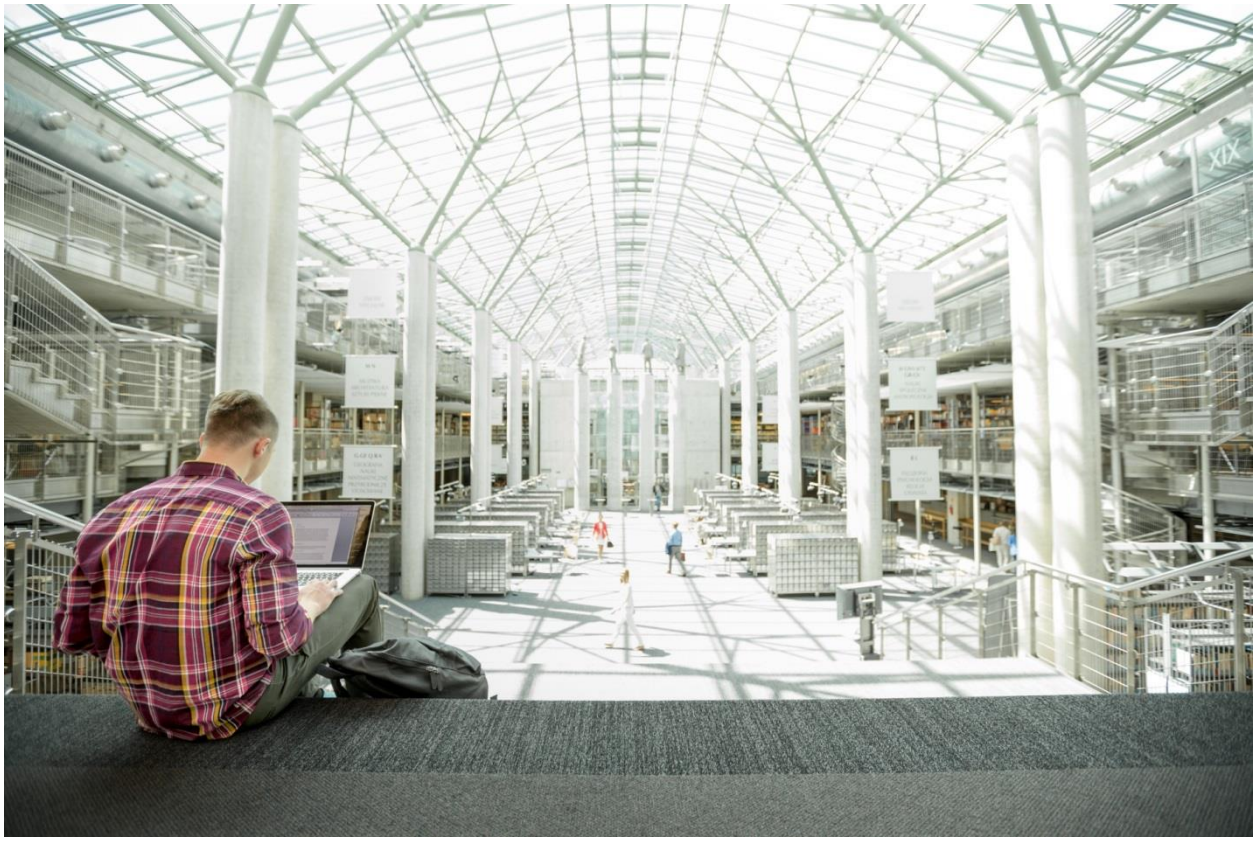

SmartStack with Cisco UCS and Nimble AF7000 All Flash Array

Deploying a SmartStack Integrated Infrastructure Platform based on Cisco UCS, Nimble AF7000 All Flash Array, Cisco MDS Switches for the SAN Fabric and Cisco Nexus Switches for the Datacenter LAN Network

Last Updated: August 22, 2016

About Cisco Validated Designs

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.

Table of Contents

Executive Summary

Cisco Validated Designs (CVD) are systems and solutions that have been designed, tested and documented to facilitate and accelerate customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of a customer. CVDs deliver a validated design, documentation and support information to guide customers from design to deployment.

Cisco, and Nimble Storage have partnered to deliver a series of best of breed SmartStack™ solutions that combine Cisco Unified Computing System servers, Cisco Nexus family of switches, and Nimble Storage arrays for both Enterprise and Cloud datacenters.

SmartStack enables a data center platform with the above characteristics by delivering an integrated architecture that incorporates compute, storage and network design best practices. SmartStack minimizes IT risk by testing the integrated architecture to ensure compatibility between integrated components. SmartStack also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in all stages (planning, design, and implementation) of a deployment.

SmartStack incorporates compute, network and storage best practices to deliver a resilient, scalable and flexible datacenter architecture. The design uses Cisco UCS servers for compute, VMware vSphere 6.0U2 hypervisor, Cisco Nexus 9000 series as the network platform and Cisco MDS 9000 Series switches for the Fibre Channel (FC) network to connect to the Nimble Storage AF7000 all flash arrays.

Documentation for this CVD includes the following documents:

- SmartStack Design Guide
- SmartStack Deployment Guide

This document serves as the SmartStack Deployment Guide. SmartStack Design Guide associated with this deployment guide can be found at:
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/smartstack_AF7000_design.html

Solution Overview

Introduction

This document outlines the deployment procedures for implementing a SmartStack infrastructure solution based on VMware vSphere 6.0, Cisco UCS, Nimble AF7000 and Cisco Nexus and Cisco MDS switches.

Audience

The intended audience for this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to deploy the core SmartStack architecture with Cisco UCS and Nimble Storage all flash arrays.

Solution Design

The SmartStack architecture based on Cisco UCS and Nimble Storage AF7000 all flash array provides a converged infrastructure solution for Enterprise datacenters and cloud deployments. SmartStack integrates Cisco UCS compute, Nimble AF7000 storage, Cisco Nexus 9000 series platform switch and Cisco MDS Fabric switch to deliver a foundational platform capable of supporting multiple services and applications. SmartStack delivers compute, storage, Fibre Channel (FC) based 8G/16G SAN connectivity and 10 Gigabit Ethernet (10GbE) LAN connectivity in a highly resilient, flexible and scalable architecture. The architecture is modular with architectural flexibility to scale up by adding resources or scale out by adding multiple SmartStack units or modules. The SAN fabric from Cisco MDS to Nimble array can scale to 16G Fibre Channel with existing infrastructure. The architecture can also support 16G FC end-to-end by migrating from Cisco 6200 Series Fabric Interconnects in the current design to Cisco UCS 6300 Series Fabric Interconnects. The LAN network architecture using Cisco Nexus 9300 series has the port density for significant scale and provides a migration path to 40GbE LAN connectivity and Cisco ACI with investment protection.

The SmartStack topology provided below illustrates the SmartStack design using the following components:

- Cisco Unified Computing System (UCS) – blade and rack servers
- Cisco UCS 6200 Series Fabric Interconnects (FI) – unified access to storage and LAN networks
- Cisco Nexus 9300 series switches – connectivity to users, other LAN networks and Cisco UCS domains
- Cisco MDS fabric switches – SAN fabric providing Fibre Channel (FC) connectivity to storage
- Nimble AF7000 array – SAN boot of all flash storage array with SSDs
- Cisco Nexus 1000V (optional) – access layer switch for virtual machines
- VMware vSphere 6.0 U2 – Hypervisor

Figure 1 SmartStack Design

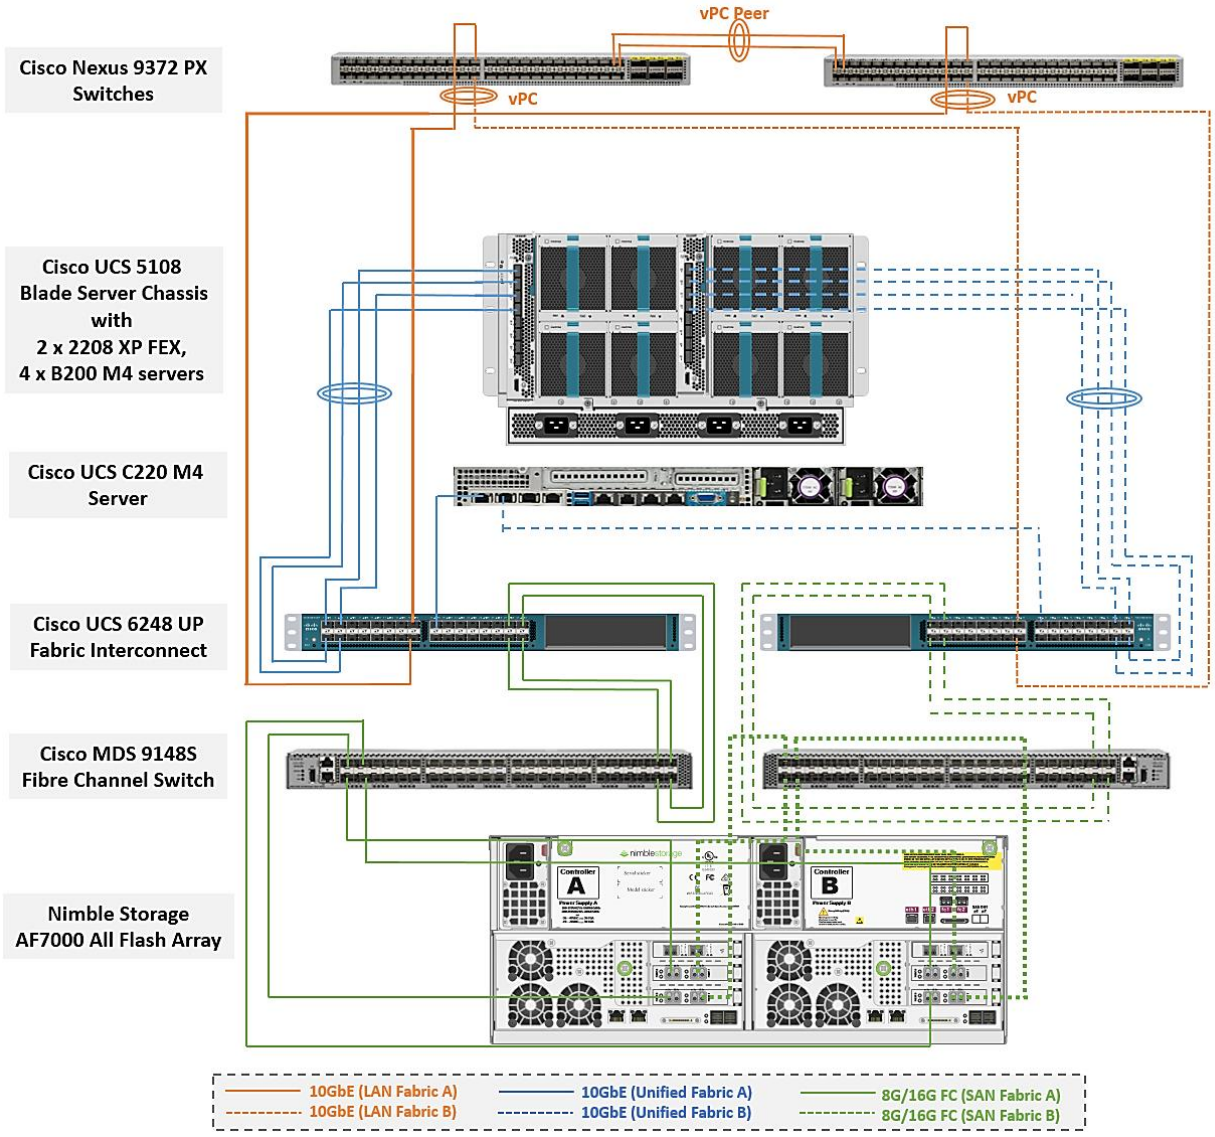

The stateless server provisioning and management capabilities of Cisco UCS Manager combined with Nimble wizard based provisioning and simplified storage management enable quick deployment and provisioning infrastructure resources. Each functional layer (compute, network, and storage) of the architecture is designed to be highly resilient. Two SAN fabrics provide redundant paths for SAN boot and block storage to shared datastores.

Depending on the size and needs of a deployment, the common services can also be deployed with Applications VMs on a SmartStack POD. This document assumes that a Common POD already exists in the deployment and therefore the focus here is on the deployment of the SmartStack POD. Note, that the Management POD can use the same design as SmartStack POD.

For a more details on the SmartStack architecture, please refer to the SmartStack Design Guide: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/smartstack_AF7000_design.html

Compute

Cisco UCS B-Series and C-Series servers provide the compute resources in this SmartStack design. Several models of these servers are supported – Cisco UCS B200M4 and C220 M4 servers were used for validation. The compute design consists of an infrastructure management POD and an application POD, each with its own hardware and running VMware ESXi 6.0 U2 hypervisor.

The infrastructure management layer consists of common infrastructure services that are necessary to deploy, operate and manage the entire deployment. For validation, common components such as Active Directory, DNS, DHCP, vCenter, Cisco Nexus 1000v virtual supervisor module (VSM) and Cisco UCS Performance Manager were deployed in an infrastructure management POD.

The application layer or the SmartStack POD consists of any virtualized application, hosted on SmartStack compute, which the business needs to fulfill its business functions. For validation, IOMeter virtual machines representing application VMs were hosted on the SmartStack POD. Depending on the size and needs of a deployment, the common services can also be deployed with Applications VMs on a SmartStack POD.

Features available at the hypervisor layer (for example, VMware clustering, high availability) are leveraged in both the infrastructure management and application PODs. The blade server chassis and rack mount servers connect into a pair of Cisco UCS 6248 Fabric Interconnects (FI). The blade server chassis connects to the FI through the Cisco FEX module located at the back of the chassis. The rack mount server uses the direct-attached design to connect directly into the FIs and does not use an optional FEX.

The two Fabric Interconnects are deployed in a cluster for redundancy and provide two fabric planes (FI-A, FI-B) that are completely independent of each other from a data plane perspective. In the event of a failure or if only one was deployed, the fabric can be fully operational with one FI.

The FIs provide 10GbE connectivity to the LAN network infrastructure and 8G/16G FC connectivity to the SAN fabric. The FI provides 40Gbps of aggregate bandwidth to the LAN network and 64Gbps to the SAN network. Link aggregation using port channels are used on the unified fabric FEX-to-FI and on the FI-to-Cisco MDS connection. Virtual port channels are used on the FI-to-Cisco Nexus uplinks to the LAN network.

Virtualization

The common infrastructure services and the associated virtual machines (VM) are part of a dedicated infrastructure cluster (`SS-Mgmt`) running VMware High Availability. A separate cluster (`SS-AppVMCluster-CVDFC`) is also enabled for VMware HA and hosts the application VMs. Cisco Nexus 1000V switches is an optional component of the design and was used to provide Layer 2 connectivity between the virtual machines. VMware virtual switches (vSwitch) or VMware distributed virtual switches could also be used.

Cisco UCS B200M4 blade servers with Cisco 1340 VIC and Cisco UCS C220M4 rack servers with Cisco 1227 VIC running vSphere 6.0 U2 are used in this SmartStack design. Cisco UCS servers were assigned to a VMware High Availability (HA) cluster to mitigate against host failures. Two VMware HA clusters were used in validation – one for infrastructure management and services (for example, VMware vCenter) and one for applications that users access. The Cisco VIC on each server presents multiple vPCIe devices to ESXi. vSphere identifies these virtual adapters as vmnics. In this SmartStack design, the following virtual adapters (vNICs) were used with –A connected to unified fabric A and –B to unified fabric B resulting in each ESXi node being dual homed to the external network.

- One vNIC for application VM, in-band management and vMotion traffic through Fabric A (`vNIC-A`)
- One vNIC for application VM, in-band management and vMotion traffic through Fabric B (`vNIC-B`)

The connectivity within each ESXi server and the vNIC presented to ESXi are shown below.

Figure 2 Cisco UCS Server Networking - vSwitch

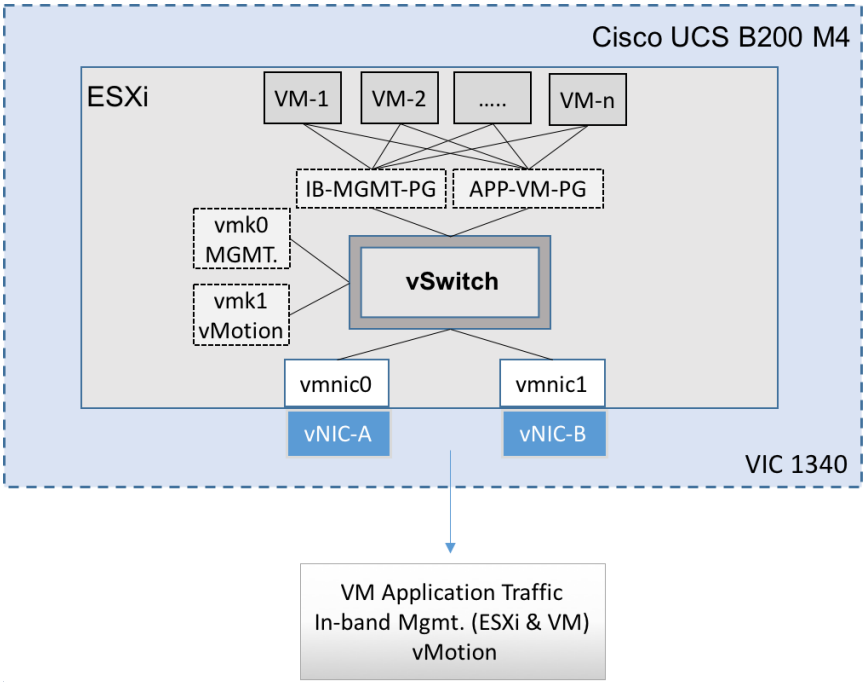

The SmartStack architecture uses two port groups (IB-MGMT-PG) for in-band management of the VMs and APP-VM-PG for application traffic. The design also used two VMkernel NICs (vmk), each with its own port group for host level functionality:

- vmk0 - ESXi management
- vmk1 - vMotion interface

The ESXi management interface is for host to vCenter connectivity, direct access to ESXi shell and VMware cluster communication. The vMotion interfaces are private subnets supporting data access and VM migration across the SmartStack infrastructure.

SmartStack design also supports using Cisco Nexus 1000V virtual switch instead of vSwitch. This is optional and if used, the connectivity within the host is as follows.

Figure 3 Cisco UCS Server Networking – Cisco Nexus 1000V (Optional)

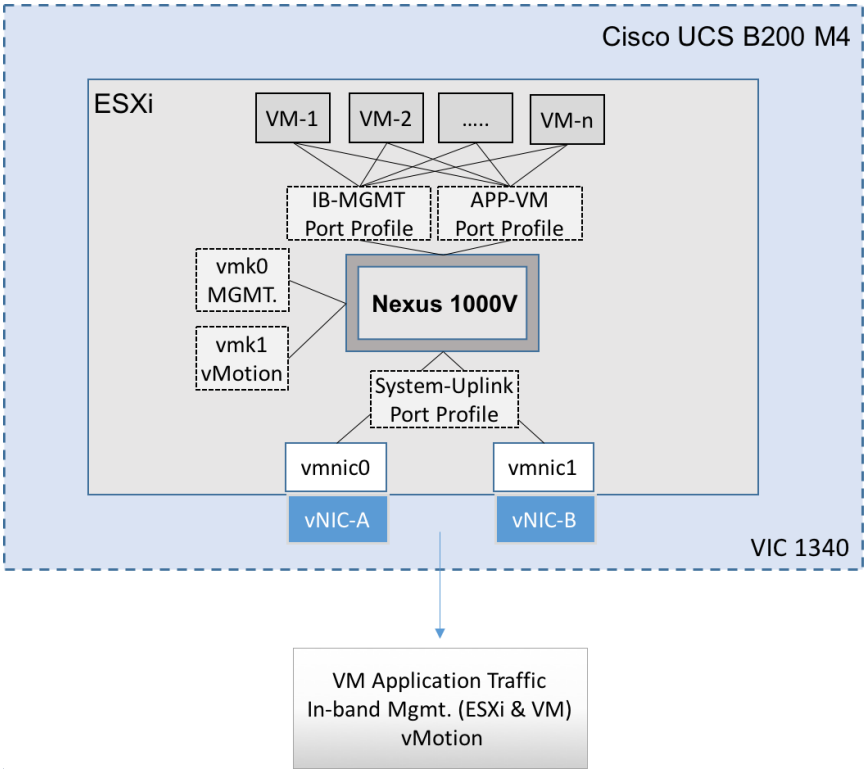

Storage

This SmartStack design uses Nimble AF7000 all flash array to provide block storage. The array used in the validation was deployed with a base configuration of 11.5TB of raw capacity (24 x 480GB SSDs).

To optimize and model the required performance, the storage array and virtual machines will be remotely monitored from the cloud using Nimble InfoSight™. This provides insight into data I/O patterns and capacity usage, and trend analysis for capacity planning and expansion. Also it allows for pro-active ticketing and notification when any issues occur. Providing this kind of deep level analytics into the data patterns and requirements, along with Nimble expandability allows an array to scale performance in exactly the desired area.

Nimble AF7000 supports the addition of up to 2 expansion shelves and up to 4 adaptive flash (CS-Series) or all flash (AF-Series) arrays (in any combination) in a scale-out cluster to increase performance and capacity. Each Nimble Storage controller supports up to 3 FC interface cards, each with dual 16G FC ports. This SmartStack design uses 16G fabric connectivity with two FC interface cards to provide 64G of FC bandwidth per controller. For additional FC bandwidth, a third FC card can be deployed on each controller but this interface is typically used for 10GbE connections to other arrays in a scale-out cluster for data replication traffic. The links between a pair of Cisco MDS and Fabric Interconnect switches are aggregated using 4x8G FC links to deliver 32G of bandwidth across the SAN fabric to each controller. Nimble Storage arrays support non-intrusive upgrades for adding additional capacity\ additional cache, controller upgrades, or adding additional arrays for scale-out.

This SmartStack design uses FC SAN boot for the primary boot device of the Cisco UCS blades. The Service Profile used to configure and deploy Cisco UCS servers is configured to include a boot policy that points to the Nimble Storage array. The boot policy specifies a primary and secondary SAN path to the two

controllers on the array where the boot volumes reside. A second boot policy is also configured but with the primary and secondary paths reversed from that of the first boot profile. The second boot policy is used to load balance SAN boot across different paths when multiple servers are booting. This is an optional aspect of the design that can be helpful in larger deployments for distributing load when multiple servers have to be simultaneously booted. Each server has a dedicated boot volume (40GB) on the Nimble storage array. Nimble Storage arrays provide an Access Control List at the initiator level to only allow connections from the appropriate Cisco UCS blade. During the initial SAN boot, the server attaches to all primary and secondary connections to both active and standby controllers. This provides for normal boot operations even when a controller or primary path is offline. The hosts are configured with the Nimble Connection Manager and Path Selection Policy which optimize MPIO (multi-pathing) settings. This will allow for proper FC path management and failover connectivity with Nimble Storage volumes.

The following section of this document provides more details on the connectivity and high availability aspects of this design.

Figure 4 Nimble Storage Array Configuration

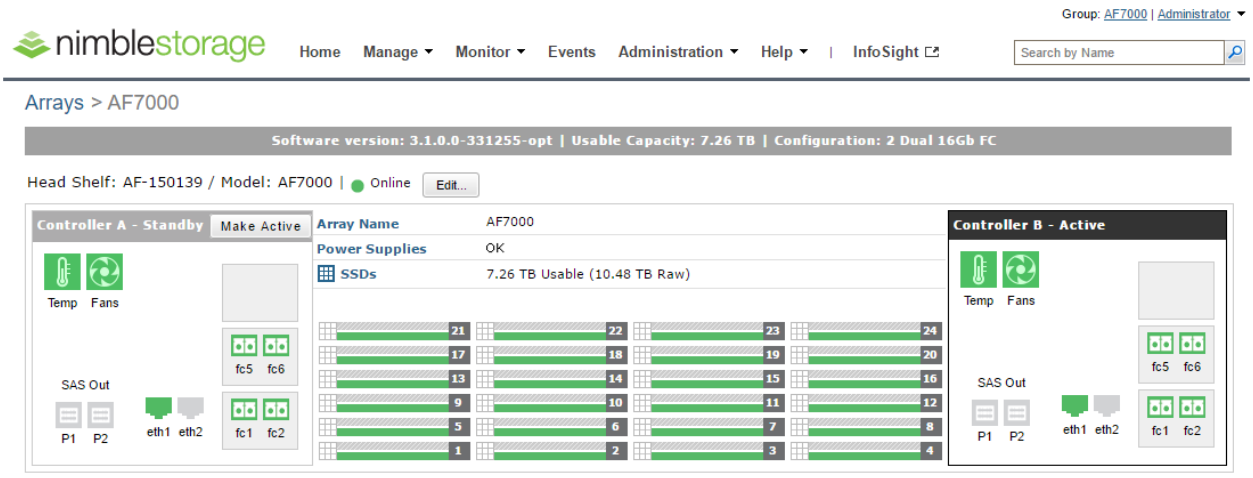

Networking

The LAN network provides network reachability to the applications hosted on Cisco UCS servers in the data center. The infrastructure consists of pair of Cisco Nexus 9372 PX switches deployed in NX-OS standalone mode. Two 10Gbps links from each Cisco Nexus switch are connected to a 10Gbps port on each FI to provide 20Gbps of uplink bandwidth through each Cisco Nexus switch. Virtual Port Channels (vPCs) are configured across these links to provide link and node redundancy while providing higher uplink bandwidth. VLAN trunking is enabled on these links as multiple application data, management and vMotion VLANs needs to traverse these links. A number of design practices are also implemented in this design – see Cisco Nexus 9000 best practices section of the SmartStack Design Guide for details.

The SAN network provides fibre channel connectivity to the Nimble storage array and consists of a pair of Cisco MDS switches. The Cisco MDS switches form completely separate fabrics (SAN fabric A, SAN fabric B) and use a dual vSAN (vSAN-A, vSAN-B) design to provide two redundant and completely diverse paths to the Nimble Storage array.

Link aggregation using port channels are used to aggregate 4 x 8G FC links to provide 32G of FC bandwidth on each SAN Fabric between Cisco FI and Cisco MDS switches. Link aggregation is not used on the links to Nimble array but 2 links from each SAN fabric connects to both controllers to provide 32G of FC bandwidth

to the active controller. Four links from the SAN fabric, 2 from SAN Fabric A and 2 from SAN Fabric B, also connect to the backup controller so that both controllers have 32B FC access to the SAN fabric.

Cisco MDS switches are deployed with N-Port ID Virtualization (NPIV) enabled to support the virtualized environment running on Cisco UCS blade and rack servers. NPIV is necessary to provide isolation in virtualized environments where multiple virtual machines are running on a single server but a LUN needs to be presented to only one VM and not all VMs running on the server. Without NPIV, LUNs would be presented to the host and as a result, all VMs running on that host. To support NPIV on the Cisco UCS servers, the Cisco UCS Fabric Interconnects that connect the servers to the SAN fabric, are enabled for N-Port Virtualization (NPV) mode by configuring to operate in end-host mode (as opposed to FC switching mode). NPV enables Cisco FIs to proxy fabric login, registration and other messages from the servers to the SAN Fabric without being a part of the SAN fabric. This is important for keeping the limited number of Domain IDs that Fabric switches require to a minimum.

SmartStack design also uses jumbo frames with an MTU of 9216 Bytes across the LAN and Unified Fabric links. Jumbo frames increase the throughput between devices by enabling larger sized frames to be sent and received on the wire while reducing the CPU resources necessary to process them. Jumbo frames were enabled during validation on the LAN network links in the Cisco Nexus switching layer and on the Unified Fabric links.

Low-Level Design

Compute

To validate this SmartStack design, a Cisco UCS with 4x Cisco B200M4 half-width blades and a Cisco C220 M4 rack mount server running VMware ESXi 6.0 U2 were deployed in the SmartStack POD to host application VMs. The servers were configured to be part of a cluster with VMware high availability enabled. The blade server chassis is connected to a pair of Cisco UCS 6248 FIs using a pair of Cisco 2208 XP fabric extenders located at the back of the chassis. Eight 10GbE links are used for FEX to FI connectivity, 4 from FEX-A to FI-A and 4 from FEX-B to FI-B to provide an aggregate access bandwidth of 80Gbps to the unified fabric.

The Fabric Interconnects in the SmartStack design are deployed in End-host Ethernet switching mode. Ethernet switching mode determines how the fabric interconnects behave as switching devices between the servers and the network. End-host mode is the default and generally recommended mode of operation. In this mode, the fabric interconnects appear to the upstream LAN devices as end hosts with multiple adapters and do not run Spanning Tree. The Cisco Nexus switch ports that connect to the FI are therefore deployed as spanning tree edge ports.

The ports on the Cisco UCS 6248 FI are unified ports that can support either Ethernet or Fibre Channel traffic based by changing the port mode.

Ethernet ports on the fabric interconnects are not configured by default and must be explicitly configured as a specific type, which determines the port's behavior. The port types used in this design are:

- Uplink ports for connecting to the Cisco Nexus 9300 series switches and external LAN network
- Fiber Channel ports for connecting to the SAN Fabric
- Server ports for connecting to external Cisco UCS C-series rack mount servers

Cisco UCS Manager (Cisco UCSM) is used to provision and manage Cisco UCS and its sub-components (chassis, FI, blade, and rack mount servers). Cisco UCSM runs on the Fabric Interconnects.

A key feature of Cisco UCSM is Service Profile Templates that enable the abstraction of policies, pools, and other aspects of a server configuration and consolidate it in the form of a template. The configuration in a service profile template includes:

- Server Identity (for example, UUID Pool, Mac Address Pool, IP Pools)
- Server Policies (for example, BIOS Policy, Host Firmware Policy, Boot Policy)
- LAN Configuration (for example, VLAN, QoS, Jumbo Frames)
- Storage Configuration (for example, IQN pool)

The template once created, can be used to generate a service profile that configure and deploy individual server or group of servers. A service profile defines the server and its storage and networking characteristics. A service profile template reduces the deployment time, and increases the operational agility and provides general ease of deployment. Service profile templates are used in this SmartStack design to rapidly configure and deploy multiple servers with minimal modification.

LAN Network

The LAN network infrastructure in this SmartStack design consists of a pair of Cisco Nexus 9372 PX switches. Each Cisco UCS FI connects to the switches using 2x10GbE links and the FI Ethernet ports used for this connectivity are configured as Uplink Ports. The uplinks ports are enabled for Link aggregation with FI-A uplinks in port channel 13, and the FI-B uplink ports in port channel 14. The uplinks ports are connected to different Cisco Nexus switches and configured to be part of a virtual PortChannel (vPC) on the Cisco Nexus switches. vPC 13 connects Cisco Nexus A to FI-A and FI-B and vPC 14 connects Cisco Nexus B to FI-A and FI-B. The Cisco Nexus vPC feature allows a device to aggregate links going to different Cisco Nexus switches. As a result, the uplink ports appear as a regular PortChannel to Cisco UCS. Link aggregation, spanning tree and other configuration parameters between Cisco UCS and Cisco Nexus switches follow Cisco recommended best practices – see Design Guide associated with this document for more details.

Multiple VLANs need to traverse the uplink ports and uplink ports are automatically configured as IEEE 802.1Q trunks for all VLANs defined on the fabric interconnect. VLANs for management, vMotion, and applications traffic are defined in the fabric interconnects and enabled on the uplinks. The VLANs used in validating the SmartStack design are summarized in the table below.

Table 1 Uplink VLANs between Cisco UCS and Cisco Nexus

VLAN Type (VLAN Name)	VLAN ID (Used in SmartStack Validation)	Description
Native VLAN (NATIVE-VLAN)	2	Untagged VLAN Traffic are forwarded on this VLAN
In-Band Management	12	VLAN used for in-band management, including ESXi hosts and

(IB-MGMT-VLAN)		Infrastructure VMs
vMotion (vMOTION-VLAN)	3000	VLAN used by ESXi for moving VMs between hosts. vMotion uses management VLAN
VM Application Traffic (APP1-VM, APP2-VM)	950, 951	VLAN used by Application Data Traffic

The detailed deployment procedures for configuring Cisco UCS and Cisco Nexus switches are provided in the SmartStack Deployment section of this document.

SAN Fabric and Storage

FC Cisco MDS Switch to Nimble Storage AF7000 Connectivity

This SmartStack design uses a dual Fabric, each its own VSAN configuration which allows for two diverse paths for FC connectivity. In this example ports FC1/13-16 are the Nimble Storage target ports. The connections on the Nimble Storage must mirror each other (for example, FC1.1 and FC5.1 from each controller are both connected to FI-A. FC2.1 and FC6.1 from each controller are connected to FI-B).

Figure 5 Cisco MDS to Nimble AF7000 Connectivity Diagram

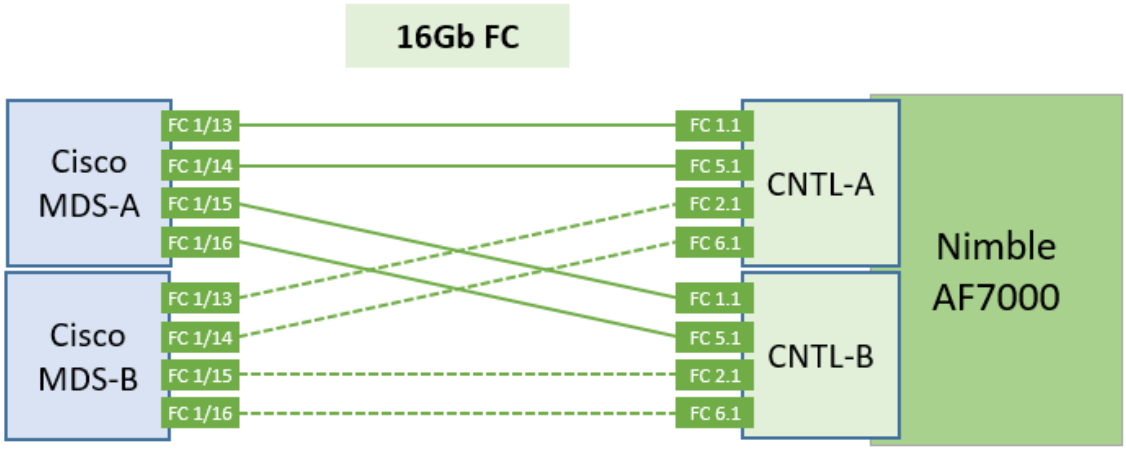

FC Cisco MDS Switch to Cisco UCS Fabric Interconnect Topology

On both Cisco MDS switches, Ports FC1/45-48 is part of a FC port-channel that connects to Fabric Interconnect ports FC1/29-1/32. FC port-channel 41 connects Fabric Interconnect A to MDS-A and FC port-channel 42 connects Fabric Interconnect B to MDS-B. By default, both MDS and FI from Cisco UCSM, VSANs 4091 needs to be created in SAN tab > SAN Cloud > Fabric A. The FC port channel needs to be defined for the ports connected to the Cisco MDS-A (for example, 29 - 32).

Figure 6 Cisco UCS Servers to Cisco UCS Fabric Interconnect Connectivity Diagram

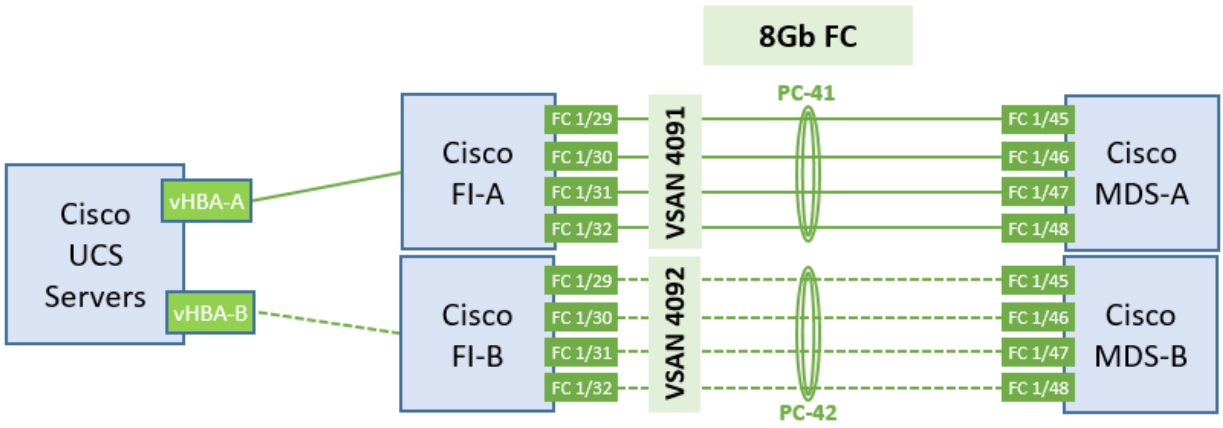

Also, VSANs 4092 needs to be created in SAN tab > SAN Cloud > Fabric B. The FC port channel needs to be defined for the ports connected to the Cisco MDS-A (for example, 29 - 32). When complete you get a similar screen as shown below.

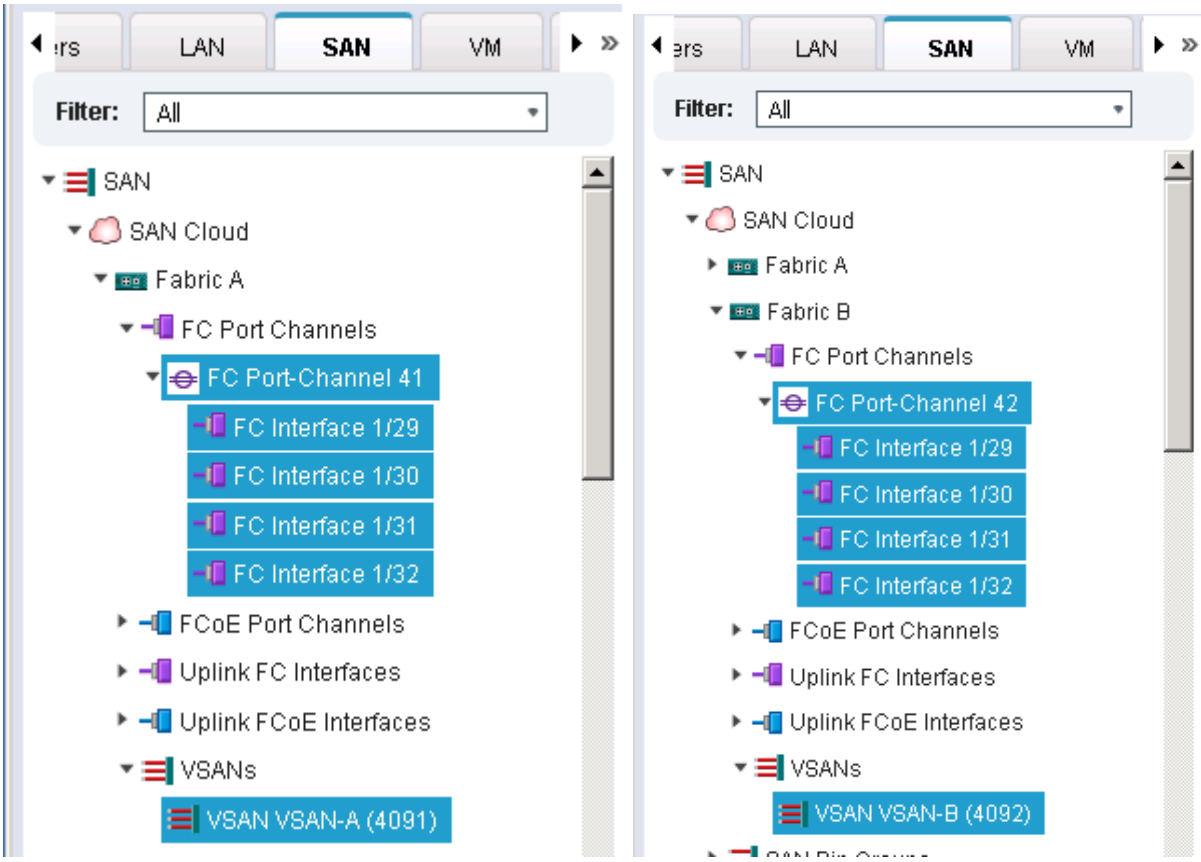

Cisco MDS Switch Fabric Configuration

Cisco MDS-A Switch VSAN requires that a port channel be created to the corresponding ports on Fabric Interconnect A ports (for example, FC1/45-48). The created port channel should match the configuration on the Cisco UCS Fabric Interconnect (for example, 41). Also, note that the VSAN membership is required for

the physical ports as well as the port-channel. In addition, confirm that the VSAN for the Nimble Storage FC ports are also part of the same VSAN (for example, 4091).

```
vsan 4091 interfaces:
    fc1/13          fc1/14          fc1/15          fc1/16
    fc1/45          fc1/46          fc1/47          fc1/48
    port-channel41
```

Cisco MDS-B Switch VSAN requires that a port channel be created to the corresponding ports on Fabric Interconnect B ports (for example, FC1/45-48). The created port channel should match the configuration on the Cisco UCS Fabric Interconnect (for example, 42). Also, note that the VSAN membership is required for the physical ports as well as the port-channel. In addition, confirm that the VSAN for the Nimble Storage FC ports are also part of the same VSAN (for example, 4092).

```
vsan 4092 interfaces:
    fc1/13          fc1/14          fc1/15          fc1/16
    fc1/45          fc1/46          fc1/47          fc1/48
    port-channel42
```

For further security and traffic isolation purposes, the zoning design allows a single initiator to multiple target ports. Note that both the Active and Standby target ports are configured in the same zone.

Figure 7 Nimble Storage AF7000 WWPN Target Ports

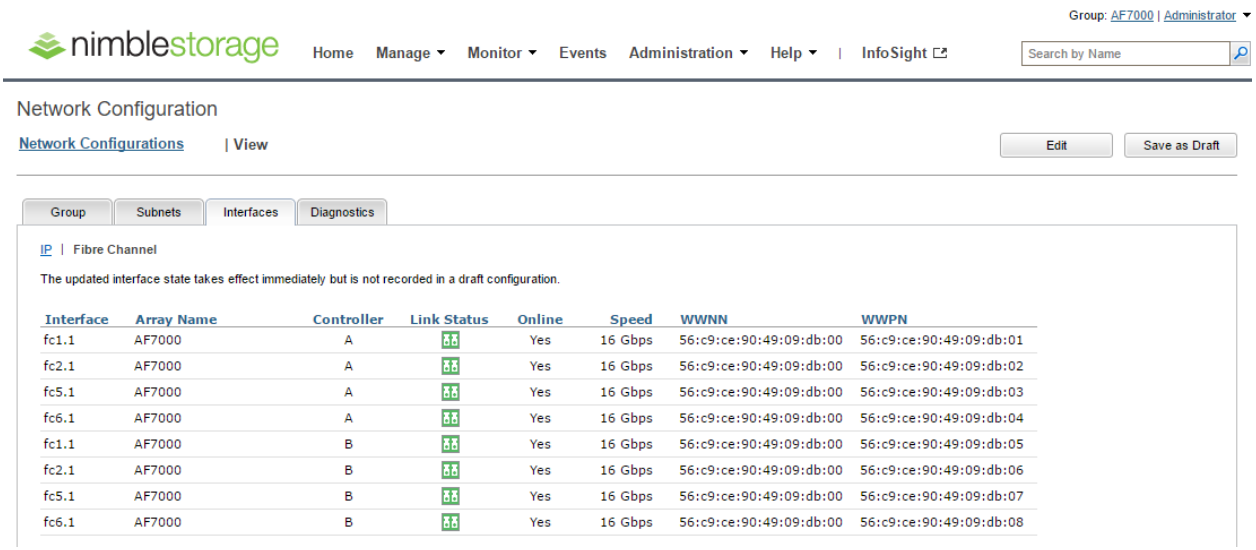

Cisco MDS-A Zoning Configuration

```
zone name AFA-AppVMHost-FIA-3-A vsan 4091
* fcid 0x940407 [pwwn 20:00:00:25:b5:11:aa:06] [AFA-AppVMHost-FIA-3-A]
* fcid 0x940a00 [pwwn 56:c9:ce:90:49:09:db:01] [AFA-CNTLA-p1]
* fcid 0x940b00 [pwwn 56:c9:ce:90:49:09:db:03] [AFA-CNTLA-p2]
* fcid 0x940c00 [pwwn 56:c9:ce:90:49:09:db:05] [AFA-CNTLB-p1]
* fcid 0x940d00 [pwwn 56:c9:ce:90:49:09:db:07] [AFA-CNTLB-p2]

zone name AFA-AppVMHost-FIB-4-A vsan 4091
* fcid 0x940408 [pwwn 20:00:00:25:b5:11:aa:07] [AFA-AppVMHost-FIB-4-A]
* fcid 0x940a00 [pwwn 56:c9:ce:90:49:09:db:01] [AFA-CNTLA-p1]
* fcid 0x940b00 [pwwn 56:c9:ce:90:49:09:db:03] [AFA-CNTLA-p2]
* fcid 0x940c00 [pwwn 56:c9:ce:90:49:09:db:05] [AFA-CNTLB-p1]
* fcid 0x940d00 [pwwn 56:c9:ce:90:49:09:db:07] [AFA-CNTLB-p2]
```

Cisco MDS-B Zoning Configuration

```
zone name AFA-AppVMHost-FIA-3-B vsan 4092
* fcid 0x4b0007 [pwwn 20:00:00:25:b5:11:bb:06] [AFA-AppVMHost-FIA-3-B]
* fcid 0x4b0700 [pwwn 56:c9:ce:90:49:09:db:02] [AFA-CNTLA-p1]
* fcid 0x4b0800 [pwwn 56:c9:ce:90:49:09:db:04] [AFA-CNTLA-p2]
* fcid 0x4b0900 [pwwn 56:c9:ce:90:49:09:db:06] [AFA-CNTLB-p1]
* fcid 0x4b0a00 [pwwn 56:c9:ce:90:49:09:db:08] [AFA-CNTLB-p2]

zone name AFA-AppVMHost-FIB-4-B vsan 4092
* fcid 0x4b0008 [pwwn 20:00:00:25:b5:11:bb:07] [AFA-AppVMHost-FIB-4-B]
* fcid 0x4b0700 [pwwn 56:c9:ce:90:49:09:db:02] [AFA-CNTLA-p1]
* fcid 0x4b0800 [pwwn 56:c9:ce:90:49:09:db:04] [AFA-CNTLA-p2]
* fcid 0x4b0900 [pwwn 56:c9:ce:90:49:09:db:06] [AFA-CNTLB-p1]
* fcid 0x4b0a00 [pwwn 56:c9:ce:90:49:09:db:08] [AFA-CNTLB-p2]
```

Cisco UCS Fabric Interconnect VSAN Configuration

Table 2 Storage VSANs between Cisco UCS and Nimble

VSAN Type (VSAN Name)	VSAN ID (Used in SmartStack Validation)	Description
FC Path A (VSAN-A)	4091	VSAN used for FC traffic on Fabric A. This vsan exists only on Fabric A.
FC Path B (VSAN-B)	4092	VSAN used for FC traffic on Fabric B. This vsan exists only on Fabric B.

For the host side configuration, a single WWNN and dual WWPN configuration was used for the two independent FC paths as identified below.

Table 3 FC Deployment Information

WWNN Pool (Block of 32)	20:00:00:25:B5:11:11:00- 20:00:00:25:B5:11:11:1F
WWPN Pool-A (Block of 128)	20:00:00:25:B5:11:AA:00- 20:00:00:25:B5:11:AA:7F
WWPN Pool-B (Block of 128)	20:00:00:25:B5:11:BB:00- 20:00:00:25:B5:11:BB:7F

Cisco UCS Service Profile Considerations

Each ESXi Service Profile host is configured with two fabric diverse vHBAs to allow connectivity into VSAN-A and VSAN-B.

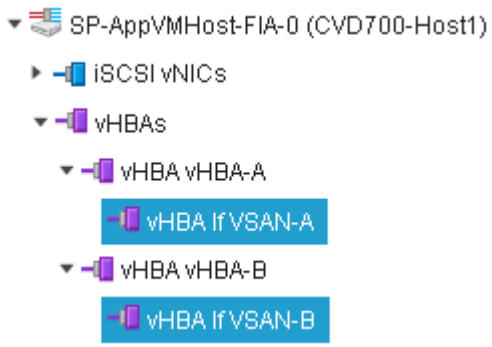

SAN Boot

Each Cisco UCS blade was deployed using FC SAN boot. Using SAN boot affords the advantages of Nimble snapshot, recovery, replication, and cloning mechanisms.

This design has each Cisco UCS blade utilizing two vHBAs that have a presence into diverse fabrics. Each blade had a boot volume created on the Nimble Storage array. The Nimble Storage array provides an initiator group to only honor connections from this single service profile. During FC SAN boot connectivity, the blade connects to both primary and secondary WWPN target for the both active and standby controllers. This provides for normal boot operations regardless of which Nimble Storage Controller is active. The host software utilized MPIO and the Nimble Connection Manager assisted with FC path management. Also the VMware hosts in question were deployed in a cluster to allow for HA failover and to avoid a single point of failure at the hypervisor layer.

Host Storage MPIO Considerations

The VMware ESXi host is configured to use the NIMBLE_PSP_DIRECTED path policy from the Nimble Connection Manager. See “below” section for the detailed procedures on installing and configuring Nimble’s multitpathing policy. For this design, 4 paths will be in Active (I/O) running state and 4 paths in Standby.

Validated Hardware and Software

The table below is a summary of all the components used for validating the SmartStack design.

Table 4 Infrastructure Components and Software Revisions

	Components	Software Version	Comments
Network	Cisco Nexus 9372PX (N9k-9372PX)	6.1(2)I3(5)	Cisco Platform Switch for ToR, MoR, EoR deployments; Provides connectivity to users and other networks and deployed in NX-OS Standalone mode
	Cisco Nexus 1000V	5.2(1)SV3(1.15)	(Optional) Distributed Virtual Switch
	Cisco UCS 6248UP FI	3.1.1g	Fabric Interconnect with embedded management
	Cisco MDS 9148S	6.2(13b)	16G Multilayer Fabric Switch
Compute	Cisco UCS 5108	3.1.1g	Blade Server Chassis
	Cisco UCS B200M4 servers with E5-2600 v4 processors	3.1.1g	Blade Servers
	Cisco UCS C220M4 servers with E5-2600 v4 processors	3.1.1g	Rack mount Servers
	Cisco ENIC Driver	2.3.0.7	Cisco VIC Ethernet driver
	Cisco FNIC Driver	1.6.0.25	Cisco VIC FCoE driver
	Adapter Firmware	4.1(1d)	Cisco VIC Adapter Firmware
Management	Cisco UCS Manager	3.1.1g	Embedded Management
	vCenter plugin for Nimble Storage	Follows Nimble OS	
	Cisco UCS Performance Manager	2.0.0	Performance Monitoring Tool

Storage	Nimble AF7000	NimbleOS 3.1.x	Build: 3.1.0.0-331255-opt
	Nimble NCM for ESXi	3.2.0.0	Build: 3.2.0-600002
	Nimble Storage Windows Toolkit	3.2.0.410	
Virtualization	VMware vSphere	6.0 U2	Custom Cisco ISO but with updated drivers
	VMware vCenter Server	6.0 U1	Appliance
Tools	Workload – IOMeter Tool		
Other	Microsoft Active Directory/DNS		

Solution Deployment – Network Configuration

Cisco Nexus 9372PX Switch

This section provides detailed procedures for deploying and configuring a Cisco Nexus 9000 switch in a SmartStack environment.

Figure 8 Cisco Nexus 9000 Configuration Workflow

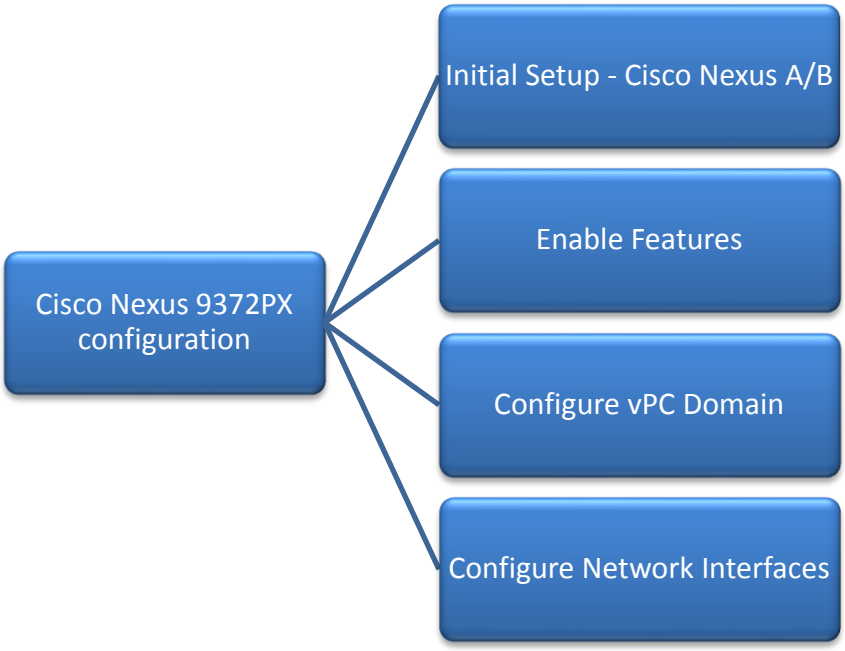

Initial Configuration and Setup

This section outlines the initial configuration necessary for bringing up a new Cisco Nexus 9000.

Cisco Nexus A

To set up the initial configuration for the first Cisco Nexus switch complete the following steps:

1. Connect to the serial or console port of the switch

```
Enter the configuration method: console
Abort Auto Provisioning and continue with normal setup? (yes/no[n]: y

---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no[y] :
Enter the password for "admin":
Confirm the password for "admin":

---- Basic System Configuration Dialog VDC: 1 ----
This setup utility will guide you through the basic configuration of the system. Setup
configures only enough connectivity for management of the system.
Please register Cisco Nexus9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. Nexus9000 devices must be
registered to receive entitled support services.
```

```
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining
dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
Create another login account (yes/no) [n]: n
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name: D01-n9k1
Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
Mgmt0 IPv4 address: 192.168.155.3
Mgmt0 IPv4 netmask: 255.255.255.0
 Configure the default gateway? (yes/no) [y]:
IPv4 address of the default gateway: 192.168.155.1
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]:
Type of ssh key you would like to generate (dsa/rsa) [rsa]:
Number of rsa key bits <1024-2048> [1024]: 2048
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: 192.168.155.254
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/no shut) [no shut]:
Configure CoPP system profile (strict/moderate/lenient/dense/skip) [strict]:
```

2. Review the settings printed to the console. If they are correct, answer yes to apply and save the configuration
3. Wait for the login prompt to make sure that the configuration has been saved prior to proceeding.

Cisco Nexus B

To set up the initial configuration for the second Cisco Nexus switch complete the following steps:

1. Connect to the serial or console port of the switch
2. The Cisco Nexus B switch should present a configuration dialog identical to that of Cisco Nexus A shown above. Provide the configuration parameters specific to Cisco Nexus B for the following configuration variables. All other parameters should be identical to that of Cisco Nexus A.
    - Admin password
    - Nexus B Hostname: D01-n9k2
    - Nexus B mgmt0 IP address: 192.168.155.4
    - Nexus B mgmt0 Netmask: 255.255.255.0
    - Nexus B mgmt0 Default Gateway: 192.168.155.1

In the next section we look at the configuration required on the Cisco Nexus Switches for LAN and management connectivity.

Enabling Global Features and Other Configuration

1. On both Cisco Nexus switches, enable the following features and best practices.

```
feature nxapi
feature udld
feature interface-vlan
feature lacp
feature vpc

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

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

'feature nxapi' is for integration with Cisco UCS Performance Manager

2. Create VLANs for SmartStack system.

```
vlan 12
name IB-MGMT

vlan 2
name Native-VLAN

vlan 950
name APP1-VM

vlan 951
name APP2-VM

vlan 3000
name vMotion
```

Configure vPC Domain

Cisco Nexus A

To configure virtual port channels (vPCs) for switch A, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 155
```

2. Make Cisco Nexus A the primary vPC peer by defining a low priority value:

```
role priority 10
```

3. Use the management interfaces on the supervisors of the Cisco Nexus switches to establish a keepalive link:

```
peer-keepalive destination 192.168.155.4 source 192.168.155.3
```

4. Enable following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
ip arp synchronize
auto-recovery
```

5. Save the configuration:

```
copy run start
```

Cisco Nexus B

To configure vPCs for switch B, complete the following steps:

1. From the global configuration mode, create a new vPC domain:

```
vpc domain 155
```

2. Make Cisco Nexus A the primary vPC peer by defining a higher priority value on this switch:

```
role priority 20
```

3. Use the management interfaces on the supervisors of the Cisco Nexus switches to establish a keepalive link:

```
peer-keepalive destination 192.168.155.3 source 192.168.155.4
```

4. Enable following features for this vPC domain:

```
peer-switch
delay restore 150
peer-gateway
ip arp synchronize
auto-recovery
```

5. Save the configuration:

```
copy run start
```

Configure Network Interfaces for VPC Peer Links

Cisco Nexus A

1. Define a port description for the interfaces connecting to VPC Peer D01-n9k2.

```
interface Eth1/53
description VPC Peer D01-n9k2:e1/53
interface Eth1/54
description VPC Peer D01-n9k2:e1/54
```

2. Apply a port channel to both VPC Peer links and bring up the interfaces.

```
interface Eth1/53,Eth1/54
channel-group 155 mode active
no shutdown
```

3. Enable UDLD on both interfaces to detect unidirectional links.

```
udld enable
```

4. Define a description for the port-channel connecting to D01-n9k2.

```
interface port-channel 155
description vPC peer-link
```

5. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, and the native VLAN.

```
switchport
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950, 951
spanning-tree port type network
```

6. Make this port-channel the VPC peer link and bring it up.

```
vpc peer-link
no shutdown
copy run start
```

Configure Network Interfaces for VPC Peer Links

Cisco Nexus B

1. Define a port description for the interfaces connecting to VPC Peer D01-n9k1.

```
interface Eth1/53
 description VPC Peer D01-n9k1:e1/53
```

```
interface  Eth1/54
description VPC Peer D01-n9k1:e1/54
```

2. Apply a port channel to both VPC Peer links and bring up the interfaces.

```
interface Eth1/53,Eth1/54
channel-group 155 mode active
no shutdown
```

3. Enable UDLD on both interfaces to detect unidirectional links.

```
udld enable
```

4. Define a description for the port-channel connecting to D01-n9k1.

```
interface port-channel 155
description vPC peer-link
```

5. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, and the native VLAN.

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
spanning-tree port type network
```

6. Make this port-channel the VPC peer link and bring it up.

```
vpc peer-link
no shutdown
```

```
copy run start
```

Configure Network Interfaces to Cisco UCS Fabric Interconnects

Cisco Nexus A

1. Define a description for the port-channel connecting to D01-FI-A.

```
interface port-channel 13
description D01-FI-A
```

2. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, and the native VLANs.

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
```

3. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree guard root
no lacp graceful-convergence
mtu 9216
```

4. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

5. Define a port description for the interface connecting to D01-FI-A.

```
interface Eth1/23
description D01-FI-A:p15
```

6. Apply it to a port channel and bring up the interface.

```
channel-group 13 mode active
no shutdown
```

7. Enable UDLD to detect unidirectional links.

```
udld enable
```

8. Define a description for the port-channel connecting to D01-FI-B.

```
interface port-channel 14
description D01-FI-B
```

9. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic VLANs and the native VLAN.

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
```

10. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree guard root
no lacp graceful-convergence
mtu 9216
```

11. Make this a VPC port-channel and bring it up.

```
vpc 14
no shutdown
```

12. Define a port description for the interface connecting to D01-FI-B

```
interface Eth1/24
description D01-FI-B:p15
```

13. Apply it to a port channel and bring up the interface.

```
channel-group 14 mode active
no shutdown
```

14. Enable UDLD to detect unidirectional links.

```
udld enable

copy run start
```

Cisco Nexus B

1. Define a description for the port-channel connecting to D01-FI-A.

```
interface port-channel 13
description D01-FI-B
```

2. Make the port-channel a switchport, and configure a trunk to allow in-band management, VM traffic, and the native VLANs.

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
```

3. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree guard root
no lacp graceful-convergence
mtu 9216
```

4. Make this a VPC port-channel and bring it up.

```
vpc 13
no shutdown
```

5. Define a port description for the interface connecting to D01-FI-A

```
interface Eth1/23
description D01-FI-A:p2
```

6. Apply it to a port channel and bring up the interface.

```
channel-group 13 mode active
no shutdown
```

7. Enable UDLD to detect unidirectional links.

```
udld enable
```

8. Define a description for the port-channel connecting to D01-FI-B

```
interface port-channel 14
description D01-FI-B
```

9. Make the port-channel a switchport, and configure a trunk to allow in-band management, and VM traffic VLANs and the native VLAN.

```
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
```

10. Make the port channel and associated interfaces spanning tree edge ports.

```
spanning-tree port type edge trunk
spanning-tree guard root
no lacp graceful-convergence
mtu 9216
```

11. Make this a VPC port-channel and bring it up.

```
vpc 14
no shutdown
```

12. Define a port description for the interface connecting to D01-FI-B

```
interface Eth1/24
description D01-FI-A:p2
```

13. Apply it to a port channel and bring up the interface.

```
channel-group 14 mode active
no shutdown
```

14. Enable UDLD to detect unidirectional links.

```
udld enable

copy run start
```

Cisco MDS 9148S Switch

Figure 9 Cisco MDS Configuration Workflow

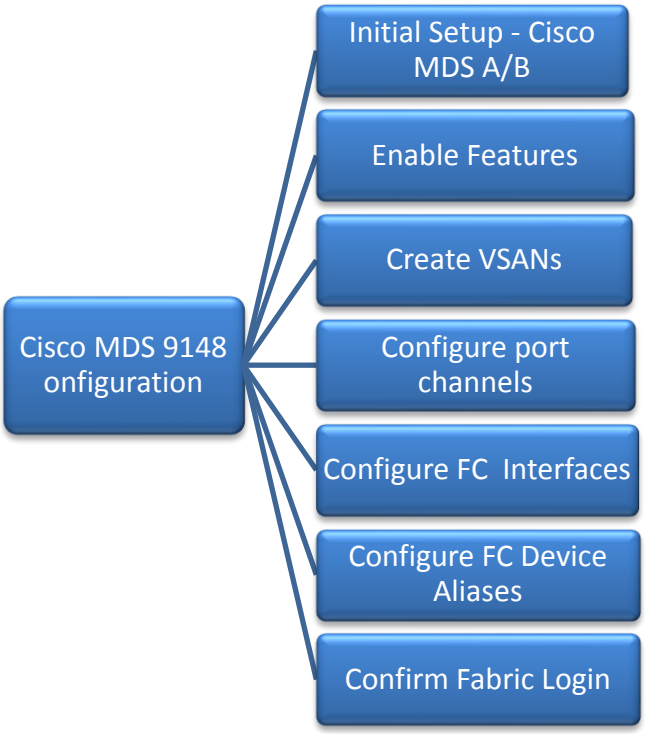

Initial Configuration and Setup

This section provides details on the initial setup of Cisco MDS Fibre Channel Switches. Two switches, Cisco MDS-A and Cisco MDS-B are deployed to provide redundancy in the event of a switch failure.

Cisco MDS A

To set up the initial configuration for the first Cisco MDS switch complete the following steps:

On initial boot, connect to the serial or console port of the switch and the switch should automatically start and attempt to enter Power ON Auto Provisioning.

1. Connect to the serial or console port of the switch

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

         ---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:
      Enter the password for "admin":
  Confirm the password for "admin":

         ---- Basic System Configuration Dialog ----
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management of the system.
Please register Cisco MDS 9000 Family devices promptly with your supplier. Failure to
register may affect response times for initial service calls. MDS devices must be
registered to receive entitled support services.
```

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

Would you like to enter the basic configuration dialog (yes/no): y
 Create another login account (yes/no) [n]:
  Configure read-only SNMP community string (yes/no) [n]:
  Configure read-write SNMP community string (yes/no) [n]:
  Enter the switch name: D01-MDS-A
  Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
    Mgmt0 IPv4 address: 192.168.155.6
    Mgmt0 IPv4 netmask: 255.255.255.0
    Configure the default gateway? (yes/no) [y]:
    IPv4 address of the default gateway: 192.168.155.1
    Configure advanced IP options? (yes/no) [n]:
  Enable the ssh service? (yes/no) [y]:
    Type of ssh key you would like to generate (dsa/rsa) [rsa]:
    Number of rsa key bits <1024-2048> [1024]: 2048
    Enable the telnet service? (yes/no) [n]:
Configure congestion/no credit drop for fc interfaces? (yes/no[y]:
Enter the type of drop to configure congestion/no_credit drop? (con/no) [c]:
Enter milliseconds in multiples of 10 for congestion-drop for port mode F in range (<100-
500>/default), where default is 500.  [d]:
Congestion-drop for port mode E must be greater than or equal to Congestion-drop for port
mode F. Hence, Congestion drop for port mode E will be set as default.
 Enable the http-server? (yes/no) [y]:
 Configure clock? (yes/no) [n]:
 Configure timezone? (yes/no) [n]: y
Enter timezone config [PST/MST/CST/EST]: EST
Enter Hrs offset from UTC [-23:+23]: -5
Enter Minutes offset from UTC [0-59]:
Configure summertime? (yes/no) [n]:
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address: 192.168.155.254
Configure default switchport interface state (shut/noshut) [shut]:
  Configure default switchport trunk mode (on/off/auto) [on]:
  Configure default switchport port mode F (yes/no) [n]:
  Configure default zone policy (permit/deny) [deny]:
  Enable full zoneset distribution? (yes/no) [n]:
  Configure default zone mode (basic/enhanced) [basic]:
```

2. Review the settings printed to the console. If they are correct, answer yes to apply and save the configuration
3. Wait for the login prompt to make sure that the configuration has been saved prior to proceeding.

Cisco MDS B

To set up the initial configuration for the second Cisco MDS switch complete the following steps:

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

1. On initial boot, connect to the serial or console port of the switch.
2. The Cisco Nexus B switch should present a configuration dialog identical to that of Cisco Nexus A shown above. Provide the configuration parameters specific to Cisco Nexus B for the following configuration variables. All other parameters should be identical to that of Cisco Nexus A.
   - Admin password

- MDS B Hostname: **D01-MDS-B**
- MDS B mgmt0 IP address: **192.168.155.7**
- MDS B mgmt0 Netmask: **255.255.255.0**
- MDS B mgmt0 Default Gateway: **192.168.155.1**
- Timezone: **EST**
- Offset from UTC: **-5**
- NTP Server IP: **192.168.155.254**

Cisco MDS Configuration

Enable Global Cisco MDS Features and Settings

The following features should be enabled globally on both Cisco MDS switches from configuration mode.

```
feature npiv
feature fport-channel-trunk
```

Create VSANs

One VSAN per fabric is used in this design – VSAN 4091 on Cisco MDS-A and VSAN 4092 on Cisco MDS-B. Configure VSAN one each switch as follows.

Cisco MDS-A

```
vsan database
vsan 4091
```

Cisco MDS-B

```
vsan database
vsan 4092
```

Configure Port Channels to Cisco UCS Fabric Interconnects

Each FI has a port channel link to one of the Cisco MDS switches for storage traffic. The port channel in this design has 4 links and is configured as follows.

Cisco MDS-A

1. Create the port channel to D01-FI-A.

```
interface port-channel 41
  channel mode active
  switchport rate-mode dedicated
```

2. Assign interfaces to the port channel.

```
interface fc1/45 - 48
  port-license acquire
  channel-group 41 force
  no shutdown
```

3. Add port channel to VSAN in the VSAN database.

```
vsan database
  vsan 4091 interface port-channel41
```

4. Save the configuration.

```
copy run start
```

Cisco MDS-B

1. Create the port channel to D01-FI-B.

```
interface port-channel 42
  channel mode active
  switchport rate-mode dedicated
```

2. Assign interfaces to the port channel.

```
interface fc1/45 – 48
  port-license acquire
  channel-group 42 force
  no shutdown
```

3. Add port channel to VSAN in the VSAN database.

```
vsan database
  vsan 4092 interface port-channel42
```

4. Save the configuration.

```
copy run start
```

Configure FC Interfaces to Nimble Array

Each Cisco MDS has four links to the Nimble AF7000 array, two links to Controller A and two to Controller B. This design uses 16G FC since both Cisco MDS and Nimble array can support up to 16G on these interfaces.

Cisco MDS-A

1. Configure the speed and license for the ports connected to Nimble array.

```
  interface fc1/13 - 16
  port-license acquire
  no shutdown
```

2. Add the interfaces to the VSAN in the VSAN database.

```
vsan database
  vsan 4091 interface fc1/13
  vsan 4091 interface fc1/14
  vsan 4091 interface fc1/15
  vsan 4091 interface fc1/16
```

3. Save the configuration.

```
copy run start
```

Cisco MDS-B

1. Configure the speed and license for the ports connected to Nimble array.

```
interface fc1/13 - 16
  port-license acquire
  no shutdown
```

2. Add the interfaces to the VSAN in the VSAN database.

```
vsan database
  vsan 4092 interface fc1/13
  vsan 4092 interface fc1/14
  vsan 4092 interface fc1/15
  vsan 4092 interface fc1/16
```

3. Save the configuration.

```
copy run start
```

Configure Device Aliases for Nimble Arrays

Using the FLOGI database, configure device-aliases for the WWPN IDs of Nimble Array ports as shown below.

Cisco MDS-A

1. Run 'show flogi database' on Cisco MDS-A to obtain the WWPN info for Nimble array.

```
D01-MDS-A# show flogi database

--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/13           4091  0x940a00  56:c9:ce:90:49:09:db:01 56:c9:ce:90:49:09:db:00
fc1/14           4091  0x940b00  56:c9:ce:90:49:09:db:03 56:c9:ce:90:49:09:db:00
fc1/15           4091  0x940c00  56:c9:ce:90:49:09:db:05 56:c9:ce:90:49:09:db:00
fc1/16           4091  0x940d00  56:c9:ce:90:49:09:db:07 56:c9:ce:90:49:09:db:00
```

2. Configure the device-aliases for the above WWPNs above and commit it as follows.

```
device-alias confirm-commit enable
device-alias database
  device-alias name AFA-CNTLA-p1 pwwn 56:c9:ce:90:49:09:db:01
  device-alias name AFA-CNTLA-p2 pwwn 56:c9:ce:90:49:09:db:03
  device-alias name AFA-CNTLB-p1 pwwn 56:c9:ce:90:49:09:db:05
  device-alias name AFA-CNTLB-p2 pwwn 56:c9:ce:90:49:09:db:07
device-alias commit
```

Cisco MDS-B

1. Run 'show flogi database' on Cisco MDS-B to obtain the WWPN info for Nimble array.

```
D01-MDS-B# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/13           4092  0x4b0700  56:c9:ce:90:49:09:db:02 56:c9:ce:90:49:09:db:00
fc1/14           4092  0x4b0800  56:c9:ce:90:49:09:db:04 56:c9:ce:90:49:09:db:00
fc1/15           4092  0x4b0900  56:c9:ce:90:49:09:db:06 56:c9:ce:90:49:09:db:00
fc1/16           4092  0x4b0a00  56:c9:ce:90:49:09:db:08 56:c9:ce:90:49:09:db:00
```

2. Configure the device-aliases for the above WWPNs above and commit it as follows.

```
device-alias confirm-commit enable
device-alias database
  device-alias name AFA-CNTLA-p1 pwwn 56:c9:ce:90:49:09:db:02
  device-alias name AFA-CNTLA-p2 pwwn 56:c9:ce:90:49:09:db:04
  device-alias name AFA-CNTLB-p1 pwwn 56:c9:ce:90:49:09:db:06
  device-alias name AFA-CNTLB-p2 pwwn 56:c9:ce:90:49:09:db:08
device-alias commit
```

Solution Deployment – Storage Array Configuration

This section provides the procedure for initializing a Nimble Storage array and setting up basic IP connectivity. Note, that the dialog below is specific for an FC array setup.

Figure 10 Nimble Storage Configuration Workflow

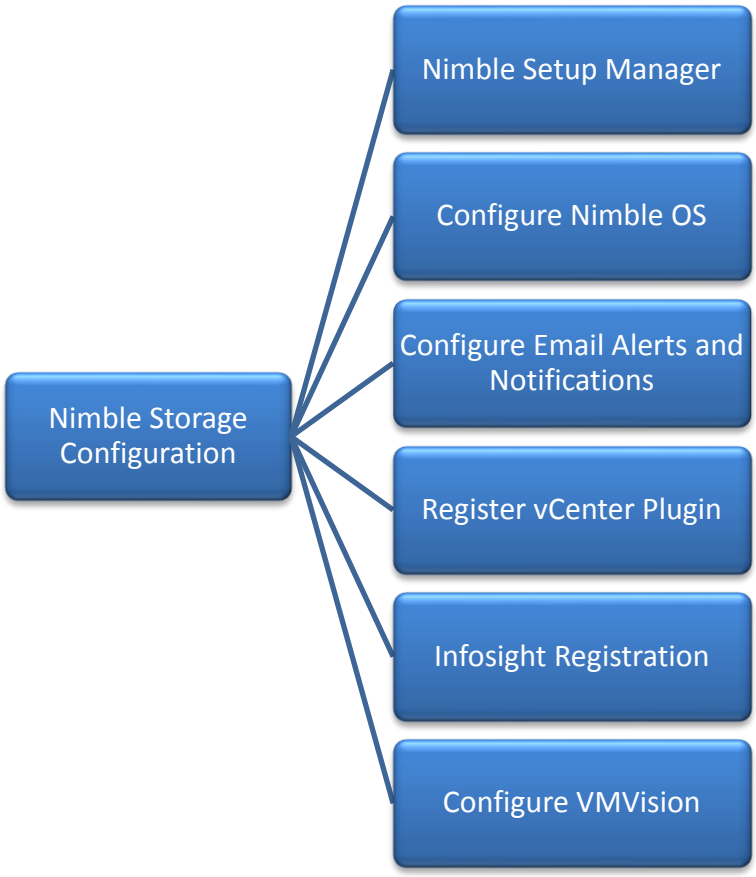

Base Setup of Nimble AF7000 All Flash Array

Nimble Setup Manager

The Nimble Setup manager is part of the Nimble Storage Windows Toolkit. In this section, the Nimble Setup Manager is the only component that needs to be installed. The Nimble Setup manager is used to do the initial setup of the array and can be downloaded from Infosight at this location: https://infosight.nimblestorage.com/InfoSight/media/software/active/1/61/Setup-NimbleNWT-x64.3.2.0.410.zip

The version of the setup manager used must be same or higher than the version of NimbleOS being deployed. Always check InfoSight to see all of the currently available versions of the Windows Toolkit.

Initialize Nimble Storage Array

1. In the Windows Start menu, click Nimble Storage > Nimble Setup Manager.

2. Select one of the uninitialized arrays from the Nimble Setup Manager list and click Next.

If the array is not visible in Nimble Setup Manager, verify that the array's eth1 ports of both controllers are on the same subnet as the Windows host.

Configure Nimble OS using the GUI

1. Choose the appropriate group option and click Next.
    a. Set up the array but do not join a group. Continue to Step 5.
    b. Add the array to an existing group.

If you chose to join an existing group, your browser automatically redirects to the login screen of the group leader array. See Add Array to Group Using the GUI to complete the configuration.

2. Provide or change the following initial management settings and click Finish:
    - Array name
    - Group name
    - Management IP address and subnet mask for the eth1 interface
    - Default gateway IP address
    - Optional. Administrator password
3. You may see a warning similar to "There is a problem with this website's security certificate". It is safe to ignore this warning and click Continue.

If prompted, you can also download and accept the certificate. Alternatively, create your own. See the cert command in the Nimble Command Line Reference Guide. Also, if Internet Explorer v7 displays a blank page, clear the browser's cache. The page should be visible after refreshing the browser.

4. In the login screen, type the password you set and click Log In. From this point forward, you are in the Nimble OS GUI. The first time you access the Nimble OS GUI, the Nimble Storage License Agreement appears.
5. In the Nimble Storage License Agreement, read the agreement, scroll to the bottom, check the acknowledgment box, and then click Proceed.
6. Provide the Subnet Configuration information for the following sections and click Next:
    a. Management IP: IP address, Network and Subnet Mask.

The Management IP is used for the GUI, CLI, and replication. It resides on the management subnet and floats across all "Mgmt only" and "Mgmt + Data" interfaces on that subnet. Note: in this instance you only need to configure the Management network. No IP data network connectivity is required.

    b. Subnet: Subnet label, Network, Netmask, Traffic Type(Data only, Mgmt Only, Mgmt +Data), MTU.
7. Maximum Transmission Unit (MTU) – Standard (1500) Provide Interface Assignment information for the following sections and click Next:

a. Interface Assignment: For each IP interface, assign it a subnet and a Data IP address within the specified network. For inactive interface, assign the "None" subnet.

b. Diagnostics:

   i. Controller A diagnostics IP address will be on the same subnet as the management IP address.
   ii. Controller B diagnostics IP address will be on the same subnet as the management IP address.

8. Provide the following Domain information and click Next:

   a. Domain Name

   b. DNS Servers: Type the hostname or IP address of your DNS server. You can list up to five servers.

9. Provide the following Time information and click Next:

   a. Time Zone: Choose the time zone the array is located in.

   b. Time (NTP) Server: Type the hostname or IP address of your NTP server.

10. Provide Support information for the following sections and click Finish.

11. Email Alerts:

   a. From Address: This is the email address used by the array when sending email alerts. It does not need to be a real email address. Include the array name for easy identification.

   b. Send to Address: Nimble recommends that you check the Send event data to Nimble Storage Support check box.

   c. SMTP server hostname or IP address

   d. AutoSupport:

      i. Checking the Send AutoSupport data to Nimble Storage check box enables Nimble Storage Support to monitor your array, notify you of problems, and provide solutions.
      ii. HTTP Proxy: AutoSupport and software updates require an HTTPS connection to the Internet, either directly or through a proxy server. If a proxy server is required, check the Use HTTP Proxy check box and provide the following information to configure proxy server details:
      iii. HTTP proxy server hostname or IP address
      iv. HTTP proxy server port
      v. Proxy server user name
      vi. Proxy server password

The system does not test the validity of the SMTP server connection or the email addresses that you provided.

12. Click Finish. The Setup Complete screen appears. Click Continue.

13. The Nimble OS home screen appears. Nimble Storage array setup is complete.

Configure Array to Send Email Notifications for Alerts (Optional)

To setup email notification of events from the Nimble Storage array, complete the following steps. This is an optional setup but highly recommended.

1. On the Wellness page, click Daily Summary Emails.
2. Check Subscribe to daily summary email.
3. Enter an email address for delivery of the email alerts.
4. (Optional) You can click Test to send a test email to the email address that you indicated.
5. Click Submit to conclude the email alerts setup.

Setup Nimble Management Tools

vCenter Plugin

The vCenter plugin from Nimble Storage allows for single pane of glass administration directly from vCenter as well as integration with Nimble Infosight analytics. Nimble Storage has integration to vCenter through plugin registration. This allows for datastore creation and management using vCenter. The vCenter plugin is supported on ESX 5.5 update 1 and later.

The plugin is not supported for:
- Multiple datastores located on one LUN
- One datastore spanning multiple LUNs
- LUNs located on a storage device not made by Nimble

For additional info, refer Nimble Storage VMware integration guide:
https://infosight.nimblestorage.com/InfoSight/media/cms/active/pubs_vmware_integration_guide_3.pdf

The procedure for registering the plugin are covered in the later section titled “Setup – Post ESXi and vCenter Install”.

InfoSight

Register and Log into InfoSight

To register and login to InfoSight, complete the following steps.

It can take up to 24 hours for the array to appear in InfoSight after the first data set is sent. Data sets are configured to be sent around midnight array local time. Changes made right after the data set is sent at midnight might not be reflected in InfoSight for up to 48 hours.

1. Log in to the InfoSight portal at https://infosight.nimblestorage.com.
2. Click Enroll now to activate your account. If your email is not already registered, contact your InfoSight Administrator. If there is no existing, InfoSight Administrator (Super User) registered against your account or you are not sure, contact Nimble Storage Support for assistance.
3. Select the appropriate InfoSight role and enter the array serial number for your customer account. If this is the first account being created for your organization, you should select the Super User role. The number of super users is limited to the total number of arrays that are associated with an account.

4. Click Submit.
5. A temporary password is sent to the email address that you specified. You must change your password the first time you log in.

Configure Array to Send Data to InfoSight

To take full advantage of the InfoSight monitoring and analysis capabilities, configure your Nimble arrays to send data sets, statistics, and events to Nimble Storage Support. InfoSight recommendations and automatic fault detection are based on InfoSight processing the data from your arrays. If you do not configure the arrays to send this data to Nimble Storage Support during the initial setup, you can change the configuration at any time from the Administration menu in the GUI.

To configure the array to send data to InfoSight, complete the following steps.

1. From the Administration menu in the array GUI, select Alerts and Monitoring > AutoSupport / HTTP Proxy.
2. On the AutoSupport page, select Send AutoSupport data to Nimble Storage Support.
3. Click Test AutoSupport Settings to confirm that AutoSupport is set up correctly.
4. Click Save.

Configure Arrays to Monitor VMware Environment using VMVision

VMVision is part of InfoSight and provides visibility into the entire virtualization stack. It provides agentless per-VM monitoring and statistics. VMVision provides visibility into VMs with the most I/O churn and resource constraints. For additional info on VMVision, refer: http://uploads.nimblestorage.com/wp-content/uploads/2015/07/12132211/nimblestorage-vmvision.pdf

In order to monitor your VMware environment using VMVision, the following steps must be completed.

- Register the vCenter plugin with each vCenter from each Nimble Array groups
- Enable streaming of performance data to InfoSight.

To verify the vCenter plugin registration completed in earlier step and enabled streaming of performance data, complete the following steps.

1. Log in to https://infosight.nimblestorage.com
2. Go to Administration > VMVision.
3. In the VMVision list, find the array group for which you want to monitor the virtual environment.
4. Verify that your software version is up to date and vCenter plugin is registered.

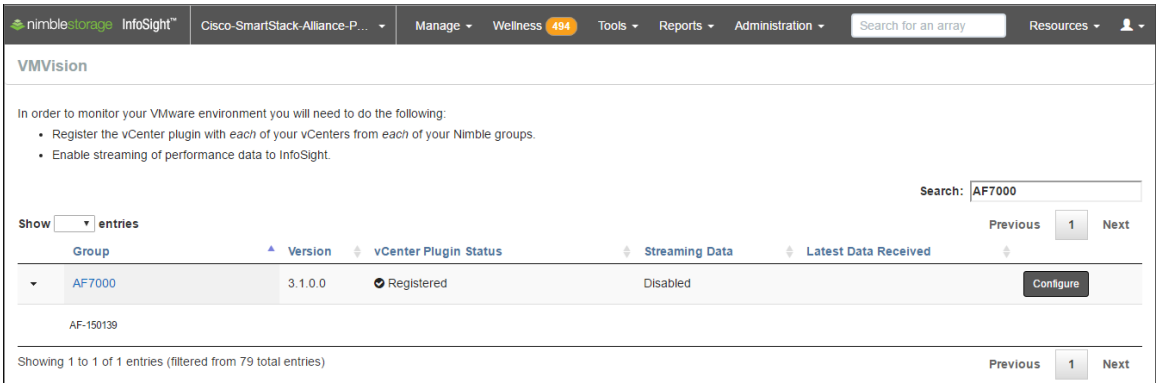

5. Click Configure. In the Configure Group dialog box that opens, select Enabled in the VM Streaming Data list and Click Update.

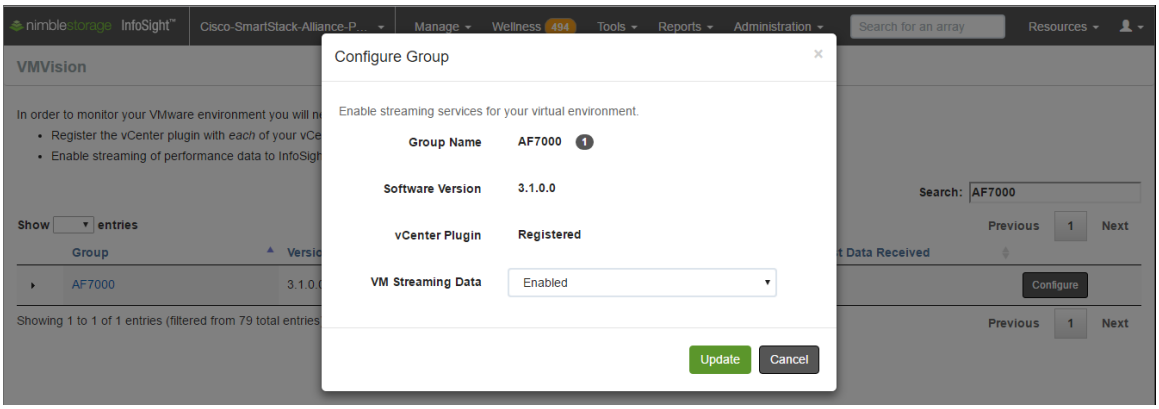

Solution Deployment – Cisco UCS Configuration

This section provides detailed procedures for deploying a Cisco Unified Computing System (Cisco UCS) in a SmartStack environment.

Cisco UCS Configuration Workflow

The figure below shows a high level workflow for deploying and configuring Cisco UCS servers using Cisco UCS Manager. Service Profile Templates enable the rapid deployment and configuration of Cisco UCS Servers in a SmartStack design by consolidating the configuration policies and parameters in a template format so that it can be used to deploy new servers without having to configure each server from scratch.

Figure 11 High Level Workflow for deploying and configuring Cisco UC

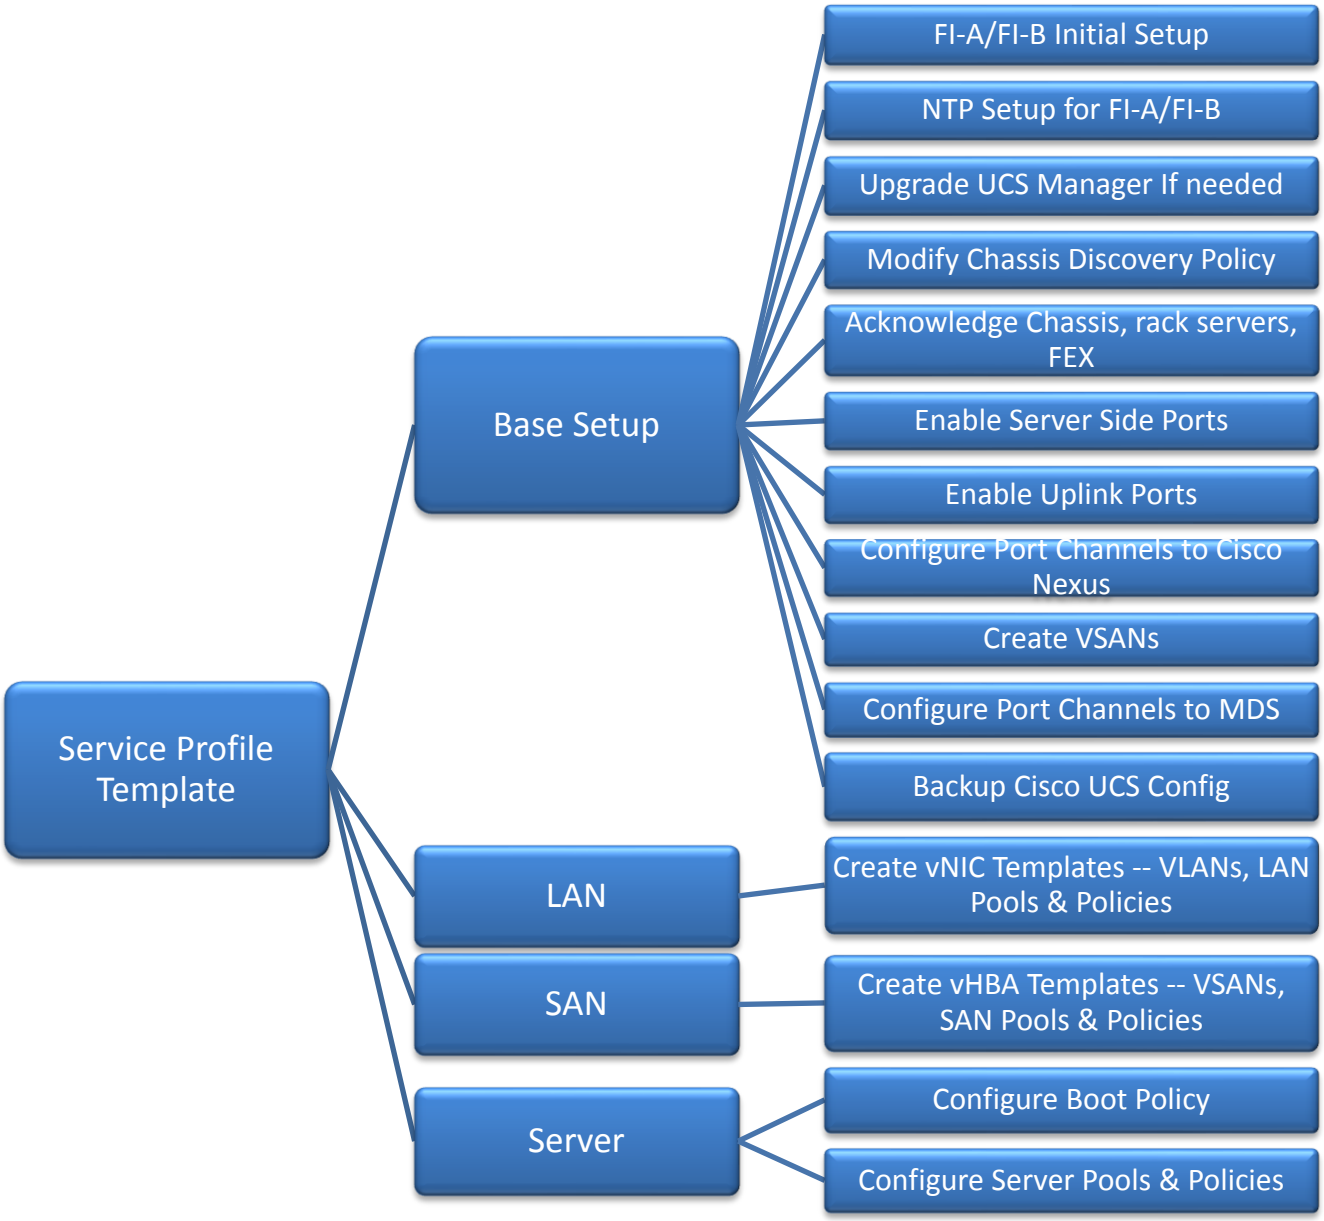

Cisco UCS Configuration – Base Setup

This section outlines the initial setup necessary to deploy a new Cisco UCS domain in a SmartStack environment using a pair of Cisco UCS Fabric interconnects (FI) with embedded Cisco UCS Manager for management.

Initial Setup of Cisco Fabric Interconnects

A pair of Cisco UCS 6248UP Fabric Interconnects is used in this SmartStack design. Minimum configuration required to bring up the FIs and embedded Cisco UCS Manager (UCSM) are outlined below. All configurations after this will be done using Cisco UCS Manager.

Cisco UCS 6248UP FI – Primary (FI-A)

1. Connect to the console port of the primary Cisco UCS FI.

```
Enter the configuration method: console
Enter the setup mode; setup newly or restore from backup.(setup/restore)? Setup You have
chosen to setup a new fabric interconnect? Continue? (y/n): y
Enforce strong passwords? (y/n) [y]: y
Enter the password for "admin": <Enter Password>
Enter the same password for "admin": <Enter Password>
Is this fabric interconnect part of a cluster (select 'no' for standalone)? (yes/no)[n]: y
Which switch fabric (A|B): A
Enter the system name: D01-FI-A
Physical switch Mgmt0 IPv4 address: 192.168.155.8
Physical switch Mgmt0 IPv4 netmask: 255.255.255.0
IPv4 address of the default gateway: 192.168.155.1
Cluster IPv4 address: 192.168.155.89
Configure DNS Server IPv4 address? (yes/no) [no]: y
DNS IPv4 address: 192.168.155.15
Configure the default domain name? y
Default domain name: smartstack.local
Join centralized management environment (UCS Central)? (yes/no) [n]: <Enter>
```

2. Review the settings printed to the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt to make sure that the configuration has been saved prior to proceeding.

Cisco UCS 6248UP FI – Secondary (FI-B)

1. Connect to the console port on the second FI on Cisco UCS 6248UP FI.

```
Enter the configuration method: console

Installer has detected the presence of a peer Fabric interconnect. This Fabric intercon-
nect will be added to the cluster.  Do you want to continue {y|n}?   y

Enter the admin password for the peer fabric interconnect: <Enter Password>

Physical switch Mgmt0 IPv4 address: 192.168.155.9

Apply and save the configuration (select ‘no’ if you want to re-enter)?(yes/no: y
```

2. Verify the above configuration by using Secure Shell (SSH) to login to each FI and verify the cluster status. Status should be as follows if the cluster is up and running properly.

```
D01-FI-A# show cluster state
   Cluster Id: 0x8cb67462eb0c11e2-0x9fff002a6a419c64
```

```
A: UP, PRIMARY

B: UP, SUBORDINATE

HA READY
D01-FI-A#
```

3. Now you're ready to log into Cisco UCS Manager using either the individual or cluster IPs of the Cisco UCS Fabric Interconnects.

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS 6248 Fabric Interconnect cluster IP address configured in earlier step.
2. Click Launch Cisco UCS Manager link to download the Cisco UCS Manager software.
3. If prompted, accept security certificates as necessary.
4. When prompted, enter admin as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.
6. Select Yes or No to authorize Anonymous Reporting if desired and click OK.

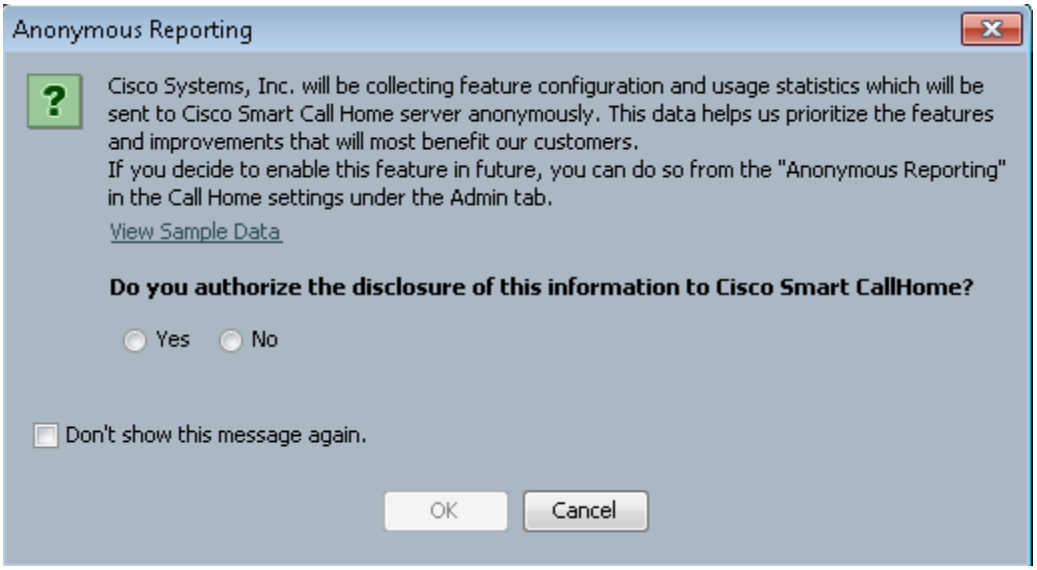

Cisco UCS Manager – Configure NTP Server

To synchronize the Cisco UCS environment to the NTP server, complete the following steps:

1. From Cisco UCS Manager, click Admin tab in the navigation pane.
2. Select All > Timezone Management > Timezone.
3. Right-click and select Add NTP Server.
4. Specify NTP Server IP (for example, `192.168.155.254`) and click OK twice to save edits. The Time Zone can also be specified in the Properties section of the Time Zone window.

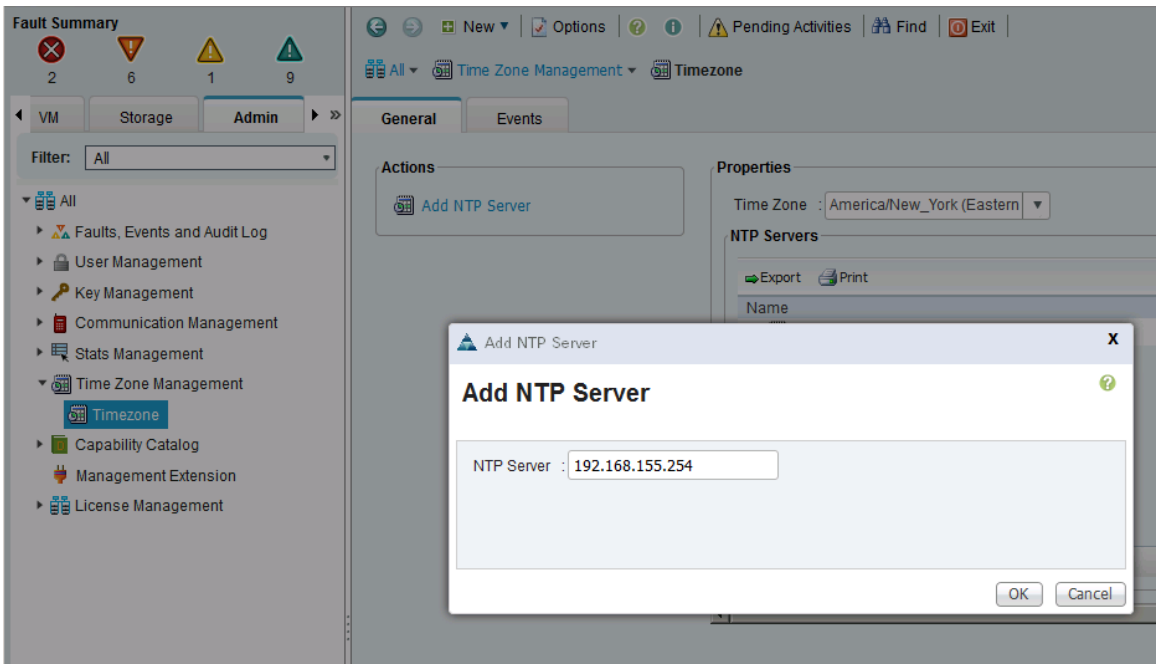

Upgrade Cisco UCS Manager

This document assumes that the Cisco UCS Manager is running the version outlined in the Software Matrix. If an upgrade is required, follow the procedures outlined in the Cisco UCS Install and Upgrade Guides.

Assign Block of IP addresses for KVM Access

To create a block of IP addresses for in-band access to servers in the Cisco UCS environment, complete the following steps. The addresses are used for Keyboard, Video, and Mouse (KVM) access to individual servers managed by Cisoc UCS Manager.

This block of IP addresses should be in the same subnet as the management IP addresses for the Cisco UCS Manager.

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > Pools > root > IP Pools.
3. Right-click and select Create IP Pool.
4. Specify a Name (for example, `ext-mgmt`) for the pool. Click Next.
5. Click [+] Add to add a new IP Block. Click Next.
6. Enter the starting IP address (From), the number of IP addresses in the block (Size), the Subnet Mask, Default Gateway and DNS information. Click OK.
7. Click Finish to create the IP block.

Edit Chassis Discovery Policy

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

1. From Cisco UCS Manager, click Equipment tab in the navigation pane and select Equipment in the list on the left.
2. In the right pane, click Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel.
5. Click Save Changes and then OK to complete.

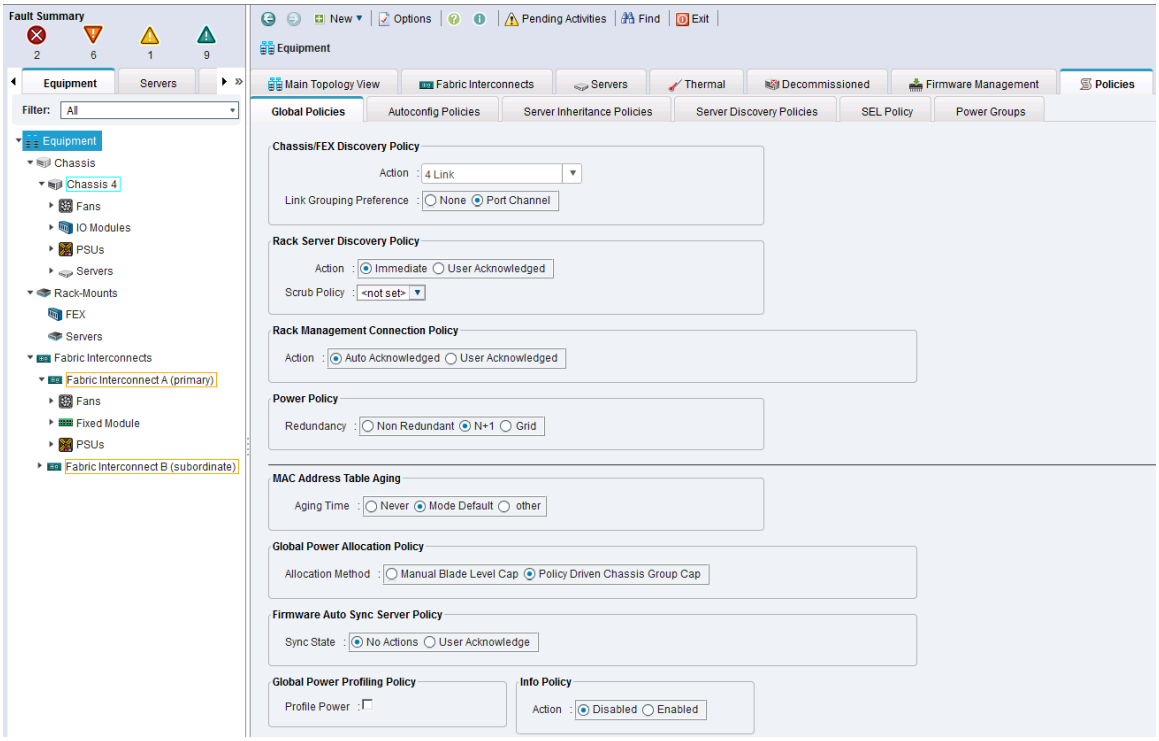

Acknowledge Cisco UCS Chassis, Cisco UCS C-series and FEX

To acknowledge all Cisco UCS chassis and C-Series Servers, complete the following steps:

1. From Cisco UCS Manager, click Equipment tab in the navigation pane.
2. Expand Chassis and for each chassis in the deployment, right-click and select Acknowledge Chassis.
3. In the Acknowledge Chassis pop-up, click Yes and then click OK.
4. If C-Series servers are part of the deployment, expand Rack Mounts.
5. Go to Rack-Mounts in the left window pane and right-click each Server listed and select Acknowledge Chassis. Using FEX for rack mount servers is a design option in the SmartStack design. If FEX is used, acknowledge each FEX.

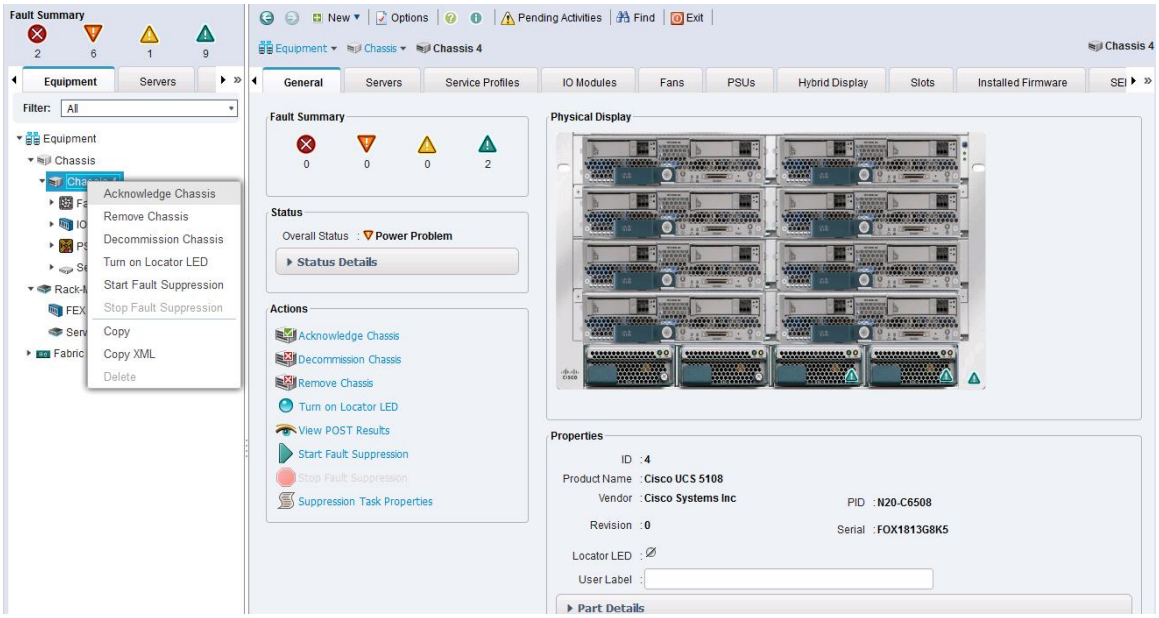

Enable Server Ports

To configure ports connected to Cisco UCS servers as Server ports, complete the following steps:

1. From Cisco UCS Manager, click Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the ports that are connected to Cisco UCS Blade server chassis, Cisco UCS C-series servers or any FEX used for C-series connectivity. Right-click and select Configure as Server Port.
5. Click Yes and then OK to confirm the changes.
6. Repeat above steps for Fabric Interconnect B (secondary) ports that connect to servers.
7. Verify that the ports connected to the servers are now configured as server ports. The view below is filtered to only show Server ports.

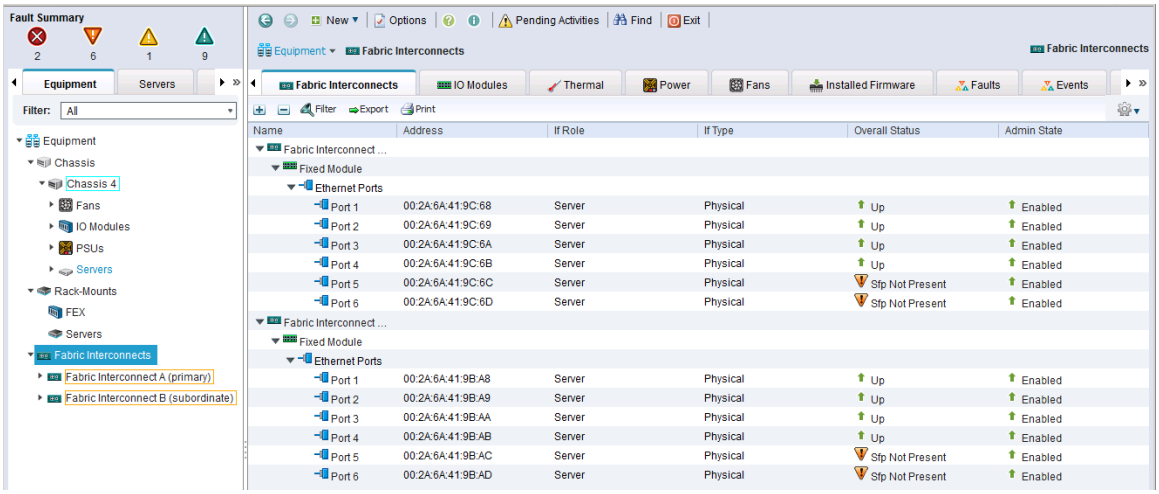

Enable Uplink Ports to Cisco Nexus 9000 Series Switches

To configure ports connected to Cisco Nexus switches as Network ports, complete the following steps:

1. From Cisco UCS Manager, click Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand Ethernet Ports.
4. Select the first port (for example, Port 15) that connects to Cisco Nexus A switch, right-click and select Configure as Uplink Port, click Yes to confirm the uplink ports and click OK. Repeat for second port (for example, Port 16) that connects to Cisco Nexus B switch.
5. Repeat above steps for Fabric Interconnect B (secondary) uplink ports that connect to Cisco Nexus A and B switches.
6. Verify that the ports connected to the servers are now configured as server ports. The view below is filtered to only show Network ports.

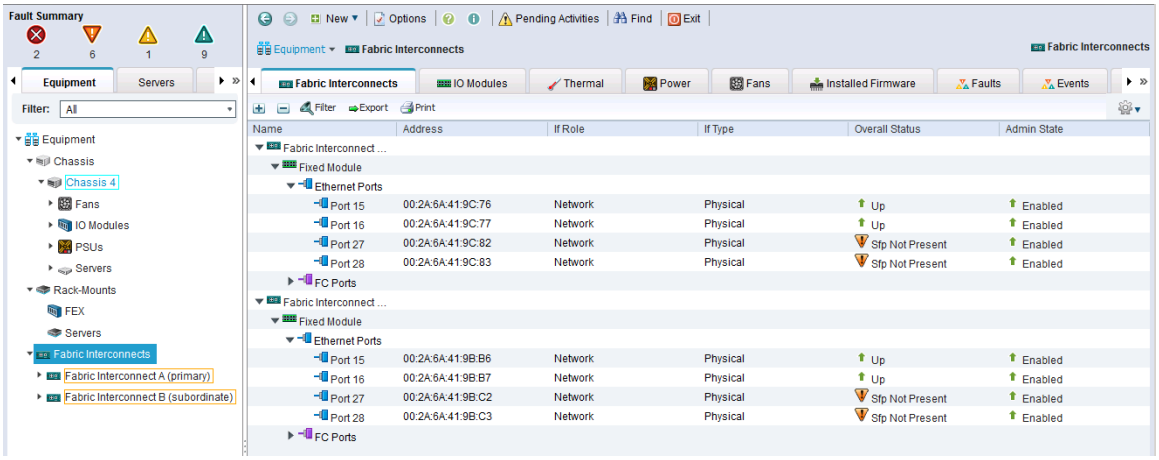

Configure Port Channels on Uplink Ports to Cisco Nexus Switches

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

To configure port channels on Uplink/Network ports connected to Cisco Nexus switches, complete the following steps:

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > Fabric A > Port Channels.
3. Right-click and select Create Port Channel.
4. In the Create Port Channel window, specify a Name and unique ID.

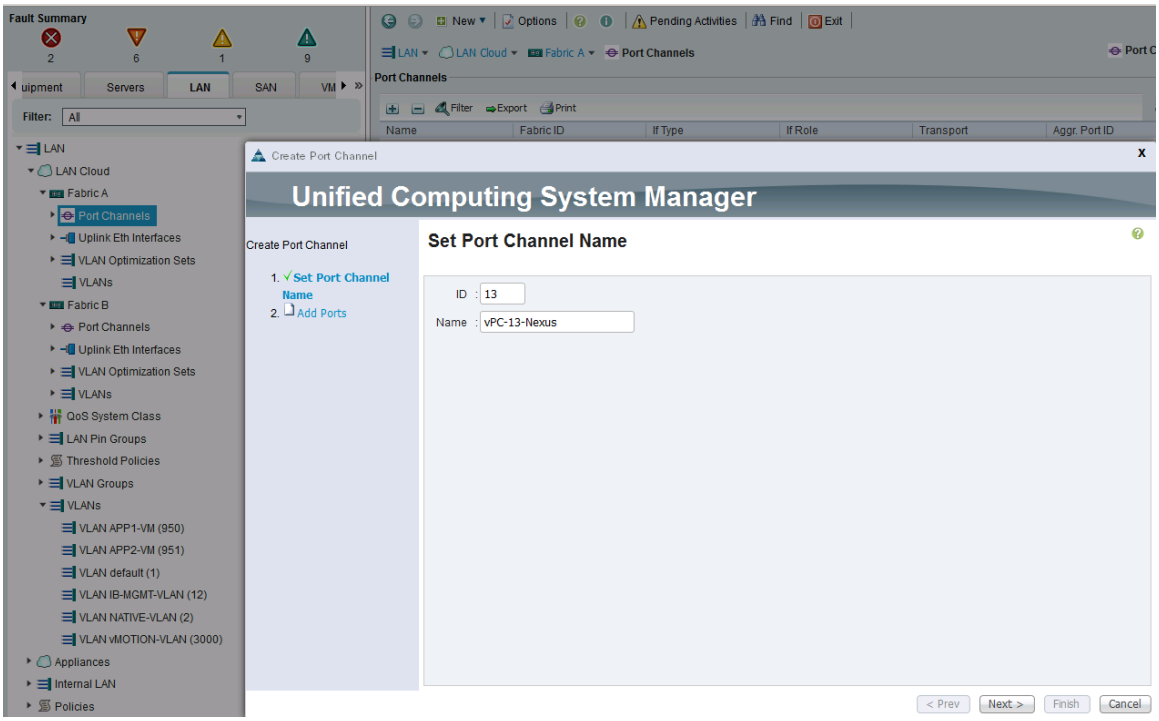

5. In the Create Port Channel window, select the ports to put in the channel (for example, Eth1/15 and Eth1/16). Click Finish to create the port channel.

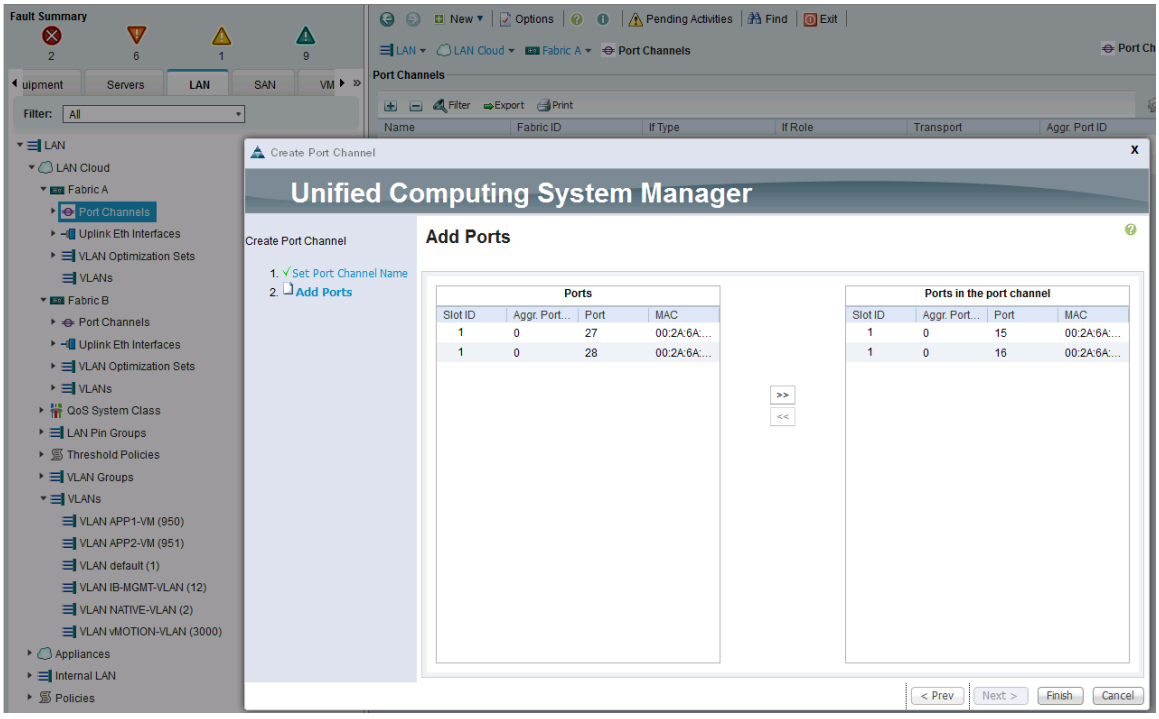

6. Verify the resulting configuration is as shown below.

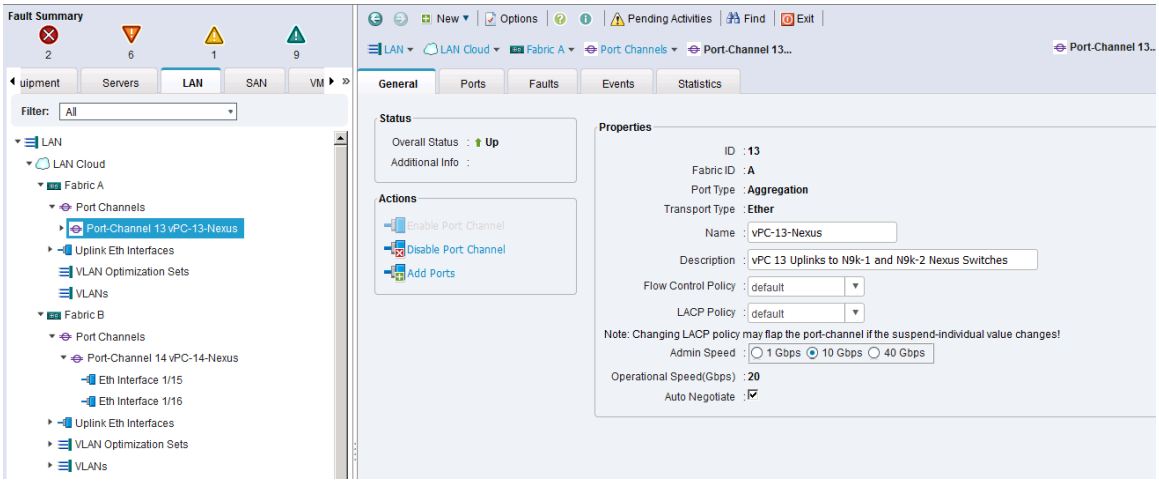

7. Repeat above steps for Fabric B and verify the configuration is as shown below.

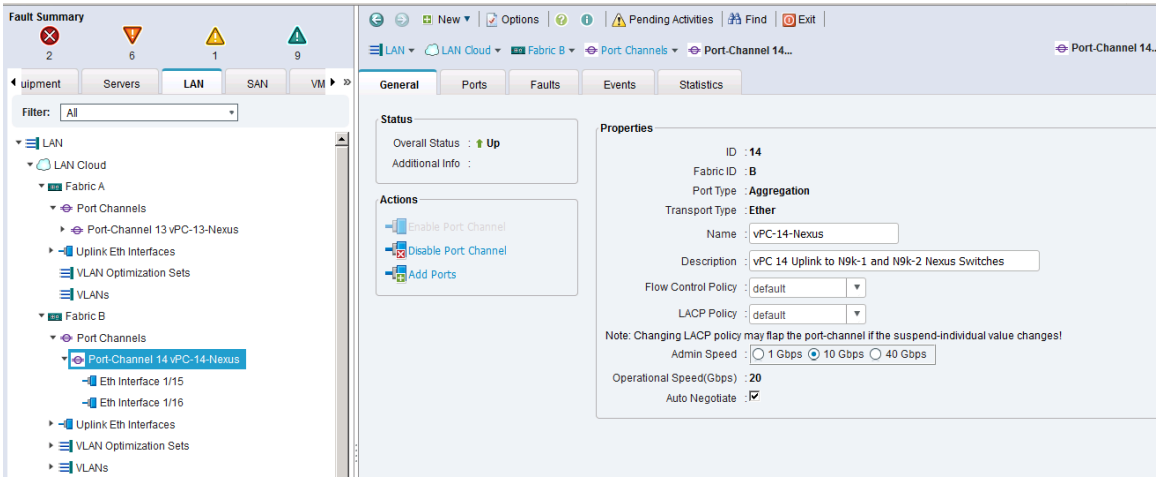

Enable Fibre Channels Ports to Cisco MDS 9100 Series Switches

Cable the following FC connections as specified in the figure below. Complete the steps below to configure the Fabric Interconnects ports that connect to upstream Cisco MDS switches.

Figure 12 FC Connections

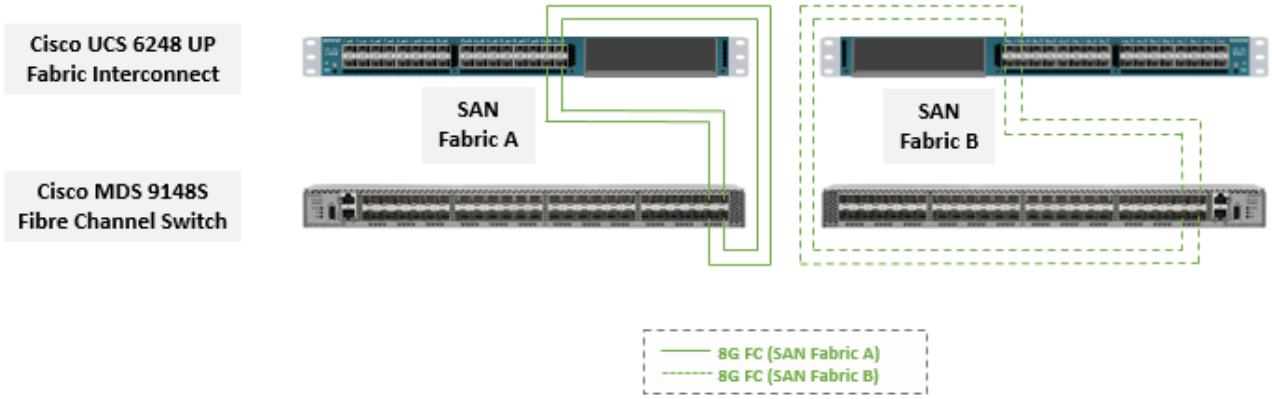

1. From Cisco UCS Manager, click Equipment tab in the navigation pane.
2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.
3. Expand FC Ports.
4. Select the first port (for example, `Port 29`) that connects to the Cisco MDS-A switch for SAN Fabric A, right-click and Configure as Uplink Port, click Yes to confirm. Repeat for the following 3 ports (for example, `Ports 30 – 32`) that connect to the same Cisco MDS switch.
5. Repeat above steps on the Fabric Interconnect B to configure the ports connected to the Cisco MDS-B switch in SAN Fabric B.

Create VSAN for Fibre Channel Interfaces

Create VSAN for SAN Fabric A

1. From Cisco UCS Manager, navigate to the SAN tab.
2. Select SAN > SAN Cloud > Fabric A > VSANs.
3. Right click and select Create VSAN. The VSAN configuration will need to match on the upstream Cisco MDS-A switch.

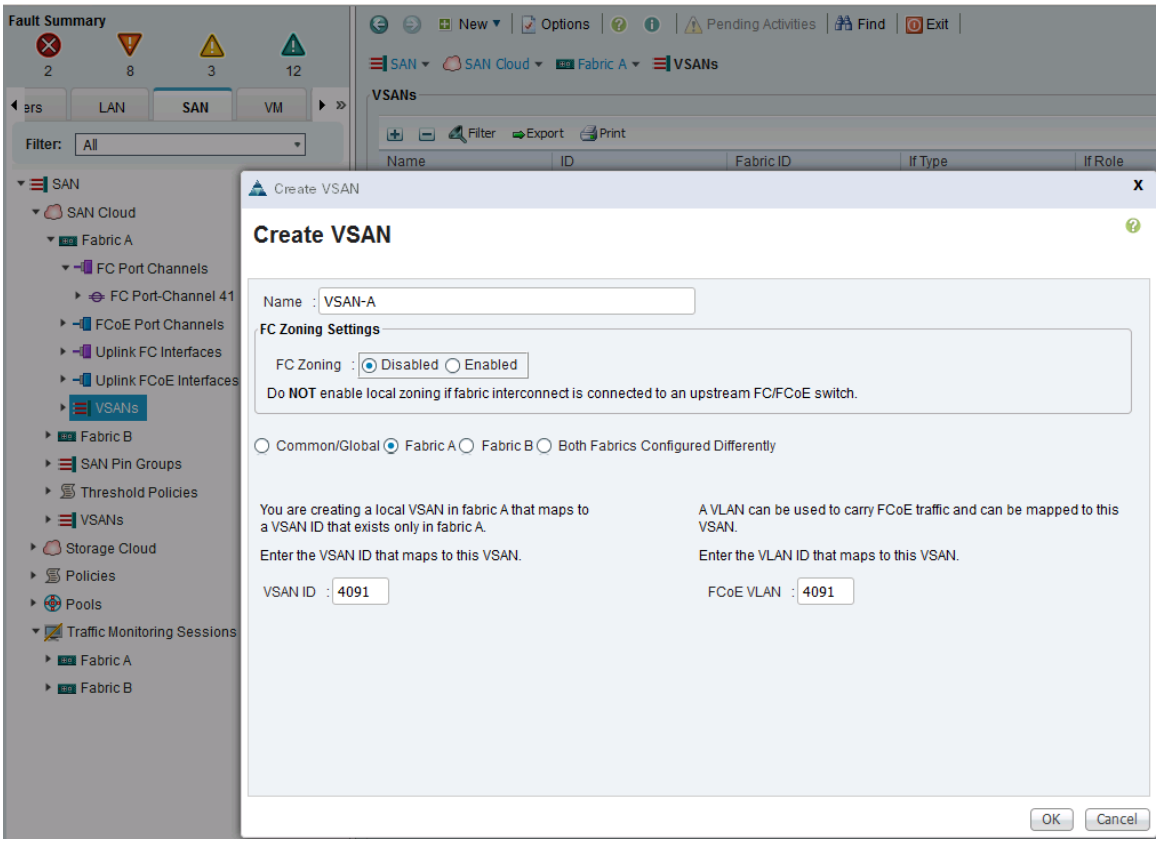

4. Enter a VSAN name for Fabric A (for example, **VSAN-A**), the VSAN ID (for example, **4091**) and FCoE VLAN ID (for example, **4091**). Keep the FC Zoning disabled. Select Fabric A radio button.

Create VSAN for SAN Fabric B

1. From Cisco UCS Manager, navigate to the SAN tab.
2. Select SAN > SAN Cloud > Fabric B > VSANs.
3. Right click and select Create VSAN. The VSAN configuration will need to match on the upstream Cisco MDS-B switch.

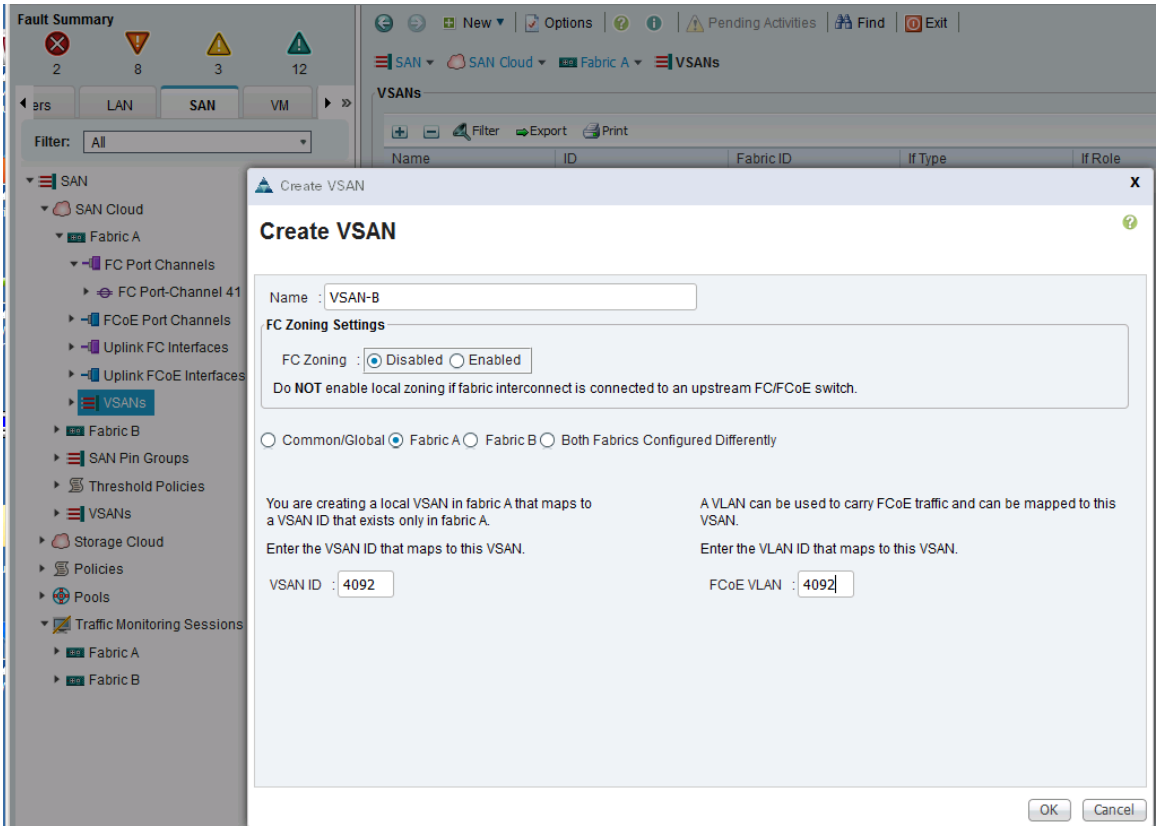

4. Enter a VSAN name for Fabric B (for example, `VSAN-B`), the VSAN ID (for example, `4092`) and FCoE VLAN ID (for example, `4092`). Keep the FC Zoning disabled. Select Fabric B radio button.

Configure Port Channels on Fibre Channel Uplinks to Cisco MDS Switches

Configure Port Channels from Fabric A to Cisco MDS-A

1. From Cisco UCS Manager, navigate to the SAN tab.
2. Select SAN > SAN Cloud > Fabric A > VSANs.
3. From Cisco UCS Manager, navigate to the SAN tab > SAN Cloud > Fabric A > FC Port Channels.
4. Right click and select Create Port Channel. The port channel configuration should match the upstream Cisco MDS-A switch.

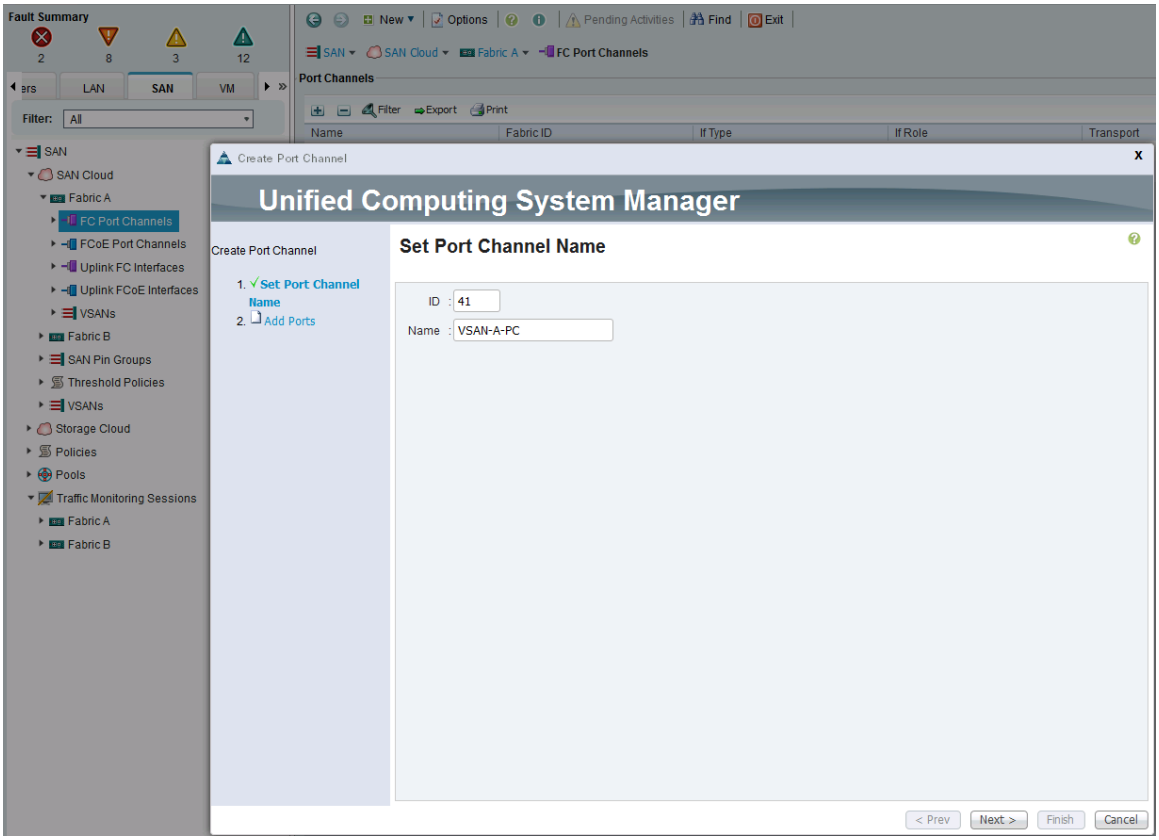

5. Select a port channel number (for example, `41`) and assign it a name (for example, `VSAN-A-PC`). Click Next.
6. In the Add Ports window, select the previously configured four ports. And then, click on the >> button.

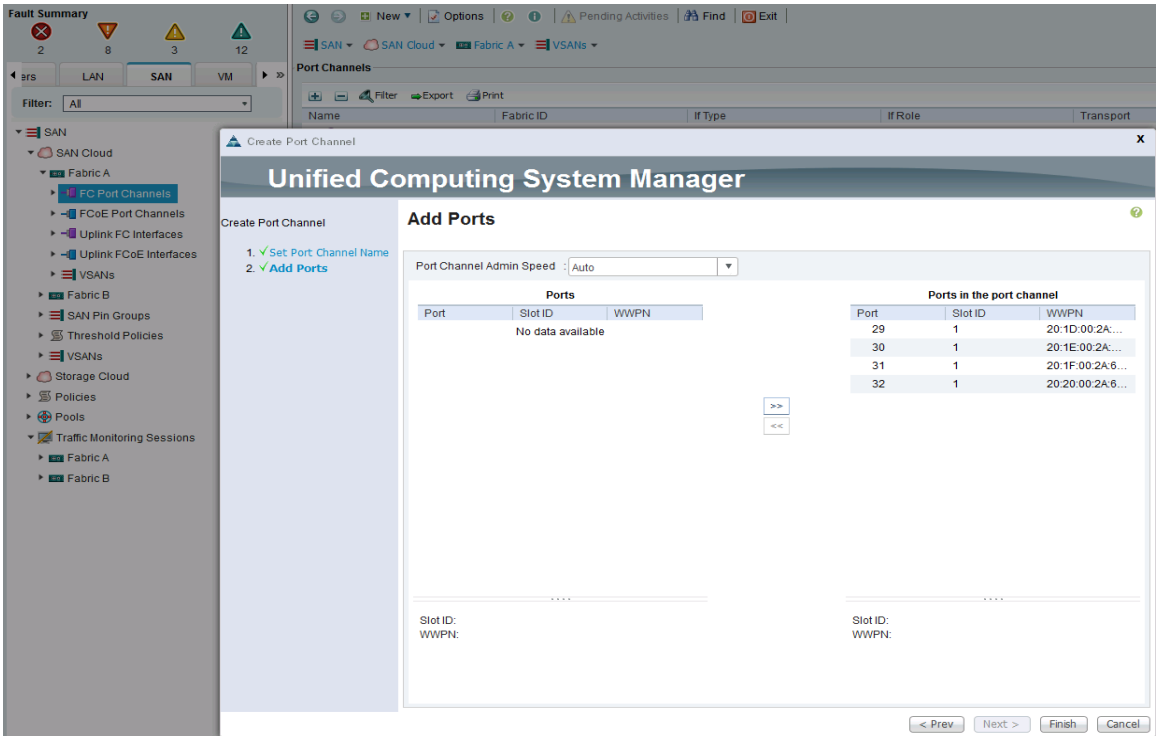

7. Click Finish to create the port channel to Cisco MDS-A switch.
8. Navigate to the newly created port channel under Fabric A and change the VSAN to be the one created earlier for Fabric A (for example, `4091`).

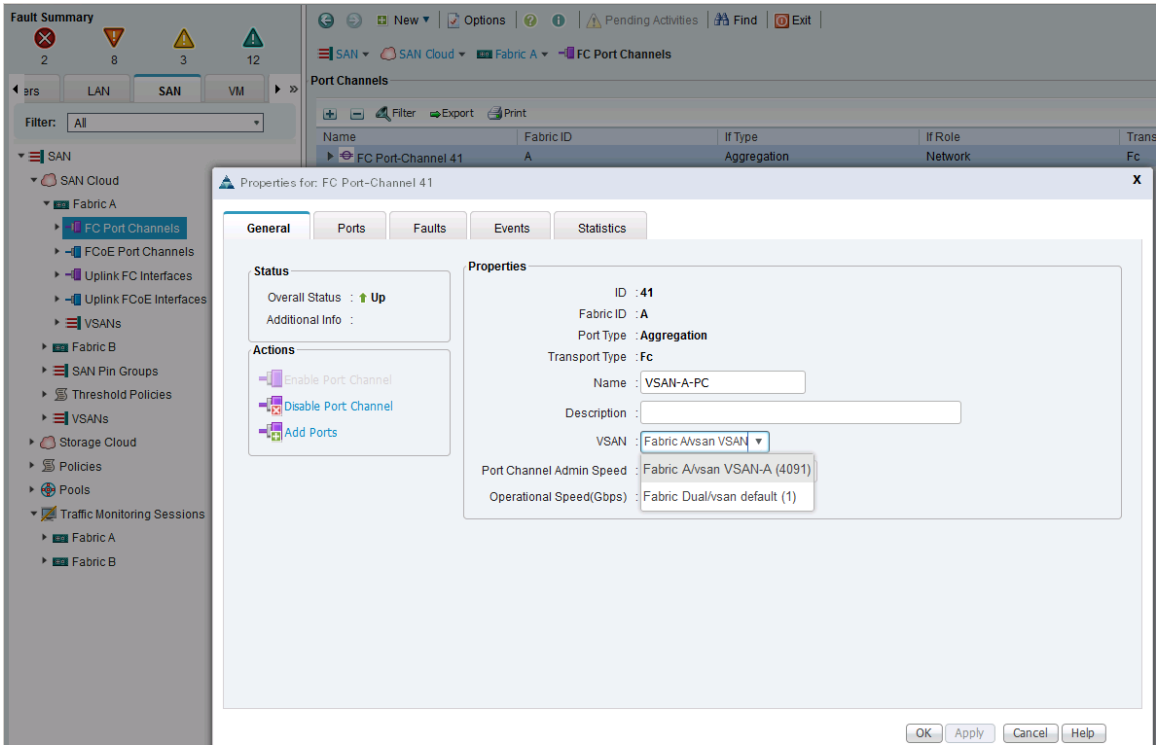

Configure Port Channels from Fabric B to Cisco MDS-B

1. From Cisco UCS Manager, navigate to SAN tab.
2. Select SAN > SAN Cloud > Fabric B > VSANs.
3. From Cisco UCS Manager, navigate to SAN tab > SAN Cloud > Fabric B > FC Port Channels.
4. Right click and select Create Port Channel. The port channel configuration should match the upstream Cisco MDS-A switch.

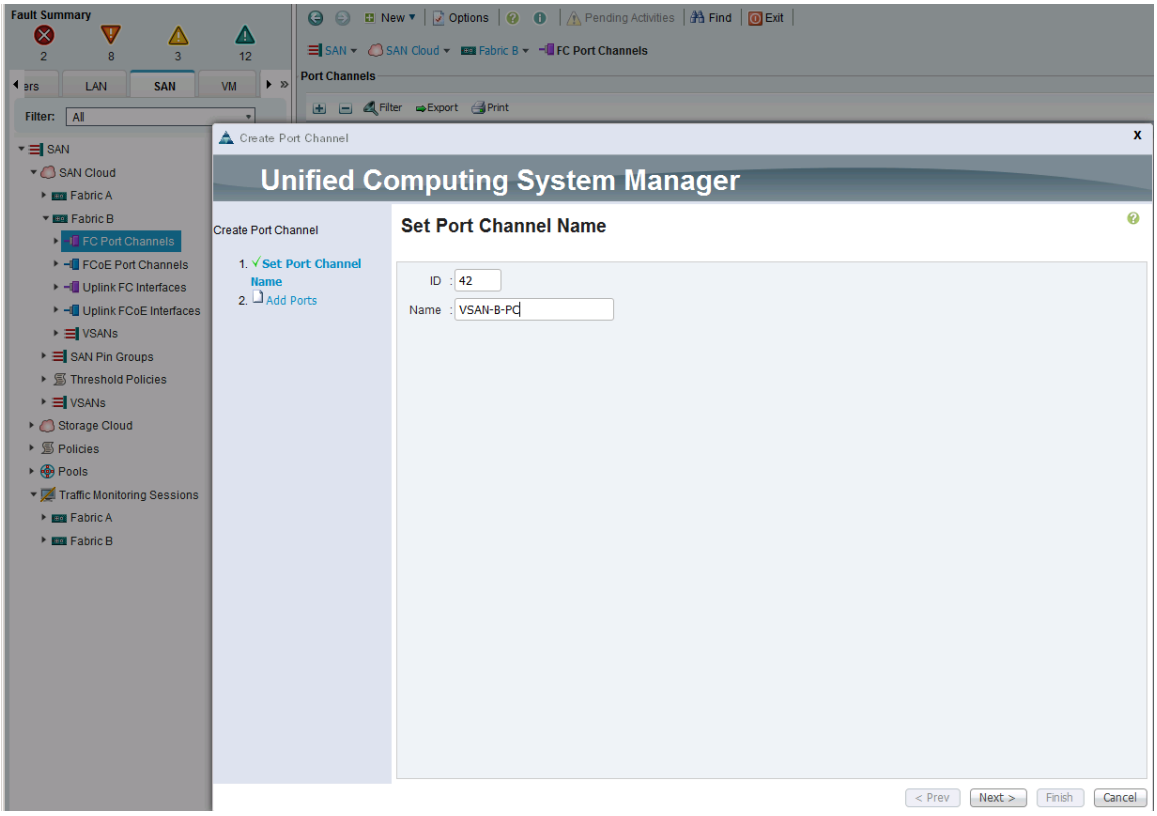

5. Select a port channel number (for example, `42`) and assign it a name (for example, `VSAN-B-PC`). Click Next.
6. In the Add Ports window, highlight the 4 previously configured ports and click on the >> button to add the ports to the port channel.

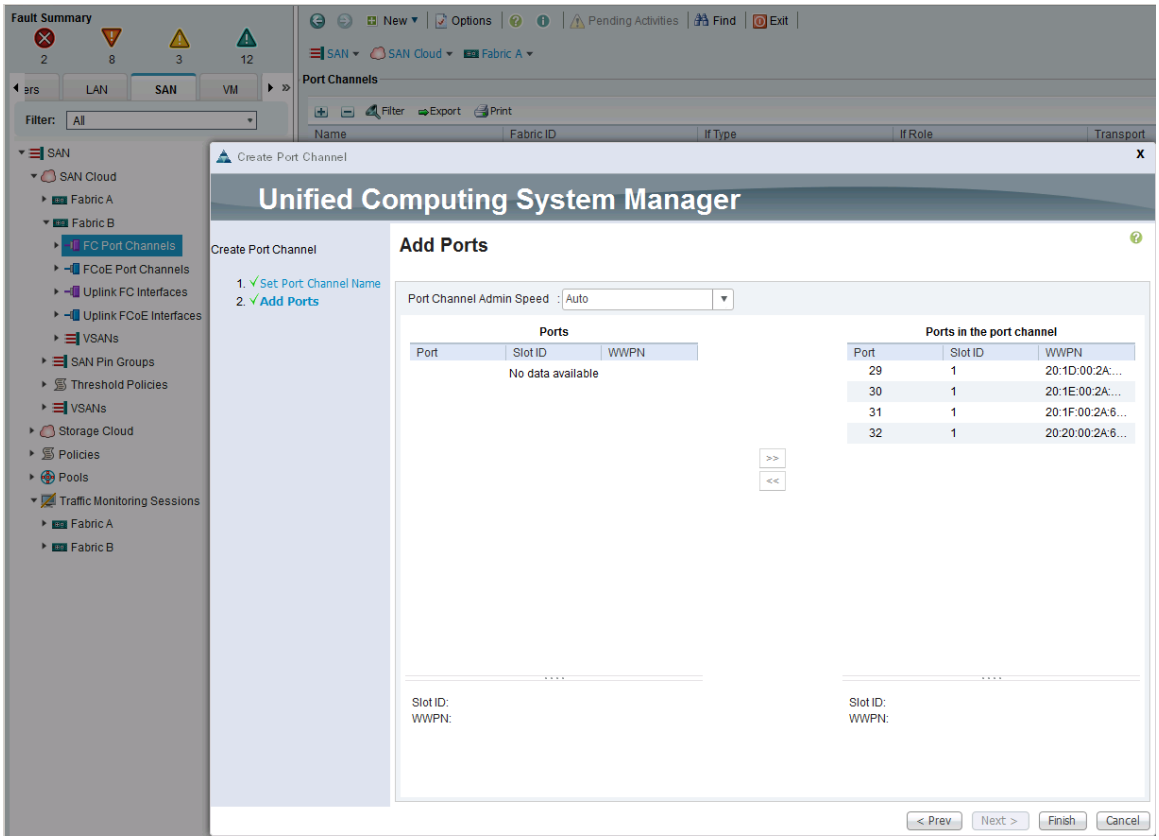

7. Click Finish to create the port channel to Cisco MDS-B switch.
8. Navigate to the newly created port channel under Fabric B and change the VSAN to be the one created earlier.

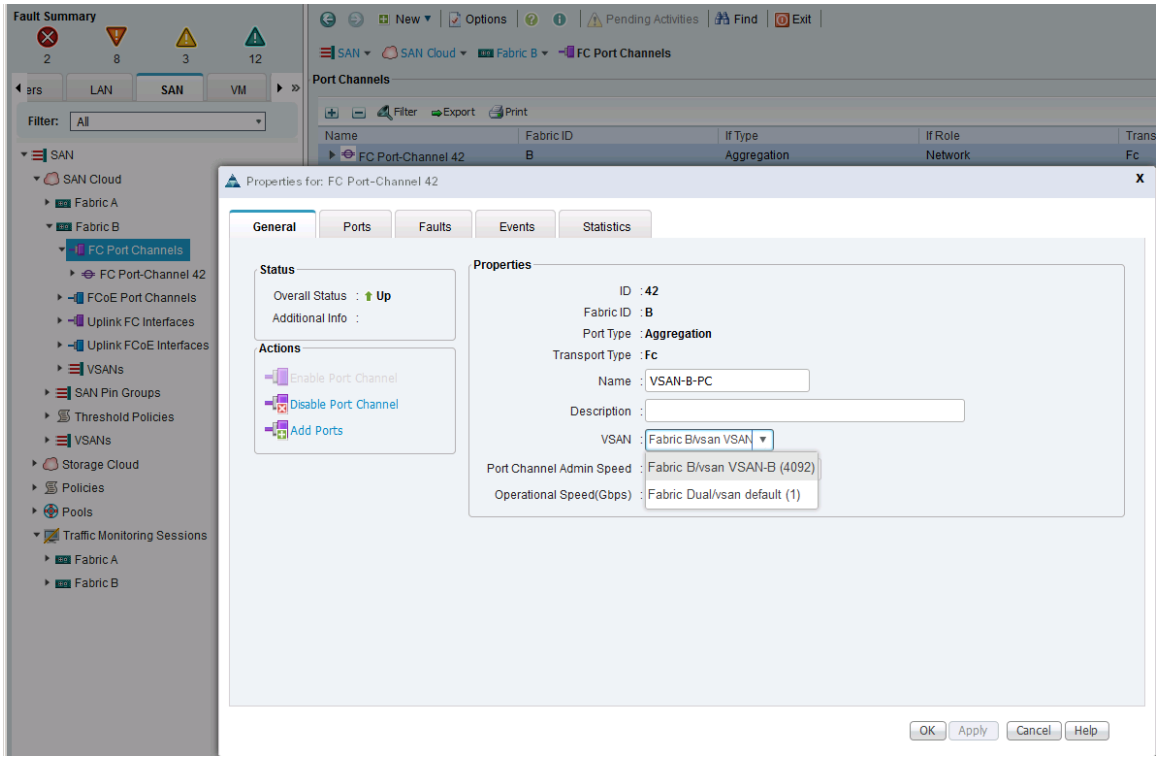

Cisco UCS Configuration Backup

The Cisco UCS Configuration should be backed up. For details on how to do the backup, see: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/Admin-Management/3-1/b_Cisco_UCS_Admin_Mgmt_Guide_3_1/b_Cisco_UCS_Admin_Mgmt_Guide_3_1_chapter_01001.html

Cisco UCS Configuration – LAN

LAN Configuration Workflow

The workflow below shows the configuration required to create a vNIC template. The vNIC Templates encapsulate the LAN configuration aspect of Cisco UCS. Two vNIC templates are created in the SmartStack design, one through Fabric A and another through Fabric B for redundancy. The subsections that follow will cover the configuration of the individual steps in the work flow below.

Figure 13 vNIC Template Creation Workflow

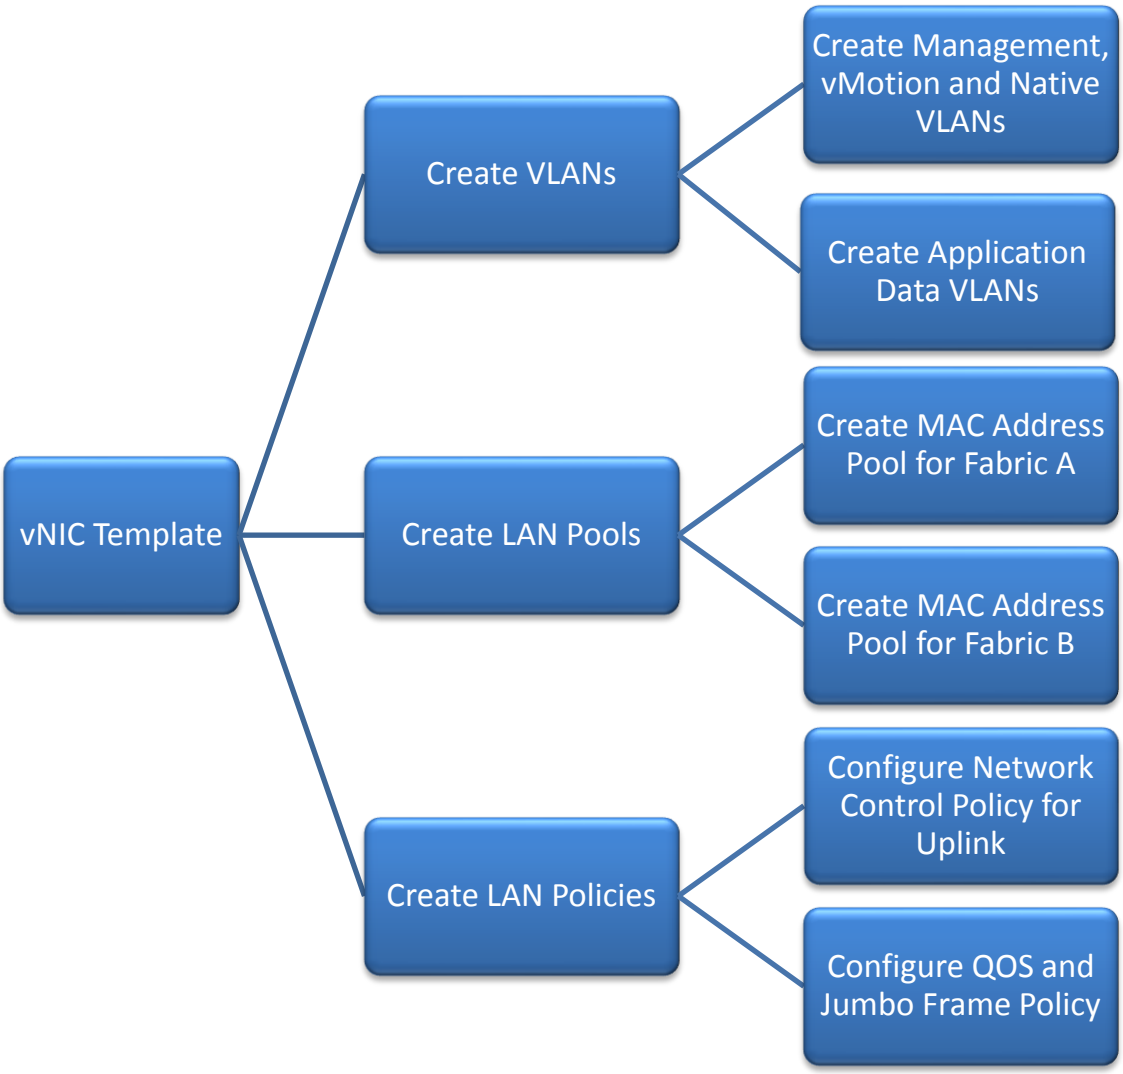

Create VLANs

Create Management VLAN on Uplink ports to Cisco Nexus 9000 Series Switches

The management VLAN is necessary for in-band management access to Cisco UCS hosts and virtual machines on hosts.

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > VLANs.
3. Right-click and select Create VLANs. Specify a name (for example, `IB-MGMT-VLAN`) and VLAN ID (for example, `12`).
4. If the newly created VLAN is a native VLAN, select VLAN, right-click and select Set as Native VLAN from the list. Either option is acceptable, but it needs to match what the upstream switch is set to.

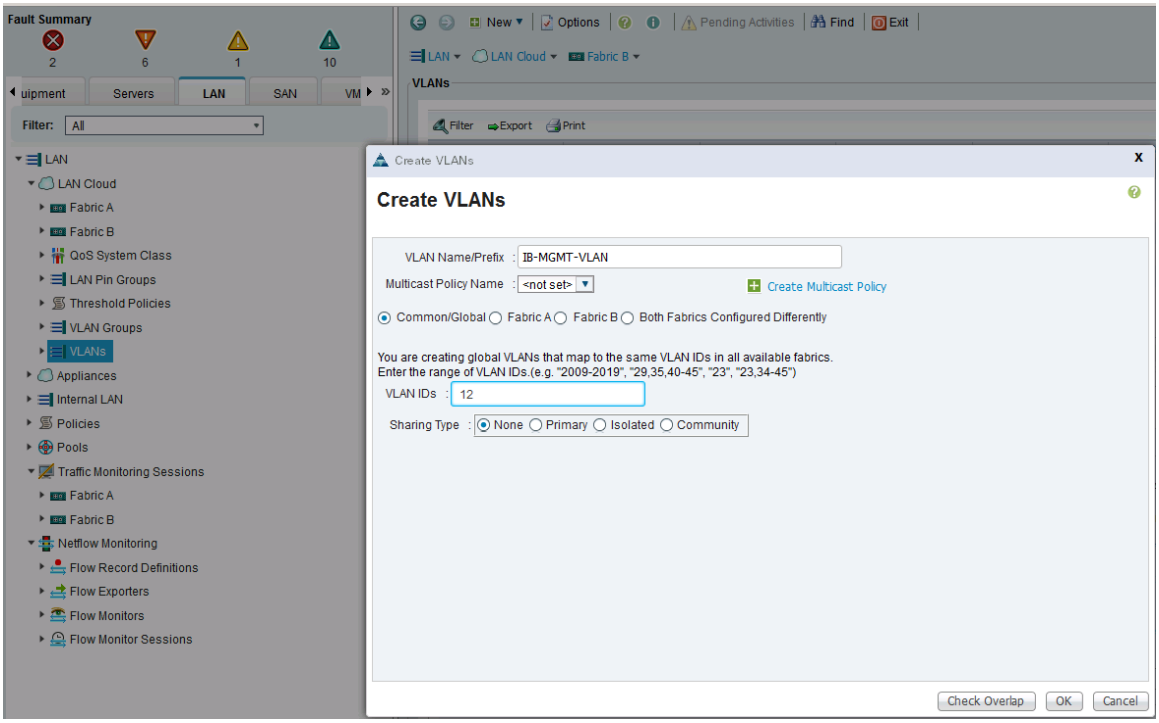

Create vMotion VLAN on Uplink ports to Cisco Nexus 9000 Series Switches

The vMotion VLAN is necessary for supporting vMotion between hosts in the Cisco UCS domain and other domains in the data center. vMotion will use a dedicated vNIC on each host per VMware vSphere best practices. Steps for creating the VLAN are identical to the previous step.

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > VLANs.
3. Right-click and select Create VLANs. Specify a name (for example, `vMOTION-VLAN`) and VLAN ID (for example, `3000`).
4. If the newly created VLAN is a native VLAN, select VLAN, right-click and select the option Set as Native VLAN from the list. Either option is acceptable and should match upstream switch setting for the VLAN.

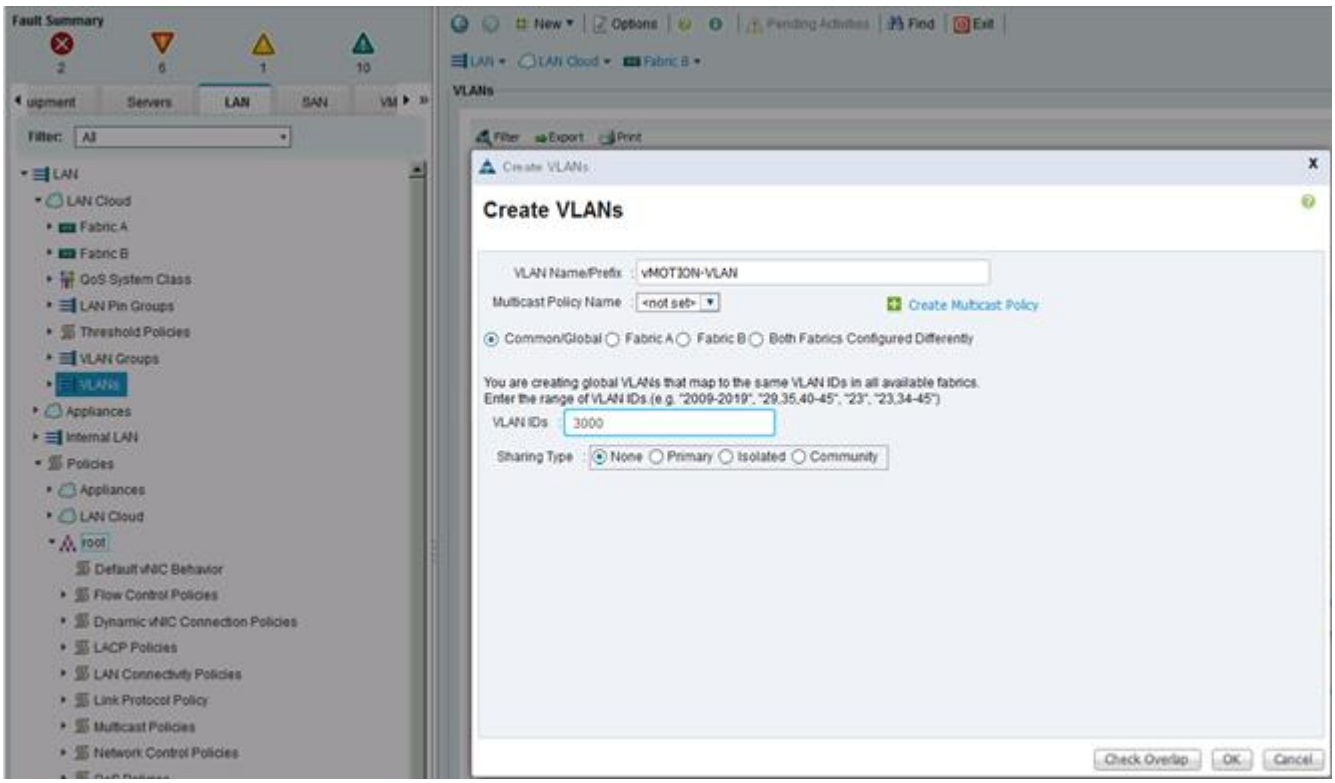

Create Application Data VLANs on Uplink ports to Cisco Nexus 9000 Series Switches

To create Application Data VLANs that needs to be trunked through the Cisco Nexus 9000 series switches to other parts of the network, repeat the steps above to create the management and vMotion VLANs. To validate this SmartStack architecture, the following applications VLANs were created.

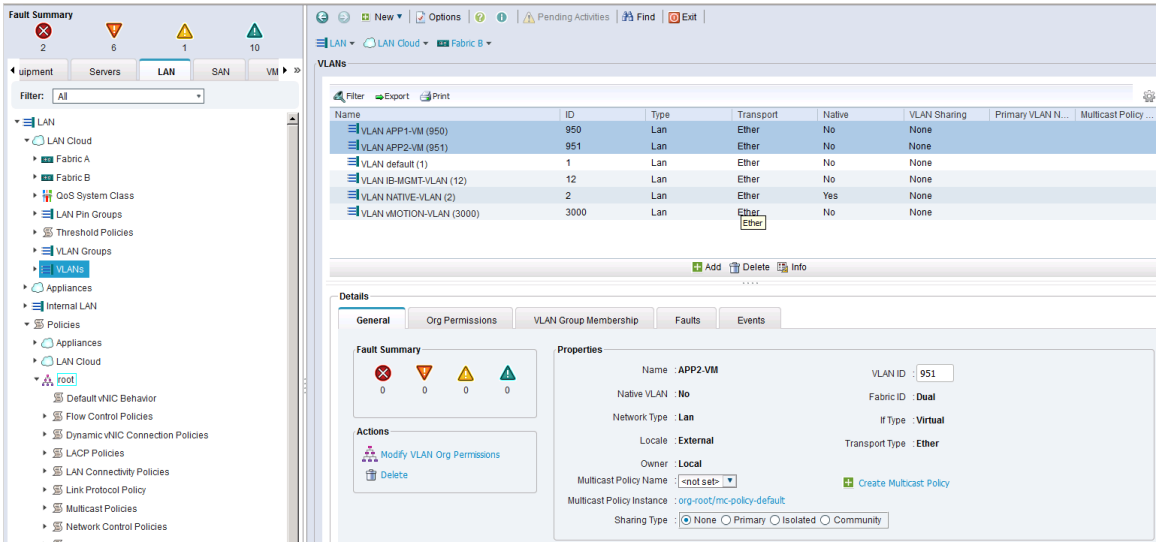

Change Default Native VLAN on Uplink ports to Cisco Nexus 9000 Series Switches

Per security best practices, the default native VLAN (for example, `VLAN 1`) was changed to a newly created native VLAN (for example, `VLAN 2`).

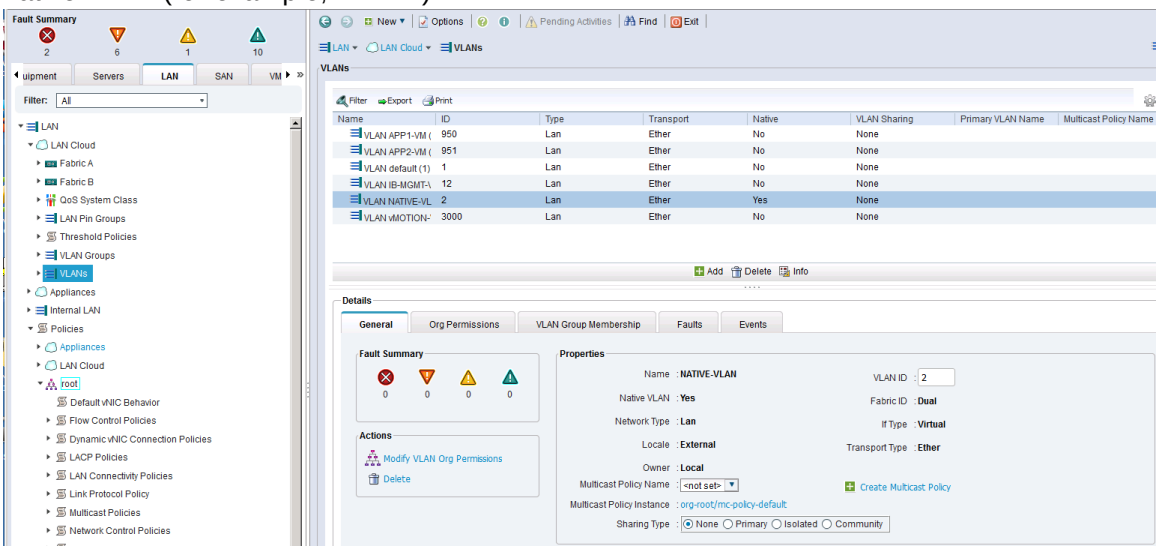

VLAN Summary View

A summary of the VLANs created in the previous steps are shown below.

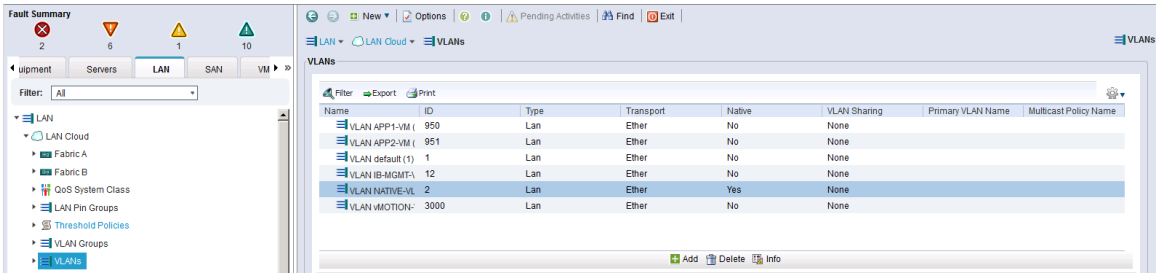

Create LAN Pools

Create MAC Address Pools

Cisco UCS allows a MAC Address Pool to be created and allocated in such a way that makes it easy for troubleshooting if needed. In this design, the 4th and 5th octet in the mac-address is modified to reflect whether the host using that mac-address will use Fabric A (for example, `AA:AA`) or Fabric B(for example, `BB:BB`).

Create MAC pool for Fabric Interconnect A

The MAC addresses in this pool will be used for traffic through Fabric Interconnect A.

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > Pools > root > MAC Pools.
3. Right-click and select Create Mac Pool.

Specify a name (for example, `MAC-POOL-FI-A`) that identifies this pool is specific to Fabric Interconnect

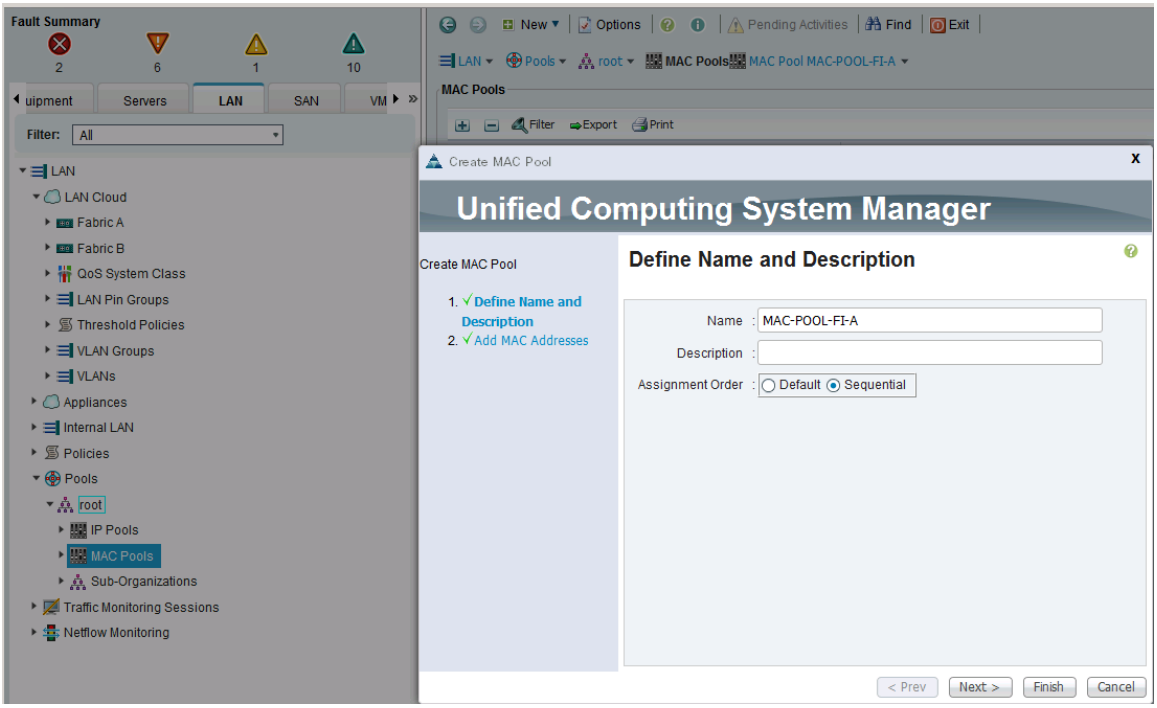

4. Select the Assignment Order as Sequential and click Next.

Click + to add a new MAC pool.

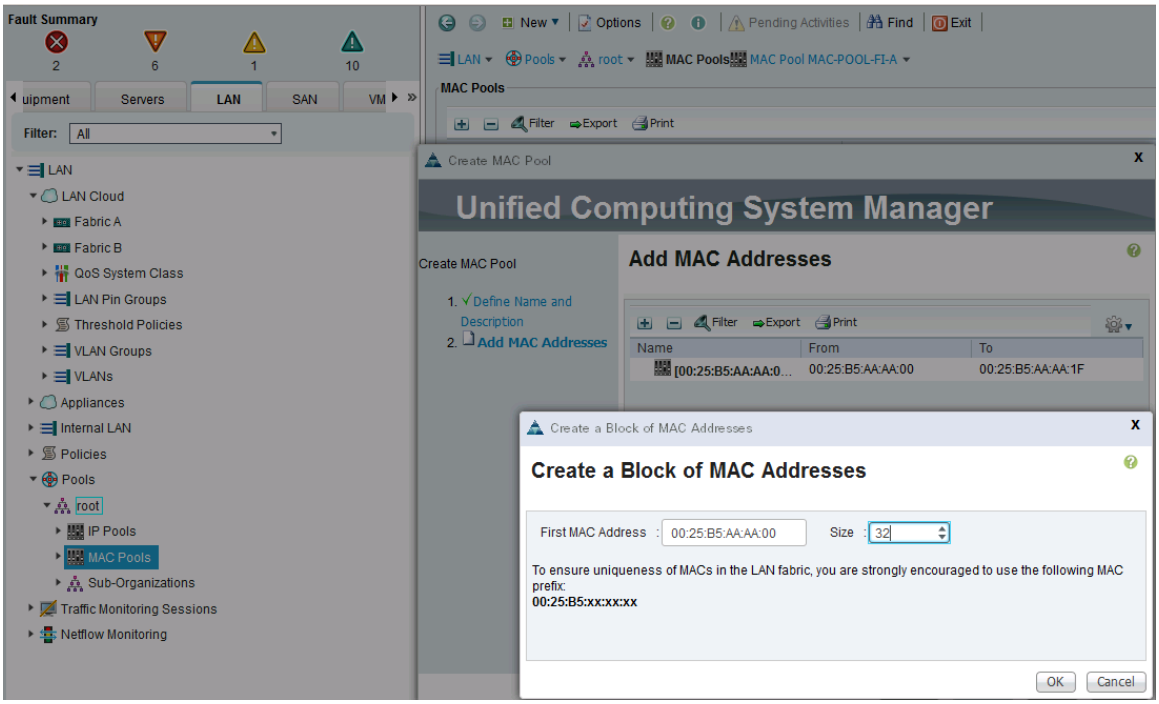

5. For ease-of-troubleshooting, change the 4th and 5th octet to AA:AA traffic using Fabric Interconnect A. Generally speaking, the first three octets of a mac-address should not be changed.
6. Select a size (for example, `32`) and select OK and then click Finish to add the MAC pool.

Create MAC pool for Fabric Interconnect B

The MAC addresses in the pool will be used for traffic using Fabric Interconnect B.

1. From Cisco UCS Manager, click LAN tab in the navigation pane.
2. Select LAN > Pools > root > Mac Pools.
3. Right-click and select Create Mac Pool.
4. Specify a name (for example, `MAC-POOL-FI-B`) that identifies this pool is specific to Fabric Interconnect B.

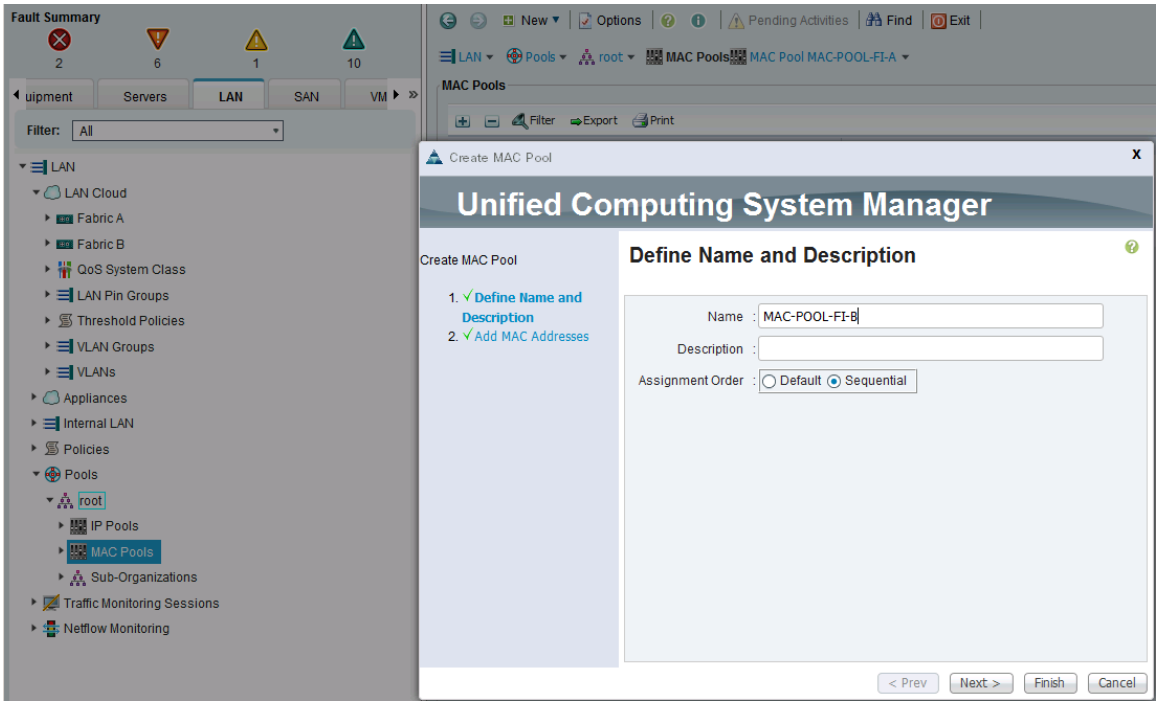

5. Select the Assignment Order as Sequential and click Next.
6. Click + to add a new MAC pool.

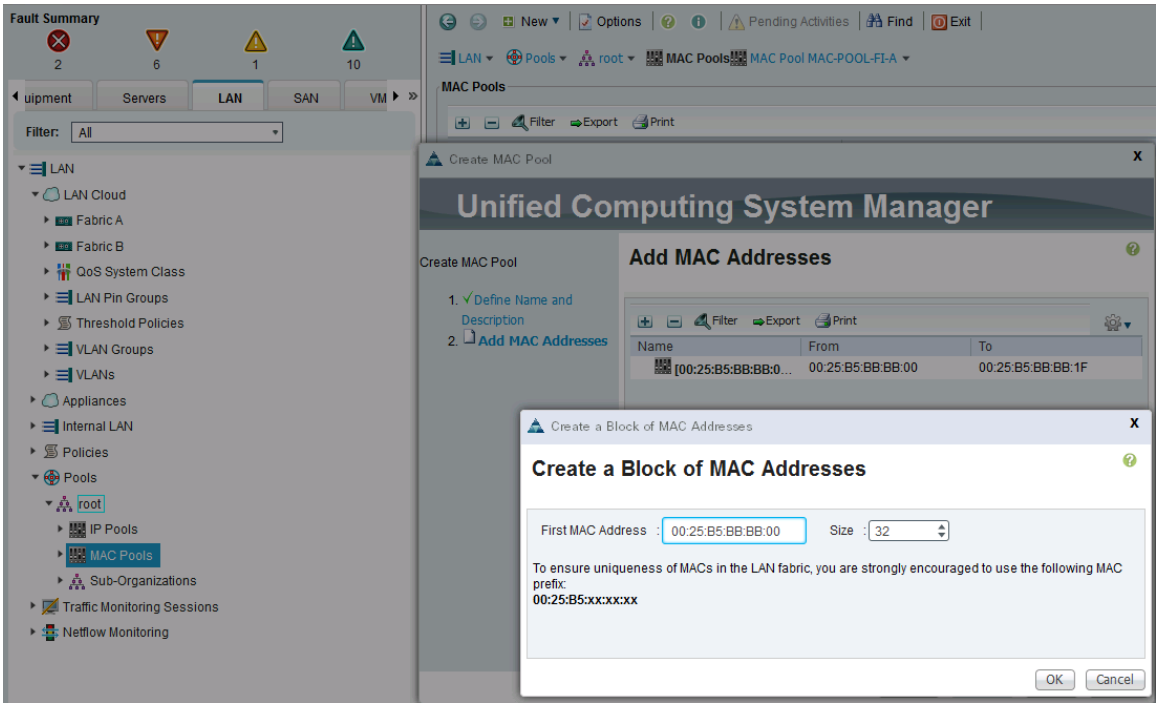

7. For ease-of-troubleshooting, change the 4th and 5th octet to BB:BB traffic using Fabric Interconnect A. Generally speaking, the first three octets of a mac-address should not be changed.
8. Select a size (for example, 32) and select OK and then click Finish to add the MAC pool.

MAC Pool Summary View

The resulting configuration is shown below.

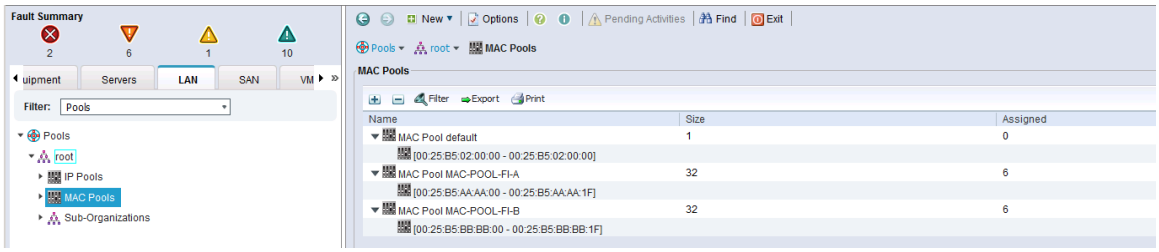

Create LAN Policies

Configure QoS and Jumbo Frame Policy

Proper QoS configuration is required to specify the priority of the traffic relative to each other. The CoS / QoS priority should be the same end-to-end from Nimble Storage array to the server blade. MTU configuration is also done through QoS policy to enable Jumbo frame support.

1. From Cisco UCS Manager, select LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QOS System Class.
3. Select the Enabled checkbox for the Best Effort.
4. In the Best Effort row, change the MTU to be a value of 9216.

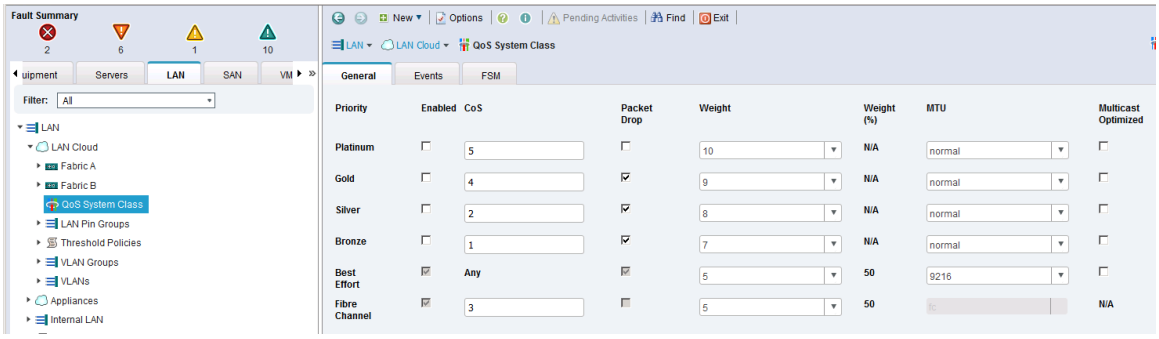

5. Click Save Changes and then OK to save.
6. Select LAN tab in the navigation pane.
7. Select LAN > Policies > root > QOS Policies.
8. Right-click and select Create QoS Policy.
9. Specify a policy Name (for example, `SmartStack-QoS`) and Priority as Best Effort. Click OK to save.

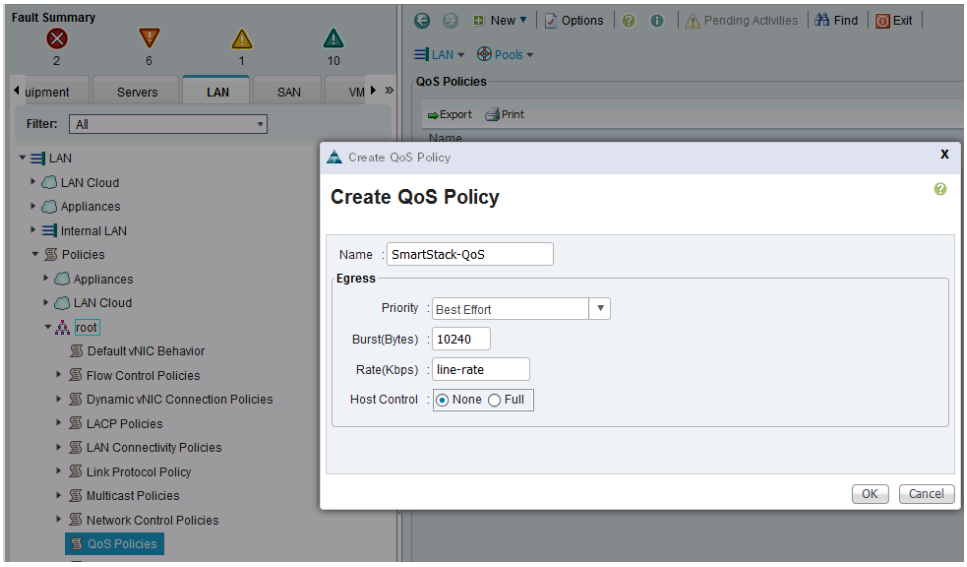

Configure Network Control Policy for Uplinks

For Uplink ports connected to Cisco Nexus switches, the policy is used to enable Cisco Discovery Protocol (CDP). CDP is useful for troubleshooting and allows devices to discover and learn about devices connected to the ports where CDP is enabled.

1. From Cisco UCS Manager, select LAN tab in the navigation pane.
2. Select LAN > Policies > LAN Cloud > root > Network Control Policies.
3. Right-click and select Create Network Control Policy.
4. Specify a Name for the network control policy (for example, `Enable_CDP`).
5. Select Enabled for CDP and default for setting for the remaining. Click OK to create network control policy.

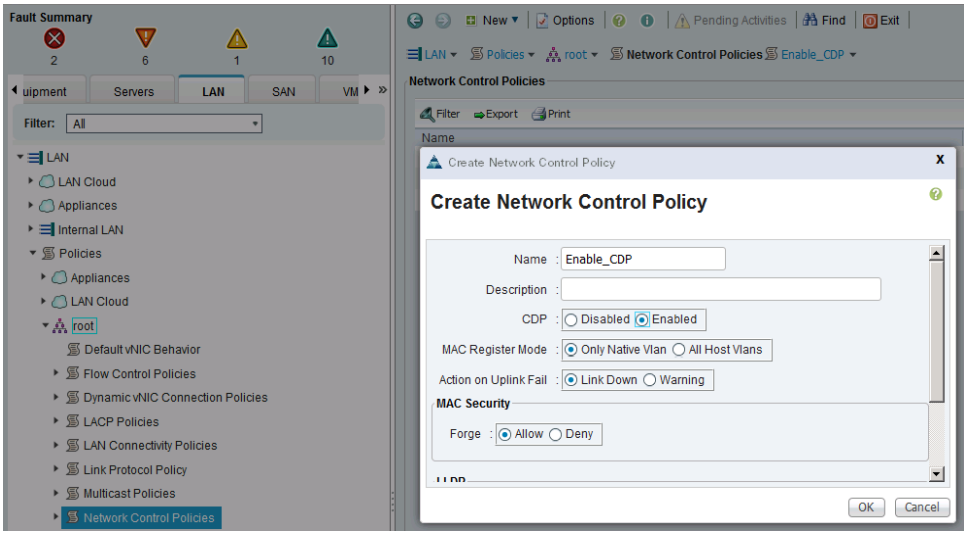

Create vNIC Templates

To create virtual network interface card (vNIC) templates for Cisco UCS hosts, complete the following steps. Two vNICs are created for redundancy – one through Fabric A and another through Fabric B. All host traffic is carried across these two vNICs in this design.

Create vNIC Template for Fabric A

Enable Failover option is left unchecked if a Cisco Nexus 1000V switch (optional) is used.

1. From Cisco UCS Manager, select LAN tab in the navigation pane.
2. Select LAN > Policies > root > vNIC Templates.

Right-click and select Create vNIC Template.

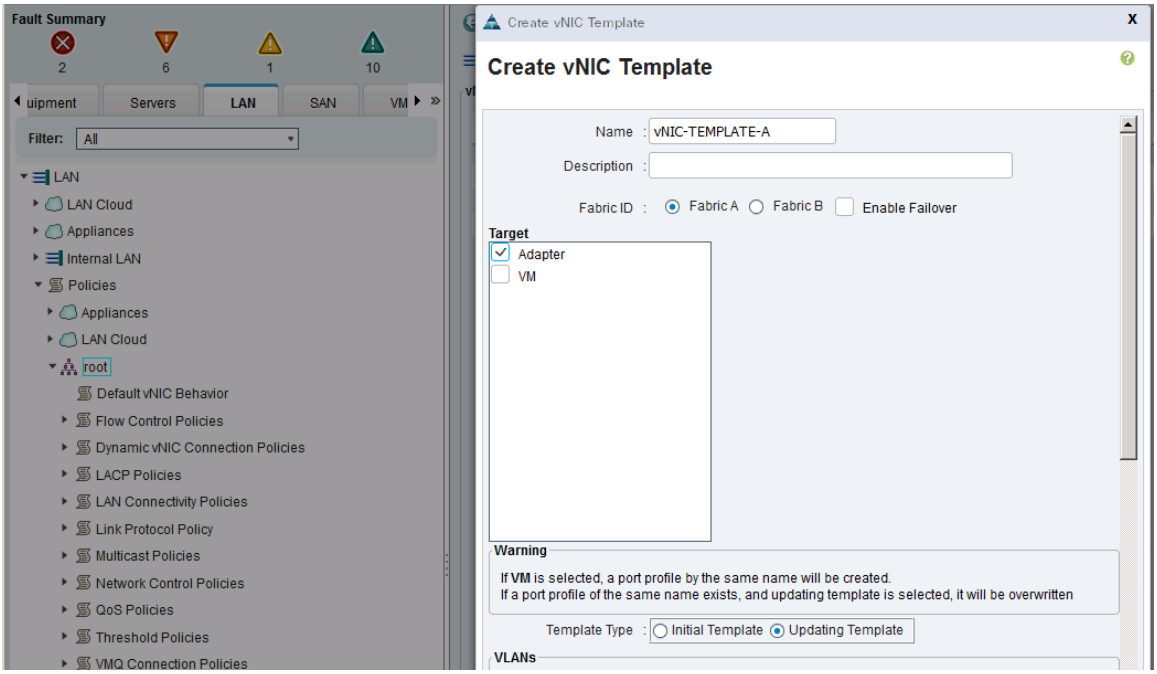

3. Specify a template Name (for example, `vNIC-TEMPLATE-A`) for the policy.
4. Keep Fabric A selected and leave Enable Failover checkbox unchecked.
5. Under Target, make sure that the VM checkbox is NOT selected.
6. Select Updating Template as the Template Type.

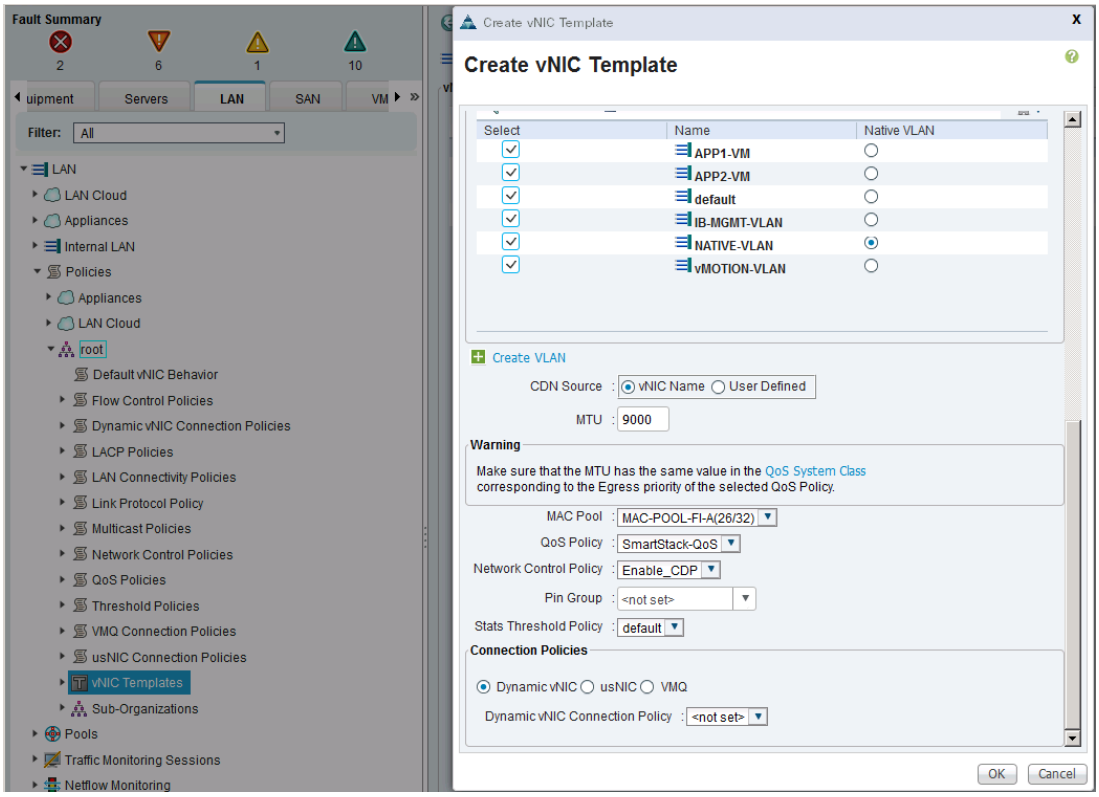

7. Under VLANs, select the checkboxes for all VLAN traffic that a host needs to see (for example, `IB-MGMT-VLAN, vMOTION-VLAN, NATIVE-VLAN, APP1-VM, APP2-VM`).
8. For MTU, enter `9000` since vMotion traffic will use this vNIC.
9. For MAC Pool, select the previously configured LAN pool (for example, `MAC-POOL-FI-A`).
10. For Network Control Policy, select the previously configured LAN policy (for example, `Enable_CDP`).
11. Click OK twice to create the vNIC template.
12. Choose the default values in the Connection Policies section.
13. The resulting configuration is as shown below.

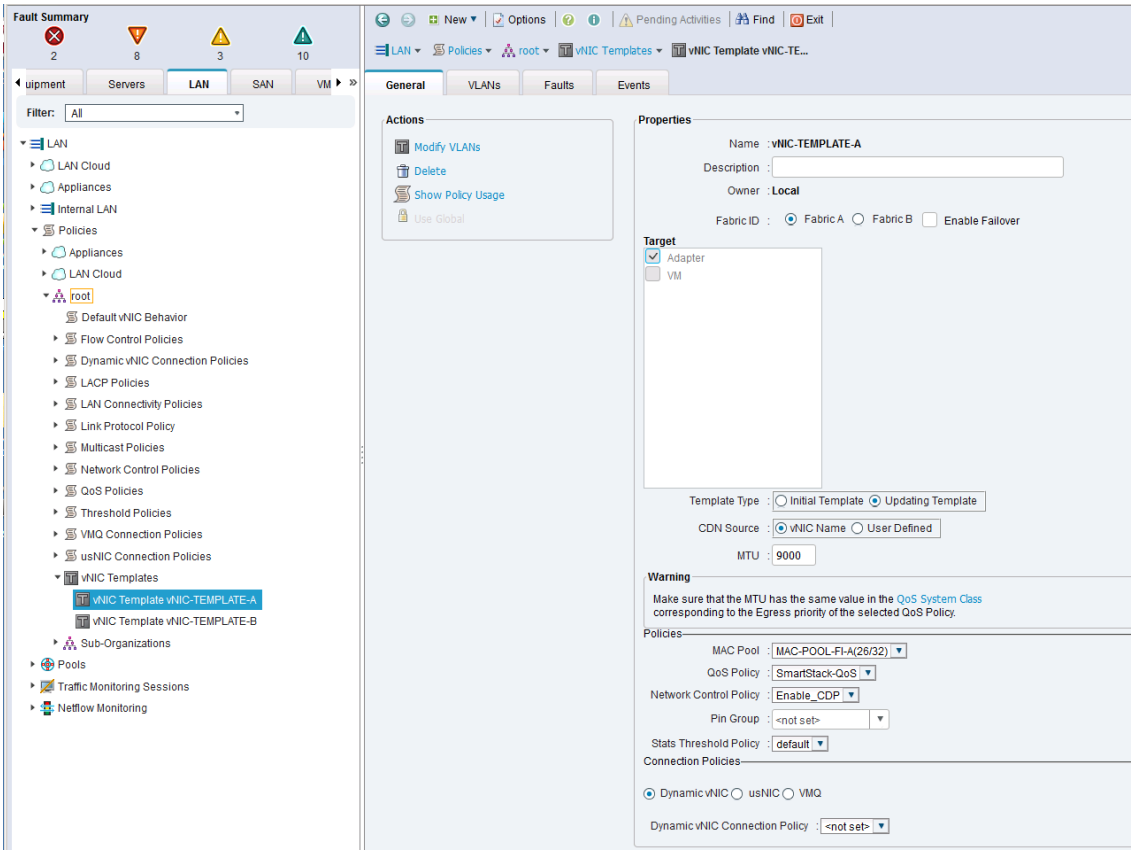

Create vNIC Template for Fabric B

Repeat above steps to create a vNIC (`vNIC-TEMPLATE-B`) template through Fabric B.

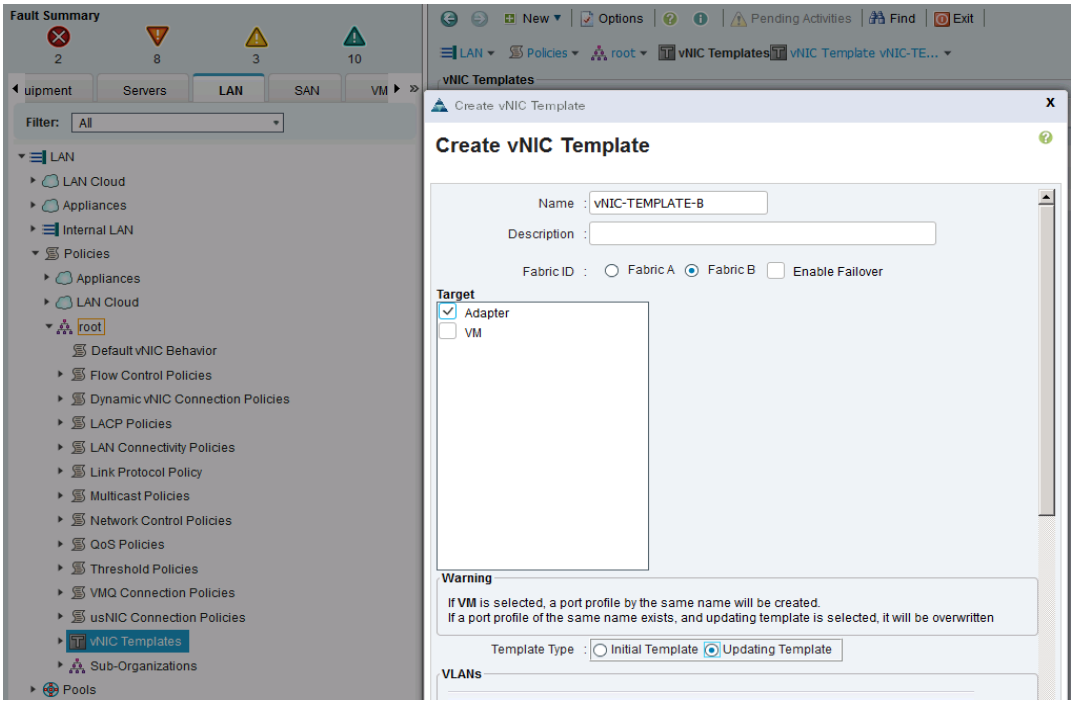

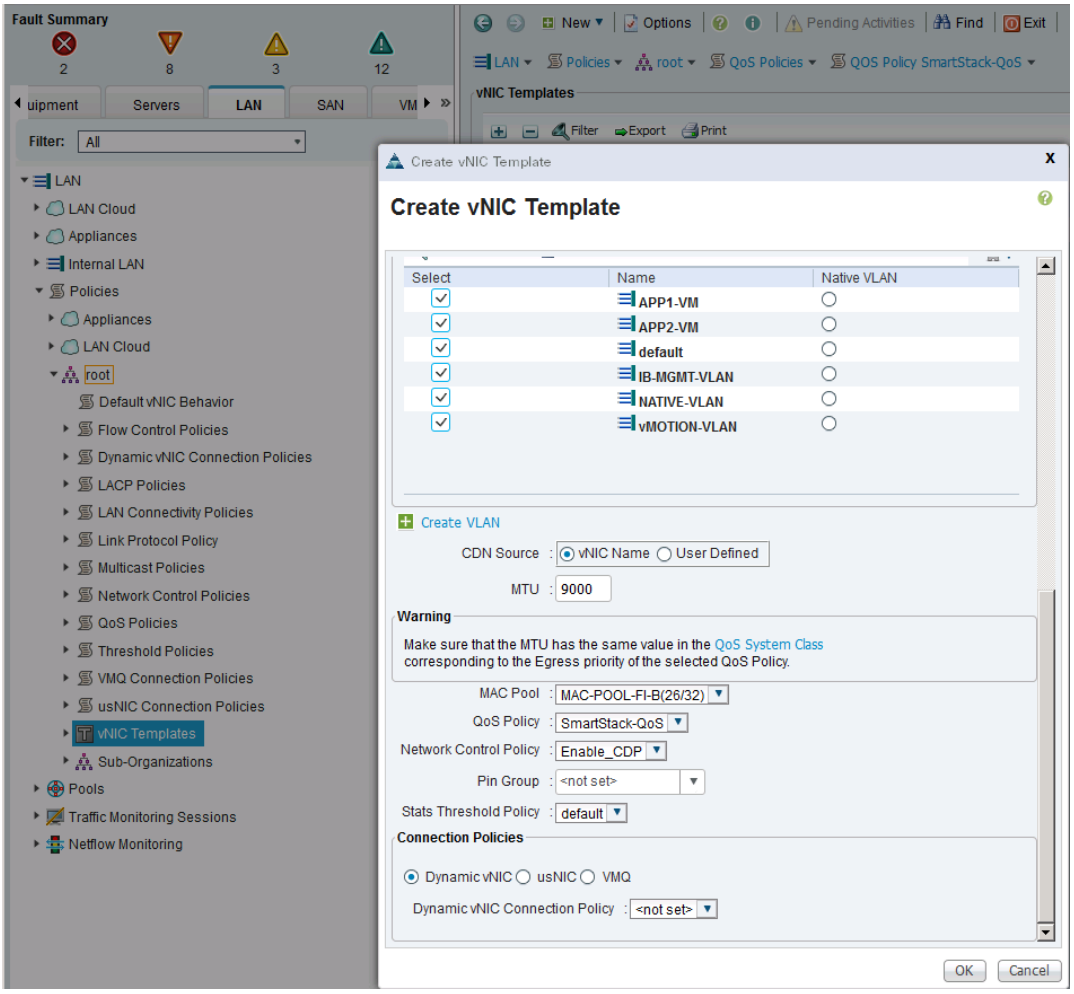

1. The resulting configuration is as shown below:

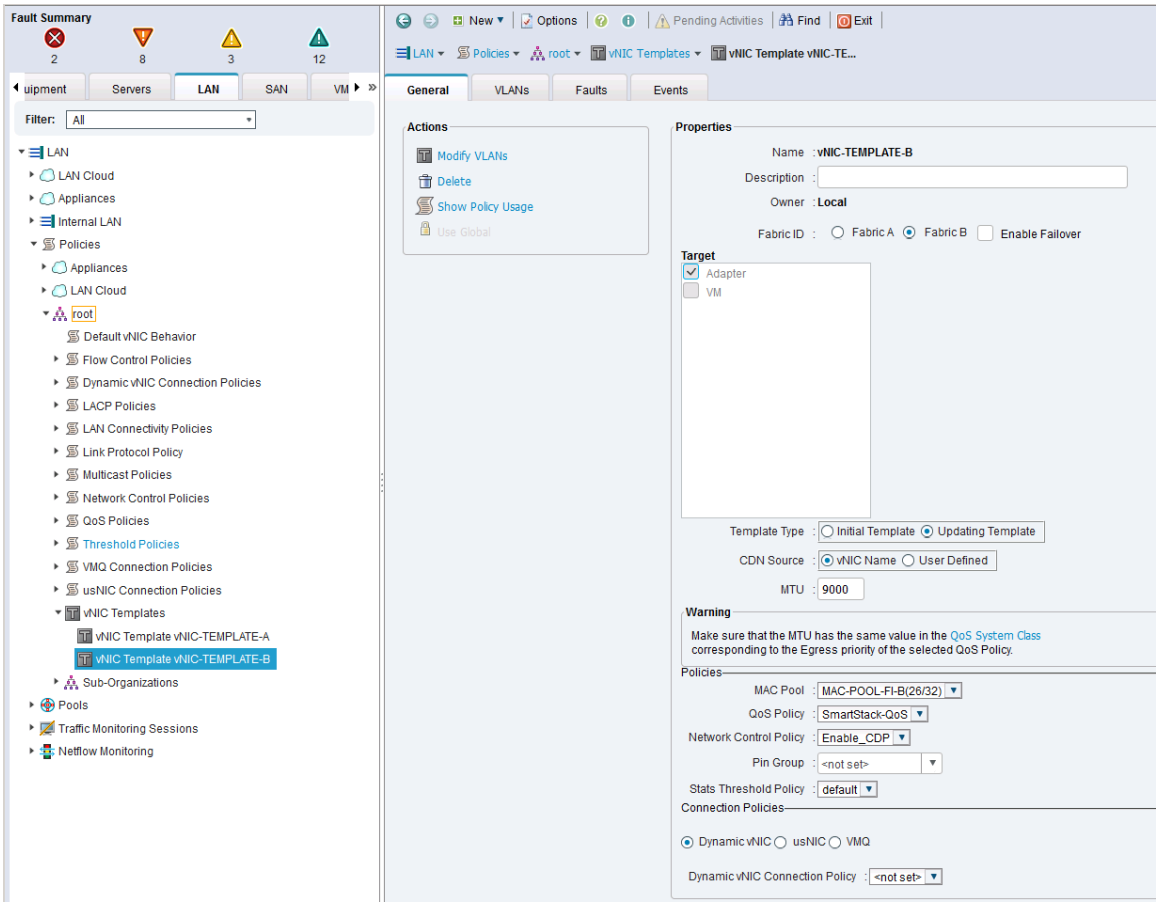

Cisco UCS Configuration – Server

Server Configuration Workflow

The workflow below shows the high level server configuration workflow. The subsections that follow will cover the configuration of the individual steps in the workflow.

Figure 14 Server Configuration Workflow

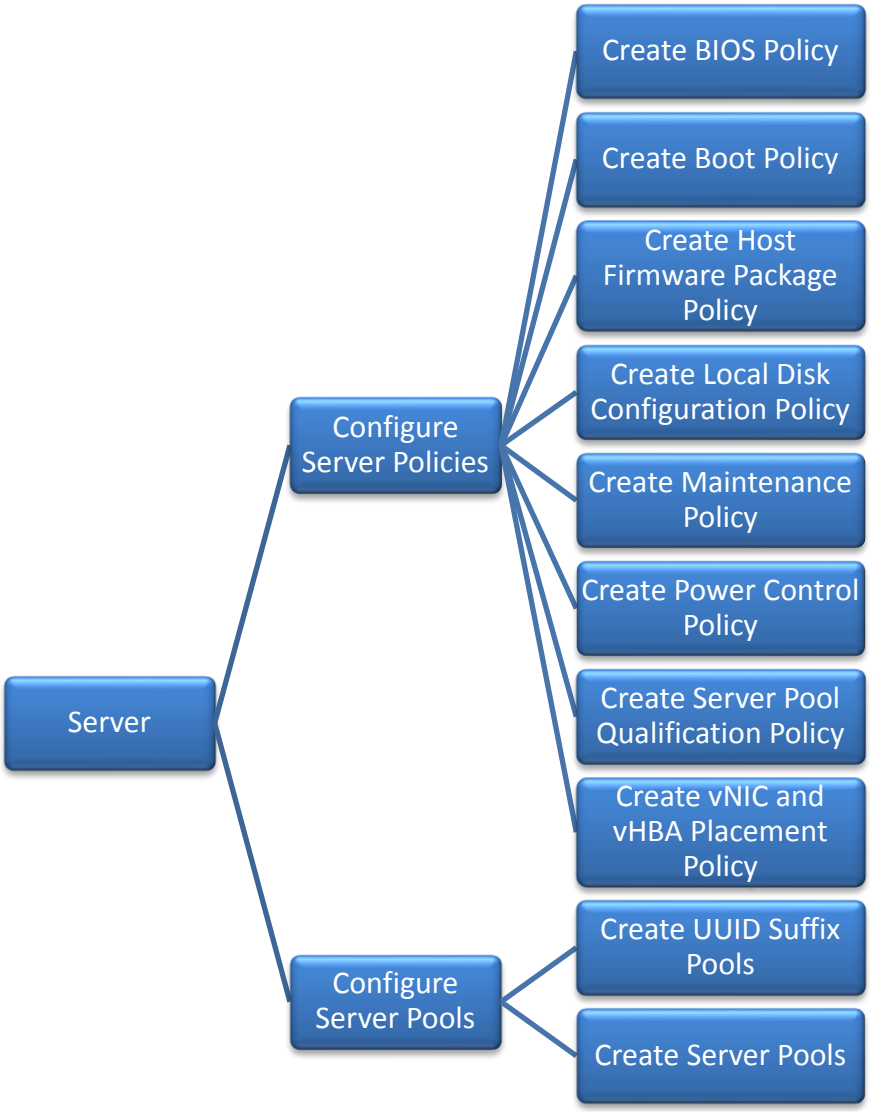

Configure Server Policies

The procedures for creating the different server policies deployed in a SmartStack environment are covered in this section.

Create BIOS Policy

To create a server BIOS policy for Cisco UCS hosts, complete the following steps:

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > BIOS Policies.
3. Right-click and select Create BIOS Policy.
4. In the Main screen, enter BIOS Policy Name (for example, `VM-Host-Infra`) and change the Quiet Boot setting to Disabled. Click Next.

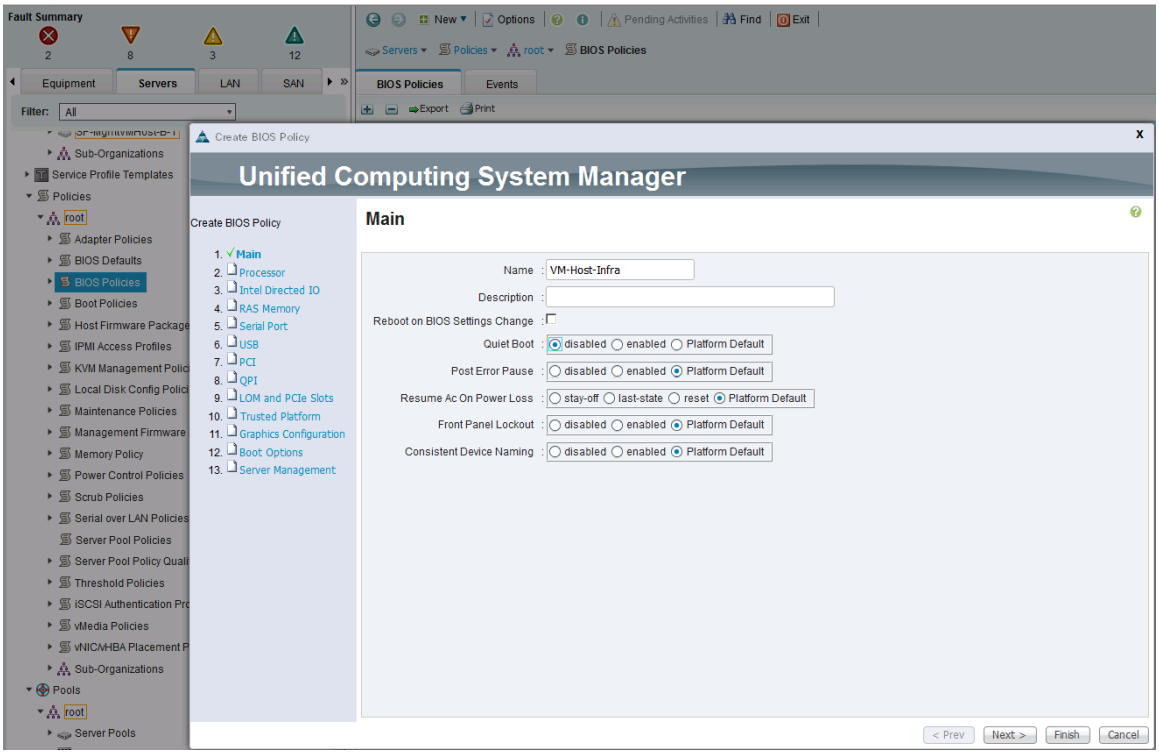

5. In the Processor screen, change the following:
    a. Turbo Boost to Enabled
    b. Enhanced Intel Speedstep to Enabled
    c. Hyper Threading to Enabled
    d. Core Multi Processing to All
    e. Execution Disabled Bit to Enabled
    f. Virtualization Technology (VT) to Enabled
    g. Direct Cache Access to Enabled
    h. CPU Performance to Enterprise

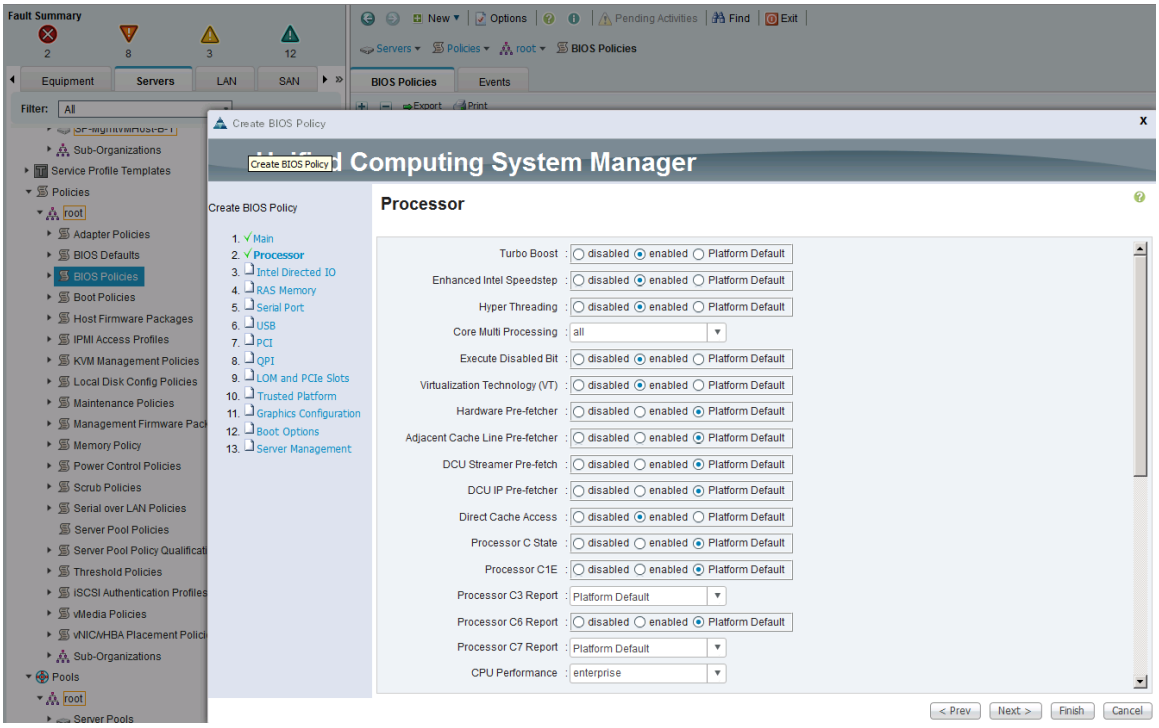

6. Click Next. In the Intel Directed IO screen, change the VT for Direct IO to Enabled.

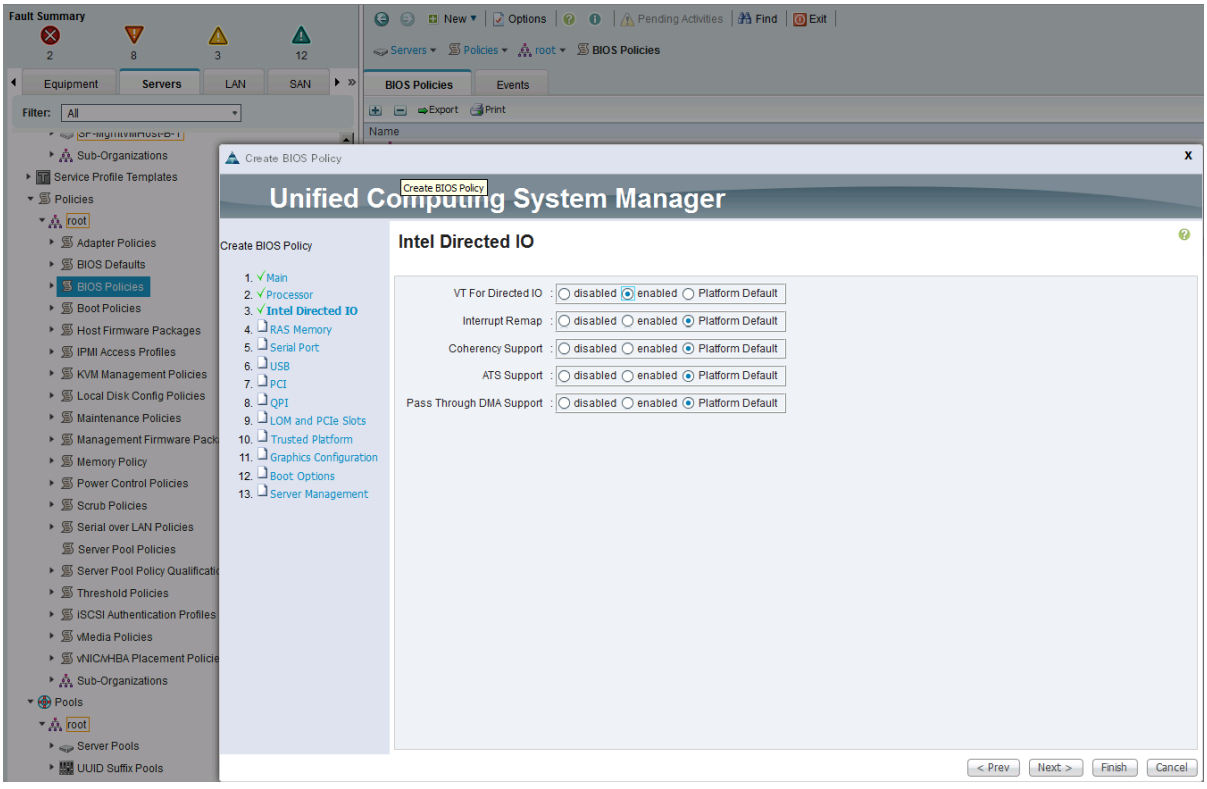

7. Click Next. In the RAS Memory screen, change the following:

    a. Memory RAS Config to maximum-performance

b. NUMA to Enabled

c. LV DDR Mode to performance-mode

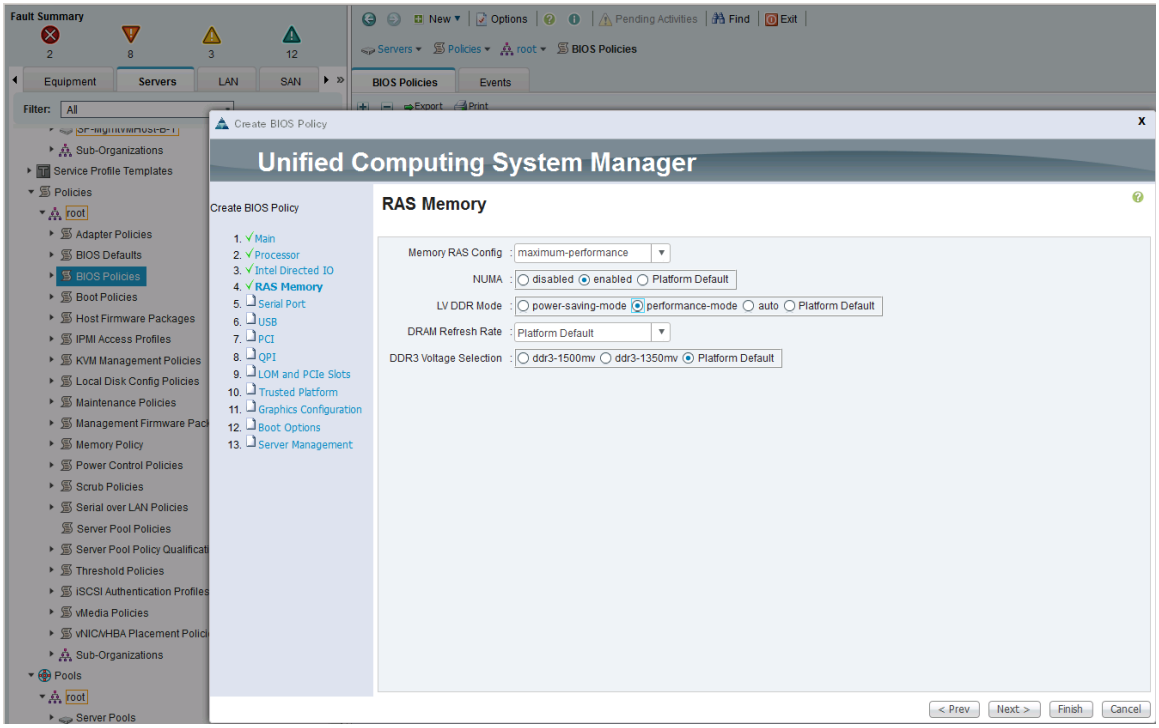

8. Click Finish and OK to create the BIOS policy.

Create Boot Policy

The procedure outlined in this section creates two boot policies, one that uses SAN Fabric A as the primary path for booting and a second one that uses SAN Fabric B as the primary path for booting. In the event of a failure in the primary path, each boot policy will use the other SAN Fabric as backup path for booting. Service Profile Templates using each boot policy are also created and distributed across hosts in the deployment. When multiple hosts are booting up, this will ensure that hosts use both SAN Fabrics for boot up.

This configuration assumes that two interface cards with 2 ports each are deployed on each controller.

Creating boot policies will require the WWPN info from the Nimble Storage array. To determine the target WWPN info from the Nimble array,

1. Login to web management interface of Nimble array.
2. Navigate to Administration > Network Configuration.
3. Click Active Settings link, followed by Interfaces tab.
4. Click Fibre Channel link to get a summary view of the WWPNs of all interfaces on the array. The WWPN of the Nimble array used in validation is used in the example below.

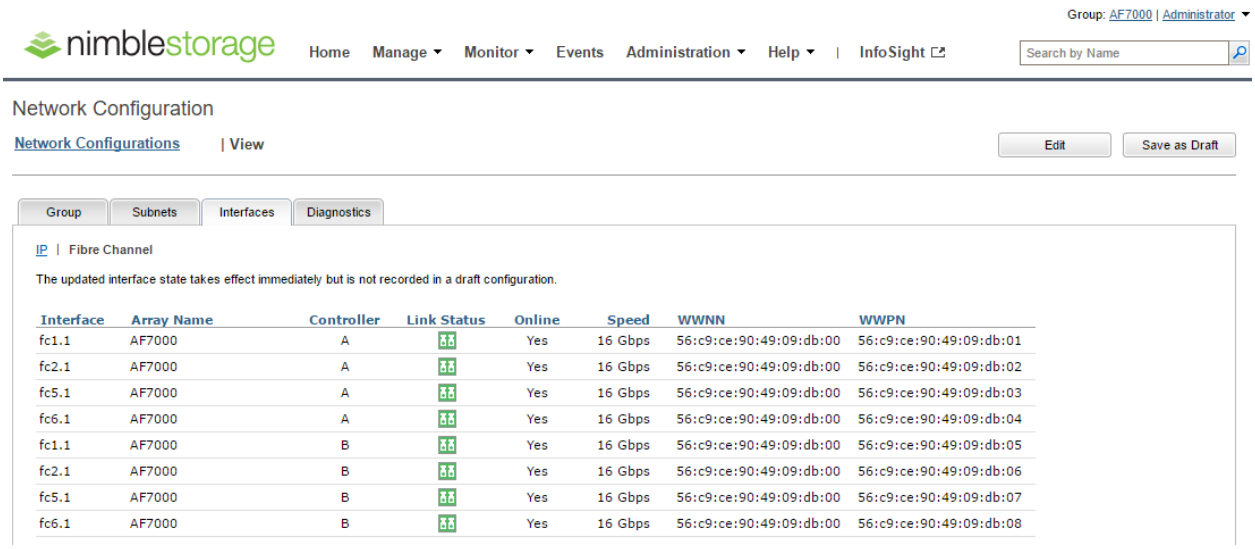

5. Above info can also be collected by ssh to the array and running “fc –list” command as shown below.

```
Nimble OS $ fc --list
Array: AF7000
------+-----+-------+------+----+-----------------------+-----------------------
Name   Ctrlr Admin   Fabric Link WWNN                    WWPN
             Status
------+-----+-------+------+----+-----------------------+-----------------------
fc1.1  A     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:01
fc2.1  A     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:02
fc5.1  A     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:03
fc6.1  A     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:04
fc1.1  B     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:05
fc2.1  B     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:06
fc5.1  B     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:07
fc6.1  B     online  yes    16G  56:c9:ce:90:49:09:db:00 56:c9:ce:90:49:09:db:08
```

Create the First Boot Policy through Fabric A

This boot policy will use CD/DVD, followed by both SAN Fabrics with the primary and secondary boot paths as shown below.

- Primary Boot Path: HBA-A → FI-A → SAN Fabric A (Cisco MDS-A)
- Secondary Boot Path: HBA-B → FI-B → SAN Fabric B (Cisco MDS-B)

To configure the first boot policy using SAN Fabric A as the primary boot path, complete the following steps.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root > Boot Policies.
3. Right-click and select Create Boot Policy.
4. In the Create Boot Policy window, enter the policy name (for example, `BOOT-AF-FI-A`).

Do not select the Reboot on Boot Order Change checkbox

5. Expand the Local Devices section of the window and select Add CD/DVD. At this point, the Local CD/DVD and Remote CD/DVD will get greyed out as shown below. Click OK to complete.

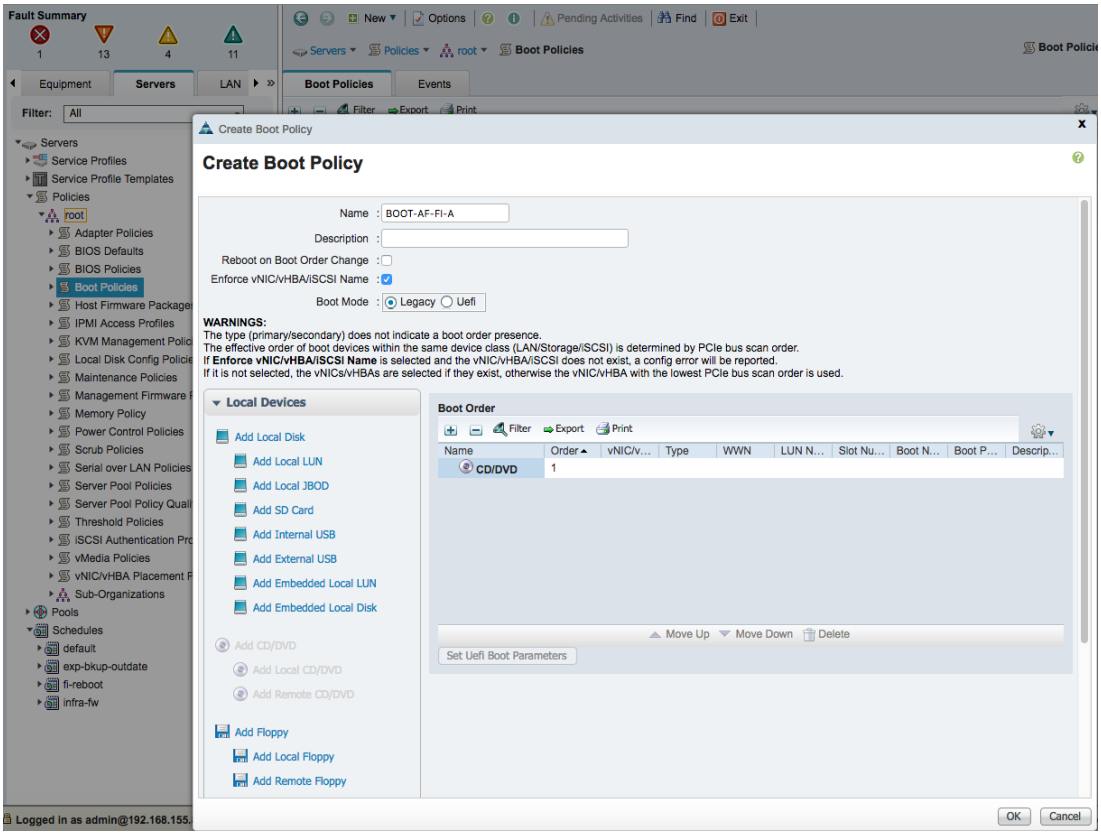

6. Collapse the Local Devices section and expand the vHBA section of the window. Click on Add SAN Boot and specify the vHBA (for example, `vHBA-A`) for SAN Fabric A as the primary boot path in the Add SAN Boot dialog box.

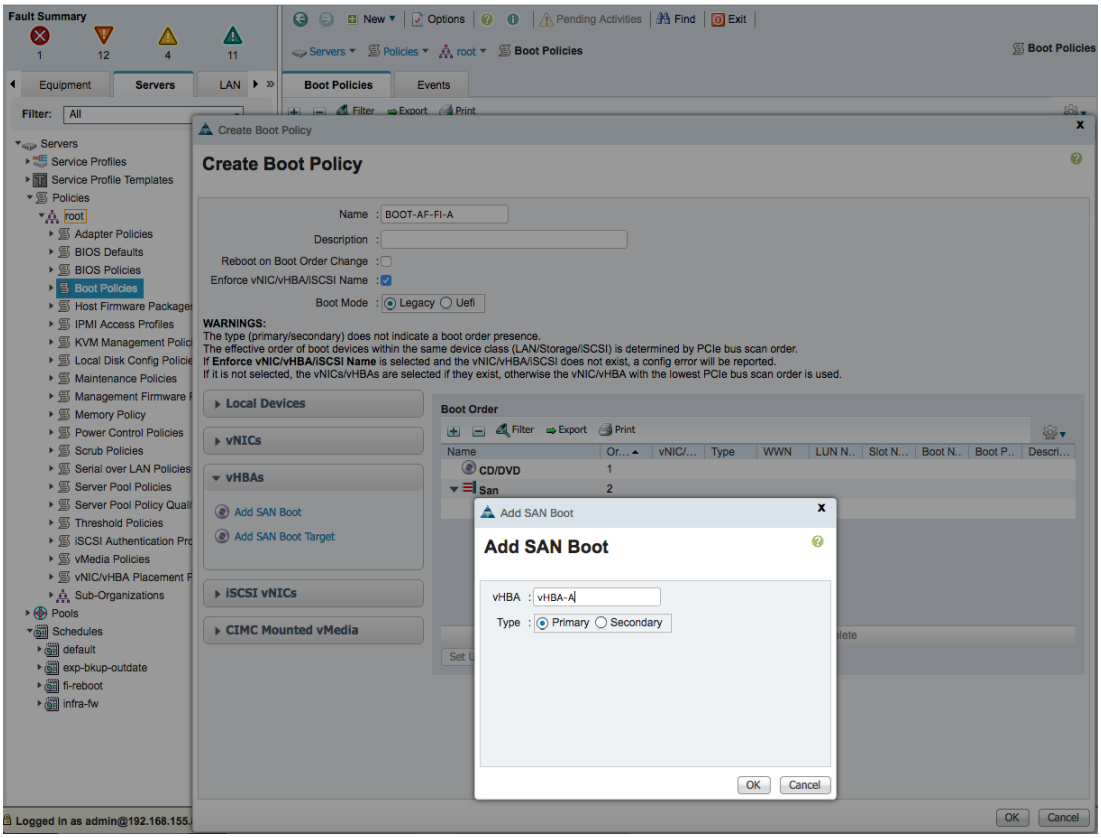

The Add SAN Boot Target under vHBAs will be greyed until this step is complete.

7. Click on Add SAN Boot Target under vHBAs and specify the primary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the primary controller reachable through the primary boot path. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:01`) of Nimble Controller A port (for example, `FC1.1`) is the SAN Boot target.

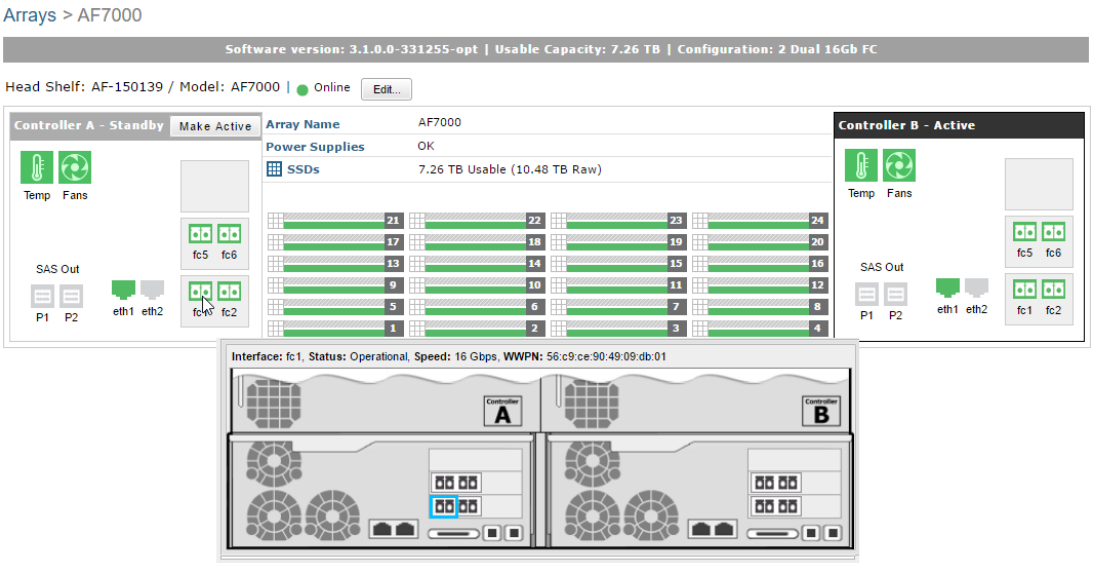

8. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

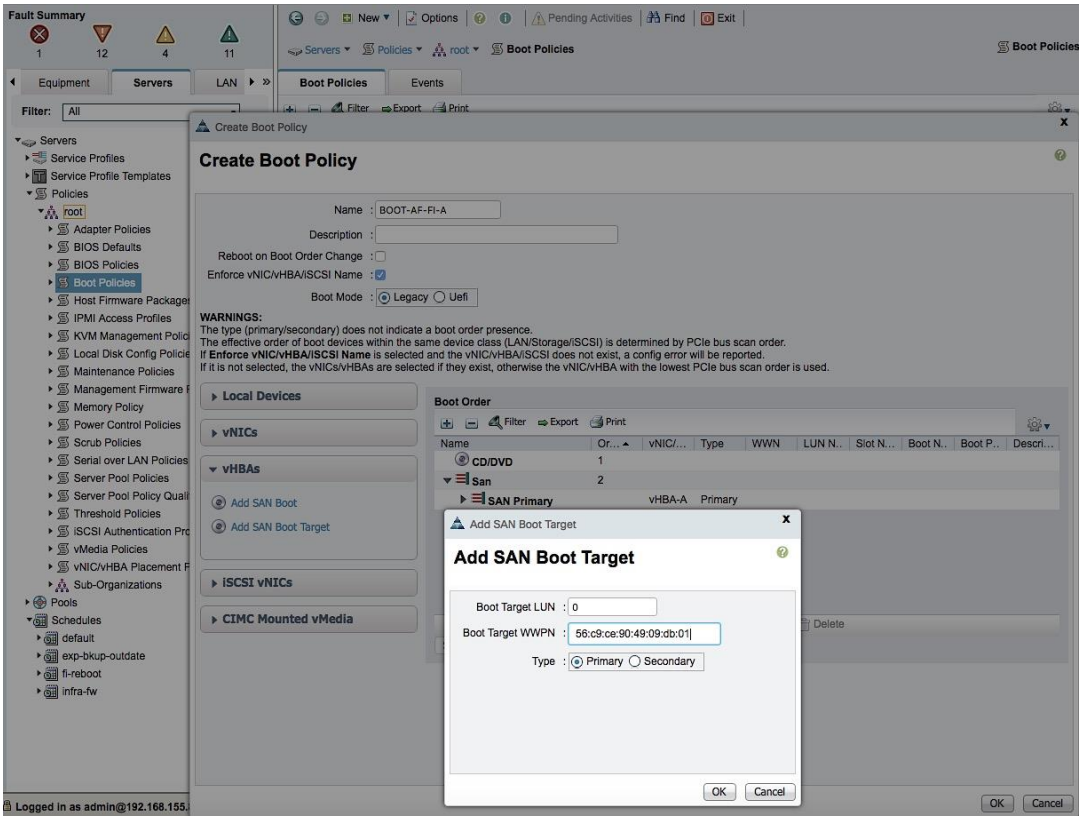

9. Click on Add SAN Boot Target under vHBAs a second time and specify the secondary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the secondary controller reachable through the primary boot path for this policy. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:05`) of Nimble Controller B port (for example, `FC1.1`) is the SAN Boot target.

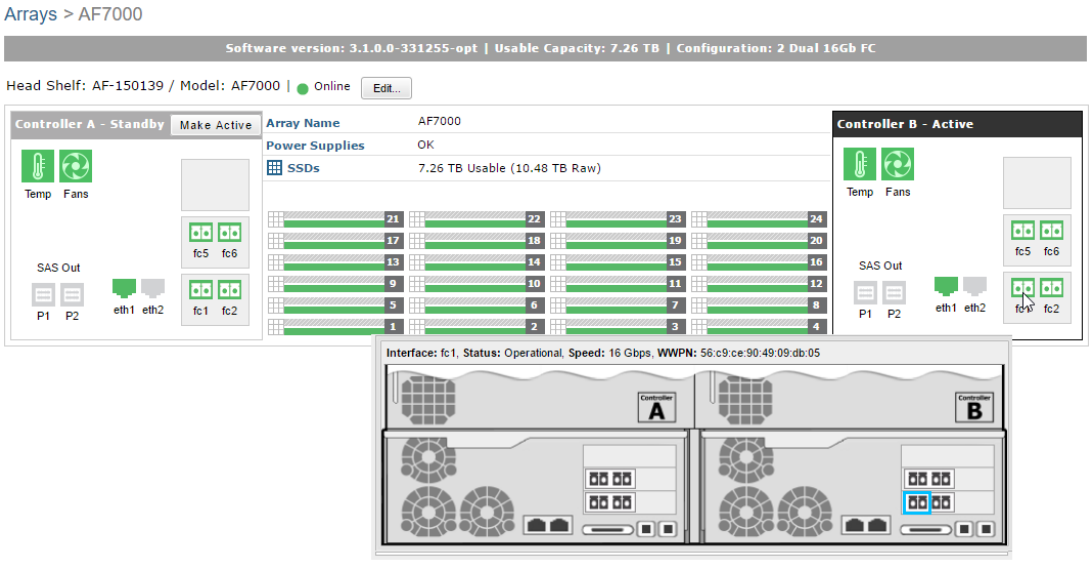

10. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

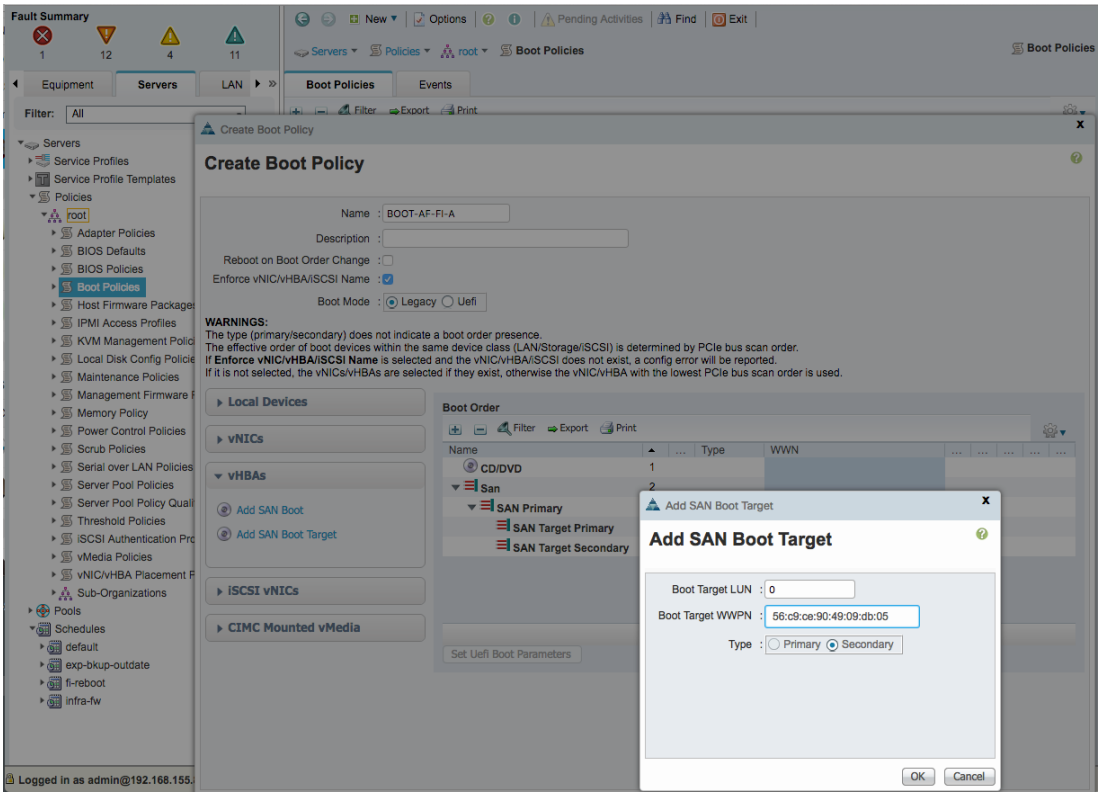

11. Repeat Step 6 and click on Add SAN Boot to specify the vHBA (for example, `vHBA-B`) for the secondary boot path in the Add SAN Boot dialog box.

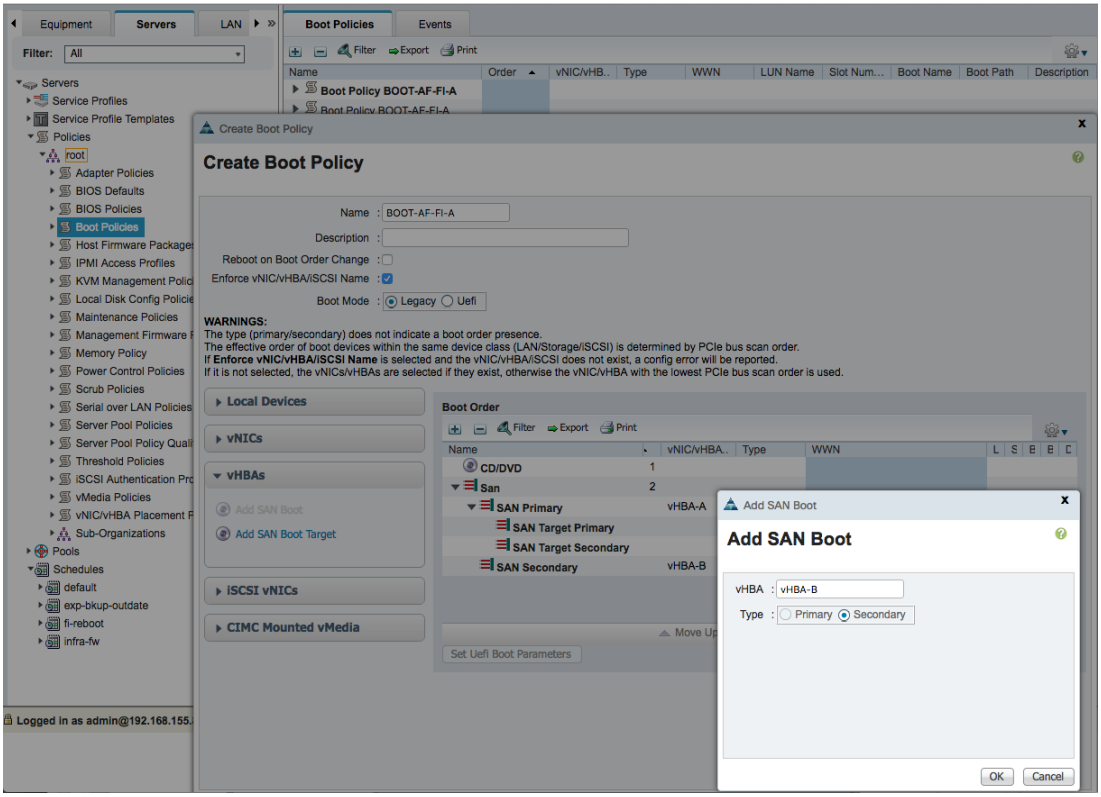

12. Repeat Step 7 and click on Add SAN Boot Target under vHBAs to specify the primary SAN boot target reachable through the secondary boot path in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the primary controller reachable through the secondary boot path. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:02`) of Nimble Controller A port (for example, `FC2.1`) is the SAN Boot target.

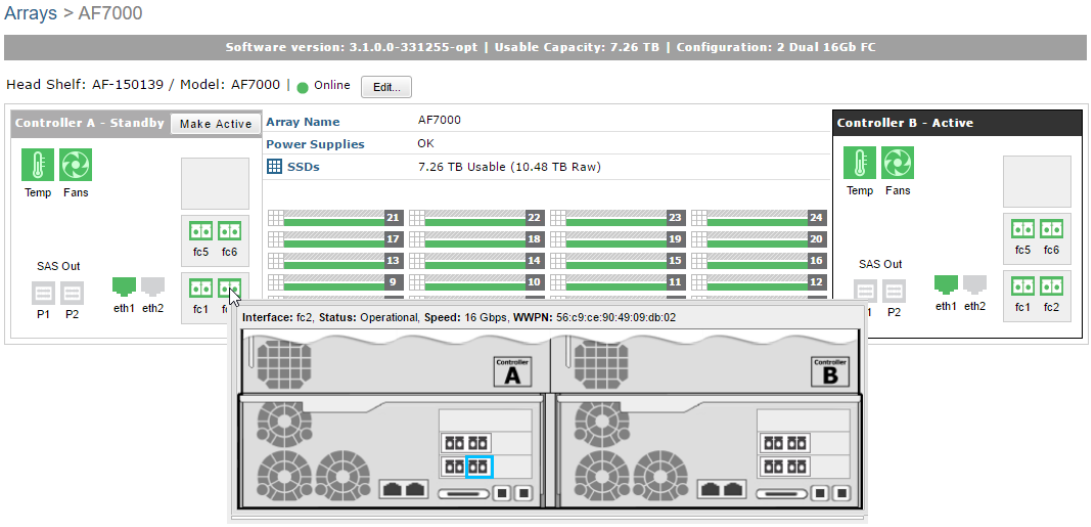

13. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

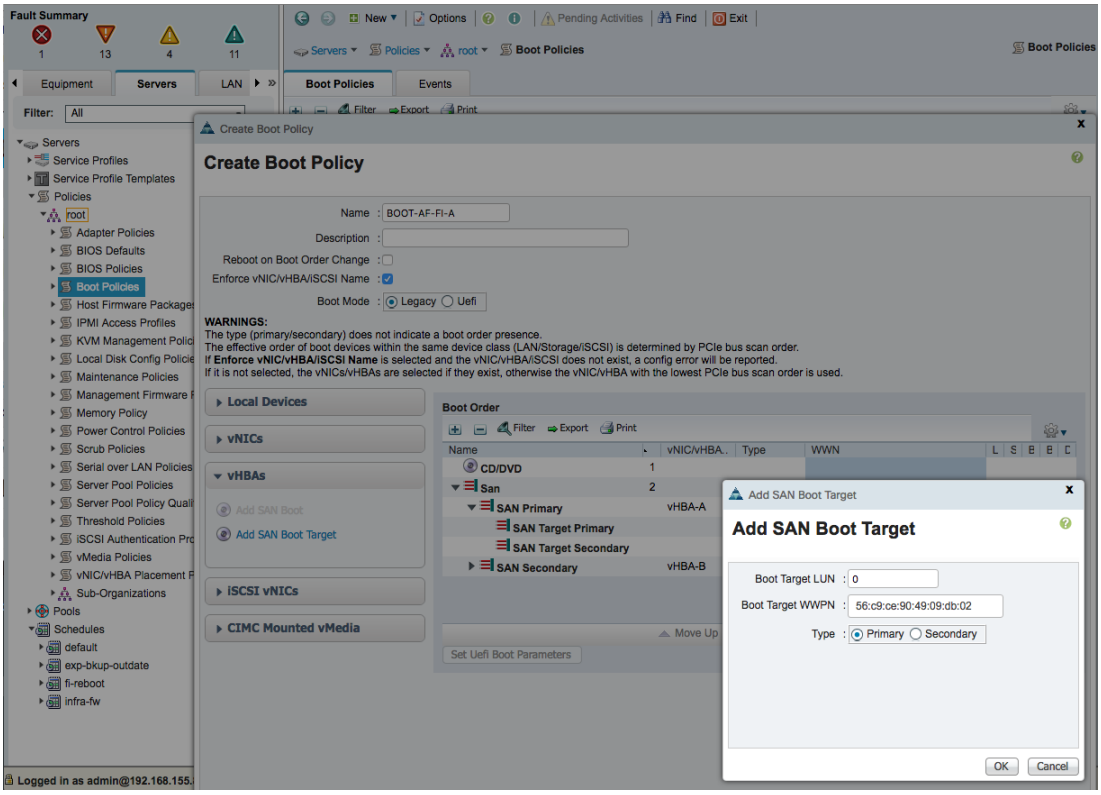

14. Repeat Step 9 and click on Add SAN Boot Target under vHBAs a second time to specify the secondary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the secondary controller reachable through the secondary boot path. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:06`) of Nimble Controller B port (for example, `FC2.1`) is the SAN Boot target.

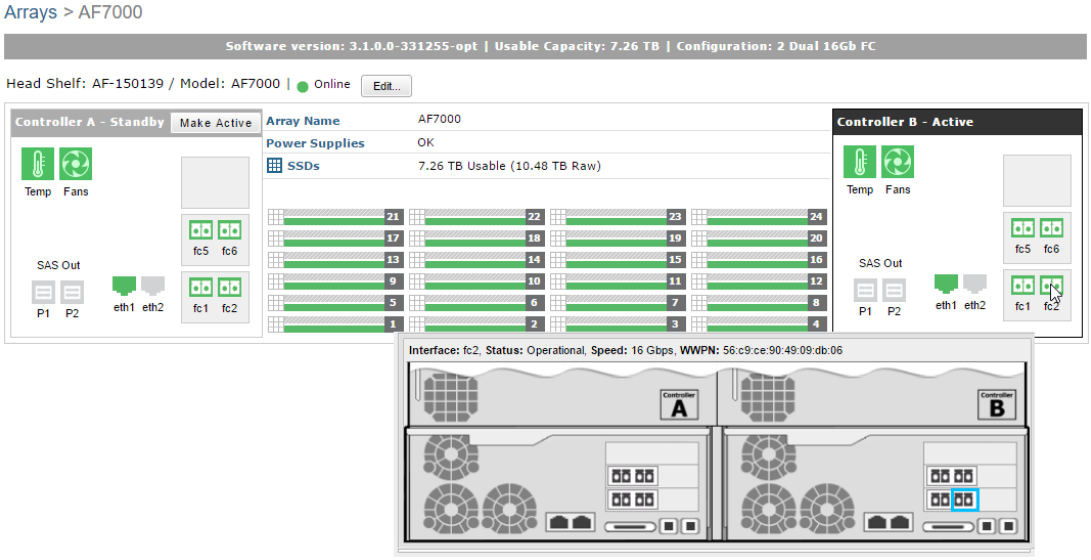

15. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

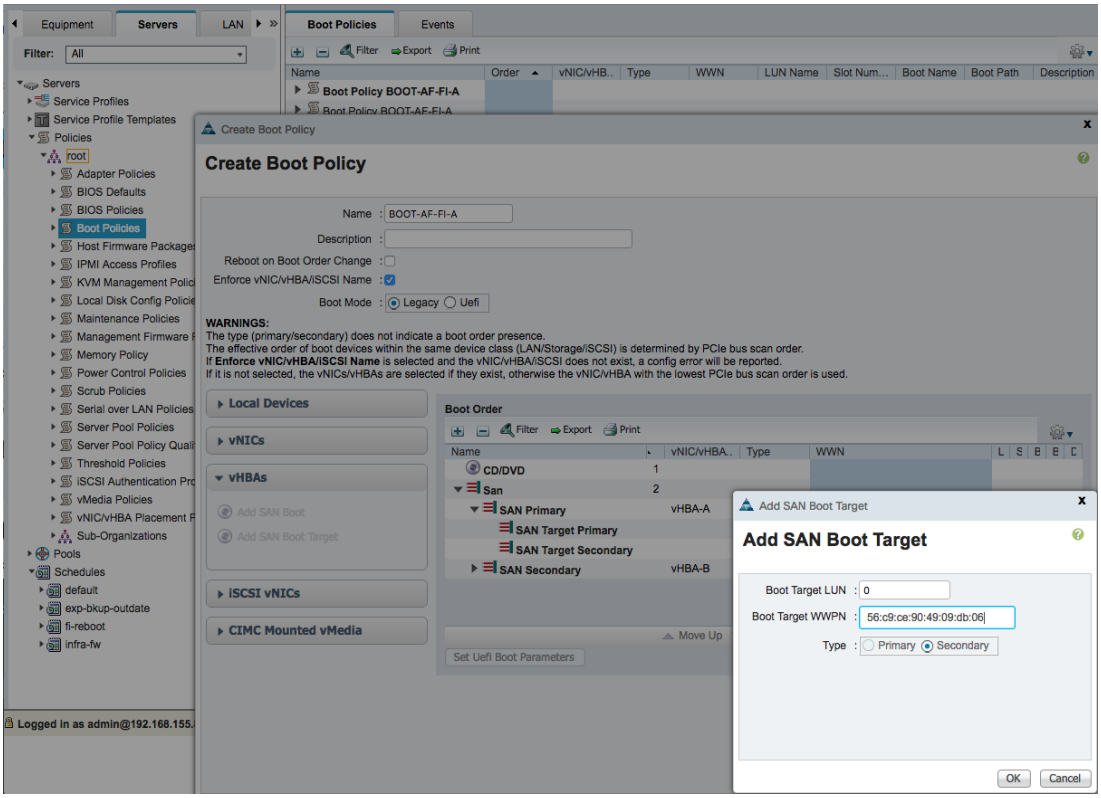

16. A summary of the first boot policy configuration is shown below.

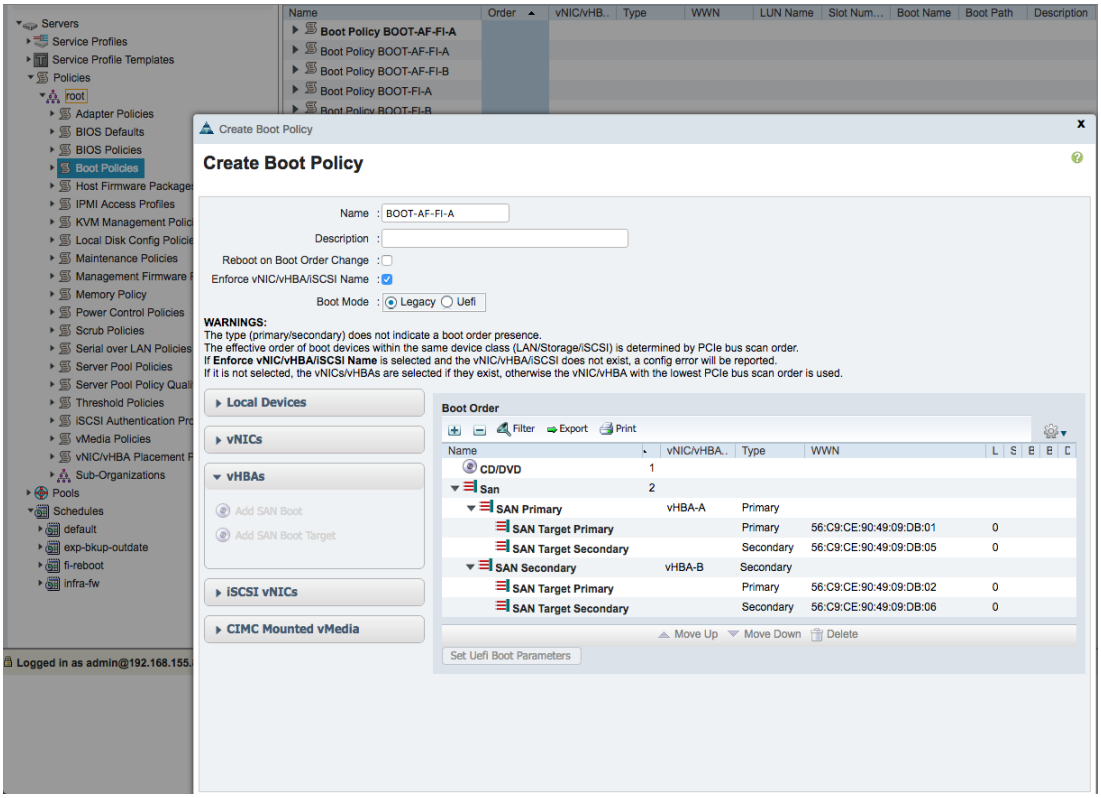

Create the Second Boot Policy through Fabric B

This boot policy will use CD/DVD, followed by both SAN Fabrics using the primary and secondary boot paths as shown below.

- Primary Boot Path: HBA-B > FI-B > SAN Fabric B (Cisco MDS-B)
- Secondary Boot Path: HBA-A > FI-A > SAN Fabric A (Cisco MDS-A)

To create the second boot policy with SAN Fabric B as the primary boot path, complete the following steps.

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > Boot Policies.
3. Right-click and select Create Boot Policy.
4. In the Create Boot Policy window, enter the policy name (for example, `BOOT-AF-FI-B`).

Do not select the Reboot on Boot Order Change checkbox

5. Expand the Local Devices section of the window and select Add CD/DVD. At this point, the Local CD/DVD and Remote CD/DVD will get greyed out. Click OK to complete.

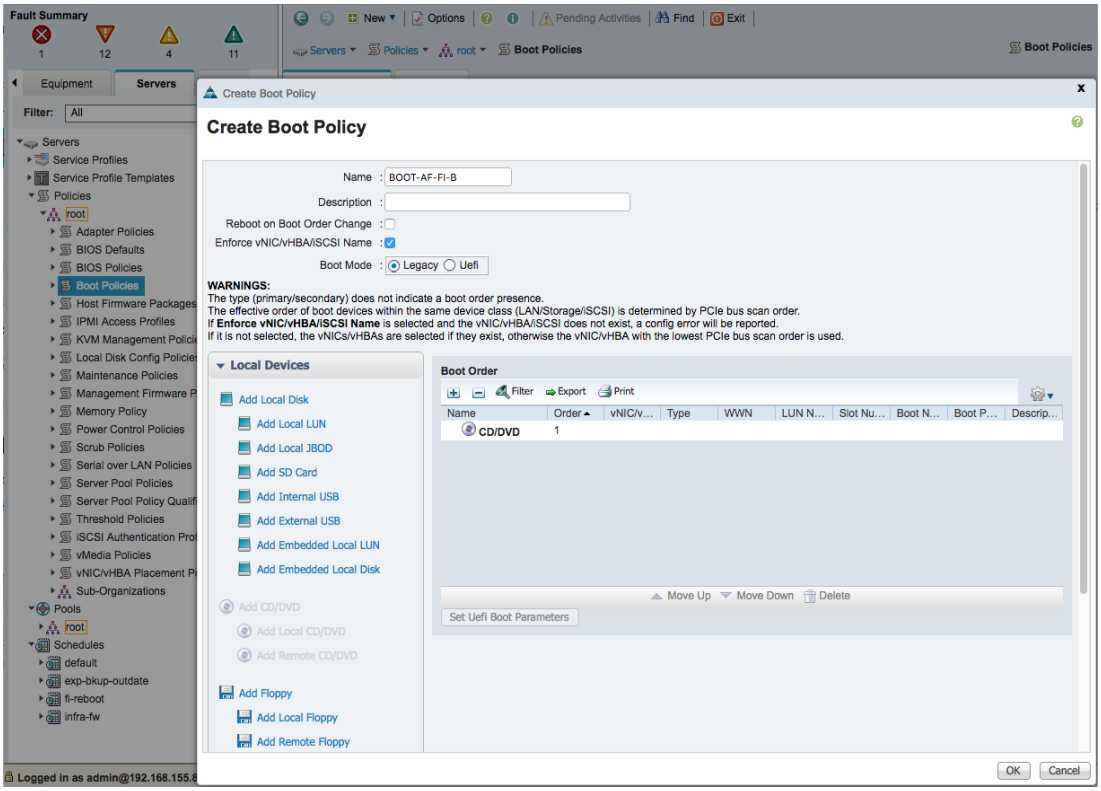

6. Collapse the Local Devices section and expand the vHBA section of the window. Click Add SAN Boot and specify the vHBA (for example, `vHBA-B`) for SAN Fabric B as the primary boot path in the Add SAN Boot dialog box.

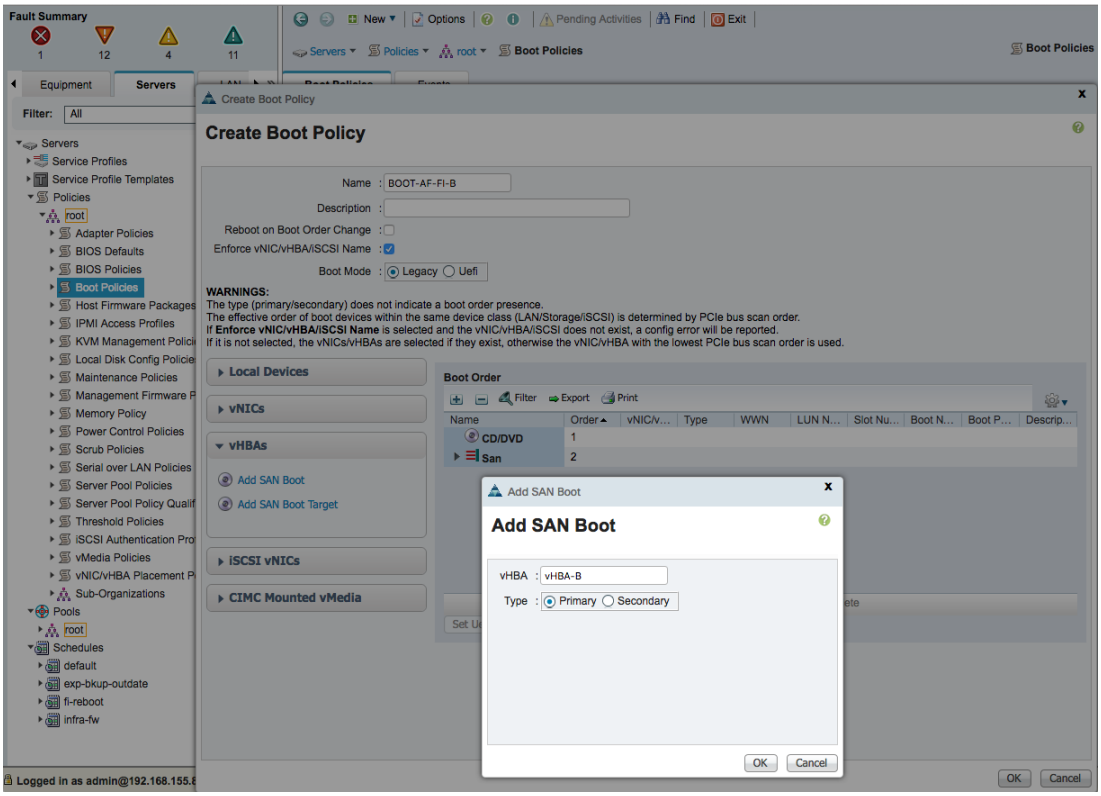

The Add SAN Boot Target will be greyed until this step is complete.

7. Click Add SAN Boot Target under vHBAs and specify the primary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the primary controller reachable through the primary boot path for this policy. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:02`) of Nimble Controller A port (for example, `FC2.1`) is the SAN Boot target.

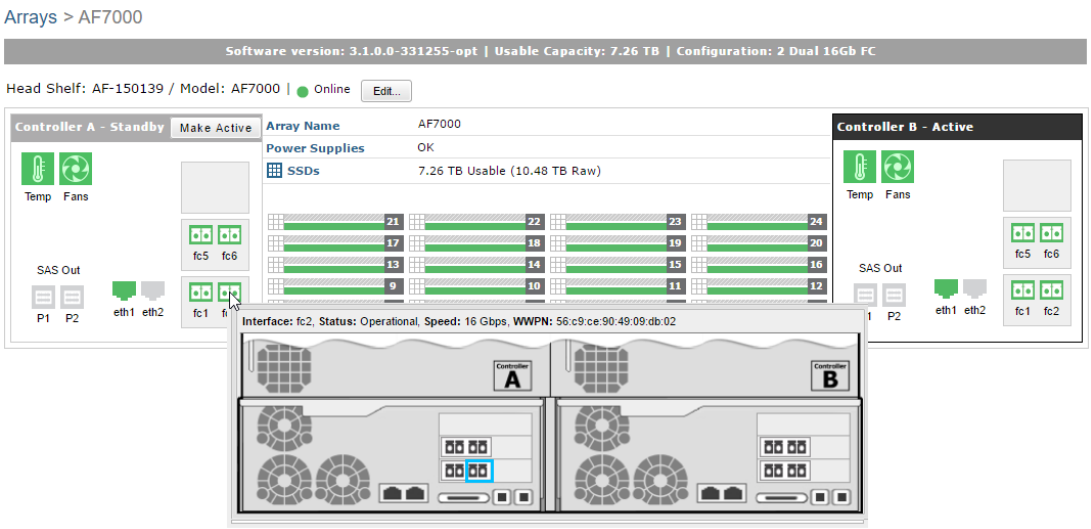

8. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

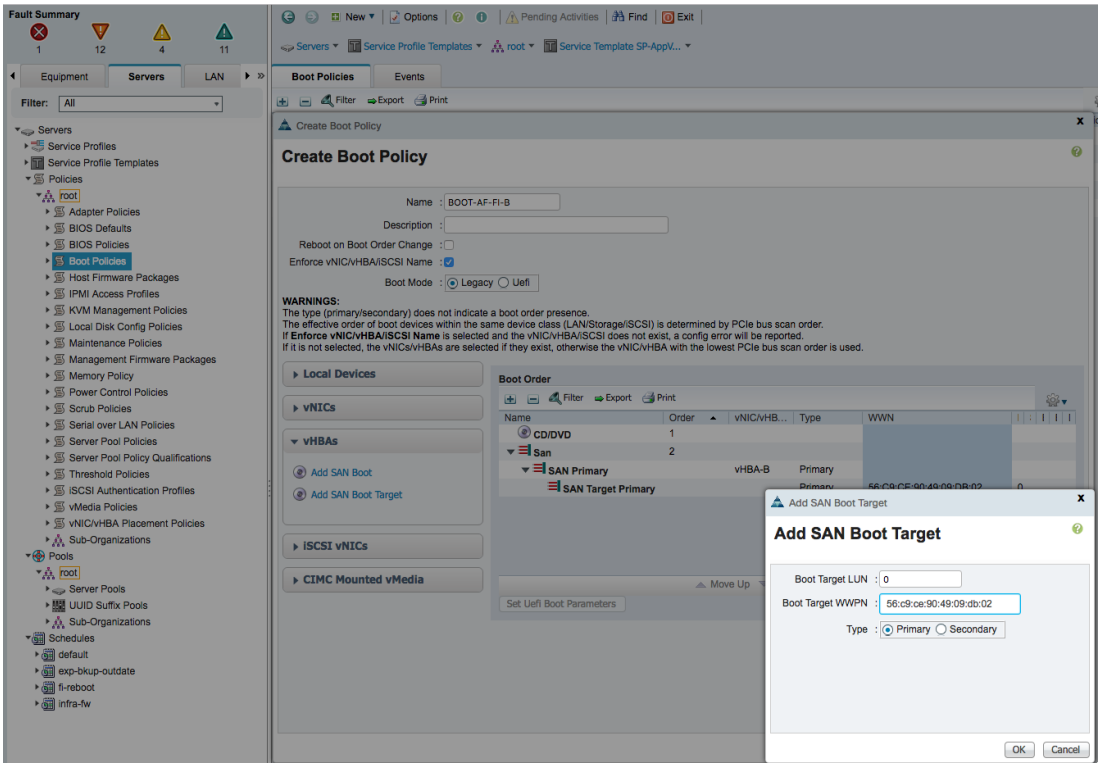

9. Click Add SAN Boot Target under vHBAs a second time and specify the secondary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the secondary controller reachable through the primary boot path for this policy. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:06`) of Nimble Controller B port (for example, `FC2.1`) is the SAN Boot target.

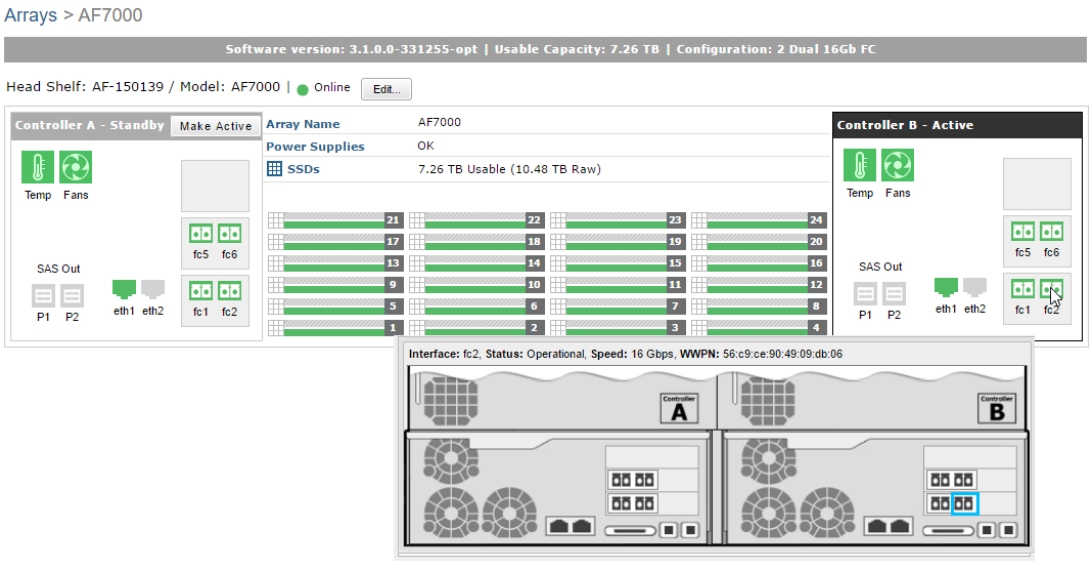

10. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

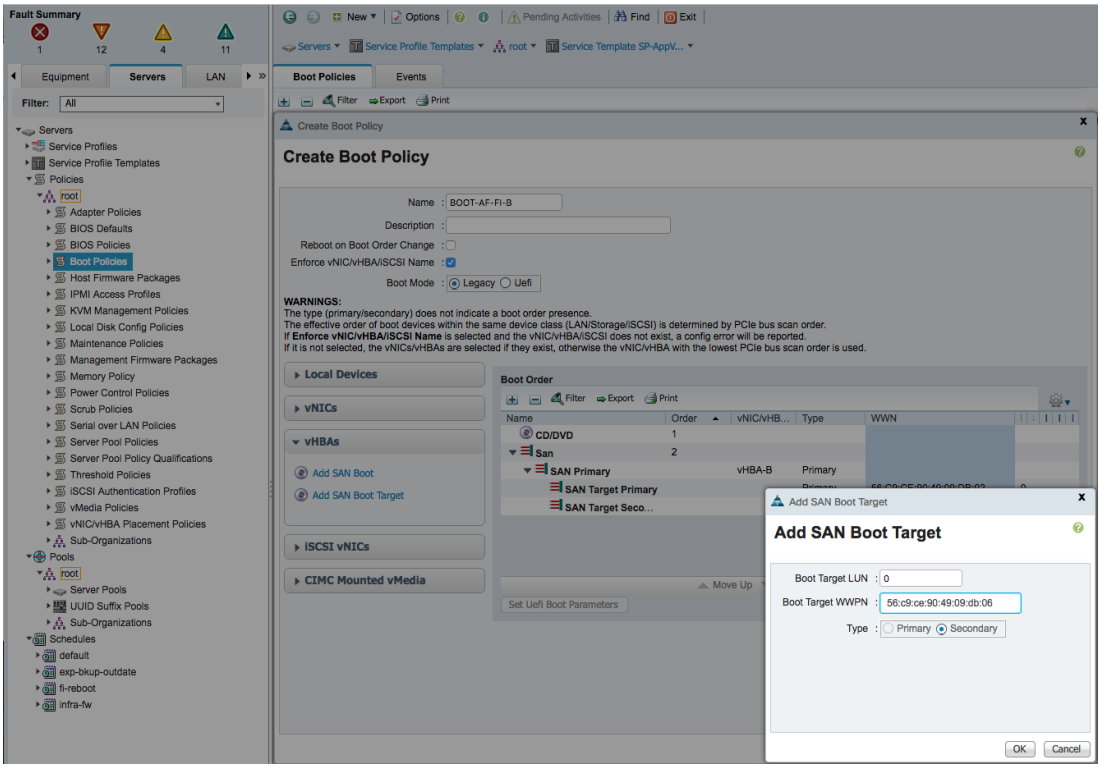

11. Repeat Step 6 and click Add SAN Boot to specify the vHBA (for example, `vHBA-A`) as the secondary boot path (using SAN Fabric A) for this policy in the Add SAN Boot dialog box.

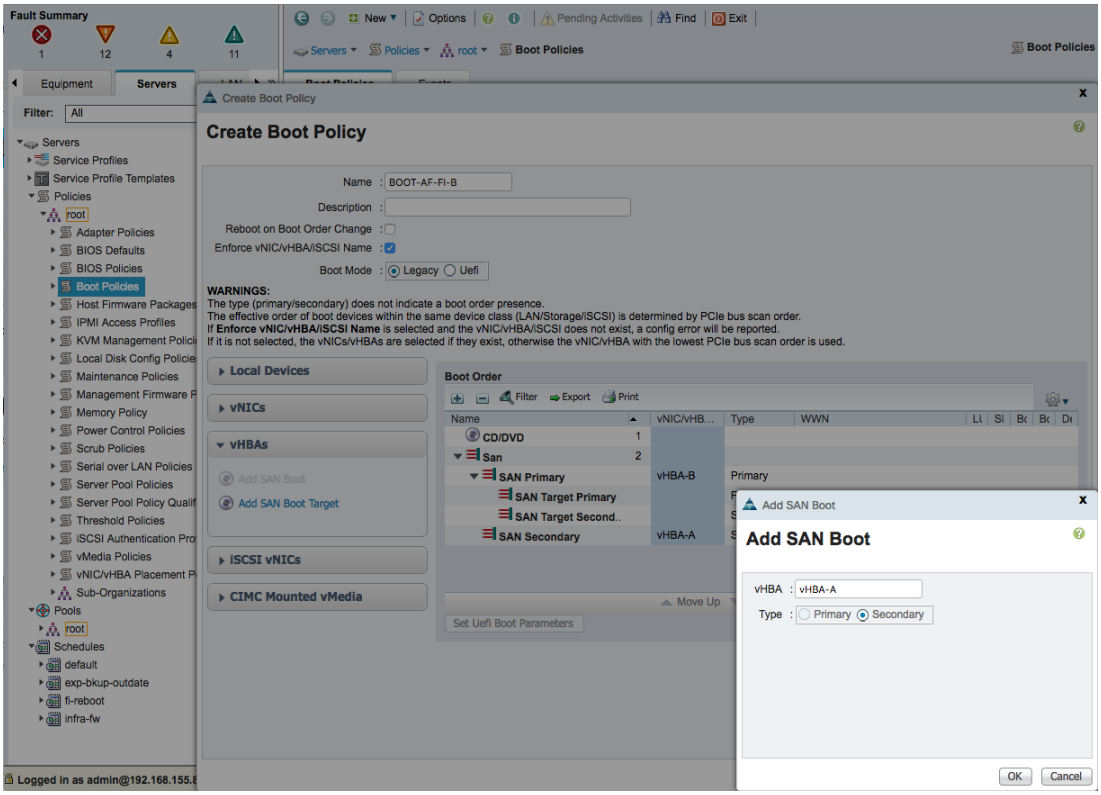

12. Repeat Step 7 and click on Add SAN Boot Target under vHBAs to specify the primary SAN boot target reachable through the secondary boot path in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the primary controller reachable through the secondary boot path for this policy. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:01`) of Nimble Controller A port (for example, `FC1.1`) is the SAN Boot target.

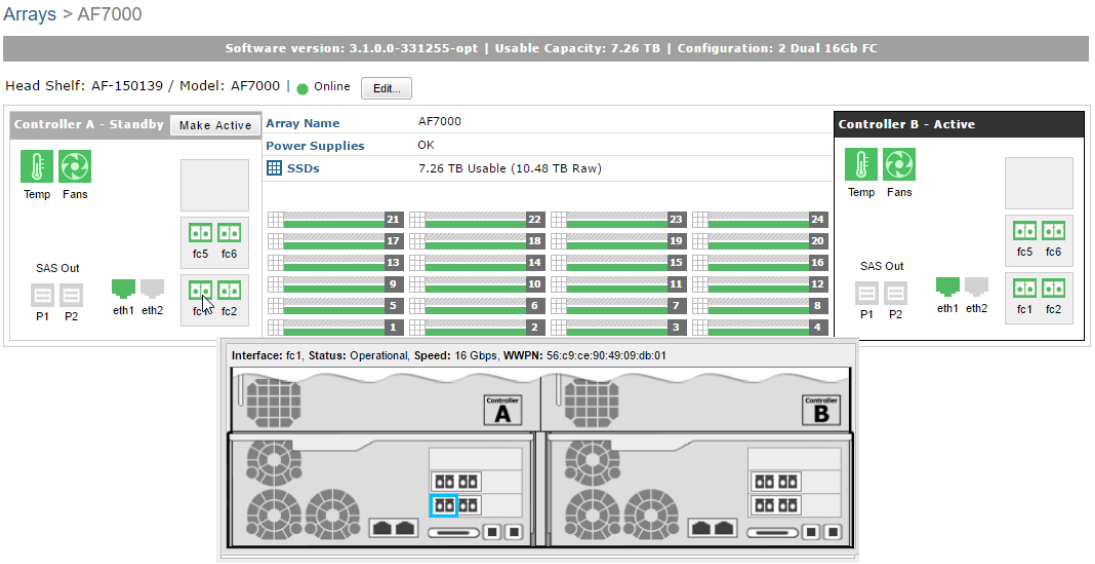

13. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

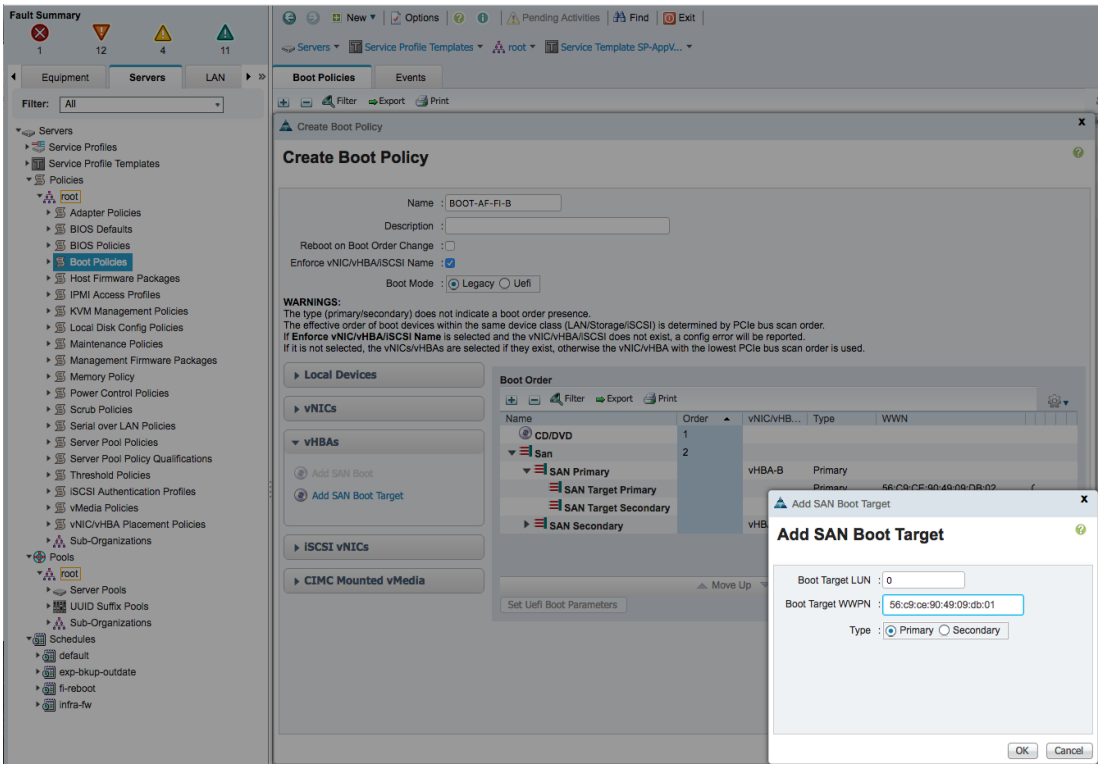

14. Repeat Step 9 and click Add SAN Boot Target under vHBAs a second time to specify the secondary SAN boot target in the Add SAN Boot Target dialog box. The SAN Boot Target is the WWPN of the secondary controller reachable through the secondary boot path for this policy. In this example, the WWPN (for example, `56:c9:ce:90:49:09:db:05`) of Nimble Controller B port (for example, `FC1.1`) is the SAN Boot target.

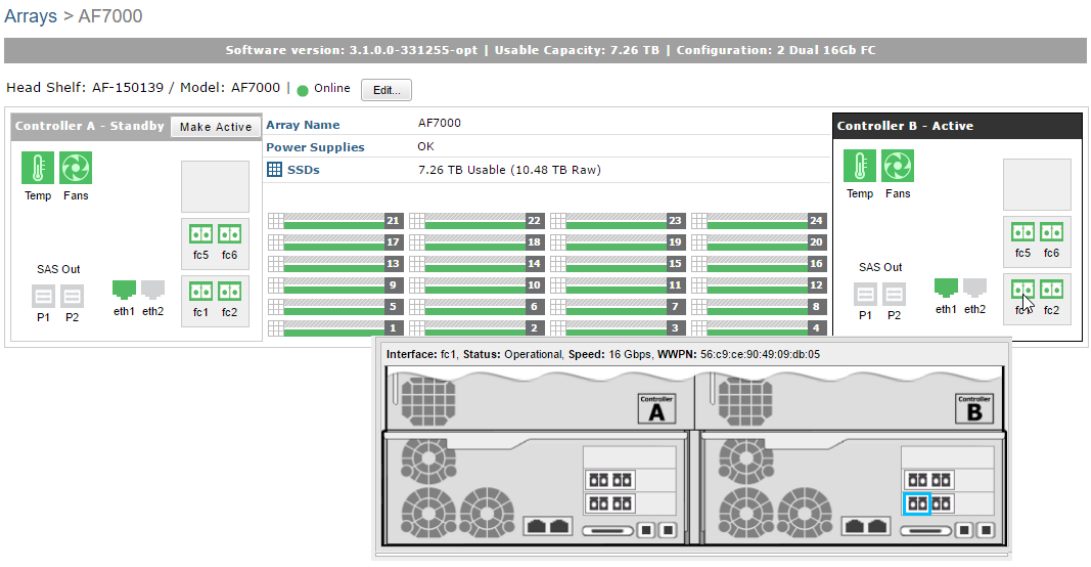

15. Specify Boot Target LUN as ‘0’. Enter the Boot Target WWPN from the previous step. Click OK to complete.

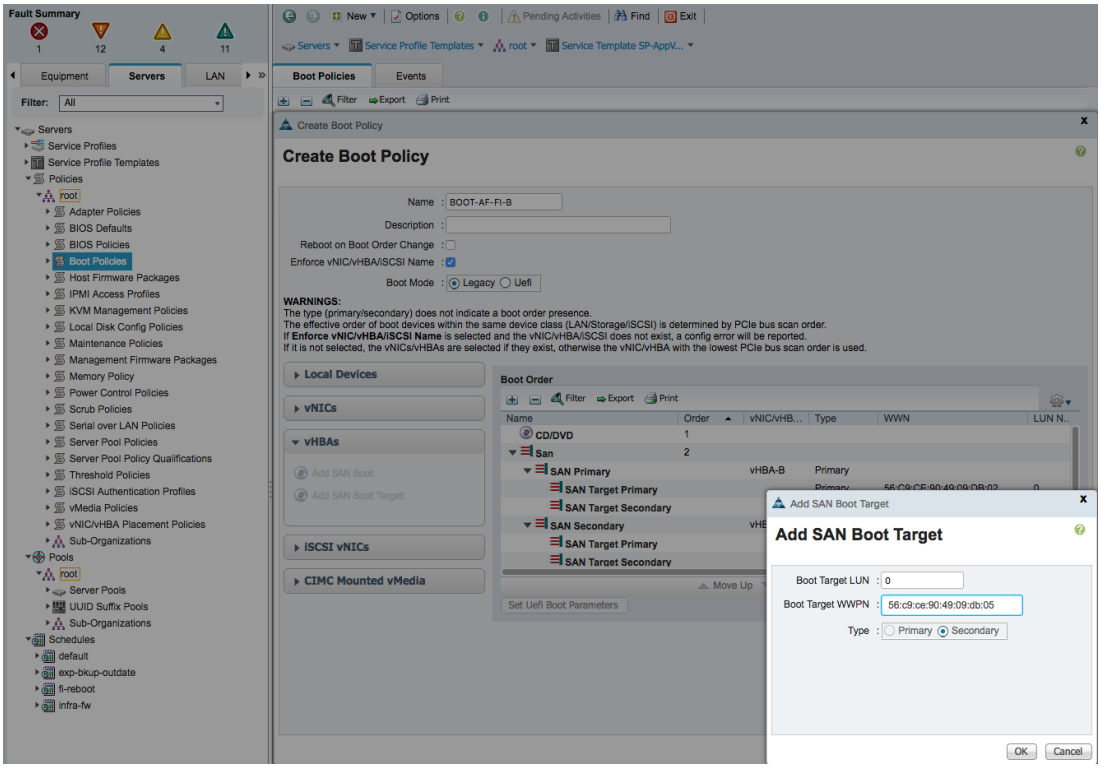

16. A complete summary of the second boot policy configuration is shown below.

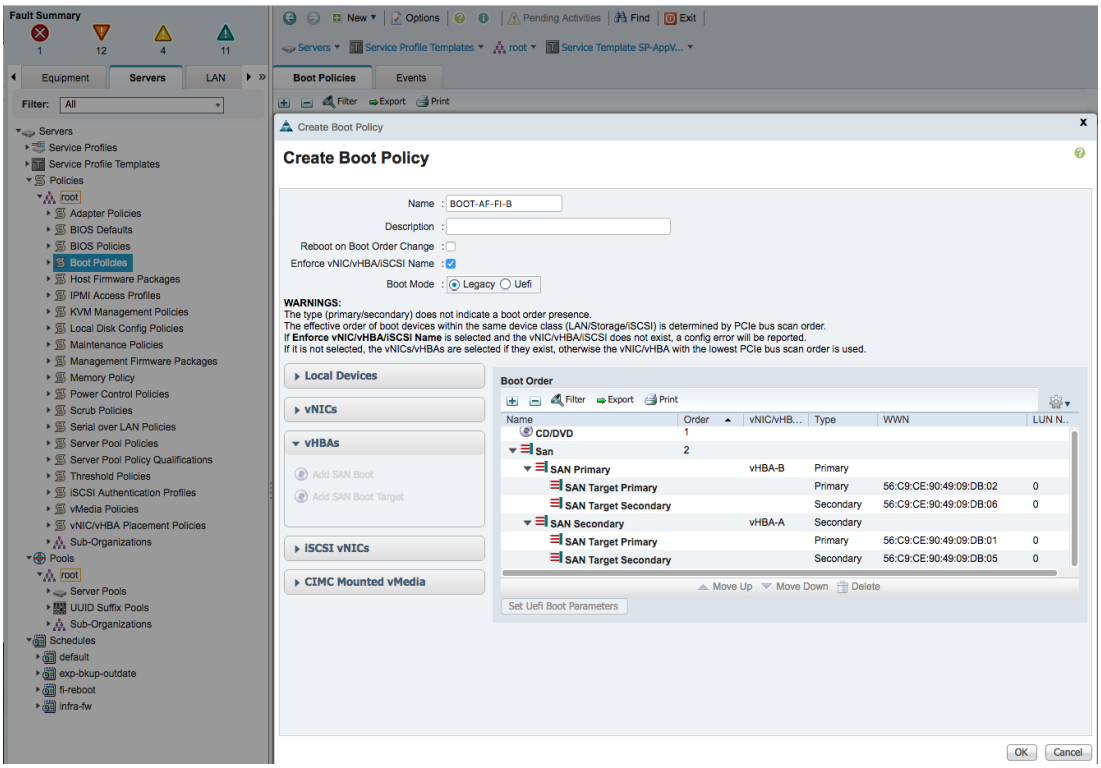

Boot Policies – Summary View

A complete summary of created boot polices are shown below.

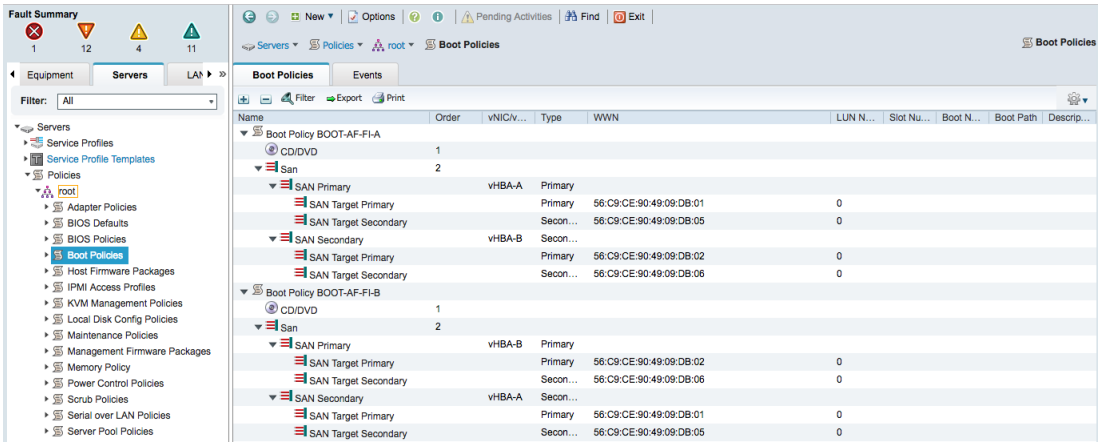

Create Host Firmware Package Policy

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties. To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > Host Firmware Packages.
3. Right-click Host Firmware Packages and select Create Host Firmware Package.
4. Enter the name of the host firmware package (for example, `HOST-FW-POLICY`).
5. Leave Simple selected.
6. Select the package versions for the different type of servers (Blade, Rack) in the deployment (for example, 3.1(1e) for Blade and Rack servers.
7. Click OK twice to create the host firmware package.

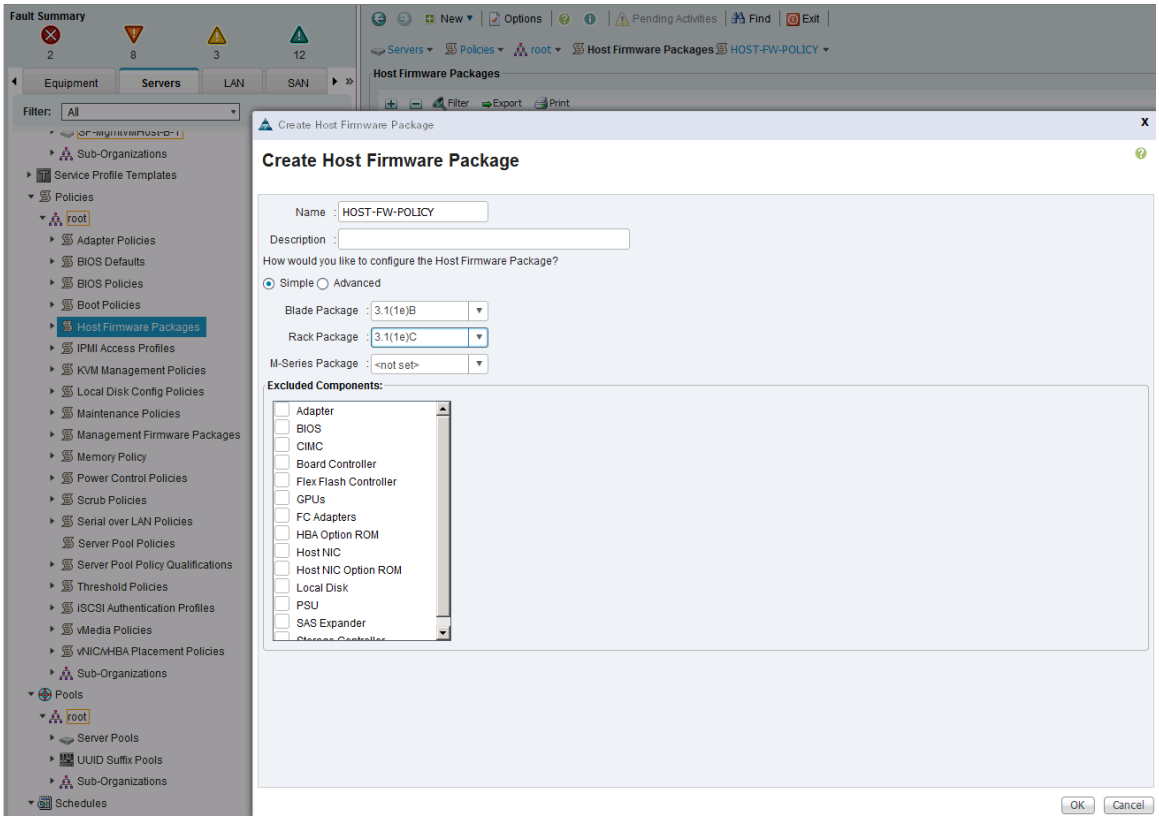

Create Local Disk Configuration Policy

A local disk configuration policy is necessary if the servers do not have a local disk. To create a local disk configuration policy (optional) for Cisco UCS hosts, complete the following steps:

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > Local Disk Configuration Policies.
3. Right-click and select Create Local Disk Configuration Policy.
4. Enter local disk configuration policy name (for example, `SAN-Boot`).
5. Click OK twice to create the local disk configuration policy.

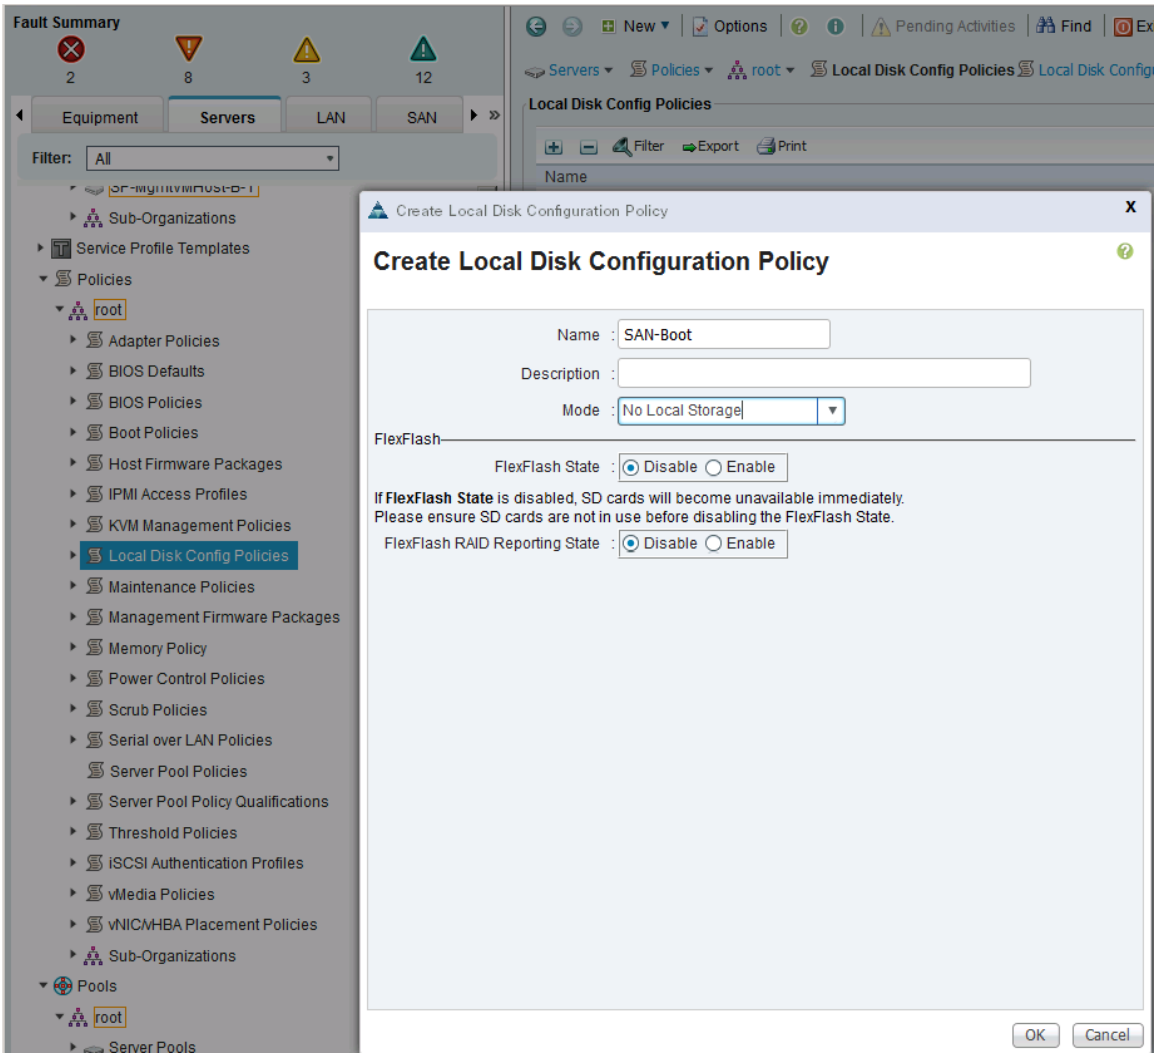

Create Maintenance Policy

To create a Maintenance Policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > Maintenance Policies.
3. Right-click and select Create Maintenance Policy.
4. Specify a name for the policy (for example, `User-ACK`).
5. Change the Reboot Policy to User Ack.
6. Click OK twice to create the maintenance policy.

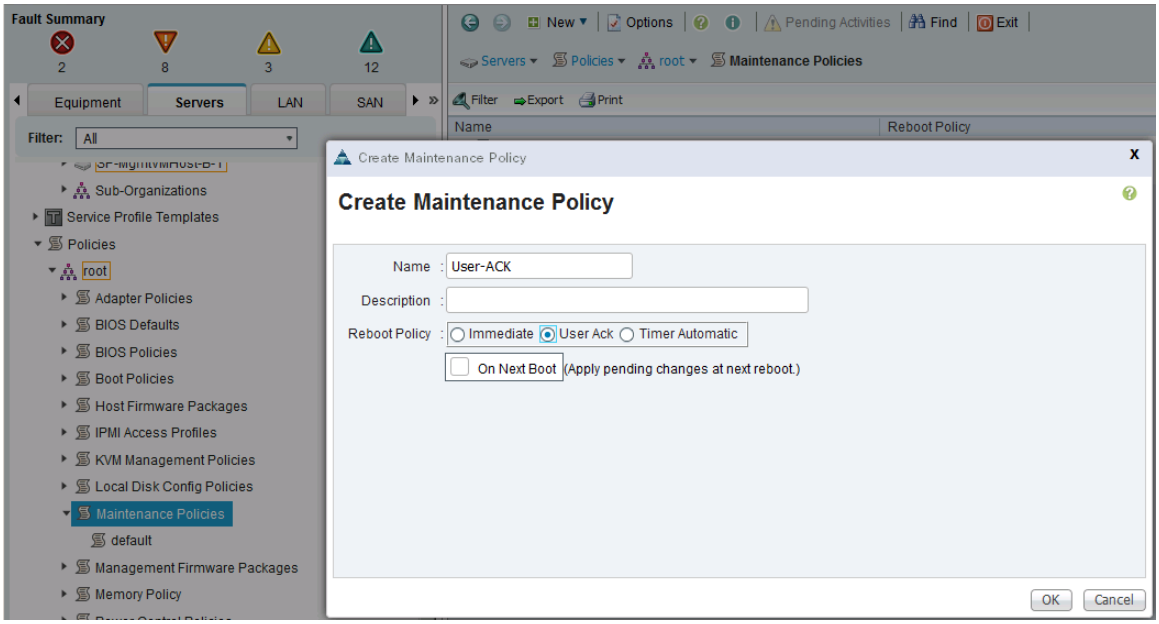

Create Power Control Policy

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

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Servers > Policies > root > Power Control Policies.
3. Right-click and select Create Power Control Policy.
4. Enter the power control policy name (for example, `No-Power-Cap`).
5. Change the power capping setting to No Cap.
6. Click OK twice to create the power control policy.

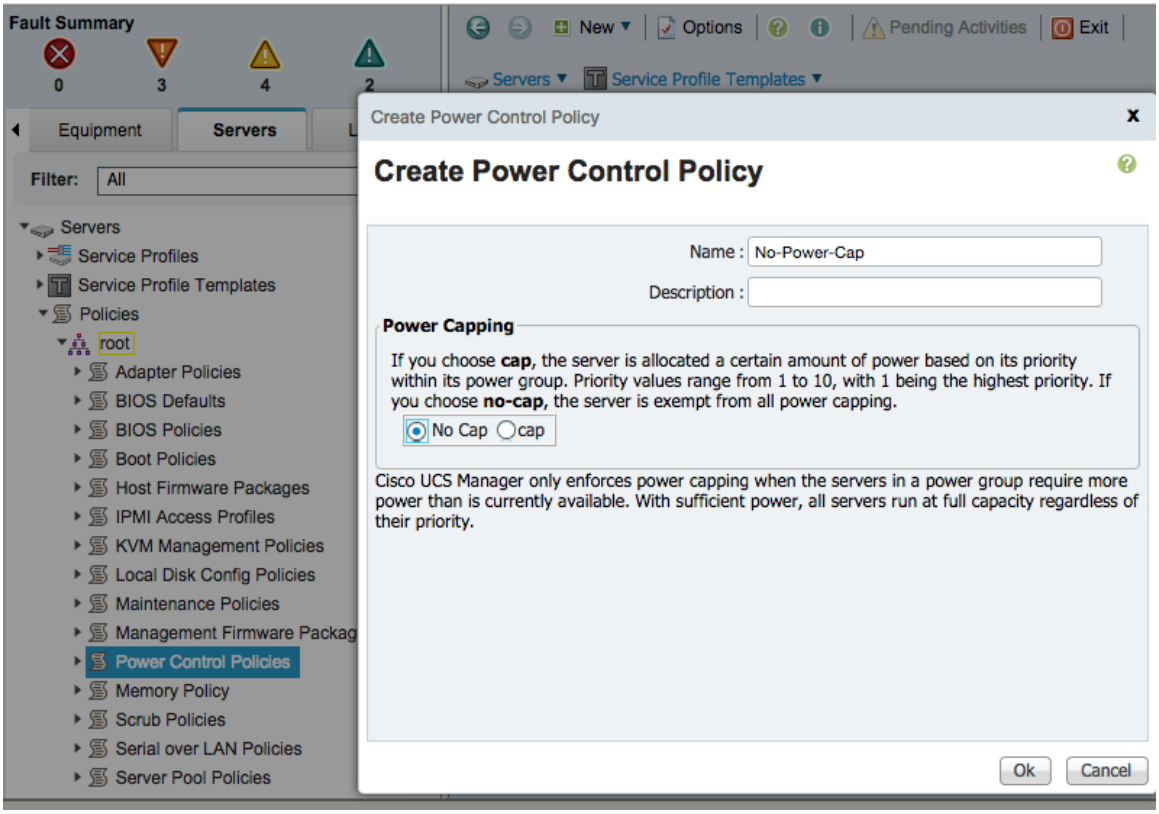

Create Server Pool Qualification Policy

To create a server pool qualification policy (optional) for the Cisco UCS environment, complete the following steps:

This example creates a policy for a Cisco UCS B200-M4 server.

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > Server Pool Policy Qualifications.
3. Right-click and select Create Server Pool Policy Qualification.

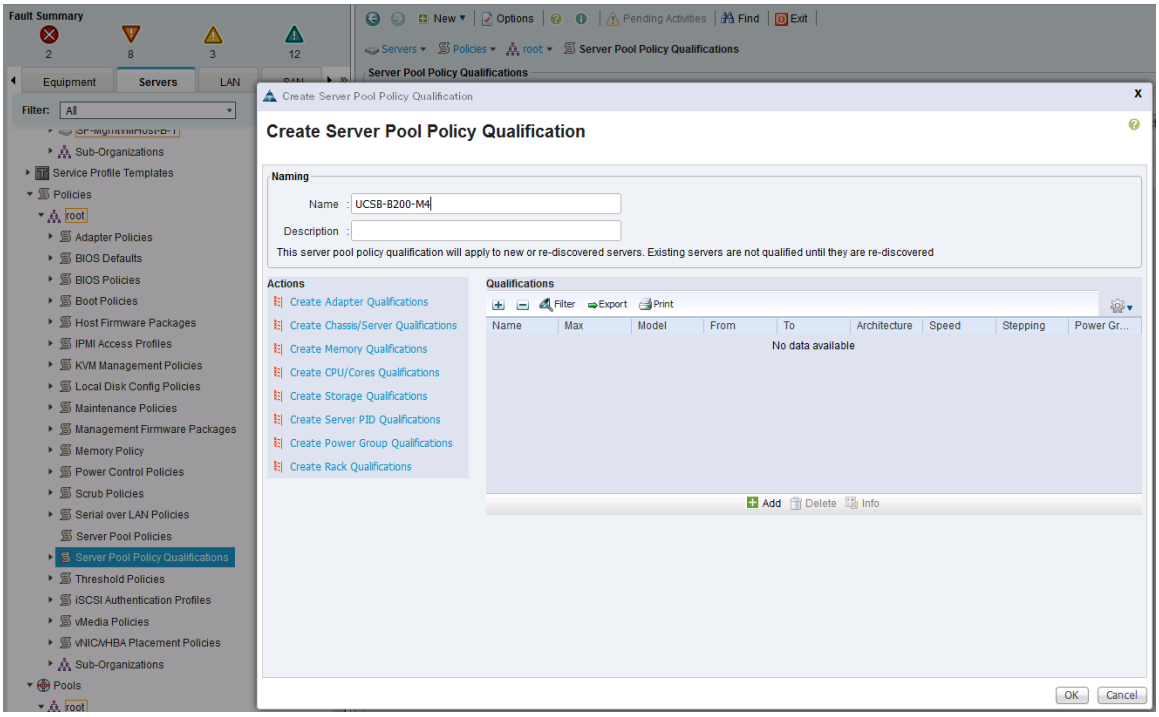

4. Enter the name for the policy (for example, `Cisco UCSB-B200-M4`).
5. Click + and select Create Server PID Qualifications.
6. Enter the PID for the server from the drop-down list (for example, `Cisco UCSB-B200-M4`).
7. Click OK twice to create the server pool qualification policy.

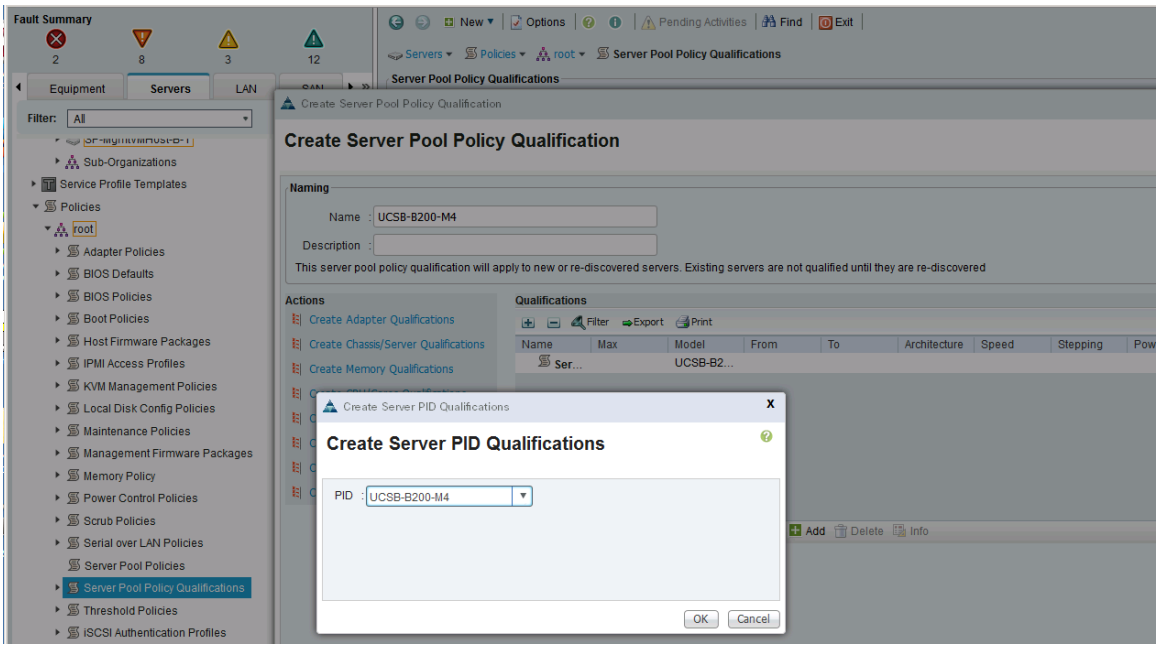

Create vNIC and vHBA Placement Policy

To create a vNIC/vHBA placement policy for Cisco UCS hosts, complete the following steps.

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Policies > root > vNIC/vHBA Placement Policies.
3. Right-click and select Create Placement Policy.
4. Enter Name of the Placement Policy (for example, `VM-Host-Infra`).
5. Go to Virtual Slot 1. Click on the associated Selection Preference and select the option Assigned Only from the drop-down list.
6. Click OK twice to create the placement policy.

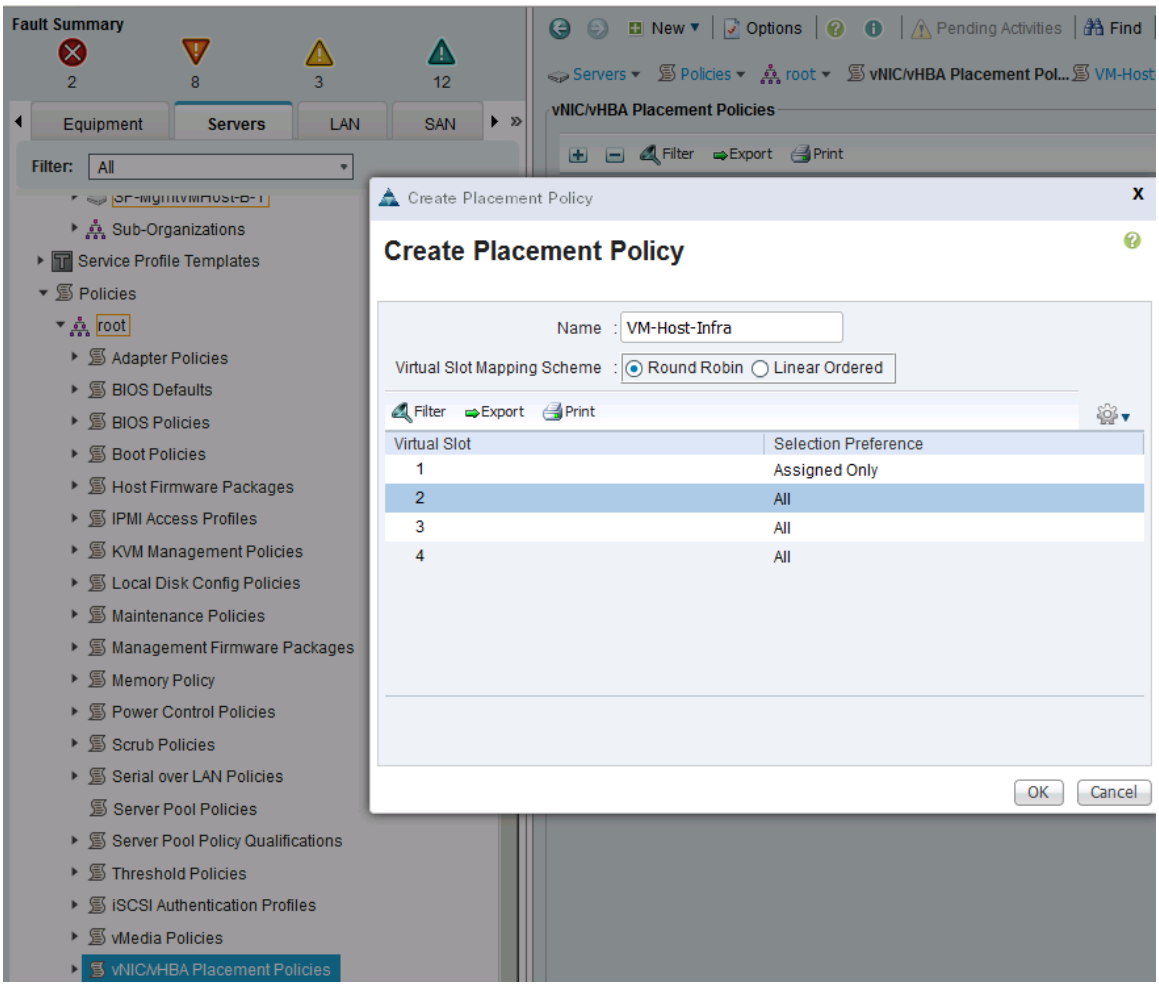

Create Server Pools

Create UUID Suffix Pool

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

1. From Cisco UCS Manager, select Servers tab in the navigation pane.

2. Select Servers > Pools > root > UUID Suffix Pools.
3. Right-click and select Create UUID Suffix Pool.
4. Specify a Name for the UUID suffix pool and click Next.

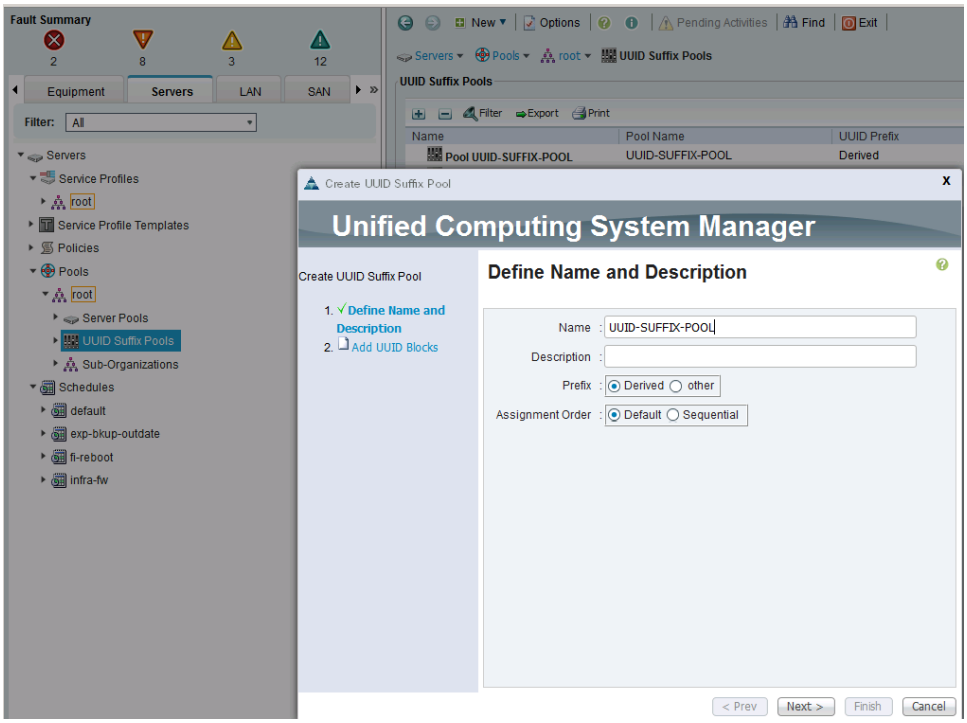

5. Click + to add a block of UUIDs. Alternatively, you can also modify the default pool and allocate/modify a UUID block.

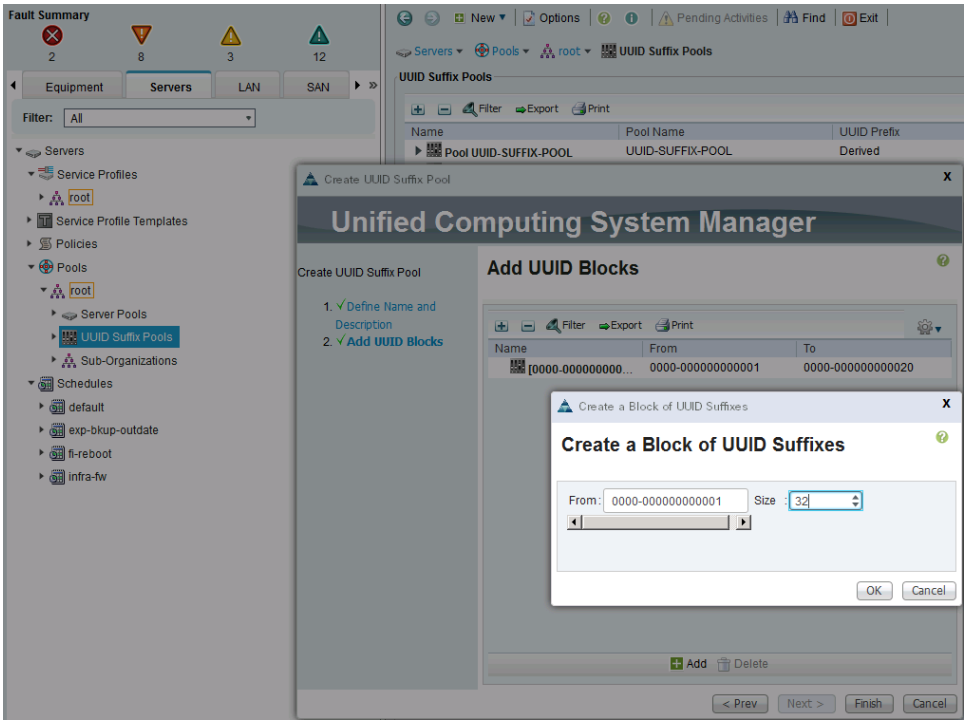

6. Keep the From field at the default setting. Specify a block size (for example, `32`) that is sufficient to support the available blade or server resources.
7. Click OK, click Finish and click OK again to create UUID Pool.

Create Server Pools

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

1. In Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.

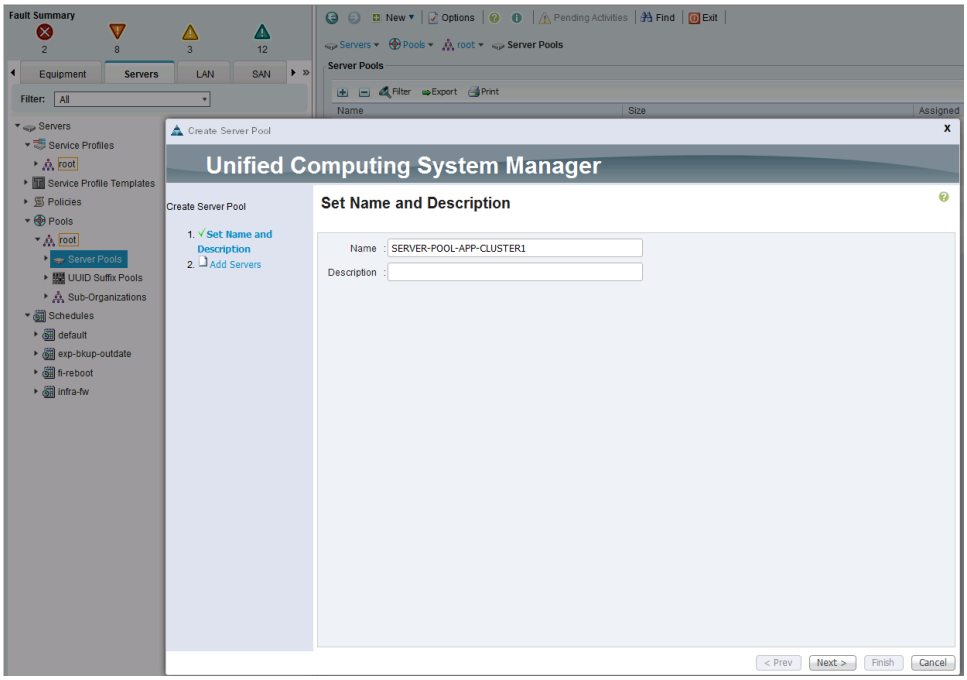

5. Enter name of the server pool (for example, `SERVER-POOL-APP-CLUSTER1`).
6. (Optional) Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the server pool. Click Finish to complete.

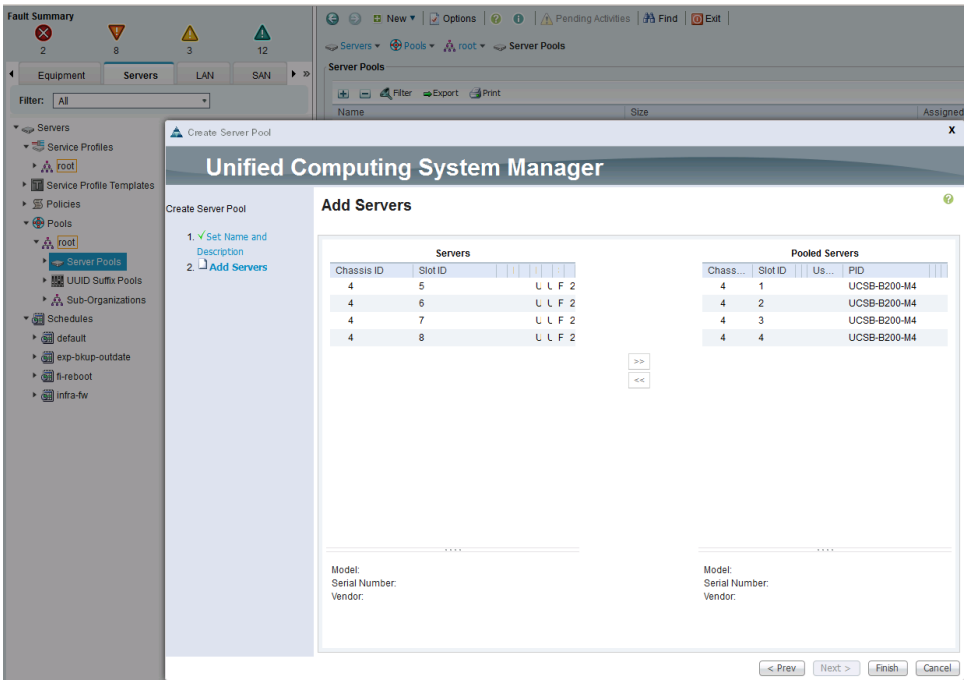

Cisco UCS Configuration – SAN

SAN Configuration Workflow

The workflow below shows the configuration required to create vHBA Templates on Cisco UCS servers. The vHBA Templates encapsulate the SAN configuration aspect of Cisco UCS. Two vHBA templates are created in the SmartStack design, one through Fabric A and another through Fabric B for redundancy. The subsections that follow will cover the configuration of the individual steps in the work flow below.

Figure 15 SAN Configuration Workflow

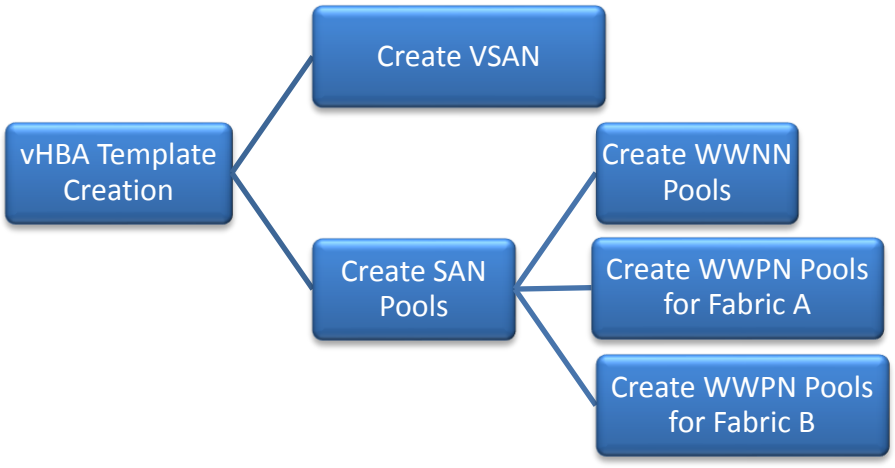

VSAN configuration was completed in an earlier section of this document.

Create SAN Pools

Create WWNN Pools

To configure the necessary World Wide Node Name (WWNN) pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN tab in the navigation pane.
2. Choose Pools > root.
3. Right-click WWNN Pools. Choose Create WWNN Pool.

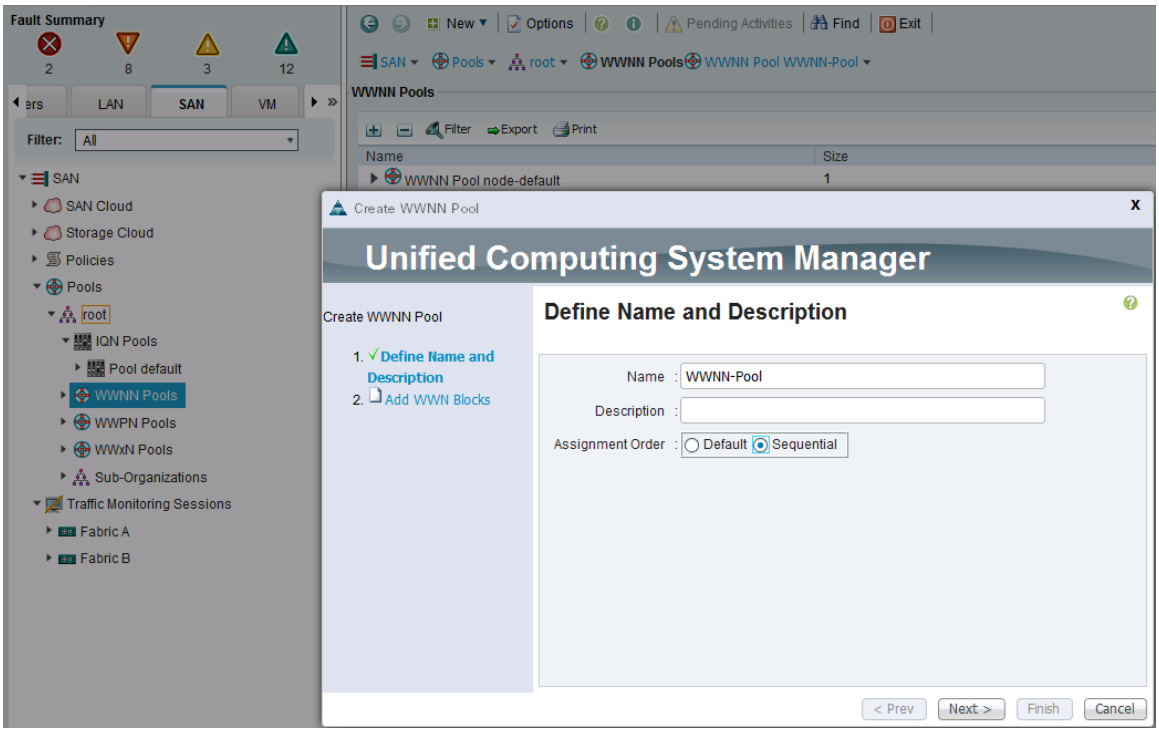

4. Enter the name (for example, `WWNN-Pool`) of the WWNN pool.
5. (Optional) Add a description for the WWNN pool.
6. (Optional) Specify assignment order as Sequential.
7. Click Next.
8. Click + to add a block of WWNNs.

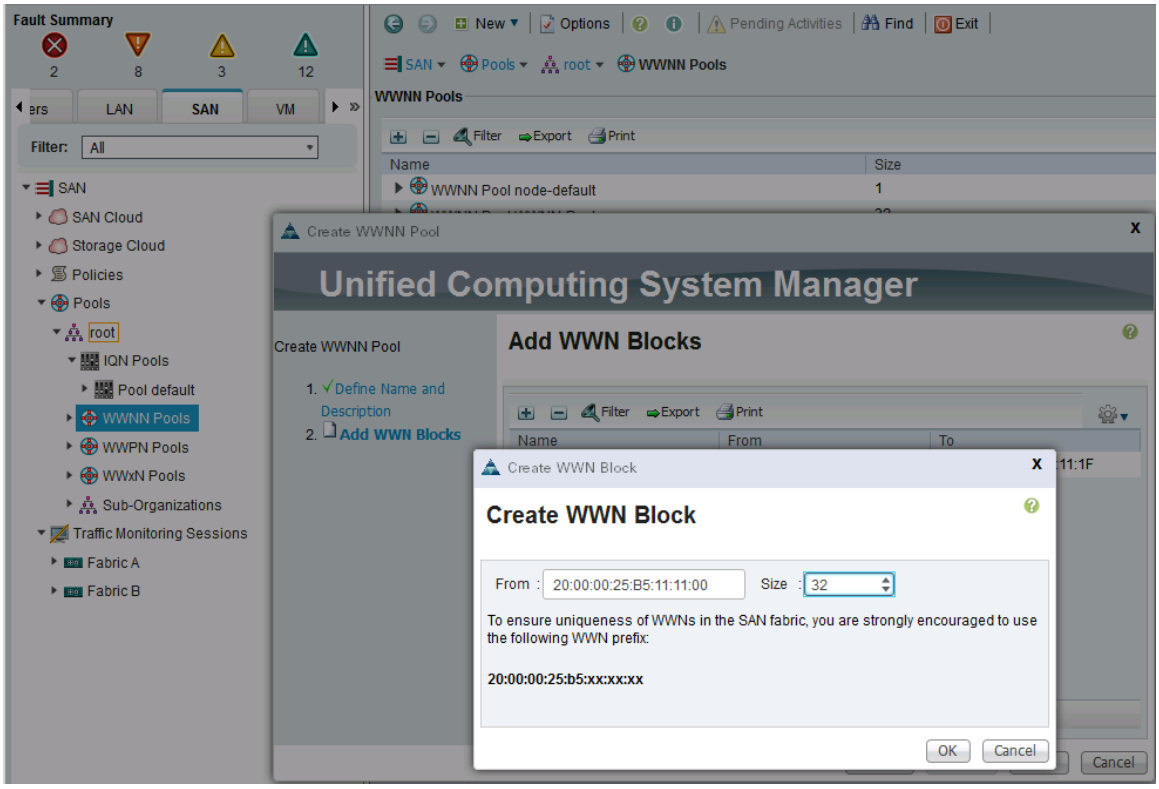

9. Keep the default block of WWNNs, or specify a base WWNN. Change the 6th and 7th octet (for example, `11:11`) for identifying traffic from this Cisco UCS domain for troubleshooting ease.
10. Specify a size for the WWNN block that is sufficient to support the available blade or server resources.
11. Click OK, click Finish. Click OK again to complete the pool configuration.

Create WWPN Pools

Two World Wide Port Name (WWPN) pools are created in this procedure, one for SAN Fabric A and another for SAN Fabric B. The 7th octet of the starting WWPN is modified to AA and BB to identify the WWPNs as Fabric A and Fabric B addresses respectively. The 6th octet is same as the value used for the WWNN pool above.

Create SAN Fabric A WWPN Pools

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

1. In Cisco UCS Manager, click SAN tab in the navigation pane.
2. Choose Pools > root.
3. Right-click WWPN Pools.
4. Choose Create WWPN Pool.

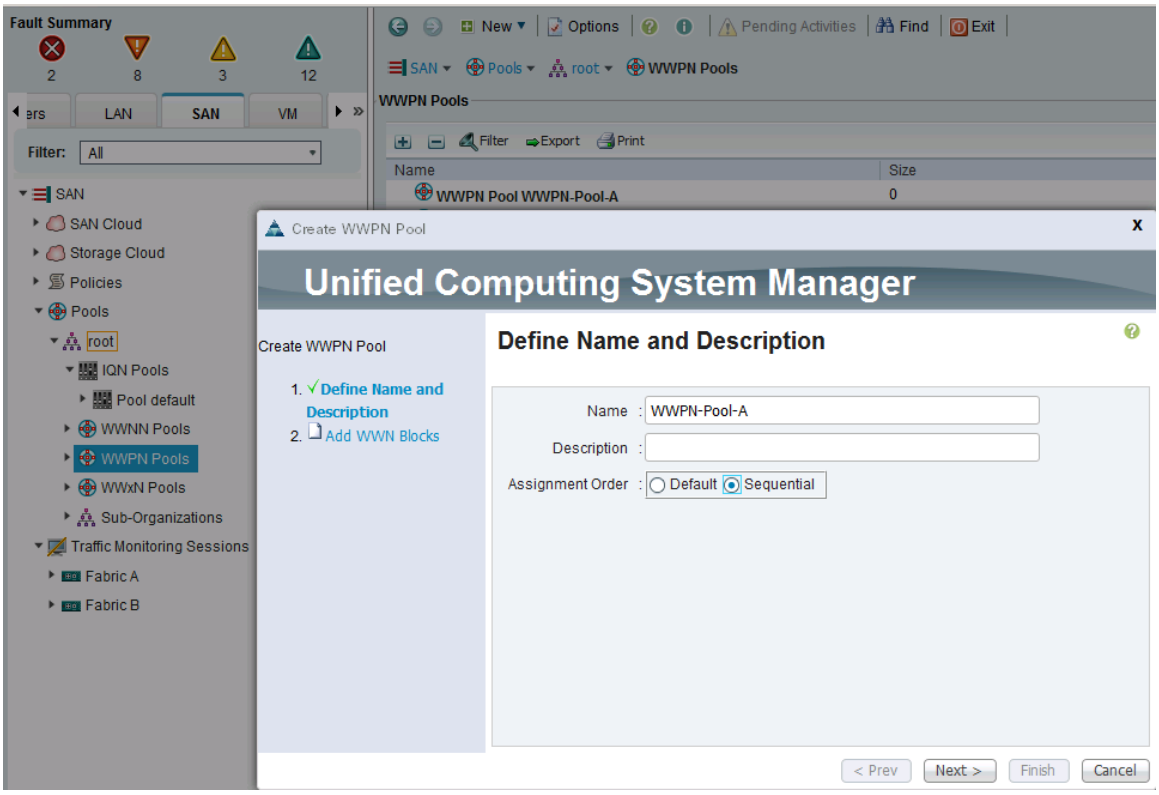

5. Enter the name (for example, `WWPN-Pool-A`) of the WWPN pool for Fabric A.
6. (Optional) Enter a description for this WWPN pool.
7. (Optional) Specify assignment order as Sequential.
8. Click Next.
9. Click + to add a block of WWPNs.

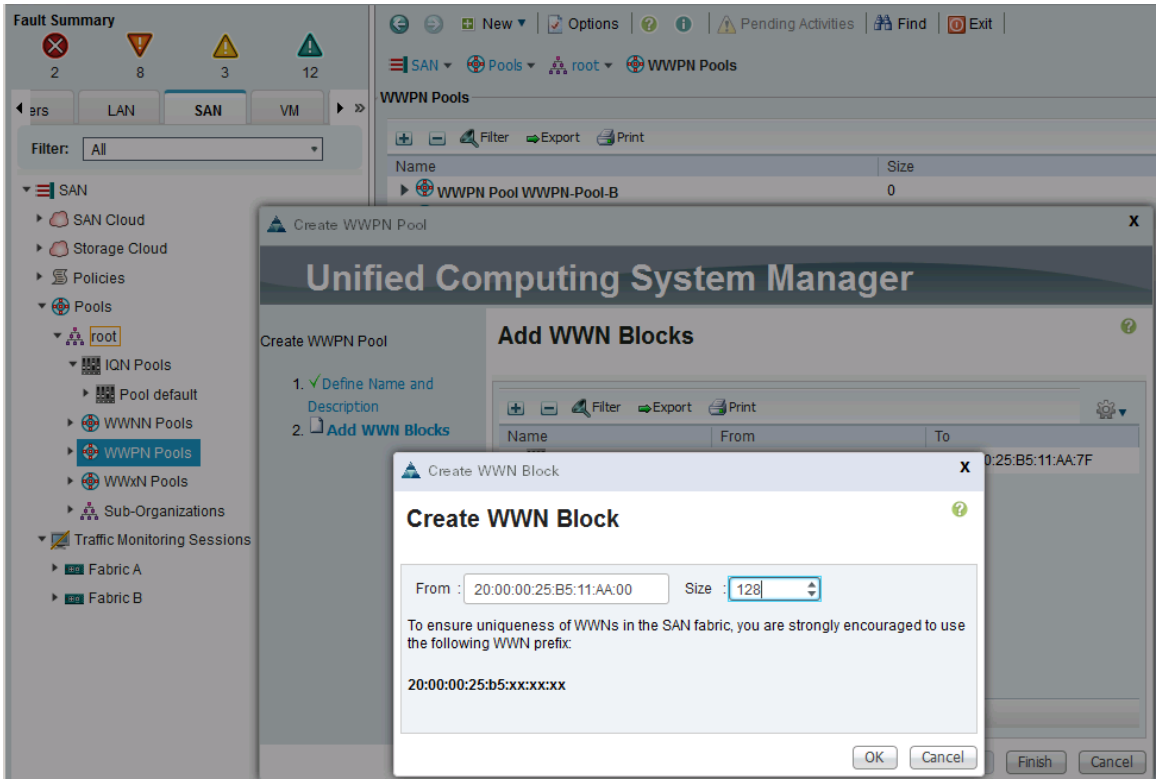

10. Specify the starting WWPN in the block for Fabric A.
11. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.
12. Click OK, click Finish. Click OK again to complete the pool configuration.

Create SAN Fabric B WWPN Pools

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

1. In Cisco UCS Manager, click SAN tab in the navigation pane.
2. Choose Pools > root.
3. Right-click WWPN Pools.

Choose Create WWPN Pool.

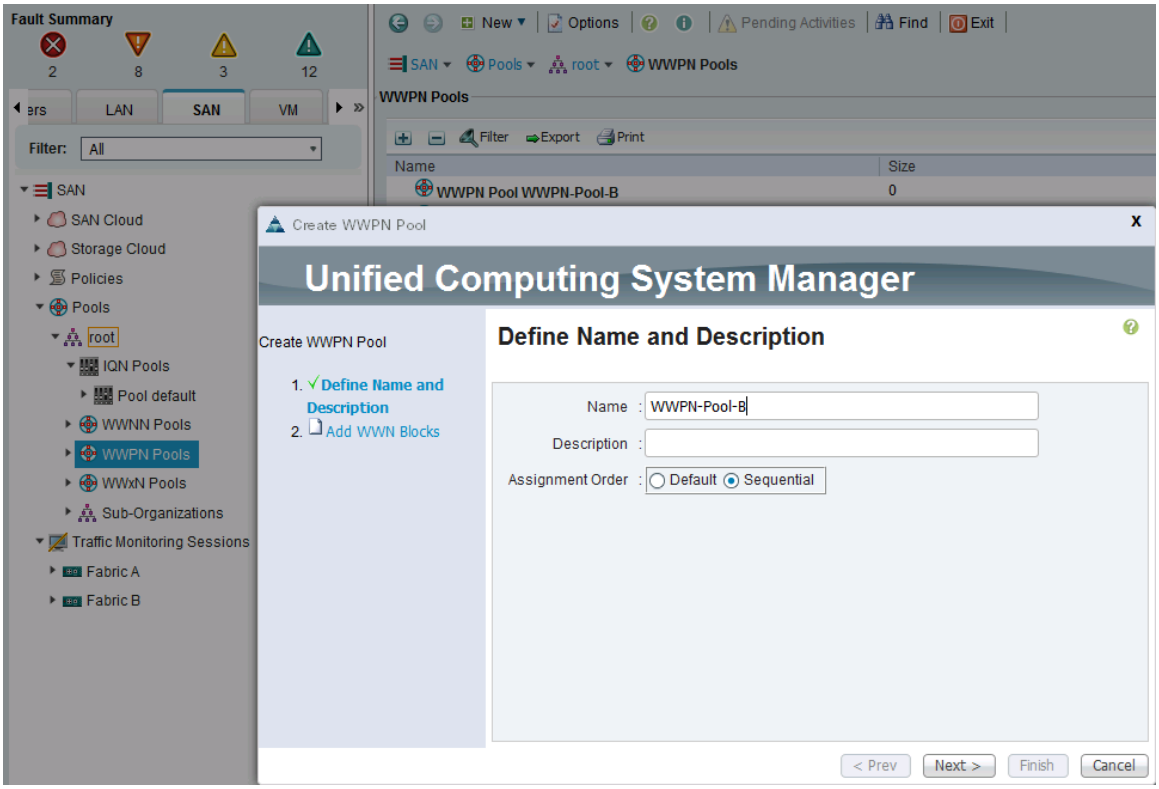

4. Enter the name (for example, `WWPN-Pool-B`) of the WWPN pool for Fabric B.
5. (Optional) Enter a description for this WWPN pool.
6. (Optional) Specify assignment order as Sequential.
7. Click Next.
8. Click + to add a block of WWPNs.

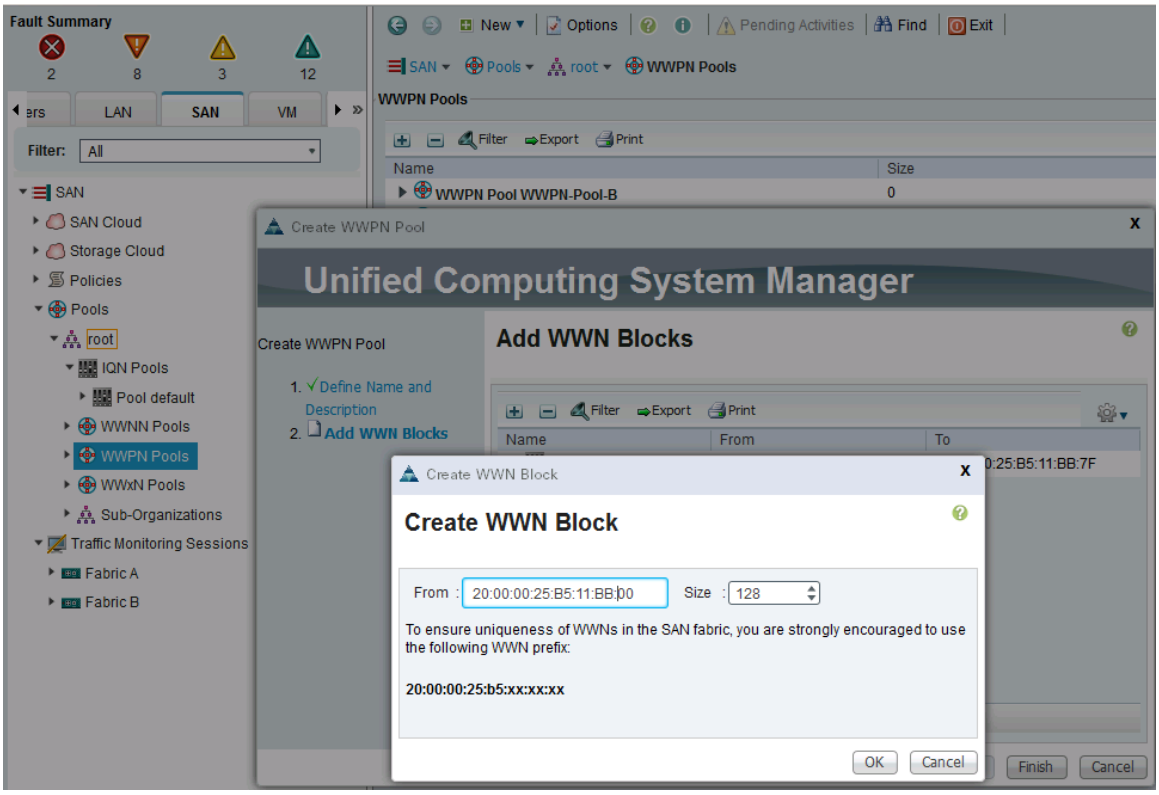

9. Specify the starting WWPN in the block for Fabric B.
10. Specify a size for the WWPN block that is sufficient to support the available blade or server resources.
11. Click OK, click Finish. Click OK again to complete the pool configuration.

Create vHBA Templates

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

Create vHBA Template for Fabric A

In Cisco UCS Manager, click SAN tab in the navigation pane.

1. Select Policies > root.
2. Right-click vHBA Templates.
3. Select Create vHBA Template.
4. Enter vHBA template name (for example, `vHBA-TEMPLATE-A`).
5. Click the radio button Fabric A.
6. For VSAN, select previously created `VSAN-A` from the drop-down list.
7. For WWPN Pool, select previously created `WWPN-POOL-A` from the drop-down list.

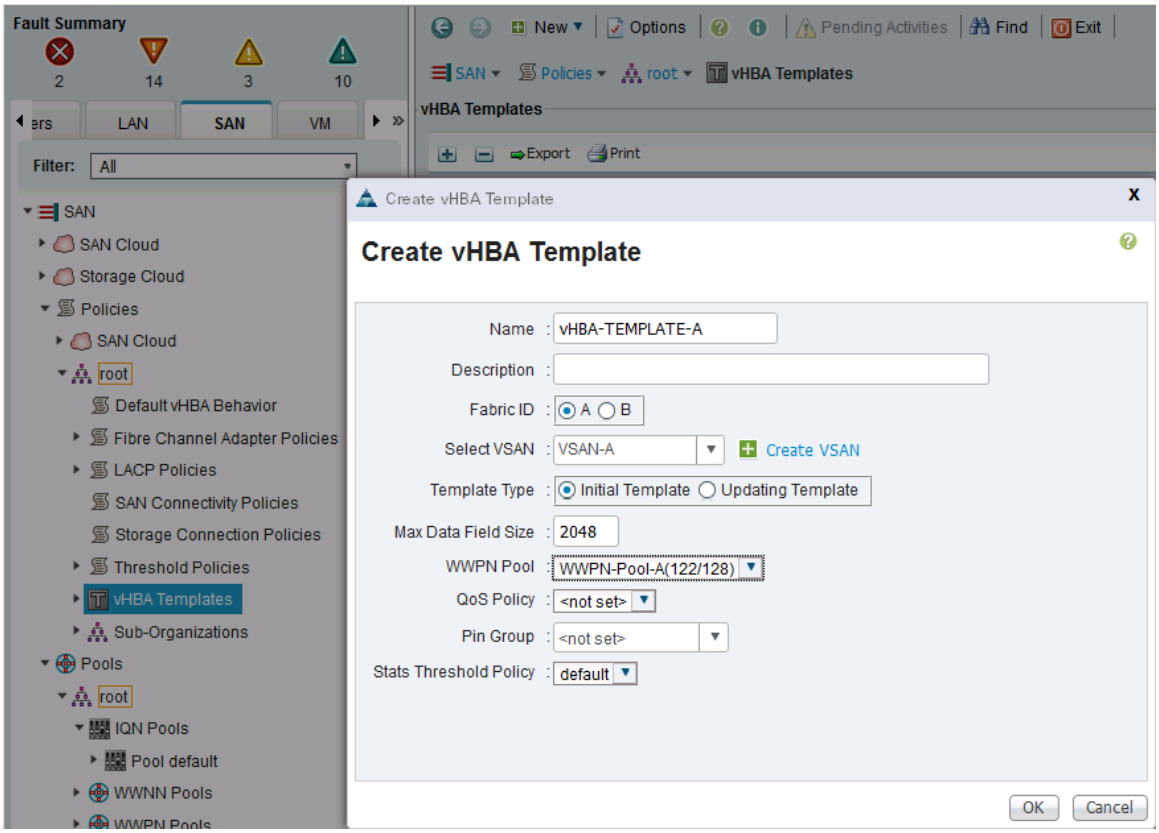

8. Click OK to create the vHBA template.
9. Click OK.

Create vHBA Template for Fabric B

1. In the navigation pane, click SAN tab.

2. Choose Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.
5. Enter vHBA template name (for example, vHBA-TEMPLATE-B).
6. Click the radio button Fabric B.
7. For VSAN, select previously created VSAN-B from the drop-down list.
8. For WWPN Pool, select previously created WWPN-POOL-B from the drop-down list.

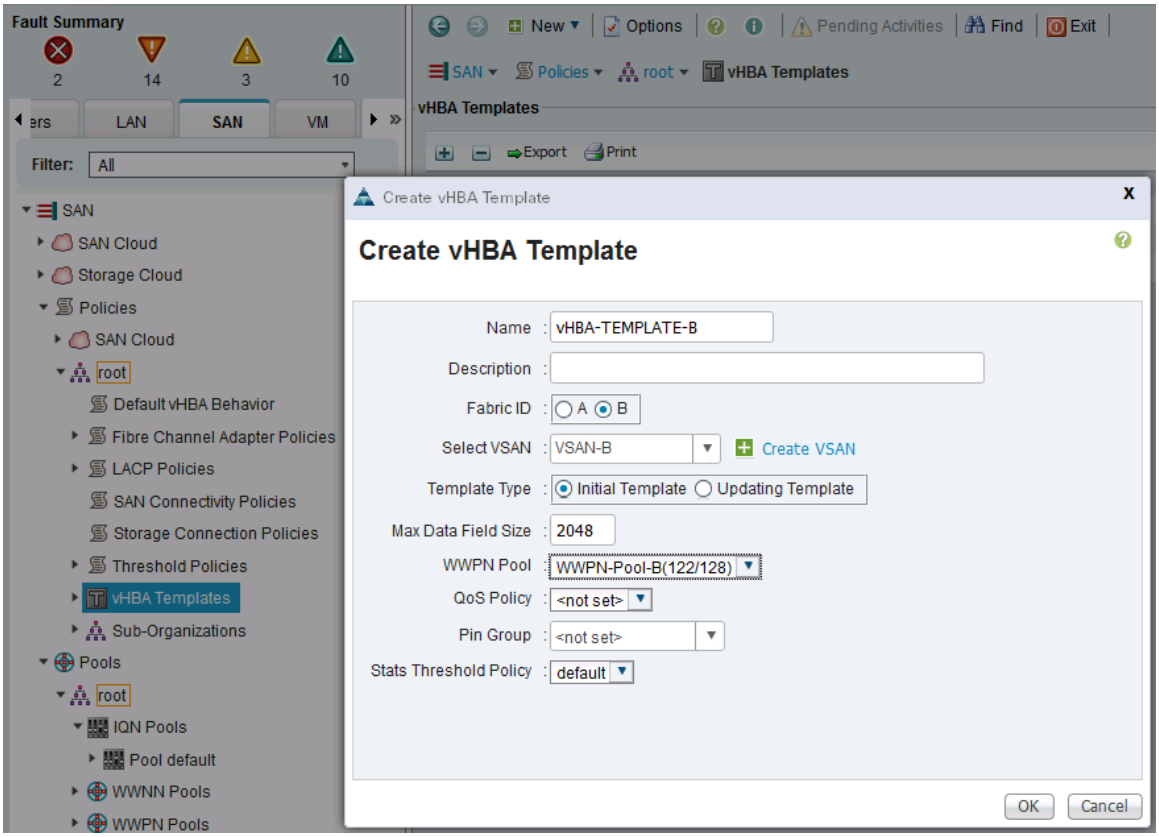

9. Click OK to create the vHBA template.
10. Click OK.

Create Service Profile Templates

In this procedure, two service profile templates are created: one for Fabric A boot and one for Fabric B boot. The first profile is created and then cloned and modified for the second host.

Create Service Profile Template for Fabric A

To create the service profile templates for Fabric A boot, complete the following steps.

1. From Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Servers > Service Profile Template > root.

3. Right-click root and select Create Service Profile Template to open the Create Service Profile Template wizard.
4. In the Identify the Service Profile Template screen, configure the following:
    a. Enter name (for example, `SP-AppVMHost-AF-FI-A`) for the service profile template.
    b. Select Updating Template radio button.
    c. Under UUID, select the previously configured UUID pool (for example, `UUID-SUFFIX-POOL`).
    d. Click Next.

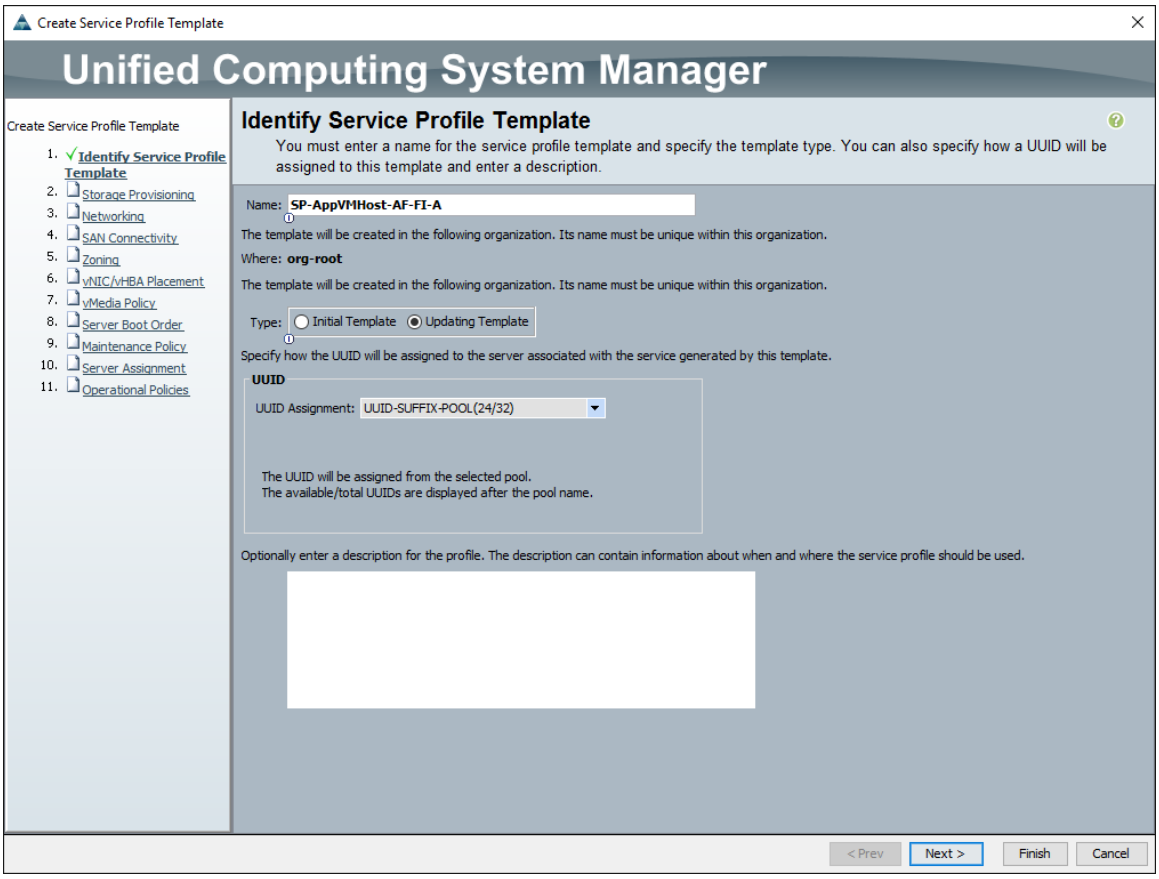

5. In the Storage Provisioning screen, configure the following:
    a. Go to Local Disk Configuration Policy tab.
    b. In the Local Storage drop-down, select a policy. If the servers will not have local disks, choose the previously configured policy (for example, `SAN-Boot`) otherwise choose the `default` policy.
    c. Click Next.

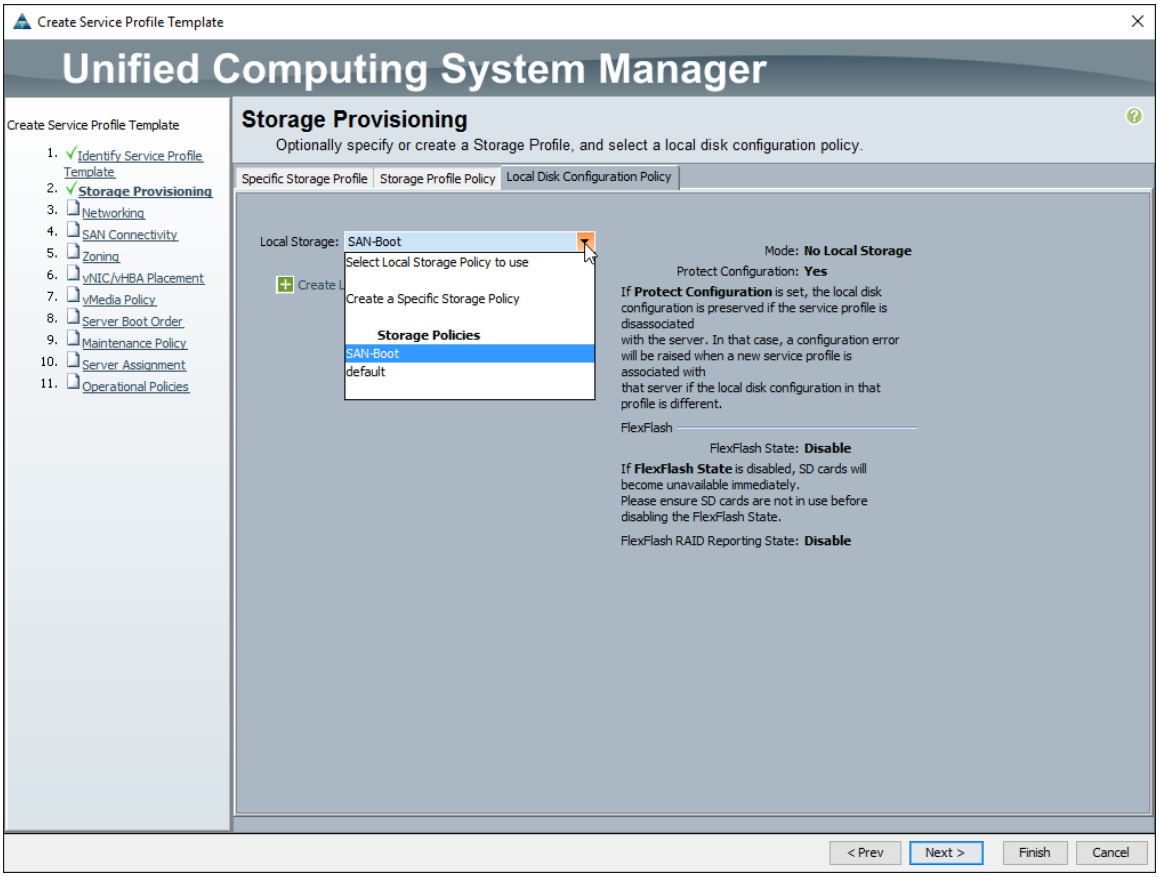

6. In the Networking screen, configure the following:
    a. For Dynamic vNIC Connection Policy, keep the default setting.
    b. Click Expert radio button to configure the LAN connectivity.

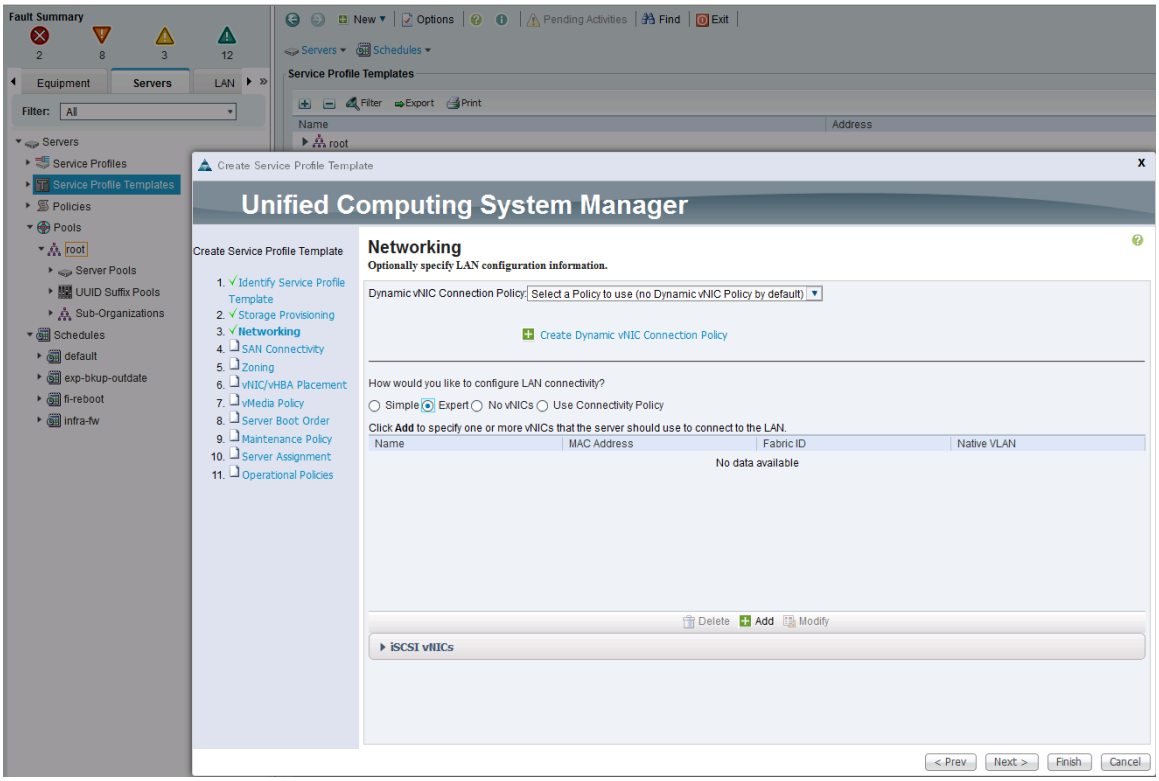

c. Click + to add a vNIC to the template.

d. In the Create vNIC dialog box:

   i. Enter the name (for example, `vNIC-A`) of the vNIC.
   ii. Check the Use vNIC Template check box.
   iii. In the vNIC Template list, choose the previously created vNIC Template for Fabric A boot (for example, `vNIC-Template-A`).
   iv. In the Adapter Policy list, choose VMWare.
   v. Click OK to add this vNIC to the template.

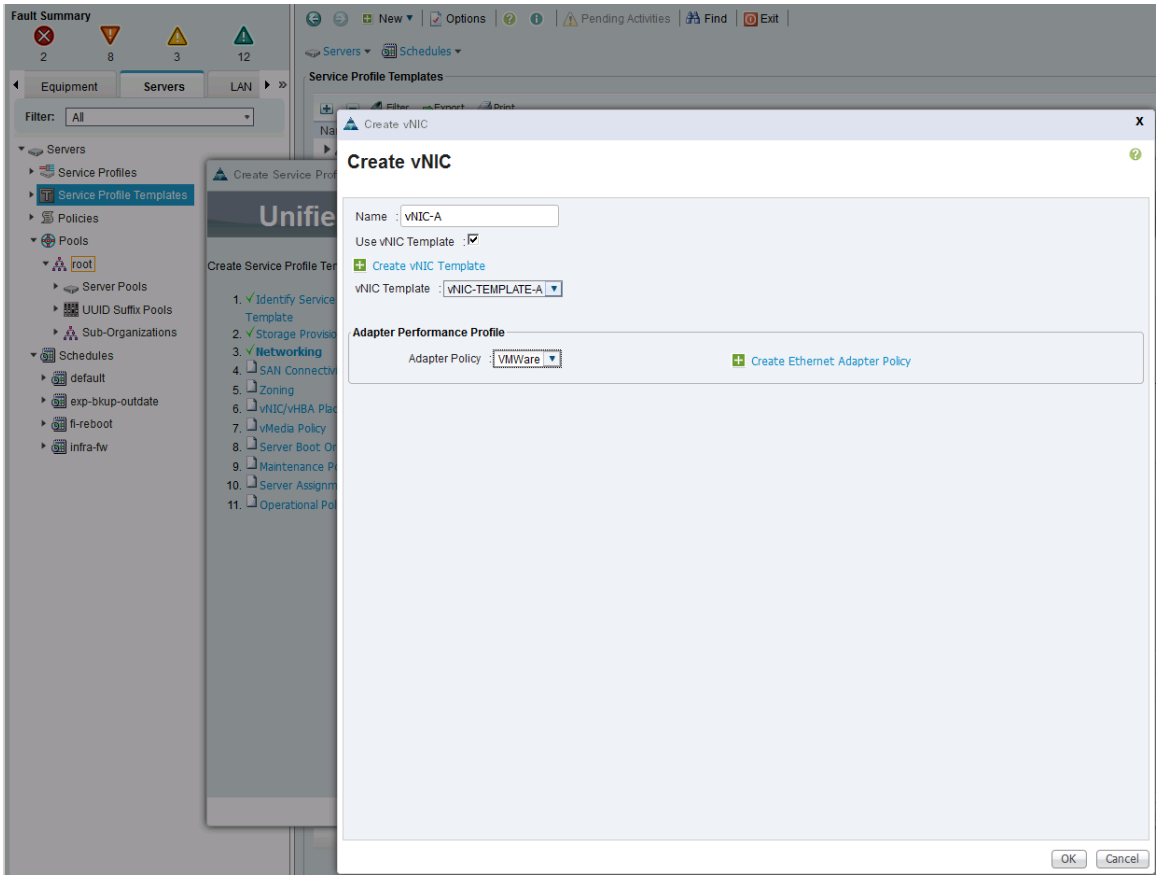

e. Click + to add a 2nd vNIC to the template.

f. In the Create vNIC dialog box:

   i. Enter the name (for example, vNIC-B) of the vNIC.
   ii. Check the Use vNIC Template check box.
   iii. In the vNIC Template list, choose the previously created vNIC Template for Fabric A boot (for example, vNIC-Template-B).
   iv. In the Adapter Policy list, choose VMWare.
   v. Click OK to add this vNIC to the template.

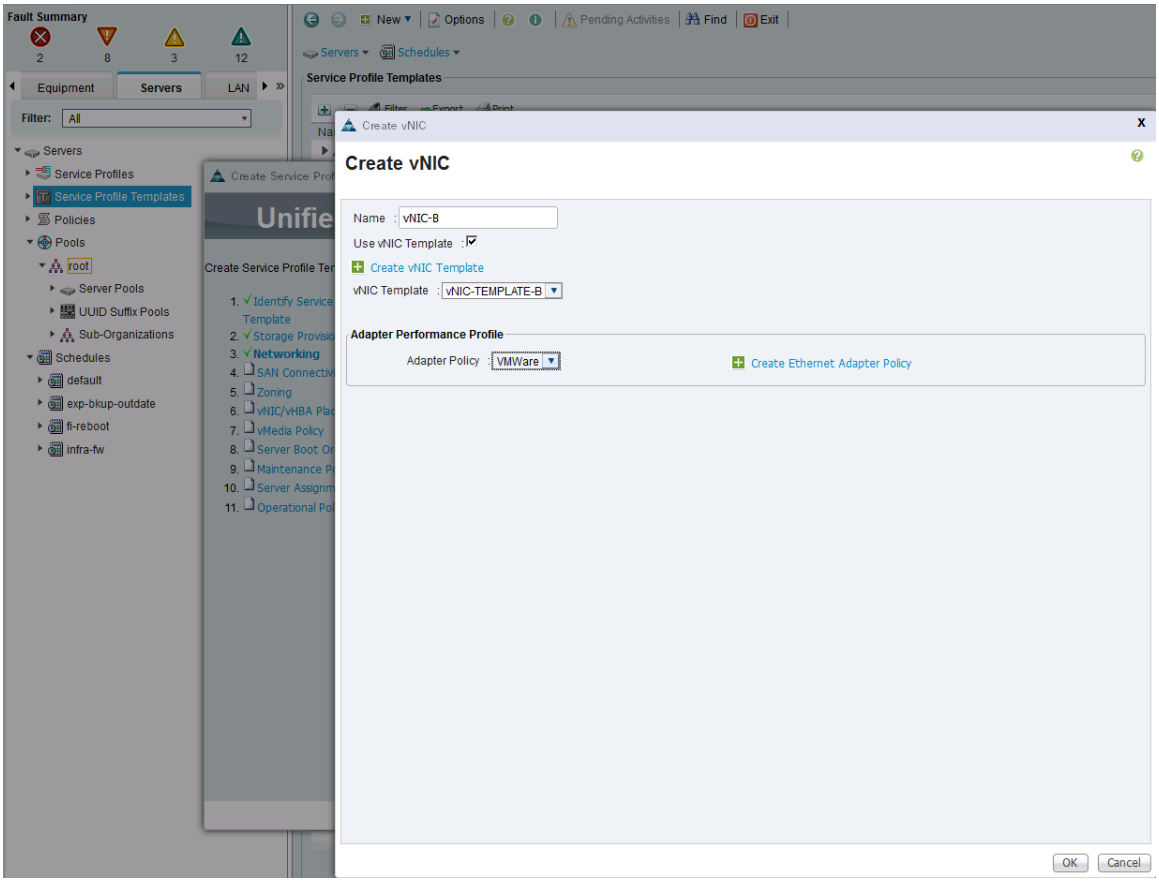

g. Review the configuration on the Networking screen of the wizard. Confirm that both vNICs were created. Click Next.

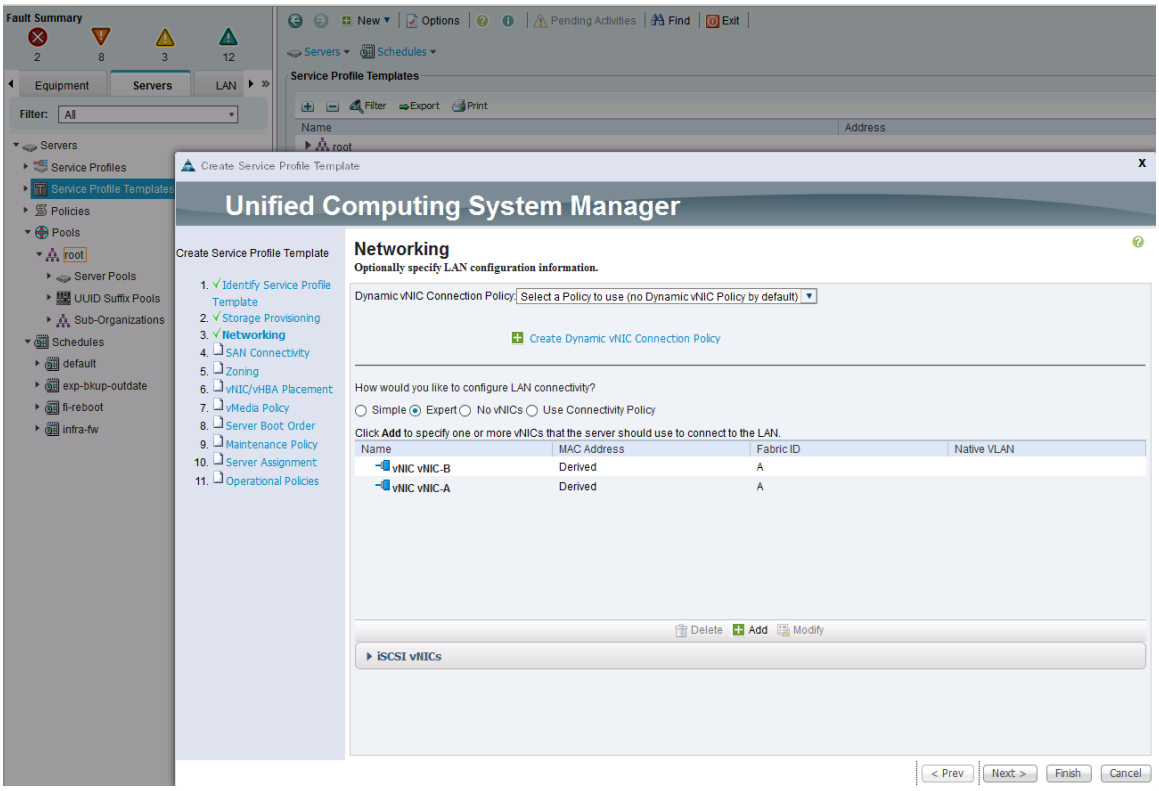

7. In the SAN Connectivity screen, configure the following:
    a. Click Expert radio button to configure the SAN connectivity.
    b. In the WWNN Assignment list, select the previously configured WWNN Pool (for example, `WWNN-Pool`).
    c. Click + to add a vHBA to the template.

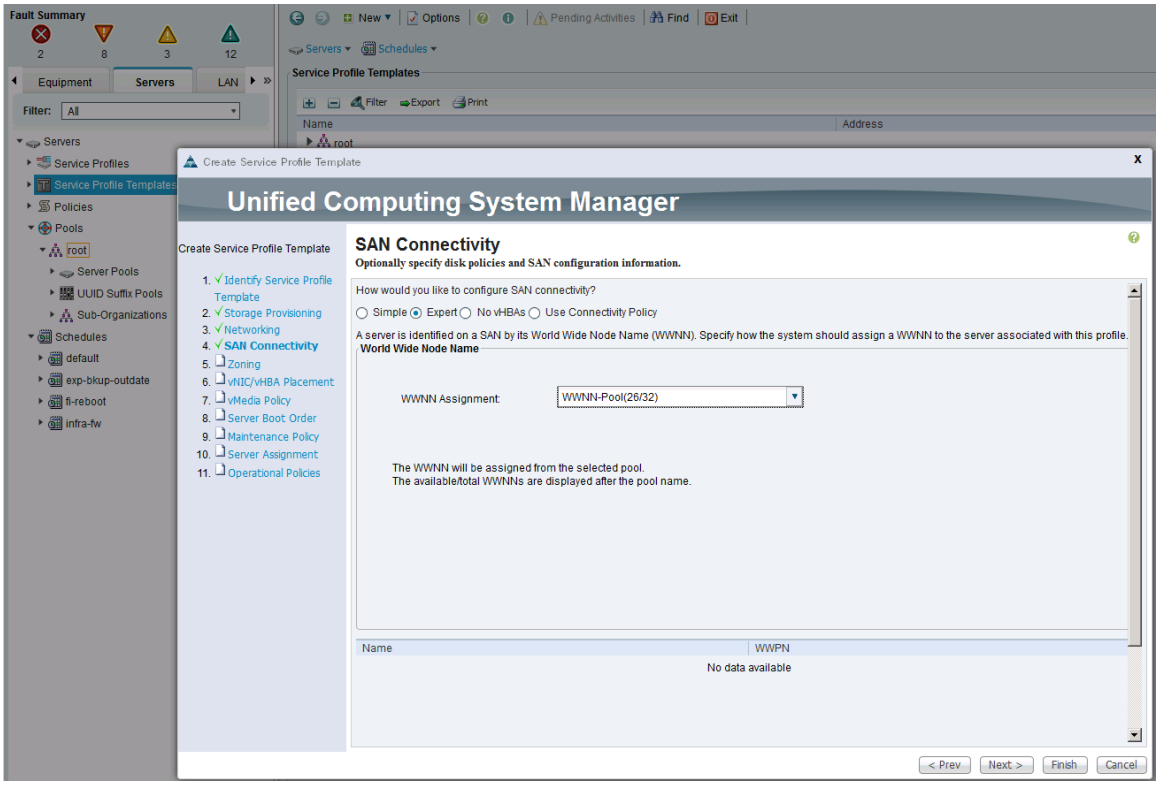

d. In the Create vHBA dialog box:

i. Enter name of the vHBA (for example, vHBA-A).
ii. Check the Use vHBA Template check box.
iii. In the vHBA Template list, choose previously created template (for example, vHBA-Template-A).
iv. In the Adapter Policy list, choose VMware.
v. Click OK to add this vHBA to the template.

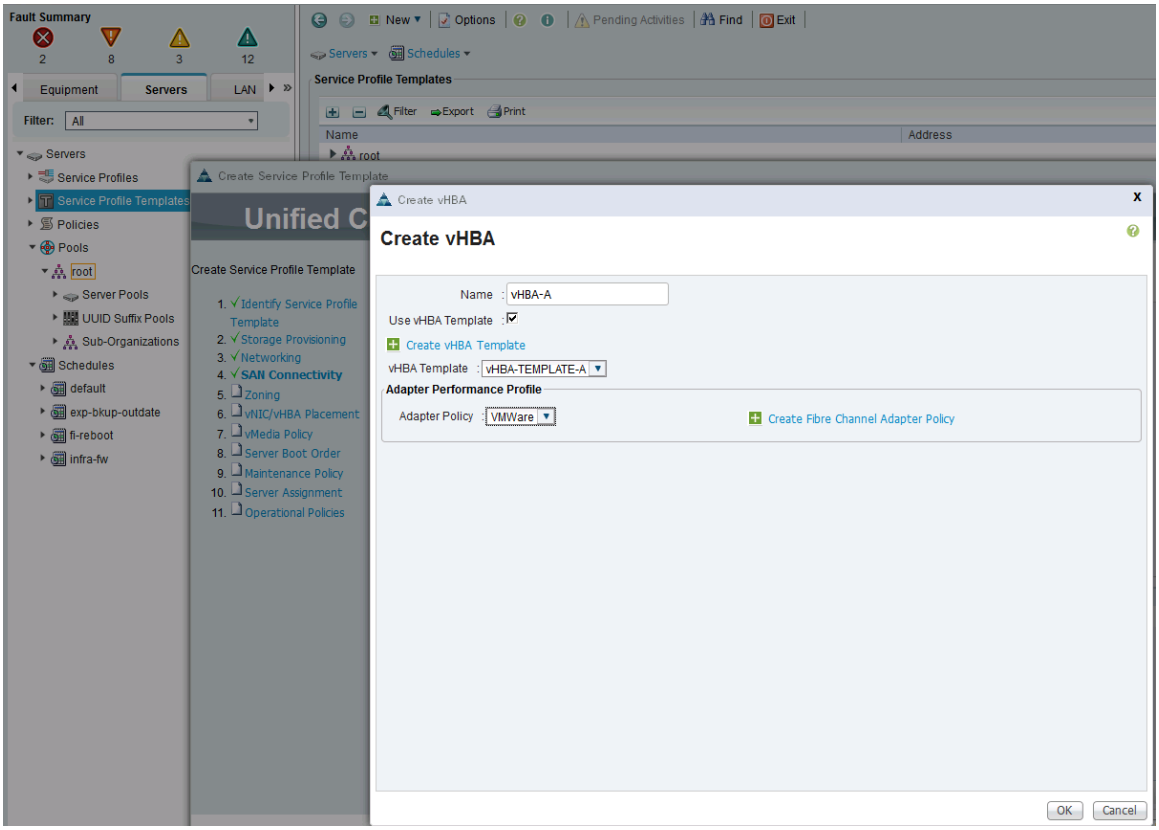

e. On the SAN Connectivity page of the wizard, click [+] Add at the bottom of the page to add a 2nd vHBA to the template.

f. In the Create vHBA dialog box:

   i. Enter the name of the vHBA (for example, vHBA-B).
   ii. Check the check box for Use HBA Template.
   iii. In the vHBA Template list, choose previously created template (for example, vHBA-Template-B).
   iv. In the Adapter Policy list, choose VMware.
   v. Click OK to add the vHBA to the template.

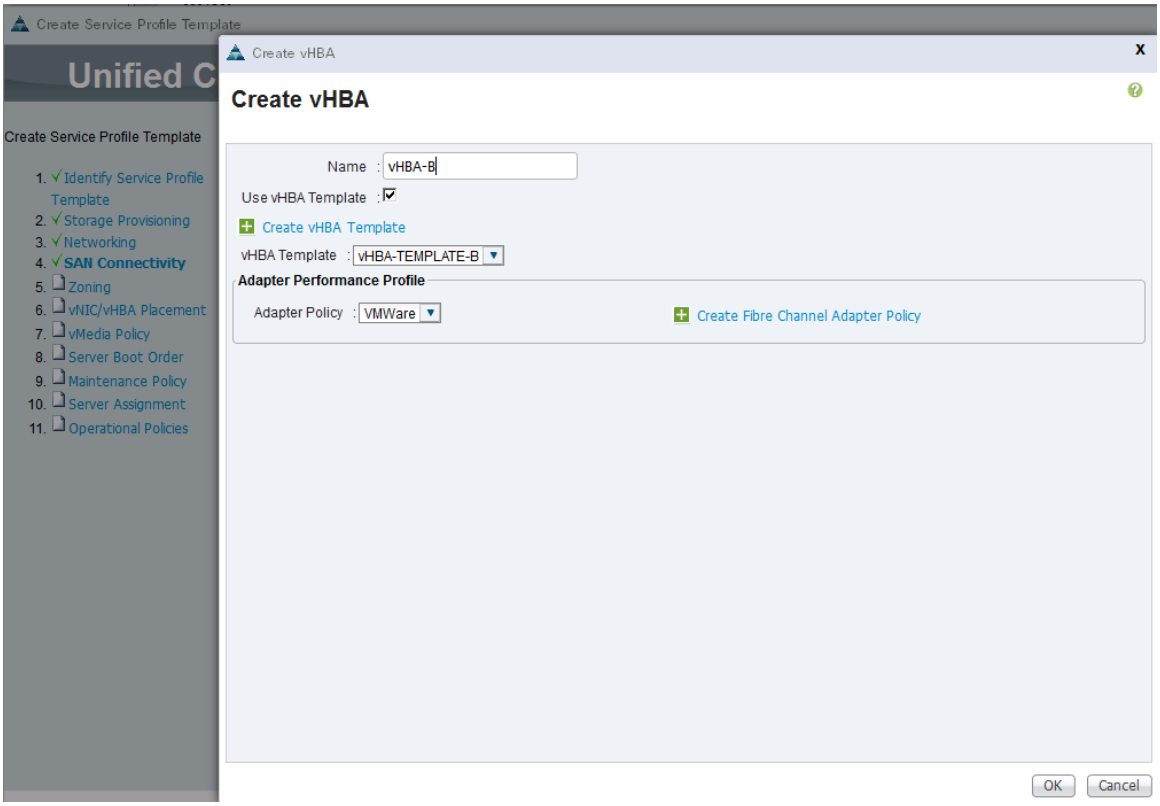

g. Review the table in the SAN Connectivity page to verify that both A and B vHBAs were created.

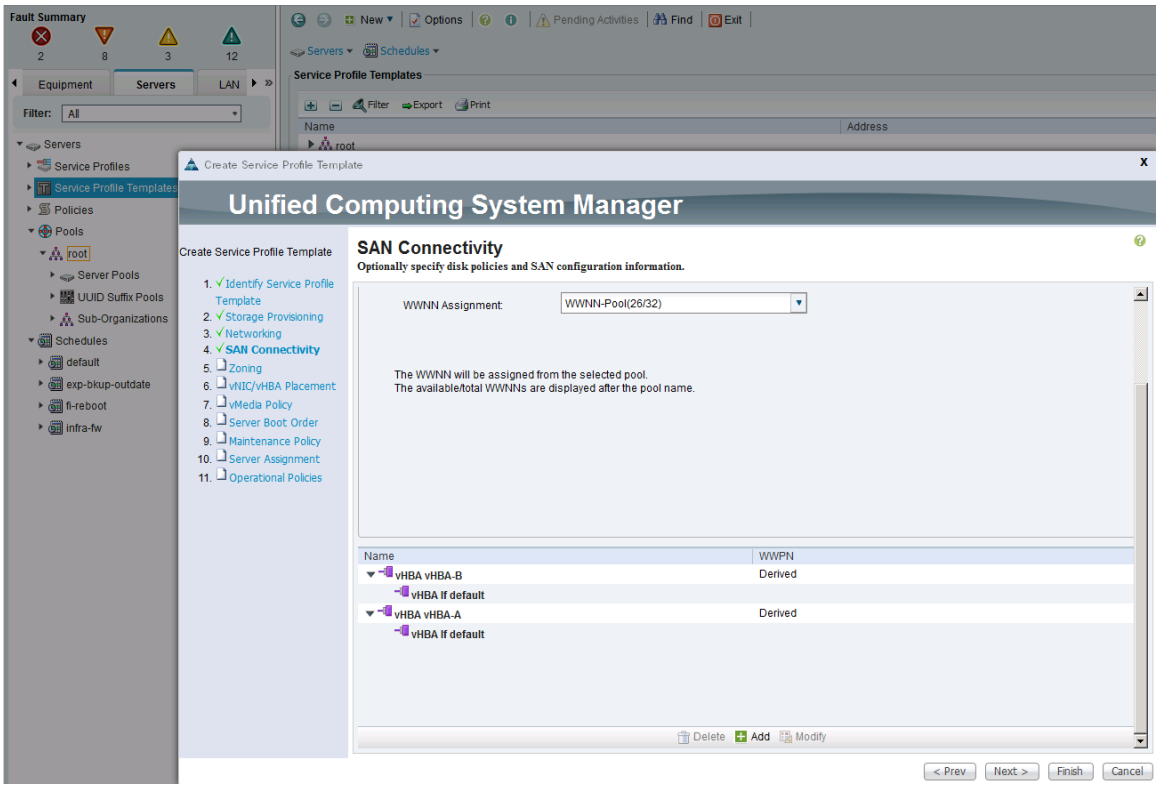

h. Click Next.

8. In the Zoning screen, take the defaults as zoning is performed at the Cisco MDS switches. Click Next.

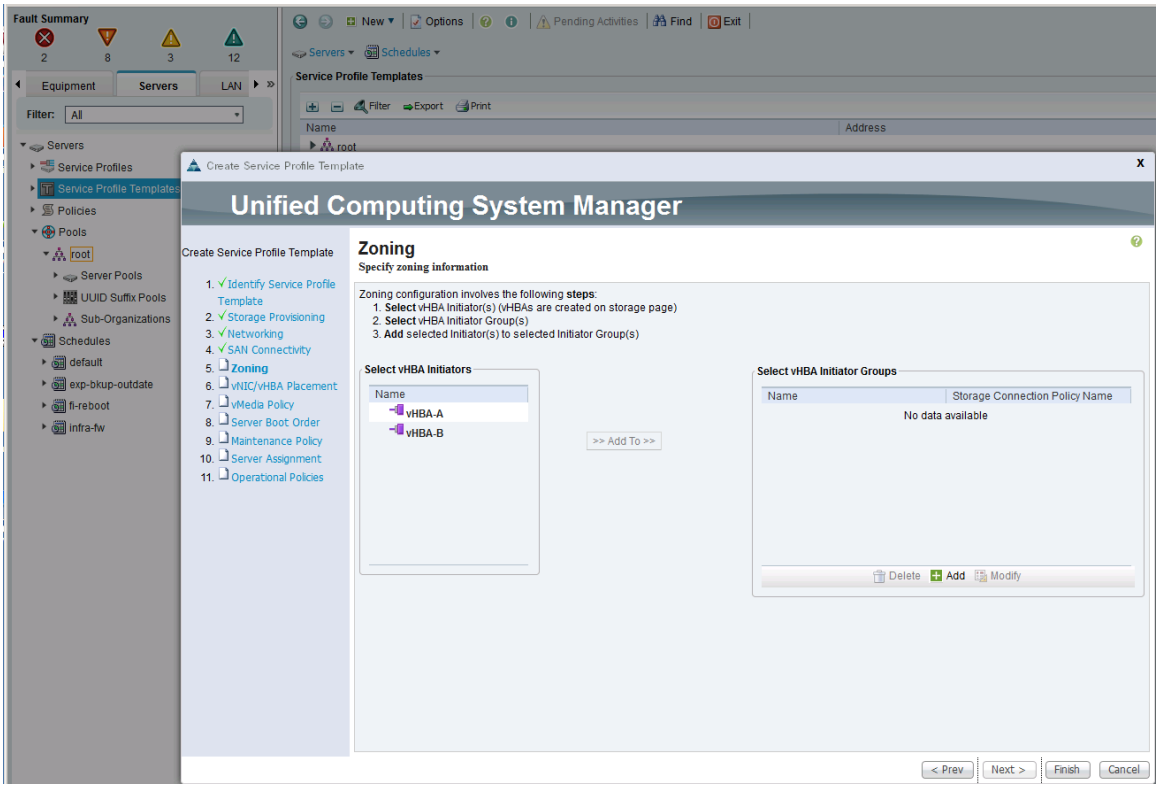

9. In the vNIC/vHBA Placement screen, configure the following:

a. In the Select Placement list, select the previously created policy (for example, `VM-Host-Infra`).

b. Choose vCon1 and assign the vHBAs/vNICs to the virtual network interfaces policy in the following order:

i. vHBA-A
ii. vHBA-B
iii. vNIC-A
iv. vNIC-B

The above step currently does not work through the web based Cisco UCSM for assigning vHBAs. Until this issue is resolved, the same configuration can be done through the Java based Cisco UCSM client.

c. Review the table to verify that all vNICs and vHBAs were assigned to the policy in the appropriate order.

d. Click Next.

10. Click Next to bypass the vMedia Policy screen.

11. In the Set Boot Order screen, select the previously created boot policy list for Fabric A (for example, `BOOT-AF-FI-A`).

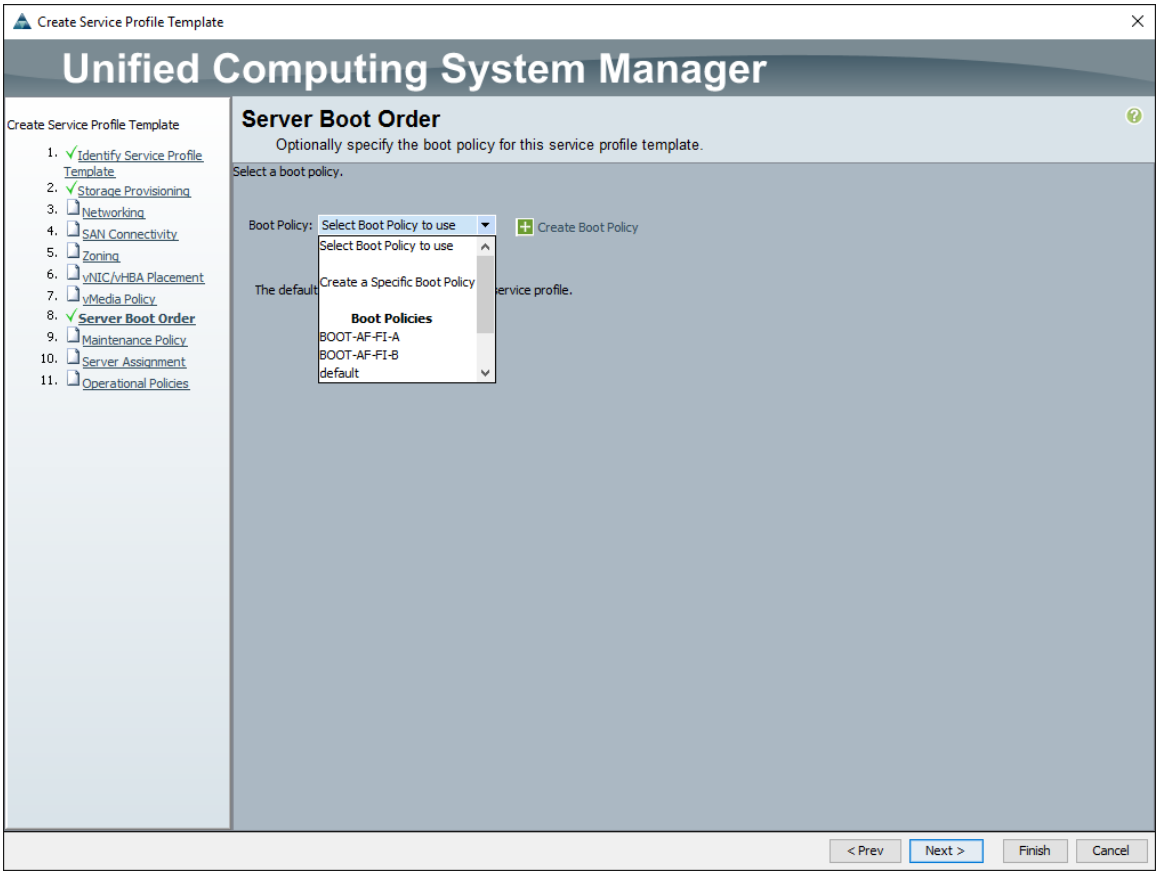

12. Review the Boot Order and verify that the boot sequence is correct.

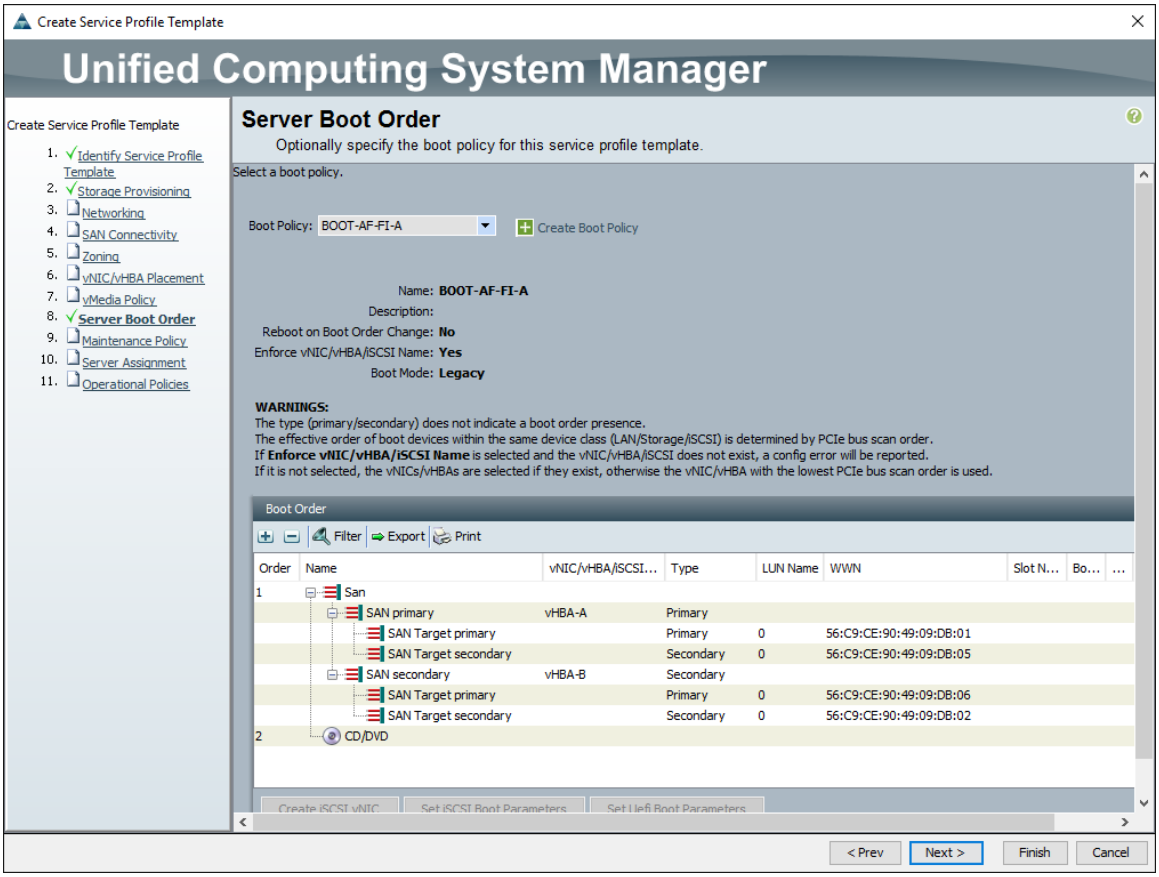

13. Click Next.
14. In the Maintenance Policy screen, choose the previously created maintenance policy (for example, `User-ACK`).

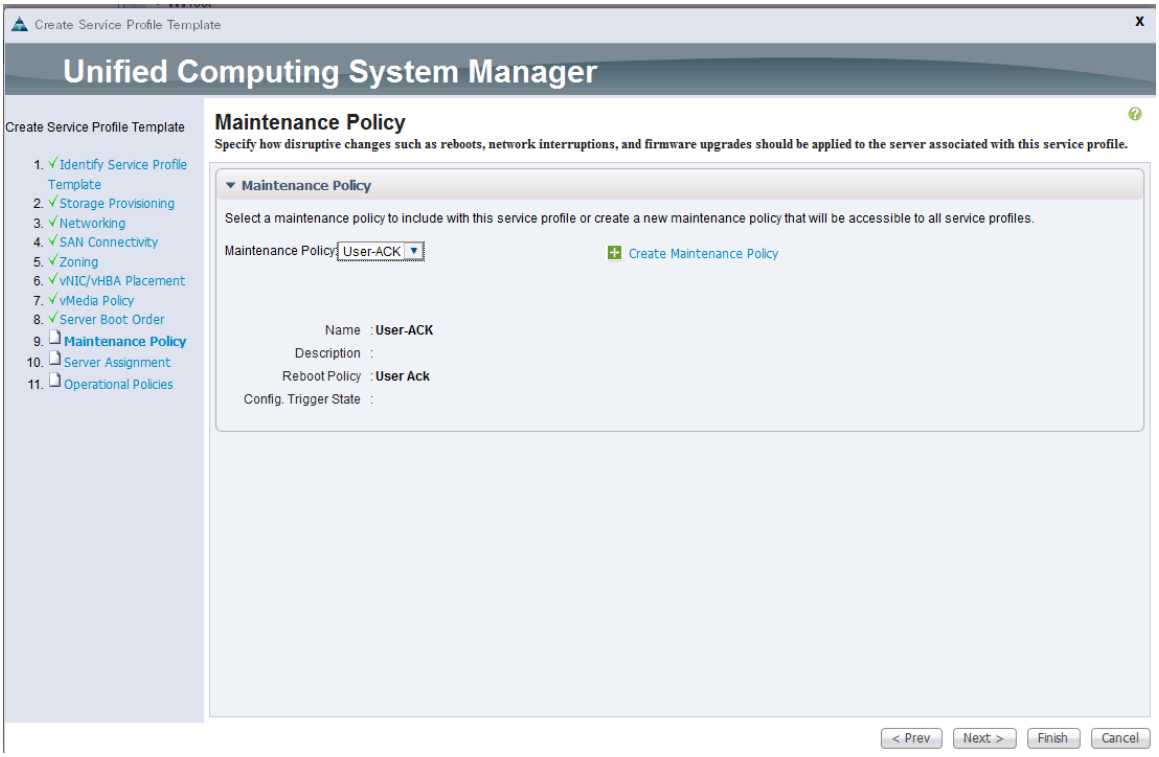

15. Click Next.

16. In the Server Assignment screen, configure the following:

    a. For Pool Assignment, choose the previously created policy from the list (for example, `SERVER-POOL-APP-CLUSTER1`).

    b. Leave the Power State as UP for when the Profile is applied to a server

    c. For Server Pool Qualification, select the previously created policy from the list (for example, `Cisco UCSB-B200-M4`).

    d. Expand the Firmware Management section. For the Host Firmware Package, select the previously selected policy from the list (for example, `HOST-FW-POLICY`).

    e. Click Next.

17. In the Operation Policies screen, configure the following:

    a. For the BIOS Policy list, select the previously configured policy (for example, `VM-Host-Infra`).

    b. Expand Power Control Policy Configuration. For IPMI Access Profile, select the previously configured policy (for example, `No-Power-Cap`).

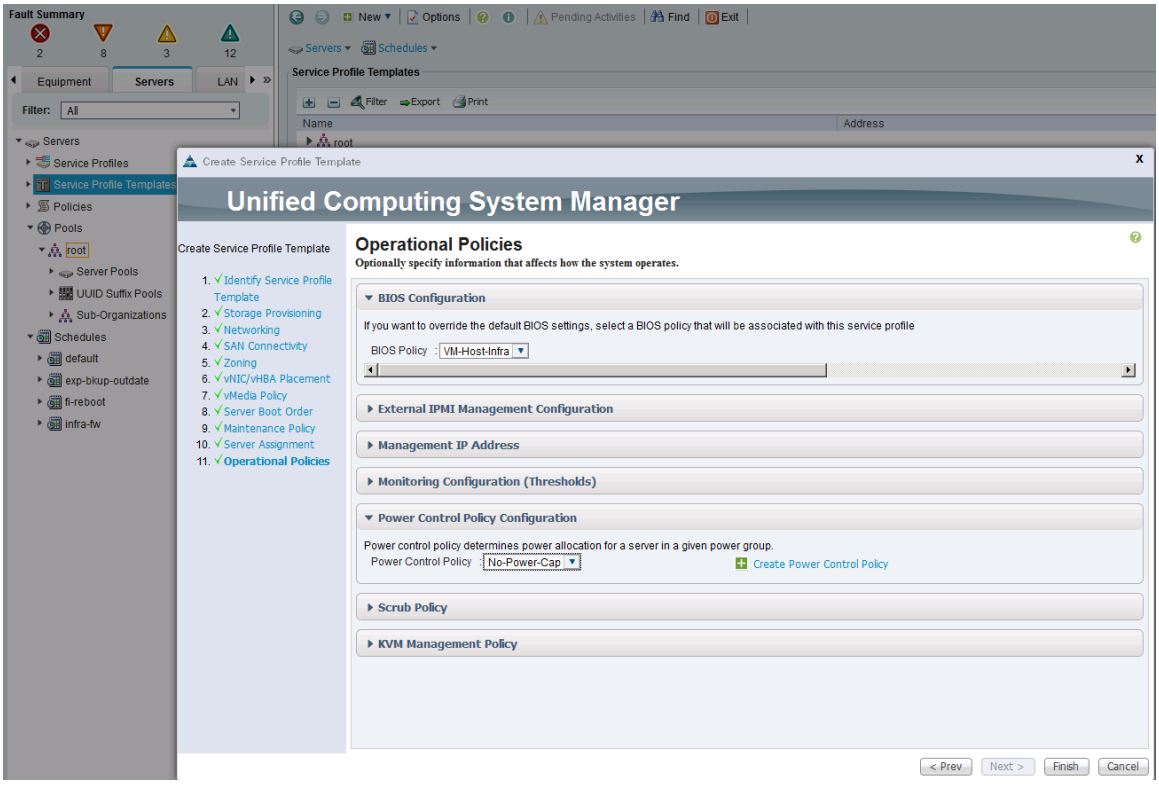

18. Click Finish to complete the creation of the Service Profile Template.

Create Service Profile Template for Fabric B

To create the service profile templates for Fabric B boot, complete the following steps.

1. From Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Servers > Service Profile Template > root.
3. Right-click the previously created `SP-AppVMHost-AF-FI-A` template. Select Create a Clone.
4. In the dialog box, enter `SP-AppVMHost-AF-FI-B` as the name of the clone, choose the root Org, and click OK.

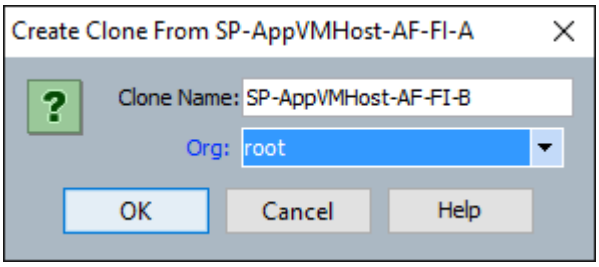

5. Choose the newly cloned service profile template and click Boot Order tab
6. Click Modify Boot Policy.

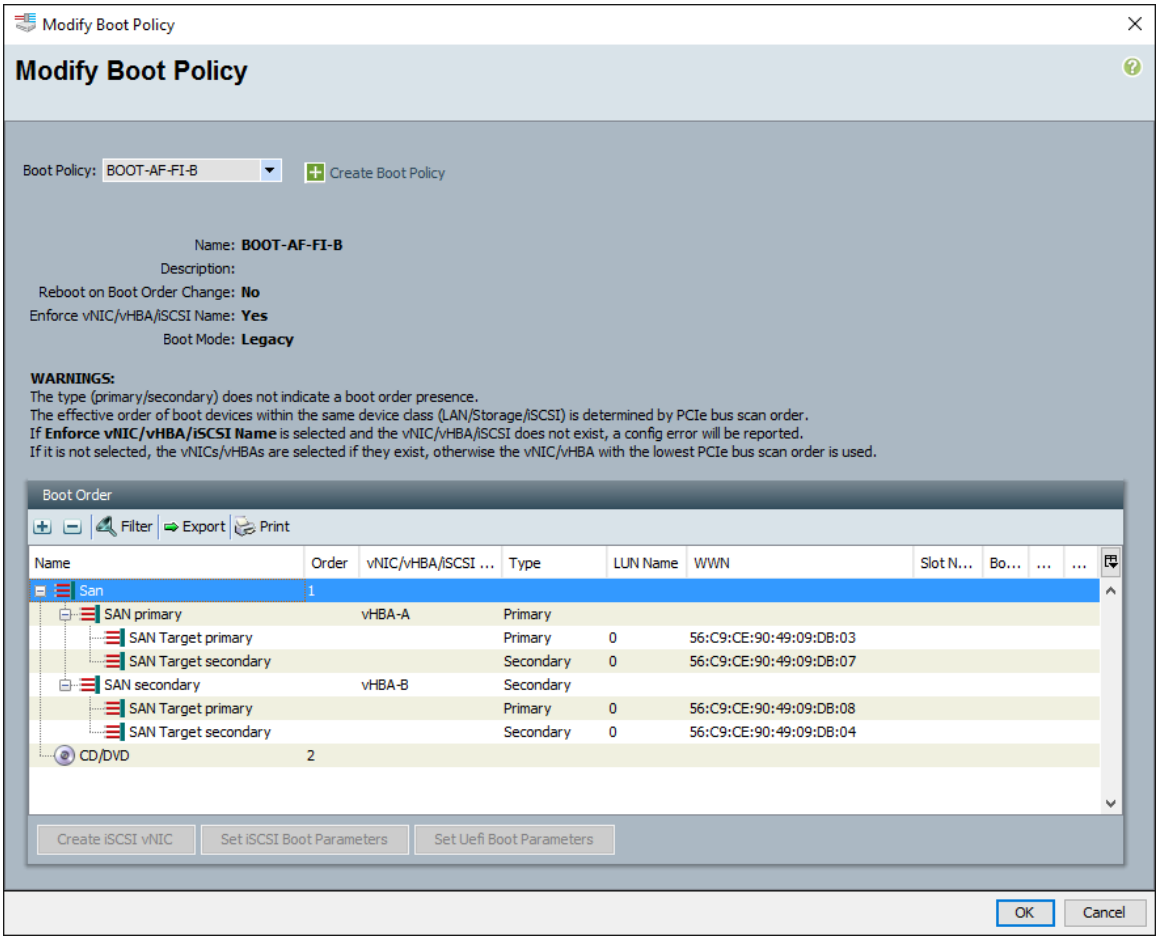

7. In the Boot Policy list, choose previously created Fabric B boot policy.
8. Click OK twice.
9. In the right pane, click the Network tab and click on Modify vNIC/vHBA Placement.
10. Select **`VM-Host-Infra`** and Expand vCon 1 and move `vHBA-B` ahead of `vHBA-A` in the placement order.

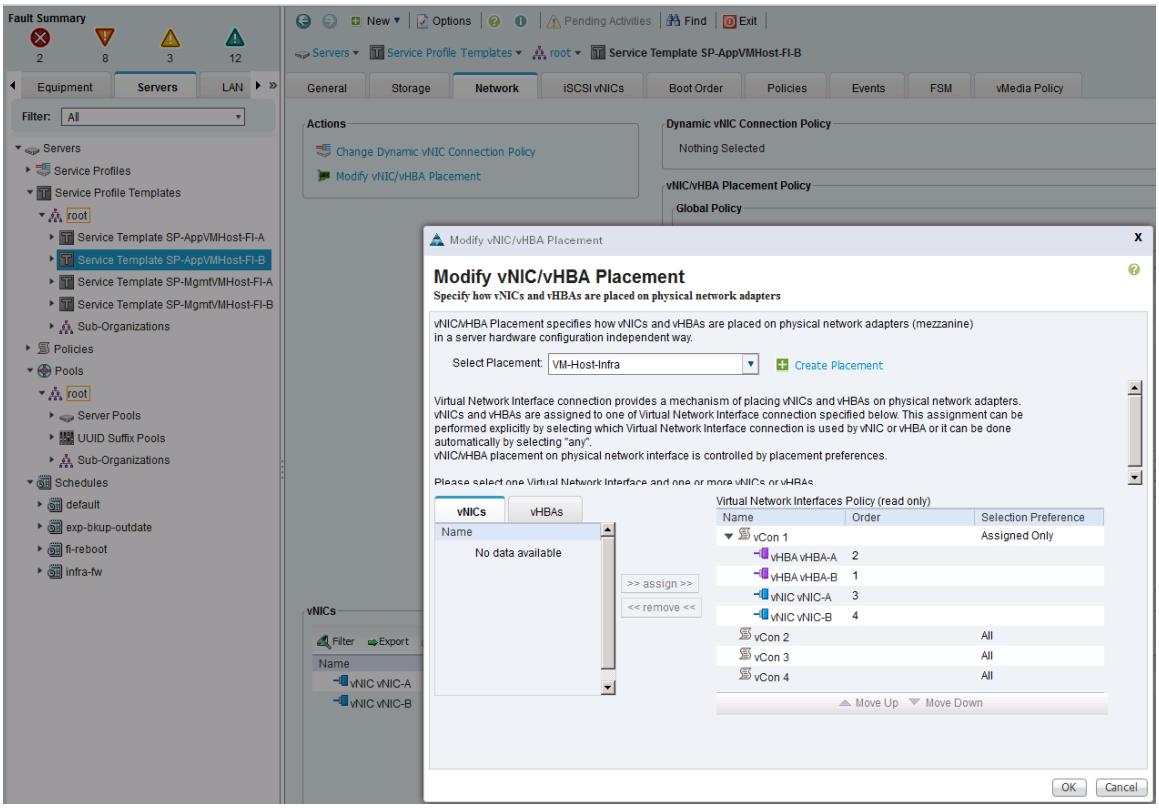

11. Click OK twice.
12. The Cisco UCS setup is complete at this point to deploy new hosts using service profiles generated from service profile templates created in earlier steps.

Optional: Deploy VMware vCenter 6.0 U1 Appliance

Skip this section if an existing VMware vCenter environment will be used to manage the SmartStack environment.

The procedures in this section are for installing VMware vCenter Server Appliance (VCSA) version 6.0U1 for managing the SmartStack environment. The following items are needed for the installation.

1. A Windows virtual machine or physical server to download the necessary files (VCSA, Client Integration Plug-In).
2. VMware vSphere client to access the Cisco UCS server where the vCenter VM will be deployed.
3. Web browser to access vCenter.

Please refer to VMware Knowledge Base 2110730 article for additional information on installing/upgrading a VMware vCenter Server Appliance (VCSA) 6.0 environment. The procedure in VCSA 6.0 is significantly different from other releases and uses an ISO file.

Download vCenter Server Appliance (VCSA) ISO from VMware

To download VMware vCenter Server Appliance, complete the following steps:

1. From a Windows machine, download VMware vCenter Server Appliance installer ISO from VMware Web site at https://my.vmware.com/web/vmware/downloads.
2. Using a VMware vSphere client, login to the host where the vCenter VM deployed.
3. Console into the host from vSphere and click 9th button (CD with a wrench) and choose Connect to ISO Image on Local Disk to mount the the VCSA Installer ISO.
4. Navigate to the VCSA ISO (VMware-VCSA-all) and click Open.

Install the Client Integration Plug-In

The Client Integration Plug-In must be installed before deploying the VCSA. Follow these steps to install the plugin.

1. On the Windows machine, go the VCSA installer directory and navigate to the VCSA directory and double-click VMware-ClientIntegrationPlugin-6.0.0.exe.
2. When the Client Integration Plug-in installation wizard comes up, click Next on the Welcome page.

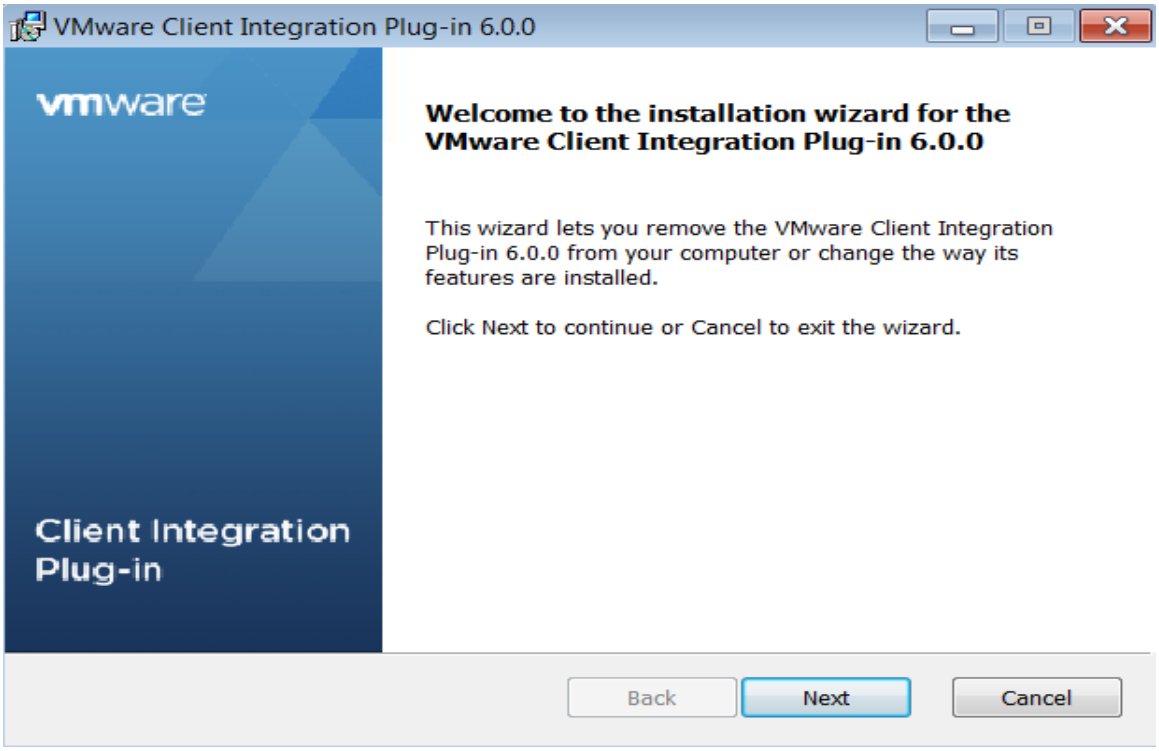

3. Read and accept the terms in the End-User License Agreement and click Next.
4. (Optional) Change the default path to the Client Integration Plug-in installation folder, and click Next.
5. On the Ready to Install the Plug-in page of the wizard, review the information and click Install.
6. After the installation completes, click Finish.

Install vCenter Server Appliance

To install vCenter 6.0 on the vCenter Server VM, complete the following steps:

1. From the VCSA installer directory, double-click vcsa-setup.html.
2. Allow the plug-in to run on the browser when prompted. You may need to wait up to three seconds for the browser to detect the Client Integration Plug-in and
3. On the Home page, click Install to start the vCenter Server Appliance deployment wizard.
4. Read and accept the End User License Agreement, and click Next.
5. On the Connect to the target server screen, provide the IP address, username and password for the host on which you want to deploy the vCenter Server Appliance appliance, and click Next.

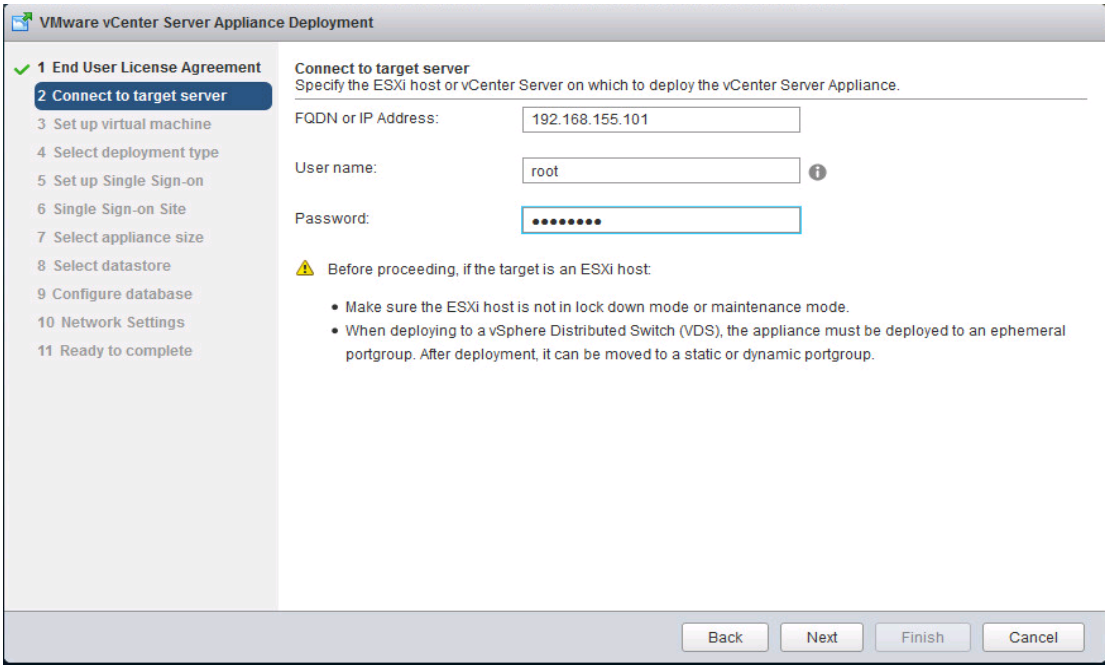

6. (Optional) Accept the certificate warning, if any, by clicking Yes.
7. On the Set up virtual machine screen, enter the vCenter Server Appliance name, set the password for the root user, and click Next.

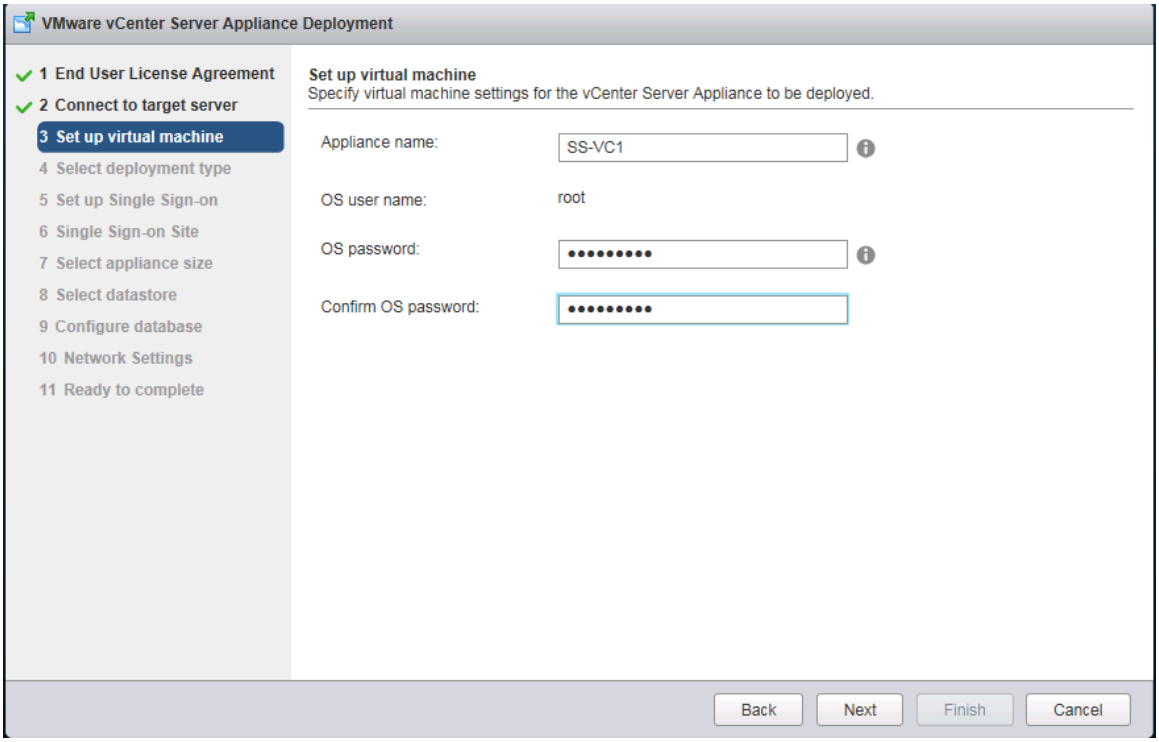

8. In the Select deployment type screen, select Install vCenter Server with an embedded Platform Serv2ices Controller and click Next.

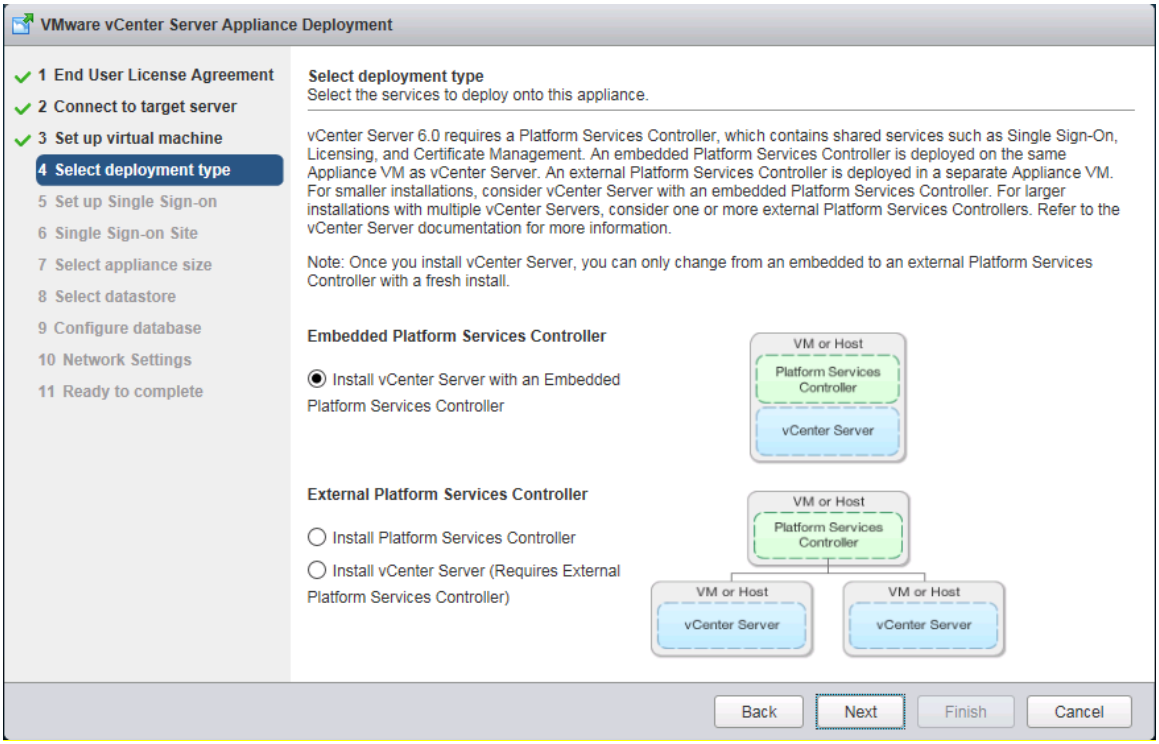

9. In the Setup Single Sign-on screen, create a new vCenter Single Sign-On domain, and click Next.

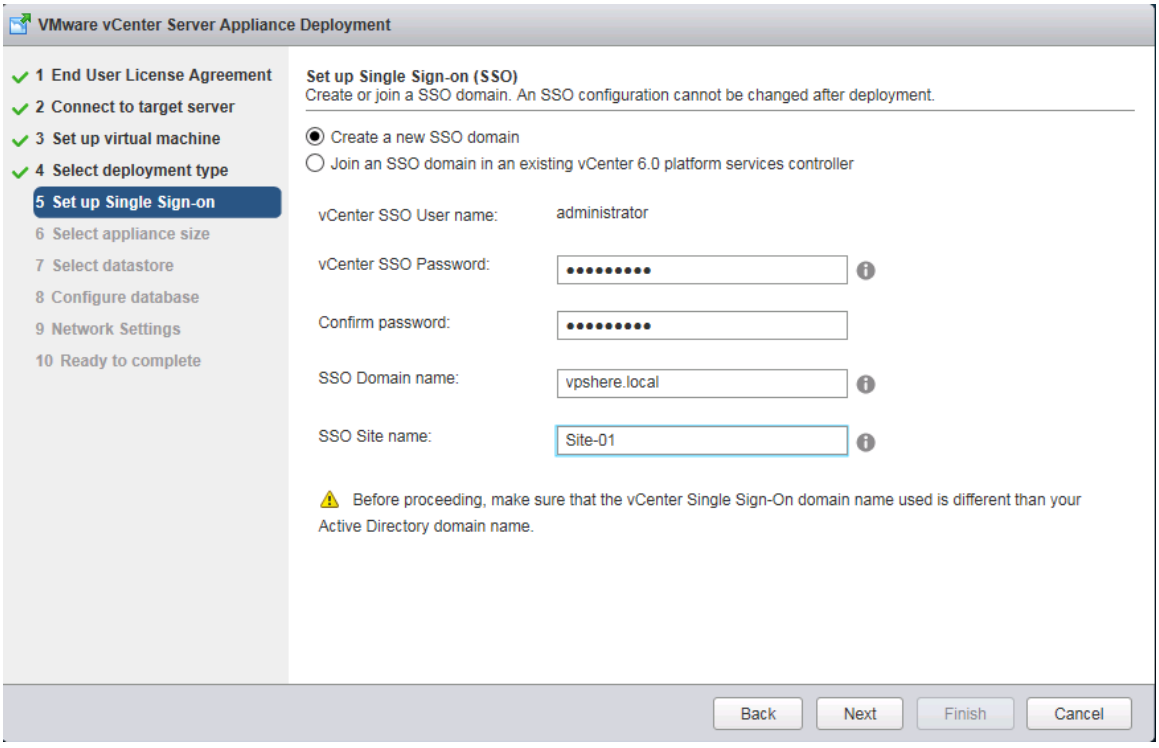

10. In the Select appliance size screen, select the size that matches your deployment, and click Next.

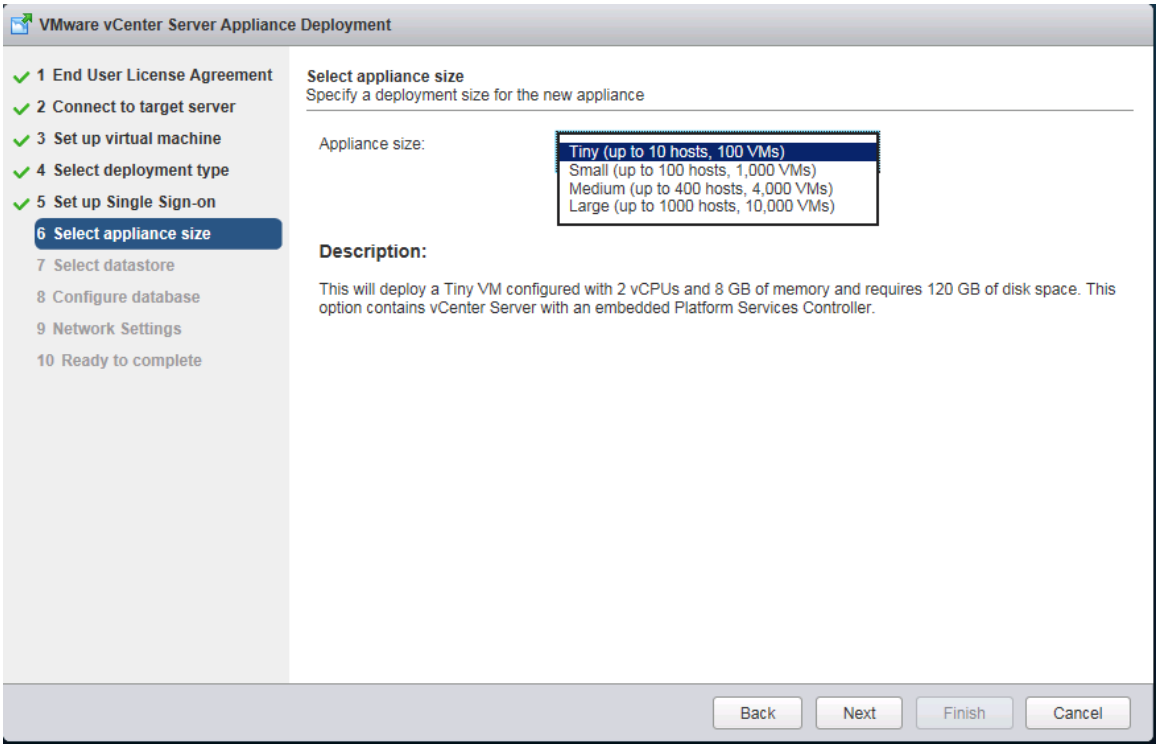

11. In the Select datastore screen, select the location for the VM configuration and virtual disks should be stored (DS-Infra), and click Next.

12. In the Configure database screen, select the database to use, and click Next.

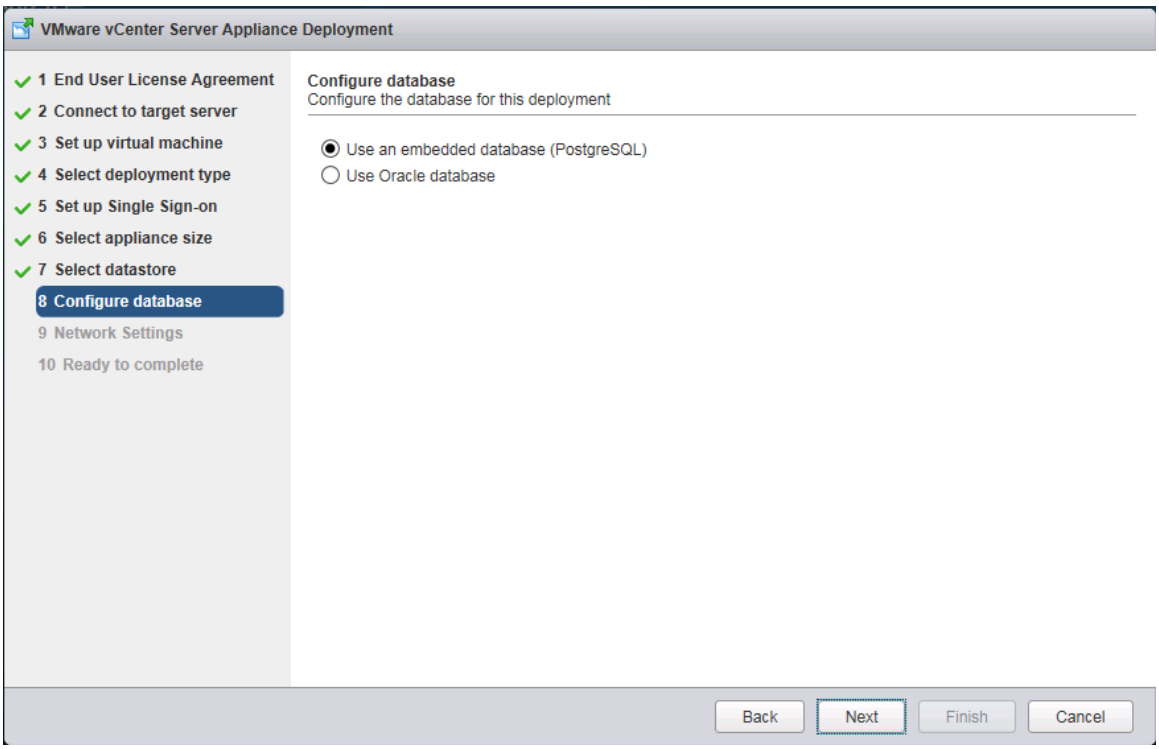

13. In the Network Settings screen, provide the VLAN network, IP address, Mask, Gateway, DNS and NTP info the vCenter Appliance should use. Optionally, you can also enable SSH to the VM. Click Next.

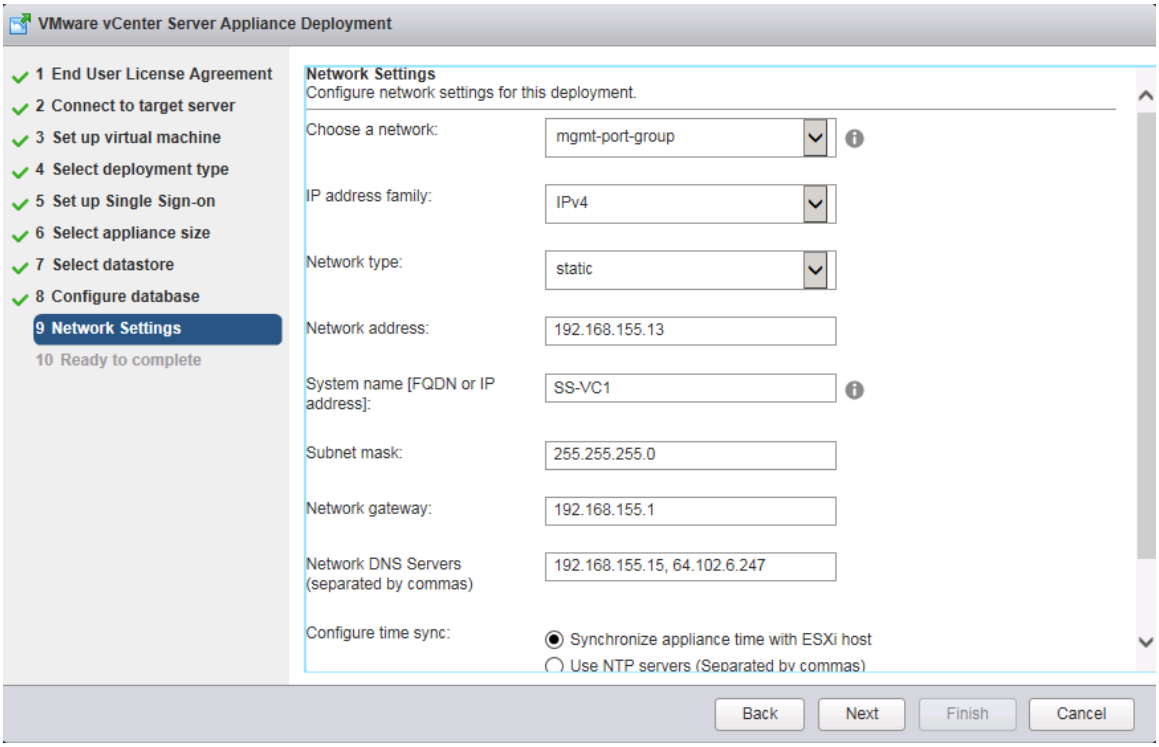

14. In the Ready to complete screen, review the deployment settings for the vCenter Server Appliance, and click Finish to complete the deployment process.

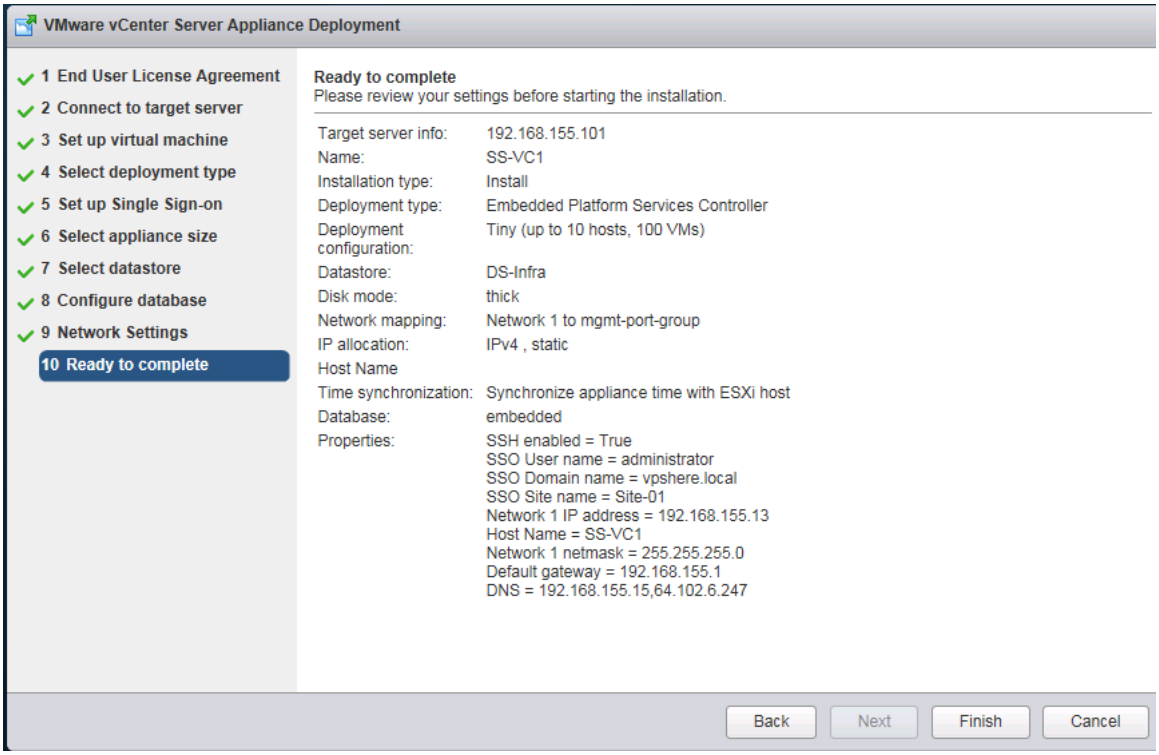

15. The vCenter appliance installation will take few minutes to complete. When complete, click Close to exit the wizard.

Log into vSphere Web Client

1. From the VMware vCenter Installer window, download
2. Using a web browser, navigate to https://192.168.155.13:9443/vsphere-client/
3. Enter the username (administrator@vsphere.local) and associated password.
4. Click Login. Now you can do some basic setup of the VCSA using the web client.

Join Active Directory Domain

1. Login in the VCSA using the vSphere Web client using administrator@vsphere.local account or another account
2. Before joining the Active Directory Domain, verify that the account (administrator@vsphere.local) used to log in to the vCenter instance is part of the Single Sign-on Administrators group. From the Navigator window, go to Administration > Single Sign-on > Users and Groups > Group Members. If the user is not a group member, add the user using the + icon below Group Members.

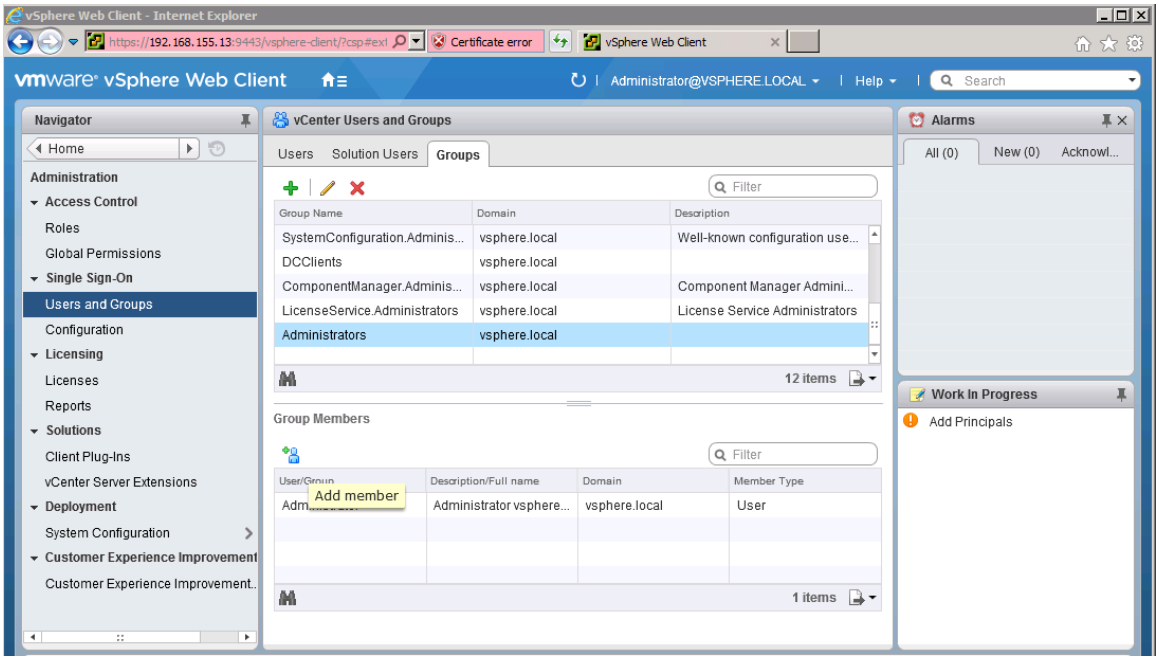

3. Now navigate to Administration > System Configuration > Nodes
4. Under Nodes, select a vCenter instance/node (192.168.155.13) and go to Manage tab
5. Click Settings and select Advanced > Active Directory. Click Join. Provide the domain and account info in the Join Active Directory window and click OK.

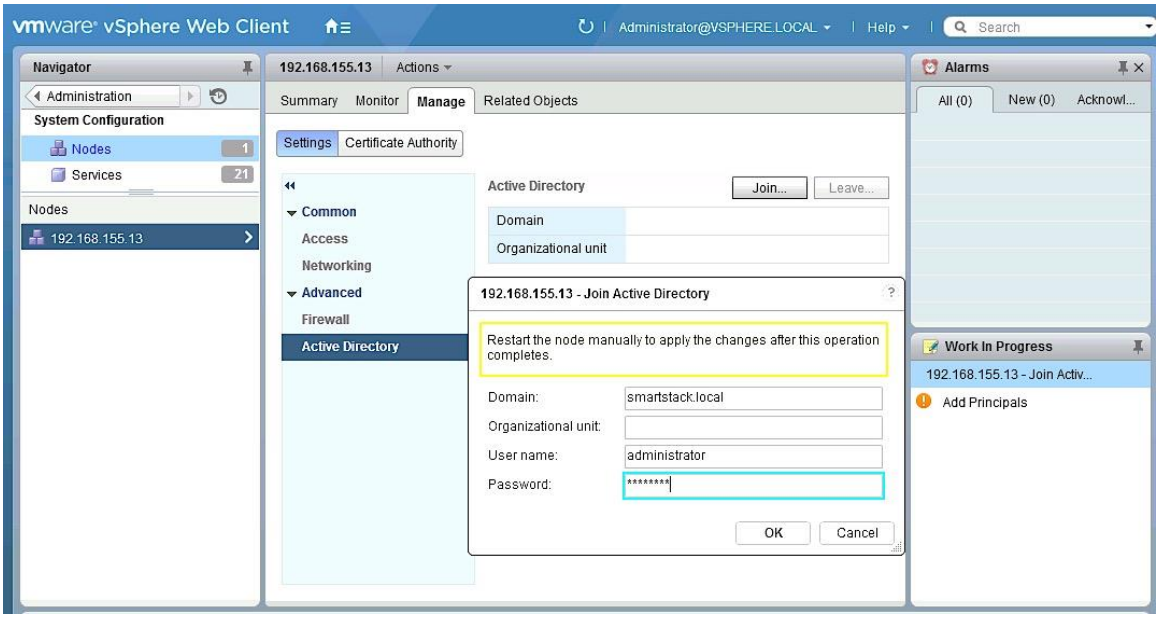

6. Right-click on the node instance (192.168.155.13) and select Reboot to restart the vCenter instance. If the changes took effect, the Join button will be greyed out and the Leave button will become available.

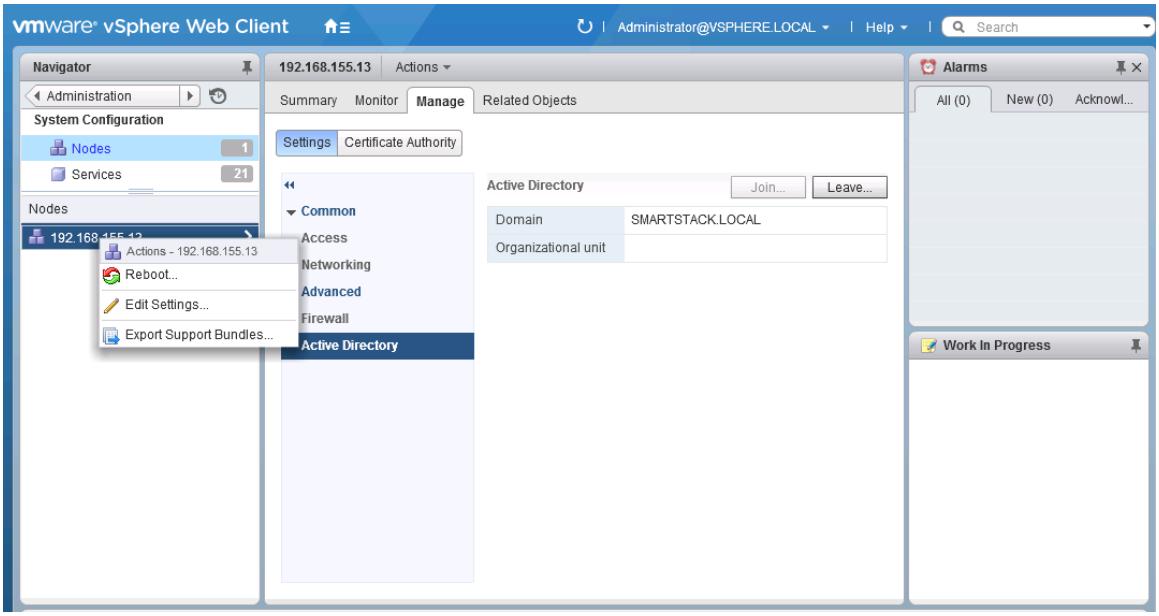

7. Navigate to Administration > Single Sign-On > Configuration.
8. On the Identity Sources tab, click the + to add Identity Source.
9. Select Active Directory (Integrated Windows Authentication) and enter the identity source settings to join the Active Directory domain. Click OK.

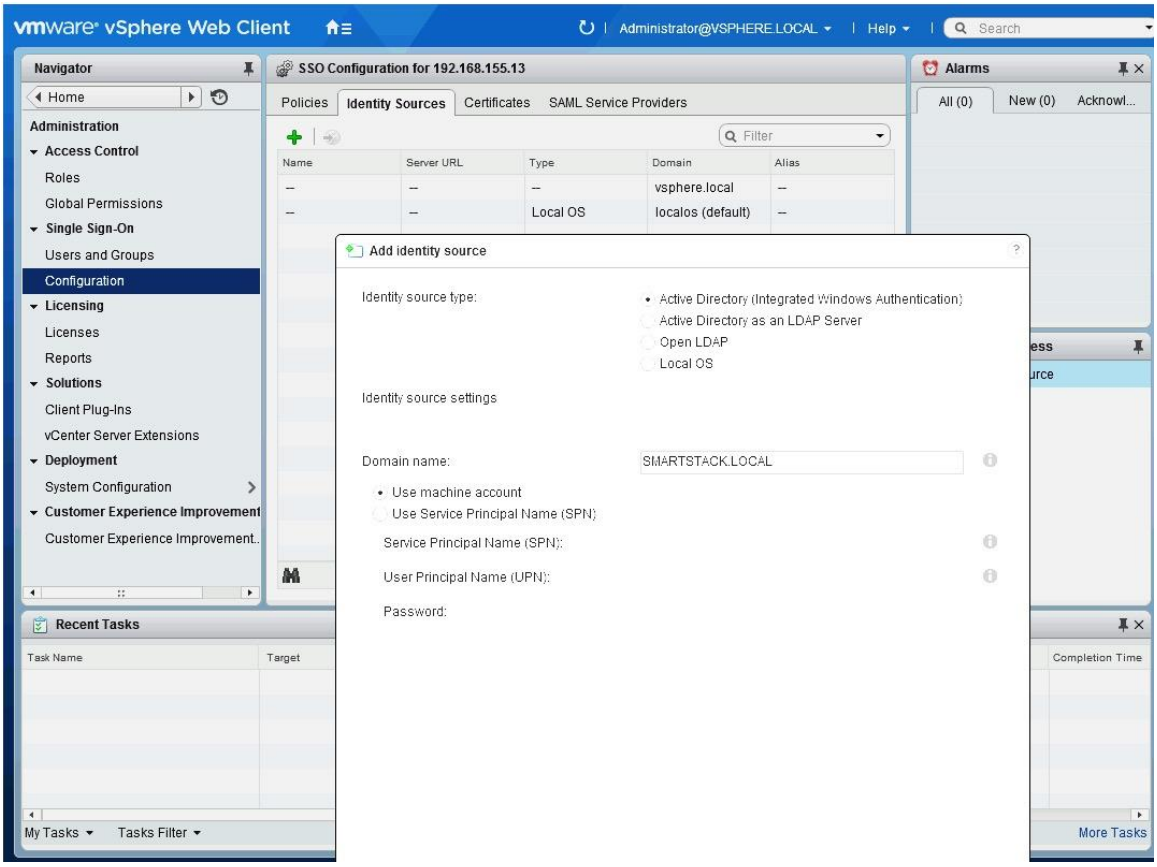

10. You should now see the joined Active Directory domain (smartstack.local).

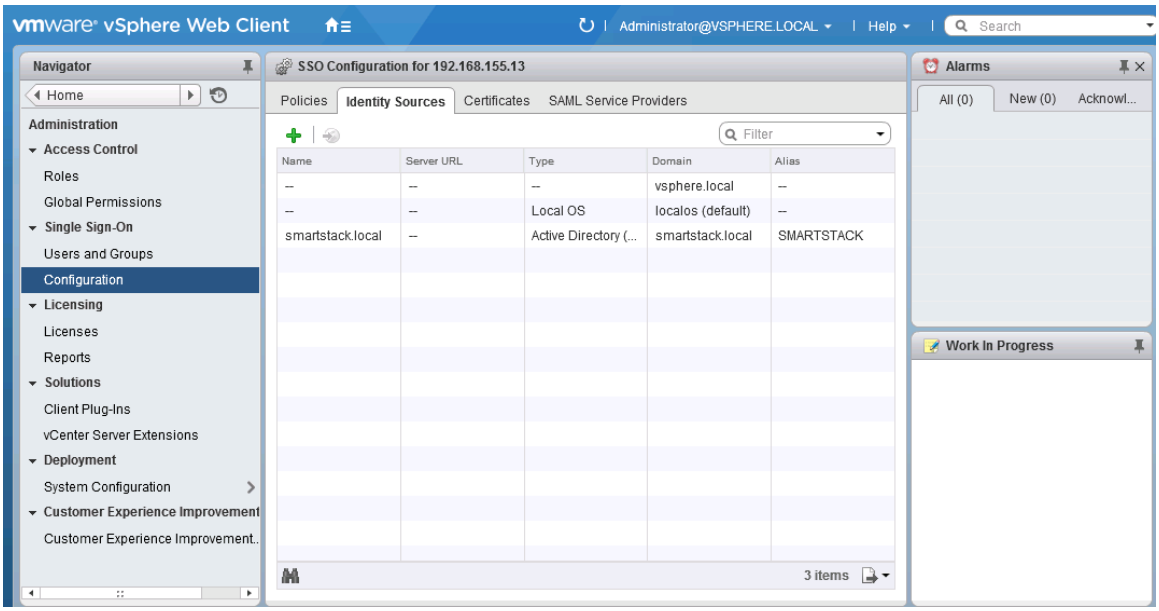

Setup vCenter for Datacenter, Cluster, DRS and HA

1. In the vSphere Web Client, navigate to the vCenter Inventory Lists > Resources > vCenter Servers.
2. Select vCenter instance (192.168.155.13).
3. Go to Actions in the toolbar and select New Datacenter from the drop-down.

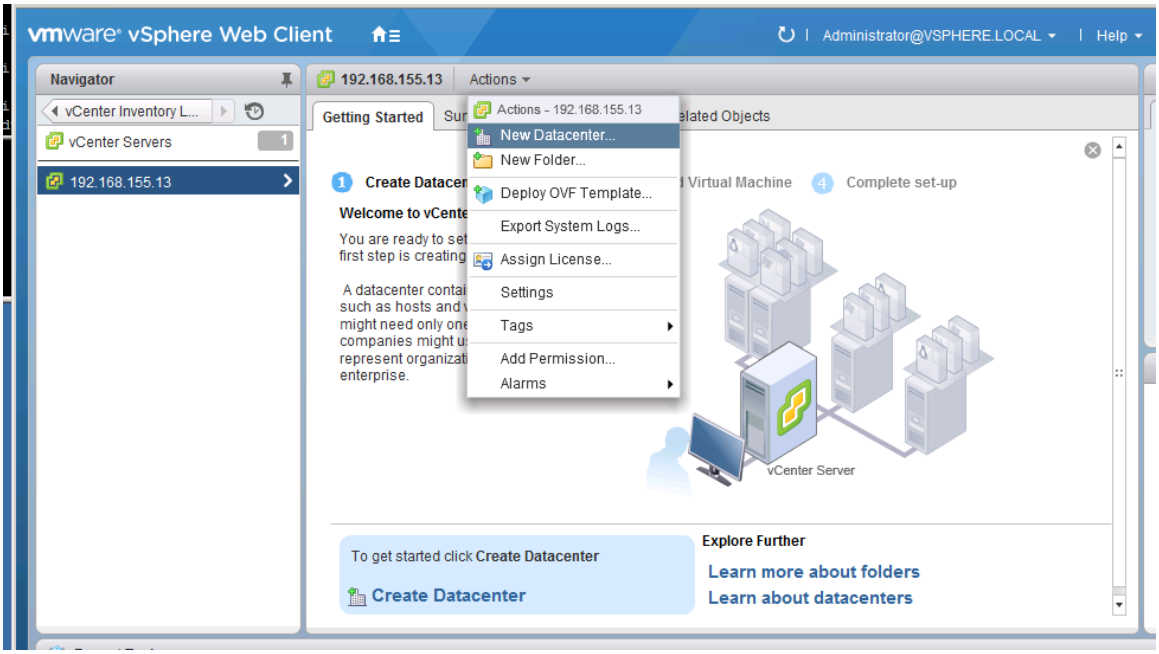

4. Rename the datacenter (for example, SS-DC-1). And then, click OK.
5. Go to Actions in the toolbar and select New Cluster from the drop-down.

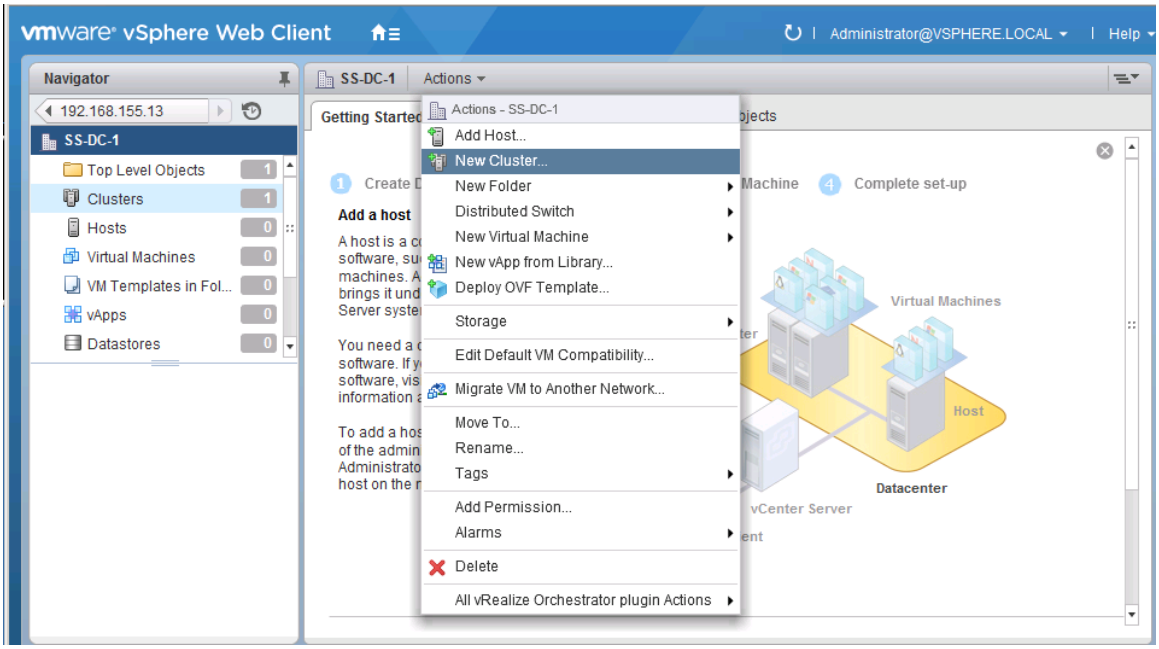

6. In the New Cluster window, provide a cluster name, enable DRS, vSphere HA and Host monitoring. Click OK.

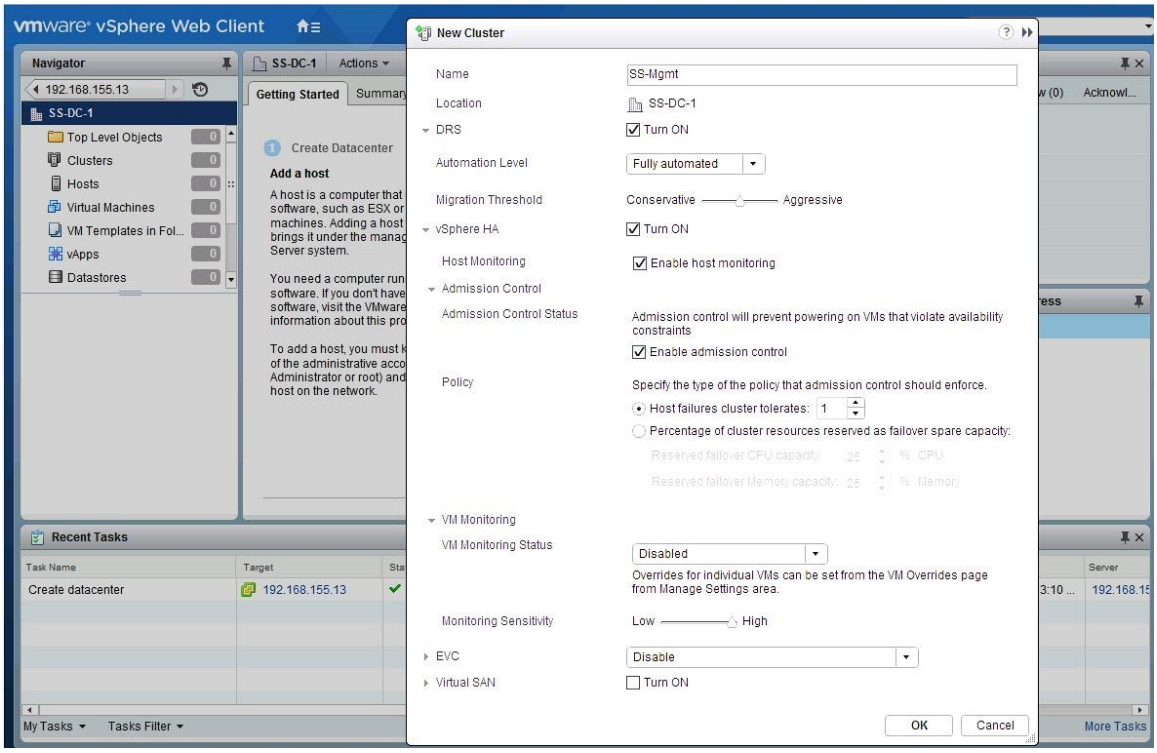

If mixing Cisco UCS M3 or M4 series servers within a vCenter cluster, it is necessary to enable EVC mode. For more details on the EVC mode, refer to the Enhanced vMotion Compatibility (EVC) Processor documentation.

Setup ESXi Dump Collector

ESXi Dump Collector is installed by default in the vCenter Server Appliance but some setup is required as outlined in the next few steps.

1. From the vSphere Web Client Home page, navigate to Administration > System Configuration > Nodes and select a node/vCenter instance (192.168.115.13) from the list.
2. Select the Related Objects tab to see the list of services running on the vCenter instance. Select the VMware vSphere ESXi Dump Collector from the list and right-click to Start the service.

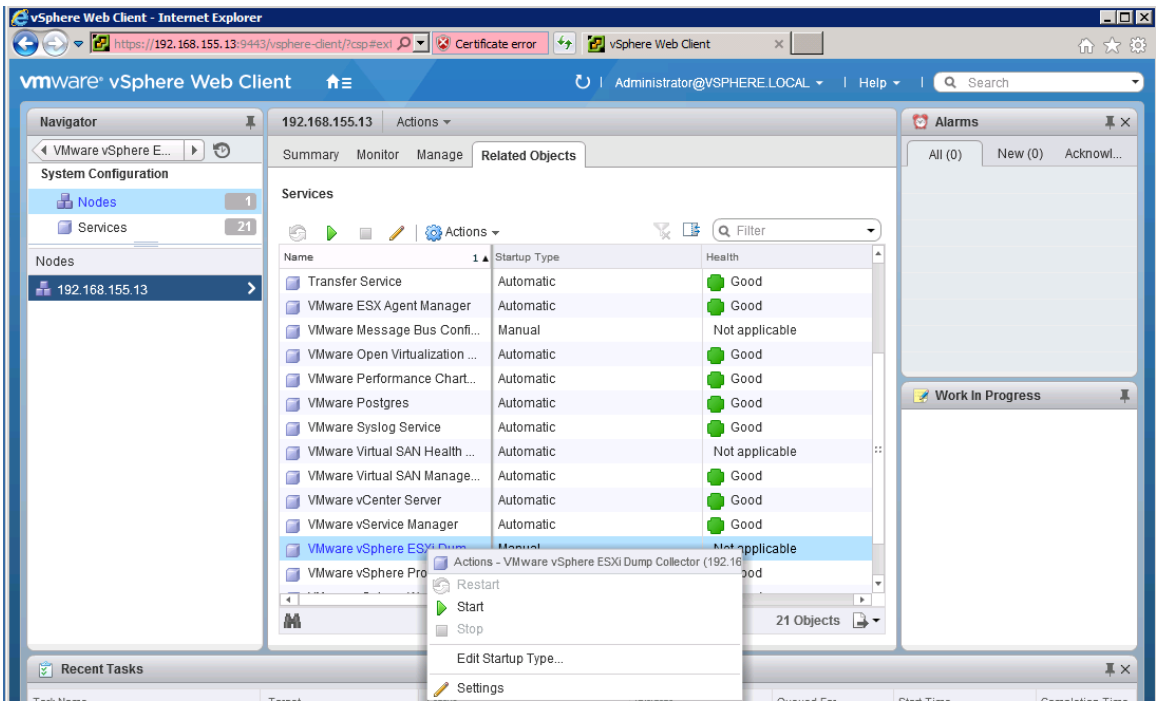

3. Right-click on the service again to Edit Startup Type to Automatic.

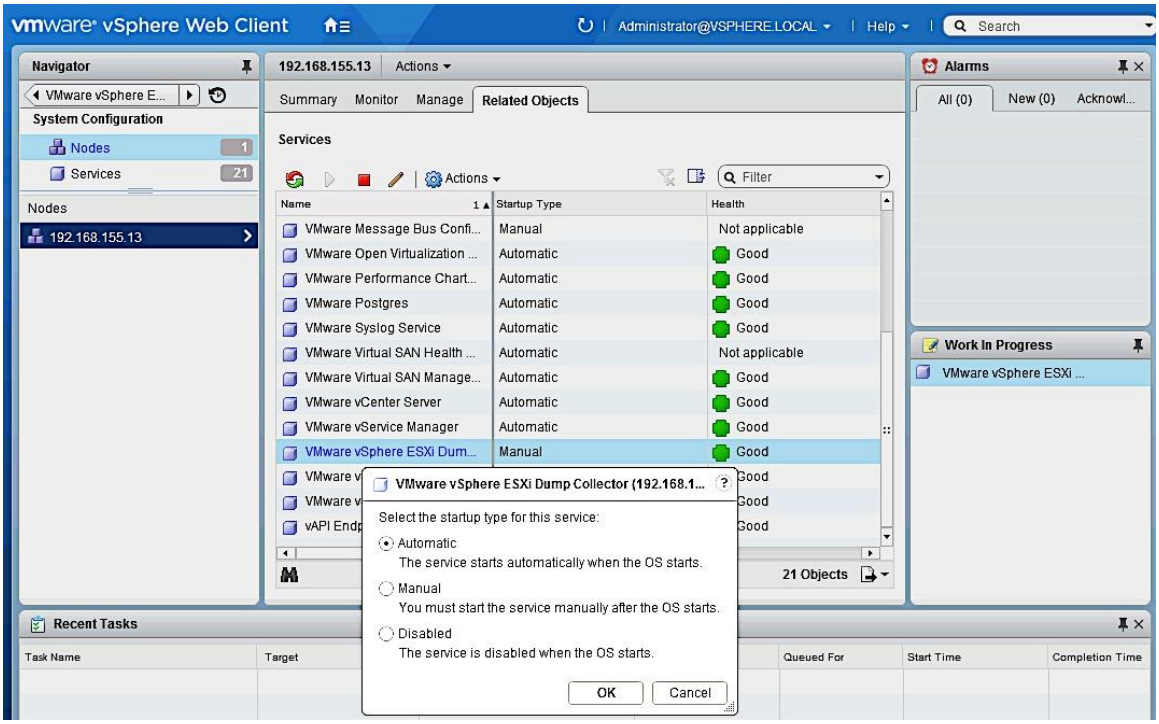

Deploy New Server

This section covers the setup procedures for deploying a new server to the SmartStack environment.

Figure 16 Workflow for Deploying New Hosts

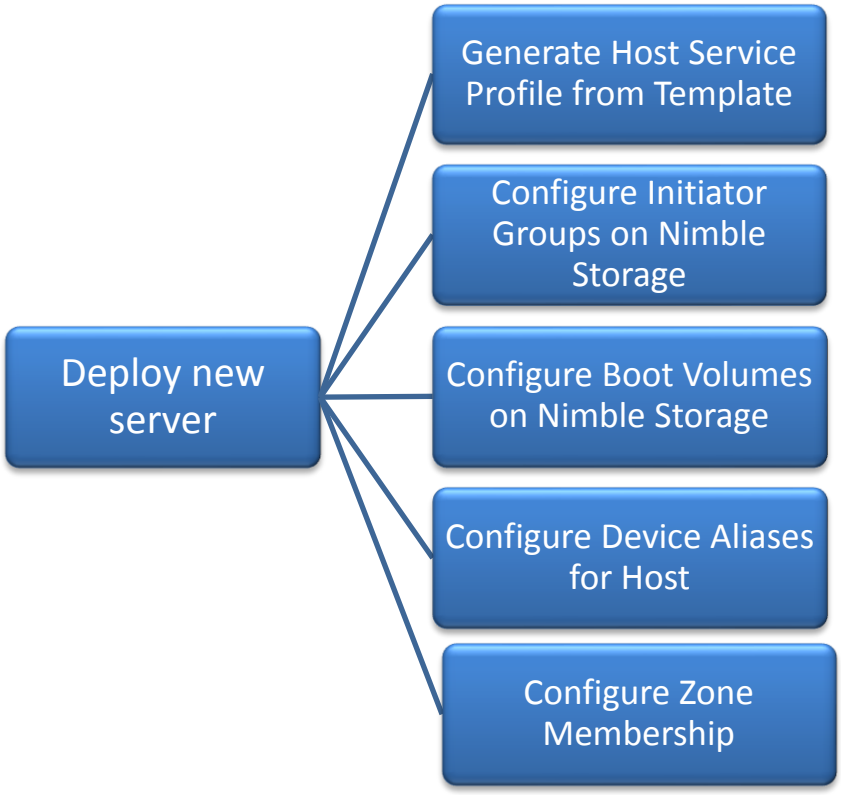

Generate Service Profile for the New Host using a Service Profile Template

Using the service profile template created earlier; generate a service profile for the newly deployed host.

1. From Cisco UCS Manager, click Servers tab in the navigation pane.
2. Select Servers > Service Profile Templates > root.
3. Select the Service Profile Template created in earlier steps (for example, `SP-AppVMHost-AF-FI-A`). Right-click on the template and select Create Service Profiles from Template.

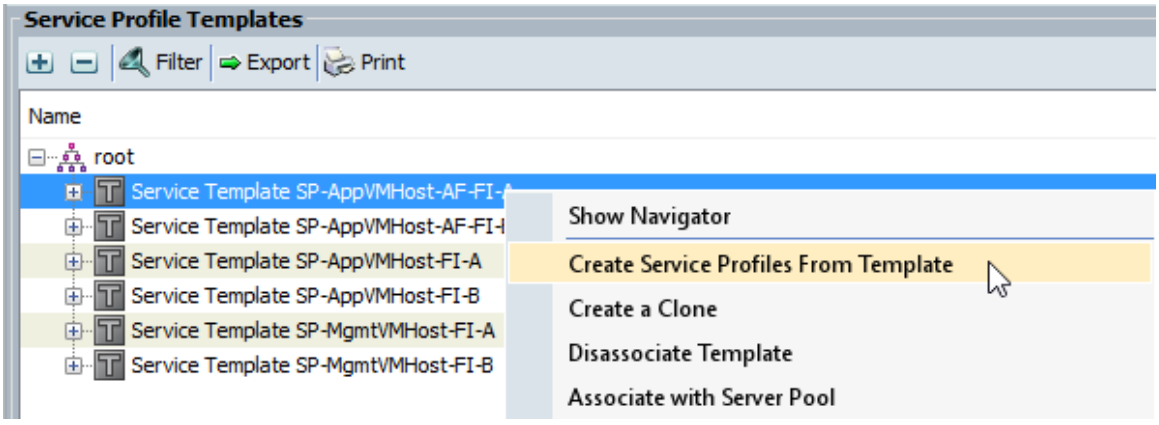

4. Enter the Naming Prefix, the Suffix Starting Number and Number of service profiles instances to create and click OK twice to complete.

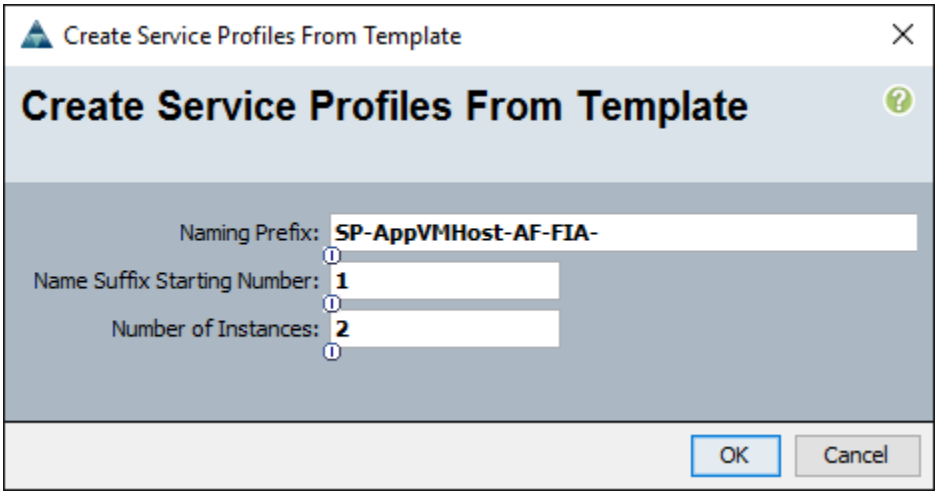

5. For the next host, repeat above steps to create a service profile that will use Fabric B as primary path for traffic. Click OK twice to complete.
6. Verify that the service profile was created and associated. The newly created service profile is automatically associated with the server if the server is part of the server pool defined in the template.
7. Create service profiles for all the servers in the server pool. The figure below shows the servers used for validation in this SmartStack design.

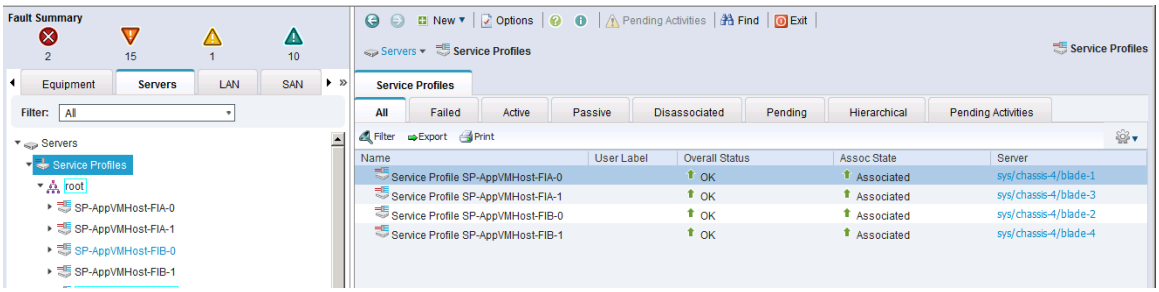

8. (Optional) Choose each newly created service profile and enter the server host name or the FQDN (for example, `CVDAFA-Host1`) in the User Label field in the General tab. Click Save Changes to map the server host name to the service profile name.

Setup Nimble Storage Array for SAN Boot of Hosts

This section provides instruction on how to create Nimble Storage volumes and initiator groups for FC SANboot operation.

Identify Initiator WWPNs from the Host Service Profile

1. Login to Cisco UCSM using a web browser – use the Cisco FI cluster IP address and admin username and password to log in.
2. Navigate to the Servers tab and then select the Service Profile in question.
3. Expand the vHBA section to determine the WWPNs are for both vHBA-A and vHBA-B.

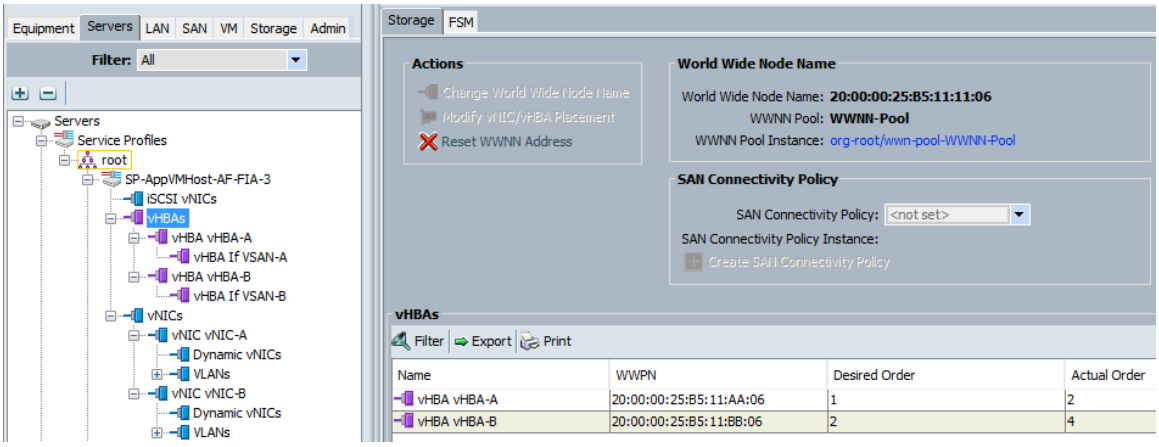

Create Initiator Group

1. Login to the Nimble Storage GUI using a web browser.
2. Navigate to Manage > Initiator Groups.
3. Select Create.
4. Fill out the form as indicated below. Note that the WWPNs come from the step 3.

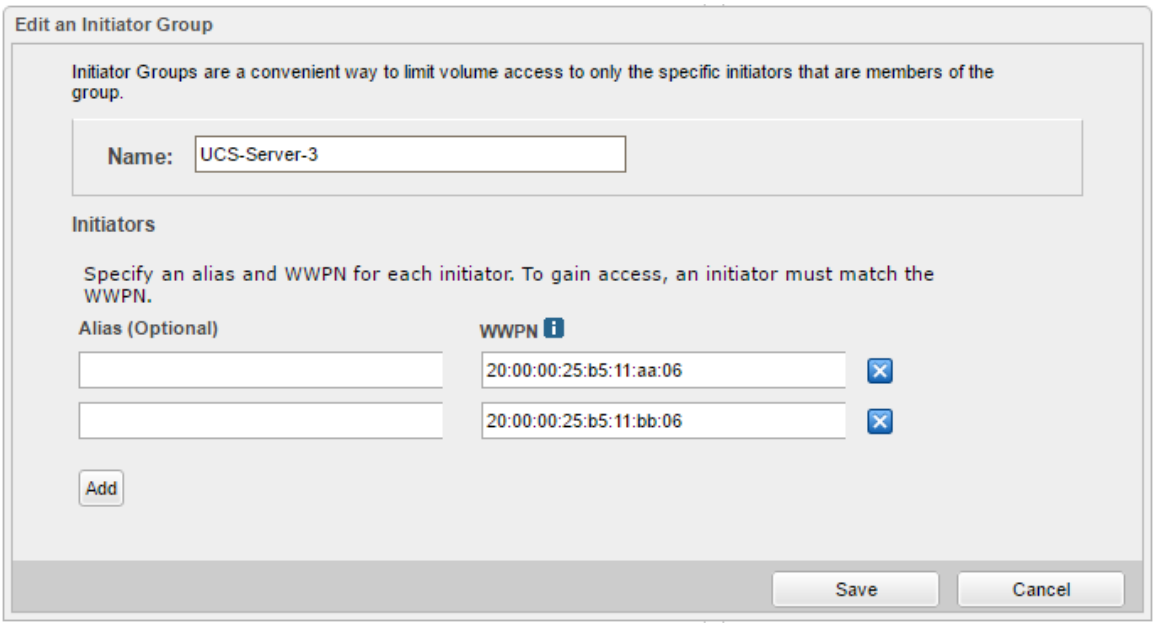

Create Boot Volumes for Host

1. Login to the Nimble Storage GUI using a web browser. Navigate to Manage > Volumes.
2. Select New Volume.
3. Fill out the form as directed below. Be sure to select the name of the initiator group that you created in step 6. Also note that the LUN ID will default to 0 which matches the boot policy for this Service Profile.

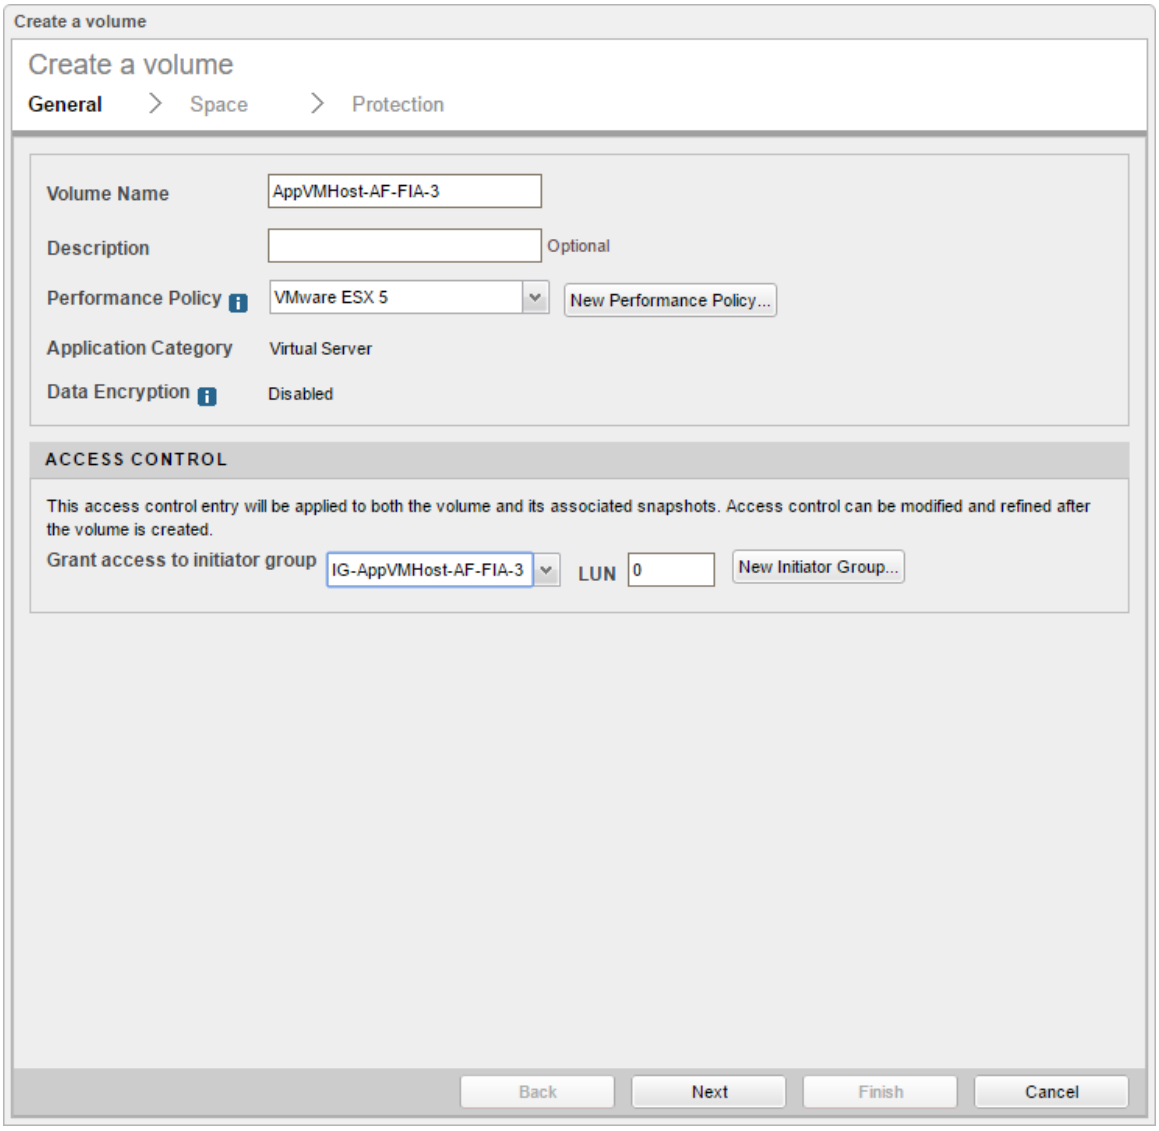

4. Repeat these steps for every new host deployed from a Service Profile. Note, that each initiator group and each volume will be unique.

Add New Hosts to Cisco MDS VSANs

Add the WWPNs of the newly deployed hosts to the VSANs of both Cisco MDS switches. Repeat these steps for each new host deployed.

Configure Device Aliases for the New Host

Cisco MDS-A

Configure the device-aliases for the above WWPNs above and commit it as follows.

```
device-alias confirm-commit enable
device-alias database
device-alias name AFA-AppVMHost-FIA-3 pwwn 20:00:00:25:b5:11:aa:06
device-alias commit
```

Cisco MDS-B

Configure the device-aliases for the above WWPNs above and commit it as follows.

```
device-alias confirm-commit enable
device-alias database
device-alias name AFA-AppVMHost-FIB-3 pwwn 20:00:00:25:b5:11:bb:06
device-alias commit
```

Create Zones and Zoneset for the New Host

Cisco MDS-A

1. Create Host Zones for each host.

```
zone name AFA-AppVMHost-FIA-3-A vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:06
    member pwwn 56:c9:ce:90:49:09:db:01
    member pwwn 56:c9:ce:90:49:09:db:03
    member pwwn 56:c9:ce:90:49:09:db:05
    member pwwn 56:c9:ce:90:49:09:db:07
```

2. Create Zoneset and Add Zones to it.

```
zoneset name Fabric-A vsan 4091
   member AFA-AppVMHost-FIA-3-A
```

3. Activate zoneset and enable distribution.

```
zoneset activate name Fabric-A vsan 4091
```

4. Save the configuration.

```
copy run start
```

Cisco MDS-B

1. Create Host Zones.

```
zone name AFA-AppVMHost-FIA-3-B vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:06
    member pwwn 56:c9:ce:90:49:09:db:02
    member pwwn 56:c9:ce:90:49:09:db:04
    member pwwn 56:c9:ce:90:49:09:db:06
    member pwwn 56:c9:ce:90:49:09:db:08
```

2. Create Zoneset and Add Zones to it.

```
zoneset name Fabric-B vsan 4092
    member AFA-AppVMHost-FIA-3-B
```

3. Activate zoneset and enable distribution.

```
zoneset activate name Fabric-B vsan 4092
```

4. Save the configuration.

```
copy run start
```

Setup Hosts to SAN Boot ESXi 6.0U2

This section provides detailed instructions for installing VMware ESXi 6.0 Update 2 in a SmartStack environment. After the procedures are completed, the ESXi hosts will be provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in Keyboard, Video, Mouse (KVM) console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot volumes. A Cisco custom ESXi 6.0 ISO file is first downloaded from VMware and positioned on the Windows machine used to KVM console into the Cisco UCS server. The custom image is used in the SmartStack deployment as it contains Cisco drivers. These drivers may need to be upgraded but using the custom image ensures a minimum supported version.

KVM Console into Host from Cisco UCSM Web Interface

1. Login to Cisco UCSM using a web browser using the IP address of the Cisco UCS FI cluster. Login using administrator (for example, `admin`) account and password.
2. Download the Cisco Custom ISO for ESXi from the VMware website.
3. From the main Cisco UCSM menu, click Servers tab.
4. Select Servers > Service Profiles > root > **`SP-AppVMHost-AF-FIA-3`**.
5. Right-click **`SP-AppVMHost-AF-FIA-3`** and select KVM Console.

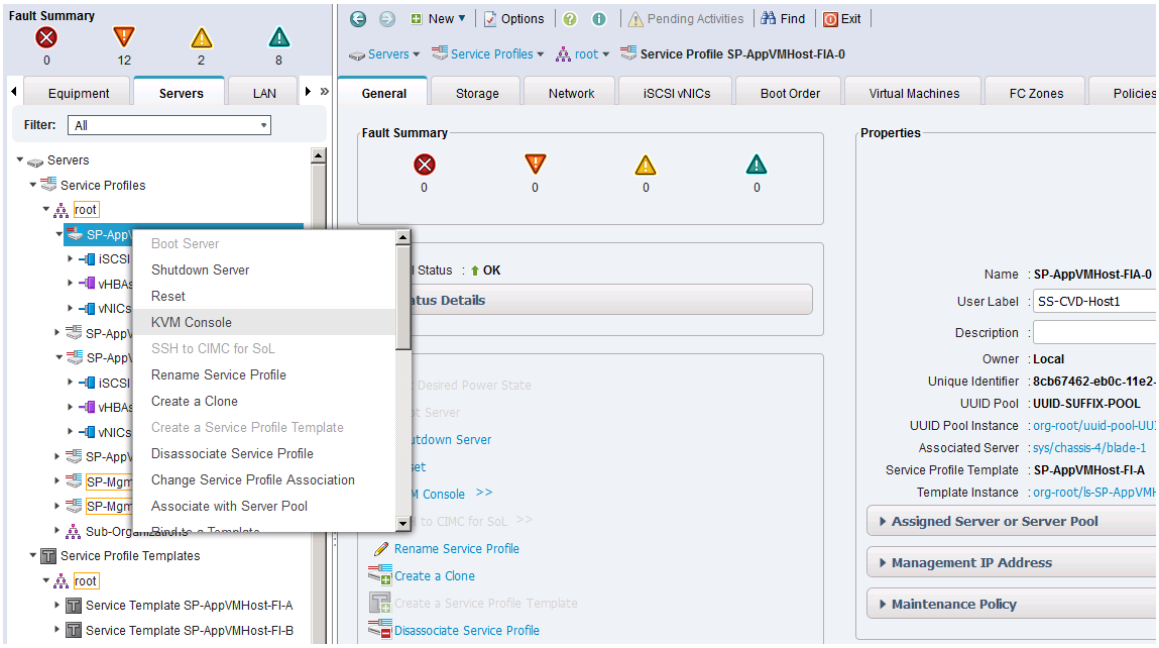

Prepare Host for ESXi Install

From the KVM window, select Virtual Media from the top menu. Select Activate Virtual Devices. In the Virtual Media Session window, select Accept this Session and click Apply.

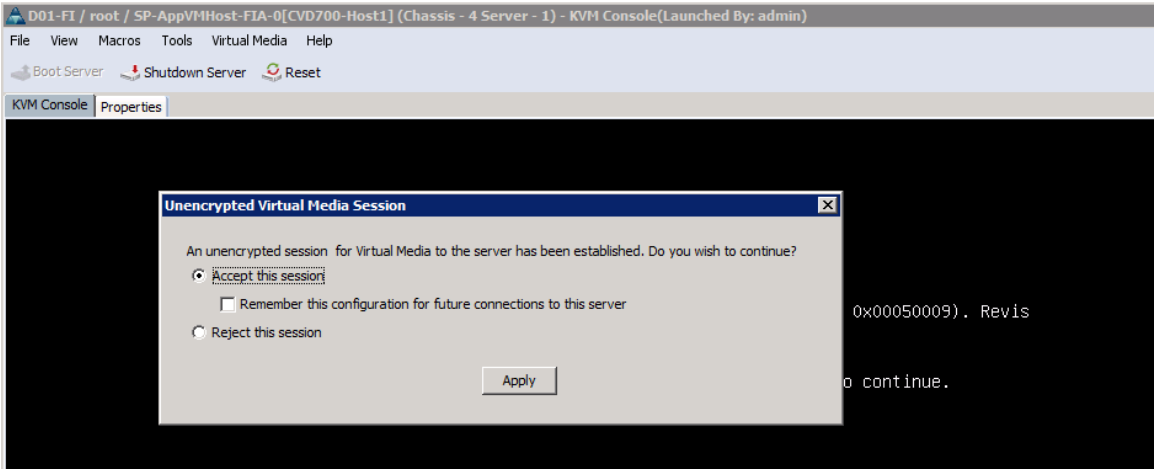

6. Select Virtual Media from the top menu again and pick the Map CD/DVD option. In the Virtual Media – Map CD/DVD window, click on Browse to select a Drive/Image File. Select the custom Cisco ESXi ISO file, previously downloaded from vmware.com
7. Click the Map Device button to map the Custom Cisco ESXi ISO image.

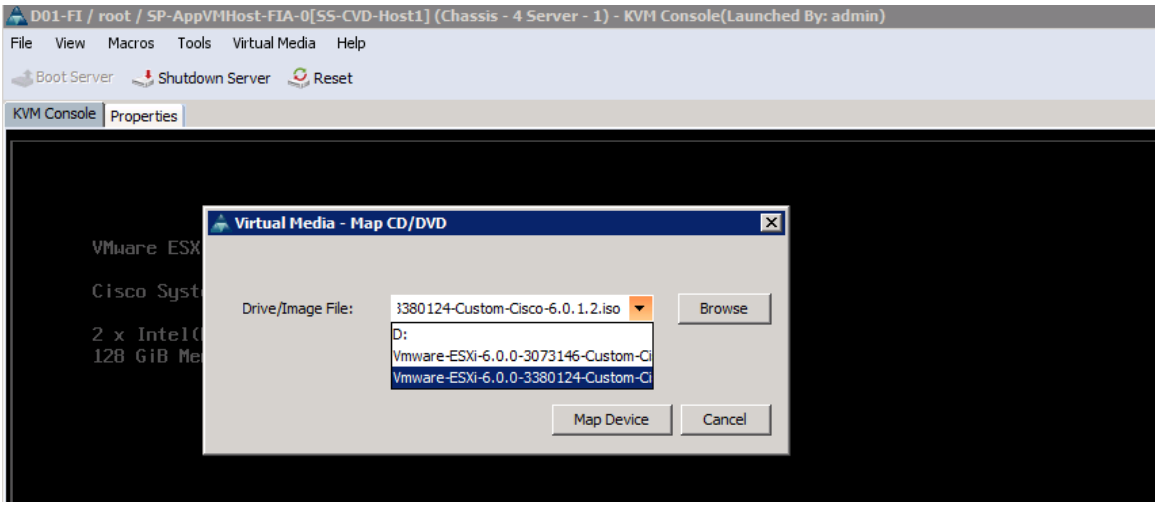

8. Click on Reset from the menu in the KVM console window to allow a power cycle and click OK 3 times. The boot process should now be visible on the KVM console window.

Install ESXi

1. When the server boots, the host detects the presence of the ESXi installation media and loads the ESXi installer as shown below. Press Enter to Continue.

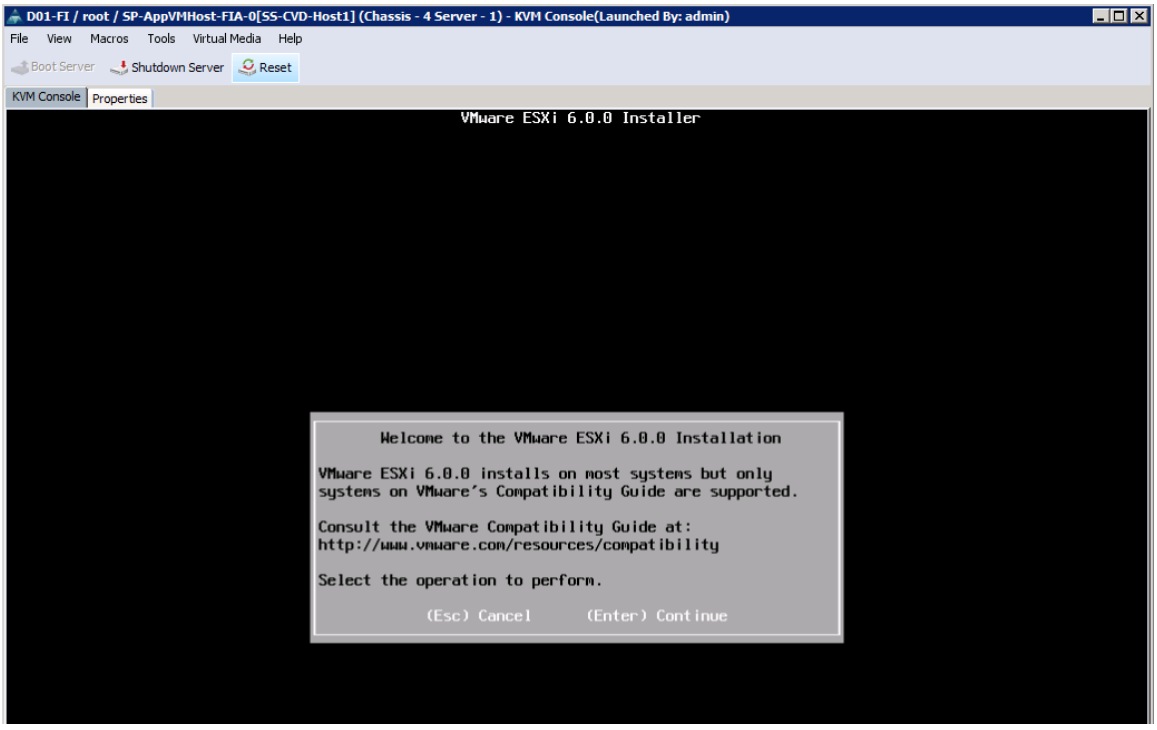

2. Read and accept the end-user license agreement (EULA). Press F11 to Accept and Continue.

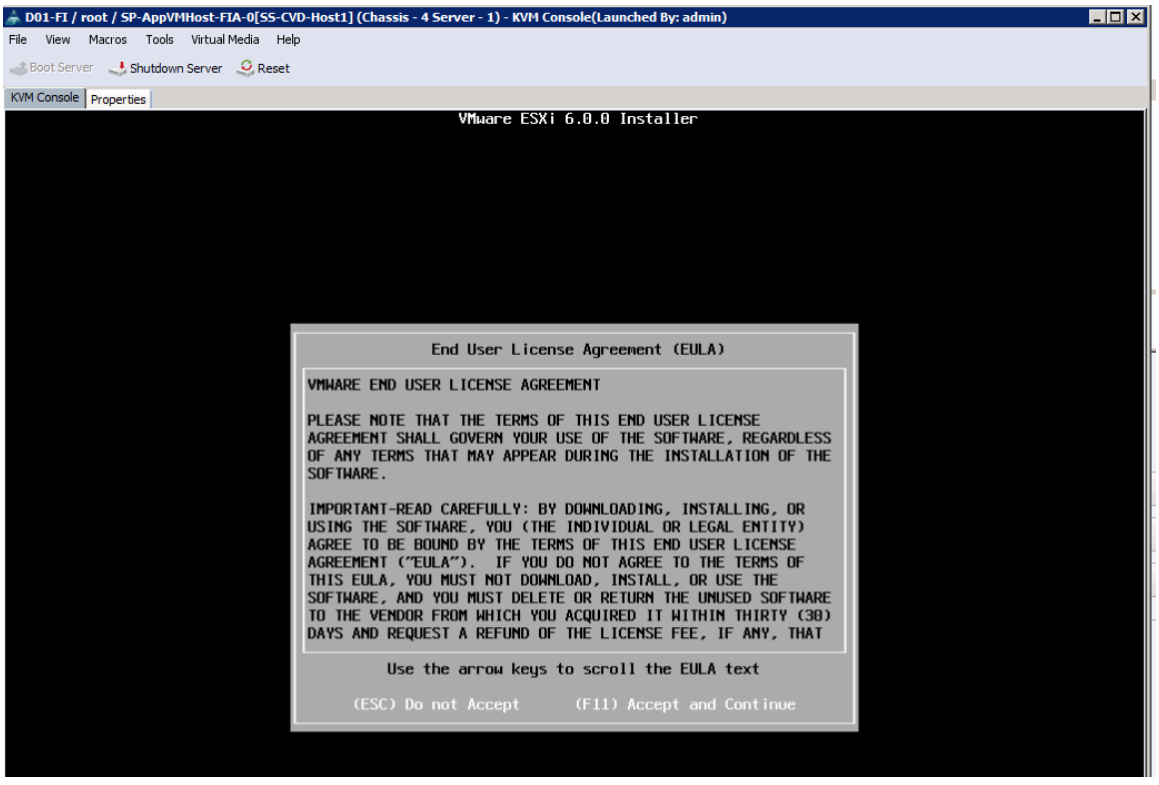

3. Select the Nimble Boot Volume that was previously created for the host. Press Enter to Continue with installation.

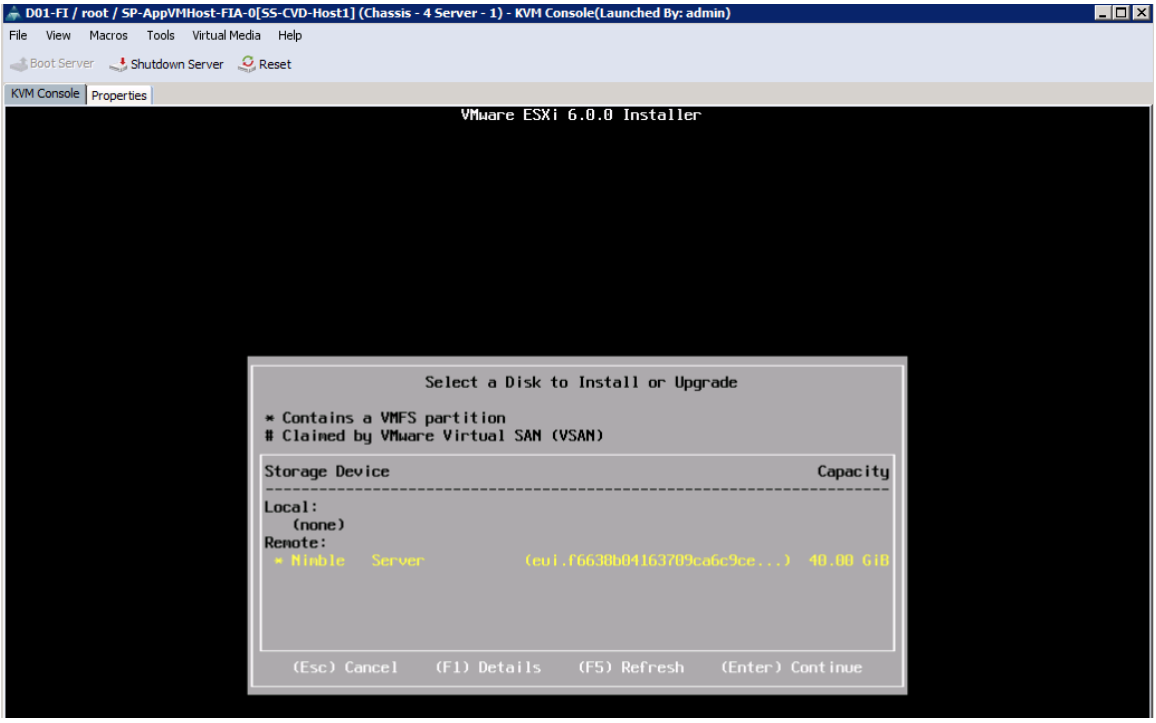

4. When prompted, select the appropriate keyboard layout (for example, US Default) in the pop-up window and press Enter to Continue.
5. When prompted, enter the root password for the host in the pop-up window and press Enter to Continue.
6. In the Confirm Install pop-up, Press F11 to Install and continue.

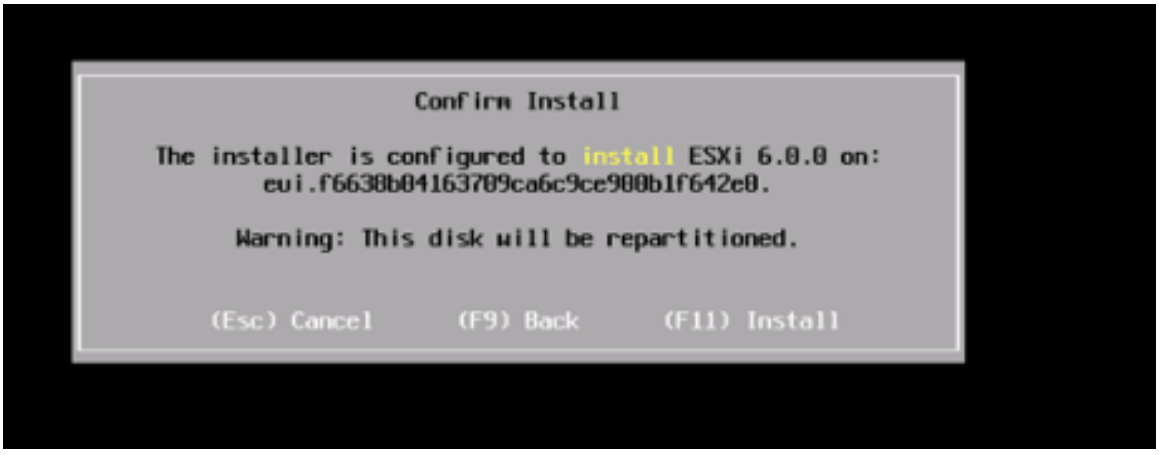

7. A progress bar will be displayed to show the ESXi installation progress on the server's boot volume on Nimble.
8. When the install is complete, select Virtual Media from the top menu and click on Activate Virtual Devices. Click Yes to un-map the Virtual Media. Press Enter to Reboot the server.

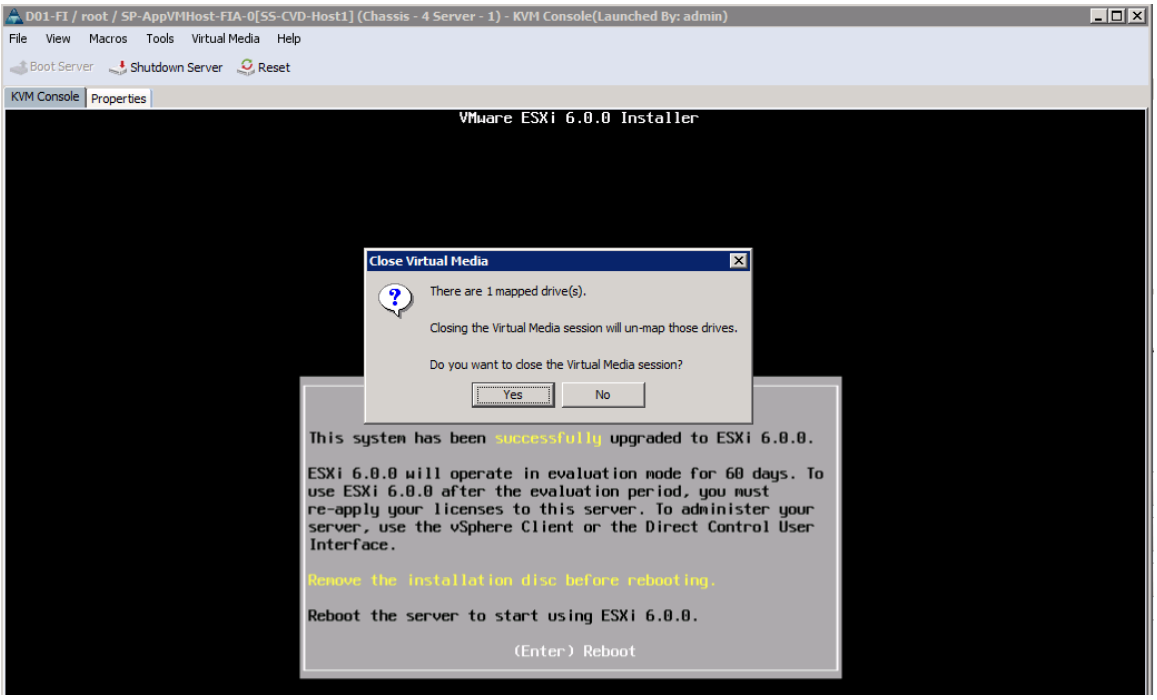

9. As the server boots, verify that the Nimble array is reachable through both SAN fabrics. The WWPN of the active controller's FC ports that connect to SAN Fabric A and SAN Fabric B should be accessible as shown below. During the ESXi install of this server, Nimble Controller A is Active. On Controller A, FC1 and FC5 ports connect to SAN Fabric A while FC2 and FC6 ports connect to SAN Fabric B. Though the active controller has dual connections to each fabric, during the boot up process, reachability to only one of the FC ports will be attempted through each Fabric as shown below.

10. Active Nimble Controller A's FC1 port (WWPN: 56:c9:ce:90:49:09:db:01) connects to SAN Fabric A. Host can reach the active controller's FC1 port (WWPN: 56:c9:ce:90:49:09:db:01) through SAN Fabric A as shown below.

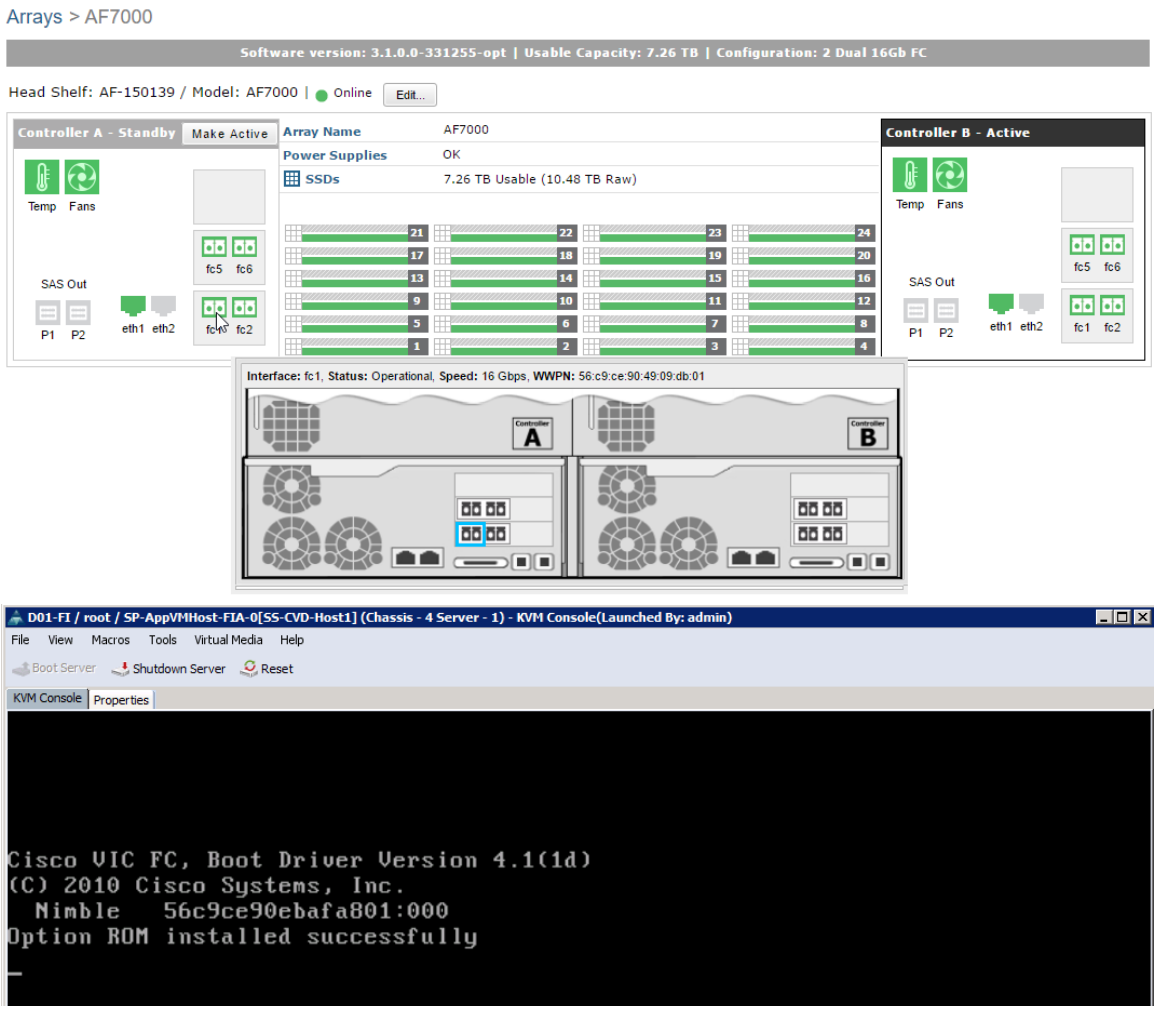

11. Active Nimble Controller A's FC2 port (WWPN: 56:c9:ce:90:49:09:db:02) connects to SAN Fabric B. Host can reach the active controller's FC2 port (WWPN: 56:c9:ce:90:49:09:db:02) through SAN Fabric B as shown below.

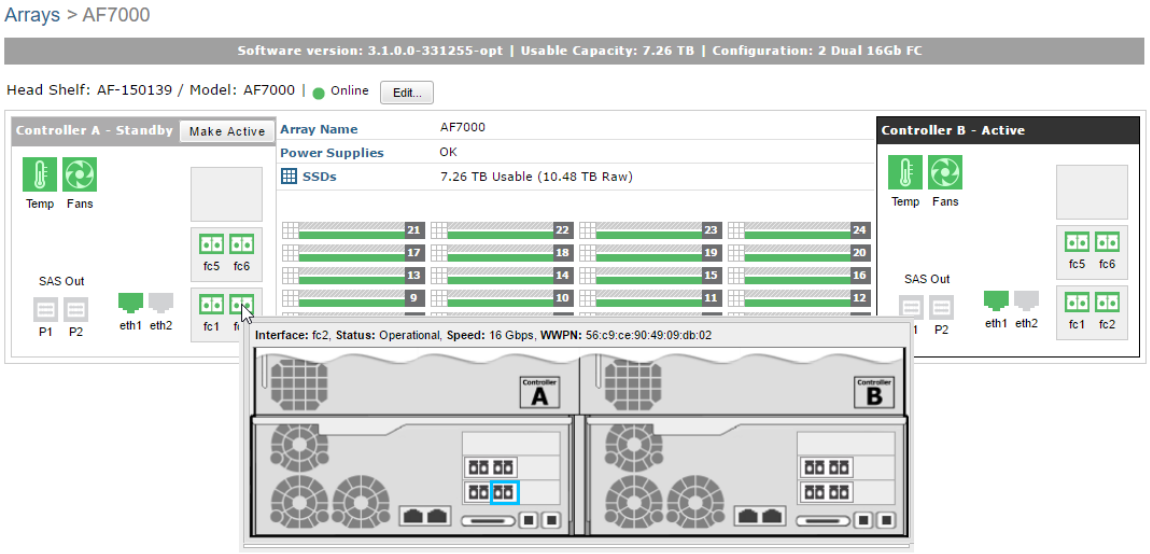

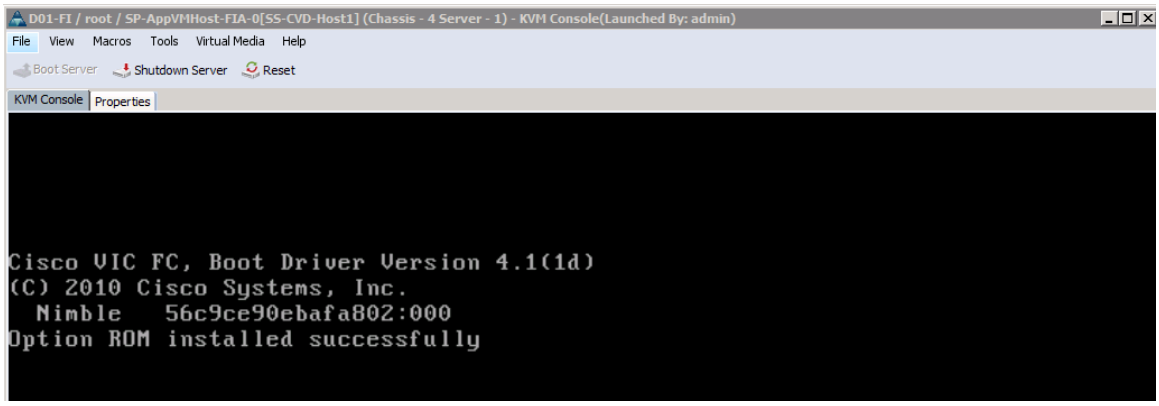

12. The Server boots the ESXi image from the Nimble boot volume and is now ready to be configured.

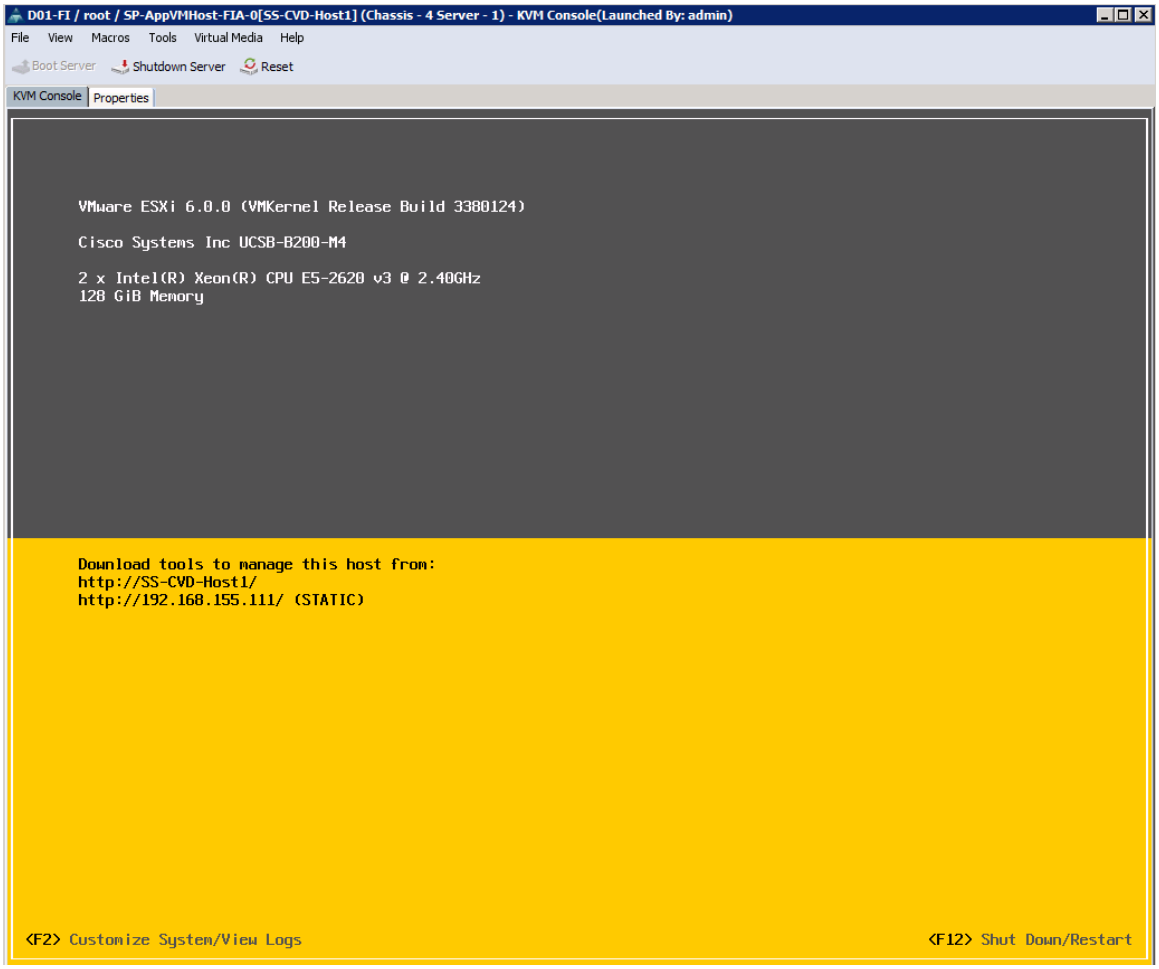

Setup ESXi Host for IP Management

1. Once the host comes up, press F2 on the KVM console to Customize System/View Logs and login as root.
2. Navigate to the Configure Management Network option and press Enter.
3. (Optional) From the Configure Management Network menu, navigate to VLAN and press Enter. Enter the in-band management VLAN ID (for example, `VLAN 12`) and press Enter.
4. From the Configure Management Network menu, select the IPv4 Configuration option and press Enter.
5. Using the space bar key, select Set Static IPv4 Address and Network Configuration option.
6. Enter the Management IP address (for example, `192.168.155.111`), subnet mask (for example, `255.255.255.0`) and default gateway for the ESXi host (for example, `192.168.155.1`). Press Enter to accept the changes to the IP configuration.
7. If IPv6 is not used for management, disable it. From the Configure Management Network menu, select IPv6 Configuration option and press Enter. Using the spacebar, select Disable IPv6 (restart required) and press Enter to accept changes.
8. From the Configure Management Network menu, select DNS Configuration option and press Enter. Enter the IP address of the primary DNS server and the secondary DNS server (optional). Enter the fully

qualified domain name (FQDN) for the first ESXi host. Press Enter to accept the changes to the DNS configuration.

9. Press Esc to exit the Configure Management Network submenu.
10. Press Y to Confirm changes and reboot host.
11. The ESXi host reboots. After reboot, press F2 and log back in as root.
12. Select Test Management Network to verify the management network is set up correctly and press Enter.
13. Press Enter to run the test.
14. Press Enter to exit the window.
15. Press Esc to log out of the VMware console.

Server Setup - Post ESXi and vCenter Install

This section covers the setup required once the ESXi host has come up using FC SAN boot. The following configurations will need to be done by directly accessing the host using vSphere client or from vCenter web interface. The setup required covered in this section are summarized below.

- Enable NTP on the hosts
- Update Host FNIC and ENIC drivers if needed
- Setup Nimble Connection manager (NCM)
- Verify Storage Configuration Post-NCM Install
- Setup ESXi Dump Collector – Host Side
- Register Nimble vCenter plugin
- Specify Virtual Machine (VM) Swap File Location
- Setup Virtual Networking using VMware vSwitch

Enable NTP on Hosts

To enable NTP on the hosts, complete the following steps on each host using VMware vSphere web client.

1. From vSphere Web Client, login to vCenter and select the host in the inventory.
2. Click Manage tab, followed by Settings.
3. Click Time Configuration in the left pane.
4. Click Edit button on the upper right side of the window.
5. In the Edit Time Configuration dialog box for the host, select radio button to Enable NTP Client.
6. Specify the NTP server to use at the bottom of the window.
7. (Optional) Select NTP Service Startup Policy to Start and Stop with the host.
8. For the NTP Service Status, Click Start button to start the service.

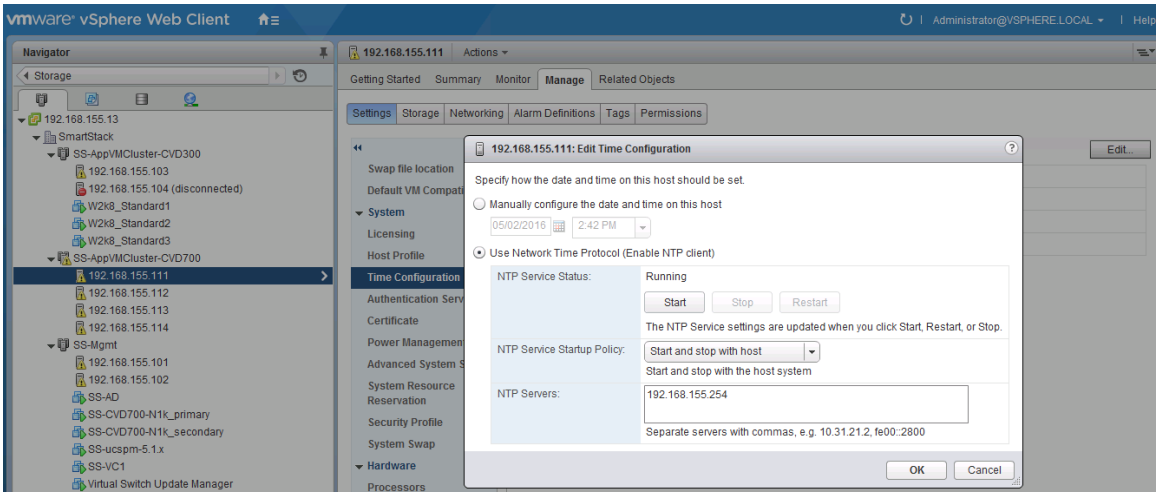

9. Click OK.
10. Repeat for all hosts in the SmartStack deployment.

The NTP server time may vary slightly from the host time.

Update ESXi Host FNIC and ENIC Drivers

The Cisco Custom ESXi ISO available from vmware.com may not have the latest FNIC and ENIC drivers that are recommended per the Cisco UCS Hardware and Software Interoperability matrix. If an upgrade is necessary, the following process can be followed to upgrade the necessary drivers.

1. To determine the latest supported drivers, search on http://www.cisco.com for the “Hardware and Software Interoperability for UCSM Managed Servers” document. This document will list the drivers required.
2. To download and install the drivers, follow the procedures outlined in the Install Guides. For the Cisco VIC used in this design, the Install Guide can be found as follows.
    a. From http://www.cisco.com, browse to Support → Downloads.
    b. In the product selector, click Products, then click Servers - Unified Computing.
    c. Select Server Software → UCS Virtual Interface Card.
    d. Select Install and Upgrade and click on Install and Upgrade Guides.
    e. Select UCS Virtual Interface Card Drivers for ESX Installation Guide and follow the download procedures outlined in the document to download and install Cisco VIC drivers for ESXi.

    Drivers can also be downloaded from vmware.com → search for “download cisco enic/fnic drivers” to find the drivers.

    ---

3. To upgrade the drivers in this SmartStack deployment, the following procedure was used.
    a. Select the driver bundle (.zip).
    b. Extract the contents of the driver zip file, and identify the *.vib file.
    c. From VMware vSphere web client, select the host from the Hosts and Clusters Inventory. Select Related Objects tab and then the datastore where you want to upload the drivers to. If possible,

use a datastore that is accessible by all hosts so that drivers can be uploaded once to upgrade multiple hosts.

d. Right-click and select Browse Files, then click on the Upload icon to upload the *.vib file to the datastore.

e. Select host and right-click and select Enter Maintenance mode to put the host into Maintenance Mode. Click Yes to Confirm and OK to acknowledge the Warning.

f. Next SSH into the ESXi host and install the driver using the following commands.

```
esxcli software vib list | grep CSCO
```
(to see all Cisco drivers)

```
esxcli software vib list | grep enic
esxcli software vib update –v /vmfs/volumes/Host1-DS0/<driver_name>
```
(Host1-DS0 is the datastore where the drivers were uploaded to previously)

g. If you get an error saying 'Could not find a trusted signer.' you can use the '--no-sig-check' option to work around it.

```
esxcli software vib update –v /vmfs/volumes/Host1-DS0/ <driver_name>  --no-sig-check
```

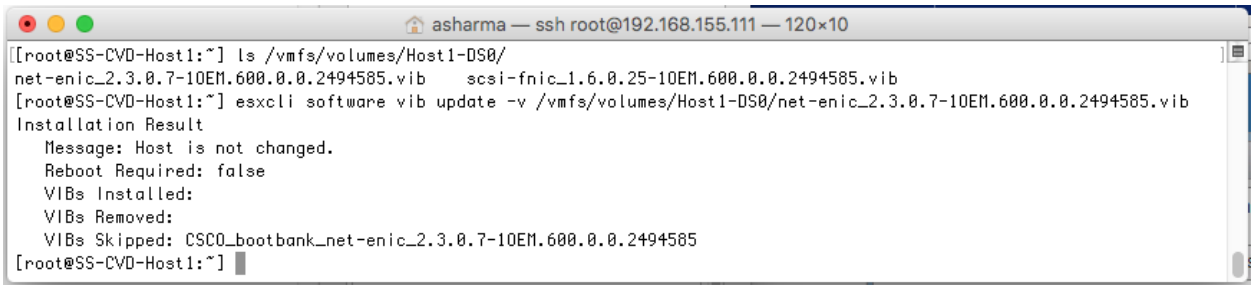

h. Repeat above steps for the FNIC drivers.

i. Reboot the host from vSphere client.

Setup Nimble Connection Manager (NCM)

NCM is required to enable optimal configuration with Nimble array such as setting up multipathing correctly, queue depth, timeout values and so on. Complete the following steps to download the appropriate NCM version for this array from Nimble Storage InfoSight website.

1. Put the newly deployed host in maintenance mode from vSphere client/vCenter before installing NCM.
2. Enable SSH (if not enabled already) on the host. This can be done via vSphere client or via KVM console into the host. From the KVM console, press F2 to Customize System/View Logs, select Troubleshooting Options and then Enable/Disable SSH will be available as an option.
3. SSH into the hosts and execute the following command:

```
esxcli software vib install --no-sig-check -d http://update.nimblestorage.com/esx6/ncm.
```

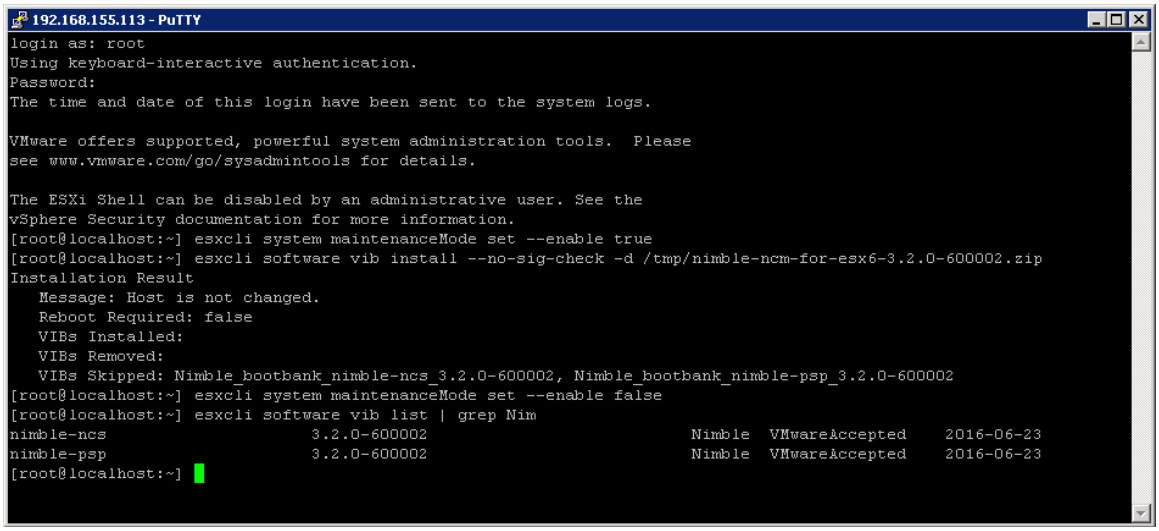

This example shows a host with the NCM software already installed. Attempting to reinstall it skips the install. For a new host, the NCM images downloaded would be displayed in the "VIBs Installed:" row. You can see what VIBs are currently install by running the "`esxcli software vib list | grep Nim`" command (to see what Nimble Storage specific VIBs are installed).

This command assumes the host can access the Nimble Storage InfoSight website. Alternatively, the image can be downloaded from another location within your site and positioned on the host.

4. Reboot host from VMware vSphere client and enable SSH (if needed) and log back into the host and verify that NCM is installed.
5. Take the host out of maintenance mode and disable SSH if needed.

Verifying Storage Configuration Post-NCM Install

In this section, the best practices implemented by installing NCM are verified. The specific NCM changes are:

1. Use vSphere web client to login to vCenter, select the Host under Hosts and Clusters. Click on the Manage tab and then Storage. Go to Storage devices in the menu and select one of the datastores.
2. In the bottom half of the window, click on the Properties tab and scroll down to the Multipathing Policies section and click Edit Multipathing. Select NIMBLE_PSP_DIRECTED from the drop-down list. Click OK.

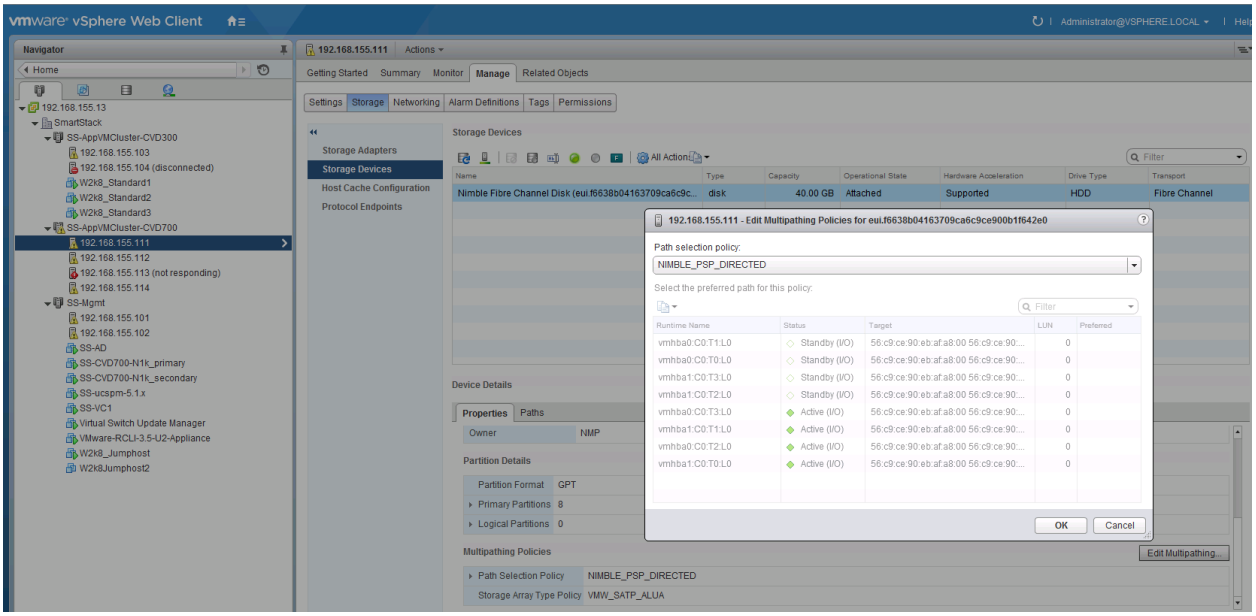

3. Click the Paths tab and verify that the info shows 4 Active Paths, 2 through Fabric A (vHBA0) and 2 through Fabric B (vHBA1). Similarly, there should also be 4 Standby Paths to the WWPN of Standby controller's ports.

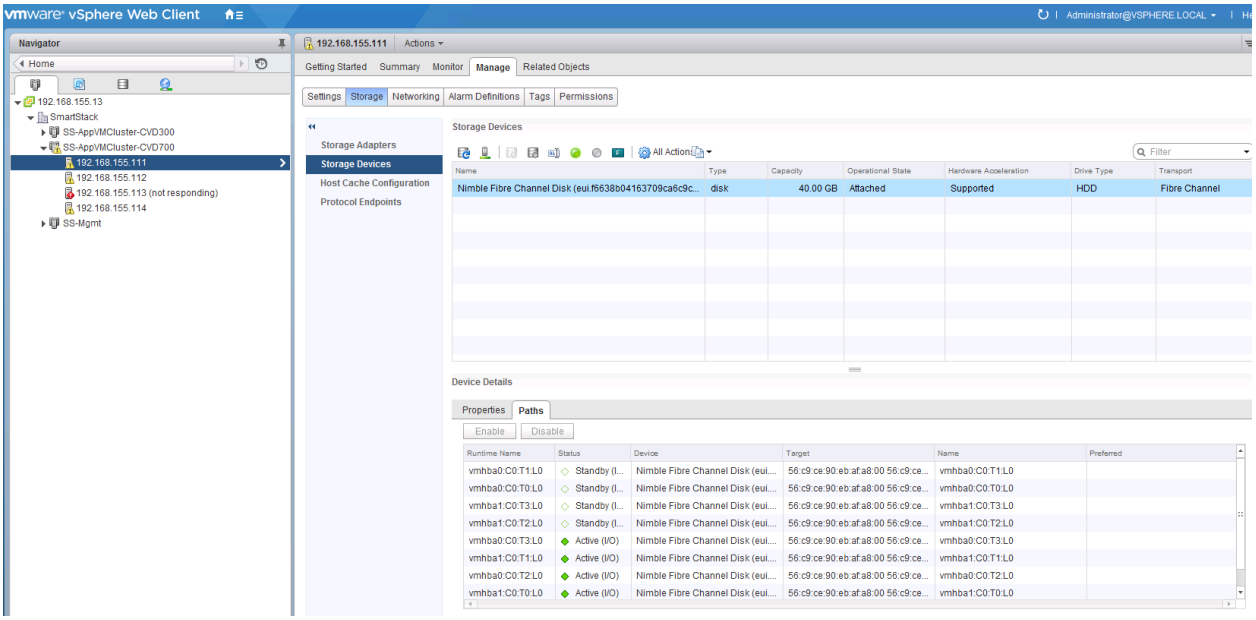

Setup ESXi Dump Collector - Host side

1. SSH into the host and login as root with associated password.
2. Execute the following commands to enable coredump to Dump Collector.
   - `esxcli system coredump network set --interface-name vmk0 --server-ip 192.168.155.13 --server-port 6500`
   - `esxcli system coredump network set --enable true`

— esxcli system coredump network check

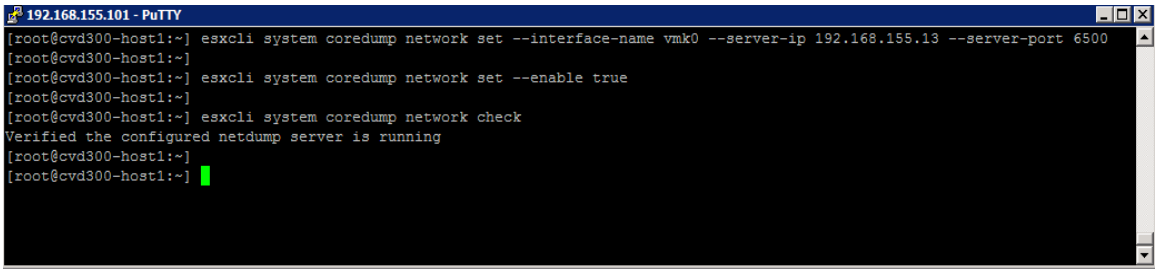

3. Repeat for each host in the topology.

Register Nimble vCenter Plugin

To register the Nimble vCenter plugin, use a vCenter account that has sufficient privileges to install a plugin. You need to know the vCenter hostname or IP address. The plugin is part of the Nimble OS. To take advantage of it, you must first register the plugin with a vCenter Server. Multiple plugins can be registered on the Nimble array. In turn, each array that registers the plugin adds a tab to the vSphere client. The tab name for the datastore page is "datacenter page-Nimble-<groupname>".

To register the vCenter plugin, complete the following steps.

1. From the Nimble OS GUI main menu, select Administration > vCenter Plugin.
2. If the fields are not already filled, enter the vCenter server host name or IP address, user name, and password.
3. Click View Status to see the current status of the plugin.
4. Click Register. If a Security Warning message appears, click Ignore.
5. Select a Subnet that is routable/ accessible to the vCenter host.

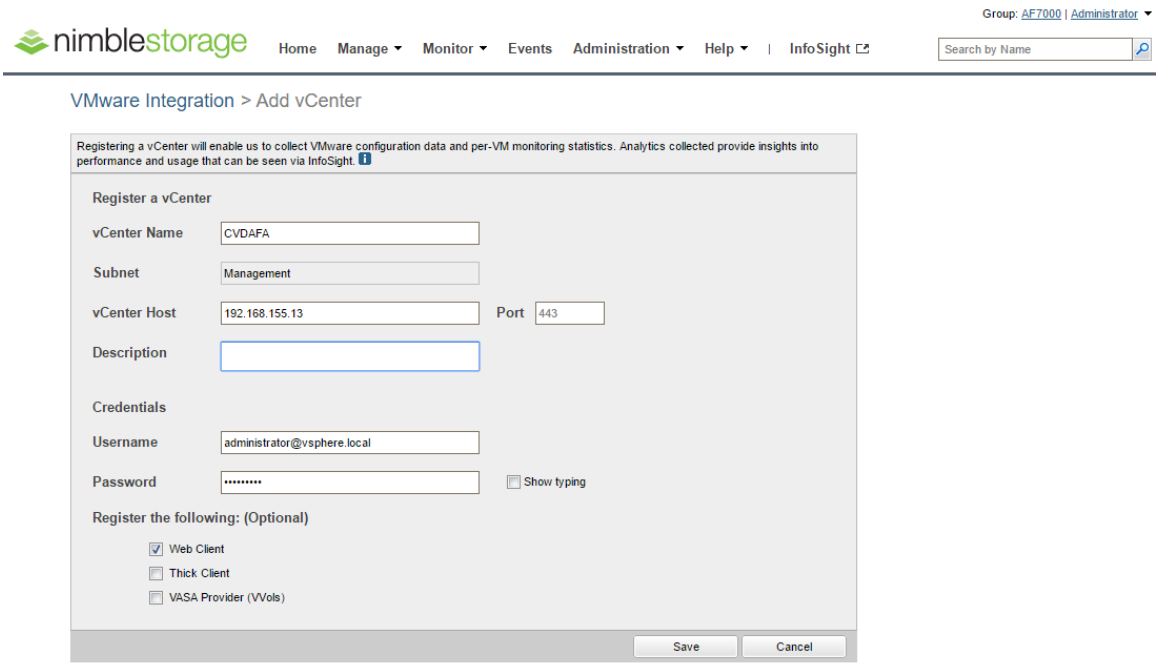

6. Click View Status again to ensure that the plugin has been registered.
7. Restart the vSphere thick client or re-login to the vCenter web client.

Verify Plugin Registration

A list of all registered plugins on the array can be discovered by doing the following:

1. Log into the Nimble OS CLI.
2. At the command prompt, enter the following command.

```
vmwplugin --list --username <username> --password <password> --server
<server_hostname-address> --port port_number <port number>
```

If no port number is specified, port 443 is selected by default.

3. The installed plugins are displayed as follows:

```
Nimble OS $ vmwplugin --list --username administrator@vsphere.local --password
 --server 192.168.155.13 --port 443


------------------+----------+---------------------------------------------
Key                 Version    Description
------------------+----------+---------------------------------------------
8036968026182778215 0.0.316666 Nimble Storage vCenter plug-in for CVDAFA
8036968026182778215 0.0.316666 Nimble Storage vCenter web-client for CVDAFA
```

Specify Virtual Machine (VM) Swap File location

SmartStack recommends that the VM swap file be located in a datastore created specifically to store VM swap files from multiple hosts. This requires a swap datastore to be created on the Nimble Array – see Storage section of this guide for creating a swap volume on the Nimble array. To change the virtual machine swap file location for a given host, configuration changes need to be done at the cluster and host level. The cluster level configuration changes the default location of storing the VM swap file in the same directory as the virtual machine to the datastore specified by the host. The host level configuration specifies the actual datastore (for example, `SWAP_DS`) the host should use to store the VM swap file.

At the cluster level, complete the following configuration steps.

1. From VMware vCenter using the vSphere web client, navigate to datacenter (for example, `SmartStack`) and then select the cluster (for example, `SS-AppVMCluster-CVDFC`).
2. On the right window pane, click the Manage tab. Select Settings > General > General and click the Edit button to edit the swap file location.
3. In the Edit Cluster Settings window, select Datastore specified by Host radio button and click OK.

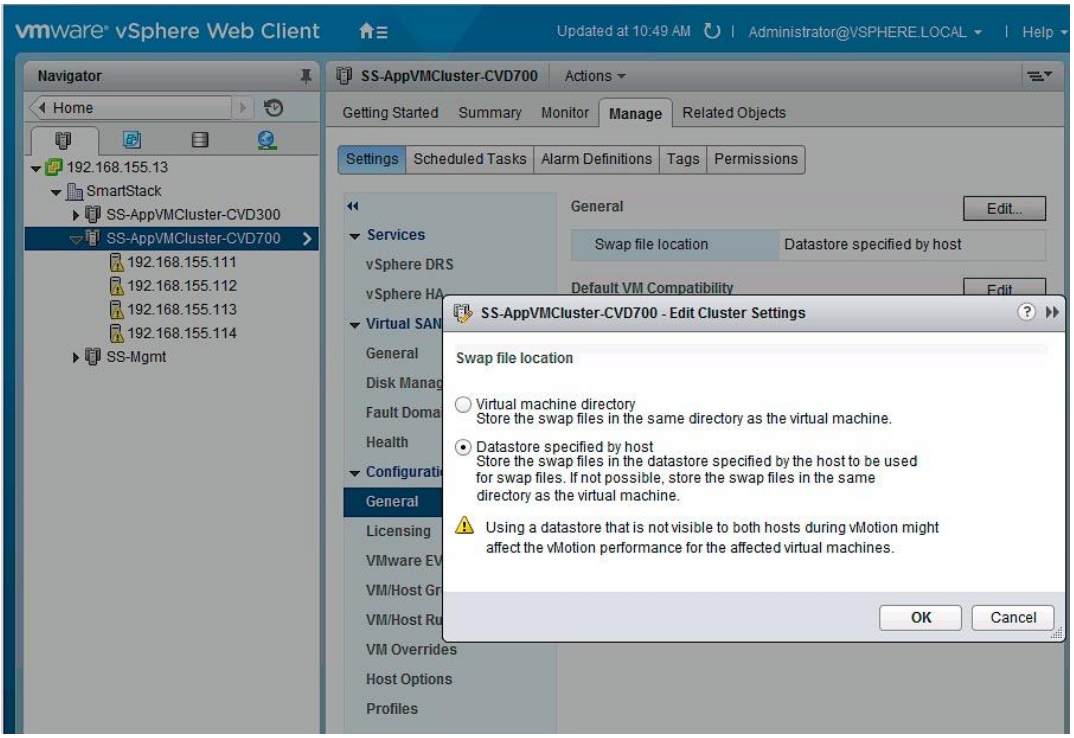

At the host level, complete the following configuration steps.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.111`). On the right window pane, click the Manage tab. Select Settings > Virtual Machines > Swap file location. Click the Edit button on the right side to edit the swap file location.
3. In the Edit Swap File Location pop-up window, select the datastore (for example, `SWAP-DS`) the host should use and click OK.

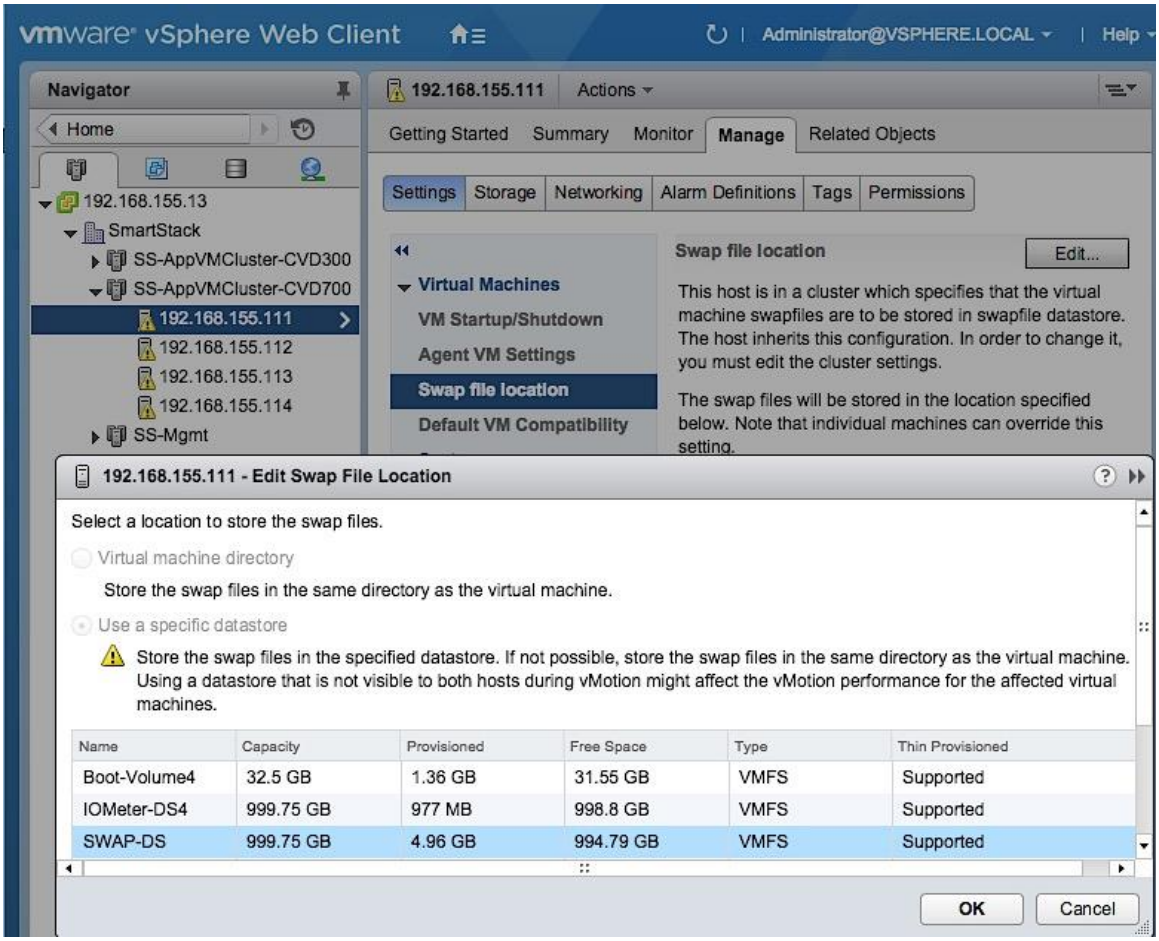

Virtual Networking Setup - using VMware vSwitch

This section covers the virtual switch (vSwitch) setup to enable network connectivity to hosts and VMs running on the hosts. The configuration workflow is as shown in the figure below.

Figure 17 Virtual Networking Setup Workflow

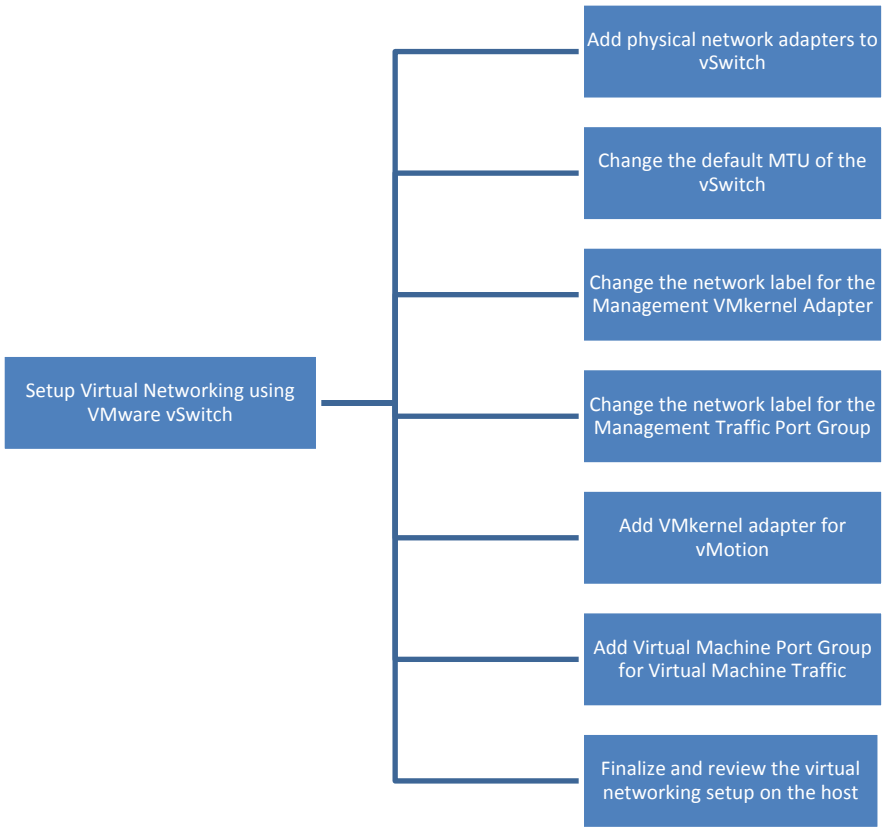

Add Physical Network Adapters to vSwitch

SmartStack design recommends the use of two vNICs (vNIC-A, vNIC-B) minimally. Traffic from these vNICs traverse in different paths across fabric providing redundancy and load balancing. The vNICs appear as vmnics to ESXi but by default, only one vmnic is associated with the default virtual switch(vSwitch0) on the host. To add the second VM NIC to vSwitch0, complete the following configuration steps for each host in the cluster.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click on the Manage tab. Click on Networking > Virtual switches and select vSwitch0 from the Virtual Switches section. Click on the **Manage physical adapter's** icon (3rd icon) to open the Manage Physical Network Adapters for vSwitch0 window.

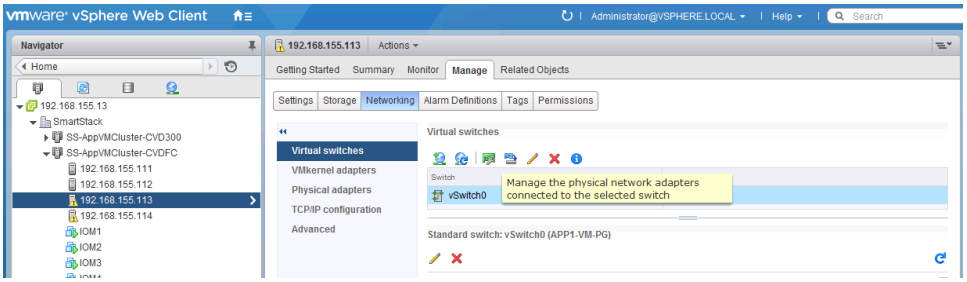

3. Click + to add Adapters. Select vmnic1 in the Add Physical Adapters to the Switch window and click OK twice to accept the edits and add vmnic1 as a second uplink/adapter to vSwitch0.

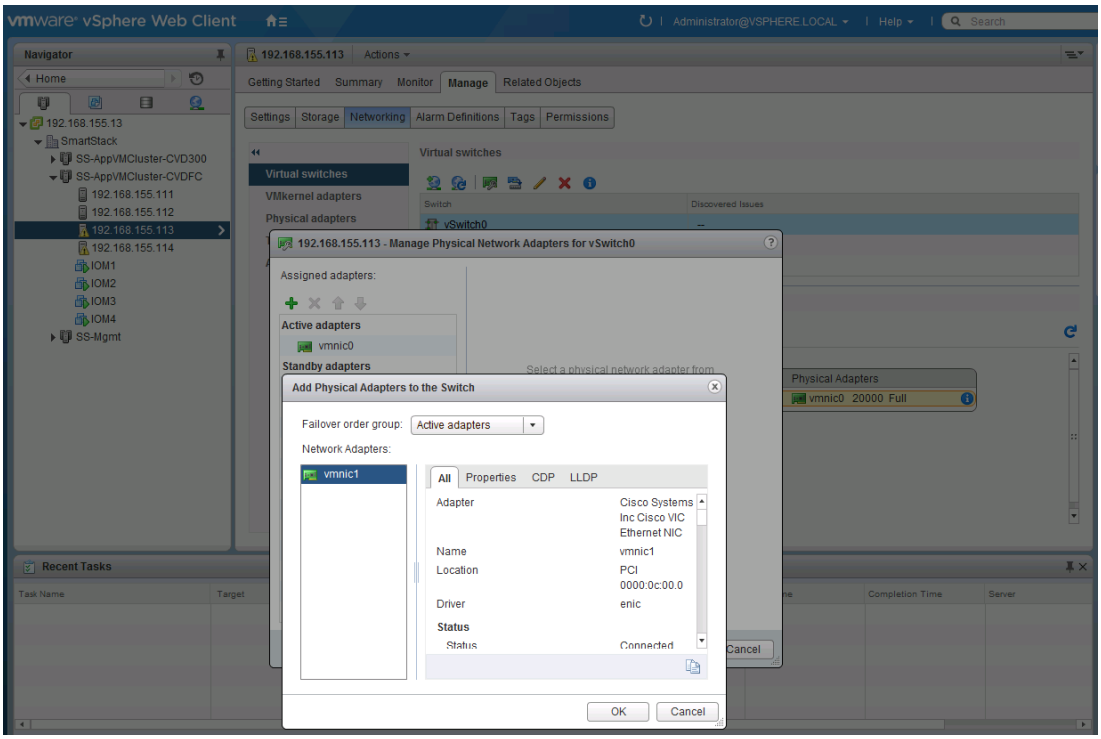

4. The resulting configuration is shown below.

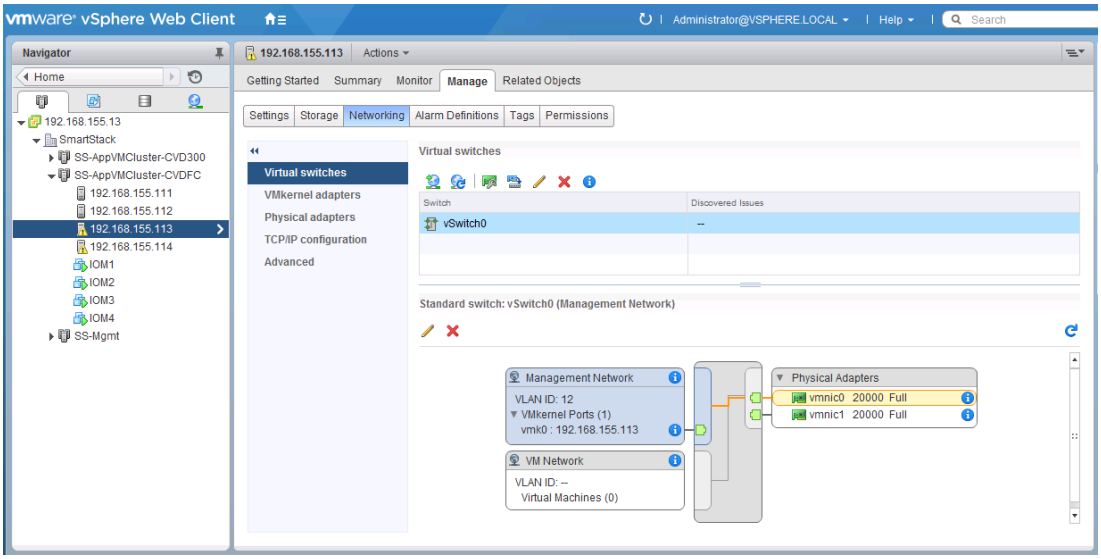

Change the Default MTU of the vSwitch

SmartStack design generally recommends an end-to-end MTU of 9000 for improved network throughput and performance to support functions such as vMotion. To change the default MTU, complete the following steps for each host in the cluster.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click on the Manage tab. Click on Networking > Virtual switches and select vSwitch0 from the Virtual Switches section. Click on the Edit Settings icon (5th icon) to open the Edit settings window. Change the MTU to 9000 as shown below. Click OK to accept the edits and close the window.

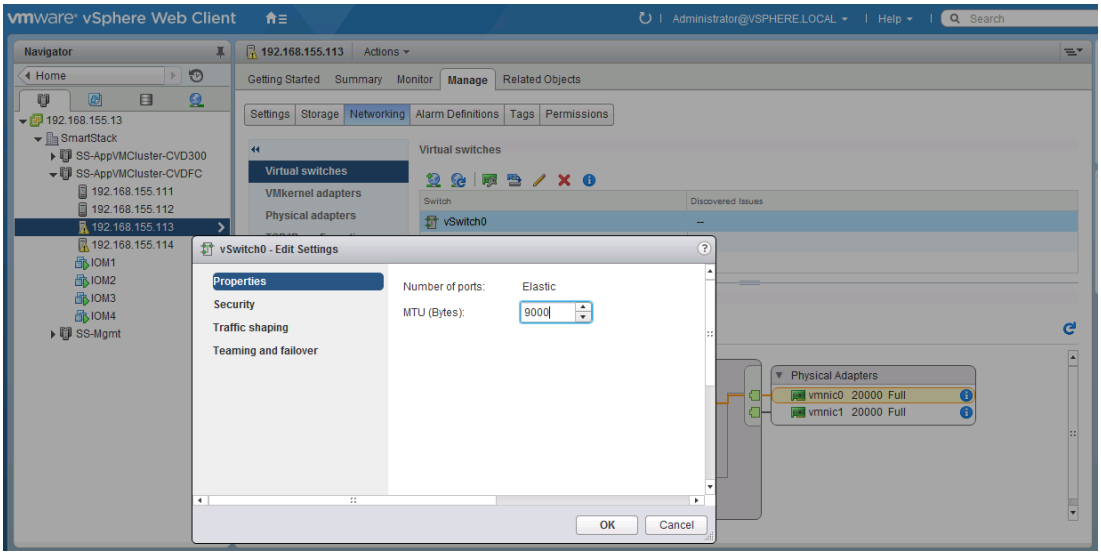

Change the Network Label for the Management VMkernel Adapter

To align with the naming scheme used in the SmartStack design, the network label for the Management VMkernel adapter was changed. For each host in the cluster, complete the following configuration steps to change the network label for the Management VMkernel adapter.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click on the Manage tab. Click on Networking > Virtual switches and select vSwitch0 from the Virtual Switches section. In the Standard Switch: vSwitch0 section of the window, select the Management Network and click on the Edit Settings icon. The Management Network: Edit Settings window will pop-up as shown below. Change the network label and VLAN ID to reflect the naming scheme and VLAN ID used. Click OK to accept the edits and close the window.

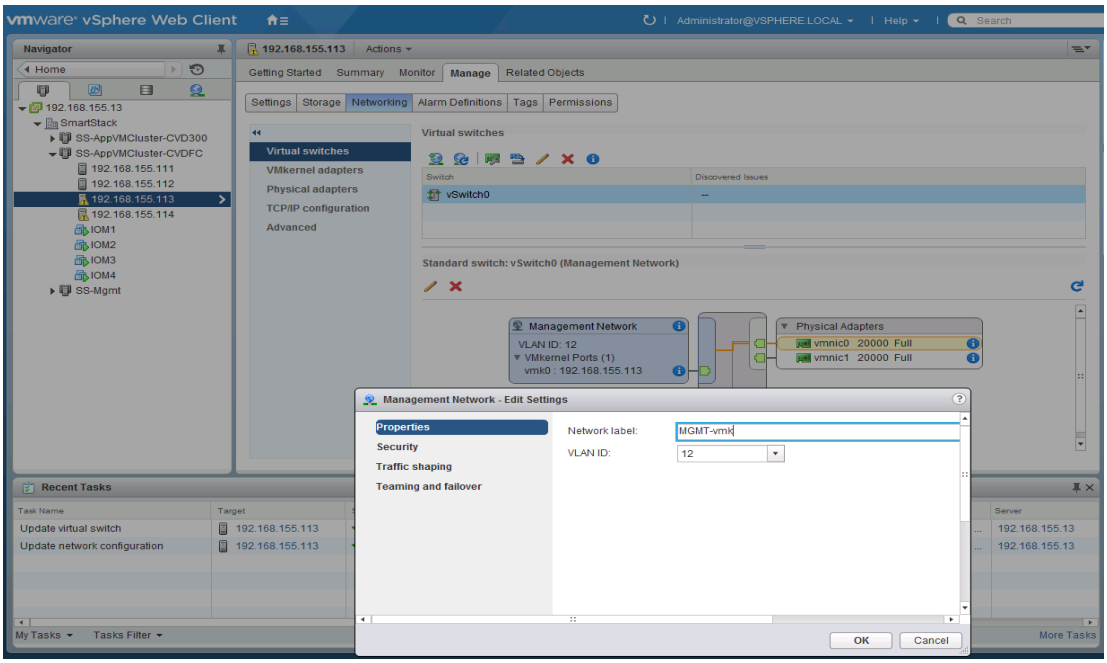

3. Click Networking > VMkernel adapters and select vmk0 from the VMkernel adapters section. Verify in the Properties table for vmk0 that Management Traffic is checked in the Enable services section.

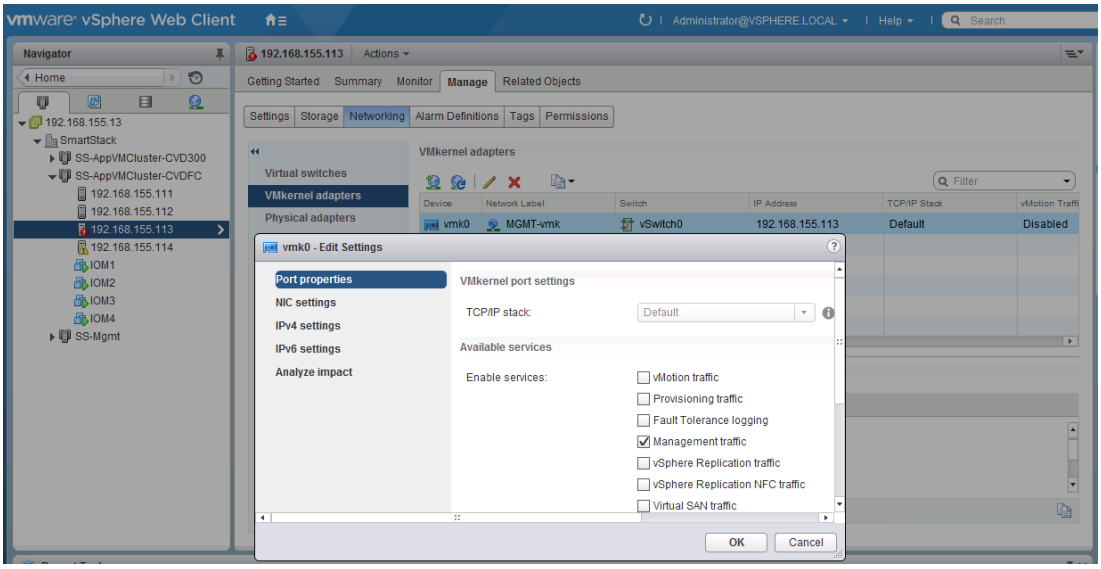

Change the Network Label for the Management Network

To align with the naming scheme, the network label for the Management Network was changed in the SmartStack design as outlined in this section.

For each host in the cluster, complete the following configuration steps:

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click the Manage tab. Click Networking > Virtual Switches and select vSwitch0 from the Virtual Switches section. In the Standard Switch: vSwitch0 section of the window, select VM Network and click the Edit Settings icon to open the Management Network: Edit Settings window as shown below. Change the network label and VLAN ID to reflect the naming scheme and VLAN ID used. Click OK to accept the edits and close the window.

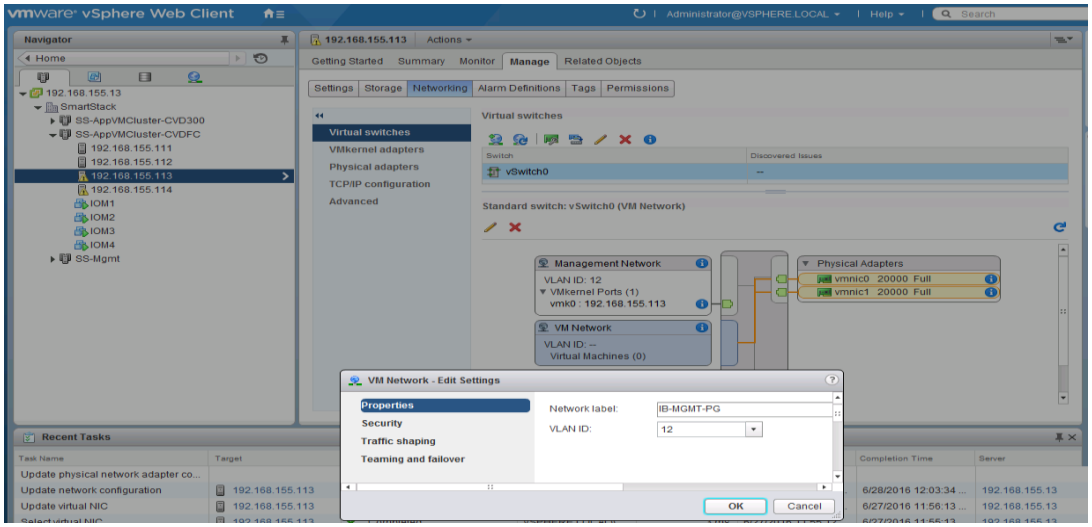

Add VMkernel Network Adapter for vMotion

To support vMotion in the SmartStack design, complete the following configuration steps for each host in the cluster.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click the Manage tab. Click on Networking > Virtual Switches and select vSwitch0 from the Virtual Switches section. Click on the Add host networking icon (1st icon). In the Add Networking window, select connection type as VMkernel Network Adapter. Click Next to continue.

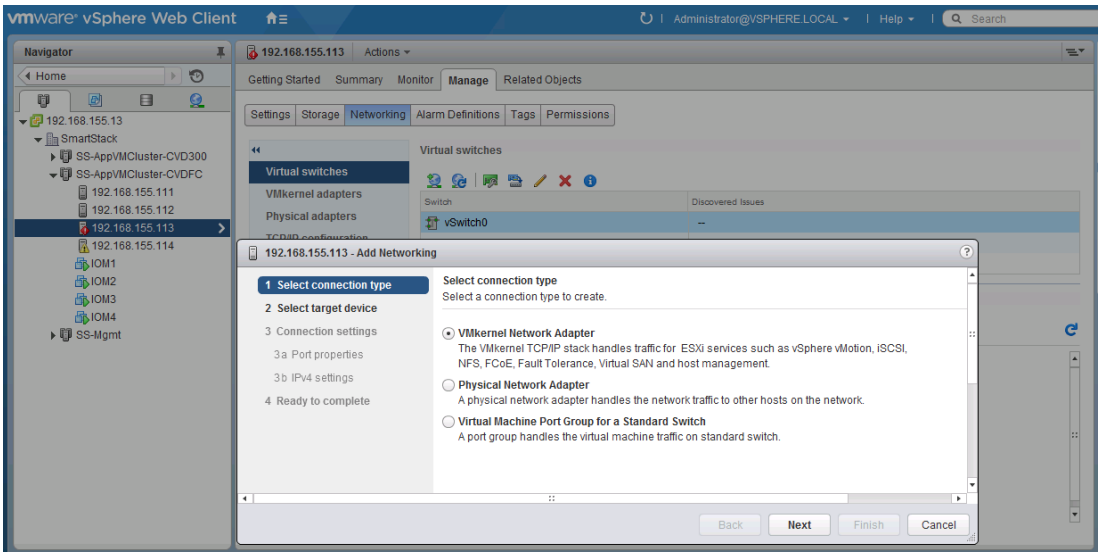

3. Select target device as an existing standard switch as shown below. Click Next to continue.

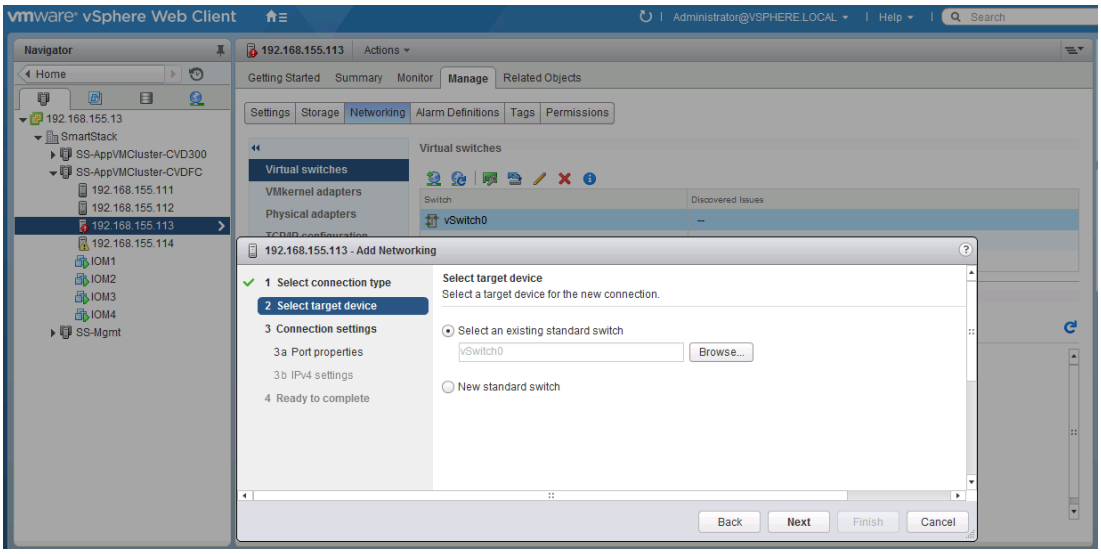

4. Under Connection Settings, for the Port Properties, specify the Network Label, VLAN ID (Optional) and enable vMotion traffic in the Enable Services section. Click Next to continue.

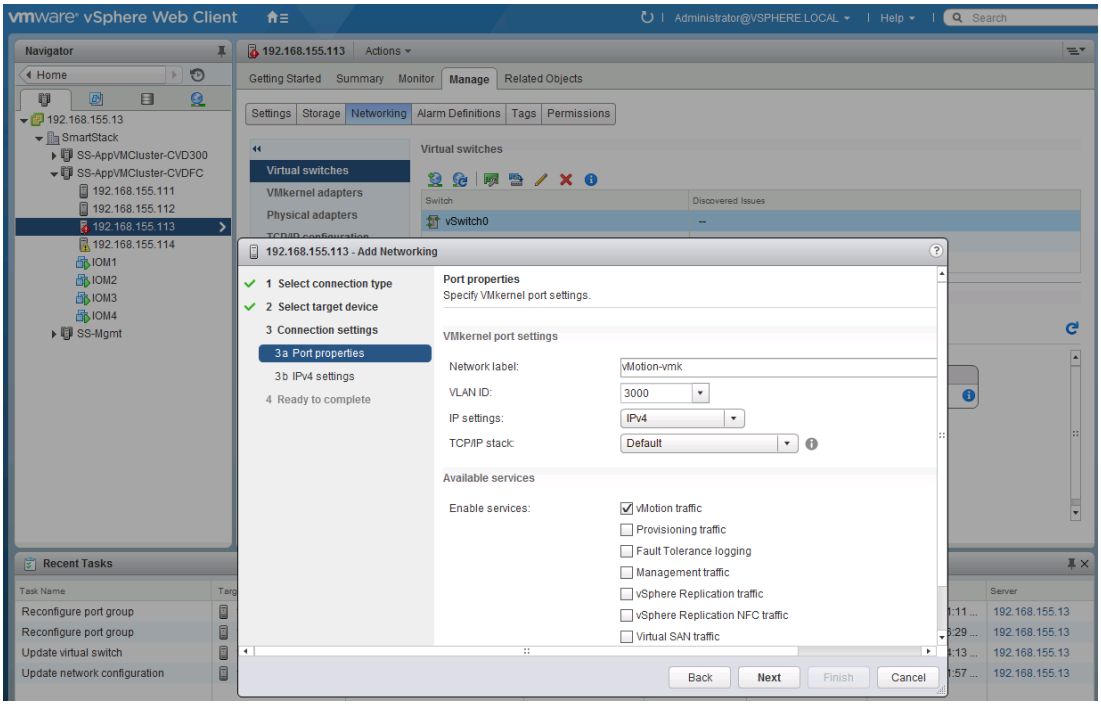

5. For the IPv4 Settings under Connection Settings, specify the IPv4 address to be used by vMotion VMkernel adapter. Click Next to continue.

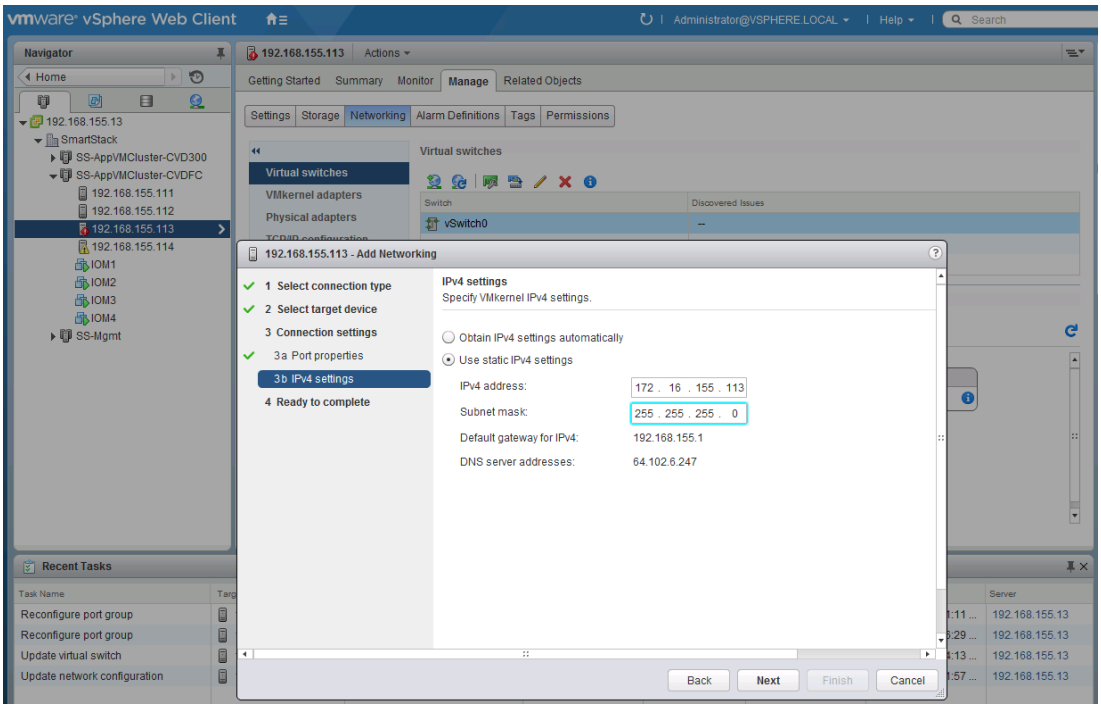

6. Review and click Finish to accept the edits and create the vMotion VMKernel adapter.

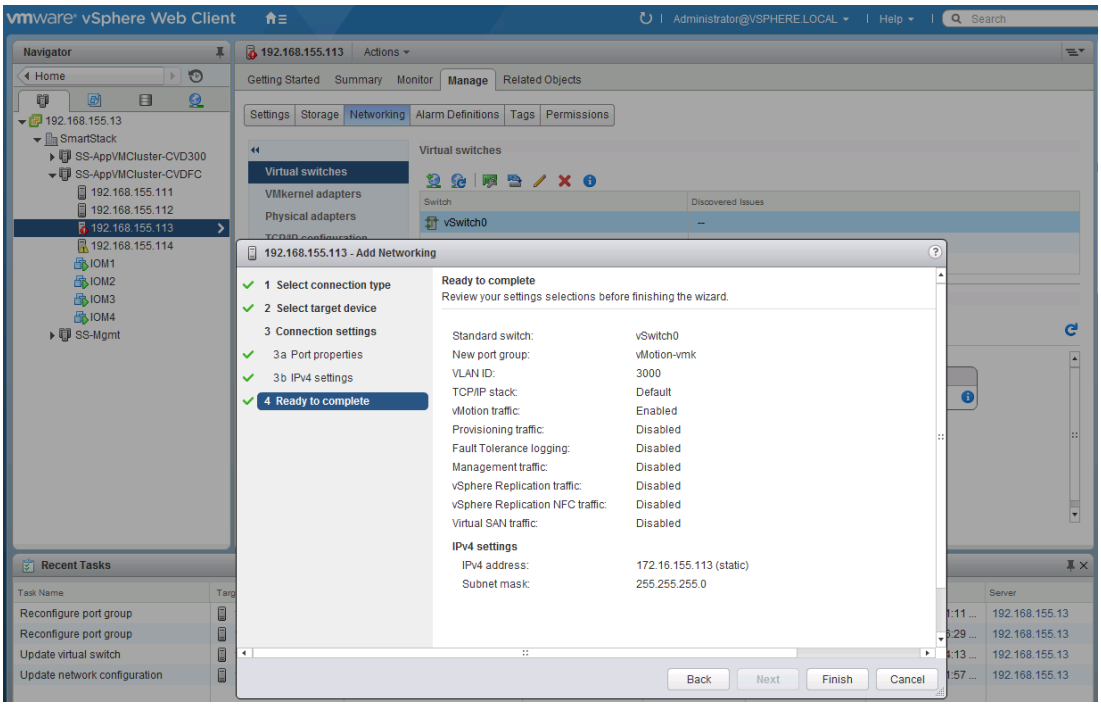

7. Click Networking > VMkernel adapters and select vmk1 from the VMkernel adapters section. Click Edit Settings icon (3rd icon). In the Edit Settings window for vmk1, select NIC settings and change the MTU to 9000. Click OK to accept the edits and close the window.

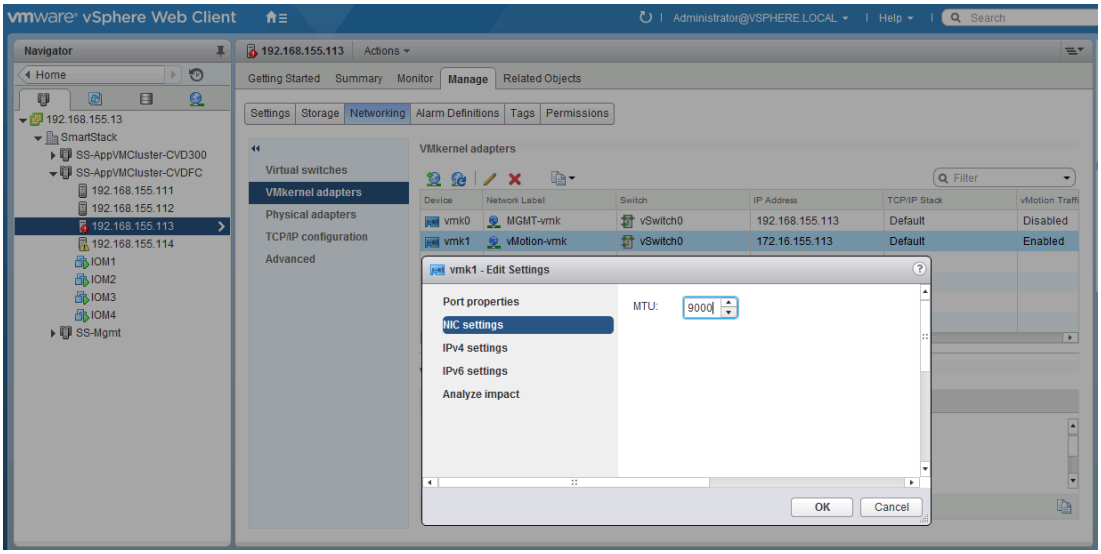

Add VM Port Group for Application Traffic

To add VM port groups for VM traffic, complete the following configuration steps for each host in the cluster.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`). On the right window pane, click the Manage tab. Click Networking > Virtual switches and select vSwitch0 from the Virtual Switches section. Click Add host networking icon (1st icon). In the Add Networking window, select connection type as Virtual Machine Port Group for a Standard Switch. Click Next to continue.

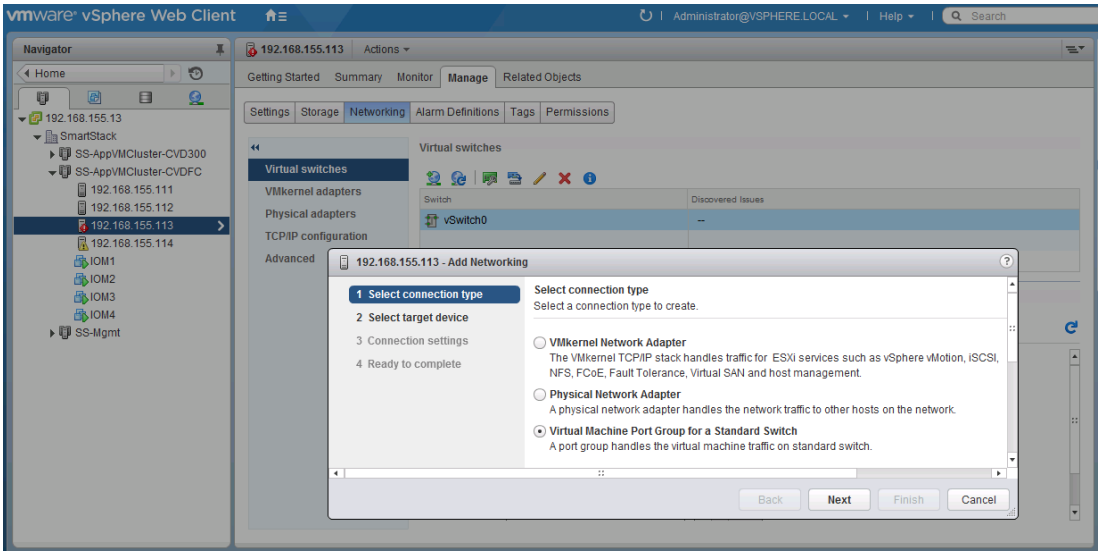

3. Select target device as an existing standard switch as shown below. Click Next to continue.

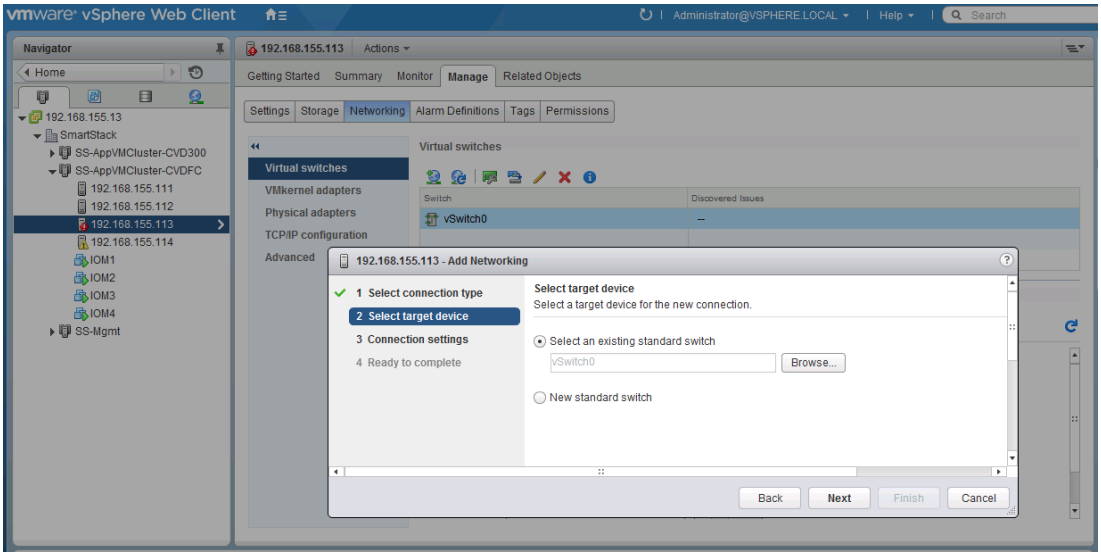

4. Under Connection Settings, for the Port Properties, specify the Network Label and VLAN ID (Optional). Click Next to continue.

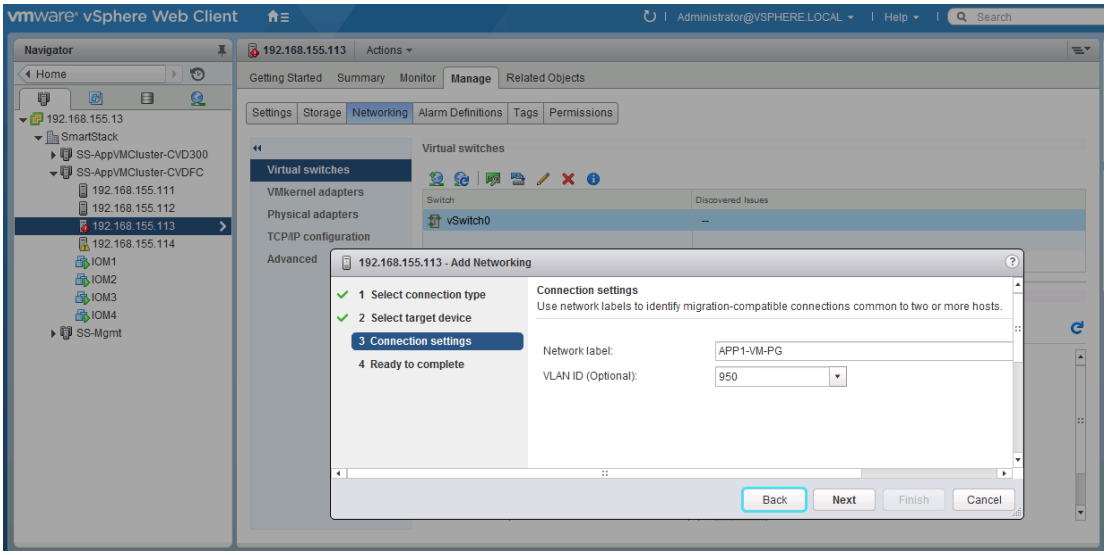

5. Review and click Finish to accept the edits and create the Virtual Machine Port Group.

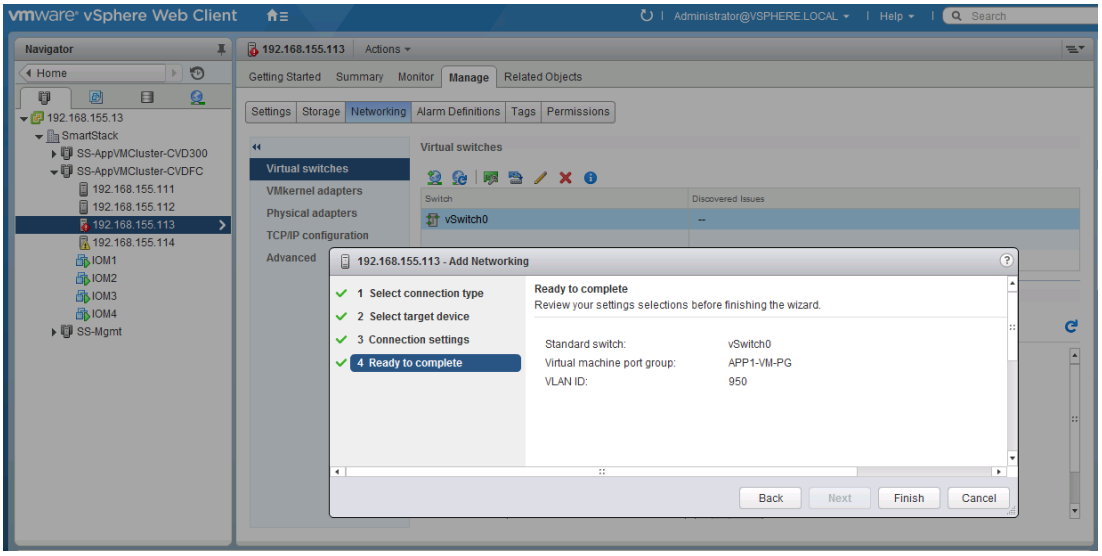

Finalize and Review the Virtual Networking setup on the ESXi Host

At this point, the networking on the ESXi host should be similar to the example below.

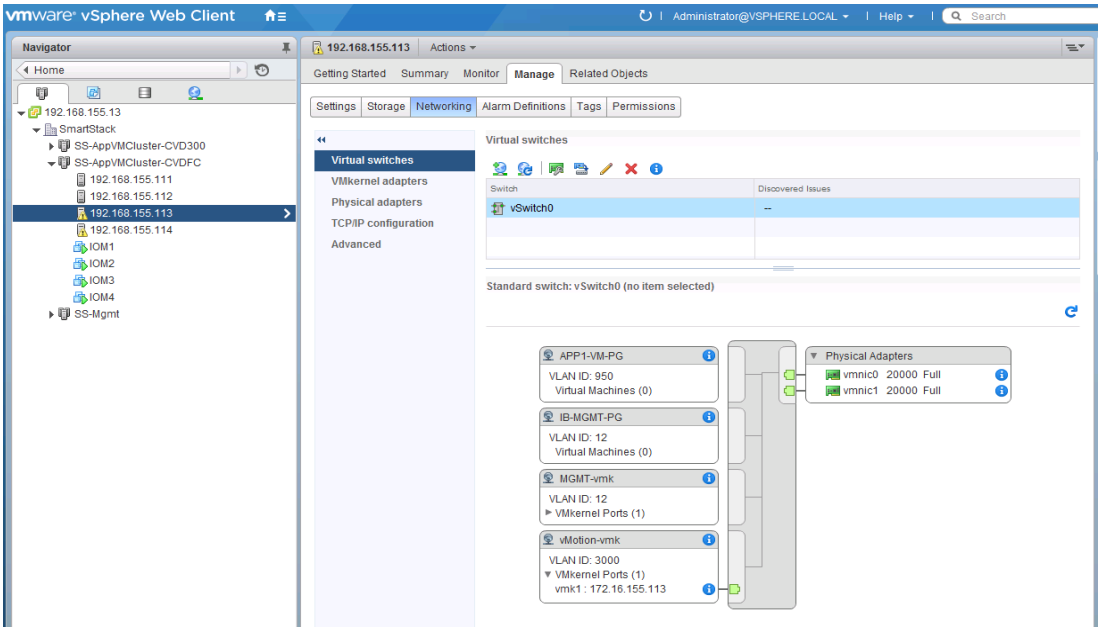

Optional: Migrate Virtual Networking from vSwitch to Cisco Nexus 1000V switch

Cisco Nexus 1000V switch can be used instead of vSwitch to provide network connectivity in the SmartStack design. SmartStack validation was done using a Cisco Nexus 1000V. Virtual Switch Update Manager (VSUM) that registers as a client plugin with VMware vCenter was used to deploy and manage the Cisco Nexus 1000V environment.

To migrate to Cisco Nexus 1000V from a VMware vSwitch environment, the following pre-migration steps are recommended.

Pre-Migration Setup: Remove Redundant Uplink from vSwitch

SmartStack design recommends two uplinks (vmnic0, vmnic1) for redundancy. However, when migrating to Cisco Nexus 1000V, remove one of the redundant vmnics from vSwitch so that it can be used as uplink for Cisco Nexus 1000V. Once the migration is successful, the first vmnic can also be moved to Cisco Nexus 1000V.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`) and click on Manage > Networking > Virtual switches. Select vSwitch0 from the Virtual Switches list. Click on the Manage Physical Adapters icon (3rd icon).

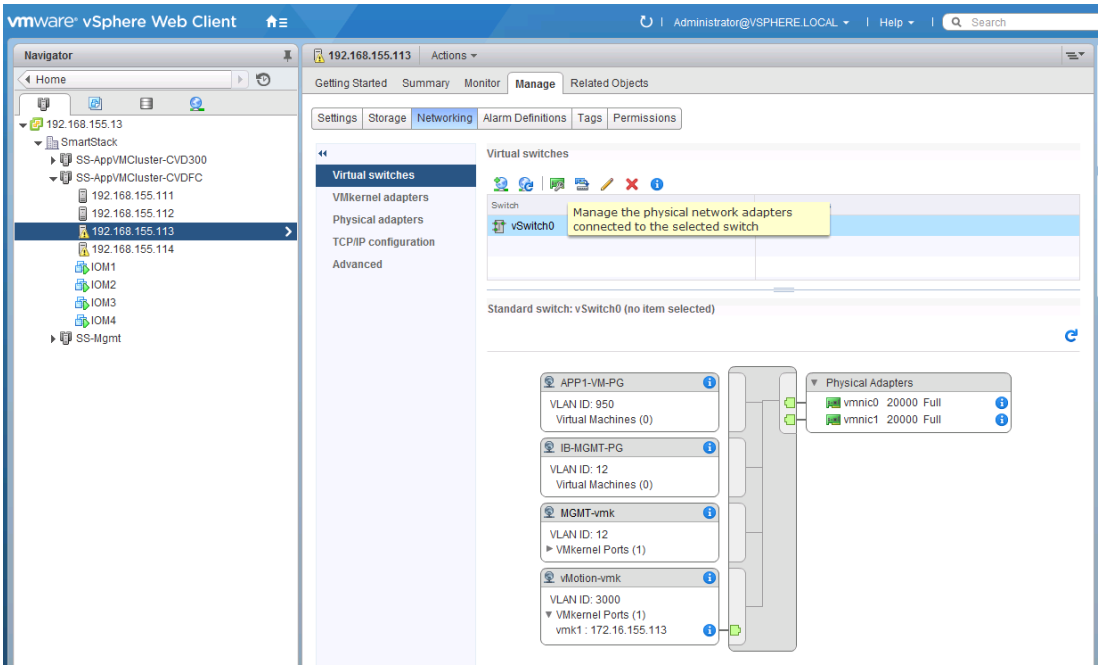

3. Select vmnic1 and click X to remove it from the list of Active Adapters. Click OK to continue.

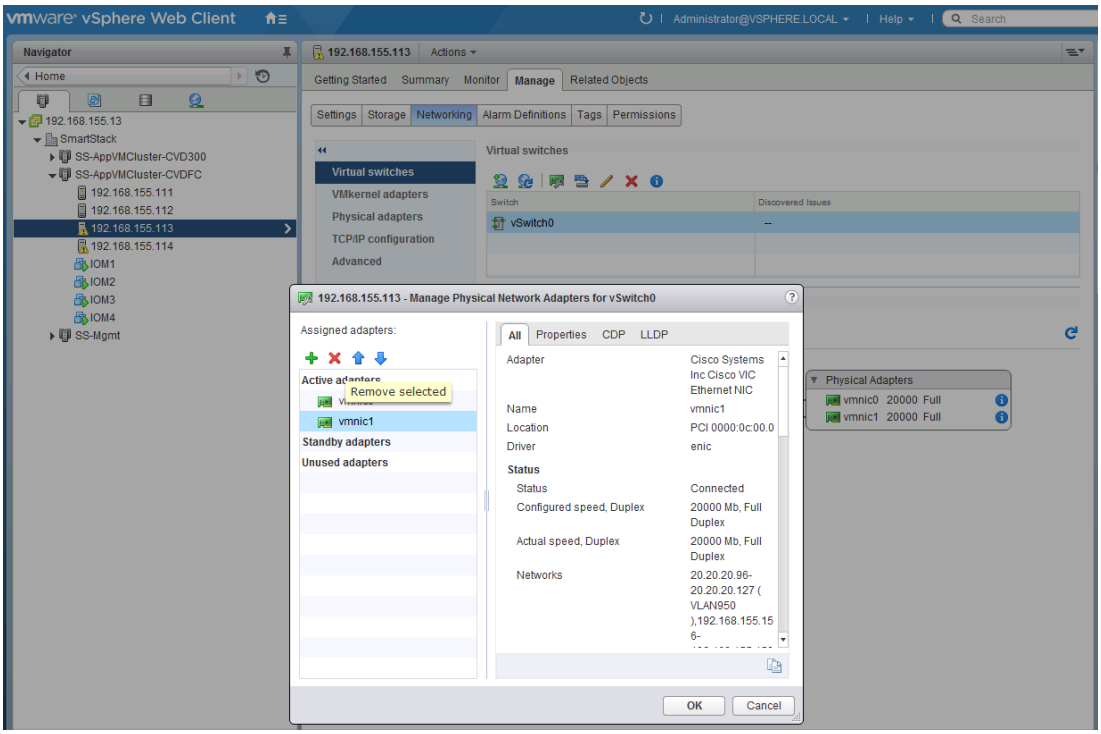

Pre-Migration Setup: Create a Temporary VMkernel Adapter

For network availability during migration, a temporary management VMkernel adapter is required. Complete the following configuration steps to create the adapter on vSwitch0. Repeat for all hosts using Cisco Nexus 1000V.

1. From VMware vCenter using the vSphere web client, navigate to the datacenter (for example, `Smart-Stack`) and cluster (for example, `SS-AppVMCluster-CVDFC`) where the host resides.
2. Select the host (for example, `192.168.155.113`) and click Manage > Networking > Virtual switches and select vSwitch0 from the Virtual Switches section. Click Add host networking icon (1st icon) and select connection type as VMkernel Network Adapter. Click Next to continue.

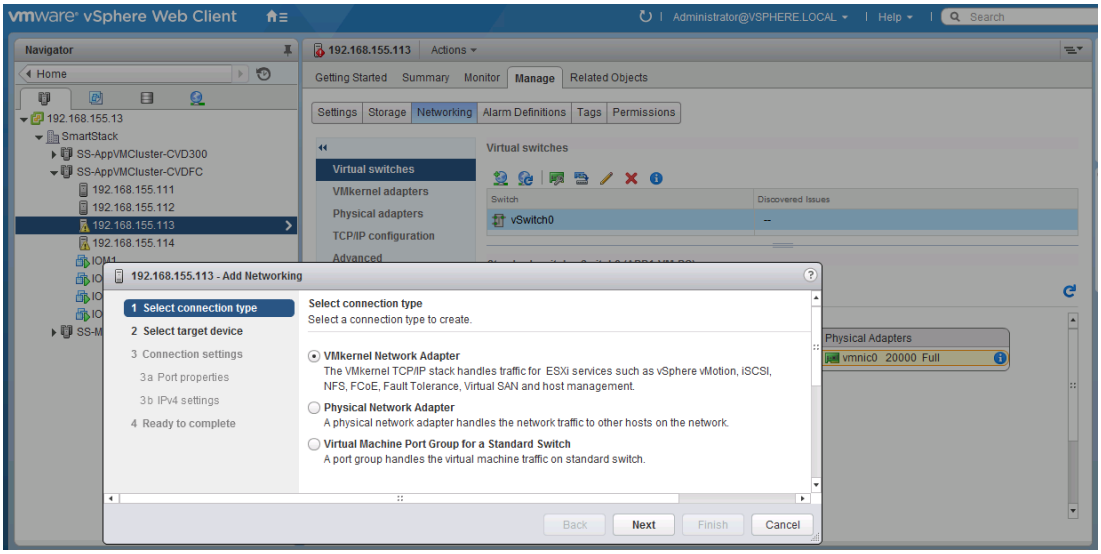

3. Select target device as the existing standard switch as shown below. Click Next to continue.

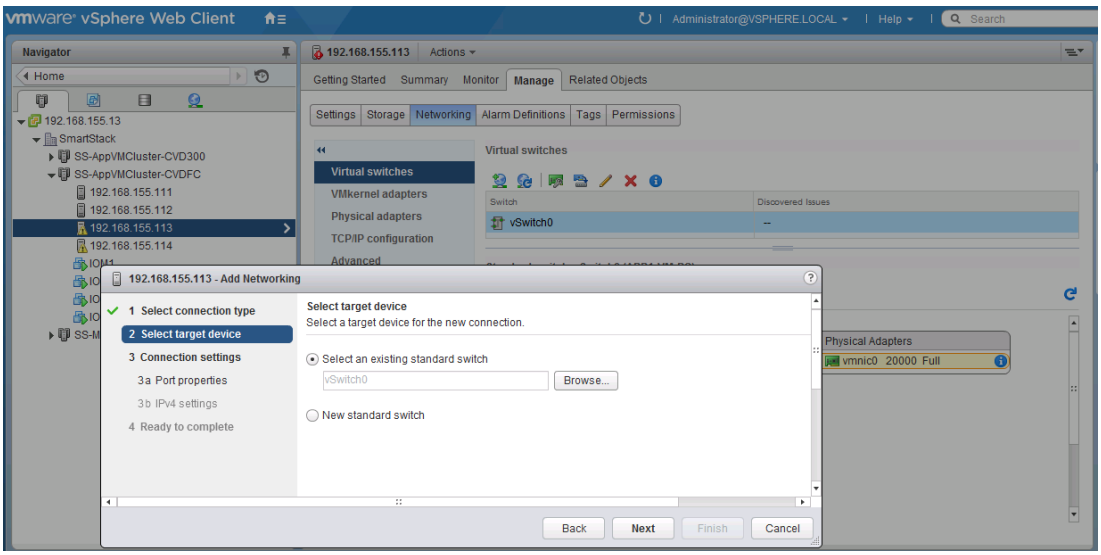

4. Under Connection Settings, for the Port Properties, specify the Network Label (for example, `TEMP-MGMT-vmk`) and VLAN ID. Select Management traffic in the Enable Services section. Click Next to continue.

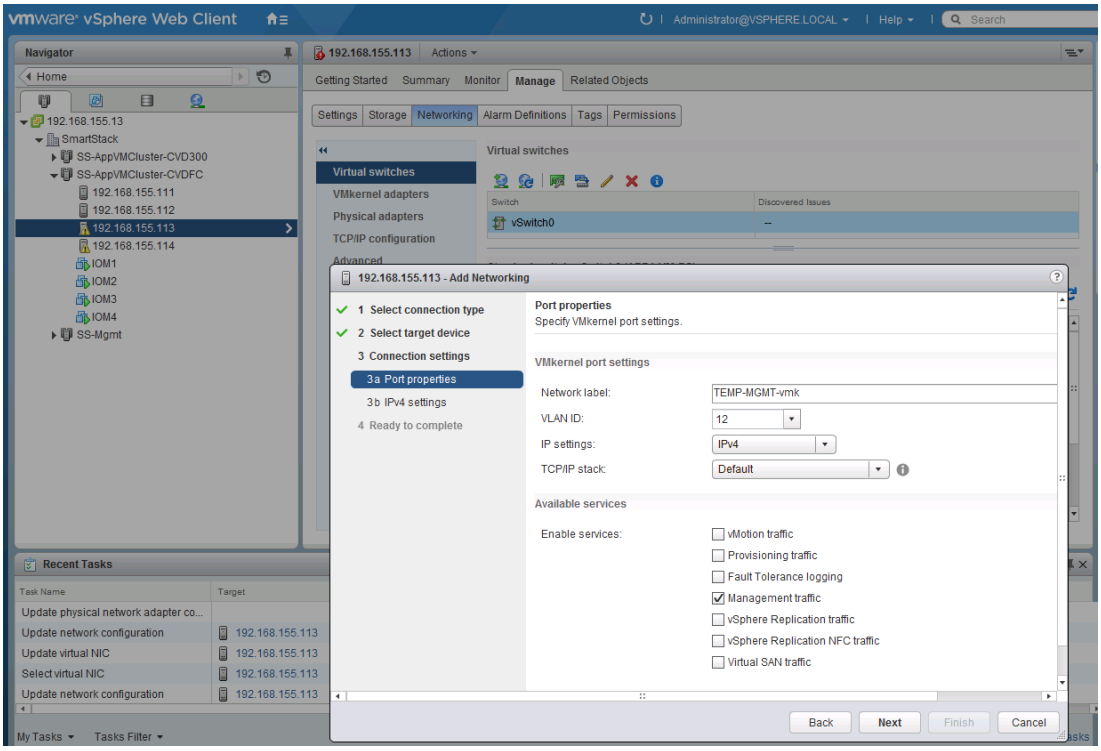

5. Specify the IPv4 settings for the temporary VMkernel adapter and click Next to continue.

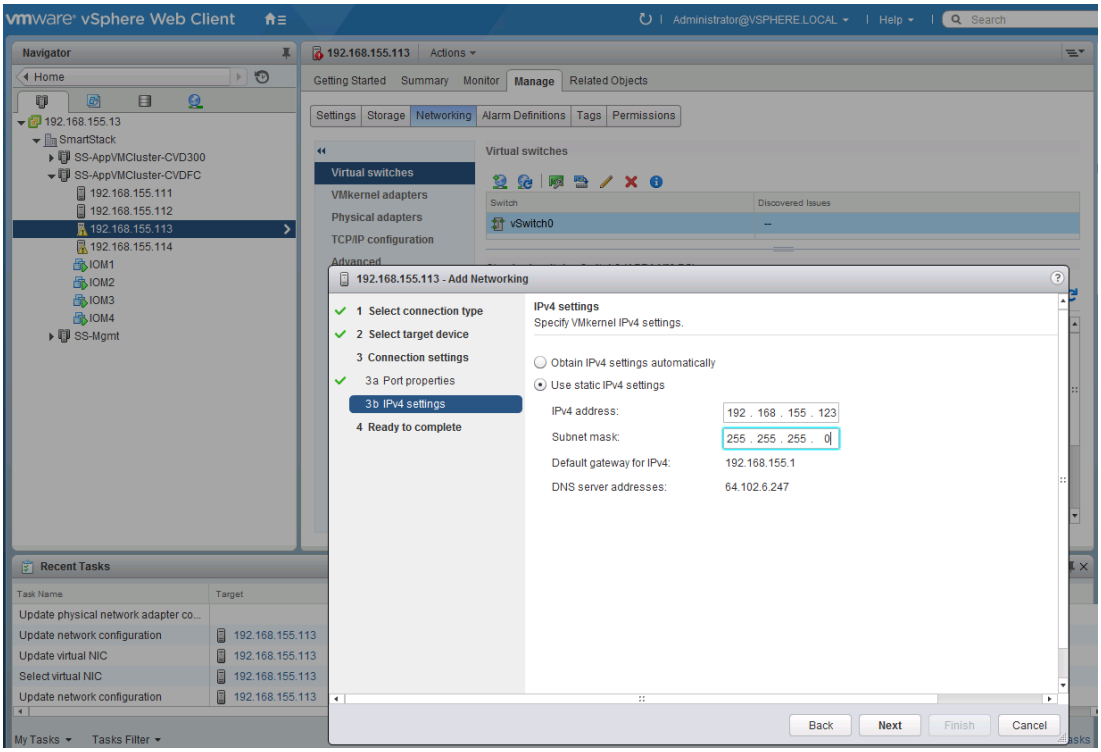

6. Review the settings and click Finish to accept the settings.

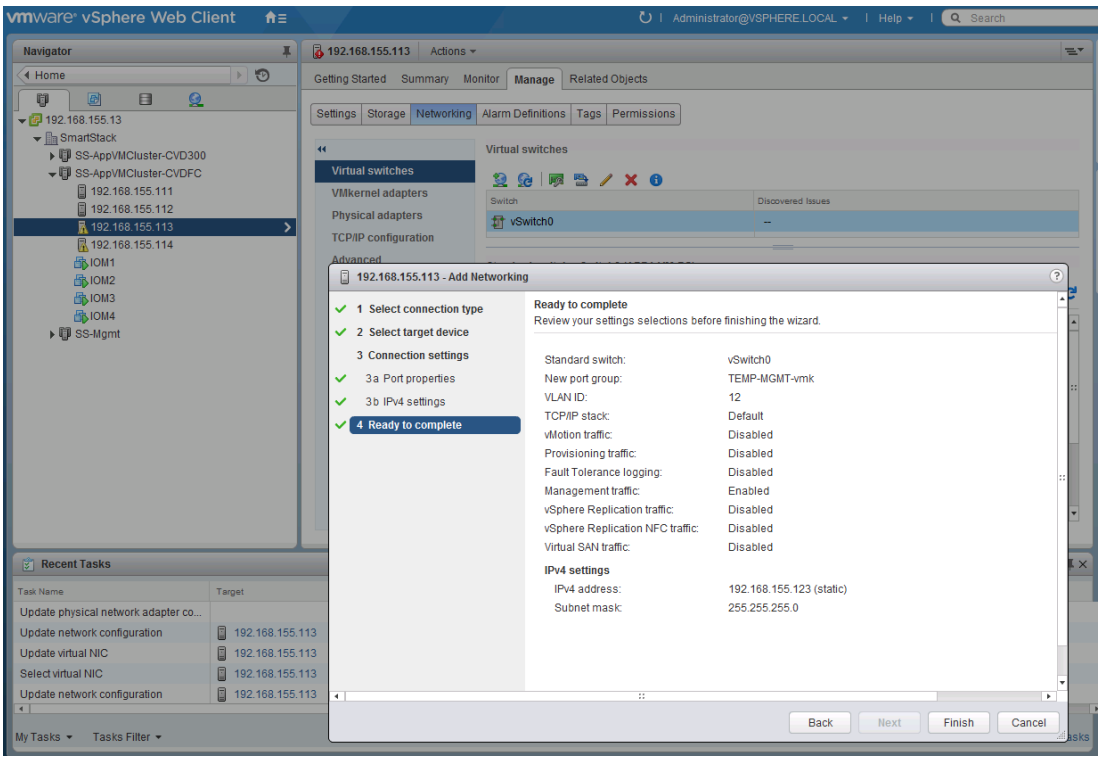

Download and Deploy the OVF Template for the Virtual Switch Update Manager (VSUM)

To deploy the OVF template, complete the following the steps:

1. Log in and Download the Cisco Nexus 1000V installation software from www.cisco.com.

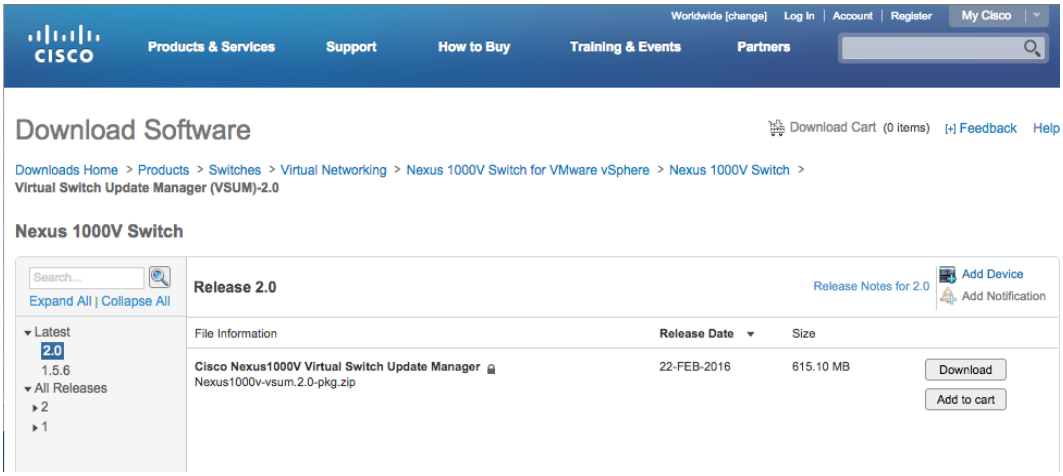

2. Unzip the package.
3. From the vSphere web client, go to Hosts and Clusters. At the top level of the navigation bar, right click and select Deploy OVF Template from the menu. Browse to the location where the package is and select the ova file to be deployed. Click Next to continue.

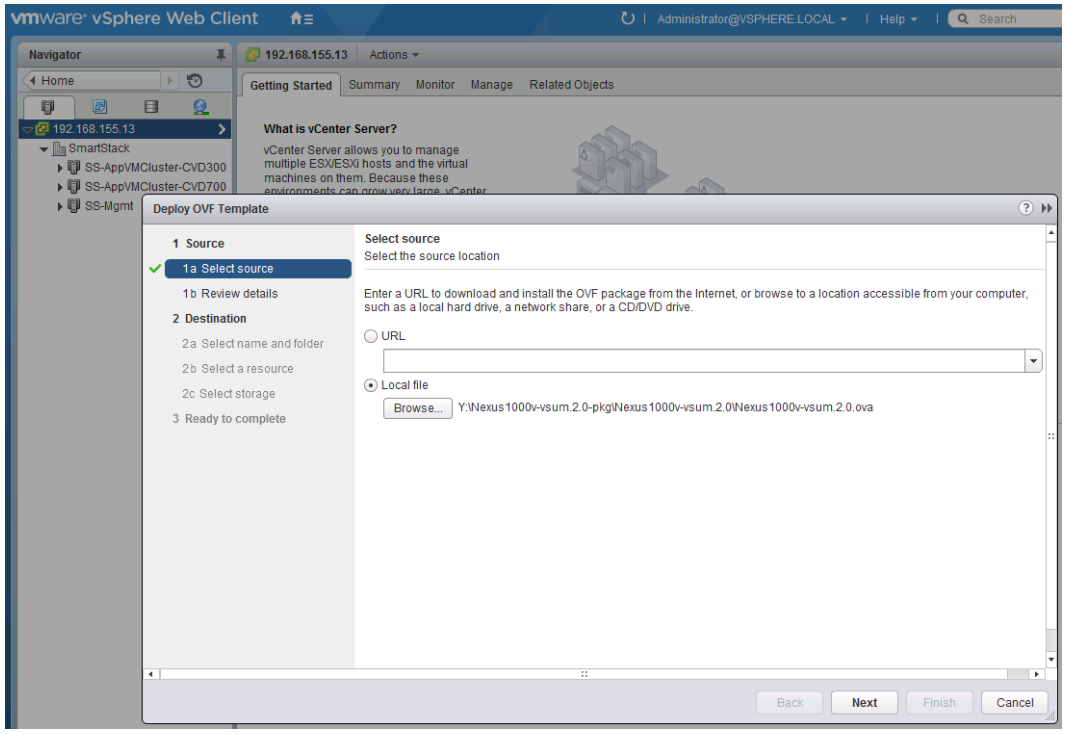

4. Review details and click Next to continue.

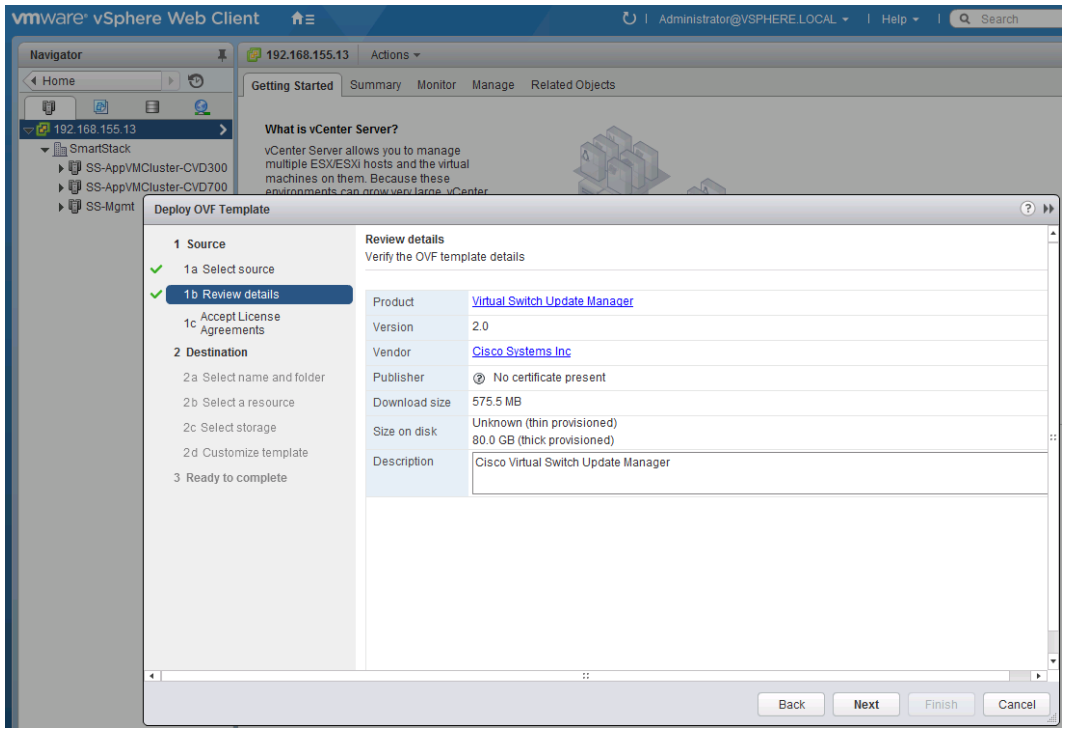

5. Accept License Agreement and click Next to continue.
6. Specify the name of the VM and datacenter where it should be deployed. Click Next to continue.

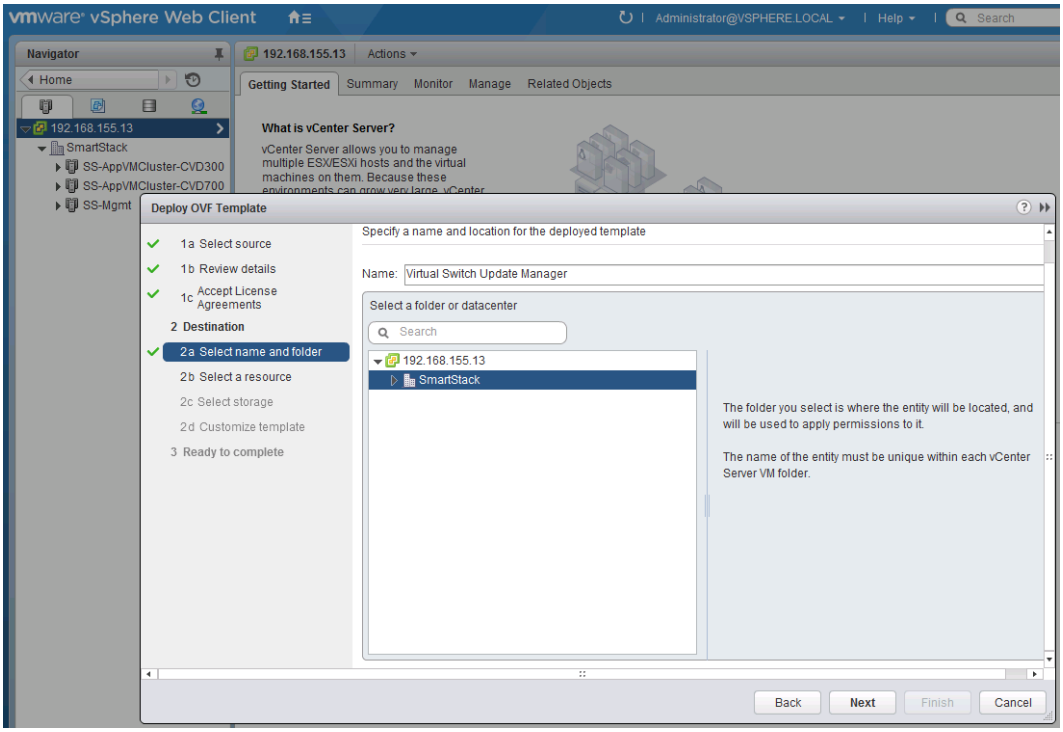

7. Specify the infrastructure management cluster where the VSUM should be deployed. Click Next to continue.

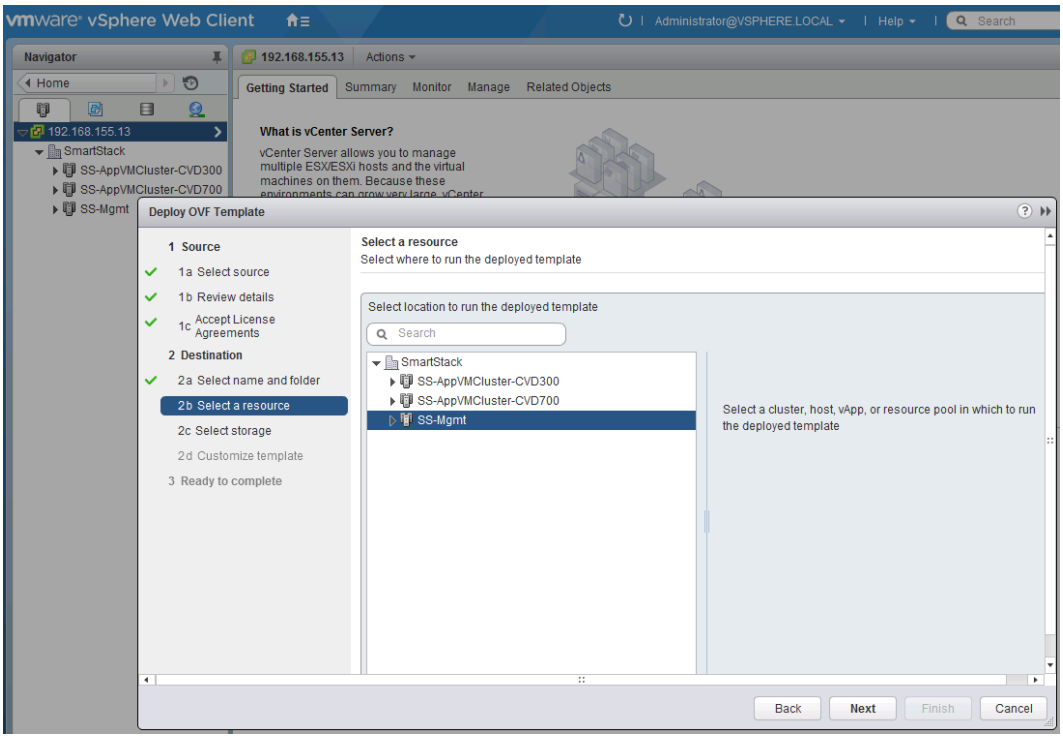

8. Specify the datastore where the VSUM should reside. Click Next to continue.

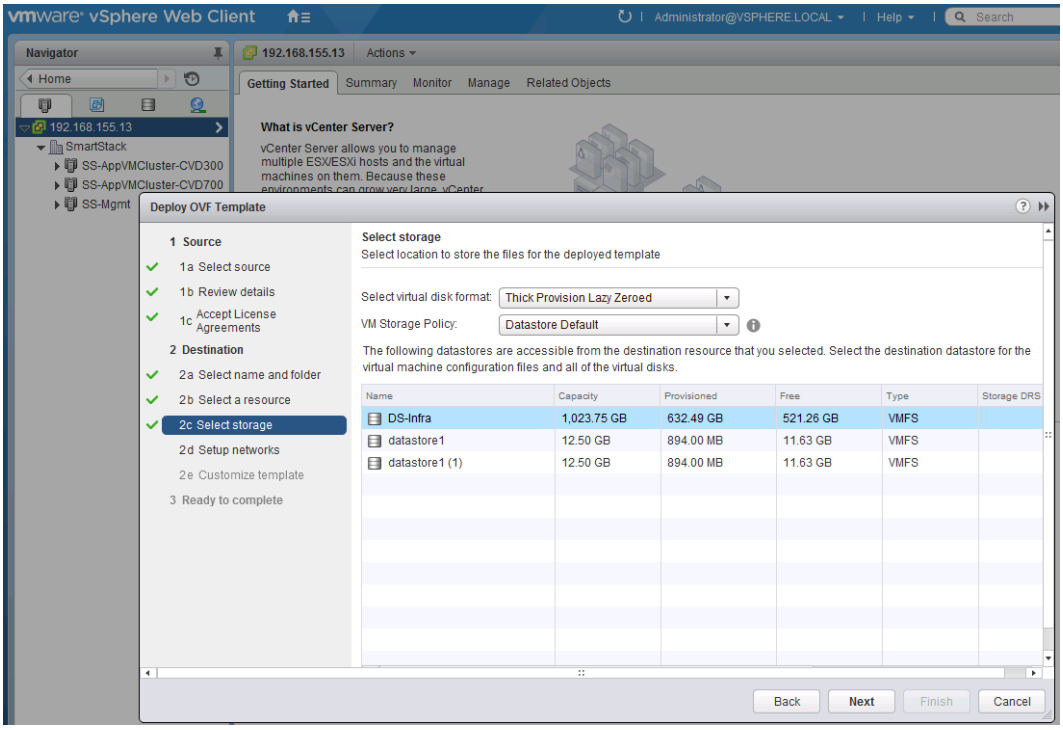

9. Specify the management network that the VSUM should reside. Click Next to continue.

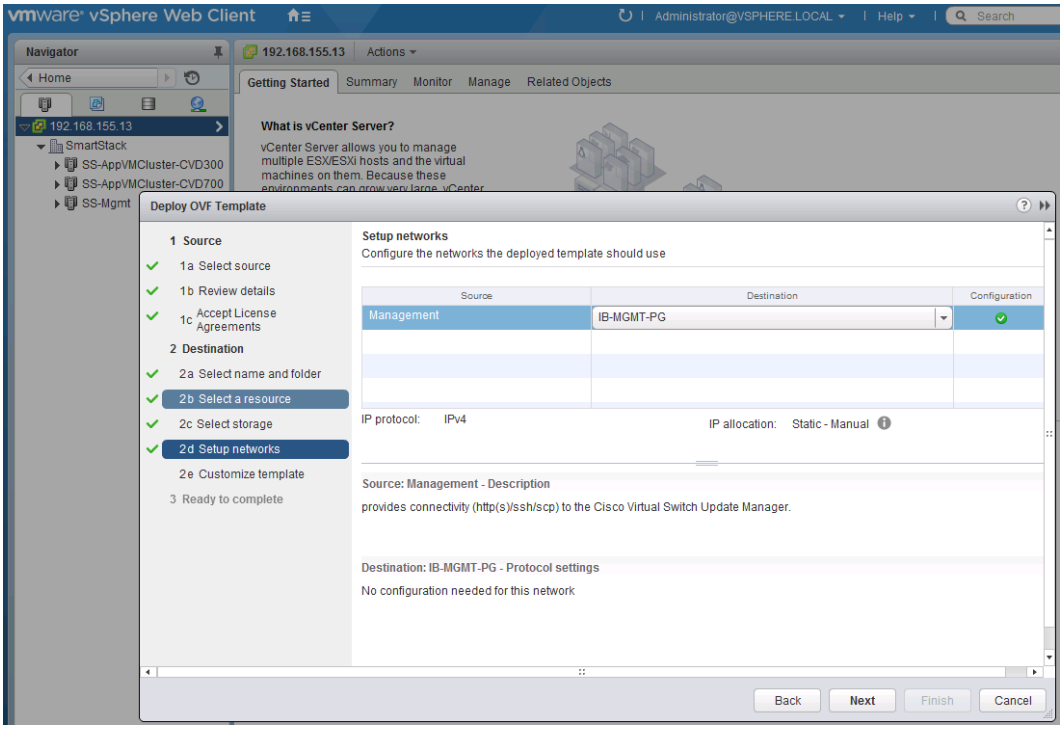

10. Specify the IP address, Mask, Default GW, DNS and vCenter info. Click Next to continue.

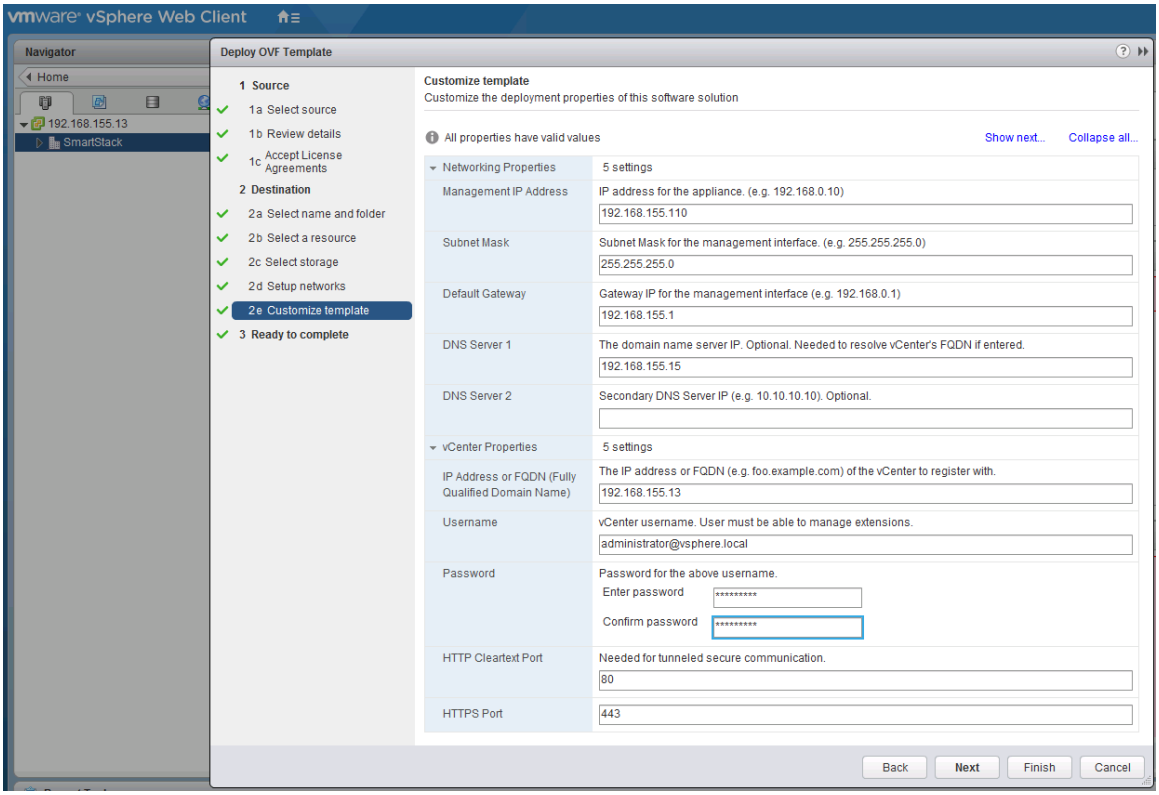

11. Review the configuration summary, select Power ON after deployment and click Finish to start deployment of VSUM. After the VSUM boots up, the plug-in will register with the vCenter after a few minutes. Verify that the plug-in is registered with vCenter by going to Home > Administration > Client Plug-ins in the vSphere Web client. Also, select Home tab and verify that the VSUM appears as an icon in the Inventories section. For any issues at this step, refer to the Cisco Virtual Switch Update Manager Troubleshooting Guide on cisco.com for assistance.

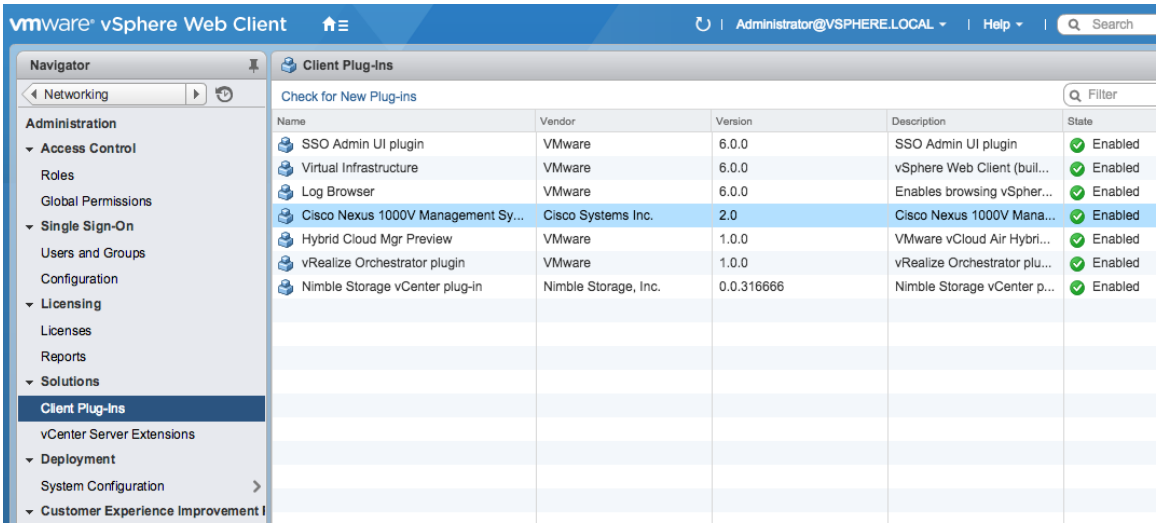

Install Cisco Nexus 1000V Virtual Supervisor Module (VSM) using Cisco VSUM

The VSUM will deploy the VSM primary and secondary to the ESXi hosts through the GUI install. You will have a VSM primary running on 1 ESXi host and a secondary running on the other ESXi host. Both of these are installed at them same time through the host selection. Complete the following steps to deploy the VSM.

To install the plug-in, click on the Download Client Integration Plug-in link at the bottom of the browser used to launch the VMware vSphere Web Client. Installing Adobe Flash may also be required.

1. Launch vSphere Web client (https://<vCenter_IP>:9443/vsphere-client) and login.
2. Select Home tab and click Cisco Virtual Switch Update Manager icon.

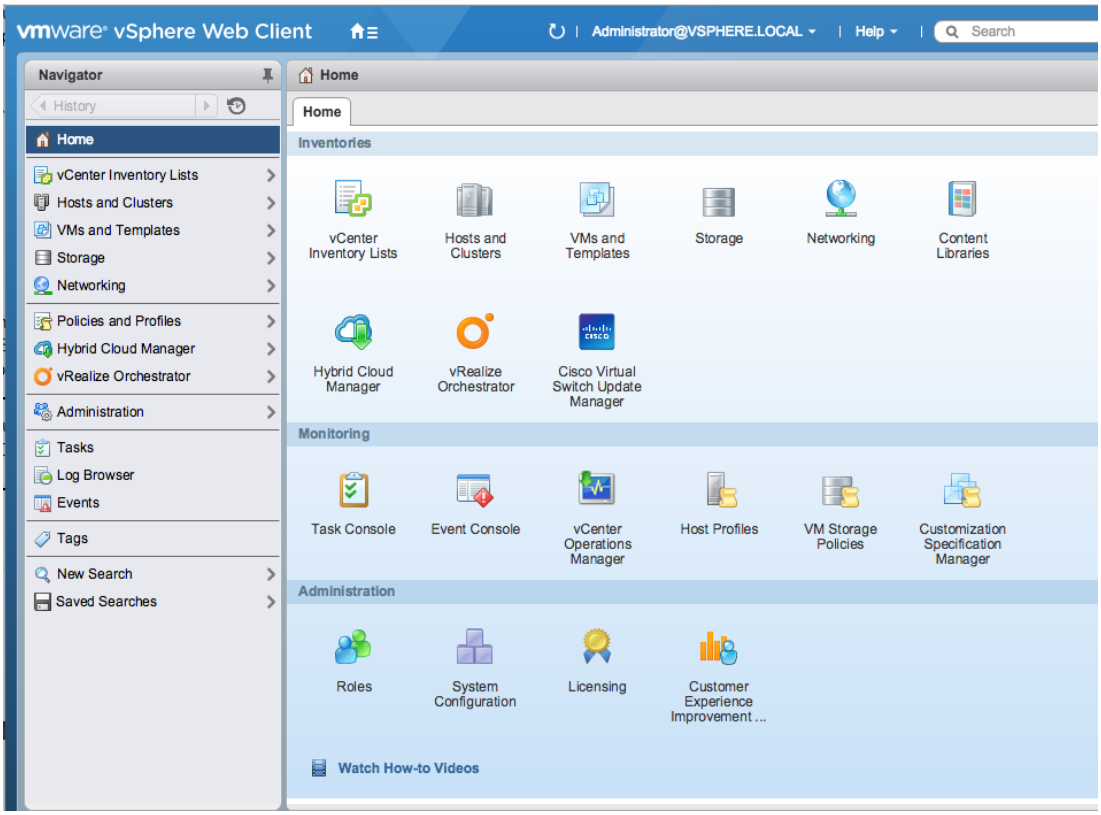

3. Select Cisco Nexus 1000V and click Install.

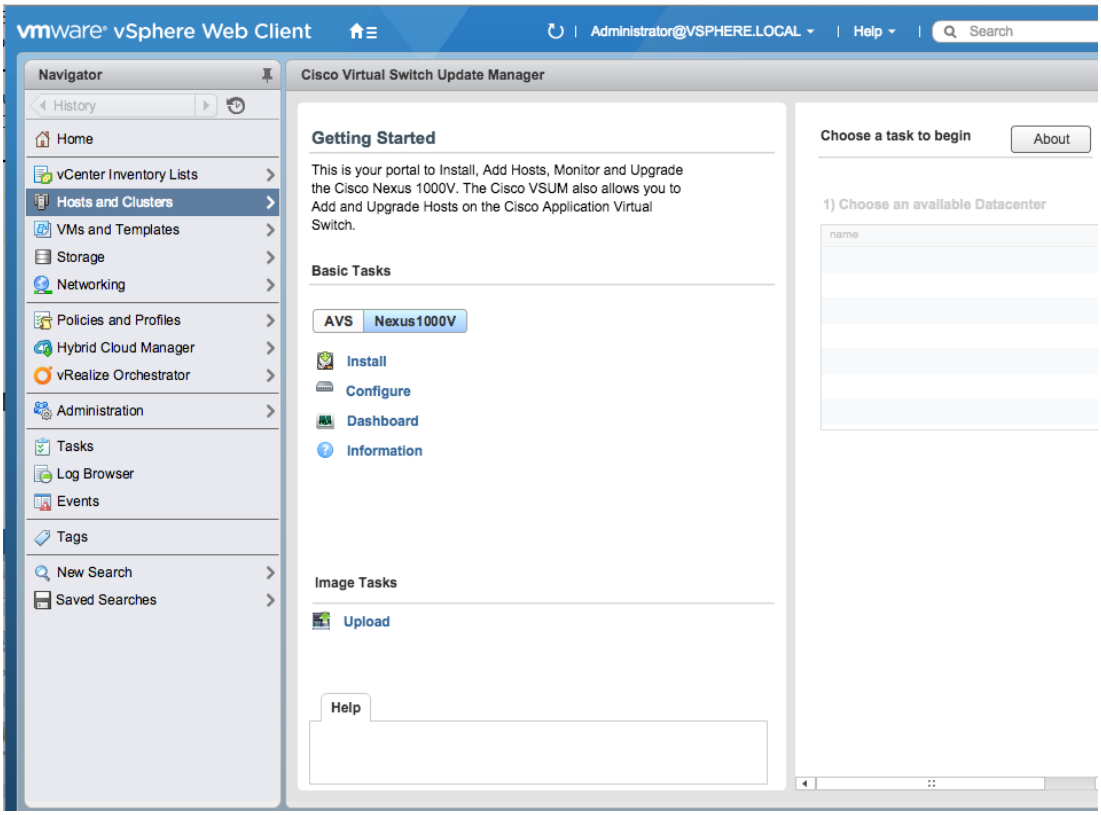

4. Select SmartStack datacenter on the right side of the screen.

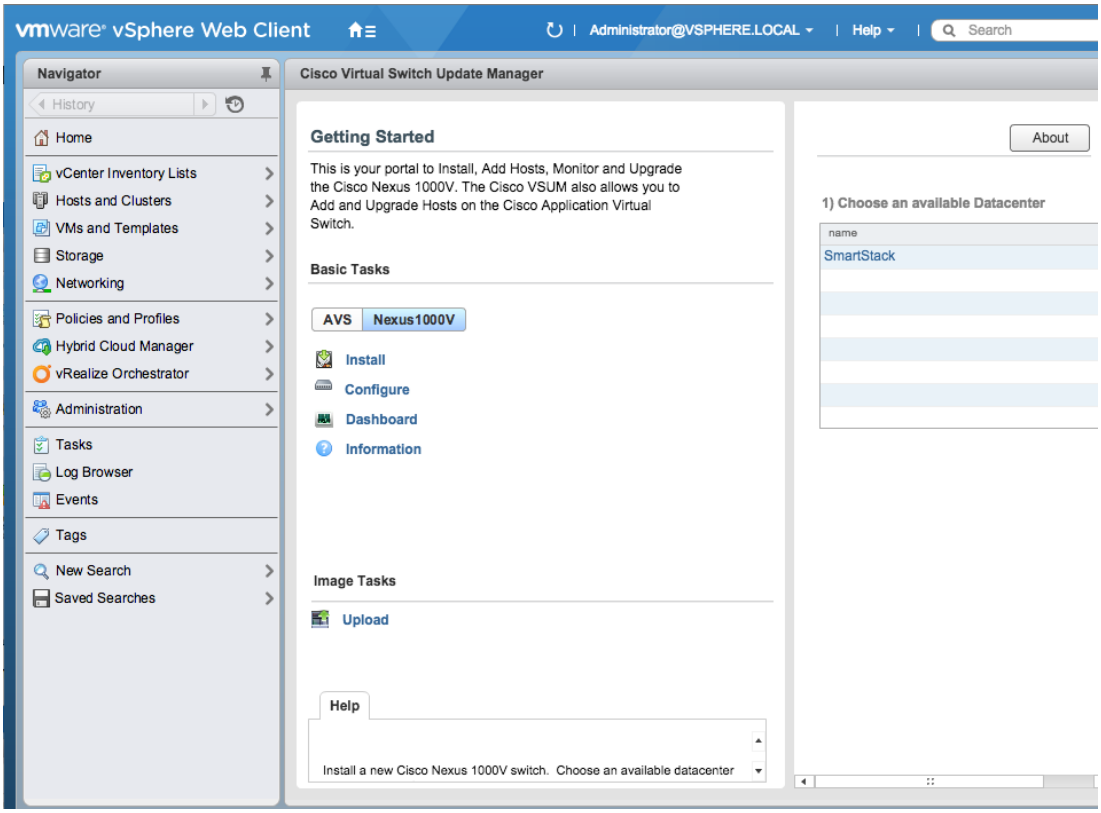

5. Keep the default for deploy new VSM and High Availability Pair. Follow the link to upload VSM version if the version is not on the list even after a refresh.

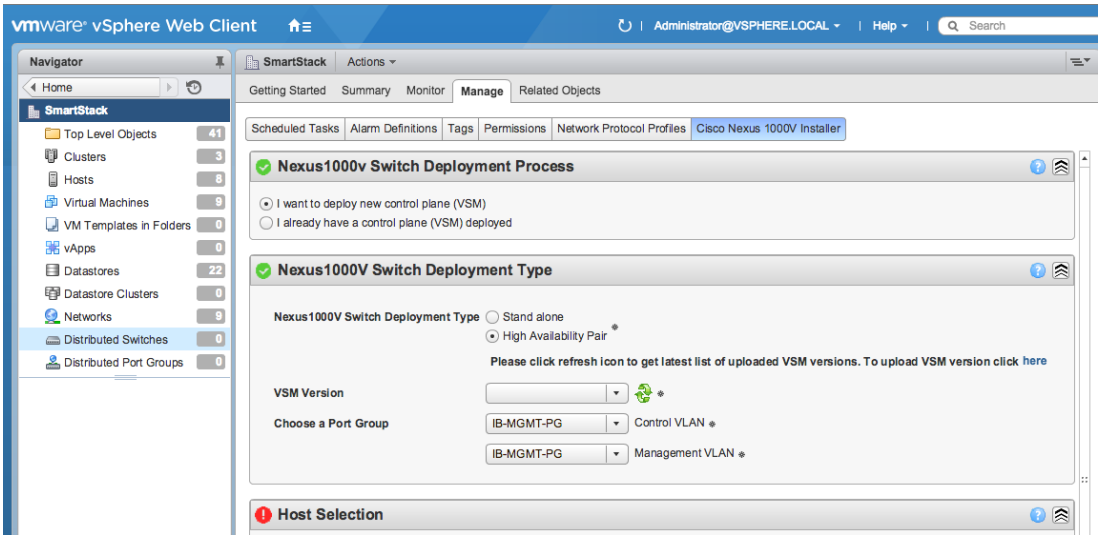

6. Click Upload VSM bundle so that it shows up in the VSM version drop down list. Click Choose File button.

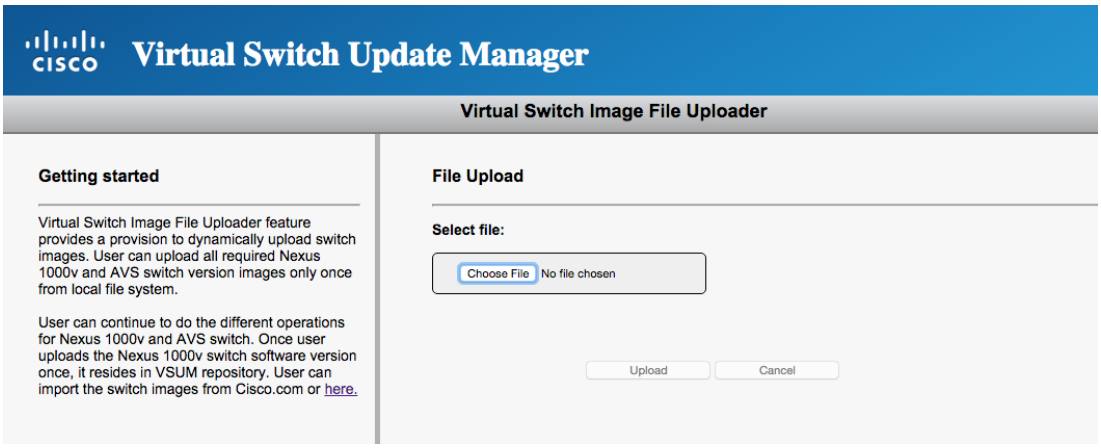

7. Select Cisco Nexus 1000V software version to use from www.cisco.com. Click Download button.

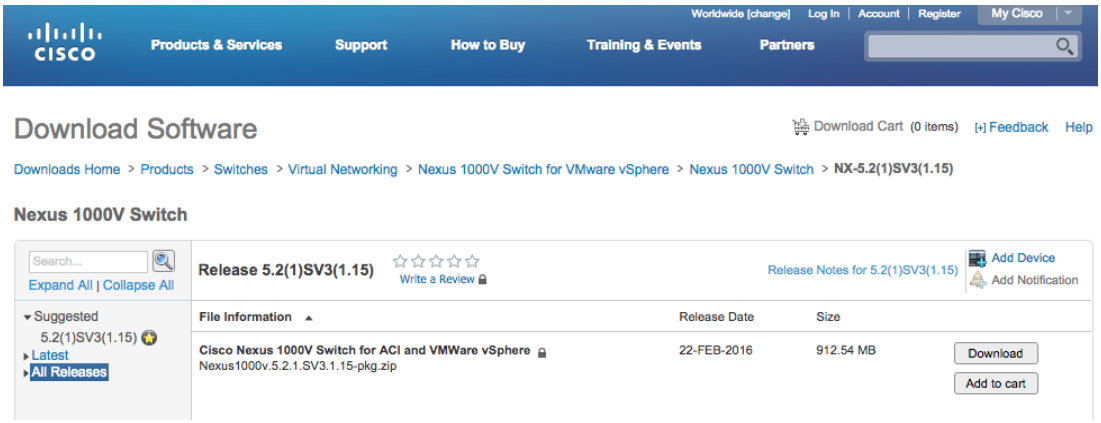

8. Once download completes, click Choose File button and browse to select the .zip image downloaded. Click Upload button to upload the Cisco Nexus 1000V image to VSUM.

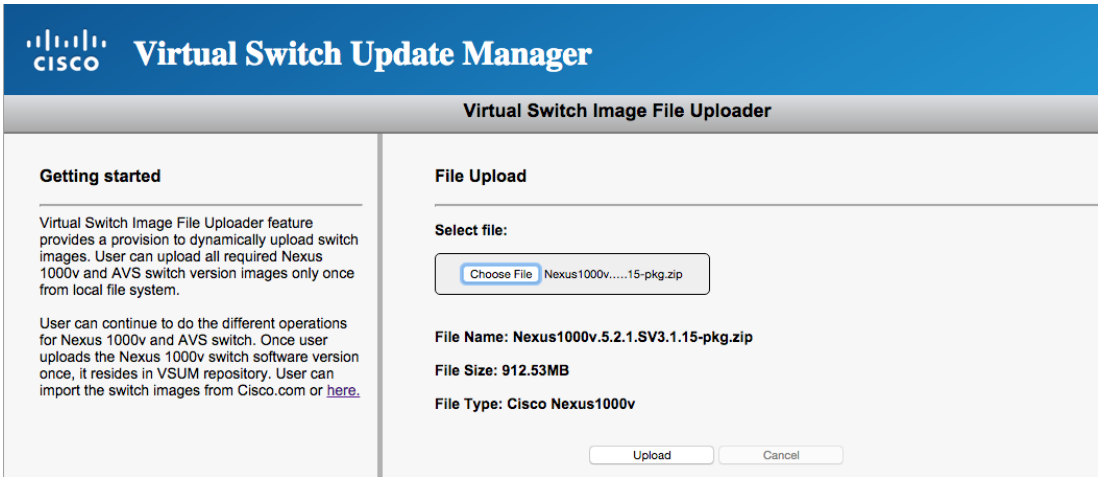

9. Return to the Cisco Nexus 1000V Installation and select the uploaded VSM version and specify the port group for the control and management VLANs. Use the same port group (for example, IB-MGMT-PG) for both the control and management VLAN.

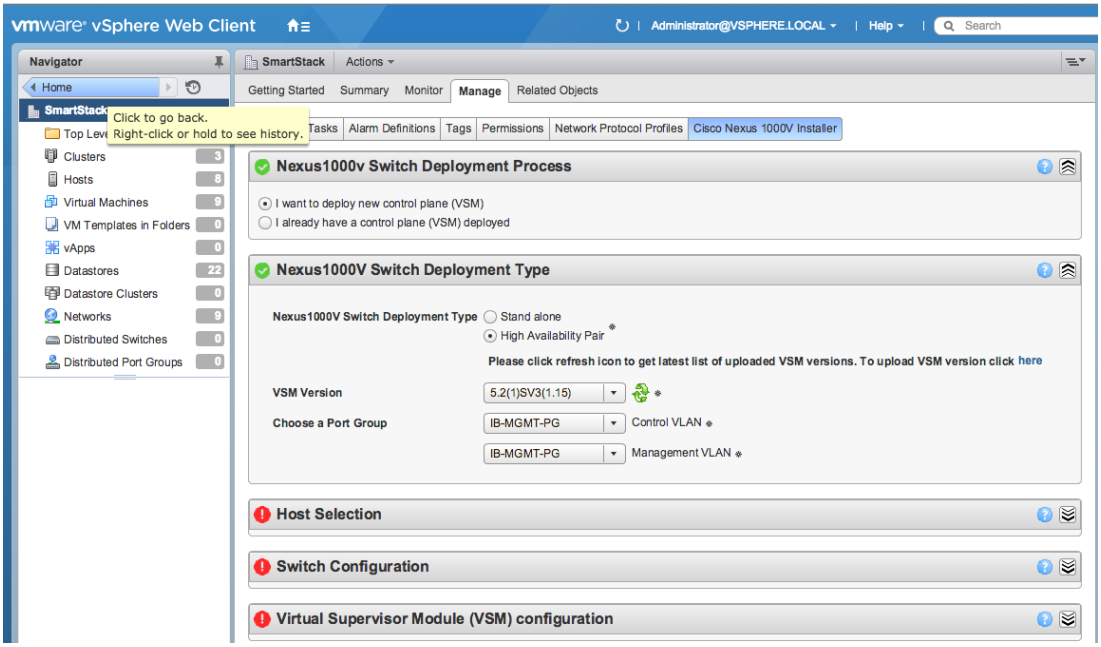

10. Go to the Host Selection section and click Suggest button.

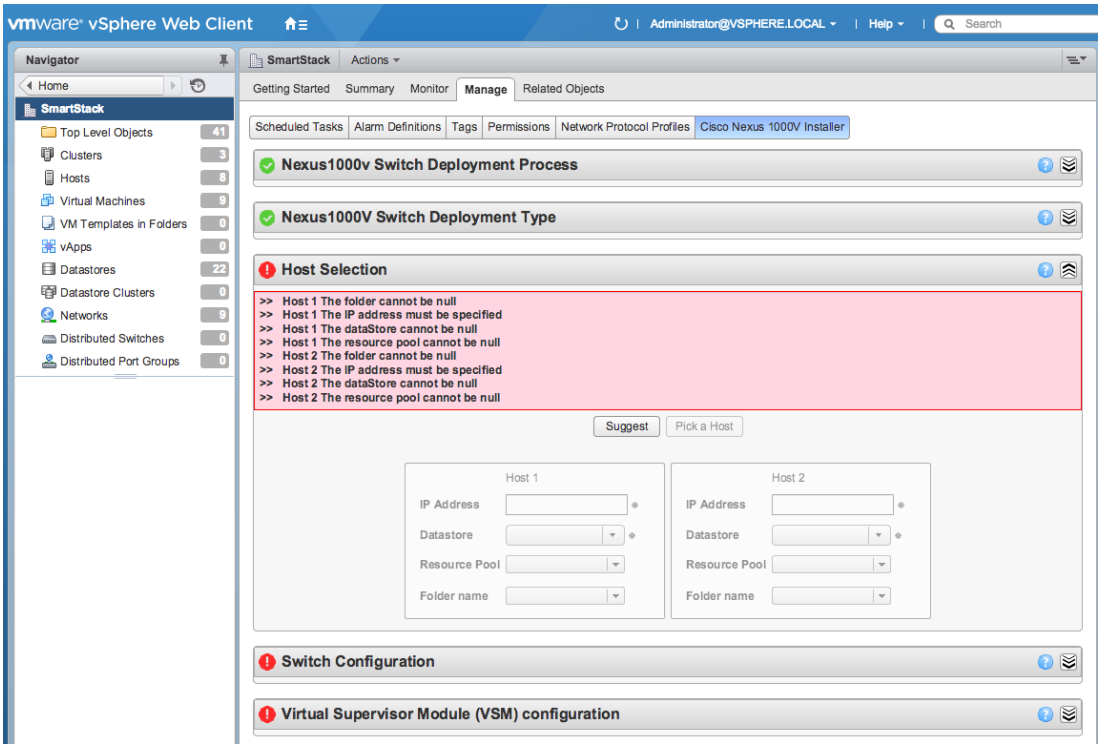

11. Click on the Pick a Host button to select servers to host the primary and secondary VSMs.

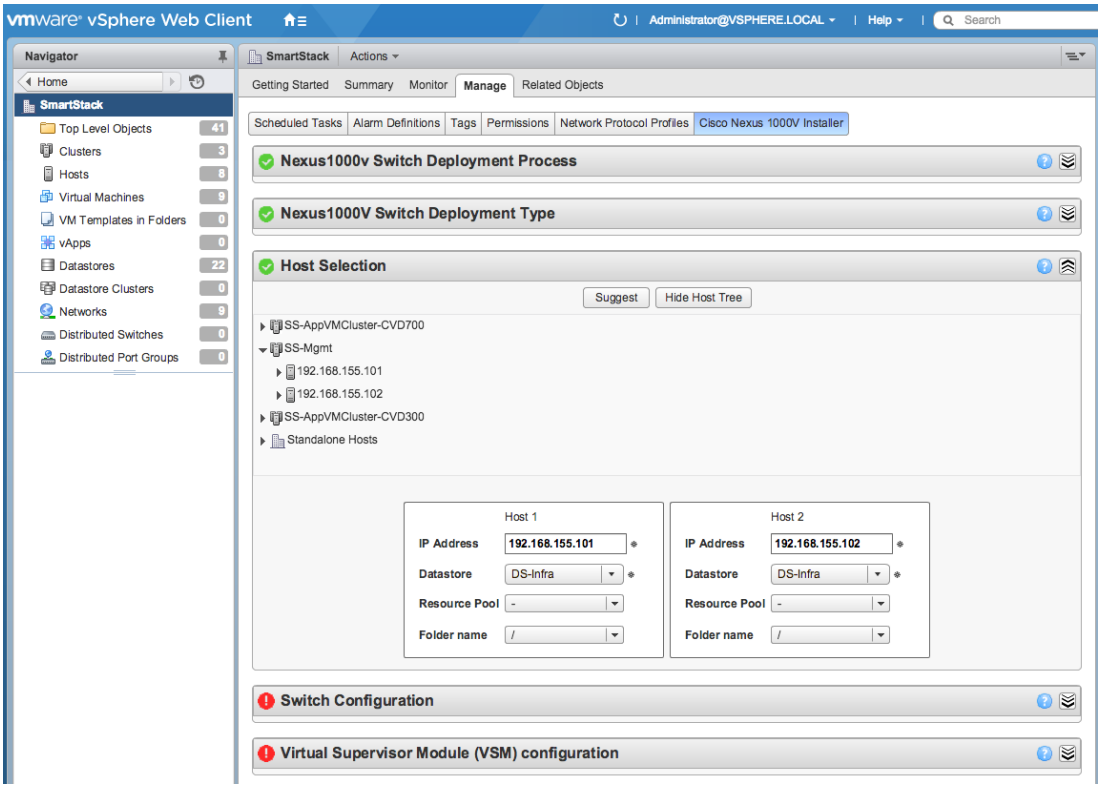

12. In the Switch Configuration section, specify the Domain ID.

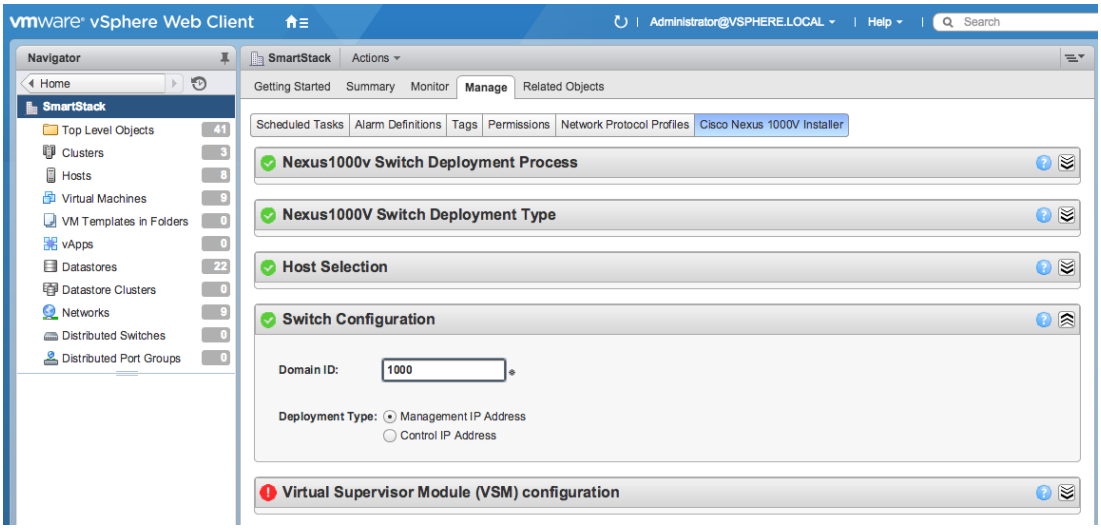

13. In the VSM Configuration section, fill out the IP address, Subnet, GW and the password the VSM should use.

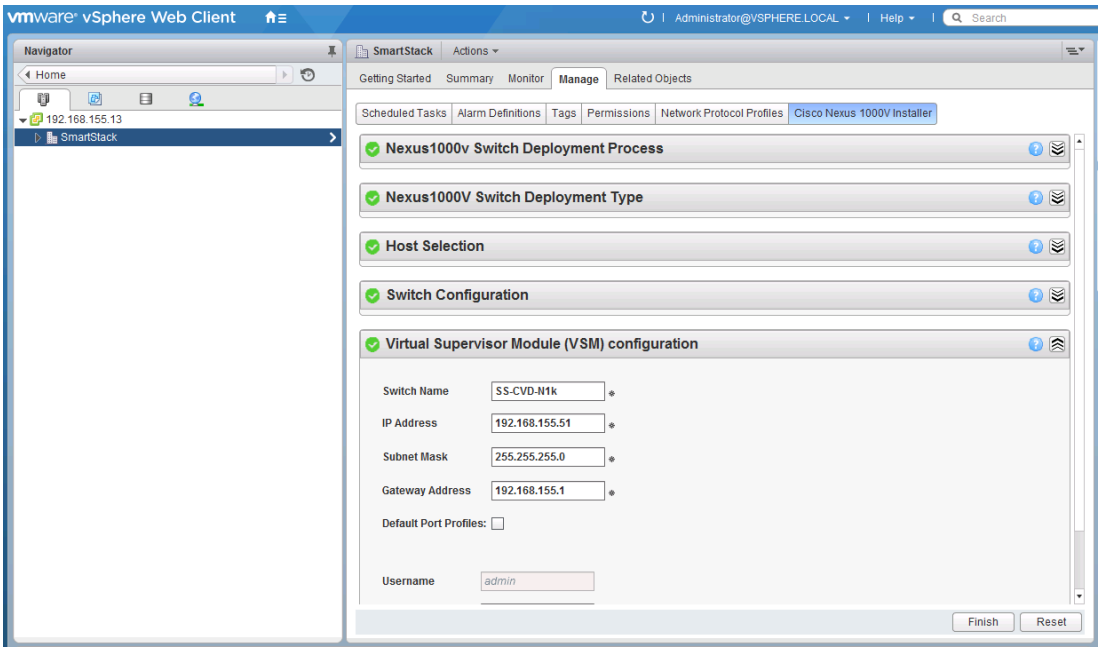

14. Click Finish to begin the installation of primary and secondary VSM on separate hosts.

Configure Primary VSM

Complete the following configuration steps to do a basic configuration of the primary VSM.

1. Using an SSH client, login to the primary Cisco Nexus 1000V VSM as admin.
2. Run the following configuration commands.

```
config t

ntp server 192.168.155.254 use-vrf management
```

```
vlan 12
name IB-MGMT-VLAN
vlan 3000
name vMotion-VLAN
vlan 950
name APP1-VM-VLAN
vlan 951
name APP2-VM-VLAN
vlan 2
name Native-VLAN
exit

port-profile type ethernet system-uplink
vmware port-group
switchport mode trunk
switchport trunk native vlan 2
switchport trunk allowed vlan 12, 950-951, 3000
channel-group auto mode on mac-pinning
no shutdown
system vlan 12, 950,951, 3000
system mtu 9000
state enabled
exit

port-profile type vethernet IB-MGMT-VLAN
vmware port-group IB-MGMT-PG
switchport mode access
switchport access vlan 12
no shutdown
system vlan 12
state enabled
exit

port-profile type vethernet vMotion-VLAN
vmware port-group vMotion-PG
switchport mode access
switchport access vlan 3000
no shutdown
system vlan 3000
state enabled
exit

port-profile type vethernet APP1-VM-VLAN
vmware port-group APP1-VM-PG
switchport mode access
```

```
switchport access vlan 950
no shutdown
system vlan 950
state enabled
exit

port-profile type vethernet APP2-VM-VLAN
vmware port-group APP2-VM-PG
switchport mode access
switchport access vlan 951
no shutdown
system vlan 951
state enabled
exit

port-profile type vethernet n1kv-L3
capability l3control
vmware port-group
switchport mode access
switchport access vlan 12
no shutdown
system vlan 12
state enabled
exit

copy run start
```

Migrate ESXi Hosts from vSphere vSwitch to Cisco Nexus 1000V

Complete the following steps on each ESXi host to migrate the host from vSwitch to Cisco Nexus 1000V distributed switch.

1. Launch vSphere Web client (https://<vCenter_IP>:9443/vsphere-client) and login.
2. Select Home tab and click Cisco Virtual Switch Update Manager icon.

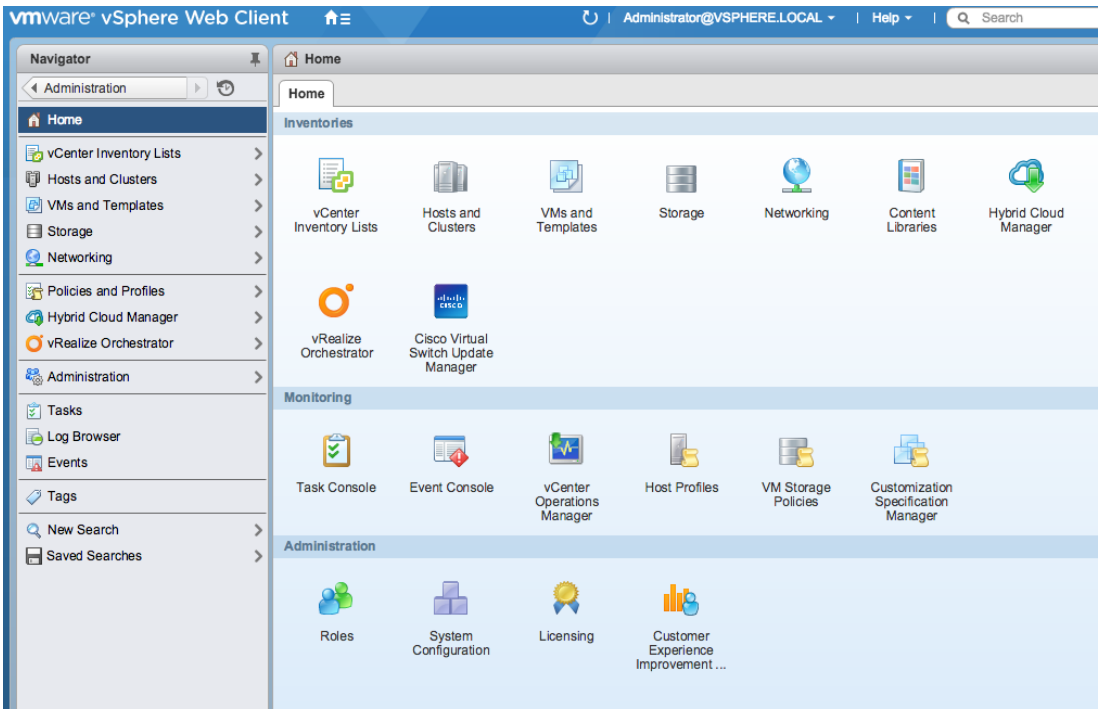

3. Click Cisco Nexus 1000V on the Cisco Virtual Update Manager pane to the right.

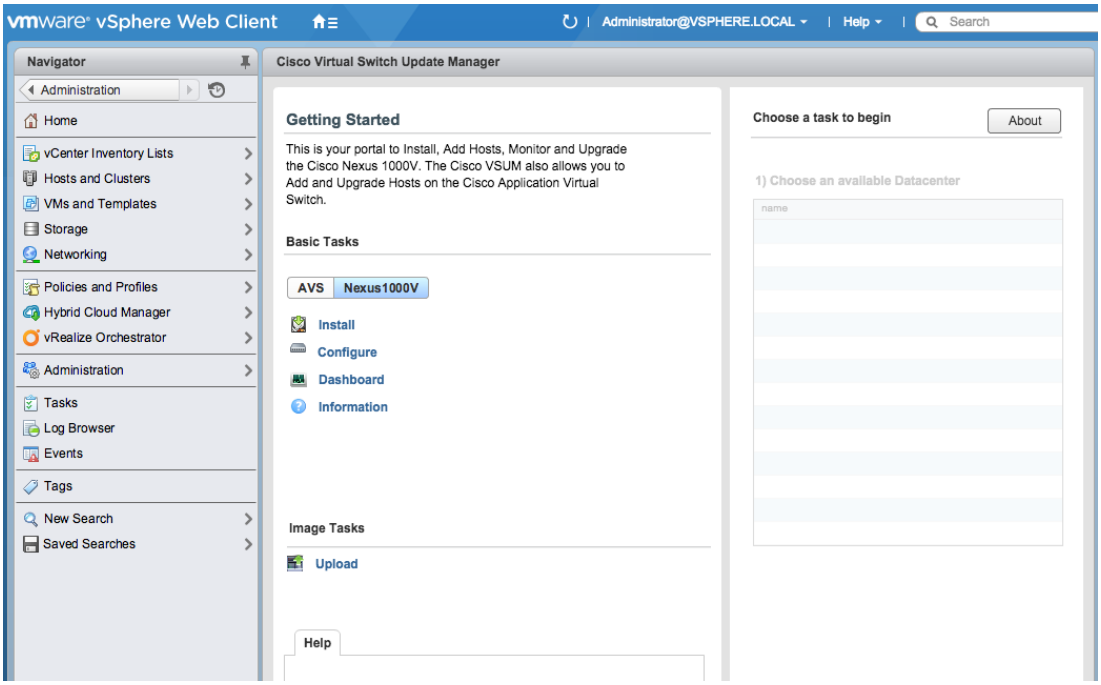

4. Click Configure and select the Datacenter (for example, `SmartStack`) .

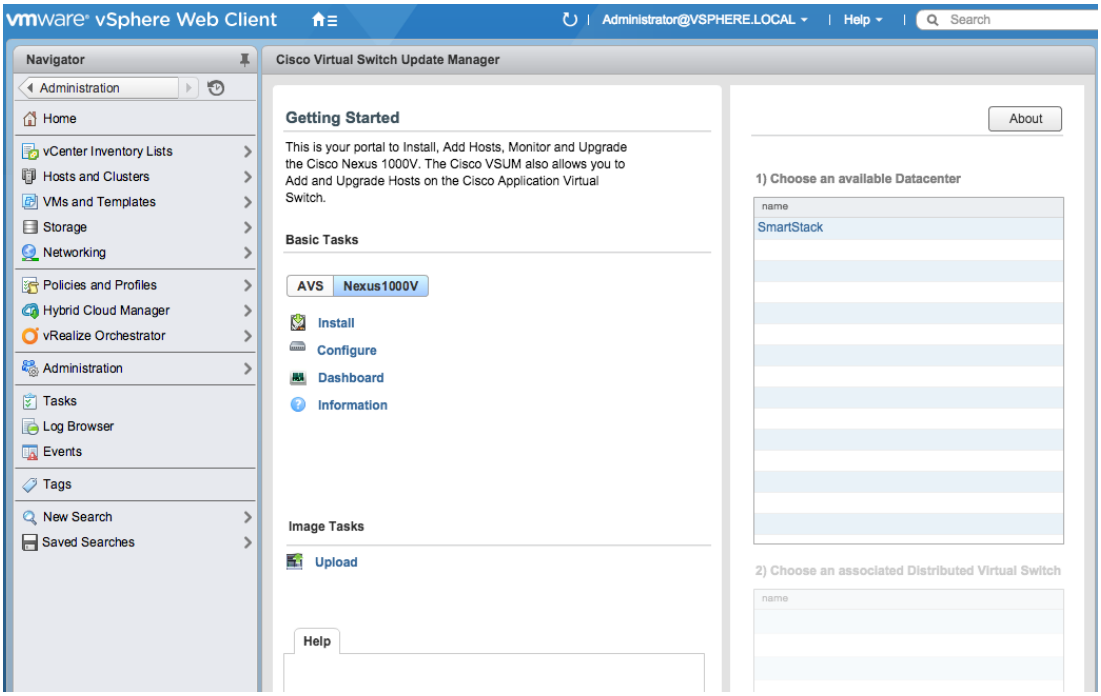

5. Select the associated Distributed Virtual Switch (for example, `SS-CVD-N1k`) and click Manage.

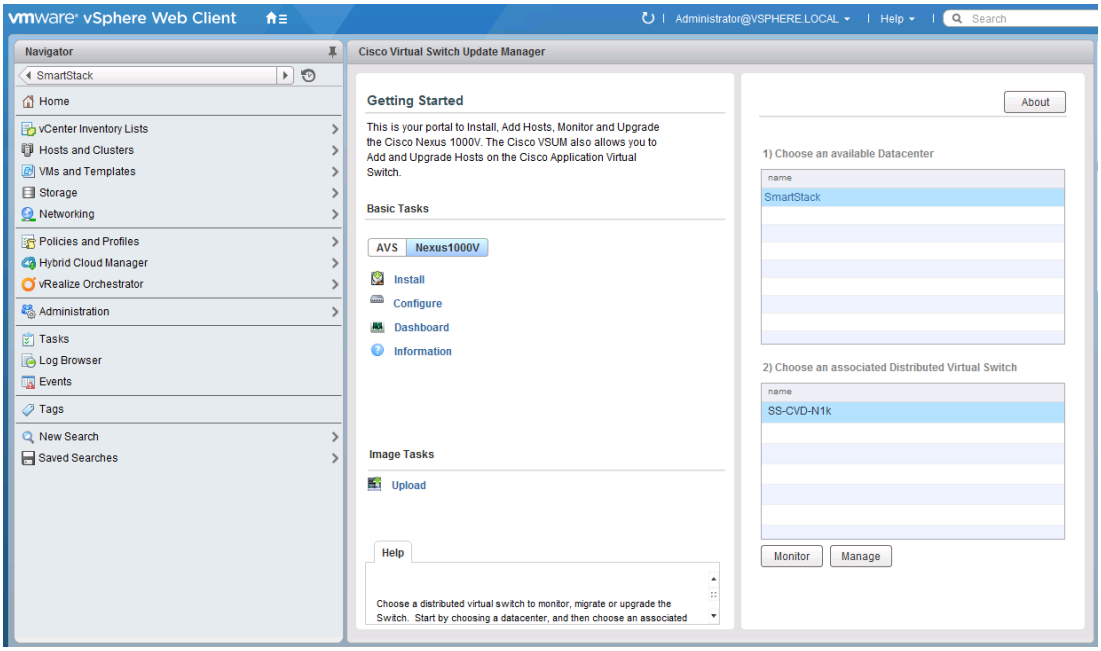

6. In the Manage Window pane to the right, click Cisco Nexus 1000V > Add Host tab. Select host(s) to migrate and click Suggest.

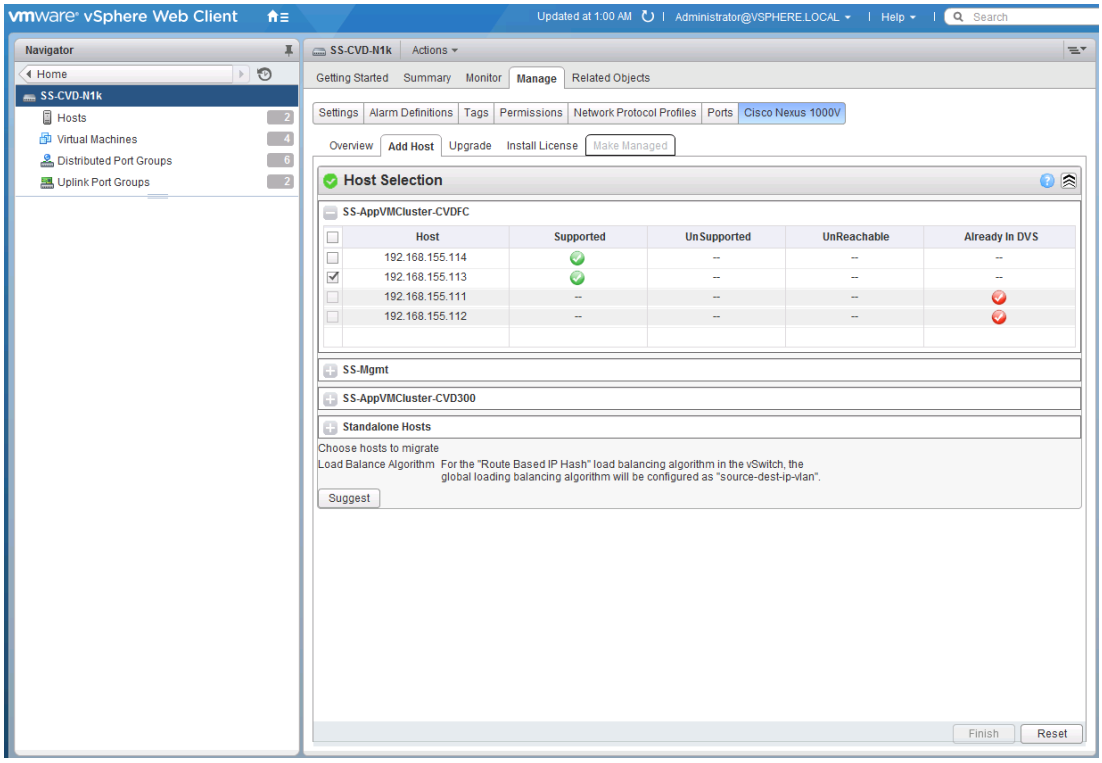

7. Minimize the Port Profile Editor section. In the Physical NIC Migration section, select vmnic1 to migrate. The first vmnic is temporarily left on vSwitch0 and will be migrated over once this migration is successful. For vmnic1, select the system-uplink Profile from the drop-down list.

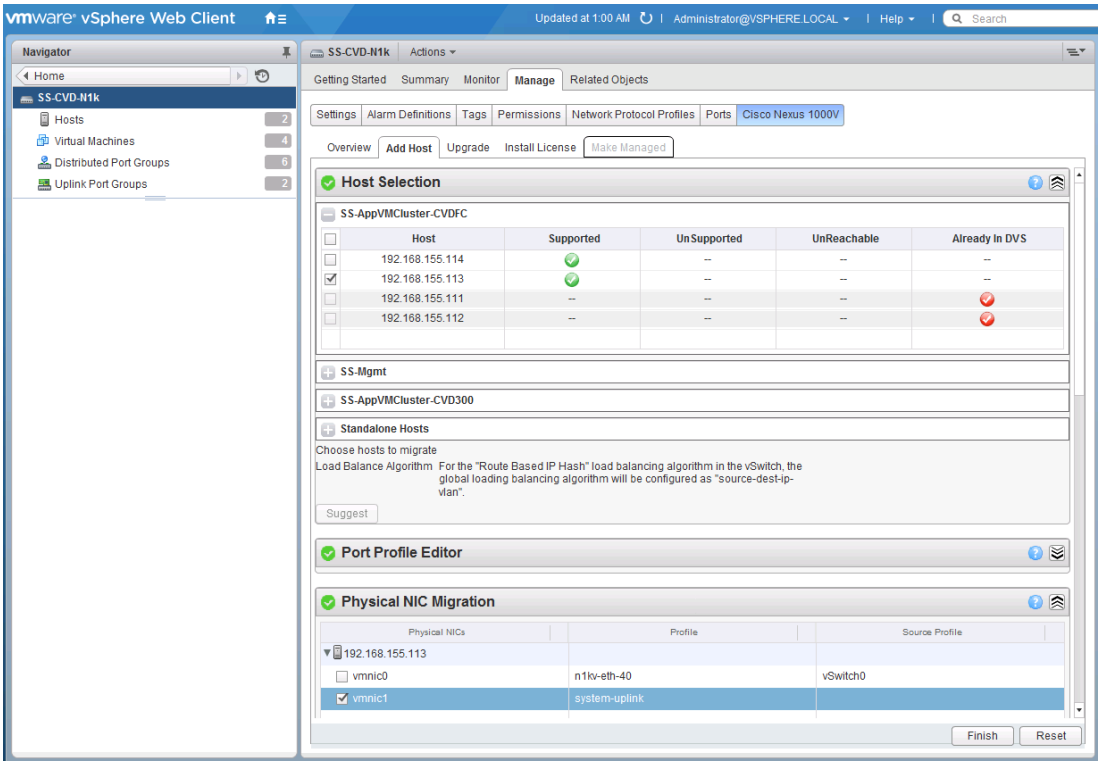

8. In the VM Kernel NIC Setup section, select the vmk(s) to migrate. Unselect the 3rd vmk (vmk2) that was specifically created for this migration. The profiles for vmk0 and vmk1 should be as shown.

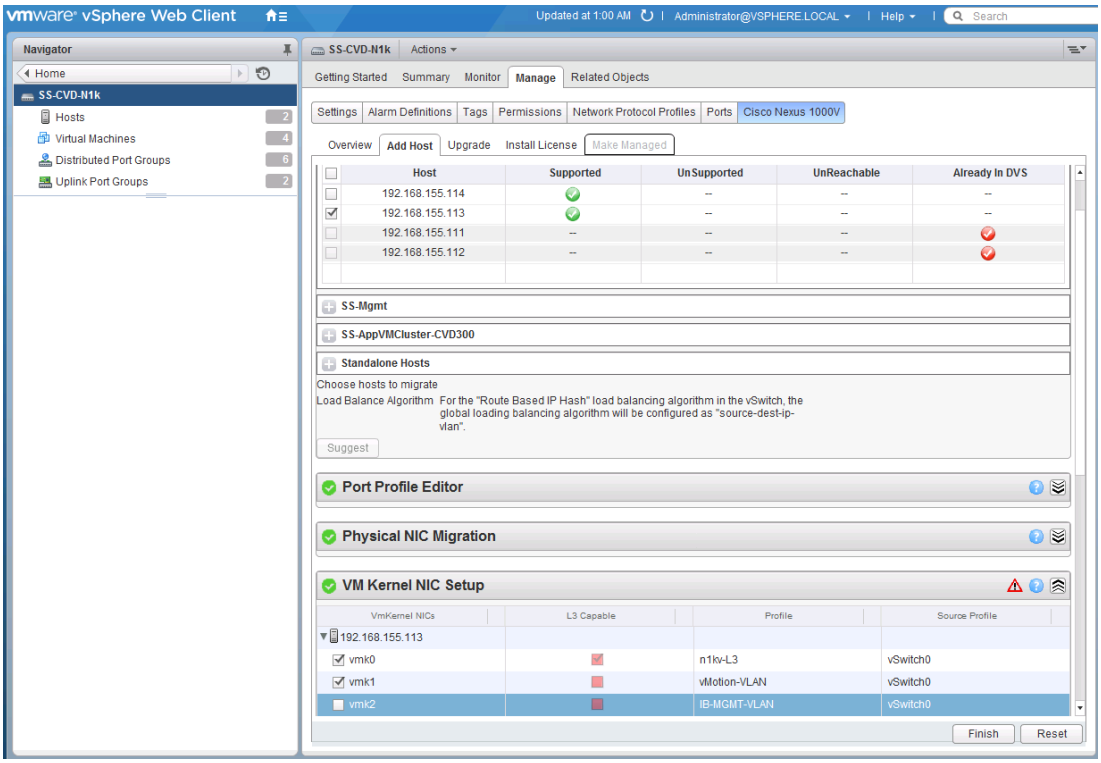

9. In the VM Migration section, select VMs (if any) to migrate. Click the button next to the virtual machine to expand the target profile and choose the correct profile (for example, IB-MGMT-VLAN). Repeat this for each Virtual Machine.

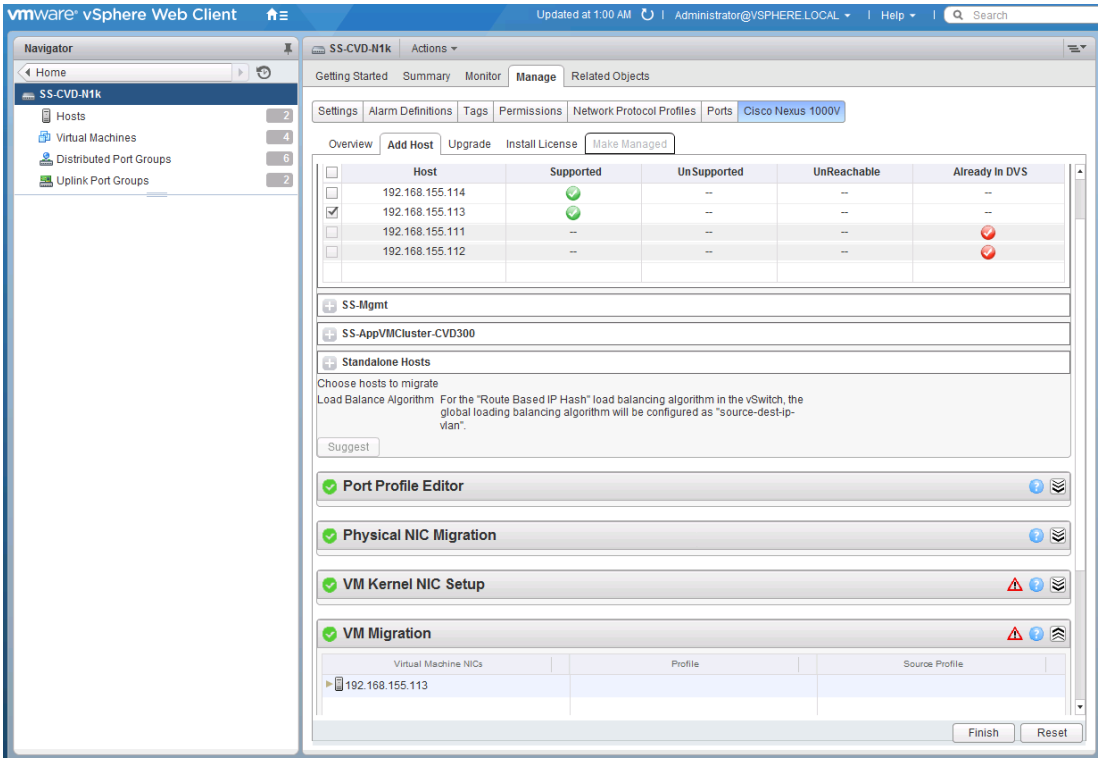

10. Click Finish to initiate the migration. View the status of the migration by clicking on More Tasks at the bottom right corner of the window. For additional help, refer to the Cisco Virtual Switch Update Manager Troubleshooting Guide for any issues seen during migration.

If a migration using VSUM is not successful for any reason, the migration can be done by manually upgrading the host using the `esxcli software vib insall -v|-b <file_name>` using the VEM vib or zip file. This file will need to be uploaded to the host datastore from VMware vCenter. Also check the Cisco Nexus 1000V and VMware Compatibility Information on cisco.com to confirm the version is compatible. Example: `esxcli software vib install -d /vmfs/volumes/Host1-DS0/VEM600-201602260119-BG-release.zip`

One issue, CSCuw55307, was seen when attempting a migration from the VMware vCenter web client by directly selecting the Cisco Nexus 1000V switch and doing a rick click to Add and Manage Hosts. For more details on this issue, login to cisco.com and do a Bug Search using the ID provided.

Migrate the Second Uplink on the host from vSwitch to Cisco Nexus 1000V

Once the migration of the host to Cisco Nexus 1000V is complete, next step is to migrate the second uplink from vSwitch to Cisco Nexus 1000V for redundancy. To do this, first remove the uplink from vSwitch 0.

1. From vSphere web Client, navigate to Hosts and Clusters. Select Host from the inventory. Click Manage > Networking > Virtual Switches. Select vSwitch0 from the list of switches. Click on the 3rd icon in the menu above the list of switches to manage the physical adapters connected to the switch.

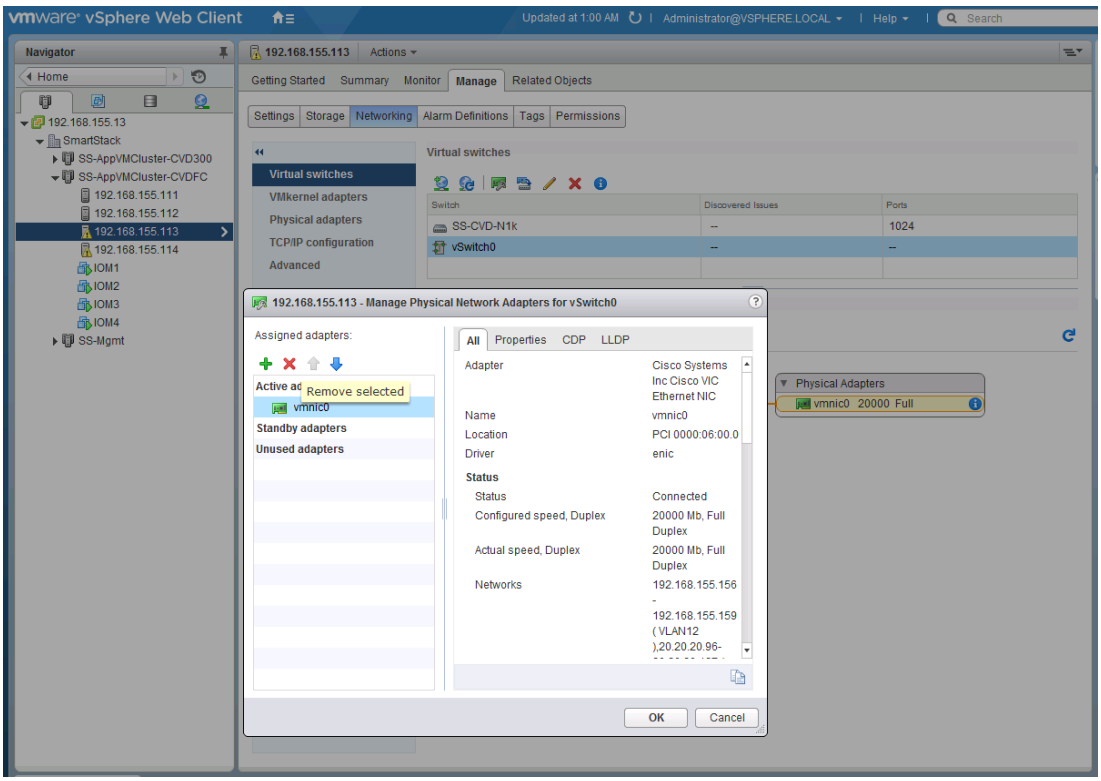

2. Select vmnic0 from the list and click ‘x’ to remove it from vSwitch0. Click Ok twice to confirm.
3. Now click on Manage > Networking > Physical Adapters. Right-click on vmnic0 and Click on the 1st icon (Add host networking) on the top menu to migrate vmnic0 to Cisco Nexus 1000V switch.

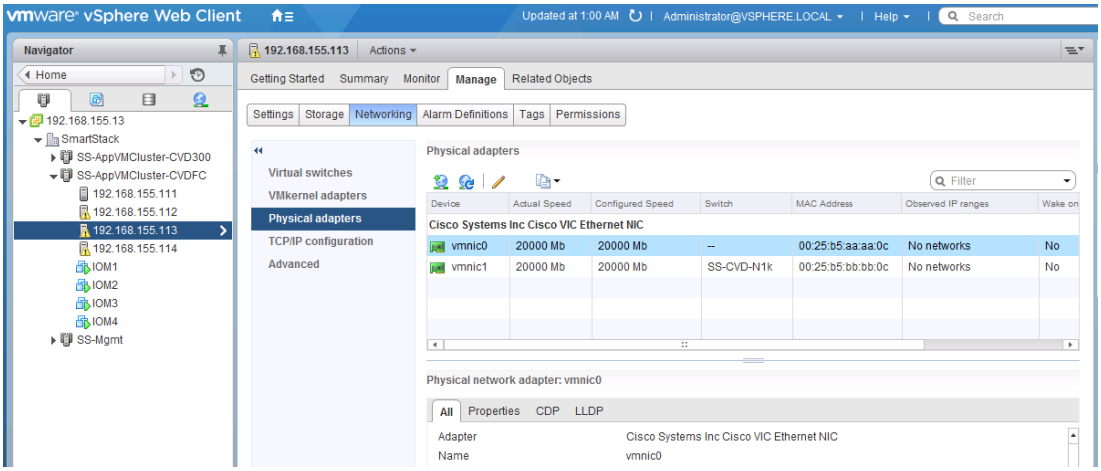

4. In the Add Networking window, select the Physical Network Adapter radio button for the connection type. Click Next to continue.

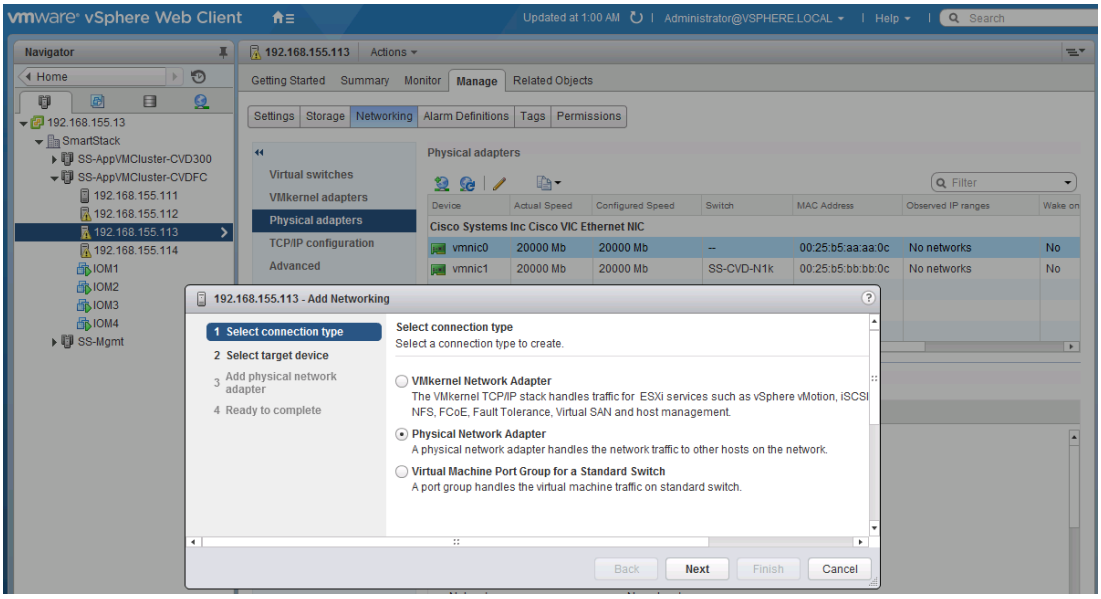

5. For the Select target device, select the radio button for Select an existing switch and click on Browse. In the Select Switch pop-up window, and then select the Cisco Nexus 1000V switch. Click Next to continue.

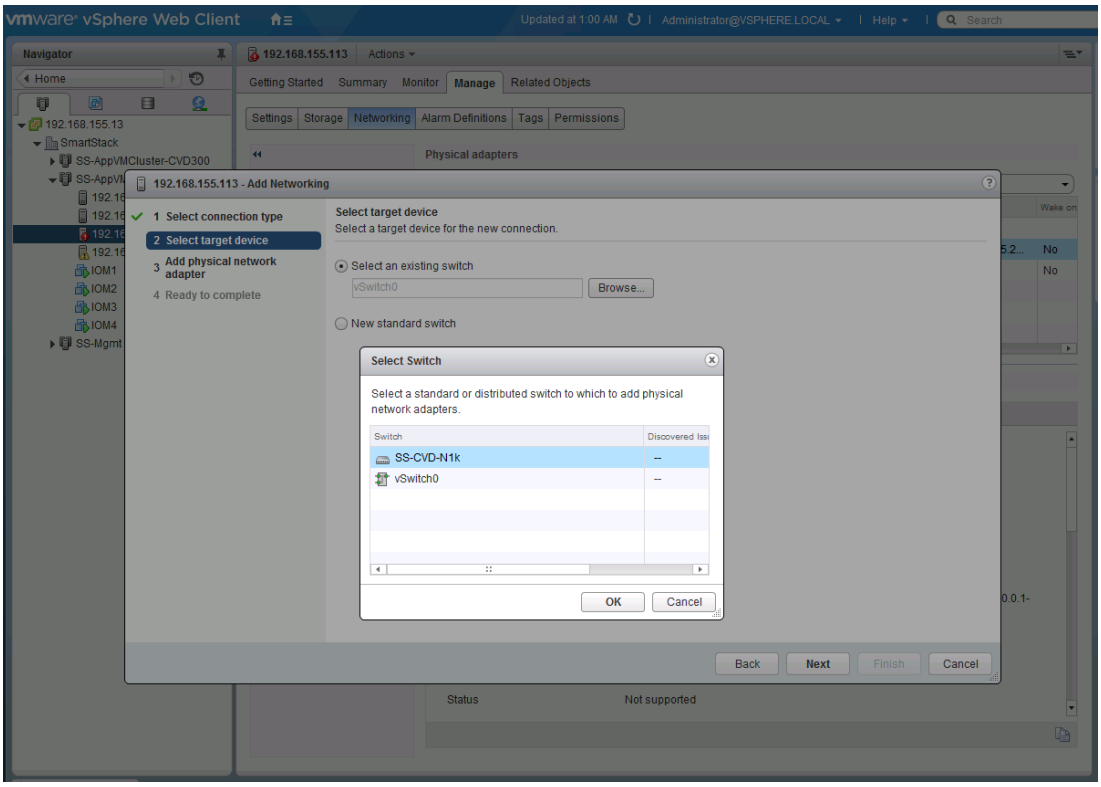

6. For the Add physical network adapter, click + and select the second Uplink name. Select system-uplink as the Uplink port group. Click Ok and then Next to continue.

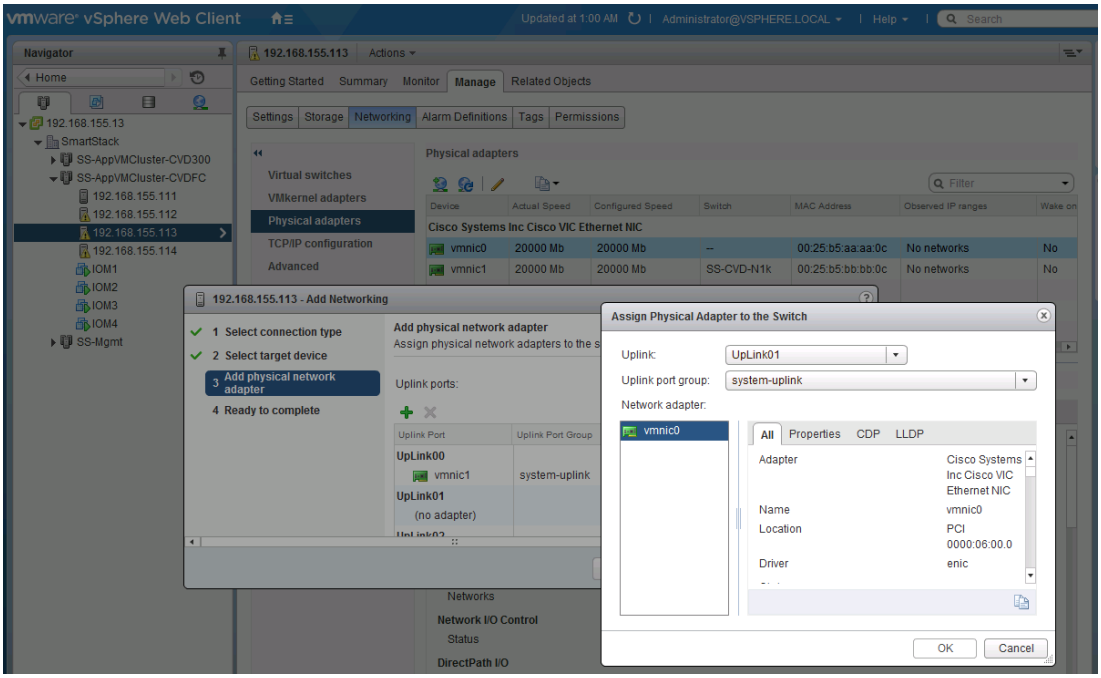

7. Click Finish to complete the migration. The vmnic0 should now show up as being part of the Cisco Nexus 1000V switch (for example, SS-CVD-N1k) as shown below.

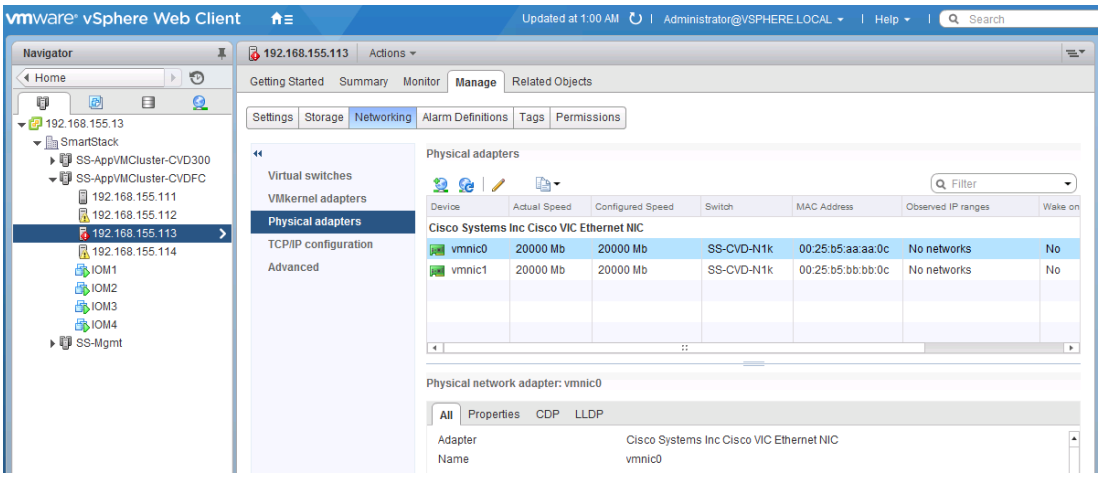

8. When the migration completes, select Manage > Settings > Topology to view the network connectivity through the Cisco Nexus 1000V.

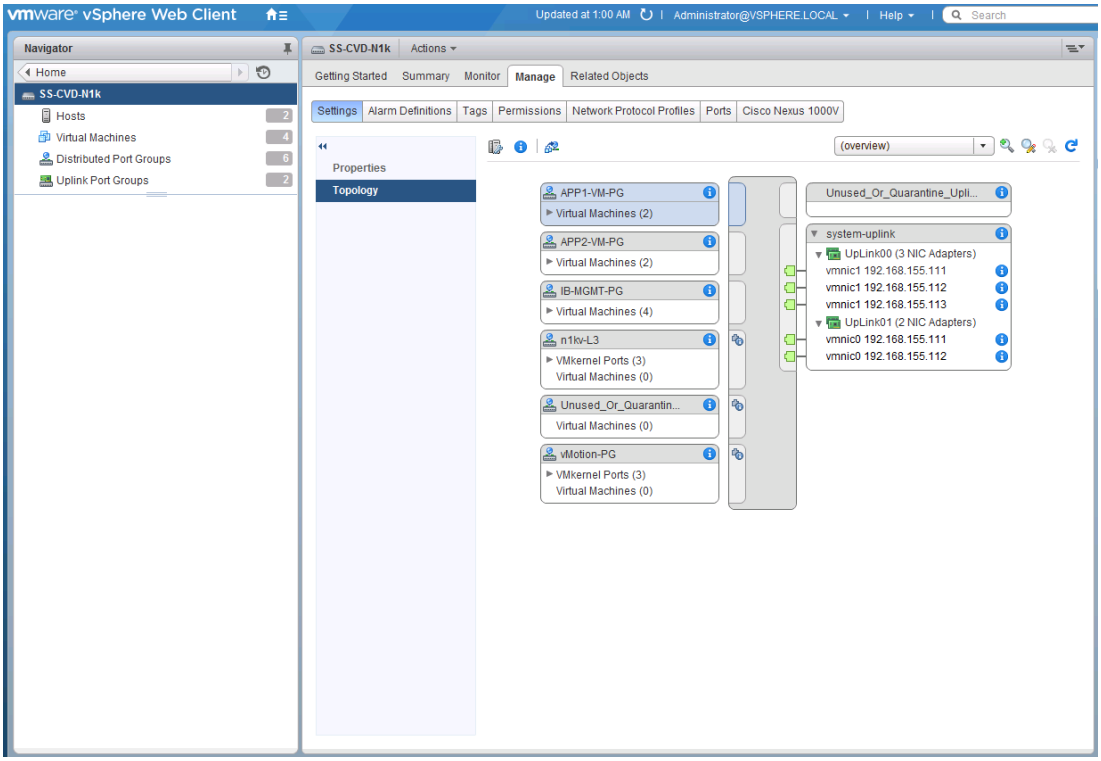

Remove vSphere vSwitch Components from Migrated ESXi Hosts

To complete the migration to Cisco Nexus 1000V distributed switch, complete the following steps on each migrated host to remove the vSwitch components remaining.

1. From vSphere web Client, navigate to Hosts and Clusters. Select Host from the inventory. Click Manage > Networking > Physical Adapters. Select vSwitch0 from the list of switches.

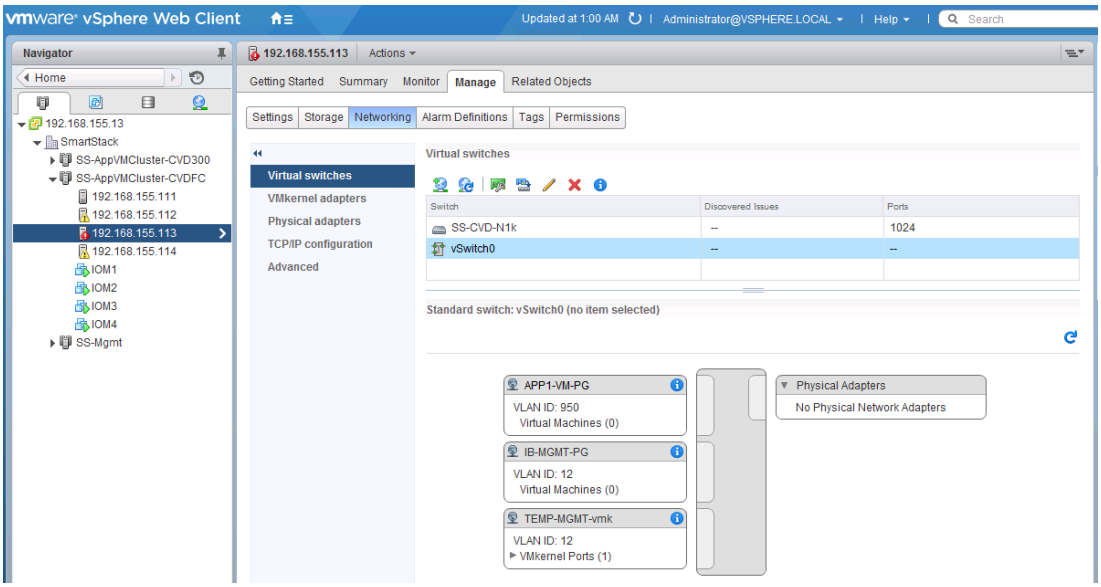

2. Select the temporary port-group created for migration (for example, `TEMP-MGMT-vmk`) and click X to remove it. Click Yes to confirm.

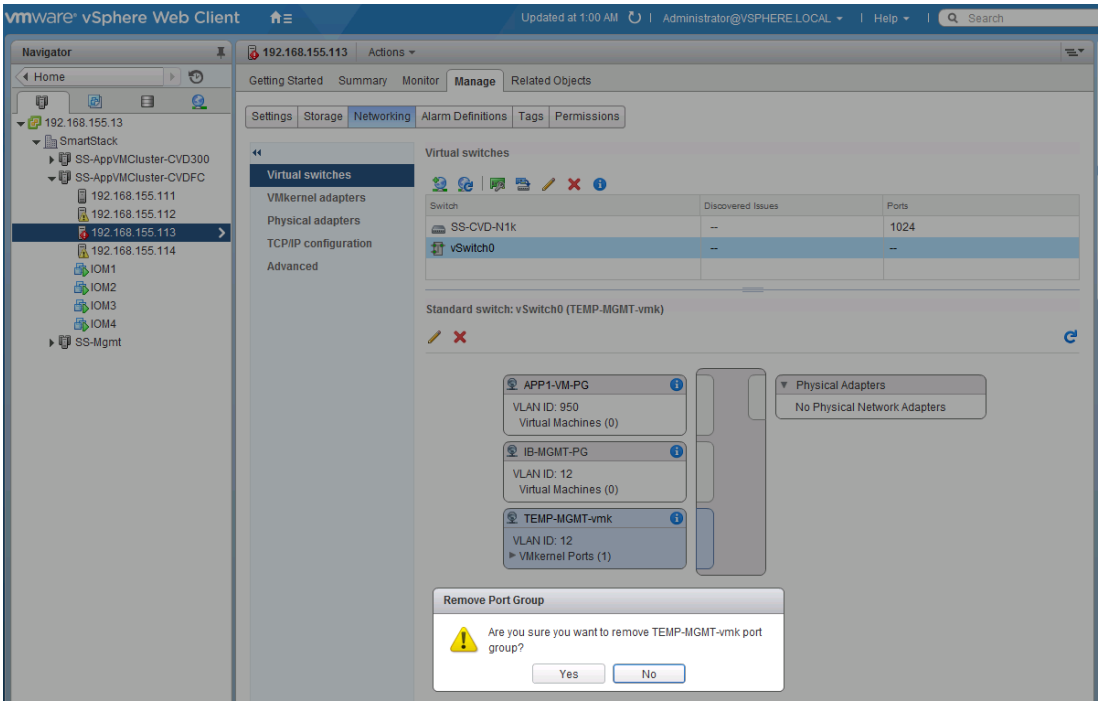

3. Confirm that vSwitch0 is still highlighted in the list of Virtual Switches and click 'X' to remove the vSwitch. Click Yes to confirm the deletion of vSwitch (for example, vSwitch0).

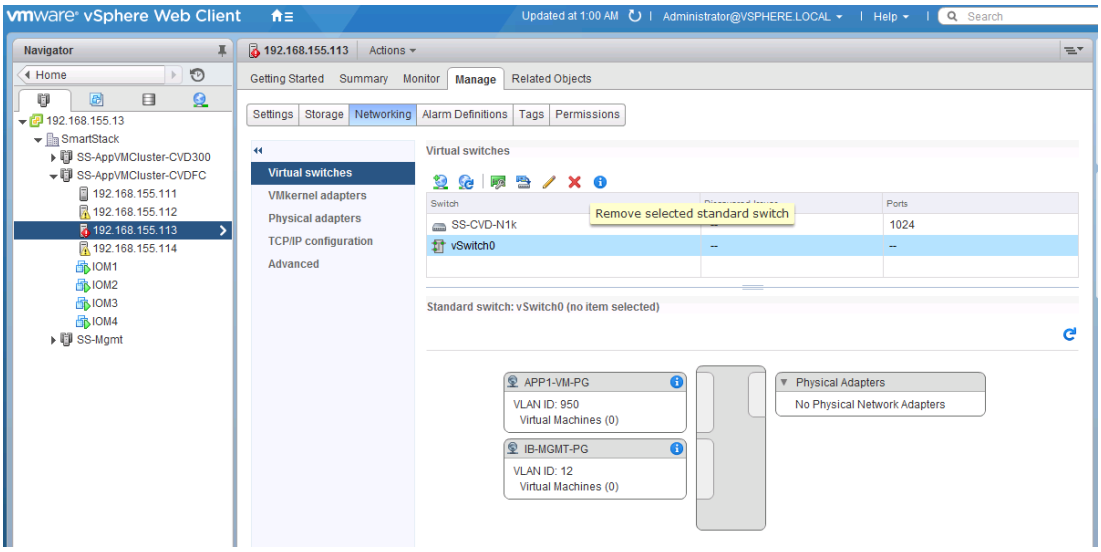

4. The resulting configuration should be as follows.

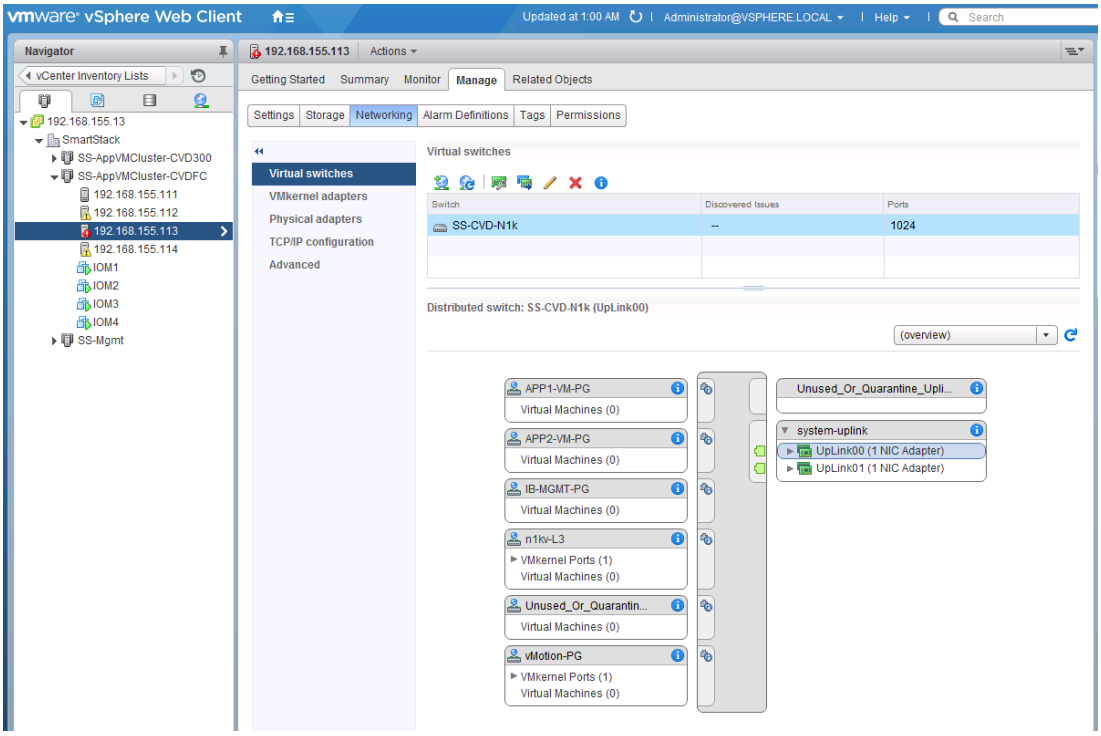

Optional: Deploy Cisco UCS Performance Manager

Download and Deploy Cisco UCS Performance Manager OVA

1. Login to www.cisco.com and download Cisco UCS Performance Manager OVA.

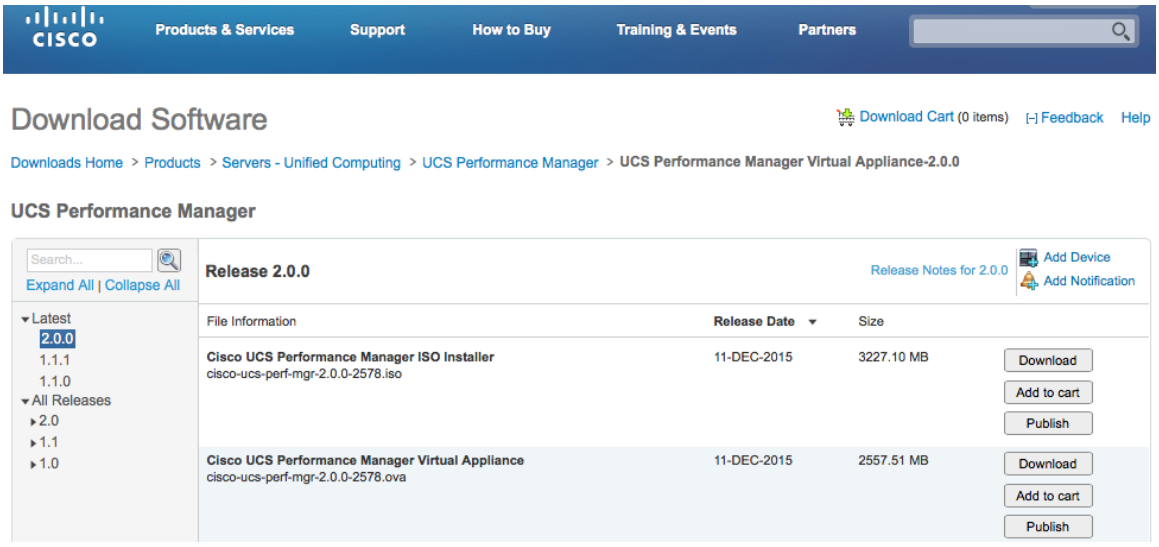

2. From VMware vSphere web client, right click on the cluster name (for example, `SS-Mgmt`) and select **Deploy OVF Template...**

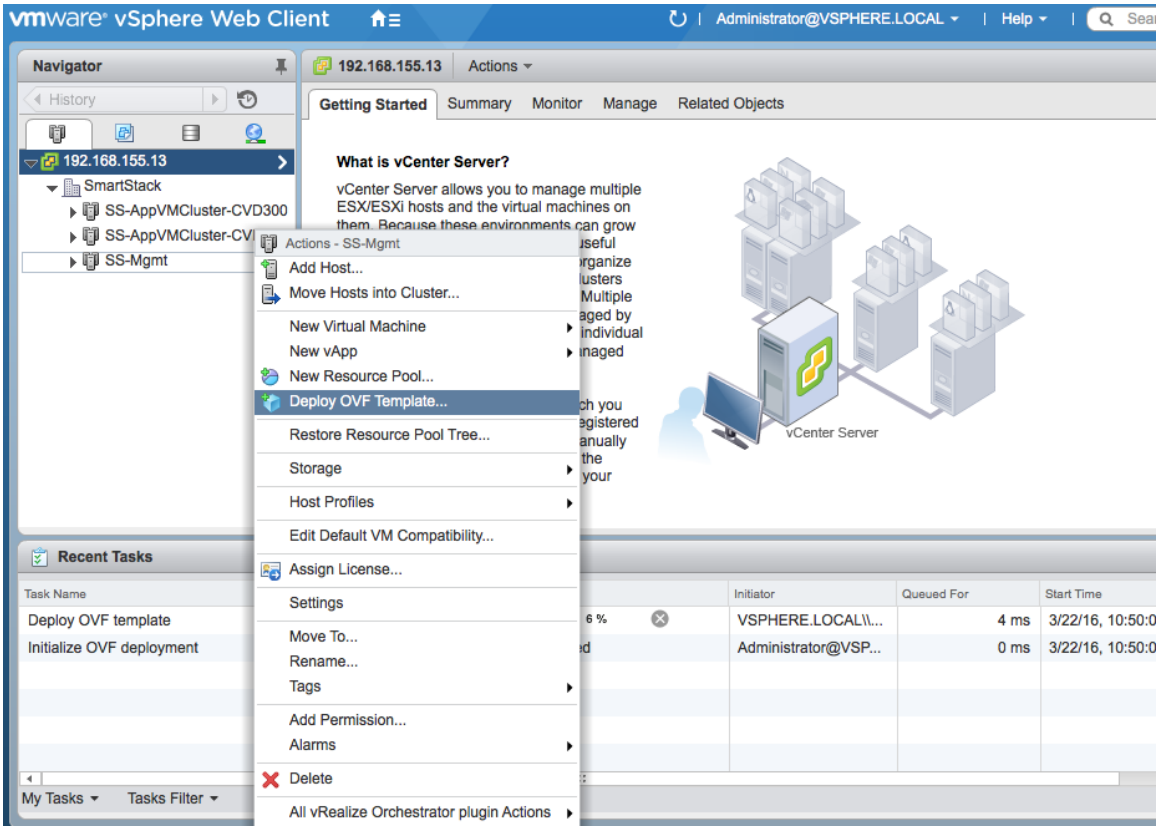

3. If the VMware Client Integration Plug-in hasn't been installed on the machine running the VMware vSphere web client, you will get the following message. Click on the Download the Client Integration Plugin to download the plug-in.

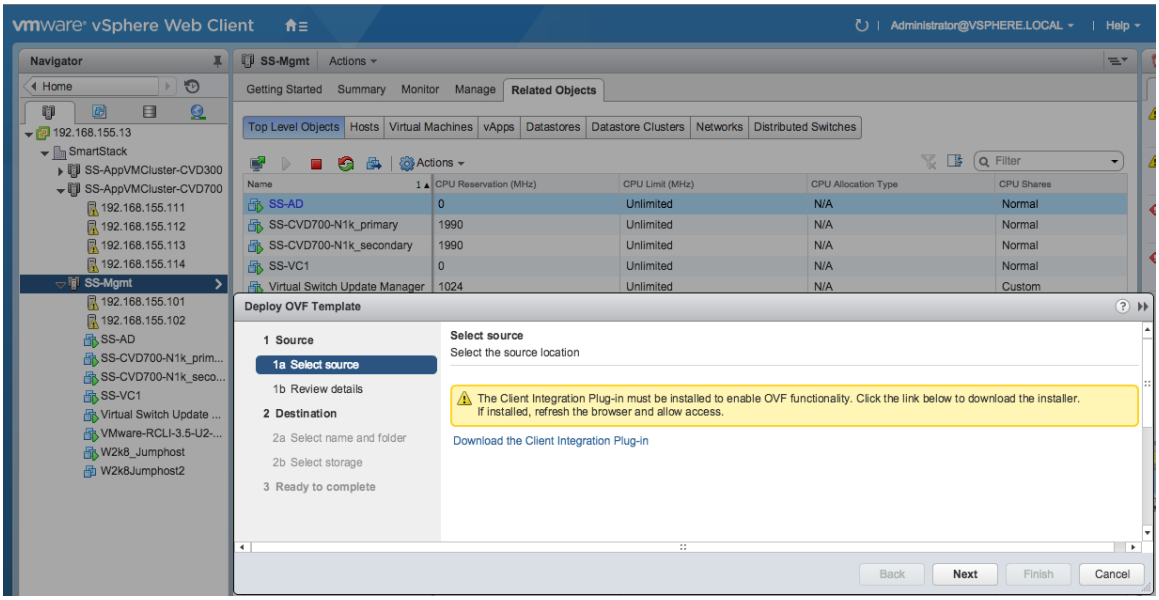

4. Install the plug-in. Refresh or restart browser and login and repeat steps to **Deploy OVF Template...**

5. Select the Local file radio button and browse to the downloaded Cisco UCS performance manager location. Click Next to continue.

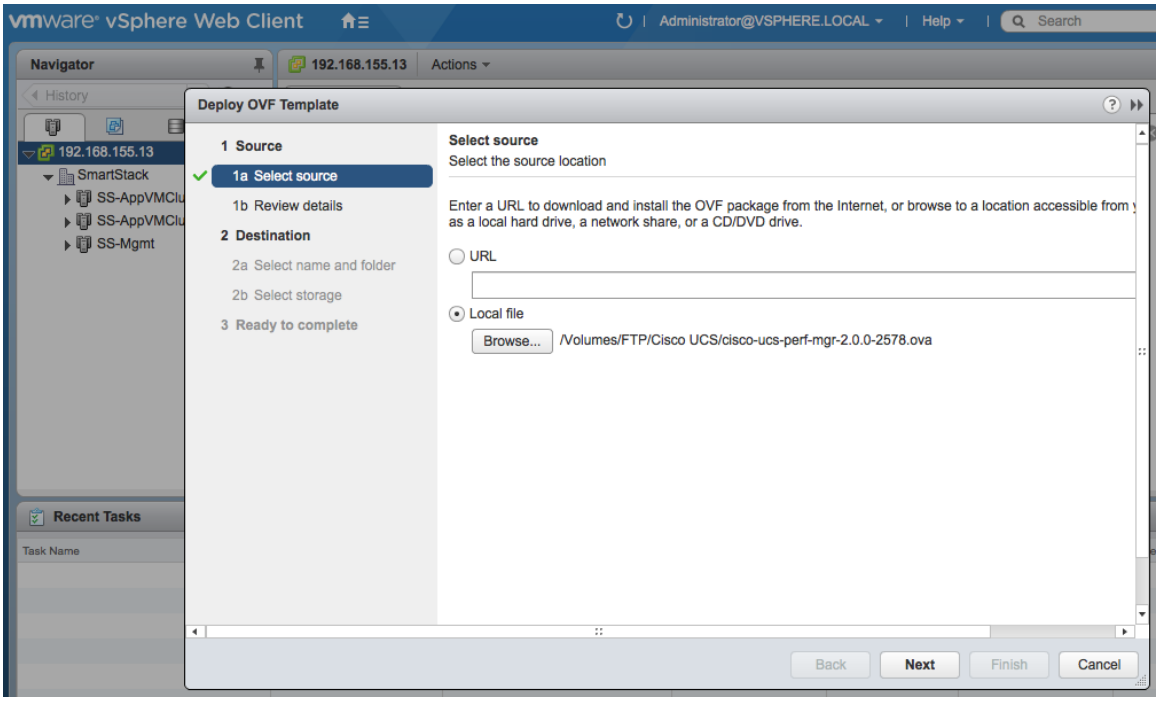

6. Review the details of the OVF template. Click Next to continue.

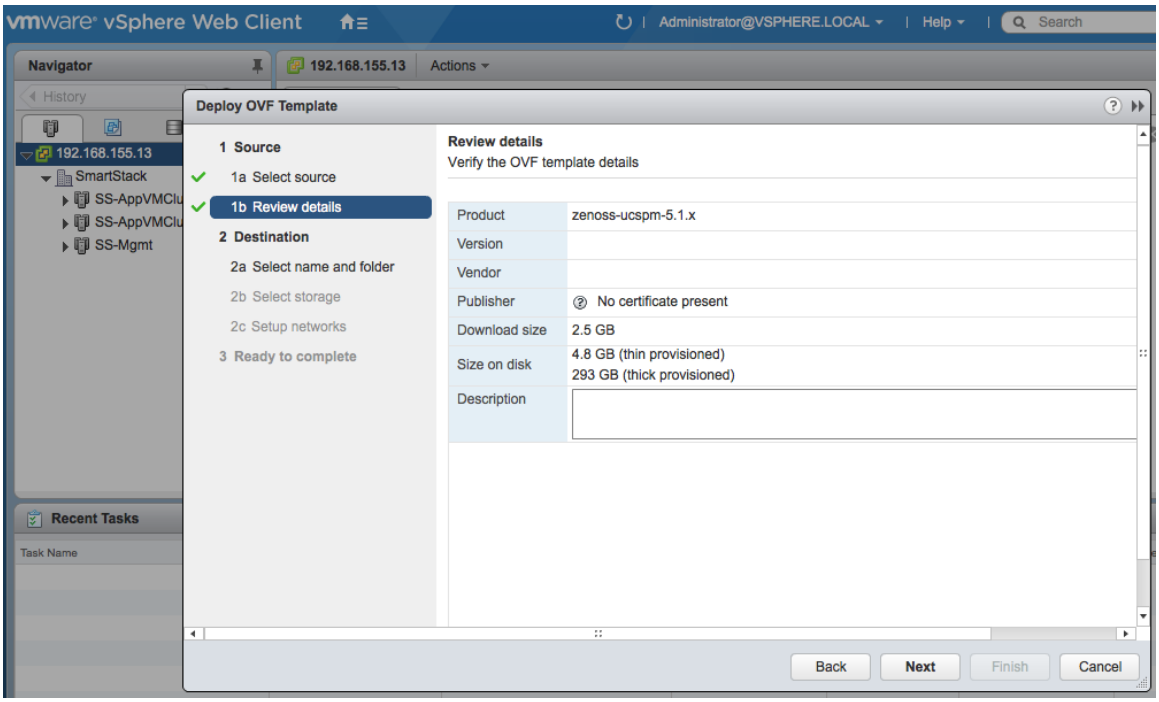

7. Select the name of the VM and location (for example, `SmartStack` datacenter) where it should be deployed. Click Next to continue.

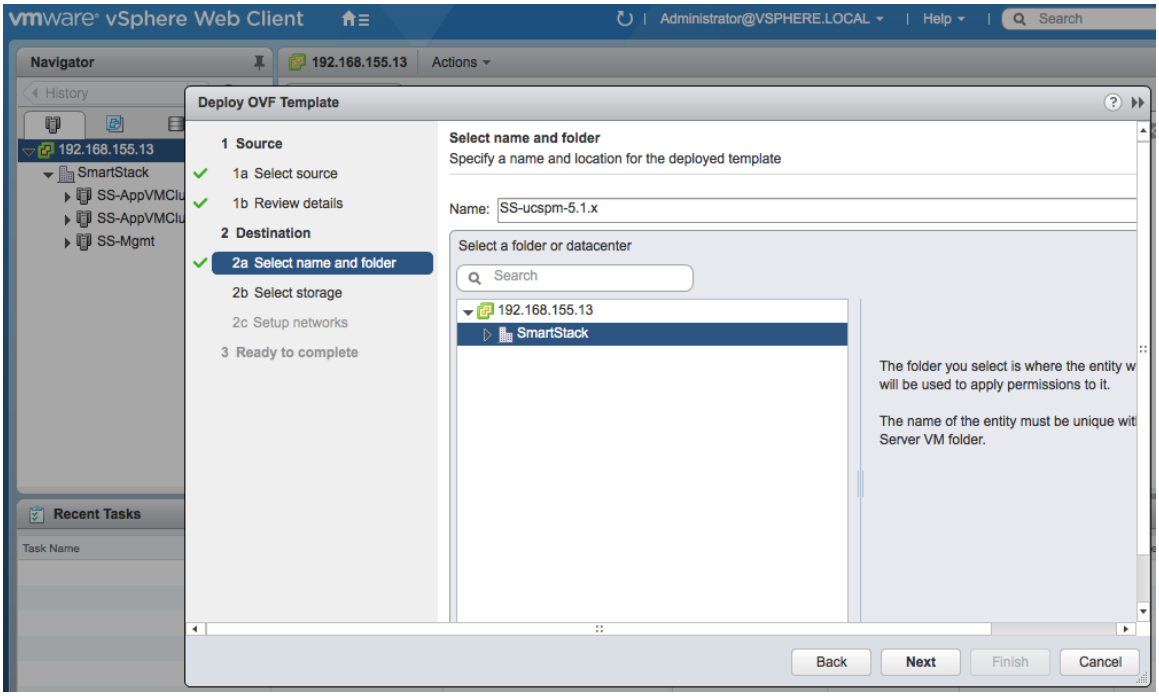

8. Select the datastore the VM should use (for example, `DS-Infra`). Click Next to continue.

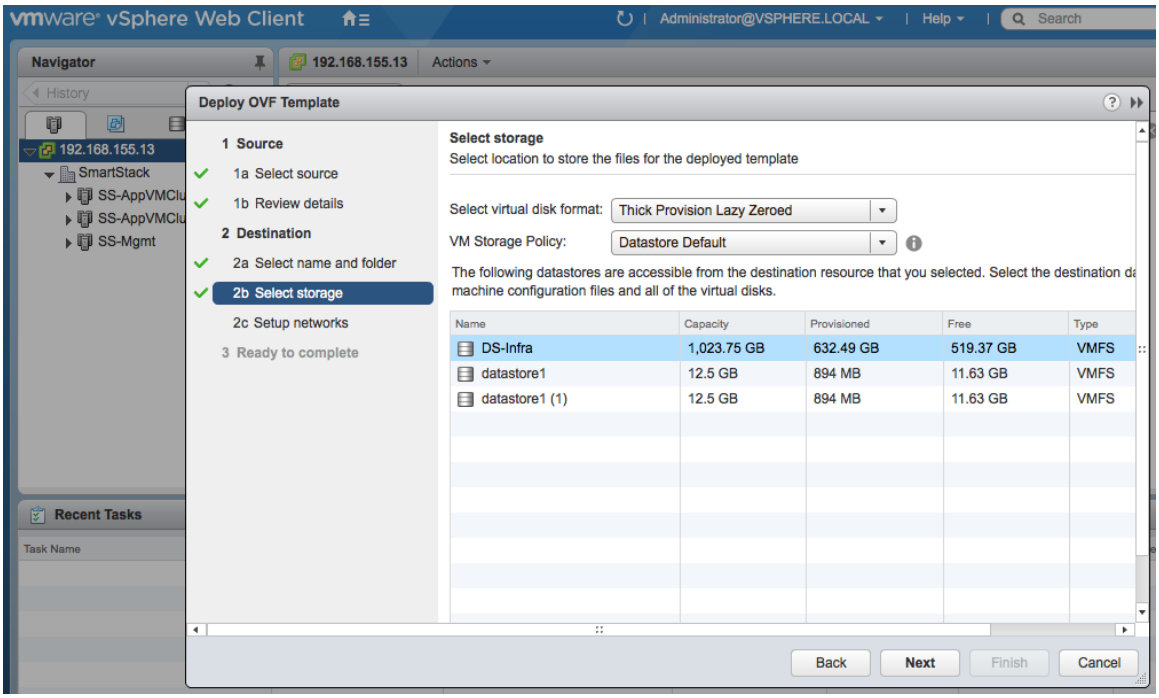

9. Select the port group/profile (for example, `IB-MGMT-PG`) that the VM should be part of.

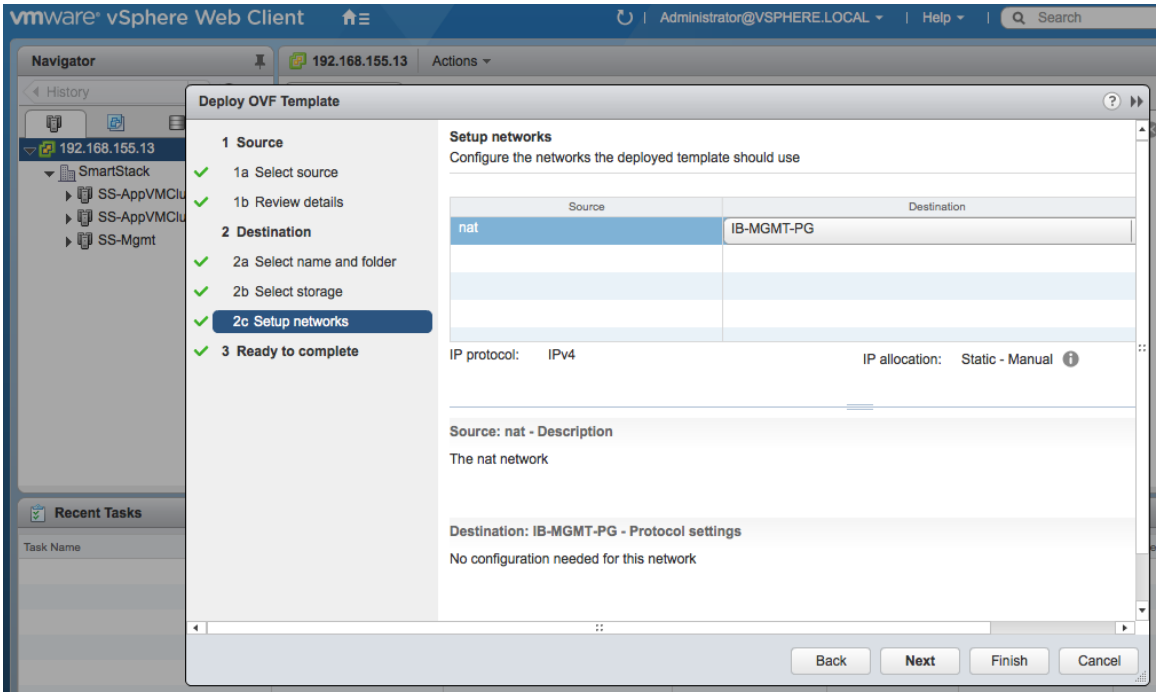

10. Review the selections. Leave Power ON after Deployment check box UNCHECKED. Click Finish to begin the installation process and power on the Cisco UCS Performance Manager VM.

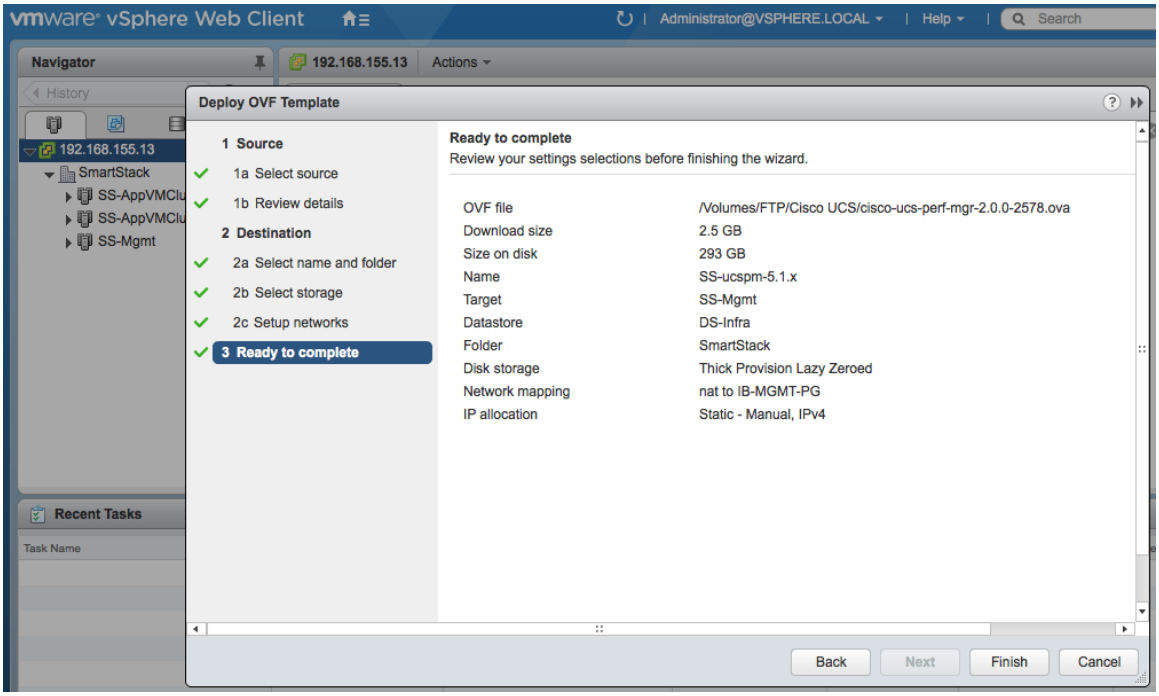

11. Navigate to Cisco UCS Performance Manager VM and select the Getting Started tab. Click Edit Settings and change the memory size from 40GB (shown below) to 64GB.

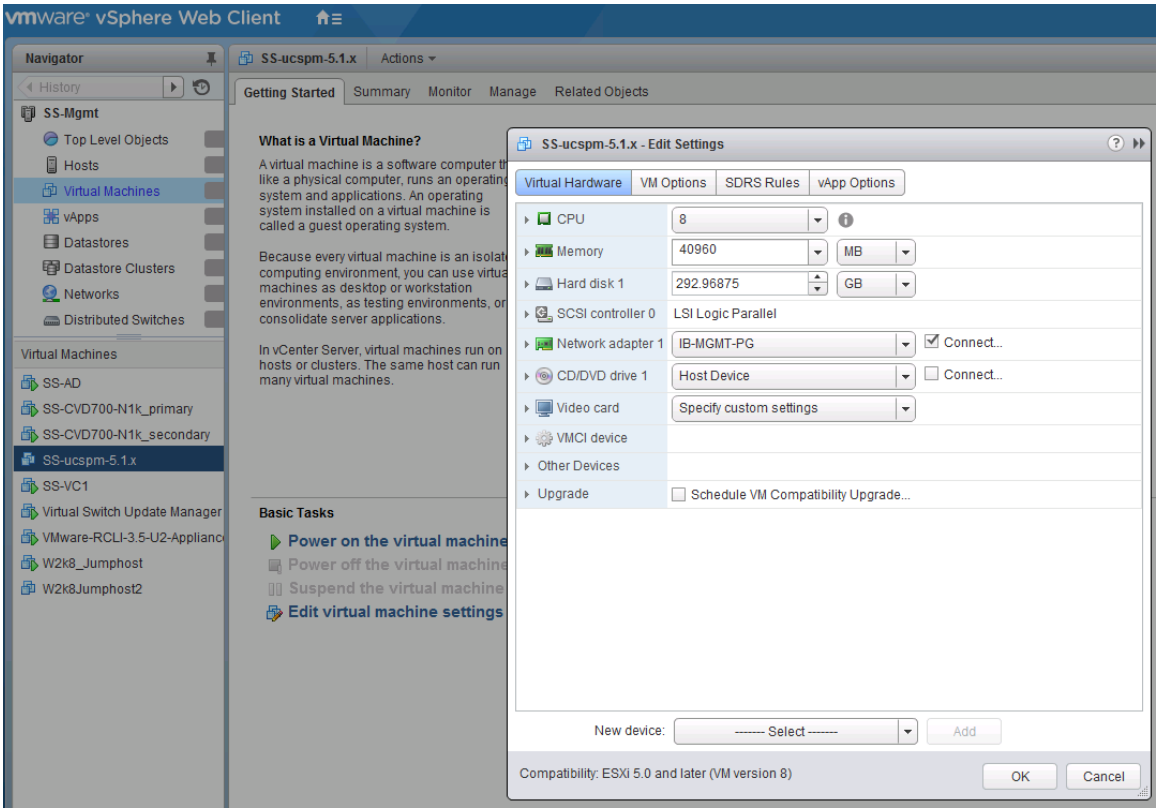

12. Power ON the Cisco UCS Performance Manager VM.

Initial Setup

1. From VMware vSphere web client, go to Summary tab and click Launch Remote Console to access the console of the Cisco UCS Performance Manager. Click Download Remote Console if needed.
2. Login as `root` and an initial password of `ucspm` (shown below). The system prompts you for new passwords for `root` and `ccuser` users. Configure the passwords for both users.

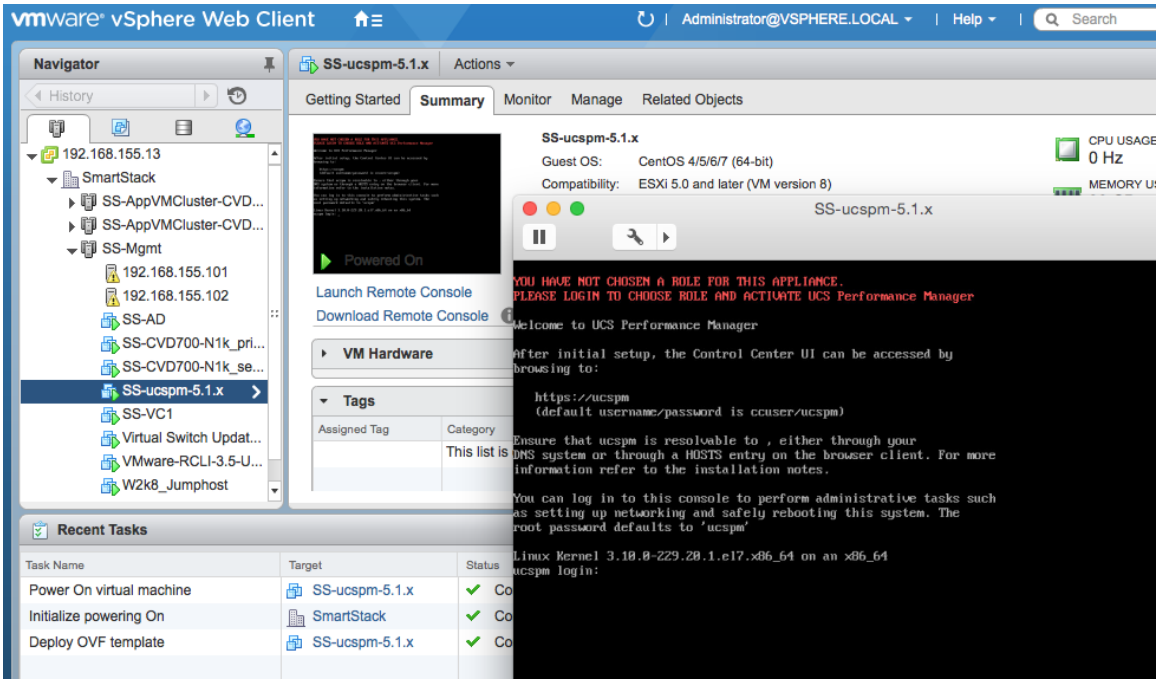

3. Select **`Master`** as the mode for this appliance and hit Enter.

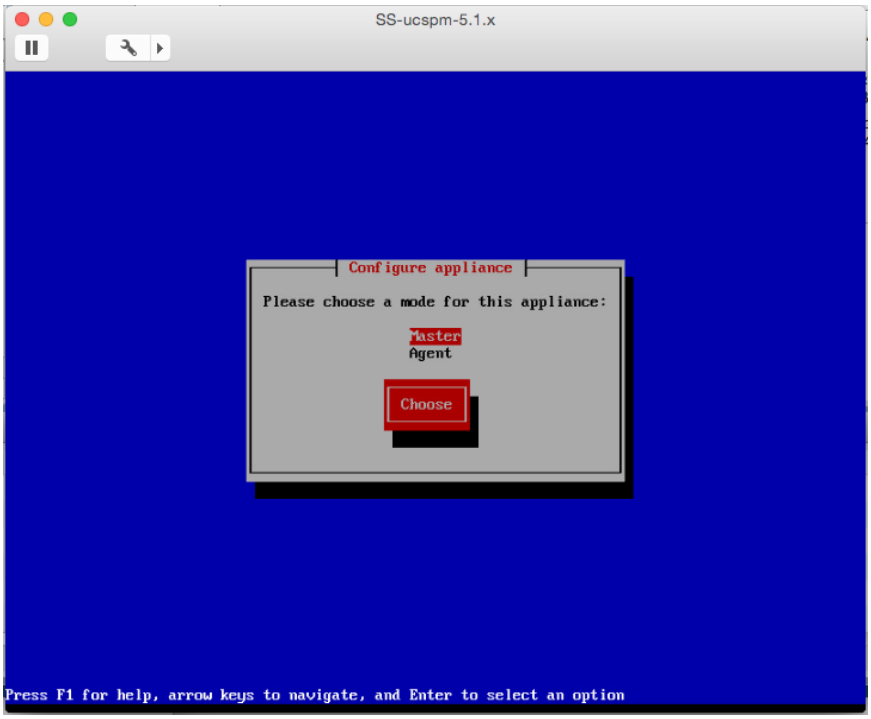

4. The system reboots and comes back up to a login prompt. Login as `root` user using the newly changed password.

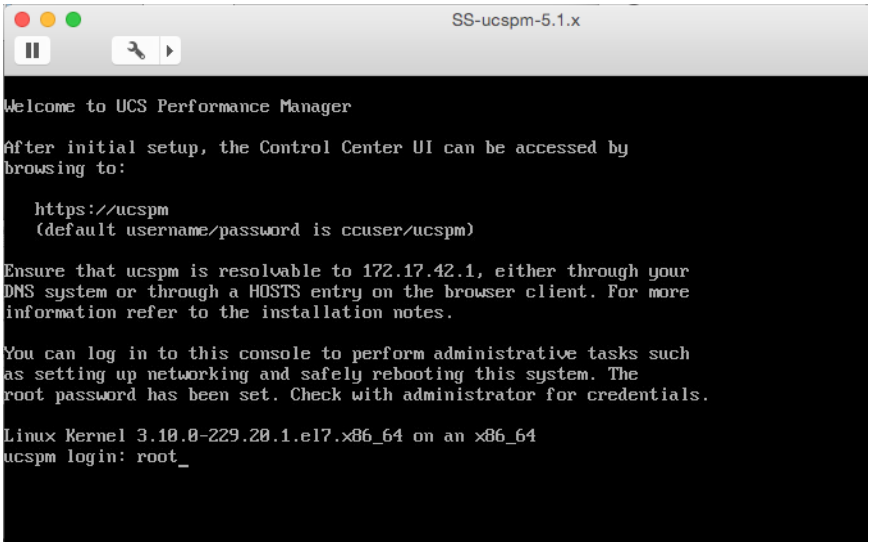

5. The default for the network connections is to use DHCP. To configure it to use static IPv4 addressing, select Configure Network and DNS and press the Enter or Return key.

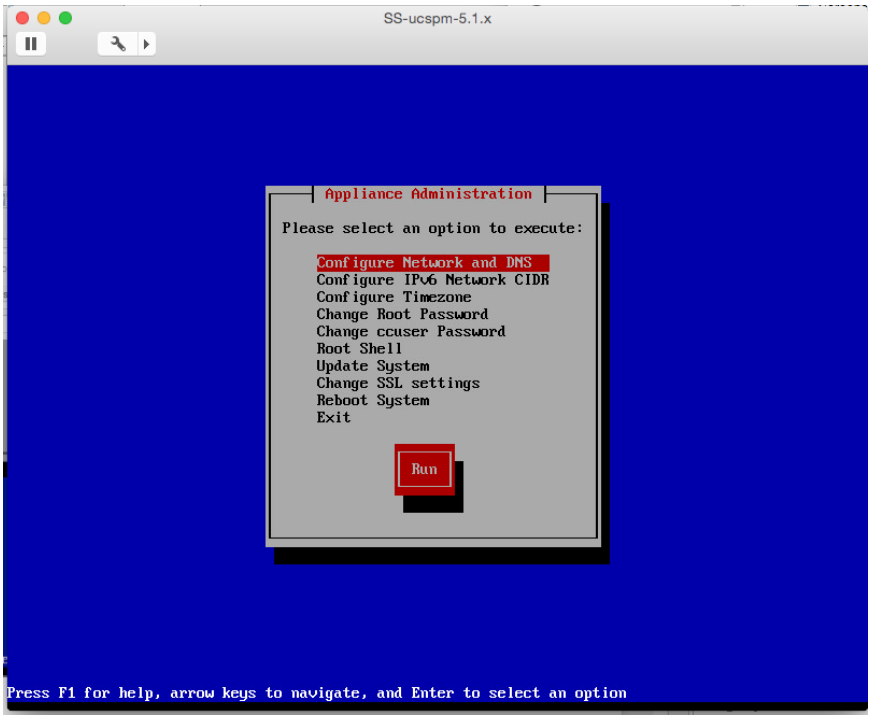

6. On the NetworkManager TUI window, select Edit a connection and then, press Enter or Return key.

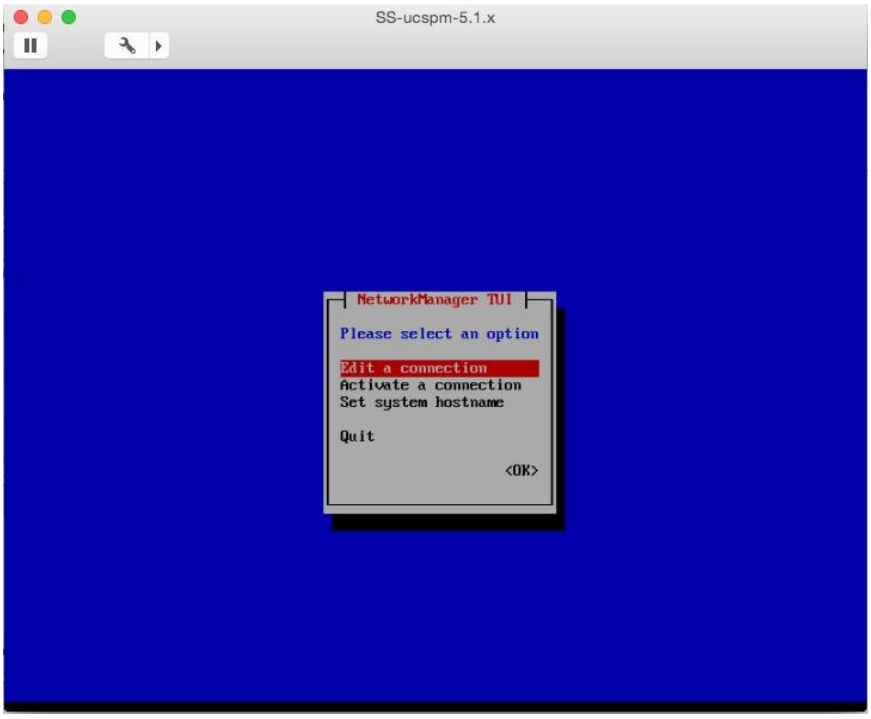

7. The available connections are shown below. Select Wired Connection 1 Tab over to <Edit...> and press Enter or Return key.

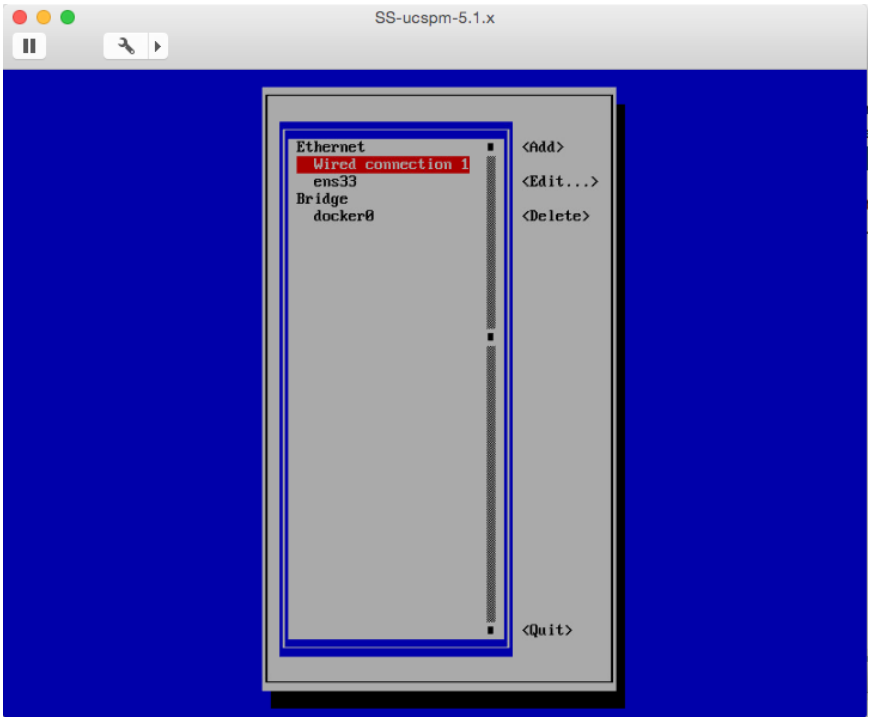

8. In the Edit Connection window for Wired Connection 1, edit the connection name (for example, `Ethernet 0`) and configure IPv4 addresses, Gateway, DNS and select IPv4 for this connection (optional). Use Tab key or Up/Down Arrow keys to navigate this window. Change the IPv4 Address from <Automatic> to Manual – hit Enter or Return key and select manual from the drop down menu.

9. Navigate to <OK> and then, press Enter or Return key.

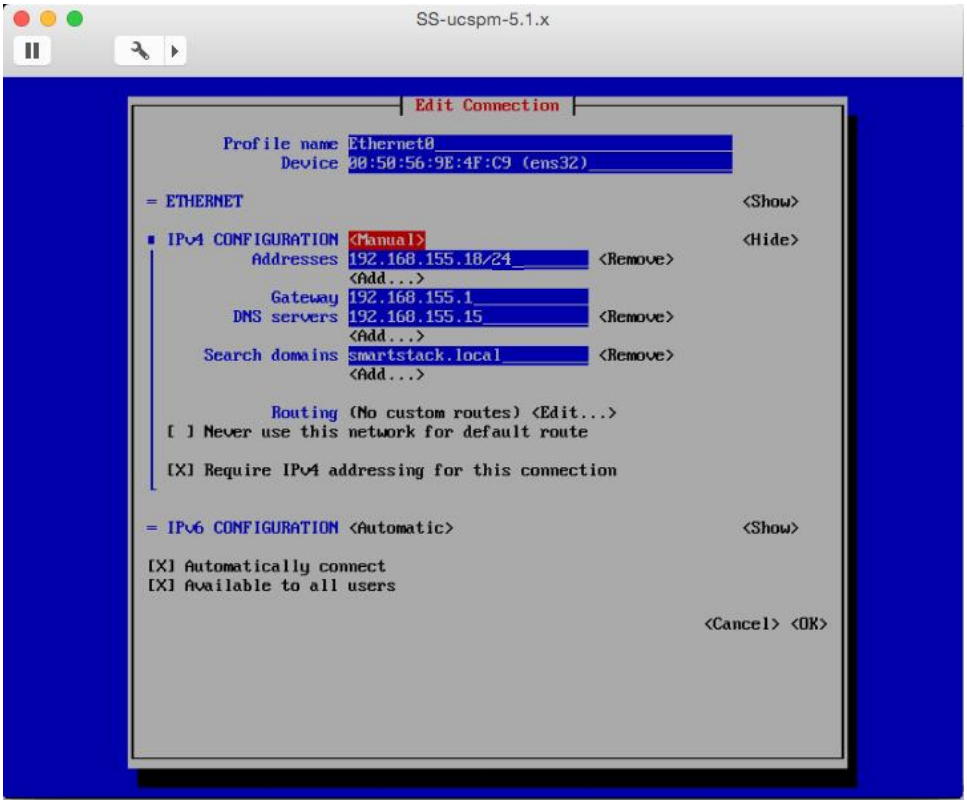

10. Navigate to <Quit> and hit Enter or Return key.

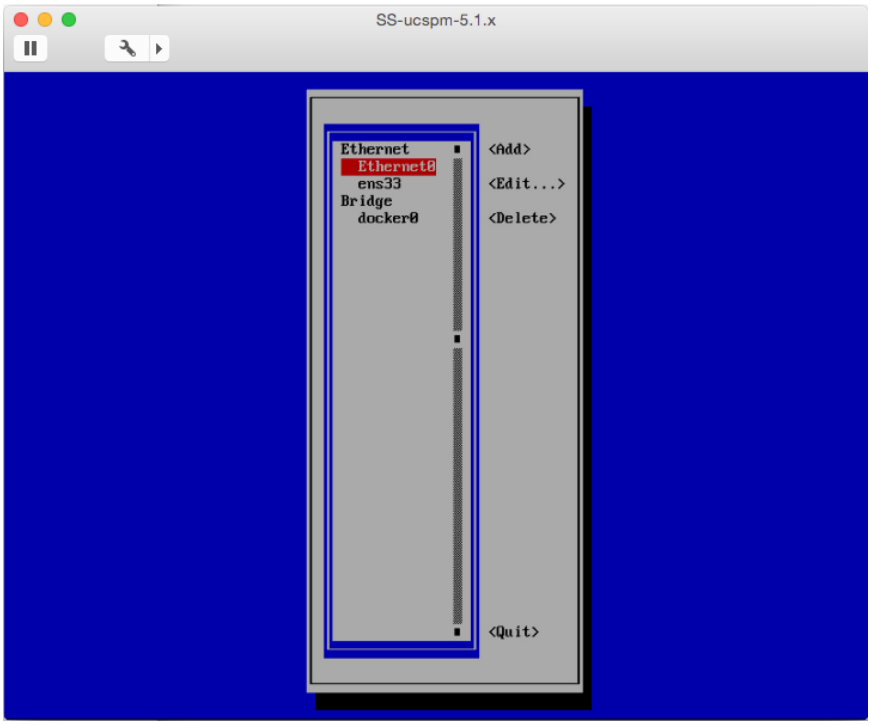

11. In the NetworkManager TUI screen, select Set system hostname to specify a hostname for the VM.

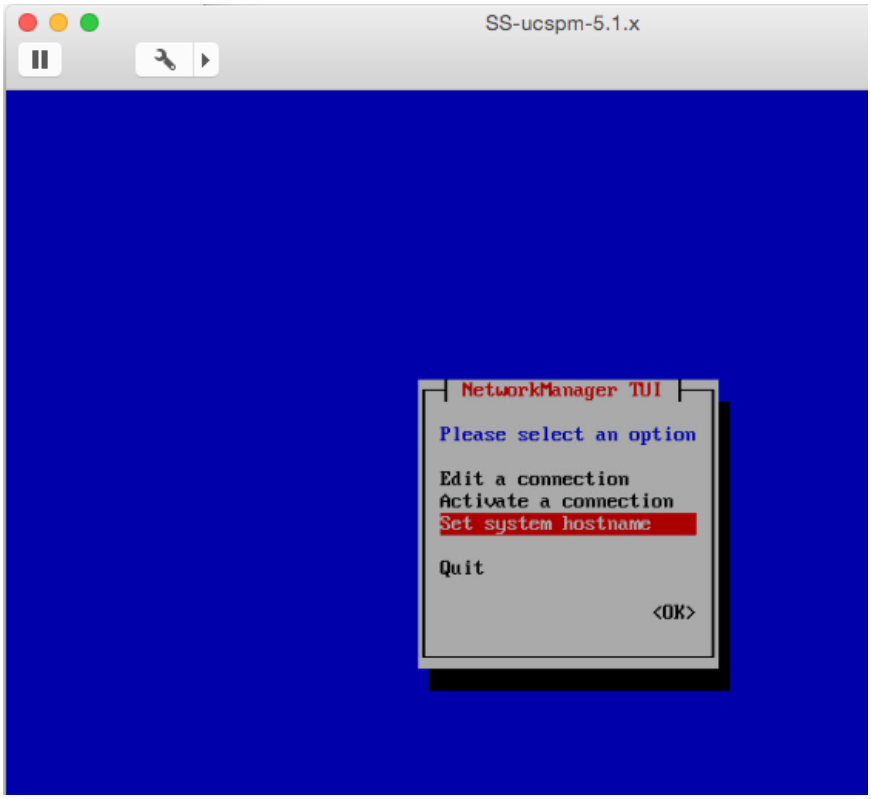

12. Specify a name (for example, `SS-ucspm`) and click <OK> to accept the change.

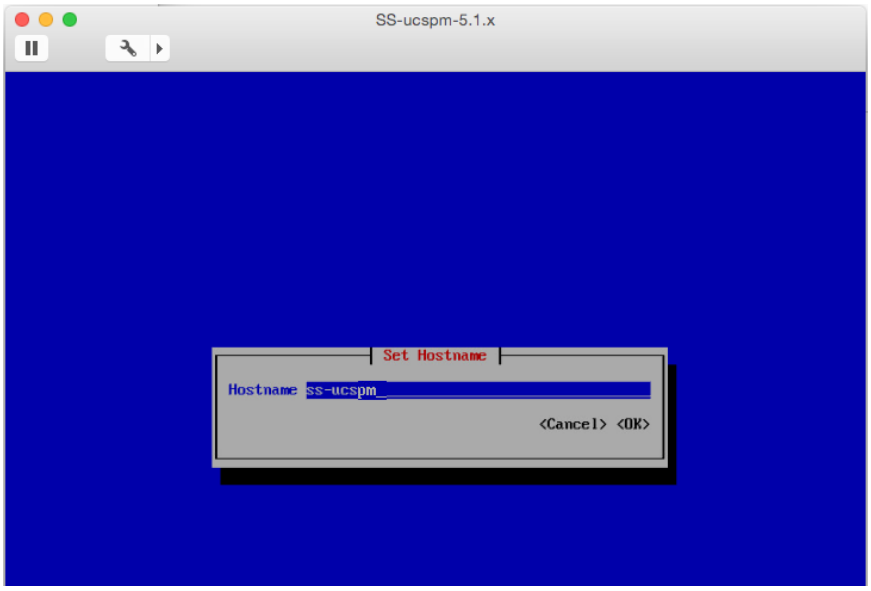

13. From the Appliance Administration menu, select Reboot the system for the changes to take effect.

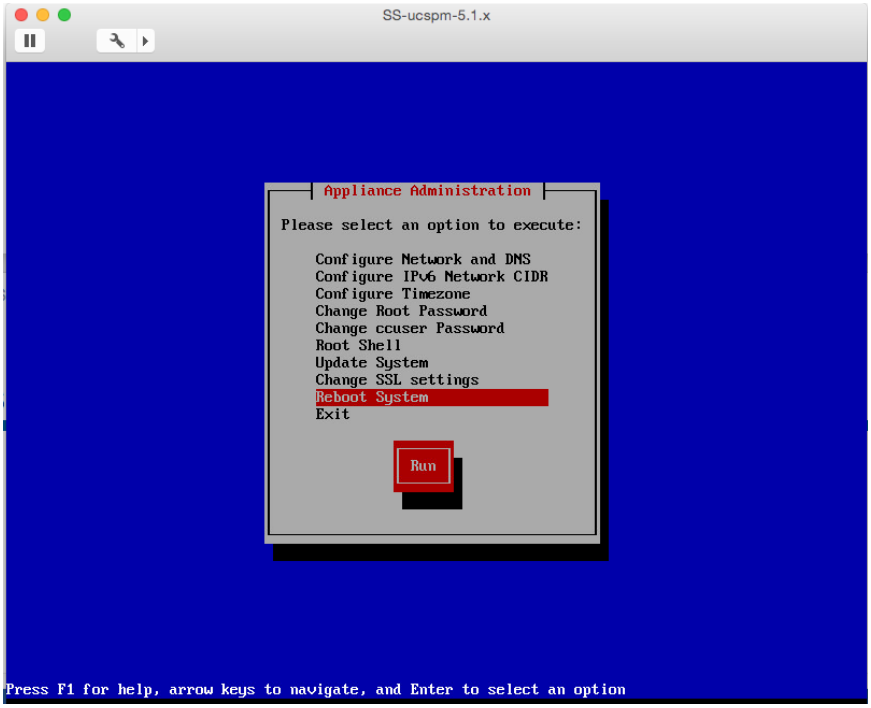

Deploy Cisco UCS Performance Manager from Control Center Master

1. Browse to Control Center web interface (for example, https://192.168.155.18). Login as ccuser with the new password.

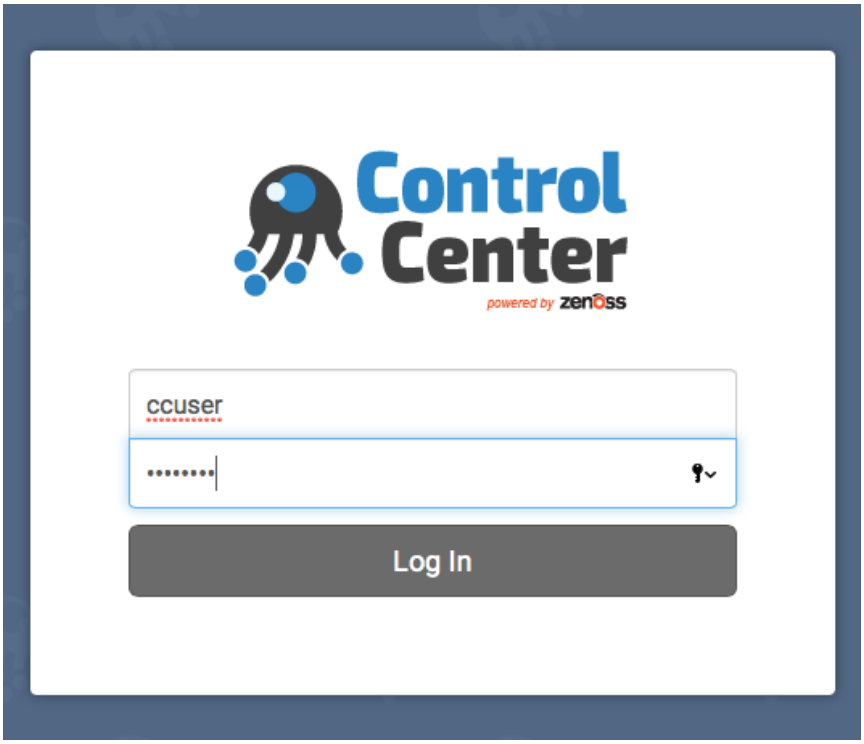

2. From the main Control Center landing page, in the Applications section, click + Application on the right side of the page.

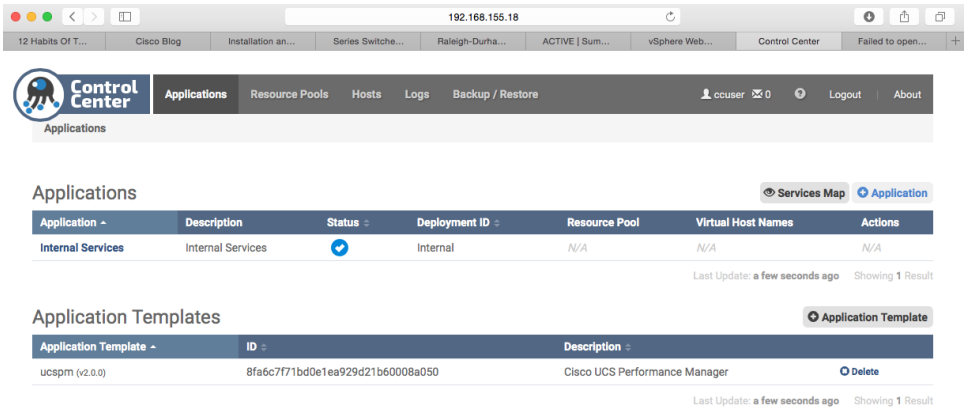

3. In the Deployment Wizard, specify the Host IP and port number (for example, 192.168.155.18:4979). Select default for the Resource Pool ID. Cisco recommends 100 for the RAM commitment. Click Next to continue.

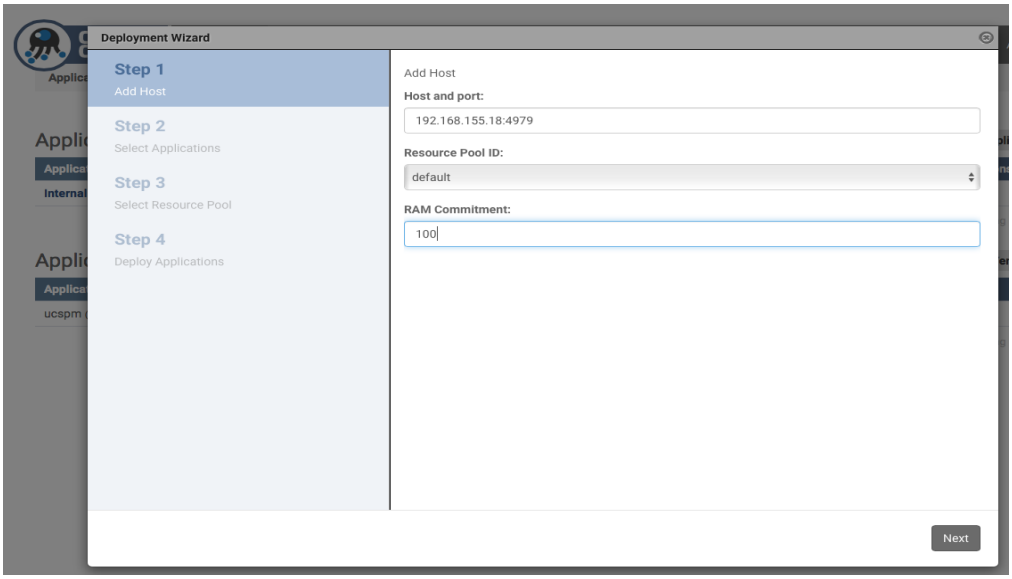

4. Select ucspm for the application to deploy. Click Next to continue.

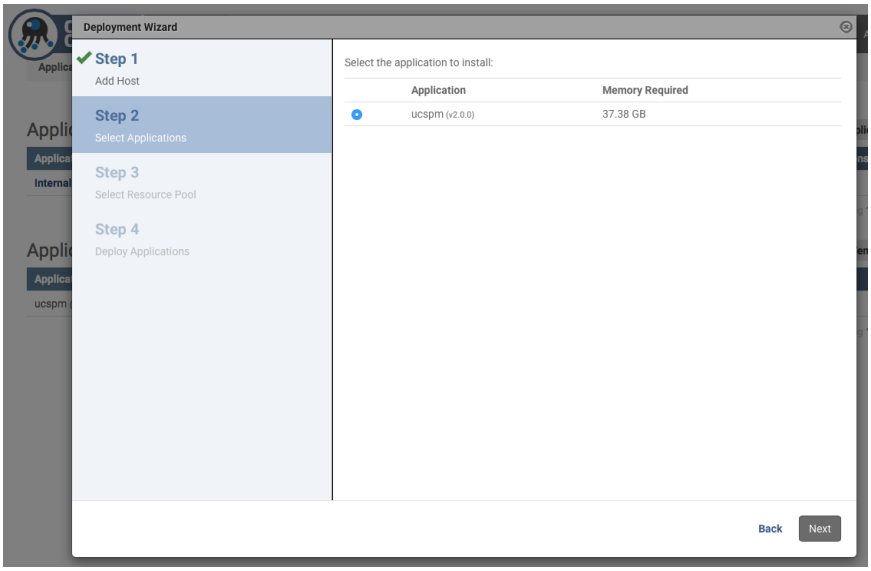

5. Select default for the Application Resource pool. Click Next to continue.

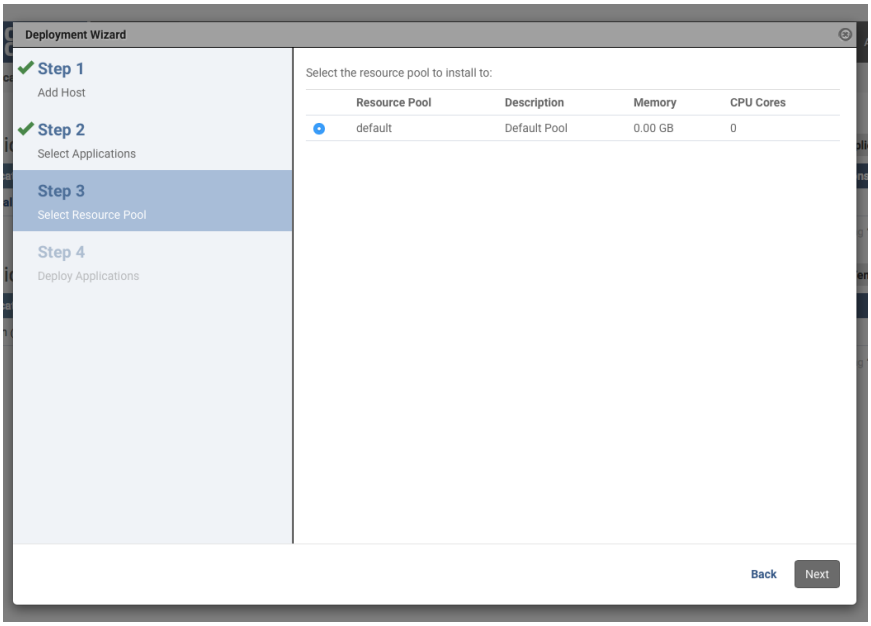

6. Select a Deployment ID (for example, `SS-UCSPM`) and click Deploy to deploy Cisco UCS Performance Manager.

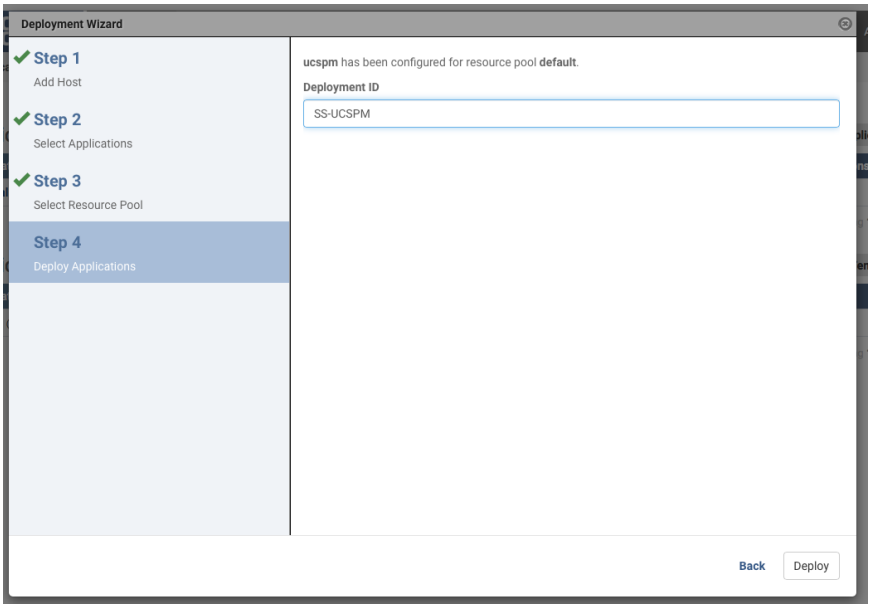

7. Cisco UCS Performance Manager is now deployed.

Manage (Start/Stop) Cisco UCS Performance Manager from Control Center

1. From the Control Center landing page (for example, https://192.168.155.18), in the Actions column of the Applications table, click Start control of the ucspm row.

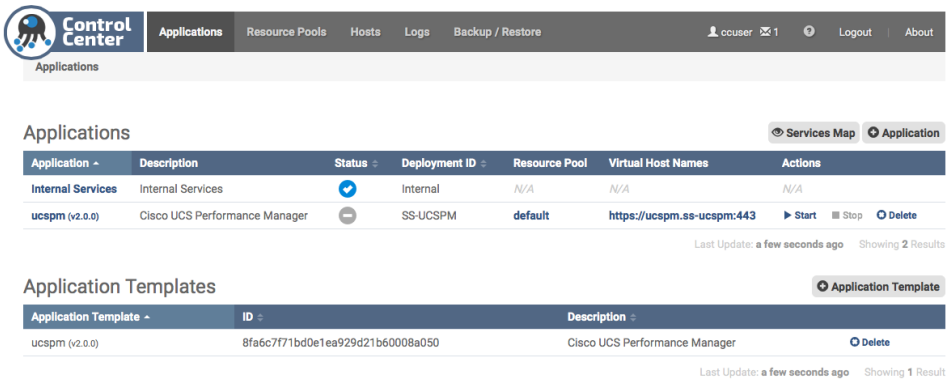

2. In the Start Service dialog, click Start Service and 46 Children button.

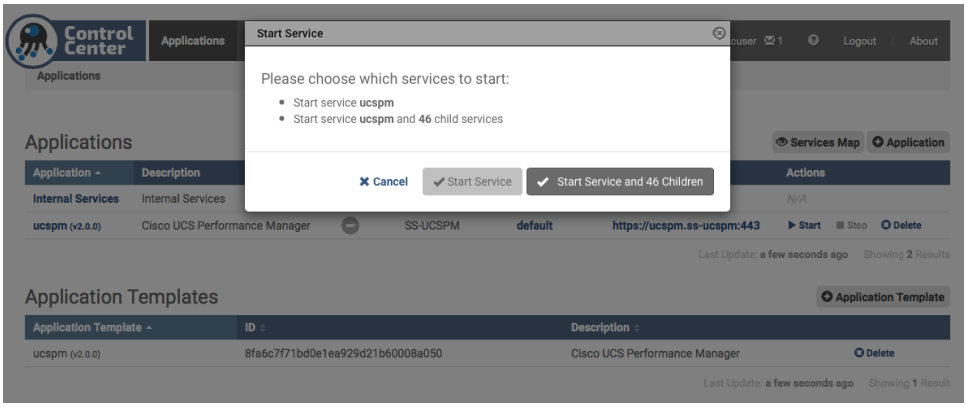

3. In the Application column of the Applications table, click ucspm in the ucspm row. Scroll down to watch child services starting.

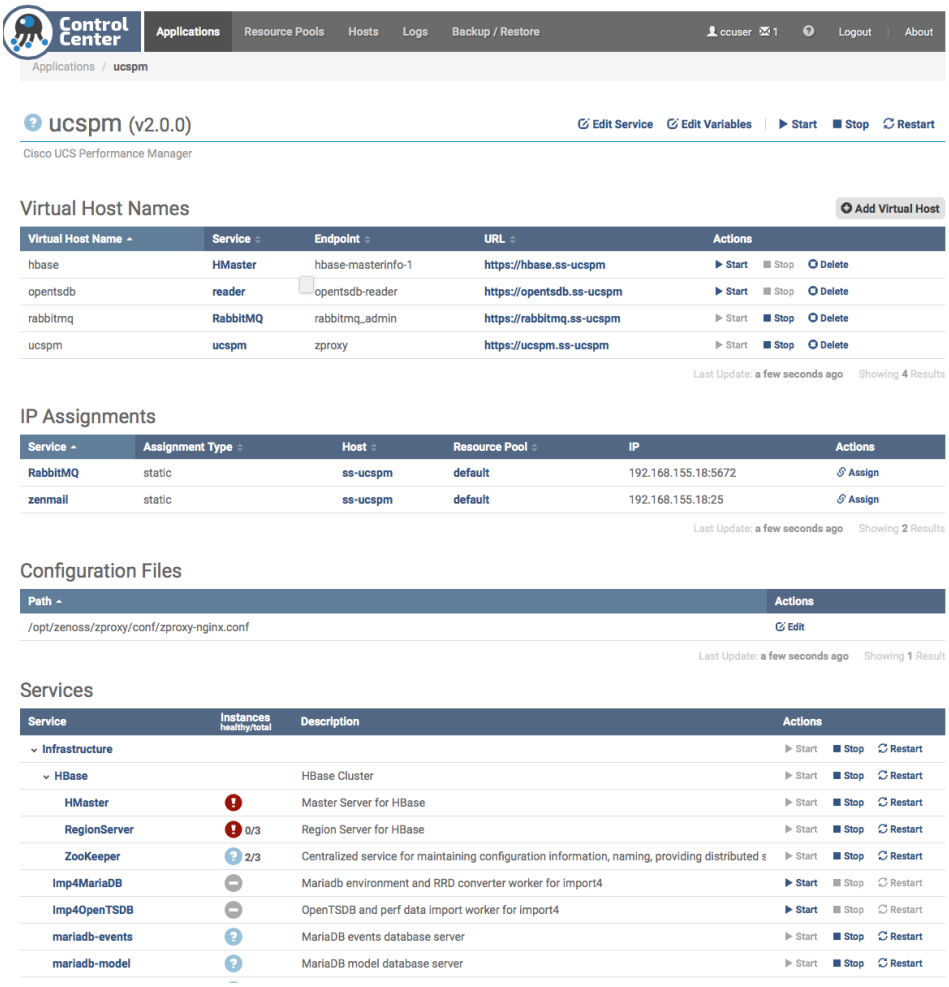

4. When a circle with a check appears next to ucspm at the top of the page instead of a circle with a '?', Cisco UCS Performance Manager is up and running.

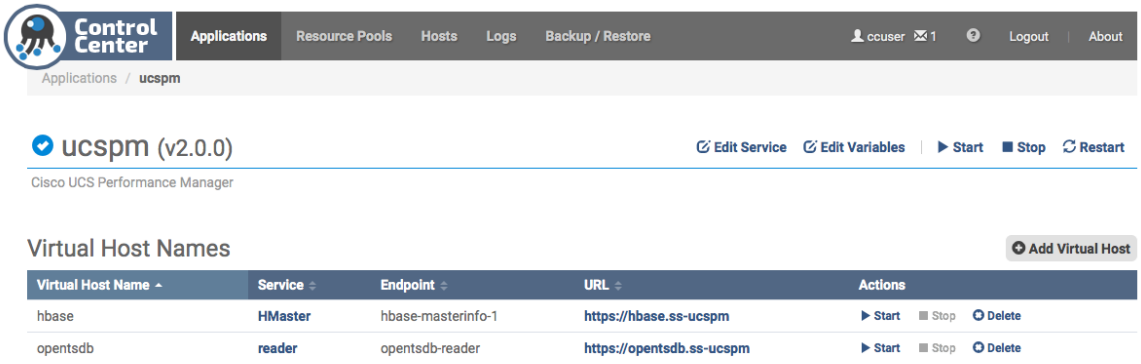

5. Optional: Log in to the Cisco UCS Performance Manager interface by scrolling to the Virtual Host Names table and click the link in the URL column of the entry with the hostname ucspm. DNS resolution must work correctly for this link to work.

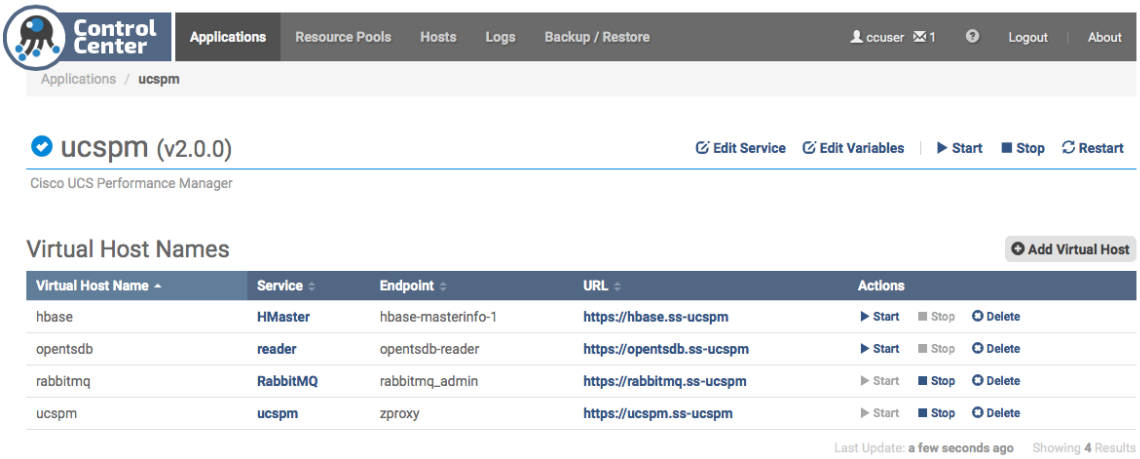

6. For more details on the initial install and setup of Control Center, see the Cisco UCS Performance Manager Installation Guide: http://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/ucs-performance-mgr/installation-guide/2-0-0/ucs-pm-installation-guide-200.pdf

Initial Setup of Cisco UCS Performance Manager

1. Browse to the hostname of the Cisco UCS Performance Manager (for example, `https:ucspm.ss-ucspm`). The first time after installing and starting Cisco UCS Performance Manager from Control Center, you'll be redirected to the Setup page and the End User License Agreement dialog will be shown. Scroll down to the bottom and check the checkbox and click on the Accept License button.
2. On the Cisco UCS Performance Manager Setup page, click Get Started!

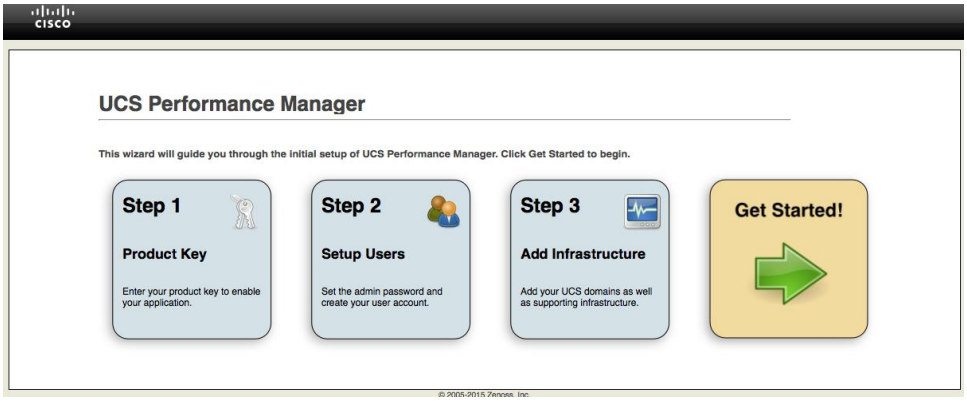

3. On the Add Licenses page, click Add License File button to upload licenses. Click Next to continue.

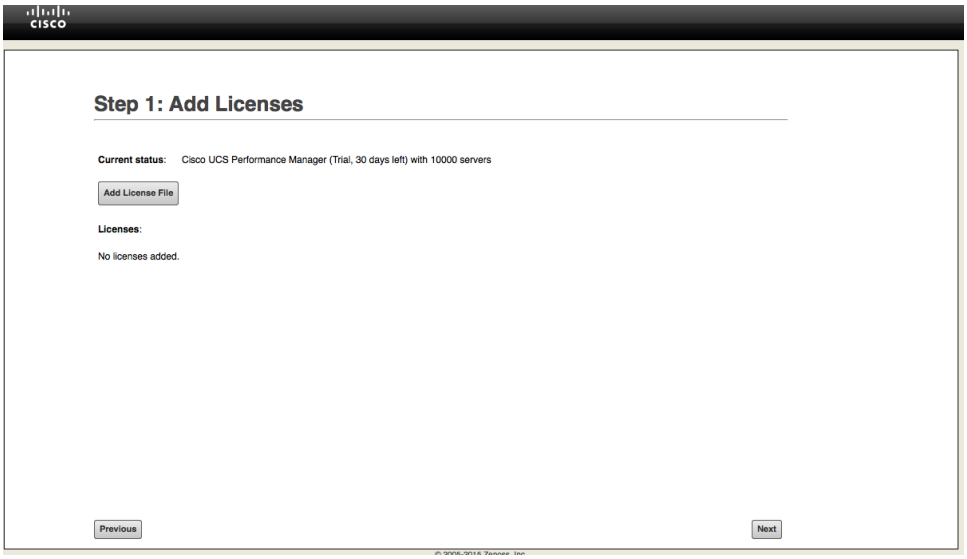

4. Setup admin and user accounts as needed. Click Next to continue.

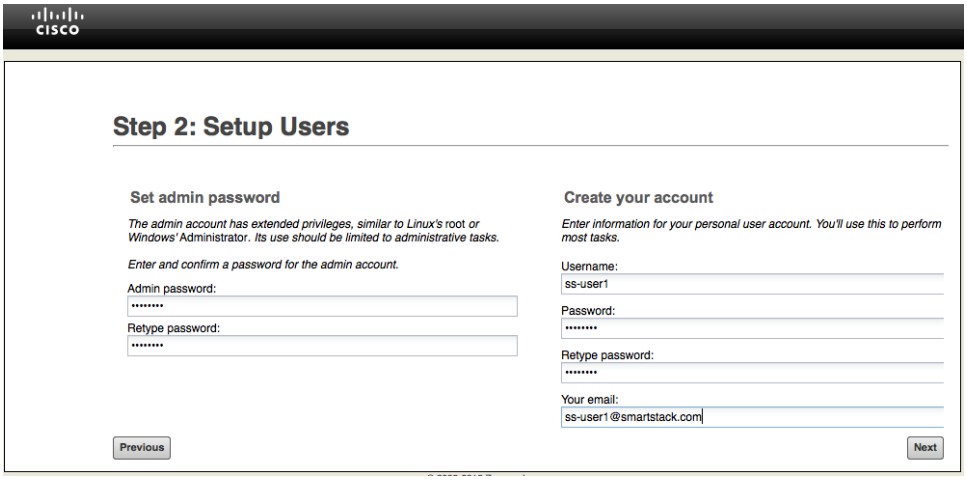

5. (Optional) If you're monitoring Cisco UCS Central, add here otherwise click Next to continue.

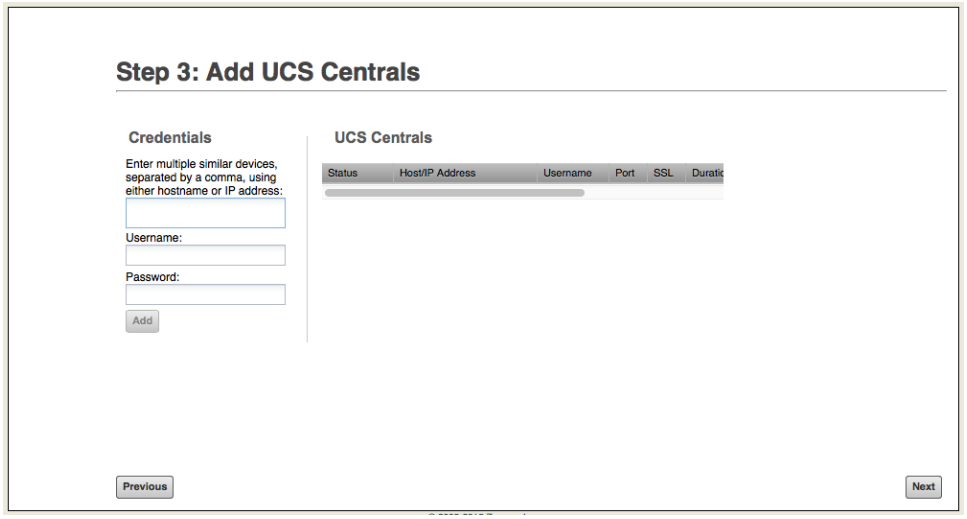

6. (Optional) Add Cisco UCS Domains to monitor. Click Next to continue.

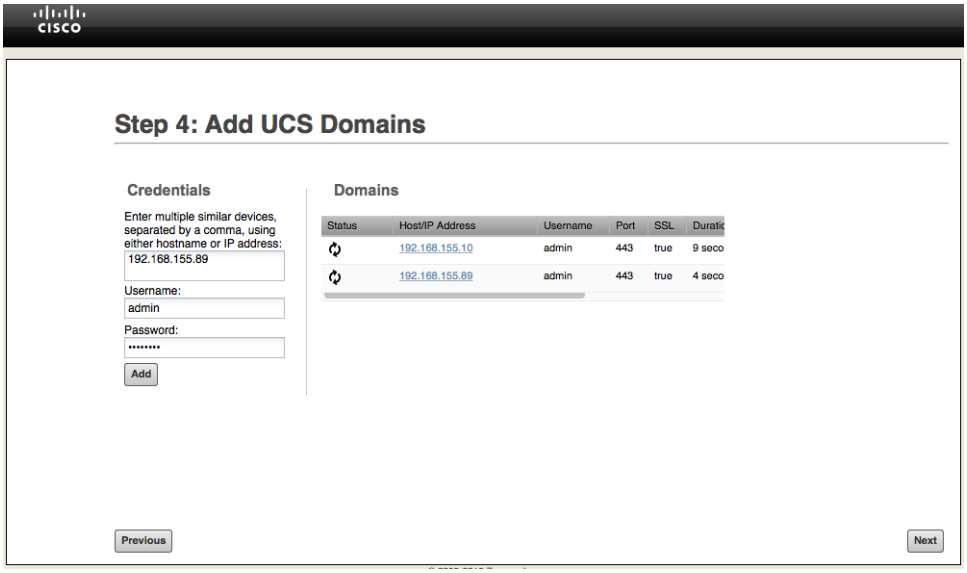

7. Add Infrastructure (Network, Storage, Hosts, Servers) components to monitor. Select Network and add Cisco Nexus and Cisco MDS components in the SmartStack design. Select the correct type when adding the devices. Cisco Nexus 9000 switches can be managed through SNMP or Netconf with NetConf preferred over SNMP if both are enabled. To monitor the switches using NetConf, '`feature nxapi`' needs to be enabled.

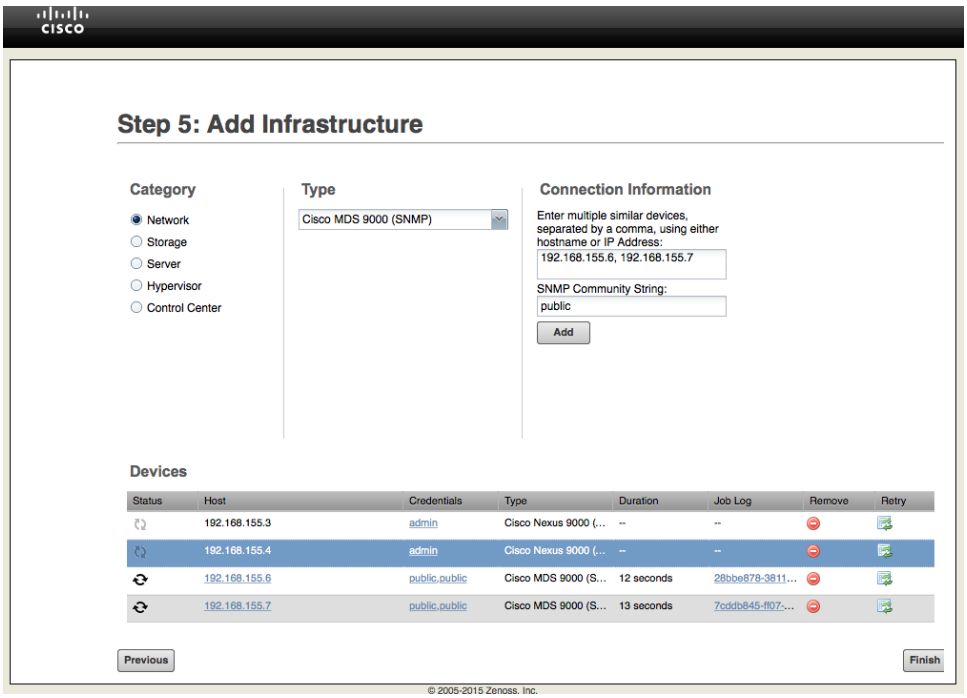

8. To monitor hosts in the SmartStack Infrastructure, select Hypervisor and add the hosts.

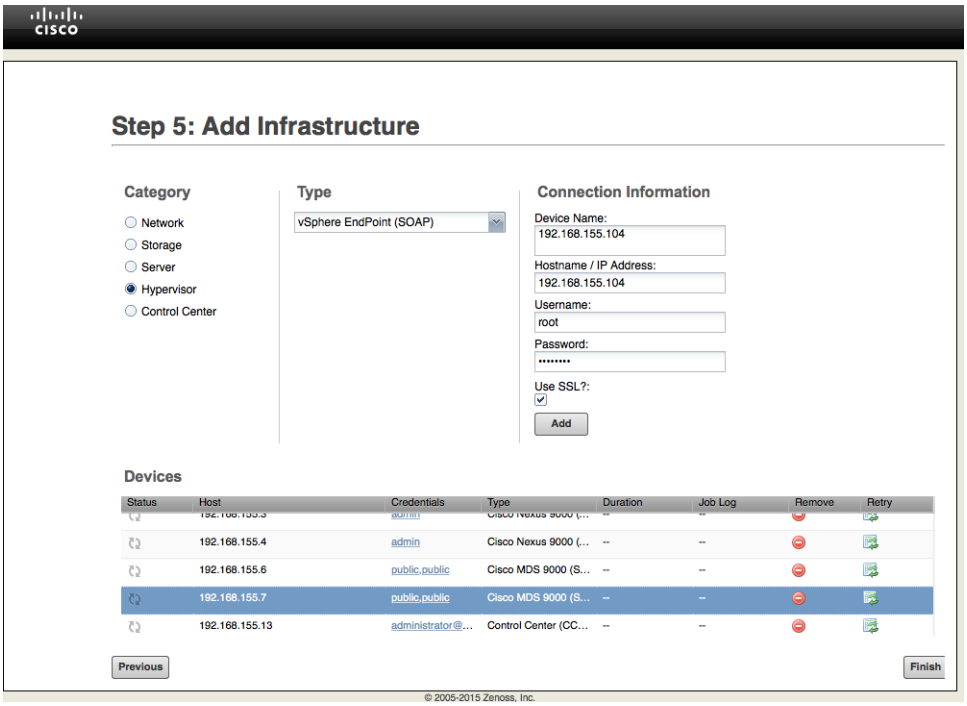

9. Click Finish when complete. The dashboard should be visible. The topology tab will show the high level topology of the setup as shown in the example below.

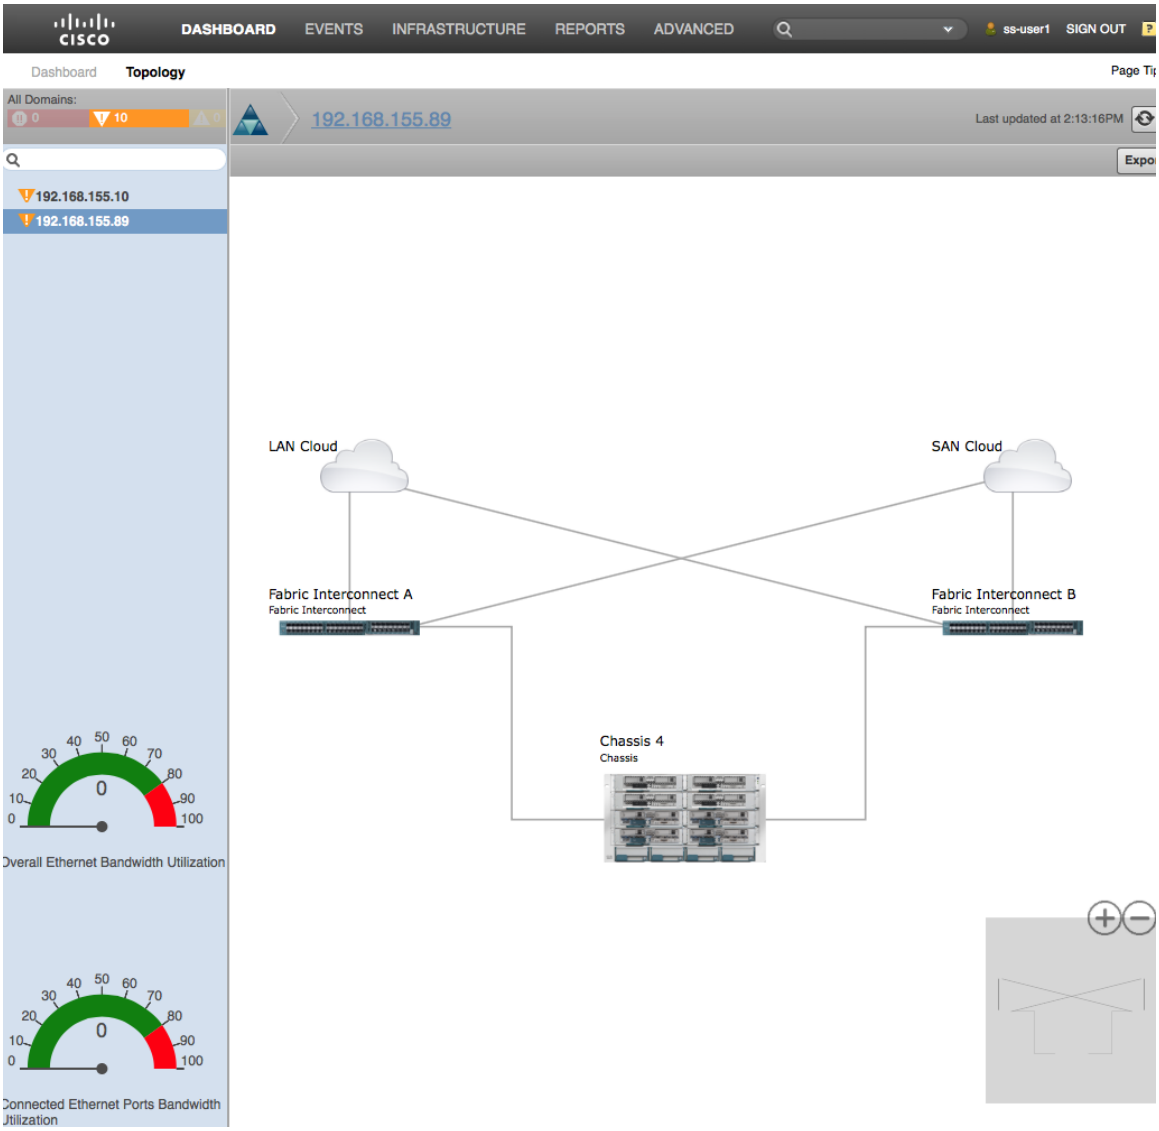

10. For more details on the initial setup of Cisco UCS Performance Manager, see the Cisco UCS Performance Manager Getting Started Guide: http://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/ucs-performance-mgr/getting-started-guide/2-0-0/ucs-pm-getting-started-guide-200.pdf

Appendix

Cisco Nexus A Configuration

```
D01-n9k1# show run

!Command: show running-config

!Time: Mon Apr 18 20:36:26 2016

version 6.1(2)I3(5)

switchname D01-n9k1

vdc D01-n9k1 id 1

  allocate interface Ethernet1/1-54

  limit-resource vlan minimum 16 maximum 4094

  limit-resource vrf minimum 2 maximum 4096

  limit-resource port-channel minimum 0 maximum 512

  limit-resource u4route-mem minimum 248 maximum 248

  limit-resource u6route-mem minimum 96 maximum 96

  limit-resource m4route-mem minimum 58 maximum 58

  limit-resource m6route-mem minimum 8 maximum 8

feature nxapi

cfs eth distribute

feature udld

feature interface-vlan

feature lacp

feature vpc

username admin password 5 $1$yjwMRDJV$ixhesPxYkTFiEPlkzje6F/  role network-admin

ip domain-lookup

system default switchport shutdown

copp profile strict

snmp-server user admin network-admin auth md5 0xa86138602a80f77232959a8e92306867 priv
0xa86138602a80f77232959a8e92306867 localizedkey

rmon event 1 log trap public description FATAL(1) owner PMON@FATAL

rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
```

```
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-admin
ip route 0.0.0.0/0 192.168.155.254
vlan 1-2,12,900-901,950-951,3000
vlan 2
  name Native-VLAN
vlan 12
  name IB-MGMT
vlan 900
  name VM-Traffic-VLAN900
vlan 901
  name VM-Traffic-VLAN901
vlan 950
  name APP1-VM
vlan 951
  name APP2-VM
vlan 3000
  name vMotion
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
vrf context management
port-channel load-balance src-dst ip-l4port-vlan
vpc domain 155
  peer-switch
  role priority 10
  peer-keepalive destination 192.168.155.4 source 192.168.155.3
  peer-gateway
  auto-recovery
  ip arp synchronize
```

```
interface Vlan1
interface Vlan12
  no shutdown
  ip address 192.168.155.252/24
interface port-channel11
  description D01-Mini1-A
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  spanning-tree port type edge trunk
  vpc 11
interface port-channel12
  description D01-Mini1-B
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  spanning-tree port type edge trunk
  vpc 12
interface port-channel13
  description D01-FI-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  spanning-tree port type edge trunk
  spanning-tree guard root
  mtu 9216
  no lacp graceful-convergence
  vpc 13
interface port-channel14
  description D01-FI-B
  switchport mode trunk
  switchport trunk native vlan 2
```

```
  switchport trunk allowed vlan 12,950-951,3000
  spanning-tree port type edge trunk
  spanning-tree guard root
  mtu 9216
  no lacp graceful-convergence
  vpc 14
interface port-channel155
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  spanning-tree port type network
  vpc peer-link
interface Ethernet1/1
  no shutdown
interface Ethernet1/2
  no shutdown
interface Ethernet1/3
  no shutdown
interface Ethernet1/4
  no shutdown
interface Ethernet1/5
  no shutdown
interface Ethernet1/6
  no shutdown
interface Ethernet1/7
  no shutdown
interface Ethernet1/8
  no shutdown
interface Ethernet1/9
  no shutdown
```

```
interface Ethernet1/10
  no shutdown
interface Ethernet1/11
  no shutdown
interface Ethernet1/12
  no shutdown

interface Ethernet1/13
  description ** D01-mini1-a:p1 **
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  channel-group 11 mode active
  no shutdown
interface Ethernet1/14
  description ** D01-mini1-b:p1 **
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  channel-group 12 mode active
  no shutdown
interface Ethernet1/15
  no shutdown
interface Ethernet1/16
  no shutdown
interface Ethernet1/17
  switchport access vlan 12
  no shutdown
interface Ethernet1/18
  no shutdown
interface Ethernet1/19
  no shutdown
interface Ethernet1/20
  no shutdown
```

```
interface Ethernet1/21
  no shutdown
interface Ethernet1/22
  no shutdown
interface Ethernet1/23
  description ** D01-FI-A:p15 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  mtu 9216
  udld enable
  channel-group 13 mode active
  no shutdown
interface Ethernet1/24
  description ** D01-FI-B:p15 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  mtu 9216
  udld enable
  channel-group 14 mode active
  no shutdown
interface Ethernet1/25
  no shutdown
interface Ethernet1/26
  no shutdown
interface Ethernet1/27
  no shutdown
interface Ethernet1/28
  no shutdown
interface Ethernet1/29
```

```
  no shutdown
interface Ethernet1/30
  no shutdown
interface Ethernet1/31
  no shutdown
interface Ethernet1/32
  no shutdown
interface Ethernet1/33
  no shutdown
interface Ethernet1/34
  no shutdown
interface Ethernet1/35
  no shutdown
interface Ethernet1/36
  no shutdown
interface Ethernet1/37
  no shutdown
interface Ethernet1/38
  no shutdown
interface Ethernet1/39
  no shutdown
interface Ethernet1/40
  no shutdown
interface Ethernet1/41
  no shutdown
interface Ethernet1/42
  no shutdown
interface Ethernet1/43
  no shutdown
interface Ethernet1/44
  no shutdown
```

```
interface Ethernet1/45
  no shutdown
interface Ethernet1/46
  no shutdown
interface Ethernet1/47
  no shutdown
interface Ethernet1/48
  description Mgmt via N55k-FEX
  switchport access vlan 12
  spanning-tree port type edge
  spanning-tree bpdufilter enable
  no shutdown
interface Ethernet1/49
  no shutdown
interface Ethernet1/50
  no shutdown
interface Ethernet1/51
  no shutdown
interface Ethernet1/52
  no shutdown
interface Ethernet1/53
  description D01-n9k2:e1/53
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  udld enable
  channel-group 155 mode active
  no shutdown
interface Ethernet1/54
  description D01-n9k2:e1/54
  switchport mode trunk
```

```
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  udld enable
  channel-group 155 mode active
  no shutdown
interface mgmt0
  vrf member management
  ip address 192.168.155.3/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.5.bin
D01-n9k1#
```

Cisco Nexus B Configuration

```
D01-n9k2# show run
!Command: show running-config
!Time: Mon Apr 18 20:42:10 2016
version 6.1(2)I3(5)
hostname D01-n9k2
vdc D01-n9k2 id 1
  allocate interface Ethernet1/1-54
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 512
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
feature nxapi
cfs eth distribute
feature udld
feature interface-vlan
```

```
feature lacp

feature vpc

username admin password 5 $1$aEaZHhoQ$CFqgw6s/wCr8c8Kzb.1DV1  role network-admin

ip domain-lookup

copp profile strict

snmp-server user admin network-admin auth md5 0x5469a3fe245f27f90497c0657c9225fe priv
0x5469a3fe245f27f90497c0657c9225fe localizedkey

rmon event 1 log trap public description FATAL(1) owner PMON@FATAL

rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL

rmon event 3 log trap public description ERROR(3) owner PMON@ERROR

rmon event 4 log trap public description WARNING(4) owner PMON@WARNING

rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

snmp-server community public group network-admin


vlan 1-2,12,900-901,950-951,3000

vlan 2

  name Native-VLAN

vlan 12

  name IB-MGMT

vlan 900

  name VM-Traffic-VLAN900

vlan 901

  name VM-Traffic-VLAN901

vlan 950

  name APP1-VM

vlan 951

  name APP2-VM

vlan 3000

  name vMotion

spanning-tree port type edge bpduguard default

spanning-tree port type edge bpdufilter default

vrf context management
```

```
port-channel load-balance src-dst ip-l4port-vlan
vpc domain 155
  peer-switch
  role priority 20
  peer-keepalive destination 192.168.155.3 source 192.168.155.4
  peer-gateway
  auto-recovery
  ip arp synchronize
interface Vlan1
interface Vlan12
  no shutdown
  no ip redirects
  ip address 192.168.155.253/24
  no ipv6 redirects
interface port-channel11
  description D01-Mini1-A
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  spanning-tree port type edge trunk
  vpc 11
interface port-channel12
  description D01-Mini1-B
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  spanning-tree port type edge trunk
  vpc 12
interface port-channel13
  description D01-FI-A
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
```

```
  spanning-tree port type edge trunk
  spanning-tree guard root
  mtu 9216
  no lacp graceful-convergence
  vpc 13
interface port-channel14
  description D01-FI-B
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  spanning-tree port type edge trunk
  spanning-tree guard root
  mtu 9216
  no lacp graceful-convergence
  vpc 14
interface port-channel155
  description vPC peer-link
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  spanning-tree port type network
  vpc peer-link
interface Ethernet1/1
interface Ethernet1/2
interface Ethernet1/3
interface Ethernet1/4
interface Ethernet1/5
interface Ethernet1/6
interface Ethernet1/7
interface Ethernet1/8
interface Ethernet1/9
```

```
interface Ethernet1/10
interface Ethernet1/11
interface Ethernet1/12
interface Ethernet1/13
  description ** D01-mini1-a:p2 **
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  channel-group 11 mode active
interface Ethernet1/14
  description ** D01-mini1-b:p2 **
  switchport mode trunk
  switchport trunk allowed vlan 12,900-901
  channel-group 12 mode active
interface Ethernet1/15
interface Ethernet1/16
interface Ethernet1/17
  switchport access vlan 12
interface Ethernet1/18
interface Ethernet1/19
interface Ethernet1/20
interface Ethernet1/21
interface Ethernet1/22
interface Ethernet1/23
  description ** D01-FI-A:p16 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  mtu 9216
  udld enable
  channel-group 13 mode active
interface Ethernet1/24
```

```
  description ** D01-FI-B:p16 **
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  mtu 9216
  udld enable
  channel-group 14 mode active
interface Ethernet1/25
interface Ethernet1/26
interface Ethernet1/27
interface Ethernet1/28

interface Ethernet1/29
interface Ethernet1/30
interface Ethernet1/32
interface Ethernet1/33
interface Ethernet1/34
interface Ethernet1/35
interface Ethernet1/36
interface Ethernet1/37
interface Ethernet1/38
interface Ethernet1/39
interface Ethernet1/40
interface Ethernet1/41
interface Ethernet1/42
interface Ethernet1/43
interface Ethernet1/44
interface Ethernet1/45
interface Ethernet1/46
interface Ethernet1/47
interface Ethernet1/48
```

```
  description Mgmt via N55k-FEX
  switchport access vlan 12
  spanning-tree port type edge
  spanning-tree bpdufilter enable
interface Ethernet1/49
interface Ethernet1/50
interface Ethernet1/51
interface Ethernet1/52
interface Ethernet1/53
  description D01-n9k1:e1/53
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  udld enable
  channel-group 155 mode active
interface Ethernet1/54
  description D01-n9k1:e1/54
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,900-901,950-951,3000
  udld enable
  channel-group 155 mode active
interface mgmt0
  vrf member management
  ip address 192.168.155.4/24
line console
line vty
boot nxos bootflash:/n9000-dk9.6.1.2.I3.5.bin
xml server validate all
D01-n9k2#
```

Cisco MDS A Configuration

```
D01-MDS-A# show run

!Command: show running-config
!Time: Wed Aug 10 14:52:27 2016

version 6.2(13b)
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $1$HwFZeqnY$j0Mhyd57z8yNHMGtYwxsn/  role network-admin
no password strength-check
ip domain-lookup
ip host D01-MDS-A  192.168.155.6
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xa1265df8e082fcaab00242a6eed11170 priv
0xa1265df8e082fcaab00242a6eed11170 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-operator
vsan database
  vsan 4091
device-alias confirm-commit enable
device-alias database
  device-alias name AFA-CNTLA-p1 pwwn 56:c9:ce:90:49:09:db:01
  device-alias name AFA-CNTLA-p2 pwwn 56:c9:ce:90:49:09:db:03
  device-alias name AFA-CNTLB-p1 pwwn 56:c9:ce:90:49:09:db:05
  device-alias name AFA-CNTLB-p2 pwwn 56:c9:ce:90:49:09:db:07
  device-alias name AppVMHost-FIA-0 pwwn 20:00:00:25:b5:11:aa:04
  device-alias name AppVMHost-FIA-1 pwwn 20:00:00:25:b5:11:aa:00
  device-alias name AppVMHost-FIB-0 pwwn 20:00:00:25:b5:11:aa:05
  device-alias name AppVMHost-FIB-1 pwwn 20:00:00:25:b5:11:aa:01
  device-alias name NIMBLE-CNTLA-p1 pwwn 56:c9:ce:90:eb:af:a8:01
  device-alias name NIMBLE-CNTLA-p2 pwwn 56:c9:ce:90:eb:af:a8:03
  device-alias name NIMBLE-CNTLB-p1 pwwn 56:c9:ce:90:eb:af:a8:05
  device-alias name NIMBLE-CNTLB-p2 pwwn 56:c9:ce:90:eb:af:a8:07
  device-alias name AFA-AppVMHost-FIA-3-A pwwn 20:00:00:25:b5:11:aa:06
  device-alias name AFA-AppVMHost-FIA-3-B pwwn 20:00:00:25:b5:11:bb:06
  device-alias name AFA-AppVMHost-FIB-4-A pwwn 20:00:00:25:b5:11:aa:07
  device-alias name AFA-AppVMHost-FIB-4-B pwwn 20:00:00:25:b5:11:bb:07

device-alias commit

fcdomain fcid database
  vsan 4091 wwn 20:1d:00:2a:6a:41:9c:80 fcid 0x940000 dynamic
  vsan 4091 wwn 20:1e:00:2a:6a:41:9c:80 fcid 0x940100 dynamic
  vsan 4091 wwn 20:1f:00:2a:6a:41:9c:80 fcid 0x940200 dynamic
  vsan 4091 wwn 20:20:00:2a:6a:41:9c:80 fcid 0x940300 dynamic
  vsan 4091 wwn 24:29:00:2a:6a:41:9c:80 fcid 0x940400 dynamic
  vsan 4091 wwn 20:1c:00:2a:6a:41:9c:80 fcid 0x940500 dynamic
  vsan 4091 wwn 56:c9:ce:90:eb:af:a8:01 fcid 0x940600 dynamic
!               [NIMBLE-CNTLA-p1]
  vsan 4091 wwn 56:c9:ce:90:eb:af:a8:03 fcid 0x940700 dynamic
!               [NIMBLE-CNTLA-p2]
```

```
  vsan 4091 wwn 56:c9:ce:90:eb:af:a8:05 fcid 0x940800 dynamic
!               [NIMBLE-CNTLB-p1]
  vsan 4091 wwn 56:c9:ce:90:eb:af:a8:07 fcid 0x940900 dynamic
!               [NIMBLE-CNTLB-p2]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:00 fcid 0x940401 dynamic
!               [AppVMHost-FIA-1]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:01 fcid 0x940402 dynamic
!               [AppVMHost-FIB-1]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:02 fcid 0x940403 dynamic
  vsan 4091 wwn 20:00:00:25:b5:11:aa:03 fcid 0x940404 dynamic
  vsan 4091 wwn 20:00:00:25:b5:11:aa:04 fcid 0x940405 dynamic
!               [AppVMHost-FIA-0]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:05 fcid 0x940406 dynamic
!               [AppVMHost-FIB-0]
  vsan 1 wwn 56:c9:ce:90:49:09:db:07 fcid 0x120000 dynamic
!            [AFA-CNTLB-p2]
  vsan 1 wwn 56:c9:ce:90:49:09:db:05 fcid 0x120100 dynamic
!            [AFA-CNTLB-p1]
  vsan 1 wwn 56:c9:ce:90:49:09:db:03 fcid 0x120200 dynamic
!            [AFA-CNTLA-p2]
  vsan 1 wwn 56:c9:ce:90:49:09:db:01 fcid 0x120300 dynamic
!            [AFA-CNTLA-p1]
  vsan 4091 wwn 56:c9:ce:90:49:09:db:01 fcid 0x940a00 dynamic
!               [AFA-CNTLA-p1]
  vsan 4091 wwn 56:c9:ce:90:49:09:db:03 fcid 0x940b00 dynamic
!               [AFA-CNTLA-p2]
  vsan 4091 wwn 56:c9:ce:90:49:09:db:05 fcid 0x940c00 dynamic
!               [AFA-CNTLB-p1]
  vsan 4091 wwn 56:c9:ce:90:49:09:db:07 fcid 0x940d00 dynamic
!               [AFA-CNTLB-p2]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:06 fcid 0x940407 dynamic
!               [AFA-AppVMHost-FIA-3-A]
  vsan 4091 wwn 20:00:00:25:b5:11:aa:07 fcid 0x940408 dynamic
!               [AFA-AppVMHost-FIB-4-A]


interface mgmt0
  ip address 192.168.155.6 255.255.255.0

interface port-channel41
  channel mode active
  switchport description 8G-PortChannel-to-FI-A via fc1/45-48
  switchport rate-mode dedicated
vsan database
  vsan 4091 interface port-channel41
  vsan 4091 interface fc1/1
  vsan 4091 interface fc1/2
  vsan 4091 interface fc1/3
  vsan 4091 interface fc1/4
  vsan 4091 interface fc1/13
  vsan 4091 interface fc1/14
  vsan 4091 interface fc1/15
  vsan 4091 interface fc1/16
clock timezone EST -5 0
switchname D01-MDS-A
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.13b.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.13b.bin
interface fc1/1
  switchport speed 8000
interface fc1/2
  switchport speed 8000
interface fc1/3
  switchport speed 8000
interface fc1/4
```

```
  switchport speed 8000
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
zoneset distribute full vsan 4091
!Active Zone Database Section for vsan 4091
zone name AppVMHost-FIA-0 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:04
!               [AppVMHost-FIA-0]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]
```

```
zone name AppVMHost-FIB-0 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:05
!               [AppVMHost-FIB-0]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIA-1 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:00
!               [AppVMHost-FIA-1]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-1 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:01
!               [AppVMHost-FIB-1]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AFA-AppVMHost-FIA-3-A vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:06
!               [AFA-AppVMHost-FIA-3-A]
    member pwwn 56:c9:ce:90:49:09:db:01
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:03
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:05
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:07
!               [AFA-CNTLB-p2]

zone name AFA-AppVMHost-FIB-4-A vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:07
!               [AFA-AppVMHost-FIB-4-A]
    member pwwn 56:c9:ce:90:49:09:db:01
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:03
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:05
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:07
!               [AFA-CNTLB-p2]

zoneset name Fabric-A vsan 4091
    member AppVMHost-FIA-0
    member AppVMHost-FIB-0
    member AppVMHost-FIA-1
    member AppVMHost-FIB-1
```

```
    member AFA-AppVMHost-FIA-3-A
    member AFA-AppVMHost-FIB-4-A

zoneset activate name Fabric-A vsan 4091
do clear zone database vsan 4091
!Full Zone Database Section for vsan 4091
zone name AppVMHost-FIA-0 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:04
!               [AppVMHost-FIA-0]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-0 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:05
!               [AppVMHost-FIB-0]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIA-1 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:00
!               [AppVMHost-FIA-1]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-1 vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:01
!               [AppVMHost-FIB-1]
    member pwwn 56:c9:ce:90:eb:af:a8:01
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:03
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:05
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:07
!               [NIMBLE-CNTLB-p2]

zone name af7000-ctrla-fc1-1 vsan 4091
zone name AFA-AppVMHost-FIA-3-A vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:06
!               [AFA-AppVMHost-FIA-3-A]
    member pwwn 56:c9:ce:90:49:09:db:01
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:03
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:05
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:07
!               [AFA-CNTLB-p2]
```

```
zone name AFA-AppVMHost-FIB-4-A vsan 4091
    member pwwn 20:00:00:25:b5:11:aa:07
!               [AFA-AppVMHost-FIB-4-A]
    member pwwn 56:c9:ce:90:49:09:db:01
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:03
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:05
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:07
!               [AFA-CNTLB-p2]

zoneset name Fabric-A vsan 4091
    member AppVMHost-FIA-0
    member AppVMHost-FIB-0
    member AppVMHost-FIA-1
    member AppVMHost-FIB-1
    member AFA-AppVMHost-FIA-3-A
    member AFA-AppVMHost-FIB-4-A


interface fc1/1
  switchport description CS700-CNTLA-fc1.1
  port-license acquire
  no shutdown

interface fc1/2
  switchport description CS700-CNTLA-fc5.1
  port-license acquire
  no shutdown

interface fc1/3
  switchport description CS700-CNTLB-fc1.1
  port-license acquire
  no shutdown

interface fc1/4
  switchport description CS700-CNTLB-fc5.1
  port-license acquire
  no shutdown

interface fc1/5
  no port-license

interface fc1/6
  no port-license

interface fc1/7
  no port-license

interface fc1/8
  no port-license

interface fc1/9
  no port-license

interface fc1/10
  no port-license

interface fc1/11
  no port-license

interface fc1/12
  no port-license
```

```
interface fc1/13
  switchport description AFA-CNTLA-fc1.1
  port-license acquire
  no shutdown

interface fc1/14
  switchport description AFA-CNTLA-fc5.1
  port-license acquire
  no shutdown

interface fc1/15
  switchport description AFA-CNTLB-fc1.1
  port-license acquire
  no shutdown

interface fc1/16
  switchport description AFA-CNTLB-fc5.1
  port-license acquire
  no shutdown

interface fc1/17
  port-license acquire

interface fc1/18
  port-license acquire

interface fc1/19
  port-license acquire

interface fc1/20
  port-license acquire

interface fc1/21
  port-license acquire

interface fc1/22
  port-license acquire

interface fc1/23
  port-license acquire

interface fc1/24
  port-license acquire

interface fc1/25

interface fc1/26

interface fc1/27

interface fc1/28

interface fc1/29

interface fc1/30

interface fc1/31

interface fc1/32

interface fc1/33

interface fc1/34

interface fc1/35
```

```
interface fc1/36

interface fc1/37

interface fc1/38

interface fc1/39

interface fc1/40

interface fc1/41
  switchport description D01-FI-A:p25
  port-license acquire
  no shutdown

interface fc1/42
  switchport description D01-FI-A:p26
  port-license acquire
  no shutdown

interface fc1/43
  switchport description D01-FI-A:p27
  port-license acquire
  no shutdown

interface fc1/44
  switchport description D01-FI-A:p28
  port-license acquire
  no shutdown

interface fc1/45
  switchport description FI-A:p1/29
  port-license acquire
  channel-group 41 force
  no shutdown

interface fc1/46
  switchport description FI-A:p1/30
  port-license acquire
  channel-group 41 force
  no shutdown

interface fc1/47
  switchport description FI-A:p1/31
  port-license acquire
  channel-group 41 force
  no shutdown

interface fc1/48
  switchport description FI-A:p1/32
  port-license acquire
  channel-group 41 force
  no shutdown
ip default-gateway 192.168.155.1

D01-MDS-A#
```

Cisco MDS B Configuration

```
D01-MDS-B# show run

!Command: show running-config
!Time: Wed Aug 10 20:04:17 2016

version 6.2(13b)
```

```
power redundancy-mode redundant
feature npiv
feature fport-channel-trunk
role name default-role
  description This is a system defined role and applies to all users.
  rule 5 permit show feature environment
  rule 4 permit show feature hardware
  rule 3 permit show feature module
  rule 2 permit show feature snmp
  rule 1 permit show feature system
username admin password 5 $1$lAA1rRTG$YLpfWjWbBldDtiGrUwXp41  role network-admin
ip domain-lookup
ip host D01-MDS-B  192.168.155.7
aaa group server radius radius
snmp-server user admin network-admin auth md5 0xf1d93b9e9d8b066cd3e7a12a29b862f4 priv
0xf1d93b9e9d8b066cd3e7a12a29b862f4 localizedkey
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO
snmp-server community public group network-operator
vsan database
  vsan 4092
device-alias confirm-commit enable
device-alias database
  device-alias name AFA-CNTLA-p1 pwwn 56:c9:ce:90:49:09:db:02
  device-alias name AFA-CNTLA-p2 pwwn 56:c9:ce:90:49:09:db:04
  device-alias name AFA-CNTLB-p1 pwwn 56:c9:ce:90:49:09:db:06
  device-alias name AFA-CNTLB-p2 pwwn 56:c9:ce:90:49:09:db:08
  device-alias name AppVMHost-FIA-0 pwwn 20:00:00:25:b5:11:bb:04
  device-alias name AppVMHost-FIA-1 pwwn 20:00:00:25:b5:11:bb:00
  device-alias name AppVMHost-FIB-0 pwwn 20:00:00:25:b5:11:bb:05
  device-alias name AppVMHost-FIB-1 pwwn 20:00:00:25:b5:11:bb:01
  device-alias name NIMBLE-CNTLA-p1 pwwn 56:c9:ce:90:eb:af:a8:02
  device-alias name NIMBLE-CNTLA-p2 pwwn 56:c9:ce:90:eb:af:a8:04
  device-alias name NIMBLE-CNTLB-p1 pwwn 56:c9:ce:90:eb:af:a8:06
  device-alias name NIMBLE-CNTLB-p2 pwwn 56:c9:ce:90:eb:af:a8:08
  device-alias name AFA-AppVMHost-FIA-3-A pwwn 20:00:00:25:b5:11:aa:06
  device-alias name AFA-AppVMHost-FIA-3-B pwwn 20:00:00:25:b5:11:bb:06
  device-alias name AFA-AppVMHost-FIB-4-A pwwn 20:00:00:25:b5:11:aa:07
  device-alias name AFA-AppVMHost-FIB-4-B pwwn 20:00:00:25:b5:11:bb:07

device-alias commit

fcdomain fcid database
  vsan 4092 wwn 24:2a:00:2a:6a:41:9b:c0 fcid 0x4b0000 dynamic
  vsan 1 wwn 20:1c:00:2a:6a:41:9b:c0 fcid 0x2b0000 dynamic
  vsan 4092 wwn 20:1d:00:2a:6a:41:9b:c0 fcid 0x4b0100 dynamic
  vsan 4092 wwn 20:1e:00:2a:6a:41:9b:c0 fcid 0x4b0200 dynamic
  vsan 1 wwn 20:1c:00:2a:6a:41:9c:80 fcid 0x2b0100 dynamic
  vsan 4092 wwn 56:c9:ce:90:eb:af:a8:08 fcid 0x4b0300 dynamic
!               [NIMBLE-CNTLB-p2]
  vsan 4092 wwn 56:c9:ce:90:eb:af:a8:04 fcid 0x4b0400 dynamic
!               [NIMBLE-CNTLA-p2]
  vsan 4092 wwn 56:c9:ce:90:eb:af:a8:06 fcid 0x4b0500 dynamic
!               [NIMBLE-CNTLB-p1]
  vsan 4092 wwn 56:c9:ce:90:eb:af:a8:02 fcid 0x4b0600 dynamic
!               [NIMBLE-CNTLA-p1]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:00 fcid 0x4b0001 dynamic
!               [AppVMHost-FIA-1]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:01 fcid 0x4b0002 dynamic
!               [AppVMHost-FIB-1]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:02 fcid 0x4b0003 dynamic
  vsan 4092 wwn 20:00:00:25:b5:11:bb:03 fcid 0x4b0004 dynamic
  vsan 4092 wwn 20:00:00:25:b5:11:bb:05 fcid 0x4b0005 dynamic
```

```
!               [AppVMHost-FIB-0]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:04 fcid 0x4b0006 dynamic
!               [AppVMHost-FIA-0]
  vsan 1 wwn 56:c9:ce:90:49:09:db:08 fcid 0x2b0200 dynamic
!             [AFA-CNTLB-p2]
  vsan 1 wwn 56:c9:ce:90:49:09:db:04 fcid 0x2b0300 dynamic
!             [AFA-CNTLA-p2]
  vsan 1 wwn 56:c9:ce:90:49:09:db:06 fcid 0x2b0400 dynamic
!             [AFA-CNTLB-p1]
  vsan 1 wwn 56:c9:ce:90:49:09:db:02 fcid 0x2b0500 dynamic
!             [AFA-CNTLA-p1]
  vsan 4092 wwn 56:c9:ce:90:49:09:db:02 fcid 0x4b0700 dynamic
!               [AFA-CNTLA-p1]
  vsan 4092 wwn 56:c9:ce:90:49:09:db:04 fcid 0x4b0800 dynamic
!               [AFA-CNTLA-p2]
  vsan 4092 wwn 56:c9:ce:90:49:09:db:06 fcid 0x4b0900 dynamic
!               [AFA-CNTLB-p1]
  vsan 4092 wwn 56:c9:ce:90:49:09:db:08 fcid 0x4b0a00 dynamic
!               [AFA-CNTLB-p2]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:06 fcid 0x4b0007 dynamic
!               [AFA-AppVMHost-FIA-3-B]
  vsan 4092 wwn 20:00:00:25:b5:11:bb:07 fcid 0x4b0008 dynamic
!               [AFA-AppVMHost-FIB-4-B]


interface mgmt0
  ip address 192.168.155.7 255.255.255.0

interface port-channel42
  channel mode active
  switchport description 8G-PortChannel-to-FI-B via fc1/45-48
  switchport rate-mode dedicated
vsan database
  vsan 4092 interface port-channel42
  vsan 4092 interface fc1/1
  vsan 4092 interface fc1/2
  vsan 4092 interface fc1/3
  vsan 4092 interface fc1/4
  vsan 4092 interface fc1/13
  vsan 4092 interface fc1/14
  vsan 4092 interface fc1/15
  vsan 4092 interface fc1/16
switchname D01-MDS-B
line console
line vty
boot kickstart bootflash:/m9100-s5ek9-kickstart-mz.6.2.13b.bin
boot system bootflash:/m9100-s5ek9-mz.6.2.13b.bin
interface fc1/1
  switchport speed 8000
interface fc1/2
  switchport speed 8000
interface fc1/3
  switchport speed 8000
interface fc1/4
  switchport speed 8000
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
interface fc1/10
interface fc1/11
```

```
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface fc1/17
interface fc1/18
interface fc1/19
interface fc1/20
interface fc1/21
interface fc1/22
interface fc1/23
interface fc1/24
interface fc1/25
interface fc1/26
interface fc1/27
interface fc1/28
interface fc1/29
interface fc1/30
interface fc1/31
interface fc1/32
interface fc1/33
interface fc1/34
interface fc1/35
interface fc1/36
interface fc1/37
interface fc1/38
interface fc1/39
interface fc1/40
interface fc1/41
interface fc1/42
interface fc1/43
interface fc1/44
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/45
interface fc1/46
interface fc1/47
interface fc1/48
zoneset distribute full vsan 4092
!Active Zone Database Section for vsan 4092
zone name AppVMHost-FIA-0 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:04
!               [AppVMHost-FIA-0]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-0 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:05
!               [AppVMHost-FIB-0]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]
```

```
zone name AppVMHost-FIA-1 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:00
!               [AppVMHost-FIA-1]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-1 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:01
!               [AppVMHost-FIB-1]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AFA-AppVMHost-FIA-3-B vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:06
!               [AFA-AppVMHost-FIA-3-B]
    member pwwn 56:c9:ce:90:49:09:db:02
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:04
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:06
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:08
!               [AFA-CNTLB-p2]

zone name AFA-AppVMHost-FIB-4-B vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:07
!               [AFA-AppVMHost-FIB-4-B]
    member pwwn 56:c9:ce:90:49:09:db:02
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:04
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:06
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:08
!               [AFA-CNTLB-p2]

zoneset name Fabric-B vsan 4092
    member AppVMHost-FIA-0
    member AppVMHost-FIB-0
    member AppVMHost-FIA-1
    member AppVMHost-FIB-1
    member AFA-AppVMHost-FIA-3-B
    member AFA-AppVMHost-FIB-4-B

zoneset activate name Fabric-B vsan 4092
do clear zone database vsan 4092
!Full Zone Database Section for vsan 4092
zone name AppVMHost-FIA-0 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:04
!               [AppVMHost-FIA-0]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
```

```
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-0 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:05
!               [AppVMHost-FIB-0]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIA-1 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:00
!               [AppVMHost-FIA-1]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AppVMHost-FIB-1 vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:01
!               [AppVMHost-FIB-1]
    member pwwn 56:c9:ce:90:eb:af:a8:02
!               [NIMBLE-CNTLA-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:04
!               [NIMBLE-CNTLA-p2]
    member pwwn 56:c9:ce:90:eb:af:a8:06
!               [NIMBLE-CNTLB-p1]
    member pwwn 56:c9:ce:90:eb:af:a8:08
!               [NIMBLE-CNTLB-p2]

zone name AFA-AppVMHost-FIA-3-B vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:06
!               [AFA-AppVMHost-FIA-3-B]
    member pwwn 56:c9:ce:90:49:09:db:02
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:04
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:06
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:08
!               [AFA-CNTLB-p2]

zone name AFA-AppVMHost-FIB-4-B vsan 4092
    member pwwn 20:00:00:25:b5:11:bb:07
!               [AFA-AppVMHost-FIB-4-B]
    member pwwn 56:c9:ce:90:49:09:db:02
!               [AFA-CNTLA-p1]
    member pwwn 56:c9:ce:90:49:09:db:04
!               [AFA-CNTLA-p2]
    member pwwn 56:c9:ce:90:49:09:db:06
!               [AFA-CNTLB-p1]
    member pwwn 56:c9:ce:90:49:09:db:08
!               [AFA-CNTLB-p2]
```

```
zoneset name Fabric-B vsan 4092
    member AppVMHost-FIA-0
    member AppVMHost-FIB-0
    member AppVMHost-FIA-1
    member AppVMHost-FIB-1
    member AFA-AppVMHost-FIA-3-B
    member AFA-AppVMHost-FIB-4-B


interface fc1/1
  switchport description CS700-CNTLA-fc2.1
  port-license acquire
  no shutdown

interface fc1/2
  switchport description CS700-CNTLA-fc6.1
  port-license acquire
  no shutdown

interface fc1/3
  switchport description CS700-CNTLB-fc2.1
  port-license acquire
  no shutdown

interface fc1/4
  switchport description CS700-CNTLB-fc6.1
  port-license acquire
  no shutdown

interface fc1/5
  no port-license

interface fc1/6
  no port-license

interface fc1/7
  no port-license

interface fc1/8
  no port-license

interface fc1/9
  no port-license

interface fc1/10
  no port-license

interface fc1/11
  no port-license

interface fc1/12
  no port-license

interface fc1/13
  switchport description AFA-CNTLA-fc2.1
  port-license acquire
  no shutdown

interface fc1/14
  switchport description AFA-CNTLA-fc6.1
  port-license acquire
  no shutdown

interface fc1/15
  switchport description AFA-CNTLB-fc2.1
  port-license acquire
```

```
  no shutdown

interface fc1/16
  switchport description AFA-CNTLB-fc6.1
  port-license acquire
  no shutdown

interface fc1/17
  port-license acquire

interface fc1/18
  port-license acquire

interface fc1/19
  port-license acquire

interface fc1/20
  port-license acquire

interface fc1/21
  port-license acquire

interface fc1/22
  port-license acquire

interface fc1/23
  port-license acquire

interface fc1/24
  port-license acquire

interface fc1/25

interface fc1/26

interface fc1/27

interface fc1/28

interface fc1/29

interface fc1/30

interface fc1/31

interface fc1/32

interface fc1/33

interface fc1/34

interface fc1/35

interface fc1/36

interface fc1/37

interface fc1/38

interface fc1/39

interface fc1/40

interface fc1/41
  switchport description D01-FI-B:p25
  port-license acquire
```

```
interface fc1/42
  switchport description D01-FI-B:p26
  port-license acquire

interface fc1/43
  switchport description D01-FI-B:p27
  port-license acquire

interface fc1/44
  switchport description D01-FI-B:p28
  port-license acquire

interface fc1/45
  switchport description FI-B:p1/29
  port-license acquire
  channel-group 42 force
  no shutdown

interface fc1/46
  switchport description FI-B:p1/30
  port-license acquire
  channel-group 42 force
  no shutdown

interface fc1/47
  switchport description FI-B:p1/31
  port-license acquire
  channel-group 42 force
  no shutdown

interface fc1/48
  switchport description FI-B:p1/32
  port-license acquire
  channel-group 42 force
  no shutdown
ip default-gateway 192.168.155.1

    D01-MDS-B#
```

Cisco Nexus 1000V Configuration

```
SS-CVD-N1k# show run

!Command: show running-config
!Time: Fri Aug 12 11:04:54 2016

version 5.2(1)SV3(1.15)
hostname SS-CVD-N1k

no feature telnet

username admin password 5 $1$.GFMx/Wt$OZ5sb.dLwU3bJTB5XyOyh/  role network-admin
username admin keypair generate rsa
username n1kvmgr password 5 $1$Yn/5np26$fhZCk0paPPxCoywhRI1Ro/  role network-operator
username n1kvmgr role network-admin

banner motd #Nexus 1000v Switch
#

ssh key rsa 2048
ip domain-lookup
ip host SS-CVD-N1k 192.168.155.51
errdisable recovery cause failed-port-state
vem 3
```

```
  host id 6274b68c-0ceb-e211-0000-00000000001d
vem 4
  host id 6274b68c-0ceb-e211-0000-00000000000c
vem 5
  host id 6274b68c-0ceb-e211-0000-00000000000d
vem 6
  host id 6274b68c-0ceb-e211-0000-00000000001c
snmp-server user admin network-admin auth md5 0xb1024c55cea9a09ad2b36e574a0a1ceb priv
0xb1024c55cea9a09ad2b36e574a0a1ceb localizedkey
snmp-server user n1kvmgr network-operator auth md5 0x247613c56d19ba3e048b13301a924047 priv
0x247613c56d19ba3e048b13301a924047 localizedkey
snmp-server user n1kvmgr network-admin
rmon event 1 log trap public description FATAL(1) owner PMON@FATAL
rmon event 2 log trap public description CRITICAL(2) owner PMON@CRITICAL
rmon event 3 log trap public description ERROR(3) owner PMON@ERROR
rmon event 4 log trap public description WARNING(4) owner PMON@WARNING
rmon event 5 log trap public description INFORMATION(5) owner PMON@INFO

vrf context management
  ip route 0.0.0.0/0 192.168.155.1
vlan 1-2,12,950-951,3000
vlan 2
  name Native-VLAN
vlan 12
  name IB-MGMT-VLAN
vlan 950
  name APP1-VM-VLAN
vlan 951
  name APP2-VM-VLAN
vlan 3000
  name vMotion-VLAN

port-channel load-balance ethernet source-mac
port-profile default max-ports 32
port-profile type ethernet Unused_Or_Quarantine_Uplink
  shutdown
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type vethernet Unused_Or_Quarantine_Veth
  shutdown
  description Port-group created for Nexus 1000V internal usage. Do not use.
  state enabled
  vmware port-group
port-profile type ethernet system-uplink
  switchport mode trunk
  switchport trunk native vlan 2
  switchport trunk allowed vlan 12,950-951,3000
  system mtu 9000
  channel-group auto mode on mac-pinning
  no shutdown
  system vlan 12,950-951,3000
  state enabled
  vmware port-group
port-profile type vethernet IB-MGMT-VLAN
  switchport mode access
  switchport access vlan 12
  no shutdown
  system vlan 12
  state enabled
  vmware port-group IB-MGMT-PG
port-profile type vethernet vMotion-VLAN
  switchport mode access
  switchport access vlan 3000
  no shutdown
  system vlan 3000
```

```
  state enabled
  vmware port-group vMotion-PG
port-profile type vethernet APP1-VM-VLAN
  switchport mode access
  switchport access vlan 950
  no shutdown
  system vlan 950
  state enabled
  vmware port-group APP1-VM-PG
port-profile type vethernet APP2-VM-VLAN
  switchport mode access
  switchport access vlan 951
  no shutdown
  system vlan 951
  state enabled
  vmware port-group APP2-VM-PG
port-profile type vethernet n1kv-L3
  switchport mode access
  switchport access vlan 12
  no shutdown
  capability l3control
  system vlan 12
  state enabled
  vmware port-group


interface port-channel1
  inherit port-profile system-uplink
  vem 3
  mtu 9000

interface port-channel2
  inherit port-profile system-uplink
  vem 5
  mtu 9000

interface port-channel3
  inherit port-profile system-uplink
  vem 6
  mtu 9000

interface port-channel4
  inherit port-profile system-uplink
  vem 4
  mtu 9000

interface mgmt0
  ip address 192.168.155.51/24

interface Vethernet1
  inherit port-profile n1kv-L3
  description VMware VMkernel, vmk0
  vmware dvport 160 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0025.B5AA.AA0D

interface Vethernet2
  inherit port-profile vMotion-VLAN
  description VMware VMkernel, vmk1
  vmware dvport 64 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.5664.95C0

interface Vethernet3
  inherit port-profile n1kv-L3
  description VMware VMkernel, vmk0
  vmware dvport 161 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0025.B5AA.AA1D
```

```
interface Vethernet4
  inherit port-profile vMotion-VLAN
  description VMware VMkernel, vmk1
  vmware dvport 65 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.566D.A653

interface Vethernet5
  inherit port-profile n1kv-L3
  description VMware VMkernel, vmk0
  vmware dvport 162 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0025.B5AA.AA0C

interface Vethernet6
  inherit port-profile vMotion-VLAN
  description VMware VMkernel, vmk1
  vmware dvport 66 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.5660.506F

interface Vethernet7
  inherit port-profile n1kv-L3
  description VMware VMkernel, vmk1
  vmware dvport 163 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0025.B5AA.AA1C

interface Vethernet9
  inherit port-profile IB-MGMT-VLAN
  description IOM1, Network Adapter 1
  vmware dvport 32 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.76D5

interface Vethernet10
  inherit port-profile IB-MGMT-VLAN
  description IOM2, Network Adapter 1
  vmware dvport 33 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.882B

interface Vethernet11
  inherit port-profile IB-MGMT-VLAN
  description IOM3, Network Adapter 1
  vmware dvport 34 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.434B

interface Vethernet12
  inherit port-profile IB-MGMT-VLAN
  description IOM4, Network Adapter 1
  vmware dvport 35 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.5505

interface Vethernet13
  inherit port-profile APP1-VM-VLAN
  description IOM2, Network Adapter 2
  vmware dvport 96 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.6EF0

interface Vethernet14
  inherit port-profile APP1-VM-VLAN
  description IOM1, Network Adapter 2
  vmware dvport 97 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.A905

interface Vethernet15
  inherit port-profile APP2-VM-VLAN
  description IOM3, Network Adapter 2
  vmware dvport 128 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.505C
```

```
interface Vethernet16
  inherit port-profile APP2-VM-VLAN
  description IOM4, Network Adapter 2
  vmware dvport 129 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.569E.36A9

interface Vethernet17
  inherit port-profile vMotion-VLAN
  description VMware VMkernel, vmk2
  vmware dvport 67 dvswitch uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e"
  vmware vm mac 0050.5667.E218

interface Ethernet3/1
  inherit port-profile system-uplink

interface Ethernet3/2
  inherit port-profile system-uplink

interface Ethernet4/1
  inherit port-profile system-uplink

interface Ethernet4/2
  inherit port-profile system-uplink

interface Ethernet5/1
  inherit port-profile system-uplink

interface Ethernet5/2
  inherit port-profile system-uplink

interface Ethernet6/1
  inherit port-profile system-uplink

interface Ethernet6/2
  inherit port-profile system-uplink

interface control0
line console
line vty
boot kickstart bootflash:/n1000v-dk9-kickstart.5.2.1.SV3.1.15.bin sup-1
boot system bootflash:/n1000v-dk9.5.2.1.SV3.1.15.bin sup-1
boot kickstart bootflash:/n1000v-dk9-kickstart.5.2.1.SV3.1.15.bin sup-2
boot system bootflash:/n1000v-dk9.5.2.1.SV3.1.15.bin sup-2
svs-domain
  domain id 1000
  control vlan 1
  packet vlan 1
  svs mode L3 interface mgmt0
  switch-guid 7c3825d2-4c2f-45c2-b873-0f310e07a482
  enable l3sec
svs connection vCenter
  protocol vmware-vim
  remote ip address 192.168.155.13 port 80
  vmware dvs uuid "37 5a 1e 50 27 a1 6a b2-6e 68 35 d9 c3 2f 55 5e" datacenter-name SmartStack
  max-ports 12000
  connect
vservice global type vsg
  no tcp state-checks invalid-ack
  no tcp state-checks seq-past-window
  no tcp state-checks window-variation
  no bypass asa-traffic
  no l3-frag
vservice global
  idle-timeout
    tcp 30
```

```
    udp 4
    icmp 4
    layer-3 4
    layer-2 2
nsc-policy-agent
  registration-ip 0.0.0.0
  shared-secret **********
  log-level
```

About Authors

Archana Sharma, Technical Leader, Cisco UCS Solutions Engineering, Cisco Systems Inc.

Archana Sharma has 20 years of experience at Cisco focused on Data Center, Desktop Virtualization, Collaboration and related technologies. Archana has been working on Enterprise and Service Provider systems and solutions and delivering Cisco Validated designs for over 10 years. Archana holds a CCIE (#3080) in Routing and Switching and a Bachelor's degree in Electrical Engineering from North Carolina State University.

Jay White, Principal Technical Marketing Engineer, Nimble Storage Inc.

Jay has 20 years of experience in both the network and storage industries, including roles in Engineering, technical Sales, data center consulting, and Technical Marketing. Jay leads the SmartStack solutions related technical activity at Nimble Storage. In the past, he has provided subject matter expertise on nearly all aspects of network storage systems, including performance, file systems, protocols, storage efficiency, disaster recovery, and more.

Acknowledgements

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

- Chris O' Brien, Director, Cisco UCS Solutions Technical Marketing Team, Cisco Systems Inc.
- Jawwad Memon, Product Manager, Cisco UCS Product Management and Solutions, Cisco Systems Inc.
- Ashish Prakash, VP Solutions Engineering, Nimble Storage Inc.
- Arun Garg, Director, Solutions Product Management, Nimble Storage Inc.
- Matt Miller, Director, Solutions Marketing, Nimble Storage Inc.
- Steve Sexton, Senior Technical Marketing Engineer, Nimble Storage Inc.