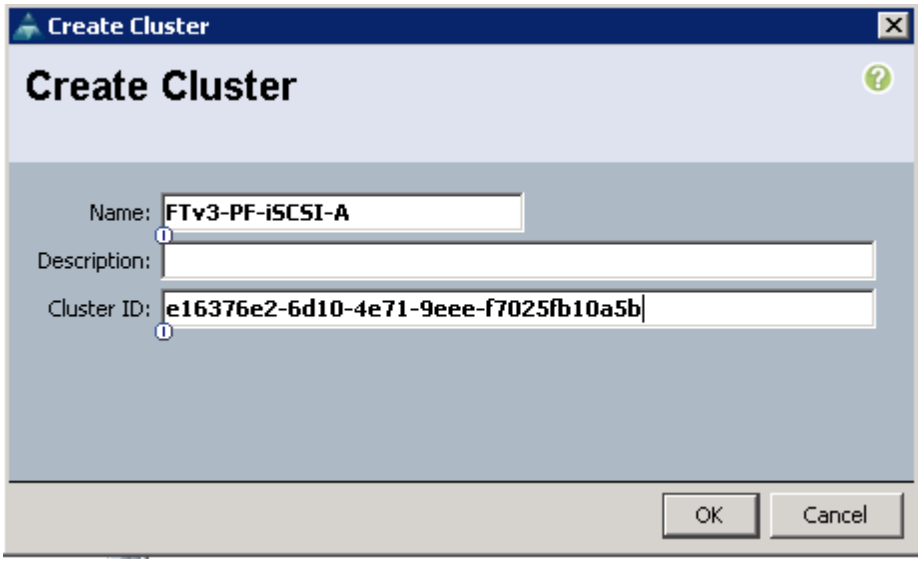

7. Click **OK** to create the cluster.
8. Repeat steps 3 through 7 to create the cluster for iSCSI-B fabric.

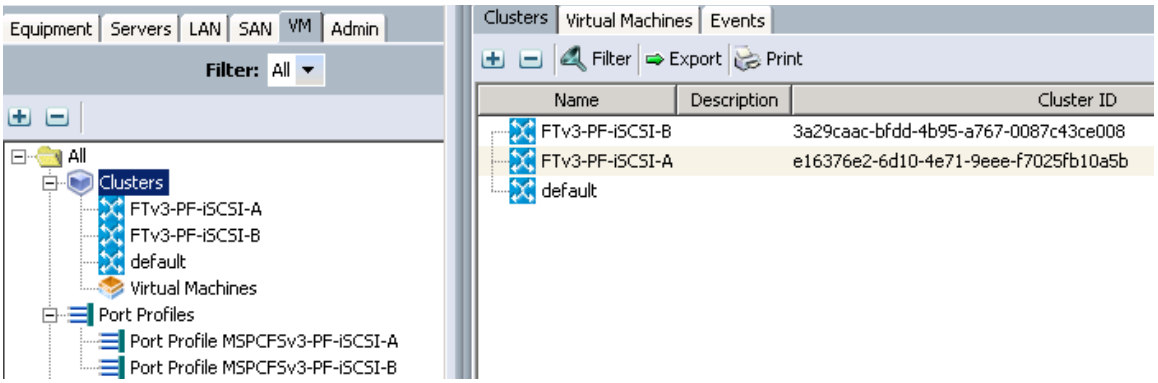

9. Right-Click on the PF-iSCSI-A Port Profile created in the previous step and select Create Profile Client.
10. Enter the name PF-iSCSI-A.
11. Leave the default setting All for the Datacenter and Folder option.
12. Select the previously created virtual switch cluster FT-v3-PF-iSCSI-A for the distributed virtual switch option.

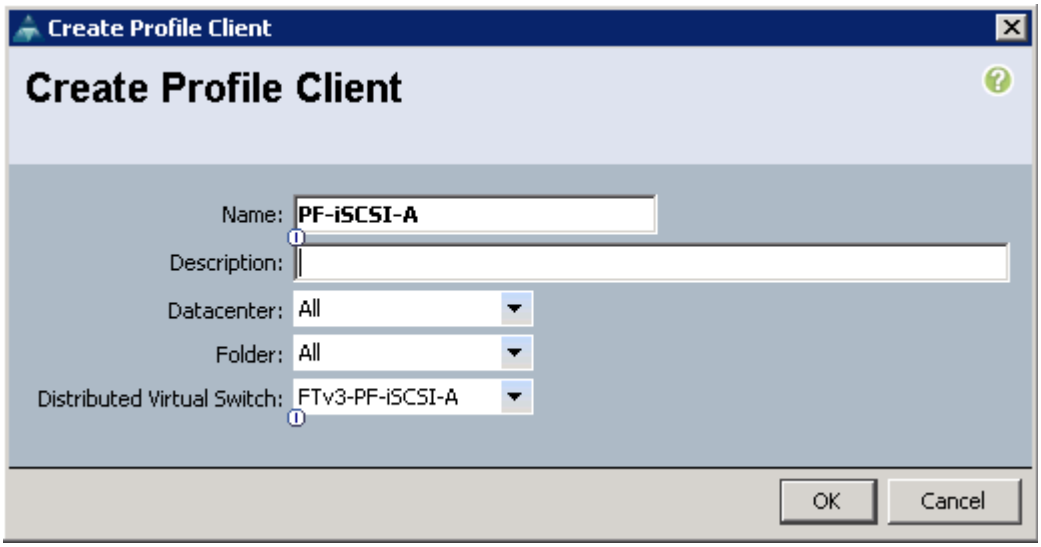

13. Click **OK** to create the port profile client.
14. Right-Click on the PF-iSCSI-B Port Profile created in the previous step and select Create Profile Client.
15. Enter the name PF-iSCSI-B.
16. Leave the default setting All for the Datacenter and Folder option.
17. Select the previously created virtual switch cluster FT-v3-PF-iSCSI-B for the distributed virtual switch option.

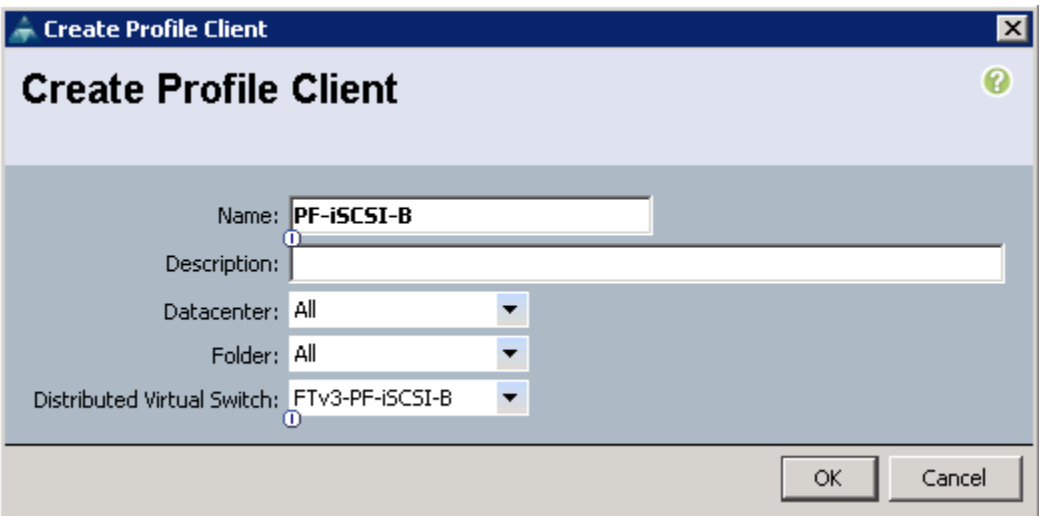

18. Click **OK** to create the port profile client.
19. Verity that the new port profile clients appear under each virtual switch cluster.

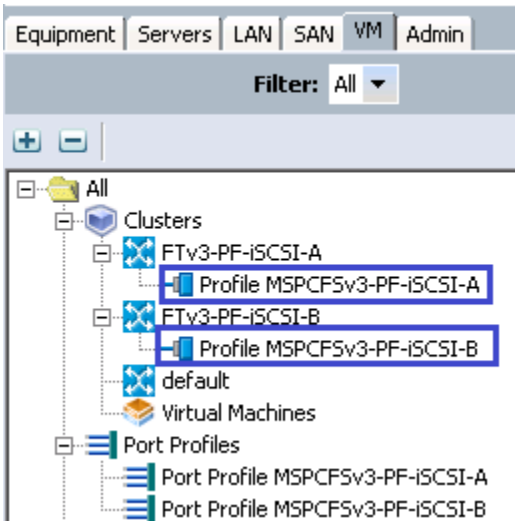

20. Select the Servers tab.
21. Find the service profile template created in the previous section and expand the vNICs object.
22. Expand the vNIC PF-iSCSI-A object and select Dynamic vNICs in the left tree view.
23. Verify that multiple vNICs appear in the right pane.

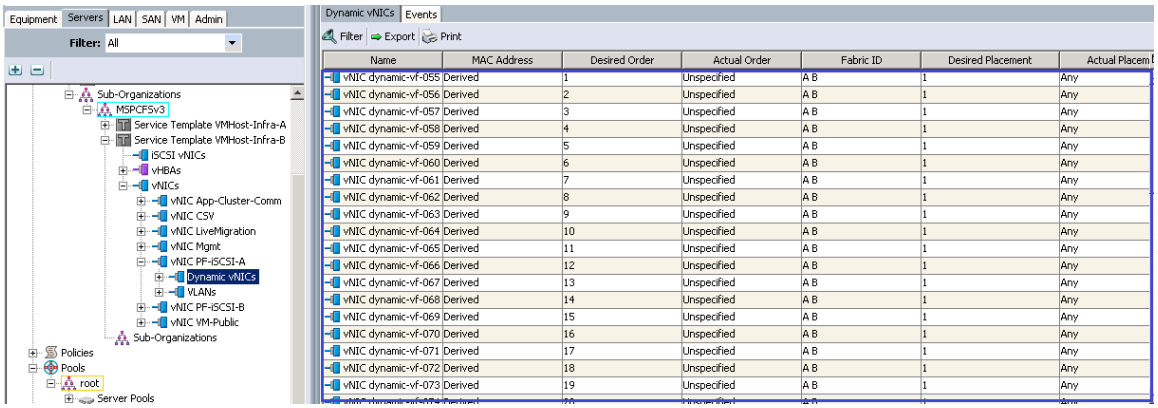

24. Repeat steps 22 and 23 for the vNIC PF-iSCSI-B.

Create Service Profiles

These steps provide details for creating a service profile from a template.

1. Select the Servers tab at the top left of the window.
2. Select Service Profile Templates VM-Host-Infra-Fabric-A
3. Right-click and select Create Service Profile From Template.
4. Enter VM-Host-Infra-0 for the service profile prefix.
5. Enter 1 for the number of service profiles to create.
6. Click **OK** to create the service profile.

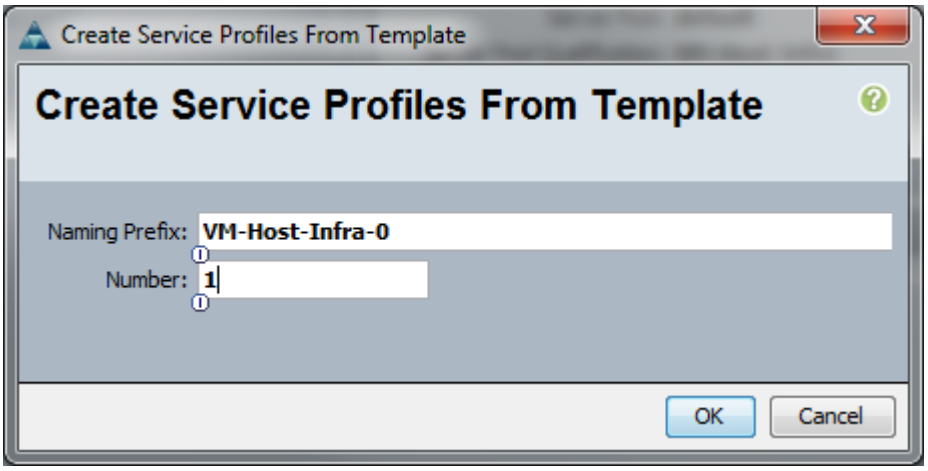

7. Click **OK** in the message box.
8. Select Service Profile Templates VM-Host-Infra-Fabric-B
9. Right-click and select Create Service Profile From Template.
10. Enter VM-Host-Infra-0 for the service profile prefix.
11. Enter 1 for the number of service profiles to create.
12. Click **OK** to create the service profile.

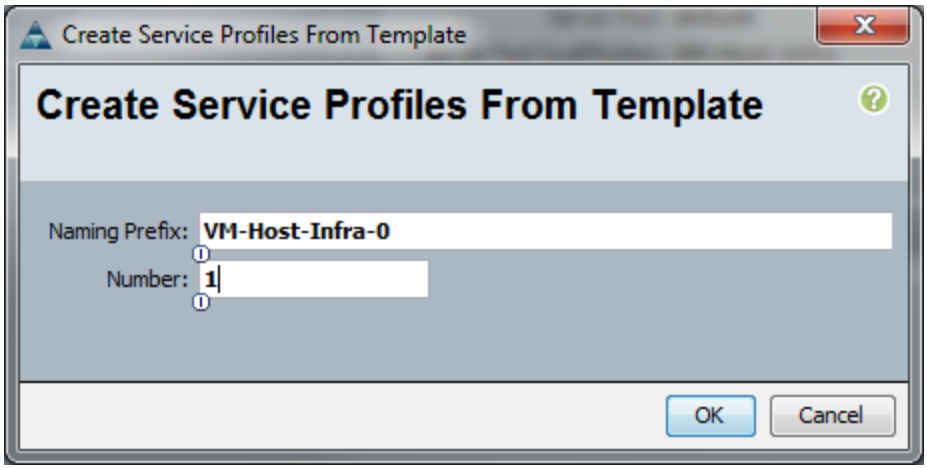

13. Click **OK** in the message box.

14. Verify that Service Profiles VM-Host-Infra-01 and VM-Host-Infra-02 are created. The service profiles will automatically be associated with the servers in their assigned server pools.

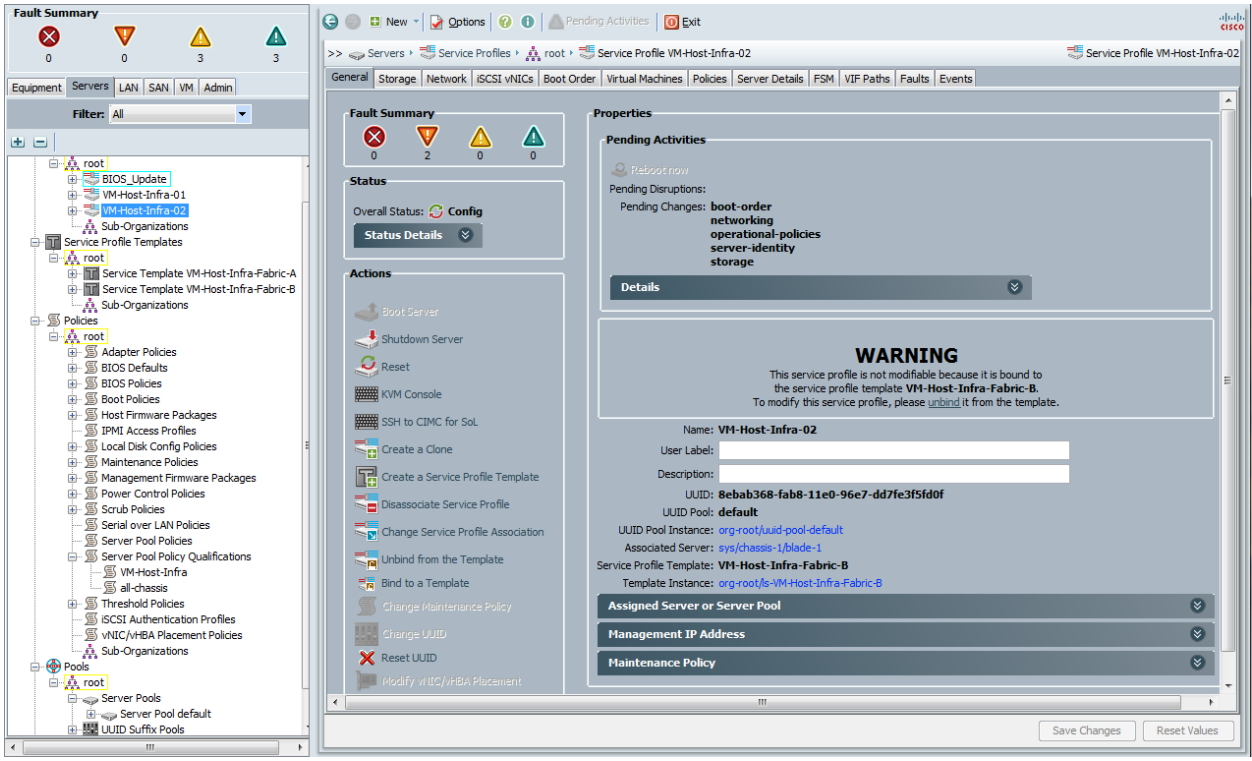

15. Click **OK** in the confirmation message.

16. Verify that the service profiles VM-Host-Infra-01 and VM-Host-Infra-02 have been created. The service profiles are automatically associated with the servers in their assigned server pools.

17. Repeat steps 1 through 17 to create Service Profiles VM-Host-Infra-03 and VM-Host-Infra-04. The odd numbered service profiles are created from service profile template VM-Host-Infra-Fabric-A and the even numbered service profiles are created from service profile template VM-Host-Infra-Fabric-B.

Add More Server Blades to the FlexPod Unit

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

Gather Necessary Information

After the Cisco UCS service profiles have been created (in the previous steps), the infrastructure blades in the environment each have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information in the tables below.

Table 10 FC Port Names for Storage Controllers 1 and 2.

NetApp Controller	FC Port	FC Portname
Controller A	2a	
	2b	
Controller B	2a	
	2b	

Table 11 vHBA WWPNs for Fabric A and Fabric B.

Cisco UCS Service Profile Name	Fabric-A WWPN	Fabric-B WWPN
VM-Host-Infra-01		
VM-Host-Infra-02		
VM-Host-Infra-03		
VM-Host-Infra-04		

Note To gather the information in the table above, launch the Cisco UCS Manager GUI, and in the left pane select the `Servers` tab. From there, expand `Servers > Service Profiles > root >`. Click each service profile and then click the `Storage` tab on the right. While doing so, record the WWPN information in the right display window for both vHBA_A and vHBA_B for each service profile in the table above.

SAN Configuration

Create Device Aliases

These steps provide details for configuring device aliases and zones for the primary boot path. Instructions are given for all target ports, however, the redundant path is enabled following operating system installation.

Nexus 5548 A

1. Using the information in FC Port Names for Storage Controllers 1 and 2.1 and vHBA WWPNs for Fabric A and Fabric B.2, Create device alias.
2. Type device-alias database.
3. Type device-alias name VM-Host-Infra-01_A pwwn <Fabric-A WWPN>.
4. Type device-alias name VM-Host-Infra-02_A pwwn <Fabric-A WWPN>.
5. Type device-alias name VM-Host-Infra-03_A pwwn <Fabric-A WWPN>.
6. Type device-alias name VM-Host-Infra-04_A pwwn <Fabric-A WWPN>.
7. Type device-alias name controller_A_2a pwwn <Controller A 2a WWPN>.
8. Type device-alias name controller_B_2a pwwn <Controller B 2a WWPN>.
9. Type exit.
10. Type device-alias commit.
11. Type copy running-config startup-config

Nexus 5548 B

1. Using the information in FC Port Names for Storage Controllers 1 and 2.1 and vHBA WWPNs for Fabric A and Fabric B.2, Create device alias.
2. Type device-alias database.
3. Type device-alias name VM-Host-Infra-01_B pwwn <Fabric-B WWPN>.
4. Type device-alias name VM-Host-Infra-02_B pwwn <Fabric-B WWPN>.
5. Type device-alias name VM-Host-Infra-03_B pwwn <Fabric-B WWPN>.
6. Type device-alias name VM-Host-Infra-04_B pwwn <Fabric-B WWPN>.
7. Type device-alias name controller_A_2b pwwn <Controller A 2b WWPN>.
8. Type device-alias name controller_B_2b pwwn <Controller B 2b WWPN>.
9. Type exit.
10. Type device-alias commit.

Create Zones for Each Service Profile (Part 1)

Note Windows requires that only a single SAN path to the boot LUN exists during installation time. The zoning procedure is broken into two parts for this reason. The first part creates a zoning configuration with only a single path to the boot LUN. The second zoning part is adds redundant paths the boot LUN and is configured after MPIO software is installed and configured.

Nexus 5548 A

Create the Zones and Add Members

1. Type zone name VM-Host-Infra-01_A vsan <Fabric A VSAN ID>.
2. Type member device-alias VM-Host-Infra-01_A.
3. Type member device-alias controller_A_2a.
4. Type exit.
5. Type zone name VM-Host-Infra-03_A vsan <Fabric A VSAN ID>.
6. Type member device-alias VM-Host-Infra-03_A.
7. Type member device-alias controller_A_2a.
8. Type exit.

Create the Zoneset and Add the Necessary Members

1. Type zoneset name Flexpod vsan <Fabric A VSAN ID>.
2. Type member VM-Host-Infra-01_A.
3. Type member VM-Host-Infra-03_A.
4. Type exit.

Activate the Zoneset

1. Type zoneset activate name flexpod vsan < Fabric A VSAN ID>.
2. Type exit.
3. Type copy run start.

Nexus 5548 B

Create the Zones and Add Members

1. Type zone name VM-Host-Infra-02_B vsan <Fabric B VSAN ID>.
2. Type member device-alias VM-Host-Infra-02_B.
3. Type member device-alias controller_B_2b.
4. Type exit.
5. Type zone name VM-Host-Infra-04_B vsan <Fabric B VSAN ID>.
6. Type member device-alias VM-Host-Infra-04_B.
7. Type member device-alias controller_B_2b.
8. Type exit.

Create the Zoneset and Add the Necessary Members

1. Type zoneset name flexpod vsan <Fabric B VSAN ID>.
2. Type member VM-Host-Infra-02_B.
3. Type member VM-Host-Infra-04_B.
4. Type exit.

Activate the Zoneset

1. Type zoneset activate name flexpod vsan < Fabric B VSAN ID>.
2. Type exit.
3. Type copy run start.

NetApp FAS3240A Deployment Procedure - Part 2

The following sections provide detailed procedures for configuring the interface groups (or igroups), creating LUNs for the service profiles on the storage controllers, and mapping those LUNs to the igroups to be accessible to the service profiles.

Create iGroups

The following steps provide details for configuring the necessary igroups on the storage controller the enable the mapping of a given host to the storage resources.

Controller A

For the odd service profile to boot off of controller A, execute the following to create igroups for each vHBA:

```
igroup create -f -t hyper_v VM-Host-Infra-01 <vHBA_A WWPN> <vHBA_B WWPN>.
igroup create -f -t hyper_v VM-Host-Infra-03 <vHBA_A WWPN> <vHBA_B WWPN>.
```

Controller B

For the even service profile to boot off of controller B, execute the following to create igroups for each vHBA:

```
igroup create -f -t hyper_v VM-Host-Infra-02 <vHBA_A WWPN> <vHBA_B WWPN>.
igroup create -f -t hyper_v VM-Host-Infra-04 <vHBA_A WWPN> <vHBA_B WWPN>.
```

Create LUNs

This section provides detailed procedure for configuring the necessary LUNs on the storage controller for deployment of the SAN booted windows operating system.

Controller A

For the odd service profile to boot off of controller A, execute the following to create the LUN for each OS installation:

```
lun create -s 150g -t hyper_v -o noreserve /vol/ucs_boot/VM-Host-Infra-01.lun
lun create -s 150g -t hyper_v -o noreserve /vol/ucs_boot/VM-Host-Infra-03.lun
```

Controller B

For the even service profile to boot off of controller B, execute the following to create the LUN for each OS installation:

```
lun create -s 150g -t hyper_v -o noreserve /vol/ucs_boot/VM-Host-Infra-02.lun
lun create -s 150g -t hyper_v -o noreserve /vol/ucs_boot/VM-Host-Infra-04.lun
```

Map LUNs to iGroup

For mapping the necessary LUNs on the storage controller to the created iGroups, execute these commands on controller A and Controller B.

Controller A

For the odd service profile to boot off of controller A map the LUN for the OS installation:

```
lun map /vol/ucs_boot/VM-Host-Infra-01.lun VM-Host-Infra-01
lun map /vol/ucs_boot/VM-Host-Infra-03.lun VM-Host-Infra-03
```

Controller B

For the even service profile to boot off of controller B map the LUN for the OS installation:

```
lun map /vol/ucs_boot/VM-Host-Infra-02.lun VM-Host-Infra-02
lun map /vol/ucs_boot/VM-Host-Infra-04.lun VM-Host-Infra-04
```

Microsoft Windows Server 2012 Hyper-V Deployment Procedure

Setup the Windows Server 2012 install

This section details the steps required to prepare the server for OS installation.

The following steps describe adding and mapping ISO image for installing OS:

All Hosts

1. Right-click on the VM-Host service profile and select KVM Console.
2. From the virtual KVM Console, select the Virtual Media tab.
3. Select Add Image in the right pane.
4. Browse to the Windows Server 2012 installation ISO image file and click **Open**.
5. Map the image that you just added by selecting Mapped.
6. To boot the server, select the KVM tab.
7. Select Power On Server in the KVM interface Summary tab, and then click **OK**.

Install Windows Server 2012

The following steps describe the installation of Windows Server 2012 to each hosts:

All Hosts

1. On boot, the machine detects the presence of the Windows installation media.
2. After the installer has finished loading, Enter the relevant region information and click **Next**.
3. Click **Install now**.
4. Enter the Product Key and click **Next**.
5. Select **Windows Server 2012 Datacenter (Server with a GUI)** and click **Next**.

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

6. After reviewing the EULA, accept the license terms and click **Next**.
7. Select Custom: Install Windows only (advanced).
8. Select Custom (advanced) installation.
9. In the Virtual Media Session manager uncheck the Mapped checkbox for the Windows ISO and select yes to confirm.
10. Click **Add Image**.
11. Browse to the Cisco fNIC driver ISO, click **Open**.
12. Check the Mapped checkbox next to the Cisco fNIC Driver ISO.
13. Back in the KVM Console, click **Load Driver** and then, click **OK**.
14. The Cisco VIC FCoE Storport Miniport driver should auto detected; Click **Next**.

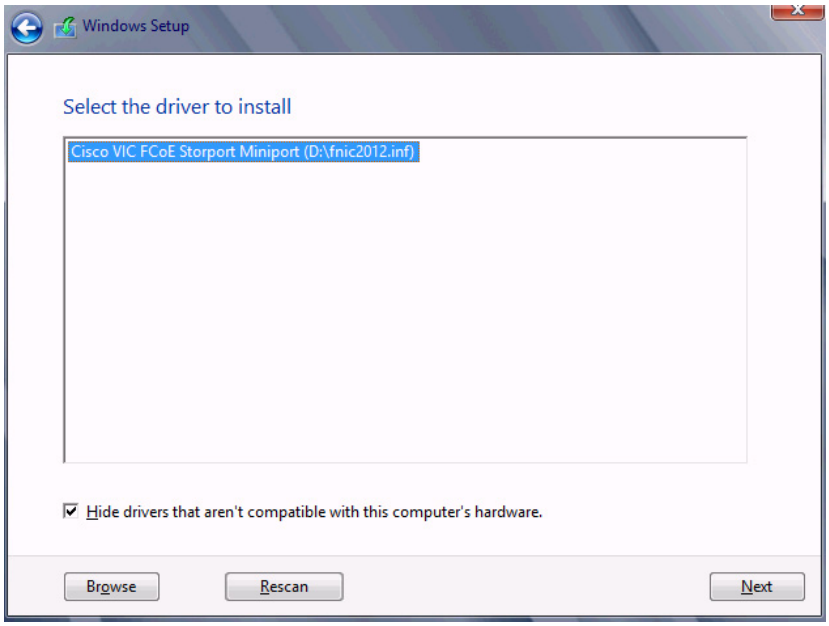

15. You should see a LUN listed in the drive selection screen.

Note

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

- The message “Windows Can’t be installed on this drive” appears because the Windows installation ISO image is not mapped at this time.
- The Cisco eNIC driver can be loaded at this point in the same way as the fNIC driver. Loading the eNIC driver at this time bypasses the need to load the eNIC driver in the section titled “Installing Windows eNIC Driver”.

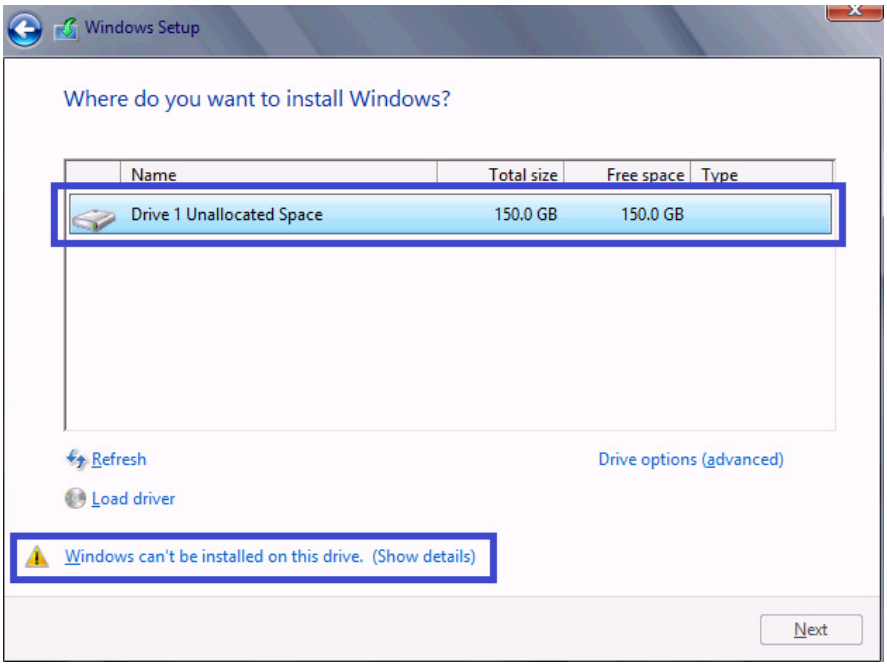

16. In the Virtual Media Session manager, uncheck the Mapped checkbox for the Cisco Driver ISO that you had recently added (fNIC driver) and choose yes to confirm.
17. Check the Mapped checkbox for the Windows ISO in the virtual media session manager.
18. In the KVM console, click **Refresh** to update the cdrom drive status.

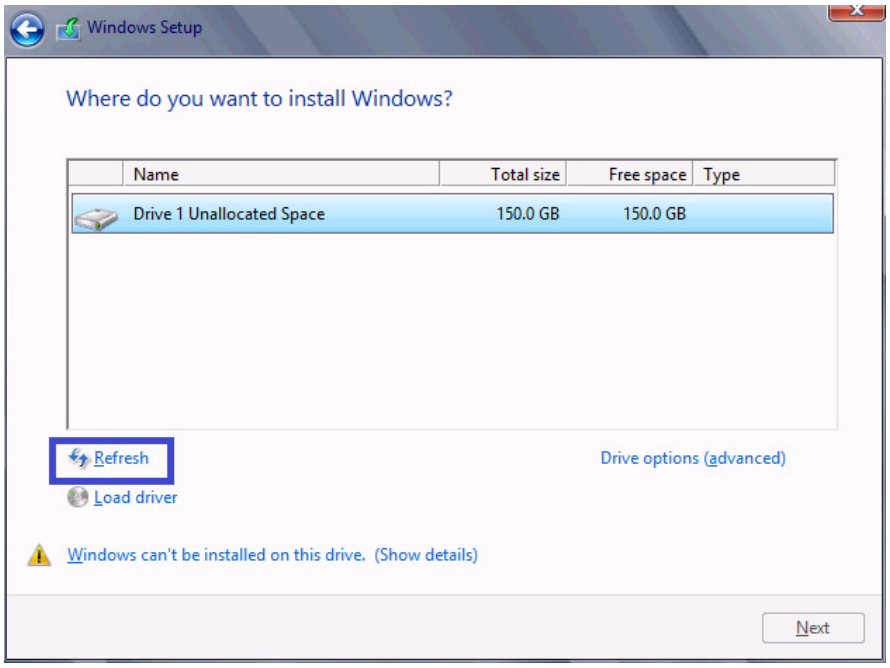

19. Select the new LUN, and click on the "Windows cannot be installed to this drive" link.
20. Click **OK** to bring the LUN online.
21. Select the LUN, and click **Next** continue with the install.
22. When Windows is finished installing, enter an administrator password on the settings page and click **Finish**.

Install Windows Roles and features

This section provides detailed information on installing all the required roles and features from Windows Server 2012 Installation media. If you have unmapped the installation ISO you will need to remap it now.

All Hosts

1. Log into Windows with the Administrator password previously entered during installation.
2. Verify that the Windows installation disk is mapped to E: drive.
3. To launch a PowerShell prompt, right-click the PowerShell icon in the taskbar, and click **Run as Administrator**.
4. To add the .Net 3.5 feature, type the following command:

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

5. To add Hyper-V, Failover-Clustering and MPIO, type the following command:

```
Add-WindowsFeature Hyper-V, Failover-Clustering, Multipath-IO,
Data-Center-Bridging -IncludeManagementTools -Restart
```

Install Windows eNIC Drivers

This section provides detailed information on installing all the required network drivers.

All Hosts

1. In the Virtual Media Session manager, uncheck the Mapped checkbox for the Windows ISO.
2. Click **Add Image**.
3. Browse to the Cisco eNIC driver ISO, click **Open**.
4. Check the Mapped checkbox for the Cisco eNIC driver ISO.
5. In the KVM console, open Server Manager, and select **Tools** > **Computer Management**.
6. In Computer Manager, select **System Tools** > **Device Manager** > **Other devices.**

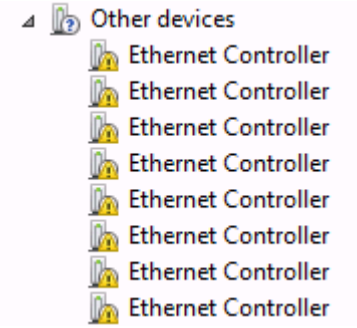

7. Right-click one of the Ethernet Controller, and select Update Driver Software.
8. Click **Browse**, and select CDROM drive, click **OK**.
9. Click **Next** > **Close**.

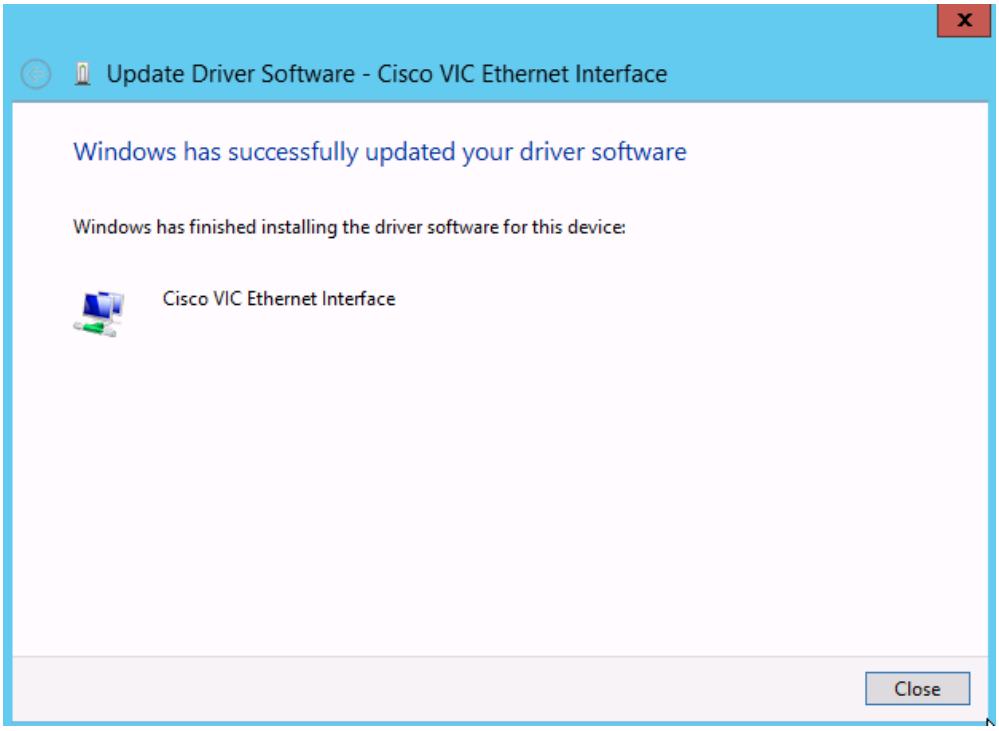

10. Right-click the next Ethernet Controller and select Update Driver Software.
11. Click Search automatically for update driver software.
12. Click **Close**.
13. Repeat these steps for the remaining Ethernet Controllers.

Note Alternatively to steps 6 to 13, the Cisco eNIC driver can be loaded for all devices at once by issuing the command: **pnputil –i –a <directory>enic6x64.inf** where <directory> is the location of the eNIC driver.

14. All Cisco VIC Ethernet devices will appear under Network Adapters.

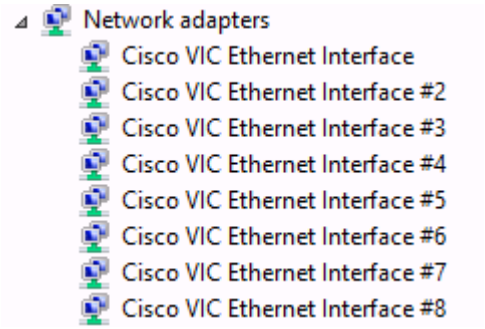

Configure Windows Networking for FlexPod

This section provides detailed information on renaming the network for each Hyper-V host.

All Hosts

1. In Server Manger, select Local Server on the left.
2. Click **IPv4 address assigned by DHCP, IPv6 enabled**, to launch the network connections control panel.
3. Right-click on each eNIC one by one, and select Status.
4. Click **Details**, and note the Physical Address.

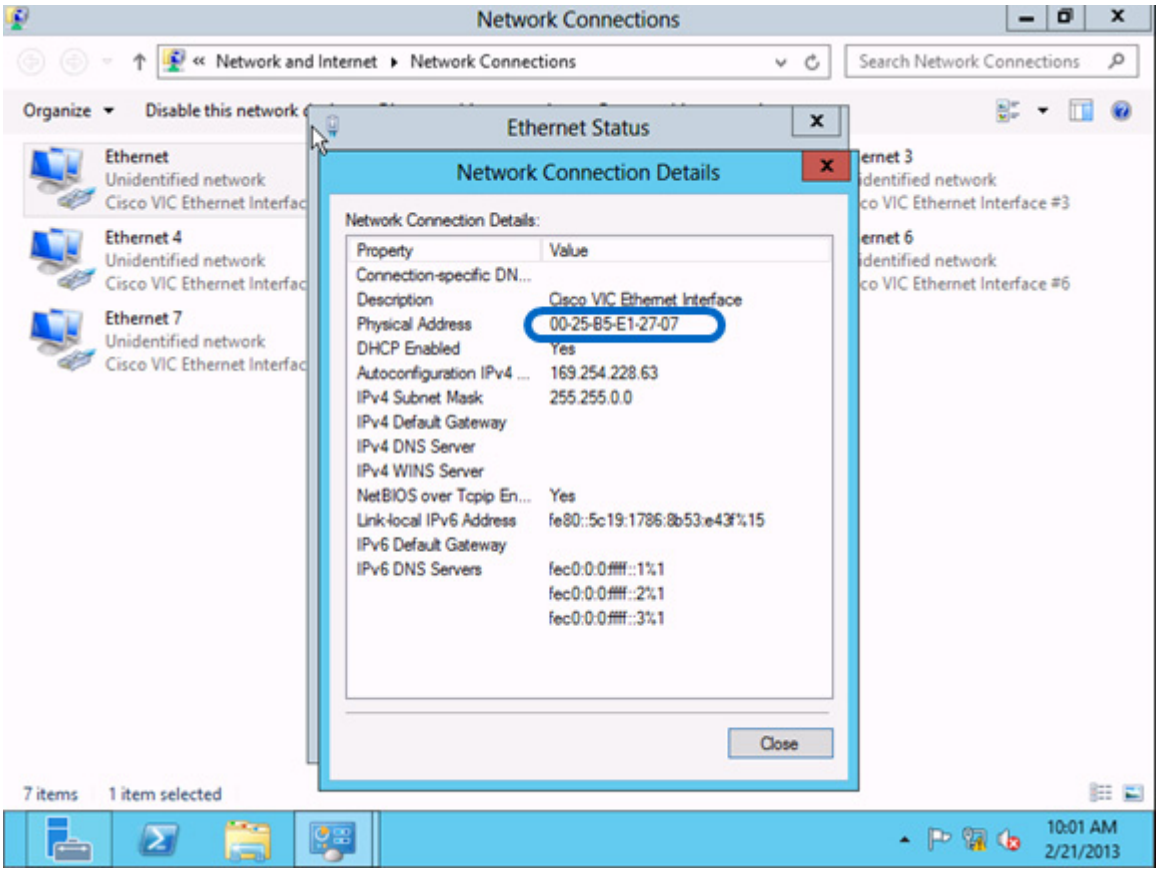

Note The following PowerShell command provides a list of the adapters with their associated MAC addresses it can be used instead of performing steps 3 through 5 for each NIC.

```
Gwmi Win32_NetworkAdapter | Where{$_.MACaddress -ne $Null} | FT NetConnectionID,
MACaddress
```

5. In the KVM console select **Properties** > **Network**. In Network window you can see all the vNICs.

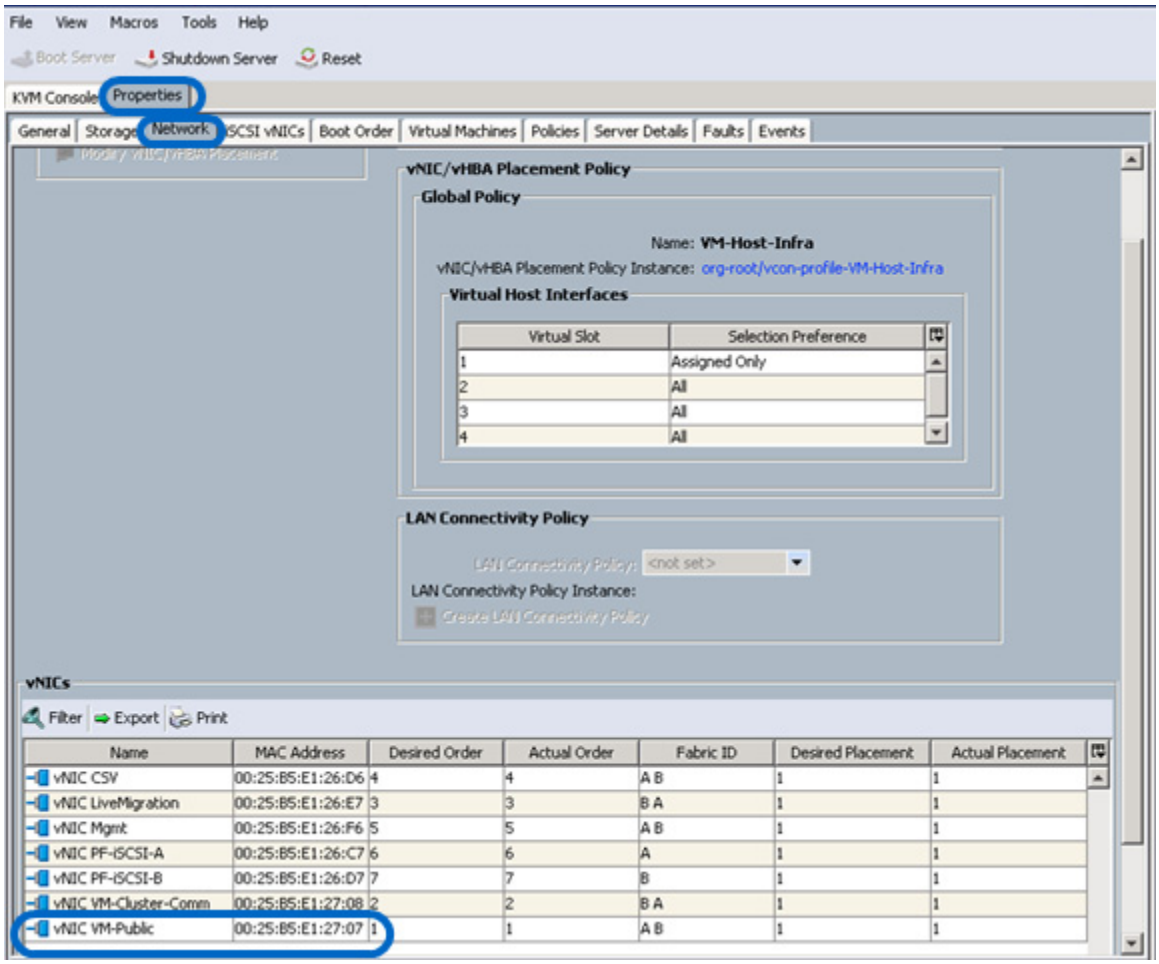

6. Identify the vNIC with the MAC Address noted in step 4.
7. In windows, rename the LAN adapter to reflect the network it is associated with.
8. Set the appropriate IP settings for that adapter.

Note Assign IP Addresses to the LiveMigration, CSV, and Management adapters.

9. Repeat for each eNIC in windows.
10. In the Network Connections Control Panel. Press **Alt** key to see the extended menu, and select **Advanced > Advanced Settings**
11. Select the adapter and use the arrows to move it up or down in binding order.
12. The recommended binding order is:
    - Mgmt
    - CSV
    - LiveMigration
    - VM-Public
    - VM-Cluster-Comm
    - PF-iSCSI-A

– PF-iSCSI-B

13. Select the PF-iSCSI-A connection.
14. In the Bindings for PF-iSCSI-A, uncheck the File and Printer Sharing for Microsoft Networks, and Client for Microsoft Networks checkboxes.
15. Repeat steps 13 and 14 for the PF-iSCSI-B connection.

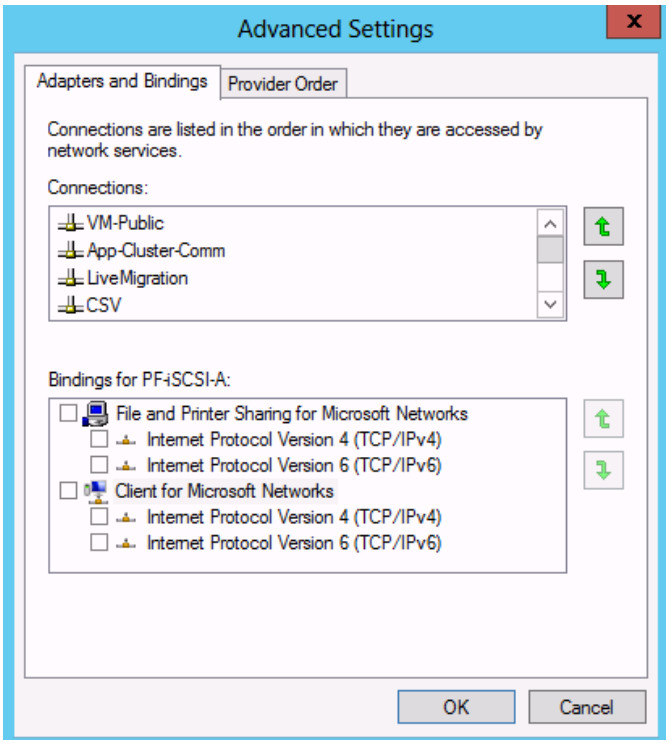

Install NetApp MPIO DSM

This section provides information on installing the NetApp Device Specific Module (DSM). For more information on NetApp DSM installation, see:

https://library.netapp.com/ecm/ecm_get_file/ECMP1141002

All Hosts

1. Download NetApp MPIO DSM at:

   http://support.netapp.com/NOW/download/software/mpio_win/4.0/ntap_win_mpio_4.0_setup_x64.msi
2. Launch the Installer, click **Next**.
3. Click **OK** acknowledging the ALUA warning.
4. Accept the EULA and Click **Next**.
5. Enter the license Key and Click **Next**.
6. Keep Use the default SYSTEM account checkbox checked and Click **Next**.
7. Select Yes, install the Hyper-V Guest Utilities, and Click **Next**.

8. Click **Next** through the driver information pane.
9. Validate the installation path and click **Next**.
10. Click **Install.**
11. When prompted reboot the host.

Create Zones for Each Service Profile (Part 2)

The following section describes how to zone in the redundant fabric paths.

Nexus 5548 A

Create the Zones and Add Members

1. Type zone name VM-Host-Infra-01_A vsan <Fabric A VSAN ID>.
2. Type member device-alias controller_B_2a.
3. Type exit.
4. Type zone name VM-Host-Infra-02_A vsan <Fabric A VSAN ID>.
5. Type member device-alias VM-Host-Infra-02_A.
6. Type member device-alias controller_A_2a.
7. Type member device-alias controller_B_2a.
8. Type exit.
9. Type zone name VM-Host-Infra-03_A vsan <Fabric A VSAN ID>.
10. Type member device-alias controller_B_2a.
11. Type exit.
12. Type zone name VM-Host-Infra-04_A vsan <Fabric A VSAN ID>.
13. Type member device-alias VM-Host-Infra-04_A.
14. Type member device-alias controller_A_2a.
15. Type member device-alias controller_B_2a.
16. Type exit.

Add the Necessary Members to the Zoneset

1. Type zoneset name flexpod vsan <Fabric A VSAN ID>.
2. Type member VM-Host-Infra-02_A.
3. Type member VM-Host-Infra-04_A.
4. Type exit.

Activate the Zoneset

1. Type zoneset activate name flexpod vsan < Fabric A VSAN ID>.
2. Type exit.
3. Type copy run start.

Nexus 5548 B

Create the Zones and Add Members

1. Type zone name VM-Host-Infra-01_B vsan <Fabric B VSAN ID>.
2. Type member device-alias VM-Host-Infra-01_B.
3. Type member device-alias controller_A_2b.
4. Type member device-alias controller_B_2b
5. Type exit.
6. Type zone name VM-Host-Infra-02_B vsan <Fabric B VSAN ID>.
7. Type member device-alias controller_A_2b.
8. Type exit.
9. Type zone name VM-Host-Infra-03_B vsan <Fabric B VSAN ID>.
10. Type member device-alias VM-Host-Infra-03_B.
11. Type member device-alias controller_A_2b.
12. Type member device-alias controller_B_2b.
13. Type exit.
14. Type zone name VM-Host-Infra-04_B vsan <Fabric B VSAN ID>.
15. Type member device-alias controller_A_2b.
16. Type exit.

Create the Zoneset and Add the Necessary Members

1. Type zoneset name flexpod vsan <Fabric B VSAN ID>.
2. Type member VM-Host-Infra-01_B.
3. Type member VM-Host-Infra-03_B.
4. Type exit.

Activate the Zoneset

1. Type zoneset activate name flexpod vsan < Fabric B VSAN ID>.
2. Type exit.
3. Type copy run start.

Install Cisco Virtual Switch Forwarding Extensions for Hyper-V

Cisco Virtual Switch Forwarding Extensions for Hyper-V enable SR-IOV capability when using VM-FEX for Hyper-V. The installation package is located in the Cisco UCS Manager drivers ISO image starting with version 2.1. The installation package file name is CSCO_VIO_INSTALLER_64_2.x.x.msi. It is located in the Windows\installers\Cisco\MLOM\W2K12\x64 directory.

All Hosts

1. Run CSCO_VIO_INSTALLER_64_2.x.x.msi with administrator rights.
2. Click **Next** in the welcome screen.

3. Accept the EULA and click **Next**.
4. Select Custom setup.
5. Clear the following options:
   - VIC Ethernet
   - VIC FCoE
   - VIC iSCSI Dump
   - VIC Management
   - VIC Ethernet Utility

Note Only VIC Vmfex Forwarding Extension need to be installed

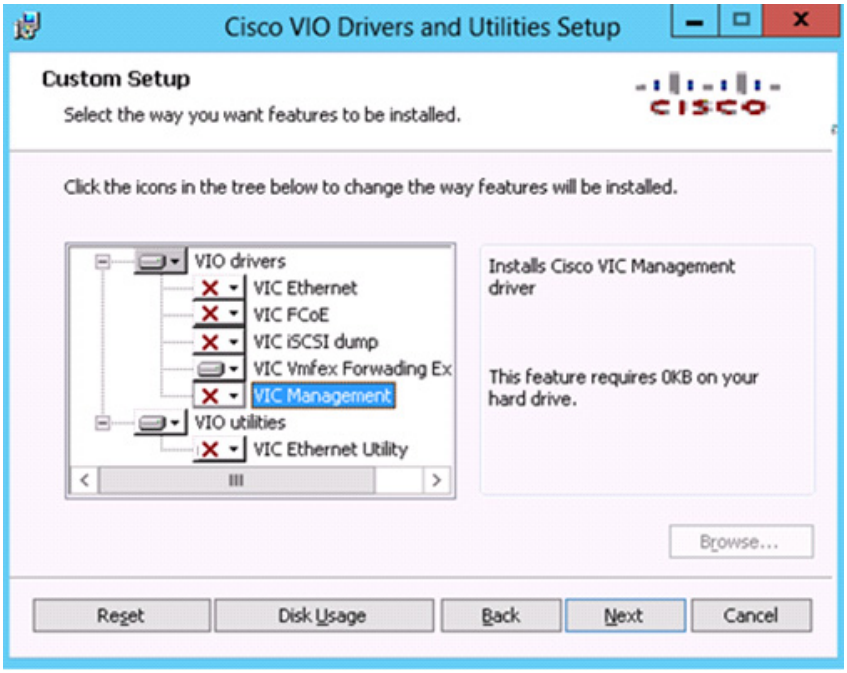

6. Click **Next**.
7. Click **Install**.
8. Click **Finish** to complete the installation.

Create Hyper-V Virtual Network Switches

All Hosts

1. Open Windows PowerShell command window.

2. Create the Hyper-V virtual switches with the following parameters:

Virtual Network Name	Connection Type	Enable SR-IOV	Interface Name	Share Network with Management Host
VM-Public	External	No	VM-Public	No
VM-Cluster-Comm	External	No	VM-Cluster-Comm	No
VF-iSCSI-A	External	Yes	PF-iSCSI-Fabric-A	No
VF-iSCSI-B	External	Yes	PF-iSCSI-Fabric-B	No

3. Create virtual switch VM-Public

   ```
   New-vmswitch -name VM-Public -NetAdapterName VM-Public -AllowManagementOS $false
   ```

4. Create virtual switch VM-Cluster-Comm.

   New-vmswitch –name VM-Cluster-Comm –NetAdapterName VM-Cluster-Comm -AllowManagementOS $false

5. Create virtual switch VF-iSCSI-A.

   New-vmswitch –name VF-iSCSI-A –NetAdapterName PF-iSCSI-A -AllowManagementOS $false –EnableIov $true

6. Create virtual switch VF-iSCSI-B.

   New-vmswitch –name VF-iSCSI-B –NetAdapterName PF-iSCSI-B -AllowManagementOS $false –EnableIov $true

7. Open Hyper-V Manager
8. Select Hyper-V Virtual Switch Manger in the action pane.
9. Expand the VF-iSCSI-A virtual switch.
10. Select Extensions under the VF-iSCSI-A virtual switch.
11. Check the Cisco VmFex Switch checkbox.

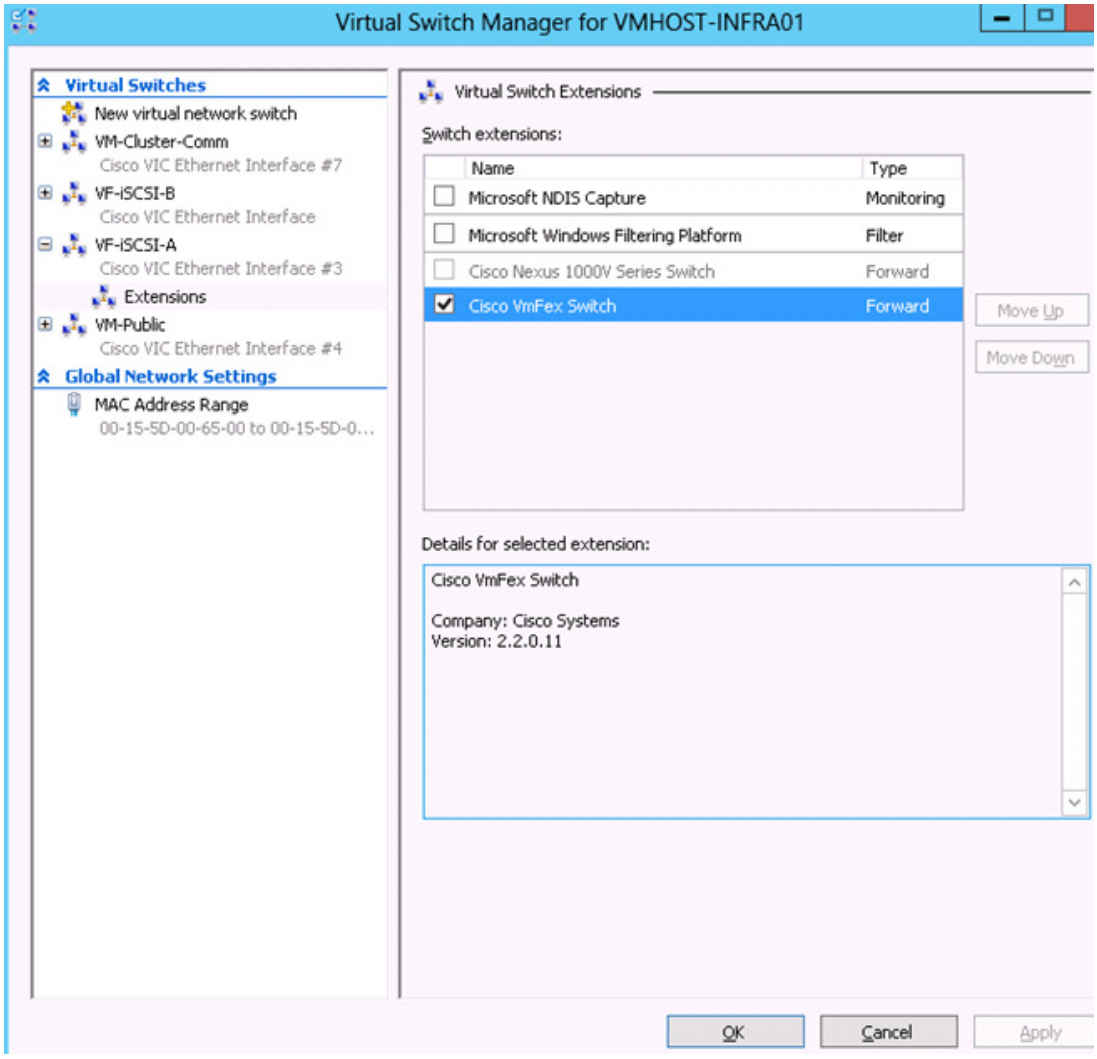

12. Click **Apply**.
13. Repeat for virtual switch VF-iSCSI-B.

Domain Controller Virtual Machines

Most environments will already have an active directory infrastructure and will not require additional domain controllers to be deployed for Flexpod with Microsoft Windows Server 2012 Hyper-V. The optional domain controllers can be omitted from the configuration in this case or used as a resource domain. The domain controller virtual machines will not be clustered because of the redundancy provided by deploying multiple domain controllers running in virtual machines on different servers. Since these virtual machines reside on Hyper-V hosts that run Windows Failover cluster, but are not clustered themselves, Hyper-V Manager should be used to manage them instead of Virtual Machine Manager.

Prepare Nodes for Clustering

This section provides details on preparing each node to be added to the Hyper-V cluster.

All Hosts

1. Rename the Host.

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

2. Add the host to Active Directory.

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

Install NetApp SnapDrive

This section provides detailed information on installing NetApp SnapDrive Windows. For more information on NetApp SnapDrive Windows installation see:

https://library.netapp.com/ecm/ecm_get_file/ECMP1141002

Service Account preparation

1. In active directory, create a SnapDrive service account.

Note This account requires no special delegation.

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

All Hosts

1. Download SnapDrive installer at:

   http://support.netapp.com/NOW/download/software/snapdrive_win/6.5/SnapDrive6.5_x64.exe

2. Launch the Installer, click **Next**.
3. Accept the EULA and click **Next**.
4. Select the Storage based Licensing method and click **Next**.

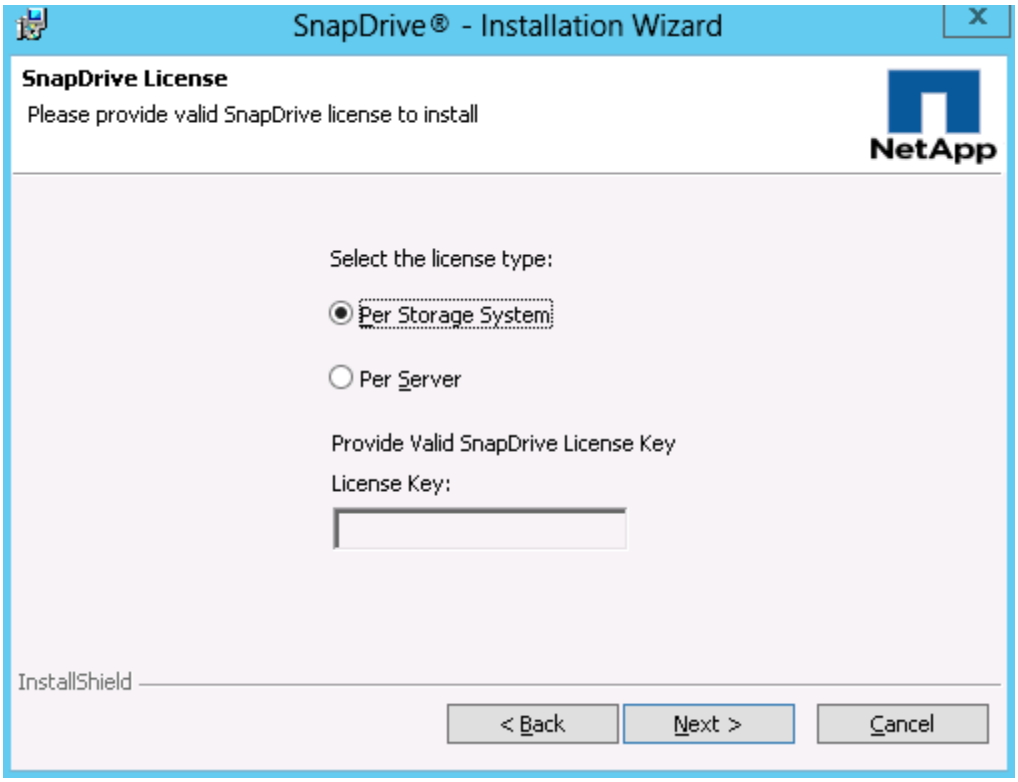

5. Enter your User Name, and Organization information, and click **Next**.
6. Validate the installation path and click **Next**.
7. Check the Enable SnapDrive to communicate through the Windows Firewall checkbox and click **Next**.
8. Enter the Account information for the Snapdrive service account, Click **Next**.
9. Click **Next**, through the SnapDrive Web Service Configuration.
10. Uncheck the Enable Transport Protocol Settings checkbox, and click **Next**.
11. Leave Enable Protection Manger Integration Unchecked, and click **Next**.
12. Click **Install**.
13. After the installation is finished. Launch a new Windows 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.

14. Configure SnapDrive Preferred IP settings for each controller.

```
sdcli preferredIP set -f <<var_controller1>> -IP <<var_controller1_e0m_ip>>
sdcli preferredIP set -f <<var_controller2>> -IP <<var_controller2_e0m_ip>>
```

15. Configure SnapDrive transport protocol authentication configuration for each controller.

```
sdcli transport_protocol set -f <<var_controller1>> -type https -user root -pwd
<<var_admin_passwd>>
sdcli transport_protocol set -f <<var_controller2>> -type https -user root -pwd
<<var_admin_passwd>>
```

Install NetApp SnapManager for Hyper-V

This section provides detailed information on installing NetApp SnapManger for Hyper-V. For more information on NetApp SnapManger for Hyper-V installation, see:

https://library.netapp.com/ecm/ecm_get_file/ECMP1141002

Service Account Preparation

1. In active directory create a SnapDrive service account.

Note This account requires no special delegation.

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

All Hosts

1. Download the SnapManger for Hyper-V installer at:

   http://support.netapp.com/NOW/download/software/snapmanager_hyperv_win/1.2/SMHV1.2_NetApp_x64.exe
2. Launch the Installer, click **Next**.
3. Accept the EULA and click **Next**.
4. Select the Storage based Licensing method and click **Next**.
5. Enter your User Name, and Organization information, and click **Next**.
6. Validate the installation path and click **Next**.
7. Enter the Account information for the SMHV service account, Click **Next**.
8. Click **Next**, through the SMHV Web Service Configuration.
9. Click **Install**.

Install VM-FEX Port Profile Management Utilities in Hyper-V

All Hosts

1. Run the VM-FEX Port Profile Management Utilities installation package VMFEX_ToolS_64_2.3.2.msi.
2. Click **Next** on the welcome screen.
3. Accept the license agreement and click **Next**.
4. Click **Custom installation**.
5. Click **Next**.
6. Click **Install**.
7. Click **Finish** to complete the installation.

Create a Cluster

One Server Only

1. Launch a PowerShell prompt with administrative permissions, by right clicking on the PowerShell icon and selecting **Run as Administrator**.
2. Create a new cluster.

```
New-Cluster -Name <cluster_name> -Node <Node1>, <Node2>, <node3>, <node4>
-NoStorage -StaticAddress <cluster_ip_address>
```

3. Rename Cluster Networks

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

4. Designate the CSV network.

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

Configure Live Migration network.

One Server Only

1. Open Failover Cluster Manger from Server Manager. Select **Tools > Failover Cluster Manger**.
2. Expand the Cluster tree on the left, and right-click on Networks, select Live Migration Settings.

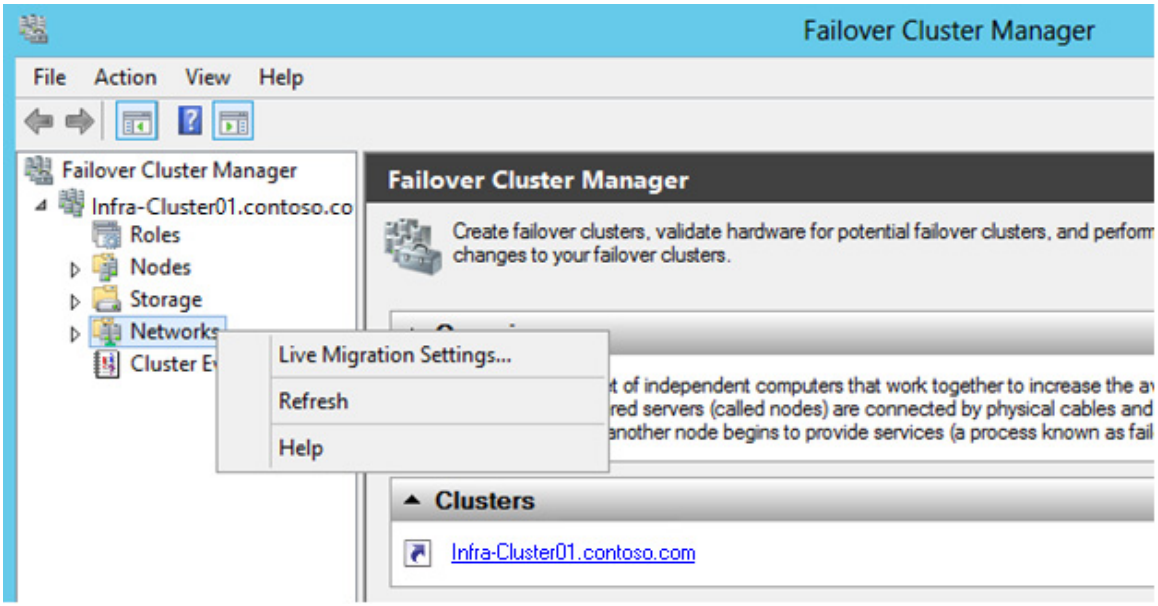

3. Clear all networks but the LiveMigration network and click **OK**.

Change the Cluster to Use a Quorum Disk.

One Server Only

1. Open Windows PowerShell prompt. To move the Available Storage cluster group, run the command:

```
Move-ClusterGroup "Available Storage" -Node $env:COMPUTERNAME | Start-ClusterGroup
```

2. Open SnapDrive from the start screen to configure cluster storage.
3. From SnapDrive, Open the Server name, then Open the Disks Icon.
4. Right-click the Disks Icon and choose Create Disk.
5. Enter the IP Address of the controller that contains the Cluster Quorum Volume.
6. Once connected, open the controller tree and select the Cluster Quorum Volume.
7. Enter the name of the LUN in the LUN name field, click **Next**.
8. Select Shared (Microsoft Cluster Services only) and click **Next**.
9. Validate that all nodes of the cluster are shown and click **Next**.
10. Change the drive letter to **W:**, set the LUN size to be 5GB and click **Next**.
11. Click **Next** through the Volume properties confirmation.
12. Select the FCP WWPN to Map the LUN to click **Next**.
13. Select Automatic igroup management and Click **Next**.
14. Select the Available Storage cluster group, and click **Next.**
15. Click **Finish**.
16. Make sure that the Q: drive is accessible on all of the nodes.
17. Expand failover cluster manager, right-click the selected cluster, select **More Actions > Configure Cluster Quorum Settings...**

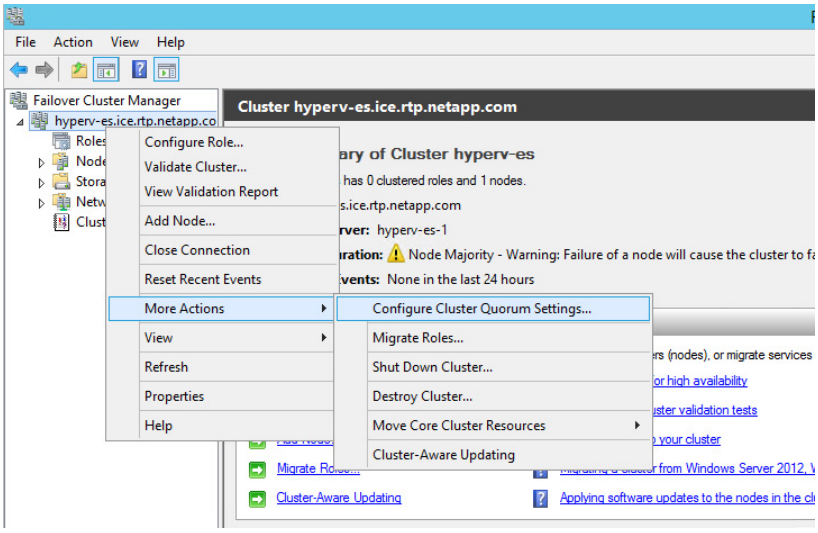

18. Select Add or Change the quorum witness, and click **Next**.
19. Select Configure a disk witness, and Click **Next**.
20. Select Disk Q: from available storage and click **Next**.

21. Click **Next** through the confirmation screen and **Finish** at the summary screen.

Create CSV LUN for VM Storage

One Server Only

1. Open a PowerShell prompt and move the Available Storage cluster group by running.

```
Move-ClusterGroup "Available Storage" -Node $env:COMPUTERNAME | Start-ClusterGroup
```

2. Open SnapDrive from the start screen to configure cluster storage.
3. From SnapDrive, Open the Server name, then Open the Disks Icon.
4. Right-click the Disks Icon and choose Create Disk.
5. Type in the IP Address of the controller that contains the infra CSV Volume.
6. Once connected, open the controller tree and select the infra CSV Volume.
7. Enter the name of the LUN in the LUN name field, click **Next**.
8. Select Shared (Microsoft Cluster Services only) and click **Next**.
9. Validate that all nodes of the cluster are shown and click **Next**.
10. Select Do not assign a Drive letter or Volume Mount Point, set the LUN size to be 500GB and click **Next**.
11. Click **Next** through the Volume properties confirmation.
12. Select the FCP WWPN to Map the LUN to click **Next**.
13. Select Automatic igroup management and Click **Next**.
14. Select Add to cluster shared volumes, and click **Next**.
15. Click Finish.

Validated the Cluster

Run the cluster validation wizard to verify that the cluster is operating correctly.

1. Open Failover Cluster Manager.
2. Click **Validate Cluster…** in the action pane.
3. Proceed through the wizard, and select the option to run all tests.
4. Review and correct any failures that are listed in the validation report.

Deploying a Virtual Machine with VM-FEX

Create Virtual Machines

Now the Windows Failover Cluster created and configured highly available virtual machines can be added to the cluster.

1. Open Failover Cluster Manager.
2. In the left-hand tree branch view, select **Roles > Virtual Machines… > New Hard Disk…**
3. Select the target cluster node and click **OK**.
4. Select the VHDX format for the virtual hard drive and click **Next**.
5. Select Fixed size for the virtual disk type and click **Next**.
6. Provide the virtual hard disk name and path the CSV volume previously created. The CSV volume is located at the path C:\ClusterStorage\. Click **Next** to proceed to the next screen.
7. Select 60GB for the virtual disk size and click **Next**.
8. Click **Finish** to create the virtual disk drive.
9. In the left-hand tree branch view for Cluster Failover Manager, select **Roles > Virtual Machines… > New Virtual Machine…**
10. Select the target cluster node and click **OK**.
11. Enter the virtual machine name.
12. Check the checkbox for storing the virtual machine in a different location.
13. Select the path to the location on the CSV volume where the virtual hard disk was created and click **Next**.

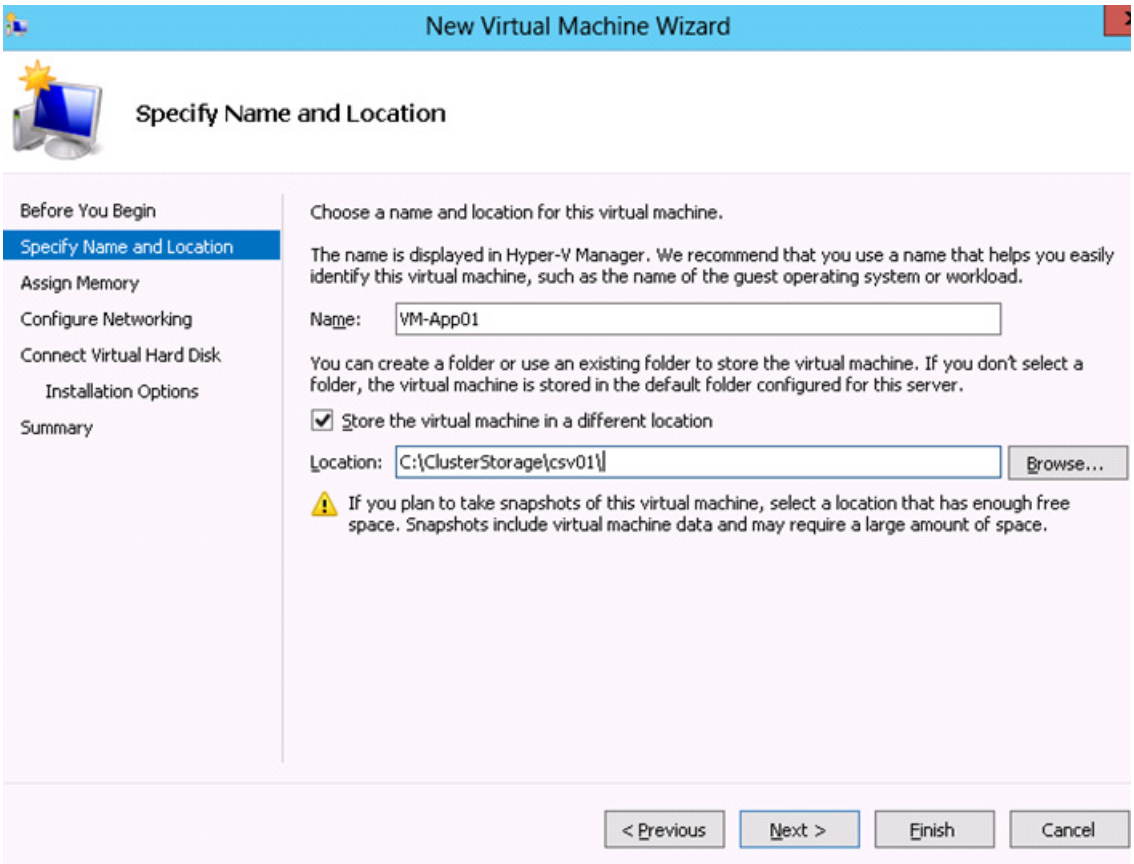

14. Set the startup memory to 8192MB or any value that is suitable for application that will run the virtual machine. Click **Next** to continue.
15. Connect the network adapter to the network switch VM-Public and click **Next**.
16. Select Use and existing virtual disk and provide the path to the virtual disk created in the previous steps. Click **Next** to continue.
17. Review the summary screen and click **Finish** to create the virtual machine.

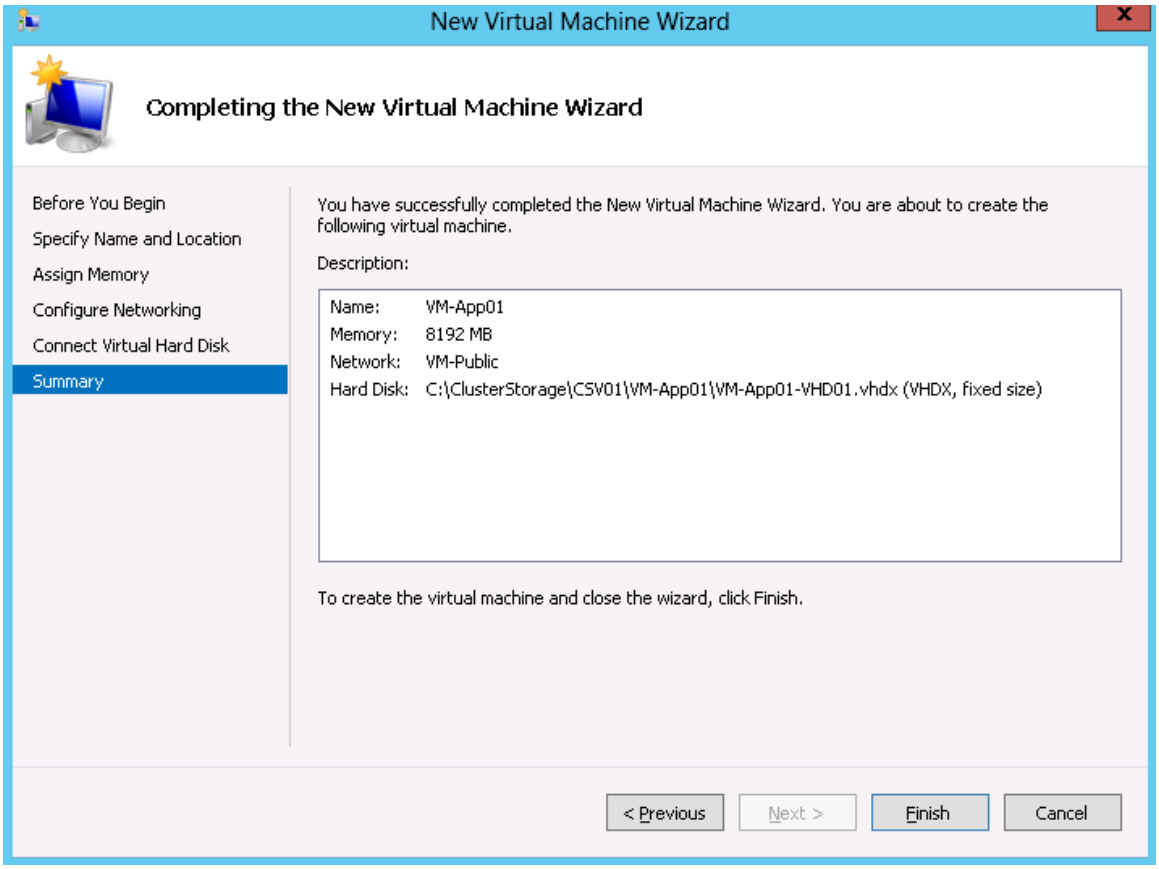

18. Click **Finish** again to close the wizard summary screen.
19. In Failover Cluster Manager Center view, select the newly created virtual machine and Settings… in the action pane.
20. Click **Processor** in the left view and select the required number of virtual processors for the application that will run in the virtual machine.
21. Click **Network Adapter** in the left pane and check the box Enable Virtual LAN identification.
22. Enter the VLAN ID for the VM-Public virtual network.
23. Click **Add Hardware** in the left pane and select Network Adapter in the right pane.
24. Select PF-iSCSI-A in the virtual switch dropdown box.
25. Leave the Enable Virtual LAN identification checkbox unchecked.
26. Expand the network adapter VF-iSCSI-A in the left pane and select Hardware Acceleration.
27. Check the checkbox Enable SR-IOV and uncheck the checkbox Enable IPSec Task offloading.

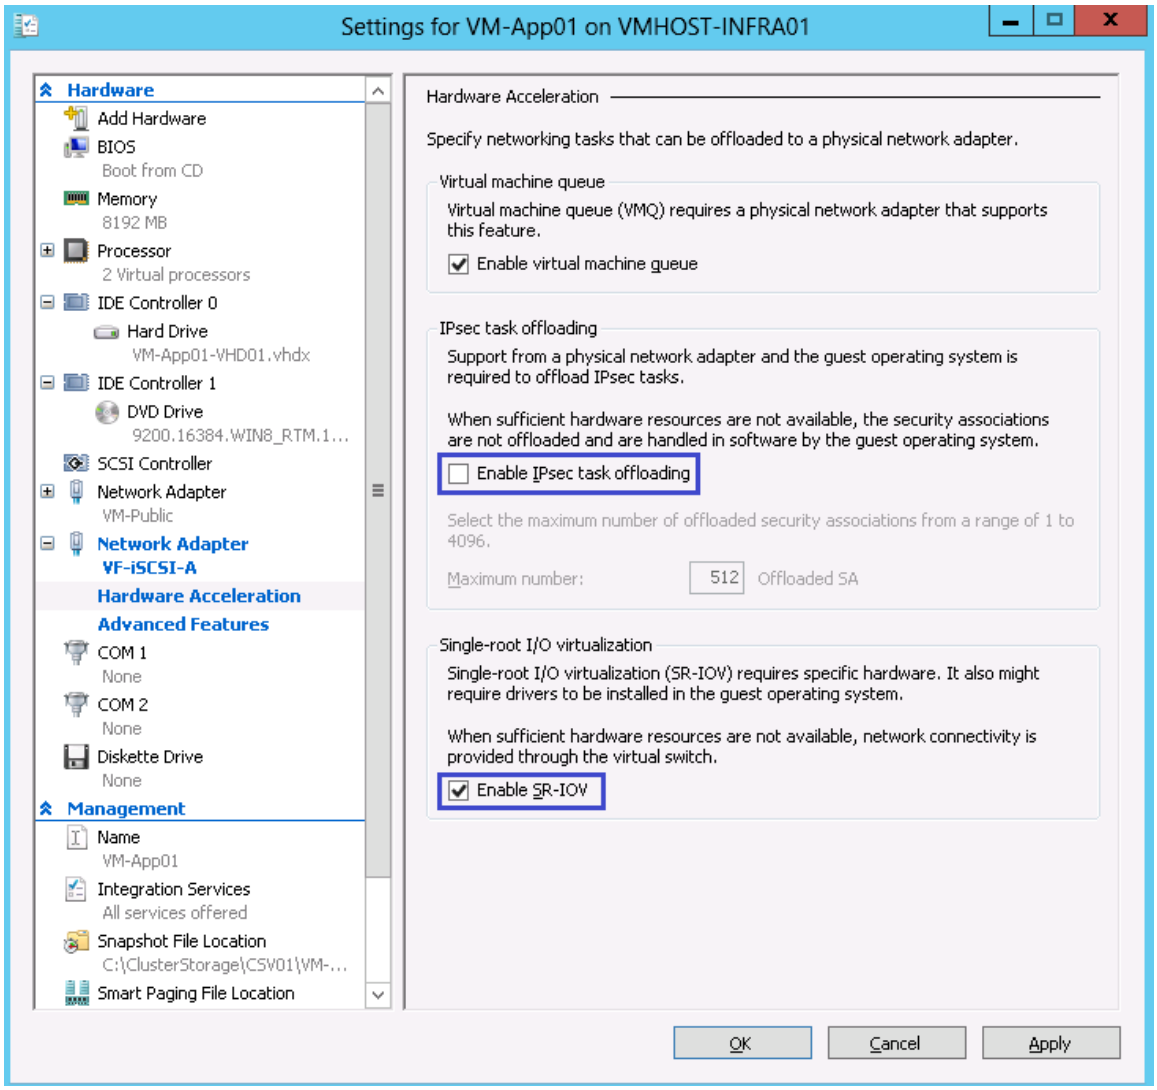

28. Click **Apply**.
29. Repeat step 23 through step 28 to create the network adapter VF-iSCSI-B.
30. Click Automatic Start Action in the left view and select the option Automatically Start if it was running when the service stopped.
31. Click **OK** to save the virtual machine settings.

Attach Port Profile to the Virtual Machine

1. Run the Cisco VM-FEX Port –Profile Manager Utility located at C:\Program Files\Cisco Systems\VIO Software\Utilities\Ethernet Utilities\Vmfex Utilities\Snapin.
2. Select CiscoVMFEX Port-Profile Manger in the left tree view and click Add UCSM in the action pane.
3. Enter the UCS Manager IP address, User Name and Password. Click **OK**.
4. Expand Vmfex Port Profile Manager in the left tree pain and click the UCS Manager instance added in the previous step.

5. Click **Add Host** in right action pane.

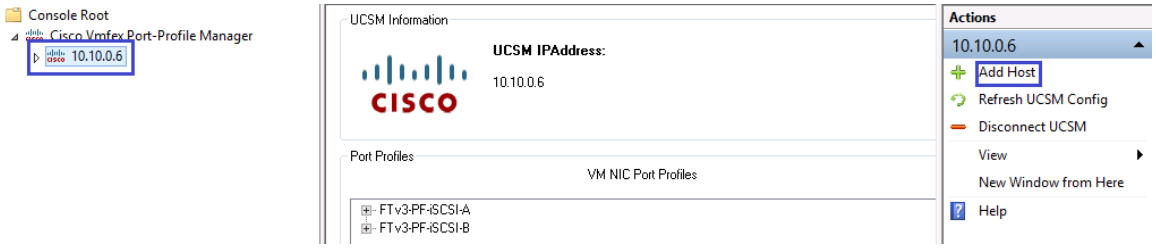

6. Select Local Computer and Click **OK**.

7. The local host is added in the tree view pane on the left. Select the host in the right tree view pane to view the status in the middle pane.

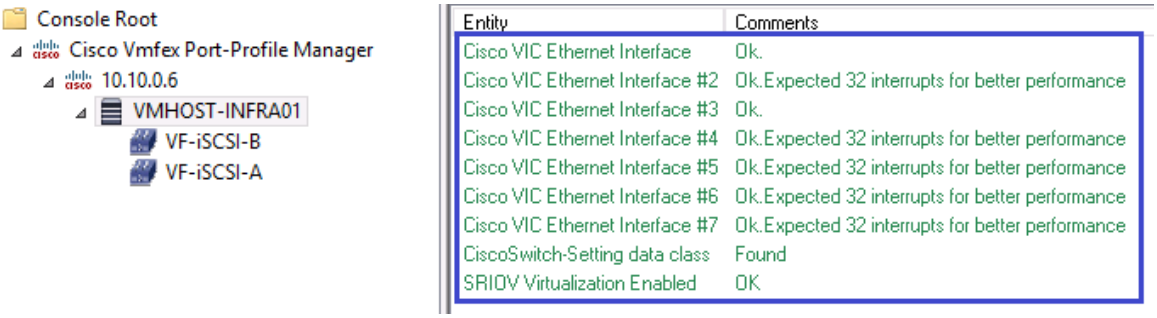

8. Select the network adapter in the left tree view pane. The adapter entry in the middle pane indicates in red text indicates that this adapter is not been attached to a port profile. Click **Attach to Cluster** in the right action pane.

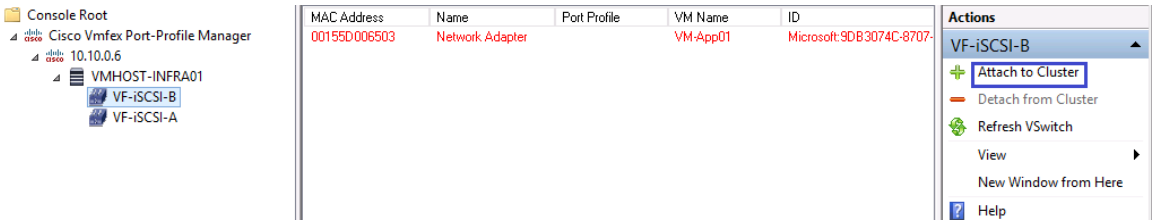

9. Select the logical switch for the corresponding adapter VF-iSCSI-B and click **OK**.

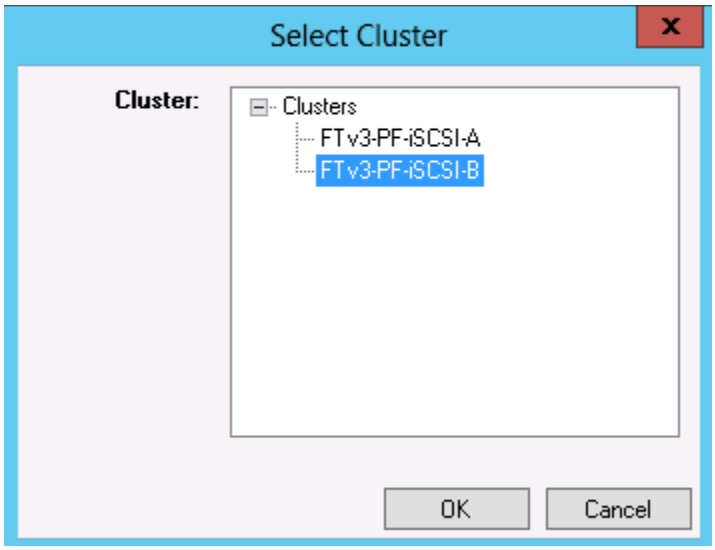

10. Click **Device** in the middle pane and click **Attach/Modify Port Profile** in the Actions pane.

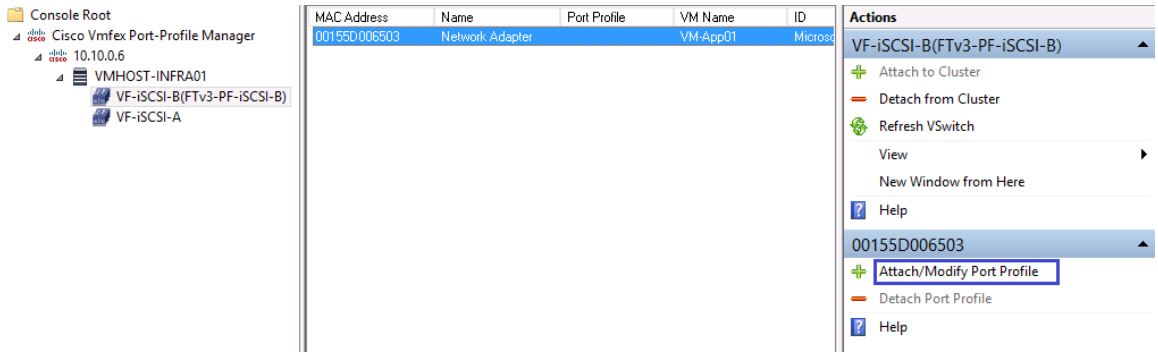

11. Select the port profile and click **OK**.

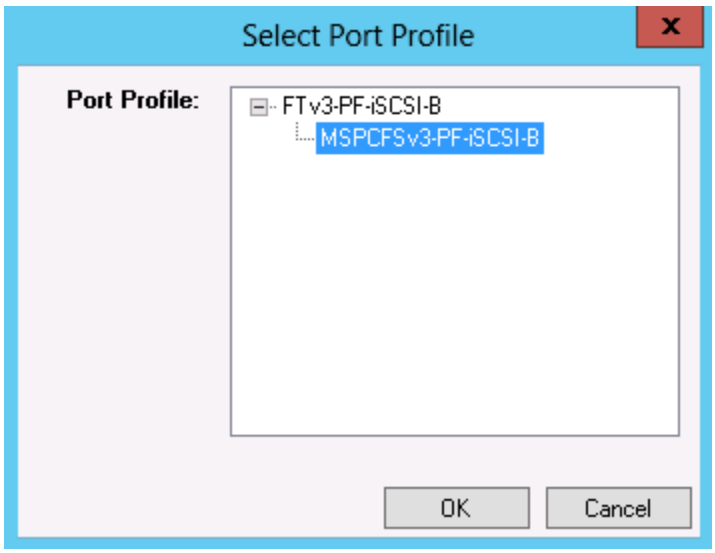

12. Check the middle pane to verify the adapter with the attached port profile in green text.

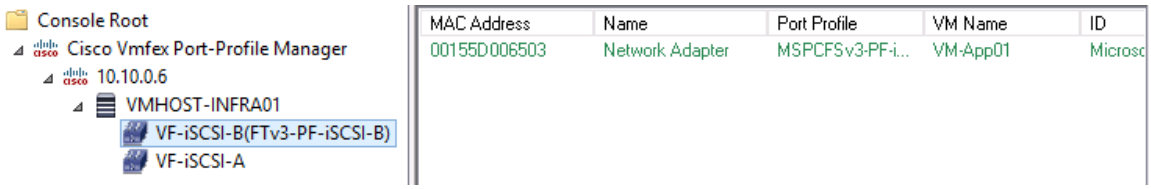

13. Repeat steps 8 through 12 for the second port profile VF-iSCSI-A.

Install Windows and Cisco VF VIC Driver

1. Start the virtual machine and Install Windows in the virtual machine using the Windows installation media or any other suitable means.
2. Log into Windows.
3. In Server Manager Select **Tools** > **Computer Management**.
4. Open Device Manager to view the Virtual Function network adapter without a driver.
5. Right-click the network controller and select Update Driver Software.
6. Install the Cisco VIC Driver.

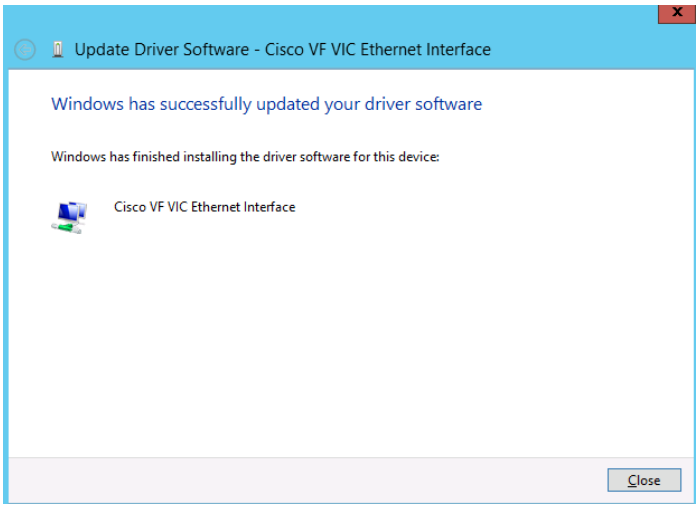

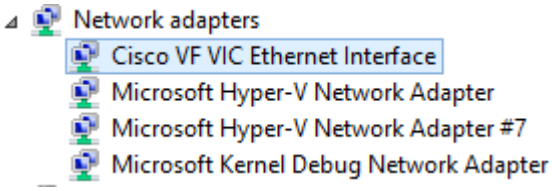

7. Repeat steps 4 through 6 for the second Network Controller under Other Devices.
8. Open Failover Cluster Manger.
9. Select Roles in the tree view and select the previously created virtual machine.
10. Select Settings… in the action pane.
11. Expand the each Network Adapter in the left had device list and select Advanced Features.
12. Record the MAC address for each network adapter.

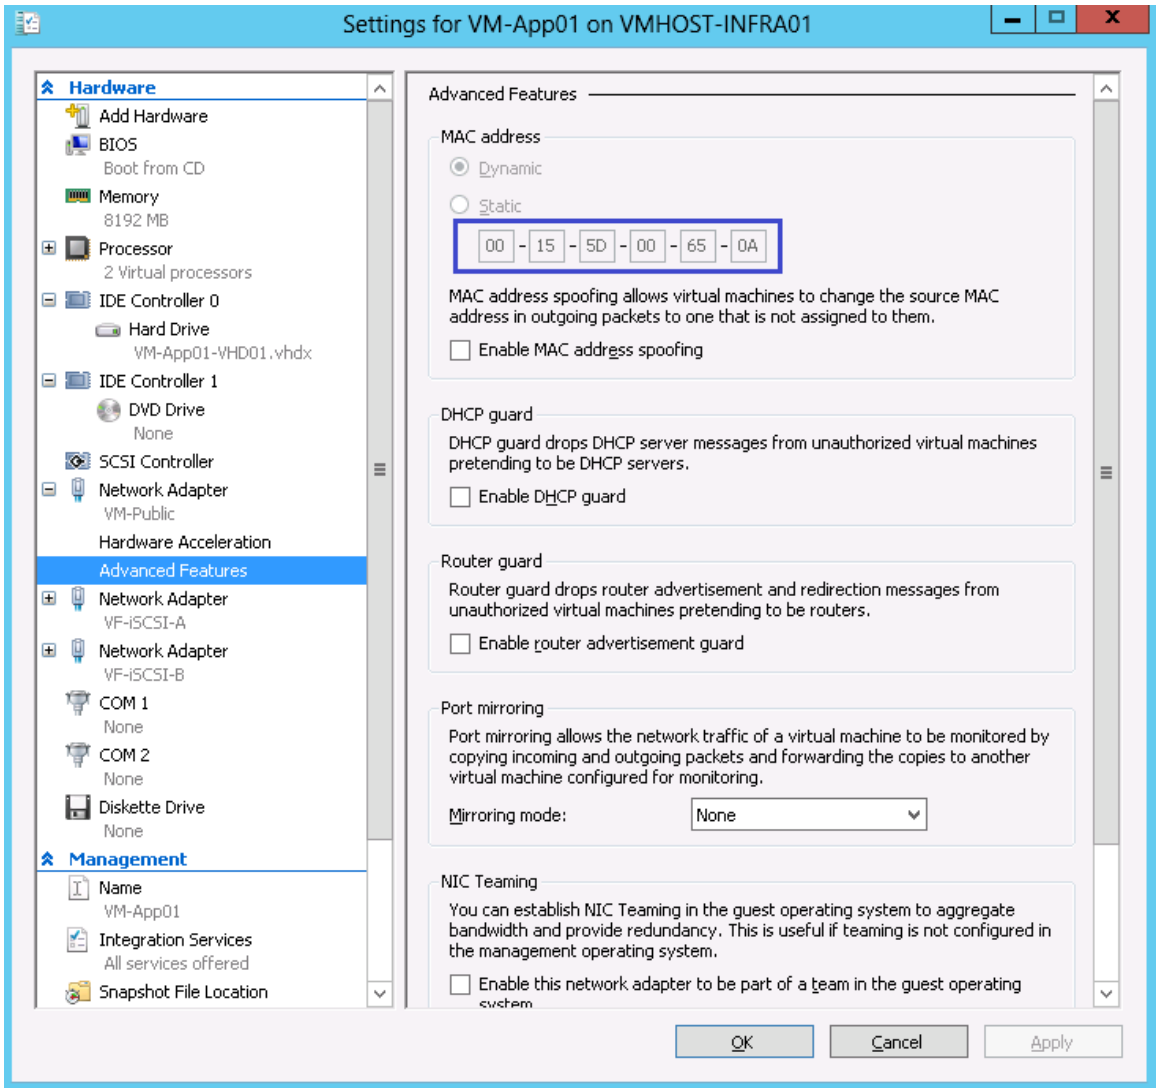

13. Click **Cancel** to close the Settings Window for the virtual machine.
14. Open the Network Connections windows and rename the LAN adapter to reflect the network it is associated with.
15. Set the appropriate IP settings for each adapter.
16. In the iSCSI-A and iSCSI-B Properties window uncheck the following

- Client for Microsoft Networks
- File and Printer Sharing for Microsoft Networks

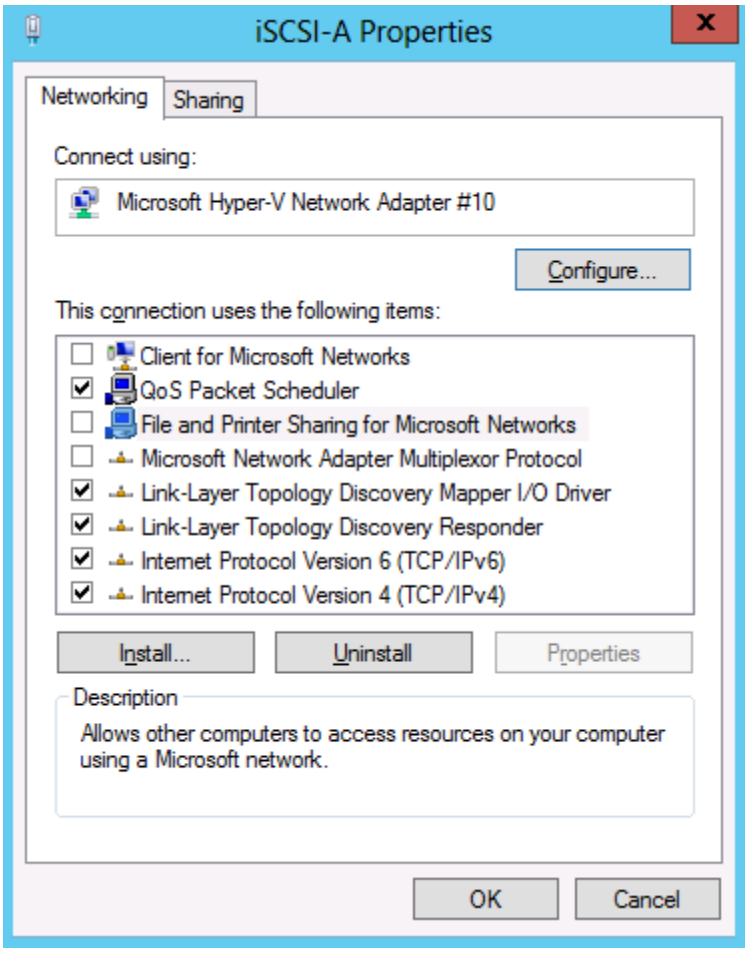

17. Click **OK** to close the window.

Install Windows Features in the Virtual Machine

1. Set the appropriate IP settings for each adapter.
2. Verify that the Windows installation disk is mapped to D: drive.
3. Launch a PowerShell prompt by right clicking the PowerShell icon in the taskbar, and selecting **Run as Administrator**.
4. Add the .Net 3.5 feature by entering the following command:

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

5. Add MPIO by entering the following command:

```
Add-WindowsFeature Multipath-IO -IncludeManagementTools -Restart
```

Configure Windows Host iSCSI initiator

The following steps describe how to configure the built in Microsoft iSCSI initiator.

All Hosts

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 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 RR
```

6. Configure an iSCSI target for each controller.

```
New-IscsiTargetPortal -TargetPortalAddress <<var_controller1_iscsia_ip>>
-InitiatorPortalAddress <iscsia_ipaddress>
New-IscsiTargetPortal -TargetPortalAddress <<var_controller2_iscsia_ip>>
-InitiatorPortalAddress <iscsia_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 Utilities in the Virtual Machine

1. Install NetApp DSM using the procedure in section 7.6
2. Install SnapDrive using the procedure in section 7.12.

Create and Map iSCSI LUNs using SnapDrive

1. Open SnapDrive from the start screen to configure cluster storage.
2. From SnapDrive, Open the Server name, then Open the Disks Icon.
3. Right-click the Disks Icon and select Create Disk.
4. Type in the IP Address of the controller.
5. Once connected, open the controller tree and select the volume.
6. Enter the name of the LUN in the LUN name field, click **Next**.
7. Select Dedicated click **Next**.
8. Change the drive letter or Mount Point, and set the LUN size, click **Next**.
9. Select the iSCSI initiators to Map the LUN to click **Next**.
10. Select Automatic igroup management and click **Next**.
11. Click **Finish**.

Appendix A

Installing Cisco UCS PowerTool

The Cisco UCS PowerTool should be installed on the FlexPod Management server.

Download the Cisco UCS PowerTool version 0.9.9.0 or newer from the Cisco Developer Network. It can be found in the Microsoft Management section:

http://developer.cisco.com/web/unifiedcomputing/microsoft

Extract the zip file and execute the extracted exe file.

Perform the following steps on the FlexPod management server:	
Launch the Cisco UCS PowerTool Installer. The **Setup Wizard** screen appears.	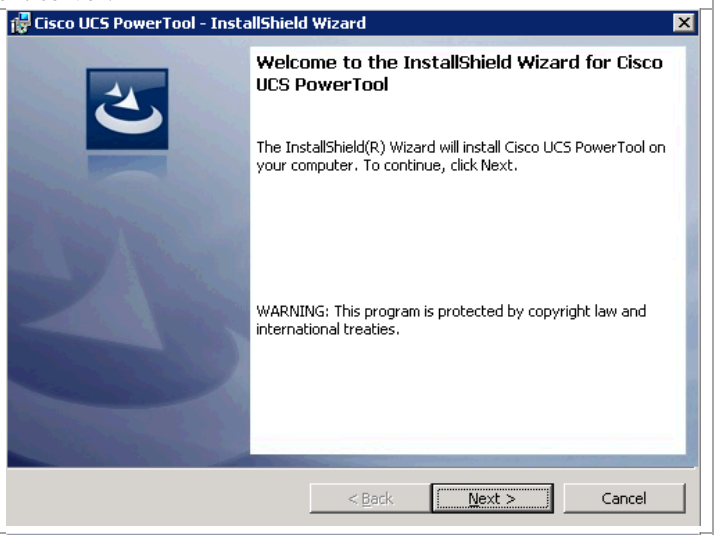
Read and accept the end user license agreement. Click **Next** to continue.	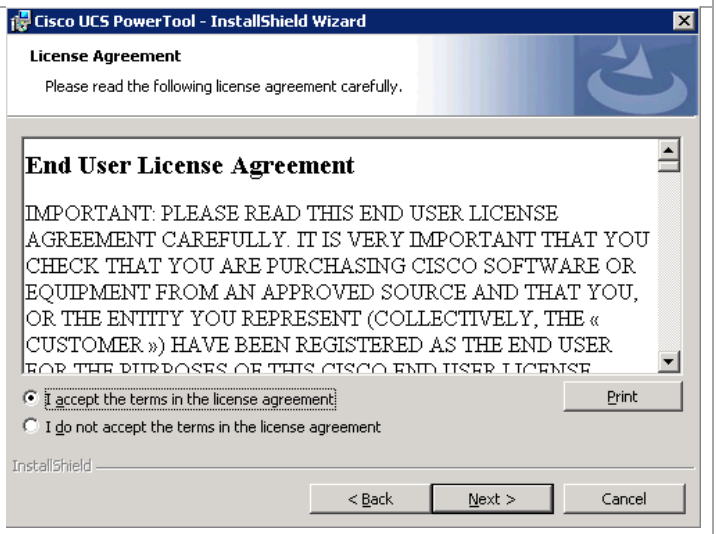
Verify proper PowerTool installation by launching the PowerTool console. Click **Start** and **All Programs**. Expand **Cisco** and select **Cisco UCS PowerTool**.	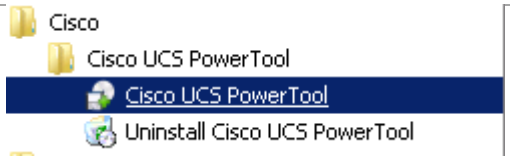

Select the **Destination Folder** and click **Next** to continue.	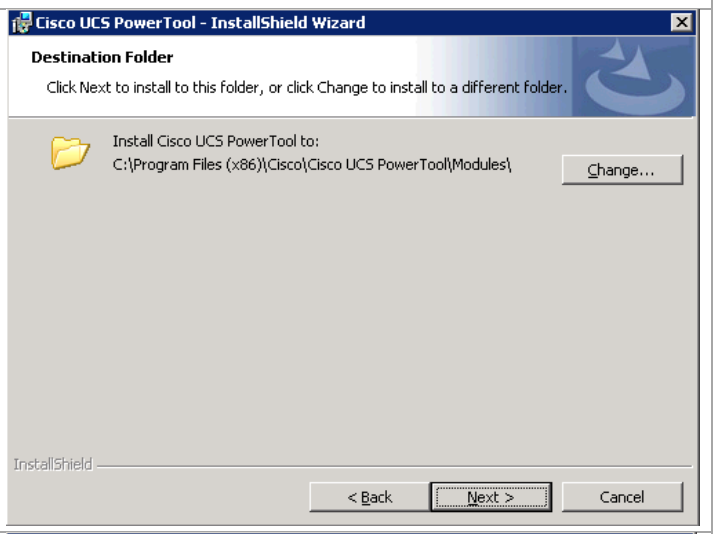
Cisco UCS PowerTool is ready to install. Click **Next** to complete the installation.	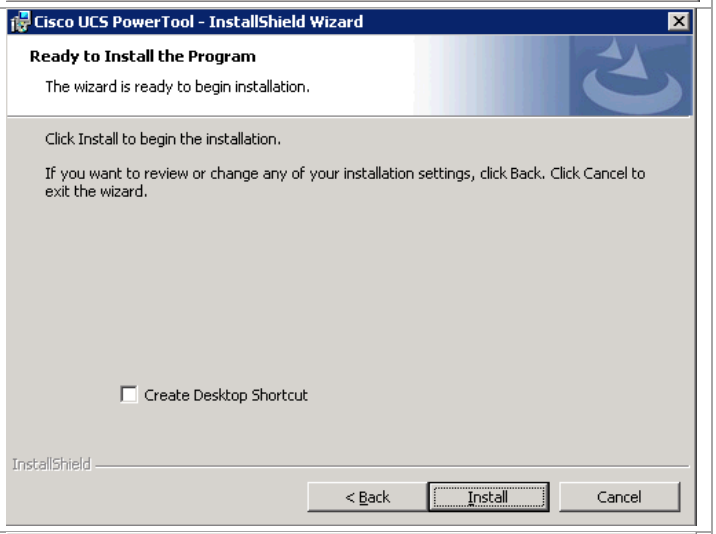
After the installation completes successfully click **Finish** to close the installation wizard.	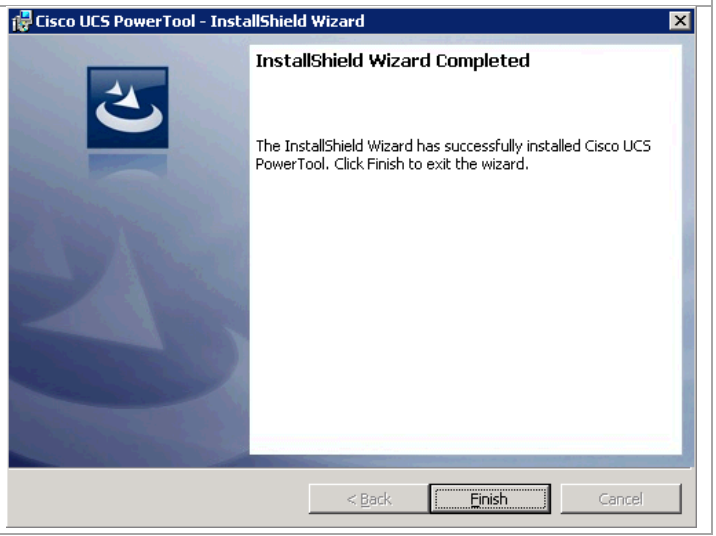

Enter the command: **Connect-ucs <FI cluster FQDN or IP>**	Administrator: Cisco UCS PowerTool C:\Program Files (x86)\Cisco\Cisco UCS PowerTool\Modules\CiscoUcsPS>C:\Windows\System32\windowspowershell\v1.0\powershell.exe -NoExit -ExecutionPolicy RemoteSigned -File .\StartUcsPS.ps1 Windows PowerShell Copyright (C) 2009 Microsoft Corporation. All rights reserved. PS C:\> Connect-Ucs ucsm.flexpod.test_
Enter the UCS Manager login credentials.	Windows PowerShell Credential Request Enter your credentials. User name: Password: OK Cancel
UCS PowerTool returns the UCS Manager session connection information. This indicates that PowerTool is properly installed and is operational.	Administrator: Cisco UCS PowerTool Proxy : Cookie : 1343256241/8be66ee0-ce3e-4591-94c8-227f2285b209 Domains : LastUpdateTime : 7/25/2012 4:13:44 PM Name : ucsm.flexpod.test NoSsl : False NumPendingConfigs : 0 NumWatchers : 0 Port : 443 Priv : {admin, read-only} RefreshPeriod : 600 SessionId : web_12934_A TransactionInProgress : False Ucs : hyper-v-flexpod-ucs Uri : https://ucsm.flexpod.test UserName : admin VirtualIpv4Address : 10.10.0.6 Version : 2.0(2q) WatchThreadStatus : None PS C:\> _

1. Open the Hyper-V Manager and select the Hyper-V server in the left pane.
2. Click **New** in the right action pane and select Hard Disk.

Appendix B: Installing the DataONTAP PowerShell Toolkit

1. Download the DataONTAP PowerShell toolkit from the NetApp Communities https://communities.netapp.com/community/products_and_solutions/microsoft/powershell
2. Run DataONTAP windows installation package.
3. Click **Next** on the welcome page.
4. Accept the ELUA and click **Next**.
5. Validate the Installation path and click **Next**.

6. Click **Install**.

Appendix C: Creating Domain Controller Virtual Machine (Optional)

Create VHD for Domain Controller Virtual Machine.

Create the following VHD storage resources that will be used by the virtual machines running active directory domain controller roles:

Table 12 VHD Storage Resources

VM Host	VM Name	Name	Location	Size	Type
Infra-VM-Host -01	Infra-DC-01	Infra-DC-01.vhdx	C:\VHD\Infra-DC-01	60GB	Fixed
Infra-VM-Host -02	Infra-DC-02	Infra-DC-01.vhdx	C:\VHD\Infra-DC-02	60GB	Fixed

1. Open the Hyper-V Manager and select the Hyper-V server in the left pane.
2. Click **New** in the right action pane and select Hard Disk.

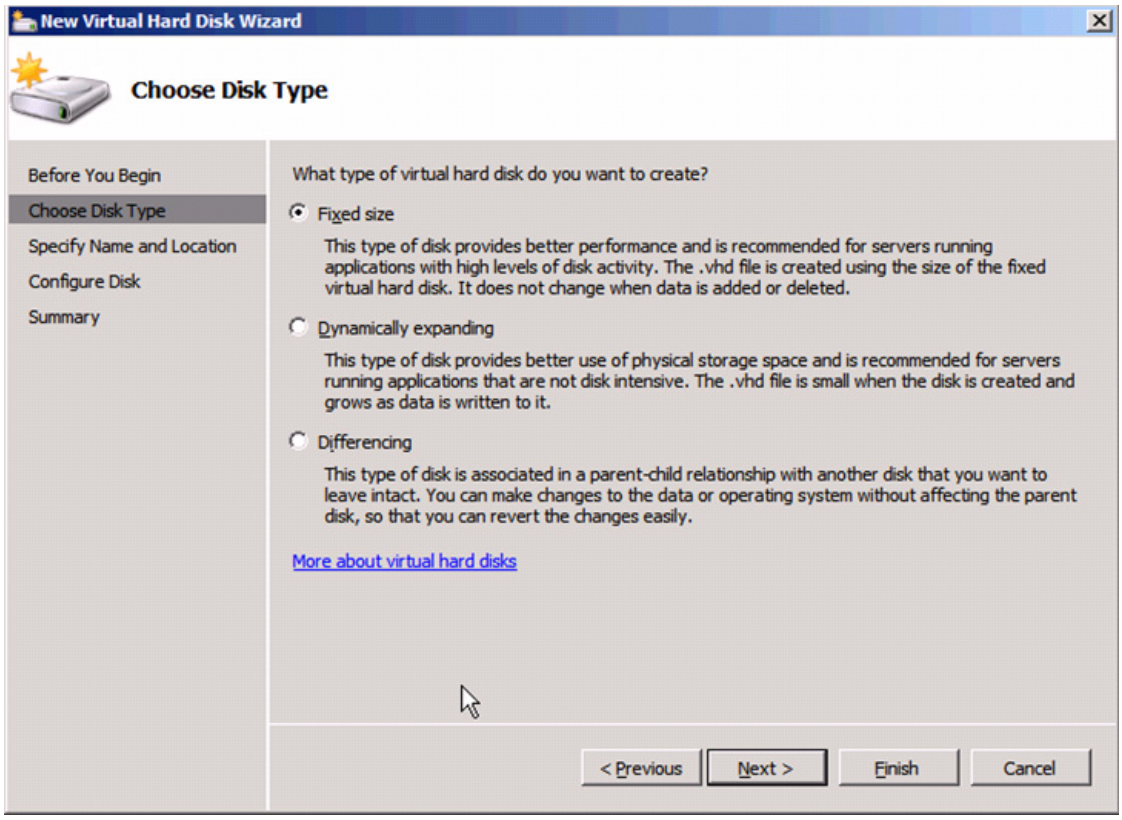

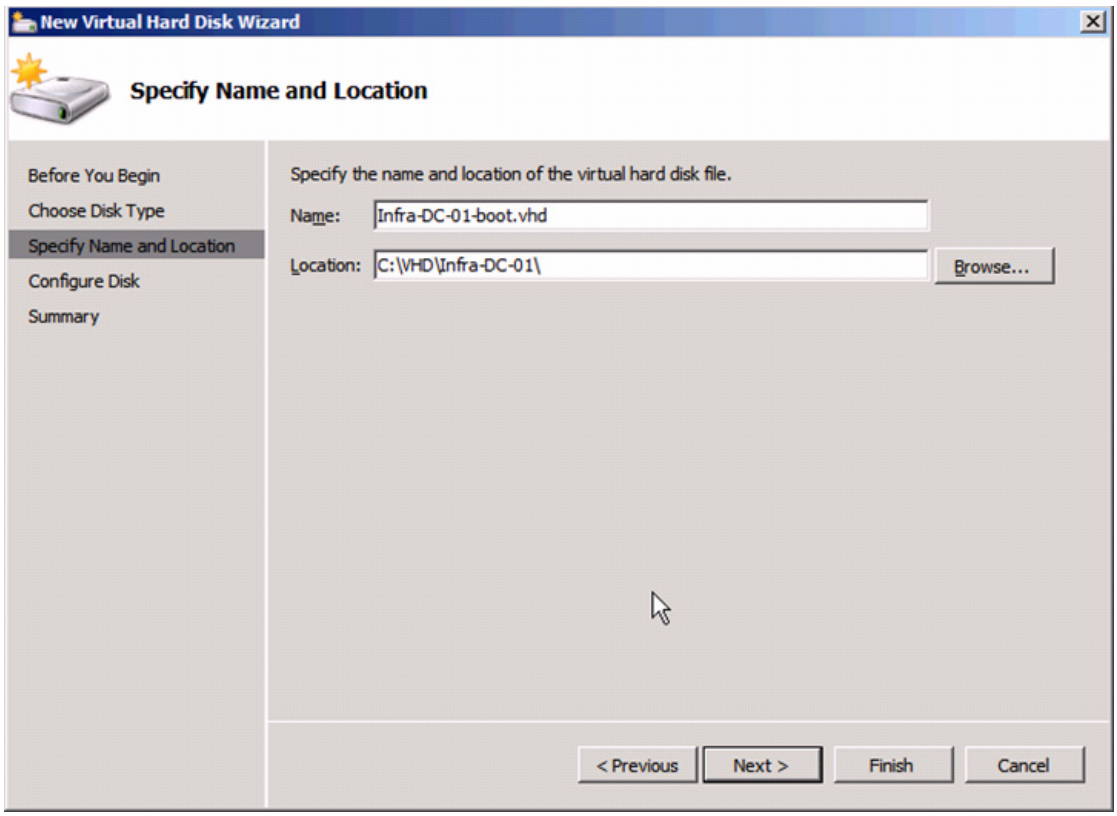
New Virtual Hard Disk Wizard
Specify Name and Location
Before You Begin
Choose Disk Type
Specify Name and Location
Configure Disk
Summary
Specify the name and location of the virtual hard disk file.
Name:
Infra-DC-01-boot.vhd
Location:
C:\VHD\Infra-DC-01\
Browse...
< Previous
Next >
Finish
Cancel

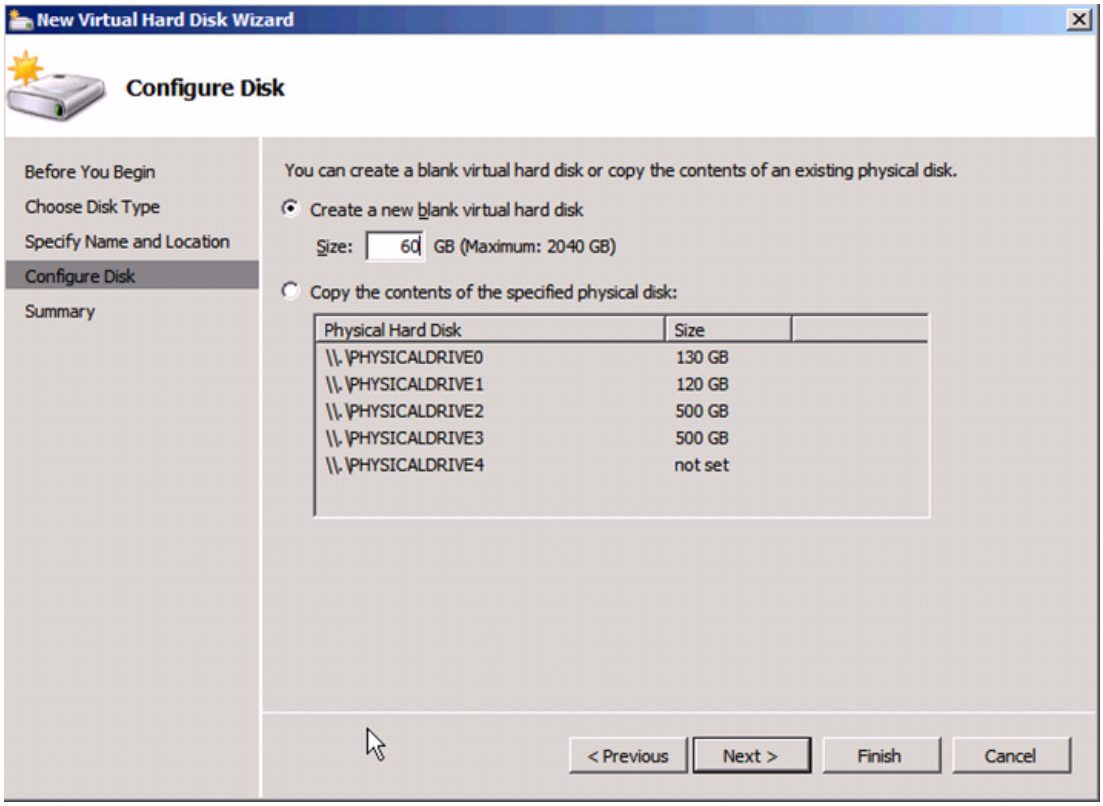
New Virtual Hard Disk Wizard
Configure Disk
Before You Begin
Choose Disk Type
Specify Name and Location
Configure Disk
Summary
You can create a blank virtual hard disk or copy the contents of an existing physical disk.
Create a new blank virtual hard disk
Size:
60
GB (Maximum: 2040 GB)
Copy the contents of the specified physical disk:
Physical Hard Disk
Size
\\.\PHYSICALDRIVE0
130 GB
\\.\PHYSICALDRIVE1
120 GB
\\.\PHYSICALDRIVE2
500 GB
\\.\PHYSICALDRIVE3
500 GB
\\.\PHYSICALDRIVE4
not set
< Previous
Next >
Finish
Cancel

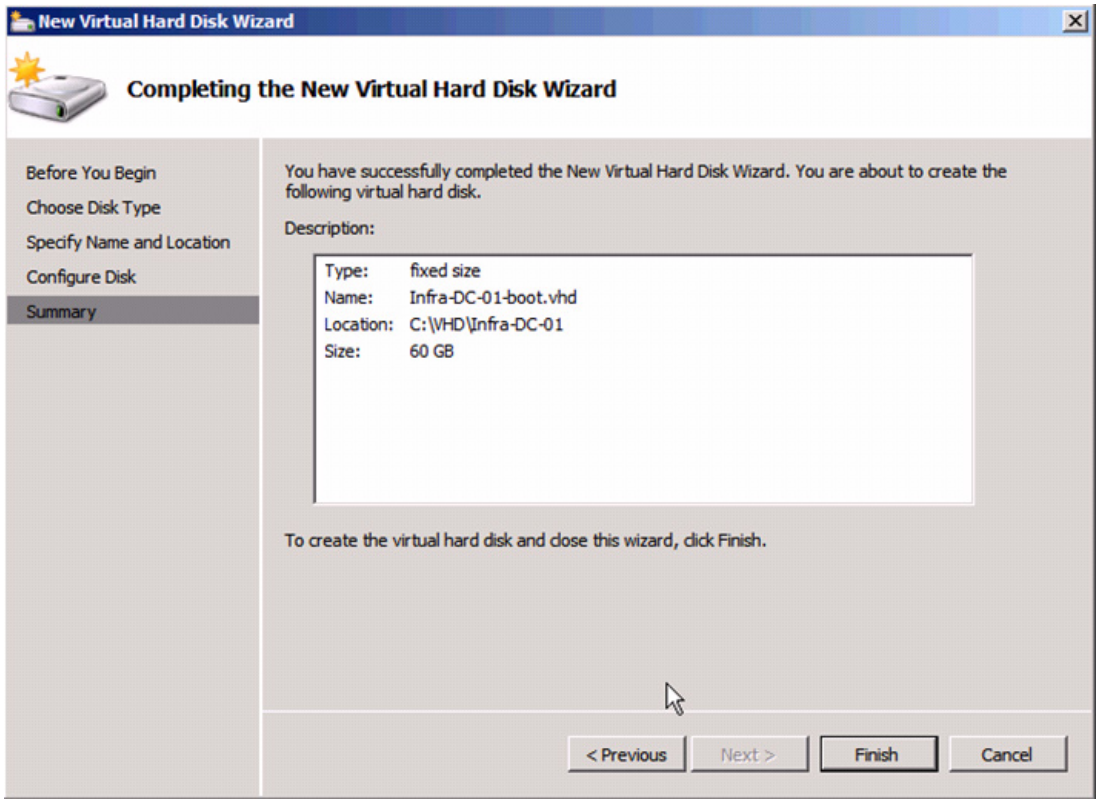
New Virtual Hard Disk Wizard
Completing the New Virtual Hard Disk Wizard
Before You Begin
Choose Disk Type
Specify Name and Location
Configure Disk
Summary
You have successfully completed the New Virtual Hard Disk Wizard. You are about to create the following virtual hard disk.
Description:
Type: fixed size
Name: Infra-DC-01-boot.vhd
Location: C:\VHD\Infra-DC-01
Size: 60 GB
To create the virtual hard disk and close this wizard, click Finish.
< Previous
Next >
Finish
Cancel

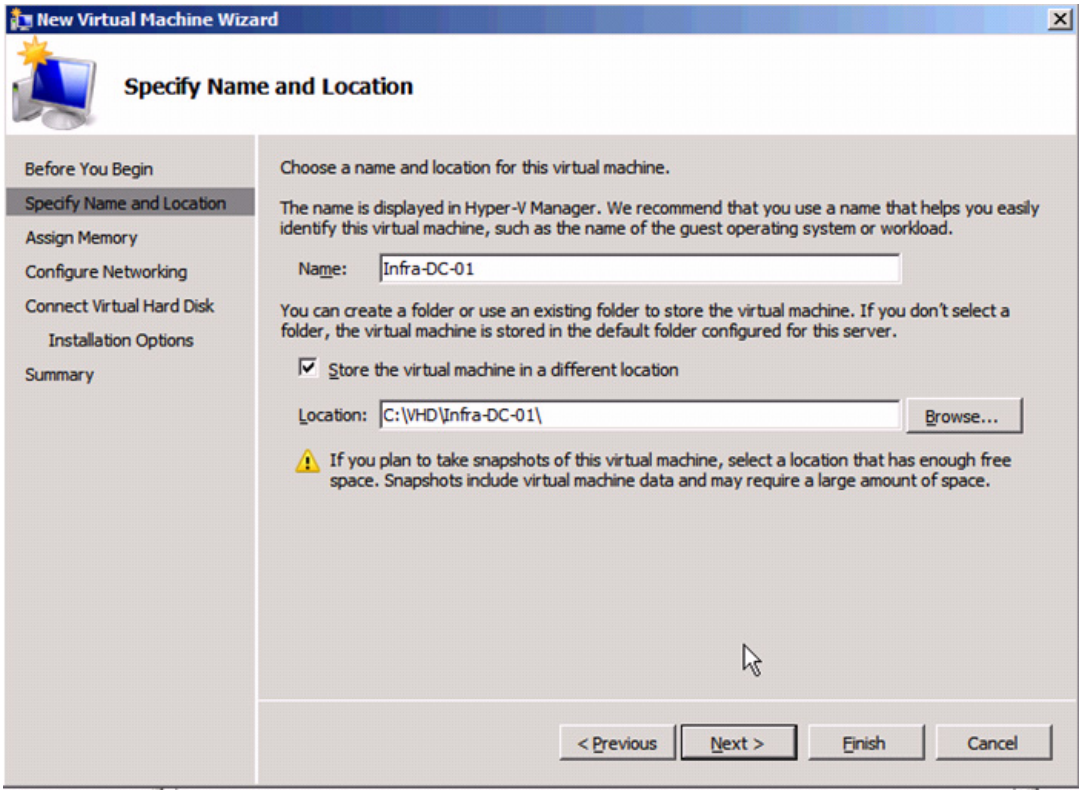
New Virtual Machine Wizard
Specify Name and Location
Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary
Choose a name and location for this virtual machine.
The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.
Name: Infra-DC-01
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: C:\VHD\Infra-DC-01\
Browse...
If you plan to take snapshots of this virtual machine, select a location that has enough free space. Snapshots include virtual machine data and may require a large amount of space.
< Previous
Next >
Finish
Cancel

Create the following virtual machines that will be used by the active directory domain controller roles.

Table 13

VM Host	VM Name	Hard Disk	Network	Memory	VLAN ID
Infra-VM-Host-01	Infra-DC-01	C:\VHD\Infra-DC-01.vhdx	VM-Data-Virtual Network	8GB	804
Infra-VM-Host-02	Infra-DC-02	C:\VHD\Infra-DC-02.vhdx	VM-Data-Virtual Network	8GB	804

1. Open Hyper-V Manager and select the Hyper-V server in the left pane.
2. Click **New** in the right action pane and select Virtual Machine.
3. Provide the name. Check the check box for storing the virtual machine in a different location and provide the path. Click **Next**.
4. Enter the memory size and click **Next**.
5. Select the Network connection VM-Data-Virtual Network. Click **Next**.
6. Select the option to use an existing virtual hard disk and specify the path to the VHD created in the previous section. Click **Next**.
7. Select the option to install the operating system later and click **Finish**.
8. Repeat steps 1 through 7 for each virtual machine.

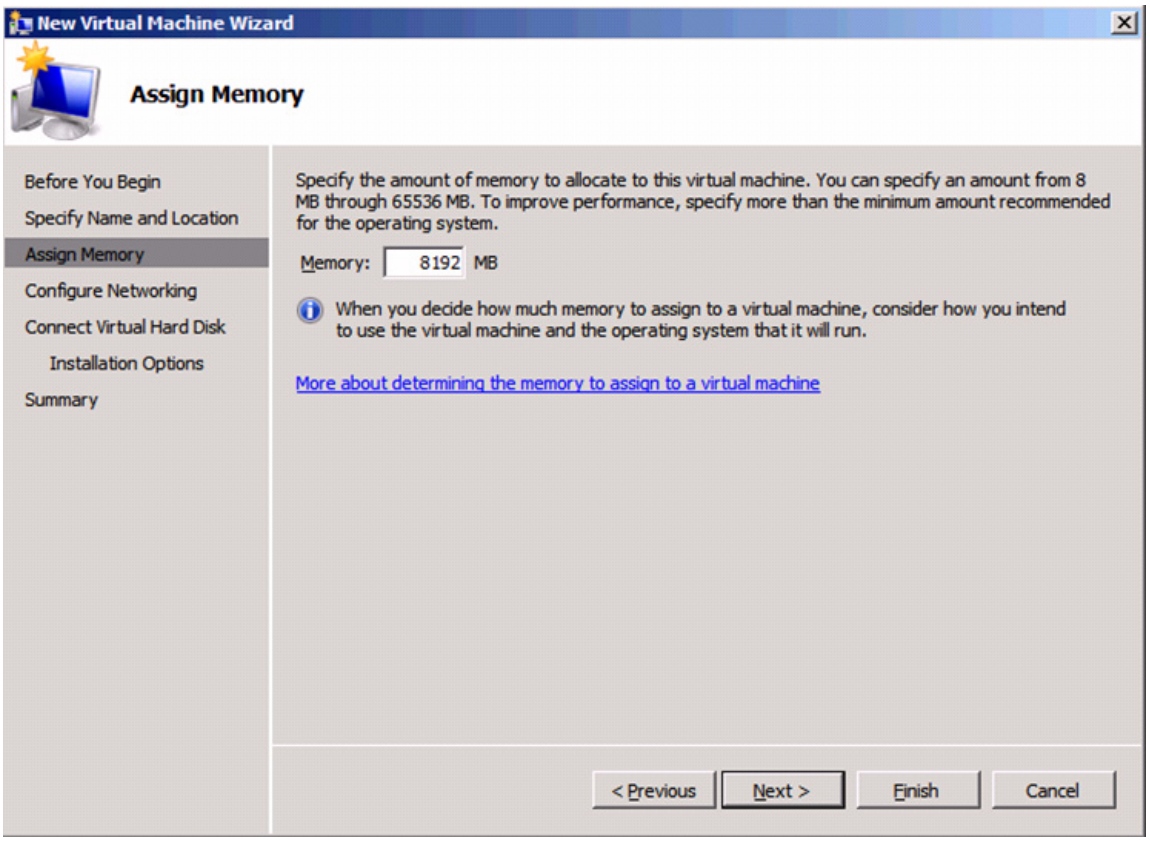

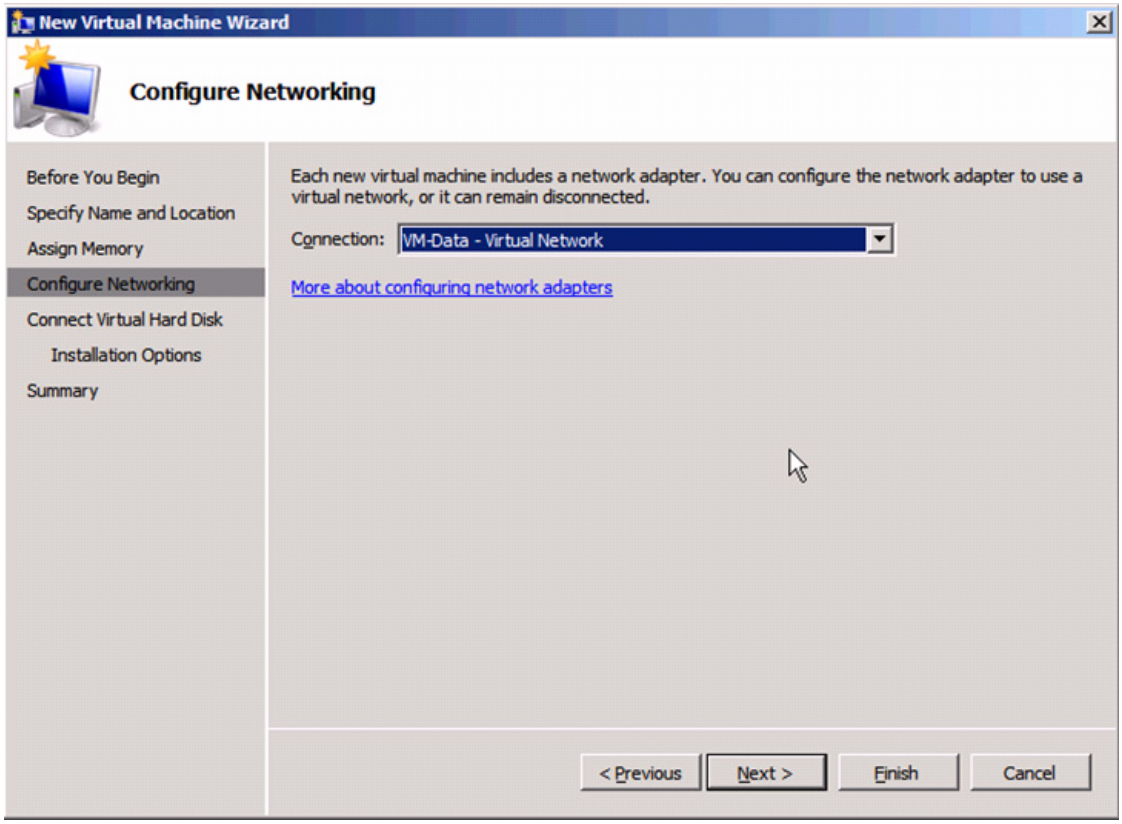
New Virtual Machine Wizard
Configure Networking
Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary
Each new virtual machine includes a network adapter. You can configure the network adapter to use a virtual network, or it can remain disconnected.
Connection:
VM-Data - Virtual Network
More about configuring network adapters
< Previous
Next >
Finish
Cancel

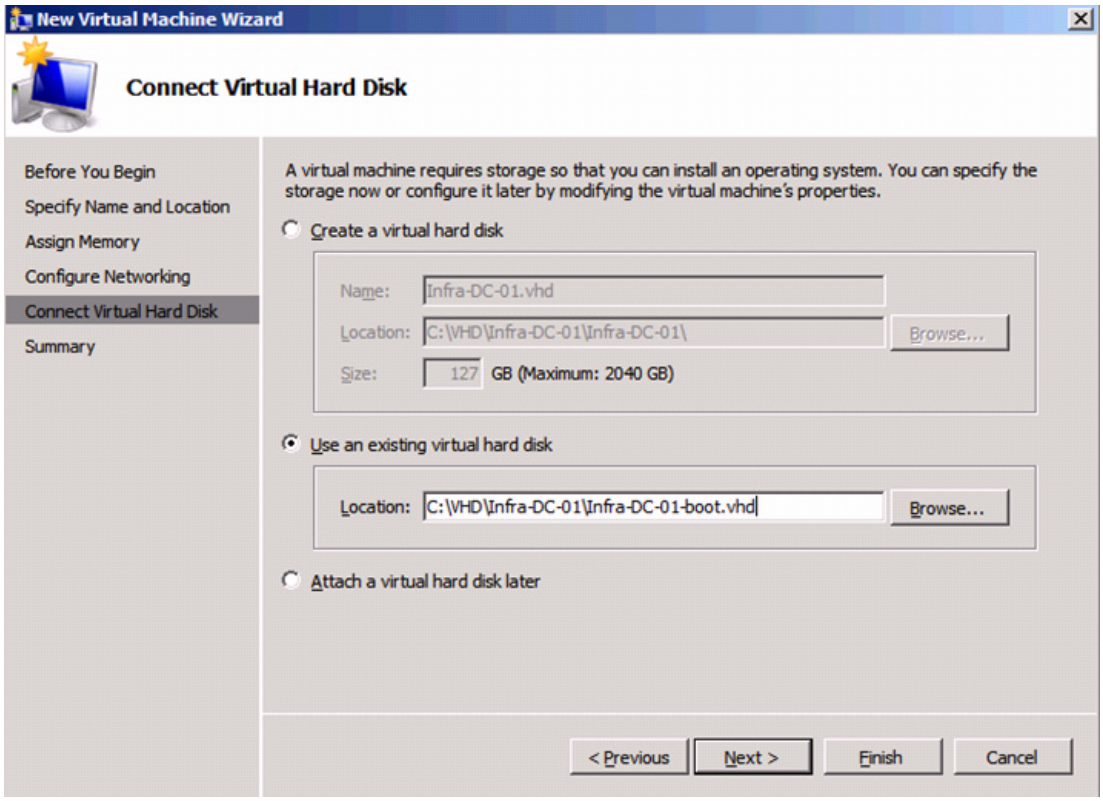
New Virtual Machine Wizard
Connect Virtual Hard Disk
Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Summary
A virtual machine requires storage so that you can install an operating system. You can specify the storage now or configure it later by modifying the virtual machine's properties.
Create a virtual hard disk
Name:
Infra-DC-01.vhd
Location:
C:\VHD\Infra-DC-01\Infra-DC-01\
Browse...
Size:
127
GB (Maximum: 2040 GB)
Use an existing virtual hard disk
Location:
C:\VHD\Infra-DC-01\Infra-DC-01-boot.vhd
Browse...
Attach a virtual hard disk later
< Previous
Next >
Finish
Cancel

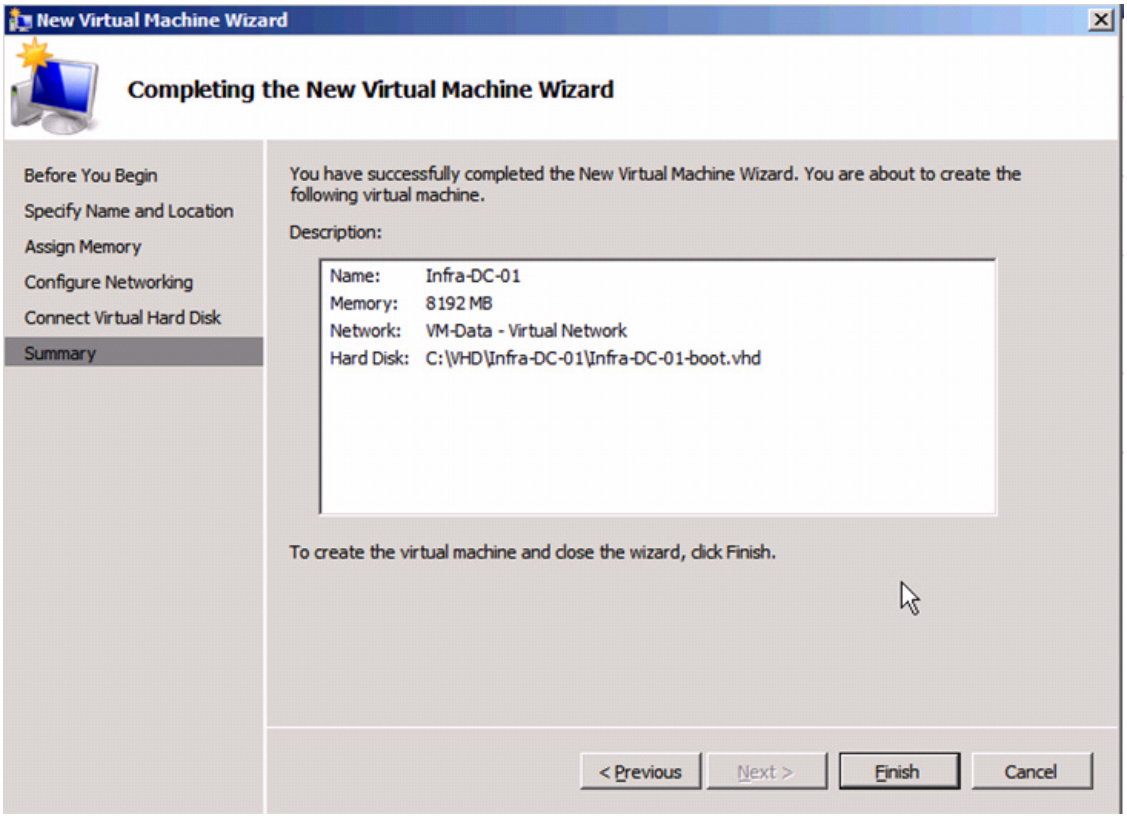
New Virtual Machine Wizard
Completing the New Virtual Machine Wizard
Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Summary
You have successfully completed the New Virtual Machine Wizard. You are about to create the following virtual machine.
Description:
Name: Infra-DC-01
Memory: 8192 MB
Network: VM-Data - Virtual Network
Hard Disk: C:\VHD\Infra-DC-01\Infra-DC-01-boot.vhd
To create the virtual machine and close the wizard, click Finish.
< Previous
Next >
Finish
Cancel

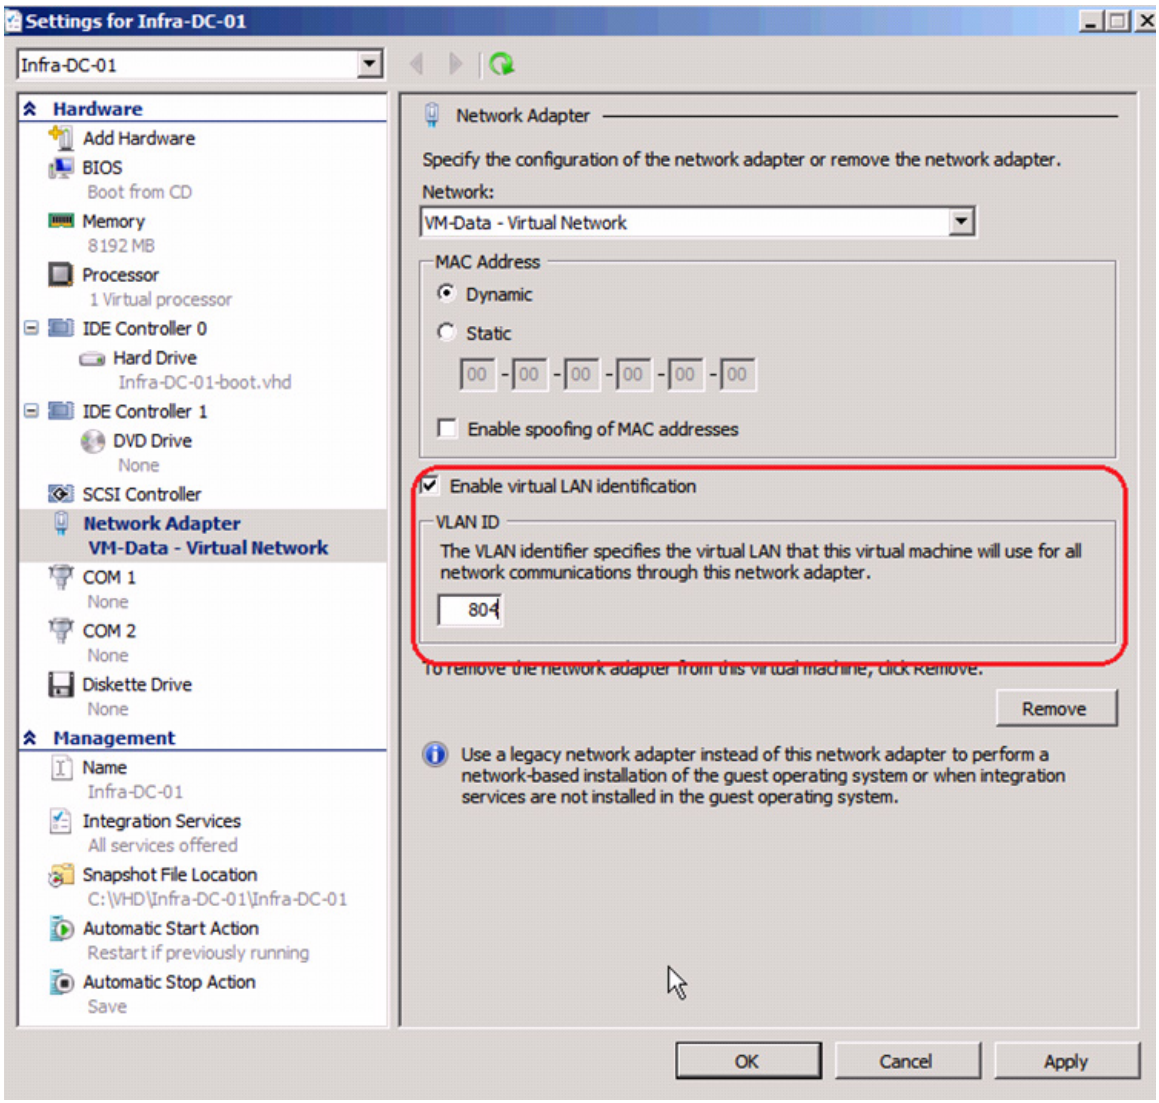

Install Windows in a Domain Controller Virtual Machine

Follow the screenshots to install windows in a Domain Controller VM.

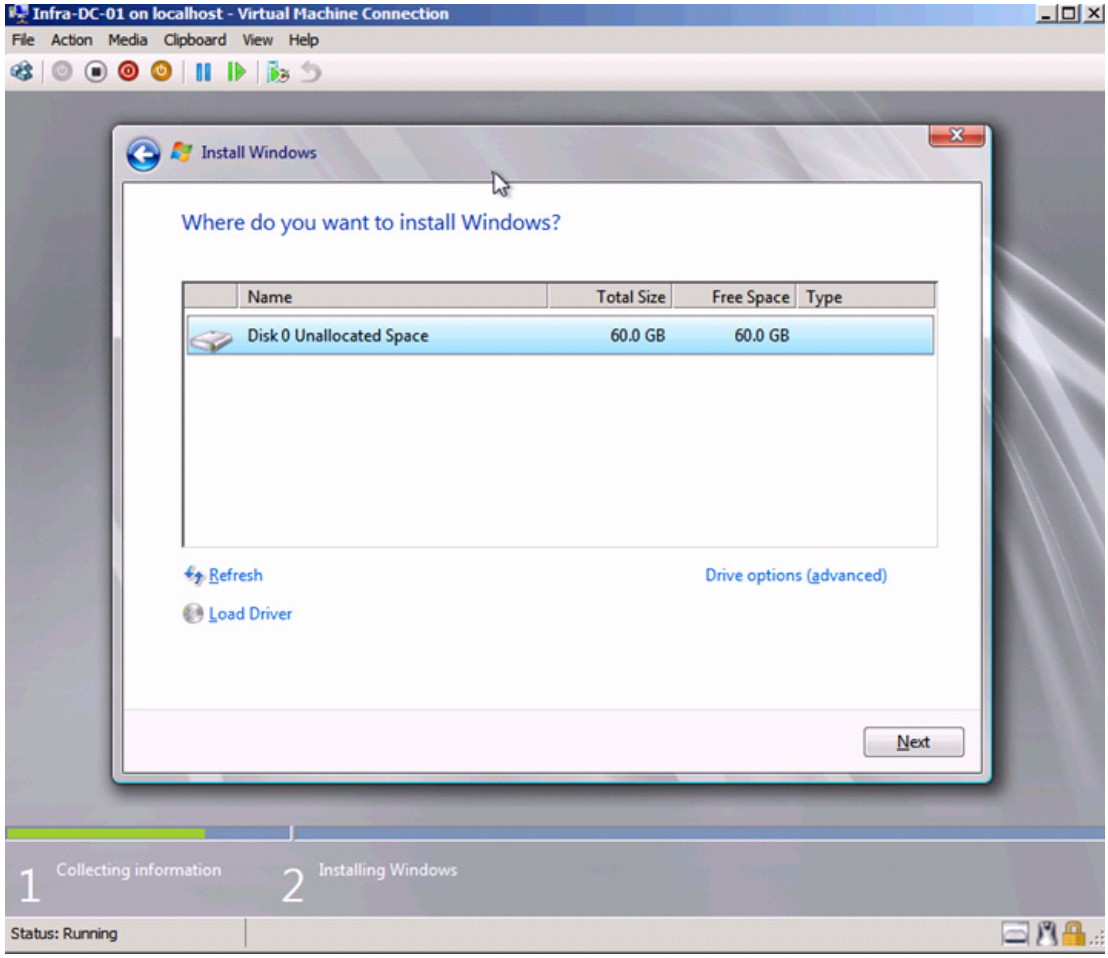
Infra-DC-01 on localhost - Virtual Machine Connection
File Action Media Clipboard View Help
Install Windows
Where do you want to install Windows?
Name
Total Size
Free Space
Type
Disk 0 Unallocated Space
60.0 GB
60.0 GB
Refresh
Load Driver
Drive options (advanced)
Next
1 Collecting information
2 Installing Windows
Status: Running

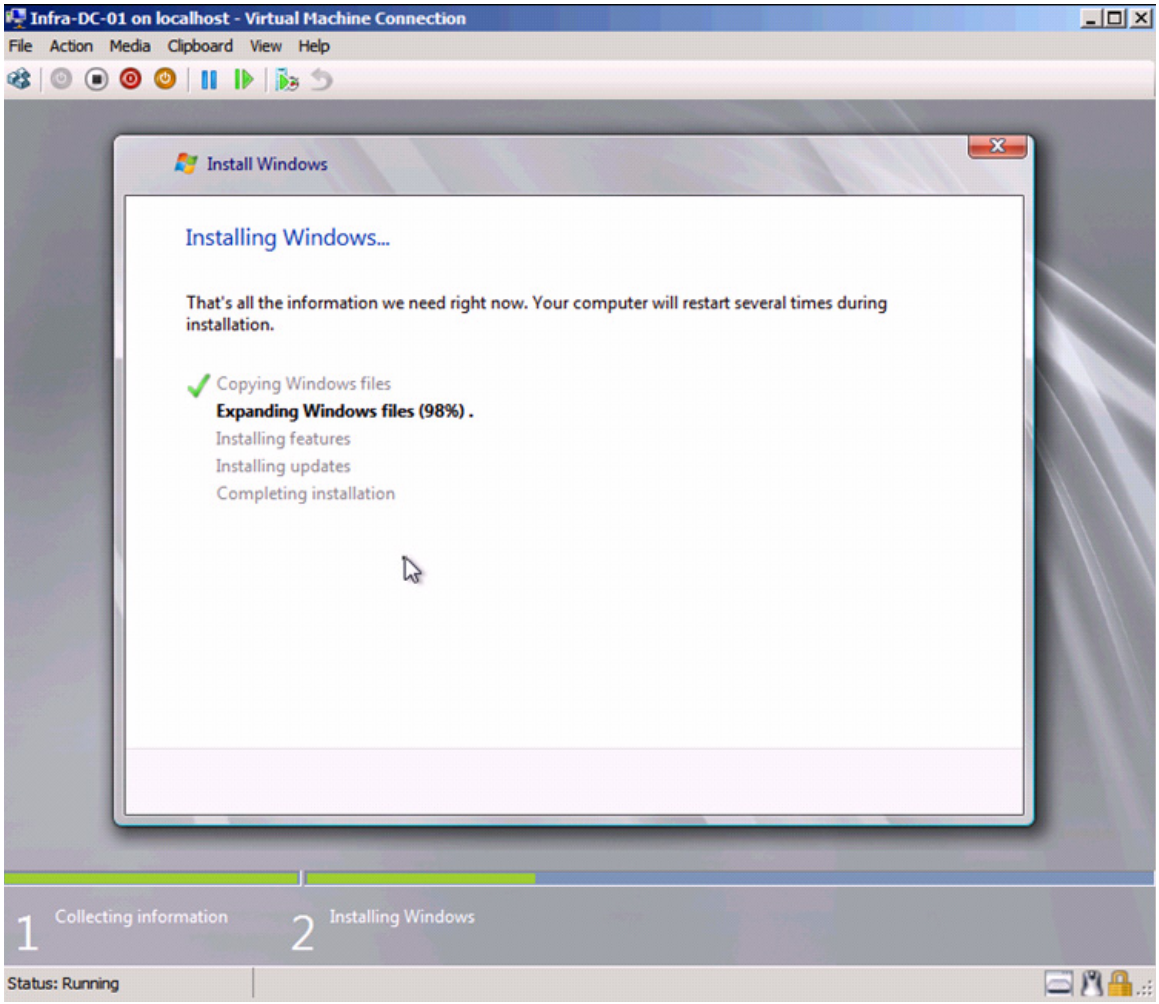

Install Active Directory Services

Follow the screenshots to install windows in a Domain Controller VM.

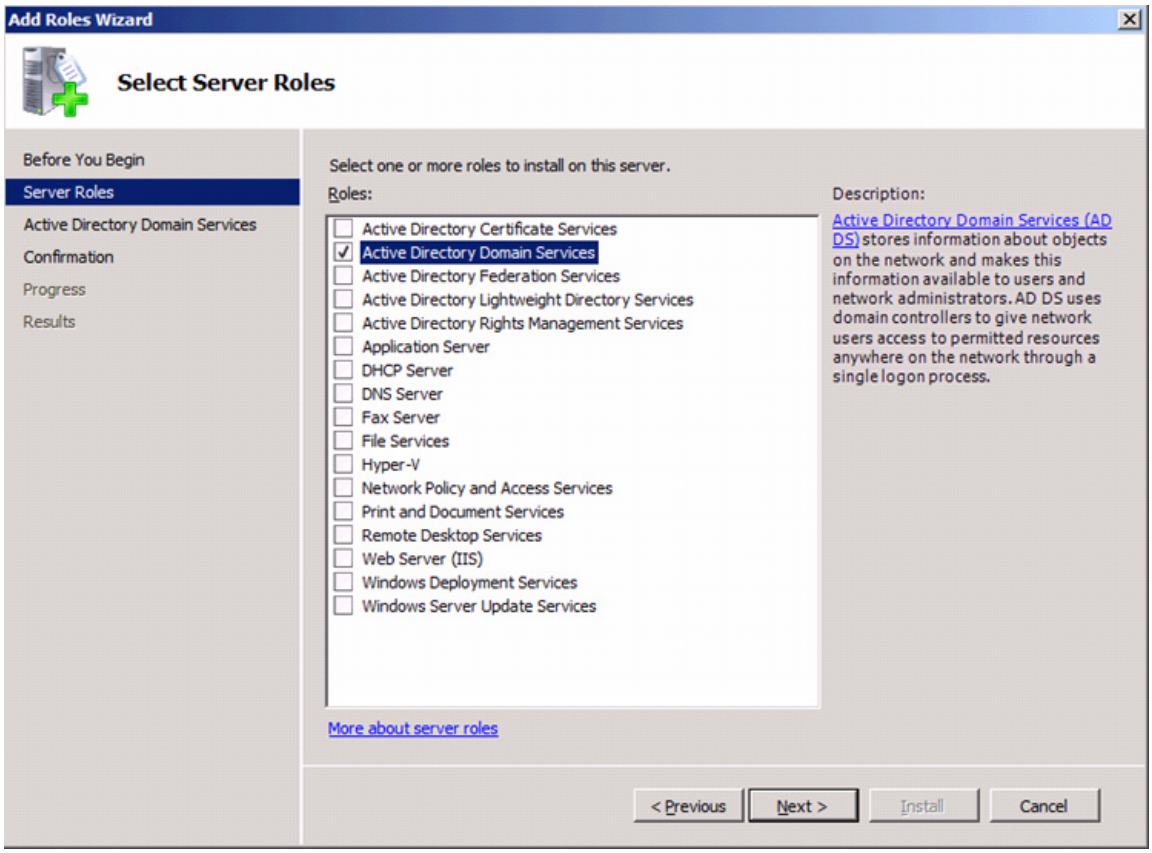
Add Roles Wizard
Select Server Roles
Before You Begin
Server Roles
Active Directory Domain Services
Confirmation
Progress
Results
Select one or more roles to install on this server.
Roles:
Active Directory Certificate Services
Active Directory Domain Services
Active Directory Federation Services
Active Directory Lightweight Directory Services
Active Directory Rights Management Services
Application Server
DHCP Server
DNS Server
Fax Server
File Services
Hyper-V
Network Policy and Access Services
Print and Document Services
Remote Desktop Services
Web Server (IIS)
Windows Deployment Services
Windows Server Update Services
Description:
Active Directory Domain Services (AD DS) stores information about objects on the network and makes this information available to users and network administrators. AD DS uses domain controllers to give network users access to permitted resources anywhere on the network through a single logon process.
More about server roles
< Previous
Next >
Install
Cancel

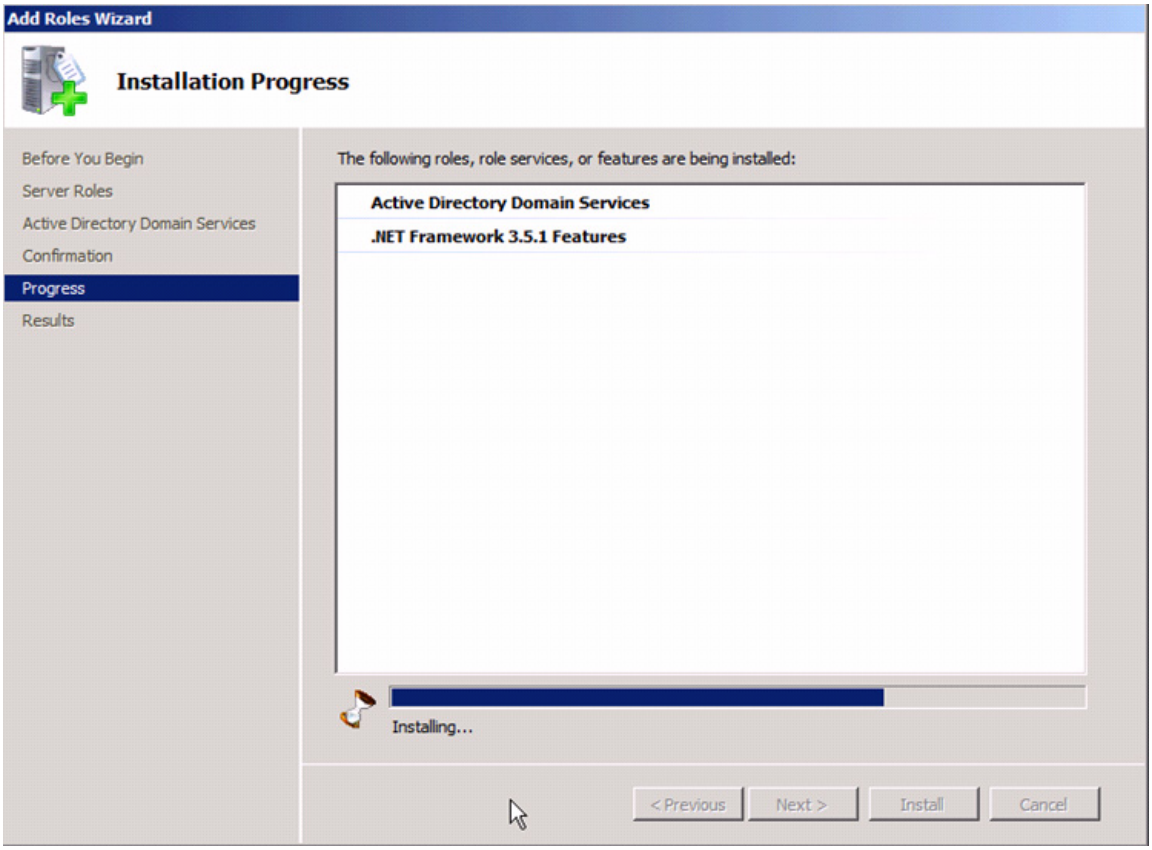

Run dcpromo to configure the Domain Controller.

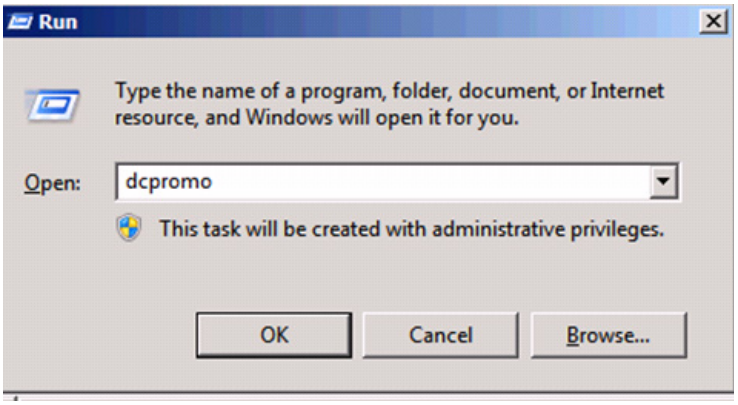

Complete the domain controller installation and repeat the process on VM-Host-Infra-02 to install the redundant domain controller.