

## Cisco UCS Mini Branch Office Solution for EMC VSPEX

With EMC VNXe3200 and VMware vSphere 5.5 for up to 100 Virtual Machines


Last Updated: September 10, 2014



Cisco  
Validated  
Design



Building Architectures to Solve Business Problems



## About the Authors



### **Vijay Durairaj, Technical Marketing Engineer, Cisco Systems**

Vijay Durairaj is a Technical Marketing Engineer with Cisco UCS Solutions and Performance group has over 10+ years of experience in UCS, network, storage and server virtualization design. Vijay has worked on performance and benchmarking on Cisco UCS servers and has delivered benchmark results on SPEC CPU2006 and SPECj ENT 2010. Vijay holds certification in VMware Certified Professional and Cisco Unified Computing systems Design specialist.



### **Mehul Bhatt, Virtualization Architect, Cisco Systems**

Mehul Bhatt is a Virtualization Architect with Server Access virtualization Business Unit (SAV) with over 13 years of Experience in virtually all layers of computer networking. His focus area includes Unified Compute Systems, network and server virtualization design. Prior to joining Cisco Technical Marketing team, Mehul was Technical Lead at Cisco, Nuova systems and Bluecoat systems. Mehul holds a Master's degree in computer systems engineering and holds various Cisco career certifications.

# Acknowledgment

For their support and contribution to the design, validation, and creation of the Cisco Validated Design, we would like to thank:

- Shiva Shastri (Cisco)
- Bathu Krishnan (Cisco)
- Kevin Phillips (EMC)



## About Cisco Validated Design (CVD) Program

---

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit <http://www.cisco.com/go/designzone>.

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.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R).

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

© 2014 Cisco Systems, Inc. All rights reserved.



# Cisco UCS Mini Branch Office Solution for EMC VSPEX

---

## Executive Summary

Cisco solution on EMC VSPEX is a pre-validated and modular architecture built with proven best of-breed technologies. Because EMC VSPEX solutions are rigorously tested, they drastically reduce server virtualization planning and configuration overhead while contributing to IT transformation through faster deployments, greater choice of components and flexibility at reduced risk.

This Cisco Validated Design (CVD) focuses on a VSPEX solution consisting of new components from Cisco such as the UCS Mini compute chassis and EMC VNXe3200 storage array with VMware vSphere 5.5 catered to the small to medium business (SMB) segment with an initial need of about 100 Virtual Machines. A typical use case would be a branch office/remote office location with centralized management to ensure consistent standards based deployment. The platform has sufficient scalability in compute and storage areas, if necessary.

This Cisco Validated Design document defines the architectural design and deployment procedure of the previously-defined VSPEX VMware platform with a focus on features and options that underscore functionality, scalability and standardized management as well as simplicity, efficiency, and flexibility for a platform which can be an extension of a data center solution or serve as a standalone platform with similar benefits.

## Introduction

Virtualization is a critical deployment strategy for reducing the Total Cost of Ownership (TCO).

It allows for consolidation for better utilization of underlying compute, network and storage components. However, selecting the appropriate platform for virtualization can be confusing given the myriad of choices at every level. Platforms should be flexible for scaling while also being reliable and cost effective to facilitate virtualization. In the VSPEX converged infrastructure, compatible components come together for a scalable reference architecture with provisions for scale within each individual component. Cisco solutions implemented as part of EMC VSPEX reference architectures leverage available flexibility and functionality for effective resource utilization while also preserving existing support structure.



## Target Audience

The reader of this document is expected to have the necessary training and background to install and configure VMware vSphere, EMC VNXe series storage arrays, and Cisco UCS Mini chassis (mini), Unified Computing Systems Manager (UCSM) and UCS Central Management. External references are provided where applicable and it is recommended that the reader be familiar with these documents.

Readers are also expected to be familiar with the infrastructure and database security policies of the customer installation.

## Purpose of this Document

This document illustrates the design and deployment steps required for implementing the Cisco branch office solution on EMC VSPEX with VMware 5.5 as the hypervisor. Centralized management of this branch office solution through UCS Central located in the data center is also highlighted. The level of details covered allows for confirmation of the correct functioning of the basic components in the solution. The solution documented is expected to cover a VMware architecture for small-to medium-sized businesses with a need for about 100 VMs. This document shows the solution with EMC VNXe 3200 series storage array using NFS for data storage and Fiber Channel (FC) for OS booting through the pair of Cisco UCS 6324 Fabric Interconnects. While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are specifically mentioned.

Following are elements within scope of this Cisco solution:

- Provide an end-to-end virtualization solution to take full advantage of unified infrastructure components.
- Show implementation progression of VMware vCenter 5.5 design and results.
- Provide a reliable, flexible and scalable reference design
- It is beyond the scope of this document to consider performance related details pertaining to the solution.

## Business Needs

Businesses have always had a need for consistent provisioning and management of remote office IT resources. The new Cisco UCS Mini presents appropriate levels of compute and connectivity options to cater to the needs of a branch office while leveraging inherent strengths of UCS Manager resident within the chassis. Complementing this setup is integration with UCS Central in the data center for a hierarchical management structure ensuring consistent standards based deployment and management of all branch office sites from a central office. Efficiencies due to converged stacks such as the VSPEX are further enhanced when integrated with centralized provisioning and management.

# Solution Overview

## Cisco UCS Mini Branch solution for EMC VSPEX VMware architecture

This solution provides an end-to-end architecture with Cisco, EMC, VMware, and Microsoft technologies that demonstrate support for up to 100 generic virtual machines and provide high availability and server redundancy.

The following are the components used for the design and deployment:

- Cisco Unified Computing System
- Cisco UCS Manager 3.0(1c)
- Cisco UCS B200M3 server
- Cisco UCS VIC adapters
- Cisco UCS Central (ROBO Implementation)
- EMC VNXe3200
- VMware vCenter 5.5
- Microsoft SQL database
- VMware DRS
- VMware HA

The solution is designed to host scalable, mixed application workloads of up to 100 reference virtual machines.

## Technology Overview

### Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain.

The main components of the Cisco UCS are:


- **Compute** - The system is based on an entirely new class of computing system that incorporates rack mount and blade servers based on Intel Xeon 2600 v2 Series Processors.
- **Network** - The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

- **Virtualization** - The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access** - The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet (NFS or iSCSI), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This provides customers with storage choices and investment protection. In addition, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.

Cisco Unified Computing System (Cisco UCS) fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

The Cisco UCS 6324 Fabric Interconnect extends the Cisco UCS architecture into environments with lesser resource requirements. Providing the same unified server and networking capabilities as the full-scale Cisco UCS solution, the Cisco UCS 6324 Fabric Interconnect embeds the connectivity within the Cisco UCS 5108 Blade Server Chassis to provide a smaller domain of up to 15 servers (8 blade servers and up to 7 direct-connect rack servers).

*Figure 1*      *Figure 1. The Cisco UCS Mini Architecture*



## Cisco UCS Manager 3.0

Cisco Unified Computing System (UCS) Manager provides unified, embedded management of all software and hardware components of the Cisco UCS through choice of an intuitive GUI, a Command Line Interface (CLI) or an XML API. The Cisco UCS Manager provides unified management domain with centralized management capabilities and controls multiple chassis and thousands of virtual machines.

The Cisco UCS 6324 Fabric Interconnect hosts and runs Cisco UCS Manager in a highly available configuration, enabling the fabric interconnects to fully manage all Cisco UCS elements. The Cisco UCS 6324 Fabric Interconnects support out-of-band management through dedicated 10/100/1000-Mbps Ethernet management ports. Cisco UCS Manager typically is deployed in a clustered active-passive configuration with two UCS 6324 Fabric Interconnects connected through the cluster interconnect built into the chassis.

Cisco UCS Manager 3.0 supports the 6324 Fabric Interconnect that integrates the FI into the UCS Chassis and provides an integrated solution for a smaller deployment environment. Cisco UCS Mini simplifies the system management and saves cost for smaller scale deployments. The hardware and software components support Cisco unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

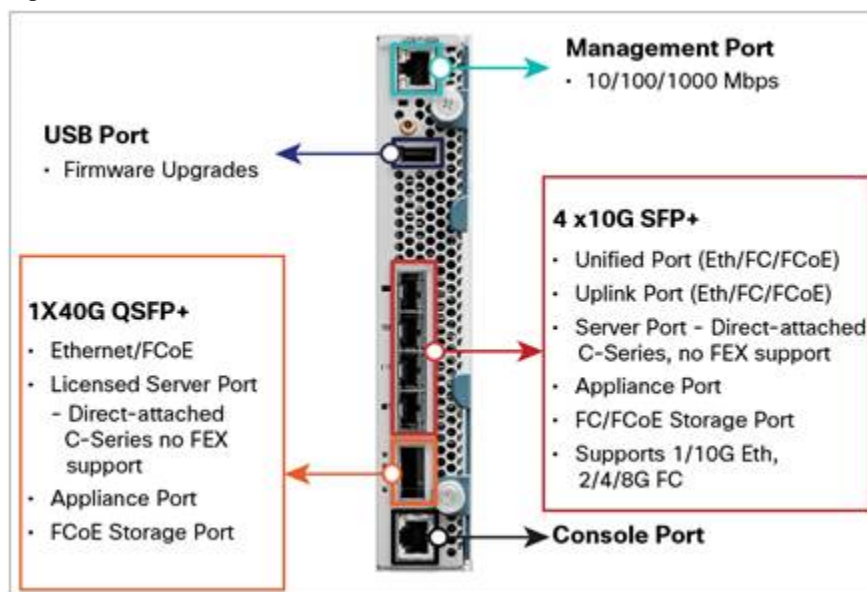
## Cisco UCS 6324UP Fabric Interconnect

The Cisco UCS 6324 Fabric Interconnect provides the management, LAN and storage connectivity for the Cisco UCS 5108 Blade Server Chassis and direct-connect rack-mount servers. It provides the same full-featured Cisco UCS management capabilities and XML API as the full-scale Cisco UCS solution in addition to integrating with Cisco UCS Central Software and Cisco UCS Director (Figure 2).

From a networking perspective, the Cisco UCS 6324 Fabric Interconnect uses a cut-through architecture supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports with switching capacity of up to 500Gbps, independent of packet size and enabled services. Sixteen 10Gbps links connect to the servers, providing a 20Gbps link from each Cisco UCS 6324 Fabric Interconnect to each server. The product family supports Cisco® low-latency, lossless 10 Gigabit Ethernet[1] unified network fabric capabilities that increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the fabric interconnect. Significant TCO savings come from Fibre Channel over Ethernet (FCoE)-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

The Cisco UCS 6324 Fabric Interconnect (Figure 2) is a 10 Gigabit Ethernet, FCoE and Fibre Channel switch offering up to 500Gbps throughput and up to four unified ports and one scalability port.

Figure 2 Cisco UCS 6234 Fabric Interconnect



## Cisco UCS B200 M3 Blade Server

Building on the success of the Cisco UCS B200 M2 Blade Servers, the enterprise-class Cisco UCS B200 M3 further extends the capabilities of the Cisco Unified Computing System portfolio in a half-blade form factor. The Cisco UCS B200 M3 Server harnesses the power of the Intel® Xeon® E5-2600 v2 processor product family, up to 786 GB of RAM, two hard drives, and up to 8 x 10GE to deliver exceptional levels of performance, memory expandability, and I/O throughput for nearly all applications.

Figure 3 Cisco UCS B200 M3 Blade Server



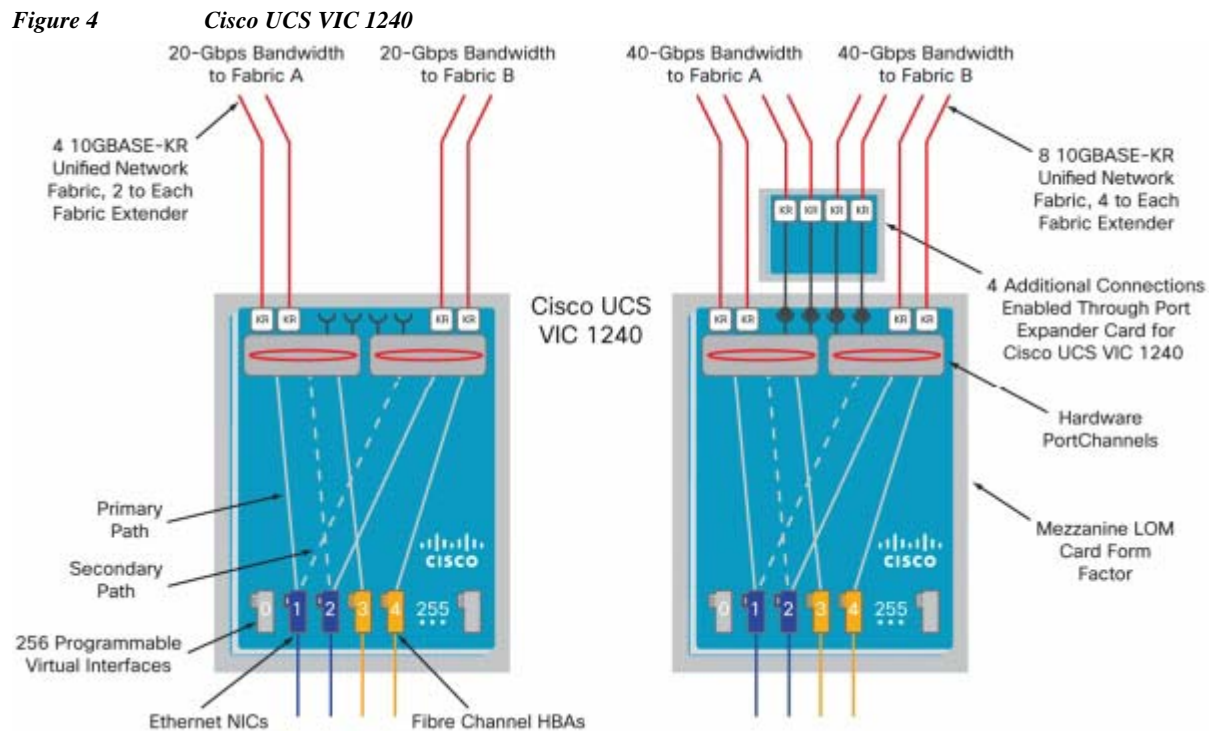
## Cisco I/O Adapters

The Cisco UCS blade server has various Converged Network Adapters (CNA) options. The UCS VIC 1240 Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

Cisco UCS VIC 1240 is a 4-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional Port Expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

The Cisco UCS VIC 1240 enables a policy-based, stateless, agile server infrastructure that can present up to 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1240 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Figure 4



## UCS Differentiators

Cisco Unified Compute System is revolutionizing the way servers are managed in data center. Following are the unique differentiators of UCS and UCS-Manager.

1. **Embedded management:** In UCS, the servers are managed by the **embedded firmware** in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
2. **Unified fabric:** The new UCS Fabric Interconnect 6324 supports unified fabric ports, which operates LAN, SAN and management traffic on the same chassis for both blade server or rack server deployment.
3. **Auto Discovery:** By simply inserting the blade server in the chassis, the discovery and inventory of compute resource occurs automatically without any management intervention. Combination of unified fabric and auto-discovery enables **wire-once architecture** of UCS, where compute capability of UCS can extend easily while keeping the existing external connectivity to LAN, SAN and management networks.
4. **Policy based resource classification:** Once a compute resource is discovered by UCSM, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing. This CVD shows the policy based resource classification of UCSM.



5. Combined Rack and Blade server management: UCSM can manage B-series blade servers and C-series rack server under the same UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic. In this CVD, we are showcasing only B series servers to demonstrate stateless and form factor independent computing work load.
6. Model based management architecture: UCSM architecture and management database is model based and data driven. Open, standard based XML API is provided to operate on the management model. This enables easy and scalable integration of UCSM with other management system, such as VMware vCloud director, Microsoft system center, and Citrix Cloud Platform.
7. Policies, Pools, Templates: Management approach in UCSM is based on defining policies, pools and templates, instead of cluttered configuration, which enables simple, loosely coupled, data driven approach in managing compute, network and storage resources.
8. Loose referential integrity: In UCSM, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each-other. This provides great flexibilities where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.
9. Policy resolution: In UCSM, a tree structure of organizational unit hierarchy can be created that mimics the real life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy. A policy referring to other policy by name is resolved in the organization hierarchy with closest policy match. If no policy with specific name is found in the hierarchy till root organization, then special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibilities to owners of different organizations.
10. Service profiles and stateless computing: Service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. **Stateless computing** enables procurement of a server within minutes, which used to take days in legacy server management systems.
11. Built-in multi-tenancy support: Combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and service profile based approach to compute resources makes UCSM inherently friendly to multi-tenant environment typically observed in private and public clouds.
12. Virtualization aware network: VM-FEX technology makes access layer of network aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-profiles defined by the network administrator’s team. VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
13. Simplified QoS: Even though fibre-channel and Ethernet are converged in UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in UCSM by representing all system classes in one GUI panel.

## VMware vSphere 5.5

VMware vSphere 5.5 is a next-generation virtualization solution from VMware which builds upon ESXi 5.1 and provides greater levels of scalability, security, and availability to virtualized environments. vSphere 5.5 offers improvements in performance and utilization of CPU, memory, and I/O. It also offers

users the option to assign up to thirty two virtual CPU to a virtual machine—giving system administrators more flexibility in their virtual server farms as processor-intensive workloads continue to increase.

The vSphere 5.5 provides the VMware vCenter Server that allows system administrators to manage their ESXi hosts and virtual machines on a centralized management platform. With the Cisco Fabric Interconnects Switch integrated into the vCenter Server, deploying and administering virtual machines is similar to deploying and administering physical servers. Network administrators can continue to own the responsibility for configuring and monitoring network resources for virtualized servers as they did with physical servers. System administrators can continue to “plug-in” their virtual machines into network ports that have Layer 2 configurations, port access and security policies, monitoring features, etc., that have been pre-defined by the network administrators; in the same way they would plug in their physical servers to a previously-configured access switch. In this virtualized environment, the system administrator has the added benefit of the network port configuration/policies moving with the virtual machine if it is ever migrated to different server hardware.

VMware vSphere 5.5 includes an expansive list of new and improved features that enhance performance, reliability, availability, and recovery of virtualized environments. Of those features, several have significant impacts upon VSPEX Private Cloud deployments, including:

- Expanded maximum memory and CPU limits for ESX hosts. Logical and virtual CPU counts have doubled in this version, as have non-uniform memory access (NUMA) node counts and maximum memory. This means host servers can support larger workloads.
- 62TB Virtual Machine Disk (VMDK) file support including Raw Device Mapping (RDM). Datastores can hold more data from more virtual machines, which simplifies storage management and leverages larger capacity NL-SAS drives.
- Enhanced Single-Root Input/Output (I/O) Virtualization (SR-IOV) support that simplifies configuration via workflows, and surfaces more properties into the virtual functions.
- 16 Gb end-to-end support for FC environments.
- Enhanced Link Aggregation Control Protocol (LACP) functions offering additional hash algorithms and up to 64 Link Aggregation Groups (LAGs).
- vSphere Data Protection (VDP), which can now replicate backup data directly to EMC Avamar.
- Virtual Machine File System (VMFS) heap improvements, which reduce memory requirements while allowing access to the full 64TB VMFS address space.

## EMC Storage Technologies and Benefits

The Storage layer is also a key component of any cloud infrastructure solution that serves data generated by applications and operating system in the data center storage processing systems. In this VSPEX solution, EMC VNXe series arrays provide features and performance to enable and enhance any virtualization environment. This increases storage efficiency, management flexibility, and reduces total cost of ownership.

The EMC VNXe series is optimized for virtual applications and delivers industry-leading innovation and enterprise capabilities for file and block storage in a scalable easy-to-use solution. This next-generation storage platform combines powerful and flexible hardware with advanced efficiency, management, and protection software to meet the demanding needs of today’s enterprises.

Intel Xeon Processors power the VNXe series for intelligent storage that automatically and efficiently scales in performance, while ensuring data integrity and security. The VNXe series is designed to meet the high performance, high-scalability requirements of small and midsize enterprises.

[Table 1](#) shows the customer benefits that are provided by the VNXe series.

**Table 1**      **VNXe Customer Benefits**

Features	Benefits
Next-generation unified storage, optimized for virtualization applications	Tight integration with VMWare allows for advanced array features and centralized management
Capacity optimization features including compression, deduplication, thin provisioning and application-consistent copies	Reduced storage costs, more efficient use of resources and easier recovery of applications
High-availability, designed to deliver five 9s availability	Higher levels of uptime and reduced outage risk
Automated tiering with FAST VP and FAST Cache that can be optimized for the highest system performance and lowest storage cost simultaneously	More efficient use of storage resources without complicated planning and configuration
Simplified management with EMC Unisphere with a single management interface for all NAS and SAN needs	Reduced management overhead and toolsets required to manage environment

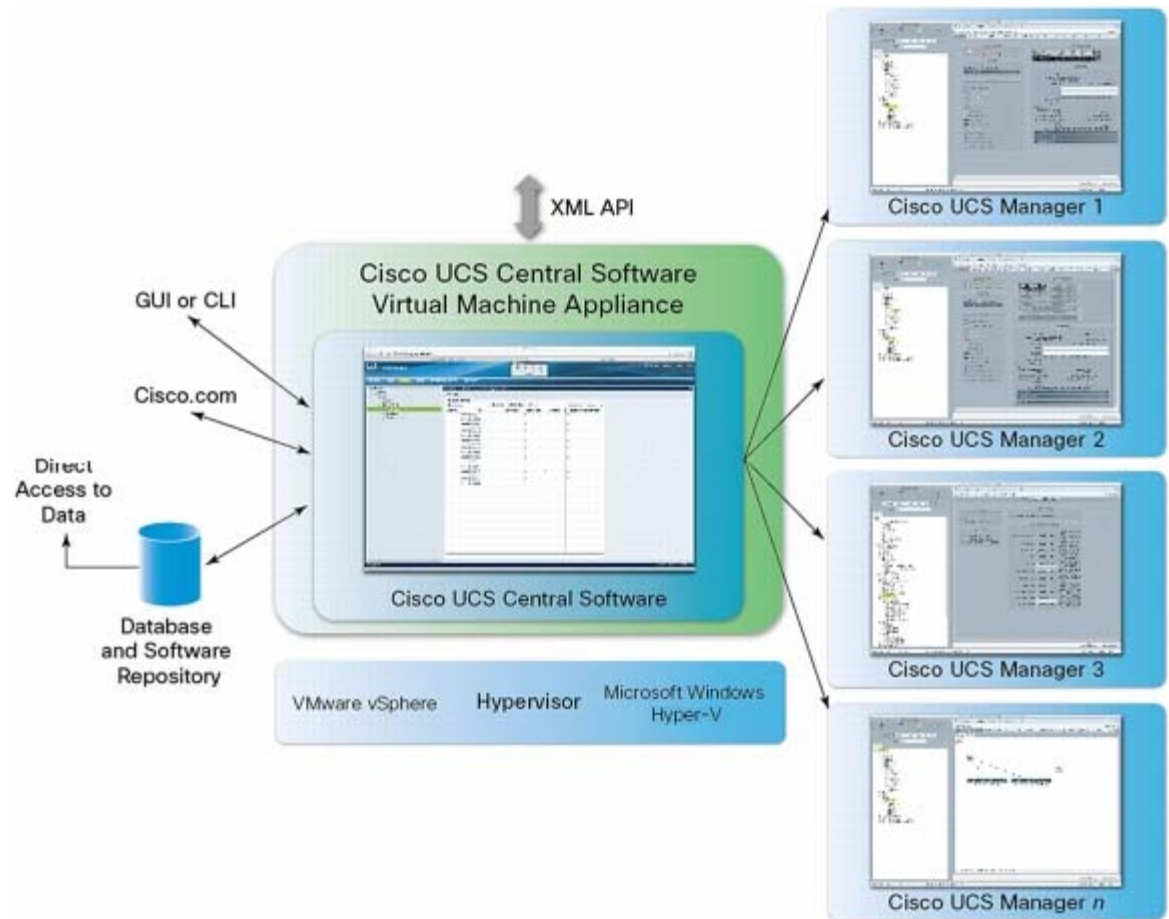
Various software suites and packs are available for the VNXe series. These provide multiple features for enhanced protections and performance. They include the following:

- **Fast Suite:** Automatically optimizes for the highest system performance and the lowest storage cost simultaneously.
- **Security and Compliance Suite:** Keeps data safe from changes, deletions, and malicious activity.

## Cisco UCS Central

For Cisco UCS customers managing growth within a single data center, growth across multiple sites, or both, Cisco UCS Central Software centrally manages multiple Cisco UCS domains using the same concepts that Cisco UCS Manager uses to support a single domain (Figure 5). Cisco UCS Central Software manages global resources (including identifiers and policies) that can be consumed within individual Cisco UCS Manager instances. It can delegate the application of policies (embodied in global service profiles) to individual domains, where Cisco UCS Manager puts the policies into effect. Cisco UCS Central software manages multiple, globally distributed Cisco UCS domains with thousands of servers from a single pane. Every instance of Cisco UCS Manager and all of the components managed by it form a domain. Cisco UCS Central integrates with Cisco UCS Manager, and utilizes it to provide global configuration capabilities for pools, policies, and firmware.

**Figure 5** Cisco UCS Central Software Architecture



Cisco UCS Central Software makes global policy and compliance easier. When Cisco UCS domains are registered with Cisco UCS Central Software, they can be configured to automatically inherit global identifiers and policies that are centrally defined and managed. Making identifiers such as universal user IDs (UUIDs), MAC addresses, and worldwide names (WWNs) global resources allows every server worldwide to be configured uniquely so that identifier conflicts are automatically avoided. Globally defined policies take this concept significantly further: by defining and enforcing server identity, configuration, and connectivity policies centrally, standards compliance is essentially ensured. The system simply will not configure a server in a way that is inconsistent with standards, and configuration drift and an entire class of errors that can cause downtime are avoided.

## UCS Manager (Mini) Management with UCS Central

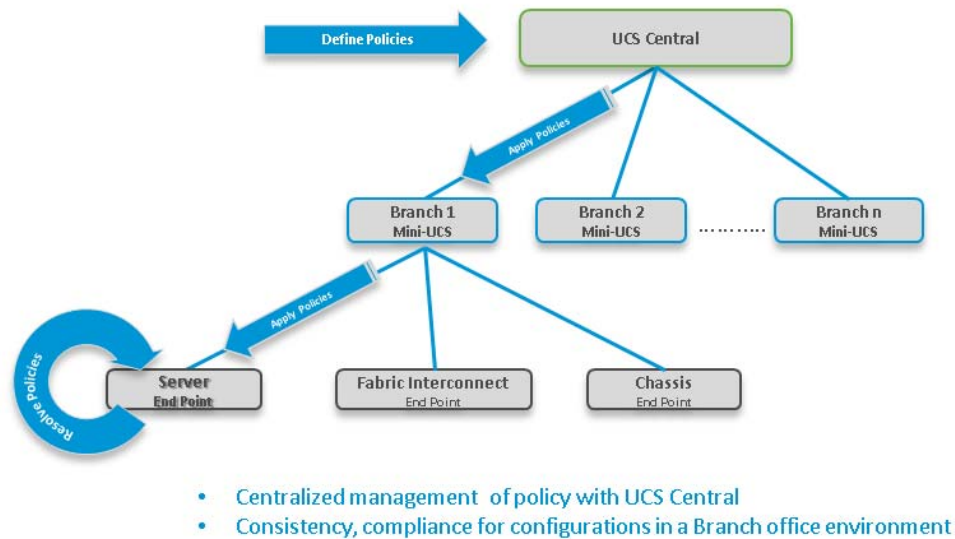
Cisco UCS Central Software is designed and operates similar to Cisco UCS Manager in that policies and configuration definitions, which make up a Cisco UCS service profile, can be created at a central location and then applied to the endpoint recipient, where they are resolved. With Cisco UCS Manager, the endpoint recipients are the Cisco UCS infrastructure (servers, network, etc.).

For Cisco UCS Central Software, the recipients are individual Cisco UCS Manager instances that have been registered with Cisco UCS Central Software. With Cisco UCS Central Software, global Cisco UCS service profiles are defined centrally and are passed to Cisco UCS Manager instances according to the way they are registered with Cisco UCS Central Software (Figure 6).

Figure 6 UCS Manager (Mini) Management with UCS Central

## UCS-Mini Management with UCS Central

### Policy Based Management



© 2010 Cisco and/or its affiliates. All rights reserved.

Cisco Confidential - Do Not Distribute 11

## Architectural overview

This CVD focuses on the architecture for EMC VSPEX for VMware private cloud, targeted for mid-market segment, using EMC VNXe storage arrays. This architecture uses UCS 3.0(1c) with combined Cisco UCS B200M3 servers with VNXe3200 directly attached to UCS 6324 fabric interconnect. VMware vSphere 5.5 is used as server virtualization architecture

**Figure 7** Illustrates the Architectural Overview of VSPEX Remote Office and Branch Office Solution

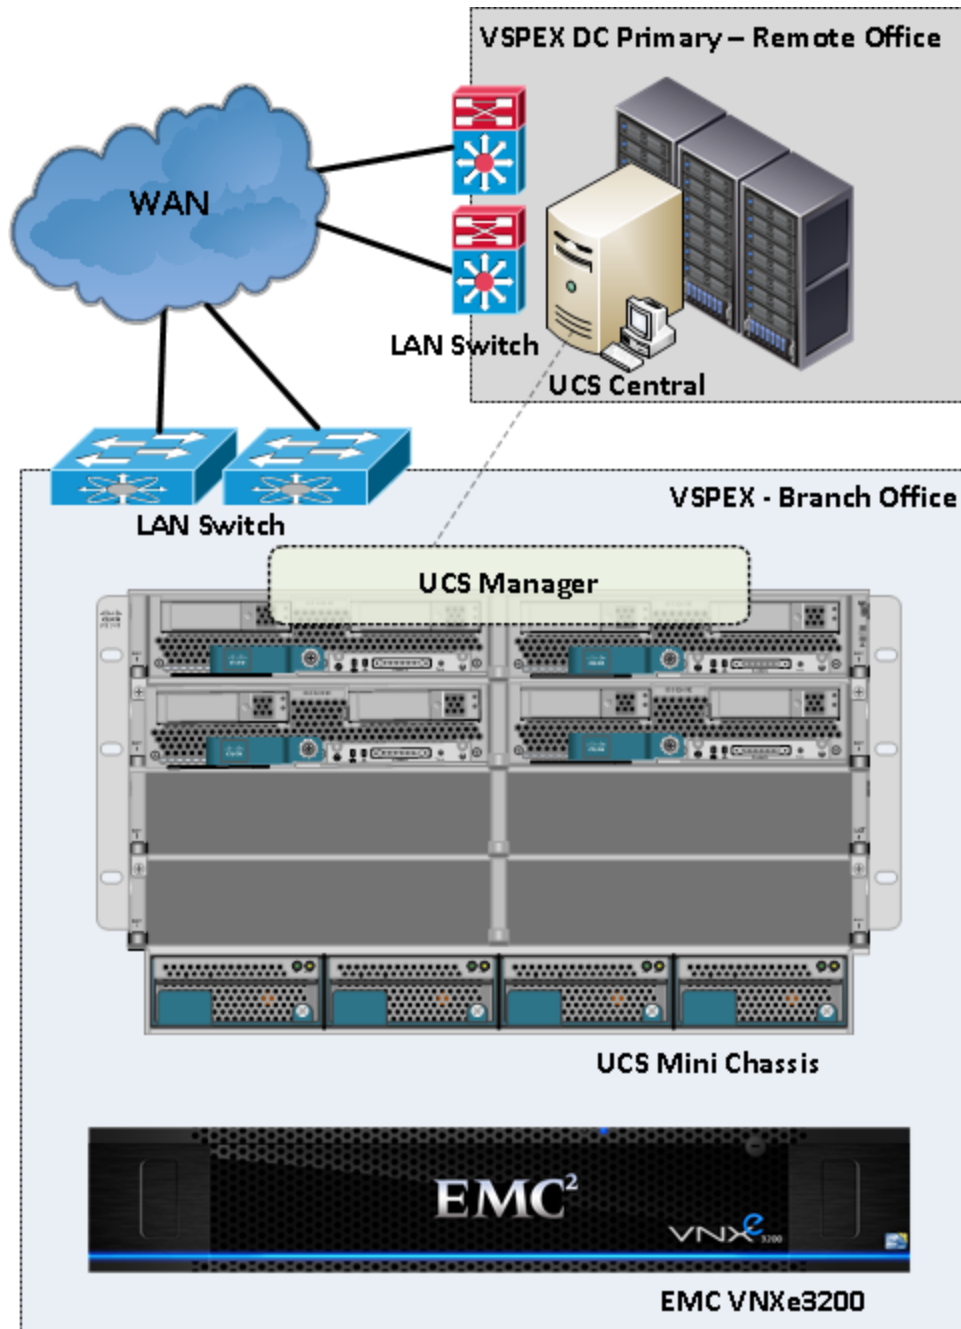


Table 2 lists the various hardware and software components which occupies different tiers of the Cisco Branch office solution for EMC VSPEX VMware architectures under test.

**Table 2** *Hardware and Software Components of the VMware Architecture*

Vendor	Name	Version	Description
Cisco	Cisco UCS Manager	3.0(1c)	UCS Manager
Cisco	Cisco UCS Fabric Interconnect 6324UP	5.0(3)N2(3.01c)	UCS Fabric Interconnects
Cisco	Cisco UCS 5108 AC2 Chassis	N/A	UCS Blade server chassis
Cisco	Cisco UCS B200 M3 Server	2.2.1a	Cisco B200 M3 Blade Servers
Cisco	Cisco UCS VIC 1240	3.0(1c)	Cisco VIC 1240 adapters
EMC	EMC VNXe 3200		VNXe Storage
VMware	VMware ESXi 5.5	5.5 build XXXXX	Hypervisor
VMware	VMware vCenter Server	5.5 build XXXXX	VMware Management
Microsoft	Microsoft Windows Server 2012 R2	2012 R2 SP1 data center	Operating System to host vCenter server & Operating System for VSPEX Virtual Machine
Microsoft	Microsoft SQL Server 2008 R2	2008 R2 Enterprise	SQL Database Server Enterprise Edition for vCenter Server

[Table 3](#) outlines the Cisco UCS B200 M3 Server configuration of this architecture. [Table 3](#) shows the configuration per server basis.

**Table 3** *Cisco UCS B200 M3 Server Configuration Details*

Components	Capacity
Memory (RAM)	256 GB (16X16 GB DIMM)
Processor	2XIntel® Xeon® E5-2640 v2 CPUs, 2GHz, 8cores, 16 thread
Local Storage	Cisco UCS RAID SAS 2008M-8i Mezzanine Card

This SMB architecture contains the infrastructure server on the same chassis in a vSphere Environment manages the virtual machine hosting vCenter server and Windows Active Directory / DNS server are present.

- This design does not dictate or require any specific layout of infrastructure network. The vCenter server, Microsoft AD server and Microsoft SQL server are hosted on infrastructure Blade Server. However, the design does require that certain VLANs are accessible from both the infrastructure blade server and the VSPEX servers.

- ESXi 5.5 is used as hypervisor operating system on each server and is installed on SAN LUNs in both architectures. All the VSPEX Virtual Machines storage is accessed thru NFS protocols. Typical load is 32 reference virtual machines per server.

## Memory Configuration Guidelines

This section provides guidelines for allocating memory to virtual machines. The guidelines outlined here take into account vSphere memory overhead and the virtual machine memory settings.

### ESXi/ESXi Memory Management Concepts

vSphere virtualizes guest physical memory by adding an extra level of address translation. Shadow page tables make it possible to provide this additional translation with little or no overhead. Managing memory in the hypervisor enables the following:

- Memory sharing across virtual machines that have similar data (that is, same guest operating systems).
- Memory overcommitment, which means allocating more memory to virtual machines than is physically available on the ESX/ESXi host.
- A memory balloon technique whereby virtual machines that do not need all the memory they were allocated will give memory to virtual machines that require additional allocated memory.


For more information about vSphere memory management concepts, see the VMware vSphere Resource Management Guide at:

[http://www.vmware.com/files/pdf/perf-vsphere-memory\\_management.pdf](http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf)

### Virtual Machine Memory Concepts

Figure 8 illustrates the use of memory settings parameters in the virtual machine.

**Figure 8** Virtual Machine Memory Settings



The vSphere memory settings for a virtual machine include the following parameters:

- Configured memory – Memory size of virtual machine assigned at creation.
- Touched memory – Memory actually used by the virtual machine. vSphere allocates only guest operating system memory on demand.
- Swappable – Virtual machine memory that can be reclaimed by the balloon driver or by vSphere swapping. Ballooning occurs before vSphere swapping. If this memory is in use by the virtual machine (that is, touched and in use), the balloon driver causes the guest operating system to swap. Also, this value is the size of the per-virtual machine swap file that is created on the VMware Virtual



Machine File System (VMFS) file system (VSWP file). If the balloon driver is unable to reclaim memory quickly enough, or is disabled or not installed, vSphere forcibly reclaims memory from the virtual machine using the VMkernel swap file.

## Allocating Memory to Virtual Machines

The proper sizing of memory for a virtual machine in VSPEX architectures is based on many factors. With the number of application services and use cases available determining a suitable configuration for an environment requires creating a baseline configuration, testing, and making adjustments, as discussed later in this paper. [Table 4](#) outlines the resources used by a single virtual machine:

**Table 4**      *Virtual Memory Details*

Characteristics	Value
Virtual processor per virtual machine (vCPU)	1
RAM per virtual machine	2GB
Available storage capacity per virtual machine	100GB
I/O operations per second (IOPS) per VM	25
I/O pattern	Random
I/O read/write ratio	2:1

The following are descriptions of recommended best practices:

- Account for memory overhead – Virtual machines require memory beyond the amount allocated, and this memory overhead is per-virtual machine. Memory overhead includes space reserved for virtual machine devices, depending on applications and internal data structures. The amount of overhead required depends on the number of vCPUs, configured memory, and whether the guest operating system is 32-bit or 64-bit. As an example, a running virtual machine with one virtual CPU and two gigabytes of memory may consume about 100 megabytes of memory overhead, where a virtual machine with two virtual CPUs and 32 gigabytes of memory may consume approximately 500 megabytes of memory overhead. This memory overhead is in addition to the memory allocated to the virtual machine and must be available on the ESXi host.
- "Right-size" memory allocations – Over-allocating memory to virtual machines can waste memory unnecessarily, but it can also increase the amount of memory overhead required to run the virtual machine, thus reducing the overall memory available for other virtual machines. Fine-tuning the memory for a virtual machine is done easily and quickly by adjusting the virtual machine properties. In most cases, hot-adding of memory is supported and can provide instant access to the additional memory if needed.
- Intelligently overcommit – Memory management features in vSphere allow for overcommitment of physical resources without severely impacting performance. Many workloads can participate in this type of resource sharing while continuing to provide the responsiveness users require of the application. When looking to scale beyond the underlying physical resources, consider the following:
- Establish a baseline before overcommitting – Note the performance characteristics of the application before and after. Some applications are consistent in how they utilize resources and may not perform as expected when vSphere memory management techniques take control. Others, such as Web servers, have periods where resources can be reclaimed and are perfect candidates for higher levels of consolidation.

- Use the default balloon driver settings – The balloon driver is installed as part of the VMware Tools suite and is used by ESXi/ESX if physical memory comes under contention. Performance tests show that the balloon driver allows ESXi/ESX to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESXi/ESX to use host-swapping to make up for the lack of available physical memory which adversely affects performance.
- Set a memory reservation for virtual machines that require dedicated resources – Virtual machines running Search or SQL services consume more memory resources than other application and Web front-end virtual machines. In these cases, memory reservations can guarantee that those services have the resources they require while still allowing high consolidation of other virtual machines.

## Storage Guidelines

VSPEX architecture for VMware virtual machines for mid-market segment uses FC for boot or NFS to store Virtual machine data in a VNXe Storage Array. vSphere provides many features that take advantage of EMC storage technologies such as auto discovery of storage resources and ESXi hosts in vCenter and VNXe respectively. Features such as VMware vMotion, VMware HA, and VMware Distributed Resource Scheduler (DRS) use these storage technologies to provide high availability, resource balancing, and uninterrupted workload migration.

## Storage Protocol Capabilities

VMware vSphere provides vSphere and storage administrators with the flexibility to use the storage protocol that meets the requirements of the business. This can be a single protocol datacenter wide, such as NFS, or multiple protocols for tiered scenarios such as using Fibre Channel for high-throughput storage pools and NFS for high-capacity storage pools. For more information, see the VMware whitepaper Comparison of Storage Protocol Performance in VMware vSphere 5 at:

[http://www.vmware.com/files/pdf/perf\\_vsphere\\_storage\\_protocols.pdf](http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf)

## Storage Best Practices

The following are vSphere storage best practices:

- Host multi-pathing – Having a redundant set of paths to the storage area network is critical to protecting the availability of your environment. This redundancy is in the form of dual adapters connected to separate fabric switches.
- Partition alignment – Partition misalignment can lead to severe performance degradation due to I/O operations having to cross track boundaries. Partition alignment is important both at the VMFS level as well as within the guest operating system. Use the vSphere Client when creating VMFS datastores to be sure they are created aligned. When formatting volumes within the guest, Windows 2008 aligns NTFS partitions on a 1024KB offset by default.
- Use shared storage – In a vSphere environment, many of the features that provide the flexibility in management and operational agility come from the use of shared storage. Features such as VMware HA, DRS, and vMotion take advantage of the ability to migrate workloads from one host to another host while reducing or eliminating the downtime required to do so.
- Calculate your total virtual machine size requirements – Each virtual machine requires more space than that used by its virtual disks. Consider a virtual machine with a 20GB OS virtual disk and 16GB of memory allocated. This virtual machine will require 20GB for the virtual disk, 16GB for the virtual machine swap file (size of allocated memory), and 100MB for log files (total virtual disk size + configured memory + 100MB) or 36.1GB total.

- Understand I/O Requirements – Under-provisioned storage can significantly slow responsiveness and performance for applications. In a multitier application, you can expect each tier of application to have different I/O requirements. As a general recommendation, pay close attention to the amount of virtual machine disk files hosted on a single VMFS volume. Over-subscription of the I/O resources can go unnoticed at first and slowly begin to degrade performance if not monitored proactively.

## Virtual Networking

Each B200 M3 blade server has one physical Cisco VIC adapter carved out with six 10 GE vNIC interfaces evenly distributed to fabric A and fabric B for high availability and also presents two virtual Host Bus Adapters (vHBAs) for the booting ESXi from SAN, one vHBA per fabric path. The MAC addresses to these vNICs are assigned using MAC address pool defined on the UCSM. The vNICs are used in active-active configuration for load-balancing and high-availability. Following are vSphere networking best practices implemented in this architecture:

- “Separate virtual machine and infrastructure traffic – Keep virtual machine and VMkernel or service console traffic separate.” This is achieved by having two vSwitches per hypervisor:
  - vSwitch (default) – used for management and vMotion traffic
  - vSwitch1 – used for Virtual Machine data traffic
- “Use NIC Teaming – Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches.” This is achieved by using two vNICs per vSwitch, each going to different Fabric Interconnects. Teaming provides redundancy against NIC failure and FI fabric failures.
- “Jumbo MTU for vMotion and Storage traffic” – this best practice is implemented in the architecture by configuring jumbo MTU end-to-end.

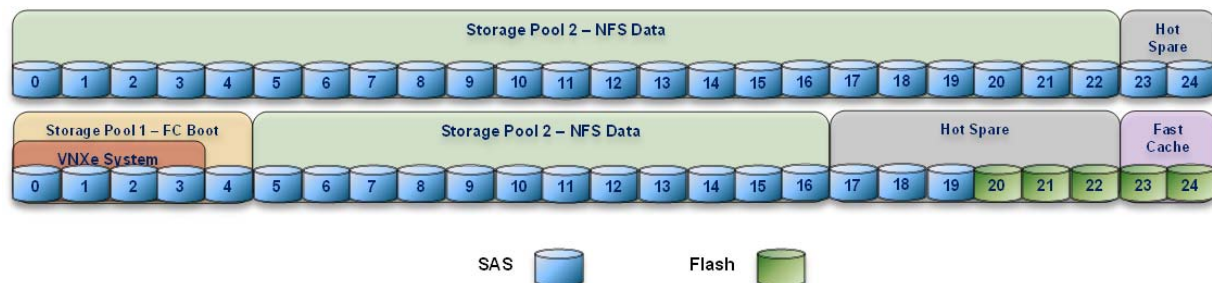
## VSPEX VMware Storage Virtualization

### Storage Layout

The architecture diagram in this section shows the physical disk layout. Disk provisioning on the VNXe series is simplified through the use of wizards, so that administrators do not choose which disks belong to a given storage pool. The wizard may choose any available disk of the proper type, regardless of where the disk physically resides in the array.

Figure 9 illustrates storage architecture for 100 virtual machines on VNXe3200 for NFS based storage architecture:

**Figure 9** Storage layout for up to 100 Reference VMs on VNXe3200



The VNX family is designed for “five 9s” availability by using redundant components throughout the array. All of the array components are capable of continued operation in case of hardware failure. The RAID disk configuration on the array provides protection against data loss due to individual disk failures, and the available hot spare drives can be dynamically allocated to replace a failing disk.

## Storage Virtualization

VMFS is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files and VMFS is the default storage system for these files on physical SCSI disks and partitions.

It is preferable to deploy virtual machine files on shared storage to take advantage of VMware VMotion, VMware High Availability™ (HA), and VMware Distributed Resource Scheduler™ (DRS). This is considered a best practice for mission-critical deployments, which are often installed on third-party, shared storage management solutions.

## Service Profile Design

This architecture implements following design steps to truly achieve stateless computing on the servers:

- Service profiles are derived from service profile template for consistency.
- The ESXi host uses following identities in this architecture:
  - Host UUID
  - Mac Addresses: one per each vNIC on the server
  - One WWNN and two WWPN for FC boot

All of these identifiers are defined in their respective identifier pools and the pool names are referred in the service profile template.

- Local disks are NOT used for booting. Boot policy in service profile template suggests host to boot from the storage devices using FC protocol.
- Server pool is defined with automatic qualification policy and criteria. Blade servers are automatically put in the pool as and when they are fully discovered by UCSM. This eliminates the need to manually assign servers to server pool.
- Service profile template is associated to the server pool. This eliminates the need to individually associating service profiles to physical servers.

Given this design and capabilities of UCS and UCSM, a new server can be procured within minutes if the scale needs to be increased or if a server needs to be replaced by different hardware. In case, if a server has physical fault (faulty memory, or PSU or fan, for example), using following steps, a new server can be procured within minutes:

- Put the faulty server in maintenance mode using vCenter. This would move VMs running on fault server to other healthy servers on the cluster.
- Disassociate the service profile from the faulty server and physically remove the server for replacement of faulty hardware (or to completely remove the faulty server).
- Physically install the new server on the chassis. Let the new server be discovered by UCSM.
- Associate the service profile to the newly deployed blade server. This would boot the same ESXi server image from the storage array as what the faulty server was running.

- The new server would assume the role of the old server with all the identifiers intact. You can now end the maintenance mode of the ESXi server in vCenter.

Thus, the architecture achieves the true statelessness of the computing in the data center. If there are enough identifiers in all the id-pools, and if more servers are attached to UCS system in future, more service profiles can be derived from the service profile template and the private cloud infrastructure can be easily expanded. We would demonstrate that blade servers can be added in the same server pool.

## Network High Availability Design

Following figure demonstrates logical layout of the architecture. Following are the key aspects of this solution:

- Cisco UCS B200 M3 servers are used, managed by UCS Manager
- VNICs on fabric A and fabric B are used for NFS based access high-availability

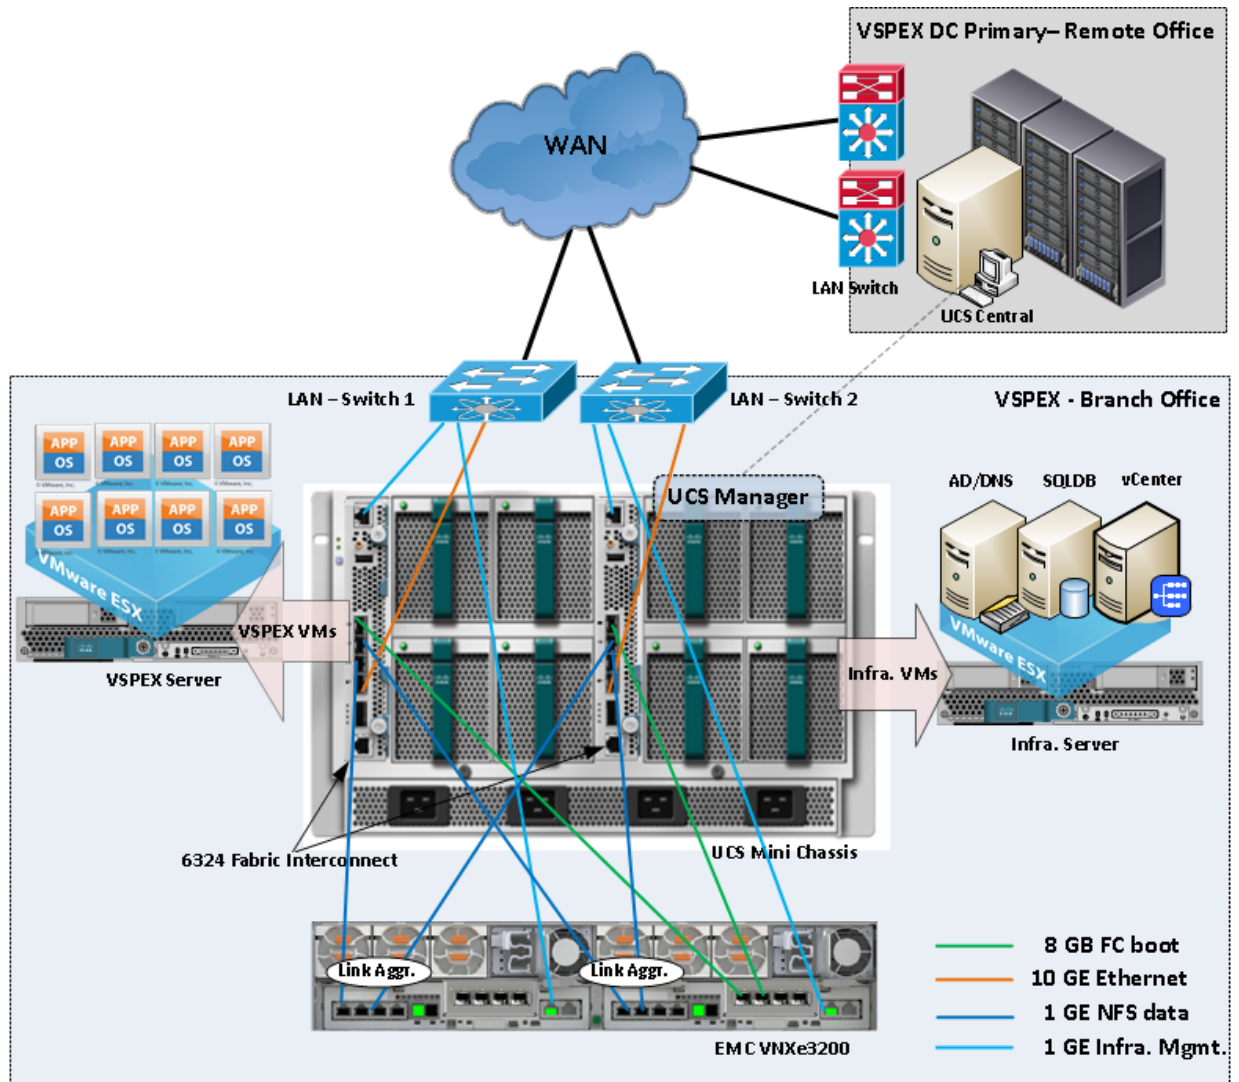


**Note** NFS would require external LAN switch connectivity from Fabric A & Fabric B in case of link failure from FI to storage server.

Storage is made highly available by deploying following practices:

- FC access is for booting each of the ESXi hypervisor images.
- VSPEX VMs are deployed on the NFS Datastore.
- VNXe storage arrays provide two Storage Processors (SPs): SP-A and SP-B for both FC(Block) & NFS (File).

Figure 10 Logical Layout of Branch Office Solution for NFS Based Architecture



## Architecture Design Layout for Optional Topology:

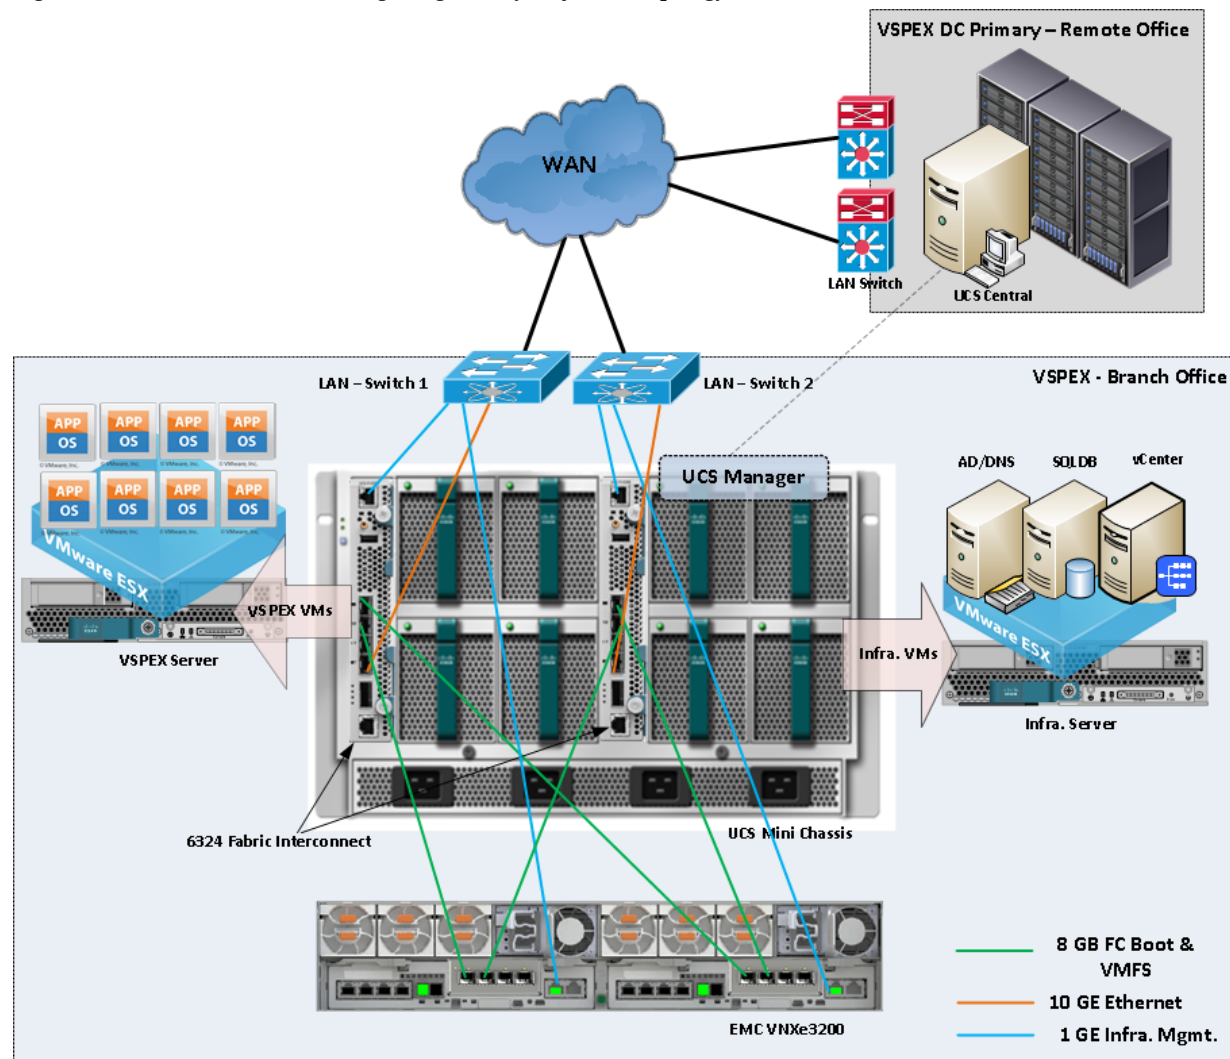
As described in Technology overview section, the UCS Mini chassis and EMC VNXe3200 supports Unified fabric design and provides support for both NAS (NFS, CIFS Protocols) and SAN (FC, iSCSI protocols) based Storage Architecture.

In this solution, we have shown only the NFS variant Architecture, which show cases FC for ESXi Hypervisor boot and NFS for Virtual Machine Datastore.

Due to UCS Mini chassis port limitations and EMC VNXe3200 storage limitations, we cannot showcase all the NAS and SAN protocols in a single architecture design.

In following sections, there are two other possible topologies using the same components:

Figure 11 Architecture Design Logical Layout for FC Topology

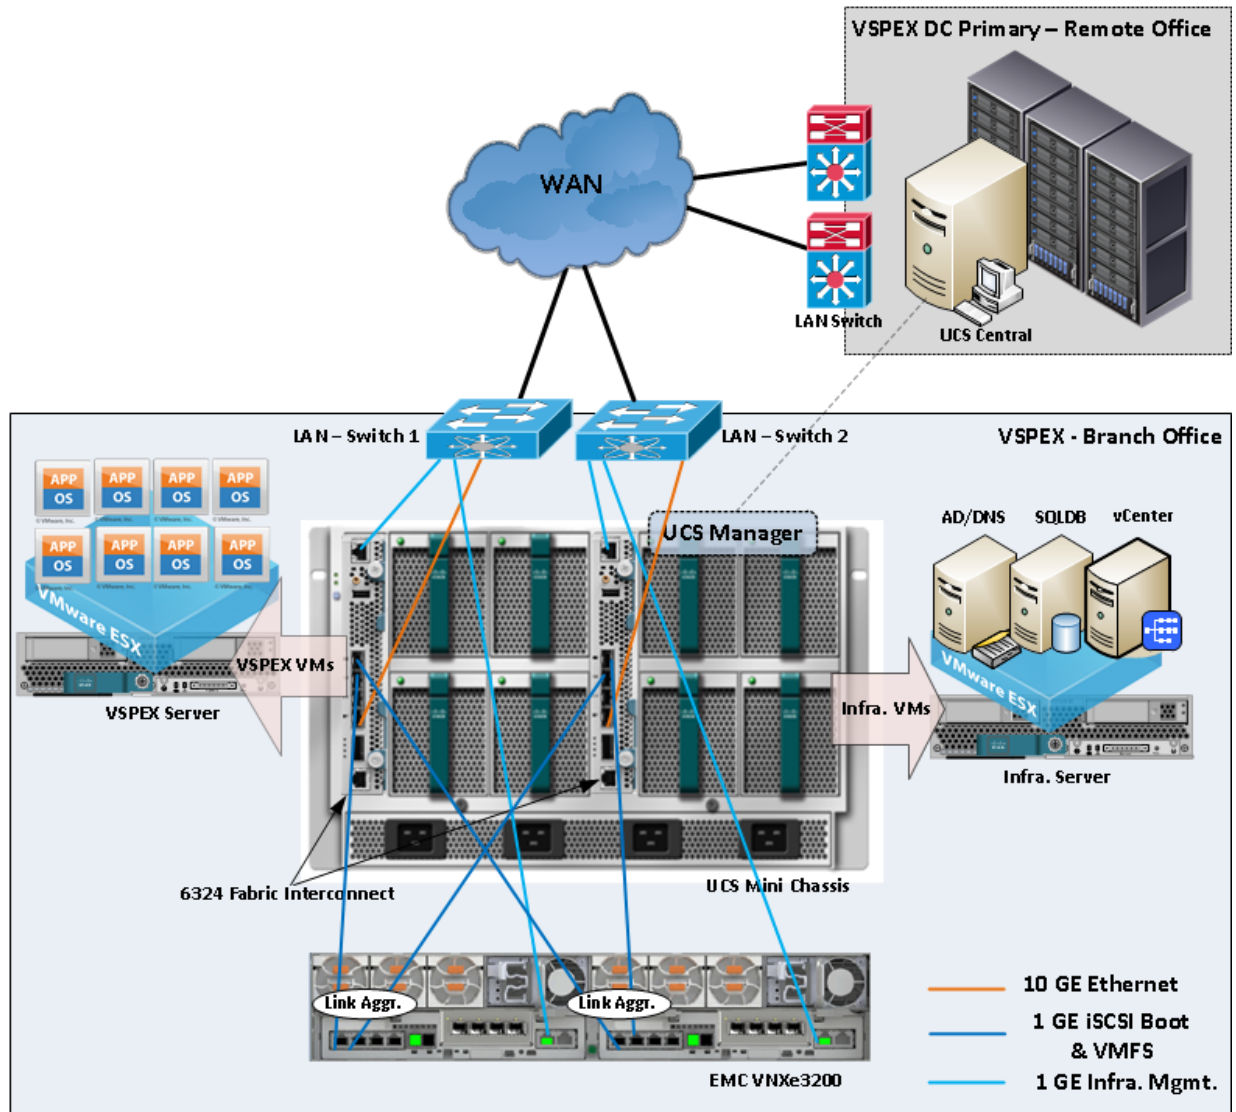


  
Note

FC only topology option will only support VMFS but free a FI port for possible C-Series expansion.



*Figure 12*



## Note

Ethernet only iSCSI topology option will also support VMFS but free a FI port for possible C-series expansion. Also, the failover is not dependent on the upstream LAN switch above but the multi-pathing capabilities of the host OS.

## Jumbo MTU:

Jumbo MTU (size 9000) is used for following two types of traffic in this architecture:

- NFS Storage access
- vMotion traffic



Both of these traffic types are “bulk transfer” traffic, and larger MTU significantly improves the performance. Jumbo MTU must be configured end-to-end to ensure that IP packets are not fragmented by intermediate network nodes. Following is the checklist of end-points where jumbo MTU needs to be configured:

1. Ethernet ports on VNX Storage Processors
2. System QoS classes in UCS Manager
3. vNICS in service profiles
4. VM-Kernel ports used for vMotion and storage access on the ESXi hosts

Next sub section goes in to sizing guidelines of the Cisco Branch solution for EMC VSPEX VMware architectures outlined here.

## Sizing Guideline

In any discussion about virtual infrastructures, it is important to first define a reference workload. Not all servers perform the same tasks, and it is impractical to build a reference that takes into account every possible combination of workload characteristics.

## Defining the Reference Workload

To simplify the discussion, we have defined a representative customer reference workload. By comparing your actual customer usage to this reference workload, you can extrapolate which reference architecture to choose.

For the VSPEX solutions, the reference workload was defined as a single virtual machine. This virtual machine has the following characteristics:

**Table 5** *Virtual Machine Characteristics*

Characteristics	Value
Virtual machine OS	Microsoft Windows Server 2012
Virtual processor per virtual machine (vCPU)	1
RAM per virtual machine	2GB
Available storage capacity per virtual machine	100GB
I/O operations per second (IOPS) per VM	25
I/O pattern	Random
I/O read/write ratio	2:1

This specification for a virtual machine is not intended to represent any specific application. Rather, it represents a single common point of reference to measure other virtual machines.

# VSPEX Configuration Guidelines

The configuration for Cisco solution for EMC VSPEX VMware architectures is divided in to following steps:

1. Pre-deployment tasks
2. Connect network cables
3. Prepare UCS FIs and configure UCSM
4. Configure data stores for ESXi images
5. Install ESXi servers and vCenter infrastructure
6. Install and configure vCenter server
7. Configure storage for VM data stores, install and instantiate VMware VMs from vCenter
8. Test the installation

Next pages go into details of each section mentioned above.

## Pre-deployment tasks

Pre-deployment tasks include procedures that do not directly relate to environment installation and configuration, but whose results will be needed at the time of installation. Examples of pre-deployment tasks are collection of hostnames, IP addresses, VLAN IDs, license keys, installation media, and so on. These tasks should be performed before the customer visit to decrease the time required onsite.

- Gather documents: Gather the related documents listed in the Preface. These are used throughout the text of this document to provide detail on setup procedures and deployment best practices for the various components of the solution.
- Gather tools: Gather the required and optional tools for the deployment. Use [Table 6](#) to confirm that all equipment, software, and appropriate licenses are available before the deployment process.
- Gather data: Collect the customer-specific configuration data for networking, naming, and required accounts. Enter this information into the Customer Configuration Data worksheet for reference during the deployment process.

**Table 6**      *Hardware and Software Requirements*

Requirements	Description	Reference
Hardware	Cisco UCS Mini chassis with 6324 Fabric Interconnect for network and compute infrastructure	See the corresponding product documentation
	Cisco UCS B200 M3 servers to host virtual machines	
	VMware vSphere™ 5.5 server to host virtual infrastructure servers	
	<b>Note</b> This requirement may be covered in the existing infrastructure	
	EMC VNXe storage: Multiprotocol storage array with the required disk layout as per architecture requirements	
Software	VMware ESXi™ 5.5 installation media	See the corresponding product documentation
	VMware vCenter Server 5.5 installation media	
	Microsoft Windows Server 2102 installation media (suggested OS for VMware vCenter)	
License	Microsoft SQL Server 2008 R2 SP1	Consult your corresponding vendor for the license
	<b>Note</b> This requirement may be covered in the existing infrastructure	
	VMware vCenter 5.5 license key	
	VMware ESXi 5.5 license keys	
	Microsoft SQL Server 2008 R2 license key <b>Note</b> This requirement may be covered in the existing infrastructure	

## Customer Configuration Data

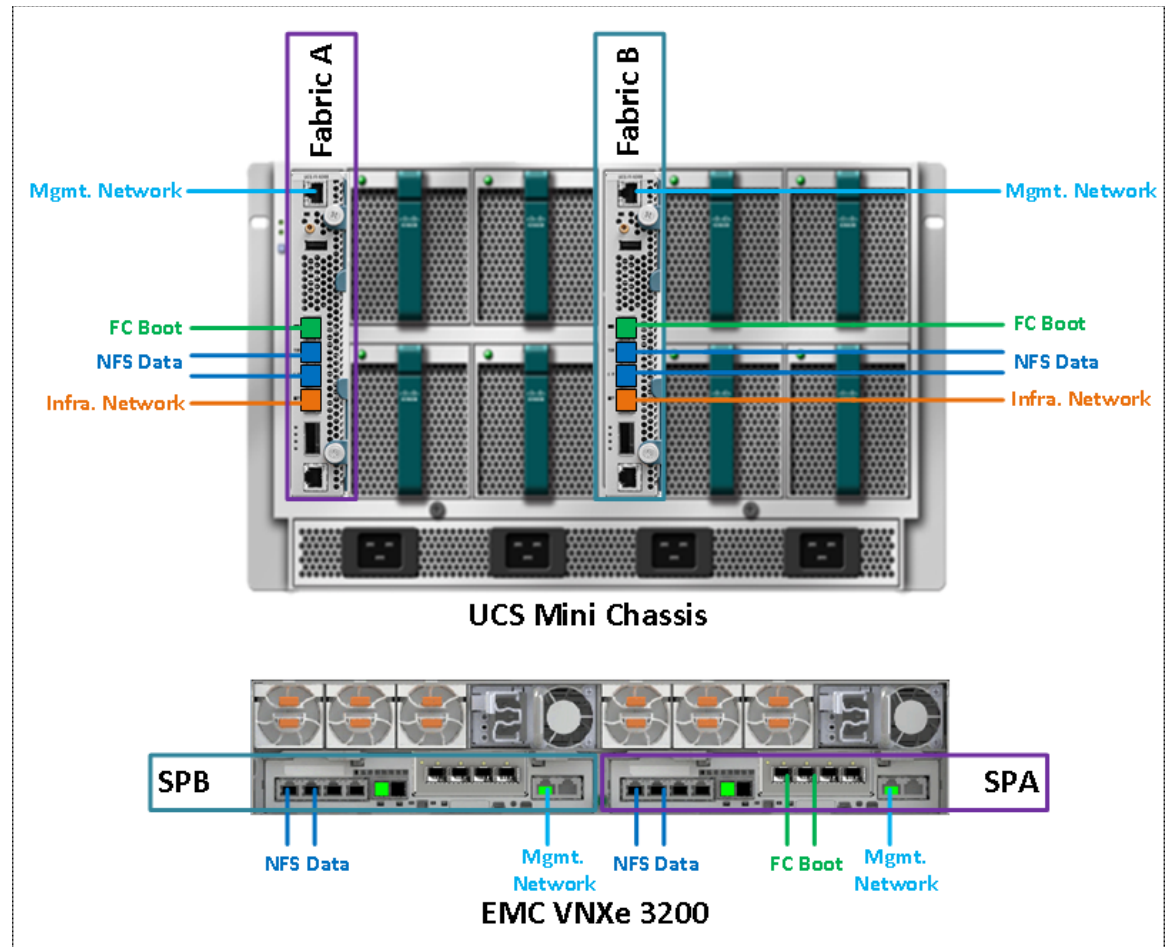
To reduce the onsite time, information such as IP addresses and hostnames should be assembled as part of the planning process. “[Customer Configuration Data Sheet](#)” section on page 174 provides a table to maintain a record of relevant information. This form can be expanded or contracted as required, and information may be added, modified, and recorded as deployment progresses.

Additionally, complete the “[Customer Configuration Data Sheet](#)” section on page 174, available on the EMC online support website, to provide the most comprehensive array-specific information.


## Connect Network Cables

See the Cisco UCS Mini Chassis, 6324 FI, B200M3 server and EMC VNXe3200 storage array configuration guide for detailed information about how to mount the hardware on the rack. Following diagrams show connectivity details for the VSPEX VMware architecture covered in this document.

**Figure 13** *Cabling Diagram for UCS Mini and VNXe Array System - Port Description*



**Figure 14**      *Cabling Diagram for UCS Mini and VNXe Array Systems - Port Number Details*



## Prepare UCS FIs and Configure UCS Manager

Next step is to configure UCS FIs and UCSM. This task can be subdivided in to following segments:

1. Initial configuration of UCS FIs
2. Configuration for server discovery on UCS Manager
3. Upstream / global network configuration
4. Configure identifier pools
5. Configure server pool and qualifying policy
6. Configure service profile template
7. Instantiate service profiles from the service profile template

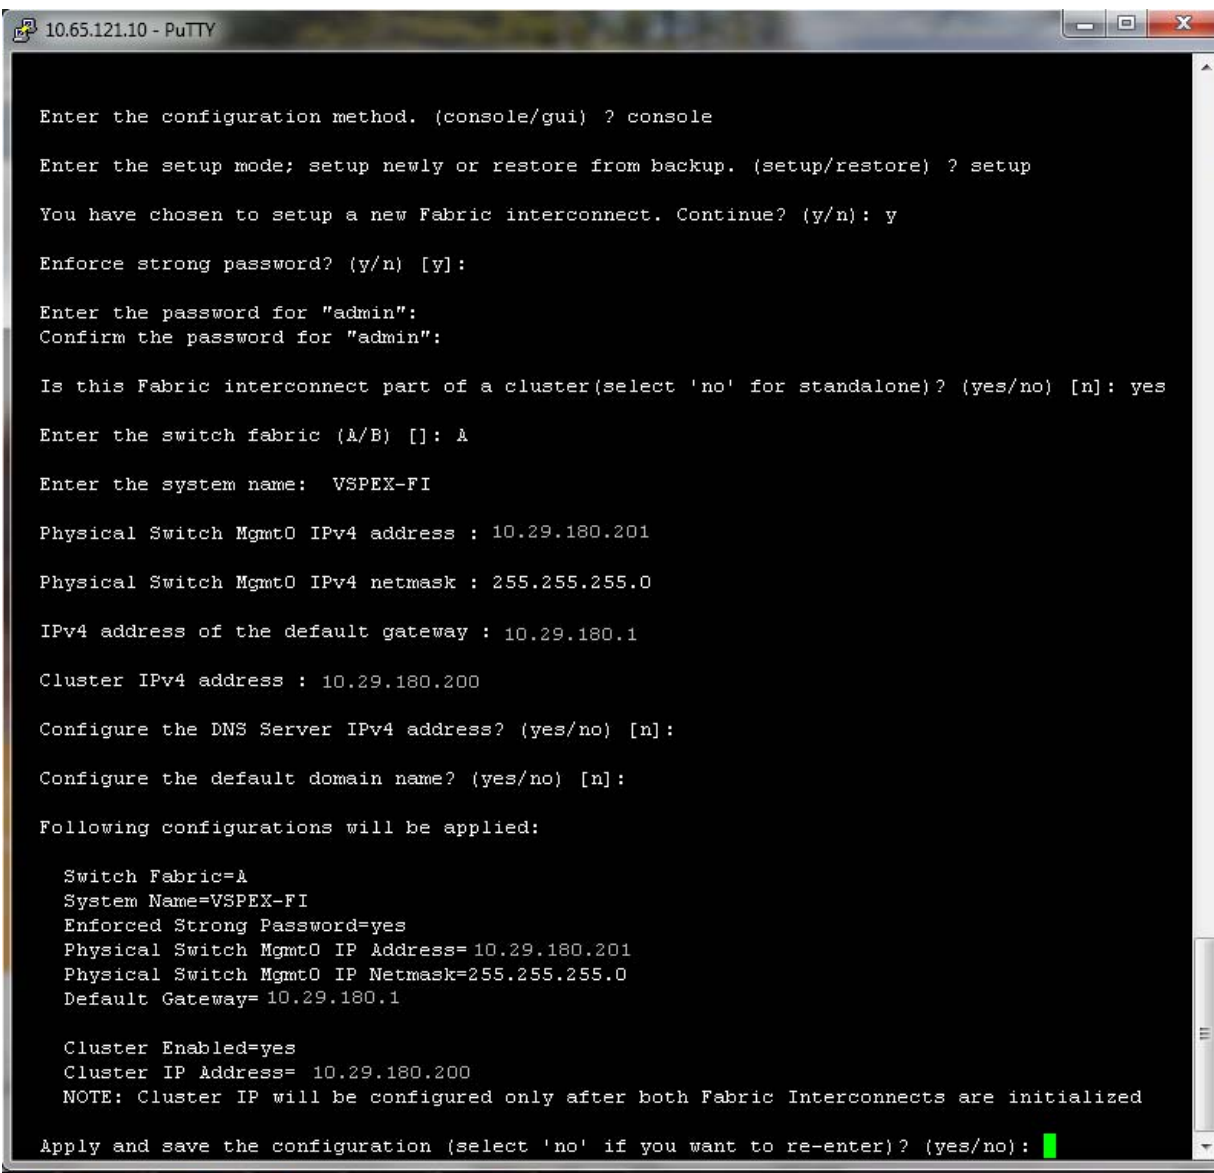
Follow the step-by-step guide to configure UCSM tasks mentioned above.

### Initial Configuration of UCS FIs

At this point of time, mount the UCS 6324 FIs and the UCS chassis with B200 M3 Blade Servers on the rack and make sure that the cable connections are appropriate as suggested in this section. Two 100 Mbps Ethernet cables must be connected between two FIs for management pairing. Two redundant power supplies are provided per FI, it is highly recommended that both are plugged in, ideally drawing power

from two different power strips. Connect mgmt0 interfaces of each FI to the infrastructure network, and put the switch port connected to FI in access mode with access VLAN as management VLAN. Now follow these steps to perform initial configuration of FIs:

1. Attach RJ-45 serial console cable to the first FI, and connect the other end to the serial port of laptop. Configure password for the “admin” account, fabric ID “A”, UCS system name, management IP address, subnet mask and default gateway and cluster IP address (or UCSM Virtual IP address), as the initial configuration script walks you through the configuration as shown in the below image. Save the configuration, which would eventually lead to UCSM CLI login prompt.



```

10.65.121.10 - PuTTY

Enter the configuration method. (console/gui) ? 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 password? (y/n) [y]:

Enter the password for "admin":
Confirm the password for "admin":

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes

Enter the switch fabric (A/B) []: A

Enter the system name: VSPEX-FI

Physical Switch Mgmt0 IPv4 address : 10.29.180.201

Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.29.180.1

Cluster IPv4 address : 10.29.180.200

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

Configure the default domain name? (yes/no) [n]:

Following configurations will be applied:

Switch Fabric=A
System Name=VSPEX-FI
Enforced Strong Password=yes
Physical Switch Mgmt0 IP Address=10.29.180.201
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway= 10.29.180.1

Cluster Enabled=yes
Cluster IP Address= 10.29.180.200
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized

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

```

2. Now disconnect the RJ-45 serial console from the FI that you just configured and attach it to the other FI. Other FI would detect that its peer has been configured, and would prompt you to just join the cluster. The only information you need to provide is the FI specific management IP address, subnet mask and default gateway, as shown in the image below. Save the configuration.

```

Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y

Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IP Address: 10.29.180.202
Peer Fabric interconnect Mgmt0 IP Netmask: 255.255.255.0
Cluster IP address      : 10.29.180.200

Physical Switch Mgmt0 IPv4 address : 10.29.180.201

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

```

3. Once initial configurations on both FIs are completed, you can disconnect the serial console cable. Now, UCSM would be accessible through web interface (<https://<ucsm-virtual-ip>/>) or SSH. Connect to UCSM using SSH, and check the HA status. As there is a common device connected between two FIs, the status would say “HA NOT READY”, but you must see both FI A and FI B in “Up” state as shown in the figure below.

```

VSPEX-FI-A# show cluster state
Cluster Id: 0xec91409a491011e2-0xb7a4547f6eaa1564

A: UP, PRIMARY
B: UP, SUBORDINATE

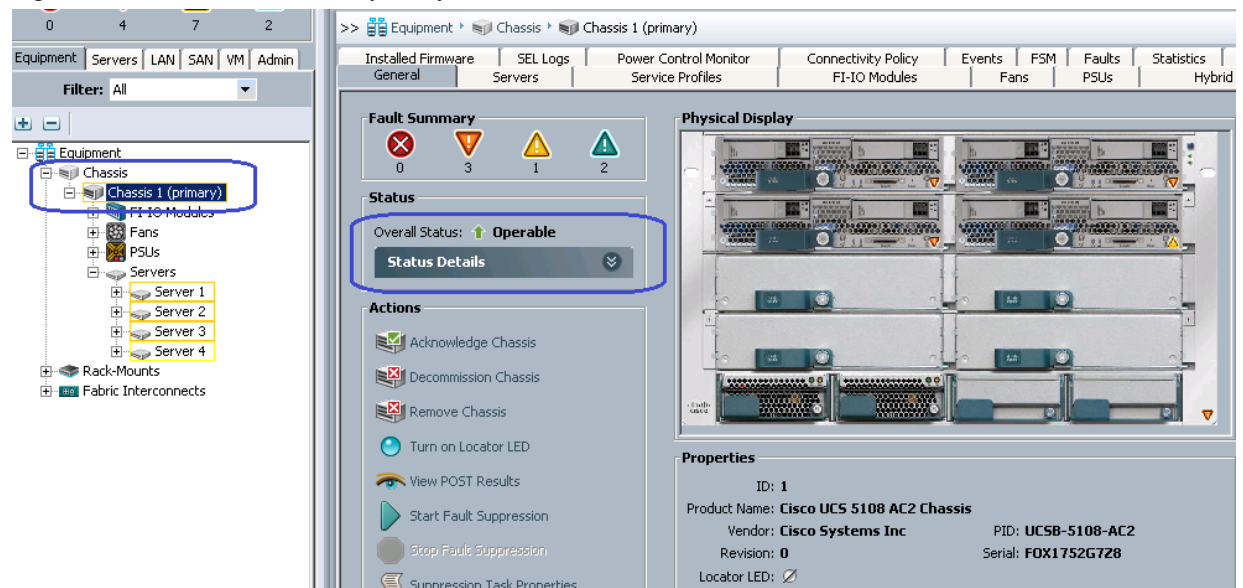
HA NOT READY
No device connected to this Fabric Interconnect
VSPEX-FI-A#

```

## Chassis and Server Discovery on UCS Manager

1. After the FI management IP configuration, launch the UCS Manager via web browser using the Virtual IP. You will see the Chassis discovered Automatically under Equipment tab and it is showing as “operable”.

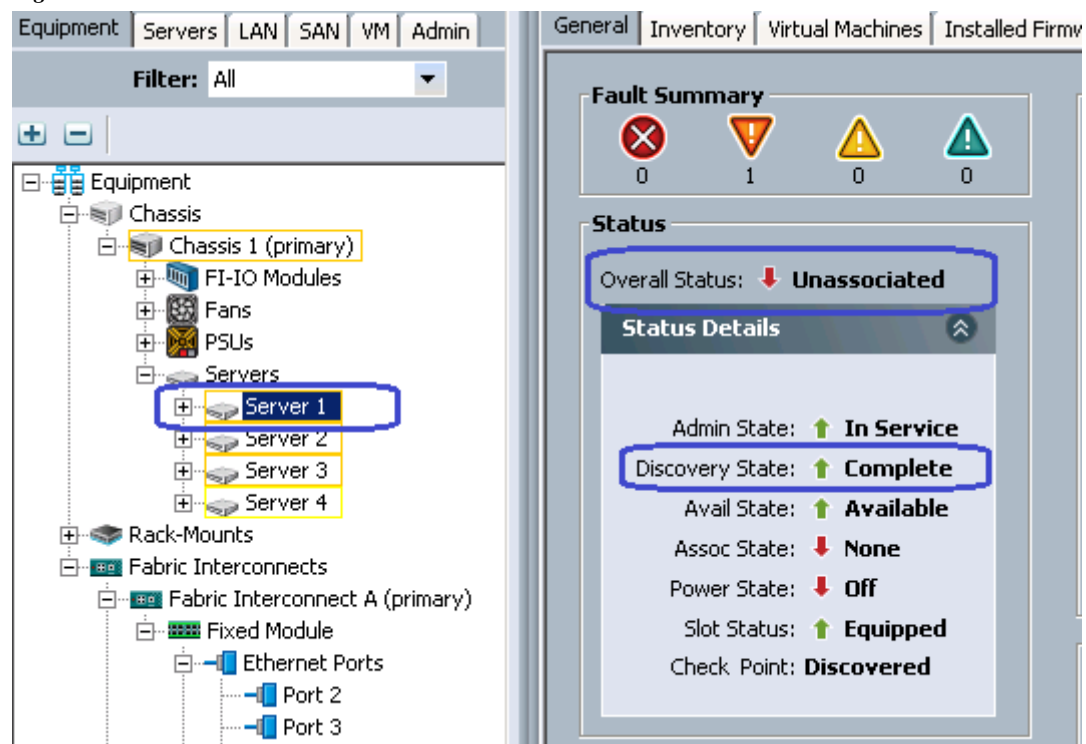
Figure 15 Chassis Discovery Policy



Also You will see 4 blades discovered automatically under **Equipment** tab with over-all status as **Unassociated** and availability state as **Available**, and discovery state as **Complete** as shown below:

On equipment tab, under **Equipment** > **Chassis** > **Chassis <id>** > **Servers** as shown below:

Figure 16 Server Association State





## Upstream / global network configuration

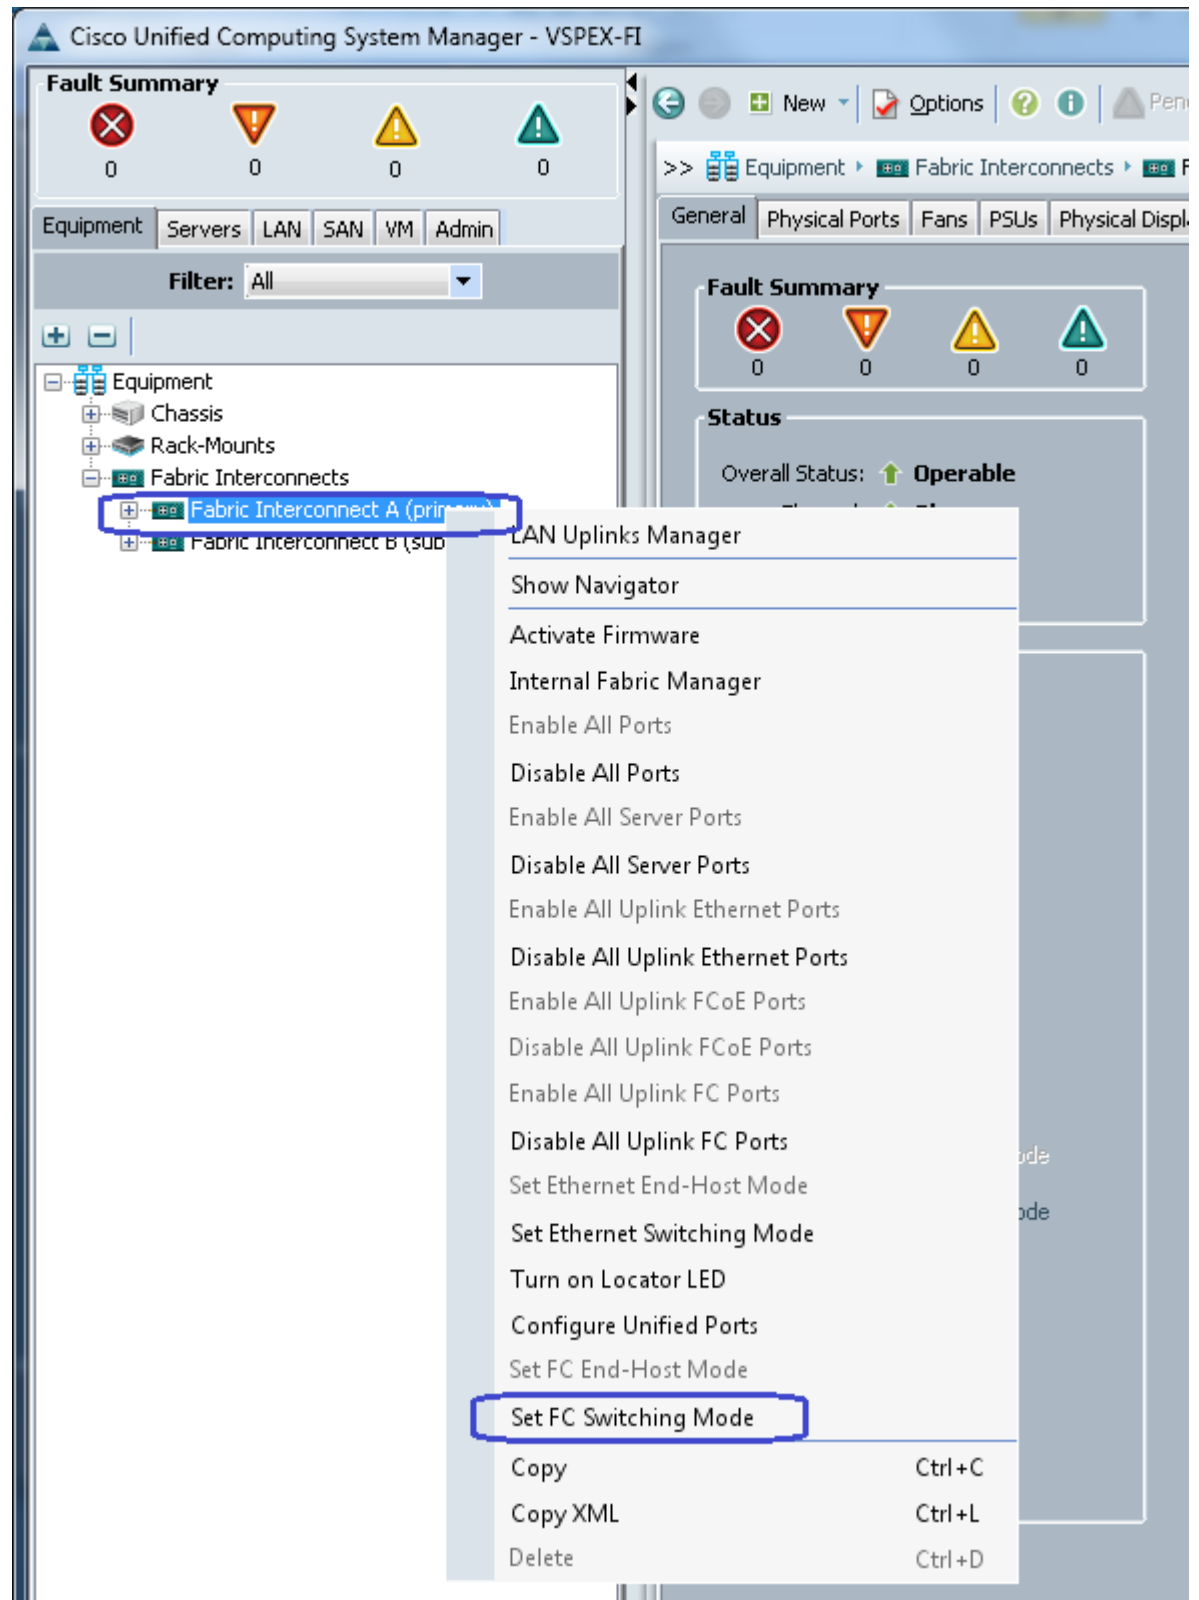
This subsection includes a few upstream/ global network configuration listed below:

1. Move to FC switching mode
2. Uplink VLAN configuration
3. Configure universal ports as FC ports
4. Configure uplink ports
5. Configure Ethernet port as Appliance Ports
6. Configure FC appliance ports
7. Configure FC Zoning policies
8. Configure QoS classes and QoS policy for jumbo MTU

Follow these steps for network configuration:

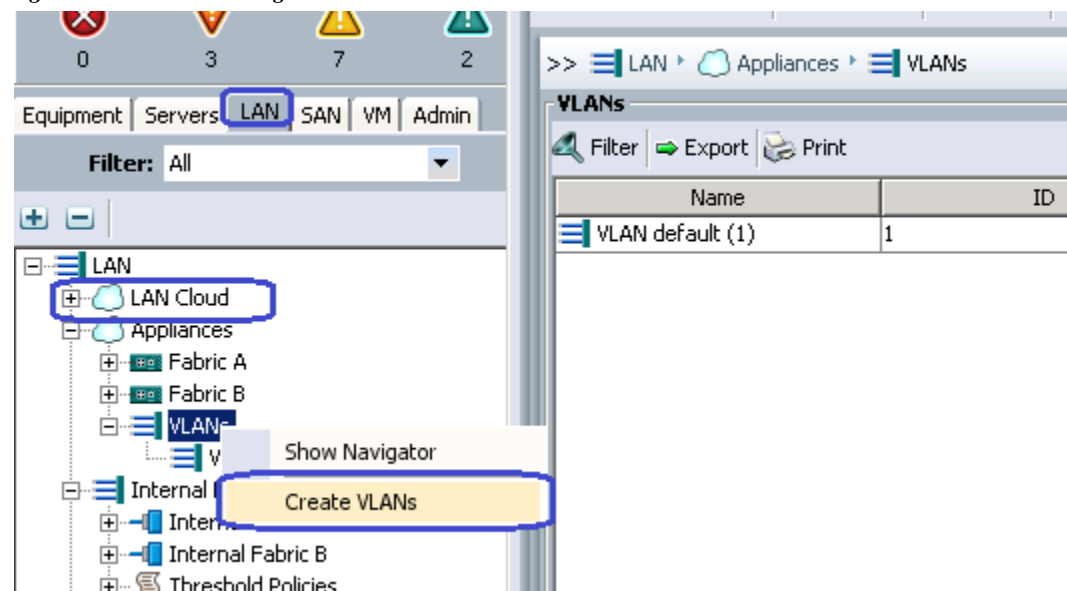
1. From the **Equipment** tab, select and right-click on **Fabric Interconnect A**, and click **Set FC Switching Mode**.

Figure 17 Setting Cisco UCS FI to FC Switching Mode



2. You would see a warning message that Fabric Interconnects would be restarted as a result of this action. Click **Yes**. Both the FIs would reboot (first the secondary FI and then the primary FI). This action is traffic disruptive, so make sure that you perform this operation during maintenance window, if you are working on a production environment.
3. From the **LAN** tab, expand **LAN > LAN Cloud**, and right-click on **VLANs**, and click **Create VLANs**.

Figure 18 Creating VLANs



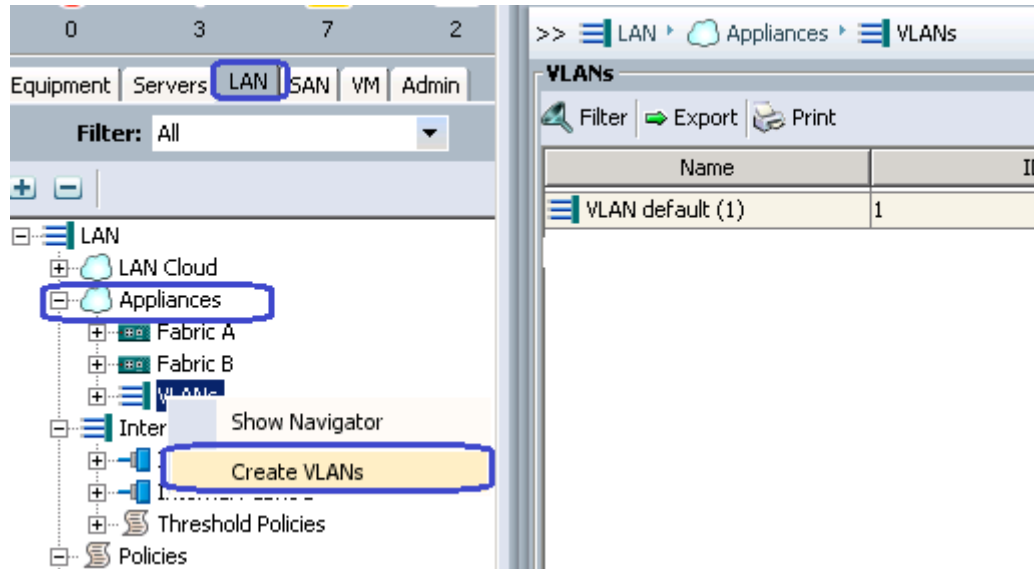
4. Give name to the VLAN and assign ID. Keep the default option **Common/Global** radio button selected.

Figure 19 Creating VLANs for Management

The screenshot shows the 'Create VLANs' form in the Cisco UCS Manager web interface. The form has a title 'Create VLANs' and a subtitle 'Create VLANs'. It contains several input fields and radio buttons. The 'VLAN Name/Prefix' field is set to 'vSphereMgmt'. The 'Multicast Policy Name' dropdown is set to '<not set>'. The 'Common/Global' radio button is selected. Below the radio buttons, there is a text box for 'Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")'. The 'VLAN IDs' field is set to '11'. The 'Sharing Type' dropdown is set to 'None'.

5. Click **OK** and deploy vSphereMgmt VLAN. Repeat these steps to create VLANs for Storage - VLAN20, VM-Data - VLAN30 and vMotion - VLAN40.
6. Click the **LAN** tab, expand **LAN > Appliance Cloud**, and right-click on VLANs, and click **Create VLANs**.

Figure 20 Creating VLANs



7. Give name to the VLAN and assign ID. Keep the default option **Common/Global** radio button selected.

Figure 21 Creating VLANs for Storage

**Create VLANs**

VLAN Name/Prefix: **Storage**

☒ 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.

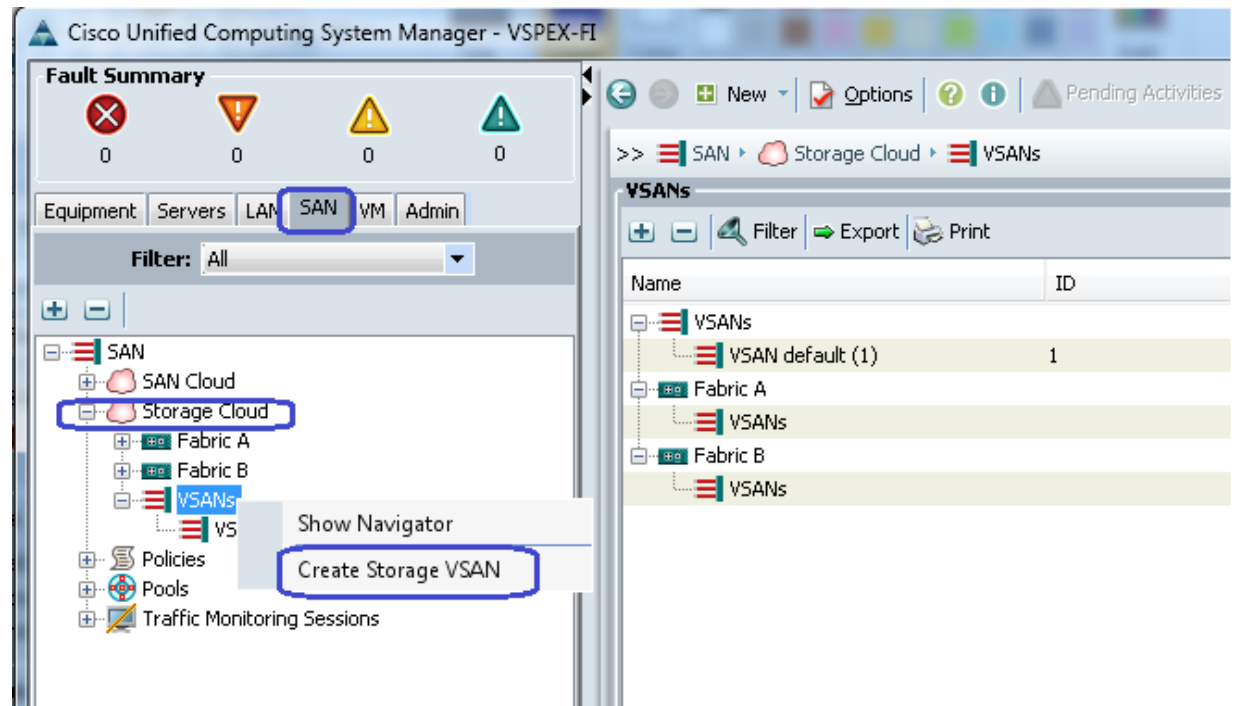
Enter the range of VLAN IDs (e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs: **20**

Check Overlap OK Cancel

8. Click the **SAN** tab, and expand **Storage Cloud**, and right-click on **VSANs**. Click **Create Storage VSAN**.

Figure 22 Creating VSAN



9. Give a name to the VSAN, enable FC zoning and provide VSAN ID and its corresponding FCoE VLAN ID. FCoE VLAN ID should not have conflict with any of the VLANs configured earlier.

Figure 23 Creating Storage VSAN

**Create Storage VSAN**

Name: **Storage**

**FC Zoning Settings**

FC Zoning: ☐ Disabled ☒ Enabled

Do **NOT** enable zoning for this VSAN if the fabric interconnect is connected to an upstream switch that has zoning enabled on the same VSAN.

☒ Common/Global ☐ Fabric A ☐ Fabric B ☐ Both Fabrics Configured Differently

You are creating a global VSAN that maps to the same VSAN ID in all available fabrics.

Enter the VSAN ID that maps to this VSAN.

VSAN ID: **10**

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: **10**

OK Cancel

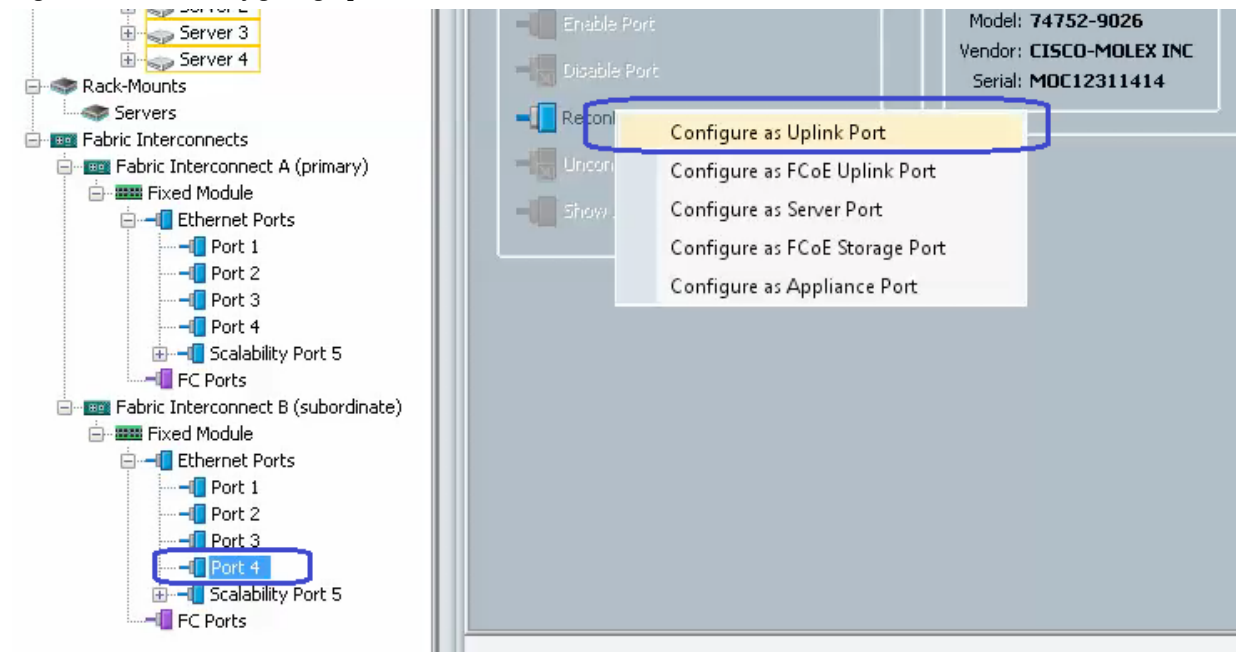
10. The new Cisco UCS 6324UP Fabric Interconnects have 4 Universal Ports. By default, 4 of the physical ports on the FI are unconfigured, but can be converted to Fibre-Channel ports, Ethernet ports or Appliance Ports. For this design we chose Port 4 as the Network Uplink Port (for LAN connection), Port 2 and Port 3 as Appliance Ports for Direct connect Storage to VNXe Array, and configured Port 1 as an FC storage port for SAN Booting the ESXi5.5 from VNXe array.

**Note**

On the 6324 Fabric Interconnect Converting the Unified Port into FC ports is supported only starting from Port 1, not from Port 4.(Configuring the FC storage port on Fabric Interconnect is shown in step15).

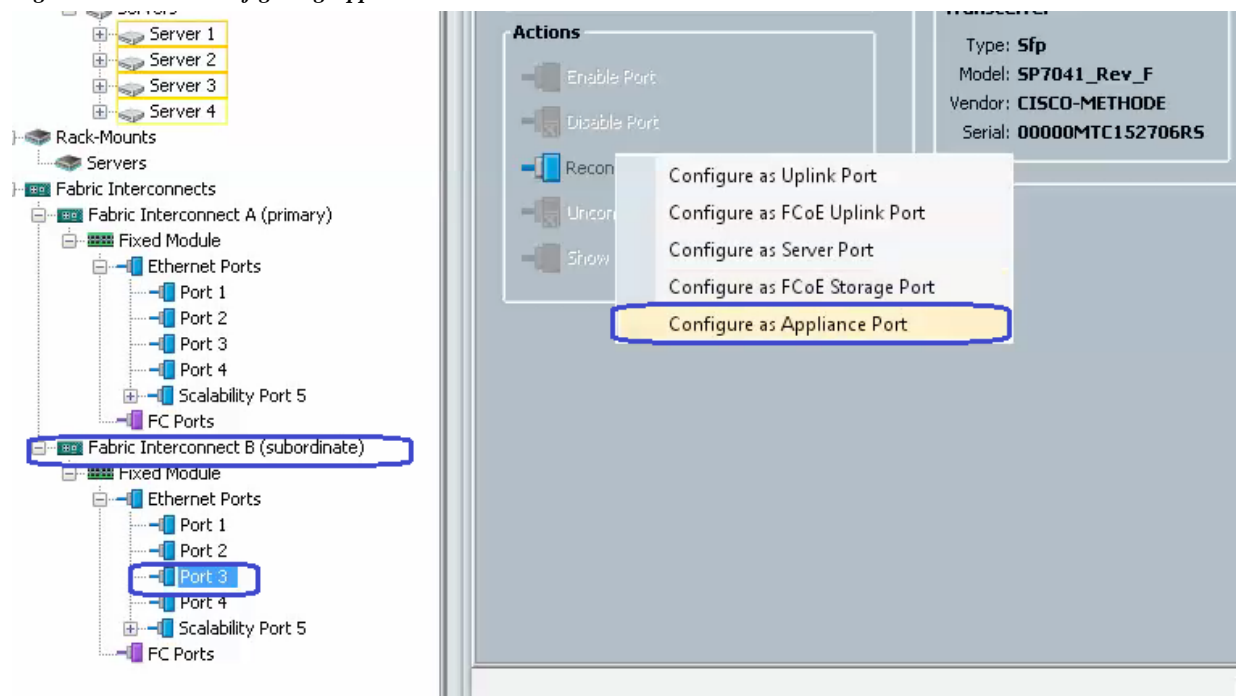
11. As shown below, Click **Fabric Interconnect B > Port 4** and Click **Configure as Uplink Port** and then click **Yes** to confirm. Repeat the same step on Fabric Interconnect A.

Figure 24 Configuring Uplink Ports



12. Click **Fabric Interconnect B > Port 2** and click **Configure as Appliance Port** and then click **Yes**.

Figure 25 Configuring Appliance Port



13. Select Priority as **Platinum** for Jumbo MTU 9000 and Speed as 1 Gbps. Click **Access** and then select VLAN as **Storage**. (In this solution, we use 1 Gbps Ethernet for NFS Storage connectivity) and click **OK** to confirm the selection.



Figure 26 Appliance Port Configuration Window

**Configure as Appliance Port**

Priority: **Platinum**

Pin Group: <not set>

Network Control Policy: default

Flow Control Policy: default

Admin Speed(gbps): ☒ 1 Gbps ☐ 10 Gbps ☐ 20 Gbps ☐ 40 Gbps

**VLANs**

Port Mode: ☐ Trunk ☒ Access

Select VLAN: Storage

+ Create VLAN

☐ Ethernet Target Endpoint

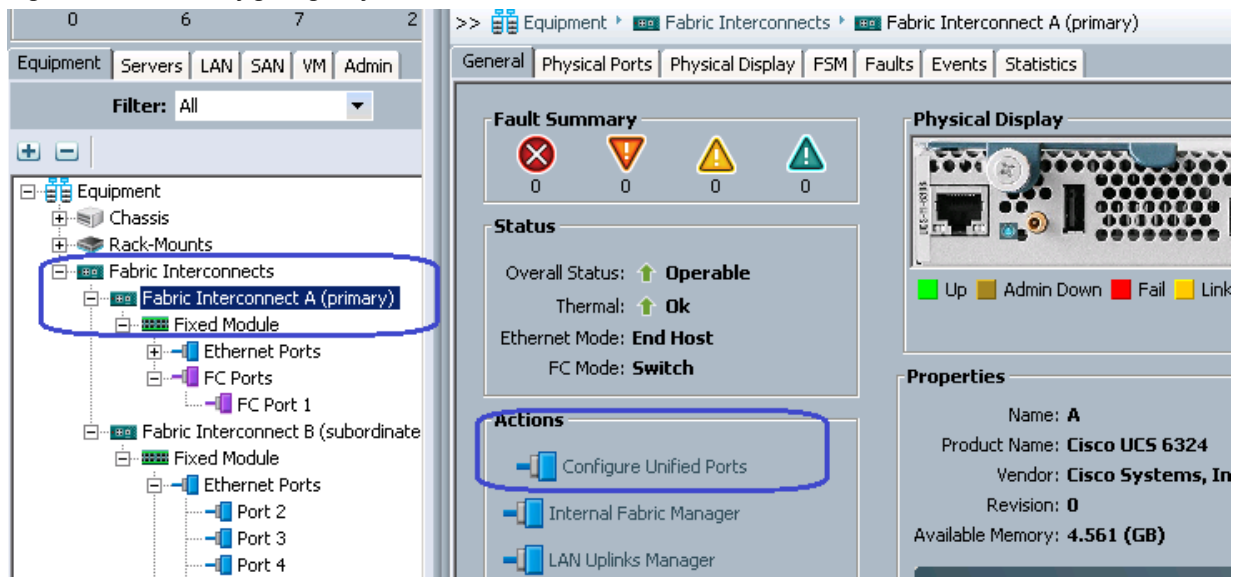
Name:

MAC Address:

OK Cancel

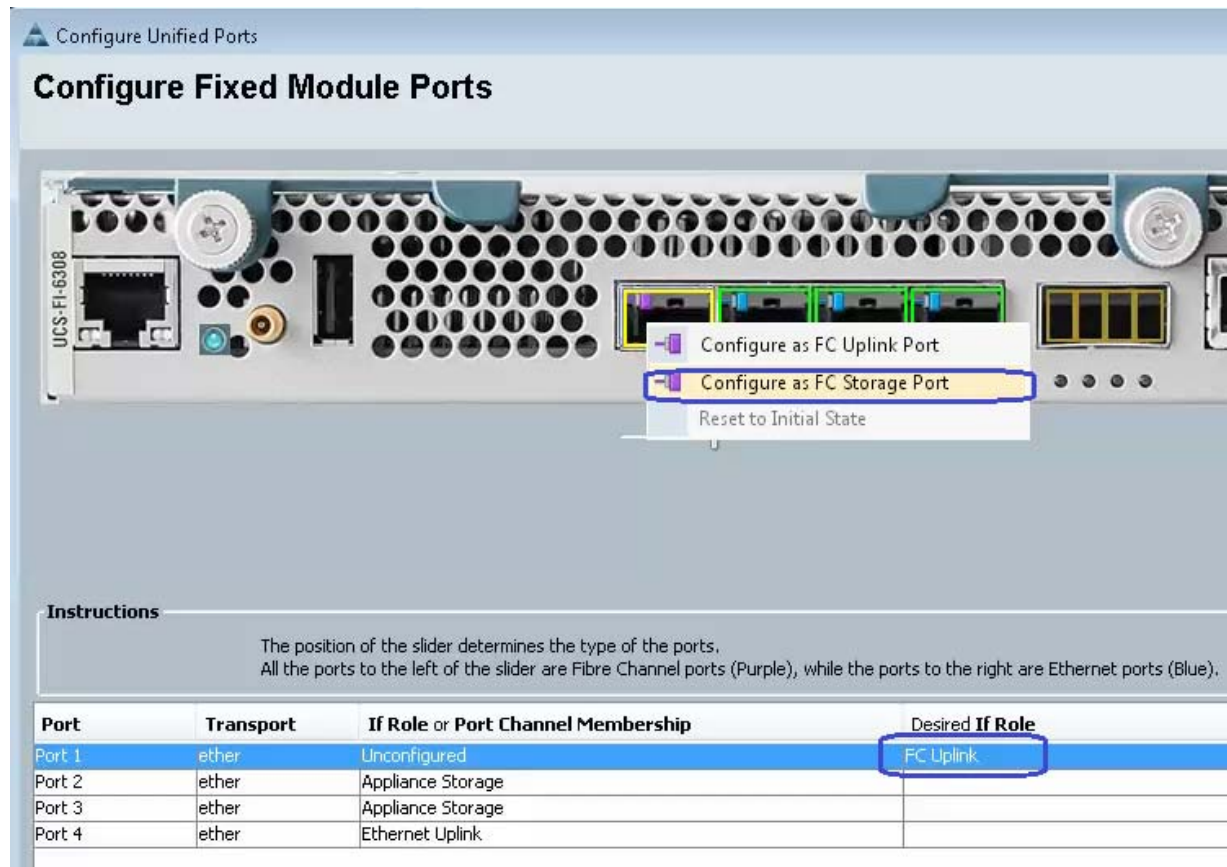
14. Repeat step 12 and 13 to configure Fabric Interconnect A Ethernet Port 2 and Port3 as Appliance Ports.
15. We need FC connectivity to EMC VNXe storage array at least for SAN boot. For configuring FC connectivity, click **Equipment** tab, expand **Fabric Interconnects** and click **Fabric Interconnect A** > **Configure Unified Ports**. Click **Yes** in the warning message window.

**Figure 27** *Configuring Unified Ports*



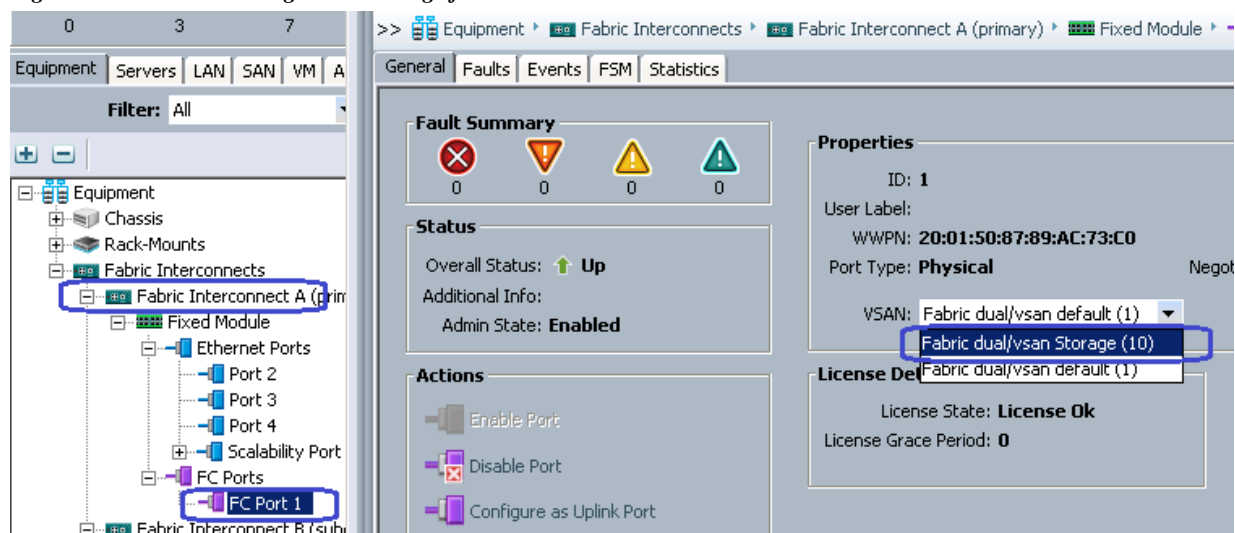
16. Click **Configure Unified Ports** to open the wizard. This will pop-up a warning message about Fabric Interconnect reboot on port changes. Click **Yes** to continue.
17. Move the slider bar at the top to select Port1. Make sure that Port 1 is showing FC Uplink. Then, right-click Port 1 and click **Configure as FC Storage Port** and then click **Finish**. Click **OK** to reboot the FI.

**Figure 28** *Configuring Unconfigured Ports as FC Storage Port*



18. Once the Fabric Interconnect A is rebooted, repeat steps 15, 16 and 17 for Fabric Interconnect B as well.
19. After the successful reboot of both FIs, physical FC ports further need to be classified as FC storage ports for directly attached storage array. Click **Fabric Interconnect A > Expand Fixed Module > Expand FC ports > FC port 1**. On the right pane, choose the VSAN we created earlier and click **Save Changes**.

Figure 29 Choosing VSAN Storage for FC Ports



Repeat the Step18, for Fabric Interconnect B FC ports.

20. After the above changes, EMC VNXe storage array would do Fibre Channel flogi into the Fabric Interconnect's. Using the WWPN of the VNXe storage array, we can carve out the SAN boot policy on the UCS Manager. Use SSH connection to the UCSM Virtual IP address, and issue "connect nxos a" command. In the read-only NXOS shell, issue "show flogi database" command and note down the WWPN of the storage array.

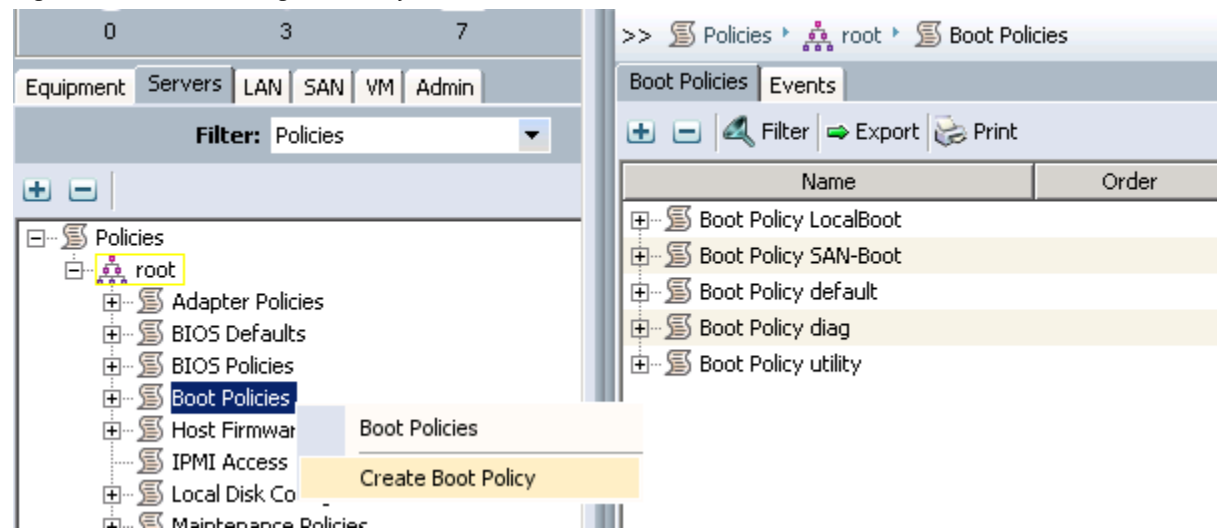
```
VSPEX-FI-A# connect nxos a
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
VSPEX-FI-A(nxos)# show flogi database

-----
INTERFACE      VSAN    FCID      PORT NAME      NODE NAME
-----
fc1/1          10      0x1b00ef  50:06:01:64:08:e0:03:68  50:06:01:60:88:e0:03:68
Total number of flogi = 1.

VSPEX-FI-A(nxos)#
```

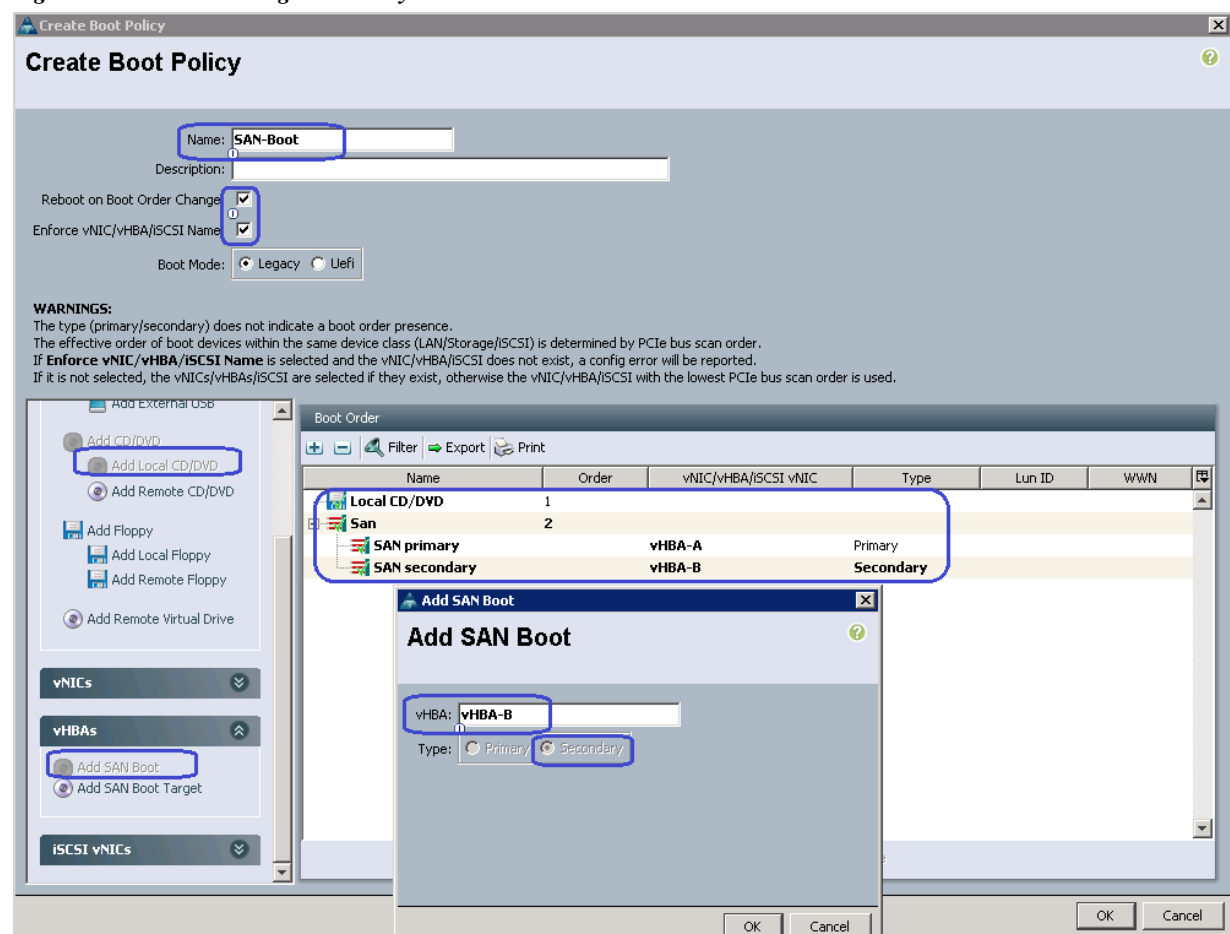
21. On UCSM GUI, click **Servers, Expand Policies > root**, and right-click Boot Policies and click **Create Boot Policy**.

Figure 30 Creating Boot Policy



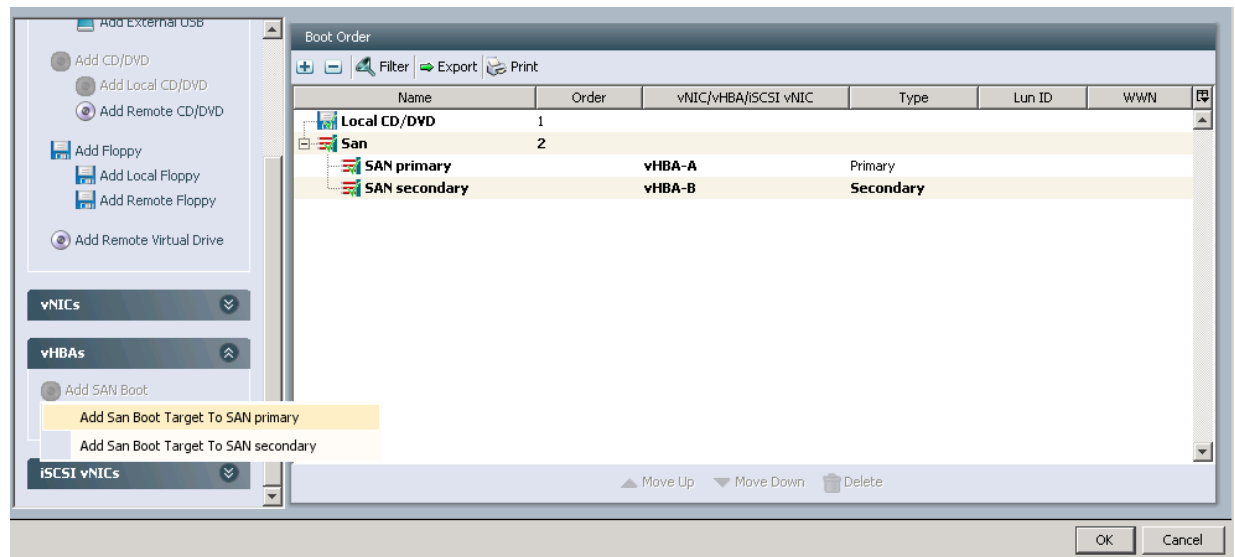
22. Specify Boot Policy Name, check the **Reboot on Boot Order Change** check box. Then click **Add Local CD/DVD** as first Boot order, click **Add SAN Boot** and specify the vHBA name for Fabric A as **Primary** and similarly specify the vHBA name for Fabric B as **Secondary**.

Figure 31 Creating Boot Policy Window



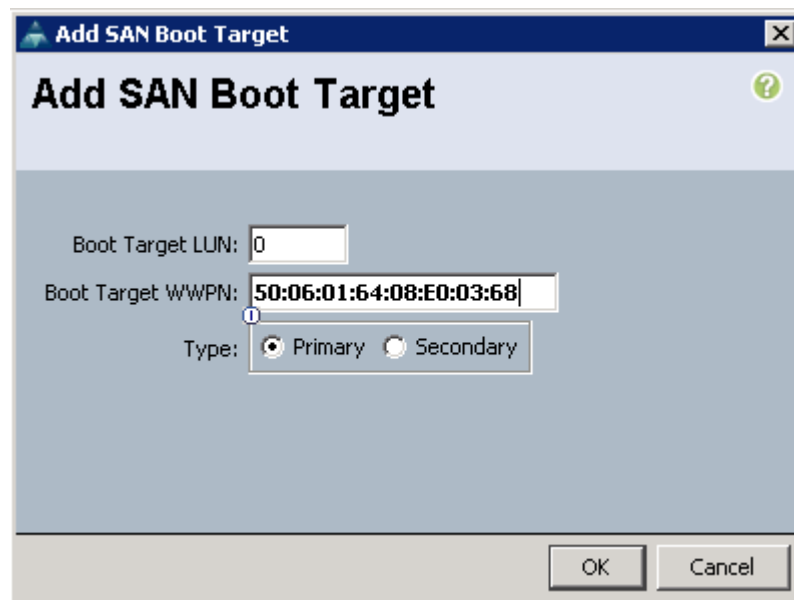
23. Once both the vHBAs are added, click **Add SAN Boot Target** under the vHBAs area, and click **Add SAN Boot Target to SAN primary**.

**Figure 32** Adding SAN Boot Target to SAN Primary



24. Provide target WWP of the VNXe storage device (which can be obtained using “show flogi database” on UCSM CLI command executed under “connect nxos {a|b}” shell as described in the previous subsection). Keep the target as **Primary**.

**Figure 33** Adding SAN Boot Target as Primary



25. Repeat step 24 to add secondary VHBA SAN boot target for the Target WWP for the fabric B too.

**Figure 34**      *Setting SAN Boot Target*

Name	Order	vNIC/vHBA/iSCSI vNIC	Type	Lun ID	WWN
Local CD/DVD	1				
San	2				
SAN primary		vHBA-A	Primary		
SAN Target primary			Primary	0	50:06:01:64:08:E0:03:68
SAN secondary		vHBA-B	Secondary		
SAN Target primary			Primary	0	50:06:01:65:08:E0:03:68

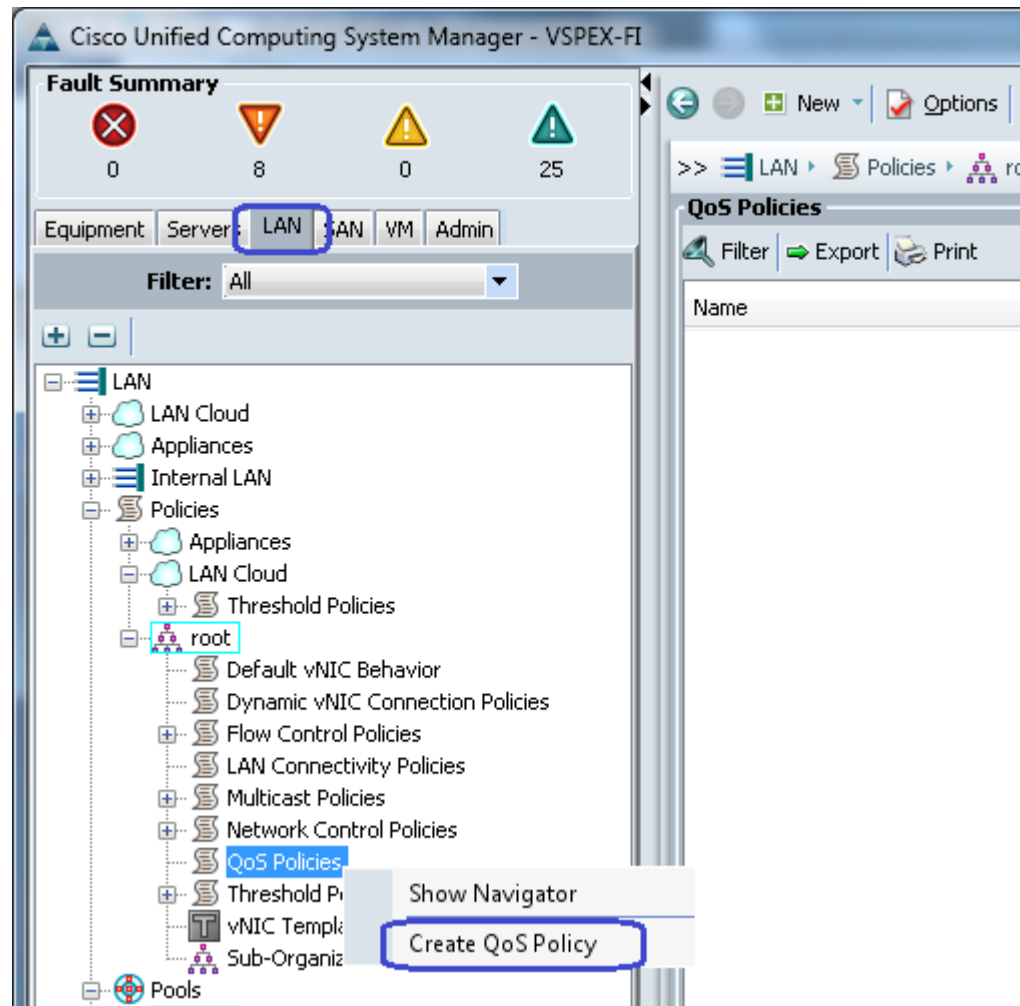
26. Next global configuration task is QoS configuration. Click the **LAN** tab, expand **LAN > LAN Cloud > QoS System Class**. Enable Platinum priority, and set MTU to 9216. Keep other configuration as default and save the configuration.

**Figure 35**      *QoS Configuration*

Priority	Enabled	CoS	Packet Drop	Weight	Weight (%)	MTU	Multicast Optimized
Platinum	<input checked="" type="checkbox"/>	5	<input type="checkbox"/>	10	50	9216	<input checked="" type="checkbox"/>
Gold	<input type="checkbox"/>	4	<input checked="" type="checkbox"/>	9	N/A	normal	<input type="checkbox"/>
Silver	<input type="checkbox"/>	2	<input checked="" type="checkbox"/>	8	N/A	normal	<input type="checkbox"/>
Bronze	<input type="checkbox"/>	1	<input checked="" type="checkbox"/>	7	N/A	normal	<input type="checkbox"/>
Best Effort	<input checked="" type="checkbox"/>	Any	<input checked="" type="checkbox"/>	5	25	normal	<input type="checkbox"/>
Fibre Channel	<input checked="" type="checkbox"/>	3	<input type="checkbox"/>	5	25	fc	N/A

27. From the **LAN** tab, expand **LAN > Policies > root**, and right-click on **Create QoS Policy**.

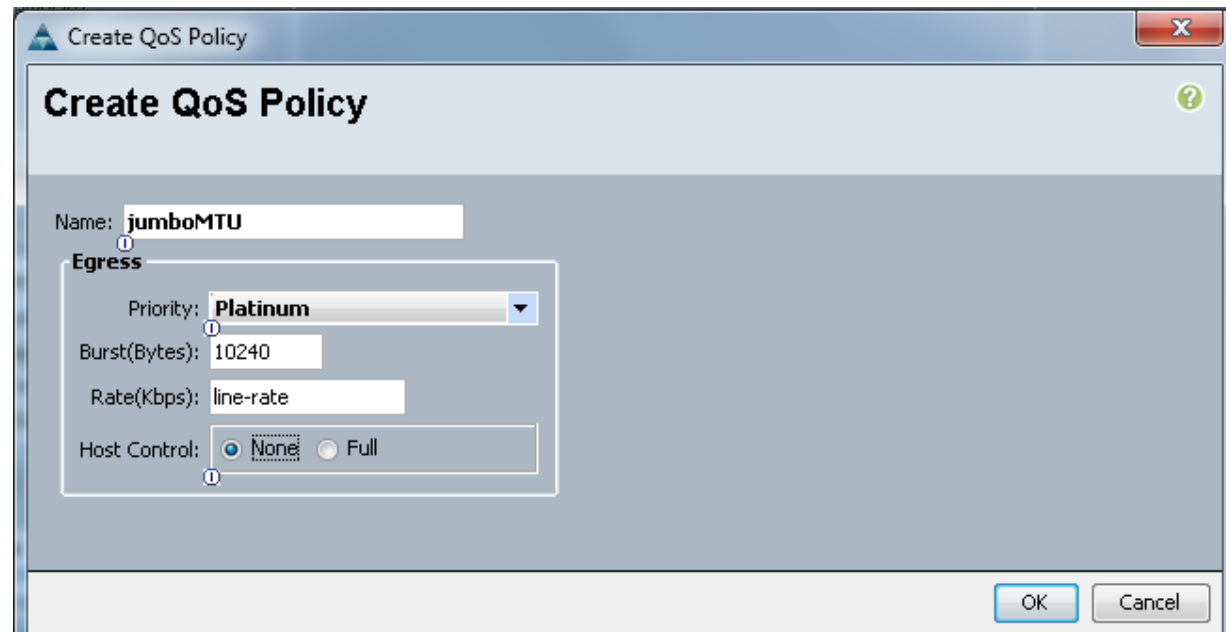
Figure 36 Creating QoS Policy



28. Create a QoS policy with name jumboMTU and select the Priority **Platinum**. Click **OK** to save the configuration.



Figure 37 Creating QoS Policy Window



## Configure identifier pools

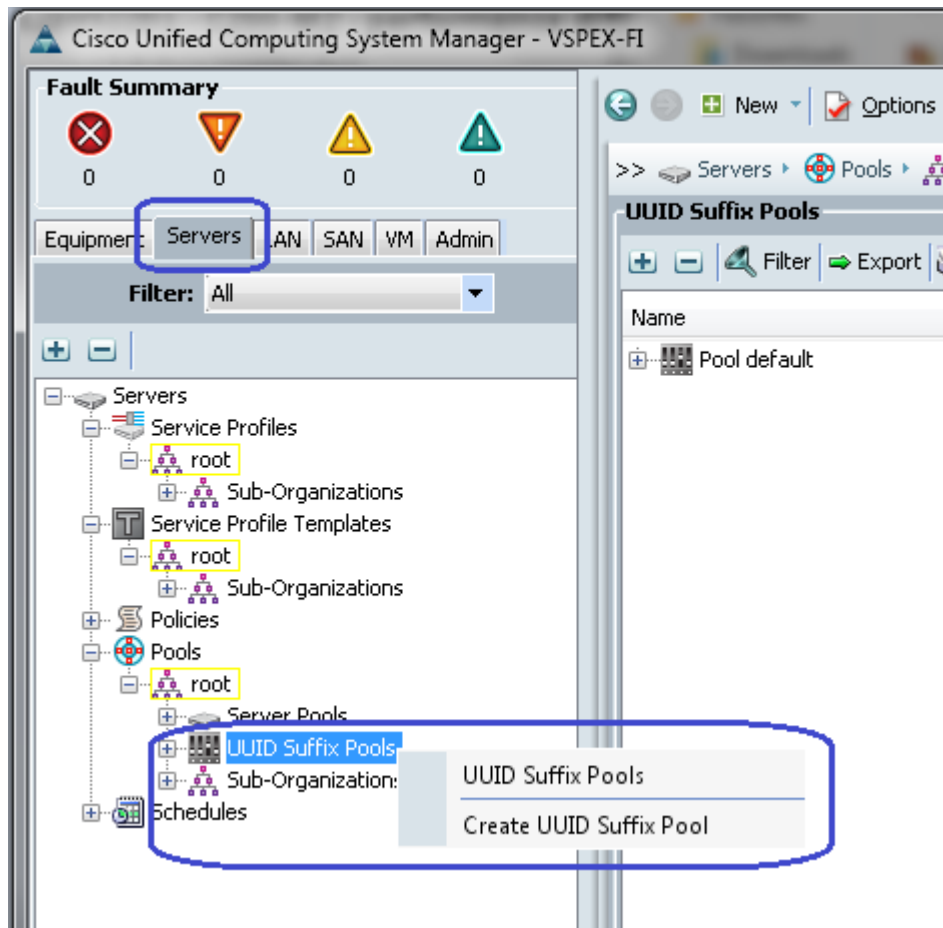
In this section, we would configure following identifier pools used by service profile:

1. Server UUID pool
2. MAC address pool
3. WWN pool
4. Management IP address pool

Follow the following steps to configure pools mentioned above.

1. From the **Servers** tab, expand **Servers > Pools > root**, and right-click on **UUID Suffix pools** and click **Create UUID Suffix Pool**.

Figure 38 Creating UUID Suffix Pool



2. Provide name and description to the UUID suffix pool. Keep other configuration as default.

**Figure 39** *Create UUID Suffix Pool*

**Create UUID Suffix Pool**

**Unified Computing System Manager**

Create UUID Suffix Pool

1. ☒ **Define Name and Description**
2. ☐ **Add UUID Blocks**

**Define Name and Description**

Name: **BO-VSPEX-UUIDs**

Description: **UUID Pool for Branch Office VSPEX Servers**

Prefix: ☒ Derived ☐ other

Assignment Order: ☒ Default ☐ Sequential

< Prev Next > Finish Cancel

3. Click **Add** to add the UUID block.

**Figure 40** *Adding a Block of UUID Suffixes*

**Create UUID Suffix Pool**

**Unified Computing System Manager**

Create UUID Suffix Pool

1. ☒ **Define Name and Description**
2. ☒ **Add UUID Blocks**

**Add UUID Blocks**

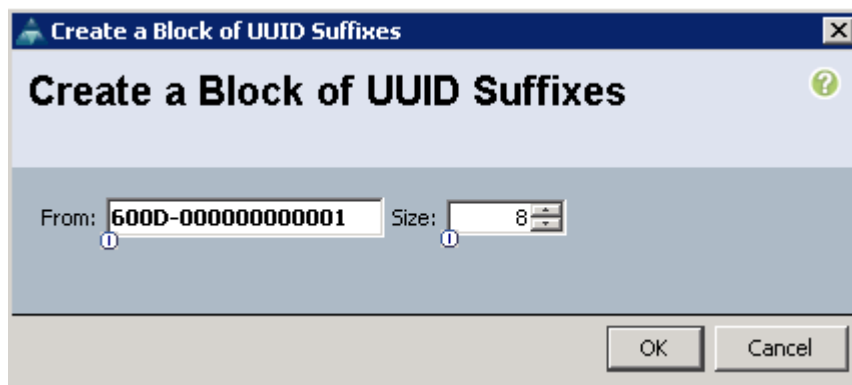
Name	From	To

+ Add Delete

< Prev Next > Finish Cancel

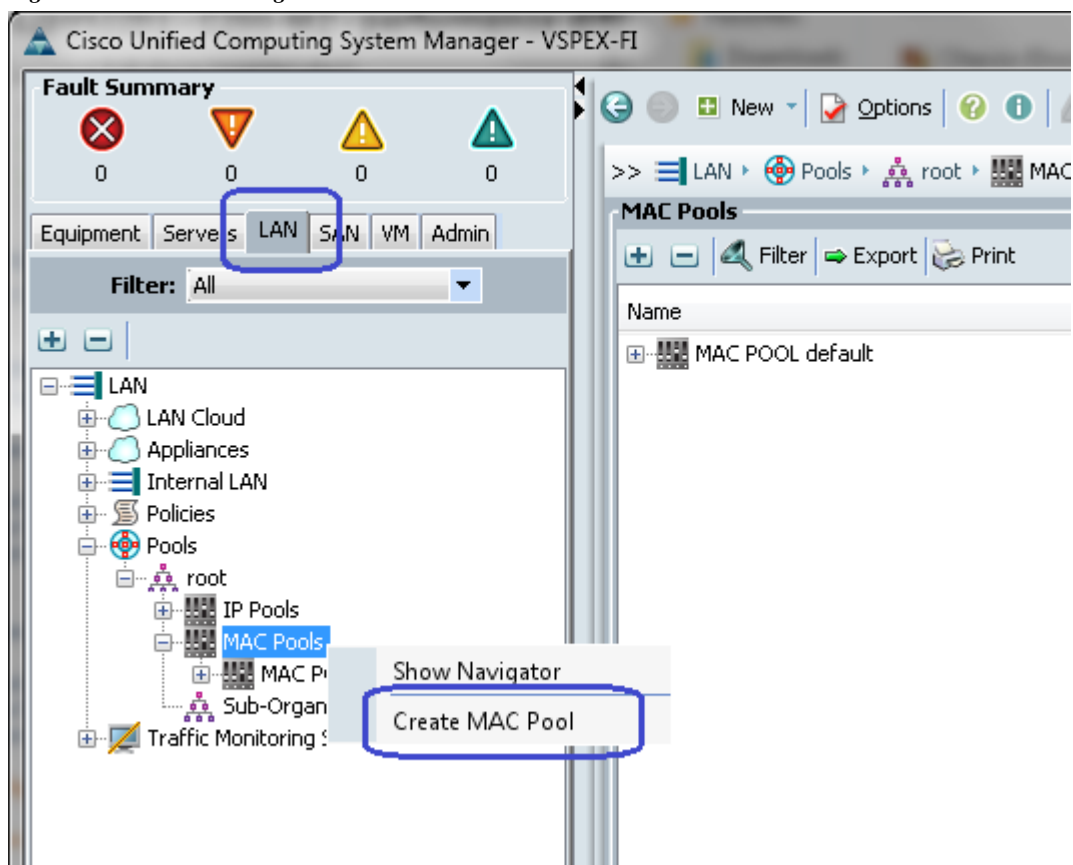
4. Specify the beginning of the UUIDs, and have a large size of UUID block to accommodate future expansion.

Figure 41 Range for UUID Block



5. Click **OK** and then **Finish** to deploy UUID pool.
6. From the **LAN** tab, expand **LAN > Pools > root**, right-click on **MAC Pools** and then click **Create MAC Pool**.

Figure 42 Creating MAC Pool



7. Provide name and description of the MAC pool and click **Next**.

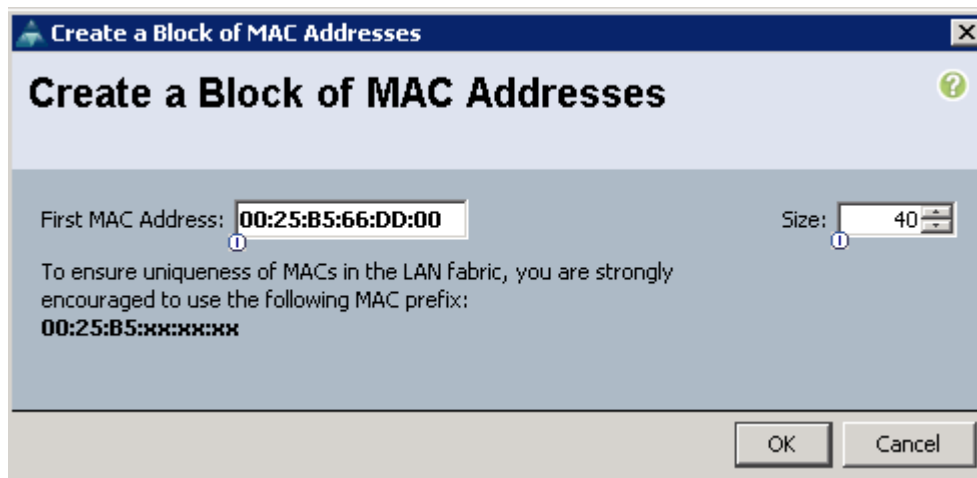
**Figure 43** *Creating MAC Pool Window*

8. Click **Add** to add MAC pool block.

**Figure 44** *Adding a Block of MAC Addresses*

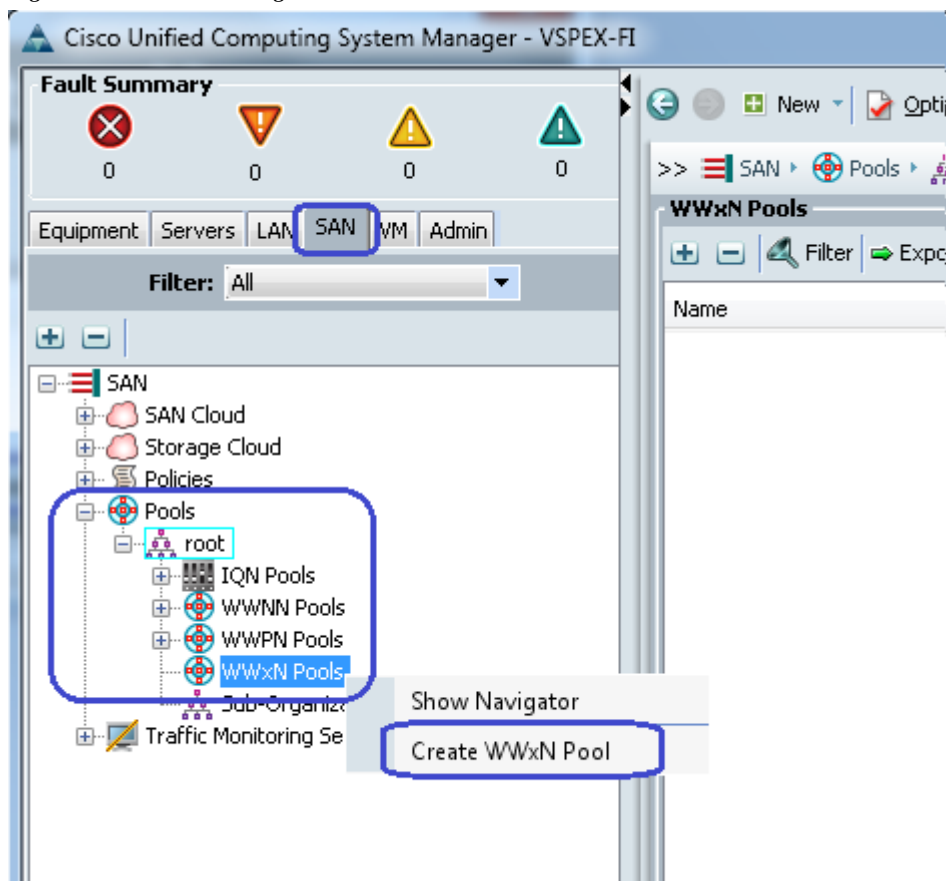
9. Provide the initial MAC address and size of the block. As always, provide large number of MAC addresses to accommodate future expansion. We require 6 MAC addresses per server.

Figure 45 Range for MAC Address Block



10. Click **OK** and **Finish** to complete configuration.
11. From the **SAN** tab, expand **SAN > Pools > root**, right-click on **WWxN Pools**, and click **Create WWxN Pool**.

Figure 46 Creating WWxN Pool



12. Provide name, description and choose **3 Ports per Node** from the drop-down menu.

**Figure 47** *Create WWxN Pool - Defining Pool*

**Create WWxN Pool**

## Unified Computing System Manager

Create WWxN Pool

1. [Define Name and Description](#)
2. [Add WWN Blocks](#)

### Define Name and Description

Name: **BO-VSPEX-WWNs**

Description: **Combined WWNN and WWPN Pool for VSPEX Servers**

Max Ports per Node: **3 Ports Per Node**

Assignment Order: ☒ Default ☐ Sequential

< Prev Next > Finish Cancel

13. Click **Add** to add a block of WWxN IDs.

**Figure 48** *Create WWxN Pool - Adding a Block of WWxN*

**Create WWxN Pool**

## Unified Computing System Manager

Create WWxN Pool

1. [Define Name and Description](#)
2. [Add WWN Blocks](#)

### Add WWN Blocks

Name	From	To

**+ Add** Delete

< Prev Next > Finish Cancel

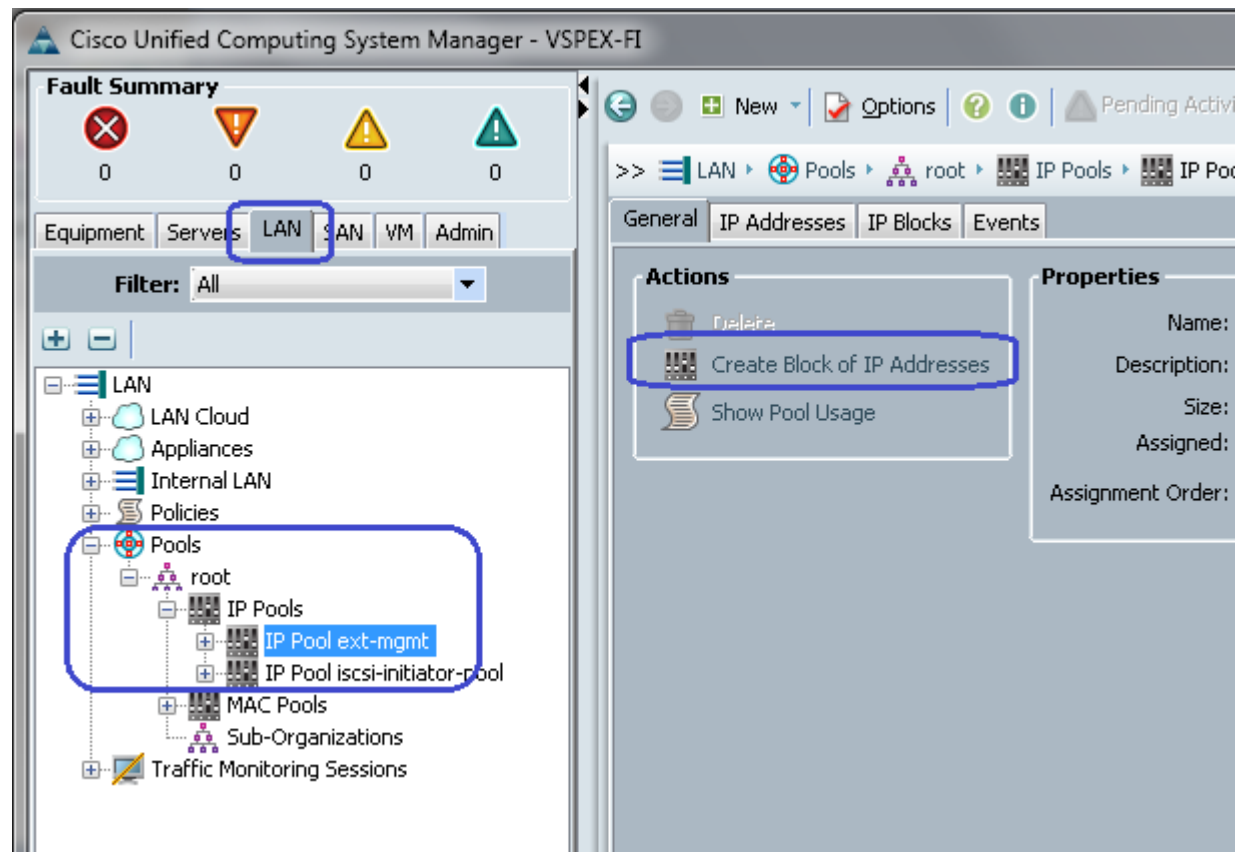
14. Provide beginning of the WWN IDs and sufficiently large number of block size. Click **OK** and **Finish**.

Figure 49 Range for WWxN Block



- Next step is creation of the management IP address block for KVM access of the servers. The default pool for server CIMC management IP addresses are created with the name ext-mgmt. From the **LAN** tab, expand **LAN > Pools > root > IP Pools** and select **IP Pool ext-mgmt**, and click the **Create Block of IP addresses** link on right hand side.

Figure 50 Creating a Block of IP Address





16. Provide initial IP address, size of the pool, default gateway and subnet mask. Click **OK** to deploy the configuration. IP addresses would be assigned to various blade server KVM or CIMC management access from this block.

Figure 51 Size of IPv4 Addresses

The screenshot shows a Windows-style dialog box titled "Create Block of IPv4 Addresses". The main heading inside is "Create a Block of IPv4 Addresses". The dialog contains the following fields and values:

- From:** 10.29.180.225
- Size:** 8
- Subnet Mask:** 255.255.255.0
- Default Gateway:** 10.29.180.1
- Primary DNS:** 0.0.0.0
- Secondary DNS:** 0.0.0.0

At the bottom right, there are "OK" and "Cancel" buttons.

That concludes configuration of all identifier pools and blocks.

## Configure server pool and qualifying policy

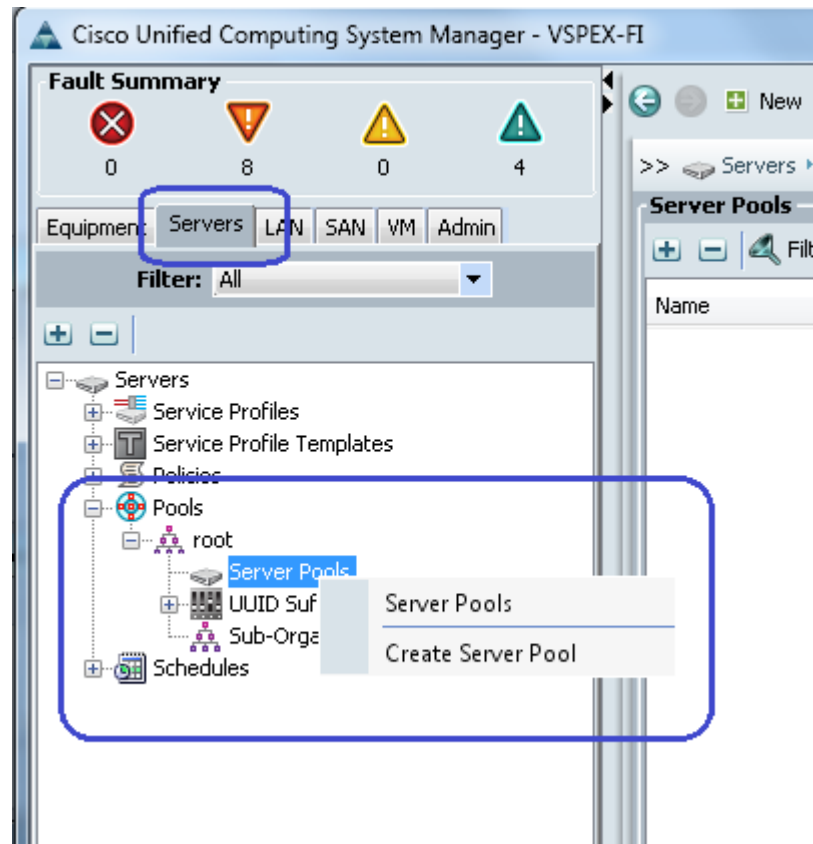
Creation and policy based auto-population of server pool can be divided in to following 3 tasks:

1. Creation of server pool
2. Creation of server pool policy qualification
3. Creation of server pool policy

Follow these steps the accomplish the above mentioned tasks:

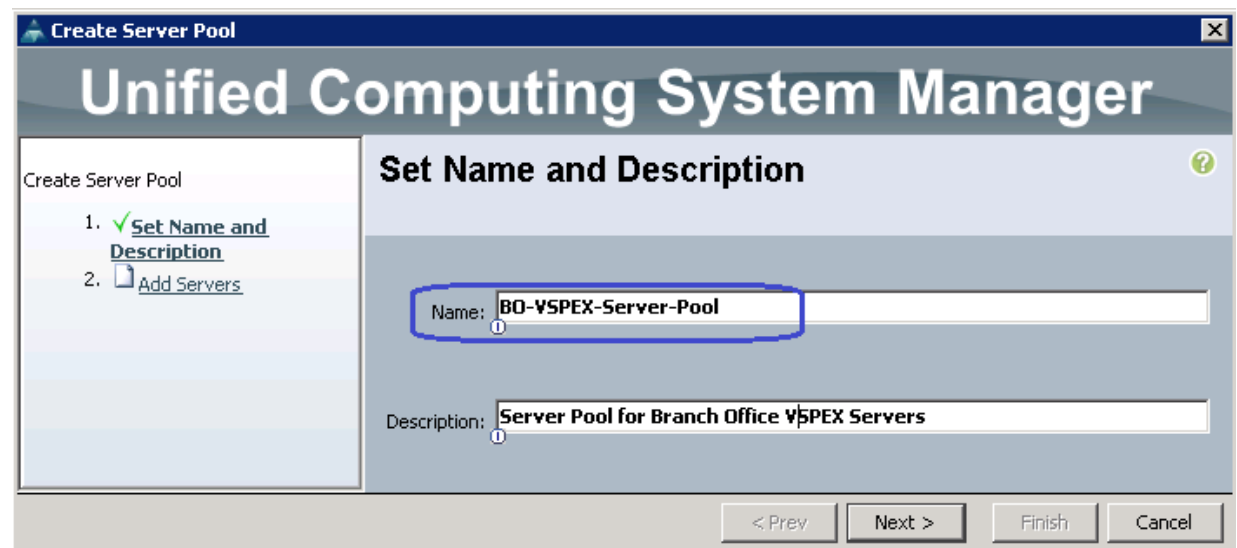
1. From the **Servers** tab, expand **Servers > Pools > root**, right-click on Server Pools and click **Create Server Pool**.

Figure 52 Creating Server Pools



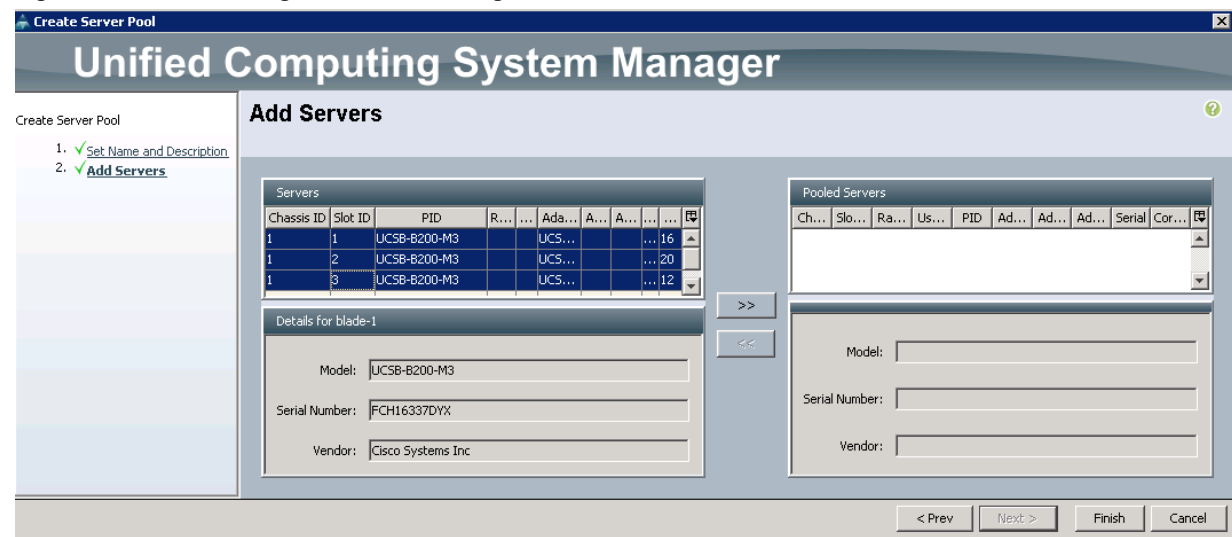
2. Enter the name and description for the server pool, and click **Next**.

Figure 53 Creating Server Pool - Defining Server Pool



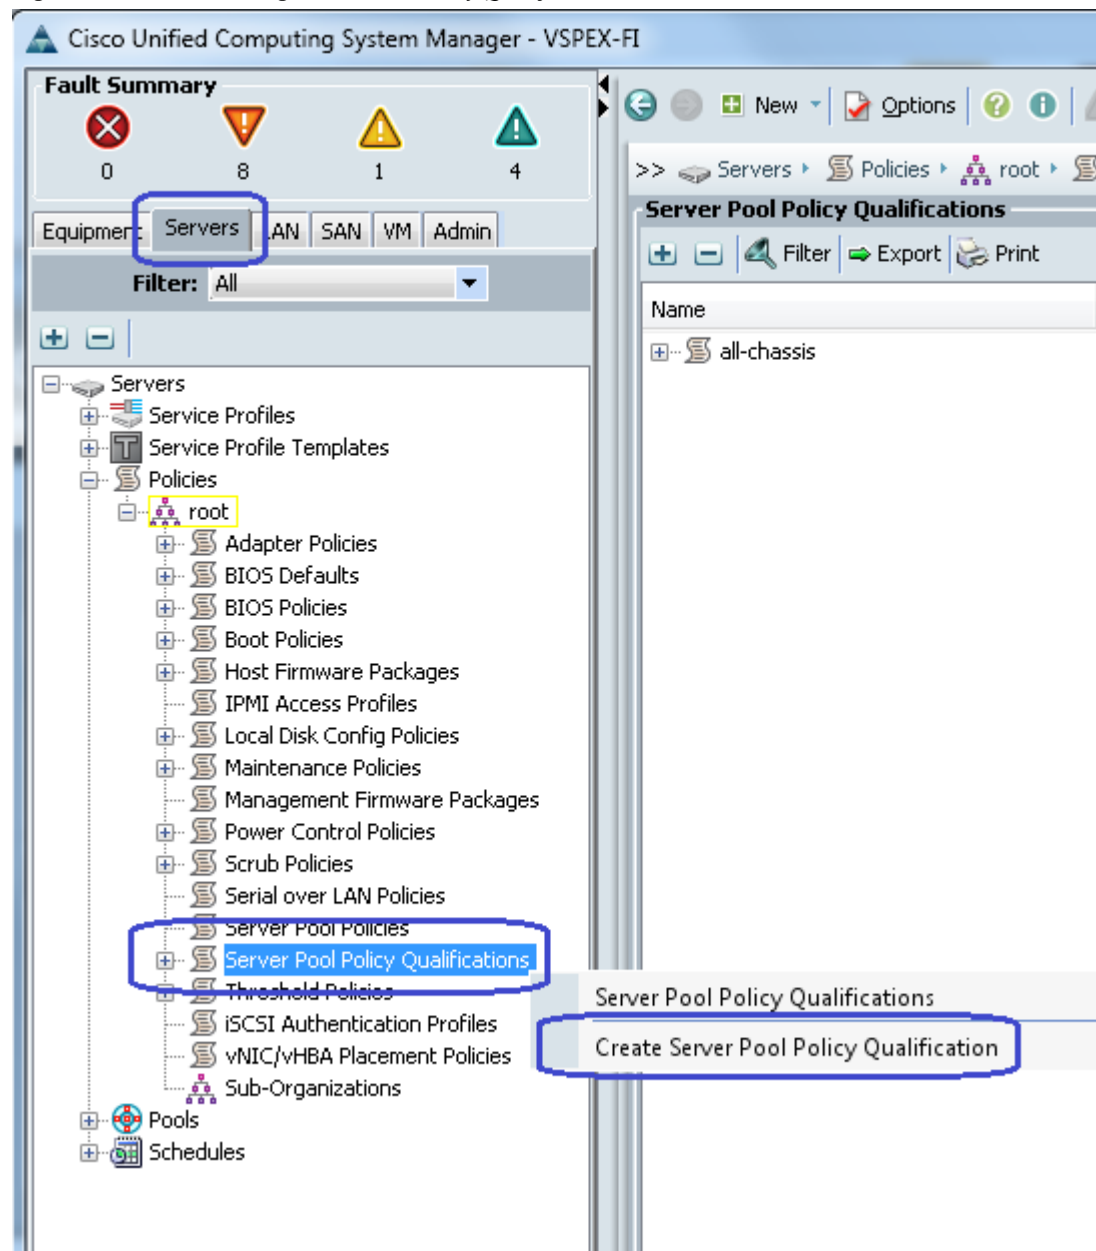
3. Click **Finish** to create the empty server pool. We can add the compute resources to this pool dynamically, based on policy.

**Figure 54** *Creating Server Pool - Adding Servers*



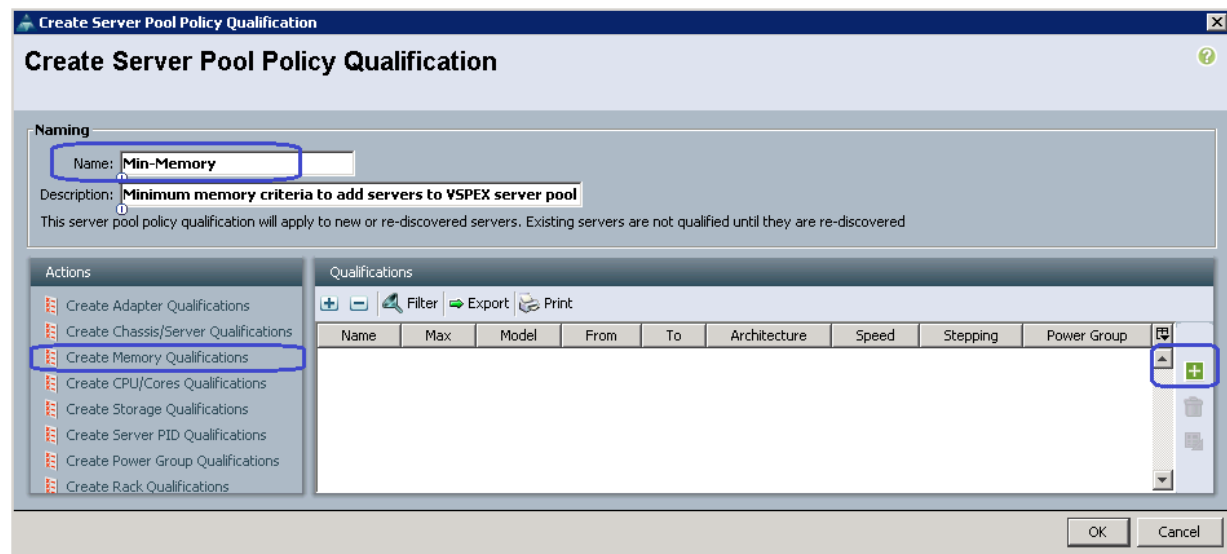
4. From the **Servers** tab, expand **Servers > Policies > root**, right-click on **Server Pool Policy Qualifications** and click **Create Server Pool Policy Qualification**.

Figure 55 Creating Server Pool Policy Qualification



5. Give a name to server policy qualification criterion. We are choosing memory qualification criterion as an example.

**Figure 56** *Creating Server Pool Policy Qualification Window*



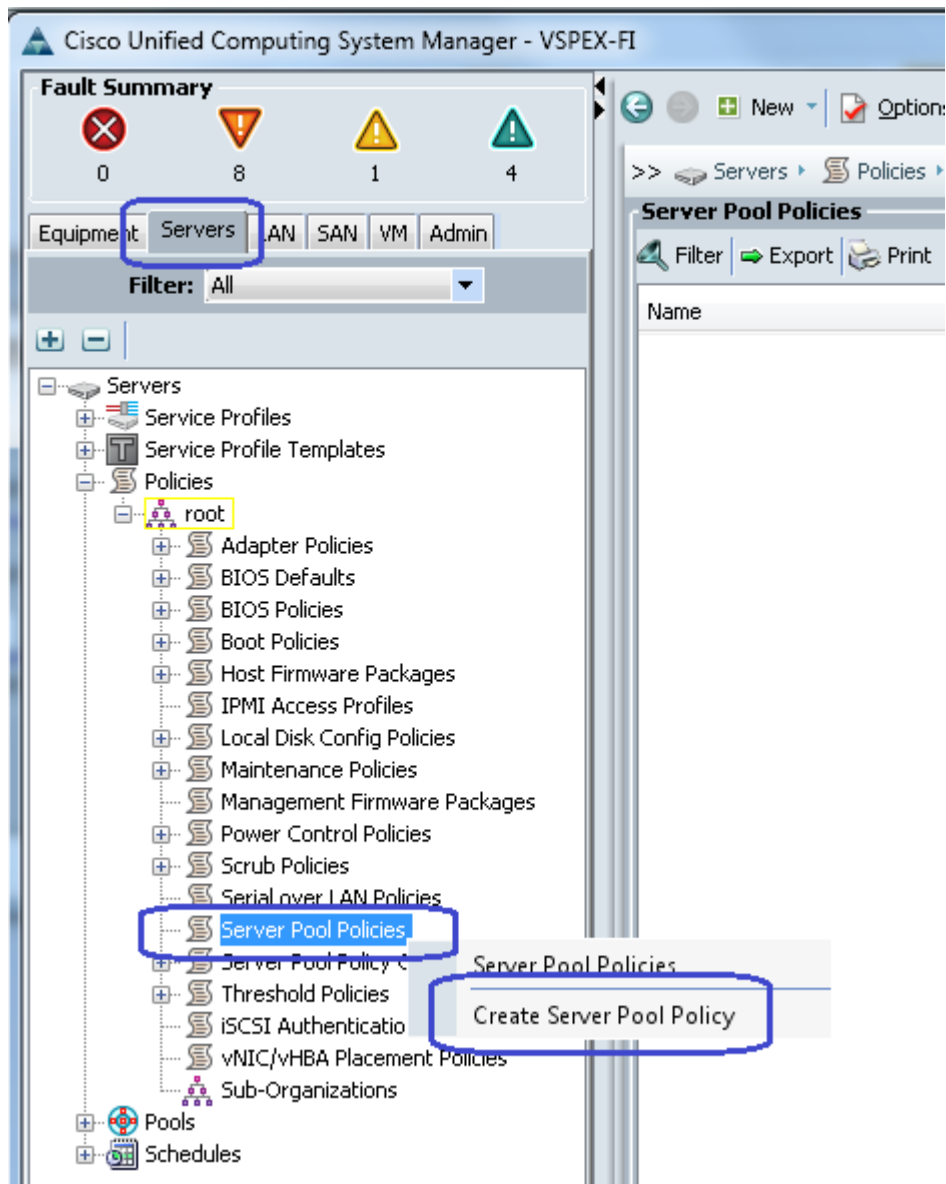
6. Set minimum 128 GB RAM for the pool qualification criterion. Note that this is an example criterion, you may choose a criterion that suites your requirement. Click **OK** twice to create the qualification.

**Figure 57** *Creating Memory Qualification*



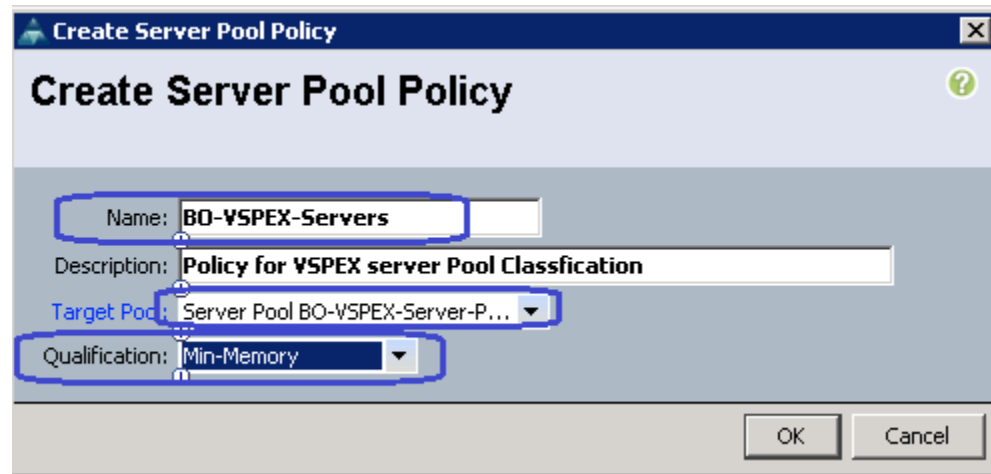
7. Finally, from the **Servers > Policies > root**, right-click on Server Pool Policies and click **Create Server Pool Policy**.

Figure 58 Creating Server Pool Policy



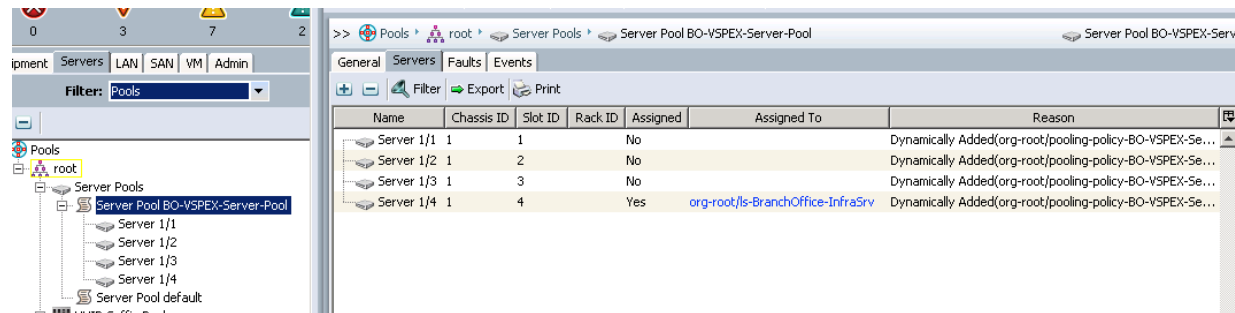
8. Give a name and description to the server pool policy. Choose recently created Target Pool and Qualification. Click **OK** to deploy the configuration.

Figure 59 Creating Server Pool Policy Window



- Click the **Servers** tab on right hand side, you will see that all the compute resources that meet the memory qualification criteria are dynamically added to the server pool.

Figure 60 Dynamically Added Servers Based on Memory Qualification

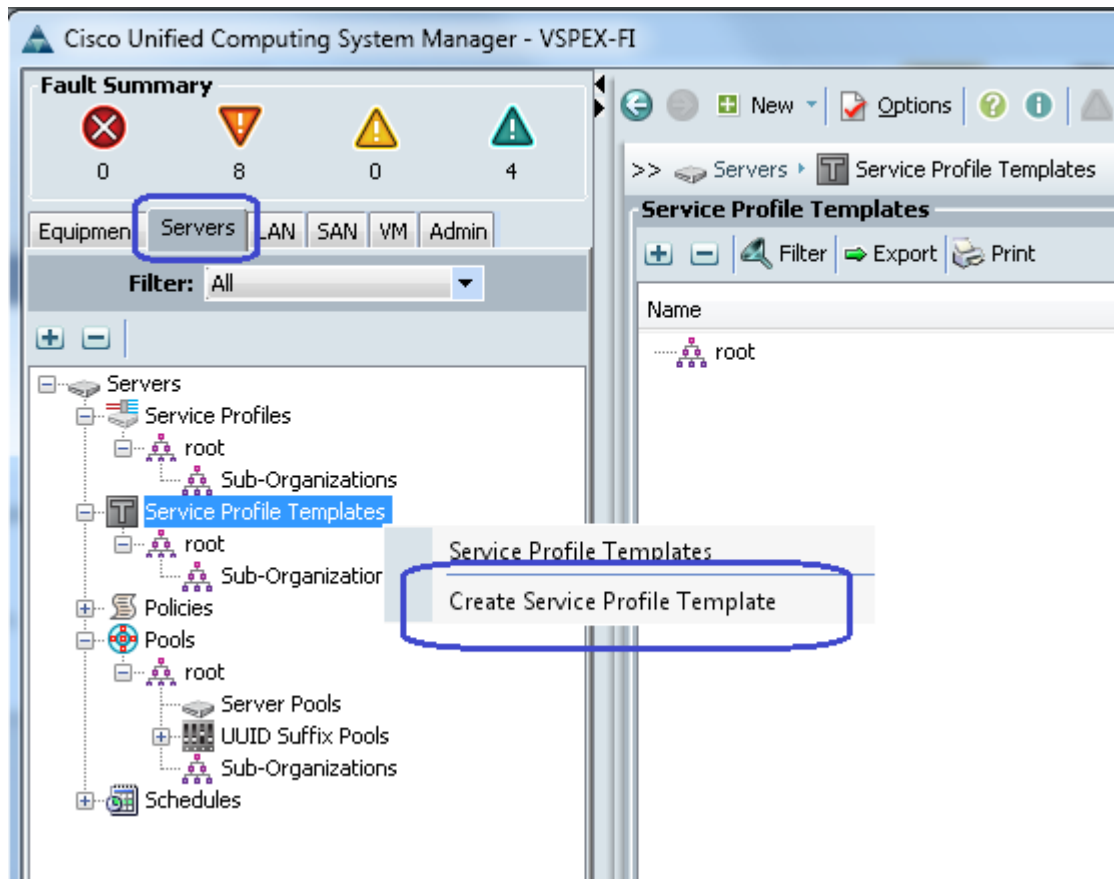


## Configure service profile template

At this point, we are ready to create service profile template, from which we can instantiate individual service profiles later. Follow the following steps to create the service profile template:

- From the **Servers** tab. Expand **Servers**, right-click on Service Profile Templates and click **Create Service Profile Template**.

Figure 61 Creating Service Profile Template



2. Provide service profile template name, keep the type as **Initial Template**, and choose **UUID pool** for UUID assignment.



**Figure 62** *Creating Service Profile Template - Identification*

**Create Service Profile Template**

**Unified Computing System Manager**

**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:

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:

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

- On next step, select configure LAN connectivity as **Expert**. Click **Add** to create a vNIC.

**Figure 63** *Creating Service Profile Template - Networking*

**Create Service Profile Template**

**Unified Computing System Manager**

**Networking**

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy:

How would you like to configure LAN connectivity? ☐ Simple ☒ Expert ☐ No vNICs ☐ Use Connectivity Policy

Click **Add** to specify one or more vNICs that the server should use to connect to the LAN.

Name	MAC Address	Fabric ID	Native VLAN

Delete  Modify

iSCSI vNICs

< Prev Next > Finish Cancel

4. Create a system VNIC for fabric A. Enter system A as the VNIC name, choose the created MAC pool, click the radio button **Fabric A** for Fabric ID, check the vMotion and vSphereMgmt check boxes for VLANs and vSphereMgmt as native VLAN. Choose MTU as 9000, VMware adapter policy and jumboMTU QoS policy.

Figure 64 Networking - Create vNIC Window for vMotion and vSphereManagement Traffic

**Create vNIC**

Name: **system-A**

Use vNIC Template: ☐

MAC Address

MAC Address Assignment: **BO-VSPEX-MACs(40/40)**

[+ Create MAC Pool](#)

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

Fabric ID: ☒ Fabric A ☐ Fabric B ☒ Enable Failover

**VLANs**

Filter Export Print

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	Storage	<input type="radio"/>
<input type="checkbox"/>	VM-Data	<input type="radio"/>
<input checked="" type="checkbox"/>	vMotion	<input type="radio"/>
<input checked="" type="checkbox"/>	vSphereMgmt	<input checked="" type="radio"/>

[+ Create VLAN](#)

MTU: **9000**

**Warning**

Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

Pin Group: **<not set>** [+ Create LAN Pin Group](#)

**Operational Parameters**

**Adapter Performance Profile**

Adapter Policy: **VMWare** [+ Create Ethernet Adapter Policy](#)

QoS Policy: **jumboMTU** [+ Create QoS Policy](#)

Network Control Policy: **<not set>** [+ Create Network Control Policy](#)

OK Cancel

5. Similarly, create one more VNIC with exact same properties on fabric B.
6. Create two more VNICs for NFS server access. Specify the name as Storage A and Storage B for VNICs on Fabric A and B respectively. Check the Storage VLAN check box and mark it as native VLAN. Choose VMware and JumboMTU for adapter policy and QoS policy respectively.

Figure 65 Networking - Create vNIC Window for Storage Traffic

**Create vNIC**

Name: **Storage-A**

Use vNIC Template: ☐

**MAC Address**

MAC Address Assignment: **BO-VSPEX-MACs(40/40)**

[+ Create MAC Pool](#)

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

**Fabric ID:** ☒ Fabric A ☐ Fabric B ☒ Enable Failover

**VLANs**

[Filter](#) [Export](#) [Print](#)

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input checked="" type="checkbox"/>	Storage	<input checked="" type="radio"/>
<input type="checkbox"/>	VM-Data	<input type="radio"/>
<input type="checkbox"/>	vMotion	<input type="radio"/>
<input type="checkbox"/>	vSphereMgmt	<input type="radio"/>

[+ Create VLAN](#)

MTU: **9000**

**Warning**  
Make sure that the MTU has the same value in the [QoS System Class](#) corresponding to the Egress priority of the selected QoS Policy.

Pin Group: **<not set>** [+ Create LAN Pin Group](#)

**Operational Parameters**

**Adapter Performance Profile**

Adapter Policy: **VMware** [+ Create Ethernet Adapter Policy](#)

QoS Policy: **JumboMTU** [+ Create QoS Policy](#)

Network Control Policy: **<not set>** [+ Create Network Control Policy](#)

**Connection Policies**

OK Cancel

7. Create a vNIC for VM data traffic. Enter data-A for the vNIC name, choose the created MAC address pool, **Fabric A** for Fabric ID. Mark VM-Data as native VLAN, and choose VMware adapter policy.

Figure 66 Networking - Create vNIC for VM Data Traffic

**Create vNIC**

Name: **data-A**

Use vNIC Template: ☐

**MAC Address**

MAC Address Assignment: **BO-VSPEX-MACs(40/40)**

[+ Create MAC Pool](#)

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

Fabric ID: ☒ **Fabric A** ☐ Fabric B ☒ **Enable Failover**

**VLANs**

[Filter](#) [Export](#) [Print](#)

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	Storage	<input type="radio"/>
<input checked="" type="checkbox"/>	<b>VM-Data</b>	<input checked="" type="radio"/>
<input type="checkbox"/>	vMotion	<input type="radio"/>
<input type="checkbox"/>	vSphereMgmt	<input type="radio"/>

[+ Create VLAN](#)

MTU: **1500**

Pin Group: **<not set>** [+ Create LAN Pin Group](#)

**Operational Parameters**

**Adapter Performance Profile**

Adapter Policy: **VMWare** [+ Create Ethernet Adapter Policy](#)

QoS Policy: **<not set>** [+ Create QoS Policy](#)

Network Control Policy: **<not set>** [+ Create Network Control Policy](#)

**OK** **Cancel**

- Similarly, create another VNIC for Fabric B for VM data traffic. Table 7 summarizes all the VNICs created on the service profile:

Table 7 Summary of all the vNICs Created

vNIC Name	MAC Address Assignment	VLANs	Native VLANs	Fabric	MTU	Adapter Policy	QoS Policy
System-A	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	A	9000	VMware	JumboMTU
System-B	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	B	9000	VMware	JumboMTU
Storage-A*	MAC pool	Storage	Storage	A	9000	VMware	JumboMTU

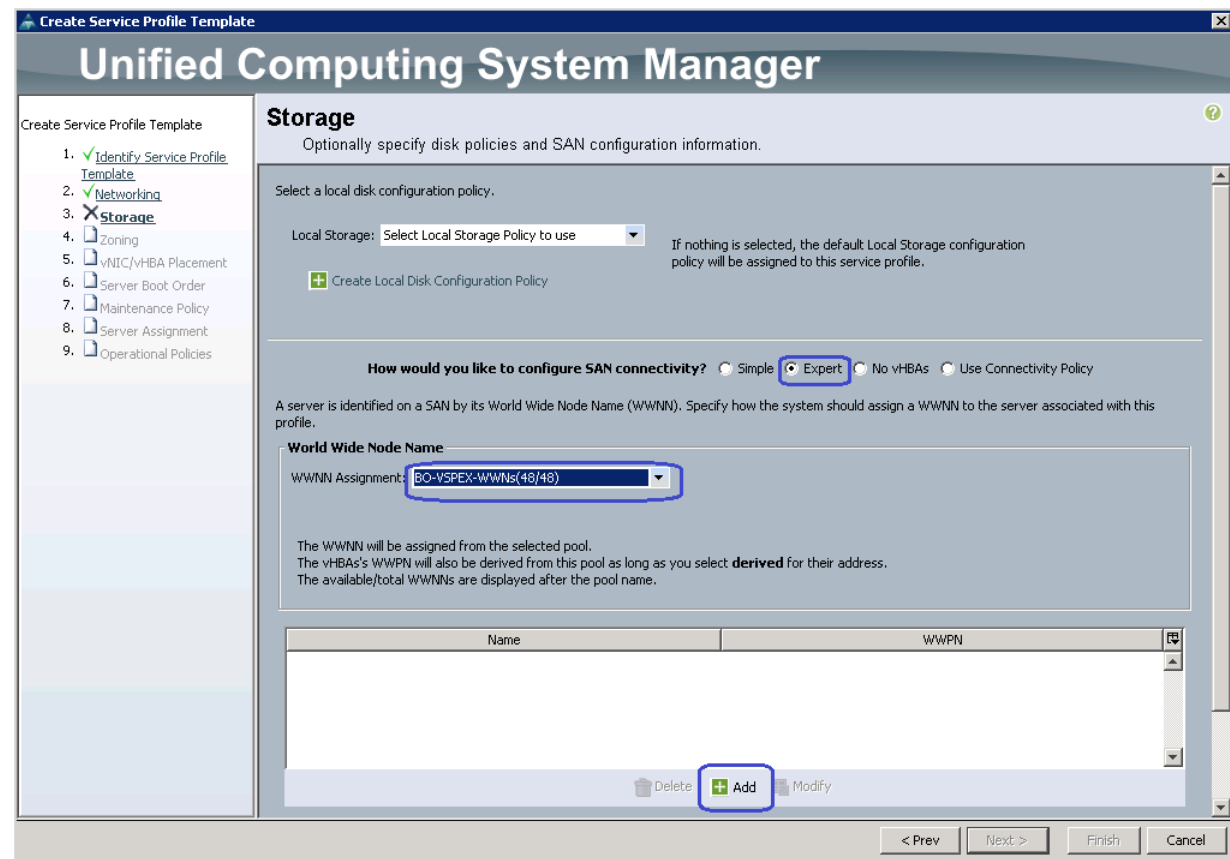
**Table 7** *Summary of all the vNICs Created*

vNIC Name	MAC Address Assignment	VLANs	Native VLANs	Fabric	MTU	Adapter Policy	QoS Policy
Storage-B*	MAC pool	Storage	Storage	B	9000	VMware	JumboMTU
Data-A	MAC pool	VM-Data	VM-Data	A	1500	VMware	-
Data-B	MAC pool	VM-Data	VM-Data	B	1500	VMware	-

\* Storage vNICs are created for NFS-variant only

- In the Storage window of the Create Service Profile Template wizard, click the **Expert** radio button for SAN connectivity and choose **VSPEX-WWNs** for WWNN pool from the drop-down menu. Click **Add** to add vHBA.

**Figure 67** *Creating Service Profile Template - Storage*



- Create a vHBA with the name vHBA-A, keep the WWPN assignment as **Derived**, select **A** as the Fabric ID, select **Storage VSAN** from drop-down menu for VSAN, and select Adapter policy as **VMWare**. Click **OK** to deploy the vHBA.

Figure 68 Storage - Creating vHBA

**Create vHBA**

Name: **vHBA-A**

Use vHBA Template: ☐

**World Wide Port Name**

WWPN Assignment: **Derived**

**+** Create vHBA Template

**+** Create WWPN Pool

If you select a WWxN Pool for the World Wide Node Name, the WWPN will be derived from that pool.  
If you did not select a WWxN Pool for the World Wide Node Name, the WWPN assigned by the manufacturer will be used.  
Note: When a manufacturer assigned WWPN is used, the WWPN will not be migrated if the service profile is moved to a new server.

Fabric ID: ☒ A ☐ B

Select VSAN: **Storage**

**+** Create VSAN

Pin Group: **<not set>**

**+** Create SAN Pin Group

Persistent Binding: ☒ Disabled ☐ Enabled

Max Data Field Size: **2048**

**Operational Parameters**

**Adapter Performance Profile**

Adapter Policy: **VMWare**

**+** Create Fibre Channel Adapter Policy

QoS Policy: **<not set>**

**+** Create QoS Policy

**OK** **Cancel**

11. Repeat step 10 for vHBA-B on Fabric B, with the same configuration. Click **Next**.
12. Keep the default configuration for Zoning and vNIC/vHBA Placement policy. Click **Next**.
13. In the Server Boot Order window, choose **SAN-Boot** for the Boot policy from the drop-down menu. Click **Next**.

**Figure 69** *Creating Service Profile Template - Server Boot Order*

**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 Boot Order

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **SAN-Boot** + Create Boot Policy

Name: **SAN-Boot**  
 Description:  
 Reboot on Boot Order Change: **Yes**  
 Enforce vNIC/vHBA/iSCSI Name: **Yes**  
 Boot Mode: **Legacy**

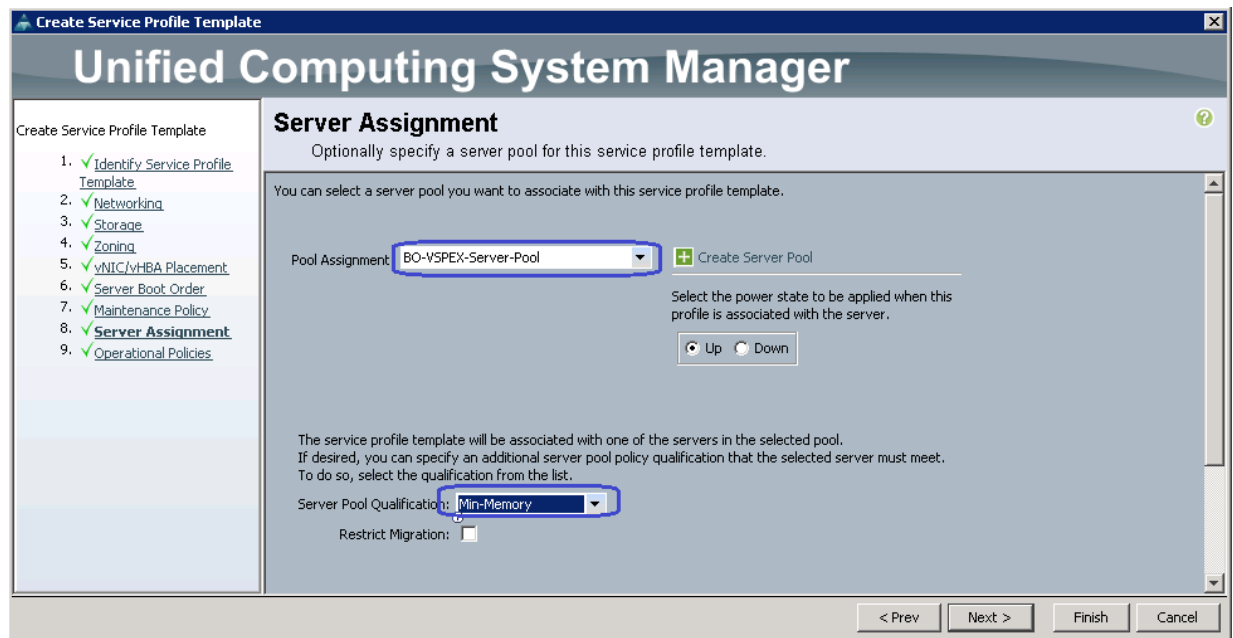
**WARNINGS:**  
 The type (primary/secondary) does not indicate a boot order presence.  
 The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.  
 If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, 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
Local CD/DVD	1				
San	2				
SAN primary		vHBA-A	Primary		
SAN Target primary			Primary	0	50:06:01:64:08:E0:03:68
SAN secondary		vHBA-B	Secondary		
SAN Target primary			Primary	0	50:06:01:6C:08:E0:03:68

< Prev **Next >** Finish Cancel

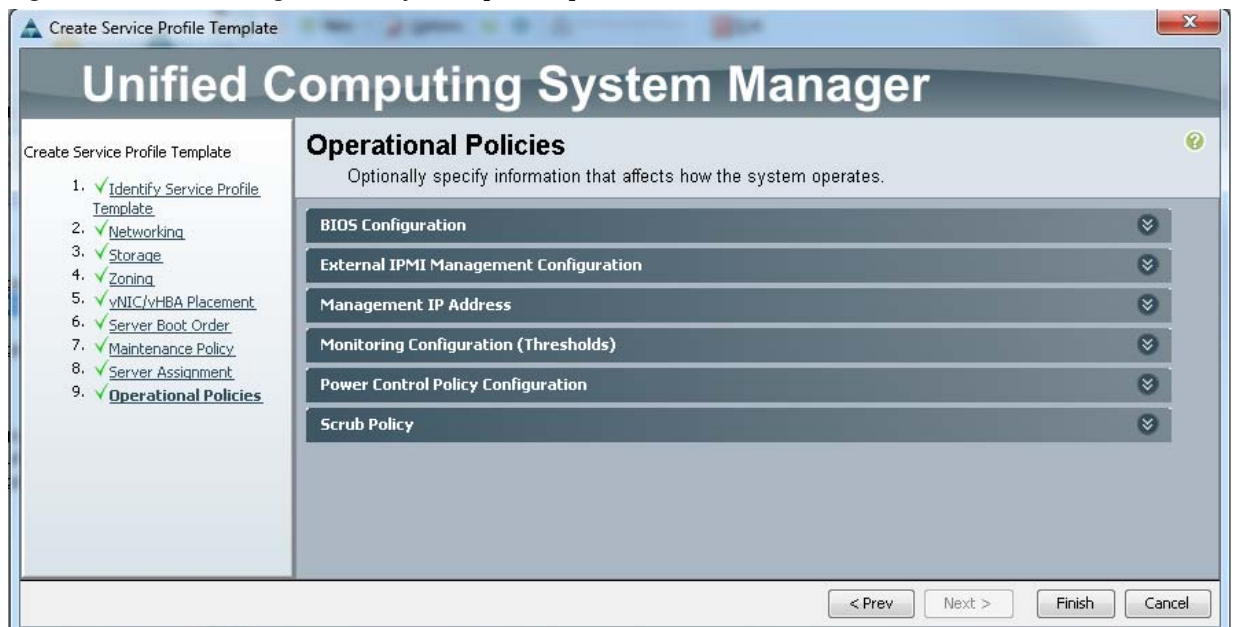
14. Click **Next** to go to the Maintenance Policy window. Keep all other fields at default and click **Next**. In the Server Assignment window, assign the server to Pool Assignment, and choose the created Server Pool and also select the Server Pool Qualification from drop-down. Click **Next**.

**Figure 70** *Creating Service Profile Template - Server Assignment*



15. In the Operation Policies window, keep all the fields at default, and click **Finish** to deploy the Service Profile Template.

**Figure 71** *Creating Service Profile Template - Operational Policies*



That concludes the service profile template creation.

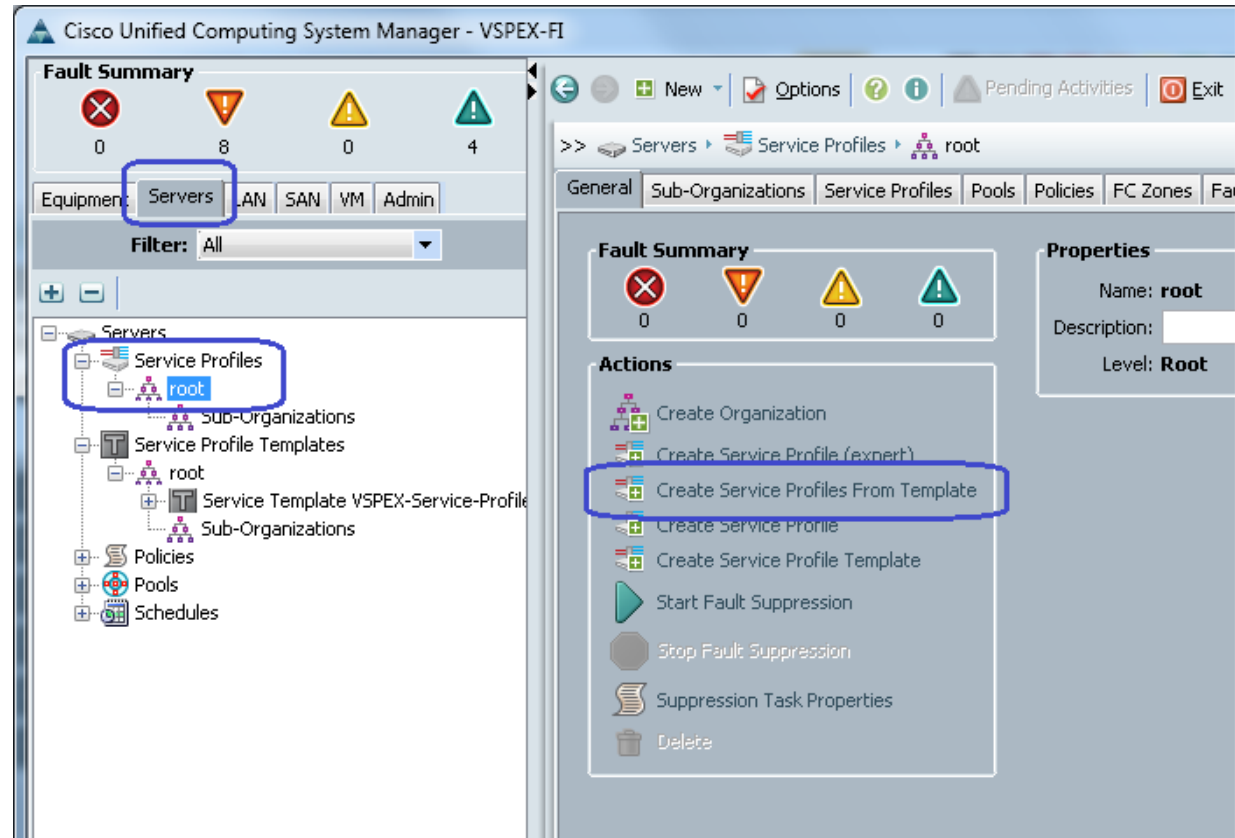
## Instantiate service profiles from the service profile template

As a final step to configure UCSM, we will instantiate service profiles from the service profile template. Follow these steps to instantiate Service Profiles:



16. From the **Servers** tab, expand **Servers > Service Profiles > root**, and click **Create Service Profile from Template** on the right pane.

Figure 72 Creating Service Profile from Template



17. Provide naming prefix, number of service profile Instances to be instantiated and choose the service profile template from the drop-down menu. Refer to the sizing guidelines for the number of servers needed for your deployment.

**Figure 73** *Creating Service Profiles from Template Window*

**Create Service Profiles From Template**

Naming Prefix: **BO-VSPEX-Server-**

Name Suffix Starting Number: **1**

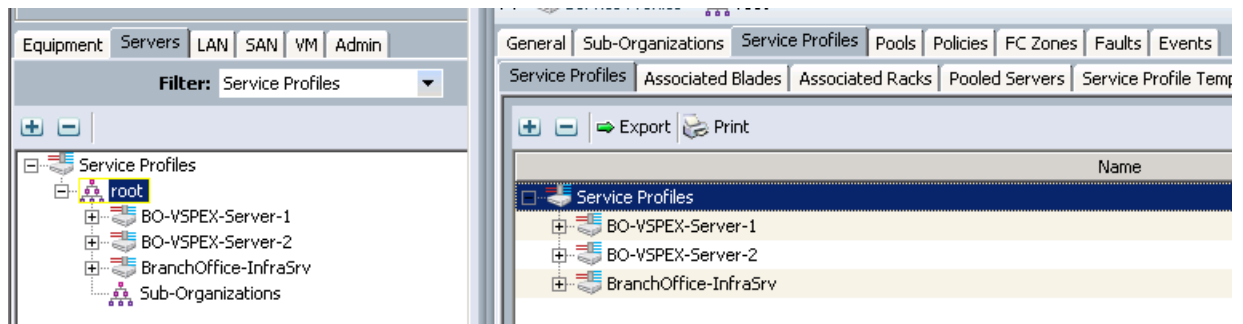
Number of Instances: **2**

Service Profile Template: **Service Template BO-VSPEX-Servers-SP-Template**

OK Cancel

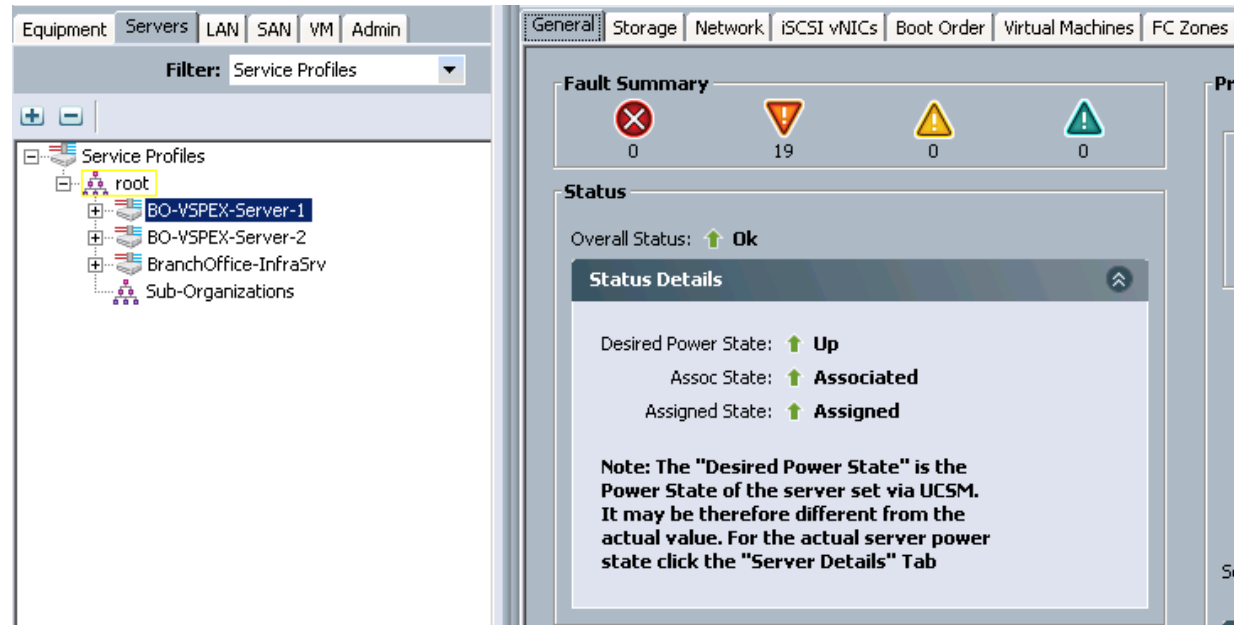
18. According to the Service profile instances you will see the Service profiles created from the template.

**Figure 74** *Service Profiles Created from Service Profile Template*



19. As Service Profile Template is assigned to a server pool, the service profiles instantiated from the template would be assigned to individual server resource from the server pool as far as they are available. You can select a service profile and see its association state, and with which server it is associated.

**Figure 75**



20. Eventually, all four servers would be associated – you can also see the summary by clicking on the **Servers** under the **Equipment** tab.

**Figure 76** Created Servers Showing Overall Status

Name	Overall Status	Model	Profile	Us...	Cores	Cores Enabled	Thre...
Server 1	Ok	Cisco UCS B200 M3	org-root/ls-BO-VSPEX-Server-1	16	16	32	
Server 2	Ok	Cisco UCS B200 M3	org-root/ls-BO-VSPEX-Server-2	16	16	32	
Server 3	Unassociated	Cisco UCS B200 M3	...	12	12	24	
Server 4	Ok	Cisco UCS B200 M3	org-root/ls-BranchOffice-InfraSrv	16	16	32	



**Note**

We have not yet carved out specific data store to install ESXi hypervisor OS image on the VNXe storage array. We need specific WWPN and WWNN addresses to allow access to the data store, and hence we needed to configure the Service Profile before we can carve out the space for each ESXi server on the storage pool.

## Configure Data Stores for ESXi Images

This section walks you through the steps to create FC accessible data stores for the ESXi boot image per server basis. This includes four steps:

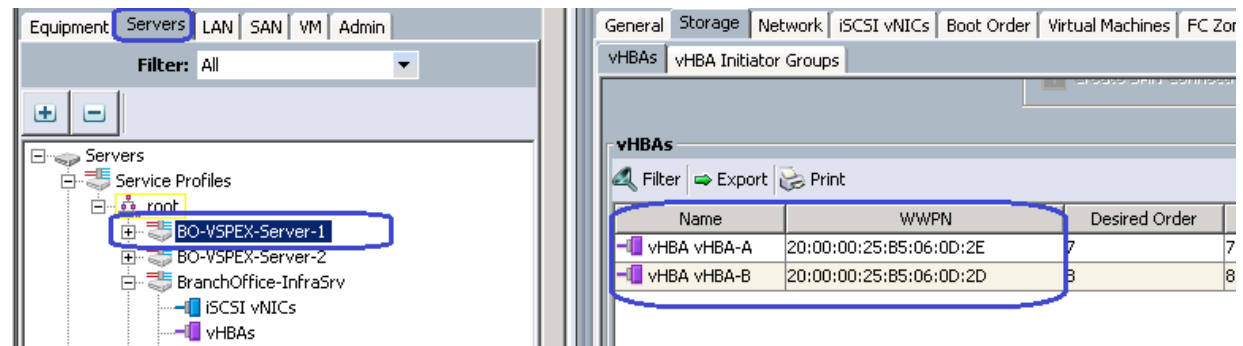
1. Host Initiator discovery on VNXe Array
2. Create Host
3. Configure Storage Pool
4. Create Boot LUNs and Configure Host Access.

## Host Initiator Discovery on VNXe Array

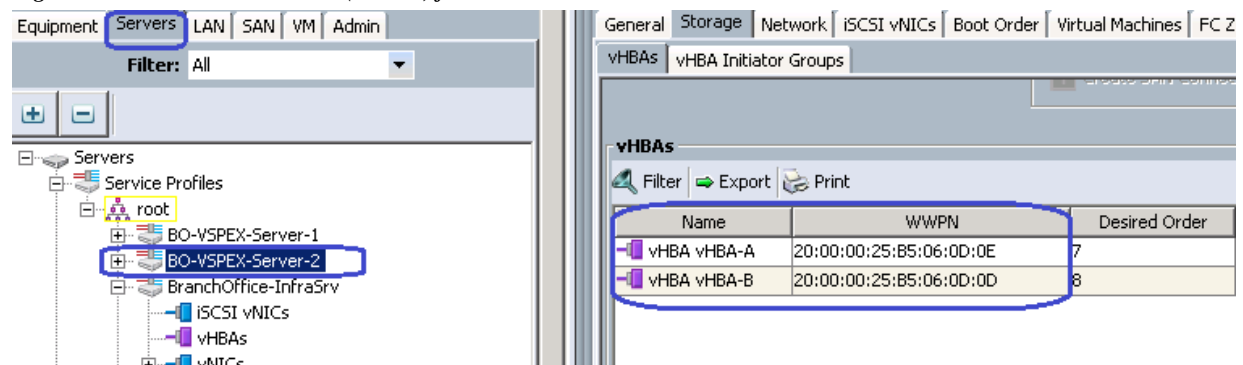
Once the UCS service profile is associated on the UCS Blades. Make sure you power on all the blades in order to discover the Host initiators (vHBA-A & vHBA-B) on the VNXe Array.

Click **Servers > Service Profiles > root**. Select the Service Profile and on the right pane click **Storage > vHBAs**.

**Figure 77** Host Initiator (vHBAs) from VSPEX Server1



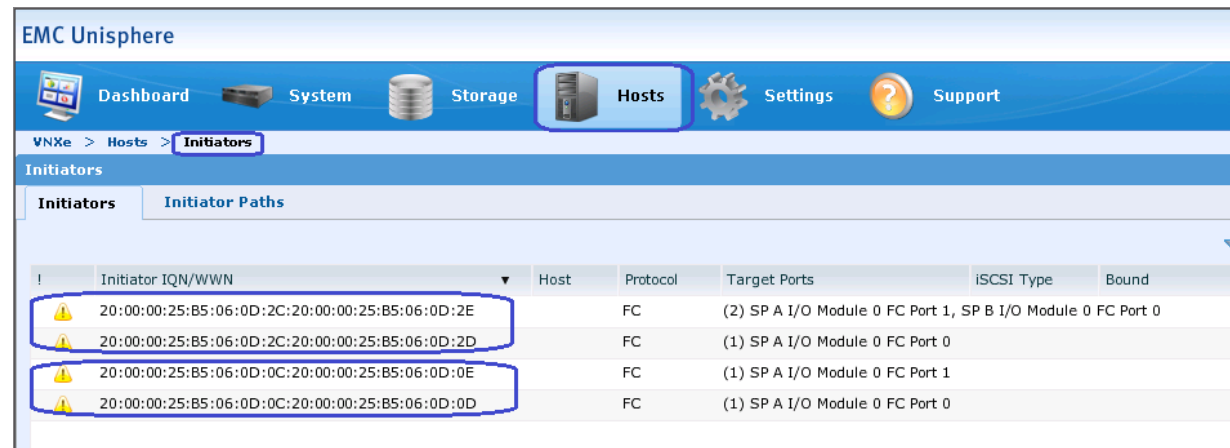
**Figure 78** Host Initiator (vHBAs) from VSPEX Server2



## Create Host

To see VSPEX Servers Host Initiator Discovery on VNXe Array, Launch Unisphere and click **Hosts** and then click **Initiators**.

**Figure 79** *EMC Unisphere Showing Host Initiators*



## Configure Storage Pool

Follow these steps to create storage pool and carve boot LUNs from that on per server basis.

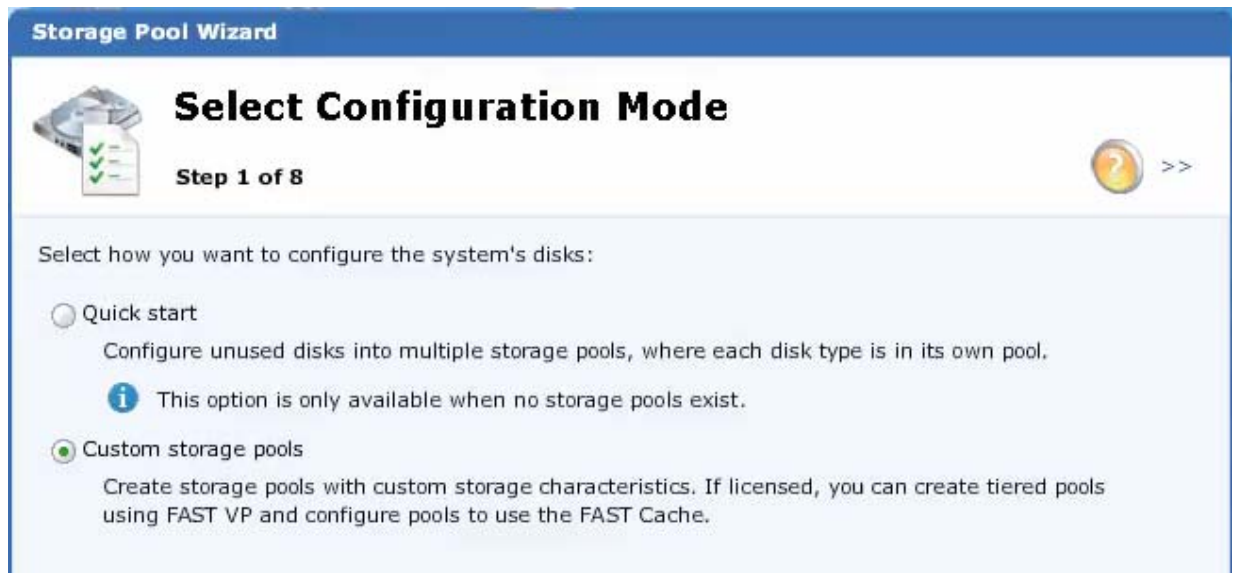
1. Connect to EMC VNXe Unisphere GUI, click the **Storage** tab. Select **Storage Configuration** and **Storage Pools** and then click **Create**.

**Figure 80** *Creating Storage Pools*



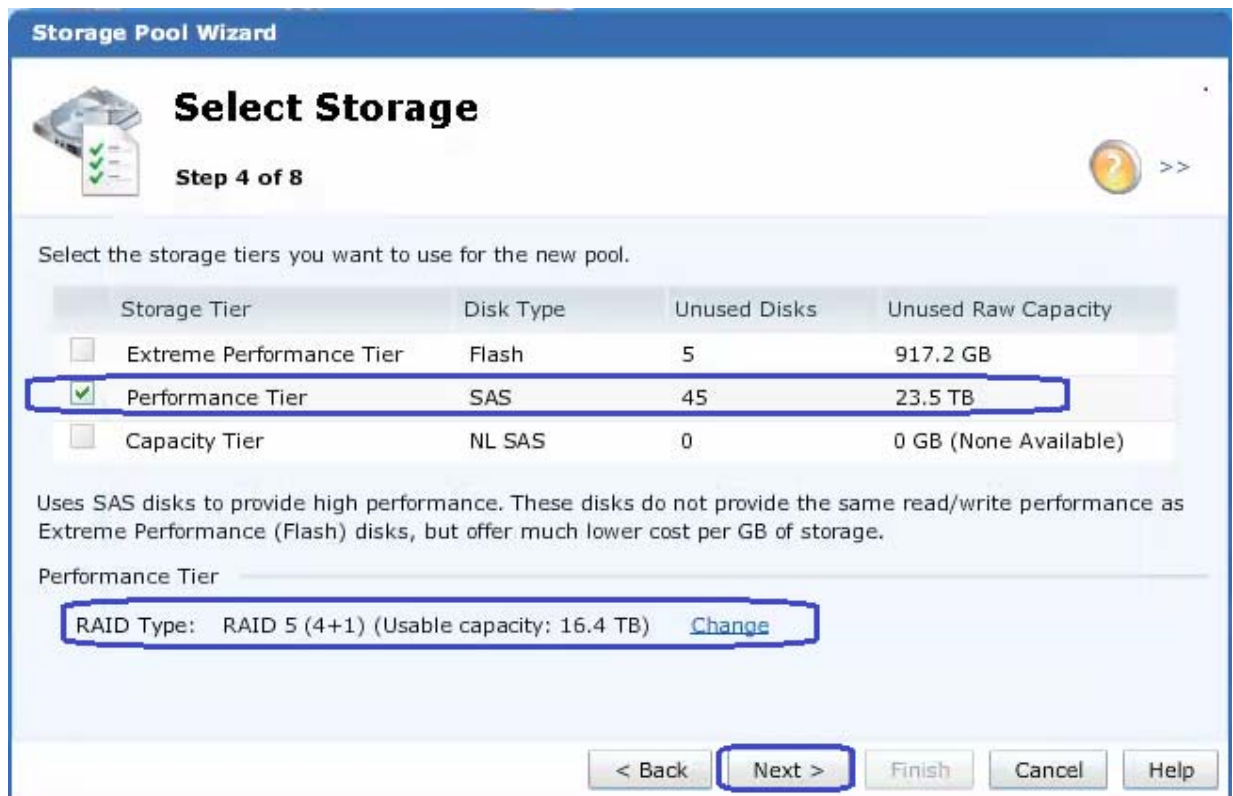
2. Click the **Custom Storage Pools** radio button. Specify the Pool name for the Storage Pool and skip the FAST VP and click **Next**.

Figure 81 Storage Pool - Configuration Mode



3. Check the Performance Tier check box for using SAS drives for SAN boot LUNs. And keep the default RAID Type: RAID 5 (4+1).

Figure 82 Storage Pool - Select Storage



4. From Performance Tier, Choose 5 of 45 disks (2.0TB) for 600GB SAS disks from the drop-down list and click **Next**. Click **Finish** to confirm the Storage Pool creation.

Figure 83 Storage Pool - Storage Size

**Storage Pool Wizard**

**Select Amount of Storage**

Step 5 of 7

Select the amount of storage for each selected tier. The number of disks you can choose is based on the RAID configuration selected. The maximum number of disks you can configure will ensure that enough disks are kept unused to satisfy the hot spare policy.  
[More information](#)

Performance Tier

600 GB (10K RPM) SAS Disks: Use 35 of 45 disks (14.6 TB) ▼

Use none of the 45 disks ▲

Use 5 of 45 disks (2.0 TB)

Use 10 of 45 disks (4.1 TB)

Use 15 of 45 disks (6.2 TB)

Use 20 of 45 disks (8.3 TB) ▼

Total Disks to Configure: 35

Total Usable Capacity: 14.6 TB

< Back Next > Finish Cancel Help

- After the successful Storage Pool creation, you will see the Storage Pool for ESXi Boot.

Figure 84 Created Storage Pool for ESXi Boot

EMC Unisphere

Dashboard System Storage Hosts Settings Support

VNXe > Storage > Storage Configuration > Storage Pools

Storage Pools

List View Graph View

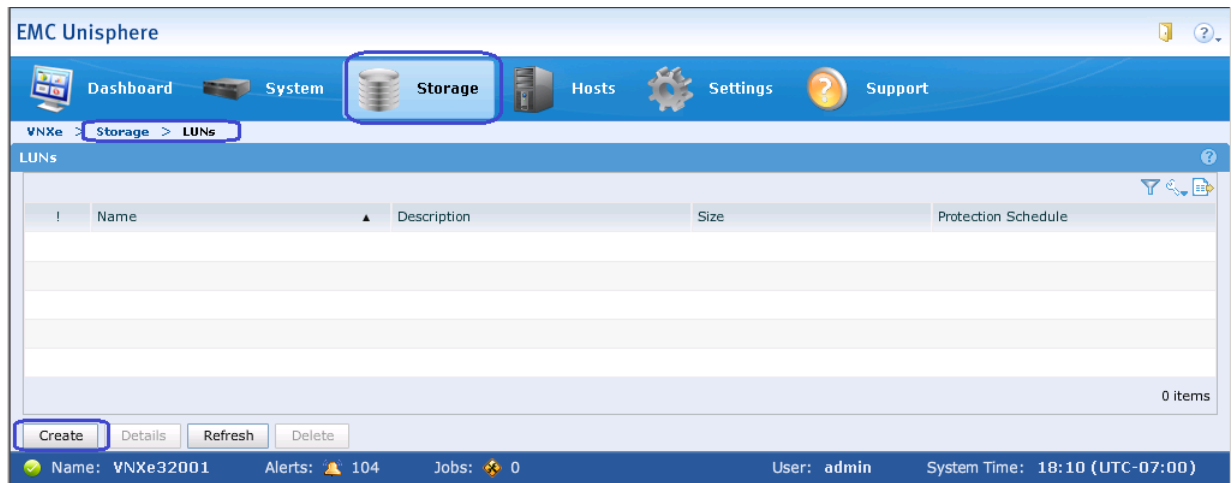
	Name	Total Space	Used Space	Percent Used	Available Space
	StoragePool-ESXi-Boot	2.0 TB	4.2 GB	0%	2.0 TB

## Create Boot LUNs and Configure Host Access

Follow these steps to create storage pool and carve boot LUNs from that on per server basis.

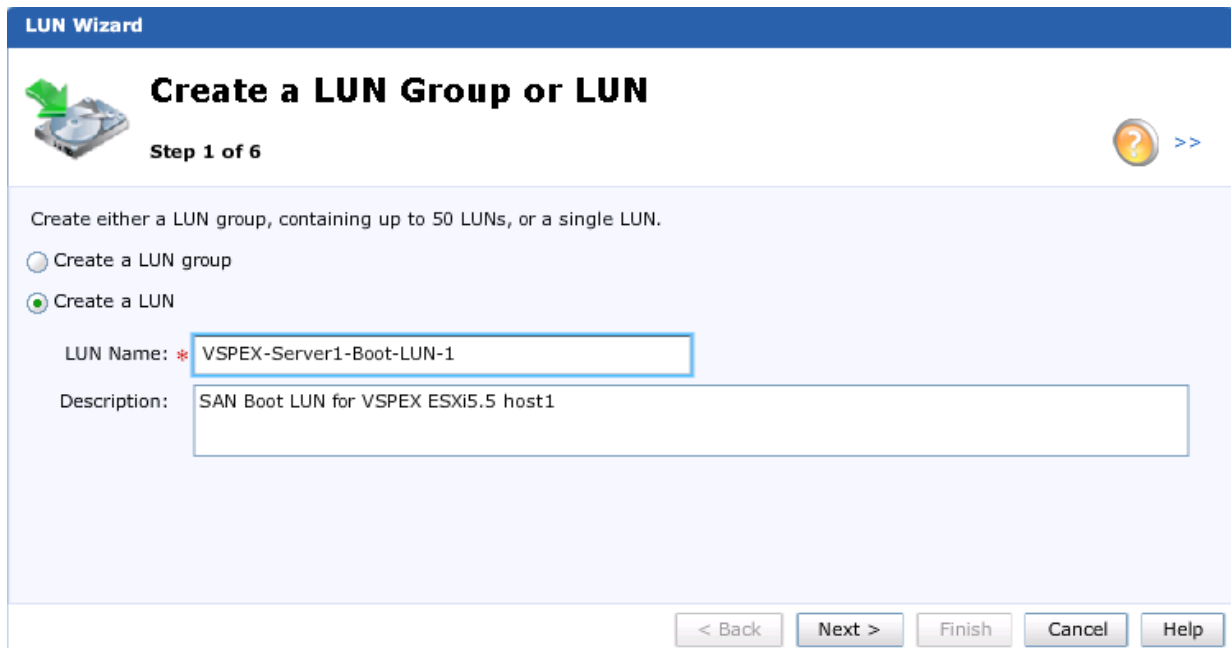
- launch EMC VNXe Unisphere GUI, click the **Storage** tab. Click **LUNs** and **Create**.

Figure 85 Creating LUNs



- Click the **Create a LUN** radio button. And specify the LUN name for the VSPEX ESXi host1 Boot LUN and click **Next**.

Figure 86 Creating LUN - Define LUN



- Choose the created Storage Pool for ESXi Boot LUN. Specify the LUN size and uncheck the check box **Thin** and click **Next**.



Figure 87 Creating LUN - Configure a LUN

**LUN Wizard**

**Configure a LUN**

Step 2 of 6

Configure the LUN's storage characteristics:

Storage Pool: StoragePool-ESXi-Boot (SAS, 2.1 TB free)

Tiering Policy: Start High Then Auto-Tier (Recommended)

*The selected pool is not tiered. The tiering policy will have no effect on the storage resource.*

Size: 100 GB ☐ Thin

< Back Next > Finish Cancel Help

4. Select **Do not configure a snapshot schedule** and click **Next**. (It is not recommended to configure snapshot for Boot LUNs)
5. In the LUN wizard, In the Configure Host Access window, choose LUN from the drop-down list only for the VSPEX ESXi host1 and click **Next**.



**Note**

Boot LUN is specific to single ESXi host, so we don't provide access to VSPEX ESXi host2 for the same Boot LUN.

Figure 88 Creating LUN - Configure Host Access

**LUN Wizard**

**Configure Host Access**

Step 4 of 6

Configure which hosts will access this storage:

Filter for:  Protocols: FC or iSCSI

!	Name	Network Address	Operating System	Protocol	Access
✓	BO-VSPEX-Serv...	10.29.180.211	VMware ESX	FC, File	LUN
✓	BO-VSPEX-Serv...	10.29.180.212	VMware ESX	FC, File	No Access

Filtered: 2 of 2

< Back **Next >** Finish Cancel Help

- Confirm the summary of the LUN configuration and click **Finish** to complete the LUN creation and host access. Repeat steps 1 to 5 to create Boot LUN for VSPEX ESXi host2.

Figure 89 Creating LUN - Summary

**LUN Wizard**

**Summary**

Step 5 of 6

Confirm the following LUN configuration:

Name: VSPEX-Server1-Boot-LUN-1

Description: SAN Boot LUN for VSPEX ESXi5.5 host1

Storage Pool: StoragePool-ESXi-Boot

Size: 100.0 GB

Thin: No

Tiering Policy: Start High Then Auto-Tier (Recommended)

Protection Schedule: None configured

LUN Access: ▼ 1 hosts configured  
BO-VSPEX-Server-1

Snapshot Access: No hosts configured

< Back **Next >** **Finish** Cancel Help

- Repeat the steps 1 to 6 to create Boot LUN for VSPEX ESXi host2.

**Figure 90** Created Boot LUNs for Server 1 and Server 2



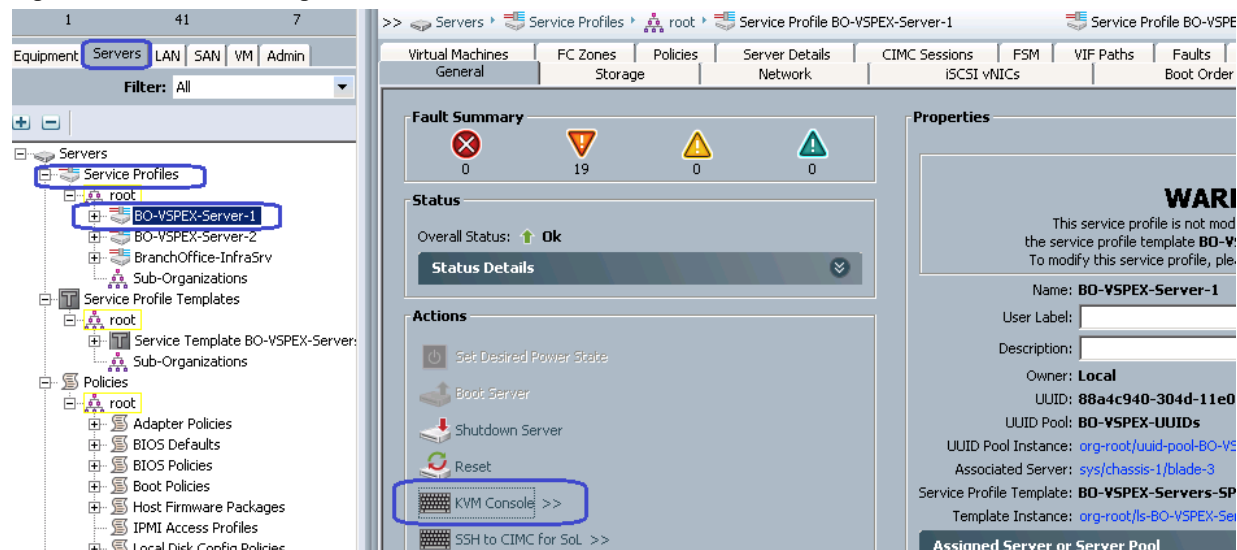
At this point, we have end-to-end FC storage access from servers in UCS to the specific boot LUN on the VNxe storage devices. We are ready to install ESXi images on the UCS server.

## Install ESXi servers and vCenter Infrastructure

Follow these steps to install ESXi5.5 image on the UCS servers:

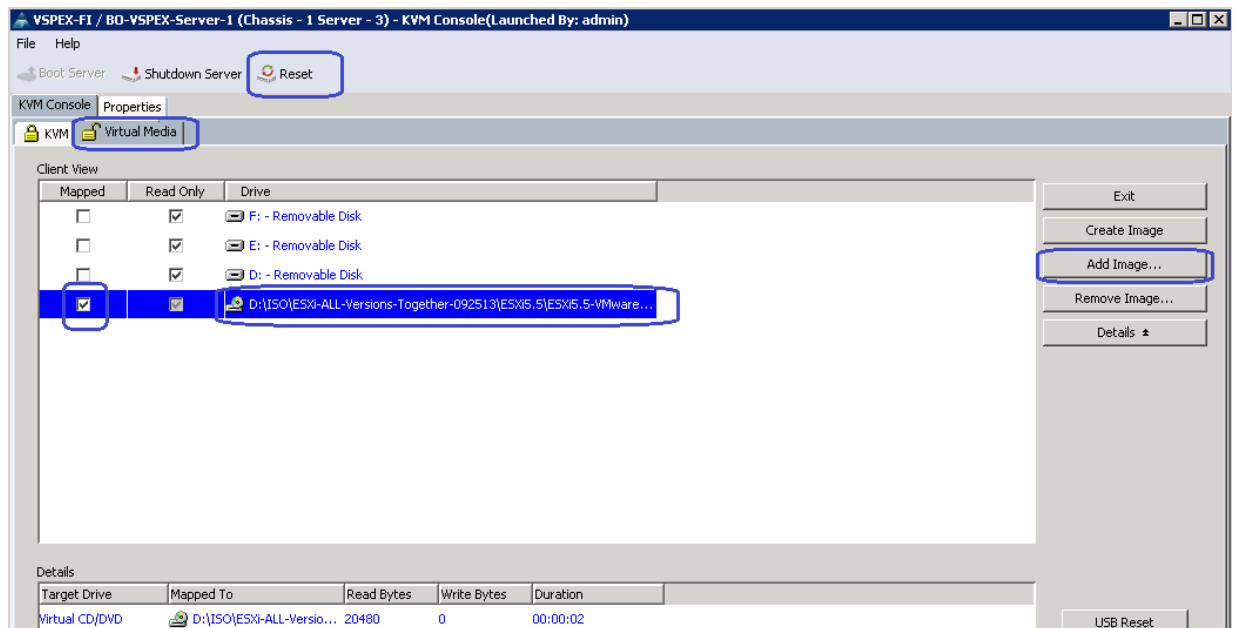
- From UCSM GUI, click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a service profile. Click **KVM Console** on right pane.

**Figure 91** Launching KVM Console



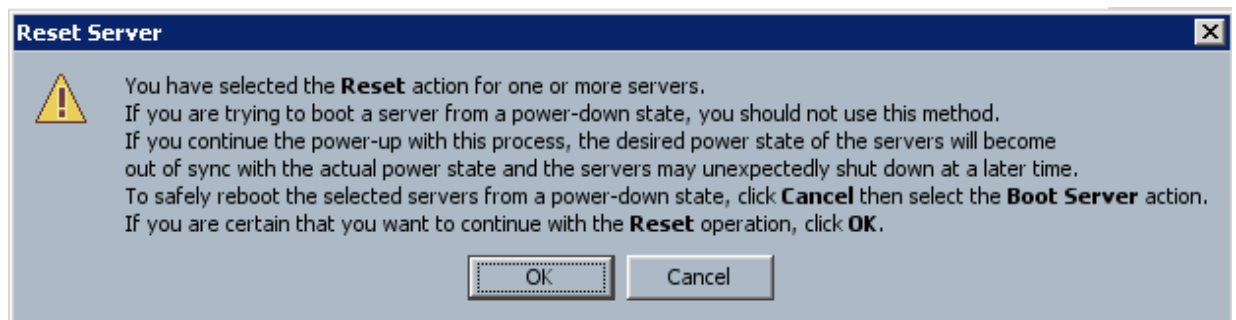
- Once the Java applet of KVM is launched, click the **Virtual Media** tab and click **Add Image**. That would open a dialog box to select an ISO image. Traverse the local directory structure and select ISO image of the ESXi 5.5 hypervisor installer media. Once the ISO image shows up in the list, check the **Mapped** check box and reset the server and click **Reset**.

**Figure 92** Adding ESXi 5.5 ISO Image



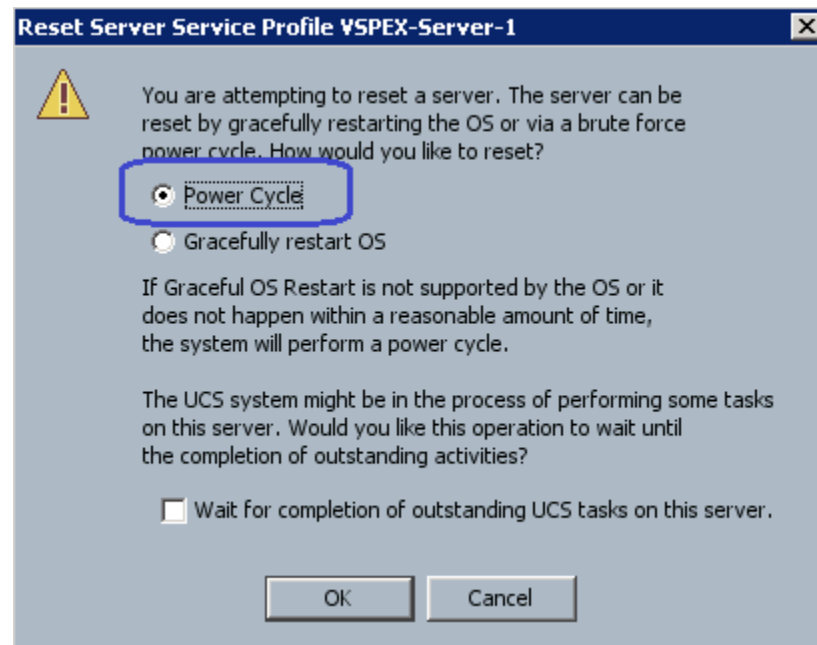
3. Click **OK** in the Reset Server warning message pop-up window.

**Figure 93** Warning Message to Reset Server




4. Click the **Power Cycle** radio button and click **OK**.

**Figure 94** *Selecting the Reset Type*




5. Click the **KVM** tab to view the boot process. During the Boot process, Press **F6** for the Boot Selection window to appear and verify that the Cisco Virtual KVM-mapped vDVD is the 1<sup>st</sup> boot device and Storage LUNs as 2<sup>nd</sup> Boot device.

**Figure 95** *Selecting the Order of the Boot Device*



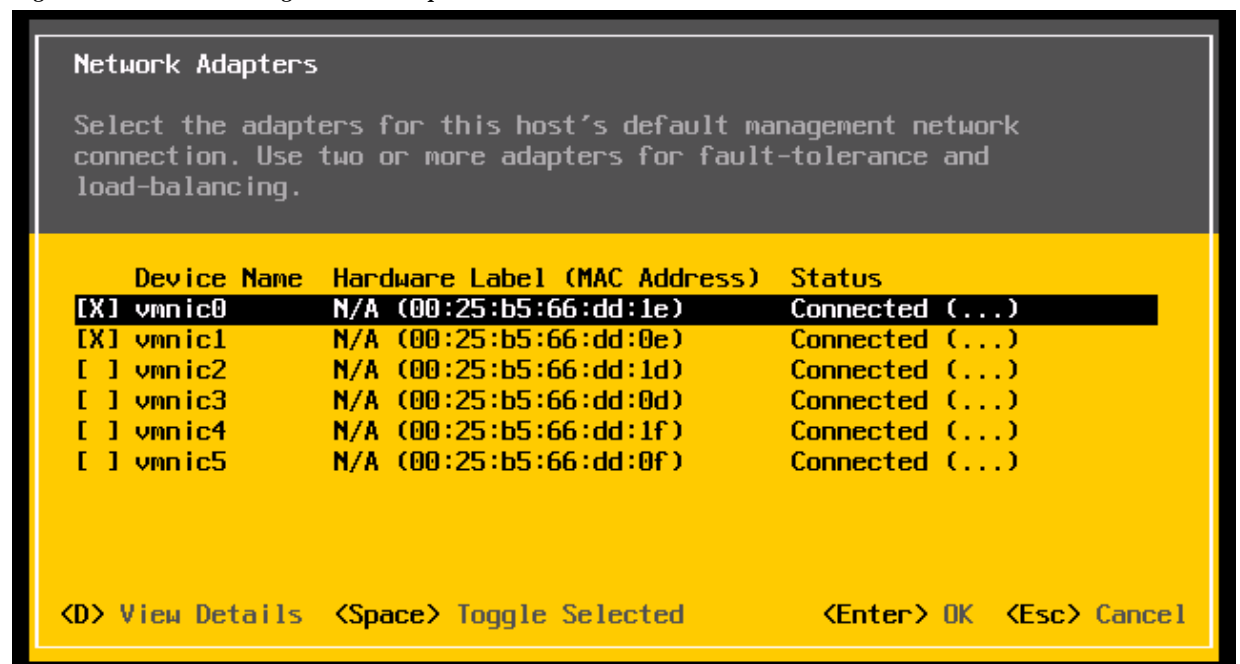
6. At this point of time, ESXi5.5 installation media would boot from the virtual disk mounted on the KVM. Follow these steps to install ESXi 5.5 hypervisor on the boot LUN. Make sure that you select the Storage boot LUN we created and not the local disk to install the hypervisor image. You can select all the default parameters or as per your requirements.

Figure 96 Selecting a Disk to Install or Upgrade



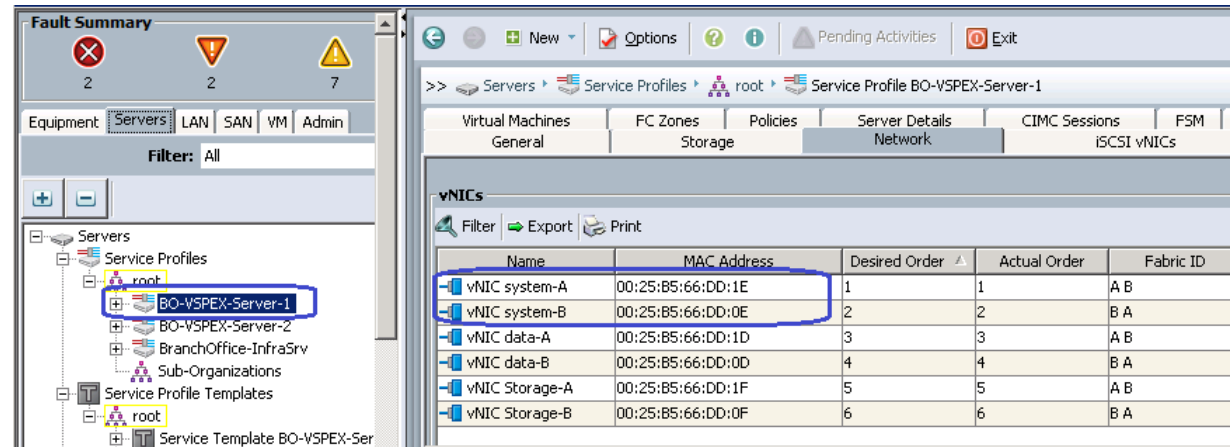
- Once the ESXi is installed, login the system by pressing **F2** on the KVM window. You need to configure basic management network for the ESXi host. Make sure that you have selected two system vNICs.

Figure 97 Selecting Network Adapters



Easiest way to figure out which vmnic adapter should be used for the vSphere management purpose, you can identify the vmnic by MAC address. The MAC addresses of the VNICs (vmnics) are summarized on the following UCSM GUI window. Click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a particular service profile and click the **VNICs**. The VNIC names and MAC addresses are listed on the right pane of the window.

**Figure 98** Created vNICs for System-A and System-B



Repeat the ESXi installation steps for the remaining VSPEX server2 similarly.

## VMware vCenter Server Deployment

This section describes the installation of VMware vCenter for VMware environment and to get the following configuration:

- A running VMware vCenter virtual machine
- VMware DRS and HA functionality enabled.

For detailed information on installing a vCenter Server, see the link:

[http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2032885](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2032885)

Following steps provides high level configuration steps to configure vCenter server:

### 1. Create the vCenter host VM

If the VMware vCenter Server is to be deployed as a virtual machine on an ESXi server installed as part of this solution, connect directly to an Infrastructure ESXi server using the vSphere Client. Create a virtual machine on the ESXi server with the customer's guest OS configuration, using the Infrastructure server datastore presented from the storage array. The memory and processor requirements for the vCenter Server are dependent on the number of ESXi hosts and virtual machines being managed. The requirements are outlined in the vSphere Installation and Setup Guide.

### 2. Install vCenter guest OS

Install the guest OS on the vCenter host virtual machine. VMware recommends using Windows Server 2012 R2 SP1. To ensure that adequate space is available on the vCenter installation drive, see vSphere Installation and Setup Guide.

### 3. Install vCenter server

Install vCenter by using the VMware VIM Setup installation media. Easiest method is to install vCenter single sign on, vCenter inventory service, and vCenter server using Simple Install. Use the customer-provided username, organization, and vCenter license key when installing vCenter.

4. Apply vSphere license keys

To perform license maintenance, log into the vCenter Server and select the Administration - Licensing menu from the vSphere client. Use the vCenter License console to enter the license keys for the ESXi hosts. After this, they can be applied to the ESXi hosts as they are imported into vCenter.

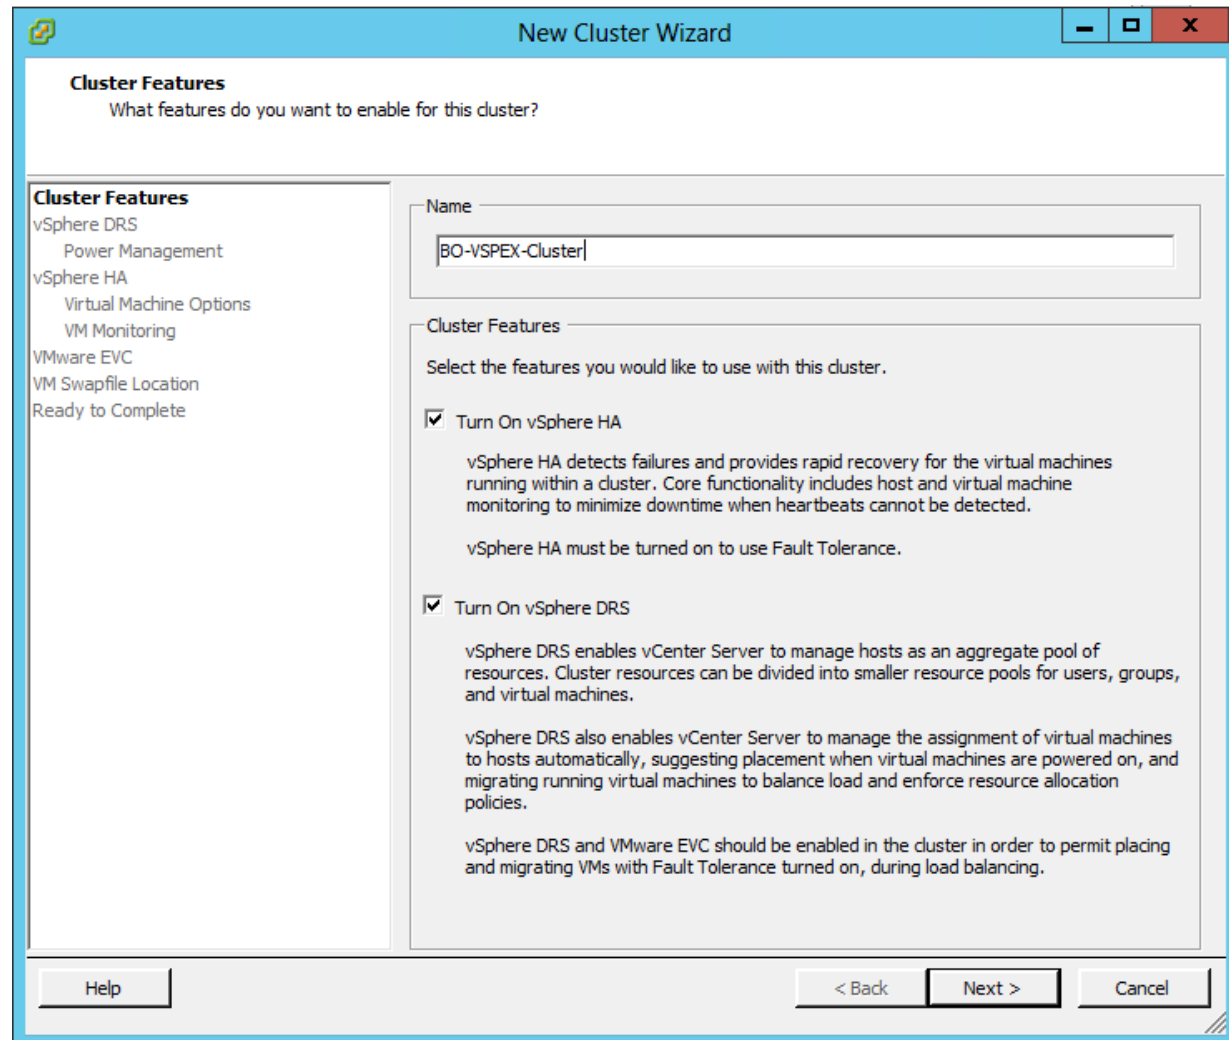
## Configuring cluster, HA and DRS on the vCenter

Perform the following steps to add all the VMware on virtual machine vCenter:

1. Log into VMware ESXi Host using VMware vSphere Client.
2. Create a vCenter Datacenter.
3. Create a new management cluster with DRS and HA enabled.
  - a. Right-click on the cluster and, in the corresponding context menu, click **Edit Settings**.
  - b. Select the checkboxes Turn On vSphere HA and Turn On vSphere DRS, as shown in the [Figure 99](#).
  - c. Click **OK**, to save changes.

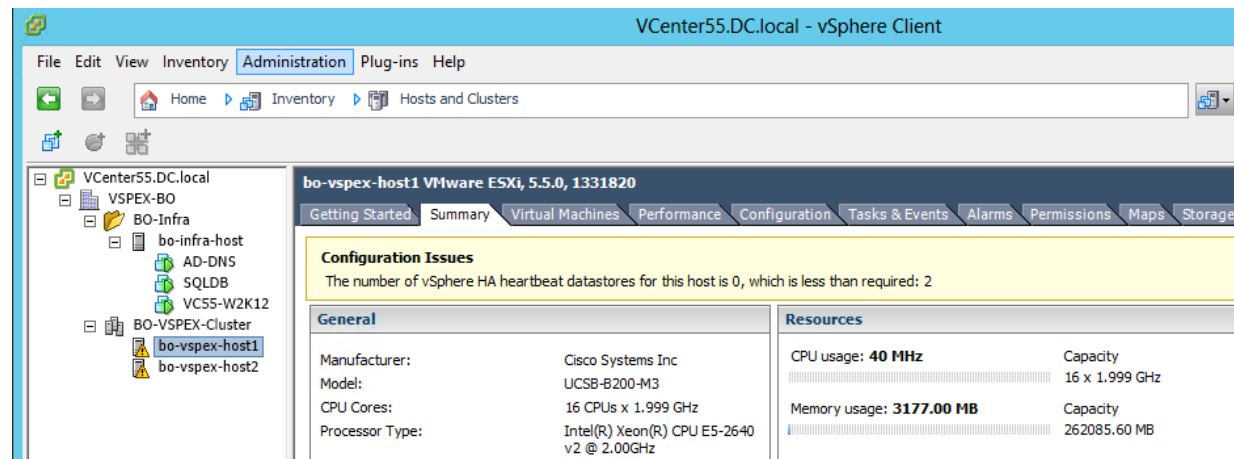


Figure 99 Configuring HA and DRS on Cluster



- d. Add all ESXi hosts to the cluster by providing servers management IP addresses and login credentials one by one.

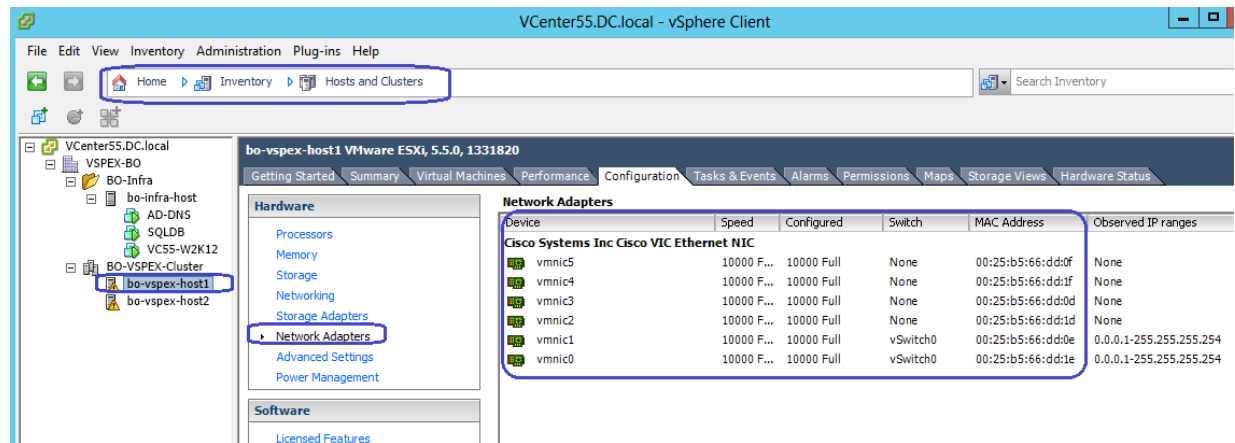
Figure 100 Adding ESXi Hosts to the Cluster



## Virtual Networking Configuration

In UCSM service profile, we created six vNICs per server for NFS-Variant. This shows up as six network adaptors or vmnics in ESXi server. You can see these adapters in the vCenter by selecting **Home > Inventory > Hosts and Clusters** view, select a server, click the **Configuration** tab on the right pane and click **Network Adapters**.

**Figure 101** Network Adapters Created Through UCS Manager



Using [Figure 101](#) and [Figure 98](#) from UCSM, it would be a good idea to make a table like following on per ESXi host basis.

For Example, below we have shown the ESXi host1 identity such as UCS Blade vNIC, vSphere vmnic & corresponding VLAN ID.

**Table 8** Service Profile vNIC and vSphere vmnic Relations

UCS Manager vNICs	vSphere NIC	MAC Address	VLAN IDs
System-A	vmnic0		11 for vSphereMgmt and 40 for vMotion
System-B	vmnic1		11 for vSphereMgmt and 40 for vMotion
Storage-A	vmnic2		20 for Storage
Storage-B	vmnic3		20 for Storage
Data-A	vmnic4		30 for VMData
Data-B	vmnic5		30 for VMData



### Note

You can use the virtual switching strategy with any variant of architecture.

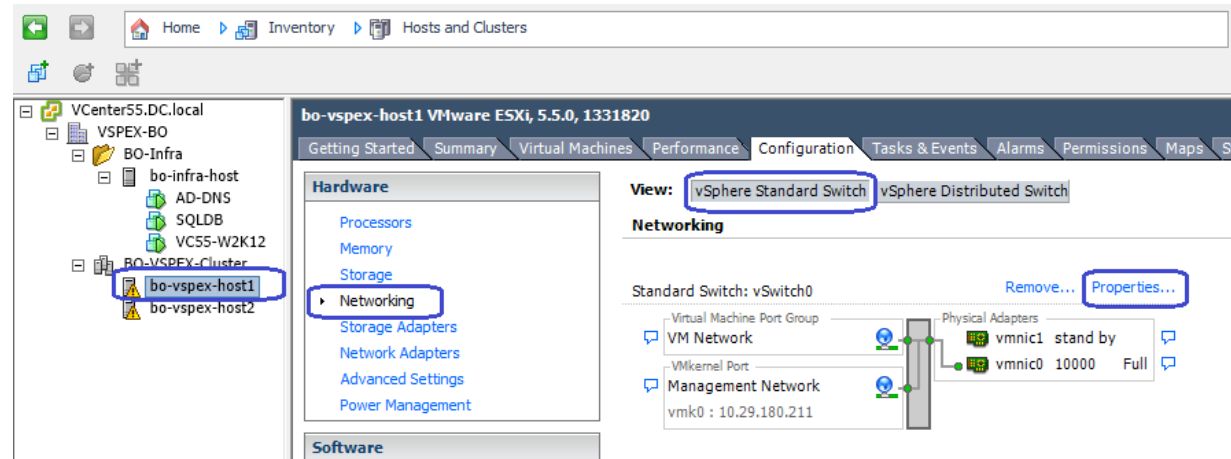
We would create 3 native Standard default virtual switches for virtual network configuration as follows:

1. vSwitch0 – For management traffic and vMotion traffic
2. vSwitch1 – For VM data traffic
3. vSwitch 2 – For NFS Storage traffic

Each vSwitch listed above would have two vmnics, one on each fabric for load balancing and high-availability. Also, for vMotion traffic, jumbo MTU needs to be configured in virtual network too. Follow these steps to configure the two vSwitches to achieve that goal.

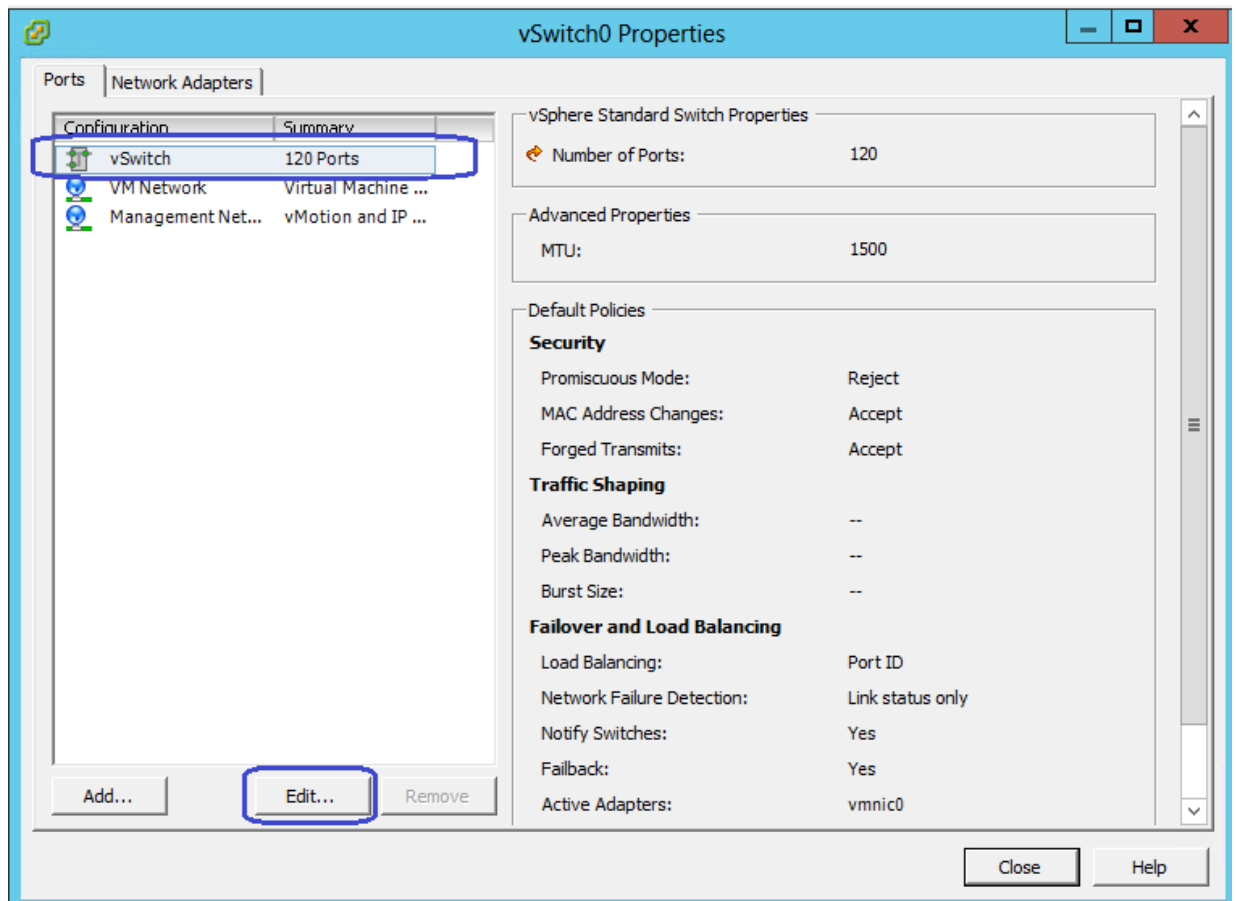
1. Choose **Home > Inventory > Hosts and Clusters** panel on vCenter, expand the VSPEX cluster and select an ESXi host. Click the **Configuration** tab > **Networking** > **Properties** to view the properties of vSwitch0.

**Figure 102** Properties of vSwitch



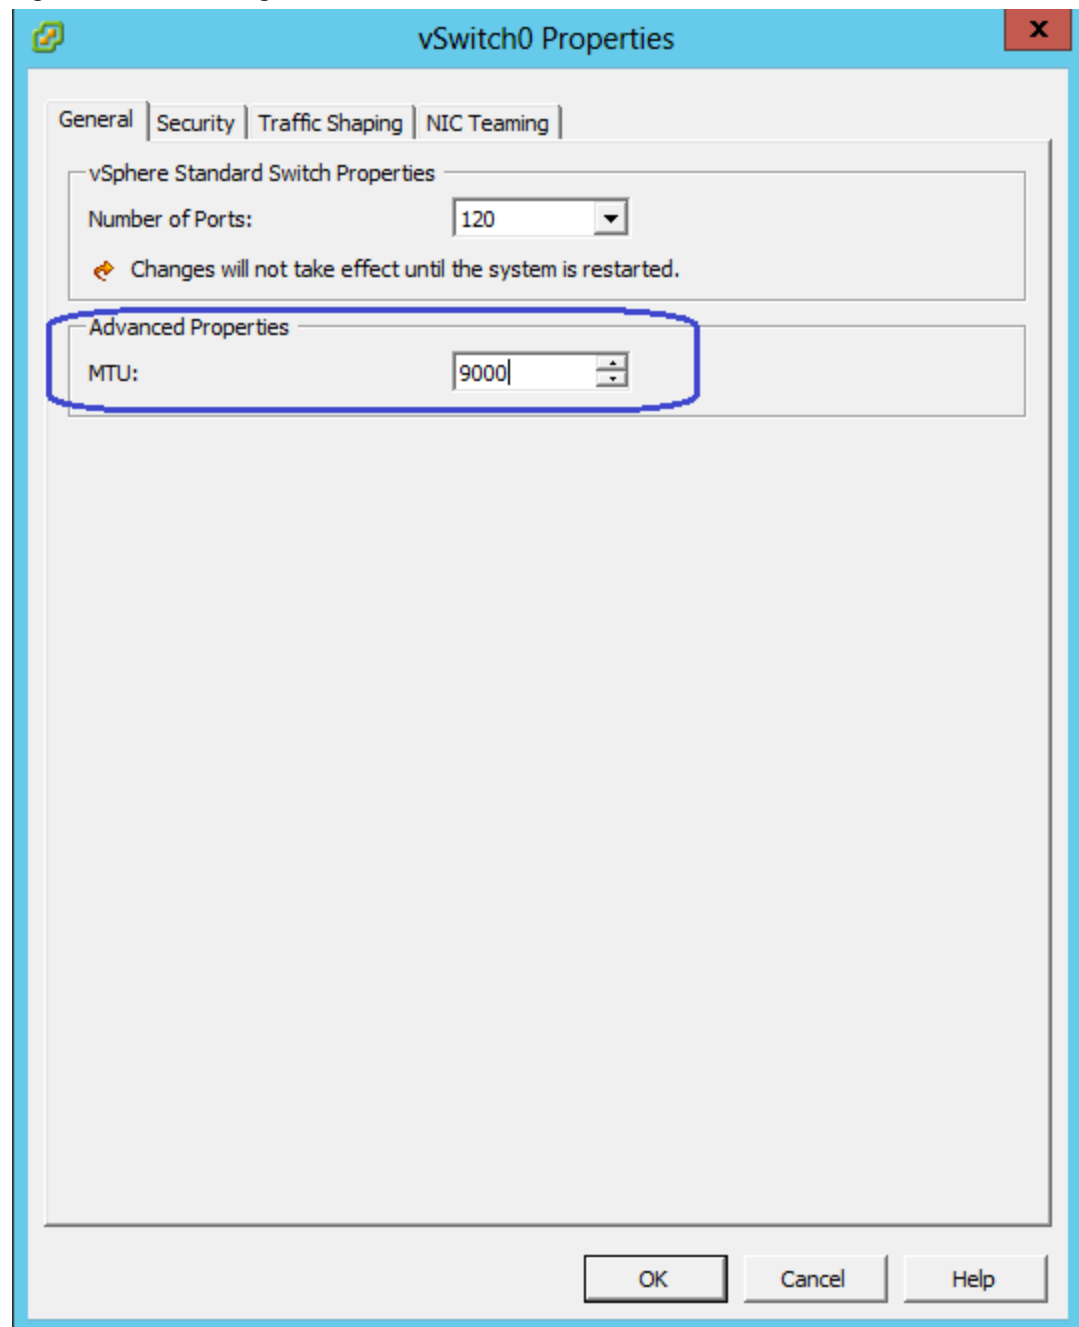
2. Select the vSwitch and then click **Edit**.

**Figure 103**      *Changing the vSwitch Properties*



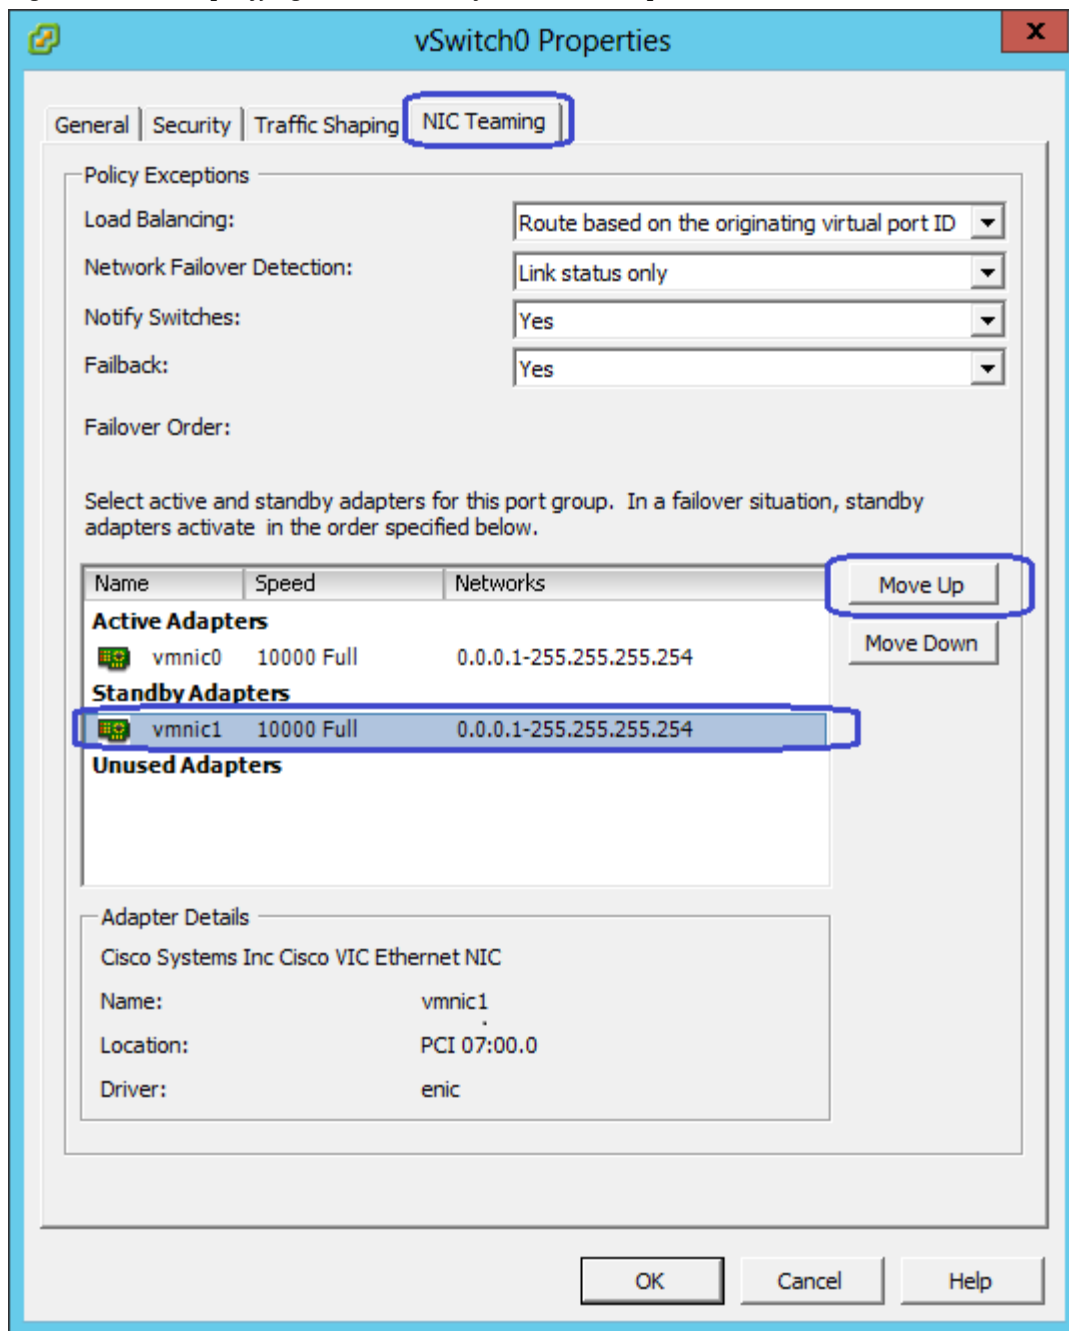
3. Change the MTU to **9000** in the **General** tab.

Figure 104 Setting JumboMTU



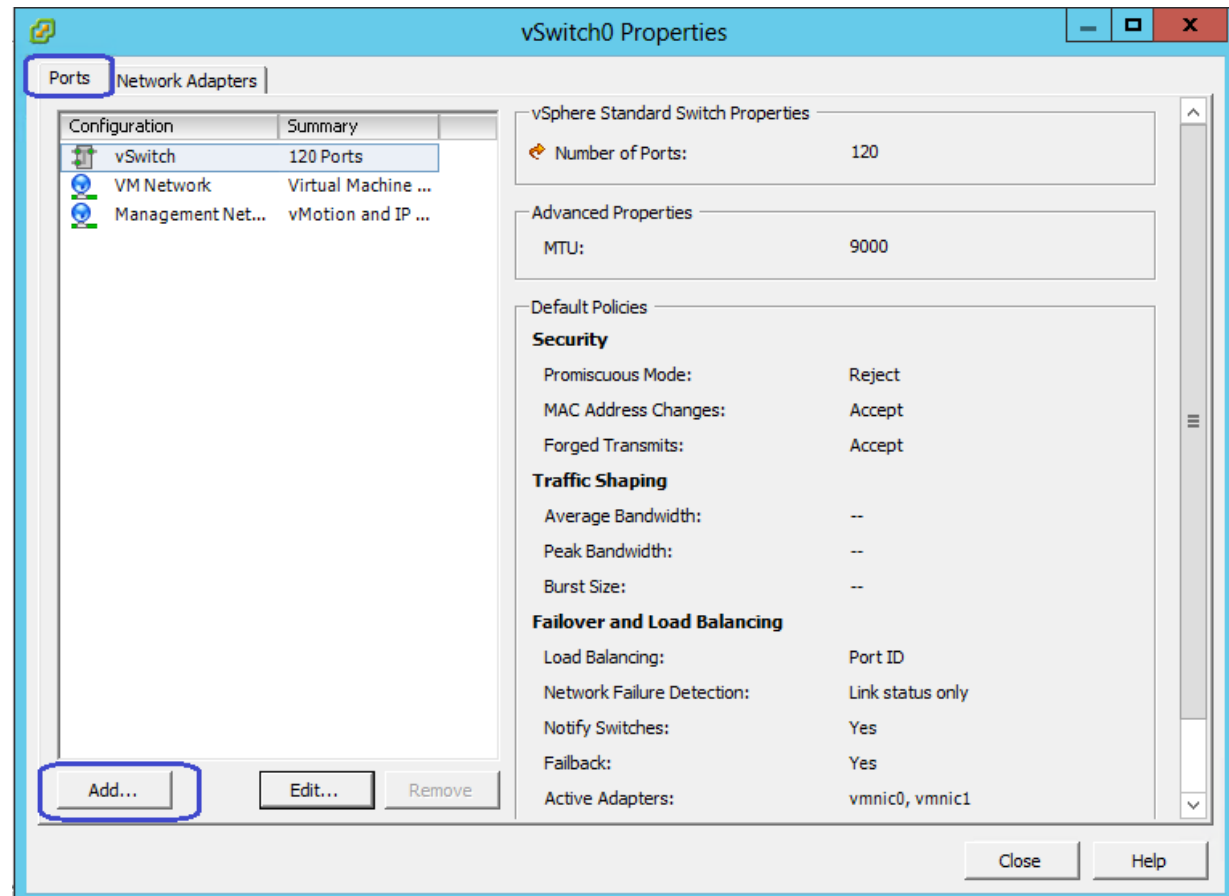
4. Click the **NIC Teaming** tab. Move up the Standby Adapter to the Active Adapters list, and then click **OK**.

Figure 105 Specifying the vmnic Order for the Port Group



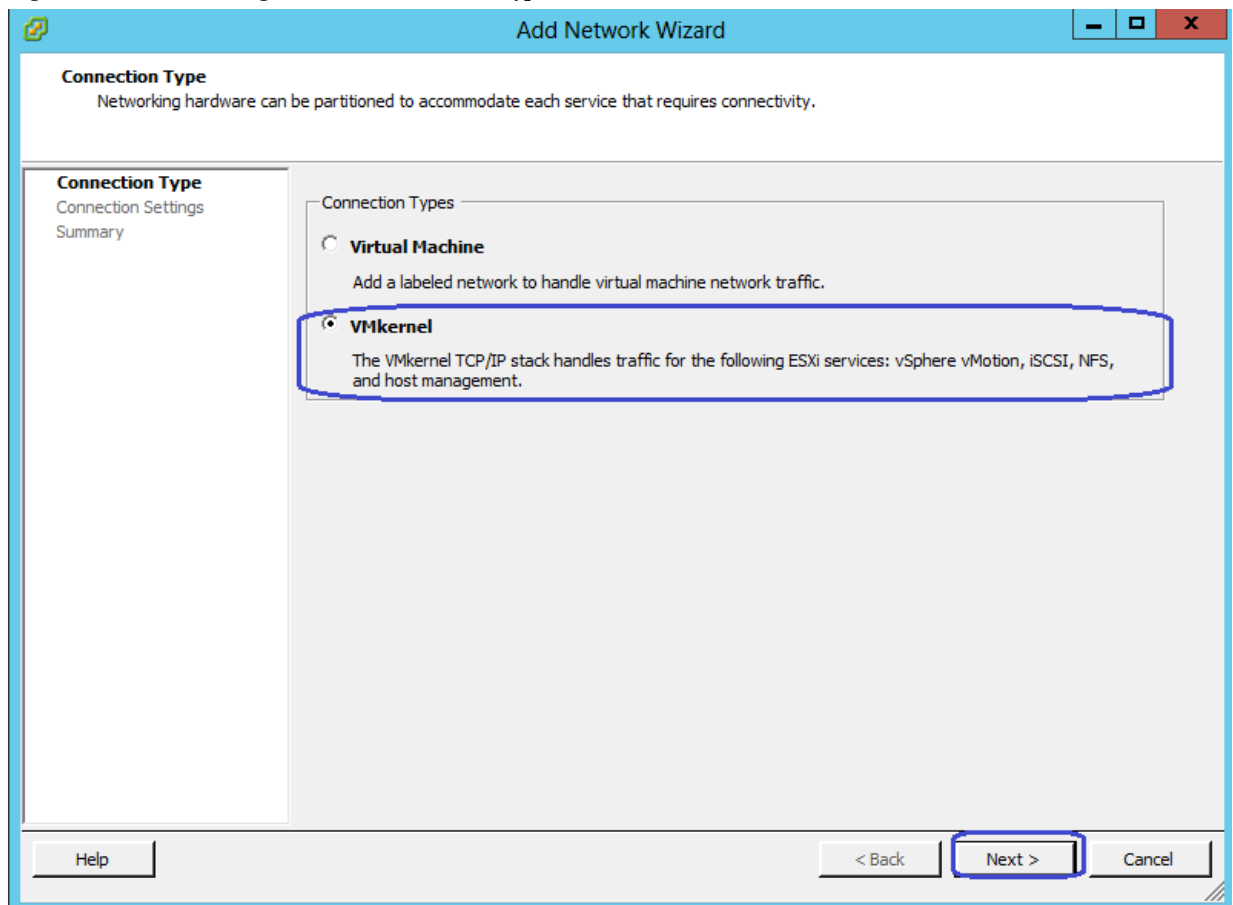
- That would bring you back to vSwitch0 configuration window. Click the **Ports** tab, and then click **Add**.

**Figure 106** Adding Ports in the vSwitch Property Window



- Click the **VMKernel** radio button and then click **Next** in the Add Network wizard window.

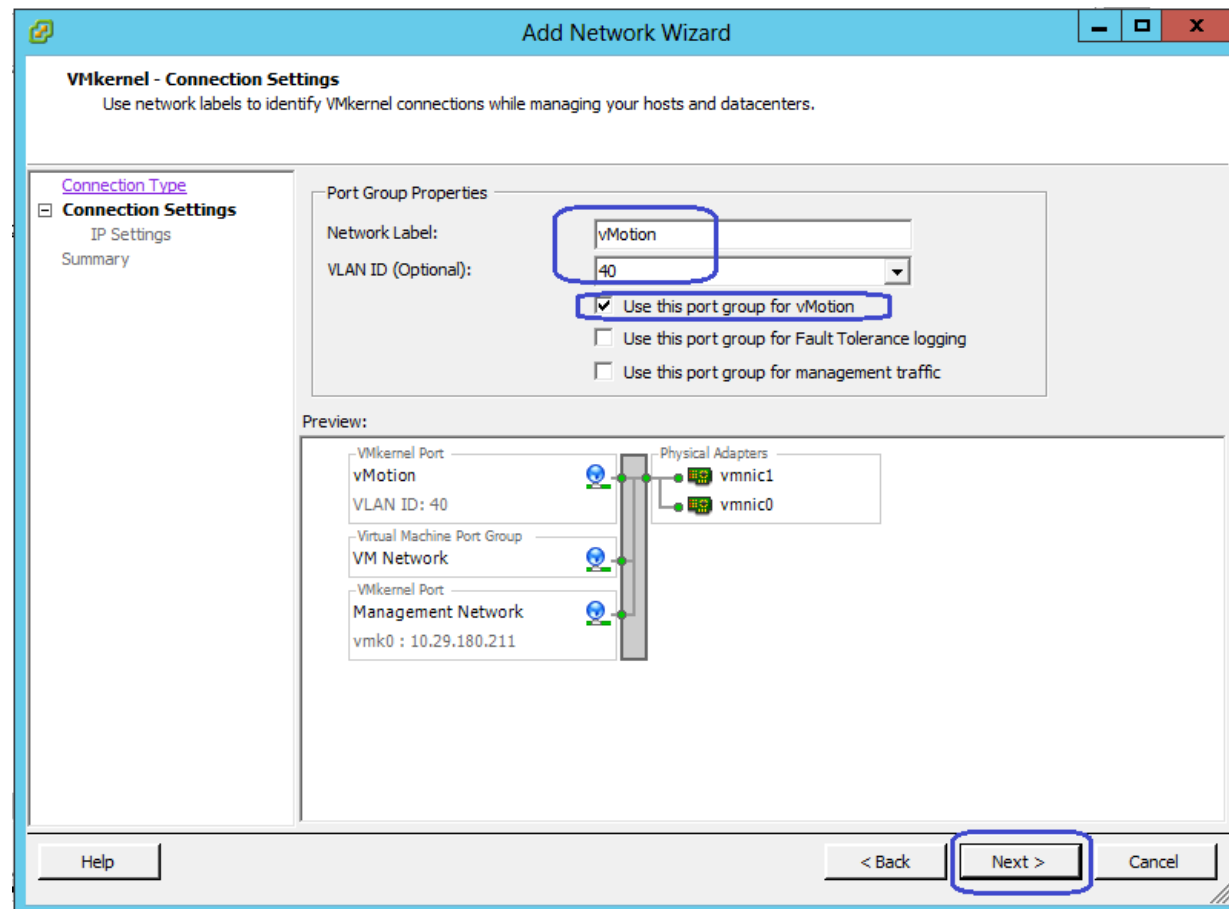
**Figure 107** Adding Network - Connection Type



7. Specify a name as vMotion in the Network Label field. Choose the VLAN ID, as standard vSwitch0 carries both management and vMotion VLANs. Management traffic leaves vSwitch0 untagged, using the native VLAN of the vNIC, but the vMotion traffic must be tagged with appropriate VLAN ID. Check the **Use this port group for vMotion** check box.

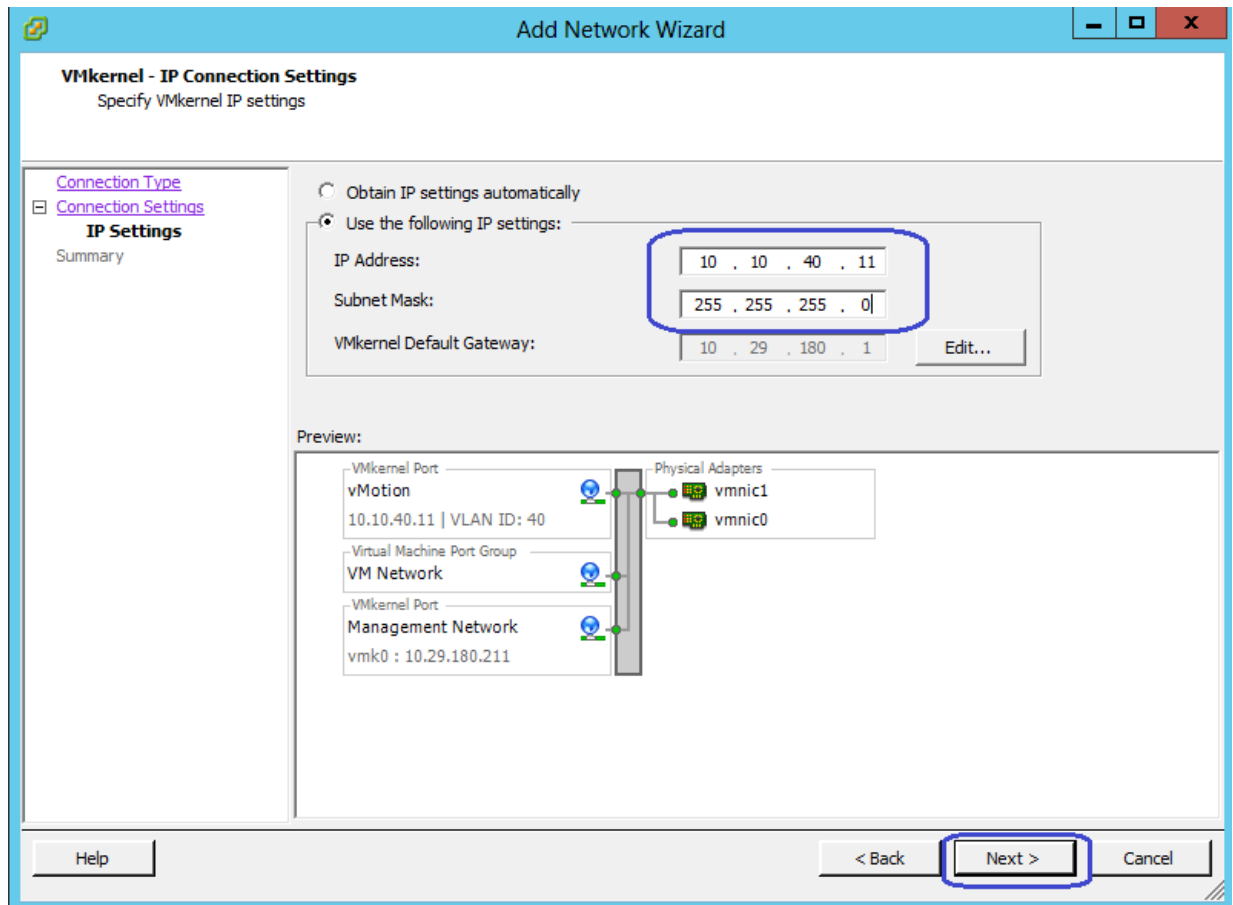


Figure 108 Adding Network - Connection Settings



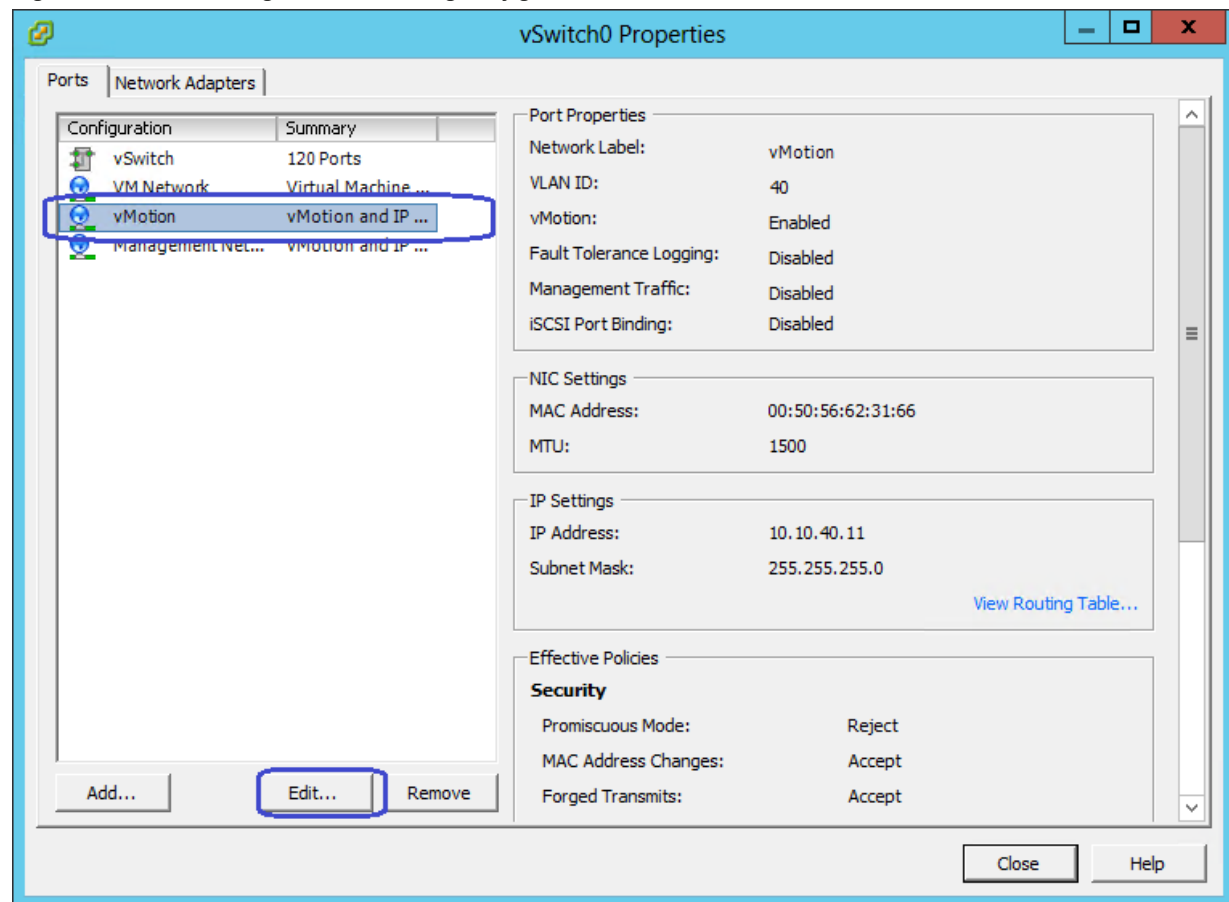
8. Configure IP address and subnet mask for the vmkernel interface in the next step.

Figure 109 Adding Network - IP Settings



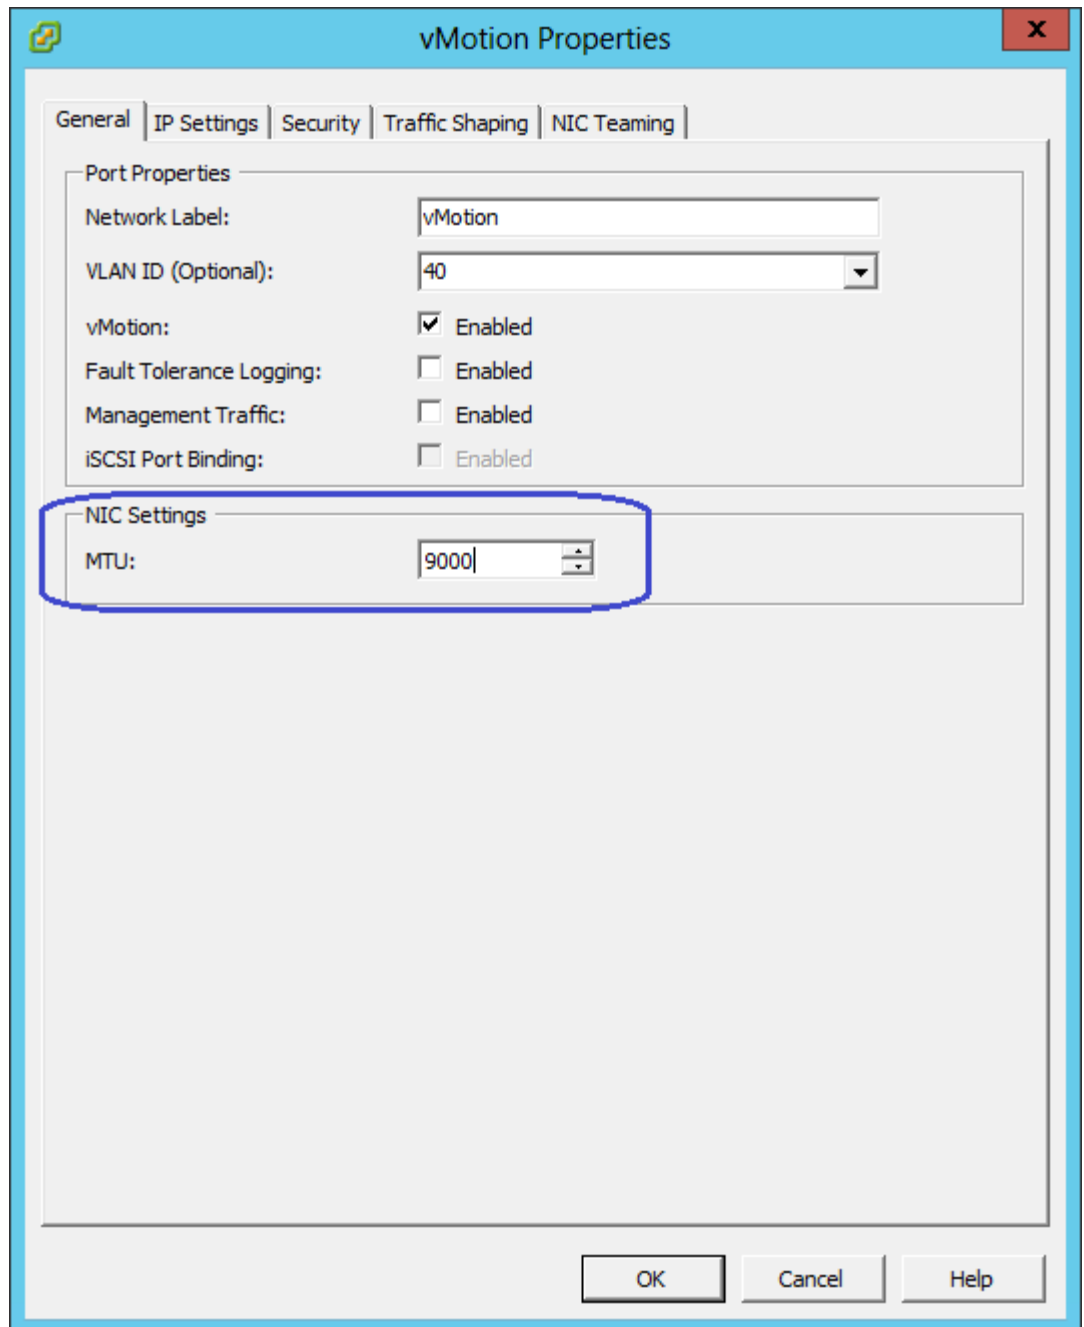
9. Click **Next** and in the next window click **Finish** to deploy the vMotion vmkernel. Back in the vSwitch0 properties window, select the newly created vMotion port group and click **Edit**.

**Figure 110** Adding Network - Editing Configuration



- Set the MTU to **9000** and click **OK**. Click **Close** in the parent window.

Figure 111 Editing Configuration - vMotion Properties Window



11. Repeat steps 1 to 10 for the remaining ESXi hosts in the cluster. Once all the ESXi hosts are configured, you must be able to ping from one host to another on the vMotion vmkernel port with jumbo MTU. Validate this by pinging the respective IPs with using the option do not fragment.

```

login as: root
Using keyboard-interactive authentication.
Password:
The time and date of this login have been sent to the system logs.

VMware offers supported, powerful system administration tools. Please
see www.vmware.com/go/sysadmintools for details.

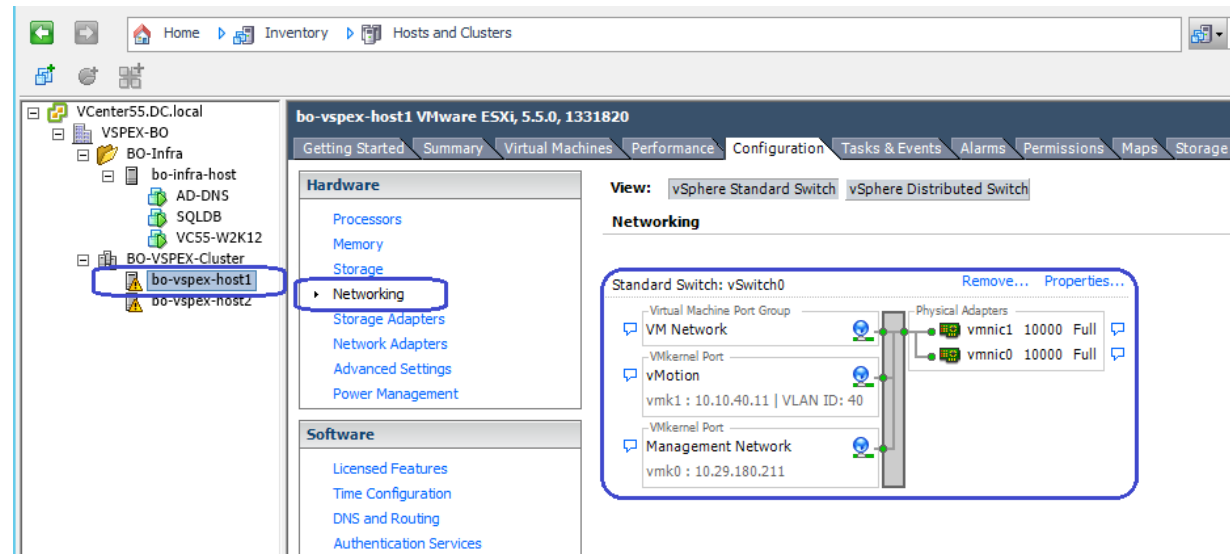
The ESXi Shell can be disabled by an administrative user. See the
vSphere Security documentation for more information.
~ # vmkping -d -s 8972 10.10.40.12
PING 10.10.40.12 (10.10.40.12): 8972 data bytes
8980 bytes from 10.10.40.12: icmp_seq=0 ttl=64 time=0.499 ms
8980 bytes from 10.10.40.12: icmp_seq=1 ttl=64 time=0.310 ms
8980 bytes from 10.10.40.12: icmp_seq=2 ttl=64 time=0.349 ms

--- 10.10.40.12 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.310/0.386/0.499 ms
~ #

```

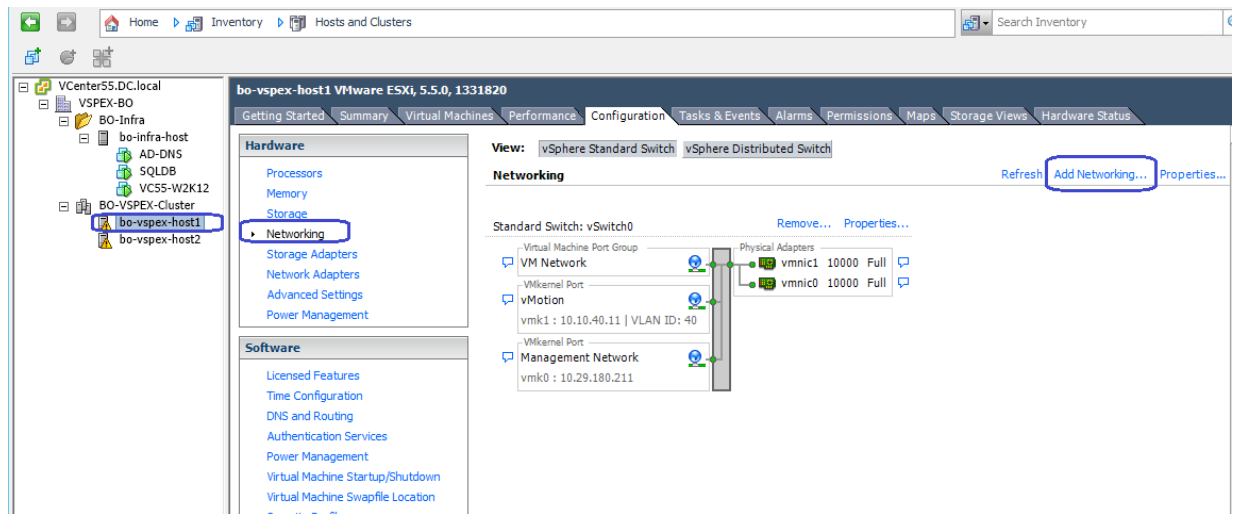
12. Back in the vCenter view, select ESXi host1 from the **Hosts and Clusters**, click the **Configuration** tab > **Networking**. Make sure vSwitch0 has both management traffic and vMotion traffic for vmnic0 and vmnic1.

**Figure 112** Verify the vmnics for vMotion Traffic



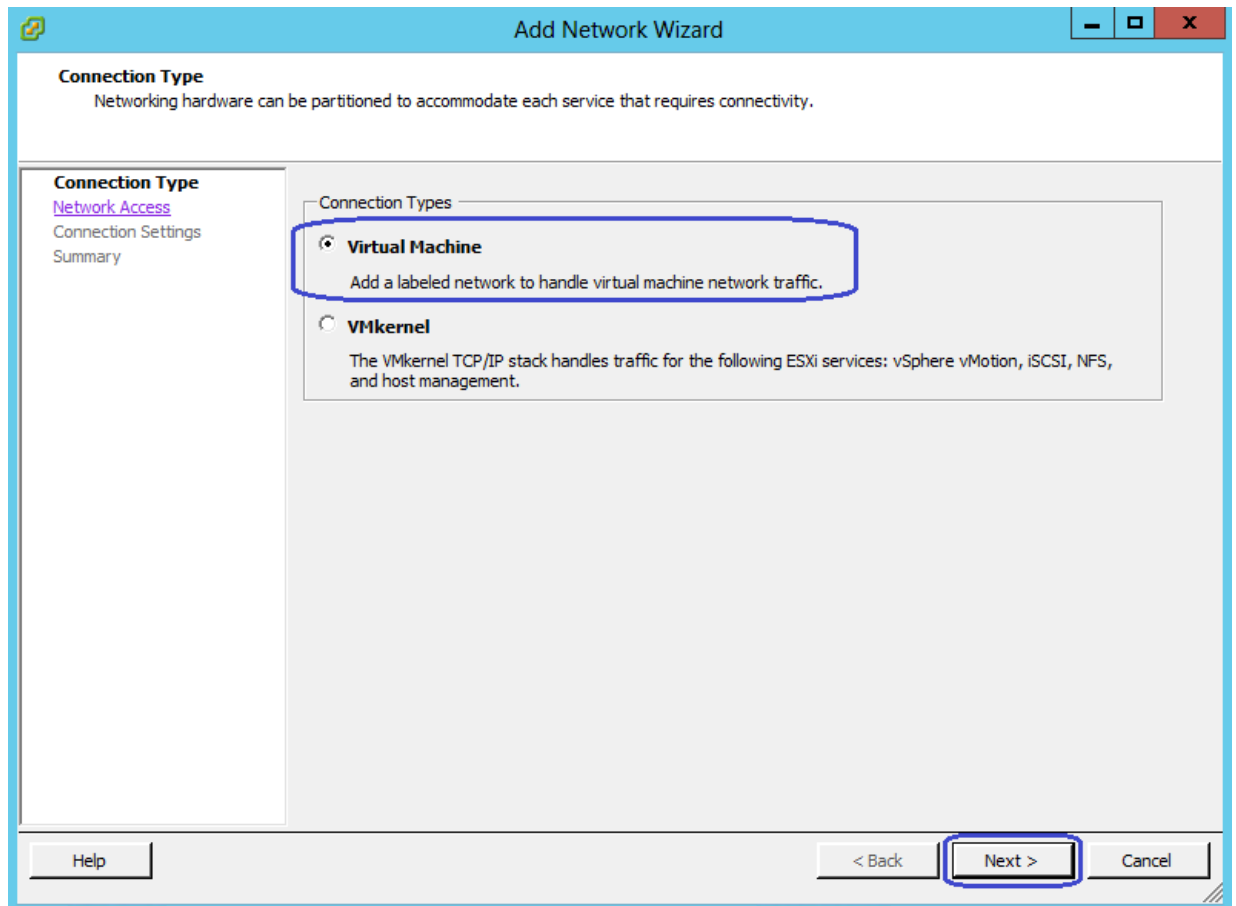
13. To Create vSwitch1 for VM-Data traffic, click **Networking** and then click **Add Networking**.

**Figure 113** Add Networking to Create vSwitch for VM-Data



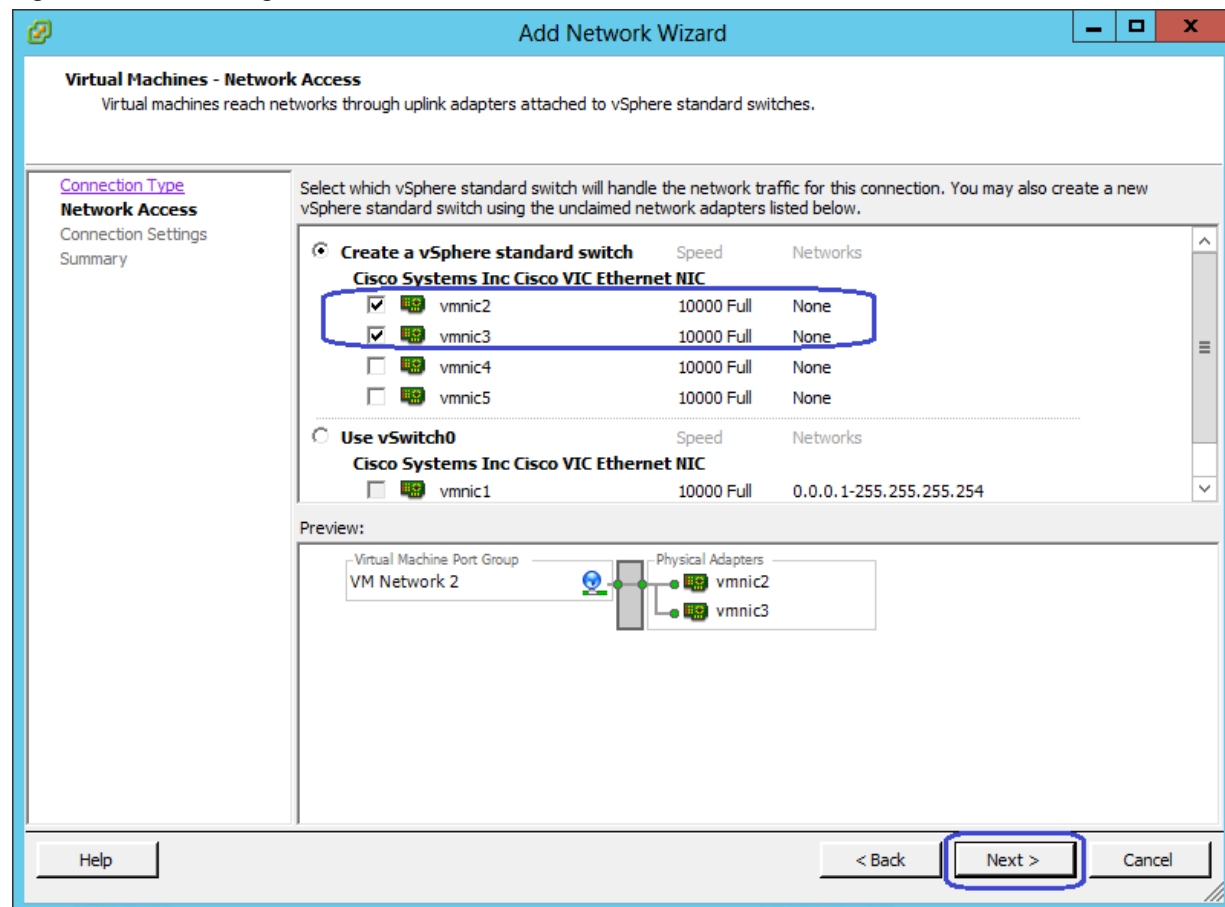
14. Click the radio button **Virtual Machine** in the Add Networking Wizard, click **Next**.

**Figure 114** Adding Network - Connection Type



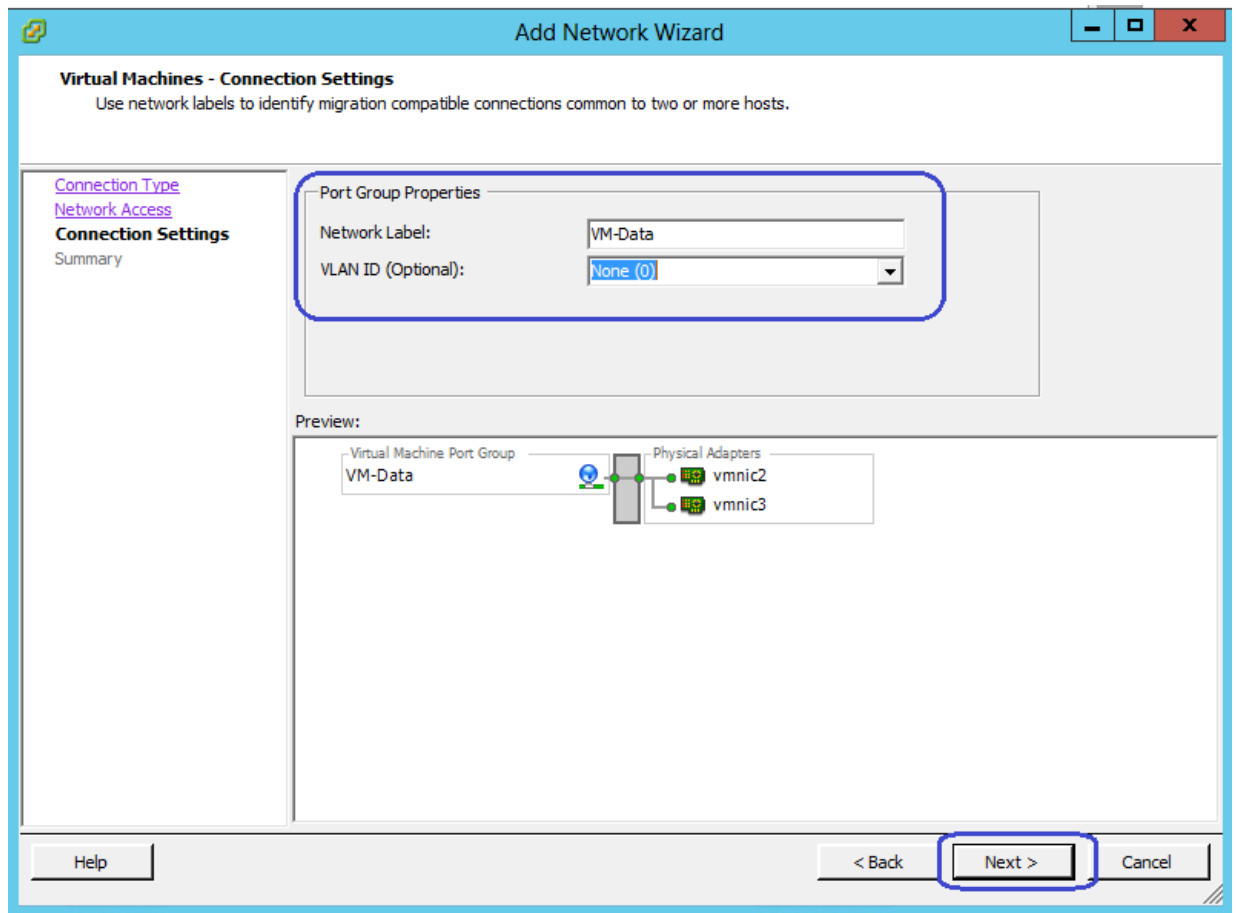
15. Select the two vmnics corresponding to the VM-Data VNICS and click **Next**.

Figure 115 Adding Network - Network Access



- Specify the name for Network Label as VM-Data, and keep VLAN ID as 0 to signify the absence of VLAN tag. Click **Next**.

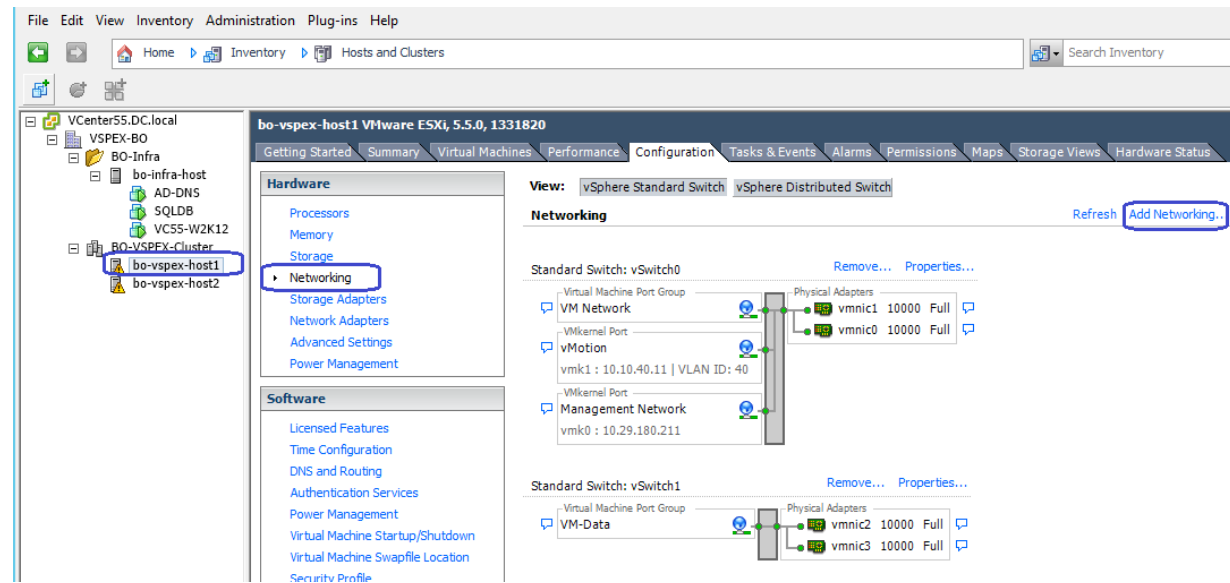
**Figure 116** Adding Network - Connection Settings



17. In the next window click **Finish** to deploy the vSwitch1 creation for VM-Data traffic. Repeat steps 13 to 16 for the remaining ESXi host in the cluster.
18. To create vSwitch2 and VMkernel for Storage traffic, click **Networking > Add Networking**.

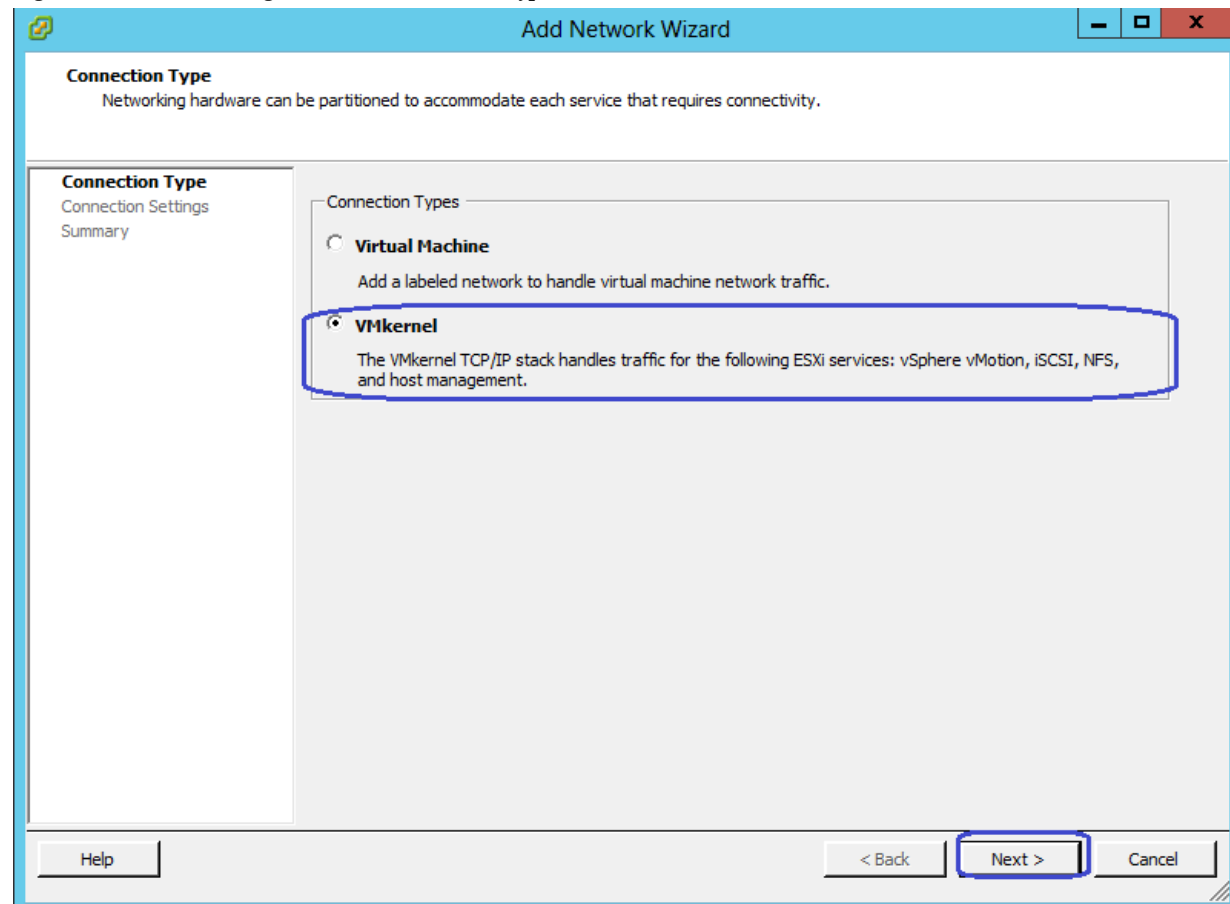


**Figure 117** Add Networking to Create vSwitch and vmKernel for Storage Traffic



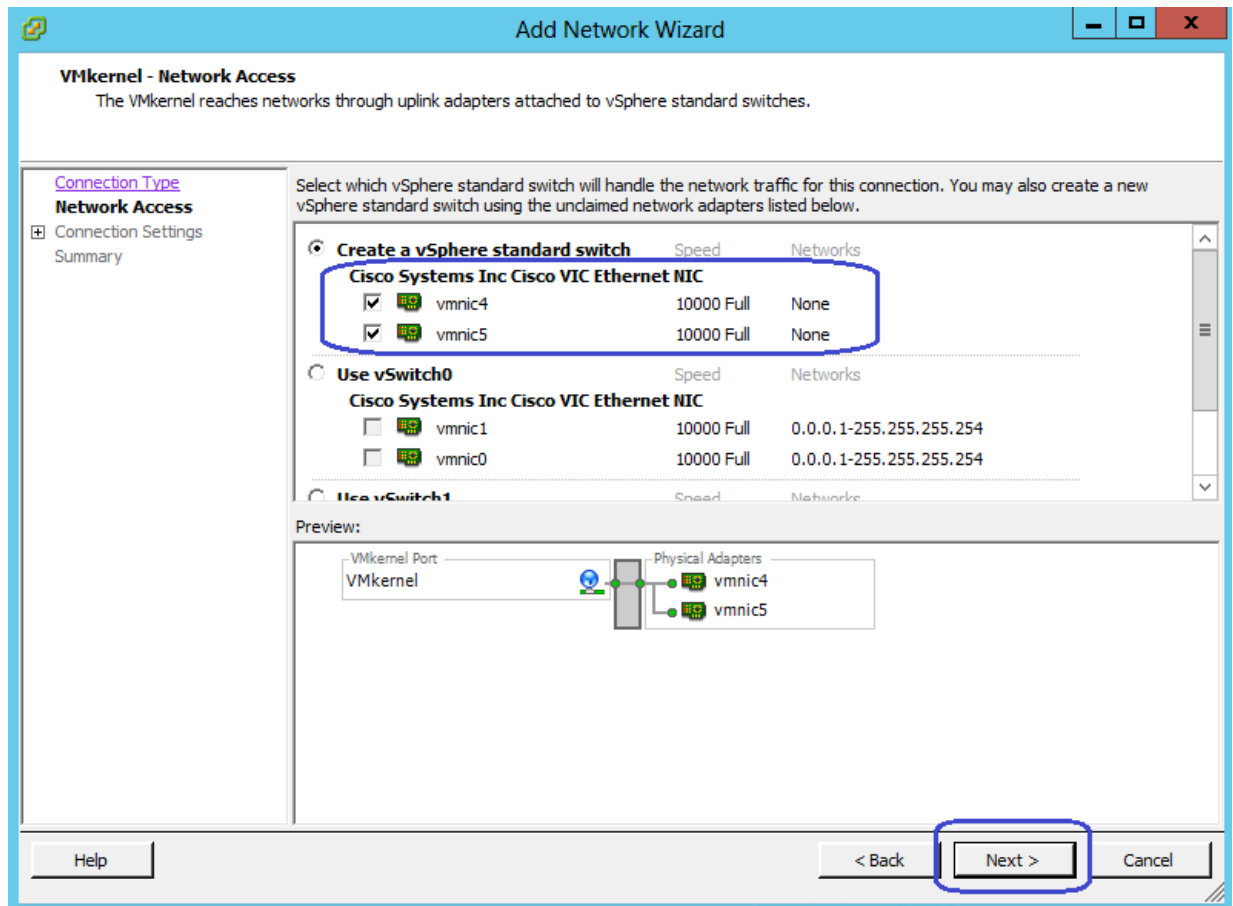
19. Click the **VMkernel** radio button and then click **Next** in the Add Network Wizard window.

**Figure 118** Adding Network - Connection Type



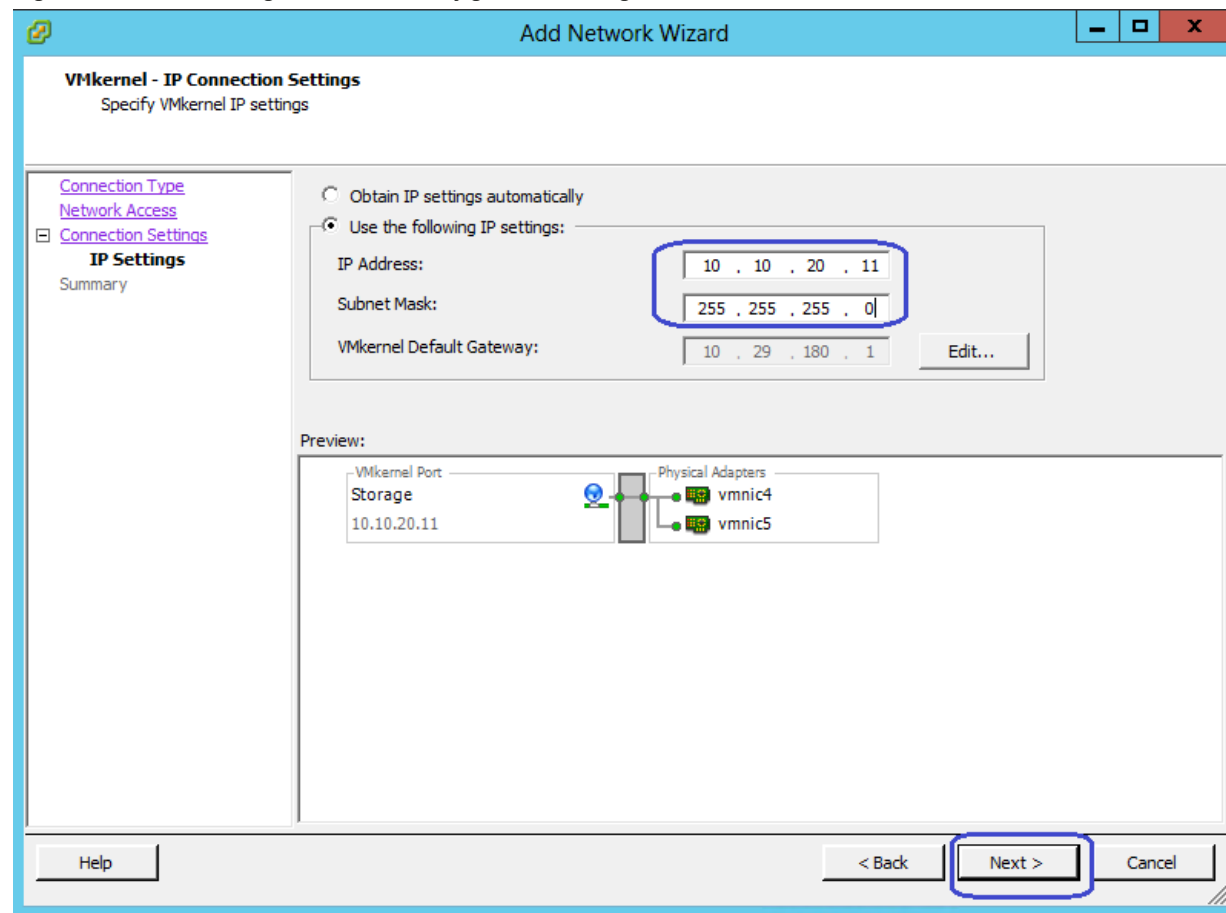
20. Select the two vmnics (vmnic4 and vmnic5) corresponding to the NFS Storage VNICs and click **Next**.

Figure 119 Adding Network - Network Access



21. Configure IP address and subnet mask for the vmkernel interface in the next step.

**Figure 120** Adding Network - IP Configuration Settings



22. Click **Next** and in the next window click **Finish** to deploy the vSwitch2 and VMkernel for Storage traffic. Repeat the steps 18 to 21 to create Storage VMkernel for the remaining ESXi host.

This concludes the Virtual Networking configuration on the vCenter.

## Configure storage for VM data stores, install and instantiate VSPEX VMs from vCenter

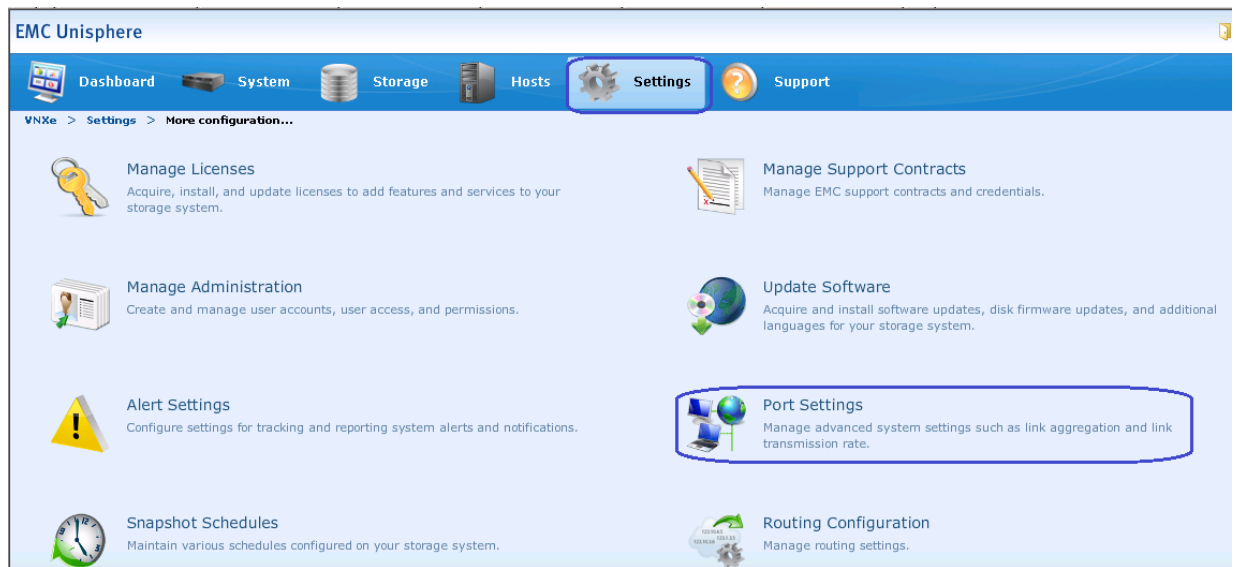
This section describes the steps to configure VNXe Storage for VM Data stores to install and instantiate VSPEX VMs from vCenter. This includes five steps:

1. Configure Link Aggregation
2. Create NAS Server
3. Configure Storage Pool for NFS-Data
4. Create VMware Datastores
5. Mount NFS share on ESXi hosts

### Configure Link Aggregation

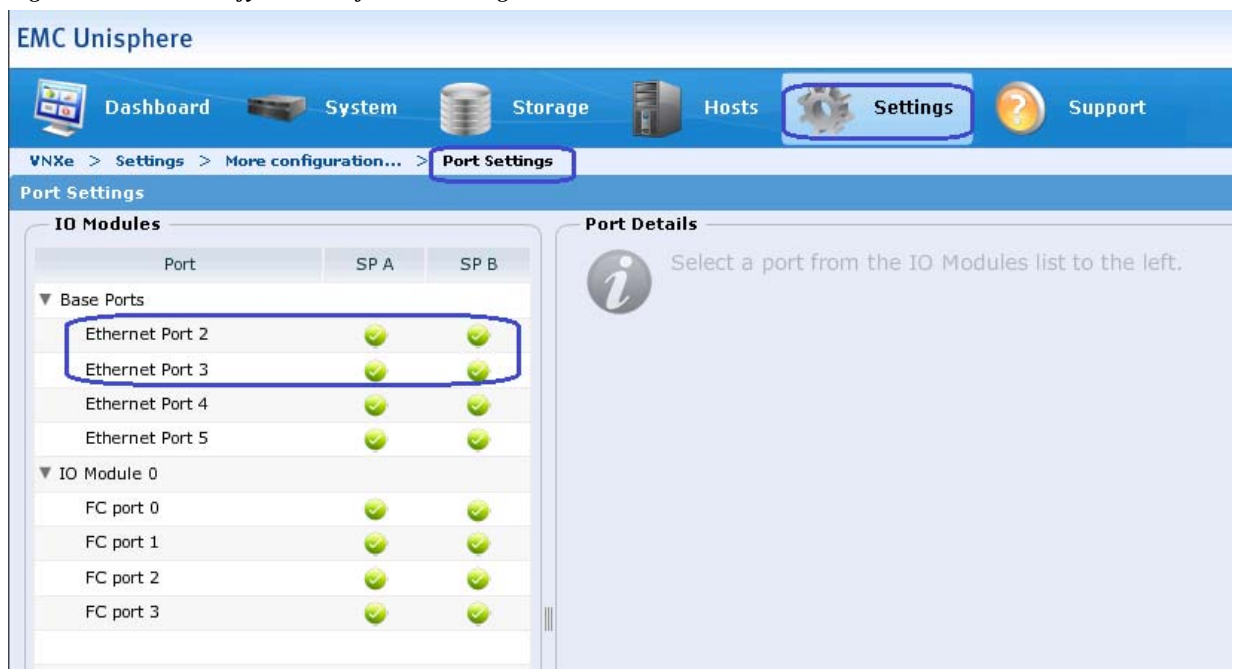
1. Connect to the EMC VNXe Unisphere GUI, click the **Settings** tab > **Port Settings**.

**Figure 121** Port Configuration in EMC Unisphere



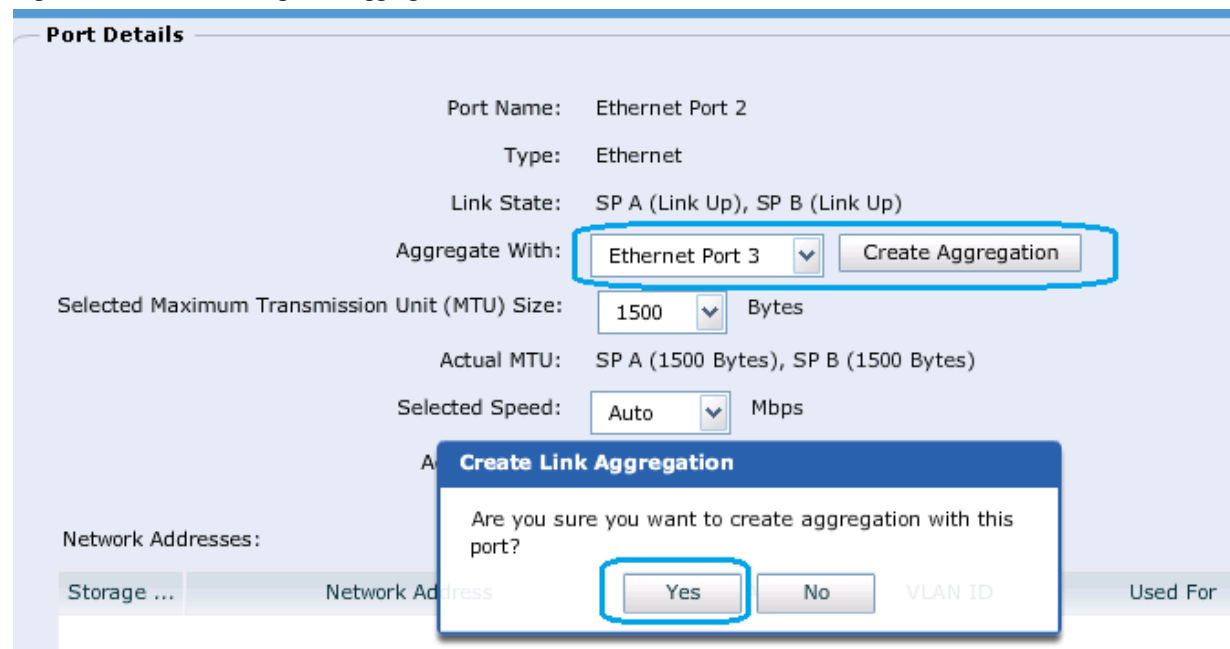
- We connected 1GbE ports Ethernet Port 2 and Port 3 for NFS Storage Access from UCS Blades.

**Figure 122** Verify the Ports for NFS Storage Access



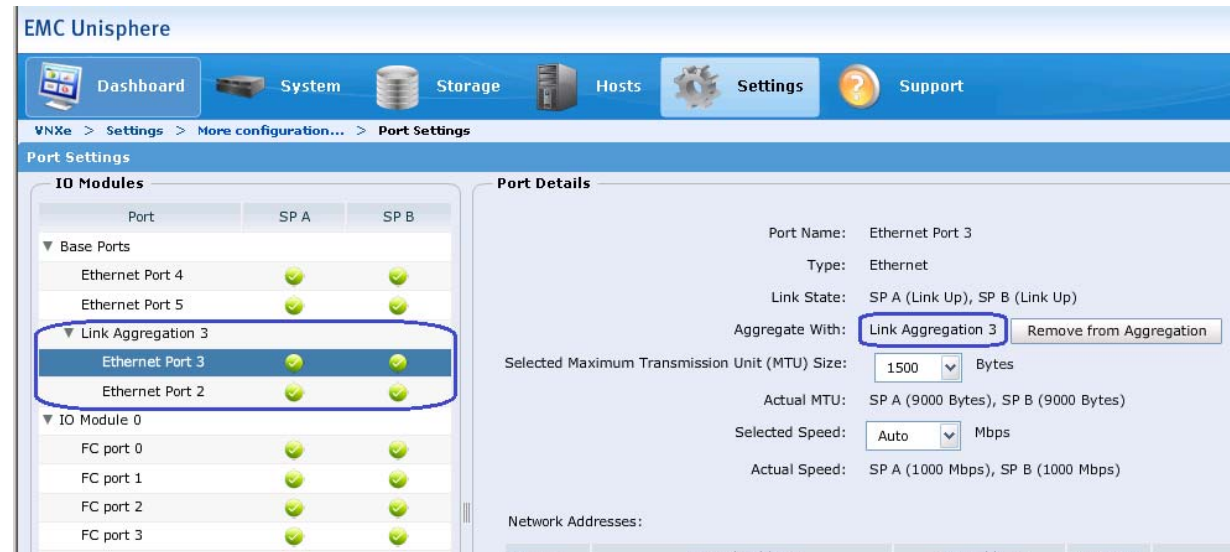
- To Create Aggregation, click **Ethernet Port 2** and choose Ethernet Port 3 from the drop-down list and click **Create Aggregation**.
- A pop-up window to confirm the Link Aggregation Creation appears. Click **Yes** to confirm.

**Figure 123**      *Creating Link Aggregation*



- After confirmation, you will see the created Link Aggregation between Ethernet Port 2 and Ethernet Port 3.

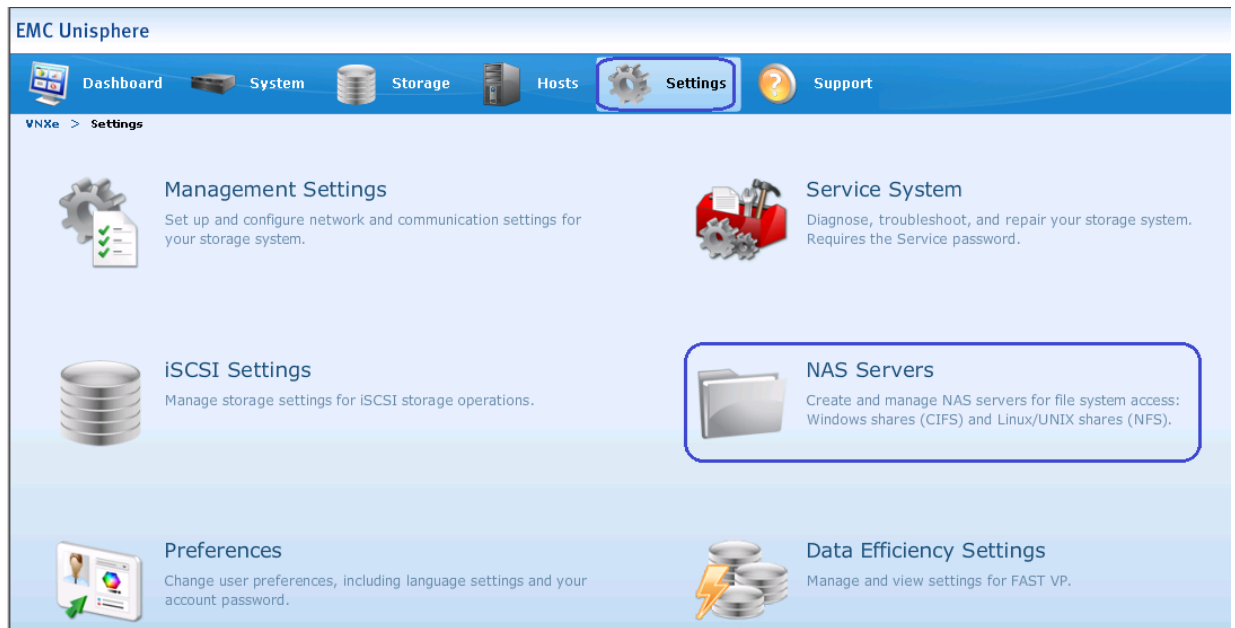
**Figure 124**      *Verify the Created Link Aggregation*



## Create NAS Server

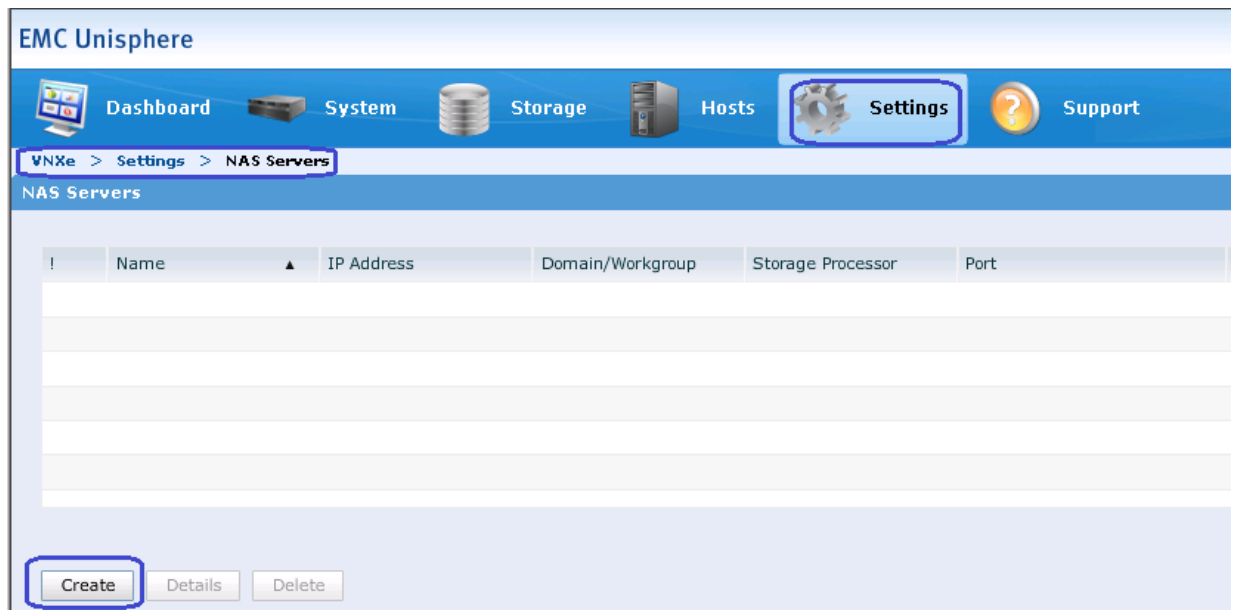
- Connect to the EMC VNXe Unisphere GUI, click the **Settings** tab > **NAS Servers**.

**Figure 125** *NAS Servers in EMC Unisphere*



2. To create NAS Server for NFS Storage, click **Create**.

**Figure 126** *Creating NAS Server for NFS Storage*



3. To configure NAS Server Address, specify the NAS Server name, and then choose the storage pool that we created for NFS-Data from the drop-down list. Specify the IP address and the Subnet Mask for NAS server and choose Link Aggregation 3 from the Ethernet port drop-down list. Click **Next** to proceed.

Figure 127 Creating NAS Server - Configure NAS Server Address

**NAS Server Wizard**

**Configure NAS Server Address**

Step 1 of 5

Server Name: \* NASServer01

Storage Pool: StoragePool-NFS-Data

IP Address: \* 10.10.20.10

Subnet Mask/Prefix Length: \* 255.255.255.0

Gateway:

Hide advanced

Storage Processor: SP A

Ethernet Port: Link Aggregation 3 (Link Up)

VLAN ID: 0 <click to edit>

< Back Next > Finish Cancel Help

- To configure Share Types, check the Linux/Unix shares (NFS) check box and click **Next**.

Figure 128 Creating NAS Server - Configure Share Types

**NAS Server Wizard**

**Configure Share Types**

Step 2 of 5

Choose the type of shares the NAS server supports:

☒ Linux/Unix shares (NFS)

☐ Windows shares (CIFS)

< Back Next > Finish Cancel Help

- Skip the Configure NAS Server DNS and NIS settings and click **Next**. Then, verify the NAS Server settings summary and click **Finish**.

Figure 129 Creating NAS Server - Summary

**NAS Server Wizard**

**Create NAS Server Summary**

Step 4 of 5

Verify the following NAS server settings:

Server Name: NASServer01

Storage Processor: SP A

IP Address: 10.10.20.10

Subnet Mask: 255.255.255.0

Gateway: Not Specified

Advanced Attributes: Custom

Storage Pool: StoragePool-NFS-Data

Support Linux/Unix Shares (NFS): Configured

Support Windows Shares (CIFS): Not configured

NIS Domain Name:

< Back Next > **Finish** Cancel Help

- After the successful creation of NAS Server, you will see the newly created NAS Server.

Figure 130 Created NAS Server with Created Link Aggregation

EMC Unisphere

Dashboard System Storage Hosts Settings Support

VNXe > Settings > **NAS Servers**

NAS Servers

	Name	IP Address	Domain/Workgroup	Storage Processor	Port	
✓	NASServer01	10.10.20.10	Not configured	SP A	Link Aggregation 3	Normal

Create Details Delete

## Configure Storage Pool for NFS-Data

- Connect to the EMC VNXe Unisphere GUI, Click the **Storage** tab > **Storage Configuration** > **Storage pools**.

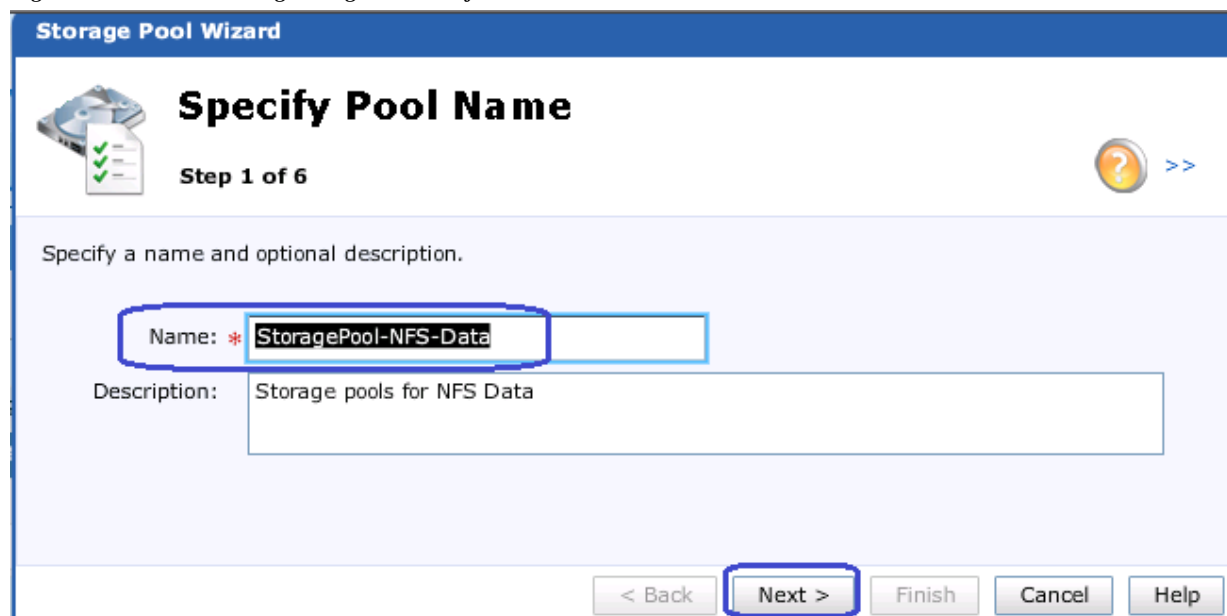


**Figure 131**      *Creating Storage Pools*



2. Specify the pool name in the Name field and Click **Next**.

**Figure 132**      *Creating Storage Pool - Define Pool*



3. Check the Performance Tier to use 600GB SAS disks check box and Click **Next**.

Figure 133 Creating Storage Pool - Selecting Storage

**Storage Pool Wizard**

**Select Storage**

Step 2 of 6

Select the storage tiers you want to use for the new pool.

	Storage Tier	Disk Type	Unused Disks	Unused Raw Capacity
<input type="checkbox"/>	Extreme Performance Tier	Flash	5	917.2 GB
<input checked="" type="checkbox"/>	Performance Tier	SAS	40	20.9 TB
<input type="checkbox"/>	Capacity Tier	NL SAS	0	0 GB (None Available)

Uses SAS disks to provide high performance. These disks do not provide the same read/write performance as Extreme Performance (Flash) disks, but offer much lower cost per GB of storage.

Performance Tier

RAID Type: RAID 5 (4+1) (Usable capacity: 14.3 TB) [Change](#)

< Back **Next >** Finish Cancel Help

- For SAS Disks, choose **Use 35 of 40 disks (14.3TB)** from the drop-down list to use 35 SAS disks for creating storage pool for NFS-Datastore and Click on **Next**.

Figure 134 Creating Storage Pool - Storage Size

**Storage Pool Wizard**

**Select Amount of Storage**

**Step 3 of 5**

Select the amount of storage for each selected tier. The number of disks you can choose is based on the RAID configuration selected. The maximum number of disks you can configure will ensure that enough disks are kept unused to satisfy the hot spare policy.  
[More information](#)

Performance Tier

600 GB (10K RPM) SAS Disks: Use 35 of 40 disks (14.3 TB)

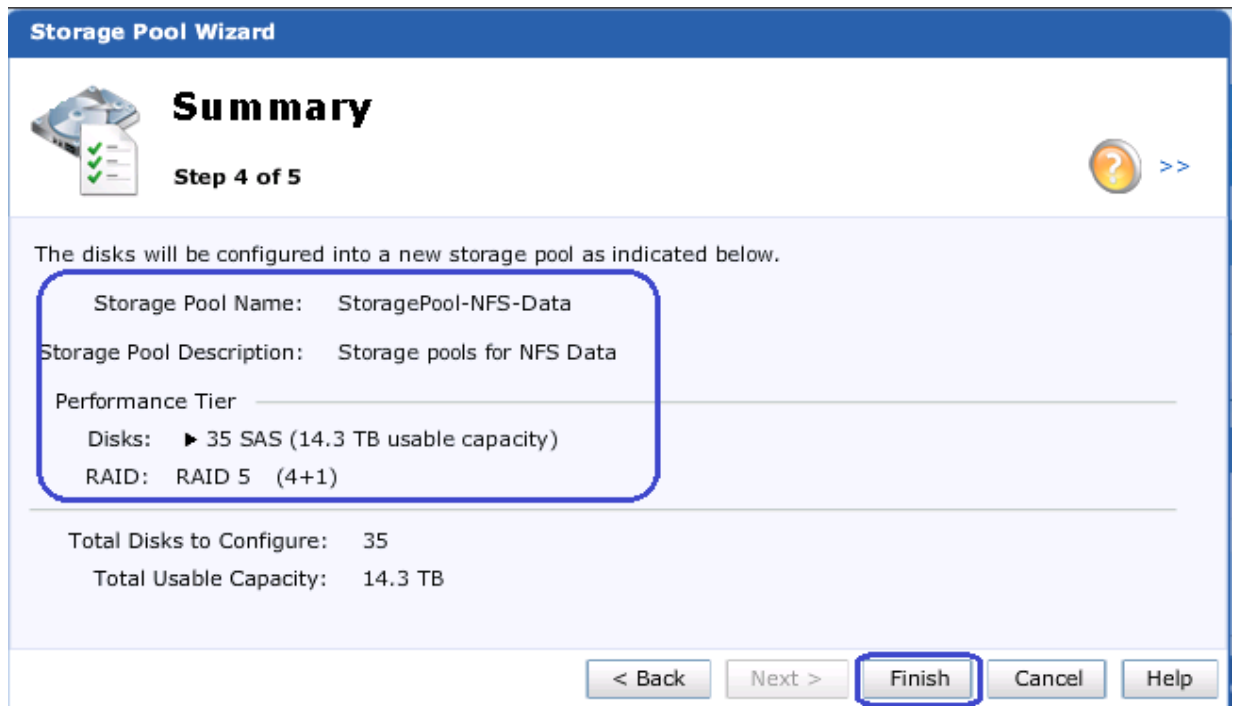
This option uses the system disks. The portion of the storage pool that uses these disks will have reduced capacity and storage resources in the pool may experience reduced performance.

Total Disks to Configure: 35  
 Total Usable Capacity: 14.3 TB

< Back **Next >** Finish Cancel Help

5. Verify the Storage pool creation and click **Finish**.

Figure 135 Creating Storage Pool - Summary



**Storage Pool Wizard**

**Summary**

Step 4 of 5

The disks will be configured into a new storage pool as indicated below.

Storage Pool Name: StoragePool-NFS-Data

Storage Pool Description: Storage pools for NFS Data

Performance Tier

Disks: ► 35 SAS (14.3 TB usable capacity)

RAID: RAID 5 (4+1)

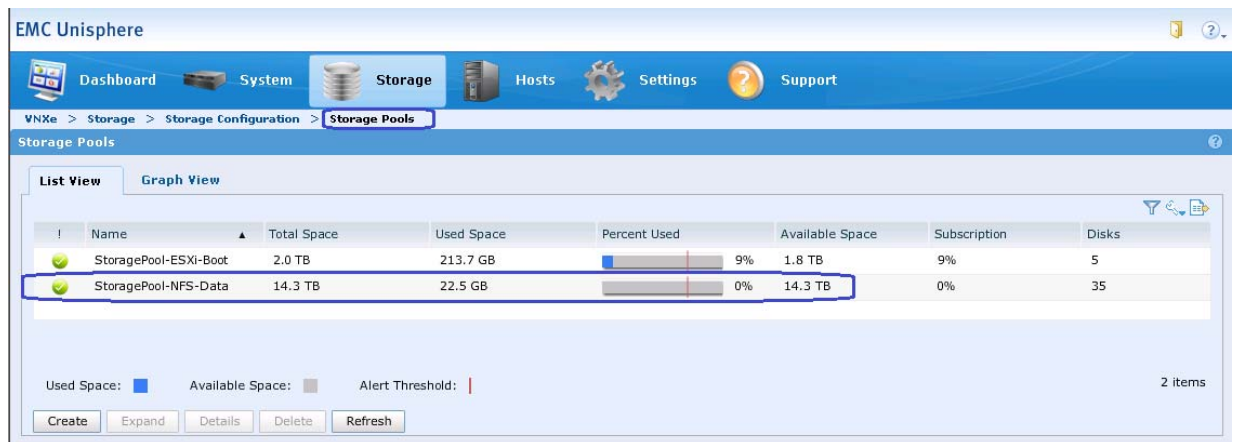
Total Disks to Configure: 35

Total Usable Capacity: 14.3 TB

< Back Next > **Finish** Cancel Help

- Now you will see the Storage pool created for NFS Data.

Figure 136 Verify the Created Storage Pool for NFS Data



EMC Unisphere

Dashboard System Storage Hosts Settings Support

VNXe > Storage > Storage Configuration > **Storage Pools**

Storage Pools

List View Graph View

	Name	Total Space	Used Space	Percent Used	Available Space	Subscription	Disks
✓	StoragePool-ESXi-Boot	2.0 TB	213.7 GB	9%	1.8 TB	9%	5
✓	StoragePool-NFS-Data	14.3 TB	22.5 GB	0%	14.3 TB	0%	35

Used Space: ■ Available Space: ■ Alert Threshold: |

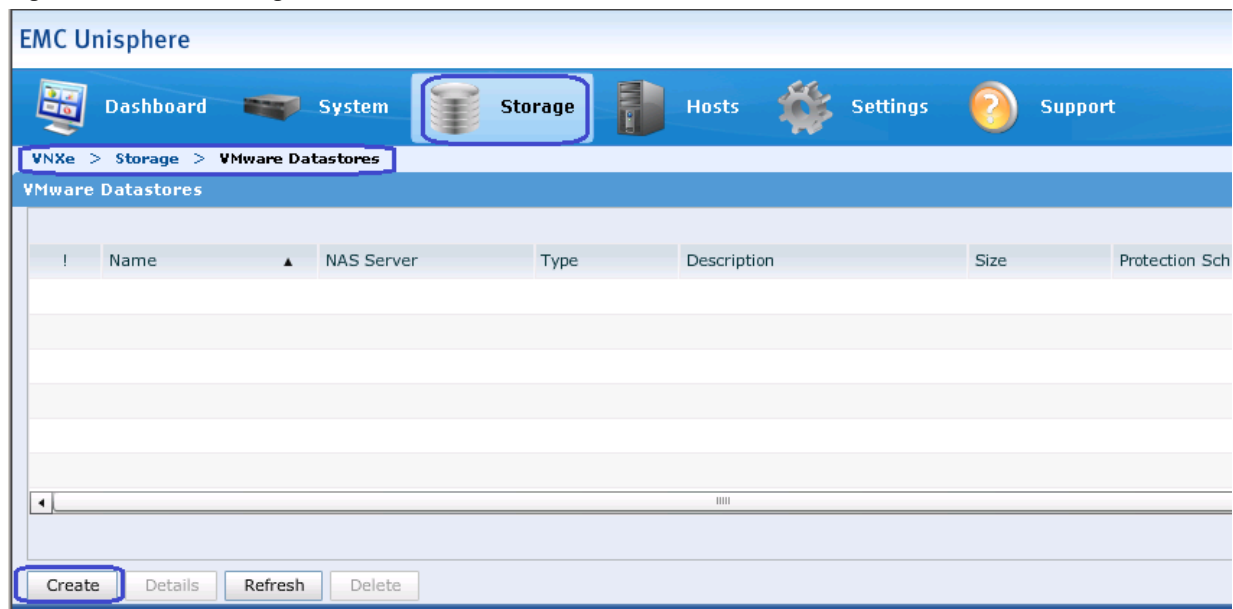
Create Expand Details Delete Refresh

2 items

## Create VMware Data Stores

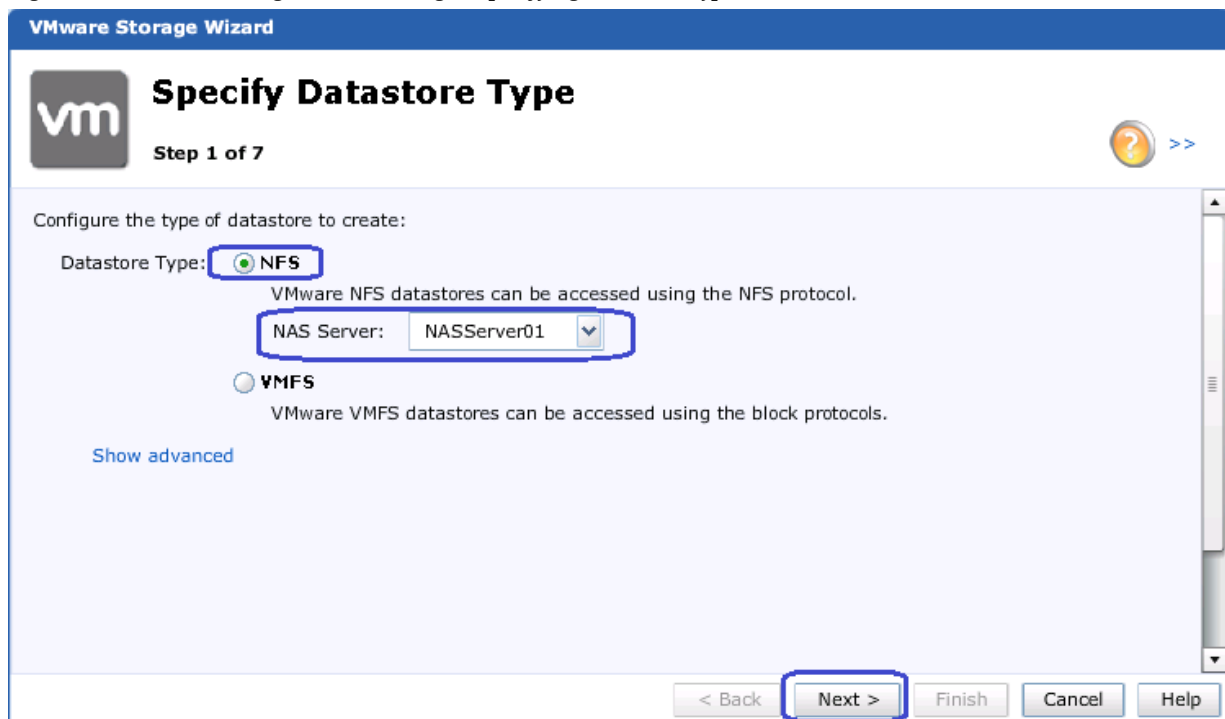
- Connect to the EMC VNXe Unisphere GUI, then click the **Storage > VMware Datastores > Create** to create VMware datastore.

Figure 137 Creating VMware Datastores



2. Click the **NFS** radio button for the Datastore Type and choose the NAS server as NASServer01 from the drop-down list.

Figure 138 Creating VMware Storage - Specifying Datastore Type



3. Specify the name for the VMware datastore and Click **Next**.

Figure 139 Creating VMware Storage - Specify Name

**VMware Storage Wizard**

**Specify Name**

Step 2 of 7

Enter a name for the VMware datastore.

Name: \* NFS-DataStore1

Description: VMFS Datastore for VSPEX Server

< Back **Next >** Finish Cancel Help

- From the Storage Pool drop down list choose the created **StoragePool-NFS-Data**. And specify the datastore size as **6TB** and check the check box **Thin** and click **Next**.

Figure 140 Creating VMware Storage - Configure Storage

**VMware Storage Wizard**

**Configure Storage**

Step 3 of 7

Configure the storage for this datastore:

Storage Pool: StoragePool-NFS-Data (SAS, 14.3 TB free)

Tiering Policy: Start High Then Auto-Tier (Recommended)

*i* The selected pool is not tiered. The tiering policy will have no effect on the storage resource.

Size: \* 6 TB ☒ Thin

< Back **Next >** Finish Cancel Help

- Skip Configure Snapshot Schedule by clicking the **Do not configure a snapshot schedule** radio button and click **Next**.

Figure 141 Creating VMware Storage - Configure Snapshot Schedule

**VMware Storage Wizard**

**Configure Snapshot Schedule**

Step 4 of 7

Configure a recurring snapshot schedule for automatic data protection:

☒ **Do not configure a snapshot schedule.**

A snapshot schedule can be selected at a later time.

☐ **Select a snapshot schedule:**

This schedule will create snapshots and synchronize data

Every day at 01:00, keep for 2 days

Note: Times are displayed in Local Time (UTC-07:00) in 24-hour format

< Back **Next >** Finish Cancel Help

- Configure the hosts access by choosing **Read/Write, allow Root** from the drop-down list for Default Access setting for both the VSPEX servers.

Figure 142 Creating VMware Storage - Configure Host Access

**VMware Storage Wizard**

**Configure Host Access**

Step 5 of 7

Configure which hosts will access this storage:

Default Access: **Read/Write, allow Root**

The Default Access setting applies to hosts configured in Unisphere with Access set to "Use Default Access" and hosts configured outside Unisphere that can reach the share.

Filter for:  Protocols: **File**

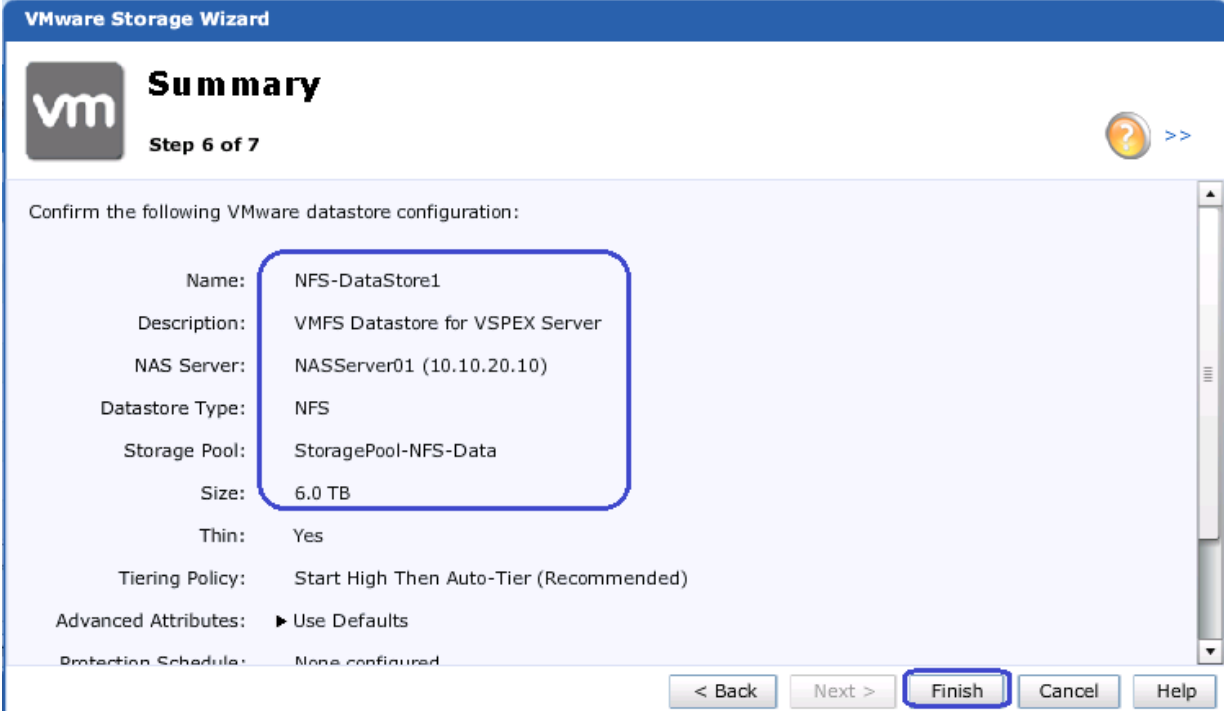
!	Name	Network Address	Protocol	Access
✓	BO-VSPEX-Serve...	10.29.180.211	FC, File	Use Default Access
✓	BO-VSPEX-Serve...	10.29.180.212	FC, File	Use Default Access

Filtered: 2 of 2

< Back **Next >** Finish Cancel Help

- Verify the VMware Datastore Configuration summary and click **Finish** to confirm.

Figure 143 Creating VMware Storage - Summary



**VMware Storage Wizard**

**Summary**

Step 6 of 7

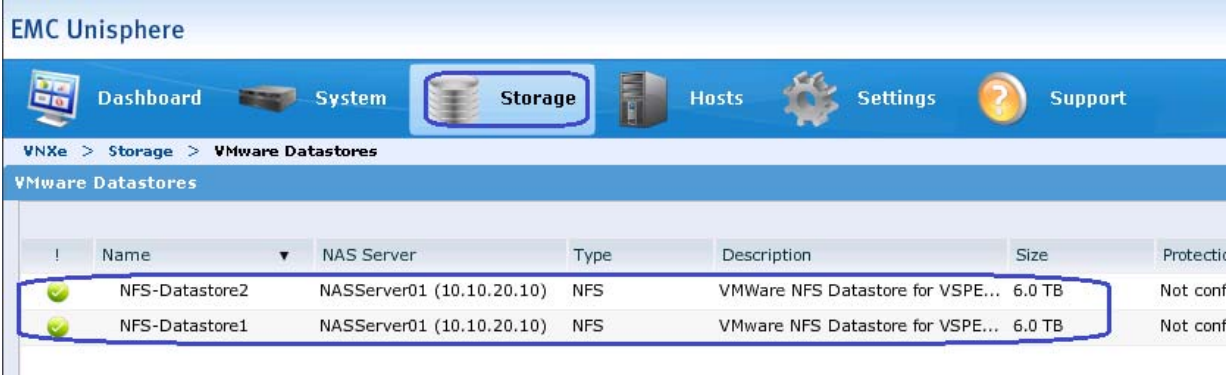
Confirm the following VMware datastore configuration:

Name:	NFS-DataStore1
Description:	VMFS Datastore for VSPEX Server
NAS Server:	NASServer01 (10.10.20.10)
Datastore Type:	NFS
Storage Pool:	StoragePool-NFS-Data
Size:	6.0 TB
Thin:	Yes
Tiering Policy:	Start High Then Auto-Tier (Recommended)
Advanced Attributes:	► Use Defaults
Protection Schedule:	None configured

< Back   Next >   **Finish**   Cancel   Help

8. Repeat the Steps 1 to 7 and create NFS-DataStore2 with 6.0TB size.
9. After successful creation, you will see both the VMware Datastores created with 6TB NFS volume.

Figure 144 Verify the Created VMware Datastores



EMC Unisphere

Dashboard   System   **Storage**   Hosts   Settings   Support

VNXe > Storage > VMware Datastores

VMware Datastores

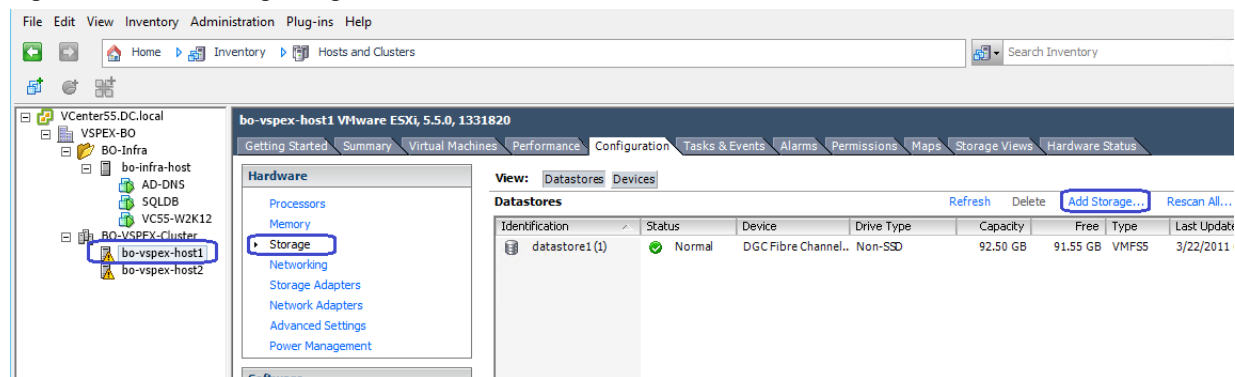
!	Name	NAS Server	Type	Description	Size	Protecti
✓	NFS-Datastore2	NASServer01 (10.10.20.10)	NFS	VMWare NFS Datastore for VSPE...	6.0 TB	Not conf
✓	NFS-Datastore1	NASServer01 (10.10.20.10)	NFS	VMware NFS Datastore for VSPE...	6.0 TB	Not conf

## Mount NFS share on ESXi hosts

1. Select the **Home > Inventory > Hosts and Clusters** tab in vCenter, expand the VSPEX cluster and select an ESXi host. Click the **Configuration** tab > **Storage > Add Storage**.

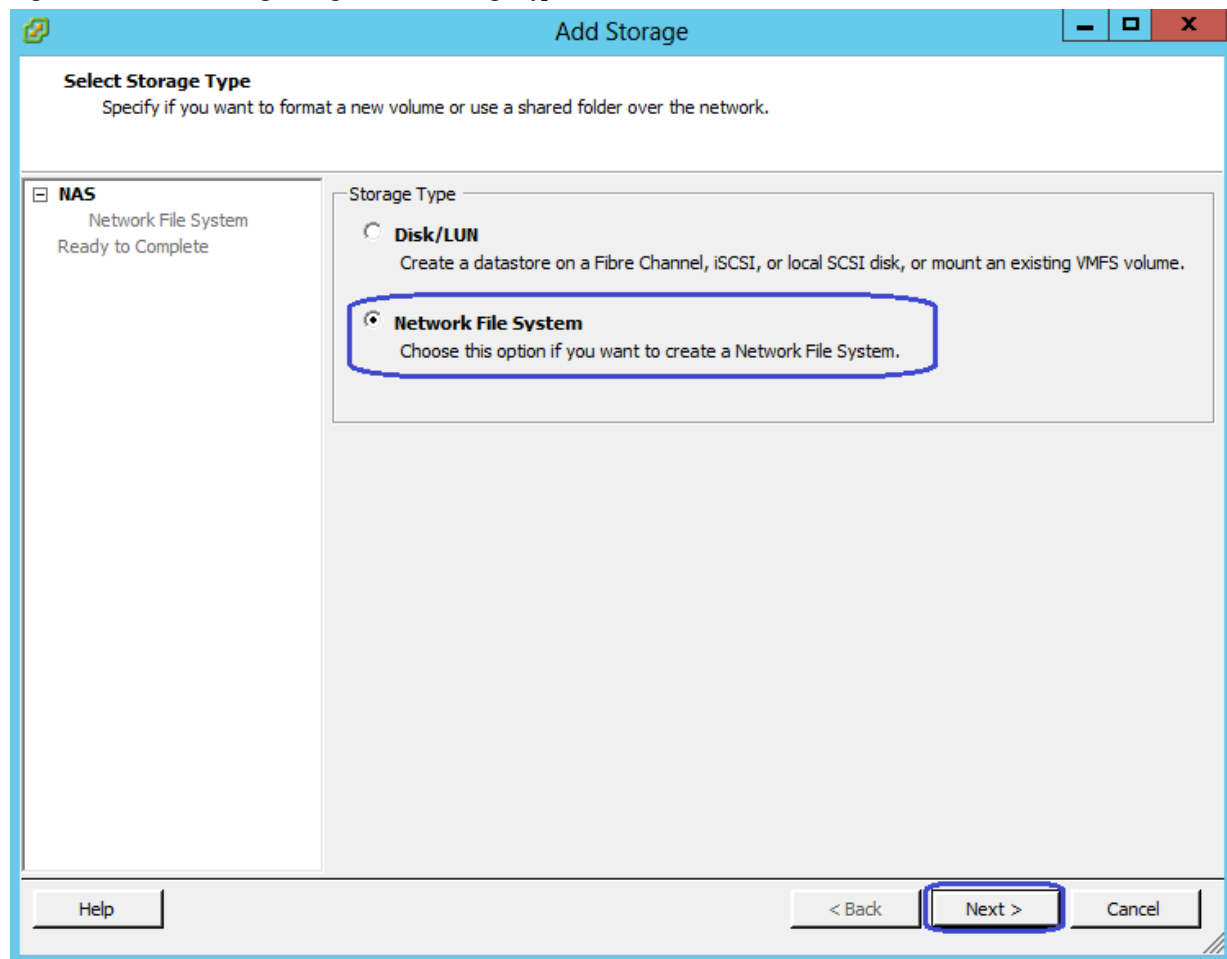


Figure 145 Adding Storage



- Click the **Network File System** radio button for the Storage Type and click **Next**.

Figure 146 Adding Storage - Select Storage Type



- Specify the NAS Server IP in the Server field and the VMware Datastore name in Folder field as we created in the EMC Unisphere wizard. Then, specify the Datastore name and click **Next**.

Figure 147 Adding Storage - Locate NFS

**Add Storage**

**Locate Network File System**  
Which shared folder will be used as a vSphere datastore?

**NAS**  
Network File System  
Ready to Complete

**Properties**

Server: 10.10.20.10  
Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:2AA:FF:FE9A:4CA2

Folder: /NFS-Datastore1  
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: NFS-DS1

Help < Back **Next >** Cancel

- After the successful creation, now you will see the new Storage for NFS-Data created.

Figure 148 Verify the Created Storage for NFS-Data

**bo-vspex-host1 VMware ESXi, 5.5.0, 1331820**

Getting Started Summary Virtual Machines Performance Configuration Tasks & Events Alarms Permissions Maps Storage Views Hardware Status

**Hardware**

- Processors
- Memory
- Storage**
- Networking
- Storage Adapters

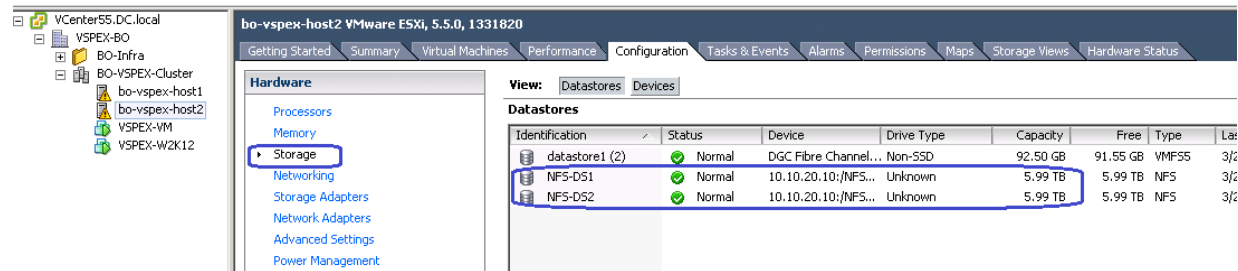
**View: Datastores Devices**

**Datastores** Refresh Delete Add Storage... R

Identification	Status	Device	Drive Type	Capacity	Free	Type
datastore1 (1)	Normal	DGC Fibre Channel..	Non-SSD	92.50 GB	91.55 GB	VMFS5
NFS-DS1	Normal	10.10.20.10:/NFS...	Unknown	5.99 TB	5.99 TB	NFS

- Repeat the steps from 1 to 3 to create NFS datastore for the remaining VMware Datastore.

**Figure 149** Verify the Created NFS Datastore for all the VMware Datastore



This concludes the NFS storage for VM datastores.

## Cisco UCS Mini Branch office Management with UCS Central

Cisco UCS Central Software extends the simplicity and agility of managing a single Cisco UCS domain to multiple Cisco UCS domains that can extend across globally distributed data centers. UCS Central provides a single point of management for thousands of UCS servers and provides centralized inventory, fault management, global ID pooling and centralized policy based firmware upgrades. UCS central therefore is ideal for managing a branch office VSPEX based solution, where it can provide the management console to drive consistency and compliance across all of the UCS domains.

UCS Central is supplied a a virtual appliance that runs on VMware and Microsoft hypervisors and is pre-packaged as a VMware ova or an ISO image for ease of installation. Redundancy can be provided with an active standby configuration (not supported across WAN links) and to ensure separation of the management plane it should be installed on separate servers that are not part of UCS Domain.

Profiles and policies defined in Cisco UCS Central (UCSC) can co-exist with the local Cisco UCS Manager defined information. Both Cisco UCS Manager and Cisco UCS Central manage the information defined in the respective tool and show the information defined in other as read-only.

For the Cisco Remote Office Branch office solutions, some of the key advantages of the Cisco UCS Central are:

- All the UCS resources, errors and warnings from two or more domains are presented in a single common interface
- Various Pools, Service Profiles and Settings are configured once, centrally
- Service Profiles can be managed and deployed from a single management pane
- Branch office setup can be managed from a series of templates ensuring that each branch has a consistent setup

In this design UCS Central will deployed as a standalone solution.

## Install and Configure UCS Central

In this architecture, we deployed UCS Central.ova appliance on the VMWare Hypervisor and configured in a standalone mode.

Following are the major steps to deploy Cisco UCS central on Remote VSPEX Primary data center.

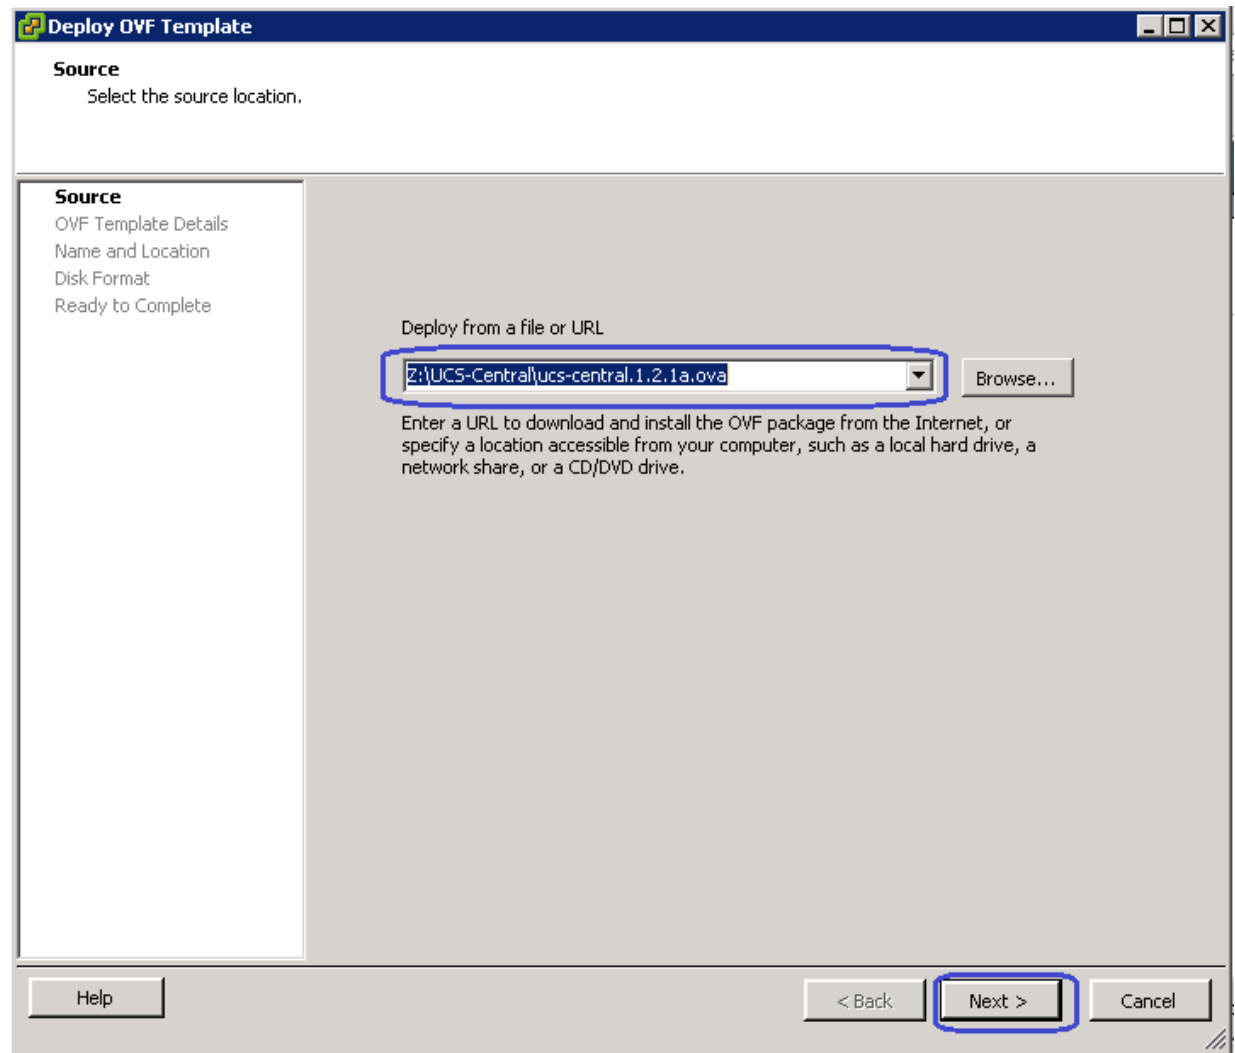
1. Install UCS Central
2. Configure UCS Central
3. Adding Cisco UCS Managers to Cisco UCS Central

## Install UCS Central

As mentioned before, the UCS Central installation media is available as VMware virtual machine OVF template. The UCS central must be deployed on the infrastructure network of the Remote VSPEX datacenter, and *not* on one of the VSPEX ESXi servers. Follow these steps to install VSM VM:

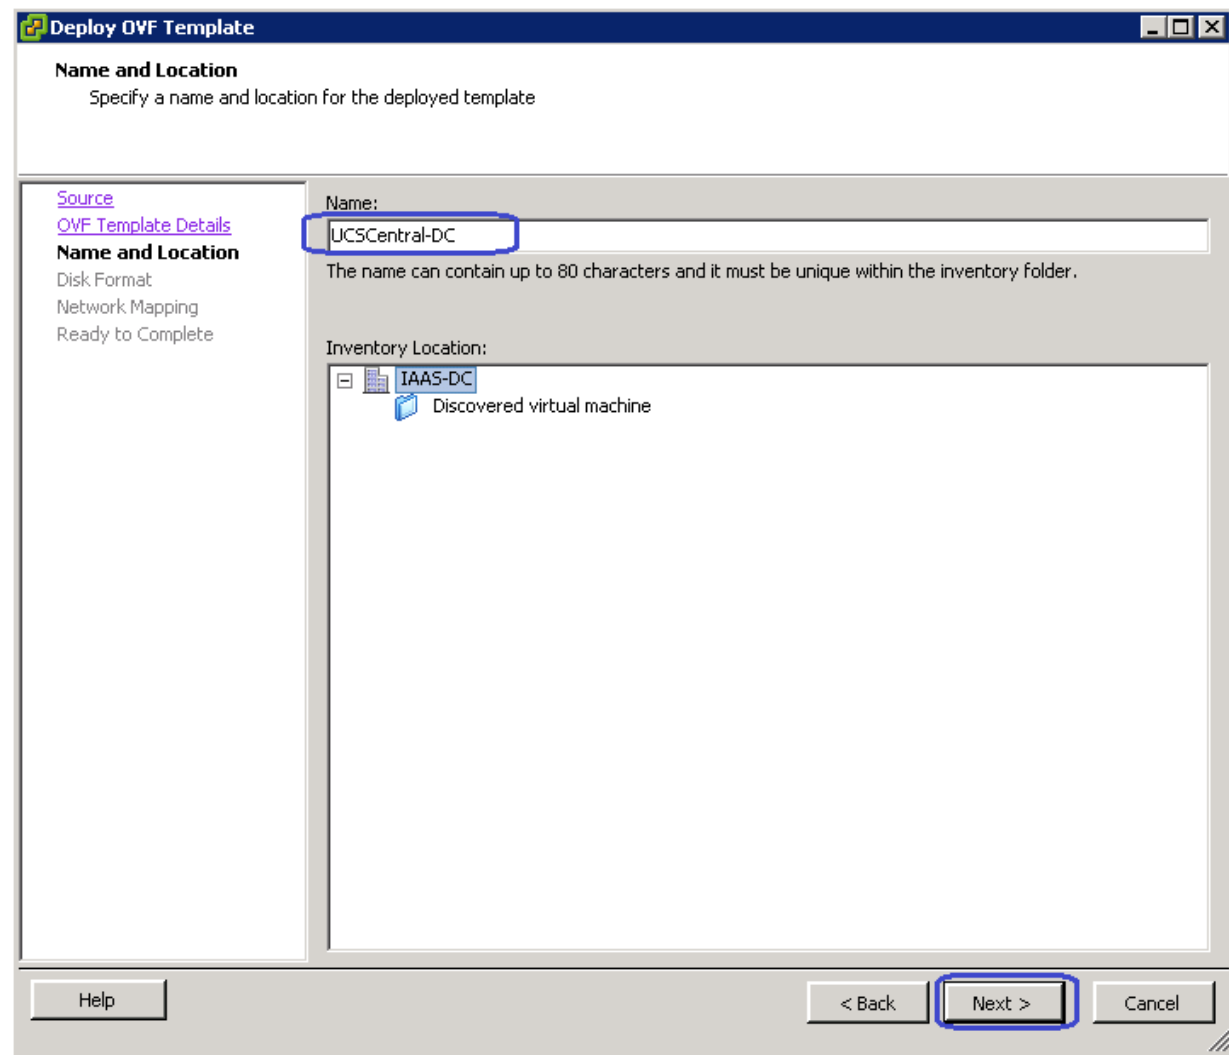
1. From the **Hosts and Cluster** tab in vCenter, select the infrastructure ESX/ESXi host and click **File** Deploy new Virtual Machine through OVF template. Select UCS central OVF, and click **Next**.

**Figure 150** Deploying OVF Template - Source



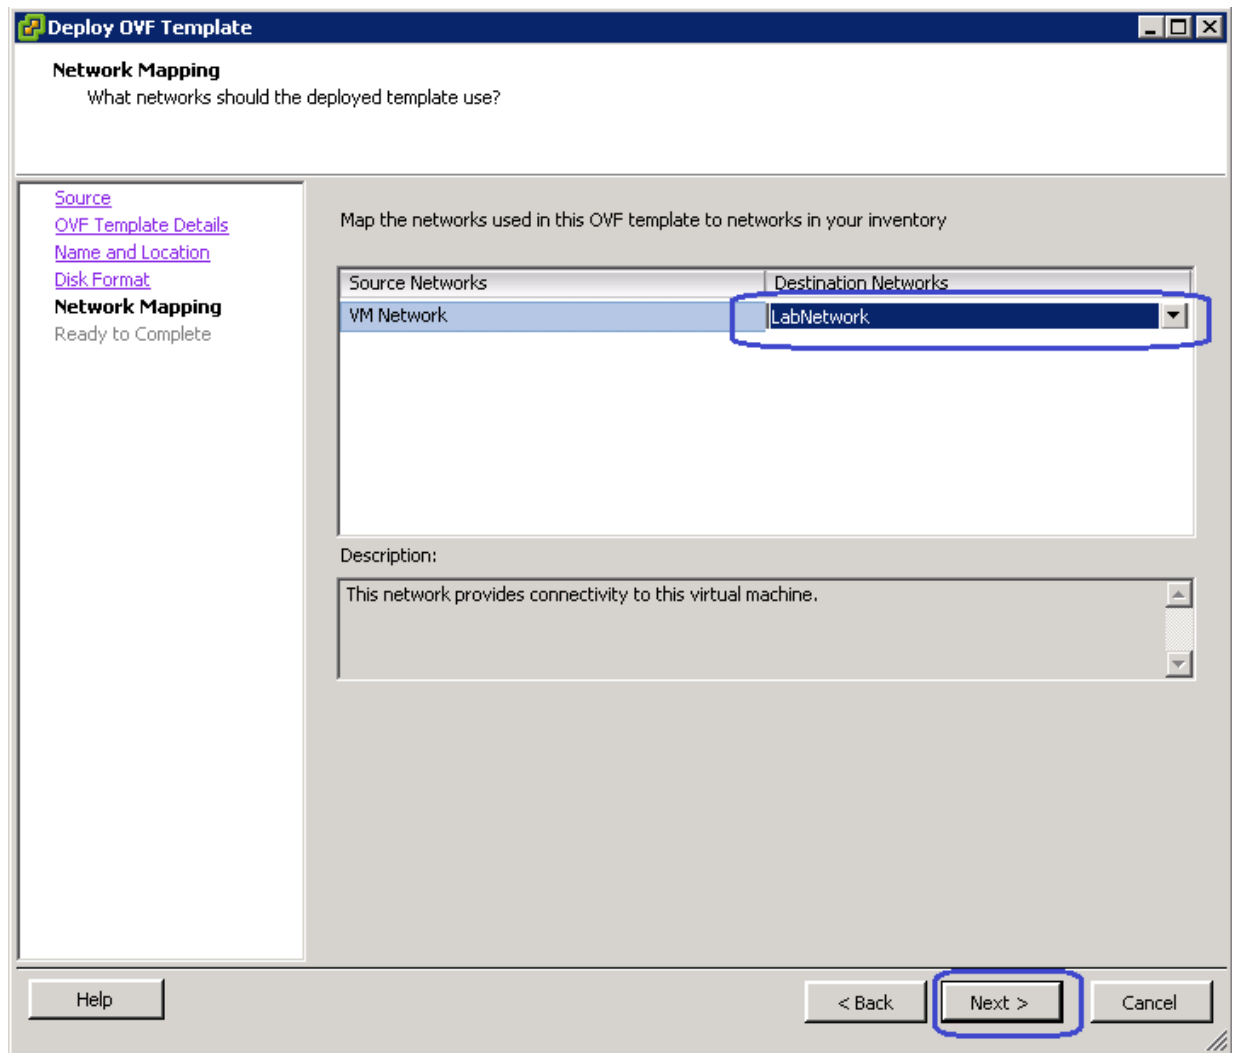
2. In the OVF template window, click **Next**. Then on the next window, Select the data center where you want to install the UCS central.

**Figure 151**      *Deploying OVF Template - Specify Name and Location*



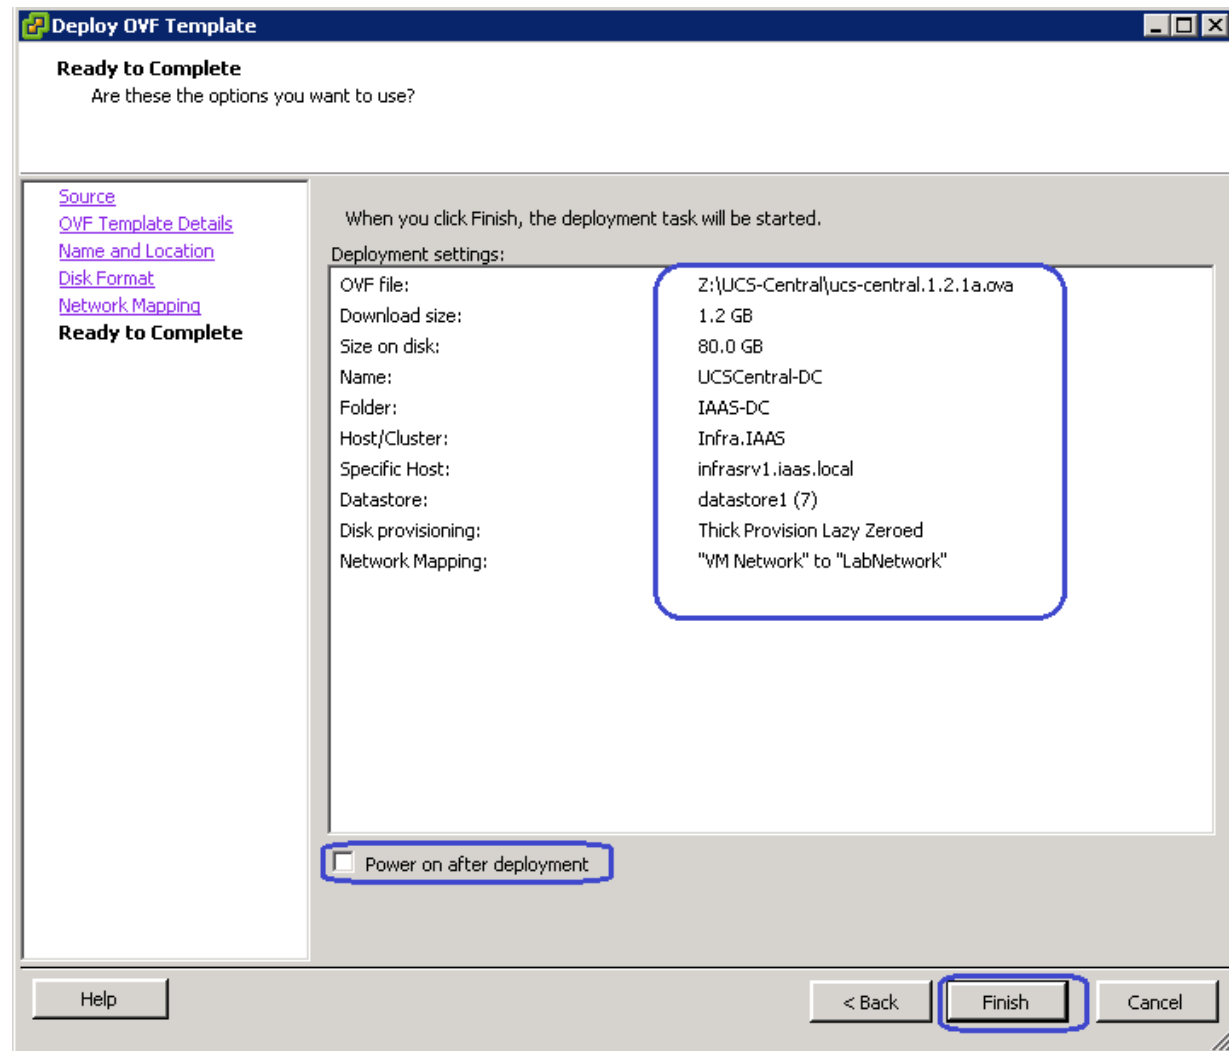
1. In the next window, let us restore the default values and click **Next**. In the next window, choose the Destination Network from the drop-down list.

Figure 152 Deploying OVF Template - Network Mapping



2. Verify the deployment settings and check the **Power on after deployment** check box and click **Finish**.

Figure 153 Deploying OVF Template - Verify in the Ready to Complete Summary



## Configure UCS Central

Open up a console window to the Cisco UCS Central VM. Once, the Cisco UCS Central VM has completed the initial part of the installation process, answer the following questions in the VM console window:

1. Setup new configuration or restore full-state configuration from backup [setup/restore] prompt, type **setup** and press **Enter**.
2. At the "Enter the UCS Central VM eth0 IPv4 Address:" Enter IP address **10.29.150.252** as assigned to Cisco UCS Central and press **Enter**.



**Note** You must enter a static IP address that is reserved for this Cisco UCS Central VM. Cisco UCS Central does not support Dynamic Host Configuration Protocol (DHCP).

3. At the "Enter the UCS Central VM eth0 IPv4 Netmask:" prompt, enter the netmask **255.255.255.0** as assigned to Cisco UCS Central and press **Enter**.

4. At the “Enter the Default Gateway:” prompt, enter the default gateway **10.29.150.1** as used by Cisco UCS Central and press **Enter**.
5. At the “Is this VM part of a cluster(select 'no' for standalone) (yes/no):” prompt, select **no** and press **Enter**.
6. At the “Enter the UCS Central VM host name:” prompt, enter the host name as **UCSCentral-DC** for the Cisco UCS Central VM and press **Enter**.
7. (Optional) At the “Enter the DNS Server IPv4 Address:” prompt, enter the IP address for the DNS server you want to use for Cisco UCS Central and press **Enter**.

If you do not plan to use a DNS server for Cisco UCS Central, leave this blank and press **Enter**.

8. (Optional) At the “Enter the Default Domain Name:” prompt, enter the domain in which you want to include Cisco UCS Central and press **Enter**.

If you do not plan to include Cisco UCS Central in a domain, leave this blank and press **Enter**. Cisco UCS Central will use the default domain named localdomain.

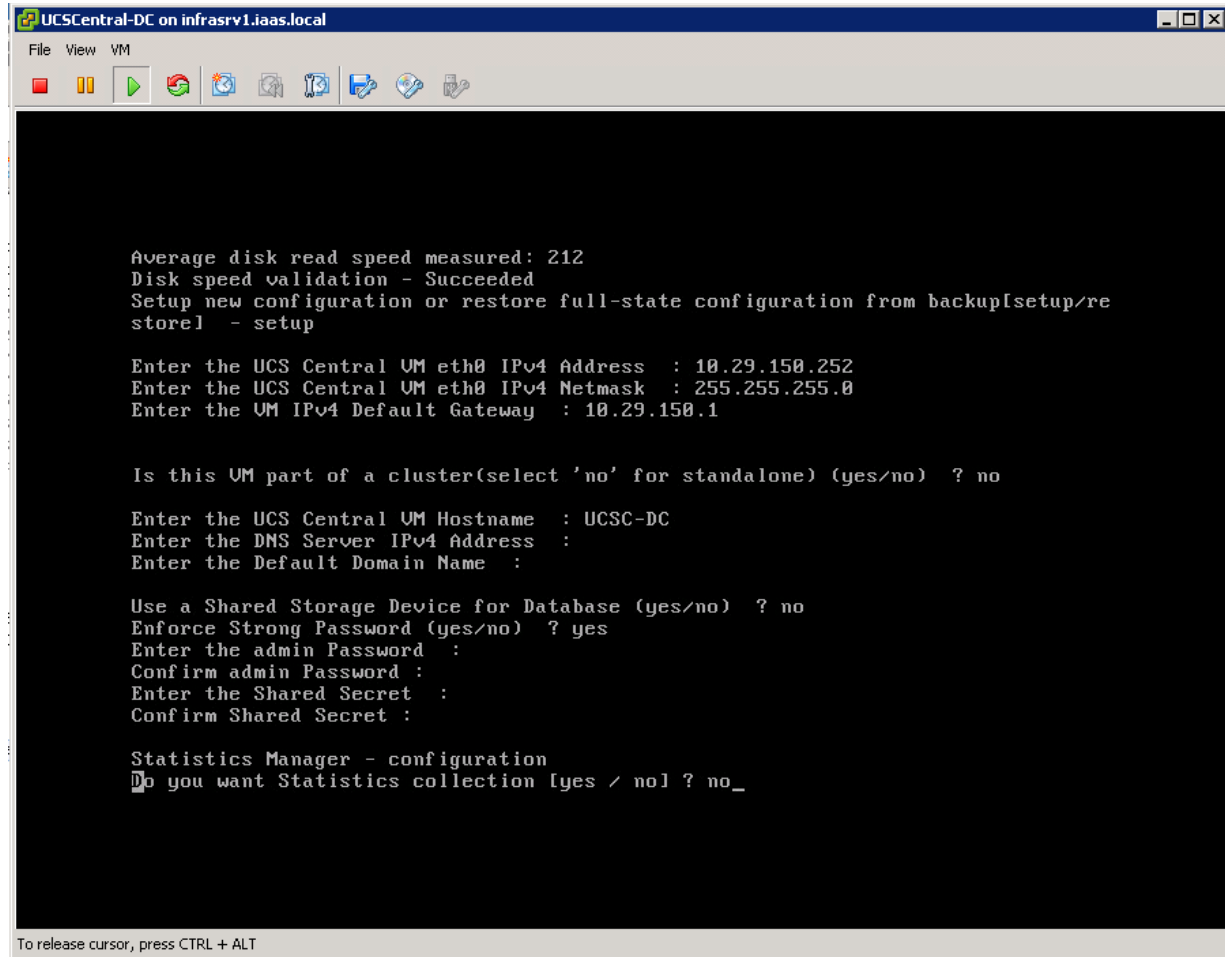
9. At the “Use Shared Storage Device for Database (yes/no):” prompt, if you want to setup shared storage, enter yes, if not enter no and press **Enter**.
10. At the “Enforce Strong Password(Yes/No):” prompt, if you want to set up strong password alert, select **yes** and press **Enter**.
11. At the “Enter the admin Password:” prompt, enter the password you want to use for the admin account and press **Enter**.
12. At the “Confirm admin Password:” prompt, re-enter the password you want to use for the admin account and press **Enter**.
13. At the “Enter the Shared Secret:” prompt, enter the shared secret (or password) that you want to use to register one or more Cisco UCS domains with Cisco UCS Central and press **Enter**.
14. At the “Confirm Shared Secret:” prompt, re-enter the shared secret and press **Enter**.
15. At the “Do you want Statistics Collection (yes/no):” prompt, if you want to enable statistics collection, enter **yes** and press **Enter**.

If you do not want to enable statistics collection now, you can enter no and proceed with the installation. You can enable the statistics collection using Cisco UCS Central CLI at any time.

16. For the “Proceed with this configuration, please confirm[yes/no]:” prompt, enter **yes** and press **Enter**.

After you confirm that you want to proceed with the configuration, the network interface reinitializes with your settings and Cisco UCS Central becomes accessible via the IP address.

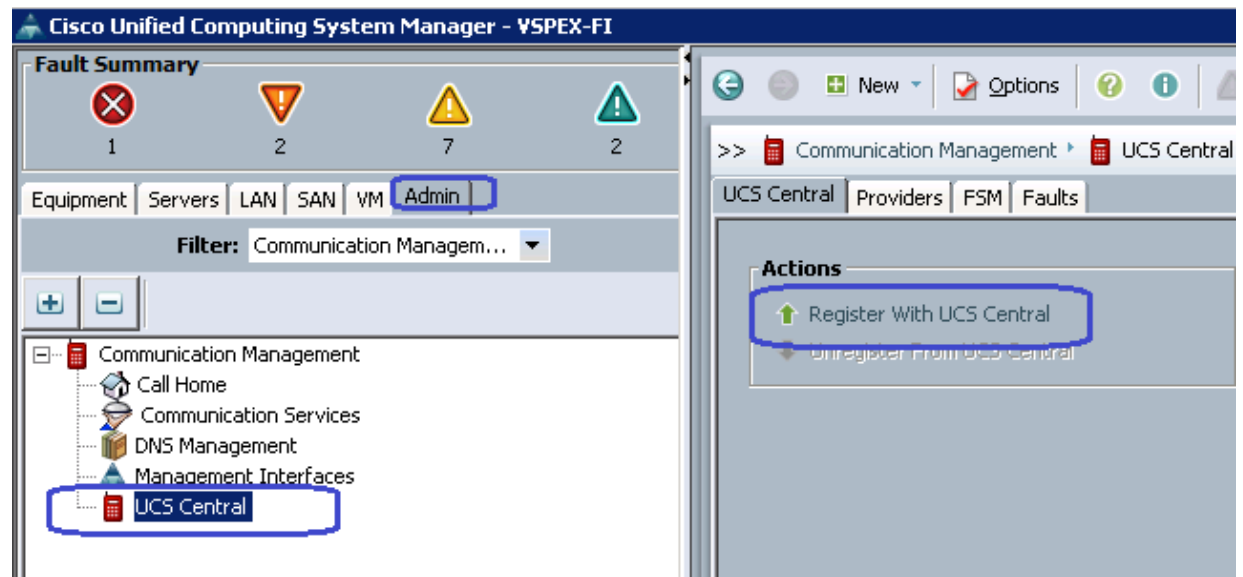




## Adding Cisco UCS Managers to UCS central

1. Launch UCS manager GUI using the UCS mini Virtual IP.

Figure 154 Registering with Cisco UCS Central



2. Specify the IP address of the UCS Central and Specify the Shared Secret password we created on the previous section.

Figure 155 Registering with UCS Central Window

**Register With UCS Central**

Hostname/IP Address: **10.29.150.252**

Shared Secret: \*\*\*\*\*

**Policy Resolution Control**


**All Global**

Infrastructure & Catalog Firmware:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the Capability Catalog and infrastructure firmware policy are defined locally or come from Cisco UCS Central.
Time Zone Management:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the time zone and NTP server settings are defined locally or comes from Cisco UCS Central.
Communication Services:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether HTTP, CIM XML, Telnet, SNMP, web session limits, and Management Interfaces Monitoring Policy settings are defined locally or in Cisco UCS Central.
Global Fault Policy:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the Global Fault Policy is defined locally or in Cisco UCS Central.
User Management:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether authentication and native domains, LDAP, RADIUS, TACACS+, trusted points, locales, and user roles are defined locally or in Cisco UCS Central.
DNS Management:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether DNS servers are defined locally or in Cisco UCS Central.
Backup & Export Policies:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the Full State Backup Policy and All Configuration Export Policy are defined locally or in Cisco UCS Central.
Monitoring:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether Call Home, Syslog, and TFTP Core Exporter settings are defined locally or in Cisco UCS Central.
SEL Policy:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the SEL Policy is defined locally or in Cisco UCS Central.
Power Allocation Policy:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the Global Power Allocation Policy is defined locally or in Cisco UCS Central.
Power Policy:	<input type="radio"/> Local <input checked="" type="radio"/> Global	Determines whether the Power Policy is defined locally or in Cisco UCS Central.

**OK** **Cancel**

- You will see a pop-up confirmation window confirming the registration with UCS central. Click **Accept**.

**Figure 156** Confirmation Window for Registering UCS Central

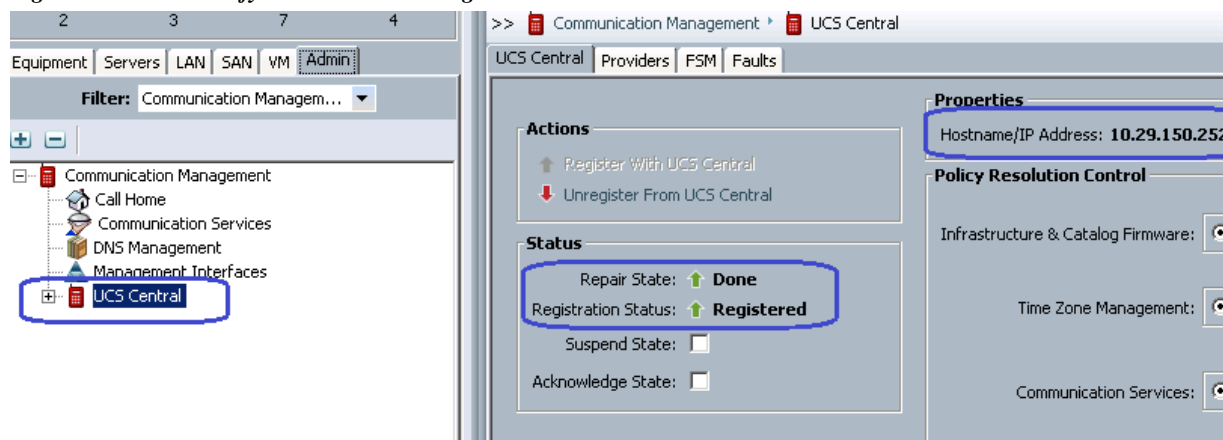


- Now you will see the UCS Central is registered successfully on the UCS Manager.



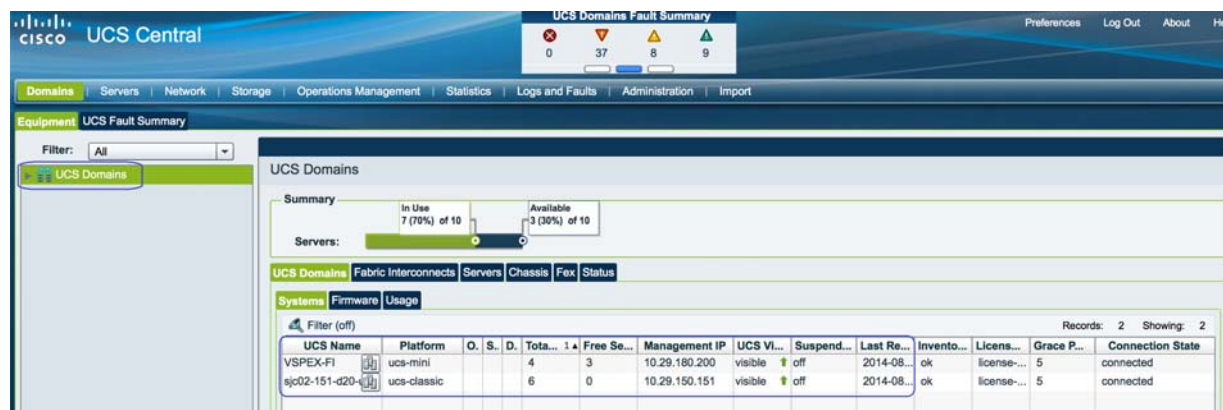
**Note** Make sure that the UCS central and UCS manager date & time is in sync. If not in sync, the UCS central registration will not be successful. (it is recommended to use NTP server for time sync).

**Figure 157** Verify the UCS Central Registration



- You can also login to UCS Central web GUI and view the connected UCS Manager for UCS mini and traditional UCS manager in Remote DC.

**Figure 158** UCS Domain in UCS Central GUI



## Configure UCS Central for Branch Office Deployment

Following are the major steps to configure UCS Mini from UCS Central.

1. Configure UCS Central Domain Group
2. Configure Pools and Policies
3. Configure Global Service Profile Template
4. Configure Global Service Profile Instance

### UCS Central Domain Group

Cisco UCS Central creates a hierarchy of Cisco UCS domain groups for managing multiple Cisco UCS domains. You will have the following categories of domain groups in Cisco UCS Central:

- **Domain Group**—A group that contains multiple Cisco UCS domains. You can group similar Cisco UCS domains under one domain group for simpler management.
- **Ungrouped Domains**—When a new Cisco UCS domain is registered in Cisco UCS Central, it is added to the ungrouped domains. You can assign the ungrouped domain to any domain group.

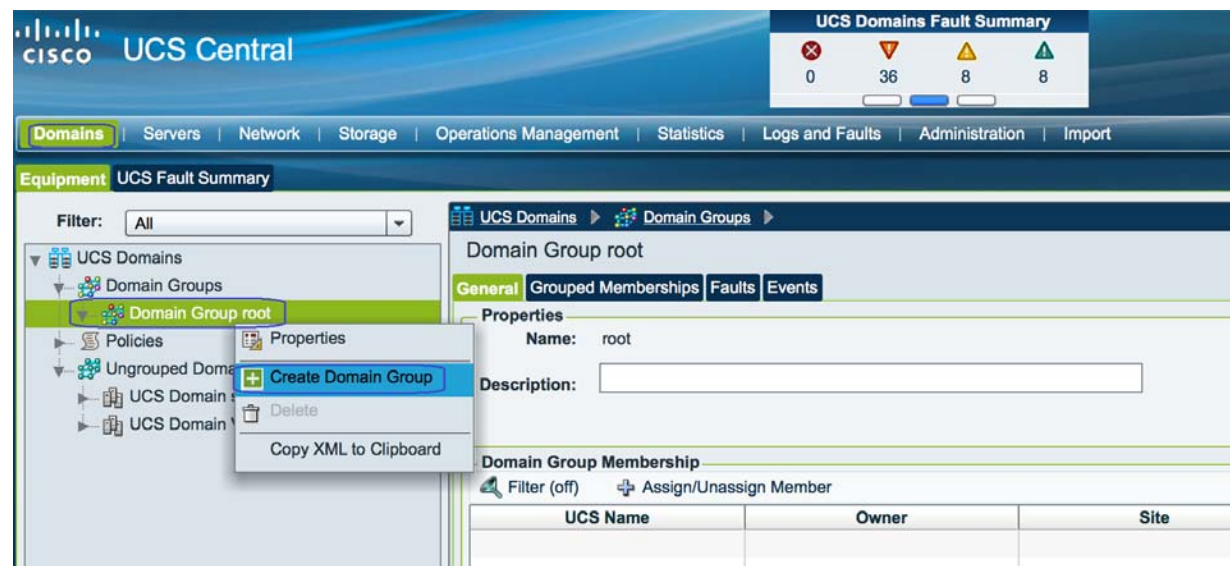
If you have created a domain group policy, a new registered Cisco UCS domain meets the qualifiers defined in the policy, it will automatically be placed under the domain group specified in the policy. If not, it will be placed in the ungrouped domains category. You can assign this ungrouped domain to a domain group.

Each Cisco UCS domain can only be assigned to one domain group. You can assign or reassign membership of the Cisco UCS domains at any time. When you assign a Cisco UCS domain to a domain group, the Cisco UCS domain will automatically inherit all management policies specified for the domain group.

Follow these steps for UCS central Domain group creation:

1. Launch UCS Central web GUI, then click **Domains > Domain Groups > Domain Group root** and then click **Create Domain Group**.

*Figure 159 Creating a Domain Group*



2. Specify the Domain group name and description and Click **OK**.

Figure 160 Creating a Domain Group - Define Domain Group

**Create**

## Create Domain Group

**Properties**

**Name:** BranchOffice-1

**Description:** Domain group for Branch office in Site 1

OK Close

- Click Ungrouped Domains and click the discovered Branch office UCS Domain **VSPEX-FI** and in the right pane click **Change Group Assignment**.

Figure 161 Changing Group Assignments

**UCS Central**

**UCS Domains Fault Summary**

0 36 8 8

**Domains** Servers Network Storage Operations Management Statistics Logs and Faults Administration Import

**Equipment: UCS Fault Summary**

Filter: All

**UCS Domains**

- Domain Groups
  - Domain Group root
  - Domain Group BranchOffice-1
- Policies
- Ungrouped Domains
  - UCS Domain sjc02-151-d20-ucs
  - UCS Domain VSPEX-FI**

**UCS Domain VSPEX-FI**

**General** Status Faults Events Audits Tech Support Files

Launch UCS Manager

**Change Group Assignment**

Re-evaluate Membership

**Fault Summary**

0 2 8 3

**Properties**

In Use: 1 (25%) of 4

Available: 3 (75%) of 4

Servers: 1 (25%) of 4

Management IP: 10.29.180.200

FW Version: 3.0(1c)A

FW Status: ready

Domain Group:

Last Modified: 2014-08-06T13:16:06

Description:

Site:

Owner:

Inventory Status: ok

**Fabric Interconnects** Servers Chassis Fex IO Modules

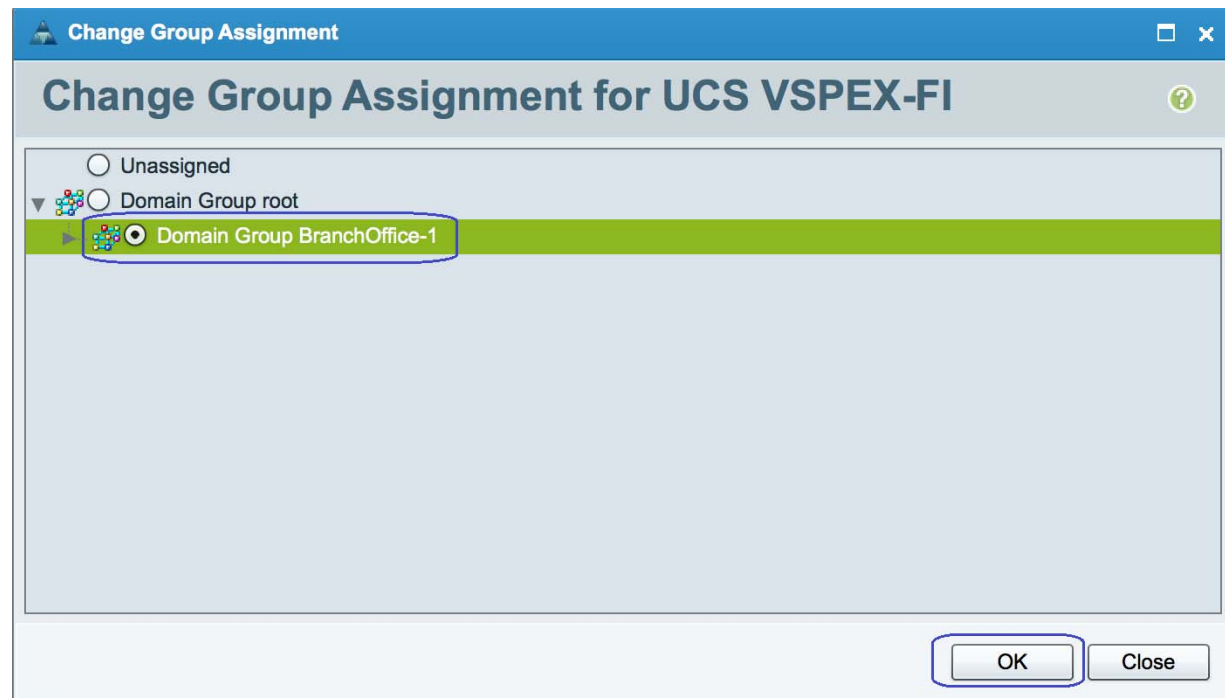
**Status** Hardware Firmware

Filter (off)

ID	Status	Operability	Power	Performance	Thermal
Fabric Interconnect A		operable			
Fabric Interconnect B		operable			

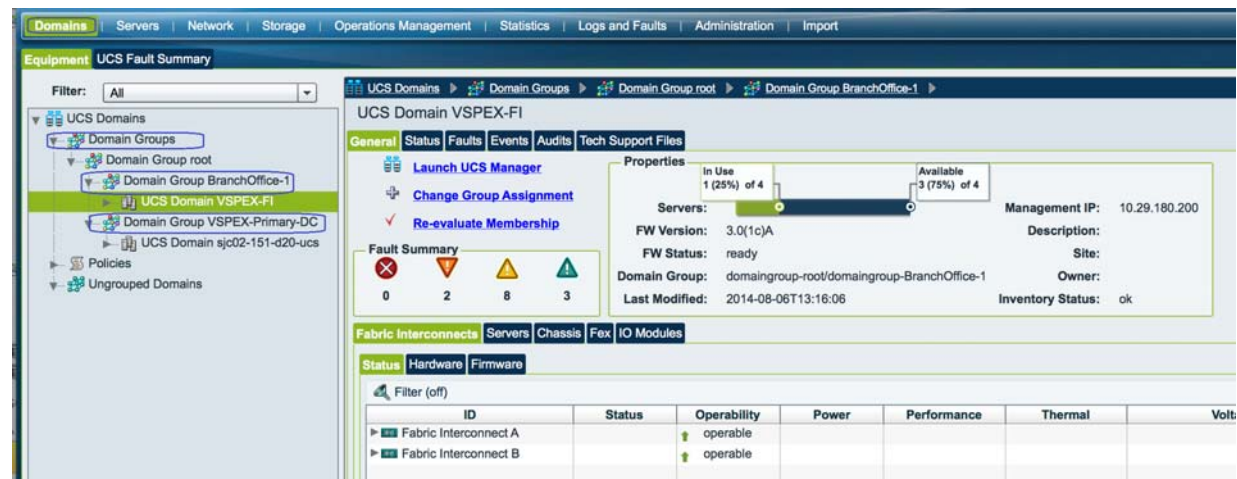
- Choose the Domain Group as **Branch Office-1** and click **OK**.

**Figure 162** *Changing Group Assignments - Select the Domain Group*



- Repeat the same steps 1, 2, 3, and 4 to add the discovered VSPEX primary UCS Domain to the domain group created for Primary DC. After the successful creation, you can see both VSPEX Primary DC domain and Branch office Domain added to their respective groups.

**Figure 163** *Summary Showing Added Domain Groups*



## Cisco UCS Central Image Management

Cisco UCS software bundles should be downloaded to Cisco UCS Central for later use in host firmware management or Cisco UCS system upgrades. The server and infrastructure images can be uploaded by navigating to Operations Management.



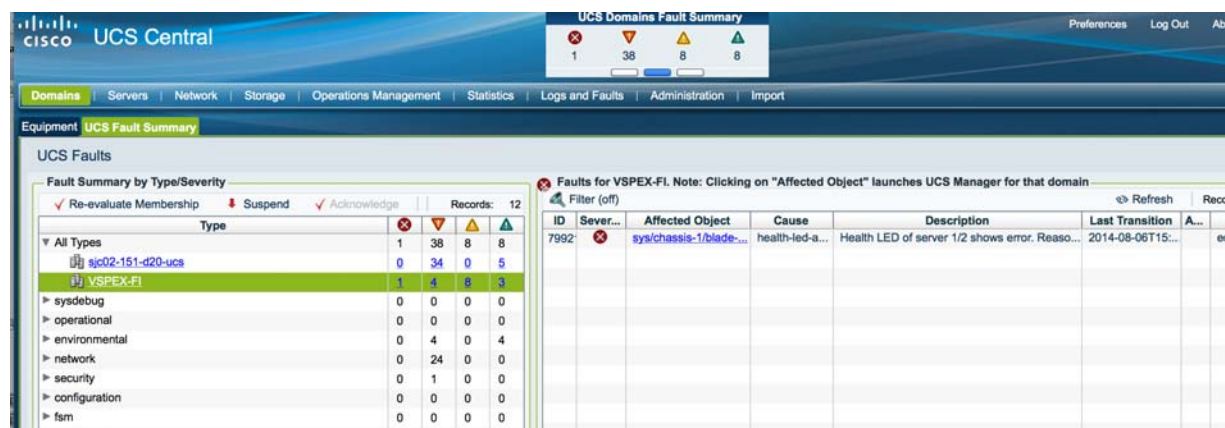
Figure 164 Downloading Software Bundles to UCS Central



## Cisco UCS Central Fault Management

Cisco UCS Central globally manages Fault and Error management for all registered UCS domains in single pane. Figure 165 shows the UCS faults for VSPEX primary DC site and VSPEX Branch office site managed by UCS central in single pane.

Figure 165 UCS Central GUI Showing Fault Summary



## Global Identifier Management

Global identifier management addresses one of the biggest challenges around multi-domain management: unique address management for system identifiers (MAC's, WWxN's, UUID's, etc.). Previously, UCS Manager best practices recommend embedding a "domain ID" within the high-order bytes of the ID pool ranges. However, this still involved manual intervention and could be error prone.

With UCS Central, all the ID pools can be defined and accessed globally across all UCS domains. Service Profile assignment can be guaranteed unique and non-overlapping with respect to ID's across all UCS domains. Global ID Pools belong to the organization structure. Global pools do not terminate on DGs, as the UCS Central "Operational Policies" do. Instead, the range of Global ID Pools extends across all UCS domains in the scope of the organization structure within UCS Central, regardless of any DG partitioning UCS Central provides visibility in to possible duplicate ID usage.



All of the pool types (UUID, MAC, WWxN) offer the ability to display duplicate IDs that may exist across UCS domains, through the “ID Usage Summary”. Duplicate ID severity will be flagged as either “Major”, for IDs that appear in multiple Service Profiles, or flagged as “Warning” for IDs that appear in multiple local pools. Note that the only way to view Local ID Pool consumption is to select an individual ID, and view the corresponding drill-down details to the right (Local Pool and Local Service Profile) conflicting pool assignments are reported as faults. Unallocated IDs that belong to overlapping pools are reported as warnings.

**Figure 166** *Detecting presence of Overlapping IDs or Duplicate ID*

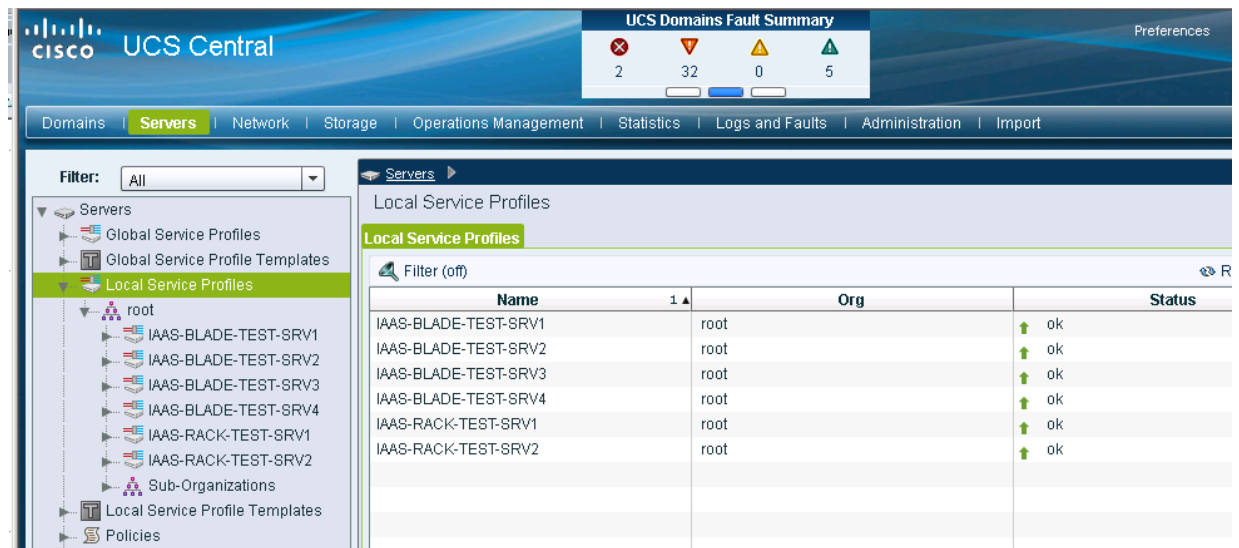
The screenshot shows the Cisco UCS Central interface. The left sidebar displays a tree view of the network configuration, with 'MAC Pools' selected. The main pane shows the 'ID Usage Summary' table. The table has columns for 'Fault Status', 'ID', 'Local Pools', 'Global Pools', 'Domains', and 'Service Profiles'. The first row shows a 'Warning' status for ID '00:25:B5:A2:00:8F' with 2 local pools, 1 global pool, 2 domains, and 2 service profiles. Subsequent rows show 'Warning' status for various other MAC addresses, each with 1 local pool, 1 global pool, 1 domain, and 1 service profile.

Fault Status	ID	Local Pools	Global Pools	Domains	Service Profiles
Warning	00:25:B5:A2:00:8F	2	1	2	2
Warning	00:25:B5:B2:00:FF	1	1	1	1
Warning	00:25:B5:FF:00:00	0	1	0	1
Warning	00:25:B5:FF:00:10	0	1	0	1
Warning	00:25:B5:FF:00:20	0	1	0	1
Warning	00:25:B5:FF:00:30	0	1	0	1
Warning	00:25:B5:FF:00:40	0	1	0	1
Warning	00:25:B5:FF:00:50	0	1	0	1
Warning	00:25:B5:FF:00:60	0	1	0	1
Warning	00:25:B5:FF:00:70	0	1	0	1
Warning	00:25:B5:FF:00:80	0	1	0	1
Warning	00:25:B5:FF:00:90	0	1	0	1

## Cisco UCS Central Service Profile Management

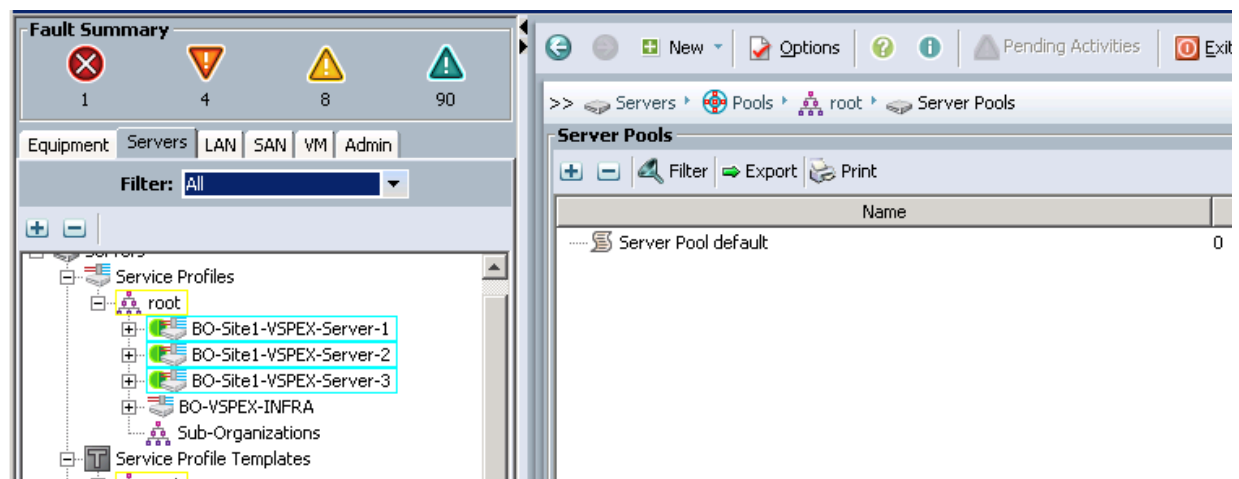
Cisco UCS Central manages both local Service profile and global service profile in a single pane. Since Local service profiles are completely managed by UCS Manager, the UCS central has limited options to manage these local service profiles. The Cisco UCS Central shows limited options and different icon for a locally defined Service Profile.

**Figure 167** *Locally Defined Service Profiles*



Cisco UCS Manager displays a green circle next to the Global Service Profiles and most of the configuration options are grayed out for a globally defined Service Profile.

**Figure 168** *Global Service Profiles Shown in UCS Manager*



The global and locally defined information conforms to the following key principles at this time:

- Existing local Services profiles templates cannot be imported into Cisco UCS Central.
- Existing local service profiles cannot be assigned to Cisco UCS Central.
- Existing local policies (for example local disk policies) can be made Global and therefore be used by global service profile templates.
- Globally defined policies can be used by local service profiles.
- Global Service Profiles can be made local but once localized, these service profiles cannot be assigned back to Cisco UCS Central.

**Note**

For this solution, unique policies, templates and pools were defined in Cisco UCS Central. If possible, a suffix should be added to the names of globally defined profiles, pools and policies to uniquely signify the value definition in Cisco UCS Central.

## Configuring Cisco UCS Central

Cisco UCS Central configuration is very similar to the Cisco UCS Manager configuration. Cisco UCS Central taps are also in line with the Cisco UCS Manager's tabs for Server, Network, and Storage. Following the steps to configure Cisco UCS Manager, the following parameters need to be configured in Cisco UCS Central:

- IP Pools for Management (Network | IP Pools | global-ext-mgmt)
- Server Pools (Servers | Pools | Server Pools)
- UUID Suffix Pools (Servers | Pools | UUID Suffix Pools)
- MAC Address Pools (Network | Pools | MAC Pools)
- WWNN Pools (Storage | Pools | WWN Pools | WWNN)
- WWPN Pools (Storage | Pools | WWN Pools | WWNN)
- Boot Policies (Servers | Policies | Boot Policies)
- BIOS Policy (Servers | Policies | BIOS Policies)
- Host Firmware Policy (Servers | Policies | Host Firmware Packages)
- Power Control Policy (Servers | Policies | Power Control Policies)
- vNIC/vHBA Placement Policy (Servers | Policies | vNIC/vHBA Placement Policies)
- vNIC Template (Network | Policies | vNIC Templates)
- vHBA Template (Storage | Policies | vHBA Templates)
- Service Profile Templates (Servers | Global Service Profile Templates)

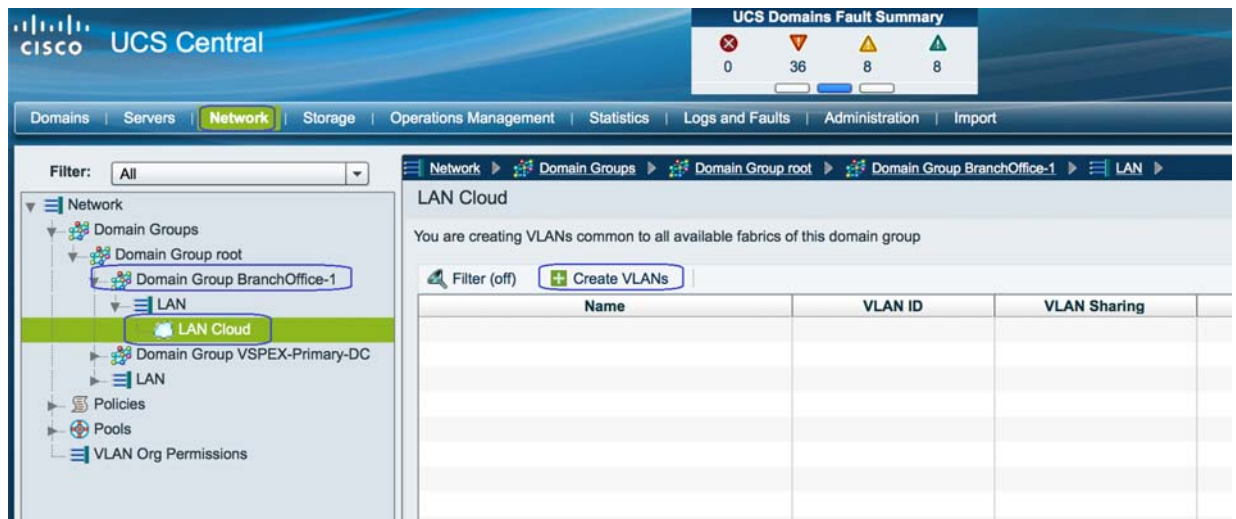
With all the configurations in place, Service Profiles can be deployed on both Primary DC Cisco UCS Domain and Branch office Cisco UCS domains from Cisco UCS Central.

## Configuring VLANs and VSANs

This section covers the creation of VLAN, VSAN, Pools, Policies, Service profile template and Service profile for Branch office VSPEX servers using UCS Central Management web GUI.

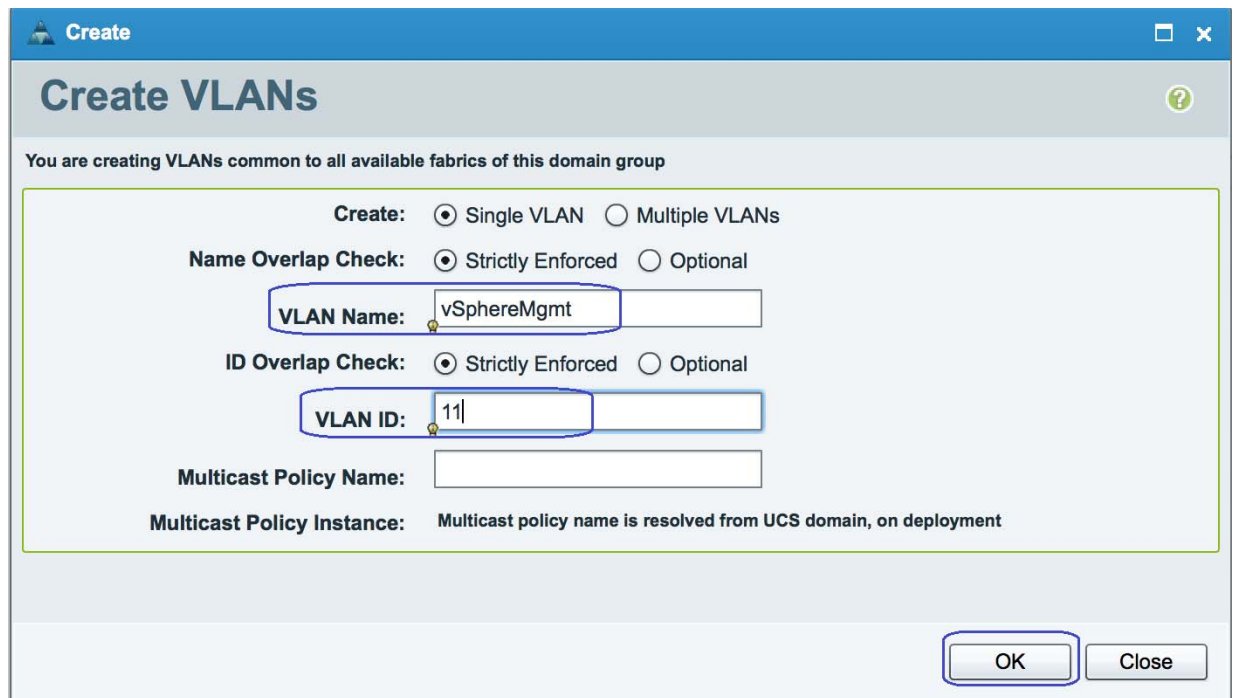
1. Launch UCS Central web GUI and click the Domain Group **BranchOffice-1**. Click **LAN Cloud > Create VLANs**.

**Figure 169** *Creating VLANs in UCS Central*



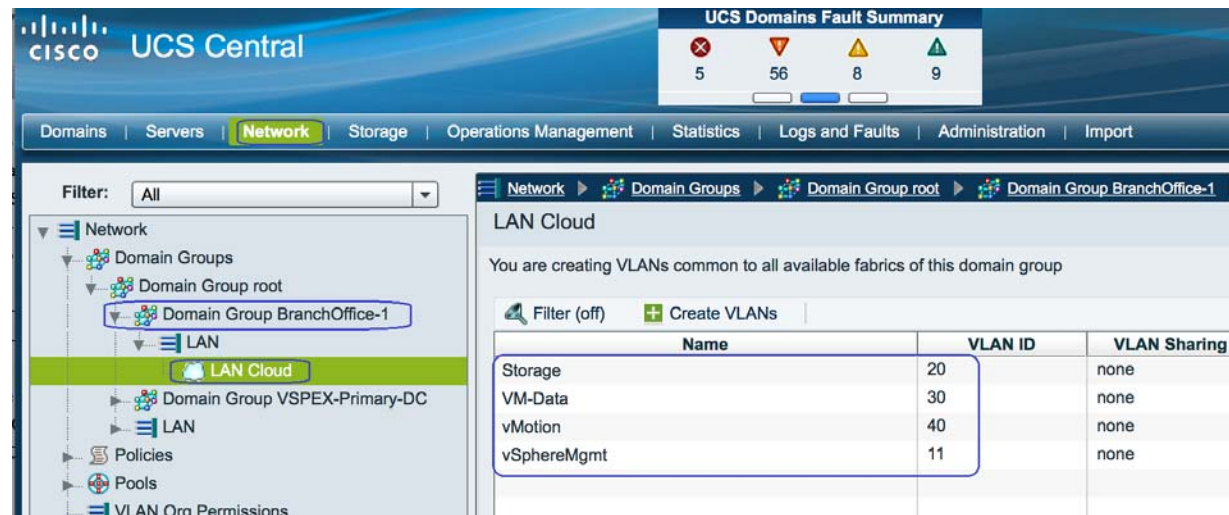
- Specify the VLAN Name and VLAN ID. Click **OK**.

**Figure 170** *Creating VLANs Window*



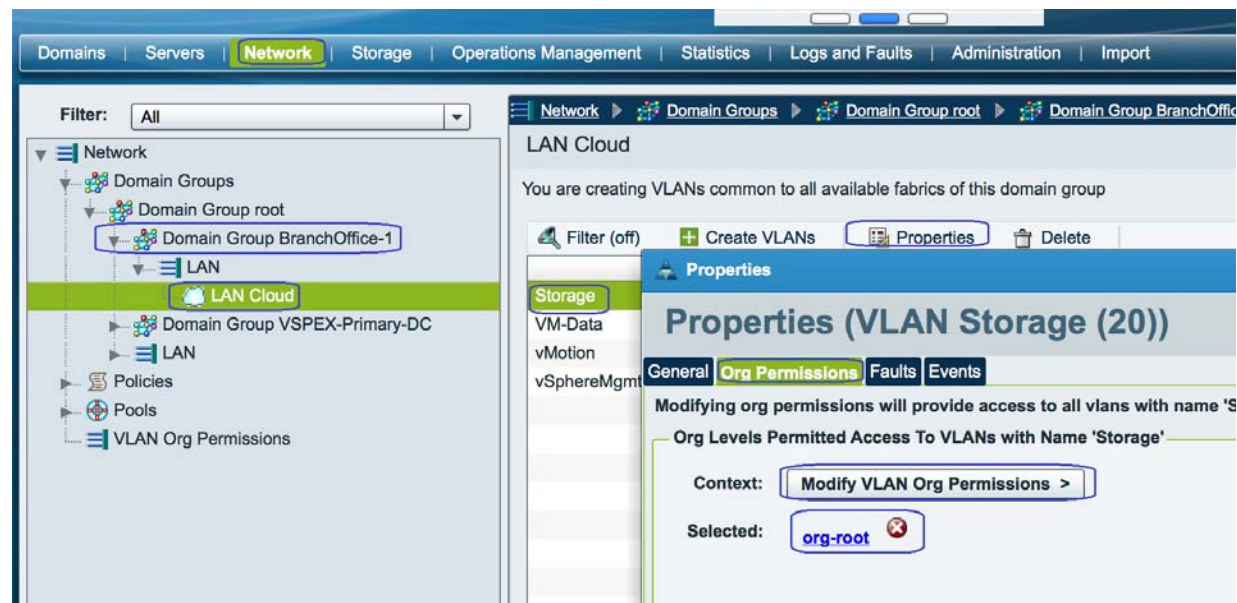
- Repeat these steps to create VLANs for Storage - VLAN20, VM-Data - VLAN30 and vMotion -VLAN40.
- After the successful VLAN creation, you can see all the created VLANs.

Figure 171 Created VLANs are Shown in UCS Central



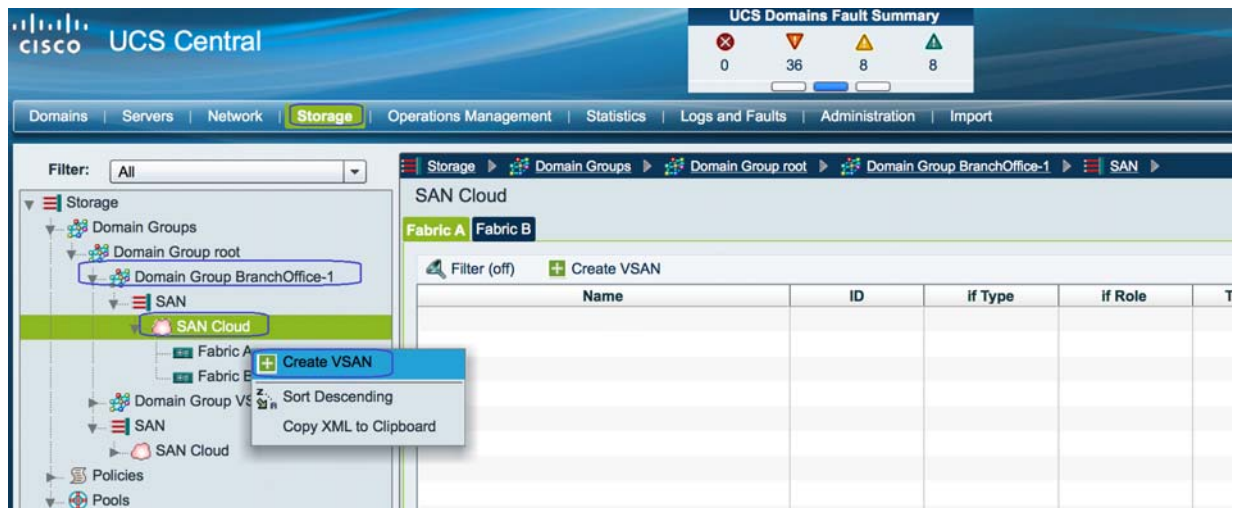
- Click the VLAN properties, we have access permissions for **Org-root**.

Figure 172 Setting Org Permissions



- To create VSAN, Click the **Domain Group BranchOffice-1** and Click **SAN Cloud > Create VSAN**.

Figure 173 Creating VSAN in UCS Central



- Specify the VSAN name and for FC zoning settings, click **Enabled** radio button and Specify the VSAN ID and FCoE ID.

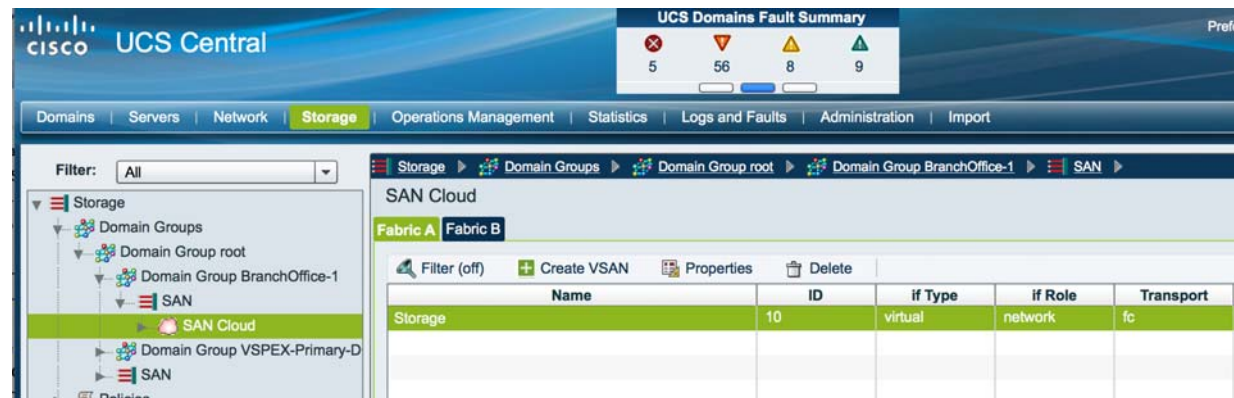
Figure 174 Creating VSAN Window

The screenshot shows the 'Create VSAN' dialog box. The 'Name' field is set to 'Storage'. Under 'FC Zoning Settings', the 'FC Zoning' radio button is set to 'Enabled'. Below this, a warning message states: 'Do NOT enable zoning for this VSAN, if the fabric interconnect is connected to an upstream switch that has zoning enabled on the same VSAN.' The 'VSAN ID' field is set to '10'. The 'FCoE VLAN' field is set to '10'. At the bottom right, there are 'OK' and 'Close' buttons.

- After the successful creation, you can see the VSAN Storage under SAN cloud.



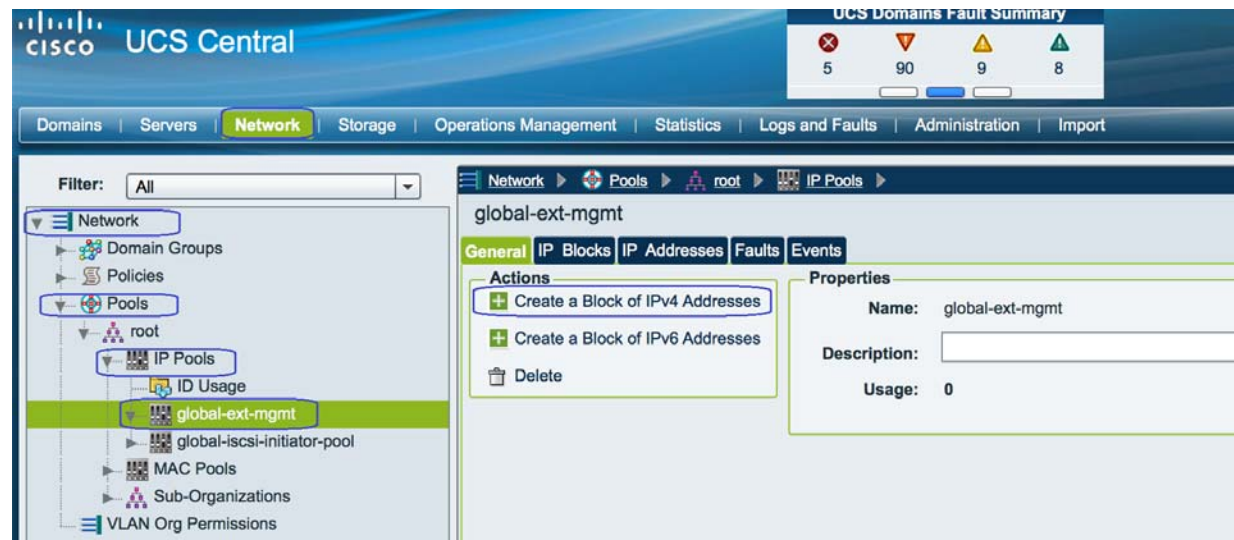
Figure 175 Verify the Created VSAN for Storage



## Configuring Pools and Policies

1. Launch UCS Central web GUI and Click **Network > Pools > IP Pools > global-ext-mgmt** and click **Create a Block of IPv4 Addresses**.

Figure 176 Creating a Block of IPv4 Addresses



2. Specify the IP range for the IP Pool for managing the UCS blades KVM console. And click **OK**.

Figure 177 Creating a Block of IPv4 Addresses Window

**Create**

## Create a Block of IPv4 Addresses

**From:** 10 . 29 . 180 . 225 **Size:** 8

**Subnet:** 255 . 255 . 255 . 0

**Default Gateway:** 10 . 29 . 180 . 1

**Primary Dns:** 0 . 0 . 0 . 0

**Secondary Dns:** 0 . 0 . 0 . 0

**Scope:** public

**ID Range Qualification Policy**

+ Create ID Range Qualification Policy

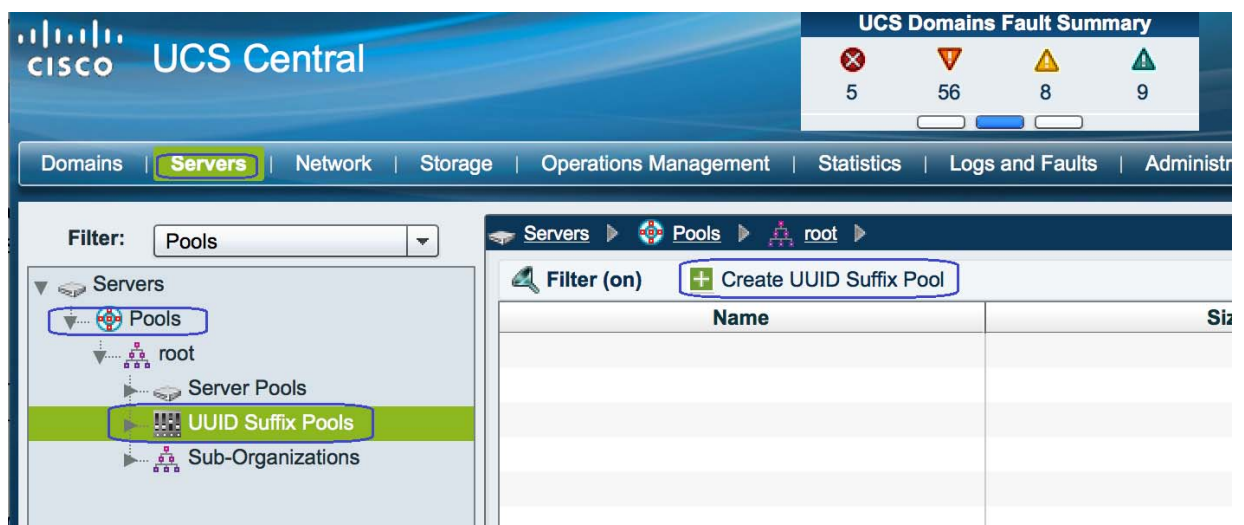
Warning: Pools containing an ID block referencing ID range qualification policy can only be used by local service profiles. Global service profiles cannot use pools referencing this policy.

**ID Range Qualification Policy:**

OK Close

- To create UUID pools, click **Servers > Pools > UUID Suffix Pools** and then click **Create UUID Suffix Pool**.

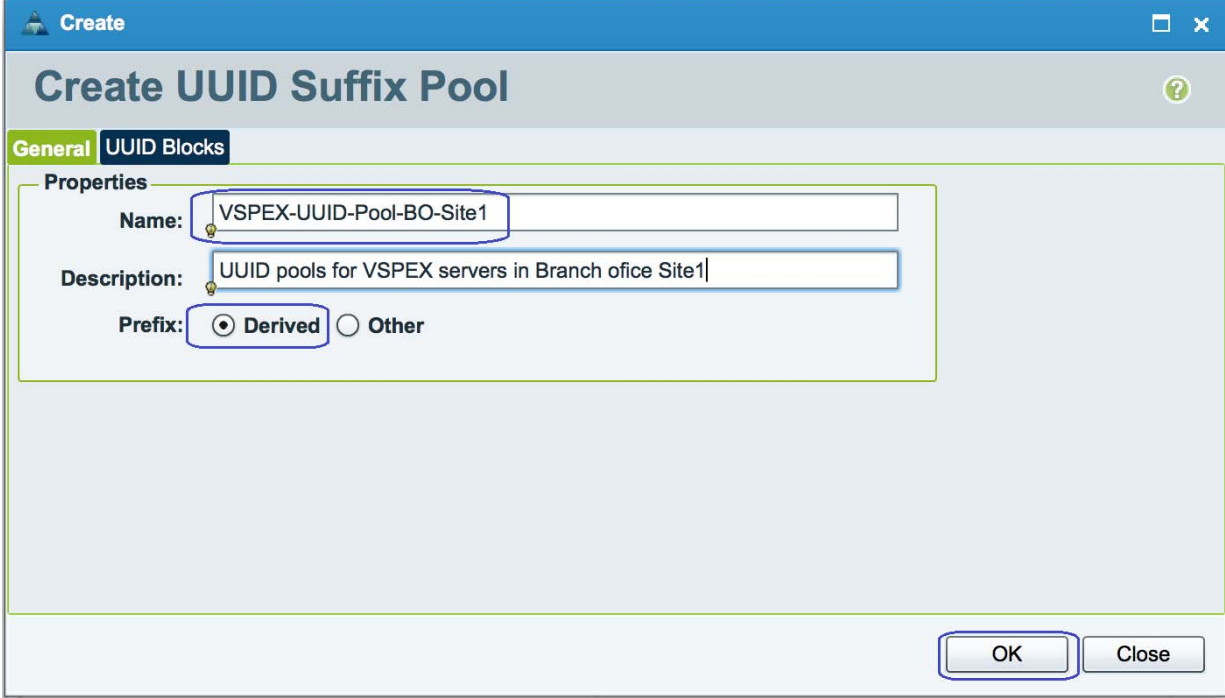
Figure 178 Creating UUID Suffix Pool



- In the **General** tab, specify the UUID Pool name and description. Click **OK**.

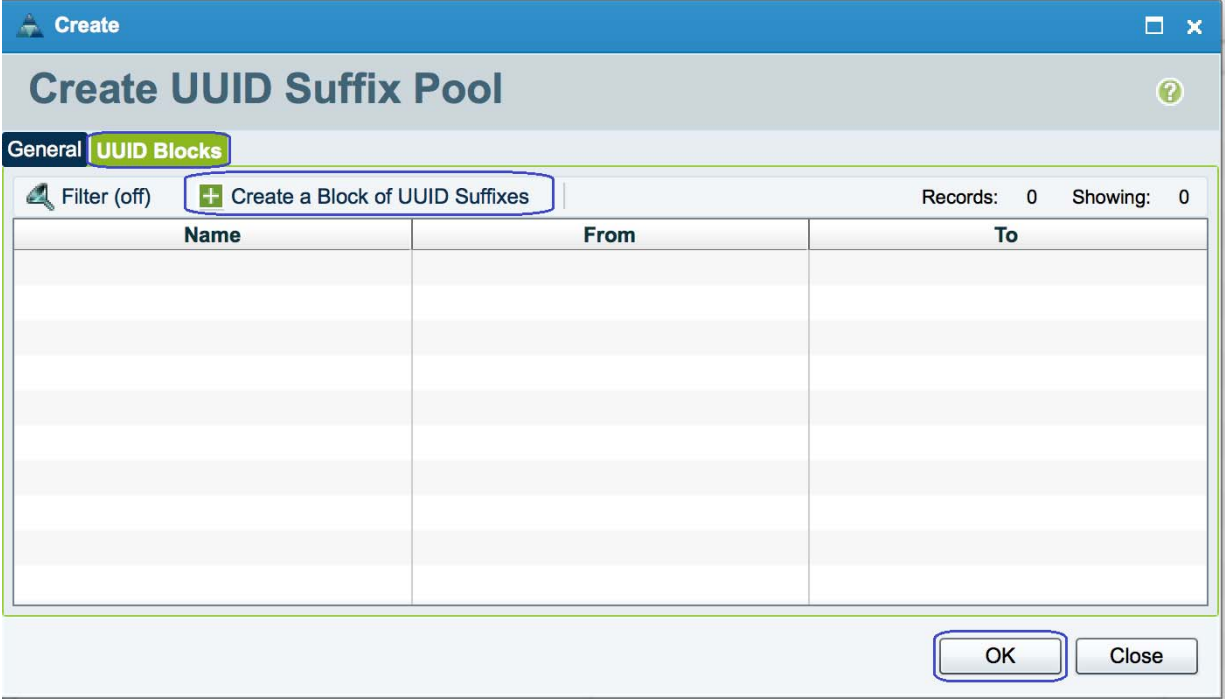


***F***



- 5

***F***



- 6

Figure 181 Creating a Block of UUID Suffixes Window

**Create**

## Create a Block of UUID Suffixes

From:  Size:

**ID Range Qualification Policy**

[+ Create ID Range Qualification Policy](#)

Warning: Pools containing an ID block referencing ID range qualification policy can only be used by local service profiles. Global service profiles cannot use pools referencing this policy.

ID Range Qualification Policy:

**OK** **Close**

- To create UUID pools, click **Network > Pools > MAC Pools** and click **Create MAC Pool**.

Figure 182 Creating MAC Pools

**CISCO UCS Central**

UCS Domains Fault Summary

0 36 8 8

Domains | Servers | **Network** | Storage | Operations Management | Statistics | Logs and Faults | Administration | Import

Filter: All

**Network**

- Domain Groups
- Policies
- Pools**
  - root
    - IP Pools
    - MAC Pools**
      - ID Usage
      - global-default
      - Sub-Organizations
    - VLAN Org Permissions

**MAC Pools**

Filter (off) [+ Create MAC Pool](#)

Name	Size
global-default	0

- In the **General** tab, specify the MAC Pool name and description and click **OK**.

**Figure 183**      **Creating MAC Pools Window**

**Create**

# Create MAC Pool

**General** **MAC Blocks**

**Properties**

**Name:** VSPEX-MAC-Pools-BO-Site1

**Description:** MAC Pools for VSPEX servers in Branch office Site1

**OK** **Close**

9. In the **MAC Blocks** tab, click **Create a Block of MAC Addresses** and then click **OK**.

**Figure 184**      *Creating a Block of MAC Addresses*

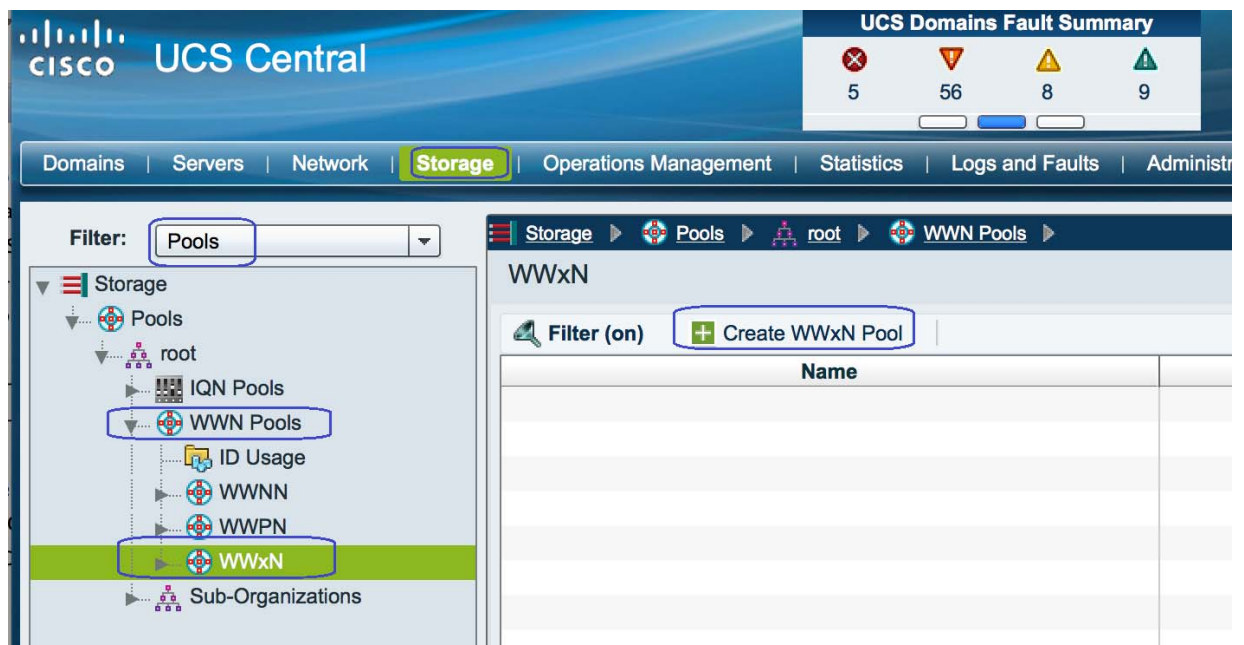
The screenshot shows a window titled "Create" with standard OS window controls. The main title bar is blue. Below it, the window has a light gray header area with the title "Create MAC Pool" and a help icon (?). The main content area has two tabs: "General" and "MAC Blocks". The "MAC Blocks" tab is selected and highlighted in yellow. Inside this tab, there's a search filter section with a magnifying glass icon, the text "Filter (off)", and a button labeled "+ Create a Block of MAC Addresses". To the right of this section, it says "Records: 0" and "Showing: 0". Below this is a large table with three columns: "Name", "From", and "To". The table has multiple rows, all of which are currently empty. At the bottom right of the window, there are two buttons: "OK" and "Close". Red boxes highlight the "MAC Blocks" tab, the "+ Create a Block of MAC Addresses" button, the table headers, and the "OK" button.

- 10. Specify the MAC range and size. Click **OK**.**

Figure 185 Creating a Block of MAC Addresses Window

11. To create WWxN Pools for WorldWideNodeName and WorldWidePortName. Click **Storage > Pools > WWN Pools > WWxN** and click **Create WWxN Pool**.

Figure 186 Creating WWxN Pool



12. In the **General** tab, specify WWxN pool name and description and choose **3 Ports per Node** from the drop-down list.

**Figure 187**      **Creating WWxN Pool Window**

**Create WWxN Pool**

**General** | WWN Initiator Blocks

**Properties**

Name: VSPEX-WWxN-Pool-BO-Site1

Description: WWNN and WWPN pools for VSPEX Server in Branch office Site1

Purpose: Node and Port WWN Assignment

Max Ports per Node: 3 Ports Per Node

OK Close

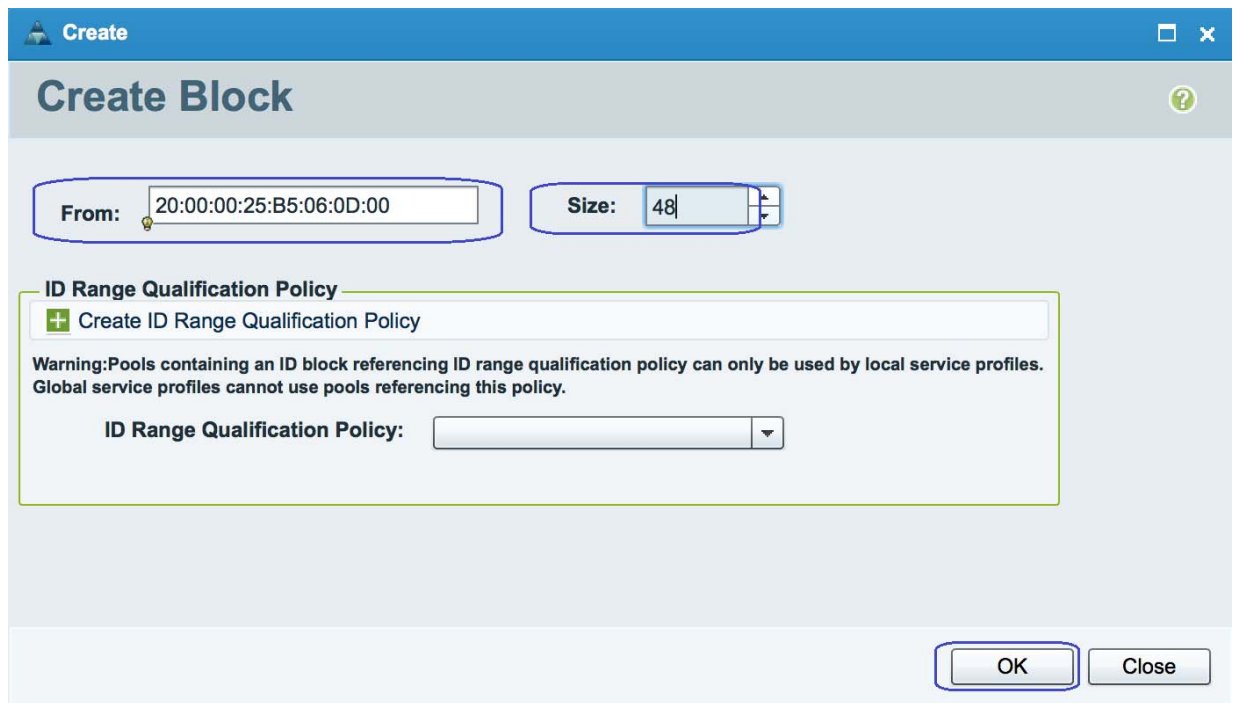
- 13. From the WWN Initiator Blocks tab, click **Create Block** and then click **OK**.**

**Figure 188**      **Creating WWxN Initiator Blocks**

The screenshot shows a software window titled "Create WWxN Pool". At the top left is a blue header bar with a small icon and the word "Create". Below this is a light gray title bar containing the main title "Create WWxN Pool" and a green question mark icon. The main area has two tabs: "General" and "WWN Initiator Blocks", with the latter being selected and highlighted in yellow. Under the "WWN Initiator Blocks" tab, there's a toolbar with a magnifying glass icon labeled "Filter (off)" and a button with a plus sign labeled "+ Create Block". To the right of the toolbar, it says "Records: 0" and "Showing: 0". Below the toolbar is a large table with three columns: "Name", "From", and "To". The table has multiple rows, all of which are currently empty. At the bottom right of the window, there are two buttons: "OK" and "Close".

- 14.** Specify the range for the Block and click **OK**.

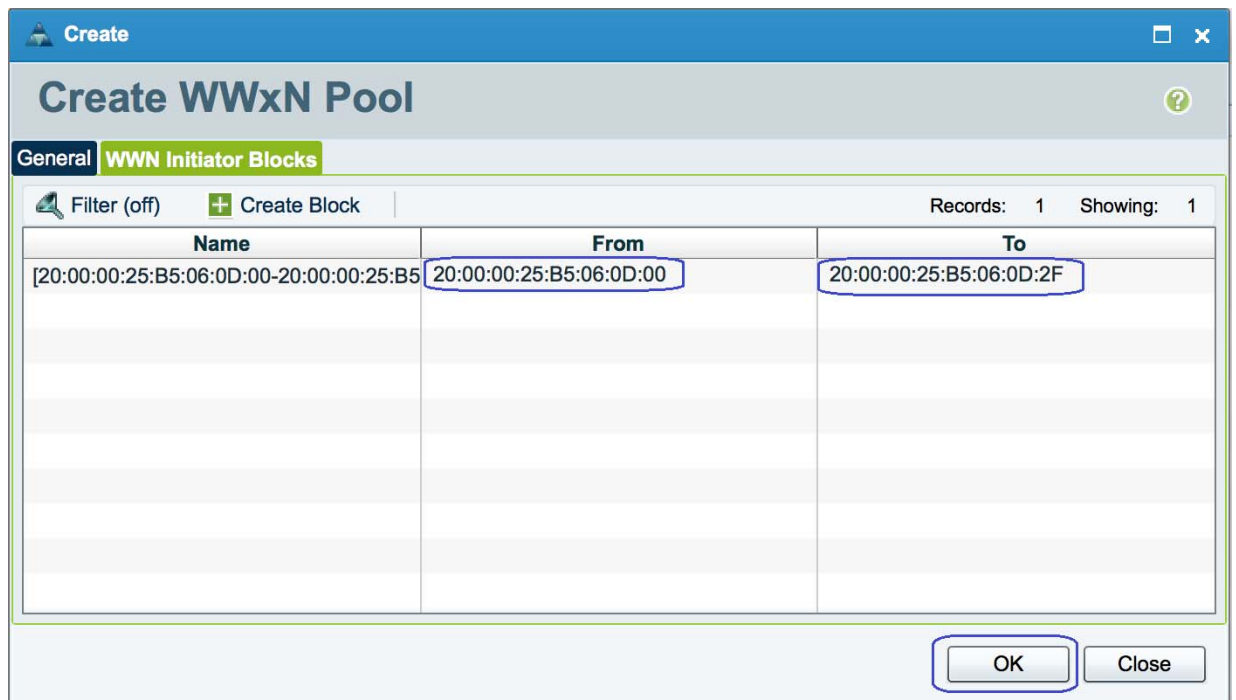
Figure 189 Creating WWxN Block Window



The 'Create Block' window has a blue title bar with a 'Create' icon and window controls. The main title is 'Create Block'. Below the title bar, there are two input fields: 'From' with the value '20:00:00:25:B5:06:0D:00' and 'Size' with the value '48'. Below these fields is a section titled 'ID Range Qualification Policy' with a '+ Create ID Range Qualification Policy' button. A warning message states: 'Warning: Pools containing an ID block referencing ID range qualification policy can only be used by local service profiles. Global service profiles cannot use pools referencing this policy.' Below the warning is a dropdown menu for 'ID Range Qualification Policy'. At the bottom right are 'OK' and 'Close' buttons.

15. After the successful creation, the WWN block will be created.

Figure 190 Verify the Created WWxN Initiator Block

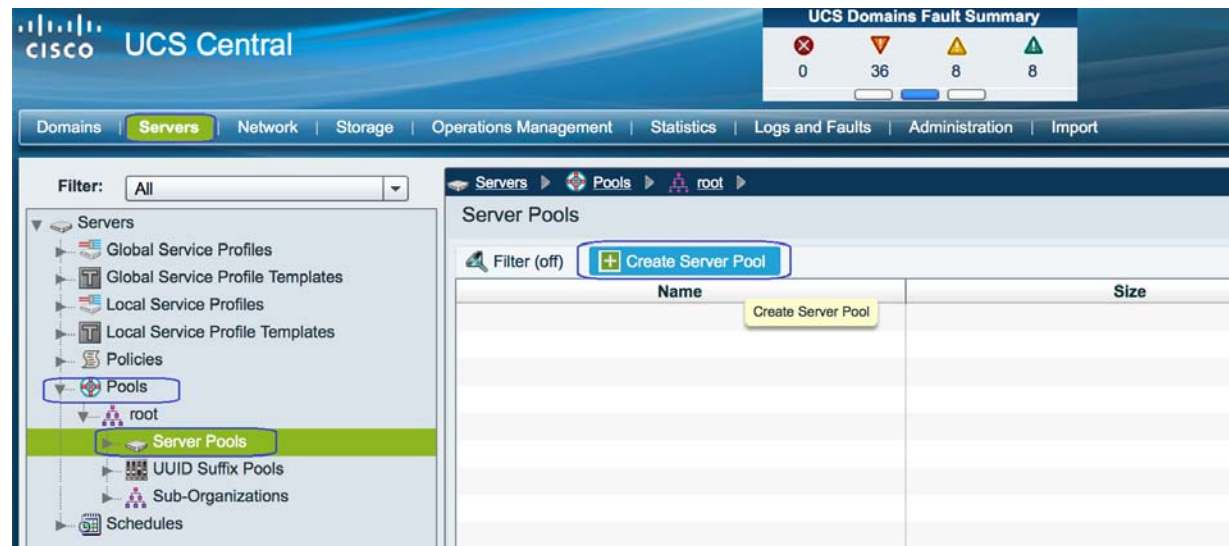


The 'Create WWxN Pool' window has a blue title bar with a 'Create' icon and window controls. The main title is 'Create WWxN Pool'. Below the title bar, there are two tabs: 'General' and 'WWN Initiator Blocks'. The 'WWN Initiator Blocks' tab is selected. Below the tabs, there is a 'Filter (off)' button and a '+ Create Block' button. To the right, it says 'Records: 1 Showing: 1'. Below this is a table with three columns: 'Name', 'From', and 'To'. The table contains one row with the following values: 'Name' is '[20:00:00:25:B5:06:0D:00-20:00:00:25:B5:06:0D:2F]', 'From' is '20:00:00:25:B5:06:0D:00', and 'To' is '20:00:00:25:B5:06:0D:2F'. At the bottom right are 'OK' and 'Close' buttons.

Name	From	To
[20:00:00:25:B5:06:0D:00-20:00:00:25:B5:06:0D:2F]	20:00:00:25:B5:06:0D:00	20:00:00:25:B5:06:0D:2F

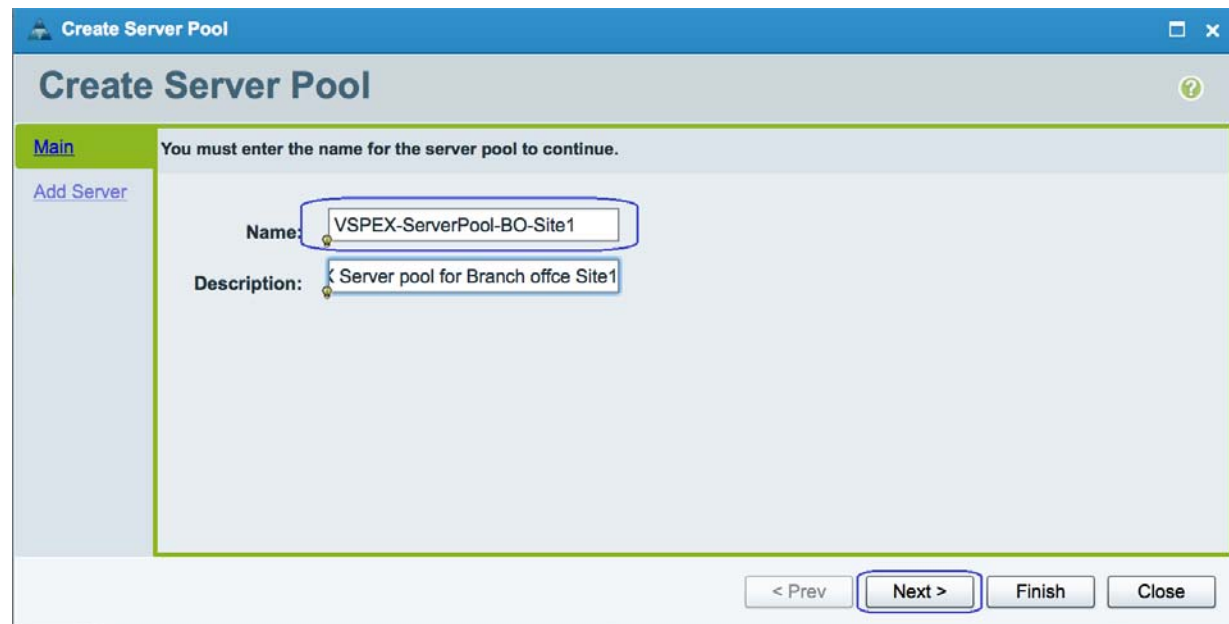
16. To create Server Pool, click **Servers > Pools > Server Pools** and click **Create Server Pool**.

Figure 191 Creating Server Pool



17. In the **General** tab, specify the Server pool name and description and click **Finish** to create a new server pool. We will add the compute resources in this pool dynamically based on the Server Pool Policy.

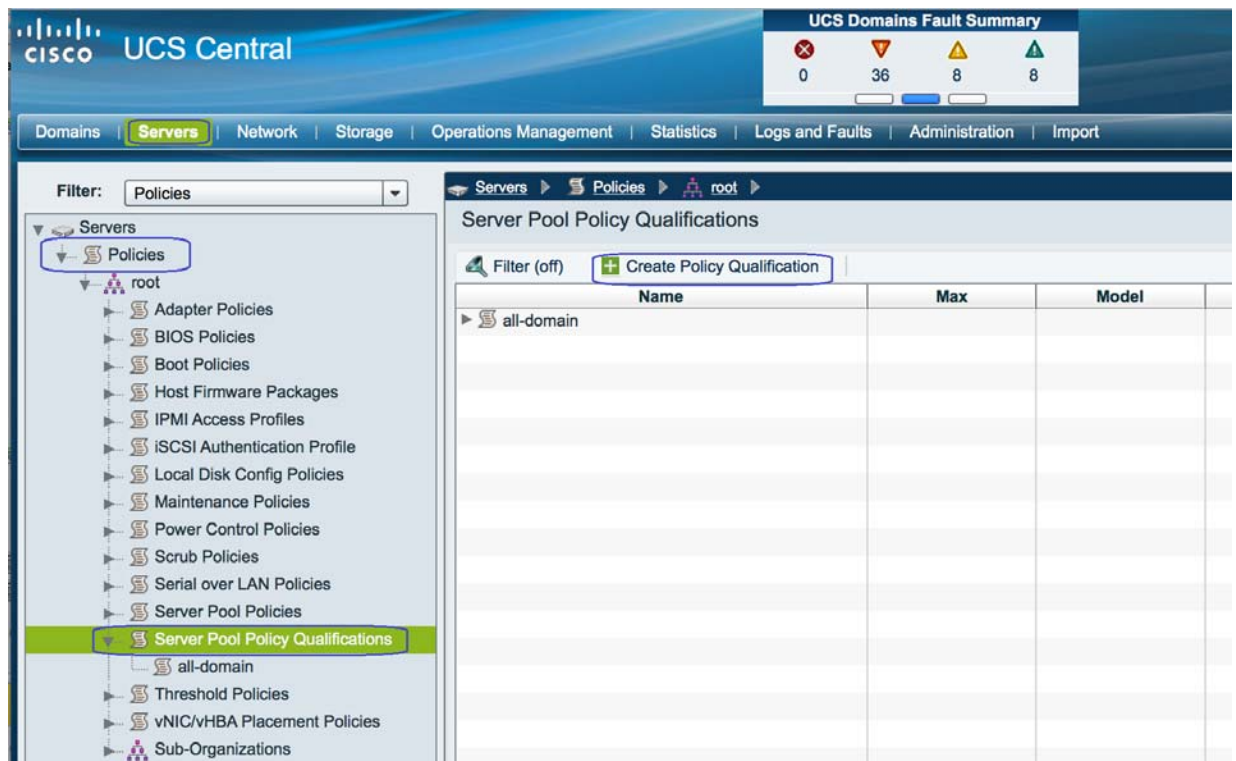
Figure 192 Creating Server Pool - Main



18. To create Server Pool Policy Qualification. Click **Servers > Policies > Server Pool Policy Qualifications** and click **Create Policy Qualification**.

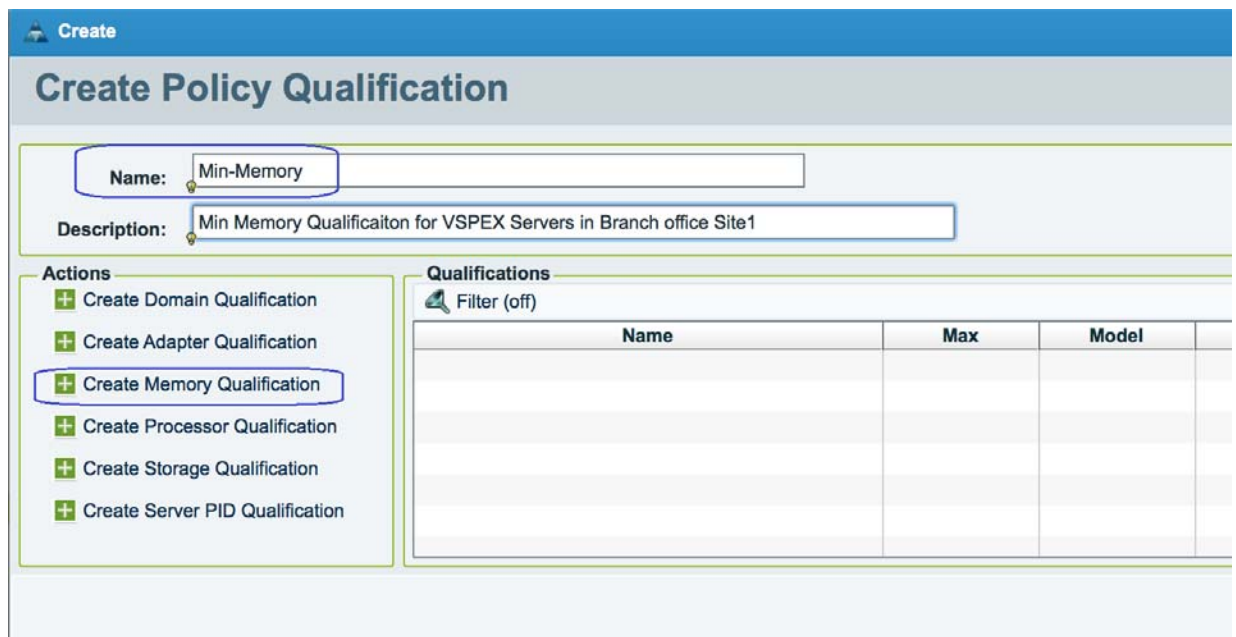


Figure 193 Creating Server Pool Policy Qualification



19. Specify Policy name and description and click **Create Memory Qualification**.

Figure 194 Creating Policy Qualification Window



20. Uncheck the check box **Min Cap(MB)** and specify the minimum memory value in MB for 128GB. Click **OK**.



Figure 195 Creating Memory Qualification Window

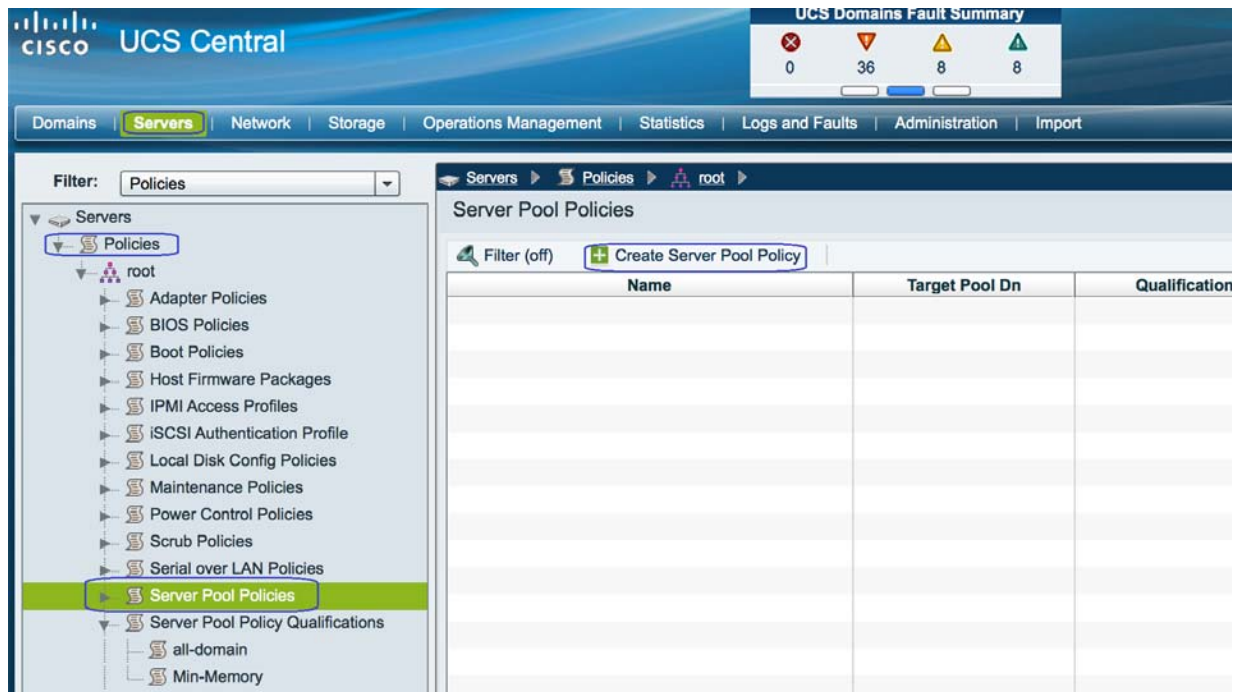
21. After the successful creation, you can see the created Memory Qualification.

Figure 196 Verify the Created Memory Qualification

Name	Max	Model	From	To	Architecture	Speed	Stepping
Memory Qualification						unspecified	

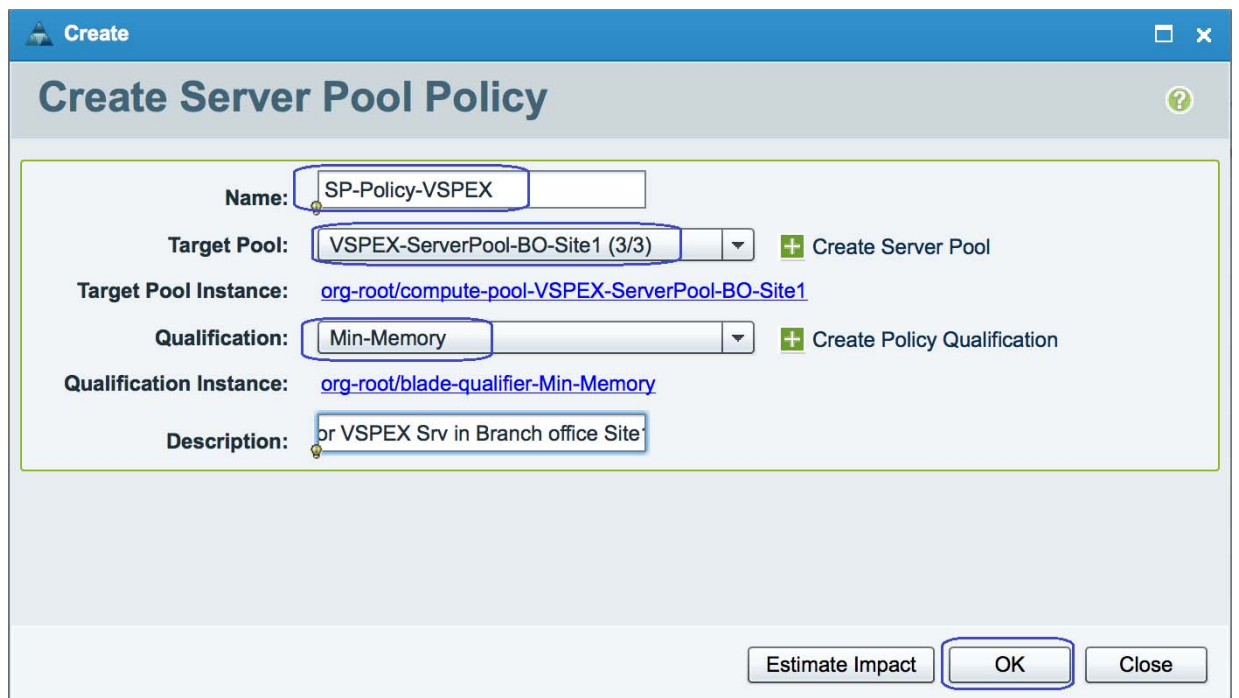
22. To create Server Pool policy, click **Servers > Policies > Server Pool Policies** and click **Create Server Pool Policy**.

*Figure 197*



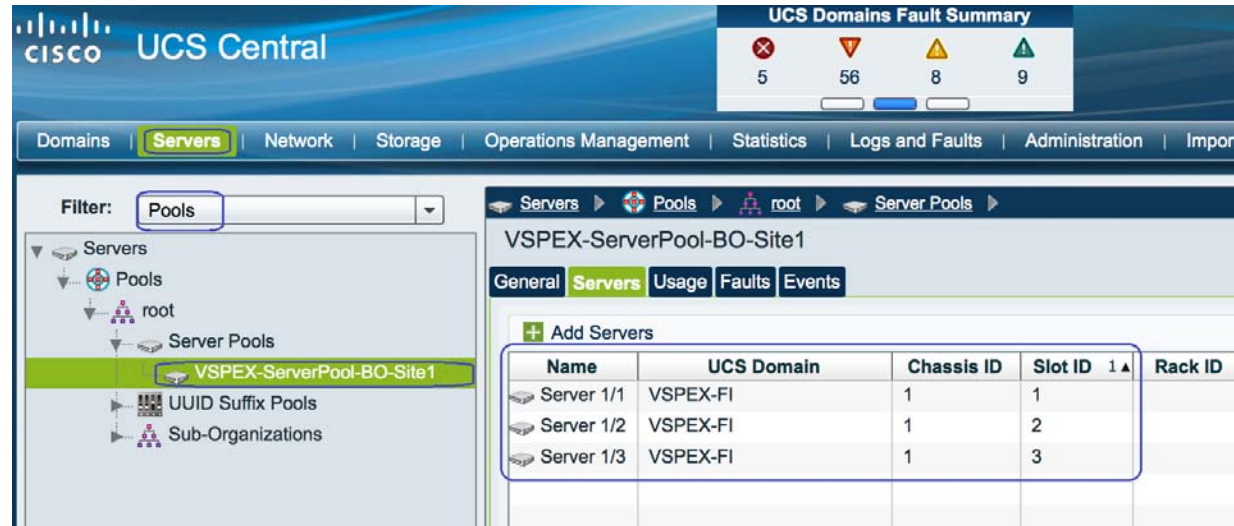
23. Specify the Server pool policy name and description. Then choose the respective Target Pool and Qualification that we have created earlier and click **OK**.

**Figure 198**



24. After the successful Server Pool Policy creation, you will see the created VSPEX Servers are added dynamically to the Branch Office Site.

**Figure 199** Verify the Dynamically Added Servers to the Branch Office Site

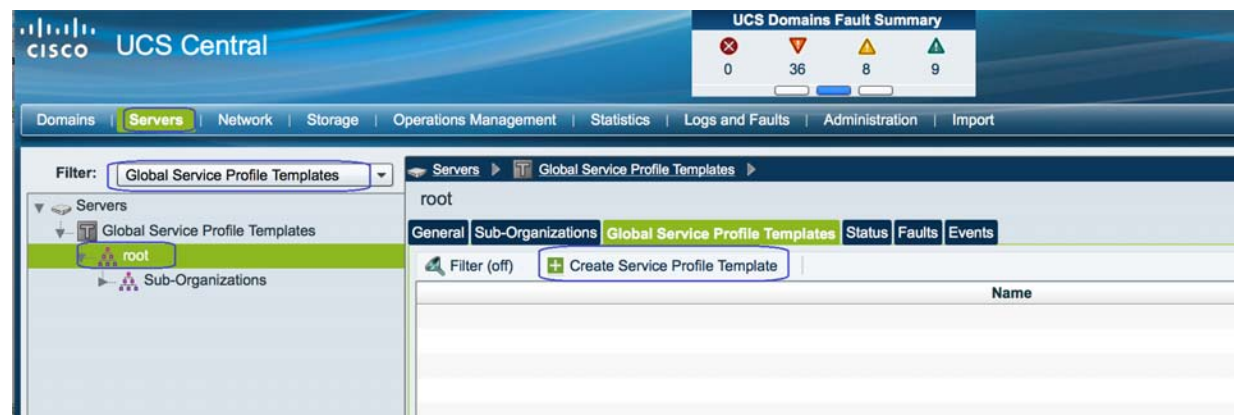


## Configuring Global Service Profile Template

Global service profile templates enable you to quickly create several service profiles with the same basic parameters, such as the number of vNICs and vHBAs, and with identity information drawn from the same pools. Service profile templates in Cisco UCS Central are similar to the service profile templates in Cisco UCS Manager.

1. To create Global Server Profile Templates for Branch office VSPEX servers, launch the UCS Central Web GUI. Click **Servers** > **Global Service Profile Templates** > **root** and then click **Create Service Profile Template**.

**Figure 200** Creating Service Profile Template



2. Specify the Service Profile Template name and description and click the **Updating Template** radio button for Type. Then, choose the UUID Pool that we have earlier and Click **Next**.

**Figure 201**      *Creating Service Profile Template - General*

**Create Service Profile Template**

You must enter the 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:** VSPEX-SP-Template-BO-Site1

**Description:** Service Profile Template for VSPEX Servers in Branch office Site1

**Type:** ☐ initial-template ☒ updating-template

**UUID Assignment**

Create UUID Suffix Pool

Indicate the method to use to assign UUIDs to the server.

☐ Hardware Default ☒ UUID Pool

global-default (0/0)  
VSPEX-UUID-Pool-BO-Site1 (8/8)

< Prev   **Next >**   Finish   Close

3. Choose the Configuration Type as **vNICs – Expert Mode** from the drop-down list and click **Create vNIC**. Click **Next**.

**Figure 202**      *Creating Service Profile Template - Networking*

**Create Service Profile Template**

Optionally specify dynamic vNIC Connection policy and LAN configuration information.

**Dynamic vNIC Connections**

Create Dynamic vNIC Connection Policy

Indicate the whether dynamic vNICs should be used and if so the policy to be used.

Dynamic vNIC Connection Usage: Do not use

**LAN Connectivity**

Create LAN Connectivity Policy

Indicate the method to use to configure LAN Connectivity.

Configuration Type: vNICs - Expert mode

**vNICs**

Click Create to specify one or more vNICs that the server should use to connect to the LAN.

Filter (off) Create vNIC

Name	MAC Address	Fabric ID

Records: 0

**ISCSI vNICs**

Modify Initiator

< Prev   **Next >**   Finish   Close

4. To create a system vNIC for Fabric A, enter System-A as the vNIC name, choose the created MAC Pool, then choose the Fabric ID as **Fabric A** and specify MTU size as **9000**. Then, Select **vMotion** and **vSphereMgmt** vLANs with **vSphereMgmt** as native VLAN. Then choose, **VMware** for the adapter policy.

Figure 203 Networking - Properties for System-A

**Properties (System-A)**

**Properties**

Name:

**MAC Address Assignment:**

☐ Hardware Default

☒ Using Pool  [+ Create MAC Pool](#) [Reset Pool](#)

Instance:

Address:

Use vNIC Template: ☐

**Details**

Fabric ID: ☒ A ☐ B

Fallover: ☒ Enable

MTU:

**Permitted VLANs**

[Filter \(off\)](#) [Refresh](#) Records: 4

Name
Storage
VM-Data
<b>vMotion</b>
vSphereMgmt

[Select >](#)

**Selected VLANs**

[Refresh](#) Records: 0

Name	Set Native
vMotion	<input type="radio"/>
<b>vSphereMgmt</b>	<input checked="" type="radio"/>

Pin Group Name:

[OK](#) [Close](#)

5. Similarly, create one more vNIC for fabric B as System-B with exact same properties on Fabric B.

Figure 204 Networking - Properties for System-B

**Properties (System-B)**

**Properties**

Name:

**MAC Address Assignment:**

☐ Hardware Default

☒ Using Pool  [+ Create MAC Pool](#) [Reset Pool](#)

Instance:

Address:

Use vNIC Template: ☐

**Details**

Fabric ID: ☐ A ☒ B

Failover: ☒ Enable

MTU:

**Permitted VLANs**

[Filter \(off\)](#) [Refresh](#) Records: 4

Name
Storage
VM-Data
<b>vMotion</b>
vSphereMgmt

[Select >](#)

**Selected VLANs**

[Refresh](#) Records: 0

Name	Set Native
vMotion	<input type="radio"/>
vSphereMgmt	<input checked="" type="radio"/>

Pin Group Name:

[OK](#) [Close](#)

6. Repeat the above steps 4 to 6 to create vNICs for VM-Data and Storage o Fabric A and Fabric B. Following table summarizes all the VNICs created on the service profile:

Table 9 Summary of all the vNICs Created

vNIC Name	MAC Address Assignment	VLANs	Native VLANs	Fabric	MTU	Adapter Policy	QoS Policy
System-A	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	A	9000	VMware	JumboMTU
System-B	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	B	9000	VMware	JumboMTU
Storage-A*	MAC pool	Storage	Storage	A	9000	VMware	JumboMTU
Storage-B*	MAC pool	Storage	Storage	B	9000	VMware	JumboMTU



**Table 9** Summary of all the vNICs Created

vNIC Name	MAC Address Assignment	VLANs	Native VLANs	Fabric	MTU	Adapter Policy	QoS Policy
Data-A	MAC pool	VM-Data	VM-Data	A	1500	VMware	-
Data-B	MAC pool	VM-Data	VM-Data	B	1500	VMware	-

7. Once after vNIC creation, click **Next**.

**Figure 205** Creating Service Profile Template - Networking

Optional specify dynamic vNIC Connection policy and LAN configuration information.

**Dynamic vNIC Connections**

Create Dynamic vNIC Connection Policy

Indicate the whether dynamic vNICs should be used and if so the policy to be used.

Dynamic vNIC Connection Usage:

**LAN Connectivity**

Create LAN Connectivity Policy

Indicate the method to use to configure LAN Connectivity.

Configuration Type:

**vNICs**

Click Create to specify one or more vNICs that the server should use to connect to the LAN.

Filter (off) Create vNIC Properties Delete Records: 6

Name	MAC Address	Fabric ID
Storage-A	derived	A-B
Storage-B	derived	B-A
System-A	derived	A-B
System-B	derived	B-A
VMData-A	derived	A-B
VMData-B	derived	B-A

**ISCSI vNICs**

Modify Initiator

< Prev Next > Finish Close

8. To create vHBAs, choose Configuration Type as **vHBAs – Simple Mode** and choose the Global Pool that we created earlier and specify the vHBA name for Fabric A and Fabric B. And then, specify the VSAN name for storage.

Figure 206 Creating Service Profile Template - Storage

Create Service Profile Template

General

Networking

Storage

vNIC/vHBA Placement

Boot Order

Maintenance Policy

Server Assignment

Policies

Optionally specify the SAN configuration information.

Note: If no selection is made, the default Storage configuration policy will be assigned.

**SAN Connectivity**

Create SAN Connectivity Policy

Indicate the method to use to configure SAN Connectivity.

Configuration Type: vHBAs - Simple Mode

**World Wide Node Name (WWNN)**

Indicate the method to use to configure WWNN.

Derived

Global Pool VSPEX-WWxN-Pool-BO-Site1

Create WWN Pool

Reset Pool

**vHBAs**

**vHBA 0 (Fabric A)**

Name: vHBA-A

VSAN: Storage

**vHBA 1 (Fabric B)**

Name: vHBA-B

VSAN: Storage

< Prev Next > Finish Close

9. In the next window, keep the Assignment Method as **Default** and Click **Next**.

Figure 207 Creating Service Profile Template - vNIC/vHBA Placement

Create Service Profile Template

General

Networking

Storage

vNIC/vHBA Placement

Boot Order

Maintenance Policy

Server Assignment

Policies

Optionally specify how vNICs and vHBAs are placed on physical network adapters.

**Placement Method**

Create vNIC/vHBA Placement Policy

Indicate the method to use to assign vNICs and vHBAs to physical network adapters.

Assignment Method: Let System Perform Placement

**PCI Order**

System will assign vNICs and vHBAs based on their PCI order. To change assignment change the order of the vNICs and vHBAs in the table below.

Filter (off)

Name	Address	PCI Order
vNIC System-B	derived	unspecified
vNIC System-A	derived	unspecified
vNIC Storage-A	derived	unspecified
vNIC VMData-B	derived	unspecified
vNIC VMData-A	derived	unspecified
vNIC Storage-B	derived	unspecified
vHBA vHBA-A	derived	unspecified

Records: 8 Showing: 8

< Prev Next > Finish Close

10. Choose the Boot policy as **SAN-boot** as created earlier and click **Next**.



Figure 208 Creating Service Profile Template - Boot Order

Optional specify the Boot Policy for this Service Profile.

**Boot Order Policy**  
 Create Boot Policy

Identify the boot order policy to be applied to the server.

Configuration Type:

Boot Policy:

Boot Policy Instance: [org-root/boot-policy-SAN-boot](#)

**Properties**

Description: ESXi SAN boot policy for VSPEX Servers in Branch office Site1

Reboot on order change: ☒

Enforce device names: ☒

Boot Mode: legacy

**Boot Order**

Name	Order	vNIC/vHBA/SCSI vNIC	Type	LUN ID	WWN
Local CD/DVD	1				
SAN	2				
SAN primary		vHBA-A	primary		
SAN Target primary			primary	0	50:06:01:64:08:E0:03:68
SAN secondary		vHBA-B	secondary		
SAN Target primary			primary	0	50:06:01:65:08:E0:03:68

Records: 2

< Prev Next > Finish Close

11. In the next window, keep the Maintenance Policy as **Default** and click **Next**.

Figure 209 Creating Service Profile Template - Maintenance Policy

Optional specify how disruptive changes (such as reboot, network interruptions, firmware upgrades) should be applied to the system.

**Maintenance Policy**  
 Create Maintenance Policy

Identify the maintenance policy to be applied to the server.

Maintenance Policy:

Policy Instance: [org-root/maint-global-default](#)

Description:

Reboot Policy: ☐ immediate ☒ user-ack ☐ timer-automatic

< Prev Next > Finish Close

12. Choose the created Server Pool and Qualification. Click **Next**.

**Figure 210** *Creating Service Profile Template - Server Assignment*

The screenshot shows the 'Create Service Profile Template' wizard in the 'Server Assignment' step. The left sidebar lists various configuration categories, with 'Server Assignment' highlighted. The main panel contains the following fields:

- Server Assignment Method:** A dropdown menu set to 'Select server from pool'.
- Power state to apply on assignment:** Radio buttons for 'down' and 'up', with 'up' selected.
- Server Pool:** A section with two links: '+ Create Server Pool' and '+ Create Policy Qualification'.
- Identify the Server Pool that the server will be assigned from:**
  - Server Pool:** A dropdown menu set to 'VSPEX-ServerPool-BO-Site1'.
  - Server Pool Instance:** A text field containing 'org-root/compute-pool-VSPEX-ServerPool-BO-Site1'.
  - Qualification:** A dropdown menu set to 'Min-Memory'.
  - Qualification Instance:** A text field containing 'org-root/blade-qualifier-Min-Memory'.
  - Restrict migration of server:** An unchecked checkbox.

At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Close'.

- Keep all the default values and click **Finish**.

**Figure 211** *Creating Service Profile Template - Policies*

The screenshot shows the 'Create Service Profile Template' wizard in the 'Policies' step. The left sidebar lists various configuration categories, with 'Policies' highlighted. The main panel contains the following fields:

- Host Firmware Management:**
  - Host Firmware Package:** A dropdown menu set to 'global-default'.
  - Host Firmware Package Instance:** A text field containing 'org-root/fw-host-pack-global-default'.
- BIOS Configuration:**
  - BIOS Policy:** A dropdown menu.
- External IPMI Management:**
  - IPMI Access Policy:** A dropdown menu.
  - Serial over LAN Policy:** A dropdown menu.
- Management IP Address:**
  - Outband IPv4 Inband:** A section with two tabs: 'Outband IPv4' (selected) and 'Inband'.
  - Management IP Address Policy:** A dropdown menu set to 'none'.

At the bottom right, there are four buttons: '< Prev', 'Next >', 'Finish', and 'Close'.

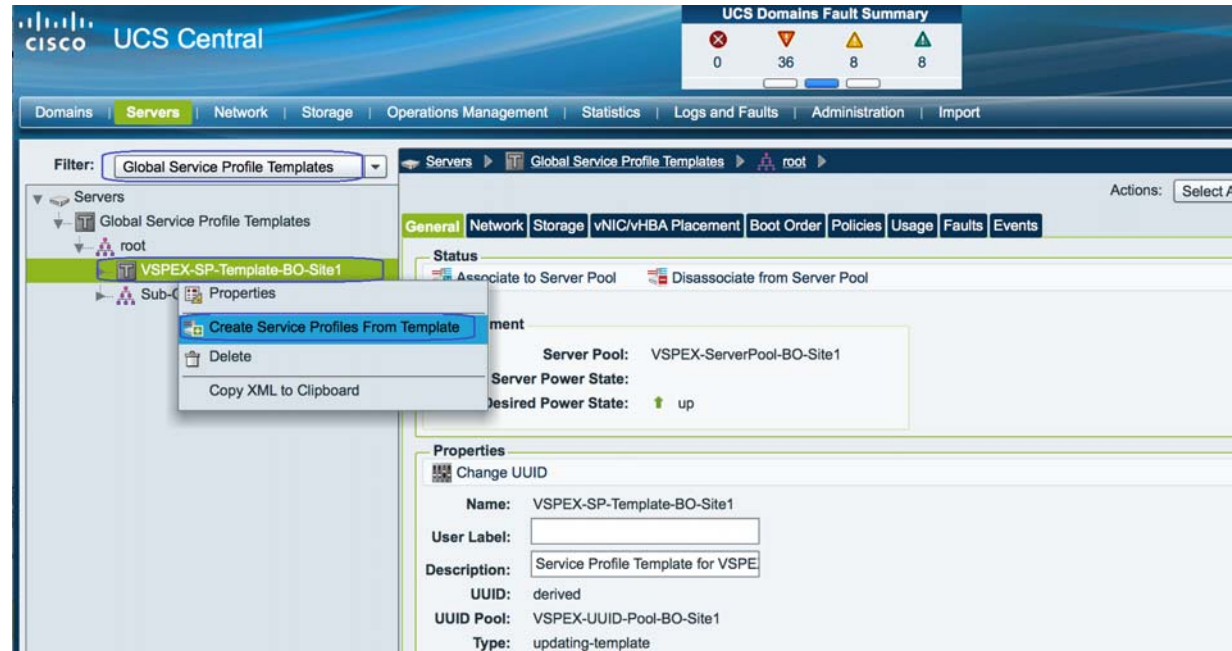
- After the successful creation of Global Service profile Template, you can proceed with the Service Profile Creation.

## Configuring Global Service Profile Instance

Global service profiles centralize the logical configuration deployed across the data center. This centralization enables the maintenance of all service profiles in the registered Cisco UCS domains from one central location, Cisco UCS Central.

1. To create Global Service Profile for Branch office VSPEX servers, launch the UCS Central Web GUI. Click **Servers > Global Service Profile Templates > root** and click **Create Service Profiles from Template**.

Figure 212 Creating Service Profile from Template



2. Specify the name for the service profile, number of service profile instances to be instantiated and choose the Org as **root**. Then, refer to the sizing guidelines for the number of servers needed for your deployment.

Figure 213 Creating Service Profile from Template - Properties

**Properties**

## Create Service Profiles From Template

Name Prefix:

Number:

Org:

Org Instance: [org-root](#)

3. You will see that three service profiles are created.

Figure 214 Verify the Service Profiles Created from the Template

UCS Central

UCS Domains Fault Summary: 0 (red), 36 (yellow), 8 (green)

Domains | **Servers** | Network | Storage | Operations Management | Statistics | Logs and Faults | Administration | Import

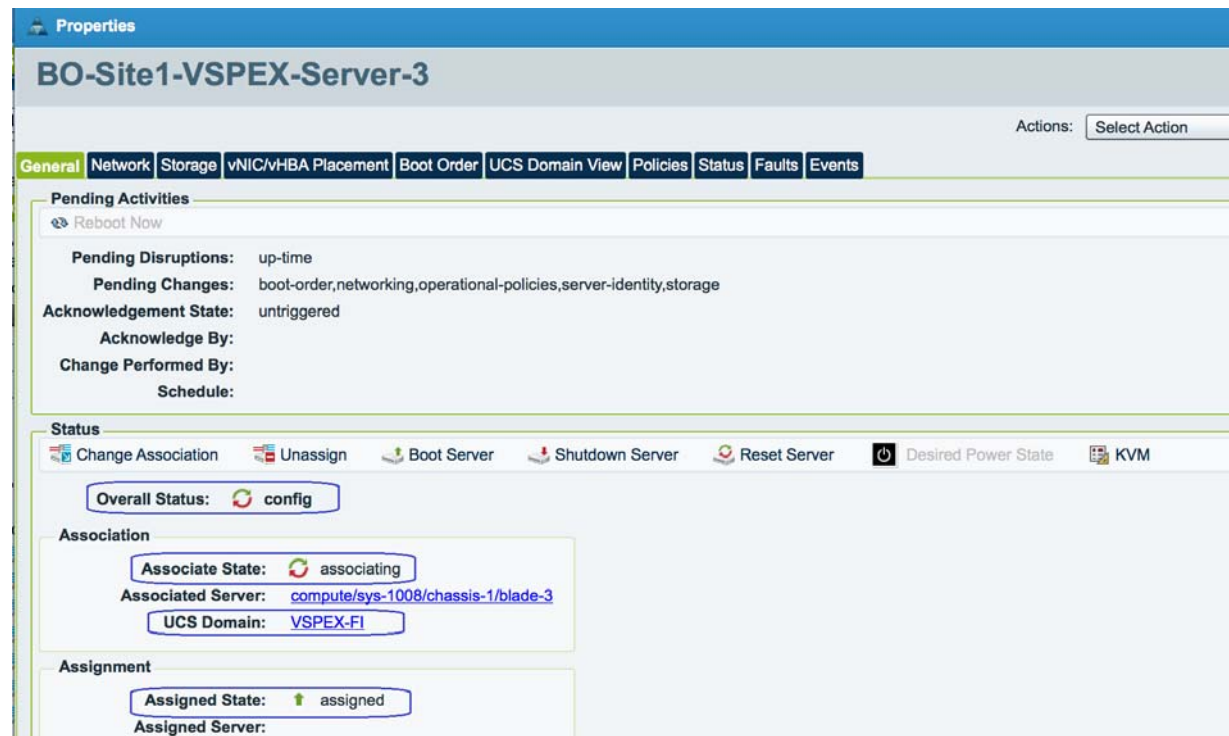
Filter: All

Servers > Global Service Profiles > root

Name	Status	Associated Server	Domain
BO-Site1-VSPEX-Server-1	unassociated	compute/sys-1008/chassis-1/blade-1	VSPEX-FI
BO-Site1-VSPEX-Server-2	unassociated	compute/sys-1008/chassis-1/blade-2	VSPEX-FI
BO-Site1-VSPEX-Server-3	unassociated	compute/sys-1008/chassis-1/blade-3	VSPEX-FI

4. As service profile template is assigned to a server pool, the service profiles instantiated from the template would be assigned to individual server resource from the server pool as far as they are available. You can select a given service profile and see its overall status and the association state.

Figure 215 Viewing Overall Status of the Server



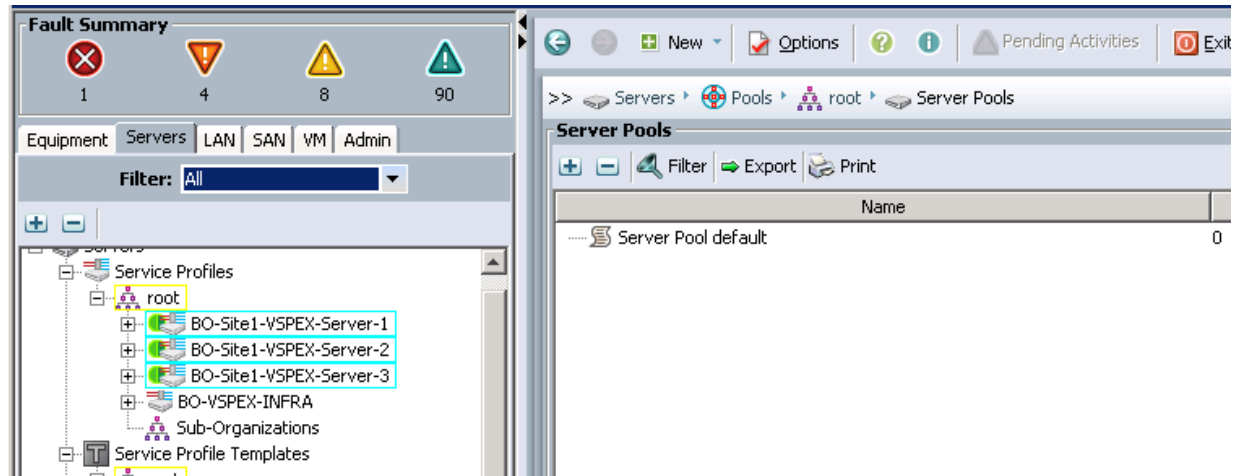
- Eventually, all the three VSPEX Servers would be associated. Click the **Servers > Global Service Profiles > root** to see the summary of all the servers.

Figure 216 Summary of the Created Global Service profiles



- By launching the UCS Manager GUI, You can also see the Global Service Profiles created on the UCS Manager GUI for the Branch office VSPEX pod.

**Figure 217** Verifying Global Service Profiles Through UCS Manager



**Note**

The above procedure presents a method for creating service profile for Branch office VSPEX servers using UCS Central management located in VSPEX Primary DC Site.

In the event of deploying this solution with UCS Central Management, follow the steps 4 to 7 shown in the [“VSPEX Configuration Guidelines”](#) section on page 30:

- Configure data stores for ESXi images
- Install ESXi servers and vCenter infrastructure
- Install and configure VMware vCenter server
- Configure storage for VM data stores, install and instantiate VSPEX VMs from vCenter


## WAN Testing for UCS Central Management

The key benefit for using UCS Central Management is to provide centralized management for multiple UCS domains whether local or remote. Given UCS Central may reside in a data center over distance, the test includes management of UCS Mini systems over the equivalent of an entry-level consumer grade DSL line – 1.5 Mbps, 500ms latency with better resiliency for temporary loss of connection between UCS Central and UCS Manager instances. While these enhancements are required for remote and branch offices, they are also useful for customers using UCS within a data center. We have used WAN emulator testing tool to simulate bandwidth, latency and packet loss measures defined above. During testing, we observed that UCS central management in primary data center to be stable and responsive during configuration changes applied to the UCS Mini system in Branch Office site.



## Template-Based Deployments for Rapid Provisioning

*Figure 218 Rapid Provisioning*

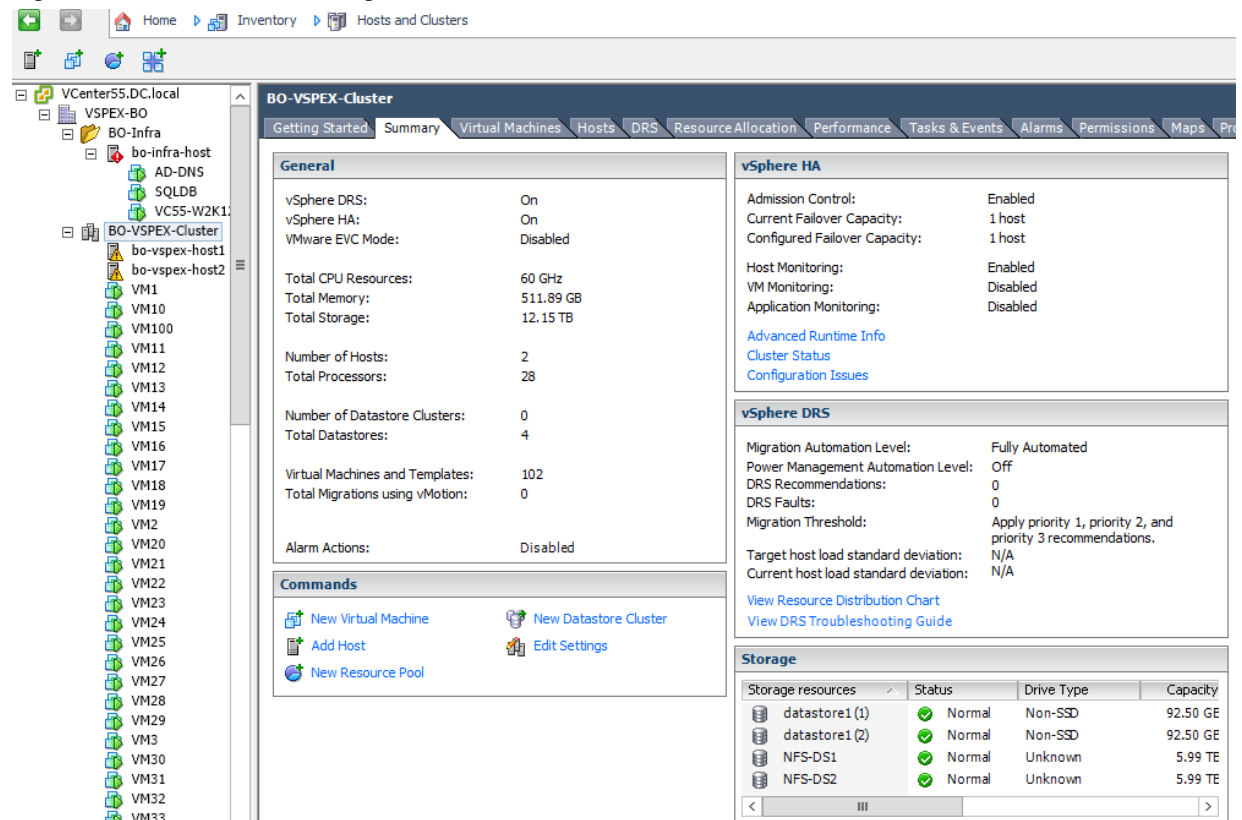


In an environment with established procedures, deploying new application servers can be streamlined, but can still take many hours or days to complete. Not only must you complete an OS installation, but downloading and installing service packs and security updates can add a significant amount of time. Many applications require features that are not installed with Windows by default and must be installed prior to installing the applications. Inevitably, those features require more security updates and patches. By the time all deployment aspects are considered, more time is spent waiting for downloads and installs than is spent configuring the application.

Virtual machine templates can help speed up this process by eliminating most of these monotonous tasks. By completing the core installation requirements, typically to the point where the application is ready to be installed, you can create a golden image which can be sealed and used as a template for all of your virtual machines. Depending on how granular you want to make a specific template, the time to deployment can be as little as the time it takes to install, configure, and validate the application. You can use PowerShell tools and VMware vSphere Power CLI to bring the time and manual effort down dramatically.

Make sure to spread VMs across different VM datastores to properly load-balance the storage usage. The final snapshot of VMs in a cluster would look similar to [Figure 219](#).

**Figure 219** Window Showing all VMs in the Cluster



## Validating Cisco Branch Office Solution for EMC VSPEX VMware Architectures

This section provides a list of items that should be reviewed once the solution has been configured. The goal of this section is to verify the configuration and functionality of specific aspects of the solution, and ensure that the configuration supports core availability requirements.

### Post Install Checklist

The following configuration items are critical to functionality of the solution, and should be verified prior to deployment into production.

- Create a test virtual machine that accesses the datastore and is able to do read/write operations. Perform the virtual machine migration (vMotion) to a different host on the cluster.
- Perform storage vMotion from one datastore to another datastore and ensure correctness of data.
- During the vMotion of the virtual machine, have a continuous ping to default gateway and make sure that network connectivity is maintained during and after the migration.



## Verify the Redundancy of the Solution Components

Following redundancy checks were performed at the Cisco lab to verify solution robustness. A continuous ping from VM to VM, and vCenter to ESXi hosts should not show significant failures (one or two ping drops might be observed at times, such as FI reboot). Also, all the data-stores must be visible and accessible from all the hosts at all the time.

1. Administratively shutdown one of the Network uplink port from Fabric Interconnect A connected to upstream LAN (Lab Network). Make sure that connectivity is not affected. Upon administratively enabling the shutdown port, the traffic should be rebalanced. Repeat this step for the Network uplink port from Fabric Interconnect B and make sure the connectivity is not affected.
2. Administratively shutdown FC port connected to Fabric Interconnect A. ESXi hosts should be able to use fabric interconnect B in this case.
3. Administratively shutdown one of the two data links connected to the storage array from FI. Make sure that storage is still available from all the ESXi hosts. Upon administratively enabling the shutdown port, the traffic should be rebalanced. Repeat this step for each link connected to the Storage Processors one after another.
4. Reboot one of the two Fabric Interconnects while storage and network access from the servers are going on. The switch reboot should not affect the operations of storage and network access from the VMs. Upon rebooting the FI, the network access load should be rebalanced across the two fabrics.
5. Fully load all the virtual machines of the solution. Put one of the ESXi host in maintenance mode. All the VMs running on that host should be migrated to other active hosts. No VM should lose any network or storage accessibility during or after the migration. This test assumes that enough RAM is available on active ESXi hosts to accommodate VMs from the host put in maintenance mode.
6. Reboot the host in maintenance mode, and put it out of the maintenance mode. This should rebalance the VM distribution across the cluster.

## Cisco Validation Test Profile

“Vdbench” testing tool (ver.5.04) was used with Windows 2012 server to test scaling of the solution in Cisco labs. Following is the detail on the test profile used.

**Table 10** *Vdbench Test Profile*

Profile Characteristics	Value
Number of virtual machines	100
Virtual machine OS	Windows Server 2012
Processors per virtual machine	1
Number of virtual processors per physical CPU core	2
RAM per virtual machine	2 GB
Average storage available for each virtual machine	100 GB
Average IOPS per virtual machine	25 IOPS

# Bill of Material

The table 6. gives details of the components used in the CVD for 100 virtual machines configuration.

**Table 11** *Component Description*

Description	Part Number
UCS Chassis 5108	UCS-5108-AC2
6324UP Fabric Interconnects	UCS-FI-M-6324
UCS B200 M3 Blade	UCSB-B200-M3
10 Gbps SFP+ multifiber mode	SFP-10G-SR
8 Gbps SFP+ fibre mode	DS-SFP-FC8G-SW
1000Base-T copper module	CIS-GLC-T-NP-OE

For more information about the part numbers and options available for customization, see Cisco UCS 6324 Fabric Interconnect datasheet at:

<http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-732207.html>

# Customer Configuration Data Sheet

Before you start the configuration, gather some customer-specific network and host configuration information. [Table 12](#), [Table 13](#), [Table 14](#), [Table 15](#), [Table 16](#), [Table 17](#), and [Table 18](#) provide information on assembling the required network and host address, numbering, and naming information. This worksheet can also be used as a “leave behind” document for future reference.

The VNXe Series Configuration Worksheet should be cross-referenced to confirm customer information.

**Table 12** *Common Server Information*

Server Name	Purpose	Primary IP
	Domain Controller	
	DNS Primary	
	DNS Secondary	
	DHCP	
	NTP	
	SMTP	
	SNMP	
	vCenter Console	
	SQL Server	

**Table 13** *ESXi Server Information*

Server Name	Purpose	Management IP	Private Net (storage) addresses		vMotion IP
	ESXi Host 1				
	...				
	ESXi Host 5				

**Table 14** *Storage Array Information*

Array name	
Admin account	
Management IP	
Storage pool name	
Datastore name	
NFS server IP	

**Table 15** *Network Infrastructure Information*

Purpose	IP	Subnet Mask	Default Gateway
Cisco UCS virtual IP address			
Cisco UCS FI A address			
Cisco UCS FI B address			

**Table 16** *VLAN Information*

Name	Network Purpose	VLAN ID	Allowed Subnets
vSphereMgmt	Virtual Machine Networking VMware ESXi Management		
Storage	NFS VLAN		
vMotion	VMware vMotion traffic network		
VM-Data (multiple)	Data VLAN of customer VMs as needed		

**Table 17** *VSAN information*

Name	Network Purpose	VSAN ID	FCoE VLAN ID
Storage	Storage Access		

**Table 18** *Service Accounts*

Account	Purpose	Password (option, secure)
Admin	UCS Manager administrator	
Admin	UCS Central	
	Windows server administrator	
Root	ESXi root	
	EMC VNXe array administrator	
	vCenter administrator	
	SQL server administrator	

## Conclusion

The Cisco branch office solution uses common components of EMC VSPEX Integrated Infrastructure with new introductions such as the Cisco UCS Mini and EMC VNXe 3200 for a version of traditional VSPEX solutions that is ideal for small and medium businesses. Additional tools such as UCS Central provide the means to address business and infrastructure requirements from a central location. These functional requirements promote uniqueness and innovation in the integrated computing stack, augmenting their original design with support for essential services such as standards based centralized management of remote instances.

## References

Cisco UCS:

[http://www.cisco.com/en/US/solutions/ns340/ns517/ns224/ns944/unified\\_computing.html](http://www.cisco.com/en/US/solutions/ns340/ns517/ns224/ns944/unified_computing.html)

Cisco UCS Mini Firmware Management:

[http://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/sw/firmware-mgmt/gui/3-0/b\\_GUI\\_Firmware\\_Management\\_30/b\\_GUI\\_Firmware\\_Management\\_30\\_chapter\\_01.html](http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/firmware-mgmt/gui/3-0/b_GUI_Firmware_Management_30/b_GUI_Firmware_Management_30_chapter_01.html)

UCS Central Software and Installation Guide:

<http://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-central-software/tsd-products-support-series-home.html>

VMware vSphere:

<http://www.vmware.com/products/vsphere/overview.html>

VMware vSphere 5.5 Documentation:

<http://pubs.vmware.com/vsphere-55/index.jsp>

EMC VNXe32xx Series Resources:

<http://www.emc.com/storage/vnx/vnxe-series.htm#!resources>

EMC Support: (requires user registration)

<http://support.emc.com>

Microsoft SQL Server Installation Guide:

<http://msdn.microsoft.com/en-us/library/ms143219.aspx>