Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 100-125 Virtual Machines

Last Updated: June 21, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Mehul Bhatt

Mehul Bhatt, Virtualization Architect, Server Access Virtualization Business Unit, Cisco Systems

Mehul Bhatt has over 12 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 Masters degree in computer systems engineering and holds various Cisco career certifications.

Acknowledgements

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

- Vadiraja Bhatt-Cisco
- Rajendra Yogendra-Cisco
- Bathu Krishnan-Cisco
- Sindhu Sudhir-Cisco
- Kevin Phillips-EMC
- John Moran-EMC
- Kathy Sharp-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.

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.

Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 100-125 Virtual Machines

Executive Summary

Cisco solution for EMC VSPEX is a pre-validated and modular architecture built with proven best of-breed technologies to create complete end-to-end virtualization solutions that enable you to make an informed decision while choosing the hypervisor, compute, storage and networking layers. VSPEX drastically reduces server virtualization planning and configuration burdens. VSPEX infrastructures accelerate your IT Transformation by enabling faster deployments, greater flexibility of choice, efficiency, and lower risk. This Cisco Validated Design document focuses on the VSPEX VMware architecture for small to medium size business segments with less than 125 typical Virtual Machines load.

Introduction

Virtualization is a key and critical strategic deployment model for reducing the Total Cost of Ownership (TCO) and achieving better utilization of the platform components like hardware, software, network and storage. However choosing the appropriate platform for virtualization can be a tricky task. The platform should be flexible, reliable and cost effective to facilitate the virtualization platform to deploy various enterprise applications. Also ability to slice and dice the underlying platform to size the application requirement is essential for a virtualization platform to utilize compute, network and storage resources effectively. In this regard, Cisco solution implementing EMC VPSEX provide a very simplistic yet fully integrated and validated infrastructure for you to deploy VMs in various sizes to suite your application needs.

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

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 Unified Computing System (UCS) and Unified Computing Systems Manager (UCSM). 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 describes the steps required to deploy and configure a Cisco solution for EMC VSPEX for VMware architectures to a level that will allow for confirmation that the basic components and connections are working correctly. The document covers VMware architectures for Small- to Medium-sized Businesses, typically having 100 to 125 VMs or less. This document show cases two variants of the solution: one involving EMC VNX series storage array using FC for storage access, and one involving EMC VNXe series storage array using iSCSI for storage access. While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to this solution's deployment s are specifically mentioned.

Business Needs

VSPEX solutions are built with proven best-of-breed technologies to create complete virtualization solutions that enable you to make an informed decision in the hypervisor, server, and networking layers. VSPEX infrastructures accelerate your IT transformation by enabling faster deployments, greater flexibility of choice, efficiency, and lower risk.

Business applications are moving into the consolidated compute, network, and storage environment. Cisco solution for EMC VSPEX for VMware helps to reduce complexity of configuring every component of a traditional deployment. The complexity of integration management is reduced while maintaining the application design and implementation options. Administration is unified, while process separation can be adequately controlled and monitored. The following are the business needs for the Cisco solution of EMC VSPEX VMware architectures:

- Provide an end-to-end virtualization solution to take full advantage of unified infrastructure components.
- Provide a Cisco VSPEX for VMware ITaaS solution for efficiently virtualizing virtual machines for varied customer use cases.
- Show implementation progression of VMware vCenter 5.1 design and results.
- Provide a reliable, flexible and scalable reference design

Solution Overview

Cisco 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 Compute System (UCS) 2.1
- Cisco B-series or C-series Unified Computing System servers, as per customer choice
- Cisco UCS VIC adapters
- EMC VNXe3300 or VNX5300 storage components as per the scale needs
- VMware vCenter 5.1
- Microsoft SQL database
- VMware DRS
- VMware HA

The solution is designed to host scalable, mixed application workloads. The scope of this CVD is limited to the Cisco solution for EMC VSPEX VMware solutions for SMB market segment only.

Technology Overview

Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims 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. It is an integrated, scalable, multi chassis platform in which all resources participate in a unified management domain.

The main components of Cisco Unified Computing System are:

- **Computing**—The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon E5-2600/4600 and E7-2800 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, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.
- **Management**—The system uniquely integrates all system components which enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.
- A cohesive, integrated system which unifies the technology in the data center. The system is managed, serviced and tested as a whole.
- Scalability through a design for hundreds of discrete servers and thousands of virtual machines and the capability to scale I/O bandwidth to match demand.
- Industry standards supported by a partner ecosystem of industry leaders.

Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management of all software and hardware components of the Cisco Unified Computing System through 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.

Cisco UCS Fabric Interconnect

The Cisco® UCS 6200 Series Fabric Interconnect is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6200 Series offers line-rate, low-latency, lossless 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE) and Fibre Channel functions.

The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available management domain. In addition, by supporting unified fabric, the Cisco UCS 6200 Series provides both the LAN and SAN connectivity for all blades within its domain.

From a networking perspective, the Cisco UCS 6200 Series uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, 1Tb switching capacity, 160 Gbps bandwidth per chassis, independent of packet size and enabled services. The product family supports Cisco low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The Fabric Interconnect supports multiple traffic classes over a lossless Ethernet fabric from a blade server through an interconnect. Significant TCO savings come from an FCoE-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Cisco UCS 6248UP Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a one-rack-unit (1RU) 10 Gigabit Ethernet, FCoE and Fiber Channel switch offering up to 960-Gbps throughput and up to 48 ports. The switch has 32 1/10-Gbps fixed Ethernet, FCoE and FC ports and one expansion slot.

Figure 1 Cisco UCS 6248UP Fabric Interconnect

Cisco UCS Fabric Extenders

Fabric Extenders are zero-management, low-cost, low-power consuming devices that distribute the system's connectivity and management planes into rack and blade chassis to scale the system without complexity. Designed never to lose a packet, Cisco fabric extenders eliminate the need for top-of-rack Ethernet and Fibre Channel switches and management modules, dramatically reducing infrastructure cost per server.

Cisco UCS 2232PP Fabric Extender

The Cisco Nexus® 2000 Series Fabric Extenders comprise a category of data center products designed to simplify data center access architecture and operations. The Cisco Nexus 2000 Series uses the Cisco® Fabric Extender architecture to provide a highly scalable unified server-access platform across a range of 100 Megabit Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet, unified fabric, copper and fiber connectivity, rack, and blade server environments. The platform is ideal to support today's traditional Gigabit Ethernet while allowing transparent migration to 10 Gigabit Ethernet, virtual machine-aware unified fabric technologies.

The Cisco Nexus 2000 Series Fabric Extenders behave as remote line cards for a parent Cisco Nexus switch or Fabric Interconnect. The fabric extenders are essentially extensions of the parent Cisco UCS Fabric Interconnect switch fabric, with the fabric extenders and the parent Cisco Nexus switch together forming a distributed modular system. This architecture enables physical topologies with the flexibility and benefits of both top-of-rack (ToR) and end-of-row (EoR) deployments.

Today's data centers must have massive scalability to manage the combination of an increasing number of servers and a higher demand for bandwidth from each server. The Cisco Nexus 2000 Series increases the scalability of the access layer to accommodate both sets of demands without increasing management points within the network.

Figure 2 Cisco UCS 2232PP Fabric Extender

Cisco C220 M3 Rack Mount Servers

Building on the success of the Cisco UCS C220 M3 Rack Servers, the enterprise-class Cisco UCS C220 M3 server further extends the capabilities of the Cisco Unified Computing System portfolio in a 1-rack-unit (1RU) form factor. And with the addition of the Intel® Xeon® processor.

Figure 3 Cisco UCS C220 M3 Rack Mount Server

The Cisco UCS C220 M3 also offers up to 256 GB of RAM, eight drives or SSDs, and two 1GE LAN interfaces built into the motherboard, delivering outstanding levels of density and performance in a compact package.

Cisco UCS Blade Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis.

The Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2204XP Fabric Extenders.

A passive mid-plane provides up to 40 Gbps of I/O bandwidth per server slot and up to 80 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards. The Cisco UCS Blade Server Chassis is shown in Figure 4.

Figure 4 Cisco Blade Server Chassis (front and back view)

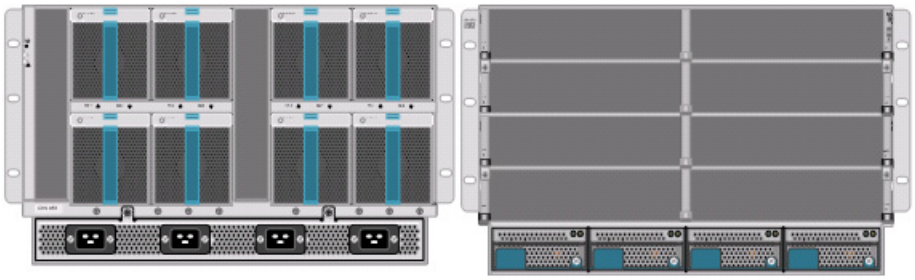

Cisco UCS Blade Servers

Delivering performance, versatility and density without compromise, the Cisco UCS B200 M3 Blade Server addresses the broadest set of workloads, from IT and Web Infrastructure through distributed database.

Building on the success of the Cisco UCS B200 M2 blade servers, the enterprise-class Cisco UCS B200 M3 server, further extends the capabilities of Cisco's Unified Computing System portfolio in a half blade form factor. The Cisco UCS B200 M3 server harnesses the power and efficiency of the Intel Xeon E5-2600 processor product family, up to 768 GB of RAM, 2 drives or SSDs and up to 2 x 20 GbE to deliver exceptional levels of performance, memory expandability and I/O throughput for nearly all applications. In addition, the Cisco UCS B200 M3 blade server offers a modern design that removes the

need for redundant switching components in every chassis in favor of a simplified top of rack design, allowing more space for server resources, providing a density, power and performance advantage over previous generation servers. The Cisco UCS B200M3 Server is shown in Figure 5.

Figure 5 Cisco UCS B200 M3 Blade Server

Cisco I/O Adapters

Cisco UCS Blade Servers support various Converged Network Adapter (CNA) options. Cisco UCS Virtual Interface Card (VIC) 1240 is used in this EMC VSPEX solution.

The Cisco UCS Virtual Interface Card 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 6 Cisco UCS VIC 1240

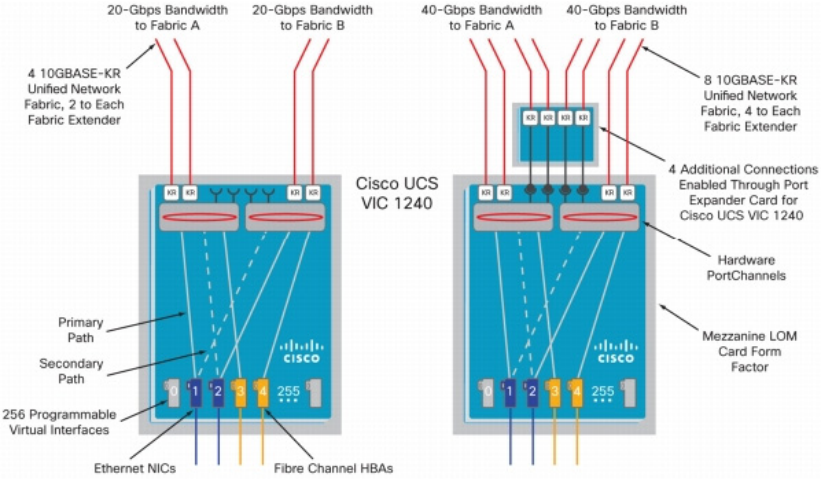

The Cisco UCS rack mount server has various Converged Network Adapters (CNA) options. The UCS 1225 Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

A Cisco® innovation, the Cisco UCS Virtual Interface Card (VIC) 1225 is a dual-port Enhanced Small Form-Factor Pluggable (SFP+) 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) card designed exclusively for Cisco UCS C-Series Rack Servers.

UCS 1225 VIC provides the capability to create multiple vNICs (up to 128) on the CNA. This allows complete I/O configurations to be provisioned in virtualized or non-virtualized environments using just-in-time provisioning, providing tremendous system flexibility and allowing consolidation of multiple physical adapters.

System security and manageability is improved by providing visibility and portability of network policies and security all the way to the virtual machines. Additional 1225 features like VM-FEX technology and pass-through switching, minimize implementation overhead and complexity.

Figure 7 Cisco UCS 1225 VIC

UCS 2.1 Singe Wire Management

Cisco UCS Manager 2.1 supports an additional option to integrate the C-Series Rack-Mount Server with Cisco UCS Manager called “single-wire management”. This option enables Cisco UCS Manager to manage the C-Series Rack-Mount Servers using a single 10 GE link for both management traffic and data traffic. When you use the single-wire management mode, one host facing port on the FEX is sufficient to manage one rack-mount server, instead of the two ports you will use in the Shared-LOM mode. Cisco VIC 1225, Cisco UCS 2232PP FEX and Single-Wire management feature of UCS 2.1 tremendously increases the scale of C-series server manageability. By consuming as little as one port on the UCS Fabric Interconnect, you can manage up to 32 C-series server using single-wire management feature.

UCS Differentiators

Cisco’s 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. Also, a pair of FIs can manage up to 40 chassis, each containing 8 blade servers. This gives enormous scaling on the management plane.
2. **Unified fabric**—In UCS, from blade server chassis or rack server fabric-extender to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.

3. **Auto Discovery**—By simply inserting the blade server in the chassis or connecting rack server to the fabric extender, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of UCS, where compute capability of UCS can be extended 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 showcases 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 combinations of B and C 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. An 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**—The management approach in UCSM is based on defining policies, pools and templates, instead of cluttered configuration, which enables a 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 flexibility 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 another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then special policy named “default” is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
10. **Service profiles and stateless computing**—A 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**—The combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and a service profiles based approach to compute resources makes UCSM inherently friendly to multi-tenant environment typically observed in private and public clouds.
12. **Extended Memory**—The extended memory architecture of UCS servers allows up to 760 GB RAM per server – allowing huge VM to physical server ratio required in many deployments, or allowing large memory operations required by certain architectures like Big-Data.
13. **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

administrators' team. VM-FEX also off loads 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.

14. **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 5.1

VMware vSphere 5.1 is a next-generation virtualization solution from VMware which builds upon ESXi 4 and provides greater levels of scalability, security, and availability to virtualized environments. vSphere 5.0 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.

vSphere 5.1 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 the network ports that have Layer 2 configurations, port access and security policies, monitoring features, and so on, that have been pre-defined by the network administrators; in the same way they need to plug in their physical servers to a previously-configured access switch. In this virtualized environment, the network port configuration/policies move with the virtual machines when the virtual machines are migrated to different server hardware.

EMC Storage Technologies and Benefits

This architecture has two variants:

- EMC VNX family based FC-variant of the solution
- EMC VNXe family based iSCSI-variant of the solution

The EMC VNX™ family is optimized for virtual applications delivering industry-leading innovation and enterprise capabilities for file, block, and object 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.

VNX series is designed to meet the high-performance, high-scalability requirements of midsize and large enterprises. The EMC VNX storage arrays are multi-protocol platform that can support the iSCSI, NFS, Fibre Channel, and CIFS/SMB protocols depending on the customer's specific needs. This solution was validated using NFS for data storage of Virtual Machines and Fibre Channel for hypervisor SAN boot.

VNX series storage arrays have the following customer benefits:

- Next-generation unified storage, optimized for virtualized applications
- Capacity optimization features including compression, deduplication, thin provisioning, and application-centric copies

- High availability, designed to deliver five 9s availability
- Multiprotocol support for file and block
- Simplified management with EMC Unisphere™ for a single management interface for all network-attached storage (NAS), storage area network (SAN), and replication needs

Software Suites

The following are the available EMC software suites:

- **Remote Protection Suite**—Protects data against localized failures, outages, and disasters.
- **Application Protection Suite**—Automates application copies and proves compliance.
- **Security and Compliance Suite**—Keeps data safe from changes, deletions, and malicious activity.

Software Packs

Total Value Pack—Includes all protection software suites, and the Security and Compliance Suite.

This is the available EMC protection software pack.

The EMC VNXe™ series is powered by Intel Xeon processor, for intelligent storage that automatically and efficiently scales in performance, while ensuring data integrity and security.

The EMC VNXe series is purpose-built for the IT manager in smaller environments. The EMC VNXe storage arrays are multi-protocol platforms that can support the iSCSI, NFS, and CIFS protocols depending on the customer's specific needs. The solution was validated using iSCSI for data storage.

EMC Avamar

EMC's Avamar® data deduplication technology seamlessly integrates into virtual environments, providing rapid backup and restoration capabilities. Avamar's deduplication results in vastly less data traversing the network, and greatly reduces the amount of data being backed up and stored; resulting in storage, bandwidth and operational savings.

The following are the two most common recovery requests used in backup and recovery:

- **File-level recovery**—Object-level recoveries account for the vast majority of user support requests. Common actions requiring file-level recovery are—individual users deleting files, applications requiring recoveries, and batch process-related erasures.
- **System recovery**—Although complete system recovery requests are less frequent in number than those for file-level recovery, this bare metal restore capability is vital to the enterprise. Some of the common root causes for full system recovery requests are viral infestation, registry corruption, or unidentifiable unrecoverable issues.

The Avamar System State protection functionality adds backup and recovery capabilities in both of these scenarios.

Architectural Overview

This CVD focuses on the architecture for EMC VSPEX for VMware private cloud, targeted for the SMB market segment, using EMC VNX and VNXe series storage arrays. There are two variants of the architecture: FC-variant and iSCSI-variant. The FC-variant of the architecture uses UCS 2.1 with combined B-series and C-series servers with VNX5300 directly attached to UCS fabric interconnect.

The iSCSI-variant of the architecture uses UCS 2.1 and C220 M3 rack mount servers with VNXe storage array directly attached to UCS fabric interconnects. In both variants, the C220 M3 servers are connected with single-wire management feature. VMware vSphere 5.1 is used as server virtualization architecture and iSCSI as the storage access protocol.

Table 1 lists the various hardware and software components which occupies different tiers of the Cisco solution for EMC VSPEX VMware architectures under test:

Table 1 Hardware and software components of VMware architectures

Vendor	Name	Version	Description
Cisco	UCSM	2.1(1)	Cisco UCS Manager
Cisco	UCS 6248UP FI	5.0(3)N2(2.11)	Cisco UCS Fabric Interconnects
Cisco	UCS 5104 Chassis	N/A	Cisco UCS Blade server chassis (FC-variant)
Cisco	UCS 2208XP FEX	2.1(1)	Cisco UCS Fabric Extenders for Blade Server chassis (FC-variant)
Cisco	UCS B200 M3 servers	2.1(1)	Cisco B200 M3 blade servers (FC-variant)
Cisco	UCS VIC 1240	2.1(1)	Cisco VIC 1240 adapters (FC-variant)
Cisco	UCS 2232PP FEX	5.0(3)N2(2.11)	UCS Fabric Extenders (iSCSI-variant)
Cisco	UCS C220 M3 servers	1.4(6c) or later – CIMC	Cisco C220 M3 rack servers (FC and iSCSI-variant)
Cisco	UCS VIC 1225	C220M3.1.5.1a.0 - BIOS	Cisco UCS VIC adapter (iSCSI-variant)
EMC	EMC VNX5300	05.32.000.5.006	EMC VNX storage array (FC-variant)
EMC	EMC VNXe3300	2.4.0.20932	VNXe storage array (iSCSI-variant)
EMC	EMC Avamar	6.1 SP1	EMC data backup software
EMC	Data Domain OS	5.3	EMC data domain operating system
VMware	ESXi 5.1	5.0 build 799733	VMware Hypervisor
VMware	vCenter Server	5.0 build 455964	VMware management
Microsoft	Microsoft Windows Server 2008 R2	2008 R2 SP1	Operating system to host vCenter server
Microsoft	Microsoft SQL server	2008 R2	Database server SQL R2 Enterprise edition for vCenter

Table 2 outlines the Cisco UCS B200 M3 or C200 M3 server configuration for the two variants of VMware architecture. Table 2 shows the configuration on per server basis.

Table 2 Server configuration details

Component	Capacity
Memory (RAM)	64 GB (8X8 MB DIMM)
Processor	2 x Intel® Xenon ® E5-2650 CPUs, 2 GHz, 8 cores, 16 threads
Local storage	Cisco UCS RAID SAS 2008M-8i Mezzanine Card, With 2 x 67 GB slots for RAID 1 configuration each

Both the architectures assume that there is an existing infrastructure/ management network available in which a virtual machine hosting vCenter server and Windows Active Directory/ DNS server are present.

Required number of C or B series servers and storage array type change depending on number of Virtual Machines. Table 3 highlights the change in the hardware components, as required by different scale points. Typically, 25 reference Virtual Machines are deployed per server.

Table 3 Hardware components for different scale

Components	VMware 100 VMs	VMware 125 VMs
Servers	4 x Cisco C220 M3 servers	5 x Cisco B200 M3 servers or C200 M3 servers
Storage	EMC VNXe3300	EMC VNX5300

Figure 8 and Figure 9 show a high level Cisco solution for EMC VSPEX VMware FC variant and iSCSI variant architectures respectively.

Figure 8 Reference Architecture for FC-variant

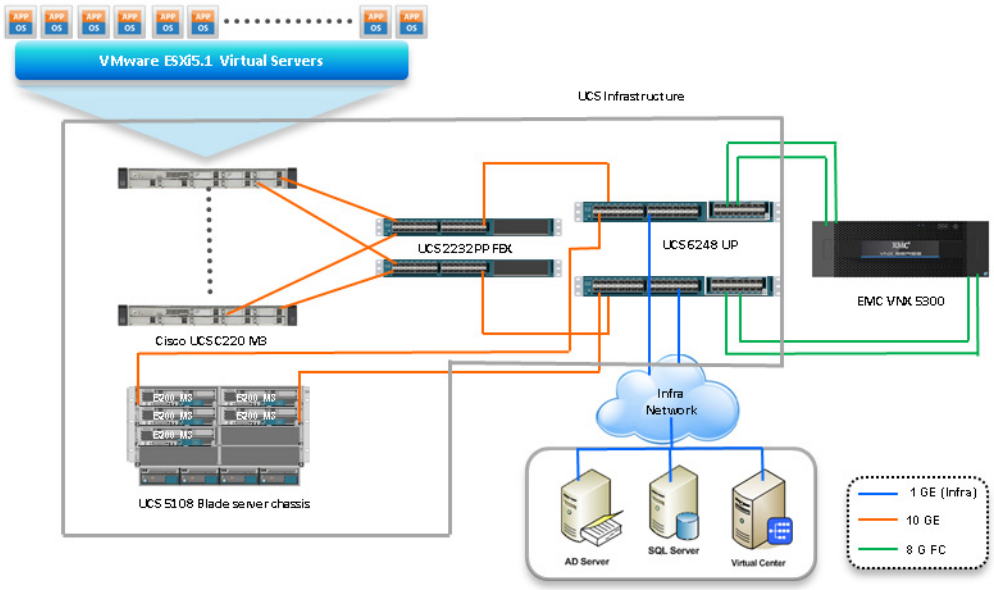

Figure 9 Reference Architecture for iSCSI-variant

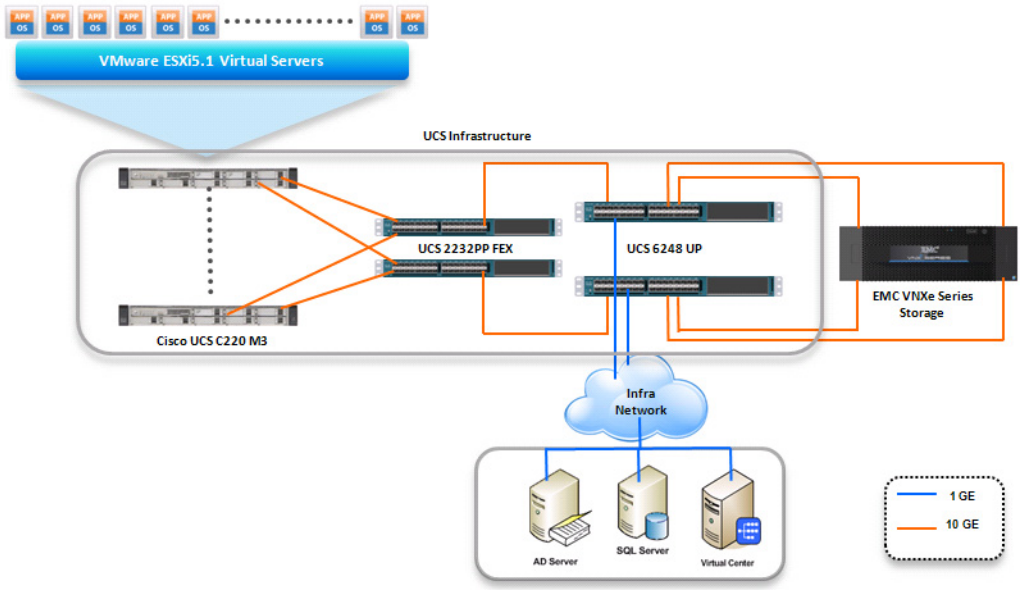

As it is evident in the above diagrams, following are the high level design points of VMware architecture for SMB market segment:

- We have directly attached storage array to UCS FIs
- iSCSI-variant uses only Ethernet as a network layer 2 media to access storage as well as the TCP/IP network
- Infrastructure network is on a separate 1GE network

- Network redundancy is built in by providing two switches, two storage controllers and redundant connectivity for data, storage and infrastructure networking.

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 network. However, design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

ESXi 5.1 is used as hypervisor operating system on each server and is installed on local hard drives. Typical load is 25 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 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

Virtual Machine Memory Concepts

The Figure 10 illustrates the use of memory settings parameters in the virtual machine.

Figure 10 Virtual Machine Memory Settings

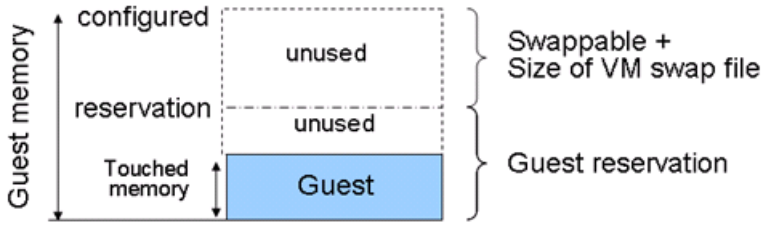

The vSphere memory settings for a virtual machine include the following parameters:

- Configured memory—Memory size of virtual machine assigned at creation.
- Touched memory—Memory actually used by the virtual machine. vSphere allocates 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 Resources used by a single VM

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

Following are the 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/ESXi if physical memory comes under contention. Performance tests show that the balloon driver allows ESXi/ESXi to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESXi/ESXi 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

The VSPEX architecture for VMware virtual machines for SMB market segment uses FC or iSCSI to access storage arrays. iSCSI is used with smaller scale points with VNXe3300 storage array, while FC is used with VNX5300 storage array. This simplifies the design and implementation for the small to medium level businesses. vSphere provides many features that take advantage of EMC storage technologies such as auto discovery of storage resources and ESXi hosts in vCenter and VNX/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 iSCSI, or multiple protocols for tiered scenarios such as using Fibre Channel for high-throughput storage pools and NFS for high-capacity storage pools. As mentioned before, this architecture uses iSCSI as storage access protocol.

For more information, see the VMware whitepaper on Comparison of Storage Protocol Performance in VMware vSphere 5 at:

http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf

Storage Best Practices

Following are the descriptions of 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

This architecture demonstrates use and benefits of Adapter-FEX technology using Cisco UCS VIC adapter. Each B200 M3 blade server and C220 M3 rack server has one physical adapter with two 10 GE links going to fabric A and fabric B for high availability. In FC-variant, Cisco UCS VIC 1225 or 1240 presents four virtual Network Interface Cards (vNICs) to the hypervisor, two vNICs per fabric path. In iSCSI-variant, the Cisco UCS VIC 1225 adapter presents six virtual Network Interface Cards (vNICs) to the hypervisor, three vNICs per fabric path. The MAC addresses to these vNICs are assigned using MAC address pool defined on the UCSM. These 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 three vSwitches per hypvervisor:
  - vSwitch (default)—Used for management and vMotion traffic
  - iSCSI-vSwitch (default)—Used for iSCSI storage traffic (iSCSI-variant only)
  - 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, switch (FI or FEX) failures, and in case of UCS, upstream switch failure (due to “End Host Mode” architecture).
- Enable PortFast on ESX/ESXi host uplinks—Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. Cisco UCS Fabric Extenders are not really separate switches – they are line cards to the Cisco UCS Fabric Interconnect, and the Cisco UCS Fabric Interconnects run in end-host-mode and avoid running Spanning Tree Protocol. Given this, there is no need to enable port-fast on the ESXi host uplinks. However, it is recommended that you enable portfast on the infrastructure switch that connects to the UCS Fabric Interconnect uplinks for faster convergence of STP in the events of FI reboots or FI uplink flaps.
- 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 11 illustrates storage architecture for 125 virtual machines on VNX5300 for FC-variant of architecture:

Figure 11 EMC VNX5300 Storage Architecture for 125 VMs

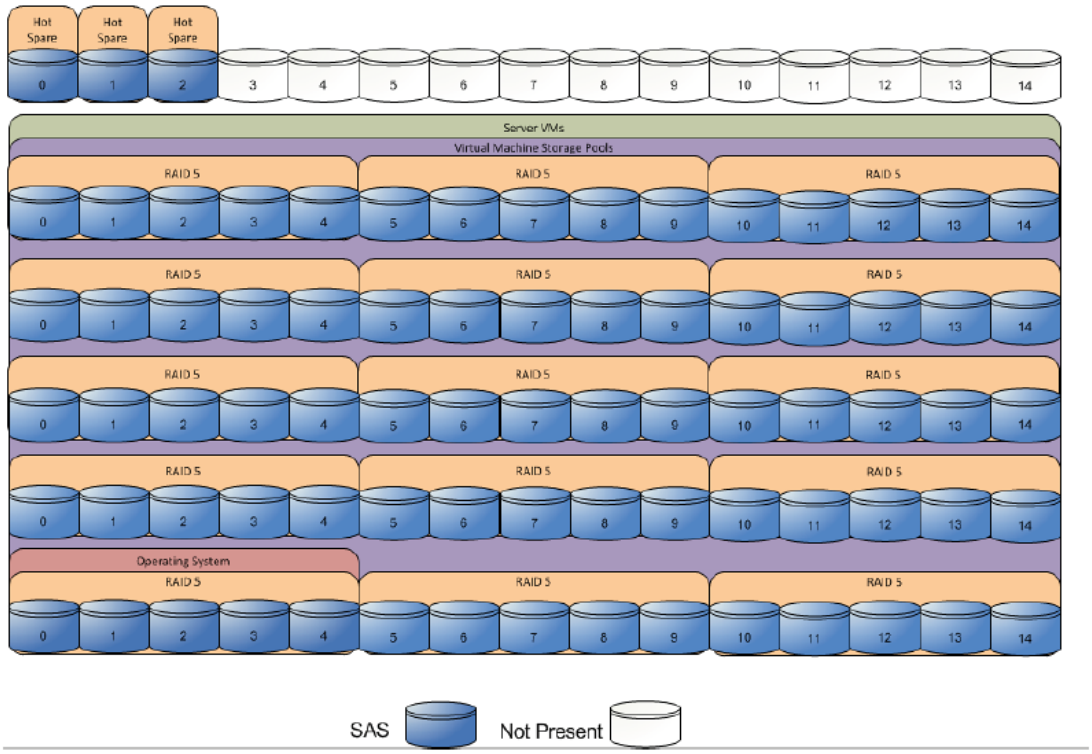

Figure 12 illustrates storage architecture for 100 virtual machines on VNXe3300 for iSCSI-variant of architecture:

Figure 12 EMC VNX3300 Storage Architecture for 100 VMs

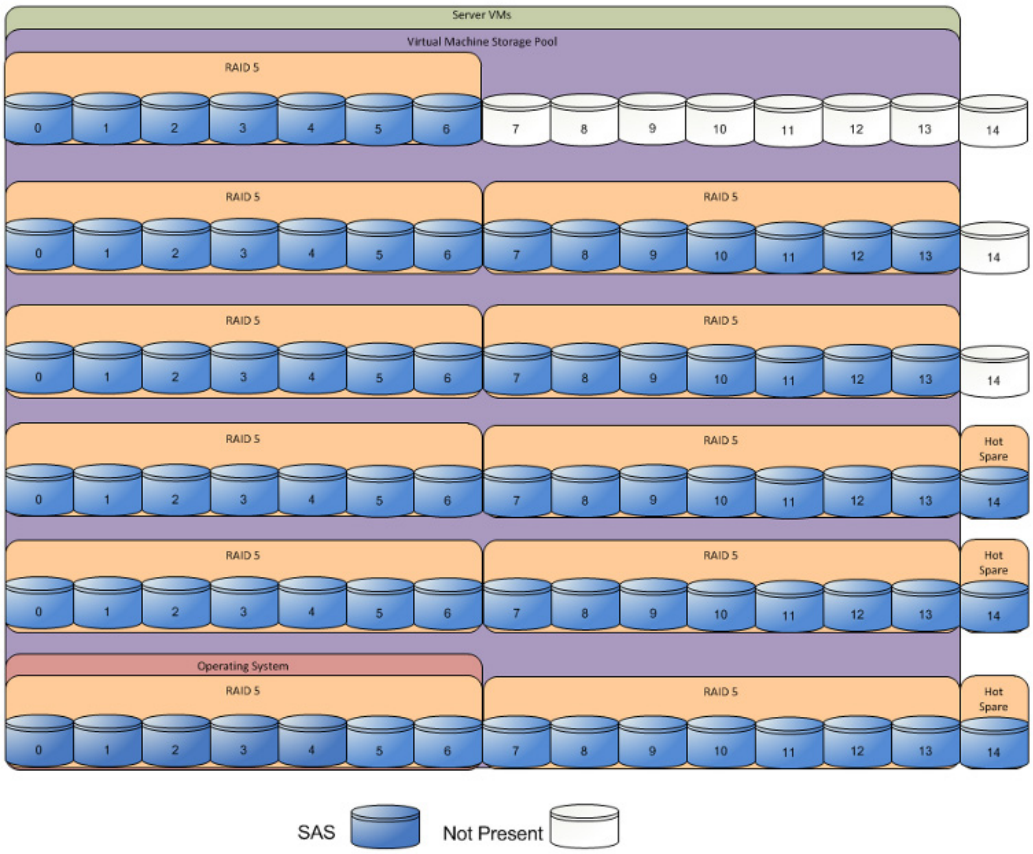

Table 5 provides the size of the datastores for both the architectures laid out Figure 11 and Figure 12:

Table 5 datastores for 100 and 125 VMs

Parameters	100 Virtual Machines	125 Virtual Machines
Disk capacity and types	600 GB SAS	600 and 300 GB SAS
Number of disks	77	60 (600 GB)
RAID type	6+1 RAID 5 groups	4+1 RAID 5 groups
Number of pools	11	15

For all the architectures, EMC recommends one hot spare disk allocated for each 30 disks of a given type.

The VNX/VNXe 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 (FC-variant)
  - Two iSCSI initiator IP addresses, one for each fabric (iSCSI-variant)
  - Two iSCSI IQN name identifiers (iSCSI-variant)

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 iSCSI or FC protocol.
- Server pool is defined with automatic qualification policy and criteria. Rack 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 rack 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 fault server and physically remove the server for replacement of fault hardware (or to completely remove the faulty server).
- Physically install the new server and connect it to the Fabric Extenders. Let the new server be discovered by UCSM.
- Associate the service profile to the newly deployed rack server. This would boot the same ESXi server image from the storage array as 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 and rack servers can be added in the same server pool.

Network High Availability Design – FC-variant

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

- Mix of Cisco UCS B200 M3 and C220 M3 servers are used, managed by Cisco UCS Manager (UCSM).
- Fabric A and Fabric B are used with host based FC multi-pathing for high availability
- EMC VNX5300 storage array is directly attached to Cisco UCS Fabric Interconnects
- Two 10GE links between FI and FEX provides enough bandwidth oversubscription for the SMB segment private cloud. The oversubscription can be reduced by adding more 10GE links between FI and FEX if needed by the VMs running on the ESXi hosts.
- Two vSwitches are used per host, as discussed in the Virtual Networking design section.

Storage is made highly available by deploying following practices:

- EMC VNX storage arrays provide two Storage Processors (SPs): SP-A and SP-B
- Cisco UCS Fabric Interconnects A and B are connected to SP-A and SP-B respectively.
- Port-channels or port-aggregation is not implemented or required in this architecture.
- Storage Processors are always in the active/active mode; if the target cannot be reached on SAN-A, server can access the LUNs through SAN-B and storage-processor inter-link.
- On hosts, boot order lists vHBA on both fabrics for high-availability.

Figure 13 Logical Layout of the FC-variant Architecture

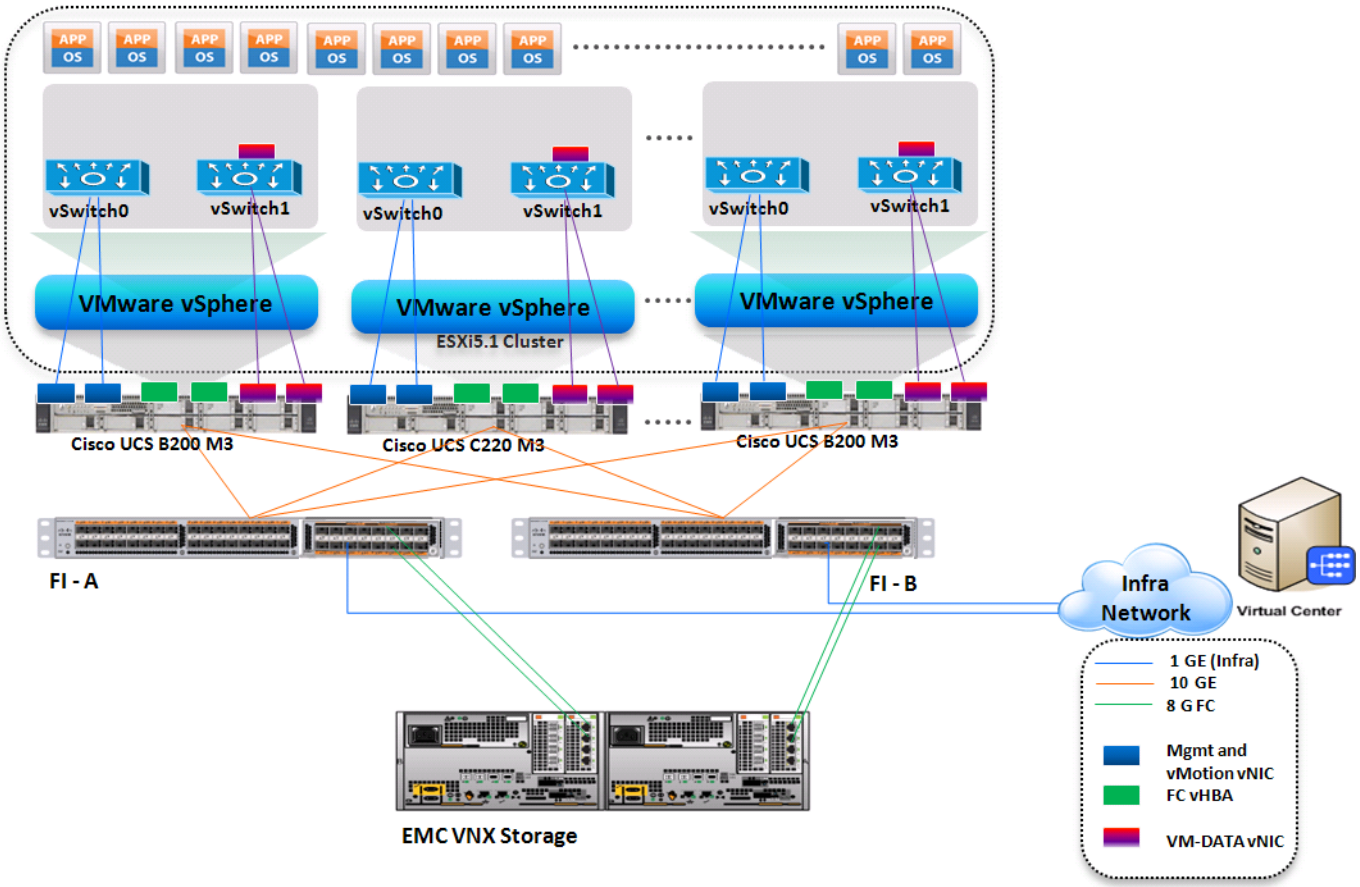

Network High Availability Design – iSCSI-variant

Figure 14 demonstrates logical layout of the iSCSI-variant architecture. Following are the key aspects of this solution:

- Cisco UCS C220 M3 servers are used, managed by Cisco UCS Manager (UCSM).
- Fabric A and Fabric B are used with host based iSCSI multi-pathing for high availability
- EMC VNXe storage array is directly attached to Cisco UCS Fabric Interconnects
- Two 10GE links between FI and FEX provides enough bandwidth oversubscription for the SMB segment private cloud. The oversubscription can be reduced by adding more 10GE links between FI and FEX if needed by the VMs running on the ESXi hosts.
- Three vSwitches are used per host, as discussed in the Virtual Networking design section.

Storage is made highly available by deploying following practices:

- VNXe storage arrays provide two Storage Processors (SPs):
  - SP-A
  - SP-B

- Both the Cisco UCS Fabric Interconnects (A and B) are connected to both Storage Processors, however, a given FI connects to the same port on each SP. FI-A connects to “eth10” port of SP-A and SP-B, while FI-B connects to “eth11” port of SP-A and SP-B.
- Port-channels or port-aggregation is not implemented or required in this architecture.
- iSCSI is implemented on different VLANs and subnets on fabric A and fabric B, as per the best practice guidelines.
- Storage Processors are always in the active/stand-by mode; the links are up on both SPs, but packets are accepted only on one of the two SPs.
- Datastore in Storage Array is accessible through only one iSCSI server, but the iSCSI server can listen on two IP addresses, one on each fabric.
- On hosts, boot order lists iSCSI initiator on both fabrics for high-availability.

Figure 14 Logical Layout of the iSCSI-variant Architecture

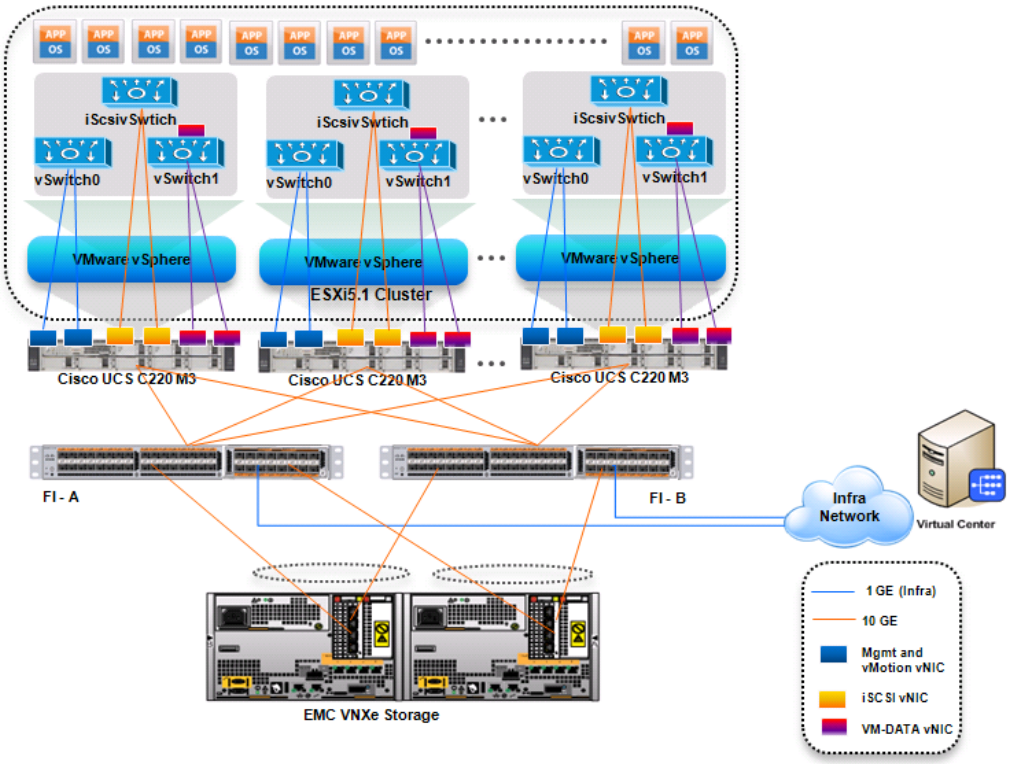

MTU Setting

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

- iSCSI 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:

- Ethernet ports on VNXe Storage Processors
- Appliance ports on FIs (through QoS policies)
- vNICS in service profiles
- System QoS classes
- vSwitches on the ESXi hosts
- VM-Kernel ports used for vMotion and storage access on the ESXi hosts

In addition to these end-points, infrastructure networks that connect FI-A and FI-B also need to support jumbo-MTU for vMotion traffic. iSCSI traffic is completely self-contained on the UCS fabrics.

The next section provides information on sizing guidelines of the Cisco solution for EMC VSPEC 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 6 Virtual Machine Characteristics

Characteristics	Value
VM operating system	Microsoft Windows Server 2008 R2
Virtual processor (vCPU) per VM	1
RAM per VM	2 GB
Available storage capacity per VM	100 GB
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.

Applying the Reference Workload

When considering an existing server that will move into a virtual infrastructure, you have the opportunity to gain efficiency by right-sizing the virtual hardware resources assigned to that system.

The reference architectures create a pool of resources sufficient to host a target number of reference virtual machines as described above. It is entirely possible that customer virtual machines may not exactly match the specifications above. In that case, you can say that a single specific customer virtual machine is the equivalent of some number of reference virtual machines, and assume that number of virtual machines have been used in the pool. You can continue to provision virtual machines from the pool of resources until it is exhausted. Consider these examples:

Example 1 Custom Built Application

A small custom-built application server needs to move into this virtual infrastructure. The physical hardware supporting the application is not being fully utilized at present. A careful analysis of the existing application reveals that the application can use one processor, and needs 3 GB of memory to run normally. The IO workload ranges between 4 IOPS at idle time to 15 IOPS when busy. The entire application is only using about 30 GB on local hard drive storage.

Based on these numbers, the following resources are needed from the resource pool:

- CPU resources for one VM
- Memory resources for two VMs
- Storage capacity for one VM
- IOPS for one VM

In this example, a single virtual machine uses the resources of two of the reference VMs. If the original pool had the capability to provide 100 VMs worth of resources, the new capability is 98 VMs.

Example 2 Point of Sale System

The database server for a customer's point-of-sale system needs to move into this virtual infrastructure. It is currently running on a physical system with four CPUs and 16 GB of memory. It uses 200 GB storage and generates 200 IOPS during an average busy cycle.

The following are the requirements to virtualize this application:

- CPUs of four reference VMs
- Memory of eight reference VMs
- Storage of two reference VMs
- IOPS of eight reference VMs

In this case the one virtual machine uses the resources of eight reference virtual machines. If this was implemented on a resource pool for 100 virtual machines, there would be 92 virtual machines of capability remaining in the pool.

Example 3 Web Server

The customer's web server needs to move into this virtual infrastructure. It is currently running on a physical system with two CPUs and 8GB of memory. It uses 25 GB of storage and generates 50 IOPS during an average busy cycle.

The following are the requirements to virtualize this application:

- CPUs of two reference VMs
- Memory of four reference VMs
- Storage of one reference VMs
- IOPS of two reference VMs

In this case the virtual machine would use the resources of four reference virtual machines. If this was implemented on a resource pool for 100 virtual machines, there would be 96 virtual machines of capability remaining in the pool.

Example 4 Decision Support Database

The database server for a customer's decision support system needs to move into this virtual infrastructure. It is currently running on a physical system with ten CPUs and 48 GB of memory. It uses 5 TB of storage and generates 700 IOPS during an average busy cycle.

The following are the requirements to virtualize this application:

- CPUs of ten reference VMs
- Memory of twenty-four reference VMs
- Storage of fifty-two reference VMs
- IOPS of twenty-eight reference VMs

In this case the one virtual machine uses the resources of fifty-two reference virtual machines. If this was implemented on a resource pool for 100 virtual machines, there would be 48 virtual machines of capability remaining in the pool.

Summary of Example

The four examples presented illustrate the flexibility of the resource pool model. In all four cases the workloads simply reduce the number of available resources in the pool. If all four examples were implemented on the same virtual infrastructure, with an initial capacity of 100 virtual machines they can all be implemented, leaving the capacity of thirty four reference virtual machines in the resource pool.

In more advanced cases, there may be tradeoffs between memory and I/O or other relationships where increasing the amount of one resource decreases the need for another. In these cases, the interactions between resource allocations become highly complex, and are outside the scope of this document. However, once the change in resource balance has been examined, and the new level of requirements is known; these virtual machines can be added to the infrastructure using the method described in the examples.

The next section provides step by step procedure for deploying the Cisco solution for EMC VSPEX VMware architectures.

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 and configure storage array for resource pools and iSCSI servers (iSCSI-variant only)
4. Prepare UCS FIs and configure UCSM
5. Configure datastores for ESXi images
6. Install ESXi servers and vCenter infrastructure
7. Install and configure vCenter server
8. Configure storage for VM datastores, install and instantiate VMs through vCenter

9. 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 following table 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 7 Deployment prerequisites

Requirement	Description	Reference
Hardware	Cisco UCS Fabric Interconnects, Fabric Extenders and UCS chassis for network and compute infrastructure	See the corresponding product documentation
	Cisco UCS B200 M3 and/or C220 M3 servers to host virtual machines	
	VMware vSphere™ 5.1 server to host virtual infrastructure servers **Note** This requirement may be covered in the existing infrastructure	
	EMC VNX/VNXe storage—Multiprotocol storage array with the required disk layout as per architecture requirements	

Table 7 Deployment prerequisites

Requirement	Description	Reference
Software	VMware ESXi™ 5.1 installation media	See the corresponding product documentation
	VMware vCenter Server 5.1 installation media	
	EMC VSI for VMware vSphere: Unified Storage Management – Product Guide	
	EMC VSI for VMware vSphere: Storage Viewer—Product Guide	
	Microsoft Windows Server 2008 R2 SP1 installation media (suggested OS for VMware vCenter)	
	Microsoft SQL Server 2008 R2 SP1 **Note** This requirement may be covered in the existing infrastructure	
Licenses	VMware vCenter 5.1 license key	Consult your corresponding vendor to obtain license keys
	VMware ESXi 5.1 license key	
	Microsoft SQL Server license key **Note** This requirement may be covered in the existing infrastructure	
	Microsoft Windows Server 2008 R2 SP1 license key	

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, page 192 provides a set of tables 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 *VNX/VNXe Series Configuration Worksheet*, available on the EMC online support website, to provide the most comprehensive array-specific information.

Connect Network Cables

See the Cisco UCS FI, FEX, Blade servers chassis, B-series and C-series server and EMC VNXe 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.

Connectivity for FC-variant:

As shown in the following figure, there are four major cabling sections in this architecture:

- Cisco UCS Fabric Interconnects to EMX storage array - Fibre Channel links (shown in yellow)
- Cisco Fabric Interconnects to Cisco UCS Fabric Extenders links (shown in blue)
- Cisco UCS Fabric Extenders to Cisco UCS C220 M3 Server links (shown in green)
- Infrastructure connectivity (not shown)

Figure 15 Detailed Connectivity Diagram of the FC-variant Architecture

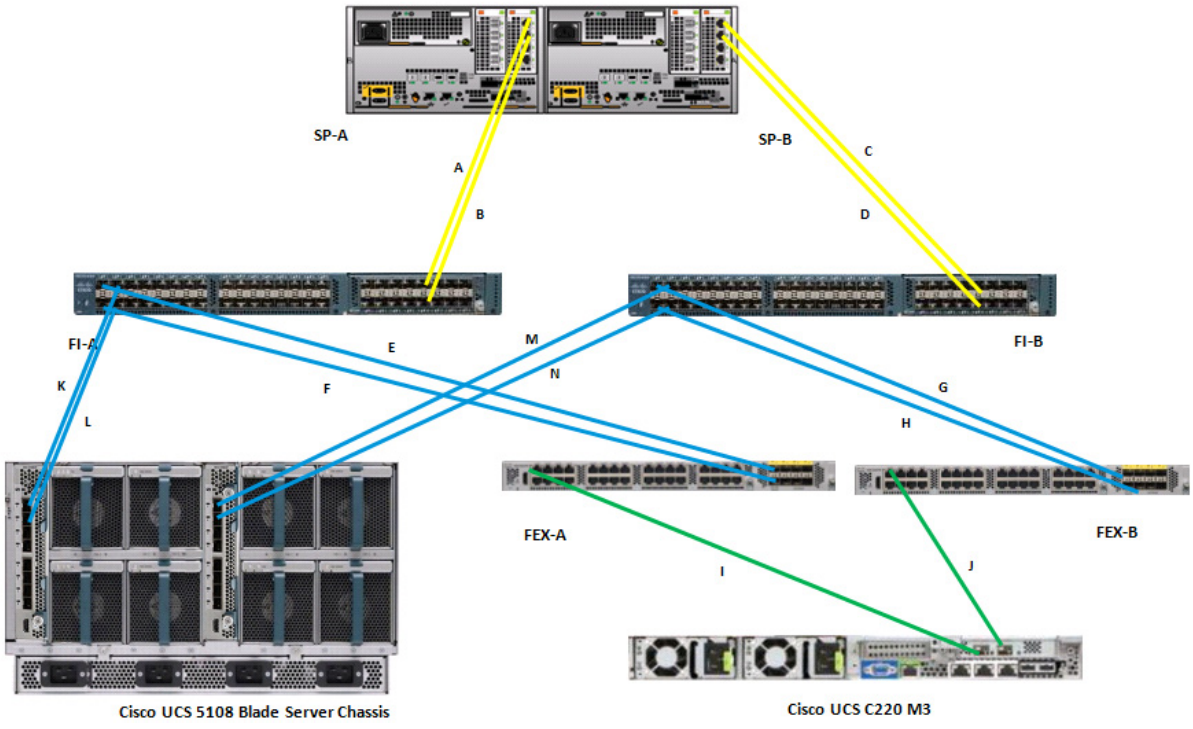

Figure 16 elaborates the detailed cable connectivity for the architecture.

Figure 16 Connectivity Details of FC-variant Architecture

Cable ID	Peer 1	Peer 2	VLAN	Mode	Description
A	FI-A, FC 2/9	SP-A,	Storage VSAN	Appliance	Directly attached storage on FI
B	FI-A, FC 2/10	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
C	FI-B, FC 2/9	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
D	FI-B, FC 2/10	SP-A,	Storage VSAN	Appliance	Directly attached storage on FI
E,F	FI-A, Eth 1/1, 1/2	FEX-A uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
G,H	FI-A, Eth 1/1, 1/2	FEX-B uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
I	FEX-A, port 1	C220-M3 VIC port 1	N/A	VNTag (internal)	Server to fabric A. VLANs are allowed on per vNIC basis
J	FEX-B, port 1	C220-M3 VIC port 2	N/A	VNTag (internal)	Server to fabric B. VLANs are allowed on per vNIC basis
K,L	FI-A, Eth 1/3, 1/4	5108 Chassis, FEX 2208 Left	N/A	Server	FI/FEX 20GE port-channel connectivity
M,N	FI-B, Eth 1/3, 1/4	5108 Chassis, FEX 2208 Right	N/A	Server	FI/FEX 20GE port-channel connectivity
(not shown)	Eth 2/1, 2/2 on FI-A and FI-B	Uplink switch	All	Uplink	Uplink to Infrastructure network

Connectivity for iSCSI-variant:

As shown in the following figure, there are four major cabling sections in this architecture:

- Cisco UCS Fabric Interconnects to EMC storage array - 10G Ethernet links (shown in yellow)
- Cisco UCS Fabric Iinterconnects to Cisco UCS Fabric Extenders links (shown in blue)
- Cisco UCS Fabric Extenders to Cisco UCS C220 M3 Server links (shown in green)
- Infrastructure connectivity (not shown)

Figure 17 Detailed Connectivity Diagram of the Architecture

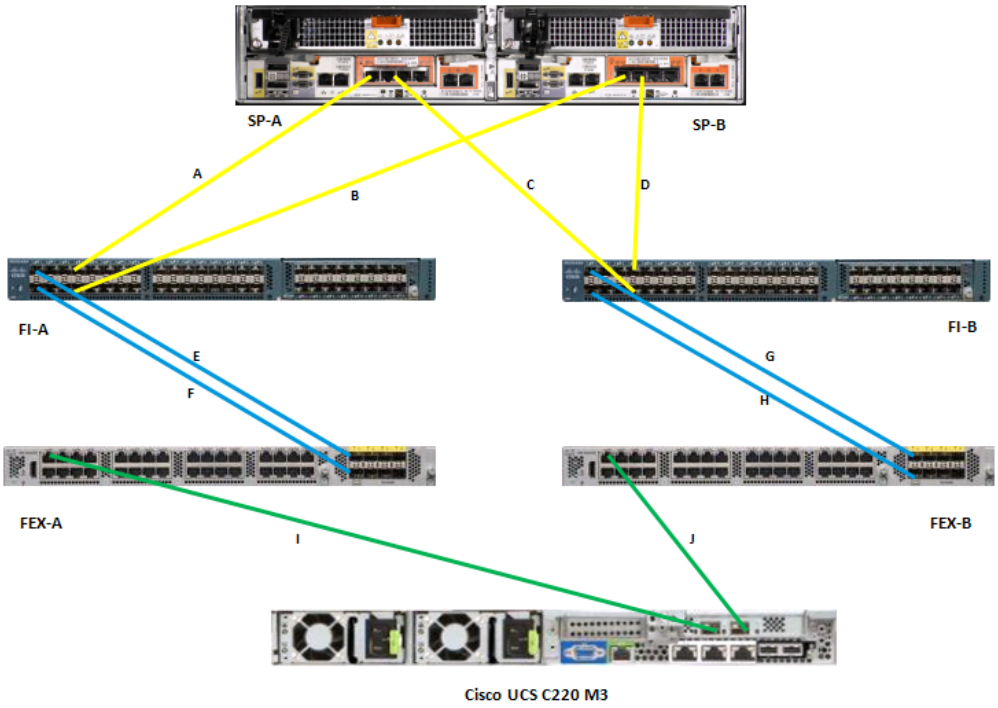

Figure 18 elaborates the detailed cable connectivity for the architecture.

Figure 18 Connectivity Details of iSCSI-variant Architecture

Cable ID	Peer 1	Peer 2	VLAN	Mode	Description
A	FI-A, Eth 1/5	SP-A, eth10	Storage (A)	Appliance	Directly attached storage on FI
B	FI-A, Eth 1/6	SP-B, eth10	Storage (A)	Appliance	Directly attached storage on FI
C	FI-B, Eth 1/5	SP-B, eth11	Storage (B)	Appliance	Directly attached storage on FI
D	FI-B, Eth 1/6	SP-A, eth11	Storage (B)	Appliance	Directly attached storage on FI
E,F	FI-A, Eth 1/1, 1/2	FEX-A uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
G,H	FI-A, Eth 1/1, 1/2	FEX-B uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
I	FEX-A, port 1	C220-M3 VIC port 1	N/A	VNTag (internal)	Server to fabric A. VLANs are allowed on per vNIC basis
J	FEX-B, port 1	C220-M3 VIC port 2	N/A	VNTag (internal)	Server to fabric B. VLANs are allowed on per vNIC basis
(not shown)	Eth 2/1, 2/2 on FI-A and FI-B	Uplink switch	All	Uplink	Uplink to Infrastructure network

An important difference to recognize between FC and iSCSI variants are transport used to connect storage array with UCS FIs. In case of FC-variant, 8G Fibre Channel cables are connected from SP-A to FI-A and SP-B to FI-B. This is required to maintain “SAN-A/SAN-B separation”. In case of iSCSI-variant, 10G Ethernet cables are crisscrossed between SP-A/SP-B and FI-A/FI-B. This design maintains high-availability in iSCSI-variant.

By connecting all the cables as outlined above, and you would be ready to configure the storage array and UCSM.

Prepare and Configure Storage Array for Resource Pools and iSCSI Servers (iSCSI-variant only)

This subsection for EMC VNXe3300 storage array is for iSCSI-variant of the architecture only. Please refer to the Configuration Guide of your respective storage array on how to perform initial configuration of storage array.

At a high level to configure storage array for this solution, you need to complete these steps:

1. Create storage resource pools for ESXi boot images, VM images and VM datastores.
2. Configure iSCSI servers.
3. Configure ESXi hosts that can access various datastores.
4. Configure VMware datastores using the resource pools created in step 1 on the iSCSI servers created on step 2.

Note that steps 1 and 2 can be completed without configuring UCS Manager. Configuration made in step 2 is used to configure boot policy for service profiles in UCS Manager. The iSCSI initiator IP address and IQN names given in the UCS Manager service profile are used to configure step 3, which is a pre-requisite for step 4. In this section, we would configure steps A and B.

Configure Resource Pools

Follow these steps to configure storage resource pools:

1. Using your browser, navigate to the management IP address of the storage array. Provide username and password.
2. Choose **Systems > Storage Pools**.
3. Click **Configure Disks** as shown in the image below.

Figure 19 Configuring Storage pools

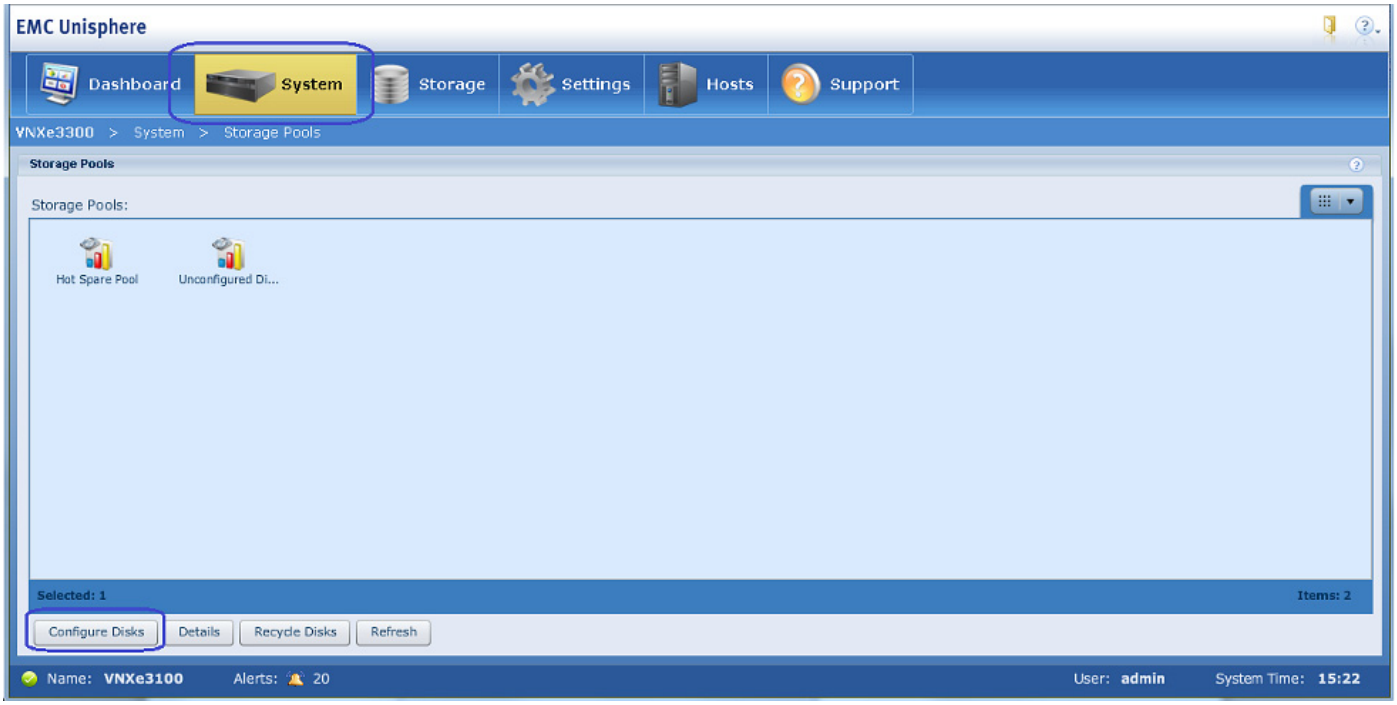

4. Click **Manually create a new pool** and choose Pool created for VMware Storage – General Purpose from drop-down menu.

Figure 20 Selecting the Configuration Mode for Pool Creation

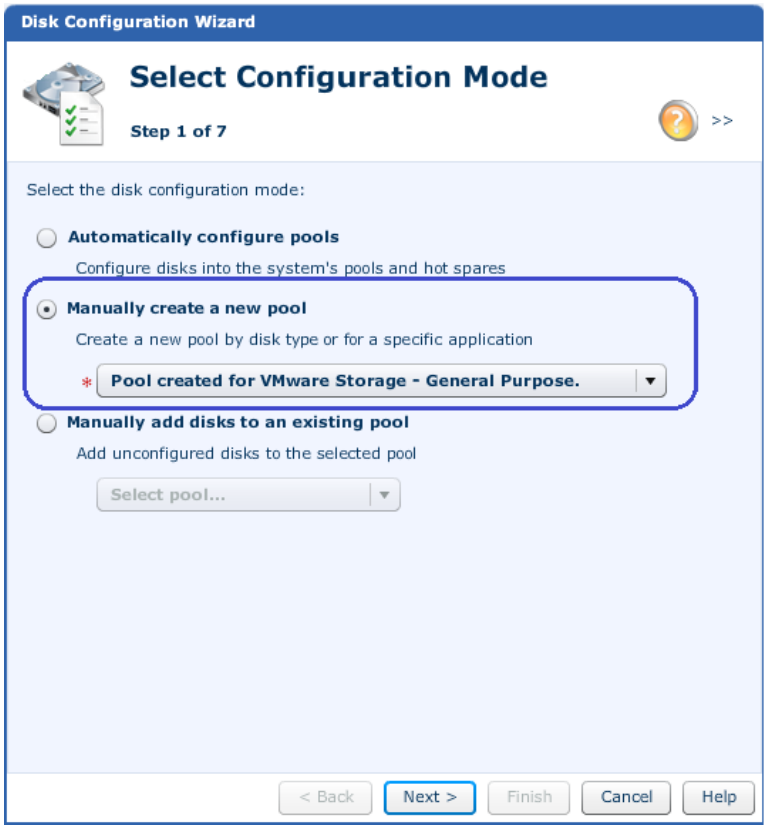

5. Enter pool name and (optional) description., click **Next**.
6. Select disk type for balanced performance storage profile as shown in Fig.

Figure 21 Specifying the Pool Name

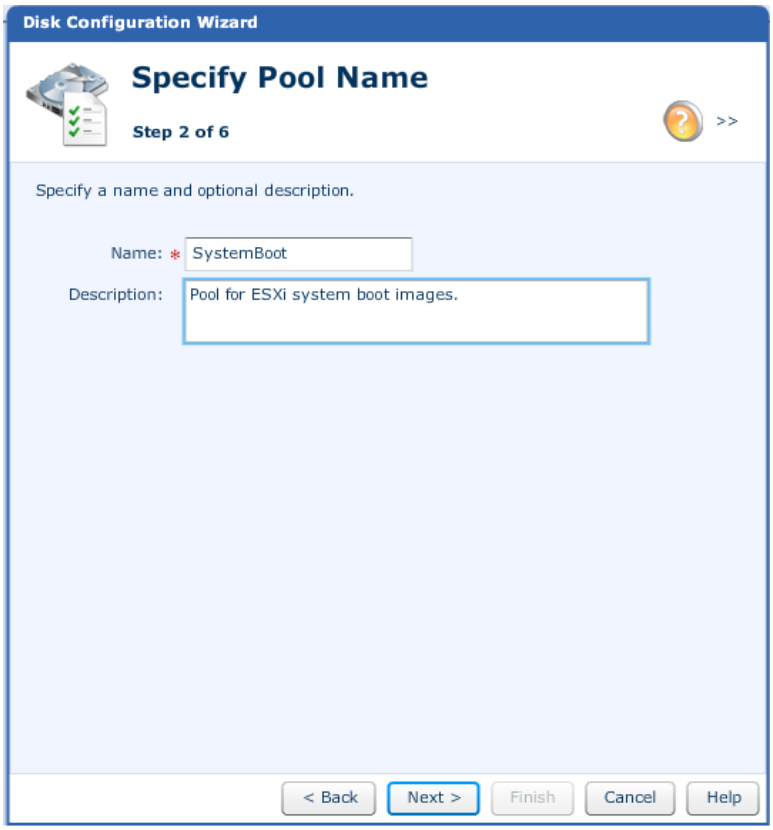

7. Choose 5 disks for the ESXi boot images (minimum disks required) for the number of storage disks required.

Figure 22 Selecting the Disk Type

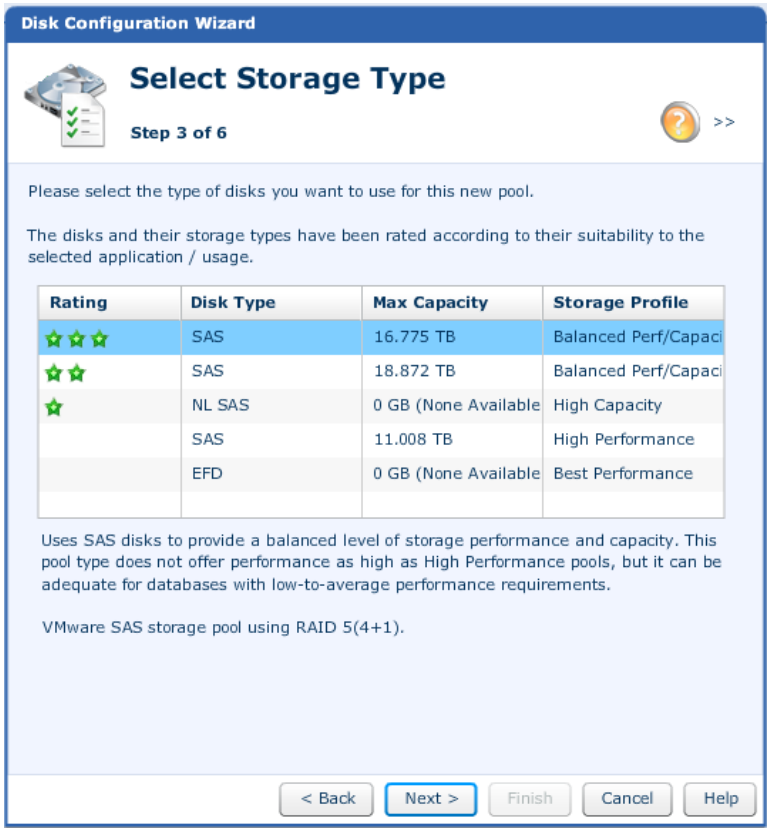

8. Click **Next** to verify configuration and click **Finish** to deploy the disk storage. Repeat these steps for three more datastores; one for VM operating system storage, and two for the VM data storage. Refer to storage architecture for the number of disks used for each type of datastores.

Configure iSCSI Servers

Next step is to create iSCSI servers on the storage array. You can configure multiple IP addresses for a given iSCSI server, one is accessible from each of the Cisco Fabric Interconnect. iSCSI server is configured on a given Storage Processor, and would automatically failover to the other Storage Processor if the primary SP fails or the link on which it is deployed goes down.

Following these steps to configure iSCSI servers on the storage array:

1. Choose **Settings > iSCSI Server Settings**.

Figure 23 Creating iSCSI Server

2. Click **Add iSCSI Server**.

Figure 24 Adding iSCSI Server

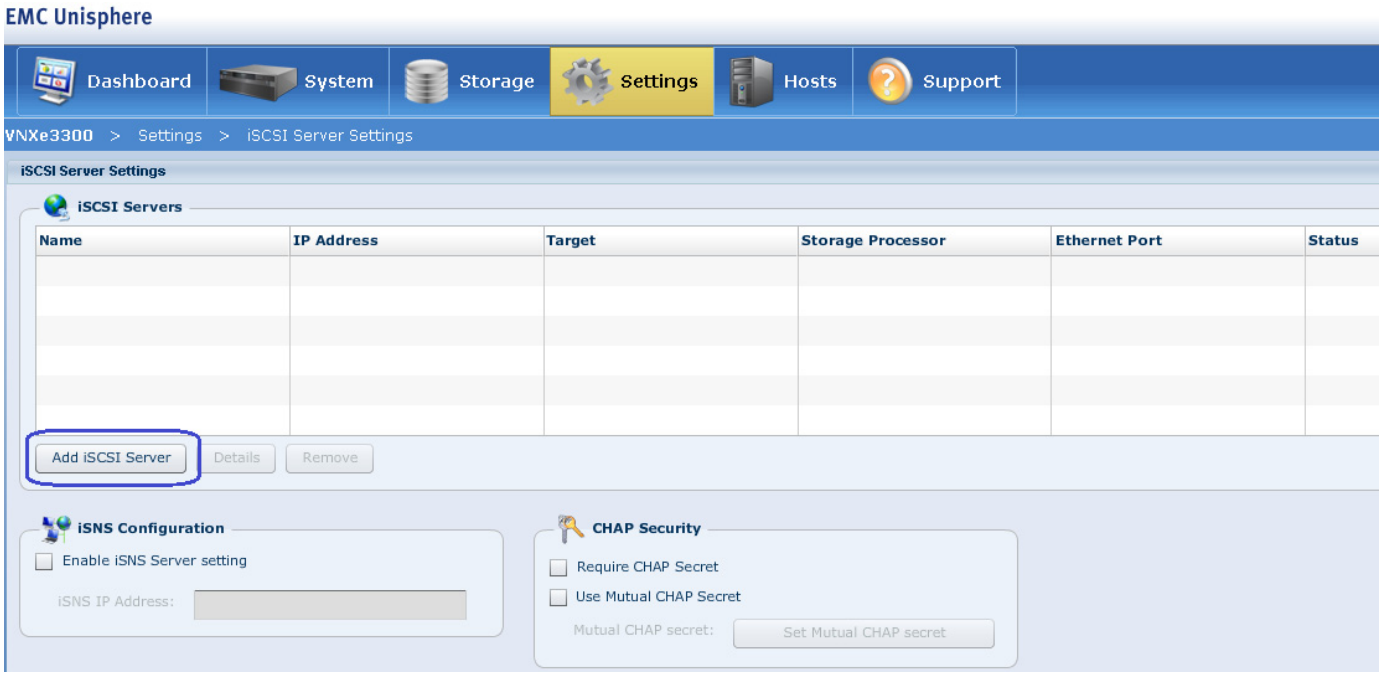

3. Enter server name, IP address and subnet mask. Make sure that the IP address is given from the subnet corresponding to the fabric-A storage VLAN subnet. See the Customer Configuration Data Sheet, page 192 for details. Click **Show advanced**, and choose SP A for Storage processor, eth10 interface for Ethernet Port to deploy the iSCSI server as shown in Figure 25.

Figure 25 Specifying the Network Interface

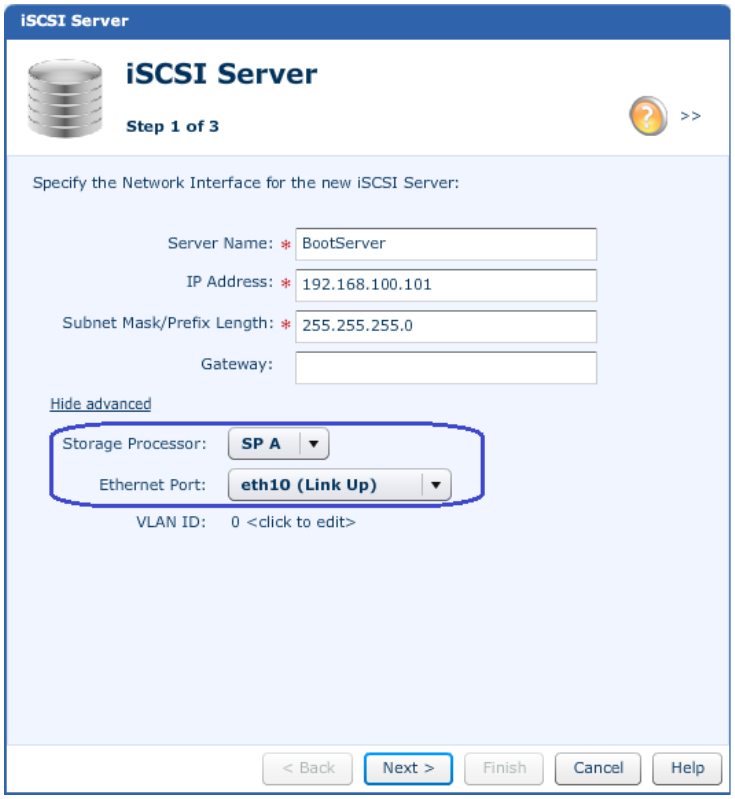

4. Click **Next** to verify the configuration and click **Finish**. Once the control goes back to the "iSCSI Server Settings" page, select the iSCSI server just created, and click **Details** to add one more IP address on the other fabric.

Figure 26 Adding iSCSI Server on Fabric B

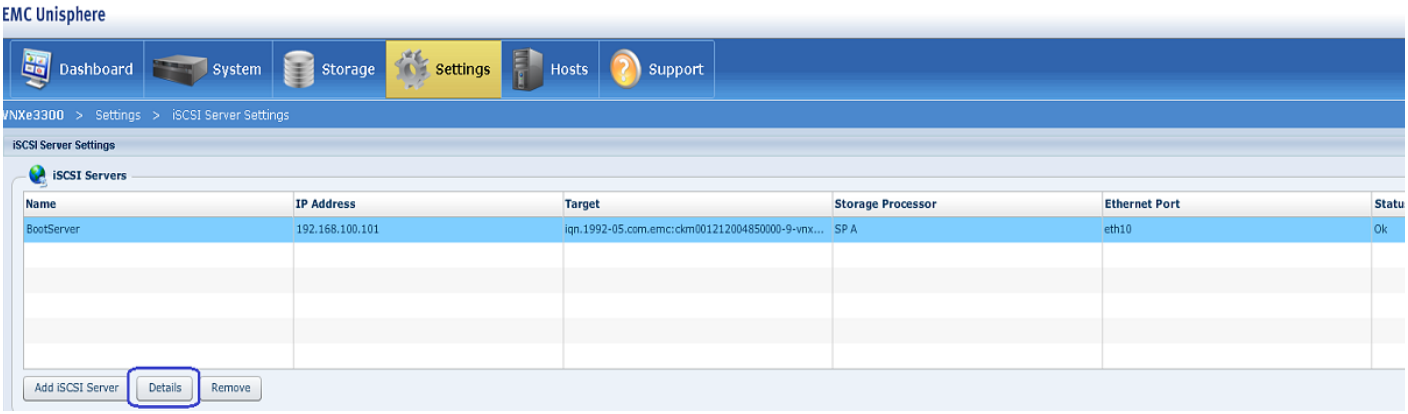

5. Click **Add Network Interface**.

Figure 27 Adding Network Interface in EMC Unisphere

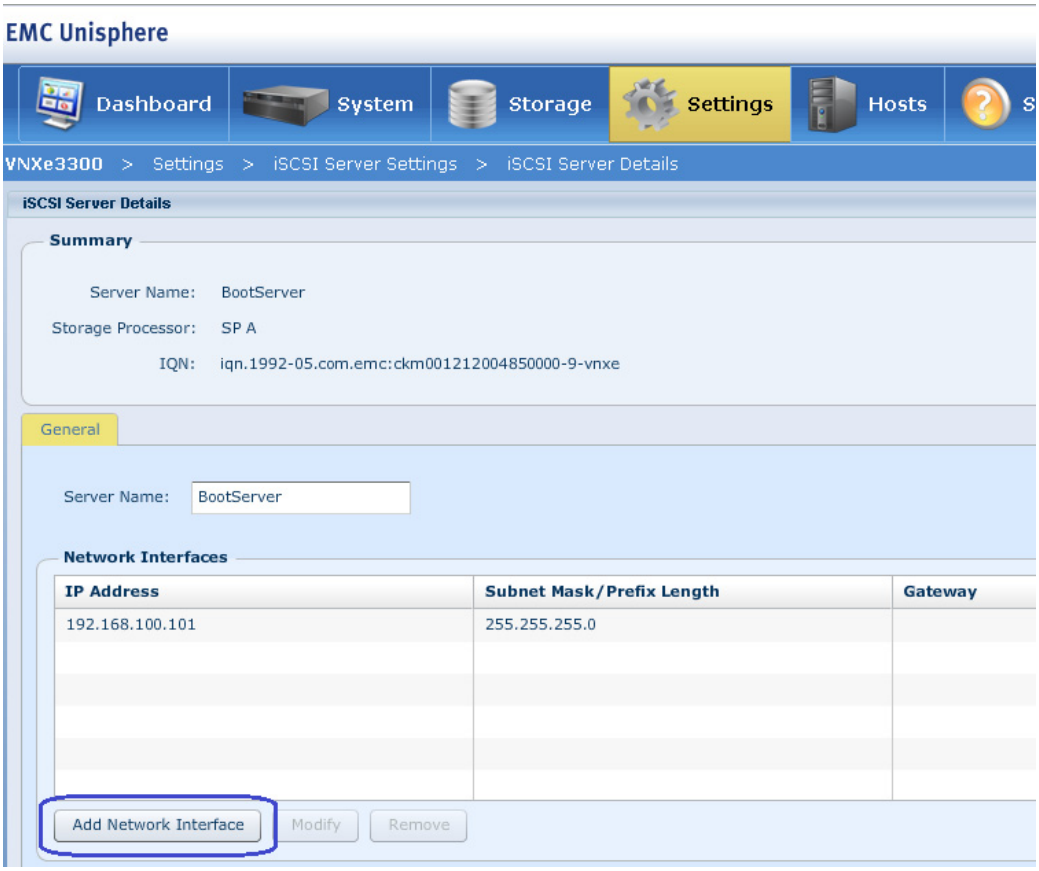

6. Specify the IP address and subnet mask corresponding to the storage VLAN subnet on fabric B. Click **Show advanced**.

Figure 28 Specifying IP Address and Subnet Mask

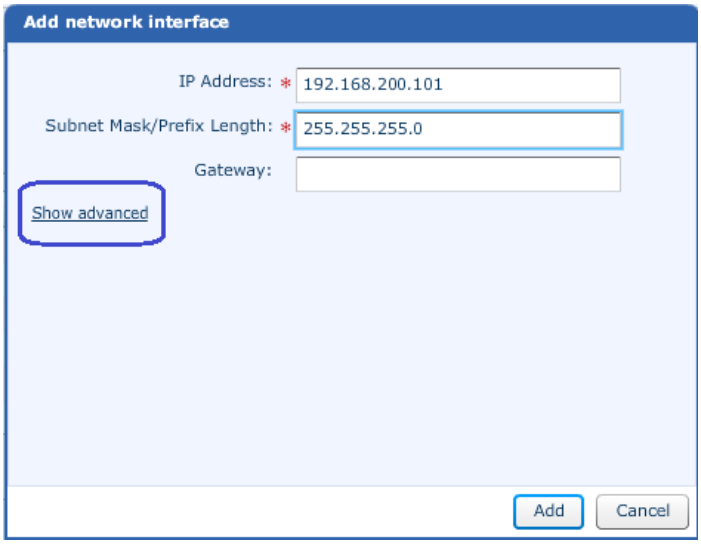

7. Since SP-A was already selected for the Storage Processor, this time only the Ethernet port option is shown. Choose eth11 as the Ethernet interface to deploy the newly configured IP address from the Ethernet Port drop-down list. Click **Add**.

Figure 29 Selecting the Ethernet Port

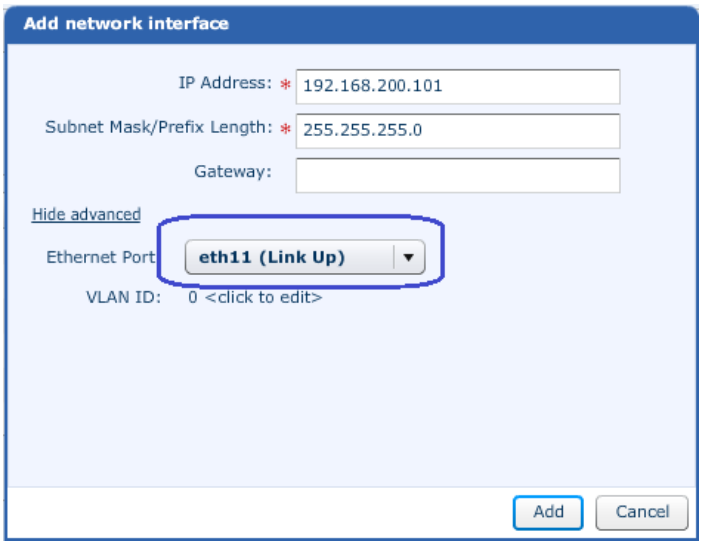

8. Once both the IP addresses are configured, IQN name and boot target IP addresses are available as shown in Figure 30. This information is used to configure iSCSI boot target in the UCS Manager service profile configuration.

Figure 30 IQN Name and IP Address Shown After the iSCSI Server Configuration

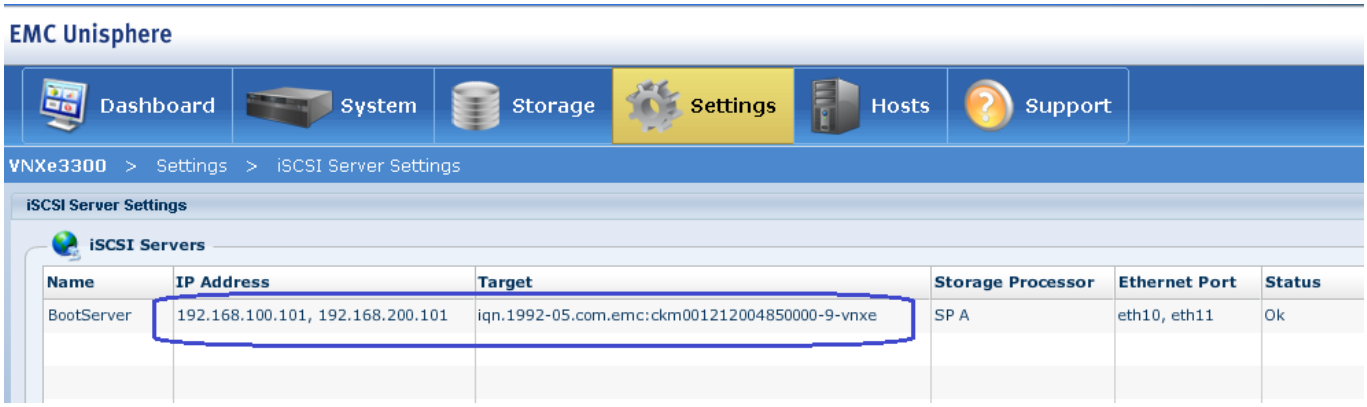

9. To increase MTU size of the iSCSI server NICs, choose **Settings > More configuration > Advanced Configuration**, and select eth10, and set MTU size to 9000 and save the configuration.

Figure 31 Setting Jumbo Frame for eth10 Interface

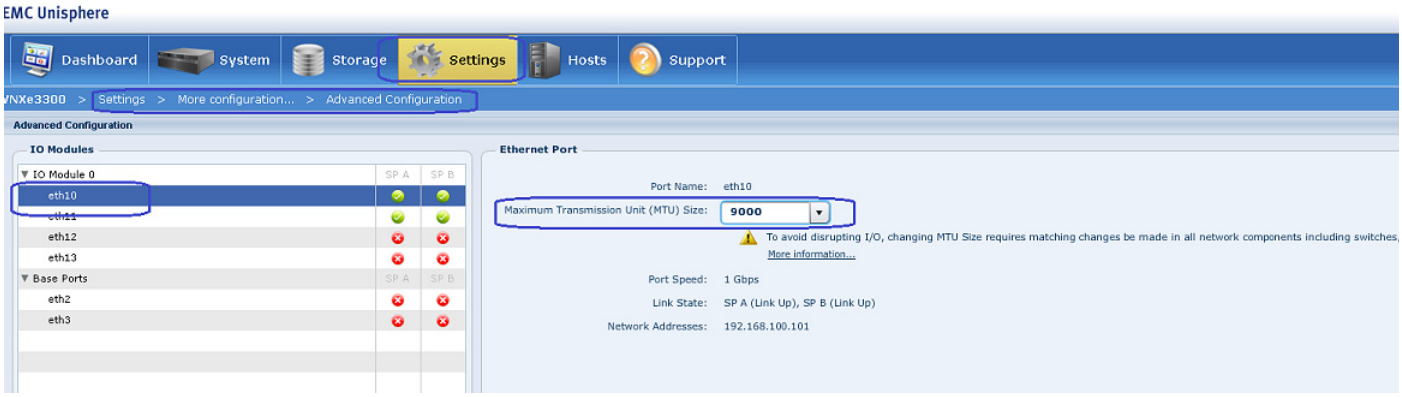

10. Repeat step 9 for eth11 interface.

Figure 32 Setting Jumbo Frame for eth11 Interface

Prepare Cisco UCS FIs and Configure Cisco UCS Manager

At a high level to configure Cisco UCS Fabric Interconnects and Cisco UCS Manager, you need to complete these steps:

1. Initial configuration of Cisco UCS Fabric Interconnects
2. Configuration for server discovery
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

The following sections provide a step-by-step procedure on configuring Cisco UCS Manager.

Initial Configuration of Cisco UCS Fabric Interconnects

At this point of time, the UCS FI, Blade servers chassis, FEX and C-series server must be mounted on the rack and appropriate cables must be connected as suggested in Architectural Overview, page 16. 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 UCS Manager 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 UCS Manager CLI login prompt.

Figure 33 Configuring Primary Cisco UCS Fabric Interconnect

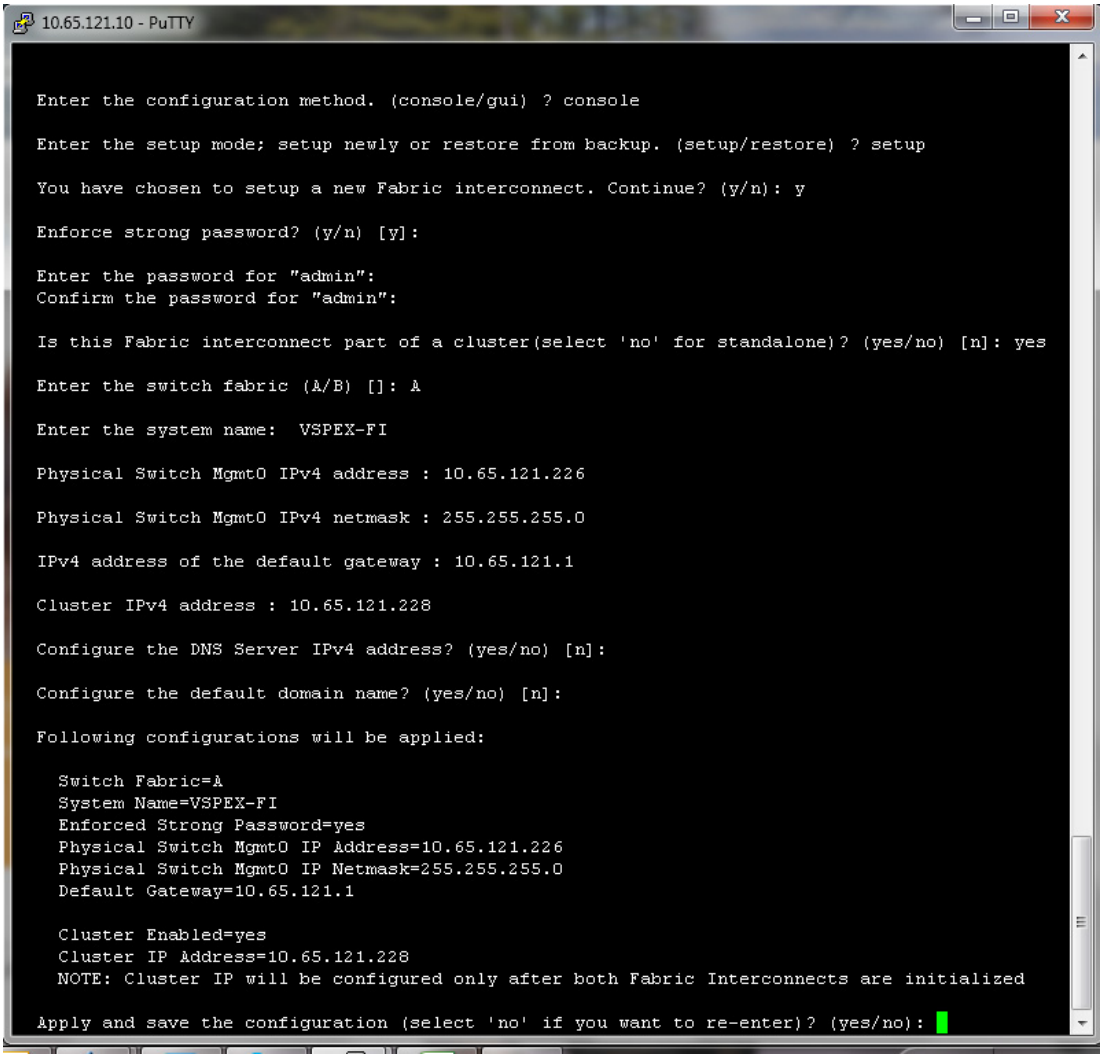

2. Now disconnect the RJ-45 serial console from the FI that you just configured and attach it to the other FI. The other FI would detect that its peer has been configured, and would prompt you to just join the cluster. The only information you need to provide is the FI specific management IP address,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

Figure 34 Configuring Secondary Cisco UCS Fabric Interconnect

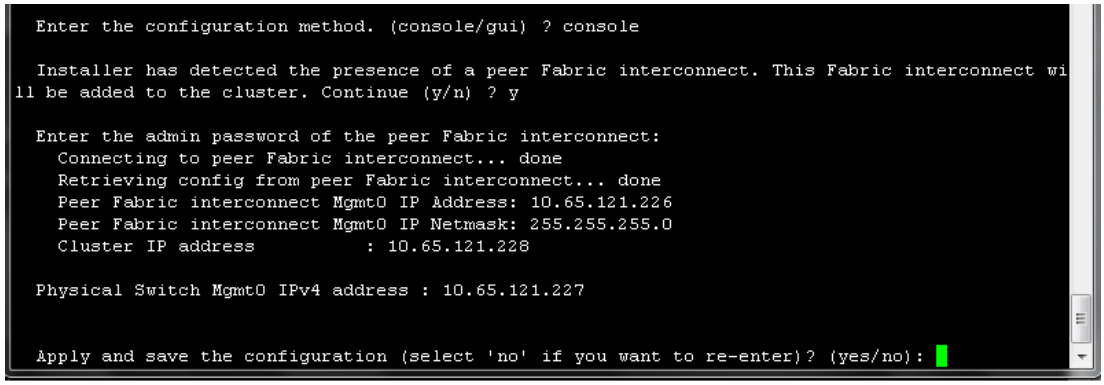

3. Once initial configurations on both FIs are completed, you can disconnect the serial console cable. Now, UCS Manager would be accessible through web interface (https://<ucsm-virtual-ip>/) or SSH. Connect to UCS Manager using SSH, and see HA status. As there is no common device connected between two FIs (a rack server or blade server chassis), the status would say "HA NOT READY", but you must see both FI A and FI B in "Up" state as shown in Figure 35.

Figure 35 Status of the Configured Cisco Fabric Interconnects

Configuration for Server Discovery

All FI Ethernet ports are Unconfigured and shut down by default. You need to classify ports as server facing ports, directly attached storage array facing ports and uplink ports. Next steps show how to configure ports for the proper server auto-discovery:

1. Firstly, we need to configure chassis discovery policy that specifies server side connectivity. Using a web browser, access the Cisco UCS Manager using the management virtual IP address and download the Java applet to launch UCS Manager GUI. In the home page, click **Equipment** tab on the left pane and then **Policies** tab on the right pane. In the Chassis Discovery Policy area, select 2 Link for Action field, as two 10 GE links are connected between Cisco UCS FI and FEX per fabric. Also, click the **Port Channel** radio button for Link Grouping Preference, for better utilization of bandwidth and link level high-availability as shown in Figure 36. Click **Save changes**.

Figure 36 Configuring Chassis Discovery Policy

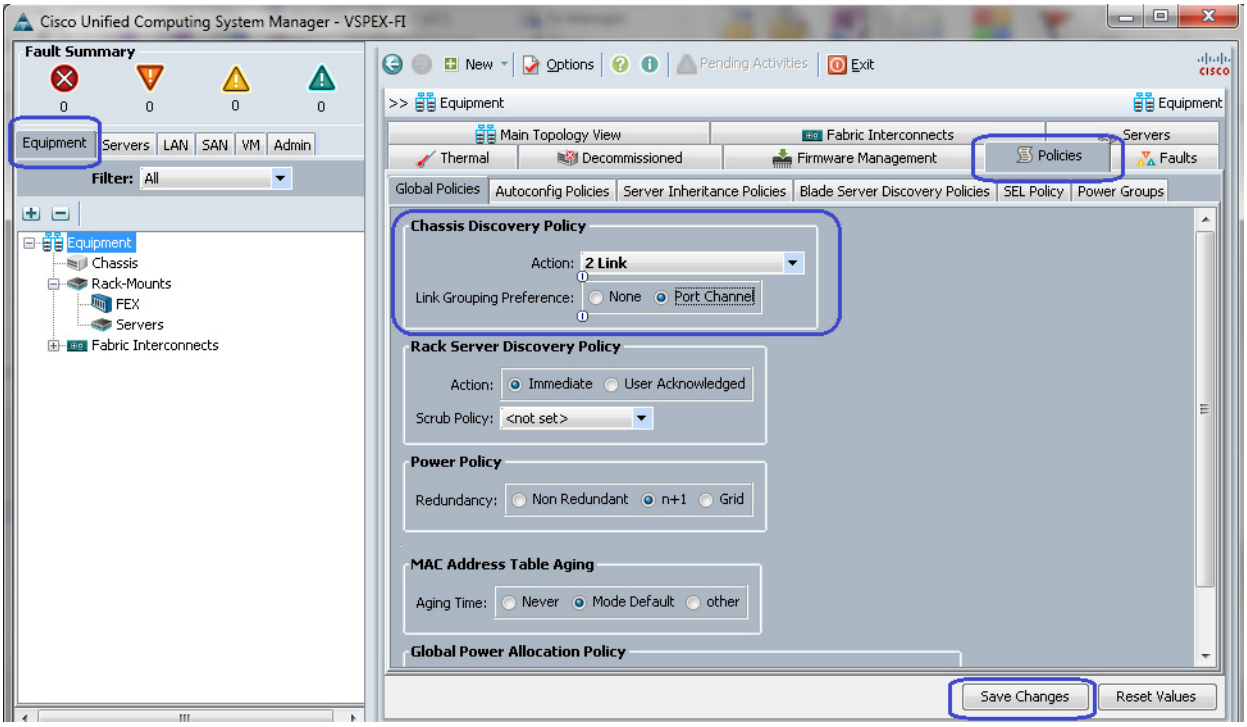

2. Next, we need to identify ports connected to the Chassis or FEX per FI basis. Click **Equipment** tab, expand Cisco UCS Fabric Interconnects, choose a particular FI, for example, Fabric Interconnect A, choose Unconfigured Ethernet Ports. A list of unconfigured Ethernet ports appears in the right pane of the window. Select the two ports connected to FEX-A, right-click on each of them and click **Configure as Server Port** as shown in Fig. Click **Yes** on the confirmation message window.

Figure 37 Configuring Ethernet Ports as Server Ports

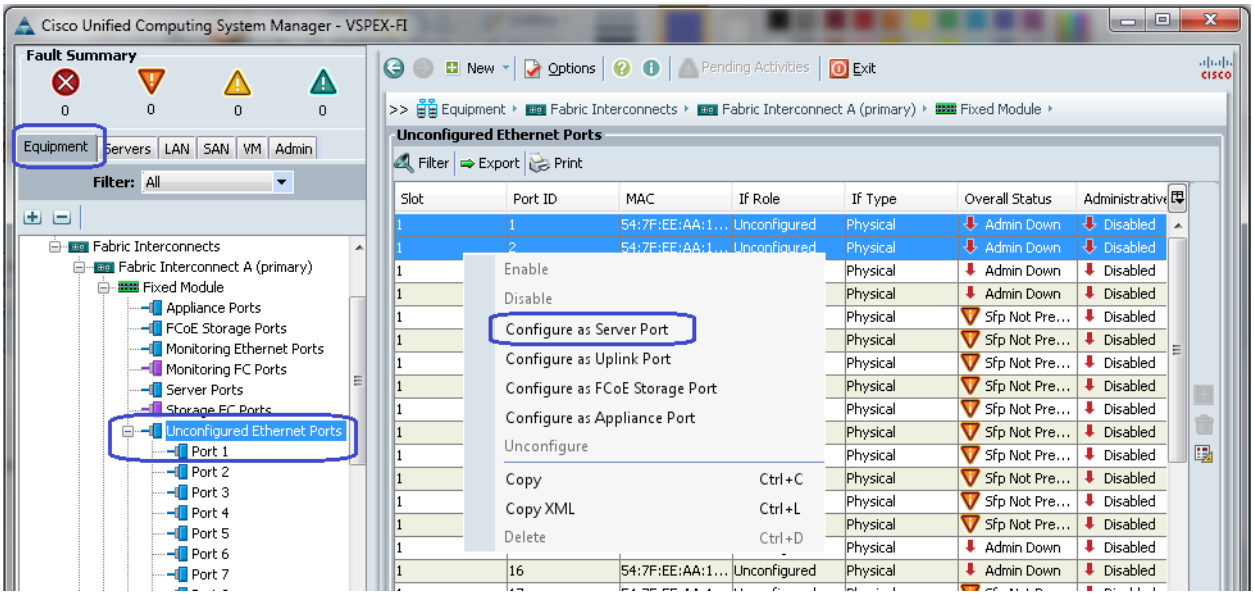

3. Repeat step 2 for the other FI.
4. Once server ports are configured on both FIs, the Chassis or FEX auto-discovery will start. In case of FEX, after the deep discovery of FEX is complete, you will be able to see two Fabric Extenders in the Equipment tab with overall status shown as Operable.

Figure 38 Status Details of Discovered Cisco UCS FEX

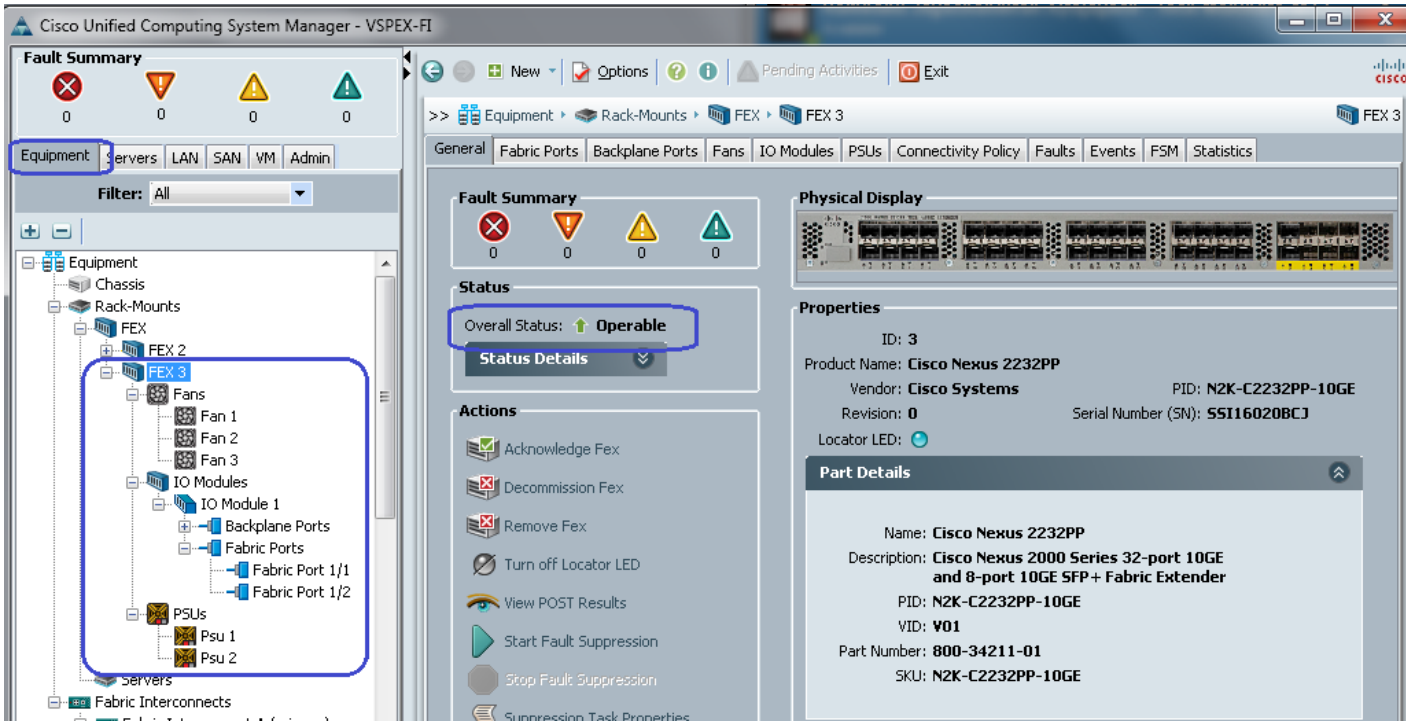

Similarly, if server ports are connected to chassis, you would see the chassis fully discovered, with all its IOMs, fans, power supplies and so on.

Figure 39 Status Details of Discovered Cisco UCS Blade Chassis

5. After the Chassis and FEX auto-discovery, the Blade Server and Rack Server auto-discovery will also start respectively. As and when the servers are discovered, you will be able to see them in the Equipment tab with over-all status shown as Unassociated and availability state shown as Available, and discovery state shown as Complete.

Figure 40 Status Details of Discovered Cisco UCS Rack Mount Server

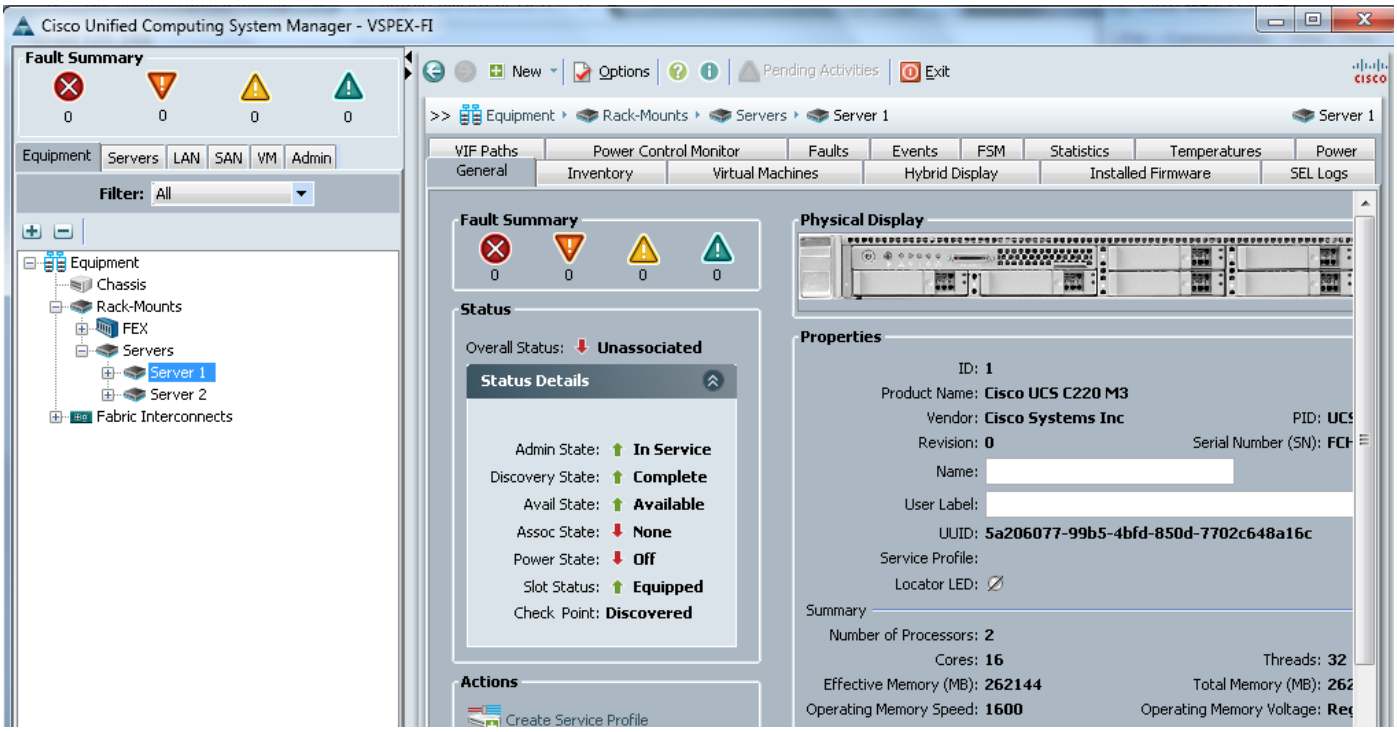

Similarly, a Blade Server's status can be observed as Figure 41:

Figure 41 Status Details of Discovered Cisco UCS Blade Server

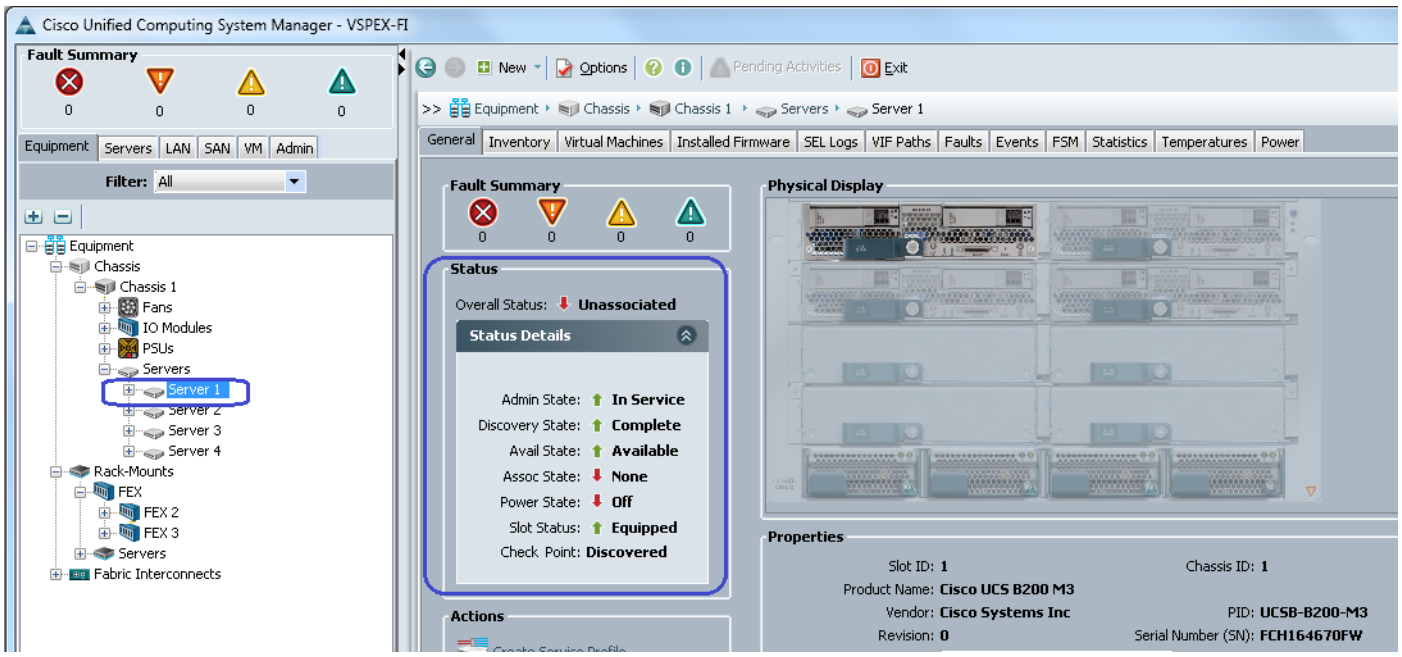

6. Once all the servers are discovered, to see summary of all the servers, choose **Equipment** tab > **Rack-Mounts > Servers.**

Figure 42 Summary of the Discovered Rack Mount Servers

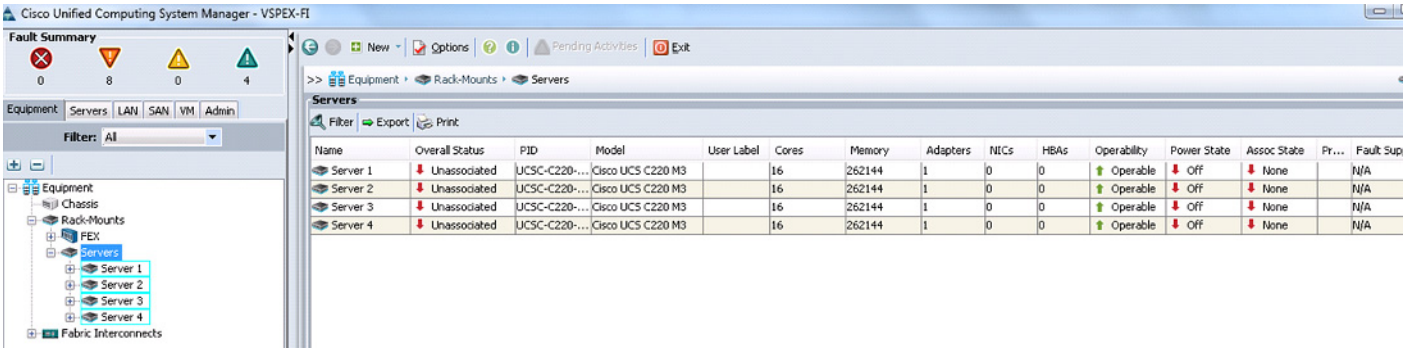

Or, in case of Blade Servers, from the **Equipment** tab, under **Equipment > Chassis > Chassis <id> > Servers.**

Figure 43 Summary of the Discovered Blade Servers

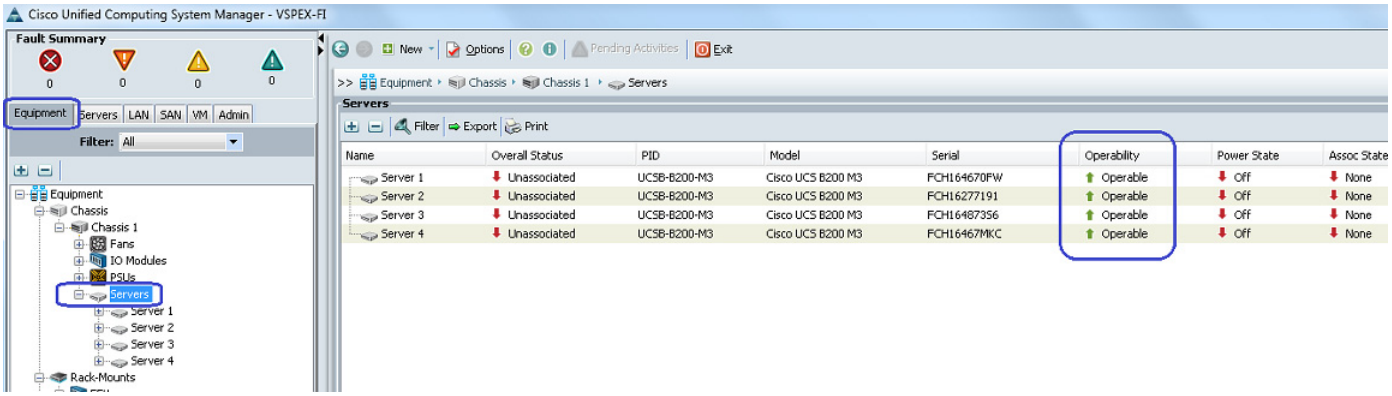

Upstream / global Network Configuration

This section provides a few upstream/global network configuration steps:

1. Move to FC switching mode (FC-variant only)
2. Uplink VLAN configuration
3. Appliance VLAN configuration (iSCSI-variant only)
4. Appliance VSAN configuration (FC-variant only)
5. Configure uplink ports
6. Configure FC appliance ports (FC-variant only)
7. Configure FC Zoning policies (FC-variant only)
8. Configure Ethernet appliance ports (iSCSI-variant only)
9. Configure QoS classes and QoS policy for jumbo MTU

To configure these upstream/global network, follow these steps:

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

Figure 44 Setting FC Switching Mode

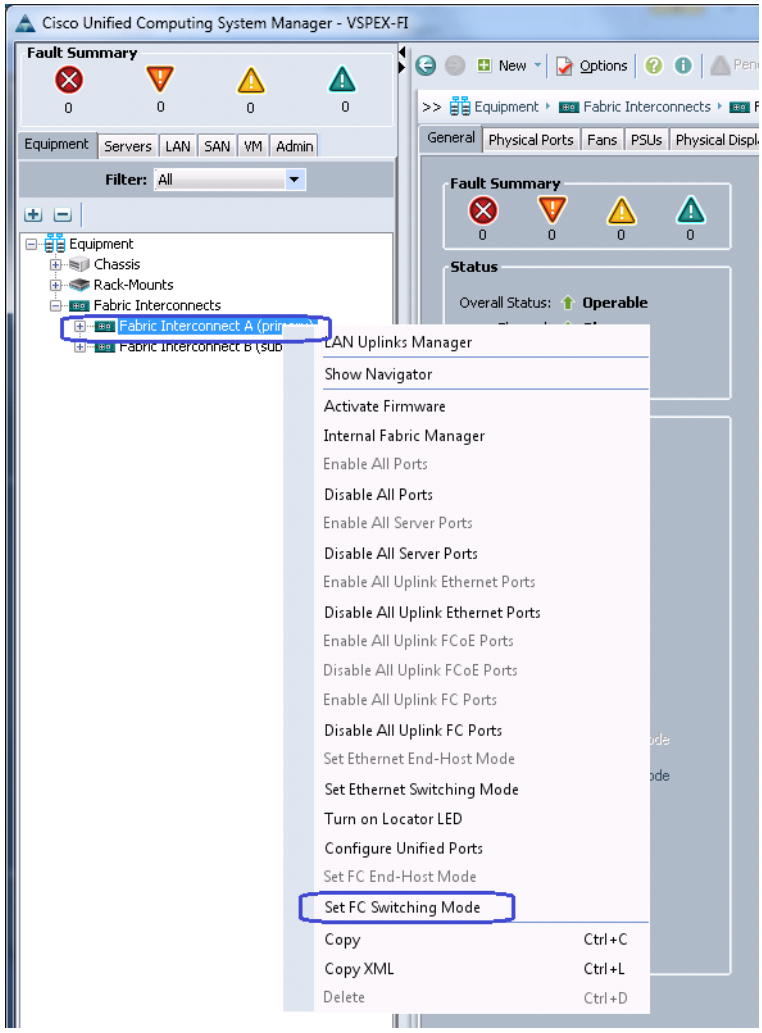

2. (FC-variant only) A message window appears with a warning that the Fabric Interconnects will be restarted as a result of this action. Click **Yes** for both the FIs to reboot (first the secondary FI and then the primary FI). This action is traffic disruptive, so make sure that you perform this operation in 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 select Create VLANs.

Figure 45 Creating VLANs

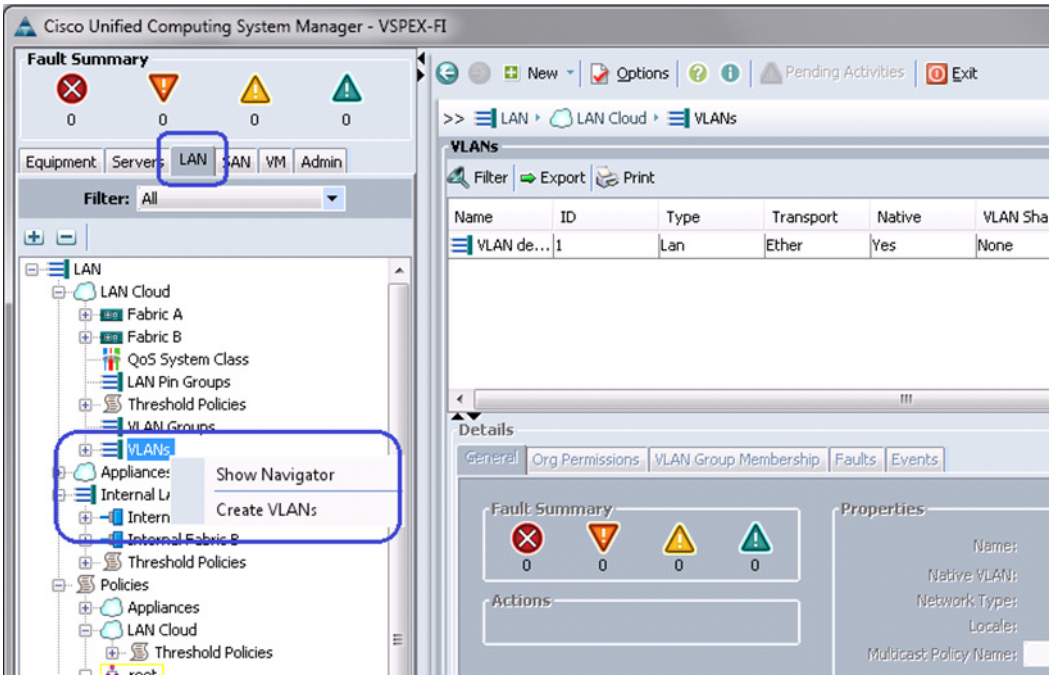

4. Enter the VLAN name in the Name field. Keep the radio button **Common/Global** selected. And Assign ID in the VLAN IDs field.

Figure 46 Creating VLAN for Management

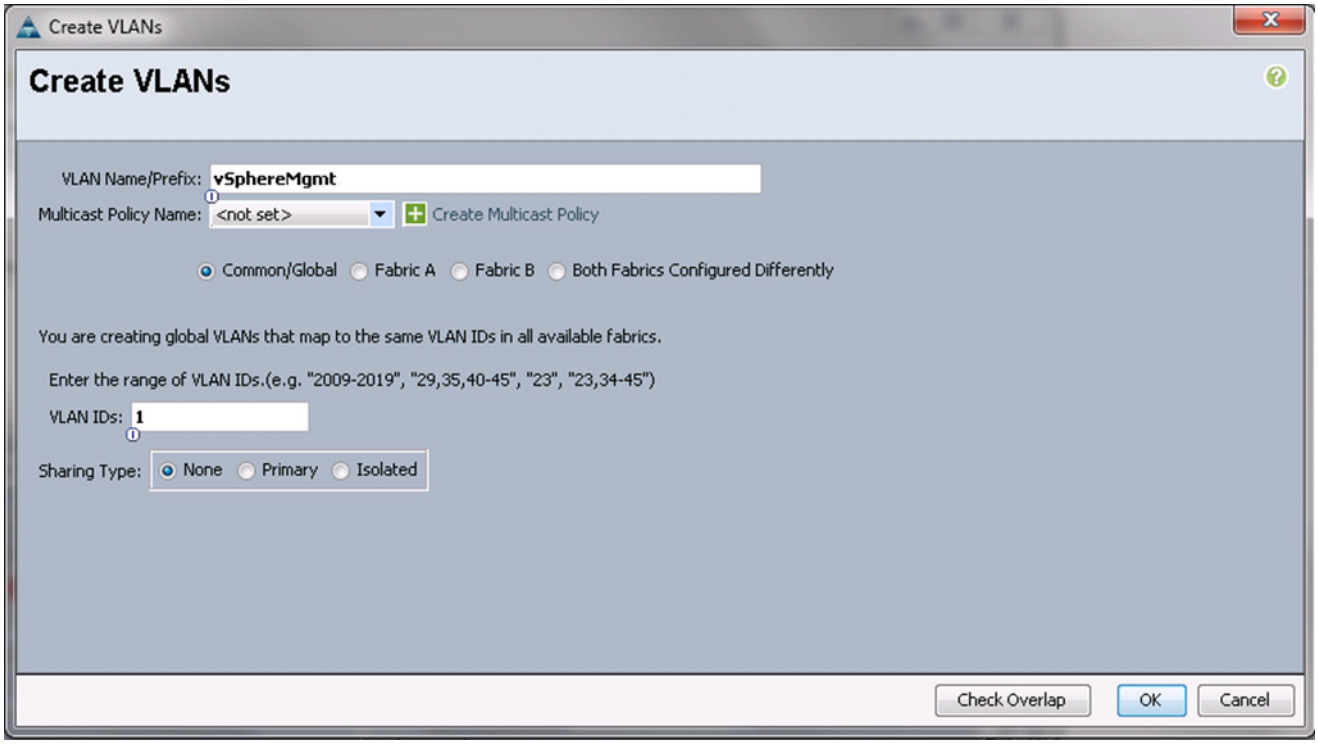

5. Click **OK** to deploy the VLAN. Repeat these steps for vSphereMgmt, VM-Data and vMotion VLANs. See Customer Configuration Data Sheet, page 192 for the VLAN values.
6. (iSCSI-variant only) Storage VLAN is deployed differently, it has different VLAN IDs on each fabric to meet the iSCSI storage access best practices. Click the radio button **Both Fabrics Configured Differently**, and assign VLAN ID for each fabric.

Figure 47 Creating VLAN for Storage

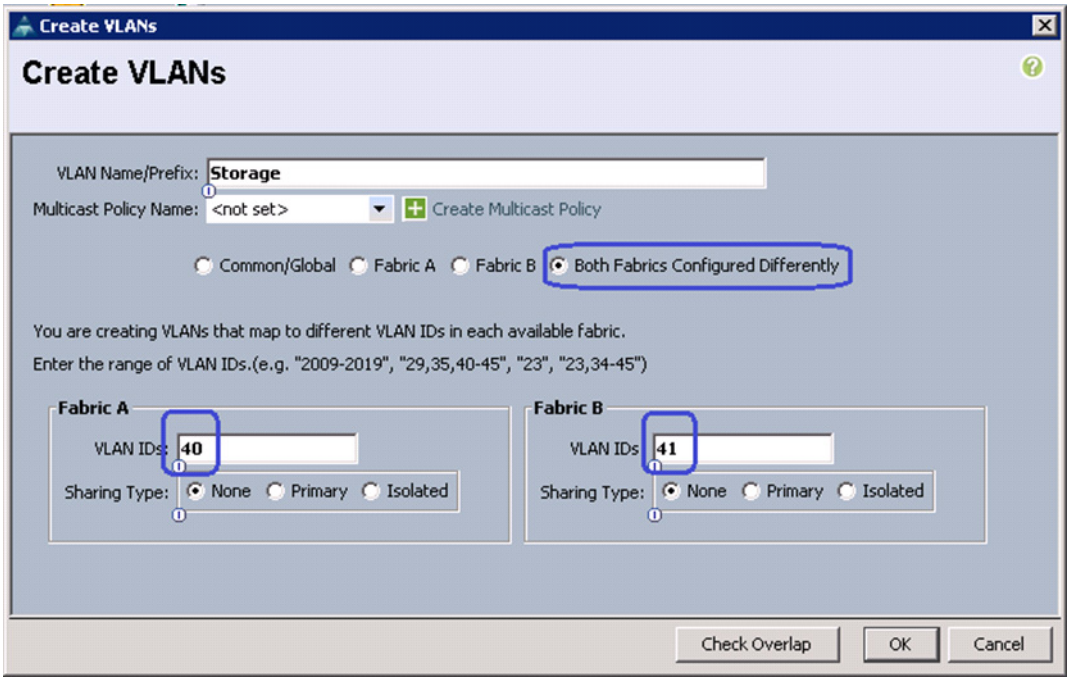

7. (iSCSI-variant only) After configuring uplink or global VLANs, next step is to create Appliance VLANs. From the LAN tab, expand **LAN > Appliances**, right-click on VLANs and select Create VLANs.

Figure 48 Creating VLANs for Appliance

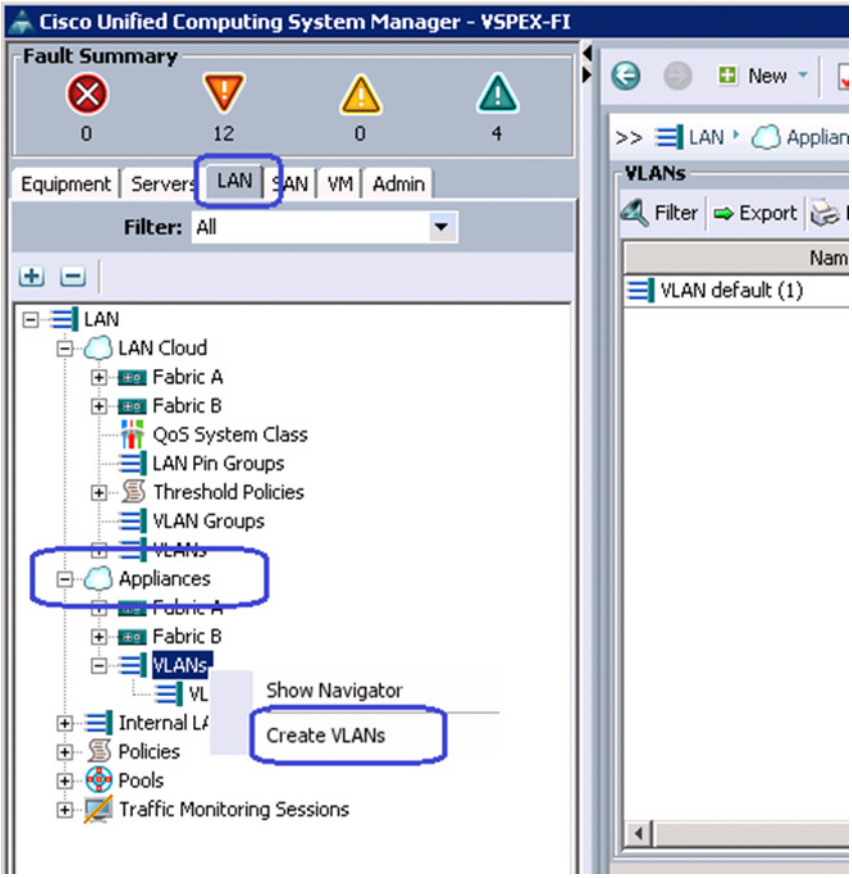

8. (iSCSI-variant only) Enter Storage for VLAN Name field and click **Both Fabrics Configured Differently** radio button, and provide the VLAN IDs for each fabric. Click **OK** to deploy the configuration.

Figure 49 Creating Storage VLAN for Appliances

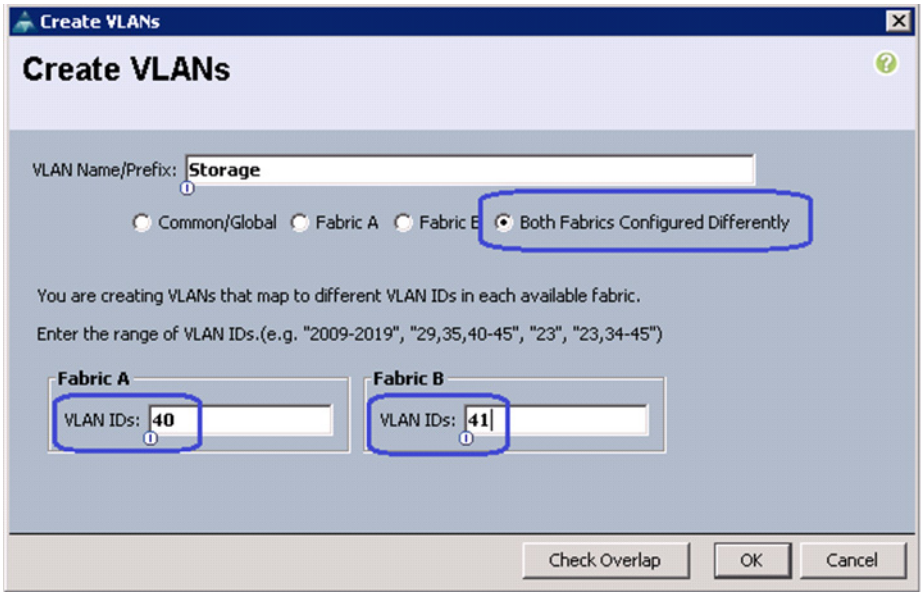

9. (FC-variant only) Click the **SAN** tab. Expand Storage Cloud and right-click on VSANs. Select Create Storage VSAN.

Figure 50 Creating Storage VSAN

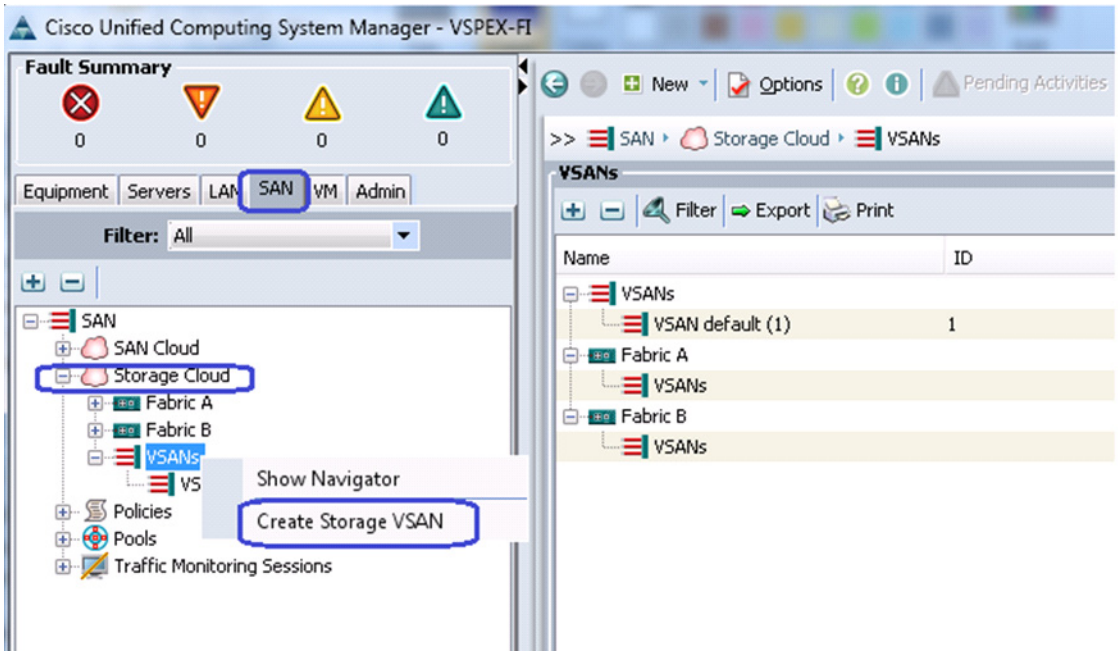

10. Enter the VSAN name in the Name field. Click the radio button **Enable** for FC Zoning and assign VSAN ID and its corresponding FCoE VLAN ID. FCoE VLAN ID should not have conflict with any of the VLANs configured before.

Figure 51 Create Storage VSAN

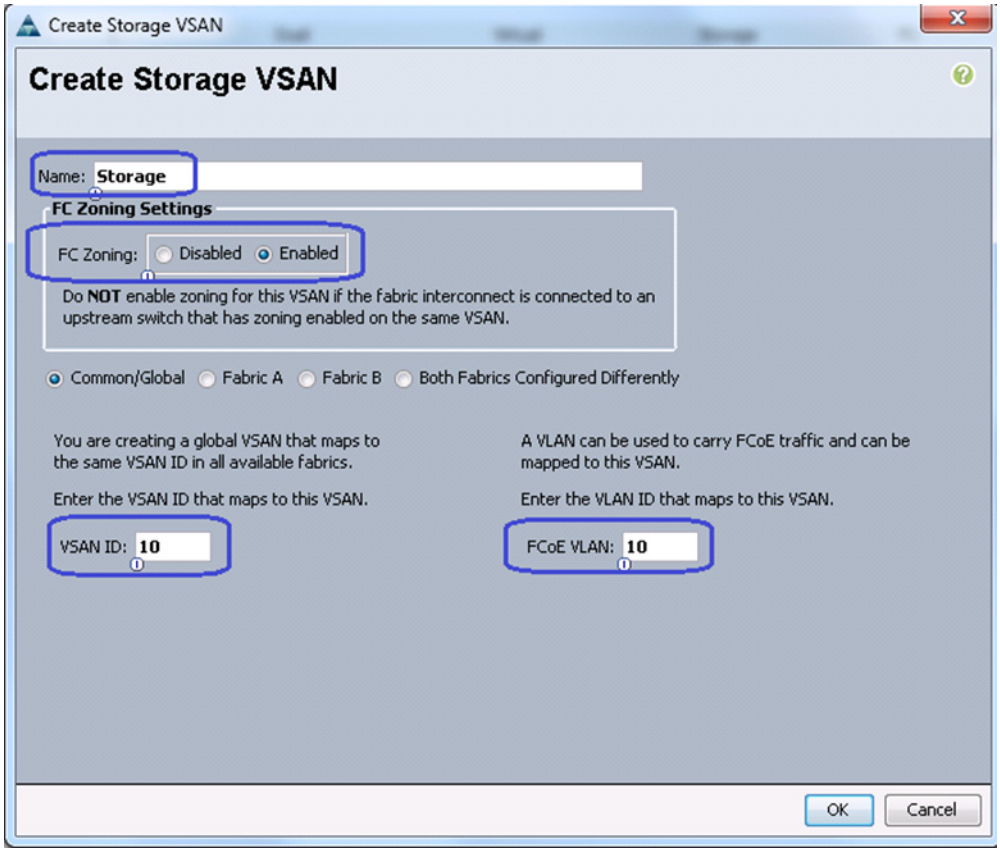

11. To configure Uplink ports connected to the infrastructure network, click **Equipment** tab. Expand Fabric Interconnects, choose a particular FI, expand Expansion Module 2 (this may vary depending on which port you have chosen as uplink port), right-click on the Ethernet port, and select Configure as Uplink Port. Repeat this step for all the uplink ports on each FI.

Figure 52 Configuring Ethernet Ports as Uplink Ports

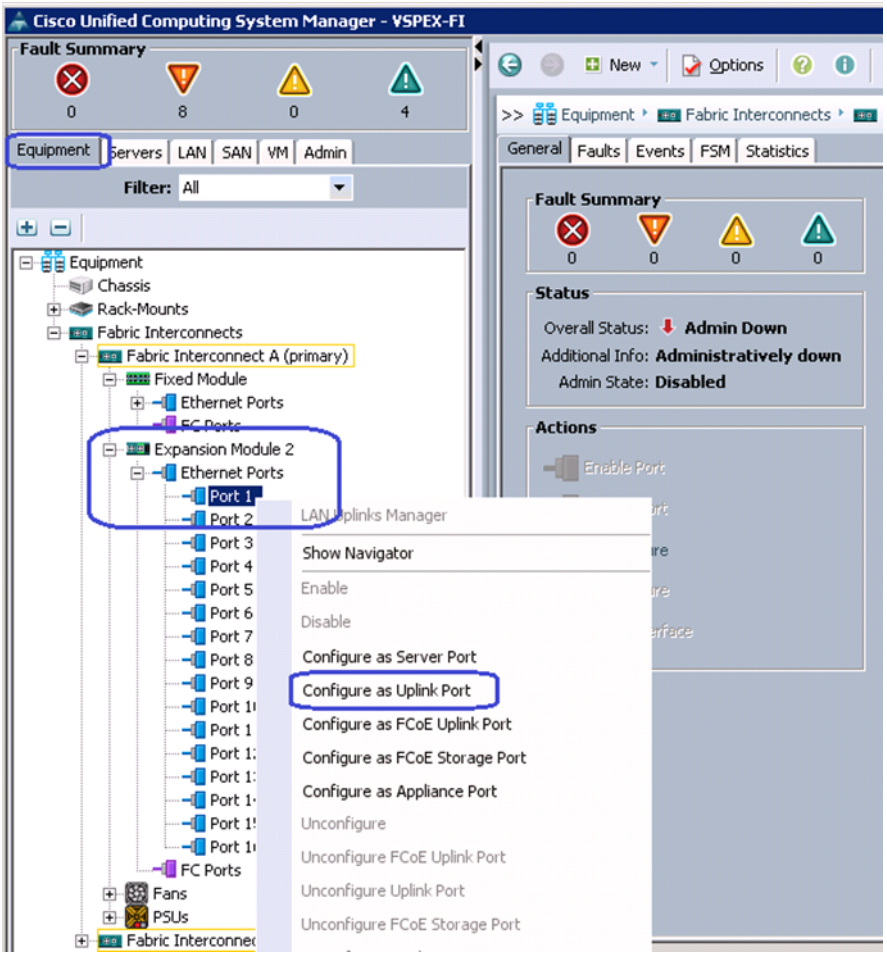

12. (FC-variant only) Cisco UCS 6248UP Fabric Interconnects have Universal Ports. The physical ports are 10G Ethernet ports by default, but can be converted in to Fibre Channel ports. For the FC-variant of the architecture, we need FC connectivity to the EMC VNX storage array. For that, some of the ports need to be converted into FC ports. We will convert ports from expansion module into FC ports. For that, click the **Equipment** tab. Expand Fabric Interconnects and select Fabric Interconnect A. In the right pane of the window, in the Actions area click **Configure Unified Ports.** Click **Yes** in the warning message window.

Figure 53 Configuring Unified Ports

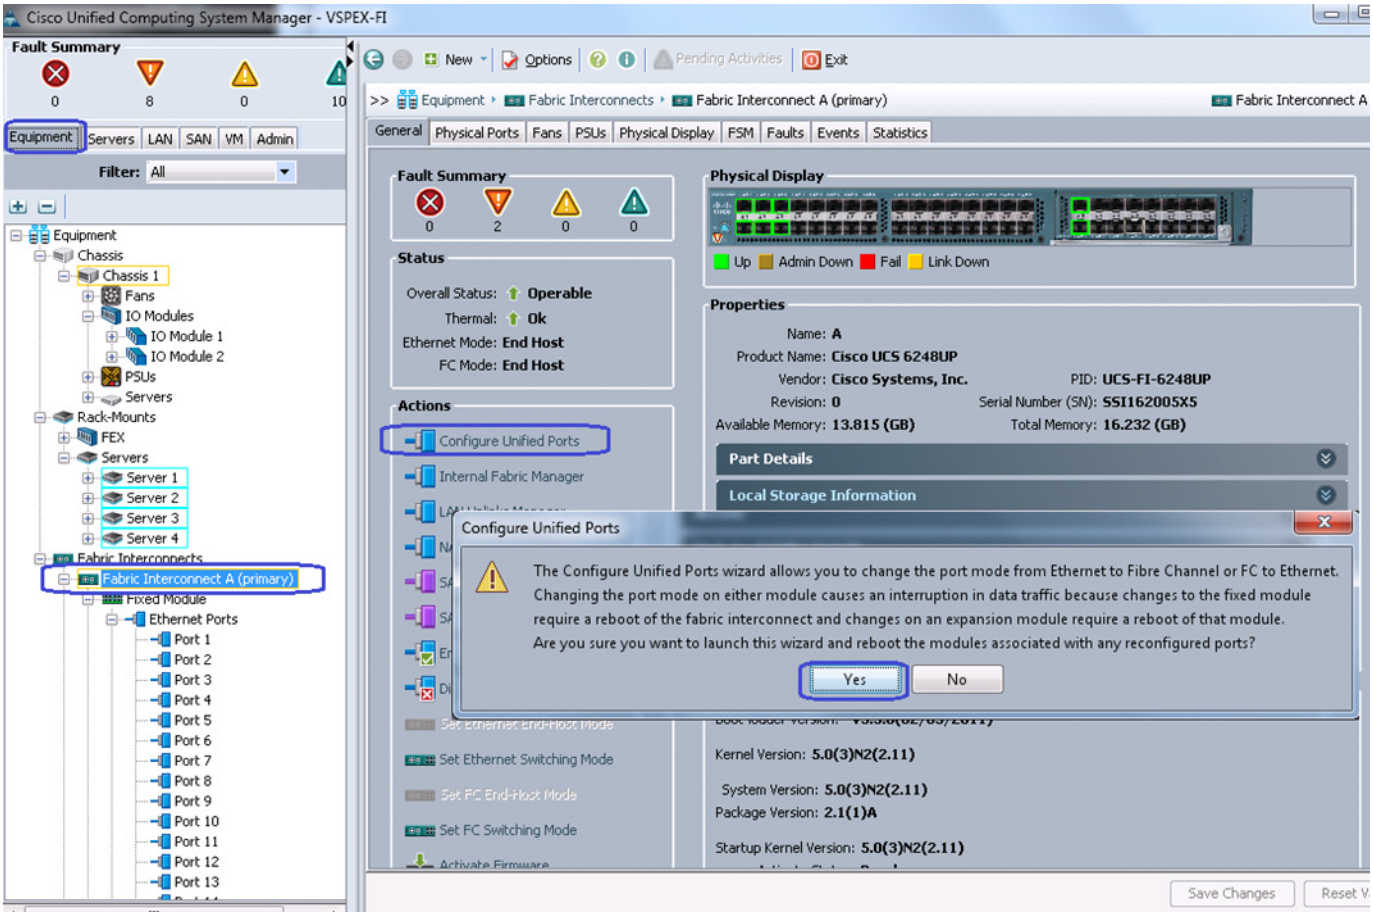

13. (FC-variant only) In the Configure Unified Ports window, click **Configure Expansion Module Ports**.

Figure 54 Configuring Expansion Module Ports

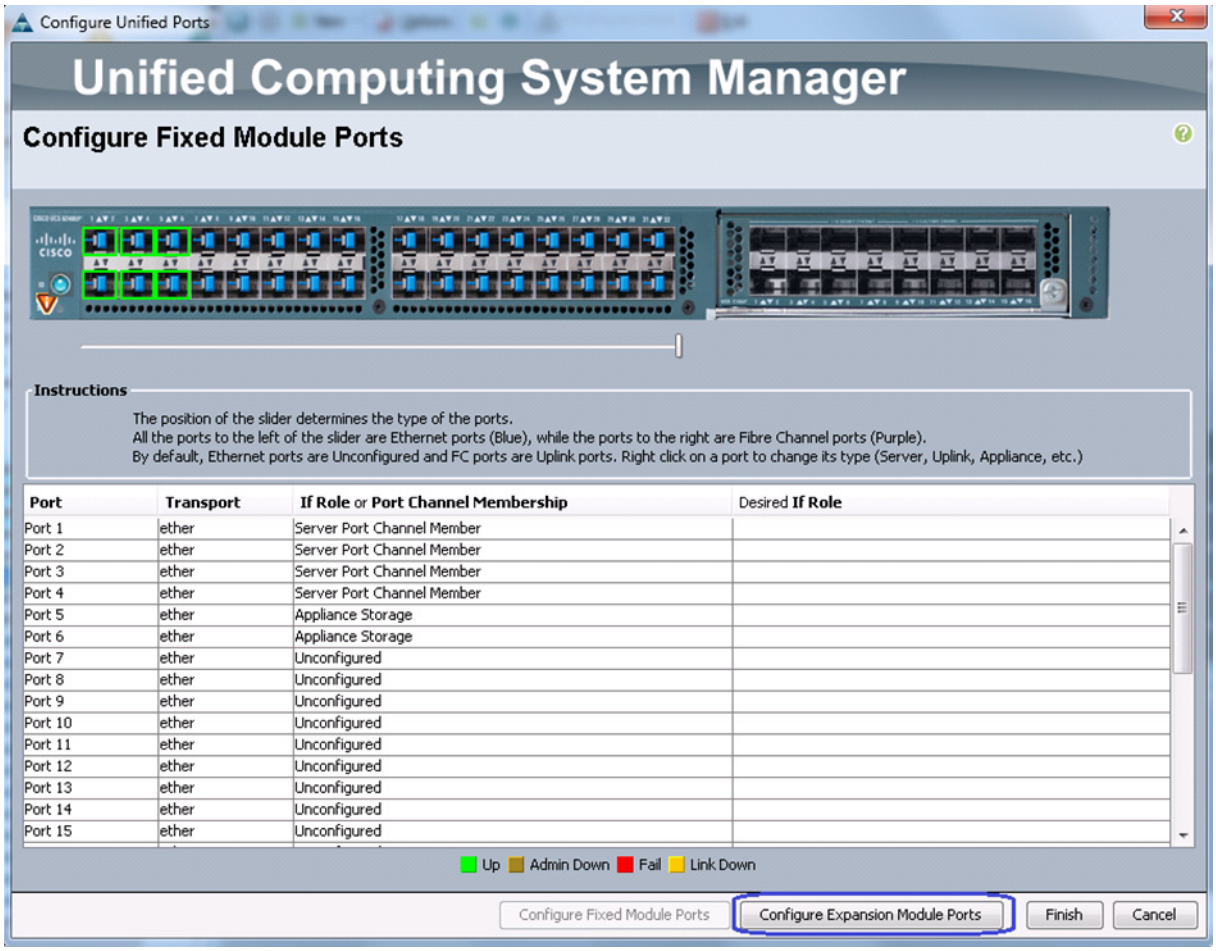

14. (FC-variant only) Select the slider bar and slide it to the middle of the slider bar. Make sure that ports 2/9 to 2/15 are showing as FC Uplink. Click **Finish**. A warning message window appears to warn on rebooting the FI. Click **OK** to reboot the FI.

Figure 55 Window Showing FC Uplink Ports

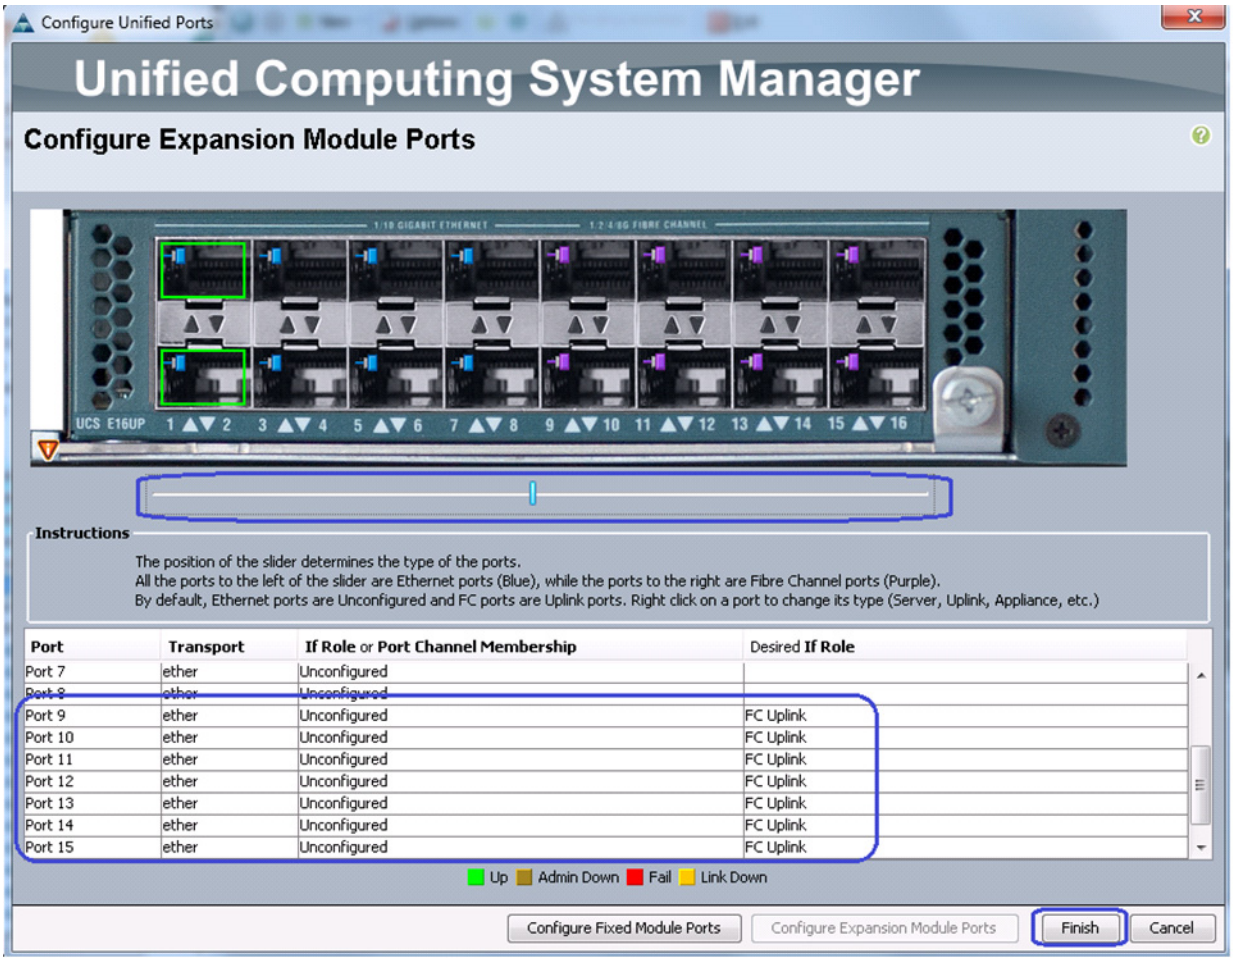

15. (FC-variant only) Once the FI is up after reboot, repeat steps 12, 13 and 14 for FI-B.
16. (FC-variant only) Physical FC ports further need to be classified as FC storage ports for directly attached storage array. Click the **Equipment** tab. Expand **Fabric Interconnect > Fabric Interconnect A > Expansion Module 2 > FC Ports.** Select each FC ports and in Actions area click **Configure as FC Storage port** in the right pane of the window.

Figure 56 Configuring Storage Ports as FC Storage Ports

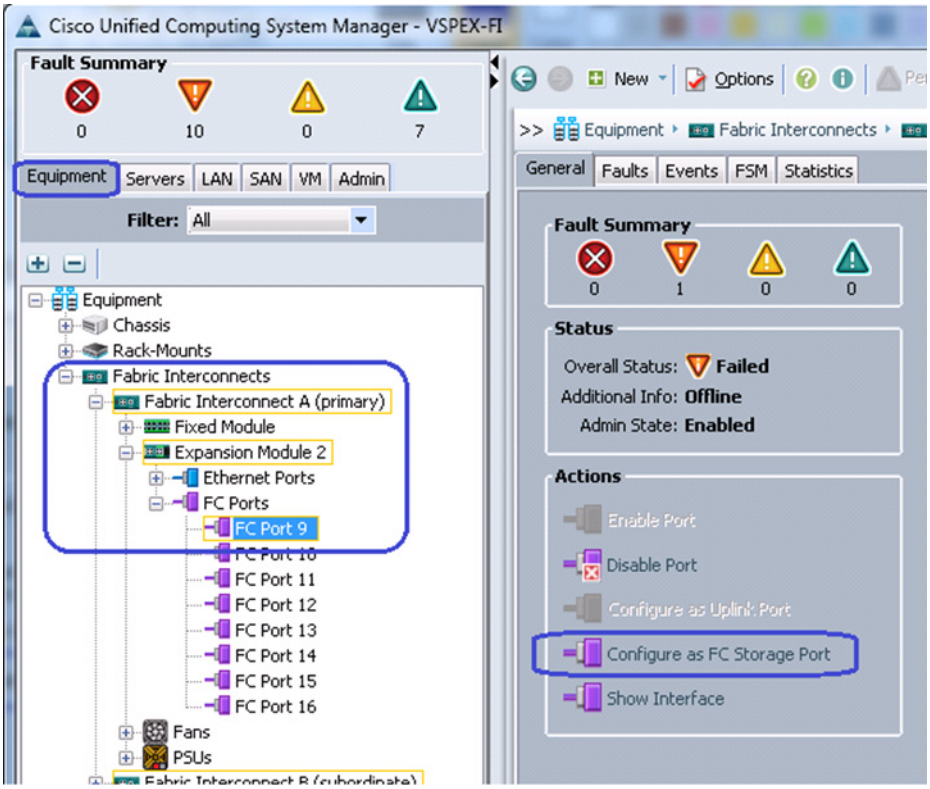

17. (FC-variant only) Make sure that the over all status of the is shown Up. In the Properties area for the VSAN field, select the Storage VSAN configured in steps 9 and 10 from the drop-down list. Click **Save Changes**.

Figure 57 Selecting the Configured Storage VSAN for the FC Port

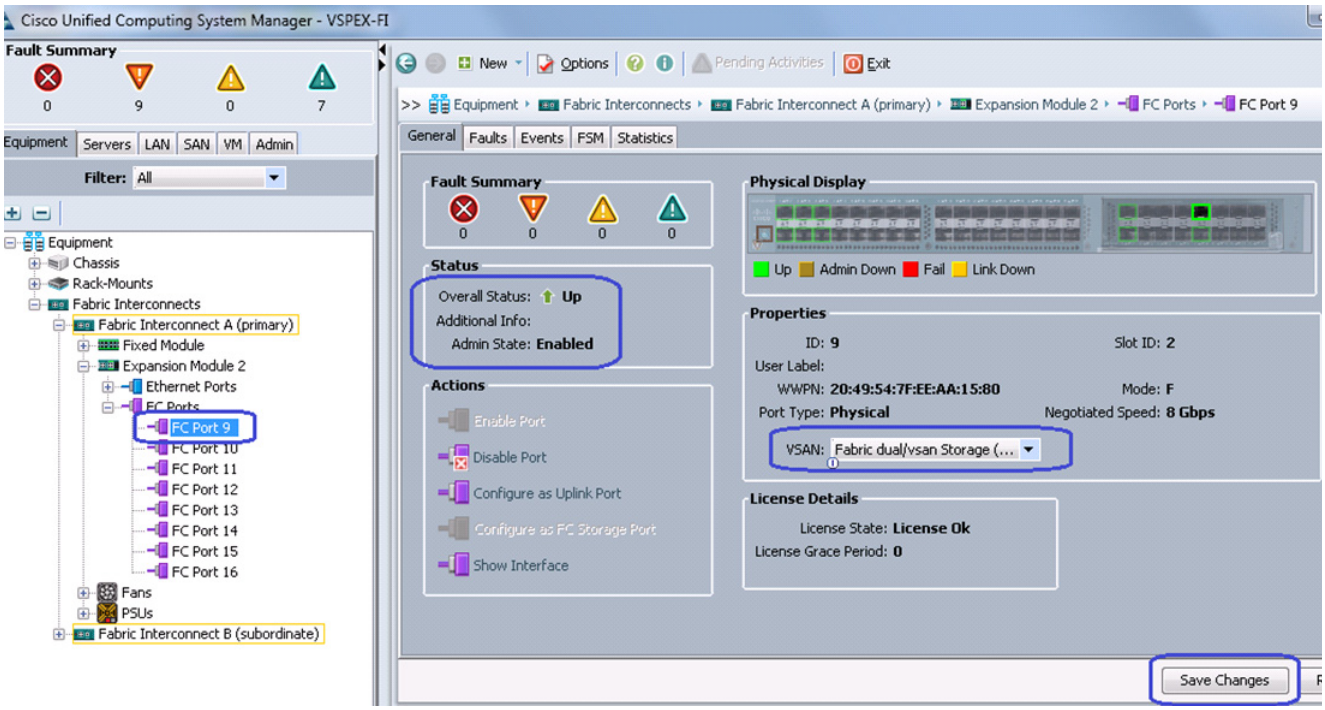

18. (FC-variant only) At this point of time, the EMC VNX storage array will do Fibre Channel flogi into the FIs. Using the WWPN of the VNX storage array, we can carve out the zoning policy on the FI. Use SSH connection to the UCS Manager Virtual IP address, and type **connect nxos a** command. In the read-only NXOS shell, type **show flogi database** command and note down the WWPN of the storage array.

Figure 58 Running flogi in Read-only NXOS Shell

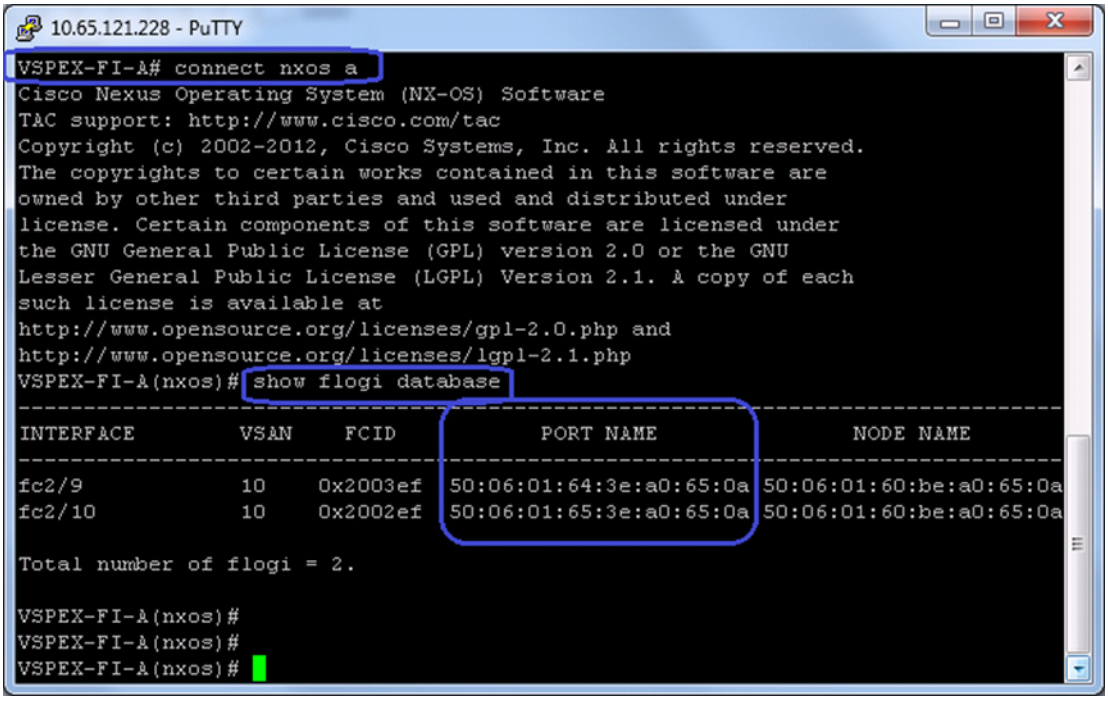

19. (FC-variant only) On UCS Manager GUI, click the **SAN** tab. Expand **SAN > Policies > root**, right-click on Storage connection policies, and choose Create Storage Connection Policy.

Figure 59 Creating Storage Connection Policy

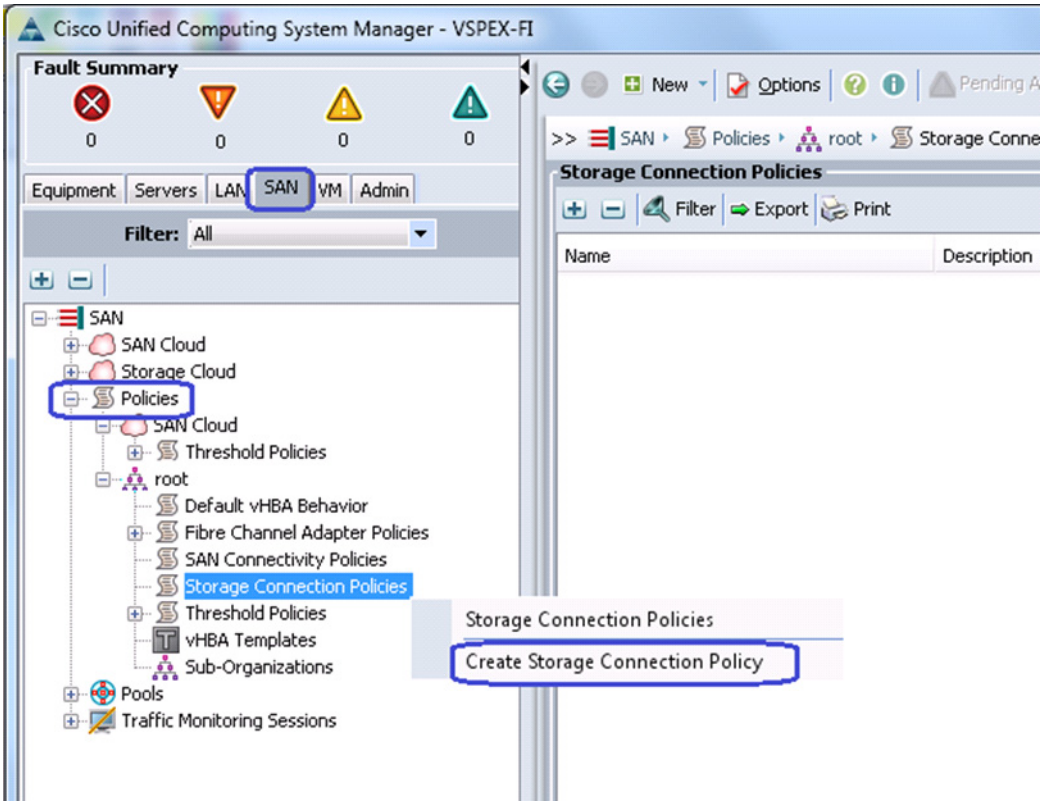

20. (FC-variant only) Enter Fabric-A in the Name field and (optional) description. Click the **Single Initiator Multiple Targets** radio button as the Zoning Type. Click + to add a new FC Target Endpoint.

Figure 60 Adding FC Target Endpoint

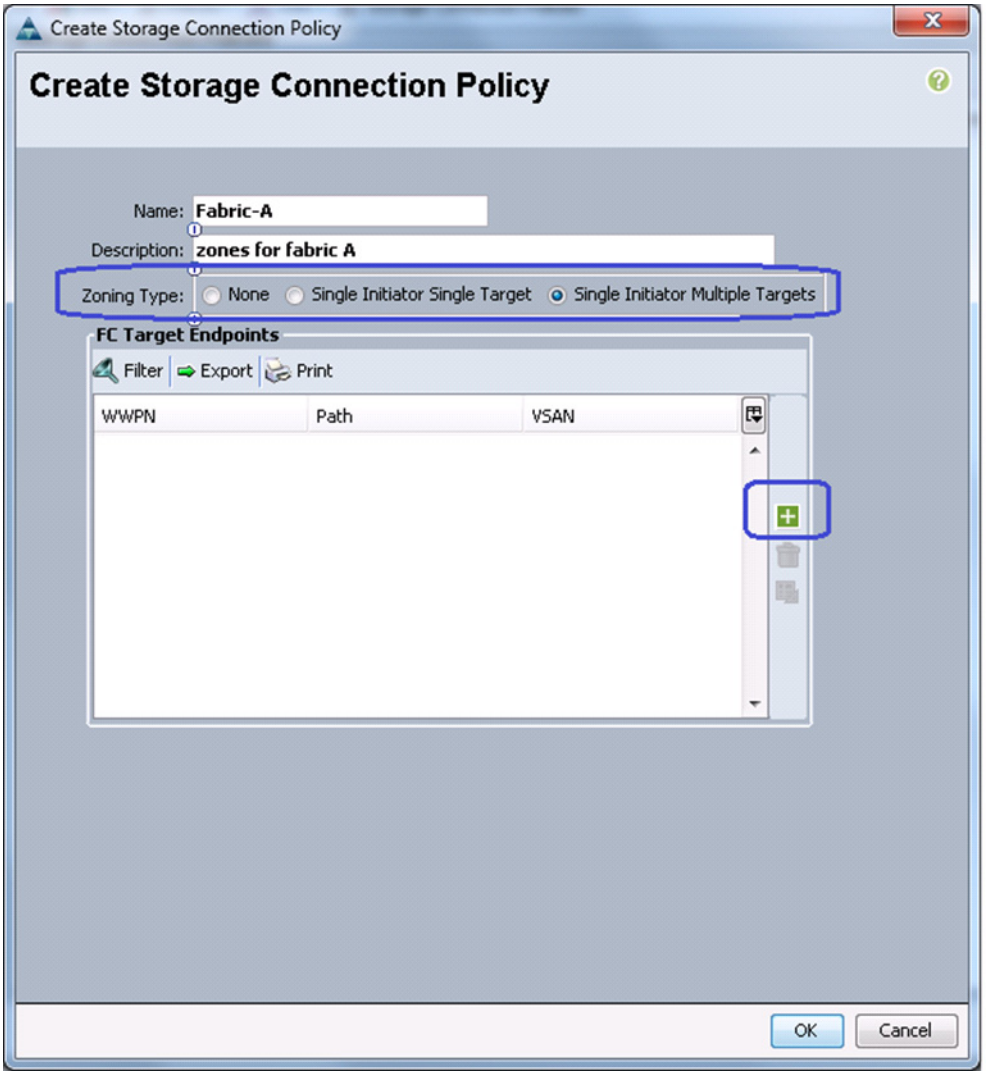

21. (FC-variant only) Copy the WWPN from the "show flogi database" output from step 18 and paste it in the WWPN field. Provide description (optional), click the **A** radio button for Path and choose Storage VSAN from VSAN drop-down list.

Figure 61 Creating FC Target Endpoint

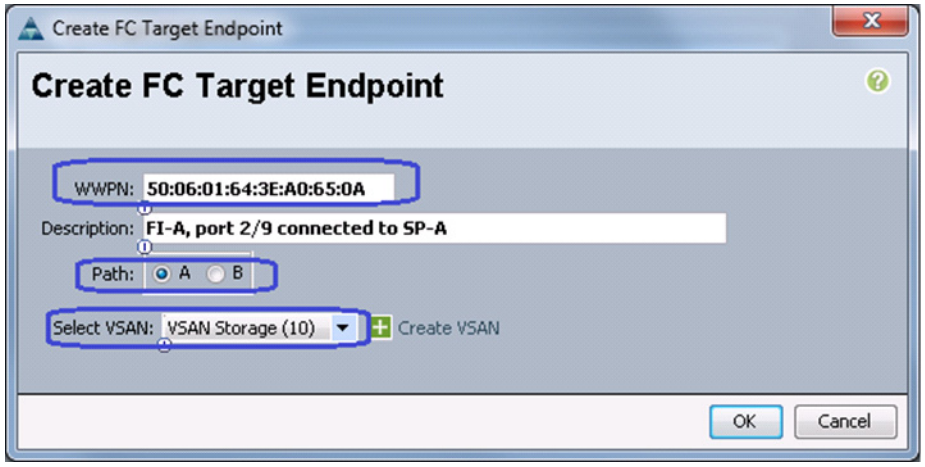

22. (FC-variant only) Similarly, add the second FC target end-point for fabric A and click **OK**.

Figure 62 Creating Second FC Target Endpoint

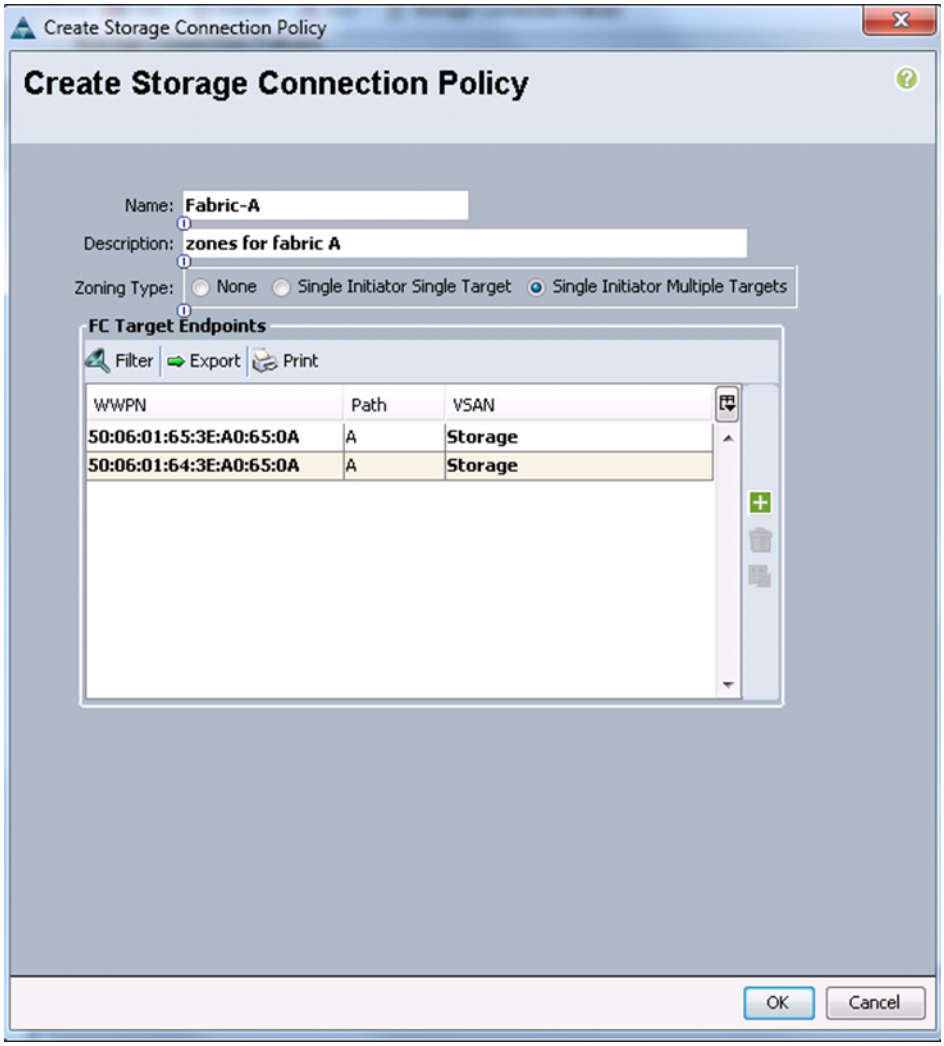

23. (FC-variant only) Repeat steps 18 to 22 for Fabric B. The end result will look similar to Figure 62.

Figure 63 Storage Connection Policies for Fabric A and Fabric B

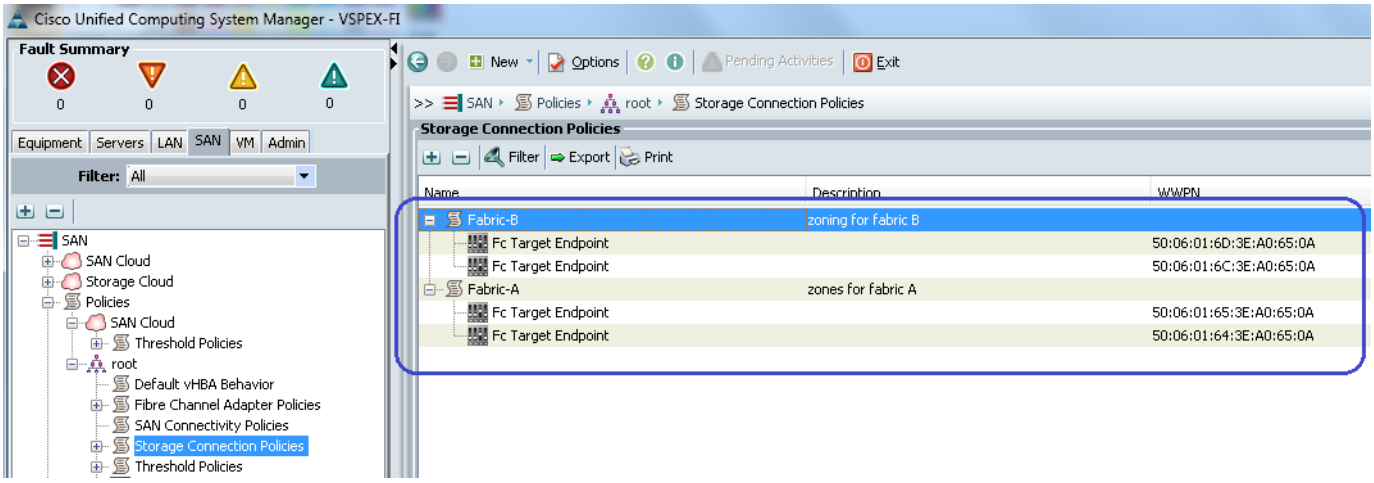

24. (iSCSI-variant only) To mark ports connected to directly attached storage arrays, from the **Equipment** tab, expand **Equipment > Fabric Interconnects**, select a particular FI. Choose the port under Fixed Module, right-click, and choose Configure as Appliance Port.

Figure 64 Configuring Ethernet Port as Appliance Port

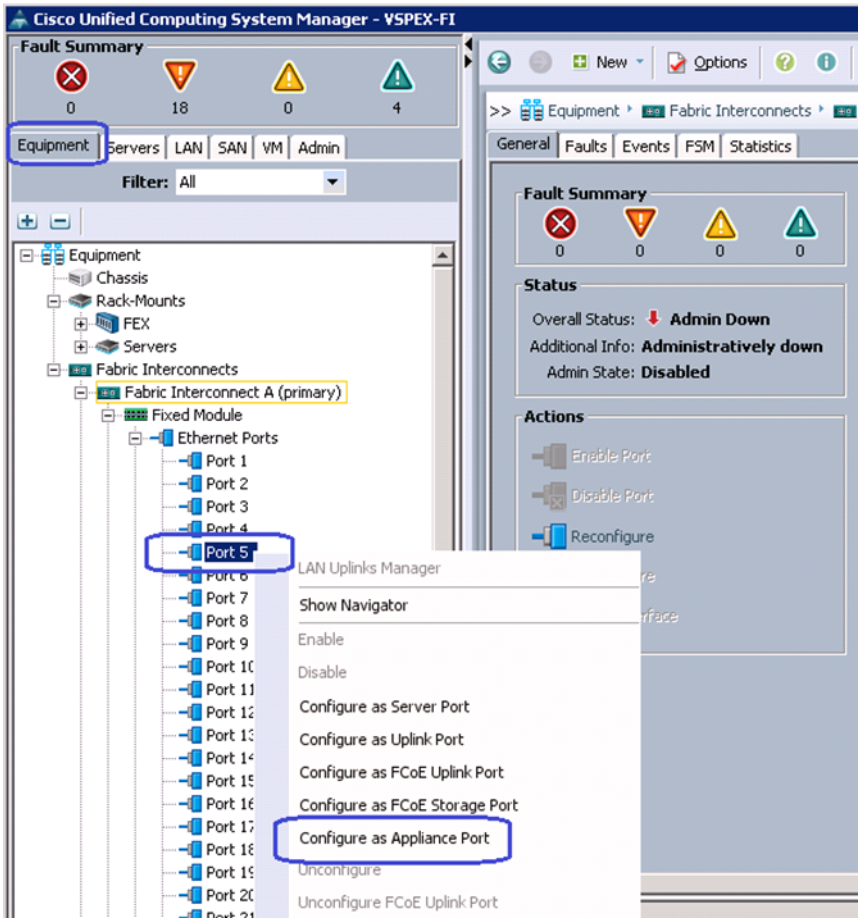

25. (iSCSI-variant only) In the Configure as Appliance Port window, choose Platinum for Priority from the drop-down list (for jumbo MTU configuration of storage access). Click the Access radio button for port-mode in the VLANs area and choose Storage from the drop-down list for Select VLAN field. Click **OK** to deploy the configuration.

Figure 65 Specifying Details in the Configure as Appliance Port Window

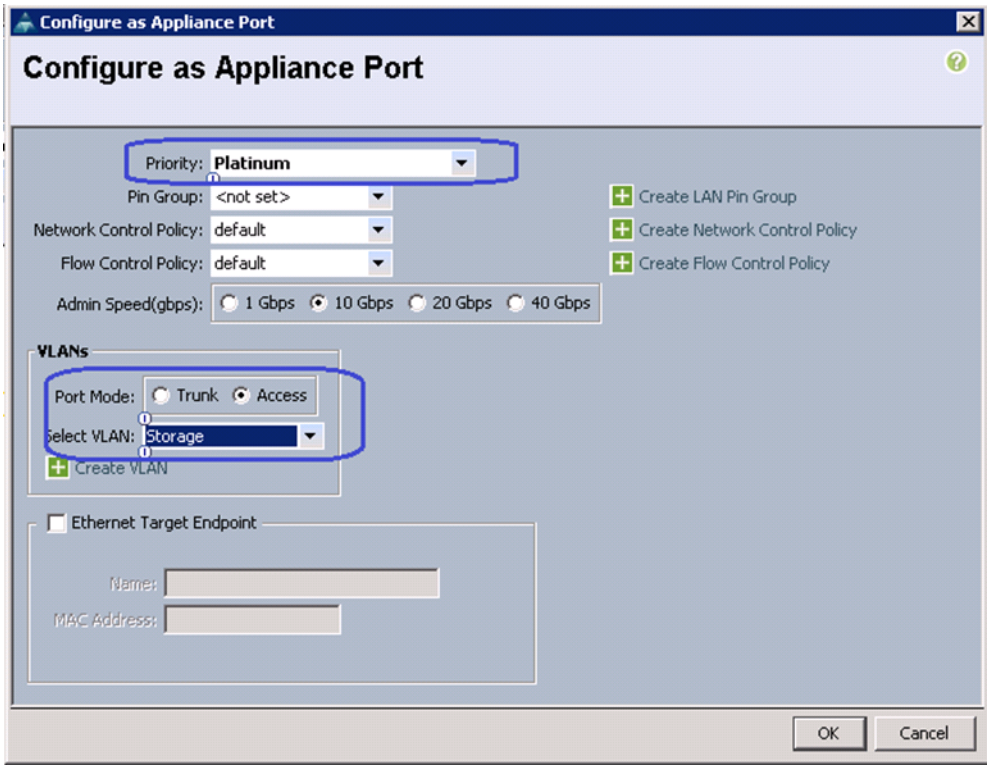

26. (iSCSI-variant only) Repeat these steps for the remaining 3 appliance ports.
27. For QoS configuration, click the **LAN** tab, expand **LAN > LAN Cloud**, and choose QoS System Class. Set the priority as Platinum, and select MTU as 9216 from the drop-down list. Keep other configuration as default and save the configuration.

Figure 66 Configuring QoS System Class

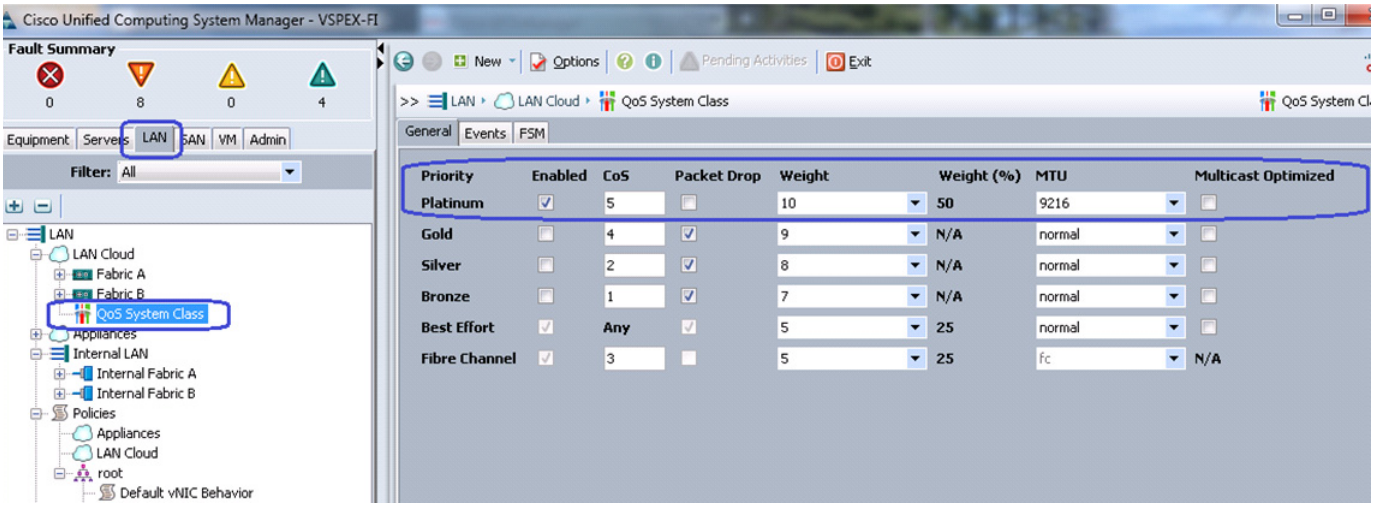

28. From the **LAN** tab, expand **LAN > Policies > root**. Right-click on QoS Policies and choose Create QoS Policy.

Figure 67 Creating QoS Policy

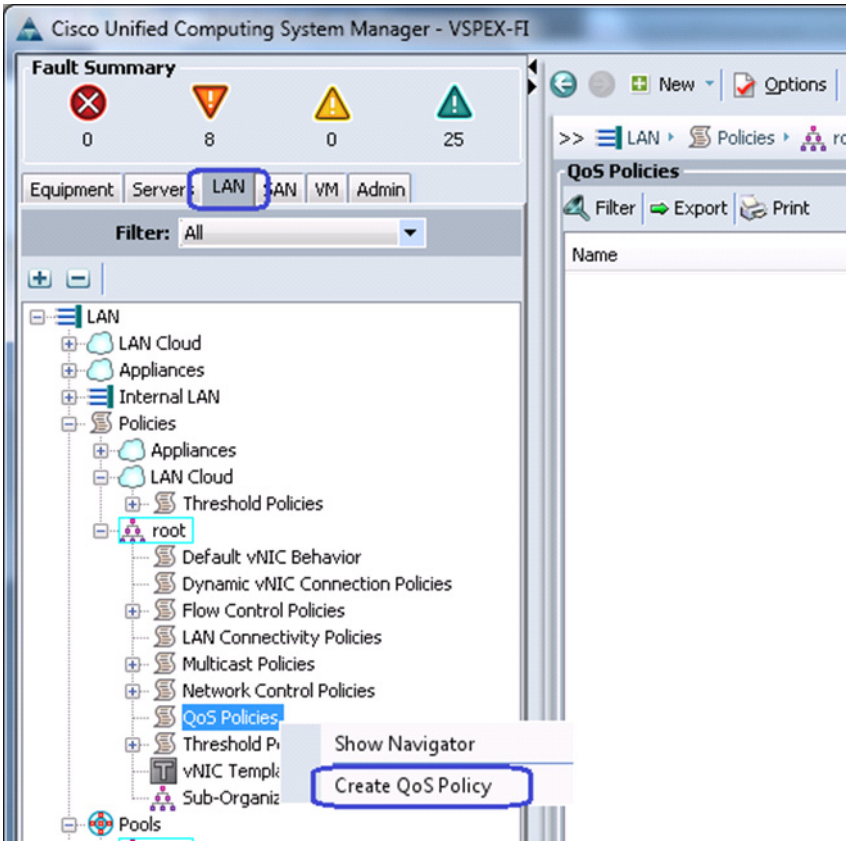

29. Enter the name of the policy as jumboMTU and choose Priority as Platinum from the drop-down list in the Egress area. Click **OK** to save the configuration.

Figure 68 Specifying Details for Creating QoS Policy

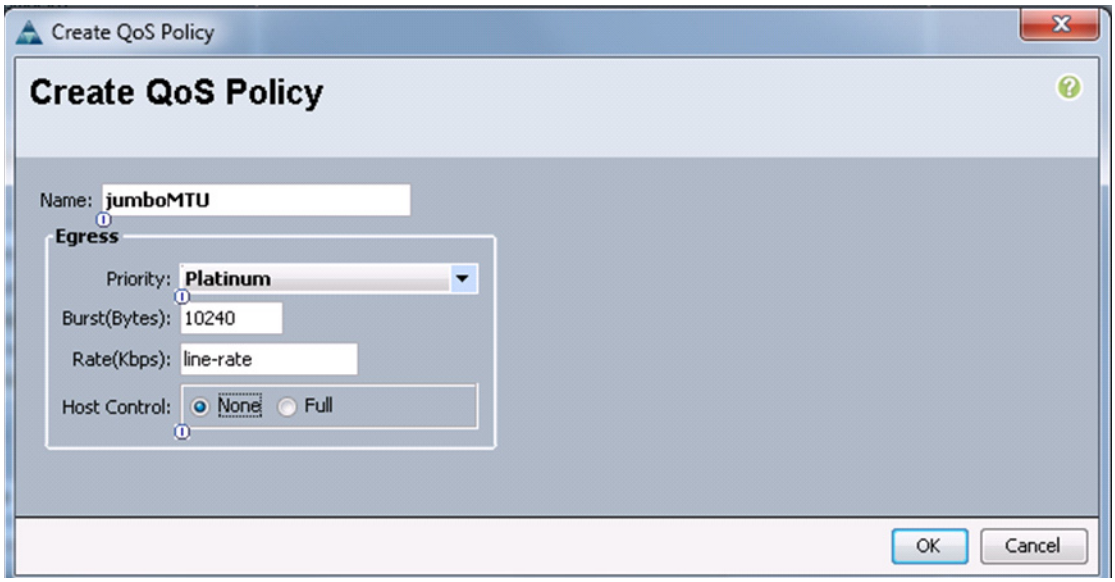

Configure Identifier Pools

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

1. Server UUID pool
2. MAC address pool
3. WWN pool (FC-variant only)
4. IQN pool (iSCSI-variant only)
5. iSCSI initiator IP address pool (iSCSI-variant only)
6. Management IP address pool

To configure these pools, follow these steps:

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

Figure 69 Creating UUID Suffix Pool

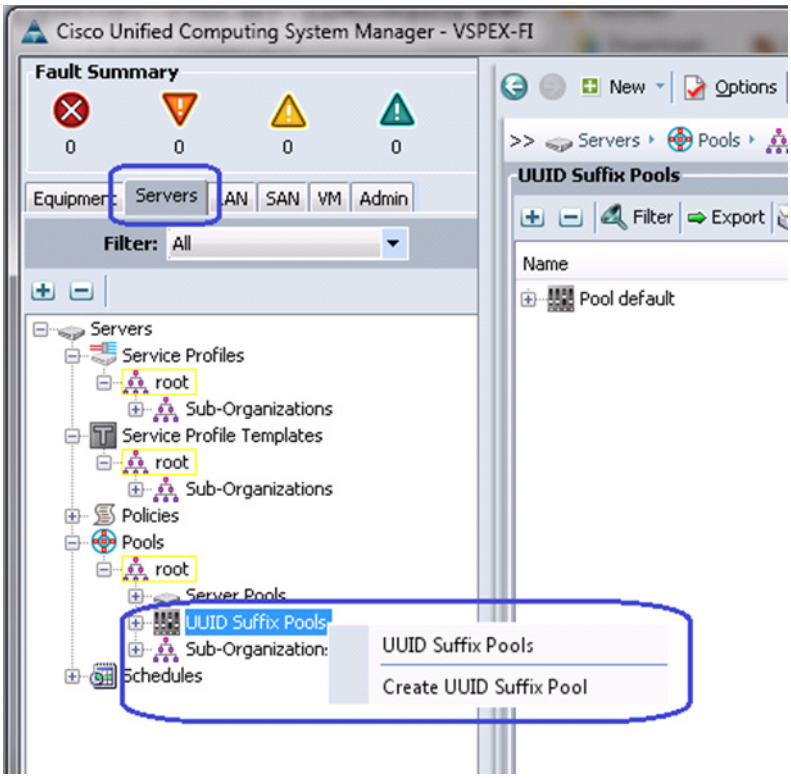

2. Enter the name of the UUID Suffix Pool and (optional) description. Keep other configurations at default.

Figure 70 Specifying Details for Creating UUID Suffix Pool

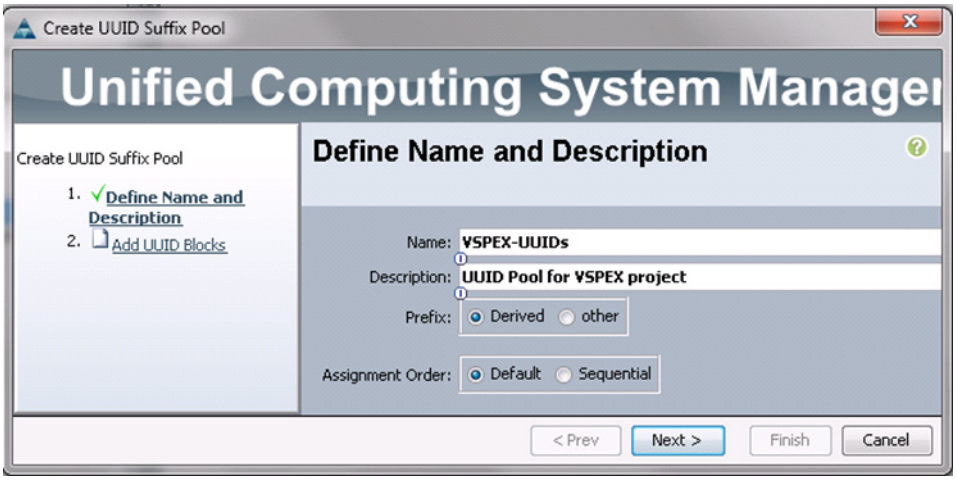

3. Click Add to add UUID block.

Figure 71 Adding UUID Block

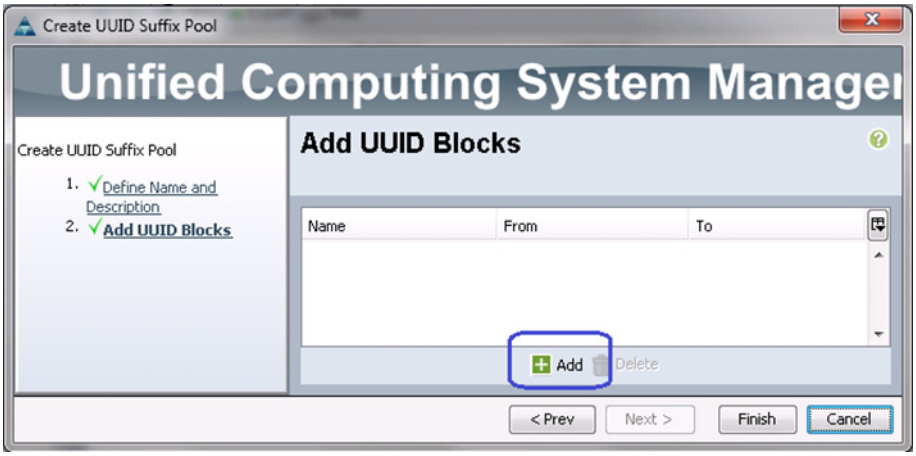

4. Specify the beginning of the UUIDs, and make sure to have a large block of UUID to accommodate future expansion.

Figure 72 Size of the UUID Suffix Block

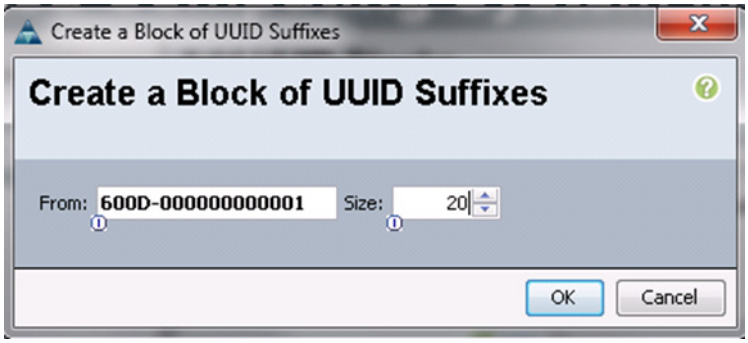

5. Click **OK** and then click **Finish** to deploy UUID pool.
6. From the **LAN** tab, expand **LAN > Pools > root**. Right-click on MAC Pools and choose Create MAC Pool.

Figure 73 Creating MAC Pool

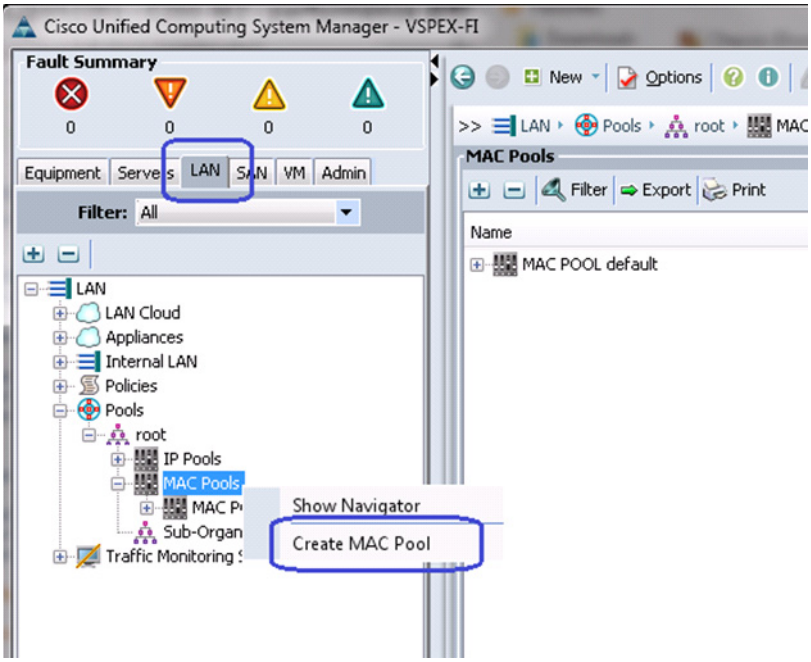

7. Enter the MAC Pool name in the Name field and (optional) description. Click **Next**.

Figure 74 Specifying Details for Creating MAC Pool

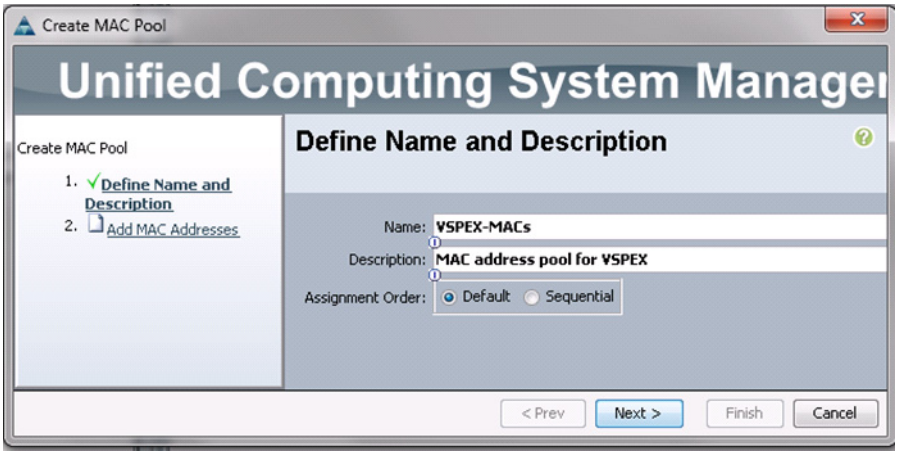

8. Click Add to add MAC pool block.

Figure 75 Adding MAC Address Block

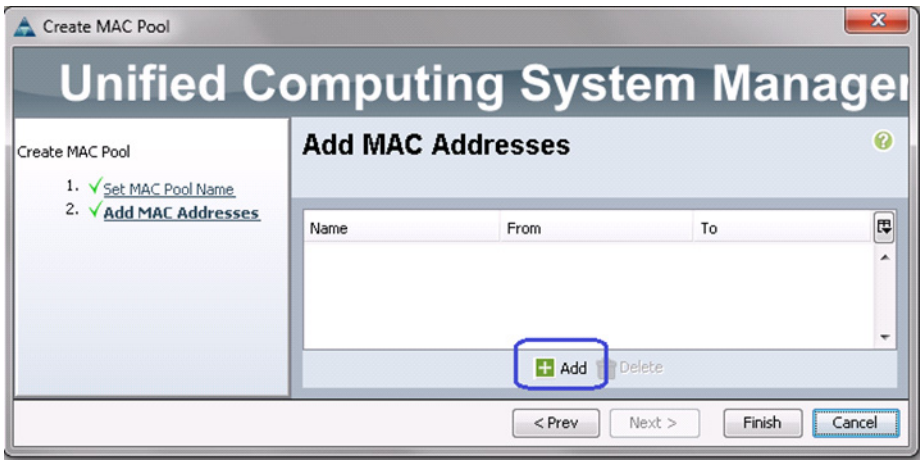

9. Provide the initial MAC address and size of the block. Make sure you provide a large block of MAC addresses to accommodate future expansion. Make sure that you have 6 MAC addresses per server.

Figure 76 Size of the MAC Address Block

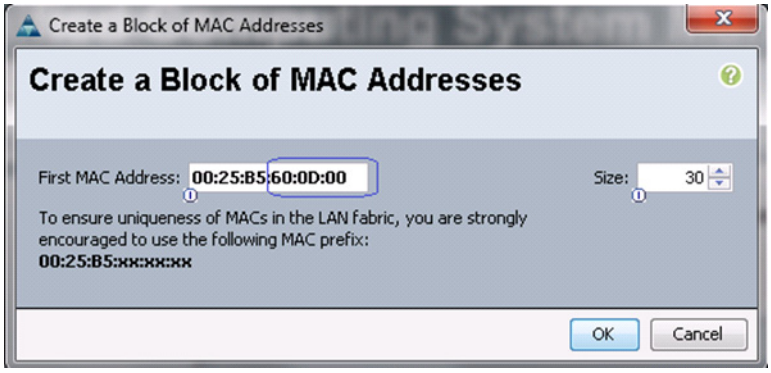

10. Click **OK** and then click **Finish** to complete the configuration.
11. (FC-variant only) Click the **SAN** tab, expand **SAN** > **Pools** > **root**, right-click WWxN Pools, and choose Create WWxN Pool.

Figure 77 Creating WWxN Pool

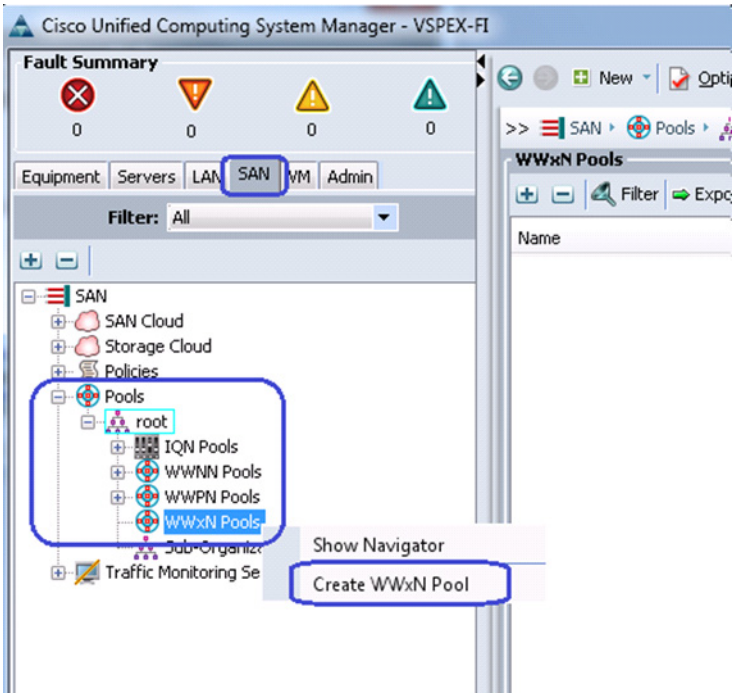

12. (FC-variant only) Enter the name of the WWxN pool in the Name field, (optional) description and select 3 Ports per Node from the drop-down for Max Ports per Node.

Figure 78 Specifying Details for Creating WWxN Pool

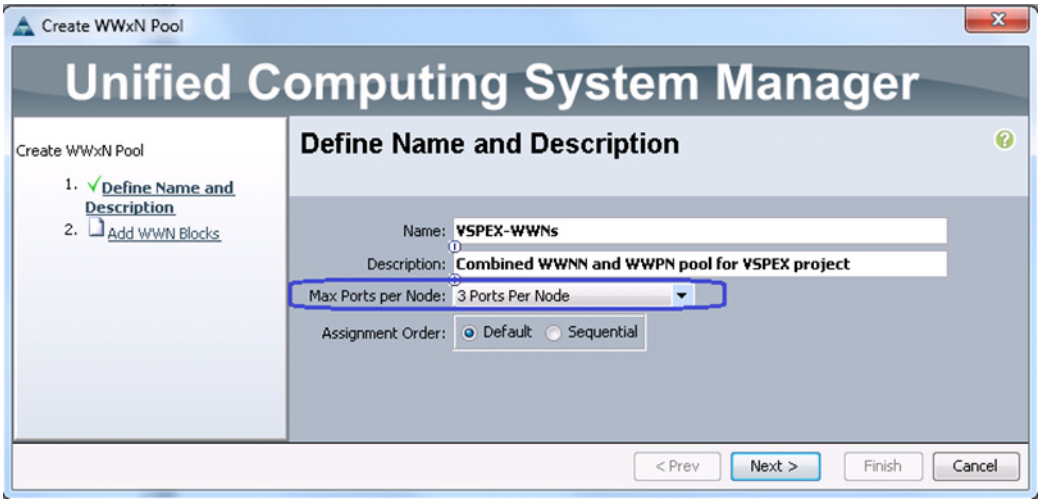

13. (FC-variant only) Click Add to add a block of WWxN IDs.

Figure 79 Adding WWxN Block

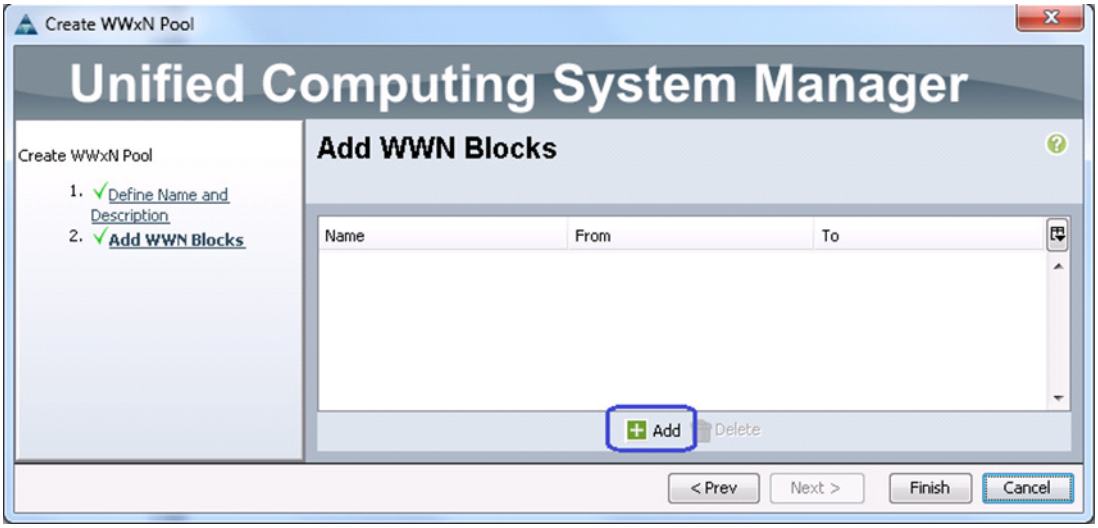

14. (FC-variant only) Provide beginning of the WWN IDs and make sure to have a larger block size. Click **OK** and then click **Finish**.

Figure 80 Size of the WWxN Block

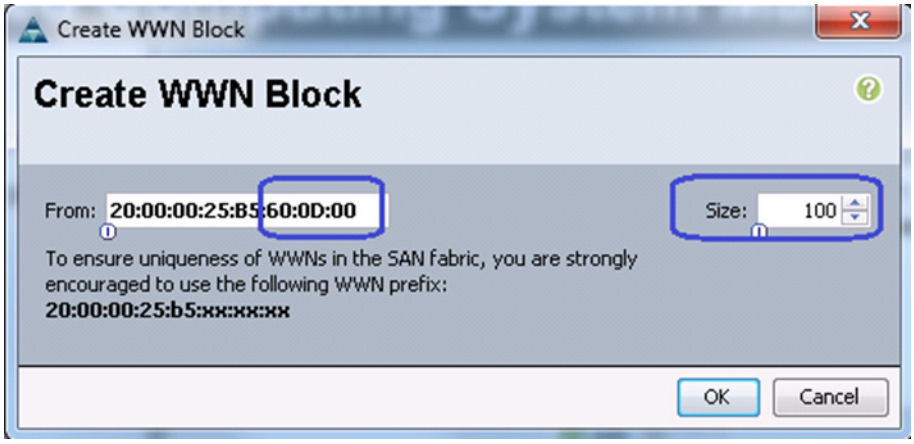

15. (iSCSI-variant only) From the **SAN** tab, expand **SAN > Pools > root**, right-click on IQN Pools, and choose Create IQN Suffix Pool.

Figure 81 Creating IQN Suffix Pool

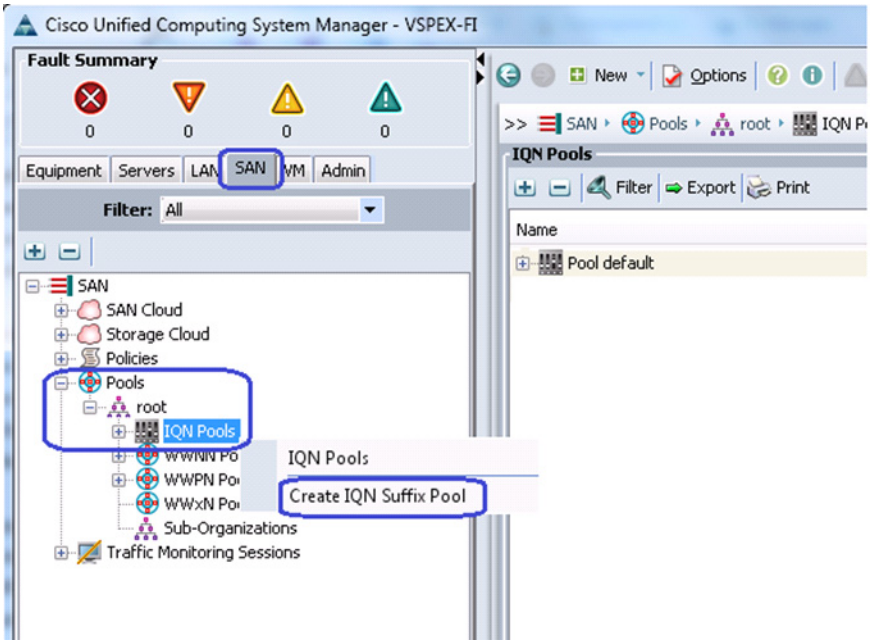

16. (iSCSI-variant only) Enter the IQN Pool name in the Name field, (optional) description and Enter the Prefix for the IQN pool. We have chosen vspex for prefix in this example.

Figure 82 Specifying Details for Creating IQN Suffix Pool

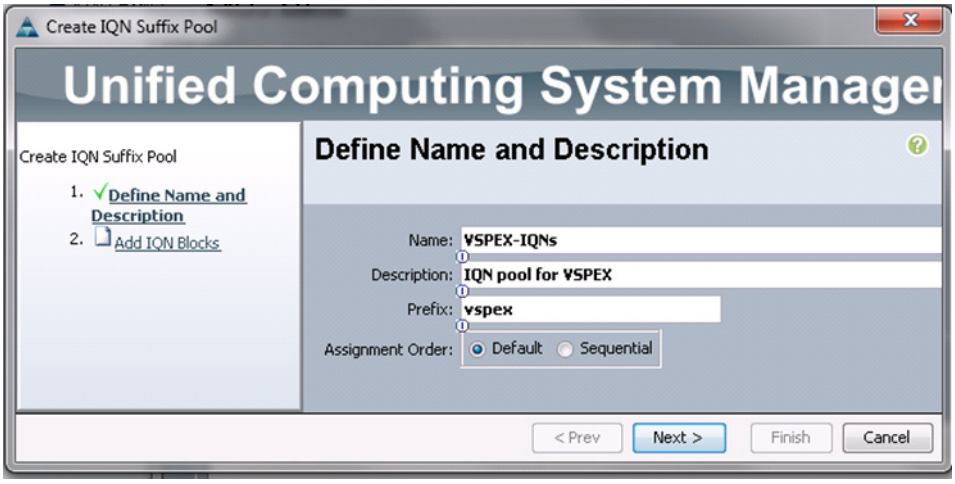

17. (iSCSI-variant only) Click **Next** and click Add to add an IQN identifier block.

Figure 83 Adding IQN Block

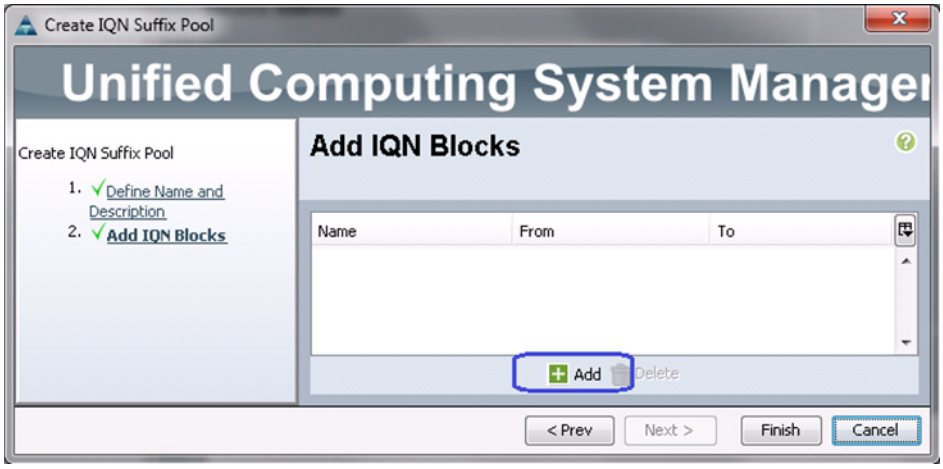

18. (iSCSI-variant only) Enter the IQN name in the Suffix field, enter a value in the From field and specify the size of the block in the Size field. The IQN identifiers are created in the following format "<prefix>:<suffix>:<from>", "<prefix>:<suffix>:<from+1>", …. "<prefix>:<suffix>:<from+size-1>".

Figure 84 Size of the IQN Block

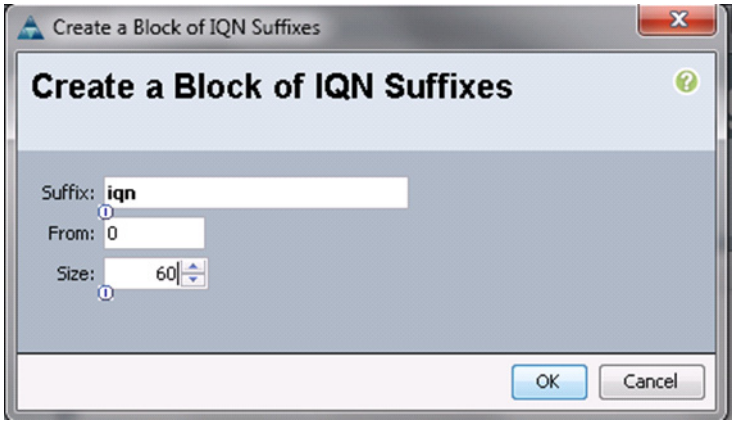

19. Provide initial IP address, size of the pool, default gateway and subnet mask as shown below. Click **OK** to deploy the configuration. IP addresses are assigned to various rack server CIMC management access from this block.

20. Next, you need to create 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 > IP Pool ext-mgmt.** Click **Create Block of IP addresses** in the Actions area on the right pane of the window.

Figure 85 Creating IP Address Block for CIMC Management

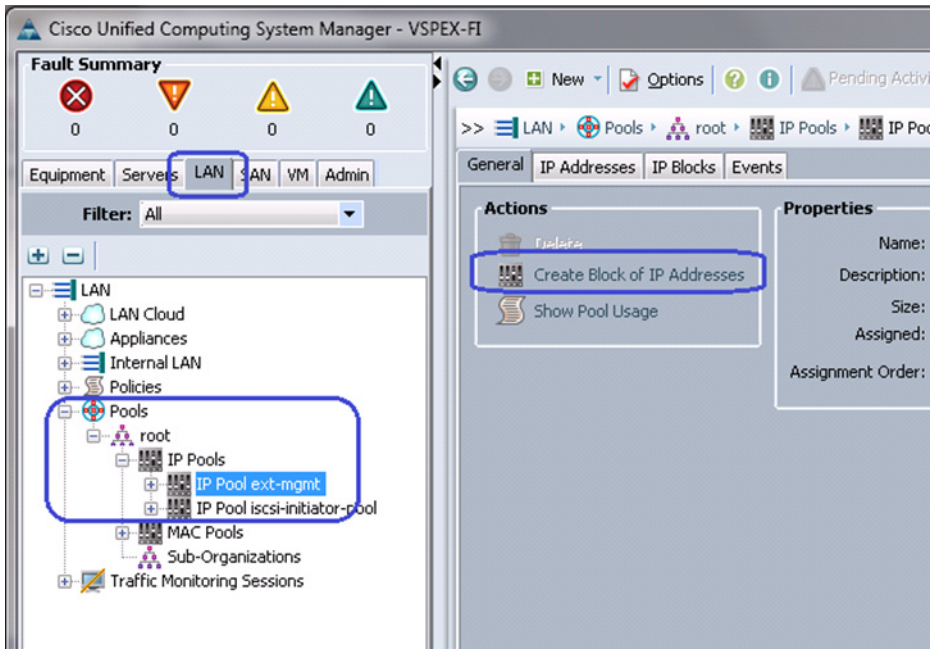

21. Specify the initial IP address in the From field, size of the pool in the Size field, Default Gateway value and Subnet Mask in the respective fields. Click **OK** to deploy the configuration. IP addresses are assigned to various rack server CIMC management access from this block.

Figure 86 Specifying Details for Creating IP Address Block for CIMC Management

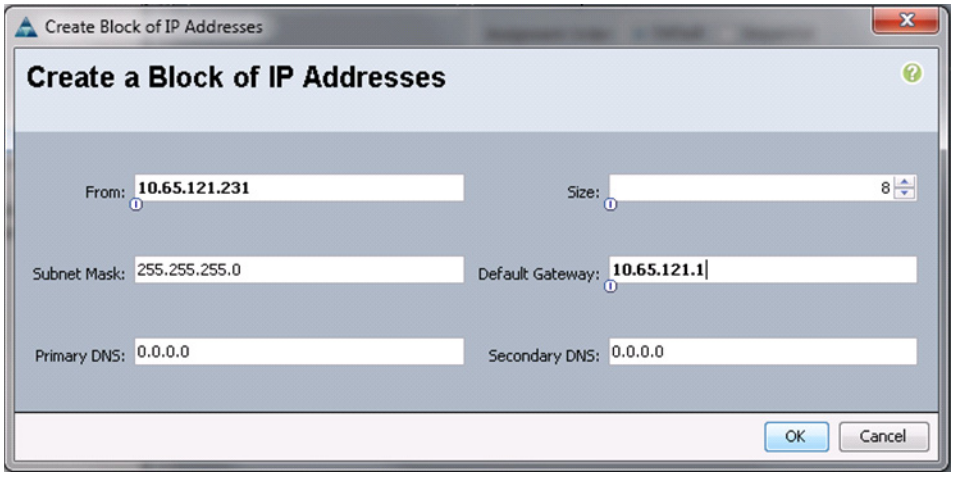

22. (iSCSI-variant only) To create iSCSI initiator IP address pool, from **LAN** tab, expand **LAN > Pools > root.** Right-click on IP Pools and choose Create IP Pool.

Figure 87 Creating IP Address Pool for iSCSI Initiator Pool

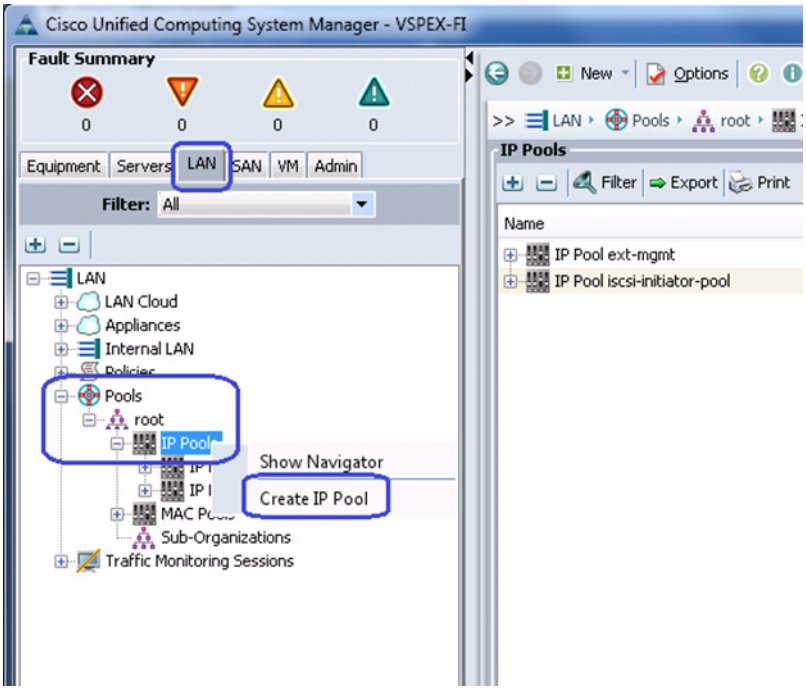

23. (iSCSI-variant only) Enter the IP pool name in the Name field and (optional) description for the IP pool, click **Next**. Pool name signifies that the pool is created for fabric A.

Figure 88 Specifying Details for Creating IP Address Pool for iSCSI Initiator Pool

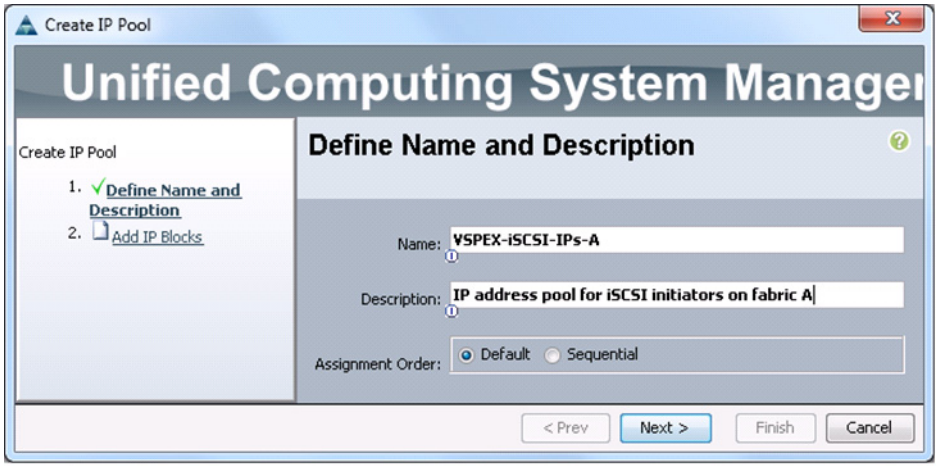

24. (iSCSI-variant only) Click Add to add iSCSI initiator IP address block.

Figure 89 Adding IP Address Block for iSCSI Initiator Pool

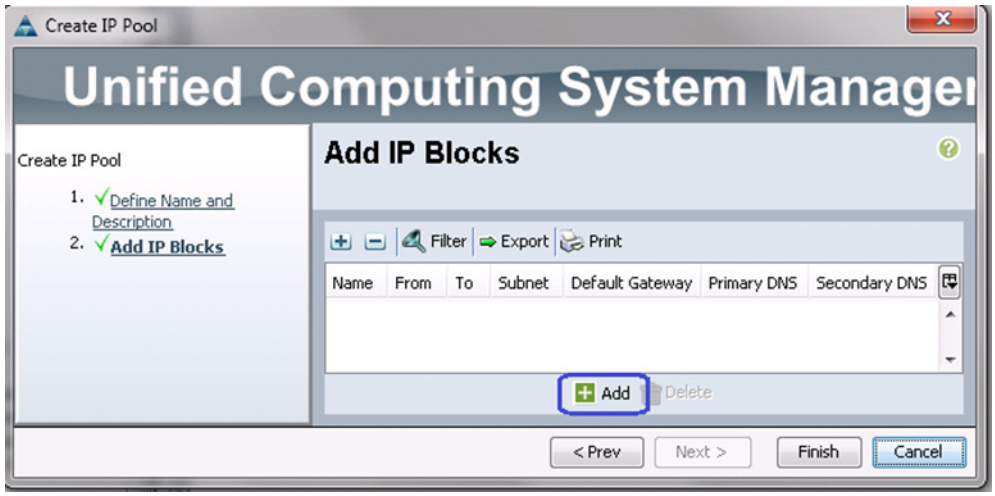

25. (iSCSI-variant only) Provide initial IP address of the block, size of the block, subnet mask and default gateway as shown below.

Figure 90 Specifying Details for Creating IP Address Block for iSCSI Initiator Pool

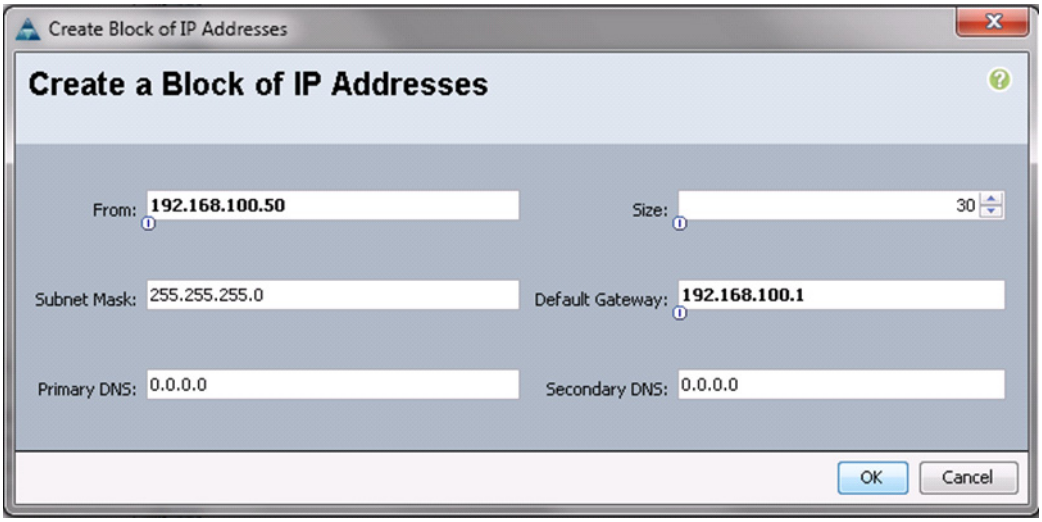

26. (iSCSI-variant only) Click **OK** and **Finish**.
27. (iSCSI-variant only) iSCSI storage access best practices suggests different VLANs and subnets on two fabrics. Given that, you need to create one more iSCSI initiator IP address pool for fabric B on a different subnet. Repeat steps 18 to 22 to create VSPEX-iSCSI-IPs-B pool.

With this all the identifier pools and block configurations are completed.

Configure Server Pool and Qualifying Policy

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

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

To complete these tasks, follow these steps:

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

Figure 91 Creating Server Pool

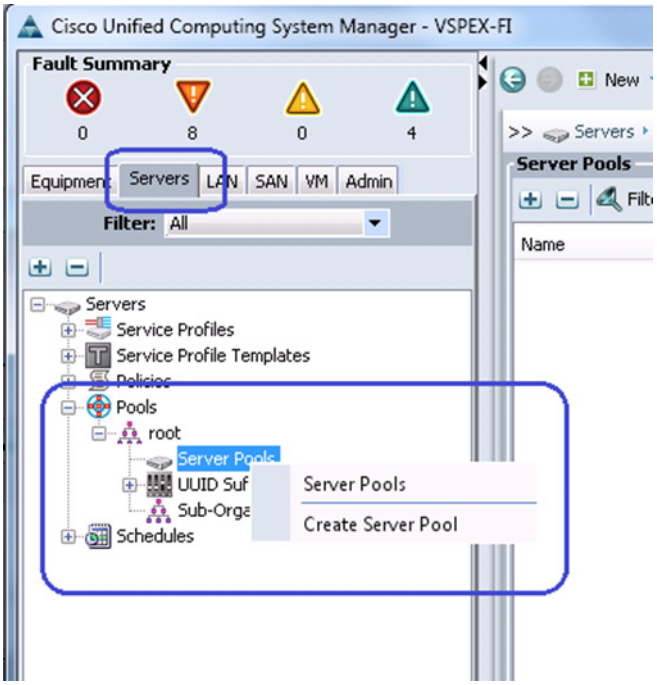

2. Enter the name of the server pool in the Name field, and click **Next**.

Figure 92 Specifying Details for Creating Server Pool

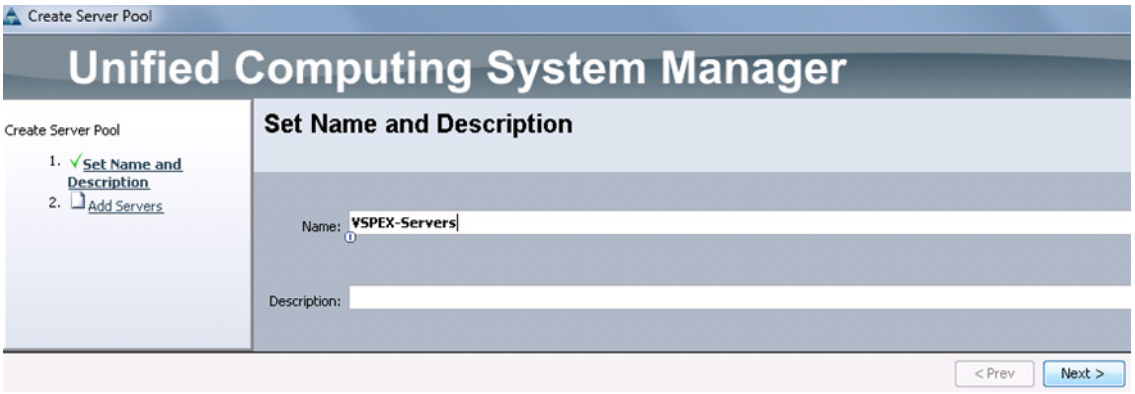

3. Click **Finish** to create the empty server pool. Compute resources need to be added to this pool dynamically, based on policy.

Figure 93 Adding Servers to the Server Pool

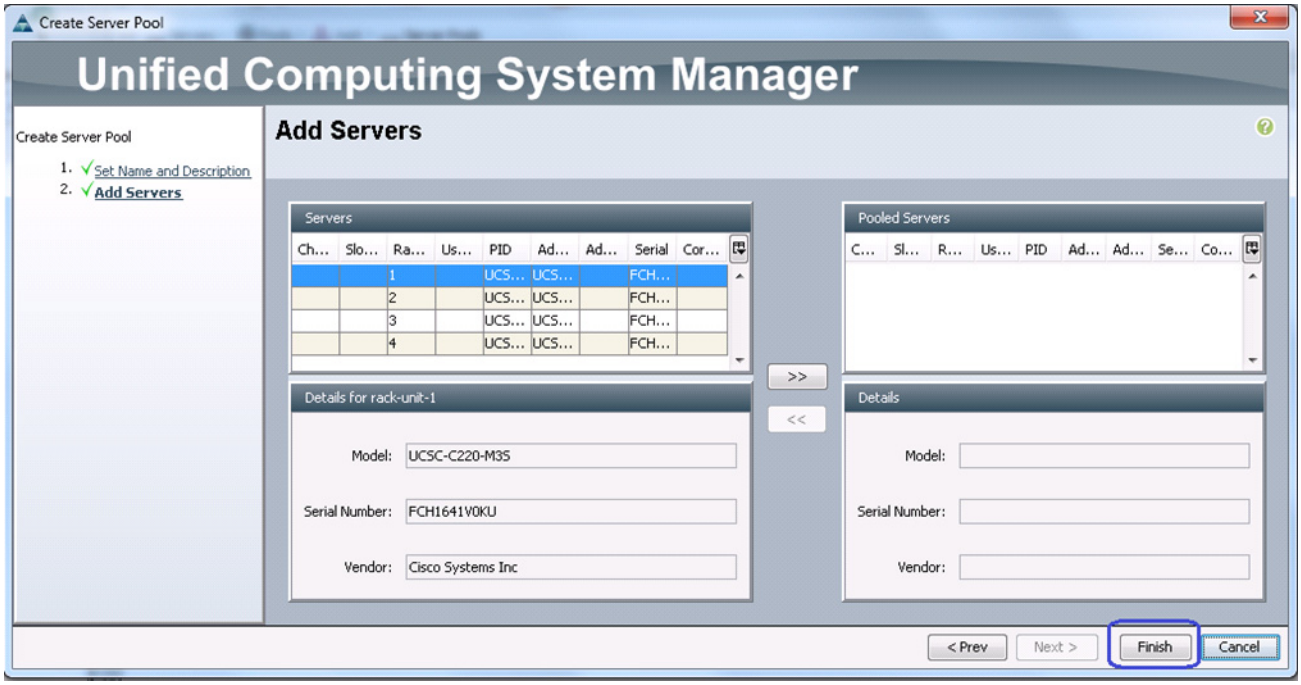

4. From the **Servers** tab, expand **Servers > Policies > root**, right-click on Server Pool Policy Qualifications and choose Create Server Pool Policy Qualification.

Figure 94 Creating Server Pool Policy Qualification

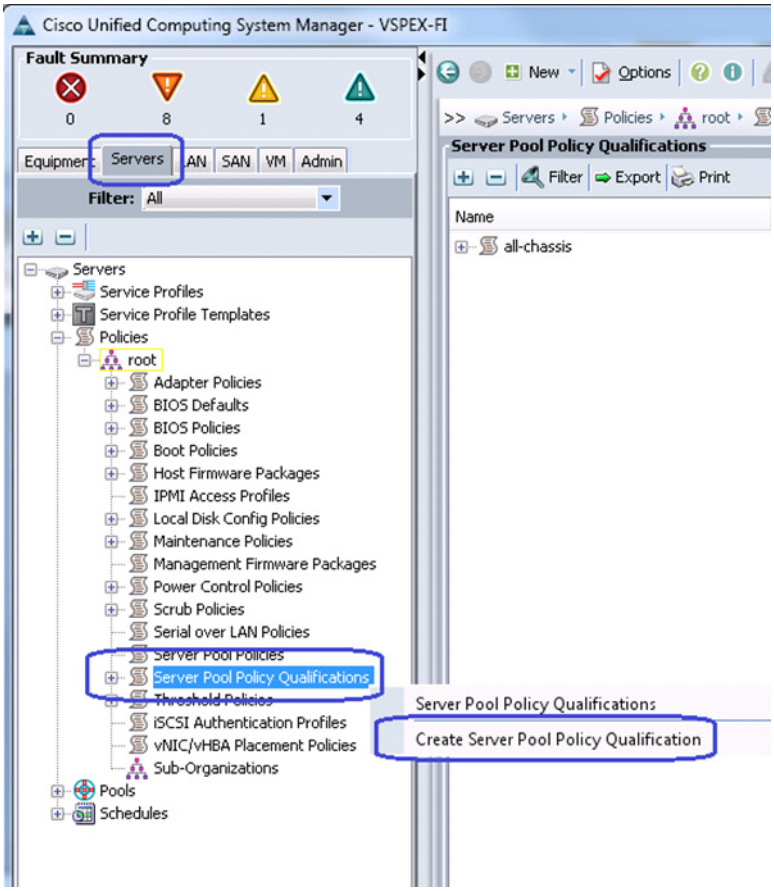

5. Enter the name of the server policy qualification criterion in the name field. We have chosen minimum memory qualification criterion in this example.

Figure 95 Creating Memory Qualifications

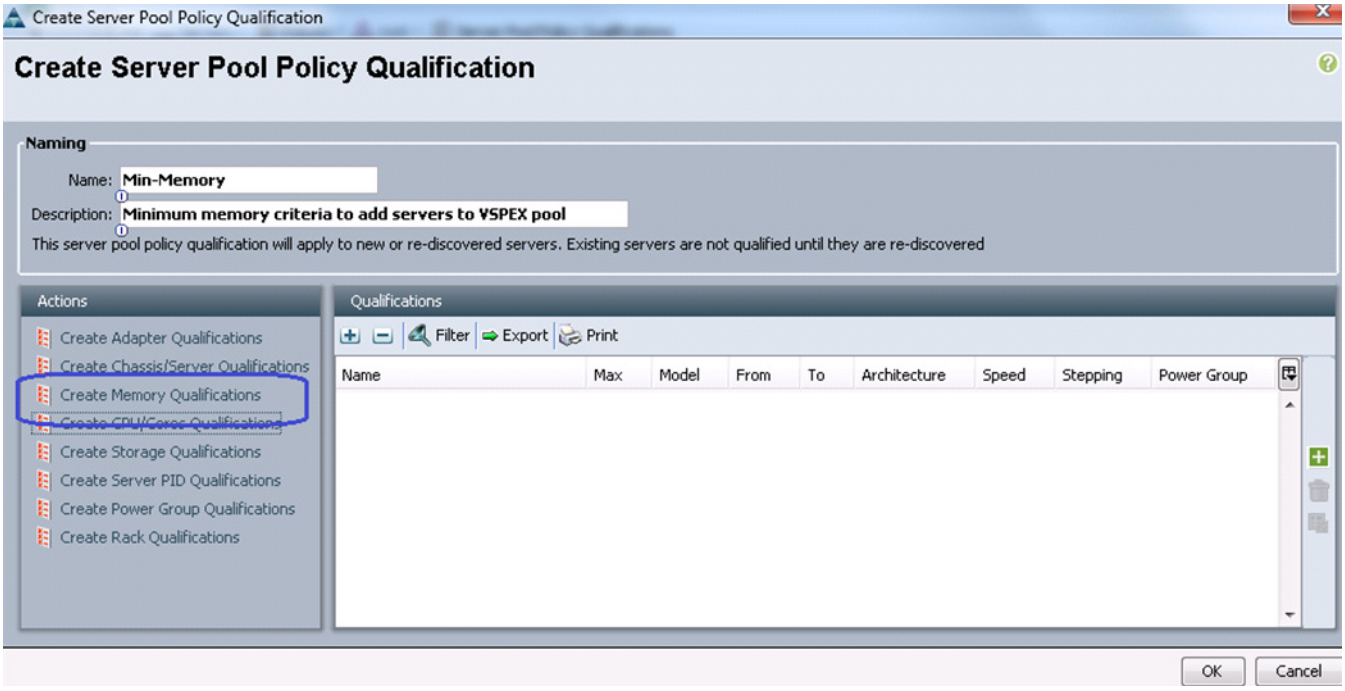

6. Set minimum 64 GB RAM as the pool qualification criterion. Click **OK** twice to create the qualification.

Figure 96 Specifying Minimum Memory Capacity

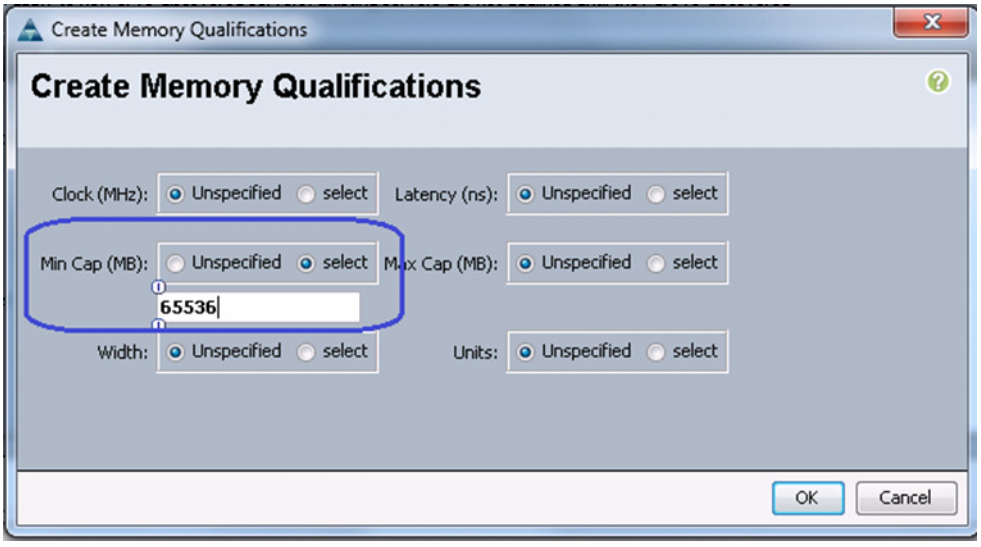

Note This is an example criterion, you need to choose a criterion that suites your requirement.

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

Figure 97 Creating Server Pool Policy

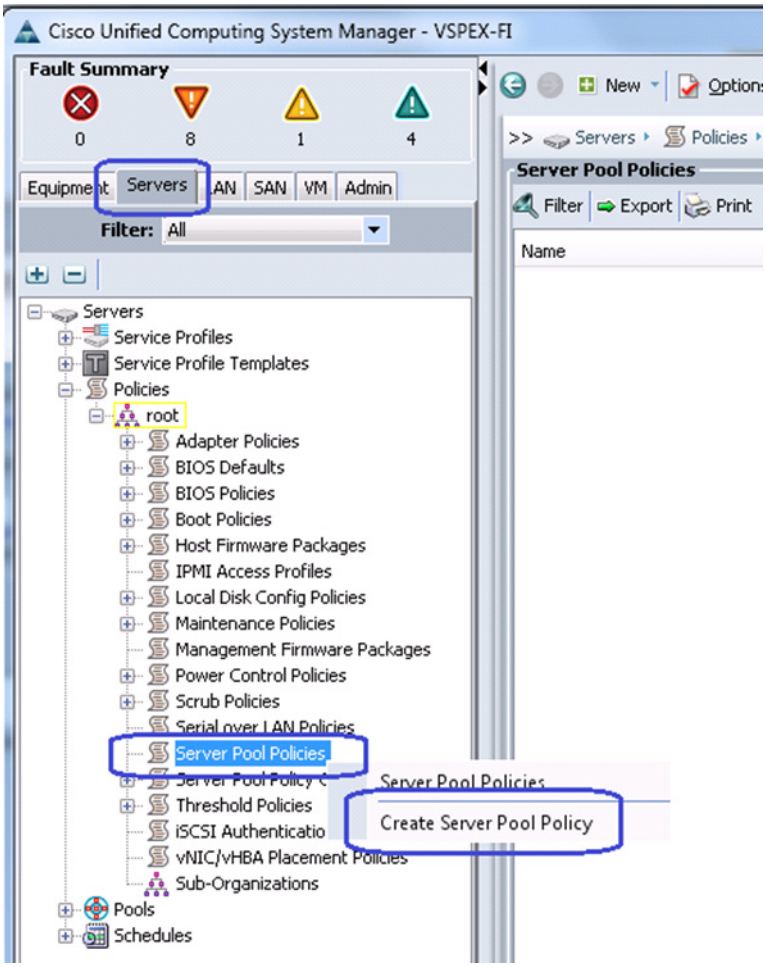

8. Enter the name of the server pool policy in the name field and (optional) description. Choose the recently created Target Pool and Qualification from the respective drop-down lists. Click **OK** to deploy the configuration.

Figure 98 Specifying Details for Creating Server Pool Policy

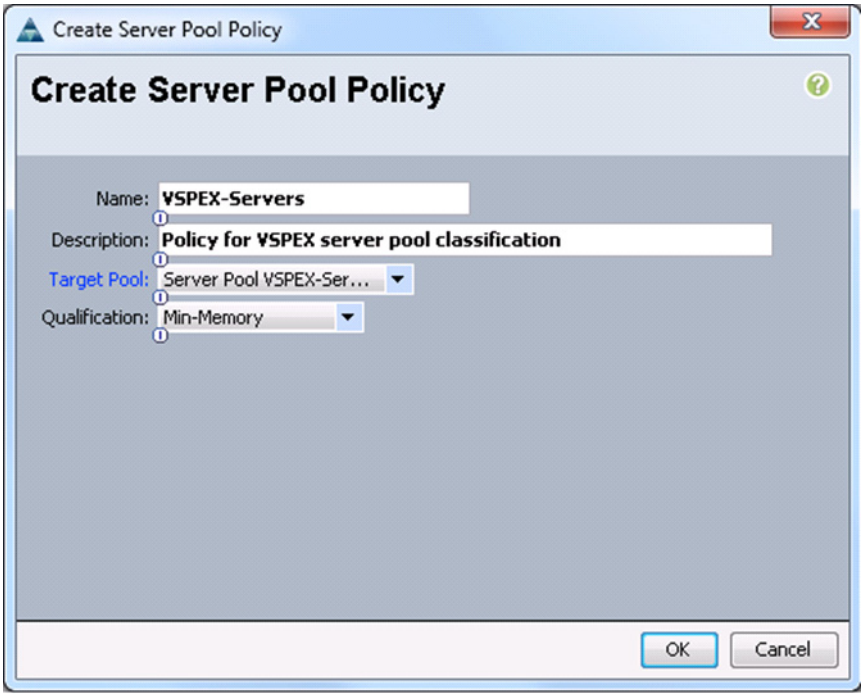

9. After creating the qualification criterion, in the **Servers** tab under Server Pools, select the created Target Pool. In the right pane of the window, you will see that all the compute resources that meet the qualification criteria are dynamically added to the server pool. Figure 99 is from the FC-variant of the architecture, where combination of UCS B200 M3 blade servers and C220 M3 rack servers are used to share the work load. This architecture showcases the form-factor independent architecture with managed servers using UCS Manager.

Figure 99 Server Pool Showing Added Cisco UCS Blade and Rack-Mount Server

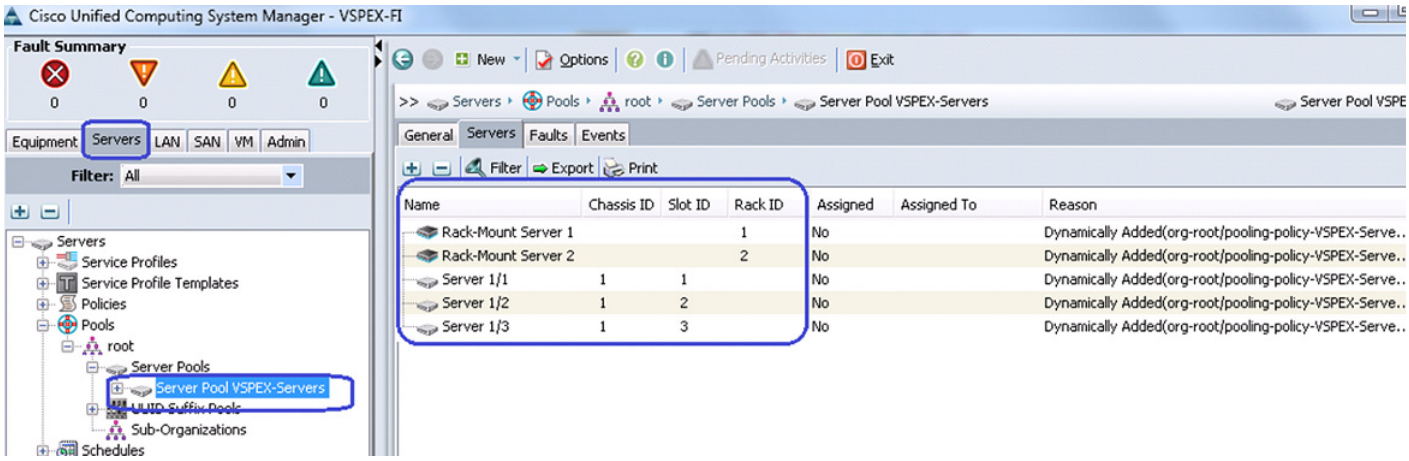

Configure Service Profile Template

Service profile template needs to be created from which we can instantiate individual service profiles. Follow these steps to create the service profile template:

1. Click the **Servers** tab, under Servers, select and right-click on Service Profile Templates. Choose Create Service Profile Template.

Figure 100 Creating Service Profile Template

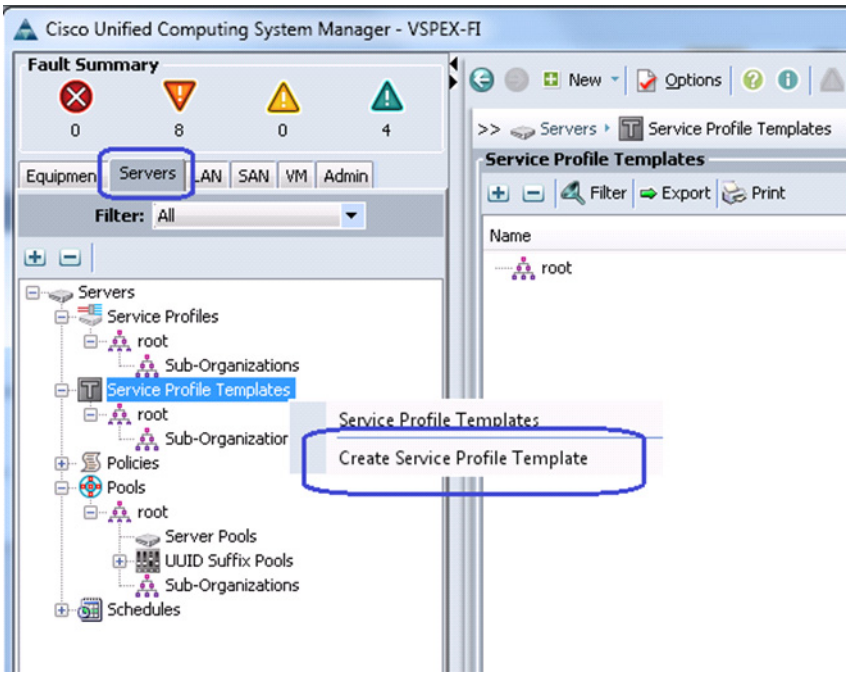

2. Enter the service profile template name in the Name field, let the Type be **Initial Template**, and choose UUID pool for UUID Assignment field from the drop-down list. Click **Next**.

Figure 101 Specifying Details for Creating a Service profile Template

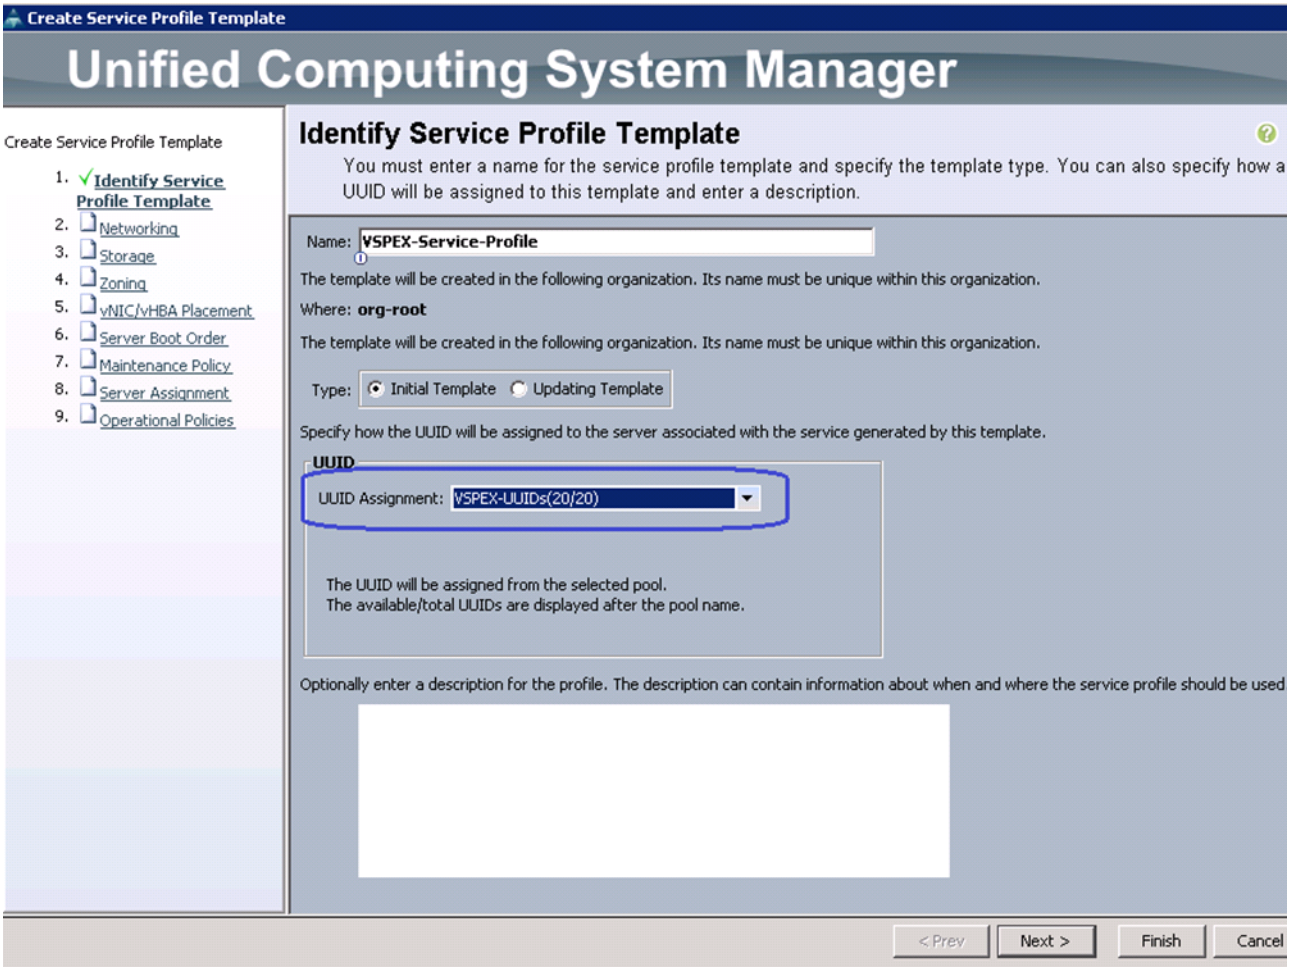

3. Click the **Expert** radio button for LAN connectivity. Click Add to create a vNIC.

Figure 102 LAN Configuration Details

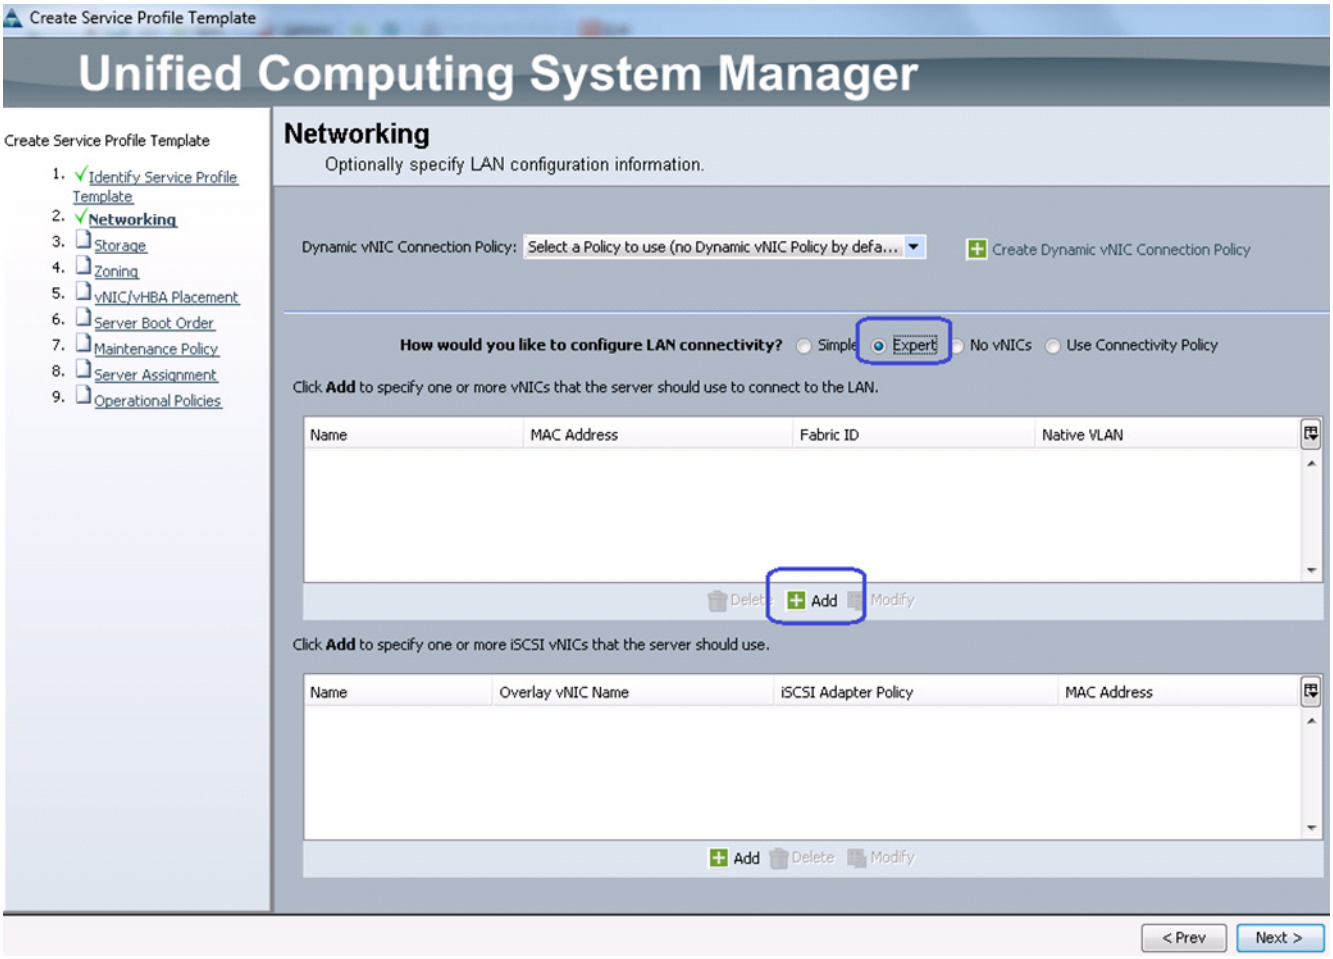

4. Create a system vNIC for fabric A. Enter the vNIC name as System A in the Name field, choose the MAC pool created from the MAC Assignment drop-down list. Click the **Fabric A** radio button for fabric ID. Check the check boxes vMotion and vSphereMgmt VLANs. Click the **vSphereMgmt** radio button to set it as the native VLAN. Enter 9000 in the MTU field, choose VMware as the adapter policy and jumboMTU as the QoS policy.

Figure 103 Specifying vNIC Details

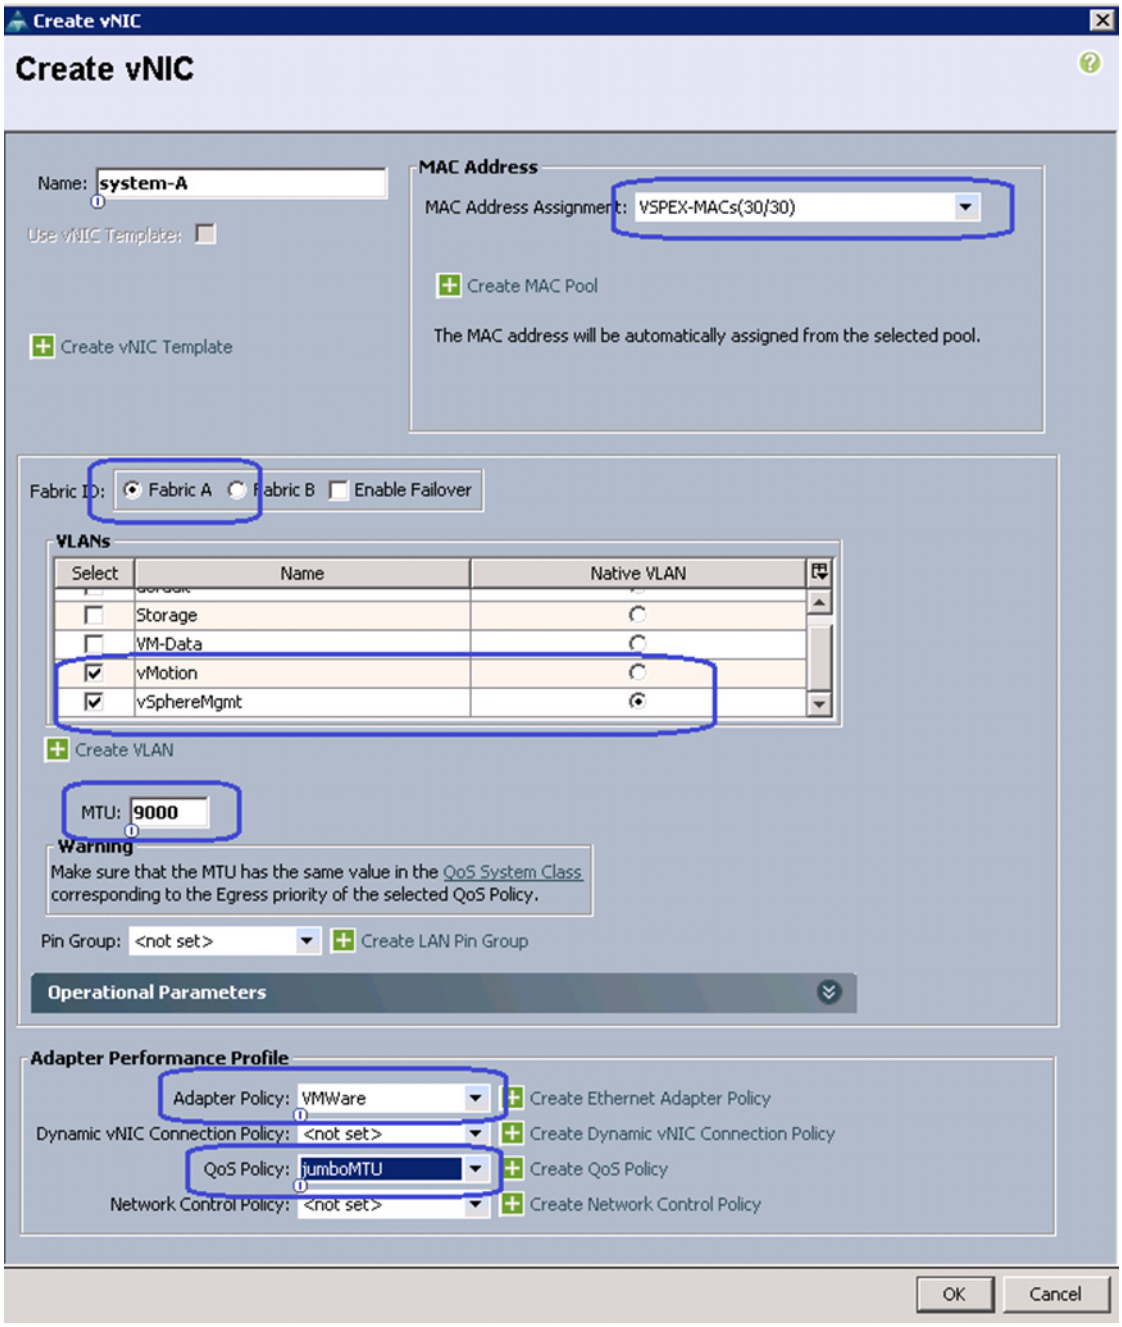

5. Similarly, create one more vNIC with the same properties on fabric B.
6. (iSCSI-variant only) Add one more vNIC for iSCSI storage access. Enter the name as "iSCSI-Overlay-A" in the name field. Choose created MAC address pool from MAC Address Assignment drop-down list, click the **Fabric A** radio button for Fabric ID. Check the check box Storage for VLAN and click the radio button to set this as the native VLAN. Enter 9000 in the MTU field, choose VMware from the Adapter Policy drop-down list and jumboMTU from QoS Policy drop-down list.

7. Similarly, create vNICs for fabric B. Table 8 summarizes all the vNICs created on the service profile.

Table 8 vNICs Created on the Service profile

vNIC Name	MAC Address Assignment	VLANS	Native VLANS	Fabric	MTU	Adapter Policy	QoS Policy
System-A	MAC Pool	vSphereMgnt, vMotion	vSphereMgnt	A	9000	VMware	jumboMTU
System-B	MAC Pool	vSphereMgnt, vMotion	vSphereMgnt	B	9000	VMware	jumboMTU
iSCSI-Overlay-A	MAC Pool	Storage	Storage	A	9000	VMware	jumboMTU
iSCSI-Overlay-B	MAC Pool	Storage	Storage	B	9000	VMware	jumboMTU
Data-A	MAC Pool	VM-Data	VM-Data	A	9000	VMware	-
Data-B	MAC Pool	VM-Data	VM-Data	B	9000	VMware	-

Note The iSCSI Overlay vNICs are created only for iSCSI variant of the solution.

8. For iSCSI variant of the solution, click Add to add iSCSI vNIC.

Figure 104 Networking Window Showing Configured vNICs

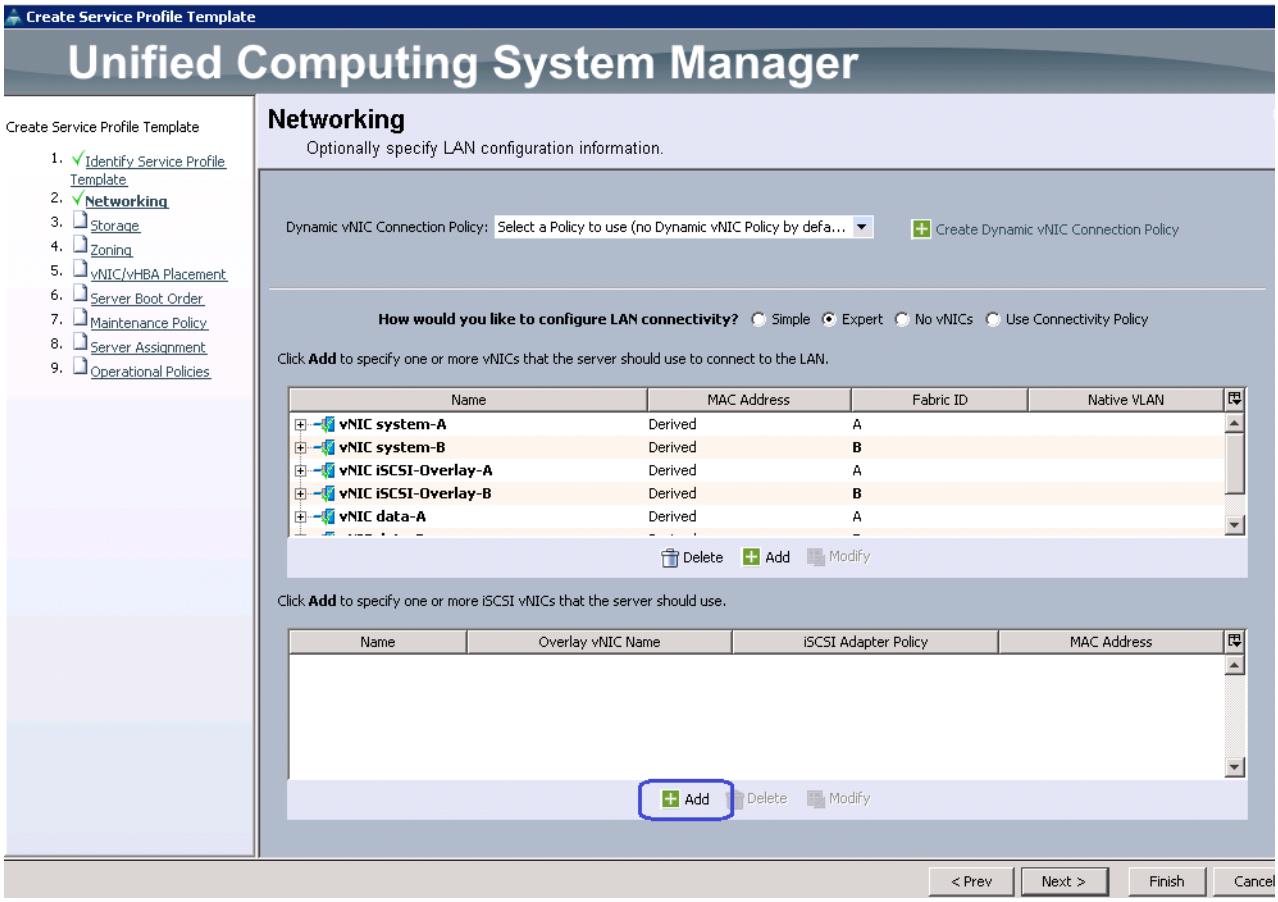

9. (iSCSI-variant only) Enter iSCSI-A in the Name field, choose iSCSI-Overlay-A from the Overlay vNIC drop-down list, choose Storage as the native VLAN from the VLAN drop-down list, and keep the default option that is, None used by default for the MAC Address Assignment field (MAC address would be taken from the overlay vNIC).

Figure 105 Creating iSCSI vNIC

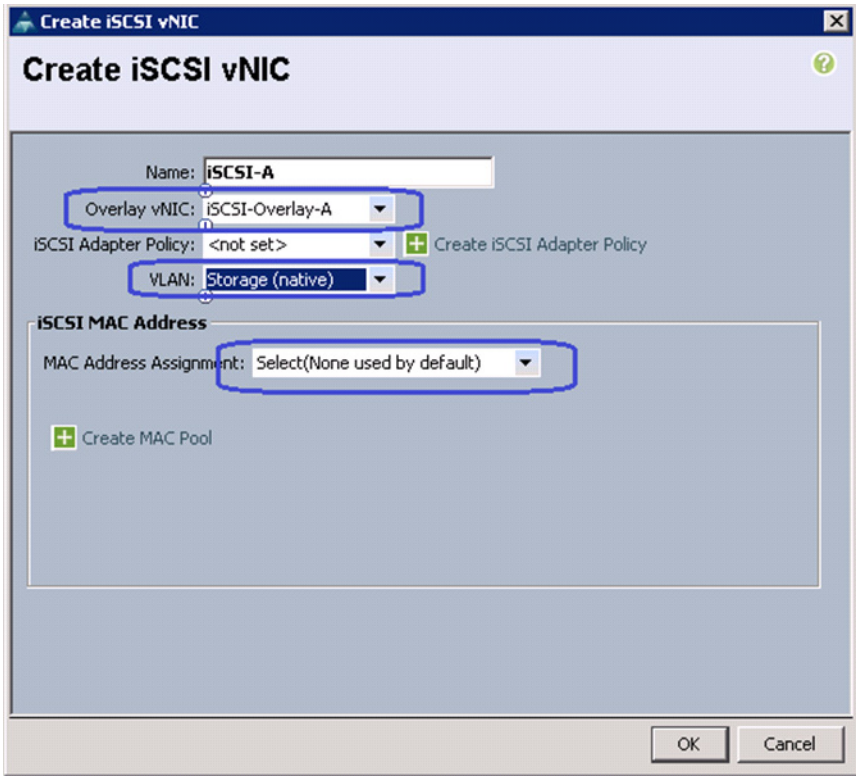

10. Similarly, create another iSCSI vNIC for fabric B. Click **Next** to configure Storage. (iSCSI-variant only)
11. For FC-variant of the solution, click the **Expert** radio button for SAN connectivity and choose VSPEX-WWNs from WWNN pool drop-down list. Click Add to add vHBA.

Figure 106 SAN Configuration Details for FC-variant of the Solution

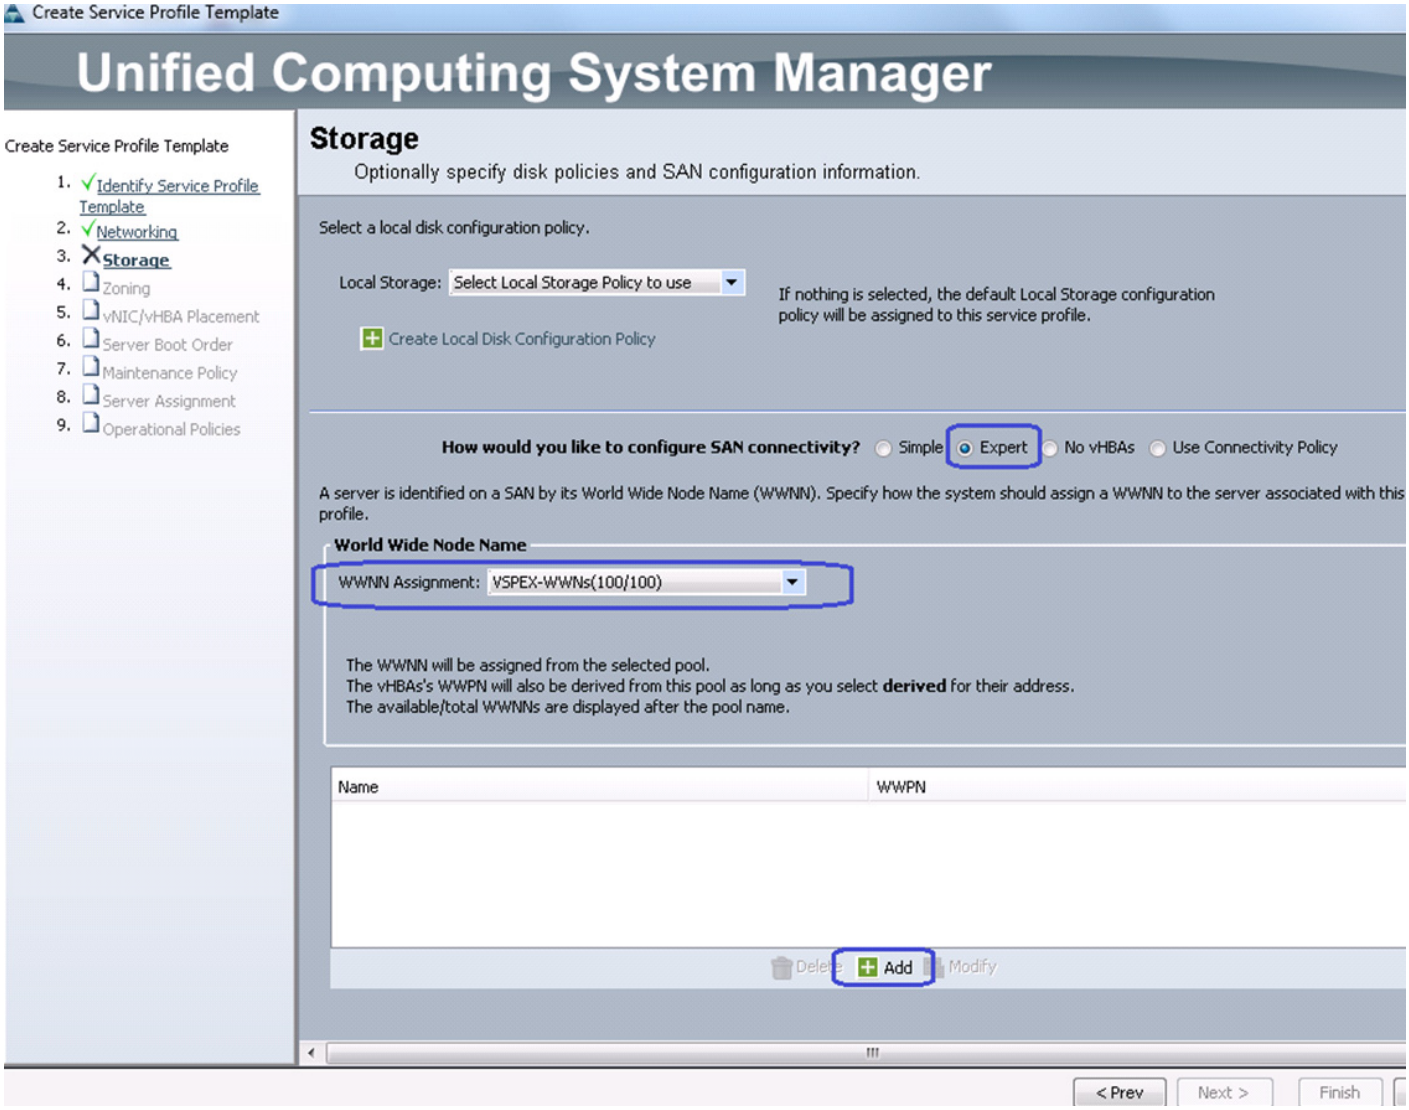

12. (FC-variant) To create a vHBA, enter vHBA-A in the Name field, choose WWPN Assignment as Derived from the drop-down list. Click the **A** radio button for Fabric ID, choose Storage from the Select VSAN drop-down list, and choose VMware from the Adapter Policy drop-down list. Click **OK** to deploy the vHBA.

Figure 107 Specifying vHBA Details

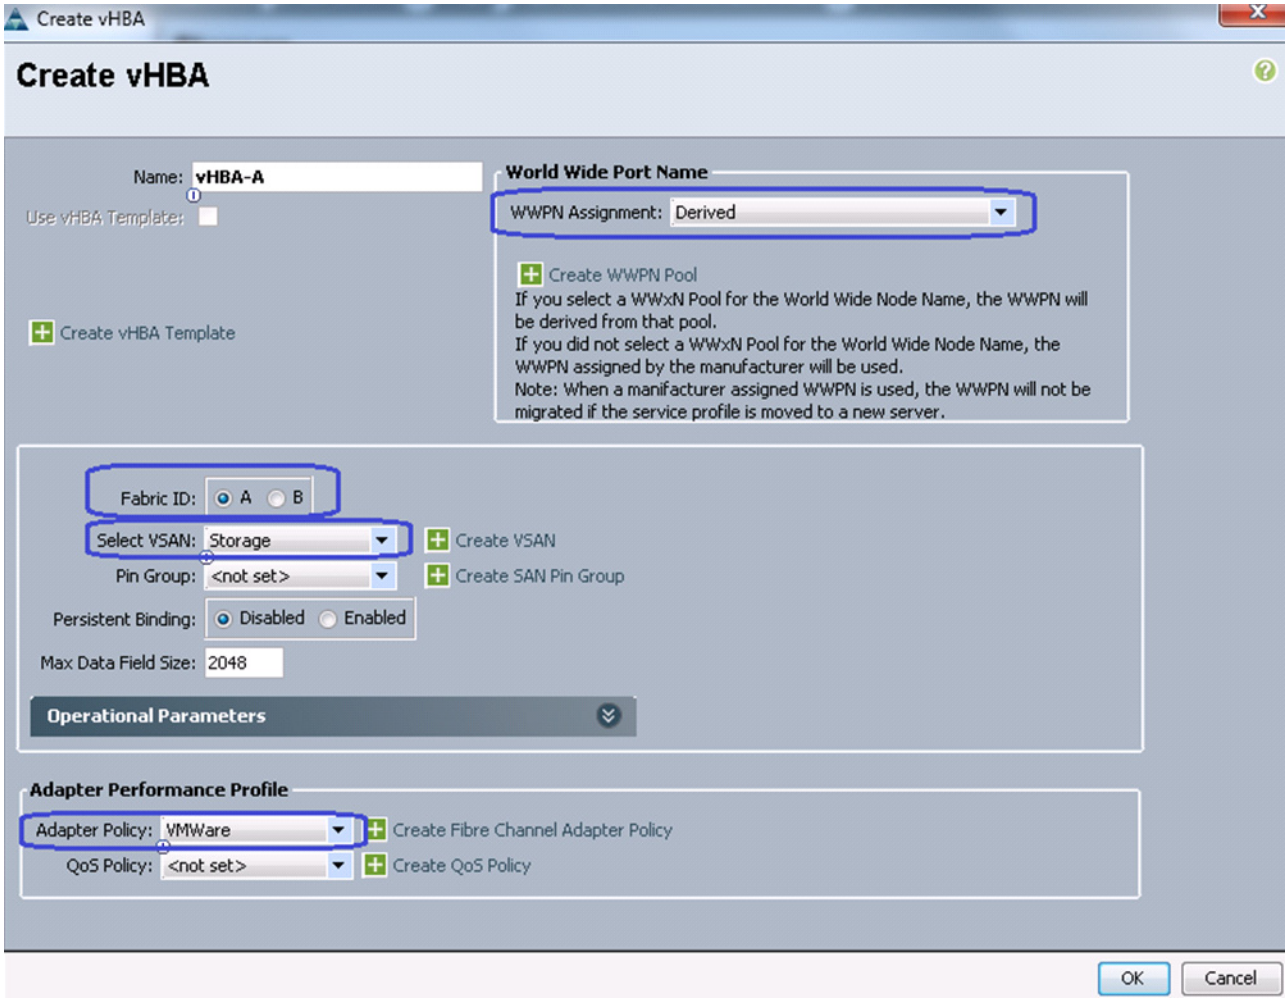

13. (FC-variant only) Repeat step 14 for vHBA-B on fabric B, with the rest of the configuration being same.
14. For iSCSI-variant of the solution, in the Storage configuration window, click the **No VHBAs** radio button for SAN connectivity and click **Next**.

Figure 108 SAN Configuration Details for iSCSI-variant of the Solution

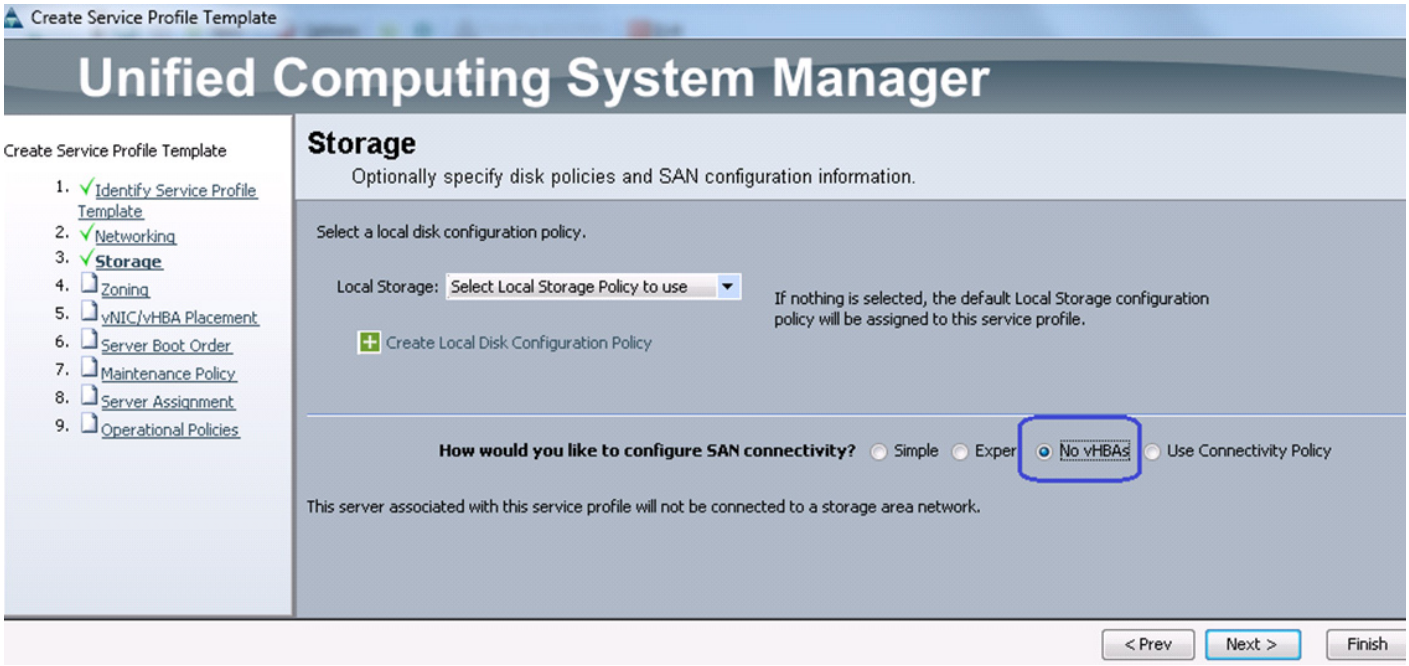

15. For FC-variant of the solution, in the Zoning window of the wizard, click Add.

Figure 109 Creating Initiator Group

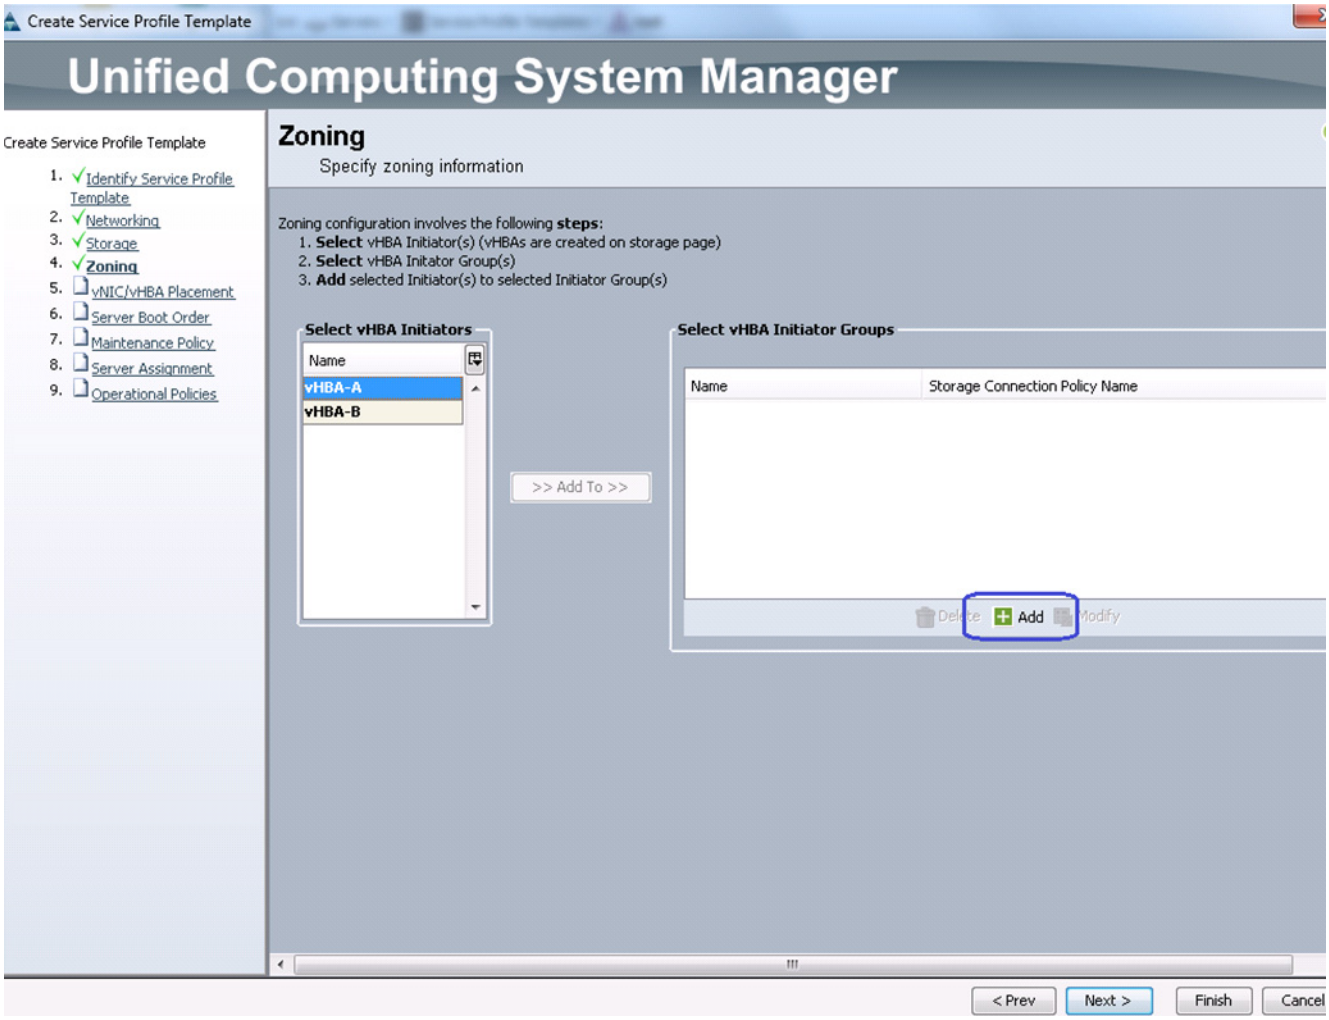

16. (FC-variant only) Enter the name of the VNBA as SAN-A in the Name field in the vHBA initiator group area, and choose previously configured Fabric-A from the Storage Connection Policy (zoning policy) drop-down list and click **OK**.

Figure 110 Specifying vHBA Initiator Group Details

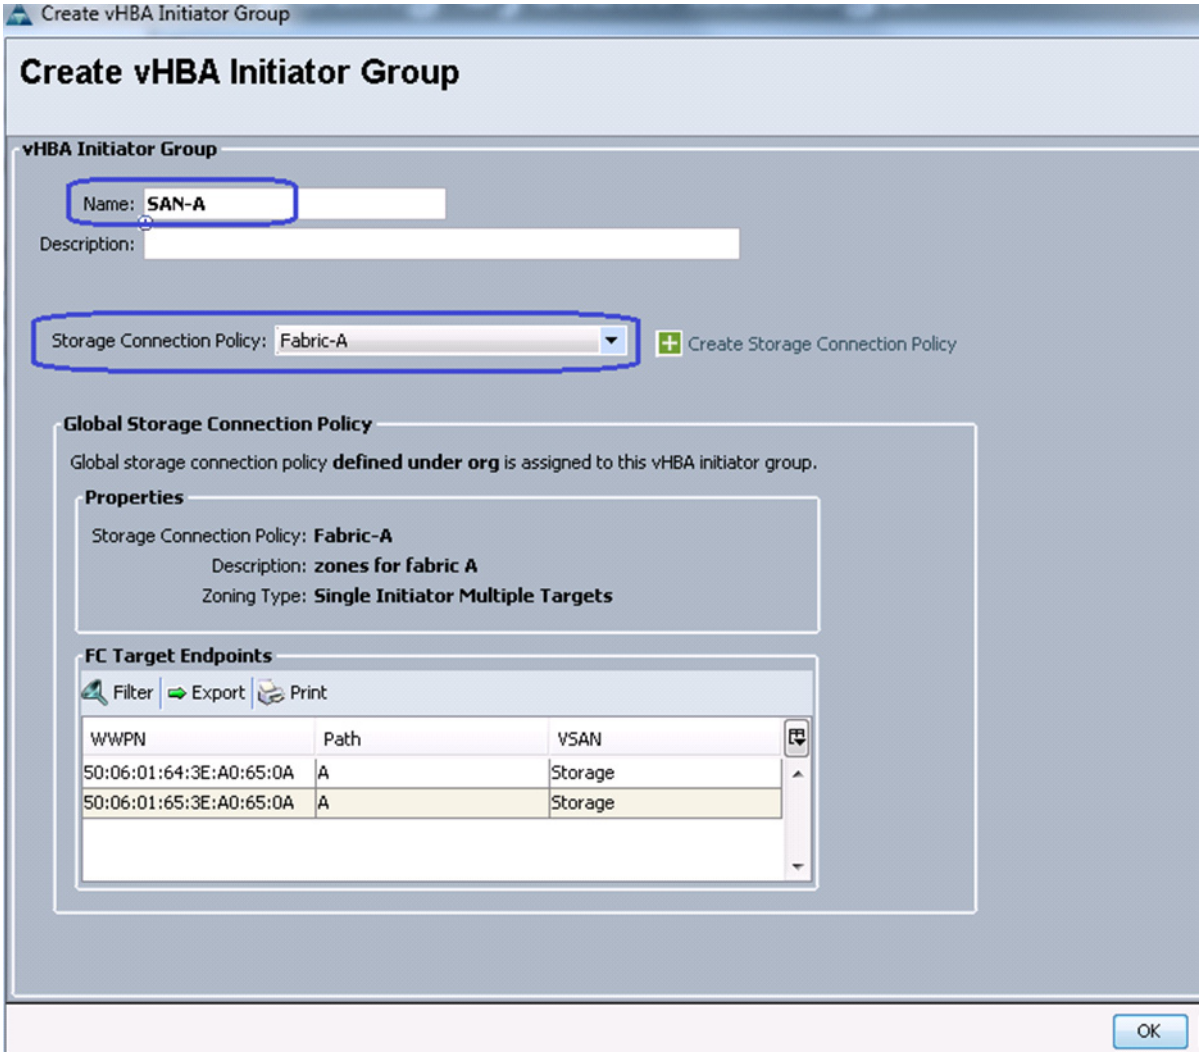

17. (FC-variant) Repeat steps 15 and 16 for zoning on fabric B.
18. Choose vHBA-A from the list of vHBA initiators and choose SAN-A from the vHBA Initiator Groups in the right pane of the window, and click >> Add To >> in the mid pane to add the initiator to the initiator group.

Figure 111 Adding vHBA Initiator to vHBA Initiator Group

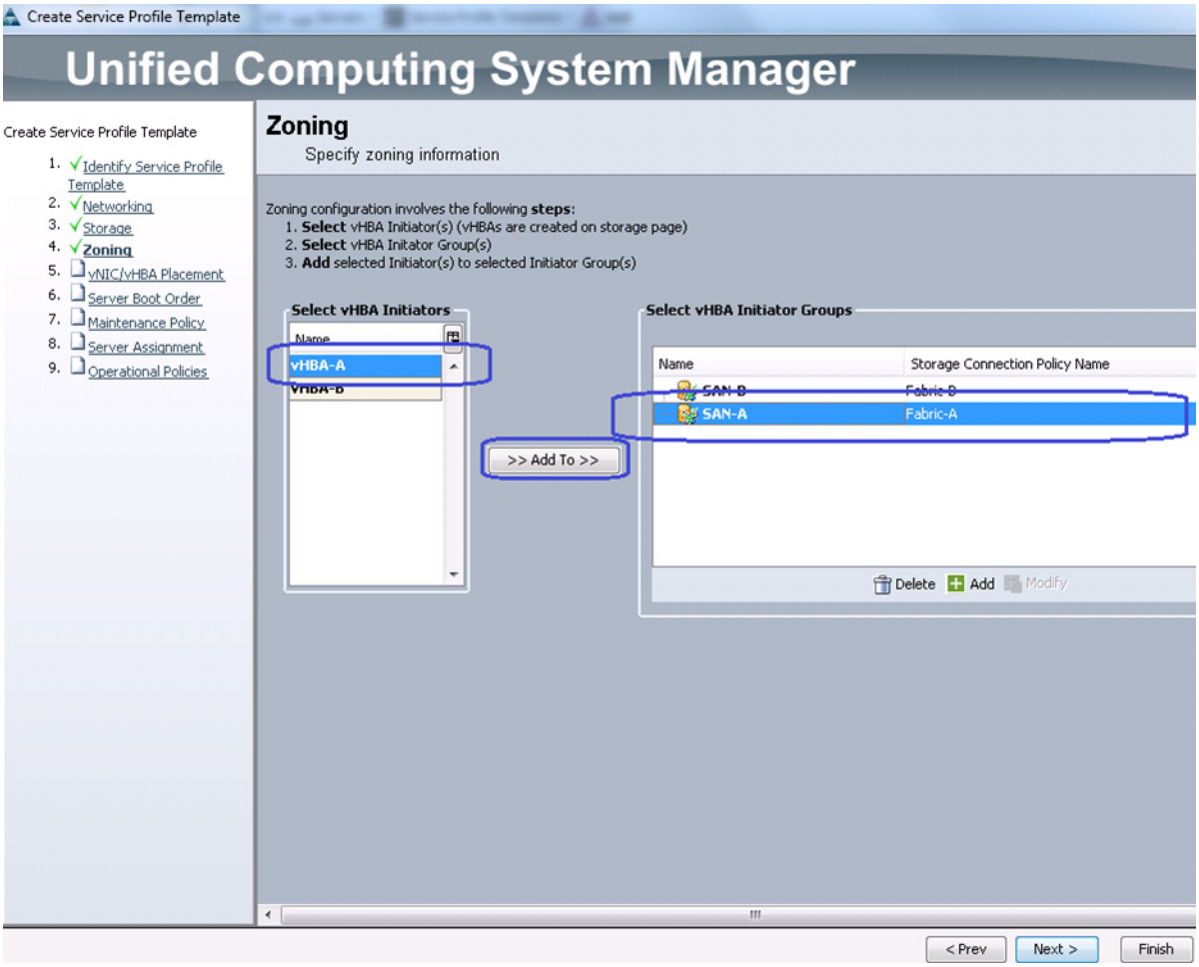

19. Repeat step 20 for fabric B. The end result looks as show in Figure 112. Click **Next**.

Figure 112 SAN Zoning Details

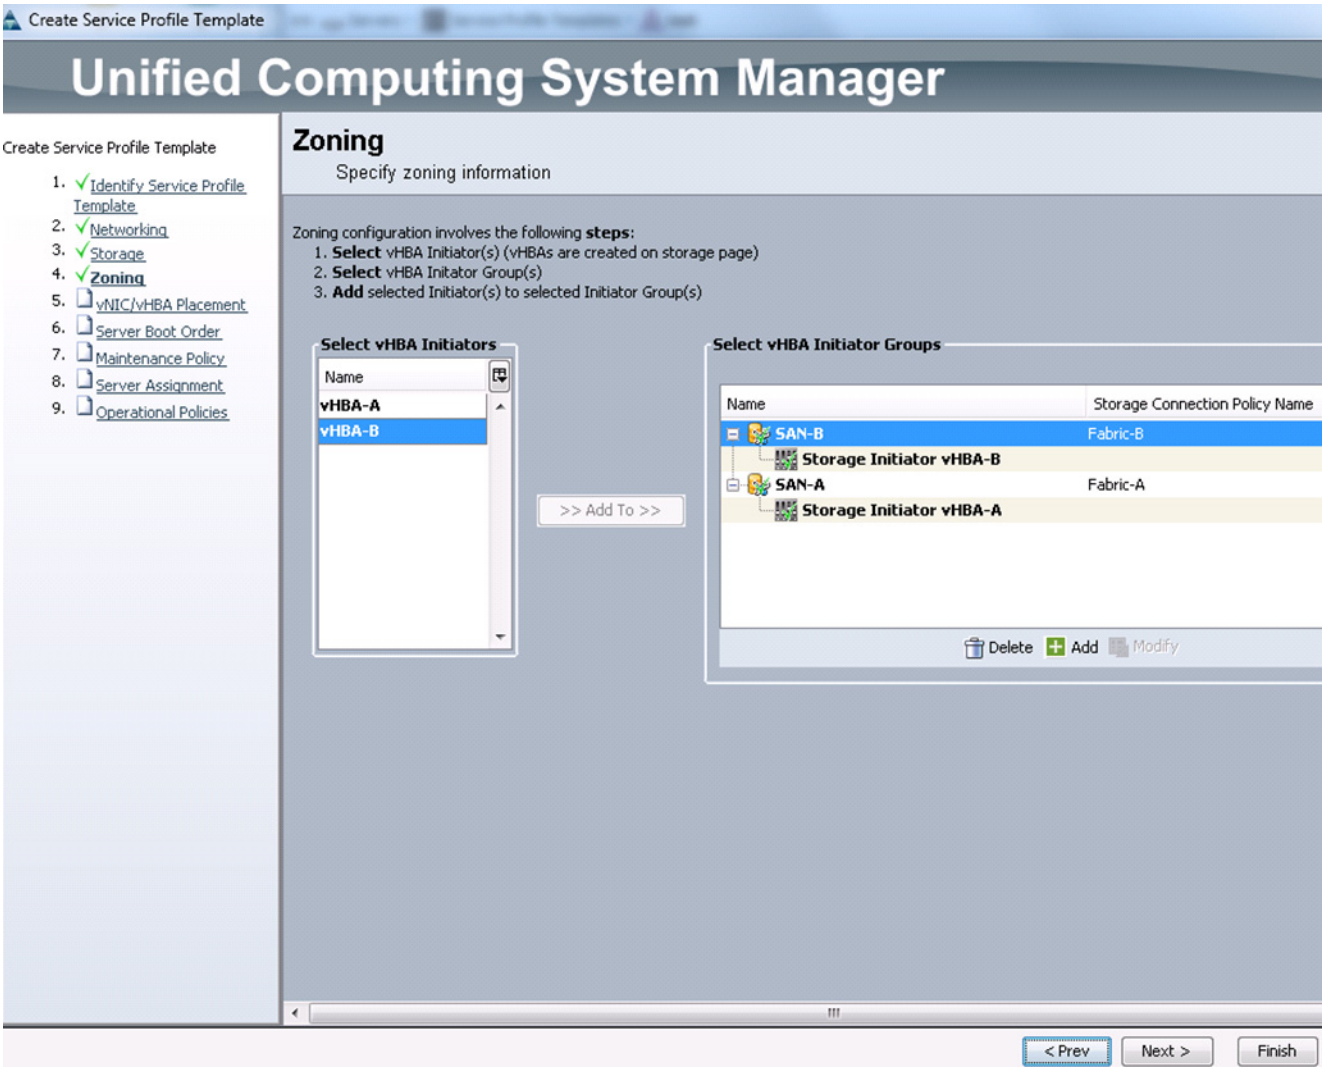

20. For the iSCSI-variant of the solution, select default configuration on Zoning. Click **Next**.
21. For FC-variant of the solution, choose Create a Specific Boot Policy from the Boot policy drop-down list. Choose Add CD-ROM in the Local Devices area as the first boot order choice. Click **vHBA-A** as the next choice in the vHBAs area, and name it as vHBA-A, keep the Type as **Primary**.

Figure 113 Primary SAN Boot Target

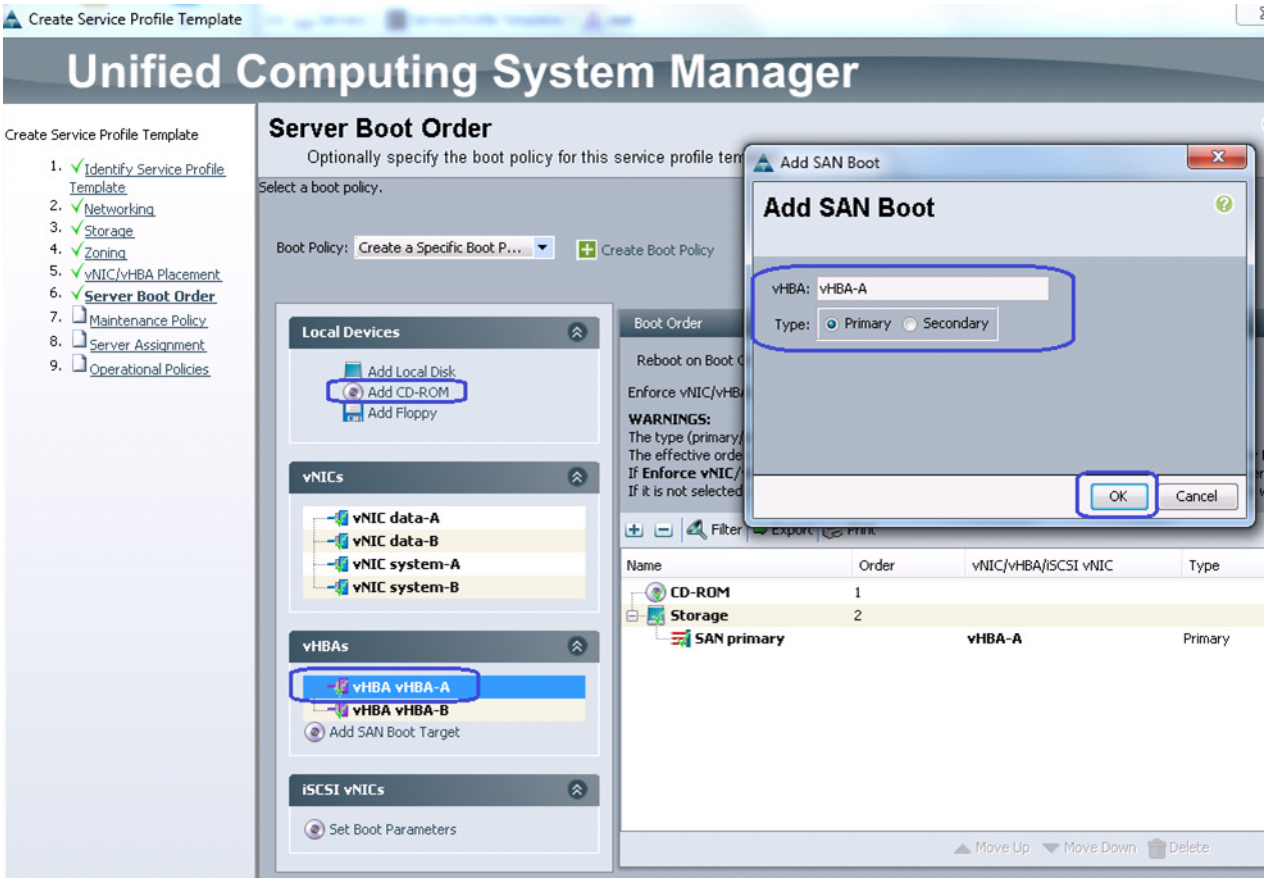

22. (FC-variant only) Similarly, select vHBA-B as the next (secondary) choice to boot from SAN. Once both vHBAs are added, make sure that Reboot on Boot Order Change and Enforce vNIC/vHBA/iSCSI name check boxes are checked. Click **Add SAN Boot Target** in the vHBAs area, and choose Add San Boot Target to SAN primary.

Figure 114 Secondary SAN Boot Target

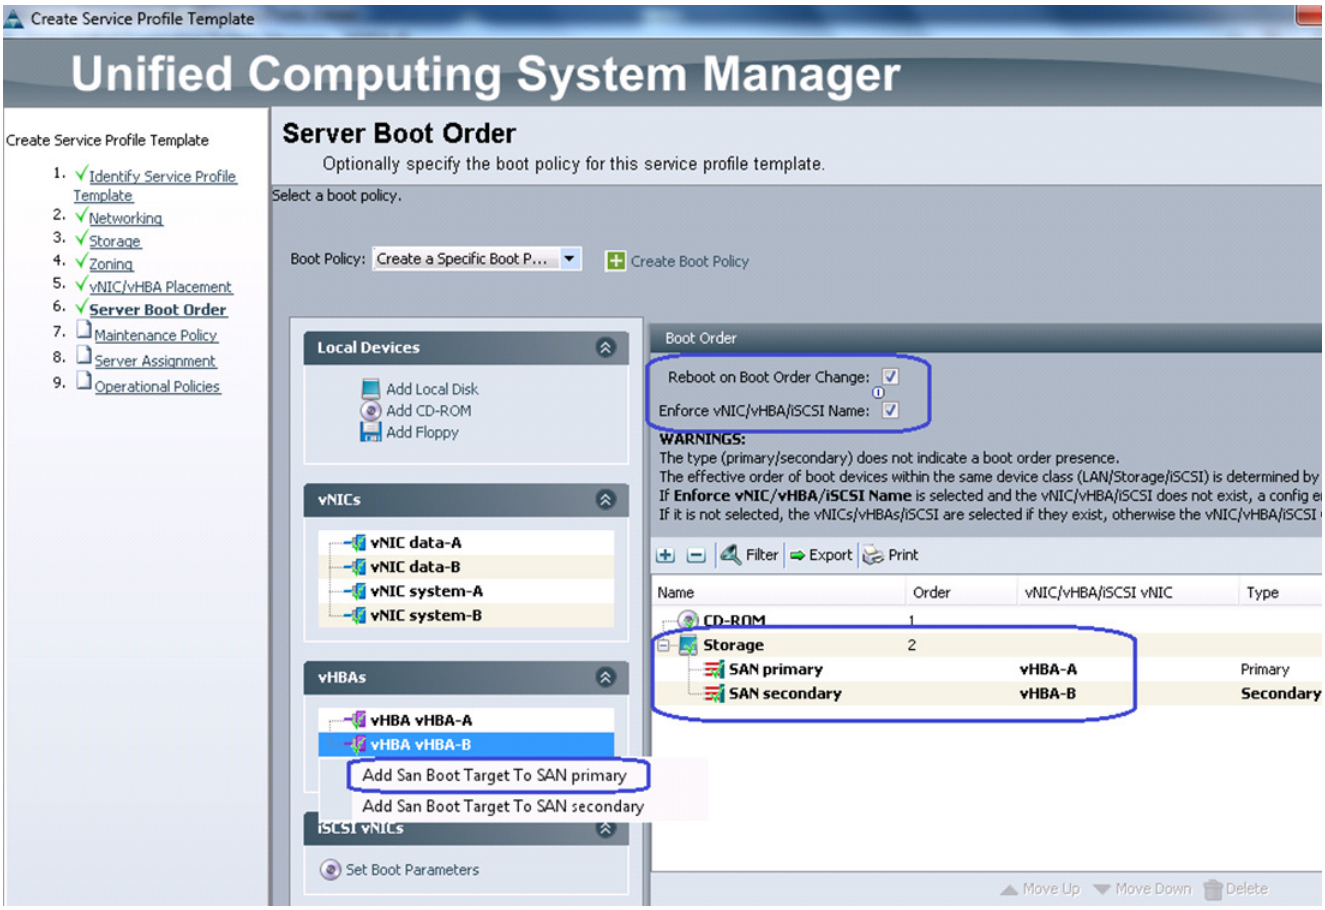

23. (FC-variant only) Enter the target WWPN of the VNX storage device (which can be obtained using "show flogi database" NXOS CLI command executed under **connect nxos {a|b}** shell). Keep the Target as Primary.

Figure 115 Specifying Boot Target WWPN

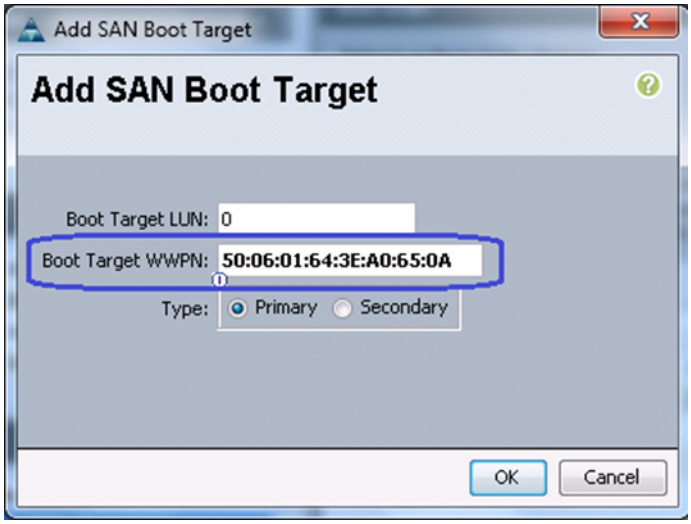

24. (FC-variant only) Repeat step 23 for the fabric B.

25. For the iSCSI-variant of the solution, choose Create a specific Boot Policy from the Boot Policy drop-down list. Check the check box Reboot on Boot Order Change. Click **Add CD-ROM** as the first boot choice. After that, iSCSI vNICs on fabric A and B as the next boot choices. Click **Set Boot Parameters** under the iSCSI vNICs area.

Figure 116 Specifying Boot Order for iSCSI variant Solution

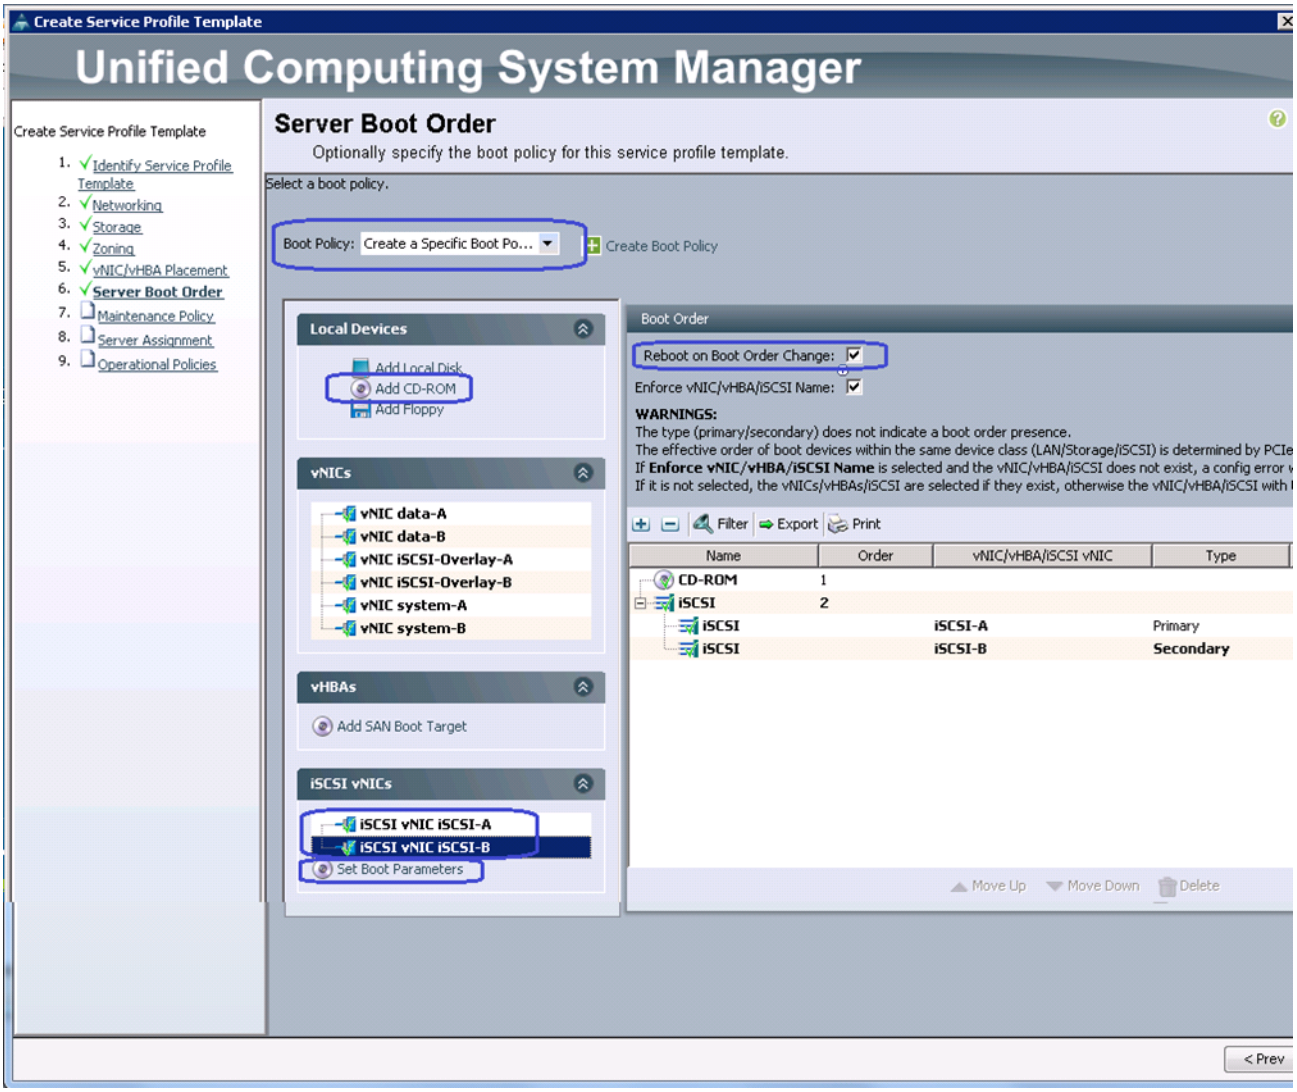

26. (iSCSI-variant only) For iSCSI vNIC on fabric A, choose the IQN pool name from Initiator Name Assignment drop-down list, choose iSCSI initiator IP Address Policy from the drop-down list. Click the **iSCSI Static Target Interface** radio button and click + to add the iSCSI target.

Figure 117 Setting iSCSI Boot Parameters for iSCSI-A

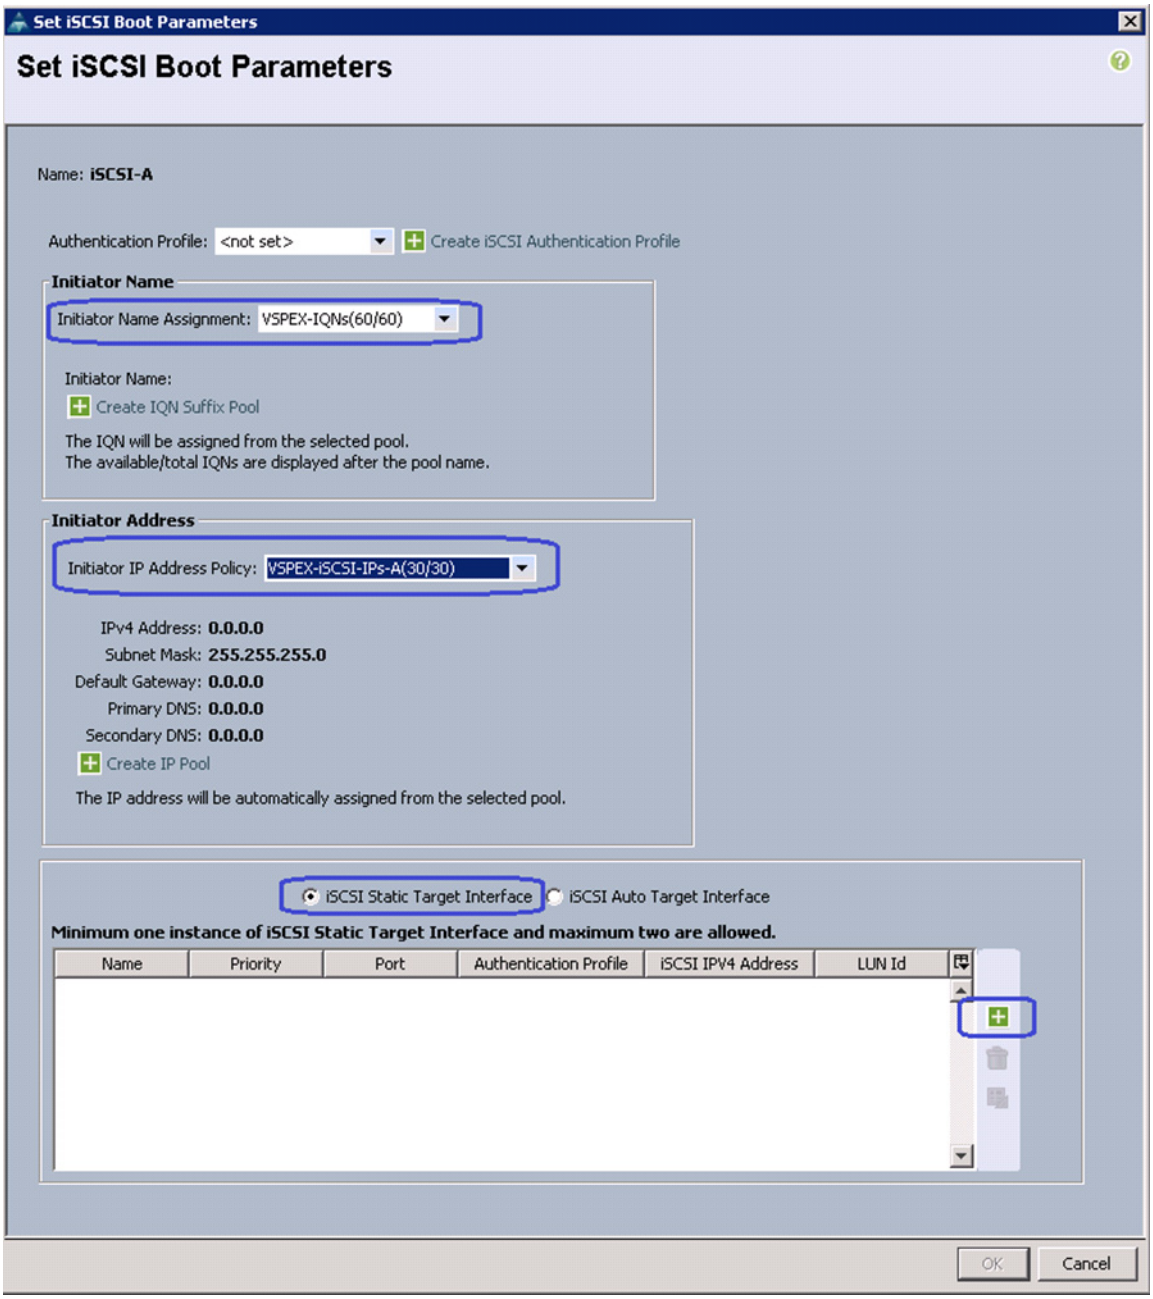

27. (iSCSI-variant only) In the Create iSCSI Static Target window, enter the iSCSI target name. Provide the target IPv4 address on the same subnet as initiator IP address and click **OK**. Let the LUN ID be 0 for now, we need to update the LUN ID for different datastores later per Service Profile basis.

Figure 118 Specifying iSCSI Static Target for iSCSI-A

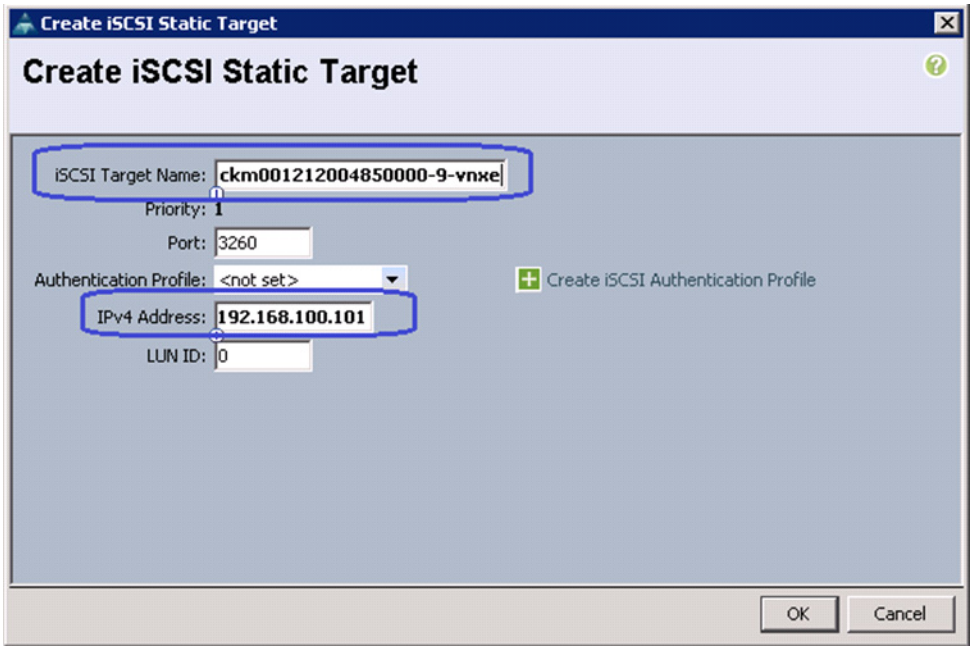

28. (iSCSI-variant only) Click **OK** and in the Server Boot Order window, click Set Boot Order Parameter again for the iSCSI vNIC on fabric B. IQN name for fabric B is taken from the same IQN pool, but the IPv4 initiator address comes from fabric B specific iSCSI initiator IP address pool. Click  to statically add the iSCSI target.

Figure 119 Setting iSCSI Boot Parameters for iSCSI-B

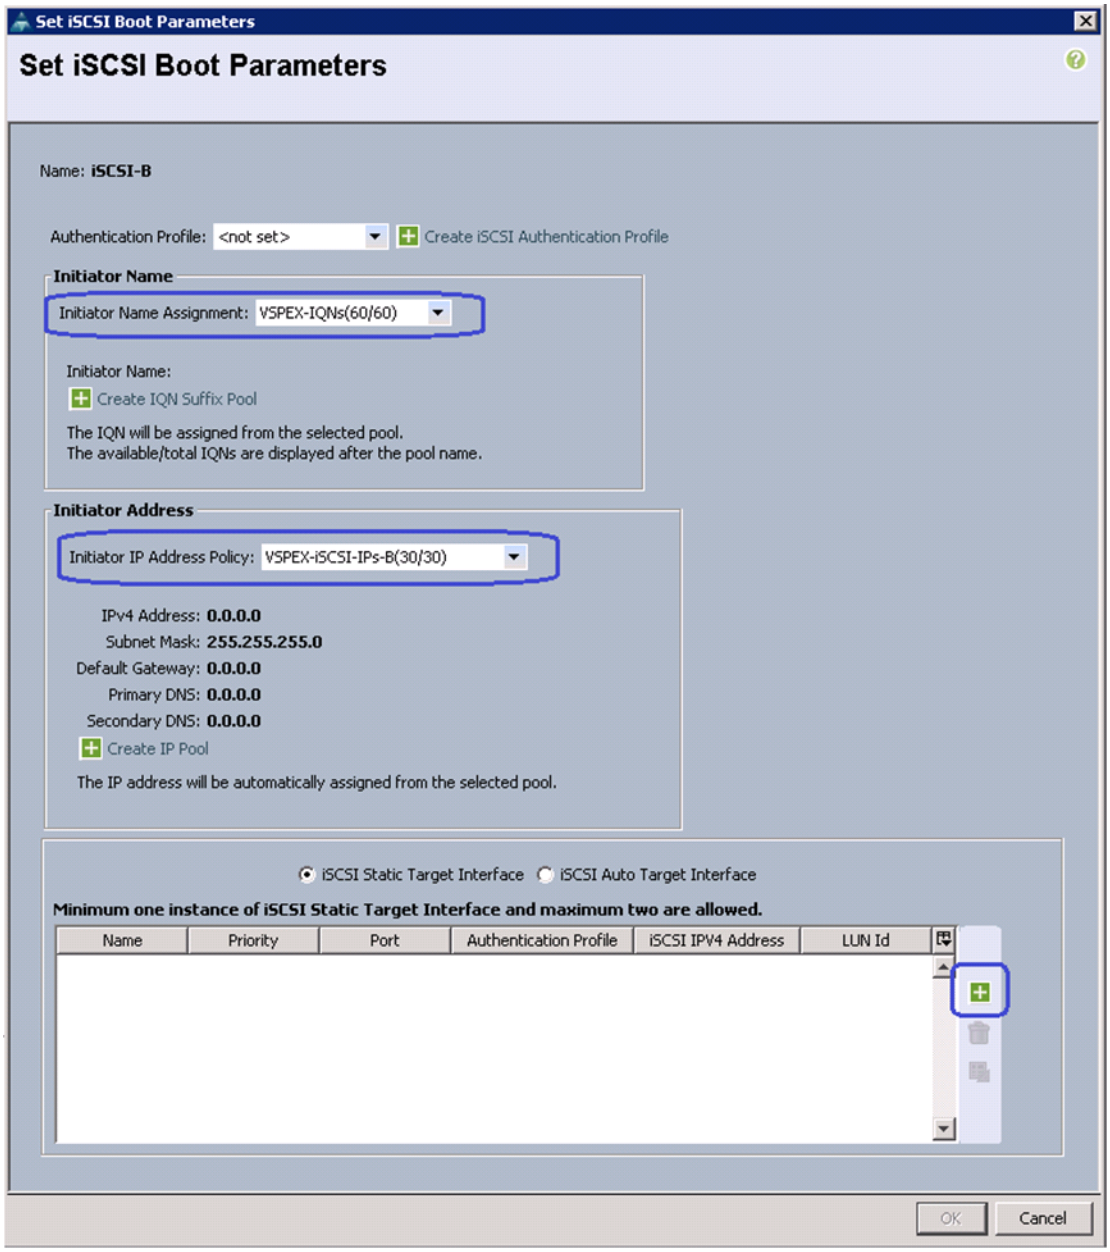

29. (iSCSI-variant only) Specify target IQN name and target IPv4 address. Make sure that the target IP address is from the same subnet as the initiator IP address, that is, pool's subnet.

Figure 120 Specifying iSCSI Static Target for iSCSI-B

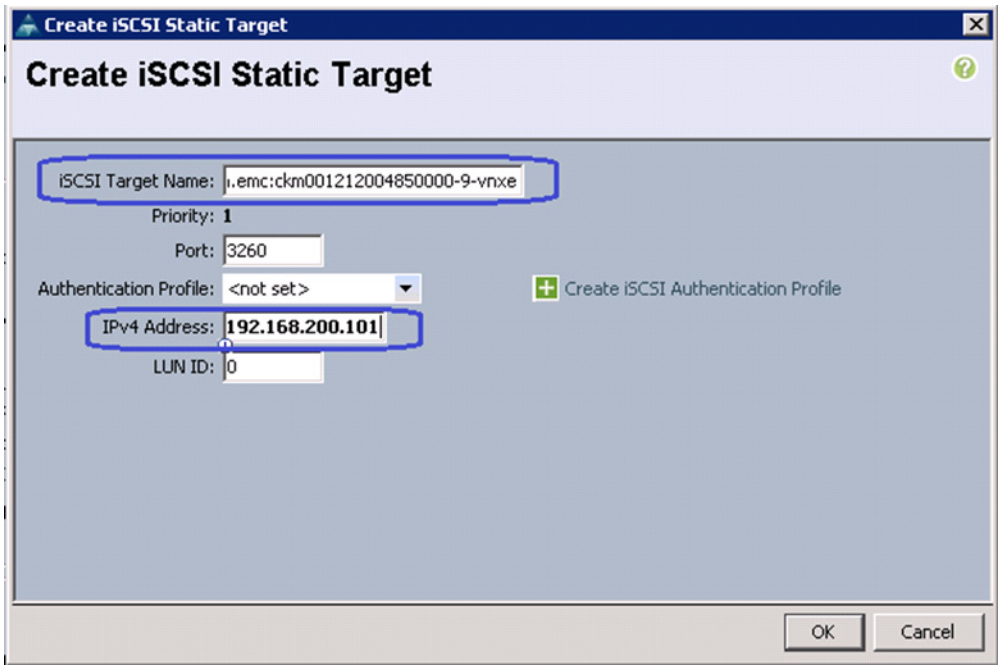

30. Click **OK** and then click **Next** to go to the Maintenance Policy window. Keep everything at default and click **Next** to go to the Server Assignment window.
31. Choose the previously created server pool from the Pool Assignment drop-down list. Click **Next**.

Figure 121 Assigning a Server Pool

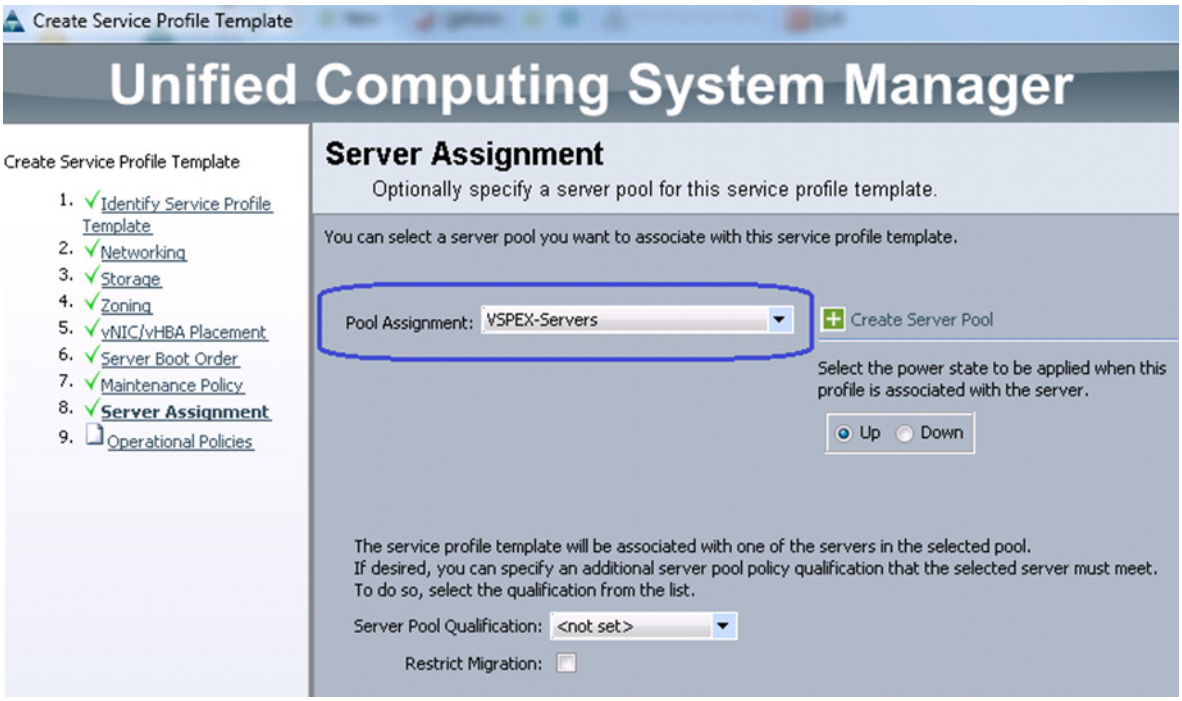

32. In the Operation Policies window, keep everything at default, and click **Finish** to deploy the Service Profile Template.

Figure 122 Operational Policies Window

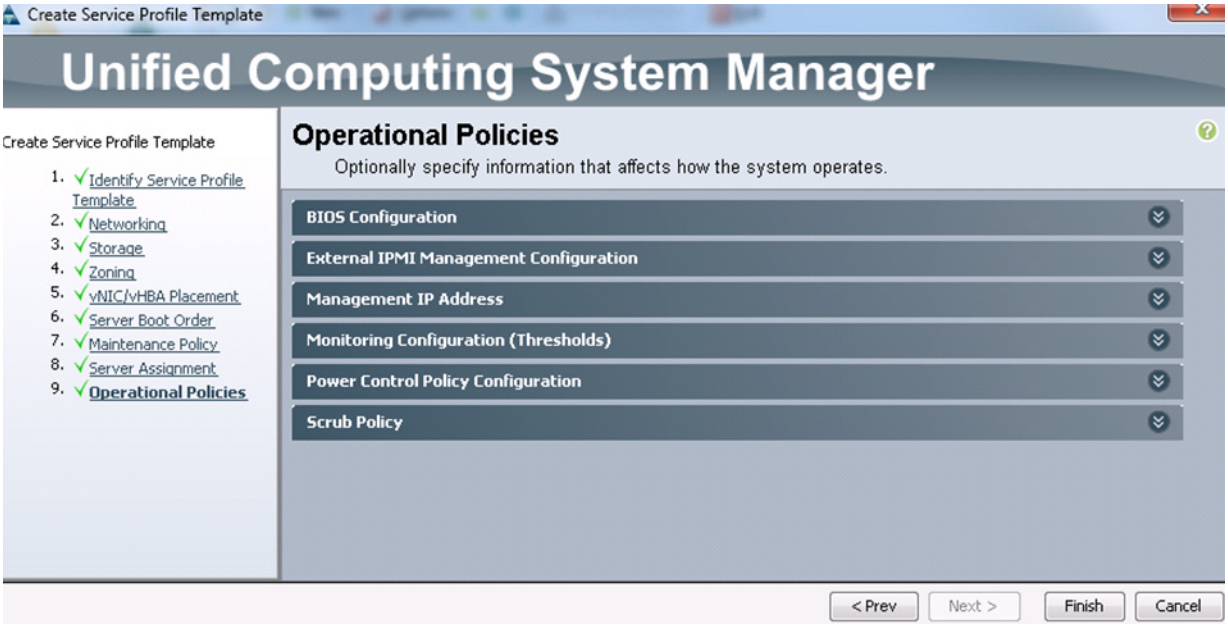

Instantiate Service Profiles from the Service Profile Template

This section details on instantiating service profiles from the service profile template created in the previous section. Follow these steps to instantiate service profiles:

1. From the **Servers** tab, expand **Servers > Service profiles > root**, and click **Create Service Profile from Template** on right pane os the window.

Figure 123 Creating Service Profiles from Template

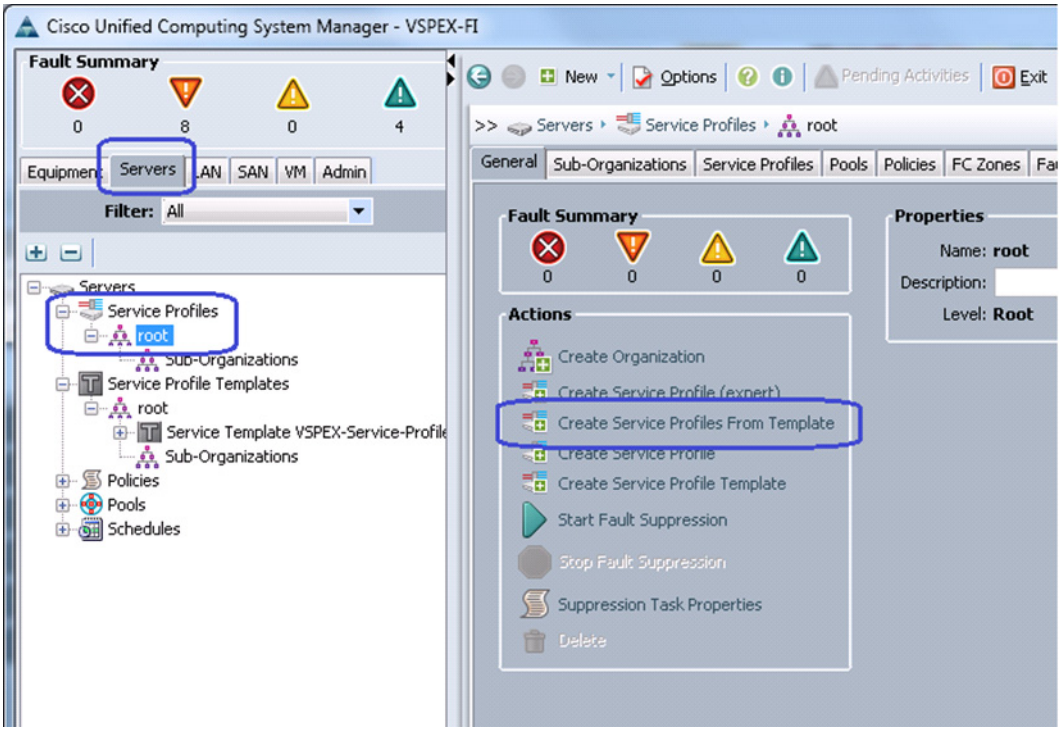

2. Enter the naming prefix, number of service profiles to be instantiated and choose the created service profile template from the drop-down list.

Figure 124 Specifying Details for Creating Service Profile

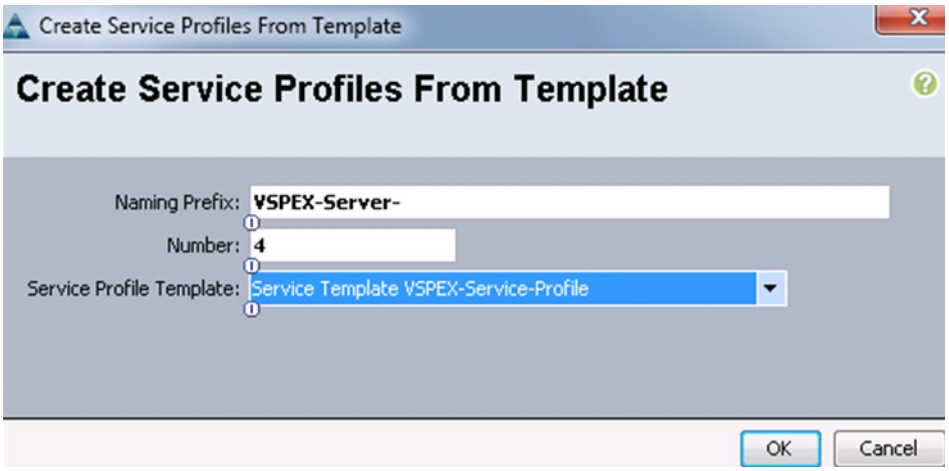

3. Four service profiles are created in this example.

Figure 125 Service Profiles Instantiated from a Service profile Template

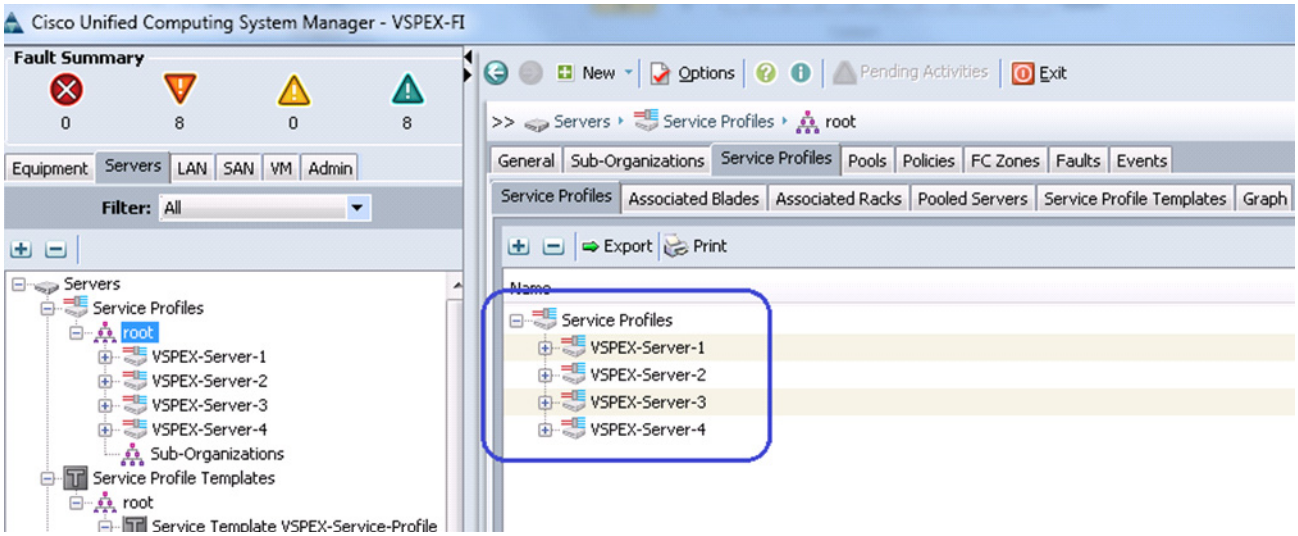

4. As the service profile template is assigned to a server pool, the service profiles instantiated from the template will be assigned to the individual server resource from the server pool as long as they are available. You can select the service profile and see its association state, and the associated server details.

Figure 126 Status Details of the Associated Servers

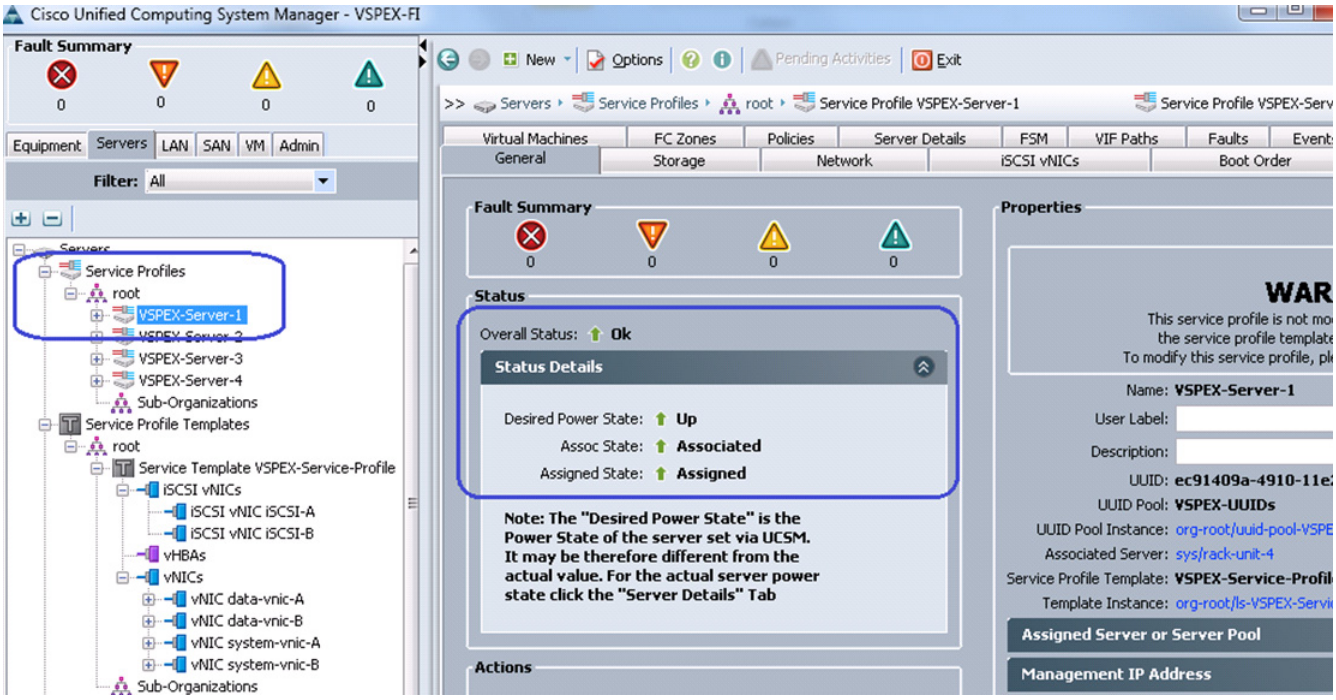

5. Eventually, all four servers will get associated. See the summary by clicking on **Servers** under the **Equipment** tab.

Figure 127 Summary of the Associated Servers

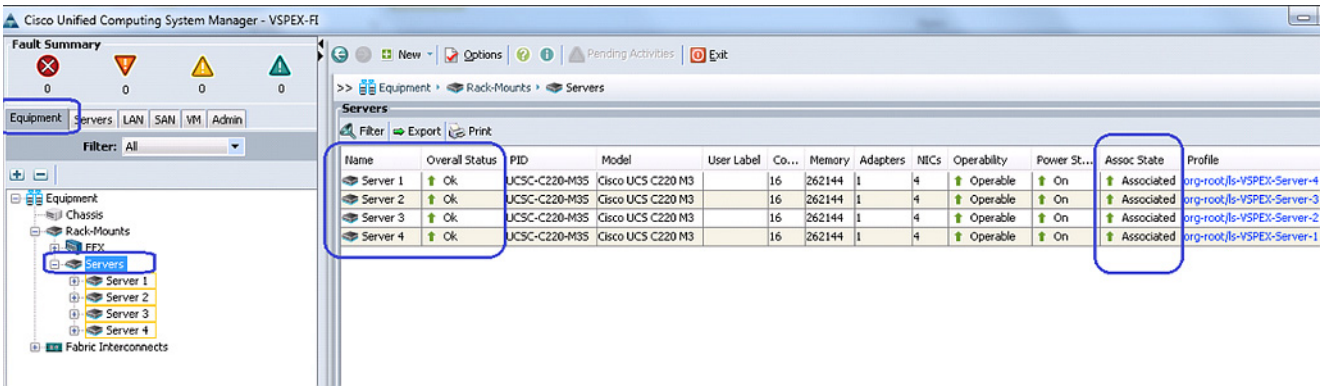

Note Note that we have not yet carved out specific datastore to install ESXi hypervisor OS image on the VNXe storage array. We need specific IQN and iSCSI initiator IP address to allow access to the datastore, and hence we needed to configure the service profile before we can carve out the space for each ESXi server on the storage pool.

Configure Datastores for ESXi images – FC-variant

This section details on creating FC accessible datastores for the ESXi boot image per server basis. This includes three steps:

1. Configure Storage Pool
2. Register hosts
3. Configure Storage Groups

A: Configure Storage Pool

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

1. Login to the EMC VNX5300 Unisphere GUI, click **Storage** tab. Choose **Storage Configuration** > **Storage Pools**. Click the **RAID Groups** tab and click **Create** to create RAID groups.

Figure 128 Creating RAID Groups

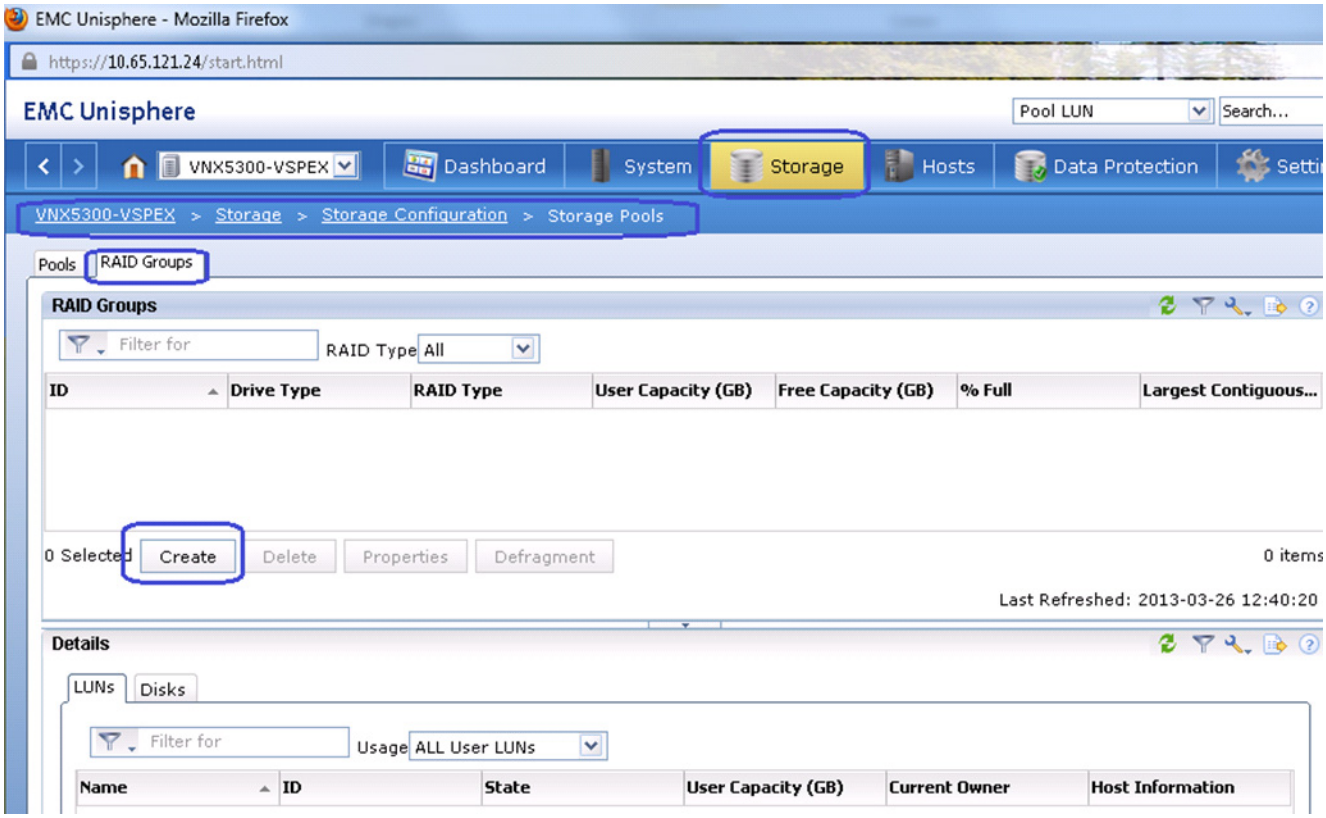

2. By default, for RAID5 RAID Group, the system would select five SAS disks. Click the **Manual** radio button, and click **Select** to manually select/deselect disks.

Figure 129 Specifying RAID Group Details

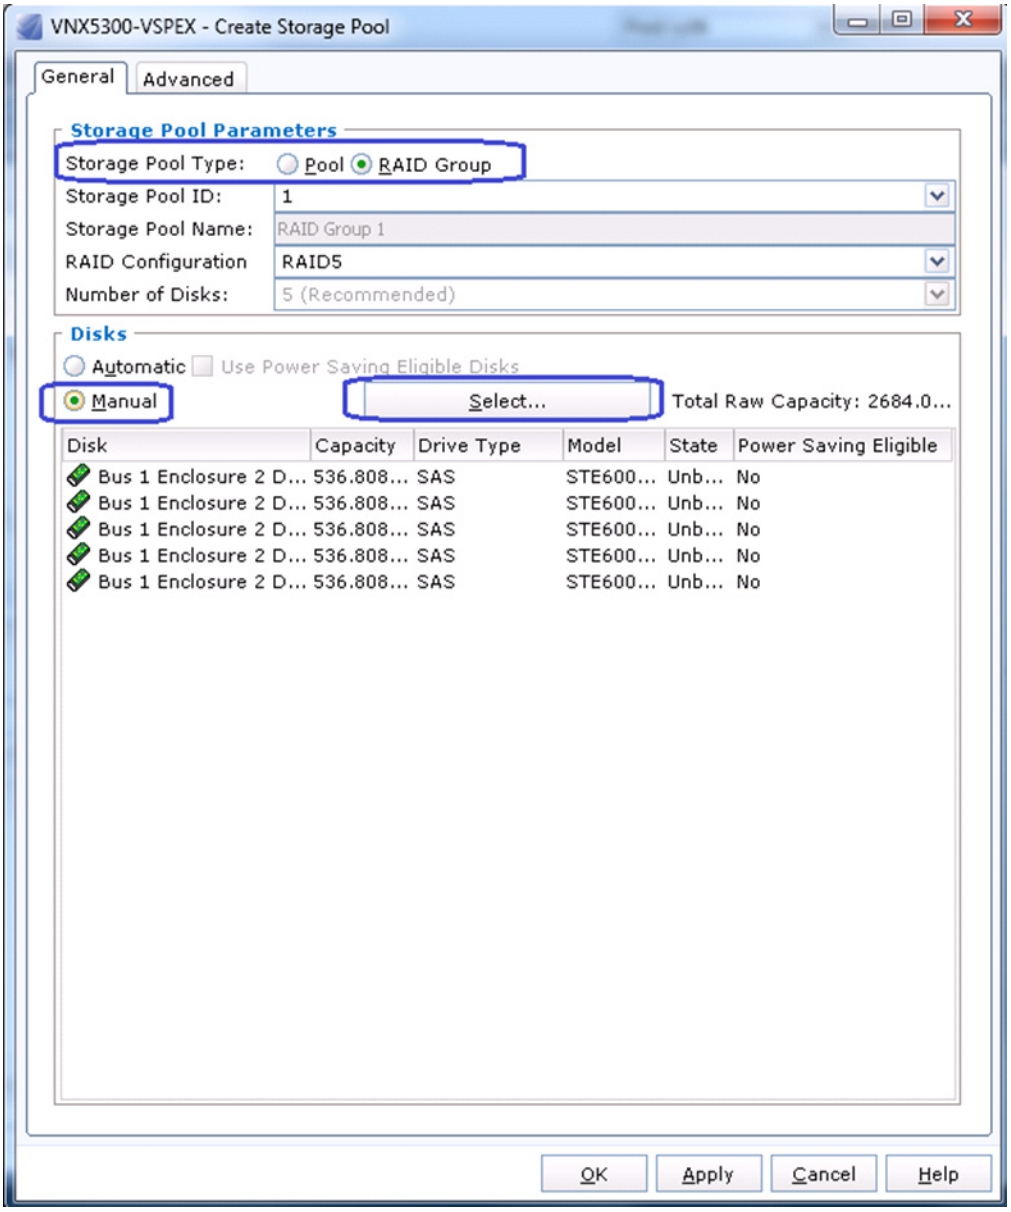

3. Deselect last two disks by clicking on  from the five disks. Click **OK** twice.

Figure 130 Removing the Disks from the Selected Disks List

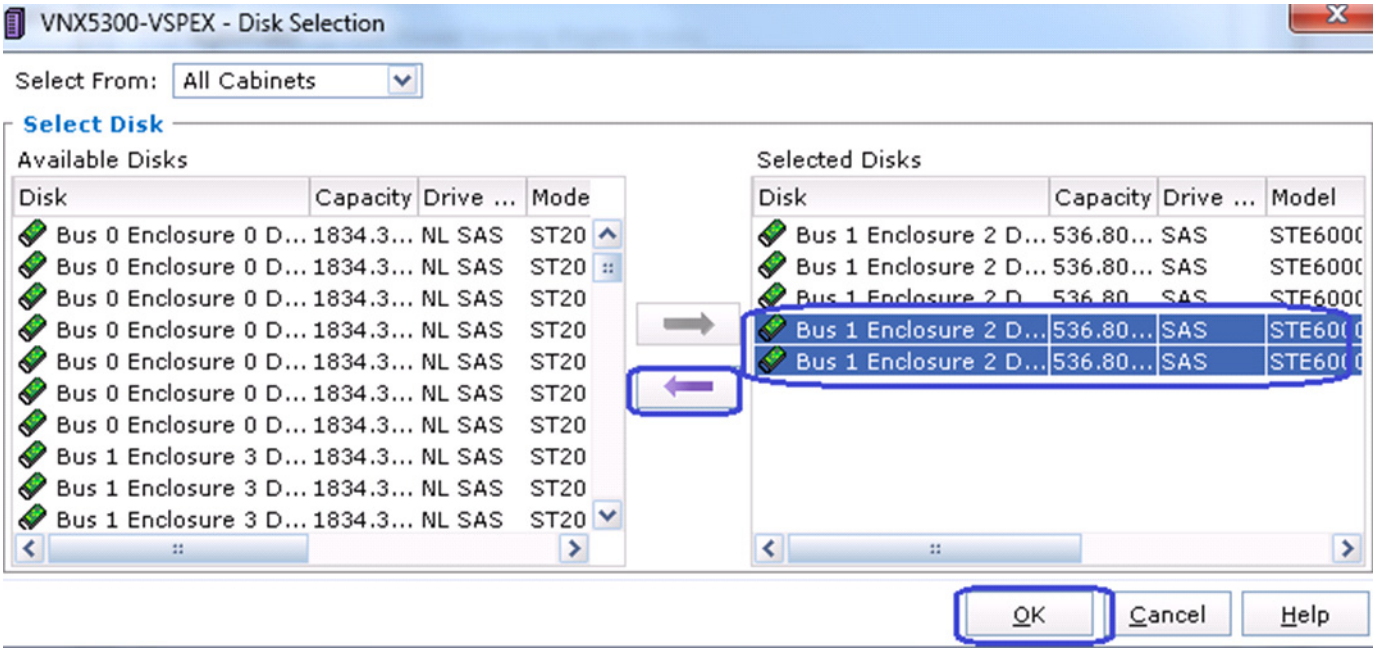

4. Select the newly created RAID group, right-click and choose Create LUN.

Figure 131 Creating LUN for the Created RAID Group

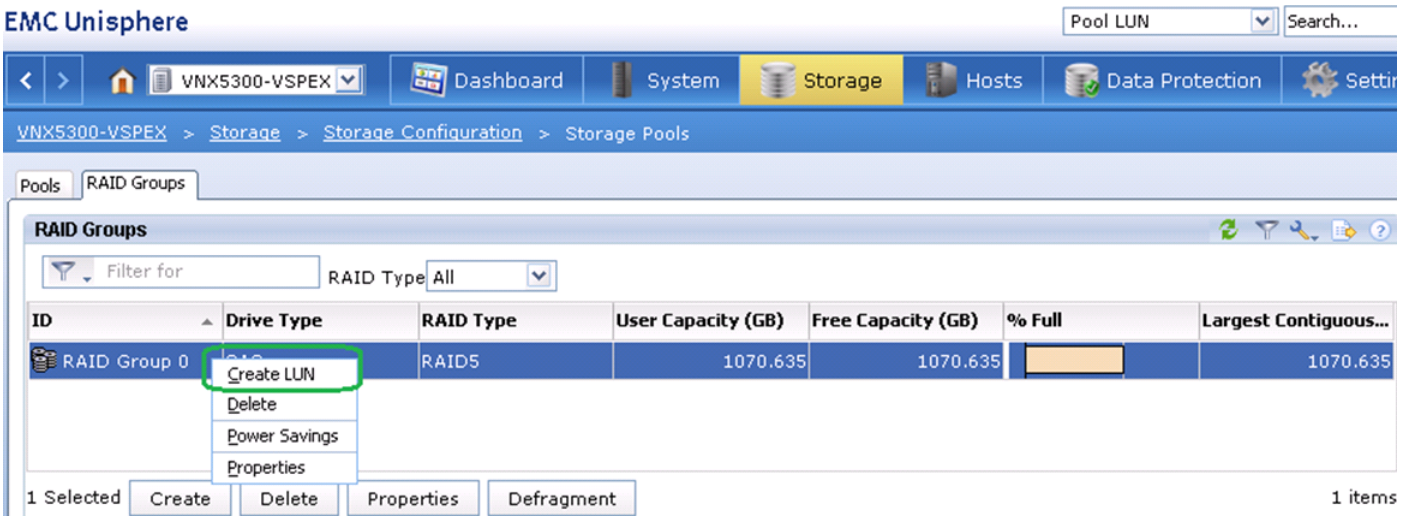

5. Choose the value 5 from Number of LUNs to create drop-down list for create 5 LUNs for five ESXi hosts, with 50 GB capacity each.

Figure 132 Specifying Details for Creating LUN

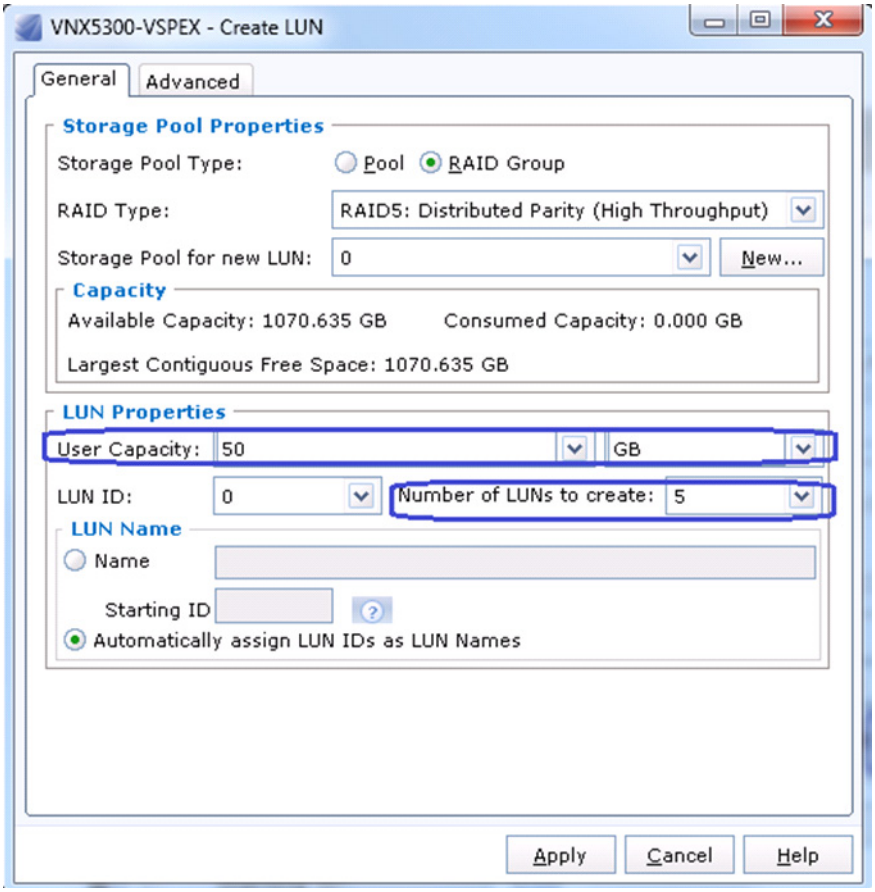

Register Hosts

As soon as the service profiles are associated in UCS Manager, the vHBAs will do flogi in the network and the SAN initiators will be identified by the VNX storage array. Follow these steps to register the hosts identified by the WWPN of the server.

1. In the Unisphere GUI, click the **Hosts** tab, and choose Initiators. Select the first unregistered initiator and click **Register**.

Figure 133 Registering Initiators in EMC Unisphere

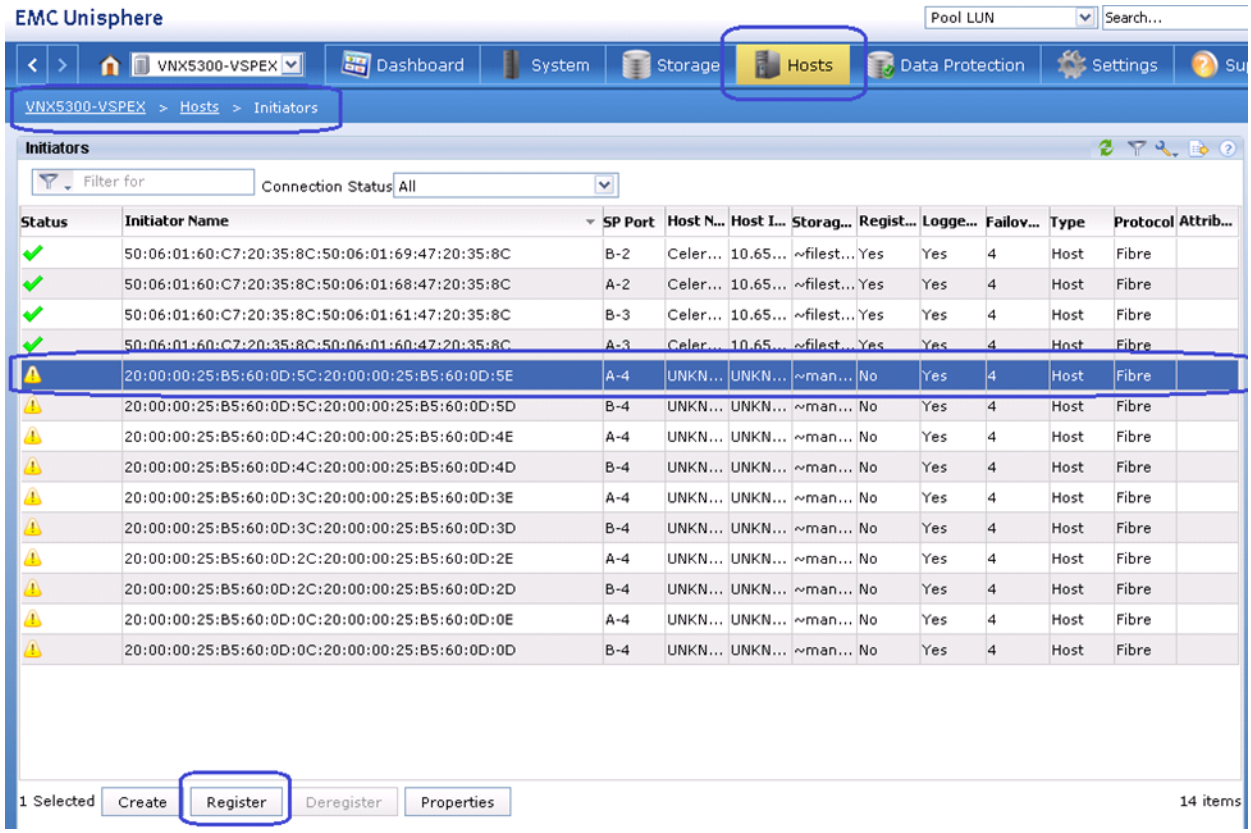

2. From the Cisco UCS Manager GUI, click the **Servers** tab, expand **Servers > Service Profiles > root**, select a specific service profile, and choose vHBAs. WWPN identities are listed. Using this IDs, you can associate WWPN to the server.

Figure 134 Associating WWPN to the Server

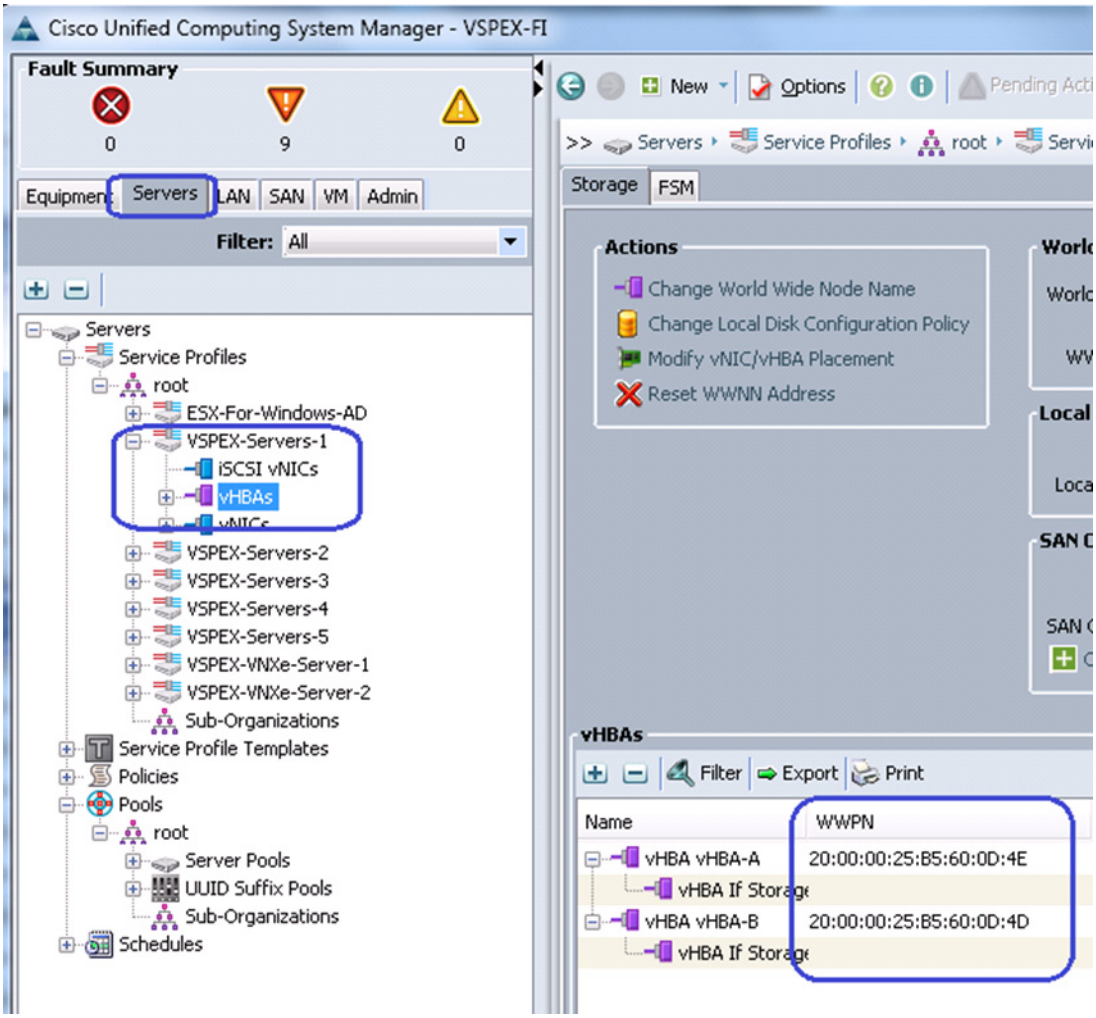

3. In the Register Initiator Record wizard, choose the Initiator Type as CLARiiON/VNX from the drop-down list and choose Failover Mode as failover mode 1 from the drop-down list. Click the **New Host** radio button, enter the hostname and (future) management IP address of the host. Click **OK**.

Figure 135 Registering Initiator Record for New Host

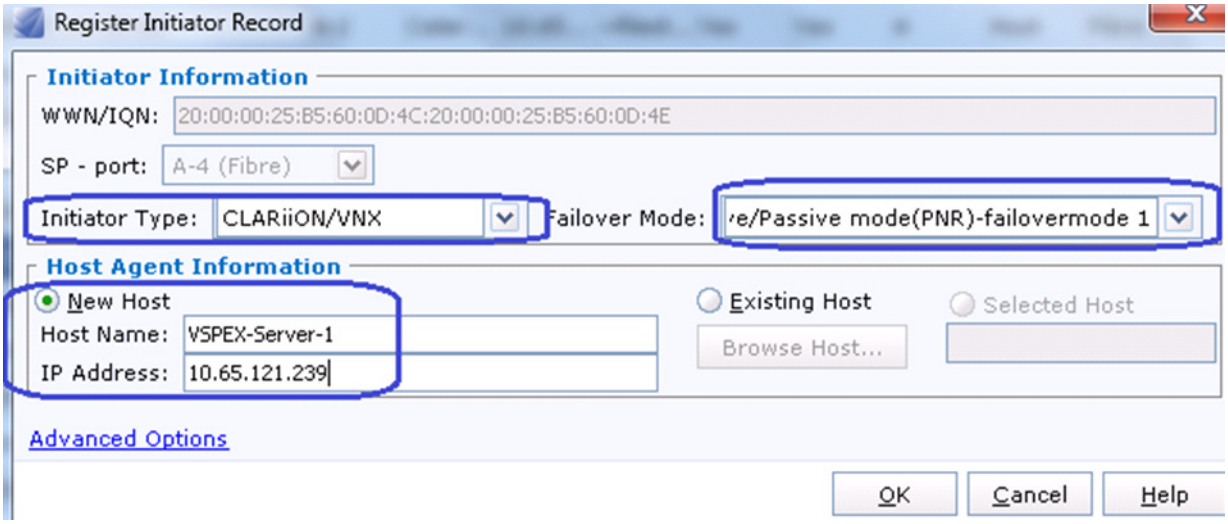

4. Select the second vHBA's WWPN from the same server, click **Register.** Click the **Existing Host** radio button in the Host Agent Information area. Click **Browse Host** to select the host.

Figure 136 Registering Initiator Record for the Second vHBA's WWPN

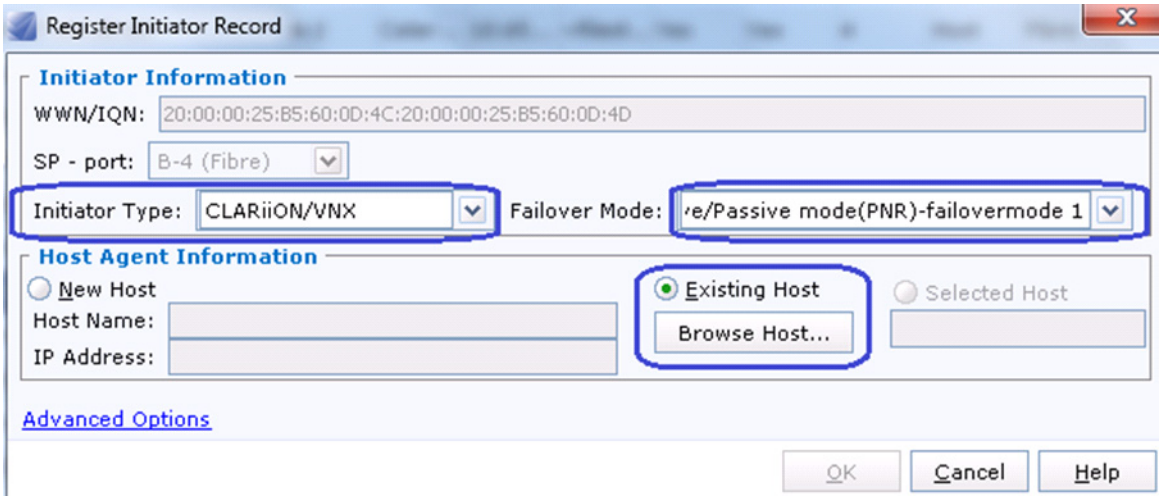

5. Select the previously registered host, and click **OK**.

Figure 137 Connectivity Status of the Registered Host

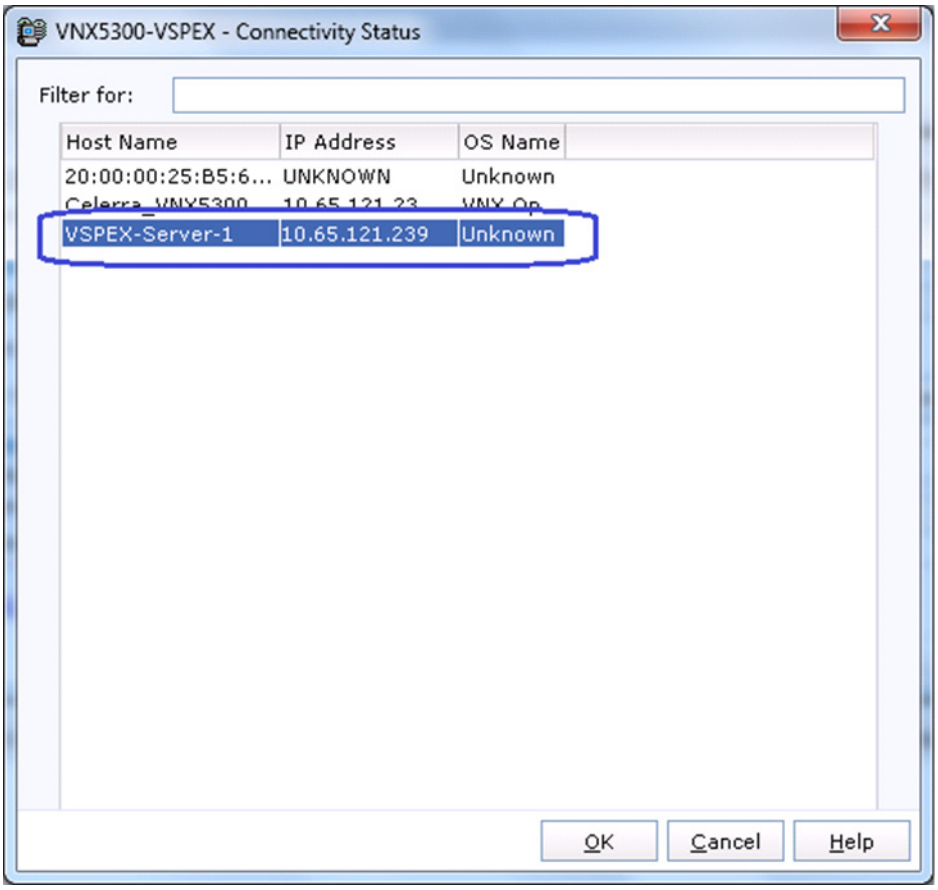

6. Repeat these steps for all the servers in the group. The final page is shown in the Figure 138.

Figure 138 Initiators Window Showing the Connection Status of All the Registered Servers

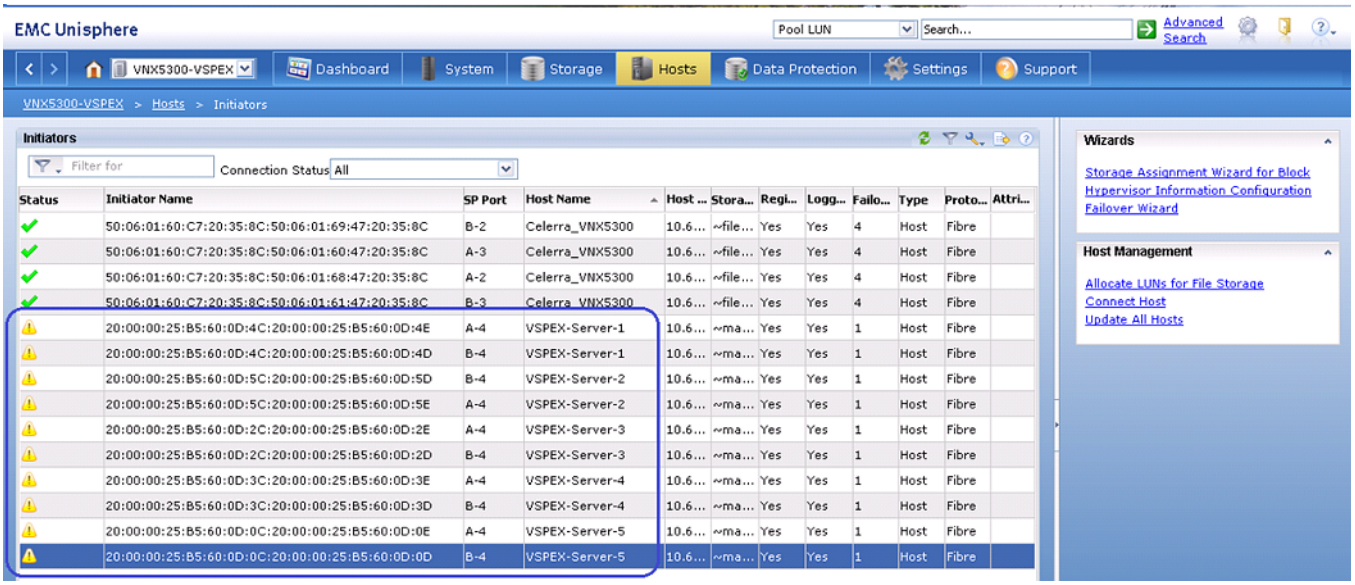

Configure Storage Groups

Now that hosts as well as LUNs are created on the VNX storage array, we need to create storage groups to assign access to LUNs for various hosts. Boot LUN will be dedicated to a specific server. Follow these steps to configure storage groups:

1. Click the **Hosts** tab in the EMC VNX Unisphere GUI, choose Storage Groups. Click **Create** to create a new storage group.

Figure 139 Creating a Storage Group

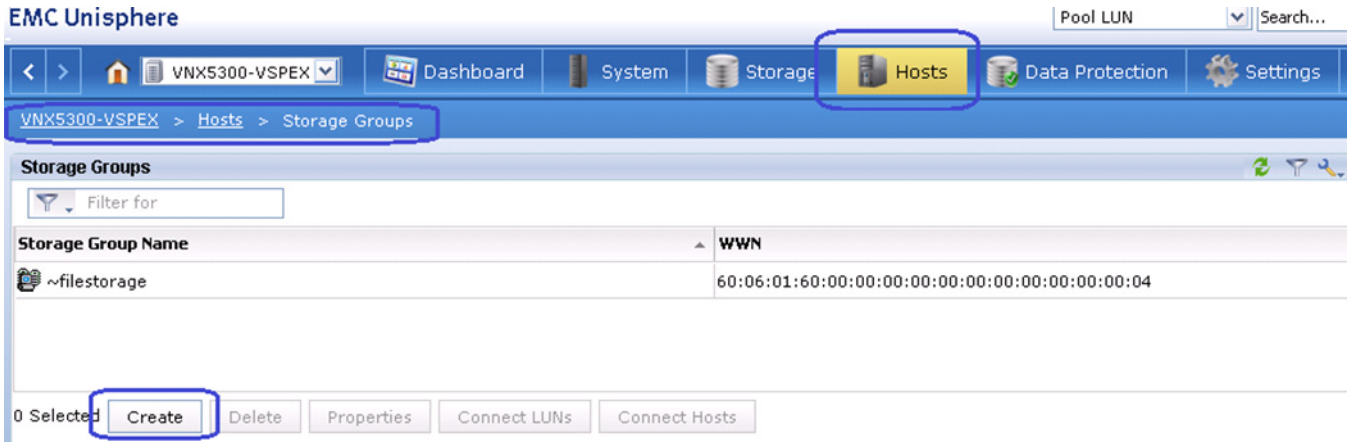

2. Enter the storage group name in the Storage Group Name field. Click **OK**.

Figure 140 Specifying Storage Group Details

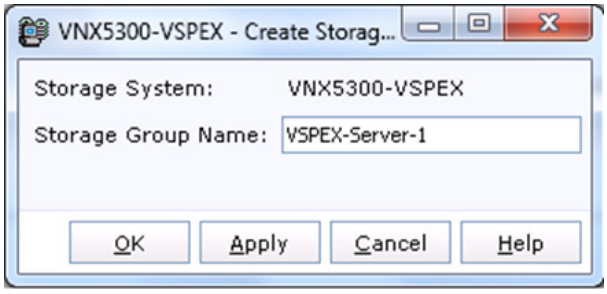

3. A confirmation message window pops up. The system prompts you to create LUNs and connect hosts. Click **Yes**.

Figure 141 Confirmation on Adding LUNs or Connecting Hosts

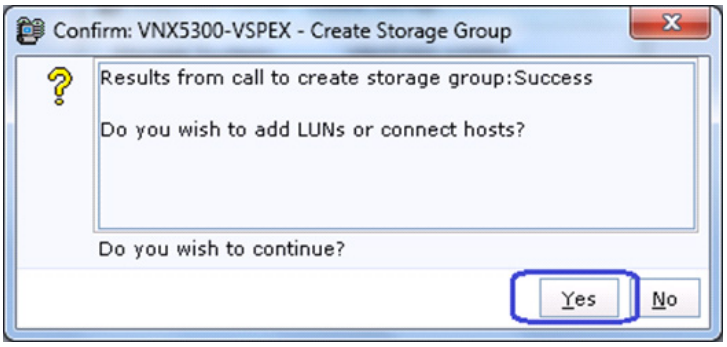

4. From the **LUNs** tab of the wizard, select a single LUN and click **Add**.

Figure 142 Adding LUN to the Storage Group

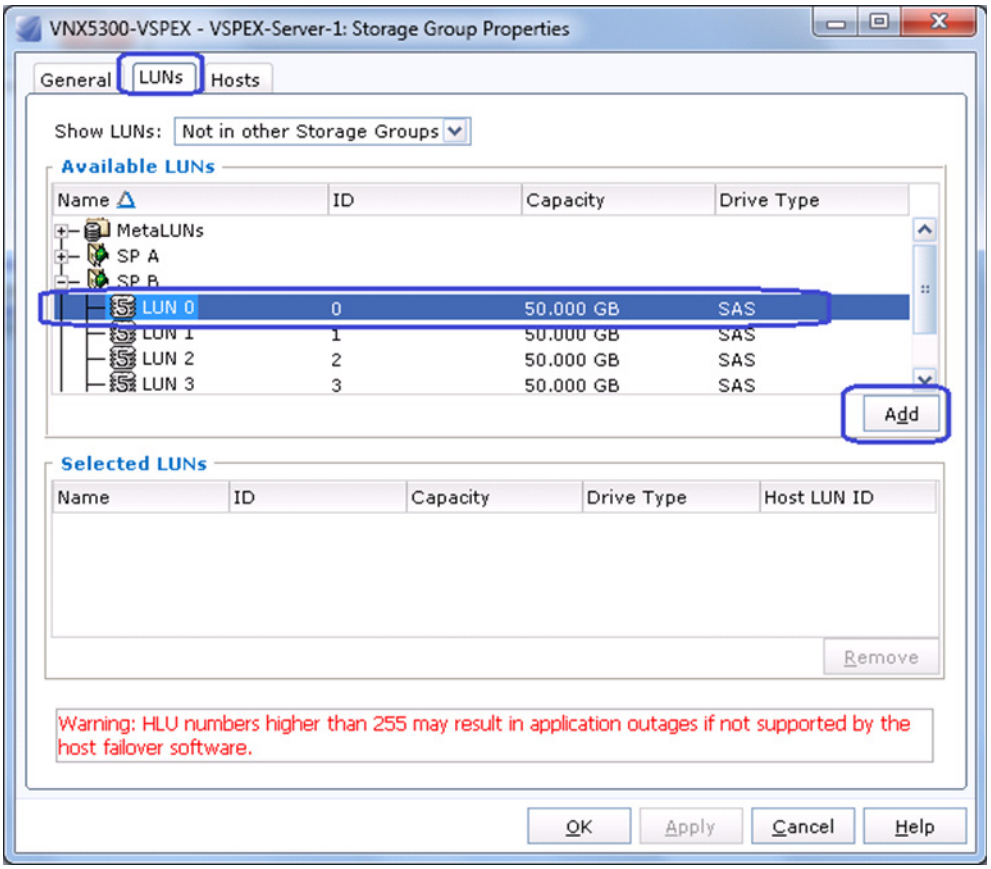

5. Click the **Hosts** tab in the wizard, and select a single server and click → to add to the storage group. Click **OK** to deploy the storage group.

Figure 143 Selecting Hosts to be Connected

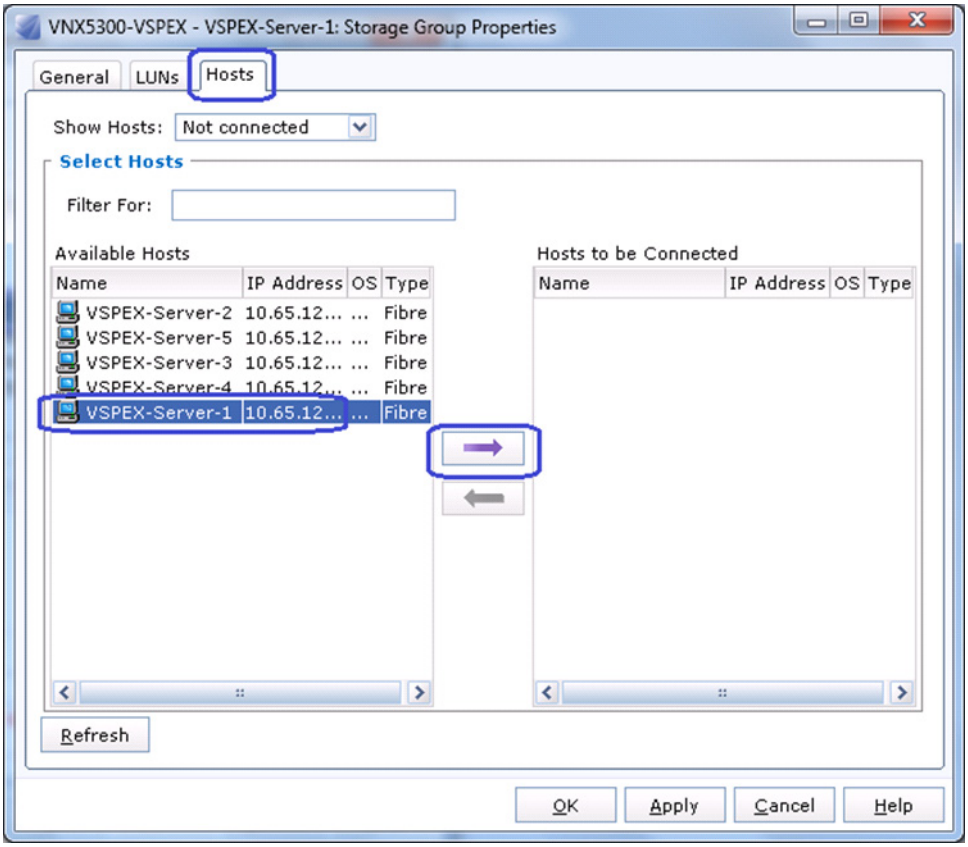

6. Repeat these steps for all the five servers. After adding all the servers to the storage group, you can see all the servers listed in the storage group as shown in Figure 144.

Figure 144 Storage Group Window Showing All the Added Servers

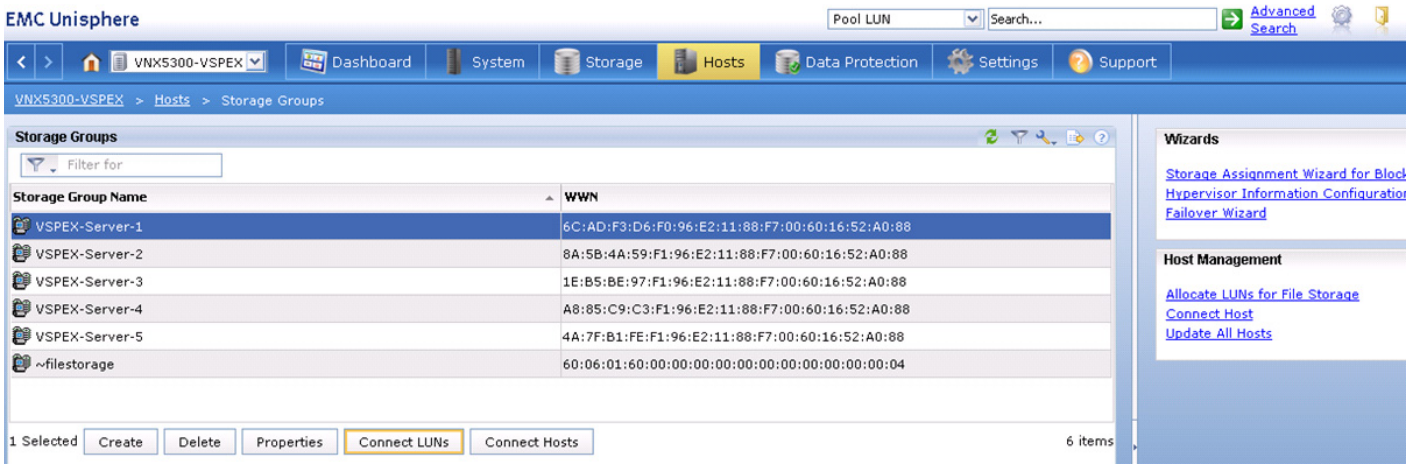

Configure Datastores for ESXi images – iSCSI-variant

This section details on creating iSCSI accessible datastore for the ESXi boot image per server basis. This includes two steps:

1. Add ESXi host information on storage
2. Create boot LUNs (datastores) for each server

Add ESXi Host Information on Storage

Follow these steps to add ESXi host information on the VNXe storage array. As the ESXi boot image is not yet installed and storage array and hypervisors are not integrated with vCenter, the ESXi host information needs to be added statically, instead of dynamically discovering the hosts through vCenter.

1. From the VNXe Unisphere web interface, click the **Hosts** tab and then choose Hosts.

Figure 145 Selecting Hosts to Add ESXi Host Information on VNXe Storage Array

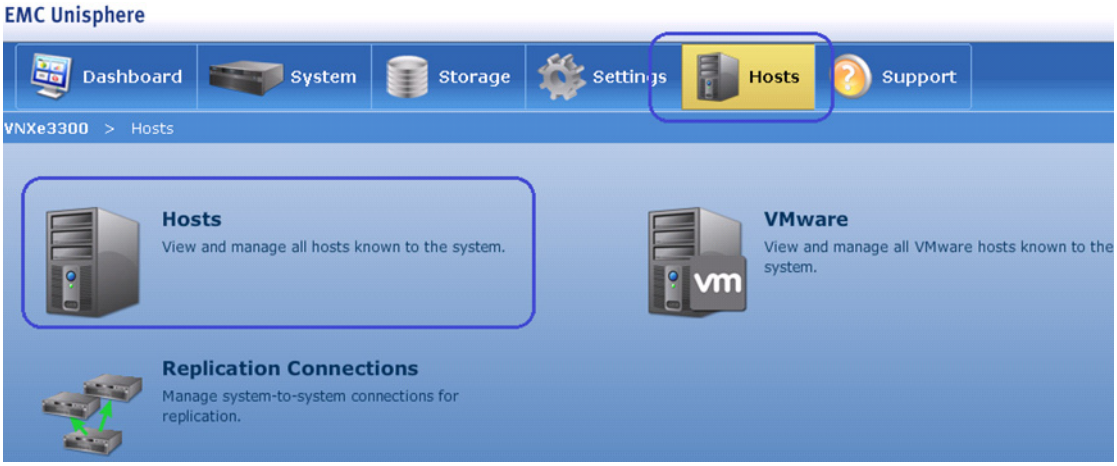

2. Click **Create Host** in the Hosts wizard.

Figure 146 Creating Host in EMC Unisphere

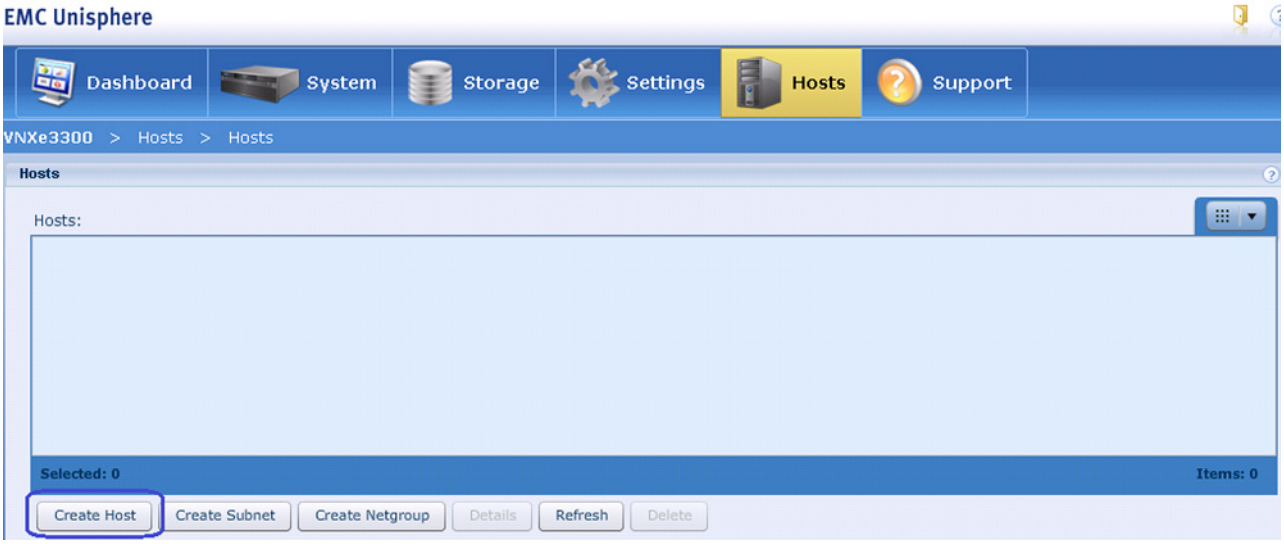

3. Enter the name of the host in the Name field and (optional) description. Click **Next**.

Figure 147 Specifying Details for Creating Host

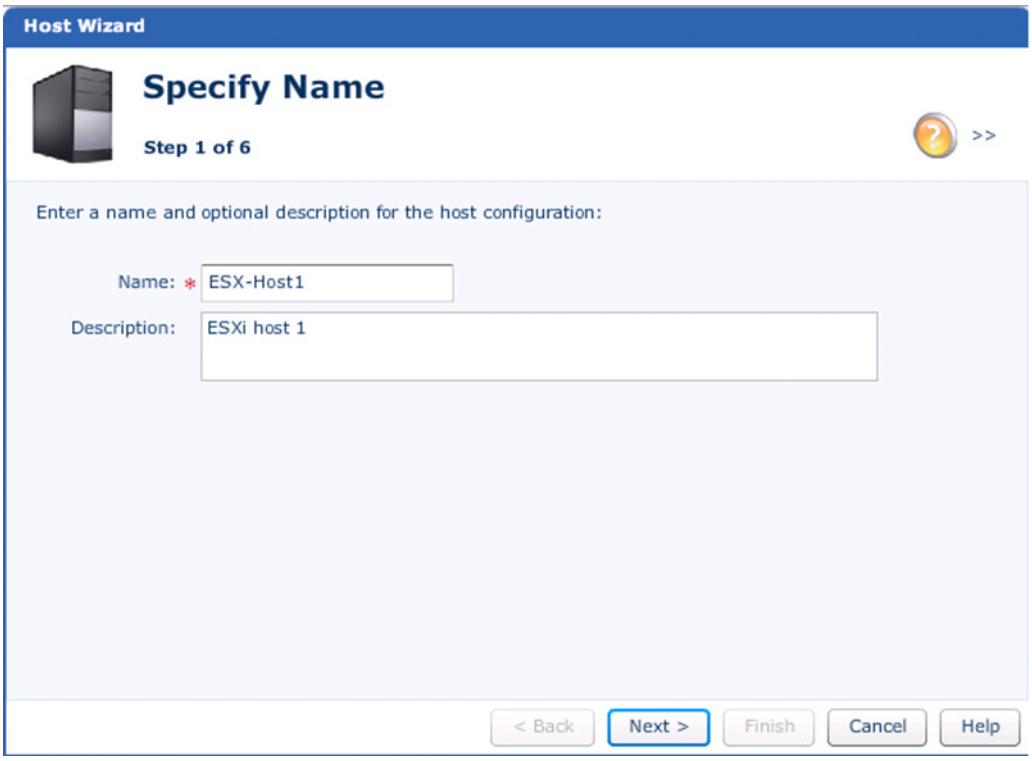

4. Choose the host operating system as VMware ESX from the drop-down list. This will generate an alert, which you can ignore. Click **Next**.

Figure 148 Specifying Operating System for Creating Host

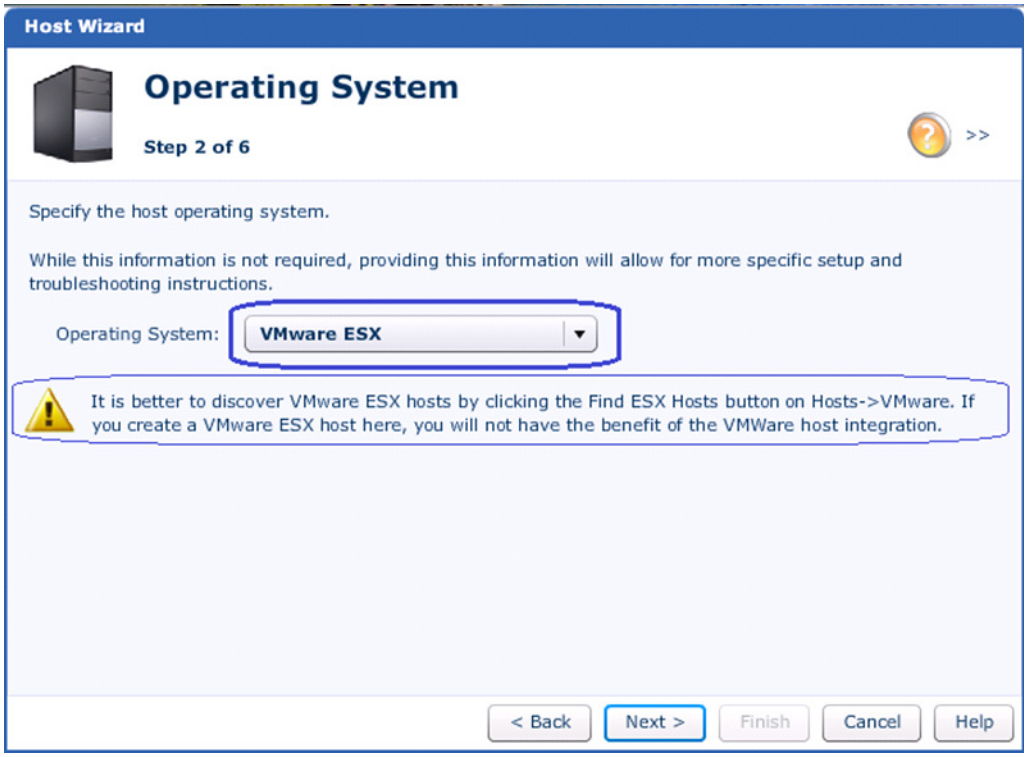

5. Provide iSCSI initiator IP address of the host's iSCSI vNIC. Click **Next**.

Figure 149 Specifying iSCSI Initiator IP Address

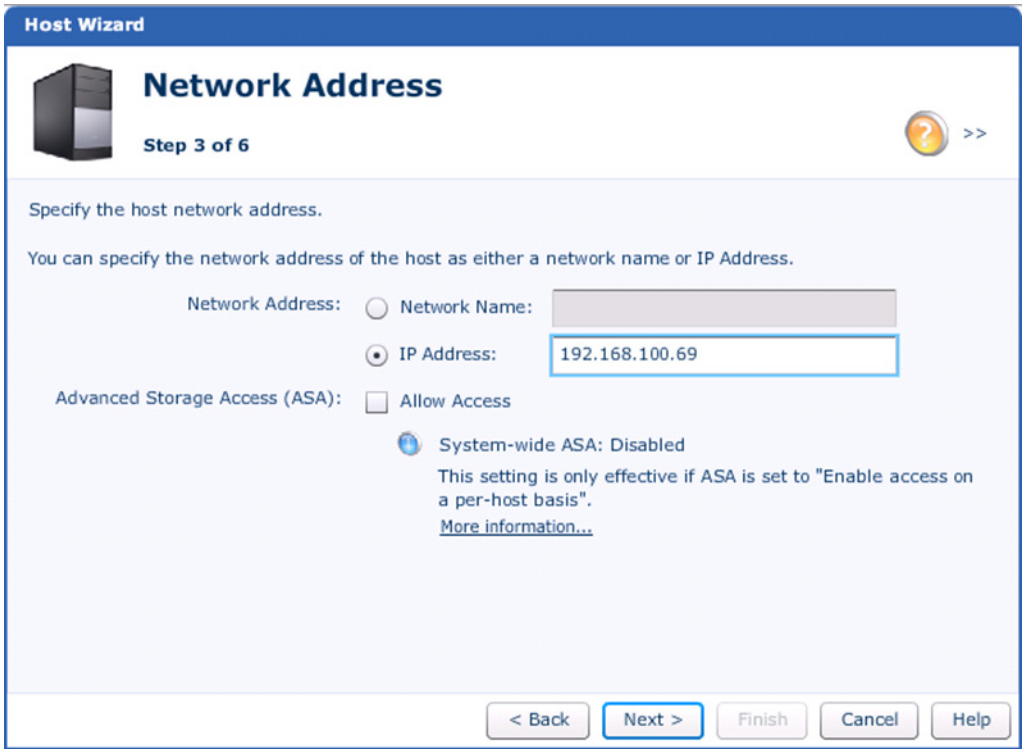

Note The iSCSI initiator IP address of each servers created on the UCS Manager can be easily obtained from the following page: From the **LAN** tab, expand **LAN > Pools > root > IP Pools**, and select the iSCSI IP pool on fabric A. Click the **IP addresses** tab on right pane of the window, and sort by Assigned To column. You can see the IP addresses configured per server basis in Figure 150.

Figure 150 IP Addresses Window Showing iSCSI Initiator IP Addresses of All the Servers

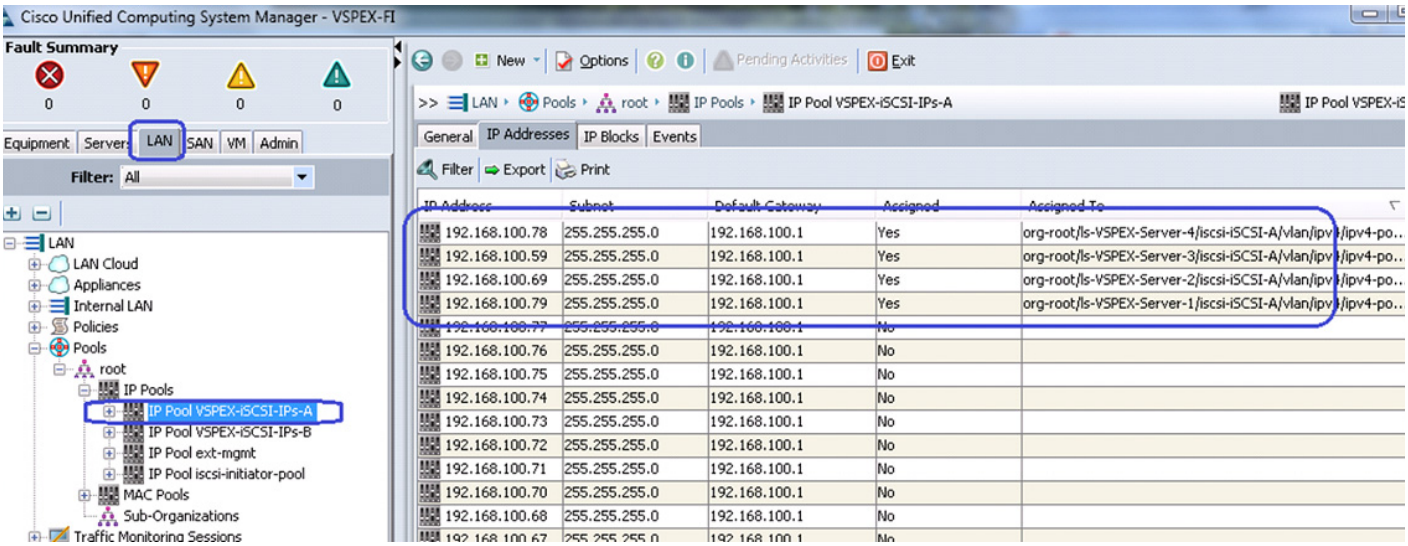

6. In the next step, enter the IQN name of the iSCSI vNIC of the service profile. Click **Next**.

Figure 151 Specifying iSCSI Address

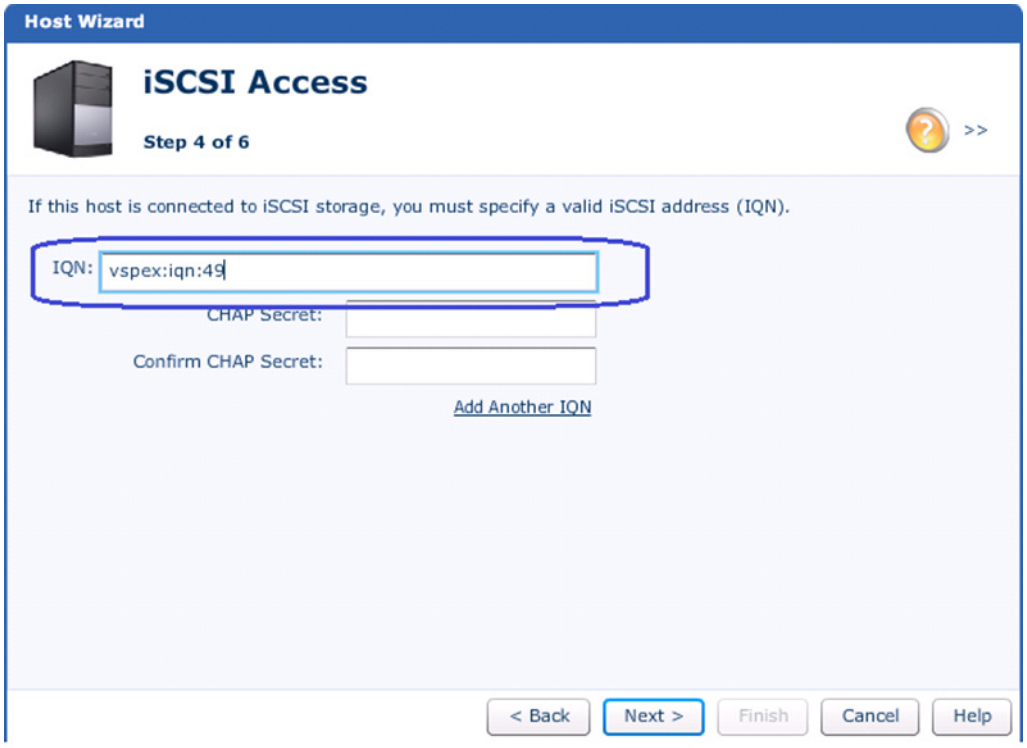

Note To find the IQN name assigned to various iSCSI vNICs, From the **SAN** tab in UCS Manager, expand **SAN > Pools > root > IQN Pools** and select the IQN pool created in previous section. Click the **IQN Suffixes** tab in the right pane of the window and sort the rows by clicking on Assigned To column title. This provides IQN name suffixes per vNIC basis as shown in Figure 152.

Figure 152 IQN Suffixes Window Showing IQN Name Assigned to iSCSI vNICs

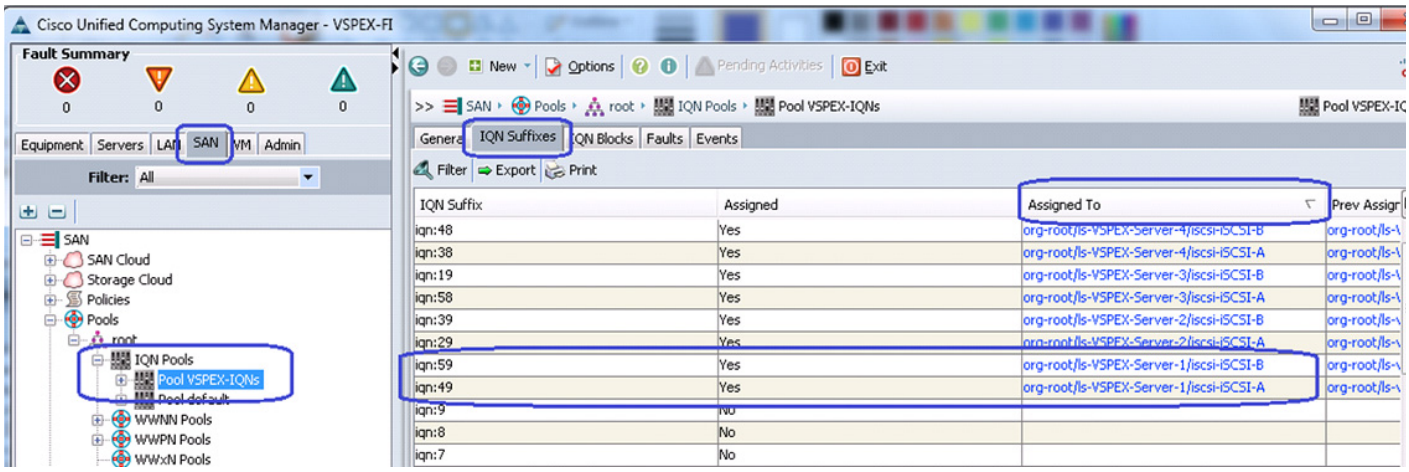

7. Click **Next** in the Host Wizard in the EMC Unisphere GUI, verify the configuration and deploy the configuration. This wizard allows us to add iSCSI initiator on one fabric only. To add the iSCSI initiator information on fabric B, select the ESX-Host1 in Hosts window, and click **Details.**

Figure 153 Adding iSCSI Initiator on Fabric B

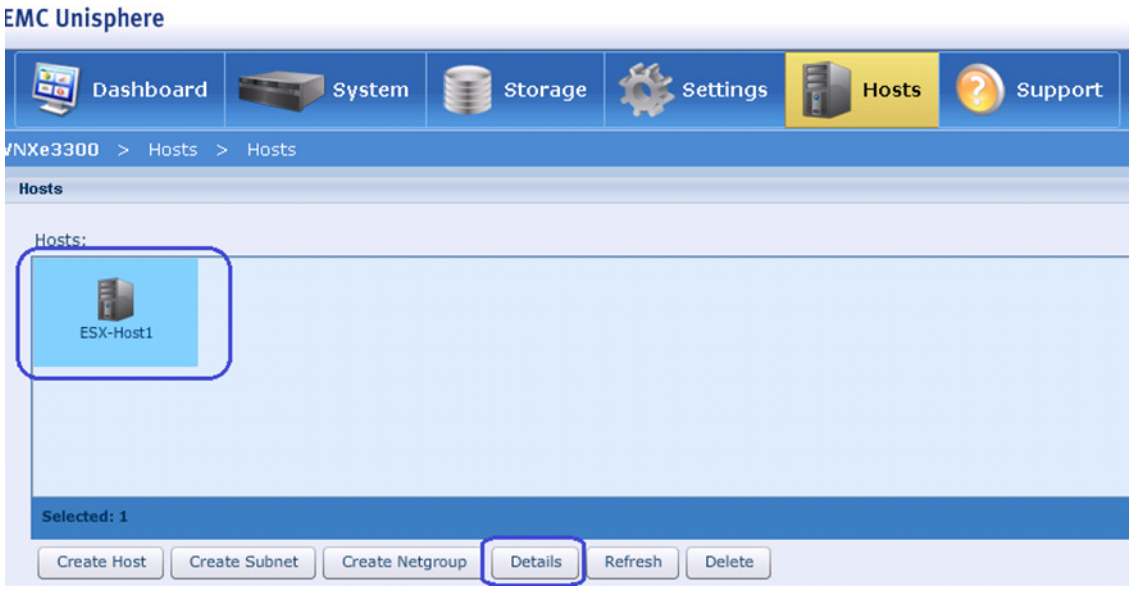

8. In the Details window, click **Add** to add IP address on fabric B.

Figure 154 Adding Data Storage Address in the Host Details Window

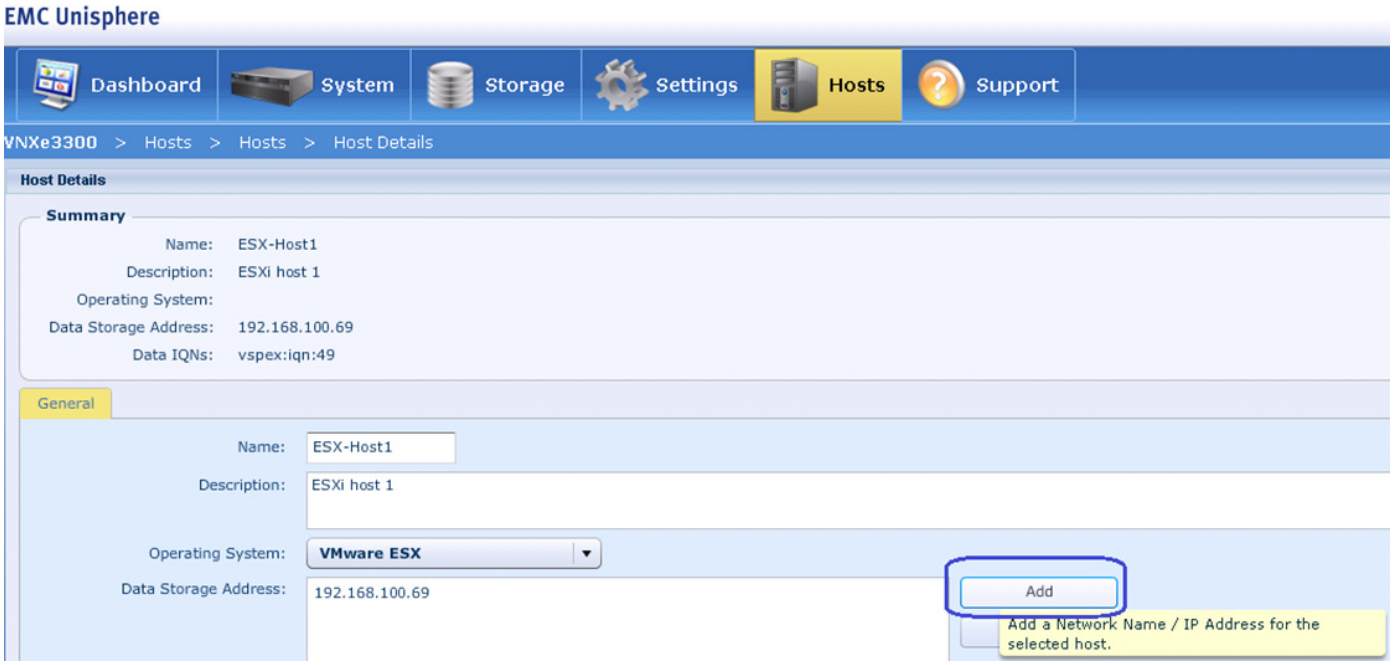

9. Similarly, in the details window click **Add** to add the IQN of iSCSI vNIC on fabric B. Click **Apply Changes** after the second IQN is added.

Figure 155 Adding IQN in the Host Details Window

10. Repeat steps 1 to 9 for all the four ESXi hosts service profiles created in UCS Manager.

Create Boot LUNs (Datastores) for Each Server

All the ESXi hosts access information is configured on VNXe, we can now carve our boot LUNs for each server using appropriate access privileges. Follow these steps to carve out boot LUNs for each server:

1. Click the **Storage** tab and click **VMware** to create datastore.

Figure 156 Selecting VMware for Carving Out Boot LUNs for Each Server

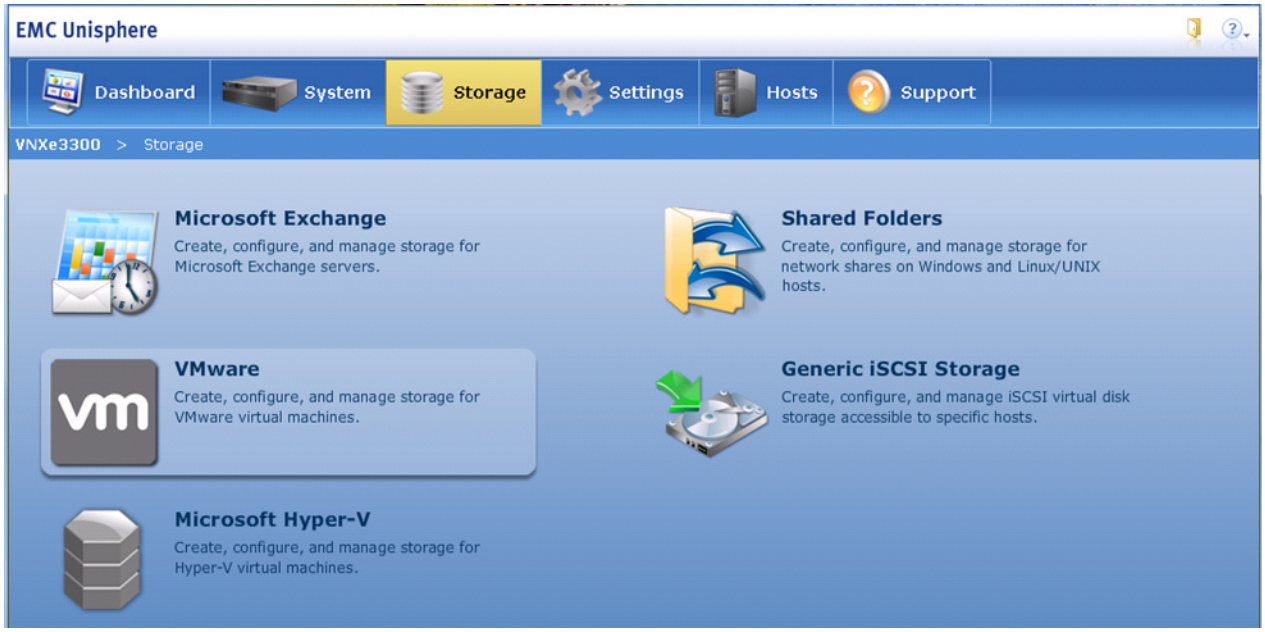

2. In the VMware Storage window, click **Create**.

Figure 157 Creating VMware Datastores

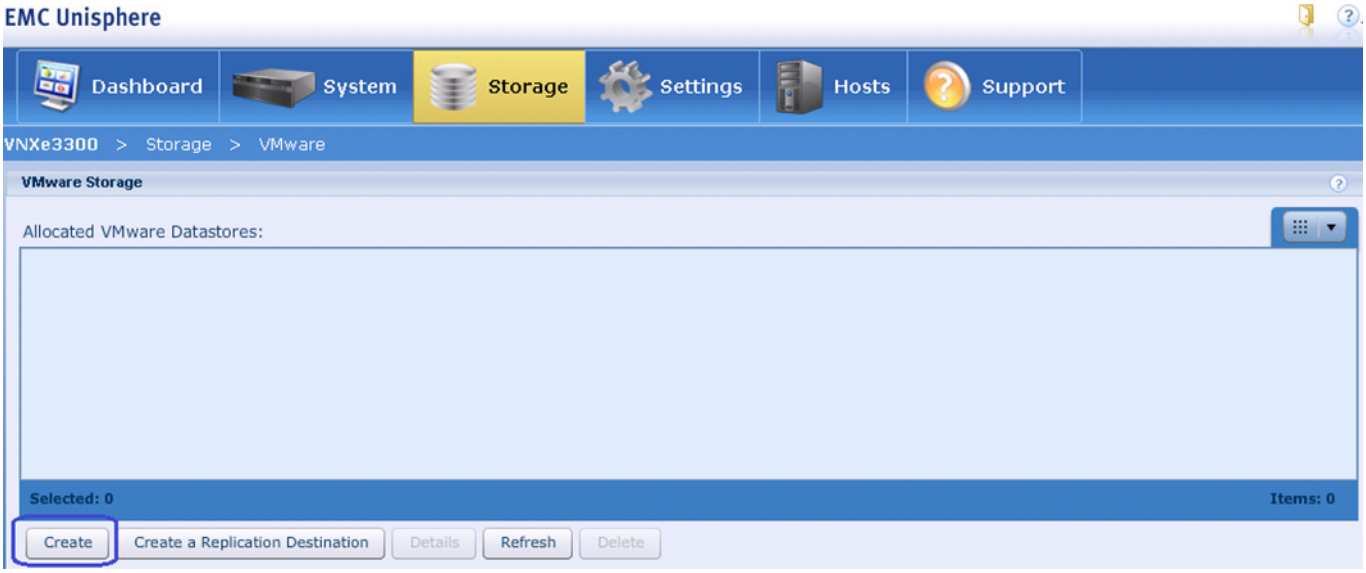

3. Enter the name of the datastore in the Name field and (optional) description. Click **Next**.

Figure 158 Specifying Datastore Name

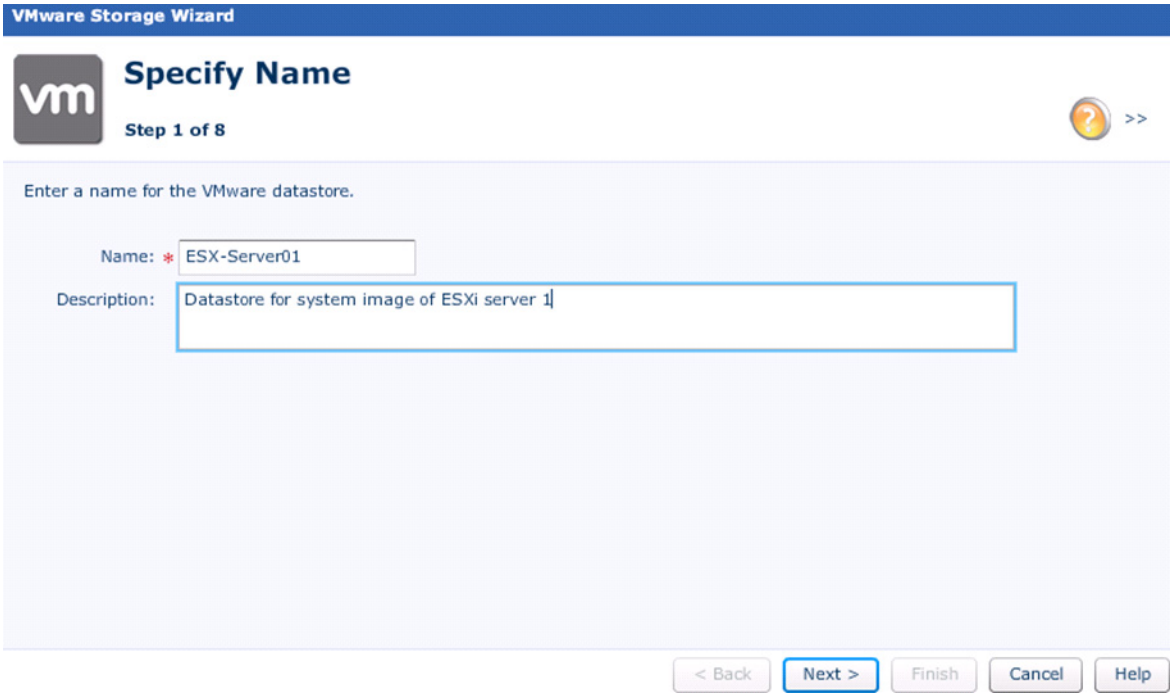

4. Click the **Virtual Disk (VMFS)** radio button in the VMware Storage wizard. Click **Next**.

Figure 159 Selecting Datastore Type

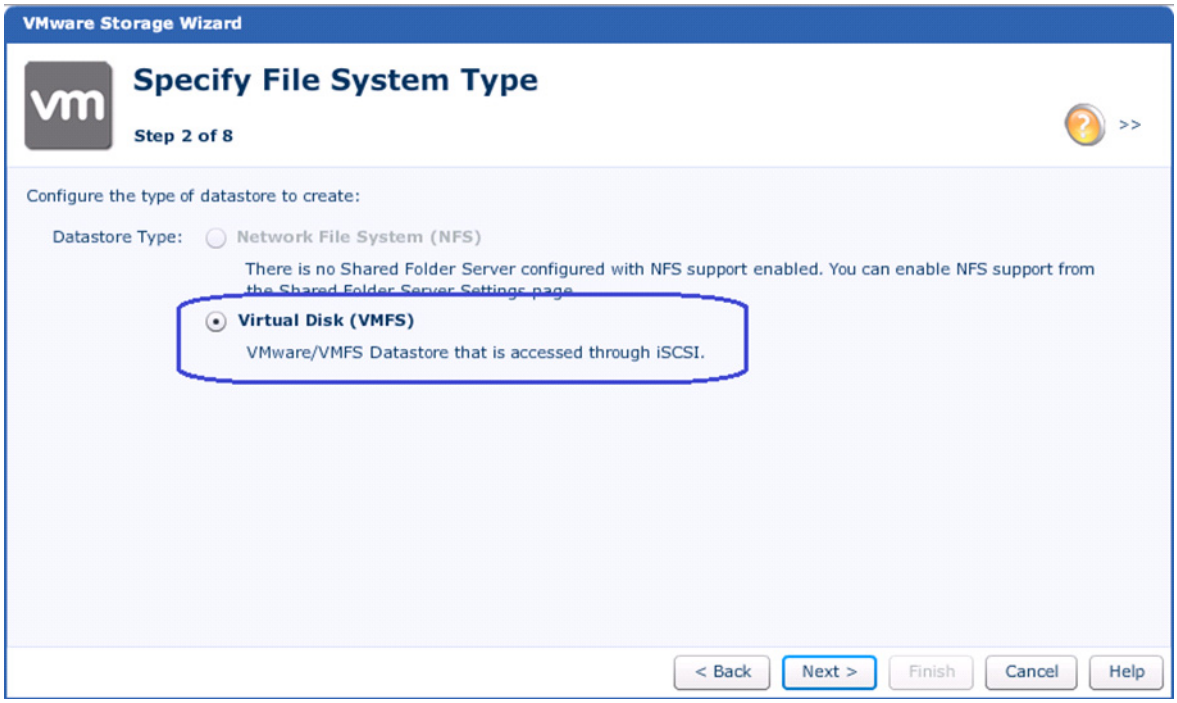

5. Choose 50 GB datastore from the Size drop-down list and select Systematic resource pool. Click **Next**.

Figure 160 Specifying the Size of the Datastore

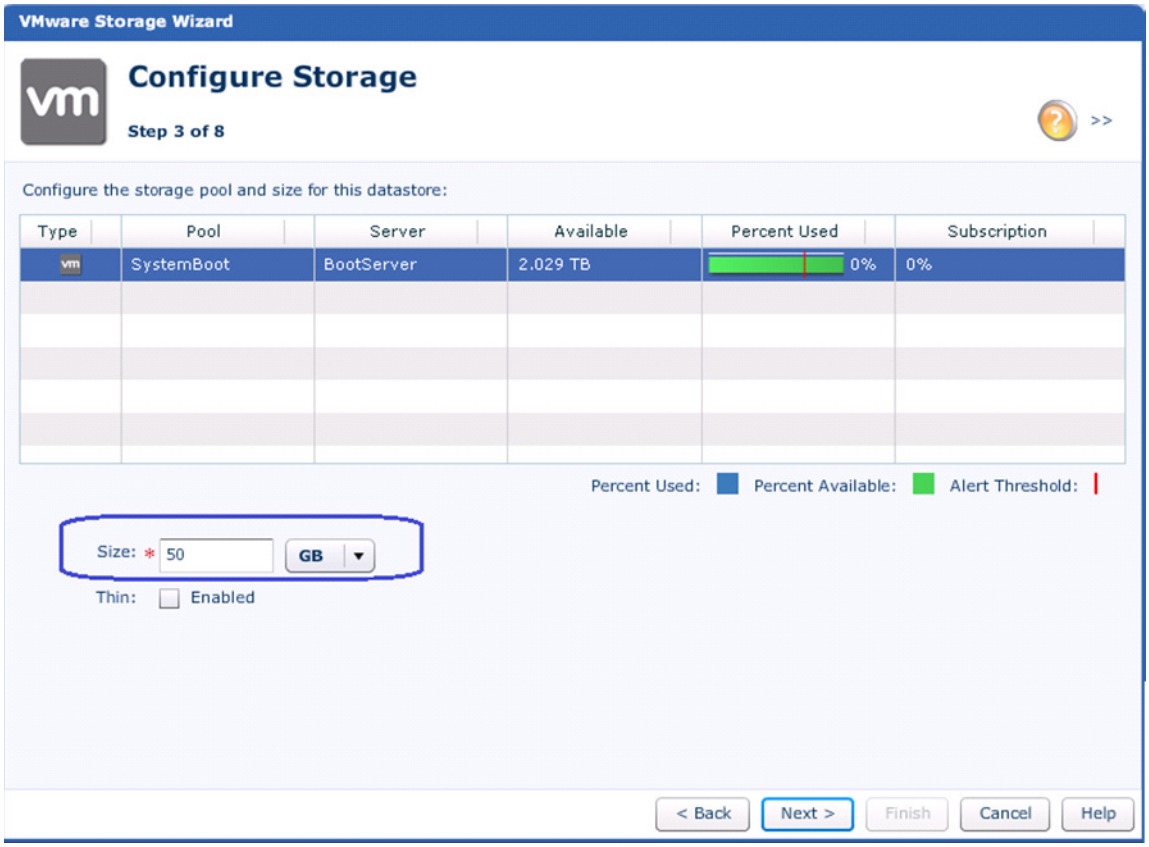

6. Only allow ESX-Host1 to access this datastore by selecting Datastore and snapshot for access permission from the drop-down list.

Figure 161 Providing Access permission to the Host

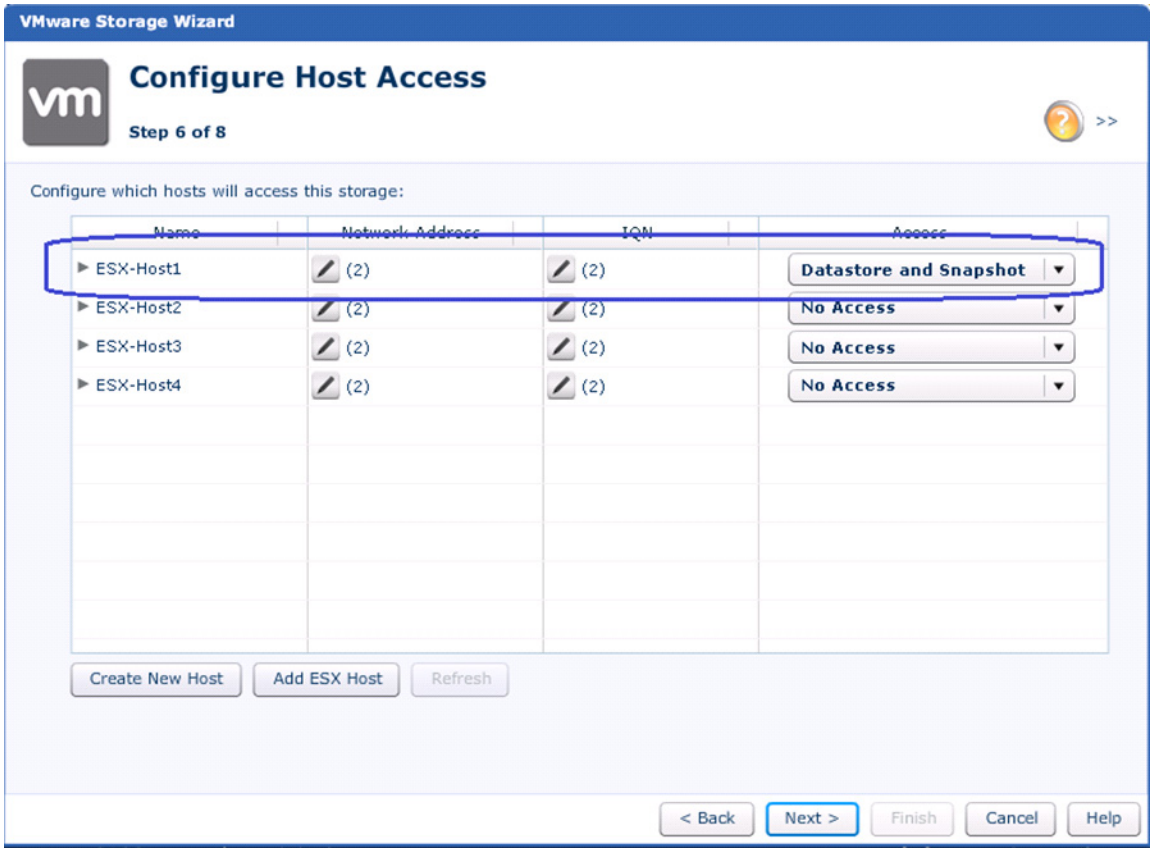

7. Verify the information and deploy the configuration. When success notification is received, click **Close** to create the datastore.

Figure 162 Verifying the Datastore Creation

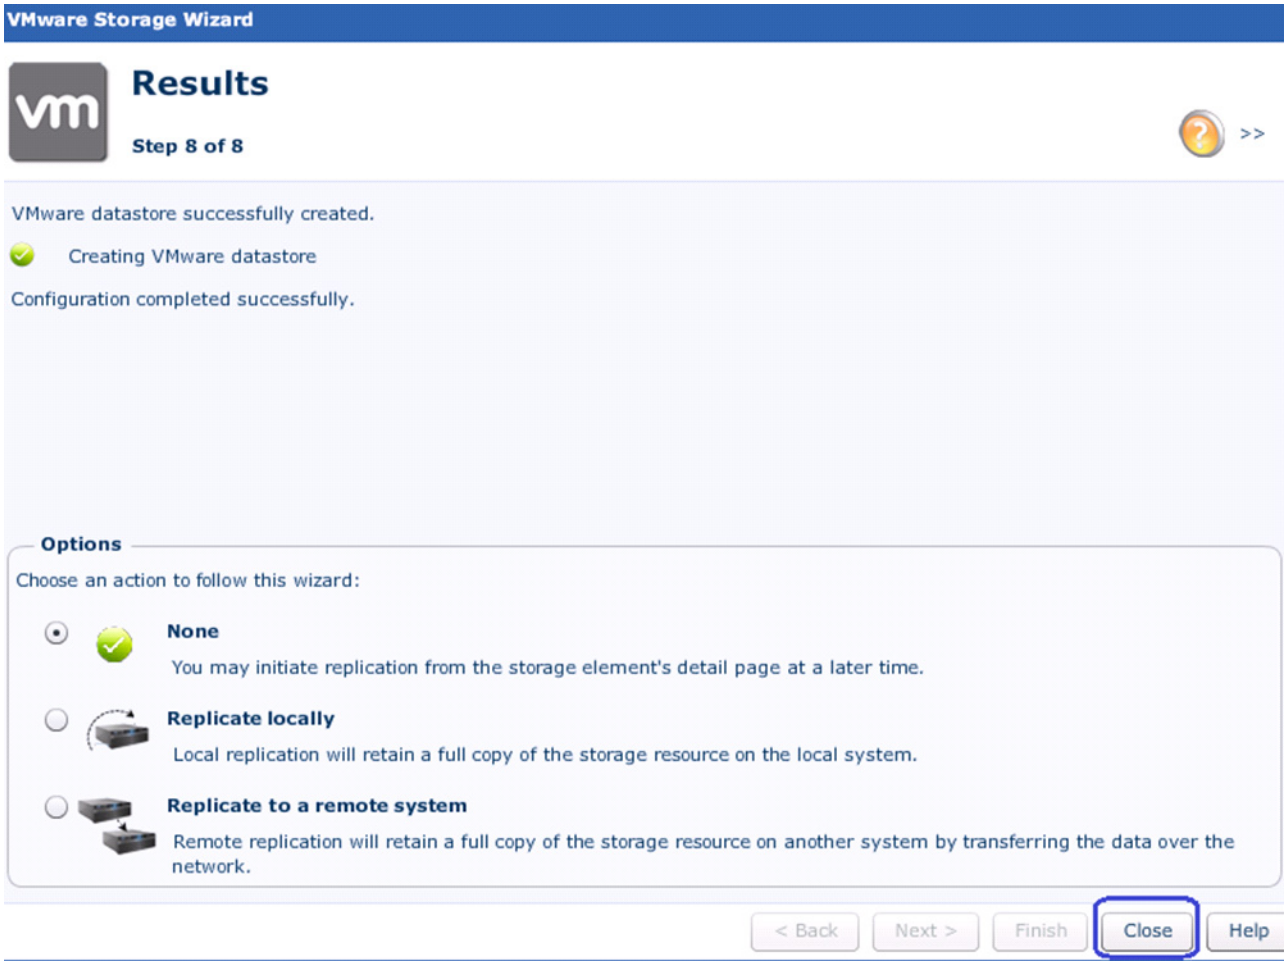

8. Repeat steps 1 to 7 for all the four ESXi hosts in the architecture. Next, select a specific ESXi host, and click **Details**.

Figure 163 Creating VMware Datastores for All the Hosts

9. Note down the LUN ID of the given server.

Figure 164 VMware Storage Details Window Showing Summary of the Selected Server

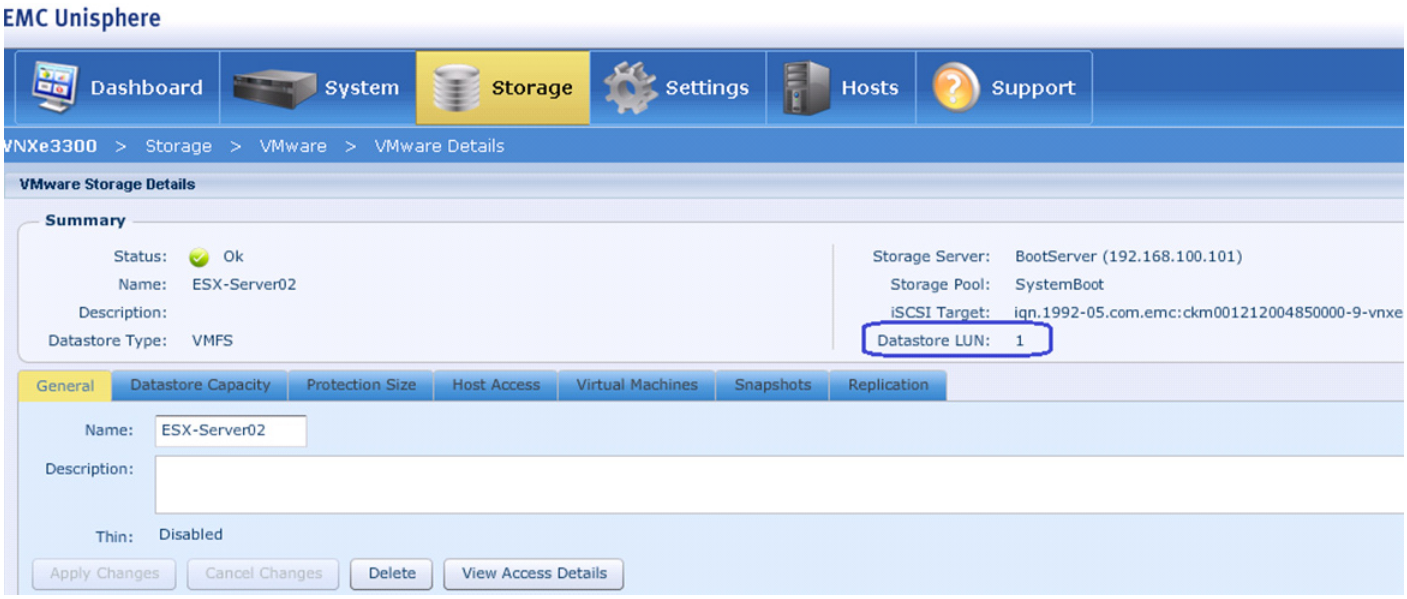

10. We need to update the iSCSI target LUN ID in the service profile boot policy. Only one of the four servers (typically the first datastore created) will have the LUN ID as 0. For all the other servers, the LUN ID needs to be updated. Click the **Servers** tab of UCS Manager, choose specific service profile, click the **Boot Order** tab in the right pane of the window and click **Set Boot Parameters** in the iSCSI vNICs area.

Figure 165 Setting Boot Parameters for the Second Server

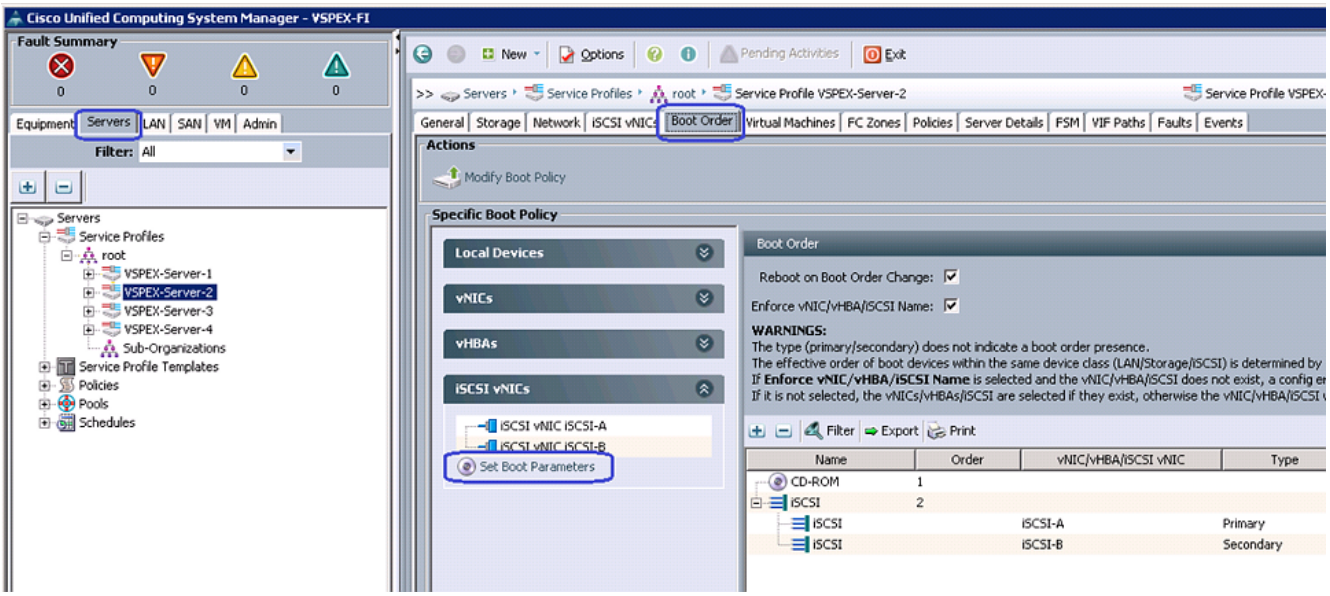

11. Select the iSCSI Static Target Interface, and click  in the right pane of the window.

Figure 166 Editing the iSCSI Static Target Interface

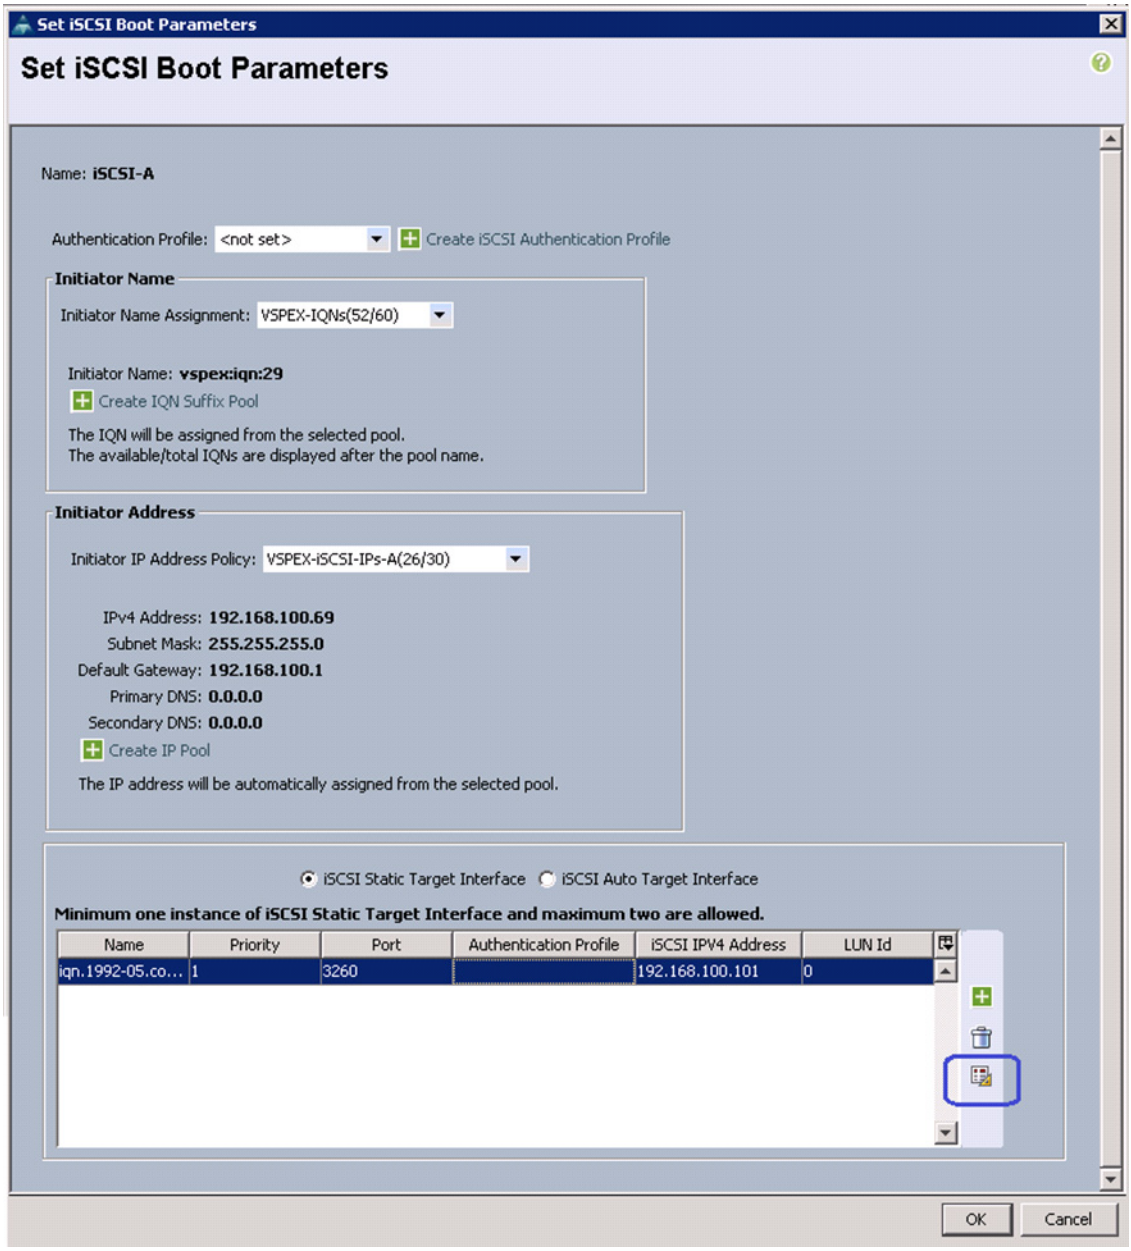

12. Update the LUN ID to match the LUN ID shown in step 9, and click **OK**.

Figure 167 Updating the LUN ID

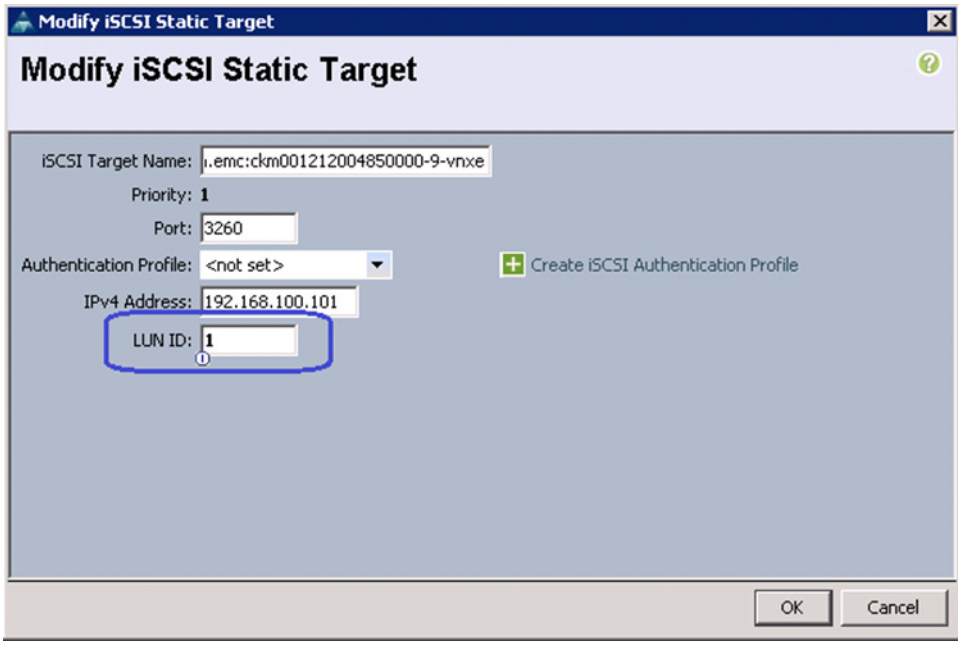

13. Click **OK** in the Boot Order window. A warning message window appears. Click **Yes** to continue.

Figure 168 Warning Message on Updating the LUN ID

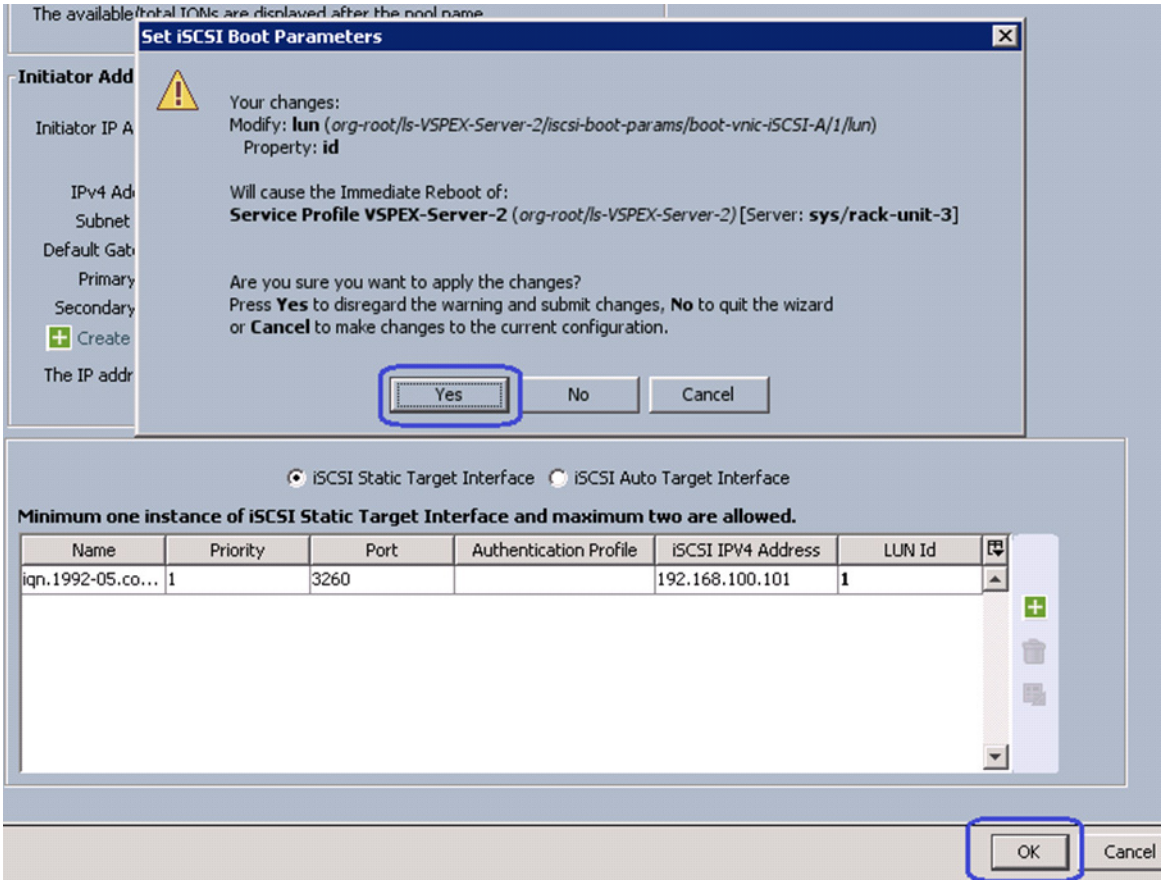

14. Repeat steps 10 to 13 for the iSCSI vNIC on fabric B using the same service profile.
15. Repeat steps 1 to 14 for all the service profiles in UCS Manager.

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

Install ESXi Servers and vCenter Infrastructure

Follow these steps to install ESXi image on Cisco UCS servers:

1. From UCS Manager GUI, click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a particular service profile. Click KVM Console in the right pane of the window.

Figure 169 Launching KVM Console

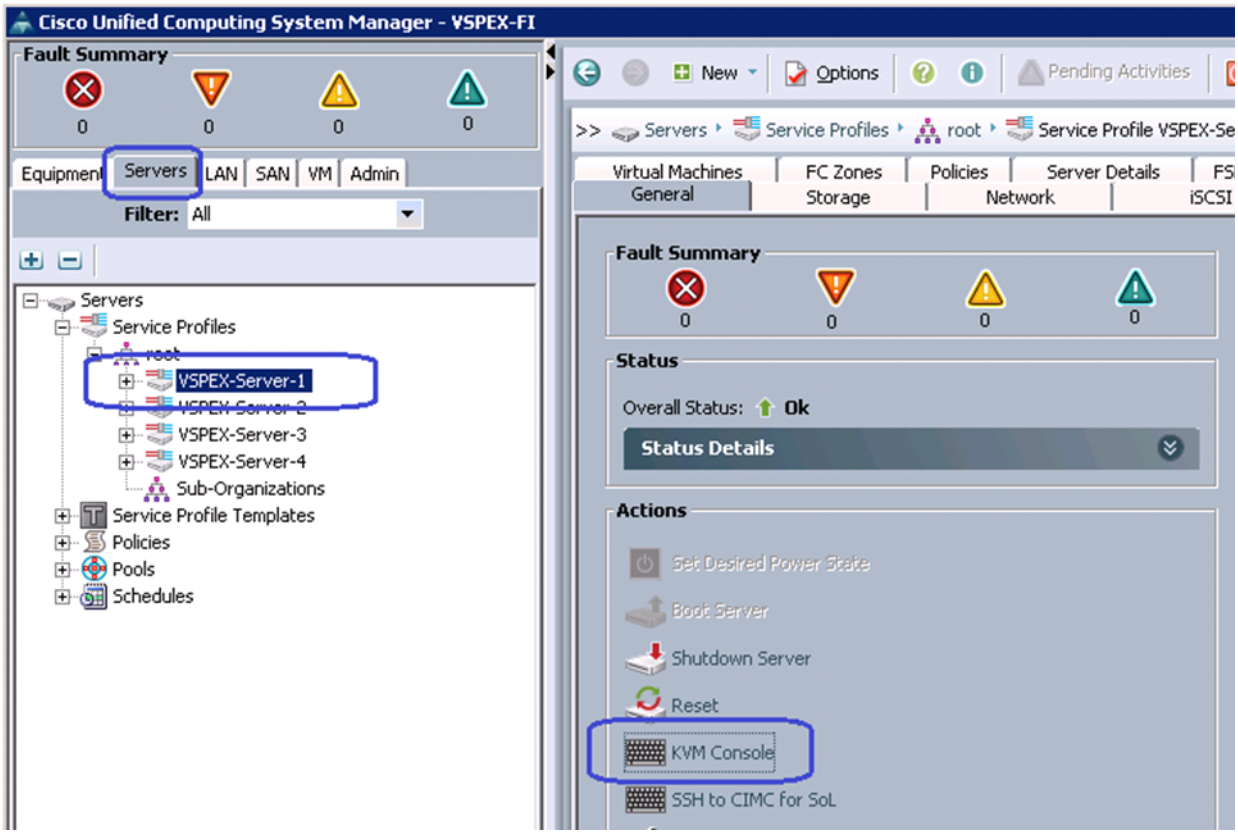

2. Once the Java pallet of KVM is launched, click the **Virtual Media** tab and click **Add Image**. A window appears to select an ISO image. Browse through the local directory structure and select ISO image of the ESXi 5.1 hypervisor installer media.

Figure 170 Adding an ISO Image of the ESXi 5.0 Hypervisor Installer Media

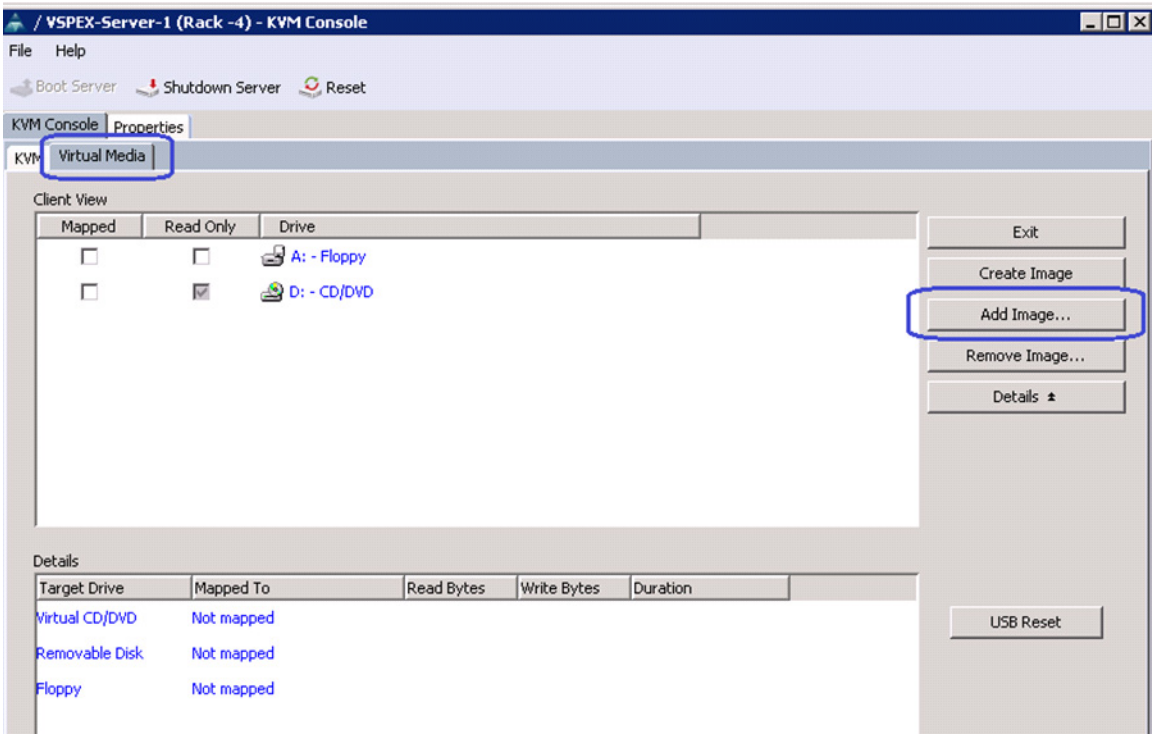

3. When the ISO image shows up in the list, check the Mapped check box and click **Reset** to reset the server.

Figure 171 Mapping the ISO Image and Resetting the Server

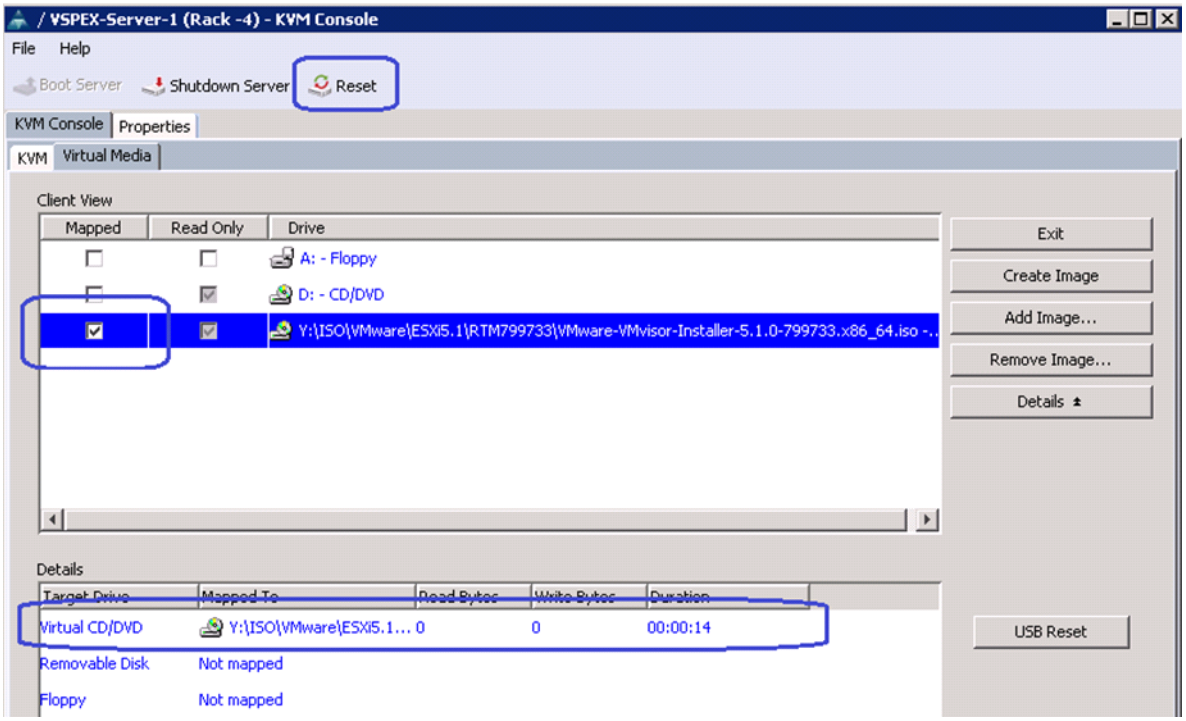

4. Click **OK** in the Reset Server warning message window.

Figure 172 Warning Message for Resetting the Server

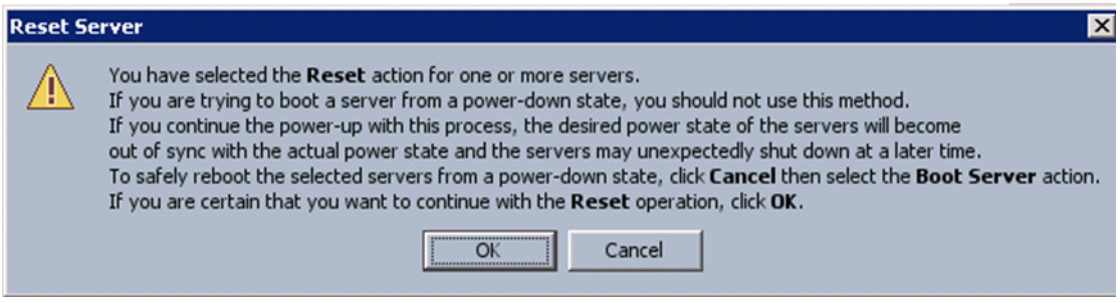

5. Click the **Power Cycle** radio button and click **OK**.

Figure 173 Selecting Resetting Option

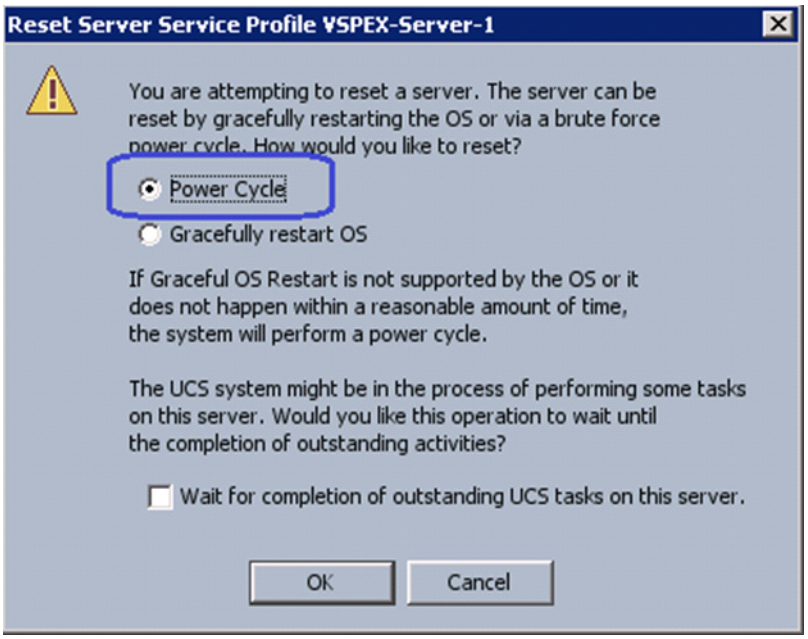

6. Click the **KVM** tab, and for iSCSI-variant, make sure iSCSI Option ROM installs successfully on both the fabrics. (For FC-variant, you will see similar Option ROM installation through vHBAs).

Figure 174 Verifying Option ROM Installed Successfully for iSCSI-variant

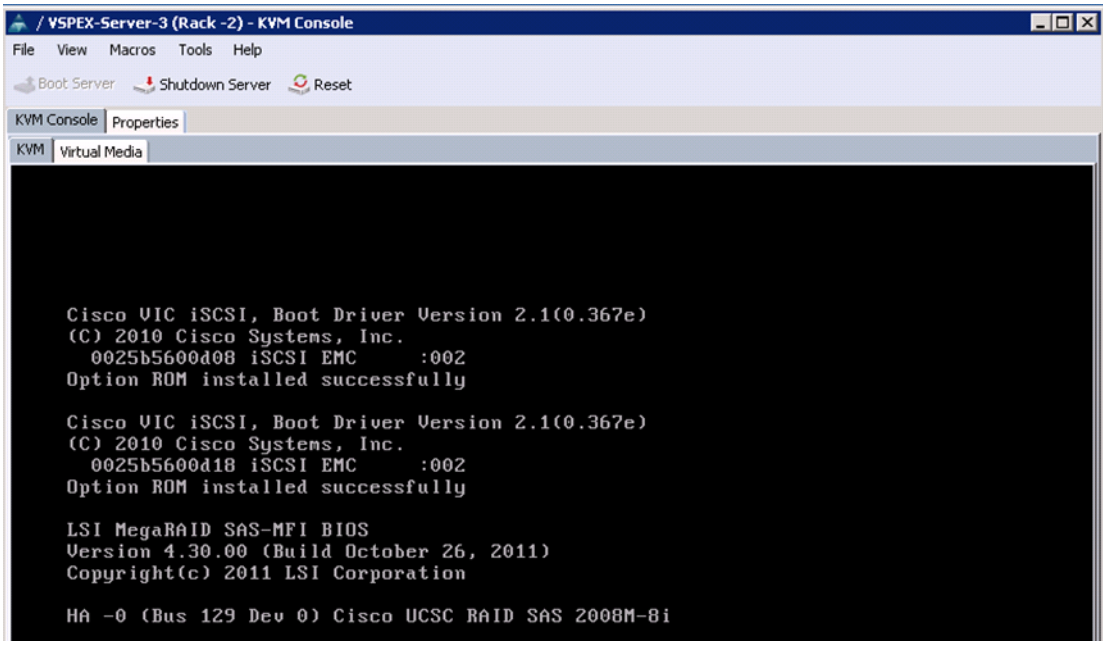

At this point of time, ESXi installation media would boot from the virtual disk mounted on the KVM. Follow the steps to install ESXi 5.1 hypervisor on the boot LUN. Make sure that you select the boot LUN and not the local disk. You can select all the default parameters or parameters settings as per your requirements.

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

Figure 175 Selecting Adapters for Default Management Network Connection

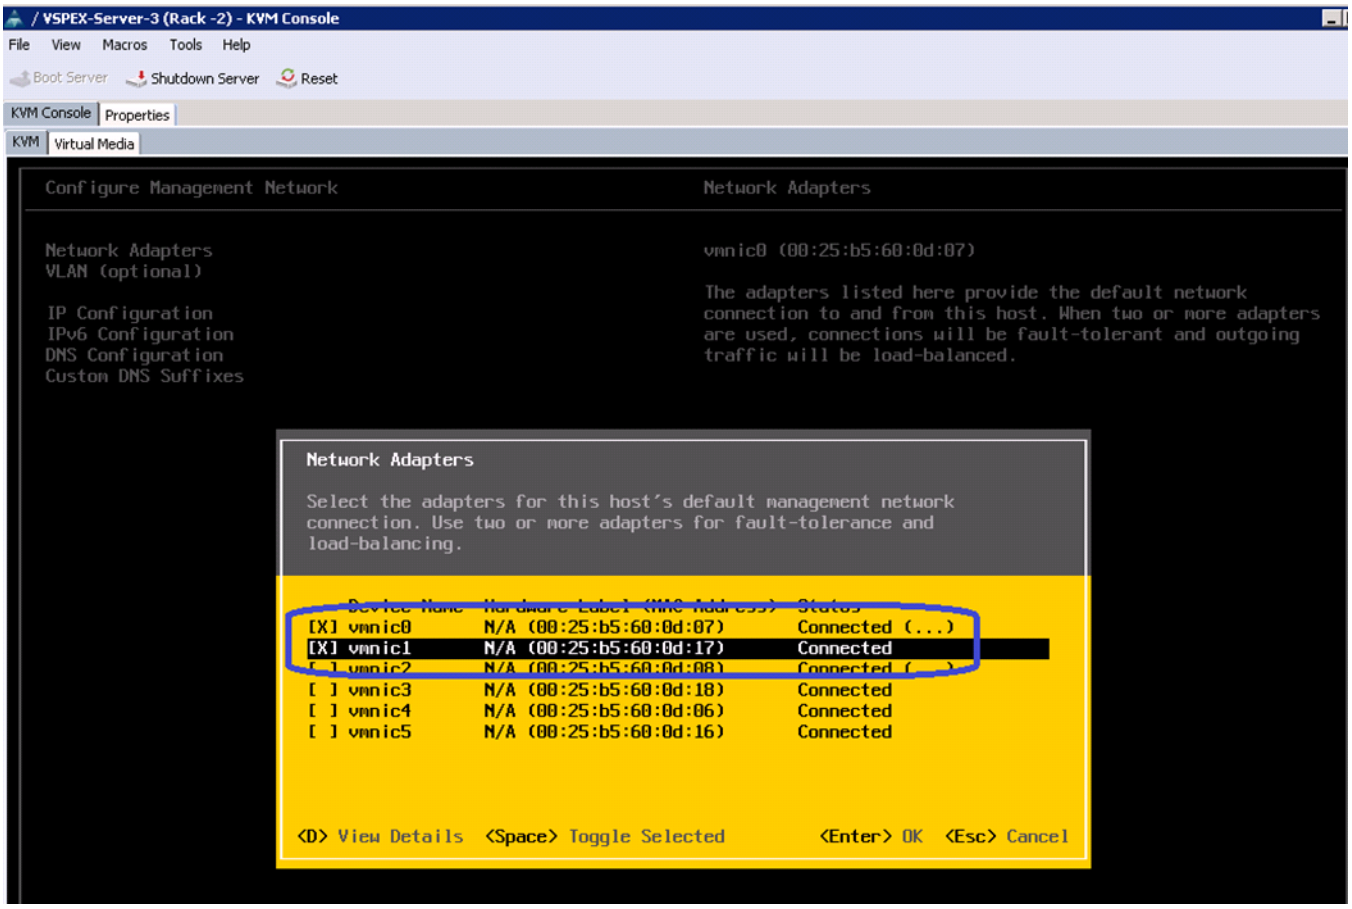

Note Easiest way to figure out which vmnic adapter should be used for the vSphere management purpose, you can identify the vmnic by MAC address. The MAC addresses of the vNICs (vmnic's) are summarized in UCS Manager GUI as show in Figure 176. Click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a particular service profile and click **VNICs**. The vNIC names and MAC addresses are listed in the right pane of the window.

Figure 176 vNIC Names and Their MAC Addresses are Shown in the vNICs Area

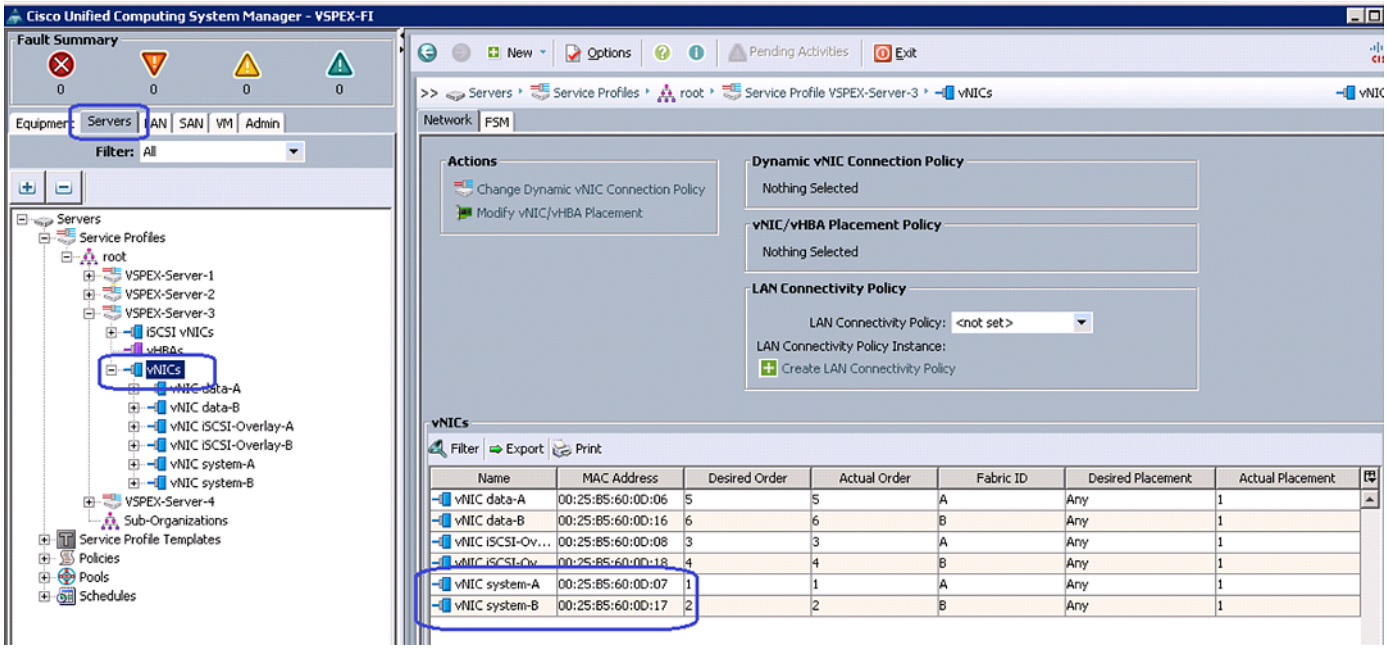

7. Repeat the ESXi installation steps for all the four servers.

VMware vCenter Server Deployment

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

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

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

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2032885

The following steps provide high level configuration procedure 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, then we need to directly connect 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 2008 R2 SP1. To ensure that adequate space is available on the vCenter and vSphere, Update Manager installation drive, see *vSphere Installation and Setup Guide*.

3. Install vCenter server

   Install vCenter by using the VMware VIMSetup 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 VMware vCenter

Follow these steps to configure cluster, HA, and DRS on 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. Check the check boxes Turn On vShpere HA and Turn On vSphere DRS, as shown in Figure 177.
   c. Click **OK**, to save changes.

Figure 177 Configuring HA and DRS on Cluster

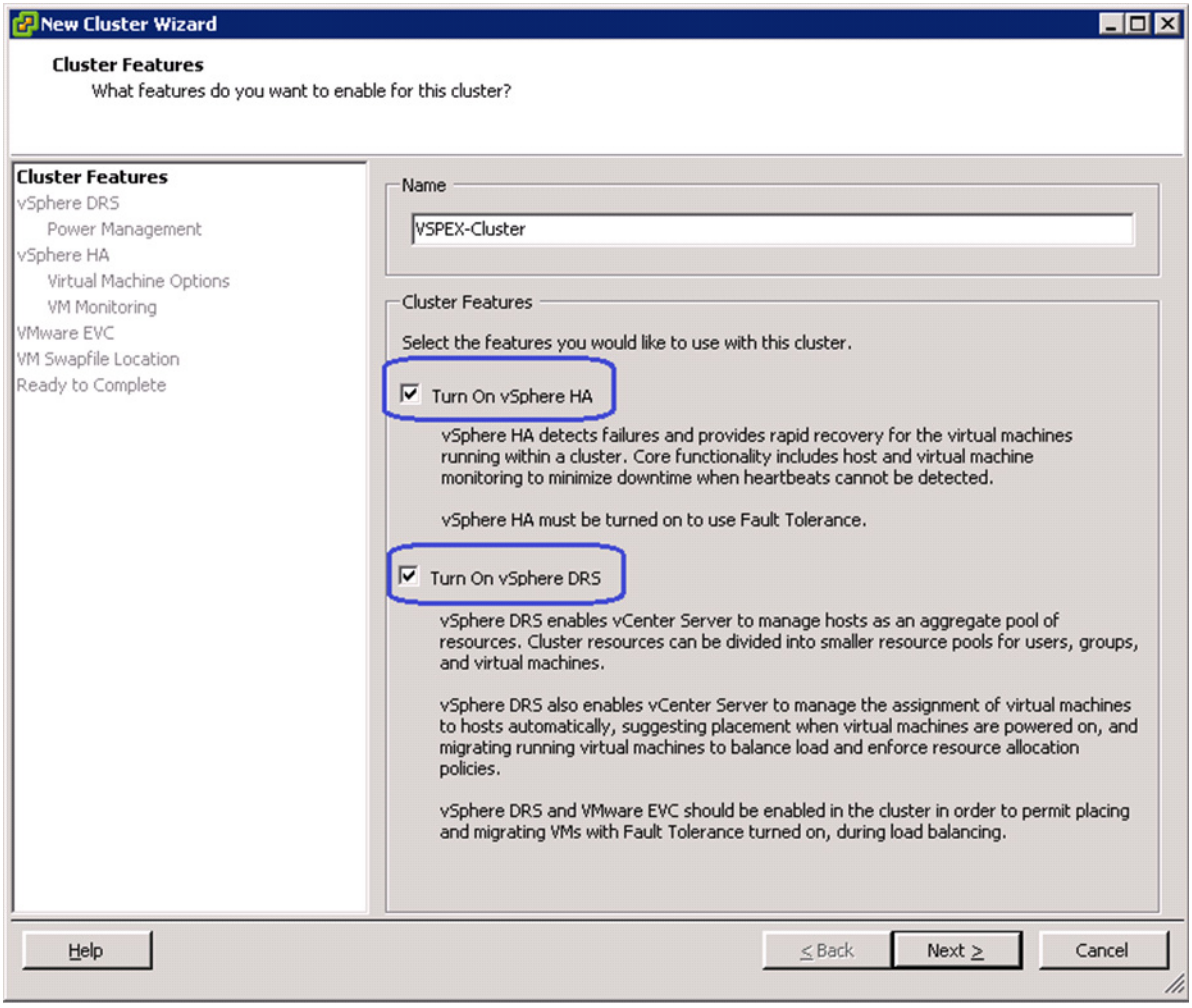

4. Add all the ESXi hosts to the cluster by providing servers' management IP addresses and login credentials one by one.

Figure 178 Adding ESXi Hosts to the Cluster

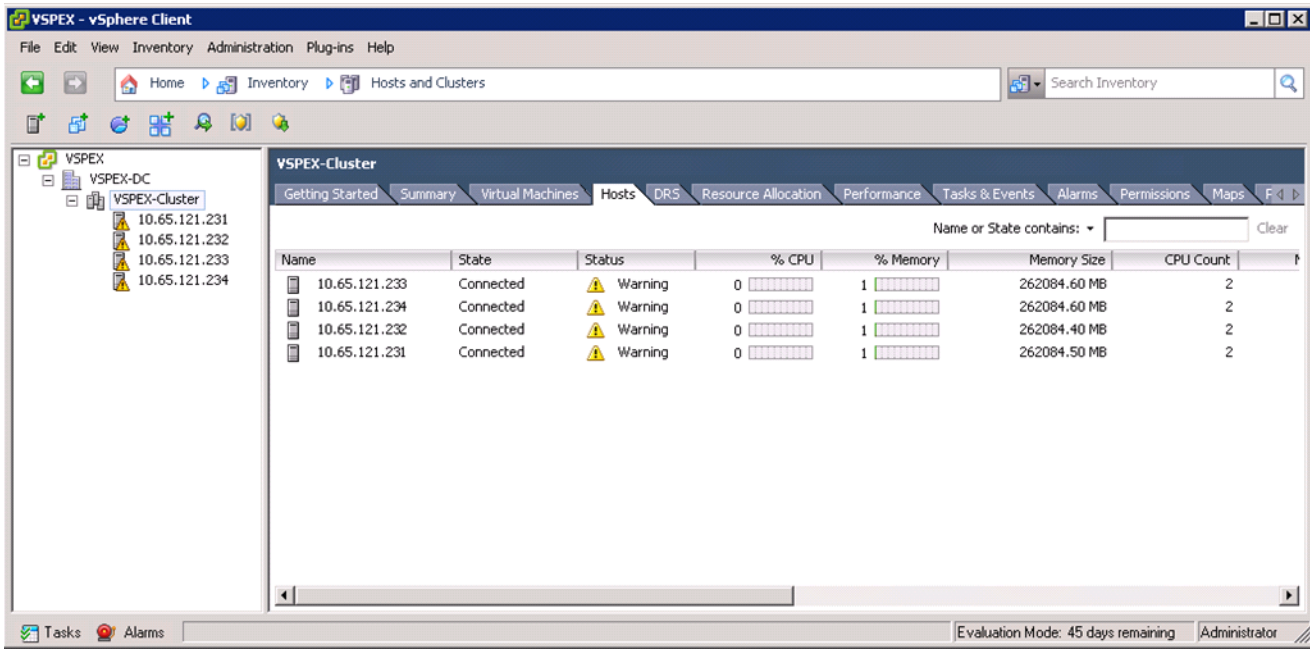

Virtual Networking Configuration

In UCS Manager service profile, we created six vNICs per server for iSCSI-variant and four vNICs per server for FC-variant. This shows up as six or four network adapters or vmnics in ESXi server. You can see these adapters in the vCenter by choosing **Home > Inventory > Hosts and Clusters**, select a particular server, click the **Configuration** tab in the right pane of the window, and click **Network Adapters**.

Figure 179 Network Adapter Showing vmnics in ESXi Server

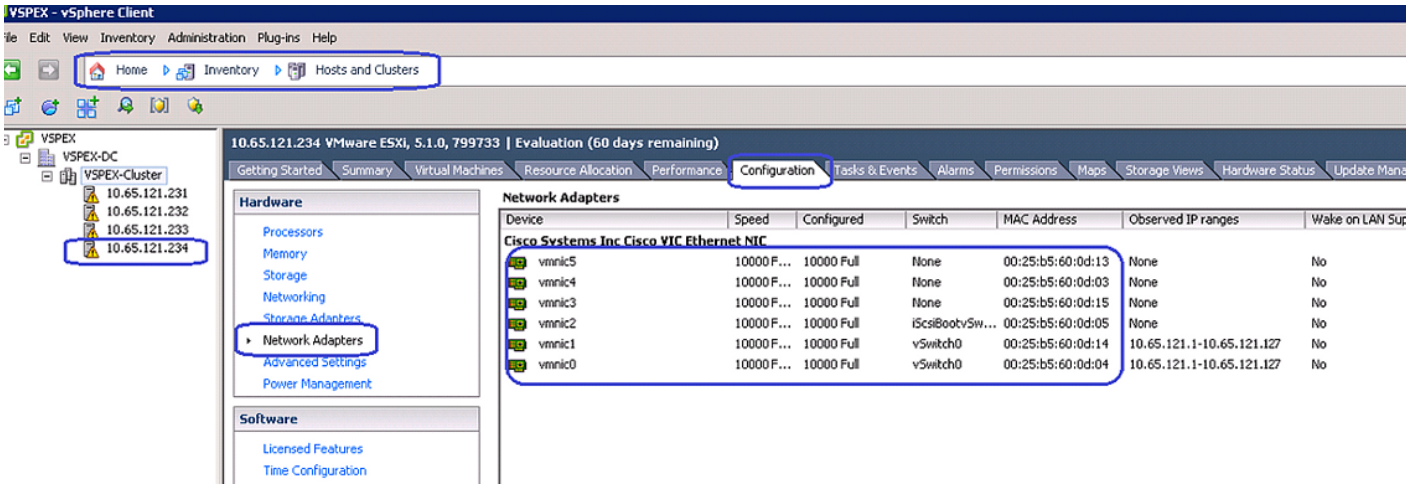

Table 9 shows UCS Manager service profile and vSphere vmnic per ESXi host basis:

Table 9 Service profile vNIC and vSphere vmnic relations

UCS Manager vNIC Names	vSphere VM NIC Names	MAC Address
System-A	vmnic0	
System-B	vmnic1	
*iSCSI-Overlay-VNIC-A	vmnic2	
*iSCSI-Overlay-VNIC-B	vmnic3	
Data-A	vmnic4	
Data-B	vmnic5	

*iSCSI Overlay vNICs are applicable for iSCSI-variant of the solution only.

We need to create three native vSwitches for virtual network configuration as follows:

1. vSwitch0—Standard, default vSwitch for management and vMotion traffic.
2. iScsiBootvSwitch—Default iSCSI boot vSwitch (iSCSI-variant only).
3. vSwitch1—For VM data traffic.

Each vSwitch listed will have two vmnics, one on each fabric for load balancing and high-availability. Also, for vMotion and iSCSI storage traffic, jumbo MTU needs to be configured in virtual network.

Follow these steps to configure the three vSwitches:

1. (iSCSI-variant only) In the vSphere client, choose **Home > Inventory > Hosts and Clusters**. In the Hosts and Clusters window, click the **Configuration** tab on the right pane of the window. Click **Networking** in the Hardware area. Click **Properties** to see the details of iScsiBootvSwitch.

Figure 180 Viewing Details of iScsiBootvSwitch

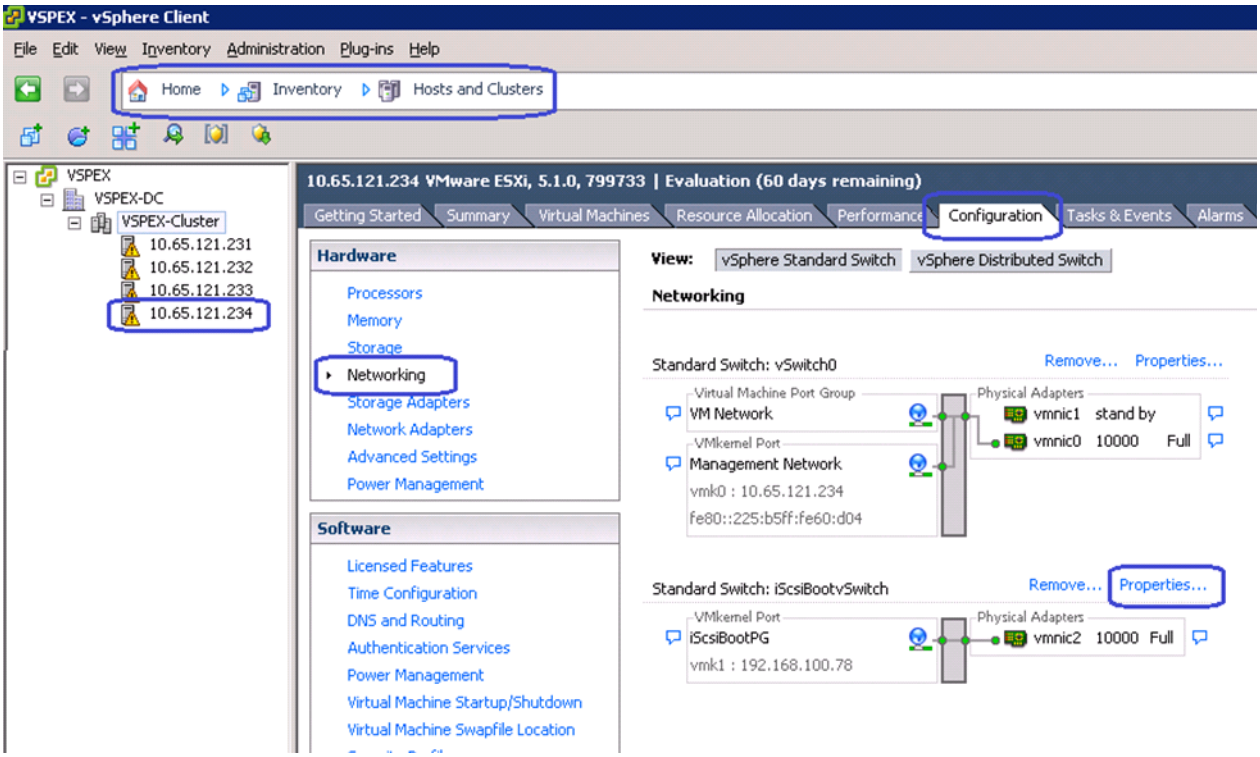

2. (iSCSI-variant only) Click the **Network Adapters** tab in the iScsiBootvSwitch Properties window, and click **Add**.

Figure 181 Adding vmnics to the iScsiBootStwitch

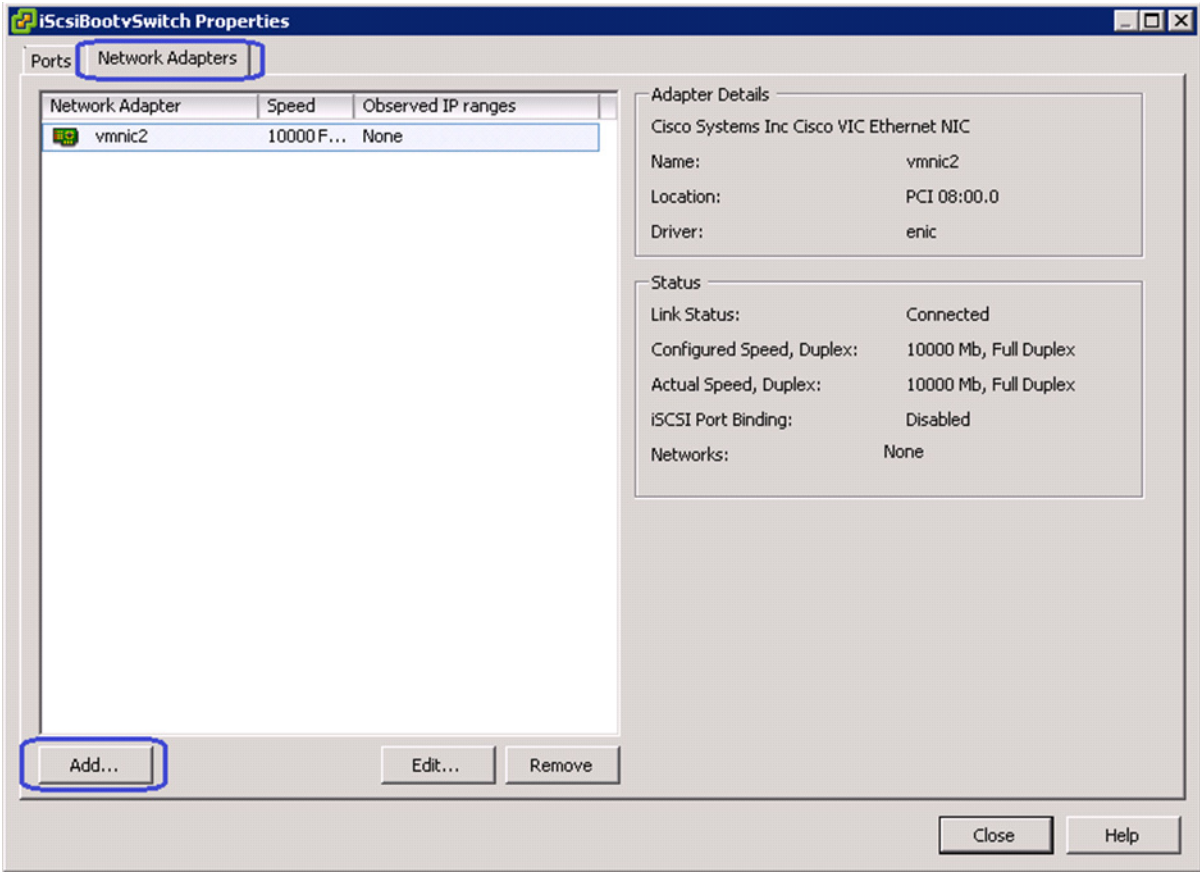

3. (iSCSI-variant only) Select the second vmnic to the iScsiBootvSwitch, use Table 9 to choose the correct vmnic. Click **Close**.

Figure 182 Selecting the Appropriate vmnic

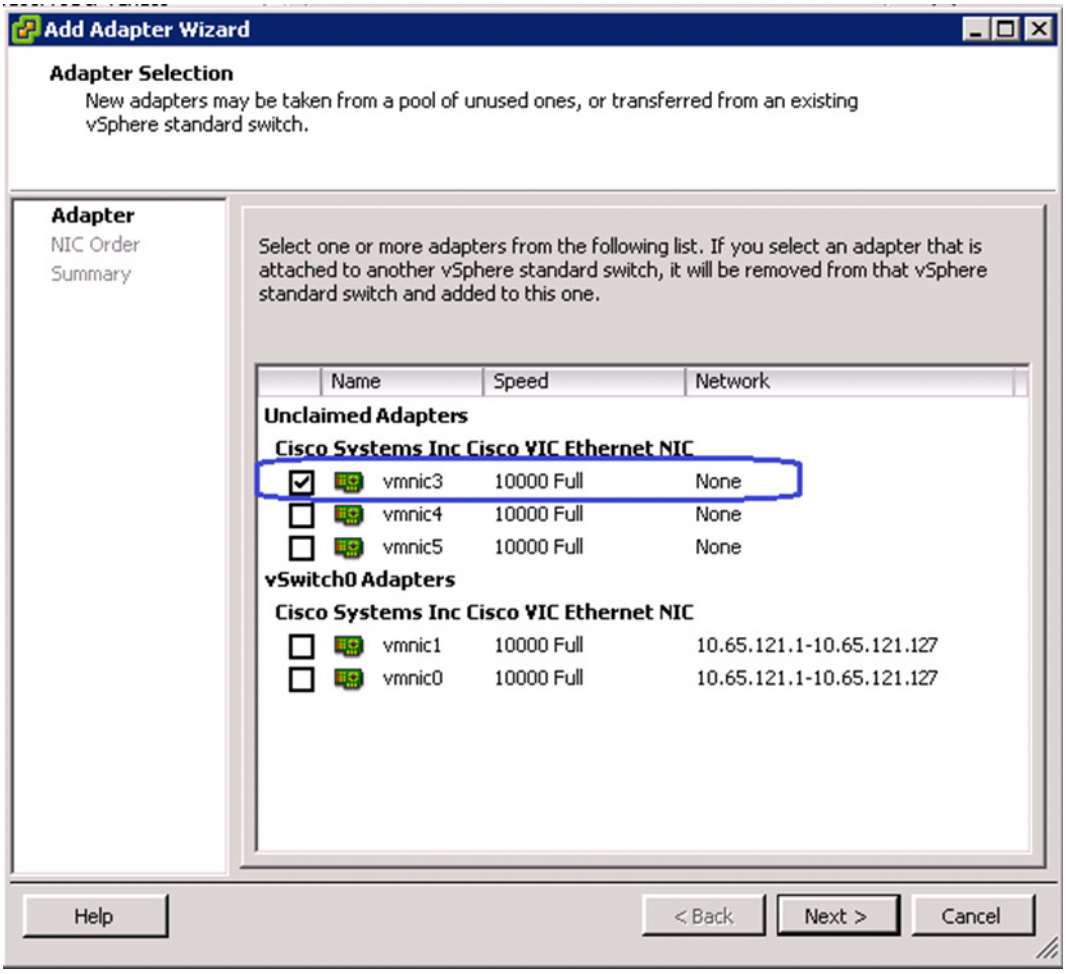

4. Now, click **Properties** to see the details of vSwitch0.

Figure 183 Viewing Details of vSwitch0

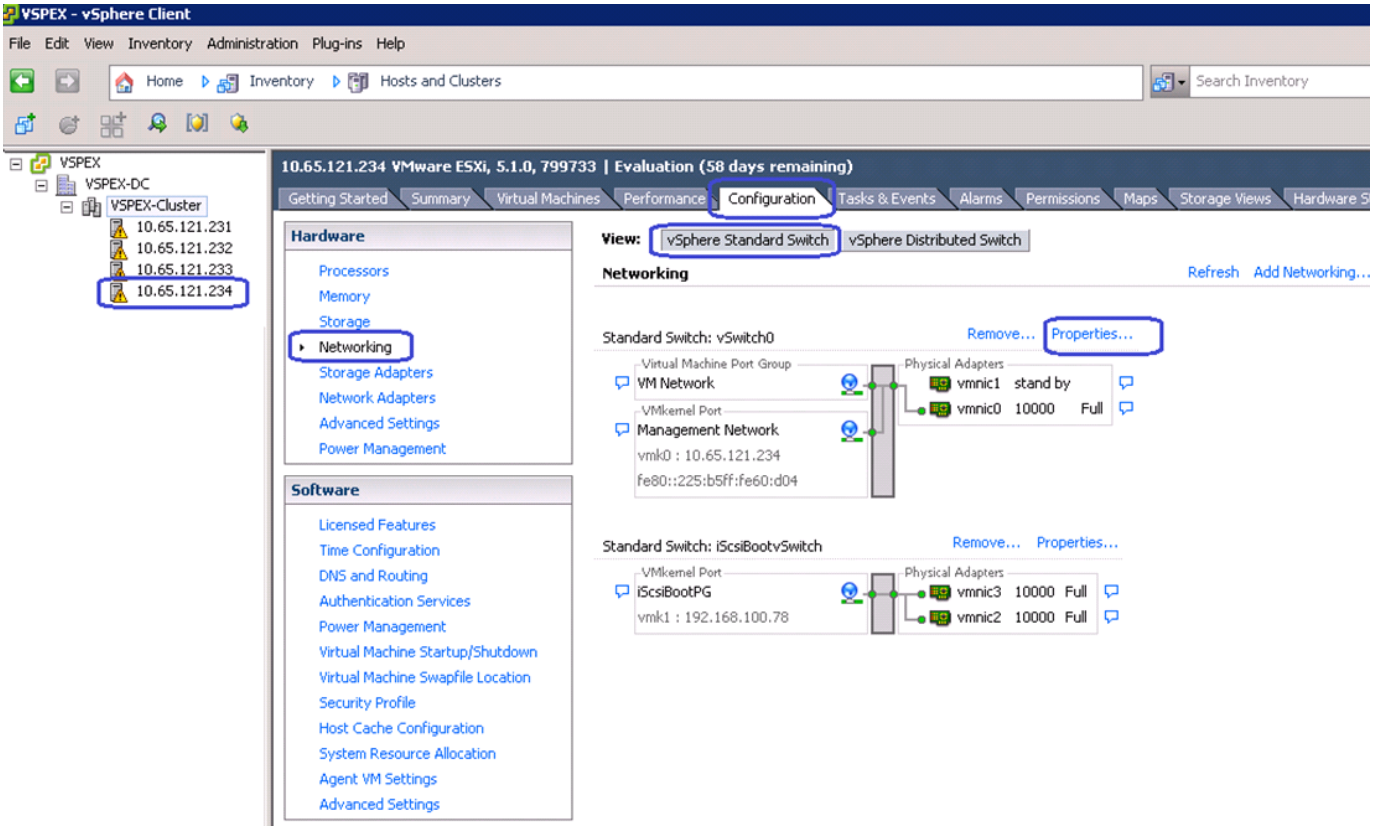

5. Select the vSwitch in the vSwitch0 Properties window and click **Edit**.

Figure 184 Editing vSwitch0 Properties

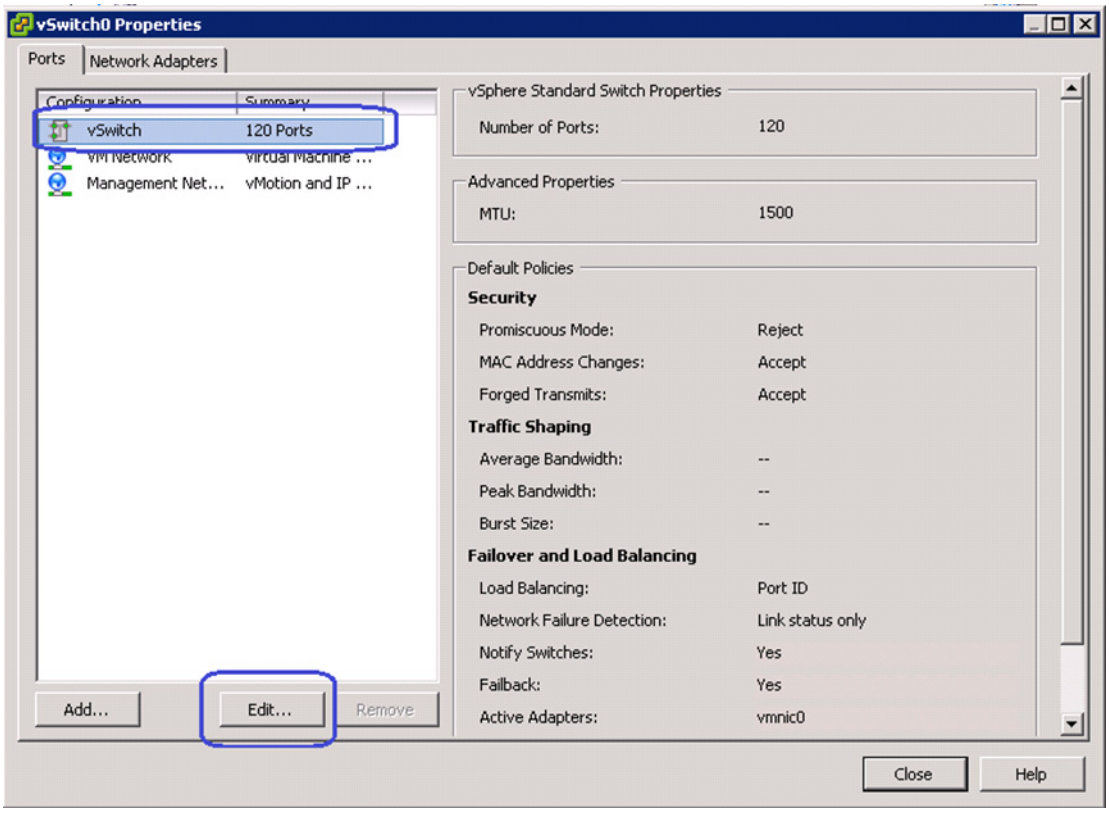

6. Click the **General** tab of the vSwitch0 Properties window, change the MTU in the Advanced Properties area to 9000.

Figure 185 Setting Jumbo MTU

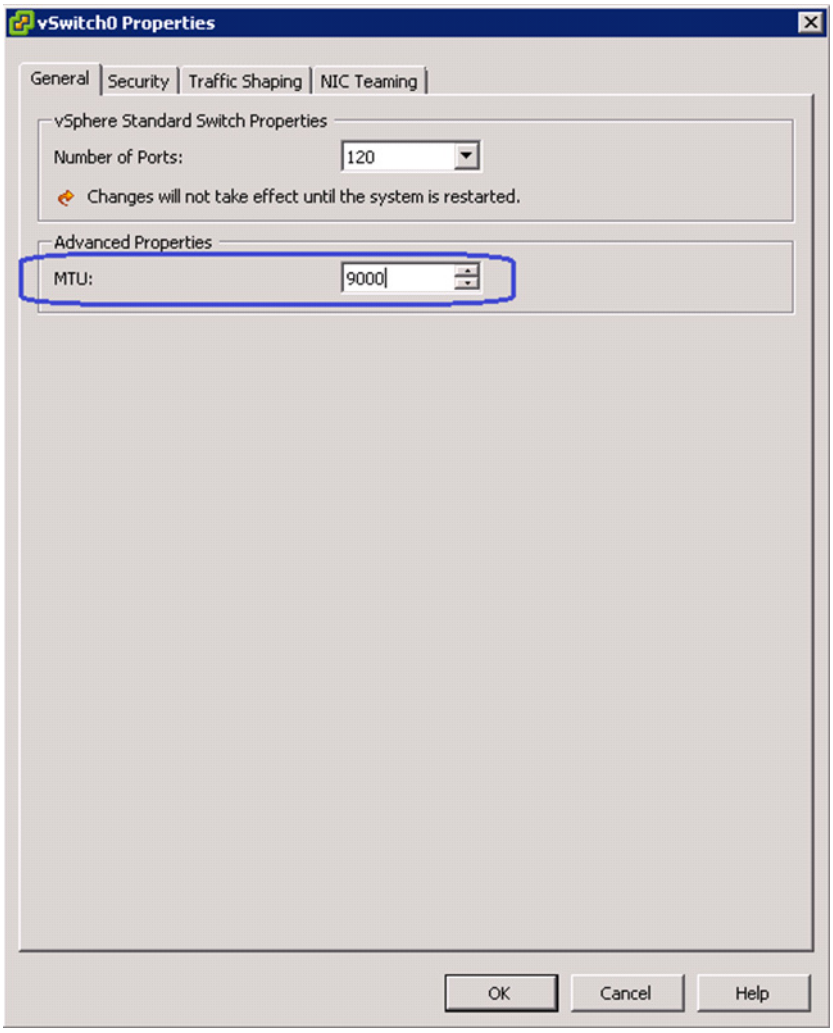

7. Click the **NIC Teaming** tab in the vSwitch0 Properties window. Select the adapter under Standby Adapters and click **Move up** to get it under Active Adapters, and click **OK**.

Figure 186 Moving Standby Adapter to Active Adapter

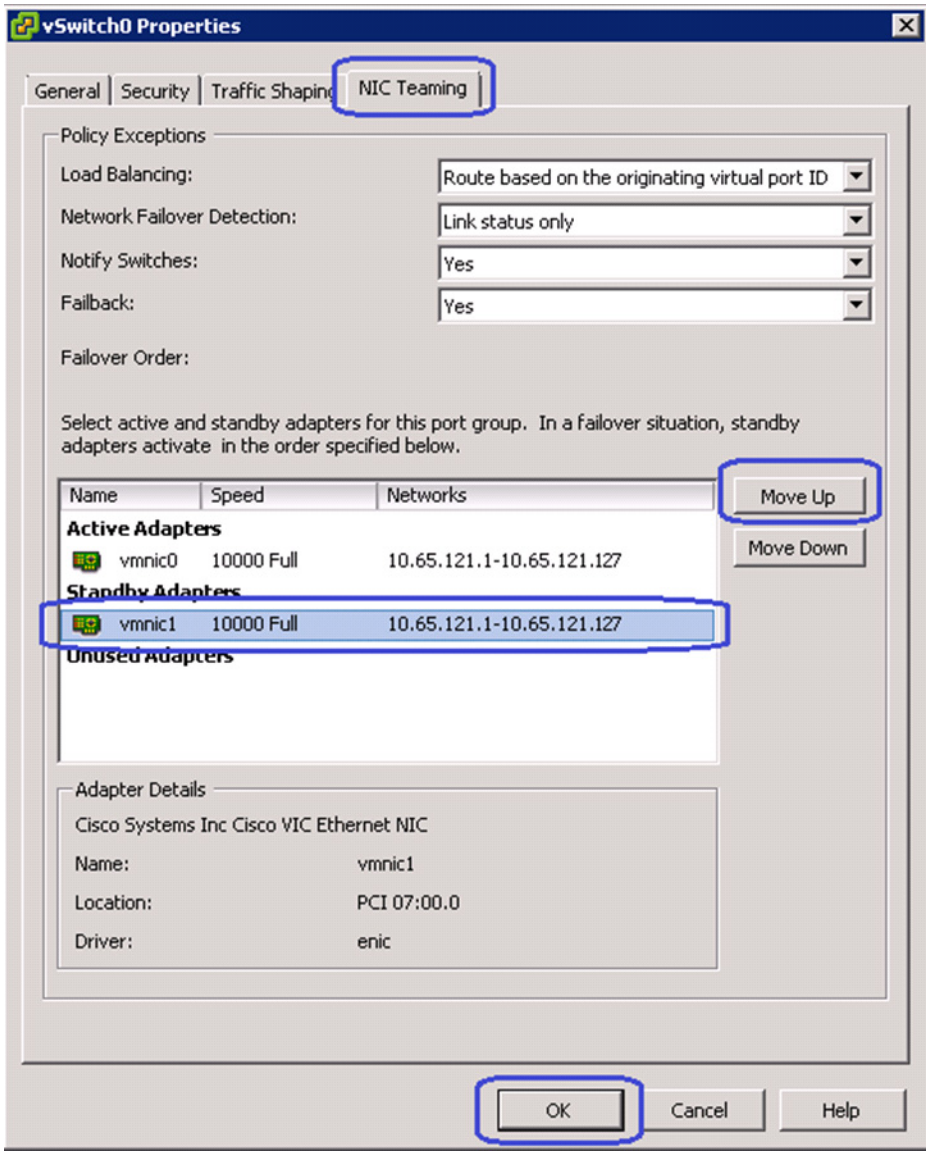

8. In the vSwitch0 configuration window, click the **Ports** tab, and click **Add**.

Figure 187 Adding Ports to vSwitch0

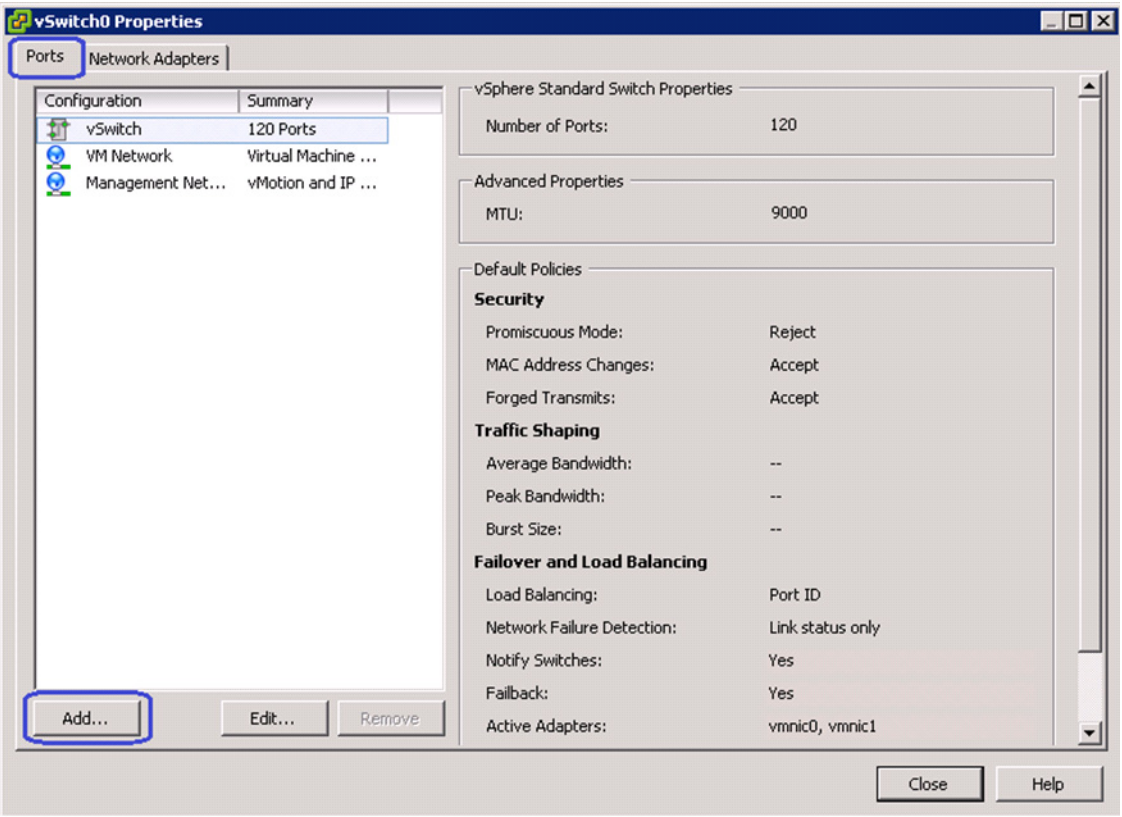

9. Click the **VMKernel** radio button in the Connection Types area and click **Next** in the Add Network Wizard window.

Figure 188 Specifying Connection Type

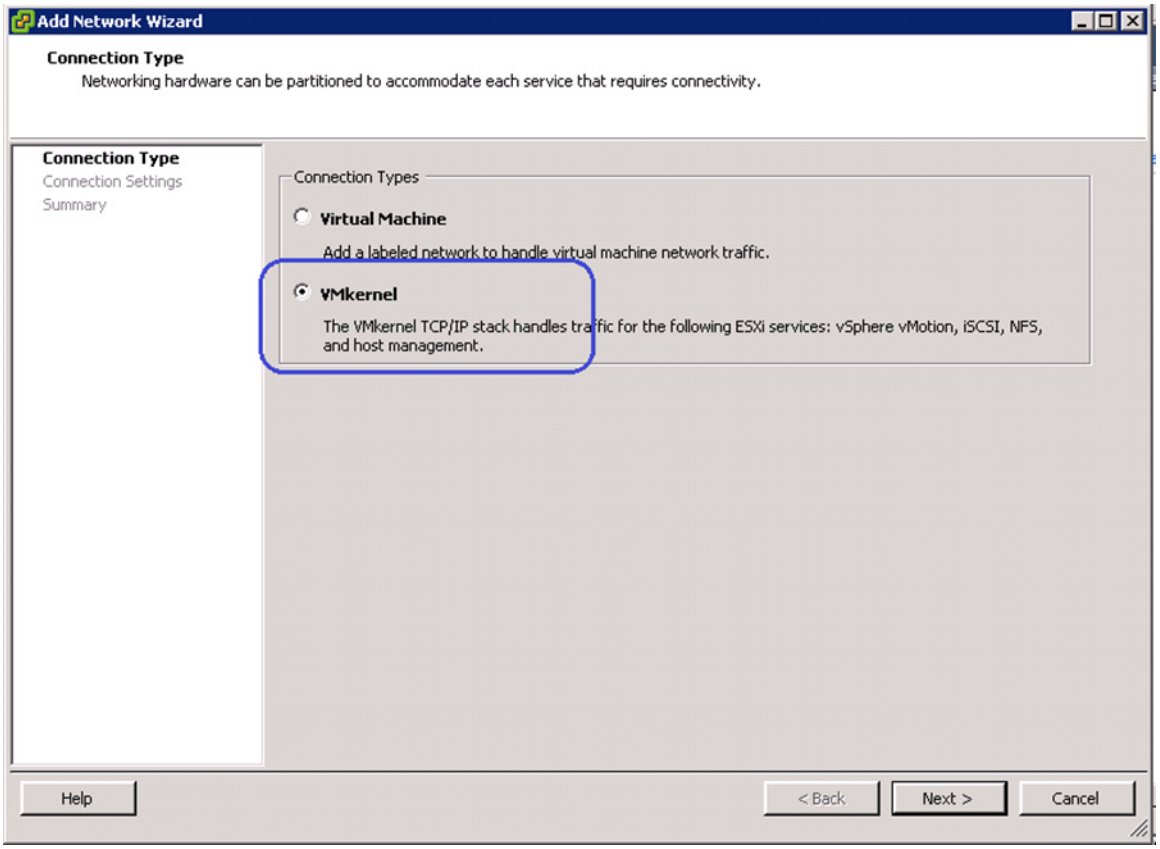

10. Enter vMotion in the Network Label field. Choose the VLAN ID from the drop-down list. The standard vSwitch0 carries both management and vMotion VLANs. Management traffic leaves vSwitch0 untagged using the native VLAN of the vNIC. The vMotion traffic must be tagged with the appropriate VLAN ID. Check the Use this port group for vMotion check box.

Figure 189 Specifying Port Group Properties for Setting VMkernel Connection

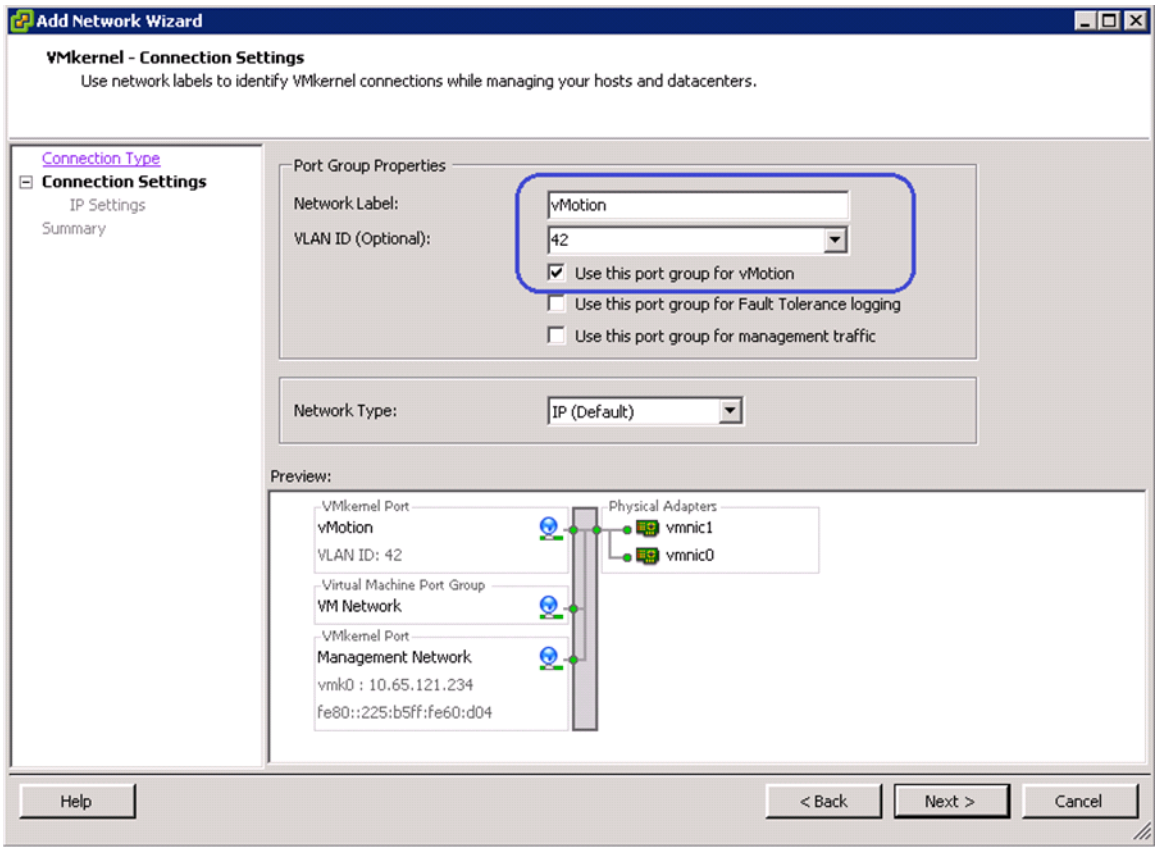

11. Configure IP address and subnet mask for the vmkernel interface. Click **Next** and deploy the vmkernel.

Figure 190 Specifying IP Address and Subnet Mask for Setting VMkernel Connection

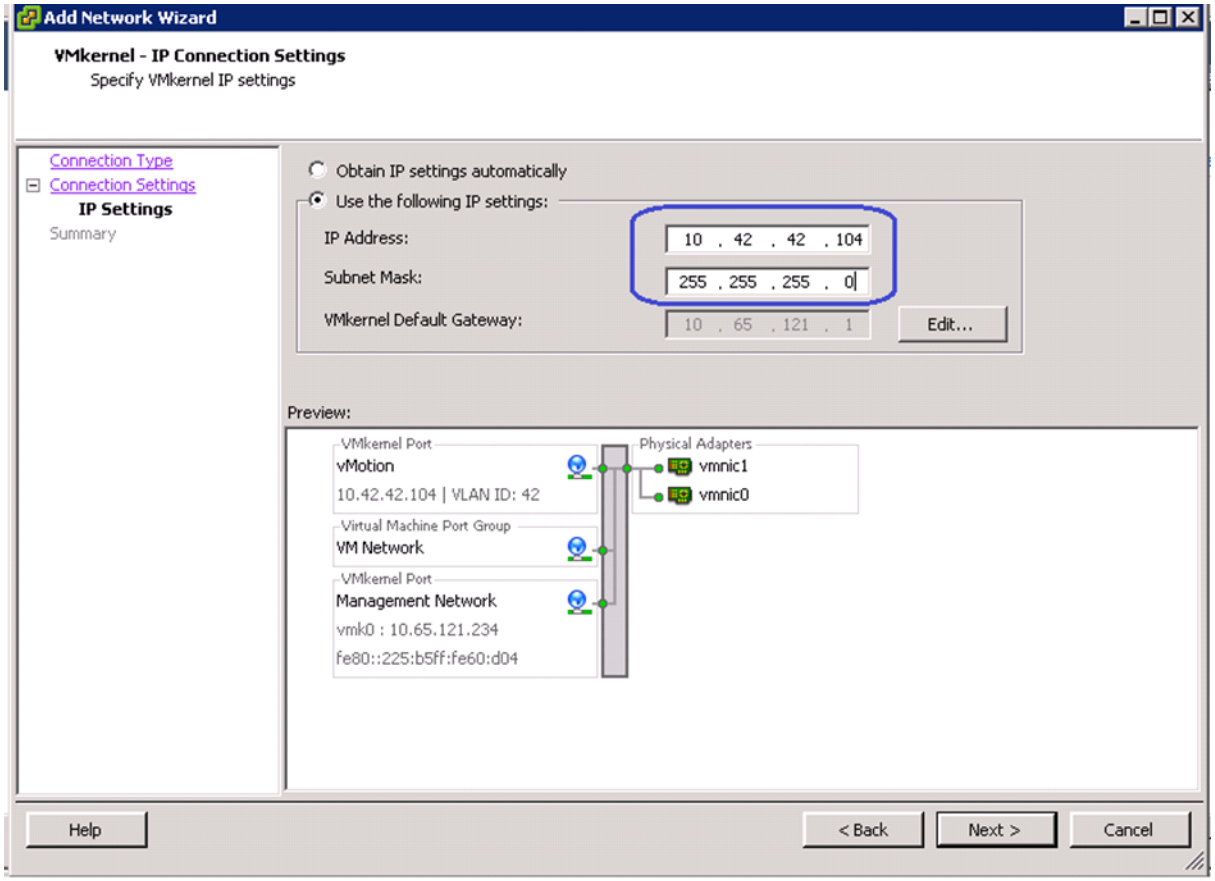

12. In the vSwitch0 properties window, select the newly created vMotion port group and click **Edit**.

Figure 191 Editing vMotion vSwitch0 to Set Jumbo MTU

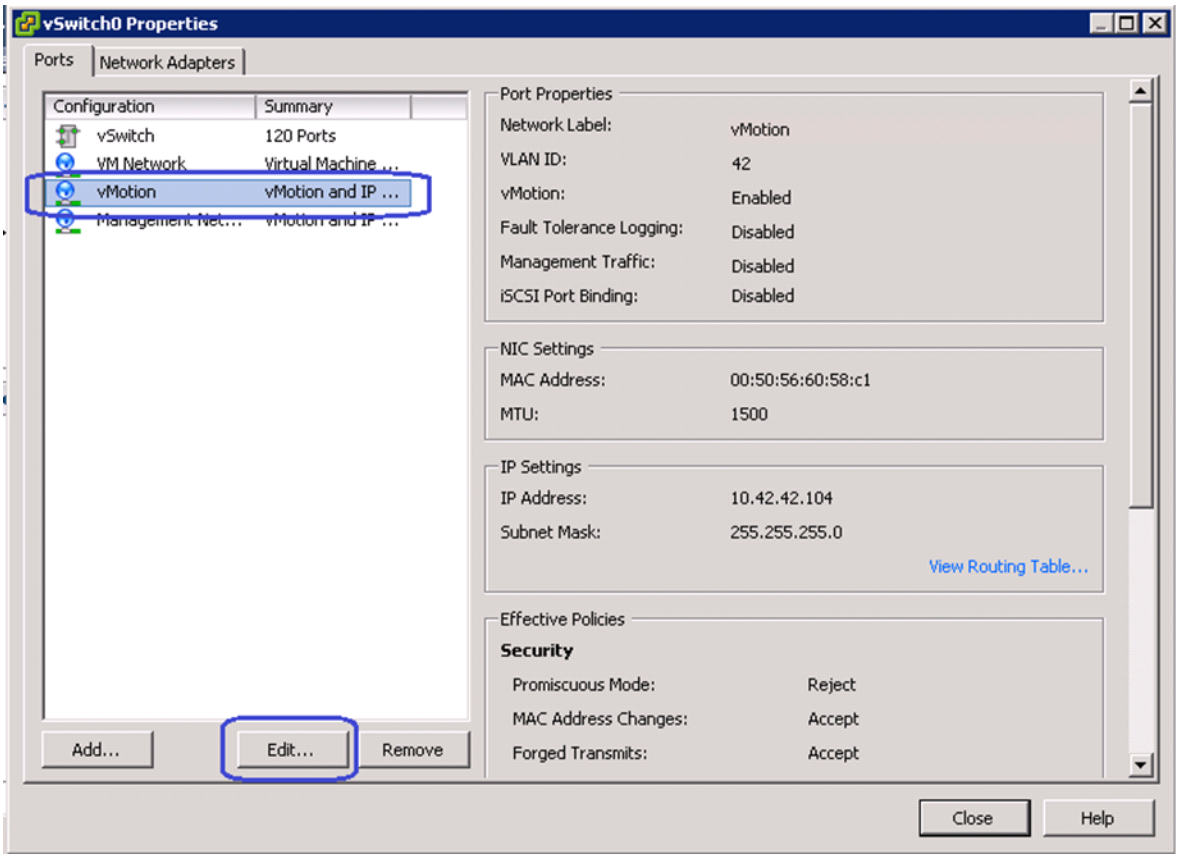

13. Set the MTU to 9000 and click **OK**. Click **Close**.
14. Repeat steps 1 to 13 for all the ESXi hosts in the cluster. Once all the ESXi hosts are configured, you must be able to ping from one host to another on the vMotion vmkernel port with jumbo MTU. Validate this by issuing ping with IP's don't fragment.

Figure 192 Pinging the VMKernel Port with Jumbo MTU

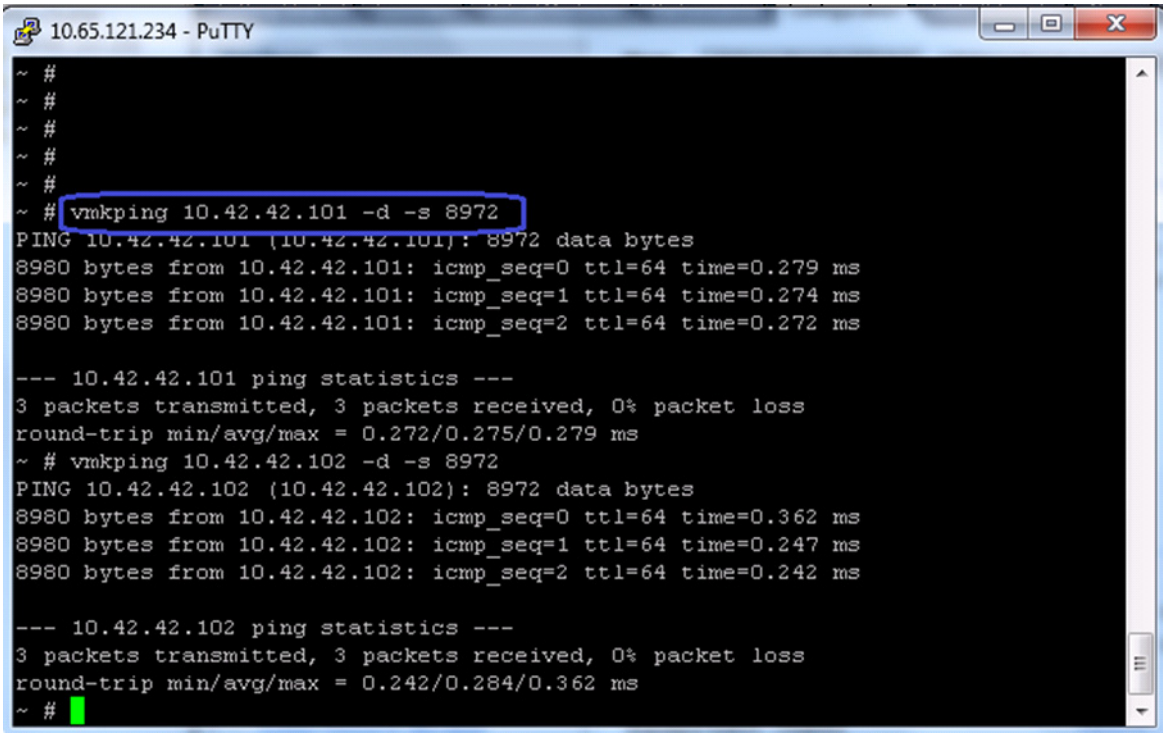

15. In the vCenter GUI, choose **Home > Inventory > Hosts and Clusters**. In the Hosts and Clusters window, click the **Configuration** tab on the right pane of the window. Click **Networking** in the Hardware area. Click **Add Networking**.

Figure 193 Add Networking in VMware vSphere Client

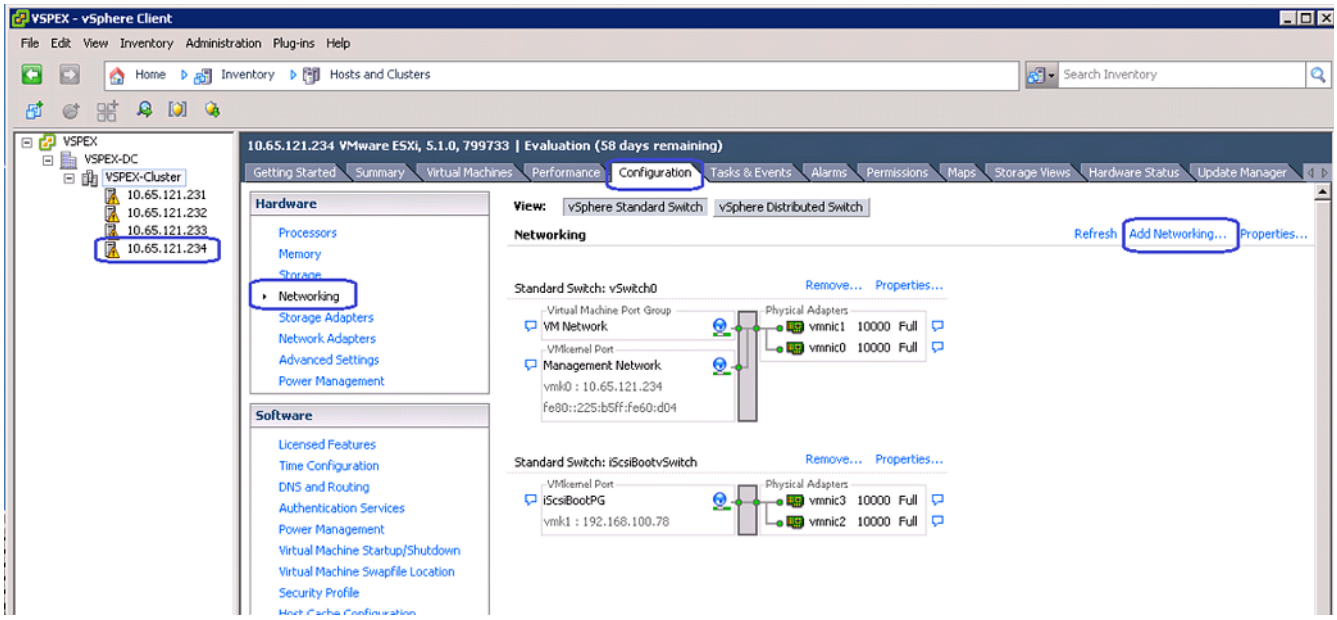

16. Click the **Virtual Machine** radio button in the Add Networking Wizard, click **Next**.

Figure 194 Specifying the Connection Type

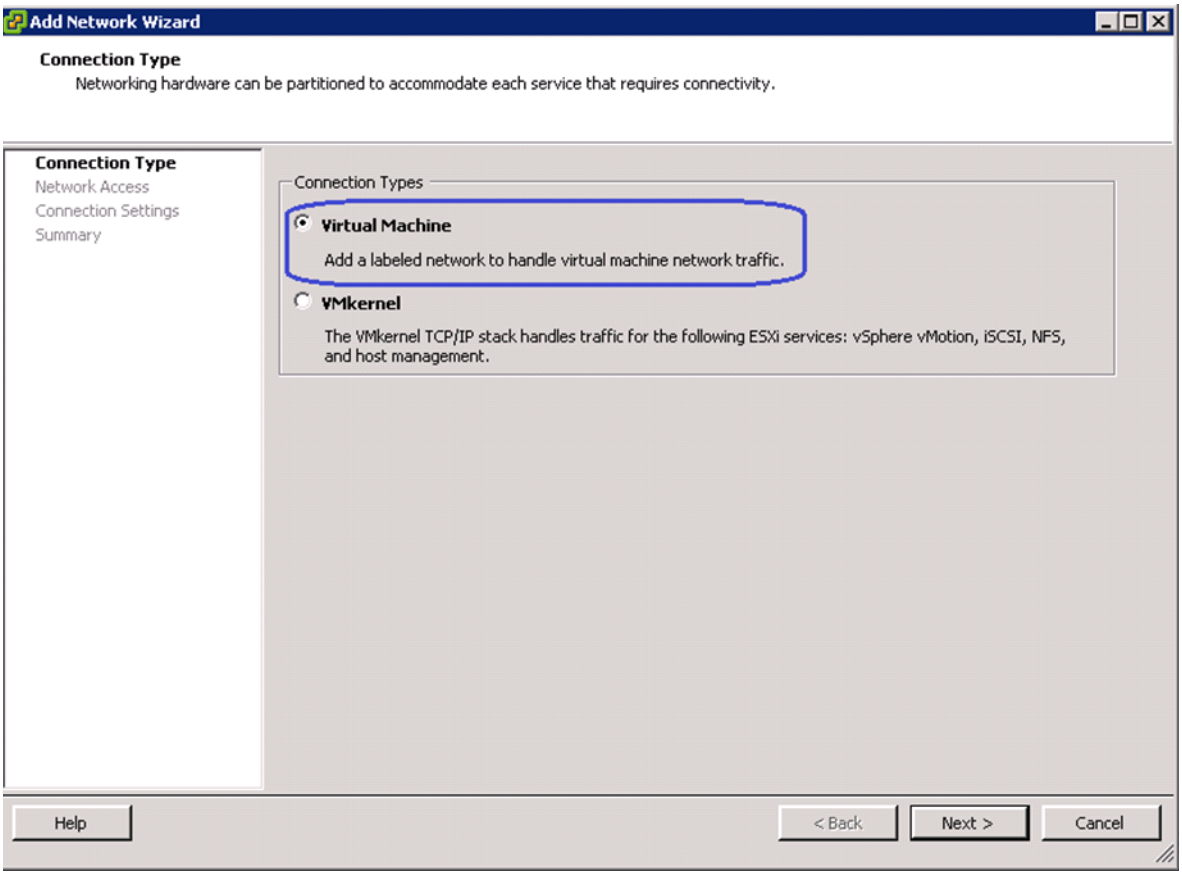

17. Select the two vmnics corresponding to the VM-Data vNICs by checking the check boxes and click **Next**.

Figure 195 Selecting the vSphere Standard Switch for Handling Network Traffic

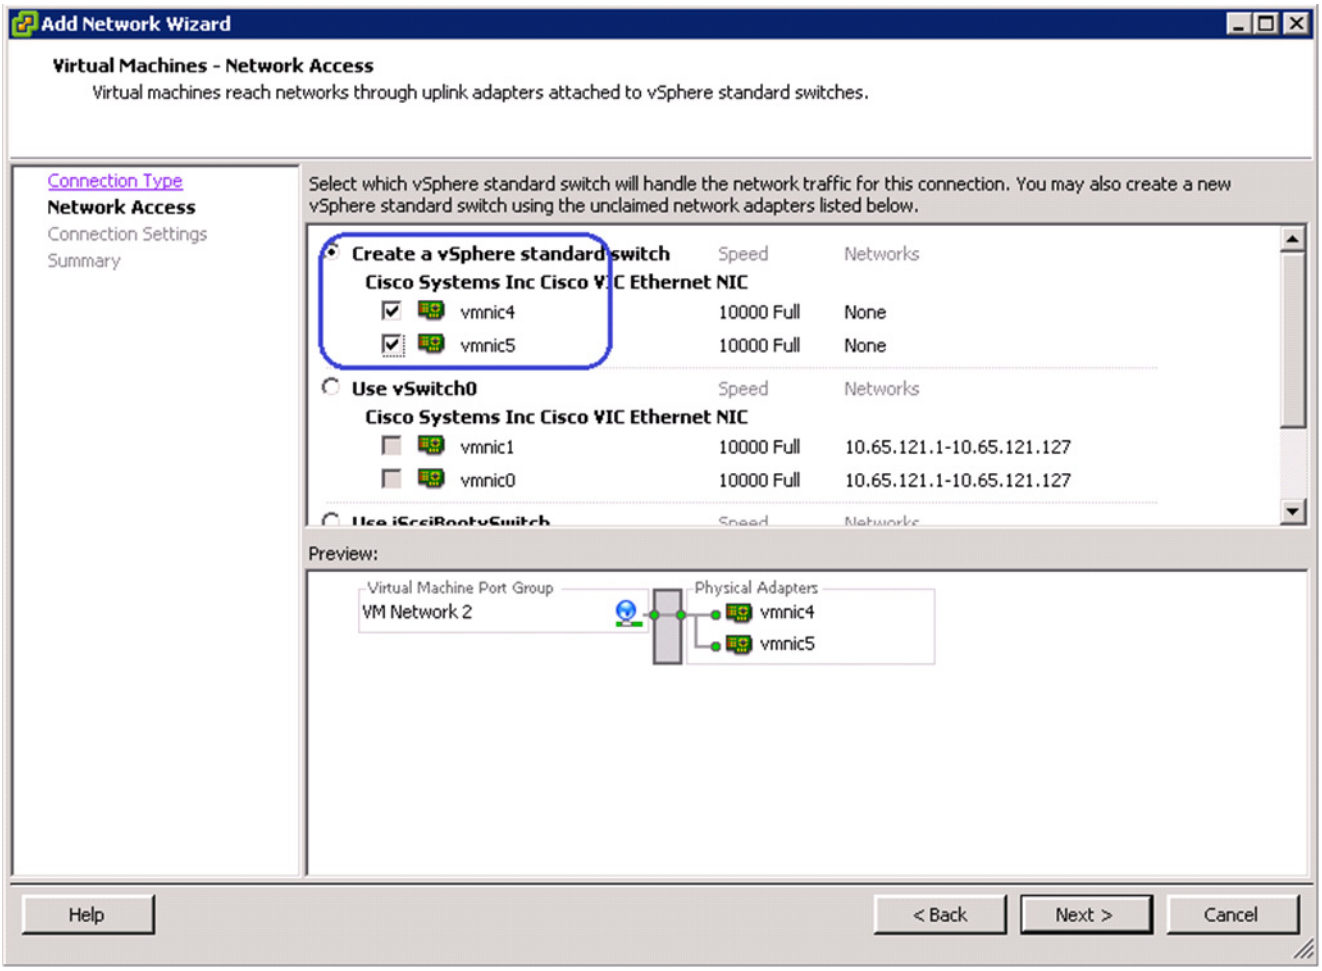

18. Enter VM-Data in the Network Label field, and keep VLAN ID as None (0) to signify absence of VLAN tag. Click **Next**.

Figure 196 Specifying Port Group Properties

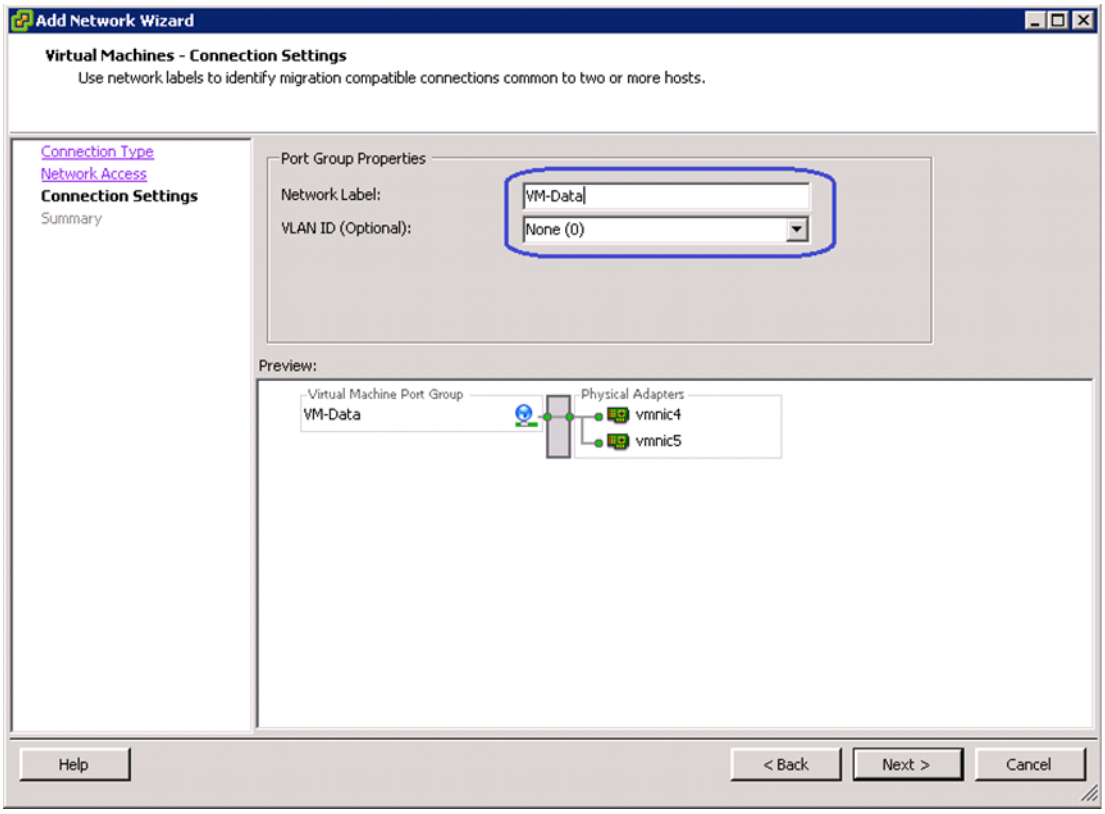

This concludes the Virtual Networking configuration on vCenter. Repeat steps 15 to 18 for all the ESXi hosts in the cluster.

Configure Storage for VM datastores, Install and Instantiate VMs from vCenter

This section details the VM datastore creation for the FC-variant of the solution. Refer to Figure 8 for Storage Architecture for 125 VMs on VNX5300 to have a high-level overview of storage architecture. Follow these steps to implement the same:

1. Login to EMC VNX Unisphere, and click the **Storage** tab. Choose **Storage Configuration** > **Storage Pools** and click the **RAID Groups** tab. Click **Create**.

Figure 197 Creating Storage Pool

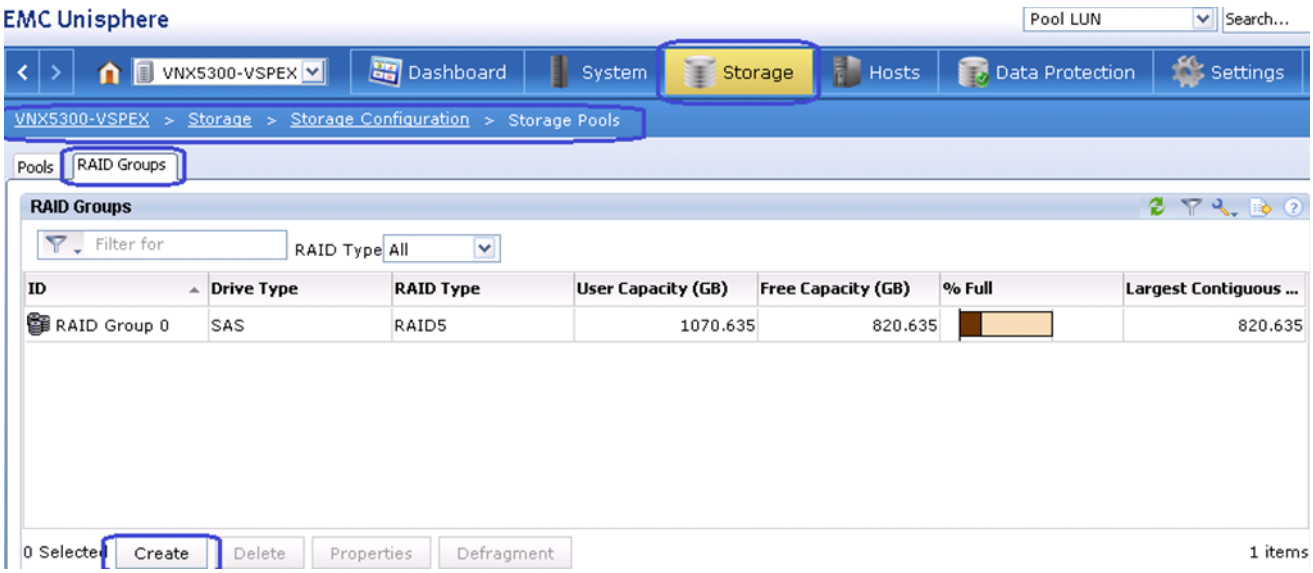

2. In the create Storage Pool wizard, click the **RAID Group** radio button for Storage Pool Type in the Storage Pool Parameters area. Keep the radio button **Automatic** selected in the Disks area.

Figure 198 Specifying Storage Pool Parameters

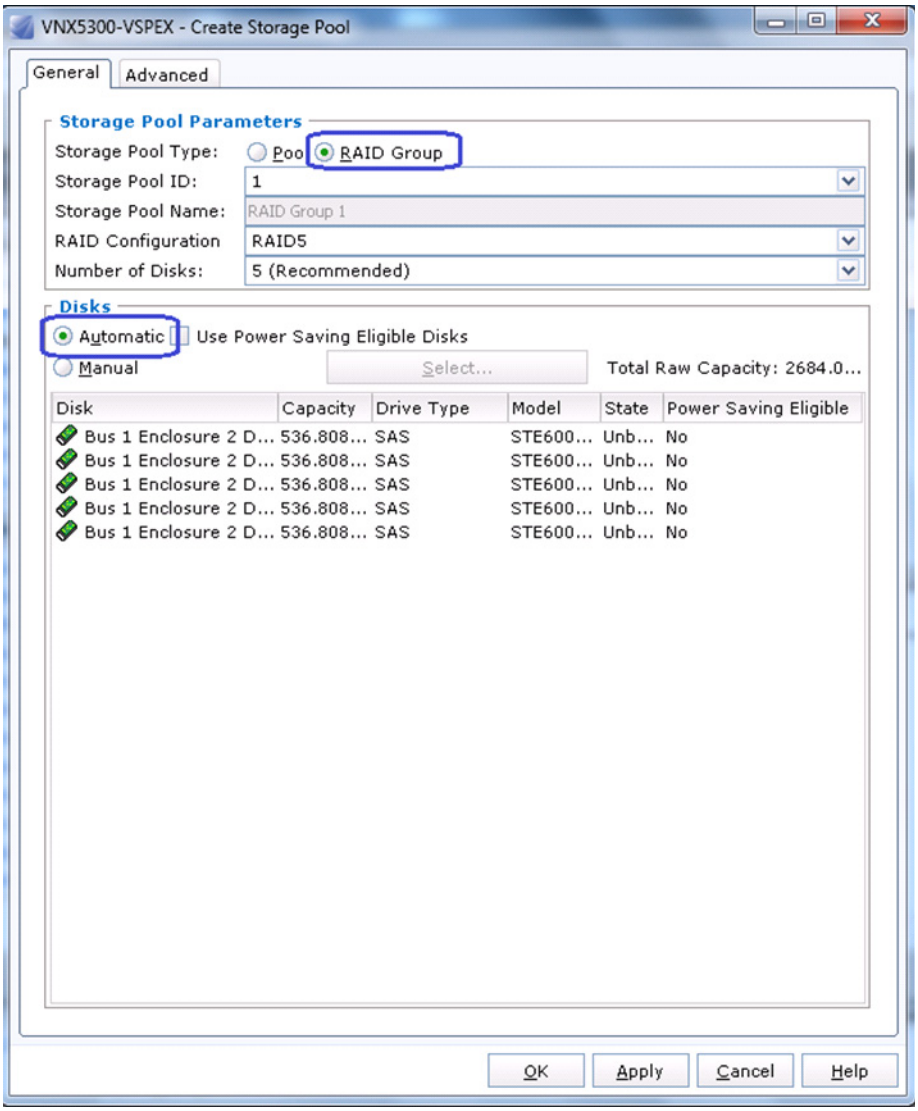

3. Repeat this step for 13 more times to create total 14 RAID5 groups for VM data storage.

Figure 199 Summary of the Configured RAID Groups

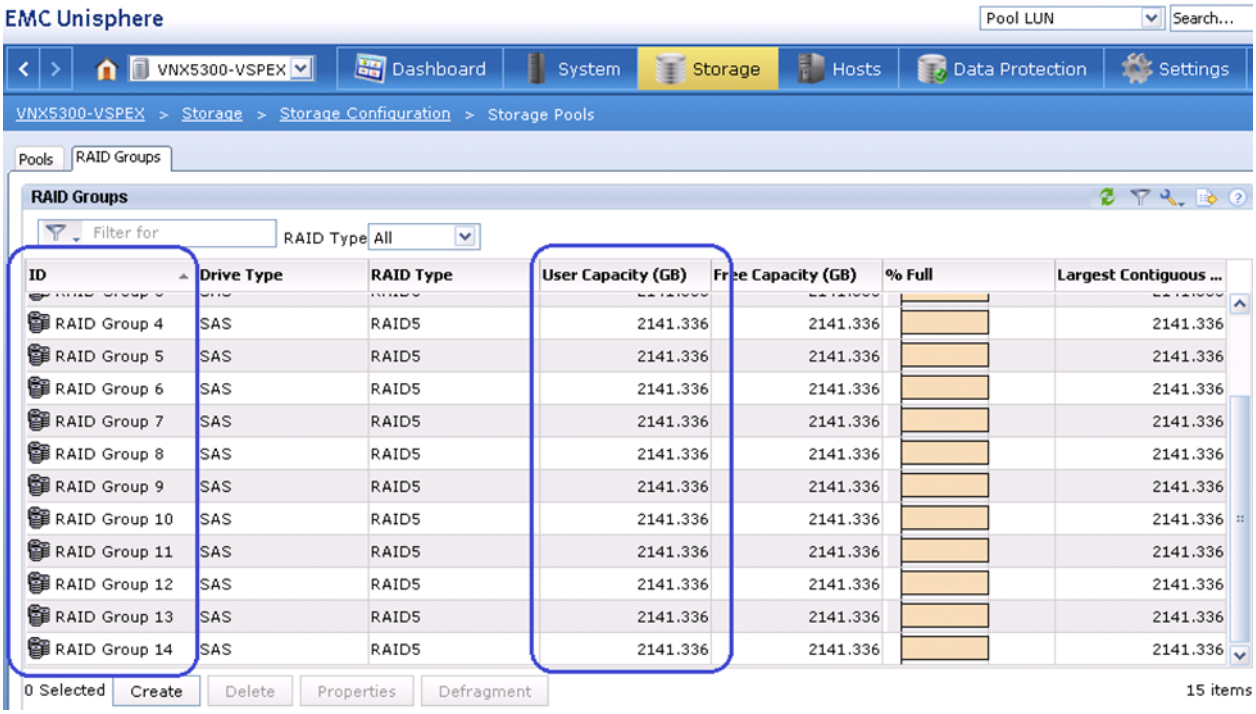

4. Choose **Storage > LUNs**. You will see five boot LUNs created for five hosts. Click **Create** to create the LUN for VM datastore.

Figure 200 Creating LUNs for VM Datastore

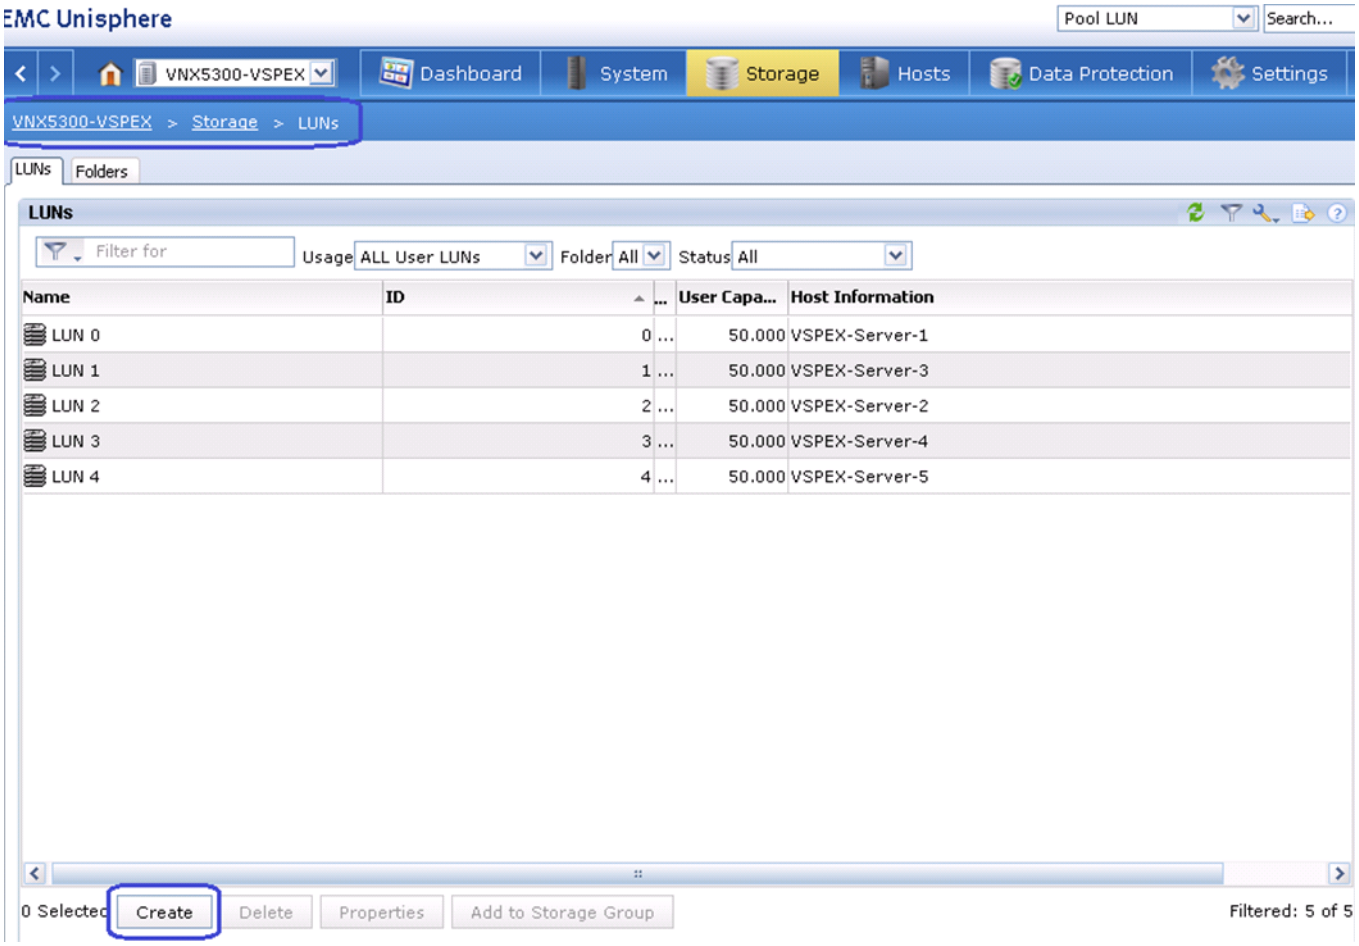

5. Click the **RAID Group** radio button for Storage Pool Type. For Storage Pool for new LUN choose 1 (VM data RAID group ID) from the drop-down list. Choose the User Capacity as MAX from the drop-down list to create one LUN per RAID group, and click **Apply**.

Figure 201 Specifying Details for Creating LUN

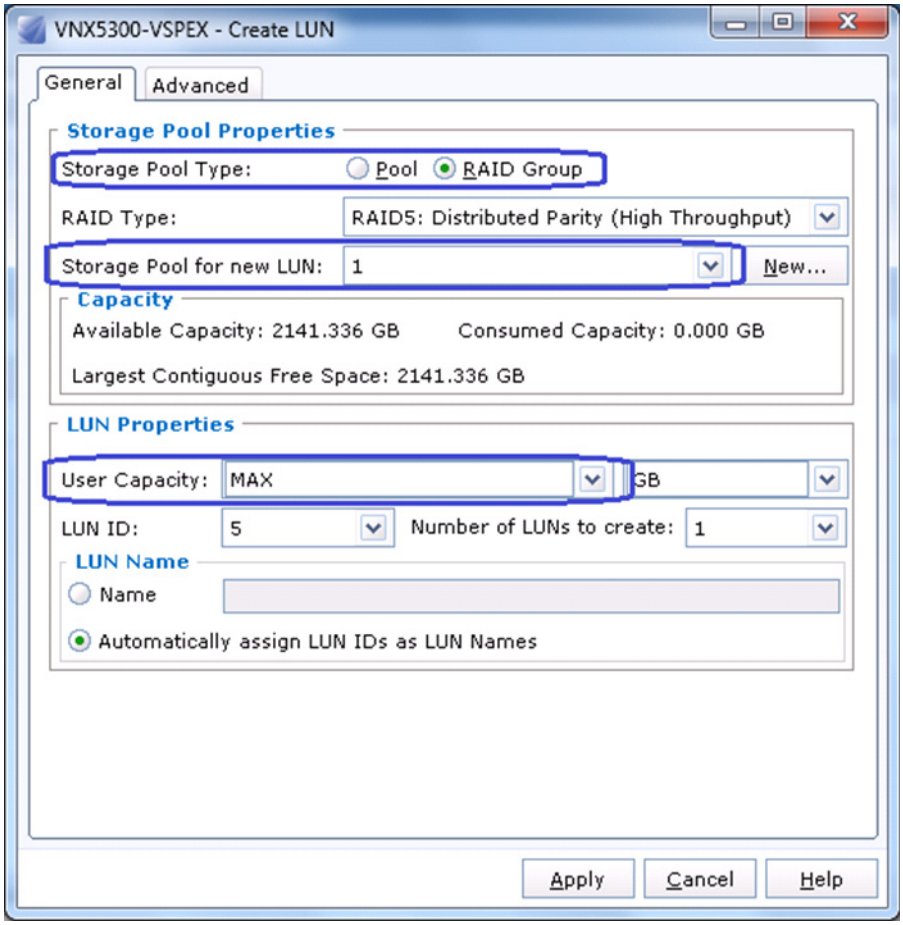

6. Repeat step 5 for all 14 RAID groups in the system. After completing this, the end result is as shown in Figure 202.

Figure 202 EMC Unisphere Showing Created LUNs

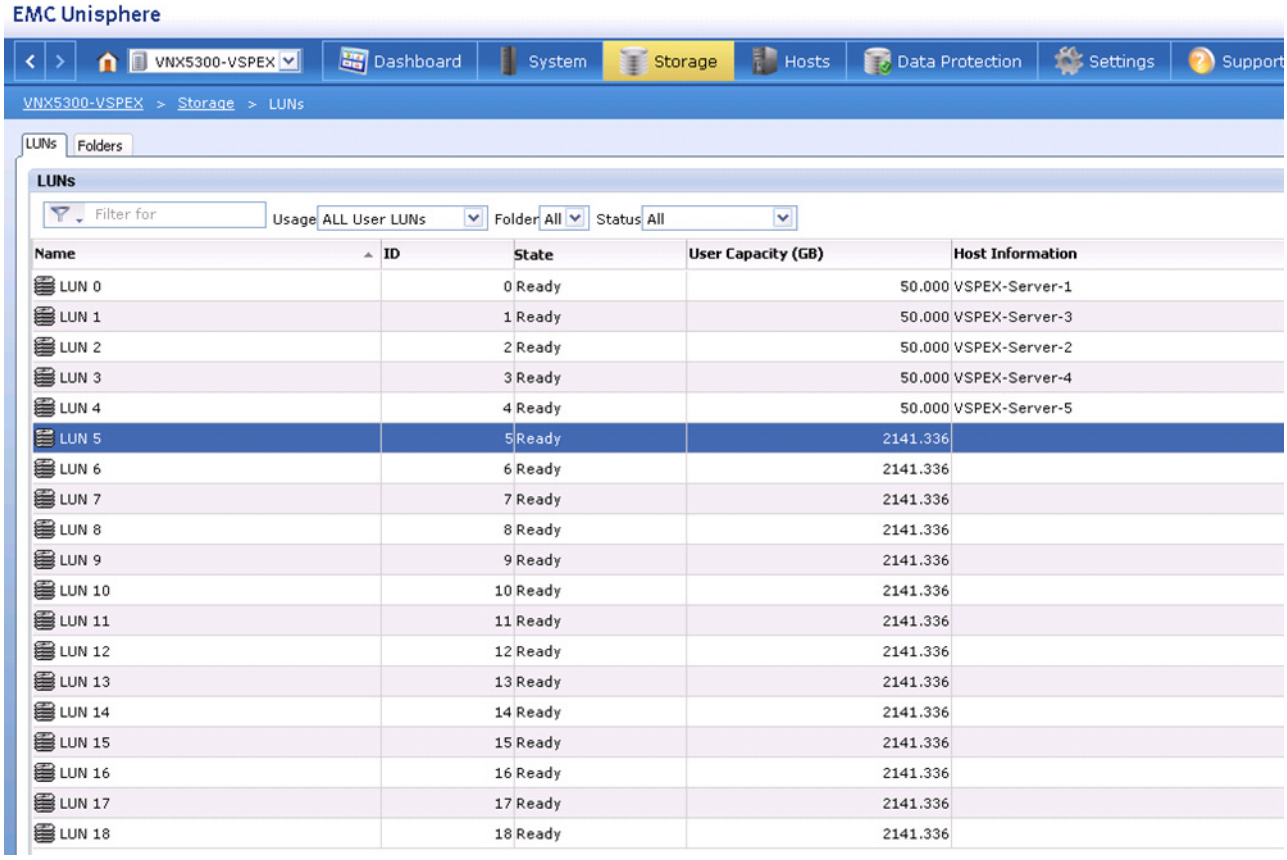

7. Select all the newly created LUNs, and click **Add to Storage Group**.

Figure 203 Adding LUNs to Storage Group

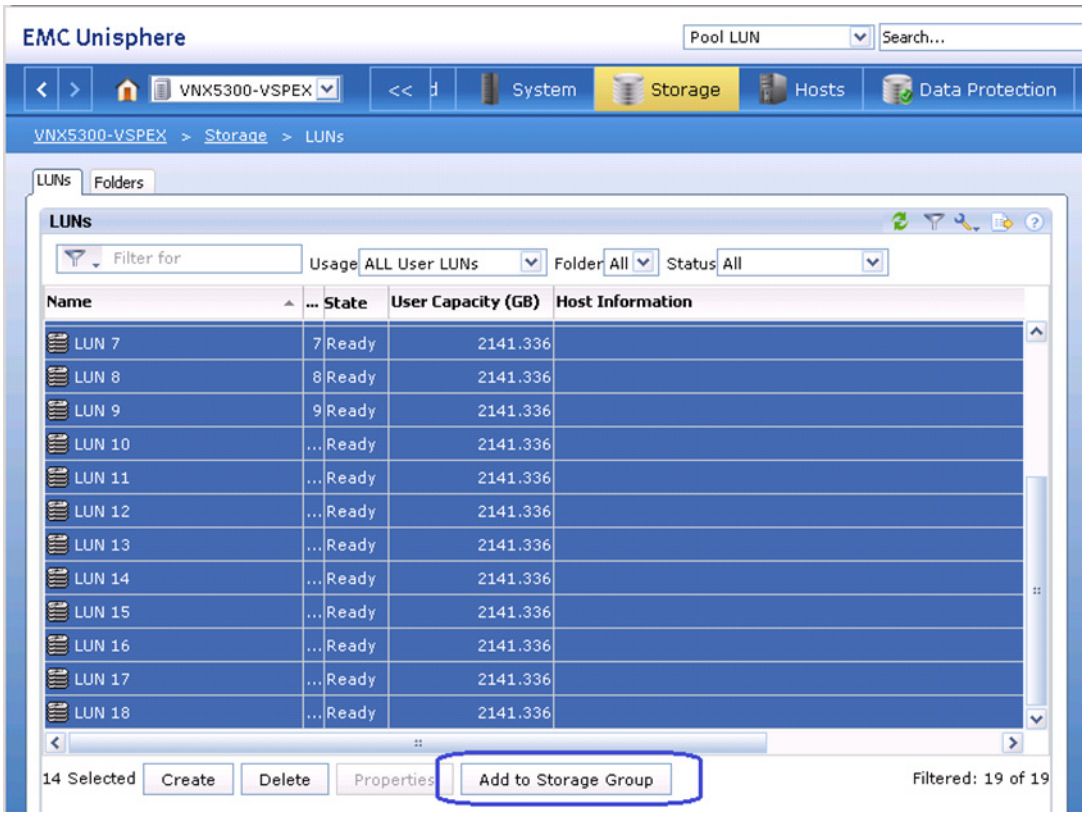

8. Select all the five servers under Available Servers, and move them under Selected Storage Groups by clicking on . This would allow all ESXi hosts to see the datastore, which is essential for the vMotion of VMs across the cluster.

Figure 204 Selecting the Storage Groups

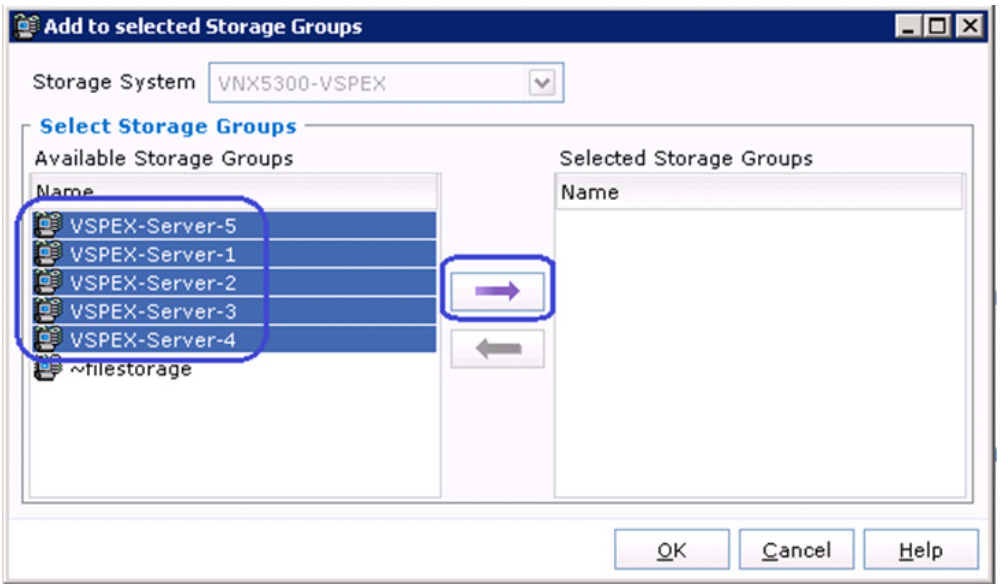

9. In the **LUNs** tab, you will see the storage group (and hence host) access for LUNs and their State showing Ready.

Figure 205 Summary of the Configured LUNs

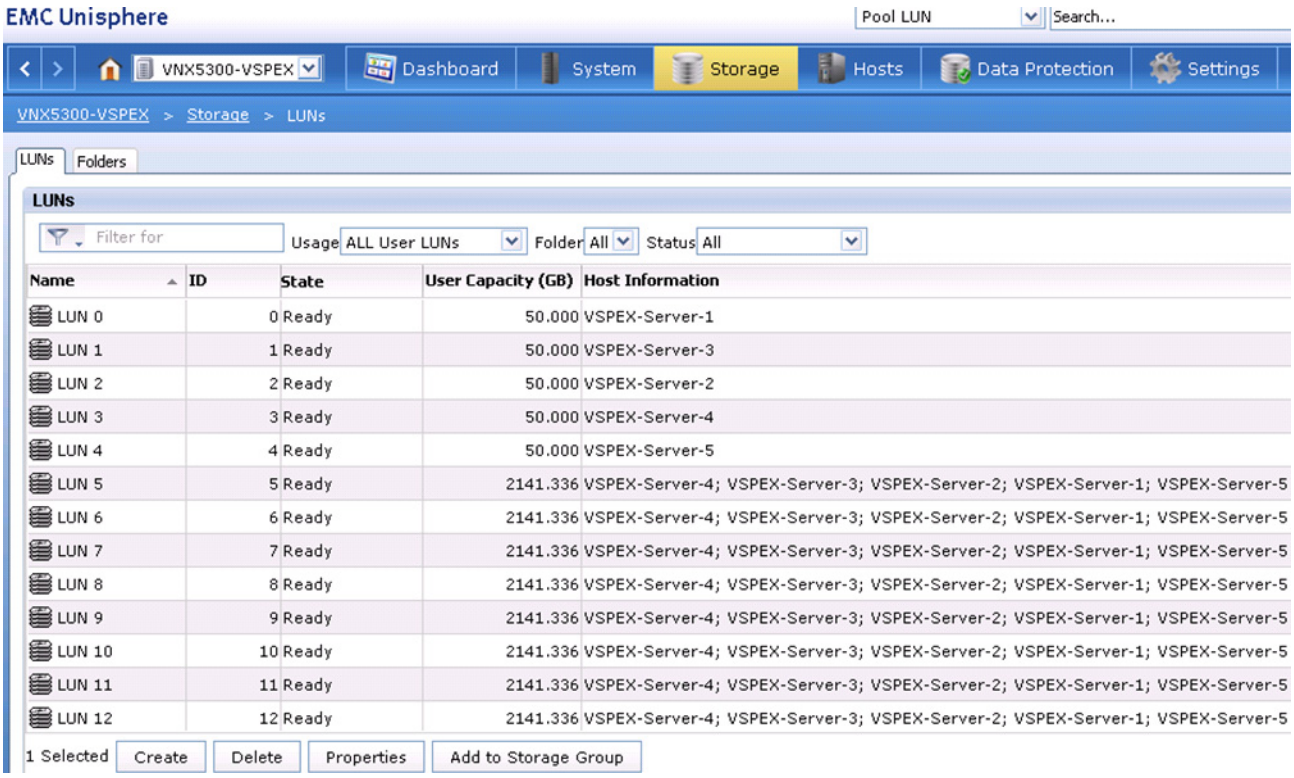

10. Login to vCenter GUI, select a host from the Hosts and Clusters window, click **Configuration** and then **Storage**. Click **Add Storage**.

Figure 206 Adding Storage in VMware vSphere Client

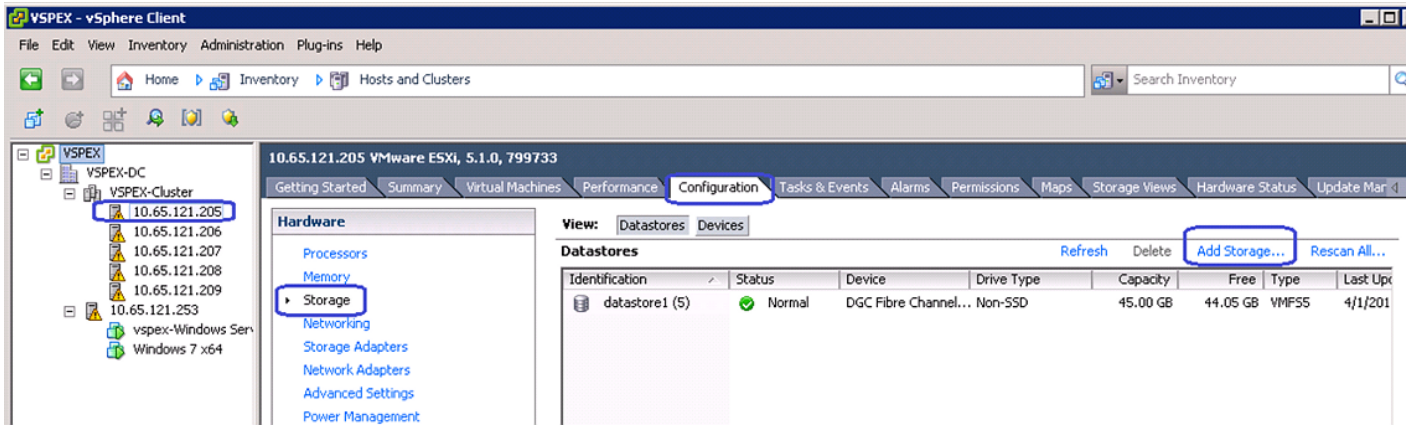

11. In the Storage Type area, click the **Disk/LUN** radio button in the Add Storage wizard. Click **Next**.

Figure 207 Selecting a Storage Type

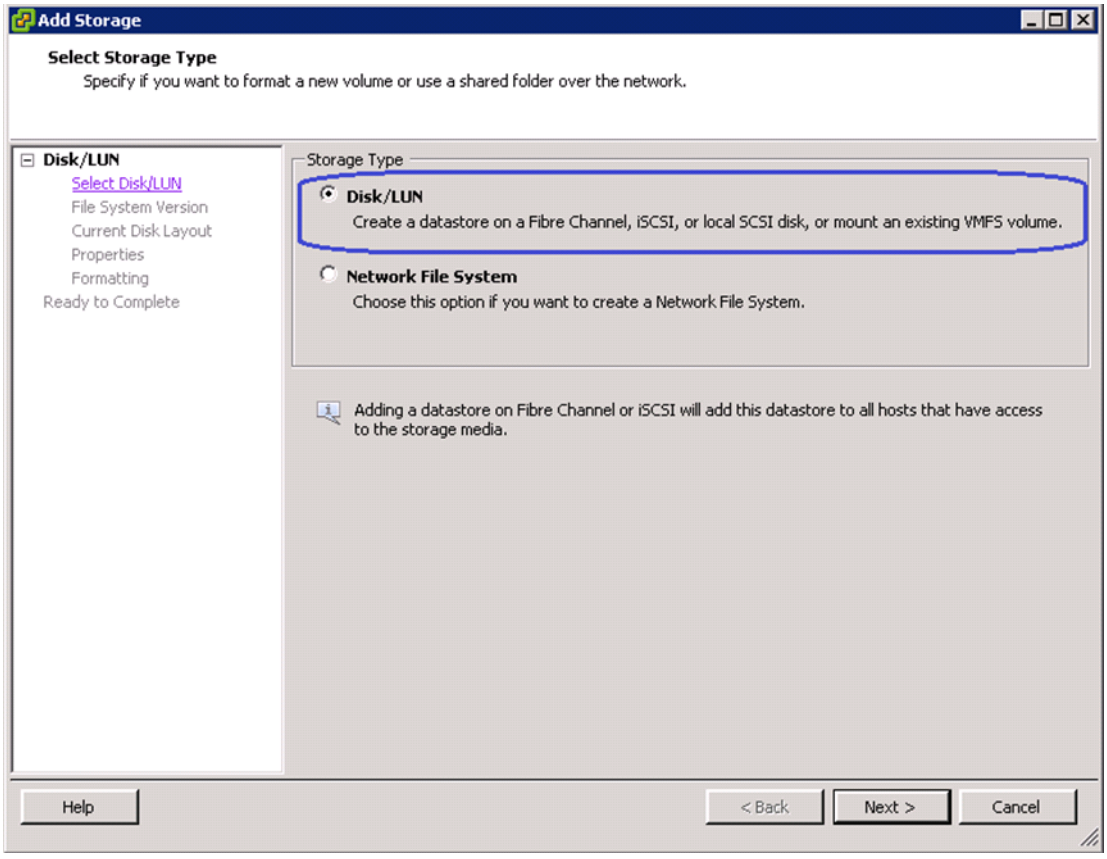

12. Select the first DGC Fibre Channel Disk (..) from the list and click **Next**.

Figure 208 Selecting the First DCG Fibre Channel Disk

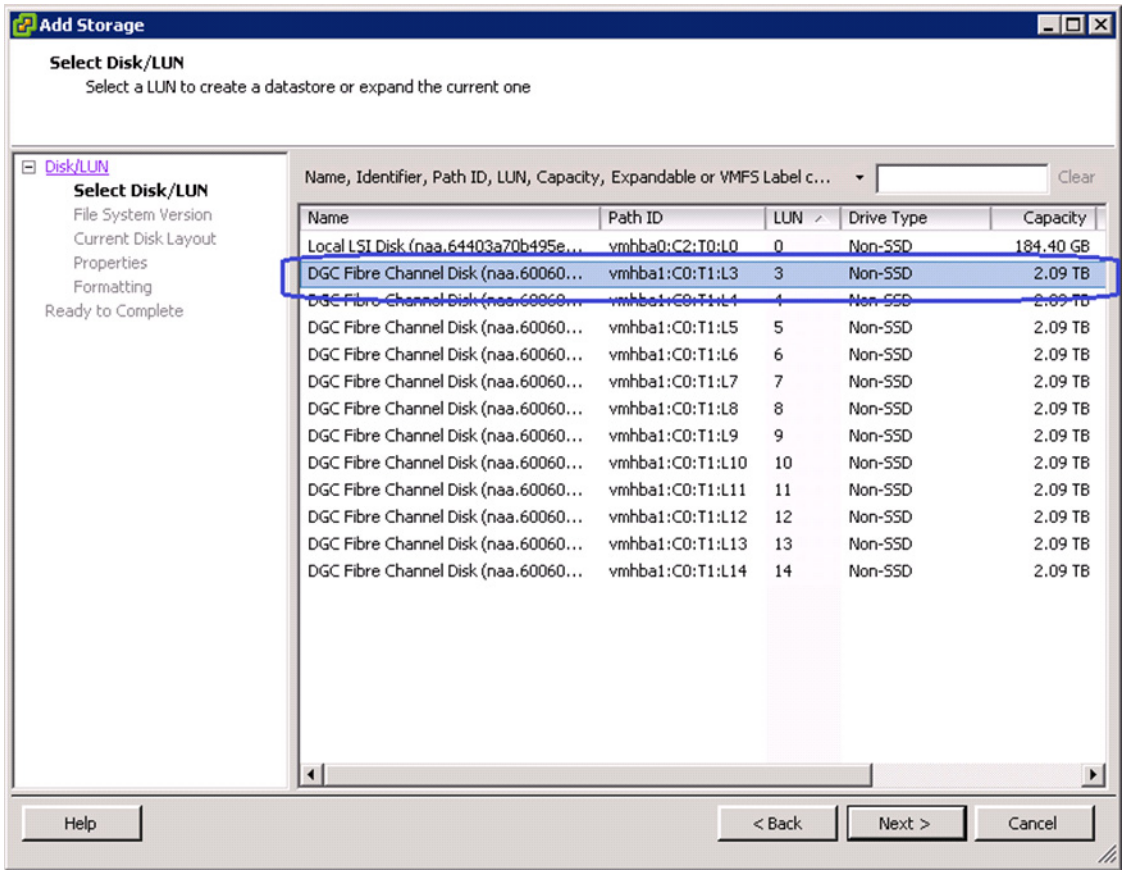

13. Enter VM-DS1 in the datastore name field as the first datastore and click **Next** and then **Finish**.

Figure 209 Specifying Datastore Name for Current Disk Layout

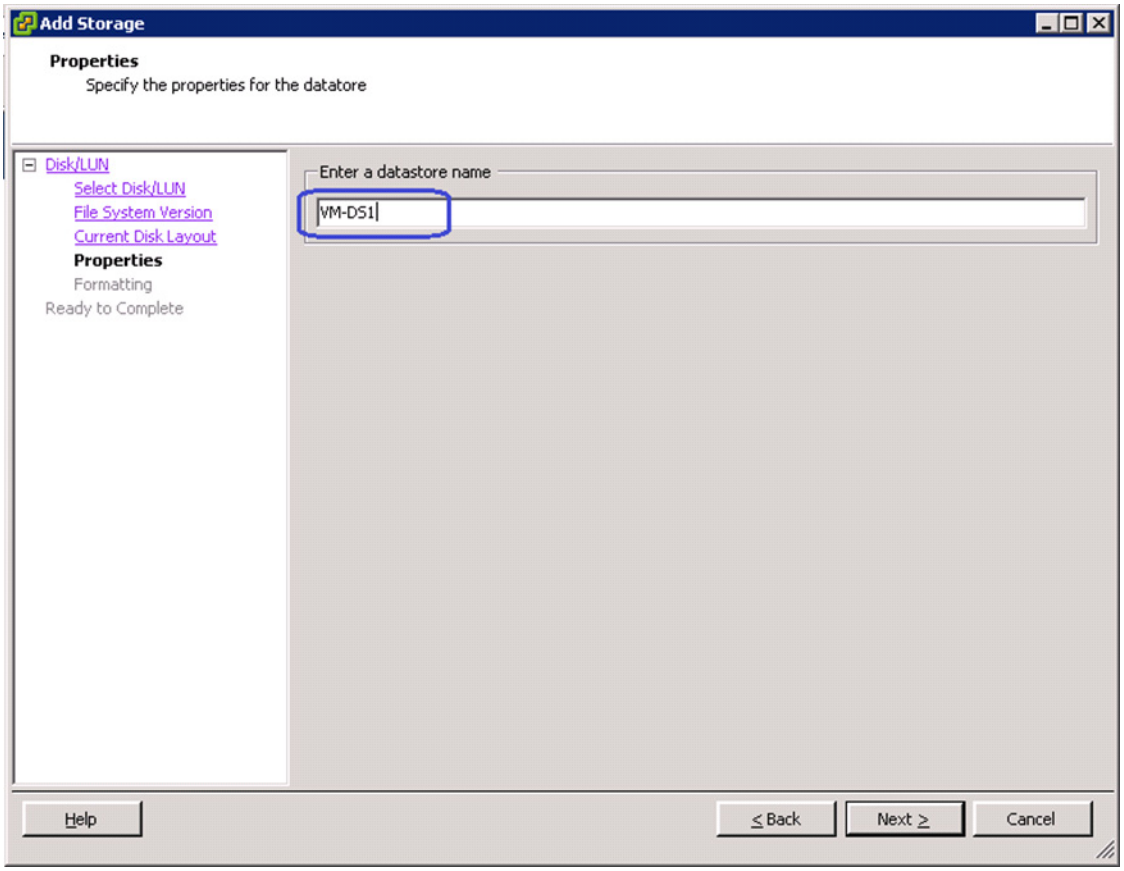

14. Repeat steps 10 to 13 for all the 14 datastores. Once datastores are added to one host, it would automatically show up for the other hosts too. The end result is as shown in Figure 210 on each host.

Figure 210 DataStores Added to a Host

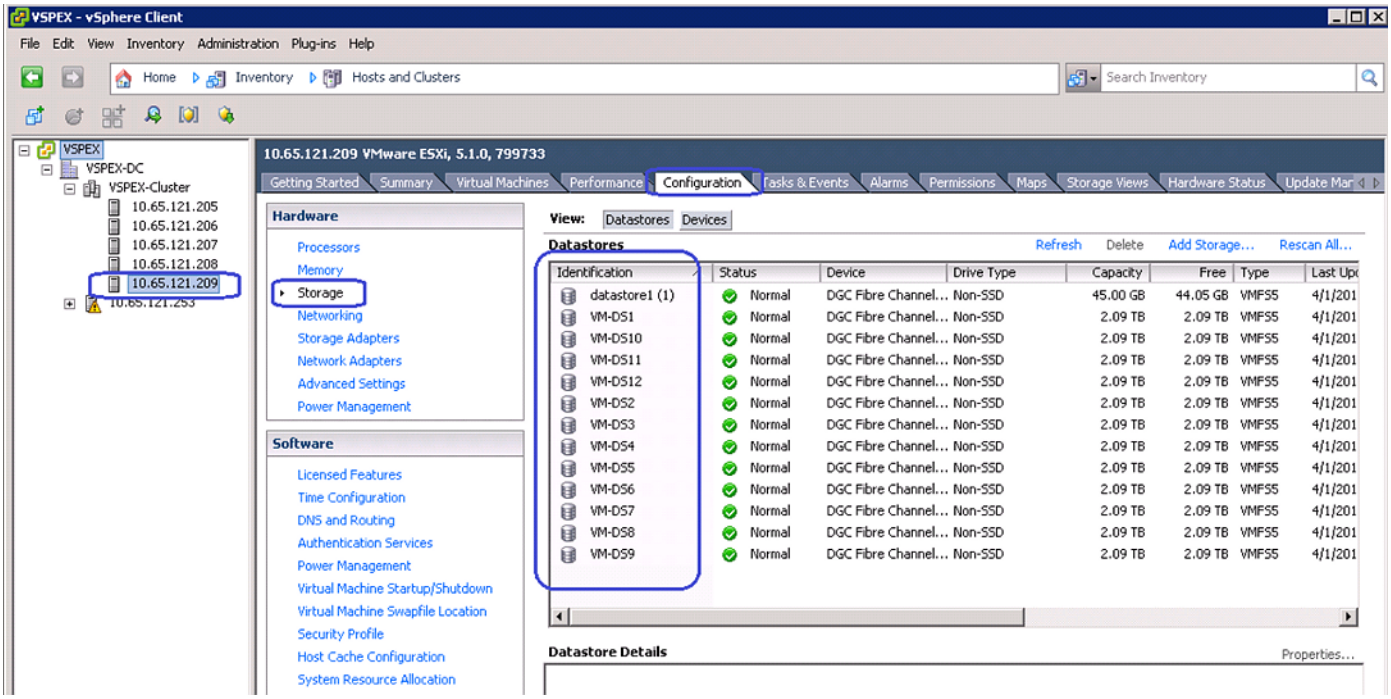

iSCSI-variant

In section "Prepare and configure storage array for resource pools and iSCSI servers", we demonstrated how to create resource pool and iSCSI servers on VNXe storage array. Later in the "Configure datastores for ESXi images" section, we demonstrated how to create datastore. You need to deploy 10 additional datastores with 6+1 RAID5 pools for Virtual Machines using the same steps. Please refer to the storage architecture section for the size of the datastore. Following table summarizes the relationship between resource pool, iSCSI servers and datastores:

Table 10 Relationship between resource pool, iSCSI servers and datastores

Data-Resource-pool 1	iSCSI Servers Data-A: Active on SP-A	IP1: eth10	Data-Store-1
		IP2: eth11	Data-Store-2
			Data-Store-3
			Data-Store-4
			Data-Store-5
	iSCSI Server Data-B: Active on SP-B	IP3: eth10	Data-Store-6
		IP4: eth11	Data-Store-7
			Data-Store-8
			Data-Store-9
			Data-Store-10

Make sure that Data-Store-1 and Data-Store-2 are available from all the ESXi hosts. Discover the newly created datastores from the vCenter 5.1 server by rescanning the iSCSI datastores from all the hosts in the cluster.

Next subsection explains how to deploy virtual machines using vCenter GUI.

Template-Based Deployments for Rapid Provisioning

Figure 211 Rapid Provisioning Using VM Templates

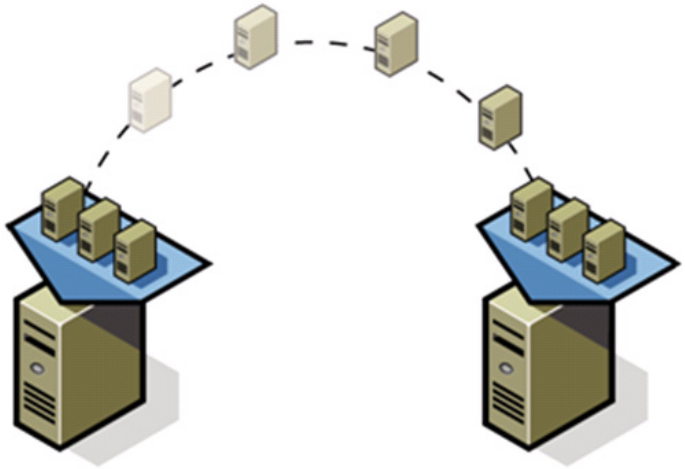

In an environment with established procedures, deploying new application servers can be streamlined, but can still take many hours or days to complete. Not only must you complete an OS installation, but downloading and installing service packs and security updates can add a significant amount of time. 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 snap-shot of VMs in a cluster would look similar to the following figure:

Figure 212 Summary of the Cluster in VMware vSphere Client

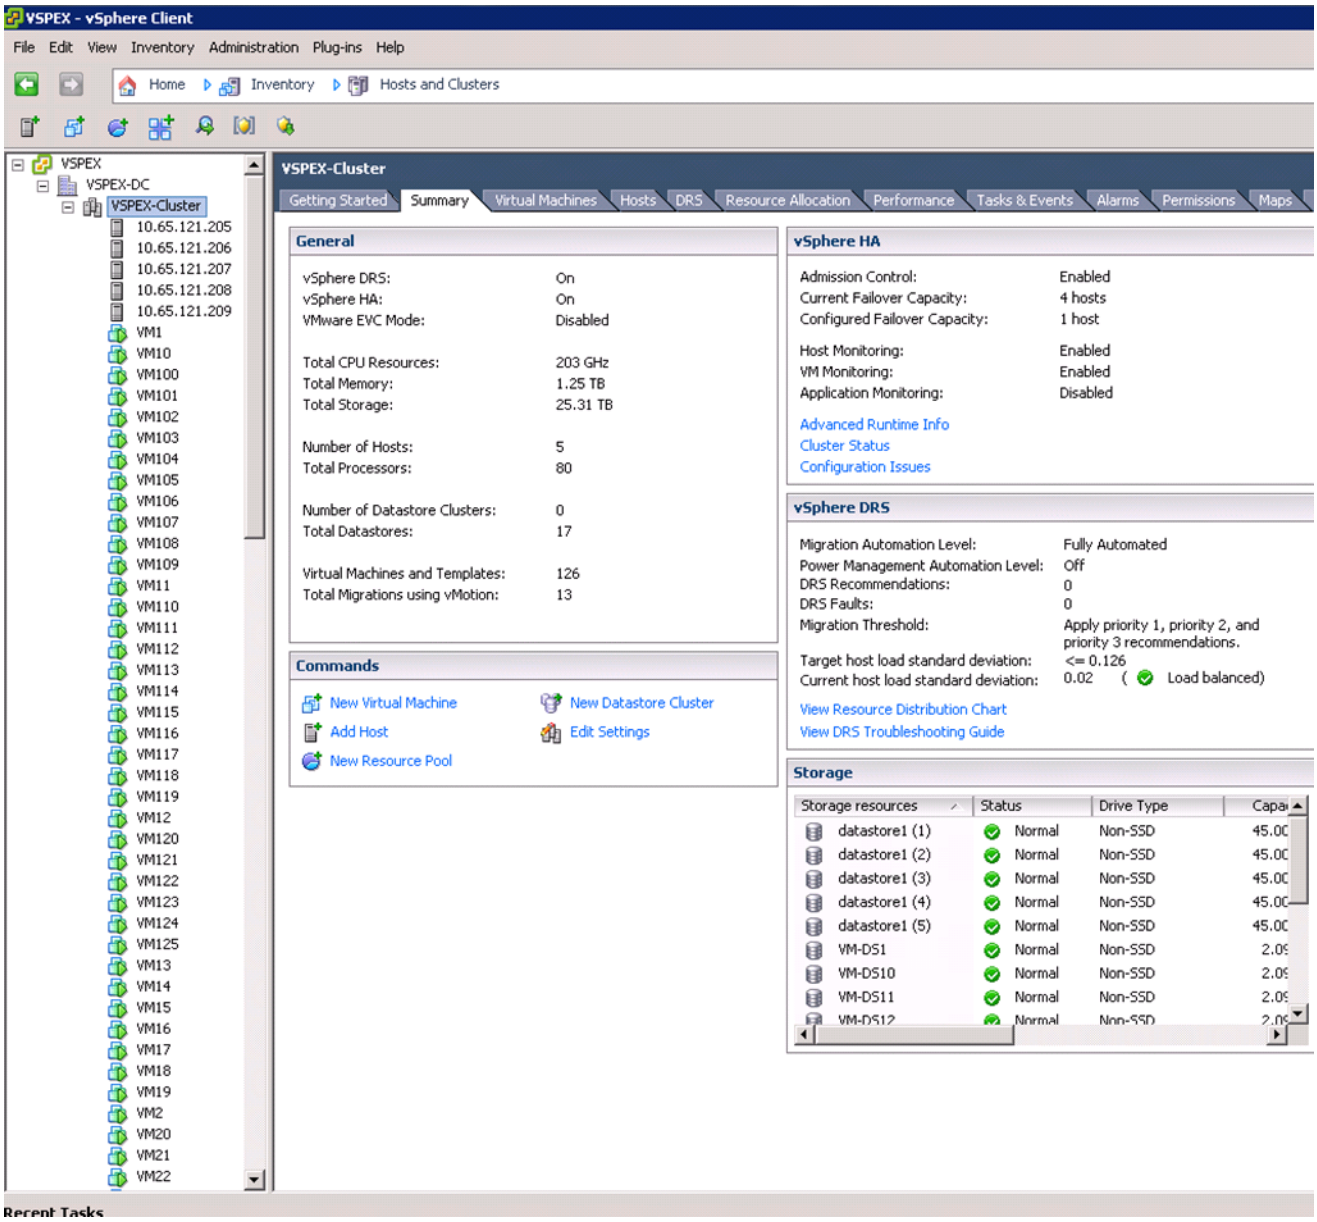

Validating Cisco Solution for EMC VSPEX VMware Architectures

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 datastores must be visible and accessible from all the hosts at all the time.

1. Administratively shutdown one of the two server ports connected to the Fabric Extender A. Make sure that connectivity is not affected. Upon administratively enabling the shutdown port, the traffic should be rebalanced. This can be validated by clearing interface counters and showing the counters after forwarding some data from virtual machines on the Nexus switches.
2. Administratively shutdown both server ports connected to Fabric Extender A. ESXi hosts should be able to use fabric 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. Reboot the active storage processor of the VNXe storage array and make sure that all the iSCSI shares are still accessible during and after the reboot of the storage processor.
6. 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.
7. 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 was used with Windows 2008 R2 SP1 server to test scaling of the solution in Cisco labs. Table 11 details on the test profile used.

Table 11 Test profile details

Profile characteristic	Value
Number of virtual machines	100 or 125 depending on the architecture
Virtual machine OS	Windows Server 2008 R2 SP1
Processors per virtual machine	1
Number of virtual processors per physical CPU core	4

Table 11 Test profile details

Profile characteristic	Value
RAM per virtual machine	2 GB
Average storage available for each virtual machine	100 GB
Average IOPS per virtual machine	25 IOPS
Number of datastores to store virtual machine disks	2
Disk and RAID type for datastores	RAID 5, 600 GB, 15k wpm, 3.5 inch SAS disks

Bill of Material

Table 12 gives the list of the components used in the CVD for 250 virtual machines configuration.

Table 12 List of hardware components used in the CVD

Description	Part #
4 x UCS C220 M3 rack servers	UCSC-C220-M3S
CPU for C220 M3 rack servers (2 per server)	UCS-CPU-E5-2650
Memory for C220 M3 rack servers (4 per server)	UCS-MR-1X162RY-A
Cisco UCS 1225 VIC adapter (1 per server)	UCSC-PCIE-CSC-02
UCS 6248UP Fabric Interconnects (2)	UCS-FI-6248UP
UCS 2232PP Fabric Extenders (2)	N2K-C2232PP-10GE
10 Gbps SFP+ multifiber mode	SFP-10G-SR

For more information on part numbers and options available for customization, see Cisco C220 M3 server specsheet at:

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/C220M3_SFF_SpecSheet.pdf

Customer Configuration Data Sheet

Before you start the configuration, gather the customer-specific network and host configuration information. Table 13, Table 14, Table 15, Table 16, Table 17, Table 18, Table 19 provide information on assembling the required network, host address, numbering, and naming information. This worksheet can also be used as a “leave behind” document for future reference.

Table 13 Common Server Information

Server Name	Purpose	Primary IP
	Domain Controller	
	DNS Primary	
	DNS Secondary	

Table 13 Common Server Information

Server Name	Purpose	Primary IP
	DHCP	
	NTP	
	SMTP	
	SNMP	
	vCenter Console	
	SQL Server	

Table 14 ESXi Server Information

Server Name	Purpose	Primary IP	Private Net (storage) addresses	VMkernel IP	vMotion IP
	ESXi Host1				

Table 15 Array Information

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

Table 16 Network Infrastructure Information

Description	IP	Subnet Mask	Default Gateway
Cisco UCSM Virtual IP address			
Cisco UCS Fabric Interconnect A			
Cisco UCS Fabric Interconnect B			

Table 17 VLAN Information

Name	Network Purpose	VLAN ID	Allowed Subnets
vSphereMgmt	Virtual Machine Networking ESXi Management		
Storage (A)	iSCSI VLAN on fabric A (iSCSI-variant only)		
Storage (B)	iSCSI VLAN on fabric B (iSCSI-variant only)		
vMotion	vMotion traffic network		
vlan-data (multiple)	Data VLAN of customer VMs as needed		

Table 18 VSAN Information

Name	Network Purpose	VSAN ID	FCoE VLAN ID
Storage	storage access		

Table 19 Service Accounts

Account	Purpose	Password (optional, secure appropriately)
	Windows Server administrator	
root	ESXi root	
	Array administrator	
	vCenter administrator	
	SQL Server administrator	

References

Cisco UCS:

http://www.cisco.com/en/US/solutions/ns340/ns517/ns224/ns944/unified_computing.html

Cisco UCSM 2.1 configuration guides:

CLI:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.1/b_UCSM_CLI_Configuration_Guide_2_1.html

GUI:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/gui/config/guide/2.1/b_UCSM_GUI_Configuration_Guide_2_1.html

VMware vSphere:

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

VMware vSphere 5.1 Documentation:

http://pubs.vmware.com/vsphere-51/index.jsp

EMC VNX5xxx series resources:

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

EMC VNXe3xxx series resources:

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

Microsoft SQL Server installation guide:

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