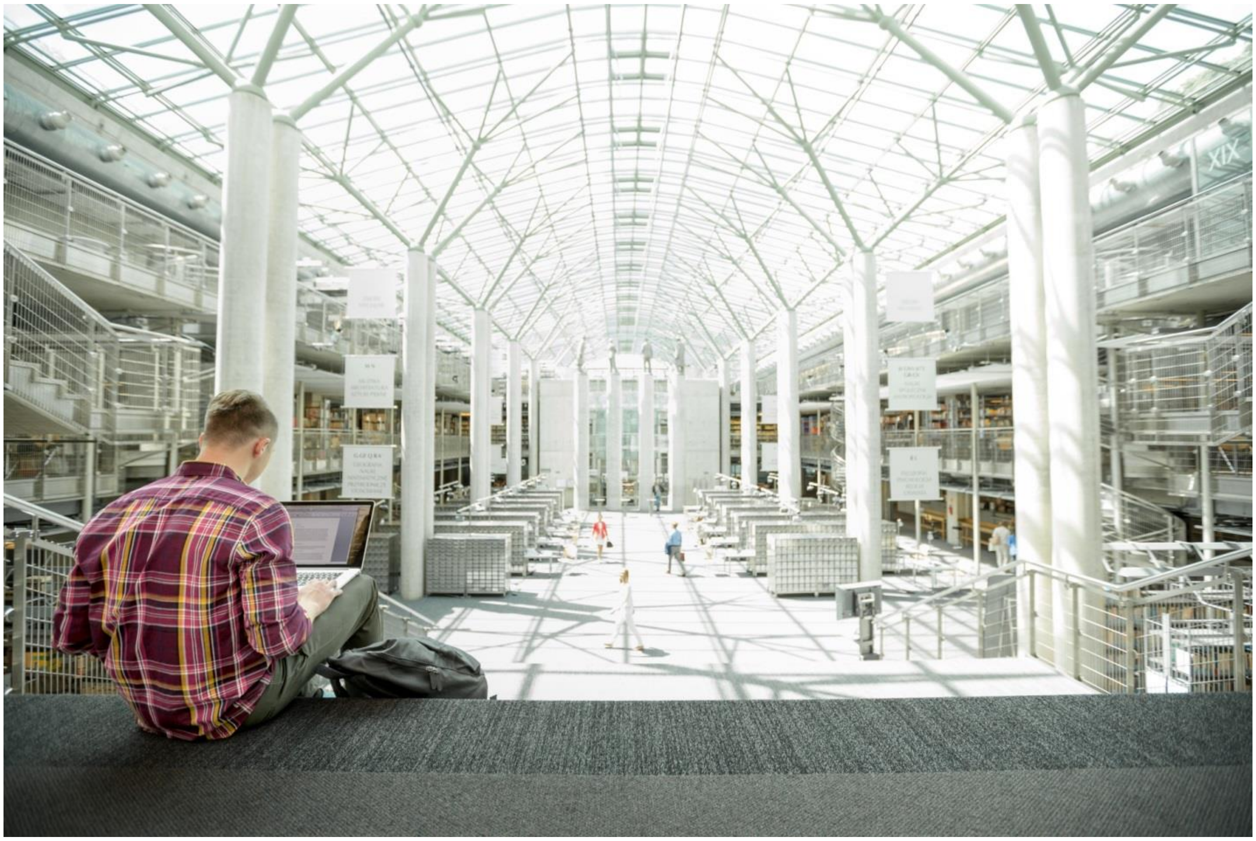

FlexPod Datacenter with NetApp All Flash FAS, Cisco Application Centric Infrastructure, and VMware vSphere

Deployment Guide for FlexPod with NetApp All Flash FAS and Cisco Application Centric Infrastructure and VMware vSphere 5.5U2

Last Updated: August 25, 2015

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.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

Table of Contents

Executive Summary

Cisco® Validated Designs include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

This document describes the Cisco and NetApp® FlexPod Datacenter with NetApp All Flash FAS (AFF), Cisco Application Centric Infrastructure (ACI), and VMware vSphere 5.5 Update 2. FlexPod Datacenter with NetApp AFF and Cisco ACI is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus® 9000 family of switches, and NetApp AFF.

Solution Overview

Introduction

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. Business agility requires application agility, so IT teams need to provision applications in hours instead of months. Resources need to scale up (or down) in minutes, not hours.

To simplify the evolution to a shared cloud infrastructure based on an application driven policy model, Cisco and NetApp have developed the solution called FlexPod Datacenter with NetApp AFF and Cisco ACI. Cisco ACI provides a holistic architecture with centralized automation and policy-driven application profiles that delivers software flexibility with hardware performance. NetApp All Flash FAS addresses enterprise storage requirements with high performance, superior flexibility, and best-in-class data management.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This document provides a step-by-step configuration and implementation guide for the FlexPod Datacenter with NetApp AFF and Cisco ACI solution. For the design decisions and technology discussion of the solution, please refer to FlexPod Datacenter with NetApp All Flash FAS, Cisco Nexus 9000 ACI, and VMware vSphere Design Guide:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi55u2_n9k_aci_aff8040_design.html

What's New?

The following design elements distinguish this version of FlexPod from previous non-ACI FlexPod models:

- Validation of the Cisco ACI with a NetApp All-Flash FAS storage array
- Support for the Cisco UCS 2.2 release and Cisco UCS B200-M4 servers
- Support for the latest release of NetApp Data ONTAP® 8.3
- An IP-based storage design supporting both NAS datastores and iSCSI based SAN LUNs
- Support for direct attached Fiber Chanel storage access for boot LUNs
- Application design guidance for multi-tiered applications using Cisco ACI application profiles and policies

Solution Design

Architecture

FlexPod is a defined set of hardware and software that serves as an integrated foundation for both virtualized and non-virtualized solutions. VMware vSphere® built on FlexPod includes NetApp storage, NetApp Data ONTAP, NetApp All Flash FAS, Cisco Nexus® networking, the Cisco Unified Computing System (Cisco UCS®), and VMware vSphere software in a single package. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. Port density enables the networking components to accommodate multiple configurations of this kind.

One benefit of the FlexPod architecture is the ability to customize or " flex" the environment to suit a customer's requirements. A FlexPod can easily be scaled as requirements and demand change. The unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture.

0 shows the VMware vSphere built on FlexPod components and the network connections for a configuration with IP-based storage. This design uses the Cisco Nexus 9000, Cisco Nexus 2232PP FEX, and Cisco UCS C-Series and B-Series servers and the NetApp AFF family of storage controllers connected in a highly available modular design. This infrastructure is deployed to provide iSCSI-booted hosts with file-level and block-level access to shared storage. The reference architecture reinforces the " wire-once" strategy, because as additional storage is added to the architecture, no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

The ACI switching architecture is laid out in a leaf-and-spine topology where every leaf connects to every spine using 40G Ethernet interface(s). The software controller, APIC, is delivered as an appliance and three or more such appliances form a cluster for high availability and enhanced performance.

Physical Topology

0 illustrates the physical architecture.

Figure 1 FlexPod Design with Cisco ACI and NetApp Data ONTAP

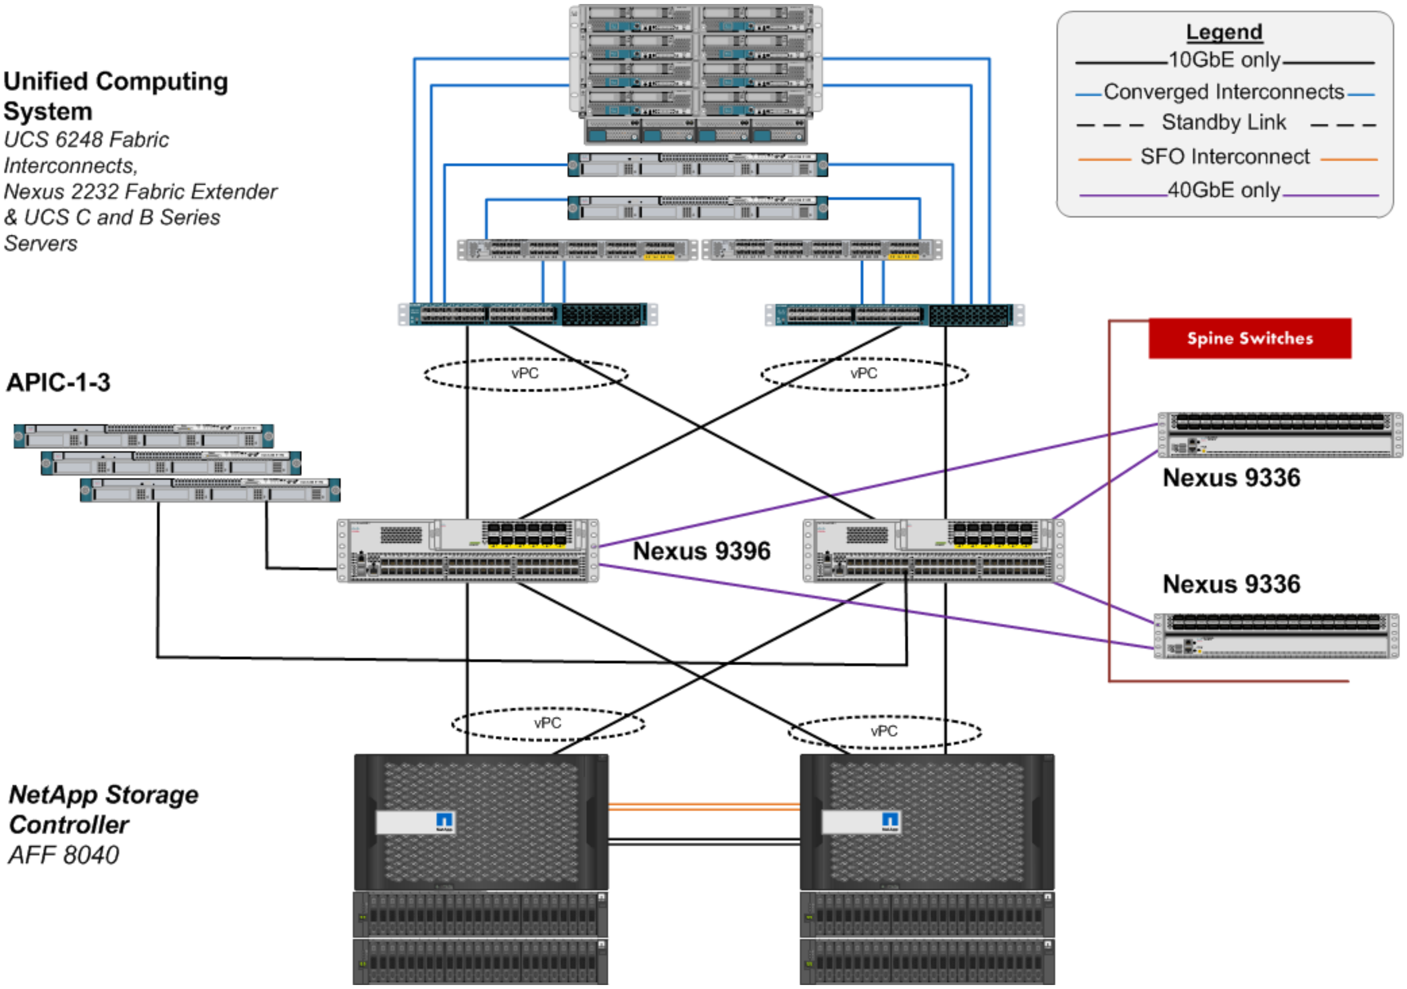

The reference hardware configuration includes:

- Two Cisco Nexus 9396 switches
- Two Cisco Nexus 2232 fabric extenders
- Two Cisco UCS 6248UP fabric interconnects
- One NetApp AFF8040 (HA pair) running clustered Data ONTAP with Disk shelves and Solid State Drives (SSD)

While not included in the FlexPod BOM, Cisco ACI spines and APIC controllers are integral part of Cisco ACI design. The following components were used in the validation efforts:

- Three APIC Controllers
- Two Cisco Nexus 9336 based spines

For server virtualization, the deployment includes VMware vSphere. Although this is the base design, each of the components can be scaled easily to support specific business requirements. For example, more (or different) servers or even blade chassis can be deployed to increase compute capacity, additional disk shelves can be deployed to improve I/O capability and throughput, and special hardware or software features can be added to introduce new features.   This document guides you through the low-level steps for deploying the base architecture, as shown in 0. These procedures cover everything from physical cabling to network, compute and storage device configurations.

Deployment Hardware and Software

Software Revisions

Table 1 lists the software revisions for this solution.

Table 1 Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M4, UCS C-220 M4	2.2(3d)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, UCS VIC 1240 and UCS VIC 1340
	Cisco eNIC	2.1.2.62	
	Cisco fNIC	1.6.0.12b	
Network	Cisco APIC	1.0(4h)*	
	Cisco Nexus 9000 iNX-OS	11.0(4h)*	
Storage	NetApp AFF 8040	Data ONTAP 8.3	
Software	VMware vSphere ESXi	5.5u2	
	VMware vCenter	5.5u2	
	OnCommand Unified Manager for clustered Data ONTAP	6.2	
	NetApp Virtual Storage Console (VSC)	6.0	
	OnCommand Performance Manager	1.1	

* Customers should always use the latest ACI software after consulting with their account team. The APIC screen captures in this Deployment Guide were captured in an earlier version and might be slightly different.

Configuration Guidelines

This document provides details for configuring a fully redundant, highly available configuration for a FlexPod unit with clustered Data ONTAP storage. Therefore, reference is made to which component is being configured with each step, either 01 or 02 or A and B. For example, node01 and node02 are used to identify the two NetApp storage controllers that are provisioned with this document, and Cisco Nexus A or Cisco Nexus B identifies the pair of Cisco Nexus switches that are configured. The Cisco UCS fabric interconnects are similarly configured. Additionally, this document details the steps for provisioning multiple Cisco UCS hosts, and these are identified sequentially: VM-Host-Infra-01, VM-Host-Infra-02, and so on. Finally, to indicate that you should include information pertinent to your environment in a given step, <text> appears as part of the command structure. See the following example for the `network port vlan create` command:

Usage:

```
network port vlan create ?
 [-node] <nodename>                   Node
 { [-vlan-name] {<netport>|<ifgrp>}   VLAN Name
 |  -port {<netport>|<ifgrp>}         Associated Network Port
 [-vlan-id] <integer> }               Network Switch VLAN Identifier
```

Example:

```
network port vlan -node <node01> -vlan-name i0a-<vlan id>
```

This document is intended to enable you to fully configure the customer environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and VLAN schemes, as well as to record appropriate MAC addresses. Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this guide. Table 2 describe the VLANs necessary for deployment as outlined in this guide.

Table 2   Necessary VLANs

VLAN Name	VLAN Purpose	ID Used in Validating This Document
Out of band Mgmt	VLAN for out-of-band management interfaces	3177
Native	VLAN to which untagged frames are assigned	2
NFS LIF	VLAN for NFS LIF (NetApp) traffic	3170
NFS VMK	VLAN for NFS VMkernel (Infrastructure ESXi hosts) traffic	3270
iSCSI-A LIF	VLAN for Fabric A iSCSI LIF	901
iSCSI-B LIF	VLAN for Fabric B iSCSI LIF	902
iSCSI-A-VMK	VLAN for iSCSI traffic for fabric A	911
iSCSI-B-VMK	VLAN for iSCSI traffic for fabric B	912
Tenant Traffic	VLAN Range defined for ACI	1101-1200

Table 3 lists the virtual machines (VMs) necessary for deployment as outlined in this document.

Table 3   Virtual Machines

Virtual Machine Description	Host Name
Active Directory	
vCenter SQL Server database	
vCenter Server	
NetApp Virtual Storage Console (VSC)	
NetApp OnCommand Unified Manager	

Virtual Machine Description	Host Name
OnCommand Performance Manager	

Table 4 lists the configuration variables that are used throughout this document. This table can be completed based on the specific site variables and used in implementing the document configuration steps.

Table 4   Configuration Variables

Variable	Value
<<var_node01_mgmt_ip>>	Out-of-band management IP for cluster node 01
<<var_node01_mgmt_mask>>	Out-of-band management network netmask
<<var_node01_mgmt_gateway>>	Out-of-band management network default gateway
<<var_url_boot_software>>	Data ONTAP 8.3 URL; format: http://
<<var_node02_mgmt_ip>>	Out-of-band management IP for cluster node 02
<<var_node02_mgmt_mask>>	Out-of-band management network netmask
<<var_node02_mgmt_gateway>>	Out-of-band management network default gateway
<<var_clustername>>	Storage cluster host name
<<var_cluster_base_license_key>>	Cluster base license key
<<var_nfs_license>>	NFS license key
<<var_iscsi_license>>	iSCSI license key
<<var_password>>	Global default administrative password
<<var_clustermgmt_ip>>	In-band management IP for the storage cluster
<<var_clustermgmt_mask>>	Out-of-band management network netmask
<<var_clustermgmt_gateway>>	Out-of-band management network default gateway
<<var_dns_domain_name>>	DNS domain name
<<var_nameserver_ip>>	DNS server IP(s)
<<var_node_location>>	Node location string for each node
<<var_node01_sp_ip>>	Out-of-band cluster node 01 service processor management IP
<<var_node01_sp_mask>>	Out-of-band management network netmask
<<var_node01_sp_gateway>	Out-of-band management network default gateway
<<var_node02_sp_ip>>	Out-of-band cluster node 02 device processor management IP

<<var_node02_sp_mask>>	Out-of-band management network netmask
<<var_node02_sp_gateway>	Out-of-band management network default gateway
<<var_node01>>	Cluster node 01 hostname
<<var_node02>>	Cluster node 02 hostname
<<var_num_disks>>	Number of disks to assign to each storage controller
<<var_nfs_vlan_id>>	Infrastructure NFS VLAN ID for LIF
<<var_nfs_vlan_tenant>>	Tenant NFS VLAN ID for LIF (only required when deploying a tenant)
<<var_iscsi_vlan_A_id>>	Infrastructure iSCSI-A VLAN ID for LIF
<<var_iscsi_vlan_B_id>>	Infrastructure iSCSI-B VLAN ID for LIF
<<var_iscsi_vlan_A_tenant>>	Tenant iSCSI-A VLAN ID for LIF (optional)
<<var_iscsi_vlan_B_tenant>>	Tenant iSCSI-B VLAN ID for LIF (optional)
<<var_nfs_vlan_vmk>>	Infrastructure NFS VLAN ID for VMkernel Port
<<var_iscsi_vlan_A_vmk>>	Infrastructure iSCSI-A VLAN ID for VMkernel Port
<<var_iscsi_vlan_B_vmk>>	Infrastructure iSCSI-B VLAN ID for VMkernel Port
<<var_ib_mgmt_vlan_id>>	In-band management network VLAN ID
<<var_oob_mgmt_vlan_id>>	Out-of-band management network VLAN ID
<<var_timezone>>	FlexPod time zone (for example, America/New_York)
<<var_global_ntp_server_ip>>	NTP server IP address
<<var_snmp_contact>>	Administrator e-mail address
<<var_snmp_location>>	Cluster location string
<<var_oncommand_server_fqdn>>	VSC or OnCommand virtual machine fully qualified domain name (FQDN)
<<var_snmp_community>>	Storage cluster SNMP v1/v2 community name
<<var_mailhost>>	Mail server host name
<<var_storage_admin_email>>	Administrator e-mail address
<<var_esxi_host1_nfs_ip>>	NFS VLAN IP address for VMware ESXi host 1
<<var_esxi_host2_nfs_ip>>	NFS VLAN IP address for VMware ESXi host 2
<<var_node01_nfs_lif_infra_swap_ip>>	IP address of Infra Swap
<<var_node01_nfs_lif_infra_swap_mask>>	Subnet Mask of Infra Swap

<<var_node02_nfs_lif_infra_datastore_1_ip>>	IP address of Datastore 1
<<var_node02_nfs_lif_infra_datastore_1_mask>>	Subnet mask of Datastore 1
<<var_vserver_mgmt_ip>>	Management IP address for Vserver
<<var_vserver_mgmt_mask>>	Subnet mask for Vserver
<<var_routing_group>>	Routing group for Vserver
<<var_vserver_mgmt_gateway>>	Default Gateway for Vserver
<<var_vsadmin_password>>	Password for VS admin account
<<var_ucs_clustername>>	Cisco UCS Manager cluster host name
<<var_ucsa_mgmt_ip>>	Cisco UCS fabric interconnect (FI) A out-of-band management IP address
<<var_ucsa_mgmt_mask>>	Out-of-band management network netmask
<<var_ucsa_mgmt_gateway>>	Out-of-band management network default gateway
<<var_ucsb_mgmt_ip>>	Cisco UCS FI B out-of-band management IP address
<<var_vm_host_infra_01_iqn>>	IQN of Infra 01
<<var_vm_host_infra_02_iqn>>	IQN of Infra 02
<<var_vm_host_infra_01_ip>>	VMware ESXi host 01 out-of-band management IP
<<var_vm_host_infra_02_ip>>	VMware ESXi host 02 out-of-band management IP
<<var_nfs_vlan_ip_host_01>>	ESXi host 1, NFS VLAN IP
<<var_nfs_vlan_ip_mask_host_01>>	ESXi host1, NFS VLAN subnet mask
<<var_nfs_vlan_ip_host_02>>	ESXi host 2, NFS VLAN IP
<<var_nfs_vlan_ip_mask_host_02>>	ESXi host2, NFS VLAN subnet mask
<<var_vcenter_server_ip>>	IP address of the vCenter Server
<<var_svm_mgmt_vlan_id>>	Infrastructure Vserver management VLAN ID
<<var_svm_mgmt_vlan_tenant>>	Tenant Vserver management VLAN ID
<<var_nfs_subnet_address>>	NFS subnet address
<<var_node02_nfs_lif_tenant_datastore_1_ip>>	Tenant Datastore 1 IP address
<<var_node02_nfs_lif_tenant_datastore_1_mask>>	Tenant Datastore 1 Subnet mask
<<var_node01_iscsi_lif01a_ip>>	iSCSI LIF 01a IP address
<<var_node01_iscsi_lif01a_mask>>	iSCSI LIF 01a subnet mask
<<var_node01_iscsi_lif01b_ip>>	iSCSI LIF 01b IP address

<<var_node01_iscsi_lif01b_mask>>	iSCSI LIF 01b subnet mask
<<var_node01_iscsi_lif02a_ip>>	iSCSI LIF 02a IP address
<<var_node01_iscsi_lif02a_mask>>	iSCSI LIF 02a subnet mask
<<var_node01_iscsi_lif02b_ip>>	iSCSI LIF 02b IP address
<<var_node01_iscsi_lif02b_mask>>	iSCSI LIF 02b subnet mask
<<var_node01_iscsi_tenant_lif01a_ip>>	Tenant iSCSI LIF 01a IP address
<<var_node01_iscsi_tenant_lif01a_mask>>	Tenant iSCSI LIF 01a subnet mask
<<var_node01_iscsi_tenant_lif01b_ip>>	Tenant iSCSI LIF 01b IP address
<<var_node01_iscsi_tenant_lif01b_mask>>	Tenant iSCSI LIF 01b subnet mask
<<var_node01_iscsi_tenant_lif02a_ip>>	Tenant iSCSI LIF 02a IP address
<<var_node01_iscsi_tenant_lif02a_mask>>	Tenant iSCSI LIF 02a subnet mask
<<var_node01_iscsi_tenant_lif02b_ip>>	Tenant iSCSI LIF 02b IP address
<<var_node01_iscsi_tenant_lif02b_mask>>	Tenant iSCSI LIF 02b subnet mask
<<var_vserver_mgmt_ip>>	Management IP address for Infrastructure Vserver
<<var_vserver_mgmt_mask>>	Management subnet mask for Infrastructure Vserver
<<var_vserver_tenant_mgmt_ip>>	Management IP address for Tenant Vserver
<<var_vserver_tenant_mgmt_mask>>	Management subnet mask for Tenant Vserver
<<var_vserver_mgmt_gateway>>	Management Gateway for Infrastructure Vserver
<<var_vserver_tenant_mgmt_gateway>>	Management Gateway for Tenant Vserver
<<var_oncommand_server_ip>>	IP address of the OnCommand Unified Manager
<<var_rule_index>>	Rule index number
<<var_vm_host_infra_01_A_wwpn>>	WWPN of Infra Datastore 01 A
<<var_vm_host_infra_01_B_wwpn>>	WWPN of Infra Datastore 01 B
<<var_vm_host_infra_02_A_wwpn>>	WWPN of Infra Datastore 02 A
<<var_vm_host_infra_02_B_wwpn>>	WWPN of Infra Datastore 02 B
<<var_server_nfs_vlan_id>>	NFS VLAN ID
<<var_nfs_lif02_ip>>	NFS LIF 02 IP Address
<<var_nfs_lif01_ip>>	NFS LIF 01 IP Address

Physical Infrastructure

FlexPod Cabling

The information in this section is provided as a reference for cabling the physical equipment in a FlexPod environment. To simplify cabling requirements, the tables include both local and remote device and port locations.

The tables in this section contain the details for the prescribed and supported configuration of the NetApp AFF8040 running clustered Data ONTAP 8.3. For any modifications of this prescribed architecture, consult the NetApp Interoperability Matrix Tool (IMT).

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site. These interfaces will be used in various configuration steps

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

0 shows a cabling diagram for a FlexPod configuration using the Cisco Nexus 9000 and NetApp storage systems with clustered Data ONTAP. The NetApp storage controller and disk shelves should be connected according to best practices for the specific storage controller and disk shelves. For disk shelf cabling, refer to the Universal SAS and ACP Cabling Guide: https://library.netapp.com/ecm/ecm_get_file/ECMM1280392.

Figure 2   FlexPod Cabling Diagram

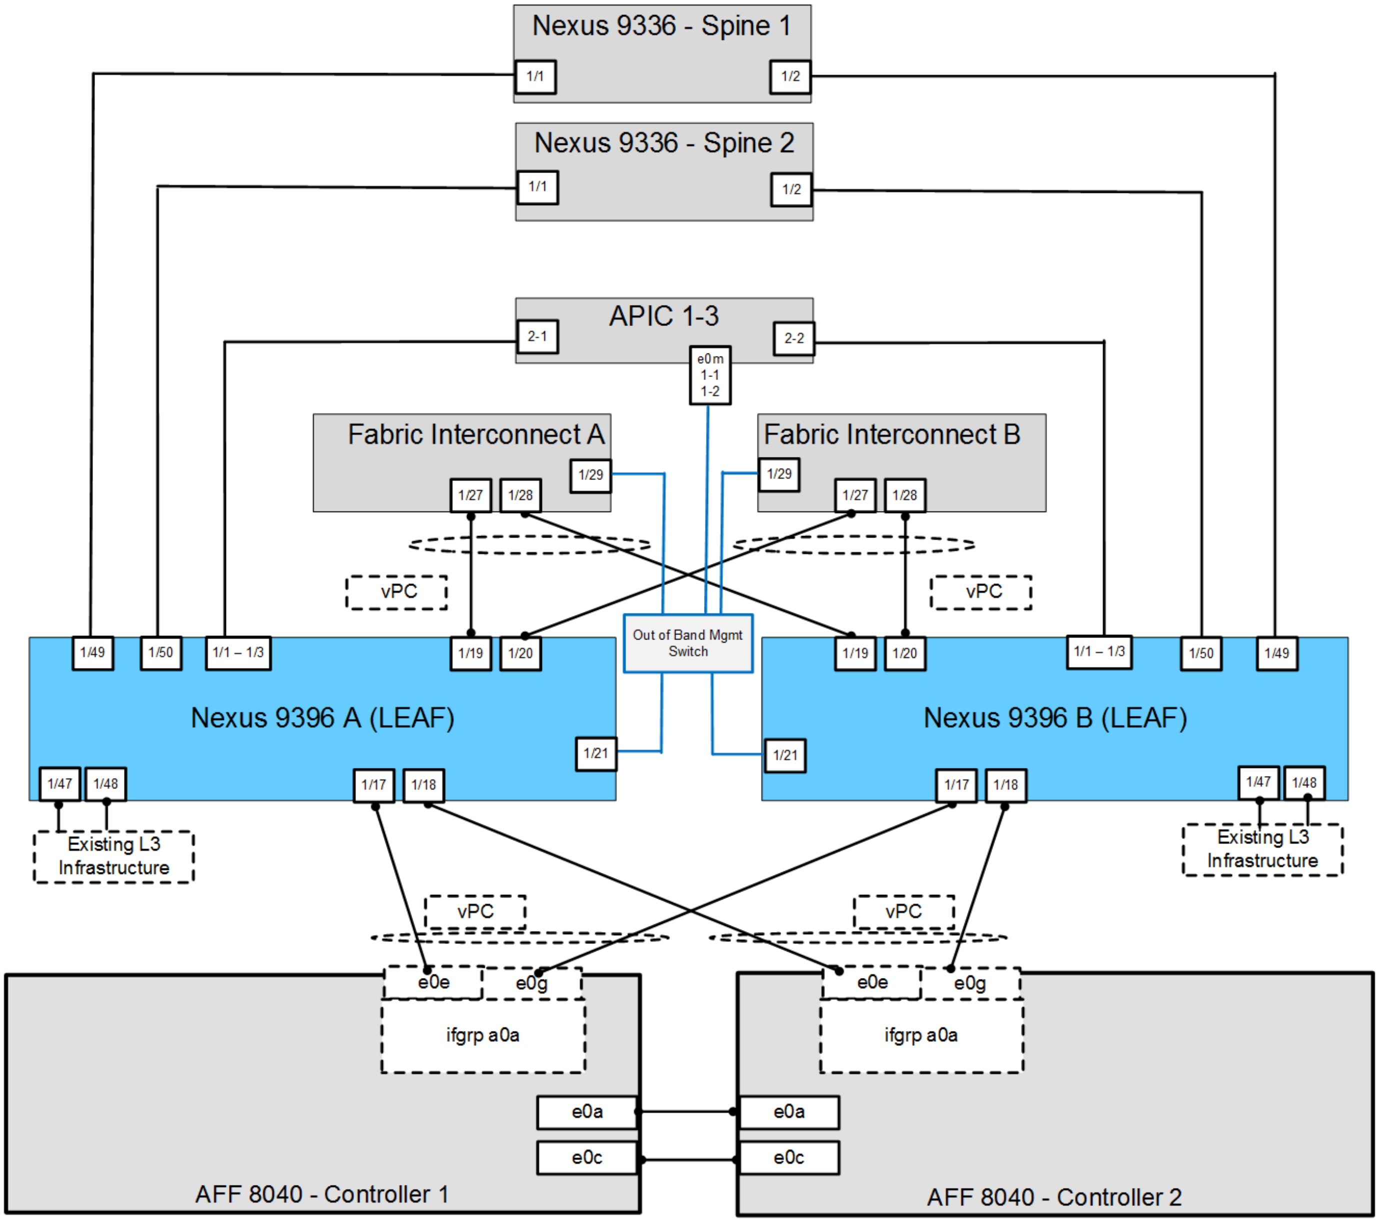

Table 5 through Table 14 provide the details of all the connections in use.

Table 5   Cisco Nexus 9396-A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396 A	Eth1/1	10GbE	APIC 1	Eth 2-1
	Eth1/2	10GbE	APIC 2	Eth 2-1
	Eth1/3	10GbE	APIC 3	Eth 2-1
	Eth1/17	10GbE	NetApp controller 1	e0e

Local Device	Local Port	Connection	Remote Device	Remote Port
	Eth1/18	10GbE	NetApp controller 2	e0e
	Eth1/19	10GbE	Cisco UCS fabric interconnect A	Eth1/31
	Eth1/20	10GbE	Cisco UCS fabric interconnect B	Eth1/31
	Eth1/21	GbE	Common Services Mgmt. Switch	Any
	Eth1/49	40GbE	Cisco 9336 Spine 1	Eth1/1
	Eth1/50	40GbE	Cisco 9336 Spine 2	Eth1/1
	MGMT0	GbE	GbE management switch	Any

Note: For devices requiring GbE connectivity, use the GbE Copper SFP+s (GLC-T=).

Table 6 Cisco Nexus 9396-B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 9396 A	Eth1/1	10GbE	APIC 1	Eth 2-2
	Eth1/2	10GbE	APIC 2	Eth 2-2
	Eth1/3	10GbE	APIC 3	Eth 2-2
	Eth1/17	10GbE	NetApp controller 1	e0g
	Eth1/18	10GbE	NetApp controller 2	e0g
	Eth1/19	10GbE	Cisco UCS fabric interconnect A	Eth1/32
	Eth1/20	10GbE	Cisco UCS fabric interconnect B	Eth1/32
	Eth1/21	GbE	Common Services Mgmt. Switch	Any
	Eth1/49	40GbE	Cisco 9336 Spine 1	Eth1/2
	Eth1/50	40GbE	Cisco 9336 Spine 2	Eth1/2
	MGMT0	GbE	GbE management switch	Any

Table 7 NetApp Controller-1 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 1	e0M	100MbE	100MbE management switch	Any
	e0a	GbE	GbE management switch	Any

Local Device	Local Port	Connection	Remote Device	Remote Port
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 2	e0a
	e0c	10GbE	NetApp Controller 2	e0c
	e0e	10GbE	Cisco Nexus 9000 A	Eth 1/17
	e0g	10GbE	Cisco Nexus 9000 B	Eth 1/17

Note: When the term e0M is used, the physical Ethernet port to which the table is referring is the port indicated by a wrench icon on the rear of the chassis.

Table 8 NetApp controller 2 Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
NetApp controller 2	e0M	100MbE	100MbE management switch	Any
	e0a	GbE	GbE management switch	Any
	e0P	GbE	SAS shelves	ACP port
	e0a	10GbE	NetApp Controller 1	e0a
	e0c	10GbE	NetApp Controller 1	e0c
	e0e	10GbE	Cisco Nexus 9000 A	Eth 1/18
	e0g	10GbE	Cisco Nexus 9000 B	Eth 1/18

Table 9 UCS Fabric Interconnect A Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/1	10GbE	Cisco UCS Chassis FEX A	IOM 1/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX A	IOM 1/2
	Eth1/27	10GbE	Cisco Nexus 9000 A	Eth 1/19
	Eth1/28	10GbE	Cisco Nexus 9000 B	Eth 1/19
	Eth1/29	10GbE	Management Switch	Any
	Eth1/31	10GbE	Cisco Nexus 2232PP FEX A	Uplink 1
	Eth1/32	10GbE	Cisco Nexus 2232PP FEX A	Uplink 2

Local Device	Local Port	Connection	Remote Device	Remote Port
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 10 UCS Fabric Interconnect B Cabling Information

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS fabric interconnect A	Eth1/1	10GbE	Cisco UCS Chassis FEX B	IOM 1/1
	Eth1/2	10GbE	Cisco UCS Chassis FEX B	IOM 1/2
	Eth1/27	10GbE	Cisco Nexus 9000 A	Eth 1/20
	Eth1/28	10GbE	Cisco Nexus 9000 B	Eth 1/20
	Eth1/29	10GbE	Management Switch	Any
	Eth1/31	10GbE	Cisco Nexus 2232PP FEX B	Uplink 1
	Eth1/32	10GbE	Cisco Nexus 2232PP FEX B	Uplink 2
	MGMT0	GbE	GbE management switch	Any
	L1	GbE	Cisco UCS fabric interconnect B	L1
	L2	GbE	Cisco UCS fabric interconnect B	L2

Table 11 Cisco Nexus 2232 FEX A—Single Wire Management

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 2232PP FEX A	Port 1	10GbE	Cisco UCS C-Series 1	Port 0
	Port 2	10GbE	Cisco UCS C-Series 2	Port 0

Table 12 Cisco Nexus 2232 FEX B—Single Wire Management

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

Table 13 Cisco UCS C-Series 1

Local Device	Local Port	Connection	Remote Device	Remote Port

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 1	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 1
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 1

Table 14 Cisco UCS C-Series 2

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS C-Series 2	Port 0	10GbE	Cisco Nexus 2232PP FEX A	Port 2
	Port 1	10GbE	Cisco Nexus 2232PP FEX B	Port 2

Storage Configuration

Controller AFF80XX Series

Refer to the Site Requirements Guide for planning the physical location of the storage systems. From the downloaded guide, refer to the following sections:

- Site Preparation
- System Connectivity Requirements
- Circuit Breaker, Power Outlet Balancing, System Cabinet Power Cord Plugs, and Console Pinout Requirements
- 80xx Series Systems

NetApp Hardware Universe

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

1. Confirm that the hardware and software components are supported with the version of Data ONTAP that you plan to install by using the NetApp Hardware Universe (HWU) application at the NetApp Support site.

2. Access the HWU application to view the System Configuration guides. Click the Controllers tab to view the compatibility between Data ONTAP software versions and NetApp storage appliances with the desired specifications.

3. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controllers

Follow the physical installation procedures for the controllers which can be found in the AFF8000 Series product documentation at the NetApp Support site.

Disk Shelves

NetApp storage systems support a wide variety of disk shelves and disk drives. The complete list of disk shelves that are supported with AFF 80xx is available at the NetApp Support site.

When using SAS disk shelves with NetApp storage controllers, refer to the SAS Disk Shelves Universal SAS and ACP Cabling Guide for proper cabling guidelines.

Clustered Data ONTAP 8.3

Complete the Configuration Worksheet

Before running the setup script, complete the cluster setup worksheet from the Clustered Data ONTAP 8.3 Software Setup Guide.

	How to Access the Configuration Worksheet Configuration Guide	Comments
Configuration Worksheet	Clustered Data ONTAP 8.3 Software Setup Guide.	Requires access to the NetApp Support site.

Configure Clustered Data ONTAP Nodes

Before running the setup script, review the configuration worksheets in the Clustered Data ONTAP 8.3 Software Setup Guide to learn about the information required to configure clustered Data ONTAP. Table 15 lists the information that you will need to configure two clustered Data ONTAP nodes. You should customize the cluster detail values with the information that is applicable to your deployment.

Table 15 Clustered Data ONTAP Software Installation Prerequisites

Cluster Detail	Cluster Detail Value
Cluster Node01 IP address	<<var_node01_mgmt_ip>>
Cluster Node01 netmask	<<var_node01_mgmt_mask>>
Cluster Node01 gateway	<<var_node01_mgmt_gateway>>
Cluster Node02 IP address	<<var_node02_mgmt_ip>>
Cluster Node02 netmask	<<var_node02_mgmt_mask>>
Cluster Node02 gateway	<<var_node02_mgmt_gateway>>
Data ONTAP 8.3 URL	<<var_url_boot_software>>

Configure Node 01

To configure node 01, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

Note: If Data ONTAP 8.3 is not the version of software being booted, continue the following steps to install new software. If Data ONTAP 8.3 is the version being booted, select option 8 and y (Yes) to reboot the node, then continue with step 14.

4. To install new software, select option 7.

```
7
```

5. Enter `y` (Yes) to perform an upgrade.

```
y
```

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

```
e0M
```

7. Enter `y` (Yes) to reboot now.

```
y
```

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

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

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

Note: This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

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

```
y
```

12. Enter `y` (Yes) to reboot the node.

```
y
```

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

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

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

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

```
4
```

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

```
y
```

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

```
Y
```

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

Configure Node 02

To configure node 02, complete the following steps:

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

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

2. Allow the system to boot up.

```
autoboot
```

3. Press Ctrl-C when prompted.

```
Ctrl-C
```

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

4. To install new software, select option 7.

```
7
```

5. Enter `y` (Yes) to perform a non-disruptive upgrade.

```
y
```

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

```
e0M
```

7. Enter `y` (Yes) to reboot now.

```
y
```

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

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

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

This web server must be pingable.

```
<<var_url_boot_software>>
```

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

```
Enter
```

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

```
y
```

12. Enter `y` (Yes) to reboot the node.

```
y
```

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

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

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

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

```
4
```

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

```
y
```

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

```
y
```

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

Set Up Node

From a console port program attached to the storage controller A (node 01) console port, run the node setup script. This script appears when Data ONTAP 8.3 boots on the node for the first time. To set up the node, complete the following steps:

1. Follow the prompts to set up node 01.

```
Welcome to node setup.


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

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

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

Type yes to confirm and continue {yes}: yes

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

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

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

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

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

2. Press Enter and log in to the node with the admin user id and no password.

3. At the node command prompt, enter the following commands:

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


::> system node reboot


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

4. After reboot, set up the node with the preassigned values.

```
Welcome to node setup.


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


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



Enter the node management interface port [e0M]: Enter
Enter the node management interface IP address [<<var_node01_mgmt_ip>>]: Enter
Enter the node management interface netmask [<<var_node01_mgmt_mask>>]: Enter
Enter the node management interface default gateway
[<<var_node01_mgmt_gateway>>]: Enter


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


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


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


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

5. Log in to the node as the admin user and no password.

6. Repeat this procedure for storage cluster node 02.

Create Cluster on Node 01

In clustered Data ONTAP, the first node in the cluster performs the cluster create operation. All other nodes perform a cluster join operation. The first node in the cluster is considered node 01.

Table 16 Cluster **`create`** in Clustered Data ONTAP Prerequisites

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

7. Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,

"back" - if you want to change previously answered questions, and

"exit" or "quit" - if you want to quit the cluster setup wizard.

Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".

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

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

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

To create a new cluster, complete the following steps:

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

```
create
```

9. Enter `no` for the single-node cluster option.

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

10. Enter `no` for cluster network using network switches.

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

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

```
Existing cluster interface configuration found:


Port    MTU     IP              Netmask

e0a     9000    169.254.118.102 255.255.0.0

e0b     9000    169.254.152.110 255.255.0.0

e0c     9000    169.254.191.92  255.255.0.0

e0d     9000    169.254.233.52  255.255.0.0


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


System Defaults:

Private cluster network ports [e0a,e0c].

Cluster port MTU values will be set to 9000.

Cluster interface IP addresses will be automatically generated.


Do you want to use these defaults? {yes, no} [yes]: yes
```

12. The steps to create a cluster are displayed.

```
Enter the cluster administrators (username “admin) password: <<var_password>>
Retype the password: <<var_password>>


It can take several minutes to create cluster interfaces...




Step 1 of 5: Create a Cluster

You can type "back", "exit", or "help" at any question.

Enter the cluster name: <<var_clustername>>
Enter the cluster base license key: <<var_cluster_base_license_key>>
Creating cluster <<var_clustername>>

Enter an additional license key []:<<var_iscsi_license>>
```

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

Note: For this validated architecture, NetApp recommends installing license keys for NetApp SnapRestore®, NetApp FlexClone®, and NetApp SnapManager® Suite. In addition, install all required storage protocol licenses and all licenses that came with the AFF bundle. After you finish entering the license keys, press Enter.

```
Enter the cluster management interface port [e0e]: e0i
Enter the cluster management interface IP address: <<var_clustermgmt_ip>>
Enter the cluster management interface netmask: <<var_clustermgmt_mask>>
Enter the cluster management interface default gateway:
<<var_clustermgmt_gateway>>
```

13. Enter the DNS domain name.

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

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

14. Set up the node.

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


The node management interface has been modified to use port e0M with IP address
<<var_node01_mgmt_ip>>.

This system will send event messages and weekly reports to NetApp Technical
Support.

To disable this feature, enter "autosupport modify -support disable" within 24
hours.

Enabling AutoSupport can significantly speed problem determination and resolution
should a problem occur on your system.

For further information on AutoSupport, please see:
http://support.netapp.com/autosupport/

Press enter to continue: Enter
Cluster "<<var_clustername>>" has been created.

To complete cluster setup, you must join each additional node to the cluster

by running "cluster setup" on each node.
```

```
Once all nodes have been joined to the cluster, see the Clustered Data ONTAP
Software Setup Guide for information about additional system configuration
tasks.  You can find the Software Setup Guide on the NetApp Support Site.

To complete system configuration, you can use either OnCommand System Manager
or the Data ONTAP command-line interface.

To access OnCommand System Manager, point your web browser to the cluster
management IP address (<<var_clustermgmt_ip>>).
To access the command-line interface, connect to the cluster management
IP address (for example, ssh admin@<<var_clustermgmt_ip>>).



<<var_clustername>>::>
```

Note: The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document it is assumed to be on the same subnet.

Join Node 02 to Cluster

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

Table 17 Cluster **`join`** in Clustered Data ONTAP Prerequisites

Cluster Detail	Cluster Detail Value
Cluster name	`<<var_clustername>>`
Cluster management IP address	`<<var_clustermgmt_ip>>`
Cluster node02 IP address	`<<var_node02_mgmt_ip>>`
Cluster node02 netmask	`<<var_node02_mgmt_mask>>`
Cluster node02 gateway	`<<var_node02_mgmt_gateway>>`

To join node 02 to the existing cluster, complete the following steps:

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

```
admin
```

2. Run the `cluster setup` command to start the Cluster Setup wizard.

```
cluster setup

This node's storage failover partner is already a member of a cluster.
Storage failover partners must be members of the same cluster.
The cluster setup wizard will default to the cluster join dialog.
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{join}:
```

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

3. Run the following command to join a cluster:

```
join
```

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

```
Existing cluster interface configuration found:

Port    MTU     IP              Netmask
e0a     9000    169.254.1.79    255.255.0.0
e0b     9000    169.254.54.223  255.255.0.0
e0c     9000    169.254.100.157 255.255.0.0
e0d     9000    169.254.138.142 255.255.0.0

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

System Defaults:
```

```
Private cluster network ports [e0a,e0c].

Cluster port MTU values will be set to 9000.

Cluster interface IP addresses will be automatically generated.


Do you want to use these defaults? {yes, no} [yes]:Enter
It can take several minutes to create cluster interfaces...
```

5. The steps to join a cluster are displayed.

```
Step 1 of 3: Join an Existing Cluster

You can type "back", "exit", or "help" at any question.


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

Starting cluster support services ..


This node has joined the cluster <<var_clustername>>.




Step 2 of 3: Configure Storage Failover (SFO)

You can type "back", "exit", or "help" at any question.




SFO is enabled.




Step 3 of 3: Set Up the Node

You can type "back", "exit", or "help" at any question.


Notice: HA is configured in management.
```

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

6. Set up the node.

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

```
The node management interface has been modified to use port e0M with IP address
<<var_node02_mgmt_ip>>.

This system will send event messages and weekly reports to NetApp Technical
Support.

To disable this feature, enter "autosupport modify -support disable" within 24
hours.

Enabling AutoSupport can significantly speed problem determination and resolution
should a problem occur on your system.

For further information on AutoSupport, please see:
http://support.netapp.com/autosupport/

Press enter to continue: Enter

This node has been joined to cluster "<<var_clustername>>".

To complete cluster setup, you must join each additional node to the cluster

by running "cluster setup" on each node.


Once all nodes have been joined to the cluster, see the Clustered Data ONTAP

Software Setup Guide for information about additional system configuration

tasks.  You can find the Software Setup Guide on the NetApp Support Site.


To complete system configuration, you can use either OnCommand System Manager

or the Data ONTAP command-line interface.


To access OnCommand System Manager, point your web browser to the cluster

management IP address (<<var_clustermgmt_ip>>).

To access the command-line interface, connect to the cluster management

IP address (for example, ssh admin@<<var_clustermgmt_ip>>).
```

Note: The node management interface can be on the same subnet as the cluster management interface, or it can be on a different subnet. In this document it is assumed to be on the same subnet.

Log In to Cluster

To log in to the cluster, complete the following steps:

1. Open an SSH connection to either the cluster IP or host name.

2. Log in to the admin user with the password you provided earlier.

Zero All Spare Disks

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

1. Run the following command:

```
disk zerospares
```

Note: Disk autoassign should have assigned half of the connected disks to each node in the HA pair. If a different disk assignment is required, disk autoassignment must be disabled on both nodes in the HA pair by running the disk option modify command. Spare disks can then be moved from one node to another by running the `disk removeowner` and `disk assign` commands.

Set Onboard UTA2 Ports Personality

To set the personality of the onboard Unified Target Adapter 2 (UTA2), complete the following steps:

1. Verify the Current Mode and Current Type of the ports by running the `ucadmin` show command.

```
ucadmin show
                       Current  Current    Pending  Pending    Admin
Node          Adapter  Mode     Type       Mode     Type       Status
------------  -------  -------  ---------  -------  ---------  -----------
<<var_node01>>
              0e       cna      target     -        -          online
<<var_node01>>
              0f       cna      target     -        -          online
<<var_node01>>
              0g       cna      target     -        -          online
<<var_node01>>
              0h       cna      target     -        -          online
<<var_node02>>
              0e       cna      target     -        -          online
<<var_node02>>
              0f       cna      target     -        -          online
<<var_node02>>
              0g       cna      target     -        -          online
<<var_node02>>
              0h       cna      target     -        -          online
8 entries were displayed.
```

2. Verify that the Current Mode of all the ports in use is `cna` and the Current Type is set to target. If not, change the port personality by running the following command:

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

Note: The ports must be offline to run this command. To take an adapter offline, run the `fcp adapter modify -node <home node of the port> -adapter <port name> -state down` command. Ports must be converted in pairs, for example, 0c and 0d, after which, a reboot is required, and the ports must be brought back to the up state.

Set Auto-Revert on Cluster Management

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

Note: The storage virtual machine (SVM) is referred to as Vserver (or `vserver`) in the GUI and CLI.

1. Run the following command:

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

Set Up Management Broadcast Domain

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

1. Run the following commands:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_node01>>:e0b,<<var_node01>>:e0d,<<var_node01>>:e0e,<<var_node01>>:e0f,<<var
_node01>>:e0g,<<var_node01>>:e0h,<<var_node01>>:e0j,<<var_node01>>:e0k,<<var_node
01>>:e0l,<<var_node02>>:e0b,<<var_node02>>:e0d,<<var_node02>>:e0e,<<var_node02>>:
e0f,<<var_node02>>:e0g,<<var_node02>>:e0h,<<var_node02>>:e0j,<<var_node02>>:e0k,<
<var_node02>>:e0l
broadcast-domain show
```

Set Up Service Processor Network Interface

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

1. Run the following commands:

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


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

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

Create Aggregates

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

To create new aggregates, complete the following steps:

1. Run the following commands:

```
aggr create -aggregate aggr1_node01 -nodes <<var_node01>> -diskcount
<<var_num_disks>>

aggr create -aggregate aggr1_node02 -nodes <<var_node02>> -diskcount
<<var_num_disks>>
```

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

Note: Start with five disks initially; you can add disks to an aggregate when additional storage is required. In an AFF configuration with a small number of SSDS, it may be desirable to create an aggregate with all but one remaining disk (spare) assigned to the controller.

Note: The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display aggregate creation status. Do not proceed until both aggr1_node1 and aggr1_node2 are online.

2. Disable NetApp Snapshot® copies for the two data aggregates recently created.

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

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

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

```
node run <<var_node01>> snap delete –A –a –f aggr1_node01

node run <<var_node02>> snap delete –A –a –f aggr1_node02
```

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

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

Verify Storage Failover

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

1. Verify the status of storage failover.

```
storage failover show
```

Note: Both the nodes `<<var_node01>>` and `<<var_node02>>` must be capable of performing a takeover. Continue with step 3 if the nodes are capable of performing a takeover.

2. Enable failover on one of the two nodes.

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

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

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

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

```
cluster ha show
```

4. Continue with step 6 if high availability is configured.

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

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

```
cluster ha modify -configured true

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

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

```
storage failover hwassist show

storage failover modify –hwassist-partner-ip <<var_node02_mgmt_ip>> -node
<<var_node01>>

storage failover modify –hwassist-partner-ip <<var_node01_mgmt_ip>> -node
<<var_node02>>
```

Disable Flow Control on UTA2 Ports

NetApp recommends disabling flow control on all of the 10GbE and UTA2 ports that are connected to external devices. To disable flow control, complete the following steps:

1. Run the following commands to configure node 01:

```
network port modify -node <<var_node01>> -port e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0h -
flowcontrol-admin none

Warning: Changing the network port settings will cause a several second
interruption in carrier.

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

2. Run the following commands to configure node 02:

```
network port modify -node <<var_node02>> -port e0a,e0b,e0c,e0d,e0e,e0f,e0g,e0hq -
flowcontrol-admin none

Warning: Changing the network port settings will cause a several second
interruption in carrier.

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

network port show –fields flowcontrol-admin
```

Disable Unused FcoE Ports

Unused data FCoE ports on active interfaces should be disabled. To disable these ports, complete the following step:

1. Run the following commands:

```
fcp adapter modify -node <<var_node01>> -adapter 0e –state down

fcp adapter modify -node <<var_node01>> -adapter 0g –state down

fcp adapter modify -node <<var_node02>> -adapter 0e –state down

fcp adapter modify -node <<var_node02>> -adapter 0g –state down

fcp adapter show –fields state
```

Configure NTP

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

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

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

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

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

```
date <ccyymmddhhmm.ss>
```

Note: The format for the date is `<[Century][Year][Month][Day][Hour][Minute].[Second]>`; for example, `201309081735.17`

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

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

Configure SNMP

To configure SNMP, complete the following steps:

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

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>
```

```
snmp init 1
```

```
options snmp.enable on
```

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

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

Configure SNMPv1 Access

To configure SNMPv1 access, complete the following step:

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

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

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

Create SNMPv3 User

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

1. Create a user called `snmpv3user`.

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

2. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.

3. Run the security `snmpusers` command to view the engine ID.

4. When prompted, enter an eight-character minimum-length password for the authentication protocol.

5. Select `des` as the privacy protocol.

6. When prompted, enter an eight-character minimum-length password for the privacy protocol.

Configure AutoSupport

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

1. Run the following command:

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

Enable Cisco Discovery Protocol

To enable Cisco Discovery Protocol (CDP) on the NetApp storage controllers, complete the following step:

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

2. Run the following command to enable CDP on Data ONTAP:

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

Create Jumbo Frame MTU Broadcast Domains in Clustered Data ONTAP

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

1. Run the following commands to create a broadcast domain on Data ONTAP:

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

Create Interface Groups

To create the LACP interface groups for the 10GbE data interfaces, complete the following step:

1. Run the following commands to create the LACP interface groups:

```
ifgrp create -node <<var_node01>> -ifgrp a0a -distr-func port -mode
multimode_lacp
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node01>> -ifgrp a0a -port e0g

ifgrp create -node <<var_node02>> -ifgrp a0a -distr-func port -mode
multimode_lacp
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0e
ifgrp add-port -node <<var_node02>> -ifgrp a0a -port e0g


ifgrp show
```

Create VLANs

To create VLANs, complete the following steps:

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

```
network port modify –node <<var_node01>> -port a0a –mtu 9000
network port modify –node <<var_node02

>> -port a0a –mtu 9000


network port vlan create –node <<var_node01>> -vlan-name a0a-<<var_nfs_vlan_id>>
network port vlan create –node <<var_node02>> -vlan-name a0a-<<var_nfs_vlan_id>>

broadcast-domain add-ports -broadcast-domain Infra_NFS -ports <<var_node01>>:a0a-
<<var_nfs_vlan_id>>, <<var_node02>>:a0a-<<var_nfs_vlan_id>>
```

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

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

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

Create Storage Virtual Machine

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

Note: The storage virtual machine (SVM) is referred to as a Vserver (or `vserver`) in the GUI and CLI.

1. Run the `vserver create` command.

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

2. Select the SVM data protocols to configure, keeping `nfs` and `iscsi`.

```
vserver remove-protocols –vserver Infra-SVM -protocols fcp,cifs,ndmp
```

3. Add the two data aggregates to the Infra-SVM aggregate list for NetApp Virtual Storage Console (VSC).

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

4. Enable and run the NFS protocol in the Infra-SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plugin.

```
vserver nfs modify –vserver Infra-SVM –vstorage enabled

vserver nfs show
```

Create Load-Sharing Mirrors of SVM Root Volume

To create a load-sharing mirror of an SVM root volume, complete the following steps:

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

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

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

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

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

3. Create the mirroring relationships.

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

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

4. Initialize the mirroring relationship.

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

Create iSCSI Service

To create the iSCSI service, complete the following step:

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

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

Configure HTTPS Access

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

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

```
set -privilege diag

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

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

```
security certificate show
```

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

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

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

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

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

5. To obtain the values for the parameters that would be required in step 6, run the `security certificate show` command.

6. Enable each certificate that was just created by using the `–server-enabled true` and `–client-enabled false` parameters. Use TAB completion to aid in the completion of these commands.

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

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

```
system services web modify -external true -sslv3-enabled true

Warning: Modifying the cluster configuration will cause pending web service
requests to be interrupted as the web servers are restarted.

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

system services firewall policy delete -policy mgmt -service http –vserver
<<var_clustername>>
```

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

9. Change back to the normal admin privilege level and set up the system to allow SVM logs to be available by web.

```
set –privilege admin

vserver services web modify –name spi|ontapi|compat –vserver * -enabled true
```

Configure NFSv3

To configure NFSv3 on the SVM, complete the following steps:

1. Create a new rule for each ESXi host in the default export policy. Assign a rule for each ESXi host created so that each host has its own rule index. For example, the first ESXi host has rule index 1, the second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create –vserver Infra-SVM -policyname default –
ruleindex 1 –protocol nfs -clientmatch <<var_esxi_host1_nfs_ip>> -rorule sys –
rwrule sys -superuser sys –allow-suid false
```

```
vserver export-policy rule create –vserver Infra-SVM -policyname default –
ruleindex 2 –protocol nfs -clientmatch <<var_esxi_host2_nfs_ip>> -rorule sys –
rwrule sys -superuser sys –allow-suid false
vserver export-policy rule show
```

2. Assign the FlexPod export policy to the infrastructure SVM root volume.

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

Create FlexVol Volumes

The following information is required to create a FlexVol volume:

- Volume name
- Volume size
- Aggregate on which the volume exists

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

1. Run the following commands:

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_node02 -size 500GB -state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent-snapshot-space 0


volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_node01 -size
100GB -state online -policy default -junction-path /infra_swap -space-guarantee
none -percent-snapshot-space 0 -snapshot-policy none


volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_node01 -size
100GB -state online -policy default -space-guarantee none -percent-snapshot-space
0


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

Create Boot LUNs

To create two boot LUNs, complete the following step:

1. Run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 15GB
-ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-02 -size 15GB
-ostype vmware -space-reserve disabled
```

Enable Deduplication

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

1. Run the following commands:

```
volume efficiency on -vserver Infra-SVM -volume infra_datastore_1
volume efficiency on -vserver Infra-SVM -volume esxi_boot
```

Create iSCSI LIFs

To create four iSCSI LIFs (two on each node), complete the following step:

1. Run the following commands:

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


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


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


network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data -data-protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iscsi_vlan_B_id>> -address <<var_node02_iscsi_lif01b_ip>> -netmask <<var_node02_iscsi_lif01b_mask>> -status-admin up -failover-policy disabled -firewall-policy data -auto-revert false

network interface show
```

Create NFS LIF

To create an NFS LIF, complete the following step:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif nfs_infra_swap -role data -data-protocol nfs -home-node <<var_node01>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_node01_nfs_lif_infra_swap_ip>> -netmask <<var_node01_nfs_lif_infra_swap_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface create -vserver Infra-SVM -lif nfs_infra_datastore_1 -role data -data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> -address <<var_node02_nfs_lif_infra_datastore_1_ip>> -netmask <<var_node02_nfs_lif_infra_datastore_1_mask>> -status-admin up -failover-policy broadcast-domain-wide -firewall-policy data -auto-revert true

network interface show
```

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

Add Infrastructure SVM Administrator

To add the infrastructure SVM administrator and SVM administration LIF in the out-of-band management network, complete the following steps:

1. Run the following commands:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data -data-protocol
none -home-node <<var_node02>> -home-port  e0i -address <<var_svm_mgmt_ip>> -
netmask <<var_svm_mgmt_mask>> -status-admin up -failover-policy broadcast-domain-
wide -firewall-policy mgmt -auto-revert true
```

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

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

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



network route show
```

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

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


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

Server Configuration

Cisco UCS Base Configuration

Perform Initial Setup of Cisco UCS 6248 Fabric Interconnect for FlexPod Environments

This section provides detailed procedures for configuring the Cisco Unified Computing System (Cisco UCS) for use in a FlexPod environment. The steps are necessary to provision the Cisco UCS C-Series and B-Series servers and should be followed precisely to avoid improper configuration.

Cisco UCS 6248 A

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

1. Connect to the console port on the first Cisco UCS 6248 fabric interconnect.

```
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": <<var_password>>

Enter the same password for "admin": <<var_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: <<var_ucs_clustername>>

Physical switch Mgmt0 IPv4 address: <<var_ucsa_mgmt_ip>>

Physical switch Mgmt0 IPv4 netmask: <<var_ucsa_mgmt_mask>>

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

Cluster IPv4 address: <<var_ucs_cluster_ip>>

Configure DNS Server IPv4 address? (yes/no) [no]: y

DNS IPv4 address: <<var_nameserver_ip>>

Configure the default domain name? y

Default domain name: <<var_dns_domain_name>>

Join centralized management environment (UCS Central)? (yes/no) [n]: Enter
```

Cisco UCS 6248 B

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

1. Connect to the console port on the second Cisco UCS 6248 fabric interconnect.

```
Enter the configuration method: console
Installer has detected the presence of a peer Fabric interconnect. This
Fabric interconnect will be added to the cluster.  Do you want to continue
{y|n}? y
Enter the admin password for the peer fabric interconnect: <<var_password>>
Physical switch Mgmt0 IPv4 address: <<var_ucsb_mgmt_ip>>
Apply and save the configuration (select 'no' if you want to re-enter)?
(yes/no): y
```

2. Wait for the login prompt to make sure that the configuration has been saved.

Cisco UCS Setup

Log in to Cisco UCS Manager

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

1. Open a web browser and navigate to the Cisco UCS 6248 fabric interconnect cluster address.
2. Click the Launch UCS Manager link to download the Cisco UCS Manager software.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. Click Login to log in to Cisco UCS Manager.

Upgrade Cisco UCS Manager Software to Version 2.2(3d)

This document assumes the use of Cisco UCS 2.2(3d). To upgrade the Cisco UCS Manager software and the Cisco UCS 6248 Fabric Interconnect software to version 2.2(3d), refer to Cisco UCS Manager Install and Upgrade Guides.

Add Block of IP Addresses for KVM Access

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

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

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select Pools > root > IP Pools > IP Pool ext-mgmt.
3. In the Actions pane, select Create Block of IP Addresses.

4. Enter the starting IP address of the block and the number of IP addresses required, and the subnet and gateway information.

5. Click OK to create the IP block.

6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

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

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

2. Select All > Timezone Management.

3. In the Properties pane, select the appropriate time zone in the Timezone menu.

4. Click Save Changes, and then click OK.

5. Click Add NTP Server.

6. Enter `<<var_global_ntp_server_ip>>` and click OK.

7. Click OK.

Edit Chassis Discovery Policy

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

1. In Cisco UCS Manager, click the Equipment tab in the navigation pane and select Equipment in the list on the left.

2. In the right pane, click the 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.

6. Click OK.

Enable Server and Uplink Ports

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

1. In Cisco UCS Manager, click the 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 the chassis, Cisco 2232 FEX (two per FEX), and direct connect UCS C-Series servers, right-click them, and select Configure as Server Port.

5. Click Yes to confirm server ports and click OK.

6. Verify that the ports connected to the chassis, C-series servers and to the Cisco 2232 FEX are now configured as server ports.

7. Select ports 27 and 28 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

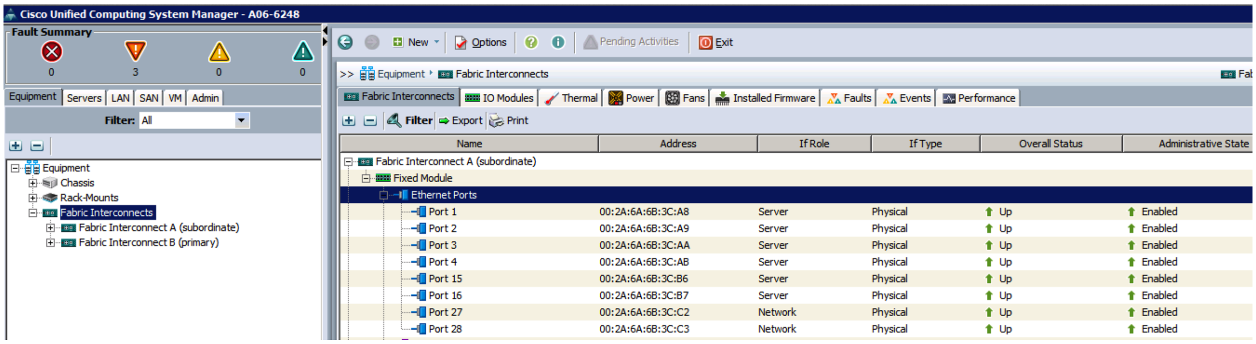

Name	Address	If Role	If Type	Overall Status	Administrative State
Fabric Interconnect A (subordinate)					
Fixed Module					
Ethernet Ports					
Port 1	00:2A:6A:6B:3C:A8	Server	Physical	Up	Enabled
Port 2	00:2A:6A:6B:3C:A9	Server	Physical	Up	Enabled
Port 3	00:2A:6A:6B:3C:AA	Server	Physical	Up	Enabled
Port 4	00:2A:6A:6B:3C:AB	Server	Physical	Up	Enabled
Port 15	00:2A:6A:6B:3C:B6	Server	Physical	Up	Enabled
Port 16	00:2A:6A:6B:3C:B7	Server	Physical	Up	Enabled
Port 27	00:2A:6A:6B:3C:C2	Network	Physical	Up	Enabled
Port 28	00:2A:6A:6B:3C:C3	Network	Physical	Up	Enabled

8. Click Yes to confirm uplink ports and click OK.

9. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

10. Expand Ethernet Ports.

11. Select the ports that are connected to the chassis, C-series servers or to the Cisco 2232 FEX (two per FEX), right-click them, and select Configure as Server Port.

12. Click Yes to confirm server ports and click OK.

13. Select ports 27 and 28 that are connected to the Cisco Nexus switches, right-click them, and select Configure as Uplink Port.

14. Click Yes to confirm the uplink ports and click OK.

15. Select port 29 connected to out of band management network switch on each Fabric Interconnect and select Configure as Uplink Port.

16. Click Yes to confirm the uplink ports and click OK.

17. Optional: If the out of band management switch is 1Gbps switch, for the port 29 on each Fabric Interconnect, right click the interface and select Show Navigator.

18. Click Show Interface in the Properties window.

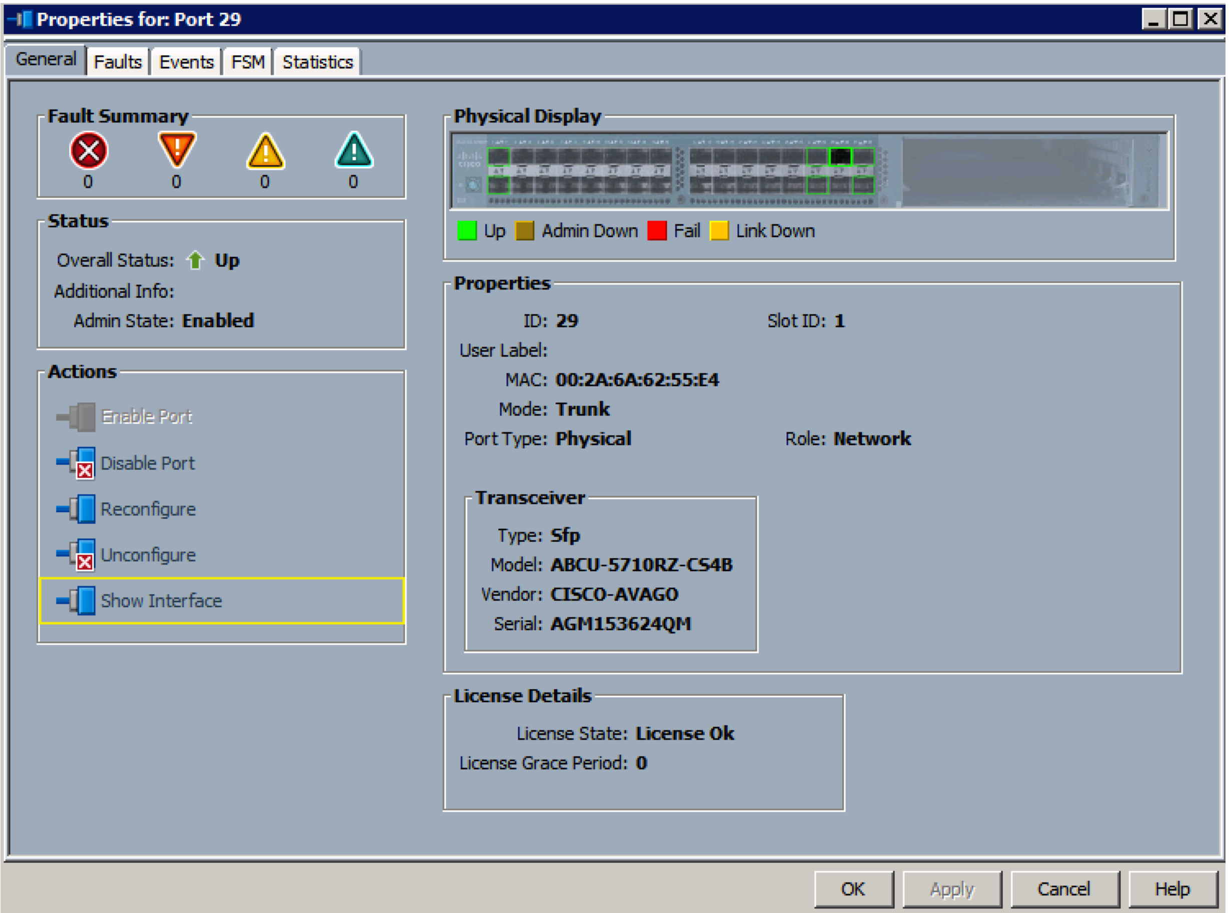

Properties for: Port 29

General | Faults | Events | FSM | Statistics

Fault Summary

0 0 0 0

Status

Overall Status: **Up**

Additional Info:

Admin State: **Enabled**

Actions

Enable Port

Disable Port

Reconfigure

Unconfigure

Show Interface

Physical Display

🟩 Up 🟫 Admin Down 🟥 Fail 🟨 Link Down

Properties

ID: **29** Slot ID: **1**

User Label:

MAC: **00:2A:6A:62:55:E4**

Mode: **Trunk**

Port Type: **Physical** Role: **Network**

Transceiver

Type: **Sfp**

Model: **ABCU-5710RZ-CS4B**

Vendor: **CISCO-AVAGO**

Serial: **AGM153624QM**

License Details

License State: **License Ok**

License Grace Period: **0**

19. Select 1Gbps as the speed of the interface.

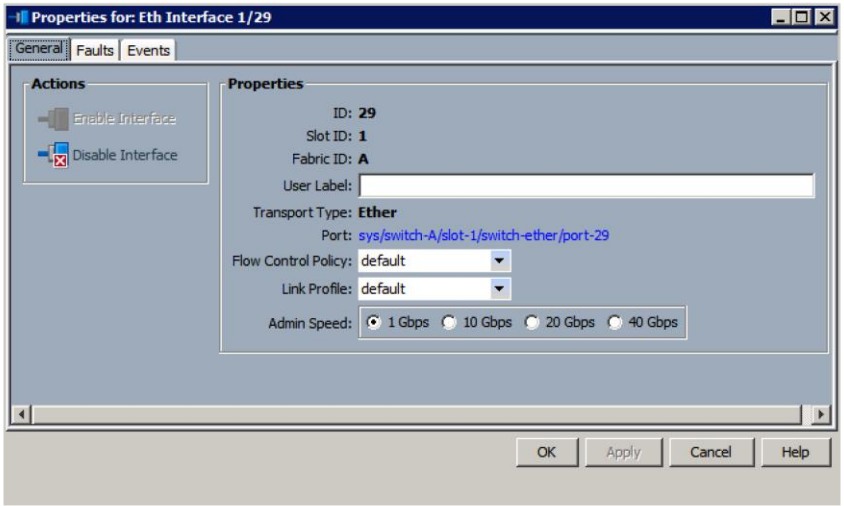

Properties for: Eth Interface 1/29

General | Faults | Events

Actions

Enable Interface

Disable Interface

Properties

ID: **29**

Slot ID: **1**

Fabric ID: **A**

User Label:

Transport Type: **Ether**

Port: sys/switch-A/slot-1/switch-ether/port-29

Flow Control Policy: default

Link Profile: default

Admin Speed: ● 1 Gbps ○ 10 Gbps ○ 20 Gbps ○ 40 Gbps

Acknowledge Cisco UCS Chassis and FEX

To acknowledge all Cisco UCS chassis and external 2232 FEX modules, complete the following steps:

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

2. Expand Chassis and select each chassis that is listed.

3. Right-click each chassis and select Acknowledge Chassis.

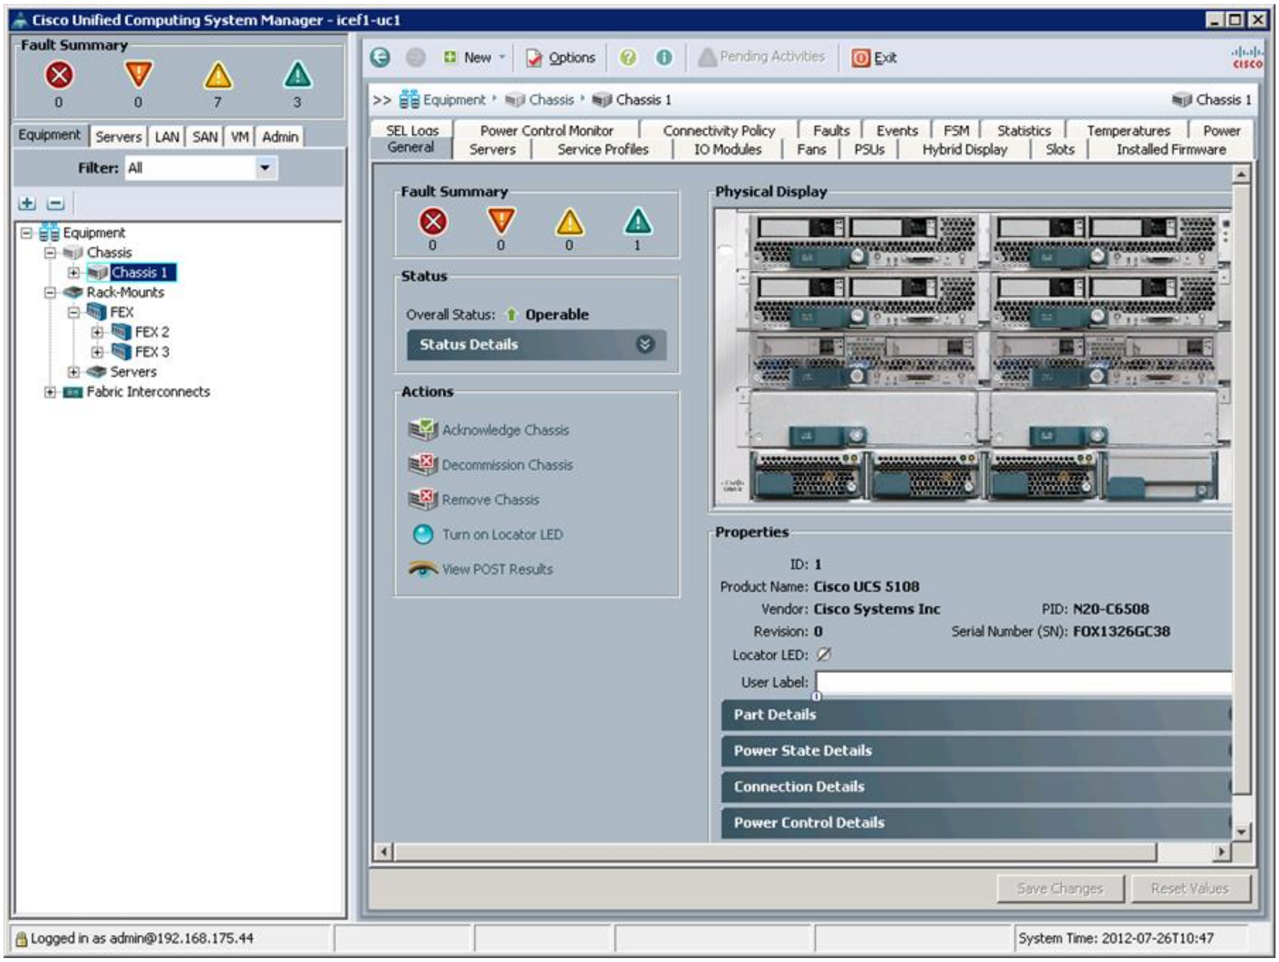

Overall Status: Operable

ID: 1
Product Name: Cisco UCS 5108
Vendor: Cisco Systems Inc
PID: N20-C6508
Revision: 0
Serial Number (SN): FOX1326GC38
Locator LED:
User Label:

Logged in as admin@192.168.175.44
System Time: 2012-07-26T10:47

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

5. If Nexus 2232 FEX is part of the configuration, expand Rack Mounts and FEX.

6. Right-click each FEX that is listed and select Acknowledge FEX.

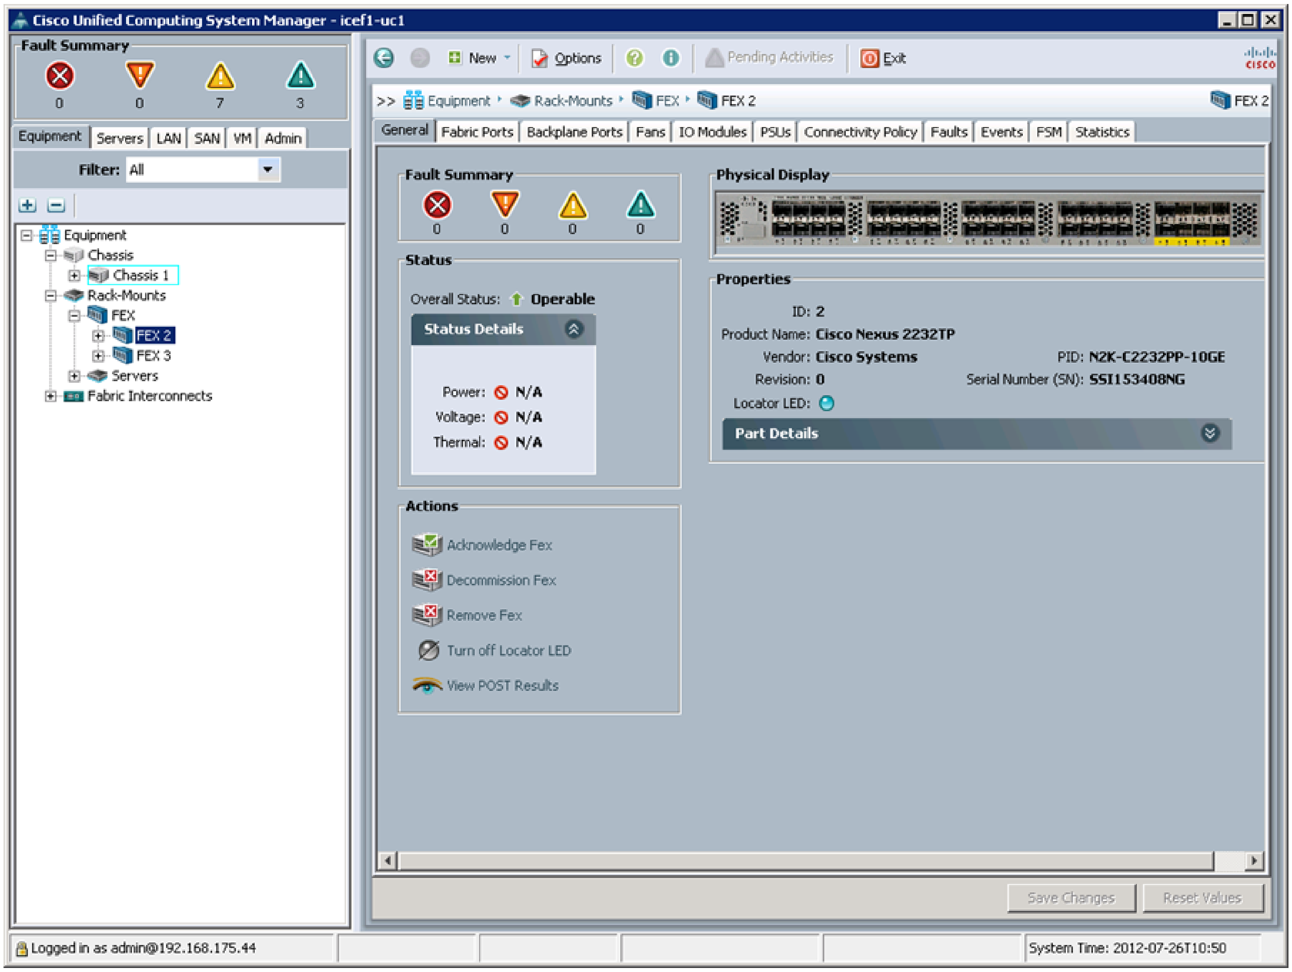

Status

Overall Status: Operable

Power: N/A
Voltage: N/A
Thermal: N/A

Properties

ID: 2
Product Name: Cisco Nexus 2232TP
Vendor: Cisco Systems
PID: N2K-C2232PP-10GE
Revision: 0
Serial Number (SN): SSI153408NG
Locator LED:

7. Click Yes and then click OK to complete acknowledging the FEX.

Create Uplink Port Channels to Cisco Nexus Switches

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

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

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

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

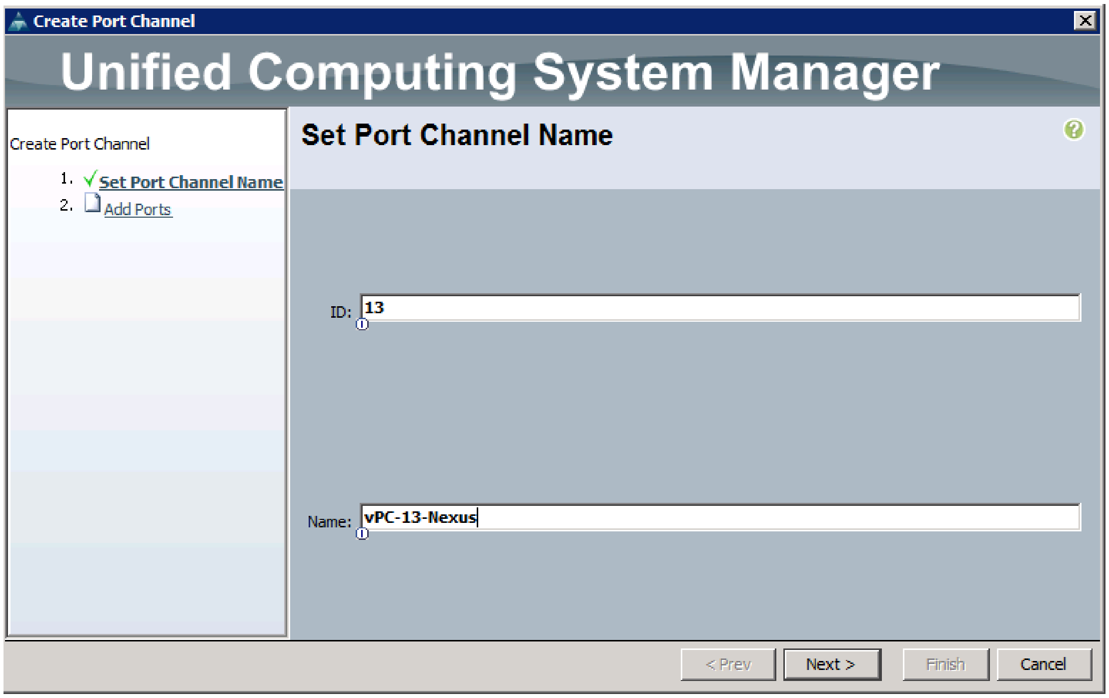

Set Port Channel Name

ID: 13

Name: vPC-13-Nexus

8. Select the following ports to be added to the port channel:

— Slot ID 1 and port 27

— Slot ID 1 and port 28

9. Click >> to add the ports to the port channel.

10. Click Finish to create the port channel.

11. Click OK.

12. In the navigation pane, under LAN > LAN Cloud, expand the fabric B tree.

13. Right-click Port Channels.

14. Select Create Port Channel.

15. Enter `14` as the unique ID of the port channel.

16. Enter `vPC-14-NEXUS` as the name of the port channel.

17. Click Next.

18. Select the following ports to be added to the port channel:

— Slot ID 1 and port 27

— Slot ID 1 and port 28

19. Click >> to add the ports to the port channel.

20. Click Finish to create the port channel.

21. Click OK.

Create MAC Address Pools

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

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

2. Select Pools > root.

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

3. Right-click MAC Pools under the root organization.

4. Select Create MAC Pool to create the MAC address pool.

5. Enter `MAC_Pool_A` as the name of the MAC pool.

6. Optional: Enter a description for the MAC pool.

7. Click Next.

8. Click Add.

9. Specify a starting MAC address.

Note: For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses.

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

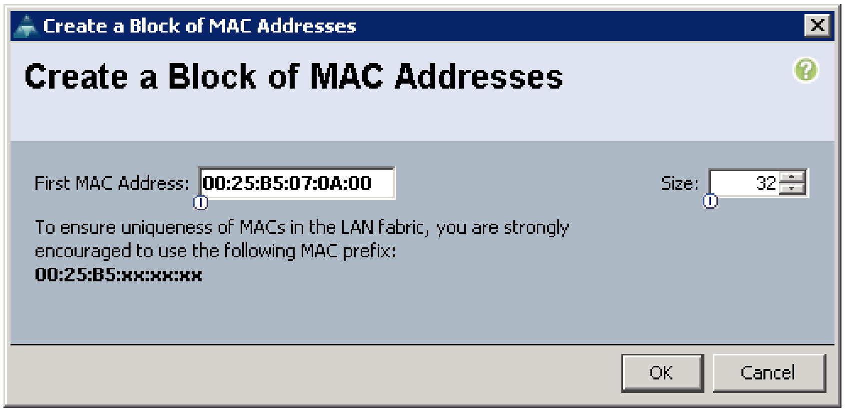

Create a Block of MAC Addresses

First MAC Address: 00:25:B5:07:0A:00

Size: 32

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

11. Click OK.

12. Click Finish.

13. In the confirmation message, click OK.

14. Right-click MAC Pools under the root organization.

15. Select Create MAC Pool to create the MAC address pool.

16. Enter `MAC_Pool_B` as the name of the MAC pool.

17. Optional: Enter a description for the MAC pool.

18. Click Next.

19. Click Add.

20. Specify a starting MAC address.

Note: For the FlexPod solution, it is recommended to place `0B` in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses.

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

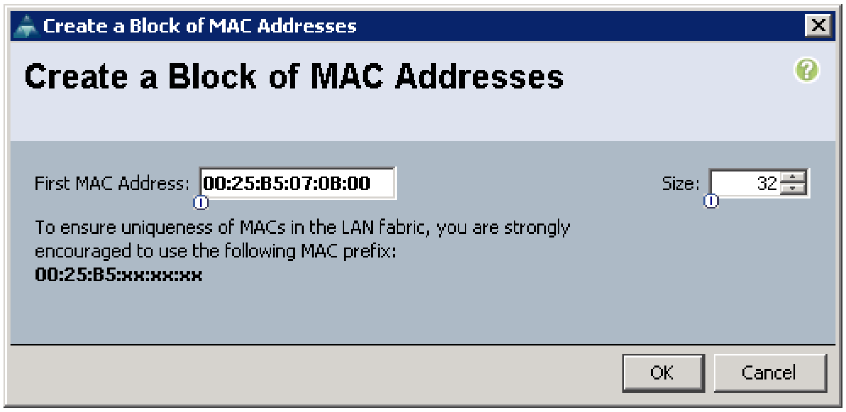

Create a Block of MAC Addresses

First MAC Address: 00:25:B5:07:0B:00

Size: 32

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

22. Click OK.

23. Click Finish.

24. In the confirmation message, click OK.

Create IQN Pools for iSCSI Boot

To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps.

1. In the UCS Manager, select the SAN tab on the left.

2. Select Pools > root.

3. Right-click IQN Pools under the root organization.

4. Select Create IQN Suffix Pool to create the IQN pool.

5. Enter `IQN_Pool` for the name of the IQN pool.

6. Optional: Enter a description for the IQN pool.

7. Enter `iqn.1992-08.com.cisco` as the prefix

8. Select Sequential for Assignment Order.

9. Click Next.

10. Click Add.

11. Enter `ucs-host` as the suffix.

12. Enter `1` in the From field.

13. Specify a size of the IQN block sufficient to support the available server resources.

14. Click OK.

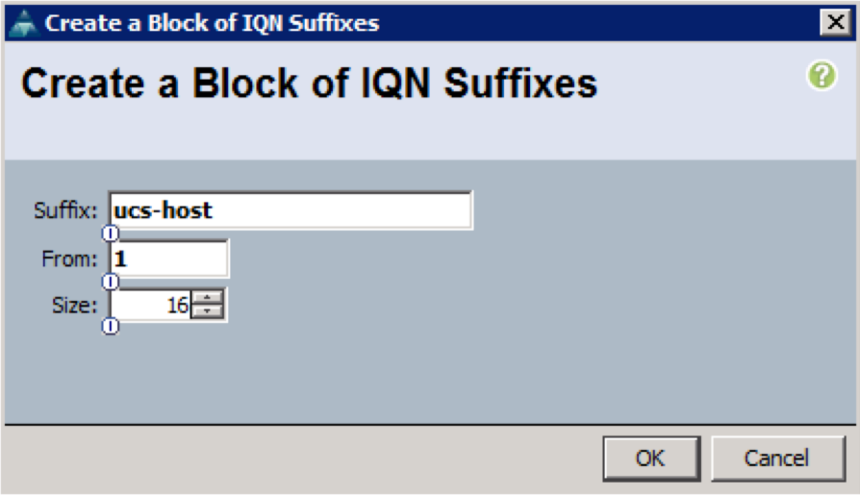

Create a Block of IQN Suffixes

Create a Block of IQN Suffixes

Suffix: ucs-host

From: 1

Size: 16

15. Click Finish.

16. In the message box that displays, click OK.

Create IP Pools for iSCSI Boot

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, select the LAN tab on the left.

2. Select Pools > root.

3. Two IP pools are created, one for each switching fabric.

4. Right-click IP Pools under the root organization.
5. Select Create IP Pool to create the IP pool.
6. Enter `iSCSI_IP_Pool_A` for the name of the IP pool.
7. Optional: Enter a description of the IQN pool.
8. Select Sequential for Assignment Order.
9. Click Next.
10. Click Add.
11. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
12. Set the size to enough addresses to accommodate the servers.
13. Click OK.
14. Click Finish.
15. Right-click IP Pools under the root organization.
16. Select Create IP Pool to create the IP pool.
17. Enter `iSCSI_IP_Pool_B` for the name of the IP pool.
18. Optional: Enter a description of the IQN pool.
19. Select Sequential for Assignment Order.
20. Click Next.
21. Click Add.
22. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
23. Set the size to enough addresses to accommodate the servers.
24. Click OK.
25. Click Finish.

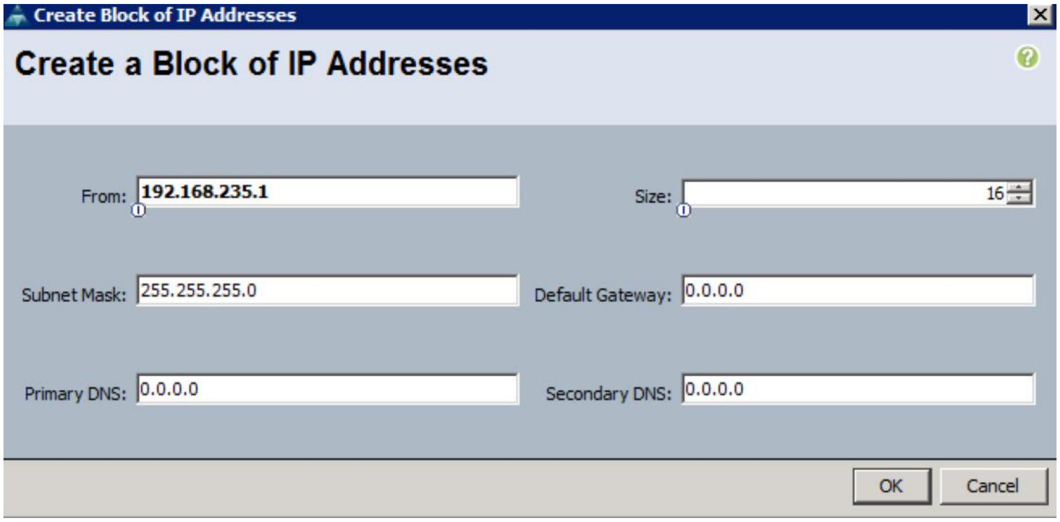

Create Block of IP Addresses

Create a Block of IP Addresses

From: 192.168.235.1

Size: 16

Subnet Mask: 255.255.255.0

Default Gateway: 0.0.0.0

Primary DNS: 0.0.0.0

Secondary DNS: 0.0.0.0

Create UUID Suffix Pool

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

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

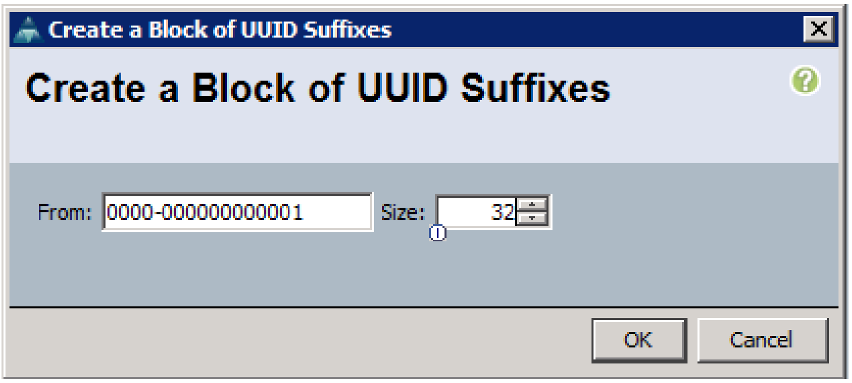

Create a Block of UUID Suffixes

From: 0000-000000000001 Size: 32

12. Click OK.

13. Click Finish.

14. Click OK.

Create Server Pool

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

1. Consider creating unique server pools to achieve the granularity that is required in your environment.
2. In Cisco UCS Manager, click the Servers tab in the navigation pane.
3. Select Pools > root.
4. Right-click Server Pools.
5. Select Create Server Pool.
6. Enter `Infra_Pool` as the name of the server pool.
7. Optional: Enter a description for the server pool.
8. Click Next.
9. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the `Infra_Pool` server pool.
10. Click Finish.
11. Click OK.

Create VLANs

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

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

Note: In this procedure, four unique VLANs and a range of 100 VLANs for APIC are created. The VLAN IDs used in this step are shown in **Error! Reference source not found.**.

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

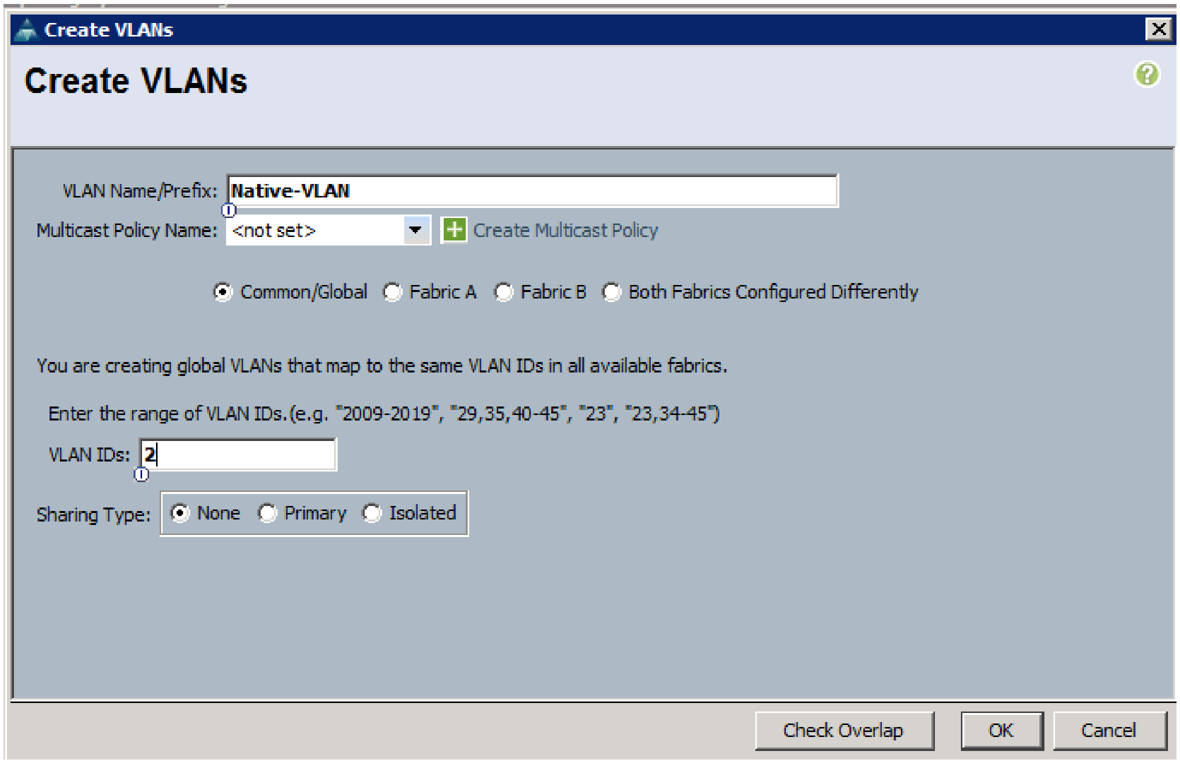

VLAN Name/Prefix: Native-VLAN

Multicast Policy Name: <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: 2

Sharing Type: ● None ○ Primary ○ Isolated

10. Expand the list of VLANs in the navigation pane, right-click the newly created `Native-VLAN` and select Set as Native VLAN.
11. Click Yes, and then click OK.
12. Right-click VLANs.
13. Select Create VLANs.

14. Enter `iSCSI-A-VLAN` as the name of the VLAN to be used for the first iSCSI VLAN.

15. Keep the Common/Global option selected for the scope of the VLAN.

16. Enter the VLAN ID for the first iSCSI VLAN.

17. Click OK, then OK.

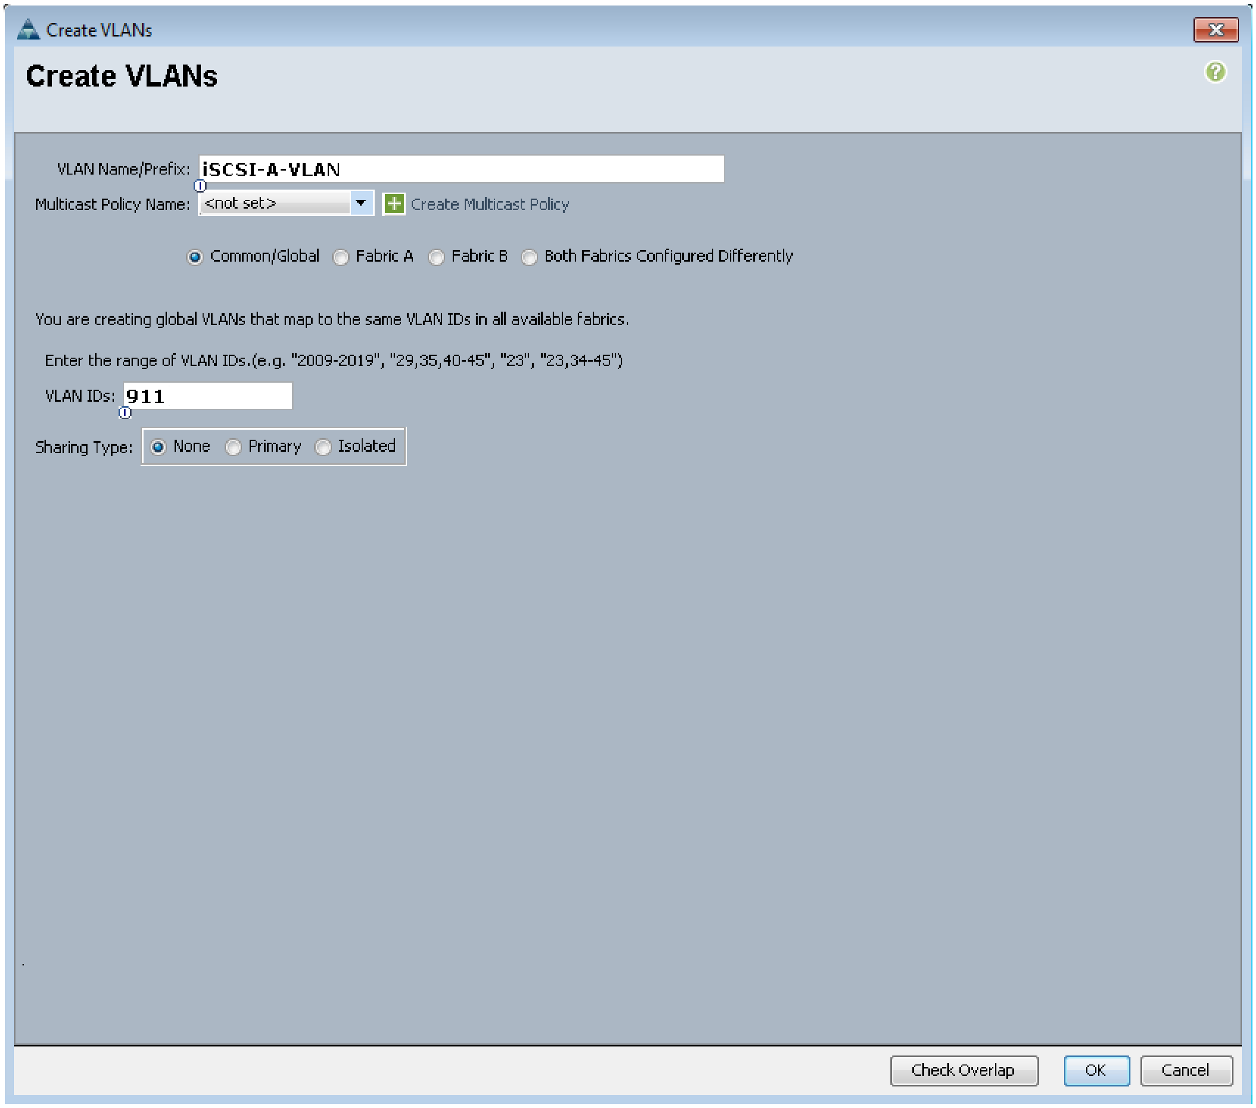

Create VLANs

VLAN Name/Prefix: iSCSI-A-VLAN

Multicast Policy Name: <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: 911

Sharing Type: ● None ○ Primary ○ Isolated

18. Right-click VLANs.

19. Select Create VLANs.

20. Enter `iSCSI-B-VLAN` as the name of the VLAN to be used for the second iSCSI VLAN.

21. Keep the Common/Global option selected for the scope of the VLAN.

22. Enter the VLAN ID for the second iSCSI VLAN.

23. Click OK, then OK.

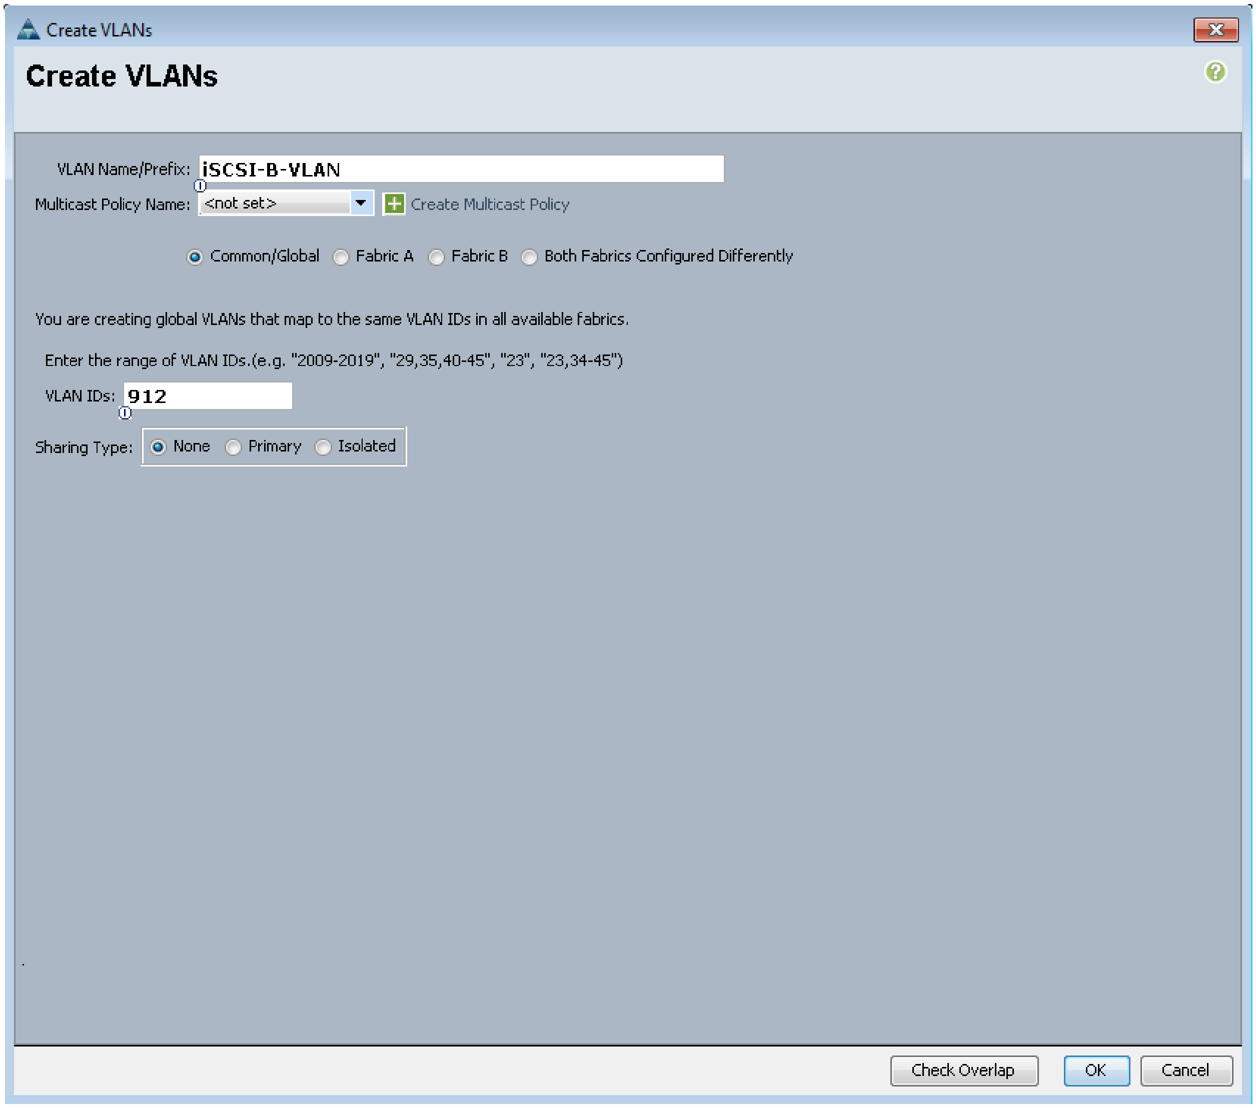

Create VLANs

VLAN Name/Prefix: iSCSI-B-VLAN

Multicast Policy Name: <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: 912

Sharing Type: ● None ○ Primary ○ Isolated

24. Right-click VLANs.

25. Select Create VLANs.

26. Enter `OOB-Mgmt` as the name of the VLAN to be used for management traffic.

27. Keep the Common/Global option selected for the scope of the VLAN.

28. Enter `<3177>` as the ID of the management VLAN.

29. Keep the Sharing Type as None.

30. Click OK, and then click OK again.

31. Right-click VLANs.

32. Select Create VLANs.

33. Enter `INFRA-NFS` as the name of the VLAN to be used for NFS VMk ports.

34. Keep the Common/Global option selected for the scope of the VLAN.

35. Enter the `<3270>` for the NFS VLAN.

36. Keep the Sharing Type as None.

37. Click OK, and then click OK again.

38. Right-click VLANs.

39. Select Create VLANs.

40. Enter `APIC-` as the prefix of the VLAN to be used for APIC.

41. Keep the Common/Global option selected for the scope of the VLAN.

42. Enter the `<1101-1200>` for VLAN IDs.

43. Keep the Sharing Type as None.

44. Click OK, and then click OK again.

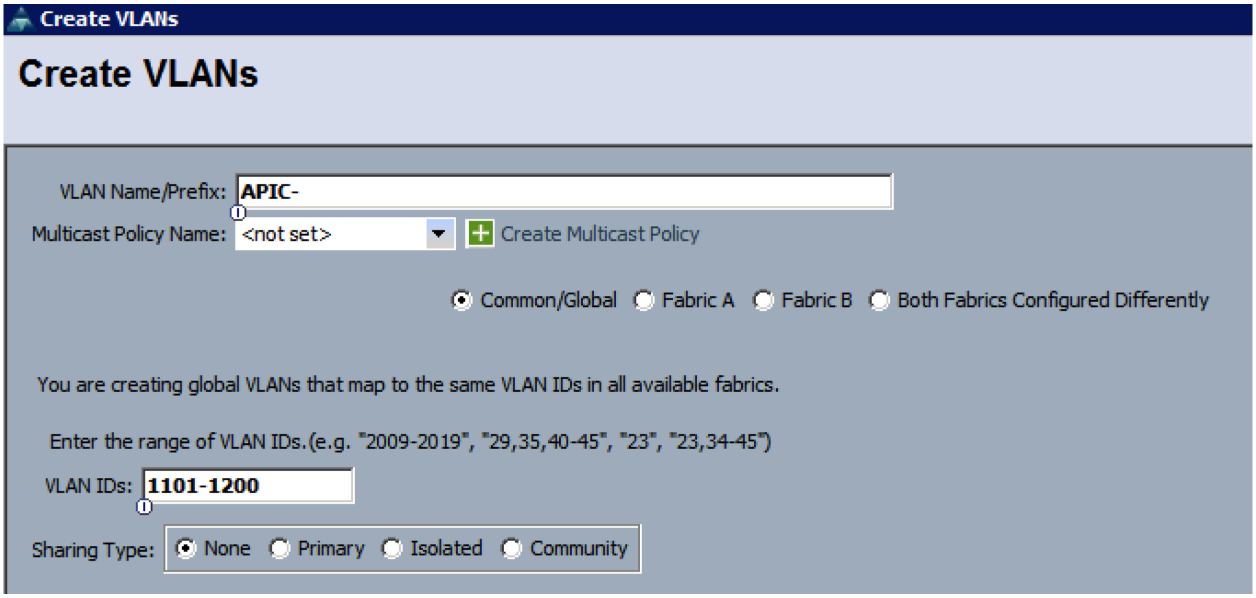

Create VLANs

VLAN Name/Prefix: APIC-

Multicast Policy Name: <not set>

● Common/Global ○ Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.

Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: 1101-1200

Sharing Type: ● None ○ Primary ○ Isolated ○ Community

Create VLAN Groups

To configure split layer-2 domain in UCS, two VLAN groups need to be created and attached to different uplink ports. In the procedure below, a VLAN group `OOB-Mgmt` is configured only with out of band management VLAN (3177) and is attached to port 19 on each Fabric Interconnect. VLAN group `Uplink-PortChannel` is configured to carry all the remaining VLANs and is attached to port-channel 13 and 14. To create the VLAN groups, complete the following steps:

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

2. Select LAN > LAN Cloud.

3. Right-click VLAN Groups.

4. Select Create VLAN Group.

5. Use `OOB-Mgmt` as the VLAN Group Name

6. Select the `Native-VLAN` (2) and `OOB-Mgmt` VLAN (3177). Click the radio button to set `Native-VLAN` as native VLAN.

7. Click Next to add Uplink Ports. Select Port 29 on both Fabric Interconnect A and B.

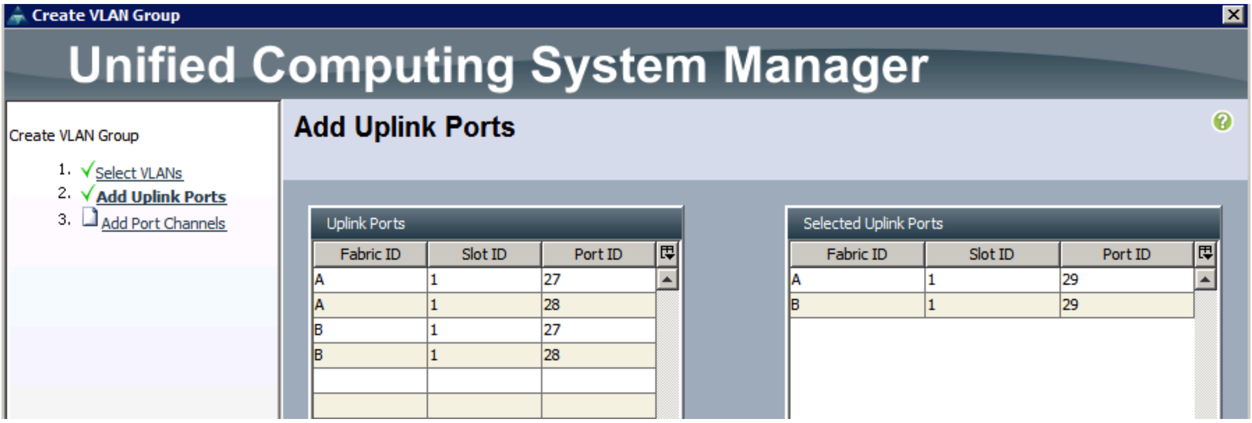

Add Uplink Ports

Uplink Ports

Fabric ID	Slot ID	Port ID
A	1	27
A	1	28
B	1	27
B	1	28

Selected Uplink Ports

Fabric ID	Slot ID	Port ID
A	1	29
B	1	29

8. Click Finish.

9. Right-click VLAN Groups.

10. Select Create VLAN Group.

11. Use `Uplink-PortChannel` as the VLAN Group Name.

12. Select the `Infra-NFS` (3270), `iSCSI-A-VLAN` (911), `iSCSI-B-VLAN` (912) and `APIC-1101` through `APIC-1200` (all 100) VLANs.

13. Click Next twice to add Uplink Port Channels. Select Port Channel 13 and 14.

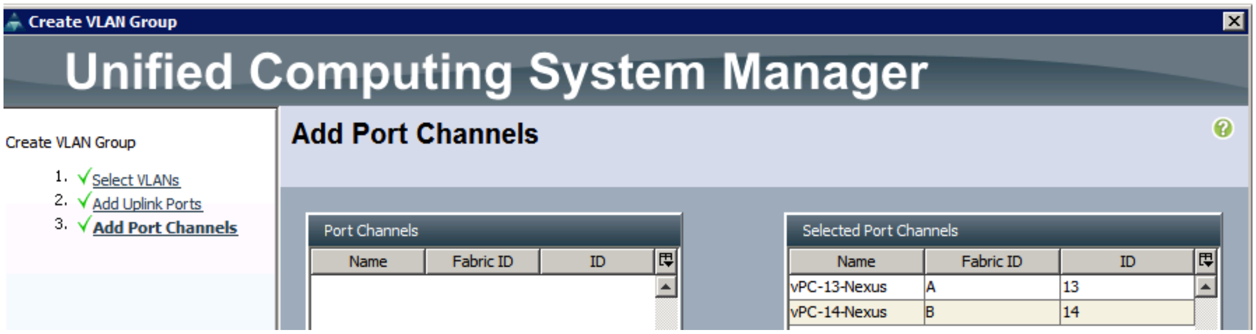

Add Port Channels

Port Channels

Name	Fabric ID	ID

Selected Port Channels

Name	Fabric ID	ID
vPC-13-Nexus	A	13
vPC-14-Nexus	B	14

14. Click Finish.

Create Host Firmware Package

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

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Host Firmware Packages.
4. Select Create Host Firmware Package.
5. Enter `VM-Host-Infra` as the name of the host firmware package.
6. Leave Simple selected.
7. Select the version 2.2(3d) for both the Blade and Rack Packages.
8. Click OK to create the host firmware package.
9. Click OK.

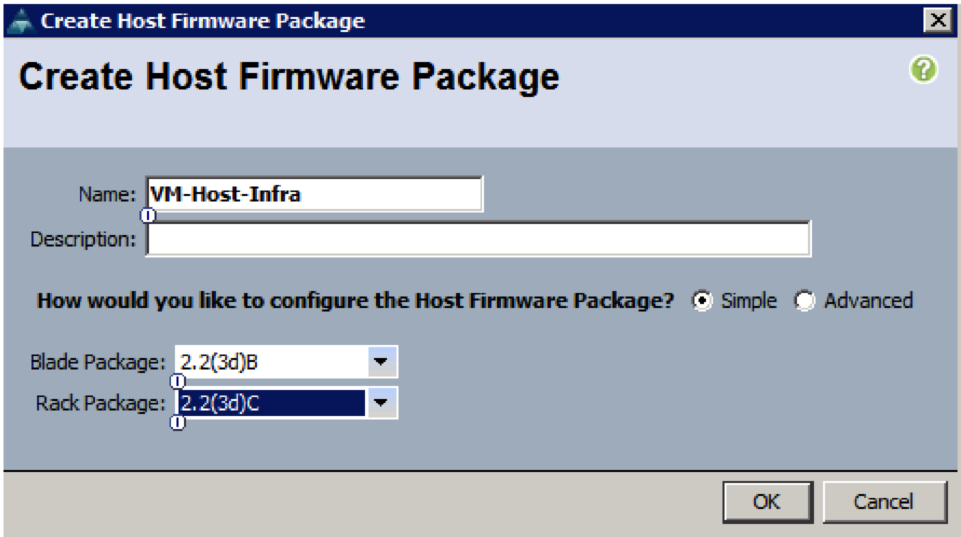

Create Host Firmware Package

Name: VM-Host-Infra

Description:

How would you like to configure the Host Firmware Package? ● Simple ○ Advanced

Blade Package: 2.2(3d)B

Rack Package: 2.2(3d)C

Set Jumbo Frames in Cisco UCS Fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, click the LAN tab in the navigation pane.
2. Select LAN > LAN Cloud > QoS System Class.

3. In the right pane, click the General tab.

4. On the Best Effort row, enter `9216` in the box under the MTU column.

5. Click Save Changes in the bottom of the window.

6. Click OK.

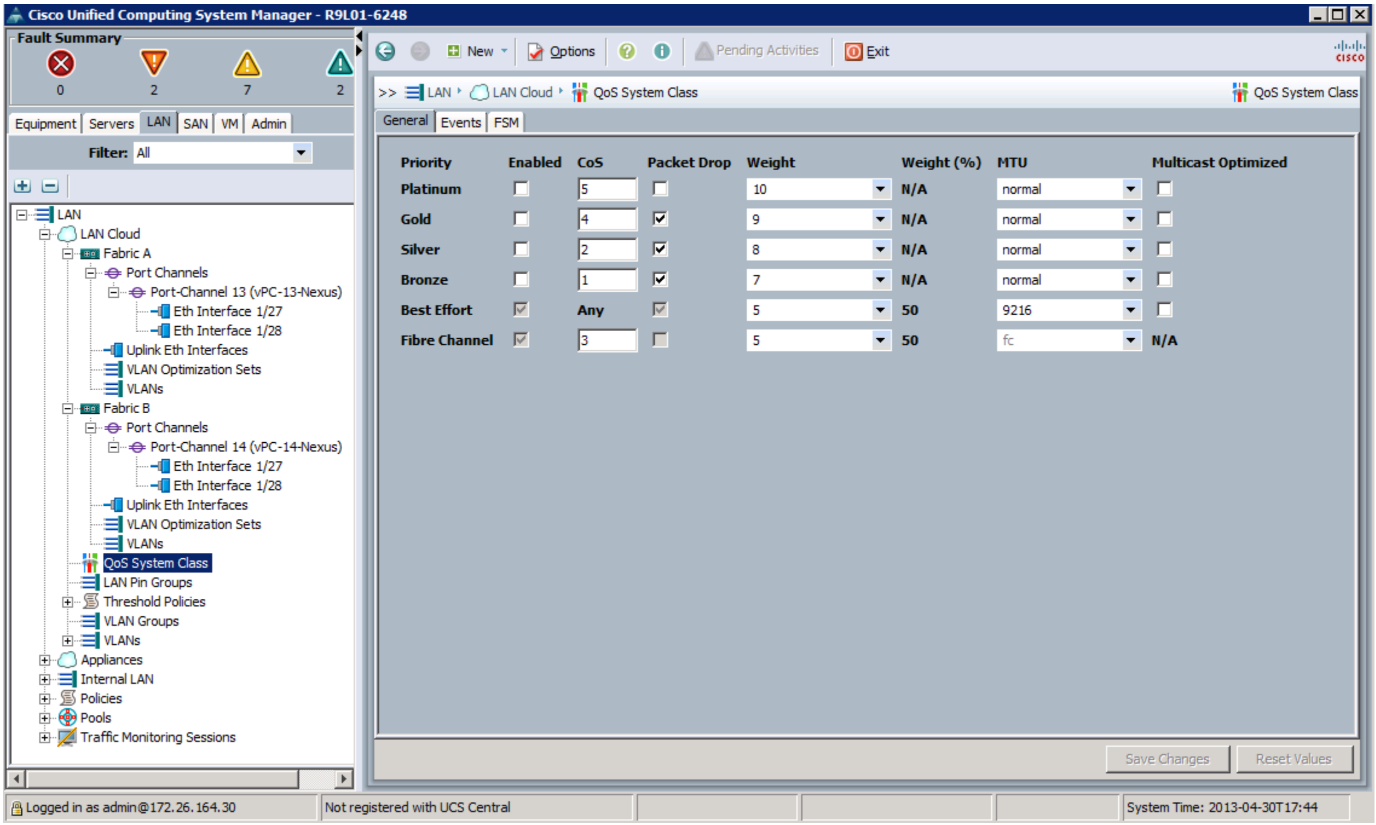

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	☐	5	☐	10	N/A	normal	☐
Gold	☐	4	☑	9	N/A	normal	☐
Silver	☐	2	☑	8	N/A	normal	☐
Bronze	☐	1	☑	7	N/A	normal	☐
Best Effort	☑	Any	☑	5	50	9216	☐
Fibre Channel	☑	3	☐	5	50	fc	N/A

Create Local Disk Configuration Policy (Optional)

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

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

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

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

2. Select Policies > root.

3. Right-click Local Disk Config Policies.

4. Select Create Local Disk Configuration Policy.

5. Enter `iSCSI-Boot` as the local disk configuration policy name.

6. Change the mode to No Local Storage.

7. Click OK to create the local disk configuration policy.

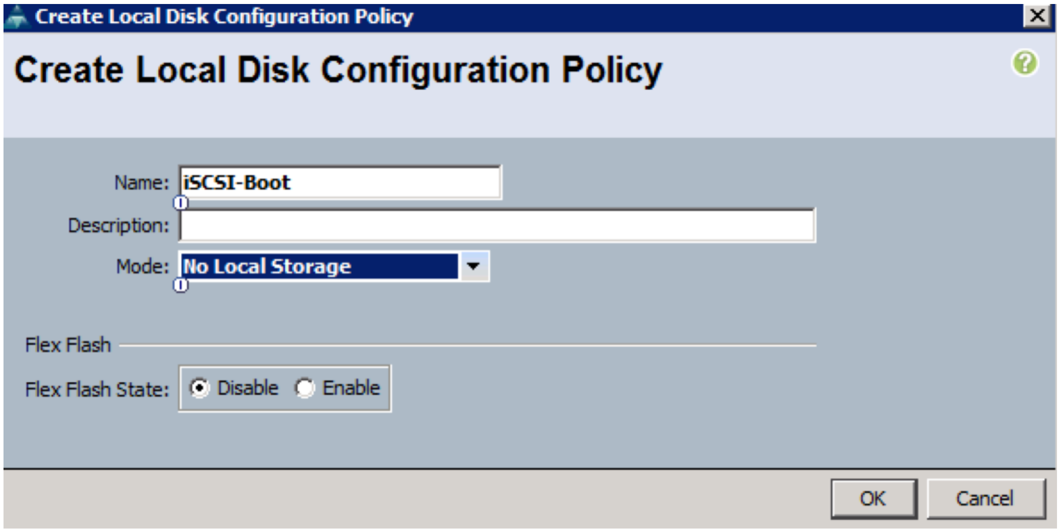

Create Local Disk Configuration Policy

Name: iSCSI-Boot

Description:

Mode: No Local Storage

Flex Flash

Flex Flash State: ● Disable ○ Enable

8. Click OK.

Create Network Control Policy for Cisco Discovery Protocol

To create a network control policy that enables Cisco Discovery Protocol (CDP) on virtual network ports, complete the following steps:

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

2. Select Policies > root.

3. Right-click Network Control Policies.

4. Select Create Network Control Policy.

5. Enter `Enable_CDP` as the policy name.

6. For CDP, select the Enabled option.

7. Click OK to create the network control policy.

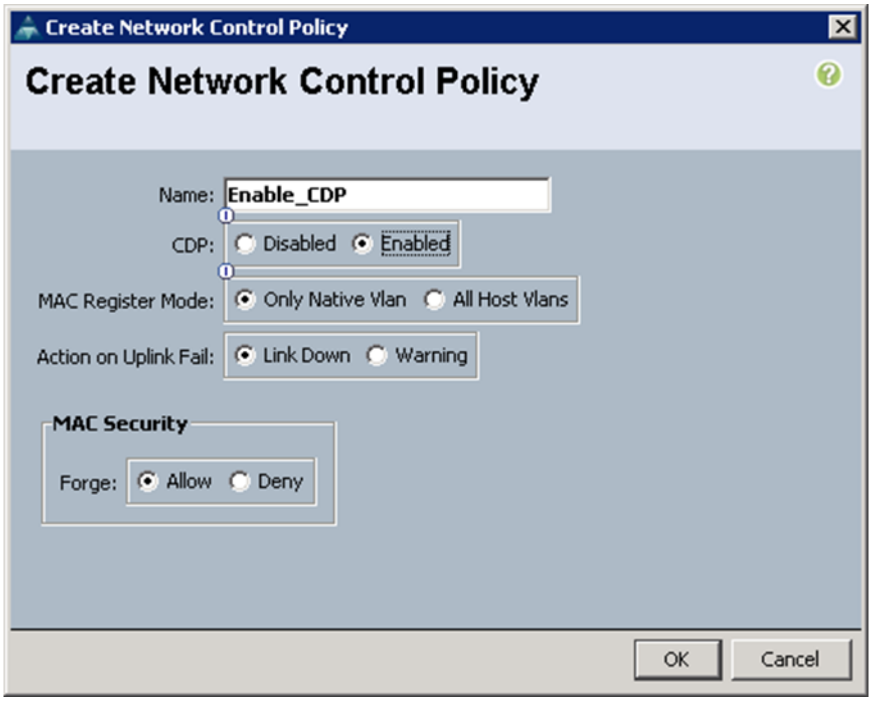

Create Network Control Policy

Name: Enable_CDP

CDP: ○ Disabled ● Enabled

MAC Register Mode: ● Only Native Vlan ○ All Host Vlans

Action on Uplink Fail: ● Link Down ○ Warning

MAC Security

Forge: ● Allow ○ Deny

8. Click OK.

Create Power Control Policy

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

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

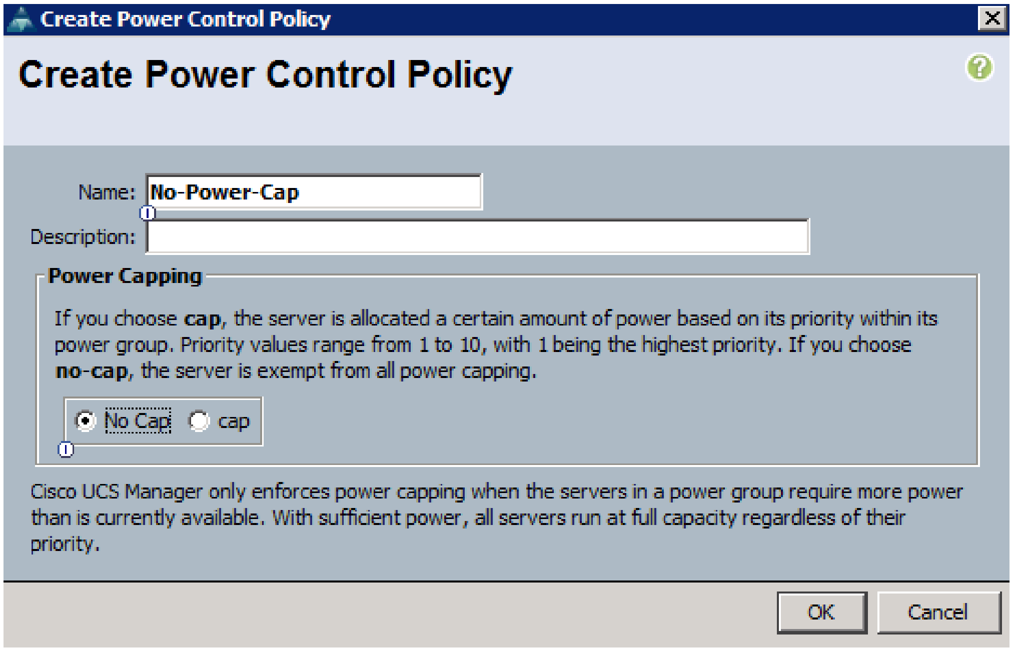

Create Power Control Policy

Name: No-Power-Cap

Description:

Power Capping

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

● No Cap ○ cap

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

Create Server Pool Qualification Policy (Optional)

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

Note: This example creates a policy for a Cisco UCS B200-M3 server.

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification.
5. Enter `UCSB-B200-M3` as the name for the policy.
6. Select Create Server PID Qualifications.
7. Enter `UCSB-B200-M3` as the PID.
8. Click OK to create the server pool qualification policy.
9. Click OK, and then click OK again.

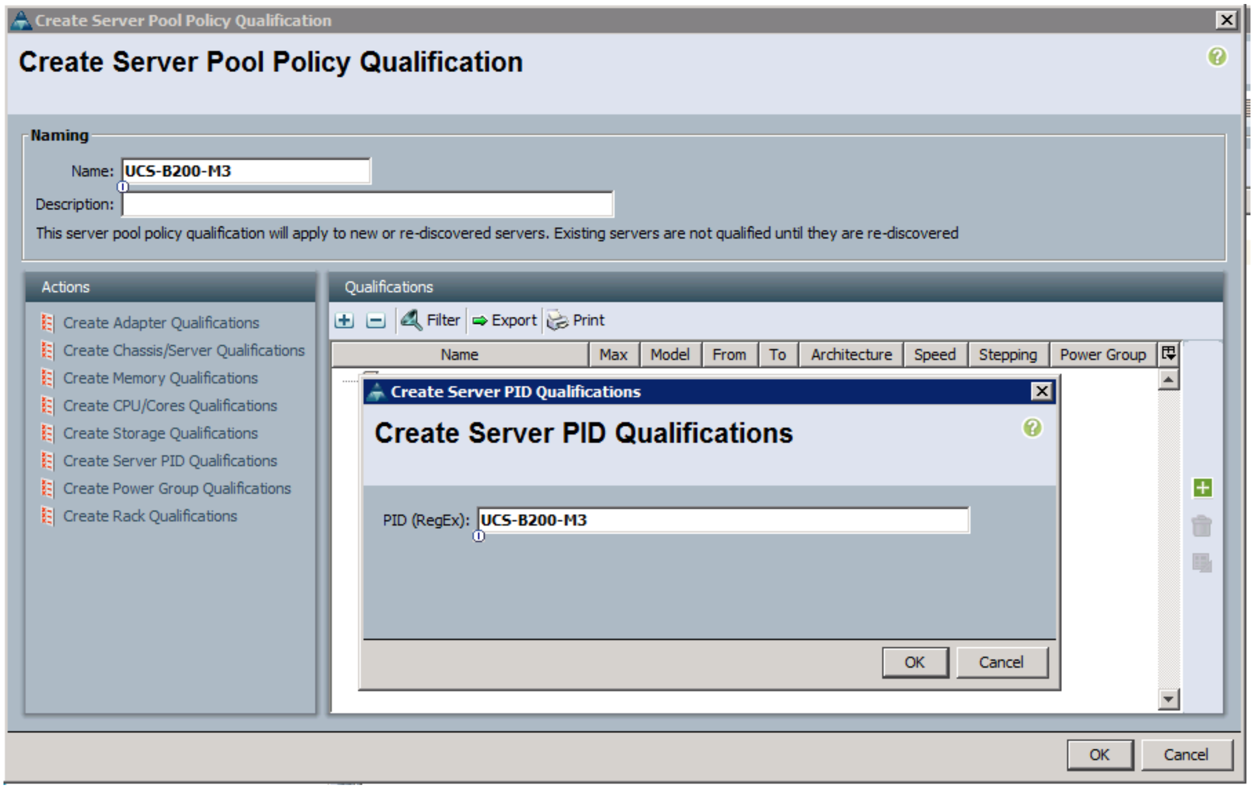

Create Server Pool Policy Qualification

Naming

Name: UCS-B200-M3

Description:

This server pool policy qualification will apply to new or re-discovered servers. Existing servers are not qualified until they are re-discovered

Create Server PID Qualifications

PID (RegEx): UCS-B200-M3

Actions
Create Adapter Qualifications
Create Chassis/Server Qualifications
Create Memory Qualifications
Create CPU/Cores Qualifications
Create Storage Qualifications
Create Server PID Qualifications
Create Power Group Qualifications
Create Rack Qualifications
Qualifications
Filter
Export
Print
Name
Max
Model
From
To
Architecture
Speed
Stepping
Power Group
OK
Cancel
OK
Cancel

Create Server BIOS Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter `VM-Host-Infra` as the BIOS policy name.
6. Change the Quiet Boot setting to Disabled.
7. Click Finish to create the BIOS policy.

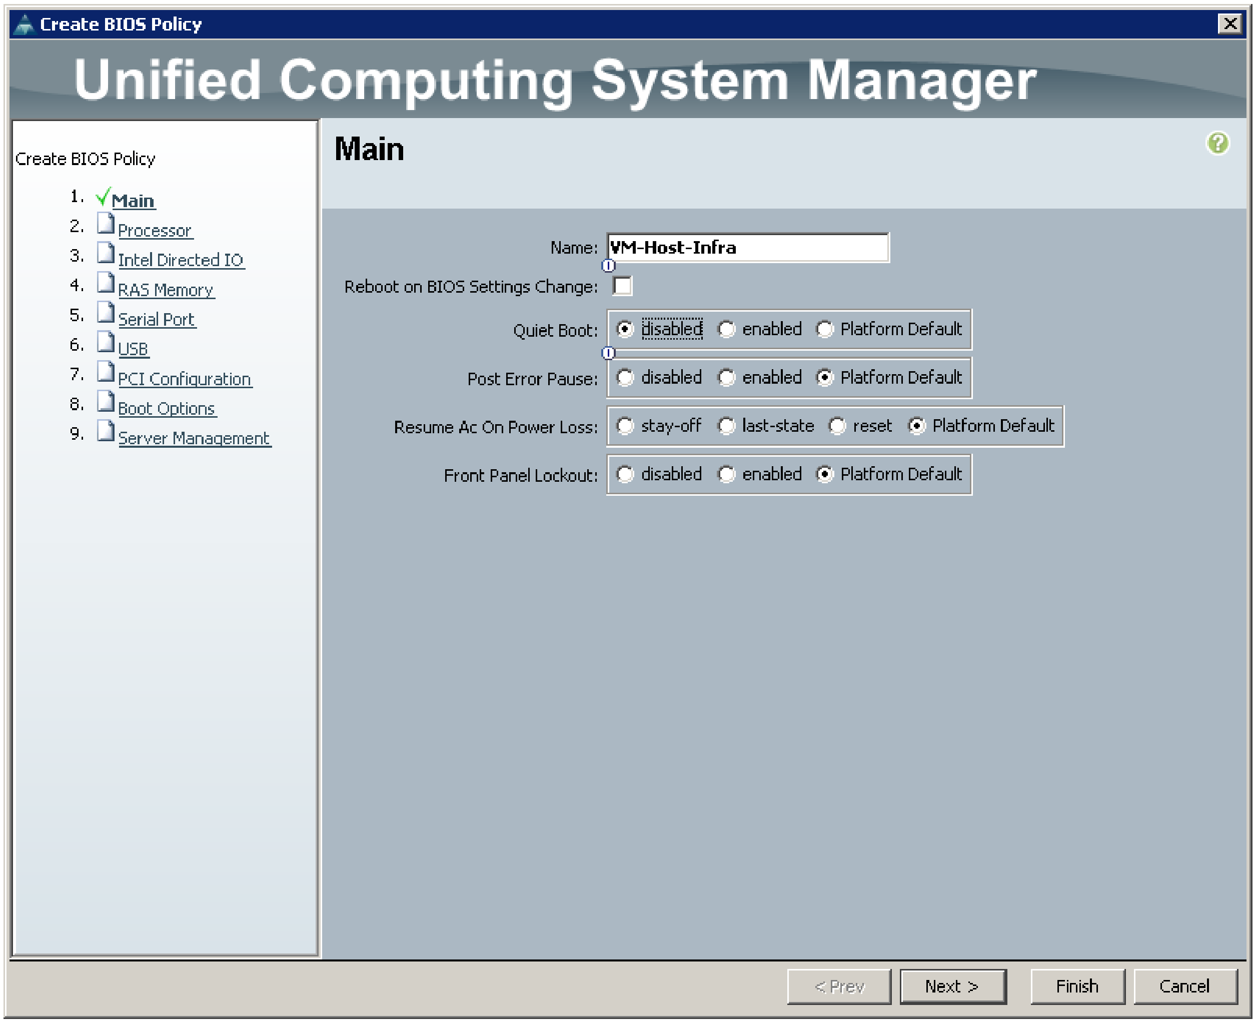

Name: VM-Host-Infra

Reboot on BIOS Settings Change: □

Quiet Boot: ● disabled ○ enabled ○ Platform Default

Post Error Pause: ○ disabled ○ enabled ● Platform Default

Resume Ac On Power Loss: ○ stay-off ○ last-state ○ reset ● Platform Default

Front Panel Lockout: ○ disabled ○ enabled ● Platform Default

8. Click OK.

Create vNIC/vHBA Placement Policy for Virtual Machine Infrastructure Hosts

To create a vNIC/vHBA placement policy for the infrastructure hosts, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click vNIC/vHBA Placement Policies.
4. Select Create Placement Policy.
5. Enter `VM-Host-Infra` as the name of the placement policy.
6. Click 1 and select Assigned Only.
7. Click OK, and then click OK again.

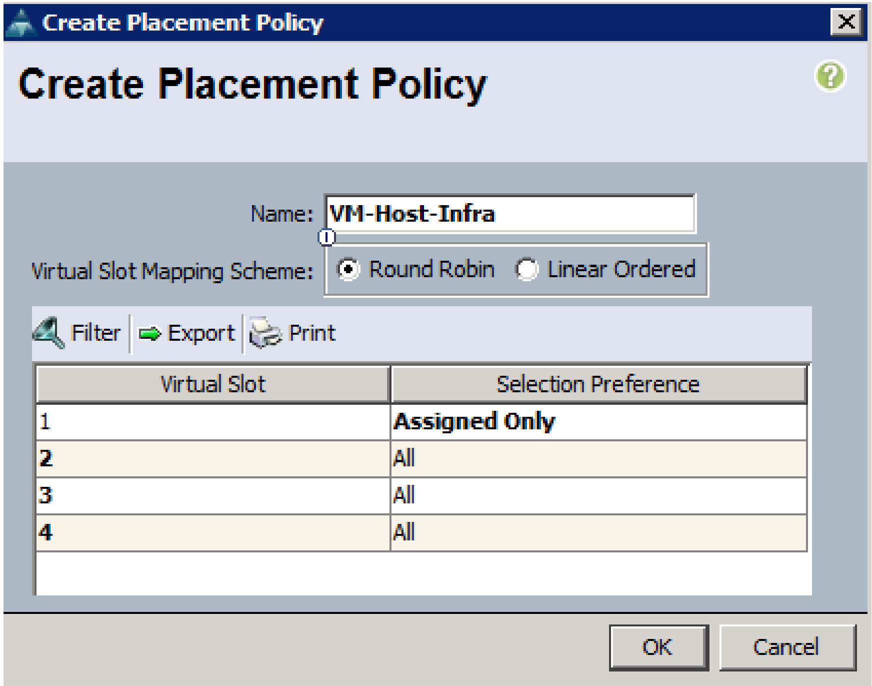

Create Placement Policy

Name: VM-Host-Infra

Virtual Slot Mapping Scheme: ● Round Robin ○ Linear Ordered

Virtual Slot	Selection Preference
1	**Assigned Only**
2	All
3	All
4	All

Update Default Maintenance Policy

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Click Save Changes.
6. Click OK to accept the change.

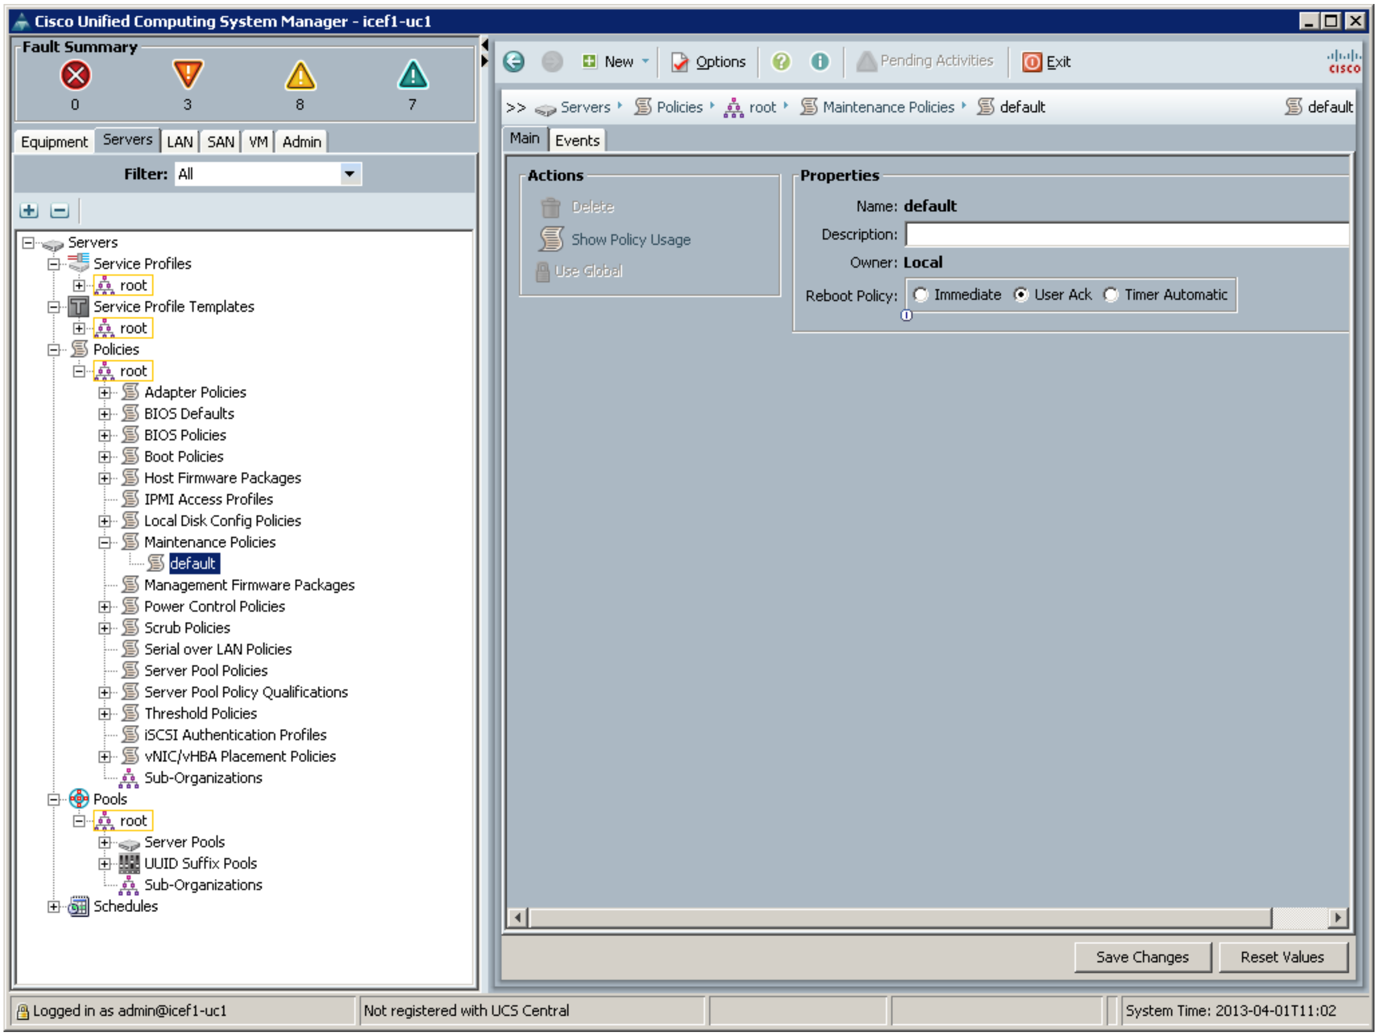

Name: **default**

Description:

Owner: **Local**

Reboot Policy: ○ Immediate ● User Ack ○ Timer Automatic

Logged in as admin@icef1-uc1 | Not registered with UCS Central | System Time: 2013-04-01T11:02

Create vNIC Templates

A total of 8 vNIC Templates will be created. Infrastructure ESXi hosts use all 8 templates (8 vNICS) while the application ESXi servers utilize only six (6 vNICs). To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the following steps.

Create Data vNICs

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

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template.

5. Enter `vNIC_Template_A` as the vNIC template name.

6. Keep Fabric A selected.

7. Do not select the Enable Failover checkbox.

8. Under Target, make sure that the VM checkbox is not selected.

9. Select Updating Template as the Template Type.

10. Under VLANs, select the checkboxes for `APIC-1101` through `APIC-1200`, and `default` VLANs.

11. Set `default` as the native VLAN.

12. For MTU, enter `9000`.

13. In the MAC Pool list, select `MAC_Pool_A`.

14. In the Network Control Policy list, select `Enable_CDP`.

15. Click OK to create the vNIC template.

16. Click OK.

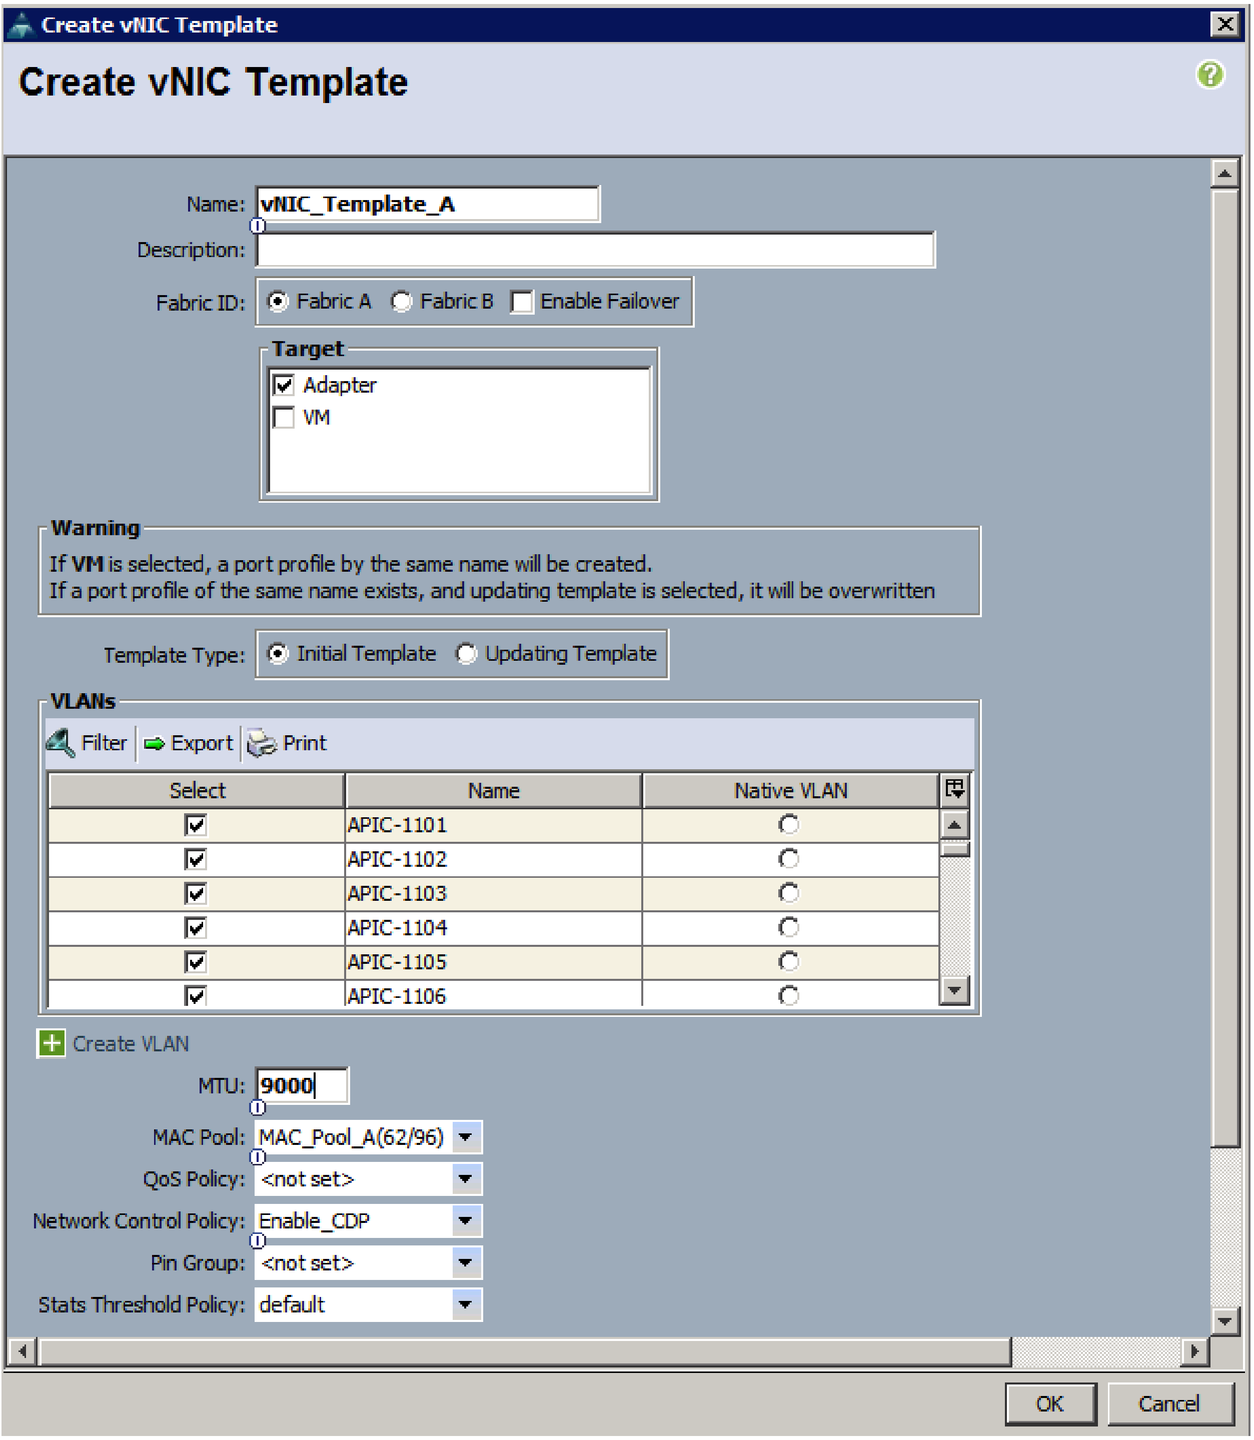

Create vNIC Template

Name: vNIC_Template_A

Description:

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

Target

☑ Adapter

□ VM

Warning

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

Template Type: ● Initial Template ○ Updating Template

VLANs

Select	Name	Native VLAN
☑	APIC-1101	○
☑	APIC-1102	○
☑	APIC-1103	○
☑	APIC-1104	○
☑	APIC-1105	○
☑	APIC-1106	○

MTU: 9000

MAC Pool: MAC_Pool_A(62/96)

QoS Policy: <not set>

Network Control Policy: Enable_CDP

Pin Group: <not set>

Stats Threshold Policy: default

17. In the navigation pane, select the LAN tab.

18. Select Policies > root.

19. Right-click vNIC Templates.

20. Select Create vNIC Template

21. Enter `vNIC_Template_B` as the vNIC template name.

22. Select Fabric B.

23. Do not select the Enable Failover checkbox.

24. Under Target, make sure the VM checkbox is not selected.

25. Select Updating Template as the template type.

26. Under VLANs, select the checkboxes for `APIC-1101` through `APIC-1200`, and `default` VLANs.

27. Set `default` as the native VLAN.

28. For MTU, enter `9000`.

29. In the MAC Pool list, select `MAC_Pool_B`.

30. In the Network Control Policy list, select `Enable_CDP`.

31. Click OK to create the vNIC template.

32. Click OK.

Create iSCSI vNICSs

1. Select the LAN tab on the left.

2. Select Policies > root.

3. Right-click vNIC Templates.

4. Select Create vNIC Template.

5. Enter `iSCSI_Template_A` as the vNIC template name.

6. Leave Fabric A selected. Do not select the Enable Failover checkbox.

7. Under Target, make sure that the VM checkbox is not selected.

8. Select Updating Template for Template Type.

9. Under VLANs, select `iSCSI-A-VLAN` (911).

10. Set `iSCSI-A-VLAN` as the native VLAN.

11. Under MTU, enter 9000.

12. From the MAC Pool list, select `MAC_Pool_A`.

13. From the Network Control Policy list, select `Enable_CDP`.

14. Click OK to complete creating the vNIC template.

15. Click OK.

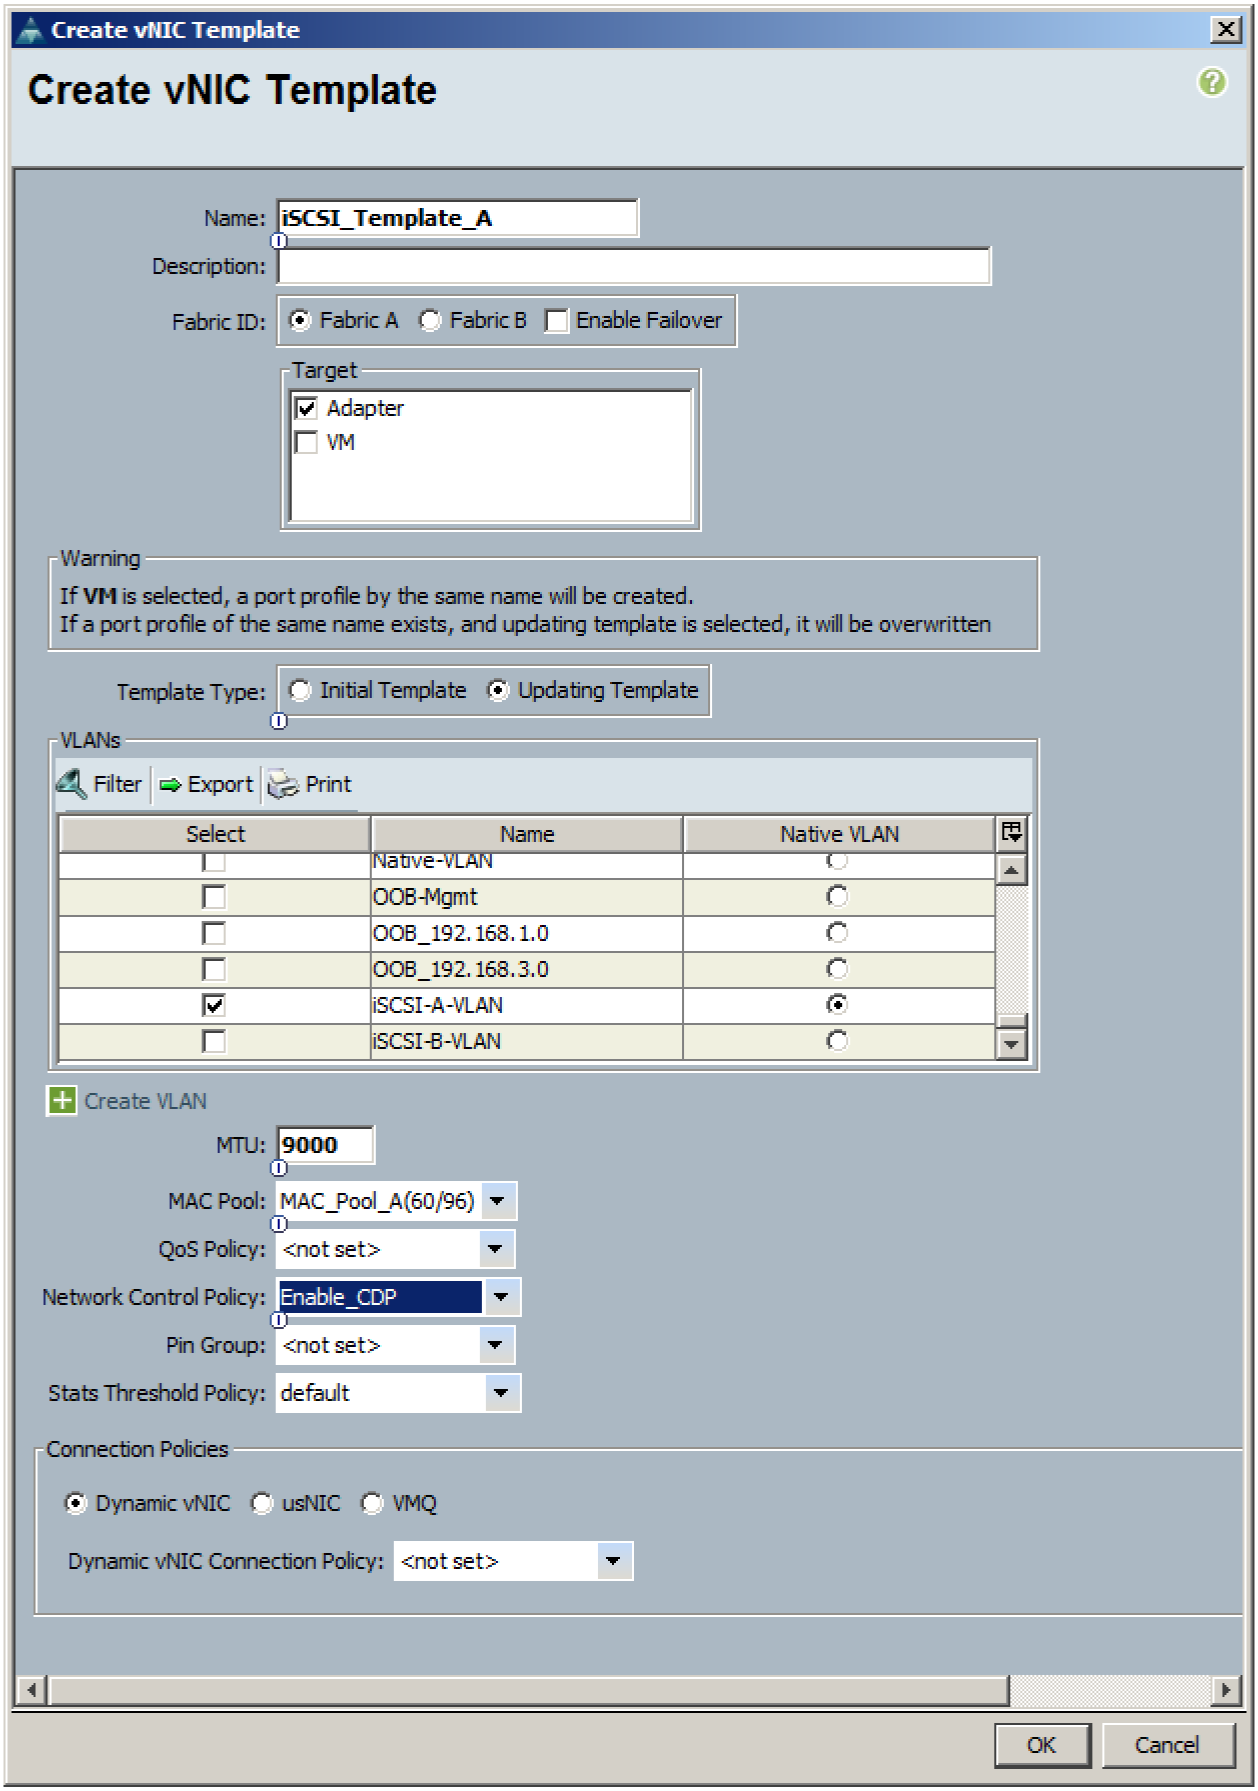

Create vNIC Template

Name: iSCSI_Template_A

Description:

Fabric ID: ● Fabric A ○ Fabric B □ Enable Failover

Target

☑ Adapter

□ VM

Warning

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

Template Type: ○ Initial Template ● Updating Template

VLANs

Select	Name	Native VLAN
□	Native-VLAN	○
□	OOB-Mgmt	○
□	OOB_192.168.1.0	○
□	OOB_192.168.3.0	○
☑	iSCSI-A-VLAN	●
□	iSCSI-B-VLAN	○

MTU: 9000

MAC Pool: MAC_Pool_A(60/96)

QoS Policy: <not set>

Network Control Policy: Enable_CDP

Pin Group: <not set>

Stats Threshold Policy: default

Connection Policies

● Dynamic vNIC ○ usNIC ○ VMQ

Dynamic vNIC Connection Policy: <not set>

16. Select the LAN tab on the left.
17. Select Policies > root.
18. Right-click vNIC Templates.
19. Select Create vNIC Template.
20. Enter `iSCSI_Template_B` as the vNIC template name.
21. Select Fabric B. Do not select the Enable Failover checkbox.
22. Under Target, make sure that the VM checkbox is not selected.
23. Select Updating Template for Template Type.
24. Under VLANs, select `iSCSI-B-VLAN` (912).
25. Set `iSCSI-B-VLAN` as the native VLAN.
26. Under MTU, enter 9000.
27. From the MAC Pool list, select `MAC_Pool_B`.
28. From the Network Control Policy list, select `Enable_CDP`.
29. Click OK to complete creating the vNIC template.
30. Click OK.

Create OOB Mgmt vNICSs

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `OOB-A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.
8. Select Updating Template for Template Type.
9. Under VLANs, select `OOB-Mgmt` VLAN and `Native-VLAN`.
10. Set `Native-VLAN` as the native VLAN.
11. Under MTU, enter 1500. From the MAC Pool list, select `MAC_Pool_A`.

12. From the Network Control Policy list, select `Enable_CDP`.
13. Click OK to complete creating the vNIC template.
14. Click OK.
15. Select the LAN tab on the left.
16. Select Policies > root.
17. Right-click vNIC Templates.
18. Select Create vNIC Template.
19. Enter `OOB-B` as the vNIC template name.
20. Select Fabric B. Do not select the Enable Failover checkbox.
21. Under Target, make sure that the VM checkbox is not selected.
22. Select Updating Template for Template Type.
23. Under VLANs, select `OOB-Mgmt` VLAN and `Native-VLAN`.
24. Set `Native-VLAN` as the native VLAN.
25. Under MTU, enter 1500. From the MAC Pool list, select `MAC_Pool_B`.
26. From the Network Control Policy list, select `Enable_CDP`.
27. Click OK to complete creating the vNIC template.
28. Click OK.

Create Infrastructure NFS vNICSs

Note: These vNICs will only be utilized on ESXi servers hosting Infrastructure services

1. Select the LAN tab on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Infra_NFS_A` as the vNIC template name.
6. Leave Fabric A selected. Do not select the Enable Failover checkbox.
7. Under Target, make sure that the VM checkbox is not selected.

8. Select Updating Template for Template Type.
9. Under VLANs, select `Infra-NFS` (3270) VLAN and `default`.
10. Set `default` as the native VLAN.
11. Under MTU, enter 9000. From the MAC Pool list, select `MAC_Pool_A`.
12. From the Network Control Policy list, select `Enable_CDP`.
13. Click OK to complete creating the vNIC template.
14. Click OK.
15. Select the LAN tab on the left.
16. Select Policies > root.
17. Right-click vNIC Templates.
18. Select Create vNIC Template.
19. Enter `Infra_NFS_B` as the vNIC template name.
20. Select Fabric B. Do not select the Enable Failover checkbox.
21. Under Target, make sure that the VM checkbox is not selected.
22. Select Updating Template for Template Type.
23. Under VLANs, select `Infra-NFS` VLAN and `default` VLAN.
24. Set `default` as the native VLAN.
25. Under MTU, enter 9000. From the MAC Pool list, select `MAC_Pool_B`.
26. From the Network Control Policy list, select `Enable_CDP`.
27. Click OK to complete creating the vNIC template.
28. Click OK.

Note: At the end of the vNIC template creation, there should be 8 vNIC templates available in the UCS Manager as shown in Figure 3 below.

Figure 3 vNIC Templates

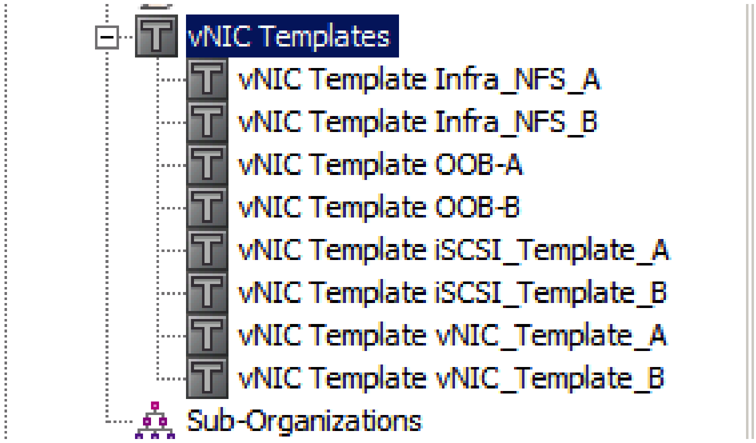

vNIC Templates
- vNIC Template Infra_NFS_A
- vNIC Template Infra_NFS_B
- vNIC Template OOB-A
- vNIC Template OOB-B
- vNIC Template iSCSI_Template_A
- vNIC Template iSCSI_Template_B
- vNIC Template vNIC_Template_A
- vNIC Template vNIC_Template_B

Sub-Organizations

Create Boot Policies

This procedure applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (`iscsi lif01a` and `iscsi lif01b`) and two iSCSI LIFs are on cluster node 2 (`iscsi lif02a` and `iscsi lif02b`). One boot policy is configured in this procedure. This policy configures the primary target to be iscsi_lif01a.

To create boot policies for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.
5. Enter `Boot-Fabric-A` as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Expand the Local Devices drop-down menu and select `Add CD-ROM`.
9. Expand the `iSCSI vNICs` section and select `Add iSCSI Boot`.
10. In the `Add iSCSI Boot` dialog box, enter `iSCSI-A-vNIC`.
11. Click OK.
12. Select `Add iSCSI Boot`.
13. In the `Add iSCSI Boot` dialog box, enter `iSCSI-B-vNIC`.
14. Click OK.

15. Click OK to save the boot policy. Click OK to close the Boot Policy window.

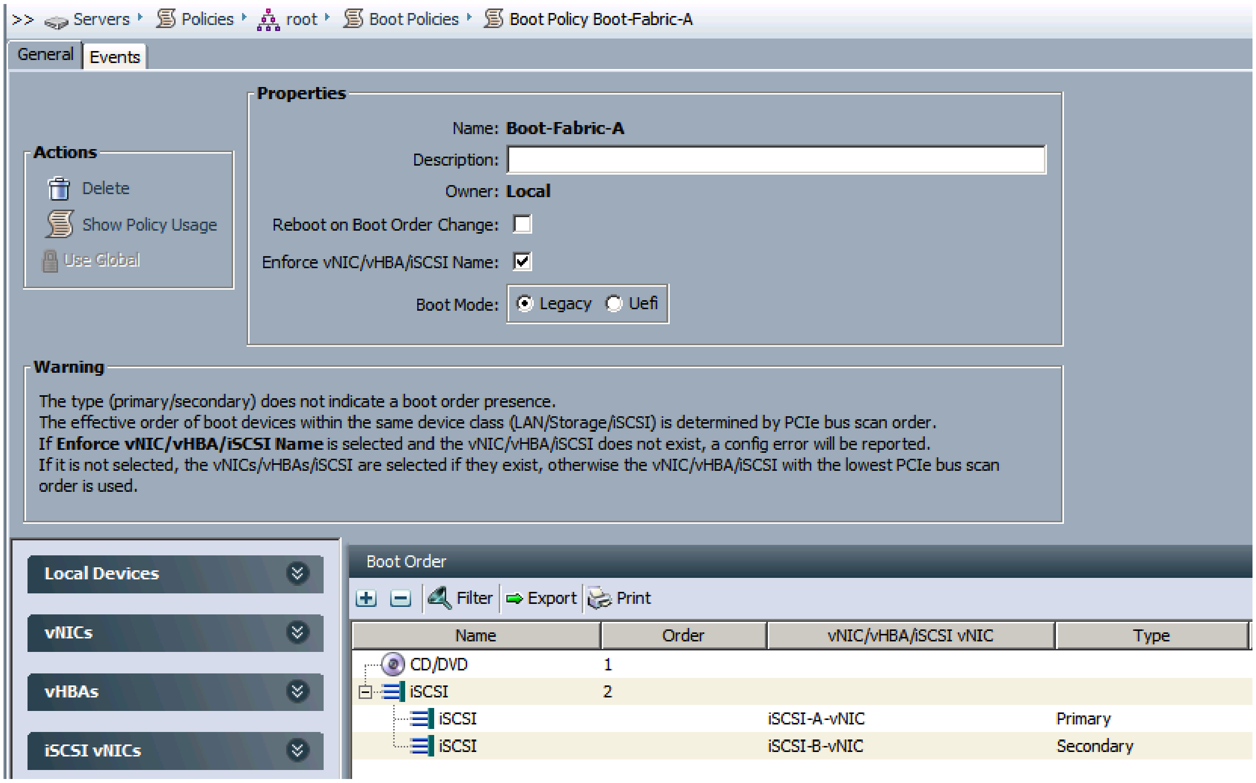

>> Servers ▸ Policies ▸ root ▸ Boot Policies ▸ Boot Policy Boot-Fabric-A
General
Events
Actions
Delete
Show Policy Usage
Use Global

Properties

Name: **Boot-Fabric-A**

Description:

Owner: **Local**

Reboot on Boot Order Change: □

Enforce vNIC/vHBA/iSCSI Name: ☑

Boot Mode: ● Legacy ○ Uefi

Warning

The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Name	Order	vNIC/vHBA/iSCSI vNIC	Type
CD/DVD	1		
iSCSI	2		
iSCSI		iSCSI-A-vNIC	Primary
iSCSI		iSCSI-B-vNIC	Secondary

Create Service Profile Template

In this procedure, one service profile template for Infrastructure ESXi hosts is created for fabric A boot.

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
6. Enter VM-Host-Infra-Fabric-A as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
7. Select the Updating Template option.
8. Under UUID, select UUID_Pool as the UUID pool.
9. Click Next.

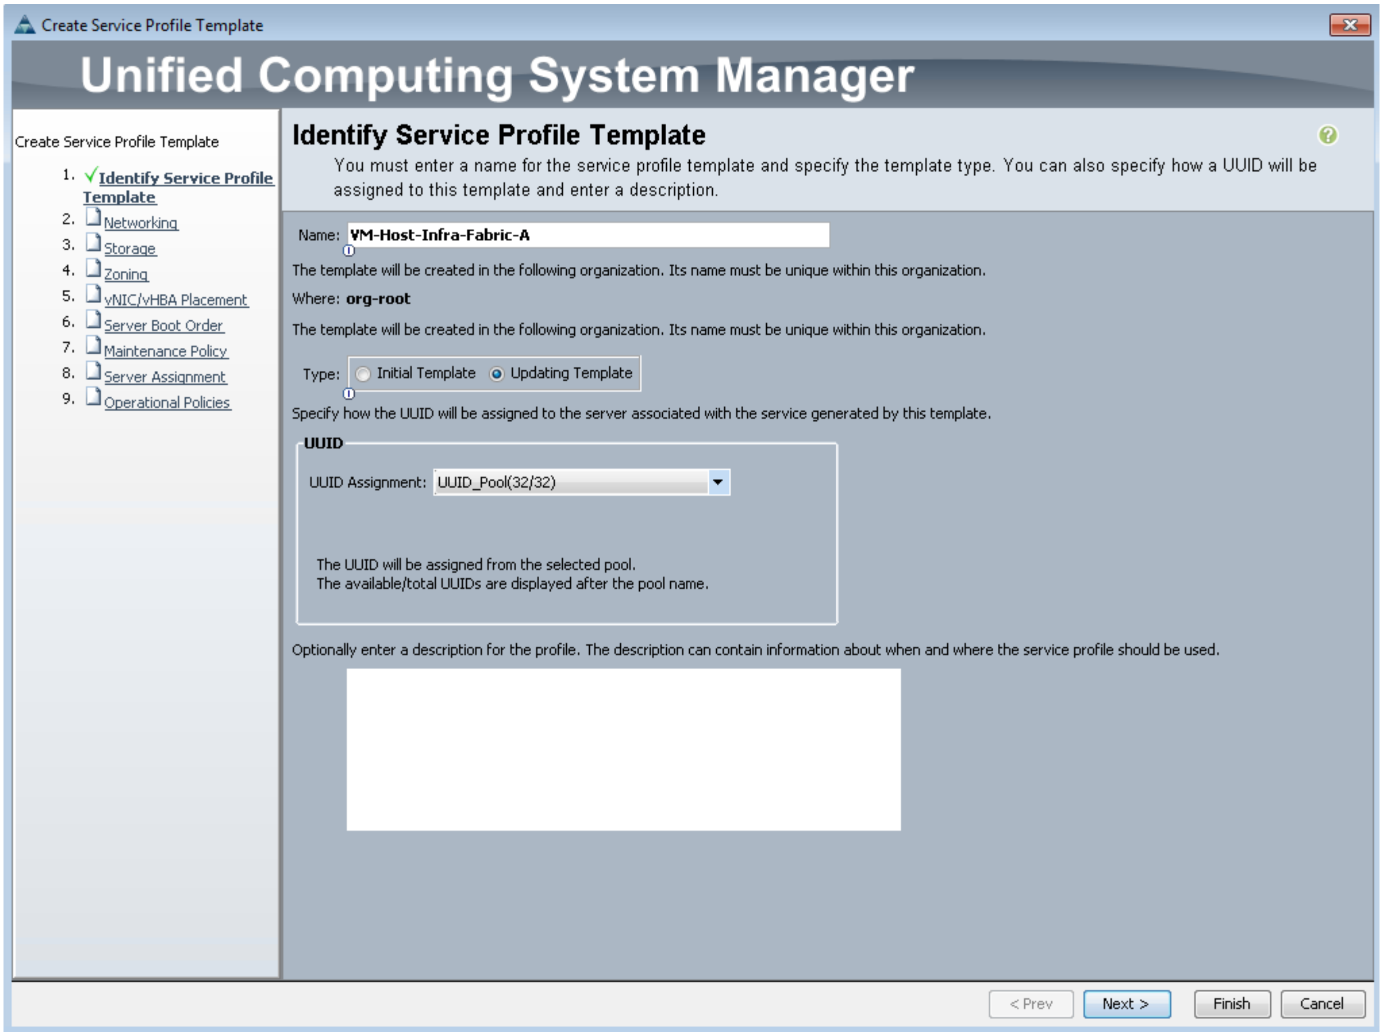

Identify Service Profile Template

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

Name: VM-Host-Infra-Fabric-A

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: UUID_Pool(32/32)

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

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

Note: To configure the networking options, 8 vNIC interfaces will be added for Infrastructure ESXi hosts:

Configure Networking options:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Expert option to configure the LAN connectivity.
3. Click the upper Add button to add a vNIC to the template.
4. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.
5. Select the Use vNIC Template checkbox.
6. In the vNIC Template list, select vNIC_Template_A.
7. In the Adapter Policy list, select VMWare.
8. Click OK to add this vNIC to the template.

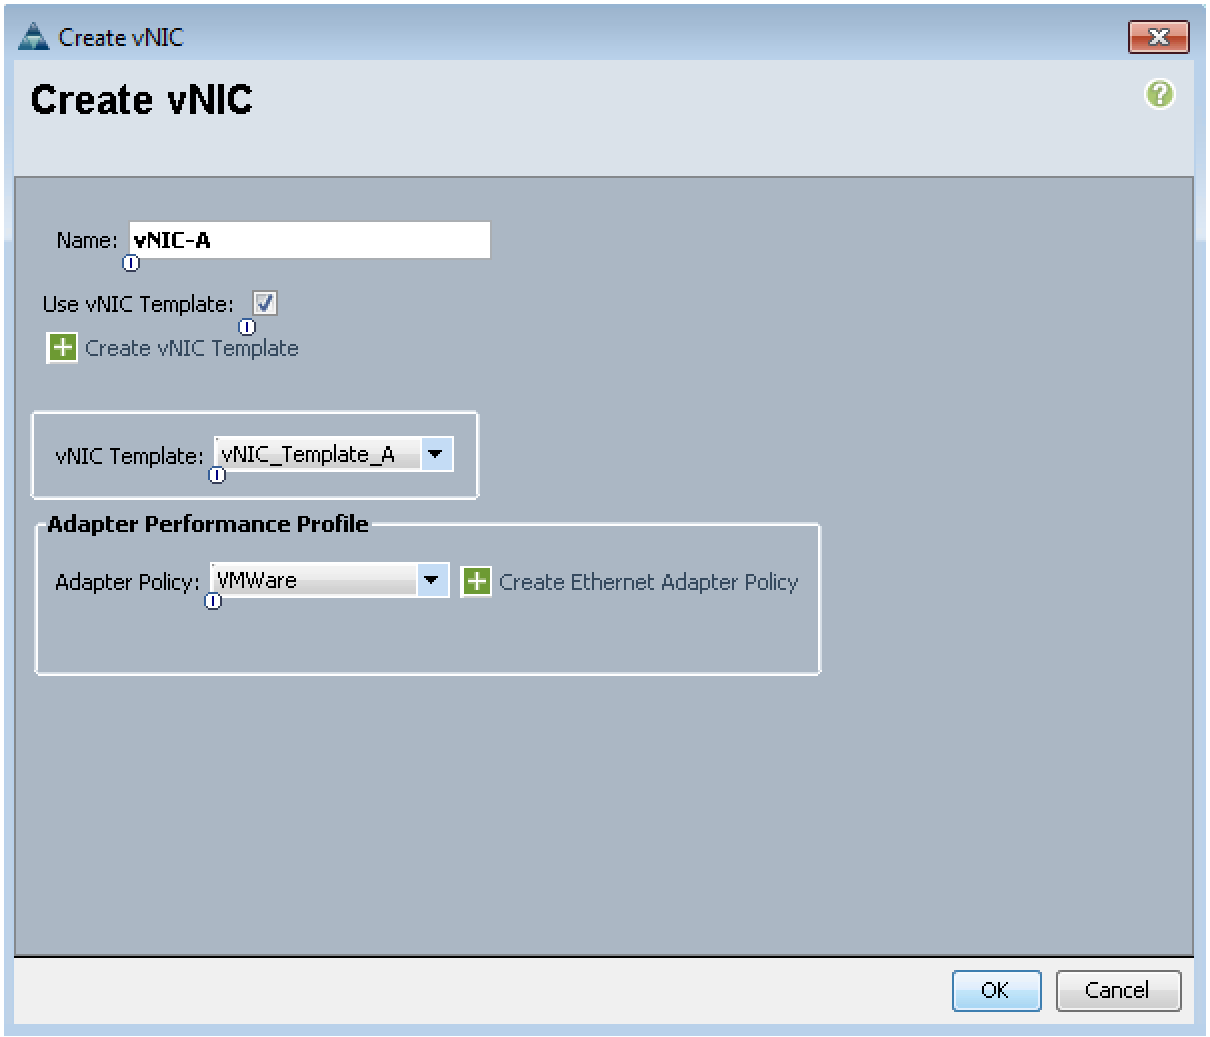

Create vNIC

Name: vNIC-A

Use vNIC Template: ☑

vNIC Template: vNIC_Template_A

Adapter Performance Profile

Adapter Policy: VMWare

9. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

10. In the Create vNIC box, enter vNIC-B as the name of the vNIC.

11. Select the Use vNIC Template checkbox.

12. In the vNIC Template list, select vNIC_Template_B.

13. In the Adapter Policy list, select VMWare.

14. Click OK to add the vNIC to the template.

15. Click the upper Add button to add a vNIC to the template.

16. In the Create vNIC dialog box, enter `iSCSI-A-vNIC` as the name of the vNIC.

17. Select the Use vNIC Template checkbox.

18. In the vNIC Template list, select iSCSI_Template_A.

19. In the Adapter Policy list, select VMWare.

20. Click OK to add this vNIC to the template.

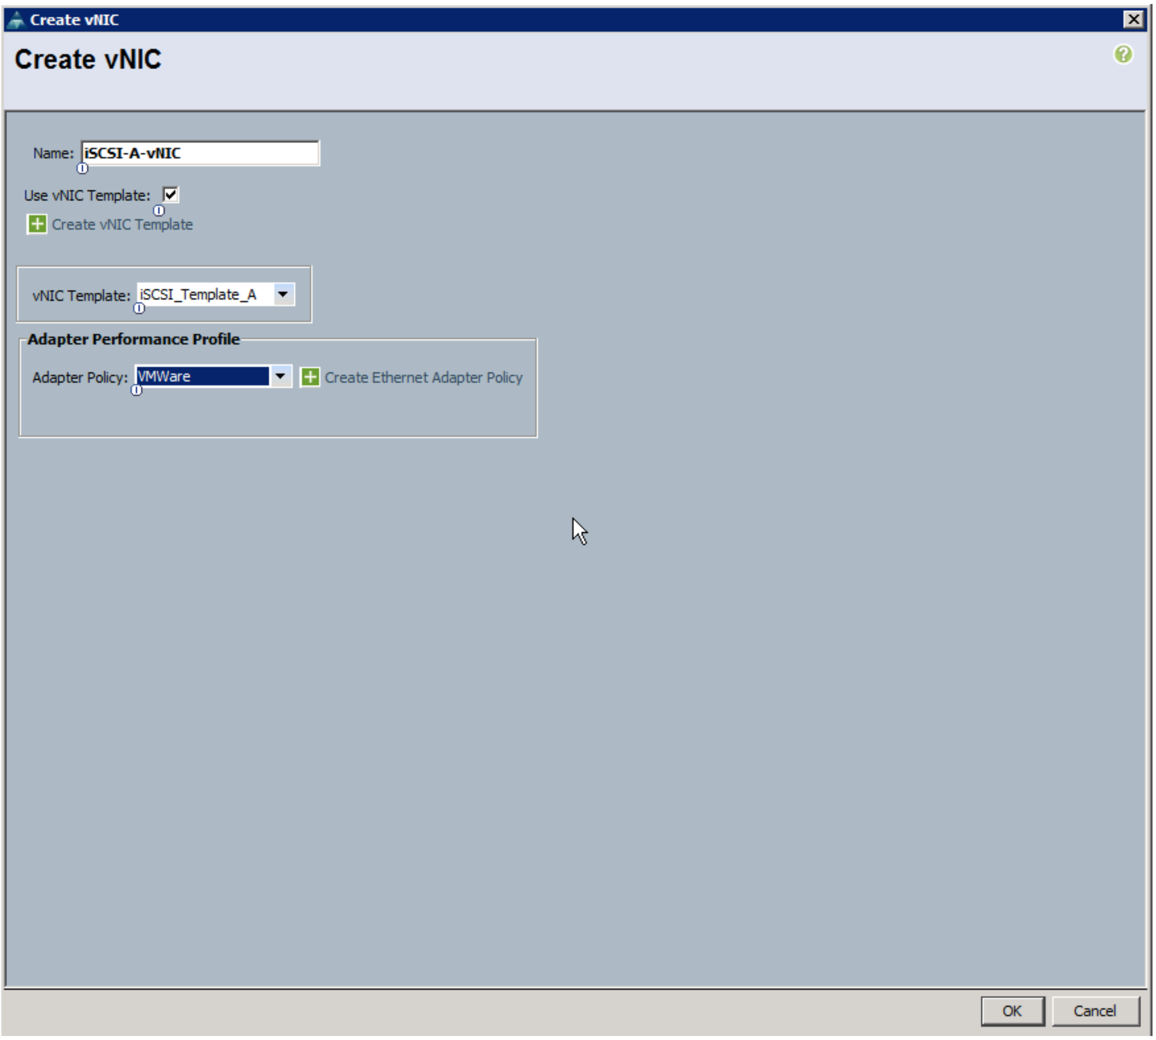

Create vNIC

Name: iSCSI-A-vNIC

Use vNIC Template: ☑

Create vNIC Template

vNIC Template: iSCSI_Template_A

Adapter Performance Profile

Adapter Policy: VMWare

Create Ethernet Adapter Policy

OK Cancel

21. Click the upper Add button to add a vNIC to the template.

22. In the Create vNIC dialog box, enter iSCSI-B-vNIC as the name of the vNIC.

23. Select the Use vNIC Template checkbox.

24. In the vNIC Template list, select iSCSI_Template_B.

25. In the Adapter Policy list, select VMWare.

26. Click OK to add this vNIC to the template.

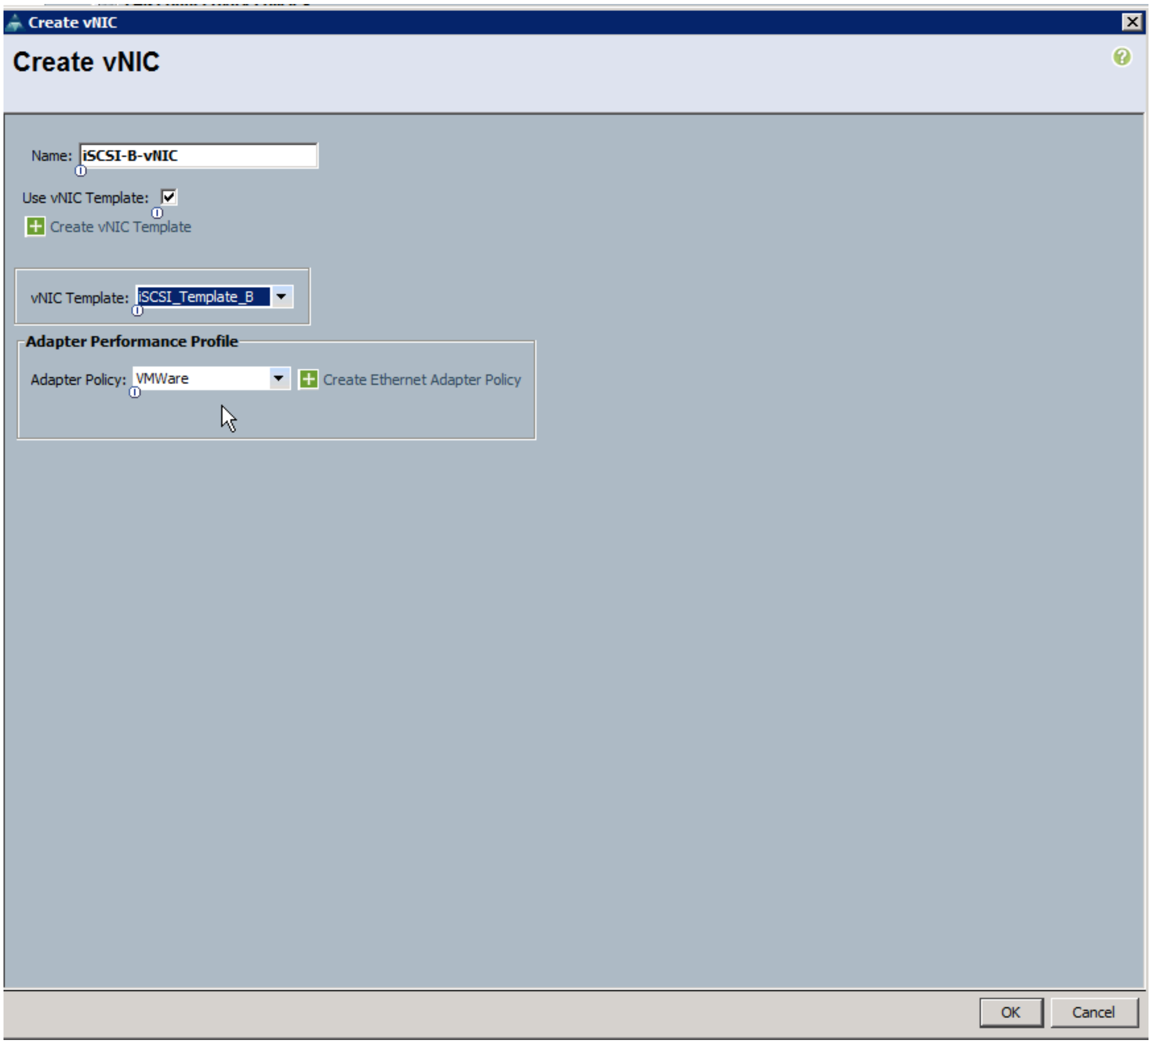

Create vNIC

Name: iSCSI-B-vNIC

Use vNIC Template: ☑

vNIC Template: iSCSI_Template_B

Adapter Performance Profile

Adapter Policy: VMWare

27. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

28. In the Create vNIC box, enter OOB-A as the name of the vNIC.

29. Select the Use vNIC Template checkbox.

30. In the vNIC Template list, select OOB-A.

31. In the Adapter Policy list, select VMWare.

32. Click OK to add the vNIC to the template.

33. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

34. In the Create vNIC box, enter OOB-B as the name of the vNIC.

35. Select the Use vNIC Template checkbox.

36. In the vNIC Template list, select OOB-B.

37. In the Adapter Policy list, select VMWare.

38. Click OK to add the vNIC to the template.

Note: The next two vNIC interfaces are only needed for Infrastructure ESXi Hosts. These interfaces enable NFS access using NFS specific vSwitch and static EPG mapping. ESXi servers not hosting infrastructure VMs, use VDS for NFS access to application specific SVMs.

39. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

40. In the Create vNIC box, enter NFS-A as the name of the vNIC.

41. Select the Use vNIC Template checkbox.

42. In the vNIC Template list, select Infra_NFS_A.

43. In the Adapter Policy list, select VMWare.

44. Click OK to add the vNIC to the template.

45. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

46. In the Create vNIC box, enter NFS-B as the name of the vNIC.

47. Select the Use vNIC Template checkbox.

48. In the vNIC Template list, select Infra_NFS_B.

49. In the Adapter Policy list, select VMWare.

50. Click OK to add the vNIC to the template.

51. Verify 8 vNIC interfaces are present.

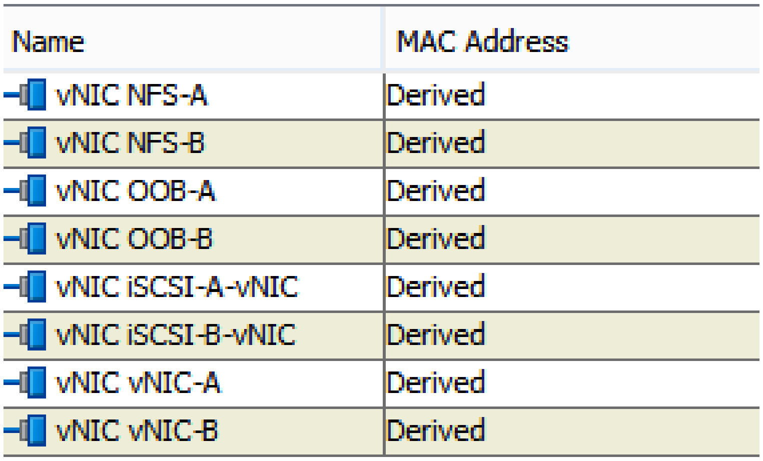

Name	MAC Address
vNIC NFS-A	Derived
vNIC NFS-B	Derived
vNIC OOB-A	Derived
vNIC OOB-B	Derived
vNIC iSCSI-A-vNIC	Derived
vNIC iSCSI-B-vNIC	Derived
vNIC vNIC-A	Derived
vNIC vNIC-B	Derived

52. Expand the iSCSI vNICs section (if not already expanded).

53. Select `iqn-pool` under Initiator Name Assignment.

54. Click the **lower** Add button in the iSCSI vNIC section to define a vNIC.

55. Enter `iSCSI-A-vNIC` as the name of the vNIC.

56. Select `iSCSI-A-vNIC` for Overlay vNIC.

57. Set the iSCSI Adapter Policy to `default`.

58. Set the VLAN to `iSCSI-A-VLAN`.

59. Leave the MAC Address set to `None`.

60. Click OK.

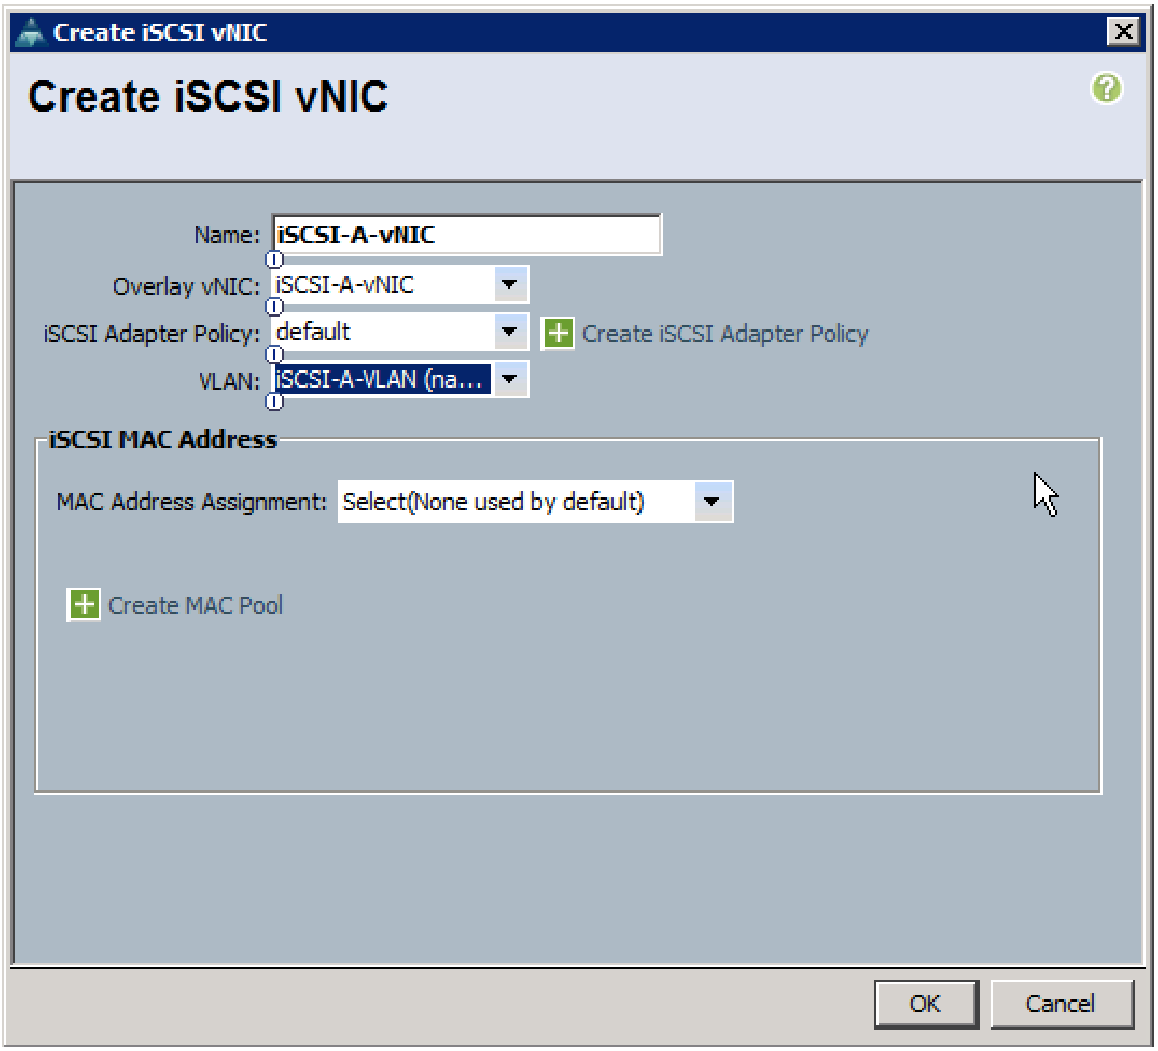

Create iSCSI vNIC

Name: iSCSI-A-vNIC

Overlay vNIC: iSCSI-A-vNIC

iSCSI Adapter Policy: default

VLAN: iSCSI-A-VLAN (na...

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

61. Click the lower Add button in the iSCSI vNIC section to define a vNIC.

62. Enter `iSCSI-B-vNIC` as the name of the vNIC.

63. Set the Overlay vNIC to `iSCSI-B-vNIC`

64. Set the iSCSI Adapter Policy to `default`.

65. Set the VLAN to `iSCSI-B-VLAN`

66. Leave the MAC Address set to `None`.

67. Click OK.

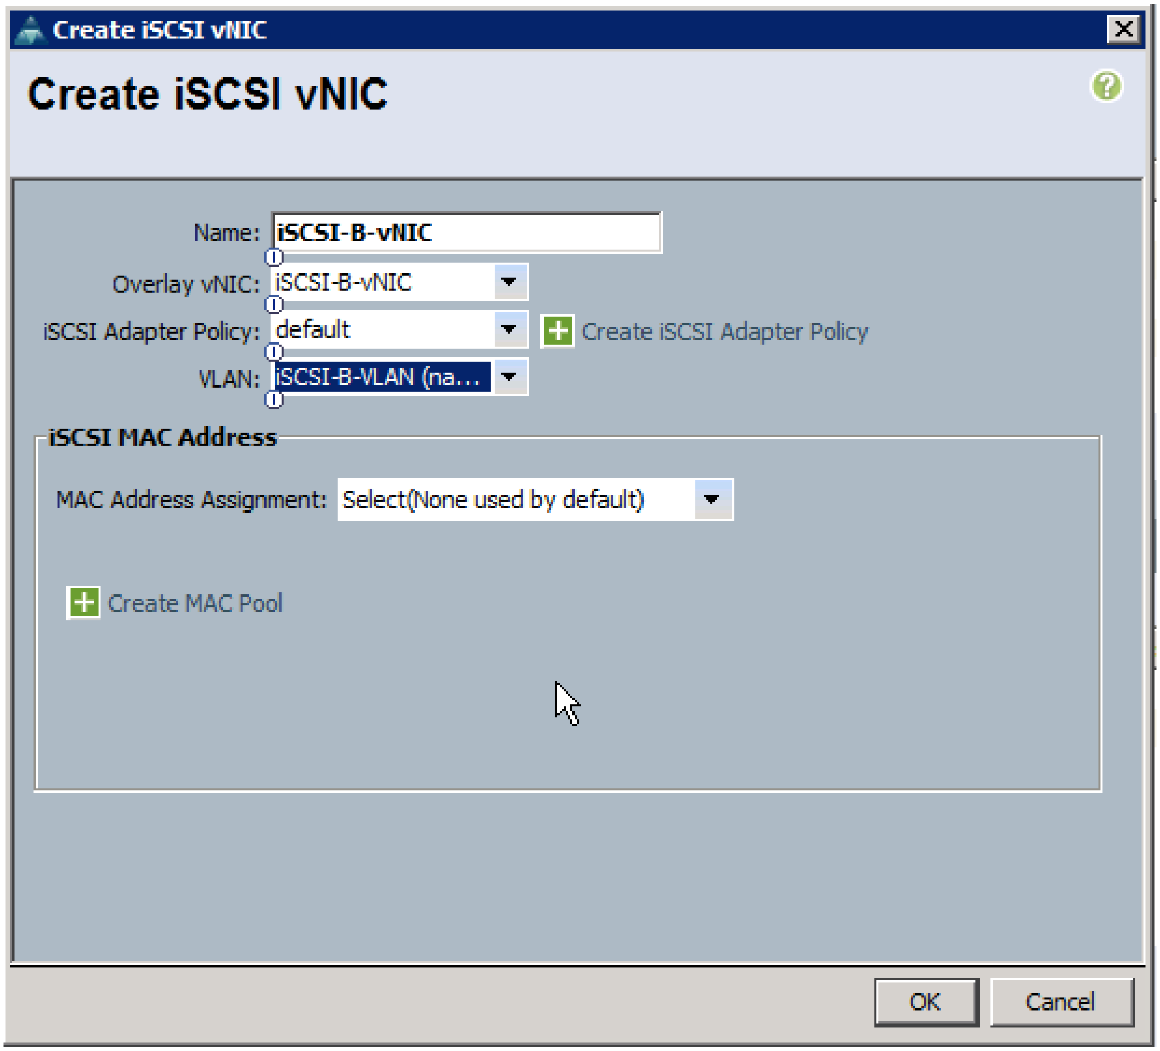

Create iSCSI vNIC

Name: iSCSI-B-vNIC

Overlay vNIC: iSCSI-B-vNIC

iSCSI Adapter Policy: default

VLAN: iSCSI-B-VLAN (na…

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

68. Click OK.

69. Review the table in the Networking page to make sure that all vNICs were created.

70. Click Next.

Configure Storage options:

1. Select a local disk configuration policy:

If the server in question has local disks, select default in the Local Storage list.

If the server in question does not have local disks, select `iSCSI-Boot`.

2. Select the `No vHBAs` option for the How would you like to configure SAN connectivity? field.

3. Click Next.

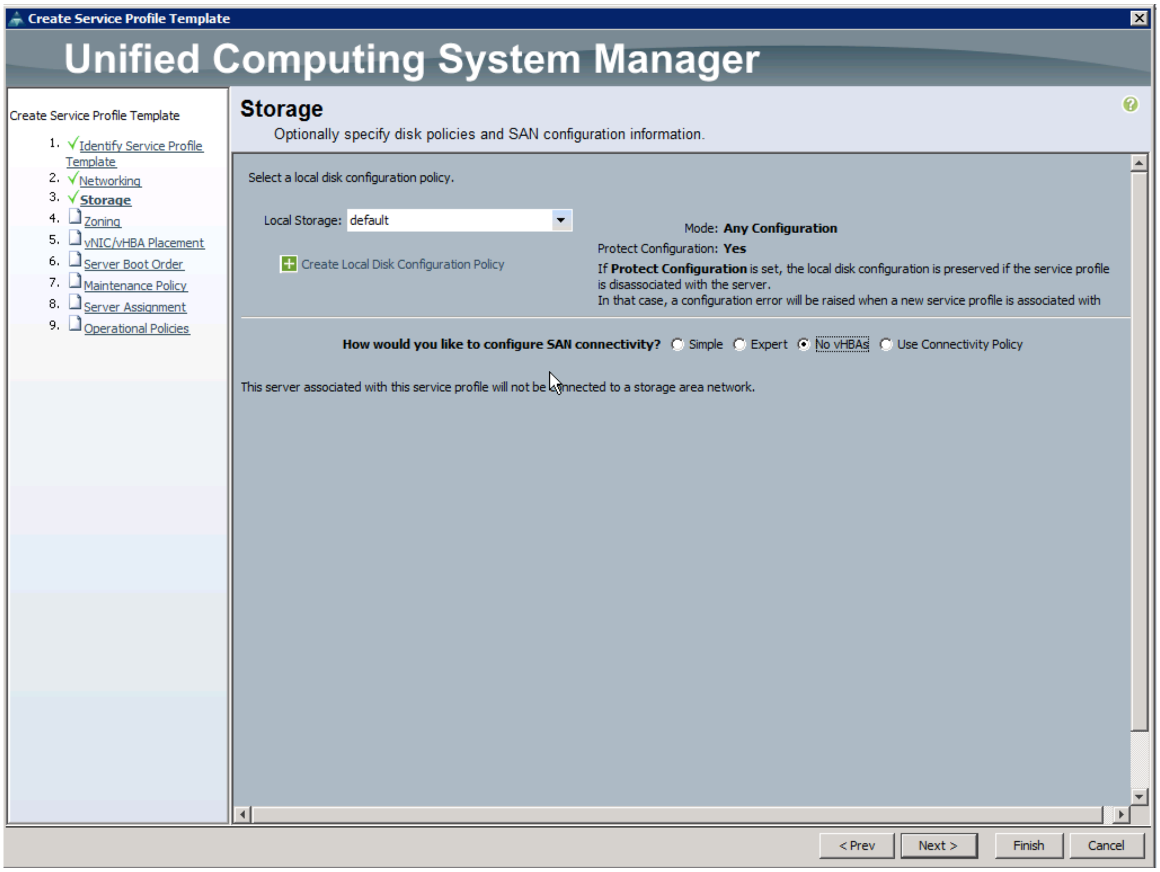
Create Service Profile Template
Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

Storage

Optionally specify disk policies and SAN configuration information.

Select a local disk configuration policy.

Local Storage: default

Create Local Disk Configuration Policy

Mode: Any Configuration

Protect Configuration: Yes

If **Protect Configuration** is set, the local disk configuration is preserved if the service profile is disassociated with the server.
In that case, a configuration error will be raised when a new service profile is associated with

How would you like to configure SAN connectivity? ○ Simple ○ Expert ● No vHBAs ○ Use Connectivity Policy

This server associated with this service profile will not be connected to a storage area network.

Configure Zoning Options

1. Set no Zoning options and click Next.

Configure vNIC/HBA Placement

1. Set the vNIC/vHBA placement options (**Error! Reference source not found.**).

2. In the Select Placement list, select the VM-Host-Infra placement policy.

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

— `vNIC-A`

— `vNIC-B`

— `iSCSI-vNIC-A`

— `iSCSI-vNIC-B`

— `OOB-A`

— `OOB-B`

- NFS-A
- NFS-B

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

5. Click Next.

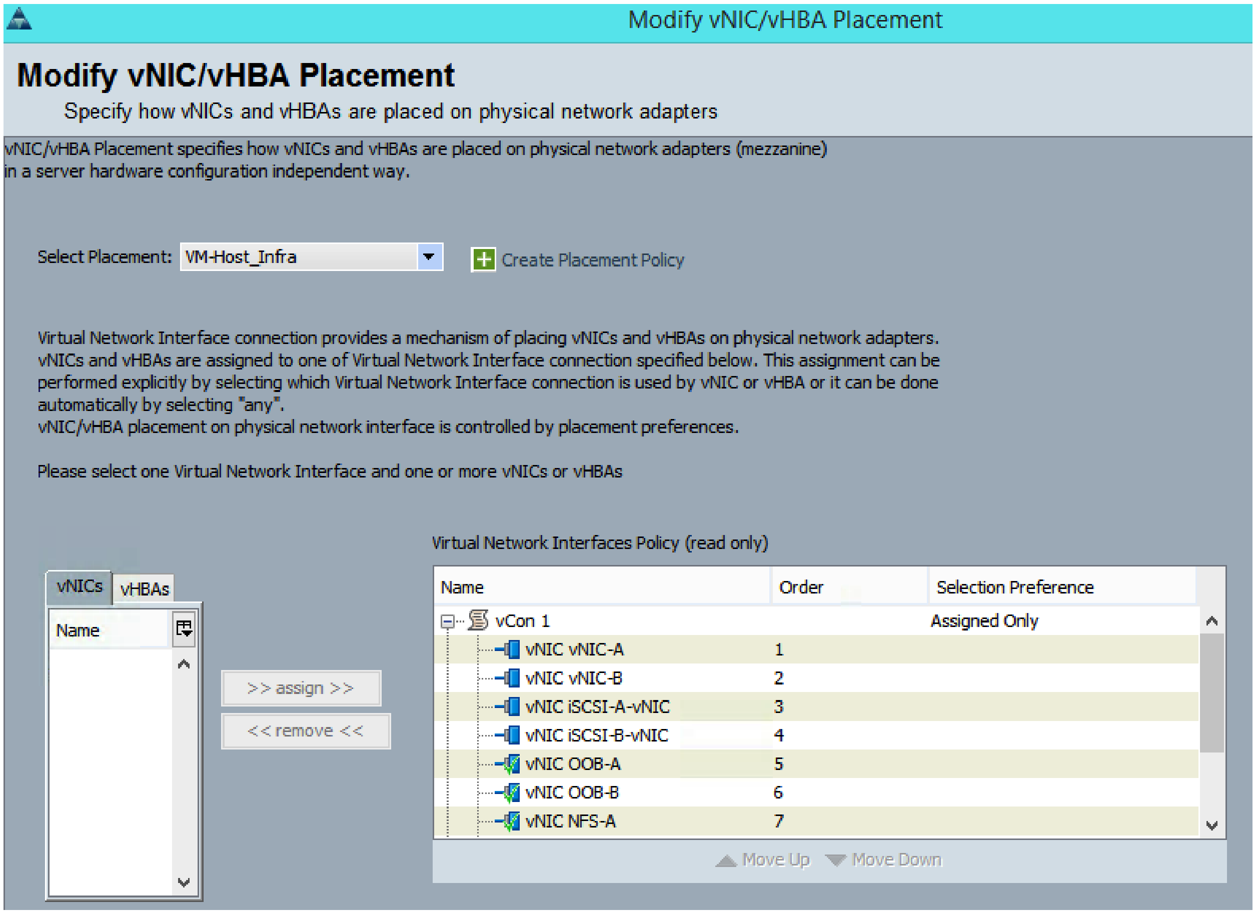

Modify vNIC/vHBA Placement

Specify how vNICs and vHBAs are placed on physical network adapters

vNIC/vHBA Placement specifies how vNICs and vHBAs are placed on physical network adapters (mezzanine) in a server hardware configuration independent way.

Select Placement: VM-Host_Infra

Virtual Network Interface connection provides a mechanism of placing vNICs and vHBAs on physical network adapters. vNICs and vHBAs are assigned to one of Virtual Network Interface connection specified below. This assignment can be performed explicitly by selecting which Virtual Network Interface connection is used by vNIC or vHBA or it can be done automatically by selecting "any".
vNIC/vHBA placement on physical network interface is controlled by placement preferences.

Please select one Virtual Network Interface and one or more vNICs or vHBAs

Virtual Network Interfaces Policy (read only)

Name	Order	Selection Preference
vCon 1		Assigned Only
vNIC vNIC-A	1	
vNIC vNIC-B	2	
vNIC iSCSI-A-vNIC	3	
vNIC iSCSI-B-vNIC	4	
vNIC OOB-A	5	
vNIC OOB-B	6	
vNIC NFS-A	7	

Configure Server Boot Order

1. Select Boot-Fabric-A for Boot Policy.

2. In the Boot Order pane, select iSCSI-A-vNIC.

3. Click the Set iSCSI Boot Parameters button.

4. Leave the Set iSCSI Boot Parameters dialog box <not set> to use the single Service Profile Initiator Name defined in the previous steps.

5. Set iSCSI_IP_Pool_A as the Initiator IP address Policy.

6. Keep the iSCSI Static Target Interface button selected and click the + button.

7. Log in to the storage cluster management interface and run the following command:

```
iscsi nodename
```

8. Note or copy the iSCSI target name for `Infra-SVM`.

9. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from `Infra-SVM`.

10. Enter the IP address of `iSCSI_lif02a` for the IPv4 Address field.

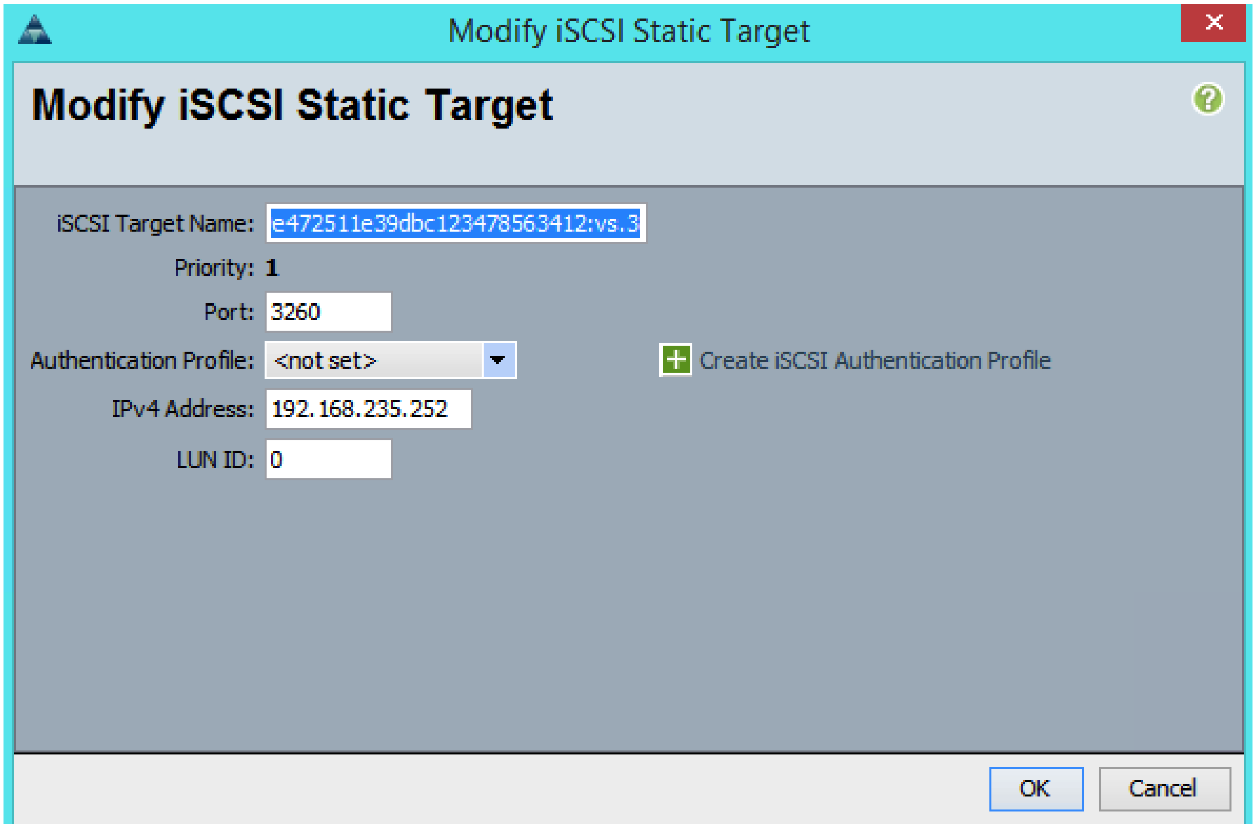

Modify iSCSI Static Target

Modify iSCSI Static Target

iSCSI Target Name: e472511e39dbc123478563412:vs.3

Priority: 1

Port: 3260

Authentication Profile: <not set>

IPv4 Address: 192.168.235.252

LUN ID: 0

Create iSCSI Authentication Profile
OK
Cancel

11. Click OK to add the iSCSI static target.

12. Keep the iSCSI Static Target Interface option selected and click the ➕ button.

13. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field.

14. Enter the IP address of `iscsi_lif01a` in the IPv4 Address field.

15. Click OK.

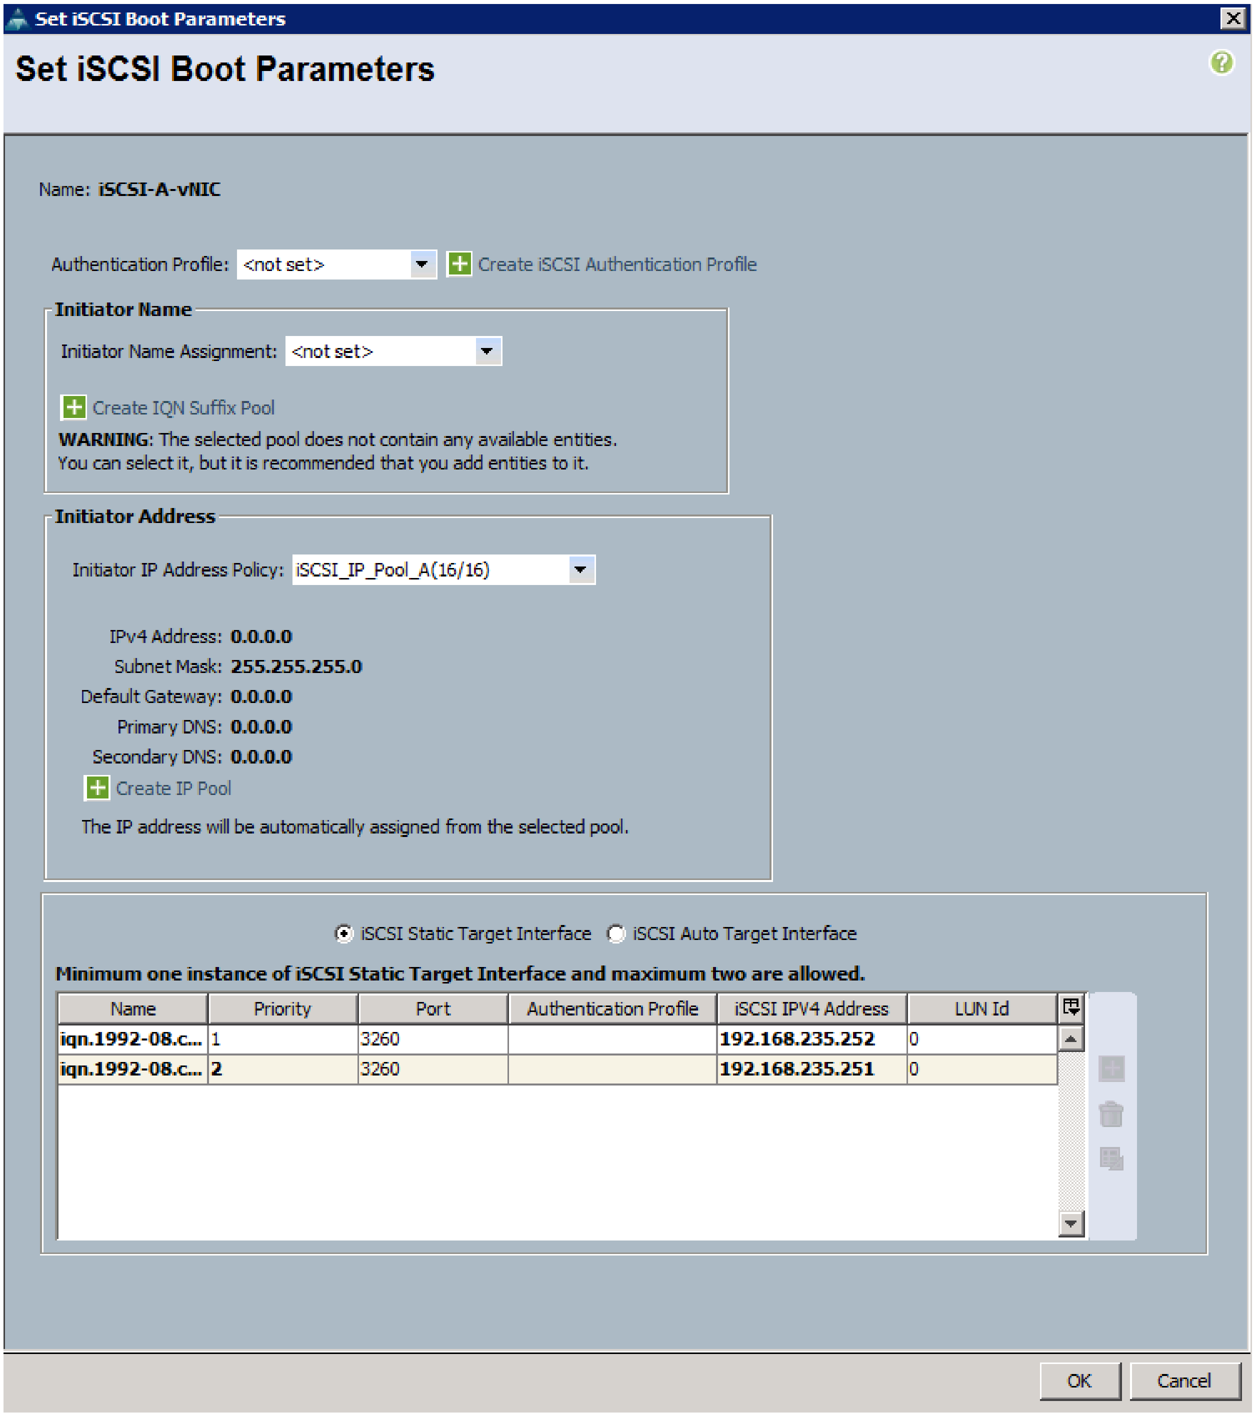

Set iSCSI Boot Parameters

Name: **iSCSI-A-vNIC**

Authentication Profile: <not set>

Initiator Name

Initiator Name Assignment: <not set>

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_A(16/16)

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

The IP address will be automatically assigned from the selected pool.

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPV4 Address	LUN Id
iqn.1992-08.c...	1	3260		**192.168.235.252**	0
iqn.1992-08.c...	**2**	3260		**192.168.235.251**	0

Set iSCSI Boot Parameters
Create iSCSI Authentication Profile
Create IQN Suffix Pool
Create IP Pool
OK
Cancel

16. Click OK.

17. In the Boot Order pane, select `iSCSI-vNIC-B`.

18. Click the Set iSCSI Boot Parameters button.

19. In the Set iSCSI Boot Parameters dialog box, set the leave the Initiator Name Assignment to <not set>.

20. In the Set iSCSI Boot Parameters dialog box, set the initiator IP address policy to iSCSI_IP_Pool_B.

21. Keep the iSCSI Static Target Interface option selected and click the + button.

22. In the Create iSCSI Static Target window, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field (same target name as above).

23. Enter the IP address of iscsi_lif02b in the IPv4 address field.

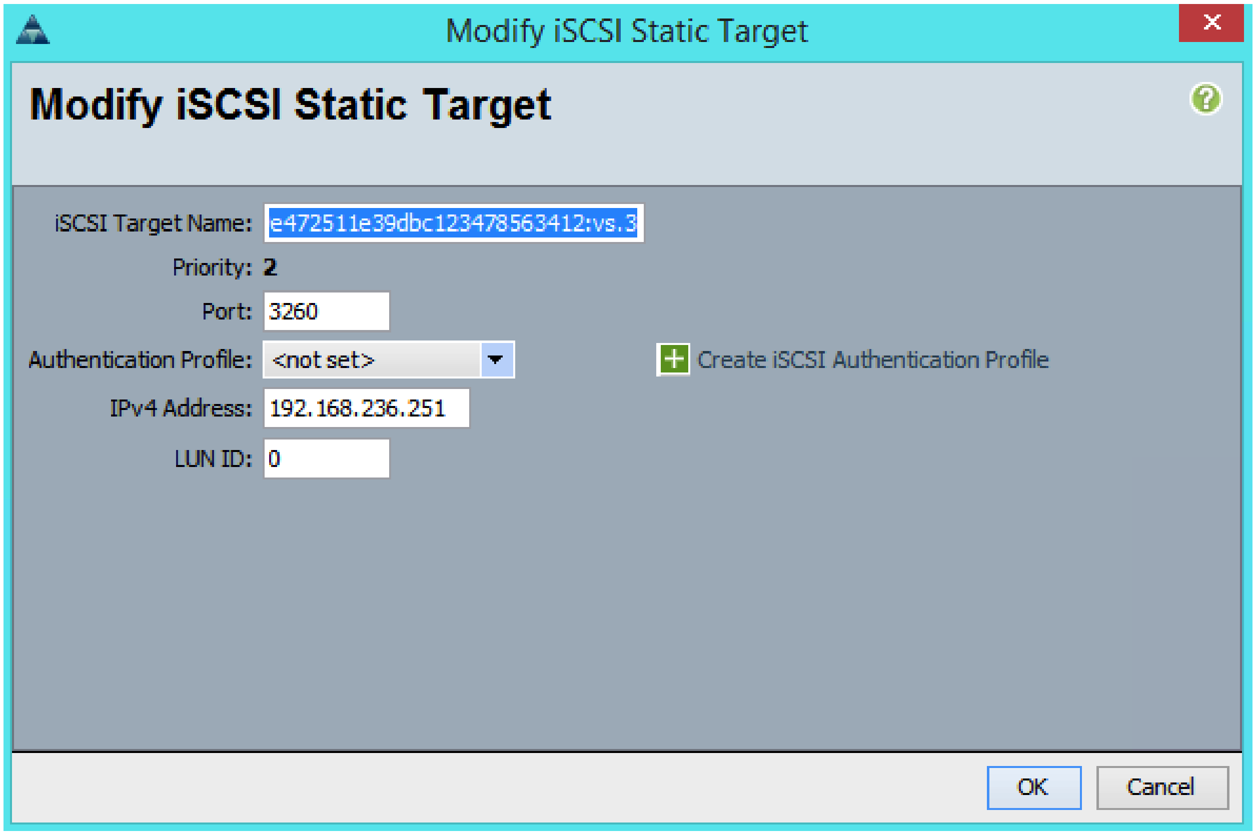
Modify iSCSI Static Target
Create iSCSI Authentication Profile
OK
Cancel

Modify iSCSI Static Target

iSCSI Target Name: e472511e39dbc123478563412:vs.3

Priority: 2

Port: 3260

Authentication Profile: <not set>

IPv4 Address: 192.168.236.251

LUN ID: 0

24. Click OK to add the iSCSI static target.

25. Keep the iSCSI Static Target Interface option selected and click the + button.

26. In the Create iSCSI Static Target dialog box, paste the iSCSI target node name from Infra-SVM into the iSCSI Target Name field.

27. Enter the IP address of iscsi_lif01b in the IPv4 Address field.

28. Click OK.

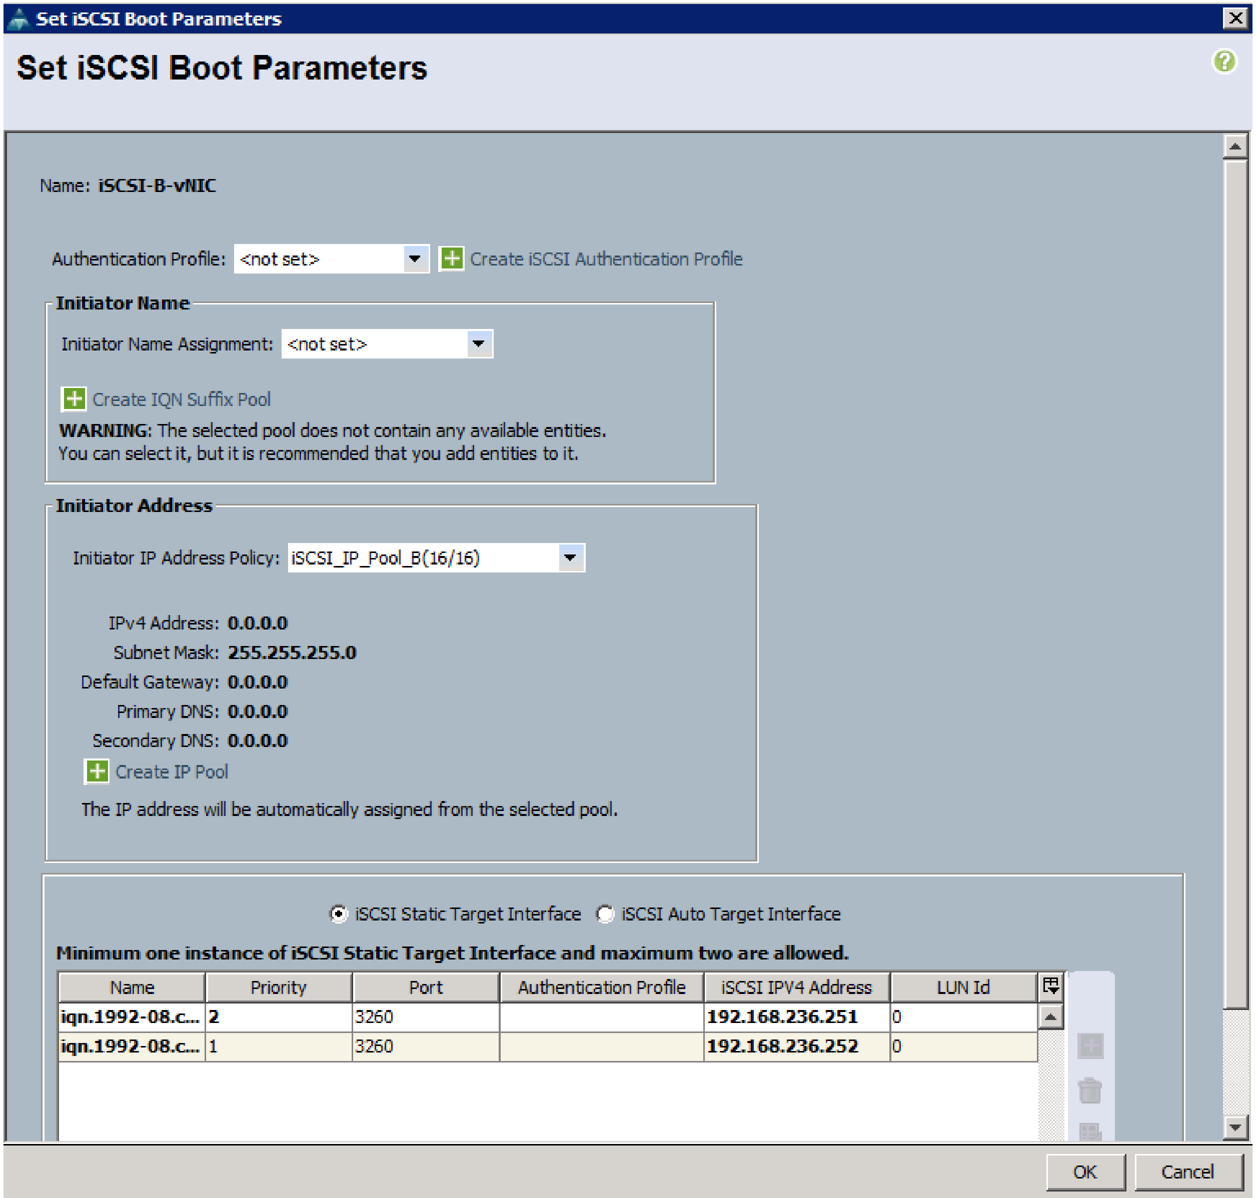

Set iSCSI Boot Parameters

Name: **iSCSI-B-vNIC**

Authentication Profile: <not set> Create iSCSI Authentication Profile

Initiator Name

Initiator Name Assignment: <not set>

Create IQN Suffix Pool

WARNING: The selected pool does not contain any available entities.
You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_B(16/16)

IPv4 Address: **0.0.0.0**
Subnet Mask: **255.255.255.0**
Default Gateway: **0.0.0.0**
Primary DNS: **0.0.0.0**
Secondary DNS: **0.0.0.0**

Create IP Pool

The IP address will be automatically assigned from the selected pool.

● iSCSI Static Target Interface ○ iSCSI Auto Target Interface

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

Name	Priority	Port	Authentication Profile	iSCSI IPV4 Address	LUN Id
iqn.1992-08.c...	2	3260		192.168.236.251	0
iqn.1992-08.c...	1	3260		192.168.236.252	0

29. Click OK.

30. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

31. Click Next to continue to the next section.

Configure Maintenance Policy

1. Select the default Maintenance Policy.

2. Click Next.

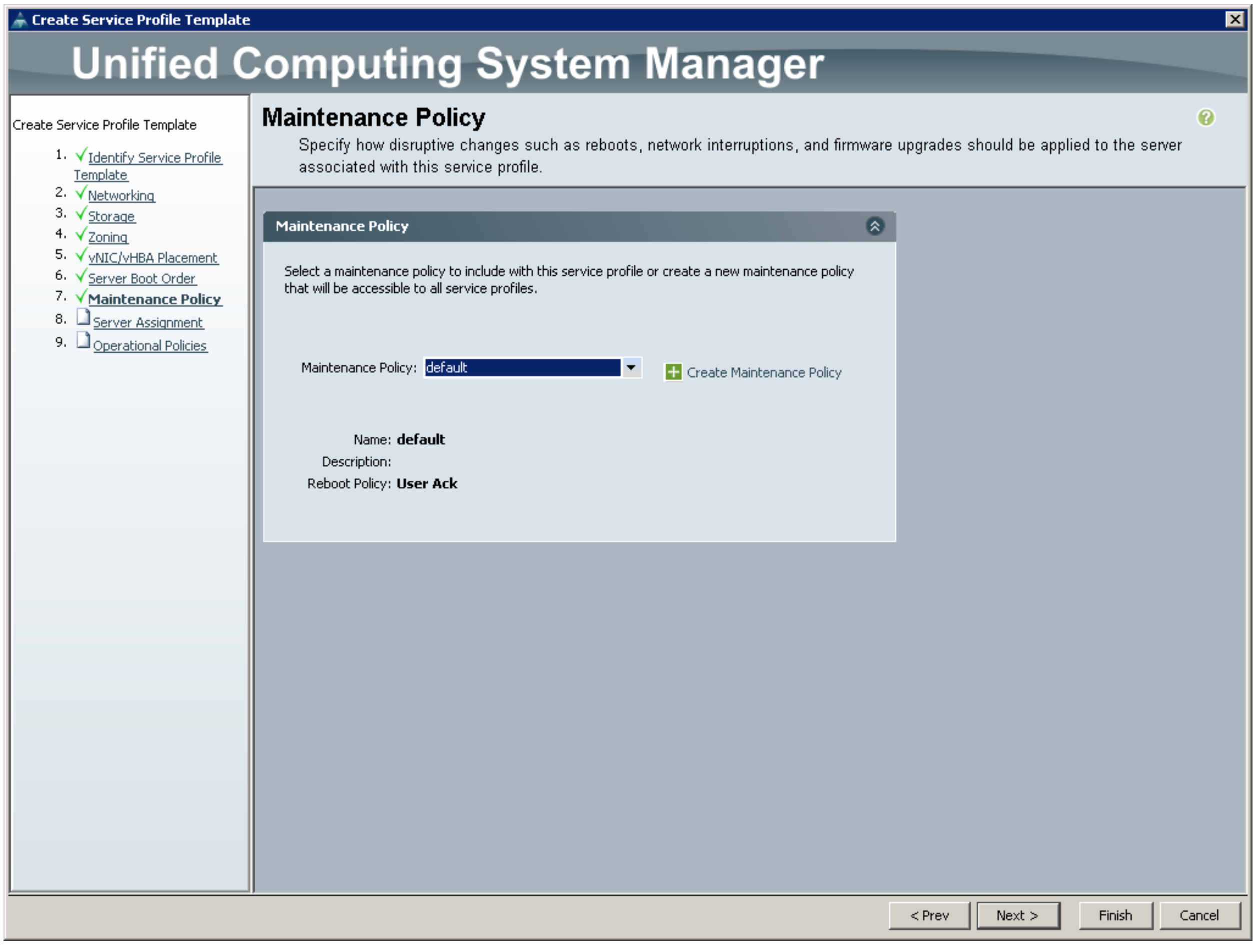

Maintenance Policy

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: default

Name: **default**

Description:

Reboot Policy: **User Ack**

Configure Server Assignment

1. In the Pool Assignment list, select `Infra_Pool`.
2. Optional: Select a Server Pool Qualification policy.
3. Select Down as the power state to be applied when the profile is associated with the server.
4. Expand Firmware Management at the bottom of the page and select `VM-Host-Infra` from the Host Firmware list.
5. Click Next.

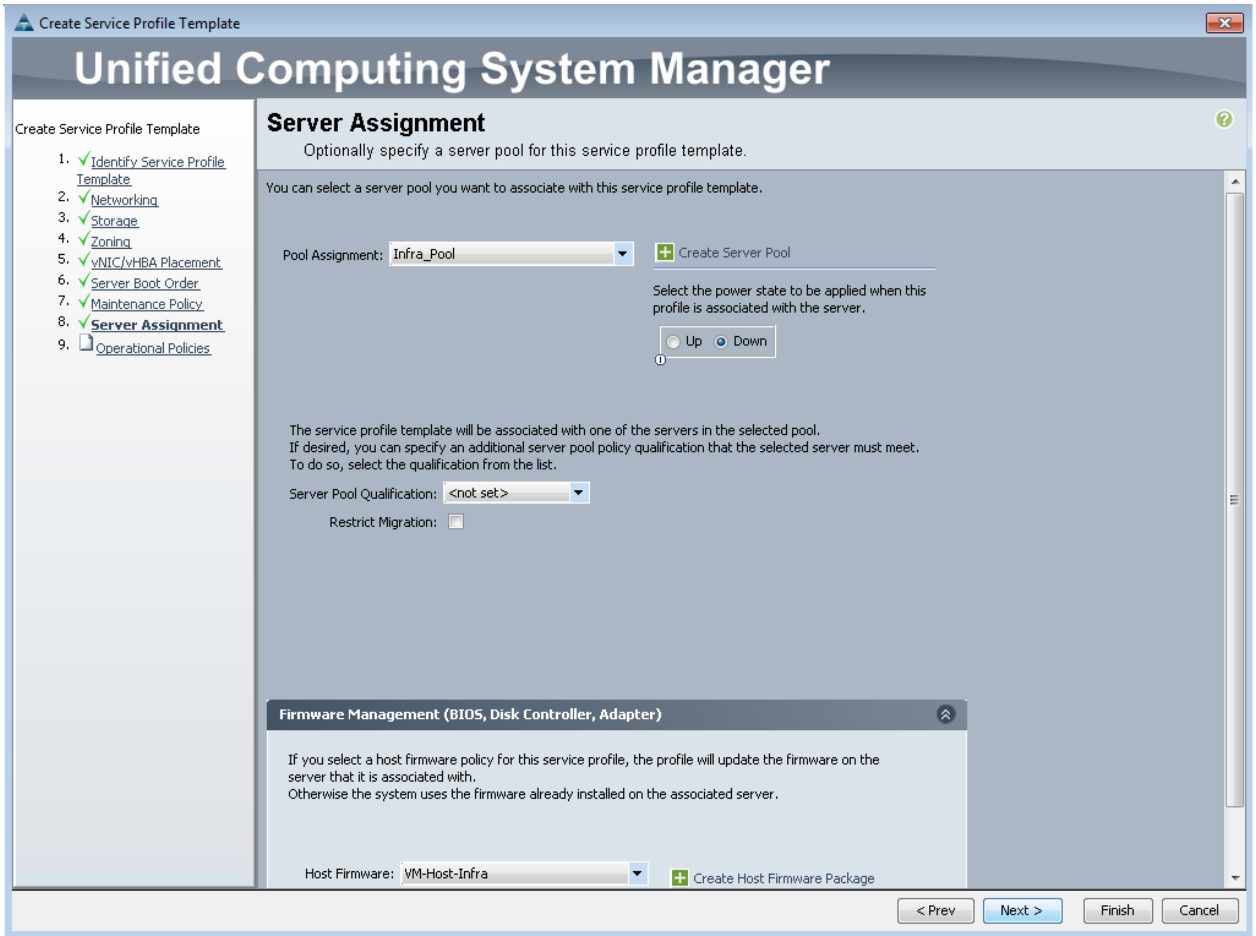

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ Zoning
5. ✓ vNIC/vHBA Placement
6. ✓ Server Boot Order
7. ✓ Maintenance Policy
8. ✓ **Server Assignment**
9. Operational Policies

Server Assignment

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Infra_Pool

Create Server Pool

Select the power state to be applied when this profile is associated with the server.

○ Up ● Down

The service profile template will be associated with one of the servers in the selected pool.
If desired, you can specify an additional server pool policy qualification that the selected server must meet.
To do so, select the qualification from the list.

Server Pool Qualification: <not set>

Restrict Migration: □

Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with.
Otherwise the system uses the firmware already installed on the associated server.

Host Firmware: VM-Host-Infra

Create Host Firmware Package

Configure Operational Policies

1. In the BIOS Policy list, select VM-Host-Infra.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

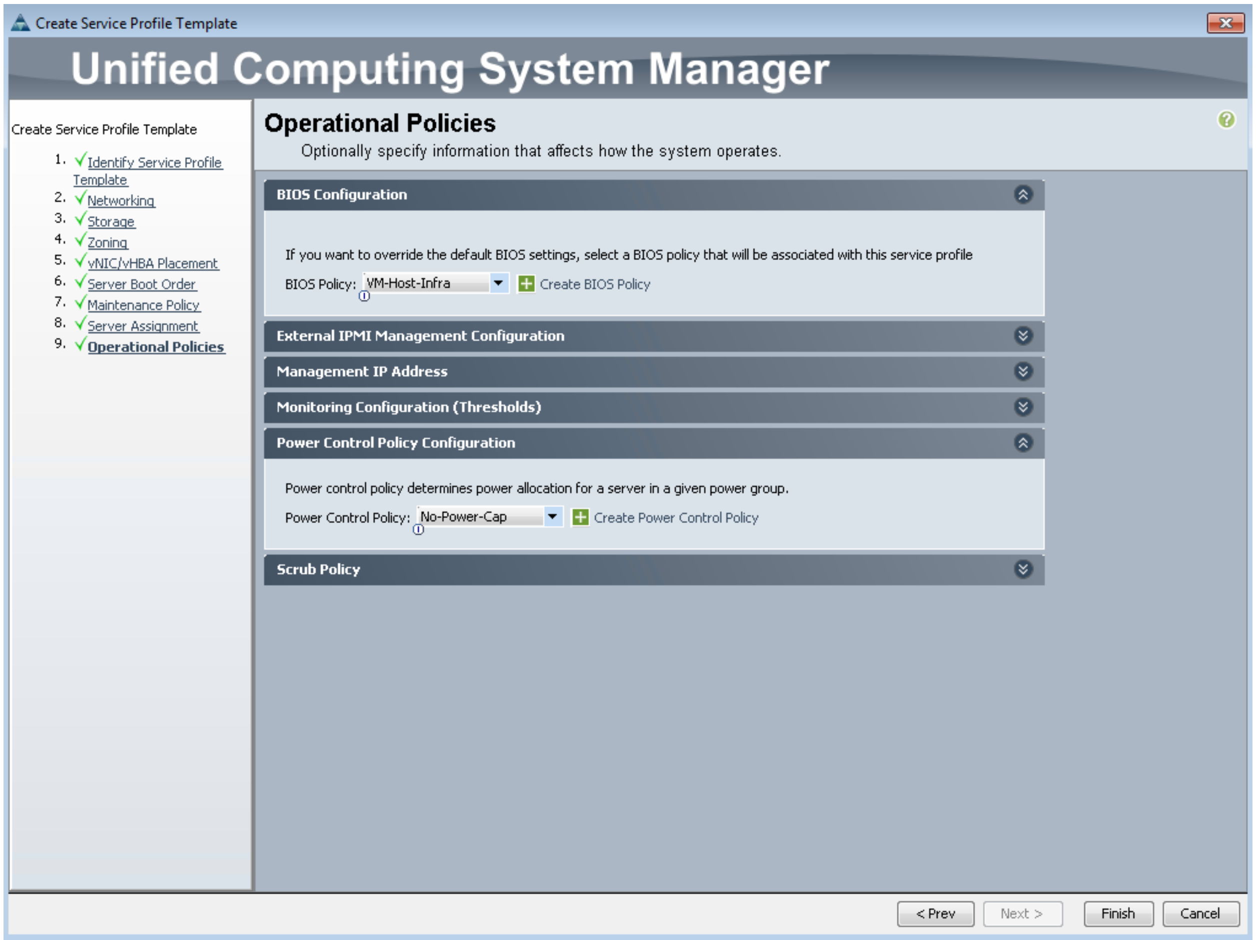
Create Service Profile Template
Unified Computing System Manager
Create Service Profile Template
1. Identify Service Profile Template
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies
BIOS Configuration
Create BIOS Policy
External IPMI Management Configuration
Management IP Address
Monitoring Configuration (Thresholds)
Power Control Policy Configuration
Create Power Control Policy
Scrub Policy
< Prev
Next >
Finish
Cancel

Operational Policies

Optionally specify information that affects how the system operates.

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

BIOS Policy: VM-Host-Infra

Power control policy determines power allocation for a server in a given power group.

Power Control Policy: No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create Service Profiles

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

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

2. Select Service Profile Templates > `root > Service Template VM-Host-Infra-Fabric-A`.

3. Right-click `VM-Host-Infra-Fabric-A` and select Create Service Profiles from Template.

4. Enter `VM-Host-Infra-0` as the service profile prefix.

5. Enter `1` as Name Suffix Starting Number.

6. Enter `1` as the Number of Instances.

7. Click OK to create the service profile.

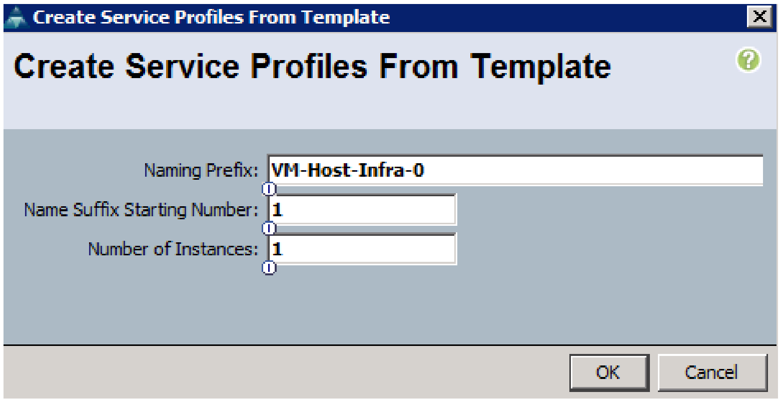
Create Service Profiles From Template

Naming Prefix: VM-Host-Infra-0

Name Suffix Starting Number: 1

Number of Instances: 1

8. Click OK in the confirmation message.

Add More Servers to FlexPod Unit

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

Gather Necessary Information

After the Cisco UCS service profiles have been created, each infrastructure blade in the environment will have a unique configuration. To proceed with the FlexPod deployment, specific information must be gathered from each Cisco UCS blade and from the NetApp controllers. Insert the required information into Table 18 and Table 19 .

Table 18 iSCSI LIFs for iSCSI IQN.

Vserver	iSCSI Target IQN
Infra-SVM	

Note: To gather the iSCSI IQN, run the `iscsi show` command on the storage cluster management interface. For 7-Mode storage, run the `iscsi nodename` command on each storage controller.

Table 19 vNIC iSCSI IQNs for fabric A and fabric B

Cisco UCS Service Profile Name	iSCSI IQN	Variables
VM-Host-Infra-01		<< var_vm_host_infra_01_iqn>>
VM-Host-Infra-02		<< var_vm_host_infra_02_iqn>>

Note: To gather the vNIC IQN information, launch the Cisco UCS Manager GUI. In the navigation pane, click the Servers tab. Expand Servers > Service Profiles > root. Click each service profile and then click the iSCSI vNICs tab on the right. Note Initiator Name displayed at the top of the page under Service Profile Initiator Name

ACI Infrastructure Configuration

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

Physical Connectivity

Follow the physical connectivity guidelines for FlexPod in the section **Error! Reference source not found.**In ACI, both spine and leaf switches are configured using APIC; individual configuration of the switches is not required. APIC discovers the ACI infrastructure switches using LLDP and acts as the central point for the entire configuration.

Cisco APIC Initial Configuration Setup

To configure the Cisco APIC initial setup, complete the following steps:

1. Log into the APIC CIMC using a web browser and launch the KVM.

1. Browse to https://<cimc_ip_address>.

2. Log in using admin as username and use the password defined during CIMC setup.

3. From the Server tab on the left, select Summary and click Launch KVM Console.

4. KVM application will be launched and initial APIC setup screen should be visible.

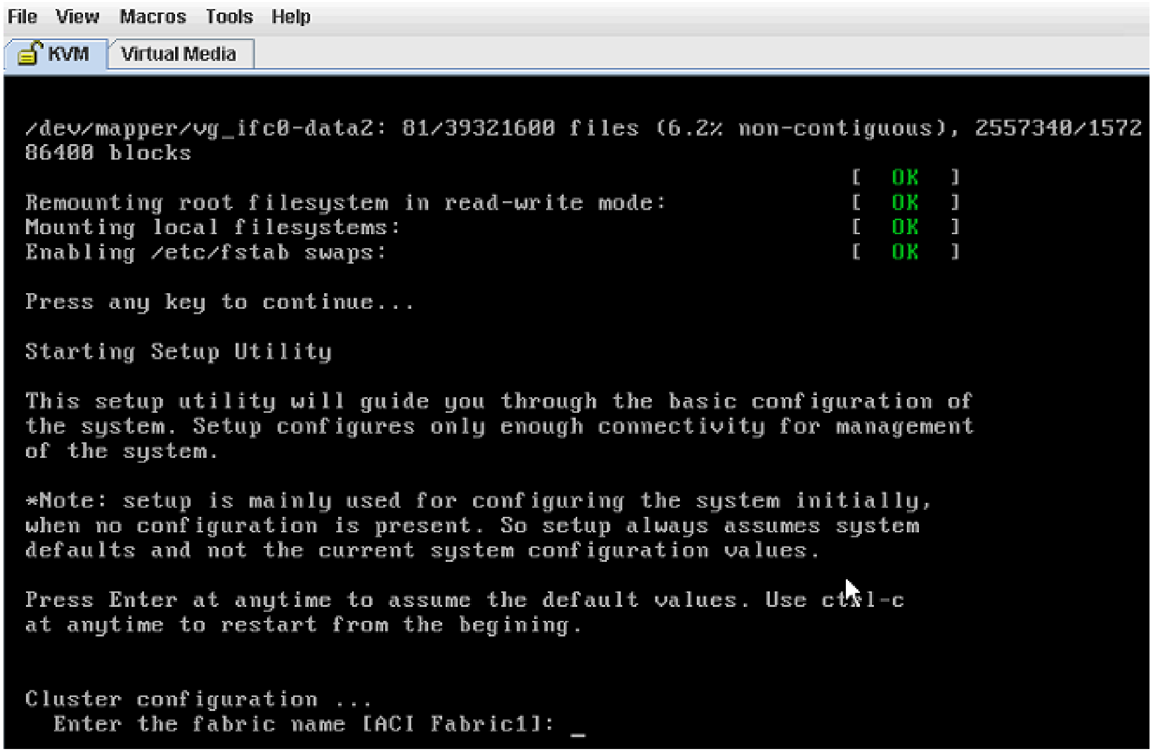

```
/dev/mapper/vg_ifc0-data2: 81/39321600 files (6.2% non-contiguous), 2557340/1572
86400 blocks
                                                         [  OK  ]
Remounting root filesystem in read-write mode:           [  OK  ]
Mounting local filesystems:                              [  OK  ]
Enabling /etc/fstab swaps:                               [  OK  ]

Press any key to continue...

Starting Setup Utility

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

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

Press Enter at anytime to assume the default values. Use ctrl-c
at anytime to restart from the begining.


Cluster configuration ...
  Enter the fabric name [ACI Fabric1]: _
```

5. Press <return> to select the default value for Enter the fabric name. This value can be changed if desired.

6. Press <return> to select the default value for Enter the number of controllers in the fabric. While the fabric can operate with a single APIC, 3 APICs are recommended for redundancy.

7. Enter the controller number currently being set up under Enter the controller ID (1-3). **Please remember only controller number 1 will allow you to setup the admin password.** Remaining controllers and switches sync their passwords to the admin password set on the controller 1.

8. Enter the controller name or choose default name for Enter the controller name.

9. Press <return> to select the default pool under Enter the address pool for TEP addresses. If the network is already in use, please choose a different range.

10. Press <return> to select the default vlan for Enter the VLAN id for infra network.

11. Press <return> to select the default range for Enter address pool for BD multicast addresses.

12. Enter appropriate values for the out of band management network configuration. The out of band management IP address will be used to access the APIC from client browsers.

13. Enter the admin password (controller 1 only).

14. Press <return> to accept the configuration without changes.

15. Let the APIC complete its boot process.

16. Repeat the above steps for all three APIC controllers.

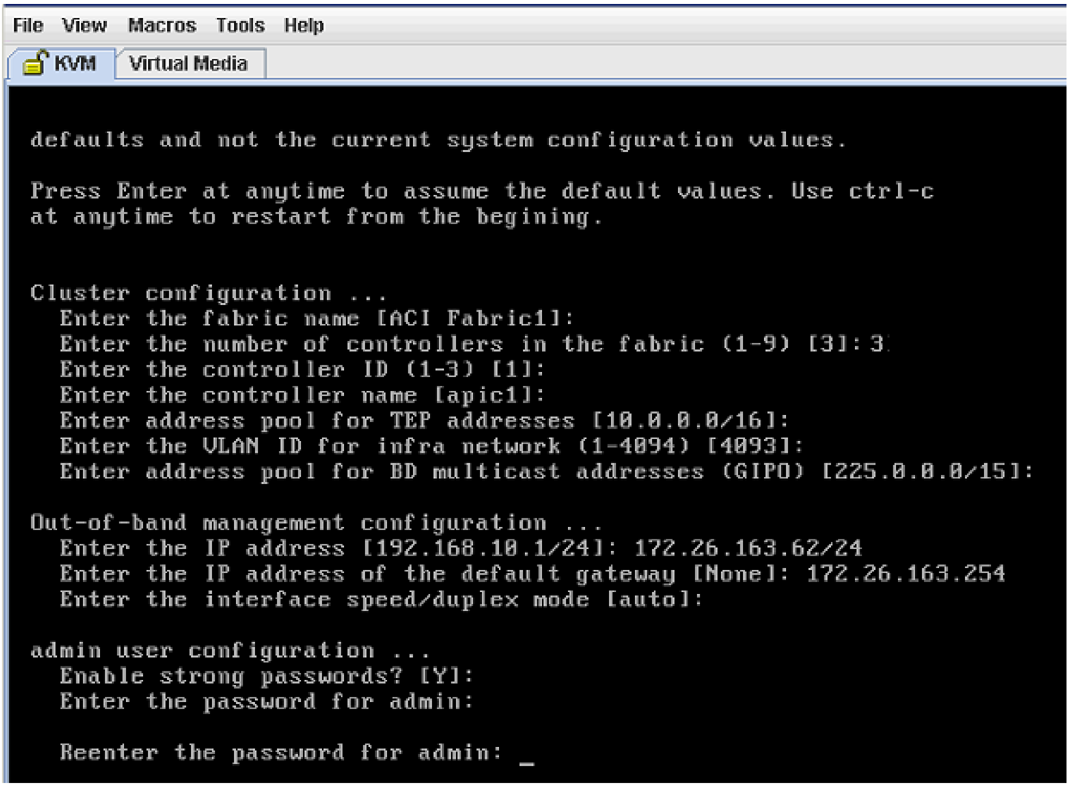

```
defaults and not the current system configuration values.

Press Enter at anytime to assume the default values. Use ctrl-c
at anytime to restart from the begining.


Cluster configuration ...
  Enter the fabric name [ACI Fabric1]:
  Enter the number of controllers in the fabric (1-9) [3]: 3
  Enter the controller ID (1-3) [1]:
  Enter the controller name [apic1]:
  Enter address pool for TEP addresses [10.0.0.0/16]:
  Enter the VLAN ID for infra network (1-4094) [4093]:
  Enter address pool for BD multicast addresses (GIPO) [225.0.0.0/15]:

Out-of-band management configuration ...
  Enter the IP address [192.168.10.1/24]: 172.26.163.62/24
  Enter the IP address of the default gateway [None]: 172.26.163.254
  Enter the interface speed/duplex mode [auto]:

admin user configuration ...
  Enable strong passwords? [Y]:
  Enter the password for admin:

  Reenter the password for admin: _
```

Note: When APIC-1 boots up for the first time, it might take up to 5 minutes to allow login using the admin password set during the setup procedure. If something went wrong during the setup, APIC does allow log-in using a special user called **rescue-user**. If admin password was never set or was not setup properly, rescue-user will allow access to APIC without any password. If an admin password was set previously, use rescue-user with the admin password.

Cisco ACI – Fabric Discovery

1. Log into the APIC GUI using a web browser*.

*Note: For accessing APIC GUI, Google Chrome was utilized during validation

2. Browse to https://<Out of Band IP address of APIC 1>.

3. Log in using admin as username and use the password defined during initial setup.

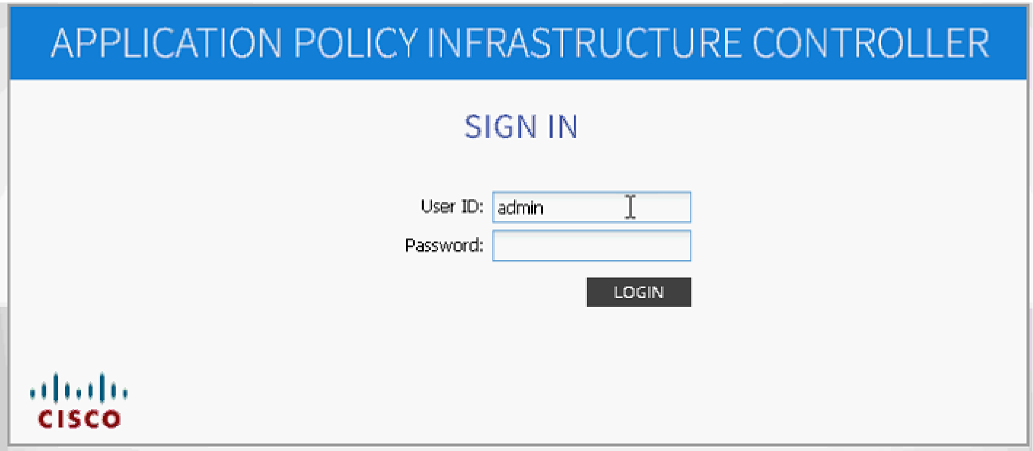

APPLICATION POLICY INFRASTRUCTURE CONTROLLER

SIGN IN

User ID: admin

Password:

4. Click FABRIC from the top bar. Under INVENTORY, expand Fabric Membership.

5. At least one of the leaves should be visible.

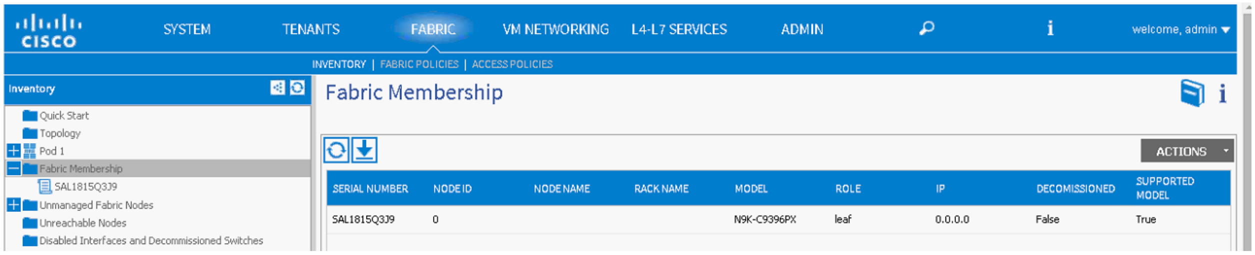

Fabric Membership

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMISSIONED	SUPPORTED MODEL
SAL1815Q3J9	0			N9K-C9396PX	leaf	0.0.0.0	False	True

6. Click FABRIC from the top bar. Under INVENTORY, expand Fabric Membership.

7. Log into the leaf using console connection (admin/<no password needed>) and use the serial number to identify discovered leaf (Leaf-1 or Leaf-2 in the physical setup).

```
switch# show inventory

NAME: "Chassis",  DESCR: "Nexus C9396PX Chassis"

PID: N9K-C9396PX          ,  VID: V02  ,  SN: SAL1815Q3J9

<snip>
```

8. Double-click the identified leaf description on the right hand side and assign a NODE ID value of 101 and NODE NAME <device name>. Click UPDATE.

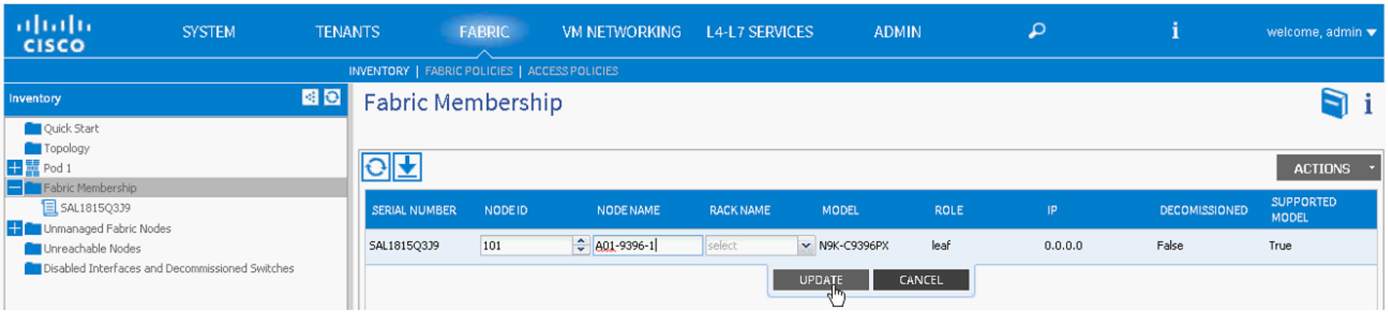

Fabric Membership

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMISSIONED	SUPPORTED MODEL
SAL1815Q3J9	101	A01-9396-1	select	N9K-C9396PX	leaf	0.0.0.0	False	True

9. As the fabric discovery continues, both spines and leaves will start appearing under the Fabic Membership window. Repeat Step 7 to assign the NODE ID and NODE NAME to these devices.

10. When the NODE ID and NODE NAME values are assigned, APIC assigns IP addresses from TEP the pool defined during initial setup.

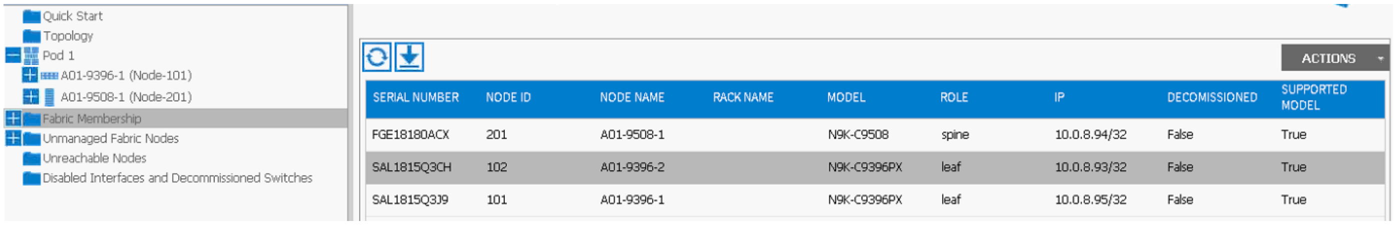

SERIAL NUMBER	NODE ID	NODE NAME	RACK NAME	MODEL	ROLE	IP	DECOMISSIONED	SUPPORTED MODEL
FGE18180ACX	201	A01-9508-1		N9K-C9508	spine	10.0.8.94/32	False	True
SAL1815Q3CH	102	A01-9396-2		N9K-C9396PX	leaf	10.0.8.93/32	False	True
SAL1815Q3J9	101	A01-9396-1		N9K-C9396PX	leaf	10.0.8.95/32	False	True

11. When both leaves and spines are added to the fabric, click Topology on the left hand side. 3 APICs, 2 leaves and 2 spines should be visible.

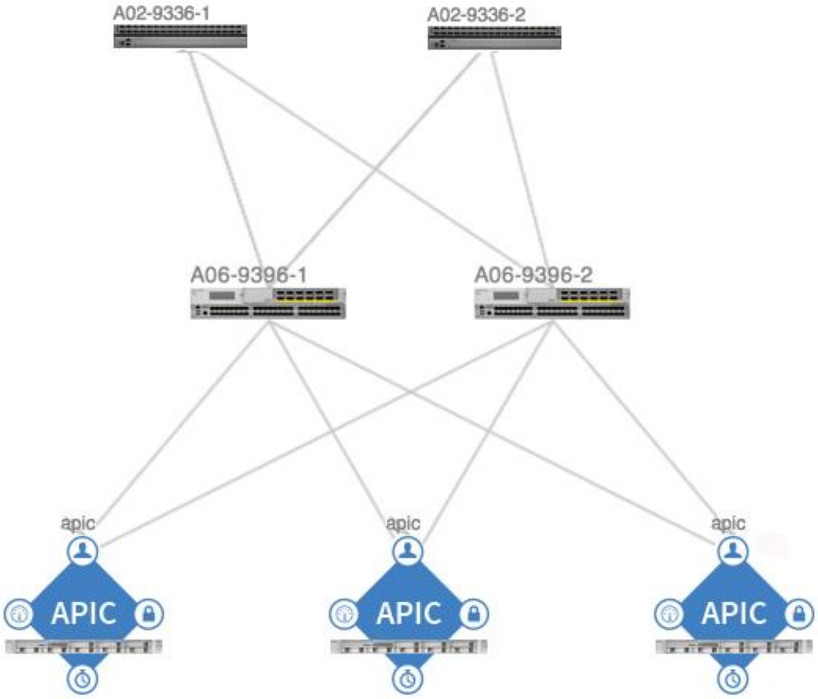

Cisco ACI – Defining Fabric Access Policies

In this section, various access policies such as CDP, LACP and LLDP etc. will be defined. These policies will be used during vPC and VM domain creation. To define fabric access policies, complete the following steps:

1. Click FABRIC and select ACCESS POLICIES under the sub-menu.

2. From the left menu bar, expand Interface Policies.

3. Expand Policies.

4. Right Click CDP Interface and select Create CDP Interface Policy.

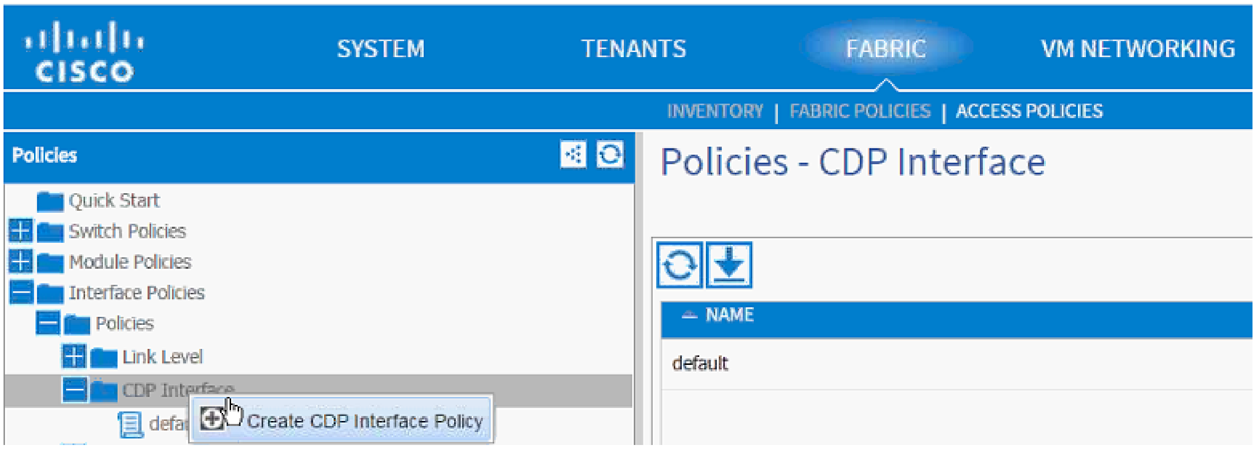

5. In the menu box, enter CDP_Enabled as the policy name and set Admin State Enabled.

6. Click SUBMIT.

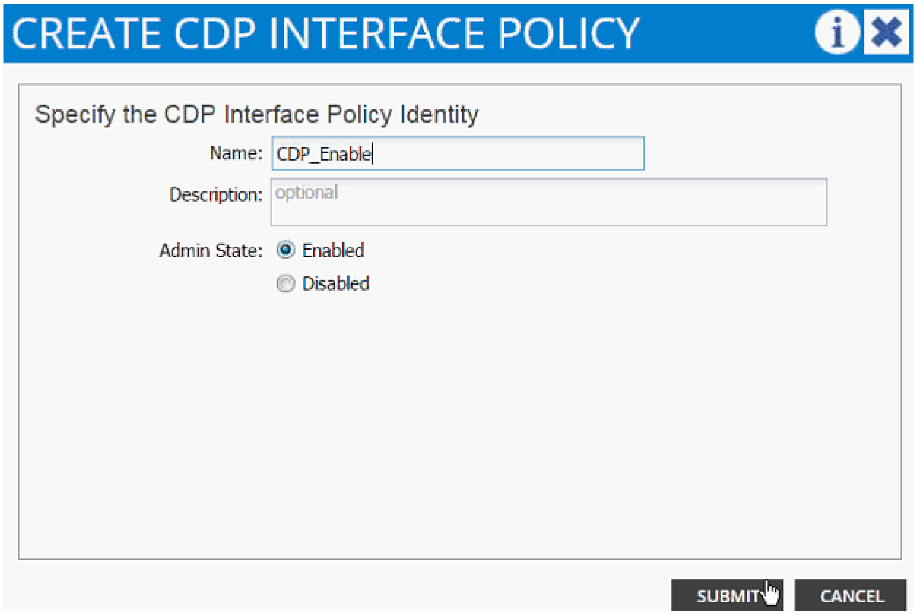
CREATE CDP INTERFACE POLICY

Specify the CDP Interface Policy Identity

Name: CDP_Enable

Description: optional

Admin State: ● Enabled
○ Disabled

7. From the left menu bar, Expand LLDP interface.

8. Right Click and select Create LLDP Interface Policy.

9. In the menu box, enter LLDP_Disabled as the policy name and set both Transmit State and Receive State Disabled.

10. Click SUBMIT.

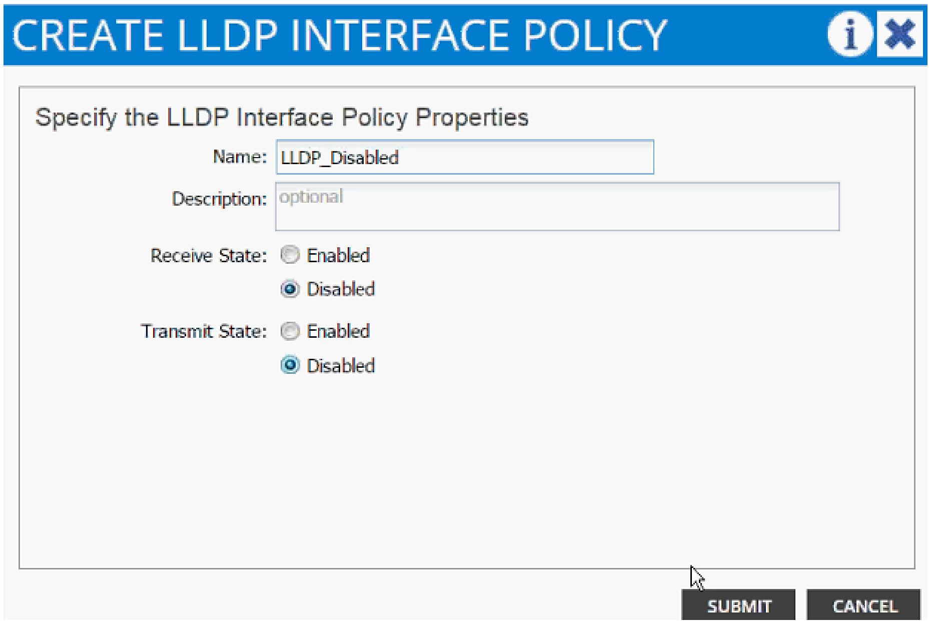

CREATE LLDP INTERFACE POLICY

Specify the LLDP Interface Policy Properties

Name: LLDP_Disabled

Description: optional

Receive State: ○ Enabled ● Disabled

Transmit State: ○ Enabled ● Disabled

11. Right-click and select Create LLDP Interface again.

12. In the menu box, enter LLDP_Enabled as the policy name and set both Transmit State and Receive State Enabled.

13. Click SUBMIT.

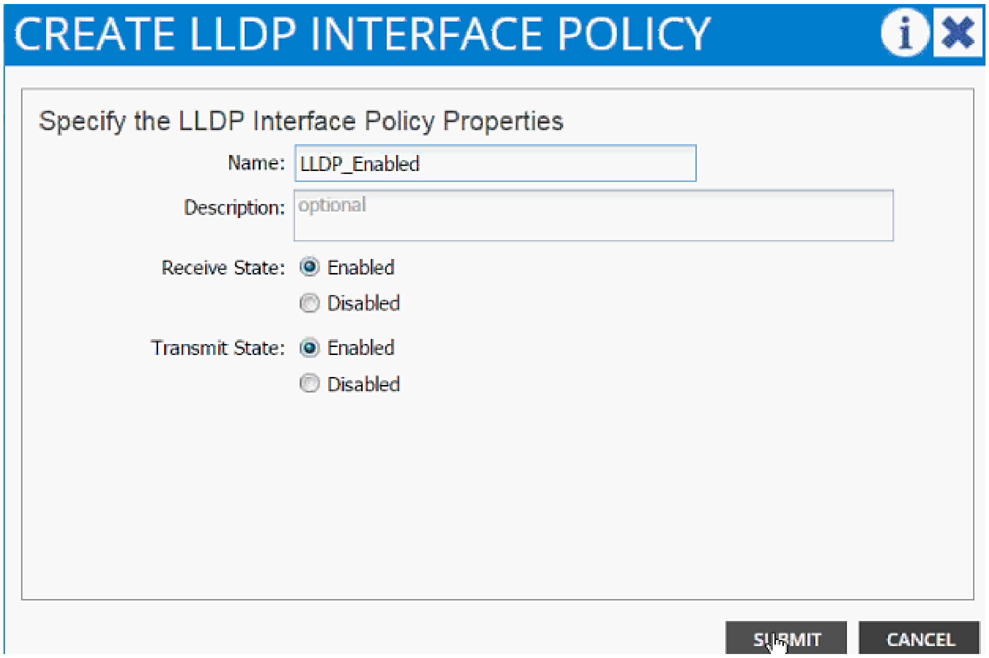

CREATE LLDP INTERFACE POLICY

Specify the LLDP Interface Policy Properties

Name: LLDP_Enabled

Description: optional

Receive State: ● Enabled ○ Disabled

Transmit State: ● Enabled ○ Disabled

14. From the left menu bar, Expand LACP.

15. Right-click and select Create LACP Policy.

16. In the menu box, enter LACP_Active as the policy name and select Mode Active. Leave remaining options as default.

17. Click SUBMIT.

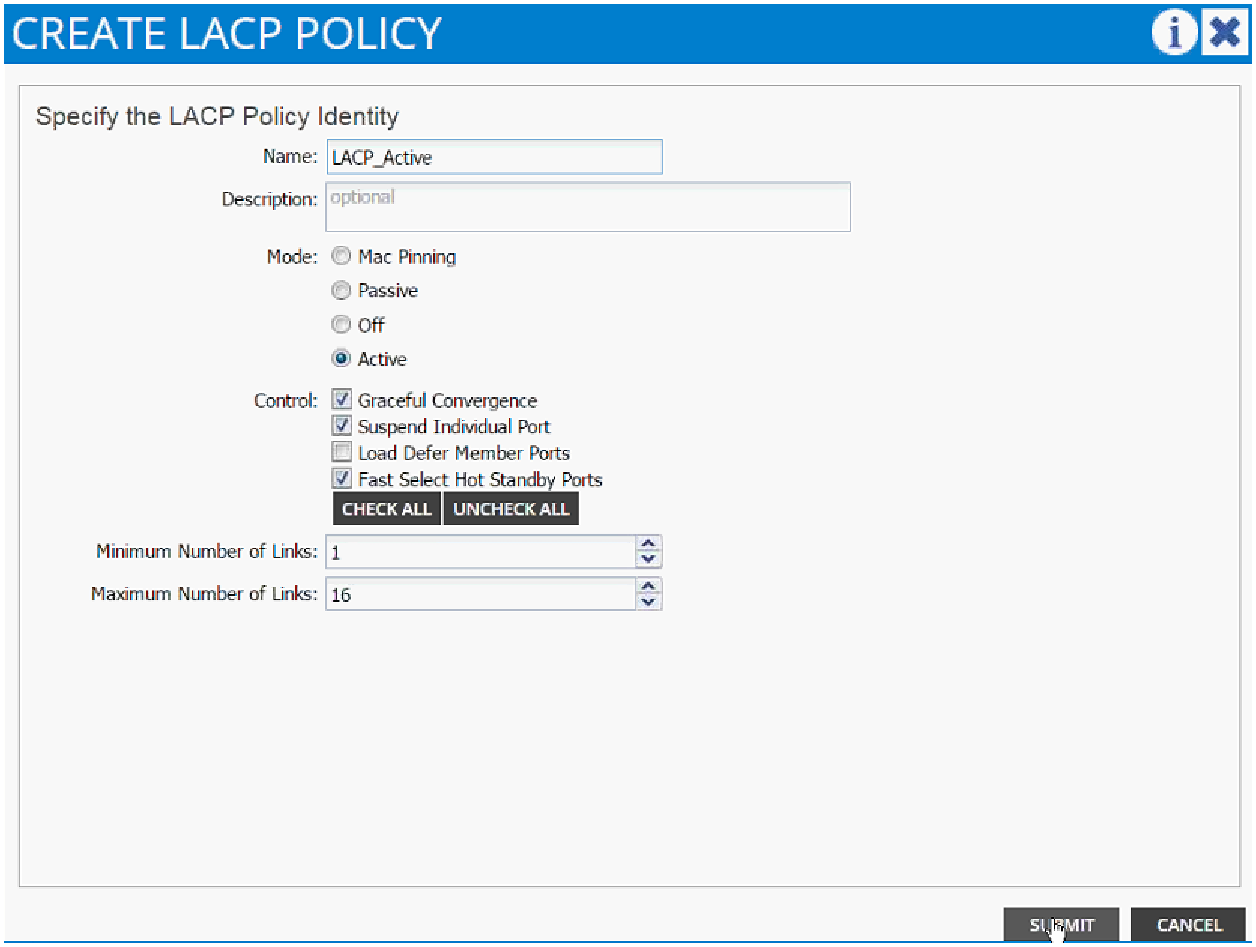

CREATE LACP POLICY

Specify the LACP Policy Identity

Name: LACP_Active

Description: optional

Mode: ○ Mac Pinning ○ Passive ○ Off ● Active

Control: ☑ Graceful Convergence ☑ Suspend Individual Port □ Load Defer Member Ports ☑ Fast Select Hot Standby Ports

Minimum Number of Links: 1

Maximum Number of Links: 16

18. From the left menu bar, expand LACP.

19. Right-click and select Create LACP Policy.

20. In the menu box, enter LACP_MAC_Pinning as the policy name and select “Mode “Mac Pinning. Leave remaining options as default.

21. Click SUBMIT.

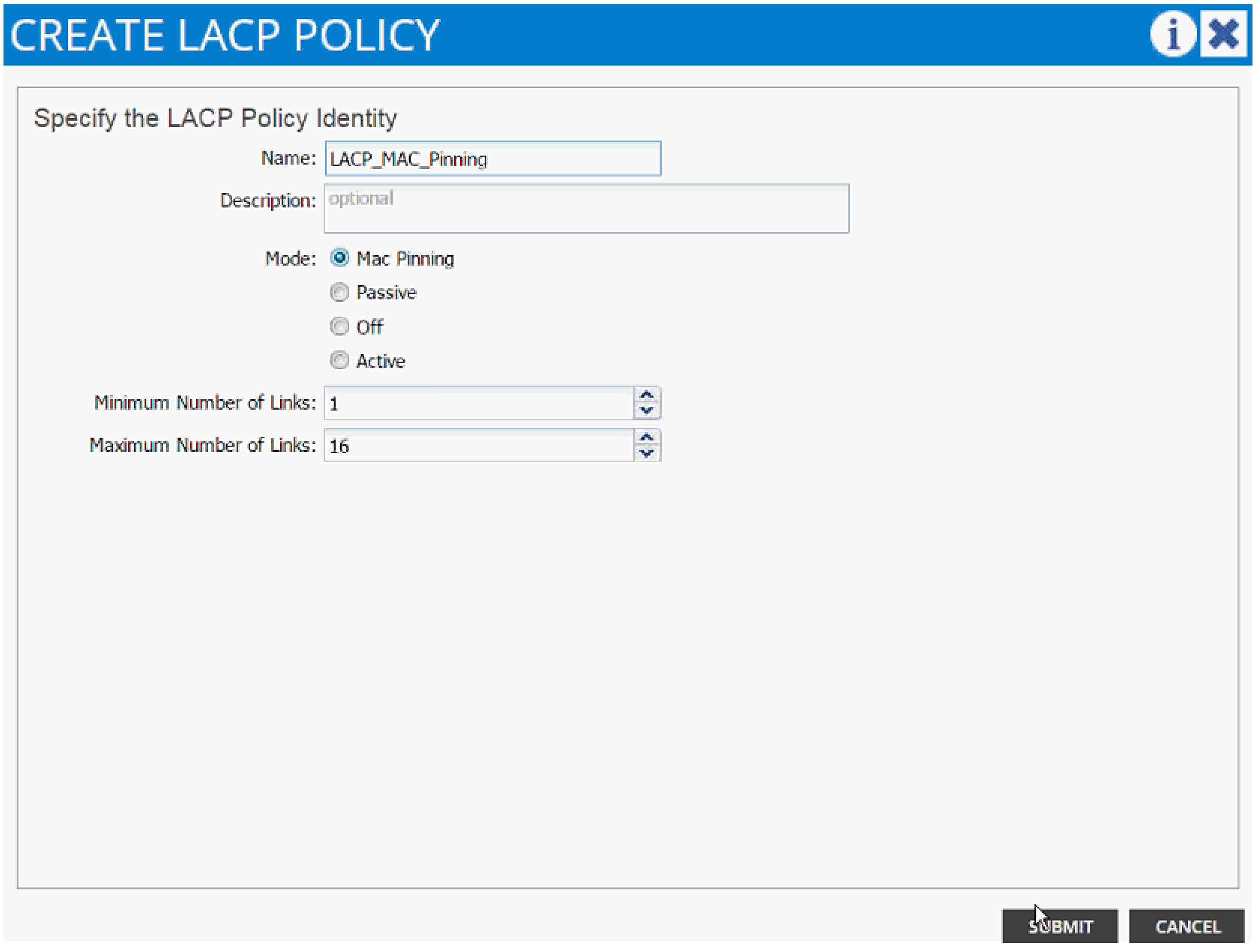

CREATE LACP POLICY

Specify the LACP Policy Identity

Name: LACP_MAC_Pinning

Description: optional

Mode: ● Mac Pinning ○ Passive ○ Off ○ Active

Minimum Number of Links: 1

Maximum Number of Links: 16

Cisco ACI – Creating vPC for Cisco UCS Fabric Interconnect A

To create a vPC for Cisco UCS Fabric Interconnect A, complete the following steps:

1. From the main menu, click FABRIC and select Access Policies.
2. Right-click Interface Policies and select Configure interface, PC and vPC.

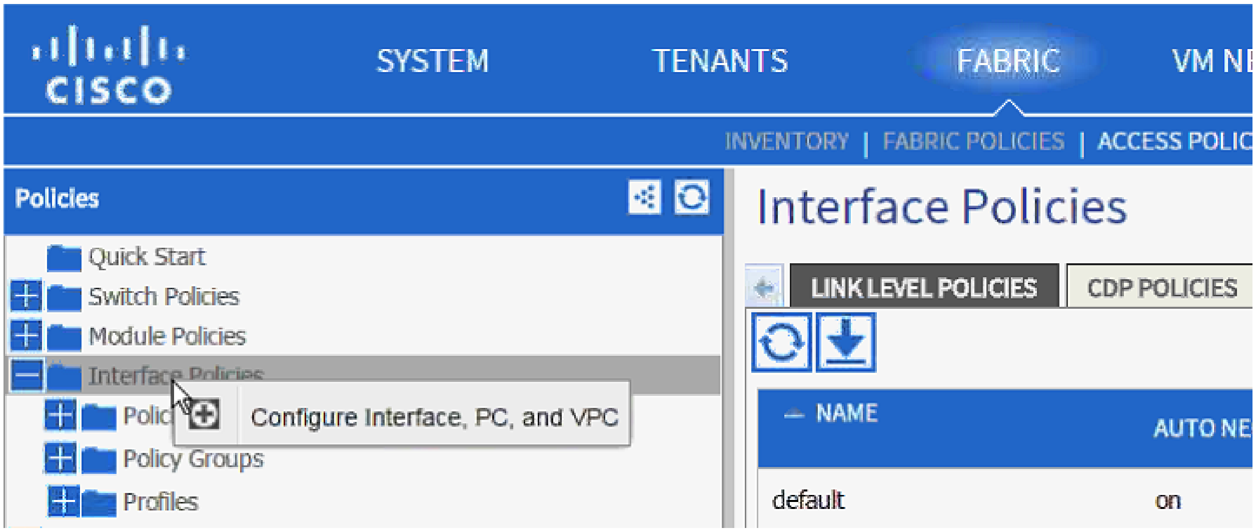

NAME	AUTO NE
default	on

3. In the dialog box, Click + under the vPC SWITCH PAIRS.

4. Enter 10 as vPC Domain ID.

5. Drop down the Switch 1 and Switch 2 values and select both leaves.

6. Click Save.

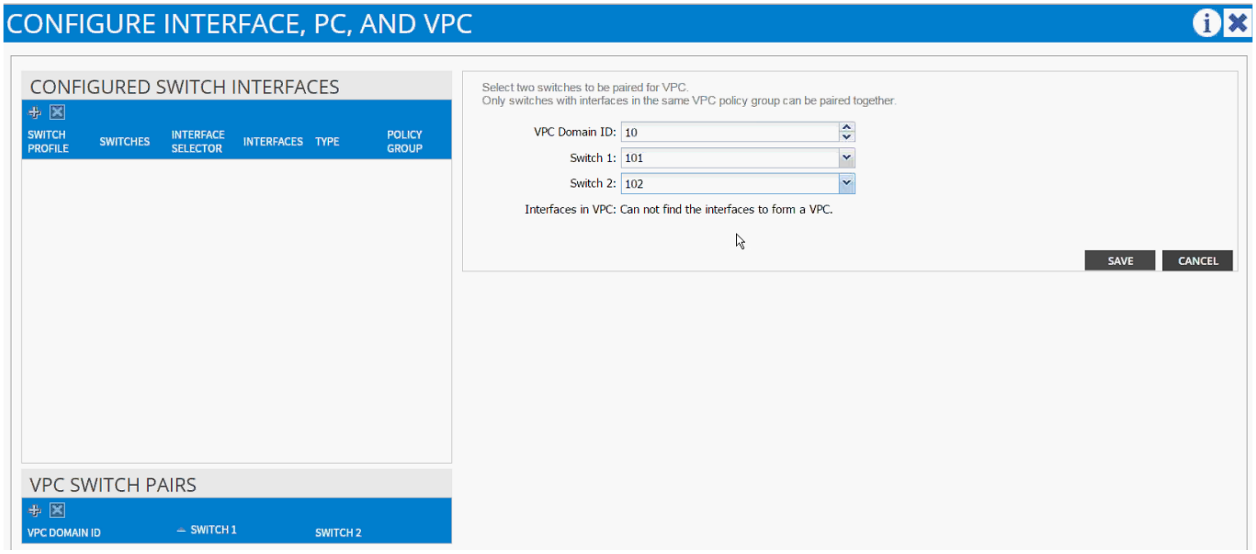

CONFIGURE INTERFACE, PC, AND VPC

CONFIGURED SWITCH INTERFACES

SWITCH PROFILE	SWITCHES	INTERFACE SELECTOR	INTERFACES	TYPE	POLICY GROUP

Select two switches to be paired for VPC.
Only switches with interfaces in the same VPC policy group can be paired together.

VPC Domain ID: 10
Switch 1: 101
Switch 2: 102
Interfaces in VPC: Can not find the interfaces to form a VPC.

VPC SWITCH PAIRS

VPC DOMAIN ID	SWITCH 1	SWITCH 2

7. Validate that the create vPC domain appears under the vPC SWITCH PAIRS.

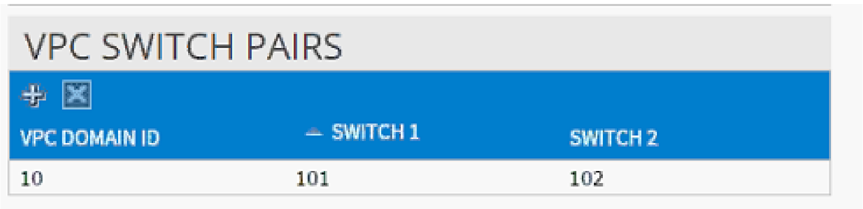

VPC SWITCH PAIRS

VPC DOMAIN ID	SWITCH 1	SWITCH 2
10	101	102

8. In the dialog box, click + under the CONFIGURED SWITCH INTERFACES.

9. Drop down the menu next to Switches and select both leaves.

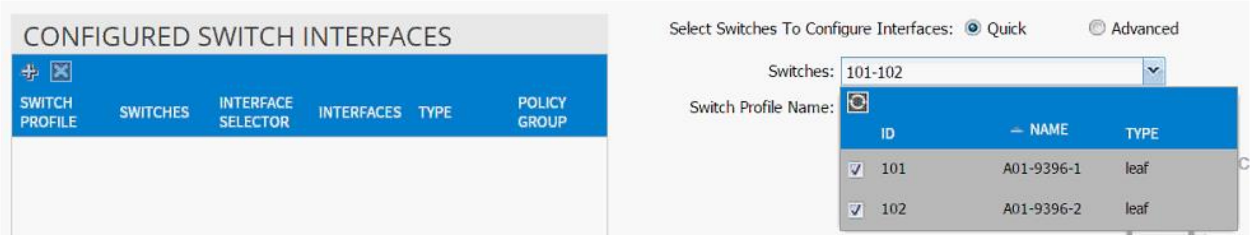

CONFIGURED SWITCH INTERFACES

SWITCH PROFILE	SWITCHES	INTERFACE SELECTOR	INTERFACES	TYPE	POLICY GROUP

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name:

	ID	NAME	TYPE
☑	101	A01-9396-1	leaf
☑	102	A01-9396-2	leaf

10. Enter <sp-UCS-FI-1> as Switch Profile Name. UCS-FI-1 is the host name for UCS Fabric Interconnect A.

11. Click the + sign to add interfaces.

12. Select vPC radio button to configure vPC.

13. Enter 1/19 under Interfaces. This is the port on both switches where Fabric Interconnect A is connected.

14. Enter <ifs- UCS-FI-1> as Interface Selector Name.

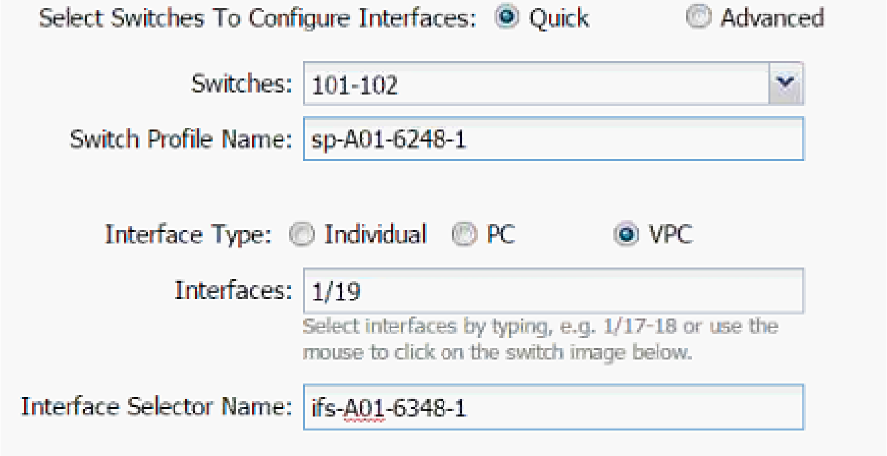

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: sp-A01-6248-1

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/19

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Interface Selector Name: ifs-A01-6348-1

15. Drop down vPC Policy Group and click Create vPC Interface Policy Group. A new dialog box will appear.

16. Enter <pg- UCS-FI-1> as the name of the vPC INTERFACE POLICY GROUP in the dialog box.

17. Select various policy values from the drop-down menus.

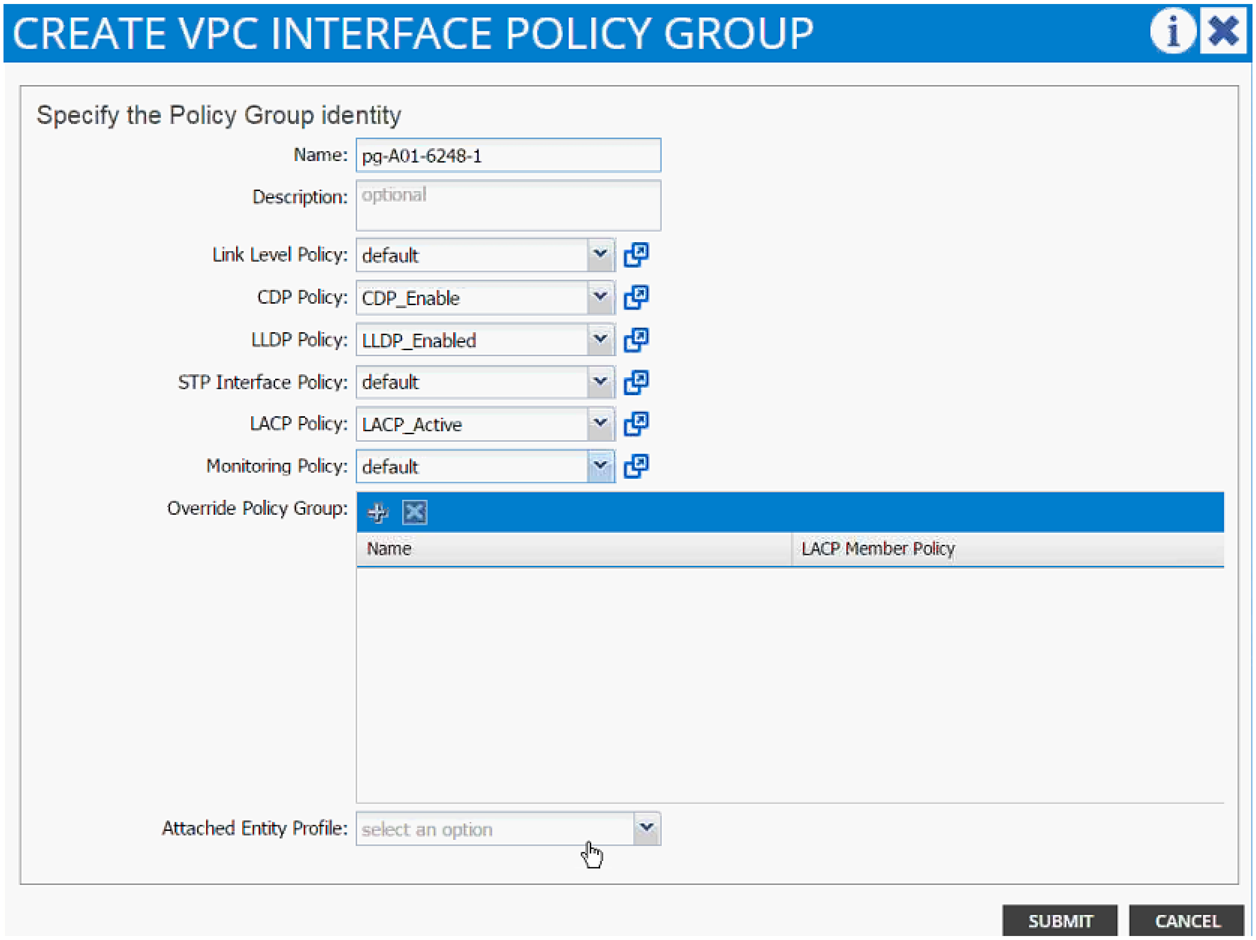

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name:	pg-A01-6248-1
Description:	optional
Link Level Policy:	default
CDP Policy:	CDP_Enable
LLDP Policy:	LLDP_Enabled
STP Interface Policy:	default
LACP Policy:	LACP_Active
Monitoring Policy:	default
Override Policy Group:	
Attached Entity Profile:	select an option

18. From the Attached Entity Profile drop-down list, select Create Attachable Access Entity Profile. A new dialog box will appear.

19. The two UCS Fabric Interconnects will share this Attachable Entity Profile (AEP). Enter <aep-UCS-FI-hostname> as the Name (avoid using A or B at the end).

20. Click + to add Domain.

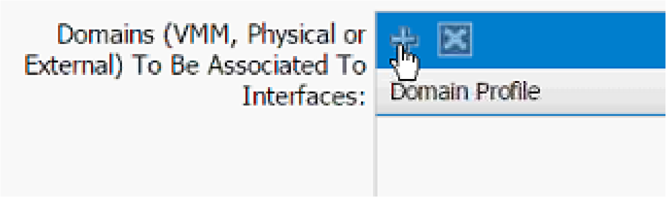

Domains (VMM, Physical or External) To Be Associated To Interfaces:
Domain Profile

21. In the added domain, drop-down the menu and select Create Physical Domain.

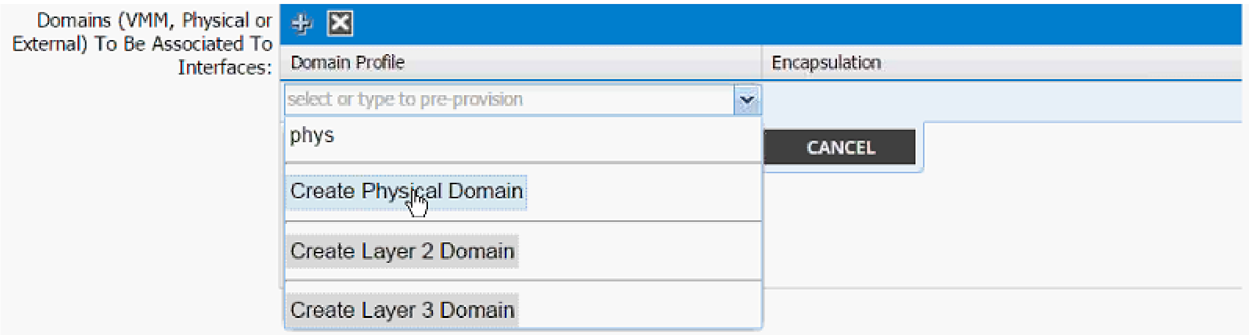

22. In the Create Physical Domain dialog box, enter <pd-UCS-FI hostname> as Name.

23. Drop down the VLAN Pool menu and select Create VLAN Pool.

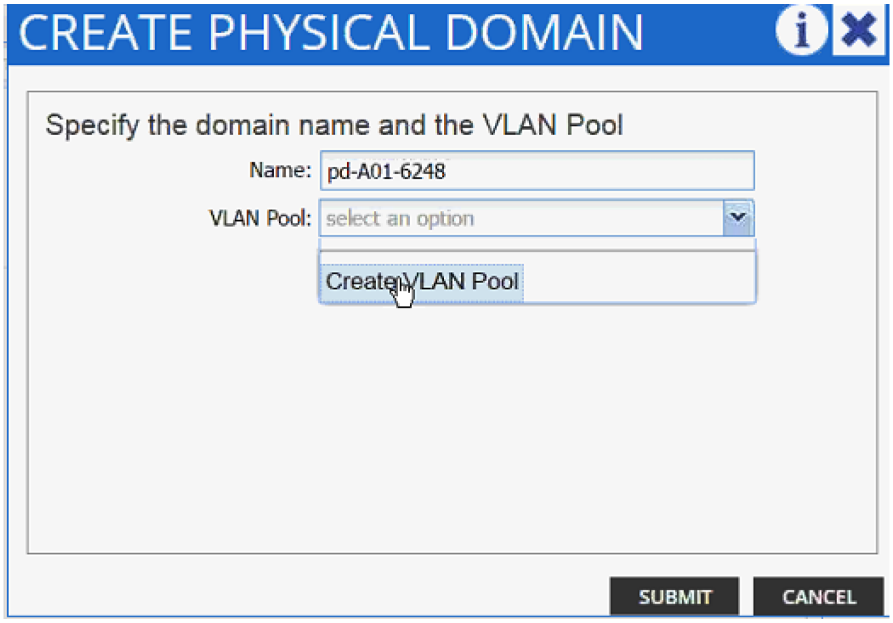
CREATE PHYSICAL DOMAIN

Specify the domain name and the VLAN Pool

Name: pd-A01-6248

VLAN Pool: select an option

24. In the Create VLAN Pool dialog box, enter <vp-UCS-FI hostname> as Name.

25. Select Allocation Mode Static Allocation.

26. Click + next to Encap Block.

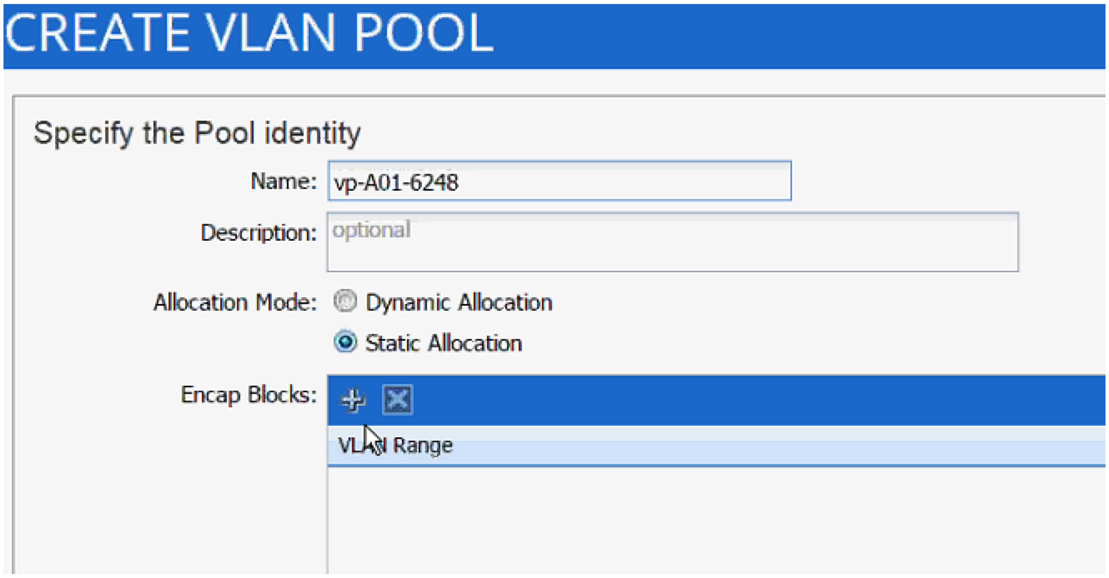

CREATE VLAN POOL

Specify the Pool identity

Name: vp-A01-6248

Description: optional

Allocation Mode: ○ Dynamic Allocation ● Static Allocation

Encap Blocks:

VLAN Range

27. In the CREATE RANGES dialog box, enter the two iSCSI VLANs and the NFS VLAN.

Note: In the screenshot below, 911, 912 are iSCSI VLANs and 3270 is the NFS VLAN configured on Cisco UCS.

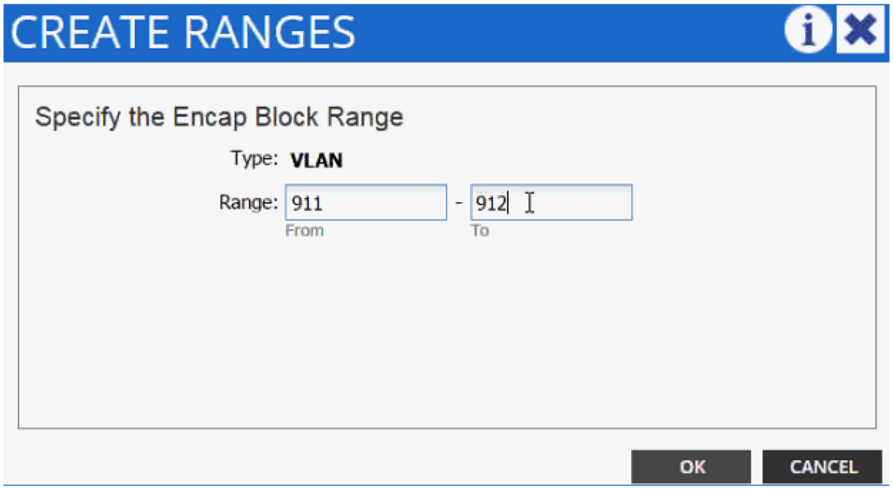

CREATE RANGES

Specify the Encap Block Range

Type: **VLAN**

Range: 911 - 912

From To

28. Click OK.

29. Click +again to add NFS VLAN and in the CREATE RANGES dialog box, enter range 3270-3270 (single VLAN).

30. Click OK.

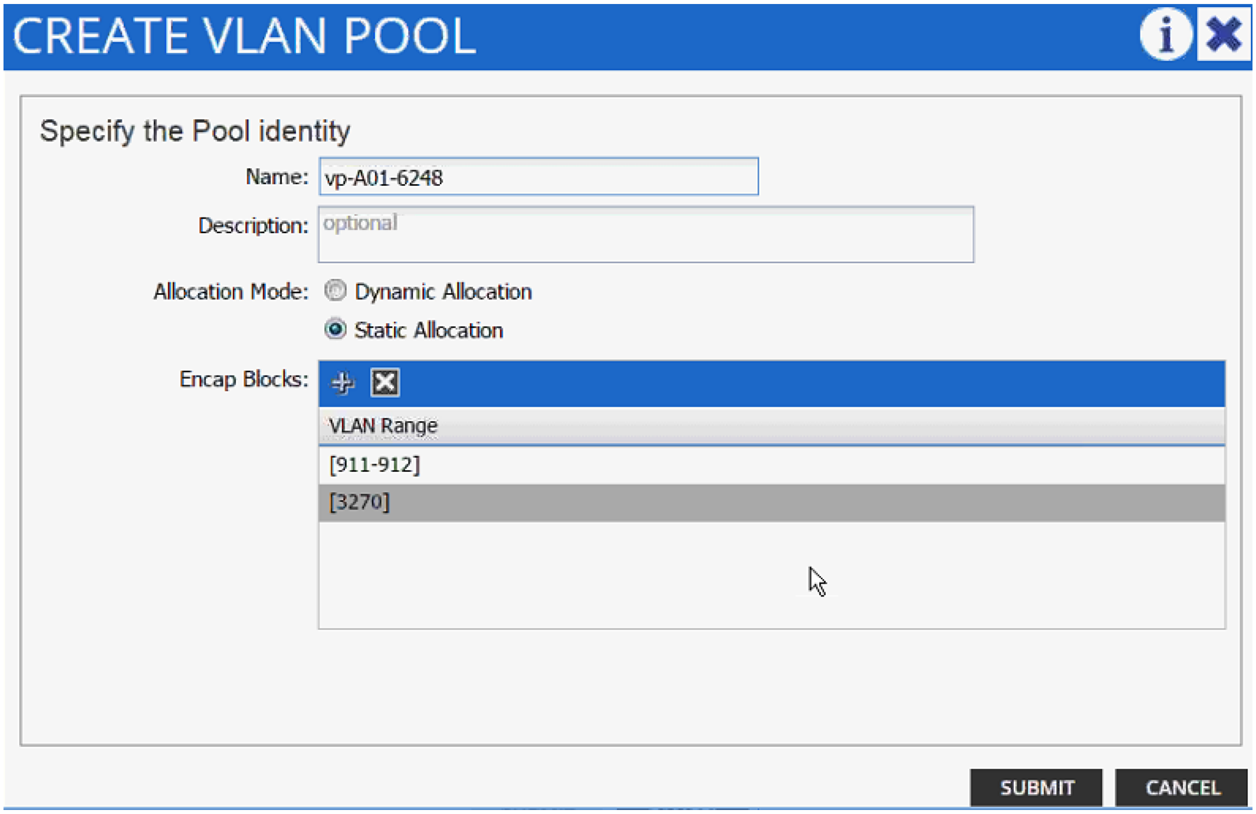

CREATE VLAN POOL

Specify the Pool identity

Name: vp-A01-6248

Description: optional

Allocation Mode: ○ Dynamic Allocation
● Static Allocation

Encap Blocks:

VLAN Range
[911-912]
[3270]

31. Click SUBMIT to finish VLAN pool creation.

32. Click SUBMIT to finish Physical Domain creation.

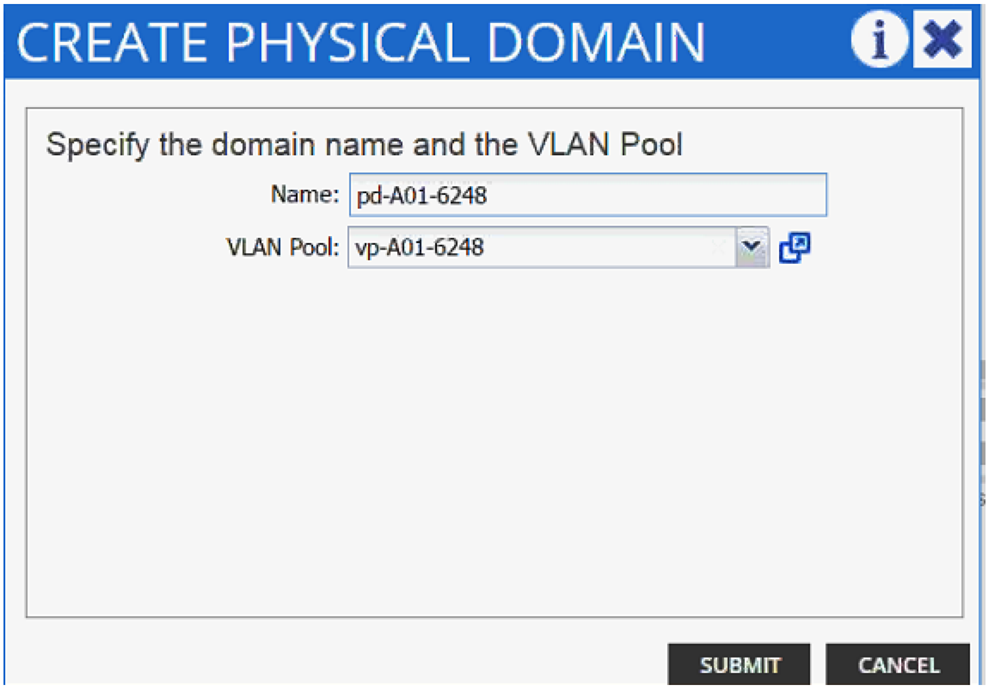

CREATE PHYSICAL DOMAIN

Specify the domain name and the VLAN Pool

Name: pd-A01-6248

VLAN Pool: vp-A01-6248

33. Click UPDATE to finish adding Physical domain to AEP.

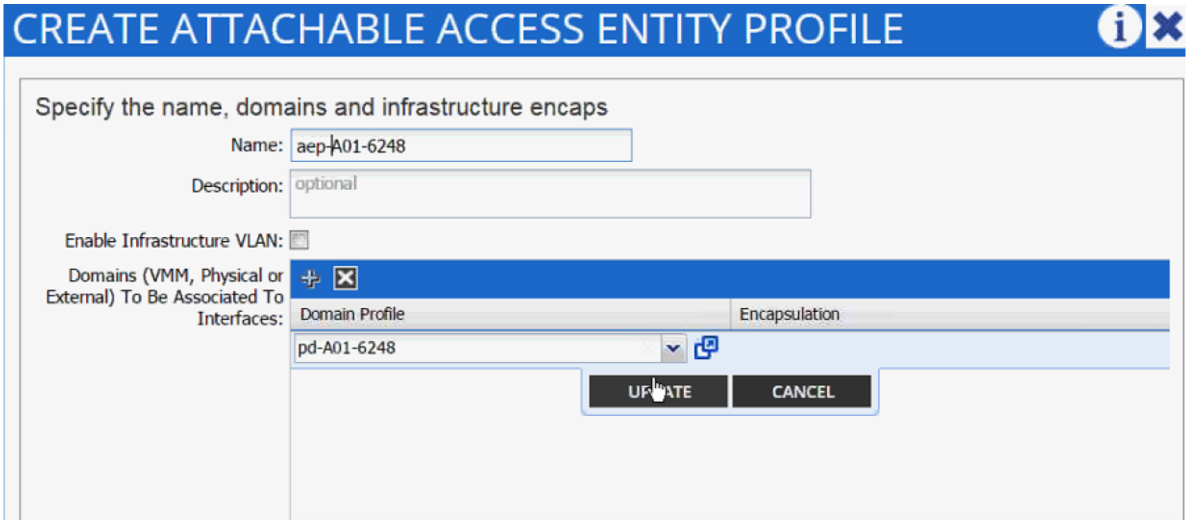

CREATE ATTACHABLE ACCESS ENTITY PROFILE

Specify the name, domains and infrastructure encaps

Name: aep-A01-6248

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation
pd-A01-6248	

34. Click SUBMIT to finish adding AEP.

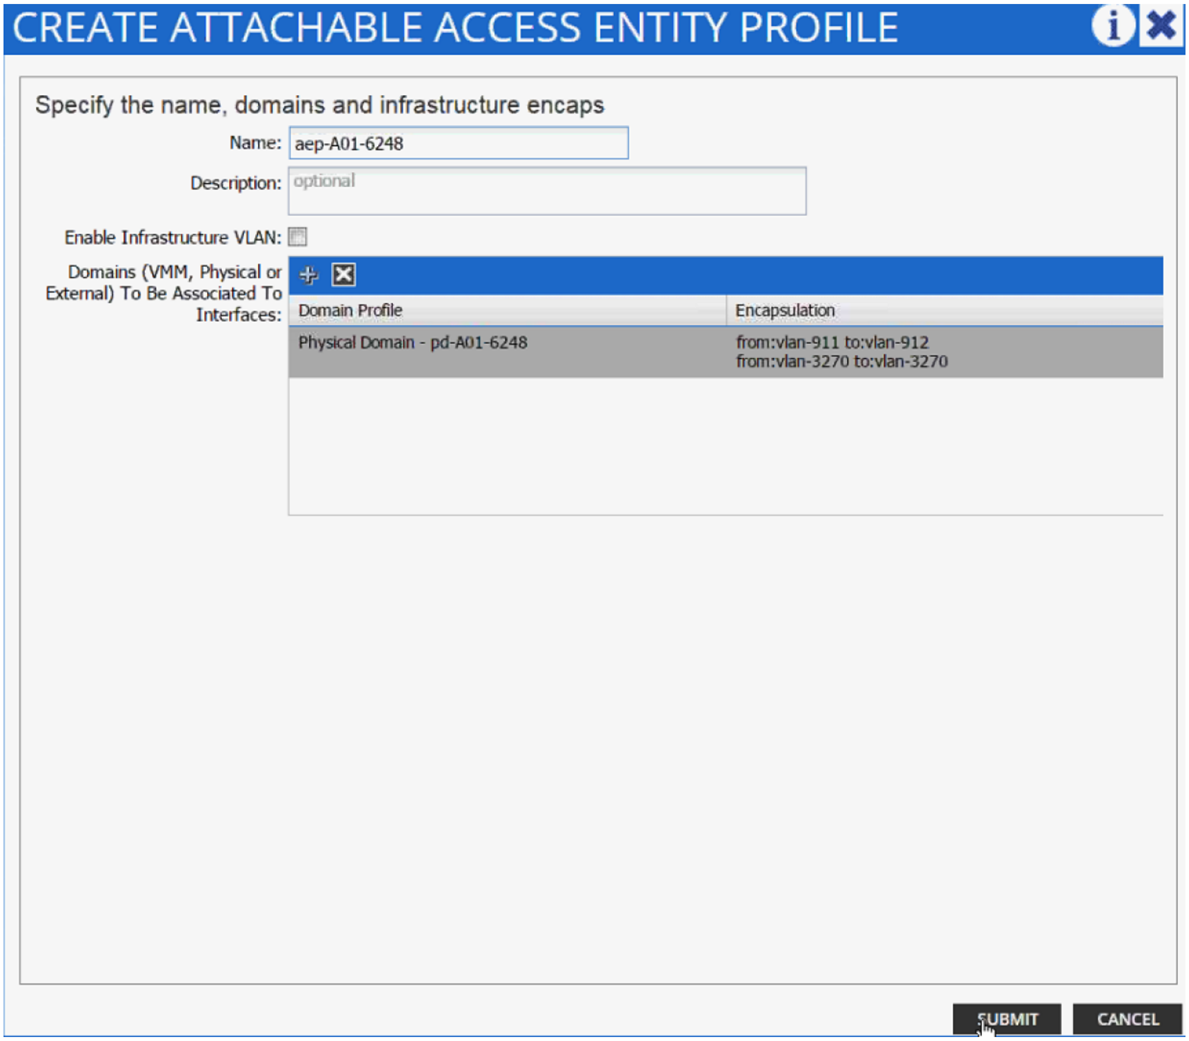

CREATE ATTACHABLE ACCESS ENTITY PROFILE

Specify the name, domains and infrastructure encaps

Name: aep-A01-6248

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation

Physical Domain - pd-A01-6248	from:vlan-911 to:vlan-912
	from:vlan-3270 to:vlan-3270

35. Click SUBMIT to finish creating vPC Interface Policy Group.

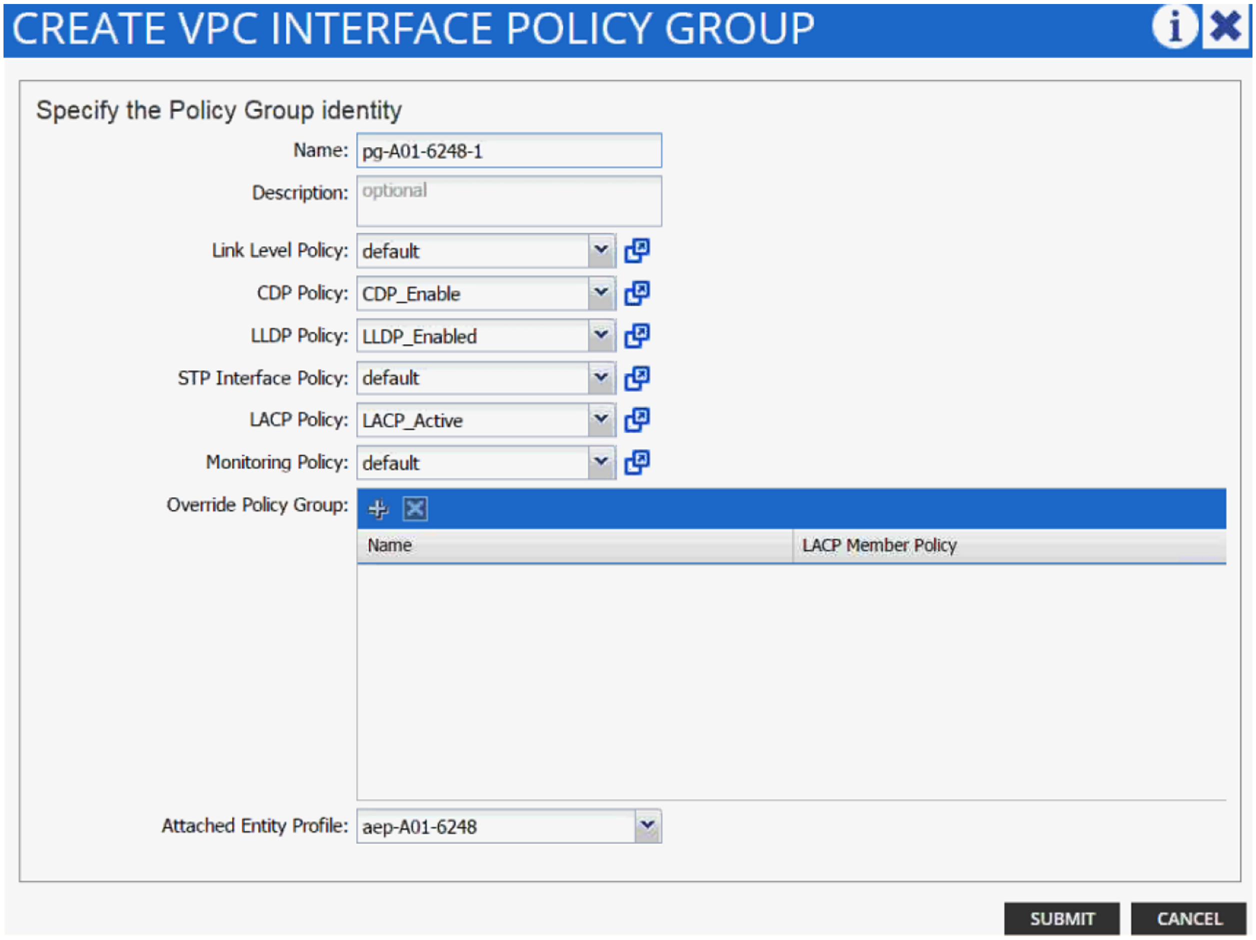

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name:	pg-A01-6248-1
Description:	optional
Link Level Policy:	default
CDP Policy:	CDP_Enable
LLDP Policy:	LLDP_Enabled
STP Interface Policy:	default
LACP Policy:	LACP_Active
Monitoring Policy:	default
Override Policy Group:	
Attached Entity Profile:	aep-A01-6248

Override Policy Group:

Name	LACP Member Policy

36. On the Configure Interface, PC, vPC screen, click SAVE.

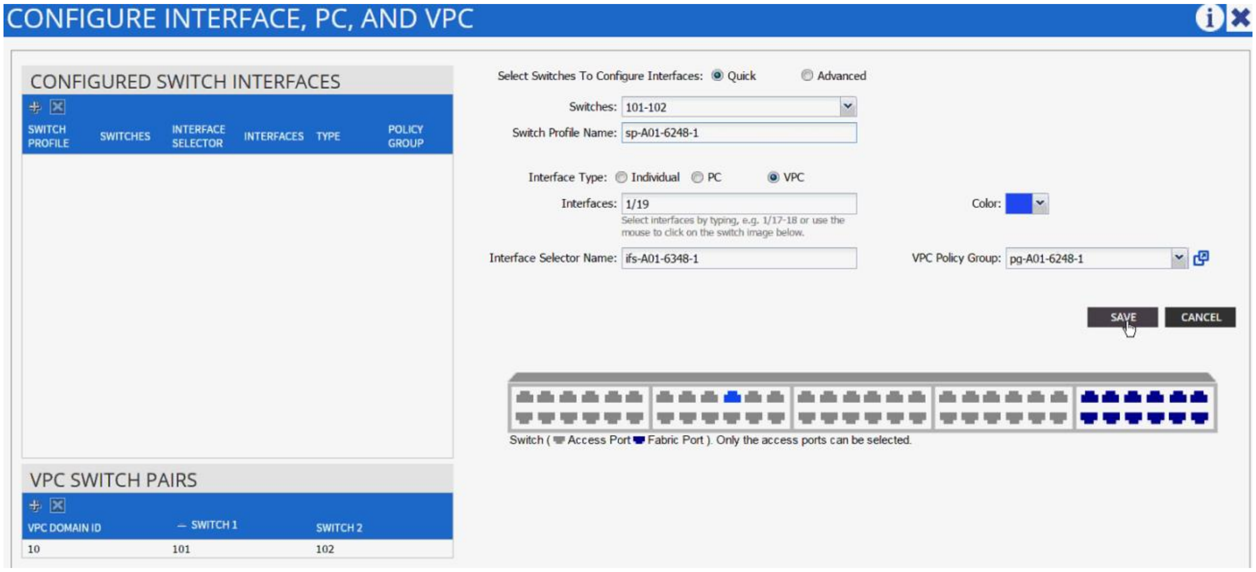

CONFIGURE INTERFACE, PC, AND VPC

CONFIGURED SWITCH INTERFACES

SWITCH PROFILE	SWITCHES	INTERFACE SELECTOR	INTERFACES	TYPE	POLICY GROUP

VPC SWITCH PAIRS

VPC DOMAIN ID	SWITCH 1	SWITCH 2
10	101	102

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches:	101-102
Switch Profile Name:	sp-A01-6248-1
Interface Type:	○ Individual ○ PC ● VPC
Interfaces:	1/19
Color:	🟦
Interface Selector Name:	ifs-A01-6348-1
VPC Policy Group:	pg-A01-6248-1

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Switch (Access Port Fabric Port). Only the access ports can be selected.

37. Click SAVE.

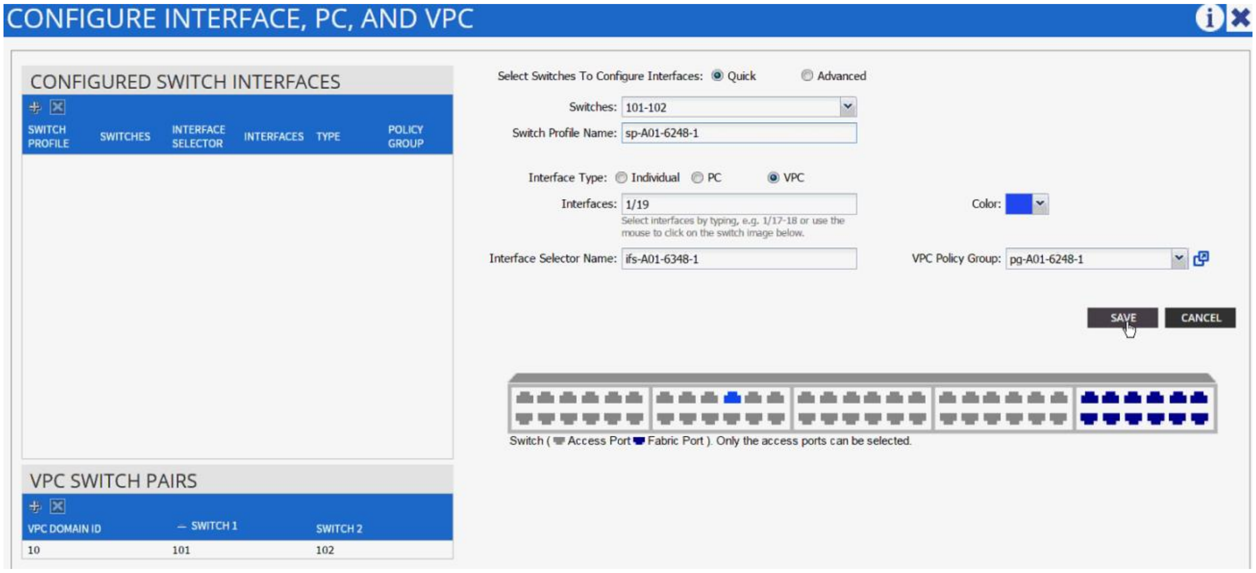

CONFIGURE INTERFACE, PC, AND VPC

CONFIGURED SWITCH INTERFACES

SWITCH PROFILE	SWITCHES	INTERFACE SELECTOR	INTERFACES	TYPE	POLICY GROUP

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: sp-A01-6248-1

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/19

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Color:

Interface Selector Name: ifs-A01-6348-1

VPC Policy Group: pg-A01-6248-1

SAVE
CANCEL
Switch (Access Port Fabric Port). Only the access ports can be selected.

VPC SWITCH PAIRS

VPC DOMAIN ID	SWITCH 1	SWITCH 2
10	101	102

38. Click SUBMIT to finish the vPC creation using wizard.

39. Under Fabric, select Inventory.

40. From the left menu, expand Pod 1, expand Leaf-1 and expand Interfaces followed by vPC Interfaces.

41. Validate that the vPC domain 10 and the vPC exist.

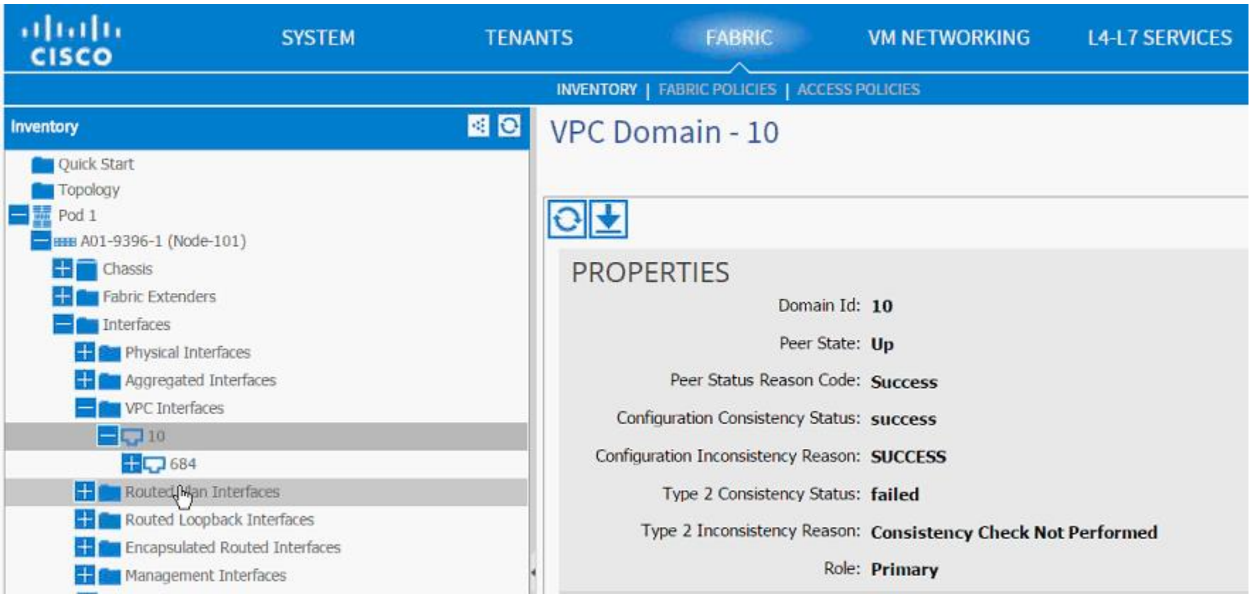

VPC Domain - 10

PROPERTIES

Domain Id: 10

Peer State: Up

Peer Status Reason Code: Success

Configuration Consistency Status: success

Configuration Inconsistency Reason: SUCCESS

Type 2 Consistency Status: failed

Type 2 Inconsistency Reason: Consistency Check Not Performed

Role: Primary

Note: Log into the switch using the console and use show port-channel summary command to verify the port-channel configuration. If Cisco UCS was configured correctly, the port-channel would show UP.

Cisco ACI – Creating vPC for UCS Fabric Interconnect B

To create the vPC for Cisco UCS Fabric Interconnect B, complete the following steps:

1. From the main menu, click FABRIC and select Access Policies.
2. Right-click Interface Policies and select Configure interface, PC and vPC.

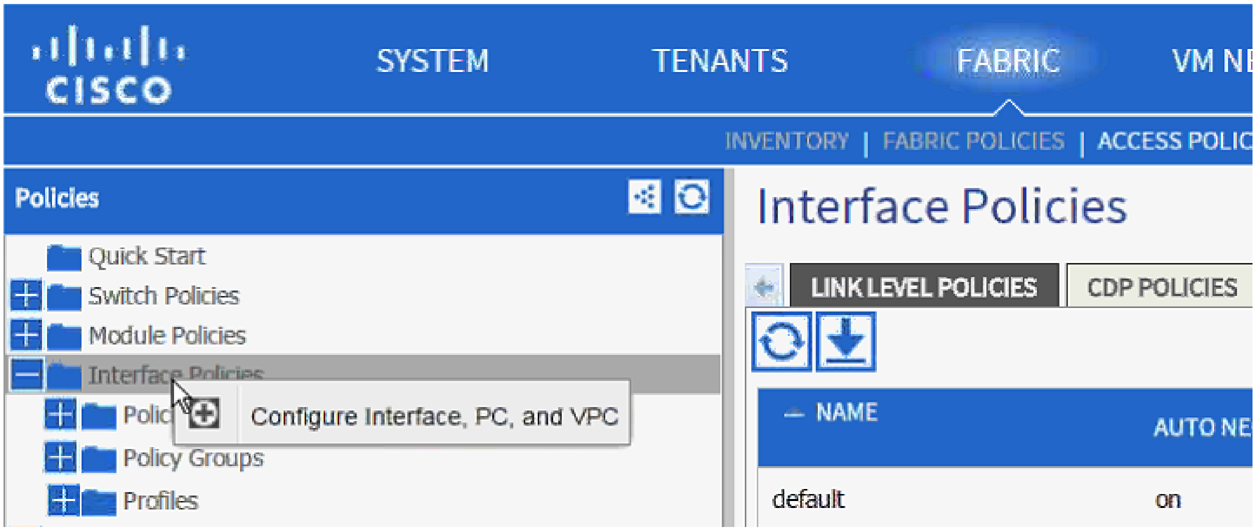

3. In the dialog box, click the + under the CONFIGURED SWITCH INTERFACES.
4. Drop down the menu next to Switches and select both leaves.

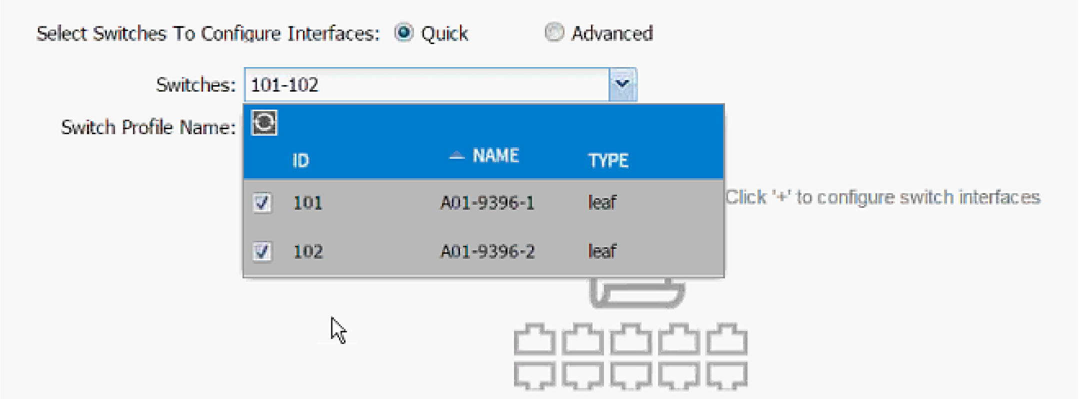
Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name:

ID	NAME	TYPE
☑ 101	A01-9396-1	leaf
☑ 102	A01-9396-2	leaf

5. Enter <sp-UCS-FI-2> as Switch Profile Name. UCS-FI-2 is the host name for UCS Fabric Interconnect B.
6. Click the + sign to add interfaces.
7. Select vPC radio button to configure vPC.
8. Enter 1/20 under Interfaces. This is the port on both switches where Fabric Interconnect B is connected.

9. Enter <ifs- UCS-FI-2> as Interface Selector Name.

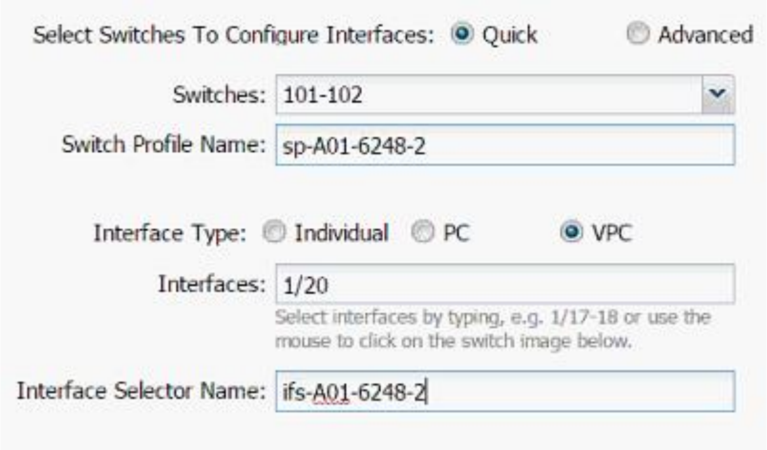

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: sp-A01-6248-2

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/20

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Interface Selector Name: ifs-A01-6248-2

10. Drop down vPC Policy Group and click Create vPC Interface Policy Group. A new dialog box will appear.

11. Enter <pg- UCS-FI-2> as the name of the vPC INTERFACE POLICY GROUP in the dialog box.

12. Select various policy values from the drop-down menus.

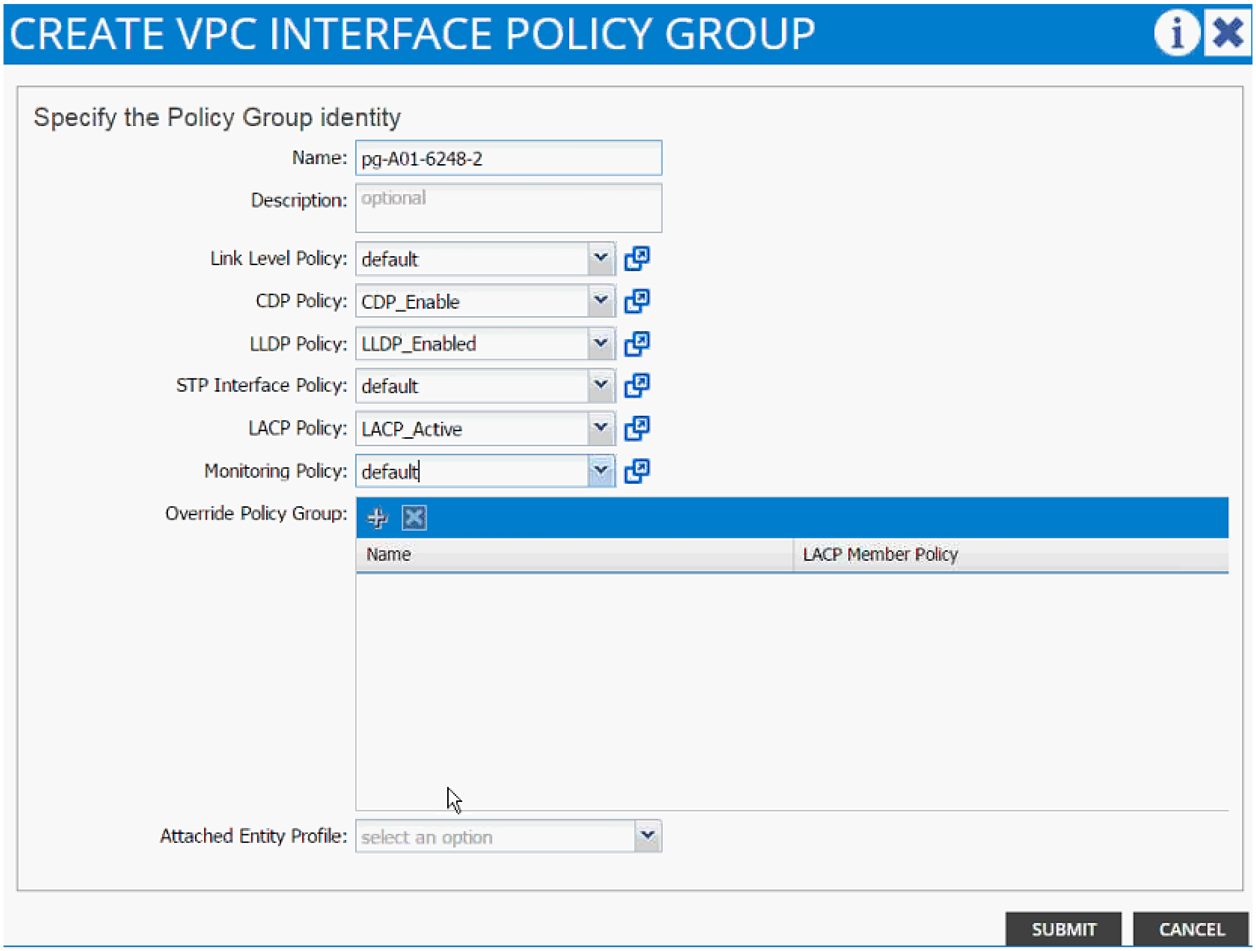

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: pg-A01-6248-2

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Enabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: select an option

13. Drop down the Attached Entity Profile and select shared AEP created in the last section.

14. Click SUBMIT to finish creating the vPC Interface Policy Group.

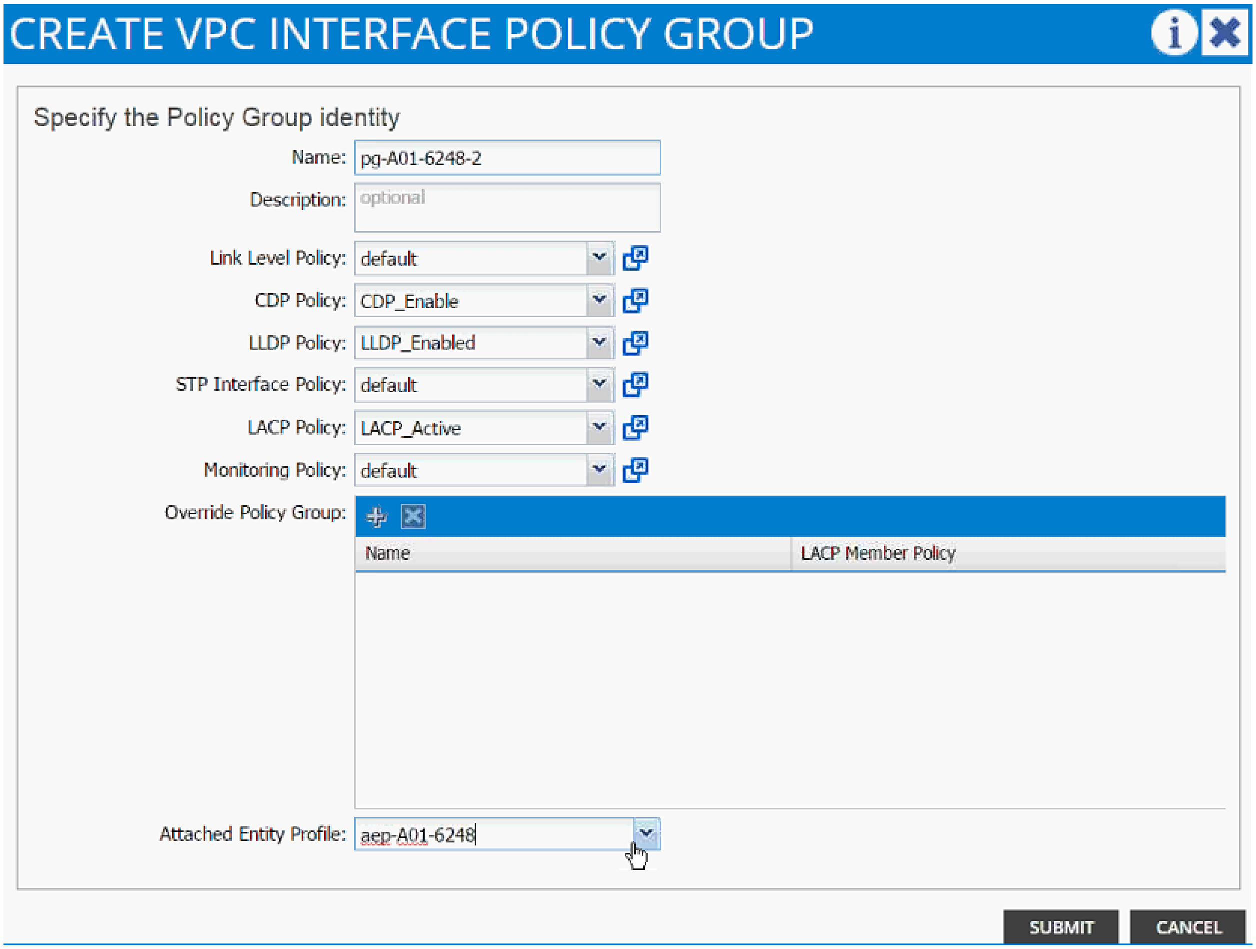

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: pg-A01-6248-2

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Enabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: aep-A01-6248

15. On the Configure Interface, PC, vPC screen, click SAVE.

16. Click SAVE.

17. Click SUBMIT to finish the vPC creation wizard.

18. (Optional) Log into the individual leaf switches and use show port-channel summary and show vpc commands to verify the configuration.

Cisco ACI – Creating vPC for NetApp Controller 1

To create the vPC for NetApp Controller 1, complete the following steps:

1. From the main menu, click FABRIC and select Access Policies.

2. Right-click Interface Policies and select Configure interface, PC and vPC.

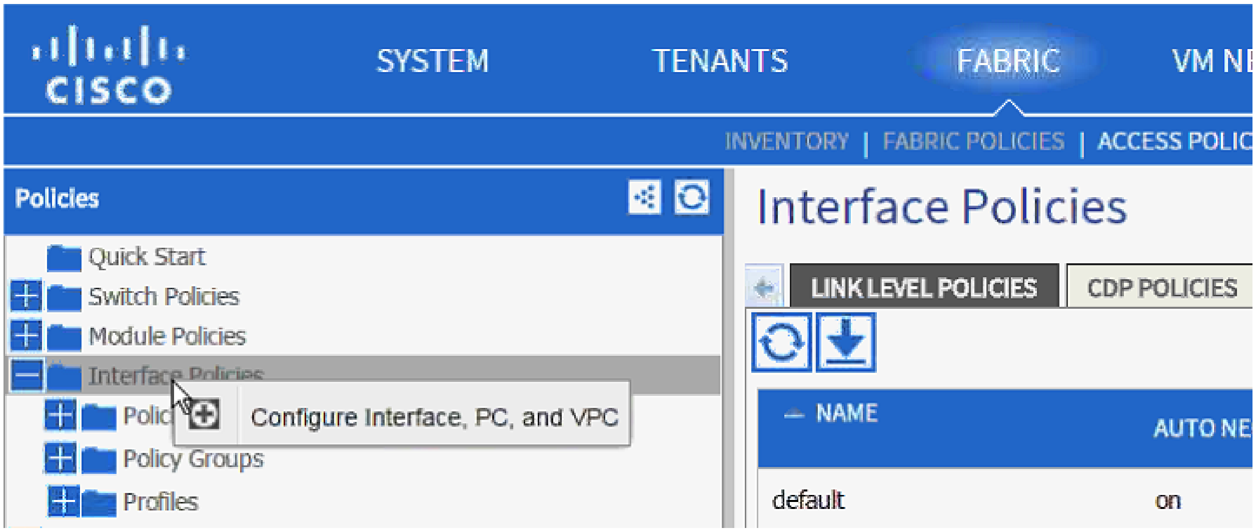

NAME	AUTO NE
default	on

3. In the dialog box, click + under the CONFIGURED SWITCH INTERFACES.

4. Drop down the menu next to Switches and select both leaves.

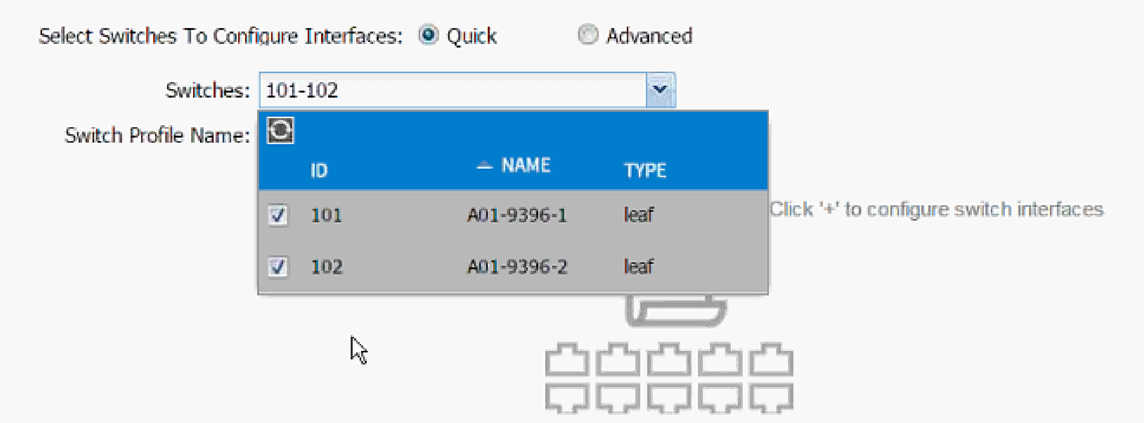

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name:

	ID	NAME	TYPE
☑	101	A01-9396-1	leaf
☑	102	A01-9396-2	leaf

Click '+' to configure switch interfaces

5. Enter <sp-NetAPP-1> as Switch Profile Name. NetApp-1 is the host name for NetApp Controller 1.

6. Click the + sign to add interfaces.

7. Select vPC radio button to configure vPC.

8. Enter 1/17 under Interfaces. This is the port on both switches where NetApp Controller 1 is connected.

9. Enter <ifs- NetApp-1> as Interface Selector Name.

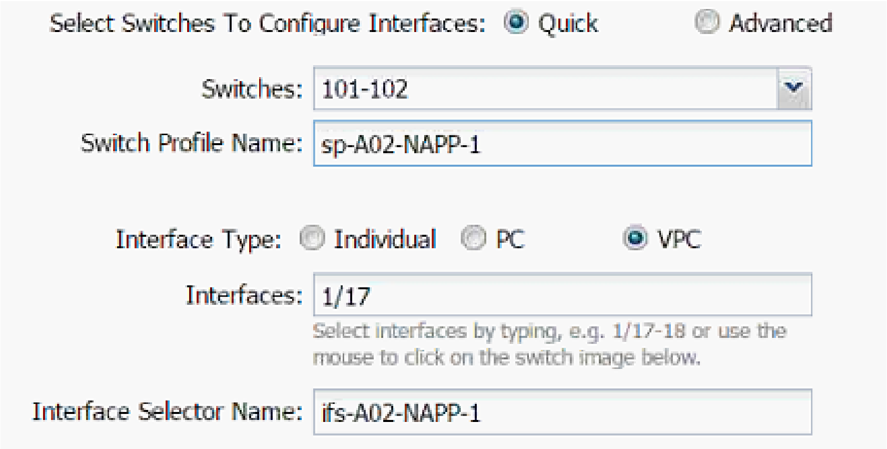

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: sp-A02-NAPP-1

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/17

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Interface Selector Name: ifs-A02-NAPP-1

10. Drop down vPC Policy Group and click Create vPC Interface Policy Group. A new dialog box will appear.

11. Enter <pg- NetApp-1> as the name of the vPC INTERFACE POLICY GROUP in the dialog box.

12. Select various policy values from the drop-down lists.

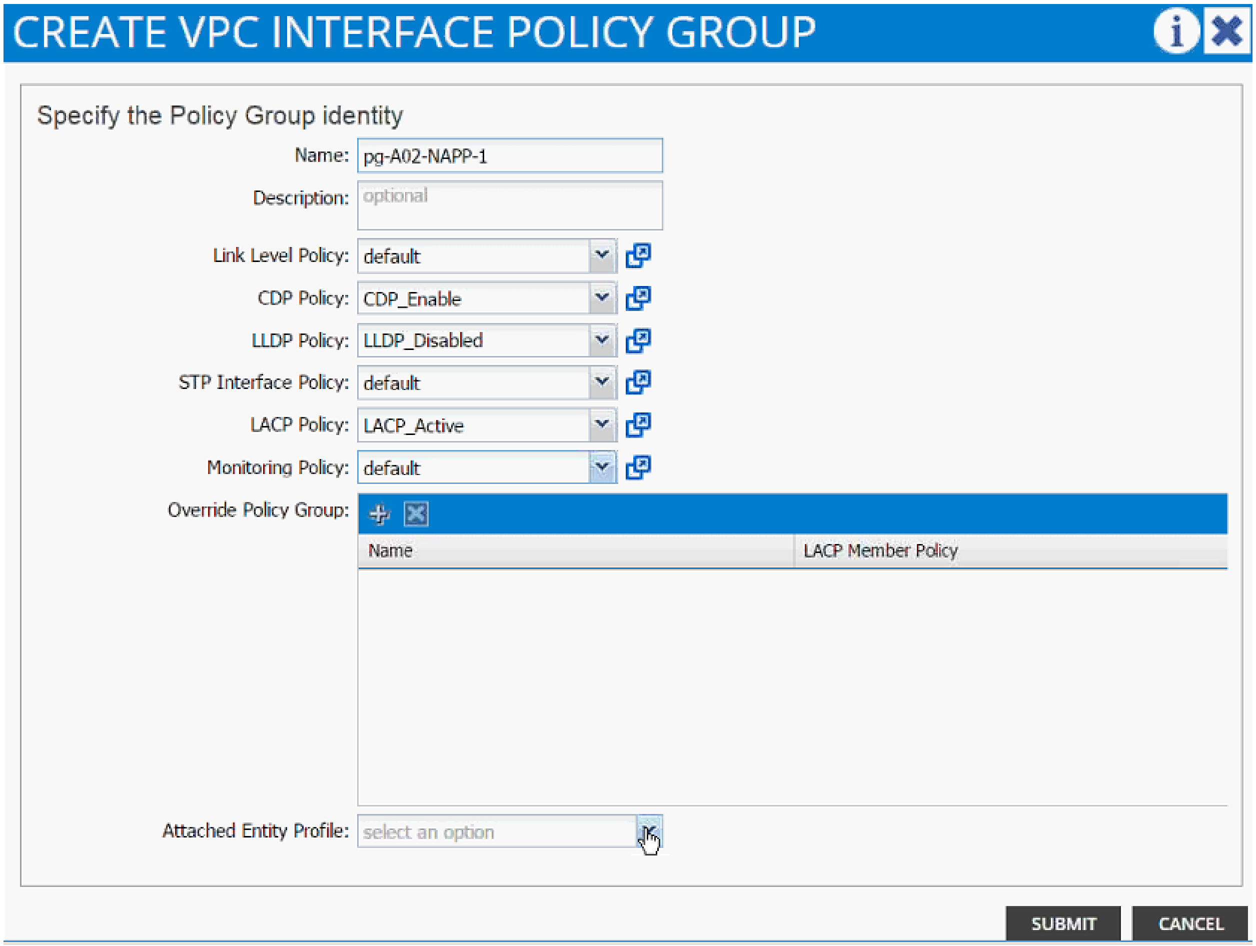

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: pg-A02-NAPP-1

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Disabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: select an option

13. Drop down the Attached Entity Profile and click Create Attachable Access Entity Profile. A new dialog box will appear.

14. The two NetApp Controllers will share the Attachable Entity Profile (AEP). Enter <aep-NetApp-hostname> as the Name. (Avoid using 1 or 2 at the end).

15. Click + to add Domain.

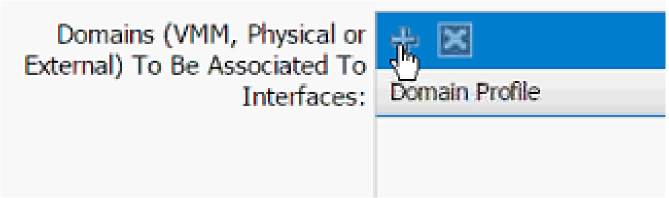

Domains (VMM, Physical or External) To Be Associated To Interfaces:

16. In the added domain, drop-down the menu and select Create Physical Domain.

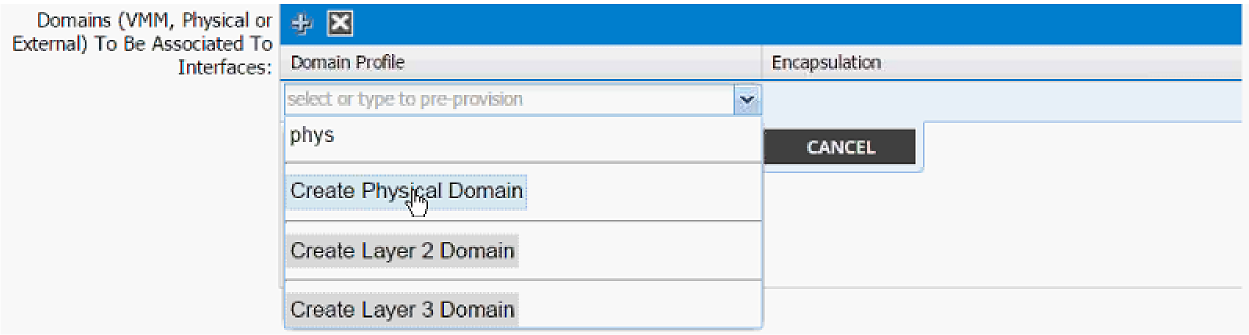

Domains (VMM, Physical or External) To Be Associated To Interfaces:

17. In the Create Physical Domain dialog box, enter <pd-NetApp-hostname> as Name.

18. Drop down the VLAN Pool menu and select Create VLAN Pool.

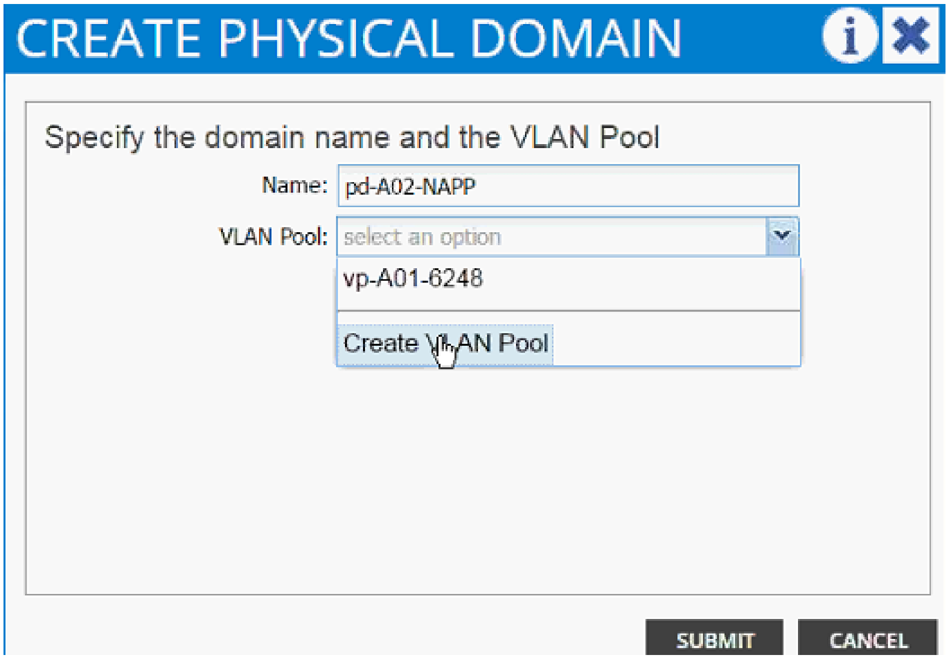

CREATE PHYSICAL DOMAIN

Specify the domain name and the VLAN Pool

Name: pd-A02-NAPP

VLAN Pool: select an option

19. In the Create VLAN Pool dialog box, enter <vp-NetApp-hostname > as Name.

20. Select Allocation Mode Static Allocation.

21. Click + next to Encap Block.

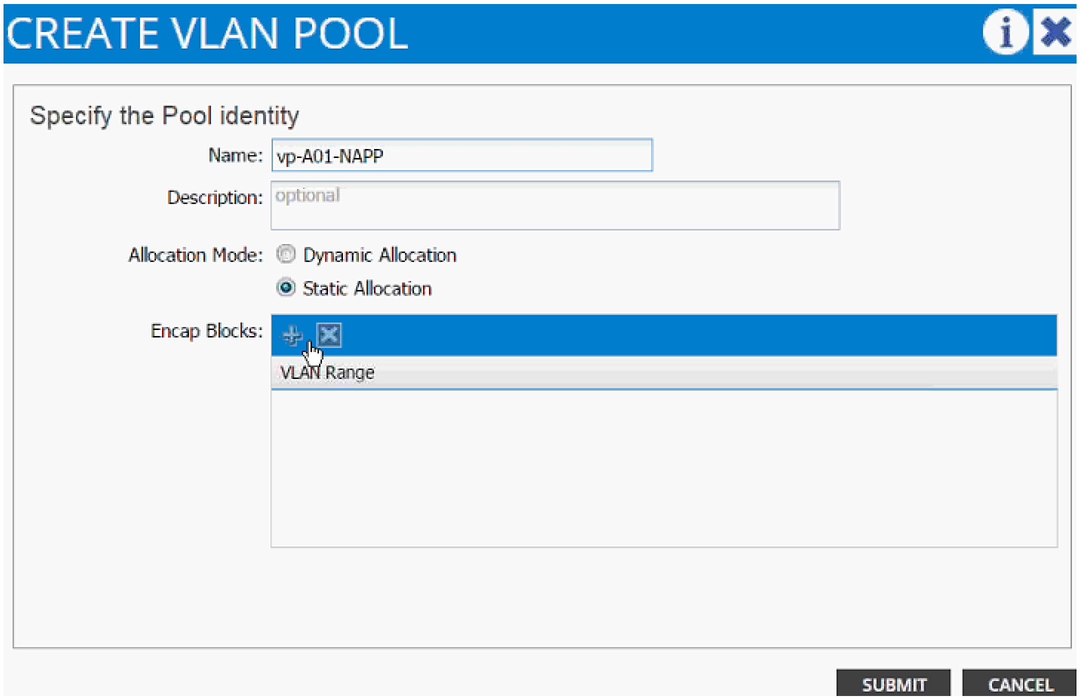

CREATE VLAN POOL

Specify the Pool identity

Name: vp-A01-NAPP

Description: optional

Allocation Mode: ○ Dynamic Allocation ● Static Allocation

Encap Blocks:

22. In the CREATE RANGES dialog box, enter the two iSCSI VLANs and the NFS VLAN.

Note: In the screenshot below, 911, 912 are iSCSI VLANs and 3170 is the NFS VLAN configured on NetApp.

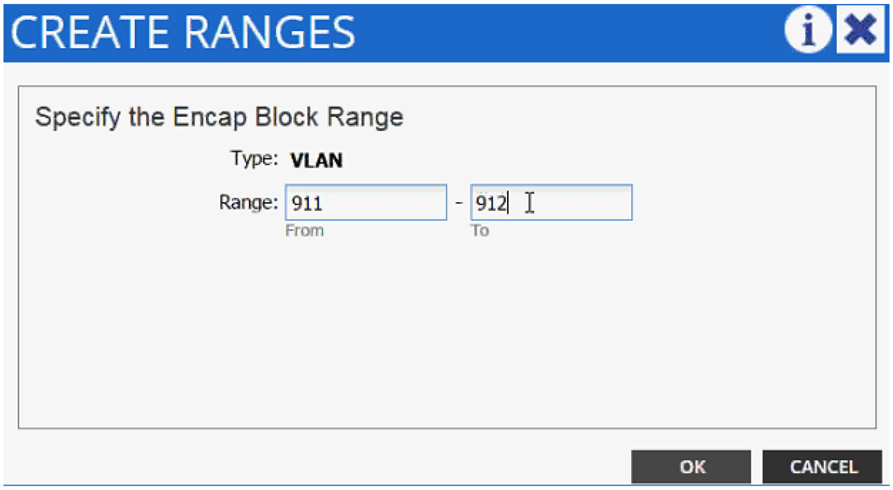

CREATE RANGES

Specify the Encap Block Range

Type: **VLAN**

Range: 911 - 912

From To

23. Click “OK.

24. Click “+ to add NFS VLAN and in the “CREATE RANGES dialog box, enter range 3170-3170 (single VLAN).

25. Click OK.

26. Click SUBMIT to finish VLAN pool creation.

27. Click SUBMIT to finish Physical Domain creation.

28. Click UPDATE to finish adding Physical domain to AEP.

29. Click SUBMIT to finish adding AEP.

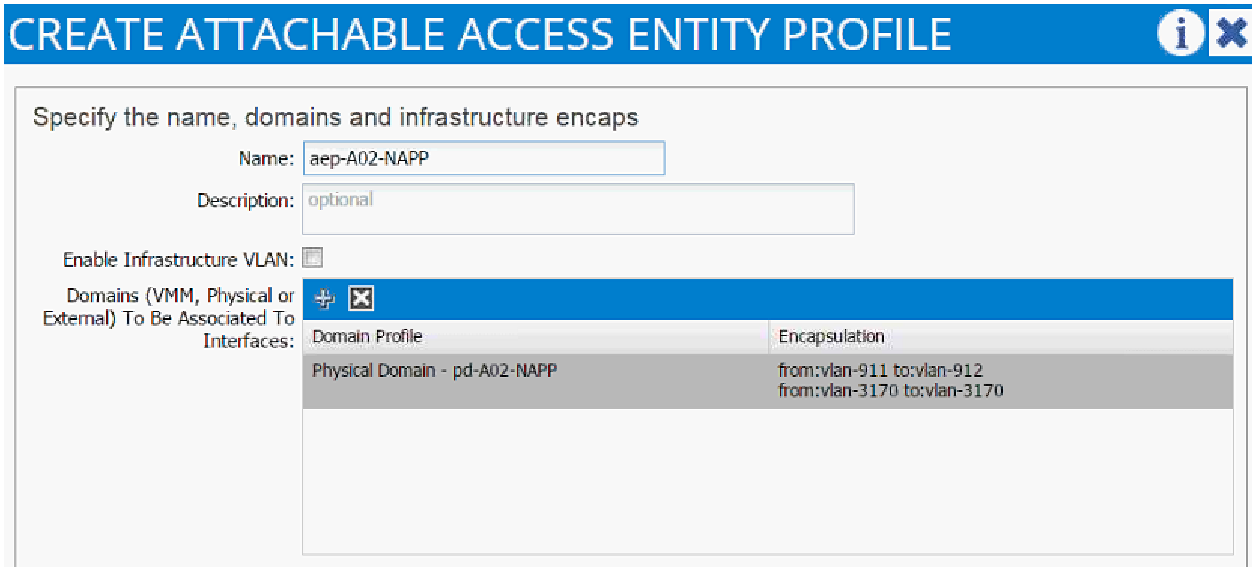

CREATE ATTACHABLE ACCESS ENTITY PROFILE

Specify the name, domains and infrastructure encaps

Name: aep-A02-NAPP

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation

Physical Domain - pd-A02-NAPP	from:vlan-911 to:vlan-912
	from:vlan-3170 to:vlan-3170

30. Click SUBMIT to finish creating vPC Interface Policy Group.

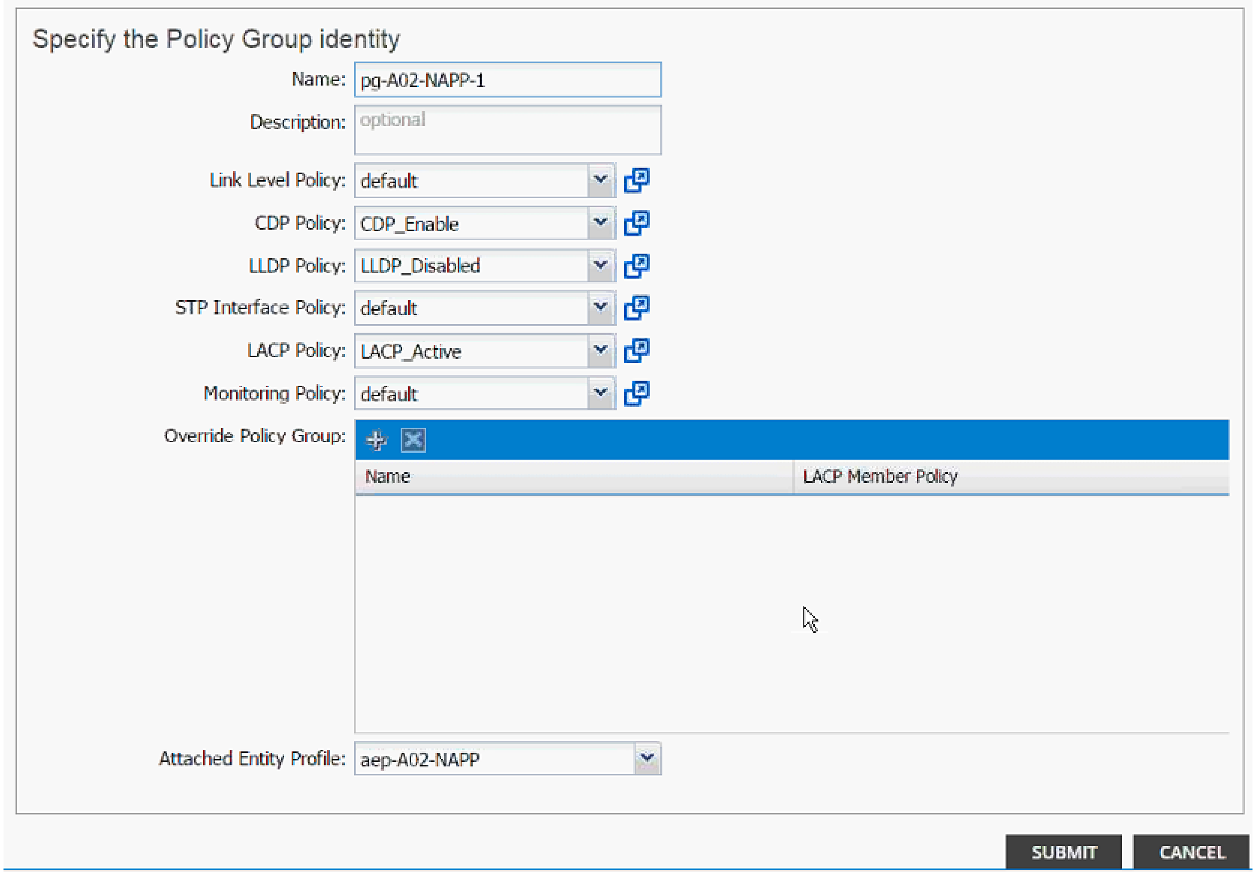

Specify the Policy Group identity

Name: pg-A02-NAPP-1

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Disabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: aep-A02-NAPP

31. On the Configure Interface, PC, vPC screen, click SAVE.

32. Click SAVE.

33. Click SUBMIT to finish the vPC creation using wizard.

Cisco ACI – Creating vPC for NetApp Controller 2

To create the vPC for NetApp Controller 2, complete the following steps:

1. From the main menu, click FABRIC and select Inventory.

2. Expand Pod 1 and right click on the first leaf and select Configure interface, PC and vPC.

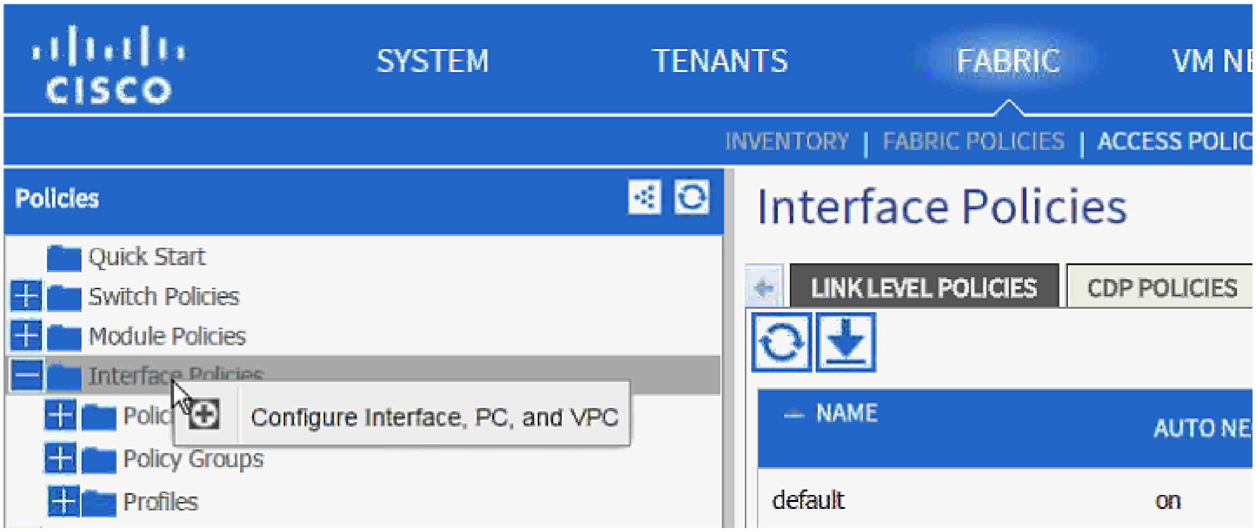

3. In the dialog box, click + under the CONFIGURED SWITCH INTERFACES.

4. From the Switches drop-down list, select both leaves.

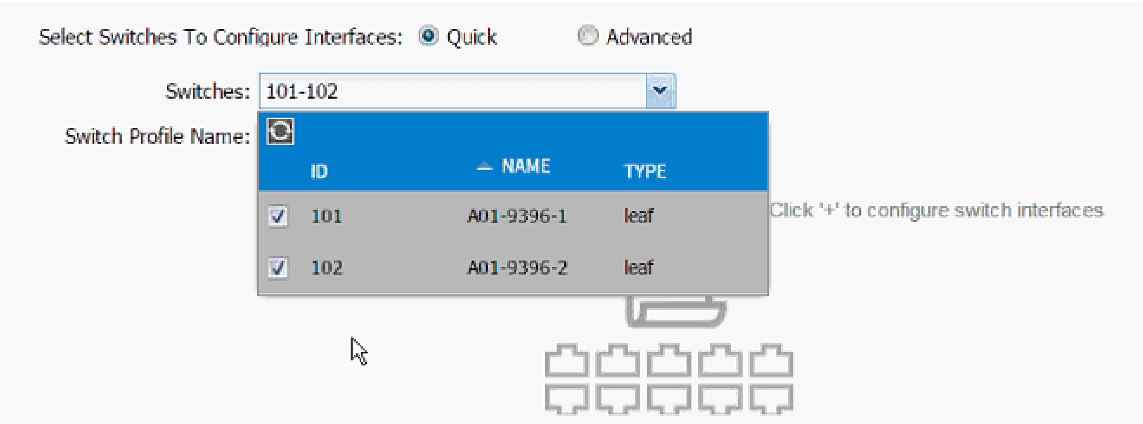

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name:

	ID	NAME	TYPE
☑	101	A01-9396-1	leaf
☑	102	A01-9396-2	leaf

Click '+' to configure switch interfaces

5. Enter <sp-NetAPP-2> as Switch Profile Name. NetApp-2 is the host name for NetApp Controller 2.

6. Click the + sign to add interfaces.

7. Select vPC radio button to configure vPC.

8. Enter 1/18 under Interfaces. This is the port on both switches where NetApp Controller 2 is connected.

9. Enter <ifs- NetApp-2> as Interface Selector Name.

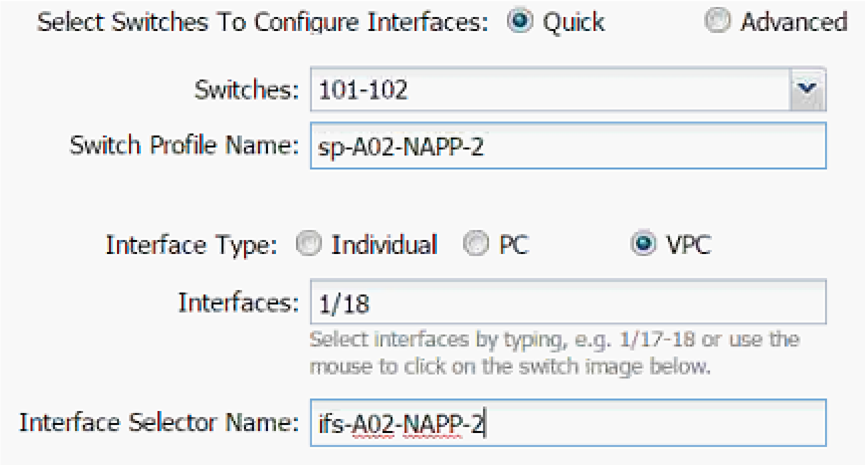

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: sp-A02-NAPP-2

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/18

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Interface Selector Name: ifs-A02-NAPP-2

10. Drop down vPC Policy Group and click Create vPC Interface Policy Group. A new dialog box will appear.

11. Enter <pg- NetApp-2> as the name of the vPC INTERFACE POLICY GROUP in the dialog box.

12. Select various policy values from the drop-down lists.

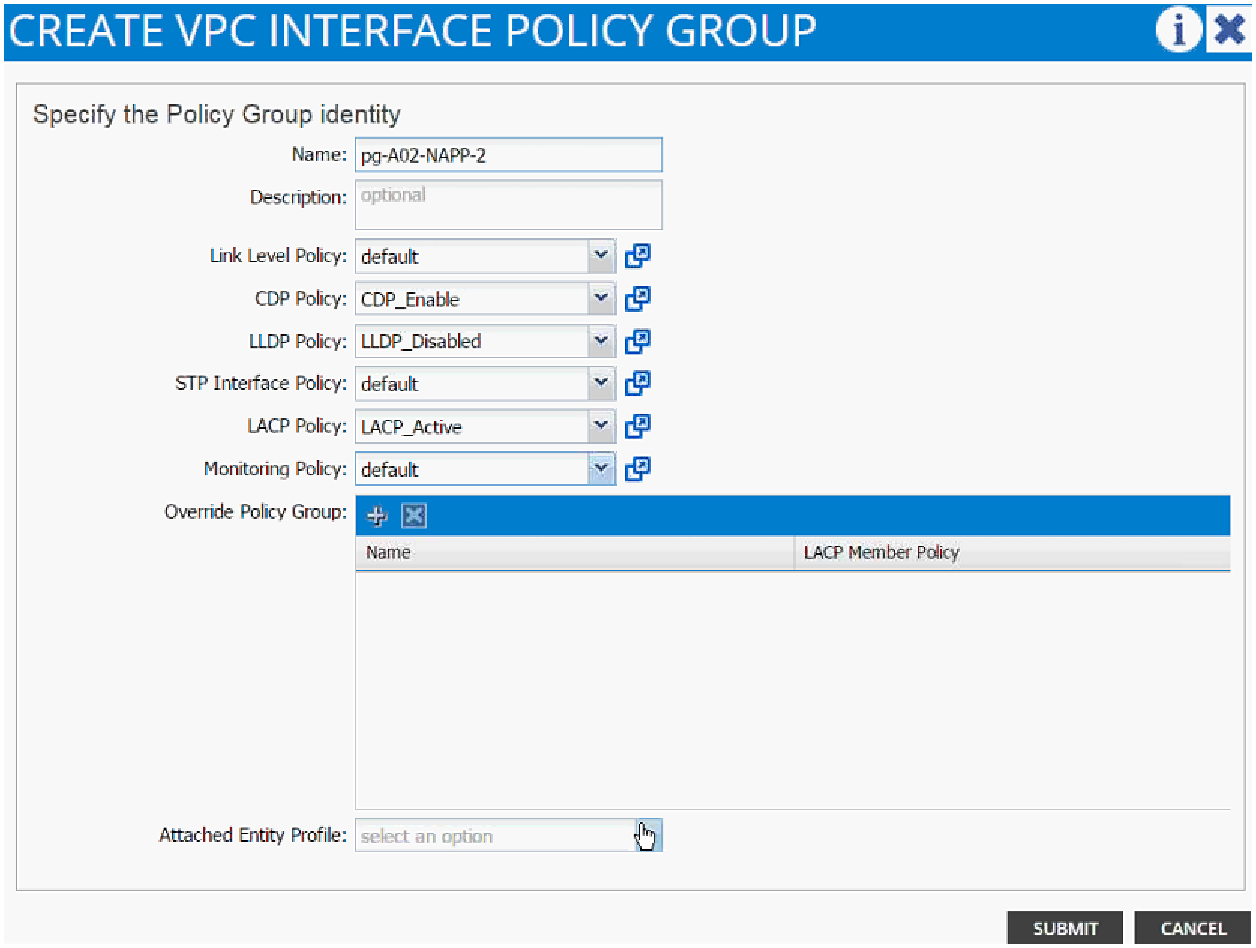

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: pg-A02-NAPP-2

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Disabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: select an option

13. Drop down the Attached Entity Profile select shared AEP created in the last section.

14. Click SUBMIT to finish creating vPC Interface Policy Group.

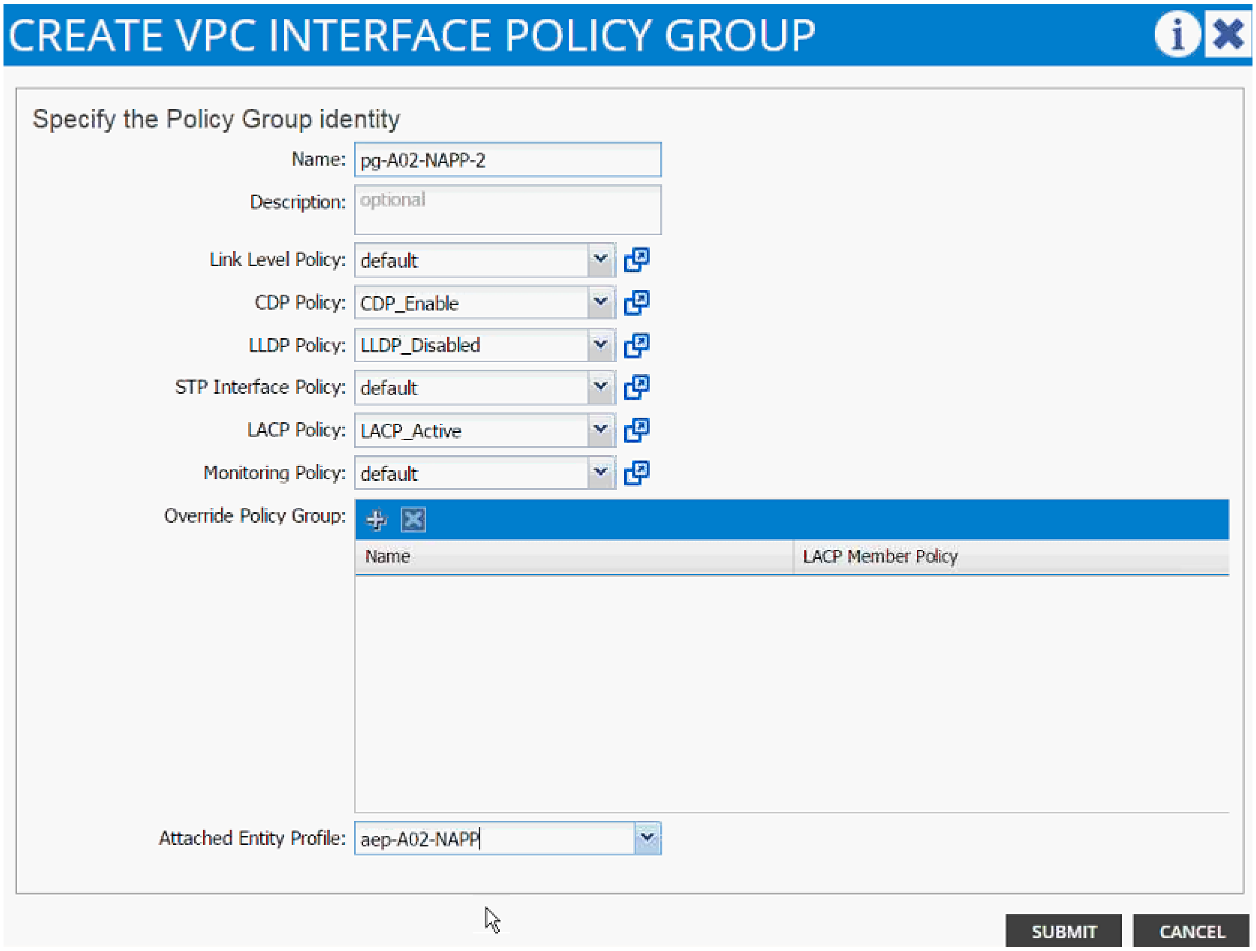

CREATE VPC INTERFACE POLICY GROUP

Specify the Policy Group identity

Name: pg-A02-NAPP-2

Description: optional

Link Level Policy: default

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Disabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: aep-A02-NAPP

15. On the Configure Interface, PC, vPC screen, click SAVE.

16. Click SAVE.

17. Click SUBMIT to finish the vPC creation.

18. Expand the Pod 1 followed by the Leaf switch(s) on the left hand menu bar and then expand the Interfaces and vPC Interfaces.

19. Expand the vPC domain 10 and validate all the vPCs are configured and up (you will not see the VLAN being forwarded at this time).

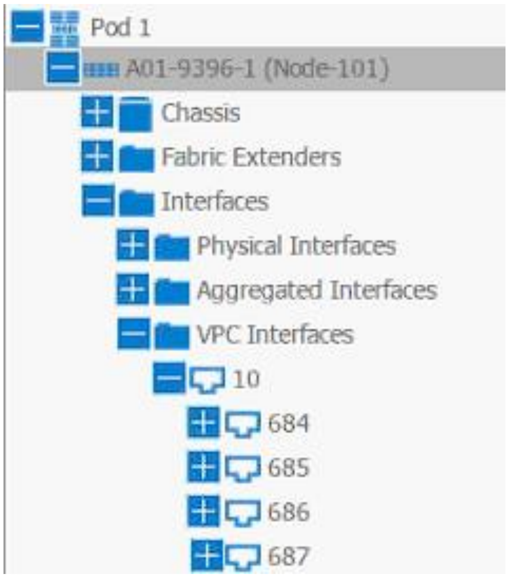

20. Optional: Log into the leaf switches using console and validate the port-channels are configured correctly. The output below assumes the NetApp and UCS port-channel configurations are in place.

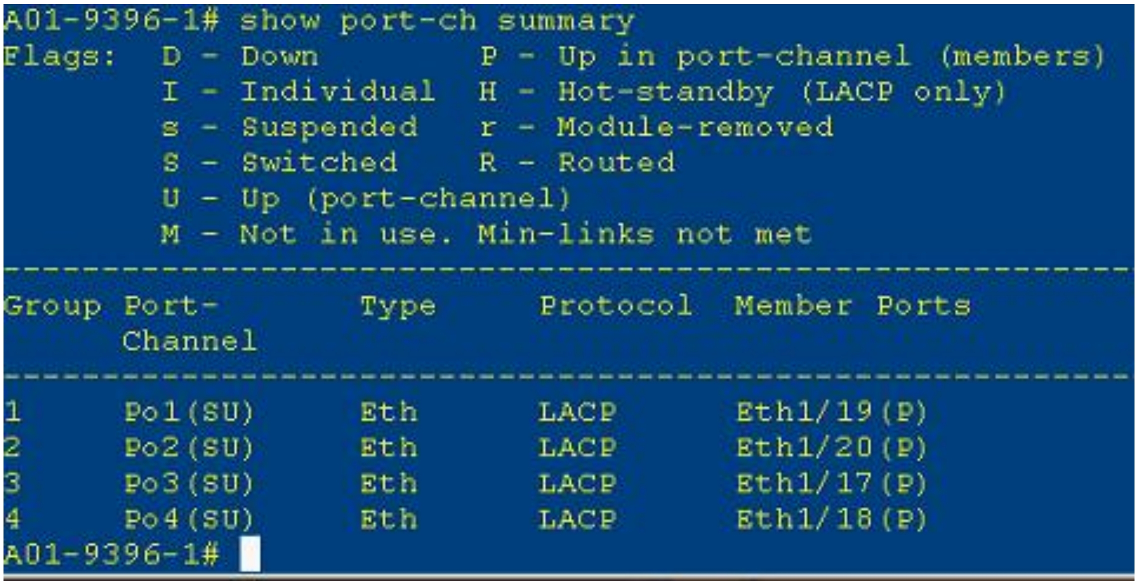

```
A01-9396-1# show port-ch summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/19(P)
2     Po2(SU)     Eth      LACP      Eth1/20(P)
3     Po3(SU)     Eth      LACP      Eth1/17(P)
4     Po4(SU)     Eth      LACP      Eth1/18(P)
A01-9396-1# 
```

Cisco ACI – Deploying Infrastructure (Foundation) Tenant

In this section, a new tenant will be deployed to host the infrastructure connectivity between the compute (VMware) and Storage (NetApp) environments. To deploy a new tenant, complete the following steps:

1. From the main menu, click TENANTS and from the sub-menu click ADD TENANT.

2. In the CREATE TENANT dialog box, type Foundation as the name of the tenant.

3. Click the checkbox next to all under Security Domains.

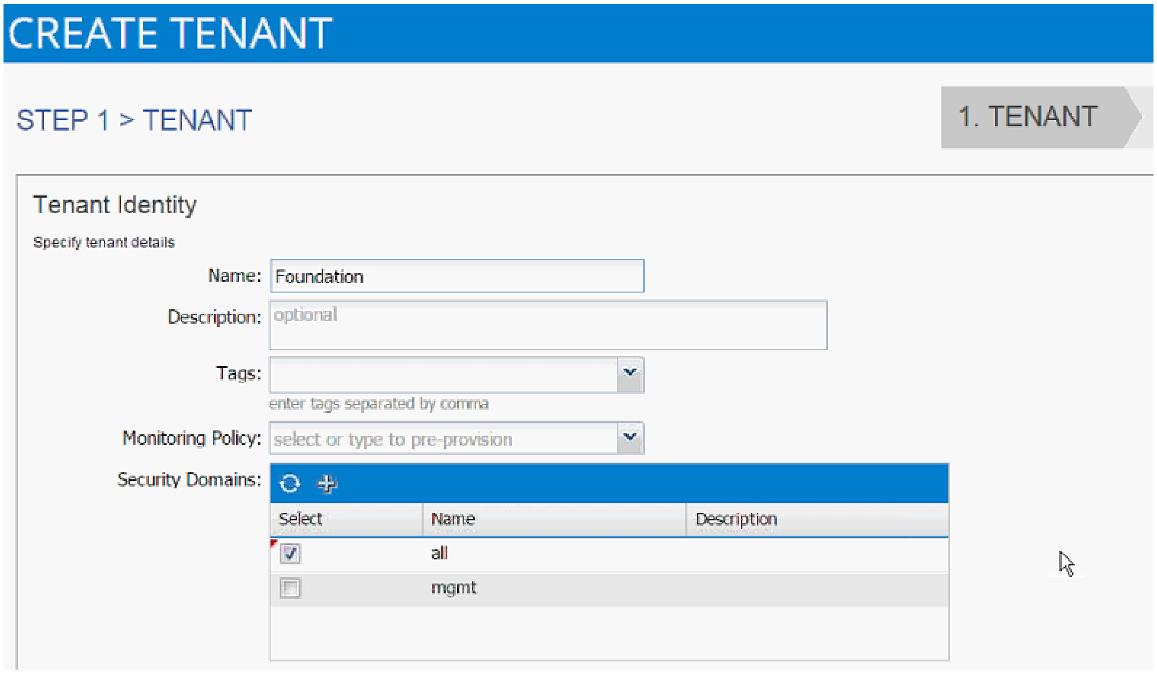

CREATE TENANT

STEP 1 > TENANT

1. TENANT

Tenant Identity

Specify tenant details

Name: Foundation

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: select or type to pre-provision

Security Domains:

Select	Name	Description
☑	all	
☐	mgmt	

4. Click Next.

5. Click + sign to add network.

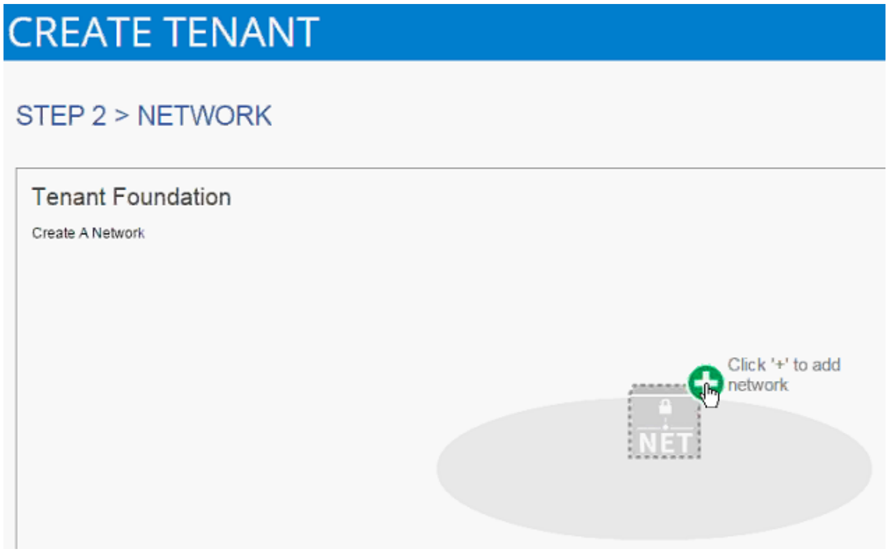

CREATE TENANT

STEP 2 > NETWORK

Tenant Foundation

Create A Network

6. In the CREATE NEW NETWORK dialog box, type Foundation as the Name. Leave everything else as default.

7. Click Next to move onto bridge domain creation.

8. Use bd-Internal as the Name of the bridge domain.

9. Drop down the menu next to Forwarding and select Custom.

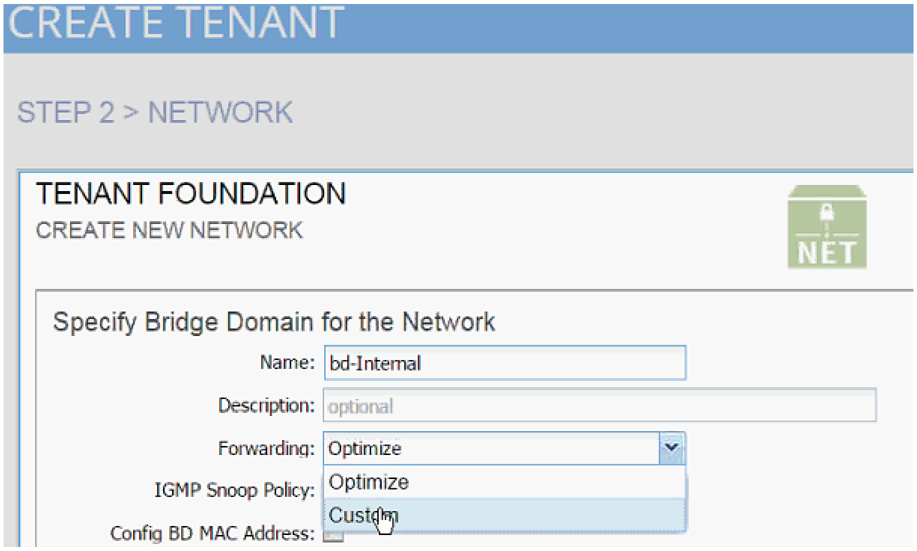

CREATE TENANT

STEP 2 > NETWORK

TENANT FOUNDATION

CREATE NEW NETWORK

Specify Bridge Domain for the Network

Name: bd-Internal

Description: optional

Forwarding: Optimize

IGMP Snoop Policy:

Config BD MAC Address:

10. Check the boxes to enable Flooding and Routing.

11. Select default for IGMP Snoop Policy.

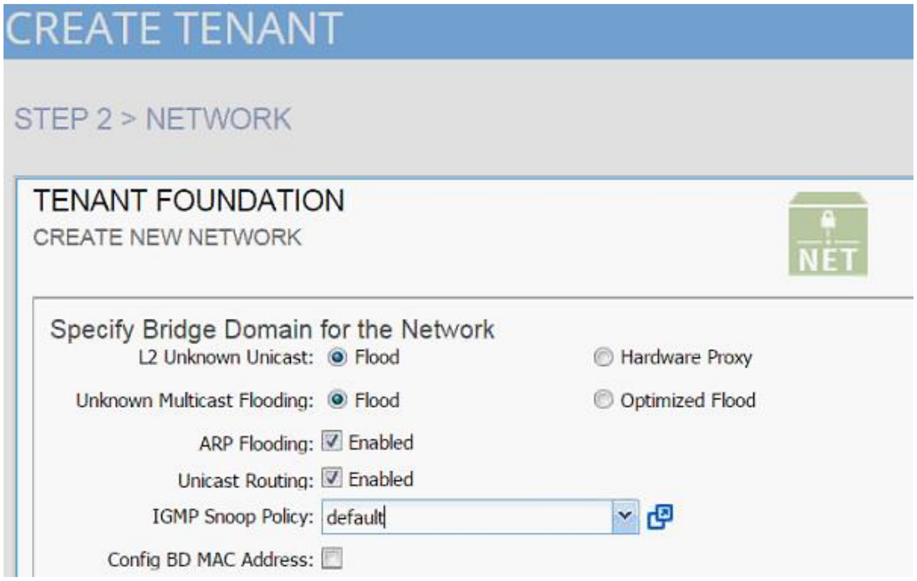

CREATE TENANT

STEP 2 > NETWORK

TENANT FOUNDATION

CREATE NEW NETWORK

Specify Bridge Domain for the Network

L2 Unknown Unicast: ● Flood ○ Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: ☑ Enabled

Unicast Routing: ☑ Enabled

IGMP Snoop Policy: default

Config BD MAC Address: □

12. Click OK.

13. Click + sign next to Bridge Domain to add another Bridge Domain.

14. Add bd-iSCSI-a as the Name.

15. Select Custom > Forwarding.

16. Enable Flooding and disable Unicast routing.

17. Set IGMP Snoop Policy to default.

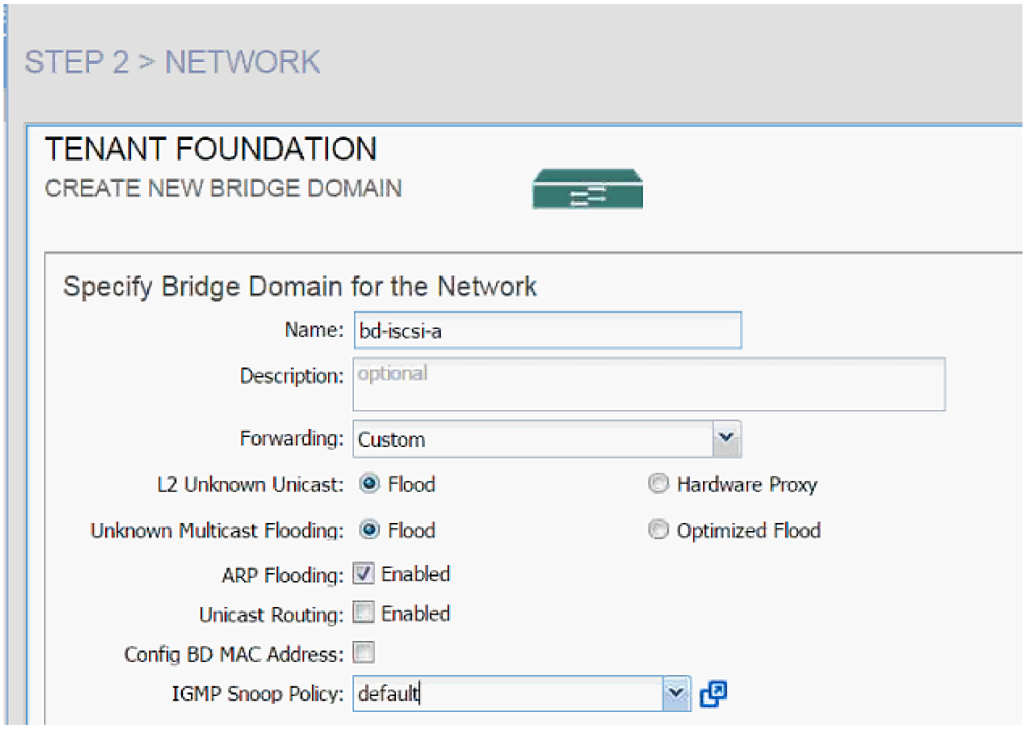

STEP 2 > NETWORK

TENANT FOUNDATION

CREATE NEW BRIDGE DOMAIN

Specify Bridge Domain for the Network

Name: bd-iscsi-a

Description: optional

Forwarding: Custom

L2 Unknown Unicast: ● Flood ○ Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: ☑ Enabled

Unicast Routing: □ Enabled

Config BD MAC Address: □

IGMP Snoop Policy: default

18. Click Next.

19. Click OK.

20. Click + sign next to Bridge Domain to add another Bridge Domain.

21. Add bd-iSCSI-b as the Name.

22. Select Custom Forwarding.

23. Enable Flooding and disable Unicast routing.

24. Set IGMP Snoop Policy to default.

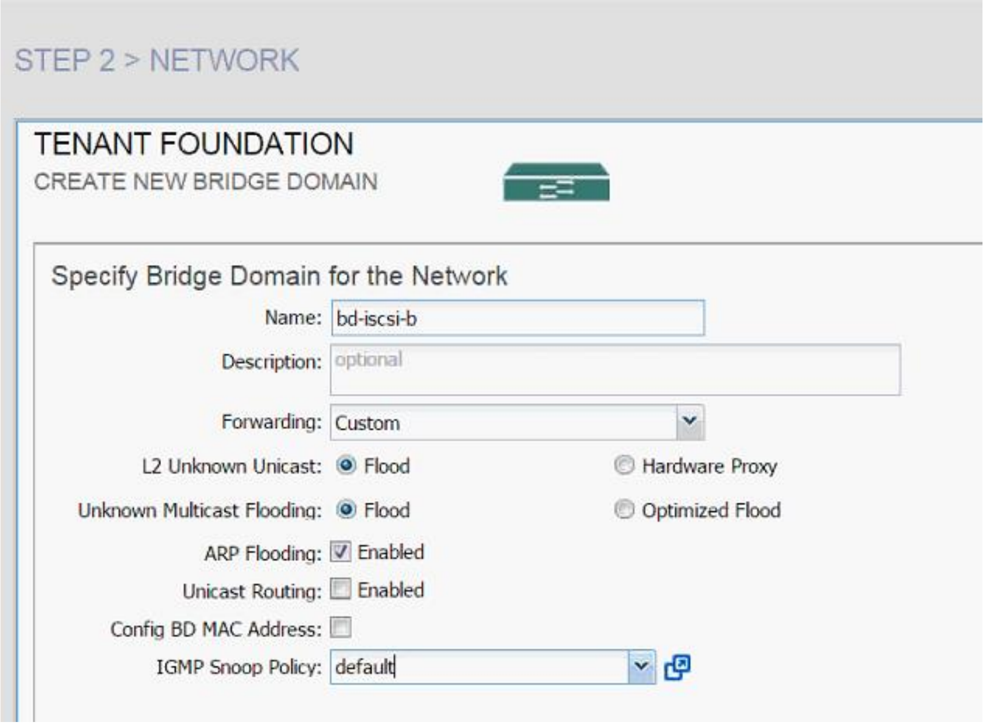

Specify Bridge Domain for the Network

Name: bd-iscsi-b

Description: optional

Forwarding: Custom

L2 Unknown Unicast: ● Flood ○ Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: ☑ Enabled

Unicast Routing: □ Enabled

Config BD MAC Address: □

IGMP Snoop Policy: default

25. Click Next.

26. Click OK.

27. Three Bridge Domains and Foundation network should be visible in the CREATE TENANT dialog box.

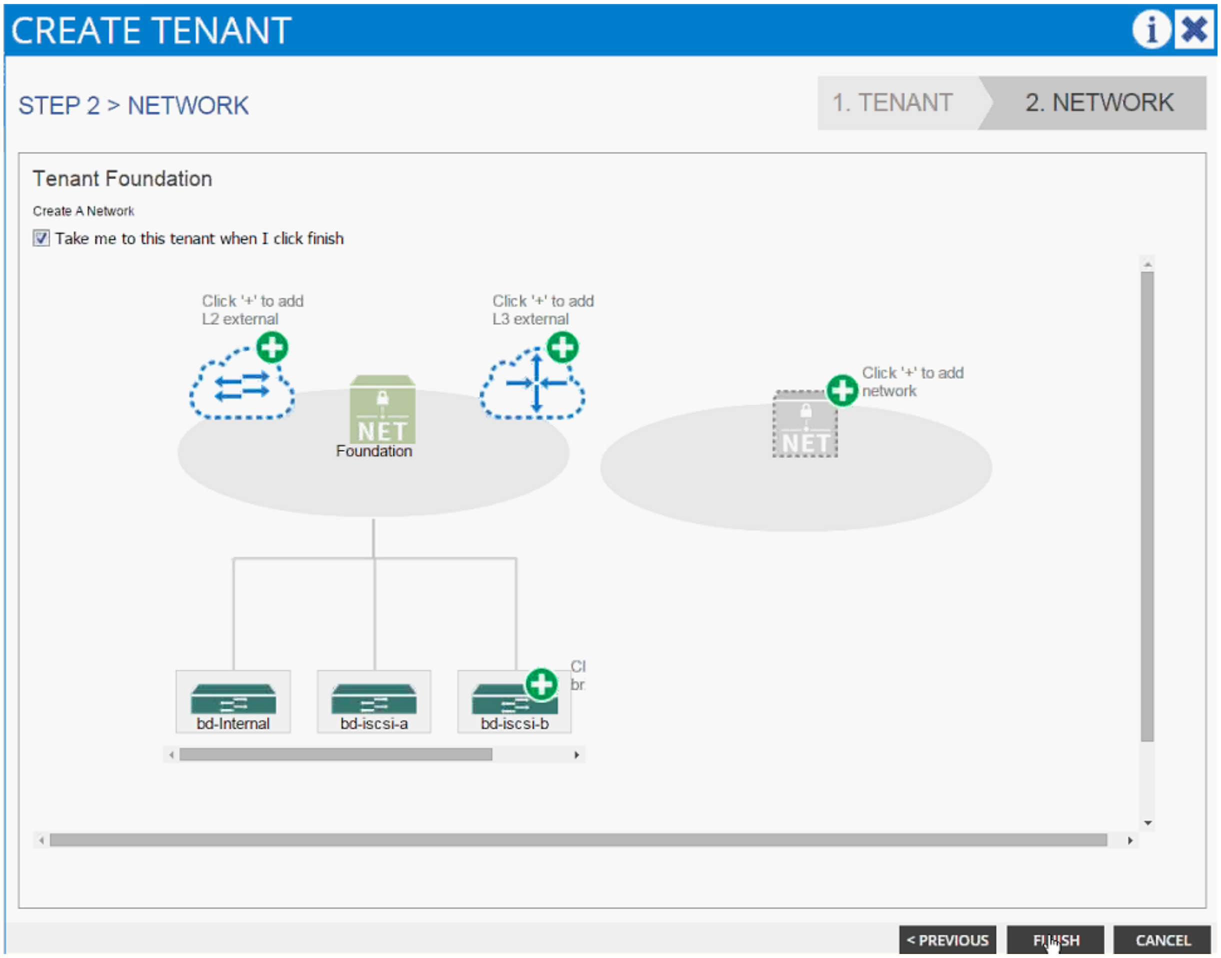

28. Click Finish.

29. Verify the selected tenant is the newly created Foundation tenant by looking at the items highlighted in the top menu.

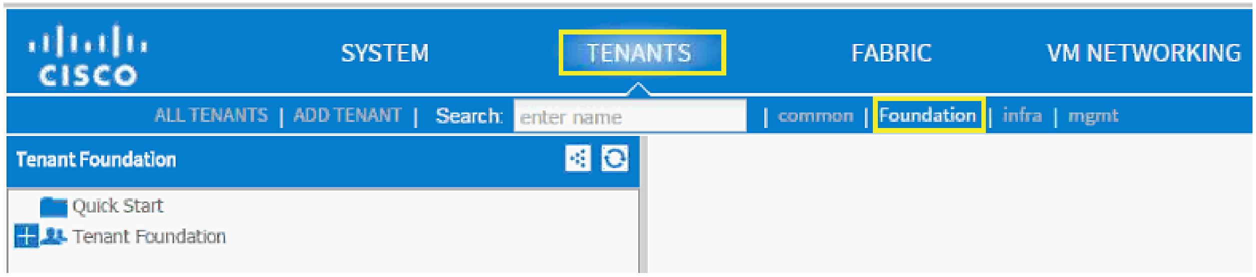
SYSTEM
TENANTS
FABRIC
VM NETWORKING
ALL TENANTS | ADD TENANT | Search: enter name | common | Foundation | infra | mgmt
Tenant Foundation
Quick Start
Tenant Foundation

Application Profile Creation

In this section, two Application Profiles, iSCSI and NFS will be created.

iSCSI Application Profile Creation

1. Select Tenant and newly created Foundation tenant from the top menu.

2. Expand Tenant Foundation in the left menu bar.

3. Right-click Application Profile and click Create Application Profiles.

4. In the CREATE APPLICATION PROFILE dialog box, enter iSCSI as the Name.

5. From the drop-down menu, select default for Monitoring Policy.

6. Click + next to EPG to add an EPG.

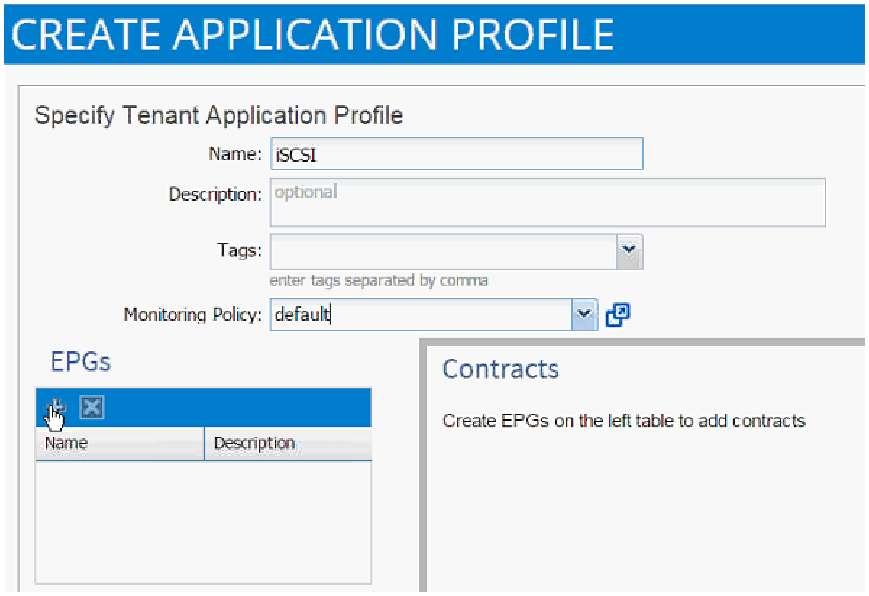

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: iSCSI

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

Create EPGs on the left table to add contracts

7. In the CREATE APPLICATION EPG dialog box, enter iscsi-a-lif as the Name.

8. From the drop-down menu, select bd-iscsi-a as the Bridge Domain.

9. From the drop-down menu, select default for Monitoring Policy.

10. Click FINISH.

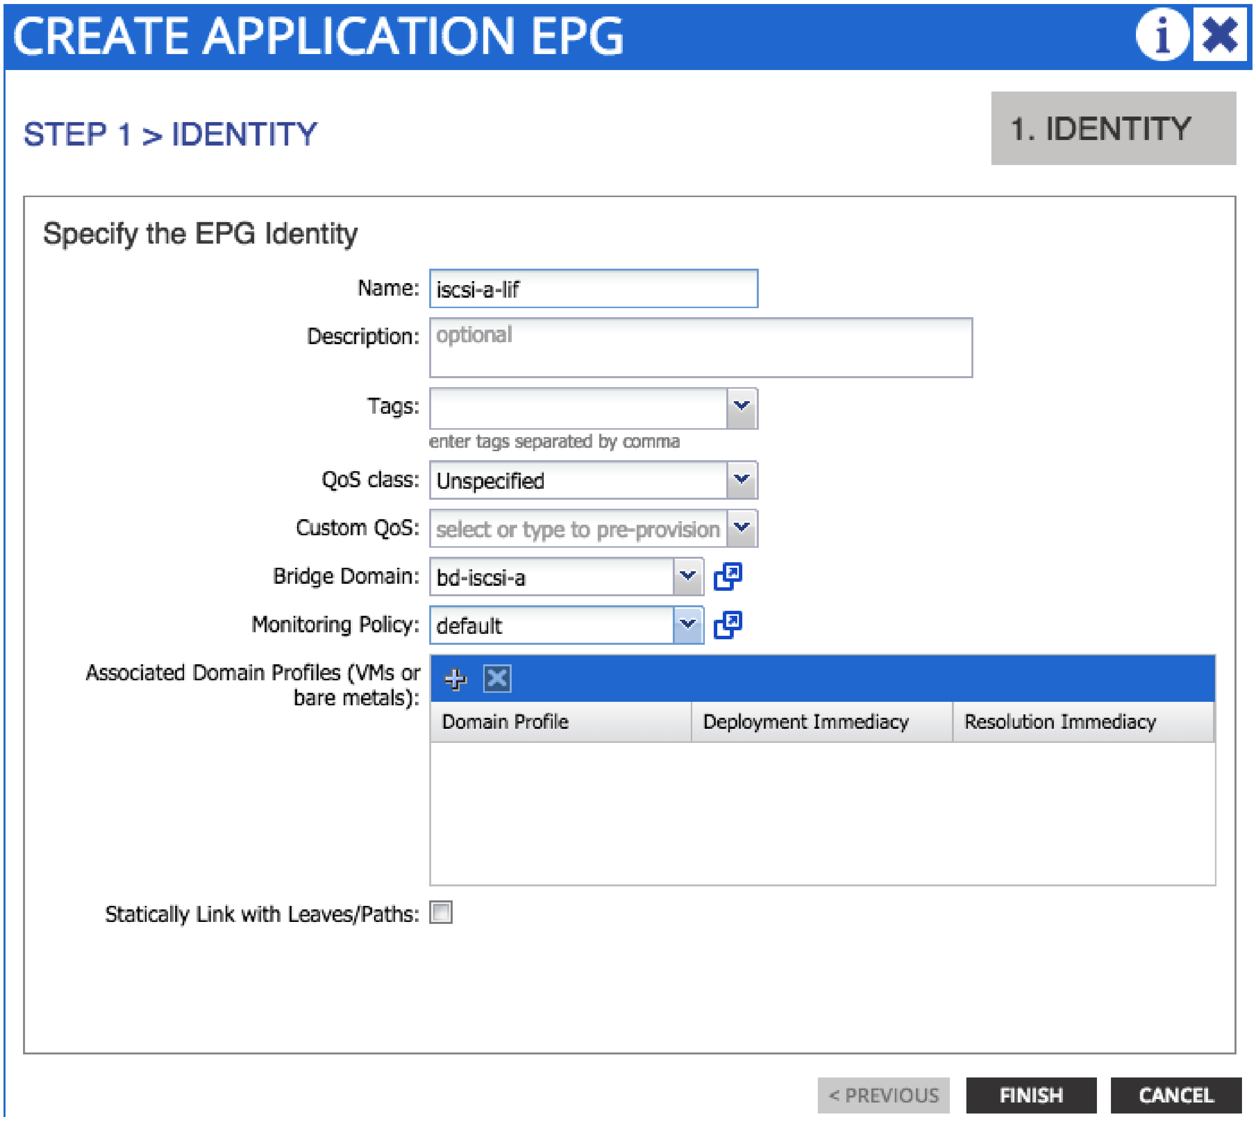

CREATE APPLICATION EPG

STEP 1 > IDENTITY

Specify the EPG Identity

Name: iscsi-a-lif

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-iscsi-a

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy

Statically Link with Leaves/Paths: □

11. Click + next to EPG to add another EPG.

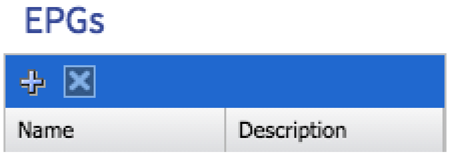

EPGs

Name	Description

12. In the CREATE APPLICATION EPG dialog box, enter iscsi-b-lif as the Name.

13. From the drop-down menu, select bd-iscsi-b as the Bridge Domain.

14. From the drop-down menu, select default Monitoring Policy.

15. Click FINISH.

16. Click + next to EPG to add another EPG.

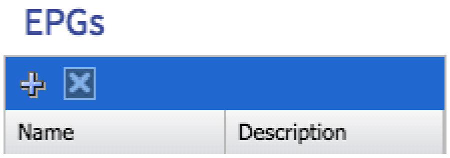

17. In the CREATE APPLICATION EPG dialog box, enter iscsi-a-vmk as the Name.

18. From the drop-down menu, select bd-iscsi-a as the Bridge Domain.

19. From the drop-down menu, select default Monitoring Policy.

20. Click FINISH.

21. Click + next to EPG to add another EPG.

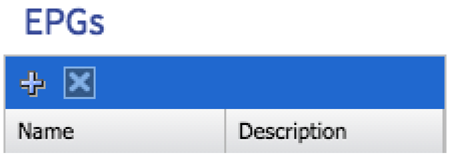

22. In the CREATE APPLICATION EPG dialog box, enter iscsi-b-vmk as the Name.

23. From the drop-down menu, select bd-iscsi-b as the Bridge Domain.

24. From the drop-down menu, select default Monitoring Policy.

25. Click FINISH.

26. Click SUBMIT to finish creating the Application Profile.

Setting up EPG iscsi-a-lif

To set up the EPG iscsi-a-lif, complete the following steps:

1. Expand the newly created iSCSI Application profile from the menu bar on the left.

2. Expand iSCSI, expand Application EPGs and expand EPG iscsi-a-lif.

3. Click Static Bindings (Paths).

4. Click Action on the right hand work area.

5. Click Deploy Static EPG on PC, vPC, or Interface.

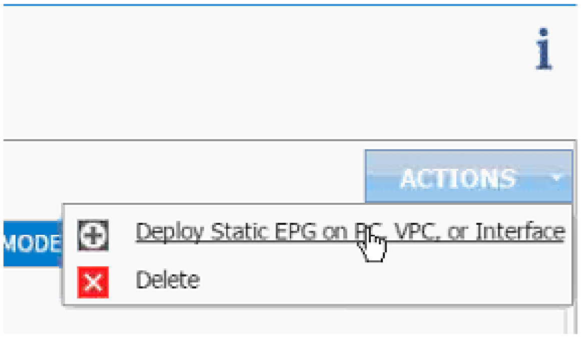

6. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

7. From the drop-down menu Path, select NetApp Controller 1.

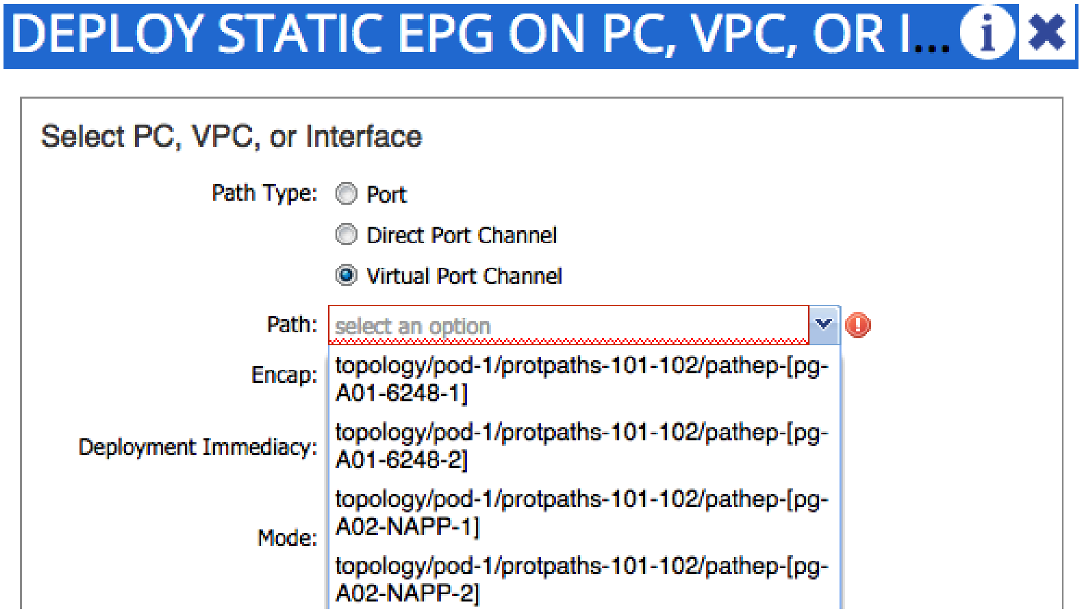

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port
○ Direct Port Channel
● Virtual Port Channel

Path: select an option

- topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]
- topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]
- topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]
- topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]

Encap:

Deployment Immediacy:

Mode:

8. Enter vlan-<<var_iscsi_vlan_A_id >> for Encap (VLAN 901 is the iSCSI-A VLAN in the screen capture below).

9. Change Deployment Immediacy to Immediate.

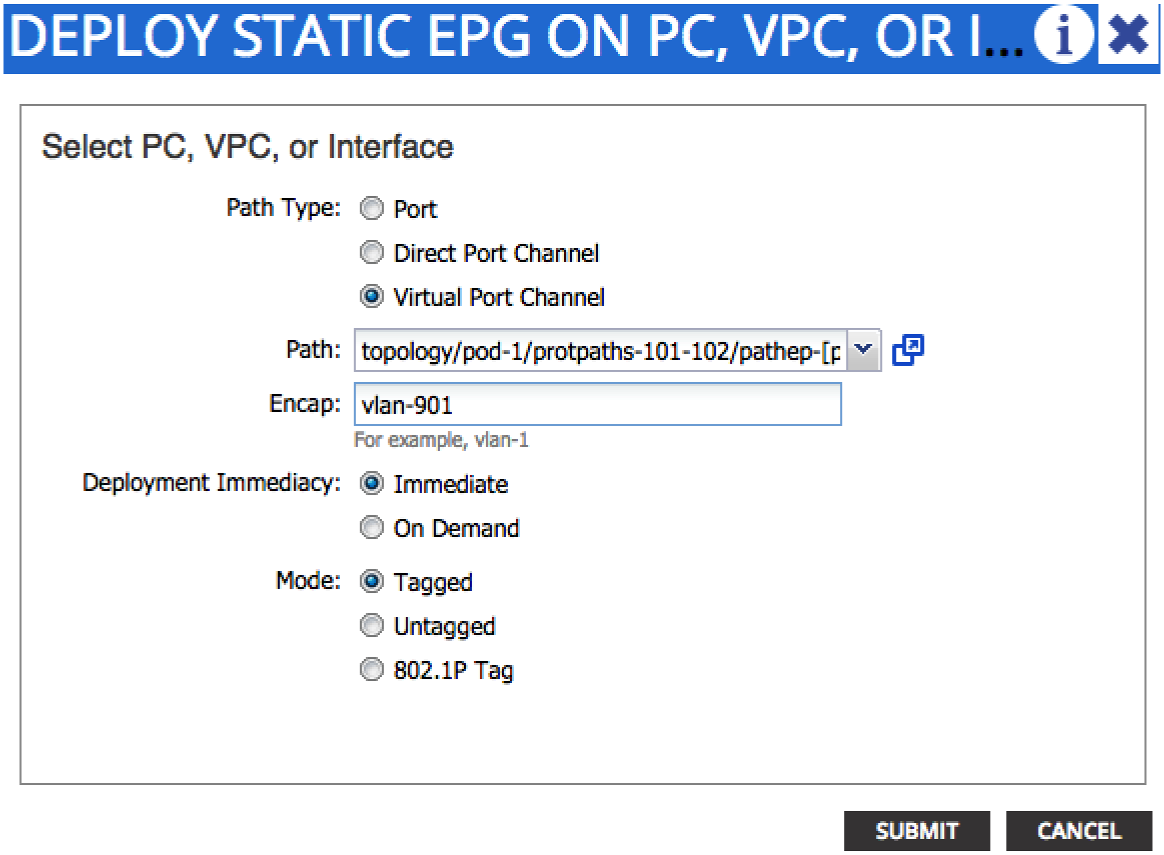

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[p

Encap: vlan-901

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

10. Click Submit.

11. Repeat steps 4-10 for mapping NetApp Controller 2 path.

12. Static bindings should display as shown below:

Figure 4 APIC - Application Profile iSCSI-A - Static Path Bindings

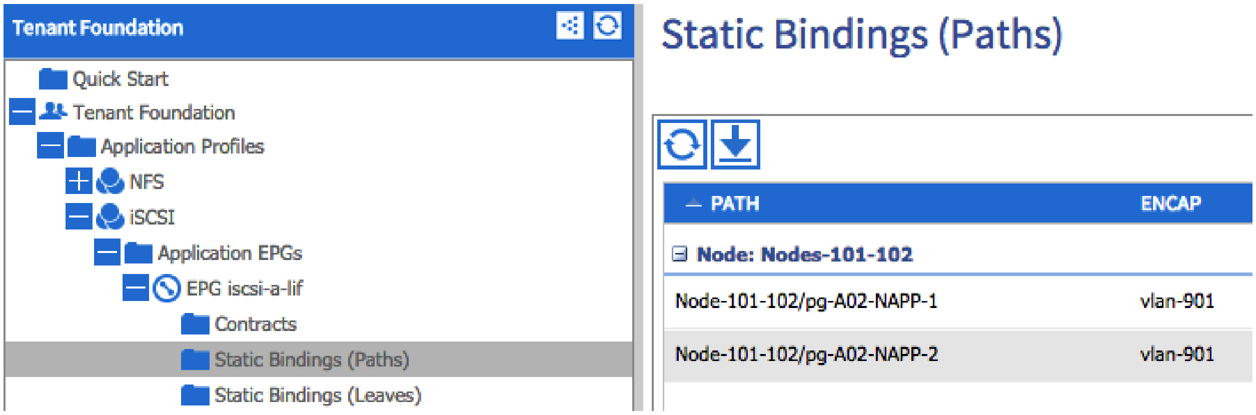

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A02-NAPP-1	vlan-901
Node-101-102/pg-A02-NAPP-2	vlan-901

Setting up EPG iscsi-b-lif

1. Expand the iSCSI Application profile from the menu bar on the left.

2. Expand iSCSI, expand Application EPGs and expand EPG iscsi-b-lif.

3. Click Static Bindings (Paths).

4. Repeat Step 4-10 to add two static bindings for both NetApp controllers using vlan-<<var_iscsi_vlan_B_id>> as Encapsulation.

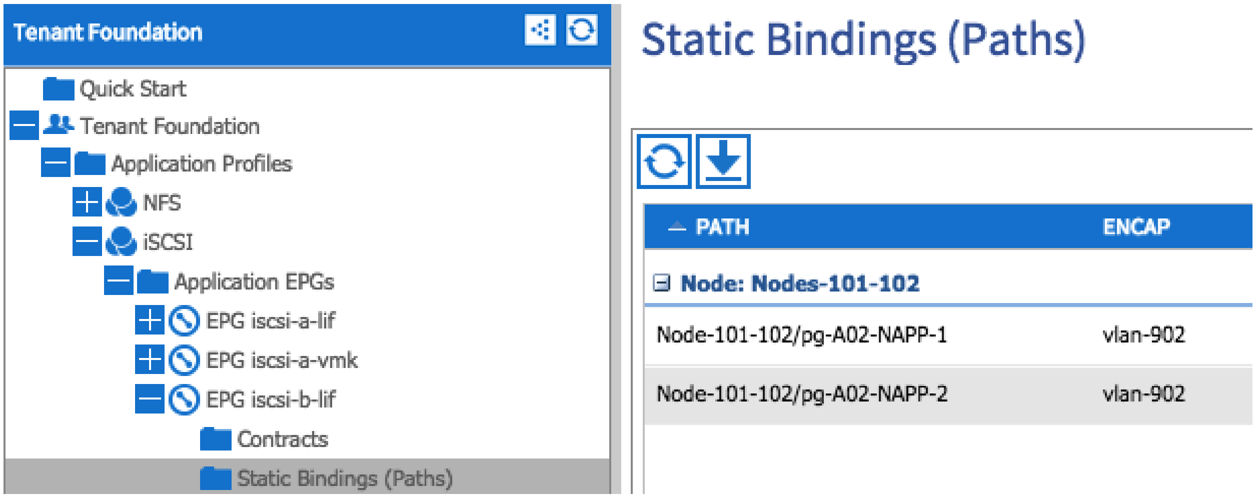

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A02-NAPP-1	vlan-902
Node-101-102/pg-A02-NAPP-2	vlan-902

Setting up EPG iscsi-a-vmk

1. Expand the iSCSI Application profile from the menu bar on the left.

2. Expand iSCSI, expand Application EPGs and expand EPG iscsi-a-vmk.

3. Click Static Bindings (Paths).

4. Repeat Step 4 -10 to add two static bindings for both UCS Fabric Interconnects using vlan-<<var_iscsi_vlan_A_vmk>> as Encapsulation.

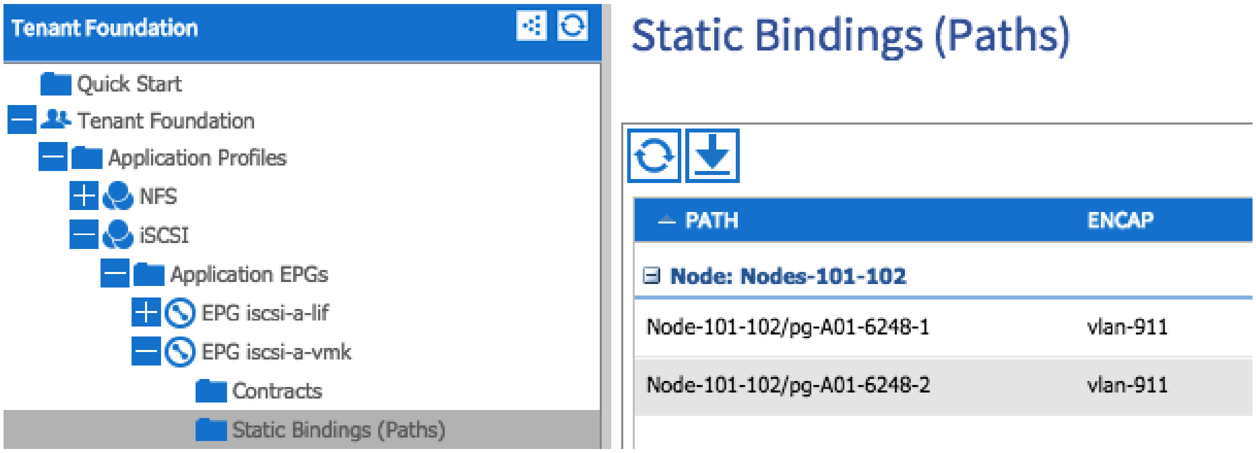

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A01-6248-1	vlan-911
Node-101-102/pg-A01-6248-2	vlan-911

Setting up EPG iscsi-b-vmk

1. Expand the iSCSI Application profile from the menu bar on the left.

2. Expand iSCSI, expand Application EPGs and expand EPG iscsi-b-vmk.

3. Click Static Bindings (Paths).

4. Repeat Step 4-10 to add two static bindings for both UCS Fabric Interconnects using vlan-<<var_iscsi_vlan_B_vmk>> as Encapsulation.

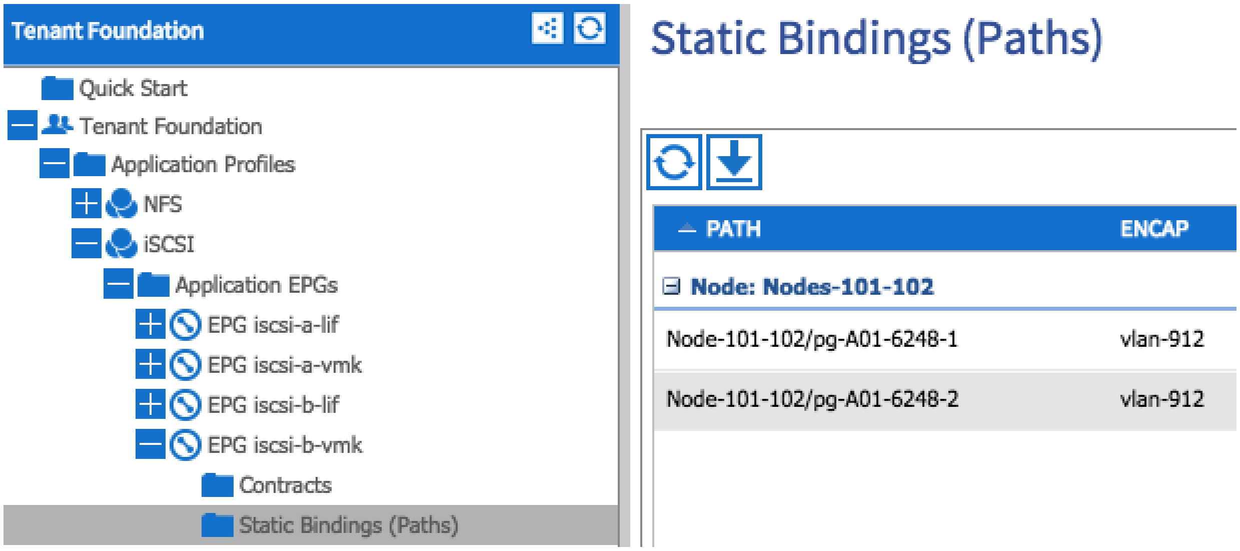

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A01-6248-1	vlan-912
Node-101-102/pg-A01-6248-2	vlan-912

Setting up Provided Contracts

To set up the provided contracts, complete the following steps:

1. Click EPG iscsi-a-lif in the left menu.

2. Click Contracts under the EPG.

3. Click ACTIONS on the right and select Add Provided Contract.

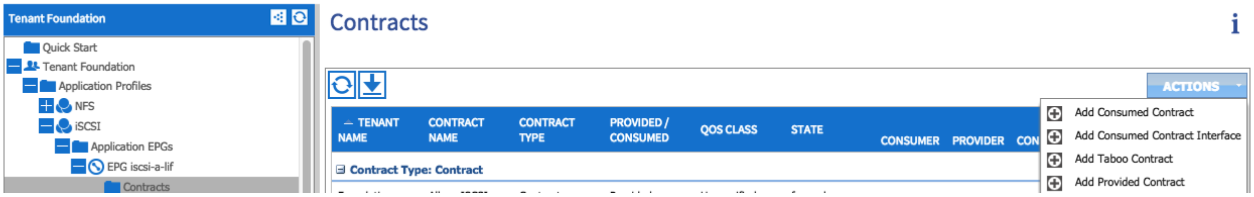

4. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

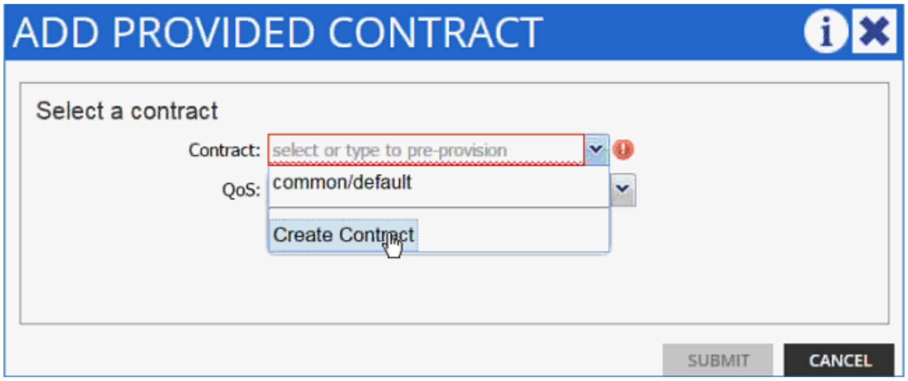

ADD PROVIDED CONTRACT

Select a contract

Contract: select or type to pre-provision

QoS: common/default

5. Enter Allow_ISCSI as Name in the CREATE CONTRACT dialog box.

6. Set Scope as Tenant.

7. Click + next to Subjects to add a new contract subject.

8. In the CREATE CONTRACT SUBJECT dialog box, enter Allow_ISCSI as the Name.

9. Click + under Filter Chain to add a new filter.

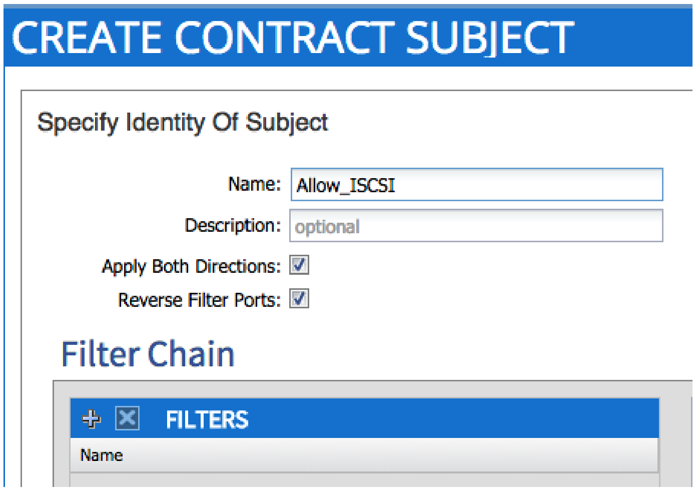

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow_ISCSI

Description: optional

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

10. Drop down the menu under FILTERS and click +.

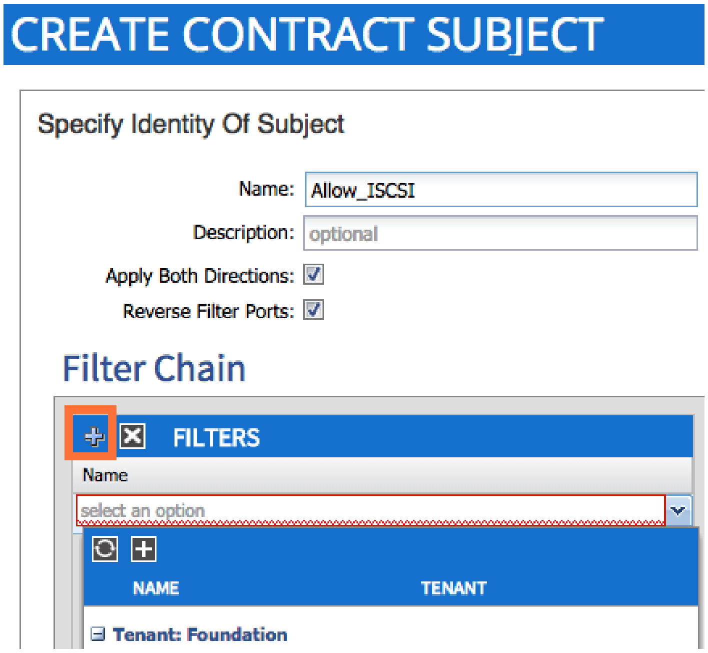

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow_ISCSI

Description: optional

Apply Both Directions: ☑

Reverse Filter Ports: ☑

Filter Chain

Name
select an option

NAME	TENANT
Tenant: Foundation	

11. In the CREATE FILTER dialog box, enter Allow_ISCSI as the Name. In this example, allow all the traffic for this contract.

12. Click + to add a filter.

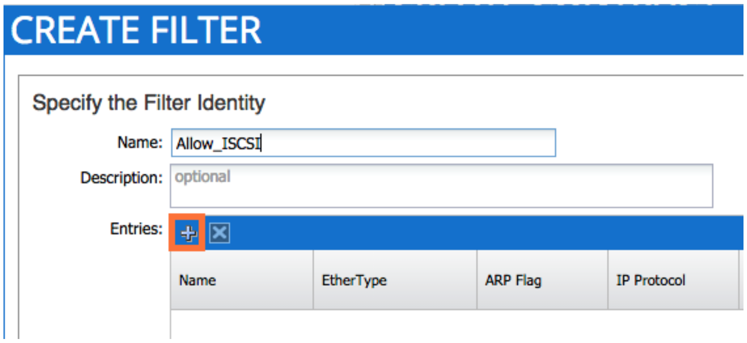

CREATE FILTER

Specify the Filter Identity

Name: Allow_ISCSI

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol

13. Enter tcp3260 as the name of the filter.

14. From drop-down menu, select IP as Ethertype.

15. From drop-down menu, select TCP as IP Protocol.

16. In Destination Port/Range type 3260 as both From and To ports.

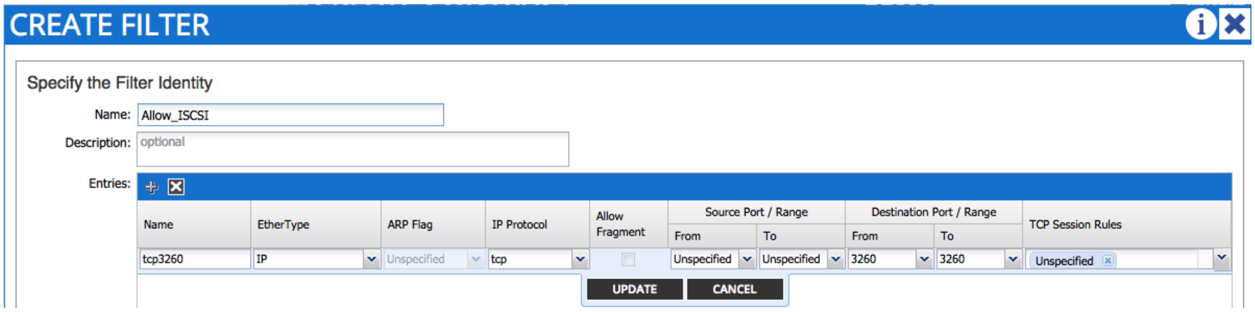

CREATE FILTER

Specify the Filter Identity

Name: Allow_ISCSI

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol	Allow Fragment	Source Port / Range		Destination Port / Range		TCP Session Rules
					From	To	From	To	
tcp3260	IP	Unspecified	tcp	□	Unspecified	Unspecified	3260	3260	Unspecified

17. Click Update.

18. Click SUBMIT to create the filter.

19. Click UPDATE to add the newly created filter to the filter chain.

20. Click OK to finish creating the Contract Subject.

21. Click SUBMIT.

22. Click SUBMIT again to finish adding a provided contract.

23. Verify the Provided Contract appears under the Contracts as shown below:

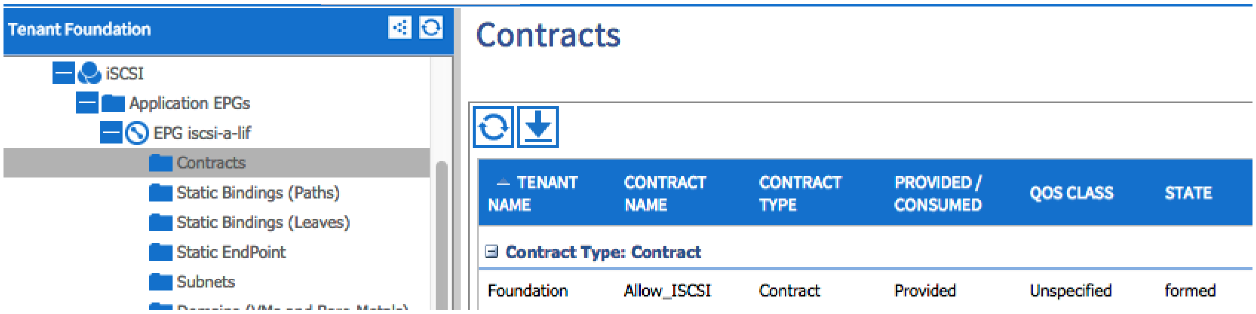

Contracts

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED	QOS CLASS	STATE
⊟ Contract Type: Contract					
Foundation	Allow_ISCSI	Contract	Provided	Unspecified	formed

24. Click EPG iscsi-b-lif in the left menu.

25. Click Contracts under the EPG.

26. Click ACTIONS on the right and select Add Provided Contract.

27. From the drop-down menu, select the recently created Allow_ISCSI contract.

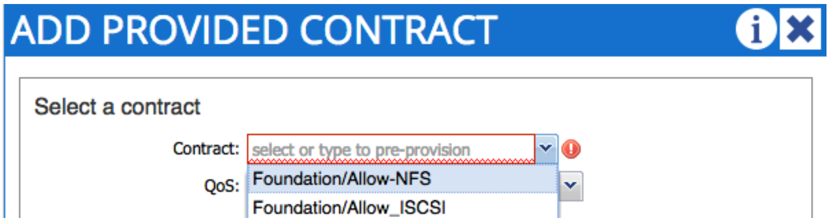

ADD PROVIDED CONTRACT

Select a contract

Contract: select or type to pre-provision

QoS:

Foundation/Allow-NFS

Foundation/Allow_ISCSI

28. Click SUBMIT to add the Allow_ISCSI contract as the Provided contract for EPG iscsi-b-lif as well.

Setting up Consumes Contracts

To set up the Consumed contract, complete the following steps:

1. Click EPG iscsi-a-vmk in the left menu.

2. Click Contracts under the EPG.

3. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select Foundation/Allow_ISCSI contract (defined in the previous step).

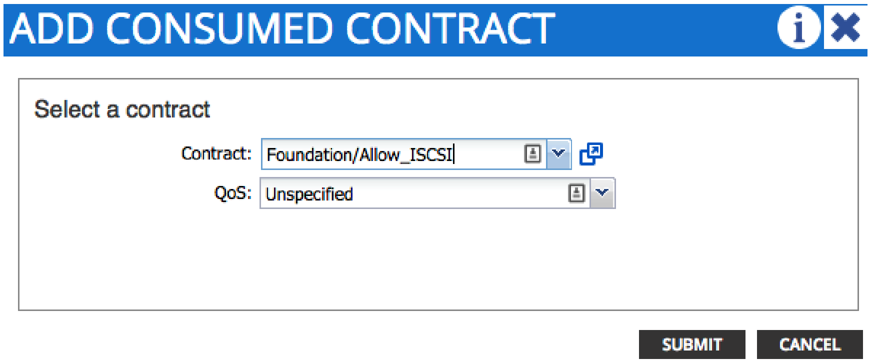

ADD CONSUMED CONTRACT

Select a contract

Contract: Foundation/Allow_ISCSI

QoS: Unspecified

4. Click SUBMIT.

5. Click EPG iscsi-b-vmk in the left menu.

6. Click Contracts under the EPG.

7. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select Foundation/Allow_ISCSI contract.

8. To validate the contract definition, click Application EPGs under Application Profile iSCSI in the left menu bar. The contract should appear as shown below:

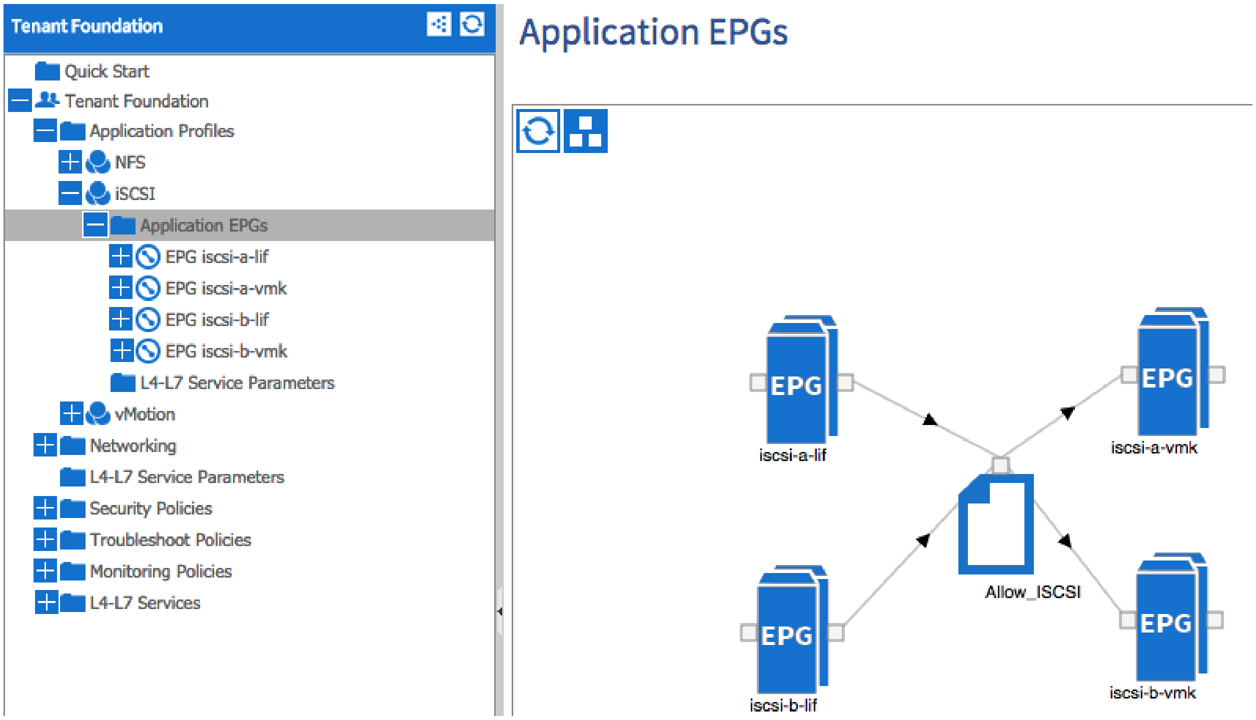

NFS Application Profile Creation

To create the NFS application profile, complete the following steps:

1. Expand Tenant Foundation in the left menu bar.
2. Right-click Application Profile and click Create Application Profile.
3. In the CREATE APPLICATION PROFILE dialog box, enter NFS as the Name.
4. From the drop-down menu, select default for Monitoring Policy.
5. Click + next to EPG to add an EPG.

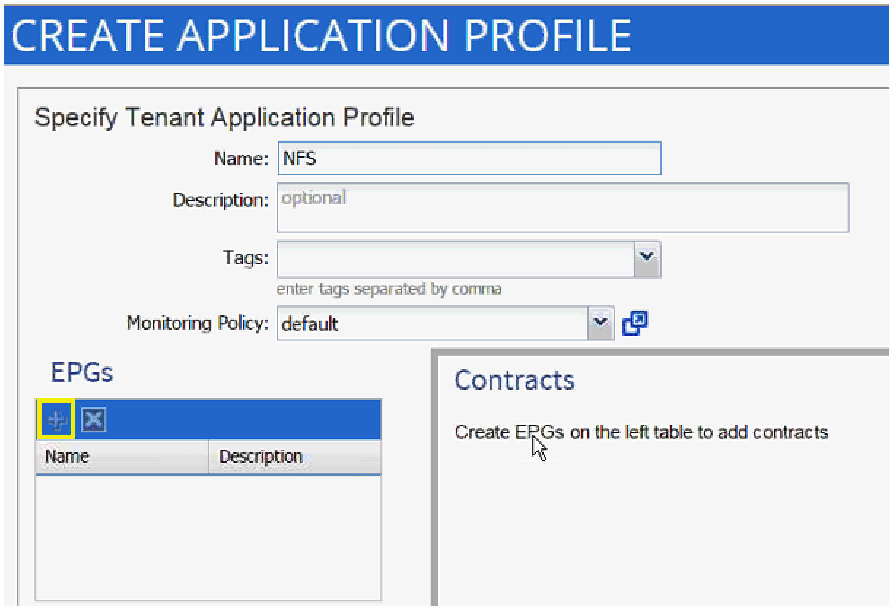

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: NFS

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

6. In the CREATE APPLICATION EPG dialog box, enter lif-nfs as the Name.
7. From the drop-down menu, select bd-internal as the Bridge Domain.
8. From the drop-down menu, select default for Monitoring Policy.

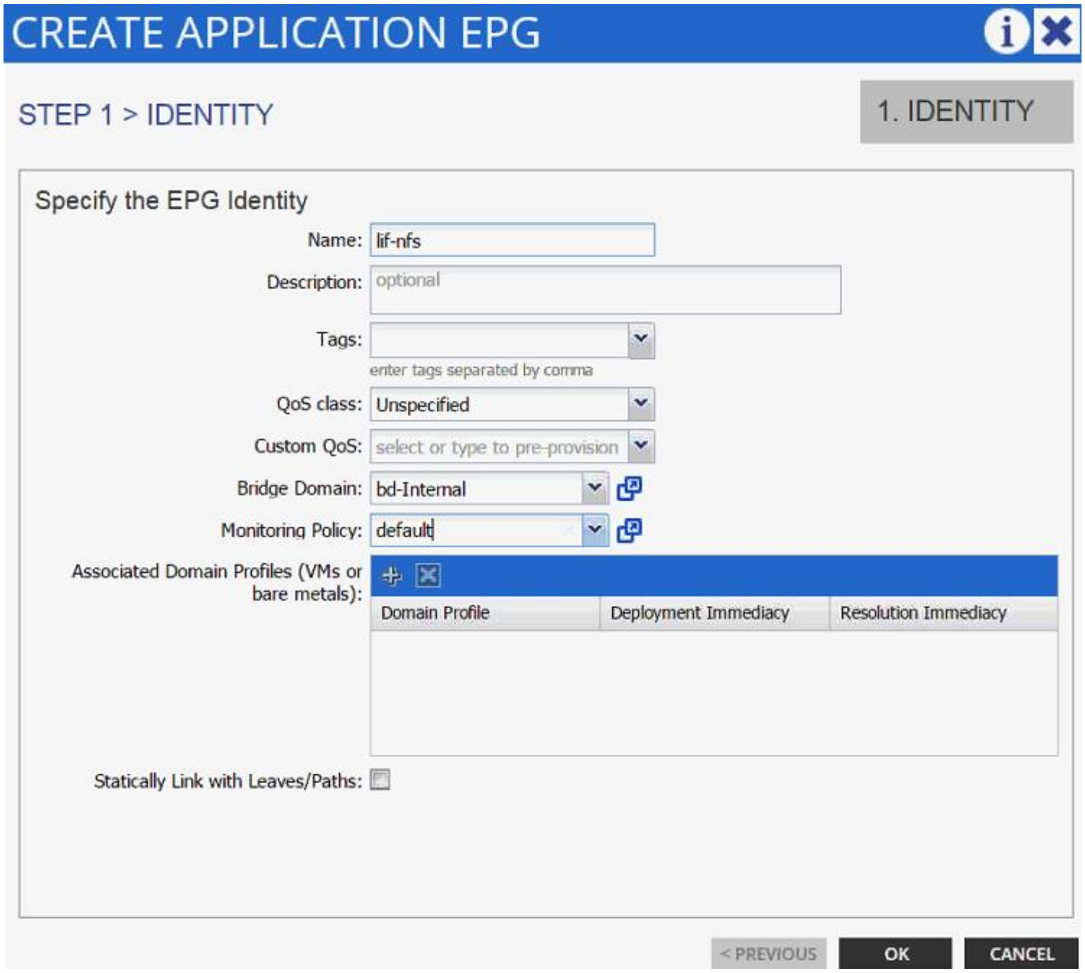

CREATE APPLICATION EPG

STEP 1 > IDENTITY

Specify the EPG Identity

Name: lif-nfs

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy

Statically Link with Leaves/Paths: □

9. Click OK.

10. Click + next to EPG to another EPG.

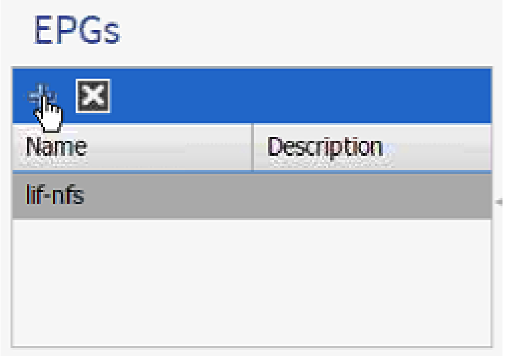

EPGs

Name	Description
lif-nfs	

11. In the CREATE APPLICATION EPG dialog box, enter vmk-nfs as the Name.

12. From the drop-down menu, select bd-internal as the Bridge Domain.

13. From the drop-down menu, select default Monitoring Policy.

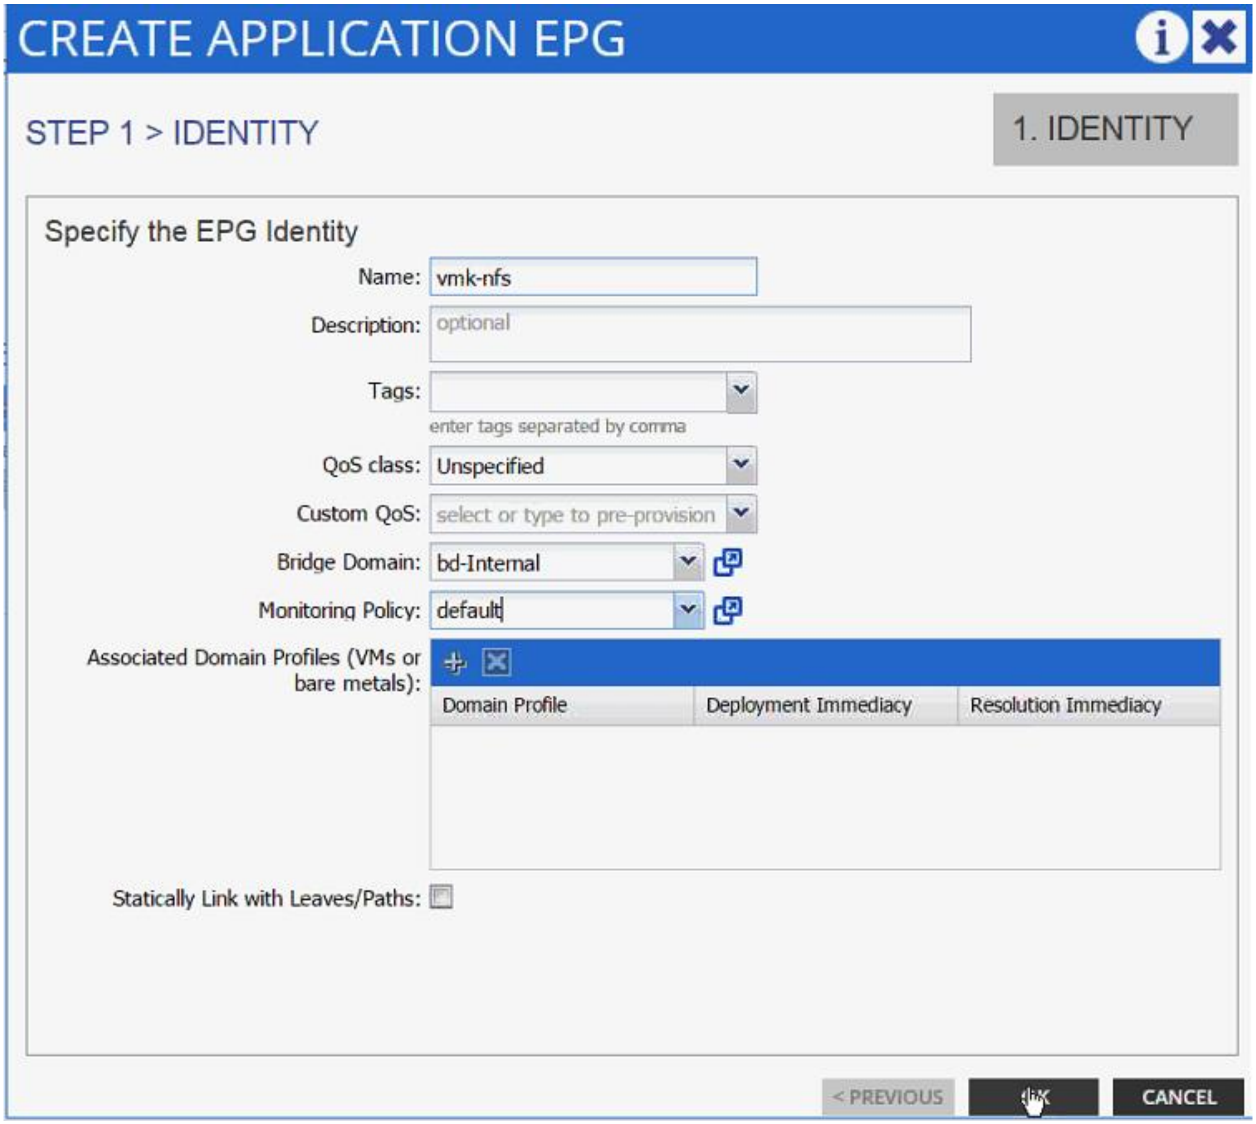

CREATE APPLICATION EPG

STEP 1 > IDENTITY

1. IDENTITY

Specify the EPG Identity

Name: vmk-nfs

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy

Statically Link with Leaves/Paths: □

14. Click OK.

15. Click SUBMIT to finish creating the Application Profile.

16. Expand the newly created NFS Application profile from the menu bar on the left.

17. Expand NFS, expand Application EPGs and expand EPG lif-nfs.

18. Click Static Bindings (Paths).

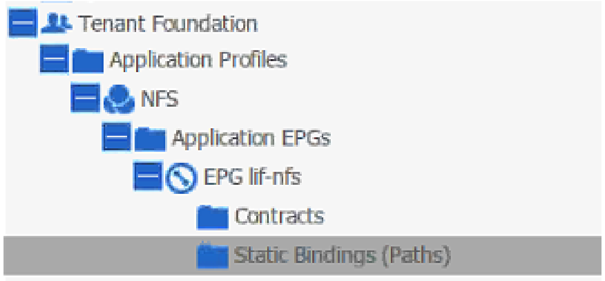

- Tenant Foundation
  - Application Profiles
    - NFS
      - Application EPGs
        - EPG lif-nfs
          - Contracts
          - Static Bindings (Paths)

19. Click Action on the right hand work area.

20. Click Deploy Static EPG on PC, vPC, or Interface.

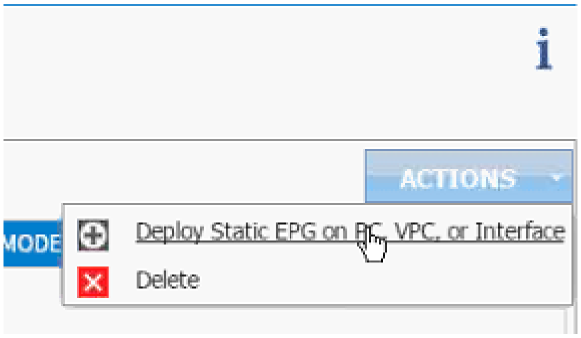

21. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

22. From the drop-down menu Path, select NetApp Controller 1.

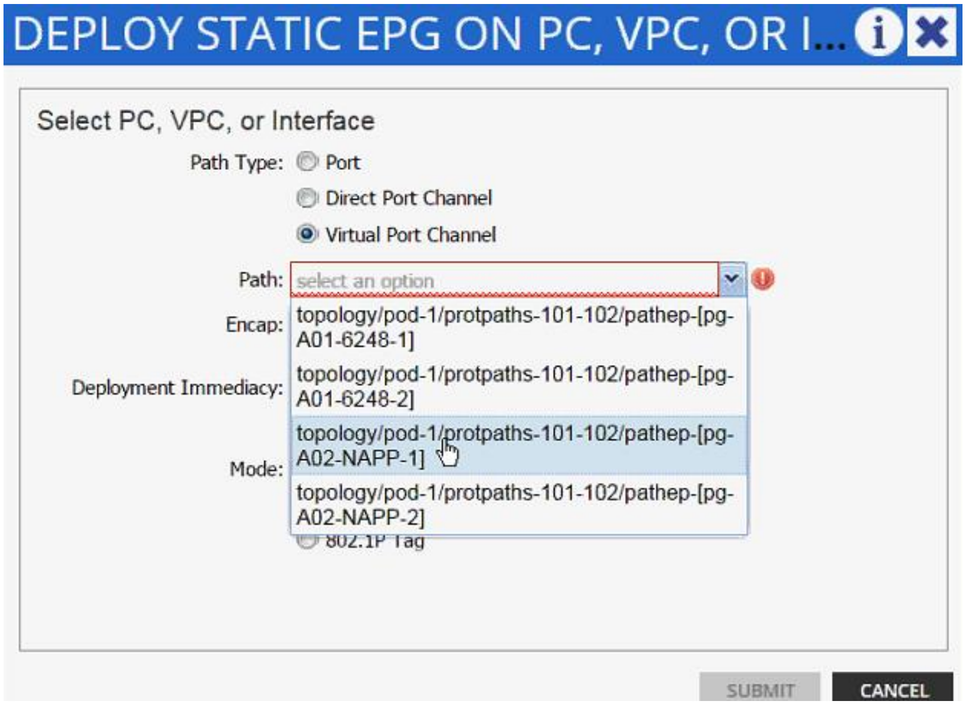

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port
○ Direct Port Channel
● Virtual Port Channel

Path: select an option

topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]

Encap:

Deployment Immediacy:

Mode: ○ 802.1P Tag

SUBMIT CANCEL

23. Enter vlan-<NFS LIF VLAN> for Encap (VLAN 3170 is the NFS VLAN on NetApp Controller in the screen capture below).

24. Change Deployment Immediacy to Immediate.

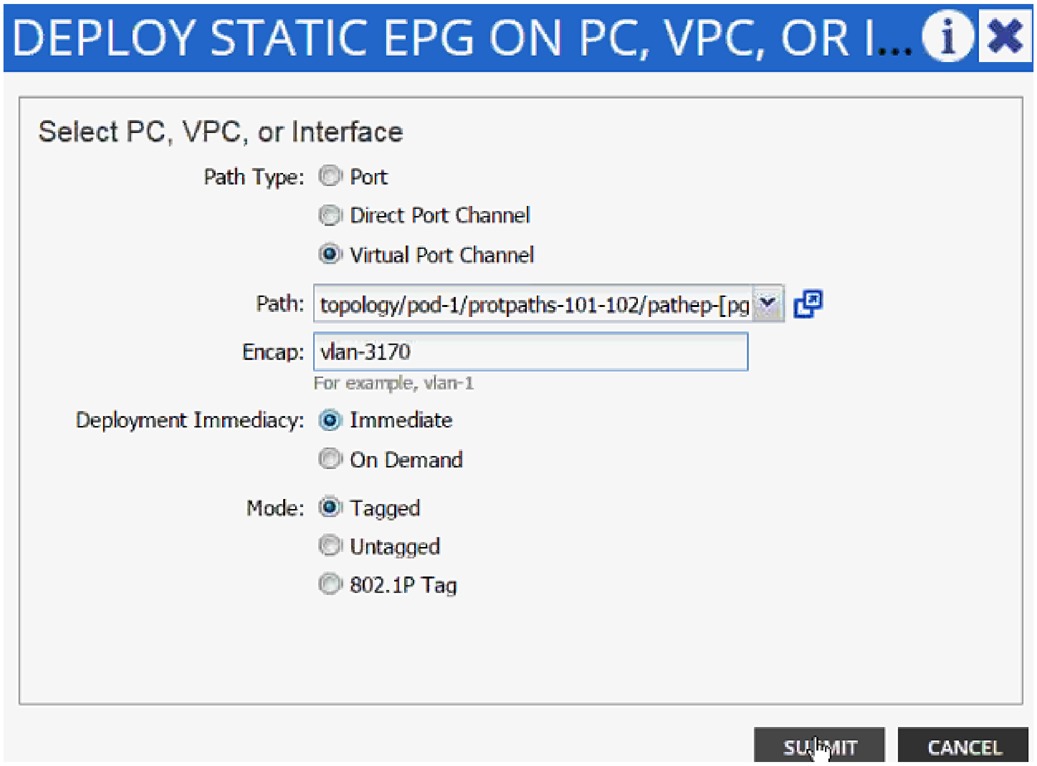

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[pg

Encap: vlan-3170

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

25. Click Submit.

26. Validate the path appears in the work area on the right.

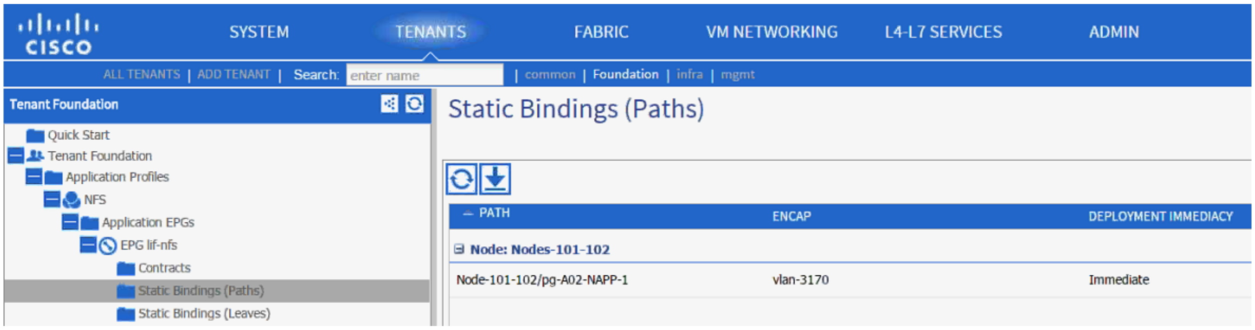

Static Bindings (Paths)

PATH	ENCAP	DEPLOYMENT IMMEDIACY
Node: Nodes-101-102		
Node-101-102/pg-A02-NAPP-1	vlan-3170	Immediate

27. Repeat these steps for mapping NetApp Controller 2 path.

28. Static bindings should be similar to screenshot below:

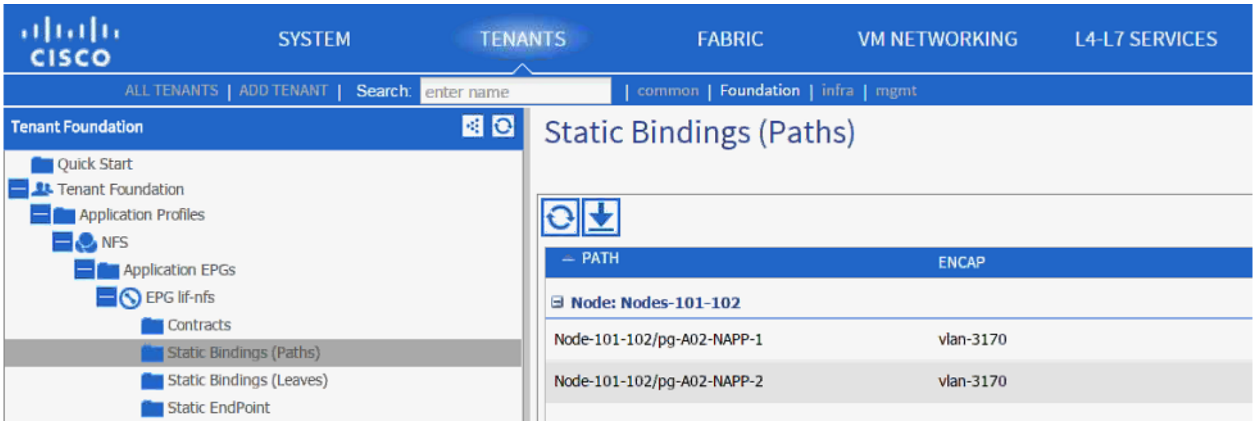

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A02-NAPP-1	vlan-3170
Node-101-102/pg-A02-NAPP-2	vlan-3170

29. Click Contracts under the EPG lif-nfs.

30. Click Action and select Add Provided Contract.

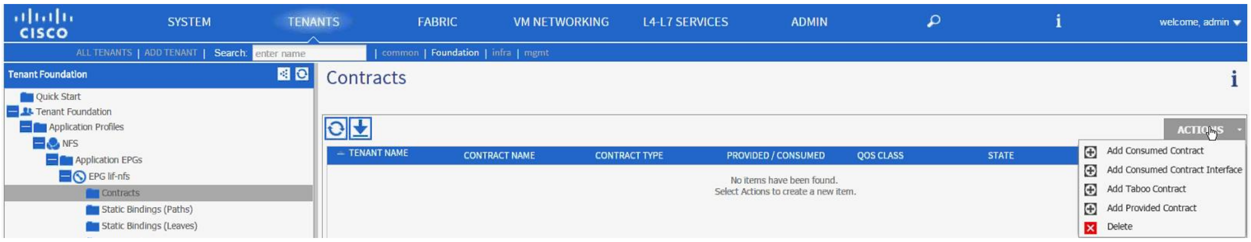

Contracts

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED	QOS CLASS	STATE

No items have been found.
Select Actions to create a new item.

31. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

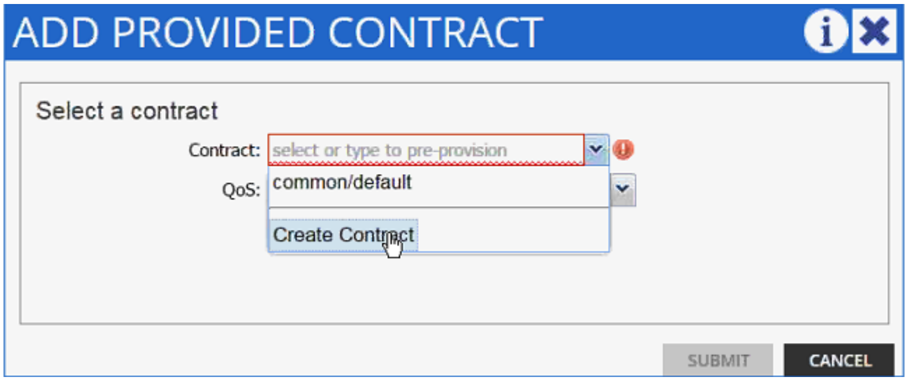

Select a contract

Contract: select or type to pre-provision

QoS: common/default

32. Enter Allow-NFS as Name in the CREATE CONTRACT dialog box.

33. Click + next to Subjects to add a new contract subject.

34. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

35. Click + under Filter Chain to add a new filter.

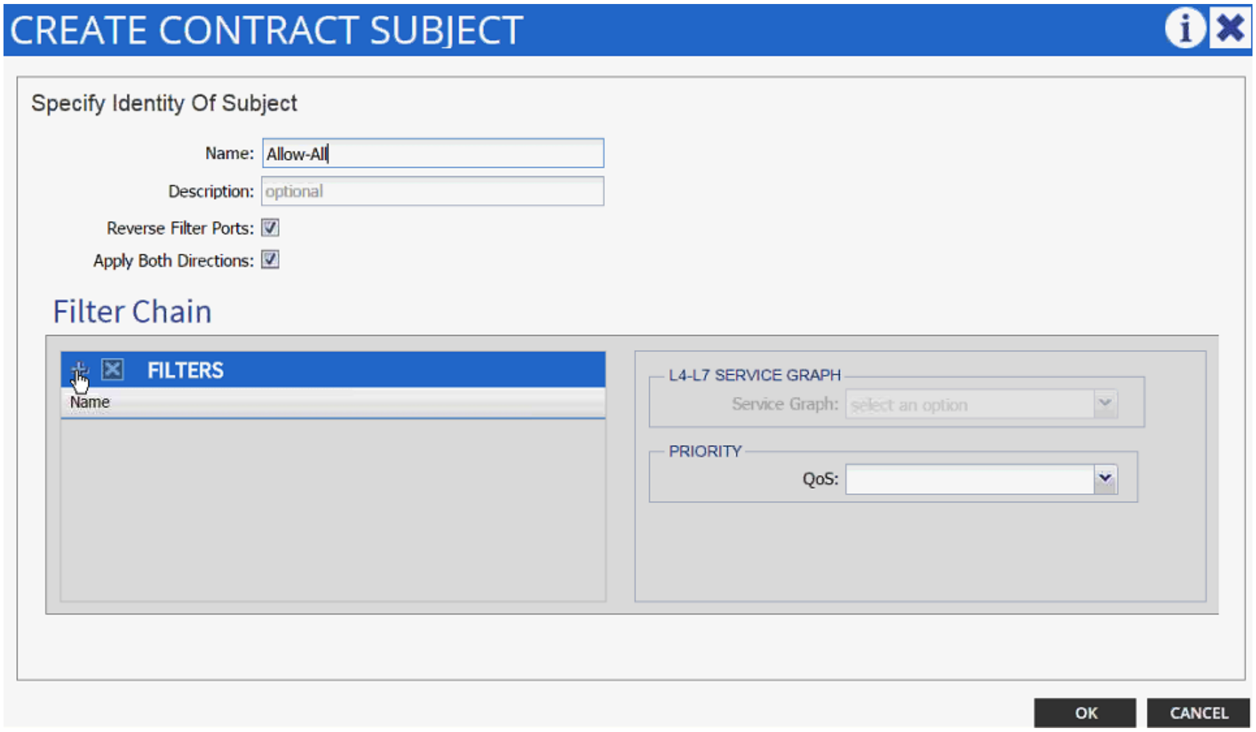

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS

Name

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

36. Drop down the menu under FILTERS and click +.

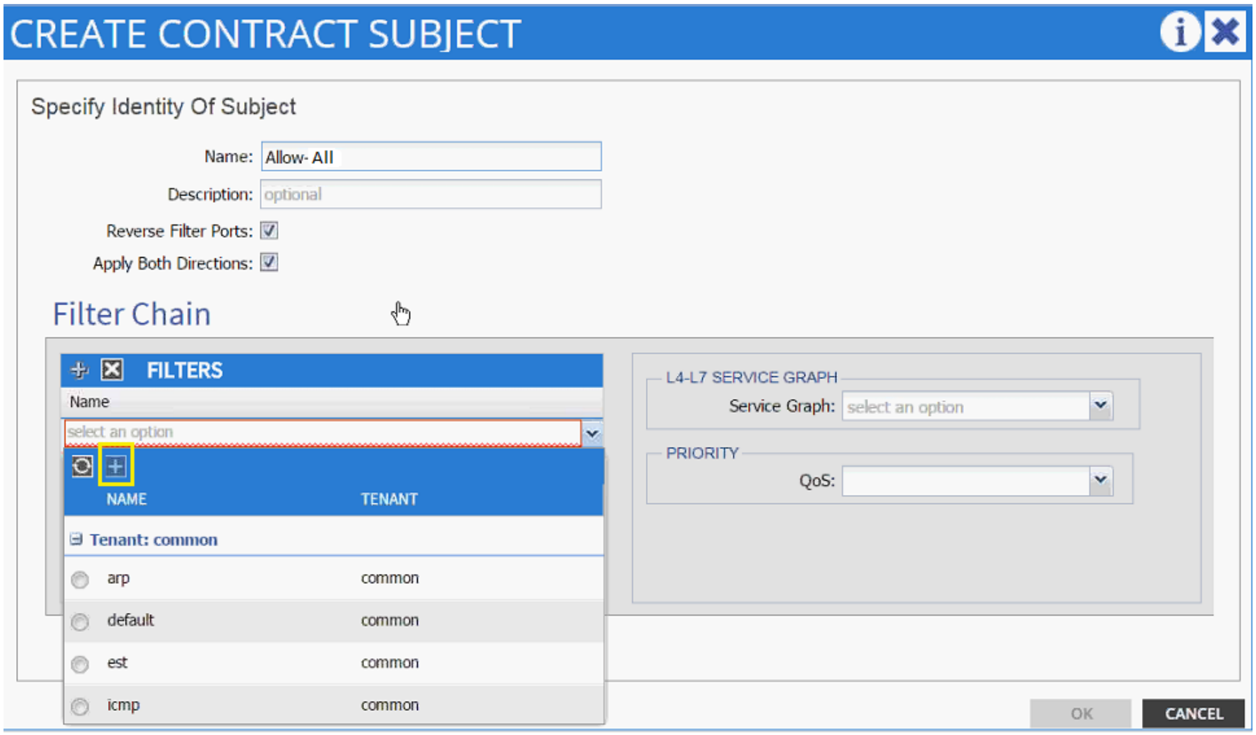

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow- All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS

Name

select an option

NAME	TENANT
⊟ Tenant: common	
○ arp	common
○ default	common
○ est	common
○ icmp	common

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

37. In the CREATE FILTER dialog box, enter Allow-All as the Name. In this example, allow all the traffic for this contract.

38. Click + to add a filter.

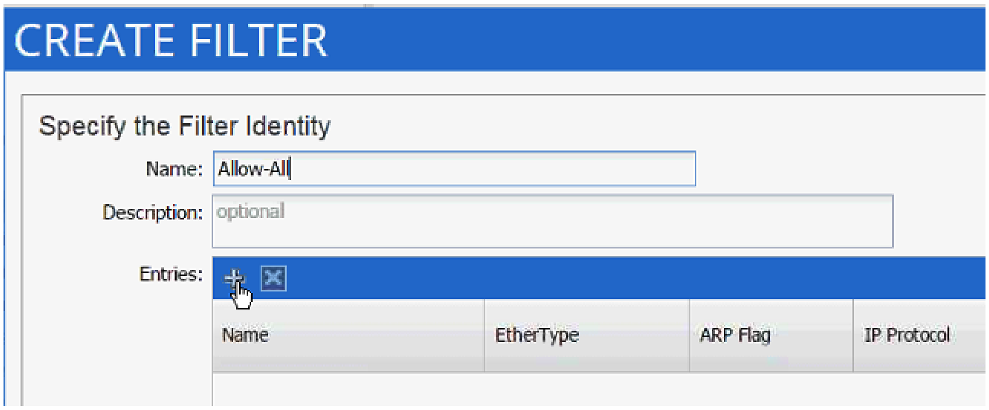

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol

39. Enter Allow-All as the name of the filter.

40. From drop-down menu, select IP as Ethertype.

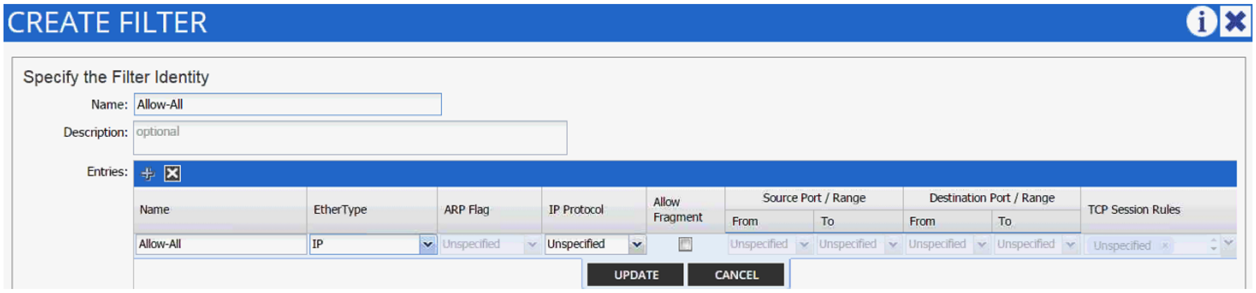

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol	Allow Fragment	Source Port / Range		Destination Port / Range		TCP Session Rules
					From	To	From	To	
Allow-All	IP	Unspecified	Unspecified	□	Unspecified	Unspecified	Unspecified	Unspecified	Unspecified

41. Click Update.

42. Click SUBMIT to create the filter.

43. Click UPDATE to add the newly created filter to the filter chain.

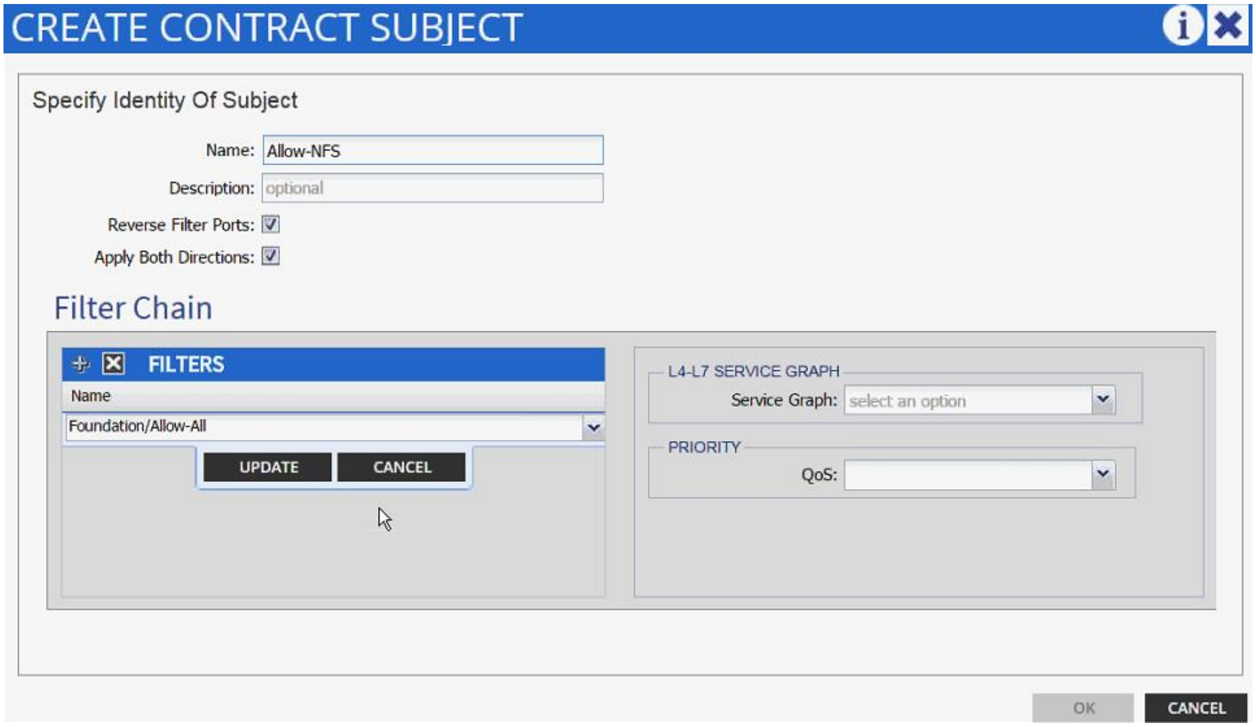

Name: Allow-NFS

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Name
Foundation/Allow-All

Service Graph: select an option

QoS:

44. Click OK to finish creating the Contract Subject.

45. Change the Scope to tenant from the drop-down list.

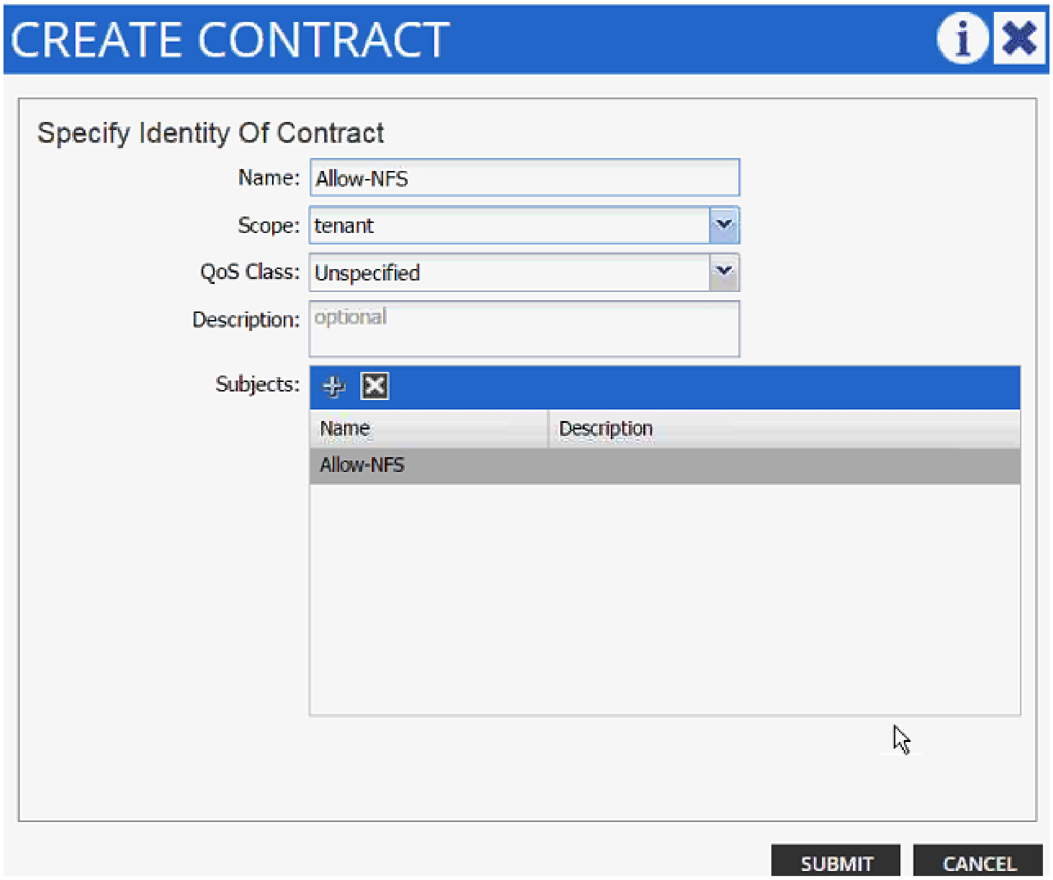

Name: Allow-NFS

Scope: tenant

QoS Class: Unspecified

Description: optional

Subjects:

Name	Description
Allow-NFS	

46. Click SUBMIT.

47. Click SUBMIT again to finish adding a provided contract.

48. Verify the Provided Contract appears under the Contracts.

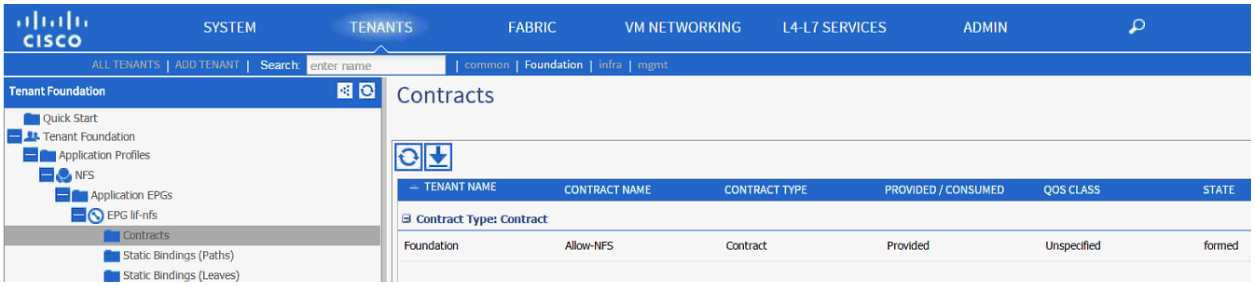

Contracts

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED	QOS CLASS	STATE
Contract Type: Contract					
Foundation	Allow-NFS	Contract	Provided	Unspecified	formed

49. Expand the NFS Application profile (again) from the menu bar on the left.

50. Expand NFS, expand Application EPGs and expand EPG vmk-nfs.

51. Click Static Bindings (Paths).

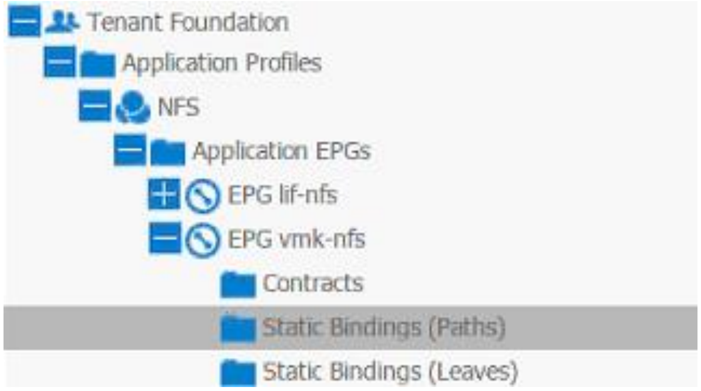

52. Click Action in the right hand work area.

53. Click Deploy Static EPG on PC, vPC, or Interface.

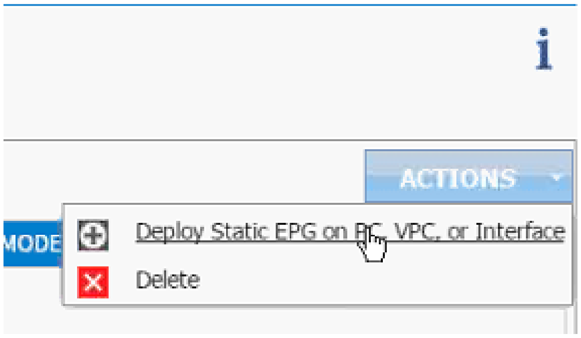

54. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

55. From the drop-down menu Path, select UCS Fabric Interconnect A.

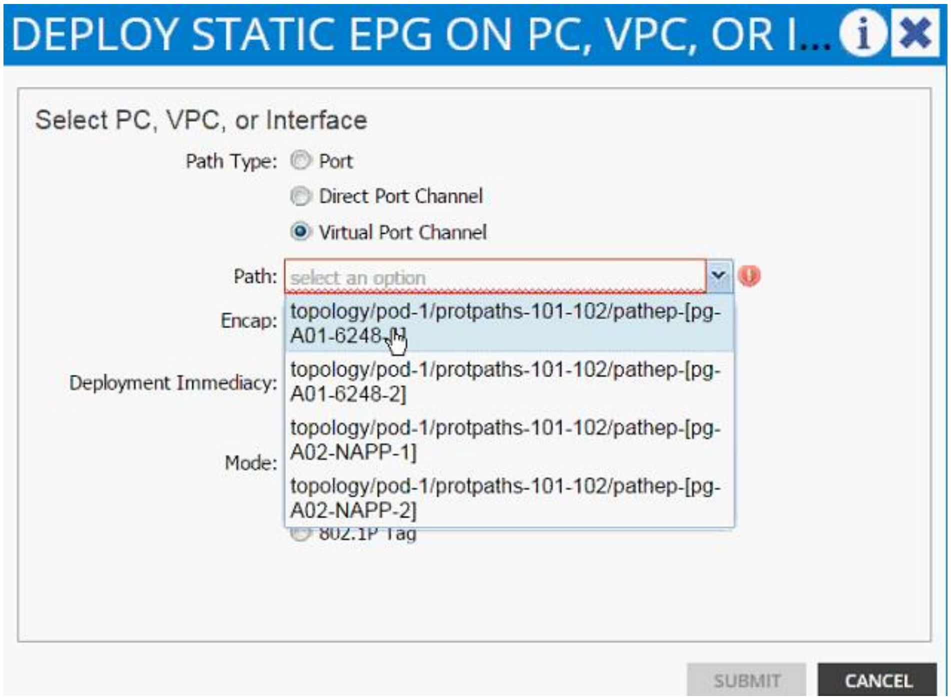

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: select an option

- topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]
- topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]
- topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]
- topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]

56. Enter vlan-<NFS VMK VLAN> for Encap; VLAN 3270 is the NFS VLAN on UCS Fabric Interconnect in the screenshot below.

Note: A VLAN on a certain path can only be mapped to a single EPG. Since VLAN 3170 (NFS VLAN on NetApp) is already mapped to EPG lif-NFS, VLAN 3270 was selected as the VLAN to host ESXi VMKernel ports. The VMKernel ports and the NetApp LIFs will still be defined in the same IP subnet; ACI Fabric will enable seamless IP connectivity when contracts are defined between the two EPGs.

57. Change Deployment Immediacy to Immediate.

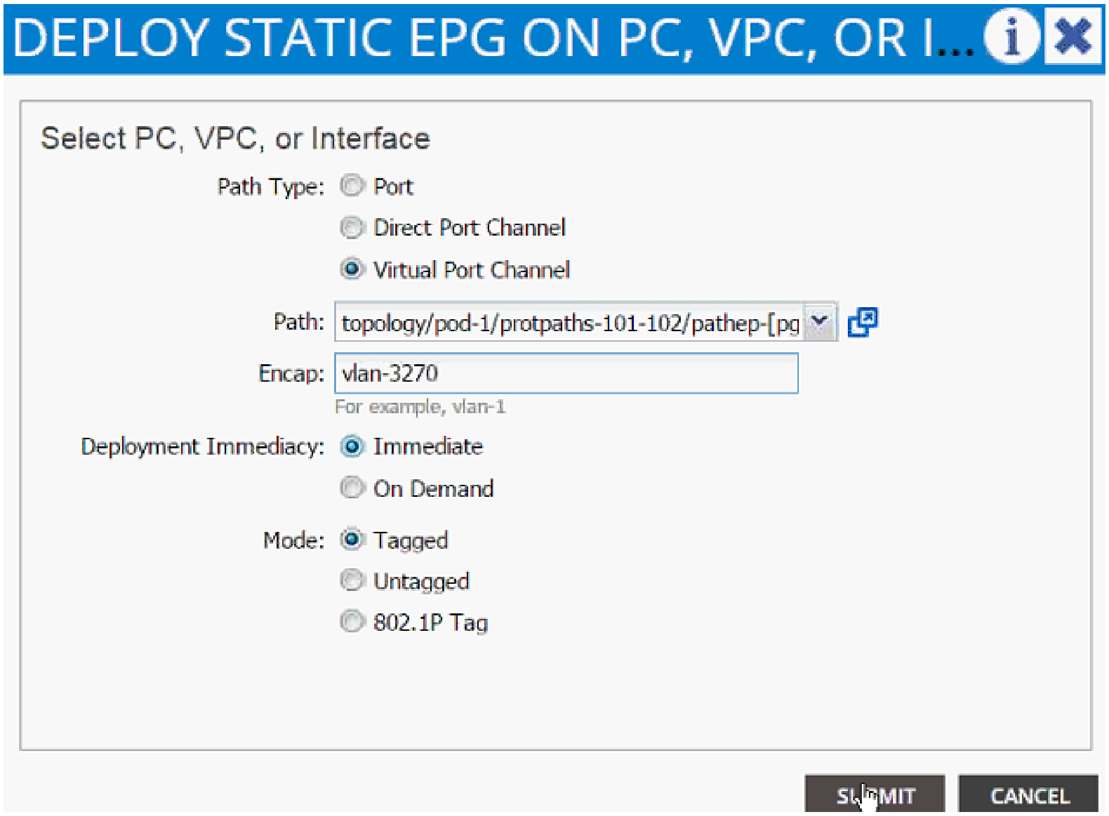

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[pg

Encap: vlan-3270

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

58. Click Submit.

59. Validate the path appears in the work area.

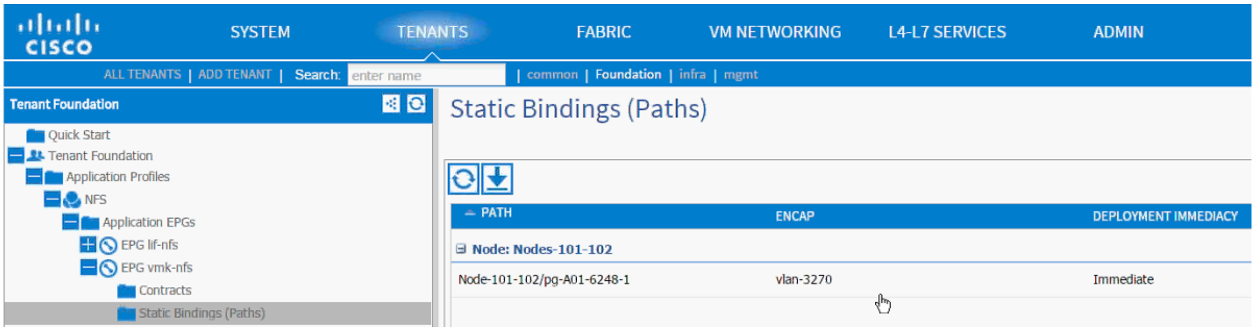

Static Bindings (Paths)

PATH	ENCAP	DEPLOYMENT IMMEDIACY
Node: Nodes-101-102		
Node-101-102/pg-A01-6248-1	vlan-3270	Immediate

60. Repeat these steps for mapping UCS Fabric Interconnect B path.

61. Static bindings should be similar to screenshot below.

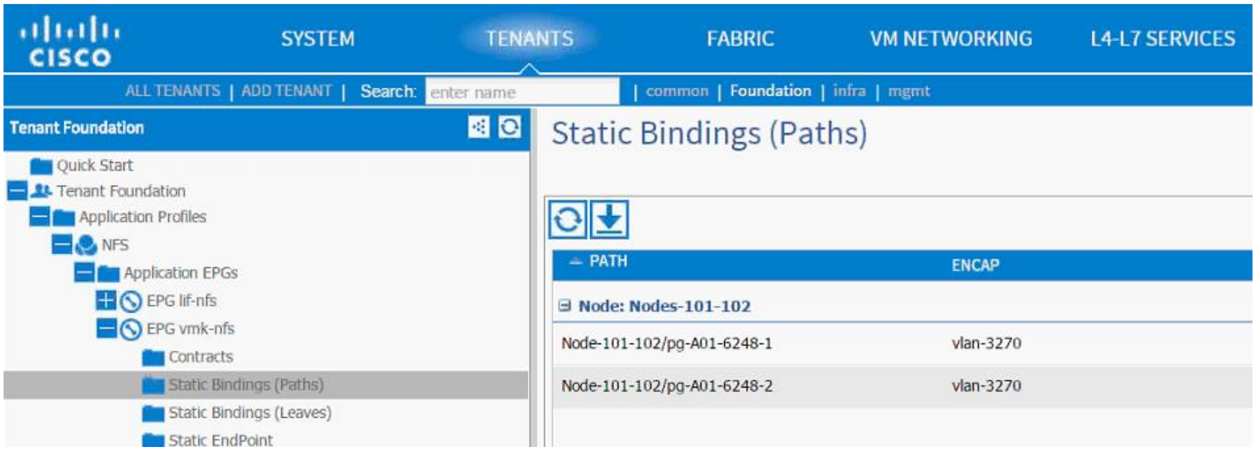

Static Bindings (Paths)

PATH	ENCAP
Node: Nodes-101-102	
Node-101-102/pg-A01-6248-1	vlan-3270
Node-101-102/pg-A01-6248-2	vlan-3270

62. Click Contracts in the left menu.

63. Click ACTIONS on the right and select Add Consumed Contract.

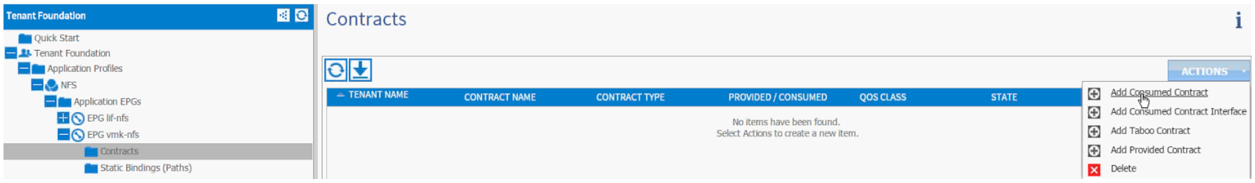

Contracts

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED	QOS CLASS	STATE

No items have been found.
Select Actions to create a new item.

64. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select Foundation/Allow-NFS contract (defined previously).

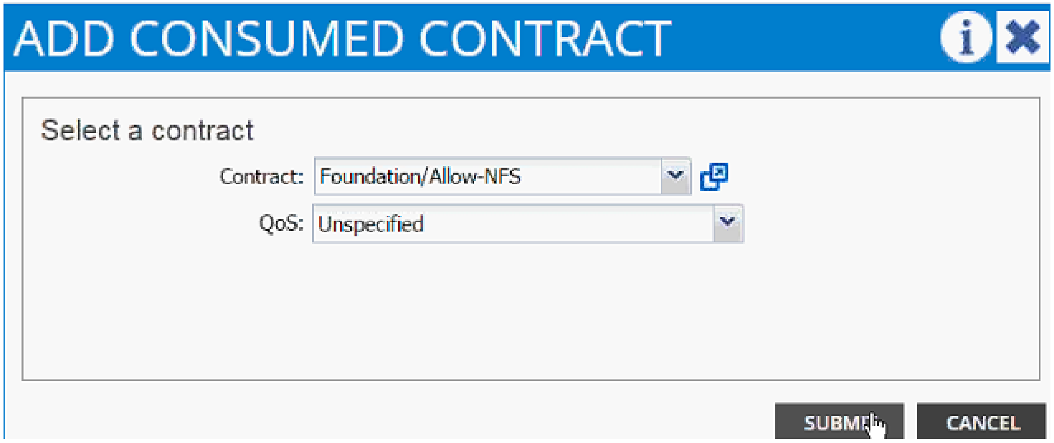

ADD CONSUMED CONTRACT

Select a contract

Contract: Foundation/Allow-NFS

QoS: Unspecified

65. Click SUBMIT.

66. To validate the contract definition, click Application EPGs under Application Profile NFS in the left menu bar. The contract should appear as shown in screenshot below:

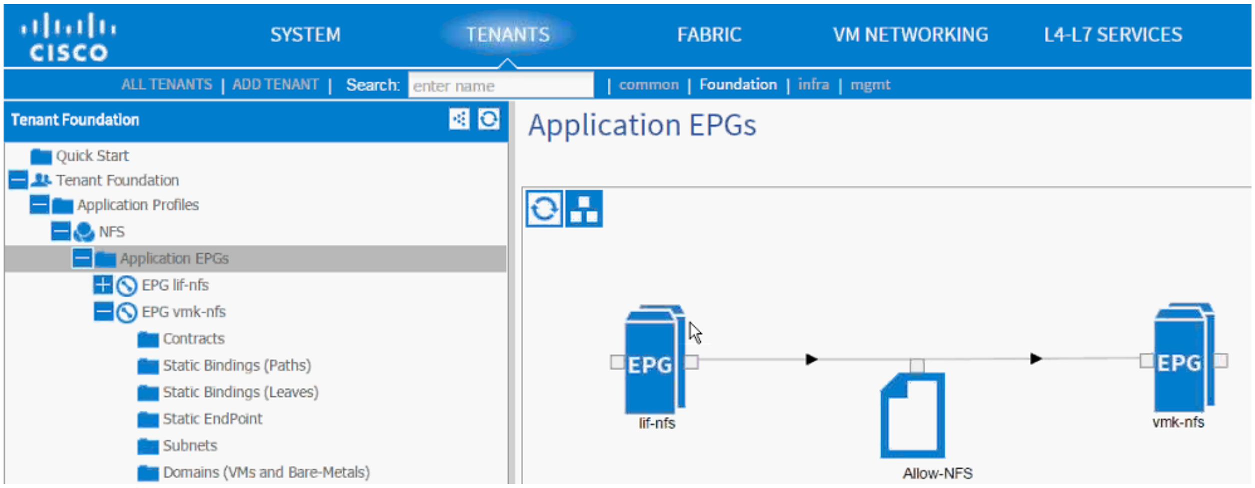

Path (vPC) Validation

Previously in this section, both iSCSI and NFS paths and VLANs were mapped to appropriate EPGs. These VLANs were also defined in the physical domains associated with the VPCs. At this point the Foundation tenant is deployed and should provide connectivity between the ESXi hosts and NetApp controllers. To validate connectivity, VPCs can be checked for appropriate VLAN forwarding.

To validate the path, complete the following steps:

67. To validate the VLAN forwarding on the vPC, select FABRIC from the top menu and select INVENTORY from the sub-menu.

68. Expand Pod 1, Leaf switch, Interfaces and then vPC Interfaces.

69. Expand the vPC domain (10) and click on a vPC.

70. As shown in the screenshot below, the vPC should show both the iSCSI and NFS VLANs being forwarded.

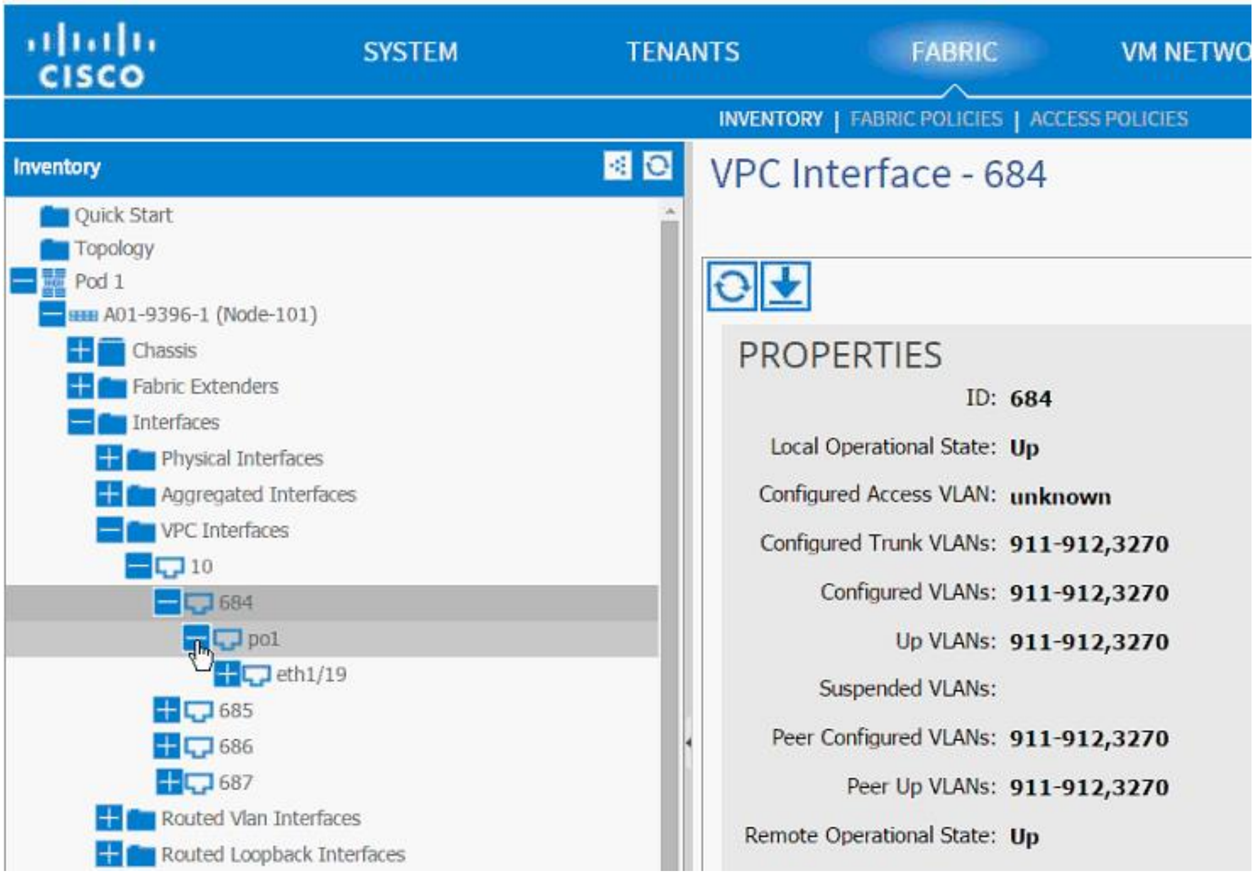

VPC Interface - 684

PROPERTIES

ID: 684

Local Operational State: Up

Configured Access VLAN: unknown

Configured Trunk VLANs: 911-912,3270

Configured VLANs: 911-912,3270

Up VLANs: 911-912,3270

Suspended VLANs:

Peer Configured VLANs: 911-912,3270

Peer Up VLANs: 911-912,3270

Remote Operational State: Up

71. Repeat these steps to validate all the VPCs.

72. Optional: Log into the Leaf using CLI and issue a show vpc command.

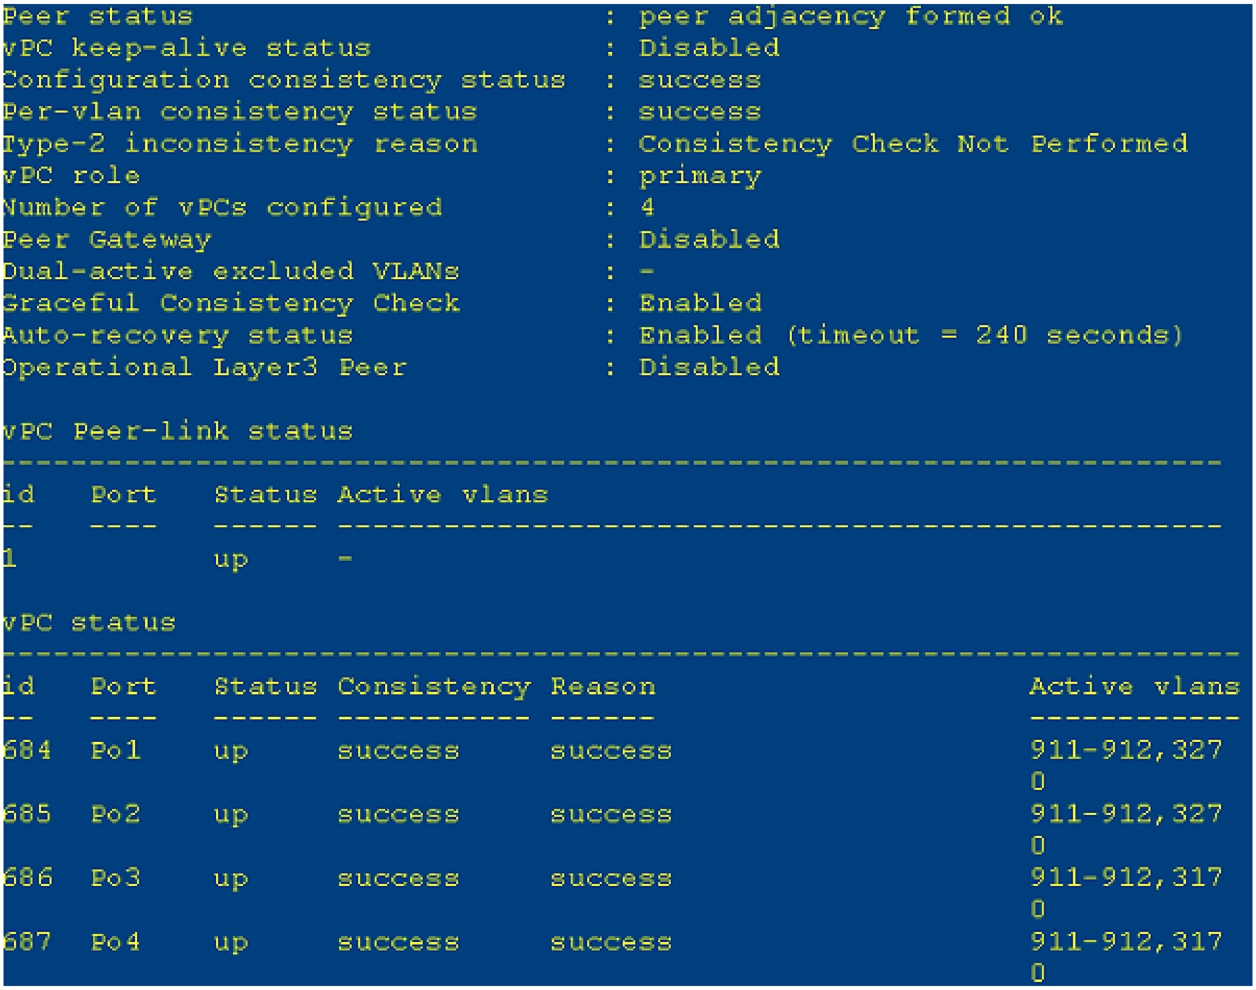

```
Peer status                              : peer adjacency formed ok
vPC keep-alive status                    : Disabled
Configuration consistency status         : success
Per-vlan consistency status              : success
Type-2 inconsistency reason              : Consistency Check Not Performed
vPC role                                 : primary
Number of vPCs configured                : 4
Peer Gateway                             : Disabled
Dual-active excluded VLANs               : -
Graceful Consistency Check               : Enabled
Auto-recovery status                     : Enabled (timeout = 240 seconds)
Operational Layer3 Peer                  : Disabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1           up     -

vPC status
----------------------------------------------------------------------
id   Port   Status Consistency Reason                     Active vlans
--   ----   ------ ----------- ------                     ------------
684  Po1    up     success     success                    911-912,327
                                                          0
685  Po2    up     success     success                    911-912,327
                                                          0
686  Po3    up     success     success                    911-912,317
                                                          0
687  Po4    up     success     success                    911-912,317
                                                          0
```

Storage Configuration – SAN Boot

Clustered Data ONTAP SAN Boot Storage Setup

Create igroups

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

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol iscsi –ostype
vmware –initiator <<var_vm_host_infra_01_iqn>>

igroup create –vserver Infra-SVM –igroup VM-Host-Infra-02 –protocol iscsi –ostype
vmware –initiator <<var_vm_host_infra_02_iqn>>

igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol iscsi –ostype
vmware –initiator <<var_vm_host_infra_01_iqn>>, <<var_vm_host_infra_02_iqn>>
```

Note: Use the values listed in **Error! Reference source not found.** for the IQN information.

Note: To view the three igroups just created, type `igroup` show.

Map Boot LUNs to igroups

1. From the storage cluster management SSH connection, enter the following:

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-01 –igroup VM-
Host-Infra-01 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-02 –igroup VM-
Host-Infra-02 –lun-id 0
```

VMware vSphere 5.5 Setup

VMware ESXi 5.5 Update 2

This section provides detailed instructions for installing VMware ESXi 5.5 Update 2 in an environment. After the procedures are completed, two booted 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 logical unit numbers (LUNs).

Download Cisco Custom Image for ESXi 5.5.0 U2

1. Click the following link vmware login page.
2. Type your email or customer number and the password and then click Log in.
3. Click the following link CiscoCustomImage5.5.0U2.
4. Click Download Now.
5. Save it to your destination folder.

Note: This ESXi 5.5.0 U2 Cisco custom image includes updates for the fnic and enic drivers. The versions that are part of this image are: Enic: 2.1.2.59; Fnic: 1.6.0.12

Log in to Cisco UCS 6200 Fabric Interconnect

Cisco UCS Manager

The IP KVM enables the administrator to begin the installation of the operating system (OS) through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

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

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. To download the Cisco UCS Manager software, click the Launch UCS Manager link.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click the Servers tab.
7. Select Servers > Service Profiles > root > `VM-Host-Infra-01`.

8. Right-click `VM-Host-Infra-01` and select KVM Console.

9. If prompted to accept an Unencrypted KVM session, accept as necessary.

10. Select Servers > Service Profiles > root > `VM-Host-Infra-02`.

11. Right-click `VM-Host-Infra-02` and select KVM Console.

12. If prompted to accept an Unencrypted KVM session, accept as necessary.

Set Up VMware ESXi Installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.

2. Click Activate Virtual Devices

3. If prompted to accept an Unencrypted KVM session, accept as necessary.

4. Click Virtual Media and select Map CD/DVD.

5. Browse to the ESXi installer ISO image file and click Open.

6. Click Map Device.

7. Click the KVM tab to monitor the server boot.

8. Boot the server by selecting Boot Server and clicking OK. Click OK again.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware ESXi to the SAN-bootable LUN of the hosts, complete the following steps on each host:

1. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.

2. After the installer is finished loading, press Enter to continue with the installation.

3. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.

4. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.

5. Select the appropriate keyboard layout and press Enter.

6. Enter and confirm the root password and press Enter.

7. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
8. After the installation is complete, click on the Virtual Media tab and clear the ✓ mark next to the ESXi installation media. Click Yes.

Note: The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

9. From the KVM tab, press Enter to reboot the server.

Set Up Management Networking for ESXi Hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi Host VM-Host-Infra-01

To configure the `VM-Host-Infra-01` ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select the Configure the Management Network option and press Enter.
4. Select the VLAN (Optional) option and press Enter.
5. Enter the `<<var_ib_mgmt_vlan_id>>` and press Enter.
6. Select Network Adapters option and select vmnic4 (defined earlier as OOB vNIC) and press Enter.
7. From the Configure Management Network menu, select IP Configuration and press Enter.
8. Select the Set Static IP Address and Network Configuration option by using the space bar.
9. Enter the IP address for managing the first ESXi host: `<<var_vm_host_infra_01_ip>>`.
10. Enter the subnet mask for the first ESXi host.
11. Enter the default gateway for the first ESXi host.
12. Press Enter to accept the changes to the IP configuration.
13. Select the IPv6 Configuration option and press Enter.
14. Using the spacebar, unselect Enable IPv6 (restart required) and press Enter.
15. Select the DNS Configuration option and press Enter.

Note: Because the IP address is assigned manually, the DNS information must also be entered manually.

16. Enter the IP address of the primary DNS server.

17. Optional: Enter the IP address of the secondary DNS server.

18. Enter the fully qualified domain name (FQDN) for the first ESXi host.

19. Press Enter to accept the changes to the DNS configuration.

20. Press Esc to exit the Configure Management Network submenu.

21. Press Y to confirm the changes and return to the main menu.

22. The ESXi host reboots. After reboot, press F2 and log back in as root.

23. Select Test Management Network to verify that the management network is set up correctly and press Enter.

24. Press Enter to run the test.

25. Press Enter to exit the window.

26. Press Esc to log out of the VMware console.

ESXi Host VM-Host-Infra-02

To configure the `VM-Host-Infra-02` ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.

2. Log in as `root` and enter the corresponding password.

3. Select the Configure the Management Network option and press Enter.

4. Select the VLAN (Optional) option and press Enter.

5. Enter the `<<var_ib-mgmt_vlan_id>>` and press Enter.

6. Select Network Adapters option and select vmnic4 (defined earlier as OOB vNIC) and press Enter.

7. From the Configure Management Network menu, select IP Configuration and press Enter.

8. Select the Set Static IP Address and Network Configuration option by using the space bar.

9. Enter the IP address for managing the second ESXi host: `<<var_vm_host_infra_02_ip>>`.

10. Enter the subnet mask for the second ESXi host.

11. Enter the default gateway for the second ESXi host.

12. Press Enter to accept the changes to the IP configuration.

13. Select the IPv6 Configuration option and press Enter.

14. Using the spacebar, clear Enable IPv6 (restart required) and press Enter.

15. Select the DNS Configuration option and press Enter.

Note: Because the IP address is assigned manually, the DNS information must also be entered manually.

16. Enter the IP address of the primary DNS server.

17. Optional: Enter the IP address of the secondary DNS server.

18. Enter the FQDN for the second ESXi host.

19. Press Enter to accept the changes to the DNS configuration.

20. Press Esc to exit the Configure Management Network submenu.

21. Press Y to confirm the changes and return to the main menu.

22. The ESXi host reboots. After reboot, press F2 and log back in as root.

23. Select Test Management Network to verify that the management network is set up correctly and press Enter.

24. Press Enter to run the test.

25. Press Enter to exit the window.

26. Press Esc to log out of the VMware console.

Download VMware vSphere Client

To download the VMware vSphere Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the `VM-Host-Infra-01` management IP address.

2. Download and install the vSphere Client.

Note: This application is downloaded from the VMware website and Internet access is required on the management workstation.

Download VMware vSphere CLI 5.5

1. Click the following link VMware vSphere CLI 5.5

2. Select your OS and click Download.

3. Save it to destination folder.

4. Run the VMware-vSphere-CLI-5.5.0.exe.

5. Click Next.

6. Accept the terms for the license and click Next.

7. Click Next on the Destination Folder screen.

8. Click Install.

9. Click Finish.

Note: Install VMware vSphere CLI 5.5 on the management workstation.

Log in to VMware ESXi Hosts by Using VMware vSphere Client

ESXi Host VM-Host-Infra-01

To log in to the `VM-Host-Infra-01` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-Infra-01` as the host you are trying to connect to: `<<var_vm_host_infra_01_ip>>`.

2. Enter `root` for the user name.

3. Enter the root password.

4. Click Login to connect.

ESXi Host VM-Host-Infra-02

To log in to the `VM-Host-Infra-02` ESXi host by using the VMware vSphere Client, complete the following steps:

1. Open the recently downloaded VMware vSphere Client and enter the IP address of `VM-Host-Infra-02` as the host you are trying to connect to: `<<var_vm_host_infra_02_ip>>`.

2. Enter `root` for the user name.

3. Enter the root password.

Set Up VMkernel Ports and Virtual Switch

ESXi Host VM-Host-Infra-01

To set up the VMkernel ports and the virtual switches on the `VM-Host-Infra-01`. ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.

2. Click the Configuration tab.

3. In the Hardware pane, click Networking.

4. On the right side of `vSwitch0`, click Properties.

5. Select the vSwitch configuration and click Edit.

6. From the General tab, change the MTU to `9000`.

7. Click OK.

8. Click Network Adapters tab, click Add.

9. Select vmnic5 and click Next.

10. Click Next and then click Finish.

11. Click the Ports tab.

12. Select the Management Network configuration and click Edit.

13. Change the network label to <`VMkernel-MGMT>` and select the Management Traffic checkbox.

14. Click OK to finalize the edits for Management Network.

15. Select the VM Network configuration and click Edit.

16. Change the network label to <`MGMT Network>` and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.

17. Click OK to finalize the edits for VM Network.

18. Click Close.

19. On the right side of `iScsiBootvSwitch`, click Properties.

20. Select the vSwitch configuration and click Edit

21. Change the MTU to 9000.

22. Click OK.

23. Select iScsiBootPG and click Edit.

24. Change the Network Label to <`VMkernel-iSCSI-A>.`

25. Change the MTU to 9000.

26. Click OK.

27. Click Close.

28. In the vSphere Standard Switch view, click Add Networking.

29. Select VMkernel and click Next.

30. Select Create a vSphere standard switch to create a new vSphere standard switch.

31. Select the check boxes for the network adapter vmnic3.

32. Click Next.

33. Change the network label to <VMkernel-iSCSI-B>.

34. Click Next.

35. Enter the IP address and the subnet mask for the iSCSI VLAN B interface for VM-Host-Infra-01.

Note: To obtain the iSCSI IP address information; login to the Cisco UCS Manager and in the servers tab select the corresponding service profiles. In the right pane, click the boot order and select the iSCSI-B-vNIC; click set iSCSI boot parameters; the IP address should appear as the initiator IP address.

36. Click Next.

37. Click Finish.

38. On the right side of vSwitch1, click Properties.

39. Select the vSwitch configuration and click Edit.

40. Change the MTU to 9000.

41. Click OK

42. Select VMkernel-iSCSI-B and click Edit.

43. Change the MTU to 9000.

44. Click OK.

45. Click Close.

46. In the vSphere Standard Switch view, click Add Networking.

47. Select VMkernel and click Next.

48. Select Create a vSphere standard switch to create a new vSphere standard switch.

49. Select vmnic6 and vmnic7 and click Next.

50. Change the network label to <VMkernel-NFS> and enter <var_nfs_vlan_id> in the VLAN ID (Optional) field.

51. Click Next.

52. Enter the IP address `<<var_nfs_vlan_ip_host_01>>` and the subnet mask `<<var_nfs_vlan_ip_mask_host_01>>` for the NFS VLAN interface for `VM-Host-Infra-01`.

53. To continue with the NFS VMkernel creation, click Next.

54. To finalize the creation of the NFS VMkernel interface, click Finish.

55. Select the `<vSwitch>` configuration and click Edit.

56. Change the MTU to `9000`.

57. Click OK.

58. Select the `<VMkernel-NFS>` configuration and click Edit.

59. Change the MTU to `9000`.

60. Click OK to finalize the edits for the VMkernel-NFS network. The properties vSwitch2 should be similar to the following example:

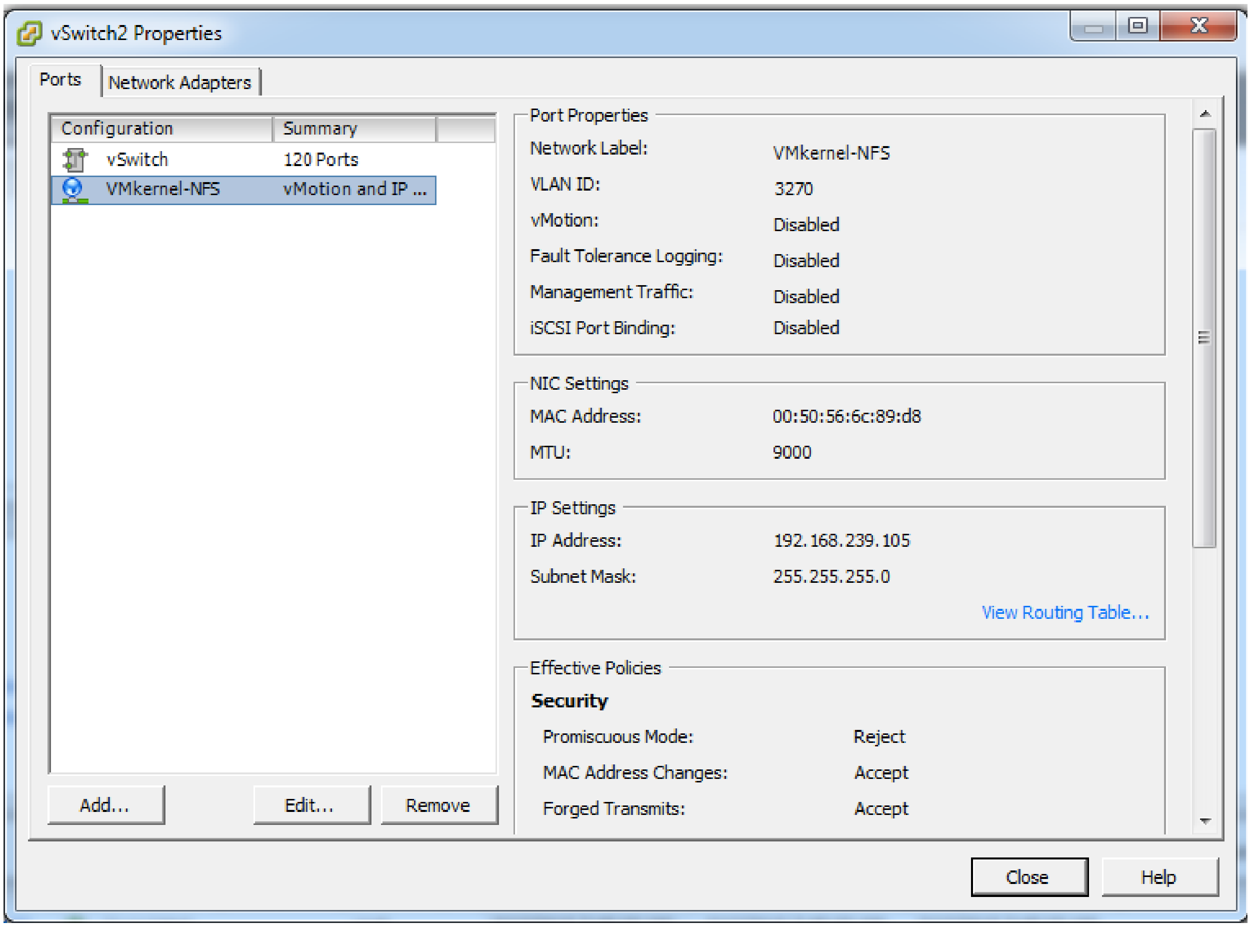

Configuration	Summary
vSwitch	120 Ports
VMkernel-NFS	vMotion and IP ...

Port Properties

Network Label:	VMkernel-NFS
VLAN ID:	3270
vMotion:	Disabled
Fault Tolerance Logging:	Disabled
Management Traffic:	Disabled
iSCSI Port Binding:	Disabled

NIC Settings

MAC Address:	00:50:56:6c:89:d8
MTU:	9000

IP Settings

IP Address:	192.168.239.105
Subnet Mask:	255.255.255.0

Effective Policies

Security

Promiscuous Mode:	Reject
MAC Address Changes:	Accept
Forged Transmits:	Accept

61. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

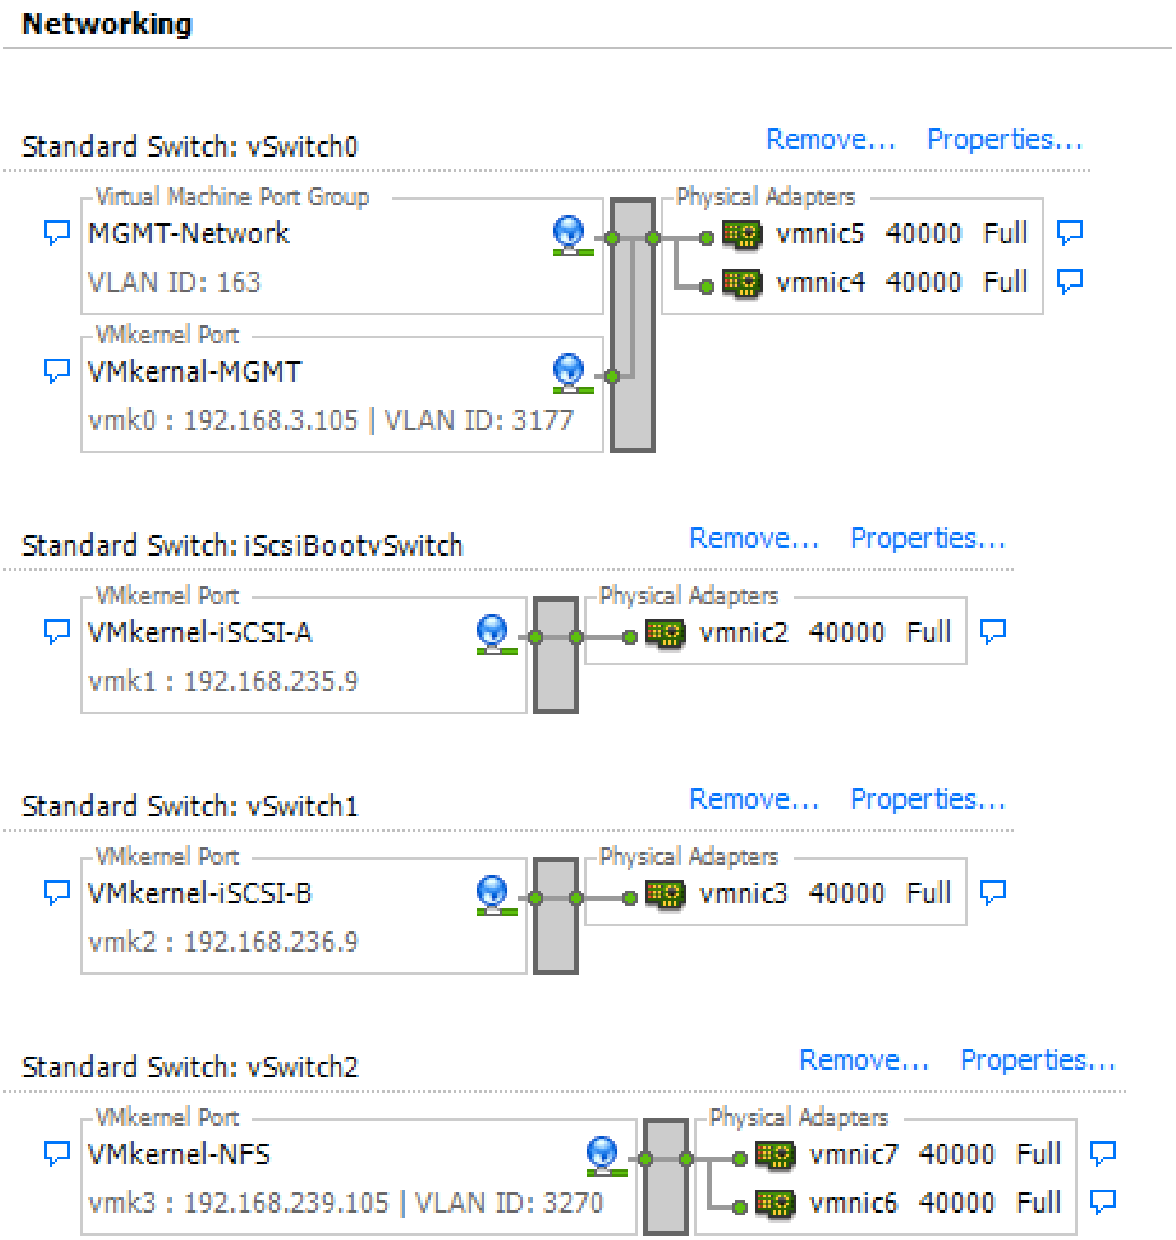

Networking

Standard Switch: vSwitch0 Remove... Properties...

Virtual Machine Port Group
MGMT-Network
VLAN ID: 163

VMkernel Port
VMkernal-MGMT
vmk0 : 192.168.3.105 | VLAN ID: 3177

Physical Adapters
vmnic5 40000 Full
vmnic4 40000 Full

Standard Switch: iScsiBootvSwitch Remove... Properties...

VMkernel Port
VMkernel-iSCSI-A
vmk1 : 192.168.235.9

Physical Adapters
vmnic2 40000 Full

Standard Switch: vSwitch1 Remove... Properties...

VMkernel Port
VMkernel-iSCSI-B
vmk2 : 192.168.236.9

Physical Adapters
vmnic3 40000 Full

Standard Switch: vSwitch2 Remove... Properties...

VMkernel Port
VMkernel-NFS
vmk3 : 192.168.239.105 | VLAN ID: 3270

Physical Adapters
vmnic7 40000 Full
vmnic6 40000 Full

ESXi Host VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the `VM-Host-Infra-02`. ESXi host, complete the following steps:

1. From the vSphere Client, select the host in the inventory.
2. Click the Configuration tab.
3. In the Hardware pane, click Networking.
4. On the right side of `vSwitch0`, click Properties.

5. Select the vSwitch configuration and click Edit.
6. From the General tab, change the MTU to `9000`.
7. Click OK.
8. Click Network Adapters tab, click Add.
9. Select vmnic5 and click Next.
10. Click Next and then click Finish.
11. Click the Ports tab.
12. Select the Management Network configuration and click Edit.
13. Change the network label to <`VMkernel-MGMT>` and select the Management Traffic checkbox.
14. Click OK to finalize the edits for Management Network.
15. Select the VM Network configuration and click Edit.
16. Change the network label to <`MGMT Network>` and enter `<<var_ib-mgmt_vlan_id>>` in the VLAN ID (Optional) field.
17. Click OK to finalize the edits for VM Network.
18. Click Close.
19. On the right side of `iScsiBootvSwitch`, click Properties.
20. Select the vSwitch configuration and click Edit.
21. Change the MTU to 9000.
22. Click OK.
23. Select iScsiBootPG and click Edit.
24. Change the Network Label to <`VMkernel-iSCSI-A>.`
25. Change the MTU to 9000.
26. Click OK.
27. Click Close.
28. In the vSphere Standard Switch view, click Add Networking.
29. Select VMkernel and click Next.
30. Select Create a vSphere standard switch to create a new vSphere standard switch.

31. Select the check boxes for the network adapter vmnic3.

32. Click Next.

33. Change the network label to <VMkernel-iSCSI-B>.

34. Click Next.

35. Enter the IP address and the subnet mask for the iSCSI VLAN interface for VM-Host-Infra-02.

36. Click Next.

37. Click Finish.

38. On the right side of vSwitch1, click Properties.

39. Select the vSwitch configuration and click Edit.

40. Change the MTU to 9000.

41. Click OK.

42. Select VMkernel-iSCSI-B and click Edit.

43. Change the MTU to 9000.

44. Click OK.

45. Click Close.

46. In the vSphere Standard Switch view, click Add Networking.

47. Select VMkernel and click Next.

48. Select Create a vSphere standard switch to create a new vSphere standard switch.

49. Select vmnic6 and vmnic7 and click Next.

50. Change the network label to <VMkernel-NFS> and enter <<var_nfs_vlan_id>> in the VLAN ID (Optional) field.

51. Click Next.

52. Enter the IP address <<var_nfs_vlan_ip_host_02>> and the subnet mask <<var_nfs_vlan_ip_mask_host_02>> for the NFS VLAN interface for VM-Host-Infra-02.

53. To continue with the NFS VMkernel creation, click Next.

54. To finalize the creation of the NFS VMkernel interface, click Finish.

55. Select the <<vSwitch>> configuration and click Edit.

56. Change the MTU to `9000`.

57. Click OK.

58. Select the `<<VMkernel-NFS>>` configuration and click Edit.

59. Change the MTU to `9000`.

60. Click OK to finalize the edits for the VMkernel-NFS network. The properties vSwitch2 should be similar to the following example:

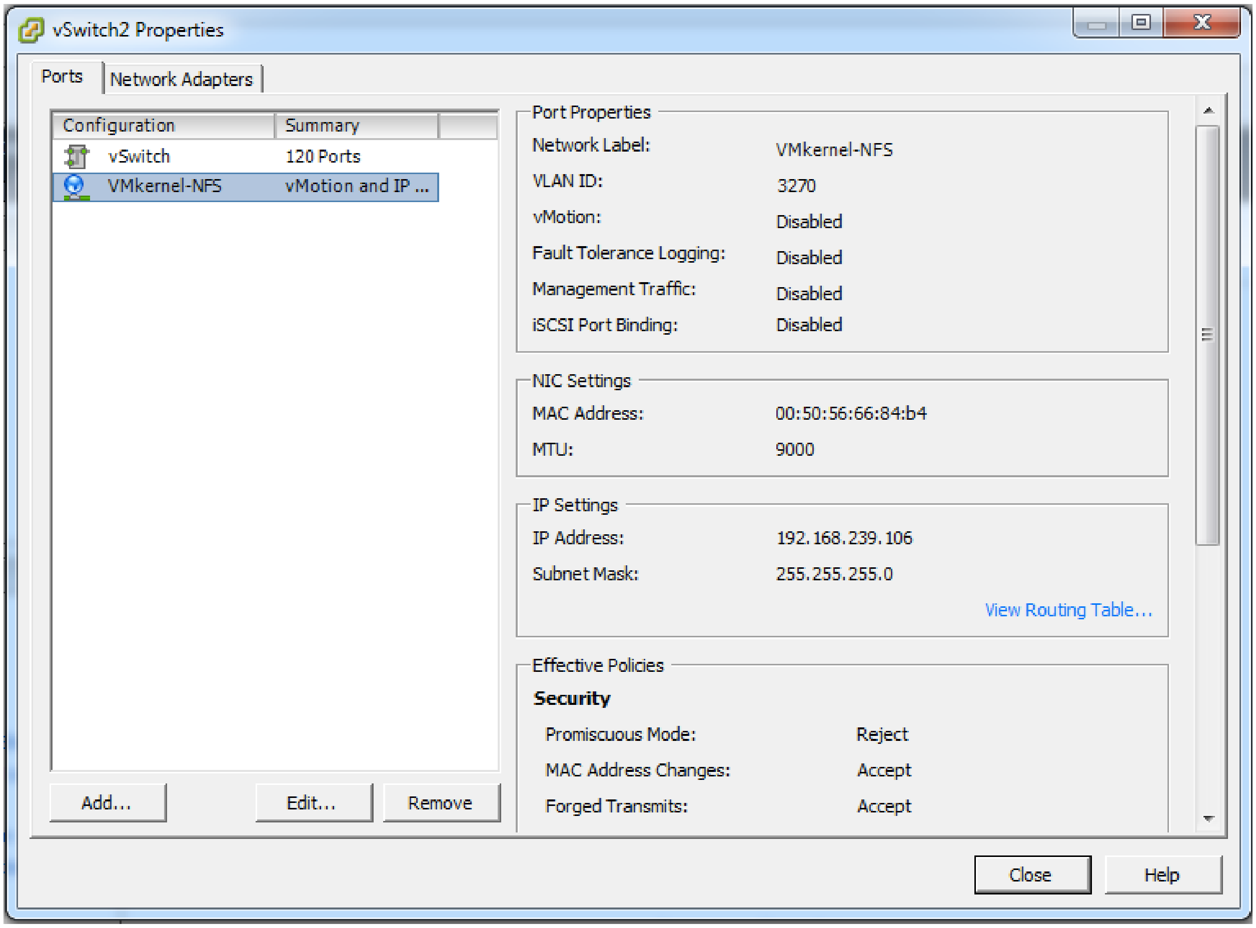

vSwitch2 Properties

Configuration	Summary
vSwitch	120 Ports
VMkernel-NFS	vMotion and IP ...

Port Properties

Network Label: VMkernel-NFS
VLAN ID: 3270
vMotion: Disabled
Fault Tolerance Logging: Disabled
Management Traffic: Disabled
iSCSI Port Binding: Disabled

NIC Settings

MAC Address: 00:50:56:66:84:b4
MTU: 9000

IP Settings

IP Address: 192.168.239.106
Subnet Mask: 255.255.255.0

Effective Policies

Security

Promiscuous Mode: Reject
MAC Address Changes: Accept
Forged Transmits: Accept

61. To finalize the ESXi host networking setup, close the dialog box. The networking for the ESXi host should be similar to the following example:

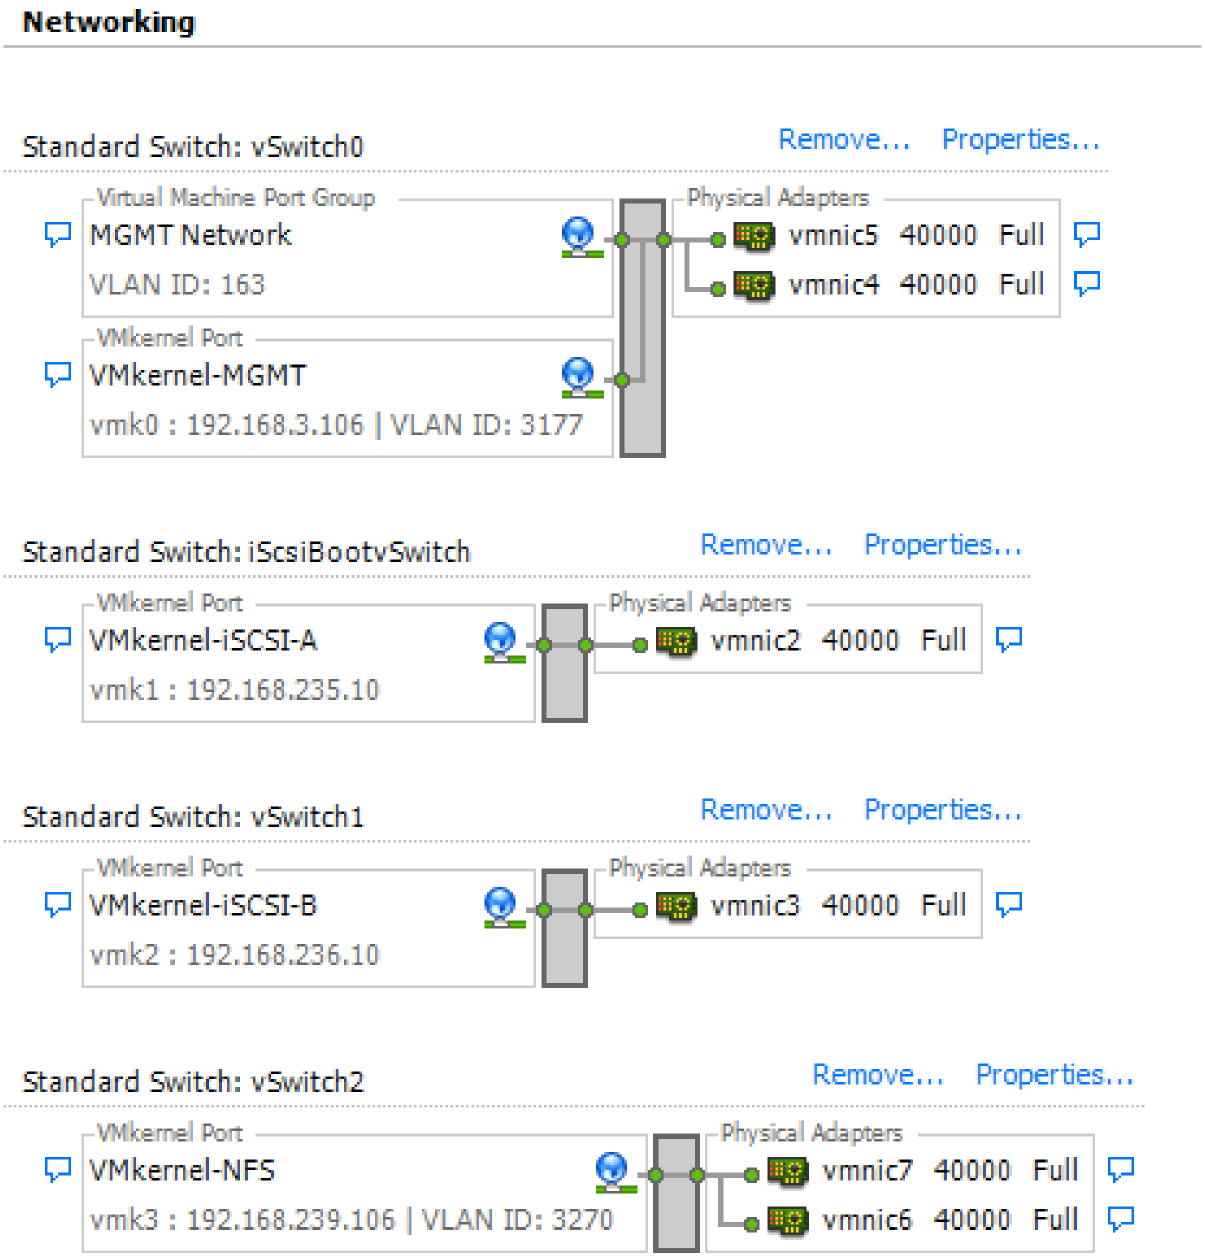

Networking

Standard Switch: vSwitch0 Remove... Properties...

Virtual Machine Port Group
MGMT Network
VLAN ID: 163

VMkernel Port
VMkernel-MGMT
vmk0 : 192.168.3.106 | VLAN ID: 3177

Physical Adapters
vmnic5 40000 Full
vmnic4 40000 Full

Standard Switch: iScsiBootvSwitch Remove... Properties...

VMkernel Port
VMkernel-iSCSI-A
vmk1 : 192.168.235.10

Physical Adapters
vmnic2 40000 Full

Standard Switch: vSwitch1 Remove... Properties...

VMkernel Port
VMkernel-iSCSI-B
vmk2 : 192.168.236.10

Physical Adapters
vmnic3 40000 Full

Standard Switch: vSwitch2 Remove... Properties...

VMkernel Port
VMkernel-NFS
vmk3 : 192.168.239.106 | VLAN ID: 3270

Physical Adapters
vmnic7 40000 Full
vmnic6 40000 Full

Setup iSCSI Multipathing

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To setup 4 iSCSI paths between storage and the ESXi host, complete the following steps on each ESXi host:

1. From the vSphere Client, click Storage Adapters in the Hardware pane.
2. Select the iSCSI Software Adapter and click Properties.
3. Select the Dynamic Discovery tab and click Add.
4. Enter the IP address of iscsi_lif01a.

5. Click OK.

6. Repeat putting in the IP addresses of iscsi_lif01b, iscsi_lif02a and iscsi_lif02b.

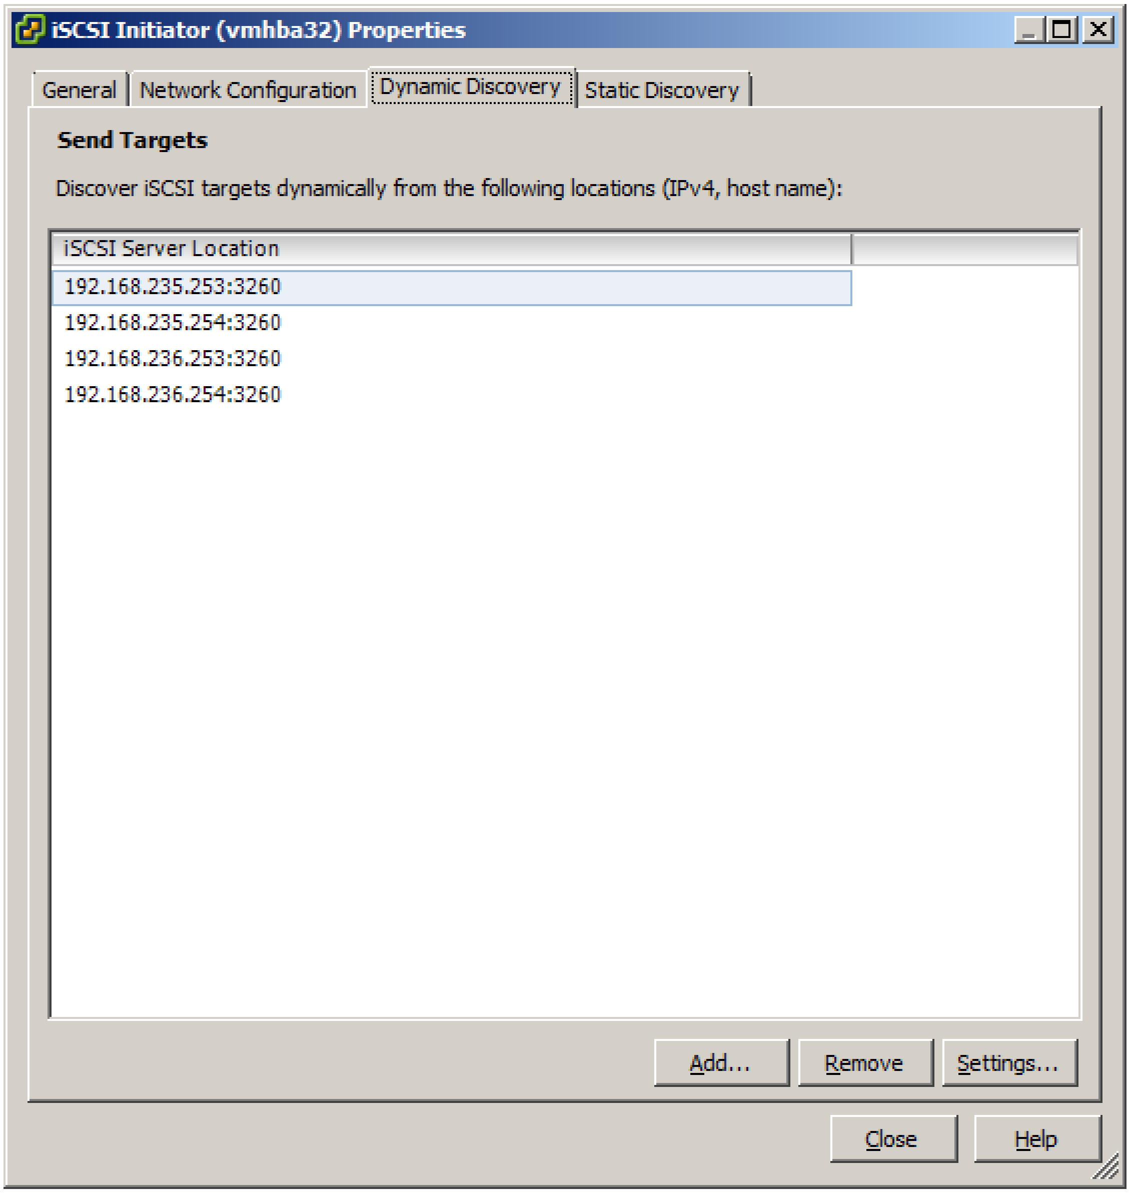

Send Targets

Discover iSCSI targets dynamically from the following locations (IPv4, host name):

iSCSI Server Location
192.168.235.253:3260
192.168.235.254:3260
192.168.236.253:3260
192.168.236.254:3260

7. Click Close and then click yes to rescan the host bus adapter.

8. You should now see 4 connected paths in the Details pane.

Install VMware Drivers for the Cisco Virtual Interface Card (VIC)

Download and extract the following VMware VIC Drivers to the Management workstation:

fnic Driver version 1.6.0.12b

enic Driver version 2.2.2.62

Note: Click this link for driver download instructions: http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/vic_drivers/install/ESX/2-0/b_Cisco_VIC_Drivers_for_ESX_Installation_Guide/Cisco_VIC_Drivers_for_ESX_Installation_Guide_chapter2.html.

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each vSphere Client, select the host in the inventory.
2. Click the Summary tab to view the environment summary.
3. From Resources > Storage, right-click datastore1 and select Browse Datastore.
4. Click the fourth button and select Upload File.
5. Navigate to the saved location for the downloaded VIC drivers and select fnic_driver_1.6.0.12b-offline_bundle-2340688.zip.
6. Click Open and Yes to upload the file to datastore1.
7. Click the fourth button and select Upload File.
8. Navigate to the saved location for the downloaded VIC drivers and select enic-2.1.2.62-esx55-offline_bundle-2340678.zip.
9. Click Open and Yes to upload the file to datastore1.
10. Make sure the files have been uploaded to both ESXi hosts.
11. From the management workstation, open the VMware vSphere Remote CLI that was previously installed.
12. At the command prompt, run the following commands to account for each host

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
update -d /vmfs/volumes/datastore1/ fnic_driver_1.6.0.12b-offline_bundle-
2340688.zip

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
update -d /vmfs/volumes/datastore1/ fnic_driver_1.6.0.12b-offline_bundle-
2340688.zip

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> software vib
update -d /vmfs/volumes/datastore1/ enic-2.1.2.62-esx55-offline_bundle-
2340678.zip

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> software vib
update -d /vmfs/volumes/datastore1/ enic-2.1.2.62-esx55-offline_bundle-
2340678.zip
```

13. Back in the vSphere Client for each host, right-click the host and select Reboot.

14. Click Yes and OK to reboot the host.

15. Log back into each host with vSphere Client.

Mount Required Datastores

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To mount the required datastores, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.

2. To enable configurations, click the Configuration tab.

3. Click Storage in the Hardware pane.

4. From the Datastores area, click Add Storage to open the Add Storage wizard.

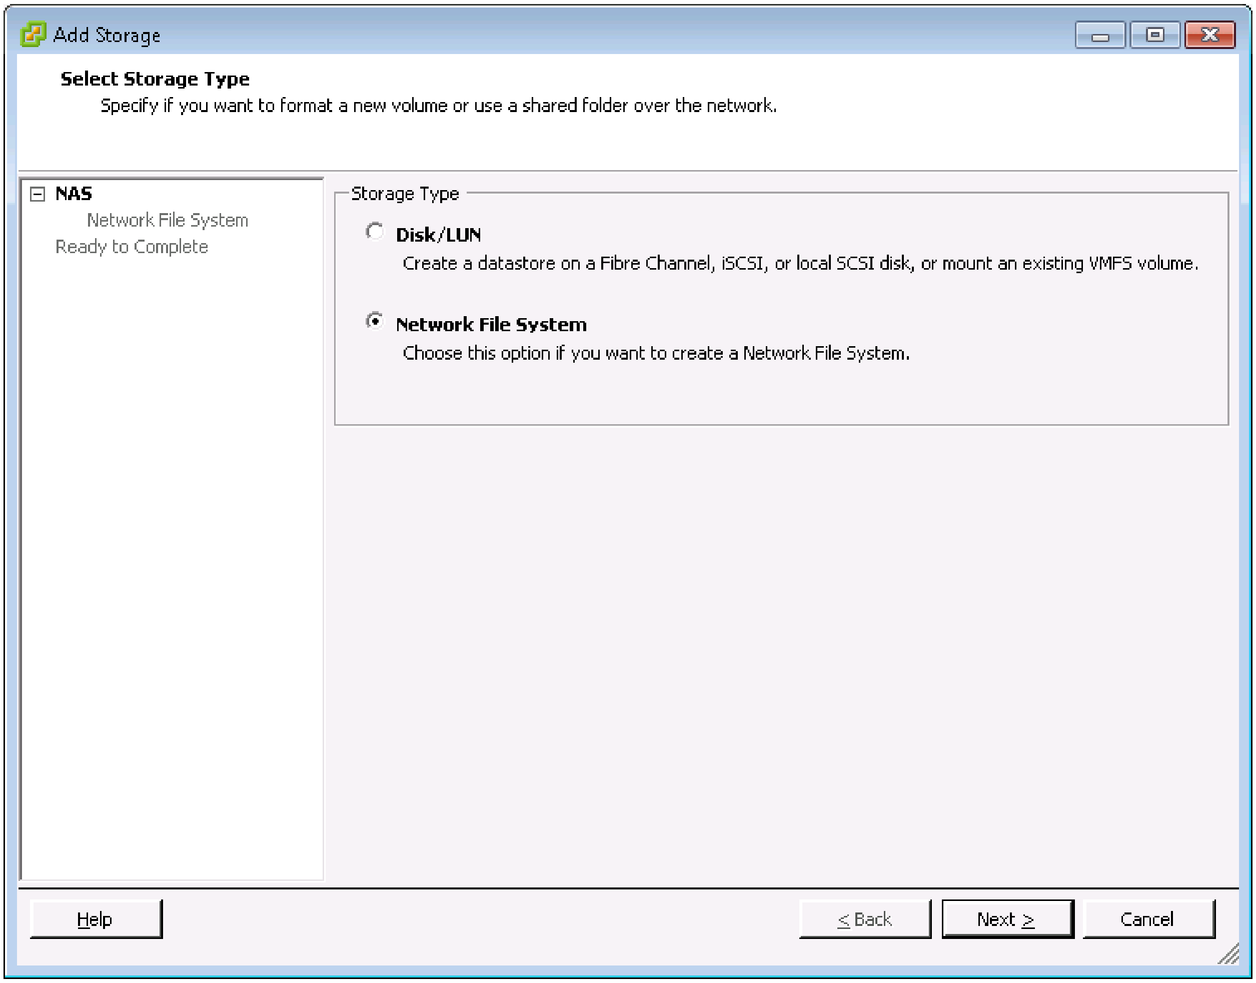

Add Storage

Select Storage Type

Specify if you want to format a new volume or use a shared folder over the network.

Storage Type

○ **Disk/LUN**
Create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume.

● **Network File System**
Choose this option if you want to create a Network File System.

5. Select Network File System and click Next.

6. The wizard prompts for the location of the NFS export. Enter `<<var_node02_nfs_lif_infra_datastore_1_ip>>` as the IP address for `nfs_lif_infra_datastore_1`.

7. Enter `/infra_datastore_1` as the path for the NFS export.

8. Confirm that the Mount NFS read only checkbox is not selected.

9. Enter `infra_datastore_1` as the datastore name.

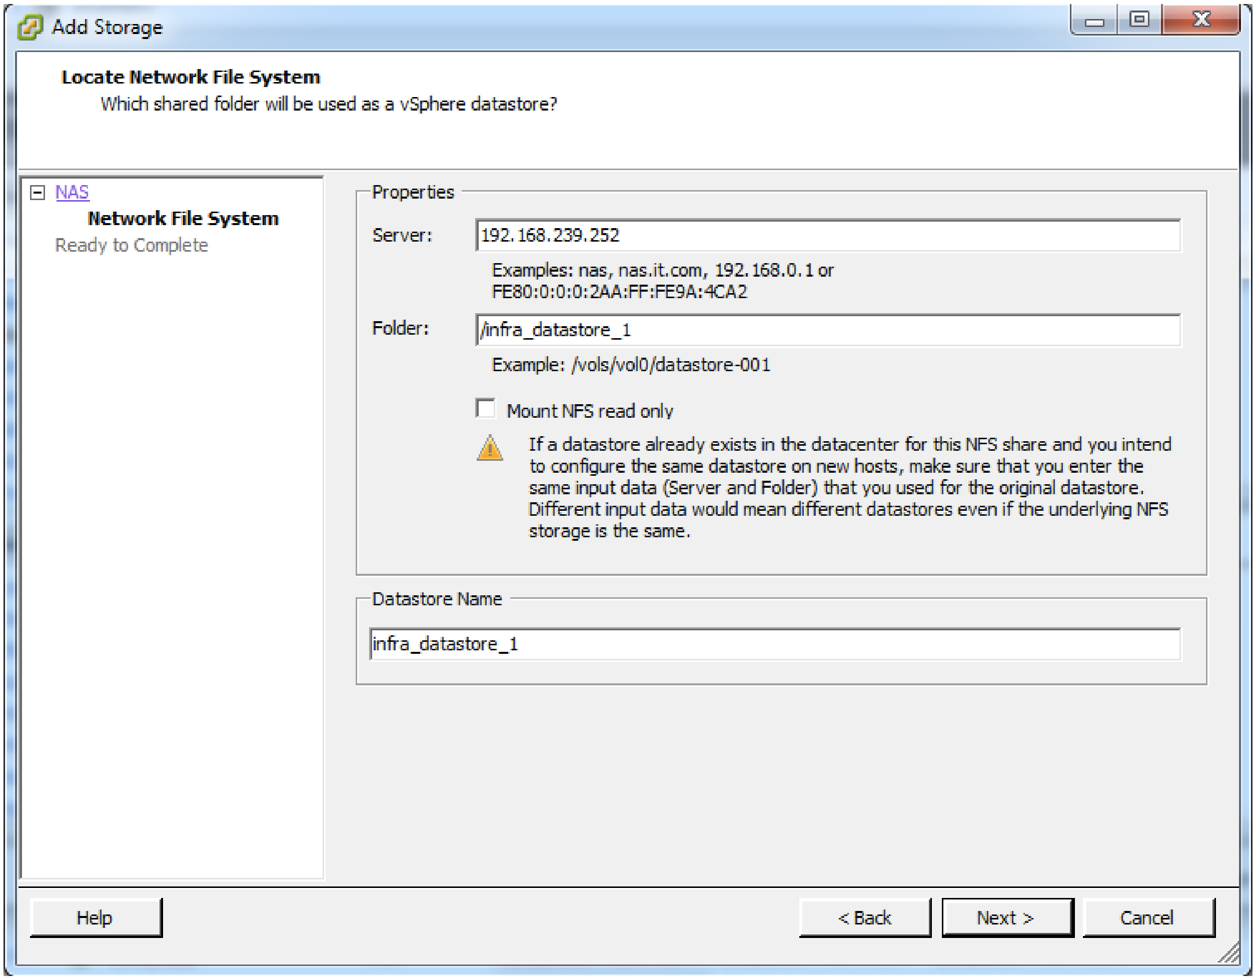

Add Storage

Locate Network File System

Which shared folder will be used as a vSphere datastore?

Server: 192.168.239.252

Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: /infra_datastore_1

Example: /vols/vol0/datastore-001

□ Mount NFS read only

If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

Datastore Name: infra_datastore_1

10. To continue with the NFS datastore creation, click Next.

11. To finalize the creation of the NFS datastore, click Finish.

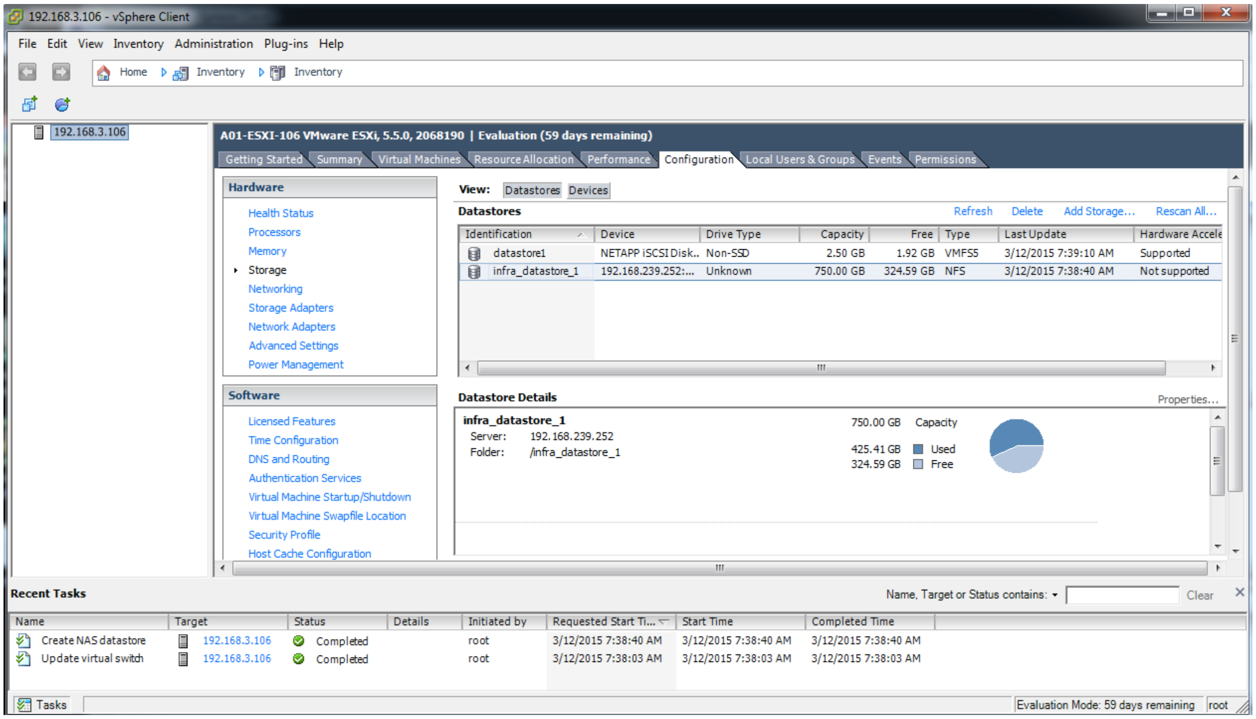

A01-ESXI-106 VMware ESXi, 5.5.0, 2068190 | Evaluation (59 days remaining)

Datastores

Identification	Device	Drive Type	Capacity	Free	Type	Last Update	Hardware Accele
datastore1	NETAPP iSCSI Disk...	Non-SSD	2.50 GB	1.92 GB	VMFS5	3/12/2015 7:39:10 AM	Supported
infra_datastore_1	192.168.239.252:...	Unknown	750.00 GB	324.59 GB	NFS	3/12/2015 7:38:40 AM	Not supported

Datastore Details

infra_datastore_1
Server: 192.168.239.252
Folder: /infra_datastore_1

750.00 GB Capacity
425.41 GB ▇ Used
324.59 GB ▢ Free

Recent Tasks

Name, Target or Status contains:

Name	Target	Status	Details	Initiated by	Requested Start Ti...	Start Time	Completed Time
Create NAS datastore	192.168.3.106	Completed		root	3/12/2015 7:38:40 AM	3/12/2015 7:38:40 AM	3/12/2015 7:38:40 AM
Update virtual switch	192.168.3.106	Completed		root	3/12/2015 7:38:03 AM	3/12/2015 7:38:03 AM	3/12/2015 7:38:03 AM

Evaluation Mode: 59 days remaining | root

12. From the Datastores area, click Add Storage to open the Add Storage wizard.

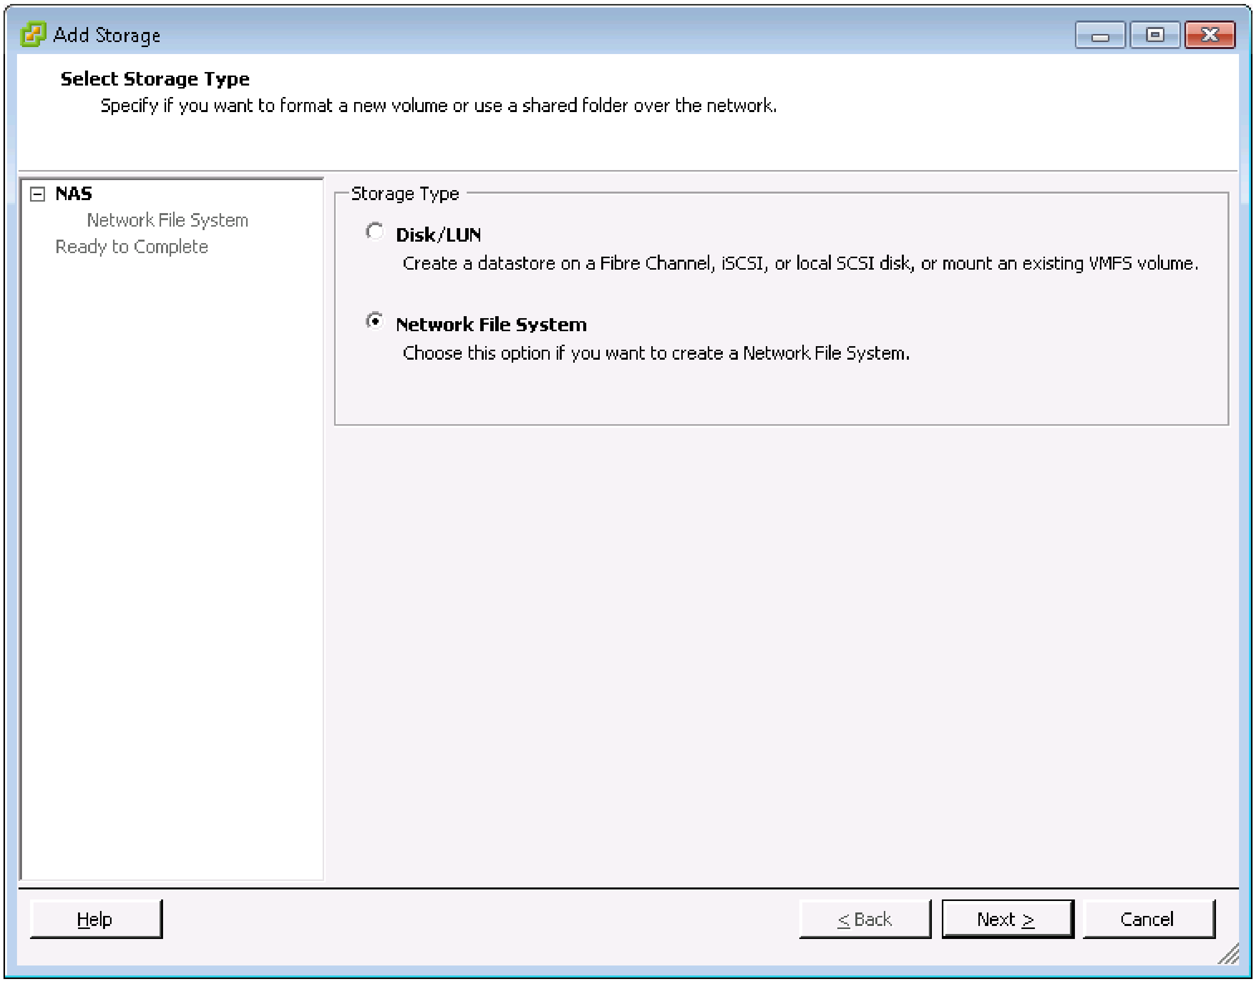

Select Storage Type

Specify if you want to format a new volume or use a shared folder over the network.

Storage Type

○ **Disk/LUN**
Create a datastore on a Fibre Channel, iSCSI, or local SCSI disk, or mount an existing VMFS volume.

● **Network File System**
Choose this option if you want to create a Network File System.

13. Select Network File System and click Next.

14. The wizard prompts for the location of the NFS export. Enter `<<var_node01_nfs_lif_infra_swap_ip>>` as the IP address for `nfs_lif_infra_swap`.

15. Enter `/infra_swap` as the path for the NFS export.

16. Confirm that the Mount NFS read only checkbox is not selected.

17. Enter `infra_swap` as the datastore name.

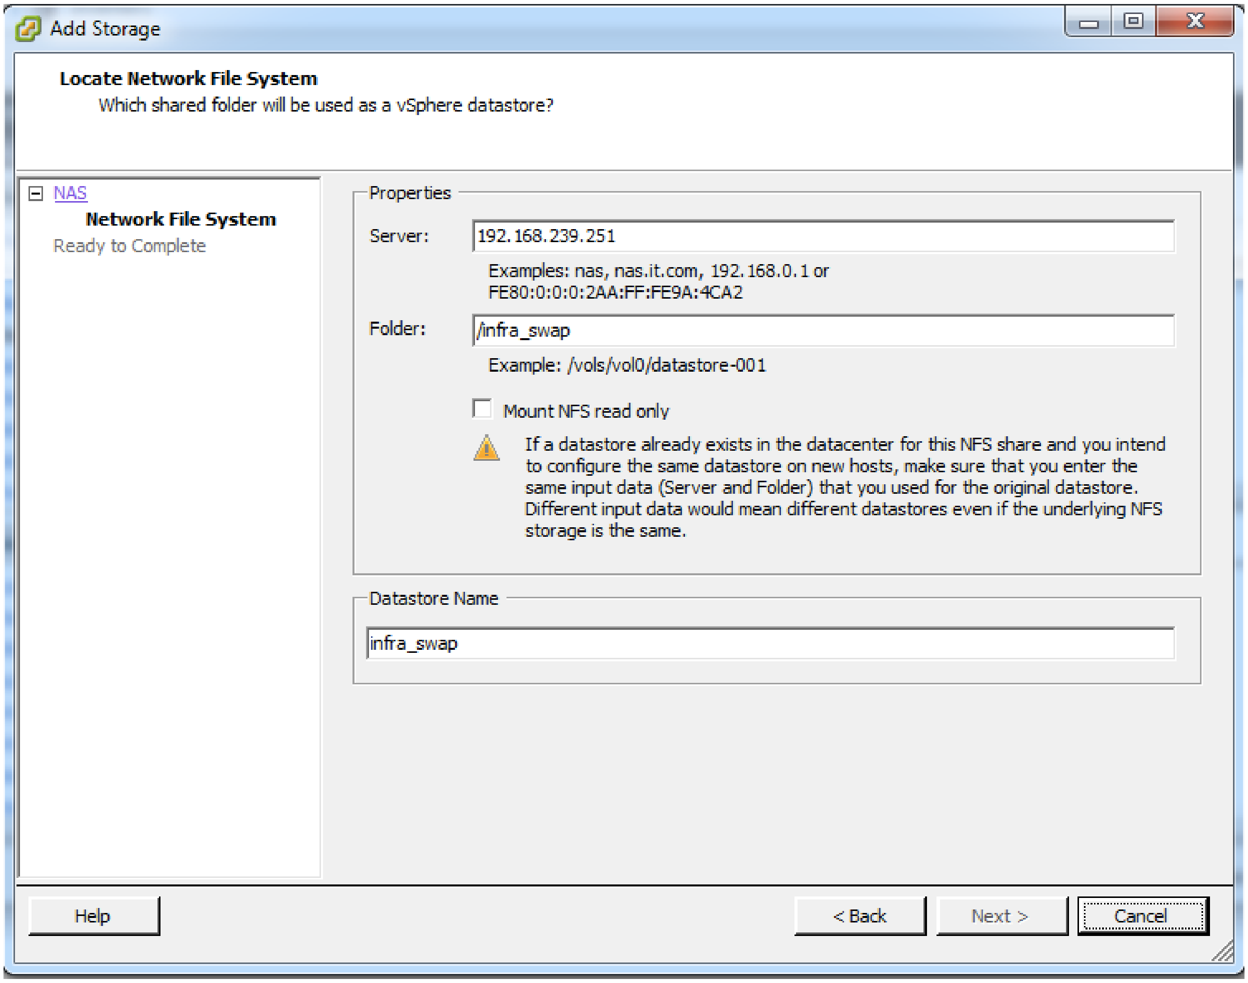

Locate Network File System

Which shared folder will be used as a vSphere datastore?

Properties

Server: 192.168.239.251

Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: /infra_swap

Example: /vols/vol0/datastore-001

□ Mount NFS read only

⚠ If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

Datastore Name

infra_swap

18. To continue with the NFS datastore creation, click Next.

19. To finalize the creation of the NFS datastore, click Finish.

Configure NTP on ESXi Hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure Network Time Protocol (NTP) on the ESXi hosts, complete the following steps on each host:

1. From the vSphere Client, select the host in the inventory.

2. To enable configurations, click the Configuration tab.

3. Click Time Configuration in the Software pane.

4. Click Properties at the upper-right side of the window.

5. At the bottom of the Time Configuration dialog box, click Options.

6. In the NTP Daemon (ntpd) Options dialog box, complete the following steps:

   a. Click General in the left pane and select Start and stop with host.

   b. Click NTP Settings in the left pane and click Add.

7. In the Add NTP Server dialog box, enter `<<var_global_ntp_server_ip>>` as the IP address of the NTP server and click OK.

8. In the NTP Daemon Options dialog box, select the Restart NTP service to apply changes checkbox and click OK.

9. In the Time Configuration dialog box, complete the following steps:

   a. Select the NTP Client Enabled checkbox and click OK.

   b. Verify that the clock is now set to approximately the correct time.

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

Move VM Swap File Location

ESXi VM-Host-Infra-01 and VM-Host-Infra-02

To move the VM swap file location, complete the following steps on each ESXi host:

1. From the vSphere Client, select the host in the inventory.

2. To enable configurations, click the Configuration tab.

3. Click Virtual Machine Swapfile Location in the Software pane.

4. Click Edit at the upper-right side of the window.

5. Select Store the swapfile in a swapfile datastore selected below.

6. Select the `<datastore_name>` datastore in which to house the swap files.

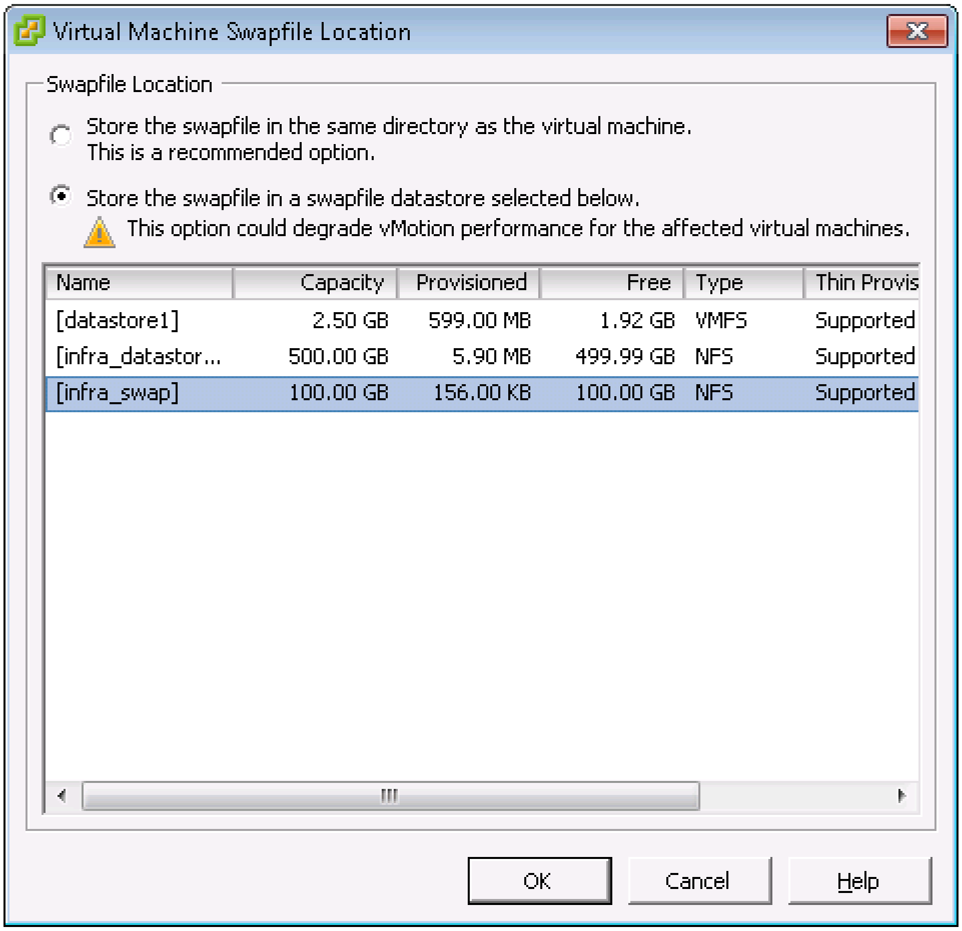

Virtual Machine Swapfile Location

Swapfile Location

○ Store the swapfile in the same directory as the virtual machine.
This is a recommended option.

● Store the swapfile in a swapfile datastore selected below.
⚠ This option could degrade vMotion performance for the affected virtual machines.

Name	Capacity	Provisioned	Free	Type	Thin Provis
[datastore1]	2.50 GB	599.00 MB	1.92 GB	VMFS	Supported
[infra_datastor...	500.00 GB	5.90 MB	499.99 GB	NFS	Supported
[infra_swap]	100.00 GB	156.00 KB	100.00 GB	NFS	Supported

7. Click OK to finalize moving the swap file location.

VMware vCenter 5.5 Update 2

The procedures in the following subsections provide detailed instructions for installing VMware vCenter 5.5 Update 1 in an environment. After the procedures are completed, a VMware vCenter Server will be configured.

Build Microsoft SQL Server VM

To build a SQL Server Virtual Machine (VM) for the `VM-Host-Infra-01`, complete the following steps:

1. Log in to the host by using the VMware vSphere Client.
2. In the vSphere Client, choose the host in the inventory pane.
3. Right-click the host and choose New Virtual Machine.
4. Choose Custom and click Next.
5. Enter a name for the VM. Click Next.
6. Choose <<`datastore_name`>>. Click Next.

7. Choose Virtual Machine Version: 8. Click Next.

8. Verify that the Windows option and the Microsoft Windows Server 2012 (64 Bit) version are selected. Click Next.

9. Choose two virtual sockets and one core per virtual socket. Click Next.

10. Choose 8GB of memory. Click Next.

11. Choose one network interface card (NIC).

12. For NIC 1, choose the `<<Network>>` Network option and the VMXNET 3 adapter. Click Next.

13. Keep the LSI Logic SAS option for the SCSI controller Selected. Click Next.

14. Keep the Create a New Virtual Disk Option selected. Click Next.

15. Make the disk size at least 80GB. Click Next.

16. Click Next.

17. Check the edit the virtual machine setting before completion. Click Continue.

18. Click the Options tab.

19. Choose Boot Options.

20. Check the Force Bios Setup check box.

21. Click Finish.

22. From the left pane, expand the host filed by click the plus sign (+).

23. Right-click the newly created SQL Server VM and click Open Console.

24. Click the third button (green right arrow) to power on the VM.

25. Click the ninth button (CD with a wrench) to map the Windows Server 2012 R2 ISO, and then choose Connect to ISO Image on Local Disk.

26. Navigate to Windows Server 2012 R2 ISO, select it, and click Open.

27. In the BIOS Setup Utility window and use the right arrow key to navigate to the Boot menu. Use the down arrow key to choose CD-ROM Drive. Press the plus (+) key twice to move CD-ROM Drive to the top of the list. Press F10 and Enter to save the selection and exit the BIOS Setup Utility.

28. The Windows Installer boots. Choose the appropriate language, time and currency format and keyboard. Click Next.

29. Click Install Now.

30. Make sure that Windows 2012 R2 Standard (Server with a GUI) option is selected. Click Next.

31. Read and accept the license terms and click Next.

32. Choose Custom (Advanced). Make sure that Disk0 Unallocated Space is selected. Click Next to allow the Windows installation to complete

33. After the Windows installation is complete and the VM has rebooted, Click OK to set the Administrator password.

34. Enter and confirm the Administrator password and click Finish.

35. After logging into the VM desktop, from the VM console window, choose the VM menu. Under Guest, choose Install/Upgrade VMware Tools. Click OK.

36. Click OK Installing the VMware tools package will greatly enhance graphics and mouse performance in your virtual machine.

37. If prompted to eject the Windows installation media before running the setup for the VMware tools, Click OK, then click OK.

38. Navigate on the CD-ROM drive, choose Run setup64.exe.

39. In the VMware Tools installer window, click Next.

40. Make sure that Typical is selected and click Next.

41. Click Install.

42. Click Finish.

43. Click Yes to restart the VM**.**

44. After the reboot is complete, choose the VM menu. Under Guest, choose Ctrl+Alt+Del and then enter the password to log into the VM.

45. Set the time zone for the VM, IP address, gateway, and host name. Add the VM to the Windows AD domain.

46. If necessary, activate Windows.

47. Log back into the VM and download and install all required Windows updates.

Install Microsoft SQL Server 2012 SP1

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

1. Connect to an AD Domain Controller in the Windows Domain and add an admin user in Active Directory Users and Computer tool. This user should be member of the Domain Administrator Security Group.

2. Log into the vCenter SQL Server VM as the admin user.

3. Navigate to the c: drive and create a new folder called database.

4. Open Server Manager.

5. Click Manage and then select Add Roles and Features to start the Add roles and Features Wizard.

6. Click Next.

7. On the Select installation screen, select Role-based or feature-based installation.

8. Select target Server and click Next.

9. Click Next on Server Roles.

10. On the Select Features screen, check the box .Net Framework 3.5 Features and click Next.

11. On the Confirm installation selections screen, a warning will be displayed asking Do you need to specify an alternate source path? If the target computer does not have access to Windows Update, click the Specify an alternate source path link to specify the path to the \sources\sxs folder on the installation media and then click OK. After you have specified the alternate source, or if the target has access to Windows update, click the X next to the warning, and then click Install.

12. Click Close.

13. Open Server Manager click Tools and then select Windows Firewall with Advanced Security.

14. Choose Inbound Rules and click New Rule.

15. Choose Port and click Next.

16. Choose TCP and enter specific local port 1433. Click Next.

17. Choose Allow the Connection. Click Next, and the click Next again.

18. Name the rule SQL Server and click Finish.

19. Close the Windows Firewall with Advanced Security.

20. In the vCenter SQL Server VMware console, click the ninth button (CD with a wrench) to map the Microsoft SQL Server 2012 SP1 ISO. Choose Connect to ISO Image on Local Disk.

21. Navigate to the SQL Server 2012 SP1, select it, and click Open.

22. In the dialog box, click Run Setup.exe.

23. In the SQL Server Installation Center window, click **Installation** on the left**.**

24. Click New SQL Server stand-alone installation or add features to an existing installation.

25. On Setup Support Rules screen, click OK.

26. Choose Enter the Product Key. Enter a product key and click Next.

27. Read and accept the license terms and choose whether to check the second check box. Click Next.

28. On the Product Updates screen, Click Next.

29. Click Show details>> Address any warning except for the Windows Firewall Warning. Click Next on Setup Support Rules screen.

Note: The Windows Firewall issue was addressed in Step 14.

30. Choose SQL Server Feature Installation and click Next on the Setup Role screen.

31. Under Instance Features, Choose Only Database Engine Services.

32. Under Shared Features, choose Management Tools - Basic and Management Tools – Complete and click Next.

33. On the Installation Rules screen click Show details>> Address any warning and click Next.

34. Keep the Default instance selected. Click Next.

35. Click Next on the Disk Space Requirements screen.

36. For the SQL Server Agent Service and SQL Server Database Engine choose the first cell in the account Name column and then click <<Browse...>>.

37. Enter the local Machine Administrator name (for example, systemname\Administrator), click Check names, and click OK.

38. In the password field enter your password.

39. Change the startup type for SQL Server Agent to Automatic, and click Next.

40. Choose Mixed Mode (SQL Server and Windows Authentication. Enter and confirm the password for the SQL Server System Administrator (sa) account, click Add Current User, and click Next.

41. Choose whether to send error reports to Microsoft and click Next.

42. On the Installation Configuration Rules screen, click Show details>> Address any warning. Click Next.

43. Click Install.

44. After the installation is complete, click Close.

45. Close the SQL Server Installation Center.

46. Install all available Microsoft updates by going to Control Panel and select Windows Updates.

47. Open SQL Server Management Studio.

48. Under Server Name, choose the local machine name. Under Authentication, choose SQL Server Authentication. Enter sa in the Login field and enter the sa password. Click Connect.

49. Click New Query on the toolbar.

50. Run the following script, substituting the vpxuser password for <Password>.

```
use [master]

go

CREATE DATABASE [VCDB] ON PRIMARY

(NAME = 'vcdb', FILENAME = 'C:\database\VCDB.mdf', SIZE = 4000KB, FILEGROWTH =
10% )

LOG ON

(NAME = 'vcdb_log', FILENAME = 'C:\database\VCDB.ldf', SIZE = 1000KB, FILEGROWTH
= 10%)

COLLATE SQL_Latin1_General_CP1_CI_AS

go

ALTER DATABASE [VCDB] SET RECOVERY SIMPLE

use VCDB

go

sp_addlogin @loginame=[vpxuser], @passwd='<password>', @defdb='VCDB',
@deflanguage='us_english'

go

ALTER LOGIN [vpxuser] WITH CHECK_POLICY = OFF

go

CREATE USER [vpxuser] for LOGIN [vpxuser]

go

use MSDB

go

CREATE USER [vpxuser] for LOGIN [vpxuser]

go

sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'

go

use VCDB

go

sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'

go
```

Note: This example illustrates the script.

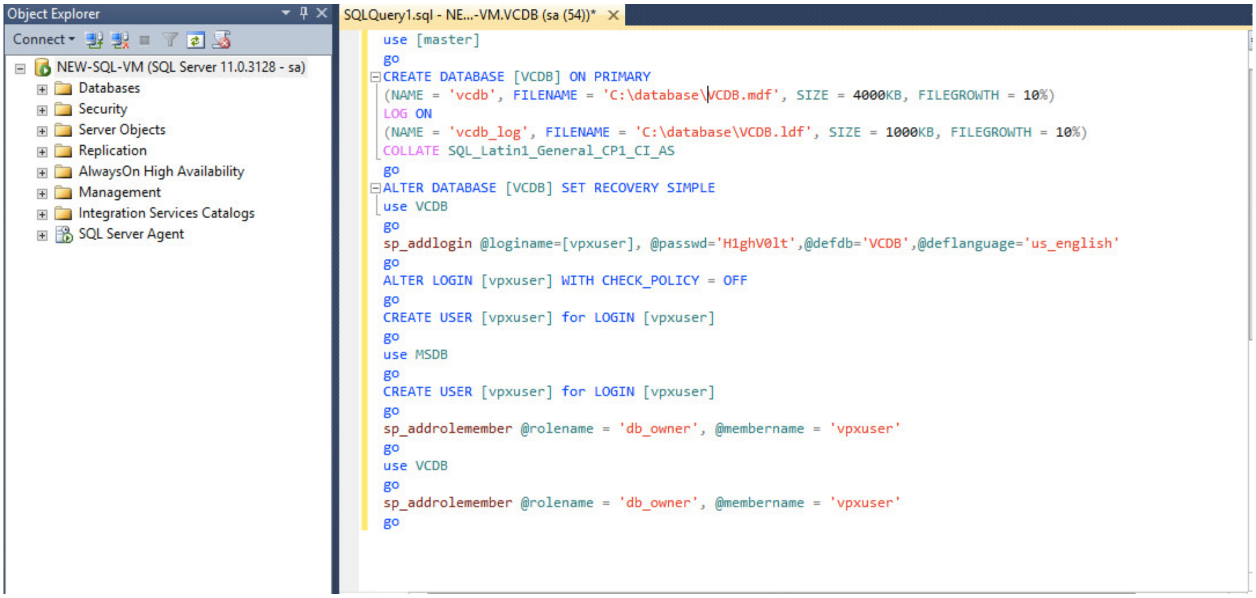

Object Explorer
Connect
NEW-SQL-VM (SQL Server 11.0.3128 - sa)
Databases
Security
Server Objects
Replication
AlwaysOn High Availability
Management
Integration Services Catalogs
SQL Server Agent
SQLQuery1.sql - NE...-VM.VCDB (sa (54))*

```
use [master]
go
CREATE DATABASE [VCDB] ON PRIMARY
(NAME = 'vcdb', FILENAME = 'C:\database\VCDB.mdf', SIZE = 4000KB, FILEGROWTH = 10%)
LOG ON
(NAME = 'vcdb_log', FILENAME = 'C:\database\VCDB.ldf', SIZE = 1000KB, FILEGROWTH = 10%)
COLLATE SQL_Latin1_General_CP1_CI_AS
go
ALTER DATABASE [VCDB] SET RECOVERY SIMPLE
use VCDB
go
sp_addlogin @loginame=[vpxuser], @passwd='H1ghV0lt',@defdb='VCDB',@deflanguage='us_english'
go
ALTER LOGIN [vpxuser] WITH CHECK_POLICY = OFF
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
use MSDB
go
CREATE USER [vpxuser] for LOGIN [vpxuser]
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
use VCDB
go
sp_addrolemember @rolename = 'db_owner', @membername = 'vpxuser'
go
```

51. Click Execute and verify that the query executes successfully.

52. Close Microsoft SQL Server Management Studio.

53. Disconnect the Microsoft SQL Server 2012 ISO from the SQL Server VM.

Build and Set Up VMware vCenter Virtual Machine

Build VMware vCenter Virtual Machine

To build the VMware vCenter VM, follow these steps

1. Using the instructions for building a SQL Server VM provided in the section Build Microsoft SQL Server VM, build a VMware vCenter VM with the following configuration in the `<<var_ib-mgmt_vlan_id>>` VLAN:

— 12GB RAM

— Two CPUs

— One virtual network interface

2. Start the VM, install VMware Tools, and assign an IP address and host name to it in the Active Directory domain.

Set UP VMware vCenter Virtual Machine

To setup the newly built VMware vCenter VM, follow these steps:

1. Log into the vCenter VM as the admin user and open Server Manager.

2. Click Manage and then select Add Roles and Features to start the Add roles and Features Wizard.

3. Click Next.

4. On the Select installation screen, select Role-based or feature-based installation.

5. Select target Server and click Next.

6. Click Next on Server Roles.

7. On the Select Features screen, check the box .Net Framework 3.5 Features and click Next.

8. On the Confirm installation selections screen, a warning will be displayed asking Do you need to specify an alternate source path? If the target computer does not have access to Windows Update, click the Specify an alternate source path link to specify the path to the \sources\sxs folder on the installation media and then click OK. After you have specified the alternate source, or if the target has access to Windows update, Click the X next to the warning, and then click Install.

9. Click Close.

10. Click the following link Windows SQL Server 2012 SP1 Feature Pack.

11. Click Download.

12. Select the file name ENU\x64\sqlncli.msi and click Next.

13. Save it to a destination folder and run it.

14. On the Microsoft SQL Server 2012 Native Client Setup click Next.

15. Accept the license terms and click Next.

16. Click Next.

17. Click Install.

18. Click Finish.

19. Create the vCenter database data source name (DSN). Open Data Source (ODBC) by selecting Server Manager > Tools > ODBC Data Sources (64-bit).

20. Click the System DSN tab.

21. Click Add.

22. Choose SQL Server Native Client 11.0 and click Finish.

23. Name the data source VCDB. In the server , enter the IP address of the vCenter SQL server and

24. Click Next.

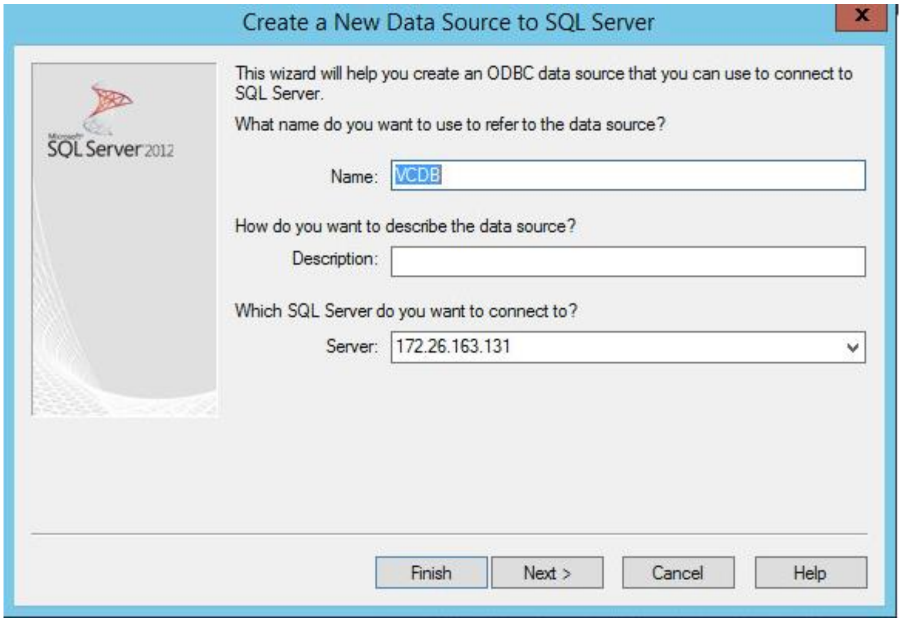

This wizard will help you create an ODBC data source that you can use to connect to SQL Server.

What name do you want to use to refer to the data source?

Name: VCDB

How do you want to describe the data source?

Description:

Which SQL Server do you want to connect to?

Server: 172.26.163.131

25. Choose With SQL Server authentication using a login ID and password entered by user. Enter vpxuser as the login ID and vpxuser password. Click Next.

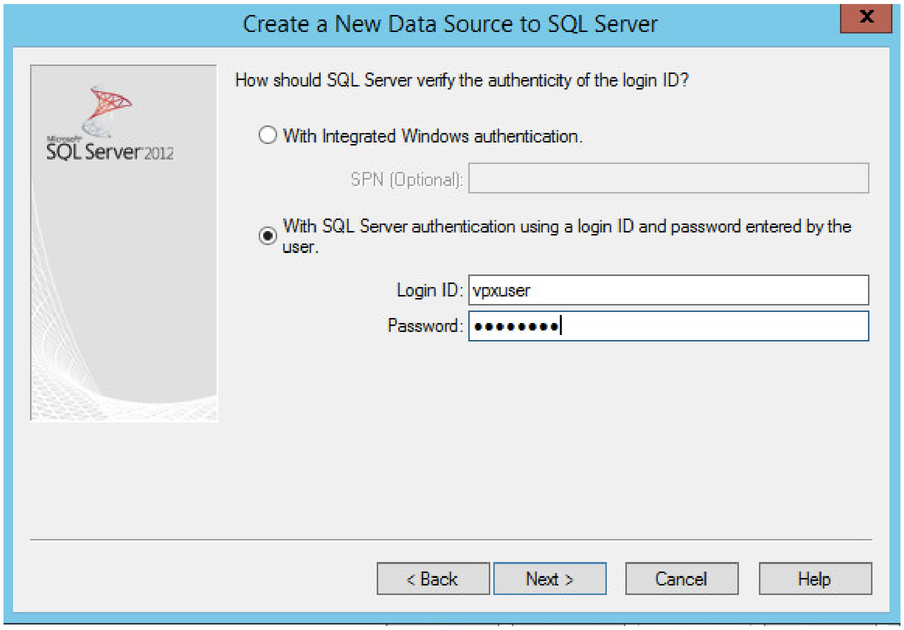

How should SQL Server verify the authenticity of the login ID?

○ With Integrated Windows authentication.

SPN (Optional):

● With SQL Server authentication using a login ID and password entered by the user.

Login ID: vpxuser

Password: ••••••••

26. Choose Change the Default Database to and choose VCDB from the list. Click Next.

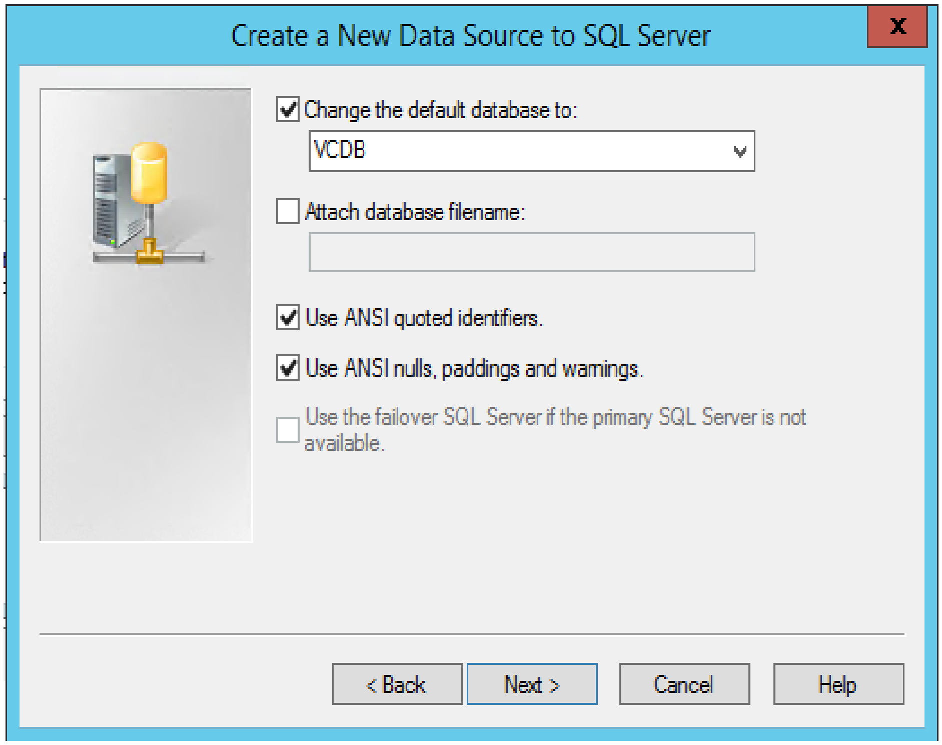

Create a New Data Source to SQL Server

☑ Change the default database to:

VCDB

□ Attach database filename:

☑ Use ANSI quoted identifiers.

☑ Use ANSI nulls, paddings and warnings.

□ Use the failover SQL Server if the primary SQL Server is not available.

27. Click Finish.

28. Click Test Data Source. Verify that the test completes successfully.

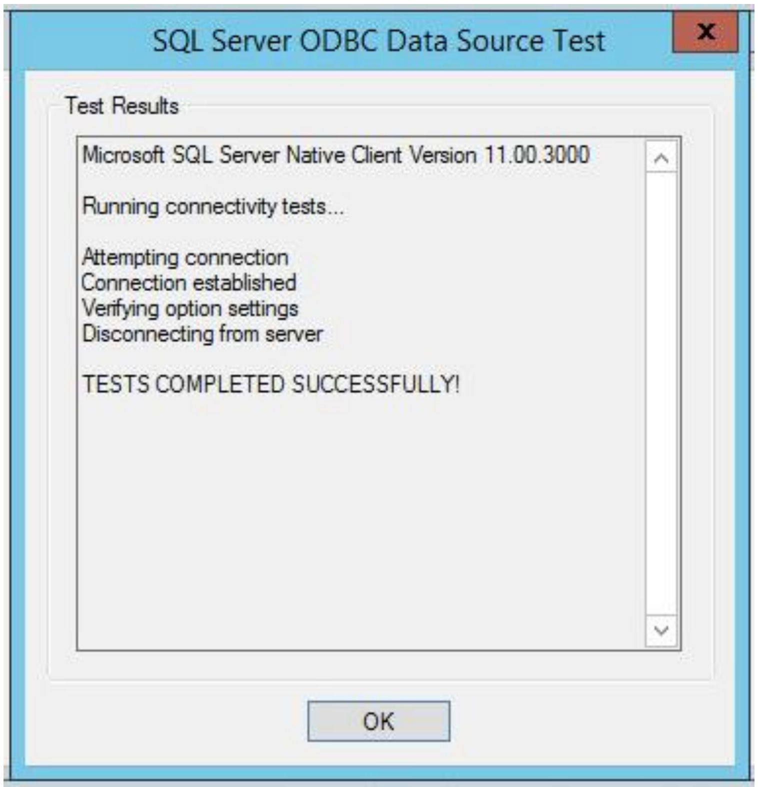

SQL Server ODBC Data Source Test

Test Results

```
Microsoft SQL Server Native Client Version 11.00.3000

Running connectivity tests...

Attempting connection
Connection established
Verifying option settings
Disconnecting from server

TESTS COMPLETED SUCCESSFULLY!
```

29. Click OK and then click OK again.

30. Click OK to close the ODBC Data Source Administrator (64-bit) window.

31. Install all available Microsoft Windows updates by right-clicking Start > Control Panel > Windows Update.

Install VMware vCenter Server

vCenter Server Virtual Machine

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

1. In the vCenter Server VMware console, click the ninth button (CD with a wrench) to map the VMware vCenter ISO and choose Connect to ISO Image on Local Disk.

2. Navigate to the VMware vCenter 5.5 (VIMSetup) ISO, select it, and click Open.

3. In the dialog box, click Run **autorun.**exe

4. In the VMware vCenter installer window, make sure that VMware vCenter Simple Install is selected and click Install**.**

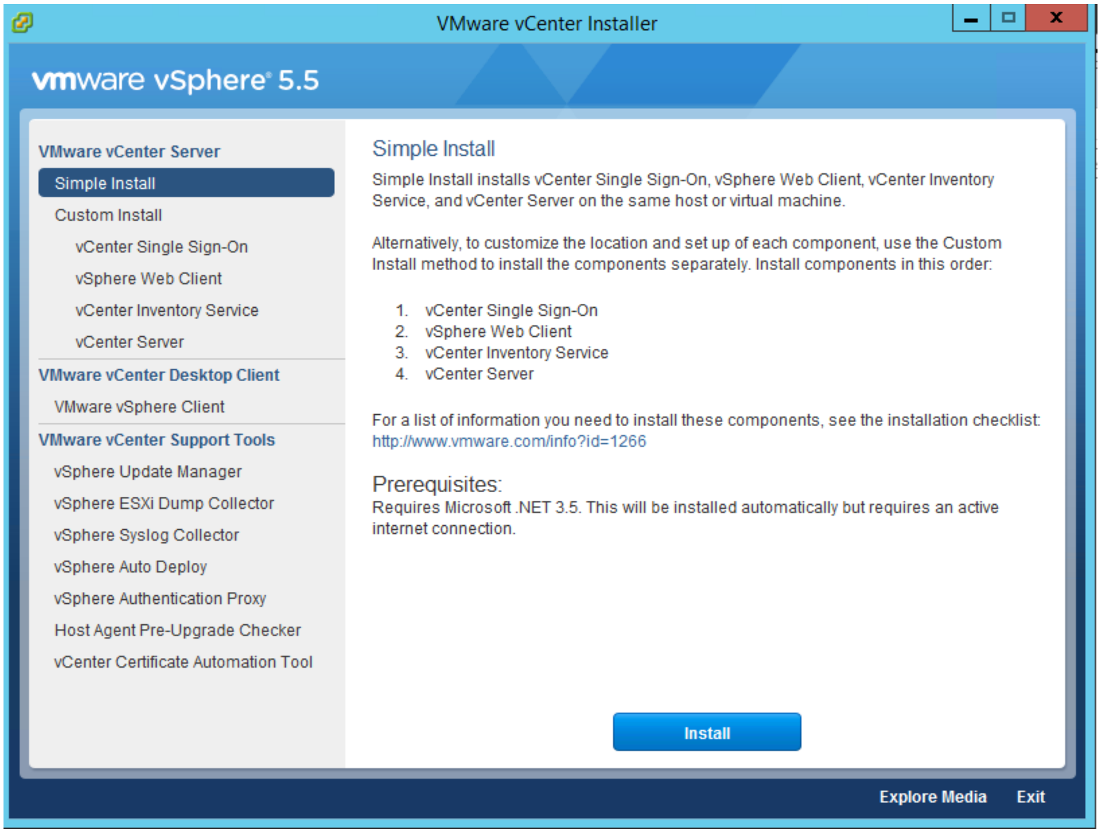

Simple Install

Simple Install installs vCenter Single Sign-On, vSphere Web Client, vCenter Inventory Service, and vCenter Server on the same host or virtual machine.

Alternatively, to customize the location and set up of each component, use the Custom Install method to install the components separately. Install components in this order:

1. vCenter Single Sign-On
2. vSphere Web Client
3. vCenter Inventory Service
4. vCenter Server

For a list of information you need to install these components, see the installation checklist: http://www.vmware.com/info?id=1266

Prerequisites:

Requires Microsoft .NET 3.5. This will be installed automatically but requires an active internet connection.

5. Click Next on the Welcome to the vCenter Single Sign-On Setup screen.
6. Accept the terms of the license agreement and click Next.
7. Click Next on Simple Install Prerequisites check screen.
8. Enter and confirm <<var_password>> for administrator@vsphere.local . Click Next.
9. Click Next on Simple Install Configure Site window.
10. Click Next on Simple Install Port Setting window.
11. Click Next on Change destination folder.
12. Review the installation option and click Install.
13. Enter the vCenter 5.5 license key and click Next.
14. Choose Use an Existing Supported Database. Choose VCDB from the Data Source Name list and click Next.

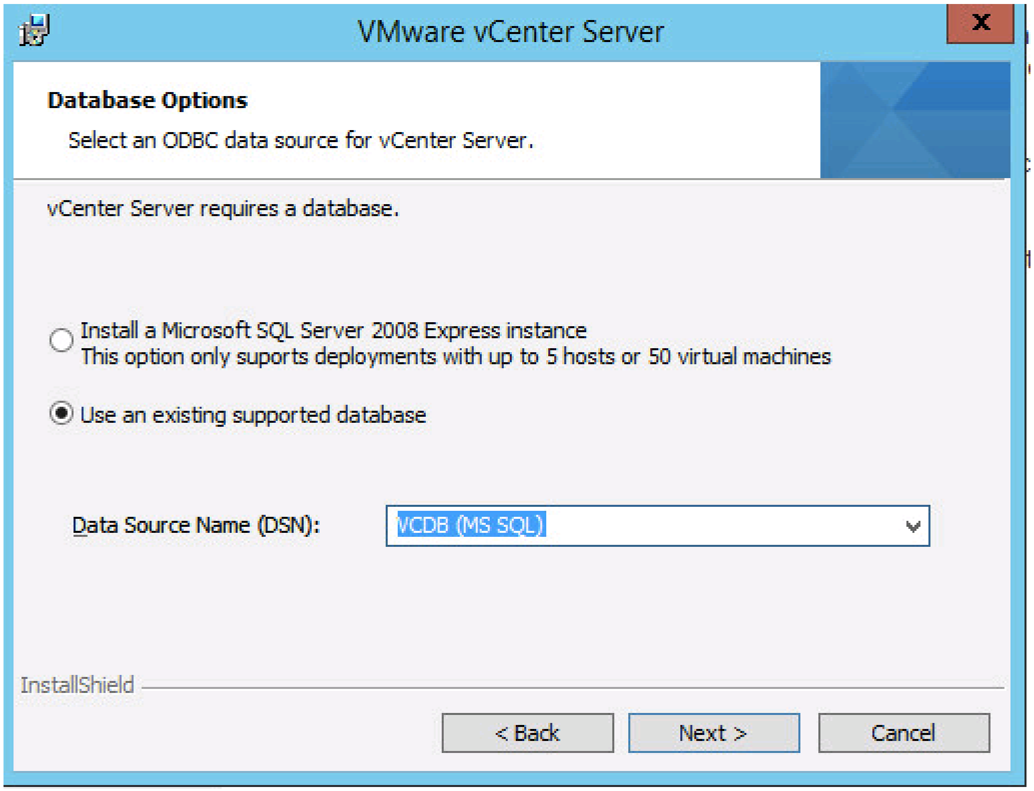

VMware vCenter Server

Database Options

Select an ODBC data source for vCenter Server.

vCenter Server requires a database.

○ Install a Microsoft SQL Server 2008 Express instance
This option only suports deployments with up to 5 hosts or 50 virtual machines

● Use an existing supported database

Data Source Name (DSN): VCDB (MS SQL)

15. Enter the vpxuser password and click Next.
16. Click Next in the vCenter Server Service Window.
17. Click Next on Configure Ports screen

18. Choose the vCenter Server Configuration that best describes your setup. Click Next.

19. Choose the inventory size that that best describes your setup. Click Next.

20. Click Install.

21. Click Yes to accept and continue with SSL SHA1 SSO lookup Service Leaf certificate.

22. Click Install certificates.

23. Click Finish.

24. Click OK.

ESXI Dump Collector Setup

1. In the VMware vCenter Installer window, under vCenter Support Tools, select vSphere ESXi Dump Collector.

2. Click Install.

3. Select the appropriate language and click OK.

4. In the vSphere ESXi Dump Collector Installation Wizard, click Next.

5. Accept the terms in the License Agreement and click Next.

6. Click Next to accept the default Destination Folders.

7. Click Next to accept a Standalone installation

8. Click Next to accept the default ESXi Dump Collector Server Port (6500).

9. Select the VMware vCenter Server IP address from the drop-down menu. Click Next.

10. Click Install to complete the installation.

11. Click Finish.

12. Click Exit in the VMware vCenter Installer window.

13. Disconnect the VMware vCenter ISO from the vCenter VM.

Note: A restart might be required.

14. Back on the Management Workstation, search for Command Prompt and do a right-click the Command Prompt entry and then click Run as administrator option to open elevated Command Prompt.

15. For 64 bit OS go to C:\Program Files (x86)\VMware\Vmware vSphere CLI\bin directory and for 32 bit OS is C:\ProgramFiles\VMware\Vmware vSphere CLI\bin.

16. Set each ESXi Host to coredump to ESXi Dump Collector by running the following commands:

```
esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump
network set --interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip>> --
server-port 6500

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system coredump
network set --interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip>> --
server-port 6500

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump
network set --enable true

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system coredump
network set --enable true

esxcli -s <<var_vm_host_infra_01_ip>> -u root -p <<var_password>> system coredump
network check

esxcli -s <<var_vm_host_infra_02_ip>> -u root -p <<var_password>> system coredump
network check
```

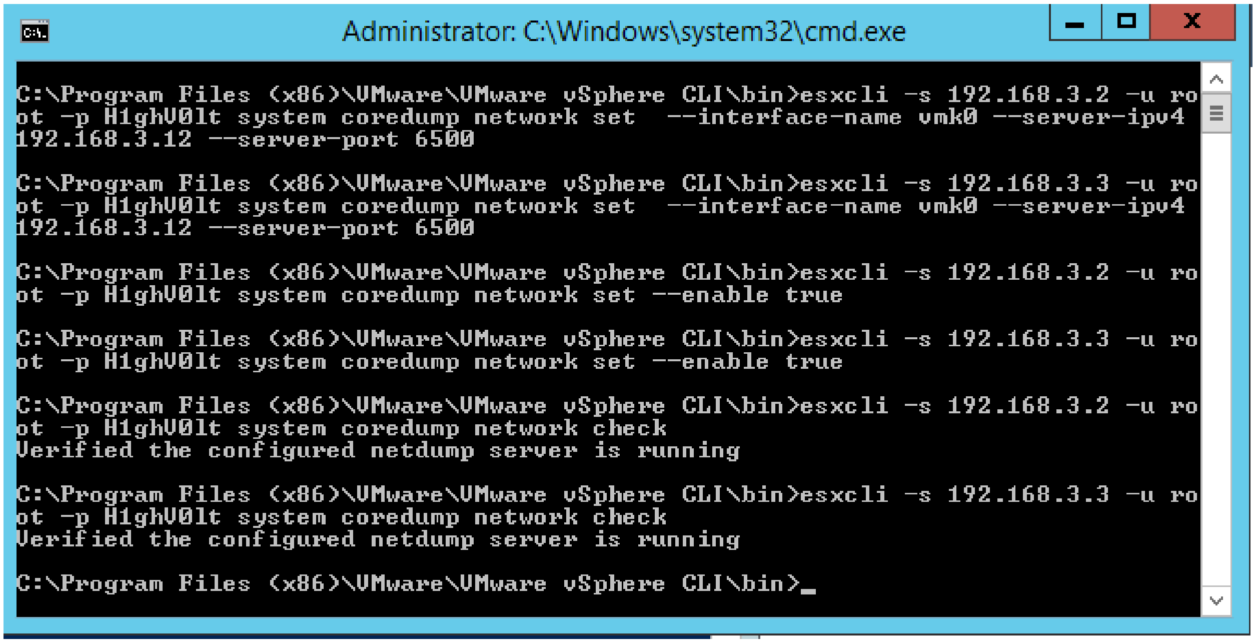

```
C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.2 -u ro
ot -p H1ghV0lt system coredump network set  --interface-name vmk0 --server-ipv4
192.168.3.12 --server-port 6500

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.3 -u ro
ot -p H1ghV0lt system coredump network set  --interface-name vmk0 --server-ipv4
192.168.3.12 --server-port 6500

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.2 -u ro
ot -p H1ghV0lt system coredump network set --enable true

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.3 -u ro
ot -p H1ghV0lt system coredump network set --enable true

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.2 -u ro
ot -p H1ghV0lt system coredump network check
Verified the configured netdump server is running

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>esxcli -s 192.168.3.3 -u ro
ot -p H1ghV0lt system coredump network check
Verified the configured netdump server is running

C:\Program Files (x86)\VMware\VMware vSphere CLI\bin>_
```

Log in to vSphere Web Client

1. Using a web browser, navigate to https:// <<var_vcenter_server_ip>>:9443/vsphere-client/#

2. Click Download the Client Integration Plug-in.

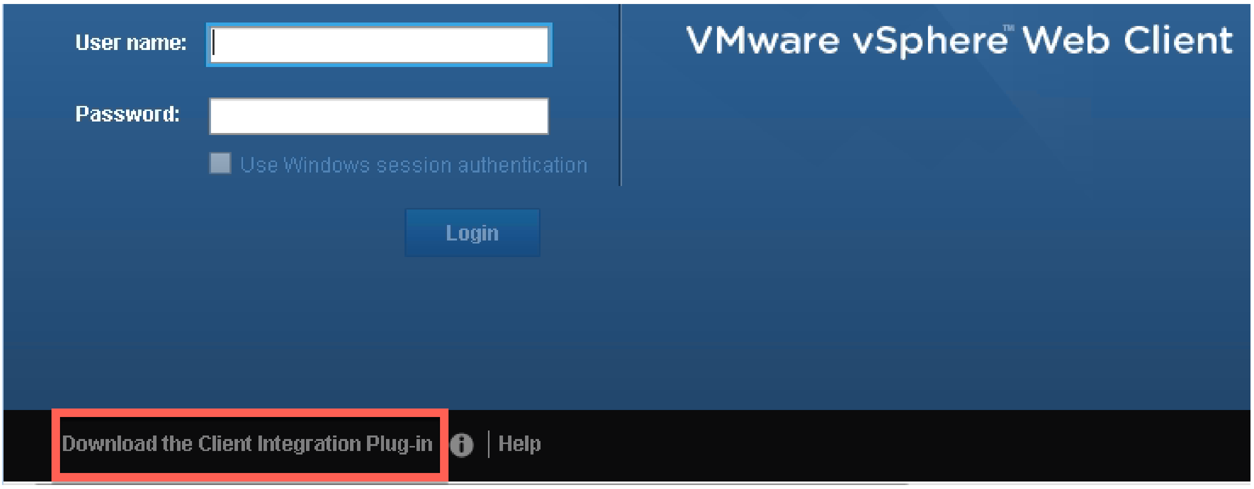
User name:
Password:
Use Windows session authentication
Login
VMware vSphere Web Client
Download the Client Integration Plug-in
Help

3. Click Run.

4. Closed the indicate browser and click Retry.

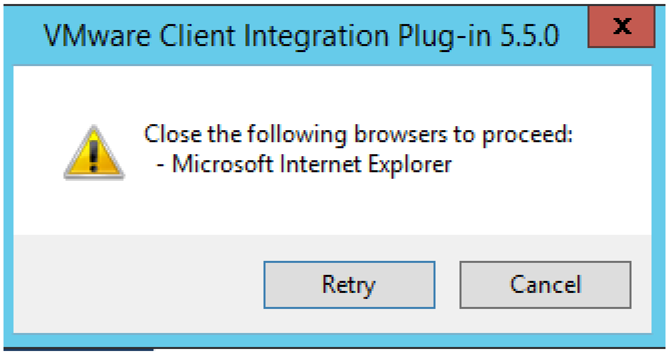
VMware Client Integration Plug-in 5.5.0

Close the following browsers to proceed:
- Microsoft Internet Explorer

5. Click Next

6. Accept the license terms and click Next.

7. Click Next on Destination Folder window.

8. Click Install.

9. Click Finish.

10. Using a web browser, navigate to https:// <<var_vcenter_server_ip>>:9443/vsphere-client/#

11. In the user name type in administrator@vsphere.local and the administrator Password.

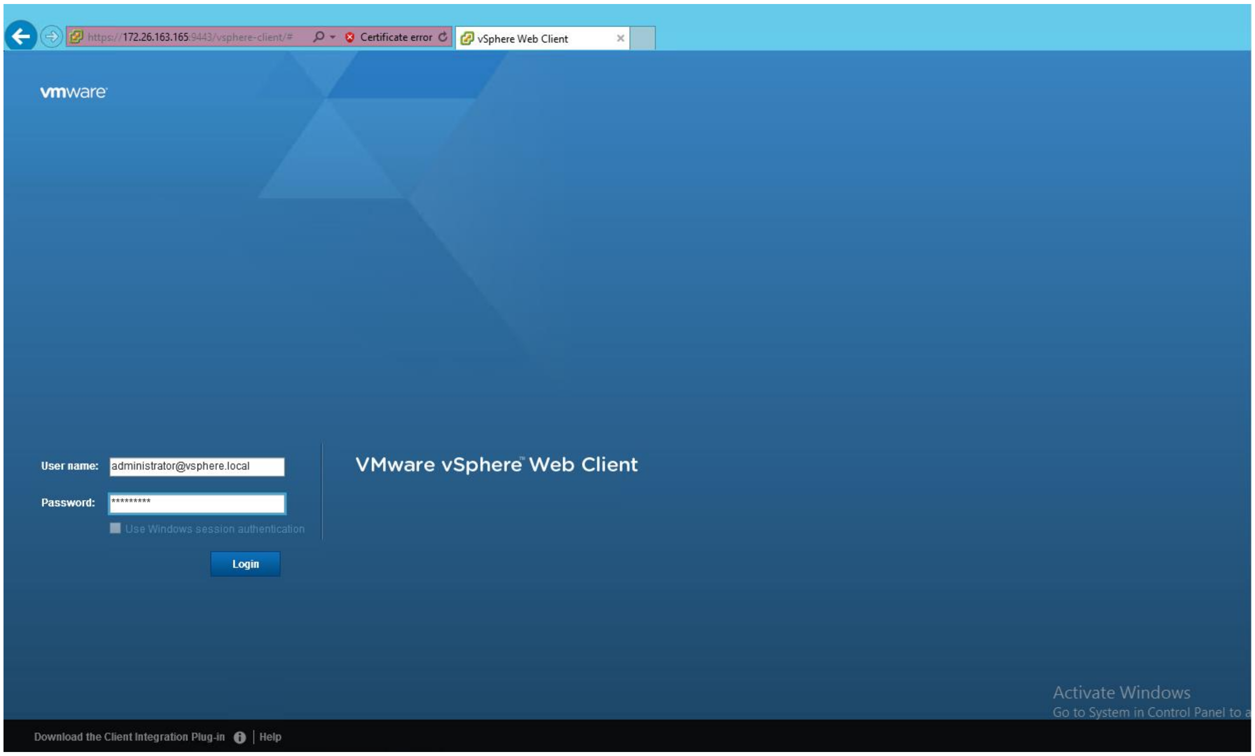

User name: administrator@vsphere.local

Password: *********

https://172.26.163.165:9443/vsphere-client/#
Certificate error
vSphere Web Client
vmware
VMware vSphere Web Client
Use Windows session authentication
Login
Activate Windows
Go to System in Control Panel to a
Download the Client Integration Plug-in
Help

12. Click Login.

Adding the AD Account to Administrator Group

1. Log in to the vSphere Web Client as Administrator@vsphere.local.

2. From the home Location, navigate to >>Administration>>Single Sign-ON>>User and Groups>>.

3. In the right pane, select Groups.

4. Click Administrators on Group Name.

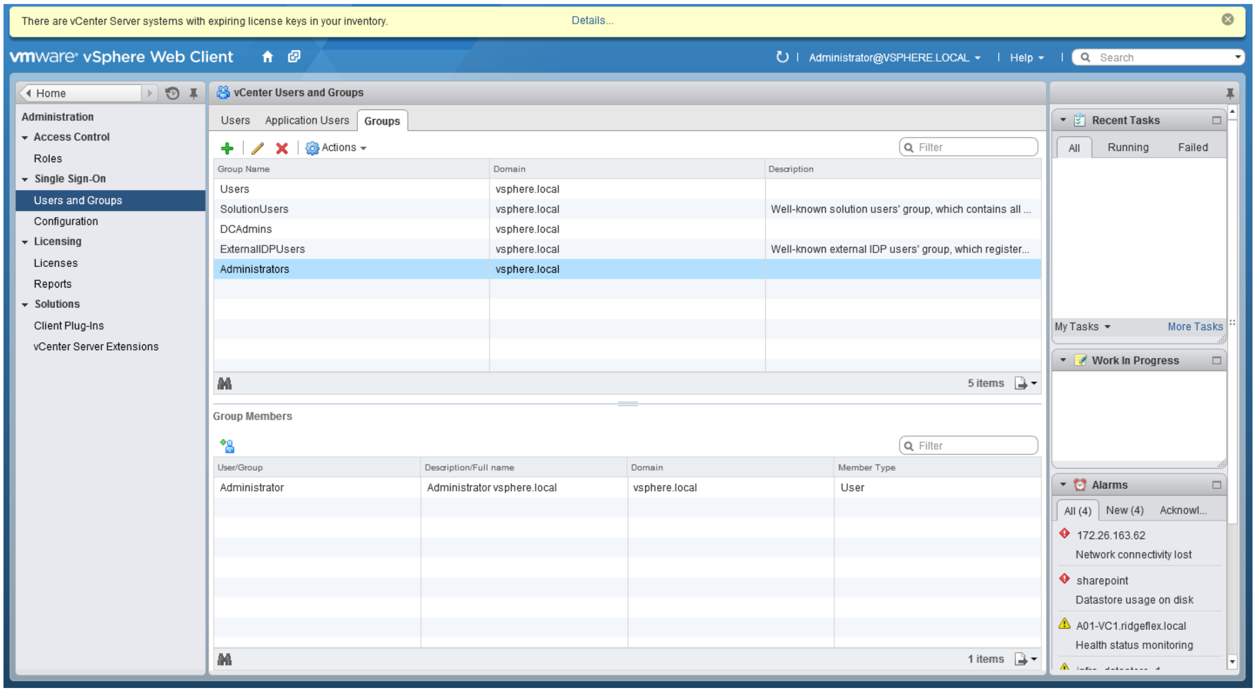

There are vCenter Server systems with expiring license keys in your inventory. Details...

Group Name	Domain	Description
Users	vsphere.local	
SolutionUsers	vsphere.local	Well-known solution users' group, which contains all ...
DCAdmins	vsphere.local	
ExternalIDPUsers	vsphere.local	Well-known external IDP users' group, which register...
Administrators	vsphere.local	

5 items

User/Group	Description/Full name	Domain	Member Type
Administrator	Administrator vsphere.local	vsphere.local	User

1 items

172.26.163.62
Network connectivity lost

sharepoint
Datastore usage on disk

A01-VC1.ridgeflex.local
Health status monitoring

5. Click the **+** on Group Members.

6. Change the Domain to <<domain>>.

7. Highlight the Administrator and click Add.

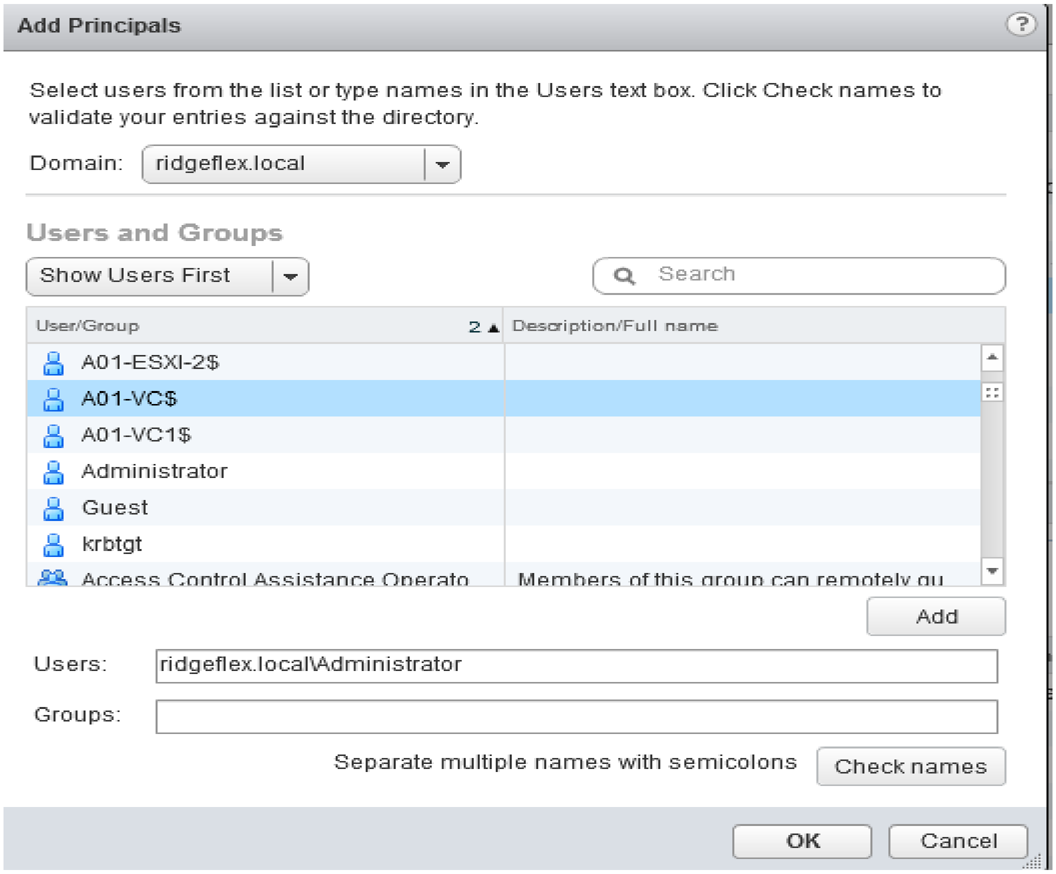

Select users from the list or type names in the Users text box. Click Check names to validate your entries against the directory.

Domain: ridgeflex.local

Users and Groups

Show Users First

User/Group	Description/Full name
A01-ESXI-2$	
A01-VC$	
A01-VC1$	
Administrator	
Guest	
krbtgt	
Access Control Assistance Operato	Members of this group can remotely qu

Users: ridgeflex.local\Administrator

Groups:

Separate multiple names with semicolons

8. Click OK.

9. Navigate to >>Home>>vCenterServers and click your <<vcenter_server>>.

10. In the right pane, click the Manage tab.

11. Click Permissions tab.

12. Click the + sign.

13. Click Add.

14. Select your domain.

15. Highlight Administrator and double-clickit.

16. Click OK.

17. Change the Assigned Role to Administrator and click OK.

18. Log out from vSphere Web Client.

19. Login to the vSphere Web Client as Administrator@<<domain>>.

Set Up vCenter Center with a Datacenter, Cluster, DRS and HA

1. In the vSphere Web Client, navigate to the >>vCenter>>vCenter Servers>>`vCenter_name`.
2. Right-click the vCenter server and Select Actions > New Datacenter.
3. Rename the datacenter and click OK.
4. Browse to a datacenter in the vSphere Web Client navigator.
5. Right-click the datacenter and select New Cluster.
6. Select DRS and vSphere HA cluster features.
7. Select the DRS Turn ON check box.
8. Select the vSphere HA Turn ON check box.
9. Click OK.

Add Host to vCenter

1. In the vSphere Web Client, navigate to a datacenter, cluster, or folder within a datacenter.
2. Right-click the datacenter, cluster, or folder and select Add Host.
3. Type the IP address or the name of the host and click Next.
4. Type root credentials and click Next.
5. Click **Yes** to accept the certificate.
6. Review the host summary and click Next.
7. Assign a license key to the host Click Next.
8. (Optional) Select Enable Lockdown Mode to disable remote access for the administrator account after vCenter Server takes control of this host and click Next.
9. (Optional) If you add the host to a datacenter or a folder, select a location for the virtual machines that reside on the host and click Next.
10. Review the summary and click Finish.

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is enabled by default on the vCenter Appliance.

1. On the Management Workstation, open the VMware vSphere CLI command prompt.

2. Set each iSCSI-booted ESXi Host to coredump to the ESXi Dump Collector by running the following commands:

```
esxcli –s <<var_vm_host_infra_01_ip>> -u root –p <<var_password>> system coredump
network set --interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --
server-port 6500
```

```
esxcli –s <<var_vm_host_infra_02_ip>> -u root –p <<var_password>> system coredump
network set --interface-name vmk0 --server-ipv4 <<var_vcenter_server_ip> --
server-port 6500
```

```
esxcli –s <<var_vm_host_infra_01_ip>> -u root –p <<var_password>> system coredump
network set --enable true
```

```
esxcli –s <<var_vm_host_infra_02_ip>> -u root –p <<var_password>> system coredump
network set --enable true
```

```
esxcli –s <<var_vm_host_infra_01_ip>> -u root –p <<var_password>> system coredump
network check
```

```
esxcli –s <<var_vm_host_infra_02_ip>> -u root –p <<var_password>> system coredump
network check
```

Cisco ACI – Virtual Machine Manager

This section provides a detailed procedure for configuring the Cisco APIC to communicate and control VMware Distributed Switch (VDS). In this section, VMware vCenter attachment to ACI will be covered.

Defining VMware Distributed Switch policies

In the Cisco UCS environment, VDS uplink configuration requires following three parameters to be explicitly set:

- No Port-Channel for uplink ports
- CDP used as the discovery protocol
- LLDP disabled as the discovery protocol

This configuration requires modifying the Access Entity Profiles (AEP) defined for Cisco UCS Fabric Interconnects.

1. From the top menu, select FABRIC.
2. Select ACCESS POLICIES from the sub menu.
3. From the left menu bar, expand Global Policies and Attachable Access Entity Profiles.
4. Select aep-<UCS_FI> where UCS_FI is the name used for UCS Fabric Interconnect profile in the previous sections.
5. On the right hand side, click Action and select Config vSwitch Policies.

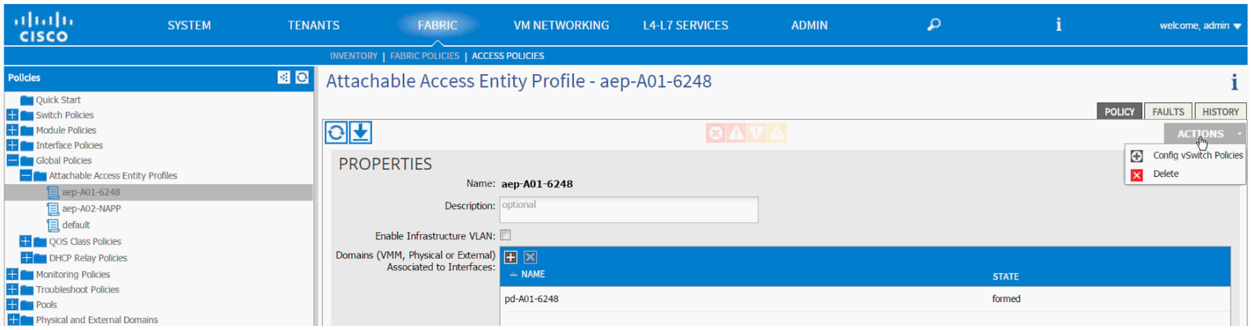

Attachable Access Entity Profile - aep-A01-6248

PROPERTIES

Name: aep-A01-6248

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) Associated to Interfaces:

NAME	STATE
pd-A01-6248	formed

6. In the ‘CONFIG VSWITCH POLICIES dialog box, select CDP_Enable as the CDP Policy.
7. Select LACP_MAC_Pinning as the LACP Policy.
8. Select LLDP_Disabled as the LLDP Policy.

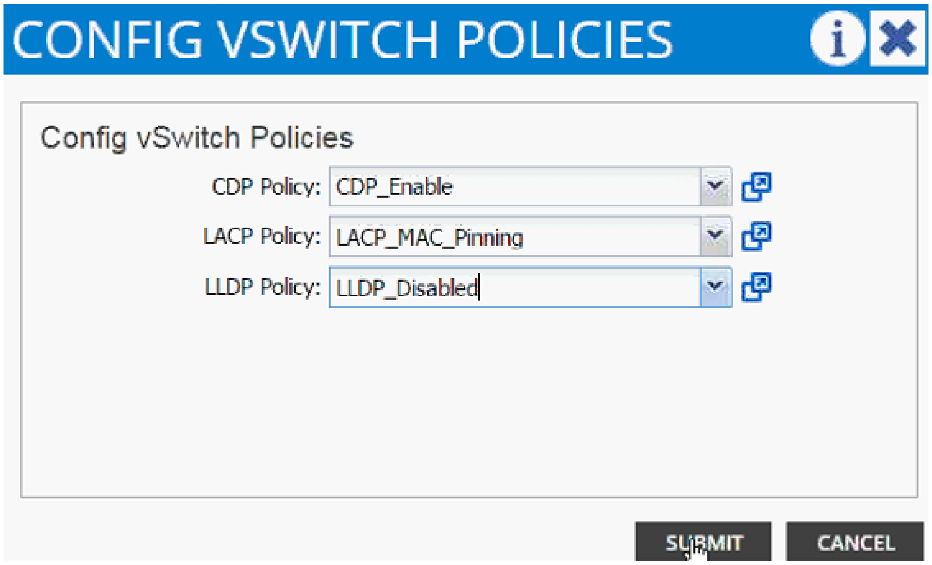

CONFIG VSWITCH POLICIES

Config vSwitch Policies

CDP Policy: CDP_Enable

LACP Policy: LACP_MAC_Pinning

LLDP Policy: LLDP_Disabled

9. Click SUBMIT.

10. Verify the resulting AEP configuration.

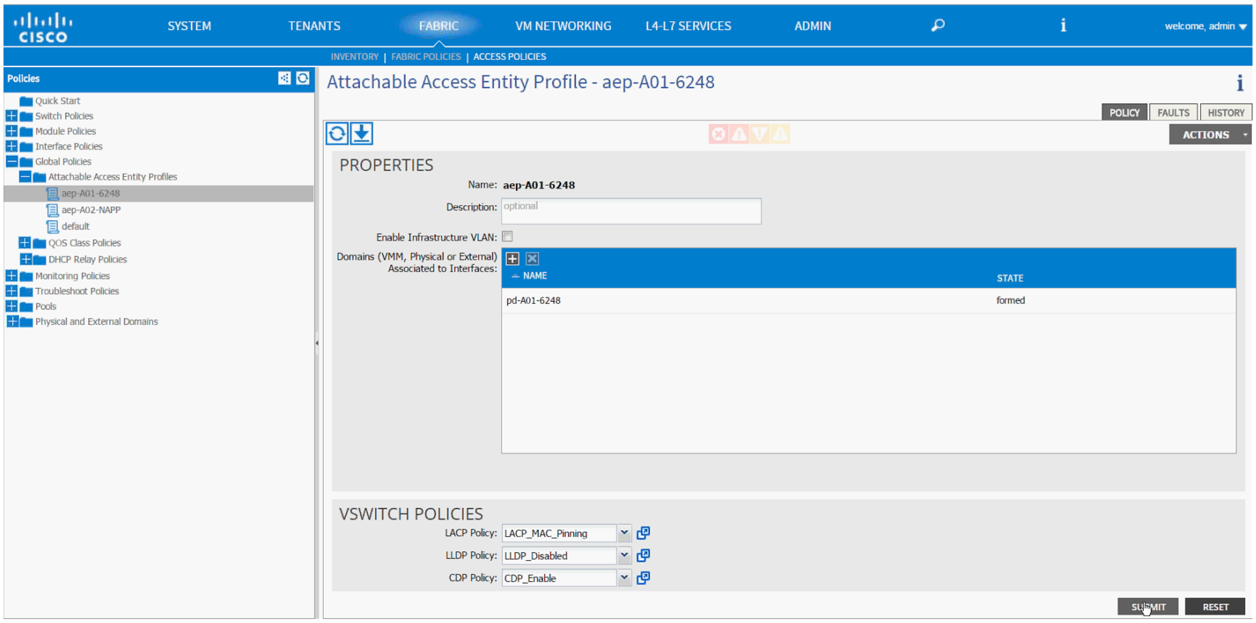

Attachable Access Entity Profile - aep-A01-6248

PROPERTIES

Name: aep-A01-6248

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) Associated to Interfaces:

NAME	STATE
pd-A01-6248	formed

VSWITCH POLICIES

LACP Policy: LACP_MAC_Pinning

LLDP Policy: LLDP_Disabled

CDP Policy: CDP_Enable

11. Click SUBMIT.

12. From the top menu, select VM NETWORKING.

13. Select POLICIES from the sub-menu.

14. From the left menu bar, click VM Provider VMware.

15. On the right side (under PROPERTIES) click + to add vCenter Domains.

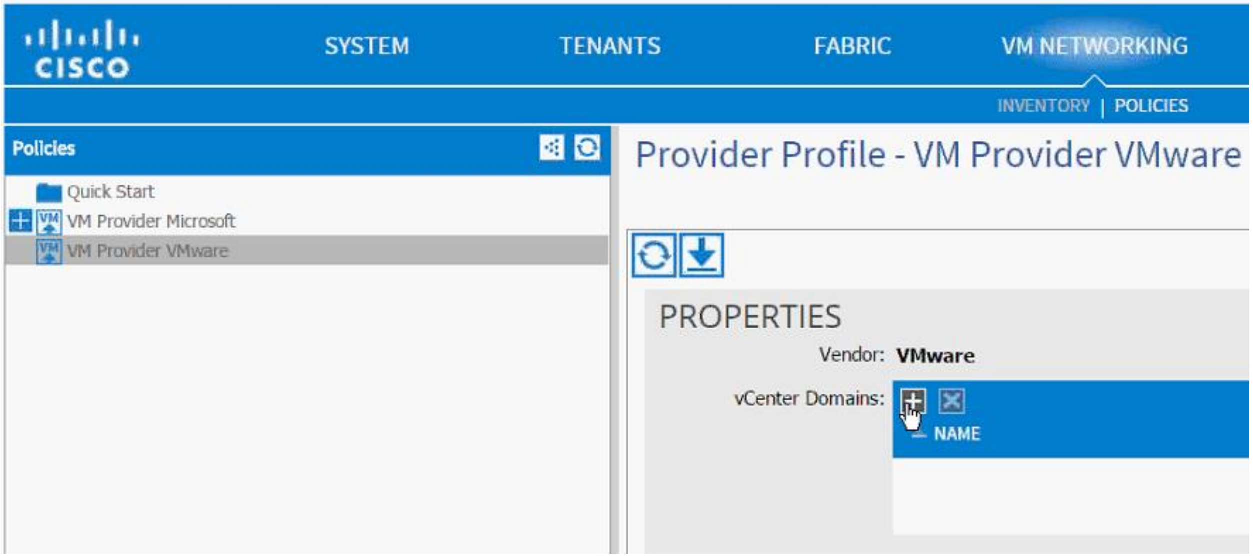

Provider Profile - VM Provider VMware

PROPERTIES

Vendor: VMware

vCenter Domains:

NAME

16. In the CREATE VCENTER DOMAIN dialog box, provide a Name to identify the vCenter.

17. Make sure VMWare vSphere Distributed Switch is selected as Virtual Switch.

18. From the drop-down menu for Associated Attachable Entity, select the AEP previously defined for UCS Fabric Interconnect.

19. From the drop-down menu next to VLAN Pool, select Create VLAN Pool.

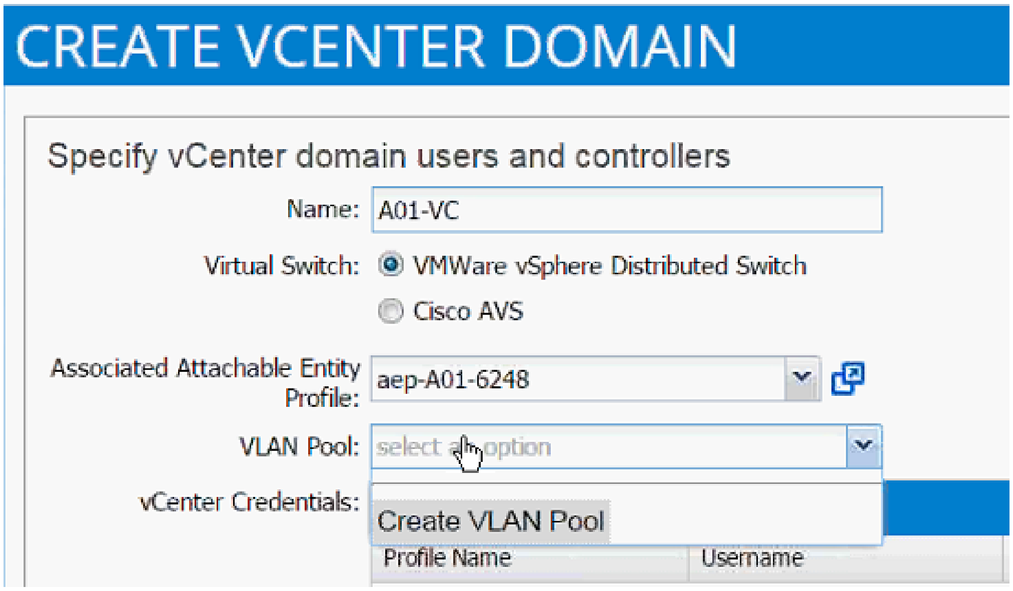

CREATE VCENTER DOMAIN

Specify vCenter domain users and controllers

Name: A01-VC

Virtual Switch: ● VMWare vSphere Distributed Switch

○ Cisco AVS

Associated Attachable Entity Profile: aep-A01-6248

VLAN Pool: select an option

vCenter Credentials:

Profile Name	Username

20. In the CREATE VLAN POOL dialog box, provide a name (e.g. vp-<NAME_of_vCenter> of the VLAN pool to be used for dynamically allocating VLANs to the EPGs.

21. Select Dynamic Allocation as the Allocation Mode.

22. Click + next to the Encap Blocks.

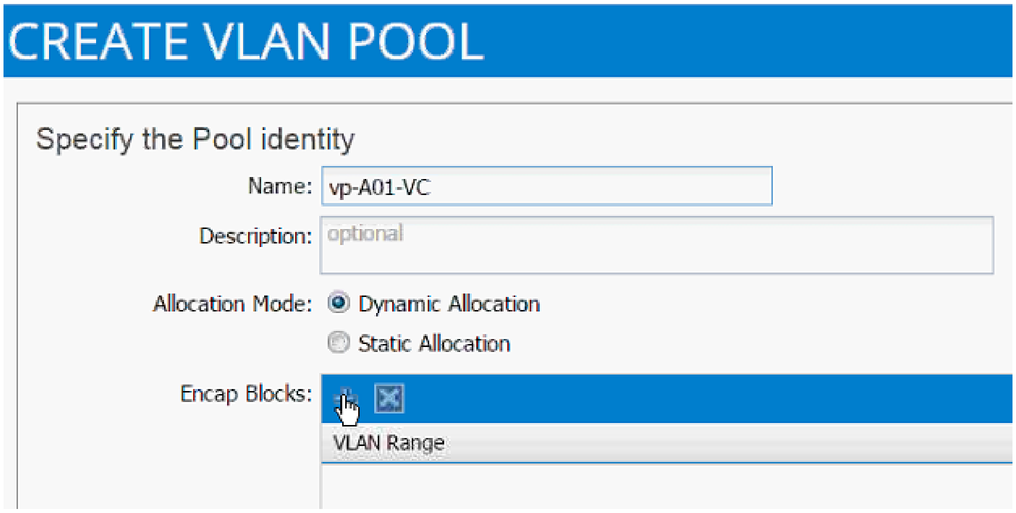

CREATE VLAN POOL

Specify the Pool identity

Name: vp-A01-VC

Description: optional

Allocation Mode: ● Dynamic Allocation
○ Static Allocation

Encap Blocks:

23. In the CREATE RANGES dialogue box add the VLAN range 1101 to 1200. (This range can be different depending on customer requirements).

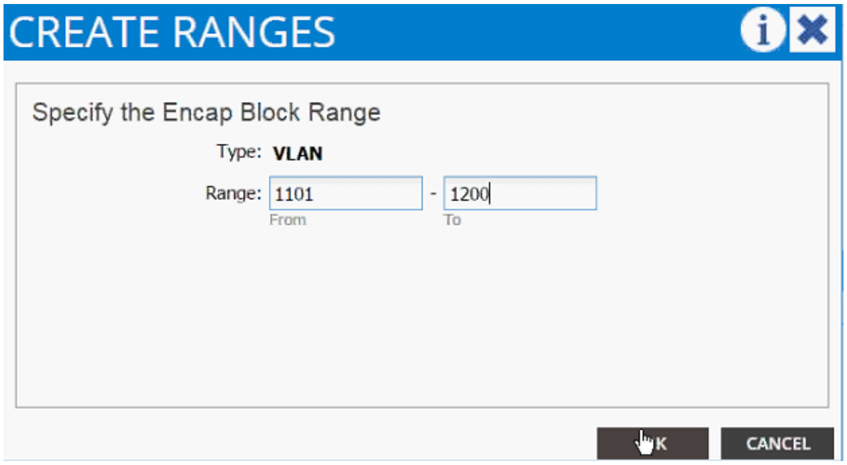

CREATE RANGES

Specify the Encap Block Range

Type: VLAN

Range: 1101 - 1200
From To

24. Click OK.

25. Click SUBMIT.

26. Click the + sign next to vCenter Credentials.

27. In the CREATE VCENTER CREDENTIALS dialog box, add vCenter Name.

28. Add admin username for vCenter in the following format: user@DOMAIN.

29. Add the password for the admin user and confirm the password.

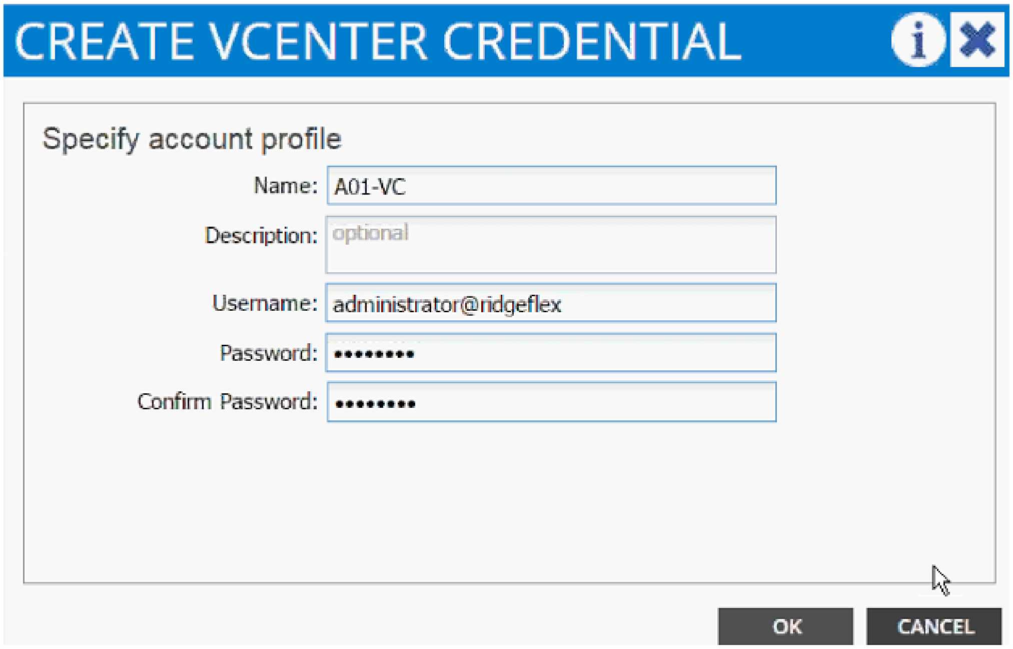

CREATE VCENTER CREDENTIAL

Specify account profile

Name: A01-VC

Description: optional

Username: administrator@ridgeflex

Password: ••••••••

Confirm Password: ••••••••

30. Click OK.

31. Click + next to vCenter/vShield.

32. In the CREATE VCENTER/VSHIELD CONTROLLER dialog box, select vCenter as the Type.

33. Add Name of the vCenter Controller.

34. Add DNS name or IP address in Hostname (or IP Address).

35. Select the DVS Version from the drop-down menu.

36. Select Enabled for the Stats Collection.

37. Type the name of the vCenter Datacenter – verify the name in the vCenter.

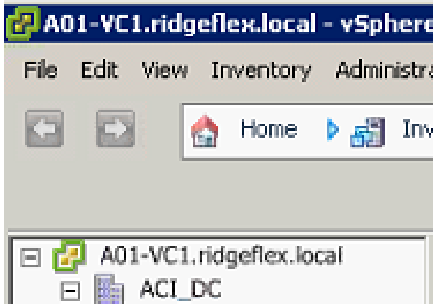

38. From the dropdown menu next to Associated Credentials, select the vCenter credentials defined in the previous steps.

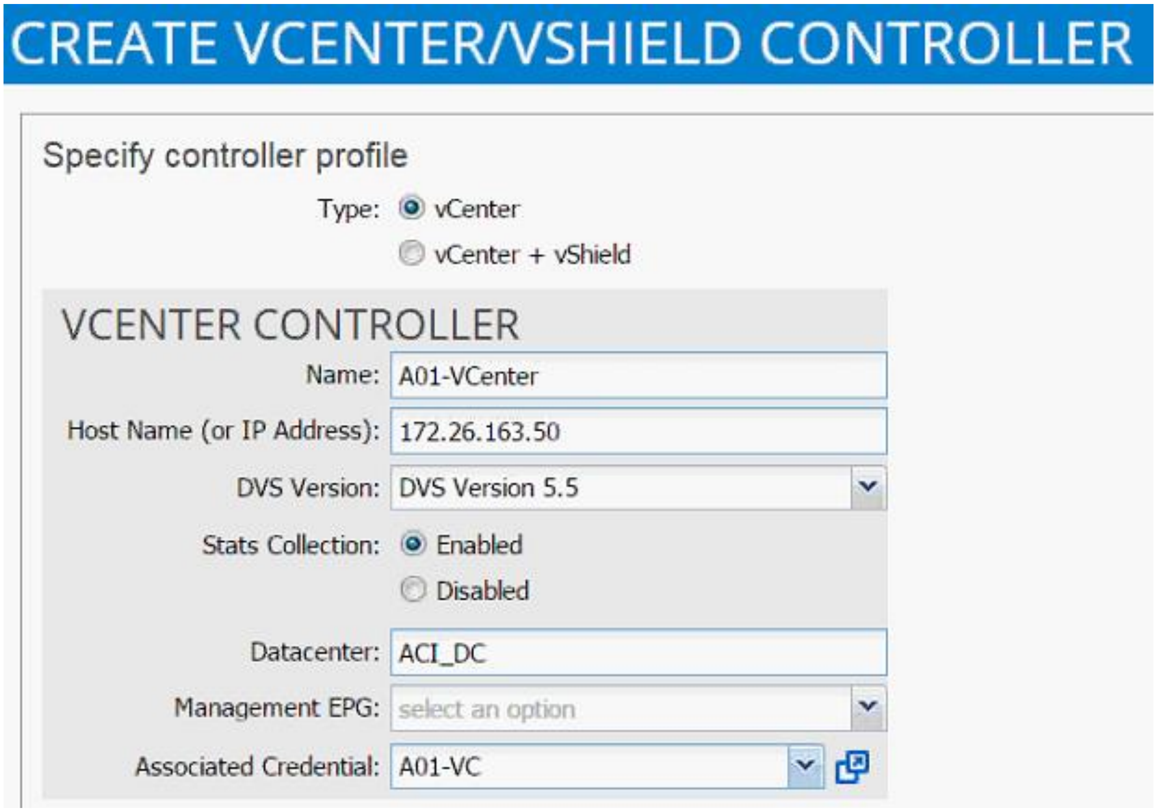

CREATE VCENTER/VSHIELD CONTROLLER

Specify controller profile

Type: ● vCenter ○ vCenter + vShield

VCENTER CONTROLLER

Name: A01-VCenter

Host Name (or IP Address): 172.26.163.50

DVS Version: DVS Version 5.5

Stats Collection: ● Enabled ○ Disabled

Datacenter: ACI_DC

Management EPG: select an option

Associated Credential: A01-VC

39. Click OK.

40. vCenter domain is now defined.

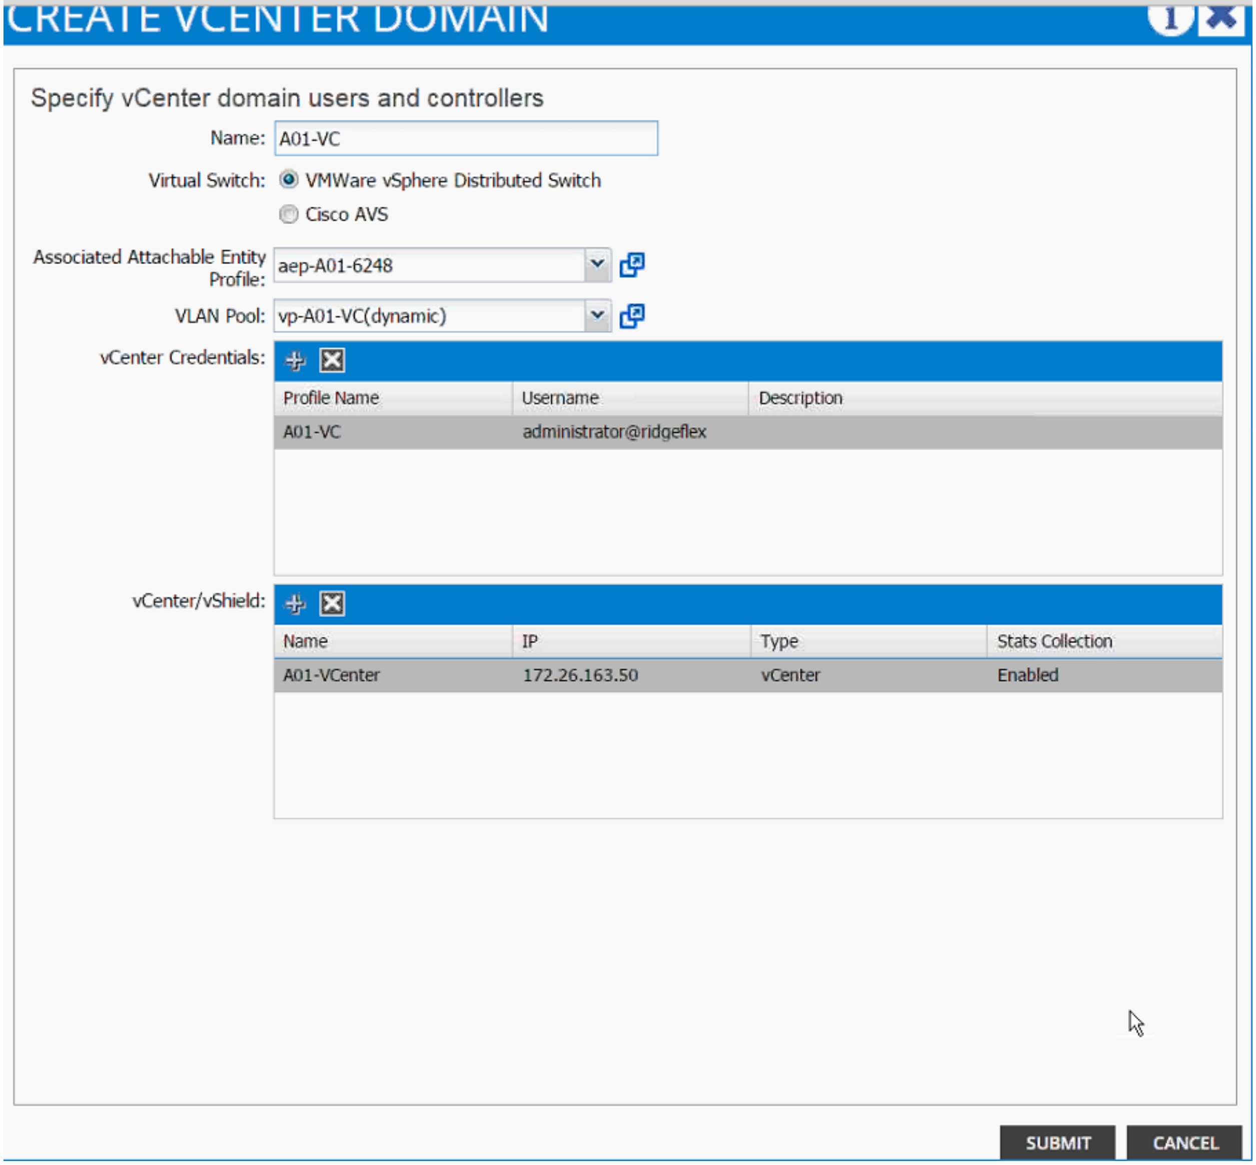

CREATE VCENTER DOMAIN

Specify vCenter domain users and controllers

Name: A01-VC

Virtual Switch: ● VMWare vSphere Distributed Switch ○ Cisco AVS

Associated Attachable Entity Profile: aep-A01-6248

VLAN Pool: vp-A01-VC(dynamic)

vCenter Credentials:

Profile Name	Username	Description
A01-VC	administrator@ridgeflex	

vCenter/vShield:

Name	IP	Type	Stats Collection
A01-VCenter	172.26.163.50	vCenter	Enabled

41. Click SUBMIT.

VMware Distributed Switch – Deployment Validation

Based on the communication defined in last section, Cisco APIC would define a new VDS in the vCenter. This can be verified using both Cisco APIC as well as VMware vCenter.

1. Log in to APIC, select VM NETWORKING from the top menu and INVENTORY from the submenu.

2. Expand VMware, vCenter and then DVS.

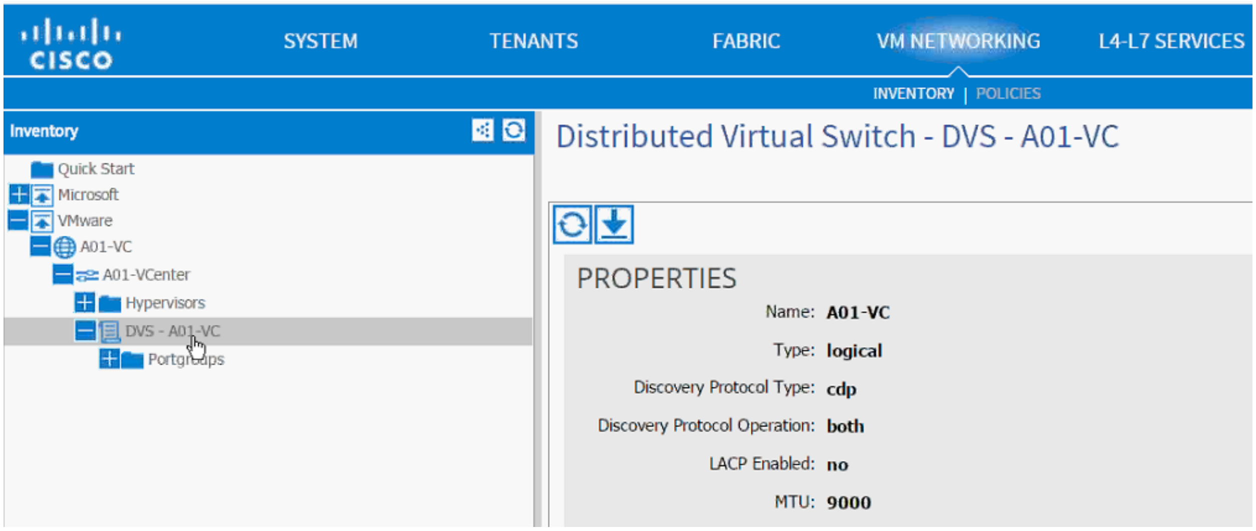

Distributed Virtual Switch - DVS - A01-VC

PROPERTIES

Name:	A01-VC
Type:	logical
Discovery Protocol Type:	cdp
Discovery Protocol Operation:	both
LACP Enabled:	no
MTU:	9000

3. Validate the MTU size, LACP value (should be disabled) and Discovery Protocol (should be CDP).

4. Log in to the vCenter using the vSphere Client.

5. Browse to Networking. A new folder and VDS is available.

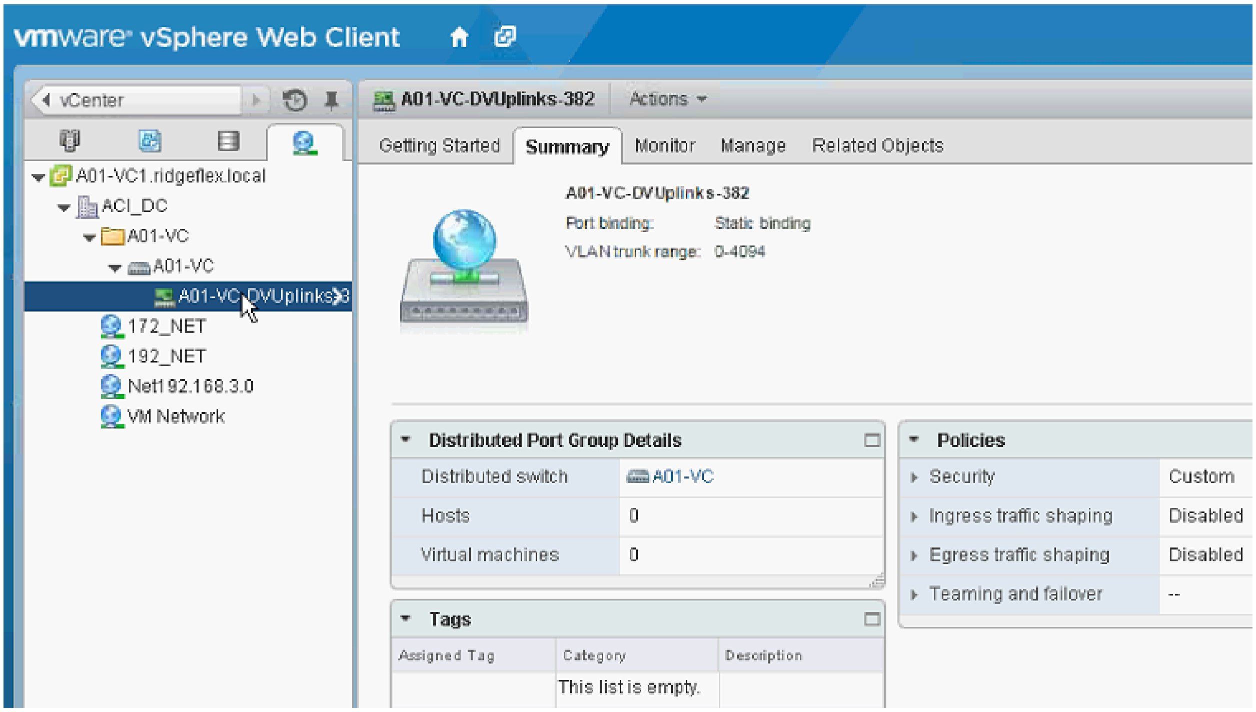

A01-VC-DVUplinks-382

Port binding: Static binding

VLAN trunk range: 0-4094

Distributed Port Group Details

Distributed switch	A01-VC
Hosts	0
Virtual machines	0

Policies

Security	Custom
Ingress traffic shaping	Disabled
Egress traffic shaping	Disabled
Teaming and failover	--

Tags

Assigned Tag	Category	Description
	This list is empty.	

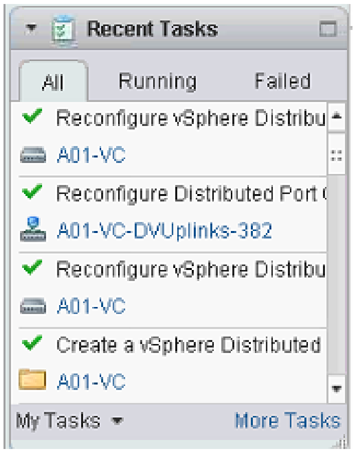

Reconfigure vSphere Distribu
A01-VC
Reconfigure Distributed Port
A01-VC-DVUplinks-382
Reconfigure vSphere Distribu
A01-VC
Create a vSphere Distributed
A01-VC

6. Right Click the Uplink PortGroup and click Edit Settings.

7. Click LACP.

8. Validate the Status is Disabled.

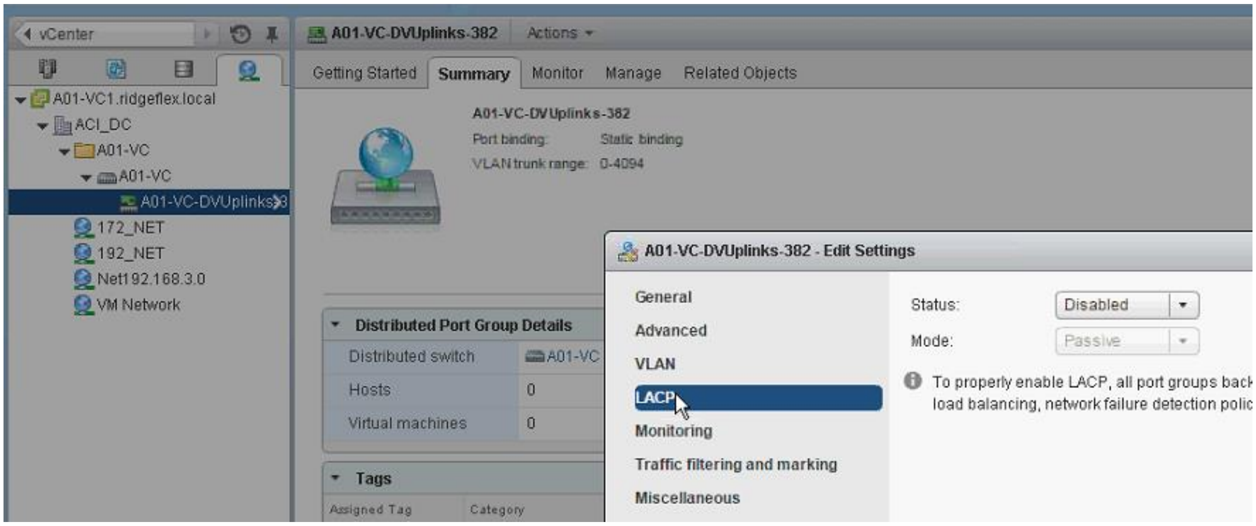

A01-VC-DVUplinks-382
Port binding: Static binding
VLAN trunk range: 0-4094

Distributed switch: A01-VC
Hosts: 0
Virtual machines: 0

Status: Disabled
Mode: Passive

To properly enable LACP, all port groups bacl
load balancing, network failure detection polic

Adding Hosts to VMware Distributed Switch

While Cisco APIC defines and configured the VDS automatically based on the user configuration, The ESXi hosts need to be added to the VDS and Uplinks need to be defined manually. For this configuration, VMware provided wizard would be utilized

1. From the networking tab in vSphere web client, right click on the VDS and click Add and Manage Hosts…

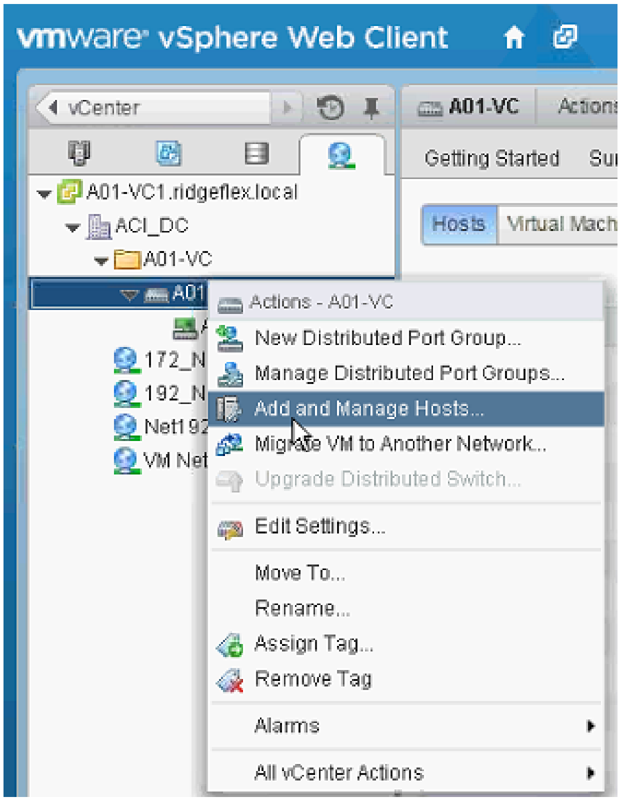

2. In the Add and Manage Hosts wizard, select Add hosts and click Next.

3. Click + to add New hosts.

4. Select all the hosts that need to be part of the VDS.

5. Click OK.

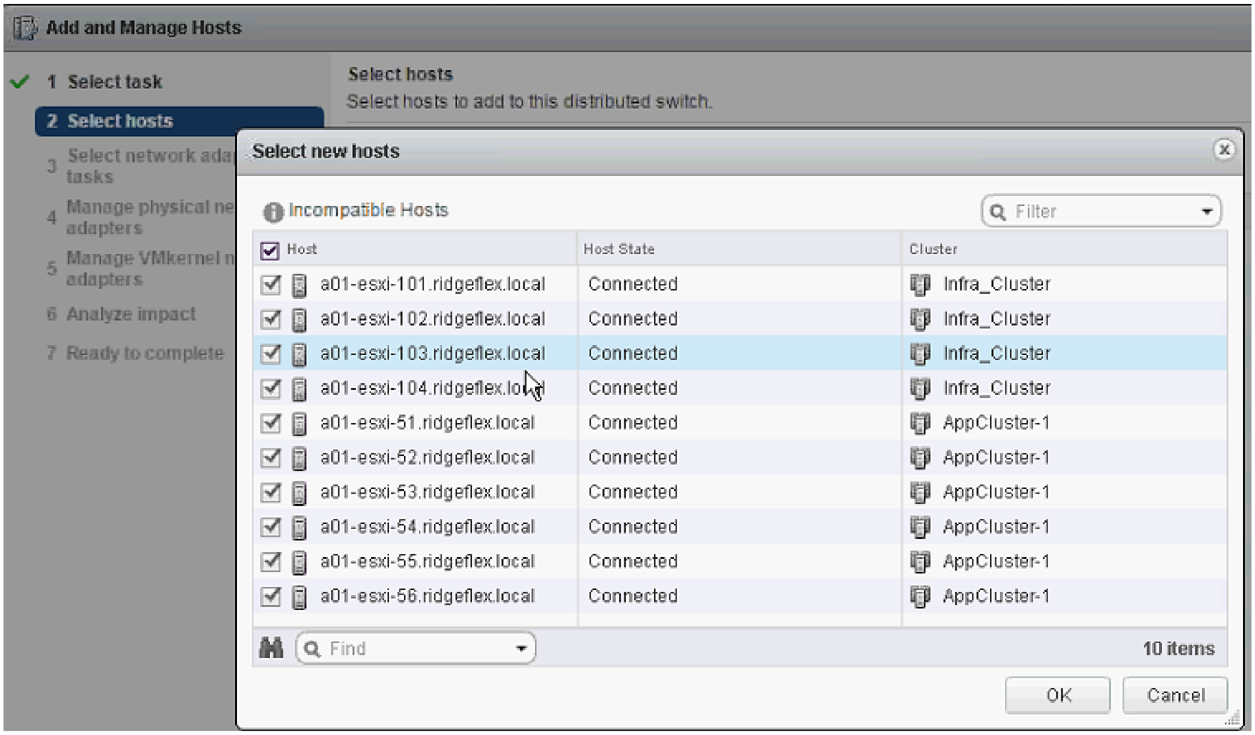

Add and Manage Hosts

- ✓ 1 Select task
- 2 Select hosts
- 3 Select network ada tasks
- 4 Manage physical ne adapters
- 5 Manage VMkernel n adapters
- 6 Analyze impact
- 7 Ready to complete

Select hosts

Select hosts to add to this distributed switch.

Select new hosts

Incompatible Hosts

☑ Host	Host State	Cluster
☑ a01-esxi-101.ridgeflex.local	Connected	Infra_Cluster
☑ a01-esxi-102.ridgeflex.local	Connected	Infra_Cluster
☑ a01-esxi-103.ridgeflex.local	Connected	Infra_Cluster
☑ a01-esxi-104.ridgeflex.local	Connected	Infra_Cluster
☑ a01-esxi-51.ridgeflex.local	Connected	AppCluster-1
☑ a01-esxi-52.ridgeflex.local	Connected	AppCluster-1
☑ a01-esxi-53.ridgeflex.local	Connected	AppCluster-1
☑ a01-esxi-54.ridgeflex.local	Connected	AppCluster-1
☑ a01-esxi-55.ridgeflex.local	Connected	AppCluster-1
☑ a01-esxi-56.ridgeflex.local	Connected	AppCluster-1

10 items

6. Click Next.

7. In the Select network adapter tasks, make sure only Manage Physical adapter tasks is selected.

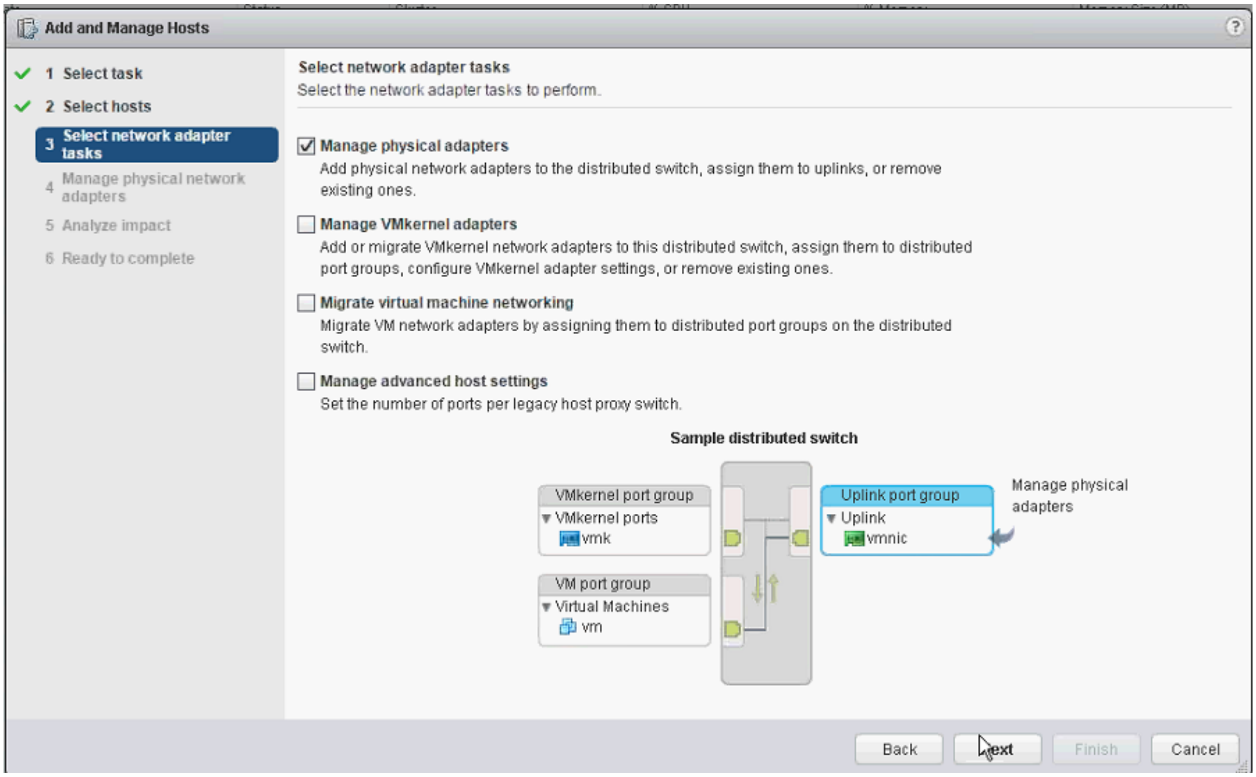

Add and Manage Hosts

- ✓ 1 Select task
- ✓ 2 Select hosts
- 3 Select network adapter tasks
- 4 Manage physical network adapters
- 5 Analyze impact
- 6 Ready to complete

Select network adapter tasks

Select the network adapter tasks to perform.

☑ Manage physical adapters

Add physical network adapters to the distributed switch, assign them to uplinks, or remove existing ones.

☐ Manage VMkernel adapters

Add or migrate VMkernel network adapters to this distributed switch, assign them to distributed port groups, configure VMkernel adapter settings, or remove existing ones.

☐ Migrate virtual machine networking

Migrate VM network adapters by assigning them to distributed port groups on the distributed switch.

☐ Manage advanced host settings

Set the number of ports per legacy host proxy switch.

Sample distributed switch

8. Click Next.

9. On the Manage physical network adapters screen, for all the hosts select vmnic 0 and vmnic 1 to the uplink ports.

10. Click Next.

11. Analyze the impact (number of adapters being added) and click Finish.

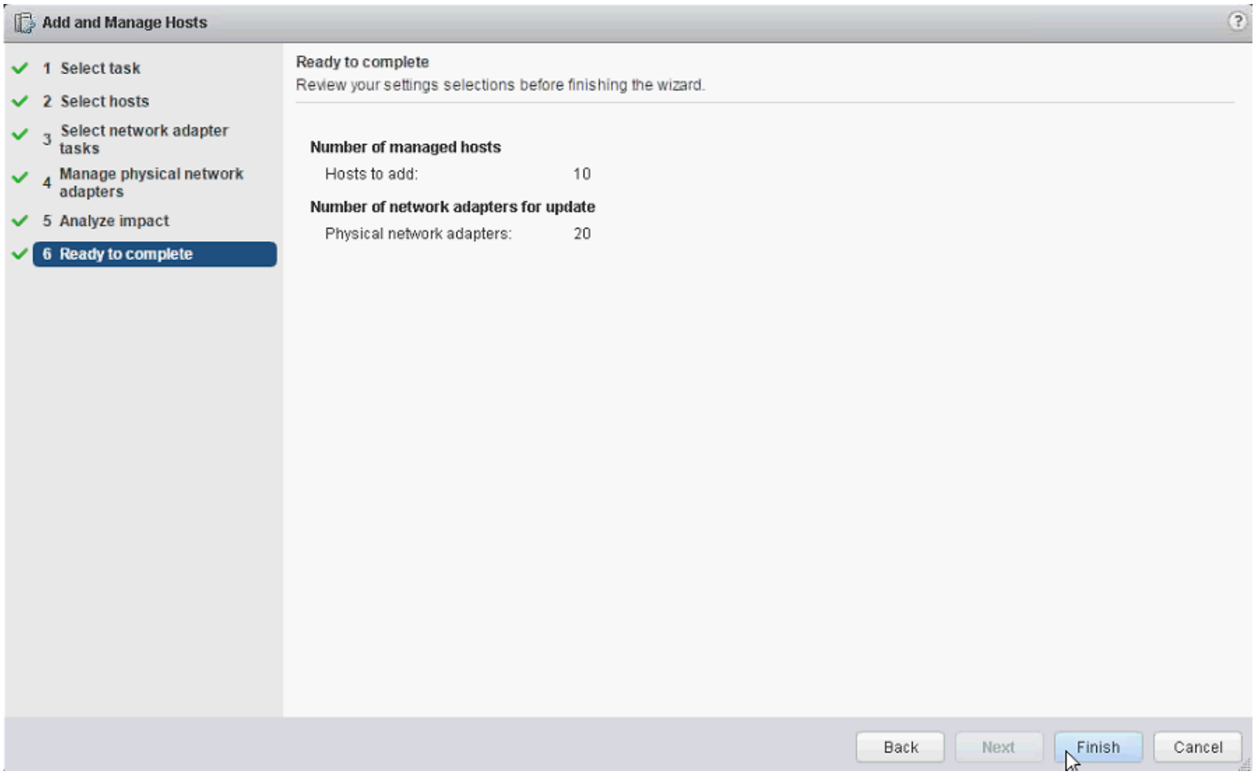

Ready to complete

Review your settings selections before finishing the wizard.

Number of managed hosts

Hosts to add: 10

Number of network adapters for update

Physical network adapters: 20

12. When the update is completed, click the Uplink port-group and validate the correct number of hosts were added to VDS.

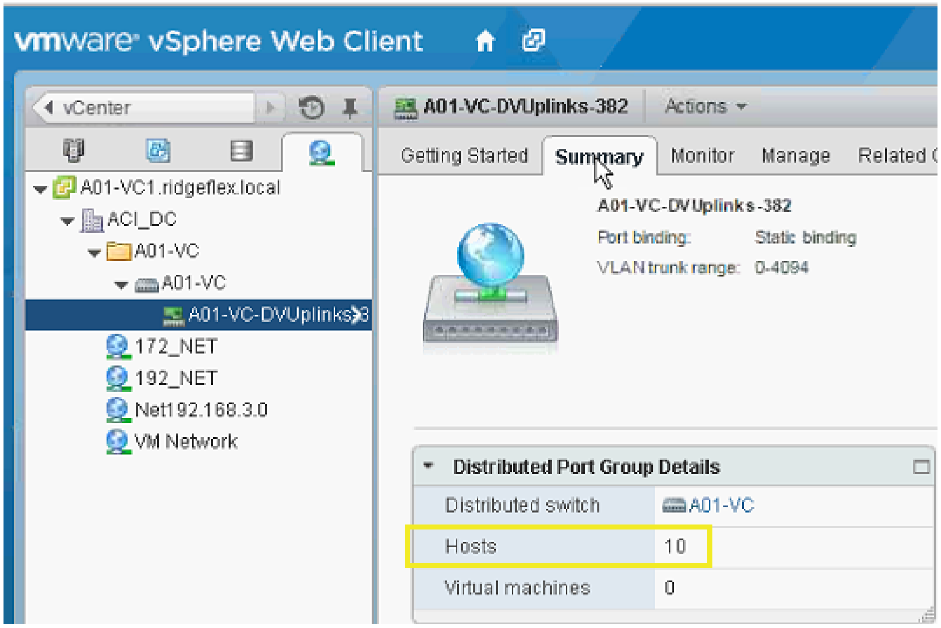

A01-VC-DVUplinks-382

Port binding: Static binding

VLAN trunk range: 0-4094

Distributed Port Group Details

Distributed switch	A01-VC
Hosts	10
Virtual machines	0

Application Profile vMotion

In this section, an Application Profiles for vMotion traffic will be created. All ESXi hosts will be configured with a VMkernel Port to act as vMotion interface.

vMotion – Application Profile Creation

1. Go to Tenant menu and select Foundation tenant from the top menu.
2. Expand Tenant Foundation in the left menu bar.
3. Right-click Application Profile and click Create Application Profiles.
4. In the CREATE APPLICATION PROFILE dialog box, enter vMotion as the Name.
5. From the drop-down menu, select default for Monitoring Policy.
6. Click + next to EPG to add an EPG.
7. Enter vmk-vmotion as the EPG Name.
8. Select bd_Internal as the Bridge Domain.
9. Select default as the Monitoring Policy.
10. Click + next to Associated Domain Profile (VMs or bare metals):
11. From the drop-down menu, select A01-VC (recently created VMM domain).
12. Set Deployment Immediacy and Resolution Immediacy to Immediate.

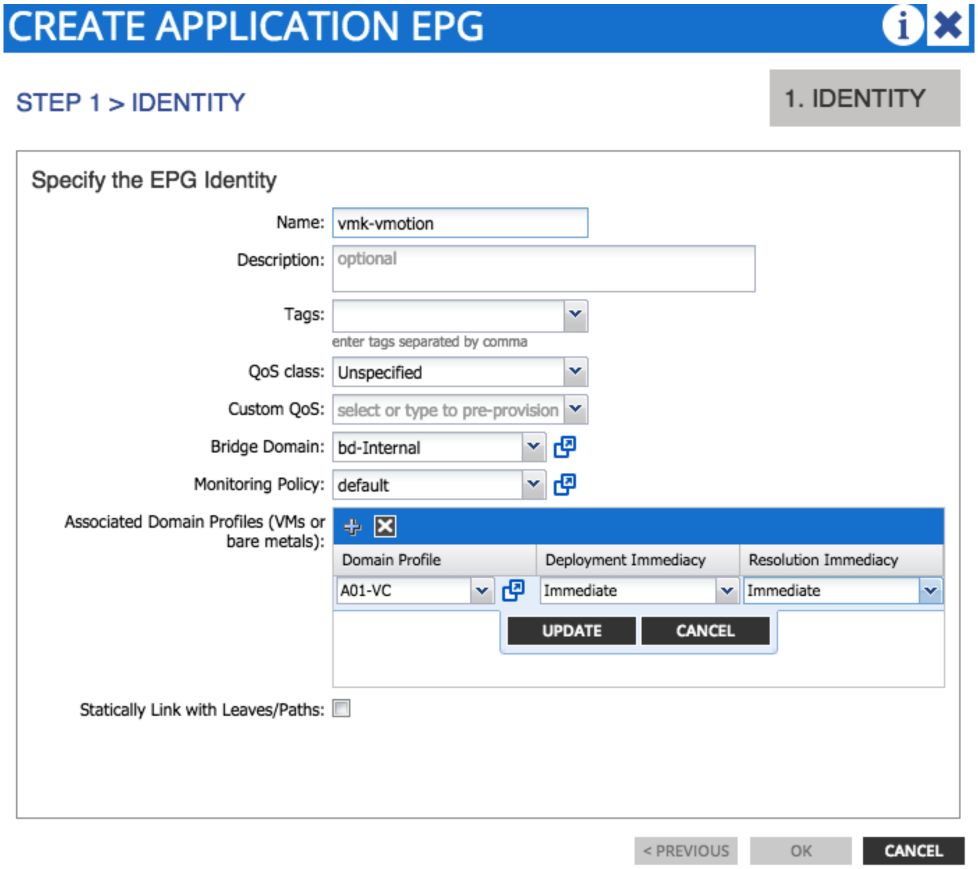

CREATE APPLICATION EPG

STEP 1 > IDENTITY

Specify the EPG Identity

Name:	vmk-vmotion
Description:	optional
Tags:	enter tags separated by comma
QoS class:	Unspecified
Custom QoS:	select or type to pre-provision
Bridge Domain:	bd-Internal
Monitoring Policy:	default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
A01-VC	Immediate	Immediate

Statically Link with Leaves/Paths: □

13. Click Update.

14. Click OK.

15. Click SUBMIT.

16. On vCenter Client, under Home -> Inventory -> Networking, make sure a port group Foundation|vMotion|vmk-vmotion is added.

17. On all the ESXi servers, add a VMKernel port, attach it to the above port-group, set the MTU to 9000 and enable vMotion traffic.

Cisco ACI – Deploying a Tenant

This section provides a detailed procedure for configuring a tenant (Business Unit or Application) using ACI as follows:

- An SVM will be deployed for tenant (named App-A)
- An Application tenant will be created on APIC (named App-A)
- NFS access will be established
- NFS Datastore will be mounted onto the Application ESXi servers for VM deployment

The procedures above will showcase interworking of Cisco APIC and VMware VDS. Communication between this tenant and an existing infrastructure (outside ACI fabric) will be enabled using OSPF routing.

Application Specific SVM Creation

VLAN in Clustered Data ONTAP

1. Create NFS VLANs.

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

2. Create iSCSI VLANs.

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

3. Create SVM Management VLANs.

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

Broadcast Domains in Clustered Data ONTAP

1. Create NFS Broadcast Domain.

```
broadcast-domain create -broadcast-domain App-A-NFS -mtu 9000 -ports
<<var_node01>>:a0a-<<var_nfs_vlan_id>>,<<var_node02>>:a0a-<<var_nfs_vlan_id>>
```

2. Create iSCSI Broadcast Domains.

```
broadcast-domain create -broadcast-domain App-A-iSCSI-A -mtu 9000 -ports
<<var_node01>>:a0a-<<var_iscsi_vlan_A_id>>,<<var_node02>>:a0a-
```

```
<<var_iscsi_vlan_A_id>>
broadcast-domain create -broadcast-domain App-A-iSCSI-B -mtu 9000 -ports
<<var_node01>>:a0a-<<var_iscsi_vlan_B_id>>,<<var_node02>>:a0a-
<<var_iscsi_vlan_B_id>>
```

3. Create SVM Management Broadcast Domain.

```
broadcast-domain create -broadcast-domain App-A-MGMT -mtu 1500 -ports
<<var_node01>>:a0a-<<var_svm_mgmt_vlan_id>>,<<var_node02>>:a0a-
<<var_svm_mgmt_vlan_id>>
```

Storage Virtual Machine (Vserver)

Note: The storage virtual machine (SVM) is referred to as Vserver (or `vserver`) in the GUI and CLI.

1. Run the `vserver create` command.

```
vserver create -vserver App-A-SVM -rootvolume rootvol -aggregate aggr1_node01 –
rootvolume-security-style unix
```

2. Select the SVM data protocols to configure, keeping `nfs` and `iscsi`.

```
vserver remove-protocols -vserver App-A-SVM -protocols fcp,cifs,ndmp
```

3. Add the two data aggregates to the Infra-SVM aggregate list for NetApp Virtual Storage Console (VSC).

```
vserver modify -vserver App-A-SVM -aggr-list aggr1_node01,aggr1_node02
```

4. Enable and run the NFS protocol in the Infra-SVM.

```
nfs create -vserver App-A-SVM -udp disabled
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plugin.

```
vserver nfs modify -vserver App-A-SVM -vstorage enabled

vserver nfs show
```

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

To create a load-sharing mirror of an SVM root volume, complete the following steps:

Note: The storage virtual machine (SVM) is referred to as Vserver (or `vserver`) in the GUI and CLI.

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

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

2. Create the mirroring relationships.

```
snapmirror create –source-path //App-A-SVM/rootvol –destination-path //App-A-
SVM/rootvol_m01 –type LS -schedule 15min
```

```
snapmirror create –source-path //App-A-SVM /rootvol –destination-path //App-A-
SVM/rootvol_m02 –type LS -schedule 15min
```

3. Initialize the mirroring relationship.

```
snapmirror initialize-ls-set –source-path //App-A-SVM/rootvol
snapmirror show
```

iSCSI Service in Clustered Data ONTAP

1. Create the iSCSI service on each Vserver. This command also starts the iSCSI service and sets the iSCSI alias to the name of the Vserver.

```
iscsi create -vserver App-A-SVM
iscsi show
```

HTTPS Access in Clustered Data ONTAP

Secure access to the storage controller must be configured. To configure secure access, complete the following steps:

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

```
set -privilege diag
```

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

2. Generally, a self-signed certificate is already in place. Check it by using the following command:

```
security certificate show
```

3. For the App-A Vserver, the certificate common name should match the DNS FQDN of the Vserver. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a Certificate Authority (CA) To delete the default certificates, run the following commands:

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

```
security certificate delete [TAB] …
Example: security certificate delete -vserver App-A-SVM -common-name
3.cert.1414163766 -ca 3.cert.1414163766 -type server -serial 544A6D36
```

4. To generate and install a self-signed certificate, run the following command as a one-time command. Generate a server certificate for the App-A Vserver. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] …
Example: security certificate create -common-name app-a-svm.ridgeflex.local -type
server -size 2048 -country US -state "North Carolina" -locality "RTP" -
```

```
organization "Cisco" -unit "SAVBU" -email-addr "abc@cisco.com" -expire-days 365 –
protocol SSL -hash-function SHA256 –vserver App-A-SVM
```

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

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

```
security ssl modify [TAB] …
Example: security ssl modify -vserver App-A-SVM -server-enabled true -client-
enabled false -ca app-a-svm.ridgeflex.local -serial 544A71D7 -common-name app-
a.ridgeflex.local
```

7. Change back to normal admin privilege level and set up to allow Vserver logs to be available by web.

```
set –privilege admin

vserver services web modify –name spi|ontapi|compat –vserver * -enabled true
```

NFSv3 in Clustered Data ONTAP

To configure NFS on the Vserver, run all commands.

1. Modify the initial default rule for the SVM NFS subnet in the default export policy.

```
vserver export-policy rule create –vserver App-A-SVM -policyname default –
ruleindex 1 –protocol nfs -clientmatch <<var_nfs_subnet_address>> -rorule sys –
rwrule sys -superuser sys –allow-suid false
vserver export-policy rule show
```

2. Assign the FlexPod export policy to the App-A Vserver root volume.

```
volume modify –vserver App-A-SVM –volume rootvol –policy default
```

FlexVol in Clustered Data ONTAP

1. The following information is required to create a FlexVol® volume: the volume’s name and size, and the aggregate on which it will exist. Create two NFS VMware datastore volumes and an iSCSI LUN volume. Also, update the Vserver root volume load sharing mirrors to make the NFS mounts accessible.

```
volume create –vserver App-A-SVM –volume app_a_datastore_1 –aggregate aggr1_node2
-size 500GB –state online –policy default –junction-path /app_a_datastore_1 -
space-guarantee none –percent-snapshot-space 0



volume create –vserver App-A-SVM –volume iSCSI_LUN –aggregate aggr1_node1 –size
100GB –state online –policy default –space-guarantee none –percent-snapshot-space
0


snapmirror update-ls-set -source-path //App-A-SVM/rootvol
```

Deduplication in Clustered Data ONTAP

1. Enable deduplication on appropriate volumes.

```
volume efficiency on –vserver App-A-SVM -volume app_a_datastore_1
volume efficiency on –vserver App-A-SVM -volume iSCSI_LUN
```

NFS LIF in Clustered Data ONTAP

1. Create an NFS logical interface (LIF).

```
network interface create -vserver App-A-SVM -lif nfs_app_a_datastore_1 -role data
-data-protocol nfs -home-node <<var_node02>> -home-port a0a-<<var_nfs_vlan_id>> –
address <<var_node02_nfs_lif_app_a_datastore_1_ip>> -netmask
<<var_node02_nfs_lif_app_a_datastore_1_mask>> -status-admin up –failover-policy
broadcast-domain-wide –firewall-policy data –auto-revert true –failover-group
App-A-NFS
```

Note: It is recommended to create a new lif for each datastore.

iSCSI LIF in Clustered Data ONTAP (Optional)

1. Create iSCSI logical interfaces (LIFs).

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

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


network interface create -vserver App-A-SVM -lif iscsi_lif02a -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iscsi_vlan_A_id>> –
address <<var_node02_iscsi_lif02a_ip>> -netmask <<var_node02_iscsi_lif02a_mask>>
-status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false


network interface create -vserver App-A-SVM -lif iscsi_lif02b -role data -data-
protocol iscsi -home-node <<var_node02>> -home-port a0a-<<var_iscsi_vlan_B_id>> –
address <<var_node02_iscsi_lif02b_ip>> -netmask <<var_node02_iscsi_lif02b_mask>>
-status-admin up -failover-policy disabled -firewall-policy data -auto-revert
false

network interface show –vserver App-A-SVM
```

Add App-A Vserver Administrator

1. Add the App-A Vserver administrator and Vserver administration logical interface in the SVM management network with the following commands:

```
network interface create –vserver App-A-SVM –lif vsmgmt –role data –data-protocol
none –home-node <<var_node02>> -home-port a0a-<<var_svm_mgmt_vlan_id>> –address
<<var_vserver_mgmt_ip>> -netmask <<var_vserver_mgmt_mask>> -status-admin up –
failover-policy broadcast-domain-wide –firewall-policy mgmt –auto-revert true –
failover-group App-A-MGMT



network route create –vserver App-A-SVM -destination 0.0.0.0/0 –gateway
<<var_vserver_mgmt_gateway>>
network route show



security login password –username vsadmin –vserver App-A-SVM
Enter a new password:  <<var_vsadmin_password>>
Enter it again:  <<var_vsadmin_password>>


security login unlock –username vsadmin –vserver App-A-VSM
```

Application Tenant Creation on APIC

1. From the main menu, click TENANTS and from the sub-menu click ADD TENANT.

2. In the CREATE TENANT dialog box, type <Name of Application> as the name of the tenant. This example used App-A as the name of the tenant.

3. Click the checkbox next to all under Security Domains.

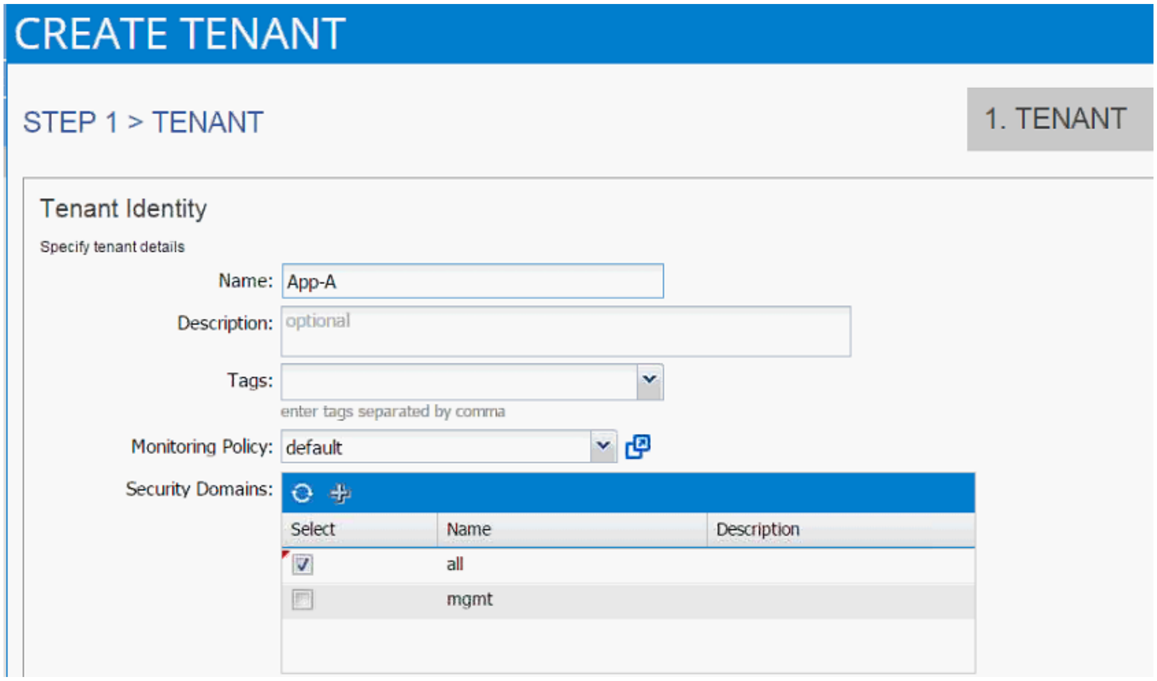

CREATE TENANT

STEP 1 > TENANT

Tenant Identity

Specify tenant details

Name: App-A

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

Security Domains:

Select	Name	Description
☑	all	
☐	mgmt	

4. Click Next.

5. Click the + sign to add network.

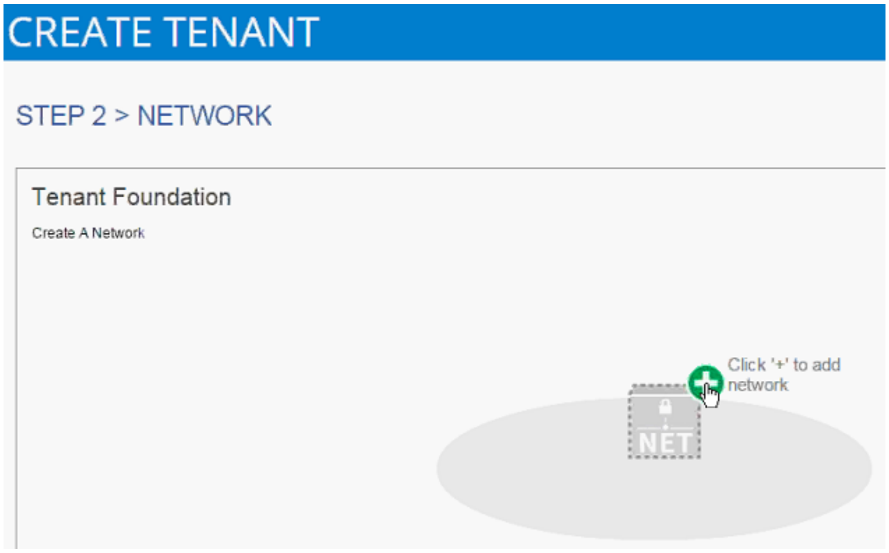

CREATE TENANT

STEP 2 > NETWORK

Tenant Foundation

Create A Network

6. In the CREATE NEW NETWORK dialog box, type App-A as the Name. Leave everything else as default.

7. Click Next to move onto bridge domain creation.

8. Use bd-Internal as the Name of the bridge domain.

9. Select default for IGMP Snoop Policy.

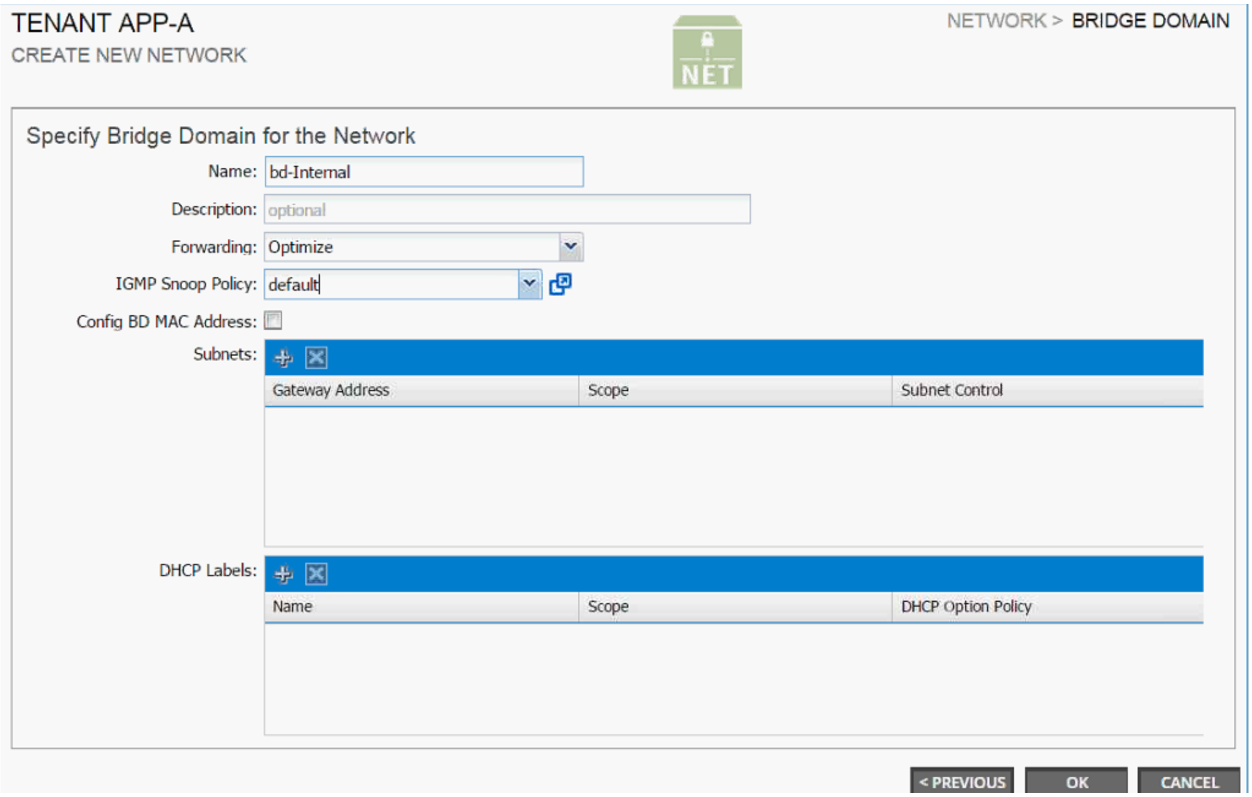

TENANT APP-A

CREATE NEW NETWORK

NETWORK > BRIDGE DOMAIN

Specify Bridge Domain for the Network

Name: bd-Internal

Description: optional

Forwarding: Optimize

IGMP Snoop Policy: default

Config BD MAC Address: □

Subnets:

Gateway Address	Scope	Subnet Control

DHCP Labels:

Name	Scope	DHCP Option Policy

10. Click OK.

11. The Bridge Domains and App-A network should be visible in the CREATE TENANT dialog box.

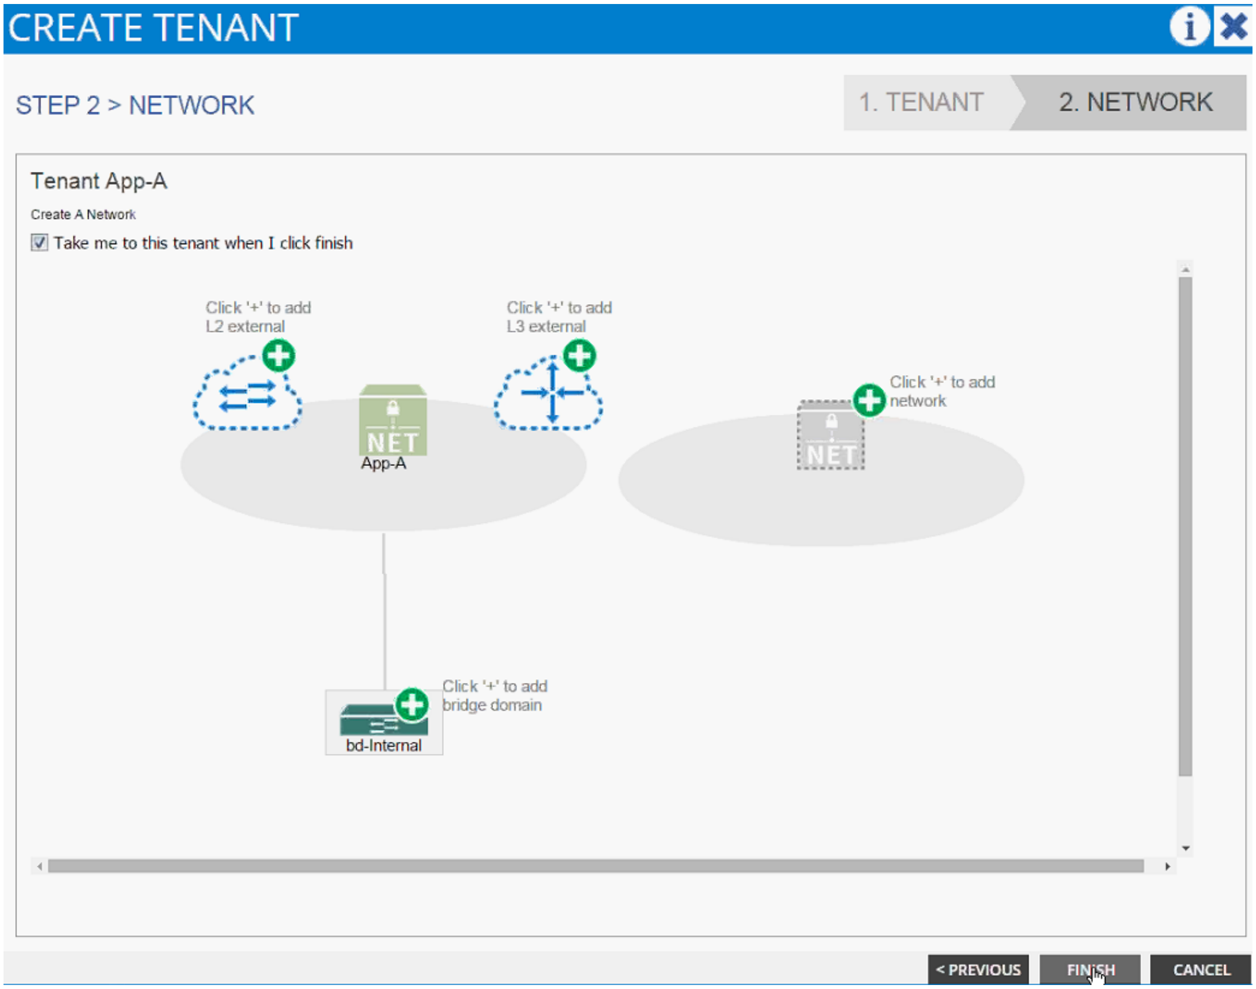

☑ Take me to this tenant when I click finish

12. Click Finish.

13. Verify the selected tenant is the newly created App-A tenant by looking at the items highlighted in the top menu.

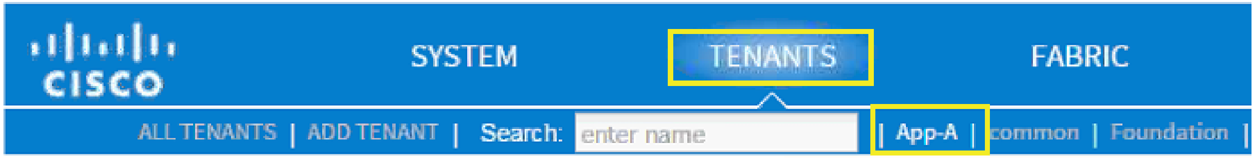
cisco
SYSTEM
TENANTS
FABRIC
ALL TENANTS | ADD TENANT | Search: enter name | App-A | common | Foundation |

3-Tier App - Application Profile Creation and Adding EPG for Web Tier

To create an Application Profile to host the application, complete the following steps:

1. Select Tenant and newly created App-A tenant from the top menu.

2. Expand Tenant App-A in the left menu bar.

3. Right-click Application Profile and click Create Application Profile.

4. In the CREATE APPLICATION PROFILE dialog box, enter 3-Tier-App as the Name.

5. From the drop-down menu, select default for Monitoring Policy.

6. Click + next to EPG to add an EPG.

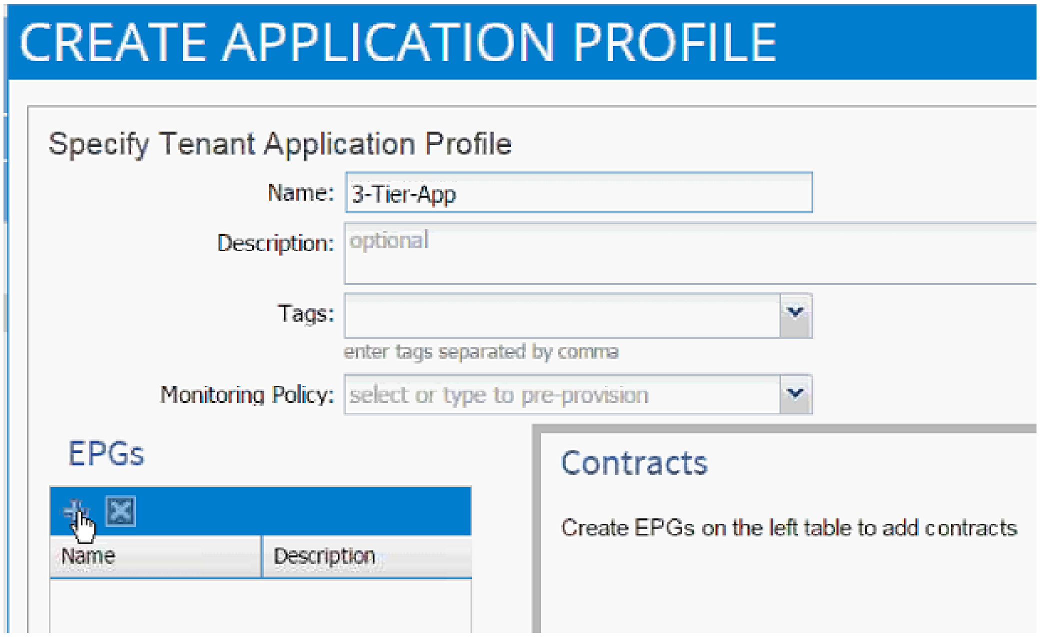

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: 3-Tier-App

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: select or type to pre-provision

EPGs

Name	Description

Contracts

Create EPGs on the left table to add contracts

7. In the CREATE APPLICATION EPG dialog box, enter Web as the Name.

8. From the drop-down menu, select bd-Internal as the Bridge Domain.

9. From the drop-down menu, select default for Monitoring Policy.

10. Click OK.

11. Click + next to Associated Domain Profiles (VMs or Bare metals).

12. From the drop-down menu, select the VMM domain previously defined.

13. Select Immediate for Deployment Immediacy.

14. Select Immediate for Resolution Immediacy.

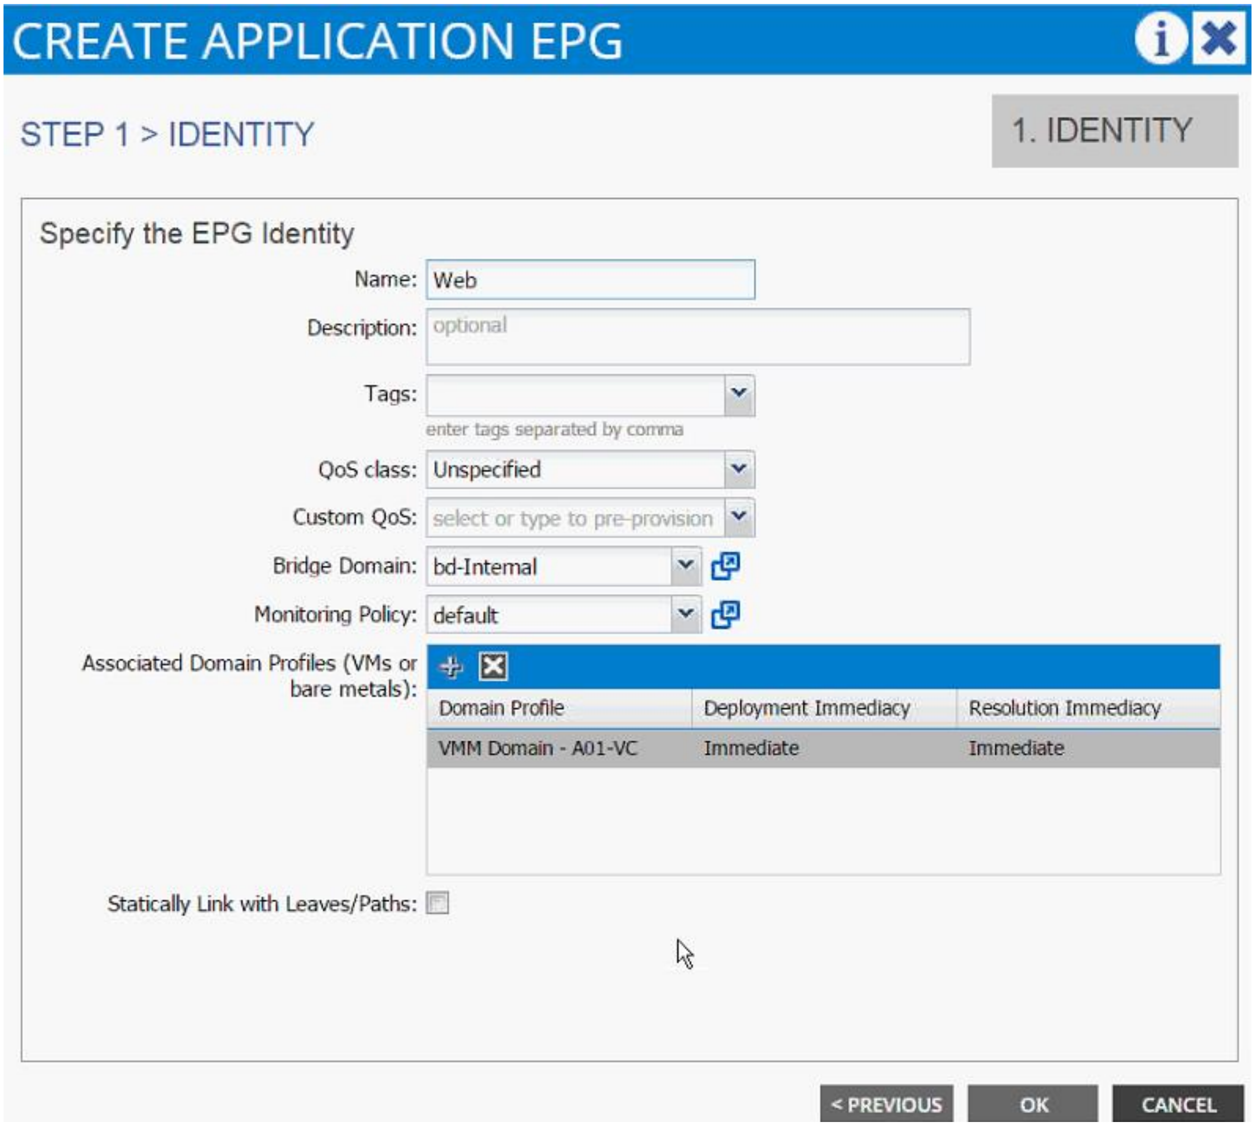

CREATE APPLICATION EPG

STEP 1 > IDENTITY

1. IDENTITY

Specify the EPG Identity

Name: Web

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
VMM Domain - A01-VC	Immediate	Immediate

Statically Link with Leaves/Paths: □

15. Click UPDATE.

16. Click OK.

17. Click SUBMIT to finish creating Application Profile.

18. Expand the newly created EPG Web and click Subnets.

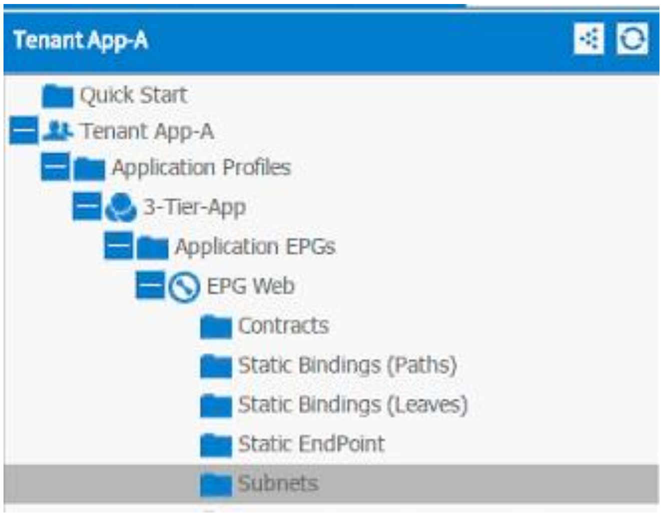

19. Click the Action menu on the right and select Create EPG Subnet.

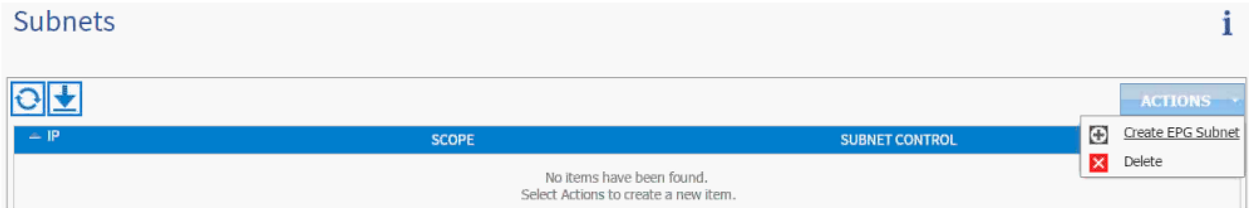

IP	SCOPE	SUBNET CONTROL

No items have been found.
Select Actions to create a new item.

20. Enter 10.10.1.254/24 for the Default Gateway IP. This IP address is the gateway that all the Web VMs will use.

21. Change scope to only Shared Subnet.

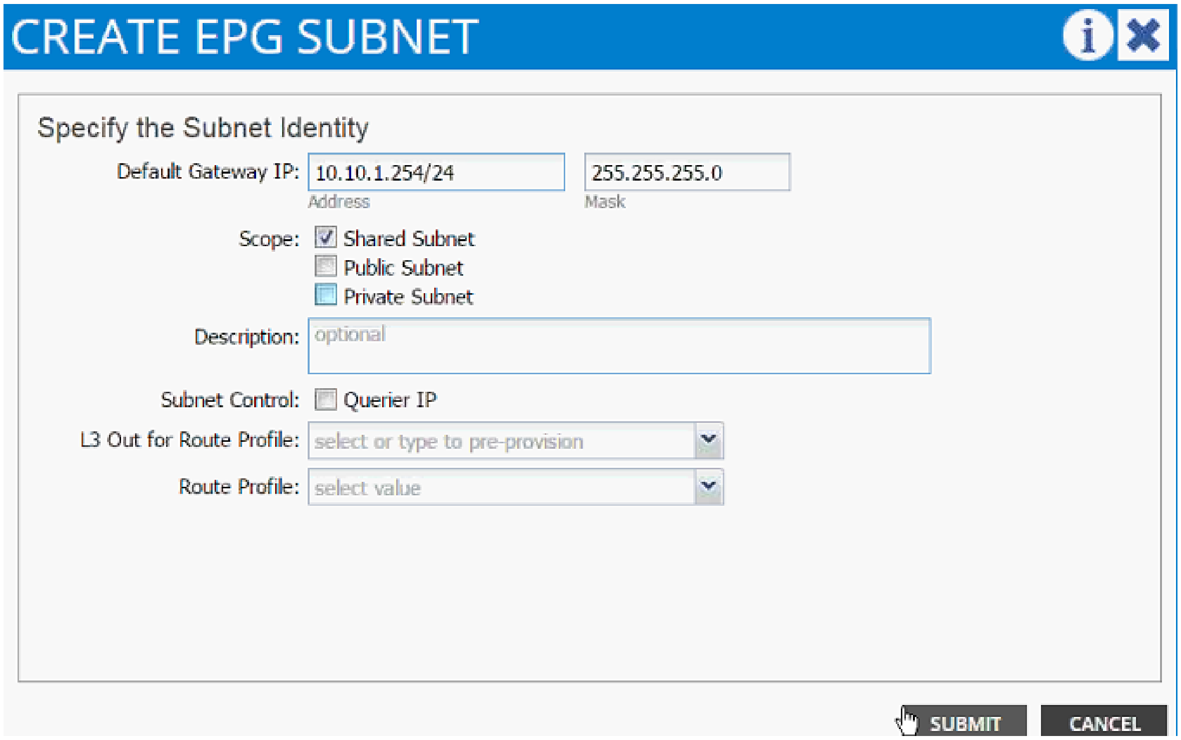

CREATE EPG SUBNET

Specify the Subnet Identity

Default Gateway IP: 10.10.1.254/24 (Address) 255.255.255.0 (Mask)

Scope: ☑ Shared Subnet
□ Public Subnet
□ Private Subnet

Description: optional

Subnet Control: □ Querier IP

L3 Out for Route Profile: select or type to pre-provision

Route Profile: select value

22. Click SUBMIT.

Port-Group Deployment Validation

When an EPG is tied to a VMM domain (previous section), a port-group gets created on the VDS so that application admin can deploy his virtual machine and make it part of the application tier just defined. A dynamic VLAN is associated with this newly created EPG/port-group. To validate the configuration, complete the following steps:

1. Browse to the VM NETWORKING and on the left menu bar, drill down to VDS to list the port-groups. As seen from the figure, VLAN 1135 from the pre-defined range 1101-1200 was assigned to the newly create port-group.

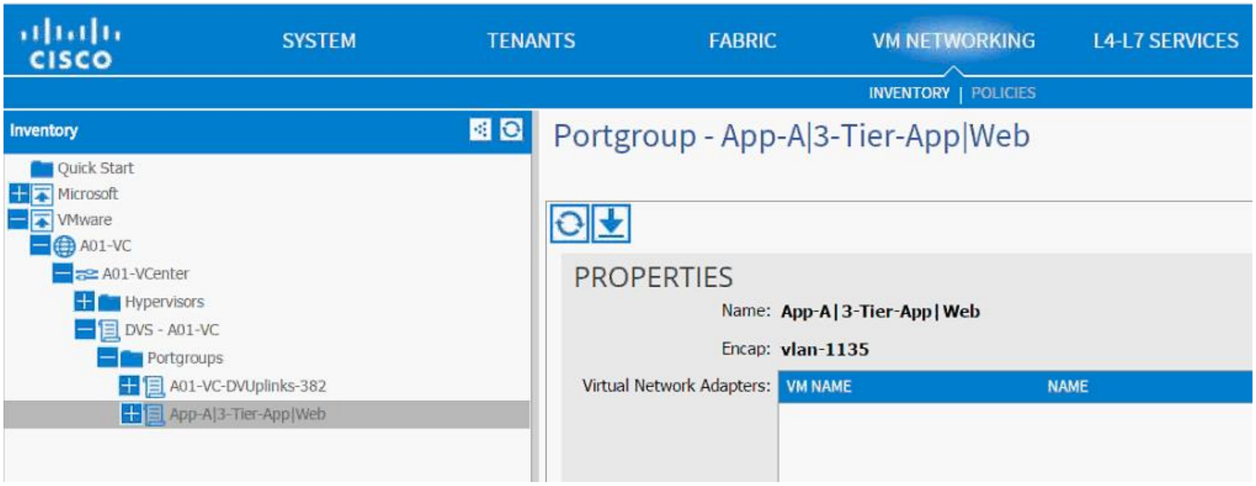

Portgroup - App-A|3-Tier-App|Web

PROPERTIES

Name: App-A|3-Tier-App|Web

Encap: vlan-1135

Virtual Network Adapters:

VM NAME	NAME

2. Browse to FABRIC and in the left menu bar, expand the Leaf, Interface, vPC Interfaces and domain 10. Click the vPC associated with UCS Fabric Interconnect.

3. VLAN 1135 should have been added to the vPC VLANs.

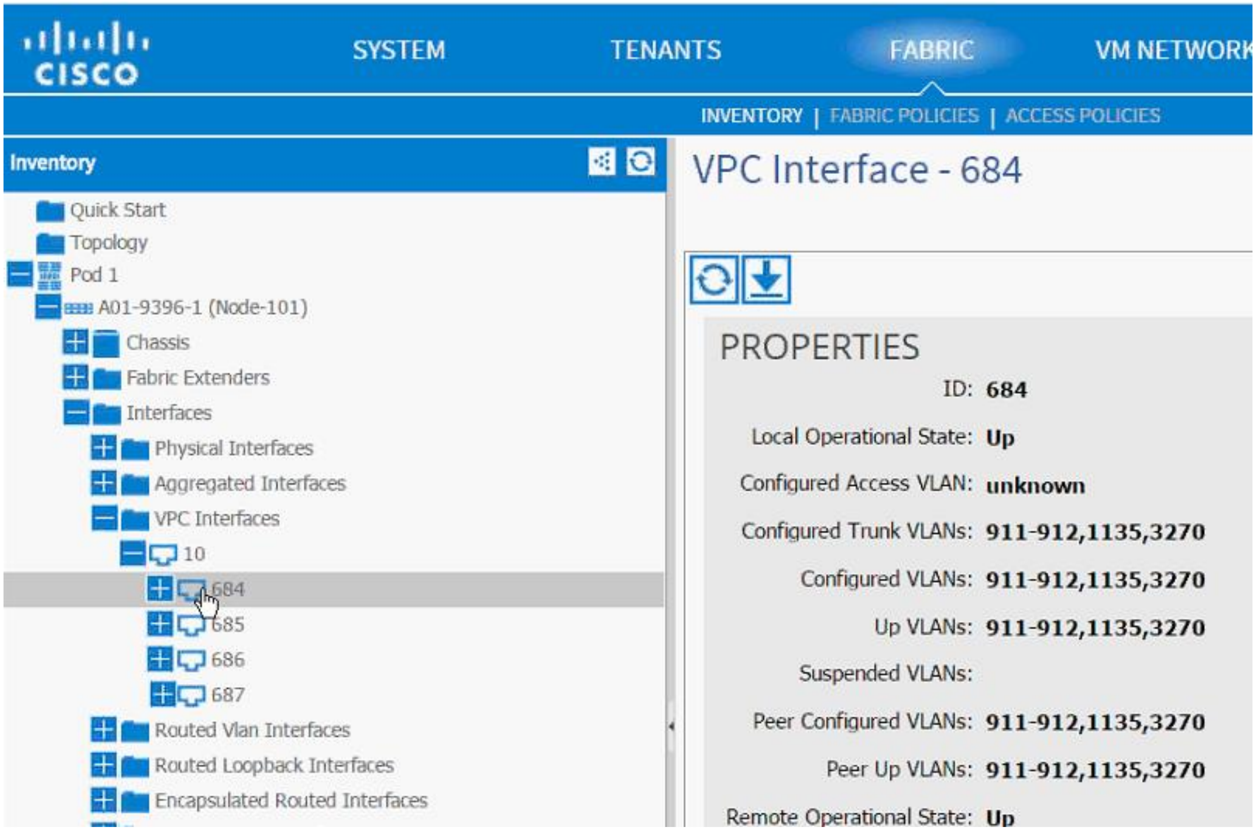

VPC Interface - 684

PROPERTIES

ID: 684
Local Operational State: Up
Configured Access VLAN: unknown
Configured Trunk VLANs: 911-912,1135,3270
Configured VLANs: 911-912,1135,3270
Up VLANs: 911-912,1135,3270
Suspended VLANs:
Peer Configured VLANs: 911-912,1135,3270
Peer Up VLANs: 911-912,1135,3270
Remote Operational State: Up

4. Log into the vCenter and browse to the VDS. Right-click the newly added port-group and click Teaming and Failover.

5. Validate Load balancing method is Route based on originating virtual port (VMware default).

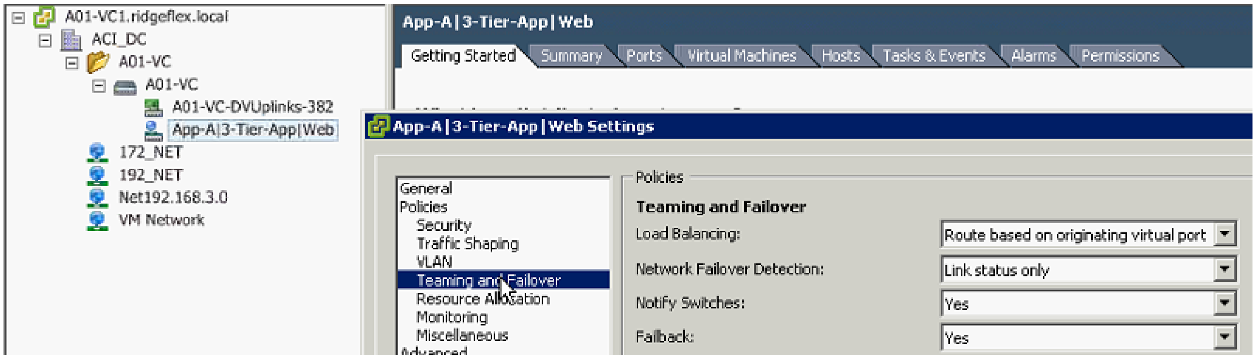

Policies

Teaming and Failover

Load Balancing: Route based on originating virtual port
Network Failover Detection: Link status only
Notify Switches: Yes
Failback: Yes

6. Validate the VLAN assigned is 1135.

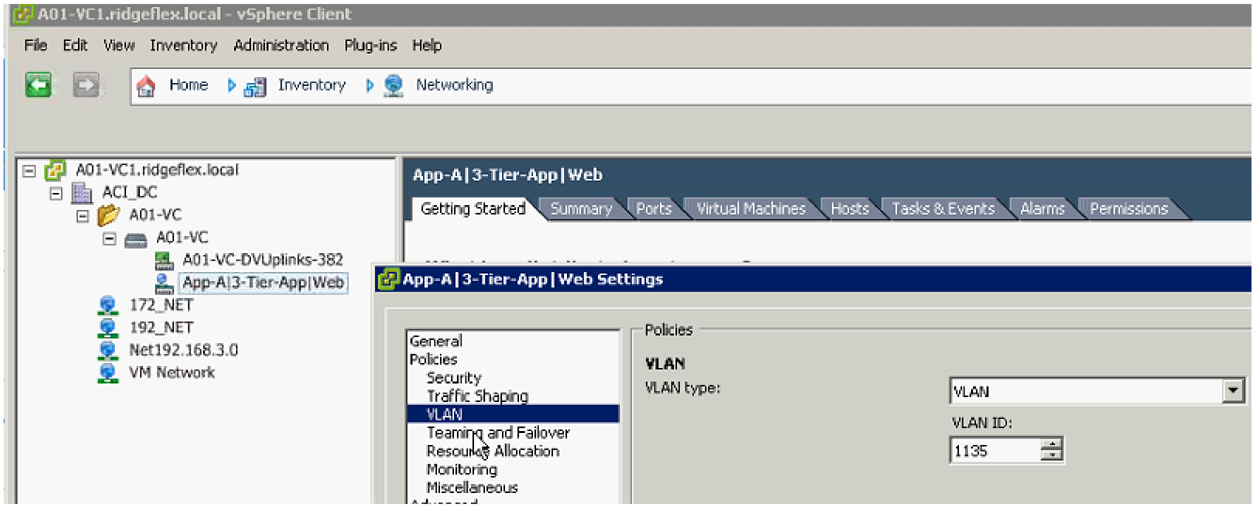

VLAN type: VLAN

VLAN ID: 1135

Modifying the Storage (NetApp) Physical Domain VLANs

Physical domain associated with NetApp storage was initially configured for infrastructure iSCSI and NFS VLANs. In this step, NFS and iSCSI VLANs associated with App-A SVM will be added to the physical domain. To add to the physical domain, complete the following steps:

1. Select Fabric and Access Policies from the top menu.

2. Expand Pools and then VLAN.

3. Click the pool name associated with NetApp controllers (vp-A01-NetApp in this example) and click + to add another Encap Block.

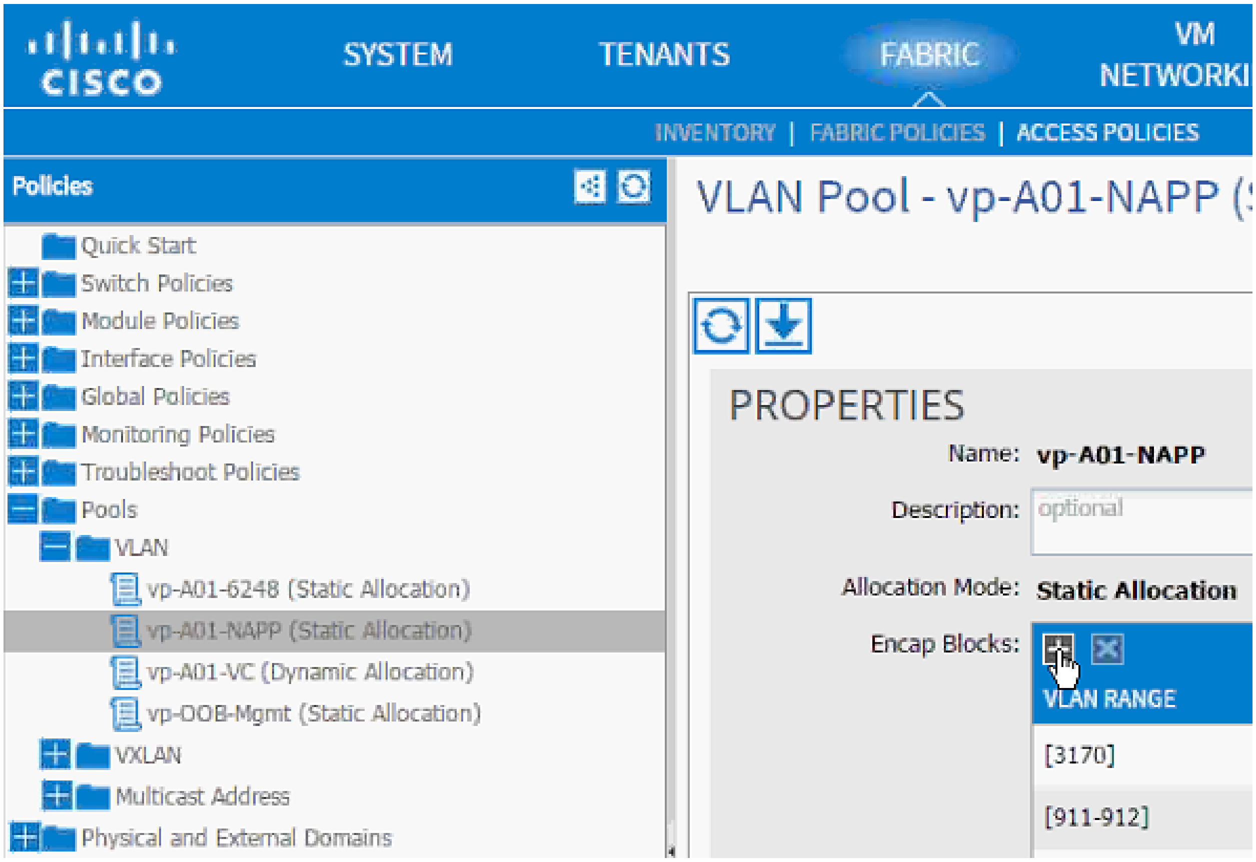
SYSTEM
TENANTS
FABRIC
VM NETWORKI
INVENTORY | FABRIC POLICIES | ACCESS POLICIES
Policies
Quick Start
Switch Policies
Module Policies
Interface Policies
Global Policies
Monitoring Policies
Troubleshoot Policies
Pools
VLAN
vp-A01-6248 (Static Allocation)
vp-A01-NAPP (Static Allocation)
vp-A01-VC (Dynamic Allocation)
vp-OOB-Mgmt (Static Allocation)
VXLAN
Multicast Address
Physical and External Domains

VLAN Pool - vp-A01-NAPP (S

PROPERTIES

Name: vp-A01-NAPP

Description: optional

Allocation Mode: Static Allocation

Encap Blocks:

VLAN RANGE
[3170]
[911-912]

4. Enter <<var_nfs_vlan_tenant>> (3180 to 3180 in this example) as the range of VLANs for Tenant NFS.

5. (Optional) Add the range of VLANs for iSCSI-a and iSCSI-b. If configuring the tenant for iSCSI access.

6. Click SUBMIT.

NFS - Application Profile Creation

In this section, you will create an Application Profile to setup NFS connectivity between the ESXi servers and the Application specific SVM. An NFS datastore will then be mounted to host application specific VMs. To create the application profile, complete the following steps:

1. Select Tenant and App-A tenant from the top menu.

2. Expand Tenant App-A in the left menu bar.

3. Right-click Application Profile and click Create Application Profile.

4. In the CREATE APPLICATION PROFILE dialog box, enter NFS as the Name.

5. From the drop-down menu, select default for Monitoring Policy.

6. Click + next to EPG to add an EPG.

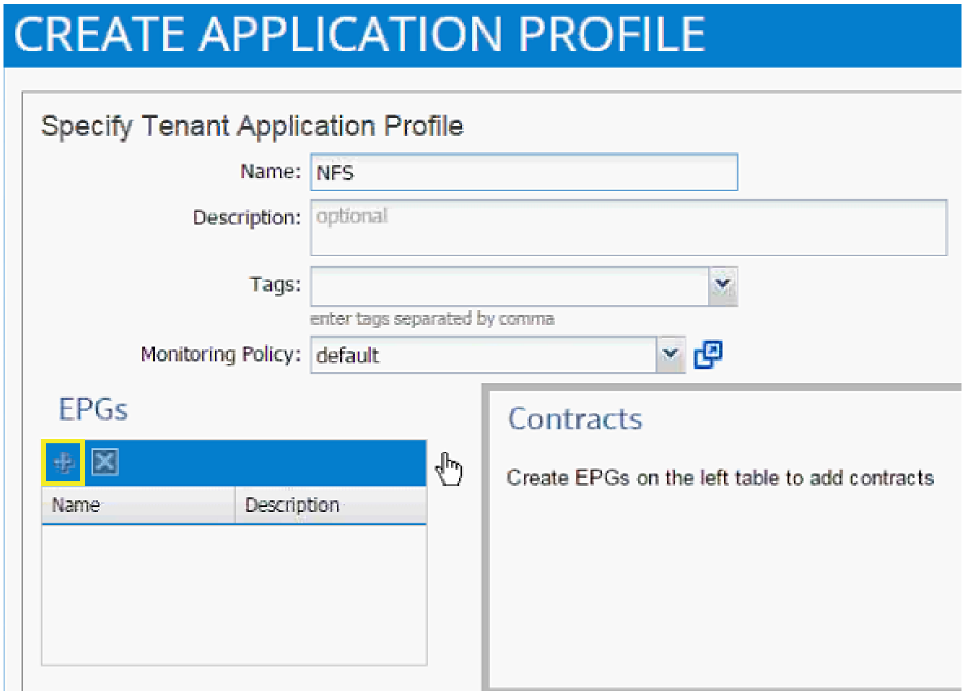

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: NFS

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

EPGs

Name	Description

Contracts

Create EPGs on the left table to add contracts

7. In the CREATE APPLICATION EPG dialog box, enter vmk-nfs as the Name.

8. From the drop-down menu, select bd-Internal as the Bridge Domain.

9. From the drop-down menu, select default for Monitoring Policy.

10. Click OK.

11. Click + next to Associated Domain Profiles (VMs or Bare metals).

12. From the drop-down menu, select the VMM domain previously defined.

13. Select Immediate for Deployment Immediacy.

14. Select Immediate for Resolution Immediacy.

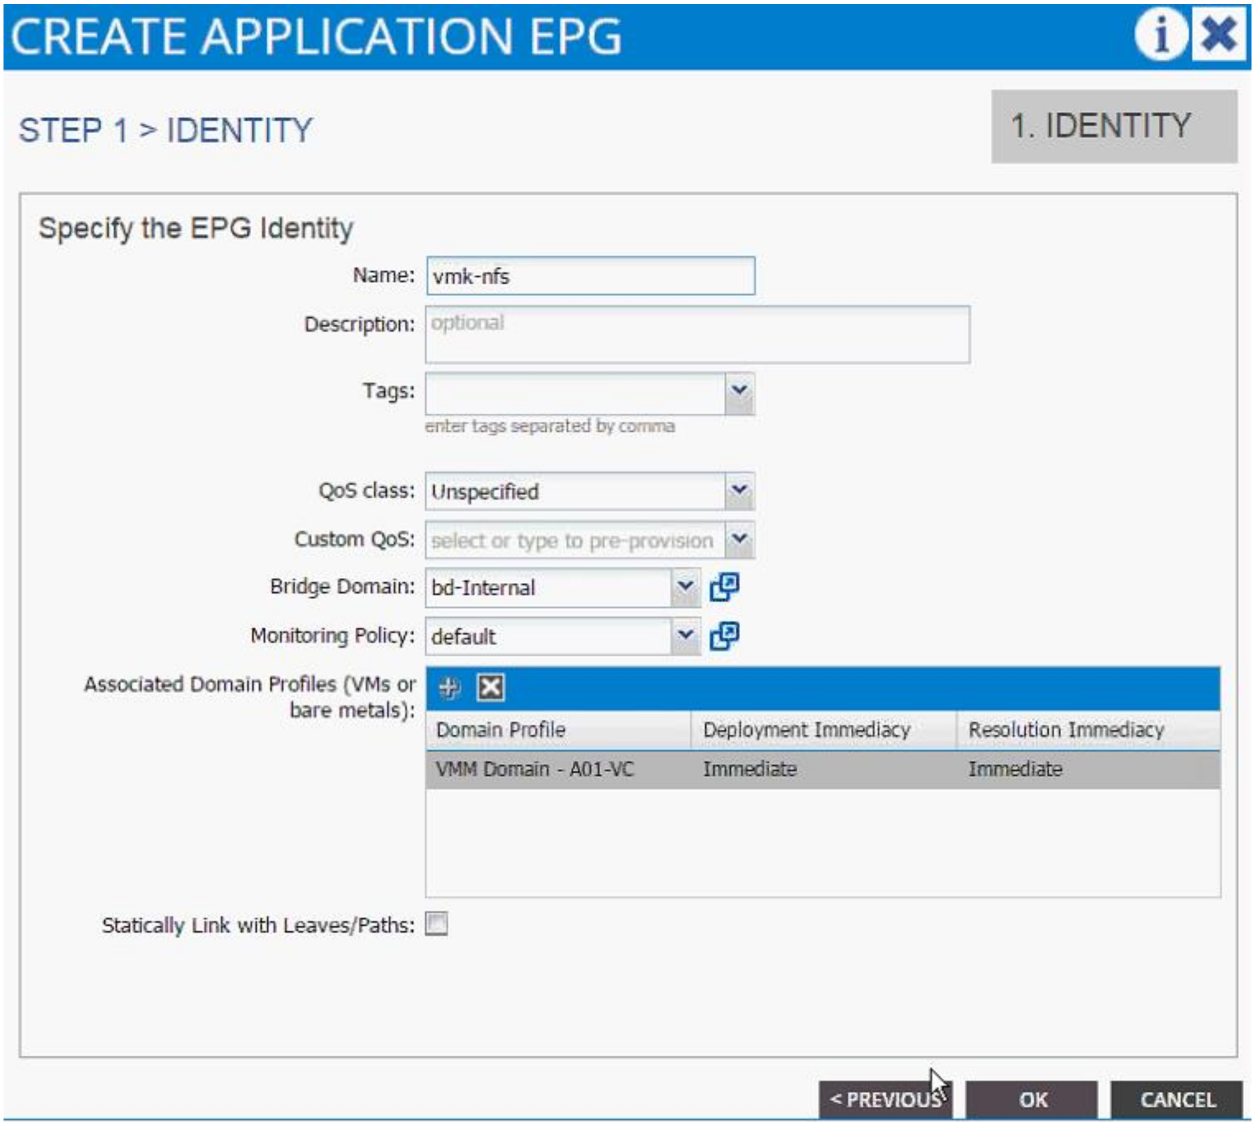

CREATE APPLICATION EPG

STEP 1 > IDENTITY

1. IDENTITY

Specify the EPG Identity

Name: vmk-nfs

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
VMM Domain - A01-VC	Immediate	Immediate

Statically Link with Leaves/Paths: □

15. Click UPDATE.

16. Click OK.

17. Click + next to EPG to add another EPG.

18. In the CREATE APPLICATION EPG dialog box, enter lif-nfs as the Name.

19. From the drop-down menu, select bd-internal as the Bridge Domain.

20. From the drop-down menu, select default for Monitoring Policy.

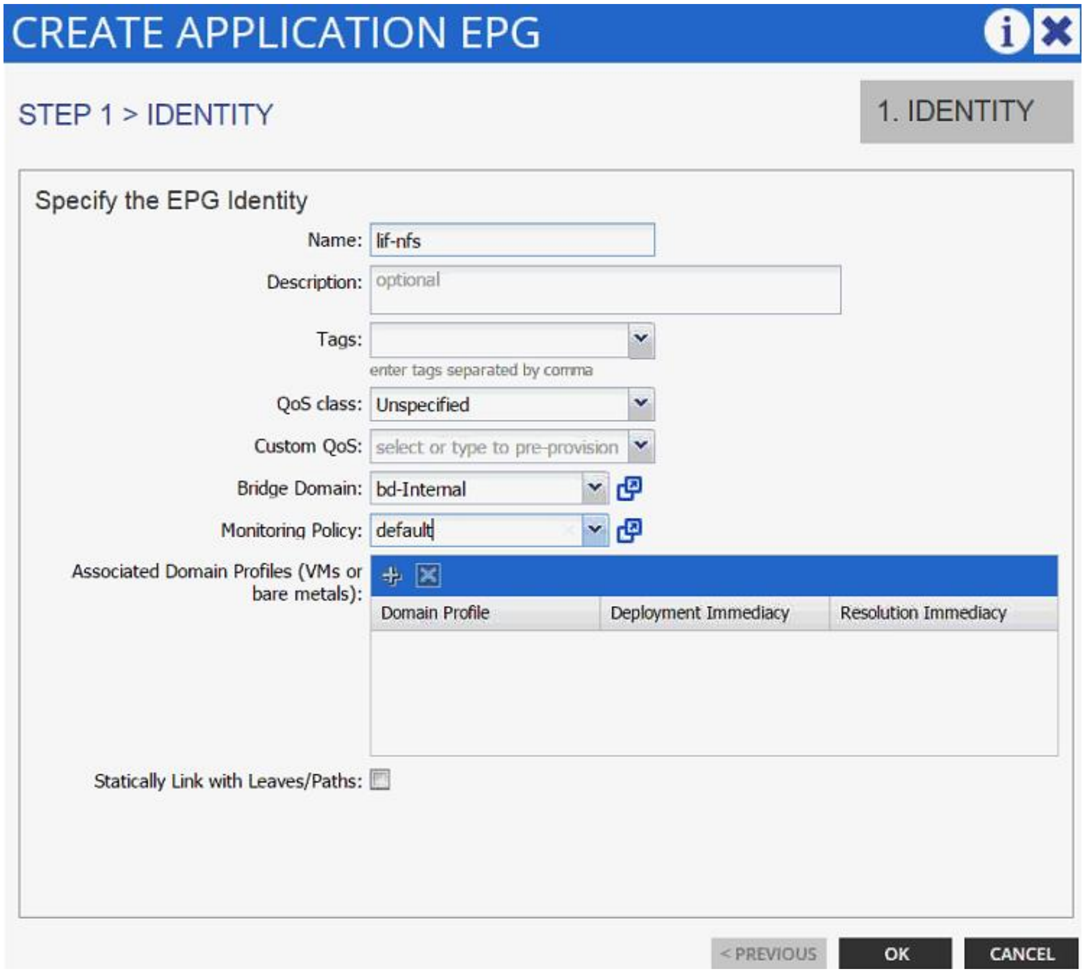

Name: lif-nfs

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy

Statically Link with Leaves/Paths: □

21. Click OK.

22. Click SUBMIT to finish creating Application Profile.

23. Expand the newly created NFS Application profile from the menu bar on the left.

24. Expand NFS, expand Application EPGs and expand EPG lif-nfs.

25. Click Static Bindings (Paths).

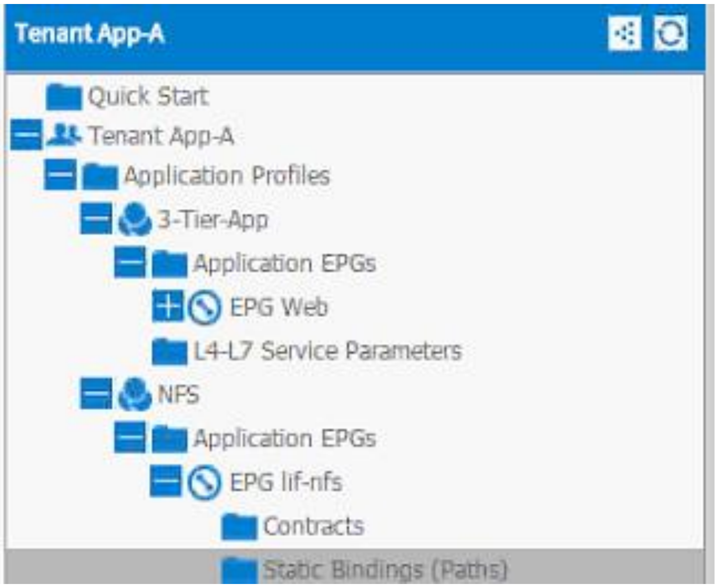

26. Click Action.

27. Click Deploy Static EPG on PC, vPC, or Interface.

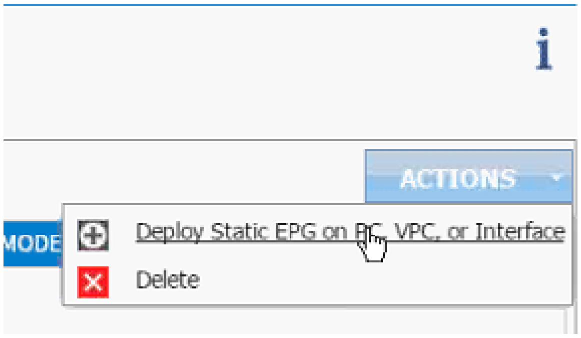

28. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

29. From the drop-down menu Path, select NetApp Controller 1.

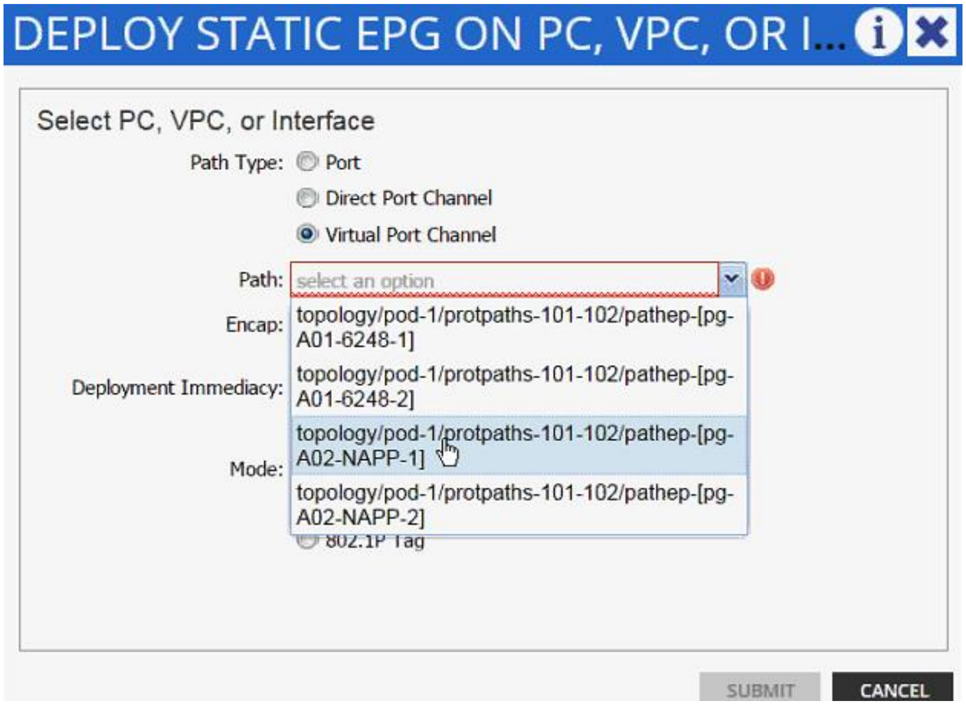

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port

○ Direct Port Channel

● Virtual Port Channel

Path: select an option

topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]

topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]

topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]

topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]

Encap:

Deployment Immediacy:

Mode:

○ 802.1P Tag

SUBMIT CANCEL

30. Enter vlan-<App-A-NFS LIF VLAN> for Encap (VLAN 3180 is the NFS VLAN on NetApp Controller in the screen capture below).

31. Change Deployment Immediacy to Immediate.

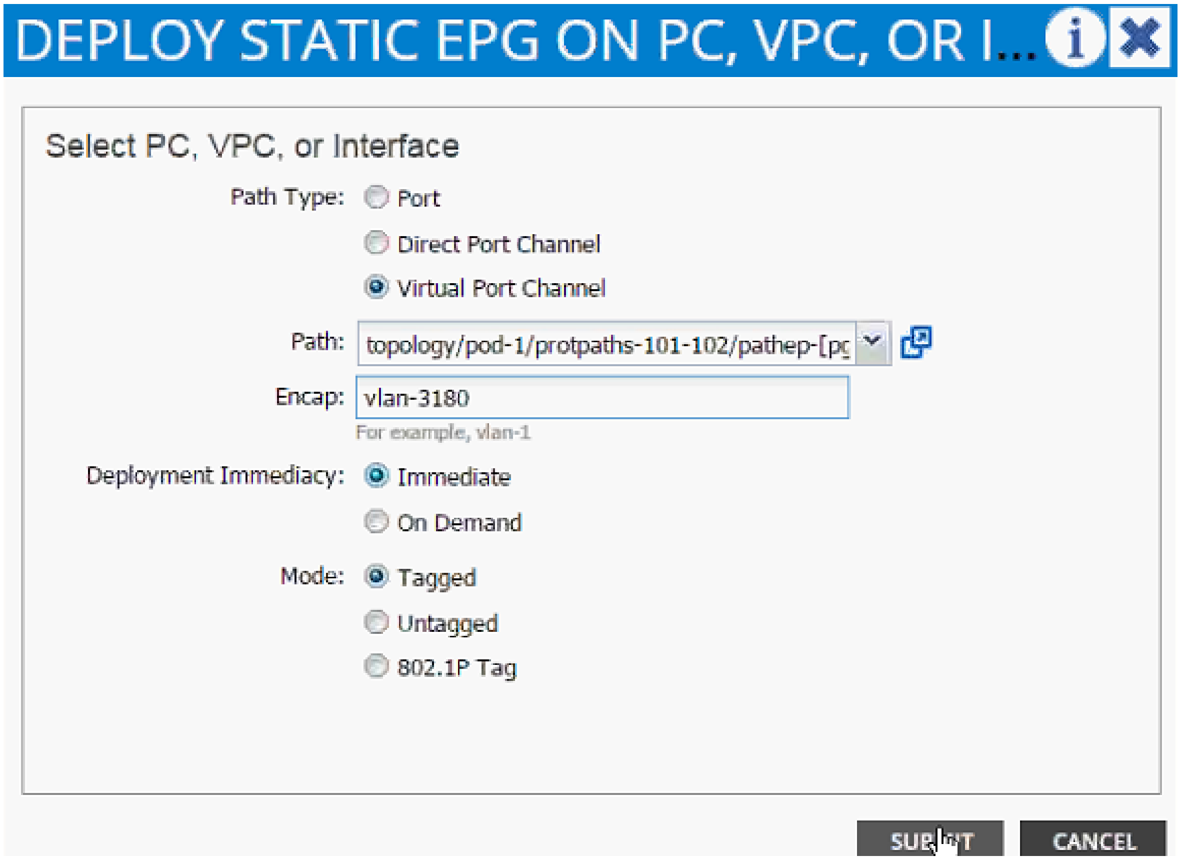

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[pc

Encap: vlan-3180

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

32. Click Submit.

33. Repeat these steps for mapping NetApp Controller 2 path.

34. Static bindings should be similar to screenshot below.

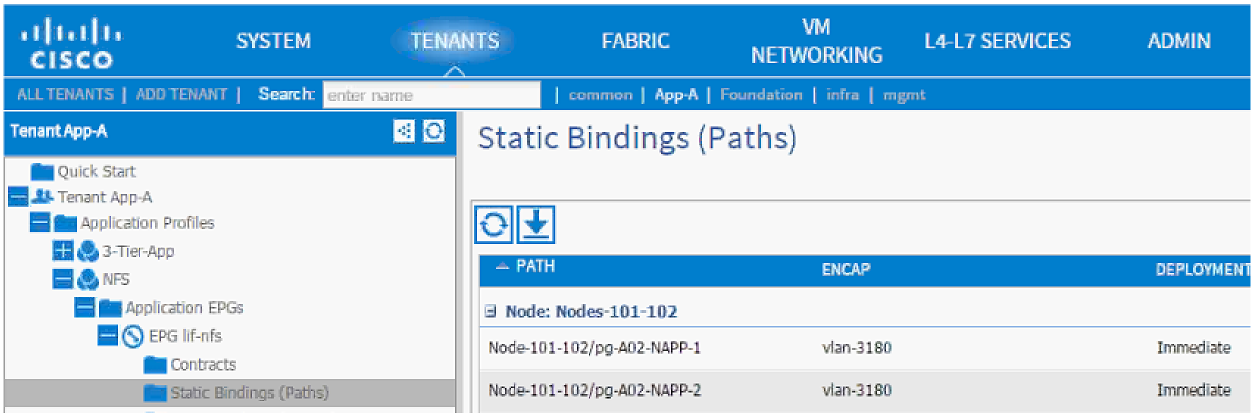

Static Bindings (Paths)

PATH	ENCAP	DEPLOYMENT
Node: Nodes-101-102		
Node-101-102/pg-A02-NAPP-1	vlan-3180	Immediate
Node-101-102/pg-A02-NAPP-2	vlan-3180	Immediate

35. Click Contracts under the EPG lif-nfs.

36. Click Action and select Add Provided Contract.

37. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

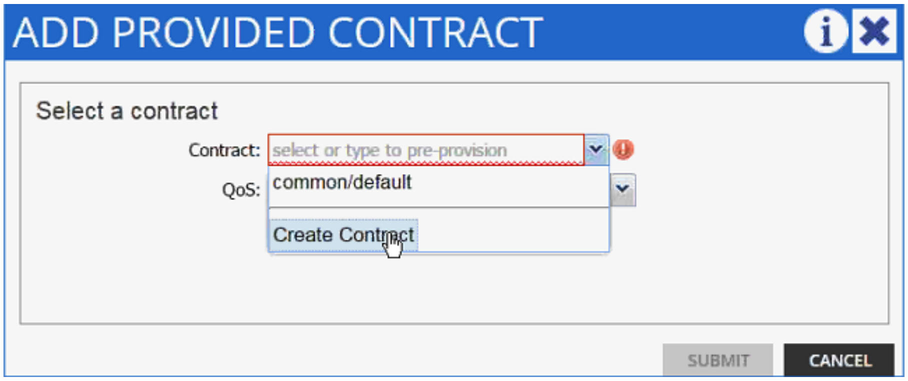

38. Enter Allow-NFS as Name in the CREATE CONTRACT dialog box.

39. Click + next to Subjects to add a new contract subject.

40. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

41. Click + under Filter Chain to add a new filter.

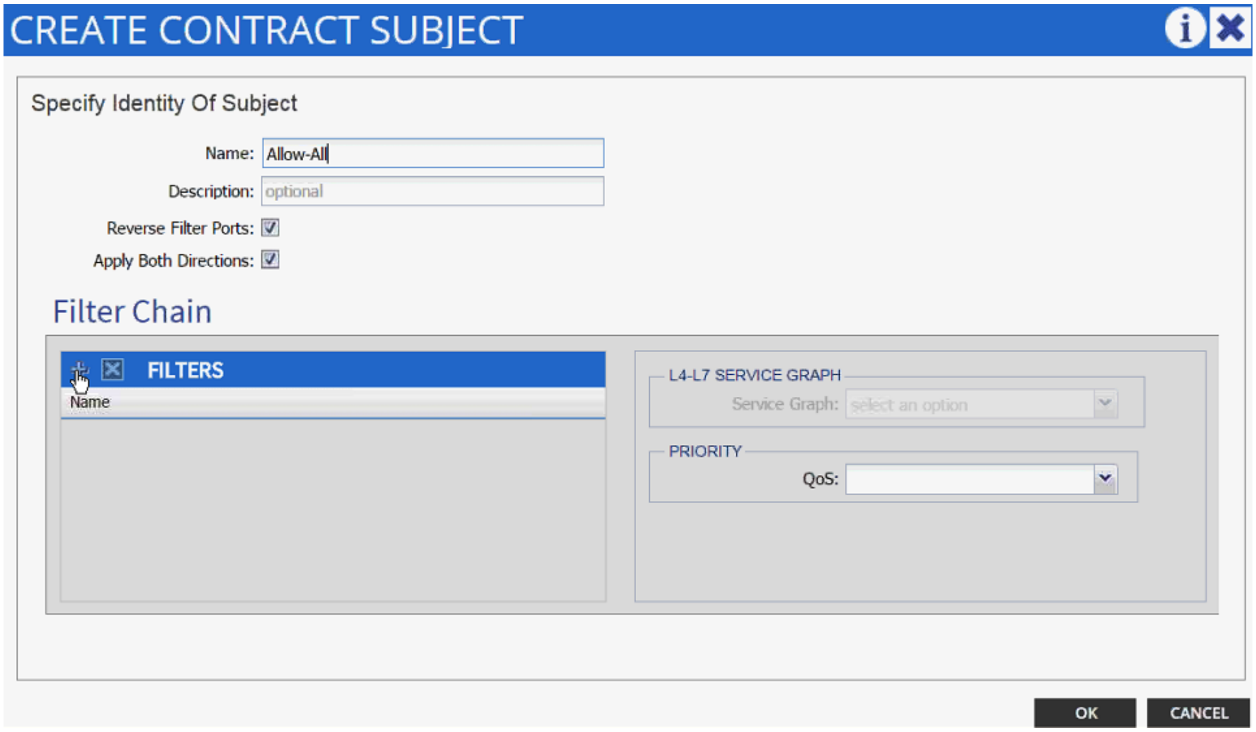

Name: Allow-All

Reverse Filter Ports: ☑

Apply Both Directions: ☑

42. Drop down the menu under FILTERS and click +.

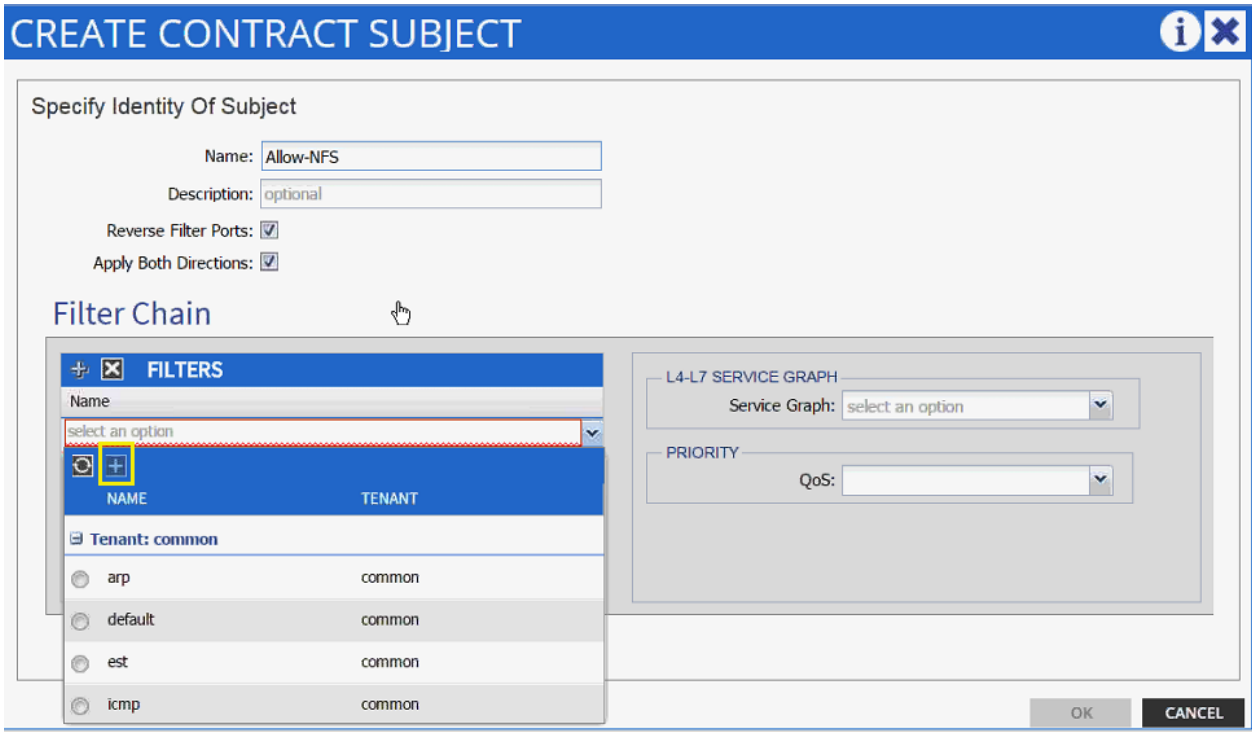

Specify Identity Of Subject

Name: Allow-NFS

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Name: select an option

Service Graph: select an option

QoS:

NAME	TENANT
Tenant: common	
○ arp	common
○ default	common
○ est	common
○ icmp	common

43. In the CREATE FILTER dialog box, enter Allow-All as the Name. In this example, we will allow all the traffic for this contract.

44. Click + to add a filter.

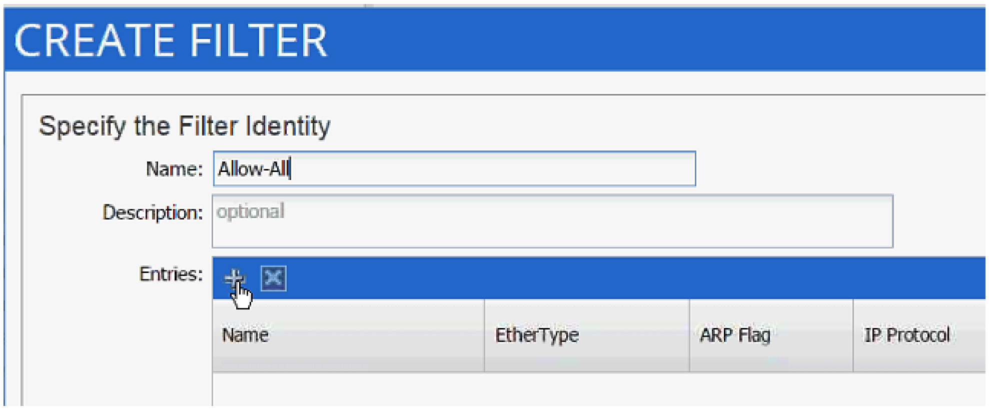

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol

45. Enter Allow-All as the name of the filter.

46. From drop-down menu, select IP as Ethertype.

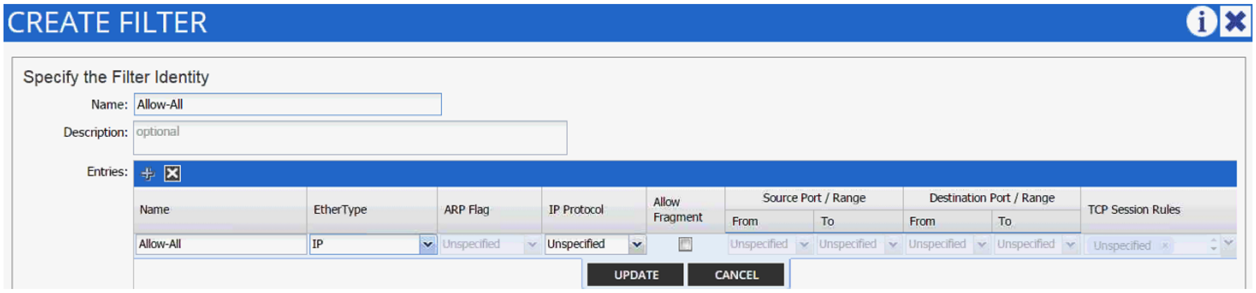

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol	Allow Fragment	Source Port / Range		Destination Port / Range		TCP Session Rules
					From	To	From	To	
Allow-All	IP	Unspecified	Unspecified	□	Unspecified	Unspecified	Unspecified	Unspecified	Unspecified

47. Click Update.

48. Click SUBMIT to create the filter.

49. Click UPDATE to add the newly created filter to the filter chain.

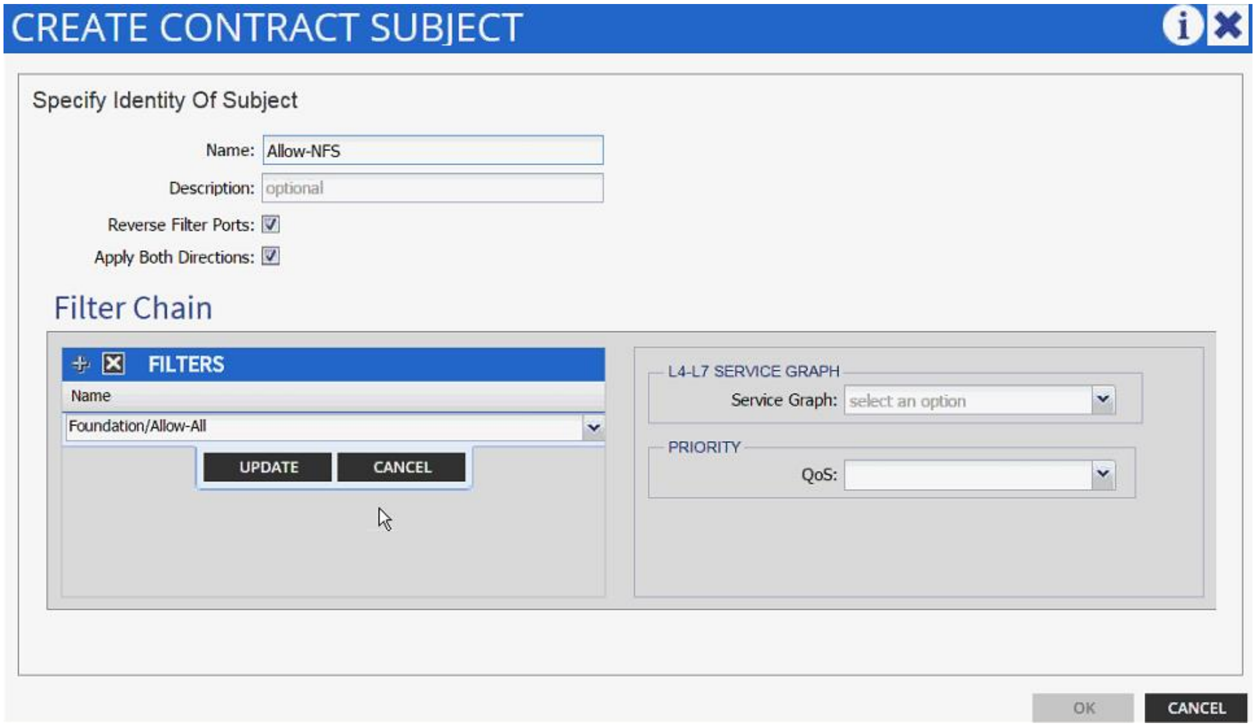

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-NFS

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

Name
Foundation/Allow-All

50. Click OK to finish creating the Contract Subject.

51. Click SUBMIT.

52. Click SUBMIT again to finish adding a provided contract.

53. Verify the Provided Contract appears under the Contracts.

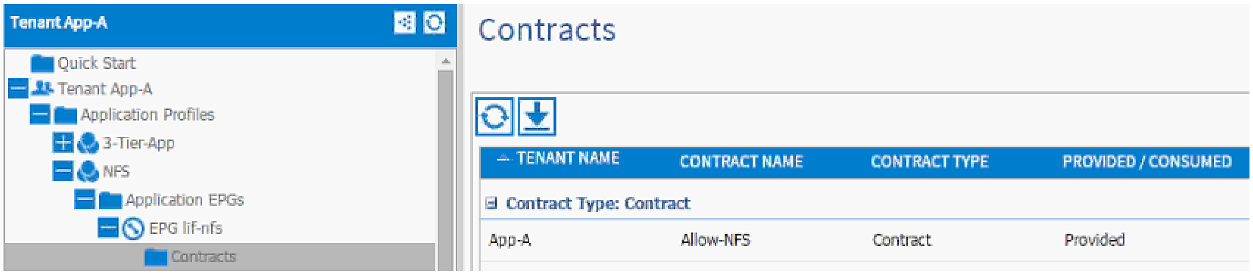

Contracts

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED
Contract Type: Contract			
App-A	Allow-NFS	Contract	Provided

54. Expand EPG vmk-nfs.

55. Click Contracts in the left menu.

56. Click ACTIONS on the right and select Add Consumed Contract.

57. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select App-A/Allow-NFS contract.

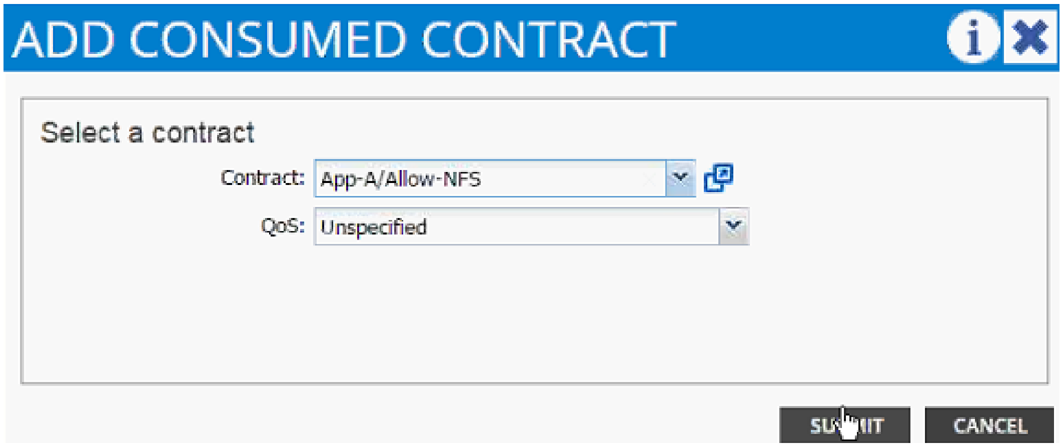

ADD CONSUMED CONTRACT

Select a contract

Contract: App-A/Allow-NFS

QoS: Unspecified

58. Click SUBMIT.

iSCSI Application Profile Creation (Optional)

If an application requires access to block based storage using iSCSI, steps from previous sections can be repeated to:

- Define two Bridge Domains – bd-iSCSI-a and bd-iSCSI-b
- Define an Application Profile called iSCSI
- Define four EPGs called lif-iSCSI-a, lif-iSCSI-b, vmk-iSCSI-a and vmk-iSCSI-b
- Use separate bridge domains for iSCSI-a and iSCSI-b EPGs
- Attach the vmk specific EPGs to the VMM domain
- Attach the lif specific EPGs to static VLAN path mappings
- Define the contracts to enable communication between iSCSI-a and iSCSI-b vmk and lif EPGs
- Define VMkernel ports on the VMware vDS for iSCSI-a and iSCSI-b.
- Add the storage SVM targets in the VMware iSCSI software initiator.

Defining NFS VMKernel Port and Mounting the Datastore

In the previous section, NFS communication between the Application SVM and ESXi servers was set up. In this section, a VMkernel port will be defined on ESXi servers and NFS datastore will be mounted to host application specific VMs.

1. Select ESXi server in the vSphere Client. Click Configuration and then Networking.

2. Click vSphere Distributed Switch.

3. Click Manage Virtual Adapters.

4. Click Add to add VMkernel Port.

5. Select New virtual adapter. Click Next.

6. Click Next.

7. Select App-A|NFS|vnk-nfs from drop-down menu for Select port group.

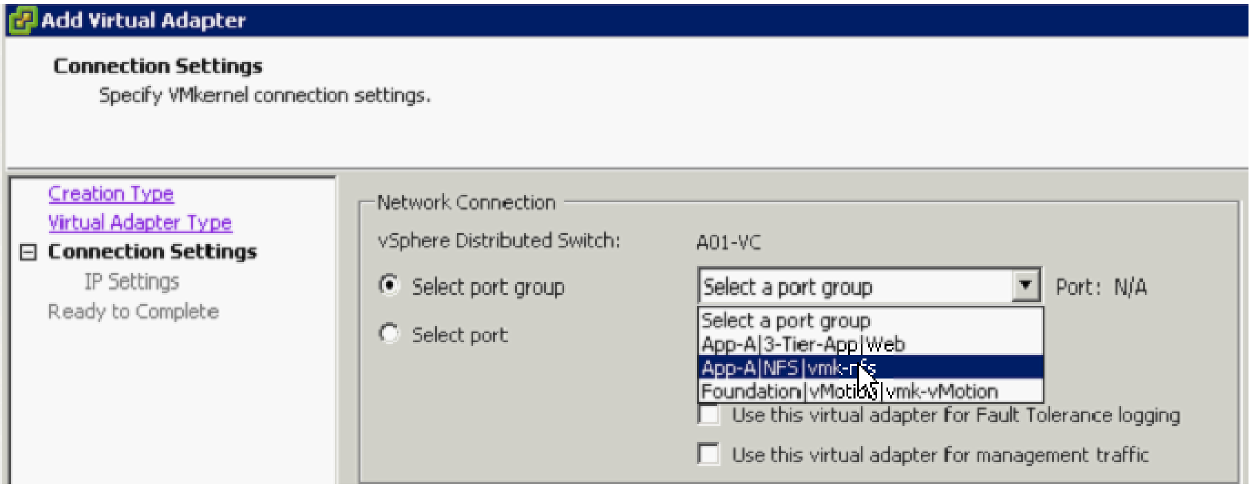

Add Virtual Adapter

Connection Settings

Specify VMkernel connection settings.

Network Connection

vSphere Distributed Switch: A01-VC

● Select port group: Select a port group ▼ Port: N/A

○ Select port

□ Use this virtual adapter for Fault Tolerance logging

□ Use this virtual adapter for management traffic

8. Click Next.

9. Enter the NFS VMkernel Port IP address in the same subnet as the NetApp LIF defined earlier.

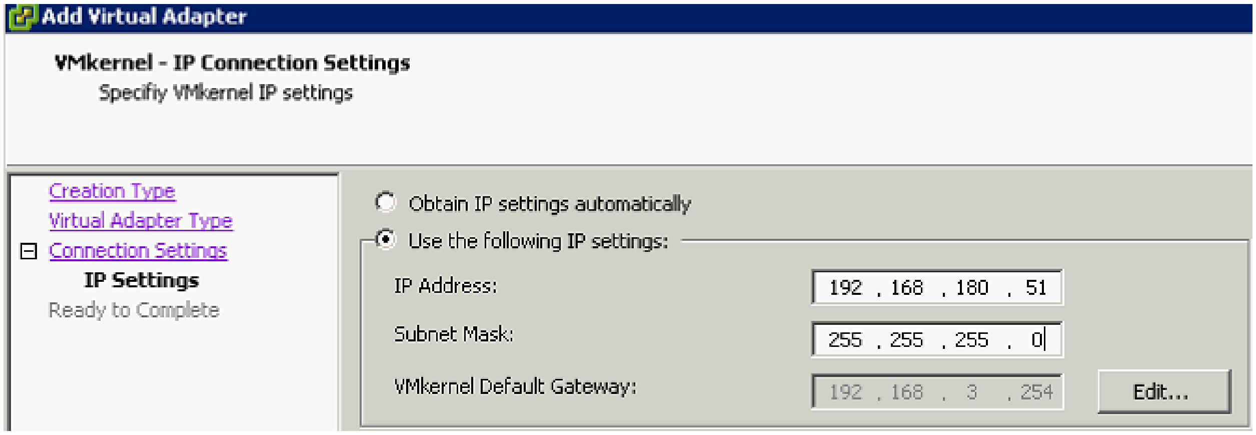

Add Virtual Adapter

VMkernel - IP Connection Settings

Specifiy VMkernel IP settings

○ Obtain IP settings automatically

● Use the following IP settings:

IP Address: 192 . 168 . 180 . 51

Subnet Mask: 255 . 255 . 255 . 0

VMkernel Default Gateway: 192 . 168 . 3 . 254

10. Click Next.

11. Click Finish.

12. Click newly created vmk port and click Edit.

13. Set the MTU to 9000.

14. Click OK.

15. Click Close to finish setting up the VMkernel port.

16. Click Storage under settings.

17. Click Add Storage.

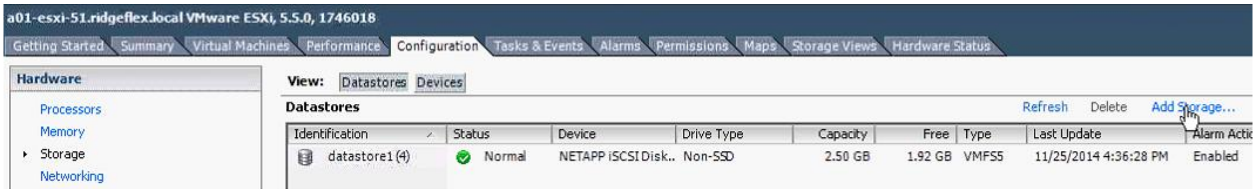

Datastores

Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update	Alarm Actic
datastore1 (4)	Normal	NETAPP iSCSI Disk..	Non-SSD	2.50 GB	1.92 GB	VMFS5	11/25/2014 4:36:28 PM	Enabled

18. Select Network File System.

19. Click Next.

20. Provide the SVM LIF address, path and datastore name for the predefined NFS datastore.

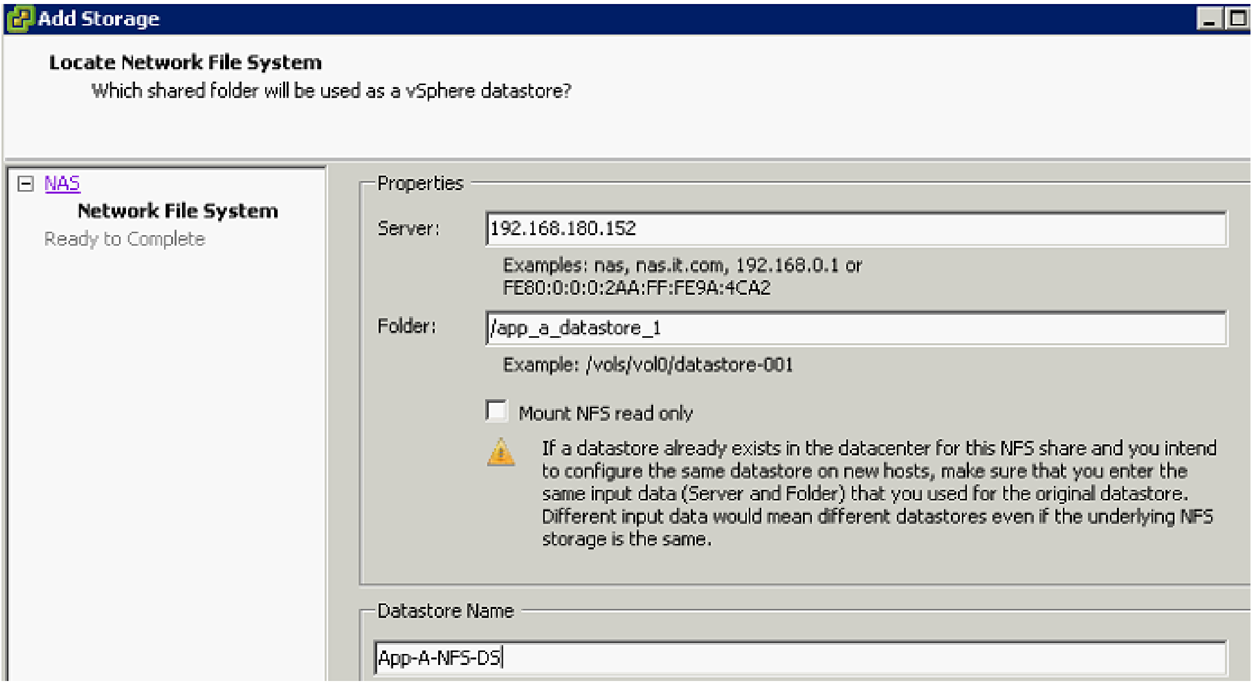

Locate Network File System

Which shared folder will be used as a vSphere datastore?

Server: 192.168.180.152

Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: /app_a_datastore_1

Example: /vols/vol0/datastore-001

☐ Mount NFS read only

If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

Datastore Name: App-A-NFS-DS

21. Click Next.

22. Verify information and click Finish.

Defining EPGs for additional Application Tiers (Optional)

In the previous section, an application profile 3-Tier-App was defined and an EPG Web was deployed. To define EPGs for additional tiers of the application, complete the following steps:

1. Expand Application Profiles for Tenant App-A.

2. Right-click 3-Tier-App and select Create Application EPG.

3. Provide the name of the Application Tier EPG (App in this example).

4. From the drop-down menu, select bd-Internal as the Bridge Domain.

5. Select Monitoring Policy as default.

6. Click + to add Associated Domain Profiles (VMs or bare metals) and select vCenter domain.

7. Change Deployment Immediacy and Resolution Immediacy to Immediate.

8. Click UPDATE.

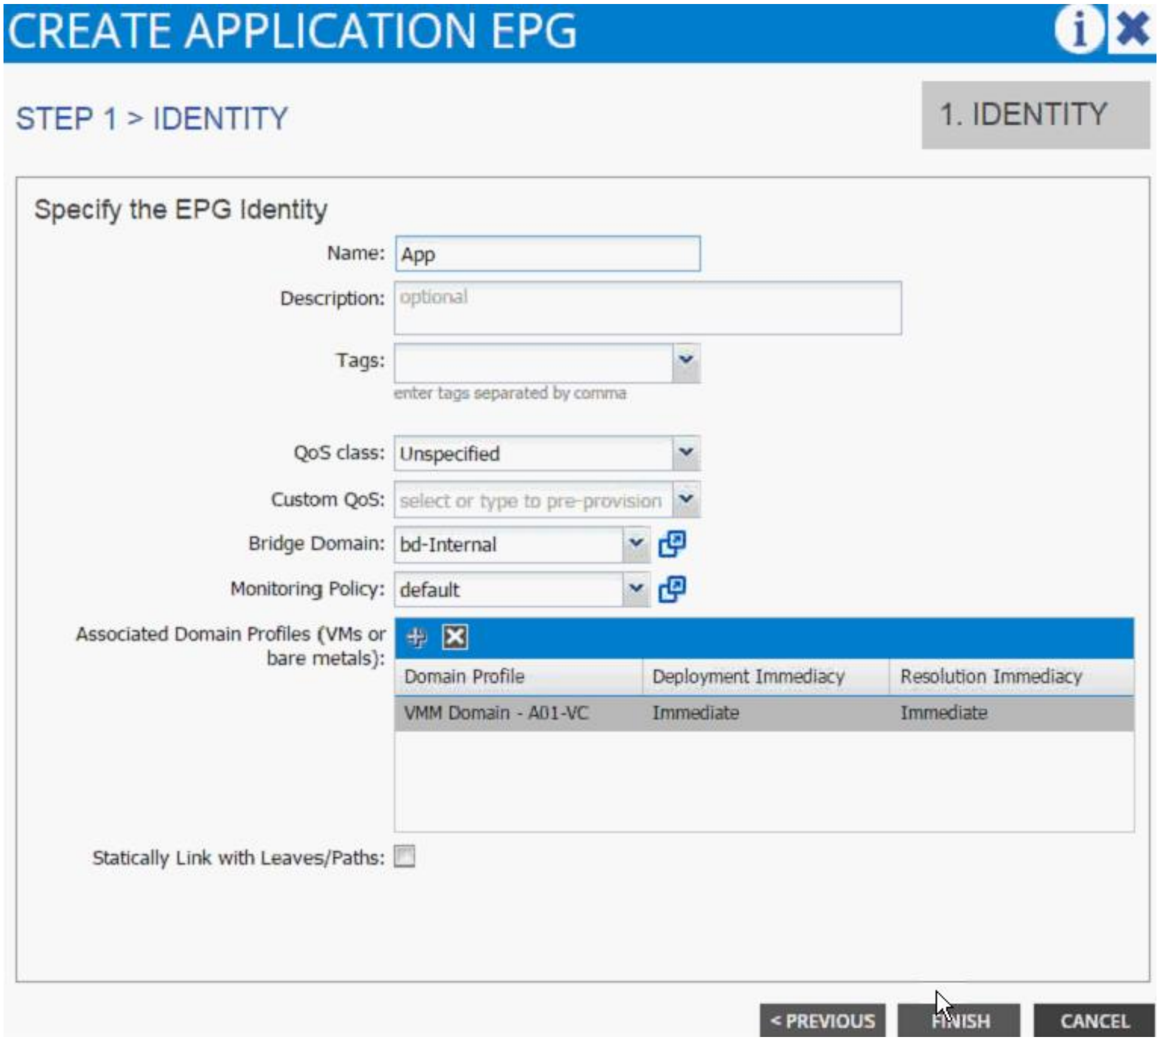

CREATE APPLICATION EPG

STEP 1 > IDENTITY

1. IDENTITY

Specify the EPG Identity

Name: App

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
VMM Domain - A01-VC	Immediate	Immediate

Statically Link with Leaves/Paths: □

9. Click FINISH.

10. Expand the newly created EPG App and click Subnets.

11. Click the Action menu on the right and select Create EPG Subnet.

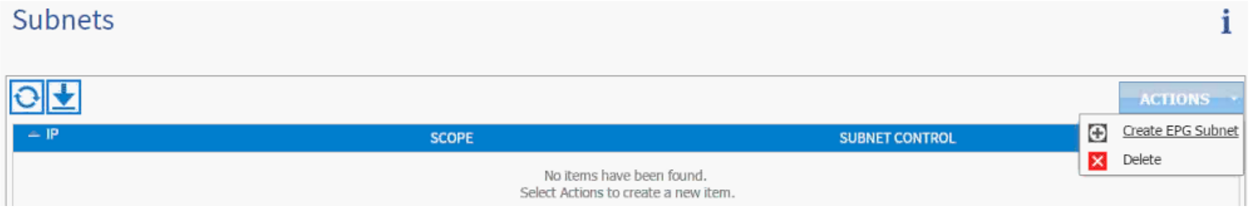

No items have been found.
Select Actions to create a new item.

12. Enter 10.10.2.254/24 for the Default Gateway IP. This IP address is the gateway that all the App VMs will use. (Please adjust this subnet according to your implementation).

13. Change scope to only Shared Subnet.

14. Click FINISH.

15. Repeat these steps for additional EPG (application tier) definitions.

Enabling Communication Between Application Tiers

If an application tier needs to communicate to another application tier, a contract needs to be provided by one EPG and consumed by the other. In this example, previously defined App EPG will provide a contract and Web EPG will consume the contract. The ports on which the two application tiers can be limited in the contract subject but for this example, all communication will be allowed between the two tiers.

To enable communication between application tiers, complete the following steps:

1. Expand the Application Profile 3-Tier-App, Application EPGs and EPG App.

2. Click Contracts under the EPG App.

3. Click Action and select Add Provided Contract.

4. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

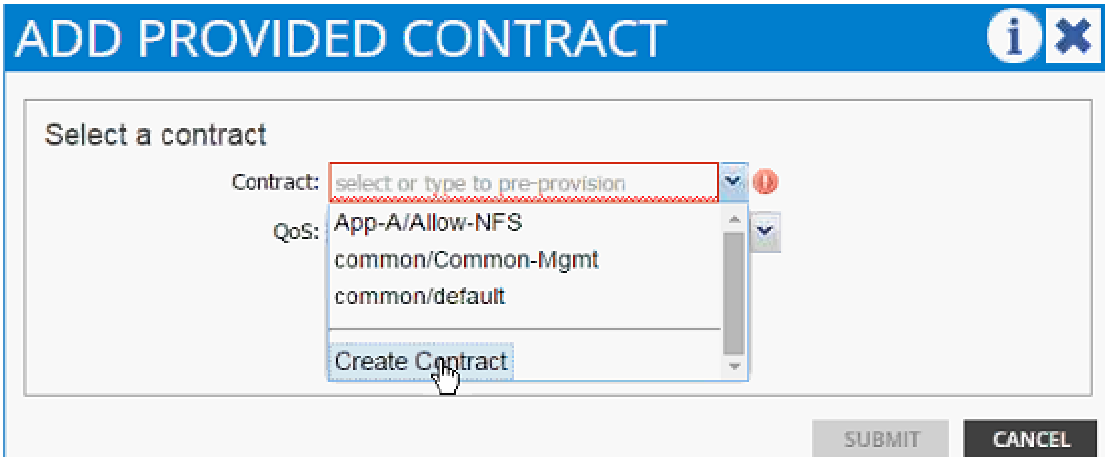

5. Enter Allow-App-Web as Name in the CREATE CONTRACT dialog box.

6. Click + next to Subjects to add a new contract subject.

7. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

8. Click + under Filter Chain to add a new filter.

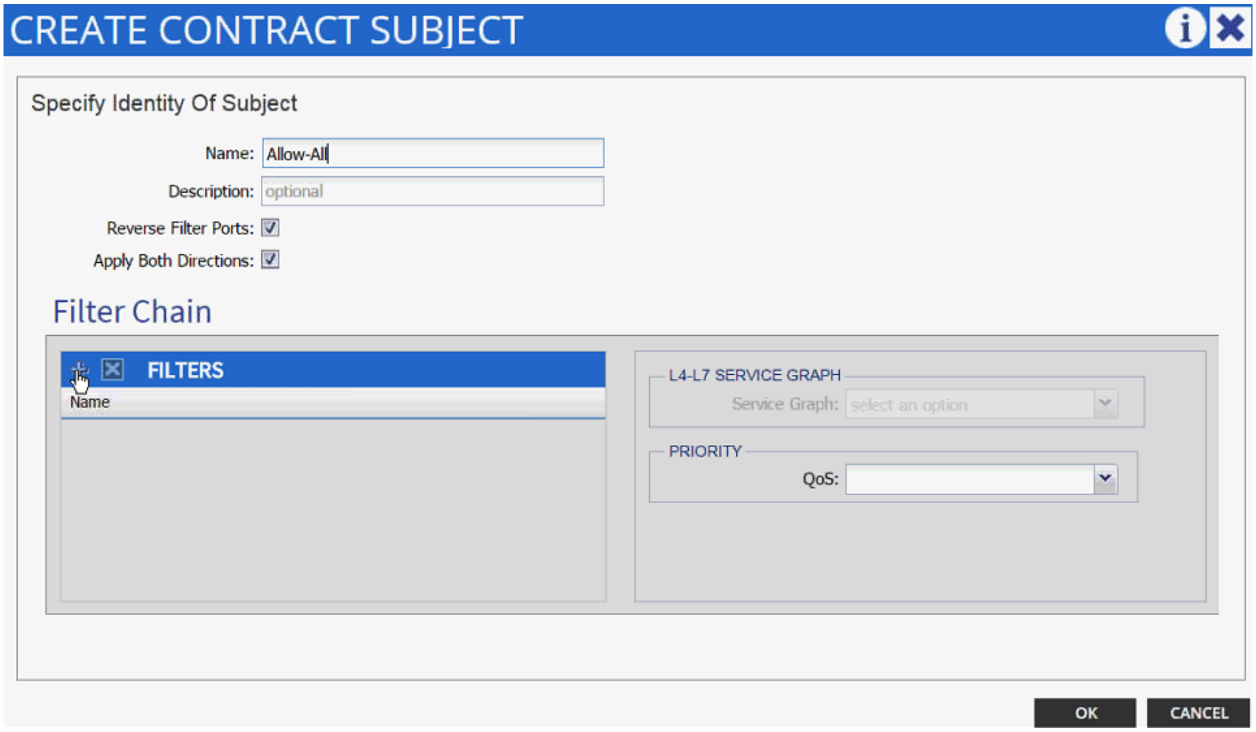

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

9. Drop down the menu under FILTERS and click +.

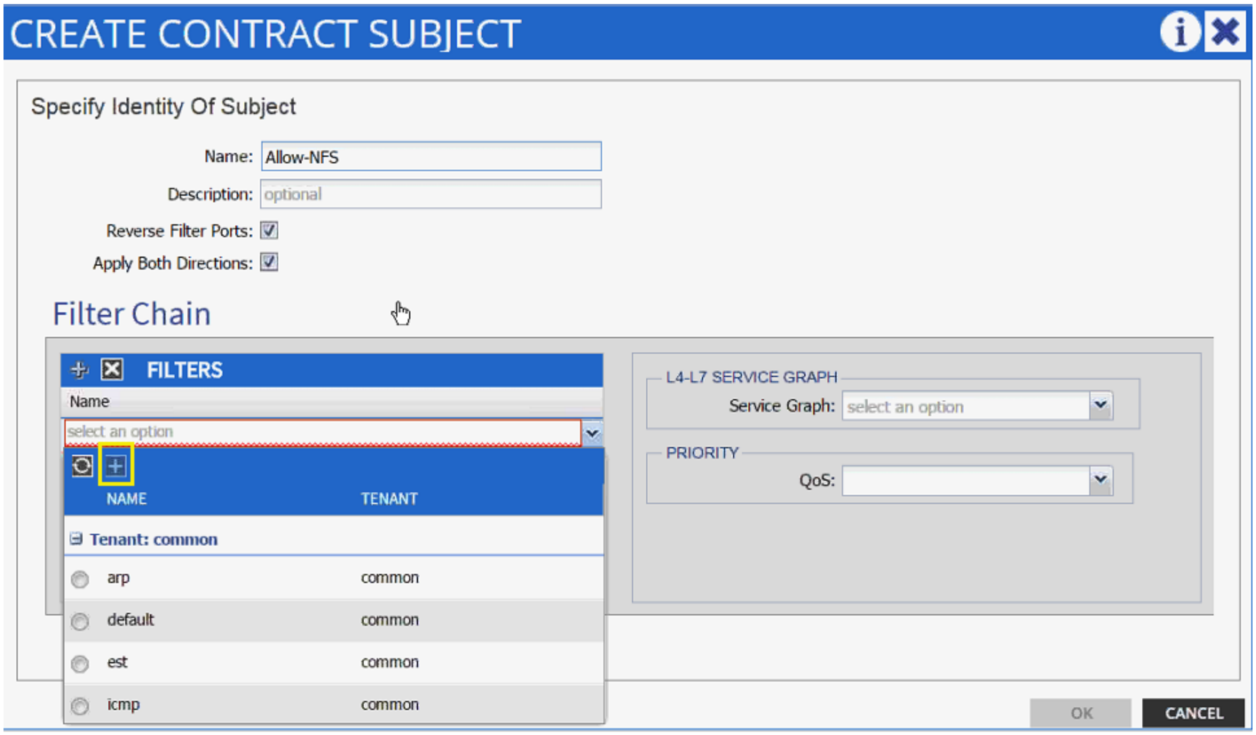

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-NFS

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

NAME	TENANT
Tenant: common	
○ arp	common
○ default	common
○ est	common
○ icmp	common

10. In the CREATE FILTER dialog box, enter Allow-All as the Name. In this example, we will allow all the traffic for this contract.

11. Click + to add a filter.

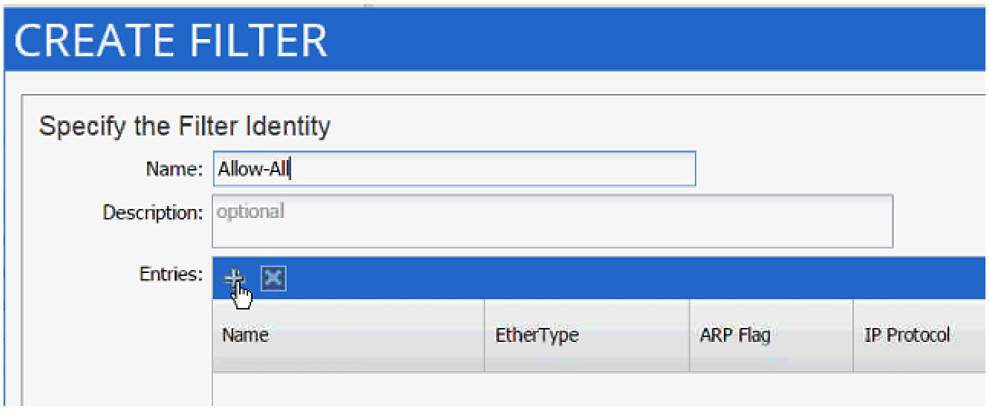

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol

12. Enter Allow-All as the name of the filter.

13. From drop-down menu, select IP as Ethertype.

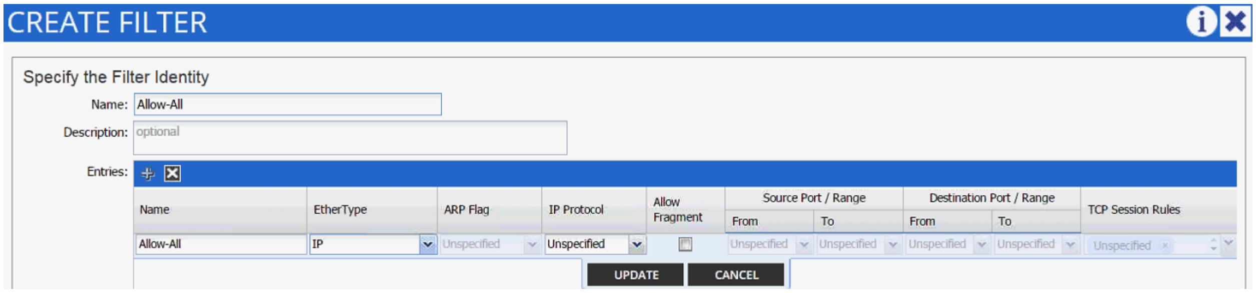

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol	Allow Fragment	Source Port / Range		Destination Port / Range		TCP Session Rules
					From	To	From	To	
Allow-All	IP	Unspecified	Unspecified	□	Unspecified	Unspecified	Unspecified	Unspecified	Unspecified

14. Click Update.

15. Click SUBMIT to create the filter.

16. Click UPDATE to add the newly created filter to the filter chain.

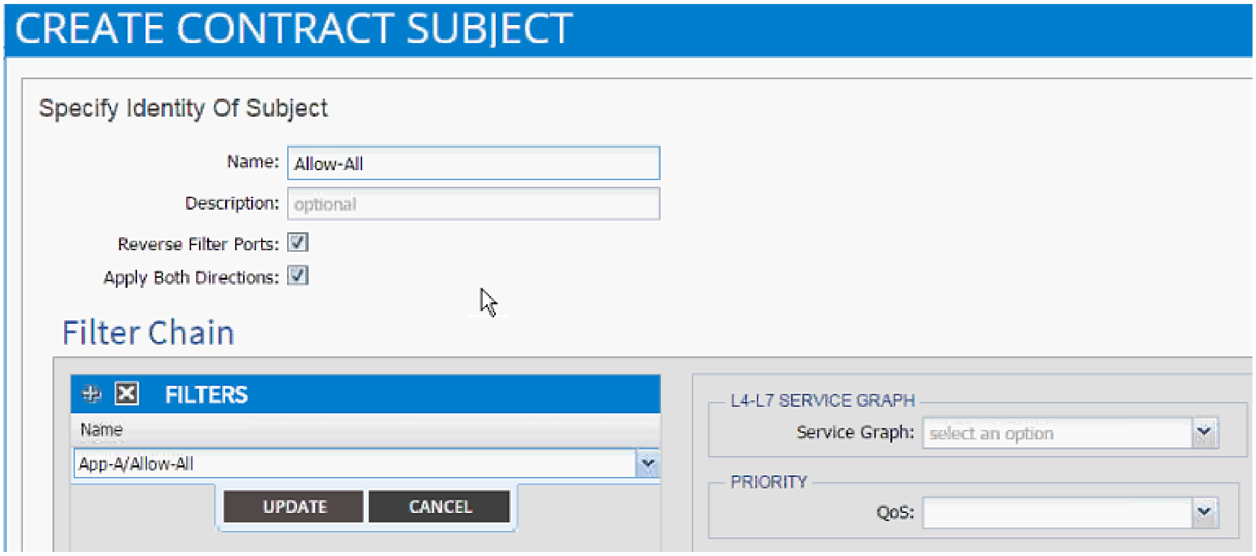

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS

Name
App-A/Allow-All

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

17. Click OK to finish creating the Contract Subject.

18. Click SUBMIT.

19. Click SUBMIT again to finish adding a provided contract.

20. Expand EPG Web.

21. Click Contracts in the left menu.

22. Click ACTIONS on the right and select Add Consumed Contract.

23. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select App-A/Allow-App-Web contract.

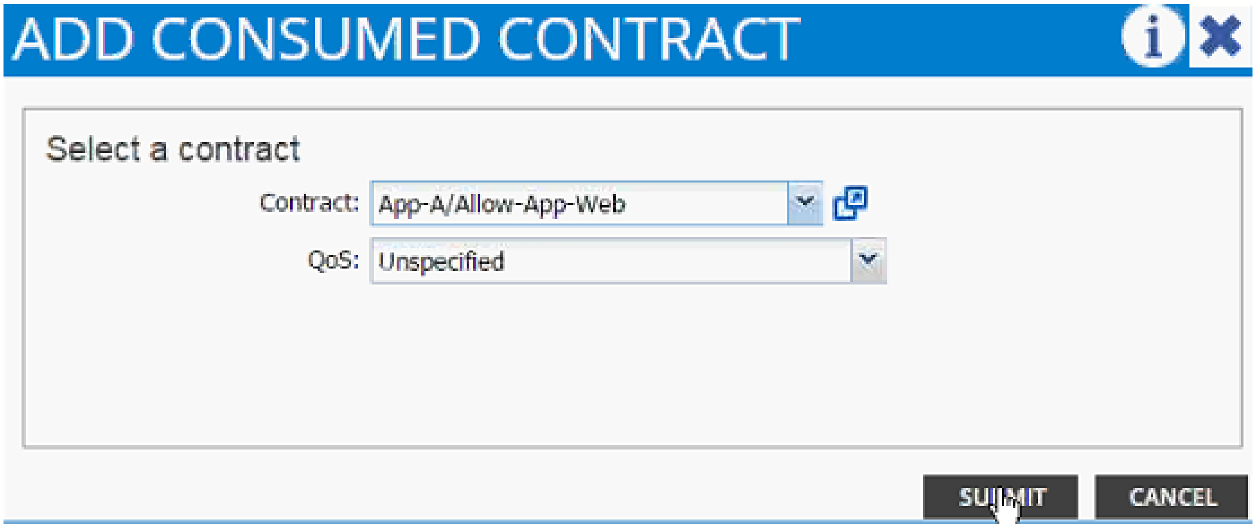

Select a contract

Contract: App-A/Allow-App-Web

QoS: Unspecified

24. Click SUBMIT.

Deploying Virtual Machines

After going through several configuration steps above, application related VMs could be deployed using vSphere client. VMware distributed switched managed using APIC will show port-groups for the Application EPGs defined in previous steps. The Web-tier VMs will be deployed and connected to the port-group labeled App-a|3-Tier-App-Web. VM related to other application tiers will be deployed in their respective port-groups. The resulting VDS configuration for an ESXi host will look like Figure 5. As the name suggests, App-A-Web is the web-VM while App-A-App is the application VM.

Figure 5 ESXi – VDS Deployment Example for Application VMs

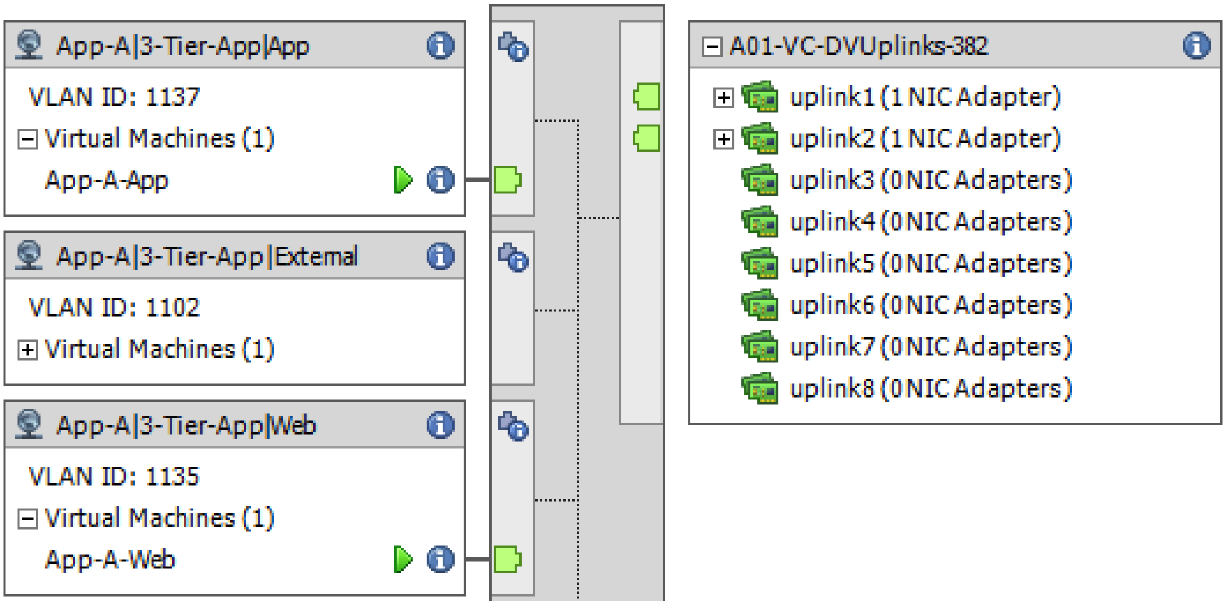

App-A|3-Tier-App|App

VLAN ID: 1137

⊟ Virtual Machines (1)

App-A-App

App-A|3-Tier-App|External

VLAN ID: 1102

⊞ Virtual Machines (1)

App-A|3-Tier-App|Web

VLAN ID: 1135

⊟ Virtual Machines (1)

App-A-Web

⊟ A01-VC-DVUplinks-382

⊞ uplink1 (1 NIC Adapter)

⊞ uplink2 (1 NIC Adapter)

uplink3 (0 NIC Adapters)

uplink4 (0 NIC Adapters)

uplink5 (0 NIC Adapters)

uplink6 (0 NIC Adapters)

uplink7 (0 NIC Adapters)

uplink8 (0 NIC Adapters)

Cisco ACI – Accessing Common Services

This section provides a detailed procedure for configuring access to common services such as AD, DNS and in some cases vCenter etc. The leaf switches in the ACI fabric connect to an existing management switch where common service servers or VMs are connected using a dedicated LAN segment (VLAN 3177 in this case). This management segment is then mapped to an EPG in the common tenant using static VLAN mappings. In ACI, any contracts defined and provided in the common tenant can be consumed in any all other tenants and this makes common tenant an ideal candidate to host the common services management EPGs.

Configuring ACI for Management Switch

To configure vPC for management switch, complete the following steps:

1. From the main menu, click FABRIC and select Access Policies.

2. Right-click Interface Policies and select Configure interface, PC and vPC.

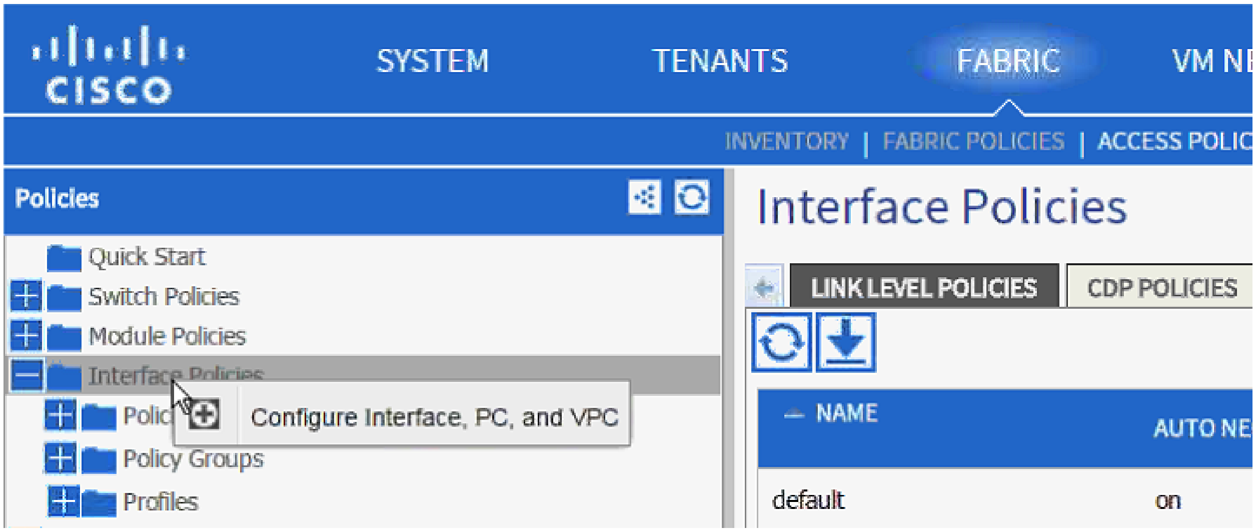

3. In the dialog box, click + under the CONFIGURED SWITCH INTERFACES.

4. Drop down the menu next to Switches and select both leaves.

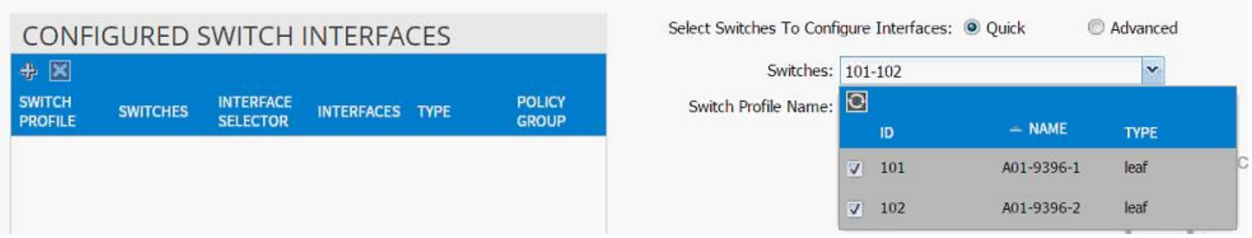

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name:

	ID	NAME	TYPE
☑	101	A01-9396-1	leaf
☑	102	A01-9396-2	leaf

5. Enter <sp-OOB-Mgmt> as Switch Profile Name. OOB-Mgmt is the host name for management switch.

6. Click the + sign to add interfaces.

7. Select vPC radio button to configure vPC.

8. Enter 1/21 under Interfaces. This is the port on both switches where existing management switch is connected.

9. Enter <ifs- OOB-Mgmt> as Interface Selector Name.

10. Drop down vPC Policy Group and click Create vPC Interface Policy Group. A new dialog box will appear.

11. Enter <pg- OOB-Mgmt> as the name of the vPC INTERFACE POLICY GROUP in the dialog box.

12. From the drop-down menu Link Level Policy, select Create Link Level Policy.

13. In the CREATE LINK LEVEL POLICY, enter 1_GE as the Name.

14. Select 1 Gbps as the Speed.

15. Click SUBMIT.

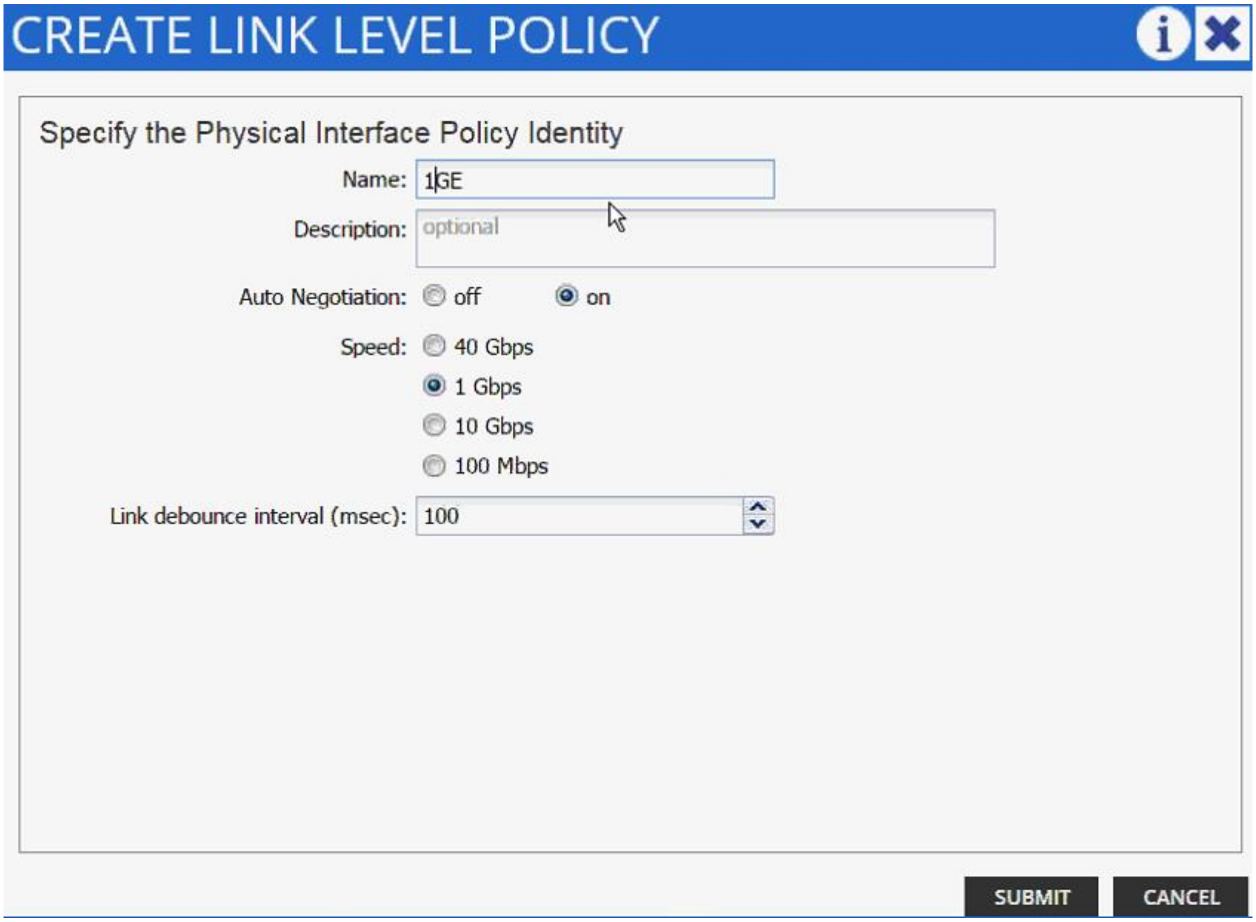

CREATE LINK LEVEL POLICY

Specify the Physical Interface Policy Identity

Name: 1GE

Description: optional

Auto Negotiation: ○ off ● on

Speed: ○ 40 Gbps ● 1 Gbps ○ 10 Gbps ○ 100 Mbps

Link debounce interval (msec): 100

16. Select CDP_Enabled as CDP Policy.

17. Select LLDP_Disabled as LLDP Policy.

18. Select default as STP Interface Policy.

19. Select LACP_ACTIVE as LACP Policy.

20. Select default as Monitoring Policy.

21. Drop down Attached Entity Profile and click Create Attachable Access Entity Profile. A new dialog box will appear.

22. Enter <aep-OOB-Mgmt> as the Name.

23. Click + to add Domain.

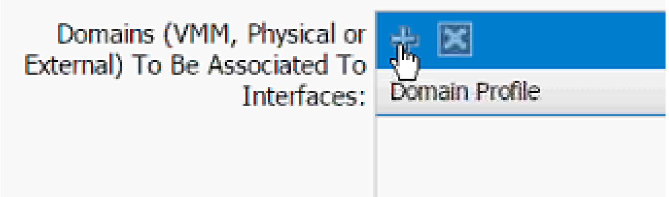

Domains (VMM, Physical or External) To Be Associated To Interfaces:

24. In the added domain, drop-down the menu and select Create Layer 2 Domain.

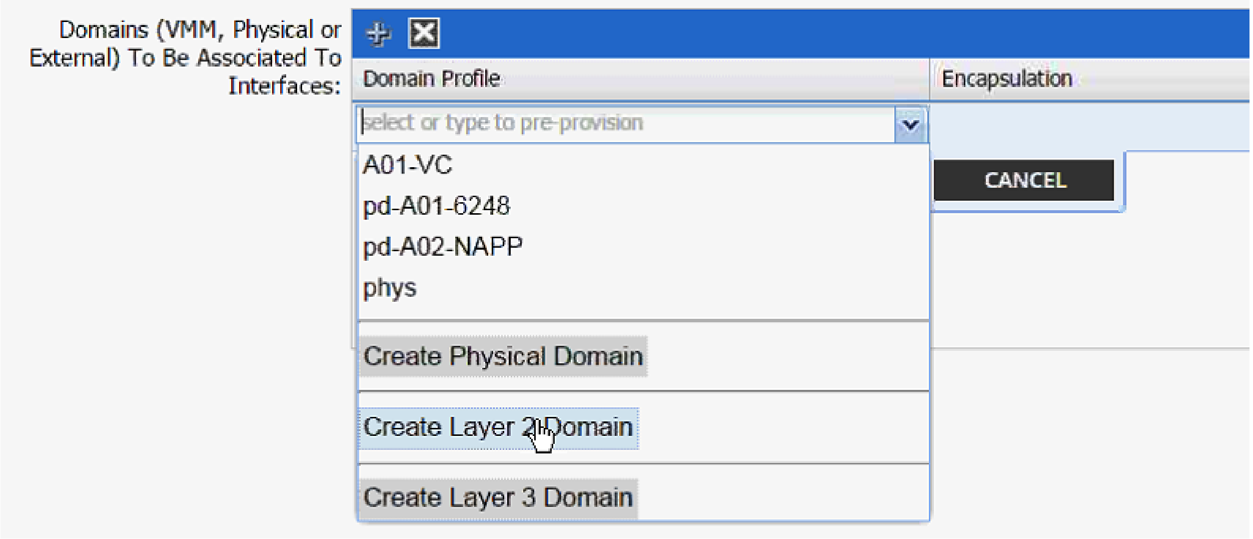

Domains (VMM, Physical or External) To Be Associated To Interfaces:

25. In the Create Layer 2 Domain dialog box, enter <L2-OOB-Mgmt> as Name.

26. Drop down VLAN Pool menu and select Create VLAN Pool.

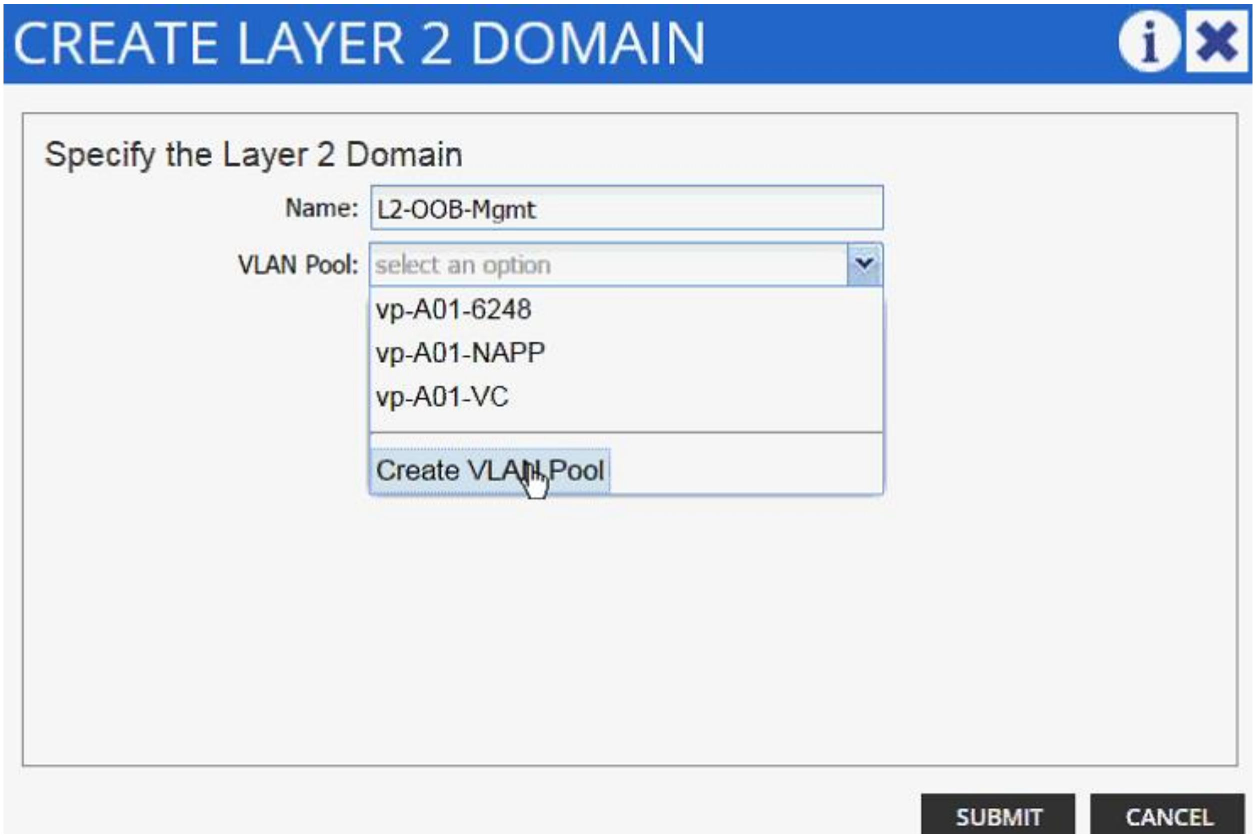

CREATE LAYER 2 DOMAIN

Specify the Layer 2 Domain

Name: L2-OOB-Mgmt

VLAN Pool: select an option

27. In the Create VLAN Pool dialog box, enter <vp-OOB-Mgmt> as Name.

28. Select Allocation Mode Static Allocation.

29. Click + next to Encap Block.

30. In the CREATE RANGES dialog box, enter the two iSCSI VLANs and the NFS VLAN.

Note: In the screenshot below, 3177 is the management VLAN utilized for common services segment.

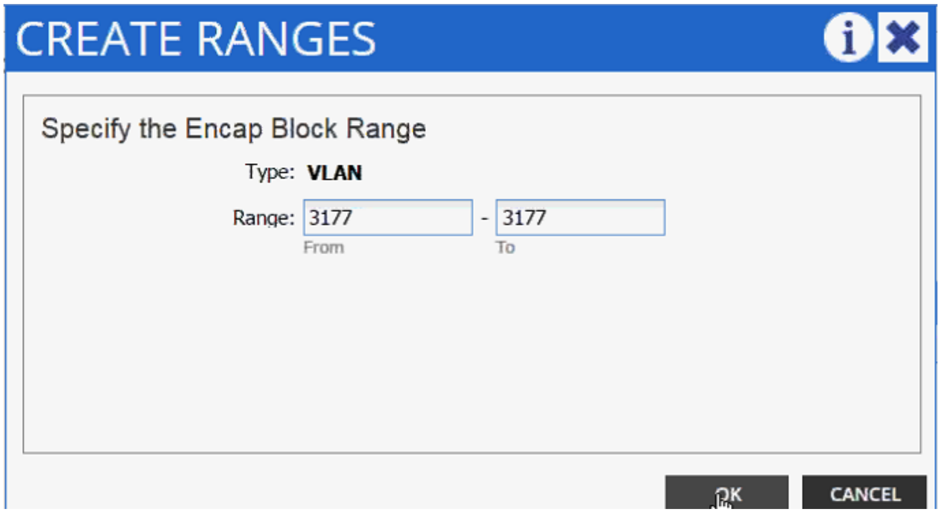

CREATE RANGES

Specify the Encap Block Range

Type: VLAN

Range: 3177 - 3177

31. Click OK.

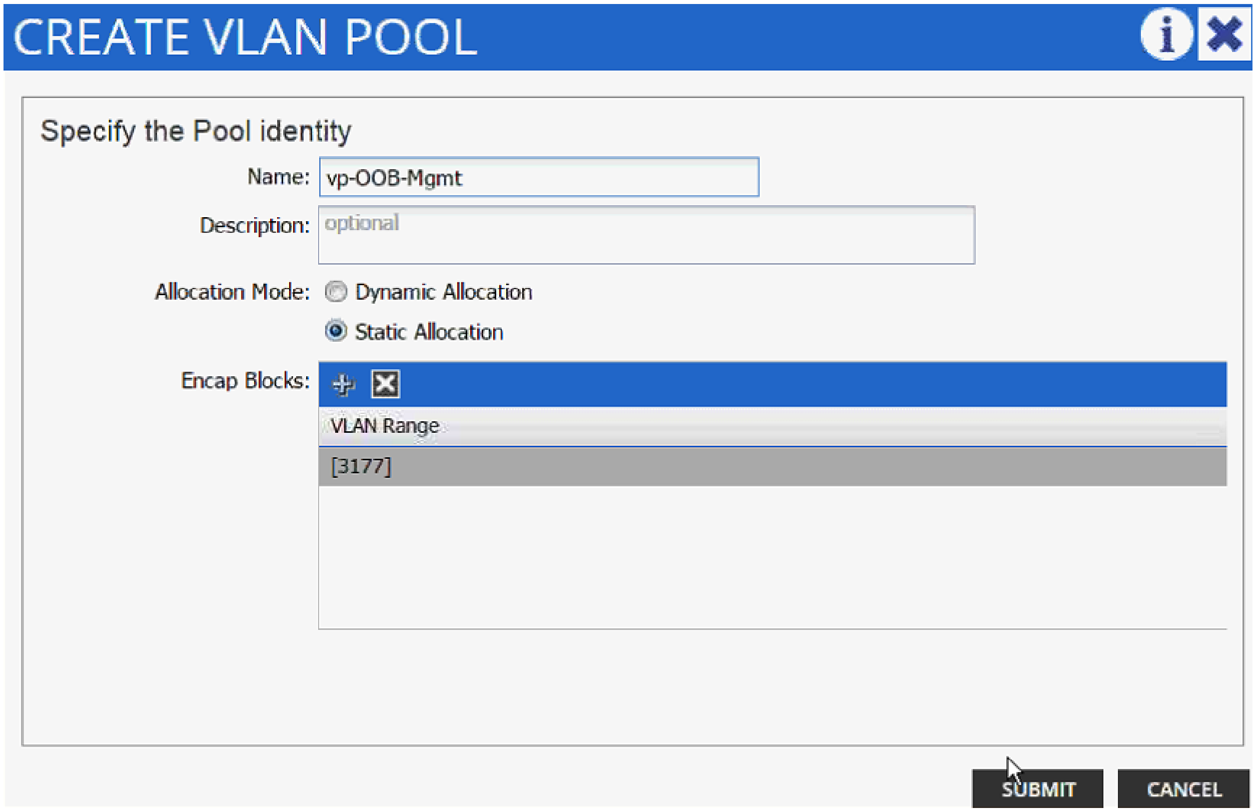

CREATE VLAN POOL

Specify the Pool identity

Name: vp-OOB-Mgmt

Description: optional

Allocation Mode: ○ Dynamic Allocation
● Static Allocation

Encap Blocks:

VLAN Range
[3177]

32. Click SUBMIT to finish VLAN pool creation.

33. Click SUBMIT to finish Layer 2 Domain creation.

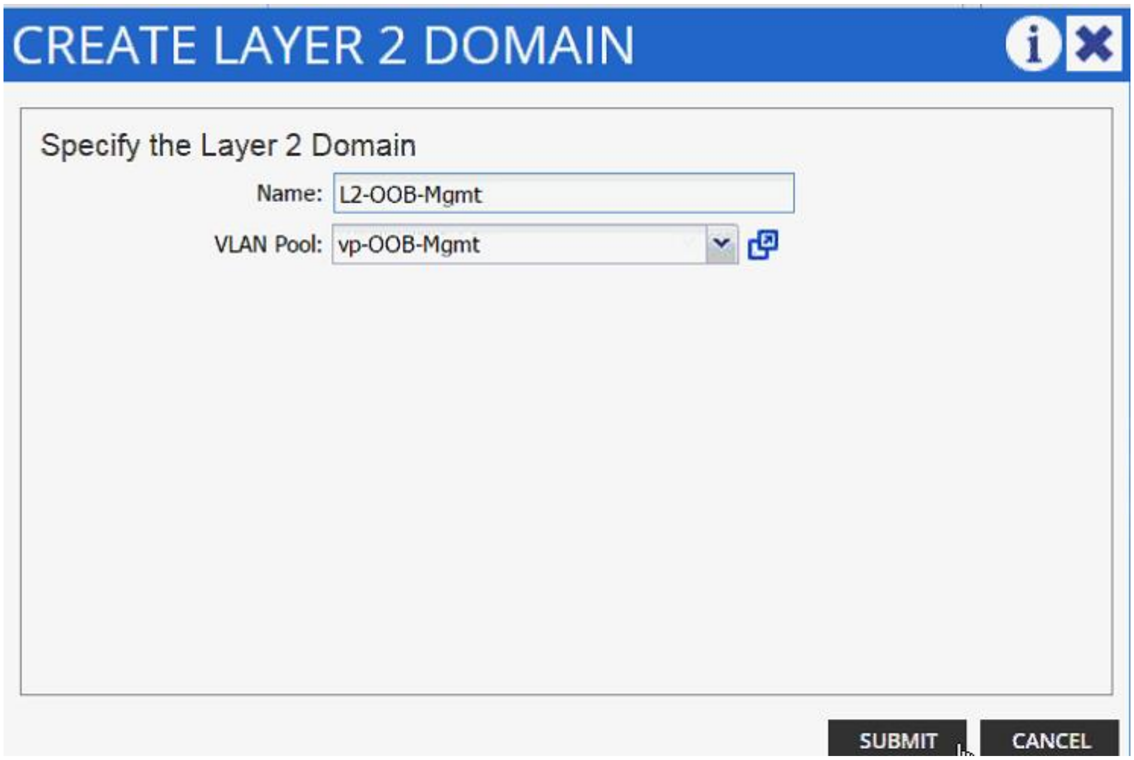

CREATE LAYER 2 DOMAIN

Specify the Layer 2 Domain

Name: L2-OOB-Mgmt

VLAN Pool: vp-OOB-Mgmt

34. Click UPDATE to finish adding Layer 2 domain to AEP.

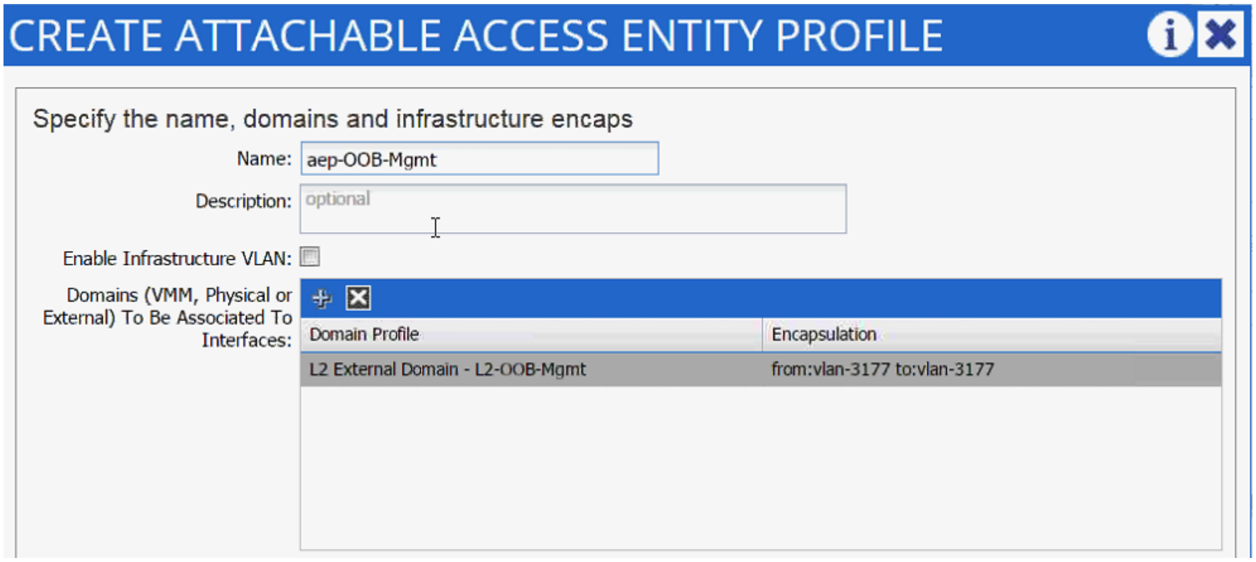

CREATE ATTACHABLE ACCESS ENTITY PROFILE

Specify the name, domains and infrastructure encaps

Name: aep-OOB-Mgmt

Description: optional

Enable Infrastructure VLAN: □

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation
L2 External Domain - L2-OOB-Mgmt	from:vlan-3177 to:vlan-3177

35. Click SUBMIT to finish adding AEP.

36. Click SUBMIT to finish creating vPC Interface Policy Group.

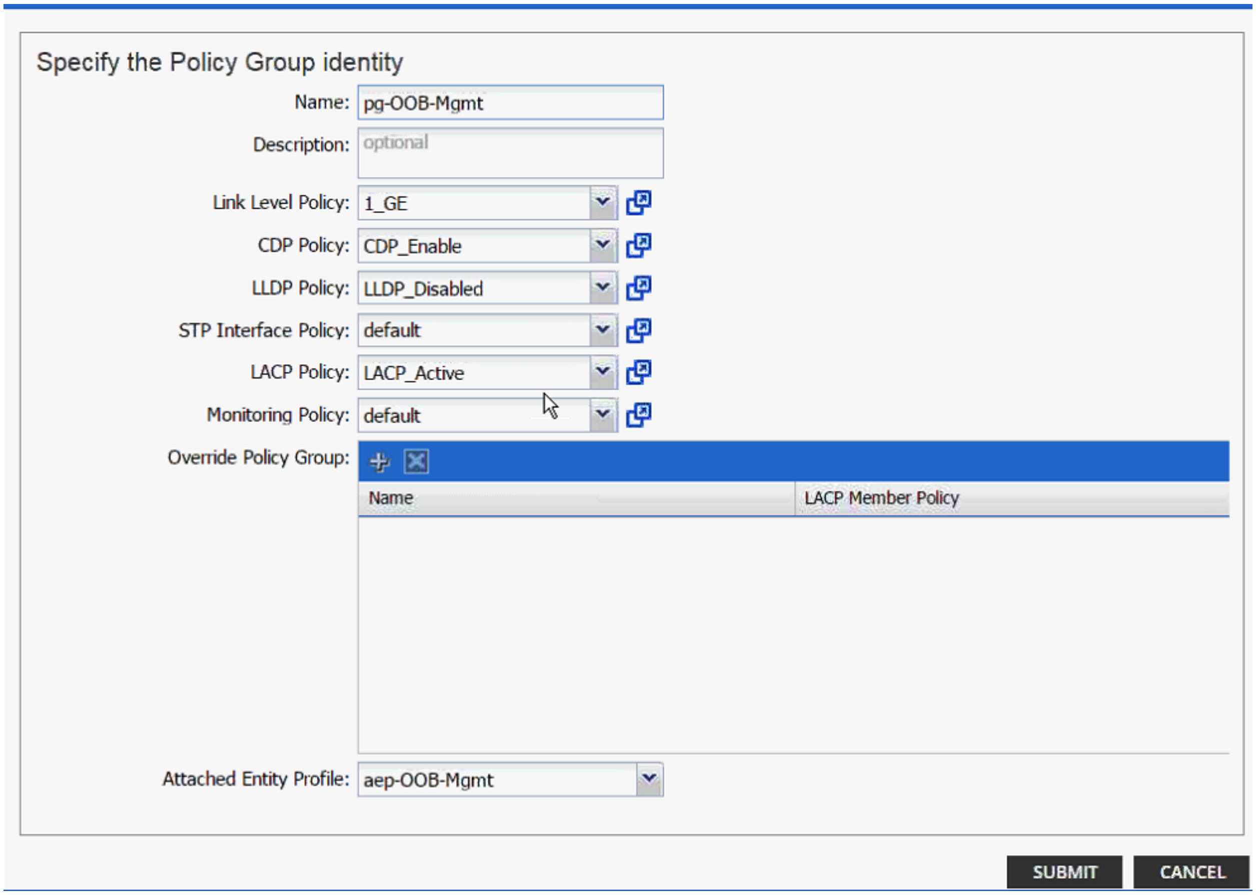

Specify the Policy Group identity

Name: pg-OOB-Mgmt

Description: optional

Link Level Policy: 1_GE

CDP Policy: CDP_Enable

LLDP Policy: LLDP_Disabled

STP Interface Policy: default

LACP Policy: LACP_Active

Monitoring Policy: default

Override Policy Group:

Name	LACP Member Policy

Attached Entity Profile: aep-OOB-Mgmt

37. On the Configure Interface, PC, vPC screen, click SAVE.

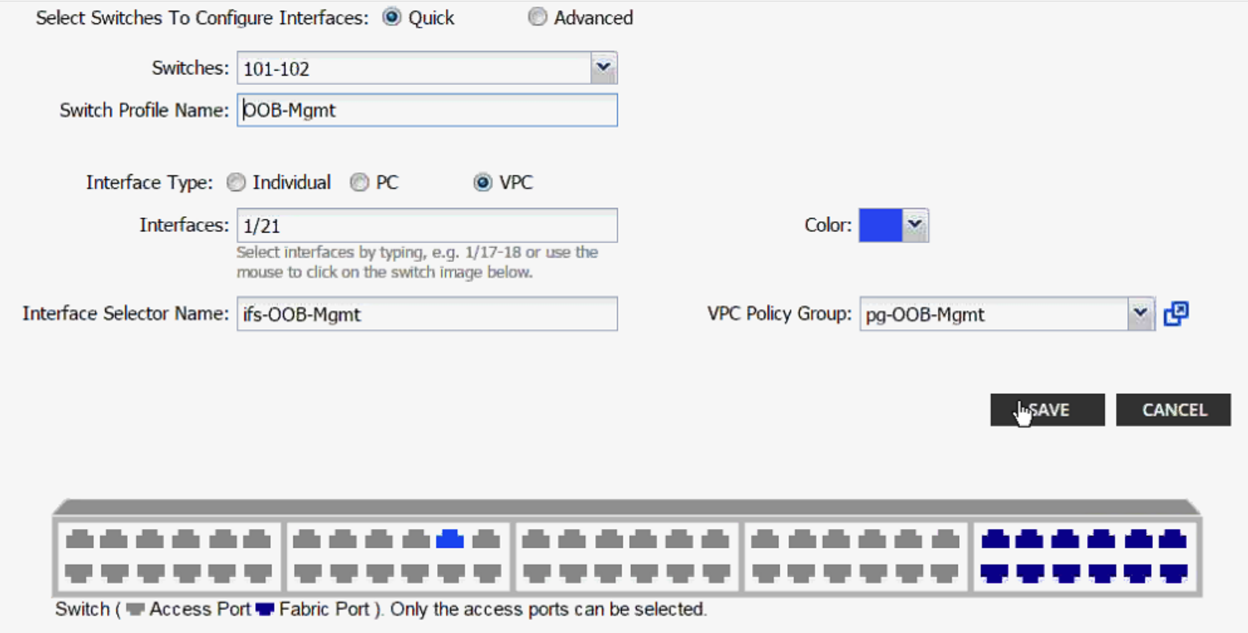

Select Switches To Configure Interfaces: ● Quick ○ Advanced

Switches: 101-102

Switch Profile Name: OOB-Mgmt

Interface Type: ○ Individual ○ PC ● VPC

Interfaces: 1/21

Select interfaces by typing, e.g. 1/17-18 or use the mouse to click on the switch image below.

Color:

Interface Selector Name: ifs-OOB-Mgmt

VPC Policy Group: pg-OOB-Mgmt

Switch (Access Port Fabric Port). Only the access ports can be selected.

38. Click SAVE.

39. Click SUBMIT to finish the vPC creation.

Configuring the Management Switch

The management switch configuration is covered below. Management switch connects to both Leaf switches using port Gig0/37 and Gig 0/38. These ports form a LACP port-channel Po1.

```
interface Port-channel1
 description *** To ACI Fabric for Common Segment Connectivity ***
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3177
 switchport mode trunk
end
!
interface GigabitEthernet0/37
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3177
 switchport mode trunk
 channel-group 1 mode active
end
!
interface GigabitEthernet0/38
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3177
 switchport mode trunk
 channel-group 1 mode active
end
!
MGMTSW# show etherchannel summary
Flags:  D - down        P - in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port



Number of channel-groups in use: 1

Number of aggregators:           1


Group  Port-channel  Protocol    Ports

------+-------------+-----------+------------------------------------------------

1      Po1(SU)         LACP      Gi0/37(P)   Gi0/38(P)
```

Configuring Common Tenant

An EPG and a contract is configured in the common tenant. This Contract will be consumed in the all the application tenants, which require access to common services.

Common Management – Application Profile Creation

1. Select Tenant and common tenant from the top menu.
2. Expand Tenant common in the left menu bar.
3. Right-click Application Profile and click Create Application Profile.
4. In the CREATE APPLICATION PROFILE dialog box, enter Management for the Name.
5. From the drop-down menu, select default for Monitoring Policy.
6. Click + next to EPG to add an EPG.

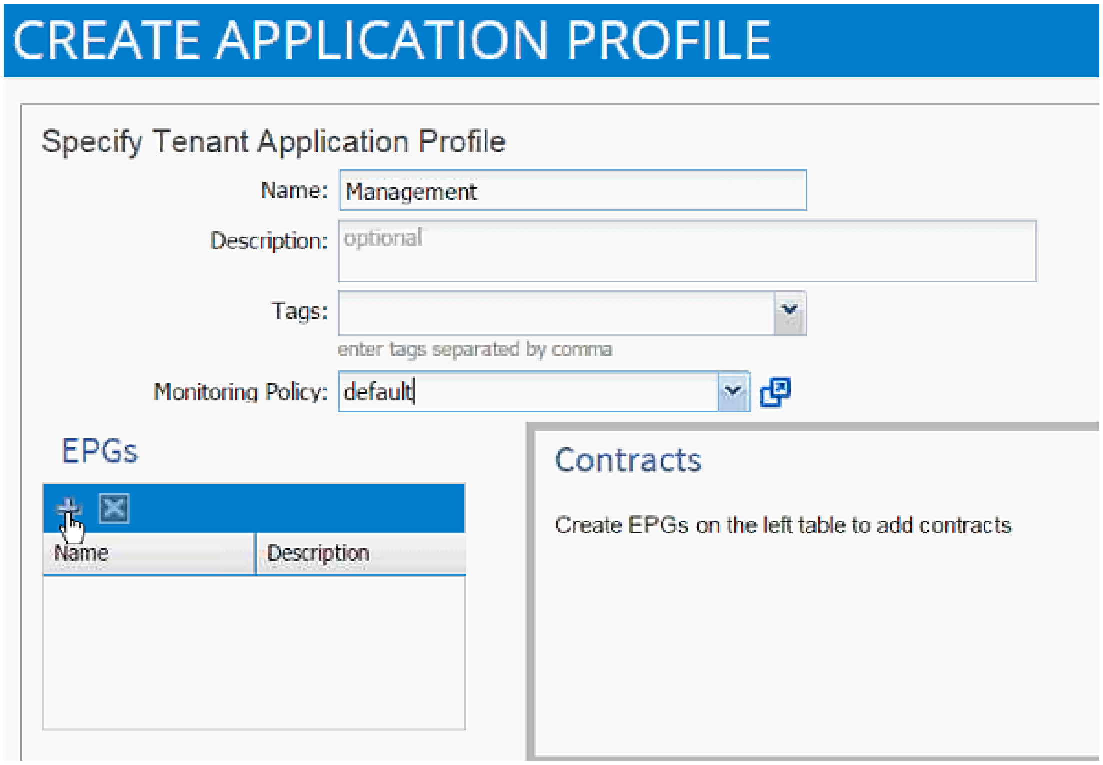

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: Management

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

EPGs

Name	Description

Contracts

Create EPGs on the left table to add contracts

7. In the CREATE APPLICATION EPG dialog box, enter Mgmt_Access for the Name.

8. From the drop-down menu, select Create Bridge Domain as the Bridge Domain.

9. In the CREATE BRIDGE DOMAIN dialog box, use bd-Internal as the Name of the bridge domain.

10. From the drop-down menu next to Network, select Create Private Network.

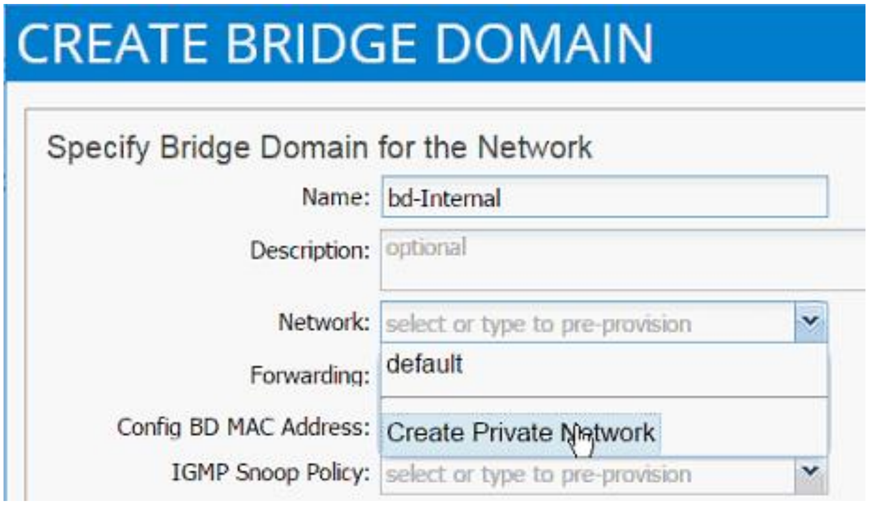

CREATE BRIDGE DOMAIN

Specify Bridge Domain for the Network

Name: bd-Internal

Description: optional

Network: select or type to pre-provision

Forwarding:

Config BD MAC Address:

IGMP Snoop Policy: select or type to pre-provision

11. In the CREATE PRIVATE NETWORK dialog box, enter Common-Mgmt for the Name.

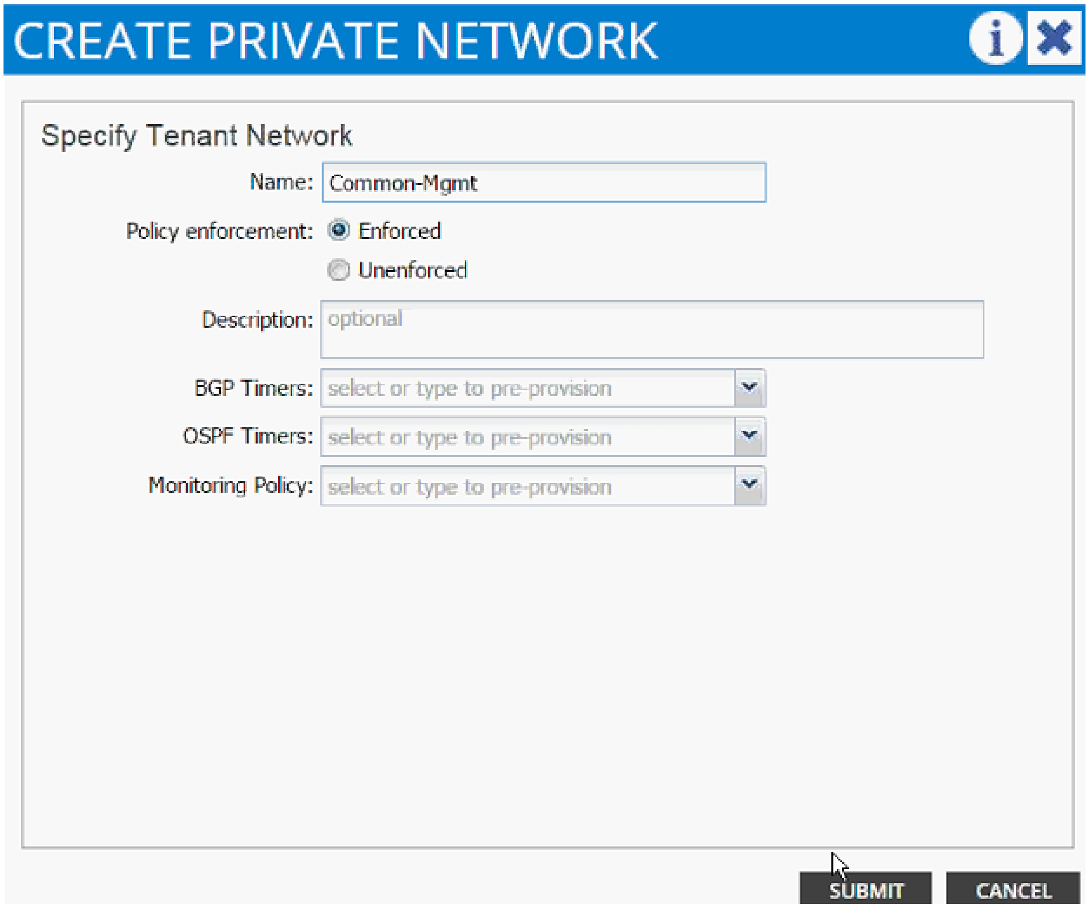

CREATE PRIVATE NETWORK

Specify Tenant Network

Name: Common-Mgmt

Policy enforcement: ● Enforced ○ Unenforced

Description: optional

BGP Timers: select or type to pre-provision

OSPF Timers: select or type to pre-provision

Monitoring Policy: select or type to pre-provision

12. Click SUBMIT.

13. Drop down the menu next to Forwarding and select Custom.

14. Check the boxes to enable Flooding and Unicast Routing.

15. Select default for IGMP Snoop Policy.

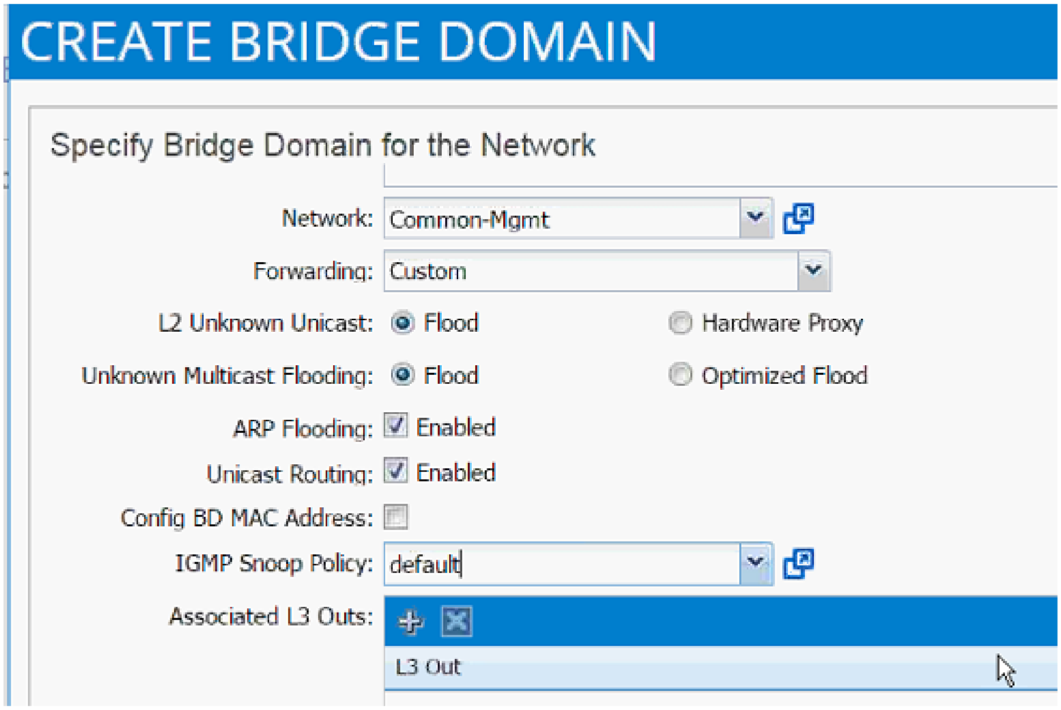

CREATE BRIDGE DOMAIN

Specify Bridge Domain for the Network

Network: Common-Mgmt

Forwarding: Custom

L2 Unknown Unicast: ● Flood ○ Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: ☑ Enabled

Unicast Routing: ☑ Enabled

Config BD MAC Address: □

IGMP Snoop Policy: default

Associated L3 Outs:

L3 Out

16. Click SUBMIT to finish bridge domain creation.

17. Back in the CREATE APPLICATION EPG dialog box, select the newly created bridge domain.

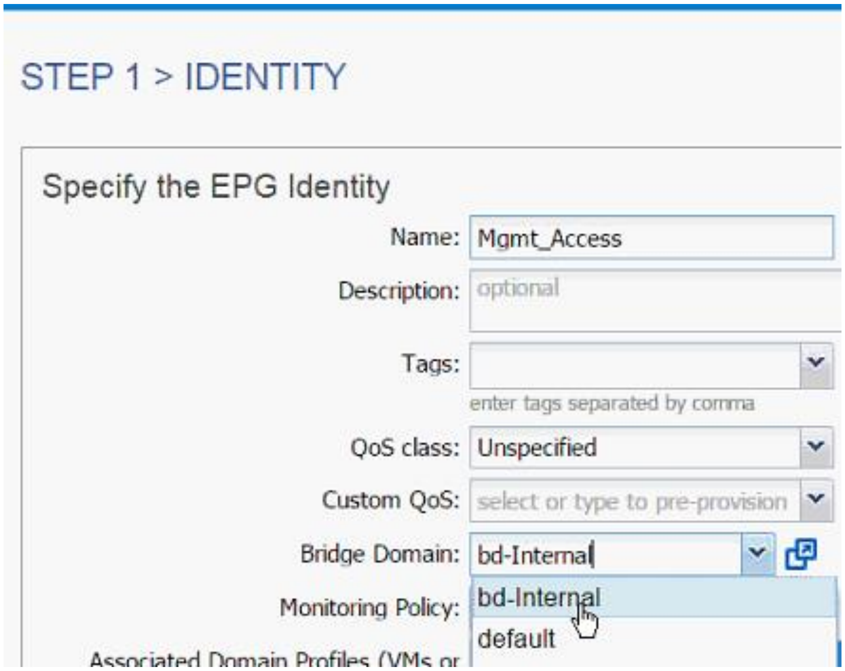

STEP 1 > IDENTITY

Specify the EPG Identity

Name: Mgmt_Access

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-Internal

Monitoring Policy:

Associated Domain Profiles (VMs or

18. From the drop-down menu, select default for Monitoring Policy.

19. Click OK to finish EPG creation.

20. Click SUBMIT to finish creating Application Profile.

21. Expand the newly created Management Application profile from the menu bar on the left.

22. Expand Management, expand Application EPGs and expand EPG Mgmt_Access.

23. Click Static Bindings (Paths).

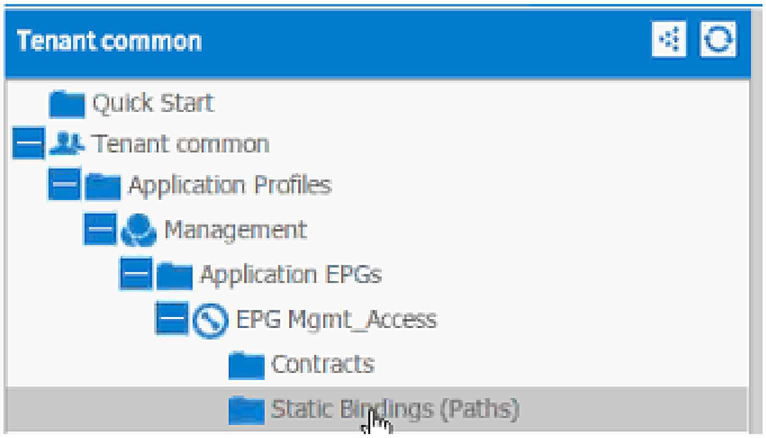

24. Click Action on the right hand work area.

25. Click Deploy Static EPG on PC, vPC, or Interface.

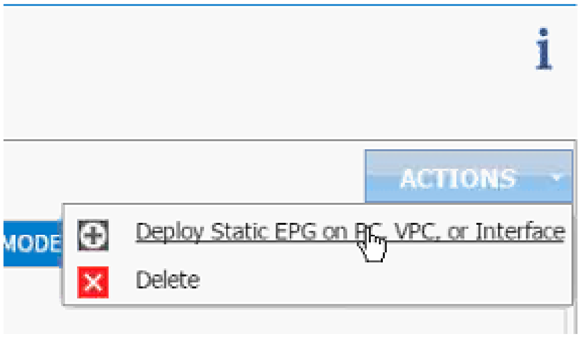

26. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

27. From the drop-down menu Path, select OUT OF BAND management switch VPC.

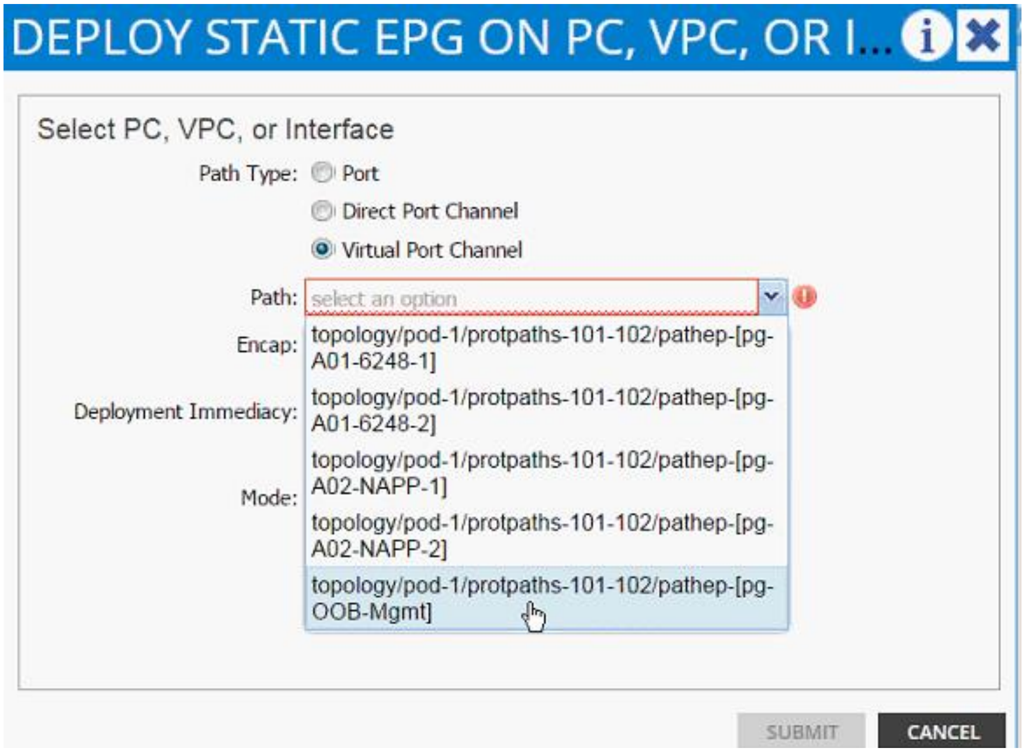

DEPLOY STATIC EPG ON PC, VPC, OR I...
Select PC, VPC, or Interface
Path Type:
○ Port
○ Direct Port Channel
● Virtual Port Channel
Path:
select an option
Encap:
Deployment Immediacy:
Mode:
topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]
topology/pod-1/protpaths-101-102/pathep-[pg-OOB-Mgmt]
SUBMIT
CANCEL

28. Enter vlan-< common_mgmt_segment> for Encap (VLAN 3177 is the common management segment VLAN in the screen capture below).

29. Change Deployment Immediacy to Immediate.

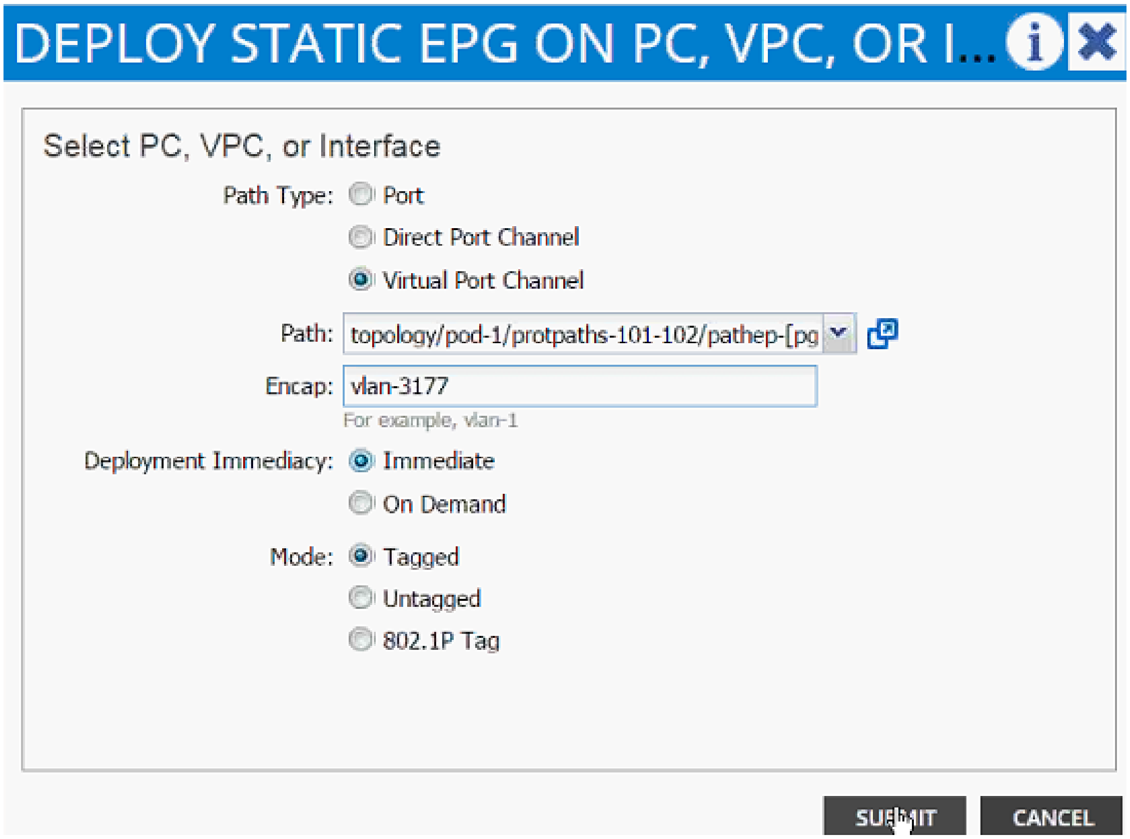

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[pg

Encap: vlan-3177

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

SUBMIT CANCEL

30. Click SUBMIT.

31. On the left hand menu bar, click on Subnet.

32. From the right ACTIONS menu, select Create EPG Subnet.

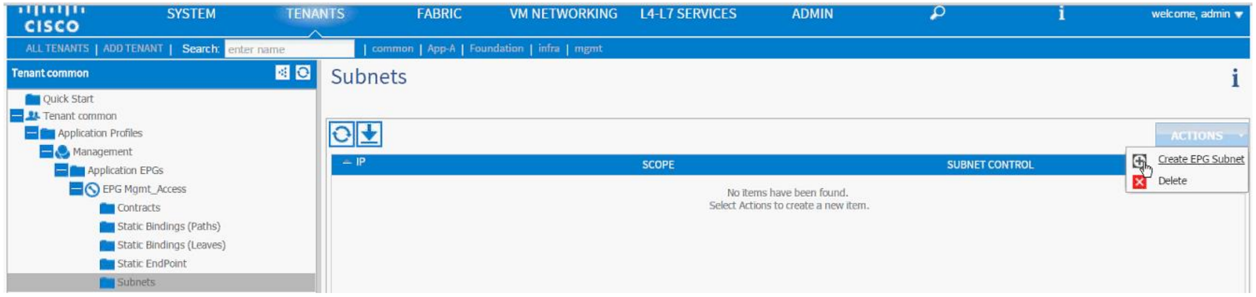

No items have been found.
Select Actions to create a new item.

33. In the CREATE EPG SUBNET dialog box, enter 192.168.3.253/24 for the Default Gateway IP. The Mask filed should be auto populated with 255.255.255.0.

34. From the scope, select only Shared Subnet.

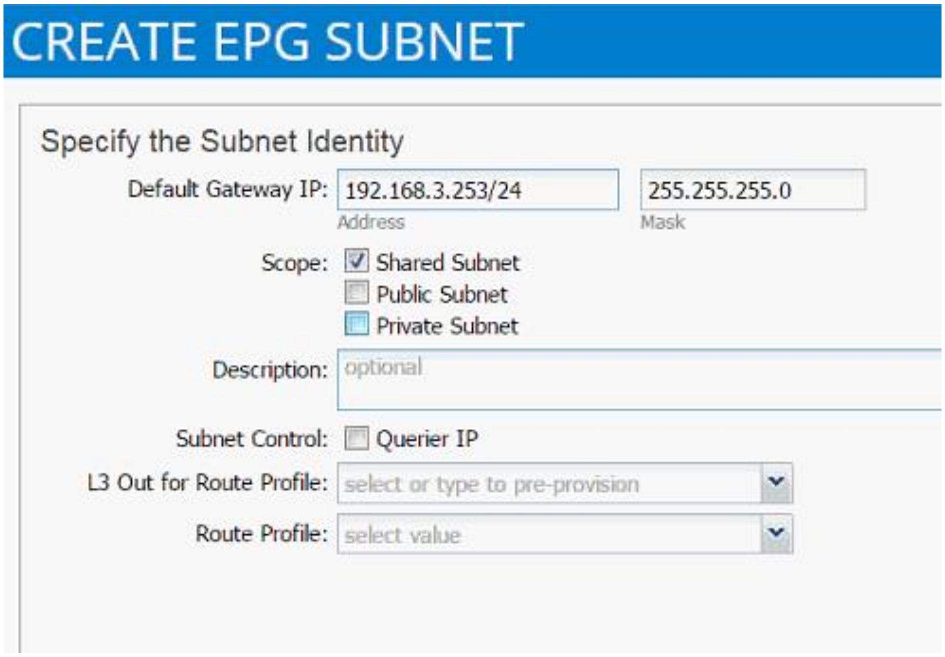

CREATE EPG SUBNET

Specify the Subnet Identity

Default Gateway IP: 192.168.3.253/24 (Address) 255.255.255.0 (Mask)

Scope: ☑ Shared Subnet
☐ Public Subnet
☐ Private Subnet

Description: optional

Subnet Control: ☐ Querier IP

L3 Out for Route Profile: select or type to pre-provision

Route Profile: select value

35. Click SUBMIT.

36. Click Contracts under the EPG Mgmt_Access.

37. Click Action and select Add Provided Contract.

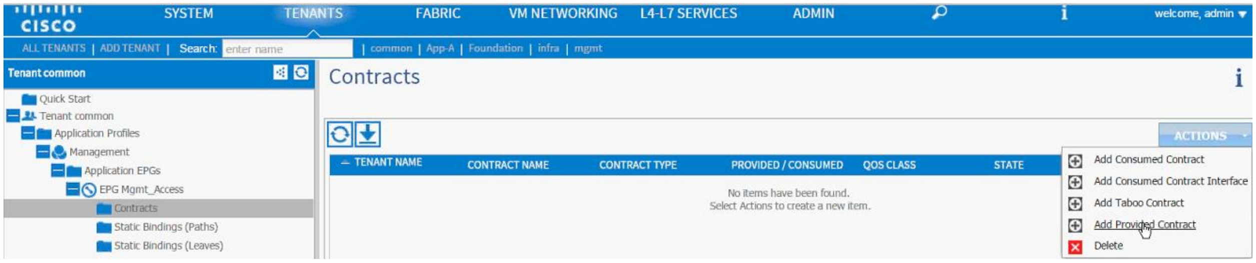

No items have been found.
Select Actions to create a new item.

38. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

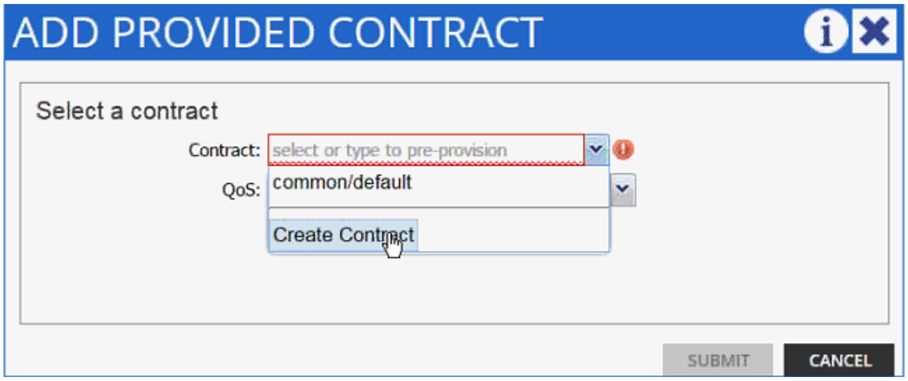

ADD PROVIDED CONTRACT

Select a contract

Contract: select or type to pre-provision

QoS: common/default

39. Enter Common-Mgmt as Name in the CREATE CONTRACT dialog box.

40. Change the Scope to global from the drop-down menu.

41. Click + next to Subjects to add a new contract subject.

42. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

43. Click + under Filter Chain to add a new filter.

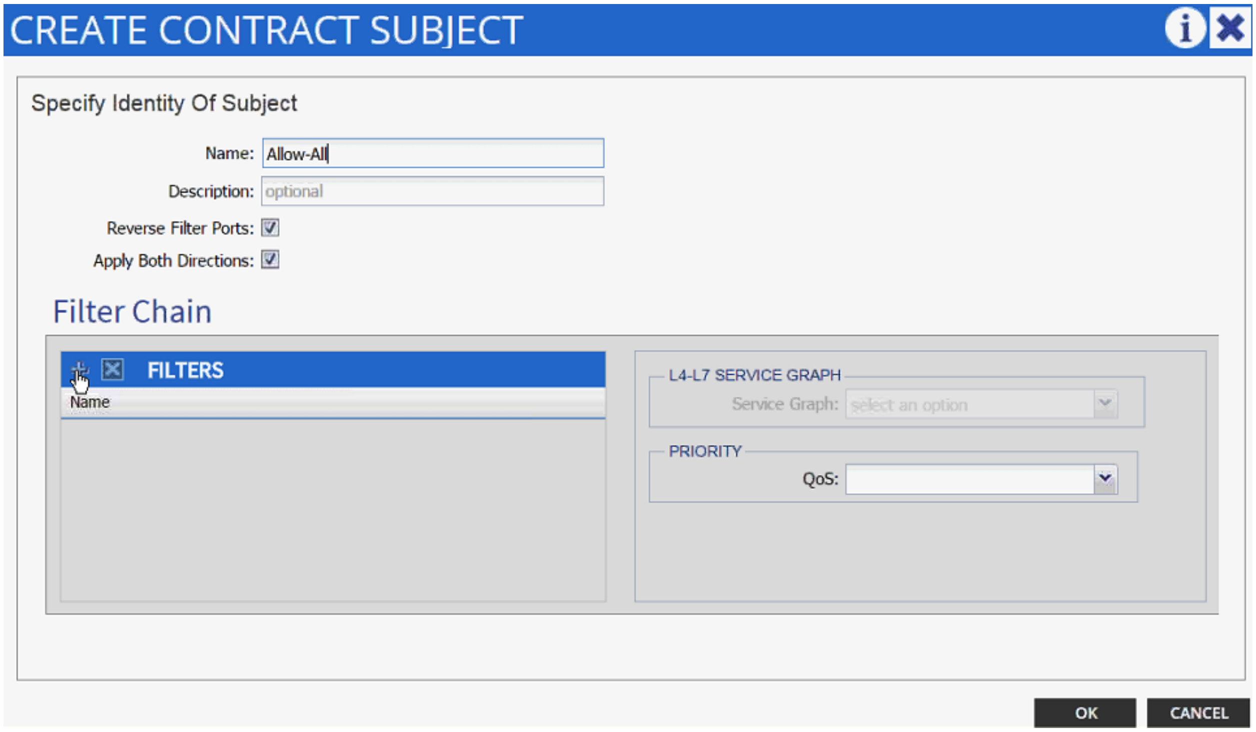

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS
Name

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

44. Drop down the menu under FILTERS and click +.

45. In the CREATE FILTER dialog box, enter Allow-All as the Name. In this example, we will allow all the traffic for this contract.

46. Click + to add a filter.

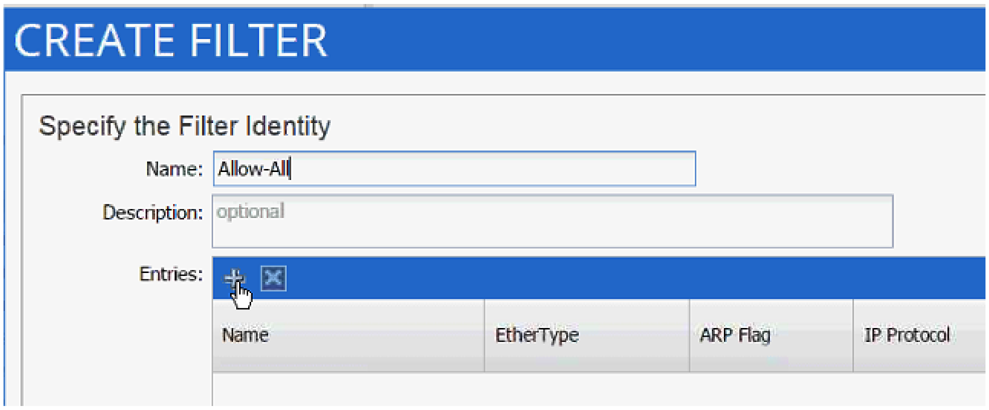

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol

47. Enter Allow-All as the name of the filter.

48. From drop-down menu, select IP as Ethertype.

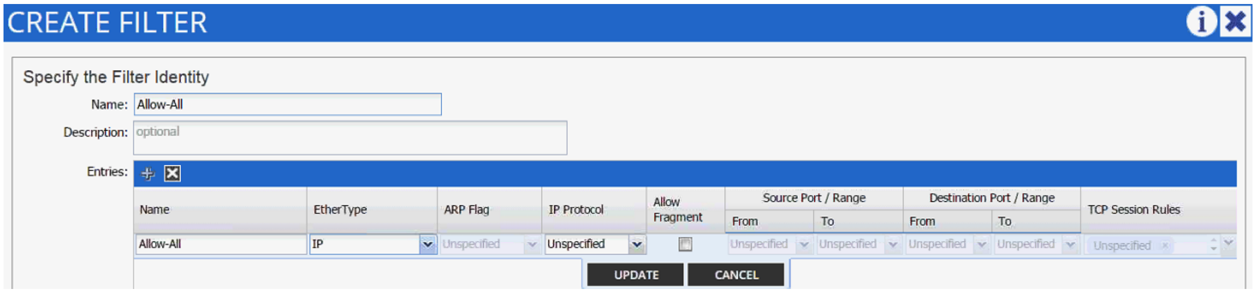

CREATE FILTER

Specify the Filter Identity

Name: Allow-All

Description: optional

Entries:

Name	EtherType	ARP Flag	IP Protocol	Allow Fragment	Source Port / Range		Destination Port / Range		TCP Session Rules
					From	To	From	To	
Allow-All	IP	Unspecified	Unspecified	□	Unspecified	Unspecified	Unspecified	Unspecified	Unspecified

49. Click Update.

50. Click SUBMIT to create the filter.

51. Click UPDATE to add the newly created filter to the filter chain.

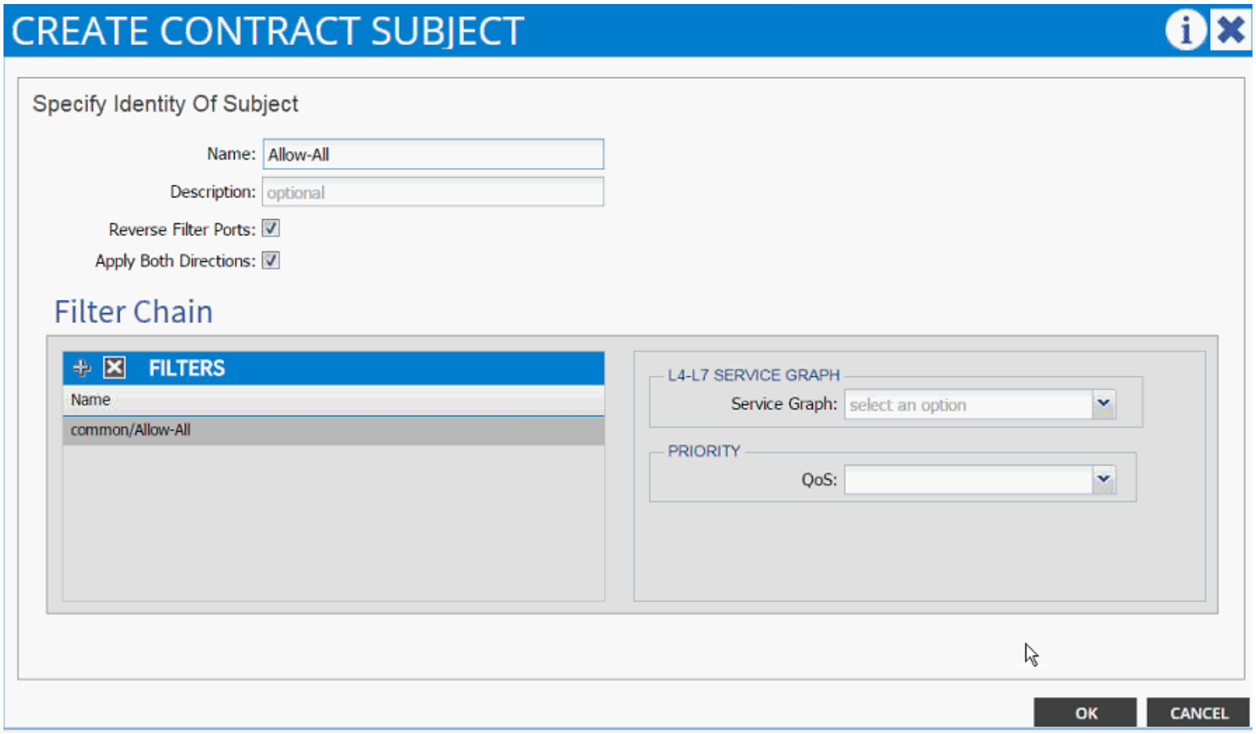

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS

Name
common/Allow-All

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

52. Click OK to finish creating the Contract Subject.

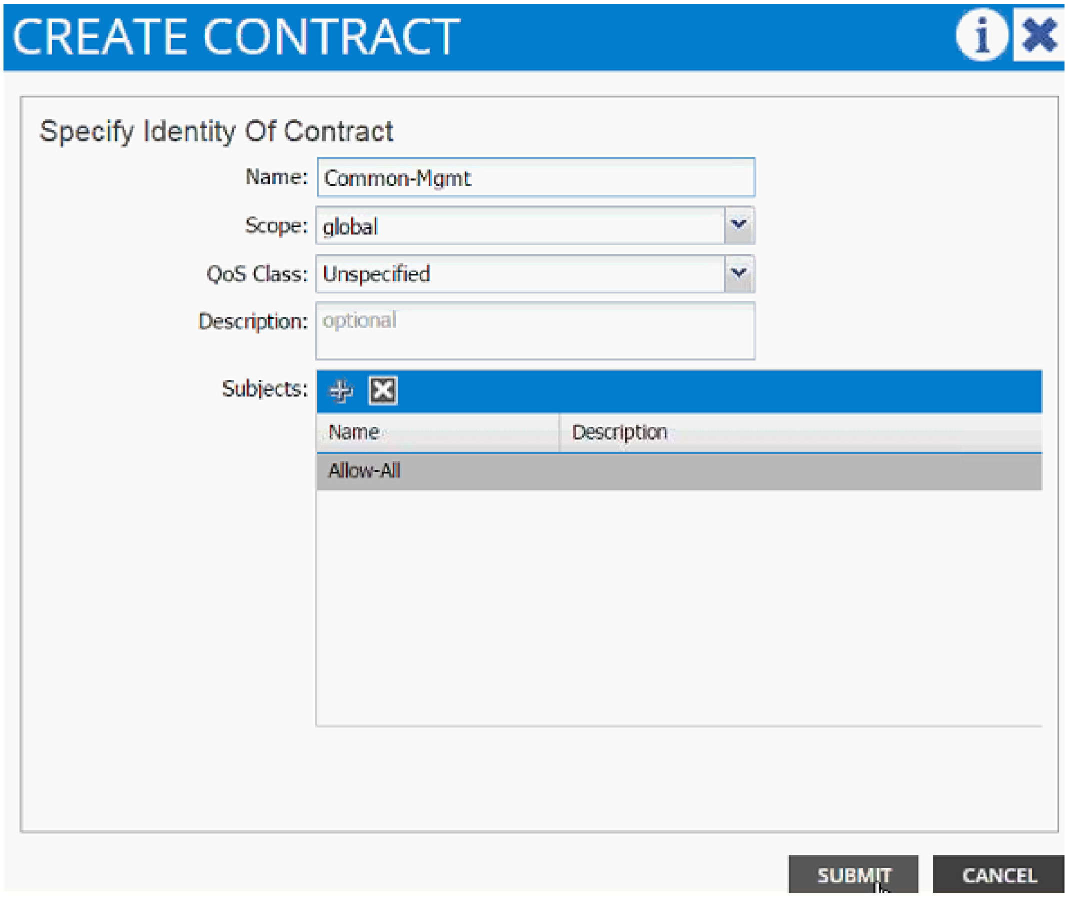

CREATE CONTRACT

Specify Identity Of Contract

Name: Common-Mgmt

Scope: global

QoS Class: Unspecified

Description: optional

Subjects:

Name	Description
Allow-All	

53. Click SUBMIT.

54. Click SUBMIT again to finish adding a provided contract.

55. To validate access to the EPG gateway just added, ping 192.168.3.253 from a common services VM. The VM should be able to ping the address.

Consuming the Common Contract in Application EPGs

1. Select Tenant and App-A tenant from the top menu.

2. Expand Tenant App-A in the left menu bar.

3. Expand the Application Profiles, 3-Tier-App, Application EPGs and EPG Web.

4. Click Contracts.

5. Click ACTIONS and select Add Consumed Contract.

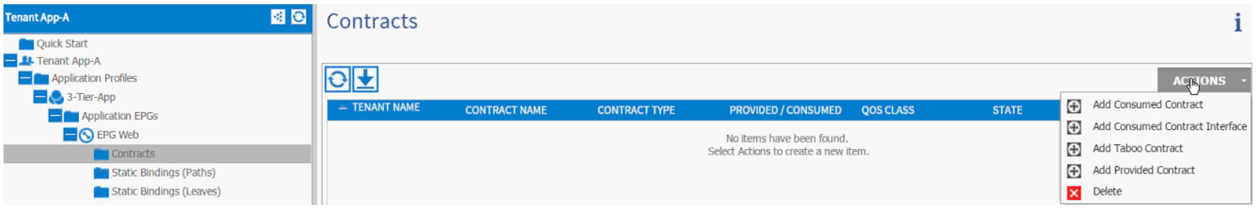

TENANT NAME	CONTRACT NAME	CONTRACT TYPE	PROVIDED / CONSUMED	QOS CLASS	STATE

No items have been found.
Select Actions to create a new item.

6. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select Common/Common-Mgmt contract.

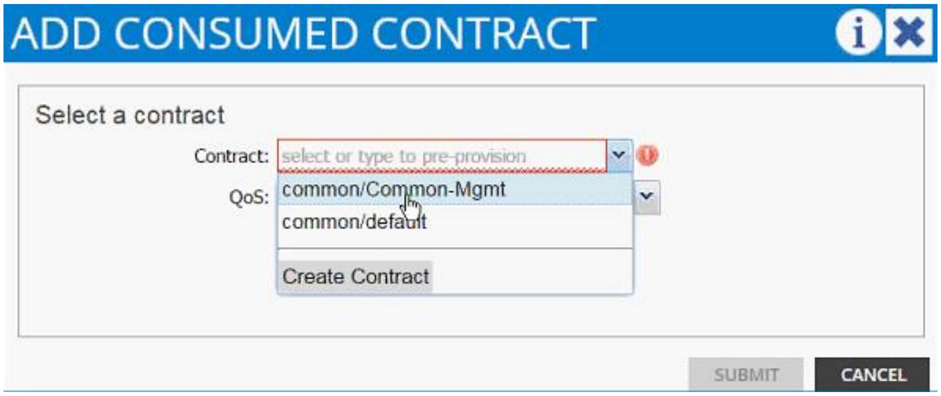
ADD CONSUMED CONTRACT

Select a contract

Contract: select or type to pre-provision

QoS:

common/Common-Mgmt
common/default

7. Click SUBMIT.

8. To validate the contract definition, click on Application EPGs under Application Profile NFS in the left menu bar.

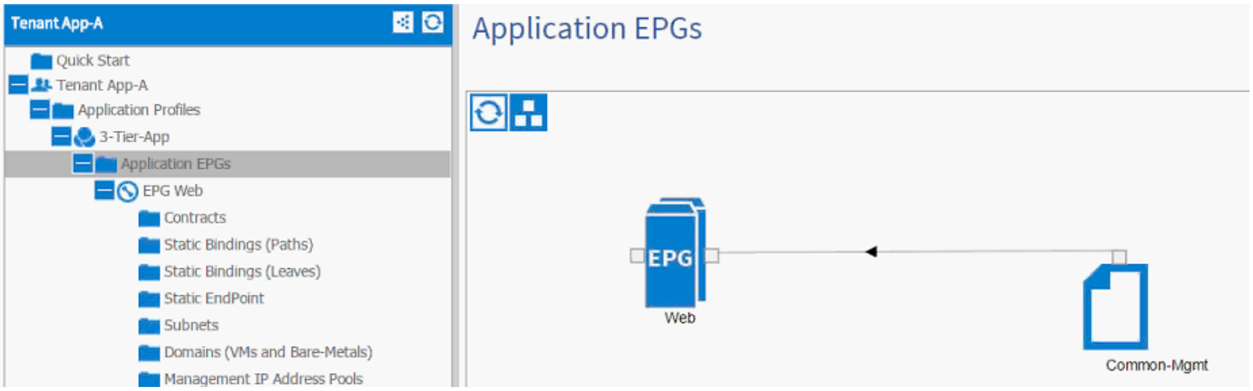

9. Repeat above steps for all the application tiers that need access to the common services.

10. For a Web VM with shown Network Parameters, communication to the 192.168.3.0 subnet should be established and ping should work from 10.1.1.1 to 192.168.3.11.

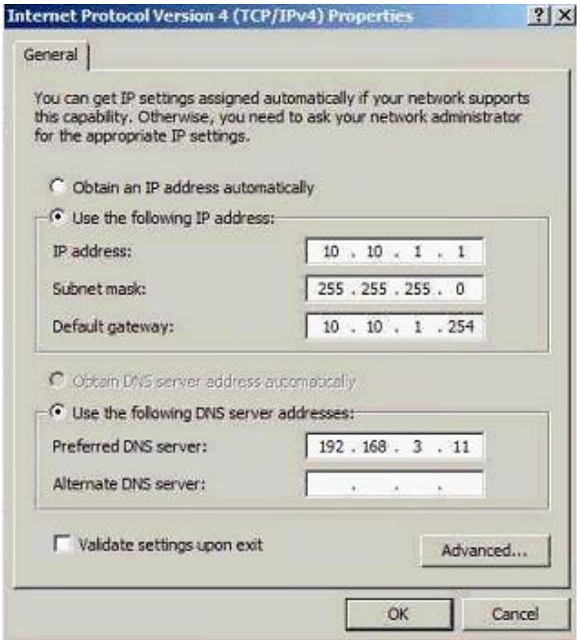

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address:	10 . 10 . 1 . 1
Subnet mask:	255 . 255 . 255 . 0
Default gateway:	10 . 10 . 1 . 254

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server:	192 . 168 . 3 . 11
Alternate DNS server:	. . .

□ Validate settings upon exit

Note: Make sure the common services VMs either use 192.168.3.253 as their default gateway or have a persistent route added for the 10.10.1.0/24 subnet with the gateway set as 192.168.3.253.

Cisco ACI – Accessing SVM Management Interface from Application Virtual Machines (Optional)

When configuring NetApp SnapManager and SnapDrive, access to SVM management LIF is required. This configuration can be enabled on a per tenant (per application) basis. A new application profile called SVM-Access is defined under the application tenant (App-A) and an EPG (svm-mgmt) is statically mapped to SVM management LIF VLAN. Access to the svm-mgmt from various application tiers (EPGs) is then enabled using contracts. The contracts are provided by the EPG svm-mgmt and consumed but the application EPGs such as Web and App as defined in the previous sections.

SVM-Access - Application Profile Creation

In this section, a bridge domain will be created and an Application Profile to setup SVM management interface connectivity between the application VMs the SVM management LIF. Since all the LIFs sharing the same uplink port-channel share the same MAC address, a unique bridge domain is required for SVM access.

To create a bridge domain, complete the following steps:

1. Select Tenant and App-A tenant from the top menu.

2. Expand Tenant App-A in the left menu bar.

3. Expand Networking, right click Bridge Domains and select Create Bridge Domain.

4. Use bd-svm-mgmt as Name of the bridge domain.

5. Select App-A from the drop-down menu as the Network.

6. Select Custom from the drop-down menu for Forwarding and enable Flood and ARP Flooding.

7. Select default as the IGMP Snoop Policy.

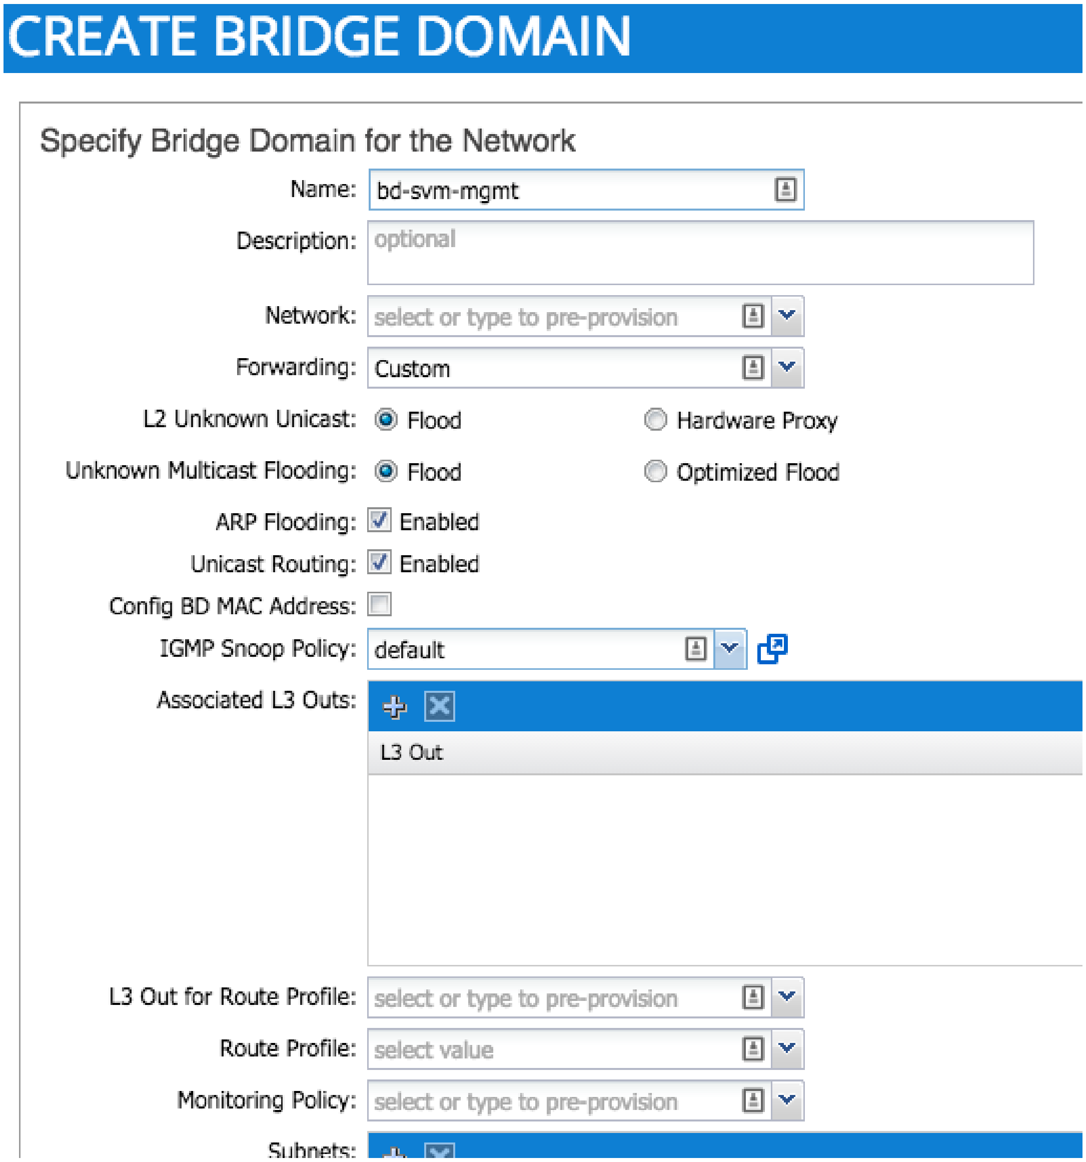

CREATE BRIDGE DOMAIN

Specify Bridge Domain for the Network

Name: bd-svm-mgmt

Description: optional

Network: select or type to pre-provision

Forwarding: Custom

L2 Unknown Unicast: ● Flood ○ Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: ☑ Enabled

Unicast Routing: ☑ Enabled

Config BD MAC Address: □

IGMP Snoop Policy: default

Associated L3 Outs:

L3 Out

L3 Out for Route Profile: select or type to pre-provision

Route Profile: select value

Monitoring Policy: select or type to pre-provision

Subnets:

8. Click SUBMIT.

9. In the menu on the left, right click Application Profile and click Create Application Profile.

10. In the CREATE APPLICATION PROFILE dialog box, enter SVM-Access as the Name.

11. From the drop-down menu, select default for Monitoring Policy.

12. Click + next to EPG to add an EPG.

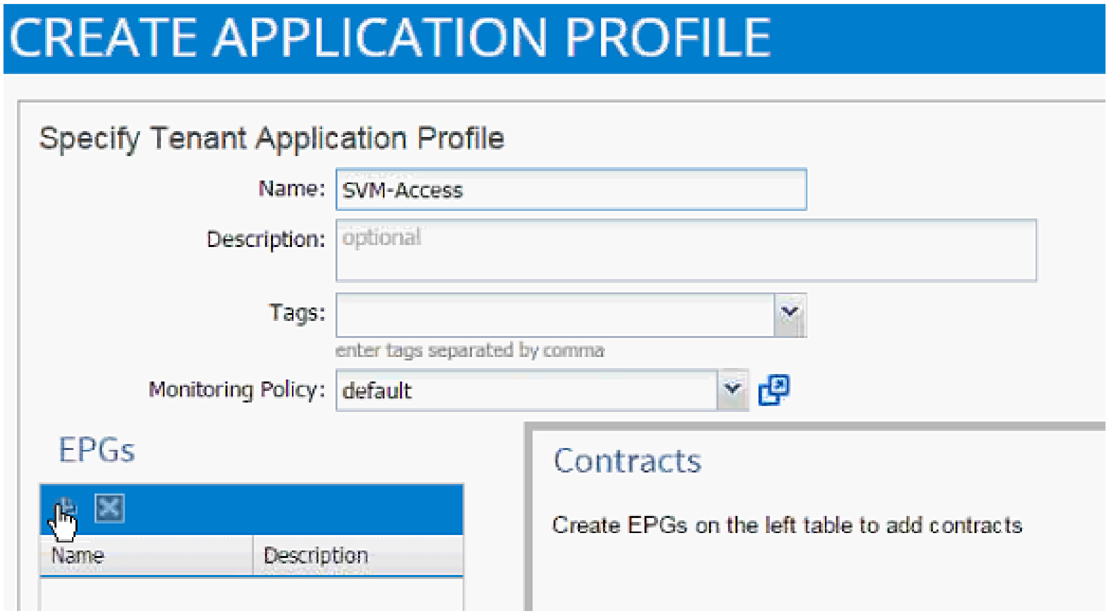

CREATE APPLICATION PROFILE

Specify Tenant Application Profile

Name: SVM-Access

Description: optional

Tags:

enter tags separated by comma

Monitoring Policy: default

Create EPGs on the left table to add contracts

13. In the CREATE APPLICATION EPG dialog box, enter svm-mgmt as the Name.

14. From the drop-down menu, select bd-svm-mgmt as the Bridge Domain.

15. From the drop-down menu, select default for Monitoring Policy.

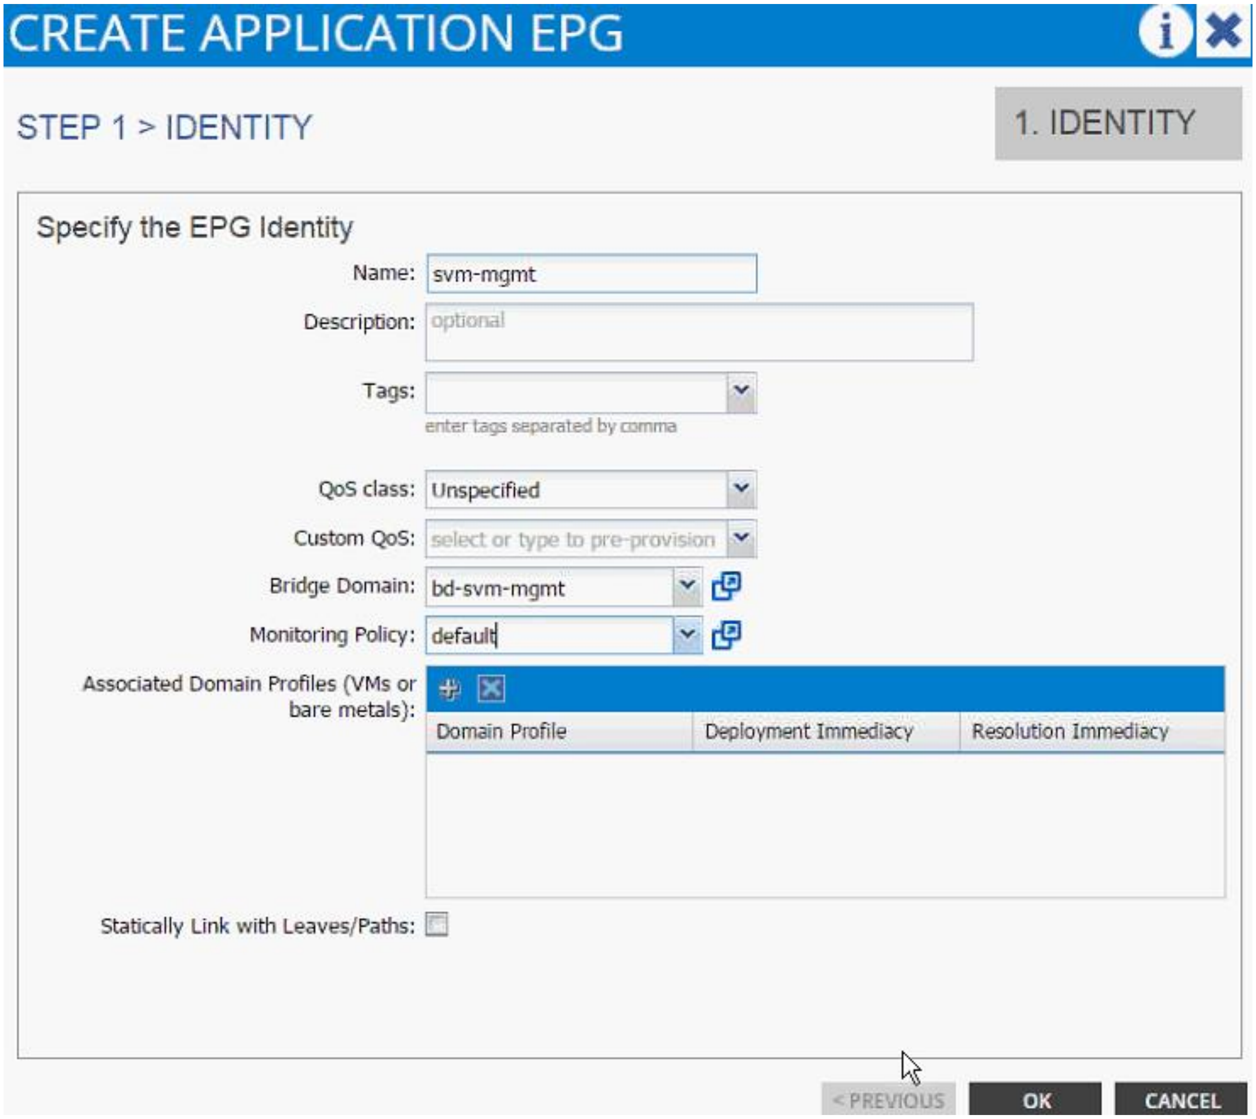

CREATE APPLICATION EPG

STEP 1 > IDENTITY

Specify the EPG Identity

Name: svm-mgmt

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-svm-mgmt

Monitoring Policy: default

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy

Statically Link with Leaves/Paths: □

16. Click OK.

17. Click SUBMIT to finish creating Application Profile.

Modifying Physical Domain

A Physical domain associated with NetApp storage needs to be modified and the SVM Management LIF VLAN associated with App-A SVM needs to be added to the physical domain. To modify the physical domain, complete the following steps:

1. Select Fabric and Access Policies from the top menu.

2. Expand Pools and then VLAN.

3. Click the pool name associated with NetApp controllers (vp-A01-NetApp in this example) and click + to add another Encap Block.

4. Enter <<var_svm_mgmt_vlan_tenant>> (a range of 3181 to 3181 in this document) for SVM Management LIF VLAN.

5. Click SUBMIT.

EPG and Contract Configuration

To configure EPG and the contract, complete the following steps:

1. Select Tenant and then App-A from the top menu.

2. Expand the newly created SVM-Access Application profile from the menu bar on the left.

3. Expand SVM-Access, expand Application EPGs and expand EPG svm-mgmt.

4. Click Static Bindings (Paths).

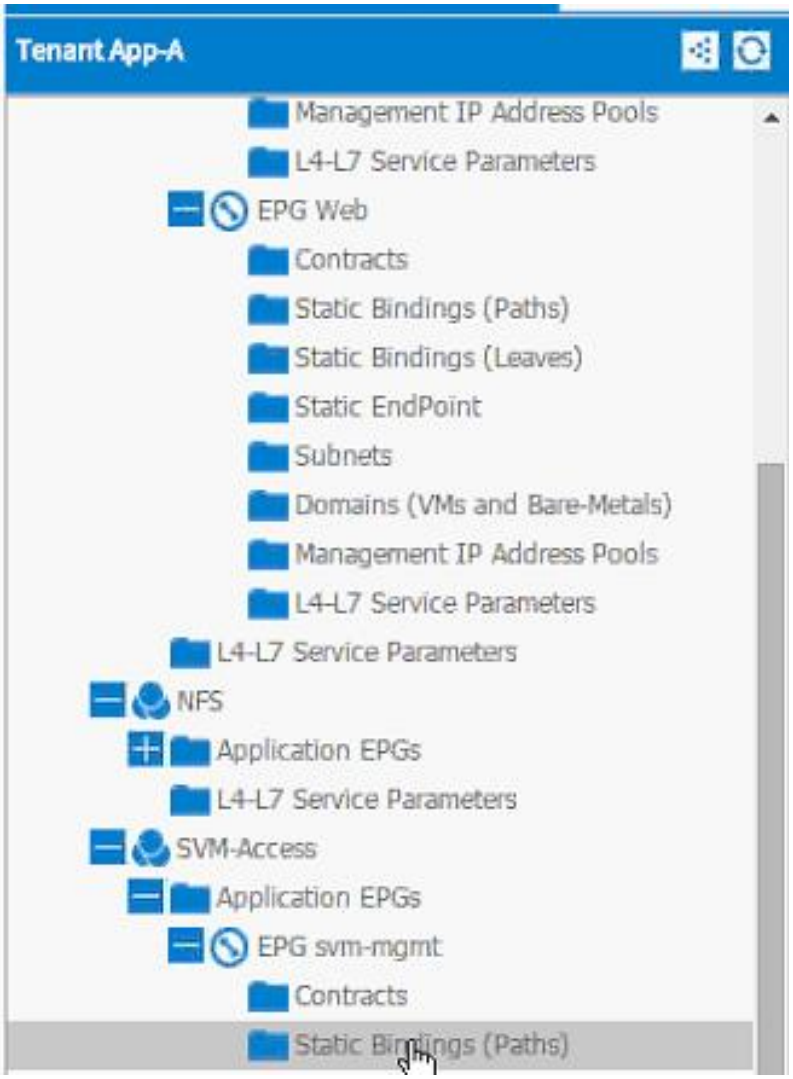

5. Click Action on the right hand work area.

6. Click Deploy Static EPG on PC, vPC, or Interface.

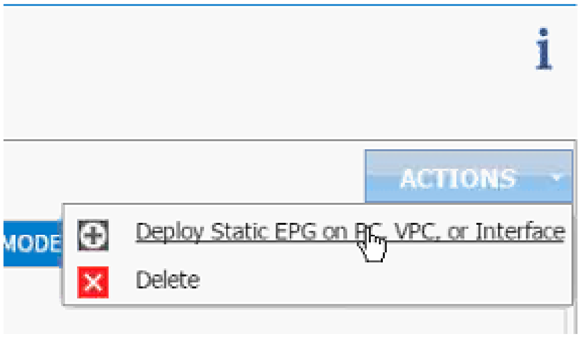

7. In the DEPLOY STATIC EPG ON PC, vPC OR INTERFACE dialog box, select Virtual Port Channel as the Path Type.

8. From the drop-down menu Path, select NetApp Controller 1.

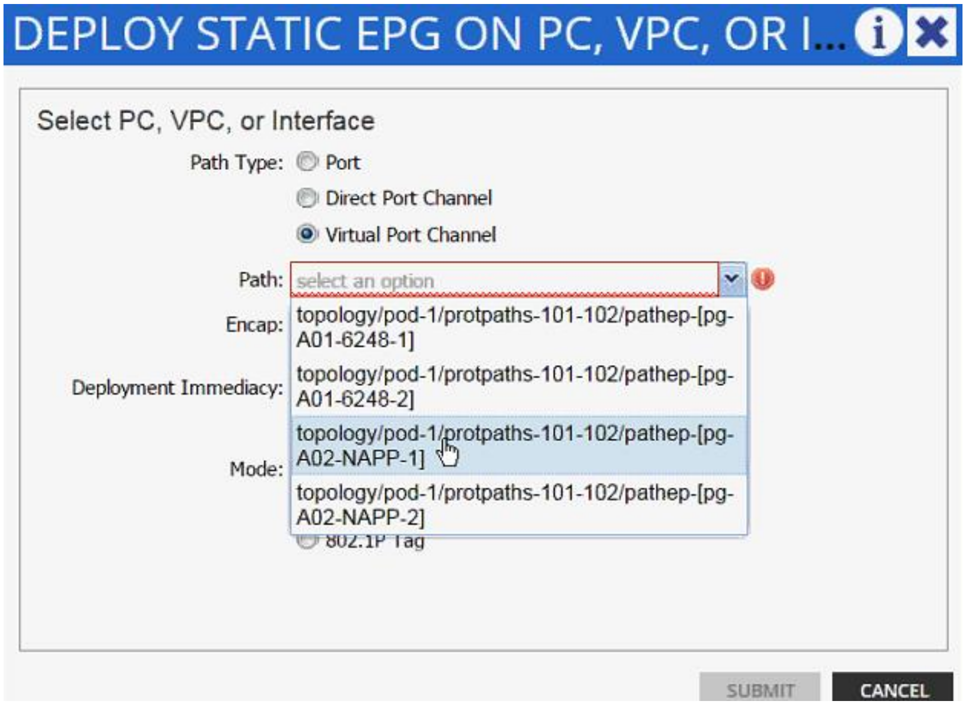
DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port
○ Direct Port Channel
● Virtual Port Channel

Path: select an option
topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A01-6248-2]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-1]
topology/pod-1/protpaths-101-102/pathep-[pg-A02-NAPP-2]

Encap:

Deployment Immediacy:

Mode: ○ 802.1P Tag

9. Enter vlan-< var_svm_mgmt_vlan_tenant > for Encap (VLAN 3181 is the Mgmt VLAN on NetApp Controller in the screen capture below).

10. Change Deployment Immediacy to Immediate.

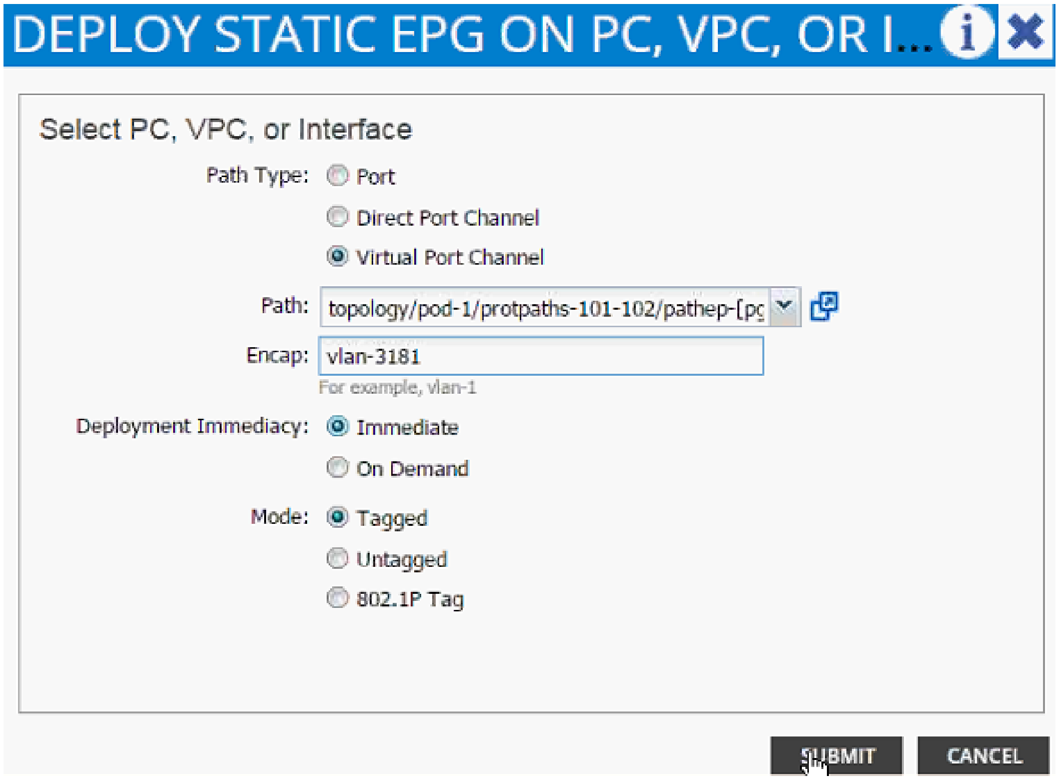

DEPLOY STATIC EPG ON PC, VPC, OR I...

Select PC, VPC, or Interface

Path Type: ○ Port ○ Direct Port Channel ● Virtual Port Channel

Path: topology/pod-1/protpaths-101-102/pathep-[pc

Encap: vlan-3181

For example, vlan-1

Deployment Immediacy: ● Immediate ○ On Demand

Mode: ● Tagged ○ Untagged ○ 802.1P Tag

11. Click Submit.

12. Repeat these steps for mapping NetApp Controller 2 path.

13. Click Subnet under the EPG svm-mgmt.

14. Click Contracts under the EPG svm-mgmt.

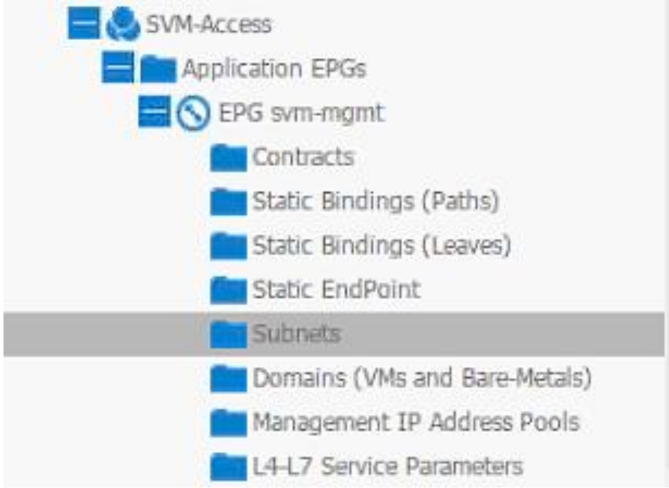

15. Click the Action menu on the right and select Create EPG Subnet.

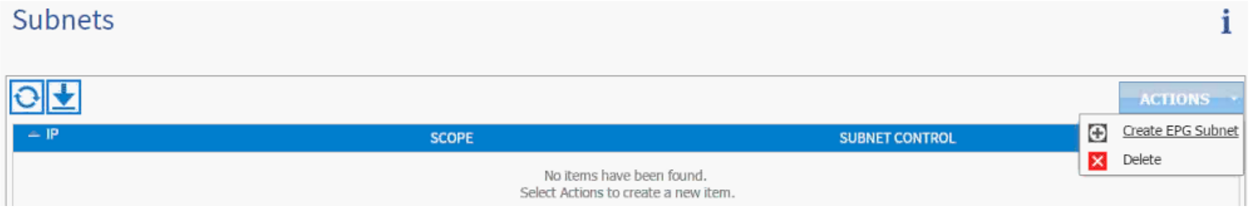
Subnets

IP	SCOPE	SUBNET CONTROL

No items have been found.
Select Actions to create a new item.

16. Enter 192.168.181.254/24 for the Default Gateway IP. This IP address is the gateway that SVM management LIF will use.

17. Change scope to only Private Subnet.

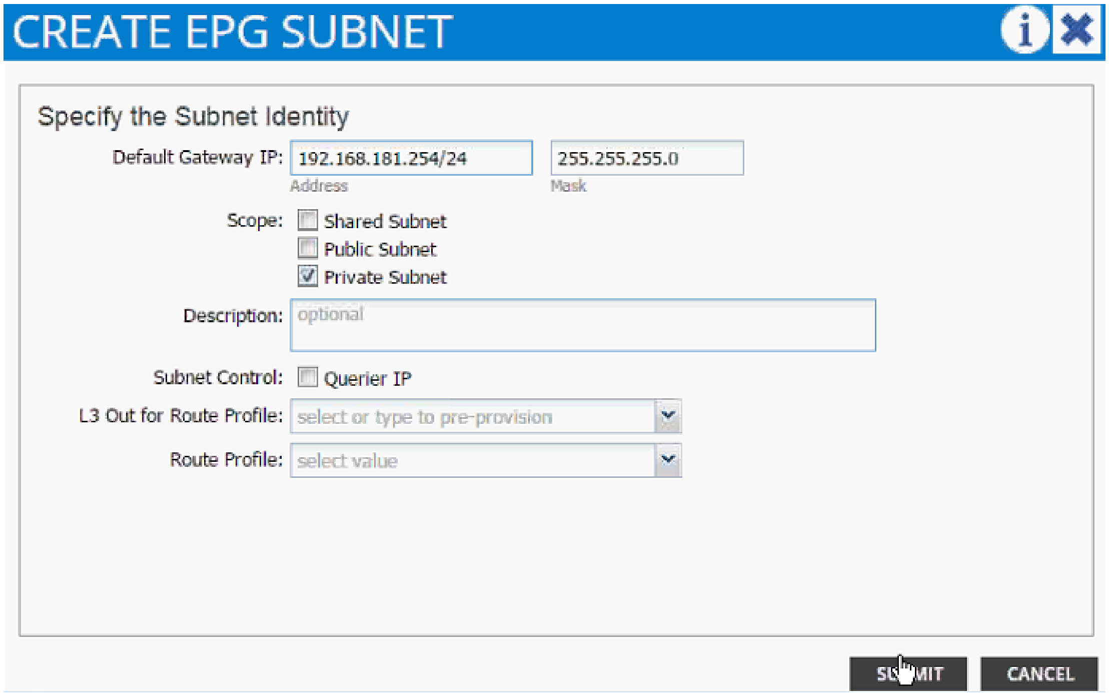
CREATE EPG SUBNET

Specify the Subnet Identity

Default Gateway IP: 192.168.181.254/24 (Address) 255.255.255.0 (Mask)

Scope: □ Shared Subnet
□ Public Subnet
☑ Private Subnet

Description: optional

Subnet Control: □ Querier IP

L3 Out for Route Profile: select or type to pre-provision

Route Profile: select value

18. Click SUBMIT.

19. Click Action and select Add Provided Contract.

20. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

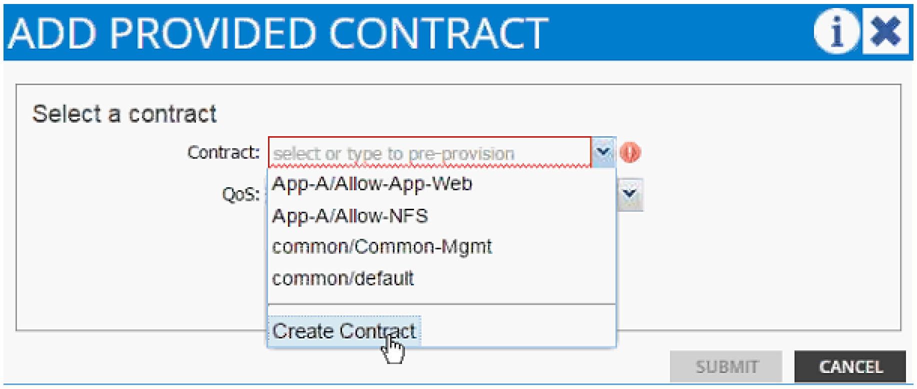
ADD PROVIDED CONTRACT

Select a contract

Contract: select or type to pre-provision

QoS:

App-A/Allow-App-Web
App-A/Allow-NFS
common/Common-Mgmt
common/default

Create Contract

21. Enter Allow-SVM-Access as Name in the CREATE CONTRACT dialog box.

22. Click + next to Subjects to add a new contract subject.

23. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

24. Click + under Filter Chain to add a new filter.

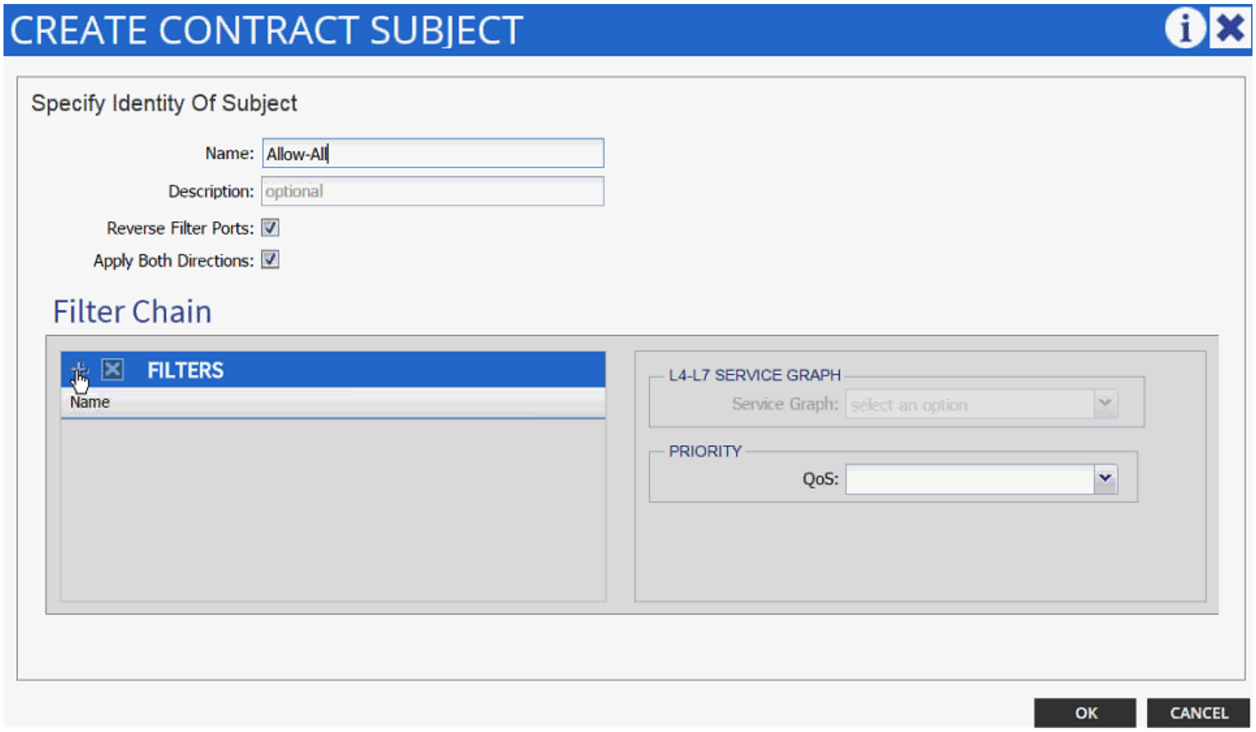

CREATE CONTRACT SUBJECT

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

FILTERS

Name

L4-L7 SERVICE GRAPH

Service Graph: select an option

PRIORITY

QoS:

25. Drop down the menu under FILTERS and select Allow-All filter under Tenant App-A.

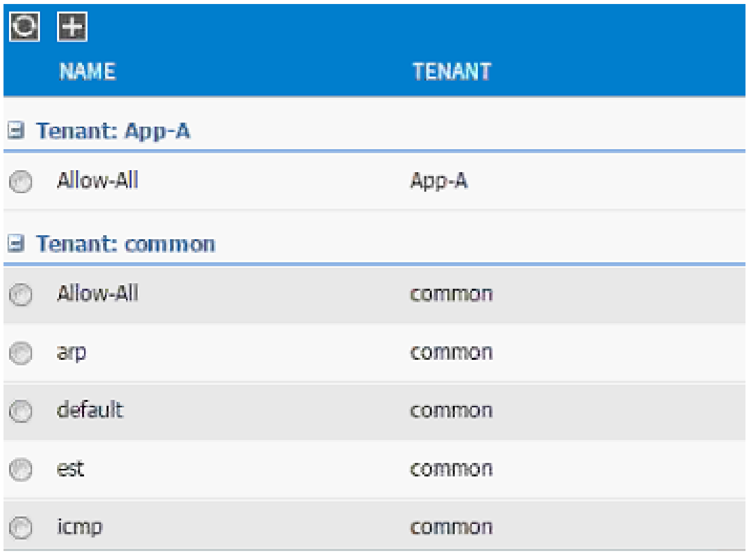

NAME	TENANT
Tenant: App-A	
○ Allow-All	App-A
Tenant: common	
○ Allow-All	common
○ arp	common
○ default	common
○ est	common
○ icmp	common

26. Click UPDATE to add the newly created filter to the filter chain.

27. Click OK to finish creating the Contract Subject.

28. Click SUBMIT.

29. Click SUBMIT to finish adding a provided contract.

Consuming the SVM Management Contract in an Application Tier (Web)

1. Expand Application Profile 3-Tier-App.

2. Expand Application EPGs.

3. Expand EPG Web.

4. Click Contracts in the left menu.

5. Click ACTIONS on the right and select Add Consumed Contract.

6. In the ADD CONSUMED CONTRACT dialog box, from the drop-down menu select App-A/Allow-SVM-Access contract (defined previously).

7. Click SUBMIT.

Cisco ACI – Connecting to Existing Infrastructure

This section provides a detailed procedure for a tenant (App-A) to existing Nexus 7000 core routers using sub-interfaces and VRF aware OSPF. Following are some of the highlights of this connectivity:

- A new bridge domain and associated private network is configured in ACI for external connectivity
- The Web VM is configured with two interfaces – one to connect to an EPG attached to inside bridge domain (bd-internal) and another to connect to an EPG on the external bridge domain (bd-external)*
- Each of the two Nexus 7000s is connected to each of Nexus 9000 leaf
- Sub-interfaces are configured and used for external connectivity
- Nexus 9000 is configured to run per-VRF OSPF – Nexus 7000 does not use VRFs
- Nexus 7000 is configured to originate and send a default route to Nexus 9000 leaves

Figure 6 shows the VLANs and networks used for this connectivity.

Note: When using service graphs (load balancer), the Web-VM does not need two separate interfaces.

Figure 6 ACI – Layer-3 Connectivity Details

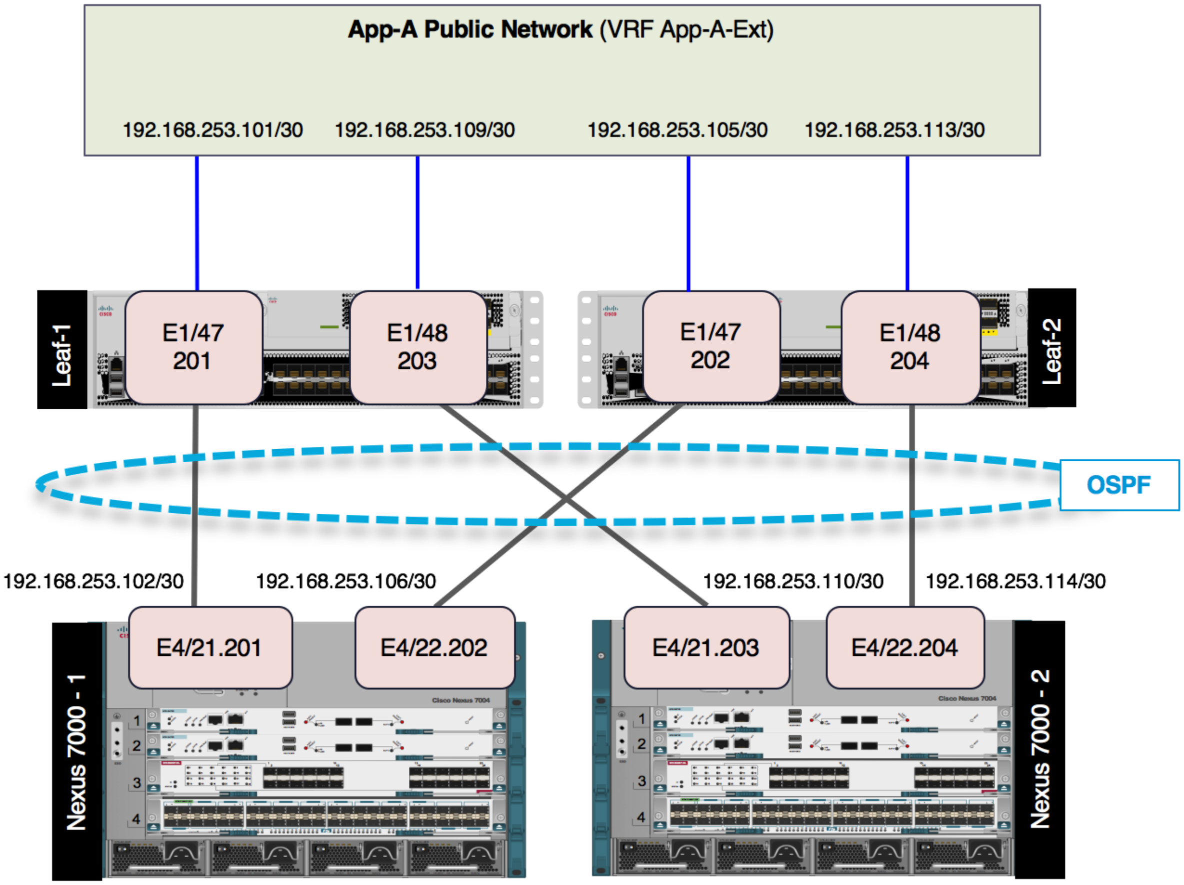

Configuring the Nexus 7000 for ACI Connectivity (Sample)

Nexus 7000-1

```
feature ospf
!
router ospf 10
  router-id 192.168.254.3
  area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
!
interface Vlan100
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.253.253/30
  no ipv6 redirects
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.0
!
interface Ethernet4/21.201
  encapsulation dot1q 201
  ip address 192.168.253.102/30
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown
!
interface Ethernet4/22.202
  encapsulation dot1q 202
  ip address 192.168.253.106/30
  ip ospf cost 5
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown
!
```

Nexus 7000-2

```
feature ospf
!
router ospf 10
  router-id 192.168.254.4
  area 0.0.0.10 nssa no-summary default-information-originate no-redistribution
!
interface Vlan100
  no shutdown
  mtu 9216
  no ip redirects
  ip address 192.168.253.254/30
  no ipv6 redirects
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.0
!
interface Ethernet4/21.203
  encapsulation dot1q 203
  ip address 192.168.253.110/30
  ip ospf cost 20
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown
!
interface Ethernet4/22.204
  encapsulation dot1q 204
  ip address 192.168.253.114/30
  ip ospf cost 30
  ip ospf mtu-ignore
  ip router ospf 10 area 0.0.0.10
  no shutdown
!
```

Configuring ACI for External Routed Domain

1. Select Tenant and App-A tenant from the top menu.
2. Expand Tenant App-A in the left menu bar.
3. Expand Networking, right click Bridge Domains and select Create Bridge Domain.
4. Use bd-external as Name of the bridge domain.
5. From the drop-down menu next to Network, select Create Private Network.
6. In the CREATE PRIVATE NETWORK dialog box, enter App-A-Ext for the Name.

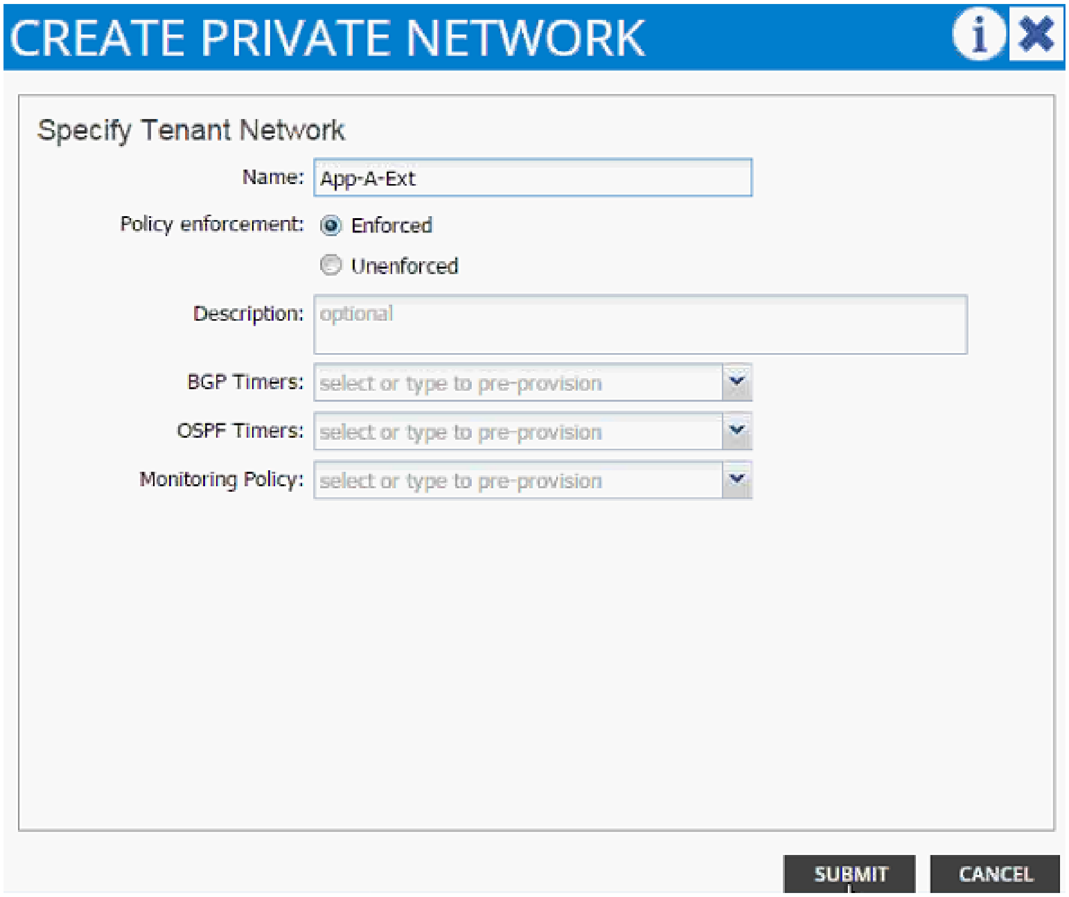

CREATE PRIVATE NETWORK

Specify Tenant Network

Name: App-A-Ext

Policy enforcement: ● Enforced ○ Unenforced

Description: optional

BGP Timers: select or type to pre-provision

OSPF Timers: select or type to pre-provision

Monitoring Policy: select or type to pre-provision

7. Click SUBMIT.
8. Select default as the IGMP Snoop Policy.
9. Click + next to Subnets to add a subnet.
10. Provide the gateway address of the subnet which will communicate to the external world.
11. Set Scope as Public.

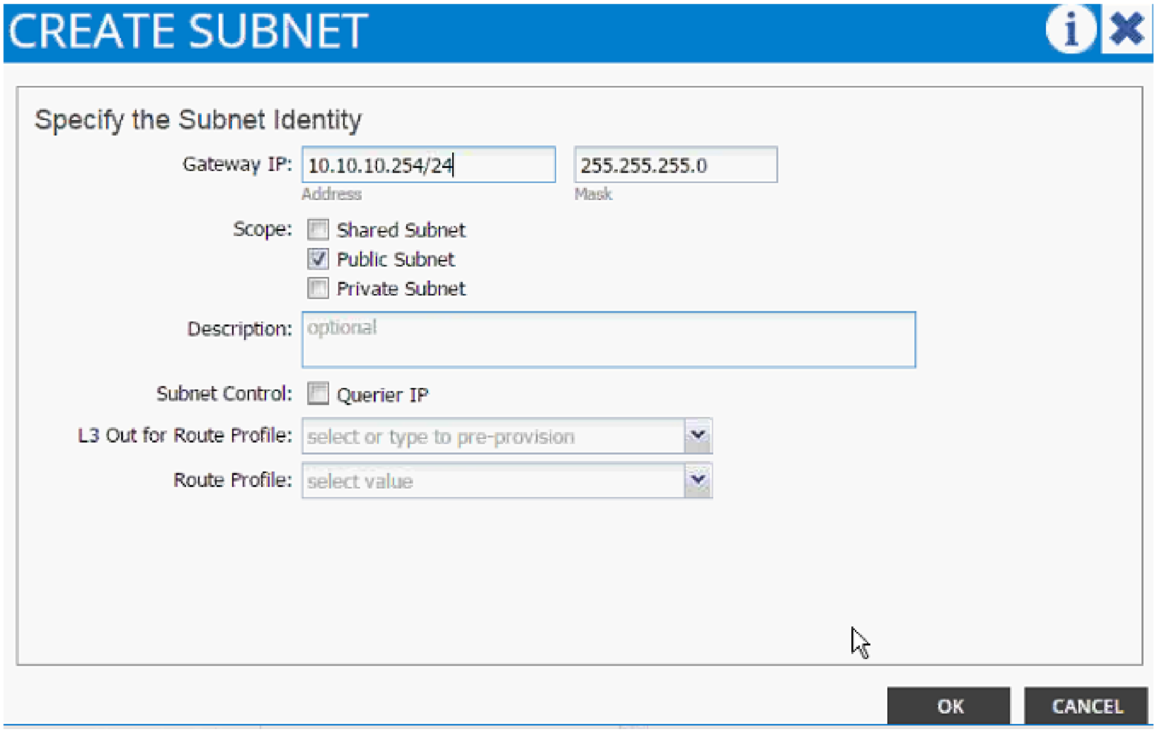

CREATE SUBNET

Specify the Subnet Identity

Gateway IP: 10.10.10.254/24 (Address) 255.255.255.0 (Mask)

Scope: □ Shared Subnet ☑ Public Subnet □ Private Subnet

Description: optional

Subnet Control: □ Querier IP

L3 Out for Route Profile: select or type to pre-provision

Route Profile: select value

12. Click OK.

13. Click SUBMIT.

14. Expand the Application Profile 3-Tier-App on the left and right click on Application EPGs and click Create Application EPG.

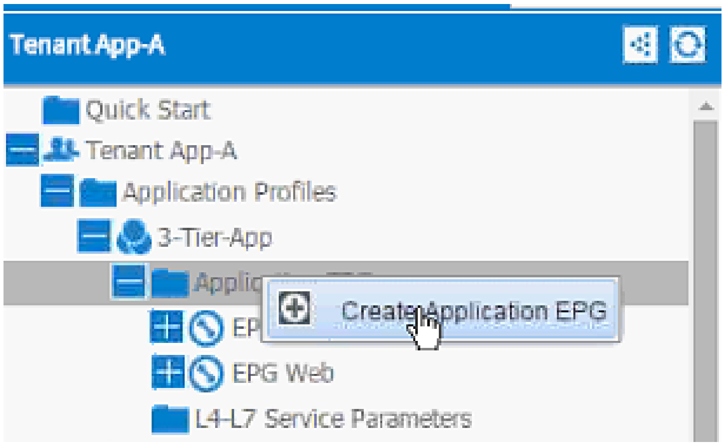

15. In the CREATE APPLICATION EPG dialog box, enter External as the Name.

16. From the drop-down menu, select bd-External as the Bridge Domain.

17. From the drop-down menu, select default for Monitoring Policy.

18. Click OK.

19. Click + next to Associated Domain Profiles (VMs or Bare metals).

20. From the drop-down menu, select the VMM domain previously defined.

21. Select Immediate for Deployment Immediacy.

22. Select Immediate for Resolution Immediacy.

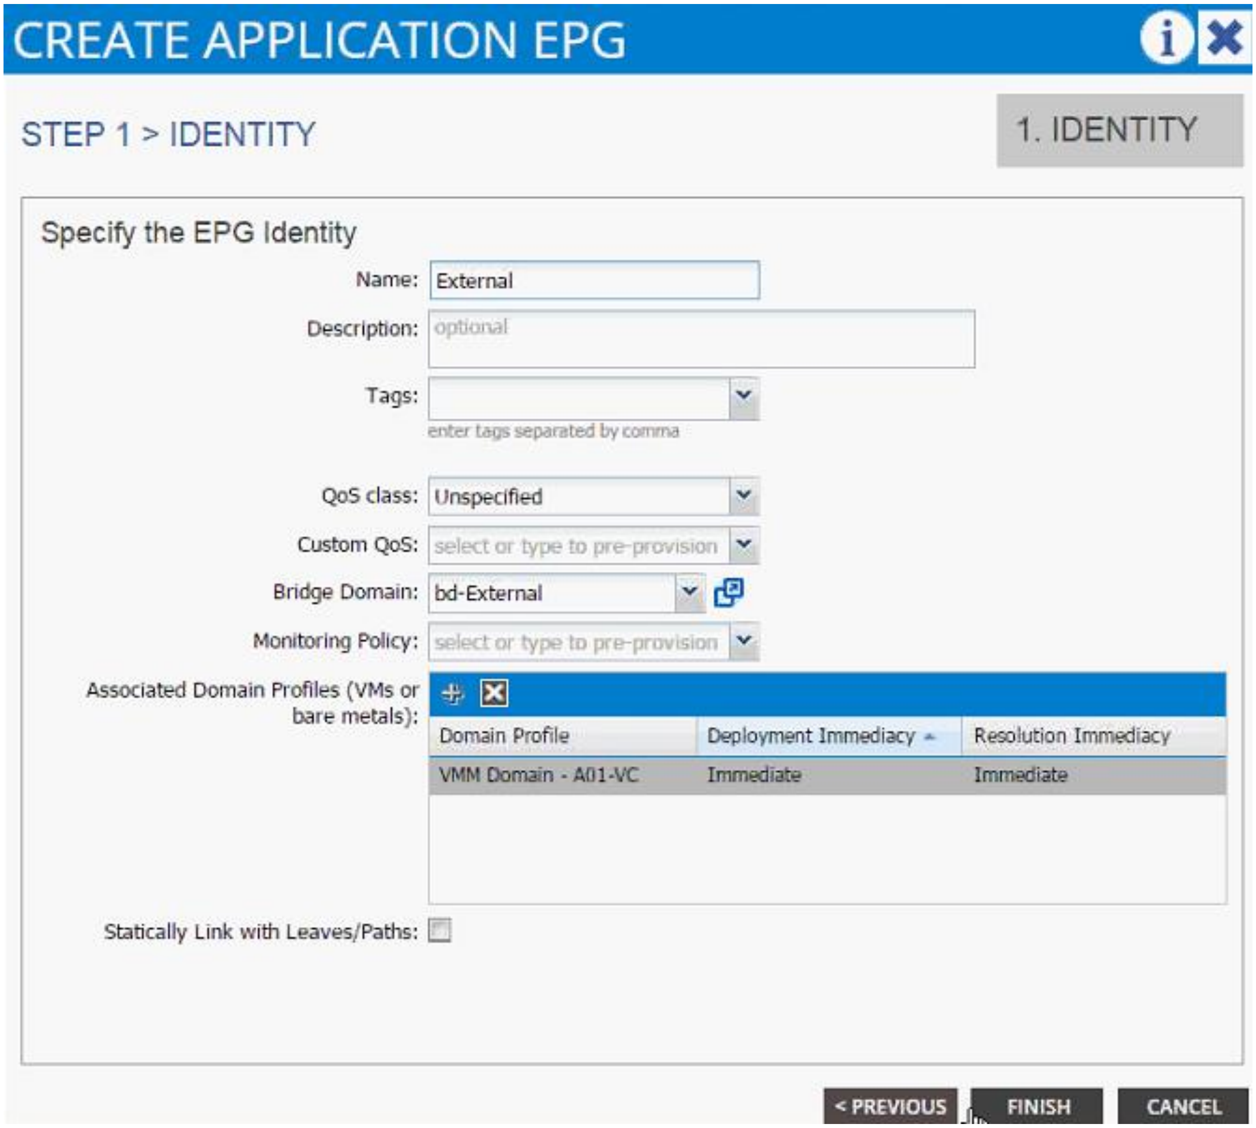

CREATE APPLICATION EPG

STEP 1 > IDENTITY

1. IDENTITY

Specify the EPG Identity

Name: External

Description: optional

Tags:

enter tags separated by comma

QoS class: Unspecified

Custom QoS: select or type to pre-provision

Bridge Domain: bd-External

Monitoring Policy: select or type to pre-provision

Associated Domain Profiles (VMs or bare metals):

Domain Profile	Deployment Immediacy	Resolution Immediacy
VMM Domain - A01-VC	Immediate	Immediate

Statically Link with Leaves/Paths: □

23. Click FINISH.

24. On the left menu, expand the EPG External and click Contracts.

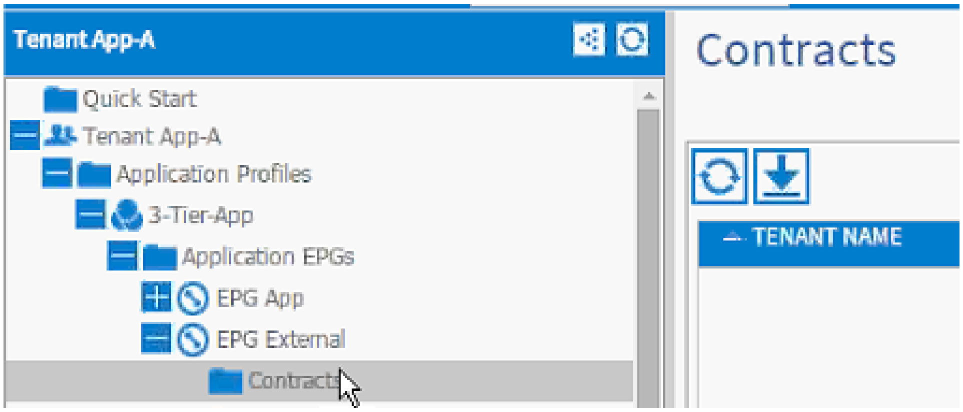

25. Click Action and select Add Provided Contract.

26. From the ADD PROVIDED CONTRACT dialog box, select Create Contract under Contract.

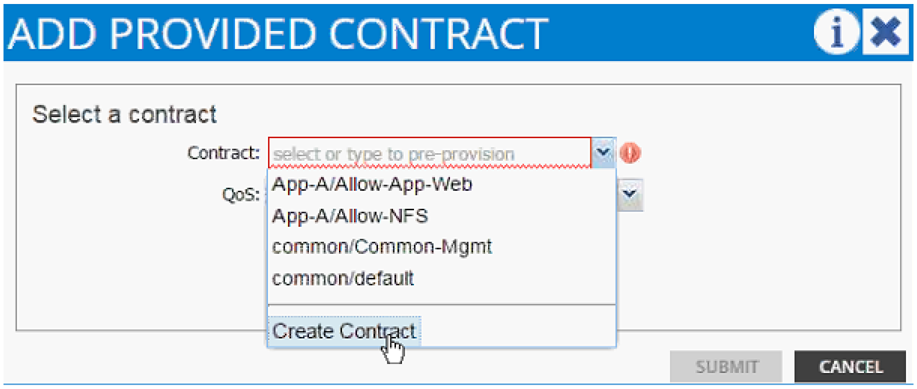

27. Enter Allow-External as Name in the CREATE CONTRACT dialog box.

28. Change the Scope to Tenant.

29. Click + next to Subjects to add a new contract subject.

30. In the CREATE CONTRACT SUBJECT dialog box, enter Allow-All as the Name.

31. Click + under Filter Chain to add a new filter.

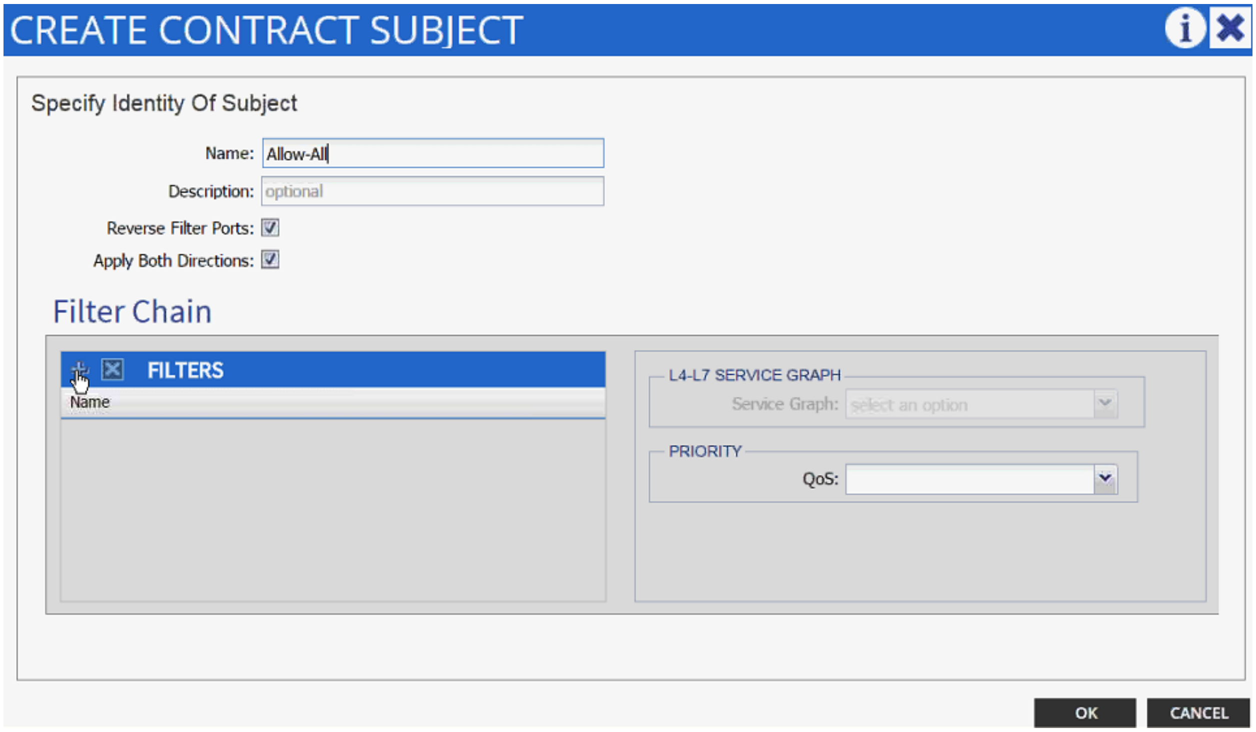

Specify Identity Of Subject

Name: Allow-All

Description: optional

Reverse Filter Ports: ☑

Apply Both Directions: ☑

Filter Chain

Name

Service Graph: select an option

QoS:

32. Drop down the menu under FILTERS and select Allow-All filter under Tenant App-A.

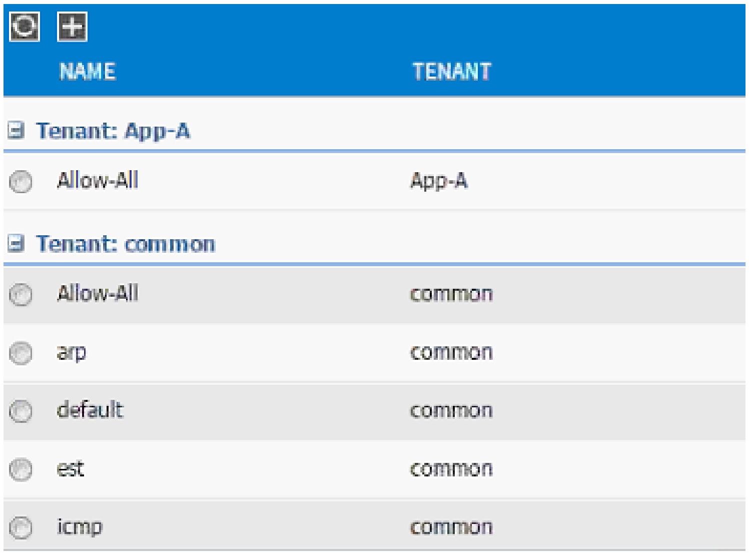

NAME	TENANT
Tenant: App-A	
○ Allow-All	App-A
Tenant: common	
○ Allow-All	common
○ arp	common
○ default	common
○ est	common
○ icmp	common

33. Click UPDATE to add the newly created filter to the filter chain.

34. Click OK to finish creating the Contract Subject.

35. Click SUBMIT.

36. Click SUBMIT again to finish adding a provided contract.

Adding External Routed Domain

1. Select Tenants and App-A from the top menu.

2. Expand Tenant App-A in the left menu.

3. Expand Networking and External Routed Networks.

4. Right-clickExternal Routed Networks and select Create Routed Outside.

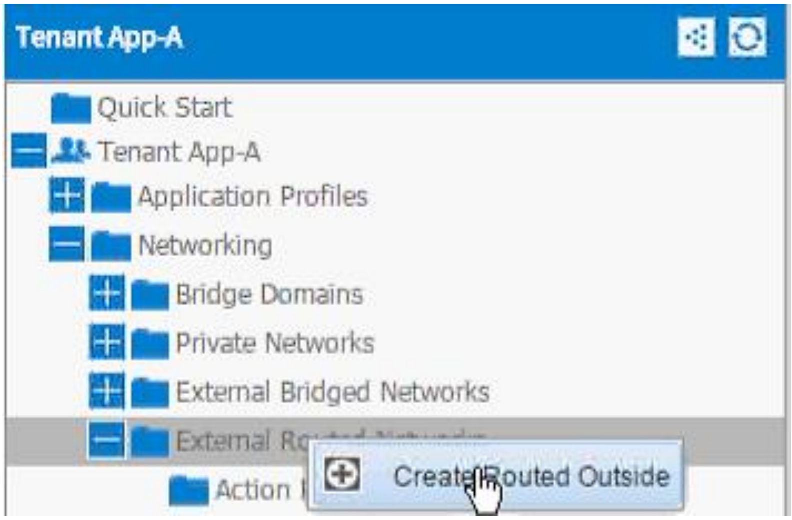

5. In the CREATE ROUTED OUTSIDE dialog box, enter App-A-L3-Out as the Name.

6. Click the check mark next to OSPF.

7. Enter 0.0.0.10 as the OSPF Area ID.

8. Select App-A-Ext from the drop-down menu as the Private Network.

9. Click + next to NODES AND INTERFACES PROTOCOL PROFILES.

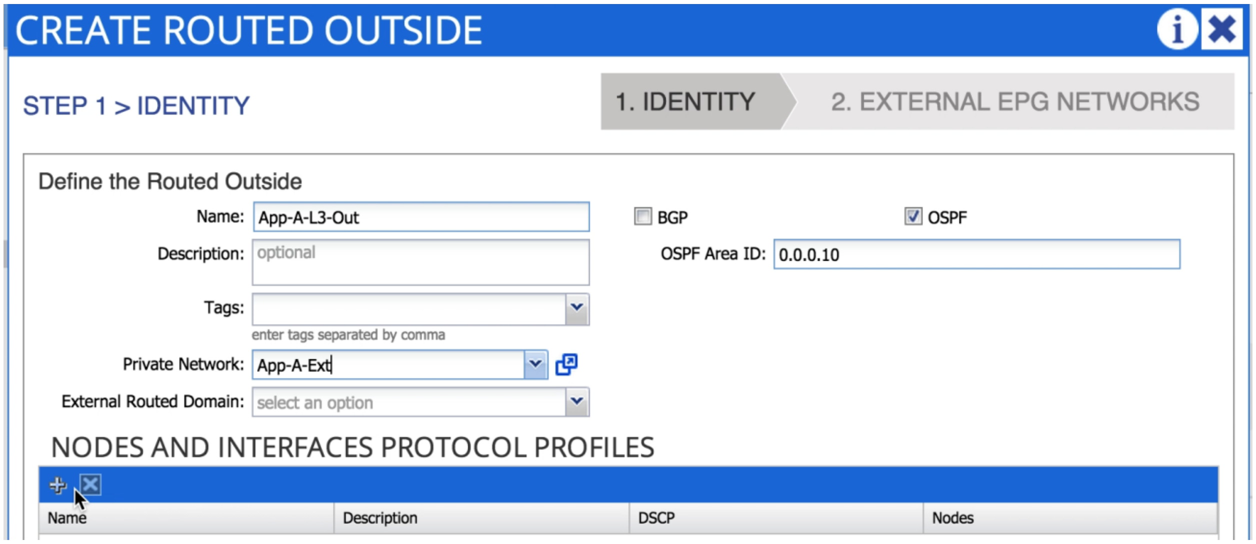
CREATE ROUTED OUTSIDE

STEP 1 > IDENTITY

1. IDENTITY | 2. EXTERNAL EPG NETWORKS

Define the Routed Outside

Name: App-A-L3-Out

Description: optional

Tags:

enter tags separated by comma

Private Network: App-A-Ext

External Routed Domain: select an option

□ BGP ☑ OSPF

OSPF Area ID: 0.0.0.10

NODES AND INTERFACES PROTOCOL PROFILES

Name	Description	DSCP	Nodes

10. In the CREATE NODE PROFILE dialog box, enter Node_101 as the Name.

11. Click + to add Nodes.

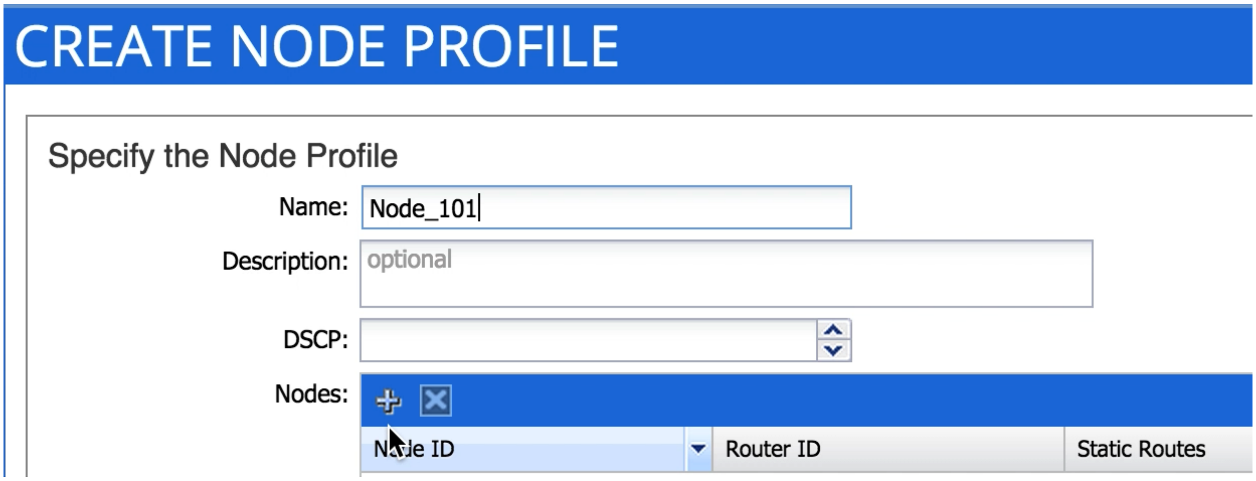

CREATE NODE PROFILE

Specify the Node Profile

Name: Node_101

Description: optional

DSCP:

Nodes:

Node ID	Router ID	Static Routes

12. In the SELECT NODE dialog box, select first Nexus 9000 switch (Node-101).

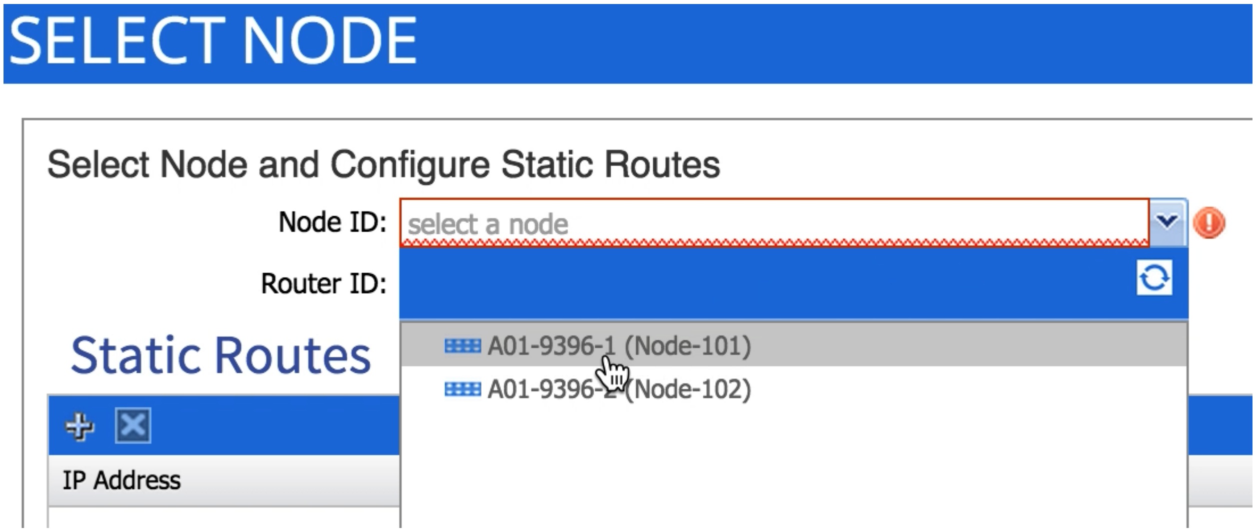

SELECT NODE

Select Node and Configure Static Routes

Node ID: select a node

Router ID:

A01-9396-1 (Node-101)

A01-9396-2 (Node-102)

Static Routes

IP Address

13. Provide a loopback address to be used for the tenant private network.

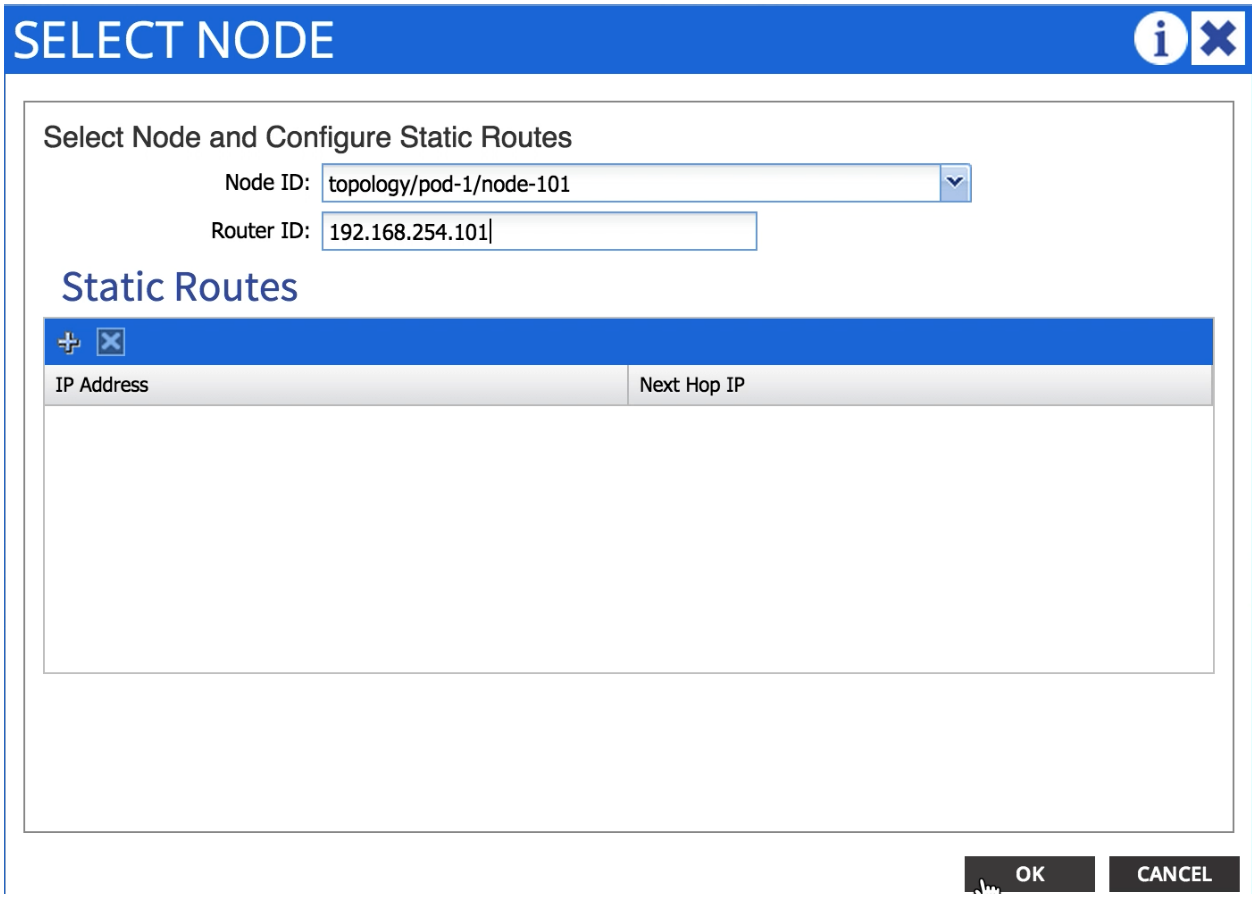

SELECT NODE

Select Node and Configure Static Routes

Node ID: topology/pod-1/node-101

Router ID: 192.168.254.101

Static Routes

IP Address	Next Hop IP

14. Click OK.

15. Click + to add OSPF INTERFACE PROFILE.

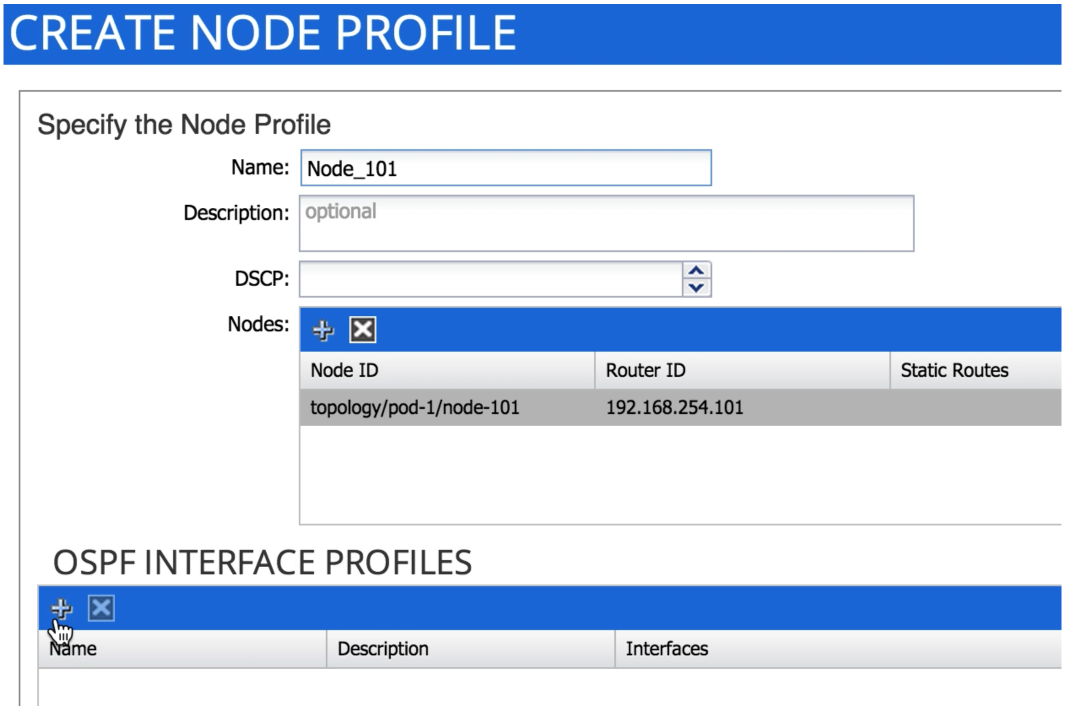

CREATE NODE PROFILE

Specify the Node Profile

Name: Node_101

Description: optional

DSCP:

Nodes:

Node ID	Router ID	Static Routes
topology/pod-1/node-101	192.168.254.101	

OSPF INTERFACE PROFILES

Name	Description	Interfaces

16. In the CREATE INTERFACE PROFILE dialog box, use a descriptive name (Node101_Int_Profile in this example provides information about interface policy for Node 101).

17. Drop down the OSPF Policy menu and select Create OSPF Interface Policy.

18. In the CREATE OSPF INTERFACE POLICY dialog box, provide a descriptive name.

19. Check MTU ignore and Advertise Subnet.

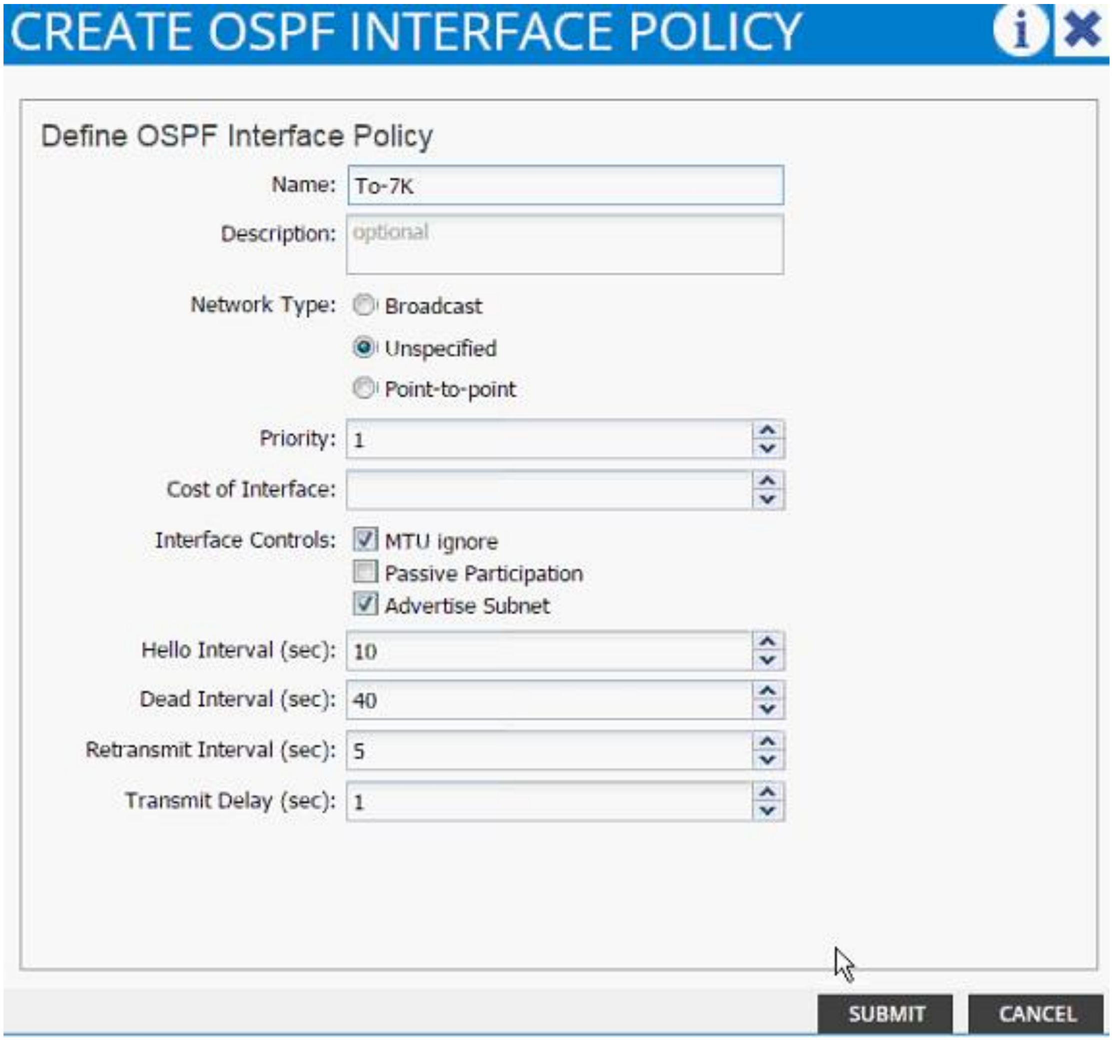

CREATE OSPF INTERFACE POLICY

Define OSPF Interface Policy

Name: To-7K

Description: optional

Network Type: ○ Broadcast ● Unspecified ○ Point-to-point

Priority: 1

Cost of Interface:

Interface Controls: ☑ MTU ignore □ Passive Participation ☑ Advertise Subnet

Hello Interval (sec): 10

Dead Interval (sec): 40

Retransmit Interval (sec): 5

Transmit Delay (sec): 1

20. Click SUBMIT.

21. Under INTERFACES, select ROUTED SUB-INTERFACE and click + to add a new sub-interface.

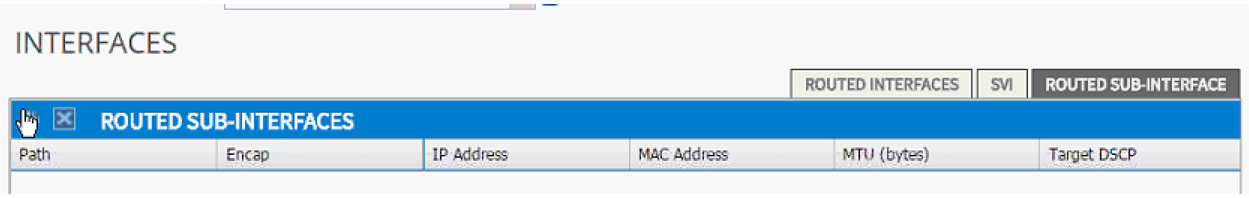

INTERFACES

ROUTED SUB-INTERFACES

Path	Encap	IP Address	MAC Address	MTU (bytes)	Target DSCP

22. In the SELECT ROUTED SUB-INTERFACE dialog box, select the Nexus 9000-1 interface connected to Nexus 7000-1 (Eth 1/47).

23. In the Encap enter vlan-201.

24. Provide the IP address 192.168.253.101/30.

25. Set MTU to 1500.

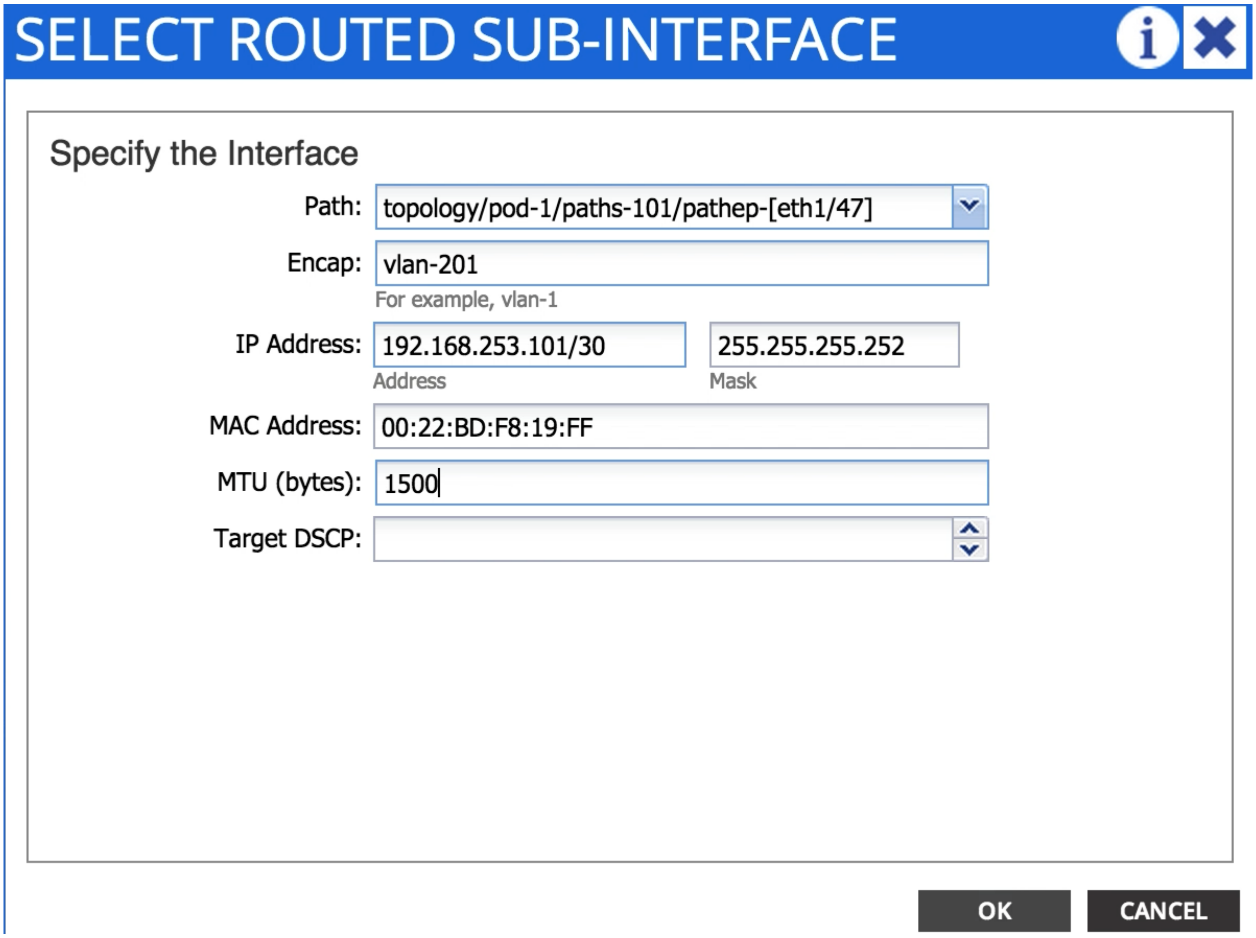

SELECT ROUTED SUB-INTERFACE

Specify the Interface

Path: topology/pod-1/paths-101/pathep-[eth1/47]

Encap: vlan-201

For example, vlan-1

IP Address: 192.168.253.101/30 (Address) 255.255.255.252 (Mask)

MAC Address: 00:22:BD:F8:19:FF

MTU (bytes): 1500

Target DSCP:

26. Click OK.

27. Repeat steps 21-26 to add a second interface profile with appropriate sub-interfaces. The OSPF Policy created in steps 18-20 can be re-used (select from the drop-down list) for the interface profiles.

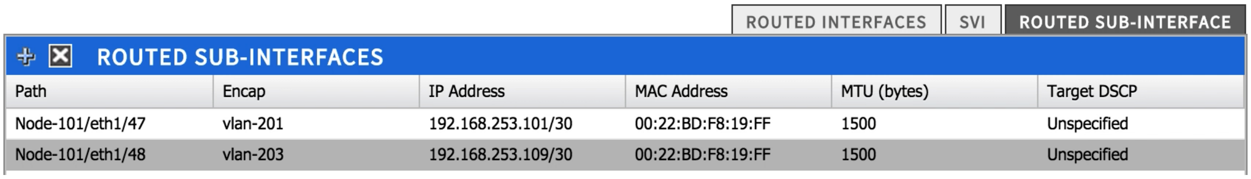

ROUTED SUB-INTERFACES

Path	Encap	IP Address	MAC Address	MTU (bytes)	Target DSCP
Node-101/eth1/47	vlan-201	192.168.253.101/30	00:22:BD:F8:19:FF	1500	Unspecified
Node-101/eth1/48	vlan-203	192.168.253.109/30	00:22:BD:F8:19:FF	1500	Unspecified

28. Click OK.

29. Click + next to NODES AND INTERFACES PROTOCOL PROFILES.

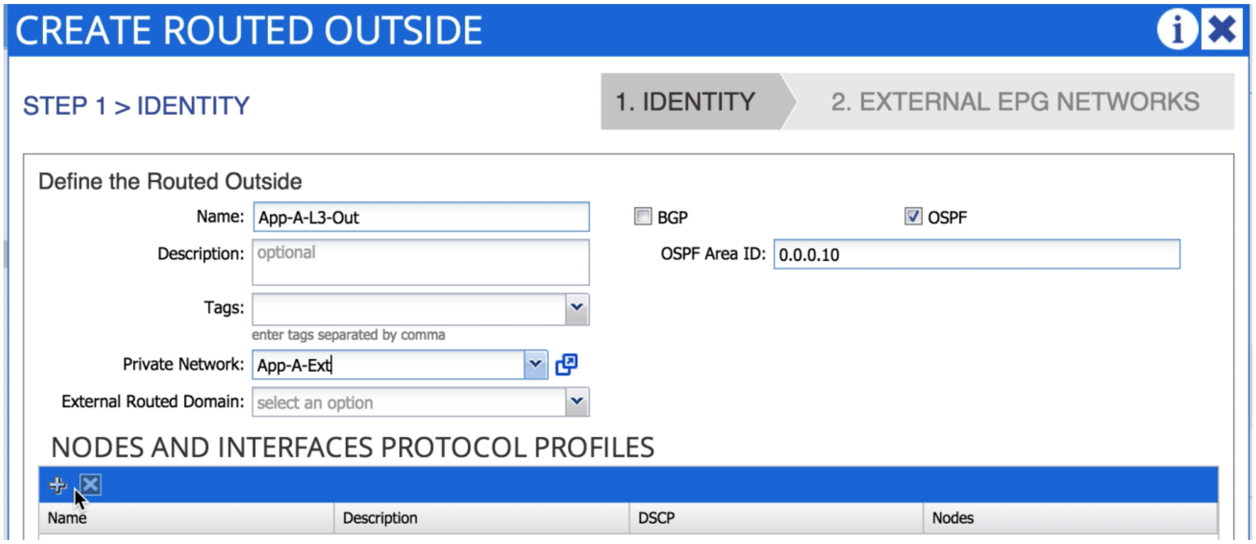

CREATE ROUTED OUTSIDE

STEP 1 > IDENTITY

1. IDENTITY | 2. EXTERNAL EPG NETWORKS

Define the Routed Outside

Name: App-A-L3-Out

Description: optional

Tags:

enter tags separated by comma

Private Network: App-A-Ext

External Routed Domain: select an option

□ BGP ☑ OSPF

OSPF Area ID: 0.0.0.10

NODES AND INTERFACES PROTOCOL PROFILES

Name	Description	DSCP	Nodes

30. In the CREATE NODE PROFILE dialog box, enter Node_102 as the Name.

31. Click + to add Nodes.

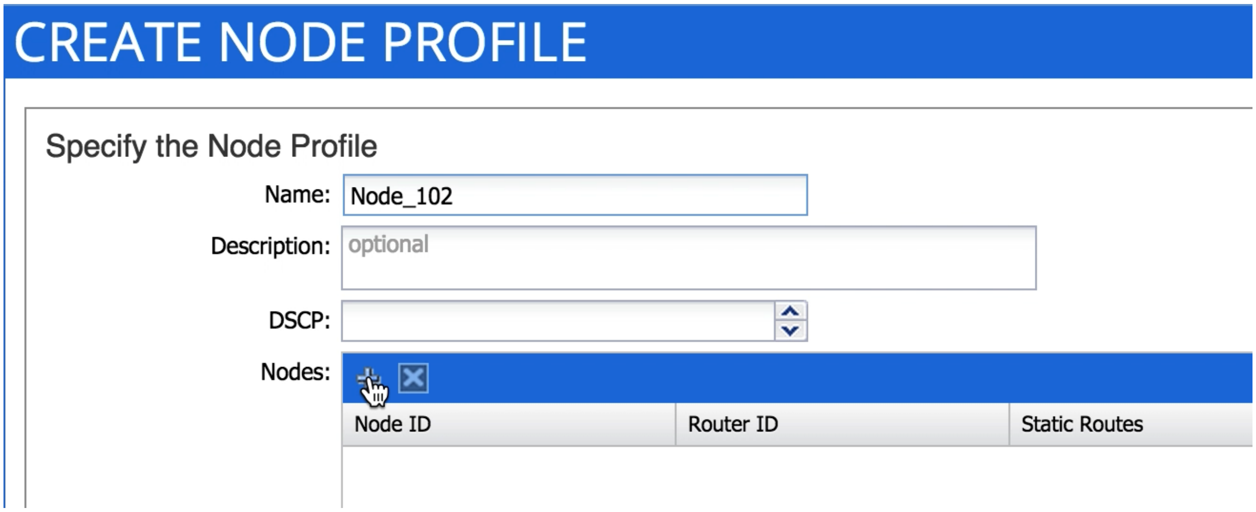

CREATE NODE PROFILE

Specify the Node Profile

Name: Node_102

Description: optional

DSCP:

Nodes:

Node ID	Router ID	Static Routes

32. In the SELECT NODE dialog box, select second Nexus 9000 switch (Node-102).

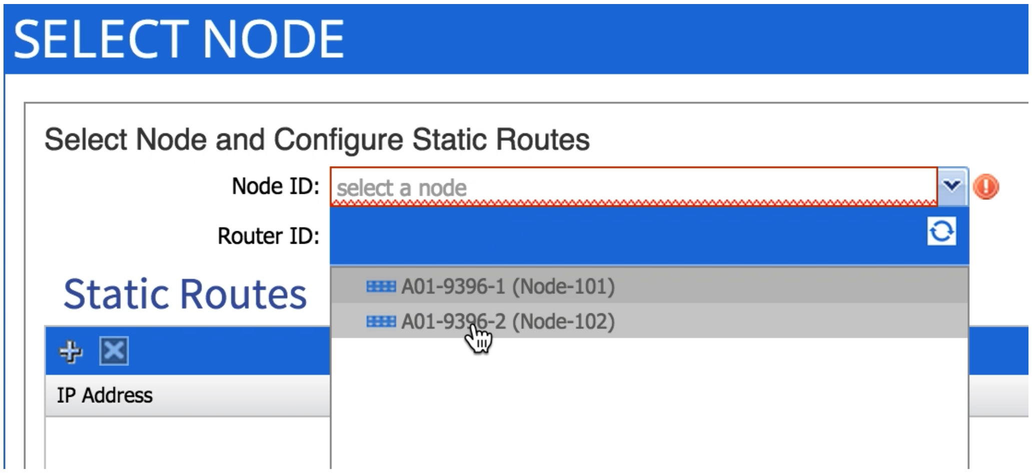

SELECT NODE

Select Node and Configure Static Routes

Node ID: select a node

Router ID:

A01-9396-1 (Node-101)
A01-9396-2 (Node-102)

Static Routes

IP Address

33. Provide a loopback address to be used for the tenant private network.

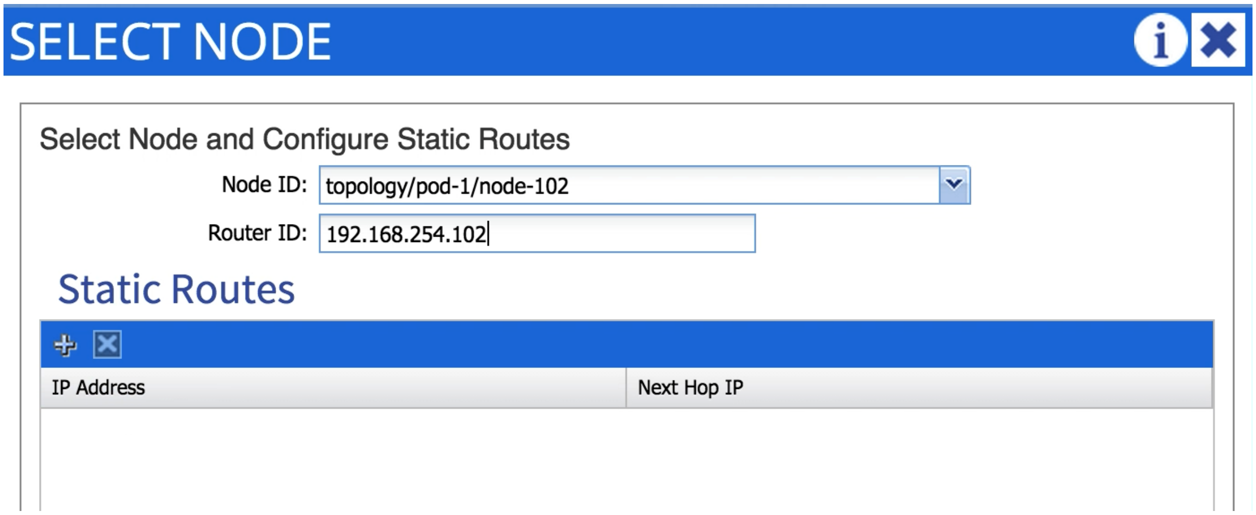

SELECT NODE

Select Node and Configure Static Routes

Node ID: topology/pod-1/node-102

Router ID: 192.168.254.102

Static Routes

IP Address	Next Hop IP

34. Click OK.

35. Click + to add OSPF INTERFACE PROFILE.

36. In the CREATE INTERFACE PROFILE dialog box, use a descriptive name (Node102_Int_Profile in this example provides information about interface policy for Node 101).

37. Drop down the OSPF Policy menu and select previously created OSPF policy.

38. Click SUBMIT.

39. Under INTERFACES, select ROUTED SUB-INTERFACE and click + to add a new sub-interface.

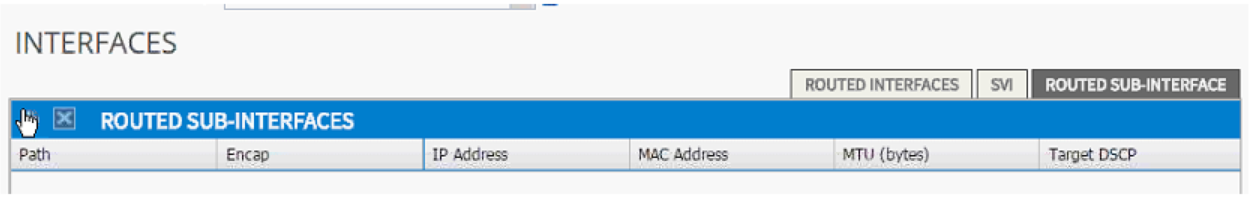

INTERFACES

ROUTED SUB-INTERFACES

Path	Encap	IP Address	MAC Address	MTU (bytes)	Target DSCP

40. In the SELECT ROUTED SUB-INTERFACE dialog box, select the Nexus 9000-2 interface connected to Nexus 7000-1 (Eth 1/47).

41. In the Encap enter vlan-202.

42. Provide the IP address.

43. Set MTU to 1500.

44. Click OK.

45. Repeat these steps to add second interface profiles with appropriate sub-interfaces.

46. Click OK.

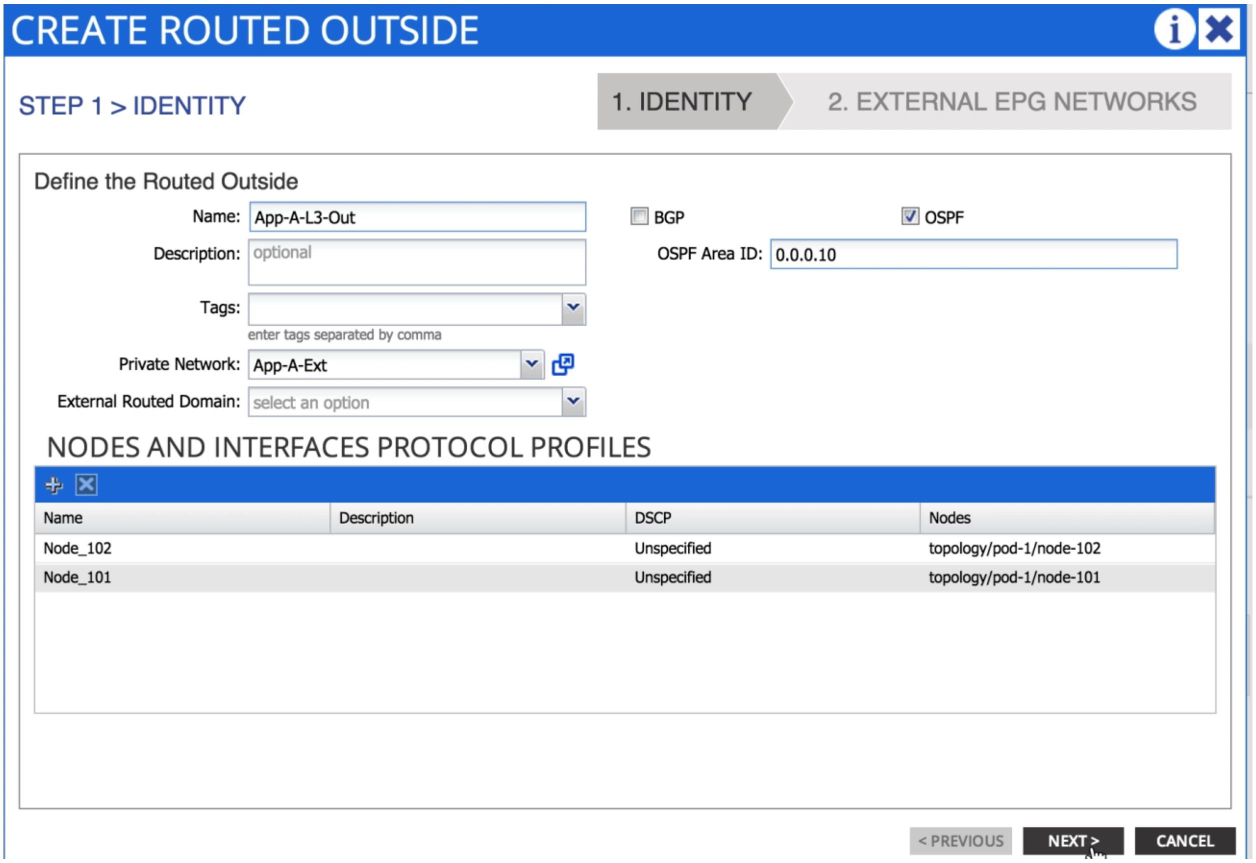

CREATE ROUTED OUTSIDE

STEP 1 > IDENTITY

Define the Routed Outside

Name: App-A-L3-Out

Description: optional

Tags:

enter tags separated by comma

Private Network: App-A-Ext

External Routed Domain: select an option

□ BGP ☑ OSPF

OSPF Area ID: 0.0.0.10

NODES AND INTERFACES PROTOCOL PROFILES

Name	Description	DSCP	Nodes
Node_102		Unspecified	topology/pod-1/node-102
Node_101		Unspecified	topology/pod-1/node-101

47. Click Next.

48. Click + to add EXTERNAL EPG NETWORKS.

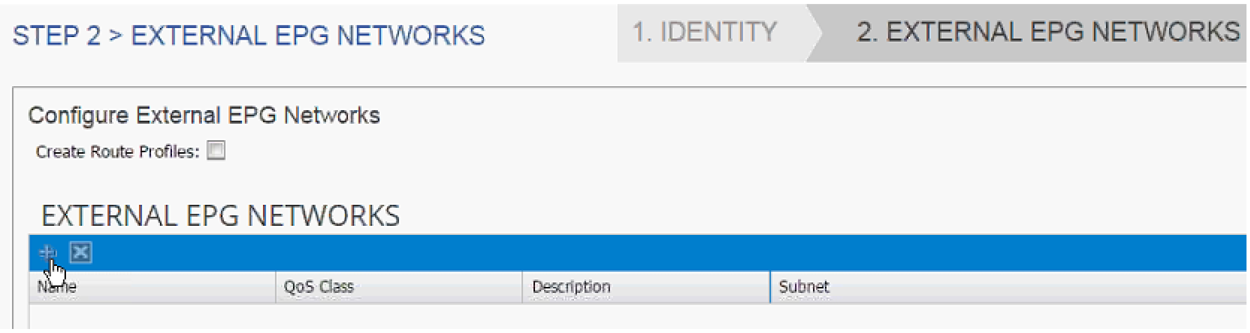

STEP 2 > EXTERNAL EPG NETWORKS

Configure External EPG Networks

Create Route Profiles: □

EXTERNAL EPG NETWORKS

Name	QoS Class	Description	Subnet

49. In the CREATE EXTERNAL NETWORK dialog box, enter Ext-Default as the Name.

50. Click + next to SUBNET to add remote subnet that will be accessed from the ACI fabric.

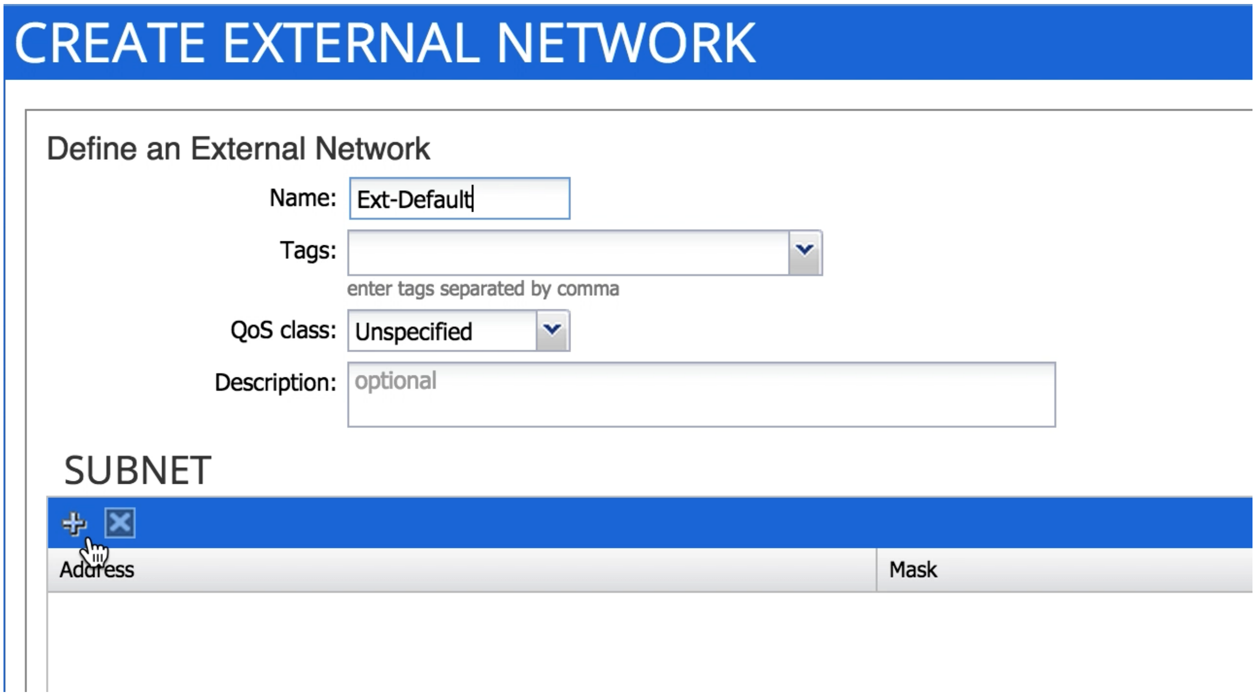

CREATE EXTERNAL NETWORK

Define an External Network

Name: Ext-Default

Tags:

enter tags separated by comma

QoS class: Unspecified

Description: optional

SUBNET

Address	Mask

51. In the CREATE SUBNET dialog box, enter 0.0.0.0/0 as the remote subnet to be accessed.

52. Click OK.

53. Click OK.

54. Click FINISH.

55. Expand the App-A-L3-Out routed network and click on the newly created network Ext-Default.

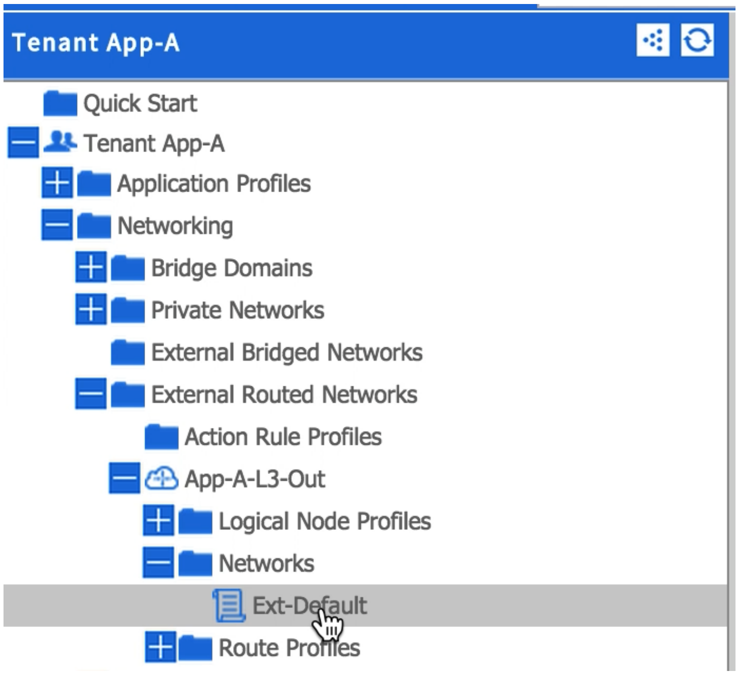

56. On the right hand side, click + next to Consumed Contracts.

57. From the drop-down menu, select App-A/Allow-External contract.

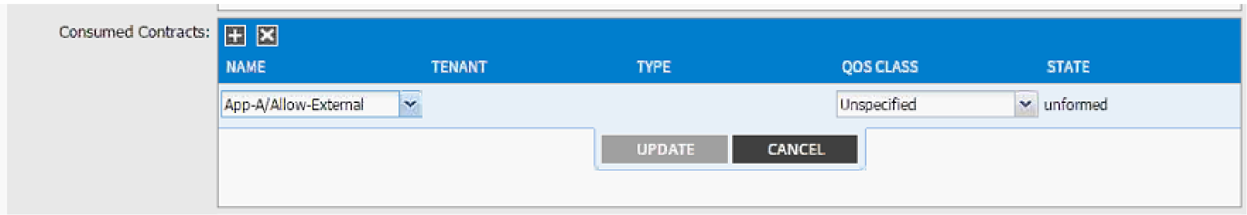

Consumed Contracts:

NAME	TENANT	TYPE	QOS CLASS	STATE
App-A/Allow-External			Unspecified	unformed

58. Click UPDATE.

59. Click SUBMIT.

60. Expand Bridge Domains in the menu on the left.

61. Click bd-External and click + next to Associated L3 Outs.

62. From the drop-down menu, select the newly created routed domain App-A/App-A-L3-Out.

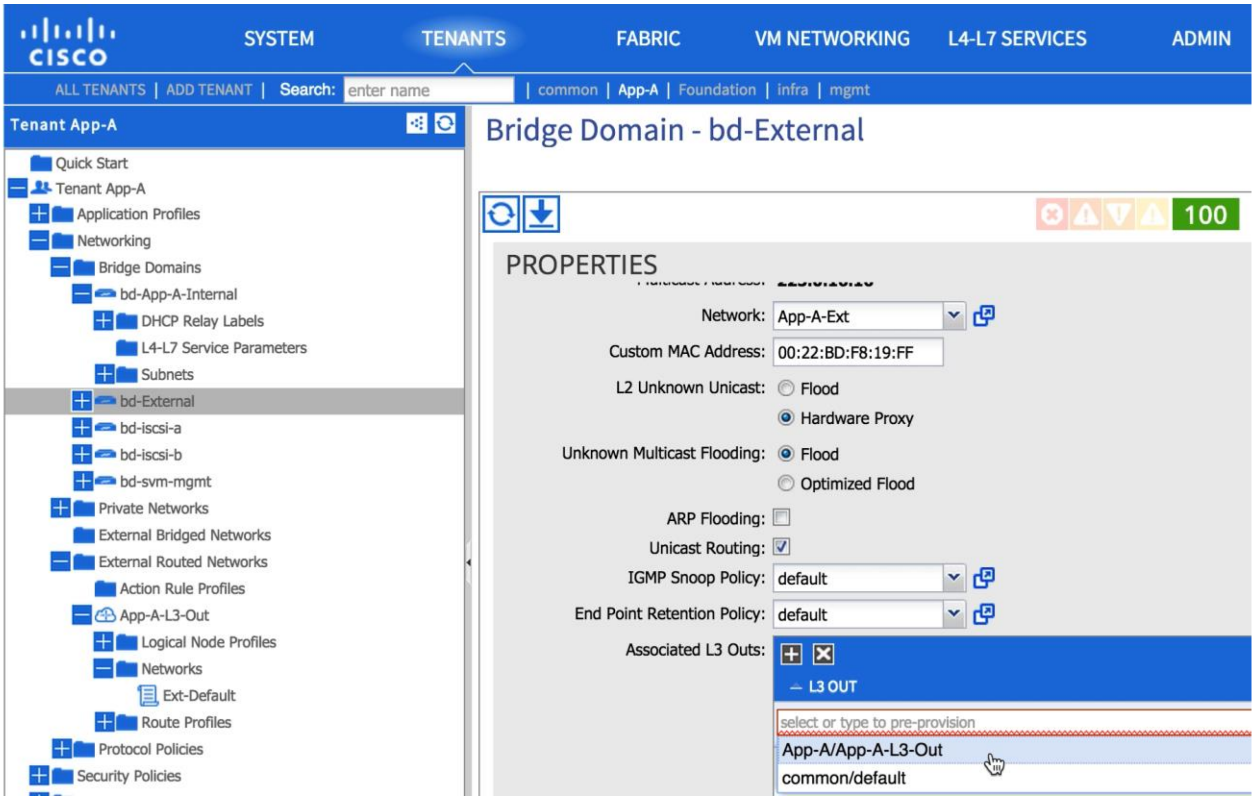

Bridge Domain - bd-External

PROPERTIES

Network: App-A-Ext

Custom MAC Address: 00:22:BD:F8:19:FF

L2 Unknown Unicast: ○ Flood ● Hardware Proxy

Unknown Multicast Flooding: ● Flood ○ Optimized Flood

ARP Flooding: □

Unicast Routing: ☑

IGMP Snoop Policy: default

End Point Retention Policy: default

Associated L3 Outs:

L3 OUT

select or type to pre-provision

App-A/App-A-L3-Out

common/default

63. Click UPDATE.

64. Click SUBMIT.

At this time, the network, 10.10.10.0/24, marked public under bridge domain bd-External should be visible in the Nexus 7000 routing table and Nexus 9000 should learn the OSPF routes including default route from Nexus 7000. The provider and consumed contracts should enable communication from a VM connected to External EPG (port-group) and any subnet on the external network.

Configuring Multi-Protocol BGP on Spines

Within the ACI fabric, Multiprotocol BGP (MP-BGP) is implemented between leaf and spine switches to propagate external routes within the ACI fabric. The BGP route reflector technology is deployed in order to support a large number of leaf switches within a single fabric. All of the leaf and spine switches are in one single BGP autonomous system (AS). When the border leaf learns the external routes, it can then redistribute the external routes of a given VRF to an MP-BGP address family VPN version 4 (or VPN version 6 when IPv6 routing is supported in ACI). With address family VPN version 4, MP-BGP maintains a separate BGP routing table for each VRF. Within MP-BGP, the border leaf advertises routes to a spine switch, which is a BGP route reflector.

MP-BGP is not enabled by default in the ACI fabric. For the deployment scenario where ACI fabric is used as L2 fabric or there is no need for L3 outside connection, MP-BGP is not required. To enable MP-BGP please configure BGP policy on the APIC to specify the BGP ASN and specify spine nodes as BGP route reflectors.

1. Select Fabric and Fabric Policies from the top menu.

2. Expand Pod Policies and Policies in the left menu bar.

3. Click BGP Route Reflector default.

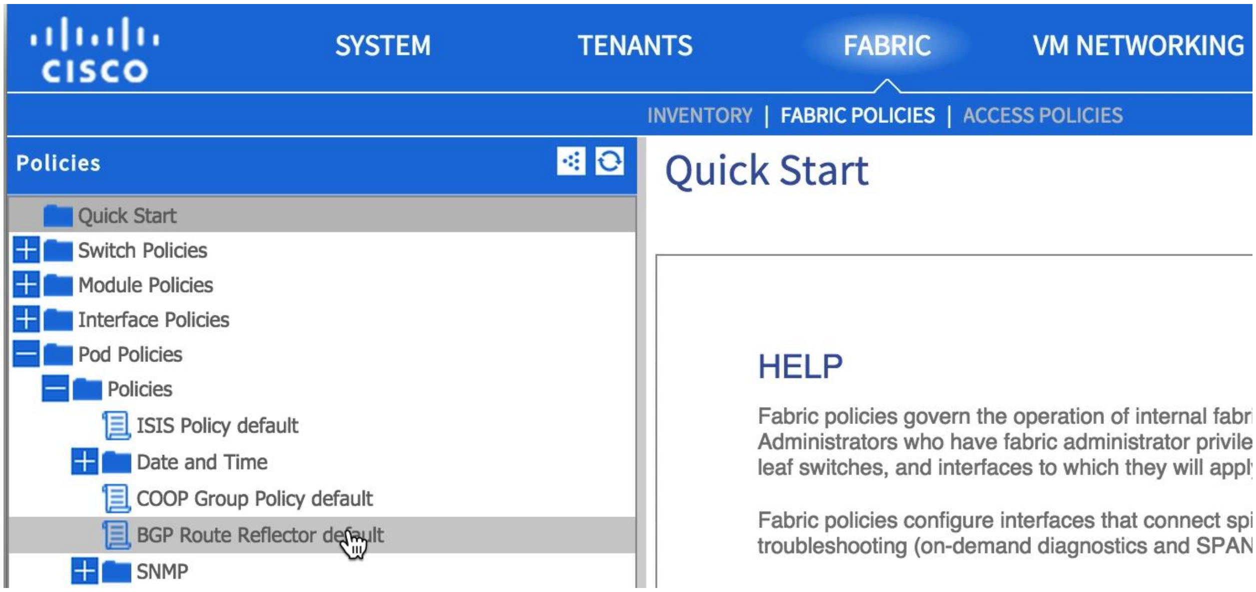

Fabric policies govern the operation of internal fabri
Administrators who have fabric administrator privile
leaf switches, and interfaces to which they will appl

Fabric policies configure interfaces that connect spi
troubleshooting (on-demand diagnostics and SPAN

4. Enter an Autonomous System Number (100 in this example).

5. Click + to select spines (one after the other) as Route Reflector Nodes.

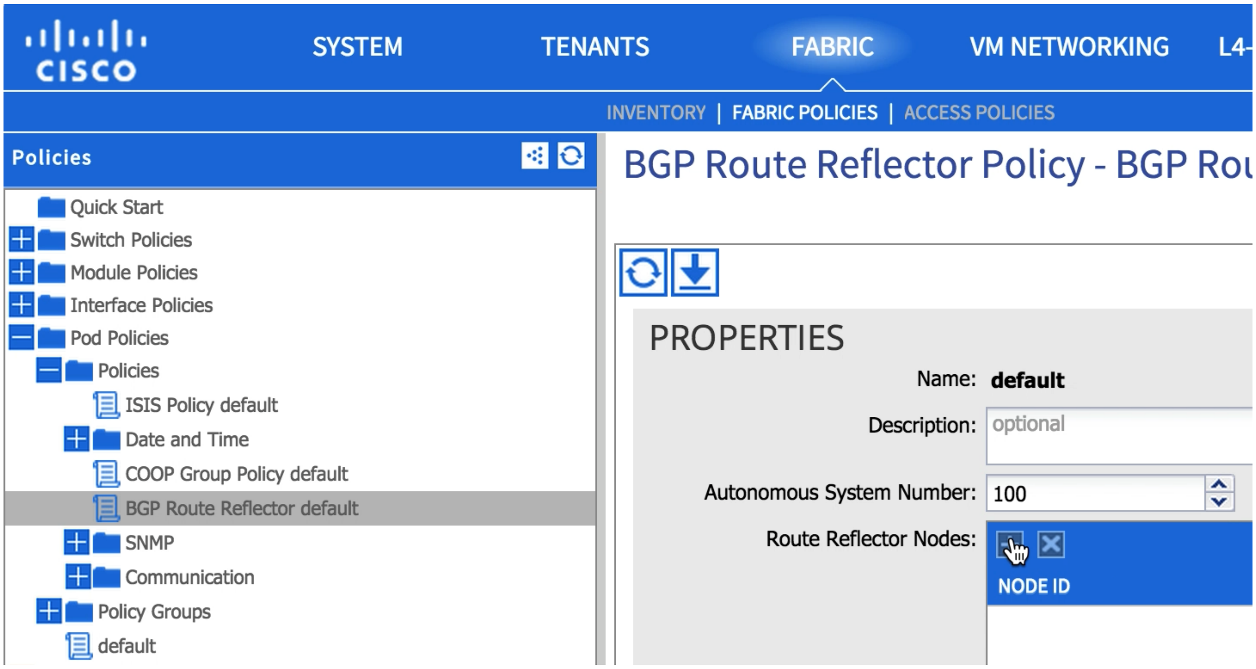

Name: **default**

Description: optional

Autonomous System Number: 100

Route Reflector Nodes:

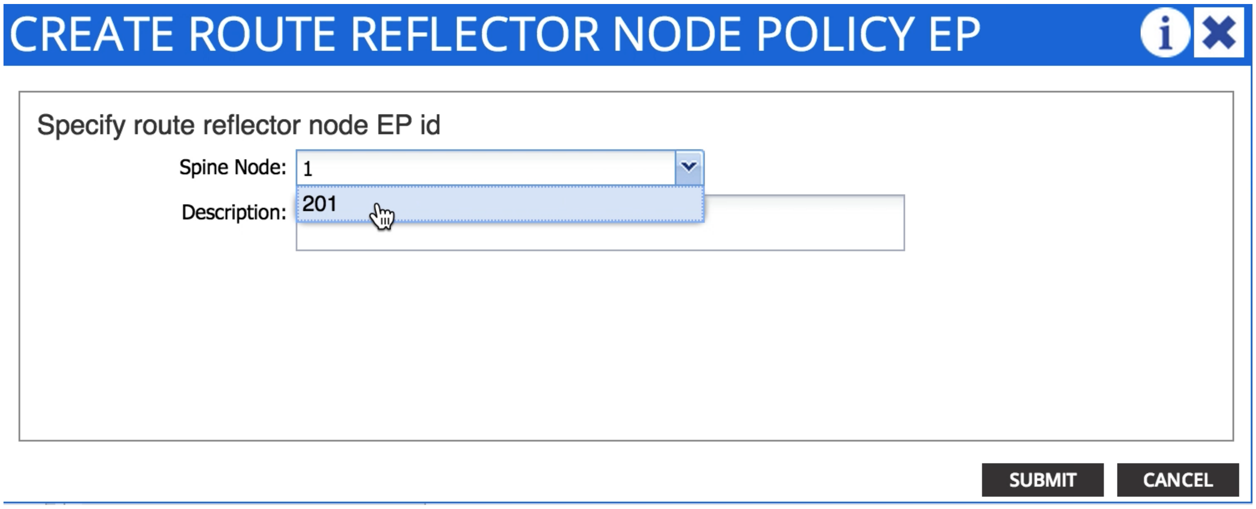

CREATE ROUTE REFLECTOR NODE POLICY EP

Specify route reflector node EP id

Spine Node: 1

201

Description:

6. Click SUBMIT.

7. Right-click Policy Groups in left menu bar and select Create POD Policy Group.

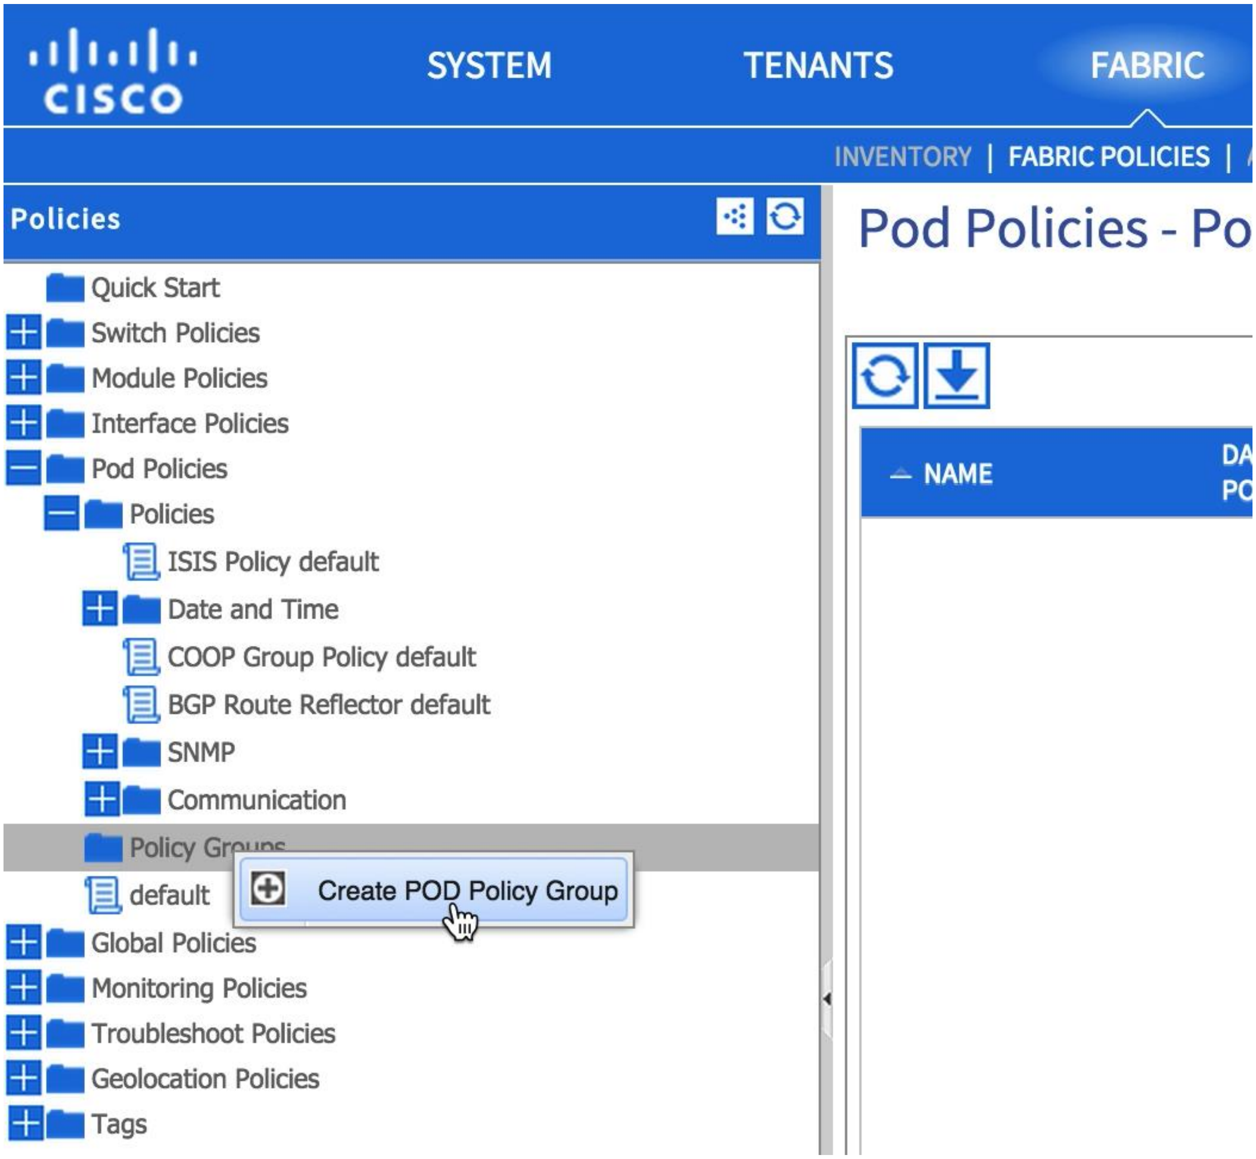

8. Enter a descriptive name for Pod1 policy.

9. From the drop-down menu BGP Route Reflector Policy, select default.

Note: The policy group allows users to combine multiple policies, such as BGP policy, Integrated System to Integrated System (IS-IS) routing protocol policy, co-operative (COOP) policy, and others, to a policy group and apply it to the POD. In this example, the policy is being utilized to only make changes to BGP.

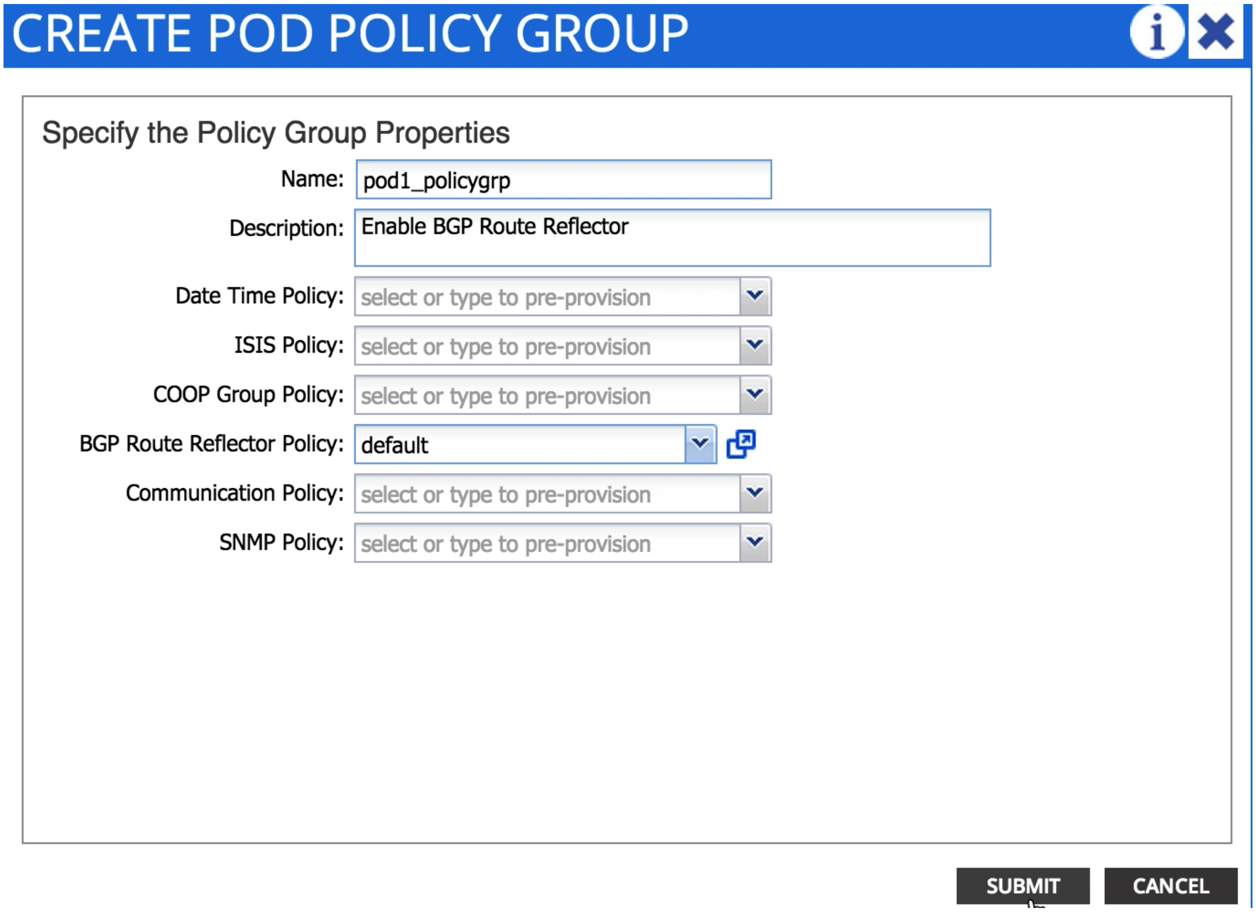

CREATE POD POLICY GROUP

Specify the Policy Group Properties

Name: pod1_policygrp

Description: Enable BGP Route Reflector

Date Time Policy: select or type to pre-provision

ISIS Policy: select or type to pre-provision

COOP Group Policy: select or type to pre-provision

BGP Route Reflector Policy: default

Communication Policy: select or type to pre-provision

SNMP Policy: select or type to pre-provision

10. Click SUBMIT.

11. Click default from Policy in the left menu bar.

12. From the drop-down menu Fabric Policy Group, select the recently created Pod policy (pod1_policygrp in this example).

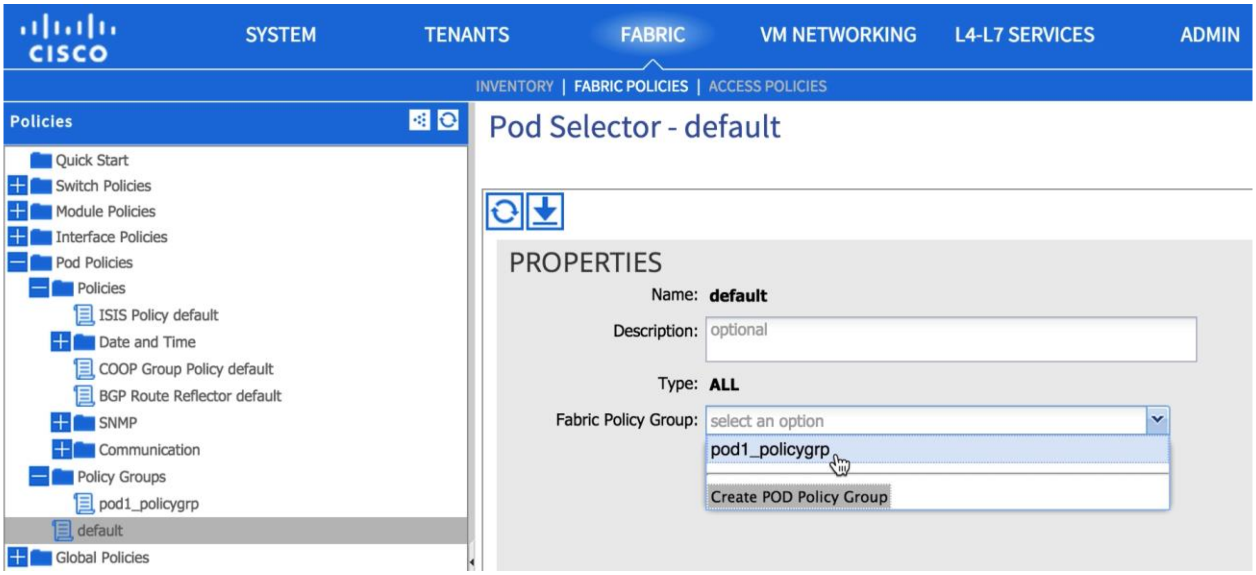

Pod Selector - default

PROPERTIES

Name: default

Description: optional

Type: ALL

Fabric Policy Group: select an option

13. Click SUBMIT.

FlexPod Management Tools Setup

NetApp Virtual Storage Console (VSC) 6.0 Deployment Procedure

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

VSC 6.0 Pre-installation Considerations

The following licenses are required for VSC on storage systems that run clustered Data ONTAP 8.3:

- Protocol licenses (NFS and FCP)
- FlexClone (for provisioning and cloning only)
- SnapRestore (for backup and recovery)
- SnapManager suite

Install VSC 6.0

To install the VSC 6.0 software, complete the following steps:

1. Build a VSC virtual machine with Windows Server 2012 R2, 4GB RAM, two CPUs, and one virtual network interface in the `<<var_ib_mgmt_vlan_id>>` VLAN. The virtual network interface should be a VMXNET 3 adapter.

2. Bring up the VM, install VMware Tools, assign IP addresses, and join the machine to the Active Directory domain.

3. Activate Adobe Flash Player in Windows Server 2012 R2 by installing Desktop Experience under the User Interfaces and Infrastructure Feature on the VM.

4. Install all Windows updates on the VM.

5. Log in to the VSC VM as FlexPod admin user.

6. Download the x64 version of the Virtual Storage Console 6.0 from the NetApp Support site.

7. From the VMware Console, right-click the VSC-6.0-win64.exe file downloaded in step 3 and select Run as administrator.

8. Select the appropriate language and click OK.

9. On the Installation wizard Welcome page, click Next.

10. Select the checkbox to accept the message, click Next.

11. Select the Backup and Recovery capability. Click Next.

Note: The Backup and Recovery capability requires an additional license.

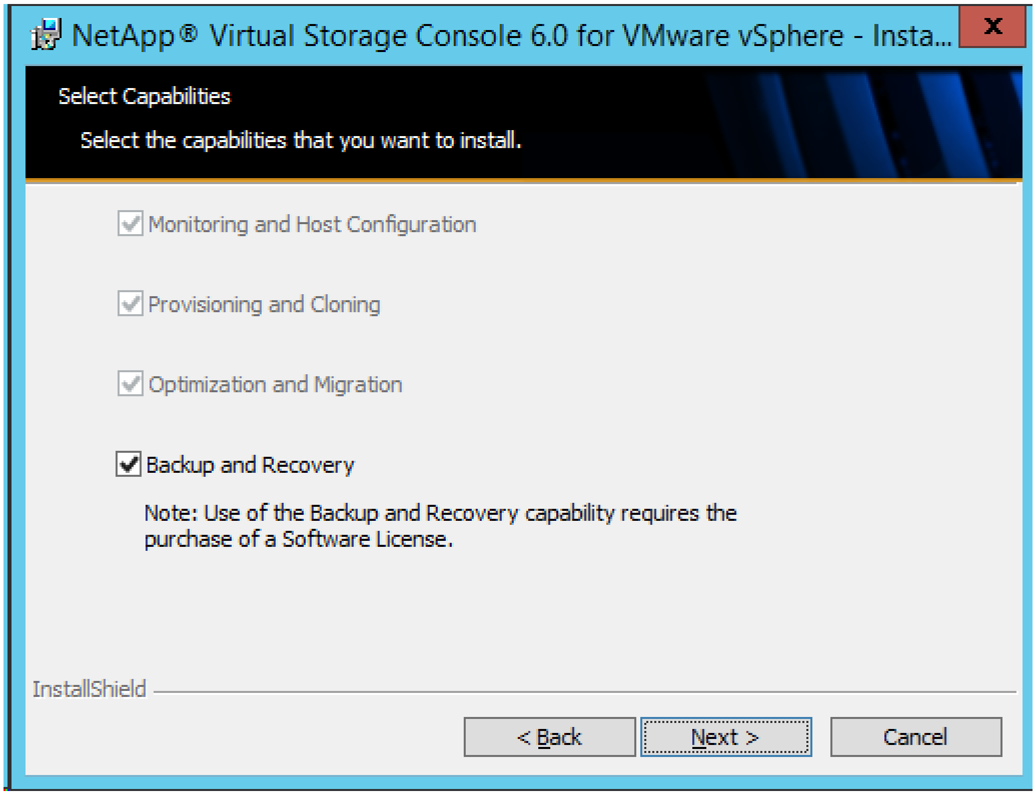

NetApp® Virtual Storage Console 6.0 for VMware vSphere - Insta...

Select Capabilities

Select the capabilities that you want to install.

☑ Monitoring and Host Configuration

☑ Provisioning and Cloning

☑ Optimization and Migration

☑ Backup and Recovery

Note: Use of the Backup and Recovery capability requires the purchase of a Software License.

12. Click Next to accept the default installation location.

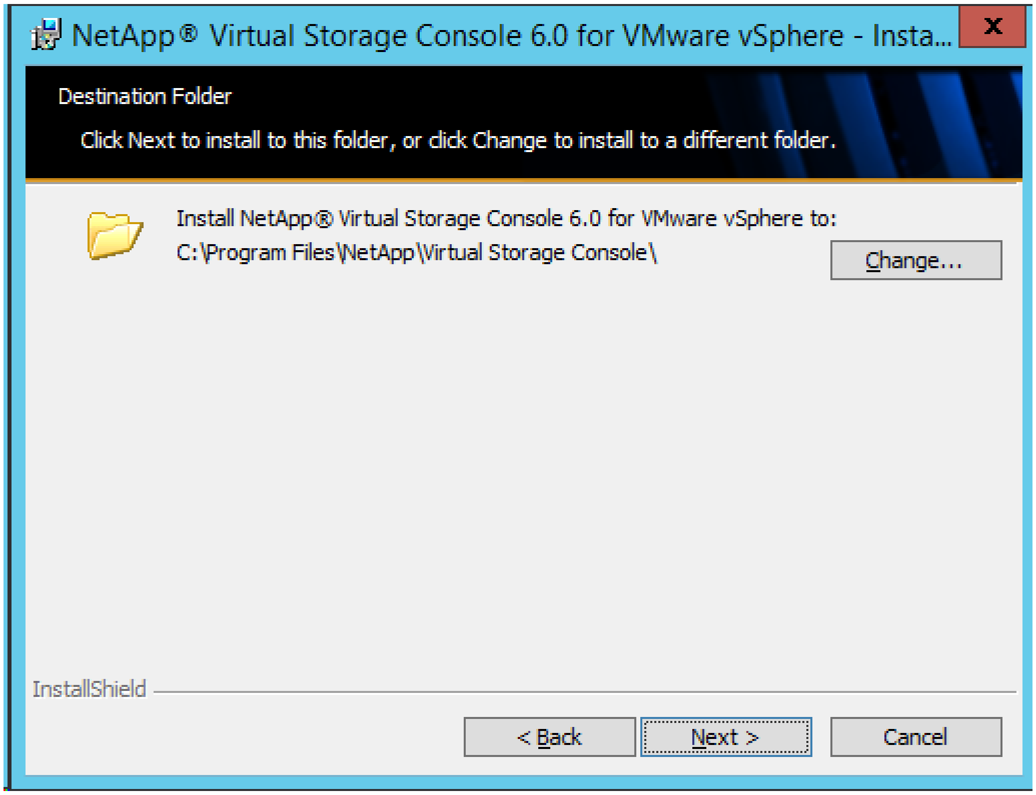

NetApp® Virtual Storage Console 6.0 for VMware vSphere - Insta...

Destination Folder

Click Next to install to this folder, or click Change to install to a different folder.

Install NetApp® Virtual Storage Console 6.0 for VMware vSphere to:
C:\Program Files\NetApp\Virtual Storage Console\

13. Click Install.

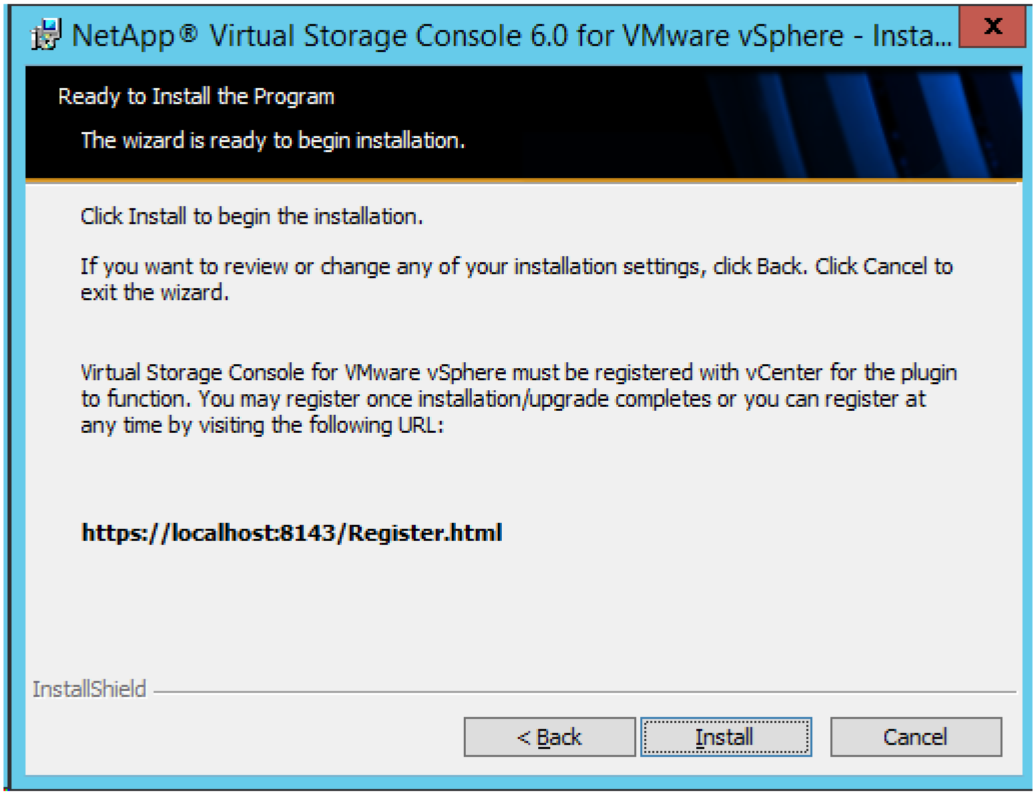

NetApp® Virtual Storage Console 6.0 for VMware vSphere - Insta...

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Virtual Storage Console for VMware vSphere must be registered with vCenter for the plugin to function. You may register once installation/upgrade completes or you can register at any time by visiting the following URL:

https://localhost:8143/Register.html

14. Click Finish.

Register VSC with vCenter Server

To register the VSC with the vCenter Server, complete the following steps:

1. A browser window with the registration URL opens automatically when the installation phase is complete. If the URL does not open automatically, open https://localhost:8143/Register.html in Internet Explorer.

2. Click Continue to this website (not recommended).

3. In the Plug-in Service Information section, select the local IP address that the vCenter Server uses to access the VSC server from the drop-down list.

4. In the vCenter Server Information section, enter the host name or IP address, user name (FlexPod admin user or root), and user password for the vCenter Server. Click Register to complete the registration.

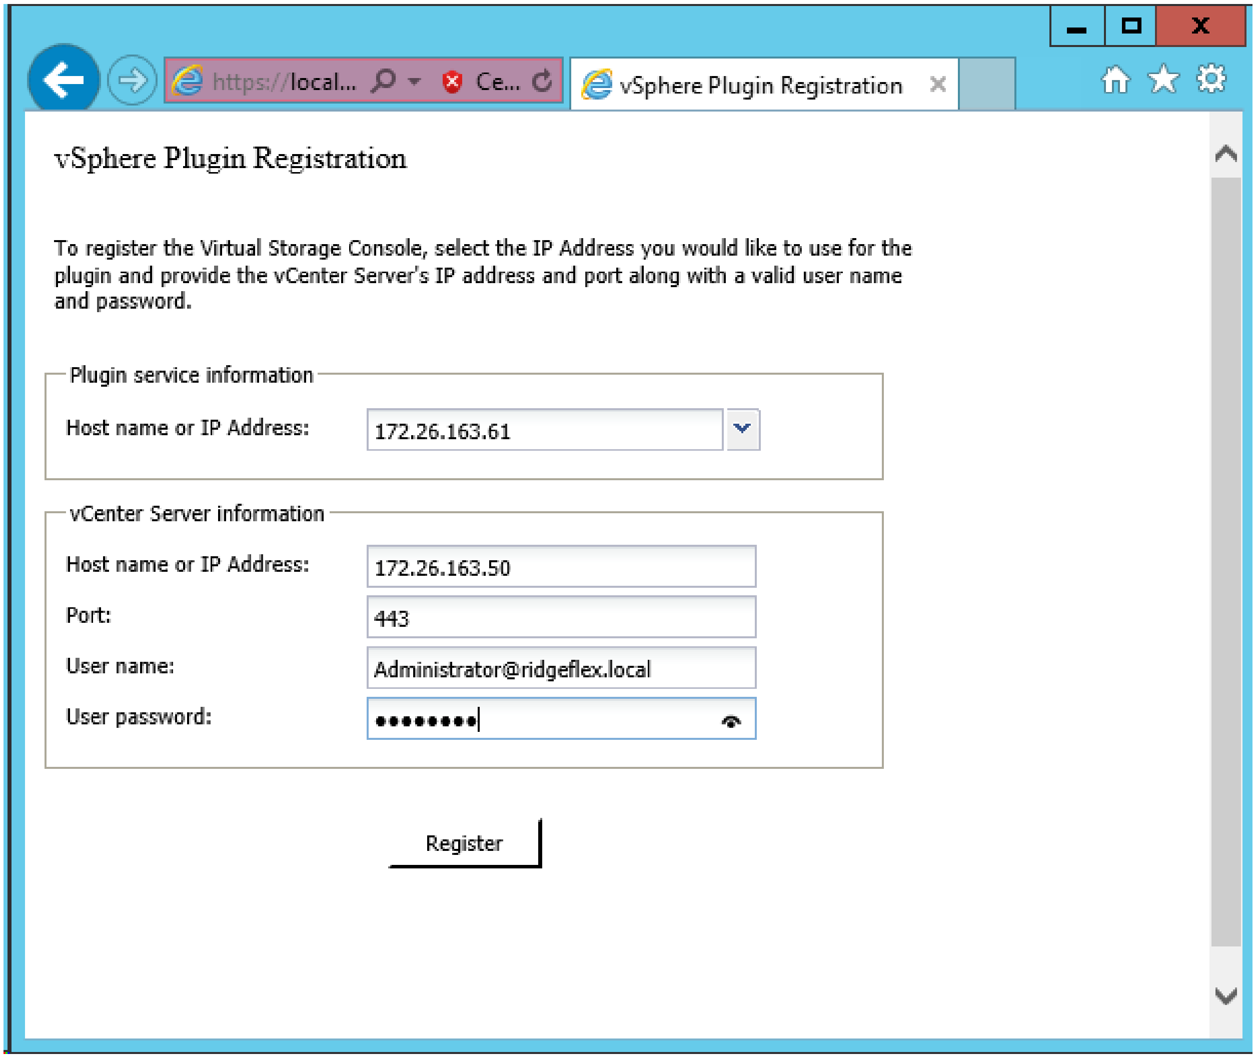

vSphere Plugin Registration

To register the Virtual Storage Console, select the IP Address you would like to use for the plugin and provide the vCenter Server's IP address and port along with a valid user name and password.

Plugin service information

Host name or IP Address: 172.26.163.61

vCenter Server information

Host name or IP Address: 172.26.163.50

Port: 443

User name: Administrator@ridgeflex.local

User password: ••••••••

5. Upon successful registration, the storage controllers are discovered automatically.

Note: Storage discovery process will take some time.

Discover and Add Storage Resources

To discover storage resources for the Monitoring and Host Configuration and the Provisioning and Cloning capabilities, complete the following steps:

1. Using the vSphere web client, log in to the vCenter Server as FlexPod admin user or root. If the vSphere web client was previously opened, close it and then reopen it.

2. In the Home screen, click the Home tab and click Virtual Storage Console.

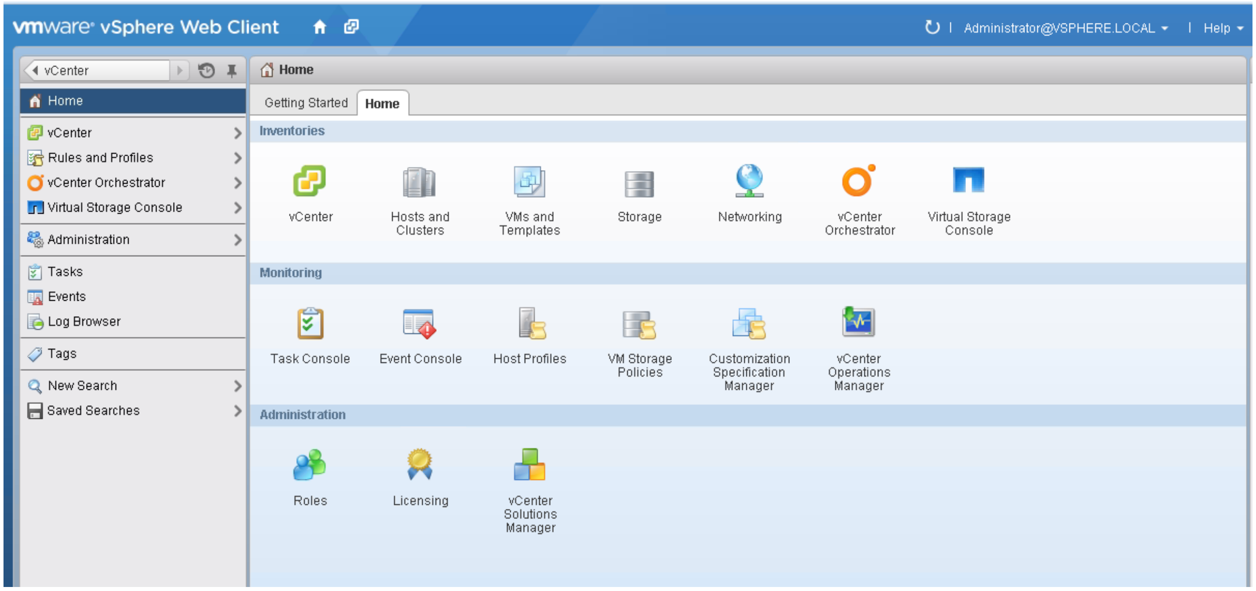

3. Select Storage Systems. Under the Objects tab, click Actions > Modify.

4. In the IP Address/Hostname field, enter the storage cluster management IP. Enter admin for the user name, and the admin password for password. Confirm that Use SSL to connect to this storage system is selected. Click OK.

5. Click OK to accept the controller privileges.

Optimal Storage Settings for ESXi Hosts

VSC allows for the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, click on vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values for these hosts.

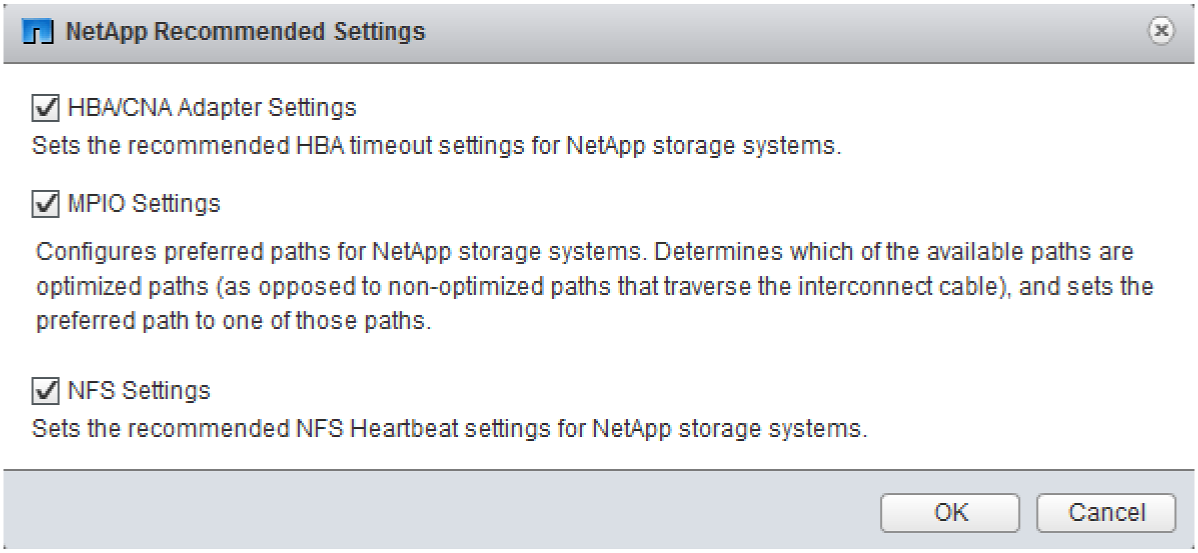

NetApp Recommended Settings

☑ HBA/CNA Adapter Settings

Sets the recommended HBA timeout settings for NetApp storage systems.

☑ MPIO Settings

Configures preferred paths for NetApp storage systems. Determines which of the available paths are optimized paths (as opposed to non-optimized paths that traverse the interconnect cable), and sets the preferred path to one of those paths.

☑ NFS Settings

Sets the recommended NFS Heartbeat settings for NetApp storage systems.

2. Check the settings that are to be applied to the selected vSphere hosts. Click OK to apply the settings.

Note: This functionality sets values for HBAs and CNAs, sets appropriate paths, and path-selection plug-ins, and verifies appropriate settings for software-based I/O (NFS and iSCSI).

3. Click OK.

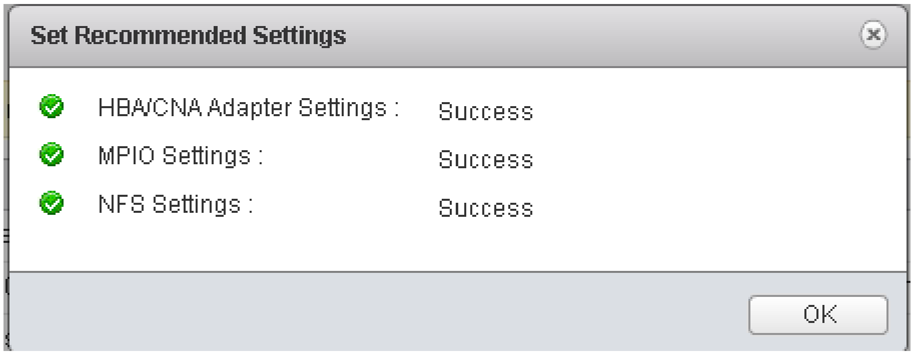

Set Recommended Settings

✓ HBA/CNA Adapter Settings : Success

✓ MPIO Settings : Success

✓ NFS Settings : Success

For each host were settings were adjusted in the previous step, place the host in Maintenance Mode, reboot the host, and exit Maintenance Mode.

VSC 6.0 Backup and Recovery

Prerequisites to Use Backup and Recovery Capability

Before you begin using the Backup and Recovery capability to schedule backups and restore your datastores, virtual machines, or virtual disk files, you must confirm that the storage systems that contain the datastores and virtual machines for which you are creating backups have valid storage credentials.

If you plan to leverage the SnapMirror update option, add all the destination storage systems with valid storage credentials.

Backup and Recovery Configuration

The following steps detail the procedure to configure a backup job for a datastore.

1. From Home screen, select the Home tab and click Storage.

2. On the left, expand the Datacenter and select Datastores.

3. Right-click the datastore which you need to backup. Select NetApp VSC > Backup > Schedule Backup Job.

Note: If you prefer a onetime backup, then choose Backup Now instead of Schedule Backup.

4. Type a backup job name and description.

Note: If you want to create a VMware snapshot for each backup, select Perform VMware consistency snapshot in the options pane.

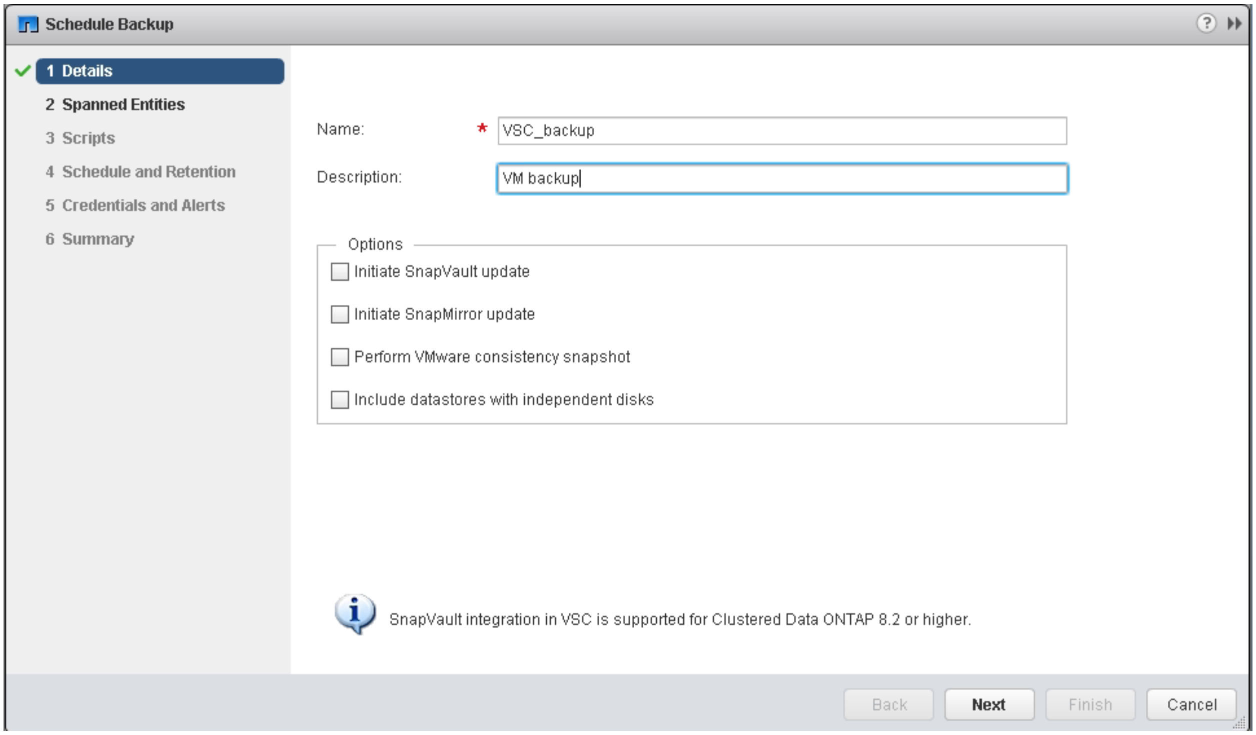

Schedule Backup

- ✓ 1 Details
- 2 Spanned Entities
- 3 Scripts
- 4 Schedule and Retention
- 5 Credentials and Alerts
- 6 Summary

Name: * VSC_backup

Description: VM backup

Options

- □ Initiate SnapVault update
- □ Initiate SnapMirror update
- □ Perform VMware consistency snapshot
- □ Include datastores with independent disks

SnapVault integration in VSC is supported for Clustered Data ONTAP 8.2 or higher.

5. Click Next.

6. Click Next.

7. Select one or more backup scripts if available and click Next.

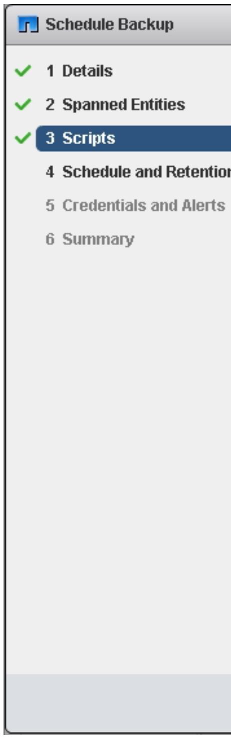

Schedule Backup

- ✓ 1 Details
- ✓ 2 Spanned Entities
- ✓ 3 Scripts
- 4 Schedule and Retentio
- 5 Credentials and Alerts
- 6 Summary

8. Select the hourly, daily, weekly, or monthly schedule that you want for this backup job and click Next.

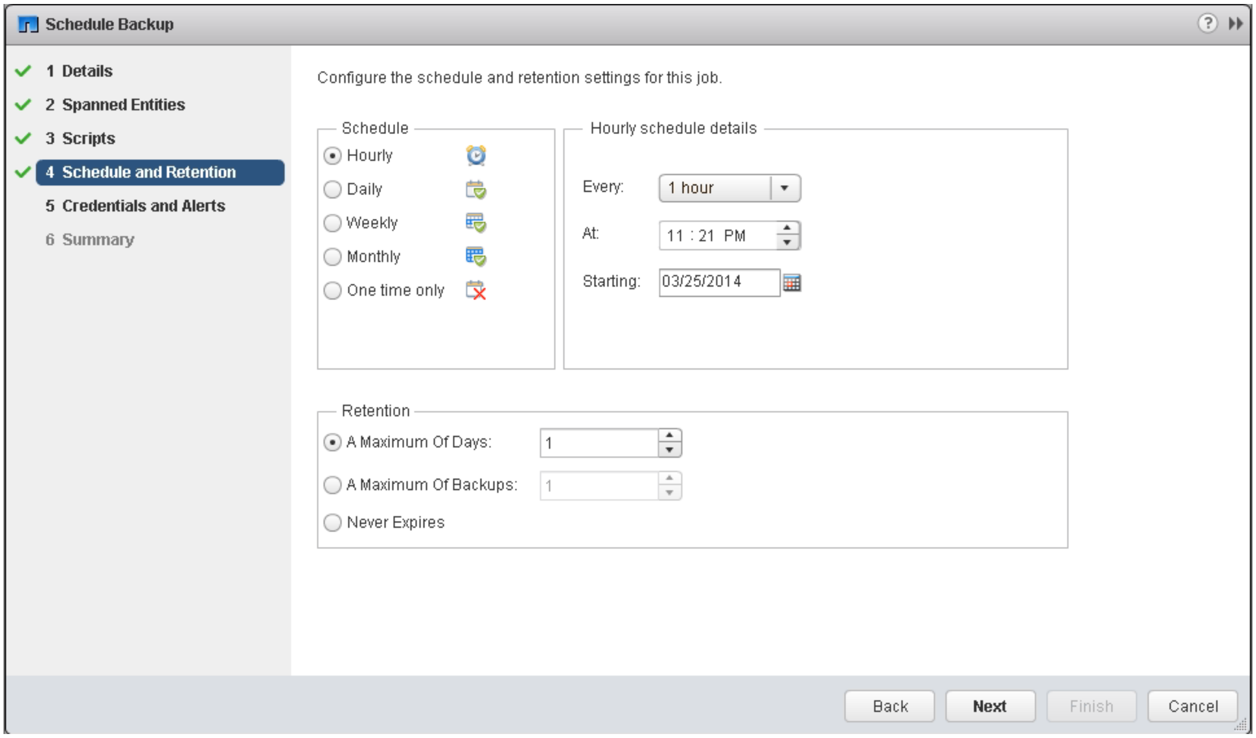

Schedule Backup

✓ 1 Details
✓ 2 Spanned Entities
✓ 3 Scripts
✓ 4 Schedule and Retention
5 Credentials and Alerts
6 Summary

Configure the schedule and retention settings for this job.

Schedule
● Hourly
○ Daily
○ Weekly
○ Monthly
○ One time only

Hourly schedule details
Every: 1 hour
At: 11 : 21 PM
Starting: 03/25/2014

Retention
● A Maximum Of Days: 1
○ A Maximum Of Backups: 1
○ Never Expires

9. Use the default vCenter credentials or type the user name and password for the vCenter Server and click Next.

10. Specify backup retention details as per requirements. Enter an e-mail address for receiving e-mail alerts. You can add multiple e-mail addresses by using semicolons to separate e-mail addresses. Click Next.

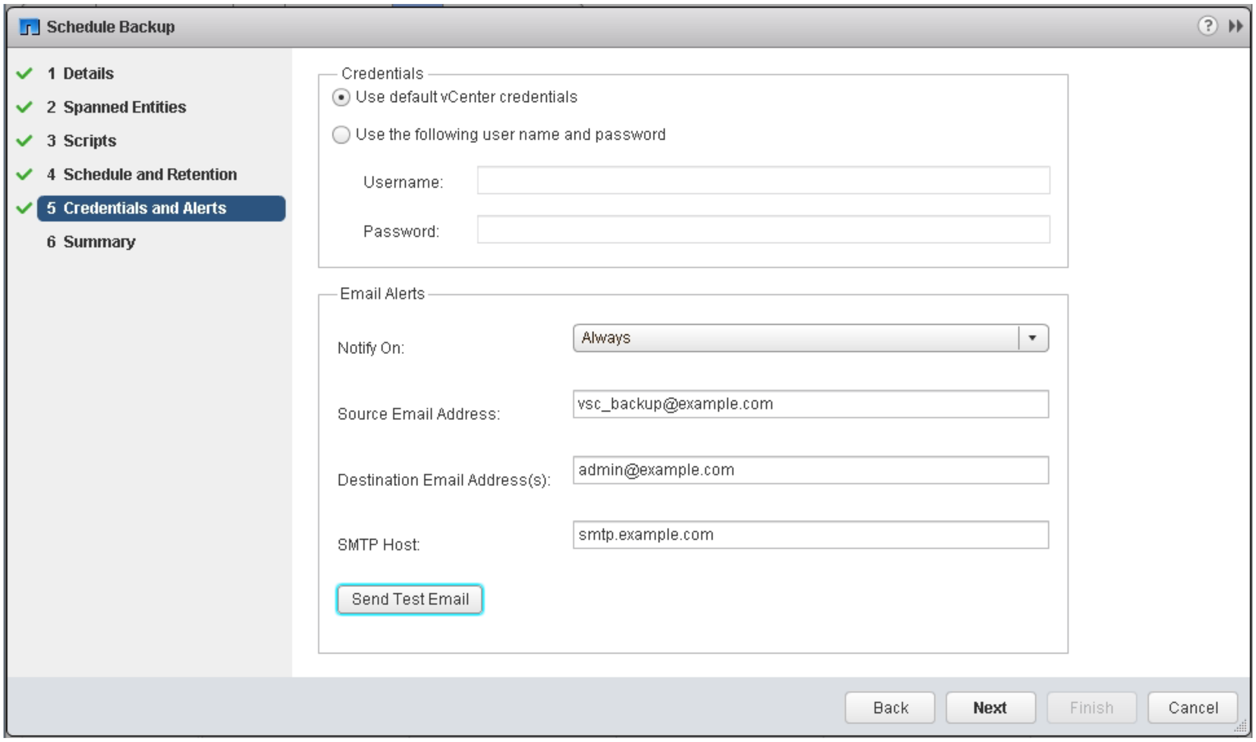

Schedule Backup

✓ 1 Details
✓ 2 Spanned Entities
✓ 3 Scripts
✓ 4 Schedule and Retention
✓ 5 Credentials and Alerts
6 Summary

Credentials
● Use default vCenter credentials
○ Use the following user name and password
Username:
Password:

Email Alerts
Notify On: Always
Source Email Address: vsc_backup@example.com
Destination Email Address(s): admin@example.com
SMTP Host: smtp.example.com

11. Review the summary page and click Finish. If you want to run the job immediately, select the Run Job Now option and then click Finish.

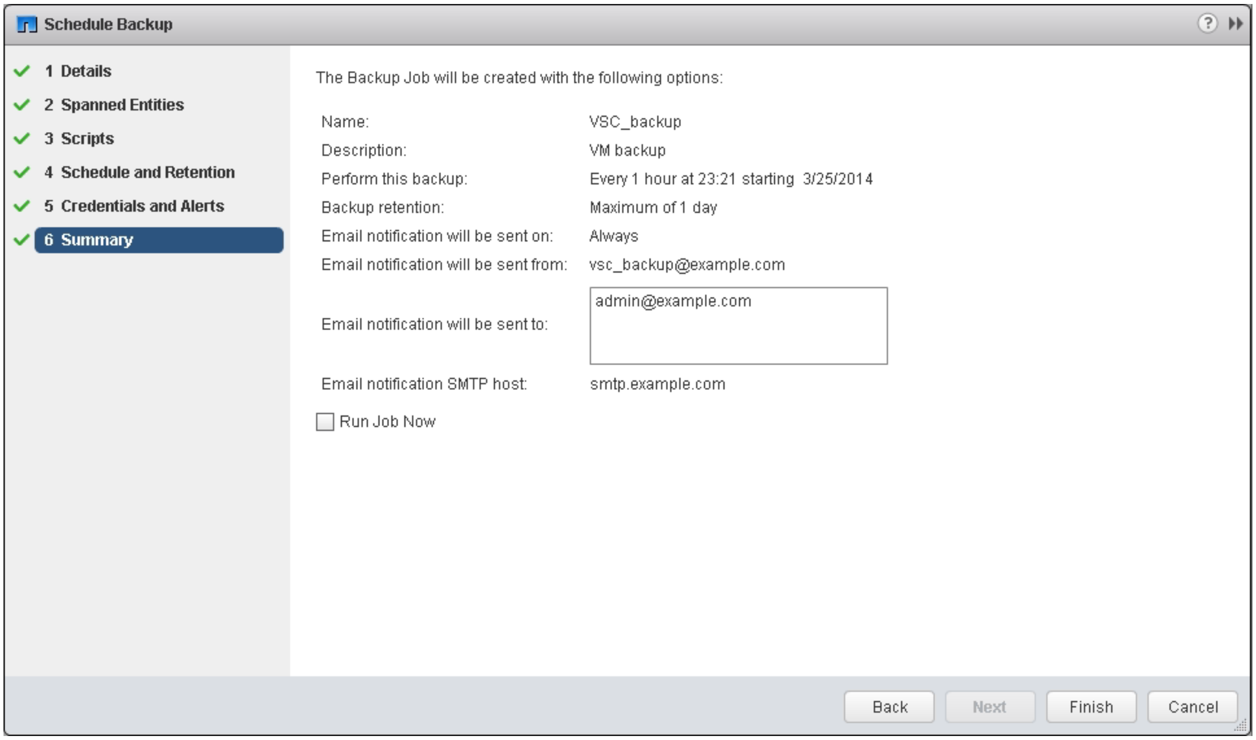

Schedule Backup

The Backup Job will be created with the following options:

Name: VSC_backup
Description: VM backup
Perform this backup: Every 1 hour at 23:21 starting 3/25/2014
Backup retention: Maximum of 1 day
Email notification will be sent on: Always
Email notification will be sent from: vsc_backup@example.com
Email notification will be sent to: admin@example.com
Email notification SMTP host: smtp.example.com

□ Run Job Now

12. Click OK.

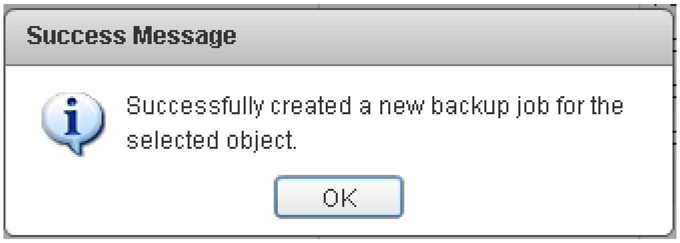

Success Message

Successfully created a new backup job for the selected object.

13. On the storage cluster interface, automatic Snapshot copies of the volume can be disabled by entering the following command:

```
volume modify -volume infra_datastore_1 -snapshot-policy none
```

14. Also, to delete any existing automatic Snapshot copies that have been created on the volume, enter the following command:

```
volume snapshot show -volume infra_datastore_1

volume snapshot delete -volume infra_datastore_1 -vserver Infra-SVM -snapshot
<snapshot name>
```

Note: The wildcard character, *, can be used in snapshot names in the previous command.

NetApp VASA Provider for Clustered Data ONTAP Deployment Procedure

VASA Provider for clustered Data ONTAP uses VMware VASA (vSphere APIs for Storage Awareness) to provide better storage management. By providing information about storage used by Virtual Storage Console for VMware vSphere to the vCenter Server, VASA Provider enables you to make more intelligent virtual machine provisioning decisions and allows the vCenter Server to warn you when certain storage conditions may affect your VMware environment. Virtual Storage Console for VMware vSphere is the management console for VASA Provider. Note that DNS must be setup properly for the VASA provider to register with vCenter. The VASA Provider will do a reverse lookup of its provisioned IP address and the hostname from the lookup should match the provisioned hostname. To install the NetApp VASA provider, complete the following steps:

1. Download the VASA Provider 6.0 for Clustered Data ONTAP OVA file from the NetApp support site.

2. Log into the vSphere Web Client. Go to vCenter > VMs and Templates

3. At the top of the center pane, click Actions > Deploy OVF Template.

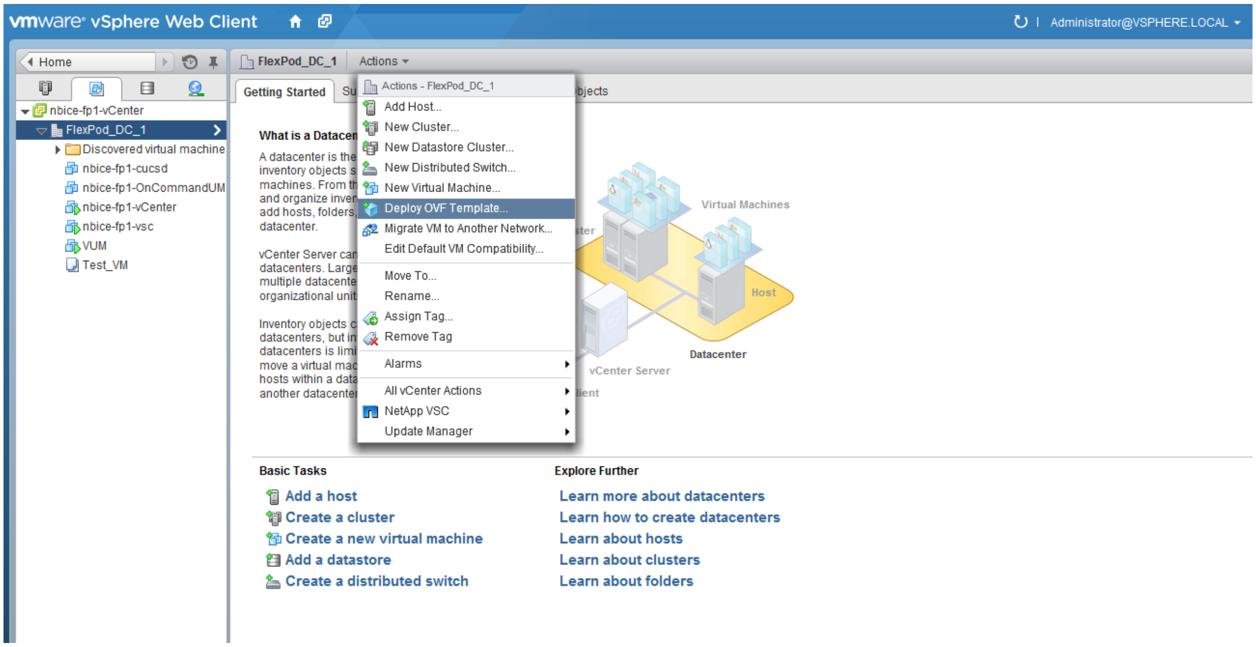

4. Browse the `.ova` file that was downloaded locally. Click Open to select the file.

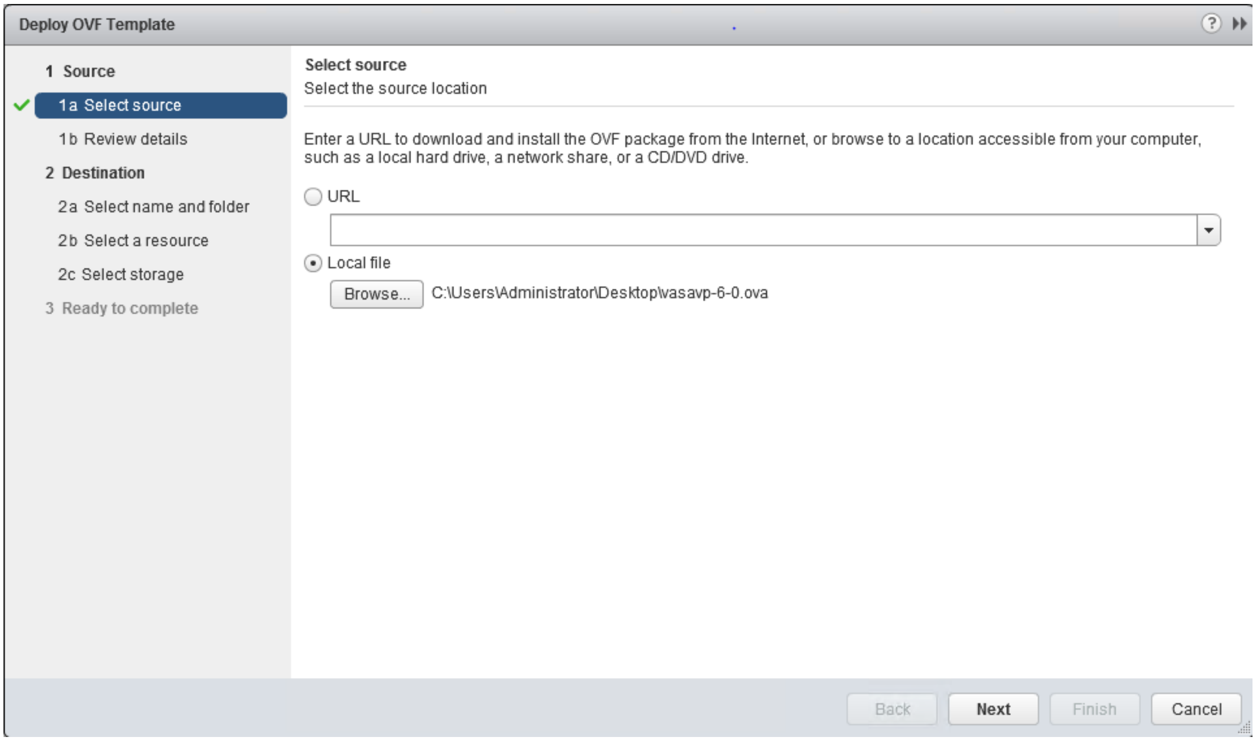

Select source

Select the source location

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

● Local file

Browse... C:\Users\Administrator\Desktop\vasavp-6-0.ova

5. Click Next to proceed with the selected file.

6. Click Next.

7. Read the EULA, then click the Accept button to accept the agreement. Click Next to continue.

8. Enter the name of the VM and select the `FlexPod_DC_1` folder to hold the VM. Click Next to continue.

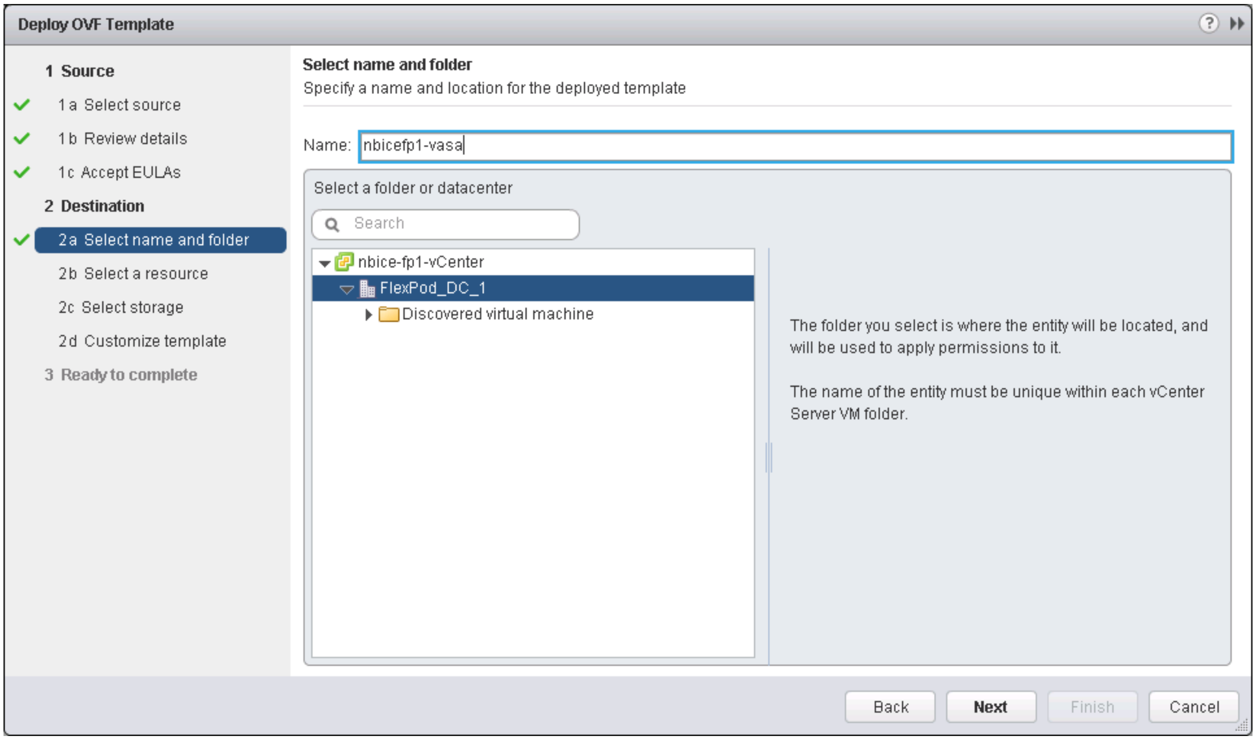

Deploy OVF Template

1 Source
✓ 1a Select source
✓ 1b Review details
✓ 1c Accept EULAs
2 Destination
✓ 2a Select name and folder
2b Select a resource
2c Select storage
2d Customize template
3 Ready to complete

Select name and folder
Specify a name and location for the deployed template

Name: nbicefp1-vasa

Select a folder or datacenter

- nbice-fp1-vCenter
  - FlexPod_DC_1
    - Discovered virtual machine

The folder you select is where the entity will be located, and will be used to apply permissions to it.

The name of the entity must be unique within each vCenter Server VM folder.

9. Select FlexPod_Management within the FlexPod_DC_1 Datacenter as the destination compute resource pool to host the VM. Click Next to continue.

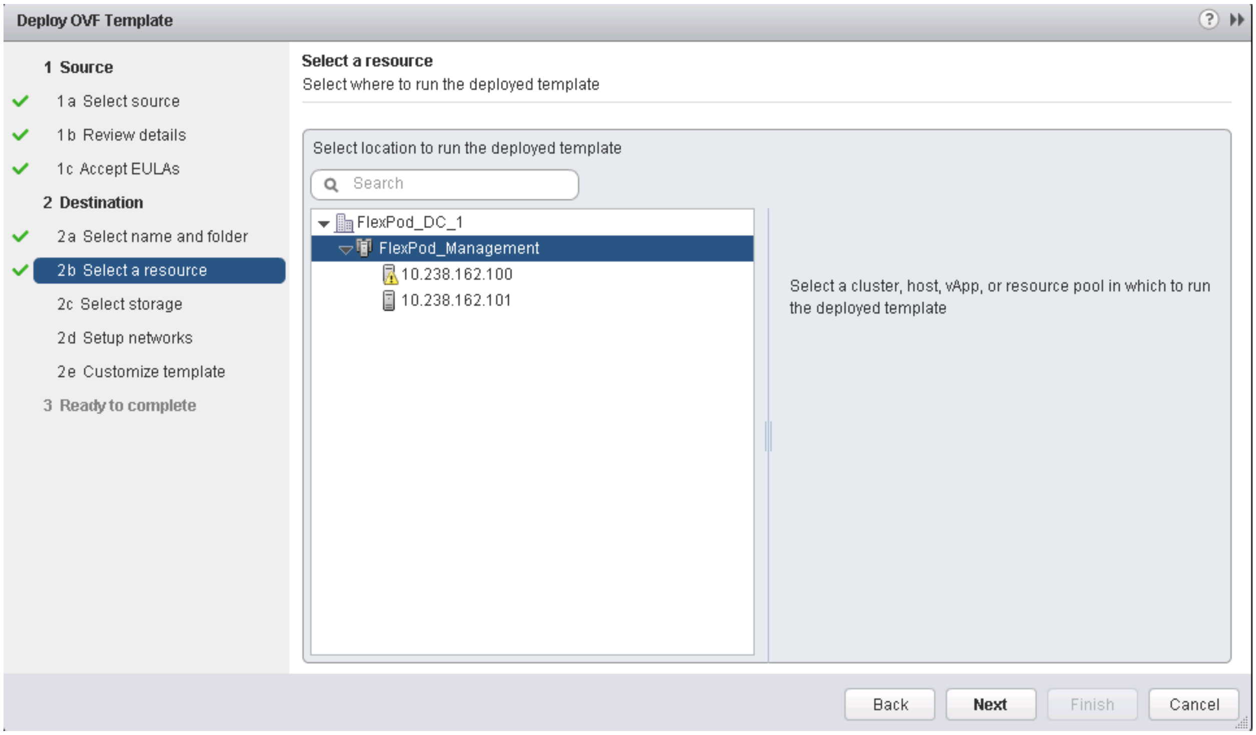

Deploy OVF Template

1 Source
✓ 1a Select source
✓ 1b Review details
✓ 1c Accept EULAs
2 Destination
✓ 2a Select name and folder
✓ 2b Select a resource
2c Select storage
2d Setup networks
2e Customize template
3 Ready to complete

Select a resource
Select where to run the deployed template

Select location to run the deployed template

- FlexPod_DC_1
  - FlexPod_Management
    - 10.238.162.100
    - 10.238.162.101

Select a cluster, host, vApp, or resource pool in which to run the deployed template

10. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the Virtual disk format. Click Next to continue.

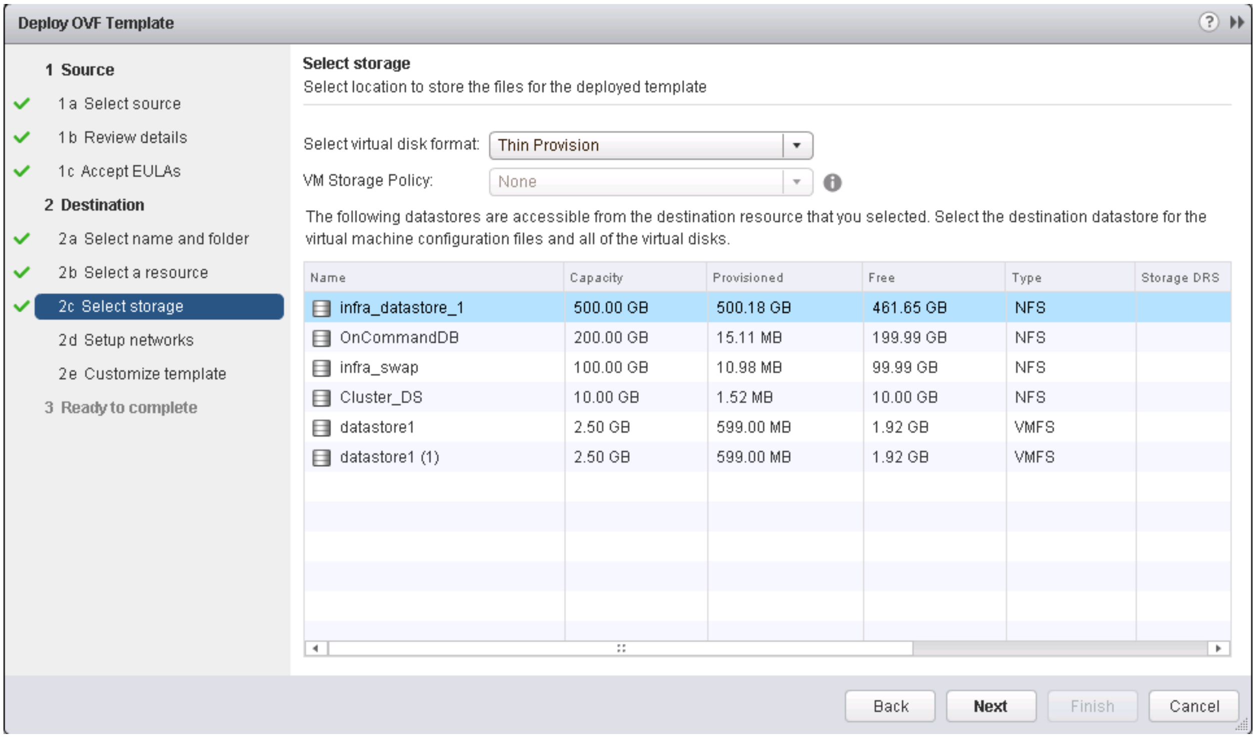

Deploy OVF Template

1 Source
- ✓ 1a Select source
- ✓ 1b Review details
- ✓ 1c Accept EULAs

2 Destination
- ✓ 2a Select name and folder
- ✓ 2b Select a resource
- ✓ 2c Select storage
- 2d Setup networks
- 2e Customize template

3 Ready to complete

Select storage

Select location to store the files for the deployed template

Select virtual disk format: Thin Provision

VM Storage Policy: None

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
infra_datastore_1	500.00 GB	500.18 GB	461.65 GB	NFS	
OnCommandDB	200.00 GB	15.11 MB	199.99 GB	NFS	
infra_swap	100.00 GB	10.98 MB	99.99 GB	NFS	
Cluster_DS	10.00 GB	1.52 MB	10.00 GB	NFS	
datastore1	2.50 GB	599.00 MB	1.92 GB	VMFS	
datastore1 (1)	2.50 GB	599.00 MB	1.92 GB	VMFS	

11. Select MGMT Network as the destination network to the nat source network. Click Next.

12. Fill out the details for the Host Name, IP Address, Netmask, gateway, Primary DNS, and Secondary DNS. Click Next to continue.

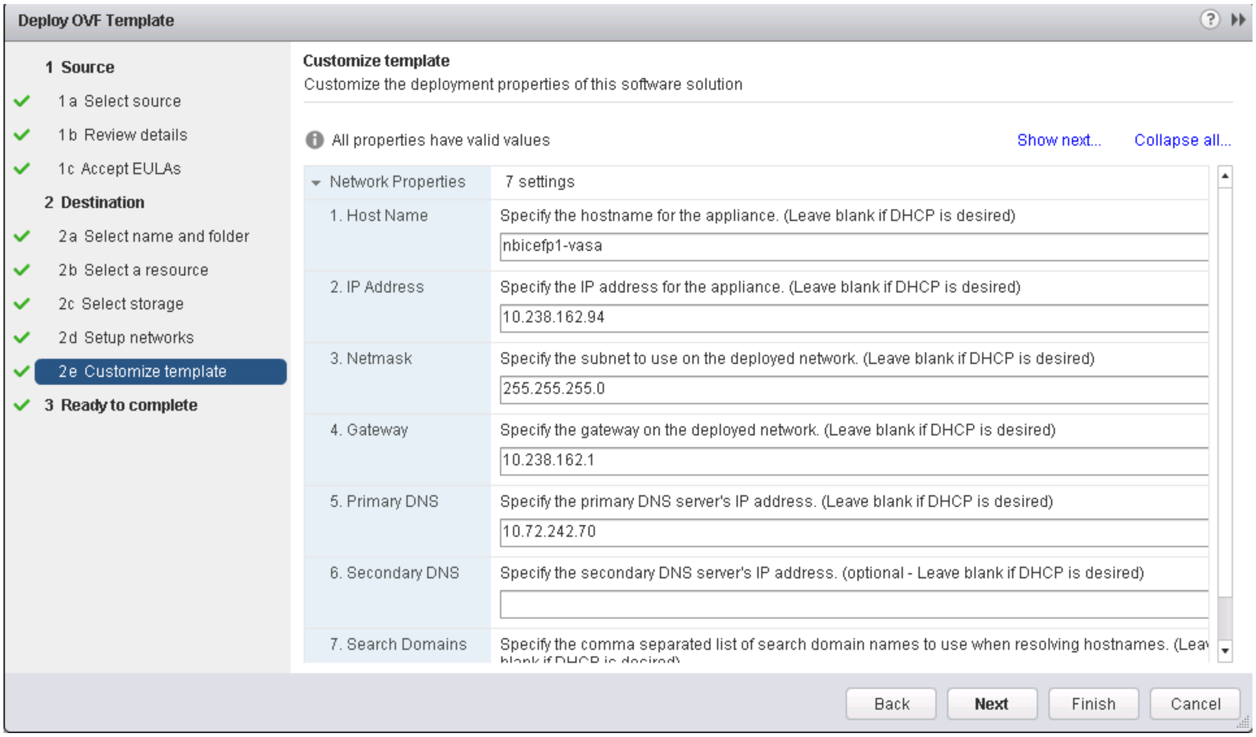

Deploy OVF Template

1 Source
- ✓ 1a Select source
- ✓ 1b Review details
- ✓ 1c Accept EULAs

2 Destination
- ✓ 2a Select name and folder
- ✓ 2b Select a resource
- ✓ 2c Select storage
- ✓ 2d Setup networks
- ✓ 2e Customize template

✓ 3 Ready to complete

Customize template

Customize the deployment properties of this software solution

All properties have valid values

Network Properties — 7 settings

1. Host Name — Specify the hostname for the appliance. (Leave blank if DHCP is desired): nbicefp1-vasa
2. IP Address — Specify the IP address for the appliance. (Leave blank if DHCP is desired): 10.238.162.94
3. Netmask — Specify the subnet to use on the deployed network. (Leave blank if DHCP is desired): 255.255.255.0
4. Gateway — Specify the gateway on the deployed network. (Leave blank if DHCP is desired): 10.238.162.1
5. Primary DNS — Specify the primary DNS server's IP address. (Leave blank if DHCP is desired): 10.72.242.70
6. Secondary DNS — Specify the secondary DNS server's IP address. (optional - Leave blank if DHCP is desired)
7. Search Domains — Specify the comma separated list of search domain names to use when resolving hostnames. (Leav blank if DHCP is desired)

13. Select `Power on after deployment` and click Finish. Wait for the OVA to deploy.

14. Select the VASA VM and click Open Console or Open with VRMC. VASA Provider installation will prompt for installing the VMware tools.

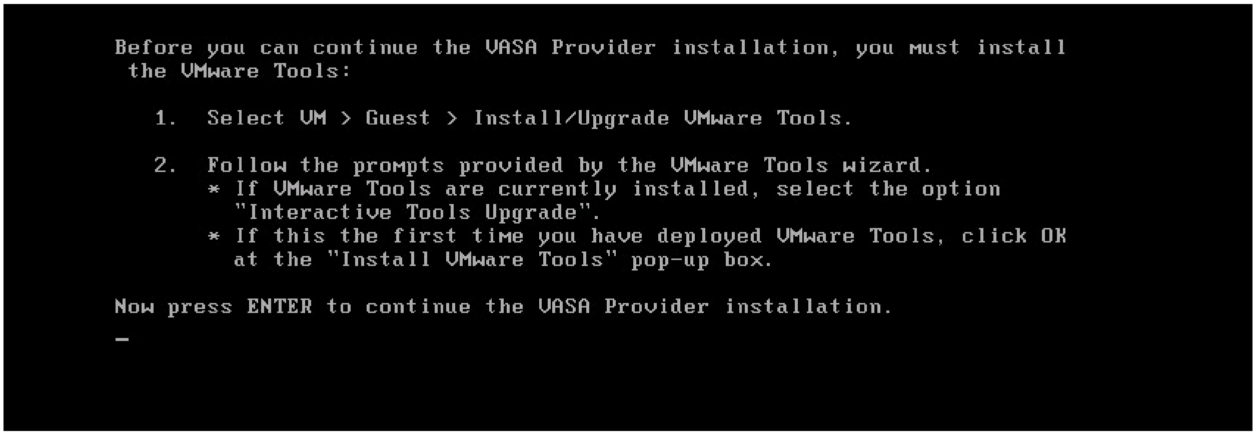

```
Before you can continue the VASA Provider installation, you must install
 the VMware Tools:

   1.  Select VM > Guest > Install/Upgrade VMware Tools.

   2.  Follow the prompts provided by the VMware Tools wizard.
       * If VMware Tools are currently installed, select the option
         "Interactive Tools Upgrade".
       * If this the first time you have deployed VMware Tools, click OK
         at the "Install VMware Tools" pop-up box.

Now press ENTER to continue the VASA Provider installation.
_
```

15. In the web client, click on the VASA VM. Click Summary tab.

16. Click Install VMware Tools and switch back to VASA VM console window.

17. In the VASA VM console, press Enter to continue the VASA provider installation and install VMware tools.

18. Press Enter to reboot the VM.

19. Upon reboot, the VASA appliance will prompt for Maint and vpserver passwords. Type the new passwords accordingly. Press Enter to continue. Wait for the VASA Provider to complete startup.

Registering VASA Provider for Clustered Data ONTAP with VSC

1. Log in to the web client. Click Virtual Storage Console.

2. Click Configuration. Click Register/Unregister VASA Vendor Provider.

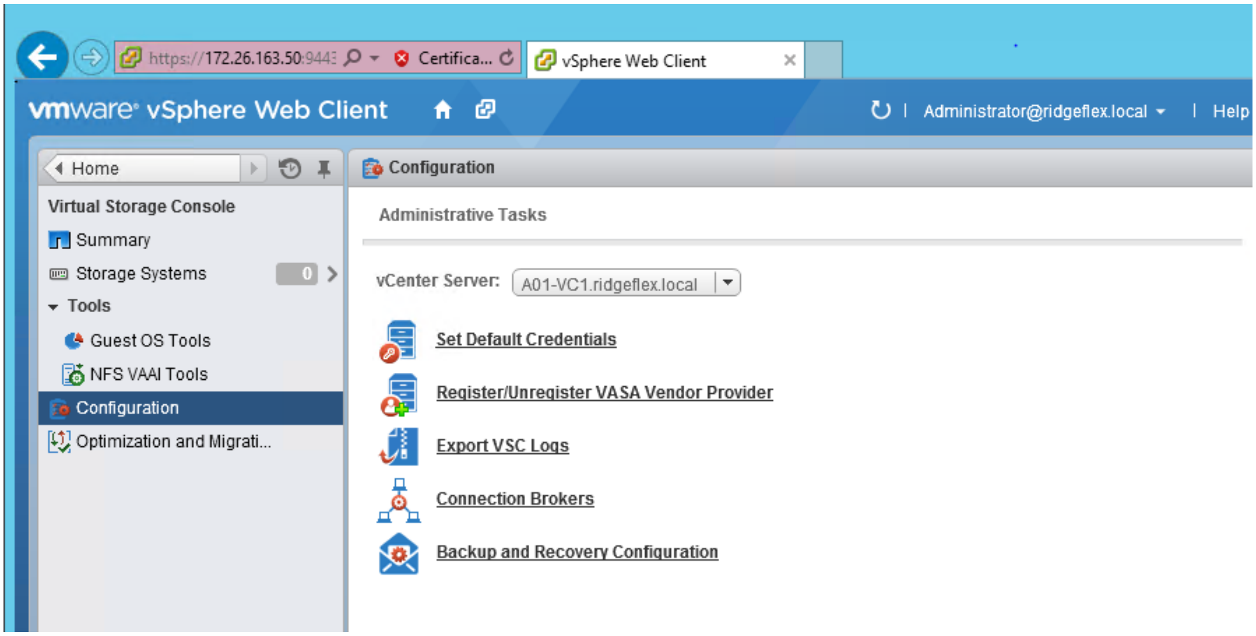

3. Click OK. Enter the IP Address and vpserver Password for the VASA VM. Click Register.

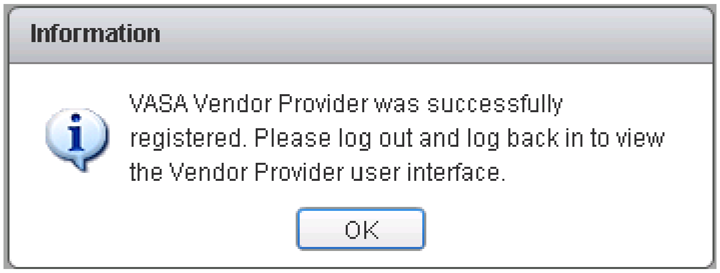

VASA Vendor Provider was successfully registered. Please log out and log back in to view the Vendor Provider user interface.

4. Log out of the web client and log back in to view the Vendor Provider user interface.

5. Select Home > Virtual Storage Console.

6. Select VASA Provider for clustered Data ONTAP.

7. Select Storage Mapping.

8. Click the first icon under the Objects tab to Global auto-generate profiles.

9. Click OK, then click OK again.

10. All of the datastores that on located on the AFF should now show SSD in their profile.

11. These profiles can now be used in datastore provisioning.

OnCommand Unified Manager 6.2

OnCommand Unified Manager OVF Deployment

To install the OnCommand Unified Manager, complete the following steps:

Download and review the OnCommand Unified Manager Installation and Setup Guide.

1. Download the OnCommand Unified Manager version 6.2P1 (OnCommandUnifiedManager-6.2P1.ova), from http://mysupport.netapp.com/NOW/download/software/oncommand_cdot/6.2P1/

2. Log in to the vSphere Web Client. Go to vCenter > VMs and Templates.

3. At the top of the center pane, click Actions > Deploy OVF Template.

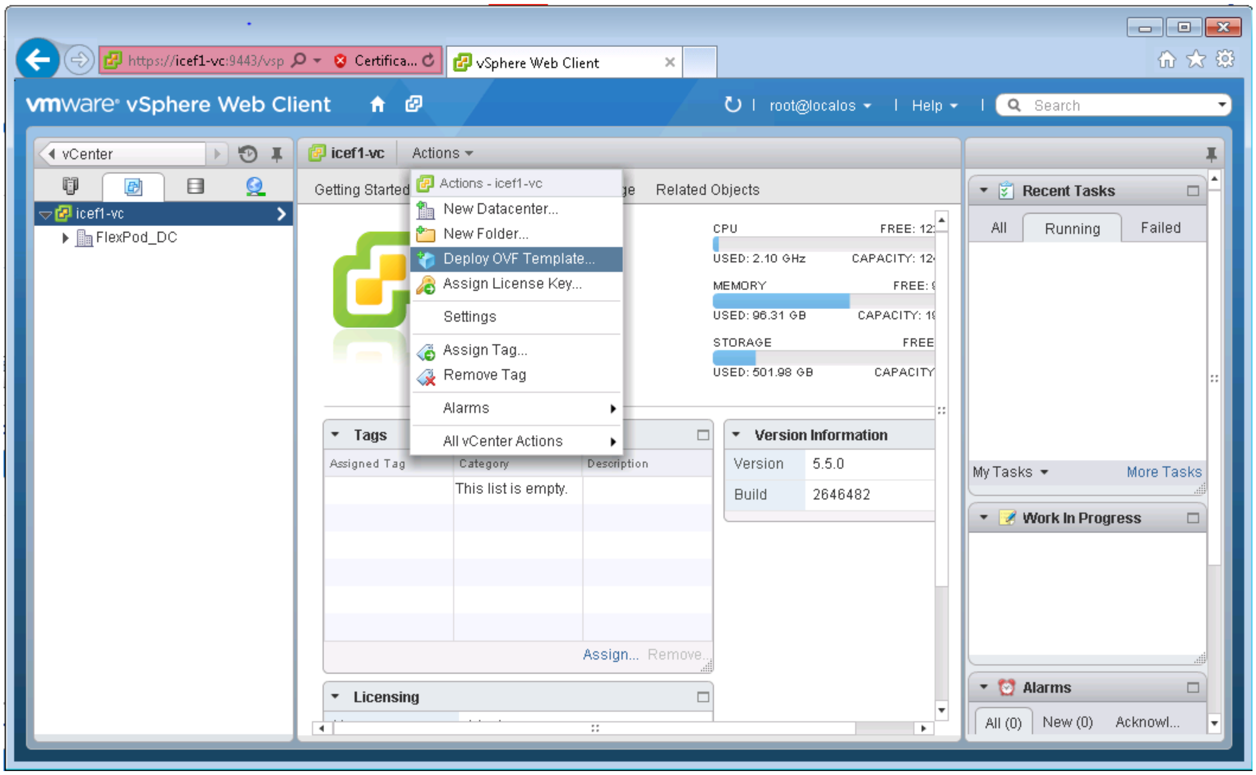

CPU FREE: 12
USED: 2.10 GHz CAPACITY: 12
MEMORY FREE:
USED: 96.31 GB CAPACITY: 1
STORAGE FREE
USED: 501.98 GB CAPACITY

This list is empty.

Version	5.5.0
Build	2646482

4. Browse the `.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.

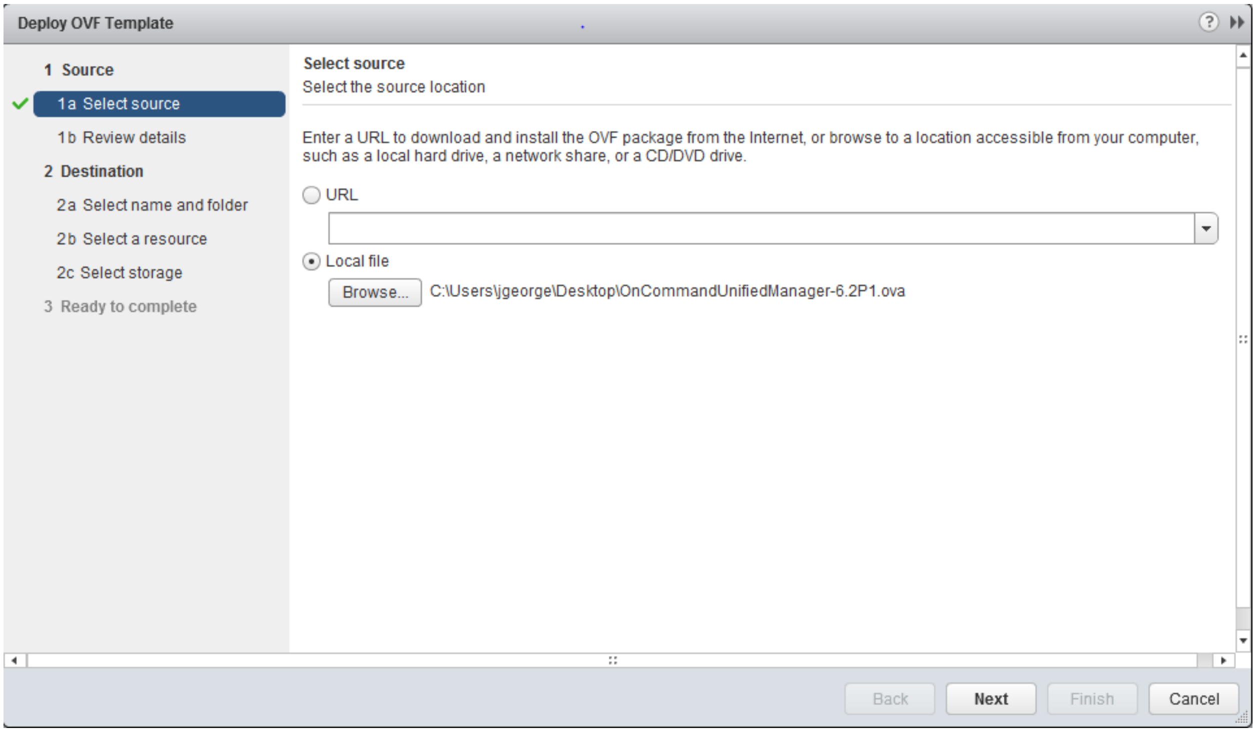

Select source

Select the source location

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

● Local file

Browse... C:\Users\jgeorge\Desktop\OnCommandUnifiedManager-6.2P1.ova

5. Click the checkbox to accept the additional configuration options and click Next.

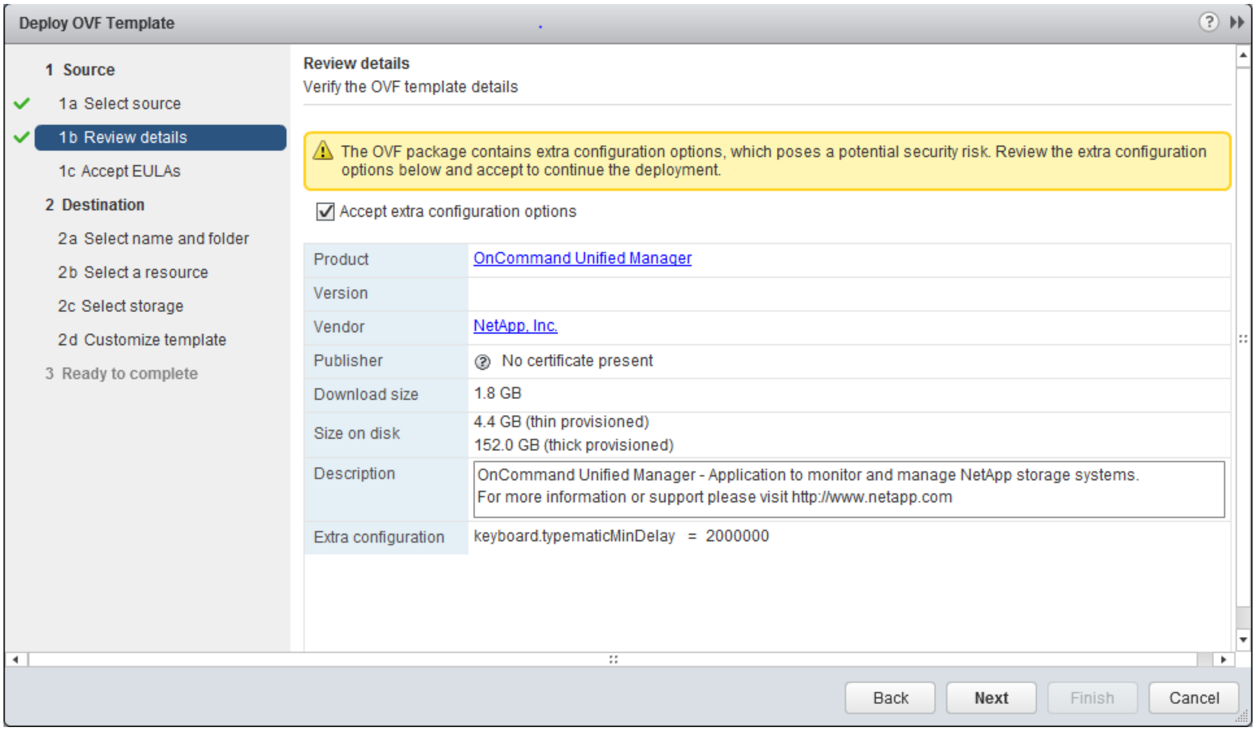

Review details

Verify the OVF template details

⚠ The OVF package contains extra configuration options, which poses a potential security risk. Review the extra configuration options below and accept to continue the deployment.

☑ Accept extra configuration options

Product	OnCommand Unified Manager
Version	
Vendor	NetApp, Inc.
Publisher	No certificate present
Download size	1.8 GB
Size on disk	4.4 GB (thin provisioned) 152.0 GB (thick provisioned)
Description	OnCommand Unified Manager - Application to monitor and manage NetApp storage systems. For more information or support please visit http://www.netapp.com
Extra configuration	keyboard.typematicMinDelay = 2000000

6. Read the EULA, then click the Accept button to accept the agreement. Click Next to continue.

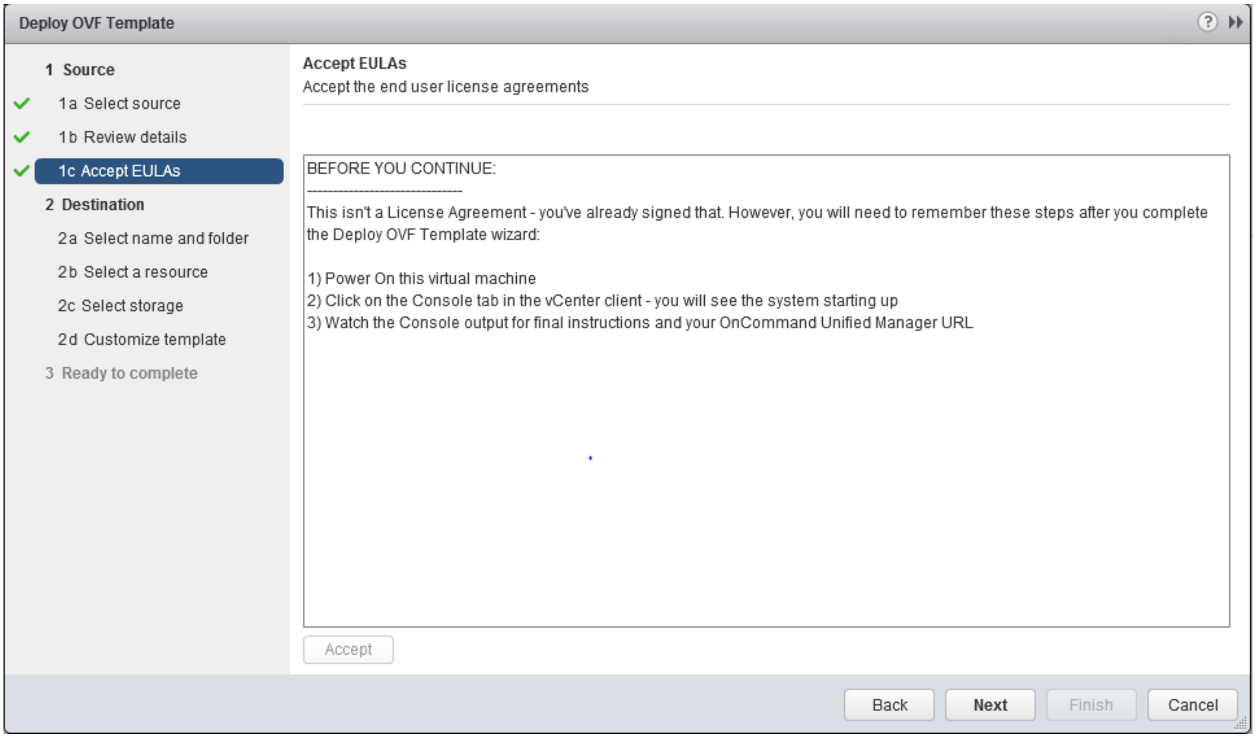

Accept EULAs

Accept the end user license agreements

BEFORE YOU CONTINUE:

This isn't a License Agreement - you've already signed that. However, you will need to remember these steps after you complete the Deploy OVF Template wizard:

1) Power On this virtual machine
2) Click on the Console tab in the vCenter client - you will see the system starting up
3) Watch the Console output for final instructions and your OnCommand Unified Manager URL

7. Enter the name of the VM and select the `FlexPod_DC_1` folder to hold the VM. Click Next to continue.

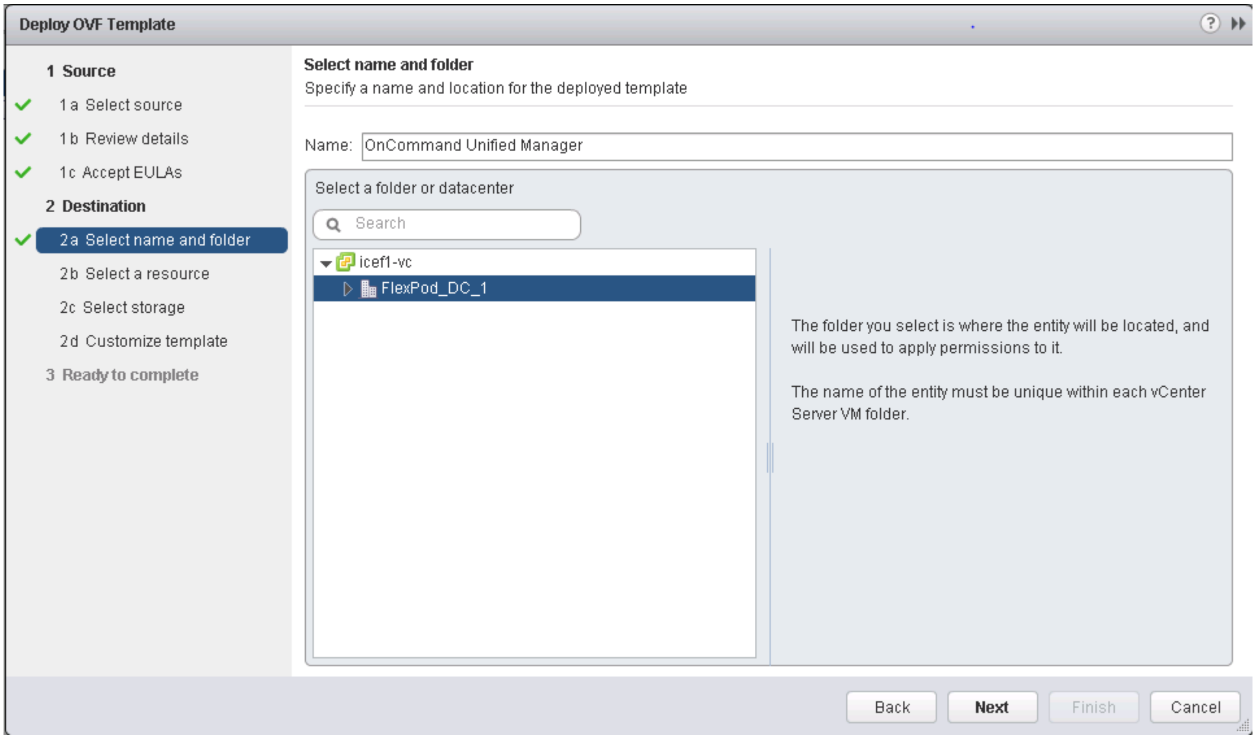

Select name and folder

Specify a name and location for the deployed template

Name: OnCommand Unified Manager

The folder you select is where the entity will be located, and will be used to apply permissions to it.

The name of the entity must be unique within each vCenter Server VM folder.

8. Select FlexPod_Management within the FlexPod_DC_1 datacenter as the destination compute resource pool to host the VM. Click Next to continue.

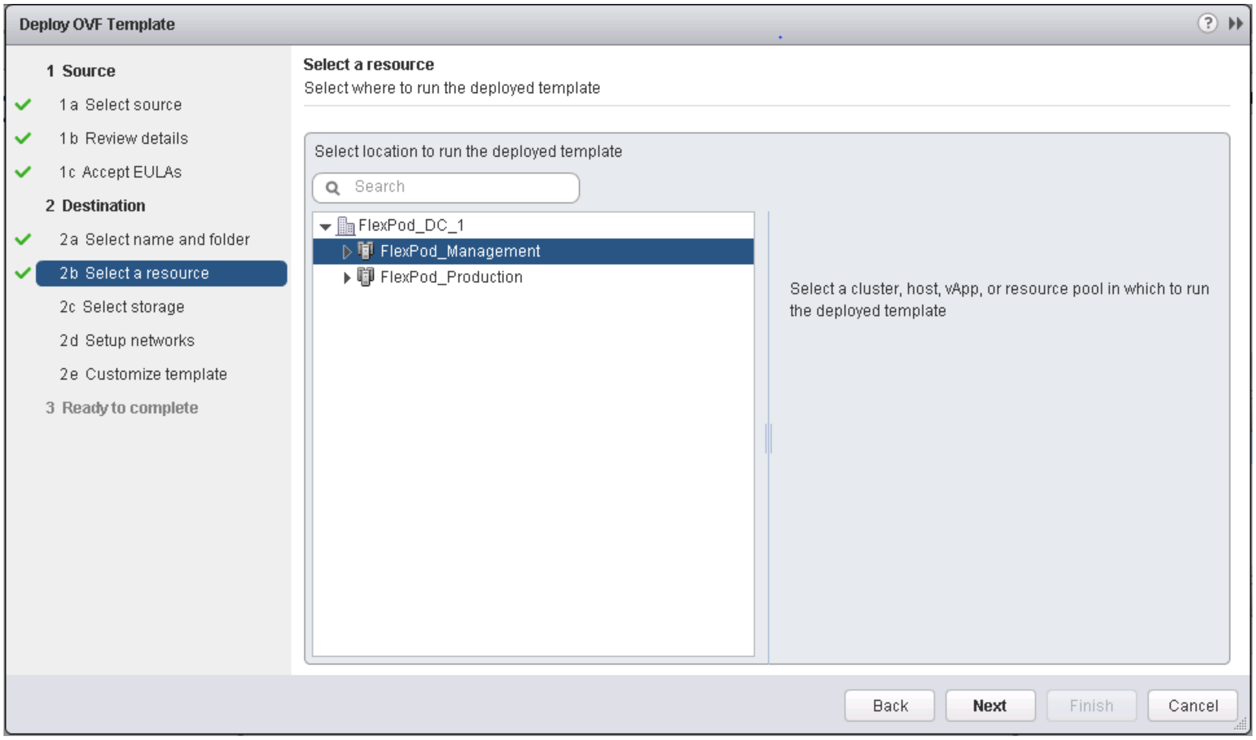

Select a resource
Select where to run the deployed template

Select location to run the deployed template

- FlexPod_DC_1
  - FlexPod_Management
  - FlexPod_Production

Select a cluster, host, vApp, or resource pool in which to run the deployed template

9. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the Virtual disk format. Click Next to continue.

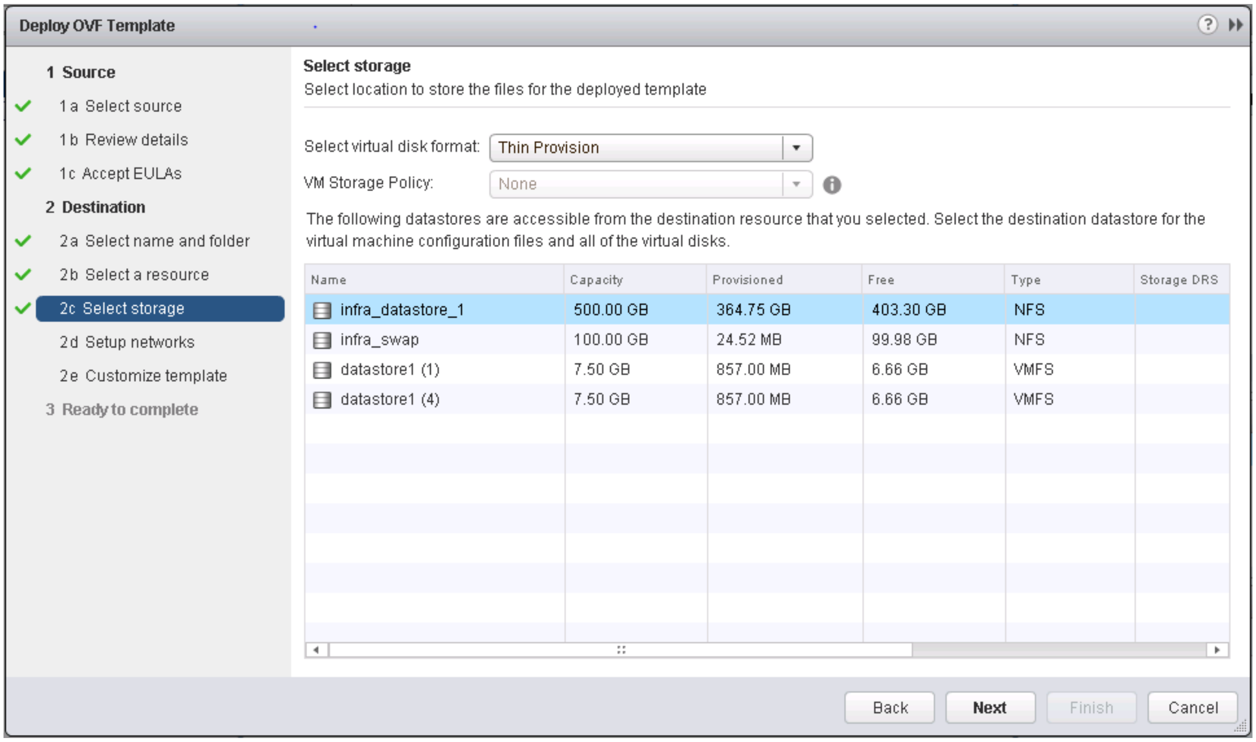

Select storage
Select location to store the files for the deployed template

Select virtual disk format: Thin Provision

VM Storage Policy: None

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
infra_datastore_1	500.00 GB	364.75 GB	403.30 GB	NFS	
infra_swap	100.00 GB	24.52 MB	99.98 GB	NFS	
datastore1 (1)	7.50 GB	857.00 MB	6.66 GB	VMFS	
datastore1 (4)	7.50 GB	857.00 MB	6.66 GB	VMFS	

10. Select `MGMT Network` as the destination network to the nat source network. Click Next.

11. Fill out the details for the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.

12. Clear the Power on after deployment checkbox.

13. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.

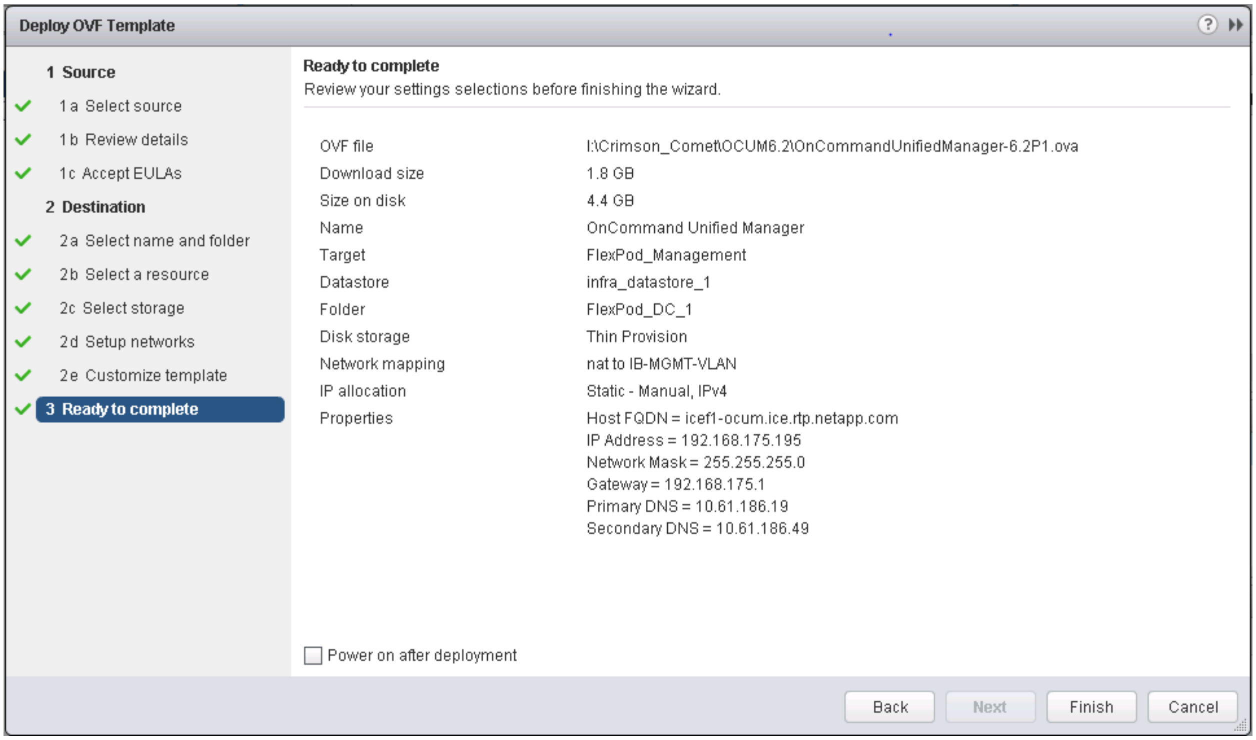

Ready to complete

Review your settings selections before finishing the wizard.

OVF file	I:\Crimson_Comet\OCUM6.2\OnCommandUnifiedManager-6.2P1.ova
Download size	1.8 GB
Size on disk	4.4 GB
Name	OnCommand Unified Manager
Target	FlexPod_Management
Datastore	infra_datastore_1
Folder	FlexPod_DC_1
Disk storage	Thin Provision
Network mapping	nat to IB-MGMT-VLAN
IP allocation	Static - Manual, IPv4
Properties	Host FQDN = icef1-ocum.ice.rtp.netapp.com IP Address = 192.168.175.195 Network Mask = 255.255.255.0 Gateway = 192.168.175.1 Primary DNS = 10.61.186.19 Secondary DNS = 10.61.186.49

□ Power on after deployment

14. On the left pane, navigate to vCenter -> Virtual machines. After OVF deployment is complete, right-click the newly created virtual machine and select Edit Settings.

15. Click the CPU tab to expand the CPU options:

    a. The minimum required CPU Reservation is 4786 MHz. Determine the CPU frequency of the host.

    b. Set the number of CPUs to the number of CPUs required (4786 / CPU Frequency of host).

    c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has 2 CPUs operating at a speed of 1999MHz, then the VM would need 4 virtual CPUs (4786 / 1999 = 2.39 – rounded to 4 virtual CPUs). If the host has 2 physical CPU sockets, 2 Cores per Socket, set the CPU Reservation and Limit to 4786 MHz

    d. The amount of memory can be set to 8 GB. Use the OnCommand Unified Manager Installation and Setup Guide for guidance on these settings.

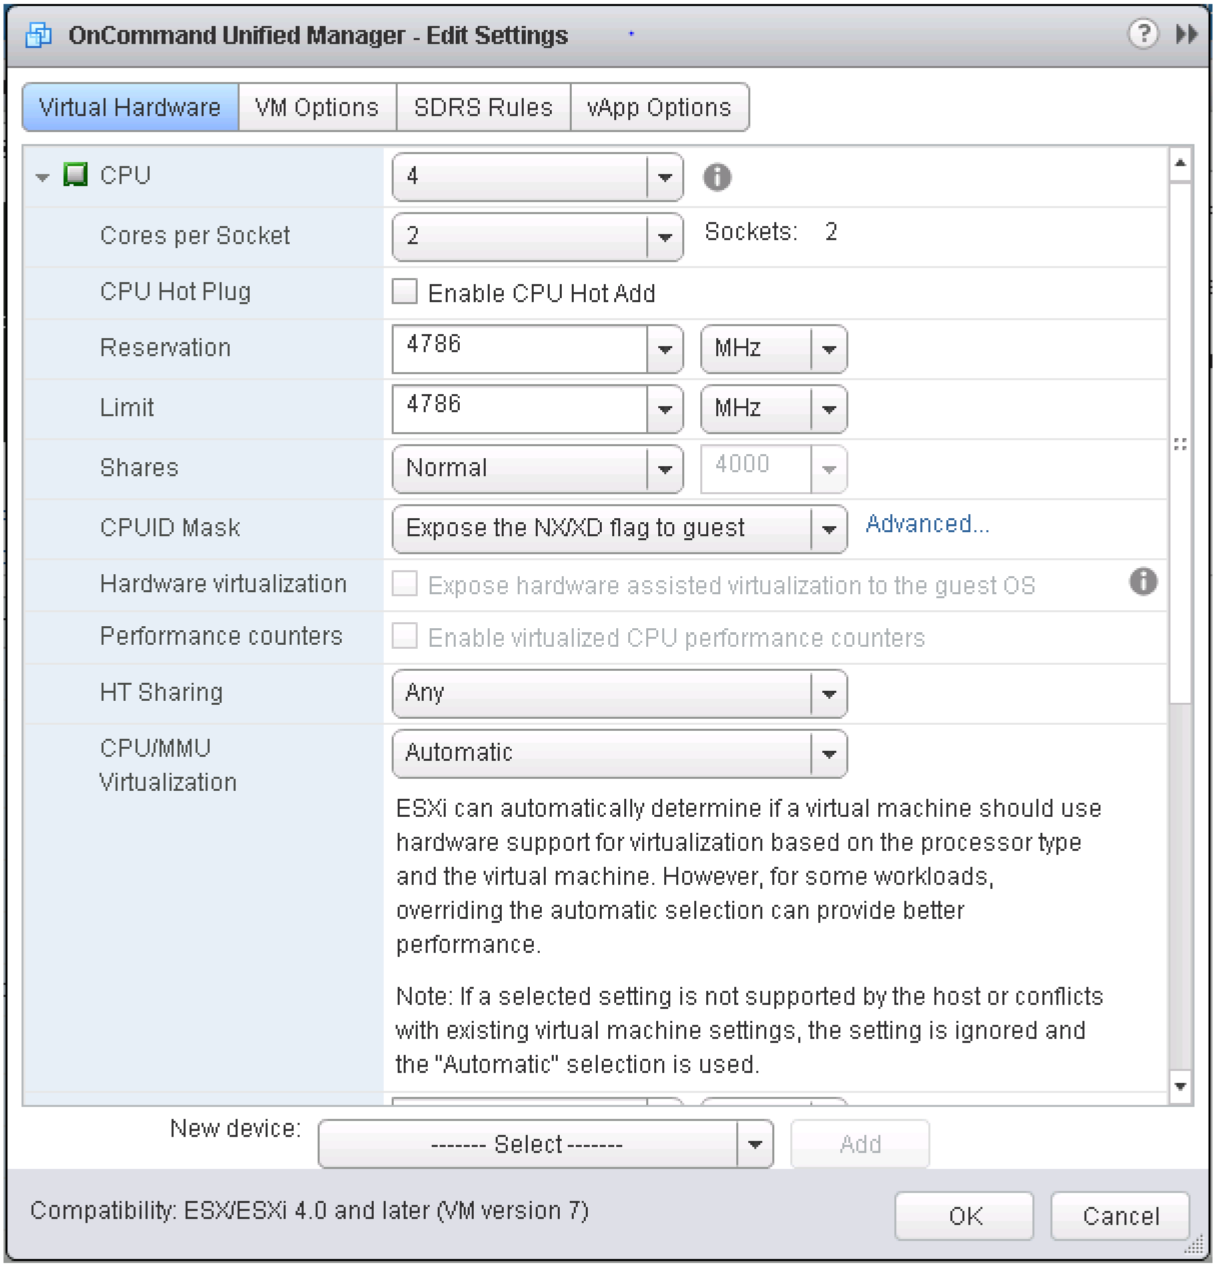

OnCommand Unified Manager - Edit Settings

CPU	4	
Cores per Socket	2	Sockets: 2
CPU Hot Plug	□ Enable CPU Hot Add	
Reservation	4786	MHz
Limit	4786	MHz
Shares	Normal	4000
CPUID Mask	Expose the NX/XD flag to guest	Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS	
Performance counters	□ Enable virtualized CPU performance counters	
HT Sharing	Any	
CPU/MMU Virtualization	Automatic	

ESXi can automatically determine if a virtual machine should use hardware support for virtualization based on the processor type and the virtual machine. However, for some workloads, overriding the automatic selection can provide better performance.

Note: If a selected setting is not supported by the host or conflicts with existing virtual machine settings, the setting is ignored and the "Automatic" selection is used.

New device: ------- Select -------

Compatibility: ESX/ESXi 4.0 and later (VM version 7)

16. Click OK to accept the changes.

17. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Unified Manager Basic Setup

1. Select the VM in the left-hand pane. In the center pane, select Open with VMRC.

2. In the VMRC window, select Manage > Install VMware Tools. VMware Tools will install in the VM.

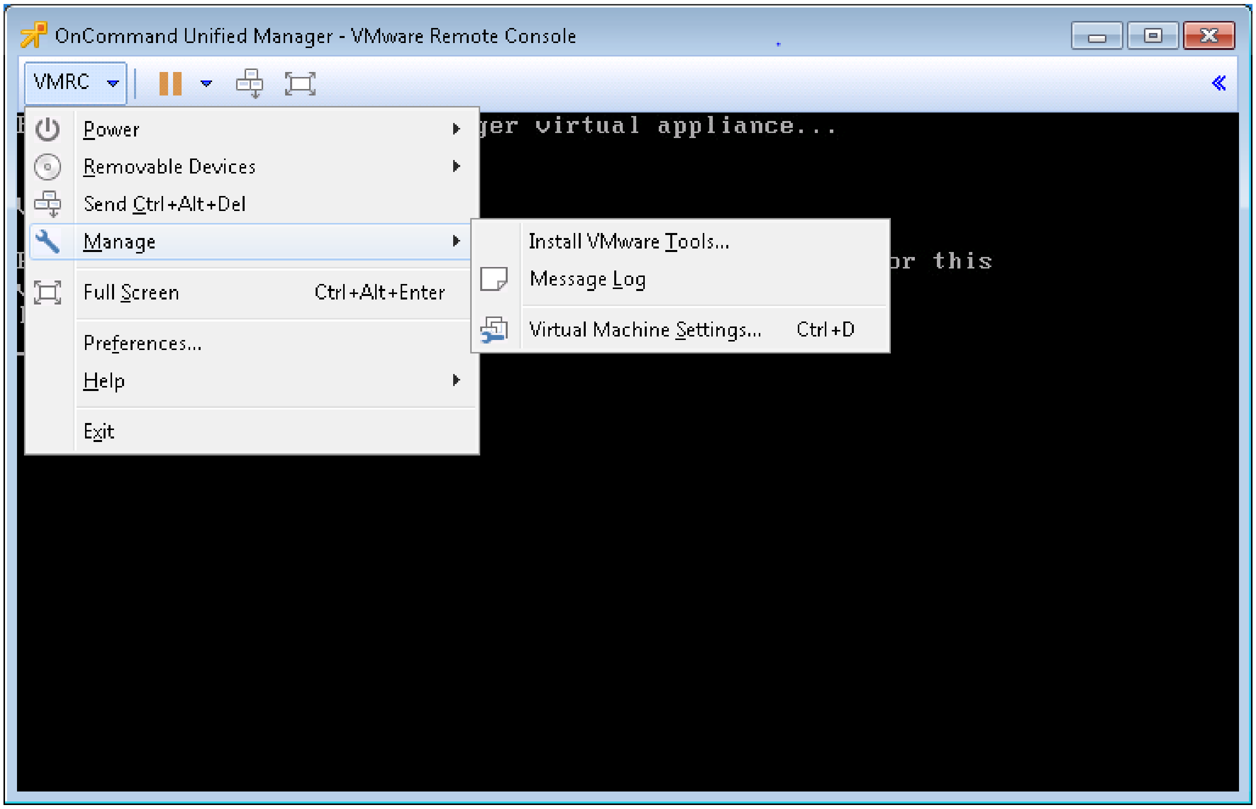

3. Set up OnCommand Unified Manager by answering the following questions in the console window:

```
Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>
```

These commands complete the network configuration checks SSL certificate generation for HTTPS and start the OnCommand Unified Manager services.

4. To Create a Maintenance User account, run the following commands:

Note: The maintenance user manages and maintains the settings on the OnCommand Unified Manager virtual appliance.

```
Username : admin

Enter new UNIX password: <<var_password>>

Retype new UNIX password: <<var_password>>
```

5. Using a web browser navigate to the OnCommand Unified Manager using URL: `https://<<var_oncommand_server_ip>>`.

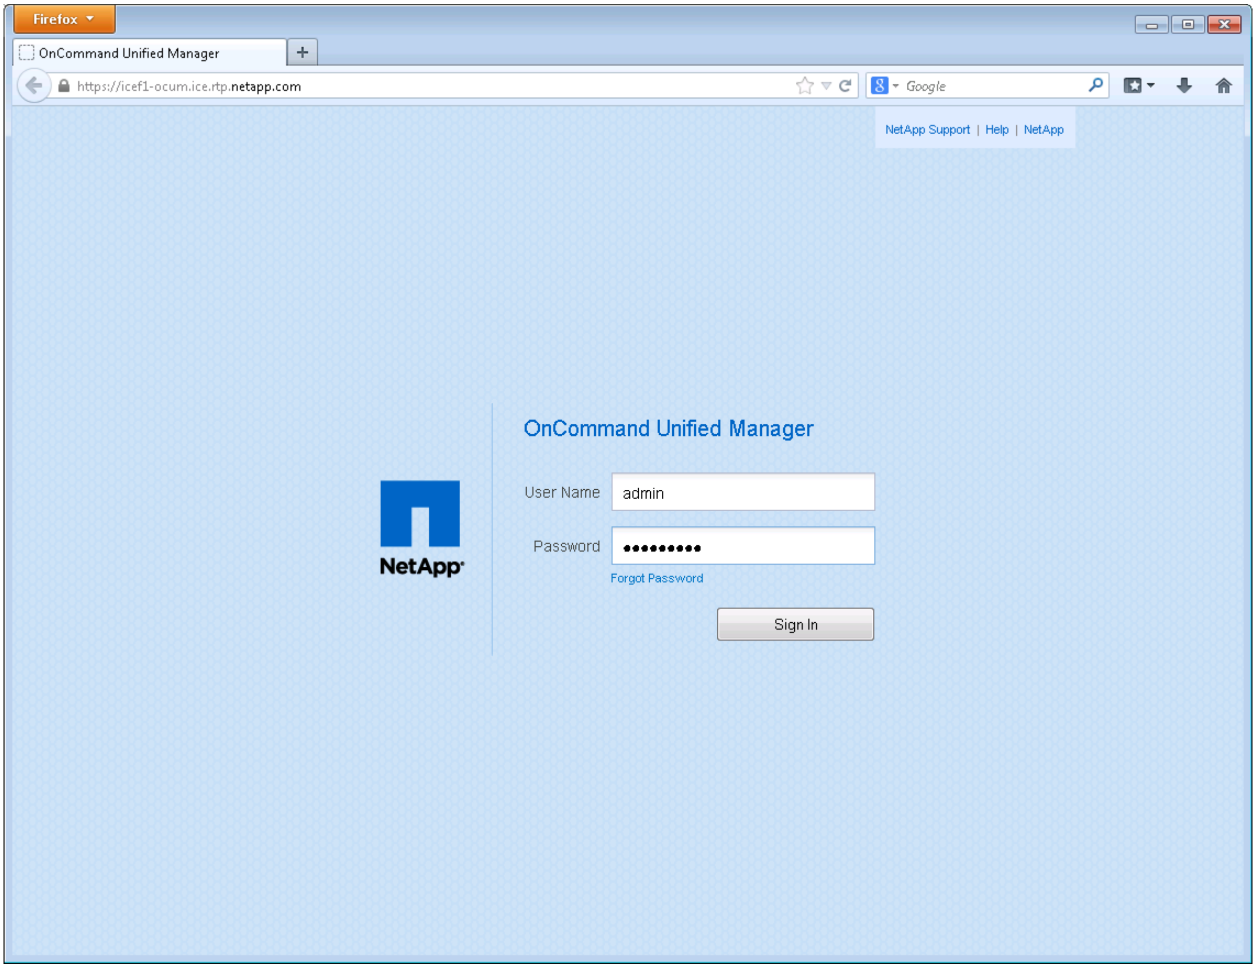

Firefox
OnCommand Unified Manager
https://icef1-ocum.ice.rtp.netapp.com
Google
NetApp Support | Help | NetApp
NetApp
OnCommand Unified Manager
Forgot Password
Sign In

User Name: admin

Password: ●●●●●●●●●

6. Log in using the Maintenance User account credentials.

7. Select `Yes` option to enable AutoSupport capabilities.

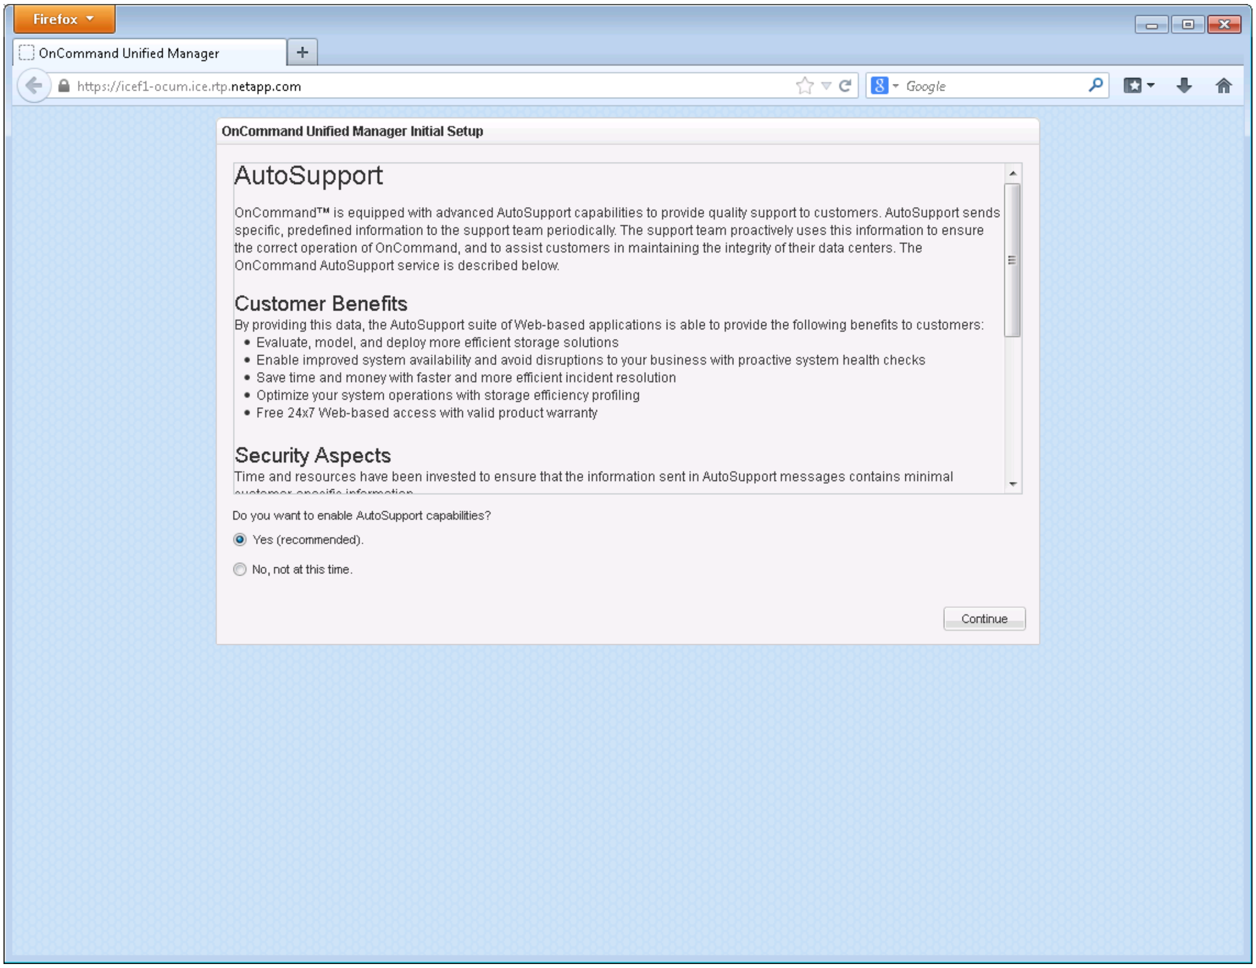
Firefox
OnCommand Unified Manager
https://icef1-ocum.ice.rtp.netapp.com
Google

OnCommand Unified Manager Initial Setup

AutoSupport

OnCommand™ is equipped with advanced AutoSupport capabilities to provide quality support to customers. AutoSupport sends specific, predefined information to the support team periodically. The support team proactively uses this information to ensure the correct operation of OnCommand, and to assist customers in maintaining the integrity of their data centers. The OnCommand AutoSupport service is described below.

Customer Benefits

By providing this data, the AutoSupport suite of Web-based applications is able to provide the following benefits to customers:

- Evaluate, model, and deploy more efficient storage solutions
- Enable improved system availability and avoid disruptions to your business with proactive system health checks
- Save time and money with faster and more efficient incident resolution
- Optimize your system operations with storage efficiency profiling
- Free 24x7 Web-based access with valid product warranty

Security Aspects

Time and resources have been invested to ensure that the information sent in AutoSupport messages contains minimal customer specific information.

Do you want to enable AutoSupport capabilities?

● Yes (recommended).

○ No, not at this time.

8. Click Continue.

9. Provide the NTP Server IP address `<<var_global_ntp_server_ip>>`

10. Provide the Maintenance User Email `<<var_storage_admin_email>>`

11. Provide the SMTP Server Hostname.

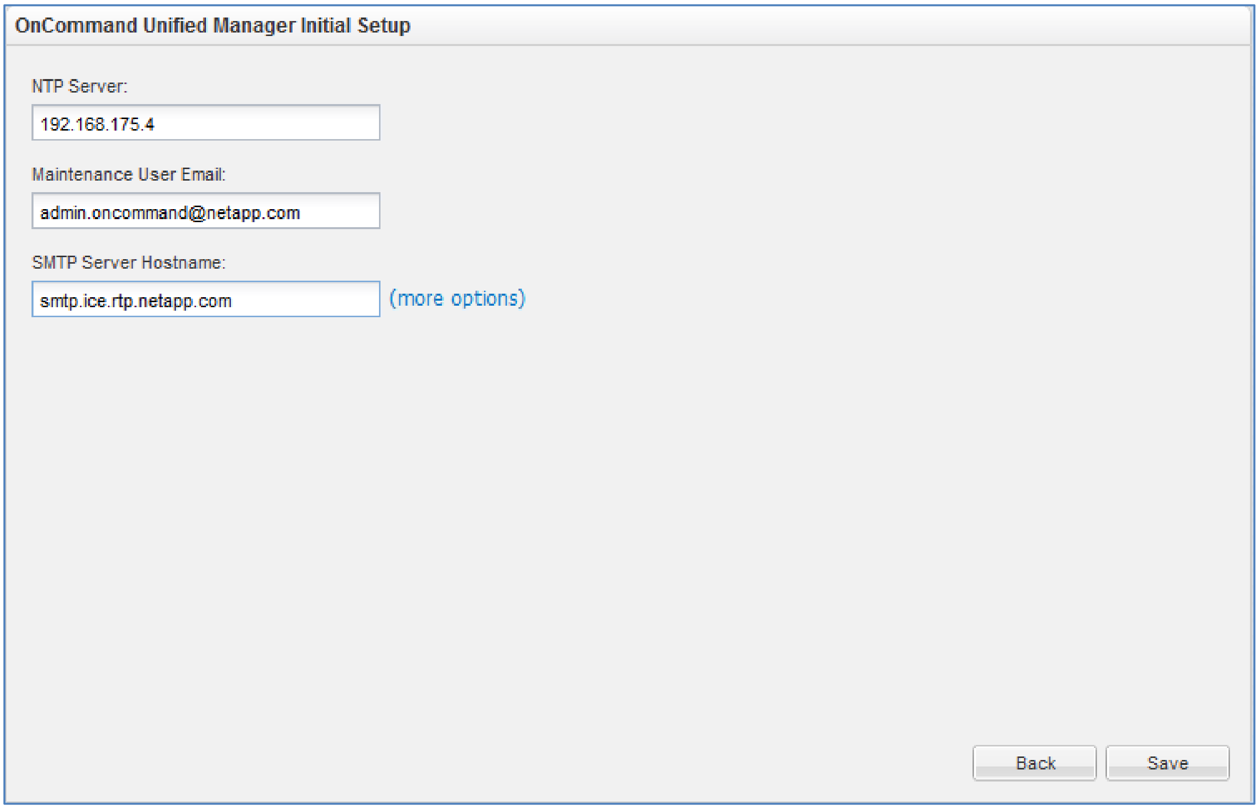

OnCommand Unified Manager Initial Setup

NTP Server: 192.168.175.4

Maintenance User Email: admin.oncommand@netapp.com

SMTP Server Hostname: smtp.ice.rtp.netapp.com

12. Click Save.

13. Click Add Cluster

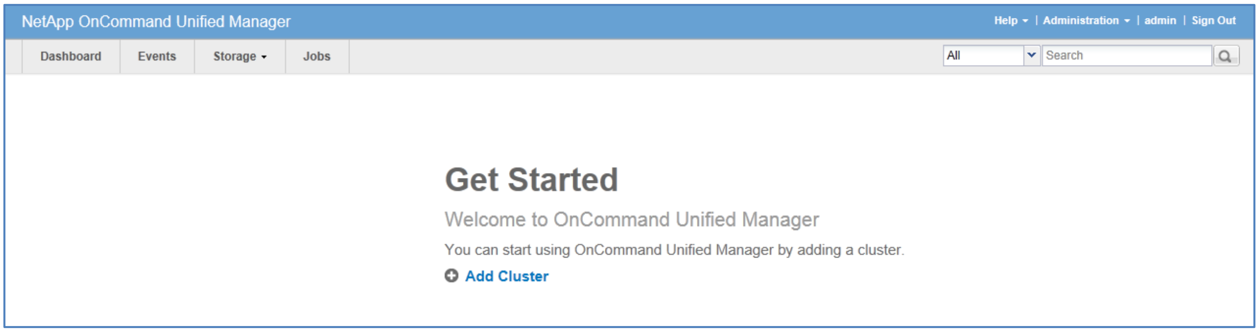

NetApp OnCommand Unified Manager

Get Started

Welcome to OnCommand Unified Manager

You can start using OnCommand Unified Manager by adding a cluster.

14. Provide the Cluster Management IP address, User Name, Password, Protocol, and Port.

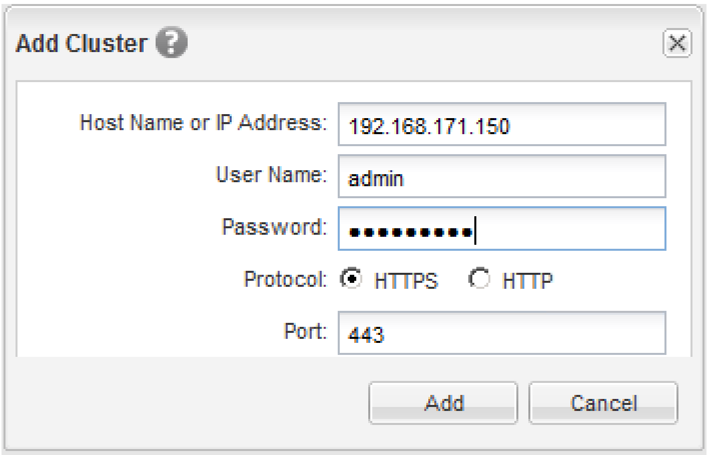

Add Cluster

Host Name or IP Address: 192.168.171.150

User Name: admin

Password: ••••••••••

Protocol: ● HTTPS ○ HTTP

Port: 443

15. Click Add.

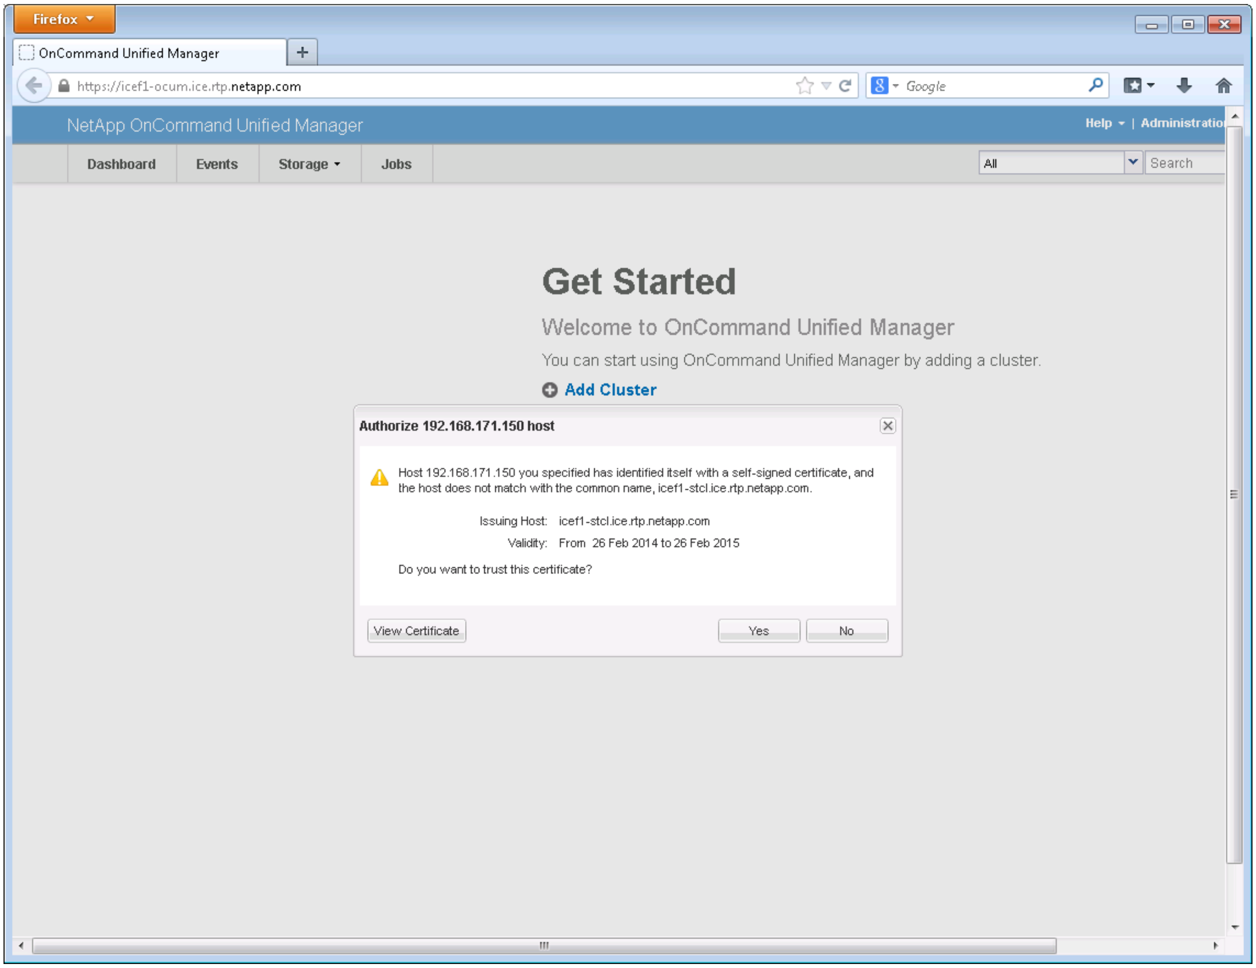

Firefox
OnCommand Unified Manager
https://icef1-ocum.ice.rtp.netapp.com
Google
NetApp OnCommand Unified Manager
Help | Administratio
Dashboard
Events
Storage
Jobs
All
Search

Get Started

Welcome to OnCommand Unified Manager

You can start using OnCommand Unified Manager by adding a cluster.

Add Cluster

Authorize 192.168.171.150 host

Host 192.168.171.150 you specified has identified itself with a self-signed certificate, and the host does not match with the common name, icef1-stcl.ice.rtp.netapp.com.

Issuing Host: icef1-stcl.ice.rtp.netapp.com

Validity: From 26 Feb 2014 to 26 Feb 2015

Do you want to trust this certificate?

16. Click Yes to trust the certificate from the controller.

Note: The Cluster Add operation might take a couple of minutes.

17. After the cluster is added it can be accessed by clicking on the Storage tab and selecting Clusters.

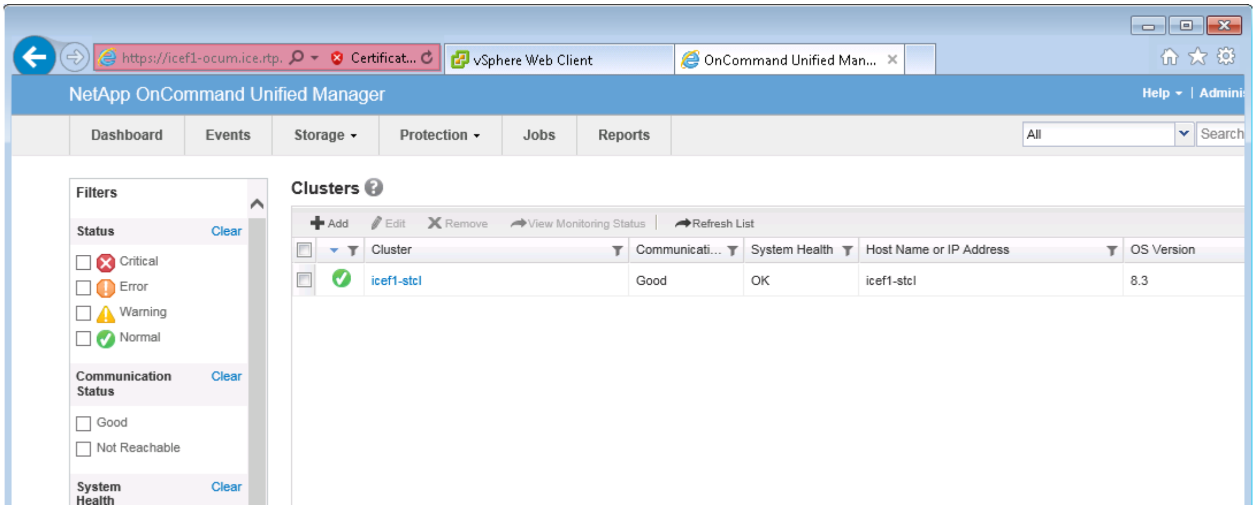

Cluster	Communicati...	System Health	Host Name or IP Address	OS Version
icef1-stcl	Good	OK	icef1-stcl	8.3

OnCommand Performance Manager 1.1

OnCommand Performance Manager OVF Deployment

To install the OnCommand Performance Manager, complete the following steps:

Download and review the OnCommand Performance Manager 1.1 Installation and Administration Guide for VMware Virtual Appliances.

1. Download the OnCommand Performance Manager version 1.1 (OnCommandPerformanceManager-netapp-1.1.0.ova), from http://mysupport.netapp.com/NOW/download/software/oncommand_pm/1.1/OnCommandPerformanceManager-netapp-1.1.0.ova

2. Log in to the vSphere Web Client. Go to vCenter > VMs and Templates.

3. At the top of the center pane, click Actions > Deploy OVF Template.

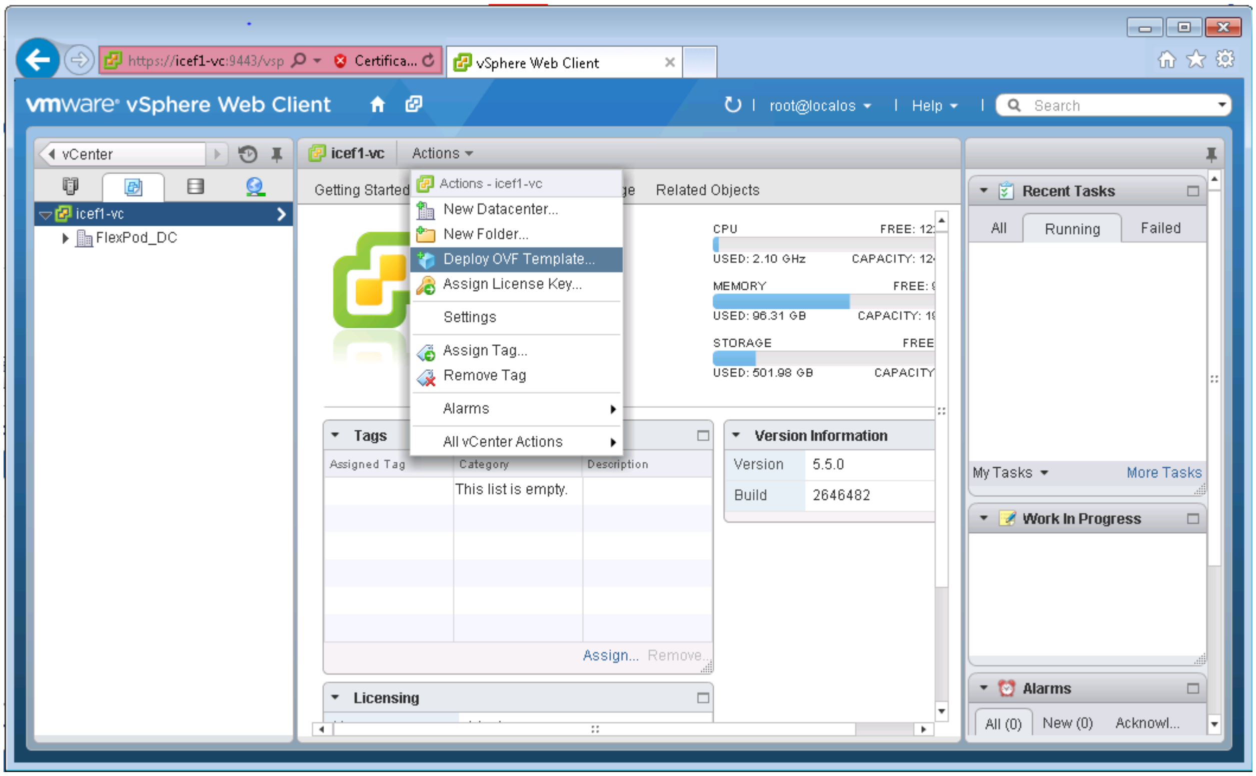

4. Browse the `.ova` file that was downloaded locally. Click Open to select the file. Click Next to proceed with the selected file.

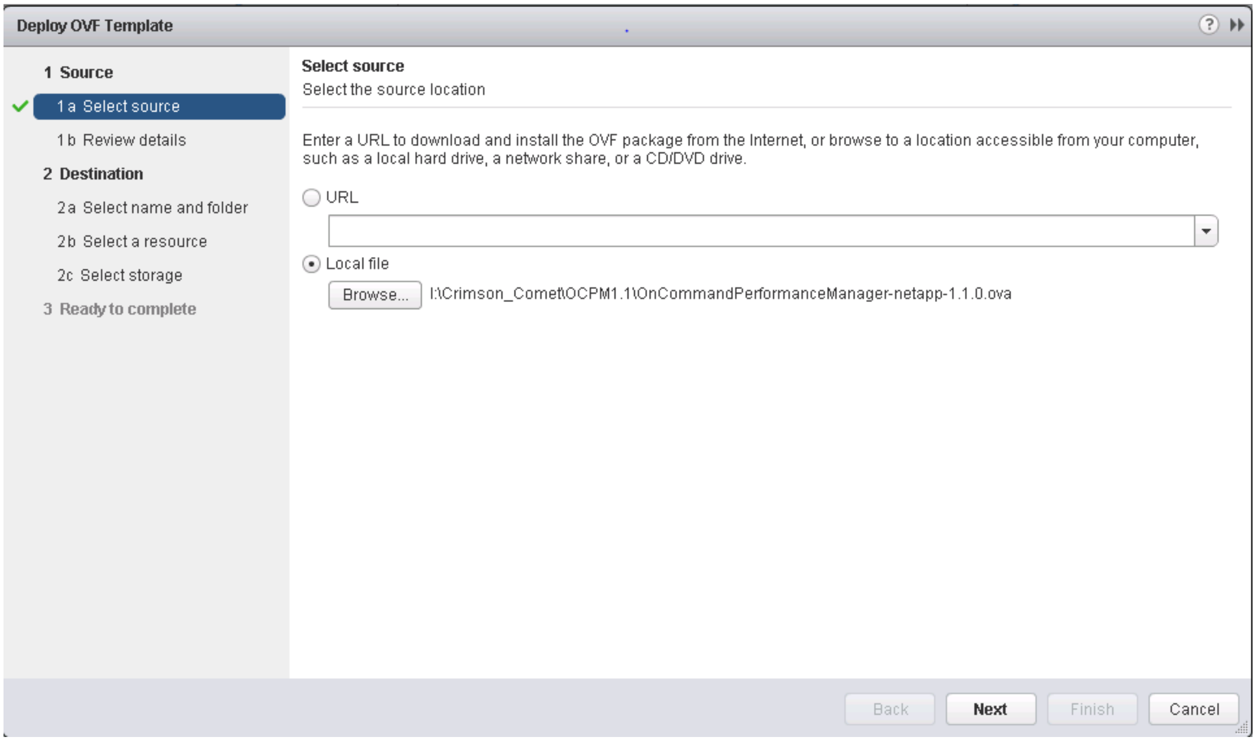

Select source

Select the source location

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

○ URL

● Local file

Browse... I:\Crimson_Comet\OCPM1.1\OnCommandPerformanceManager-netapp-1.1.0.ova

5. Click the checkbox to accept the additional configuration options and click Next.

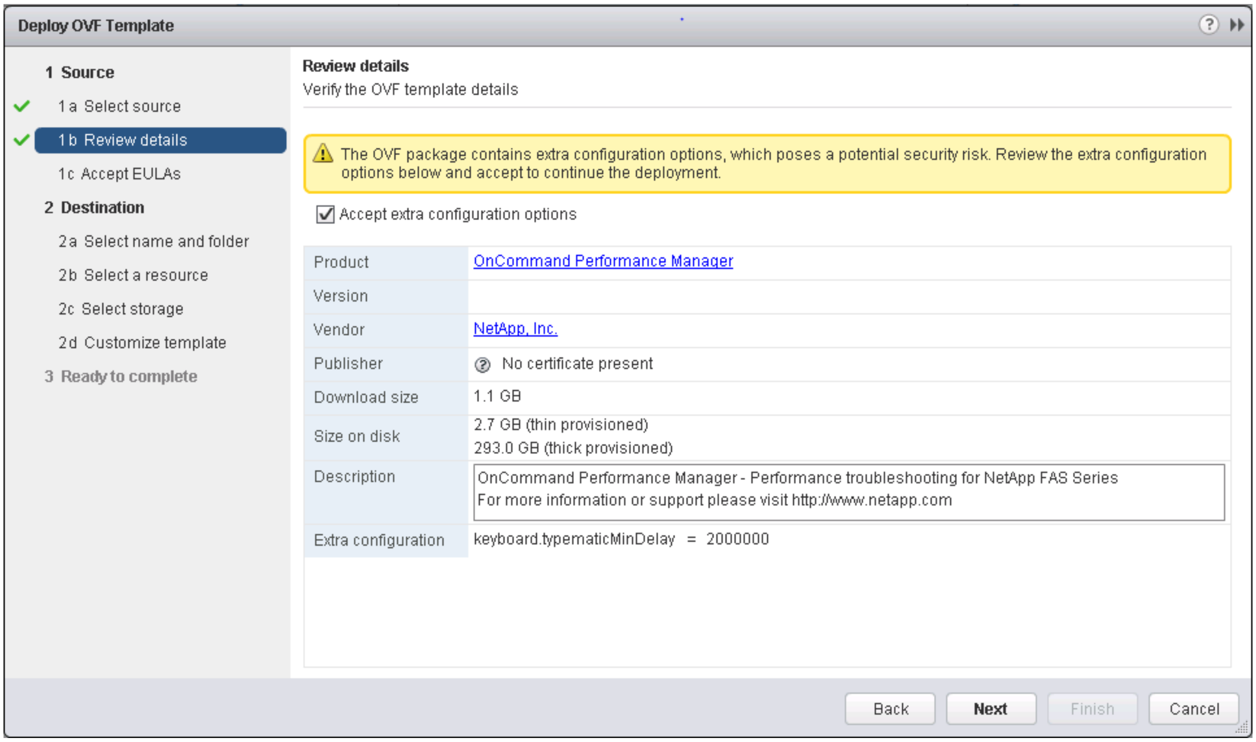

Review details

Verify the OVF template details

⚠ The OVF package contains extra configuration options, which poses a potential security risk. Review the extra configuration options below and accept to continue the deployment.

☑ Accept extra configuration options

Product	OnCommand Performance Manager
Version	
Vendor	NetApp, Inc.
Publisher	No certificate present
Download size	1.1 GB
Size on disk	2.7 GB (thin provisioned) 293.0 GB (thick provisioned)
Description	OnCommand Performance Manager - Performance troubleshooting for NetApp FAS Series For more information or support please visit http://www.netapp.com
Extra configuration	keyboard.typematicMinDelay = 2000000

6. Read the EULA, and then click the Accept button to accept the agreement. Click Next to continue.

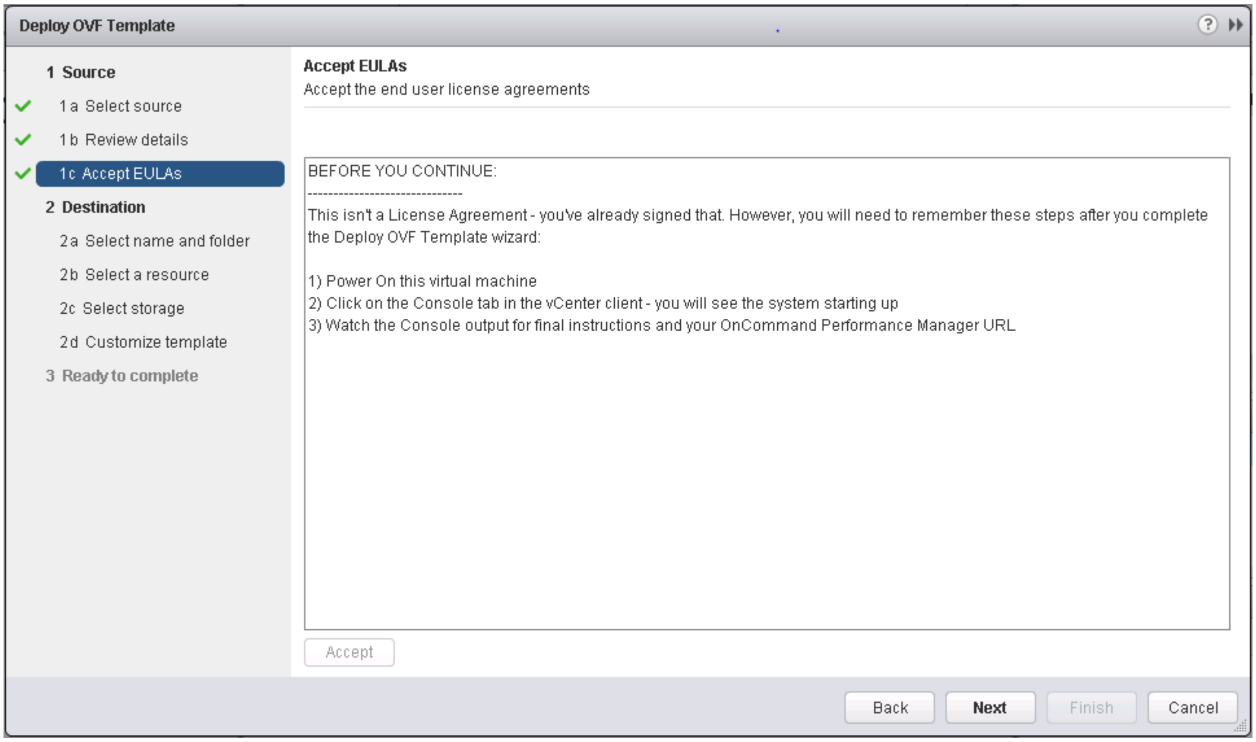

Accept EULAs

Accept the end user license agreements

```
BEFORE YOU CONTINUE:
------------------------------
This isn't a License Agreement - you've already signed that. However, you will need to remember these steps after you complete the Deploy OVF Template wizard:

1) Power On this virtual machine
2) Click on the Console tab in the vCenter client - you will see the system starting up
3) Watch the Console output for final instructions and your OnCommand Performance Manager URL
```

7. Enter the name of the VM and select the `FlexPod_DC_1` folder to hold the VM. Click Next to continue.

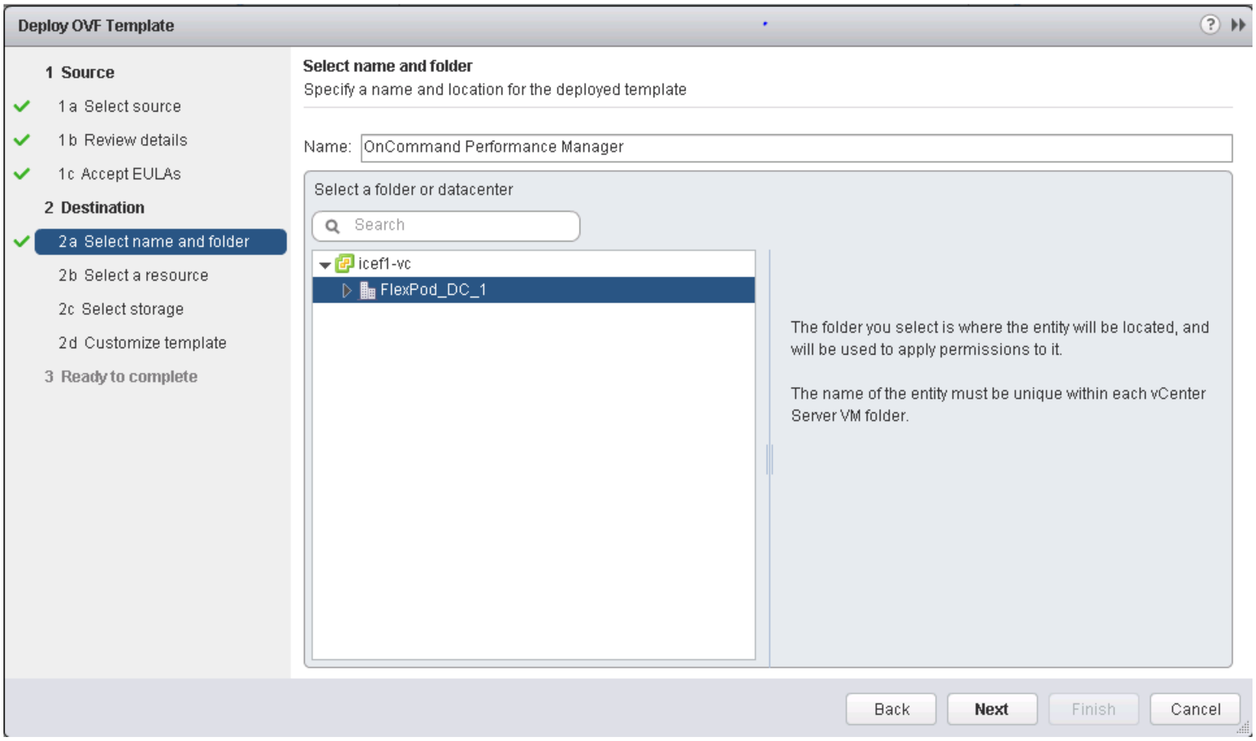

Select name and folder

Specify a name and location for the deployed template

Name: OnCommand Performance Manager

Select a folder or datacenter

The folder you select is where the entity will be located, and will be used to apply permissions to it.

The name of the entity must be unique within each vCenter Server VM folder.

8. Select FlexPod_Management within the FlexPod_DC_1 datacenter as the destination compute resource pool to host the VM. Click Next to continue.

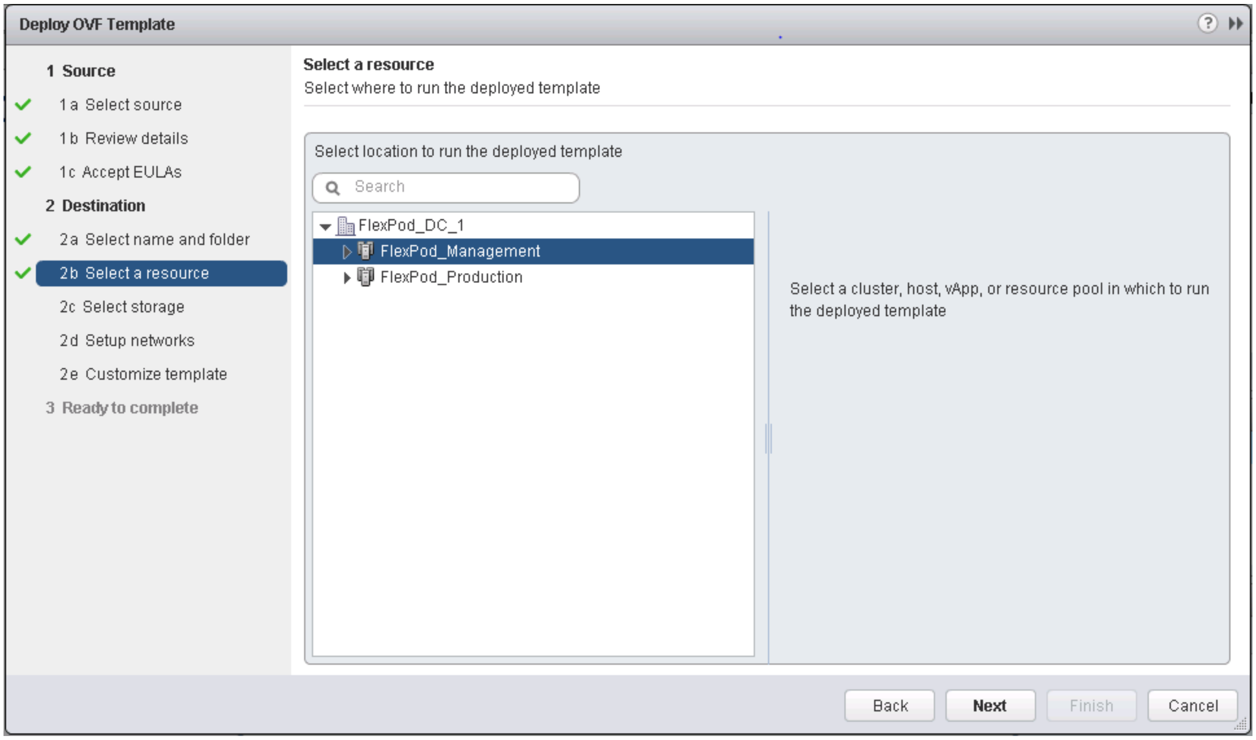

Select a resource

Select where to run the deployed template

Select location to run the deployed template

Select a cluster, host, vApp, or resource pool in which to run the deployed template

9. Select infra_datastore_1 as the storage target for the VM and select Thin Provision as the Virtual disk format. Click Next to continue.

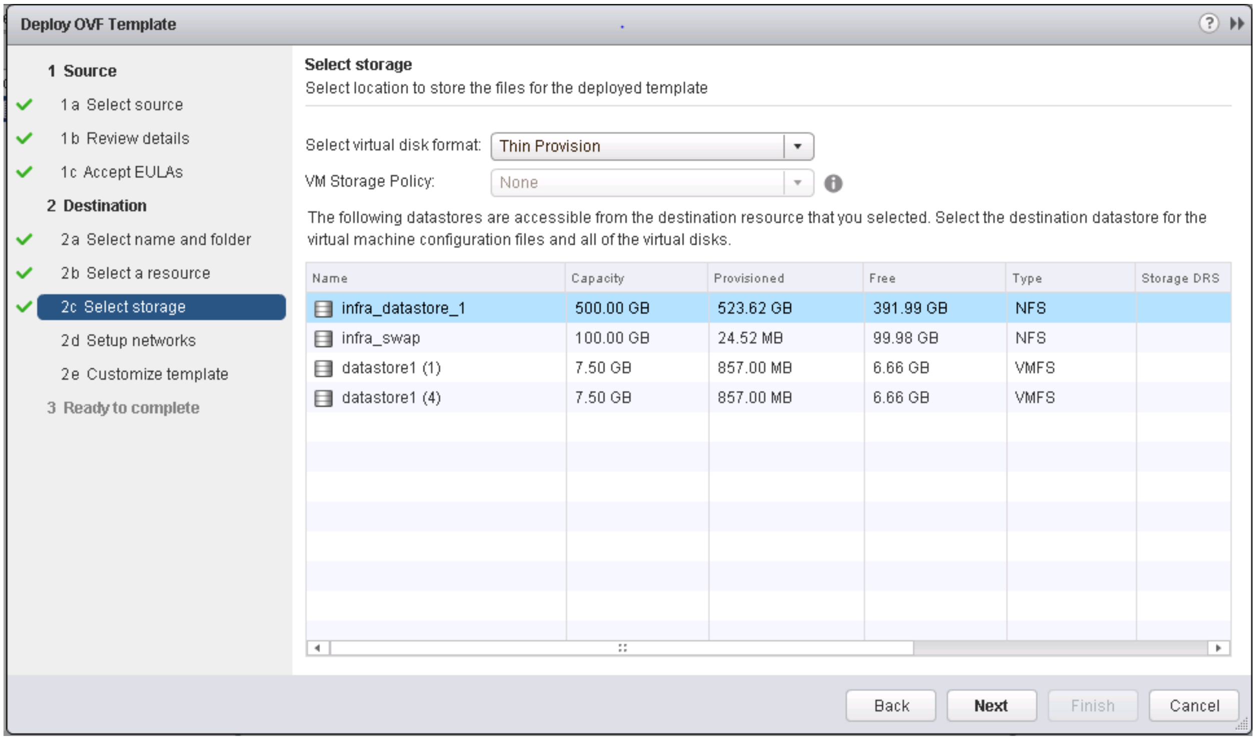

Select storage
Select location to store the files for the deployed template

Select virtual disk format: Thin Provision
VM Storage Policy: None

The following datastores are accessible from the destination resource that you selected. Select the destination datastore for the virtual machine configuration files and all of the virtual disks.

Name	Capacity	Provisioned	Free	Type	Storage DRS
infra_datastore_1	500.00 GB	523.62 GB	391.99 GB	NFS	
infra_swap	100.00 GB	24.52 MB	99.98 GB	NFS	
datastore1 (1)	7.50 GB	857.00 MB	6.66 GB	VMFS	
datastore1 (4)	7.50 GB	857.00 MB	6.66 GB	VMFS	

10. Select `MGMT Network` as the destination network to the nat source network. Click Next.

11. Fill out the details for the Host Name, IP Address, Network Mask, Gateway, Primary DNS, and Secondary DNS. Click Next to continue.

12. Clear the Power on after deployment checkbox.

13. Review the configuration details. Click Finish to begin deploying the VM with the provided configuration details.

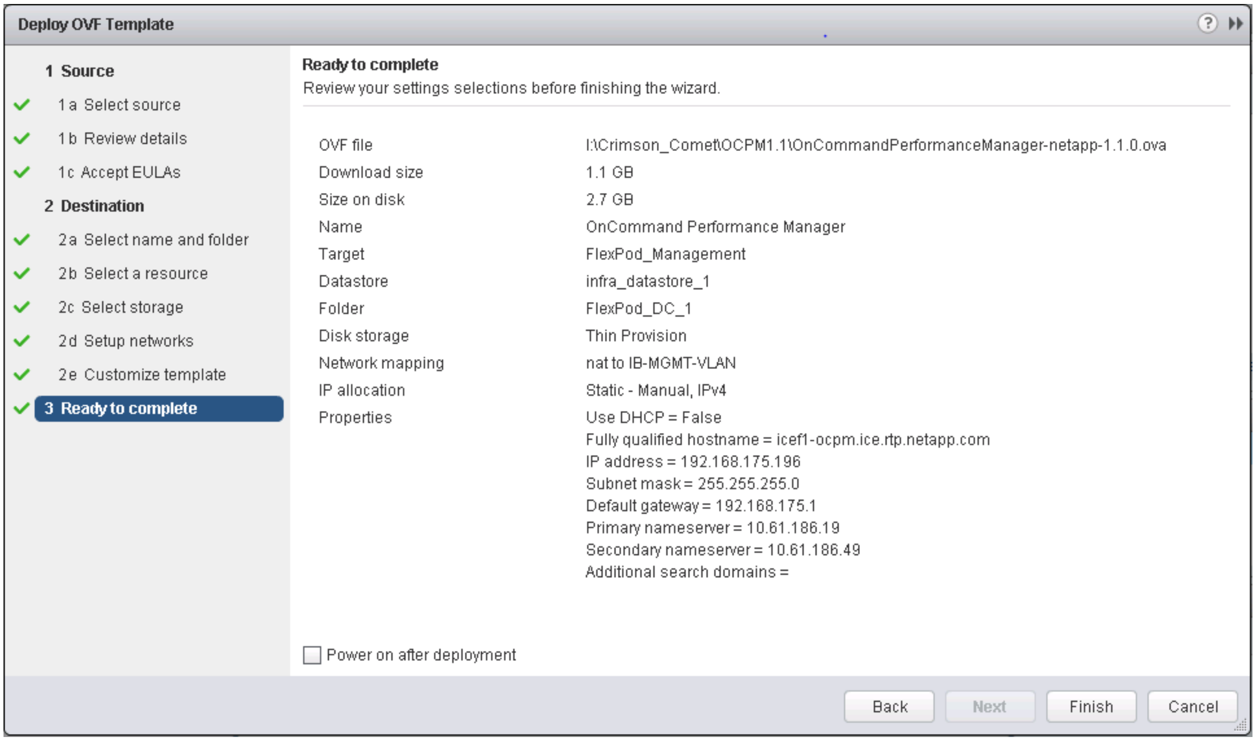

Ready to complete

Review your settings selections before finishing the wizard.

OVF file	I:\Crimson_Comet\OCPM1.1\OnCommandPerformanceManager-netapp-1.1.0.ova
Download size	1.1 GB
Size on disk	2.7 GB
Name	OnCommand Performance Manager
Target	FlexPod_Management
Datastore	infra_datastore_1
Folder	FlexPod_DC_1
Disk storage	Thin Provision
Network mapping	nat to IB-MGMT-VLAN
IP allocation	Static - Manual, IPv4
Properties	Use DHCP = False Fully qualified hostname = icef1-ocpm.ice.rtp.netapp.com IP address = 192.168.175.196 Subnet mask = 255.255.255.0 Default gateway = 192.168.175.1 Primary nameserver = 10.61.186.19 Secondary nameserver = 10.61.186.49 Additional search domains =

□ Power on after deployment

14. On the left pane, navigate to vCenter -> Virtual machines. After OVF deployment is complete, right-click the newly created virtual machine and select Edit Settings.

15. Click the CPU tab to expand the CPU options:

    a. The minimum required CPU Reservation is 9572 MHz. Determine the CPU frequency of the host.

    b. Set the number of CPUs to the number of CPUs required (9572 / CPU Frequency of host).

    c. Set the number of Cores per Socket where the Sockets number on the right matches the number of CPU sockets in the host. For example, if a host has 2 CPUs operating at a speed of 1999MHz, then the VM would need 6 virtual CPUs (9572 / 1999 = 4.79 – rounded to 6 virtual CPUs). If the host has 2 physical CPU sockets, 3 Cores per Socket.

Note: Use the OnCommand Performance Manager 1.1 Installation and Administration Guide for VMware Virtual Appliances for guidance on these settings.

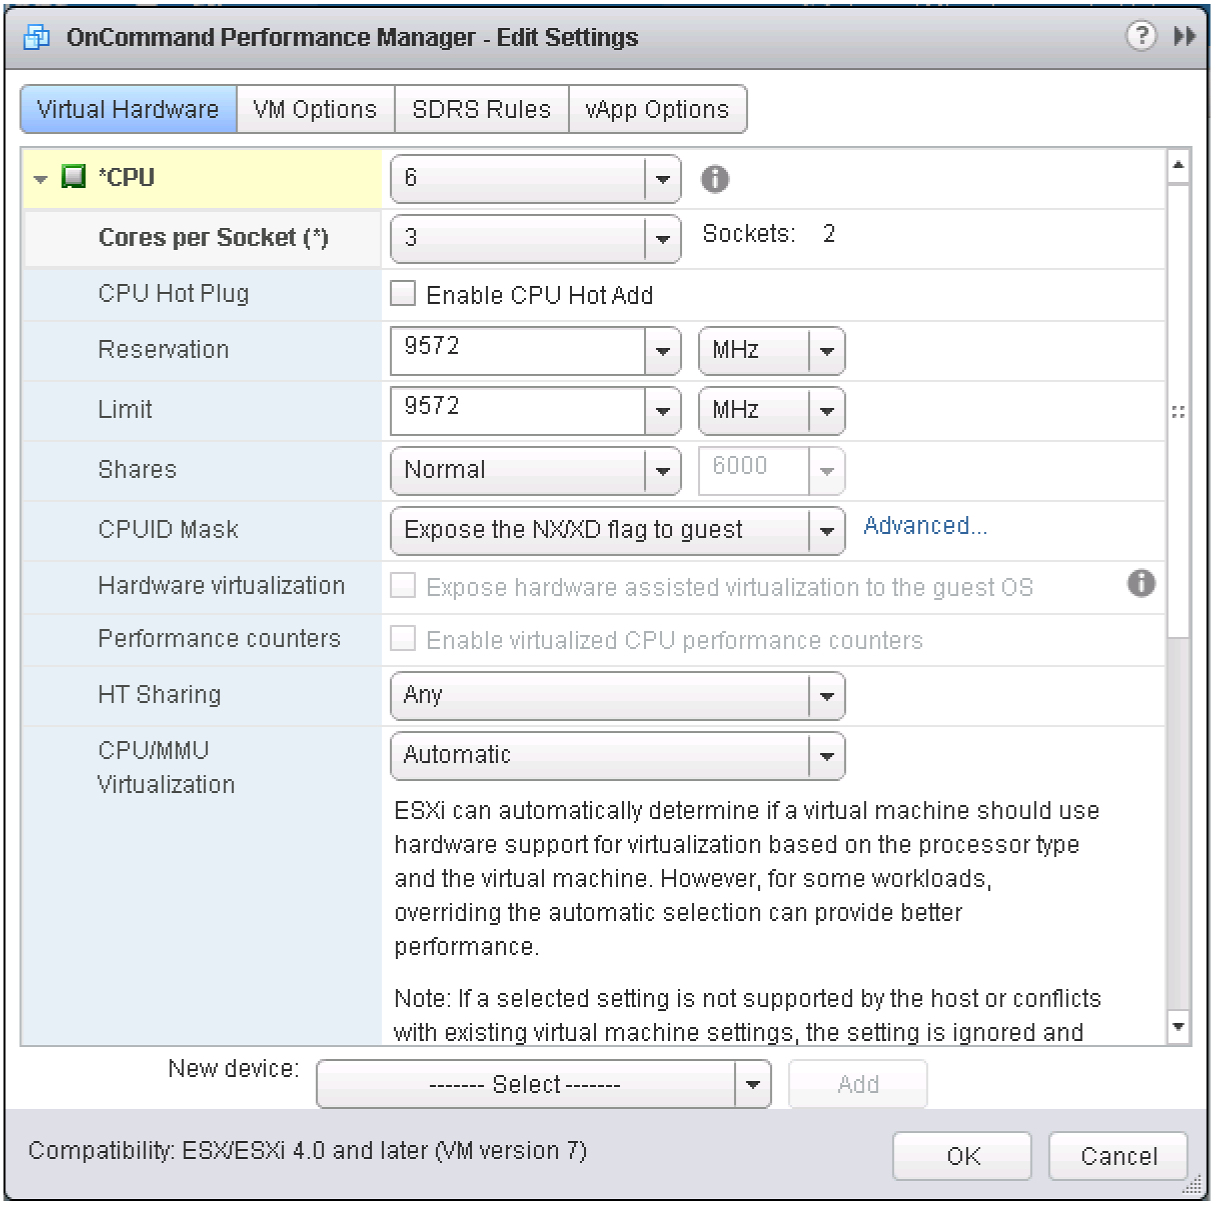

OnCommand Performance Manager - Edit Settings

*CPU	6	
Cores per Socket (*)	3	Sockets: 2
CPU Hot Plug	□ Enable CPU Hot Add	
Reservation	9572	MHz
Limit	9572	MHz
Shares	Normal	6000
CPUID Mask	Expose the NX/XD flag to guest	Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS	
Performance counters	□ Enable virtualized CPU performance counters	
HT Sharing	Any	
CPU/MMU Virtualization	Automatic	

ESXi can automatically determine if a virtual machine should use hardware support for virtualization based on the processor type and the virtual machine. However, for some workloads, overriding the automatic selection can provide better performance.

Note: If a selected setting is not supported by the host or conflicts with existing virtual machine settings, the setting is ignored and

New device: ------- Select ------- Add

Compatibility: ESX/ESXi 4.0 and later (VM version 7)

16. Click OK to accept the changes.

17. Right-click the VM in the left-hand pane. Click Power On.

OnCommand Performance Manager Basic Setup

1. Select the VM in the left-hand pane. In the center pane, select Open with VMRC.

2. In the VMRC window, select Manage > Install VMware Tools. VMware Tools will install in the VM.

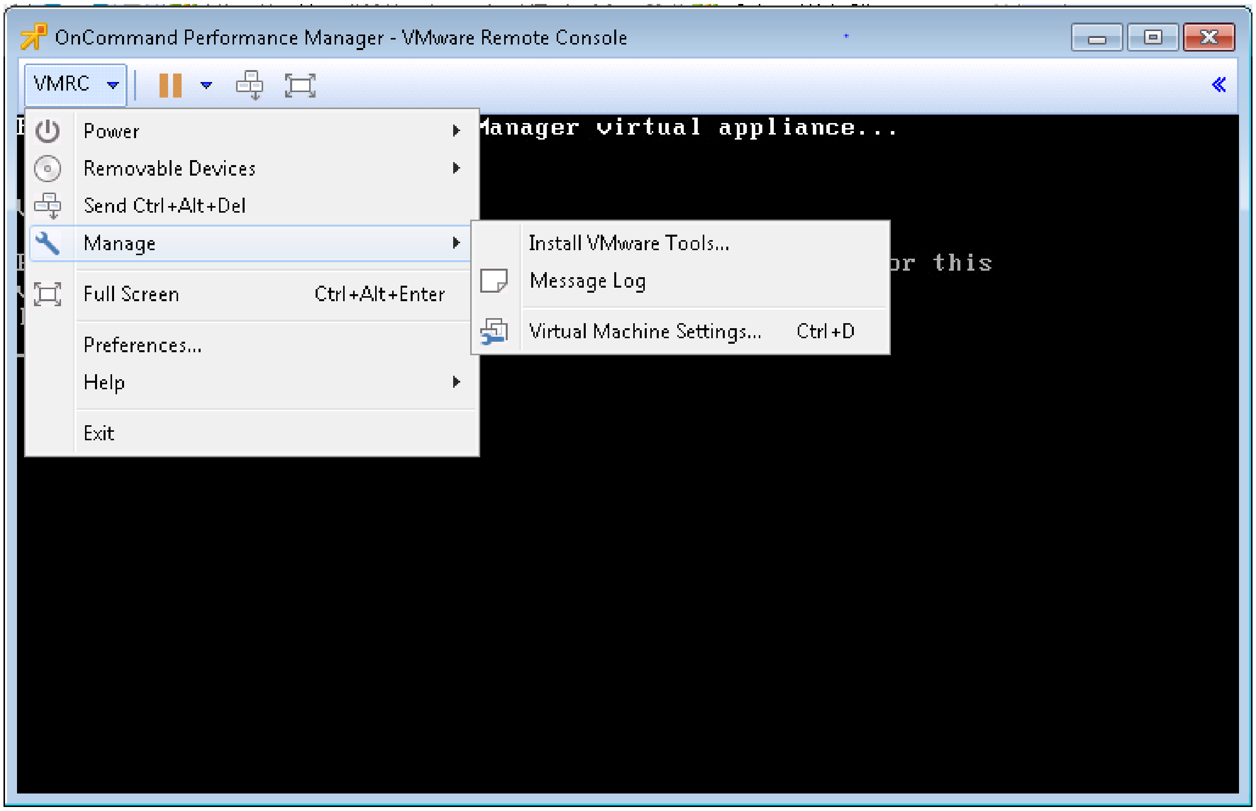
OnCommand Performance Manager - VMware Remote Console
VMRC
Power
Removable Devices
Send Ctrl+Alt+Del
Manage
Full Screen Ctrl+Alt+Enter
Preferences...
Help
Exit
Install VMware Tools...
Message Log
Virtual Machine Settings... Ctrl+D

```
Manager virtual appliance...
or this
```

3. Set up OnCommand Performance Manager by answering the following questions in the console window:

```
Geographic area: <<Enter your geographic location>>

Time zone: <<Select the city or region corresponding to your time zone>>
```

These commands complete the network configuration checks, generates SSL certificates and starts the OnCommand Performance Manager services.

1. To Create a Maintenance User account, run the following commands:

Note: The maintenance user manages and maintains the settings on the OnCommand Performance Manager virtual appliance.

```
Username : admin

Enter new UNIX password: <<var_password>>

Retype new UNIX password: <<var_password>>
```

2. Using a web browser navigate to the OnCommand Performance Manager using URL: `https://<<var_oncommand_pm_ip>>`.

3. Log in using the Maintenance User account (admin) credentials.

4. Enter a Maintenance User Email Address, SMTP Mail Server information, and the NTP server IP address. Click Save and go to next step.

5. Select `Yes` option to enable AutoSupport capabilities. Click Save and go to next step.

6. Click Save and go to next step.

7. Enter the storage cluster host name or IP address, the storage cluster admin user name, and the storage cluster admin password. Click Add Cluster, then click Save and Complete Configuration.

8. After the cluster is added it can be accessed by clicking on Administration > Manage Data Sources.

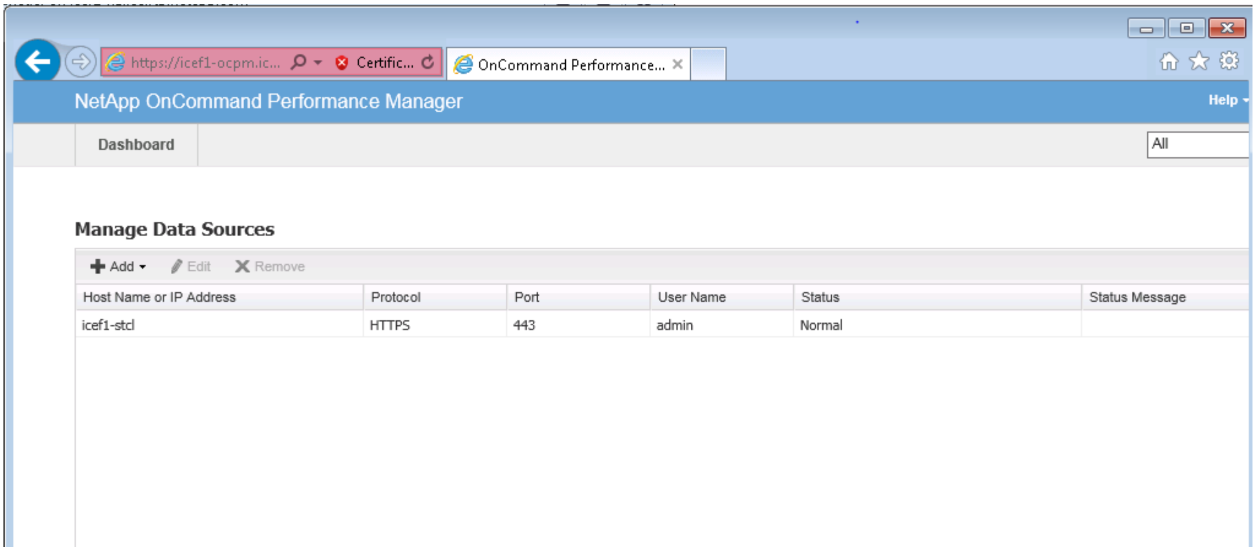

Host Name or IP Address	Protocol	Port	User Name	Status	Status Message
icef1-stcl	HTTPS	443	admin	Normal	

Link OnCommand Performance Manager to OnCommand Unified Manager

1. Using a web browser navigate to the OnCommand Unified Manager using URL: `https://<<var_oncommand_server_ip>>`. Log in with the Maintenance user id and password setup earlier.

2. In the OnCommand Unified Manager web interface, select Administrator > Manage Users to set up an Event Publication user.

3. Click Add to add a user.

4. Leave the Type set to Local User. Use eventpub as the Name and enter and confirm a password. Enter an email address for this user and set the Role to Event Publisher. Click Add.

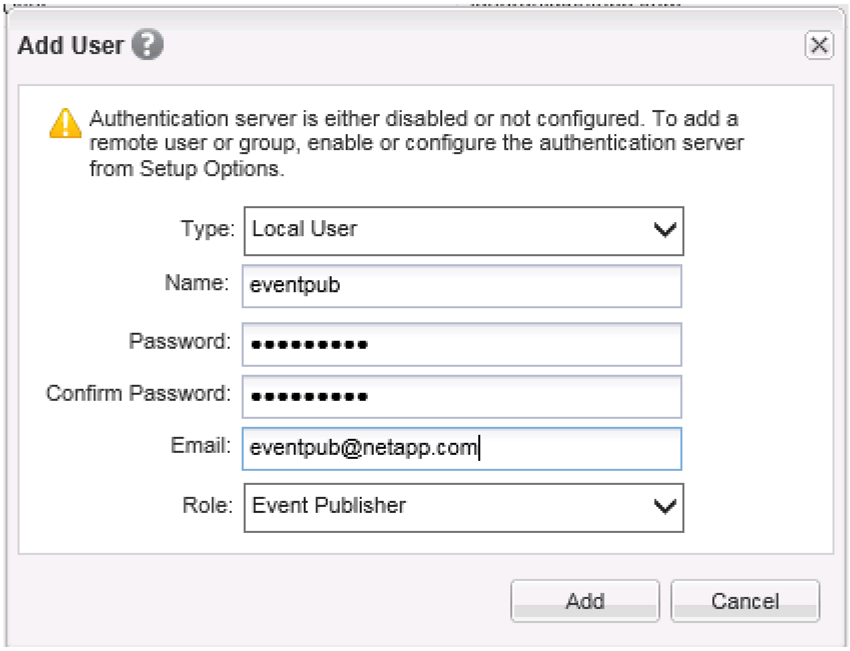

Add User

Authentication server is either disabled or not configured. To add a remote user or group, enable or configure the authentication server from Setup Options.

Type: Local User

Name: eventpub

Password: ••••••••••

Confirm Password: ••••••••••

Email: eventpub@netapp.com

Role: Event Publisher

5. At the OnCommand Performance Manager console window, log into the Command Line Interface with the Maintenance User (admin) defined earlier.

6. Enter 5 to select Unified Manager Connection.

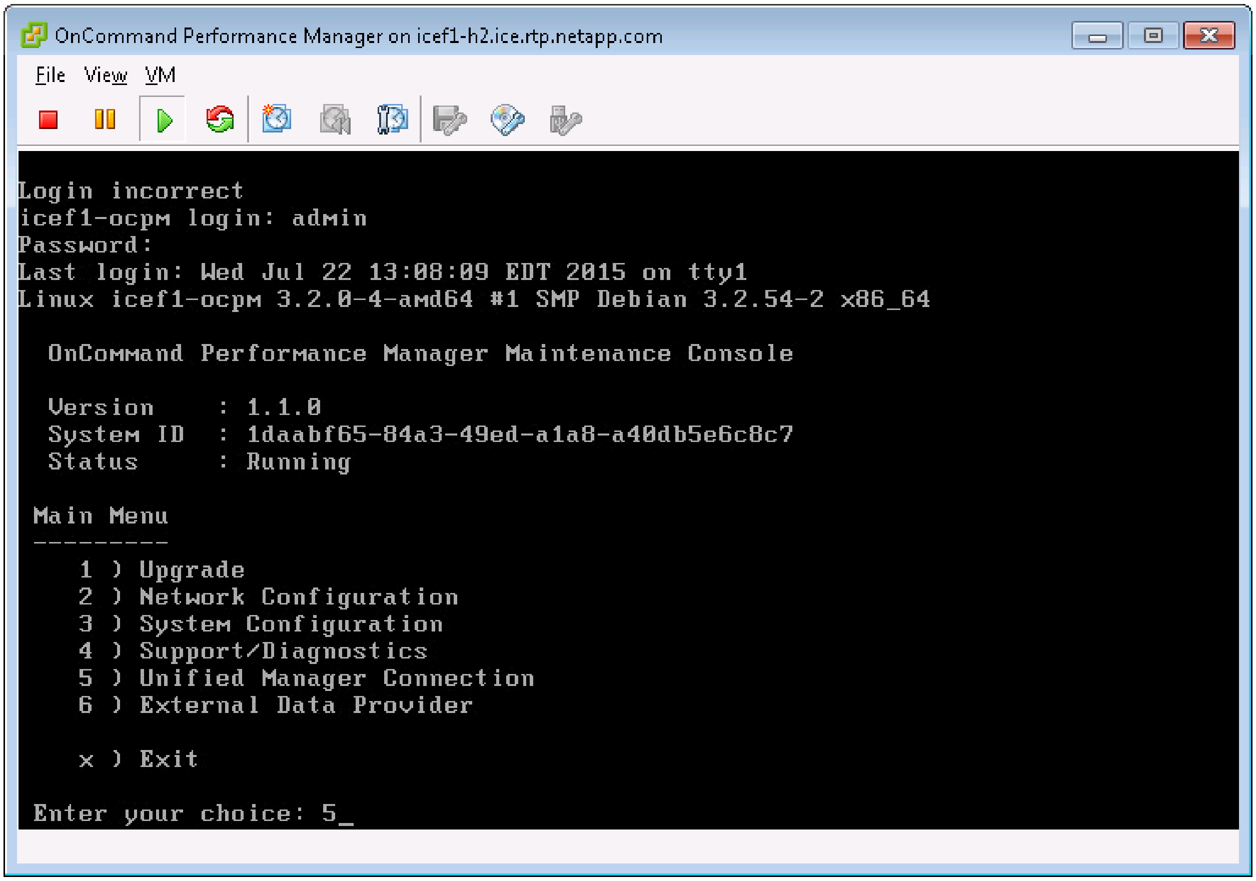

OnCommand Performance Manager on icef1-h2.ice.rtp.netapp.com

```
Login incorrect
icef1-ocpm login: admin
Password:
Last login: Wed Jul 22 13:08:09 EDT 2015 on tty1
Linux icef1-ocpm 3.2.0-4-amd64 #1 SMP Debian 3.2.54-2 x86_64

  OnCommand Performance Manager Maintenance Console

  Version     : 1.1.0
  System ID   : 1daabf65-84a3-49ed-a1a8-a40db5e6c8c7
  Status      : Running

 Main Menu
 ---------
    1 ) Upgrade
    2 ) Network Configuration
    3 ) System Configuration
    4 ) Support/Diagnostics
    5 ) Unified Manager Connection
    6 ) External Data Provider

    x ) Exit

 Enter your choice: 5_
```

7. Enter 2 to Add / Modify Unified Manager Server Connection.

8. Enter y to continue.

9. Enter the OnCommand Unified Manager FQDN or IP Address.

10. Hit Enter to accept the default port 443.

11. Enter y to accept the Unified Manager Security Certificate.

12. Enter eventpub for the Event Publisher User Name.

13. Enter the eventpub password.

14. Enter y to accept the entered settings.

15. Select `Yes` option to enable AutoSupport capabilities. Click Save and go to next step.

16. Press any key to continue.

17. Exit the OnCommand Performance Manager console. OnCommand Performance Manager events will now appear in the OnCommand Unified Manager Dashboard.

NetApp NFS Plug-In 1.0.21 for VMware VAAI

Enable VMware vStorage for NFS in Clustered Data ONTAP

To enable VMware vStorage for NFS in clustered Data ONTAP, complete the following steps:

1. From an SSH session to the storage cluster management address, log in with the admin user name and password.

2. Enable vStorage on the Vserver.

```
vserver nfs modify –vserver Infra-SVM -vstorage enabled
```

3. Verify that the export policy rules are set up correctly.

```
vserver export-policy rule show –vserver Infra-SVM
```

Sample output:

```
NetApp::> vserver export-policy rule show -vserver Infra-SVM
             Policy          Rule    Access   Client                RO
Vserver      Name            Index   Protocol Match                 Rule
------------ --------------- ------  -------- --------------------- ---------
Infra-SVM    default         1       nfs      192.168.170.61        sys
Infra-SVM    default         2       nfs      192.168.170.60        sys
Infra-SVM    default         3       nfs      192.168.170.58        sys
Infra-SVM    default         4       nfs      192.168.170.59        sys
```

```
Infra-SVM     default          5        nfs       192.168.170.62         sys
Infra-SVM     default          6        nfs       192.168.170.63         sys
6 entries were displayed.
```

4. The access protocol for the FlexPod policy name should be NFS. If the access protocol is not nfs for a given rule index, run the following command to set NFS as the access protocol:

```
vserver export-policy rule modify –vserver Infra-SVM –policyname default –
ruleindex <<var_rule_index>> -protocol nfs
```

Install NetApp NFS Plug-In for VMware VAAI

To install the NetApp NFS plug-in for VMware vStorage APIs for Array Integration (VAAI), complete the following steps:

1. From a console interface on the NetApp VSC VM, go to the Software Downloads page in the NetApp Support site.
2. Scroll down to locate the NetApp NFS Plug-in for VMware VAAI, select the ESXi5.x platform, and click Go.
3. Click View & Download.
4. Click CONTINUE.
5. Click Accept.
6. Download the `.vib` file of the most recent plug-in version to the VSC VM Desktop as `NetAppNasPlugin.vib`.

Note: It is important that the file be saved as `NetAppNasPlugin.vib`.

7. On the VSC VM Desktop, move the NetAppNasPlugin.vib file to the C:\Program Files\NetApp\Virtual Storage Console\etc\vsc\web folder.

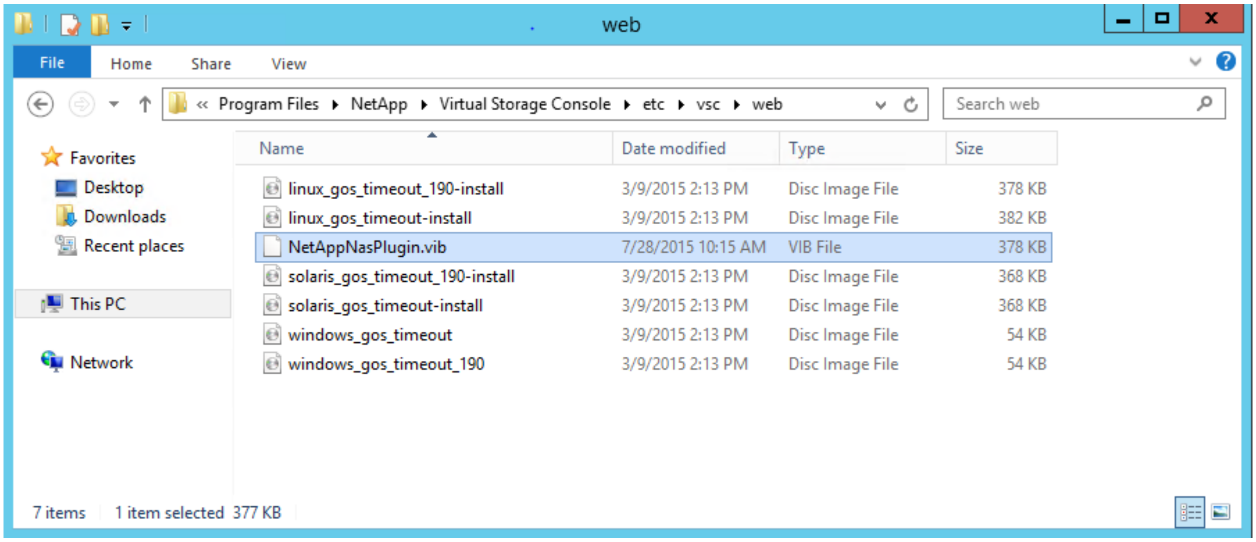

Name	Date modified	Type	Size
linux_gos_timeout_190-install	3/9/2015 2:13 PM	Disc Image File	378 KB
linux_gos_timeout-install	3/9/2015 2:13 PM	Disc Image File	382 KB
NetAppNasPlugin.vib	7/28/2015 10:15 AM	VIB File	378 KB
solaris_gos_timeout_190-install	3/9/2015 2:13 PM	Disc Image File	368 KB
solaris_gos_timeout-install	3/9/2015 2:13 PM	Disc Image File	368 KB
windows_gos_timeout	3/9/2015 2:13 PM	Disc Image File	54 KB
windows_gos_timeout_190	3/9/2015 2:13 PM	Disc Image File	54 KB

8. Go to the VMware vSphere Web Client and select VSC. Click NFS VAAI Tools. Make sure NFS Plug-in for VMware VAAI Version: 1.0-21 is shown.

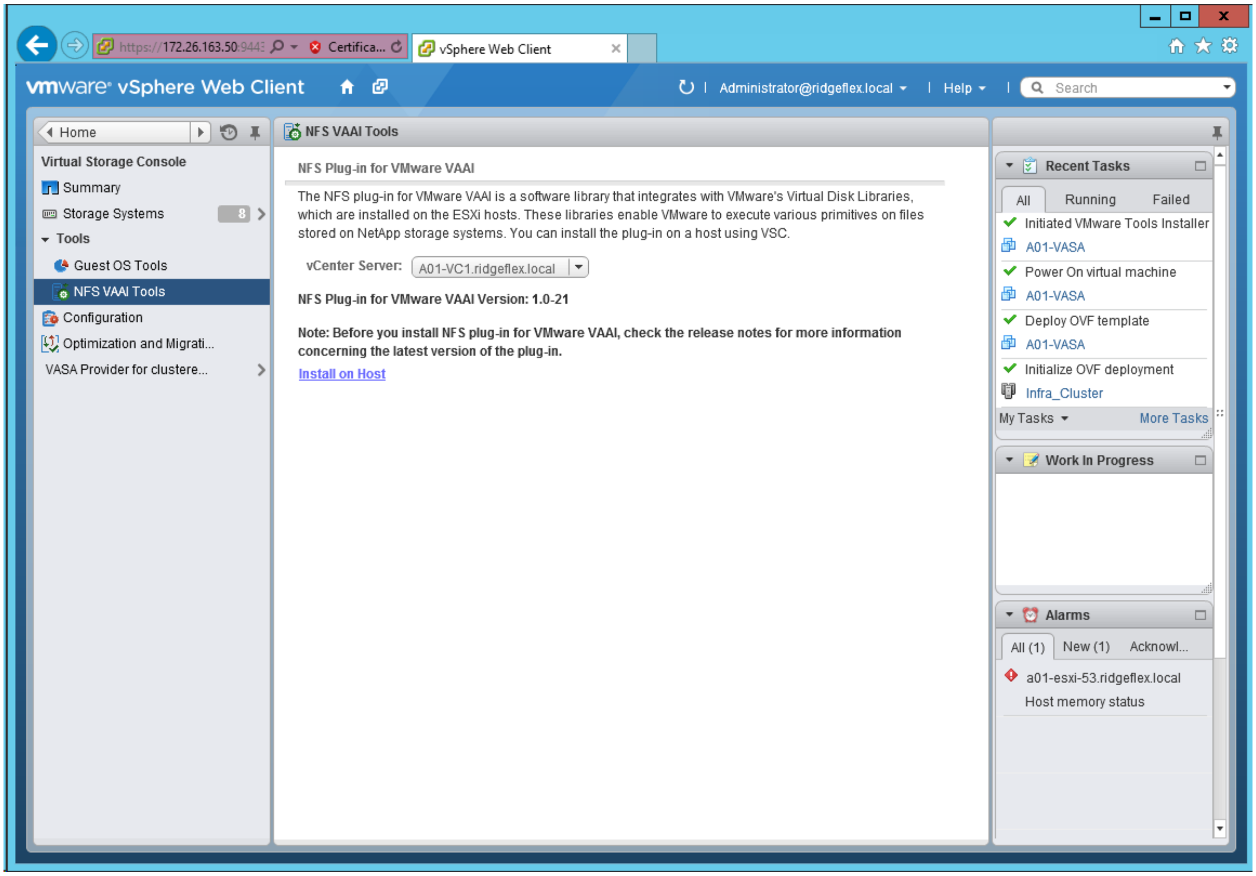

NFS Plug-in for VMware VAAI

The NFS plug-in for VMware VAAI is a software library that integrates with VMware's Virtual Disk Libraries, which are installed on the ESXi hosts. These libraries enable VMware to execute various primitives on files stored on NetApp storage systems. You can install the plug-in on a host using VSC.

vCenter Server: A01-VC1.ridgeflex.local

NFS Plug-in for VMware VAAI Version: 1.0-21

Note: Before you install NFS plug-in for VMware VAAI, check the release notes for more information concerning the latest version of the plug-in.

Install on Host

9. Click Install on Host. Select all Hosts on which you want to install the plug-in.

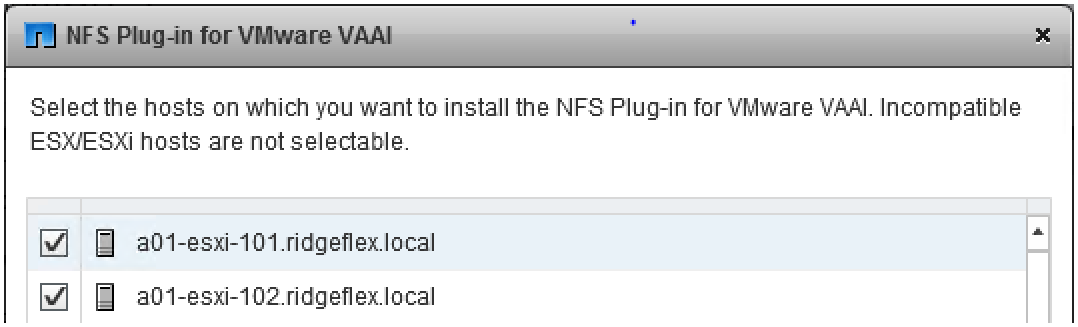

NFS Plug-in for VMware VAAI

Select the hosts on which you want to install the NFS Plug-in for VMware VAAI. Incompatible ESX/ESXi hosts are not selectable.

☑ a01-esxi-101.ridgeflex.local

☑ a01-esxi-102.ridgeflex.local

10. Click Install then click OK.

11. One at a time, put each ESXi host into Maintenance Mode, reboot the host, then Exit Maintenance Mode. It may be necessary to manually migrate VMs to the other host to allow the host to enter Maintenance Mode.

12. When the reboots have completed, in the vSphere Web Client from the Home page, click Storage, then select the infra_datastore_1 datastore. Select Settings under the Manage tab in the center pane. Hardware Acceleration should now show Supported on all hosts as shown below. All NFS datastores should now support Hardware Acceleration.

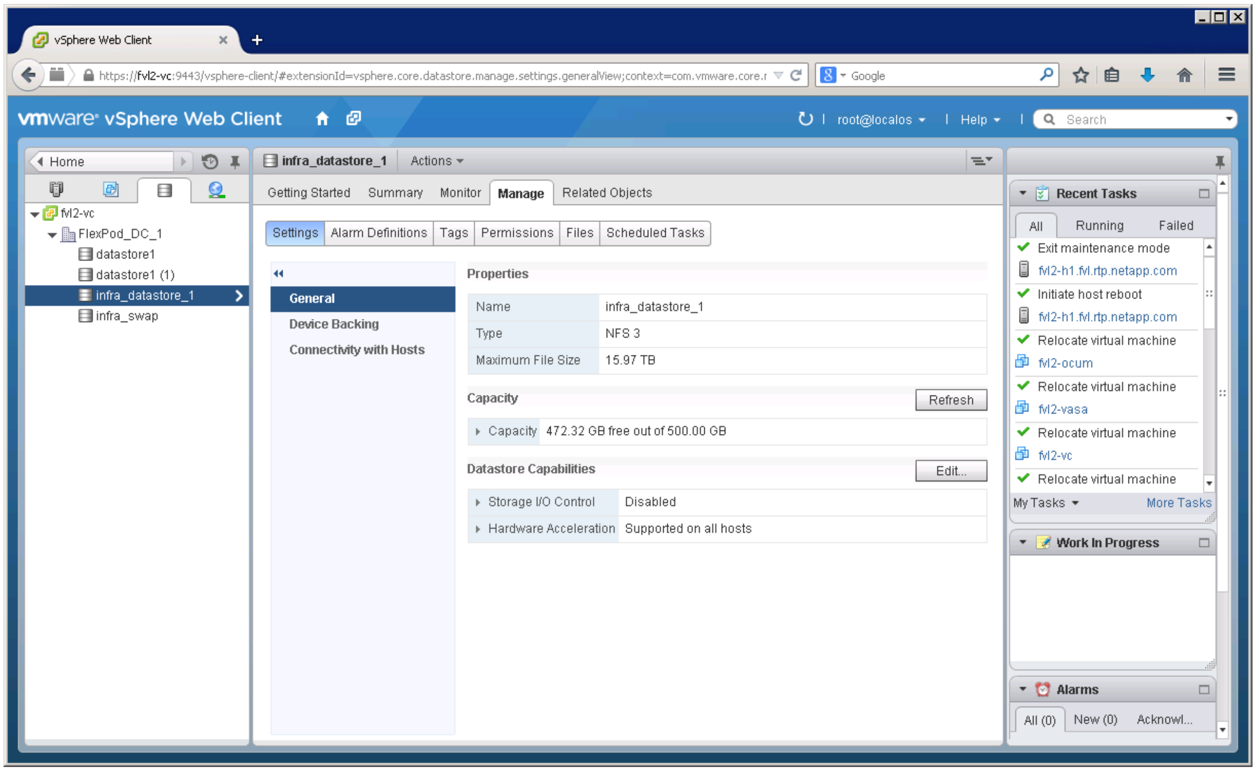

infra_datastore_1

Properties

Name	infra_datastore_1
Type	NFS 3
Maximum File Size	15.97 TB

Capacity

Capacity	472.32 GB free out of 500.00 GB

Datastore Capabilities

Storage I/O Control	Disabled
Hardware Acceleration	Supported on all hosts

Recent Tasks

Exit maintenance mode	fvl2-h1.fvl.rtp.netapp.com
Initiate host reboot	fvl2-h1.fvl.rtp.netapp.com
Relocate virtual machine	fvl2-ocum
Relocate virtual machine	fvl2-vasa
Relocate virtual machine	fvl2-vc
Relocate virtual machine	

Appendix A - Deploying Direct Connect FCoE

This section details the additional configuration to setup FCoE Boot using direct connection between AFF 8040 and UCS FI. The FCoE boot configuration can be used in addition to or instead of iSCSI Boot**.** In this appendix, FCoE boot is configured and enabled to coexist with the iSCSI boot option. This appendix assumes that the initial storage configuration section in this document has been completed. Customers can choose to replace the appropriate iSCSI specific sections of the main deployment guide with the FCoE sections in this appendix.

To support FCoE direct attached storage, the following additional physical connectivity is required:

Figure 7 FCoE Direct Attached Storage

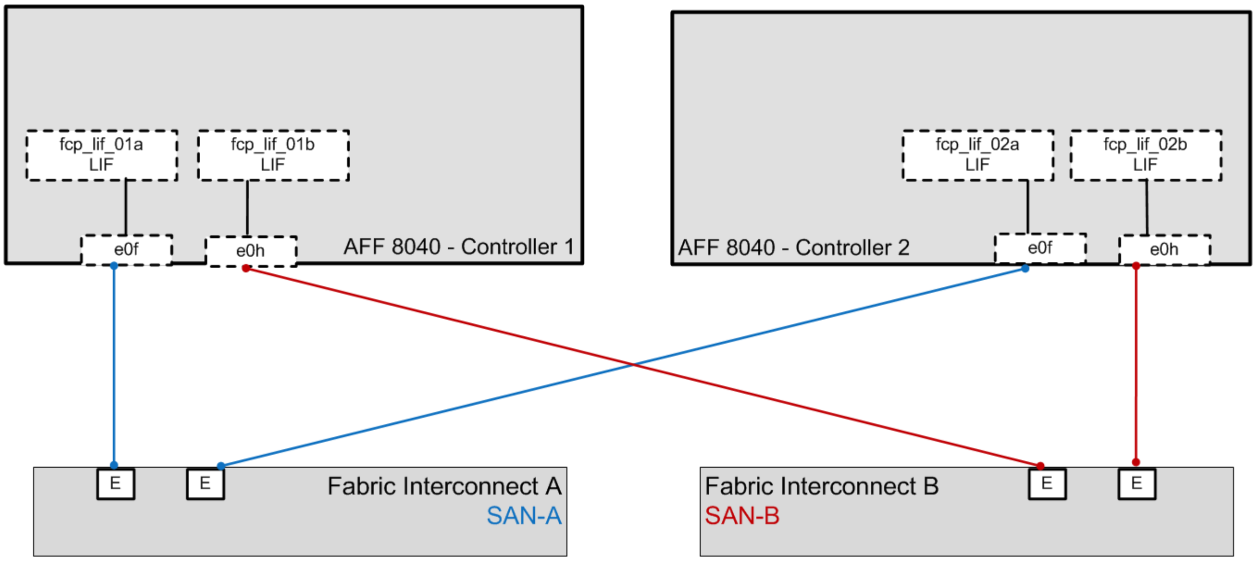

Storage Configuration

Add FCP as an Allowed Protocol in Infrastructure SVM

1. Add the FCP as an allowed protocol on the Infrastructure SVM.

```
vserver add-protocols -vserver Infra-SVM -protocols fcp
vserver show -vserver Infra-SVM -fields allowed-protocols
```

FCP Service in Clustered Data ONTAP

1. Create the FCP service on the Infrastructure SVM. This command also starts the FCP service and sets World Wide Node Name (WWNN) of the SVM.

```
fcp create -vserver Infra-SVM
fcp show
```

Create LUNs in Clustered Data ONTAP

1. Create two boot LUNS: `VM-Host-Infra-01` and `VM-Host-Infra-02`.

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-01 -size 10GB
-ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-02 -size 10GB
-ostype vmware -space-reserve disabled
```

Ensure FCoE Ports are Online

1. Check the status of the FCoE adapter ports used for Direct Connect.

```
fcp adapter show -fields state
```

2. If ports 0f and 0h on each node are not in the up state, use the following command to bring them up.

```
fcp adapter modify –node <var_node> -adapter <adapter> -state up
fcp adapter show -fields state
```

FCoE LIF in Clustered Data ONTAP

1. Create FCoE logical interfaces (LIFs).

```
network interface create -vserver Infra-SVM -lif fcp_lif01a -role data -data-
protocol fcp -home-node <<var_node01>> -home-port 0f -status-admin up

network interface create -vserver Infra-SVM -lif fcp_lif01b -role data -data-
protocol fcp -home-node <<var_node01>> -home-port 0h -status-admin up


network interface create -vserver Infra-SVM -lif fcp_lif02a -role data -data-
protocol fcp -home-node <<var_node02>> -home-port 0f -status-admin up


network interface create -vserver Infra-SVM -lif fcp_lif02b -role data -data-
protocol fcp -home-node <<var_node02>> -home-port 0h -status-admin up

network interface show –vserver Infra-SVM
```

Server (UCS) Configuration

Place Cisco UCS Fabric Interconnects in Fiber Channel Switching Mode

To use FCoE Appliance Ports, the Cisco UCS Fabric Interconnects must be placed in Fiber Channel Switching Mode. Complete the following steps:

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

2. Expand Fabric Interconnects and select Fabric Interconnect B.

Note: This next step will reboot both UCS Fabric Interconnects. If any servers are running on this system, they should be shut down before this step is executed.

3. In the Actions pane, select Set FC Switching Mode. Click Yes. Click OK.

4. After the Fabric Interconnects have rebooted, log back into UCS Manager.

5. Expand Fabric Interconnects and select Fabric Interconnects.

6. For each Fabric Interconnect, verify under Status that the FC Mode is now Switch.

Enable FCoE Storage Ports

To enable FCoE ports, complete the following steps:

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

2. Select Equipment > Fabric Interconnects > Fabric Interconnect A (primary) > Fixed Module.

3. Expand Ethernet Ports.

4. Select ports 13 and 14 that are connected to the FCoE ports on the storage controllers, right-click them, and select Configure as FCoE Storage Port. Click Yes to confirm.

5. Select Equipment > Fabric Interconnects > Fabric Interconnect B (subordinate) > Fixed Module.

6. Expand Ethernet Ports.

7. Select ports 13 and 14 that are connected to the FCoE ports on the storage controllers, right-click them, and select Configure as FCoE Storage Port. Click Yes to confirm.

Create a WWNN Pool for FCoE Boot

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

1. Select the SAN tab on the left.

2. Select Pools > root.

3. Right-click WWNN Pools under the root organization.

4. Select Create WWNN Pool to create the WWNN pool.

5. Enter `WWNN_Pool` for the name of the WWNN pool.

6. Optional: Enter a description for the WWNN pool.

7. Select Sequential for Assignment Order.

8. Click Next.

9. Click Add.

10. Modify the From field as necessary for this UCS Environment.

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

12. Click OK.

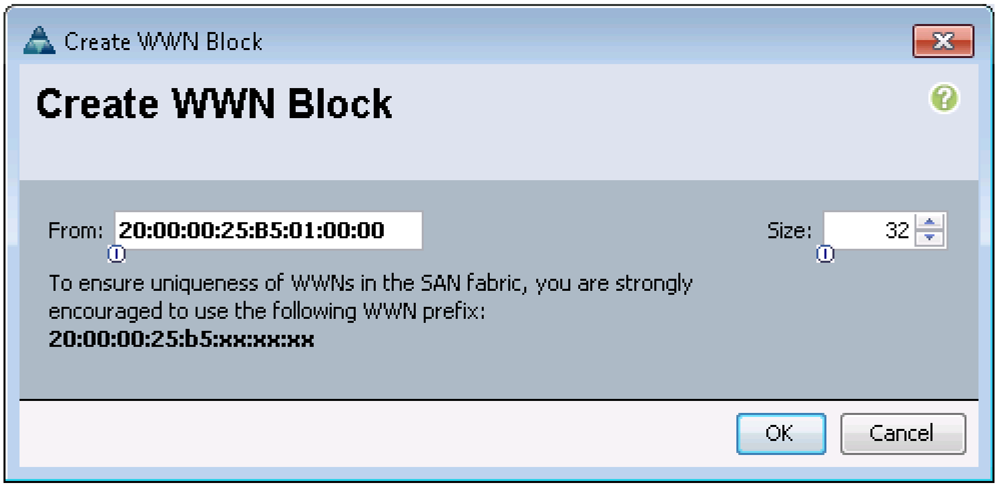

From: 20:00:00:25:B5:01:00:00

Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:
20:00:00:25:b5:xx:xx:xx

13. Click Finish.

14. In the message box that displays, click OK.

Create WWPN Pools

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

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

2. Select Pools > root.

Note: In this procedure, two WWPN pools are created, one for each switching fabric.

3. Right-click WWPN Pools under the root organization.

4. Select Create WWPN Pool to create the WWPN pool.

5. Enter `WWPN_Pool_A` as the name of the WWPN pool.

6. Optional: Enter a description for the WWPN pool.

7. Select Sequential for Assignment Order.

8. Click Next.

9. Click Add.

10. Specify a starting WWPN.

Note: For the FlexPod solution, the recommendation is to place `0A` in the next-to-last octet of the starting WWPN to identify all of the WWPNs as fabric A addresses.

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

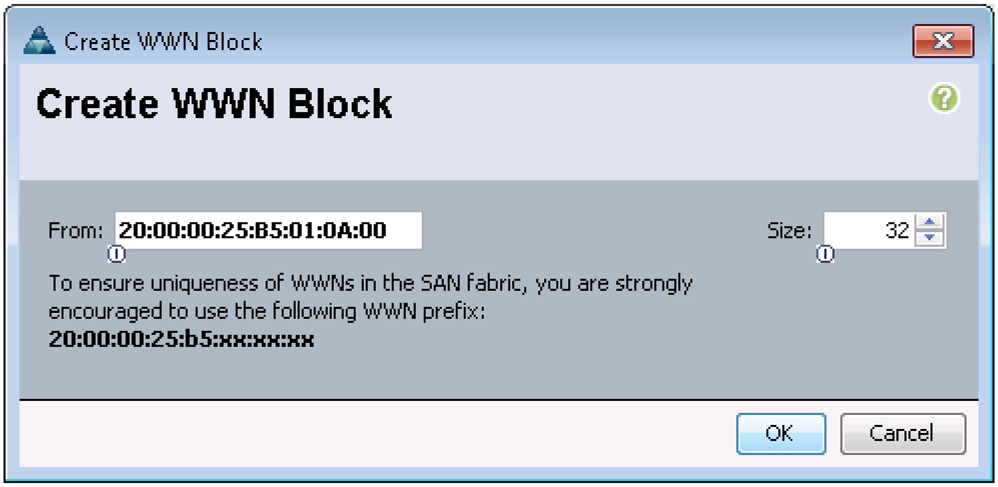

Create WWN Block

From: 20:00:00:25:B5:01:0A:00

Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:
20:00:00:25:b5:xx:xx:xx

12. Click OK.

13. Click Finish.

14. In the confirmation message, click OK.

15. Right-click WWPN Pools under the root organization.

16. Select Create WWPN Pool to create the WWPN pool.

17. Enter `WWPN_Pool_B` as the name of the WWPN pool.

18. Optional: Enter a description for the WWPN pool.

19. Select Sequential for Assignment Order.

20. Click Next.

21. Click Add.

22. Specify a starting WWPN.

Note: For the FlexPod solution, it is recommended to place `0B` in the next to last octet of the starting WWPN to identify all the WWPNs in this pool as fabric B addresses.

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

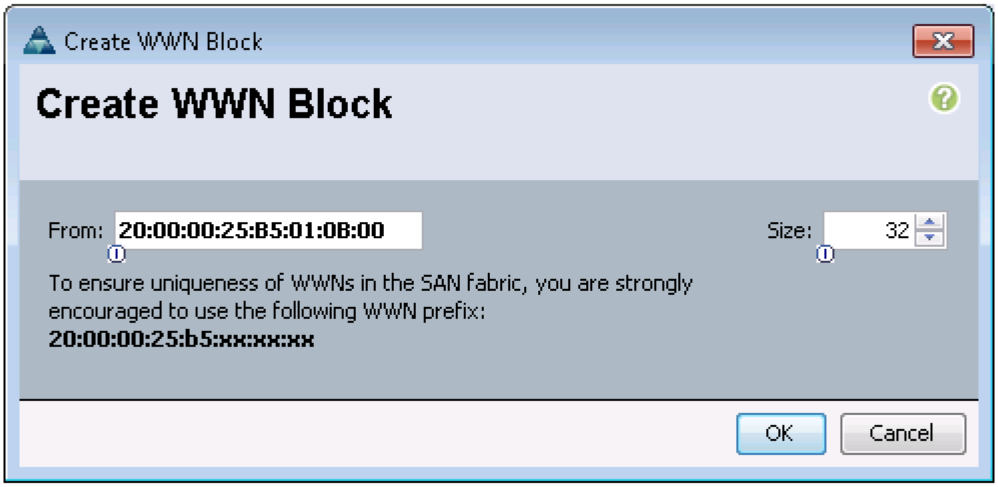

Create WWN Block

From: 20:00:00:25:B5:01:0B:00

Size: 32

To ensure uniqueness of WWNs in the SAN fabric, you are strongly encouraged to use the following WWN prefix:
20:00:00:25:b5:xx:xx:xx

24. Click OK.

25. Click Finish.

26. In the confirmation message, click OK.

Create Server Pool

To configure specific server pool for servers supporting FCoE boot, complete the following steps:

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

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Pools > root.
3. Right-click Server Pools.
4. Select Create Server Pool.
5. Enter `Infra_FCoE_Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool.
7. Click Next.
8. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the `Infra_FCoE_Pool` server pool.
9. Click Finish.
10. Click OK.

Create VSANs

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

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

Note: In this procedure, two VSANs are created.

2. Select SAN > SAN Cloud.
3. Right-click VSANs.
4. Select Create VSAN.
5. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A
6. Select Enabled for FC Zoning.
7. Select Fabric A.
8. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.
9. Click OK, and then click OK again.

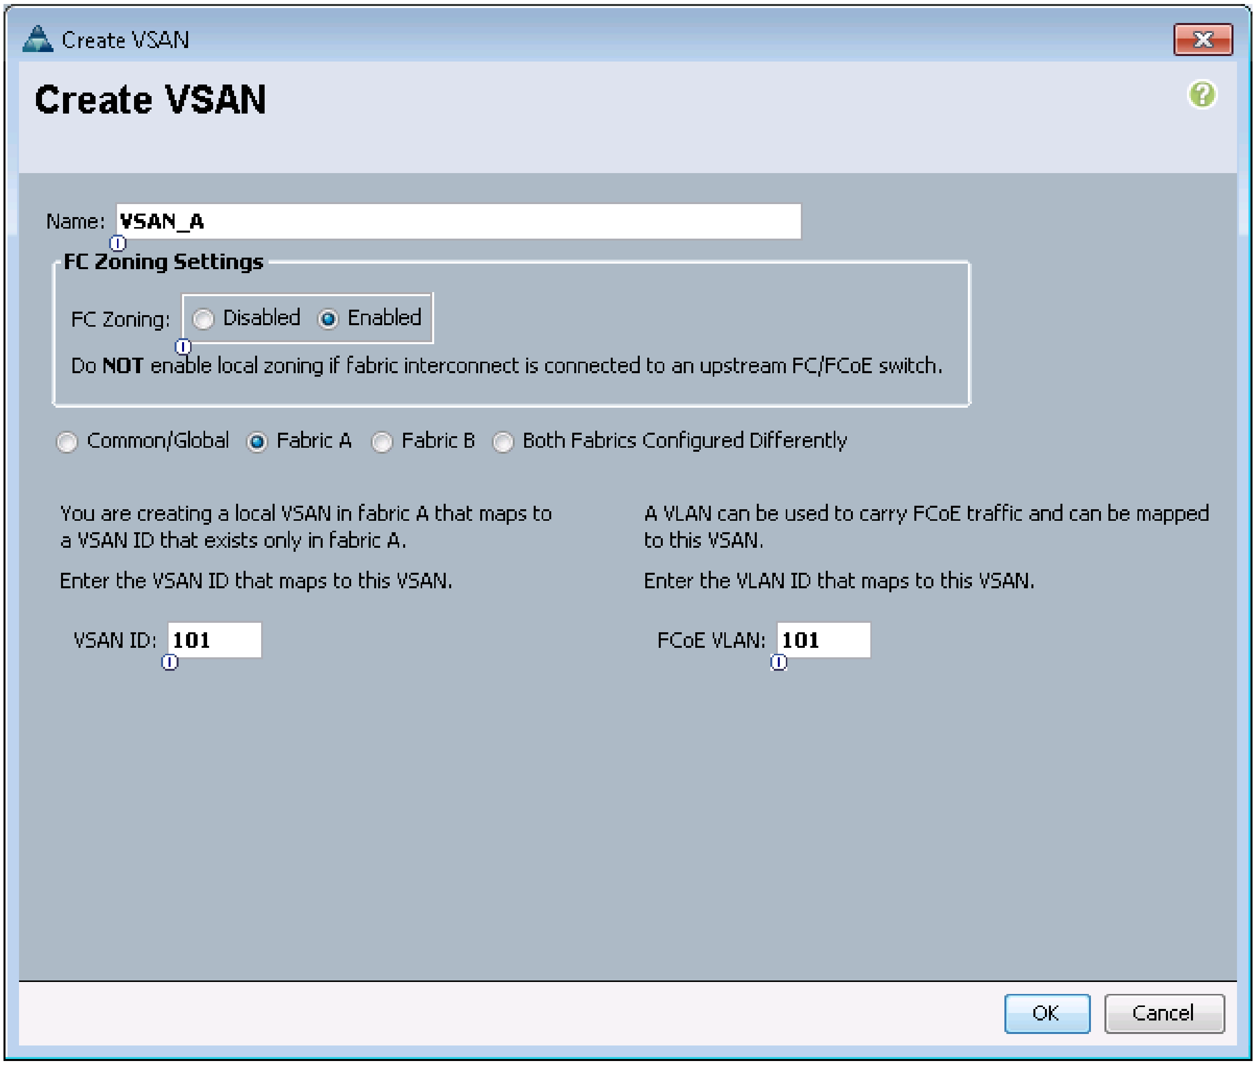

Create VSAN

Name: VSAN_A

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do NOT enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ● Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 101

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 101

10. Under SAN Cloud, right-click VSANs.

11. Select Create VSAN.

12. Enter `VSAN_B` as the name of the VSAN to be used for Fabric B.

13. Select Enabled for FC Zoning.

14. Select Fabric B.

15. Enter a unique VSAN ID and a corresponding FCoE VLAN ID. It is recommended use the same ID for both parameters and to use something other than 1.

16. Click OK, and then click OK again.

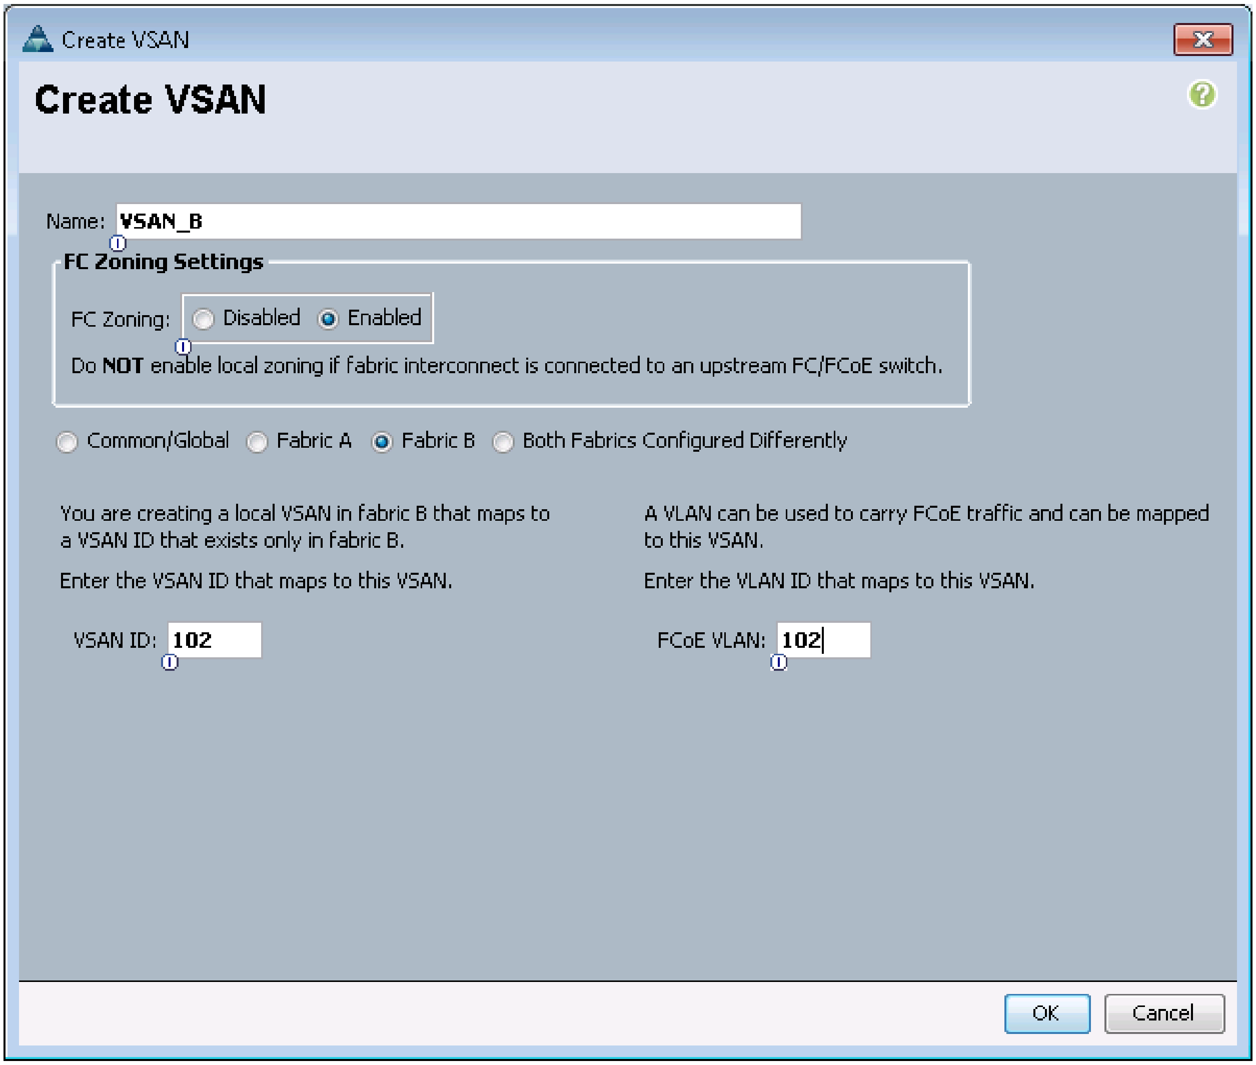

Create VSAN

Name: VSAN_B

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do NOT enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ○ Fabric A ● Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 102

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 102

17. Under Storage Cloud, right-click VSANs.

18. Select Create Storage VSAN.

19. Enter `VSAN_A` as the name of the VSAN to be used for Fabric A.

20. Select Enabled for FC Zoning.

21. Select Fabric A.

22. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric A above.

23. Click OK, and then click OK again.

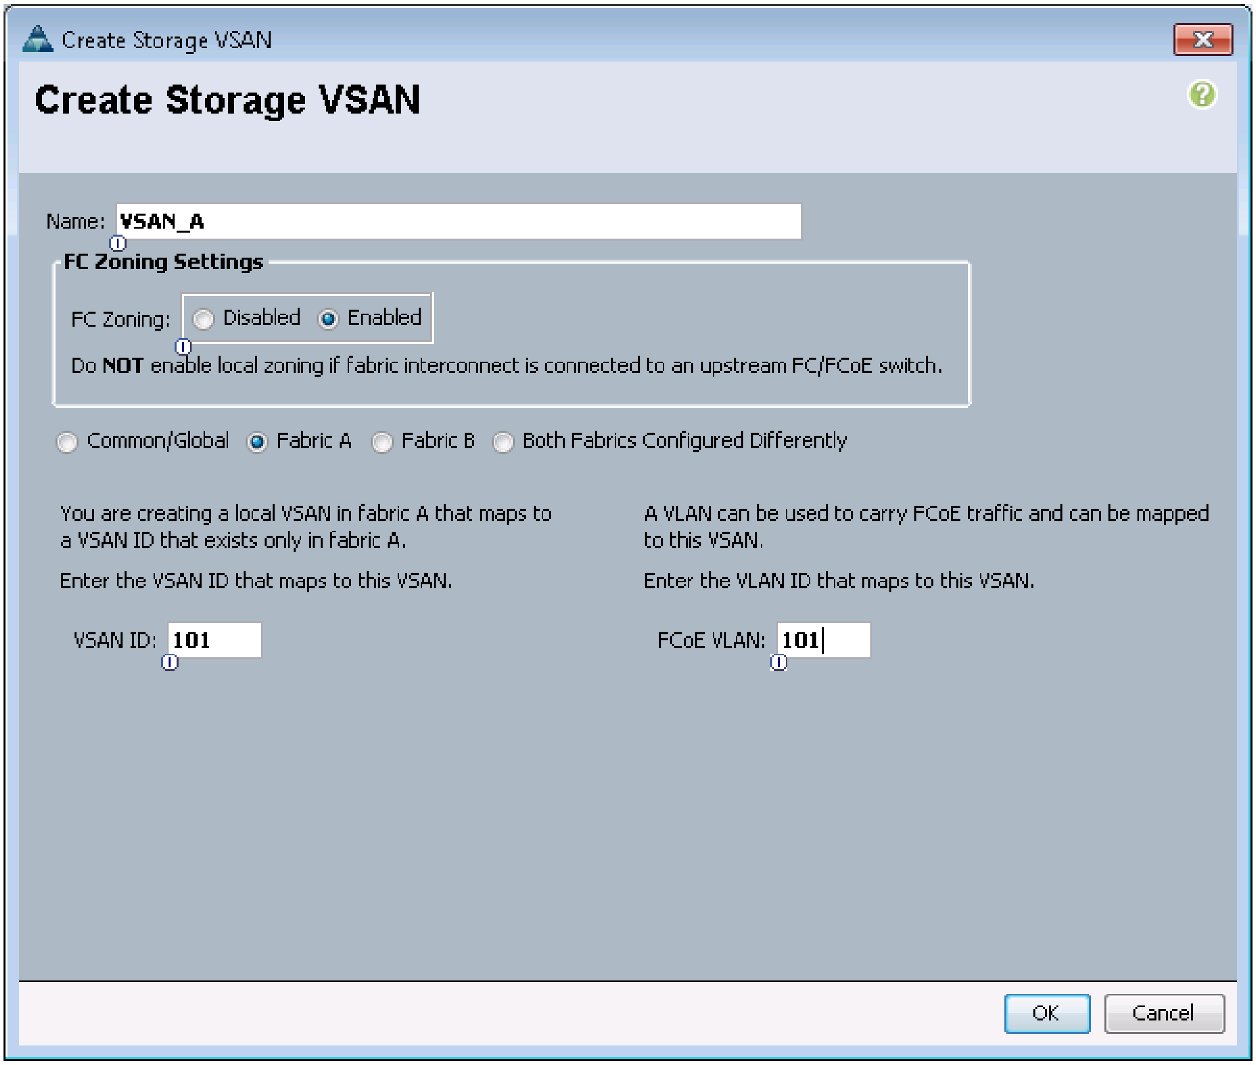

Create Storage VSAN

Name: VSAN_A

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ● Fabric A ○ Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric A that maps to a VSAN ID that exists only in fabric A.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 101

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 101

24. Under Storage Cloud, right-click VSANs.

25. Select Create Storage VSAN.

26. Enter `VSAN_B` as the name of the VSAN to be used for Fabric B.

27. Select Enabled for FC Zoning.

28. Select Fabric B.

29. Enter the same unique VSAN ID and corresponding FCoE VLAN ID that you entered for Fabric B above.

30. Click OK, and then click OK again.

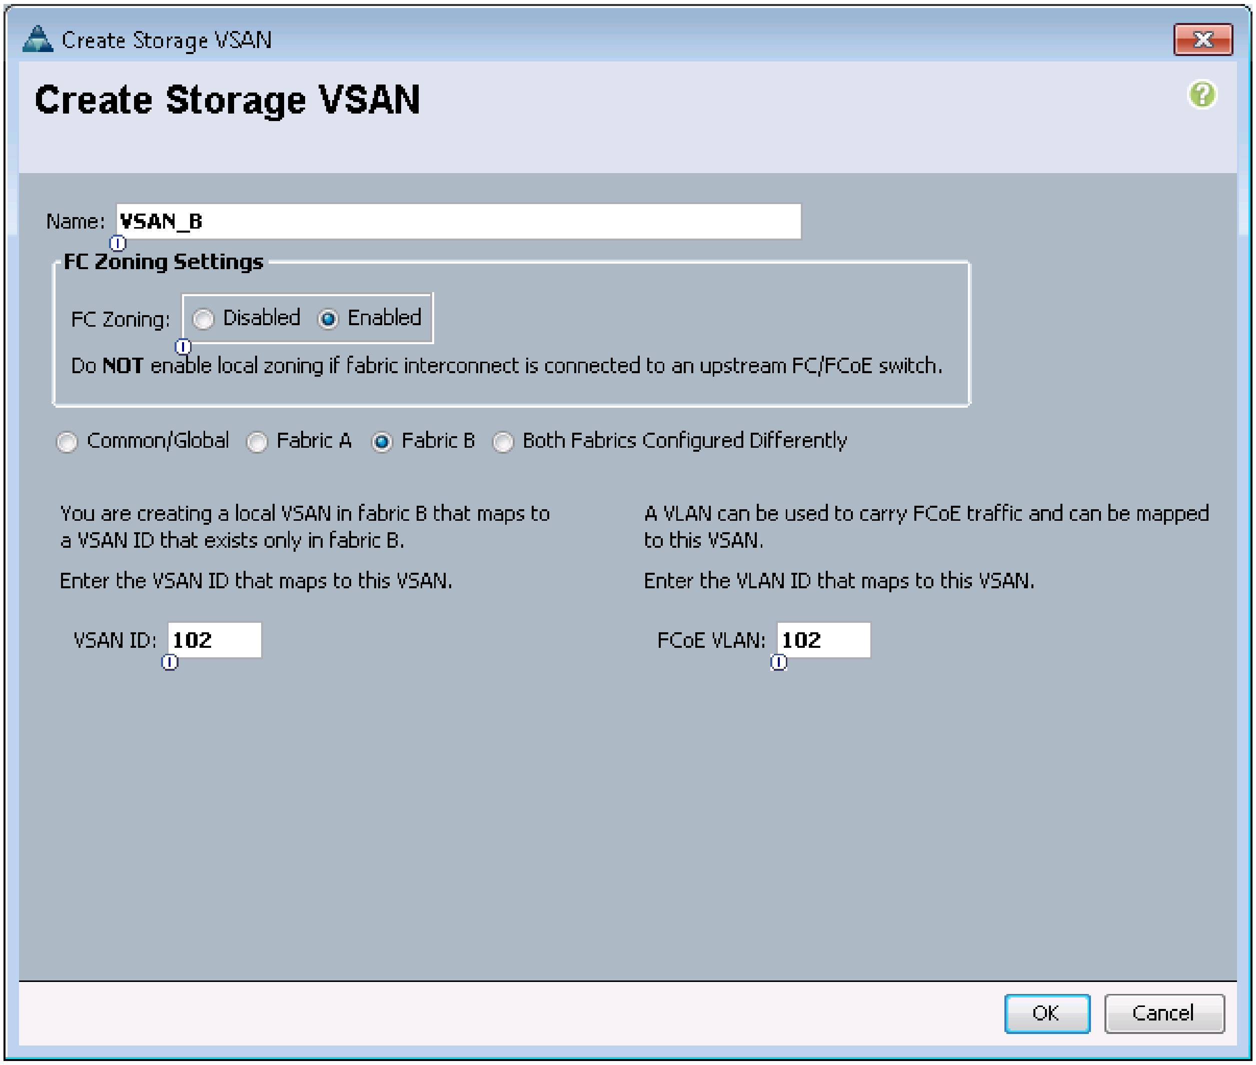

Create Storage VSAN

Name: VSAN_B

FC Zoning Settings

FC Zoning: ○ Disabled ● Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

○ Common/Global ○ Fabric A ● Fabric B ○ Both Fabrics Configured Differently

You are creating a local VSAN in fabric B that maps to a VSAN ID that exists only in fabric B.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: 102

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

Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: 102

Assign VSANs to FCoE Storage Ports

To assign the necessary virtual storage area networks (VSANs) to the FCoE Storage Ports for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select SAN > Storage Cloud.
3. Expand Fabric A and Storage FCoE Interfaces.
4. Right-click FCoE Interface 1/13 and select Storage FCoE Interface.
5. Set the User Label to the storage controller name and port that this interface is connected to.
6. Select `VSAN_A` as the VSAN.
7. Click OK.

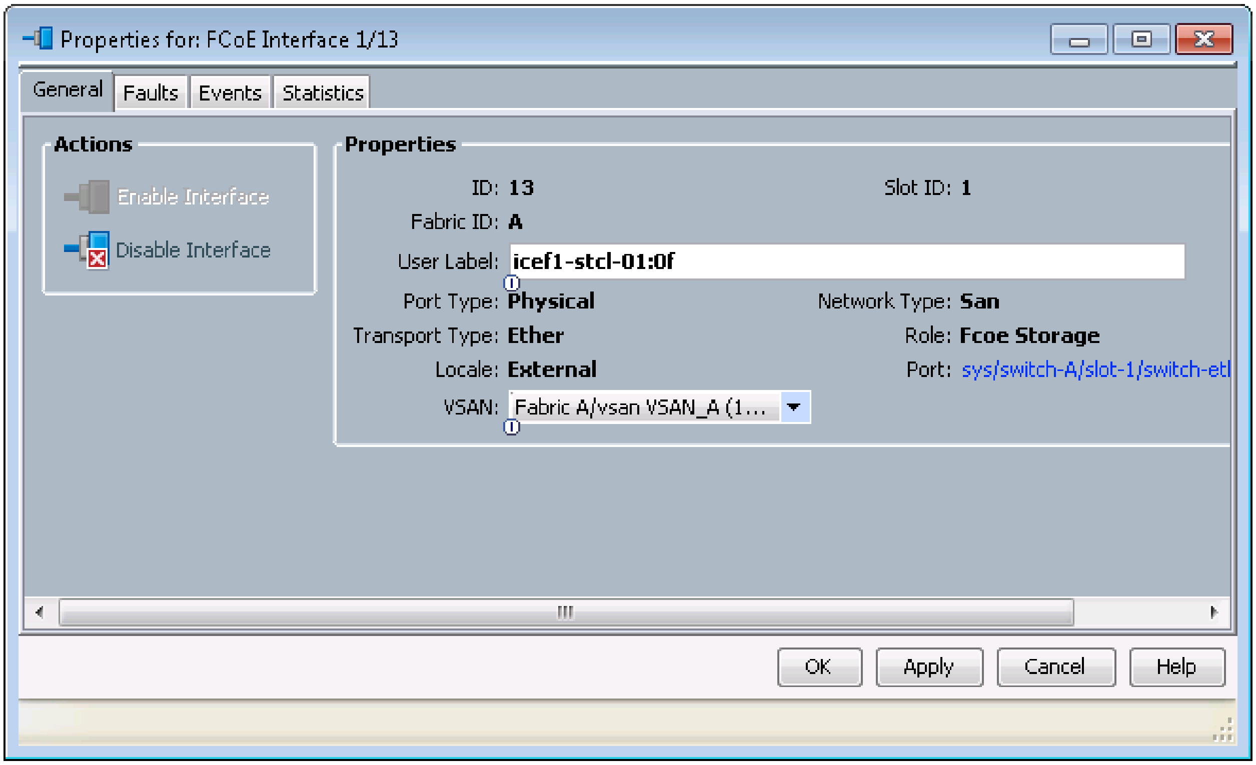

Properties for: FCoE Interface 1/13

General | Faults | Events | Statistics

Actions

Enable Interface

Disable Interface

Properties

ID: **13** Slot ID: **1**

Fabric ID: **A**

User Label: **icef1-stcl-01:0f**

Port Type: **Physical** Network Type: **San**

Transport Type: **Ether** Role: **Fcoe Storage**

Locale: **External** Port: sys/switch-A/slot-1/switch-etl

VSAN: Fabric A/vsan VSAN_A (1...

8. Expand Fabric A and Storage FCoE Interfaces.

9. Right-click FCoE Interface 1/14 and select Storage FCoE.

10. Set the User Label to the storage controller name and port that this interface is connected to.

11. Select `VSAN_A` as the VSAN.

12. Click OK.

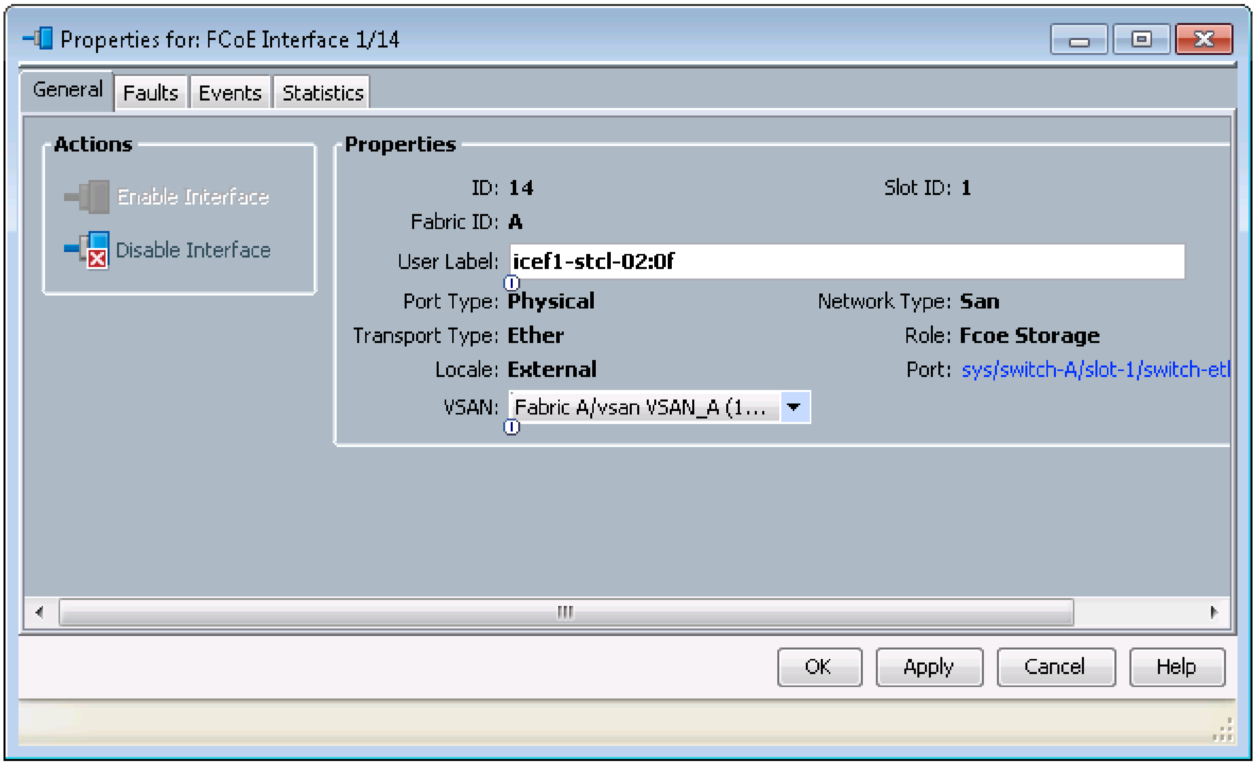

ID: **14** Slot ID: **1**

Fabric ID: **A**

User Label: icef1-stcl-02:0f

Port Type: **Physical** Network Type: **San**

Transport Type: **Ether** Role: **Fcoe Storage**

Locale: **External** Port: sys/switch-A/slot-1/switch-etl

VSAN: Fabric A/vsan VSAN_A (1...

13. Expand Fabric B and Storage FCoE Interfaces.

14. Right-click FCoE Interface 1/13 and select Storage FCoE.

15. Set the User Label to the storage controller name and port that this interface is connected to.

16. Select `VSAN_B` as the VSAN.

17. Click OK.

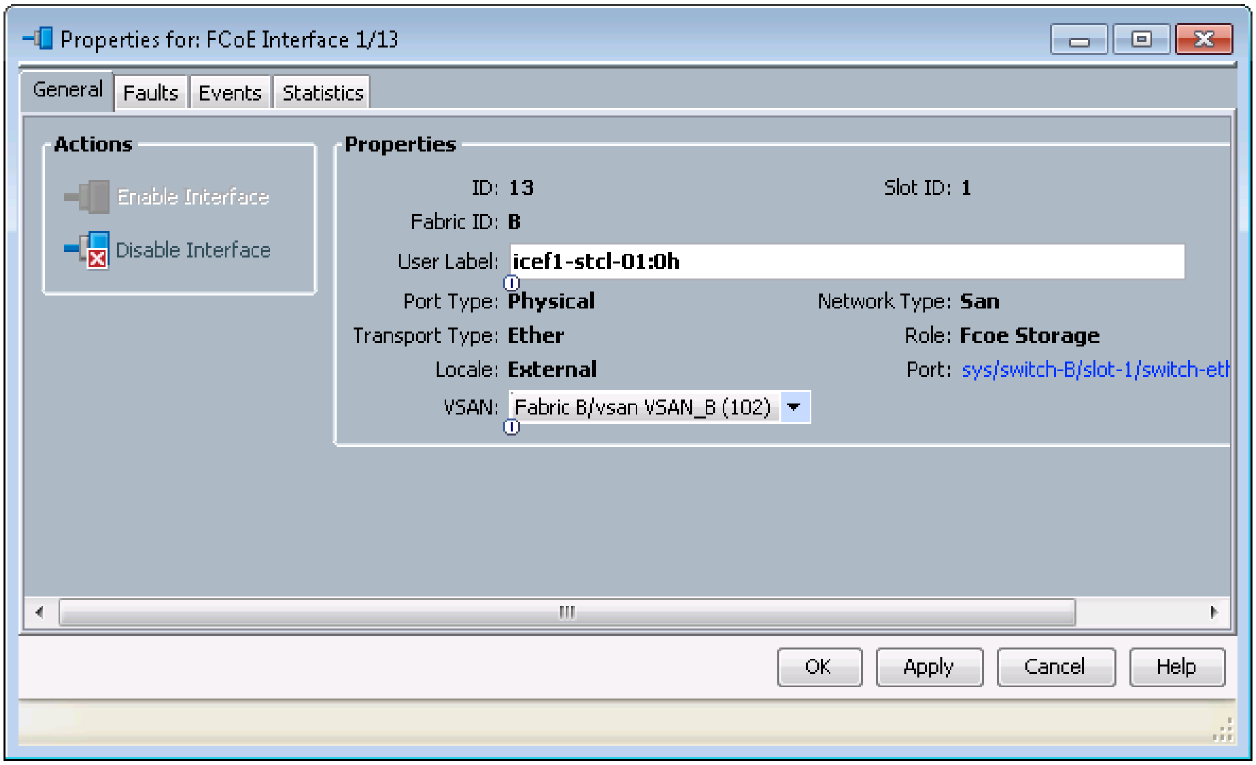

Properties for: FCoE Interface 1/13

Properties

ID: 13

Slot ID: 1

Fabric ID: B

User Label: icef1-stcl-01:0h

Port Type: Physical

Network Type: San

Transport Type: Ether

Role: Fcoe Storage

Locale: External

Port: sys/switch-B/slot-1/switch-eth

VSAN: Fabric B/vsan VSAN_B (102)

18. Expand Fabric B and Storage FCoE Interfaces.

19. Right-click FCoE Interface 1/14 and select Storage FCoE Interface.

20. Set the User Label to the storage controller name and port that this interface is connected to.

21. Select `VSAN_B` as the VSAN.

22. Click OK.

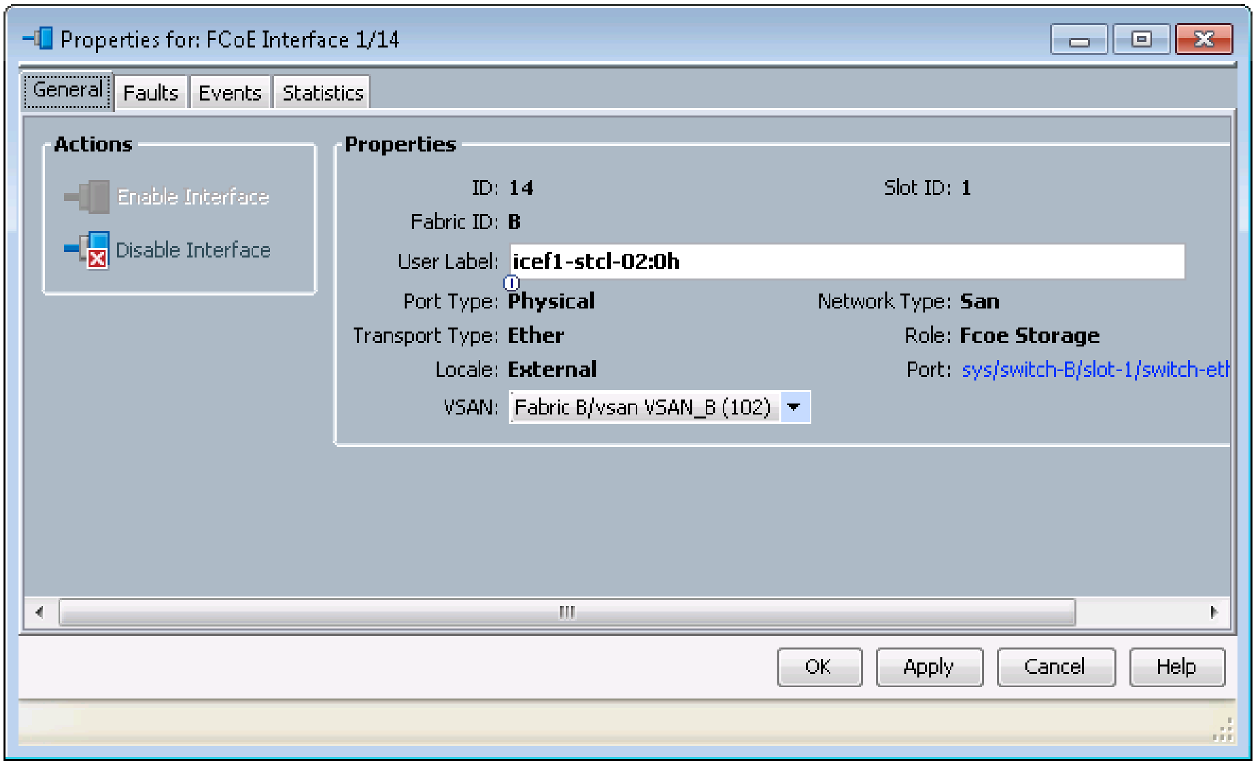

Properties for: FCoE Interface 1/14

ID: 14

Slot ID: 1

Fabric ID: B

User Label: icef1-stcl-02:0h

Port Type: Physical

Network Type: San

Transport Type: Ether

Role: Fcoe Storage

Locale: External

Port: sys/switch-B/slot-1/switch-eth

VSAN: Fabric B/vsan VSAN_B (102)

Create Storage Connection Policies for FCoE Zoning

To create Storage Connection Policies for the FCoE Zoning, complete the following steps:

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click Storage Connection Policies.
4. Select Create Storage Connection Policy.
5. Enter `Infra-Fabric-A` as the name of the policy.
6. Select the Single Initiator Multiple Targets Zoning Type.
7. Click the Plus Sign on the right to add a zoning target.
8. Enter the WWPN for fcp_lif01a from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.
9. Select Path A and VSAN_A.

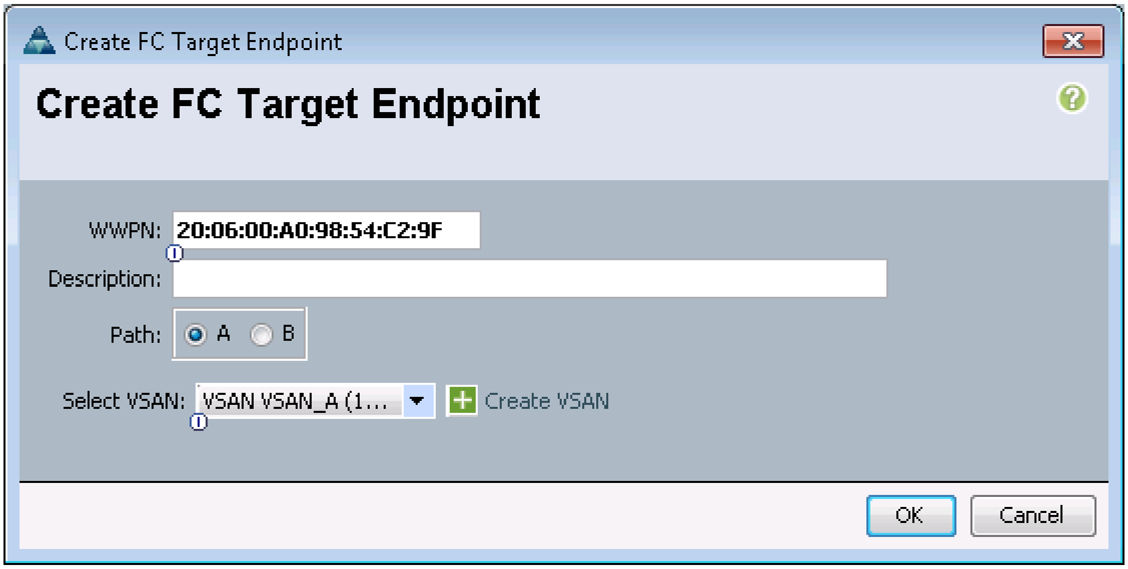
Create FC Target Endpoint

WWPN: 20:06:00:A0:98:54:C2:9F

Description:

Path: ● A ○ B

Select VSAN: VSAN VSAN_A (1...

10. Click OK.

11. Click the Plus Sign on the right to add a second zoning target.

12. Enter the WWPN for fcp_lif02a from the storage cluster. This WWPN can be obtained by logging into the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.

13. Select Path A and VSAN_A.

14. Click OK, and then click OK again.

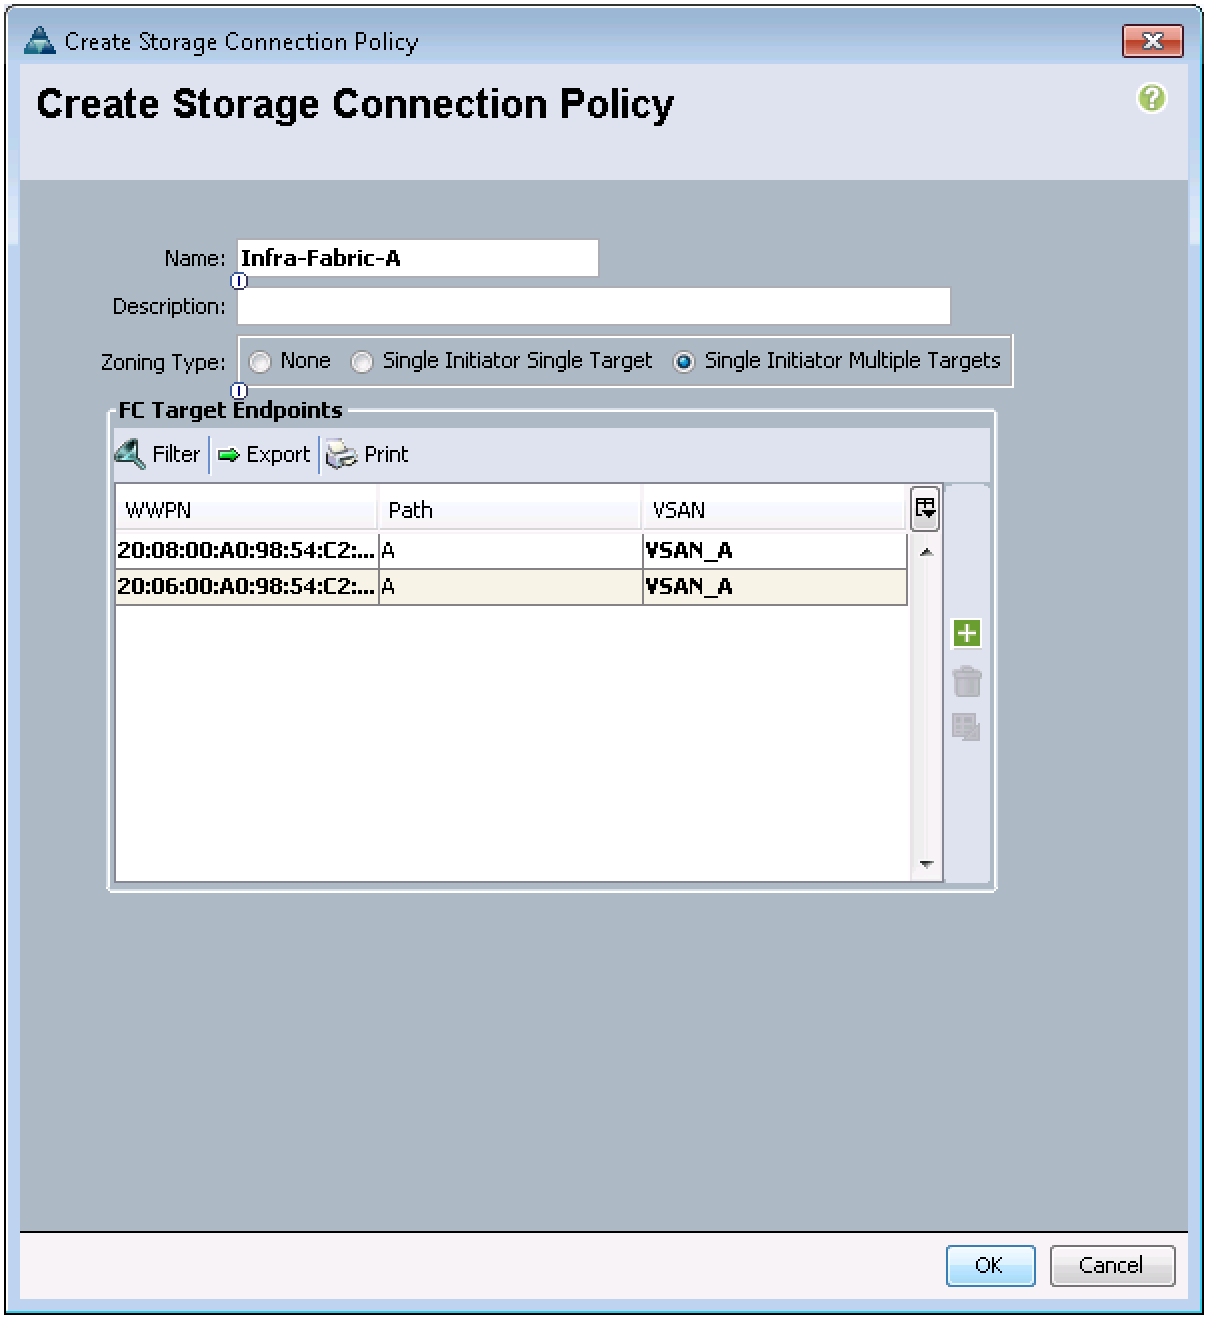

Create Storage Connection Policy

Name: Infra-Fabric-A

Description:

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN
20:08:00:A0:98:54:C2:...	A	VSAN_A
20:06:00:A0:98:54:C2:...	A	VSAN_A

15. Right-click Storage Connection Policies.

16. Select Create Storage Connection Policy

17. Enter `Infra-Fabric-B` as the name of the policy.

18. Select the Single Initiator Multiple Targets Zoning Type.

19. Click the Plus Sign on the right to add a zoning target.

20. Enter the WWPN for fcp_lif01b from the storage cluster. This WWPN can be obtained by logging in-to the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.

21. Select Path B and VSAN_B.

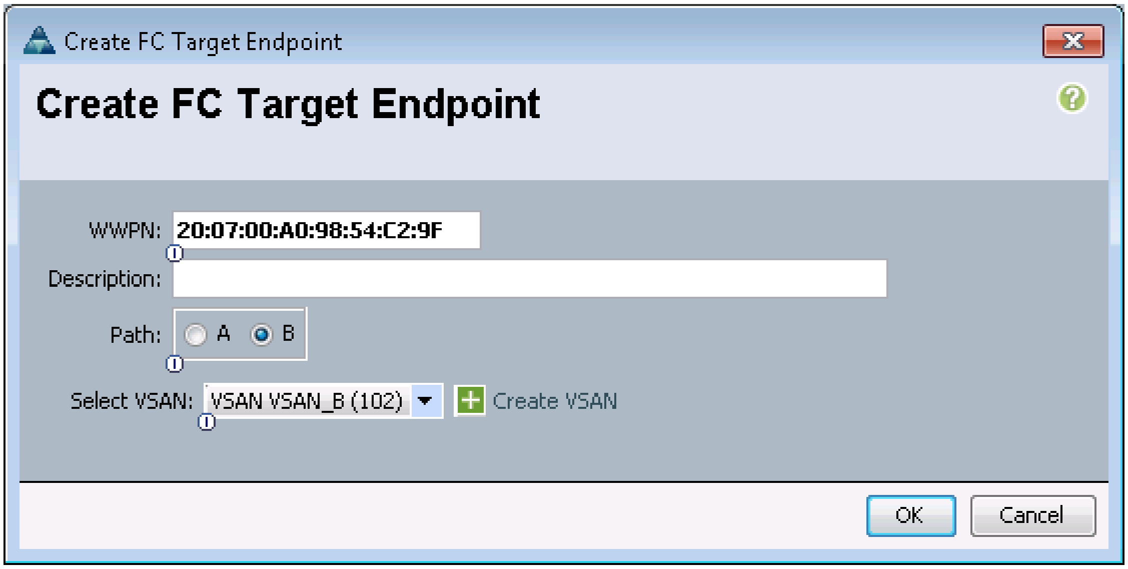

Create FC Target Endpoint

WWPN: 20:07:00:A0:98:54:C2:9F

Description:

Path: ○ A ● B

Select VSAN: VSAN VSAN_B (102)

22. Click OK.

23. Click the Plus Sign on the right to add a second zoning target.

24. Enter the WWPN for fcp_lif02b from the storage cluster. This WWPN can be obtained by logging in-to the storage cluster CLI and entering the `network interface show -vserver Infra-SVM` command.

25. Select Path B and VSAN_B.

26. Click OK, and then click OK again.

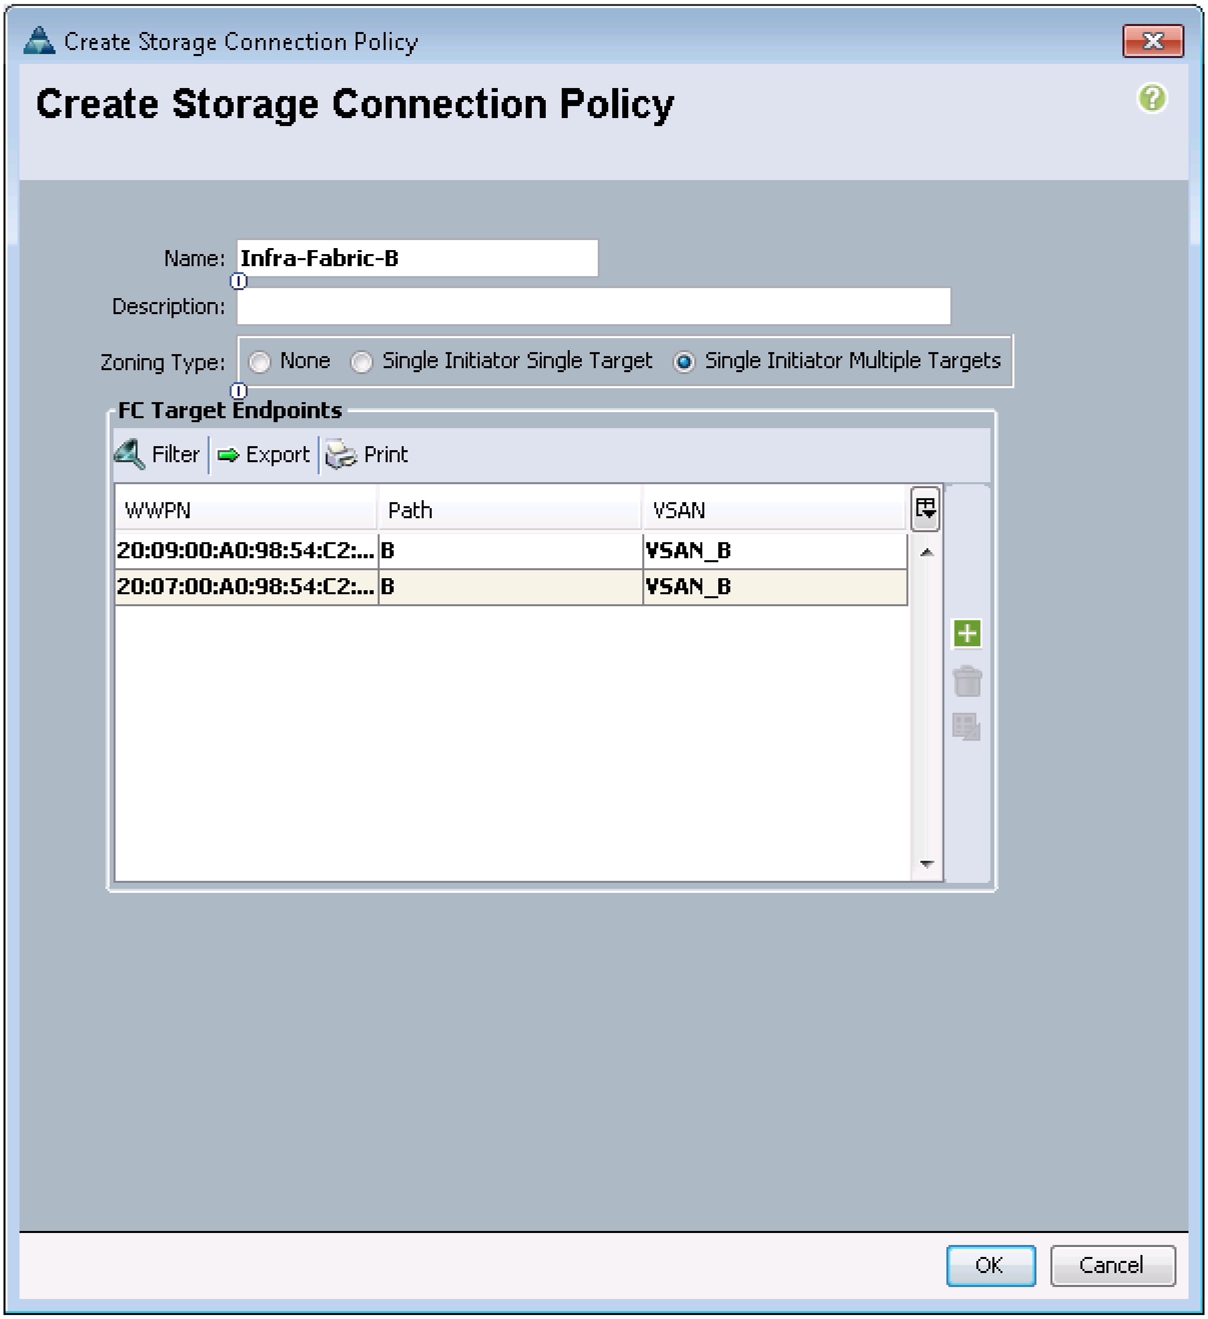

Create Storage Connection Policy

Name: Infra-Fabric-B

Description:

Zoning Type: ○ None ○ Single Initiator Single Target ● Single Initiator Multiple Targets

FC Target Endpoints

WWPN	Path	VSAN
20:09:00:A0:98:54:C2:...	B	VSAN_B
20:07:00:A0:98:54:C2:...	B	VSAN_B

Create vHBA Templates

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

1. In Cisco UCS Manager, click the SAN tab in the navigation pane.
2. Select Policies > root.
3. Right-click vHBA Templates.
4. Select Create vHBA Template.

5. Enter `vHBA_Template_A` as the vHBA template name.

6. Keep Fabric A selected.

7. Select VSAN_A.

8. Leave Initial Template as the Template Type.

9. Select `WWPN_Pool_A` as the WWPN Pool.

10. Click OK to create the vHBA template.

11. Click OK.

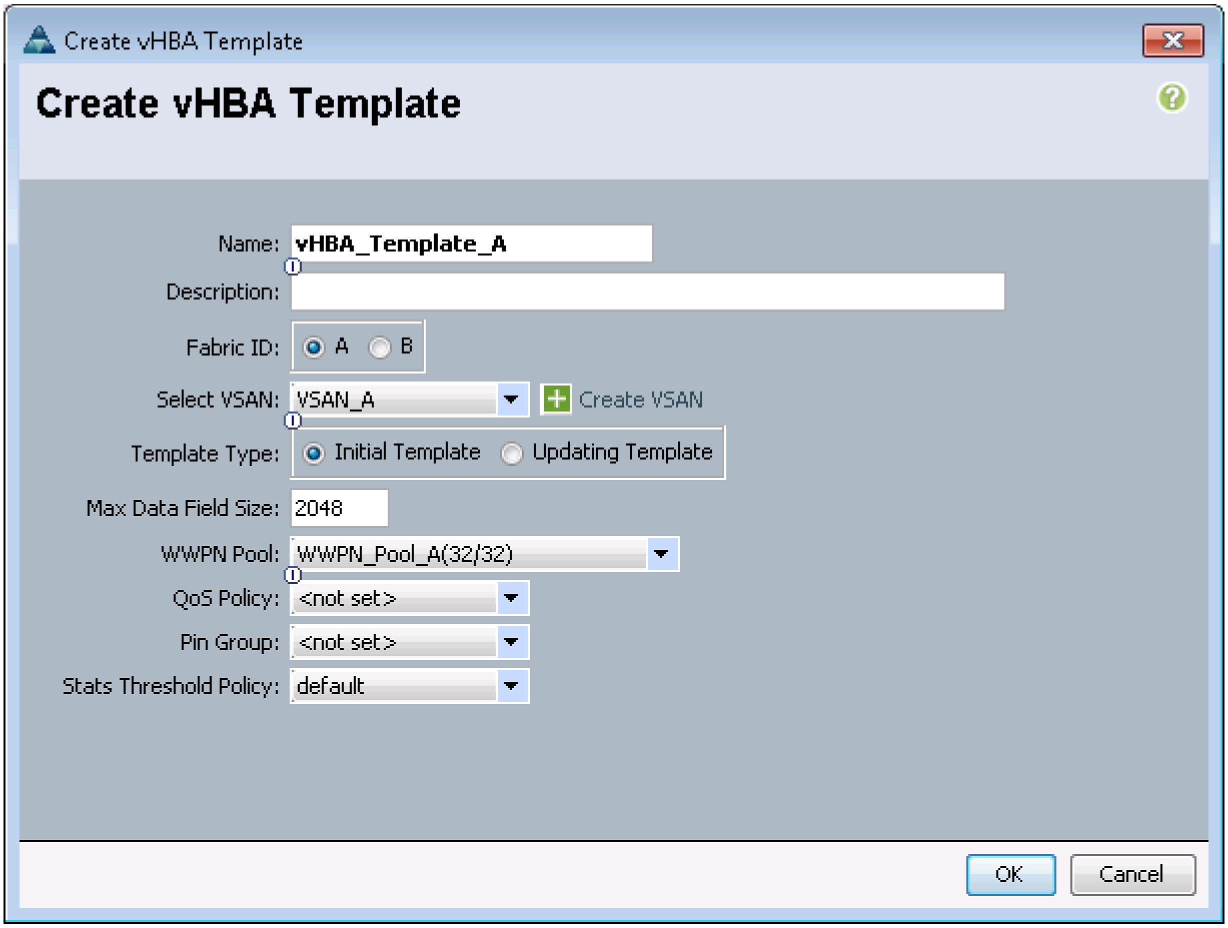

Create vHBA Template

Create vHBA Template

Name: vHBA_Template_A

Description:

Fabric ID: ● A ○ B

Select VSAN: VSAN_A

Template Type: ● Initial Template ○ Updating Template

Max Data Field Size: 2048

WWPN Pool: WWPN_Pool_A(32/32)

QoS Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

12. Right-click vHBA Templates.

13. Select Create vHBA Template.

14. Enter `vHBA_Template_B` as the vHBA template name.

15. Select Fabric B as the Fabric ID.

16. Select VSAN_B.

17. Leave Initial Template as the Template Type.

18. Select `WWPN_Pool_B` as the WWPN Pool.

19. Click OK to create the vHBA template.

20. Click OK.

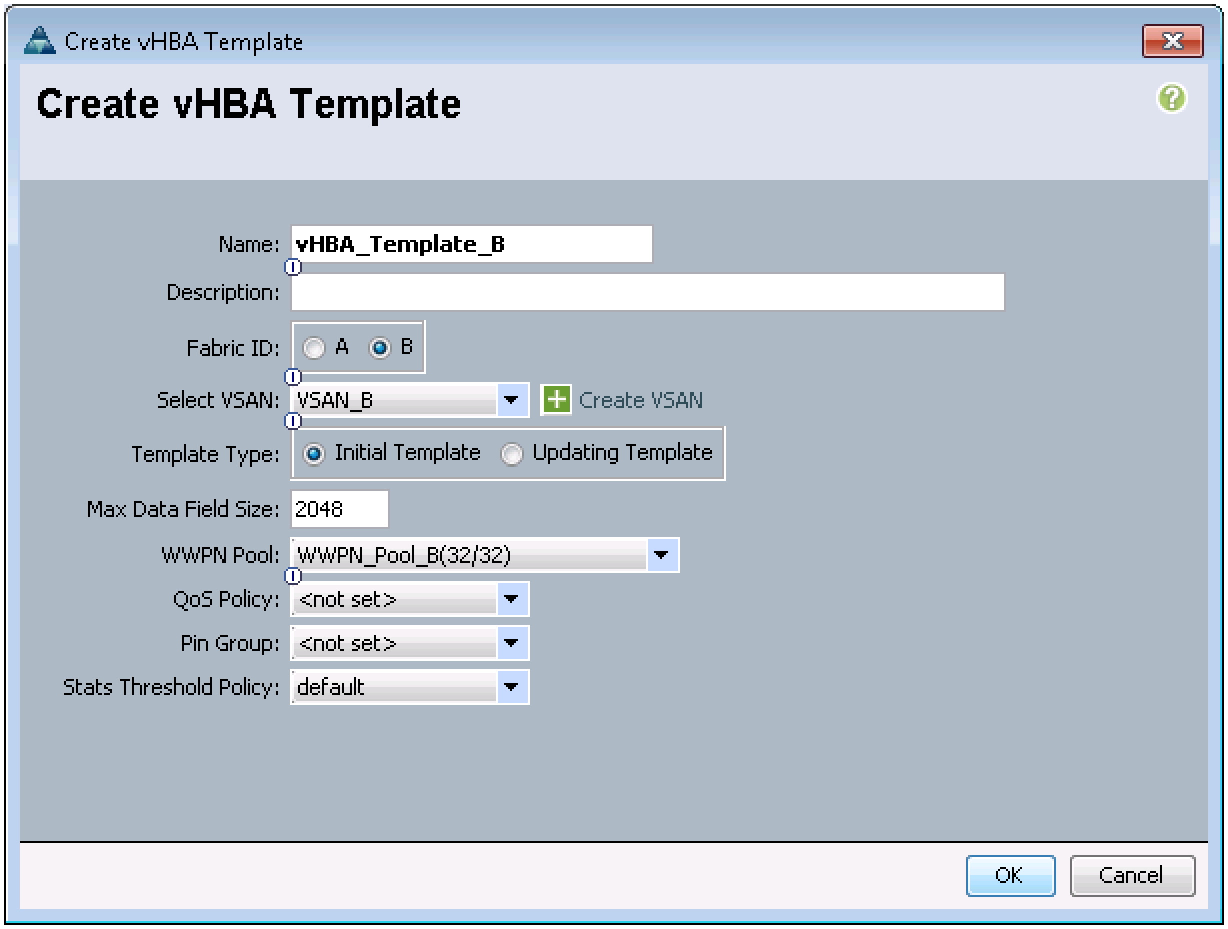

Create vHBA Template

Name: vHBA_Template_B

Description:

Fabric ID: ○ A ● B

Select VSAN: VSAN_B

Template Type: ● Initial Template ○ Updating Template

Max Data Field Size: 2048

WWPN Pool: WWPN_Pool_B(32/32)

QoS Policy: <not set>

Pin Group: <not set>

Stats Threshold Policy: default

Create Boot Policy

This procedure applies to a Cisco UCS environment in which two FCoE logical interfaces (LIFs) are on cluster node 1 (`fcp_lif01a` and `fcp_lif01b`) and two FCoE LIFs are on cluster node 2 (`fcp_lif02a` and `fcp_lif02b`). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS 6248UP A) and the B LIFs are connected to Fabric B (Cisco UCS 6248UP B).

A single boot policy is configured in this procedure. This policy configures the primary target to be fcp_lif01a.

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

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

2. Select Policies > root.

3. Right-click Boot Policies.

4. Select Create Boot Policy.

5. Enter `Boot-Fabric-A` as the name for the boot policy.

6. Optional: Enter a description for the boot policy.

Note: Do not select the Reboot on Boot Order Change checkbox.

7. Expand the Local Devices drop-down menu, select Add Remote CD/DVD.

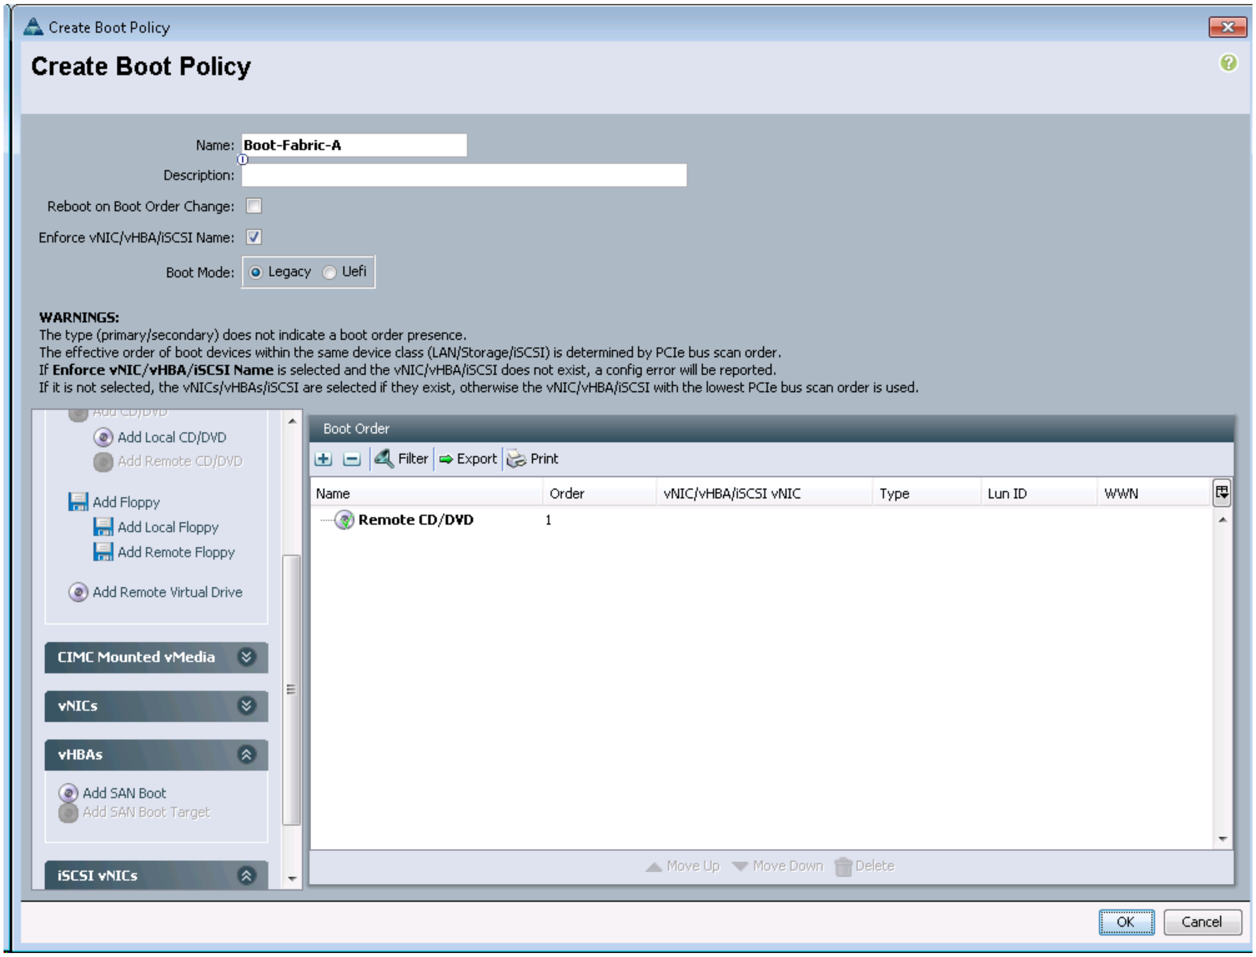

Create Boot Policy

Name: Boot-Fabric-A

Description:

Reboot on Boot Order Change: ☐

Enforce vNIC/vHBA/iSCSI Name: ☑

Boot Mode: ● Legacy ○ Uefi

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Boot Order

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Remote CD/DVD	1				

8. Expand the vHBAs drop-down menu and select Add SAN Boot.

9. In the Add SAN Boot dialog box, enter `Fabric-A` in the vHBA field.

10. Confirm that Primary is selected for the Type option.

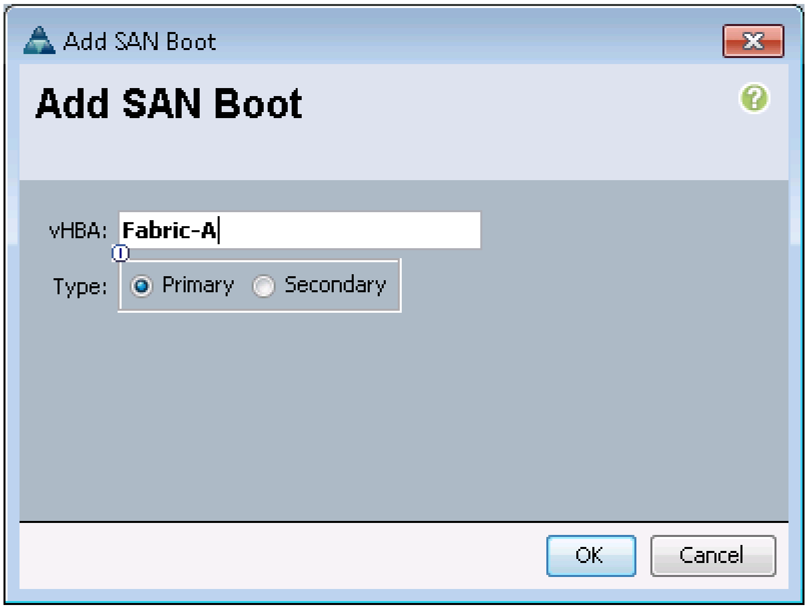

Add SAN Boot

vHBA: Fabric-A

Type: ● Primary ○ Secondary

11. Click OK to add the SAN boot initiator.

12. From the vHBA drop-down menu, select Add SAN Boot Target.

13. Keep `0` as the value for Boot Target LUN.

14. Enter the WWPN for `fcp_lif01a`.

Note: To obtain this information, log in to the storage cluster and run the `network interface show` command.

15. Select Primary for the SAN boot target type.

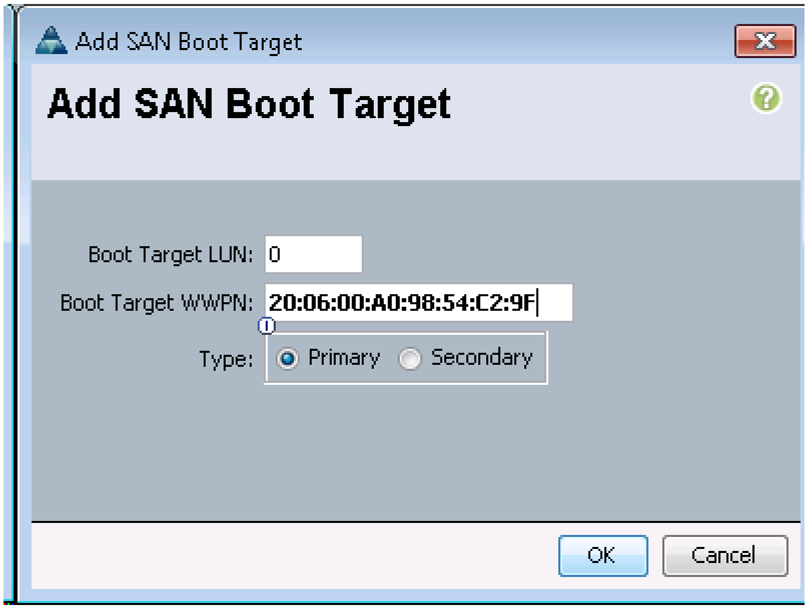

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:06:00:A0:98:54:C2:9F

Type: ● Primary ○ Secondary

16. Click OK to add the SAN boot target.

17. From the vHBA drop-down menu, select Add SAN Boot Target.

18. Enter `0` as the value for Boot Target LUN.

19. Enter the WWPN for `fcp_lif02a`.

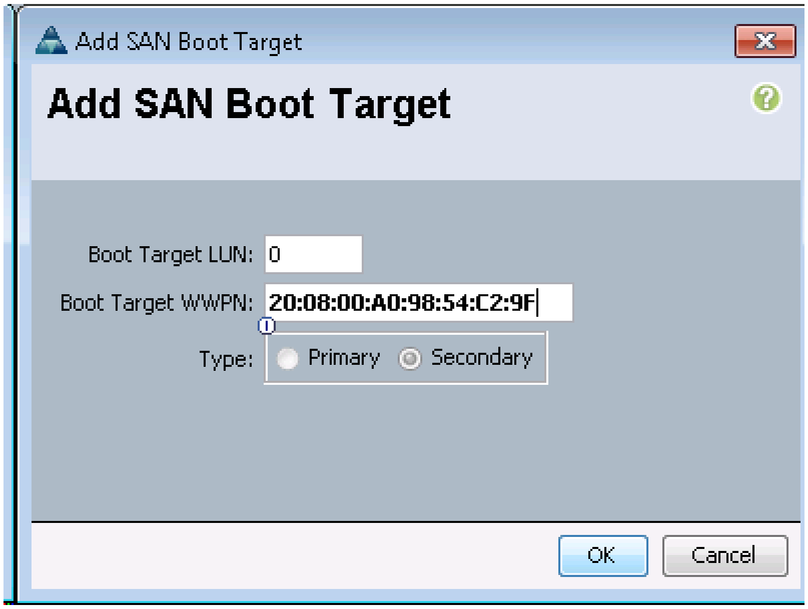

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:08:00:A0:98:54:C2:9F

Type: ○ Primary ● Secondary

20. Click OK to add the SAN boot target.

21. From the vHBA drop-down menu, select Add SAN Boot.

22. In the Add SAN Boot dialog box, enter `Fabric-B` in the vHBA box.

23. The SAN boot type should automatically be set to Secondary, and the Type option should be unavailable.

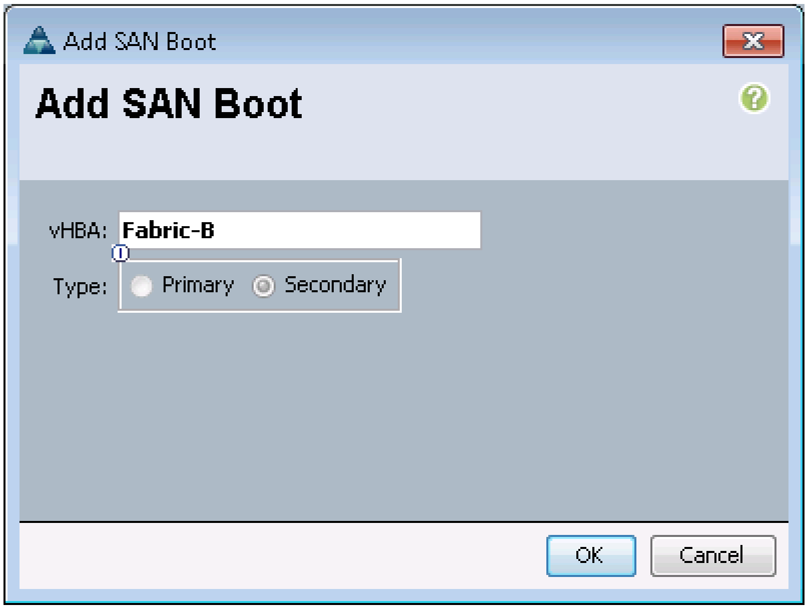

Add SAN Boot

vHBA: Fabric-B

Type: ○ Primary ● Secondary

24. Click OK to add the SAN boot initiator.

25. From the vHBA drop-down menu, select Add SAN Boot Target.

26. Keep `0` as the value for Boot Target LUN.

27. Enter the WWPN for `fcp_lif01b`.

28. Select Primary for the SAN boot target type.

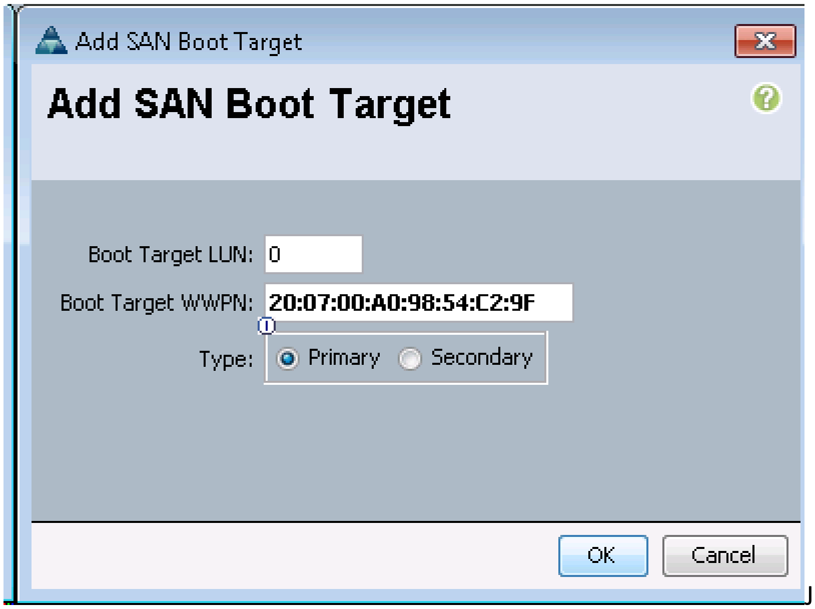

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:07:00:A0:98:54:C2:9F

Type: ● Primary ○ Secondary

John

29. Click OK to add the SAN boot target.

30. From the vHBA drop-down menu, select Add SAN Boot Target.

31. Keep `0` as the value for Boot Target LUN.

32. Enter the WWPN for `fcp_lif02b`.

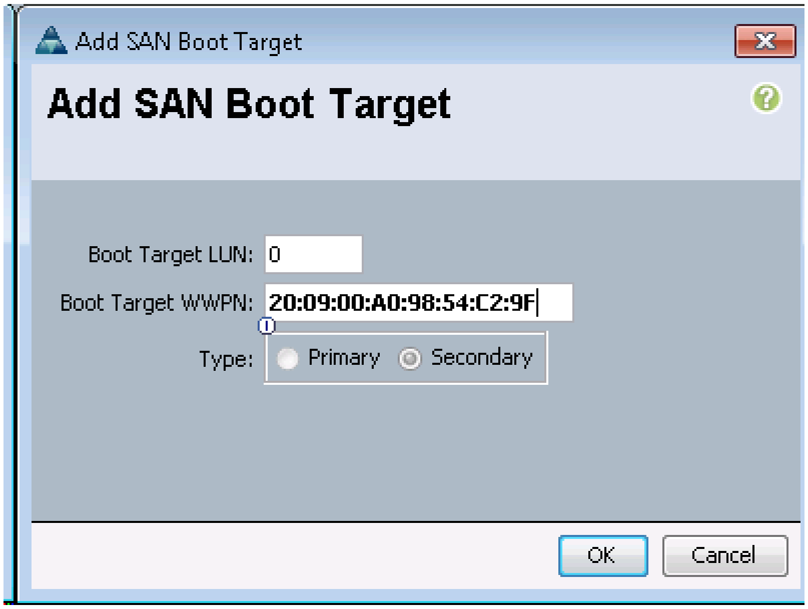

Add SAN Boot Target

Boot Target LUN: 0

Boot Target WWPN: 20:09:00:A0:98:54:C2:9F

Type: ○ Primary ● Secondary

33. Click OK to add the SAN boot target.

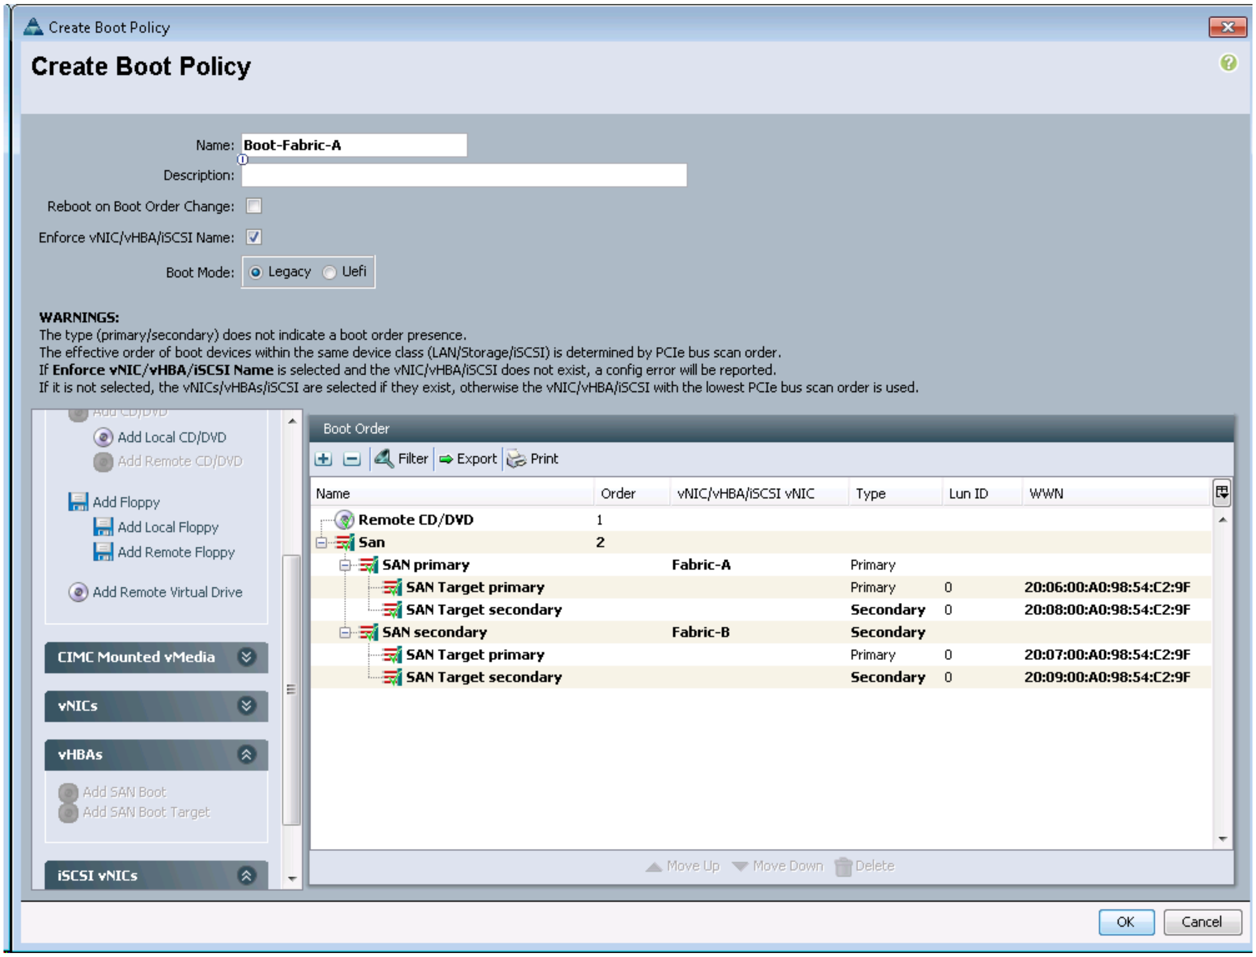

Create Boot Policy

Name: Boot-Fabric-A

Description:

Reboot on Boot Order Change: □

Enforce vNIC/vHBA/iSCSI Name: ☑

Boot Mode: ● Legacy ○ Uefi

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Remote CD/DVD	1				
San	2				
SAN primary		Fabric-A	Primary		
SAN Target primary			Primary	0	20:06:00:A0:98:54:C2:9F
SAN Target secondary			Secondary	0	20:08:00:A0:98:54:C2:9F
SAN secondary		Fabric-B	Secondary		
SAN Target primary			Primary	0	20:07:00:A0:98:54:C2:9F
SAN Target secondary			Secondary	0	20:09:00:A0:98:54:C2:9F

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

Create Service Profile Template

In this procedure, one service profile template is created for fabric A boot. To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click the Servers tab in the navigation pane.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Identify the service profile template:
    a. Enter VM-Host-Infra-Fabric-A* as the name of the service profile template. This service profile template is configured to boot from node 1 on fabric A.
    b. Select the Updating Template option.
    c. Under UUID, select UUID_Pool as the UUID pool.
    d. Click Next.

* Note: If this name has already been utilized for the Infrastructure hosts utilizing iSCSI boot, please choose a different name or append fcoe to this template name

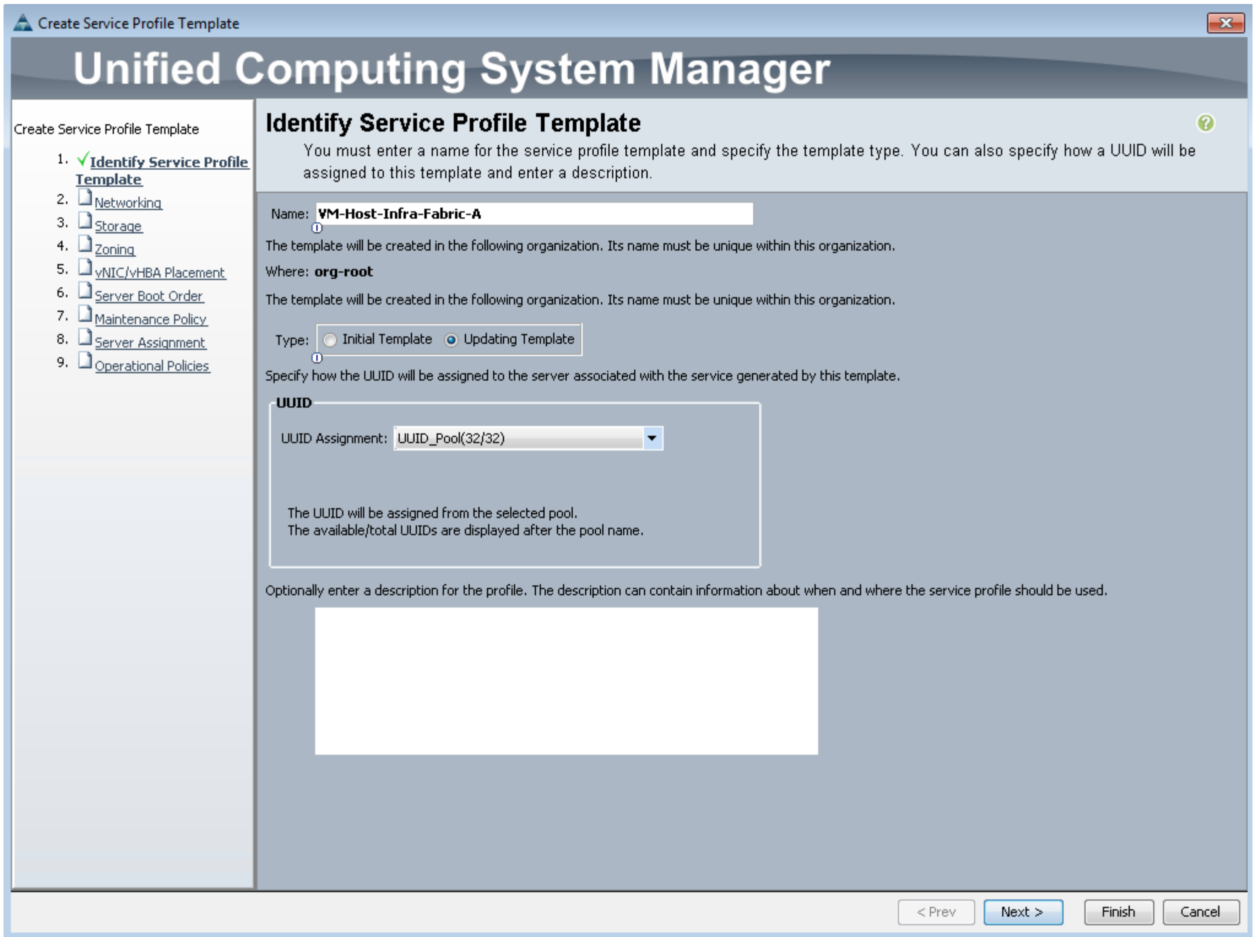

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies

Identify Service Profile Template

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

Name: VM-Host-Infra-Fabric-A

The template will be created in the following organization. Its name must be unique within this organization.

Where: **org-root**

The template will be created in the following organization. Its name must be unique within this organization.

Type: ○ Initial Template ● Updating Template

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

UUID

UUID Assignment: UUID_Pool(32/32)

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

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

< Prev | Next > | Finish | Cancel

Configure Networking Options

1. To configure the networking options, 6 vNIC interfaces will be added for Infrastructure ESXi hosts:

2. Keep the default setting for Dynamic vNIC Connection Policy.

3. Select the Expert option to configure the LAN connectivity.

4. Click the upper Add button to add a vNIC to the template.

5. In the Create vNIC dialog box, enter vNIC-A as the name of the vNIC.

6. Select the Use vNIC Template checkbox.

7. In the vNIC Template list, select vNIC_Template_A.

8. In the Adapter Policy list, select VMWare.

9. Click OK to add this vNIC to the template.

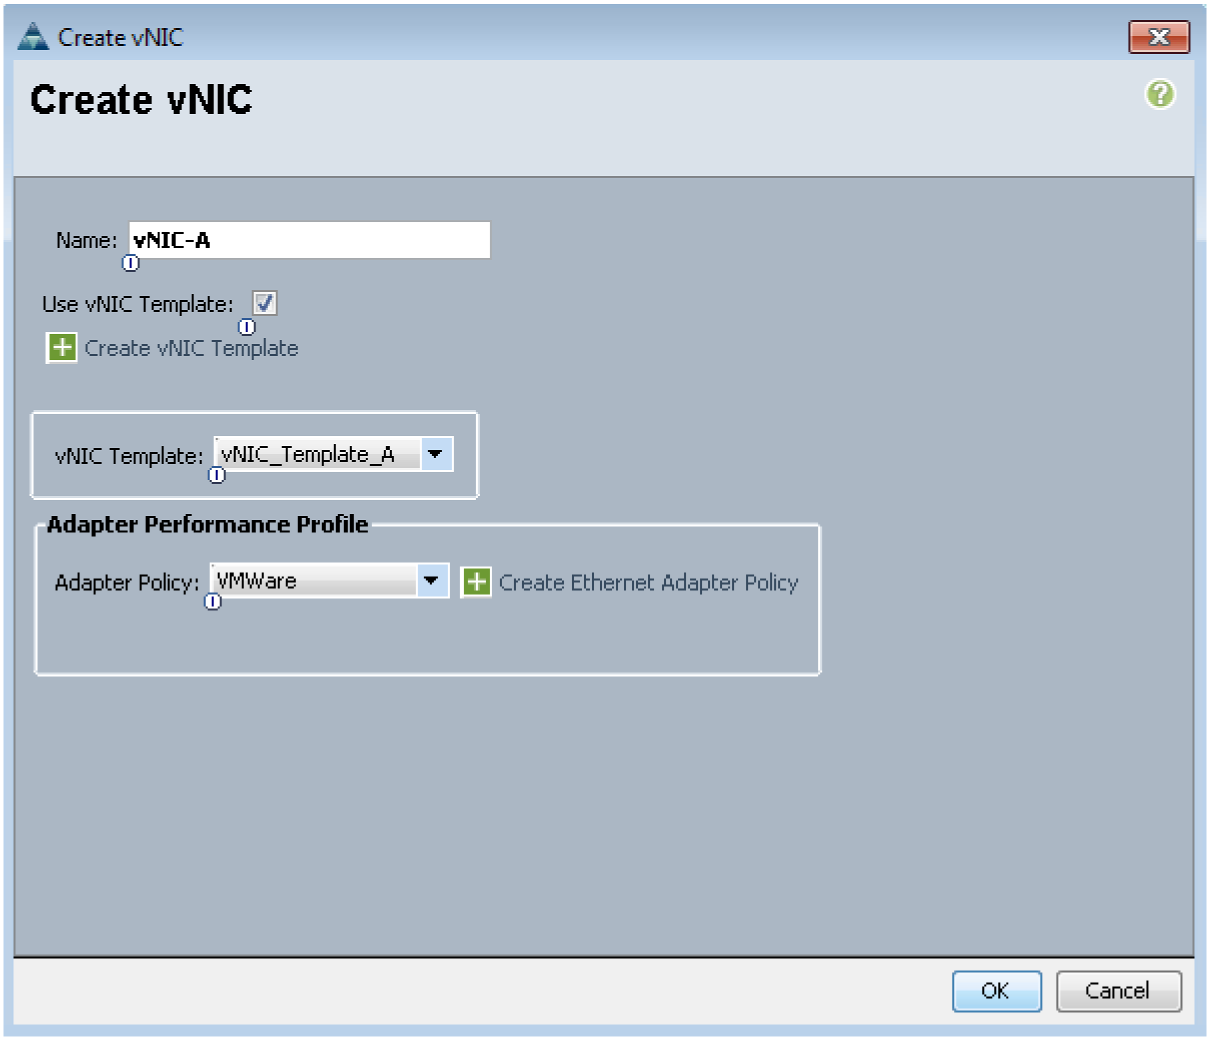

Create vNIC

Name: vNIC-A

Use vNIC Template: ☑

vNIC Template: vNIC_Template_A

Adapter Performance Profile

Adapter Policy: VMWare

10. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

11. In the Create vNIC box, enter vNIC-B as the name of the vNIC.

12. Select the Use vNIC Template checkbox.

13. In the vNIC Template list, select vNIC_Template_B.

14. In the Adapter Policy list, select VMWare.

15. Click OK to add the vNIC to the template.

16. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

17. In the Create vNIC box, enter OOB-A as the name of the vNIC.

18. Select the Use vNIC Template checkbox.

19. In the vNIC Template list, select OOB-A.

20. In the Adapter Policy list, select VMWare.

21. Click OK to add the vNIC to the template.

22. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

23. In the Create vNIC box, enter OOB-B as the name of the vNIC.

24. Select the Use vNIC Template checkbox.

25. In the vNIC Template list, select OOB-B.

26. In the Adapter Policy list, select VMWare.

27. Click OK to add the vNIC to the template.

Note: The next two vNIC interfaces are only needed for Infrastructure ESXi Hosts. These interfaces enable NFS access using NFS specific vSwitch and static EPG mapping. ESXi servers not hosting infrastructure VMs, use VDS for NFS access to application specific SVMs.

28. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

29. In the Create vNIC box, enter NFS-A as the name of the vNIC.

30. Select the Use vNIC Template checkbox.

31. In the vNIC Template list, select Infra_NFS_A.

32. In the Adapter Policy list, select VMWare.

33. Click OK to add the vNIC to the template.

34. On the Networking page of the wizard, click the upper Add button to add another vNIC to the template.

35. In the Create vNIC box, enter NFS-B as the name of the vNIC.

36. Select the Use vNIC Template checkbox.

37. In the vNIC Template list, select Infra_NFS_B.

38. In the Adapter Policy list, select VMWare.

39. Click OK to add the vNIC to the template.

40. Verify 6 vNIC interfaces are present.

41. Review the table in the Networking page to make sure that all vNICs were created.

42. Click Next.

Configure Storage Options

1. Select a local disk configuration policy:
   a. If the server in question has local disks, select default in the Local Storage list.
   b. If the server in question does not have local disks, select `SAN-Boot`.
2. Select the Expert option for the How would you like to configure SAN connectivity? field.
3. Select WWNN_Pool for the WWNN Assignment.
4. Click the Add button to add a vHBA.
5. Enter Fabric-A as the vHBA name and select the checkbox for Use vHBA Template.
6. Select the vHBA_Template_A vHBA Template and the VMWare Adapter Policy.

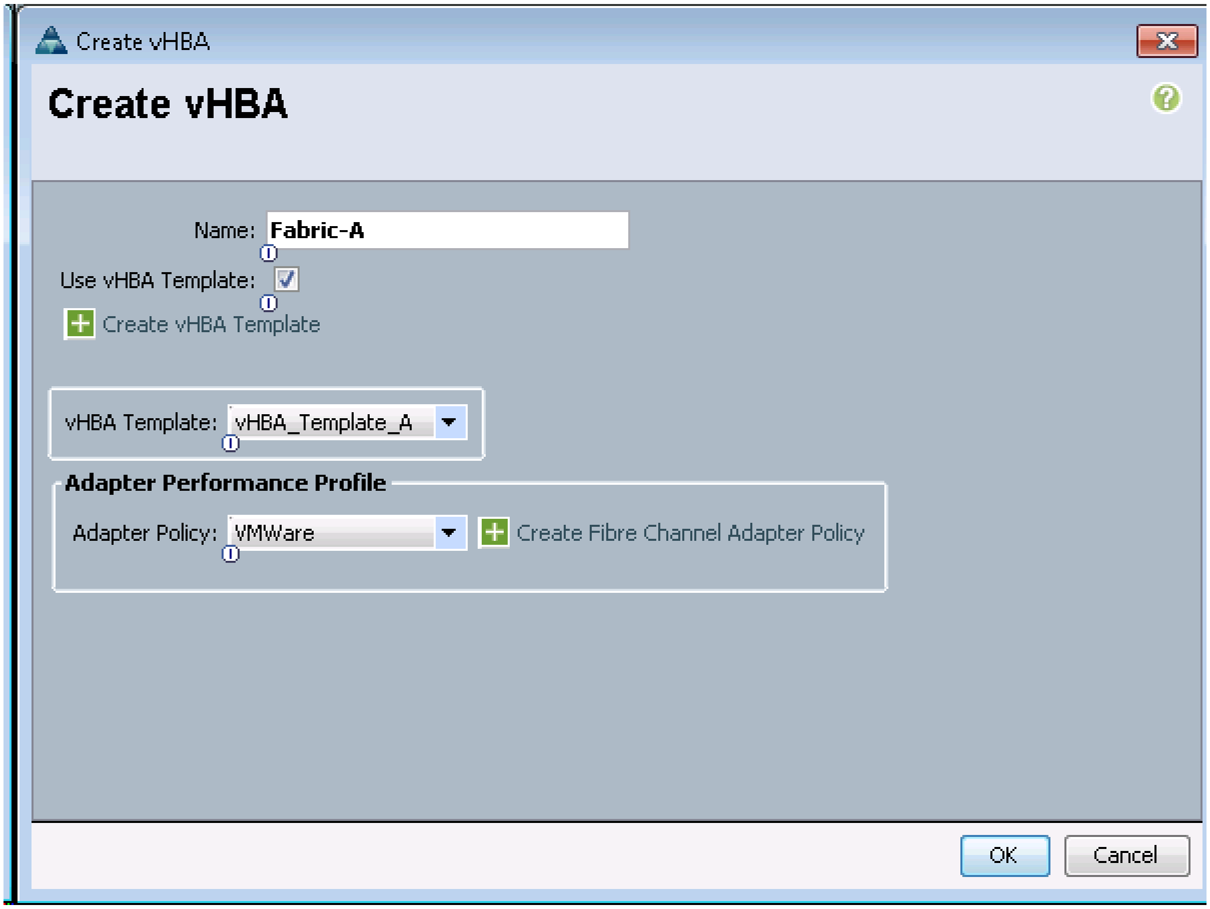
Create vHBA
Create vHBA
Create vHBA Template
Adapter Performance Profile
Create Fibre Channel Adapter Policy
OK
Cancel

Name: Fabric-A

Use vHBA Template: ☑

vHBA Template: vHBA_Template_A

Adapter Policy: VMWare

7. Click OK to add the vHBA.
8. Click the Add button to add a vHBA.
9. Enter Fabric-B as the vHBA name and select the checkbox for Use vHBA Template.
10. Select the vHBA_Template_B vHBA Template and the VMWare Adapter Policy.

11. Click OK to add the vHBA.

12. Click Next.

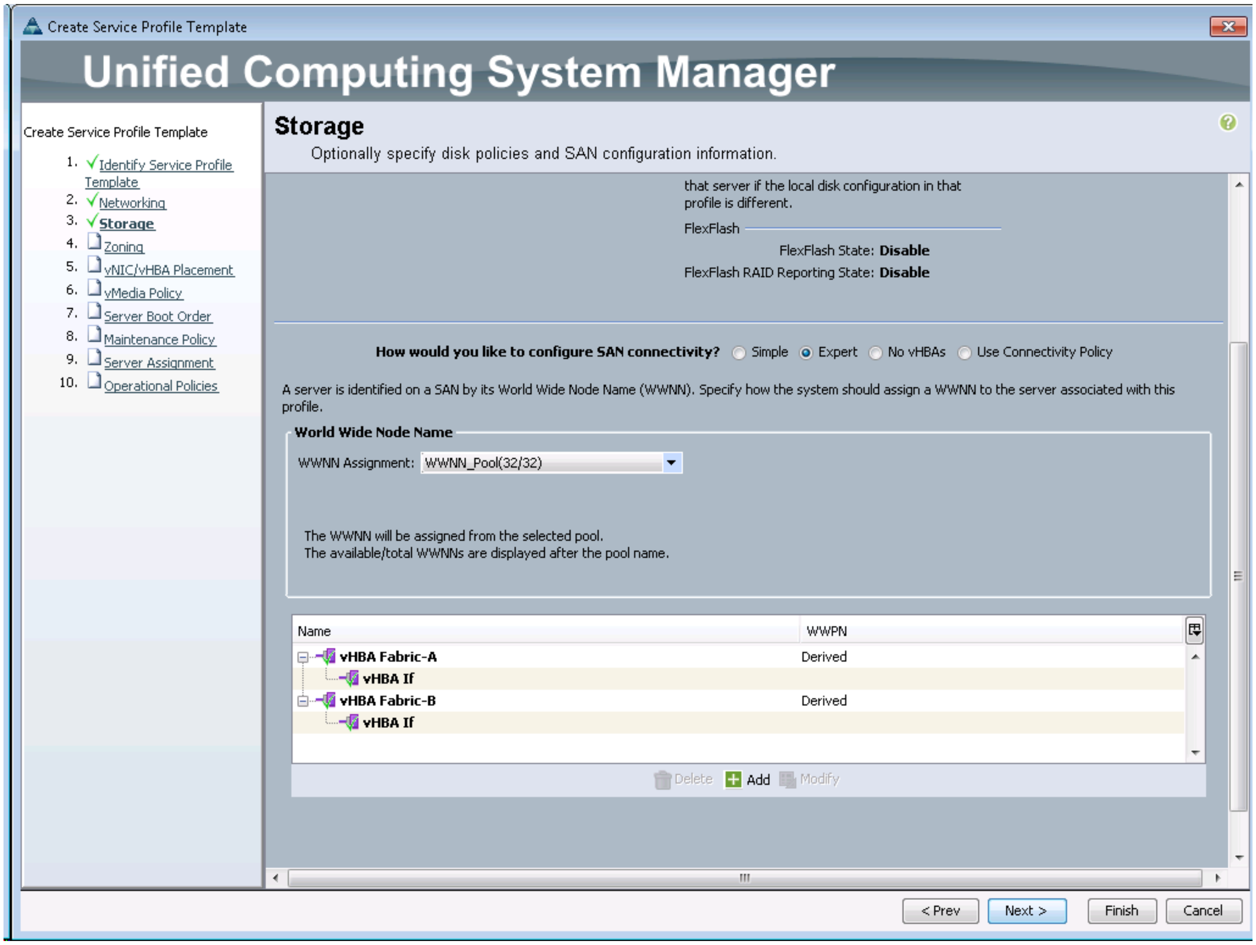

Storage

Optionally specify disk policies and SAN configuration information.

that server if the local disk configuration in that profile is different.

FlexFlash

FlexFlash State: **Disable**

FlexFlash RAID Reporting State: **Disable**

How would you like to configure SAN connectivity? ○ Simple ● Expert ○ No vHBAs ○ Use Connectivity Policy

A server is identified on a SAN by its World Wide Node Name (WWNN). Specify how the system should assign a WWNN to the server associated with this profile.

World Wide Node Name

WWNN Assignment: WWNN_Pool(32/32)

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

Name	WWPN
vHBA Fabric-A	Derived
vHBA If	
vHBA Fabric-B	Derived
vHBA If	

Configure Zoning Options

1. In the Zoning window, click the Add button.

2. Name the vHBA Initiator Group, Fabric-A and select the Infra-Fabric-A Storage Connection Policy.

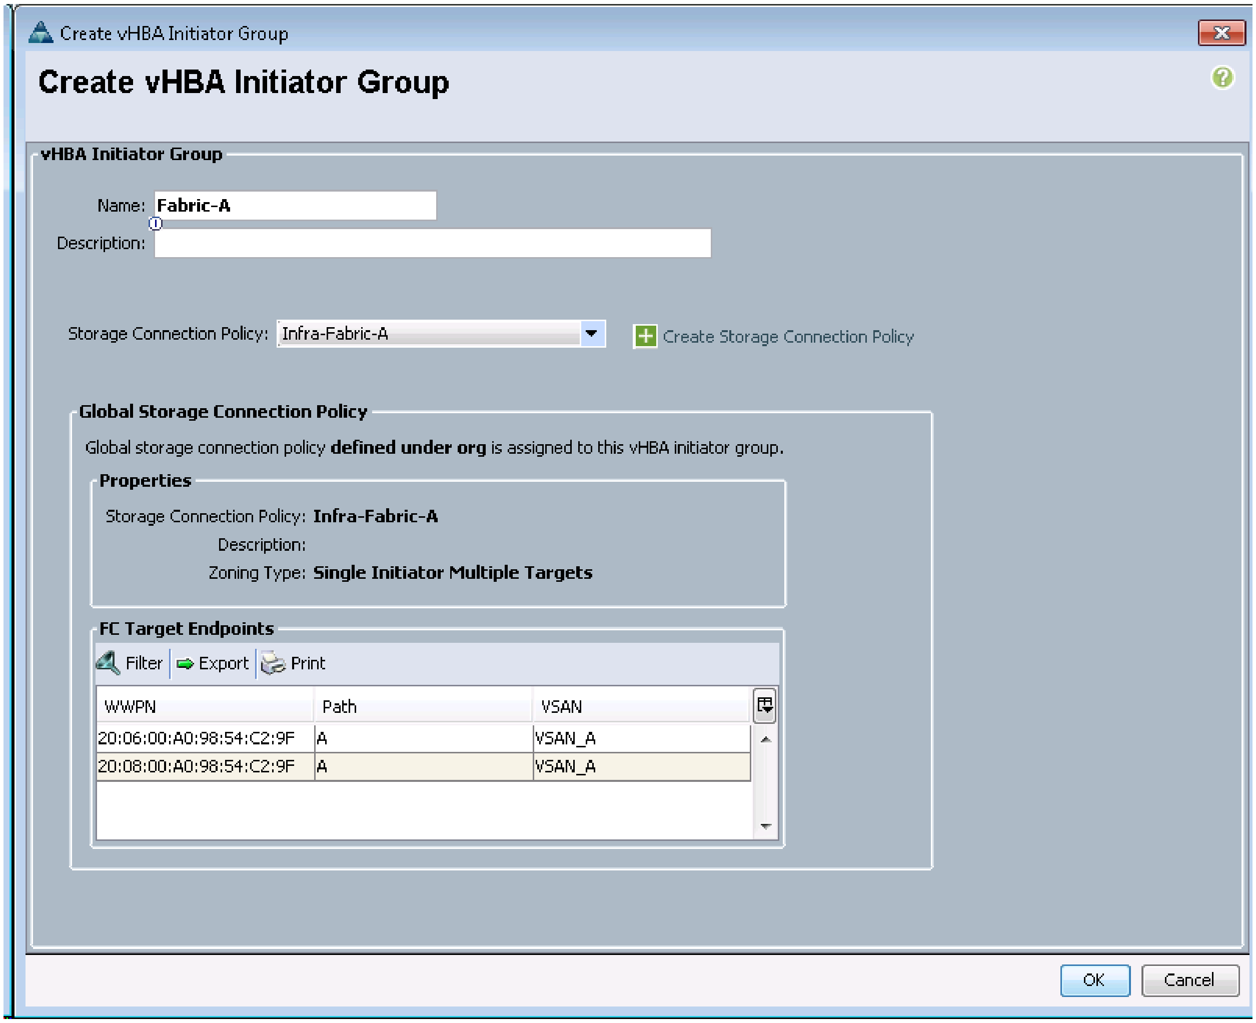
Create vHBA Initiator Group

Create vHBA Initiator Group

vHBA Initiator Group

Name: **Fabric-A**

Description:

Storage Connection Policy: Infra-Fabric-A

Global Storage Connection Policy

Global storage connection policy **defined under org** is assigned to this vHBA initiator group.

Properties

Storage Connection Policy: **Infra-Fabric-A**

Description:

Zoning Type: **Single Initiator Multiple Targets**

FC Target Endpoints

WWPN	Path	VSAN
20:06:00:A0:98:54:C2:9F	A	VSAN_A
20:08:00:A0:98:54:C2:9F	A	VSAN_A

3. Click OK to add the vHBA Initiator Group.

4. In the Zonig window, click the Add button.

5. Name the vHBA Initiator Group, Fabric-B and select the Infra-Fabric-B Storage Connection Policy.

6. Click OK to add the vHBA Initiator Group.

7. In the Zoning window, select Fabric A in the list of vHBA Initiators and Fabric-A in the list of vHBA Initiator Groups.

8. Click the >> Add To >> button.

9. In the Zoning window, select Fabric B in the list of vHBA Initiators and Fabric-B in the list of vHBA Initiator Groups.

10. Click the >> Add To >> button.

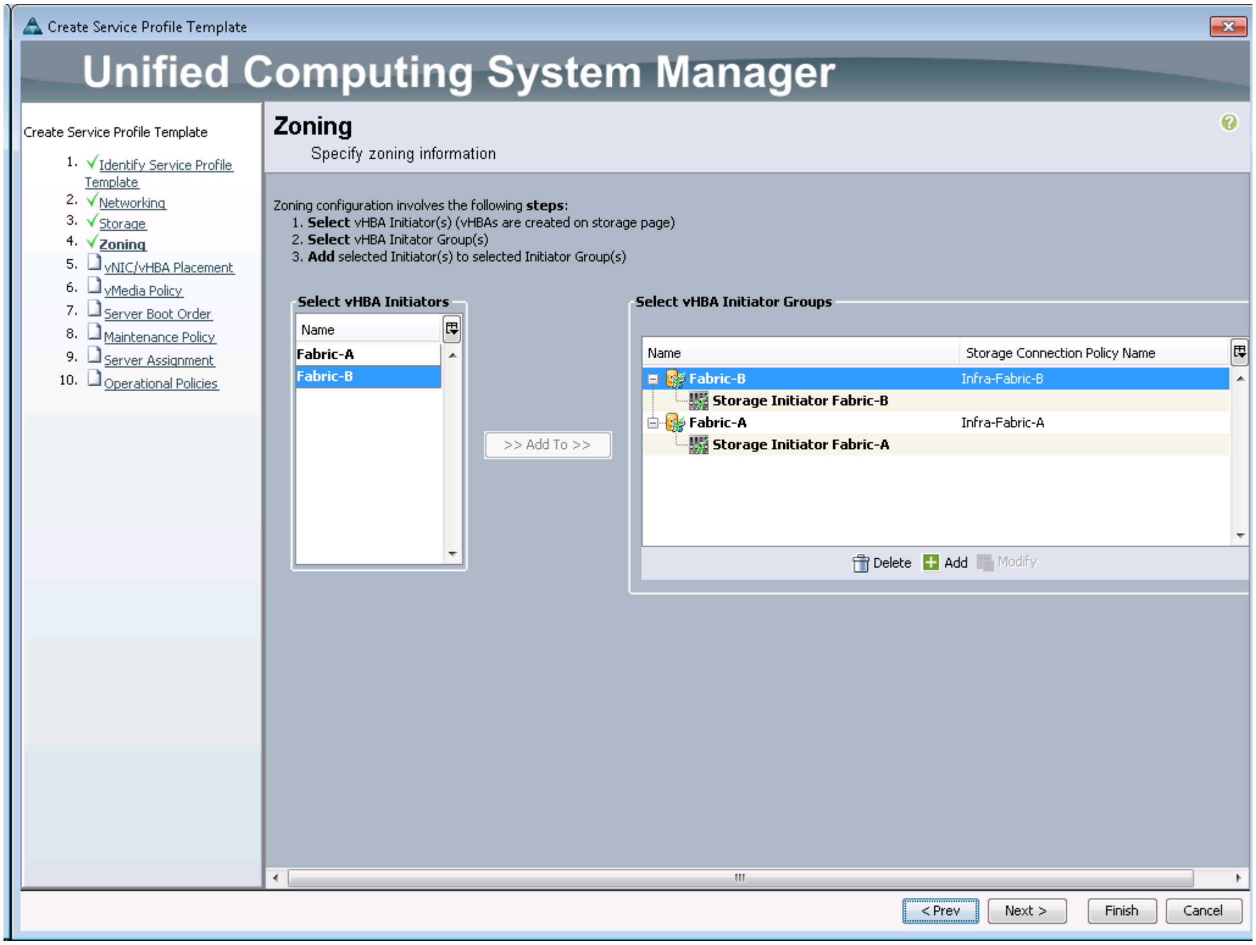

Create Service Profile Template

Unified Computing System Manager

Create Service Profile Template

1. ✓ Identify Service Profile Template
2. ✓ Networking
3. ✓ Storage
4. ✓ **Zoning**
5. □ vNIC/vHBA Placement
6. □ vMedia Policy
7. □ Server Boot Order
8. □ Maintenance Policy
9. □ Server Assignment
10. □ Operational Policies

Zoning

Specify zoning information

Zoning configuration involves the following **steps**:

1. **Select** vHBA Initiator(s) (vHBAs are created on storage page)
2. **Select** vHBA Initator Group(s)
3. **Add** selected Initiator(s) to selected Initiator Group(s)

Select vHBA Initiators

Name
Fabric-A
Fabric-B

Select vHBA Initiator Groups

Name	Storage Connection Policy Name
Fabric-B	Infra-Fabric-B
Storage Initiator Fabric-B	
Fabric-A	Infra-Fabric-A
Storage Initiator Fabric-A	

11. Click Next.

vNIC/vHBA Placement

1. In the Select Placement list, select the VM-Host-Infra placement policy.

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

   a. vNIC-A

   b. vNIC-B

   c. OOB-A

   d. OOB-B

   e. NFS-A

   f. NFS-B

   g. vHBA Fabric-A

   h. vHBA Fabric-B

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

4. Click Next.

Server Boot Order

1. Select Boot-Fabric-A for Boot Policy.

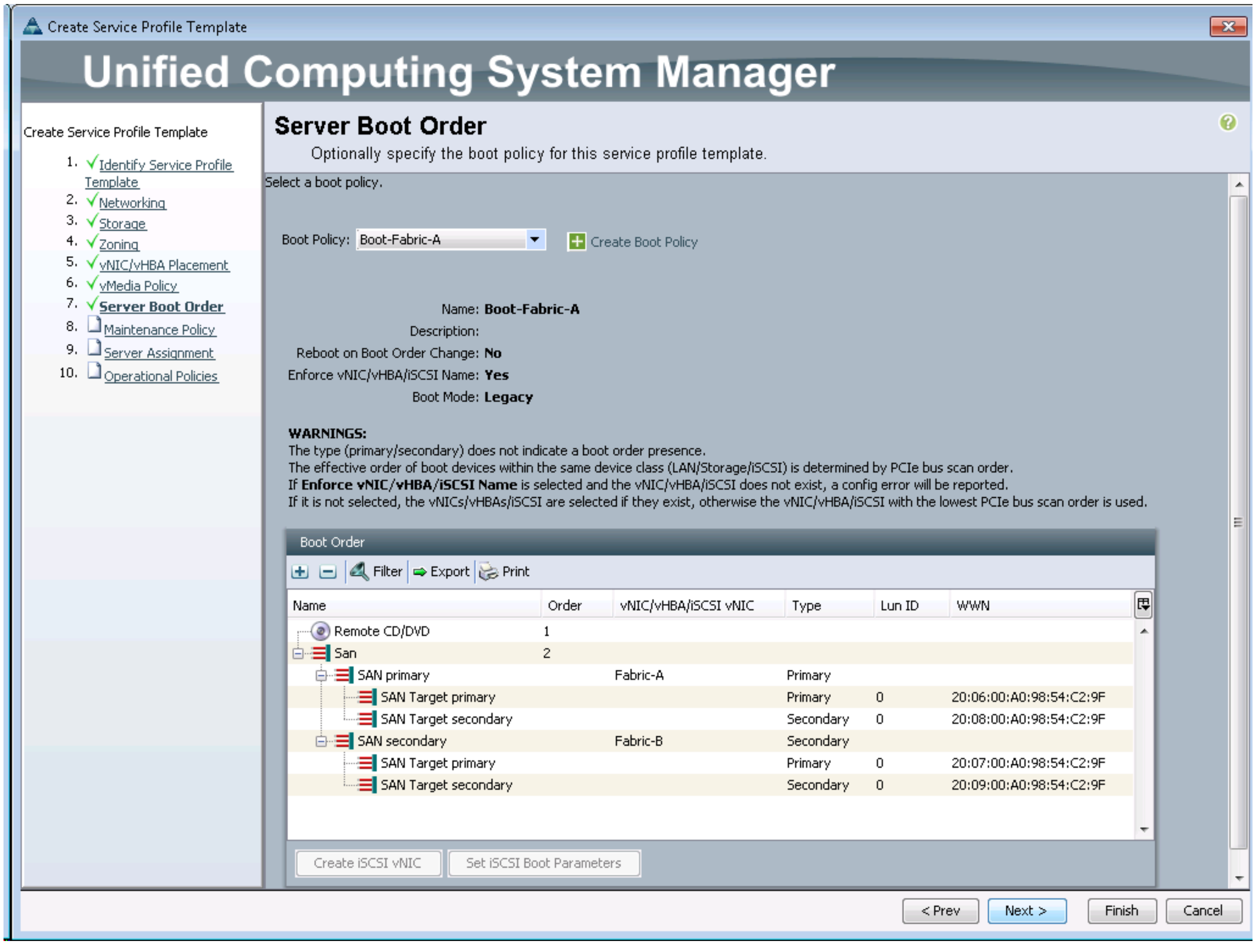

Server Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: Boot-Fabric-A

Name: **Boot-Fabric-A**
Description:
Reboot on Boot Order Change: **No**
Enforce vNIC/vHBA/iSCSI Name: **Yes**
Boot Mode: **Legacy**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs/iSCSI are selected if they exist, otherwise the vNIC/vHBA/iSCSI with the lowest PCIe bus scan order is used.

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Remote CD/DVD	1				
San	2				
SAN primary		Fabric-A	Primary		
SAN Target primary			Primary	0	20:06:00:A0:98:54:C2:9F
SAN Target secondary			Secondary	0	20:08:00:A0:98:54:C2:9F
SAN secondary		Fabric-B	Secondary		
SAN Target primary			Primary	0	20:07:00:A0:98:54:C2:9F
SAN Target secondary			Secondary	0	20:09:00:A0:98:54:C2:9F

2. Review the table to make sure that all boot devices were created and identified. Verify that the boot devices are in the correct boot sequence.

3. Click Next to continue to the next section.

Maintenance Policy

1. Select the default Maintenance Policy.

2. Click Next.

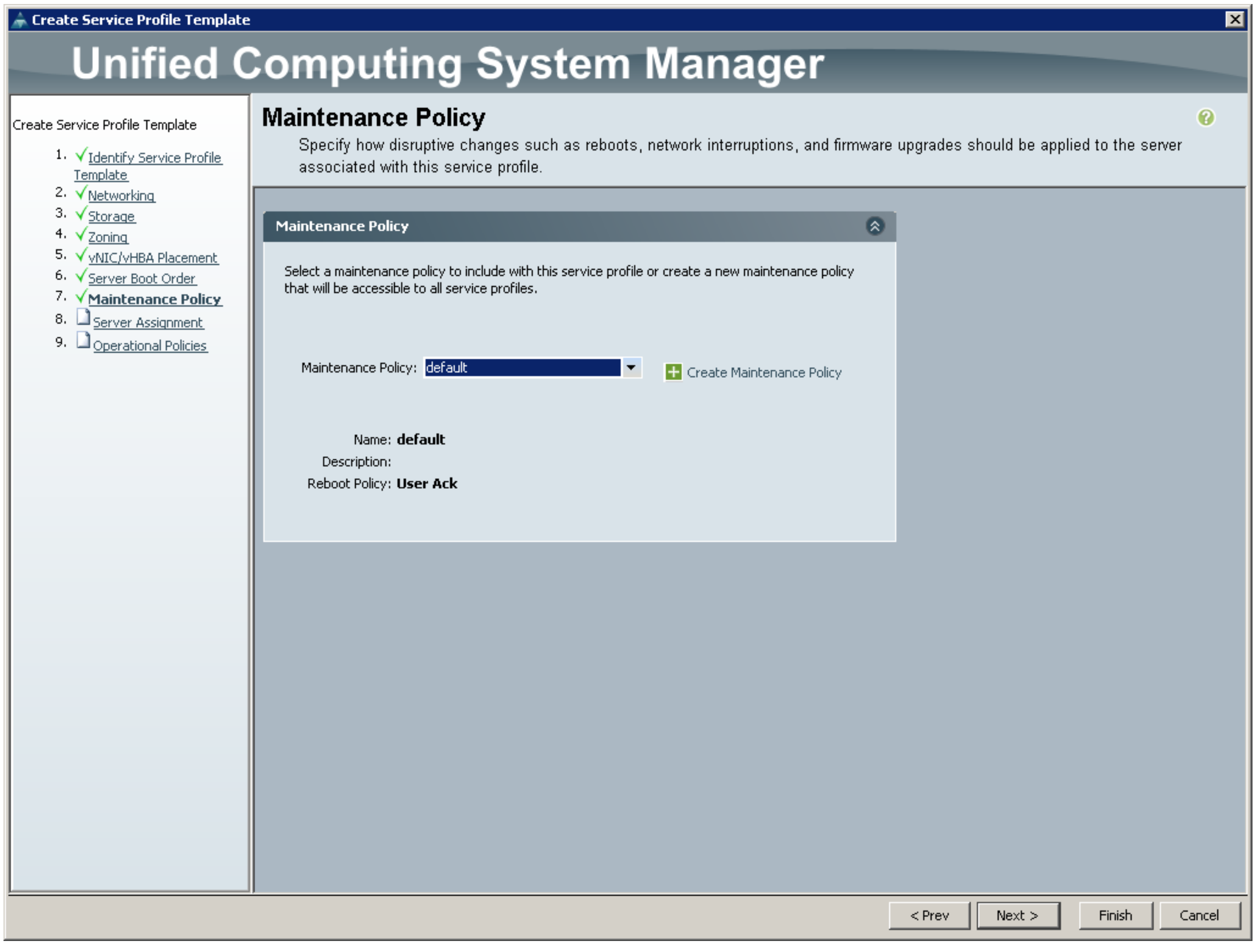

Maintenance Policy

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: default

Name: **default**
Description:
Reboot Policy: **User Ack**

Server Assignment

1. In the Pool Assignment list, select `Infra_Pool`.

2. Optional: Select a Server Pool Qualification policy.

3. Select Down as the power state to be applied when the profile is associated with the server.

4. Expand Firmware Management at the bottom of the page and select `VM-Host-Infra` from the Host Firmware list.

5. Click Next.

Operational Policies

1. In the BIOS Policy list, select VM-Host-Infra.

2. Expand Power Control Policy Configuration and select No-Power-Cap in the Power Control Policy list.

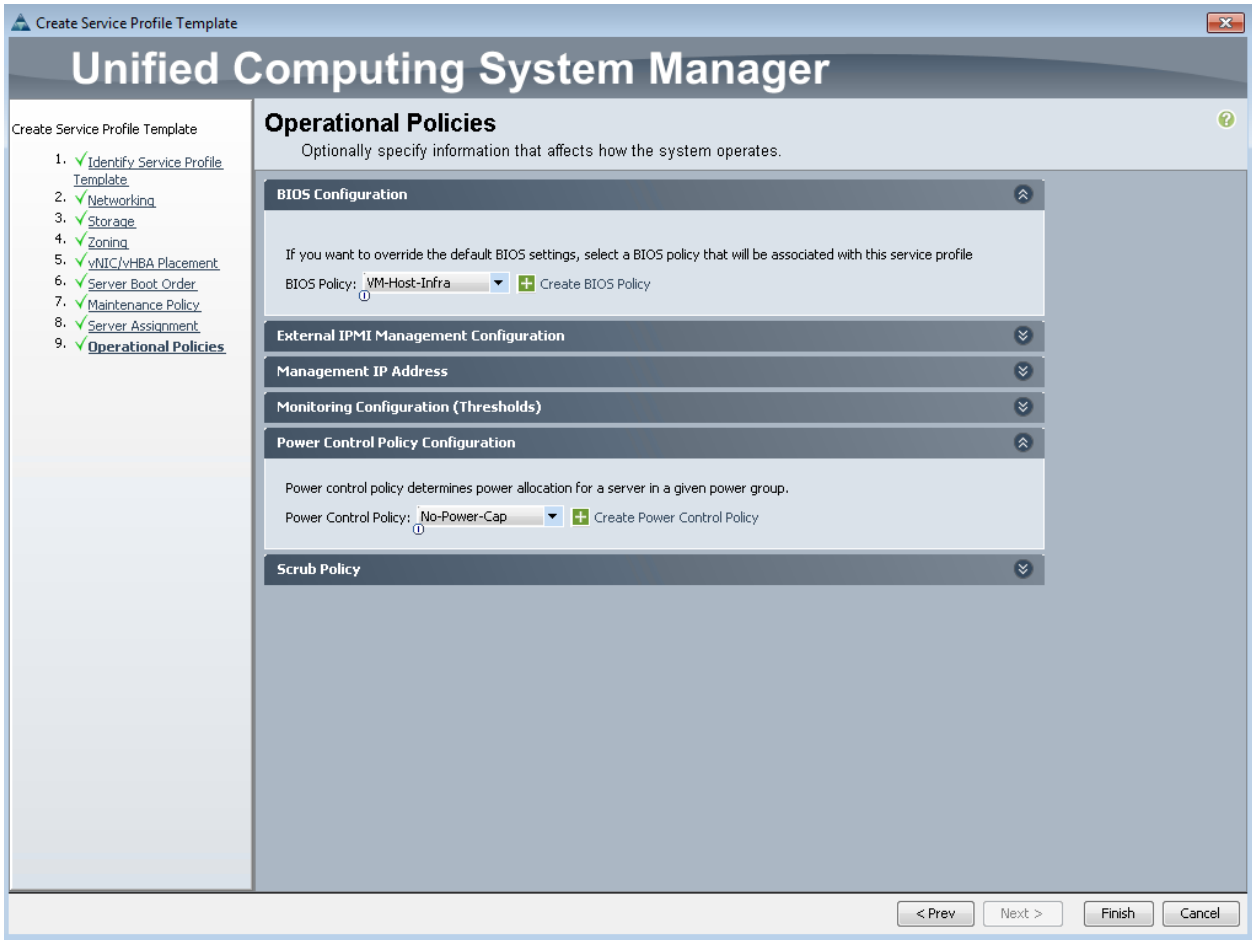
Create Service Profile Template
Unified Computing System Manager
Create Service Profile Template
1. Identify Service Profile Template
2. Networking
3. Storage
4. Zoning
5. vNIC/vHBA Placement
6. Server Boot Order
7. Maintenance Policy
8. Server Assignment
9. Operational Policies
BIOS Configuration
Create BIOS Policy
External IPMI Management Configuration
Management IP Address
Monitoring Configuration (Thresholds)
Power Control Policy Configuration
Create Power Control Policy
Scrub Policy
< Prev
Next >
Finish
Cancel

Operational Policies

Optionally specify information that affects how the system operates.

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

BIOS Policy: VM-Host-Infra

Power control policy determines power allocation for a server in a given power group.

Power Control Policy: No-Power-Cap

3. Click Finish to create the service profile template.

4. Click OK in the confirmation message.

Create Service Profiles

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

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

2. Select Service Profile Templates > root > Service Template VM-Host-Infra-Fabric-A.

3. Right-click `VM-Host-Infra-Fabric-A` and select Create Service Profiles from.

4. Enter `VM-Host-Infra-0` as the service profile prefix.

5. Enter `1` as Name Suffix Starting Number

6. Enter `1` as the Number of Instances. Adjust the number based on required number of servers.

7. Click OK to create the service profile.

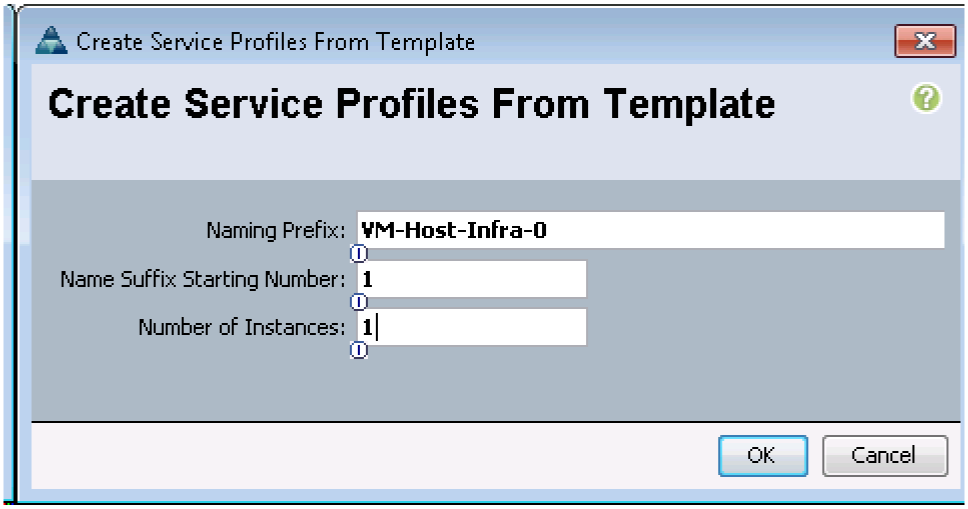
Create Service Profiles From Template
Create Service Profiles From Template
OK
Cancel

Naming Prefix: VM-Host-Infra-0
Name Suffix Starting Number: 1
Number of Instances: 1

8. Click OK in the confirmation message.

Clustered Data ONTAP SAN Boot Storage Setup

Create Igroups

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

```
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-01 –protocol fcp –ostype
vmware –initiator <<var_vm_host_infra_01_A_wwpn>>,
<<var_vm_host_infra_01_B_wwpn>>
igroup create –vserver Infra-SVM –igroup VM-Host-Infra-02 –protocol fcp –ostype
vmware –initiator <<var_vm_host_infra_02_A_wwpn>>,
<<var_vm_host_infra_02_B_wwpn>>
igroup create –vserver Infra-SVM –igroup MGMT-Hosts –protocol fcp –ostype vmware
–initiator <<var_vm_host_infra_01_A_wwpn>>, <<var_vm_host_infra_01_B_wwpn>>,
<<var_vm_host_infra_02_A_wwpn>>, <<var_vm_host_infra_02_B_wwpn>>
```

Note: To view the three igroups just recently created, enter `igroup show`.

Map Boot LUNs to Igroups

1. From the cluster management SSH connection, enter the following:

```
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-01 –igroup VM-
Host-Infra-01 –lun-id 0
lun map –vserver Infra-SVM –volume esxi_boot –lun VM-Host-Infra-02 –igroup VM-
Host-Infra-02 –lun-id 0
```

About the Authors

Haseeb Niazi, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems Inc.

Haseeb Niazi has over 16 years of experience at Cisco focused on Data Center, Security, WAN Optimization, and related technologies. As a member of various solution teams and advanced services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. Haseeb holds a master's degree in Computer Engineering from the University of Southern California.

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

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

John George, Reference Architect, Infrastructure and Cloud Engineering, NetApp

John George is a Reference Architect in the NetApp Infrastructure and Cloud Engineering team and is focused on developing, validating, and supporting cloud infrastructure solutions that include NetApp products. Before his current role, he supported and administered Nortel's worldwide training network and VPN infrastructure. John holds a Master's degree in computer engineering from Clemson University.

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, Cisco Systems Inc.