Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 250 Virtual Machines

Last Updated: June 5, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Mehul Bhatt

VijayKumar D

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.

Vijay Kumar D, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems

Vijay Kumar has over 10 years of experience in UCS, network, storage and server virtualization design. Vijay has worked on performance and benchmarking on Cisco UCS and has delivered benchmark results on SPEC CPU2006 and SPECj ENT 2010. Vijay holds certification in VMware Certified Professional and Cisco Unified Computing systems Design specialist.

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
- 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.

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

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at http://www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

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.

© 2013 Cisco Systems, Inc. All rights reserved.

Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 250 Virtual Machines

Executive Summary

Cisco solution for the EMC VSPEX is a pre-validated and modular architecture built with proven best-of-breed technologies to create and complete an end-to-end virtualization solution. The end-to-end solutions enable you to make an informed decision while choosing the hypervisor, compute, storage and networking layers. VSPEX eliminates the server virtualization planning and configuration burdens. The 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 VMware architecture for 250 virtual machines with Cisco solution for the EMC VSPEX.

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 an appropriate platform for virtualization can be challenging. Virtualization platforms should be flexible, reliable, and cost effective to facilitate the deployment of various enterprise applications. In a virtualization platform to utilize compute, network, and storage resources effectively, the ability to slice and dice the underlying platform is essential to size to the application requirements. The Cisco solution for the EMC VSPEX provides a very simplistic yet fully integrated and validated infrastructure to deploy VMs in various sizes to suit various application needs.

Target Audience

The reader of this document is expected to have the necessary training and background to install and configure VMware vSphere 5.1, EMC VNX5500, Cisco Nexus 5548UP switch, and Cisco Unified Computing (UCS) B200 M3 Blade Servers. External references are provided wherever 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.

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

Copyright 2013 Cisco Systems, Inc. All rights reserved.

Purpose of this Guide

This document describes the steps required to deploy and configure the Cisco solution for the EMC VSPEX for VMware architecture. The document provides the end-to-end solution for the VMware vSphere 5.1 for 250 Virtual Machine Architecture.

The readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment models mentioned above.

Business Needs

The 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. The 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. The Cisco solution for the EMC VSPEX using VMware reduces the complexity of configuring every component of a traditional deployment model. 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 for EMC VSPEX VMware architectures:

- Provide an end-to-end virtualization solution to utilize the capabilities of the unified infrastructure components.
- Provide a Cisco VSPEX for VMware ITaaS solution for efficiently virtualizing 250 virtual machines for varied customer use cases.
- Show implementation progression of VMware vCenter 5.1 design and the results.
- Provide a reliable, flexible and scalable reference design.

Solution Overview

The Cisco solution for EMC VSPEX using VMware vSphere 5.1 provides an end-to-end architecture with Cisco, EMC, VMware, and Microsoft technologies that demonstrate support for up to 250 generic virtual machines and provide high availability and server redundancy.

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

- Cisco B-series Unified Computing System servers
- Cisco UCS 5108 Chassis
- Cisco UCS 2204XP Fabric Extenders
- Cisco UCS 6248UP Fabric Interconnects
- Cisco Nexus 5548UP Switches
- Cisco VMFEX virtual Distributed Switch across multiple VMware ESXi hypervisors
- Cisco virtual Port Channels for network load balancing and high availability
- EMC VNX5500 storage array
- VMware vCenter 5
- Microsoft SQL database

- VMware DRS
- VMware HA

The solution is designed to host scalable, and mixed application workloads. The scope of this CVD is limited to the Cisco solution for EMC VSPEX VMware solutions for 250 virtual machines 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

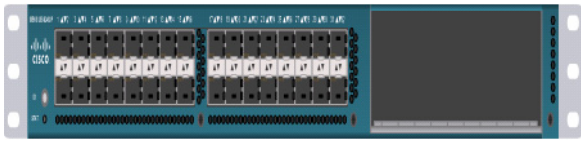

Cisco UCS Fabric Extenders

The Cisco UCS 2200 Series Fabric Extenders multiplex and forward all traffic from blade servers in a chassis to a parent Cisco UCS fabric interconnect over from 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis or virtual machines on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the fabric interconnect. At the core of the Cisco UCS fabric extender are application-specific integrated circuit (ASIC) processors developed by Cisco that multiplex all traffic.

Cisco UCS 2204XP Fabric Extender

The Cisco UCS 2204XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2204XP has sixteen 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 80 Gbps of I/O to the chassis.

Figure 2 Cisco UCS 2204XP Fabric Extender

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 3.

Figure 3 Cisco Blade Server Chassis (front and back view)

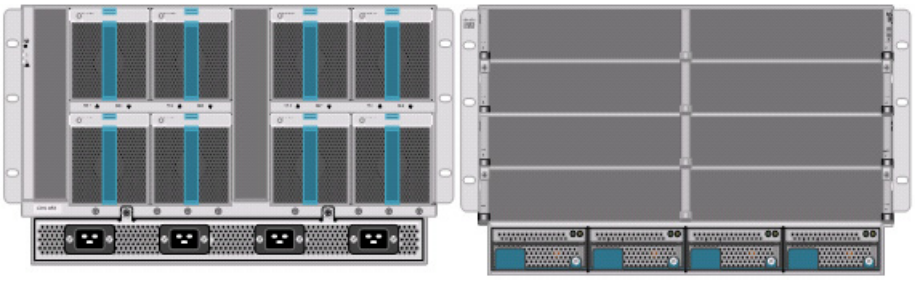

Cisco UCS Blade Servers

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 4.

Figure 4 Cisco UCS B200 M3 Blade Server

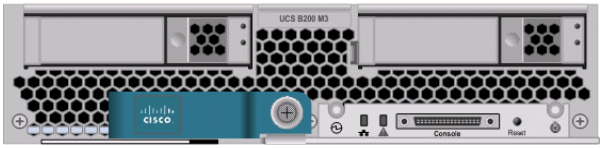

Cisco Nexus 5548UP Switch

The Cisco Nexus 5548UP is a 1RU 1 Gigabit and 10 Gigabit Ethernet switch offering up to 960 gigabits per second throughput and scaling up to 48 ports. It offers 32 1/10 Gigabit Ethernet fixed enhanced Small Form-Factor Pluggable (SFP+) Ethernet/FCoE or 1/2/4/8-Gbps native FC unified ports and three expansion slots. These slots have a combination of Ethernet/FCoE and native FC ports. The Cisco Nexus 5548UP switch is shown in Figure 5.

Figure 5 Cisco Nexus 5548UP switch

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

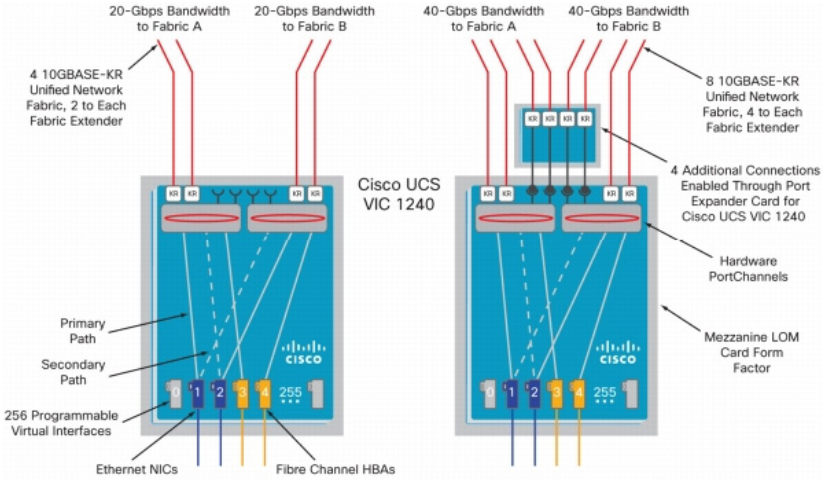

Cisco VM-FEX Technology

The Virtual Interface Card provides hardware based implementation of the Cisco VM-FEX technology. The Cisco VM-FEX technology eliminates the standard virtual switch within the hypervisor by providing individual virtual machine virtual ports on the physical network switch. Virtual machine I/O is sent directly to the upstream physical network switch, in this case, the Cisco UCS 6200 Series Fabric Interconnect, which takes full responsibility for virtual machine switching and policy enforcement.

In a VMware environment, the VIC presents itself as three distinct device types to the hypervisor OS as:

- A fibre channel interface
- A standard Ethernet interface
- A special dynamic Ethernet interface

The Fibre Channel and Ethernet interfaces are consumed by the standard VMware vmkernel components and provide standard capabilities. The dynamic Ethernet interfaces are not visible to the vmkernel layers and are preserved as raw PCIe devices.

Using the Cisco vDS VMware plug-in and Cisco VM-FEX technology, the VIC provides a solution that is capable of discovering the dynamic Ethernet interfaces and registering all of them as uplink interfaces for internal consumption of the vDS. The vDS component on each host discovers the number of uplink interfaces that it has and presents a switch to the virtual machines running on a host as shown in the Figure 7. All traffic from an interface on a virtual machine is sent to the corresponding port of the vDS switch. The traffic is mapped immediately to a unique dynamic Ethernet interface presented by the VIC. This vDS implementation guarantees the 1:1 relationship with a virtual machine interface and an uplink port. The dynamic Ethernet interface selected, is a precise proxy for the virtual machine's interface.

The dynamic Ethernet interface presented by the VIC has a corresponding virtual port on the upstream fabric interconnect.

Figure 7 VM Interfaces Showing their Virtual Ports on the Physical Switch

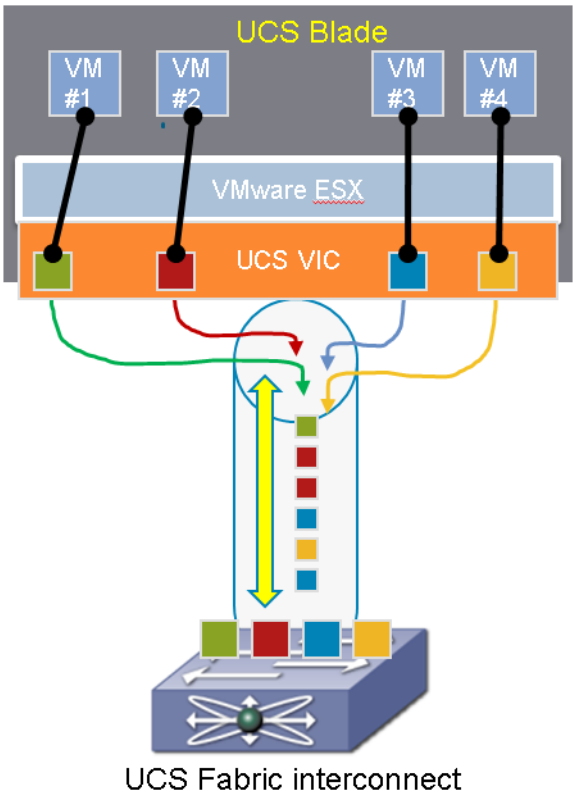

Cisco UCS Manager running on the Cisco UCS Fabric Interconnect works in conjunction with the VMware vCenter software to coordinate the creation and movement of virtual machines. Port profiles are used to describe the virtual machine interface attributes such as VLAN, port security, rate limiting, and QoS marking. Port profiles are managed and configured by network administrators using Cisco UCS Manager. To facilitate integration with the VMware vCenter, Cisco UCS Manager pushes the catalog of port profiles into VMware vCenter, where they are represented as distinct port groups. This integration allows the virtual machine administrators to simply select from a menu of port profiles as they create virtual machines. When a virtual machine is created or moved to a different host, it communicates its port group to the Virtual Interface Card. The VIC gets the port profile corresponding to the requested profile from the Cisco UCS Manager and the virtual port on the fabric interconnect switch is configured according to the attributes defined in the port profile.

The Cisco VM-FEX technology addresses the common concerns of server virtualization and virtual networking by providing the following benefits:

- **Unified virtual and physical networking**—The Cisco VM-FEX technology consolidates the virtual network and physical network into a single switching point that has a single management point. Using the Cisco VM-FEX technology, number of network management points can be reduced drastically.
- **Consistent performance and feature availability**—All the network traffic is controlled at the physical switch, which ensures consistent management of both the virtual and physical network traffic. Each virtual machine interface is coupled with a unique interface on the physical switch, which allows precise decisions to be made related to the scheduling of and operations on traffic flow from and to a virtual machine.

- **Reduced broadcast domains**—The virtual machine's identity and positioning information is known to the physical switch, so the network configuration can be precise and specific to the port in question.

Modes of Operations for VM-FEX technology

Cisco VM-FEX technology supports virtual machine interfaces that run in the following modes:

- Emulated mode

  The hypervisor emulates a NIC (also referred to as a back-end emulated device) to replicate the hardware it virtualizes for the guest virtual machine. The emulated device presents descriptors, for read and write, and interrupts to the guest virtual machine just as a real hardware NIC device would. One such NIC device that VMware ESXi emulates is the vmxnet3 device. The guest OS in turn instantiates a device driver for the emulated NIC. All the resources of the emulated devices' host interface are mapped to the address space of the guest OS.

- PCIe Pass-Through or VMDirectPath mode

  Virtual Interface Card uses PCIe standards-compliant IOMMU technology from Intel and VMware's VMDirectPath technology to implement PCIe Pass-Through across the hypervisor layer and eliminate the associated I/O overhead. The Pass-Through mode can be requested in the port profile associated with the interface using the "high-performance" attribute.

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 5.1 offers improvements in performance and utilization of CPU, memory, and I/O. It also offers

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

The vSphere 5.1 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

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 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.

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 Cisco Validated Design discusses the deployment model of the VMware solution for 250 virtual machines.

Table 1 lists the hardware components required in the solution:

Table 1 Hardware requirements

Components	Hardware required
Servers	Ten Cisco B200 M3 servers
Adapters	Ten Cisco VICs: one Cisco VIC 1240 adapter per server
Chassis	Two Cisco UCS 5108 Blade Server Chassis
Fabric extenders	Four 2204XP Fabric Extender: two fabric extenders per chassis
Fabric interconnects	Two Cisco UCS 6248UP Fabric Interconnects
Network switches	Two Cisco Nexus 5548UP Switches
Storage	One EMC VNX5500 Storage Array

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

Table 2 Firmware and software components of VMware architectures

Vendor	Name	Version	Description
Cisco	Cisco UCS B200 M3 servers	2.1(1) - CIMC B200M3.2.0.4a.0.0809 0121557 – BIOS	Cisco UCS B200 M3 blade server firmware
Cisco	Cisco VIC 1240	2.1(1)	Cisco Virtual Interface Card (adapter) firmware
Cisco	Cisco UCS 2204XP Fabric Extender	2.1(1)	Cisco UCS fabric extender firmware
Cisco	Cisco UCS 6248UP Fabric Interconnect	5.0(3)N2(2.11)	Cisco UCS fabric interconnect firmware
Cisco	Cisco UCSM	2.1(1)	Cisco UCS Manager (UCSM) software
Cisco	Cisco Nexus 5548UP Switches	5.1(3)N1(1a)	Cisco Nexus 5000 series switches running NX-OS
EMC	EMC VNX5500	05.31.000.5.704	EMC VNX storage array firmware
EMC	EMC Avamar	6.0.0-592	EMC data backup software
EMC	Data Domain OS	5.1.0.9-282511	EMC data domain operating system
VMware	ESXi 5.1	5.1 build 799733	VMware Hypervisor
VMware	vCenter Server	5.1 build 799731	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 3 outlines the B200 M3 server configuration details (per server basis) across all the VMware architectures.

Table 3 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

This architecture assumes 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. Figure 8 shows a high level Cisco solution for EMC VSPEX VMware architecture for 250 virtual machines.

Figure 8 Reference Architecture for 250 Virtual Machines

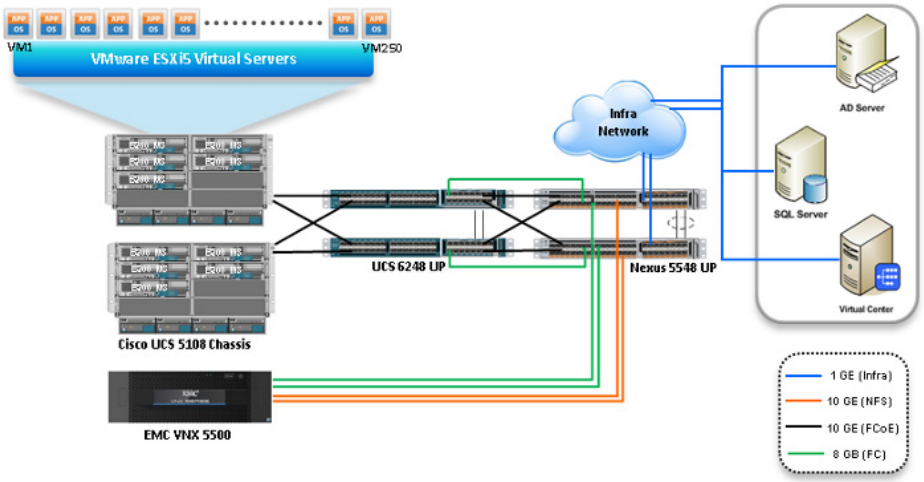

The following are the high level design points of the architecture:

- Only Ethernet is used as network layer 2 media to access Cisco UCS 6248UP from the Cisco UCS B200 M3 blade servers.
- 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 recommend or require any specific layout of infrastructure network. The vCenter server, SQL server, and AD/ DNS virtual machines are hosted on the 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 fibre channel SAN. The defined load is 25 virtual machines per server.

Memory Configuration Guidelines

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

ESX/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 over commitment, 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

Figure 9 shows the use of memory settings parameters in the virtual machine.

Figure 9 Virtual Machine Memory Settings

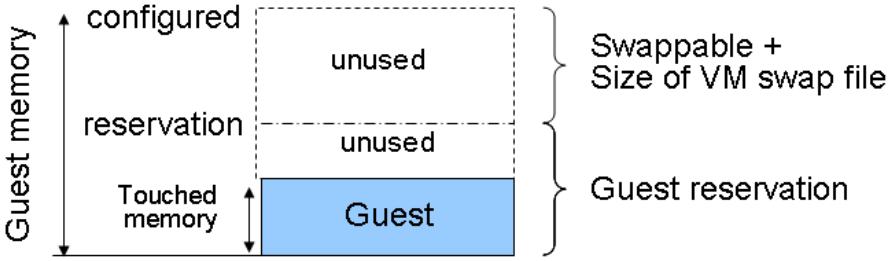

The vSphere memory settings for a virtual machine include the following parameters:

- **Configured memory**—Memory size of virtual machine assigned at creation.
- **Touched memory**—Memory actually used by the virtual machine. vSphere allocates only guest operating system memory on demand.
- **Swappable**—Virtual machine memory 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.

Allocating Memory to Virtual Machines

Memory sizing 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 for a single virtual machine

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

Table 4 Resources for a single virtual machine

Characteristics	Value
I/O pattern	Random
I/O read/write ratio	2:1

Following are the 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 GB of memory may consume about 100 MB of memory overhead, where a virtual machine with two virtual CPUs and 32 GB of memory may consume approximately 500 MB 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 over commitment 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 over committing. 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 ESX/ESXi if physical memory comes under contention. Performance tests show that the balloon driver allows ESX/ESXi to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESX/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 the services have the resources they require while still allowing high consolidation of other virtual machines.

As with over committing CPU resources, proactive monitoring is a requirement. Table 5 lists counters that can be monitored to avoid performance issues resulting from overcommitted memory.

Table 5 ESXitop memory counters

EXitop Metrics	Description	Implication
SWAP /MB: r/s, w/s	The rate at which machine memory is swapped in and out of disk.	High rates of swapping affect guest performance. If free memory is low, consider moving virtual machines to other hosts. If free memory is OK, check resource limits on the virtual machines.
MCTLSZ	The amount of guest physical memory reclaimed by the balloon driver.	If the guest working set is smaller than guest physical memory after ballooning, no performance degradation is observed. However, investigate the cause for ballooning. It could be due to low host memory or a memory limit on the virtual machine.

Storage Guidelines

VSPEX architecture for VMware 250 VMs scale, uses NFS to access storage arrays. 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 VNX VAAI plug-in for NFS storage and storage replication. 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.

Virtual Server Configuration

Figure 10 shows that the VMware storage virtualization can be categorized into three layers of storage technology:

- The Storage array is the bottom layer, consisting of physical disks presented as logical disks (storage array volumes or LUNs) to the layer above, with the vSphere virtual environment.
- Storage array LUNs that are formatted as NFS datastores provide storage for virtual disks.
- Virtual disks that are presented to the virtual machine and guest operating system as NFS attached disks can be partitioned and used in the file systems.

Figure 10 VMware Storage Virtualization Stack

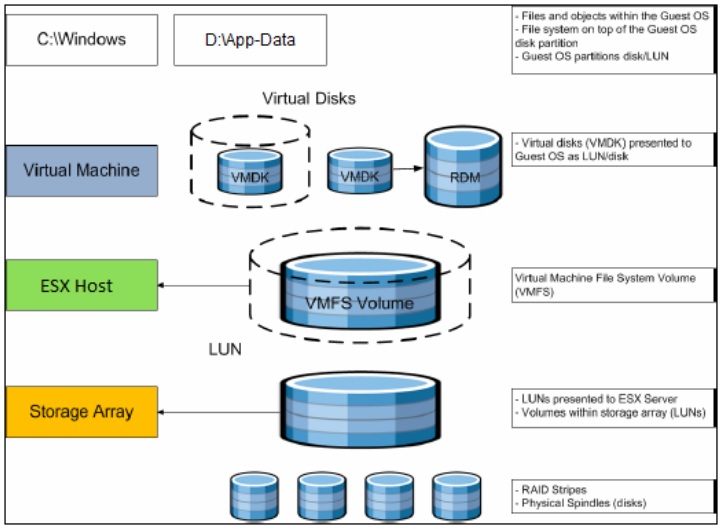

Storage Protocol Capabilities

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.

For VSPEX solution on vSphere NFS is a recommended option because of its simplicity in deployment.

For more information, see the VMware white paper Comparison of Storage Protocol Performance in VMware vSphere 5.1:

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

Storage Best Practices

Following are the 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. In this solution, the redundancy is comes from the “Fabric Failover” feature of the dynamic vNICs of Cisco UCS for NFS storage access.
- 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 NFS level as well as within the guest operating system. Use the vSphere Client when creating NFS 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 multi-tier 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 NFS volume. Over-subscription of the I/O resources can go unnoticed at first and slowly begin to degrade performance if not monitored proactively.

VSPEX VMware Memory Virtualization

VMware vSphere 5.1 has a number of advanced features that help to maximize performance and overall resources utilization. This section describes the performance benefits of some of these features for the VSPEX deployment.

Memory Compression

Memory over-commitment occurs when more memory is allocated to virtual machines than is physically present in a VMware ESXi host. Using sophisticated techniques, such as ballooning and transparent page sharing, ESXi is able to handle memory over-commitment without any performance degradation. However, if more memory than that is present on the server is being actively used, ESXi might resort to swapping out portions of a VM's memory.

For more details about Vsphere memory management concepts, see the VMware Vsphere Resource Management Guide at:

http://www.VMware.com/files/pdf/mem_mgmt_perf_Vsphere5.pdf

Virtual Networking

The Cisco VMFEX collapses virtual and physical networking into a single infrastructure. The Cisco VM-FEX technology allows data center administrators to provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare metal network traffic within a unified UCS infrastructure.

The VM-FEX technology extends Cisco data-center networking technology to the virtual machine with the following capabilities:

- Each virtual machine includes a dedicated interface on the virtual Distributed Switch (vDS).
- All virtual machine traffic is sent directly to the dedicated interface on the vDS.
- The native VMware virtual switch in the hypervisor is replaced by the vDS.
- Live migration and vMotion are also supported with the Cisco VM-FEX.

Benefits

- Simplified operations—Seamless virtual networking infrastructure through UCS Manager
- Improved network security—Contains VLAN proliferation
- Optimized network utilization—Reduces broadcast domains

- Reduced network complexity—Separation of network and server administrator's domain by providing port-profiles by name

Virtual Networking Best Practices

Following are the vSphere networking best practices:

- Separate virtual machine and infrastructure traffic—Keep virtual machine and VMkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation.
- 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. To prevent this situation, set the switch ports connected to ESX/ESXi hosts to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESX/ESXi hosts from affecting the STP topology. Loops are not possible in virtual switches.
- Converged Network and Storage I/O with 10Gbps Ethernet—Consolidating storage and network traffic can provide simplified cabling and management over maintaining separate switching infrastructures.
- Fabric Failover—Always use fabric failover feature of Cisco UCS VIC adapters for high-availability of network access.

This solution suggests 32 dynamic VNICs per ESXi host based on following assumptions and calculation:

- One vNIC per virtual machine
- With 25 VMs per hypervisor, 25 dynamic vNICs is needed
- Three vm-kernel interfaces per hypervisor:
  - One for management
  - One for vMotion
  - One for NFS storage access
- Four additional dynamic vNICs for high-availability. High availability is required when:
  - One of the hypervisor is shutdown or in maintenance mode.
  - The VMs on the hypervisor is moved to other hypervisors.

  Three dynamic vNICs are required per hypervisor, but we have provisioned one extra for head room.

vSphere VMware Performance

With every release of vSphere the overhead of running an application on the vSphere virtualized platform is reduced by the new performance improving features. Typical virtualization overhead for applications is less than 10%. Many of these features not only improve performance of the virtualized application itself, but also allow for higher consolidation ratios. Understanding these features and taking advantage of them in your environment helps guarantee the highest level of success in your virtualized deployment. Table 6 provides details on vSphere VMware performance.

Table 6 vSPhere VMware performance

ESXitop Metric	Description	Implication
NUMA Support	ESX/ESXi uses a NUMA load-balancer to assign a home node to a virtual machine. Because memory for the virtual machine is allocated from the home node, memory access is local and provides the best performance possible. Even applications that do not directly support NUMA benefit from this feature.	See The CPU Scheduler in VMware ESXi 5.1: http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf
Transparent page sharing	Virtual machines running similar operating systems and applications typically have identical sets of memory content. Page sharing allows the hypervisor to reclaim the redundant copies and keep only one copy, which frees up the total host memory consumption. If most of your application virtual machines run the same operating system and application binaries then total memory usage can be reduced to increase consolidation ratios.	See Understanding Memory Resource Management in VMware ESXi 5.1: http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf
Memory ballooning	By using a balloon driver loaded in the guest operating system, the hypervisor can reclaim host physical memory if memory resources are under contention. This is done with little to no impact to the performance of the application.	See Understanding Memory Resource Management in VMware ESXi 5.1: http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf

Table 6 vSPhere VMware performance

ESXitop Metric	Description	Implication
Memory compression	Before a virtual machine resorts to host swapping, due to memory over commitment the pages elected to be swapped attempt to be compressed. If the pages can be compressed and stored in a compression cache, located in main memory, the next access to the page causes a page decompression as opposed to a disk swap out operation, which can be an order of magnitude faster.	See Understanding Memory Resource Management in VMware ESXi 5.1: http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf
Large memory page support	An application that can benefit from large pages on native systems, such as MS SQL, can potentially achieve a similar performance improvement on a virtual machine backed with large memory pages. Enabling large pages increases the memory page size from 4KB to 2MB.	See Performance Best Practices for VMware vSphere 5.1: http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf and see Performance and Scalability of Microsoft SQL Server on VMware vSphere 4: http://www.vmware.com/files/pdf/perf_vsphere_sql_scalability.pdf

Physical and Virtual CPUs

VMware uses the terms virtual CPU (vCPU) and physical CPU to distinguish between the processors within the virtual machine and the underlying physical x86/x64-based processor cores. Virtual machines with more than one virtual CPU are also called SMP (symmetric multiprocessing) virtual machines. The virtual machine monitor (VMM), or hypervisor, is responsible for CPU virtualization. When a virtual machine starts running, control transfers to the VMM, which virtualizes the guest OS instructions.

Virtual SMP

VMware Virtual Symmetric Multiprocessing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processor cores simultaneously. vSphere supports the use of up to thirty two virtual CPUs per virtual machine. The biggest advantage of an SMP system is the ability to use multiple processors to execute multiple tasks concurrently, thereby increasing throughput (for example, the number of transactions per second). Only workloads that support parallelization (including multiple processes or multiple threads that can run in parallel) can really benefit from SMP.

The virtual processors from SMP-enabled virtual machines are co-scheduled. That is, if physical processor cores are available, the virtual processors are mapped one-to-one onto physical processors and are then run simultaneously. In other words, if one vCPU in the virtual machine is running, a second vCPU is co-scheduled so that they execute nearly synchronously. Consider the following points when using multiple vCPUs:

- Simplistically, if multiple, idle physical CPUs are not available when the virtual machine wants to run, the virtual machine remains in a special wait state. The time a virtual machine spends in this wait state is called ready time.
- Even idle processors perform a limited amount of work in an operating system. In addition to this minimal amount, the ESXi host manages these "idle" processors, resulting in some additional work by the hypervisor. These low-utilization vCPUs compete with other vCPUs for system resources.

In VMware ESXi 5 and ESXi, the CPU scheduler underwent several improvements to provide better performance and scalability; for more information, see the *CPU Scheduler in VMware ESXi 5.1*:

http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.1.pdf. For example, in VMware ESXi 5, the relaxed co-scheduling algorithm was refined so that scheduling constraints due to co-scheduling requirements are further reduced. These improvements resulted in better linear scalability and performance of the SMP virtual machines.

Overcommitment

VMware conducted tests on virtual CPU overcommitment with SAP and SQL, showing that the performance degradation inside the virtual machines is linearly reciprocal to the overcommitment. Because the performance degradation is "graceful," any virtual CPU overcommitment can be effectively managed by using VMware DRS and VMware vSphere® vMotion® to move virtual machines to other ESX/ESXi hosts to obtain more processing power. By intelligently implementing CPU overcommitment, consolidation ratios of applications Web front-end and application servers can be driven higher while maintaining acceptable performance. If it is chosen that a virtual machine not participate in overcommitment, setting a CPU reservation provides a guaranteed CPU allocation for the virtual machine. This practice is generally not recommended because the reserved resources are not available to other virtual machines and flexibility is often required to manage changing workloads. However, SLAs and multi-tenancy may require a guaranteed amount of compute resources to be available. In these cases, reservations make sure that these requirements are met.

When choosing to overcommit CPU resources, monitor vSphere and applications to be sure responsiveness is maintained at an acceptable level. Table 7 lists counters that can be monitored to help achieve higher drive consolidation while maintaining the system performance.

Table 7 List of Counters

ESXitop Metric	Description	Implication
%RDY	Percentage of time a vCPU in a run queue is waiting for the CPU scheduler to let it run on a physical CPU.	A high %RDY time (use 20% as a starting point) may indicate the virtual machine is under resource contention. Monitor this—if application speed is OK, a higher threshold may be tolerated.
%MLMTD	Percentage of time a vCPU was ready to run but was deliberately not scheduled due to CPU limits.	A high %MLMTD time may indicate a CPU limit is holding the VM in a ready to run state. If the application is running slow consider increasing or removing the CPU limit.
%CSTP	Percentage of time a vCPU spent in read, co-descheduled state. Only meaningful for SMP virtual machines.	A high %CSTP time usually means that vCPUs are not being used in a balanced fashion. Evaluate the necessity for multiple vCPUs.

Hyper-threading

Hyper-threading technology (recent versions of which are called symmetric multithreading, or SMT) enables a single physical processor core to behave like two logical processors, essentially allowing two independent threads to run simultaneously. Unlike having twice as many processor cores which can roughly double performance, hyper-threading can provide anywhere from a slight to a significant increase in system performance by keeping the processor pipeline busier.

Non-Uniform Memory Access (NUMA)

Non-Uniform Memory Access (NUMA) compatible systems contain multiple nodes that consist of a set of processors and memory. The access to memory in the same node is local, while access to the other node is remote. Remote access can take longer because it involves a multihop operation. In NUMA-aware applications, there is an attempt to keep threads local to improve performance.

The VMware ESX/ESXi provides load-balancing on NUMA systems. To achieve the best performance, it is recommended that the NUMA be enabled on compatible systems. On a NUMA-enabled ESX/ESXi host, virtual machines are assigned a home node from which the virtual machine's memory is allocated. Because it is rare for a virtual machine to migrate away from the home node, memory access is mostly kept local.

In applications that scale out well it is beneficial to size the virtual machines with the NUMA node size in mind. For example, in a system with two hexa-core processors and 64GB of memory, sizing the virtual machine to six virtual CPUs and 32GB or less, means that the virtual machine does not have to span multiple nodes.

VSPEX VMware Storage Virtualization

Disk provisioning on the EMC VNX series requires administrators to choose disks for each of the storage pools.

Storage Layout

The architecture diagram in this section shows the physical disk layout. Disk provisioning on the VNX5500 storage array 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 shows storage architecture for 250 virtual machines on VNX5500.

Figure 11 Storage Architecture for 250 VMs on EMC VNX5500

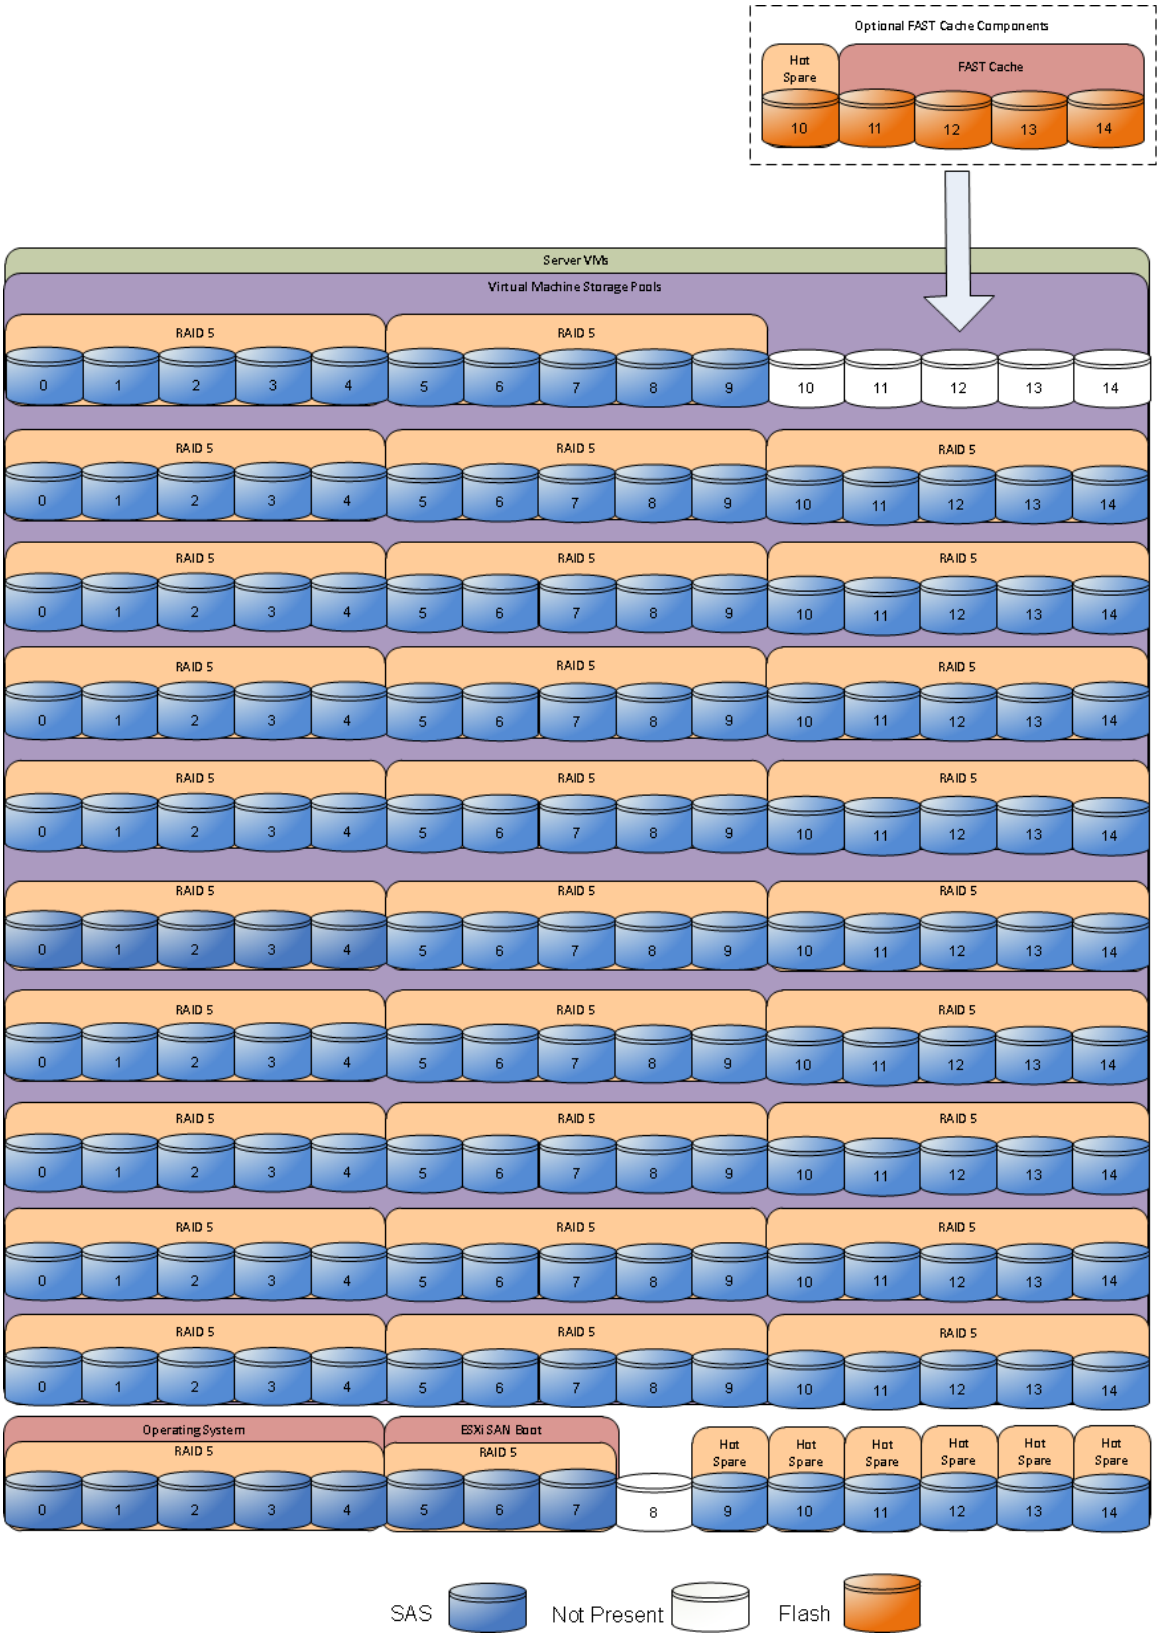

The reference architecture uses the following configuration:

- One hundred fourty-five 300 GB SAS disks are allocated to a block-based storage pool.

Note Note: System drives are specifically excluded from the pool, and not used for additional storage.

- Six 300GB SAS disks are configured as hot spares.
- Three 300GB SAS disks are configured for ESXi 5.1 hypervisor SAN Boot.
- Optionally you can configure up to 20 flash drives in the array FAST Cache. These drives are not considered to be a required part of the solution, and additional licensing may be required in order to use the FAST Suite.
- EMC recommends that at least one hot spare disk is allocated for each 30 disks of a given type.
- At least two NFS shares are allocated to the vSphere cluster from each storage pool to serve as datastores for the virtual servers.

The VNX family storage array 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

NFS is a cluster file system that provides UDP based stateless storage protocol to access storage across multiple hosts over the network. Each virtual machine is encapsulated in a small set of files and NFS datastore mount points are used for the operating system partitioning and data partitioning.

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.

Architecture for 250 VMware Virtual Machines

Figure 12 shows the logical layout of 50 VMware virtual machines. The following are the key aspects of this solution:

- Ten Cisco B200 M3 servers are used, with an average load of 25 VMs per server.
- ESXi 5 is booted from SAN disk. FCoE is used from servers up to fabric interconnect, and then native FC from fabric interconnect to storage array
- Virtual port-channels on storage side networking provide high-availability and load balancing.
- Cisco VMFEX distributed Virtual Switch provides port-profiles based virtual networking solution.
- Fabric failover capability of VMFEX dynamic vNICs provides network high availability.
- SAN boot and Cisco UCS Manager service profile provides complete stateless computing architecture. A B200 M3 server can be replaced with a very little down time window.

Figure 12 Cisco Solution VMware Architecture for 250 VMs

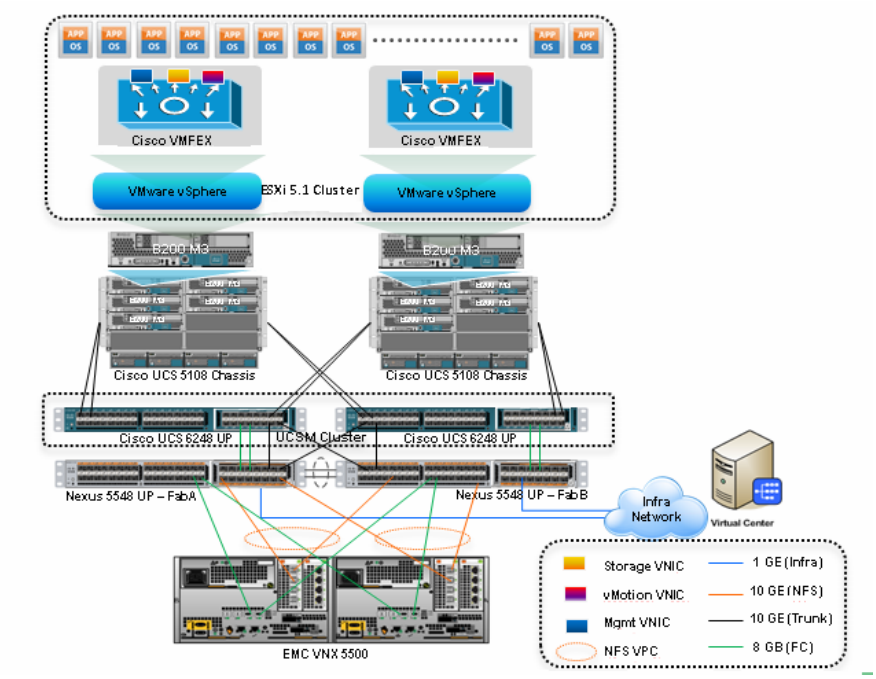

Stateless Computing

Cisco UCS Manager (UCSM) provides the concept of Service Profile for server running on a physical hardware. Service profile is a logical entity, which can be associated to the physical server. Among other things, service profile includes various identities of the server or server components, such as:

- BIOS UUID
- MAC address of virtual NIC of the server
- Node WWN (WWNN) for Fibre Channel SAN access
- Port WWN (WWPN) of the virtual HBA of the server
- IQN ID, if iSCSI protocol is used for storage access
- Management IP address for the KVM access

All these identities can be assigned to any physical server managed by the Cisco UCS Manager. All other configuration of the service profile is based on templates, pools and policies, providing immense flexibility to the administrator. This includes firmware and BIOS versions required by the server. These concepts enable Cisco UCS Manager to provide stateless computing across entire Cisco UCS Manager managed compute hardware. If remote storage is used to boot operating system of the server (such as SAN boot, PXE boot, iSCSI boot etc), then a given service profile can be associated to any physical server hardware and downtime for migrating such server can be reduced to few minutes. Solution presented in this CVD makes use of identity pools and SAN storage to simplify the server procurement and provide stateless computing capability.

Sizing Guidelines

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 8 Virtual Machine Characteristics

Characteristic	Value
Virtual machine operating system	Microsoft Windows Server 2008 R2 SP1
Virtual processor per virtual machine (vCPU)	1
RAM per virtual machine	2 GB
Available storage capacity per virtual machine	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 which 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, following resources are needed from the resource pool:
- CPU resources for 1 VM
- Memory resources for 2 VMs
- Storage capacity for 1 VM
- IOPS for 1 VM
In this example, a single virtual machine uses the resources of two of the reference VMs. Once this VM is deployed, the solution's new capability would be 248 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 resources that are needed from the resource pool to virtualize this application:
- CPUs of 4 reference VMs
- Memory of 8 reference VMs
- Storage of 2 reference VMs
- IOPS of 8 reference VMs
In this case the one virtual machine uses the resources of eight reference virtual machines. Once this VM is deployed, the solution's new capability would be 242 VMs.

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 resources that are needed from the resource pool to virtualize this application:
- CPUs of 2 reference VMs
- Memory of 4 reference VMs
- Storage of 1 reference VMs
- IOPS of 2 reference VMs
In this case the virtual machine would use the resources of four reference virtual machines. Once this VM is deployed, the solution's new capability would be 246 VMs.

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 10 CPUs and 48 GB of memory. It uses 5 TB of storage and generates 700 IOPS during an average busy cycle.
The following resources that are needed from the resource pool to virtualize this application:
- CPUs of ten reference VMs
- Memory of 24 reference VMs
- Storage of 52 reference VMs
- IOPS of 28 reference VMs
In this case the one virtual machine uses the resources of fifty-two reference virtual machines. Once this VM is deployed, the solution's new capability would be 198 VMs.

Summary of Example

The four examples show the flexibility of the resource pool model. In all the 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 250 virtual machines they can all be implemented, leaving the capacity of one hundred eighty six 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 out of the scope of this document. However, when a change in the resource balance is observed, and the new level of requirements is known; these virtual machines can be added to the infrastructure using the method described in the above examples.

VSPEX Configuration Guidelines

This sections provides the procedure to deploy the Cisco solution for EMC VSPEX VMware architecture.

Follow these steps to configure the Cisco solution for EMC VSPEX VMware architectures:

1. Pre-deployment tasks.
2. Physical setup.
3. Cable connectivity.
4. Configure Cisco Nexus switches.
5. Configure Cisco Unified Computing System using Cisco UCS Manager.
6. Prepare and configure storage array.
7. Install VMware ESXi servers and vCenter infrastructure.
8. Install and configure Microsoft SQL server database.
9. Install and configure VMware vCenter server.
10. Install and configure Cisco Nexus VM-FEX.
11. Test the installation.

These steps are described in detail in the following sections.

Pre-deployment Tasks

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

- Gather documents—Gather the related documents listed in the Preface. These are used throughout the text of this document to provide detail on setup procedures and deployment best practices for the various components of the solution.
- Gather tools—Gather the required and optional tools for the deployment. Use Table 9 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 Sheet, page 172 for reference during the deployment process.

Table 9 Customer Specific Configuration Data

Requirement	Description	Reference
Hardware	Cisco UCS B200 M3 servers to host virtual machines	EMC-Cisco Reference Architecture: *VSPEX Server Virtualization with VMware vSphere 5.1 for 250 Virtual Machines*.
	Cisco UCS 5108 Blade Server Chassis	
	Cisco UCS 2204XP Fabric Extender	
	Cisco UCS 6248UP Fabric Interconnect	
	VMware vSphere™ 5 server to host virtual infrastructure servers **Note** This requirement may be covered in the existing infrastructure	
	Cisco Nexus switches: Two Cisco Nexus 5548UP Switches for high availability	
	EMC VNX storage: EMC VNX5500 Multiprotocol storage array with the required disk layout as per architecture requirements	
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	

Table 9 Customer Specific Configuration Data

Requirement	Description	Reference
Licenses	VMware vCenter 5.1 license key	Consult your corresponding vendor obtain license keys
	VMware ESXi 5.1 license keys	
	Microsoft SQL Server license key **Note** This requirement may be covered in the existing infrastructure	

Customer Configuration Data

To reduce the onsite time, information such as IP addresses and hostnames should be assembled as part of the planning process.

The section Customer Configuration Data Sheet, page 172 provides tabulated record of relevant information (to be filled at the customer's end). This form can be expanded or contracted as required, and information may be added, modified, and recorded as the deployment progresses.

Additionally, complete the VNX Series Configuration Worksheet, available on the EMC online support website, to provide the most comprehensive array-specific information.

Physical setup

Physical setup includes the following tasks:

1. Mount all the hardware.
2. Connect power cords and management connectivity to all hardware.
3. Perform the initial setup steps for all hardware involved.

Preparing Cisco UCS Components

For information on mounting the hardware, see the *Cisco UCS B-Series Hardware Installation Guide.* Care must be taken about efficient cooling and proper airflow while mounting any equipment in the data center. Similarly, you need to pay attention to power requirements of chassis, servers and fabric interconnects.

Cisco UCS 5108 chassis, including its embedded blade servers and fabric extenders do not require management connectivity as they are managed by the fabric interconnects. Fabric interconnects are deployed in pair for high availability. Both the fabric interconnects require 100 Mbps peer connectivity for synchronizing the management plane between them. In addition, both the FIs require 1Gbps out-of-band management connectivity.

Cisco UCS Manager software runs on the Cisco UCS Fabric Interconnects. The Cisco UCS 6000 Series Fabric Interconnects expand the UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blades and Cisco UCS Blade Server Chassis. All chassis and the blade servers attached to the interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6000 Series provides the flexibility to support LAN and SAN

connectivity for all blade servers within its domain right at the configuration time. Typically deployed in redundant pairs, the Cisco UCS Fabric Interconnect provides uniform access to both network and storage, facilitating a fully virtualized environment.

Initial setup steps of Cisco UCS 6248UP Fabric Interconnects and the Cisco UCS Manager are similar to those of the Nexus 5548UP switches:

1. Connect the RJ-45 connector of the console cable to the primary fabric interconnect console port.
2. Configure the terminal emulator program on the host to match the following default port characteristics: 9600 baud, 8 data bits, 1 stop bit, and no parity.
3. Choose the CLI based initial setup configuration and provide basic information about the fabric interconnect cluster.
4. Connect two fabric interconnects using two 100 Mbps Ethernet cables to create management plane cluster.
5. Repeat steps 1, 2 and 3 for the second fabric interconnect. The initial setup for the second fabric interconnect is relatively easier, as it forms a UCS management plane cluster with the pre-configured fabric interconnect, and assumes the role of secondary fabric interconnect.

Cisco UCS 5108 Chassis, Cisco UCS 2204XP Fabric Extenders and UCS B200 M3 blade servers would be part of the Cisco UCS Manager (UCSM) management domain, so no special configuration is required for them.

Preparing Cisco Nexus Switches

Cisco Nexus 5548UP switches are 1RU top of the rack 10Gbps Ethernet and Fibre Channel switches. For information on how to deploy these switches, see *Nexus 5548UP Product Documentation*.

For initial configuration of these switches, follow these steps:

1. Connect the RJ-45 connector of the console cable to the Cisco Nexus 5548UP Switch console port.
2. Configure the terminal emulator program on the host to match the following default port characteristics: 9600 baud, 8 data bits, 1 stop bit, and no parity.
3. Type Setup at the switch prompt and follow the menu driven to configure the IP address on the management port and allow ssh to enable remote configuration of the switch.
4. Using the RJ-45 cable, connect to the upstream switch/router (or to the infrastructure network switch for managing remotely).

Preparing EMC VNX5500

For information on mounting the storage array EMC VNX5500 and initial configuration, see the EMC product documentation. Proper connectivity of storage controllers and DAEs are crucial for high availability of the storage.

Topology Diagram for 250 Virtual Machines

Following diagrams show connectivity details cable connectivity of solution covered in this document. At high level, cable connectivity can be divided in two parts:

- 10 Gbps Ethernet cables connectivity
- 8 Gbps Fibre Channel cables connectivity

As it is apparent from the following figure, there are five major cabling sections for the Ethernet connectivity:

- Chassis / fabric interconnect connectivity
- Fabric interconnect / Nexus 5548UP connectivity
- Inter-switch links
- Storage connectivity
- Infrastructure connectivity

Figure 13 Topology Diagram for 250 VMs

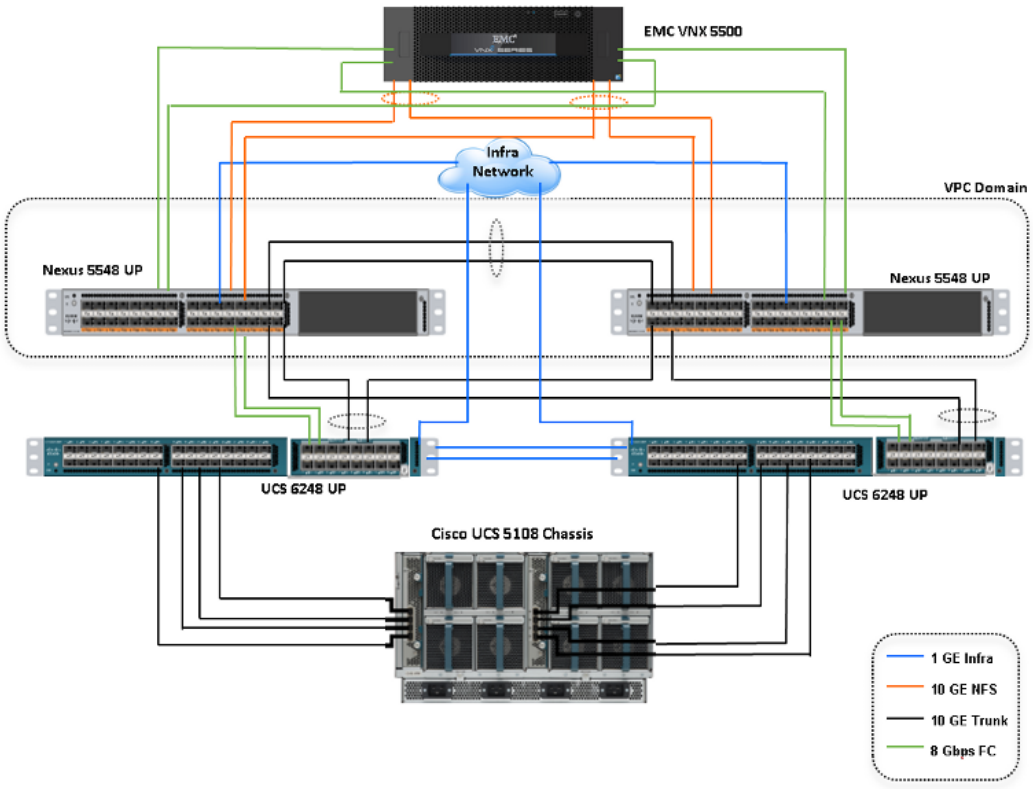

Table 10, Table 11 and Figure 14 provides the detailed cable connectivity for the EMC VSPEX 250 virtual machines configuration. Table 10 lists all the device port links from the Cisco Nexus 5548UP Switch perspective. Table 11 lists all the device port links from the Cisco UCS 6248UP Fabric Interconnects.

Table 10 Cabling Details For 250 Vms From Cisco Nexus 5548up Switch Perspective

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
A,C	Eth1/7	All	Trunk	10(D)	1	VPC peer link
B,D	Eth1/8	All	Trunk	10(D)	1	VPC peer link
E,G	Eth1/18	All	Trunk	10(D)	18	Fabric Interconnect (A)
F,H	Eth1/19	All	Trunk	10(D)	19	Fabric Interconnect (B)

Table 10 Cabling Details For 250 Vms From Cisco Nexus 5548up Switch Perspective

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
I,J	Eth1/24	40	Access	10(D)	24	VNX5500 - SP A
K,L	Eth1/25	40	Access	10(D)	25	VNX5500 - SP B
(not shown)	Eth1/15	1	Trunk	10(D)	-	Uplink to infrastructure network
(not shown)	Eth1/17	1	Trunk	10(D)	-	Uplink to infrastructure network

Figure 14 Port Connectivity from the Cisco UCS Fabric Interconnect Perspective

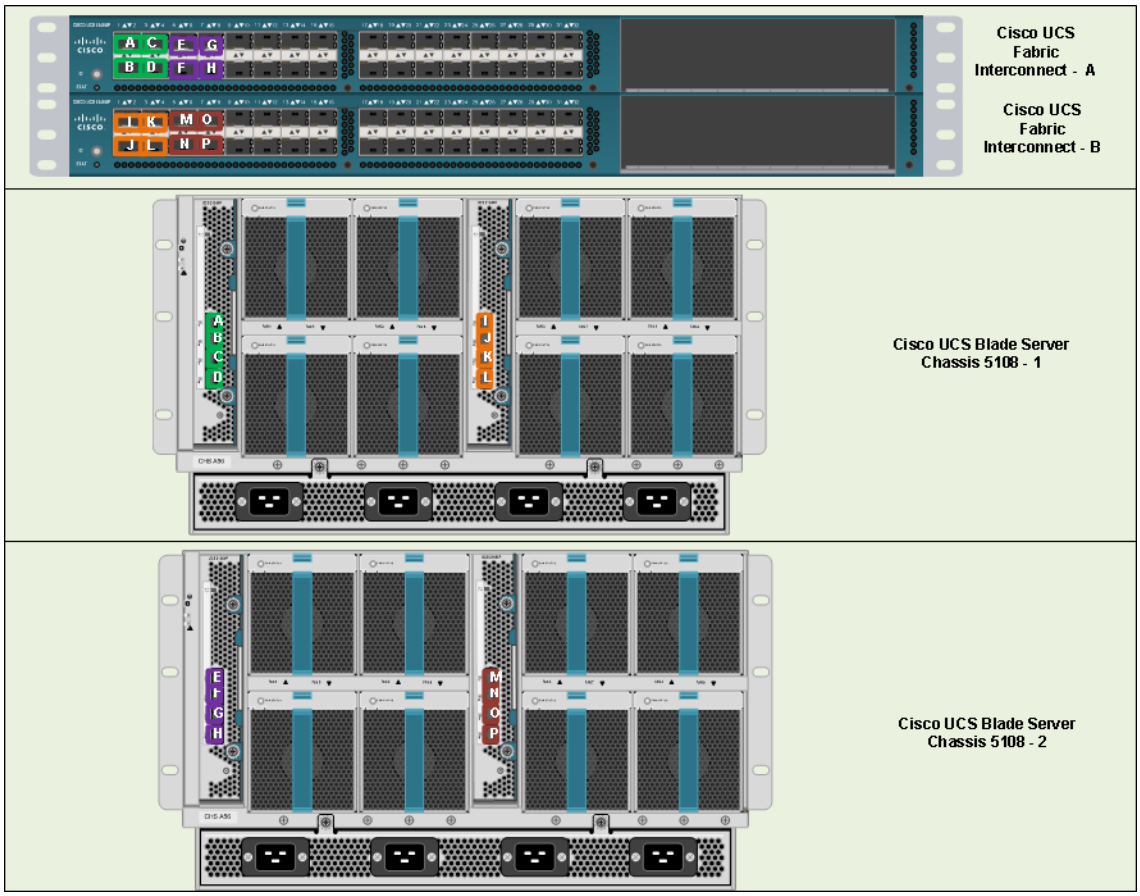

Table 11 Cabling details for 250 VMs from Cisco Fabric Interconnect 6248UP perspective

Cable ID	Fabric Interconnect Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
A,I	Eth1/1	-	Server	10(D)	-	Chassis 1
B,J	Eth1/2	-	Server	10(D)	-	Chassis 1
C,K	Eth1/3	-	Server	1(D)	-	Chassis 1
D,L	Eth1/4	-	Server	1(D)	-	Chassis 1
E,M	Eth1/5	-	Server	1(D)	-	Chassis 2
F,N	Eth1/6	-	Server	10(D)	-	Chassis 2
G,O	Eth1/7	-	Server	10(D)	-	Chassis 2
H,P	Eth1/8	-	Server	10(D)	-	Chassis 2
Q,S	Eth1/18	All	Uplink	10(D)	1	Nexus 5548UP (A)
R,T	Eth1/19	All	Uplink	10(D)	2	Nexus 5548UP (B)
(not shown)	Mgmt0	1	Access	10(D)	-	Uplink to Infrastructure network

After connecting all the cables as per Table 10 and Table 11, you can configure the switch and the fabric interconnect.

Following are the important points to note:

- There are four 10GE links between UCS 2204XP Fabric Extender and 6248UP Fabric Interconnect.
- A given fabric extender connects to only one fabric interconnect. For example, all links from left fabric extender connect to FI-A and all links from right fabric extender connect to FI-B.
- There are no direct 10GE links between two FIs.
- Each FI connect to both Nexus 5548UP switches. Nexus 5548UP switches have peer 10 GE links, and both switches connect to both storage controllers.

Fibre Channel connectivity

This solution uses Fibre Channel over Ethernet (FCoE) protocol from UCS B200 M3 servers to UCS fabric interconnects. This reduces number of cables required between fabric interconnect and UCS blade server chassis by half. Native fibre channel cables are required from FIs to Nexus 5548UP switches and from there to storage devices. Use following guideline to connect the fibre channel links:

- The Cisco UCS 6248UP Fabric Interconnects A and B run in fibre channel NPV mode, and so, Cisco UCS FI-A is connected to Cisco Nexus 5548UP A only. Similarly, Cisco UCS FI-B is connected to Cisco Nexus 5548UP B switch only.
- Both the Cisco Nexus 5548UP switches is connected to the EMC VNX Storage Controllers A and B for redundancy.

Connect all the cables as shown in Figure 14 you will be ready to configure Cisco UCS Manager and switches.

Configuring Cisco Nexus Switches

This section explains switch configuration needed for the Cisco solution for EMC VSPEX VMware architectures. For information on configuring password, and management connectivity, see the Cisco Nexus 5000 Series Configuration Guide.

Configure Global VLANs and VSANs

Figure 15 shows how to configure VLAN on a switch.

Figure 15 Creating VLAN

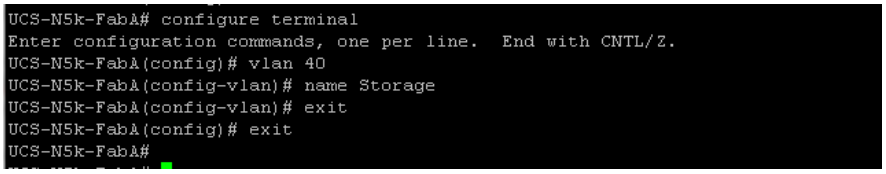

Following VLANs in Table 12 need to be configured on both switches A and B in addition to your application specific VLANs:

Table 12 Configured VLANS on Switch A and B

VLAN Name	Description
Storage	VLAN to access storage array from the servers over NFS
vMotion	VLAN for virtual machine vMotion
Infra	Management VLAN for vSphere servers to reach vCenter management plane
VM-Data	VLAN for the virtual machine (application) traffic (can be multiple VLANs)

For actual VLAN IDs of your deployment, see Customer Configuration Data Sheet, page 172.

We have used one VSAN in this solution. Table 13 gives the VSAN name and the description.

Table 13 Configured Vsan To Access Storage Array

VSAN Name	Description
Storage	VSAN to access storage array from the servers over fibre channel

For actual VSAN ID of your deployment, see Customer Configuration Data Sheet, page 172.

Figure 16 and Figure 17 show the creation of VSAN and assigning VSAN to the fibre channel interface.

Figure 16 Creating VSAN

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# vsan database
UCS-N5k-FabA(config-vsan-db)# vsan 10
UCS-N5k-FabA(config-vsan-db)# vsan 10 interface fc 1/29
UCS-N5k-FabA(config-vsan-db)# vsan 10 interface fc 1/30
UCS-N5k-FabA(config-vsan-db)# vsan 10 interface fc 1/31
UCS-N5k-FabA(config-vsan-db)# vsan 10 interface fc 1/32
UCS-N5k-FabA(config-vsan-db)# end
UCS-N5k-FabA#
```

After creating the VSAN. VSAN membership is assigned, and the peer interfaces on the links need to be configured properly, a healthy fibre channel port is shown in Figure 17.

Figure 17 Assigned VSAN Membership

```
UCS-N5k-FabA# show vsan membership
vsan 1 interfaces:
    fc1/27           fc1/28

vsan 10 interfaces:
    fc1/29           fc1/30           fc1/31           fc1/32

vsan 4079(evfp_isolated_vsan) interfaces:


vsan 4094(isolated_vsan) interfaces:


UCS-N5k-FabA# show interface fc1/29-32 brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                                (Gbps)
--------------------------------------------------------------------------------
fc1/29     10     F      on      up              swl    F       8    --
fc1/30     10     F      on      up              swl    F       8    --
fc1/31     10     F      on      up              swl    F       8    --
fc1/32     10     F      on      up              swl    F       8    --
UCS-N5k-FabA#
```

It is also crucial to enable NPIV feature on the Cisco Nexus 5548UP switches. Figure 18 show how to enable NPIV feature on Nexus 5548UP switches.

Figure 18 Enabling Npiv Feature On Cisco Nexus Switches

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# feature npiv
UCS-N5k-FabA(config)#
```

Configuring Virtual Port Channel (VPC)

Virtual port-channel effectively enables two physical switches to behave like a single virtual switch, and port-channel can be formed across the two physical switches. Following are the steps to enable vPC:

1. Enable LACP feature on both switches.
2. Enable vPC feature on both switches.
3. Configure a unique vPC domain ID, identical on both switches.
4. Configure mutual management IP addresses on both the switches and configure peer-gateway as shown in Figure 19.

Figure 19 Configuring Peer-Gateway

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# feature lacp
UCS-N5k-FabA(config)# feature vpc
UCS-N5k-FabA(config)# vpc domain 101
UCS-N5k-FabA(config-vpc-domain)# peer-keepalive destination 10.29.180.4
Note:
 --------:: Management VRF will be used as the default VRF ::--------
UCS-N5k-FabA(config-vpc-domain)# peer-gateway
UCS-N5k-FabA(config-vpc-domain)# exit
UCS-N5k-FabA(config)# exit
UCS-N5k-FabA#
```

5. Configure port-channel on the inter-switch links. Configuration for these ports is shown in Figure 20. Ensure that "vpc peer-link" is configured on this port-channel.

Figure 20 Configured VPC Peer-link on Port-Channel

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# interface port-channel 1
UCS-N5k-FabA(config-if)# switchport mode trunk
UCS-N5k-FabA(config-if)# spanning-tree port type network
UCS-N5k-FabA(config-if)# speed 10000
UCS-N5k-FabA(config-if)# vpc peer-link
Please note that spanning tree port type is changed to "network" port type on vP
C peer-link.
This will enable spanning tree Bridge Assurance on vPC peer-link provided the ST
P Bridge Assurance
(which is enabled by default) is not disabled.
UCS-N5k-FabA(config-if)# description VPC-Peerlink
UCS-N5k-FabA(config-if)# end
UCS-N5k-FabA#
```

6. Add ports with LACP protocol on the port-channel using "channel-group 1 mode active" command under the interface sub-command.
7. Verify vPC status using **show vPC** command. Successful vPC configuration is shown in Figure 21.

Figure 21 Window Showing Successful vPC Configuration

```
UCS-N5k-FabA# show vpc
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 0
Peer Gateway                      : Enabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po1    up     1,40-41,45
UCS-N5k-FabA#
```

Configuring Port-Channels Connected to Cisco UCS Fabric Interconnects

Interfaces connected to the fabric interconnects need to be in the trunk mode. Storage, vMotion, infra, and application VLANs are allowed on this port. From the switch side, interfaces connected to Cisco UCS FI-A and Cisco UCS FI-B are in a vPC, and from the FI side the links connected to Cisco Nexus 5548UP A and B switches are in LACP port-channels. Ensure that you give a right description for each port and port-channel on the switch for better diagnosis in case of any problem. Figure 22 shows the configuration commands.

Figure 22 Port-channel Configuration

```
UCS-N5k-FabA# show running-config interface port-channel 18-19

!Command: show running-config interface port-channel18-19
!Time: Thu Aug 30 11:26:00 2012

version 5.0(3)N1(1b)

interface port-channel18
  description to FI-A
  switchport mode trunk
  vpc 18
  switchport trunk allowed vlan 1,40-41,45

interface port-channel19
  description to FI-B
  switchport mode trunk
  vpc 19
  switchport trunk allowed vlan 1,40-41,45

UCS-N5k-FabA# show running-config interface ethernet 1/18-19

!Command: show running-config interface Ethernet1/18-19
!Time: Thu Aug 30 11:26:15 2012

version 5.0(3)N1(1b)

interface Ethernet1/18
  description vpc port-channel
  switchport mode trunk
  switchport trunk allowed vlan 1,40-41,45
  channel-group 18 mode active

interface Ethernet1/19
  description vpc port-channel
  switchport mode trunk
  switchport trunk allowed vlan 1,40-41,45
  channel-group 19 mode active

UCS-N5k-FabA#
```

Configuring Storage Connectivity

From each switch one link connects to each storage processor on the VNX5500 storage array. A virtual port-channel is created for the two links connected to a single storage processor, but connected to two different switches. In this example configuration, links connected to the storage processor A (SP-A) of VNX5500 storage array are connected to Ethernet port 1/26 on both the switches and links connected to the storage processor B (SP-B) are connected to Ethernet port 1/25 on both the switches. A virtual port-channel (id 26) is created for the Ethernet port 1/26 on both the switches and another virtual port-channel (id 25) is created for the Ethernet port 1/25 on both the switches.

Note The ports are in the access mode since only storage VLAN is required on these ports.

Figure 23 shows the configuration on the port-channels and interfaces.

Figure 23 Configuration of Port-channel and Interfaces

```
UCS-N5k-FabA# show running-config interface port-channel 25-26

!Command: show running-config interface port-channel25-26
!Time: Thu Aug 30 11:13:15 2012

version 5.0(3)N1(1b)

interface port-channel25
  description to VNX5500 SP-B
  vpc 25
  switchport access vlan 40

interface port-channel26
  description to VNX5500 SP-A
  vpc 26
  switchport access vlan 40

UCS-N5k-FabA# show running-config interface ethernet 1/25-26

!Command: show running-config interface Ethernet1/25-26
!Time: Thu Aug 30 11:13:26 2012

version 5.0(3)N1(1b)

interface Ethernet1/25
  switchport access vlan 40
  channel-group 25 mode active

interface Ethernet1/26
  switchport access vlan 40
  channel-group 26 mode active

UCS-N5k-FabA#
```

Configuring Ports Connected To Infrastructure Network

Port connected to infrastructure network need to be in trunk mode, and they require at least infrastructure VLAN, N1k control and packet VLANs at the minimum. You may require enabling more VLANs as required by your application domain. For example, Windows virtual machines may need to access to active directory / DNS servers deployed in the infrastructure network. You may also want to enable port-channels and virtual port-channels for high availability of infrastructure network.

Verify VLAN and Port-channel Configuration

At this point of time, all ports and port-channels are configured with necessary VLANs, switchport mode and vPC configuration. Validate this configuration using the “show vlan”, “show port-channel summary” and “show vpc” commands as shown in Figure 24.

Note The ports will be “up” only after the peer devices are configured properly, so you should revisit this subsection after configuring the EMC VNX5500 storage array and Cisco UCS fabric interconnects.

Figure 24 Validating Created Port-Channels with VLANs

```
UCS-N5k-FabA# show vlan id 40-45

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
40   Storage                          active    Po1, Po18, Po19, Po25, Po26
41   vMotion                          active    Po1, Po18, Po19
45   VM-DATA                          active    Po1, Po18, Po19
VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------

Remote SPAN VLANs
------------------------------------------------------------------------------

Primary  Secondary  Type             Ports
-------  ---------  ---------------  ------------------------------------------



UCS-N5k-FabA#
```

"show vlan" command can be restricted to a given VLAN or set of VLANs as shown in Figure 24. Ensure that on both switches, all required VLANs are in "active" status and right set of ports and port-channels are part of the necessary VLANs.

Port-channel configuration can be verified using "show port-channel summary" command. Figure 25 shows the expected output of this command.

Figure 25 Verifying Port-Channel Configuration

```
UCS-N5k-FabA# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
--------------------------------------------------------------------------------
Group Port-       Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth      LACP      Eth1/1(P)    Eth1/2(P)
18    Po18(SU)    Eth      LACP      Eth1/18(P)
19    Po19(SU)    Eth      LACP      Eth1/19(P)
25    Po25(SD)    Eth      LACP      Eth1/25(P)
26    Po26(SU)    Eth      LACP      Eth1/26(P)
UCS-N5k-FabA#
```

In this example, port-channel 1 is the vPC peer-link port-channel, port-channels 25 and 26 are connected to the storage arrays and port-channels 18 and 19 are connected to the Cisco UCS FI A and B. Make sure that the state of the member ports of each port-channel is "P" (Up in port-channel).

Note The port may not show "up" if the peer ports are not configured properly.

Common reasons for port-channel port being down are:

- Port-channel protocol mis-match across the peers (LACP v/s none)
- Inconsistencies across two vPC peer switches. Use "show vpc consistency-parameters **{global | interface {port-channel | port} <id>}** command to diagnose such inconsistencies.

vPC status can be verified using "show vpc" command. Example output is shown in Figure 26.

Figure 26 Verifying VPC Status

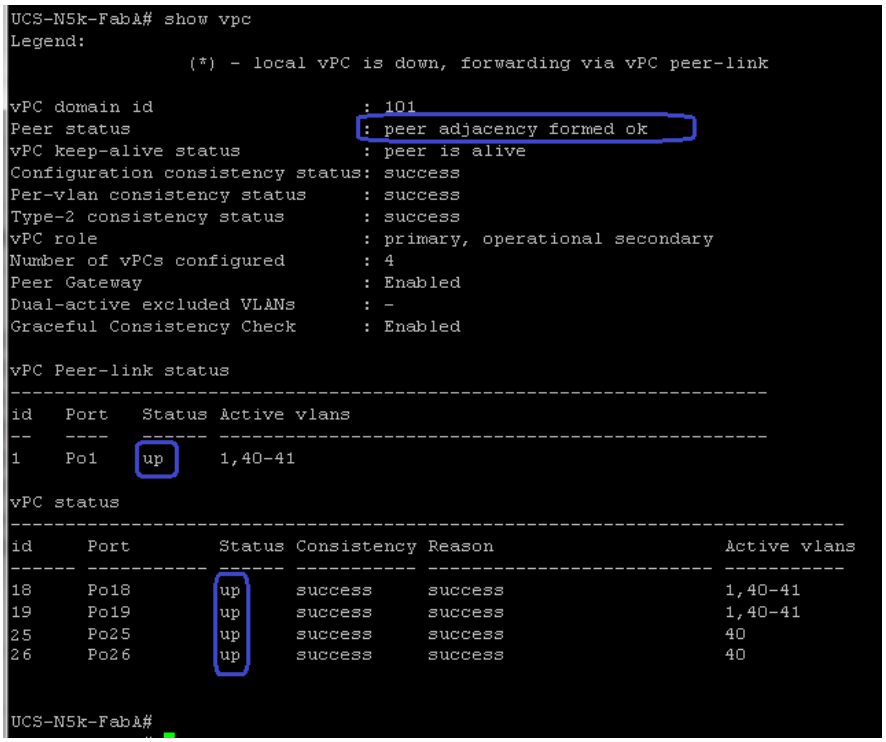

Ensure that the vPC peer status is “peer adjacency formed ok” and all the port-channels, including the peer-link port-channel status are “up”.

Configuring QoS

The Cisco solution for the EMC VSPEX VMware architectures require MTU to be set at 9216 (jumbo frames) for efficient storage and vMotion traffic. MTU configuration on Cisco Nexus 5000 fall under global QoS configuration. You may need to configure additional QoS parameters as needed by the applications. For more information on the QoS configuration, see *Cisco Nexus 5000 Series Configuration Guide*.

To configure jumbo MTU on the Cisco Nexus 5000 series switches, follow these steps on both switch A and B:

1. Create a policy map named “jumbo-mtu”.
2. As we are not creating any specific QoS classification, set 9216 MTU on the default class.
3. Configure the system level service policy to the “jumbo-mtu” under the global “system qos” sub-command.

Figure 27 shows the exact Cisco Nexus CLI for the steps mentioned above.

Figure 27 Configuring MTU on Cisco Nexus Switches

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# policy-map type network-qos jumbo-mtu
UCS-N5k-FabA(config-pmap-nq)# class type network-qos class-default
UCS-N5k-FabA(config-pmap-nq-c)# mtu 9216
UCS-N5k-FabA(config-pmap-nq-c)# exit
UCS-N5k-FabA(config-pmap-nq)# exit
UCS-N5k-FabA(config)# system qos
UCS-N5k-FabA(config-sys-qos)# service-policy type network-qos jumbo-mtu
UCS-N5k-FabA(config-sys-qos)# exit
UCS-N5k-FabA(config)#
UCS-N5k-FabA(config)#
UCS-N5k-FabA(config)# interface port-channel 25-26
UCS-N5k-FabA(config-if-range)# untagged cos 5
UCS-N5k-FabA(config-if-range)# exit
UCS-N5k-FabA(config)# exit
UCS-N5k-FabA#
```

Note Figure 27 shows the NXOS interface range CLI to configure multiple interfaces at the same time.

Configuring Cisco Unified Computing System Using Cisco UCS Manager

We would use web interface of Cisco UCS Manager (UCSM) to configure Cisco Unified Computing System. Cisco Unified Computing System configuration is broadly divided in two parts:

- Global and uplink configuration—Global configuration includes global VLAN and VSAN configuration, uplink Ethernet and Fibre Channel configuration, and server side chassis and blade server related configuration.
- Service profile configuration—Service profile configuration includes definition of various identifier pools, service profile template and instance definitions, and service profile association.

To launch Cisco UCS Manager, access https://<UCSM-Virtual-IP>/. By default, Cisco UCS Manager uses self-signed certificate, and so, browser would give untrusted SSL certificate warning. Ignore the warning and allow the initial Web UI page to load. Click **Launch UCSM**. A Java applet gets automatically downloaded and the Cisco UCS Manager login page appears. Enter the administrator's username/ password. Provide the right credential and let the Java based Cisco UCS Manager client application run.

Configuring VLANs

To create and configure VLANs, follow these steps:

1. In the Cisco UCS Manager window, select the **LAN** tab in the left pane, and right-click the **VLANs** under **LAN Cloud** as shown in Figure 28. Click **Create VLANs**.

Figure 28 Creating VLANs

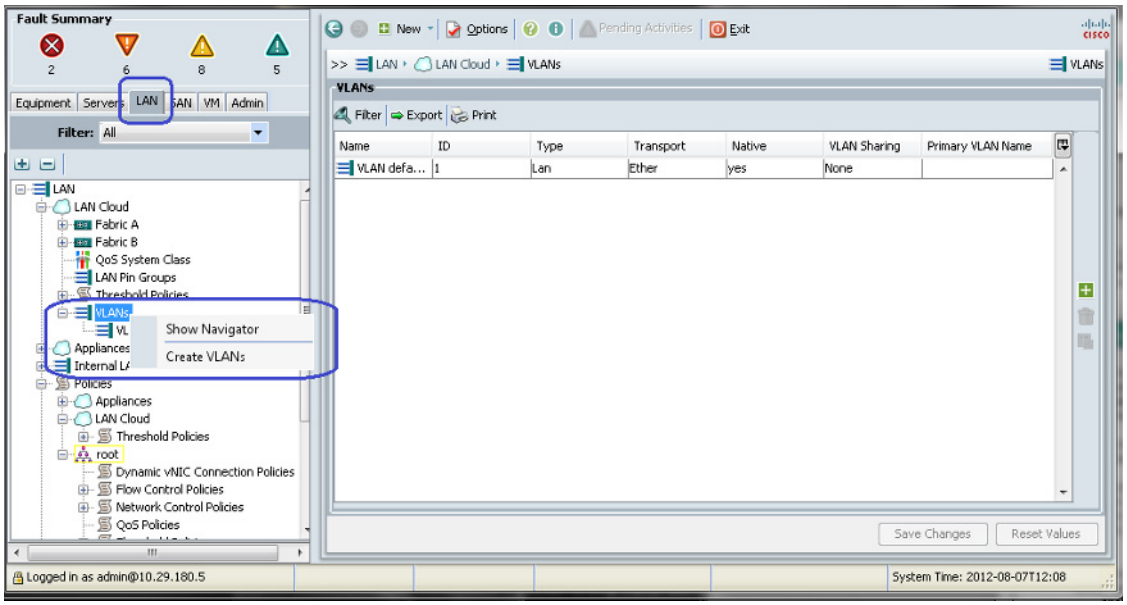

2. Enter the name of the VLAN (name cannot be changed later), VLAN ID and keep the sharing type to be default "None". Click **Ok**.

Figure 29 VLAN Details

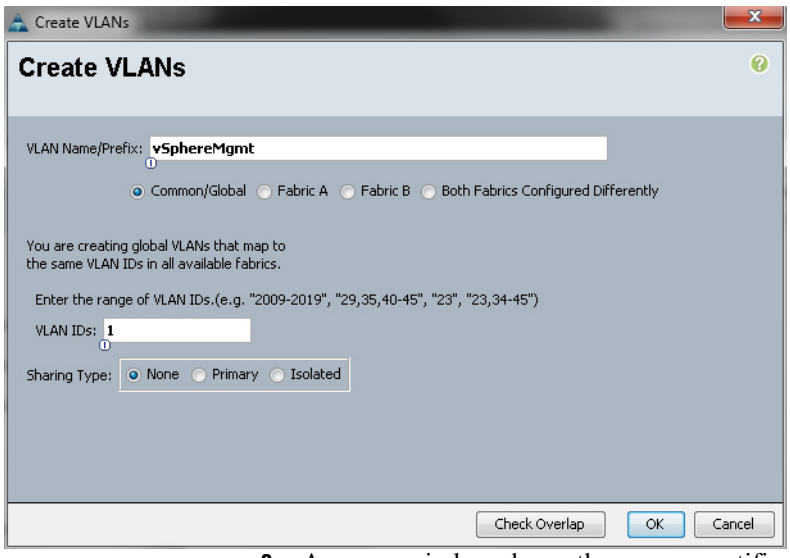

3. A popup window shows the success notification when the VLAN creation is complete.
4. Repeat steps 1 to 3 for all the VLANs required. For list of VLANs, see Configuring Cisco Nexus Switches, page 43. Figure 30 shows successful creation of all necessary VLANs.

Figure 30 Window Showing all the Created VLANs

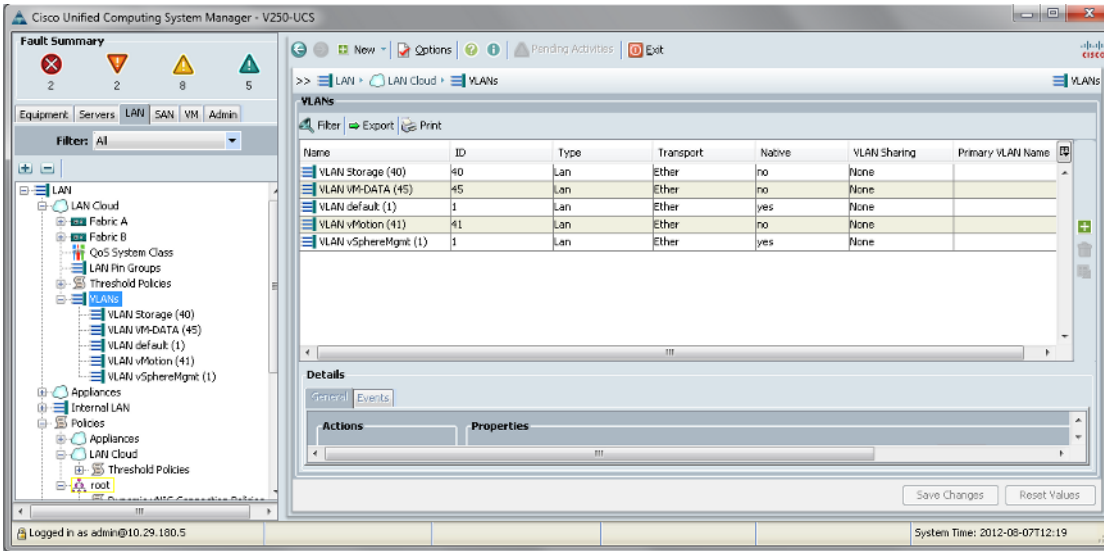

Configuring VSANs

1. In the Cisco UCS Manager window, select the **SAN** tab in the left pane, and right-click the **VSANs** under **SAN Cloud** as shown in the Figure 31. Click **Create VSAN**.

Figure 31 Creating VSANs

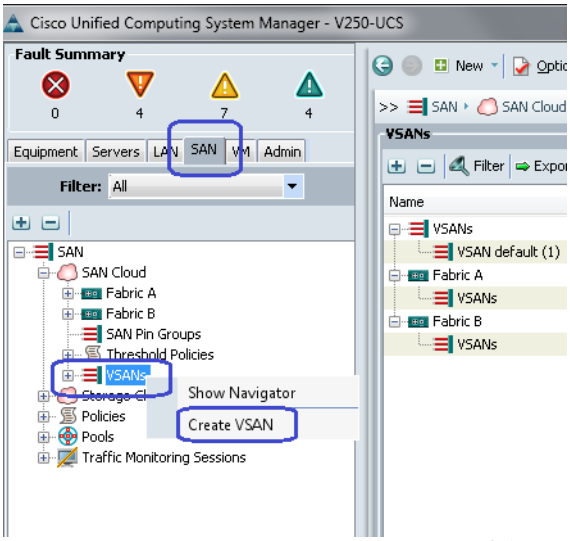

2. Enter the name of the VSAN (name cannot be changed later), enable default zoning, enter the VSAN id and the corresponding FCoE VLAN id. FCoE VLAN id can not be shared with any other VLANs defined from the Ethernet LAN domain.

Figure 32 VSAN Details

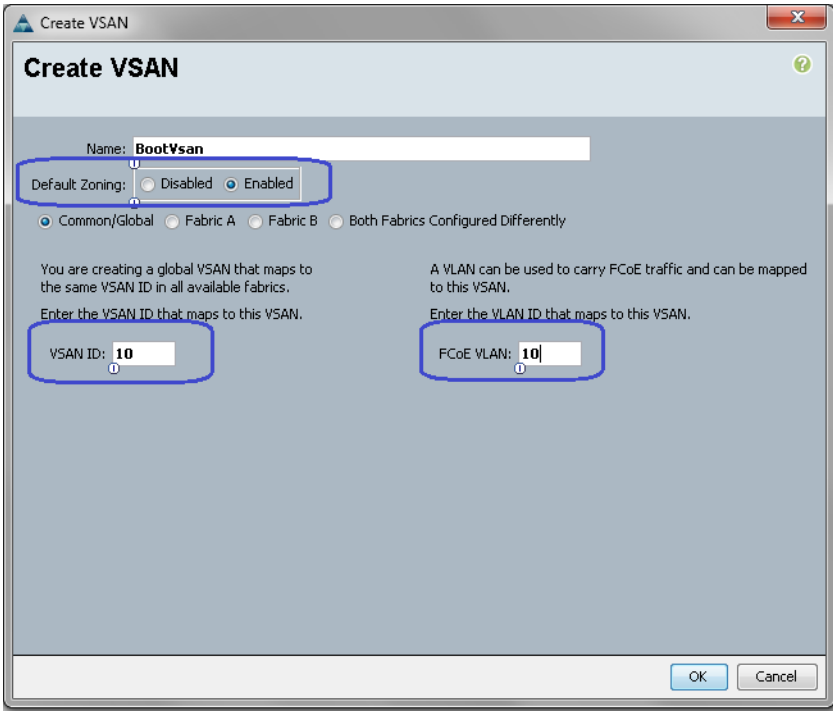

3. A popup window shows the success notification when the VSAN creation is complete.

Configure Fibre Channel Uplink Ports

Fibre Channel ports on the Cisco UCS FIs are classified as uplink ports by default, but they are under VSAN by default. Click the **SAN** tab in the UCSM window and select the uplink FC interface connected to the Cisco Nexus 5548UP switches. From the drop-down text box for VSAN, select Boot Vsan created in step 2 of Configuring VSANs as shown in Figure 33. Click **Save Changes**. Repeat this for all the uplink FC interfaces on both the Cisco UCS FIs.

Figure 33 Mapping FC Uplink Ports to Created VSAN

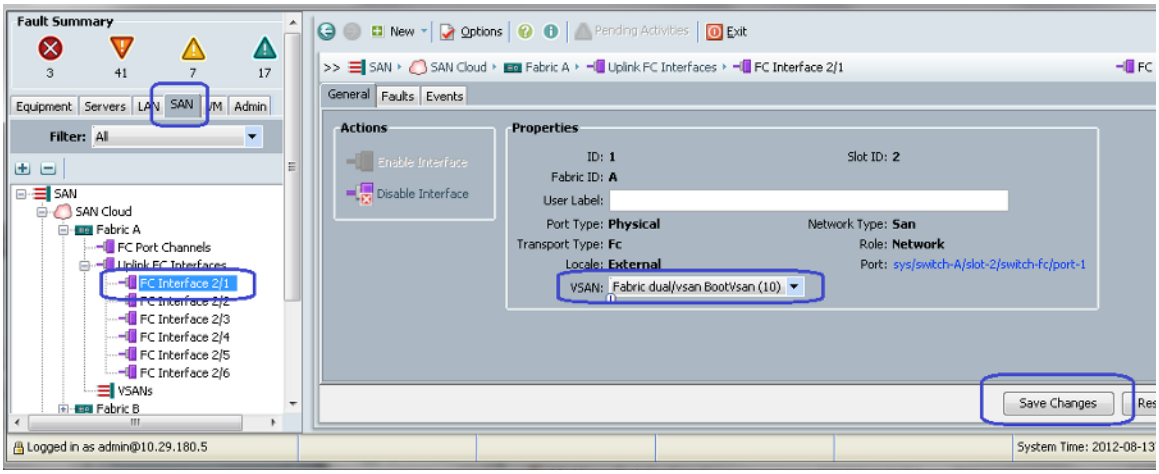

Configuring Ethernet Uplink Port-Channels

Virtual port-channels (vPC) on the Nexus 5548UP switches terminate on the UCS FIs as regular LACP port-channels. Follow these steps to configure uplink port-channels. Note that Ethernet ports on the UCS FIs are classified as "Unconfigured" by default, and need to be classified as uplink or server ports.

1. In the Cisco UCS Manager window, click the **SAN** tab. Expand **Fabric A** under **LAN Cloud** on the left pane of the Cisco UCS Manager window. Right-click on **Port Channels**. Click **Create Port Channel** to create port-channels on FI-A as shown in Figure 34.

Figure 34 Creating Port-Channel

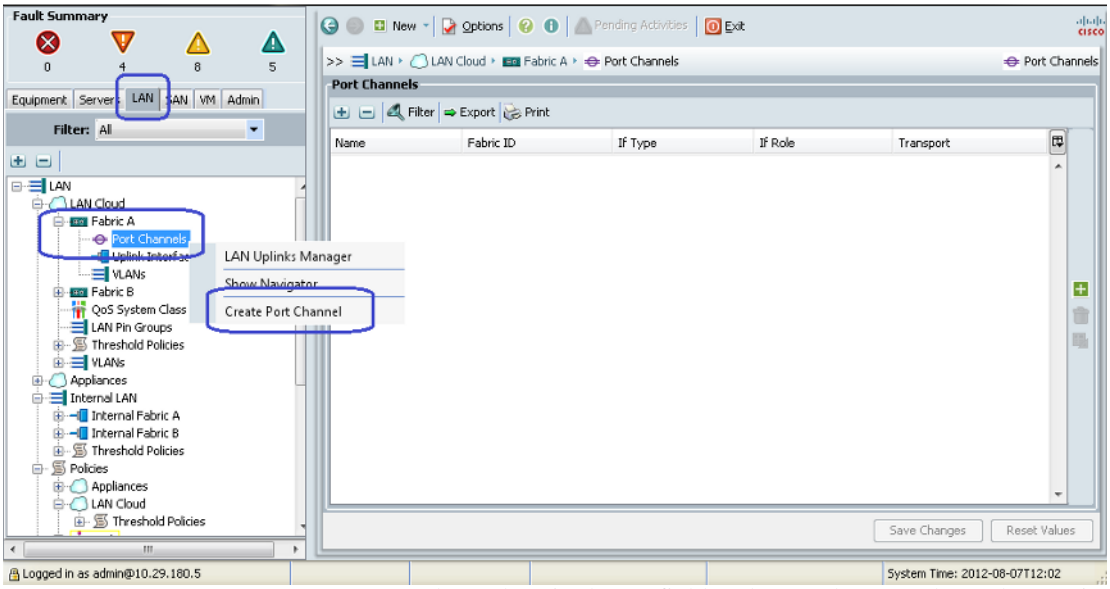

2. Enter port-channel ID in the ID field and enter the port-channel name in the Name field (optional). Port-channel ID has local significance, so it can be different from the vPC or port-channel ID configured on the Cisco Nexus 5548UP switches. Click **Next**.

Figure 35 Port-Channel Details

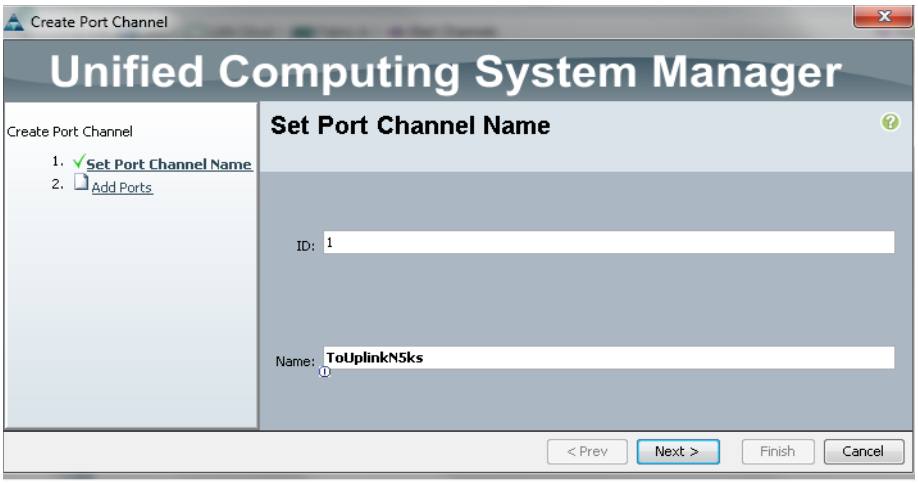

3. Select the ports that must be added as ports in the port-channel and click **Finish**.

Figure 36 Adding Ports to the Port-Channel

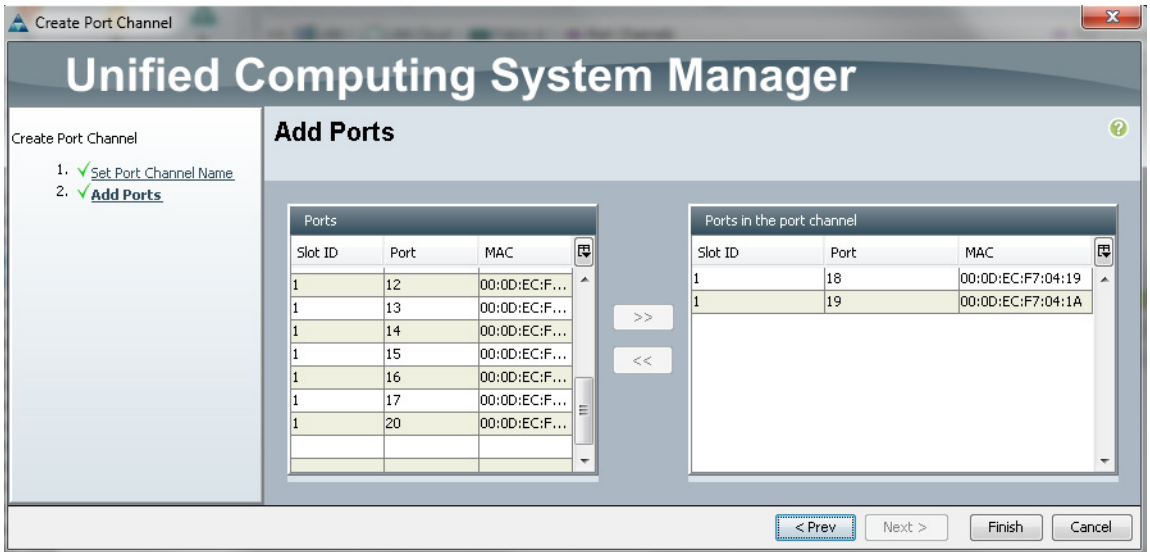

4. A popup window showing success notification will appear when the port-channel is created.
5. Port-channels are disabled by default in Cisco UCS Manager. To enable it, select the created port-channel and click **Enable Port Channel** link under Actions on the right pane of the Cisco UCS Manager window as shown in Figure 37.

Figure 37 Enabling Port-Channels in Cisco UCS Manager

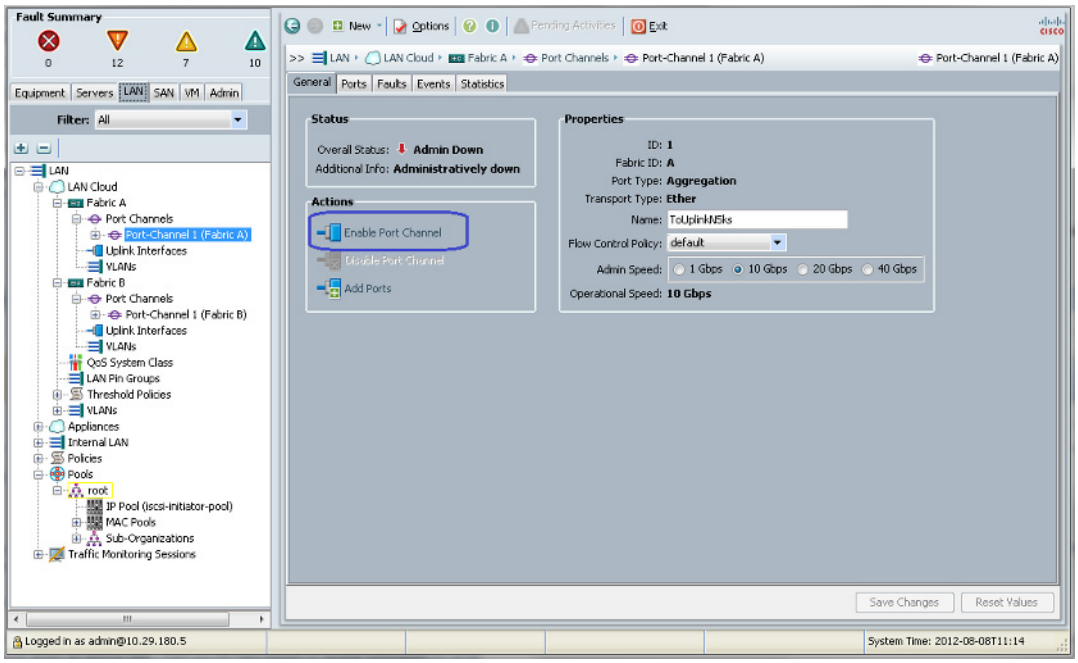

6. A popup window appears to confirm enabling of port-channels. Click **Yes** in the popup window.
7. Make sure that the Overall status is showing “UP” for all the port-channels.

Figure 38 Overall Status of all the Port-Channels

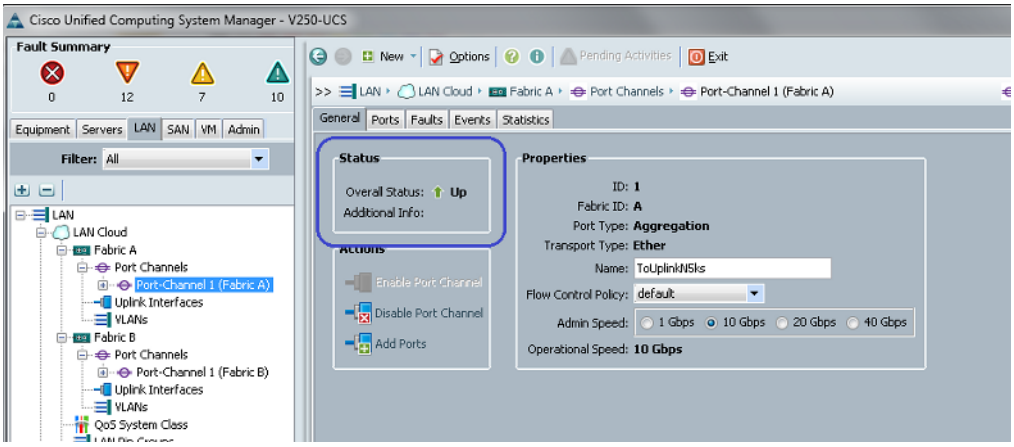

8. Repeat the steps 1 to 7 for "Fabric B" to create port-channel on FI-B.

Chassis and Server Discovery

After configuring uplink connectivity and global VLANs and VSANs, we need to configure server side connectivity for chassis and server discovery steps.

When the initial configuration in Cisco UCS Manager is completed through the serial console, the cluster state of the Cisco UCS Manager remains as "HA Not Ready" as shown in Figure 39. This is because there is no shared storage between two fabric interconnects due to lack of blade server chassis configuration on the UCS domain. Upon configuring two chassis in this solution, the HA state of the Cisco UCS Manager would transition to "HA Ready".

Figure 39 Cluster State of Cisco UCS Manager

```
Cisco UCS 6100 Series Fabric Interconnect
Using keyboard-interactive authentication.
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php

V250-UCS-A# show cluster state
Cluster Id: 0xc283335412a811df-0xac2c000decf70404

A: UP, PRIMARY
B: UP, SUBORDINATE

HA NOT READY
No device connected to this Fabric Interconnect
V250-UCS-A#
```

As this solution requires 25 VMs per server, we need 32 dynamic VNICs, two static VNICs and two more VHBAs per server, we need four links between Cisco UCS Fabric Interconnects and Cisco UCS Fabric Extenders. The default chassis discovery policy supports one link between chassis and FI, so, we need to change the chassis discovery policy to "4 Link".

To change chassis discovery policy, click the **Equipment** tab in the Cisco UCS Manager window, expand the Equipment tree root, and select the **Policies** tab on the right pane of the Cisco UCS Manager window as shown in Figure 40. Select the option "4 Link" from the "Action" drop down list in the "Chassis Discovery Policy" and click **Save Changes**.

Figure 40 Changing Chassis Discovery Policy Settings

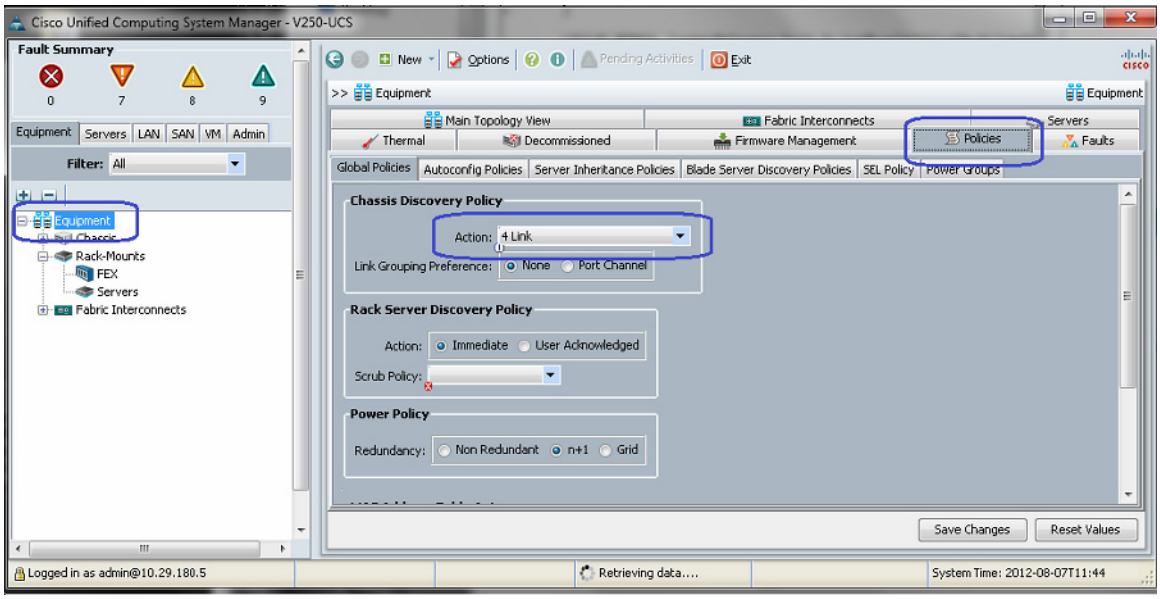

Marking Server Ports

After changing chassis discovery policy, next step is to classify interfaces connected to fabric extender as server ports on the fabric interconnect. Follow these steps for the same:

1. In the Cisco UCS Manager window, click the **Equipment** tab, expand "Fabric Interconnect A", expand "Fixed Module", and then expand "Unconfigured Ethernet Ports". Select the port that needs to be configured as the server port. On the right pane of the Cisco UCS Manager window, click the **Configure as Server Port** link as shown in Figure 41.

Figure 41 Selecting Port to Configure as Server Port

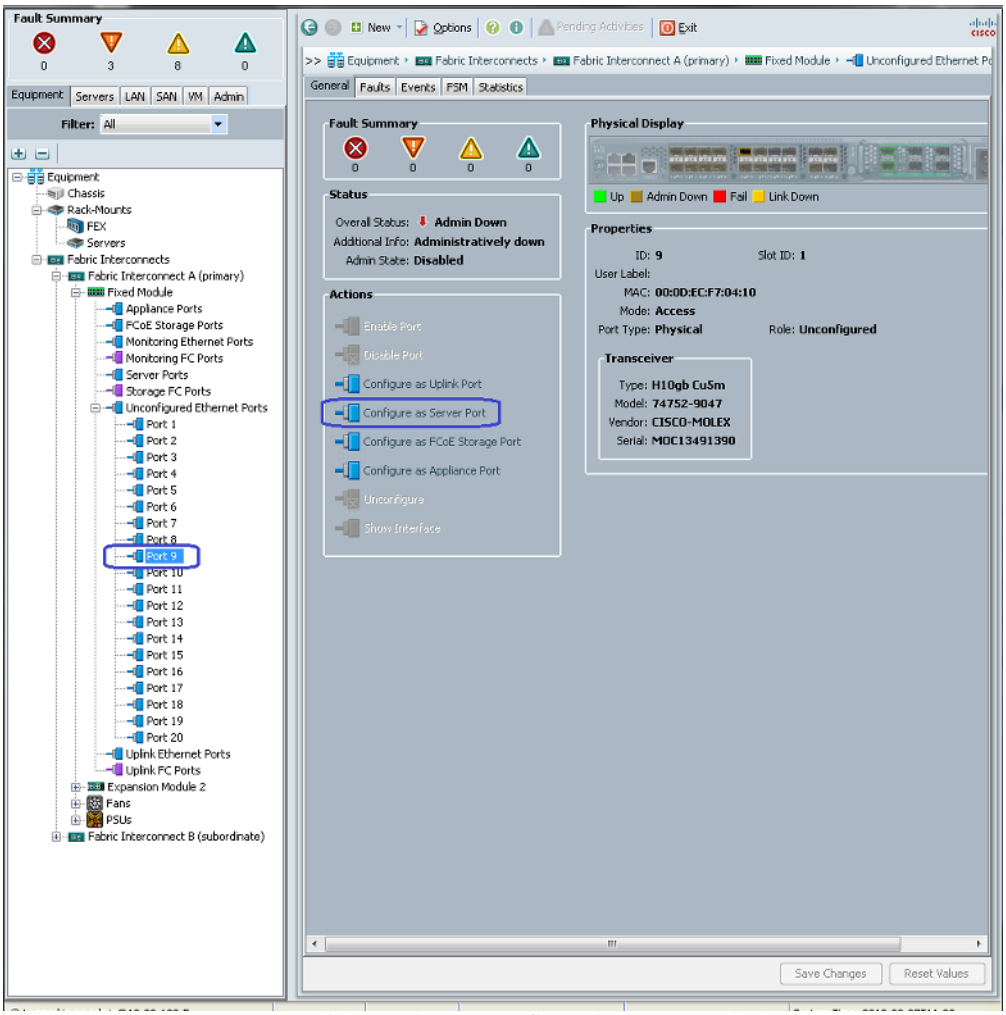

2. Click **Yes** in the confirmation popup window.
3. A success notification popup window appears when the port gets marked as a server port. Make sure that the "Overall Status" of the port shows "Up" as shown in Figure 42.

Figure 42 Overall Status of the Configured Port

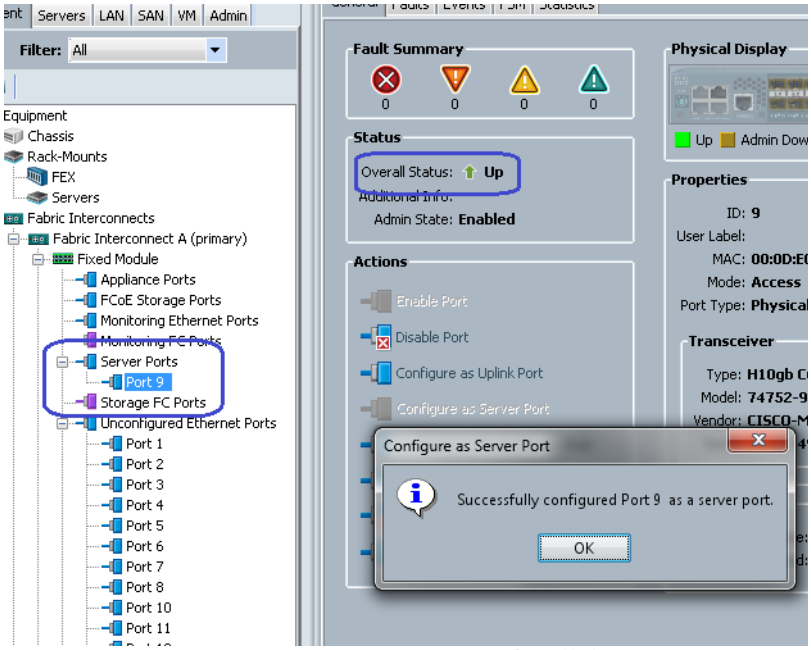

4. Repeat steps 1 to 3 for all the ports on FI-A and FI-B. Totally 16 ports will be marked as server ports.

When all the server ports are configured, Cisco UCS Manager will auto-discover the chassis connected to the Cisco UCS Fabric Interconnects. Chassis objects will show up under the **Equipment** tab in Cisco UCS Manager, and upon successful deep discovery of chassis, the "Overall Status" of the chassis will change to "Operable" as shown in Figure 43. Also ensure that the two IOMs (Fabric Extenders) are listed under each chassis by expanding the individual chassis.

Figure 43 Change in the Overall Status of the Chassis

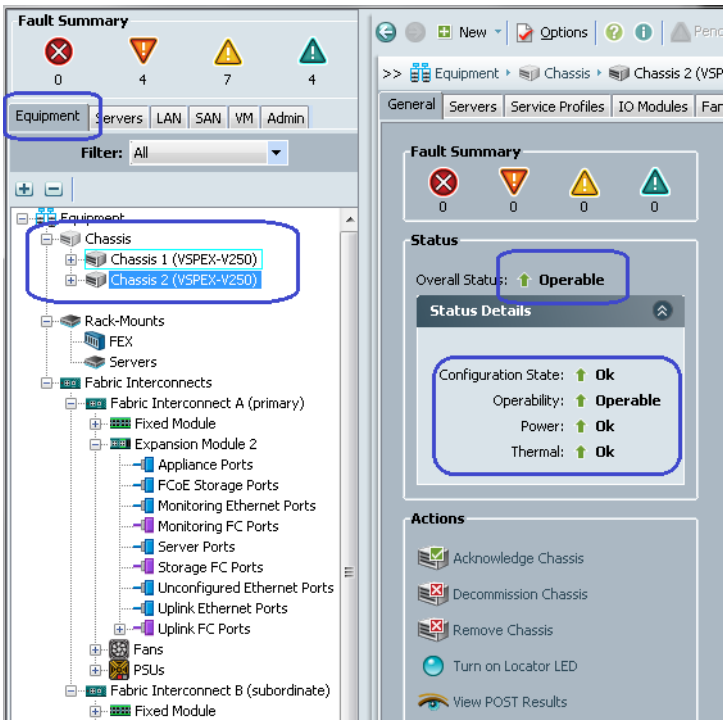

When the chassis is discovered, deep discovery of the embedded Cisco UCS B200 M3 Blade Servers would also get started. You can see the progress of the Blade Server discovery FSM. To see the discovery status, expand Chassis, expand Servers, select the required server and then click the **FSM** tab on right pane of the Cisco UCS Manager window as shown in Figure 44.

Figure 44 Discovery Status

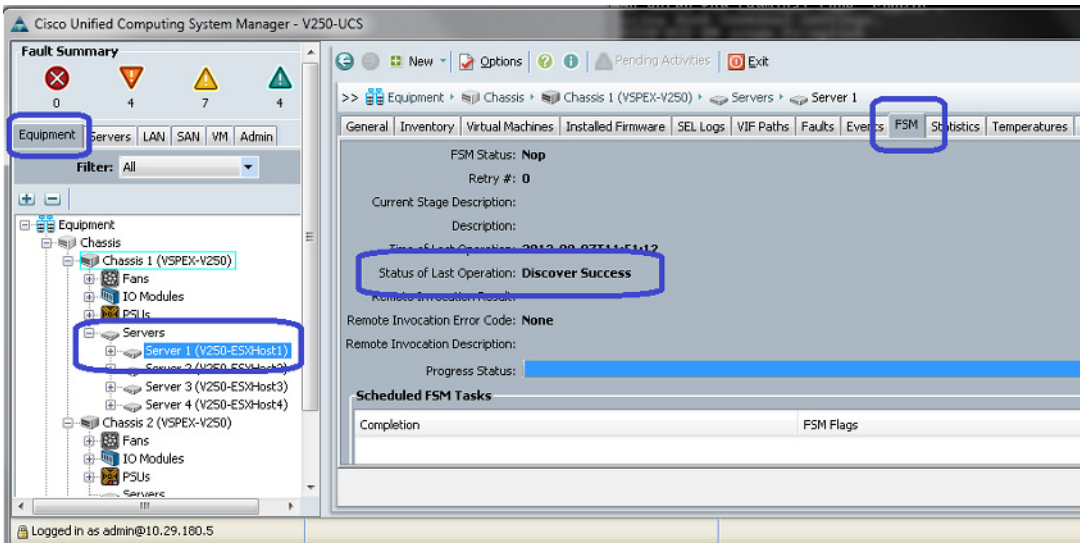

When deep discovery of the server is complete, the “Status of Last Operation” will change to “Discover Success” and “Progress Status” will reach 100%. The success notification of the server discovery will also show up and “Overall Status” becomes “Unassociated” in the **General** tab of Cisco UCS Manager.

Figure 45 Status Details of the Server

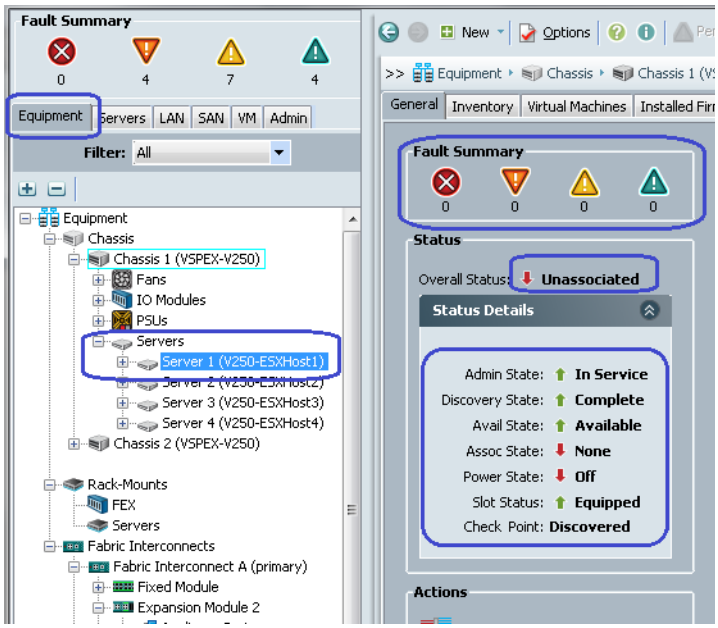

QoS Configuration

We need to configure system QoS classes and vNIC in QoS policies, which plays part in the end-to-end jumbo MTU configuration of the solution. Follow these steps to configure QoS on the Cisco UCS Manager:

1. In the Cisco UCS Manager window, click the **LAN** tab, expand LAN Cloud and click **QoS System Class**. On the right panel in the Cisco UCS Manager window, check the check box to enable "Platinum" and "Gold" classes, and keep the weight at "5" for all the classes. Set the MTU of both of these classes at 9216 bytes to support jumbo MTU as shown in Figure 46. Click **Apply Changes**.

Figure 46 Enabling Classes in Cisco UCS Manager

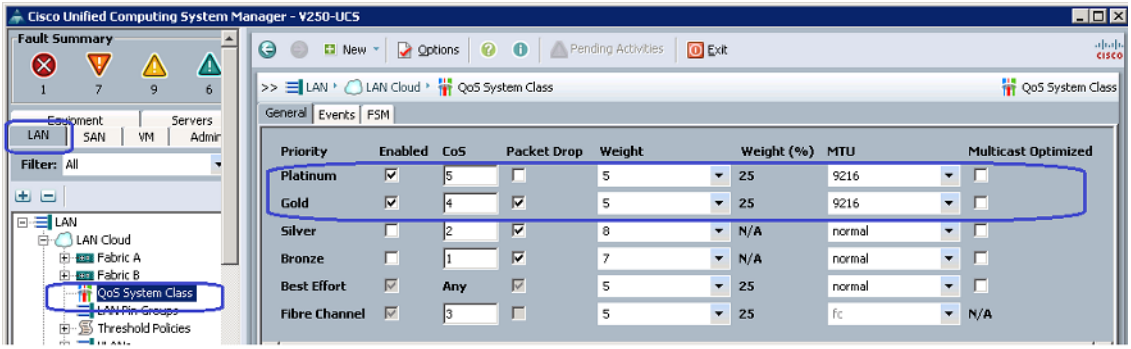

2. To configure QoS policies, in the Cisco UCS Manager window click the **LAN** tab, expand "Policies" and then expand "root". Right-click the **QoS Policies**. Click **Create QoS Policy** as shown in Figure 47.

Figure 47 Creating QoS Policies

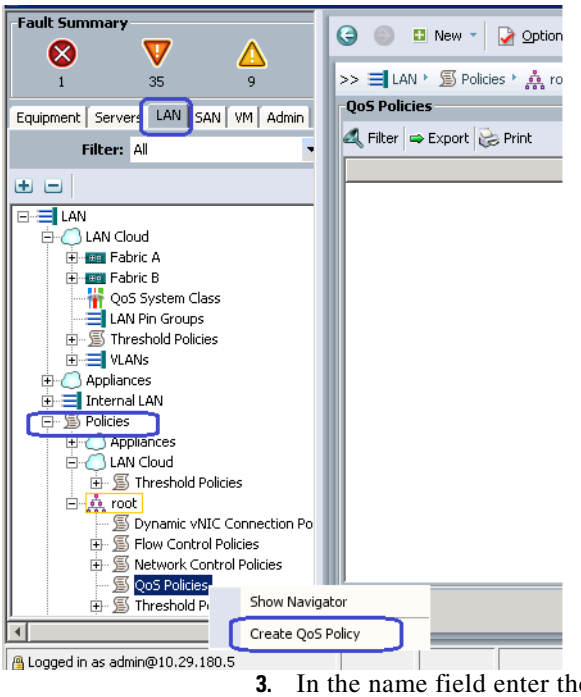

3. In the name field enter the name as "vMotion" of the Create QoS Policy window, and select the priority as "Platinum" from the drop-down list. Do not modify any other parameters and click **Ok** as shown in Figure 48.

Figure 48 Details to Create QoS Policies

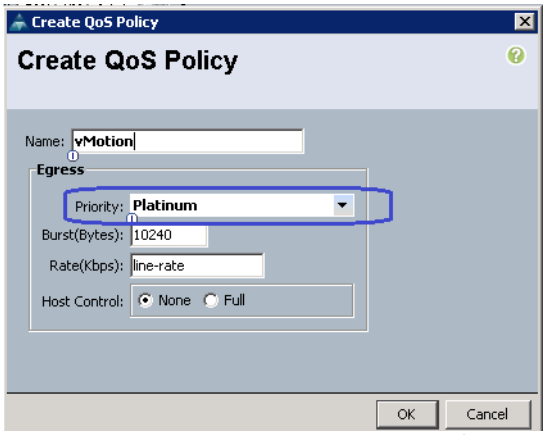

4. Repeat step 3 for creating another QoS policy with the name "NFS" and "Gold" as the priority as shown in Figure 49.

Figure 49 Details for Creating Another QoS Policy

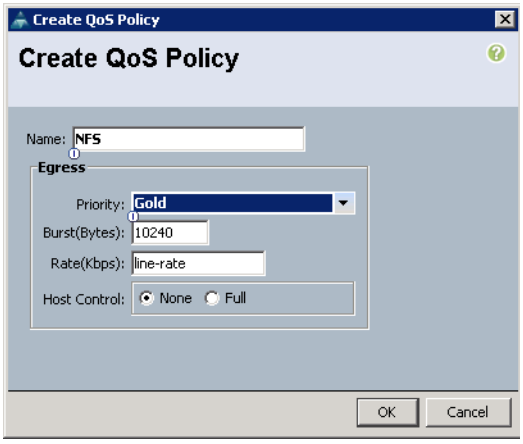

These QoS policies would be used in the port-profiles created at a later point during the solution deployment.

Configuring Identifier Pools

As described in the "Stateless Computing" section of the "Architectural Overview", Service Profile - a logical representation of the server – includes various identities. Best way to manage various identities is to configure identifier pools. We will begin with defining external management IP address pool for the servers. Most common use of external management IP address of the server is launch of KVM of the server. KVM also includes virtual CD-ROM media launch, which we would use at later point in deployment of this solution.

Configuring external management IP address pool

An IP address pool named "ext-mgmt" is predefined in Cisco UCS Manager. Follow these steps to populate the pool with IP addresses for the out-of-band management of the Cisco UCS B200 M3 blade servers.

1. Click the **Admin** tab on the left pane of Cisco UCS Manager, select "Communication Management" filter from the drop-down list. Expand the Communication Management and click **Management IP Pool (ext-mgmt)**. On the right pane of the Cisco UCS Manager window, under the "Actions" area, click **Create Block of IP Addresses** link as shown in Figure 50.

Figure 50 Creating Block of IP Addresses for the Pool

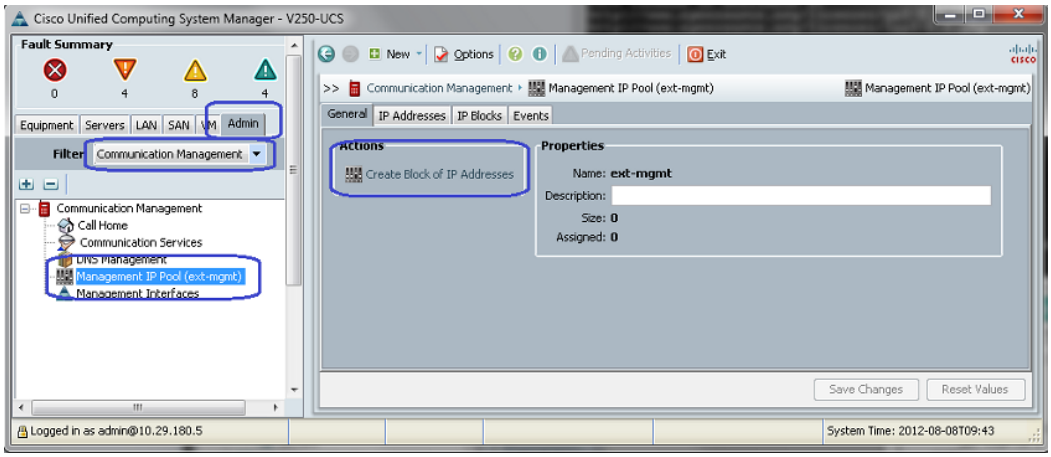

2. In the "Create Block of IP Addresses" wizard, provide the start of the IP addresses range and specify the size of the pool. You need at least 10 IP addresses, but you may want to provide larger number for the future expansion. Also provide "Subnet Mask" and "Default Gateway" associated with the IP addresses as shown in the Figure 51. Click **Ok**.

Figure 51 Entering Parameters to Create Block of IP Addresses

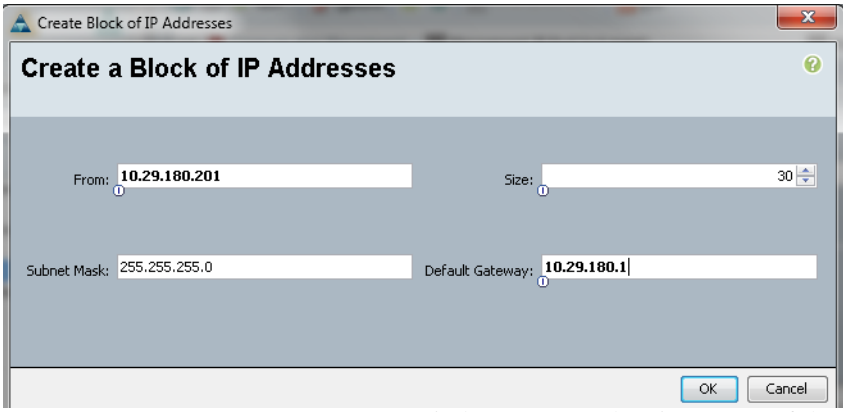

3. A pop-up window appears showing successful completion of creating a block IP addresses.
4. To see the out-of-band external management IP address, click the **IP Addresses** tab in the right pane to see the assigned IP addresses as shown in Figure 52.

Figure 52 List of Assigned IP Addresses

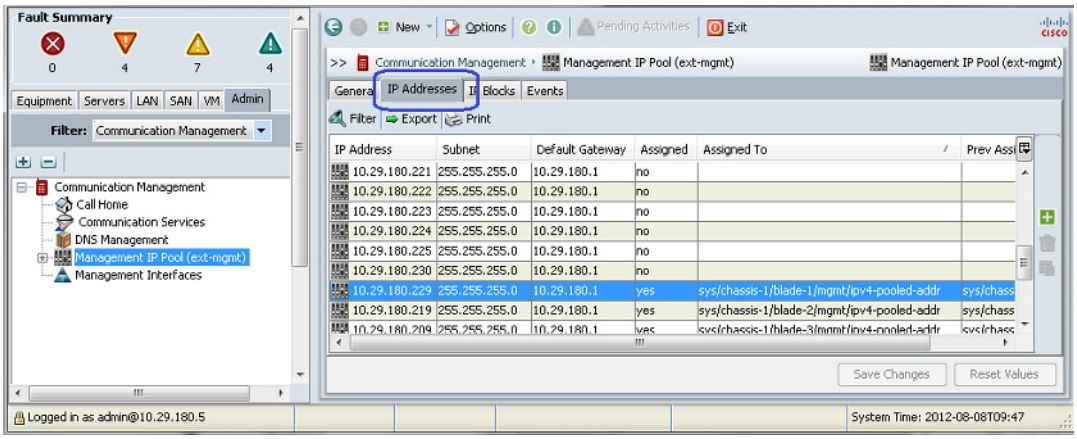

Configure UUID Pool

Follow these steps to create B200 M3 servers' BIOS UUID pool:

1. In the Cisco UCS Manager window, click the **Servers** tab, expand "Pools" and right-click on "UUID Suffix Pools" as shown in Figure 53. Click **Create UUID Suffix Pool**.

Figure 53 Creating UUID Suffix Pool

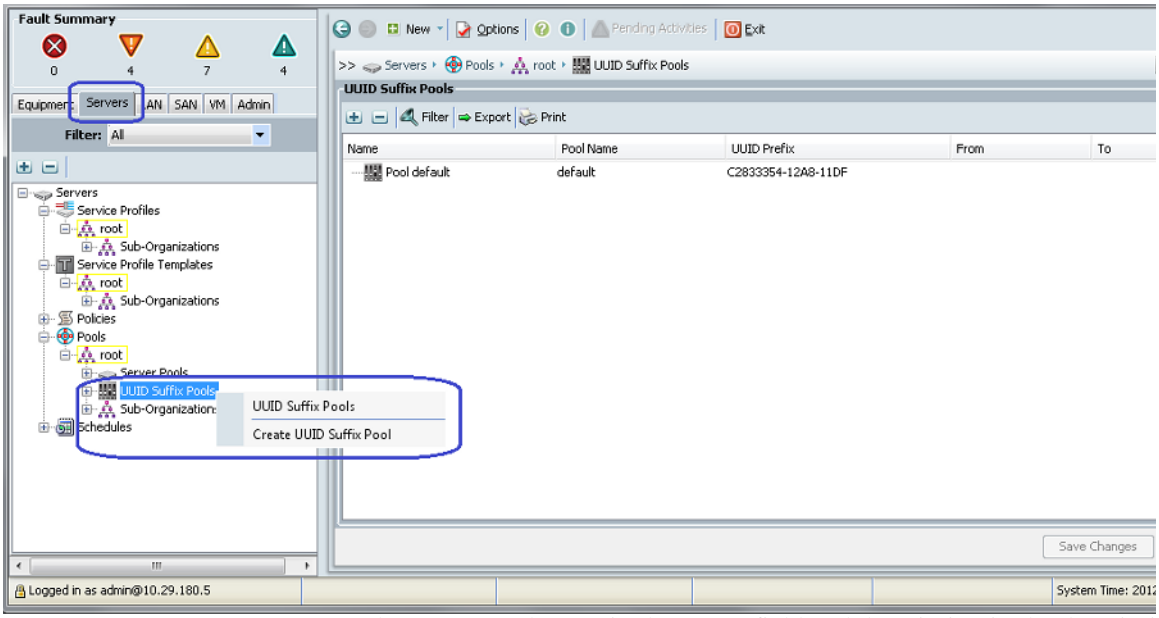

2. Enter the UUID pool name in the Name field and description in the description field (optional) as shown in Figure 54. Keep the prefix as "Derived" which is the default option. Click **Next**.

Figure 54 Entering Details for Creating UUID Suffix Pool

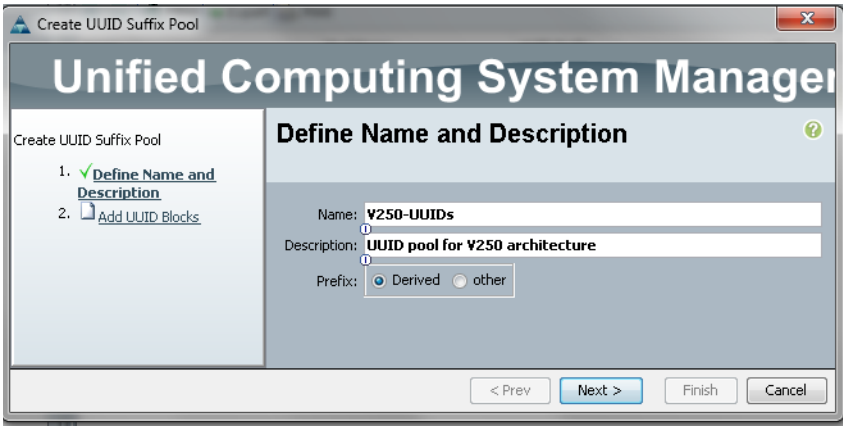

3. Click **Add** in the "Add UUID Blocks" window as shown in Figure 55.

Figure 55 Adding UUID Blocks

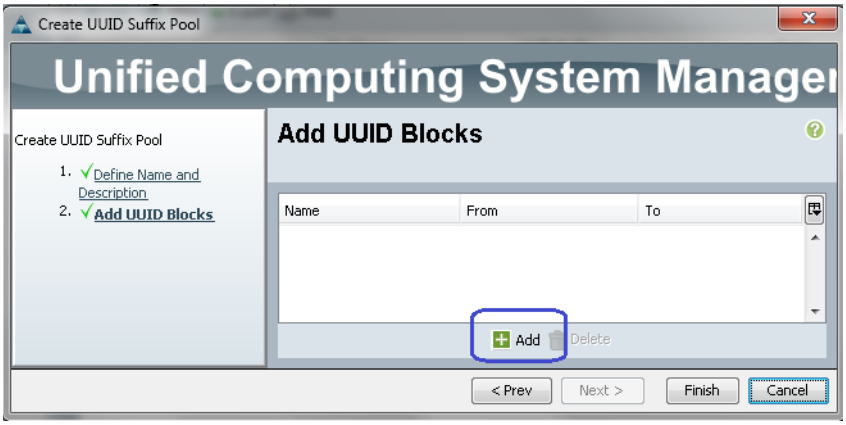

4. Enter the beginning of the UUID block range and the size of the UUID block as shown in Figure 56. You need at least 10 UUIDs in the pool, but you may want to keep larger size considering future expansion. Click **Ok** and click **Finish**.

Figure 56 Entering Parameters to Create a Block of UUID Suffixes

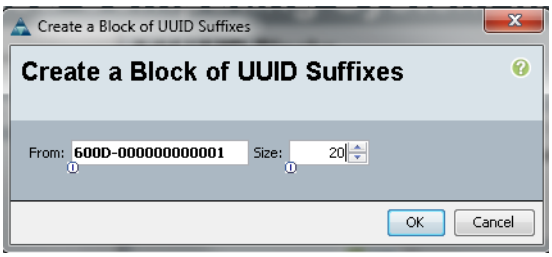

5. A pop-up window appears showing successful completion of creating a block of UUID suffixes.

Configure MAC Address Pools

For each ESXi host in the solution, we would need two vNIC cards, so we need at least 20 unique MAC addresses defined in a MAC address pool. Follow these steps to create MAC address pool.

1. In the Cisco UCS Manager window, click the **LAN** tab, expand "Pools" and right-click on "MAC Pools" and click **Create MAC Pool**.

Figure 57 Creating MAC Address Pool

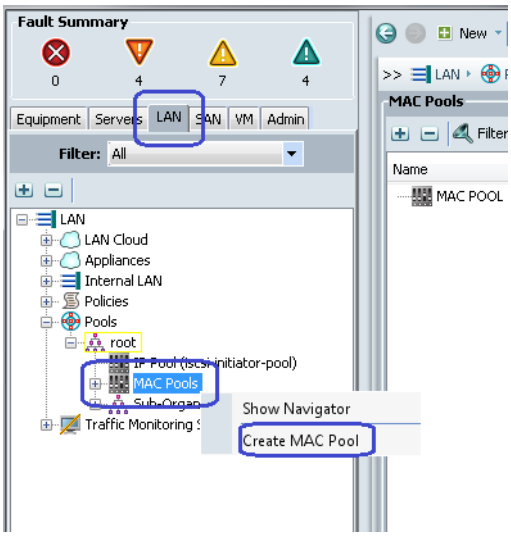

2. Enter the MAC address pool name in the Name field and description in the description field (optional) as shown in Figure 58. Click **Next**.

Figure 58 Entering Details for Creating MAC Address Pool

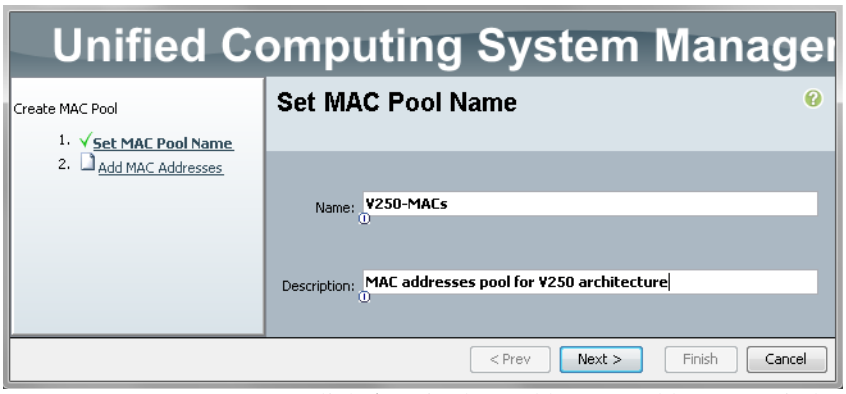

3. Click **Add** in the "Add MAC Addresses" window as shown in Figure 59.

Figure 59 Adding MAC Addresses

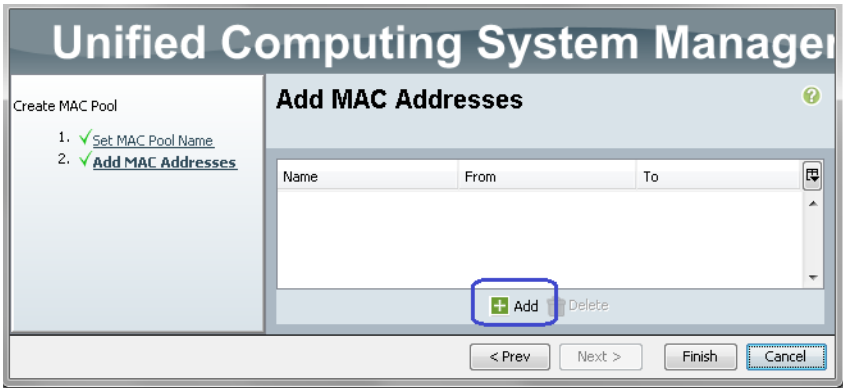

4. Enter the beginning of the MAC addresses block range and the size of the MAC addresses block as shown in Figure 60. You need at least 20 MAC addresses in the pool, but you may want to keep larger size considering future expansion. Also note that to ensure uniqueness of MAC addresses across the data center, OUI part of the MAC address must be kept 00:25:B5. Click **Ok** and click **Finish**.

Figure 60 Entering Parameters to Create a Block of MAC Addresses

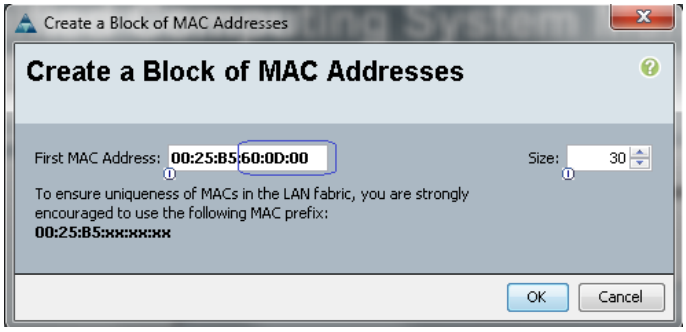

5. A pop-up window appears showing successful completion of creating a block of MAC Addresses.

Configuring WWNN Pool

Each ESXi host requires a unique Fibre Channel World Wide Node Name (WWNN), so we need at least 10 unique WWNN addresses defined in a WWN address pool. Follow these steps to create WWNN address pool.

1. In the Cisco UCS Manager window, click the **SAN** tab, expand "Pools" and right click on "WWNN Pools" and click **Create WWNN Pool**.

Figure 61 Creating WWNN Pool

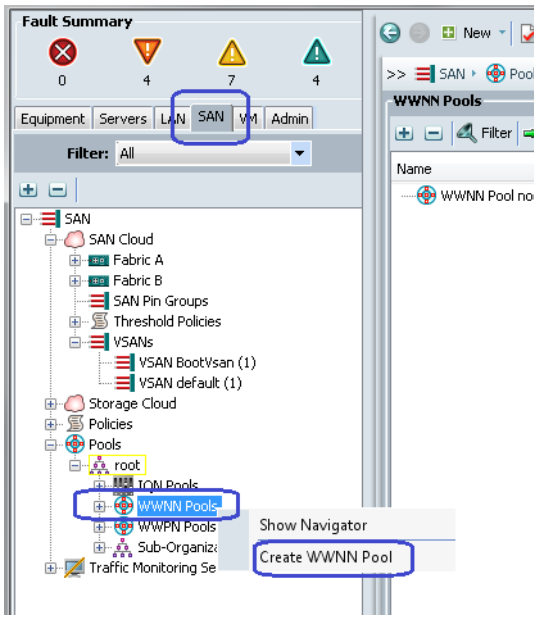

2. Enter the WWNN address pool name in the Name field and description in the description field (optional) as shown in Figure 62. Click **Next.**

Figure 62 Entering Details for Creating WWNN Address Pool

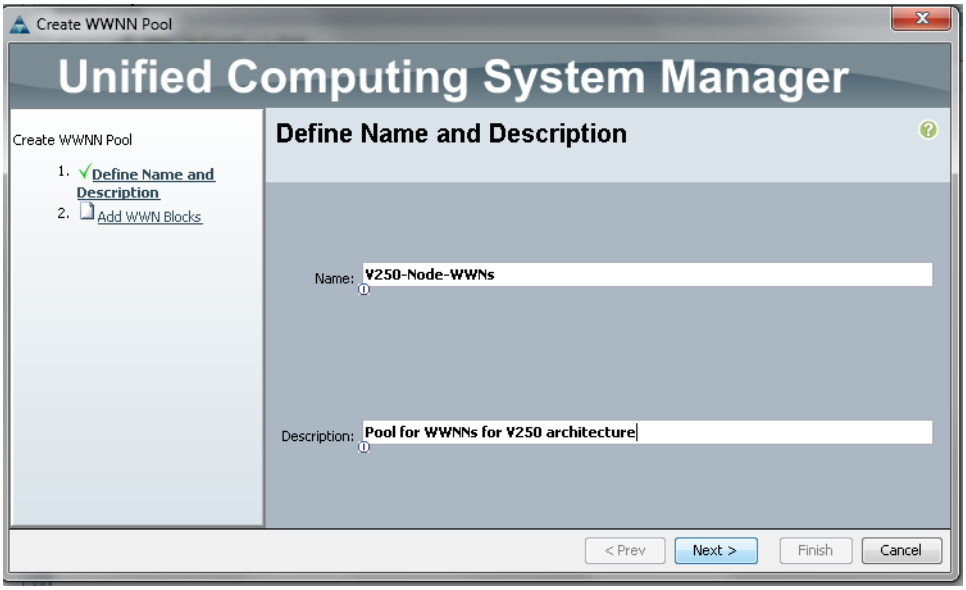

3. Click **Add** in the "Add WWNN Blocks" window as shown in Figure 63.

Figure 63 Adding WWNN Blocks

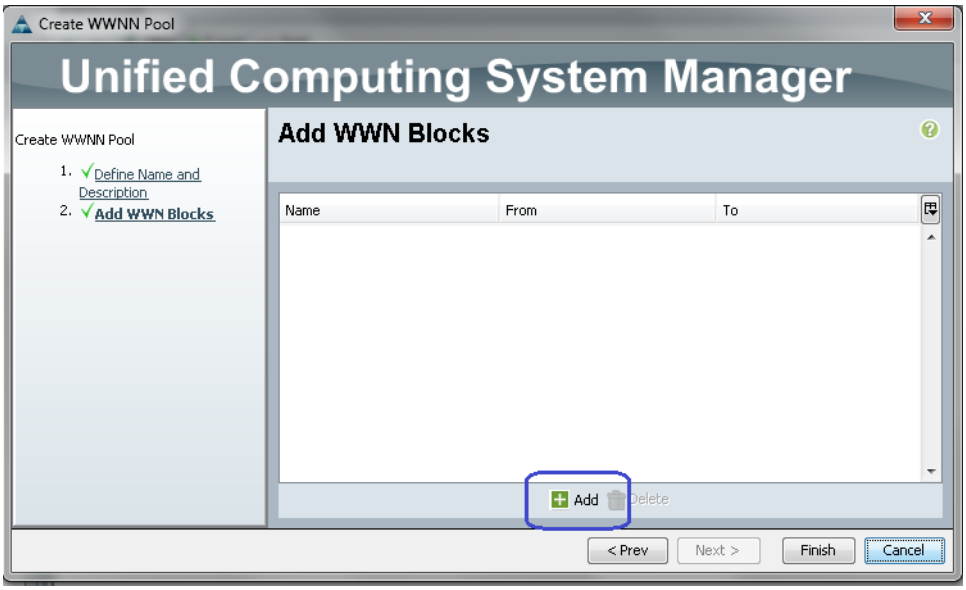

4. Enter the beginning of the WWNN addresses block range and the size of the WWNN addresses block as shown in Figure 64. You need at least 10 WWNN addresses in the pool, but you may want to keep larger size considering future expansion. Also note that to ensure uniqueness of WWNN addresses across the data center, prefix of the WWNN address must be kept 20:00:00:25:b5. Click **Ok** and click **Finish**.

Figure 64 Entering Parameters to Create WWNN Block

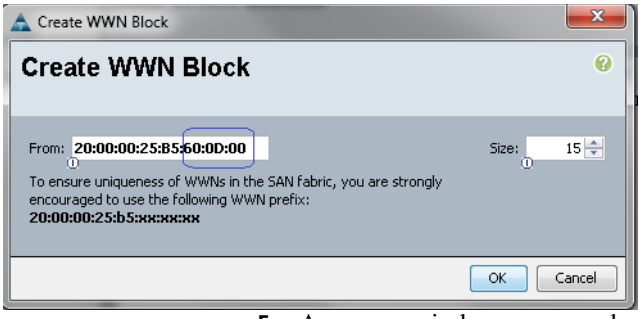

5. A pop-up window appears showing successful completion of creating a WWNN pool.

Configuring WWPN Pool

Each ESXi host in this solution has two vHBAs. Each vHBA requires a unique Fibre Channel World Wide Port Name (WWPN), so we need at least 20 unique WWPN addresses defined in a WWN address pool. Follow these steps to create WWPN address pool.

1. In the Cisco UCS Manager window, click the **SAN** tab, expand "Pools" and right click on "WWNN Pools" and click **Create WWPN Pool.**

Figure 65 Creating WWPN Pool

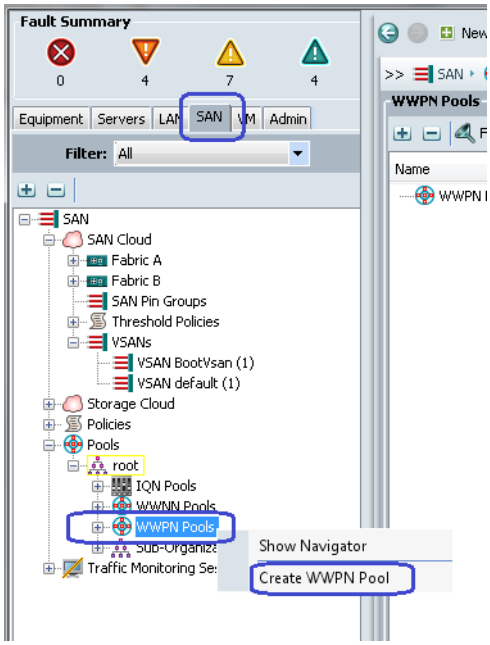

2. Enter the WWPN address pool name in the Name field and description in the description field (optional) as shown in Figure 66. Click **Next.**

Figure 66 Entering Details for Creating WWPN Address Pool

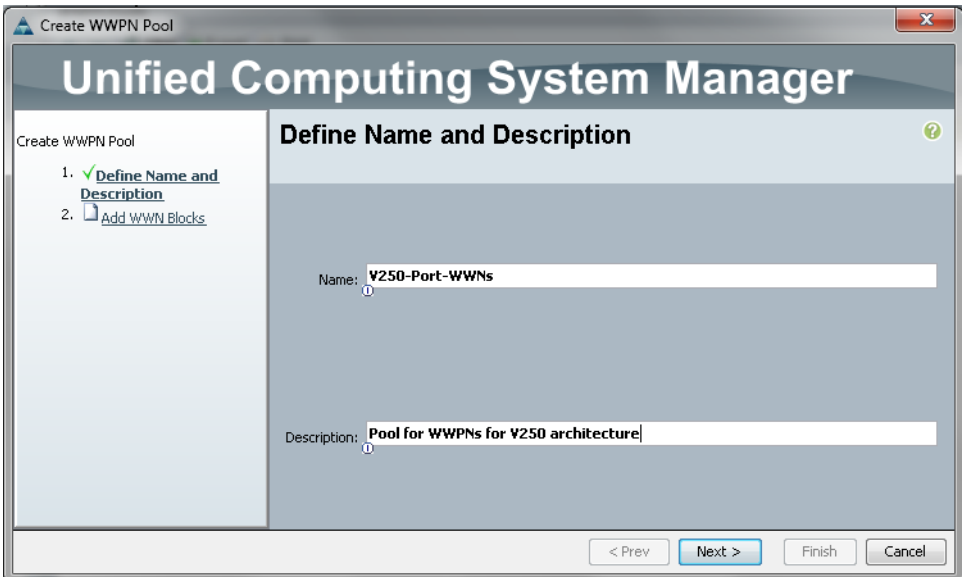

3. Click **Add** in the "Add WWPN Blocks" window as shown in Figure 67.

Figure 67 Adding WWNN Blocks

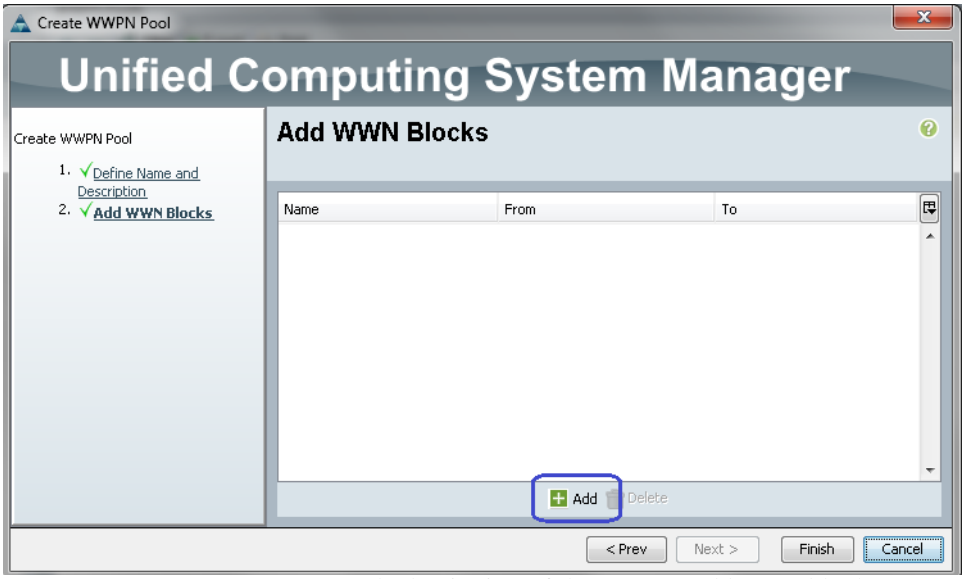

4. Enter the beginning of the WWPN addresses block range and the size of the WWNPN addresses block as shown in Figure 68. You need at least 10 WWPN addresses in the pool, but you may want to keep larger size considering future expansion. Also note that to ensure uniqueness of WWPN addresses across the data center, prefix of the WWNN address must be kept 20:00:00:25:b5. Click **Ok** and click **Finish.**

Figure 68 Entering Parameters to Create WWPN Block

5. A pop-up window appears showing successful completion of creating a WWPN pool.

Create Service Profile Template

After configuring all necessary pools, next step is to define Service Profile Template. Given that all ten B200 M3 Blade Servers have identical ESXi 5 hypervisor configuration, Service Profile Template is the most convenient approach. Follow these steps to configure service profile template.

1. In the Cisco UCS Manager window, click the **Servers** tab, expand "Service Profile Templates", right-click on "root" and click **Create Service Profile Template** as shown in Figure 69.

Figure 69 Creating Service Profile Template

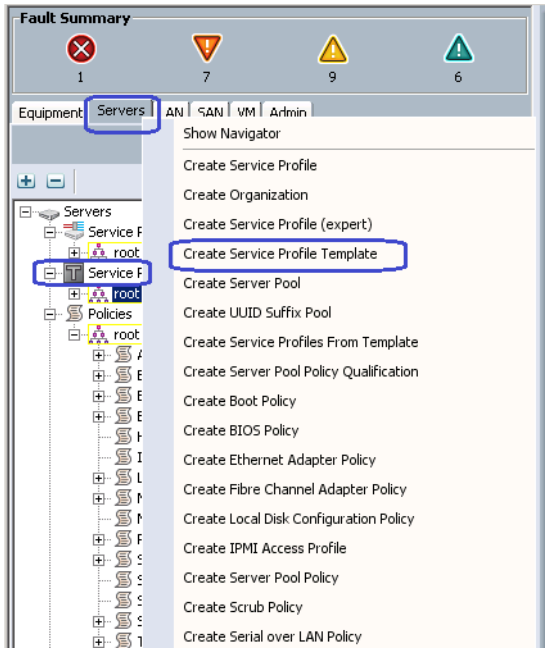

2. In the "Create Service Profile Template" wizard, enter the name of Service Profile Template. select the type as "Updating Template", select the name of UUID pool created in previous section from the "UUID Assignment" drop down list and enter description (optional) as shown in Figure 70.

Figure 70 Identify Service Profile Template Window

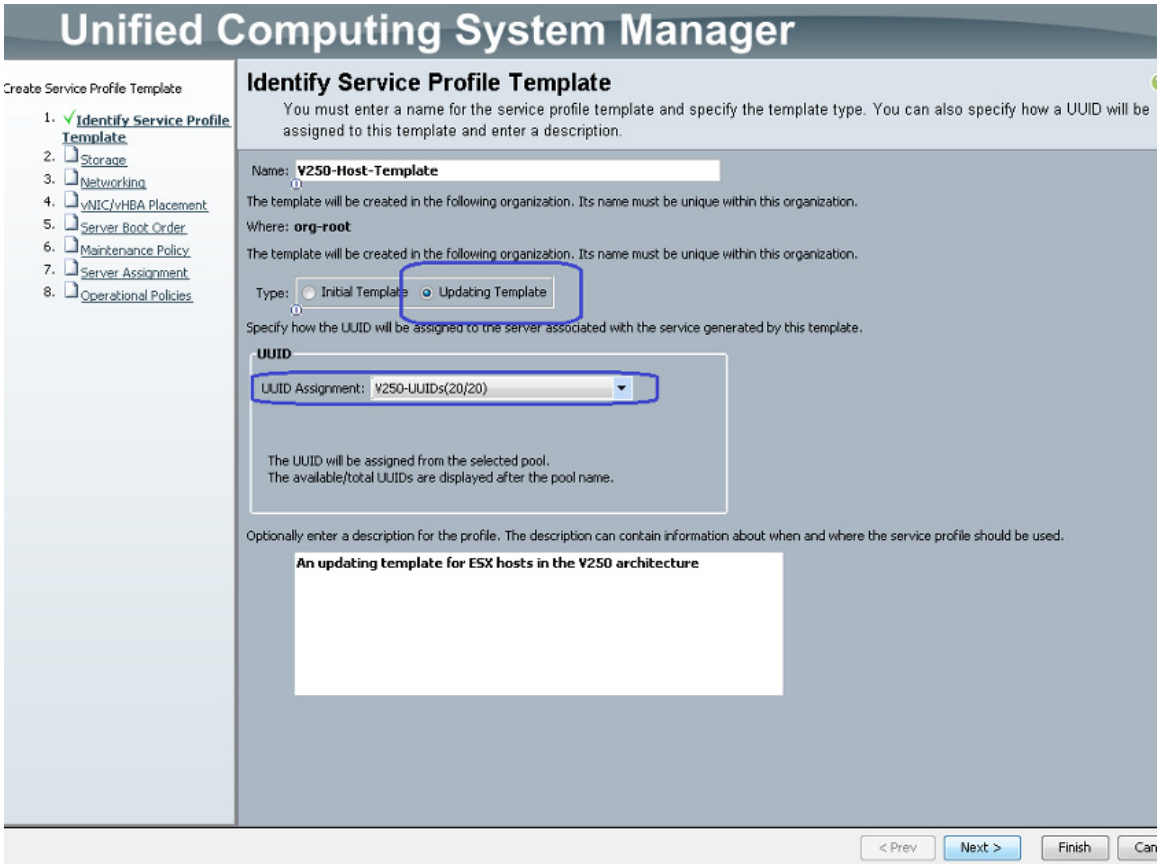

3. In the "Storage" window of the wizard, click **Expert** radio button for SAN connectivity and name of the WWNN pool created in previous step for the "WWNN Assignment" drop-down list as shown in Figure 71. Click **Add** to add vHBAs.

Figure 71 Storage Window of the Create Service Profile Template Wizard

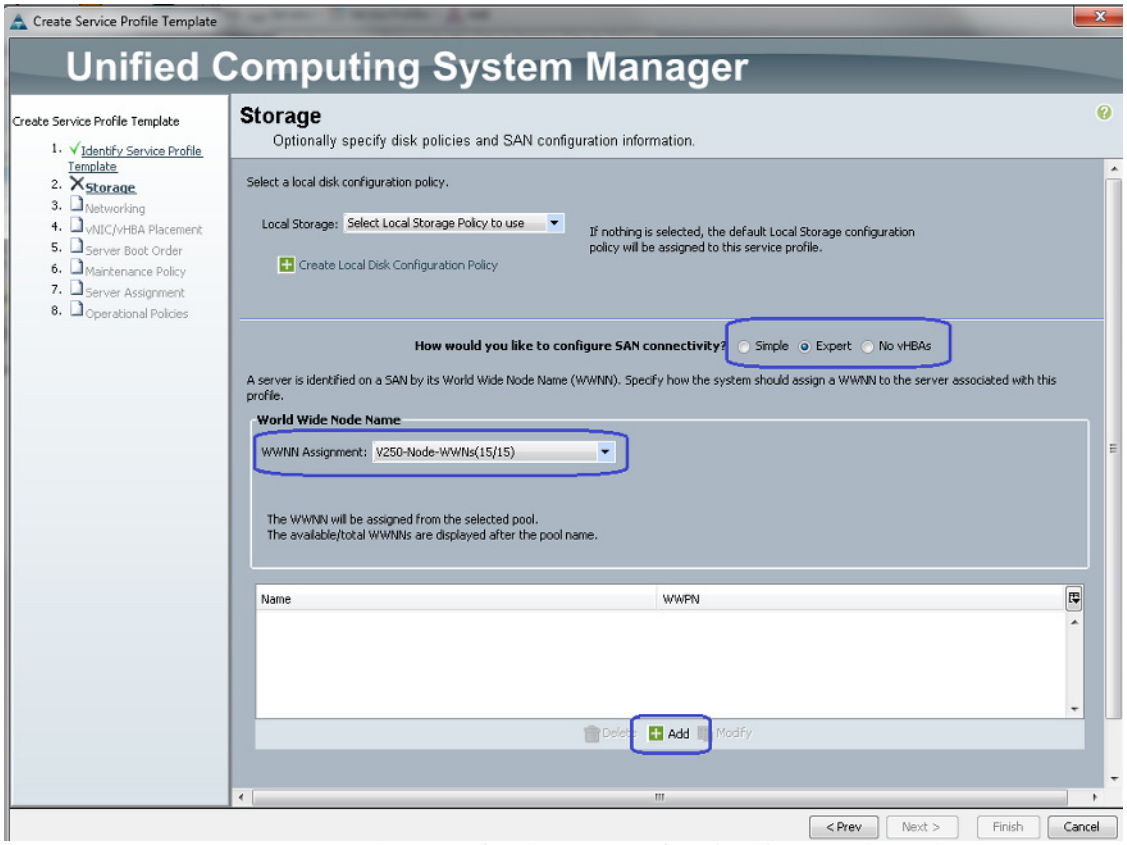

4. Enter the name for the vHBA (“fc0” in this example), select WWPN pool name created in the previous step from the “WWPN Assignment” drop down menu. Keep fabric ID “A”, select the VSAN for the vHBA by drop down names menu and change adapter policy to “VMware” as shown in Figure 72. Click **Ok**.

Figure 72 Creating vHBA on Fabric A

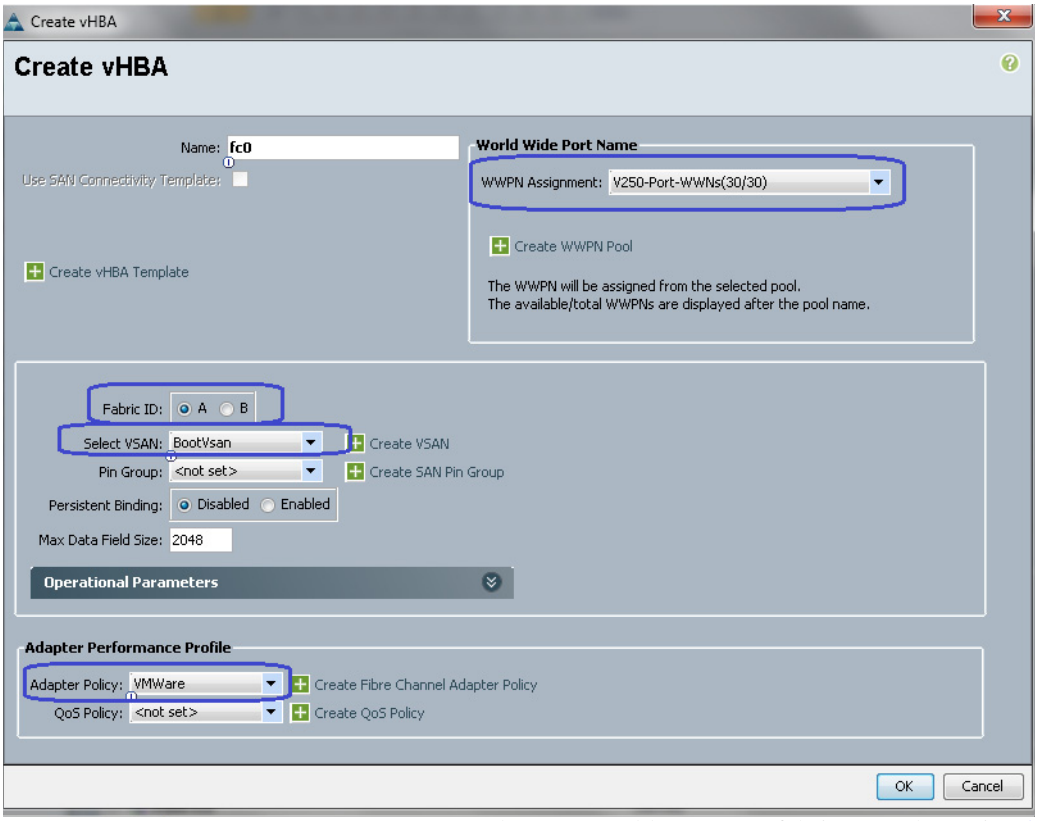

5. Repeat step 3, and step 4 to add vHBA on fabric B as shown in Figure 73. When second vHBA is added, click **Next** in the wizard.

Figure 73 Creating vHBA on Fabric B

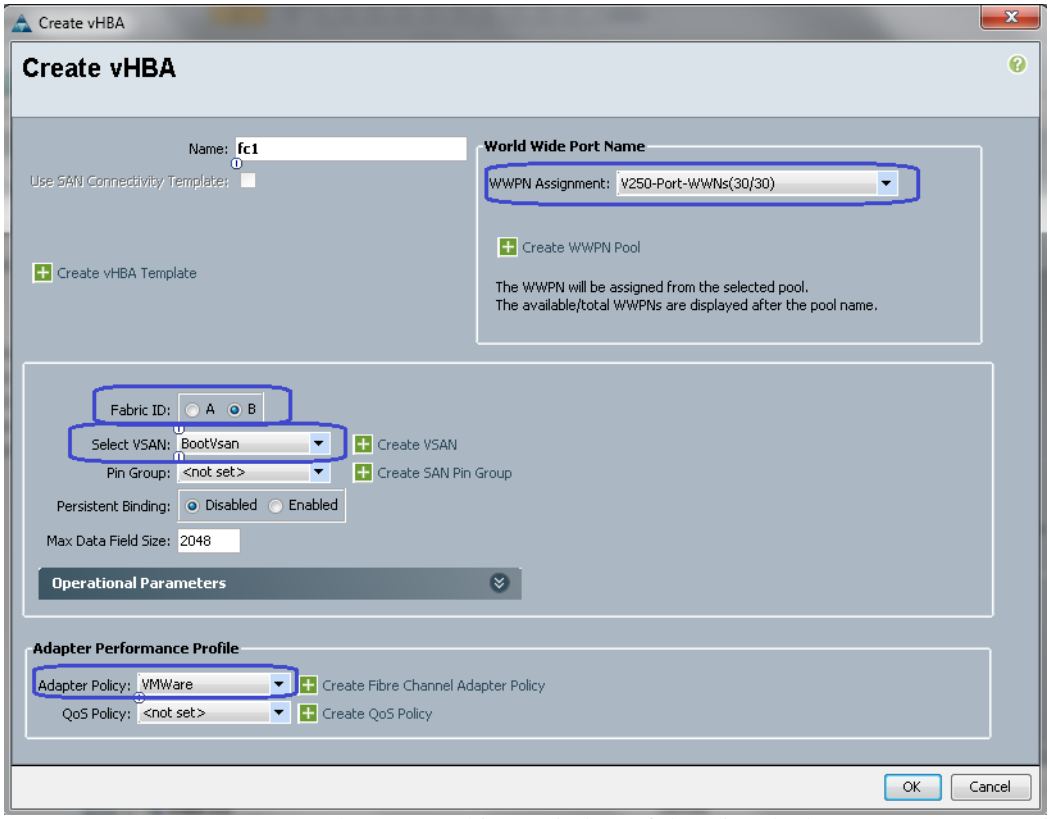

6. In the "Networking" window of the wizard, choose "Create a Specific Dynamic vNIC Connection Policy" from the drop-down list for "Dynamic vNIC Connection Policy". This will provide many options to configure various parameters of the dynamic vNIC connection policy. Enter "32" as number of dynamic vNICs (this is because we would have 25 VMs per hypervisor, 3 vmknic interfaces of hypervisor, and 4 additional dynamic VNICs for high-availability, if one of the hypervisor goes into maintenance mode and its load is distributed across 9 other hypervisors), select "VMware" for the adapter policy from the drop-down list, and keep the "Protection" radio button as **Protected** (this will provide fabric fail-over capability for dynamic vNICs). Click **Expert** radio button for LAN connectivity as shown in Figure 74, and click **Add** to add a (static) vNIC.

Figure 74 Networking Window of the Create Service Profile Template Wizard

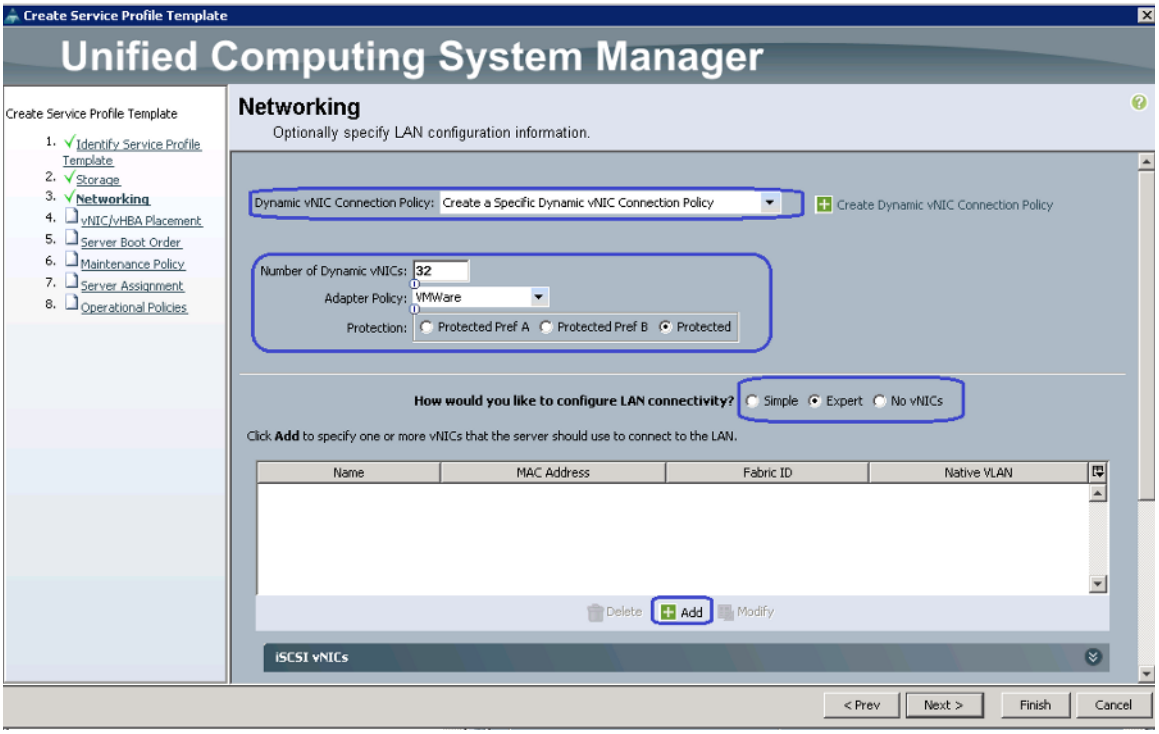

7. Enter the vHBA name (“eth0” in this example) in the name field, select name of the MAC pool created in the previous section from the “MAC Address Assignment” drop-down list, keep fabric ID as “A” and select “VMware” as “Adapter Policy” from the drop-down list as shown in Figure 75. As the static vNICs will be used only for ESXi host’s management and vCenter/ ESXi host communication through standard vSwitch on the hypervisor, you need to choose only vSphere Management (infra) VLAN for the allowed VLANs on the vNIC. Make sure that **Native VLAN** radio button is selected for this VLAN as shown in Figure 75. Click **Ok**.

Figure 75 Adding vNIC on Fabric A

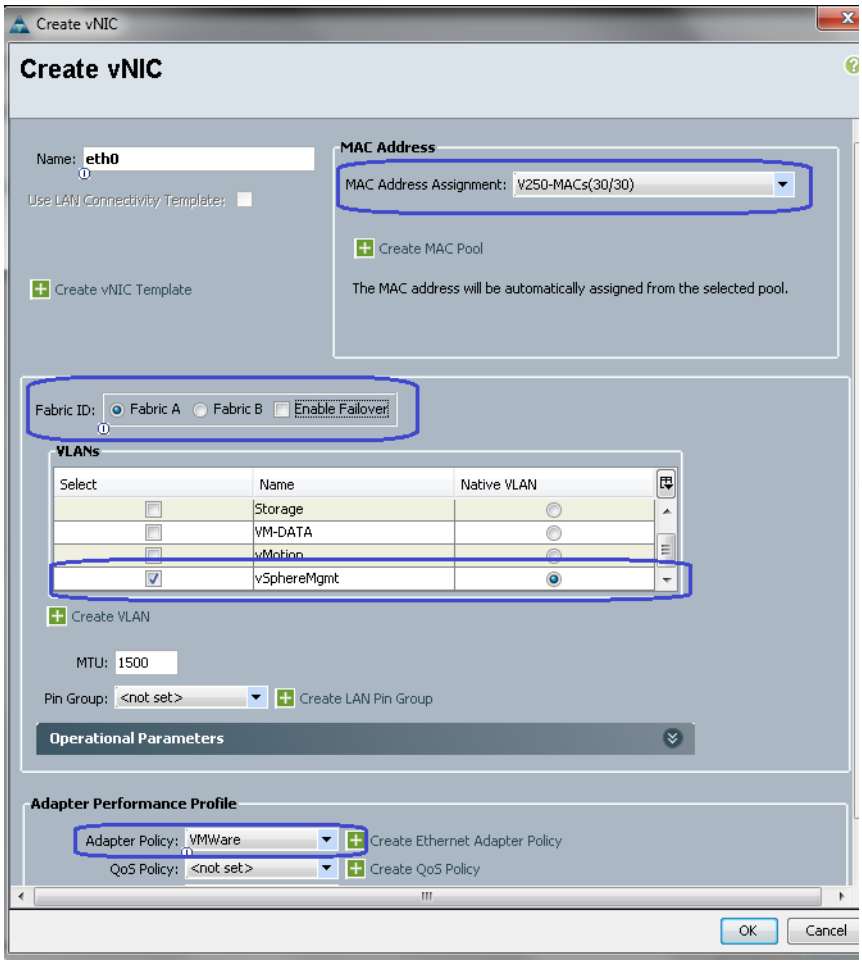

8. Repeat step 6 to add vNIC on Fabric B as shown in Figure 76. Click **Next**.

Figure 76 Adding vNIC on Fabric B

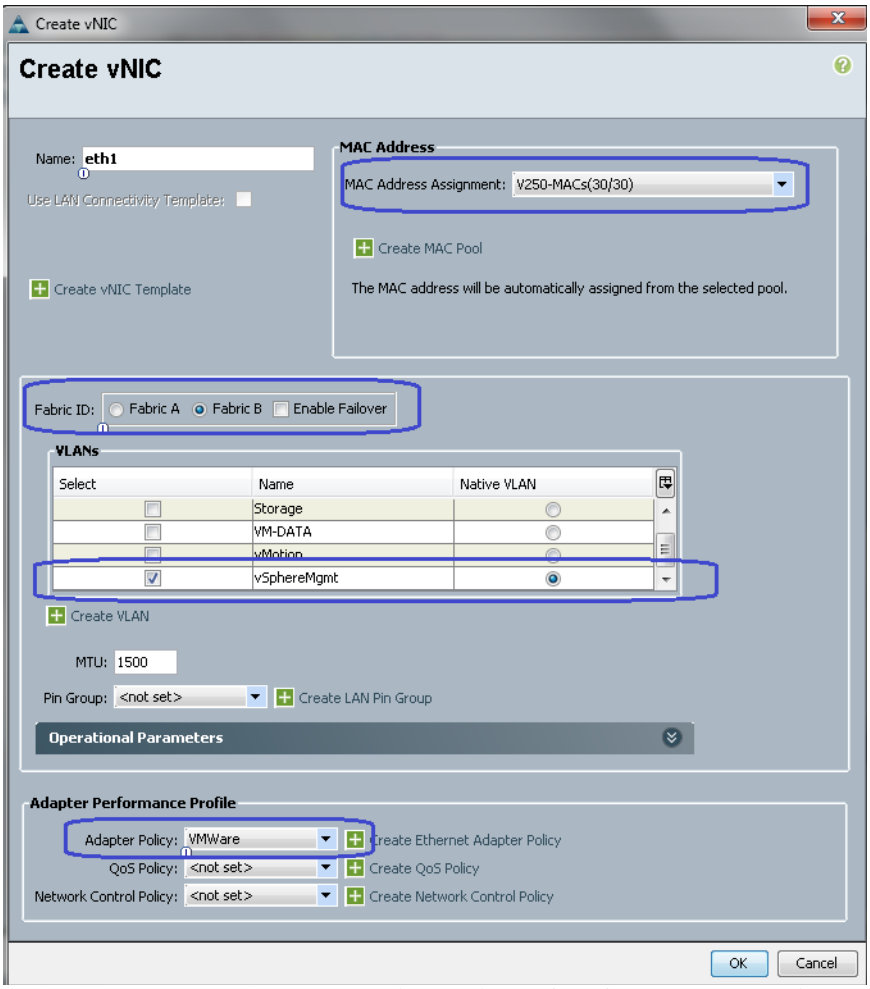

9. Do not change the settings in vNIC/vHBA Placement window of the wizard. Click **Next**.
10. In the "Server Boot Order" window of the wizard, click the **Create Boot Policy** link.
11. Enter the boot policy name in the "Name" field, check both check boxes as shown in Figure 77 and click **Add CD-ROM** radio button under "Local Devices" to choose it as 1st order for boot. Click **Add SAN Boot** radio button under "vHBAs" on the left pane of the "Create Boot Policy" window.

Figure 77 Create Boot Policy Window

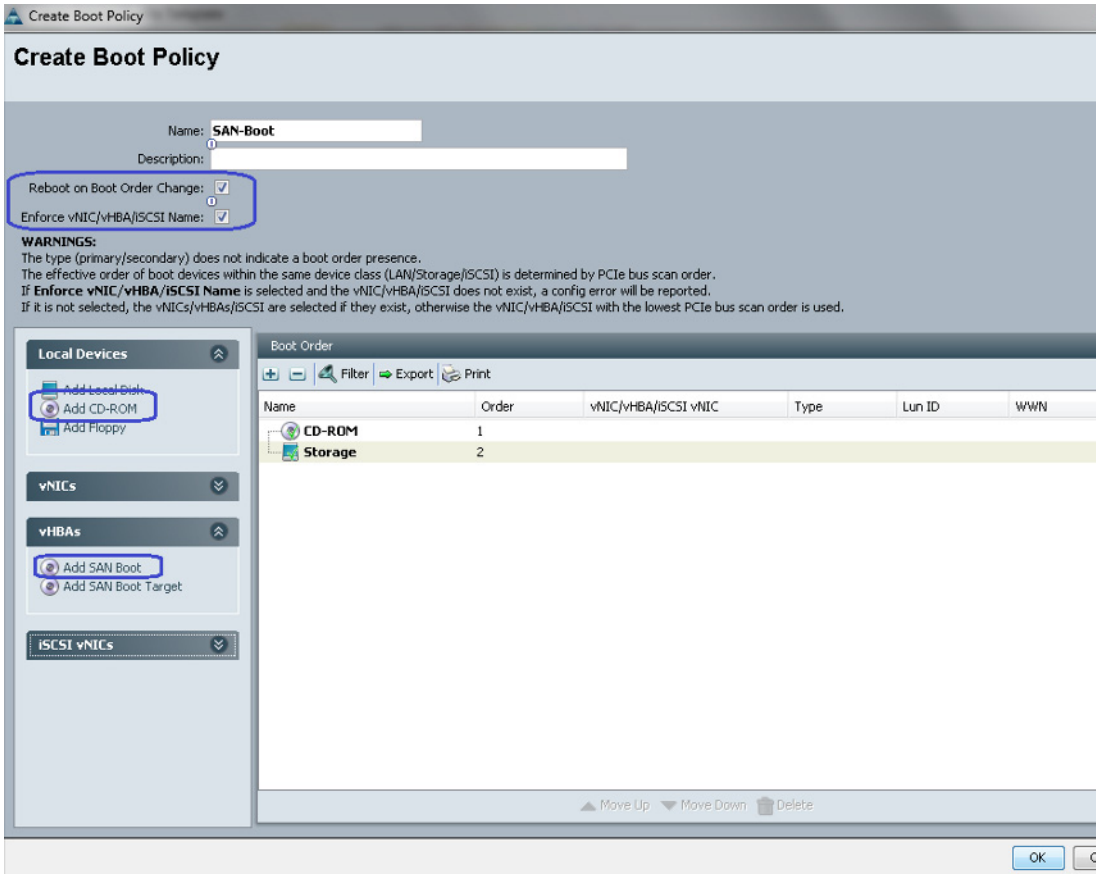

12. Provide name of the vHBA on SAN fabric A and select it as “Primary” type. Click **Ok**.

Figure 78 Adding SAN Boot

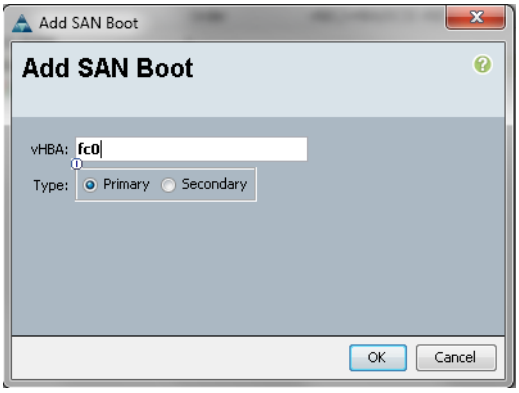

13. Click the **Add SAN Boot Target** link.

Figure 79 Adding SAN Boot Target

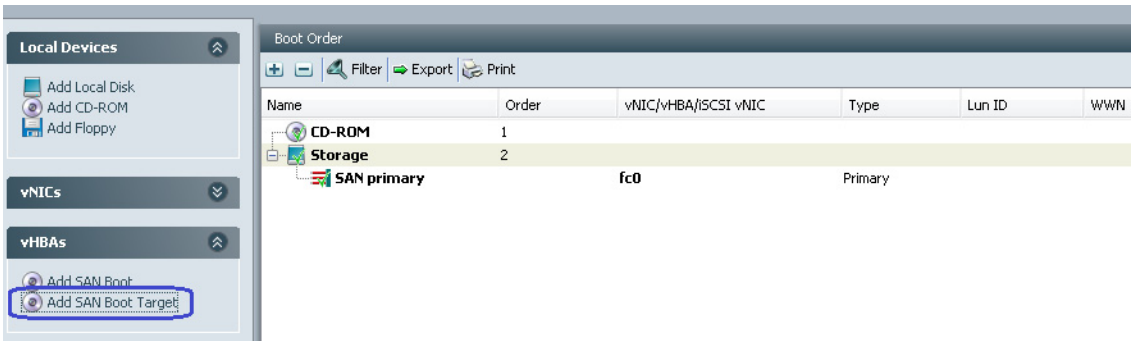

14. Enter 0 in the "Boot Target LUN" field and enter WWPN of the SP-A of the VNX5500 HBA in Boot Target WWPN field. Keep the "Type" as "Primary" and click **Ok** as shown in Figure 80.

Figure 80 Entering Details in Add SAN Boot Target Window

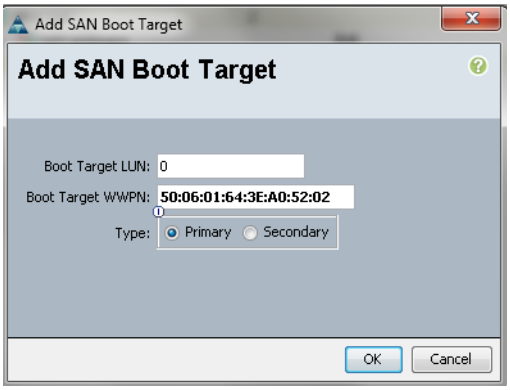

15. Repeat step 14 for secondary Boot Target on the SP-B of VNX5500. Repeat steps 12 to 14 for the secondary SAN Boot on fabric B. After configuring both primary and secondary SAN boot and boot targets the Boot Policy will look as shown in Figure 81. Click **Ok** to save the boot policy.

Figure 81 Successfully Created Boot Policy Window

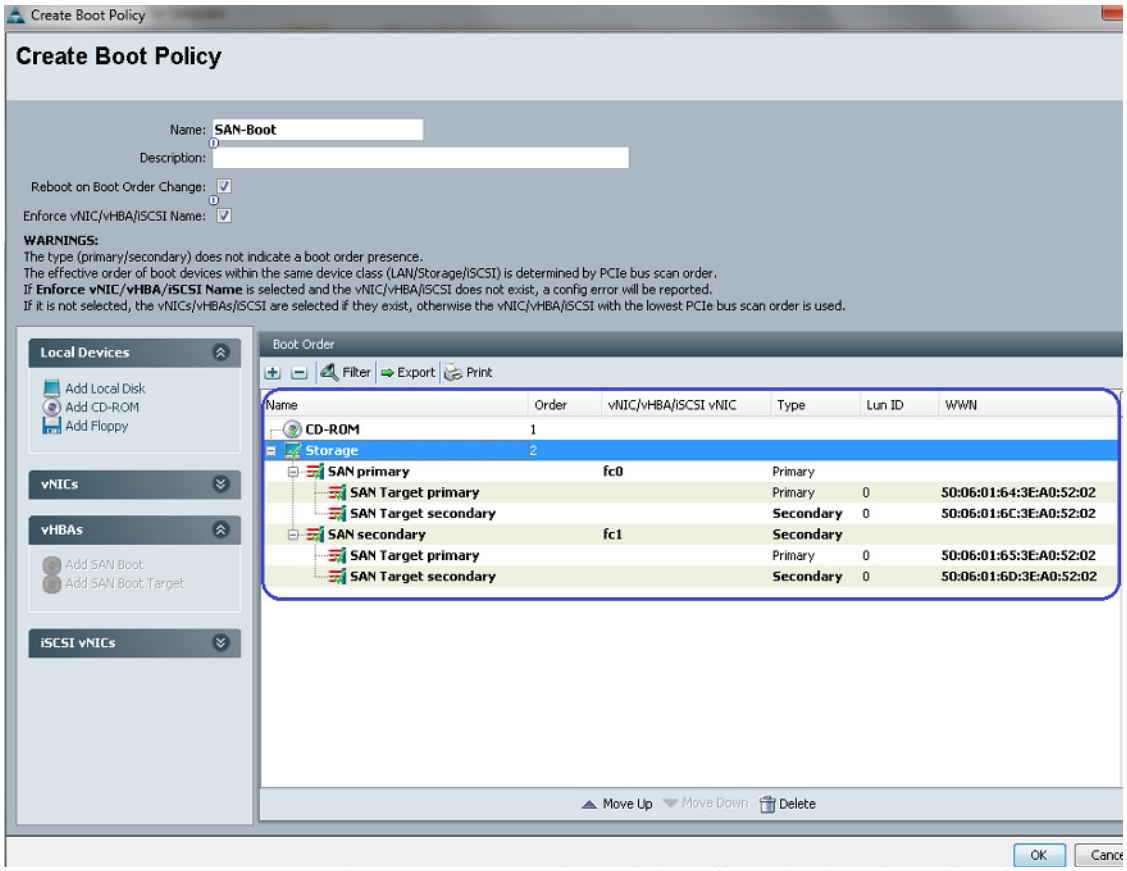

16. In the "Server Boot Order" window, from the drop-down list of the "Boot Policy" of the "Create Service Profile Template" wizard, select "SAN-Boot" a newly created boot order. click **Next**.
17. Keep the default setting in the "Management Policy" window of the "Create Service Profile Template" wizard and click **Next**.
18. Keep the default setting "Assign Later" for "Pool Assignment" in the "Server Assignment" window of the "Create Service Profile Template" wizard and click **Next**.
19. Click **Finish** in the "Operational Policies" window of the "Create Service Profile Template" wizard.
20. You will see a success notification of creation of service profile template as shown in Figure 82.

Figure 82 Success Notification of Created Service Profile Template

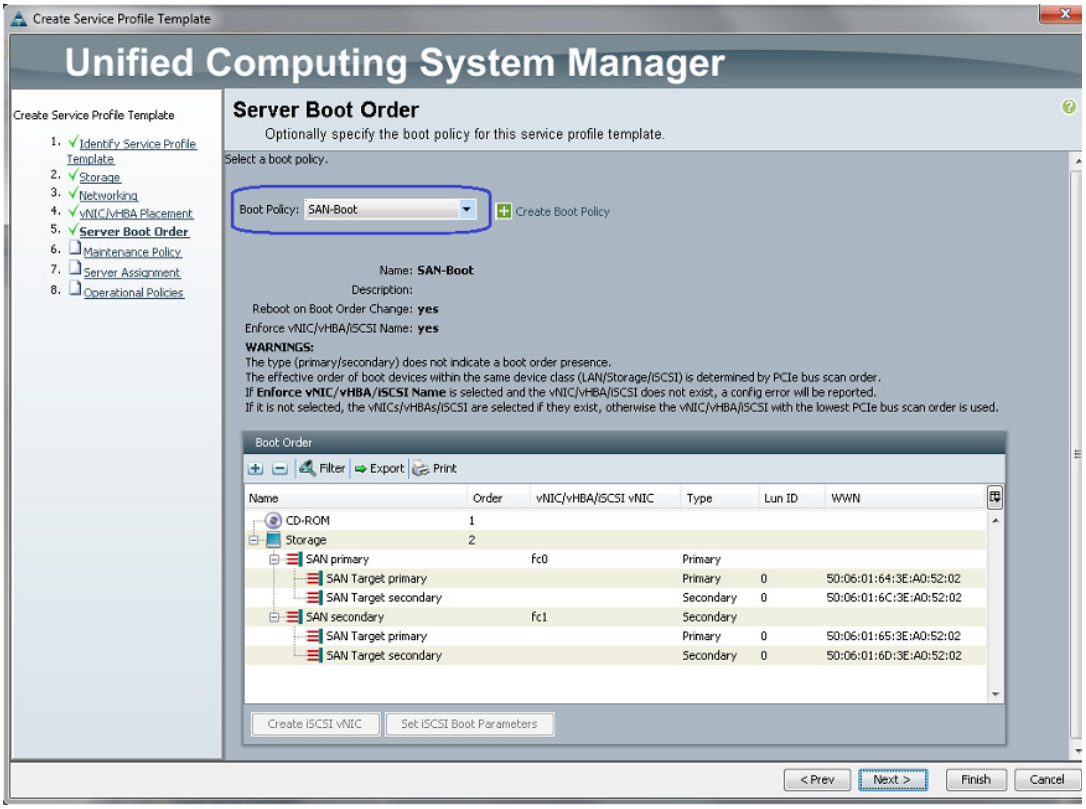

Create service profile instances from the template

In this section we will create ten service profile instances from the template created in the previous section. Follow these steps to create service profile instances:

1. In the Cisco UCS Manager window, select the **Servers** tab, expand "Service Profiles", right-click on the "root" and click **Create Service Profiles From Template** as shown in Figure 83.

Figure 83 Creating Service Profiles from Template

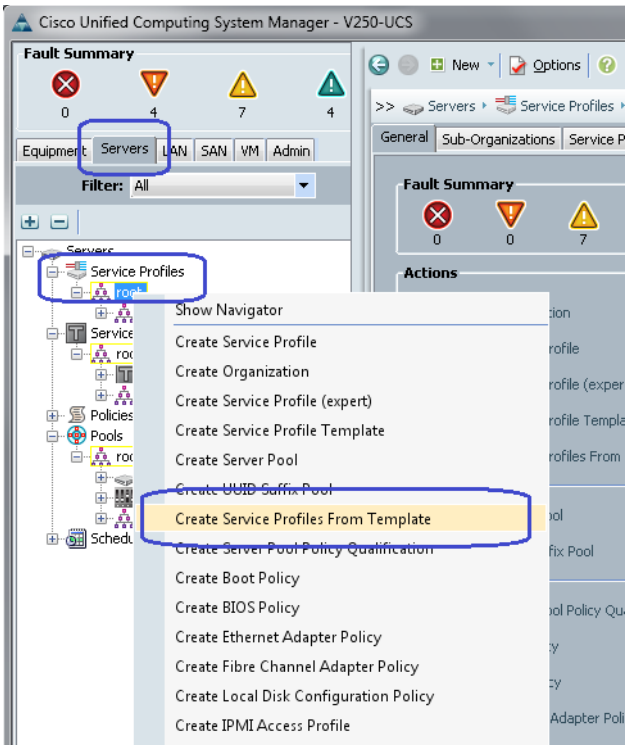

2. In the Create Service Profile from Template window, enter the name in the “Naming Prefix” field, number of Service Profiles as “10”, and select the “Service Profile Template” as “Service Profile Template V250-Host-Template” (created in the previous section) drop-down list. Click **Ok**.

Figure 84 Entering Details for Creating Service Profile instance

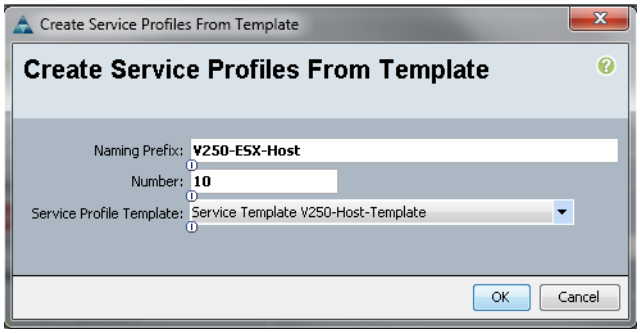

3. A pop-up window appears showing successful completion of creating service profile instances.
4. You will see 10 Service Profiles instantiated under “root” with the “Overall Status” showing “Unassociated” as in Figure 85. The window will show a warning message “Service profile is not associated”.

Figure 85 Created Service Profile Instances

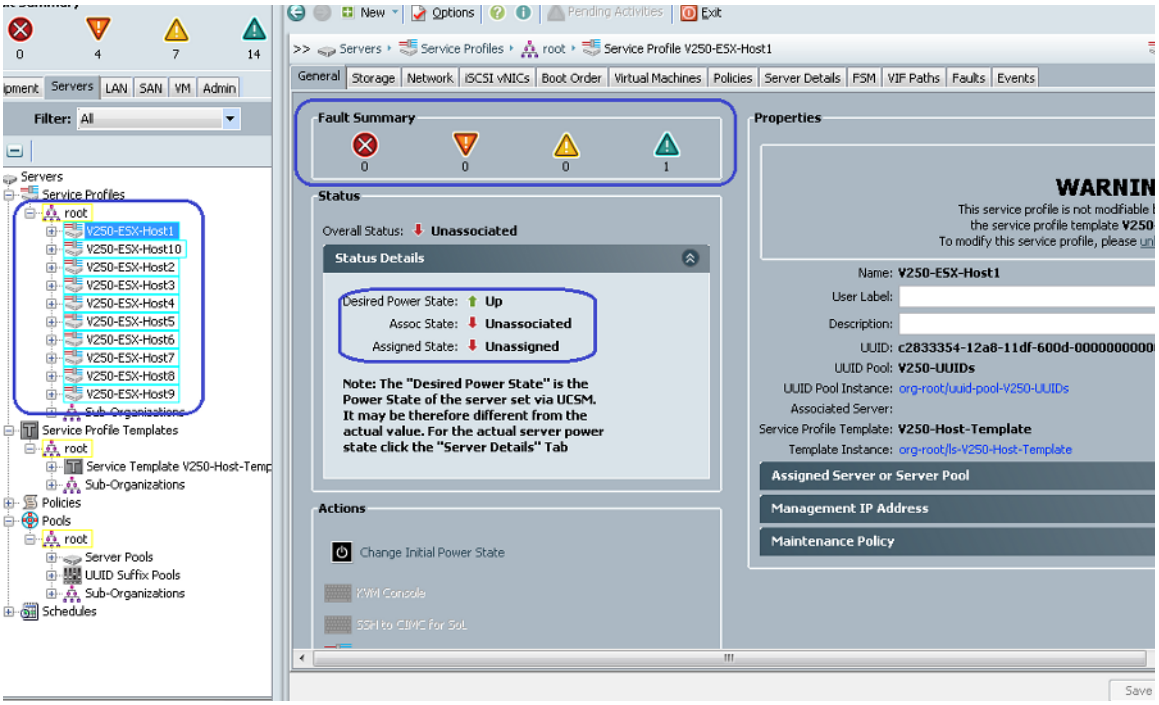

Associate Service Profiles

As mentioned before, Service Profile is a logical representation of a server. When a Service Profile is associated with available physical server, the server assumes the role described by the Service Profile and corresponding server is booted. we need to associate Service Profile instances created in previous section to the Cisco UCS B200 M3 Blade Servers available. Follow these steps to associate Service Profiles:

1. Select the first Service Profile instance out of the ten Service profiles created in previous section, and click the **Change Service Profile Association** link on the right pane of the side Cisco UCS Manager window as shown Figure 86.

Figure 86 Changing Service Profile Association

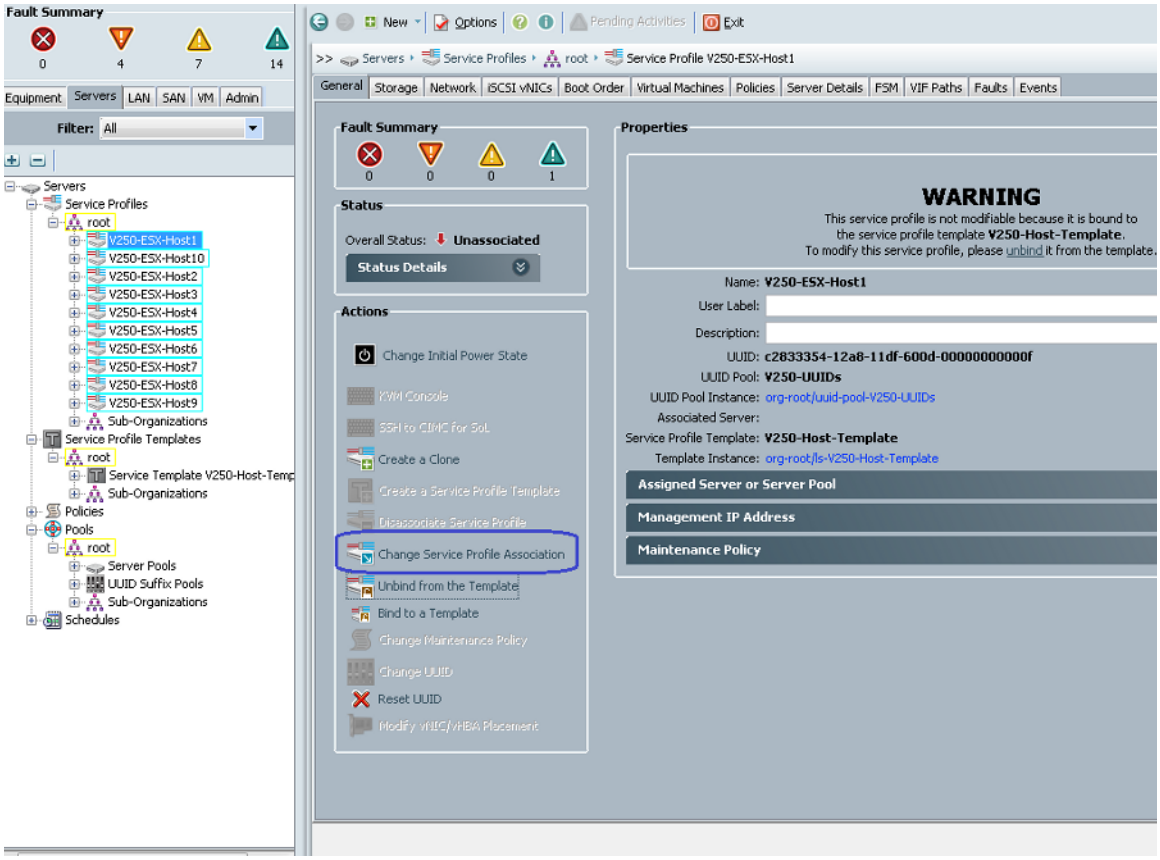

2. Select Cisco UCS B200 M3 Blade Server 1/1 and click **Ok**.

Figure 87 Associating Service Profile

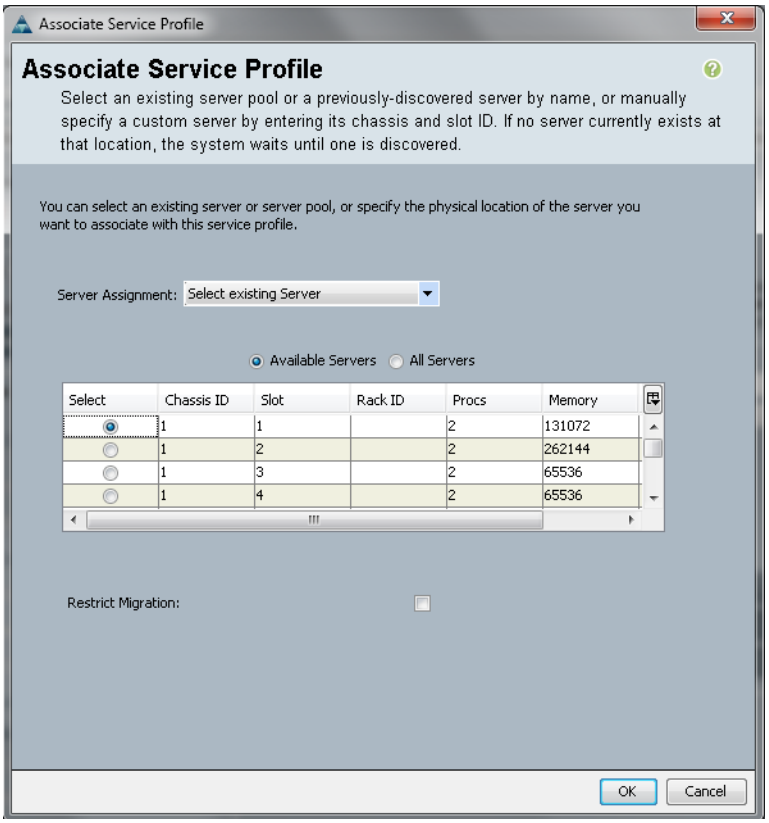

3. You will get an acknowledgement on Service Profile association process as shown in Figure 88.

Figure 88 Service Profile Association Process in Progress

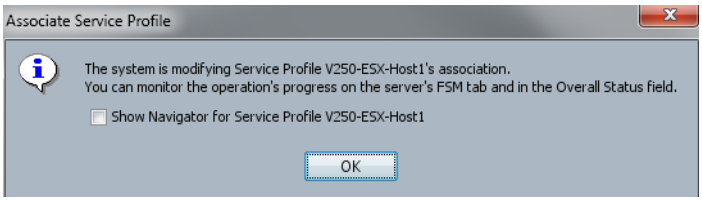

4. After the Service Profile association is complete, the "Overall Status" of the server should show "Ok" and there should be no faults under the server as shown in Figure 89.

Figure 89 Overall Status of the Server

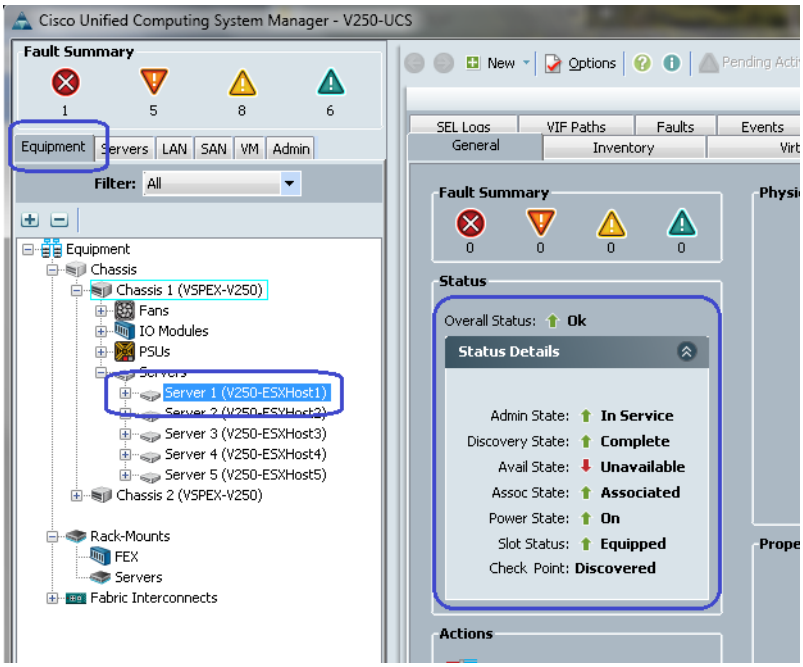

This completes the UCS server configuration. We need to configure the Cisco UCS Manager/ vCenter integration in Cisco UCS Manager and Cisco VMFEX architecture after the vSphere infrastructure is setup.

Preparing and Configuring Storage Array

To configure the EMC VNX5500 storage array follow these steps:

1. Configure end-to-end SAN boot infrastructure for ESXi hosts.
2. Create a data store for virtual machines operating systems and data, create performance pool and LUNs.
3. Configure NFS share and assign host access privileges.
4. Configure port-channel (aggregation) and jumboframe.

Configure SAN Boot Infrastructure

This section explains how to configure end-to-end SAN Boot Infrastructure for the Cisco UCS B200 M3 Blade Servers. Most of the configuration is on the EMC VNX5500, but part of it is on Cisco Nexus 5548UP switches and Cisco UCS Manager. we have the following tasks completed already:

1. VSAN configuration and FC port configuration on the Cisco Nexus 5548UP switches and Cisco UCS FIs.
2. WWPN and WWNN assignments to the proposed ESXi servers.

Follow these steps to configure SAN Boot Infrastructure:

1. In the Cisco UCS Manager window, click the **SAN** tab, expand “Pools”, under “root” click **WWPN Pools** to select the WWPN pool created for the B200 M3 servers' Service Profiles. Click the **Initiators** tab on the right pane of the Cisco UCS Manager window as shown Figure 90. The

"Assigned To" column on the right pane of the Cisco UCS Manager window provides the WWPN assignment values. This can be referred while creating the zones on the Cisco Nexus 5548UP switches.

Figure 90 Assigned WWPN Values

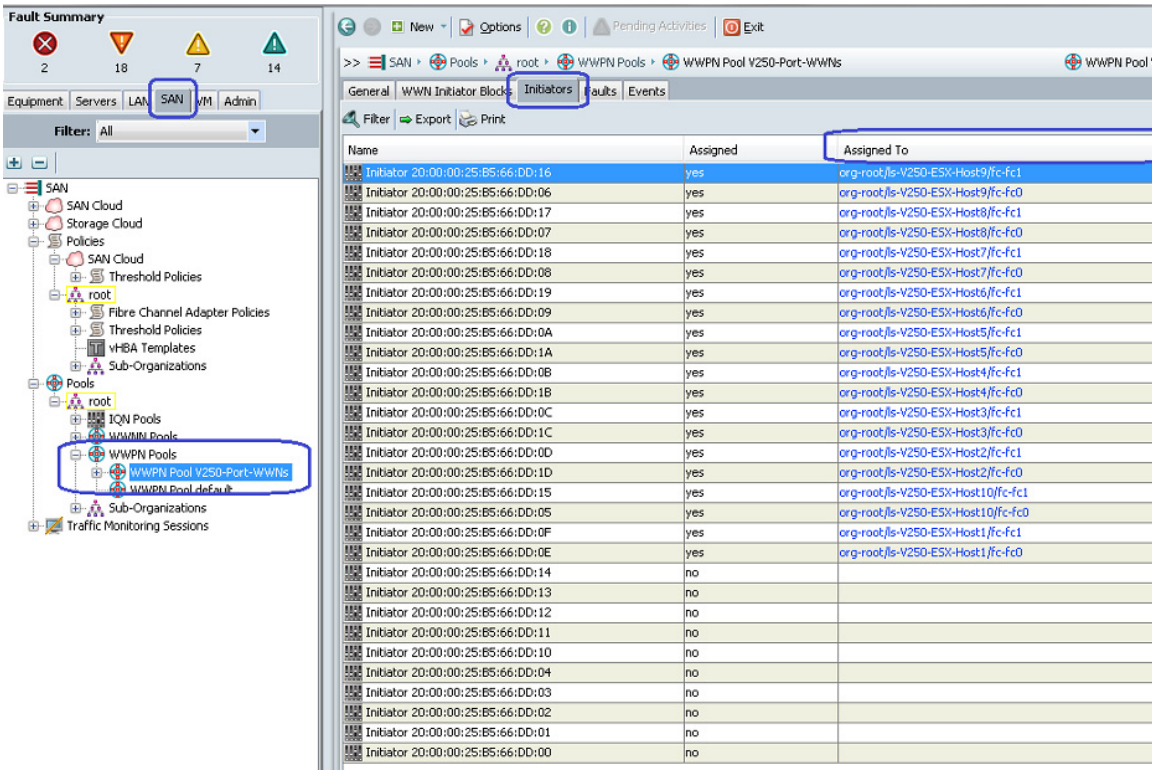

2. Login to the Nexus 5548UP switch A and configure a zoneset for SAN fabric A. Create 10 zones, one for each ESXi host, containing WPN of SP-A and SP-B of VNX5500 and WWPN of the vHBA on fabric A of the ESXi server. WWPN list in the step 1 will be helpful to verify. Entire zoneset configuration will look as shown in Figure 91. Activate the zoneset in the storage VSAN.

Figure 91 Creating Zones for Each of the ESX Hosts

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
UCS-N5k-FabA(config)# zoneset name V250-Fabric-A vsan 10
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost1-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:0e
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost2-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:1d
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost3-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:1c
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost4-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:1b
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost5-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:1a
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost6-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:09
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost7-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:08
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost8-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:07
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost9-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:06
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# zone name V250-ESXHost10-fc0
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 20:00:00:25:b5:66:dd:05
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwwn 50:06:01:6c:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# exit
UCS-N5k-FabA(config-zoneset)# exit
UCS-N5k-FabA(config)# zoneset activate name V250-Fabric-A vsan 10
Zoneset activation initiated. check zone status
UCS-N5k-FabA(config)#
```

3. Validate the successful activation of zoneset by the command **show zoneset brief** as shown below.

Figure 92 Validating the Activation of Zoneset on Fabric A

```
UCS-N5k-FabA# show zoneset brief
zoneset name V250-Fabric-A vsan 10
  zone V250-ESXHost1-fc0
  zone V250-ESXHost2-fc0
  zone V250-ESXHost3-fc0
  zone V250-ESXHost4-fc0
  zone V250-ESXHost5-fc0
  zone V250-ESXHost6-fc0
  zone V250-ESXHost7-fc0
  zone V250-ESXHost8-fc0
  zone V250-ESXHost9-fc0
  zone V250-ESXHost10-fc0
UCS-N5k-FabA#
```

4. Similarly, on the Nexus 5548UP switch B, create zoneset for fabric B and include vHBAs on fabric B on the servers. Validation of zoneset on fabric B is shown in Figure 92.

Figure 93 Validating the Activation of Zoneset on Fabric B

```
UCS-N5K-FabB# show zoneset brief
zoneset name V250-Fabric-B vsan 10
  zone V250-ESXHost1-fc1
  zone V250-ESXHost2-fc1
  zone V250-ESXHost3-fc1
  zone V250-ESXHost4-fc1
  zone V250-ESXHost5-fc1
  zone V250-ESXHost6-fc1
  zone V250-ESXHost7-fc1
  zone V250-ESXHost8-fc1
  zone V250-ESXHost9-fc1
  zone V250-ESXHost10-fc1
UCS-N5K-FabB#
```

5. To further validate the zoneset configuration across entire SAN fabric, SSH to UCS FI-A, issue **connect nxos** command, and run the command **show npv flogi-table**. It should list all the ten fLogi sessions, one from each vHBA on fabric A in storage VSAN as shown in Figure 94.

Figure 94 Validating the Created Zoneset Across SAN Fabric

```
V250-UCS-A# connect nxos
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
V250-UCS-A(nxos)# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID            PORT NAME              NODE NAME        INTERFAC
E
--------------------------------------------------------------------------------
vfc769    10   0x5c0002 20:00:00:25:b5:66:dd:0e 20:00:00:25:b5:60:0d:0e fc2/1
vfc823    10   0x5c0003 20:00:00:25:b5:66:dd:1d 20:00:00:25:b5:60:0d:0d fc2/2
vfc877    10   0x5c0004 20:00:00:25:b5:66:dd:1c 20:00:00:25:b5:60:0d:0c fc2/1
vfc931    10   0x5c0005 20:00:00:25:b5:66:dd:1b 20:00:00:25:b5:60:0d:0b fc2/2
vfc1011   10   0x5c0006 20:00:00:25:b5:66:dd:06 20:00:00:25:b5:60:0d:06 fc2/1
vfc1065   10   0x5c0007 20:00:00:25:b5:66:dd:07 20:00:00:25:b5:60:0d:07 fc2/1
vfc1119   10   0x5c0009 20:00:00:25:b5:66:dd:08 20:00:00:25:b5:60:0d:08 fc2/2
vfc1173   10   0x5c0008 20:00:00:25:b5:66:dd:09 20:00:00:25:b5:60:0d:09 fc2/2
vfc1227   10   0x5c000a 20:00:00:25:b5:66:dd:05 20:00:00:25:b5:60:0d:05 fc2/1
vfc1281   10   0x5c000b 20:00:00:25:b5:66:dd:1a 20:00:00:25:b5:60:0d:0a fc2/2

Total number of flogi = 10.

V250-UCS-A(nxos)#
```

6. Similarly, the **show flogi database** command on Nexus 5548UP switch should show 14 flogi sessions:
   - 10 from B200 M3 vHBAs
   - 2 from FI-A's FC ports
   - 2 from VNX5500 storage array's SP-A and SP-B FC ports

   Similarly, verify the FLogI entries on SAN fabric B.

Figure 95 Total Number of flogi Sessions

```
UCS-N5k-FabA# show flogi database
--------------------------------------------------------------------------------
INTERFACE        VSAN    FCID           PORT NAME               NODE NAME
--------------------------------------------------------------------------------
fc1/29           10    0x5c0000  20:41:00:0d:ec:f7:04:00 20:0a:00:0d:ec:f7:04:01
fc1/29           10    0x5c0002  20:00:00:25:b5:66:dd:0e 20:00:00:25:b5:60:0d:0e
fc1/29           10    0x5c0004  20:00:00:25:b5:66:dd:1c 20:00:00:25:b5:60:0d:0c
fc1/29           10    0x5c0006  20:00:00:25:b5:66:dd:06 20:00:00:25:b5:60:0d:06
fc1/29           10    0x5c0007  20:00:00:25:b5:66:dd:07 20:00:00:25:b5:60:0d:07
fc1/29           10    0x5c000a  20:00:00:25:b5:66:dd:05 20:00:00:25:b5:60:0d:05
fc1/30           10    0x5c0001  20:42:00:0d:ec:f7:04:00 20:0a:00:0d:ec:f7:04:01
fc1/30           10    0x5c0003  20:00:00:25:b5:66:dd:1d 20:00:00:25:b5:60:0d:0d
fc1/30           10    0x5c0005  20:00:00:25:b5:66:dd:1b 20:00:00:25:b5:60:0d:0b
fc1/30           10    0x5c0008  20:00:00:25:b5:66:dd:09 20:00:00:25:b5:60:0d:09
fc1/30           10    0x5c0009  20:00:00:25:b5:66:dd:08 20:00:00:25:b5:60:0d:08
fc1/30           10    0x5c000b  20:00:00:25:b5:66:dd:1a 20:00:00:25:b5:60:0d:0a
fc1/31           10    0x5c00ef  50:06:01:64:3e:a0:52:02 50:06:01:60:be:a0:52:02
fc1/32           10    0x5c01ef  50:06:01:6c:3e:a0:52:02 50:06:01:60:be:a0:52:02

Total number of flogi = 14.

UCS-N5k-FabA#
```

7. After the end-to-end FC SAN fabric connectivity is verified, log in to the EMC's VNX5500 Unisphere. To configure SAN storage, select VNX5500 array in the Unisphere window. Click the **Storage** tab from the menu bar, and click **Storage Configuration > Storage Pools**. Click **Create** to create a new storage pool as shown in Figure 96.

Figure 96 Creating Storage Pools in EMC VNX5500 Unisphere

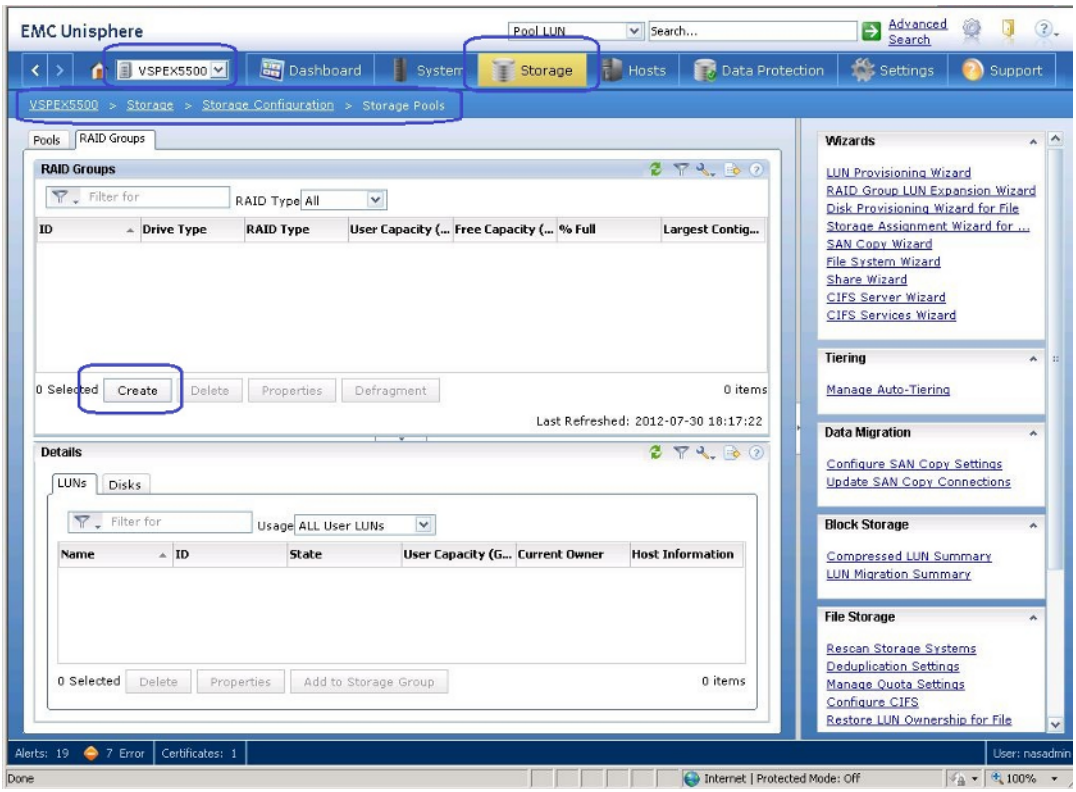

8. Click **RAID Group** radio button for Storage Pool Type, select RAID5 from the drop-down list for the RAID Type and click **Manual** radio button in the Disks area. Click **Select…** as shown Figure 97.

Figure 97 Entering Details for Creating Storage Pool

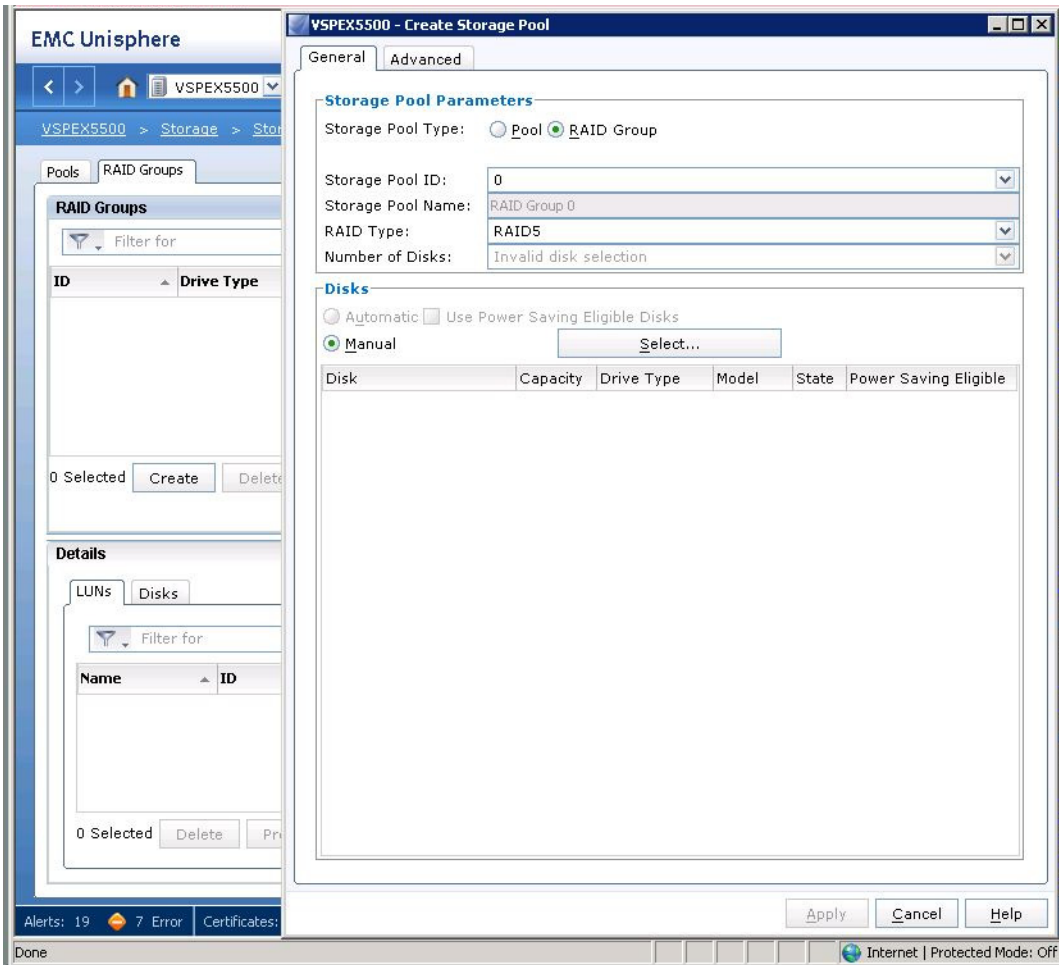

9. Select 3 disks for ESXi 5 hypervisor boot storage as shown in Figure 98. Click **Ok**.

Figure 98 Selecting Disks for ESXi 5 Hypervisor

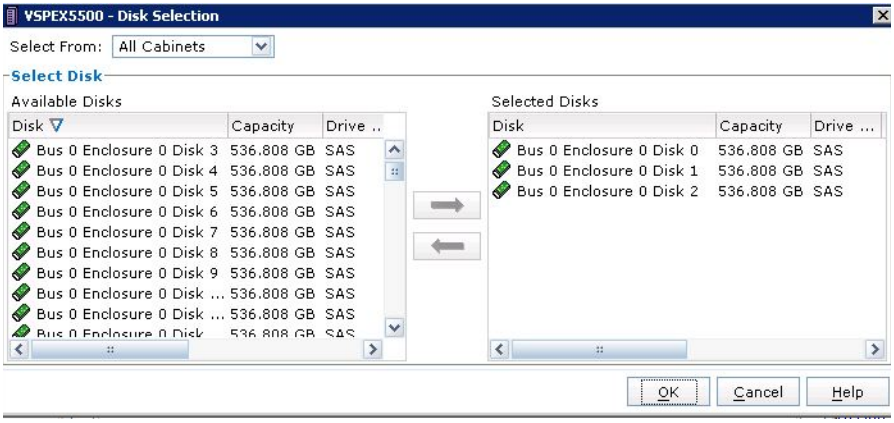

10. Click **Yes** in the pop-up window to initiate RAID group operation as shown in Figure 99.

Figure 99 Confirmation Window to Initiate RAID Group Operation

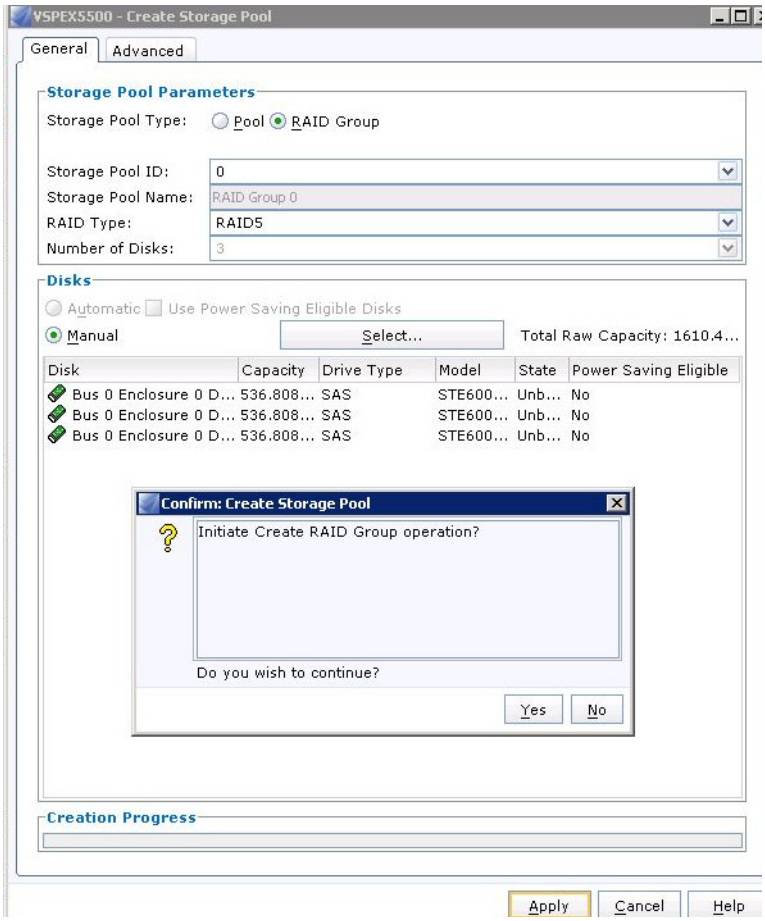

11. You will see a success notification as shown in Figure 100 upon completion of RAID group creation.

Figure 100 Success Notification of RAID Group Creation

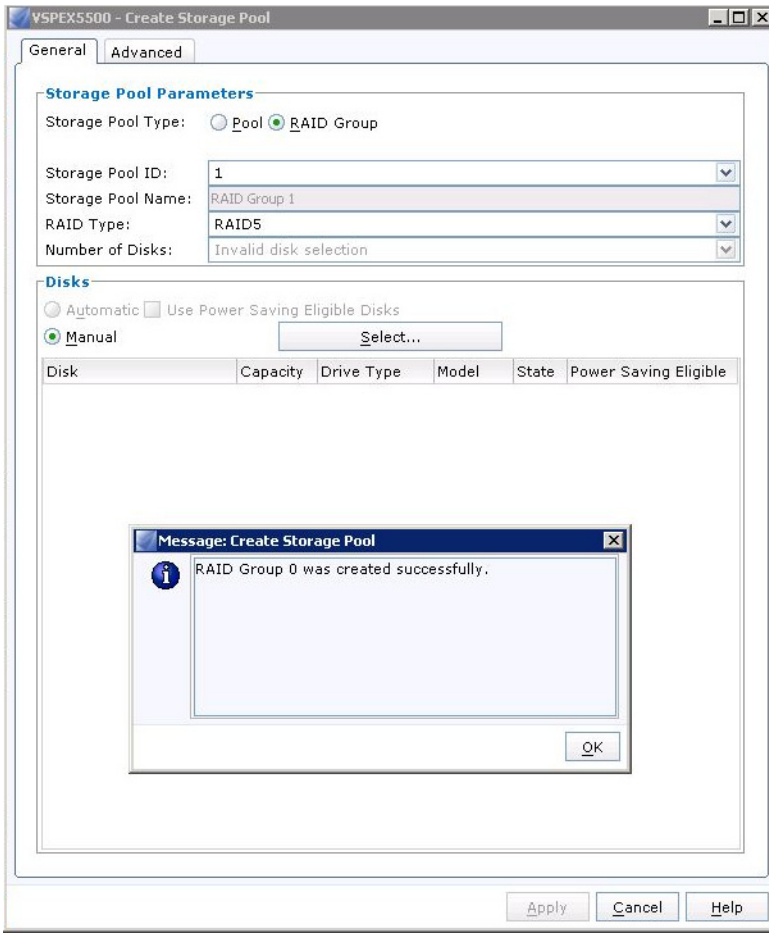

12. From the newly created RAID group, right-click and click **Create LUN** as shown in Figure 101.

Figure 101 Creating LUN in EMC Unisphere

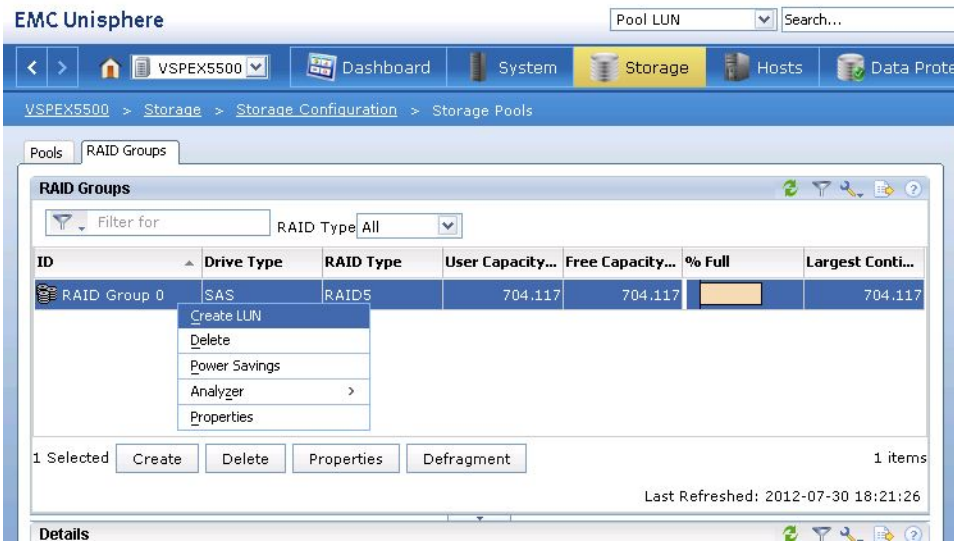

13. Create ten LUNs with 50 GB capacity each. Make sure that you click **RAID Group** radio button for "Storage Pool Type".

Figure 102 Entering Details to Create LUNs

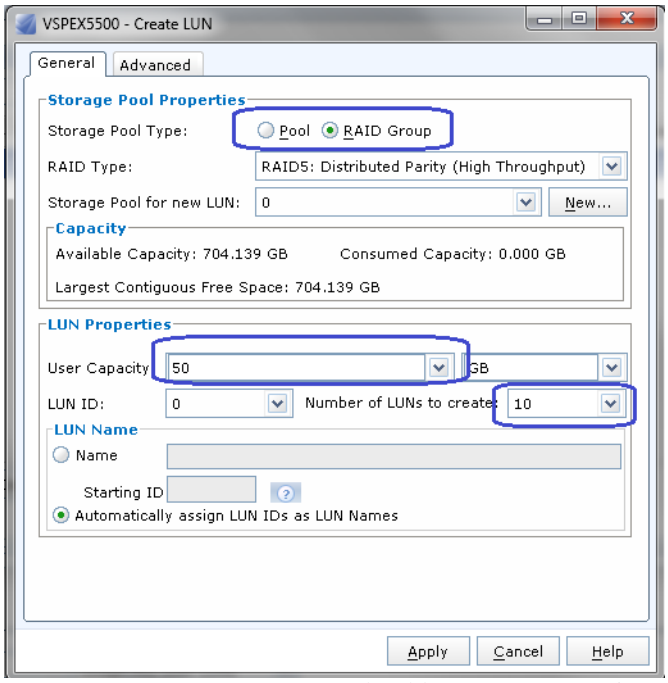

14. You should see a LUN creation notification as shown in Figure 103.

Figure 103 Window Showing LUN Creation in Progress

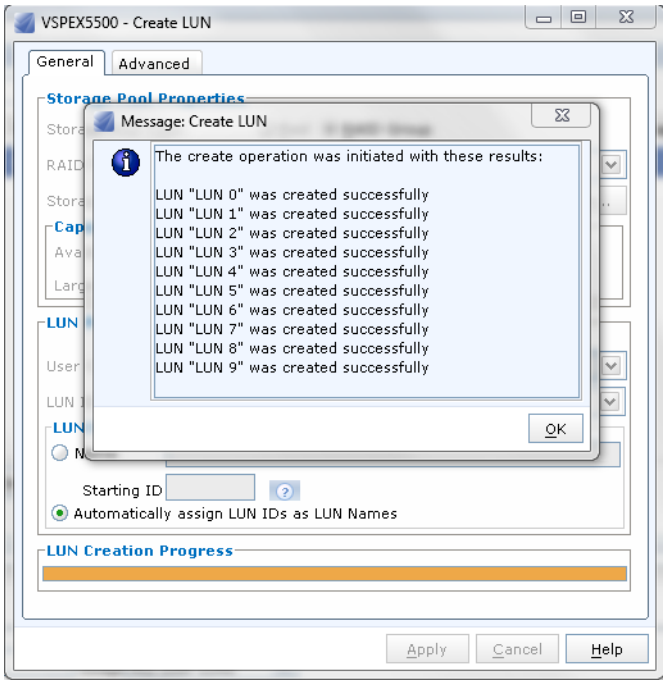

15. After the storage LUNs are created, we need to add hosts to the host list. Click **Hosts** tab from the EMC Unisphere menu bar, click the **Connectivity Status** link on the right pane of the EMC Unisphere window in the Host management Area as shown in Figure 104.

Figure 104 Adding Hosts to the Host List in the EMC Unisphere Window

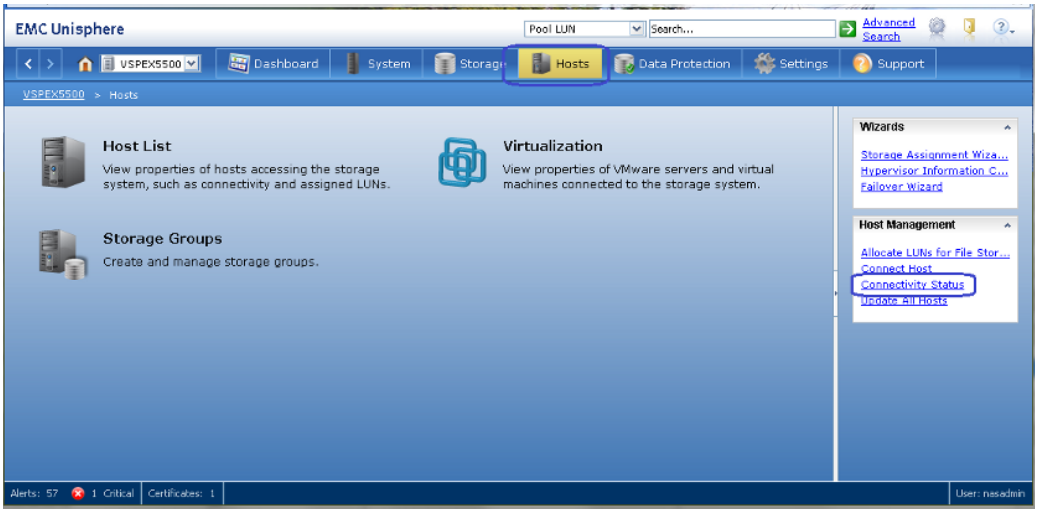

16. Select WWPN of the first ESX host and click **Edit** as shown in Figure 105. The WWPN Initiator list in Cisco UCS Manager shown in the step 1 can be used to verify.

Figure 105 Host Initiator Connectivity Status

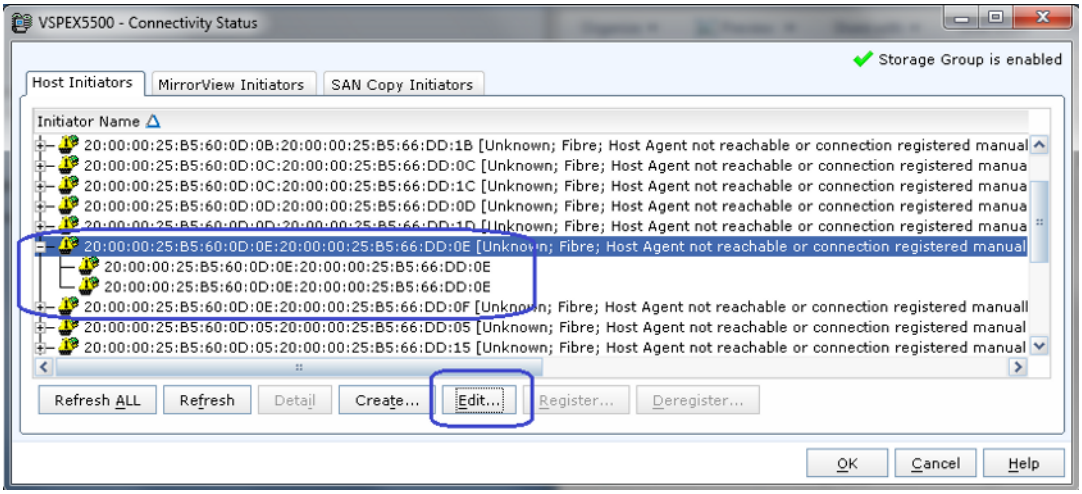

17. Click **New Host** radio button in the Edit Initiators window and provide ESXi hostname and IP address in the respective fields and Initiator information as shown in Figure 106.

Figure 106 Entering Details for Editing Initiators

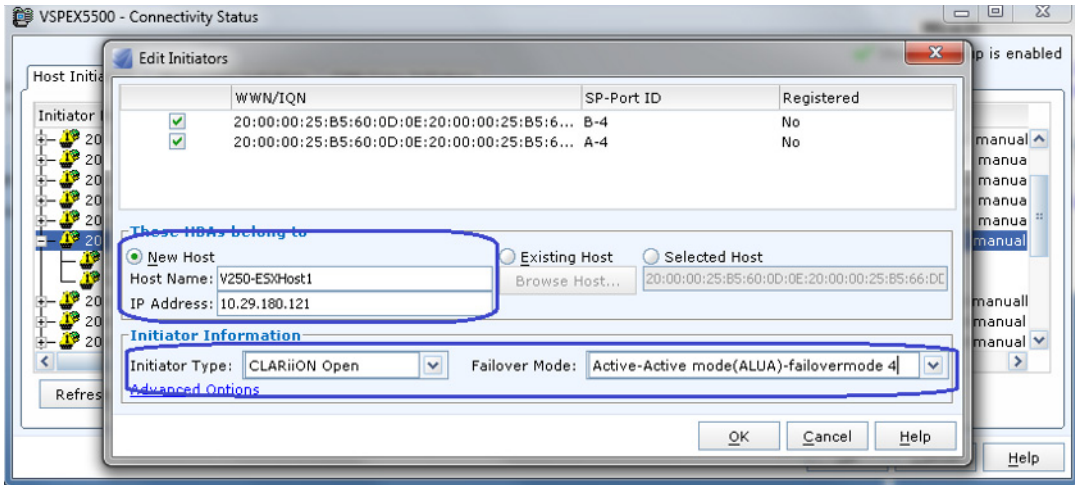

18. Click **Yes** in the confirmation popup window as shown in Figure 107.

Figure 107 Confirmation to Register Initiator with Existing Host Information

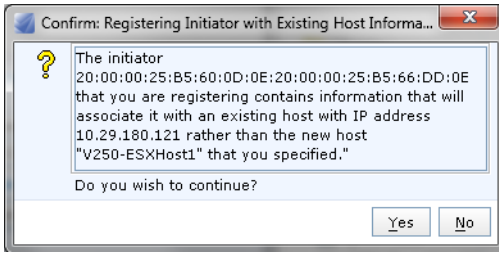

19. In the Host initiator window, click **Refresh All**, select the WWPN on fabric B of the same host, and click **Edit** as shown below.

Figure 108 Editing the Host Initiators

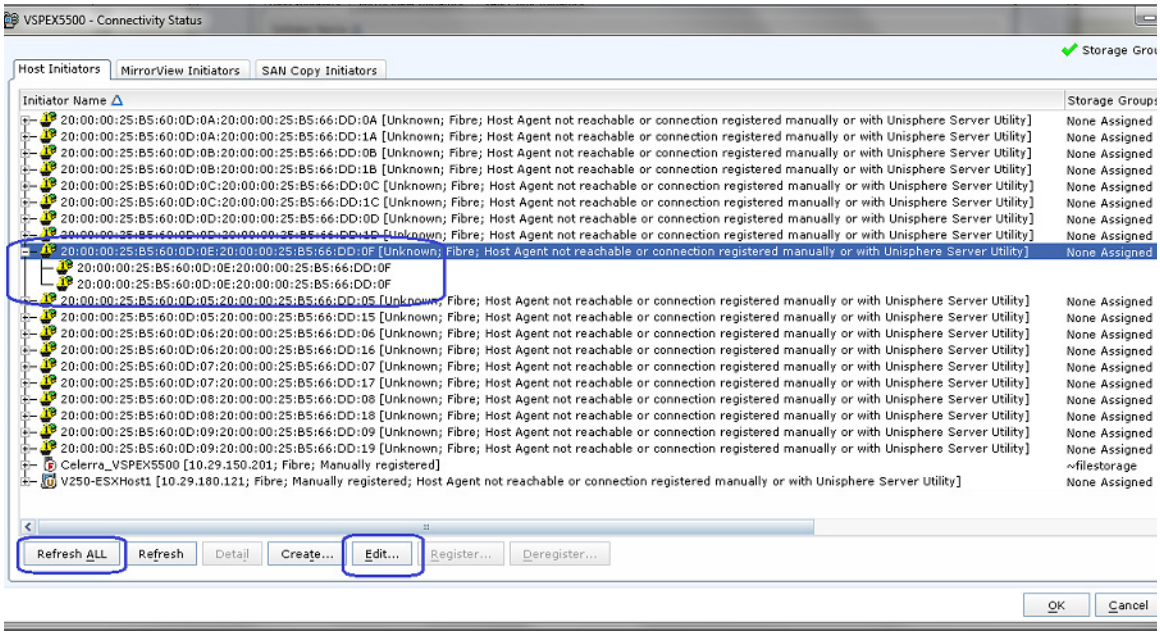

20. Click **Existing Host…** radio button as shown in Figure 109.

Figure 109 Browsing for the Existing Host

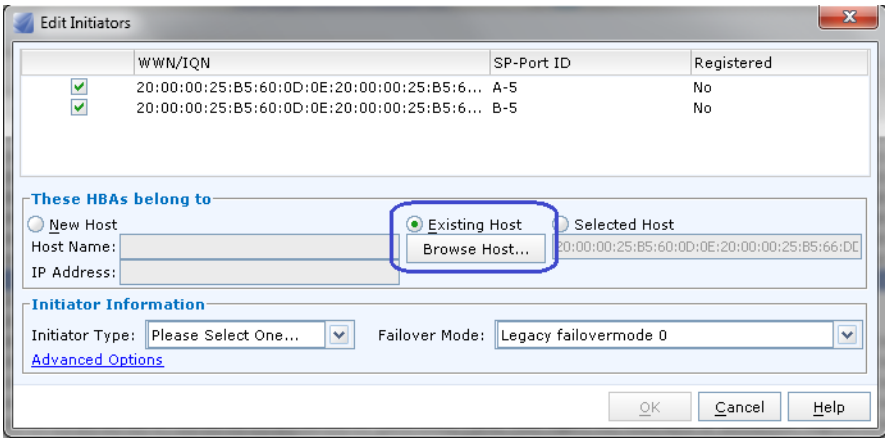

21. Select the first host created in step 17 and click **Ok** as shown in Figure 110.

Figure 110 Selecting the Existing Host Initiators

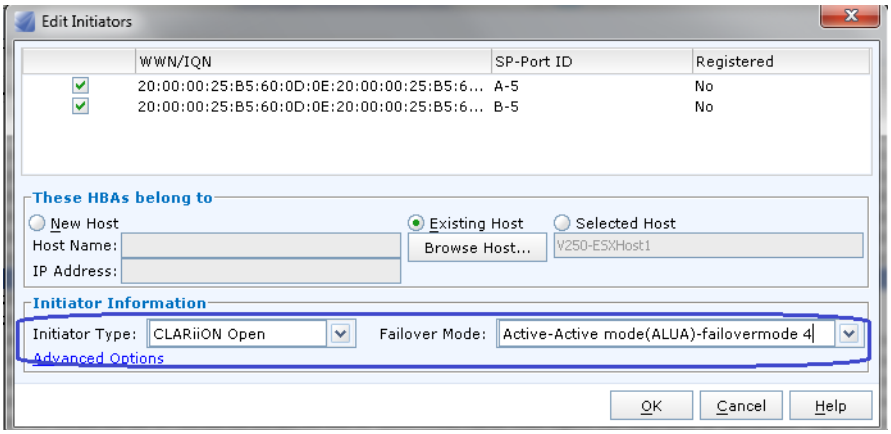

22. Select Initiator Type and Failover Mode from the respective drop-down list as shown in Figure 111 and click **Ok**.

Figure 111 Entering Initiator Information

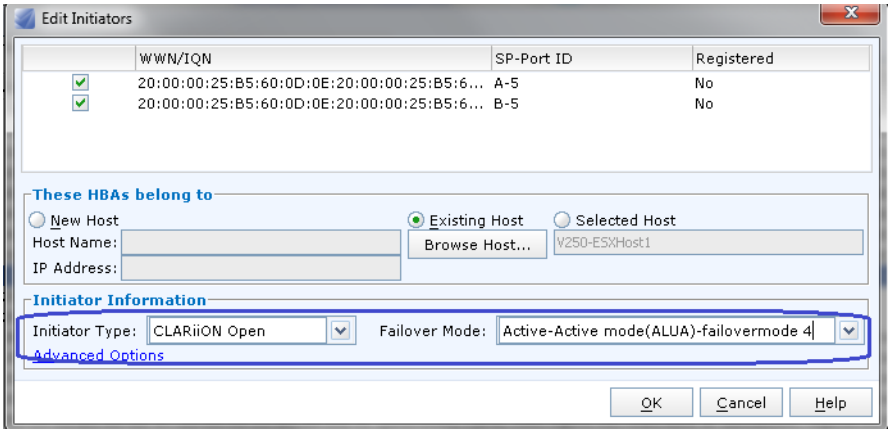

23. Click **Refresh All**. Two WWPNs of the first ESXi host should be visible from both SPs as shown in Figure 112.

Figure 112 Window Showing WWPNs of ESXi Host

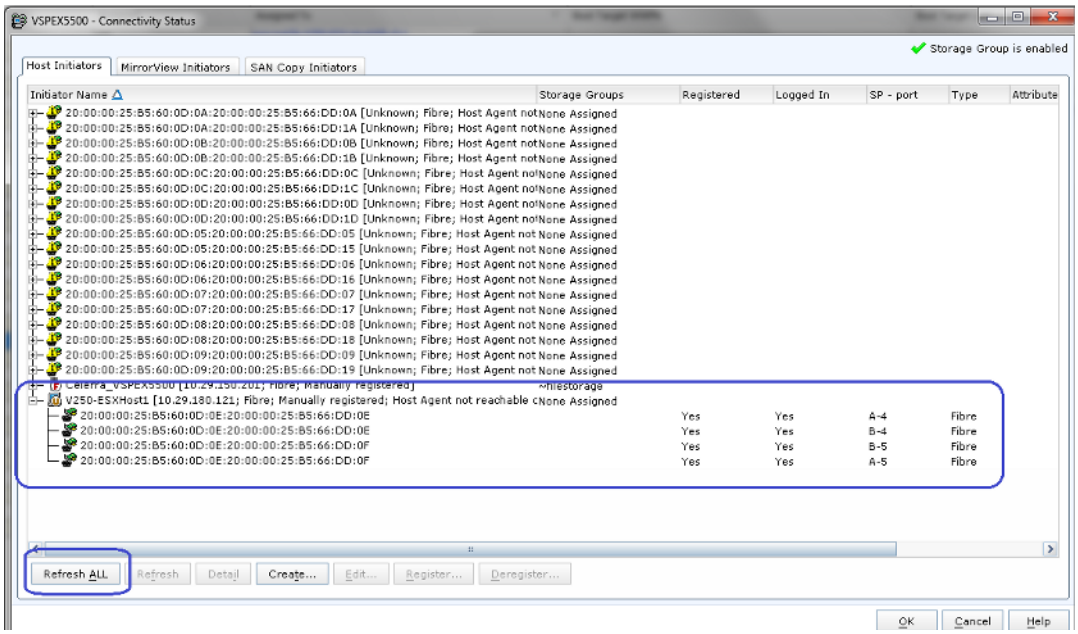

24. Repeat steps 15 to step 23 for remaining 9 servers. When all 10 servers are registered, the Host Initiators window should show all of them as in Figure 113.

Figure 113 Connectivity Status of All the TEN Servers

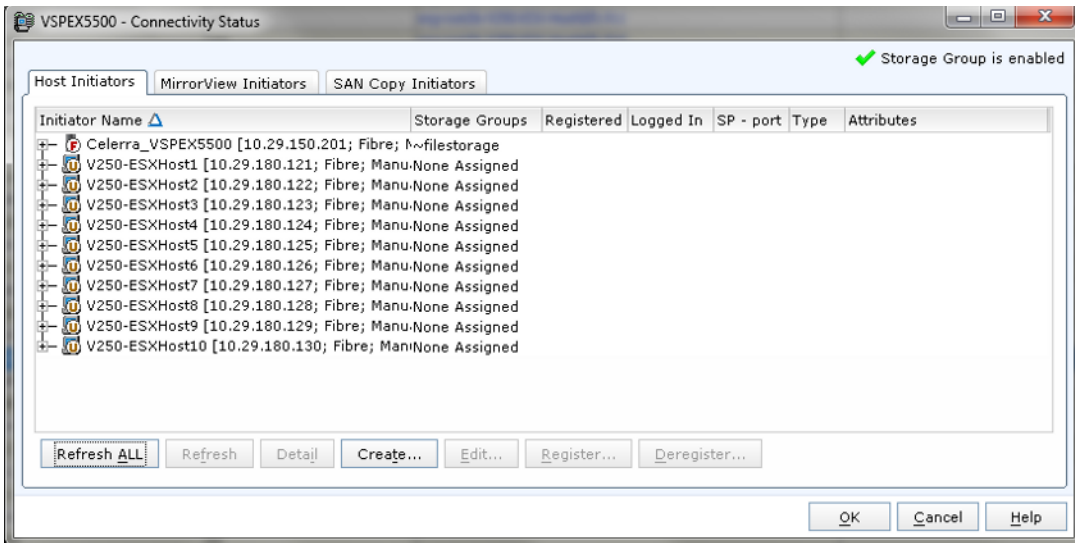

25. Click the **Hosts** tab on the menu bar in the EMC Unisphere window, click **Storage Groups** as shown in Figure 114.

Figure 114 Managing Storage Groups in the EMC Unisphere Window

26. Click **Create** as shown in Figure 115.

Figure 115 Creating Storage Groups

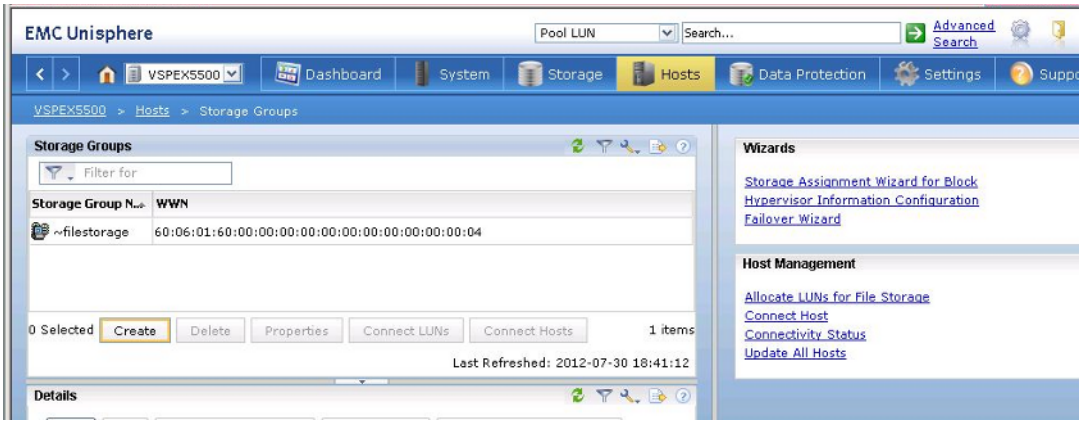

27. Create host group for the first ESXi host as shown in Figure 116.

Figure 116 Creating Host Groups

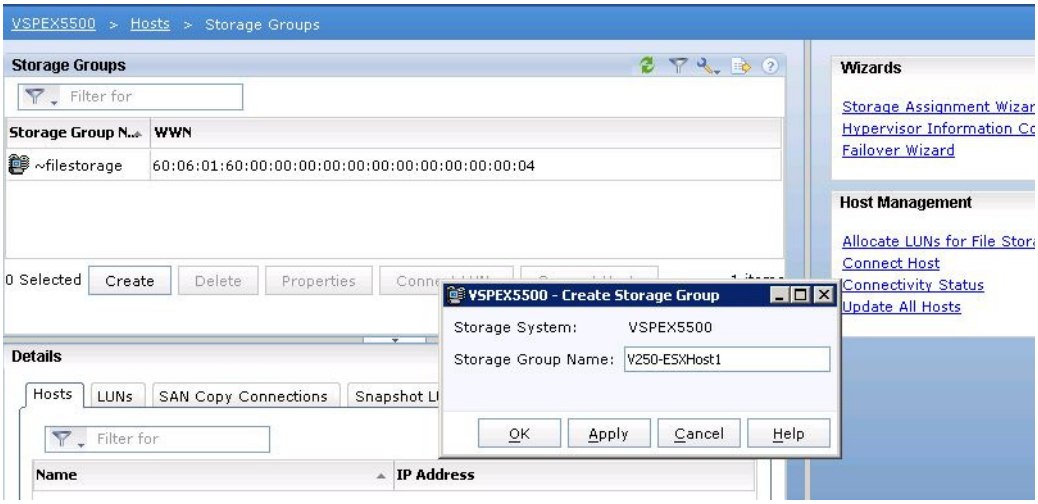

28. You will see a confirmation popup followed by an acknowledgement popup window. The acknowledgement window will ask for adding LUNs. Click **Yes** as shown in Figure 117.

Figure 117 Confirmation to Add LUNs

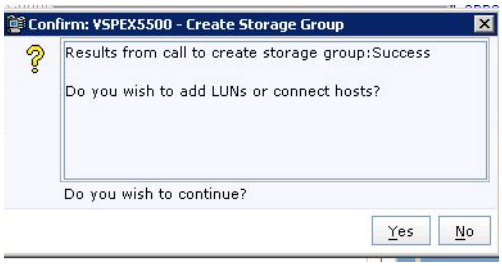

29. Expand active SP and select LUN 0. Click **Add** to add LUN as shown in Figure 118.

Figure 118 Adding LUNs

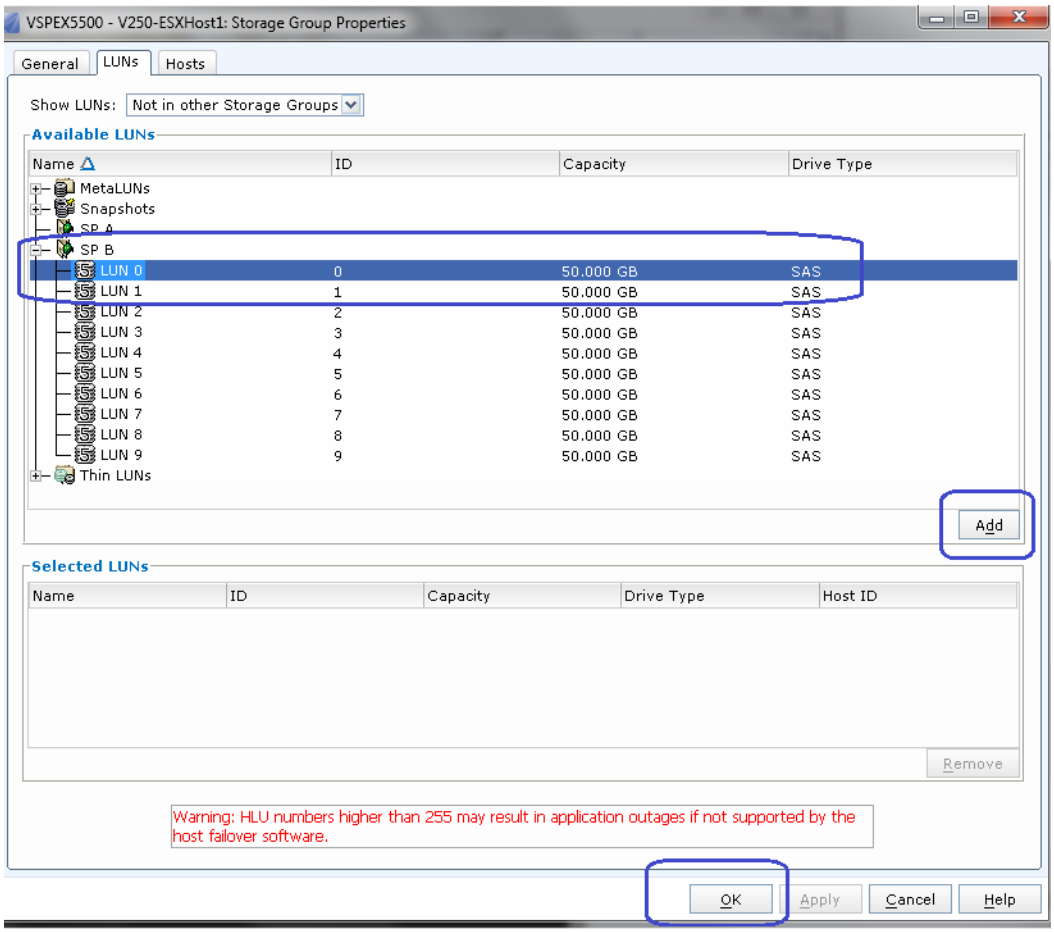

30. You will see a confirmation popup window about adding LUN 0 for the storage group. Click **Ok**. Click the **Hosts** tab in the "Storage Group Properties" window. Select ESXi host 1 as shown in Figure 119.

Figure 119 Selecting the Hosts to be Connected

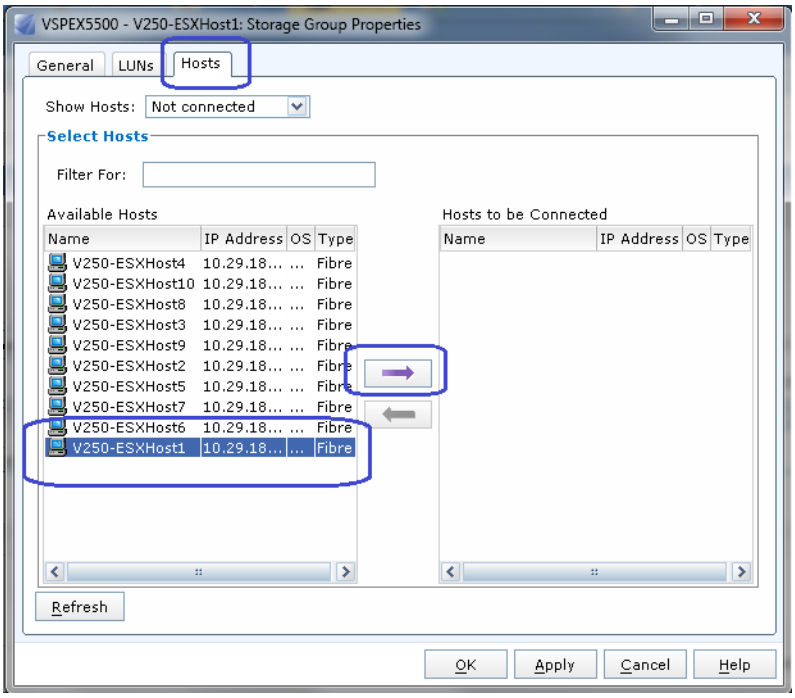

31. Click **Yes** in the confirmation popup window as shown Figure 120.

Figure 120 Confirmation to Connect the Hosts to the Storage Group

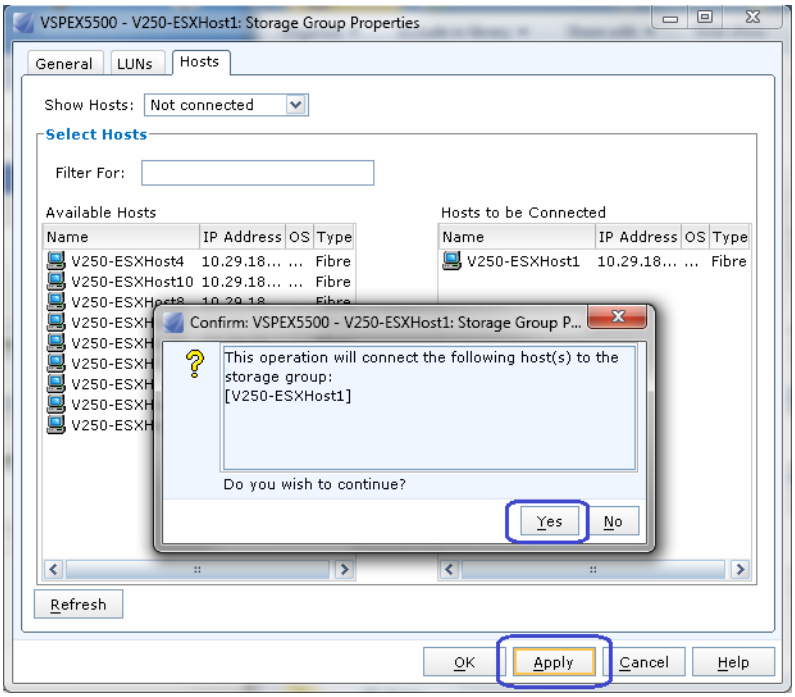

32. Repeat steps 25 to 31 for all remaining 9 hosts. After adding all the hosts, the “Storage Groups” list is as shown in Figure 121.

Figure 121 Storage Group After Adding All the Hosts

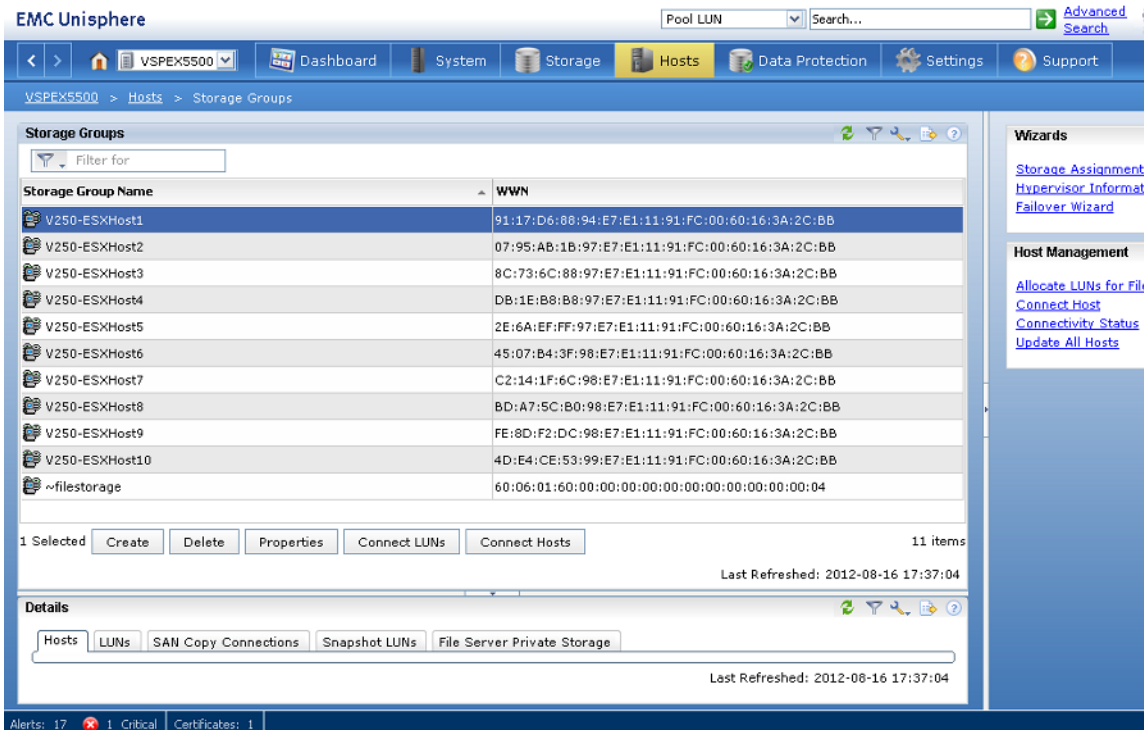

33. Launch the Cisco UCS Manager GUI again, and click the **Equipment** tab and select a server. Click the **KVM Console** link on the right pane of the Cisco UCS Manager window as shown in Figure 122.

Figure 122 Launch the KVM Console

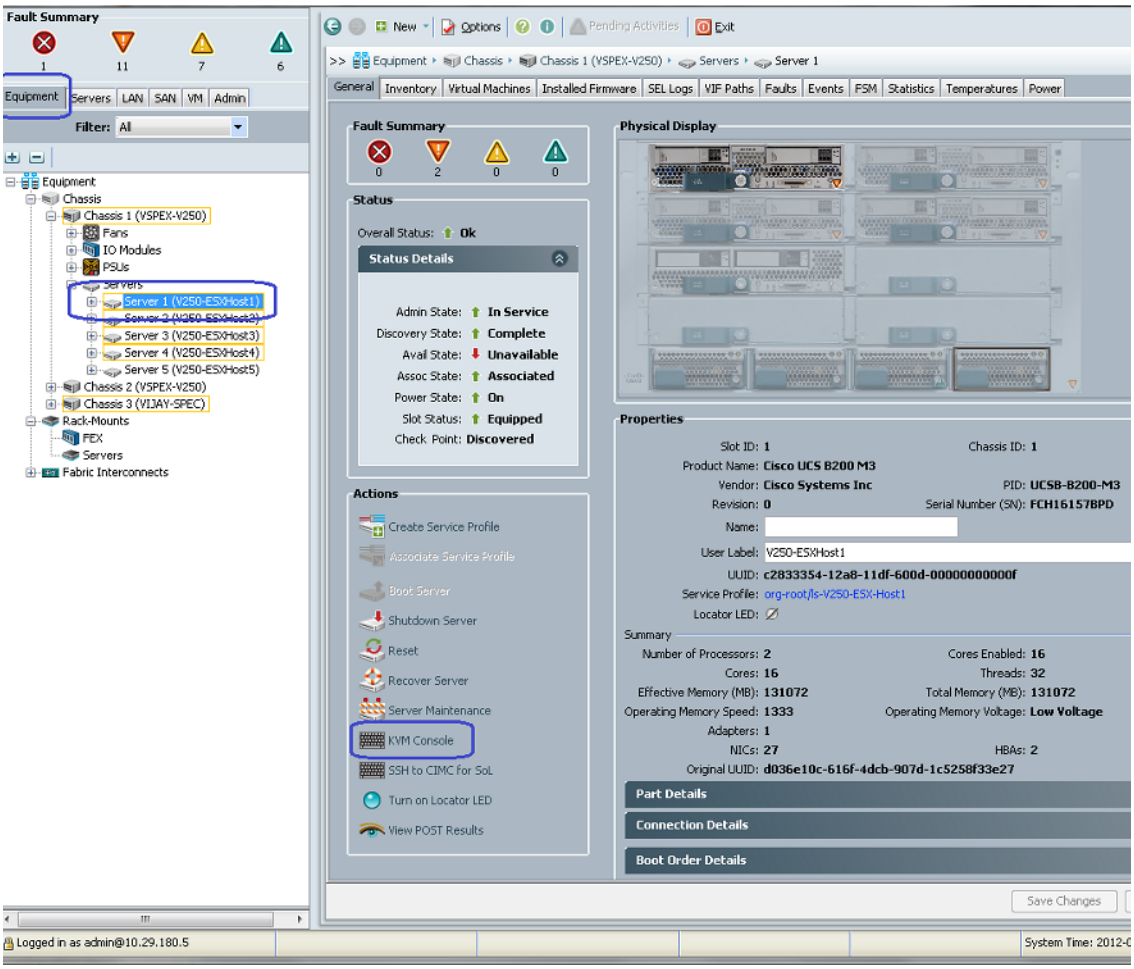

34. Click the **Macros** after KVM console is launched, and select **Ctr-Alt-Del** as shown in Figure 123.

Figure 123 Reboot the Server

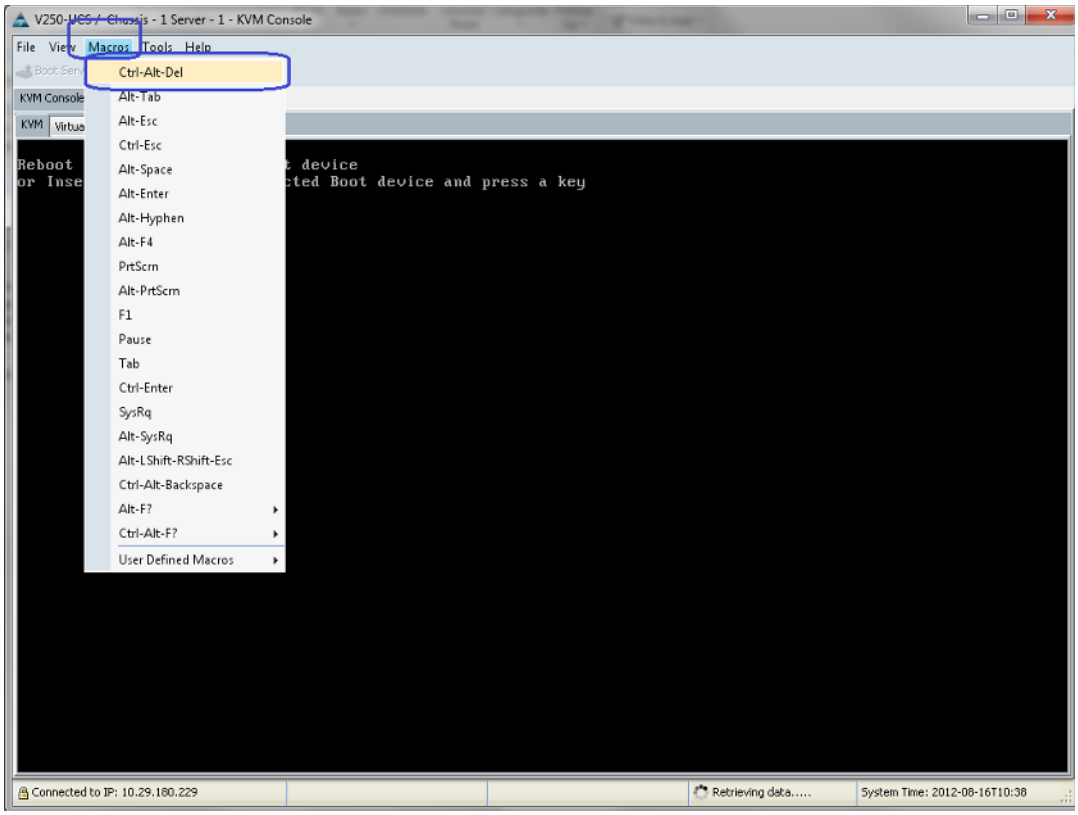

35. This will power cycle the Cisco UCS B200 M3 Blade Server. After BIOS execution, you should see each vHBA being polled for boot media and Option ROM must list the WWPN of the VNX5500 FC port of given fabric as shown in Figure 124.

Figure 124 Details of Cisco UCS B 200 M3 Blade Server

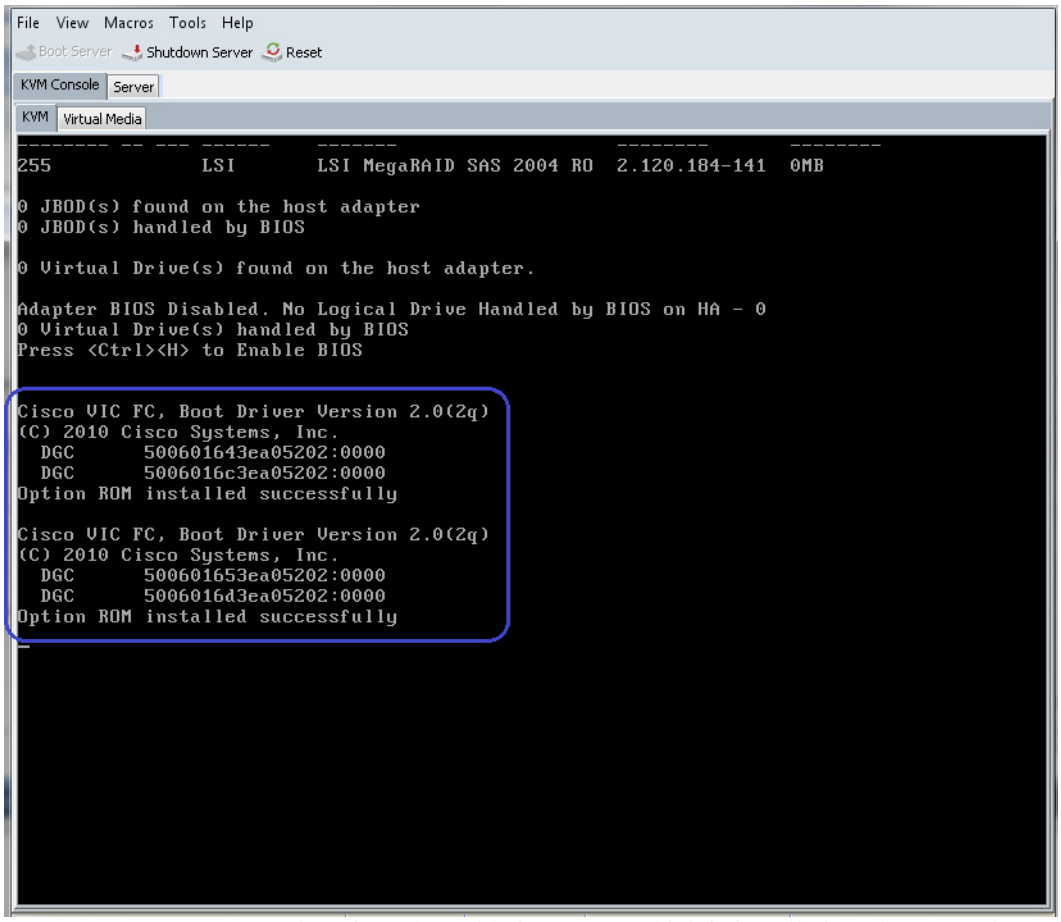

As there is no bootable image yet which is installed on the LUN, the server will not actually boot; however this is a validation of end-to-end SAN boot infrastructure from Cisco UCS B200 M3 Blade Servers to the VNX5500 LUN.

Configure NFS Storage

This section covers the configuration of NFS storage on VNX5500 storage array.

1. To Create Storage Pools for NFS Datastore. Click **Storage > Storage Configuration > Storage pools**.

Figure 125 Selecting Storage Pools in the EMC Unisphere

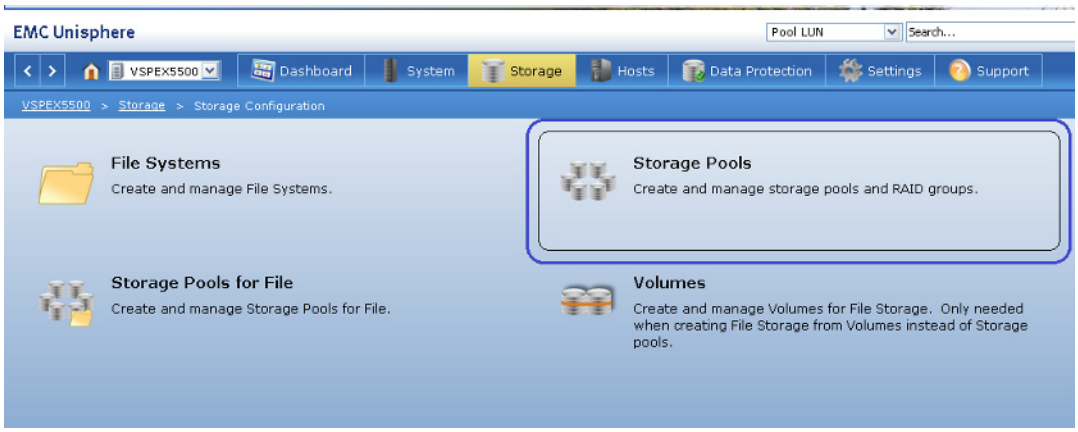

2. From the Storage Pools, click **Create**.

Figure 126 Creating Storage Pools

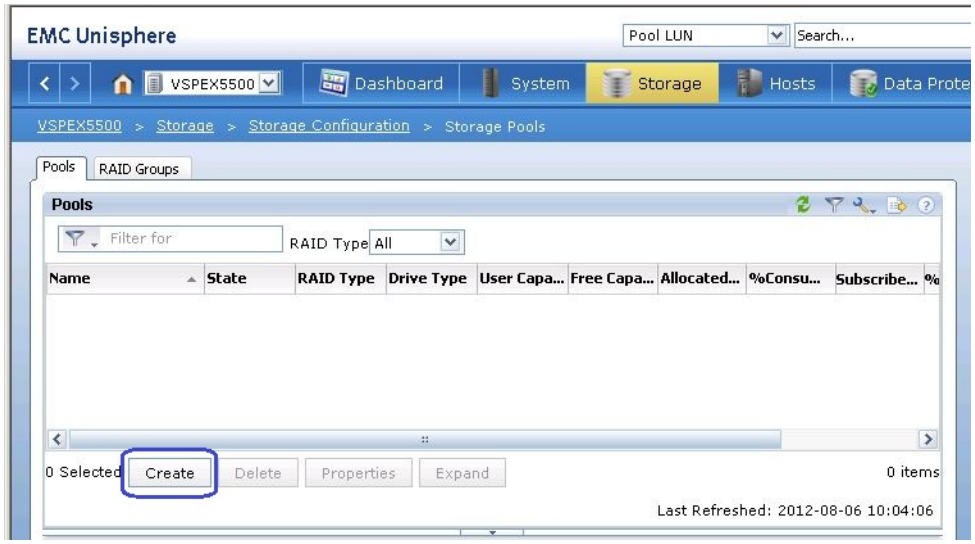

3. Enter the Storage Pool Name as “PerformancePool” and Select RAID type as RAID5 from the drop-down list. Then, Select the required SAS disks (150 Disks required for V250 validation) from the drop-down list as shown in Figure 127.

Note To Validate 250 VMs IO performance, VNX5500 storage configuration requires minimum of 165 disks. Out of 165 disks, 150 disks reserved for NFS configuration and 15 disks from “Bus 0 Enclosure 0” are reserved for three purposes: VNX OE operation system, Hot spare and SAN boot of ESXi hypervisor OS. Make sure, you did not choose “Bus0 Enclosure 0” drives during NFS Storage “PerformancePool” creation. Also, VNX5500 does not support more than 75 drives during storage pool creation. In order to choose 150 disks for the given storage pool, create the pool with 75 drives and then expand it with additional 75 drives.

Figure 127 Entering Details for Creating Storage Pools

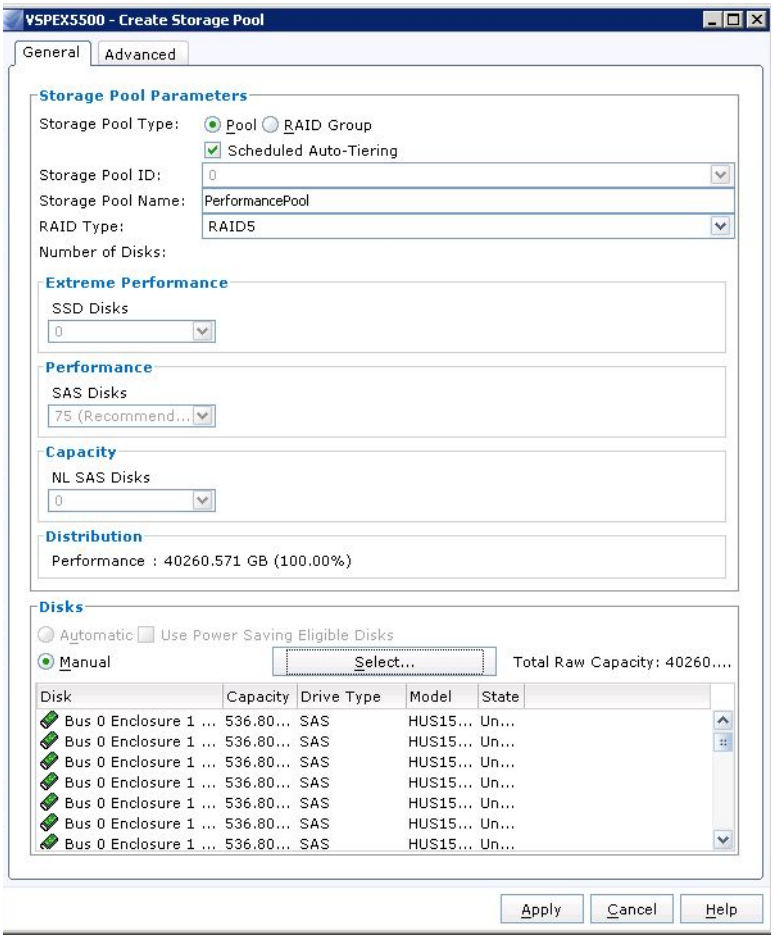

4. Manually Select 75 SAS disks and click **Apply** to initiate Pool creation and Click **Yes** to continue the Pool creation.

Figure 128 Confirmation on Creating Storage Pool

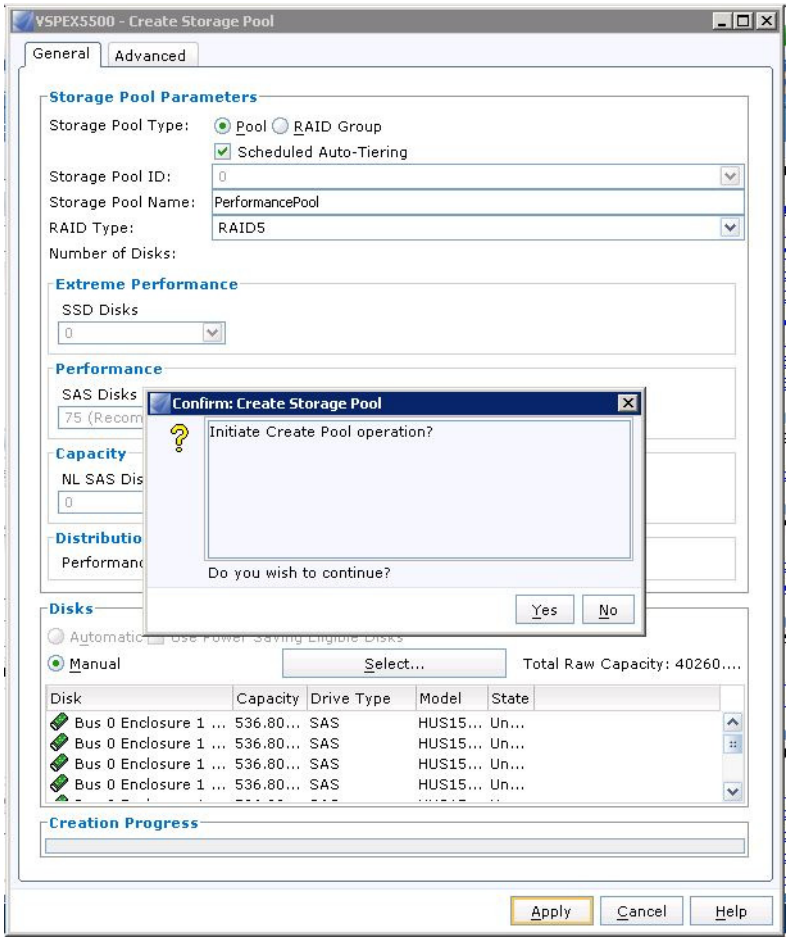

5. After the successful creation of “PerformancePool” with 75 Disks, click **Ok** in the success notification popup window.

Figure 129 Window Showing Successful Creation of Storage Pool

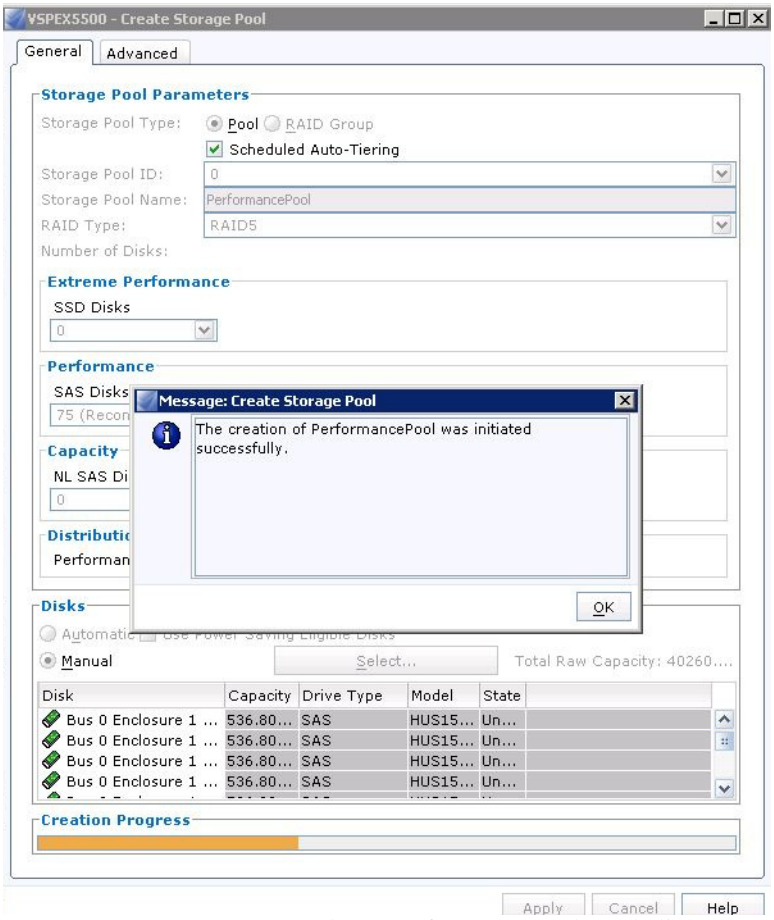

6. Select “PerformancePool” and click to refresh, until initialization state shows “Ready”. To add 75 more disks to the pool, select “PerformancePool” and click **Expand**.

Figure 130 Adding More Disks to the Pool

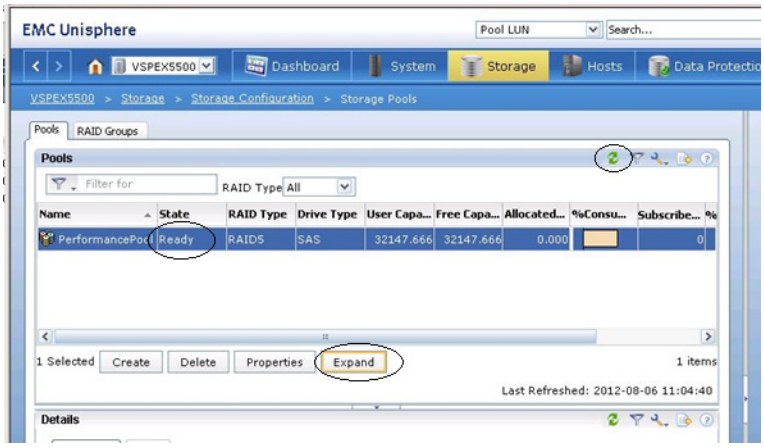

7. From the drop-down list, choose “75 (Recommended)” disks to expand. Click **Select**.

Figure 131 Expanding Storage Pool

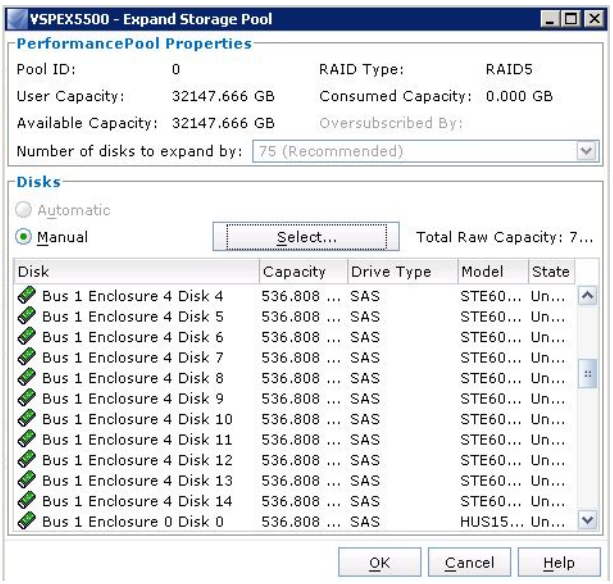

8. Click **Ok** in the popup window on successful expansion of PerformancePool.

Figure 132 Completion of Storage Pool Expansion

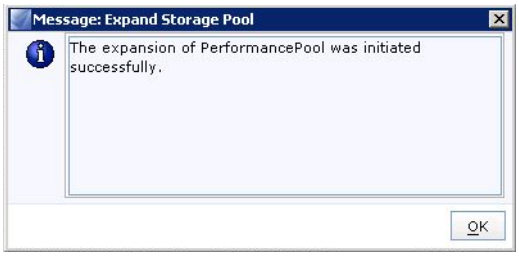

9. Wait for the expansion of the pool to be completed and the state to show “Ready”.

Figure 133 Window Showing Storage Pools After the Expansion

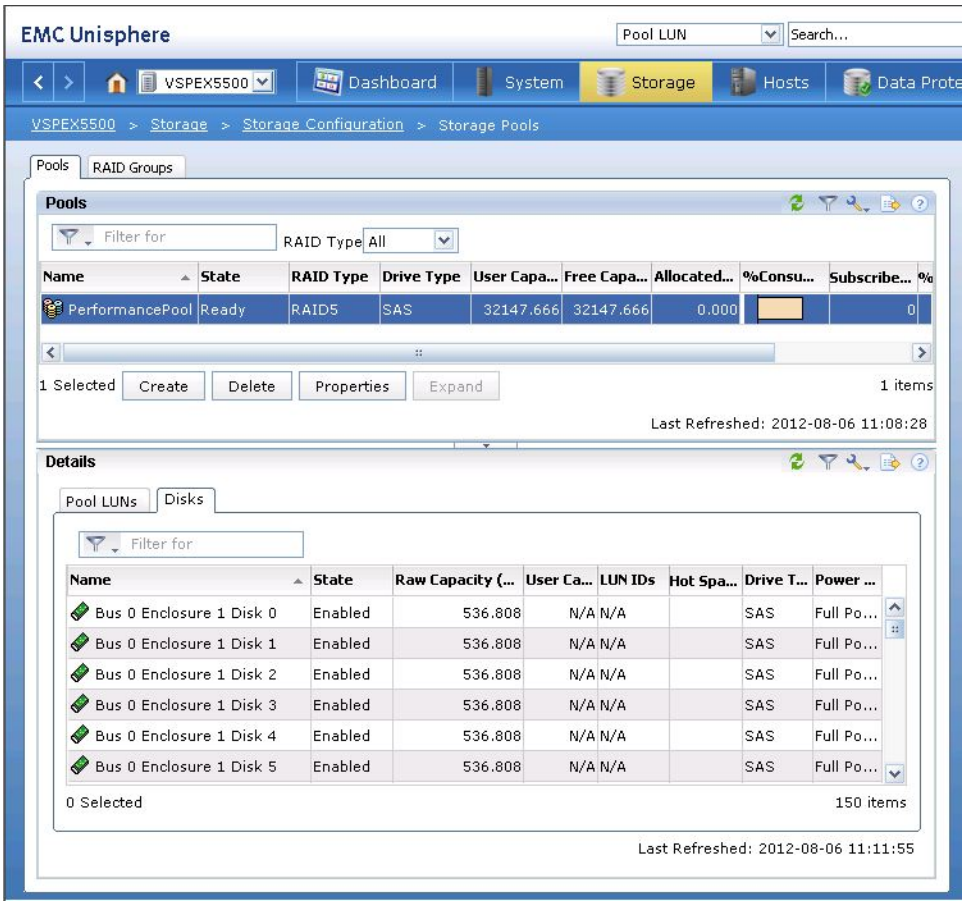

10. To create Hot Spares for the system, click **System > Hot Spares**.

Figure 134 Selecting Hot Spares in EMC Unisphere

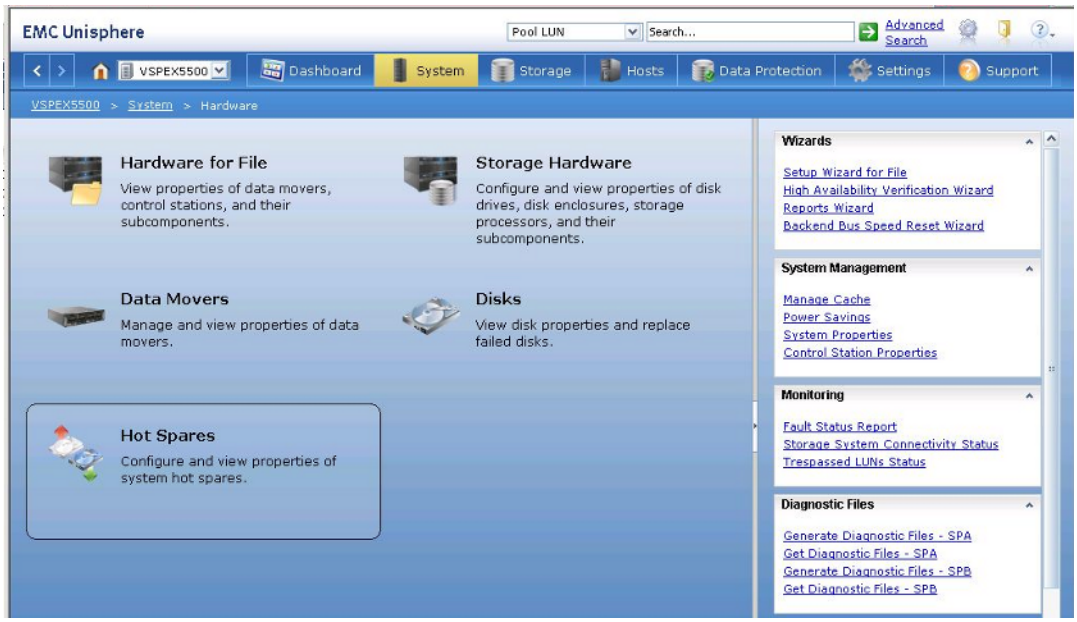

11. Click **Create** to create Hot Spares.

Figure 135 Creating Hot Spares

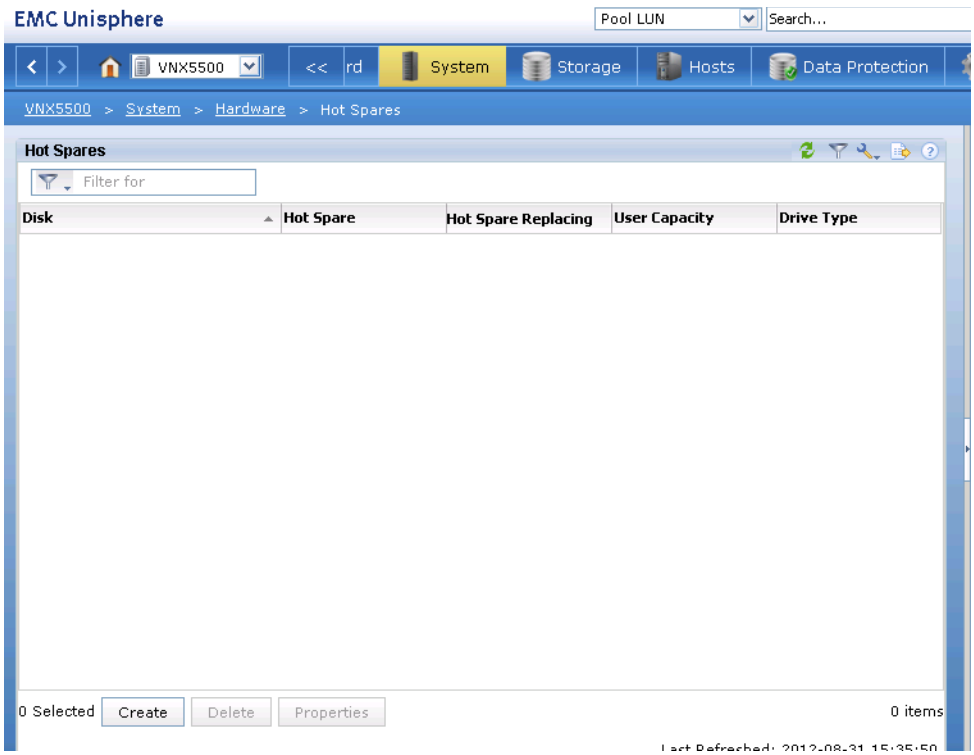

12. In the Create Hot Spare Window. Click **RAID Group** radio button for Storage Pool Type. Select Storage Pool ID as 1 from the drop-down list, enter the Storage Pool Name, select the RAID Type as “Hot Spare” from the drop-down list, and select the Number of Disks as 1 from the drop-down list. Click **Automatic** radio button for disks and click **Apply**.

Figure 136 Entering Storage Pool Parameters

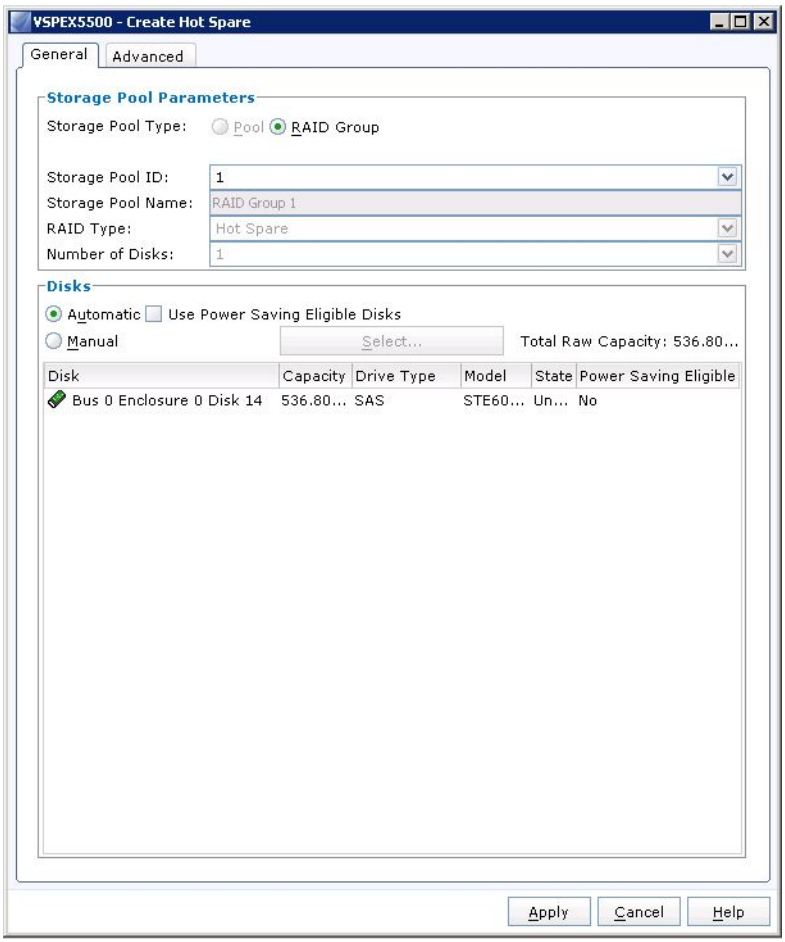

13. Figure 137 shows the RAID Group 1 has been created successfully to create the first Hot Spare for this Storage. Click **Ok** and continue creating Hot Spares.

Figure 137 Window Showing Successful Creation of RAID Group 1

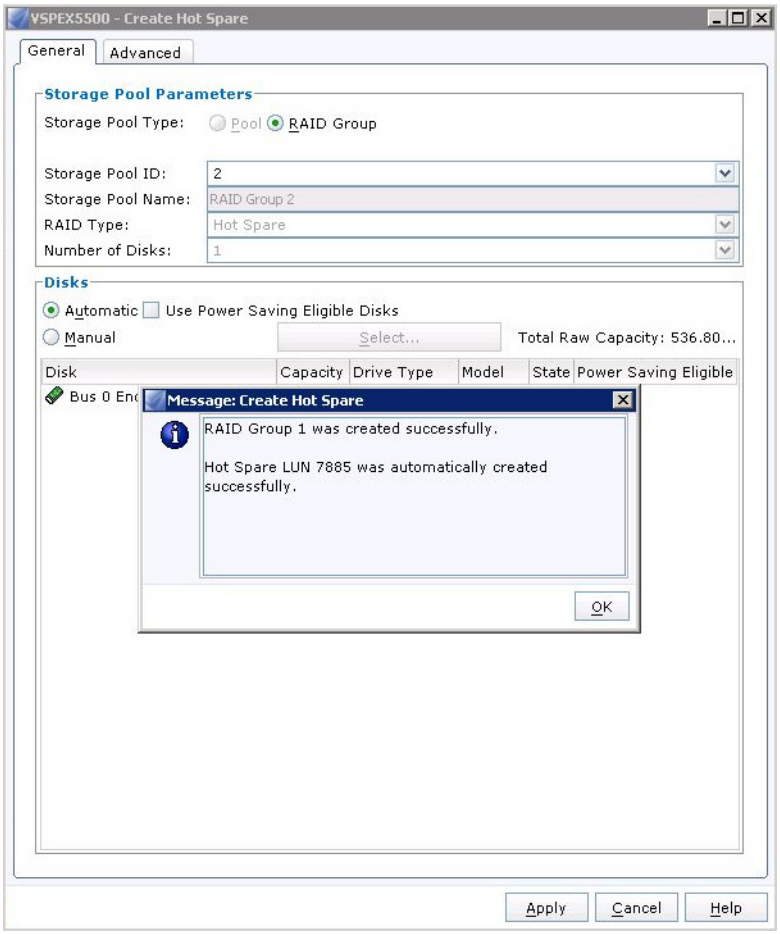

14. Repeat step 12 & Step 13 to create seven more hot spares as needed for this Storage configuration.

Figure 138 Window Showing Successful Creation of RAID Group 7

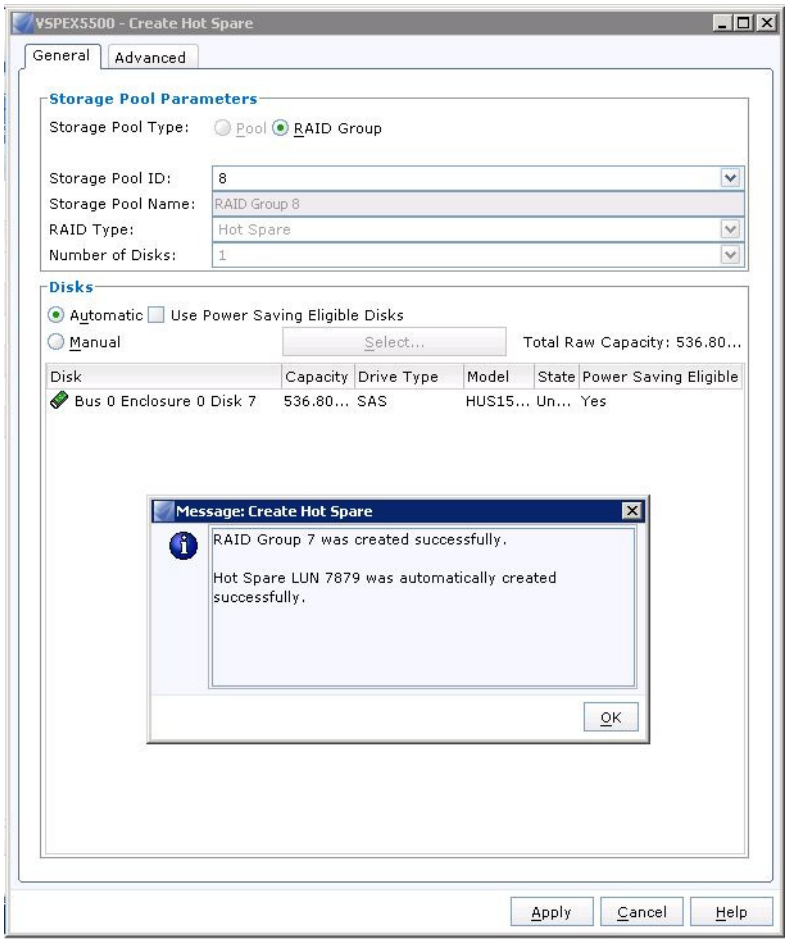

15. After creating the Hot Spares, make sure the Hot Spare state shows “Hot Spare Ready”.

Figure 139 Window Showing Hot Spare Status

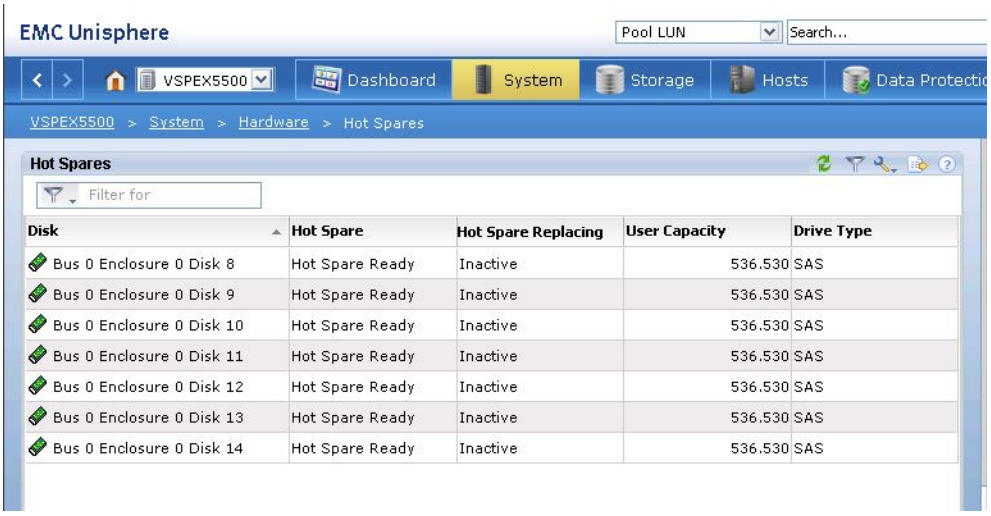

16. To create LUNs from the newly created PerformancePool for NFS Datastore; Click **Storage**, right-click on the “PerformancePool” and click **Create LUN**.

Figure 140 Creating LUN

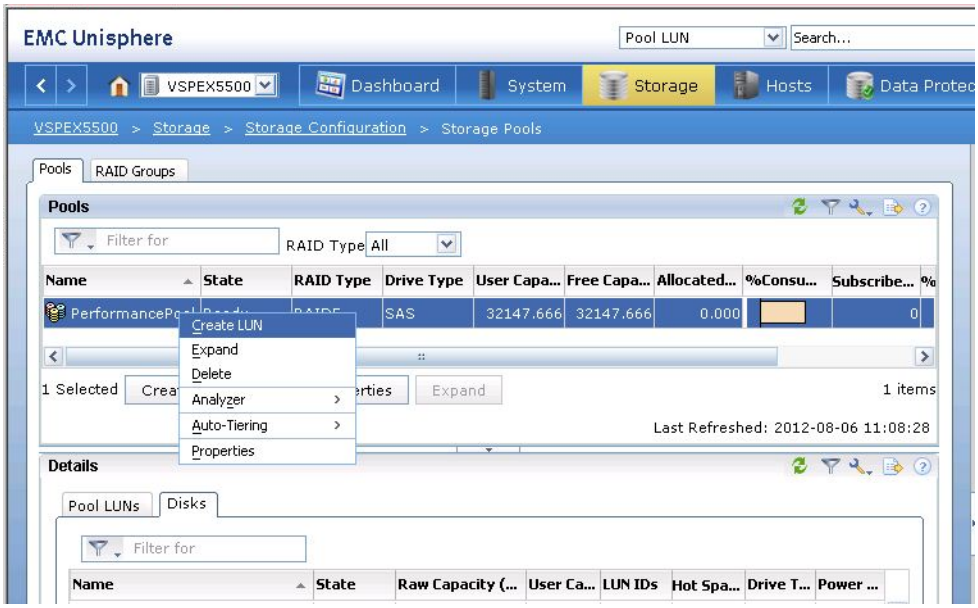

17. Click **Pool** radio button for the Storage Pool Type, Select RAID Type as "RAID5" from the drop-down list and Storage Pool for new LUN as "PerformancePool" from the drop-down list. In the LUN properties area, make sure to select User Capacity as "300GB" from the drop-down list. Select Number of LUNs to Create as "150" from the drop-down list. These 150 LUNs is equal to the number of disks selected for the "PerformancePool". Click **Automatically Assign LUN IDs as LUN Names** radio button for LUN Name. Click **Apply** to initiate the process for creating 150 LUNs.

Figure 141 Entering Details to Create LUNs

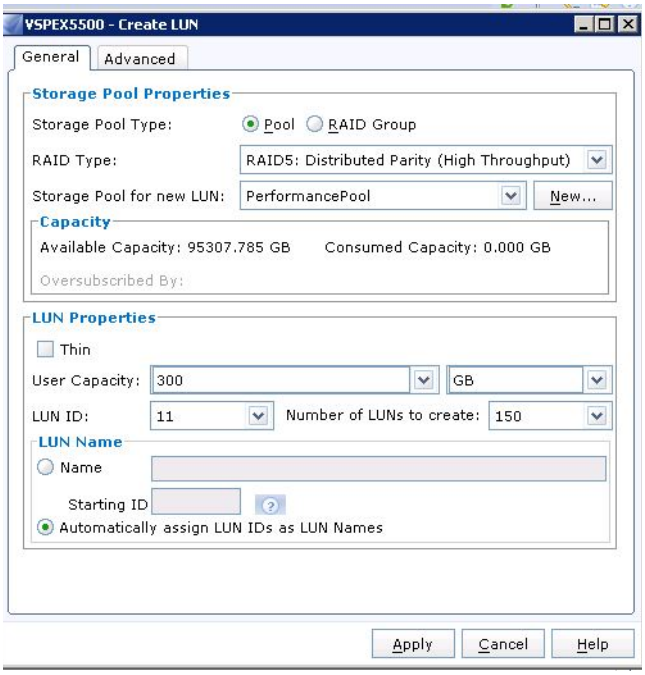

18. Click **Yes** to initiate a create LUN operation.

Figure 142 Confirmation to Create LUNs

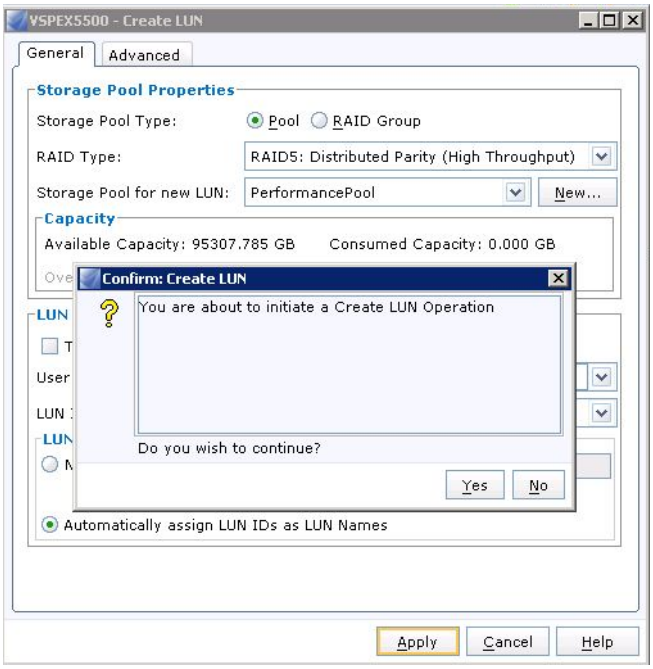

19. LUN creation is in progress. Wait for the task to be complete.

Figure 143 Window Showing LUN Creation Progress Indicator

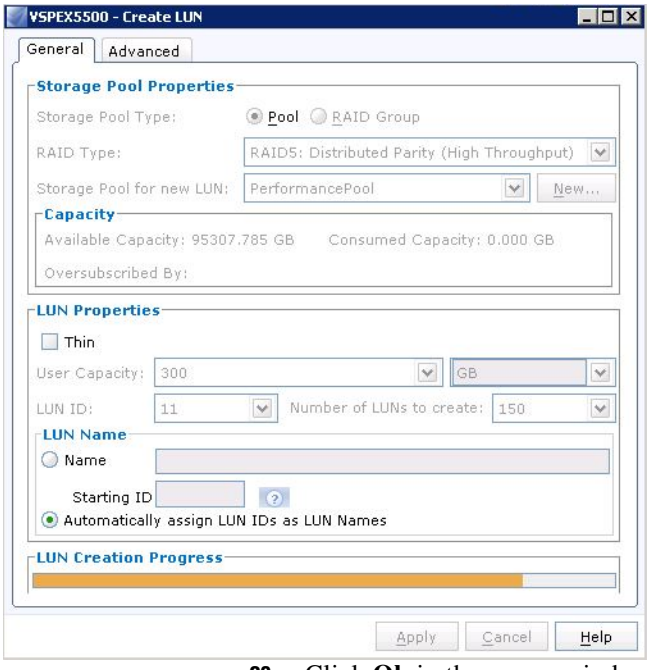

20. Click **Ok** in the popup window on successful LUN creation.

Figure 144 Window Showing Successful LUN Creation

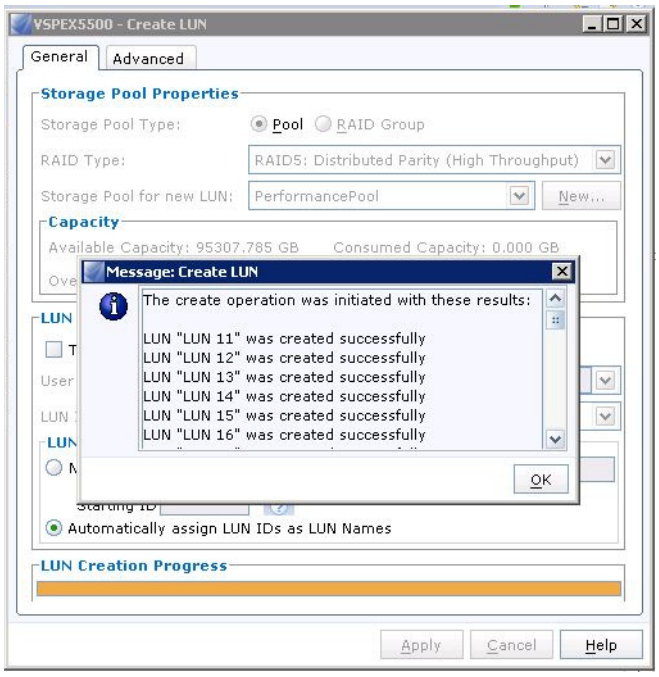

21. Select the "PerformancePool" and Select all the newly created LUNs and Click **Add to Storage Group** as shown in Figure 145. Make sure you select all the 150 LUNs from the PerformancePool.

Figure 145 Adding the Created LUNs to Storage Group

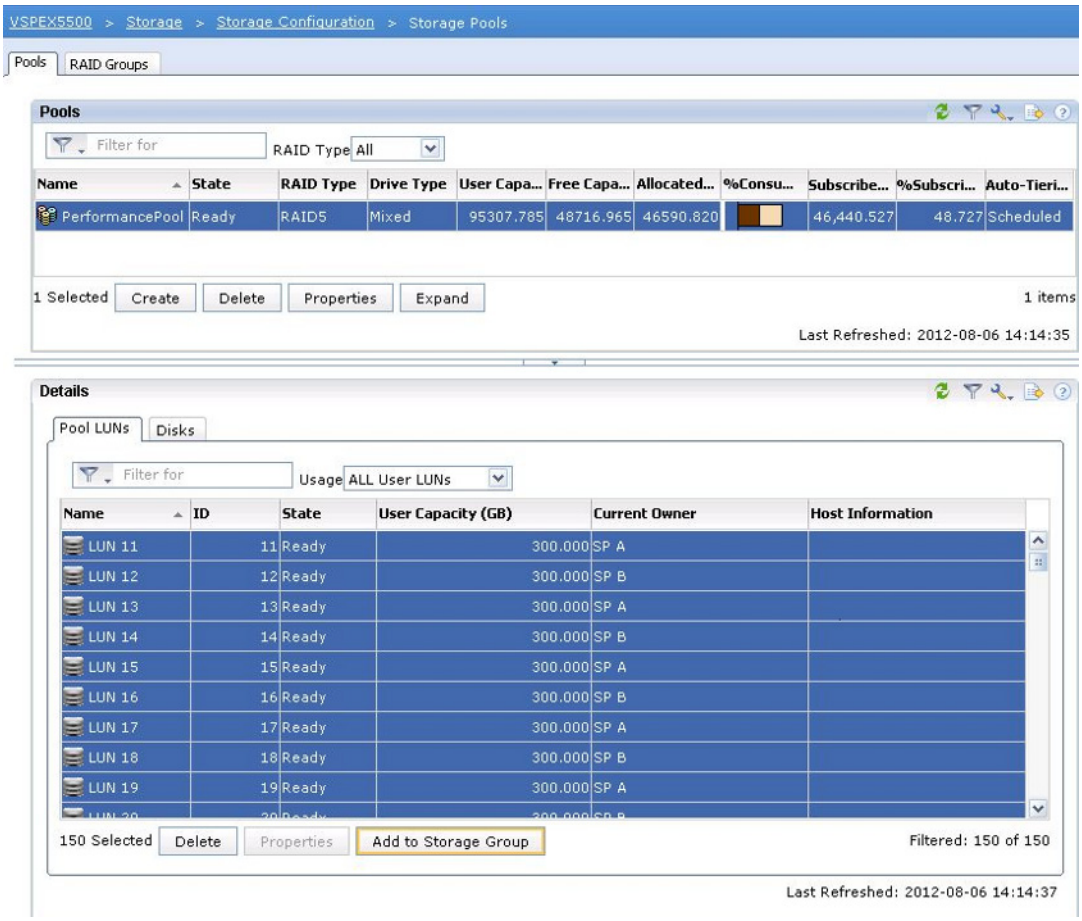

22. From the Available Storage Groups, select "~filestorage" and click the arrow highlighted in Figure 146 to add it to the Selected Storage Group. Click **Ok**.

Figure 146 Adding Storage Groups

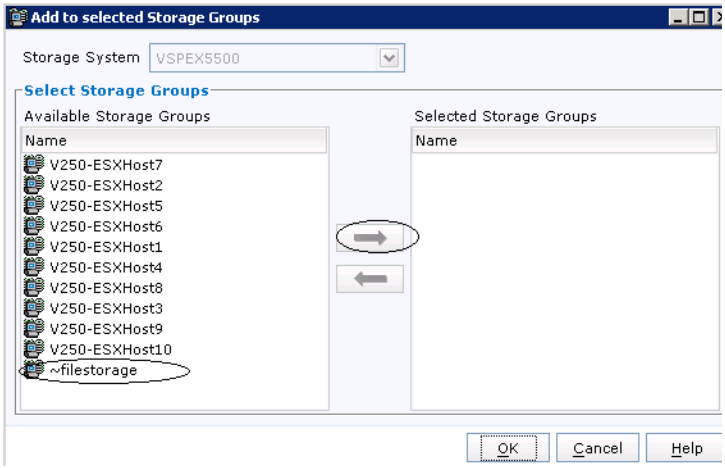

23. Make sure "~filestorage" is added to Selected Storage Groups. Click **Ok**.

Figure 147 Ensuring the Storage Group is Added

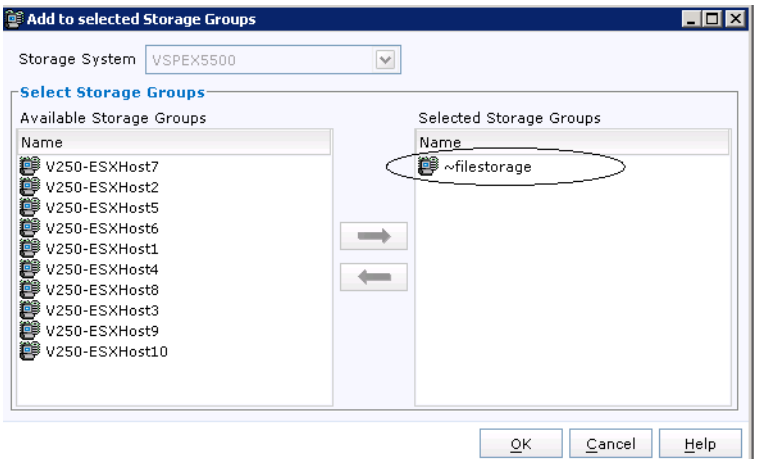

24. Click **Yes** to confirm the operation to add all the 150 LUNs to the "~filestorage" Storage group.

Figure 148 Confirmation Window to Add LUNs

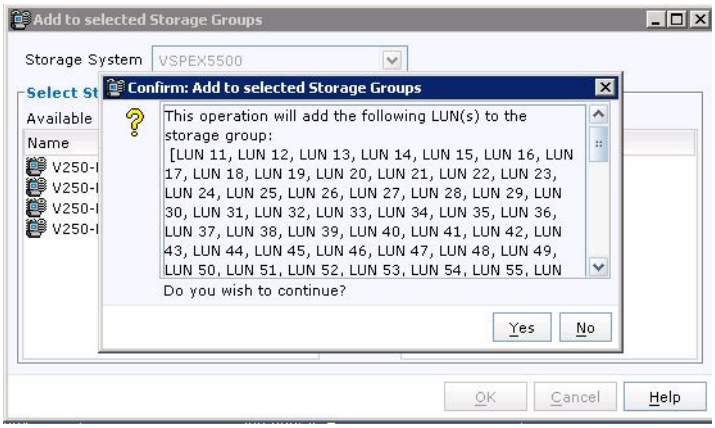

25. Make sure, all the LUNs are added to filestorage and make sure the Host Information for all the LUNs are showing "Celerra_VSPEX5500".

Figure 149 Host Information for All the Added LUNs

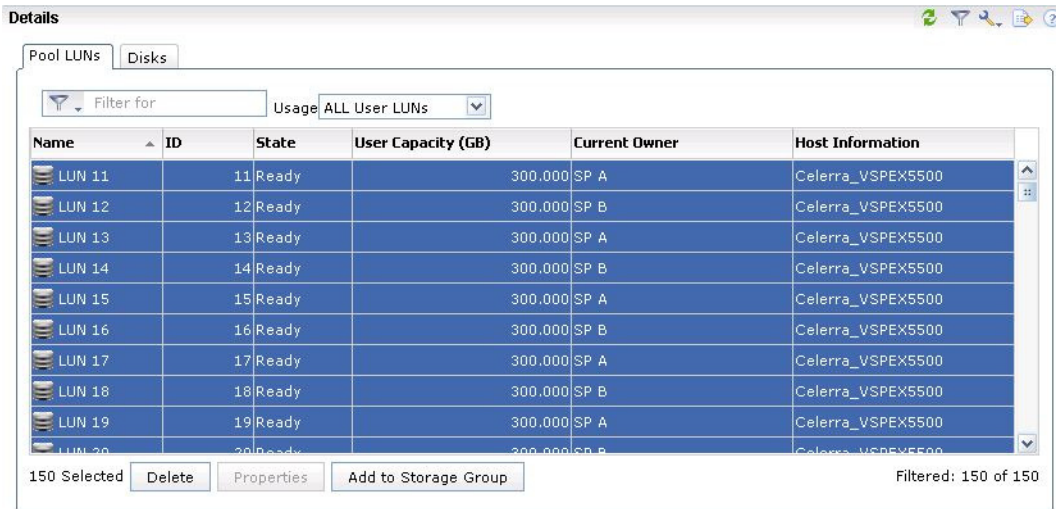

26. To discover all the 150 LUNs as volumes for NFS creation. Click **Storage > Storage Configuration > Volumes**. From Figure 150, you will see that the system volumes are created by default.

Figure 150 Window Showing System Volumes Created by Default

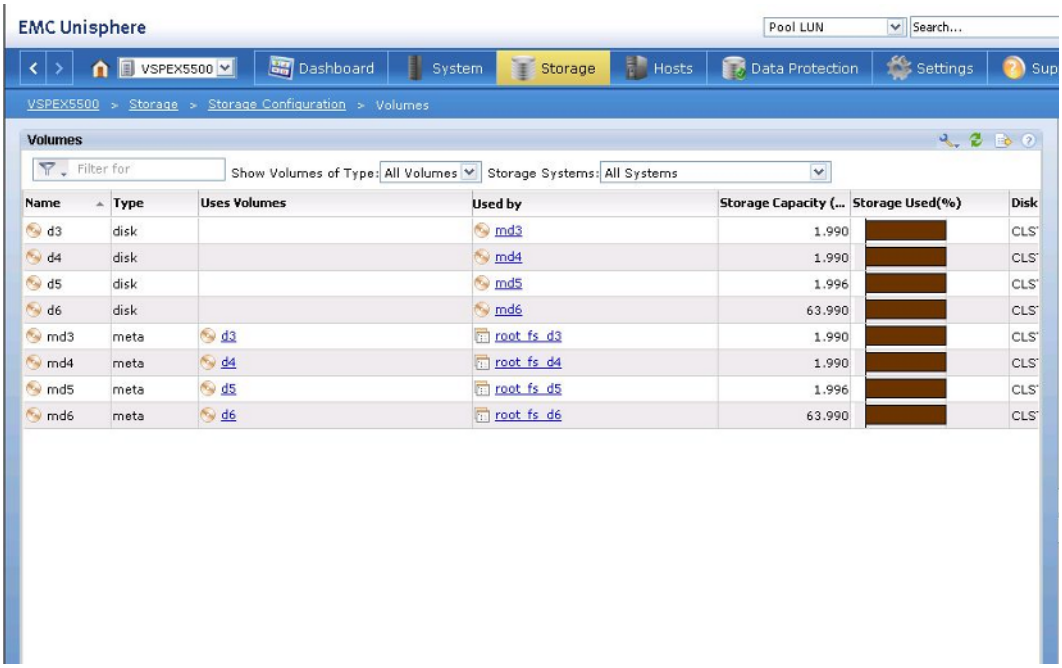

27. Log in (ssh or telnet) to the VNX Control Station IP or Storage Processor IP for the CMD line console.

Figure 151 CLI Showing List of NAS Disks

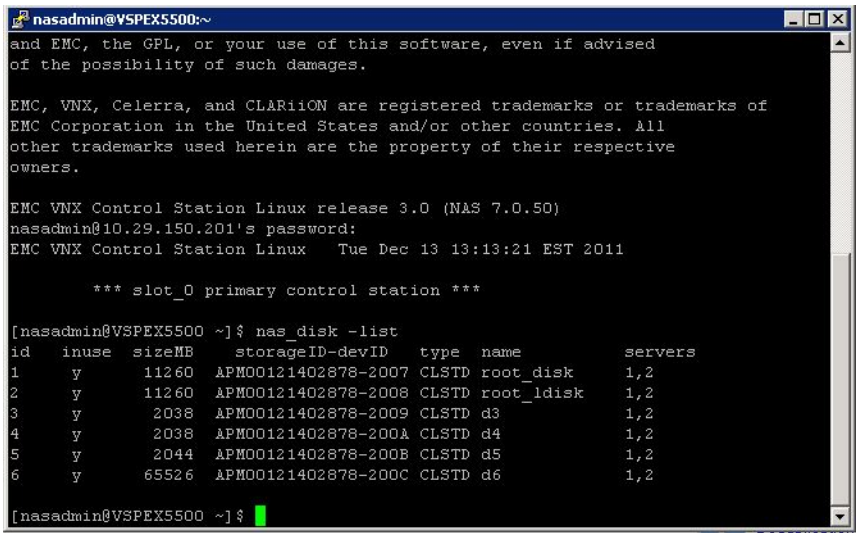

28. From the CMD line console, Type the command **nas_disk -list** to list the default volumes. Type the command **nas_diskmark -mark -all** to discover all the 150 LUNs as 150 disk volumes.

Figure 152 Command to Show All the LUNs as Disk Volumes

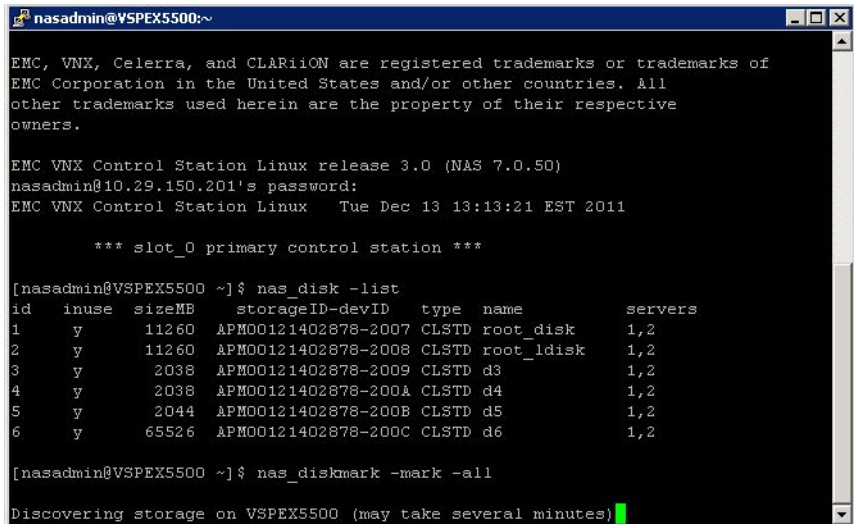

29. Wait till the discovery process is complete.

Figure 153 CLI Showing Discovery Process

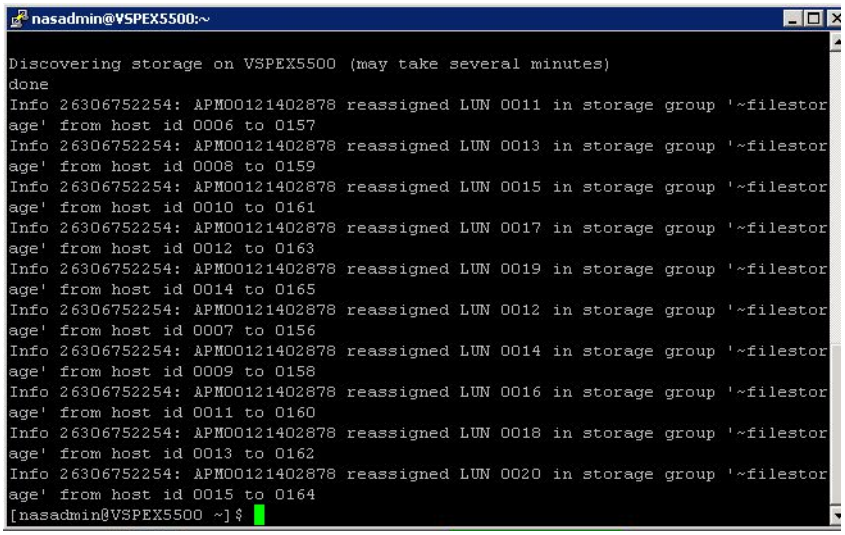

30. Type the command **nas_disk -list | grep 307 | wc -l** to make sure all the 150x300GB LUNs are discovered as 150 disk volumes.

Figure 154 Command Showing All the LUNs Discovered as Disk Volumes

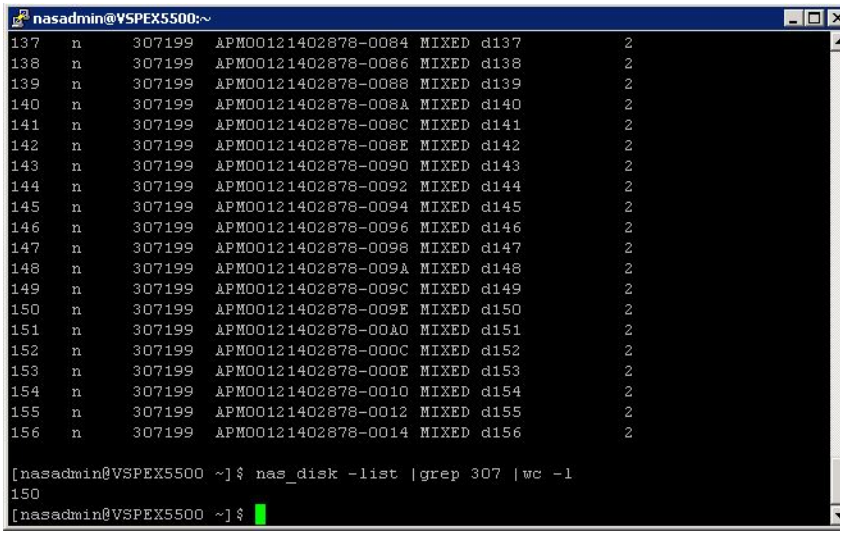

31. From the EMC Unisphere window, make sure all the new 150 disk volumes created with 300GB Storage Capacity (numbered from d7 to d156) as shown in Figure 155.

Figure 155 Verify the Storage Capacity for All the Disk Volumes

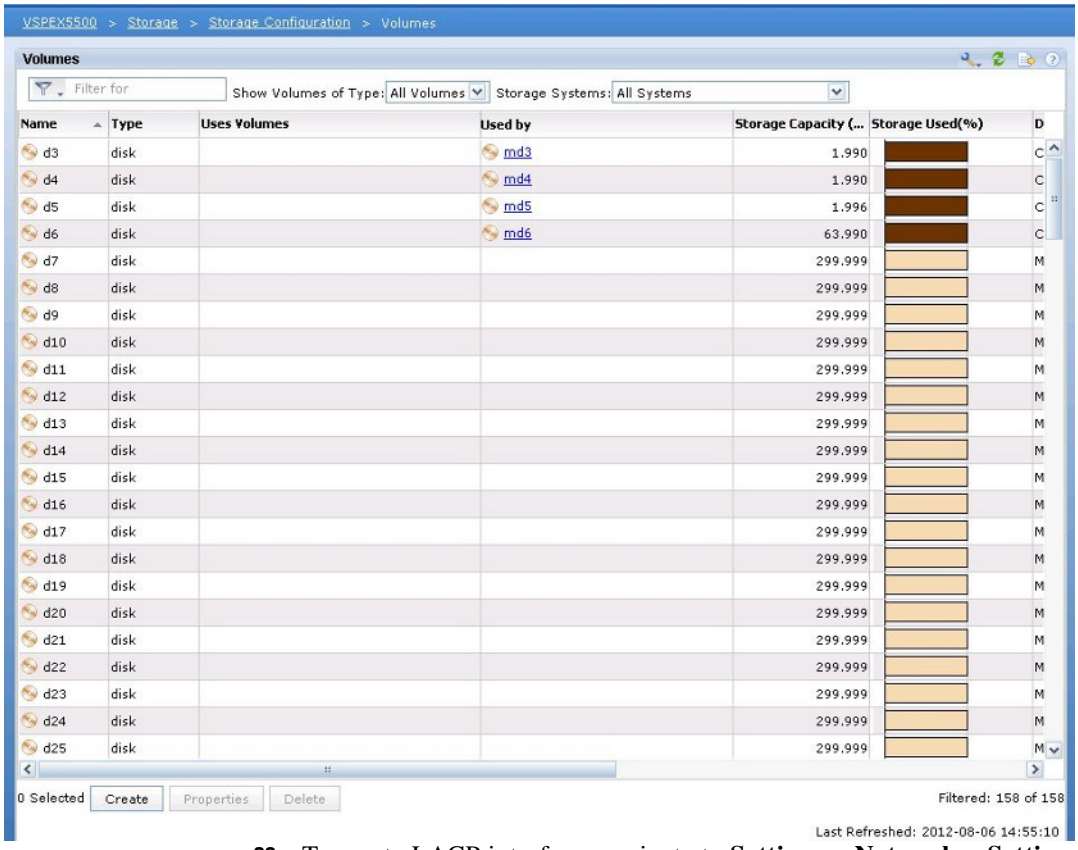

32. To create LACP interface, navigate to **Settings > Network > Settings for File**. Click **Create**.

Figure 156 Creating LACP Interface

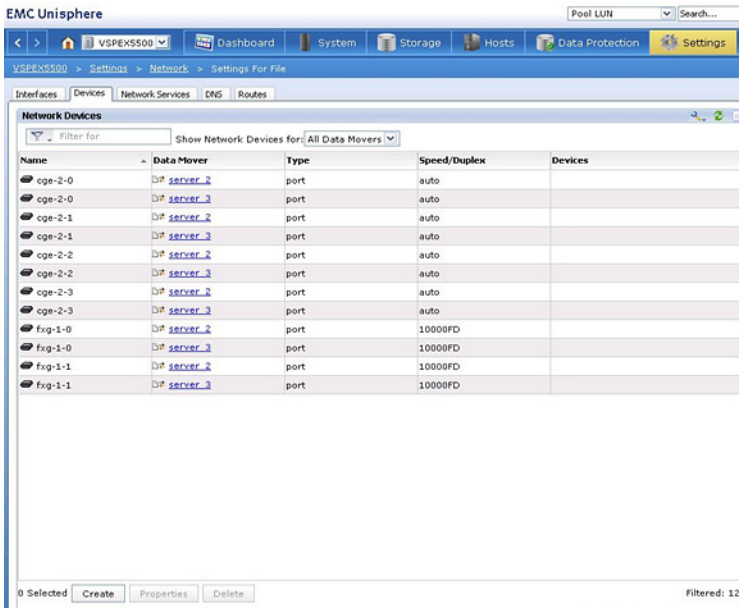

33. Select Data Mover as “All Primary Data Movers” from the drop-down list.

Figure 157 Creating Network Device

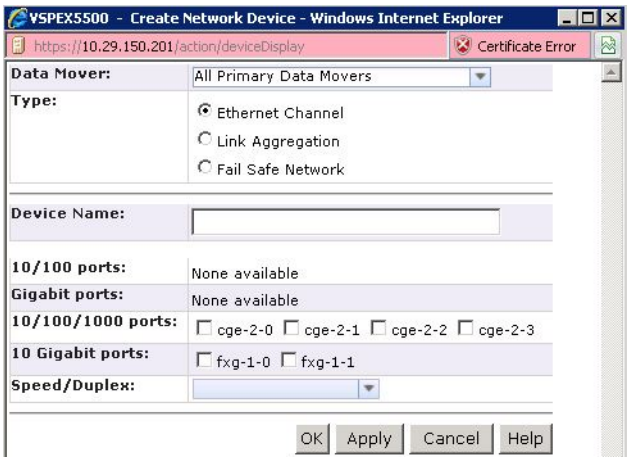

34. Click the **Link Aggregation** radio button for Type and enter Device Name as “lacp-1”. Check the check boxes for 10 Gigabit ports “fxg-1-0” and “fxg-1-1” as highlighted below. Click **Ok** to proceed the Network Device creation.

Figure 158 Entering Details to Create Network Device

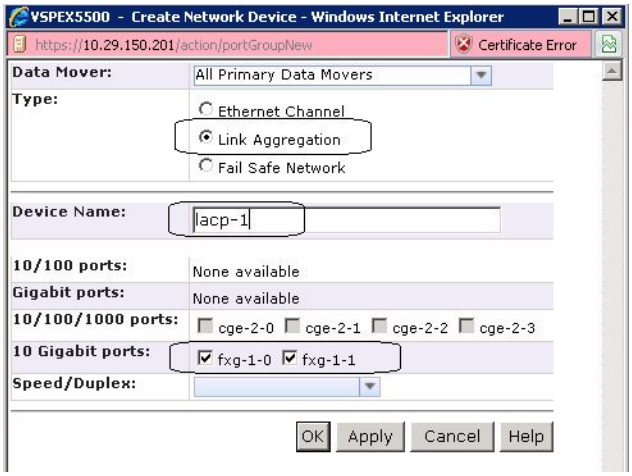

35. Figure 159 shows the creation of LACP Network device name as “lacp-1”

Figure 159 New Network Device is Created

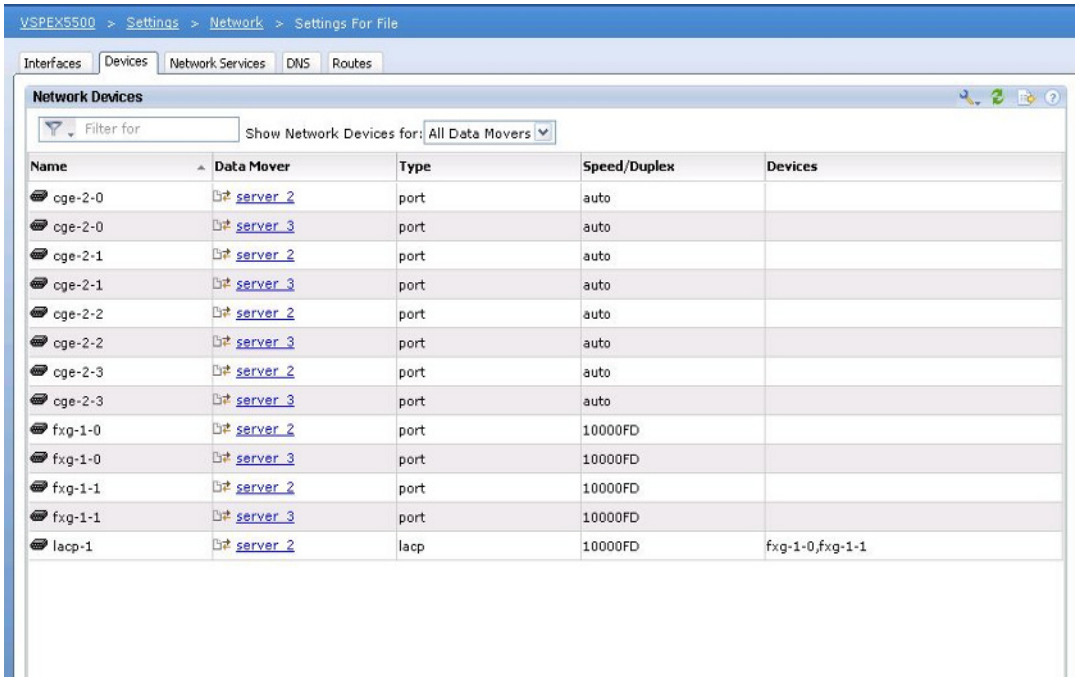

36. In the "Settings for File" window, click the **Interfaces** tab and click **Create**.

Figure 160 Creating Interfaces

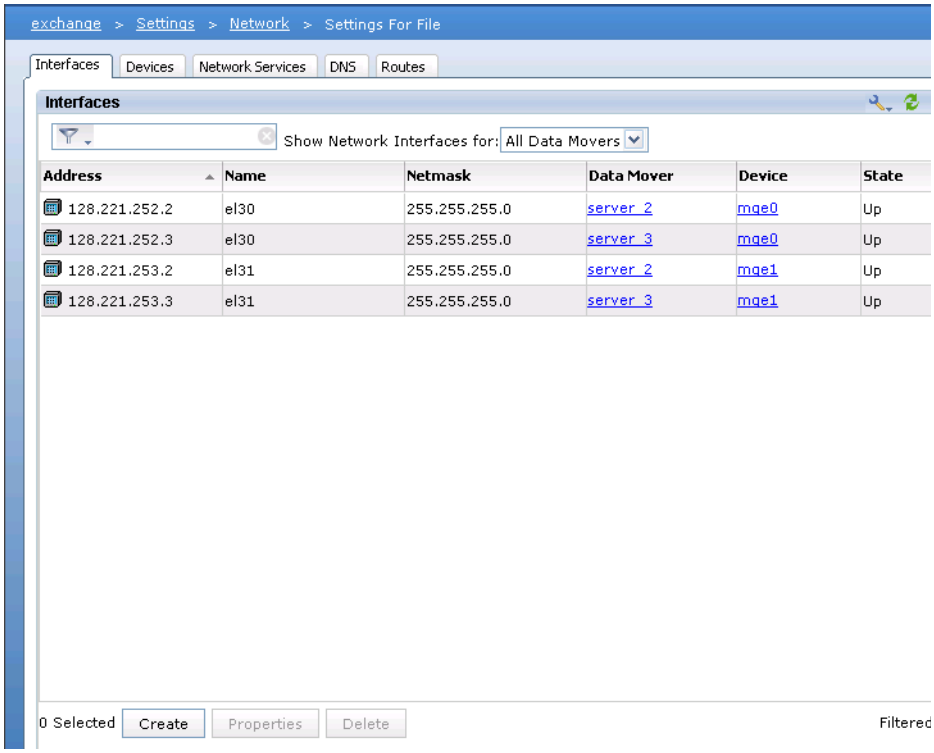

37. Select Data Mover as "server_2" from the drop-down list and select Device name as "lacp-1" from the drop down list. Enter the valid IP address, Netmask. Enter the Interface Name as "fs01" and MTU value as "9000" to allow jumbo frames for the lacp interface. Click **Ok**.

Figure 161 Entering Details to Create Network Interface

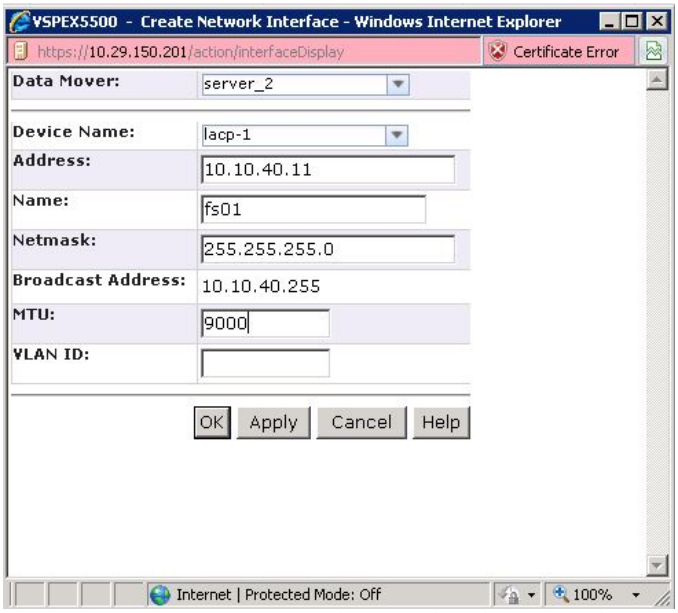

38. Make sure that the Network Interface “fs01” is created for the lacp device “lacp-1”.

Figure 162 New Network Interface is Created

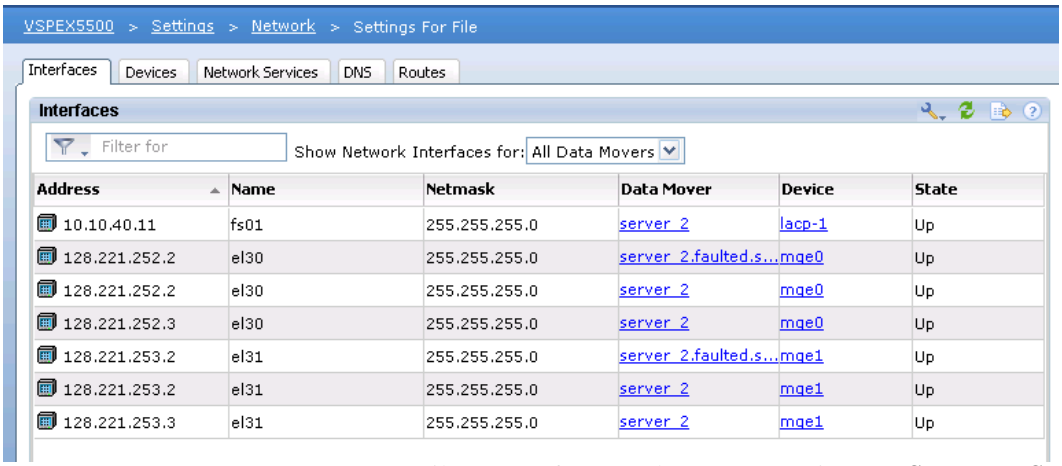

Address	Name	Netmask	Data Mover	Device	State
10.10.40.11	fs01	255.255.255.0	server_2	lacp-1	Up
128.221.252.2	el30	255.255.255.0	server_2.faulted.s...	mge0	Up
128.221.252.2	el30	255.255.255.0	server_2	mge0	Up
128.221.252.3	el30	255.255.255.0	server_2	mge0	Up
128.221.253.2	el31	255.255.255.0	server_2.faulted.s...	mge1	Up
128.221.253.2	el31	255.255.255.0	server_2	mge1	Up
128.221.253.3	el31	255.255.255.0	server_2	mge1	Up

39. To Create File system for NFS data store, Navigate to **Storage > Storage Configuration > File Systems** and Click **Create**. From the “Create File System” window, click **Storage Pool** radio button for “Create From” field and enter the File System Name as “NFS-OS” for 250 Virtual machine datastore. Select Storage Pool as “PerformancePool” from the drop down list. Enter Storage Capacity as “2 TB” for 250 VMs and Select Data Mover as “Server_2” from the drop-down list as shown in Figure 163. Click **Ok**, to create “NFS-OS” File system.

Figure 163 Entering Details to Create File System

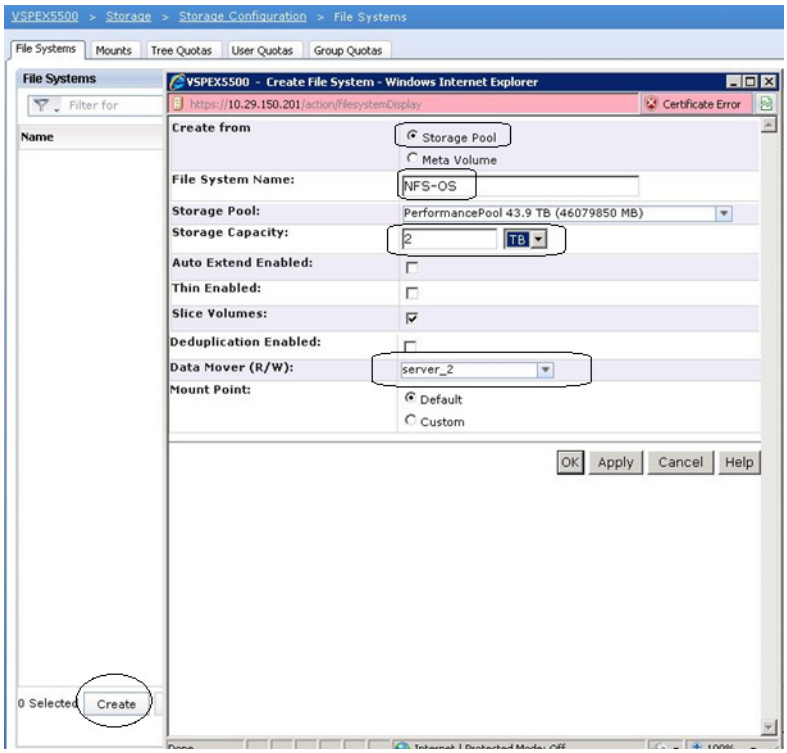

40. Wait until the "NFS-OS" File system creation process is complete. Verify the process using "Background Tasks for File".

Figure 164 Window Showing NFS-OS File System Creation in Progress

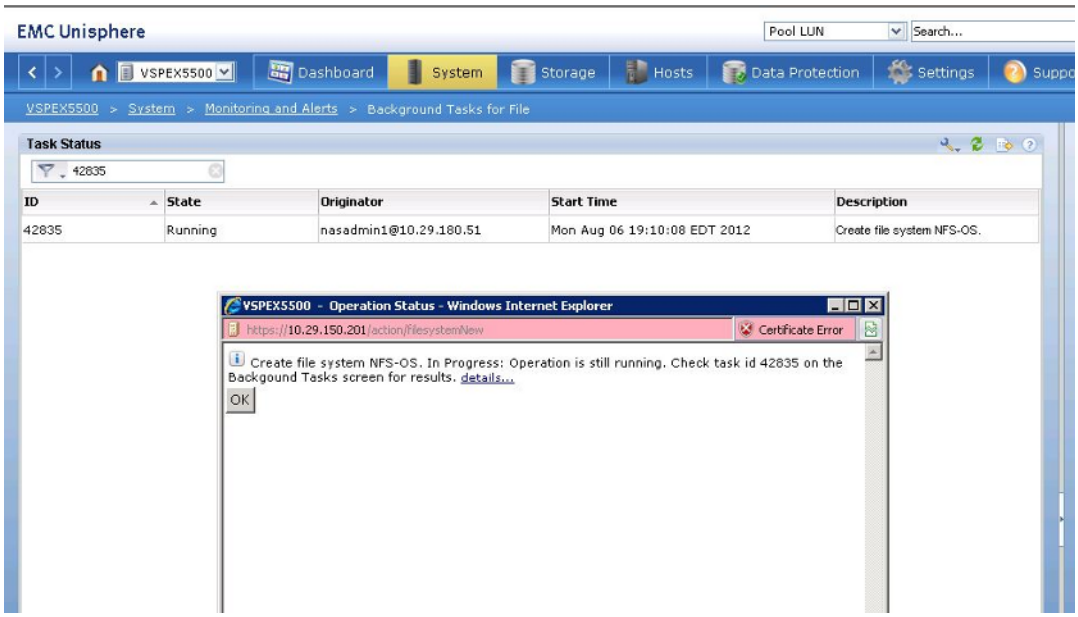

41. Make sure the File system "NFS-OS" is created with "2 TB" storage capacity as shown in Figure 165.

Figure 165 Window Showing NFS-OS Created with 2 TB Capacity

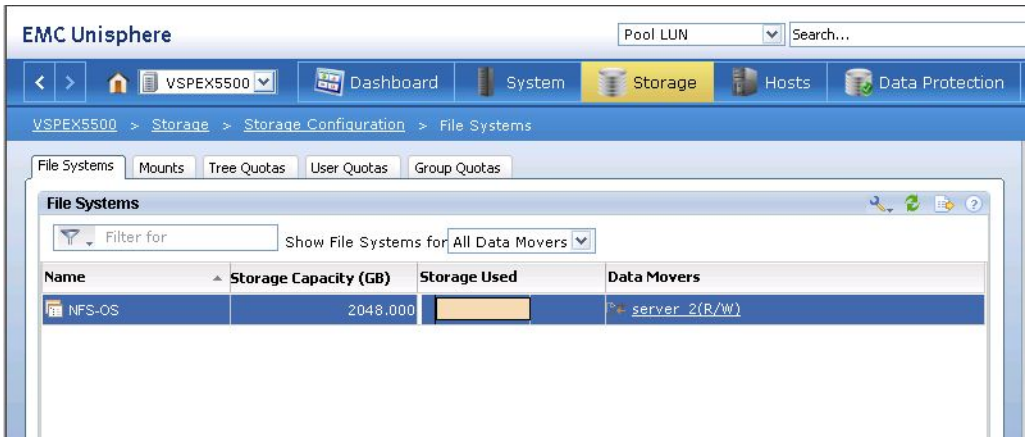

42. To validate the IO performance for 250 VMs, create 10xNFS-Data File systems with 2.5 TB capacity. From the "Create File System" window, click **Storage Pool** radio button for "Create From" field and enter the File System Name as "NFS-Data01" for 250 Virtual machine datastore. Select Storage Pool as "PerformancePool" from the drop down list. Enter Storage Capacity as "2.5 TB" for 250 VMs and Select Data Mover as "Server_2" from the drop-down list. Keep the Mount Point radio button at **Default**. Click **Ok** to create "NFS-Data01" File system.

Figure 166 Creating File System with 2.5 TB Storage Capacity

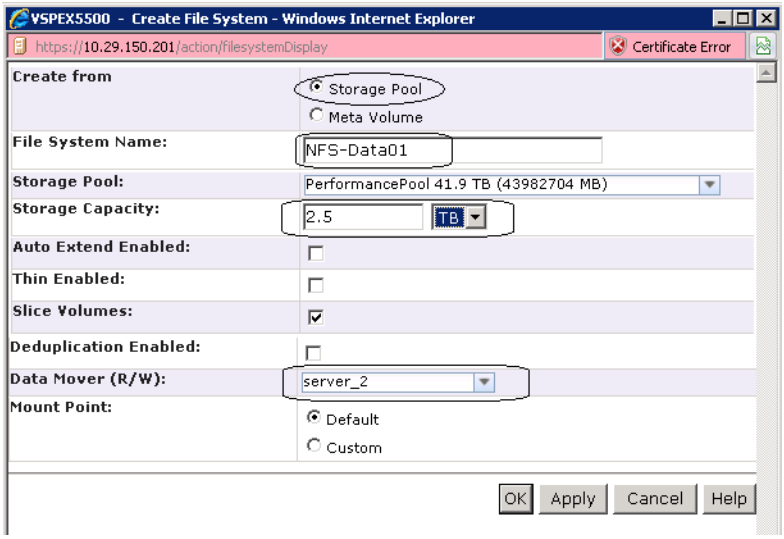

43. Follow Step 42, to create 9 more NFS Data file systems with 2.5 TB each.

Figure 167 Creating NFS Data File Systems at 2.5 TB Storage Capacity

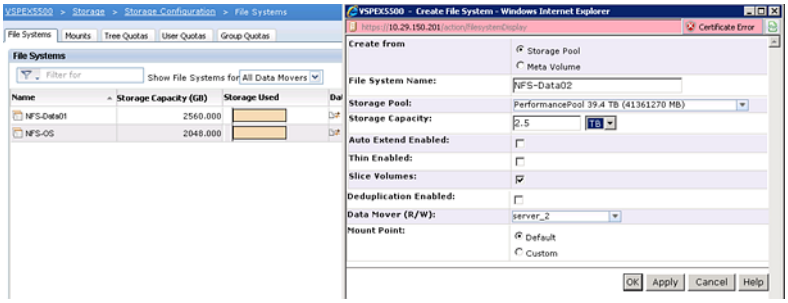

44. Make sure all the 10xNFS Data File systems are created as shown Figure 168.

Figure 168 Window Showing All the NFS Data File Systems Created

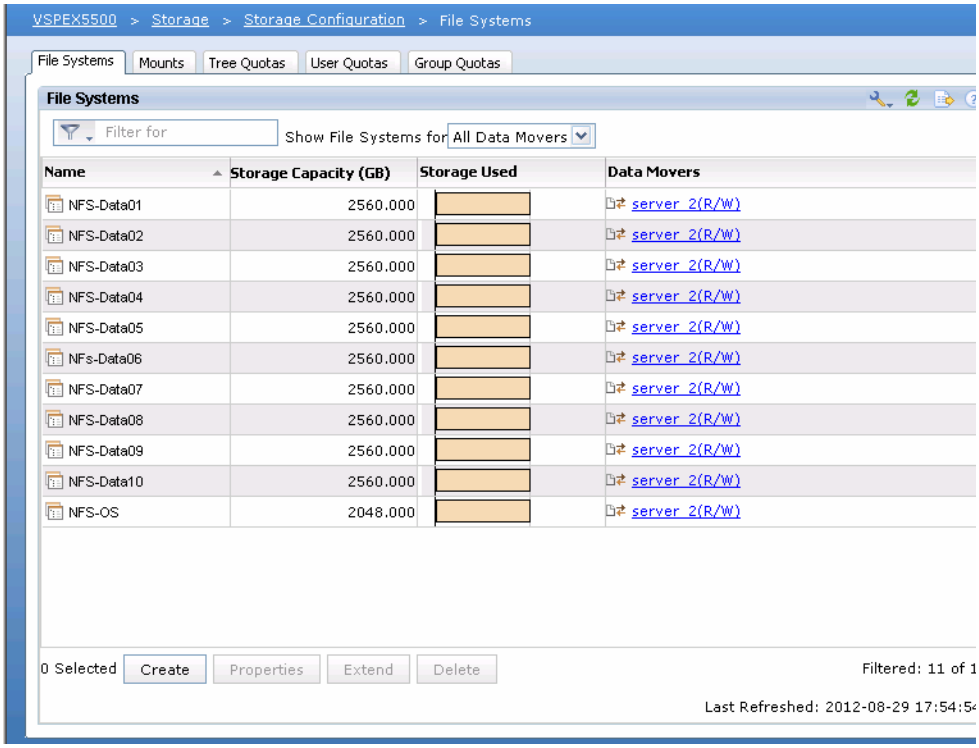

45. To enable “Direct Writes” for all the NFS File system. Select **Storage > Storage Configuration > File Systems**. Click **Mounts**.

Figure 169 Mounts Tab of File System Window

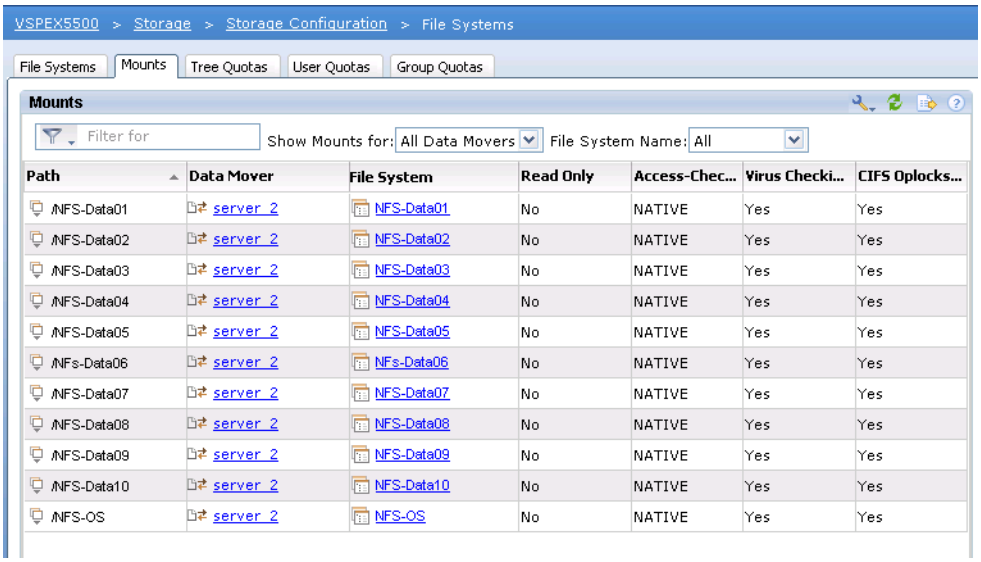

46. Select the path "/NFS-OS" for the file system "NFS-OS" and click **Properties**.

Figure 170 Window Showing the Path for NFS File Systems

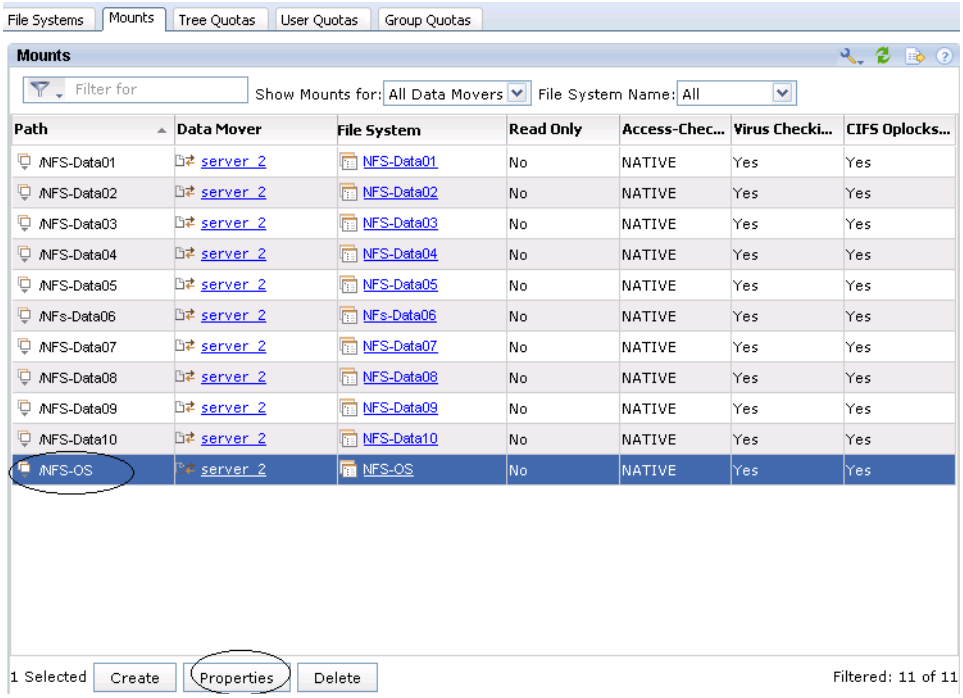

47. From the "/NFS-OS" mount properties. Make sure the radio button **Read/Write** for "Read Only" and **Native** for Access-Checking Policy is selected. Then, check the "Set Advanced Options" check box.

Figure 171 Enabling Read/Write for NFS-OS

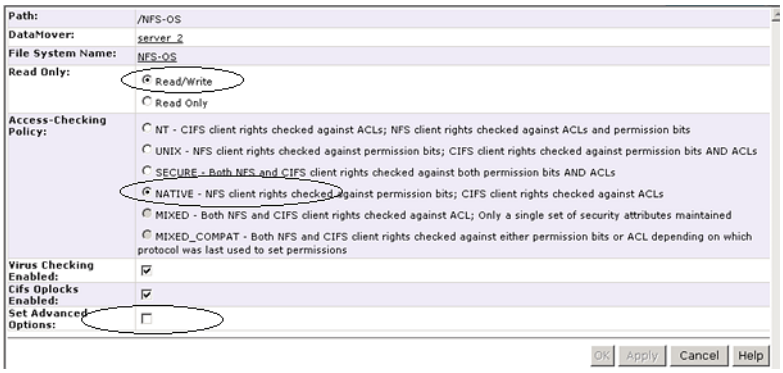

48. Check the "Advanced options" and the "Direct Writes Enabled" check boxes as shown in Figure 172 and Click **Ok**.

Figure 172 Enabling Parameters for NFS-OS

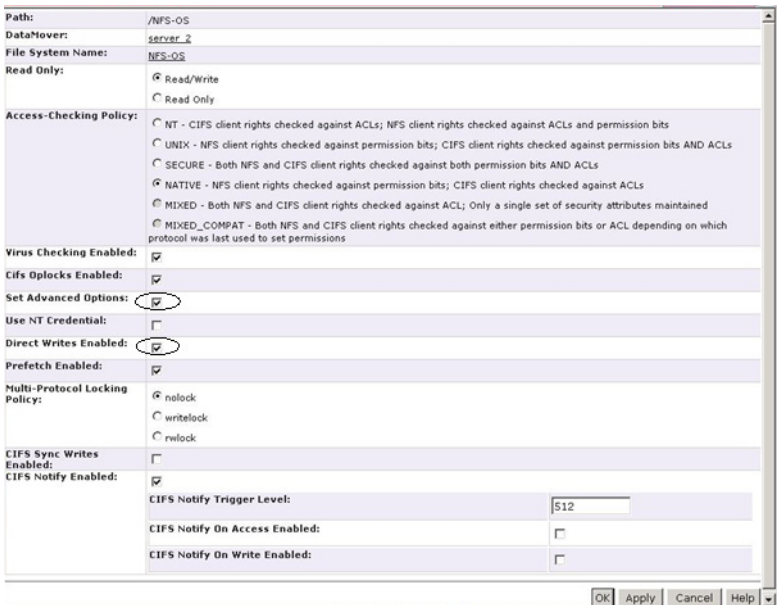

49. Follow the Steps 46, 47 & 48 to enable Direct Writes for all the remaining NFS Data file systems.

50. To Create NFS-Exports for all the NFS File systems. Click **Storage > Shared Folders > NFS** and Click **Create**.

Figure 173 Creating NFS Exports

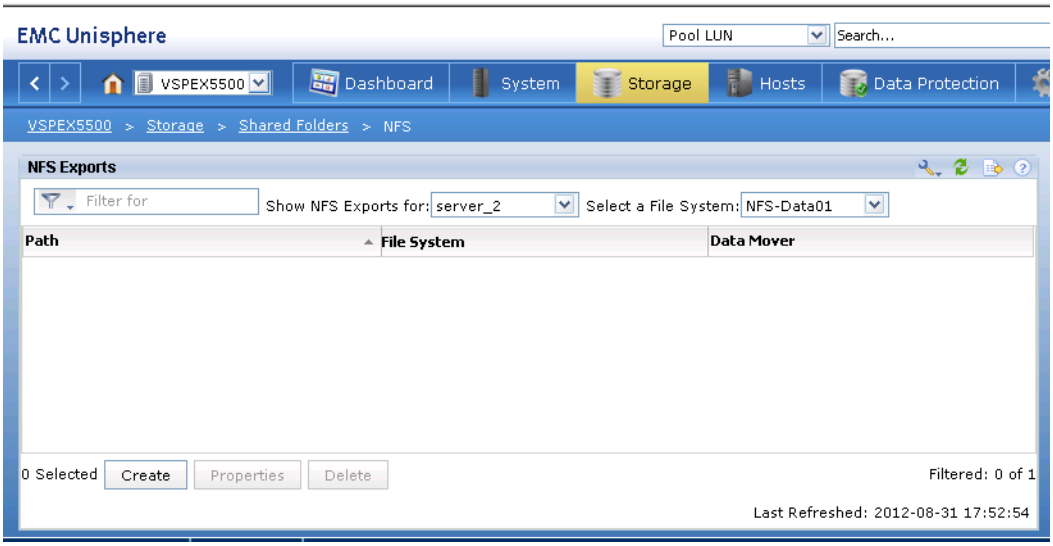

51. Select Data Mover as “server-2” from the drop-down list and select File System as “NFS-OS” and enter the Path as “/NFS-OS”. Enter the IP address of all the ESXi hosts “VMKernel Storage NIC” in “Root Hosts” and “Access Hosts” fields. Separate multiple host vmkernel IP's by “:” (colon) and Click **Ok**.

Figure 174 Entering Details for Creating NFS Exports

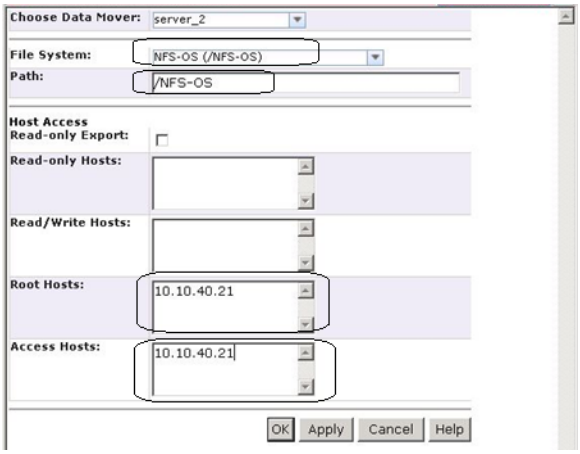

52. Make sure the NFS exports for “NFS-OS” file system is created as shown in Figure 175.

Figure 175 Window Showing Created NFS Export for NFS-OS

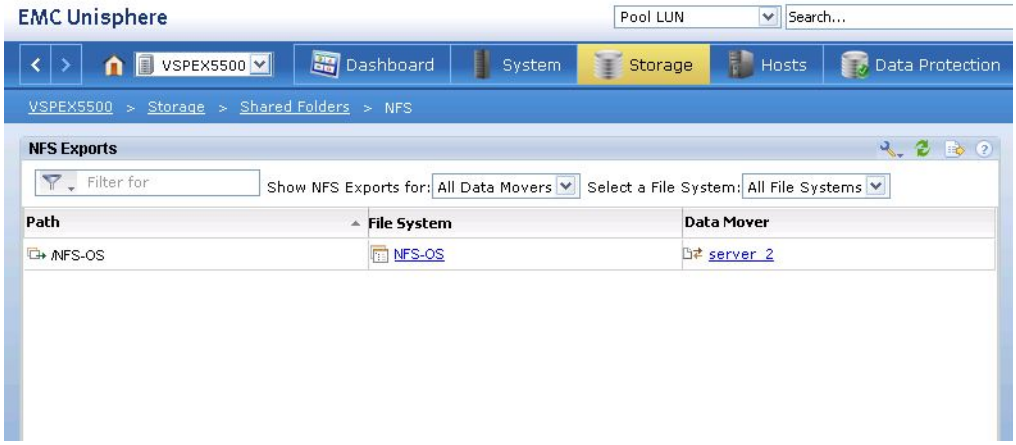

53. Repeat the Steps 50 and 51 to create NFS-Exports for all the remaining NFS-Data file systems.

Installing ESXi Servers and vCenter Infrastructure

1. In the Cisco UCS Manager window, click the **Equipment** tab, select a Cisco UCS B200 M3 Blade Server and click the **KVM Console** link to launch the KVM for the server on the right pane of the Cisco UCS Manager window.

Figure 176 Launch KVM Console

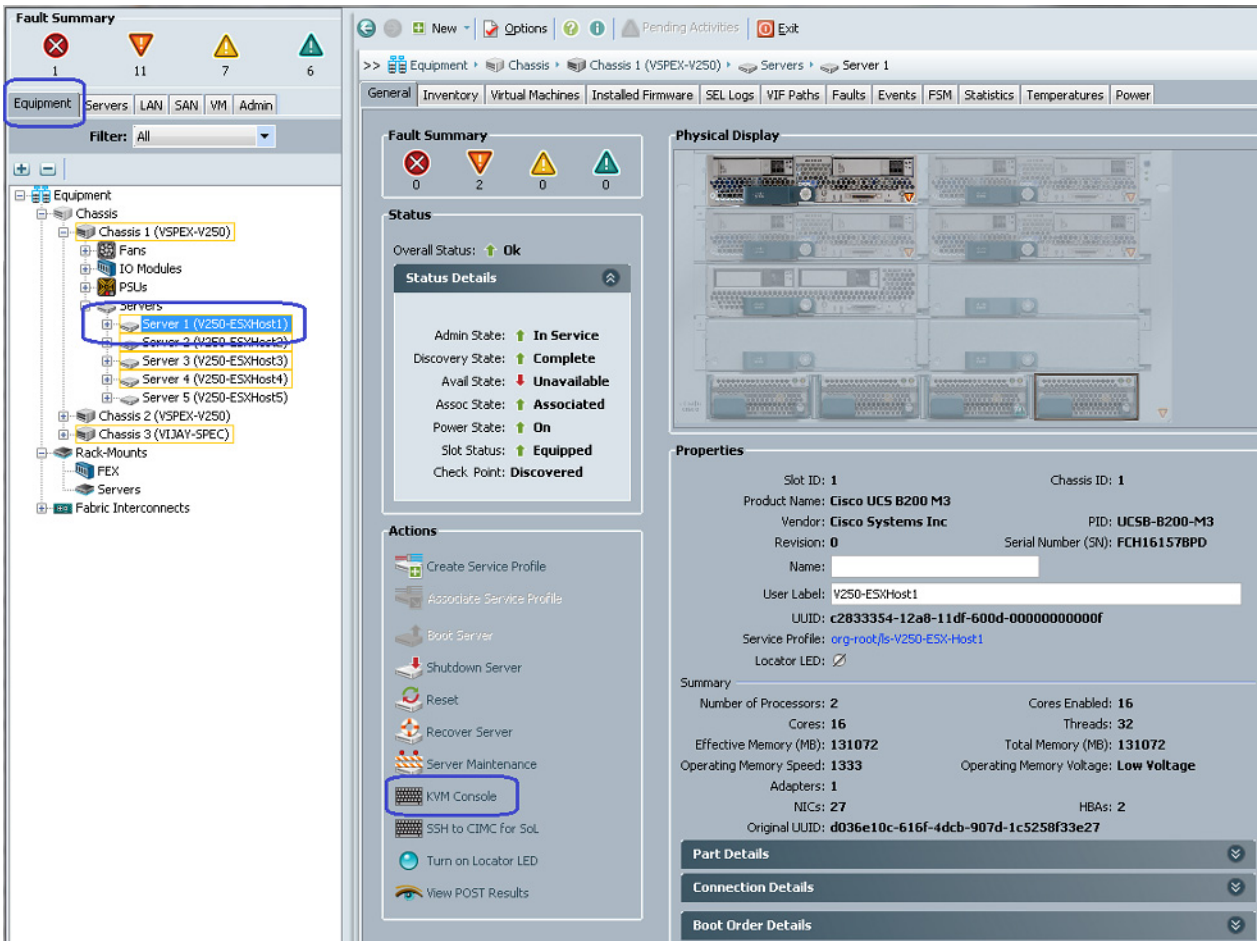

2. When the Java applet for KVM is launched, click the **Virtual Media** tab and click the **Add Image** tab as shown in Figure 177. This will open a dialog box to select an ISO image. Traverse the local directory structure and select the ISO image of ESXi 5.1 hypervisor installer media.

Figure 177 Adding Image in Virtual Media

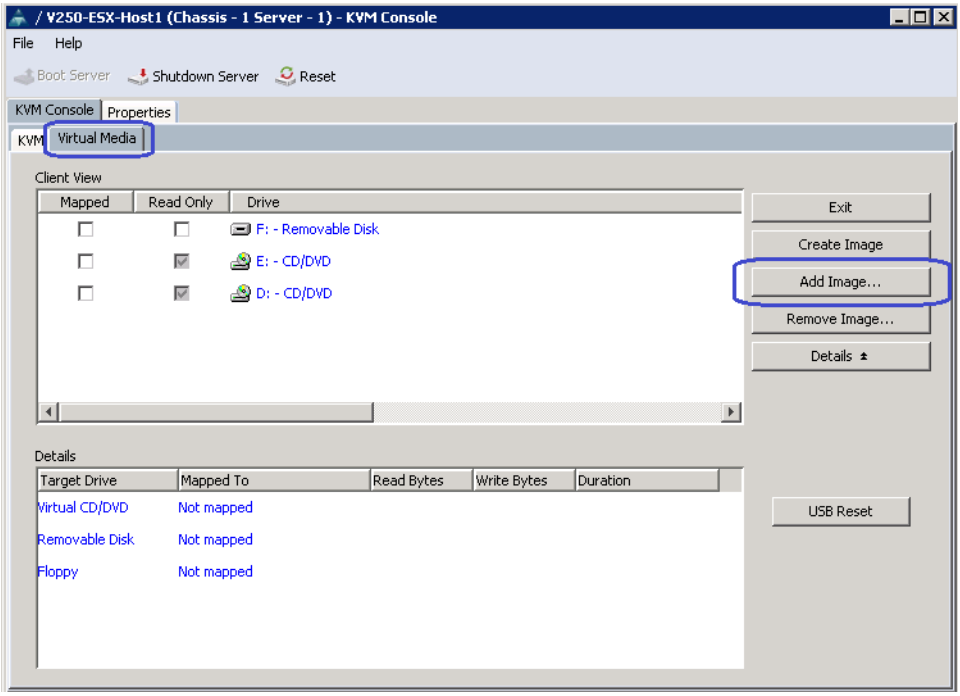

3. After the ISO image is shown in the list, check the "Mapped" check box and reset the server.

Figure 178 Reset the Server After Adding the ISO Image

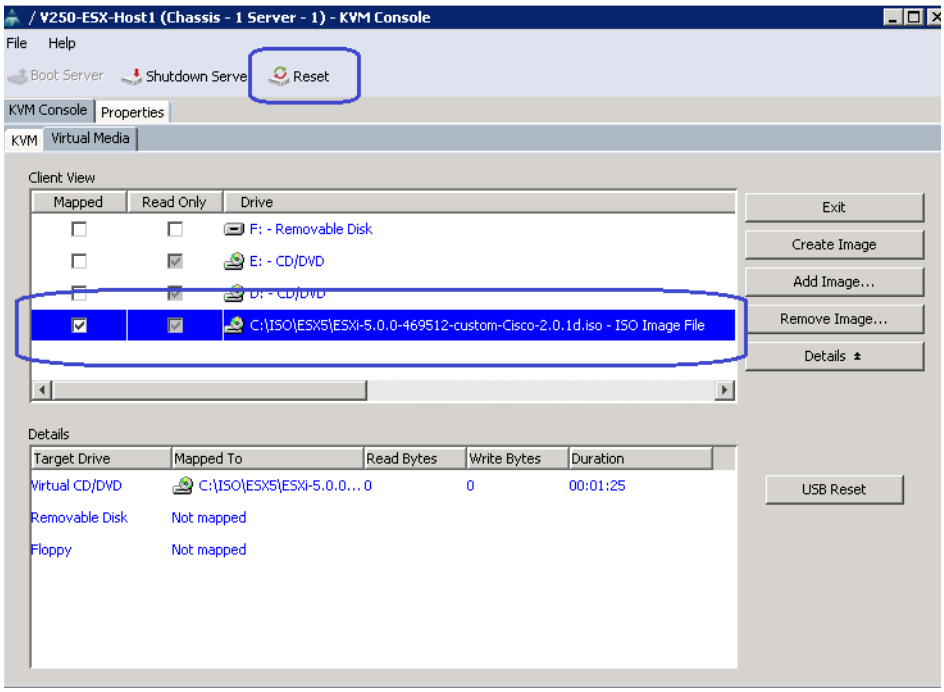

4. Click the radio button **Power Cycle** in the popup window, to immediately reboot the B200 M3 server as shown in Figure 180.

Figure 179 Selecting Power Cycle Option to Restart the Server

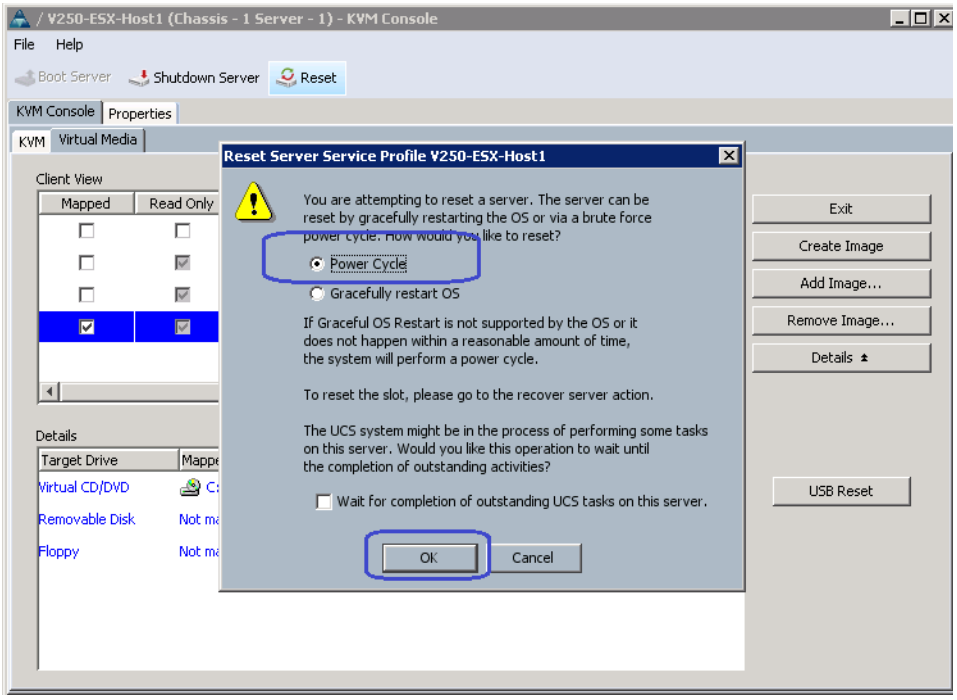

5. After rebooting the server, ESXi 5 install media will boot. Make sure to enter the following to install the hypervisor on each of the servers.
   - ESXi hostnames
   - IP addresses
   - Root password

   See, Customer Configuration Data Sheet, page 172 for appropriate values.
6. The ESXi OS should be installed on the SAN LUN of the B200 M3 servers. When the ESXi is installed, verify the network connectivity and accessibility of each server from each other.

Installing and Configuring Microsoft SQL Server Database

SQL server is used as database for the VMware vCenter server. Follow these steps to configure Microsoft SQL server.

1. Create a VM for Microsoft® SQL server—The requirements for processor, memory, and OS vary for different versions of SQL Server. To obtain the minimum requirement for each SQL Server software version, see the Microsoft technet link. The virtual machine should be created on one of the ESXi servers designated for infrastructure virtual machines, and should use the datastore designated for the shared infrastructure.

Note The customer environment may already contain an SQL Server that is designated for this role. For more information, see *Configure database for VMware vCenter.*

2. Install Microsoft® Windows on the VM—The SQL Server service must run on Microsoft Windows Server 2008 R2 SP1. Install Windows on the virtual machine by selecting the appropriate network, time, and authentication settings.

3. Install SQL server—Install SQL Server on the virtual machine from the SQL Server installation media. The Microsoft TechNet website provides information on how to install SQL Server.
4. Configure database for VMware vCenter—To use VMware vCenter in this solution, you will need to create a database for the service to use. The requirements and steps to configure the vCenter Server database correctly are covered in Preparing vCenter Server Databases. It is a best practice to create individual login accounts for each service accessing a database on SQL Server.

Note Do not use the Microsoft SQL Server Express-based database option for this solution.

5. Configure database for VMware Update Manager—To use VMware Update Manager in this solution you will need to create a database for the service to use. The requirements and steps to configure the Update Manager database correctly are covered in Preparing the Update Manager Database. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.
6. Deploy the VNX VAAI for NFS plug-in—The VAAI for NFS plug-in enables support for the vSphere 5.1 NFS primitives. These primitives reduce the load on the hypervisor from specific storage-related tasks to free resources for other operations. Additional information about the VAAI for NFS plug-in is available in the plug-in download vSphere Storage APIs for Array Integration (VAAI) Plug-in. The VAAI for NFS plug-in is installed using vSphere Update Manager. Refer process for distributing the plug demonstrated in the EMC VNX VAAI NFS plug-in - installation HOWTO video available on the www.youtube.com web site. To enable the plug-in after installation, you must reboot the ESXi server.

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 detailed information on Installing a vCenter Server, see the link:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html

For detailed information on vSphere Virtual Machine Administration, see the link:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html

For detailed information on creating a Virtual Machine in the vSphere 5.1 client, see the link:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.vm_admin.doc_50/GUID-0433C0DC-63F7-4966-9B53-0BECDDEB6420.html

Following steps provides high level configuration to configure vCenter server:

1. Create the vCenter host VM—If the VMware vCenter Server is to be deployed as a virtual machine on an ESXi server installed as part of this solution, connect directly to an Infrastructure ESXi server using the vSphere Client. Create a virtual machine on the ESXi server with the customer's guest OS configuration, using the Infrastructure server datastore presented from the storage array. The

memory and processor requirements for the vCenter Server are dependent on the number of ESXi hosts and virtual machines being managed. The requirements are outlined in the vSphere Installation and Setup Guide.

2. Install vCenter guest OS—Install the guest OS on the vCenter host virtual machine. VMware recommends using Windows Server 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. Create vCenter ODBC connection—Before installing vCenter Server and vCenter Update Manager, you must create the ODBC connections required for database communication. These ODBC connections will use SQL Server authentication for database authentication.

   For instructions on how to create the necessary ODBC connections see, vSphere Installation and Setup and Installing and Administering VMware vSphere Update Manager.
4. Install vCenter server—Install vCenter by using the VMware VIMSetup installation media. Use the customer-provided username, organization, and vCenter license key when installing vCenter.
5. Apply vSphere license keys—To perform license maintenance, log into the vCenter Server and select the Administration - Licensing menu from the vSphere client. Use the vCenter License console to enter the license keys for the ESXi hosts. After this, they can be applied to the ESXi hosts as they are imported into vCenter.

Configuring Cluster, HA and DRS on the vCenter

To add all the VMware on virtual machine vCenter, follow these steps:

1. Log into VMware ESXi Host using VMware vSphere Client.
2. Create a vCenter Datacenter.
3. Create a new management cluster with DRS and HA enabled.
   a. Right-click on the cluster and, in the corresponding context menu, click **Edit Settings**.
   b. Select the check boxes “Turn On vSphere HA” and “Turn On vSphere DRS”.
4. Click **Ok**, to save changes. Add all ESXi hosts to the cluster by providing servers' management IP addresses and login credentials one by one. After all the servers are added to the vCenter cluster, the window will look as shown in Figure 180.

Figure 180 Window Showing vCenter Cluster in VMware vSphere Client

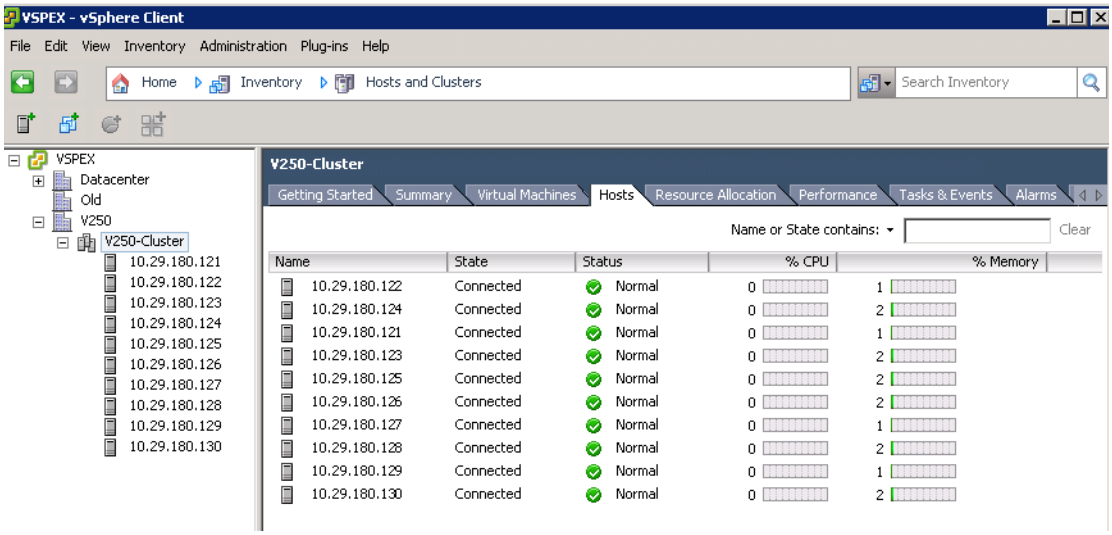

Configure Cisco VMFEX

Technology Overview section detailed about benefits of Cisco VMFEX technology. This section explains step by step configuration guide for Cisco UCS Manager/ vCenter management plane integration and Cisco VMFEX technology implementation. Follow these steps to configure Cisco VMFEX architecture.

1. Click the **VM** tab in the Cisco UCS Manager window, click **VMware** on the left pane of the Cisco UCS Manager window and click the **Modify Extension Key** link on right pane of the Cisco UCS Manager window as shown in Figure 181.

Figure 181 Modifying Extension Key in the Cisco UCS Manager Window

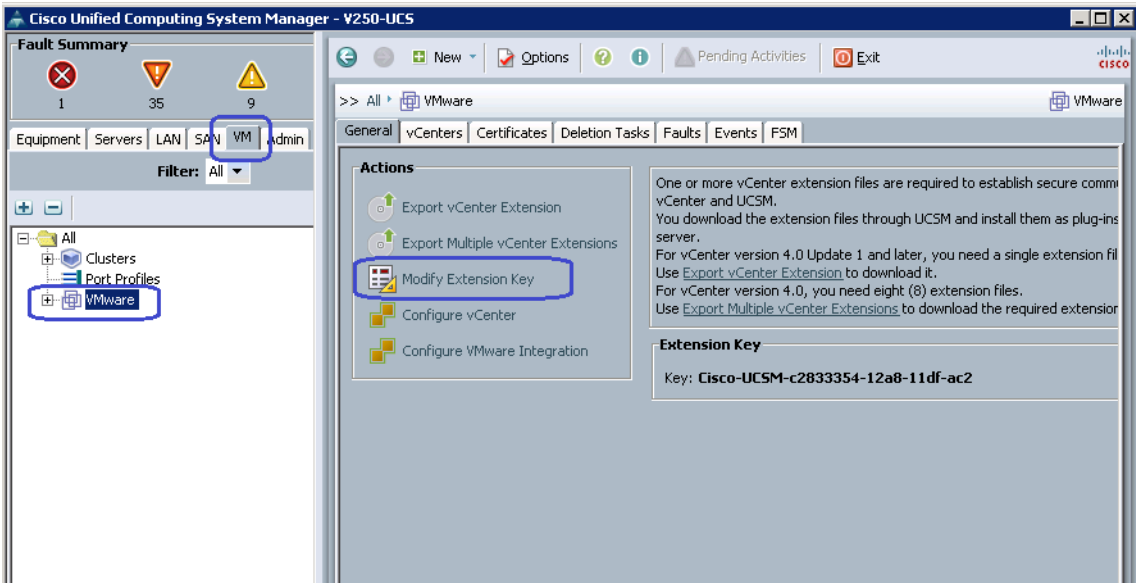

2. Change the default extension key to a value that represents the UCS pod used in this solution as shown in Figure 182 and click **Ok**.

Figure 182 Modifying Extension Key

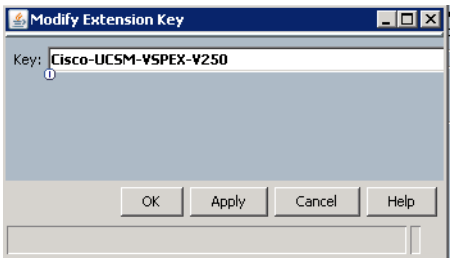

3. To establish trusted relationship between the Cisco UCS Manager and vCenter. Click the **Export vCenter Extension** link on the right pane in the **General** tab of VMware as shown in Figure 183.

Figure 183 Exporting vCenter Extension

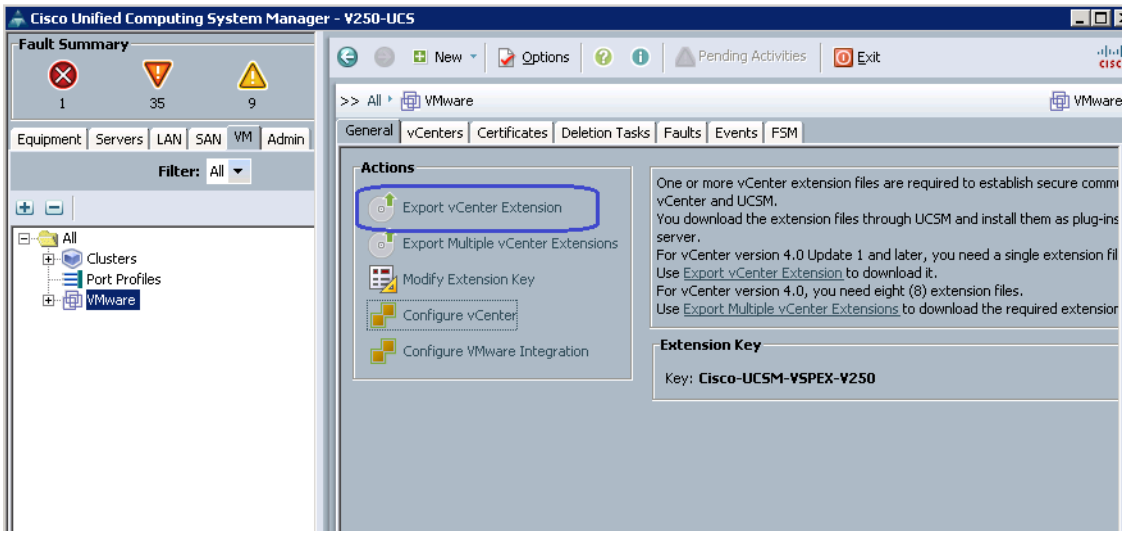

4. Specify the location where the vCenter extension XML file should be saved on the popup window as shown in Figure 184. Click **Ok**.

Figure 184 Specifying Path for vCenter Extension File

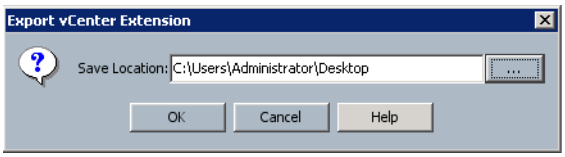

5. Using vSphere 5.1 client application, connect to the vCenter 5.1 server, click the **Plug-ins** tab in the menu bar, and click **Manage Plug-ins…** as shown in Figure 185.

Figure 185 Managing Plug-ins in VMware vSphere Client

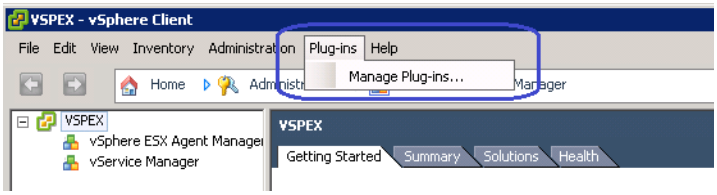

6. Right-click on the whitespace after the list of installed plug-ins populates, and click **New Plug-In…** as shown in Figure 186.

Figure 186 Creating New Plug-in—Plug-in Manager Window

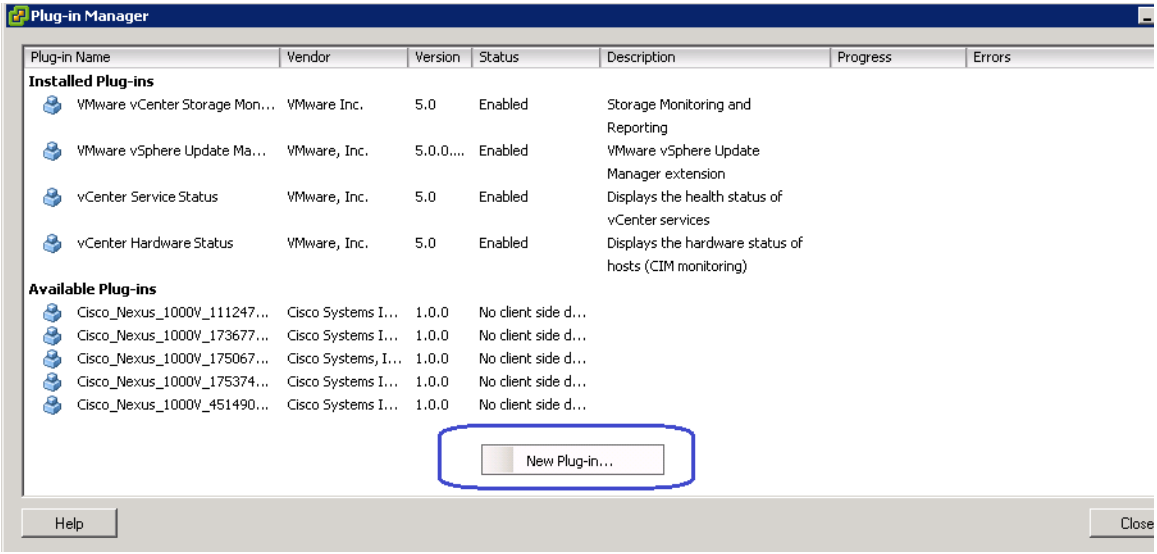

7. In the “Register Plug-In” window, browse to the Cisco UCS Manager extension XML file and select it. Make sure that the extension key set in step 1 shows up in the <key> tag in this window, and then click **Register Plug-In** as shown Figure 187.

Figure 187 Registering the Plug-in

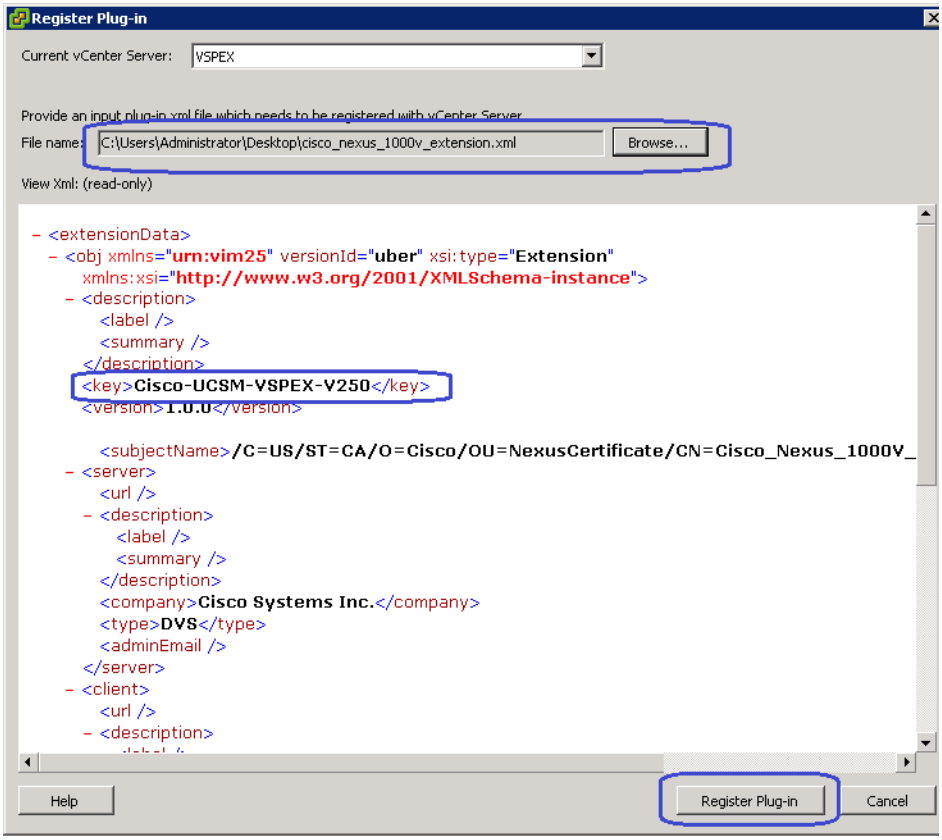

8. Given that Cisco UCS Manager has self-signed SSL certificate, you may see an untrusted certificate warning. Ignore the warning. After that, you should see a success notification as shown in Figure 188.

Figure 188 Window Showing the Plug-in Registered in vCenter

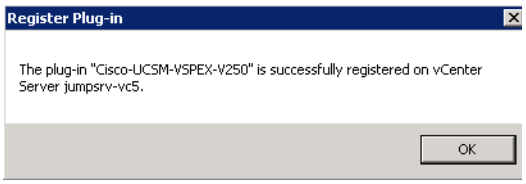

9. Cisco UCS Manager plug-in should be listed now in the “Available Plug-ins” list as shown Figure 189.

Figure 189 Cisco UCS Manager Plug-in Listed in the Available Plug-ins

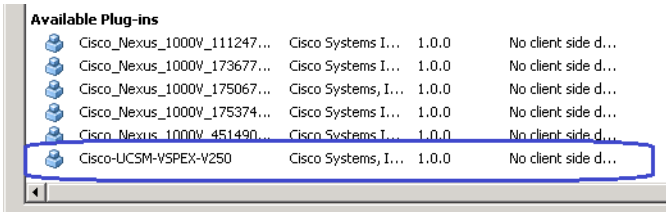

10. Trust relationship is now established between the Cisco UCS Manager and vCenter. In the VM tab of the Cisco UCS Manager and click **Configure vCenter** link on the right pane of the Cisco UCS Manager window as shown Figure 190.

Figure 190 Configuring vCenter in Cisco UCS Manager

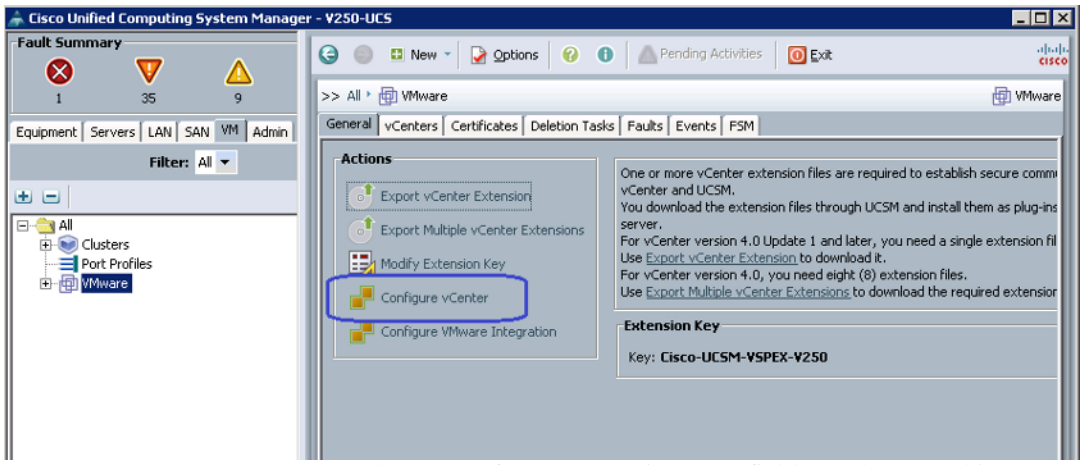

11. Enter the name of the vCenter in Name field (can be any arbitrary name), provide description (optional), and Host Name as hostname or the dotted decimal IP address of the vCenter host as shown in Figure 191. Click **Next**.

Figure 191 Entering Details in the Configure vCenter Wizard

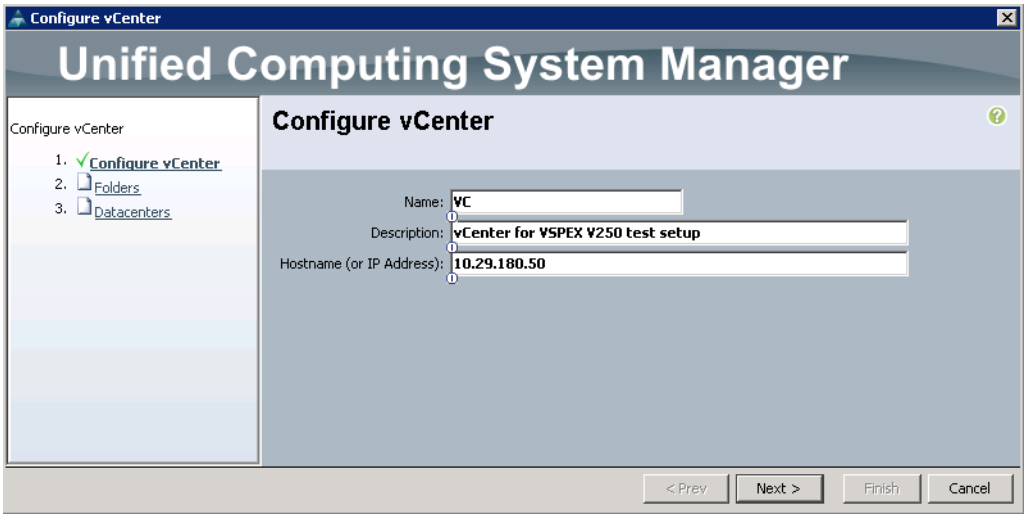

12. If your datacenter on vCenter is in a folder then you need to create same folder name in the next window. In our case, the datacenter is not contained in a folder, so simply click **Next** on this window as shown in Figure 192.

Figure 192 Folders Window of the Configure vCenter Wizard

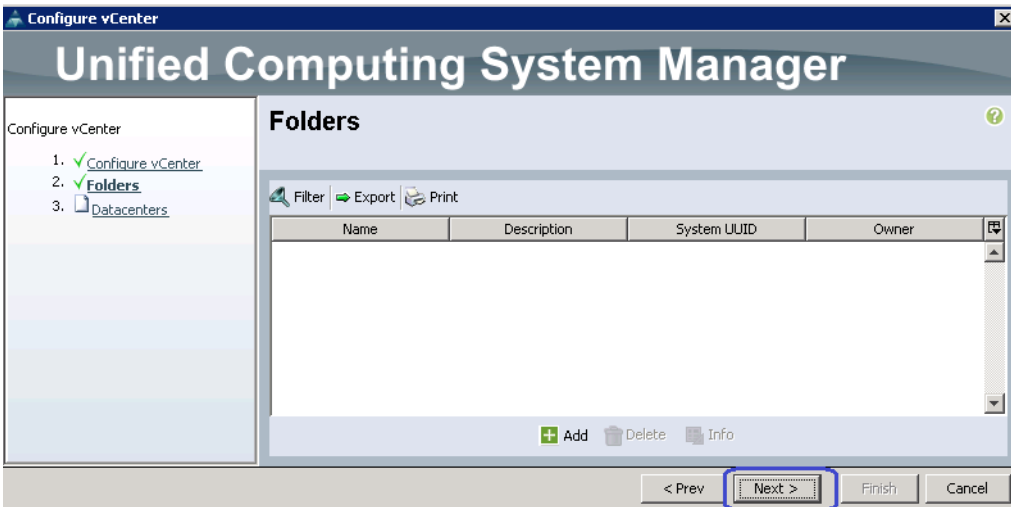

13. In the "Datacenters" window, click **Add** .

Figure 193 Datacenters Window of the Configure vCenter Wizard

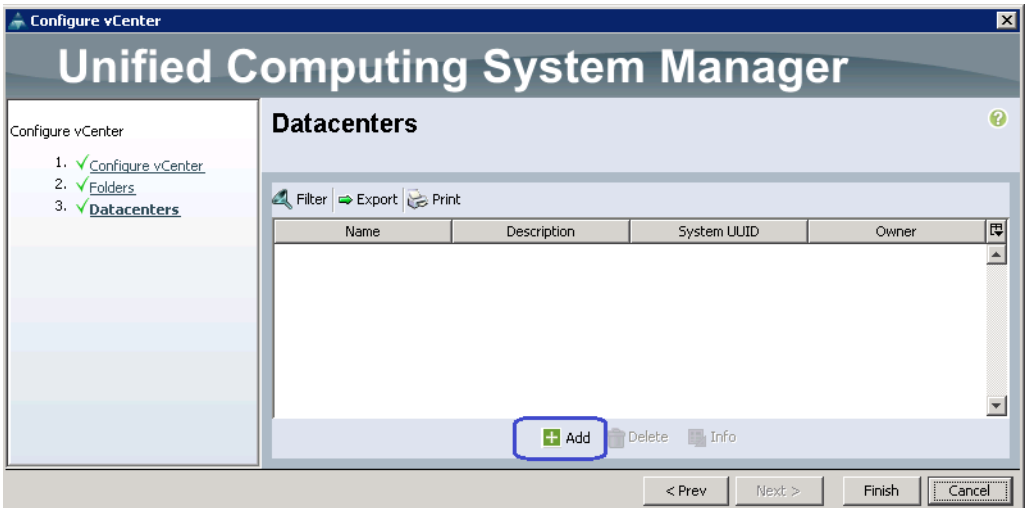

14. Enter the name of the Datacenter in the vCenter. This name must match exactly as that given in the vCenter. Description is optional. Click **Next**.

Figure 194 Identifying the Datacenter

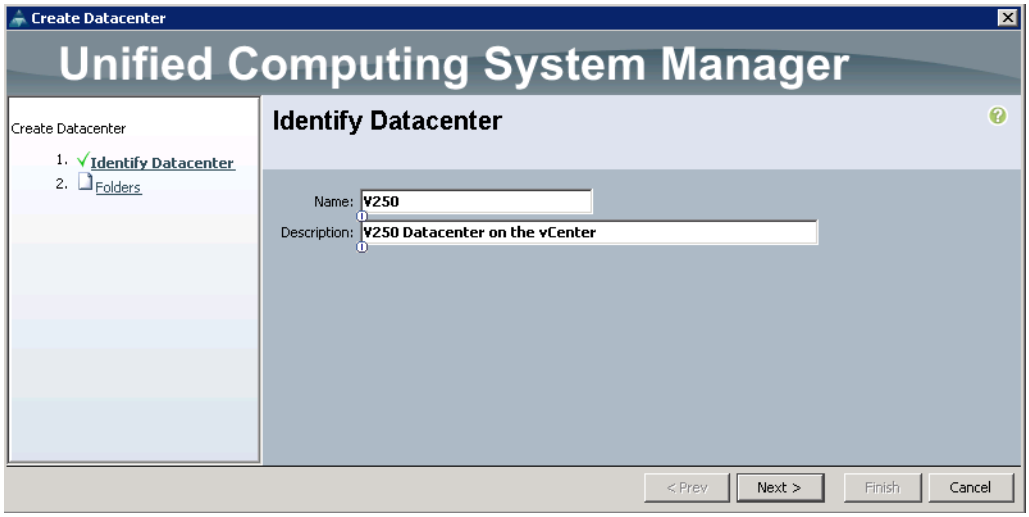

15. Now, create a folder that would contain the virtual Distributed Switch (vDS). click **Add** on this window as shown in Figure 195.

Figure 195 Adding the Datacenter

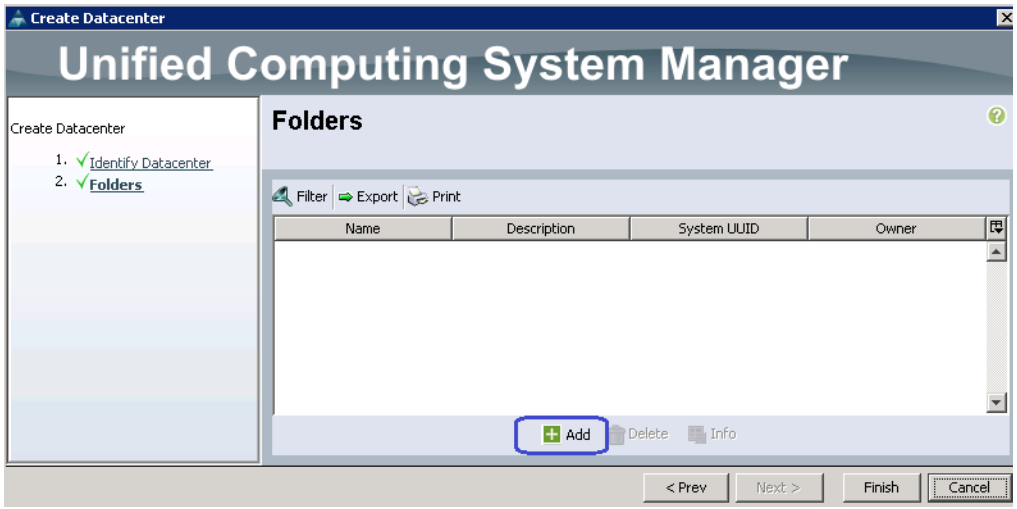

16. Enter folder name and description (optional). Click **Next**.

Figure 196 Creating Folder

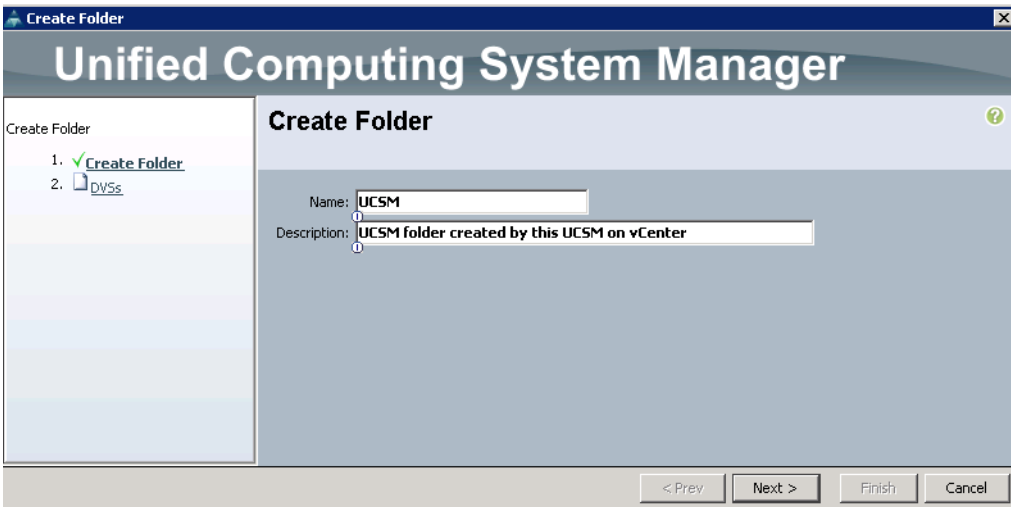

17. Click **Add** in the DVSs window, to add a Distributed Virtual Switch.

Figure 197 Adding a Distributed Virtual Switch

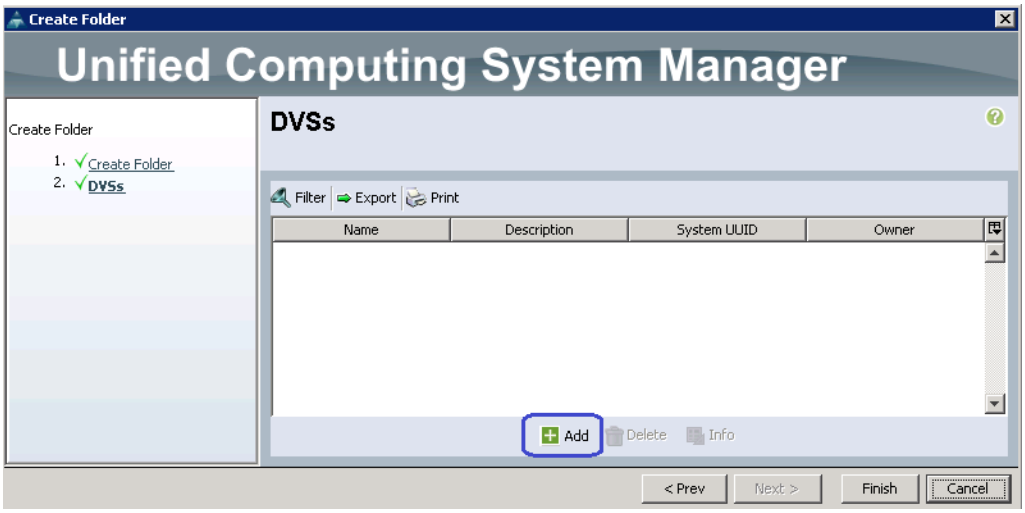

18. Enter the name of the DVS, description (optional) and click the radio button **Enable** for the "Admin State" of the DVS. Click **Ok**.

Figure 198 Creating a Distributed Virtual Switch

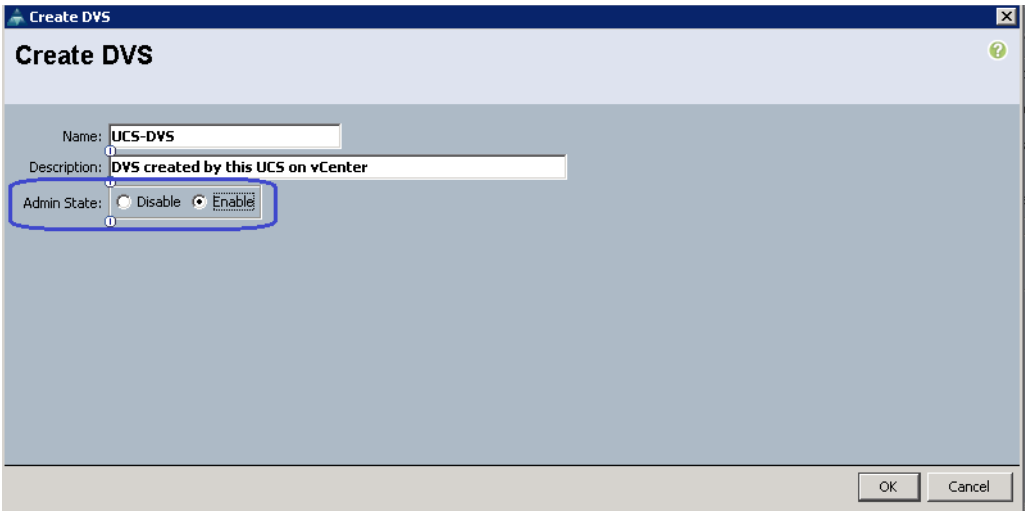

19. Click **Finish** in the "DVSs" window.

Figure 199 Wizard Showing Created UCS-DVS

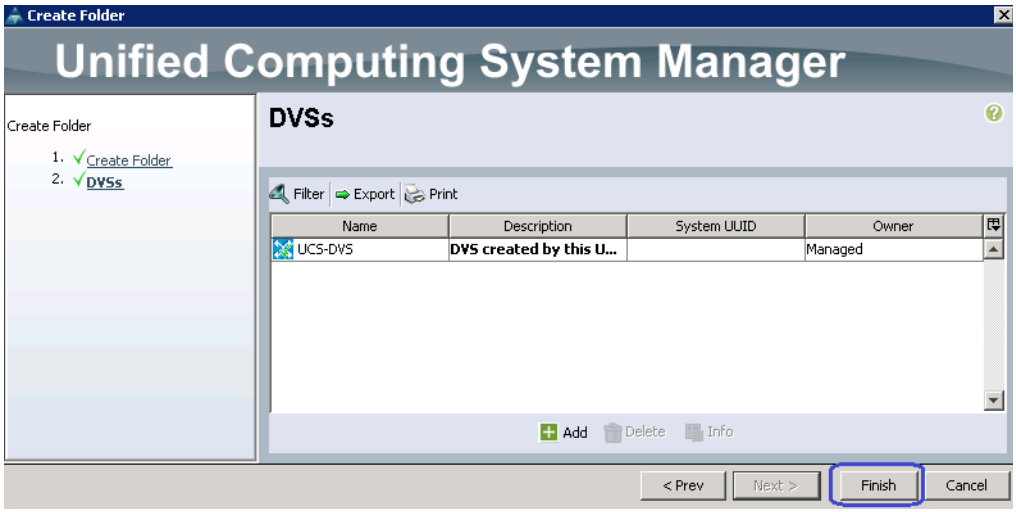

20. Click **Finish** in the "Folders" window.

Figure 200 Wizard Showing Created Folder Cisco UCS Manager

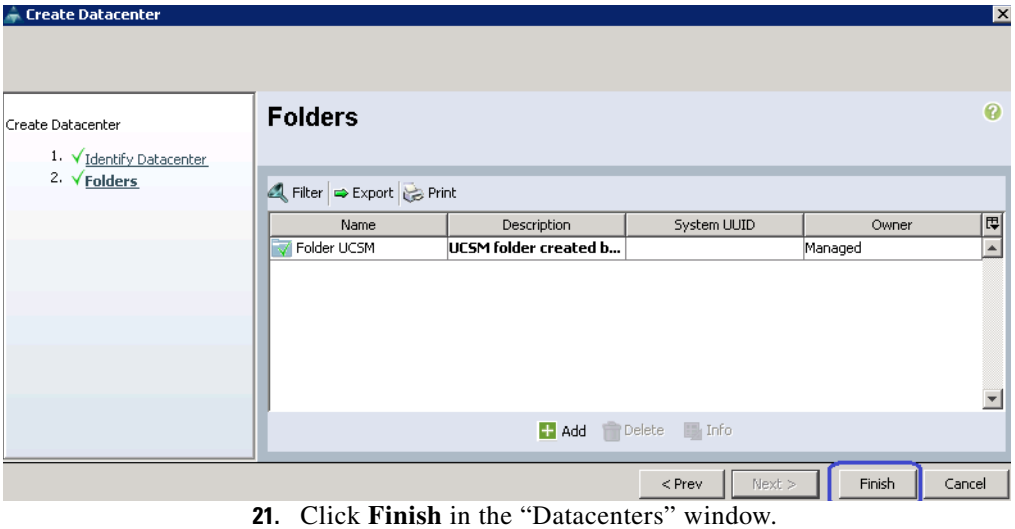

21. Click **Finish** in the "Datacenters" window.

Figure 201 Wizard Showing Created Datacenter V250

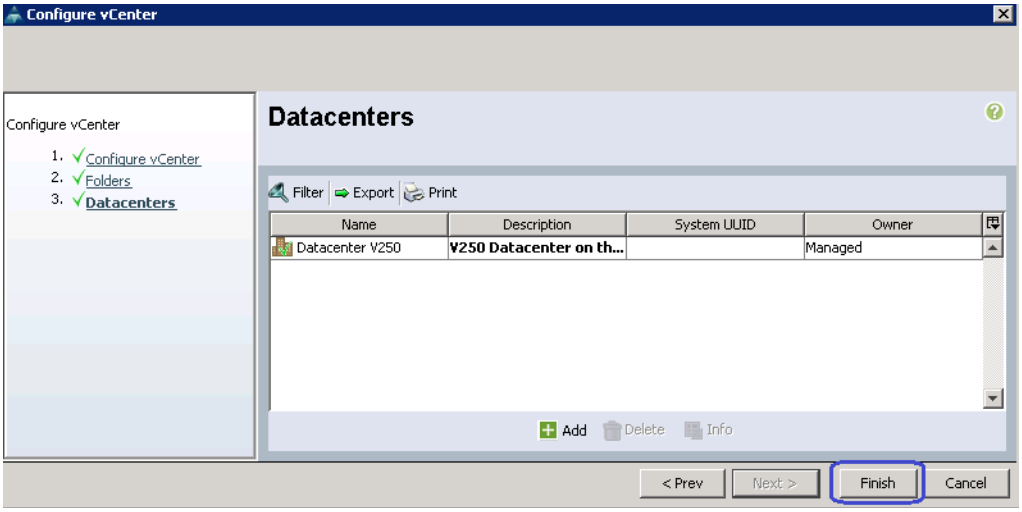

22. You will get a success notification popup window as shown Figure 202.

Figure 202 Windows Showing vCenter VC Created Successfully

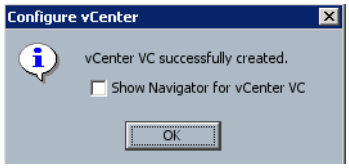

23. In the vCenter window, click **Inventory** > **Networking**, you should see the folder and DVS created, with two default port-profiles "uplink-pg-<vDS-Name>" and "deleted-pg-<vDS-Name>" as shown in Figure 203.

Figure 203 vCenter Window Showing the Folder and DVS Created

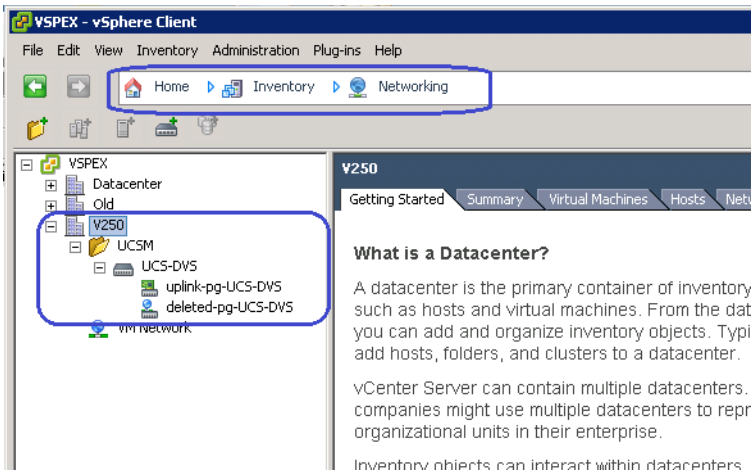

24. In the **VM** tab of the Cisco UCS Manager window, right-click on the "Port Profiles" and click **Create Port Profile** as shown in Figure 204.

Figure 204 Creating Port profile in Cisco UCS Manager

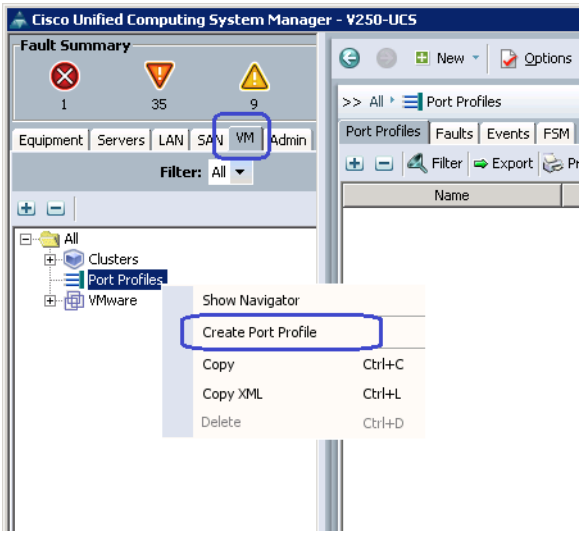

25. Create an infrastructure port profile. Provide description (optional). The name and description would show up on vCenter when the Cisco UCS Manager pushes the port profile to the vCenter. You can restrict maximum ports for the infrastructure port profile to 10. Select infra VLAN as part of the allowed VLANs list and mark it as native VLAN. Figure 205 shows infra port profile configuration. Click **Ok**.

Figure 205 Entering Details to Create Port Profile

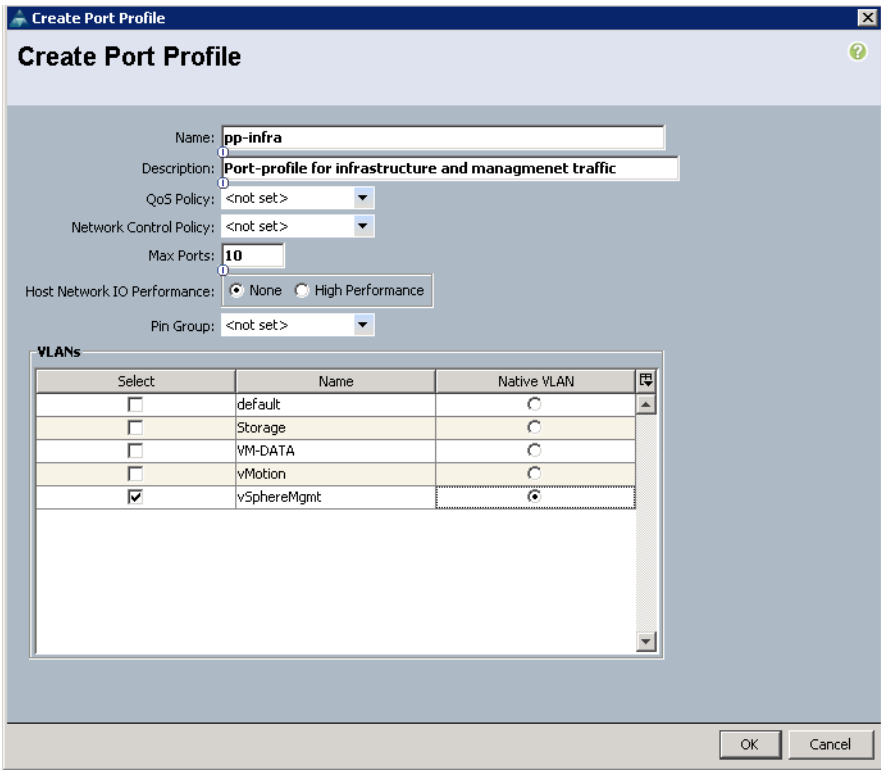

26. Select the newly created port profile and click **Create Profile Client**. In the Cisco UCS Manager window you can configure multiple vCenter and create multiple (up to 8) vDS per vCenter. A given port profile will be pushed to a set of vDS based on port profile client regex match. As we have only one vCenter and one vDS, we will simply push the port profile to **all** vDS in next step.

Figure 206 Creating Profile Client in Cisco UCS Manager

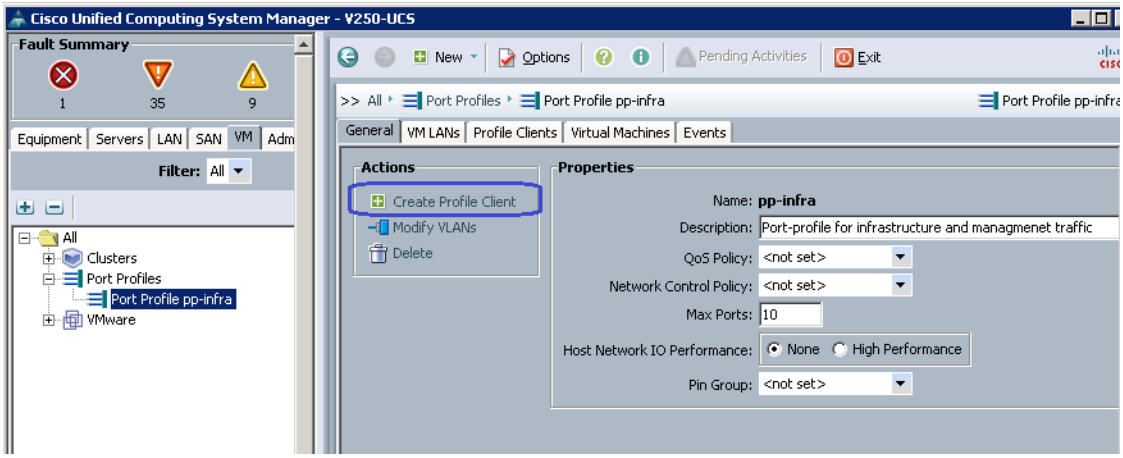

27. Default parameters for port profile client are “all”, and we will retain that. To reflect that, let us name the port profile client “all” as sown in Figure 207. Click **Ok**.

Figure 207 Entering Details for Creating Profile Client

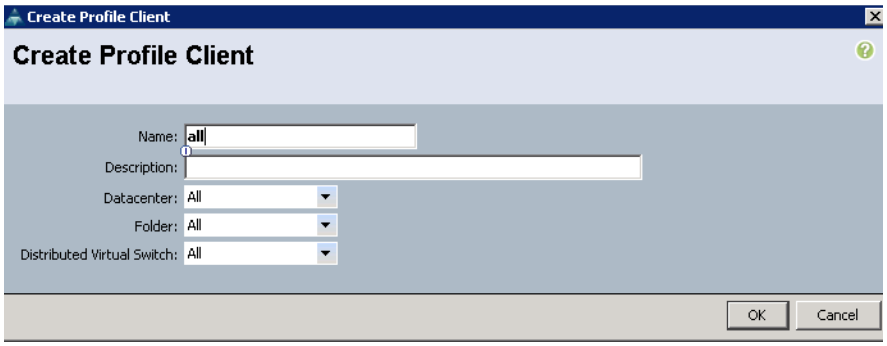

28. Similarly, create a "vMotion" port profile as shown Figure 208. Make sure that the "vMotion" is selected for the QoS policy field.

Figure 208 Creating Port Profile

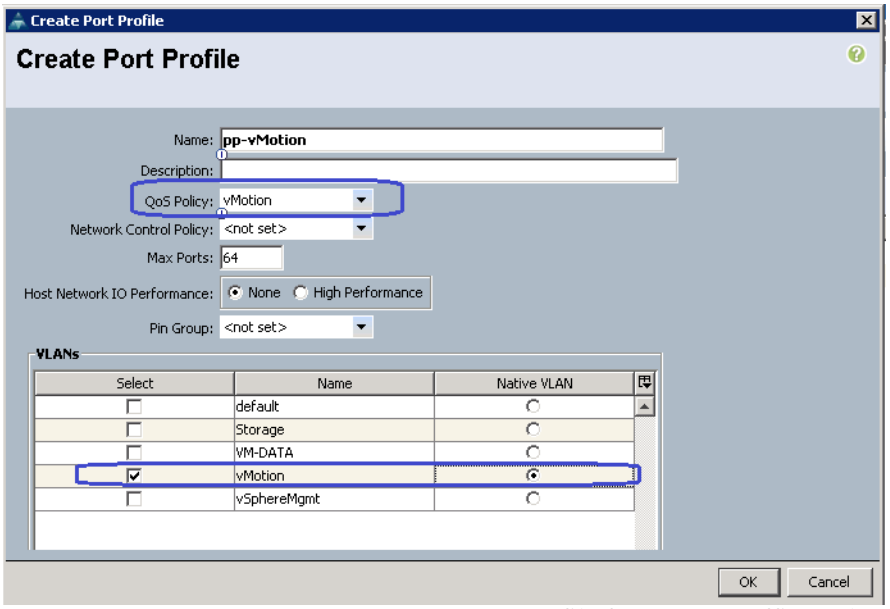

29. Create an NFS port profile for storage traffic as shown in Figure 209.

Figure 209 Creating Port Profile for Storage Traffic

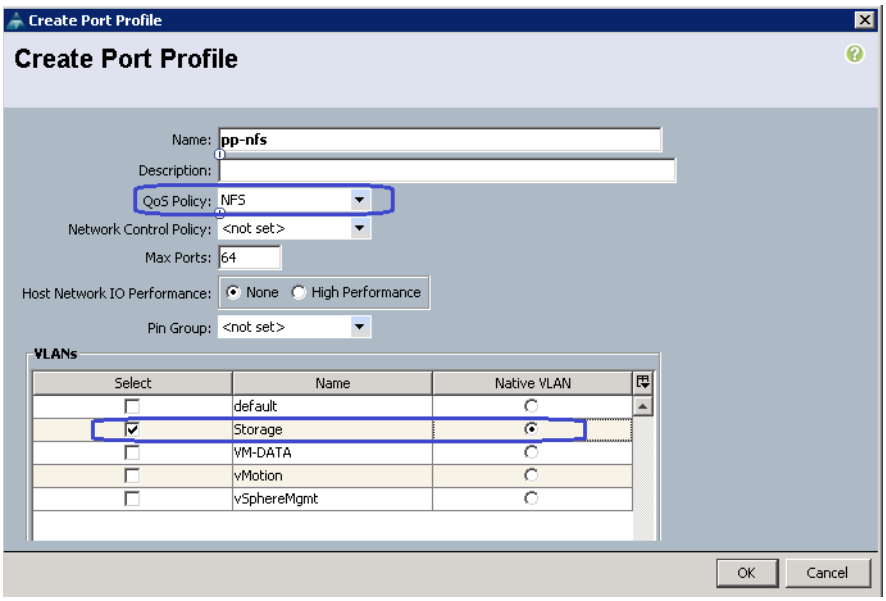

30. Figure 210 shows a sample VM application/ data port profile.

Figure 210 Window Showing Sample Port Profile Created

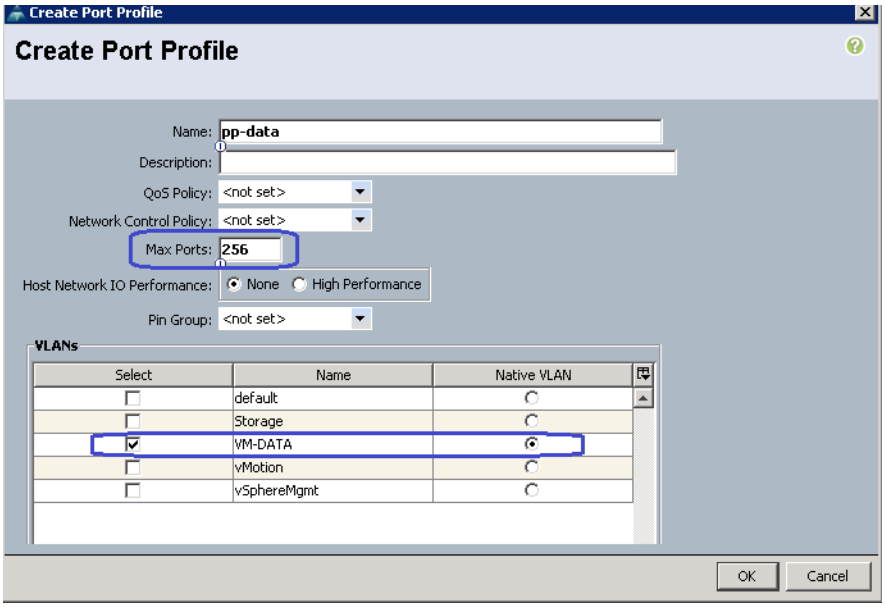

31. Create “all” port profile clients for port profiles created in steps 28 to 30. This will be shown in the vCenter server.

Figure 211 UCS-DVS Showing All the Created Port Profiles

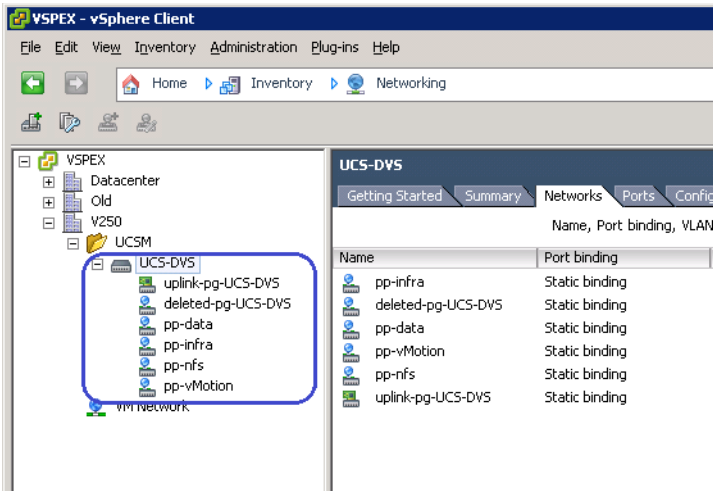

32. In the vCenter server, click **Inventory** > **Networking** in the menu bar. Right-click on the vDS created by Cisco UCS Manager and click **Add Host…** as shown in Figure 212.

Figure 212 Adding Host in UCS-DVS

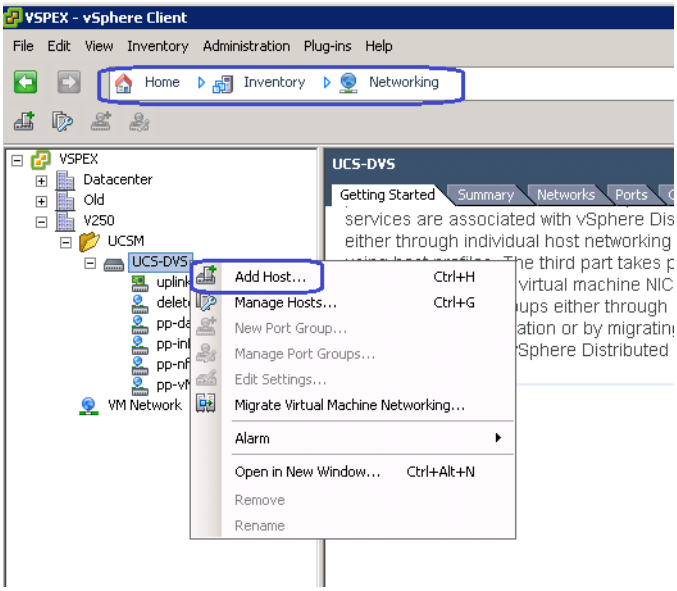

33. Select all the ESXi 5.1 hosts and all the uplinks on the servers. There is only one implicit uplink port profile created by Cisco UCS Manager and that uplink port profile is automatically selected for the migration to vDS from vSwitch as shown in Figure 213. Click **Next**.

Figure 213 Selecting Hosts and Physical Adapters

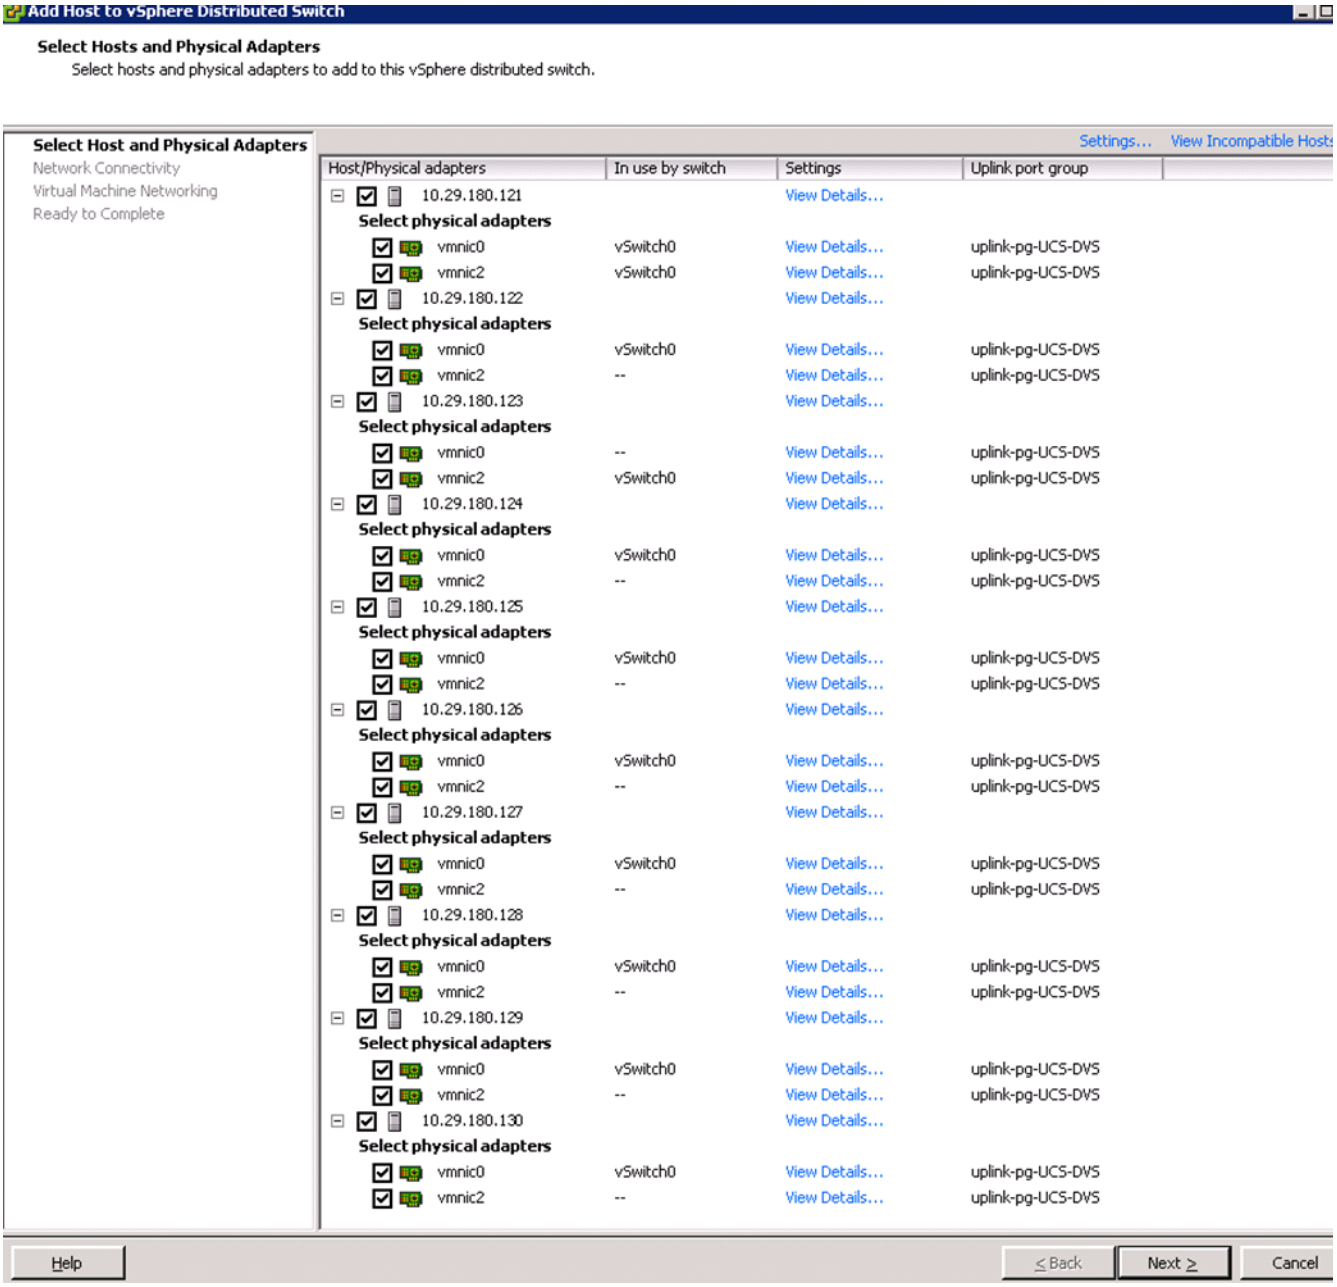

34. As we are migrating both the uplinks to vDS, any traffic going to native vSwitch will be "black-holed". Migrate the ESXi kernel management ports to the vDS. Choose the appropriate infrastructure port profiles for all the management kernel interfaces as shown in Figure 214.

Figure 214 Network Connectivity Window in Adding Hosts Wizard

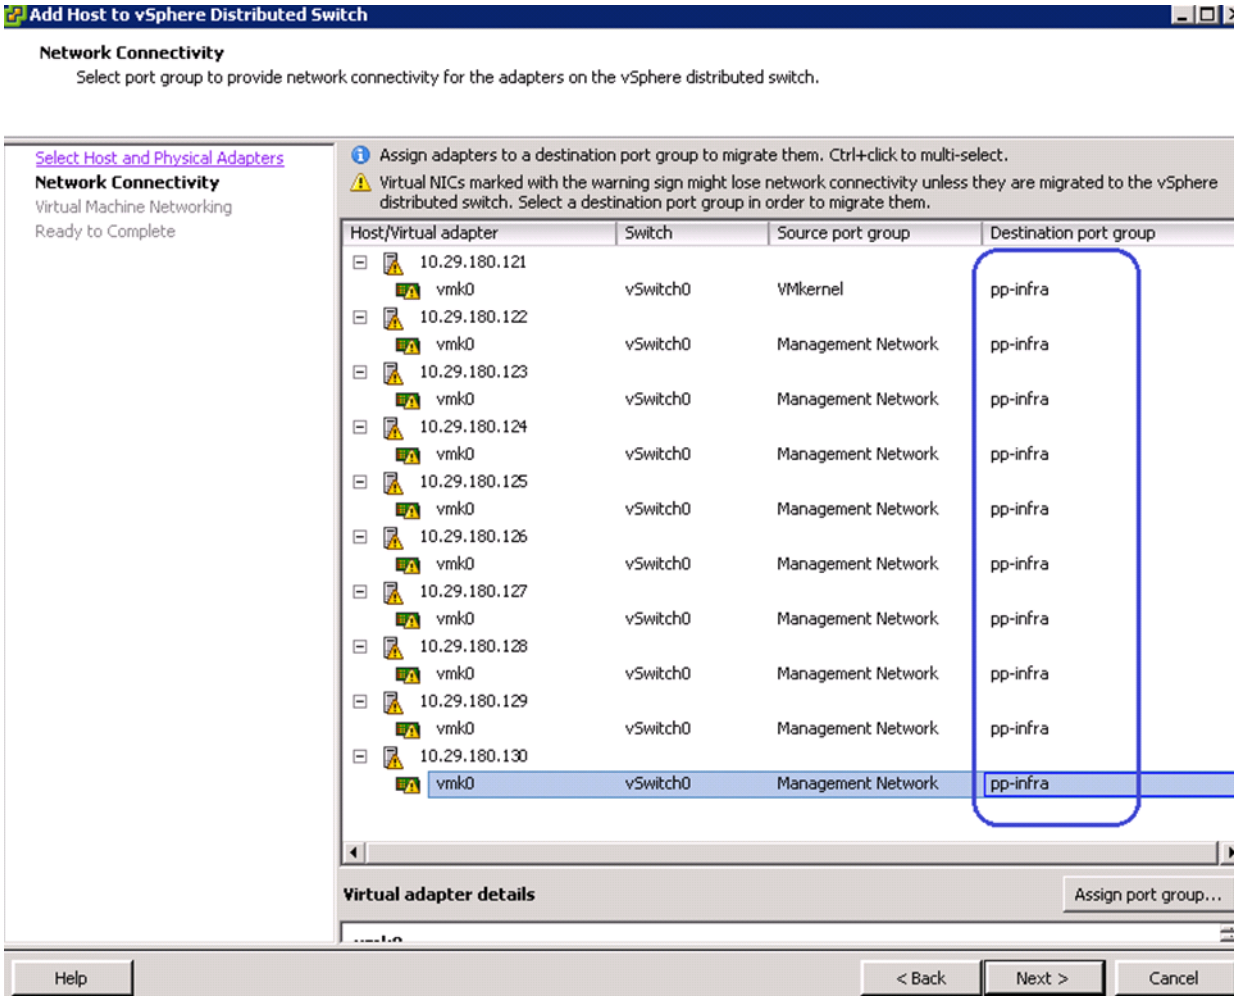

35. Click **Next** as we have not yet created any Virtual Machines on the ESXi hosts.

Figure 215 Virtual Machine Networking Window in Adding Hosts Wizard

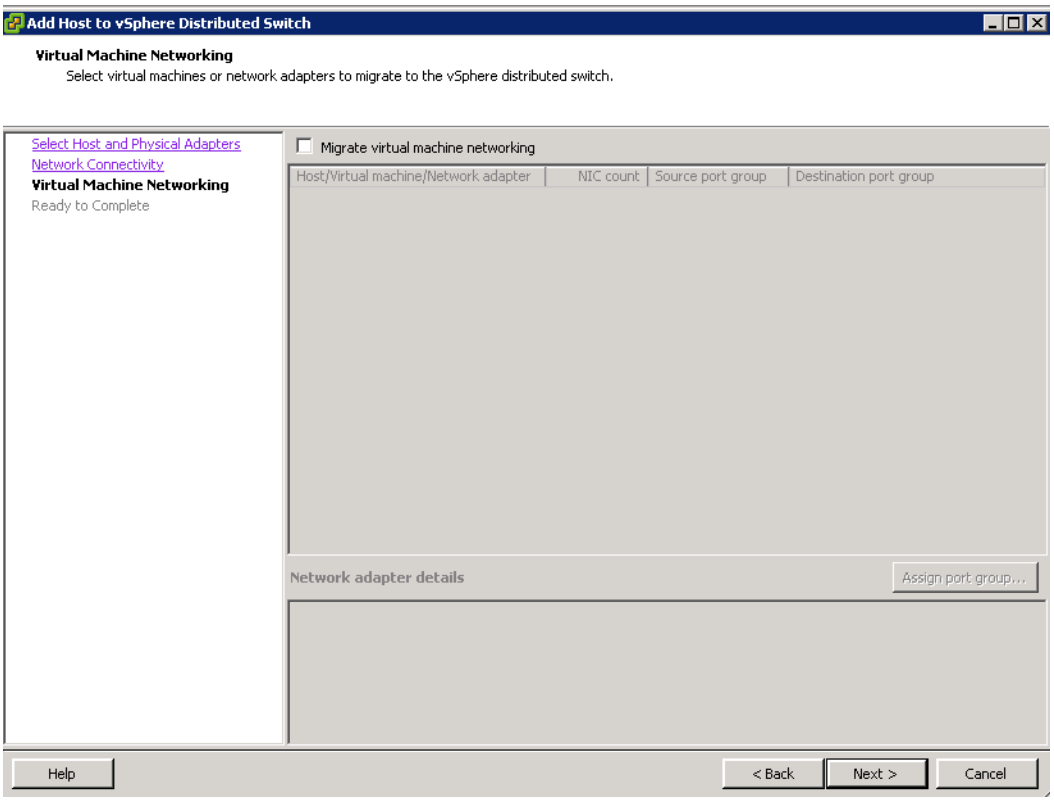

36. Verify the configuration change before submitting the configuration and click **Finish** to complete the migration as shown in Figure 216.

Figure 216 Ready to Complete Window in Adding Hosts Wizard

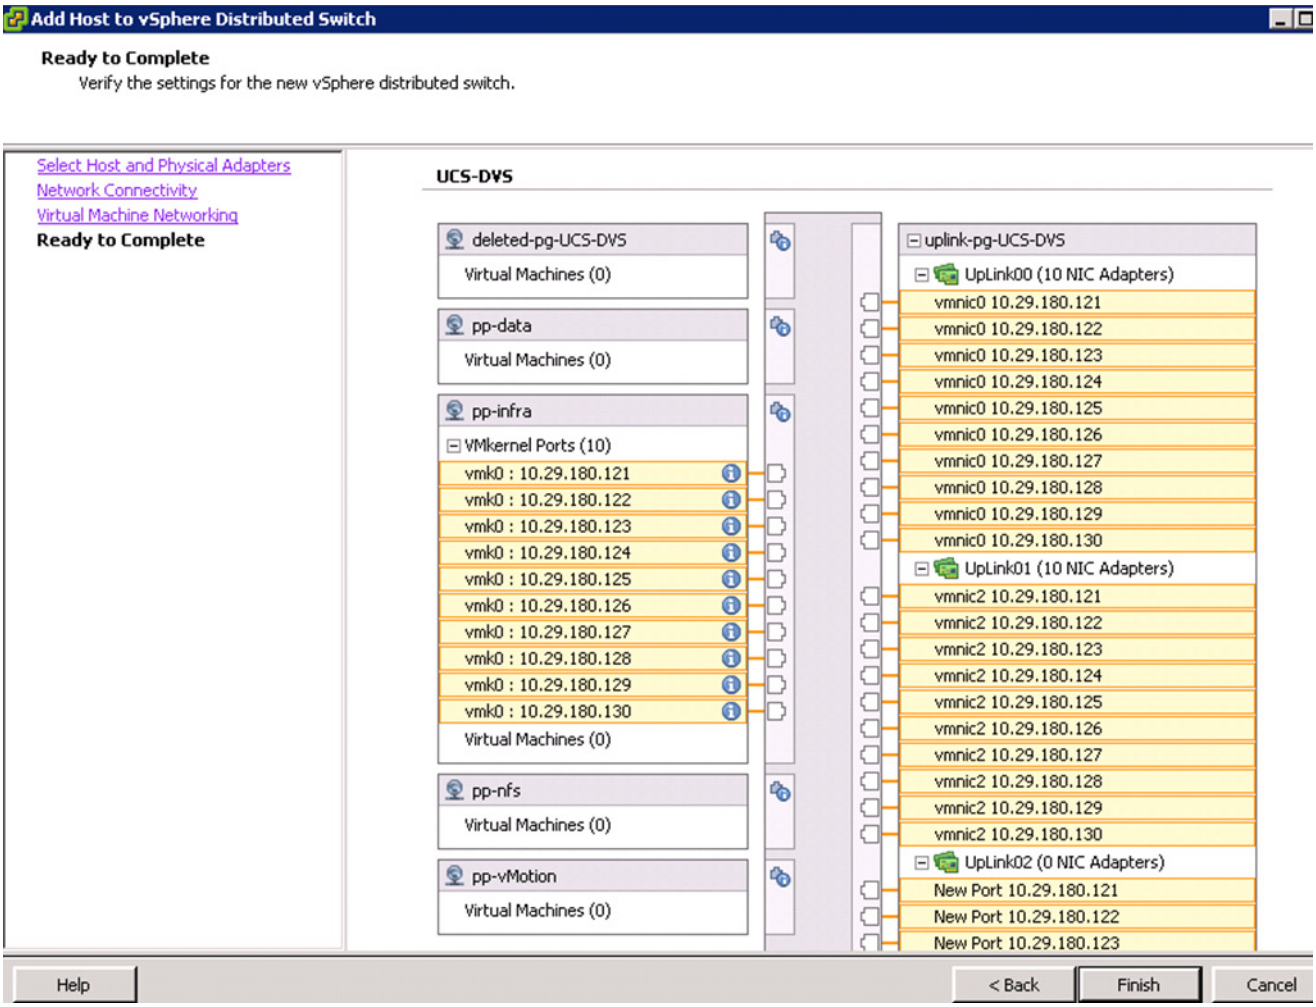

37. Verify that all the hosts shown under the "Hosts" tab and their "VDS Status" is "Up" as shown in Figure 217.

Figure 217 Cisco UCS-DVS Showing Hosts Details in vCenter

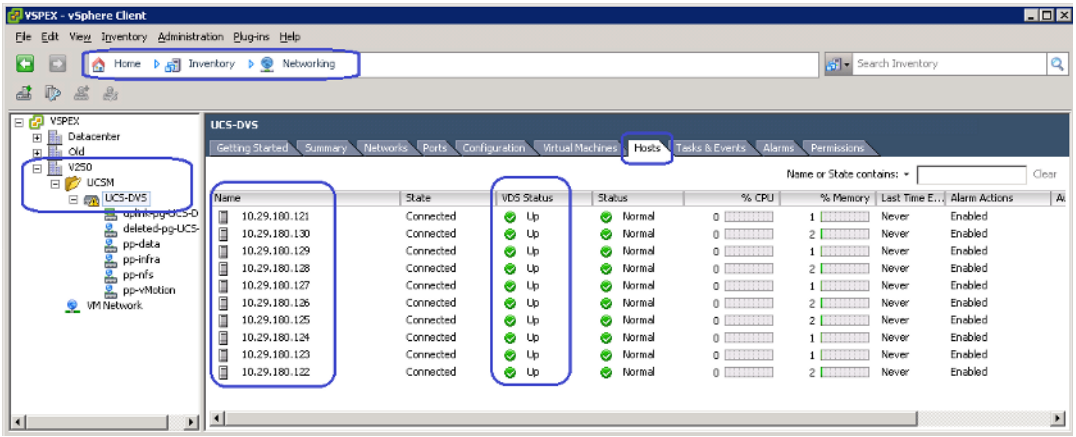

38. In the Cisco UCS Manager window, you can validate the hosts added to VM-FEX vDS by clicking on the **VM** tab, and expanding "VMware", "Virtual Machines" as shown in Figure 218. The "Status" of all the ten servers should show as "Online". Note that Cisco UCS Manager identifies hypervisor hosts by chassis ID and blade ID combination, unlike vCenter which identifies the servers by IP addresses.

Figure 218 Window Showing Host Server Status in Cisco UCS Manager

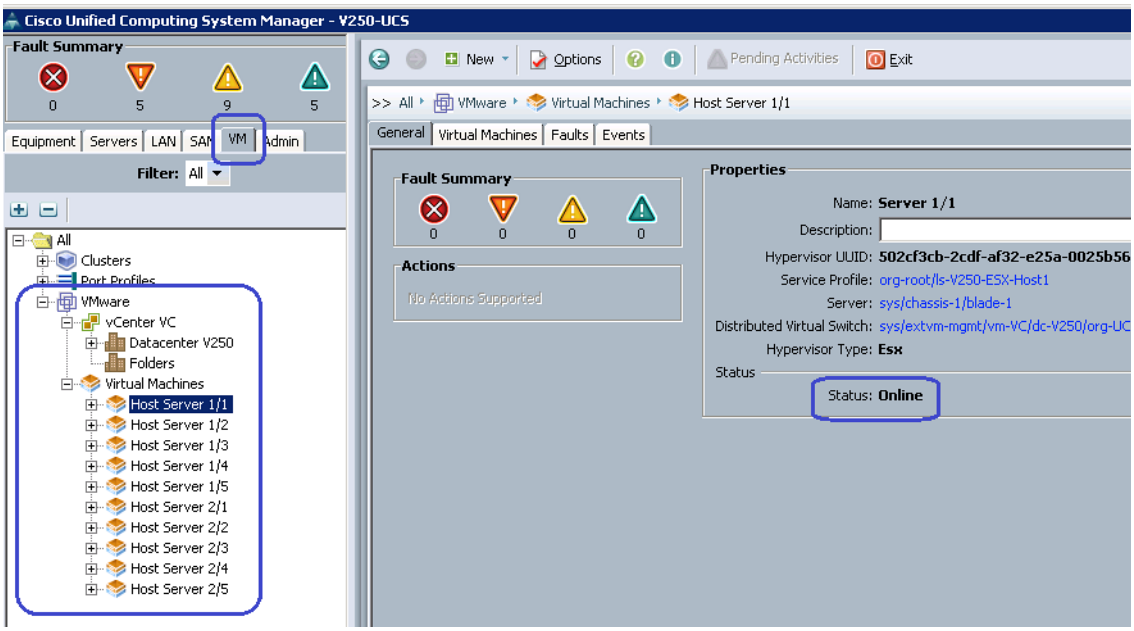

We need to create two more kernel interfaces per host, one for vMotion and one for NFS storage access by the kernel. Choose the appropriate port profiles for the same. For both the vMotion and the NFS kernel interfaces, choose the MTU to be jumbo 9000 MTS for bulk transfer efficiency. Follow these steps:

1. In the vCenter window, click **Inventory > Hosts and Clusters**. Select an individual ESX host and click **Configuration > Networking > vSphere Distributed Switch** on the right pane of the vCenter window. Click **Manage Virtual Adapters…** link as shown in Figure 219.

Figure 219 Managing Virtual Adapters in vCenter

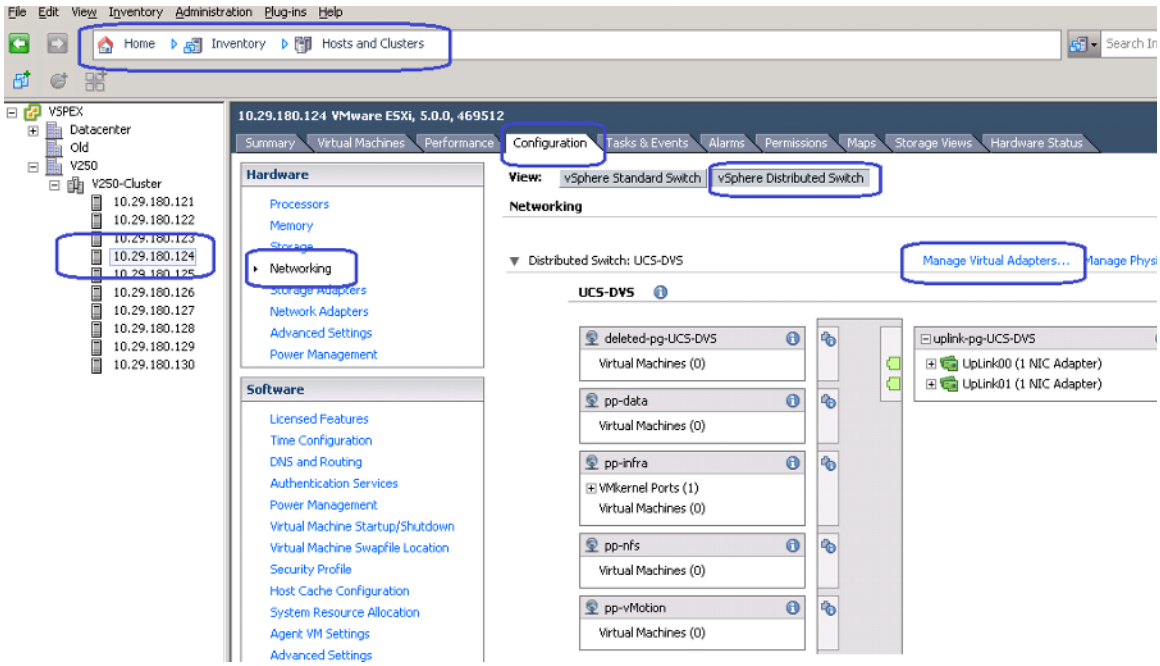

2. Click **Add** to add a new virtual kernel interface.

Figure 220 Adding Virtual Kernel Interface

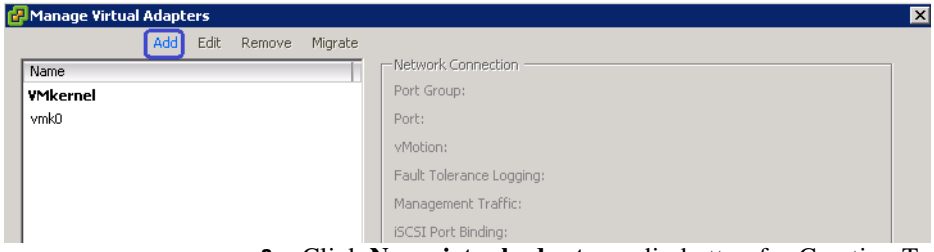

3. Click **New virtual adapter** radio button for Creation Type and click **Next** as shown in Figure 221.

Figure 221 Creation Type Window in Adding Virtual Adapter Wizard

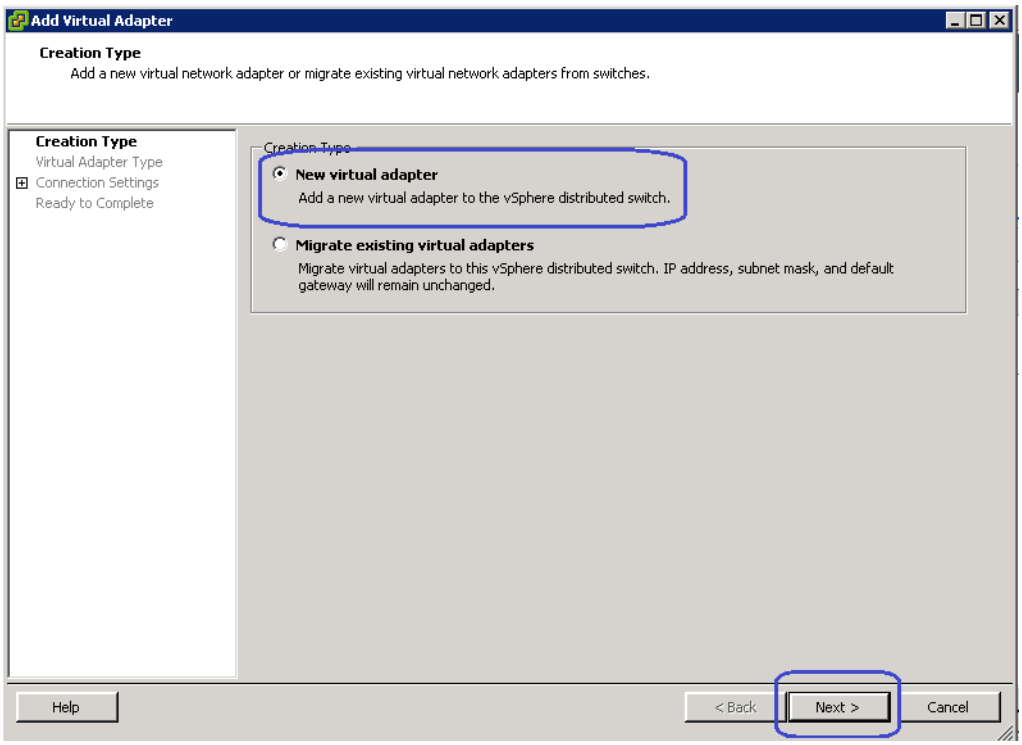

4. Select the "pp-vMotion" port profile from the drop-down list and check the check box "Use this virtual adapter for vMotion". Click **Next**.

Figure 222 Connection Settings Window in Adding Virtual Adapter Wizard

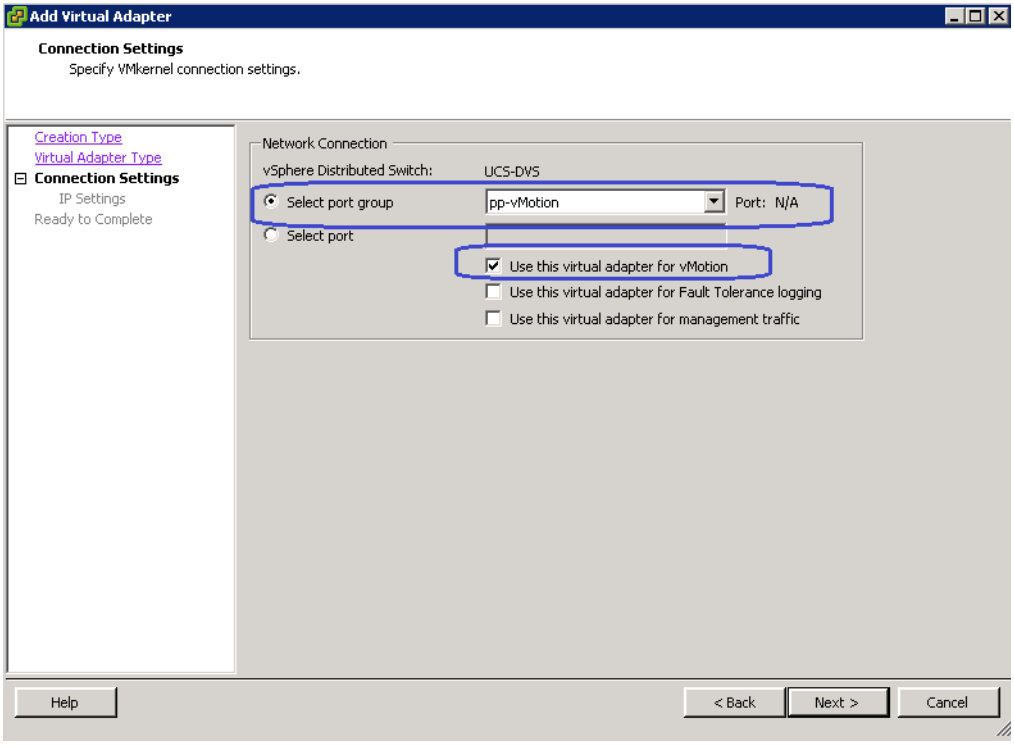

5. Enter IP address and subnet mask. See Customer Configuration Data Sheet, page 172 for assigning the IP address. Click **Next**.

Figure 223 Entering IP Address Details

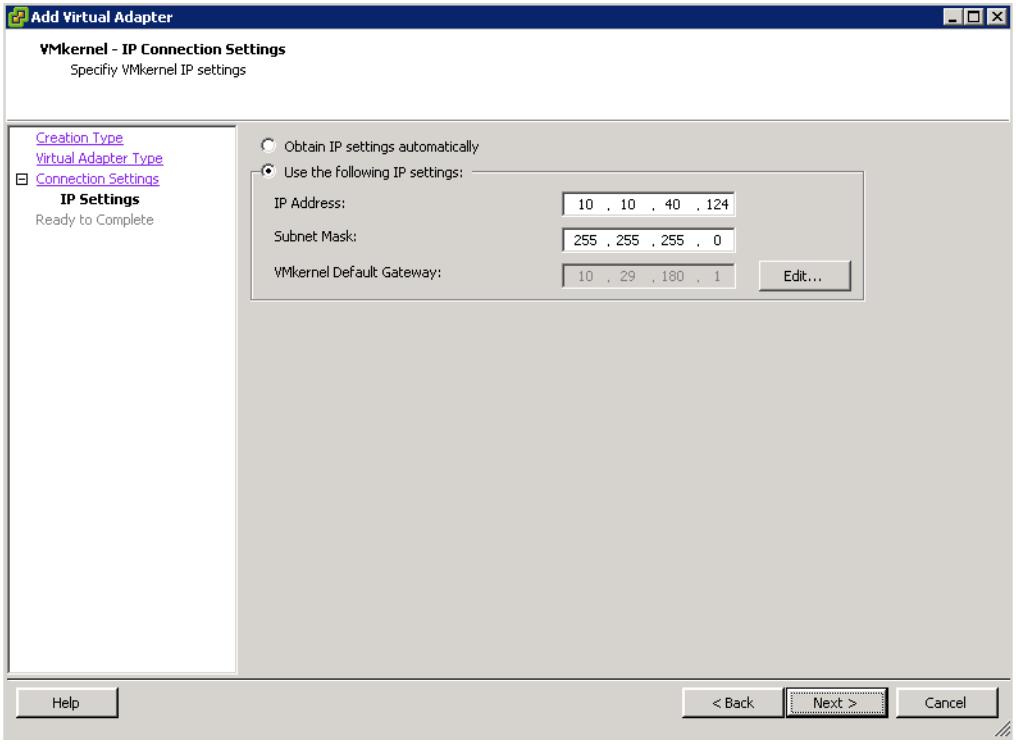

6. Click **Finish** to deploy the new virtual interface.

Figure 224 Ready to Complete Window of Adding Virtual Adapters Wizard

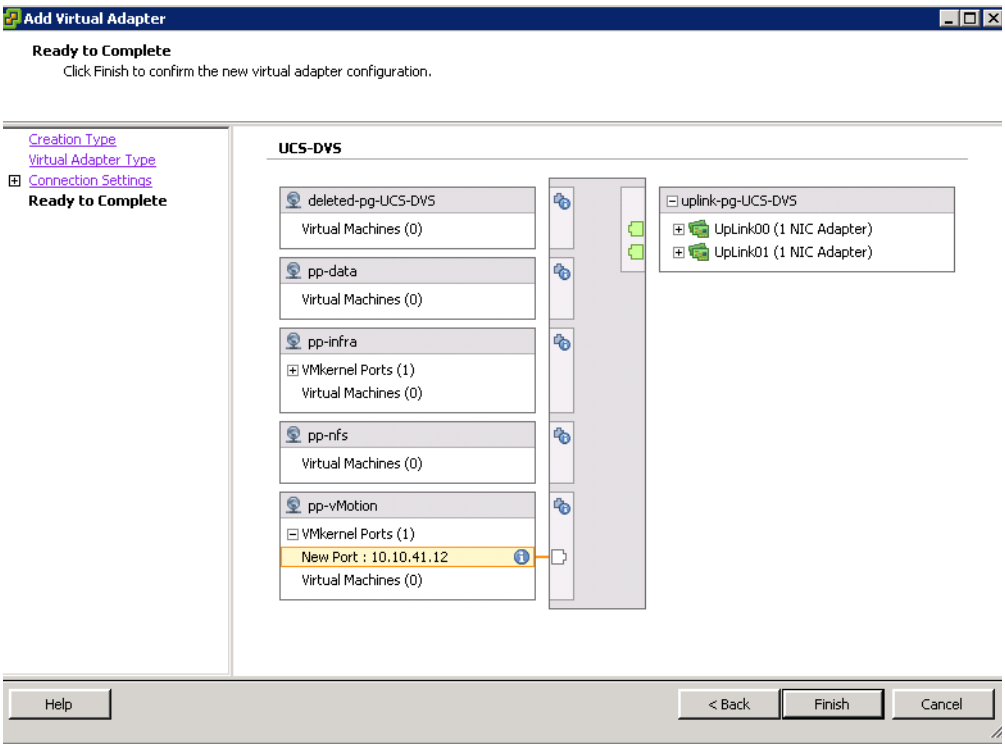

7. Repeat steps 1 to 6 for the storage access vmkernel virtual interface. On step 4, you need to select "pp-nfs" port profile and enable "vMotion".
8. Both vMotion and NFS based storage access require jumbo MTU for the efficient bulk transfer. Select the "vmk1" interface and click **Edit** as shown in Figure 225.

Figure 225 Editing the VMkernel Interface

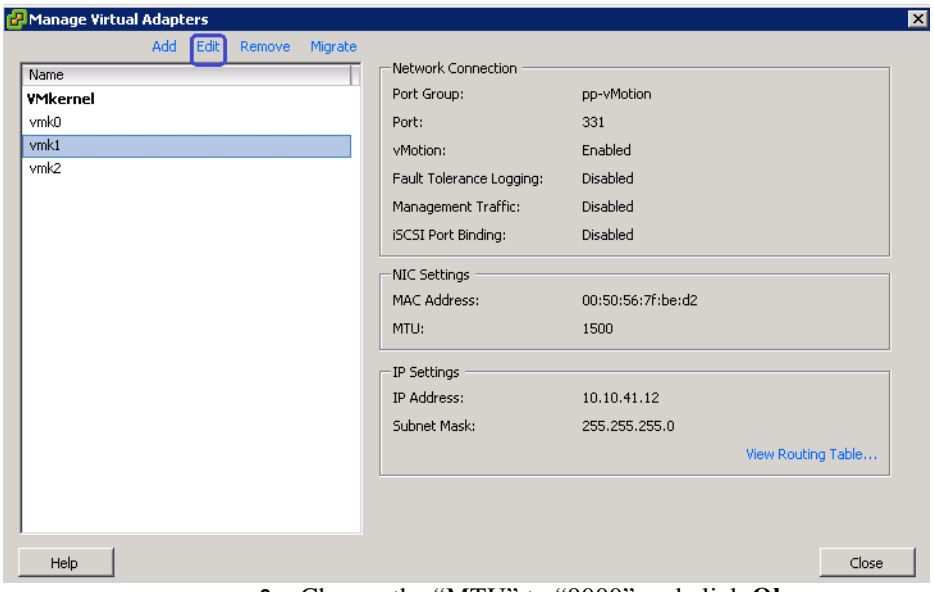

9. Change the "MTU" to "9000" and click **Ok**.

Figure 226 Changing the MTU Size for the VMkernel Interface

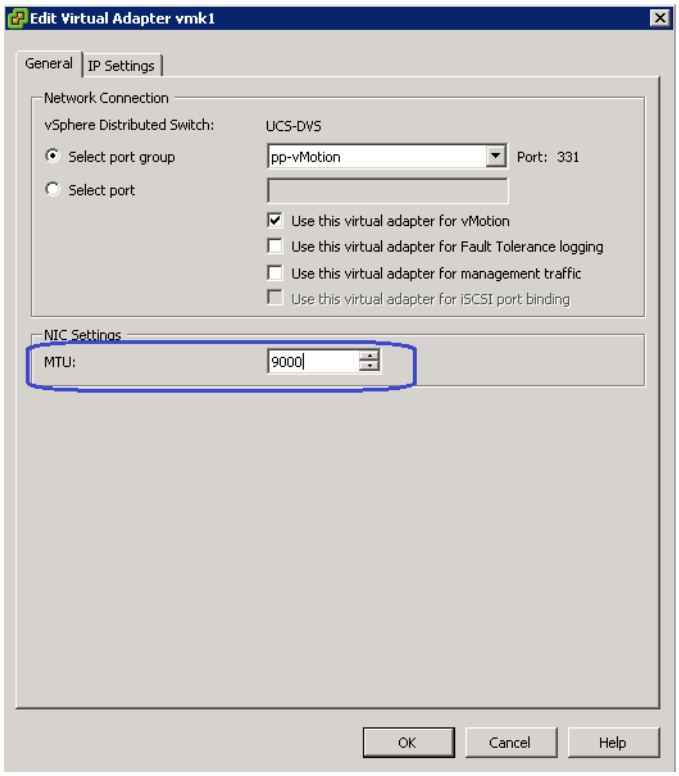

10. Repeat steps 8 and 9 for “mk2” interface as well.
11. You can verify in the Cisco UCS Manager window that every hypervisor host has now 3 vNICs, one in each “pp-vMotion”, “pp-infra” and “pp-nfs” port-profile. This can be verified on the “VM” tab by clicking **VMware > Virtual Machines**, “Host Server x/y” and “vNICs” as highlighted in Figure 227.

Figure 227 VIrtual Machine Properties in Cisco UCS Manager

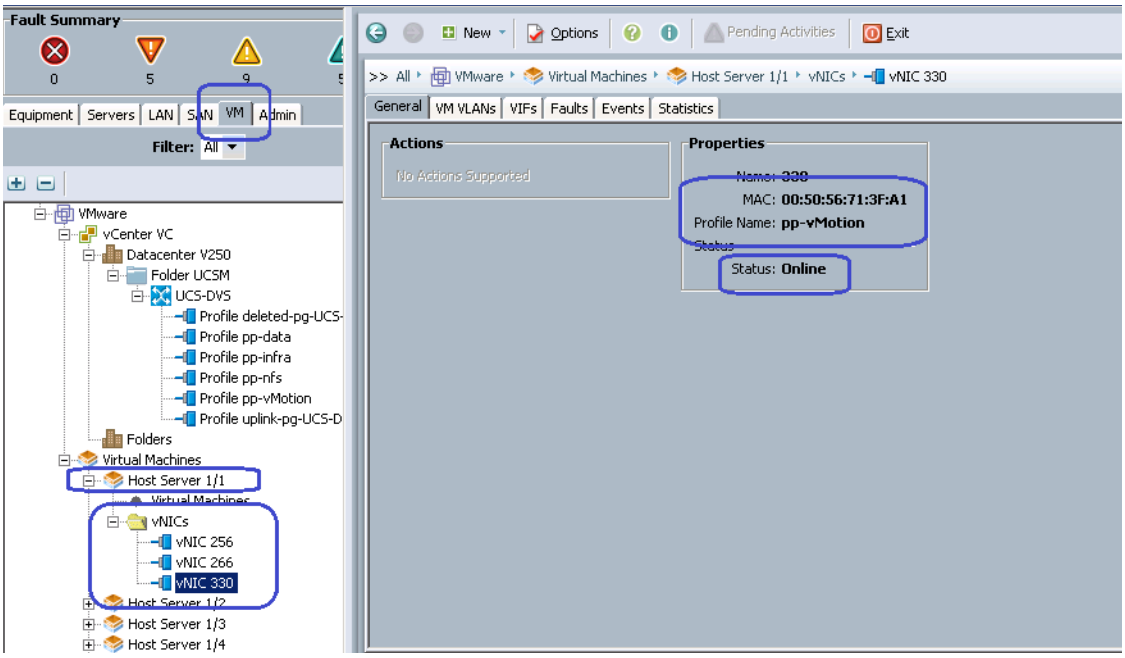

Jumbo MTU Validation and Diagnostics

To validate the jumbo MTU from end-to-end, SSH to the ESXi host. By default, SSH access is disabled to ESXi hosts. Enable SSH to ESXi host by editing hosts' security profile under “Configuration” tab.

When connected to the ESXi host through SSH, initiate ping to the NFS storage server with large MTU size and set the “Do Not Fragment” bit of IP packet to 1. Use the vmkping command as shown in the example:

Example 5

```
~ # vmkping -d -s 8972 10.10.40.6411
PING 10.10.40.64 (10.10.40.64): 8972 data bytes
8980 bytes from 10.10.40.64: icmp_seq=0 ttl=64 time=0.417 ms
8980 bytes from 10.10.40.64: icmp_seq=1 ttl=64 time=0.518 ms
8980 bytes from 10.10.40.64: icmp_seq=2 ttl=64 time=0.392 ms

--- 10.10.40.64 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.392/0.442/0.518 ms
~ #
```

Ensure that the packet size is 8972 due to various L2/L3 overhead. Also ping all other hosts' vMotion and NFS vmkernel interfaces. Ping must be successful. If ping is not successful verify that 9000 MTU configured. Follow these steps to verify:

1. 9000 MTU on the NFS share IP address on the VNX5500 storage device(s).
2. Make sure that a “jumbo-mtu” policy map is created at Nexus 5000 series servers with default class having MTU 9216. Make sure that the “jumbo-mtu” policy is applied to the system classes on the ingress traffic.
3. Make sure that the traffic from storage array to Cisco UCS B200 M3 Blade Servers are marked properly.

4. Make sure that the MTU 9216 is set in the Cisco UCS Manager system class configuration, and QoS policy is correctly set in the port-profiles.
5. Make sure that the 9000 MTU is set for vmkernel ports used for vMotion as well as storage access VNICs.

Template-Based Deployments for Rapid Provisioning

Figure 228 Rapid Provisioning

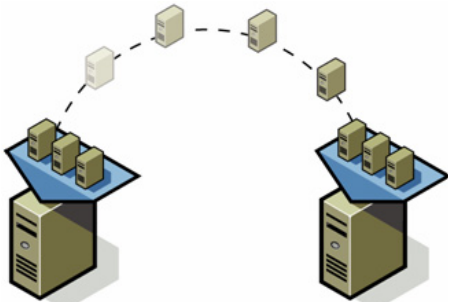

In an environment with established procedures, deploying new application servers can be streamlined, but this can still take many hours or days to complete. Not only should you complete an OS installation, but downloading and installing service packs and security updates can add a significant amount of time. Many applications require features that are not installed with Microsoft Windows by default and must be installed before installing the applications. Inevitably, those features require more security updates and patches. By the time all the deployment aspects are considered, more time is spent waiting for downloads and installs than that spent on 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 the virtual machines. Depending on how granular you want to make a specific template, the time to deploy can be as less as the time it takes to install, configure, and validate the application. You can use PowerShell tools and VMware vSphere Power CLI to bring down the time and manual effort dramatically, especially when you have a large number of VMs to deploy.

Validating Cisco Solution for EMC VSPEX VMware Architectures

This section provides a list of items that should be reviewed when 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 before deploying for production.

- On each vSphere server, verify that the port-profile of virtual Distributed Switch that hosts the client VLANs is configured with sufficient ports to accommodate the maximum number of virtual machines it may host.
- On each vSphere server used, as part of this solution, verify that all the required virtual machine port-profiles is configured and that each server has access to the required VMware datastores.
- On each vSphere server used in the solution, verify that an interface is configured correctly for vMotion, using the correct port-profile and jumbo MTU.
- Create a test virtual machine that accesses the datastore and does read/write operations. Perform the virtual machine migration (vMotion) to a different host on the cluster. Also perform storage vMotion from one datastore to another datastore and ensure correctness of data. During the vMotion of the virtual machine, you need to have a continuous ping to the default gateway and make sure that the 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:

1. Administratively shutdown one of the four links connected between Cisco UCS FI-A and Cisco UCS Fabric Extender. 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 sending some of the data from the virtual machines.
2. Administratively shutdown one of the two uplinks connected to the Cisco Nexus 5548UP switches from FIs. 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 sending some of the data from the virtual machines.
3. Administratively shutdown one of the two data links connected to the storage array from the Cisco Nexus 5548UP switches. Make sure that storage is still available from all the ESXi hosts. Upon administratively enabling the shutdown port, the traffic should be rebalanced.
4. Reboot one of the two Cisco Nexus 5548UP switches 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 switch, the network access load should be rebalanced across the two Cisco Nexus switches.
5. Reboot one of the two UCS 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 switch, the network access load should be rebalanced across the two switches.
6. Reboot the active storage processor of the VNX storage array and make sure that all the NFS shares are still accessible during and after the reboot of the storage processor.
7. 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. There is enough head room for memory in other servers to accommodate 25 additional virtual machines.

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 14 details on the test profile used.

Table 14 Test profile details

Profile characteristic	Value
Number of virtual machines	250
Virtual machine OS	Windows Server 2008 R2 SP1
Processors per virtual machine	1
Number of virtual processors per physical CPU core	4
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

Bill of Material

Table 15 gives the list of the components used in the CVD for 250 virtual machines configuration

Table 15 List of hardware components used in the CVD

Description	Part #
UCS B200 M3 blade servers	UCSB-B200-M3
CPU for B200 M3 blade servers	UCS-CPU-E5-2630
Memory for B200 M3 blade servers	UCS-MR-1X082RY-A
Cisco VIC adapter	UCSB-MLOM-40G-01
UCS 5108 Chassis	N20-C6508
UCS 2204XP Fabric Extenders	UCS-IOM-2204XP
UCS 6248UP Fabric Interconnects	UCS-FI-6248UP
Nexus 5548UP switches	N5K-C5548UP-FA
10 Gbps SFP+ multifiber mode	SFP-10G-SR

For more information on details of the hardware components, see:

http://www.cisco.com/en/US/prod/collateral/ps10265/ps10280/B200M3_SpecSheet.pdf

Customer Configuration Data Sheet

Before you start the configuration, gather the customer-specific network and host configuration information. Table 16, Table 17, Table 18, Table 19, Table 20, Table 21, Table 22 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 16 Common Server Information

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

Table 17 ESXi Server Information

Server Name	Purpose	Primary IP	Private Net (storage) addresses	VMkernel IP	vMotion IP
Host 1	ESXi				
....	….				
Host 10	ESXi				

Table 18 Array Information

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

Table 19 Network Infrastructure Information

Name	Purpose	IP	Subnet Mask	Default Gateway
	Cisco Nexus 5548UP A			
	Cisco Nexus 5548UP B			
	Cisco UCSM Virtual IP			
	Cisco UCS FI-A			
	Cisco UCS FI-B			

Table 20 VLAN Information

Name	Network Purpose	VLAN ID	Allowed Subnets
vlan-infra	Virtual Machine Networking		
ESXi Management			
vlan-nfs	NFS Storage Network		
vlan-vMotion	vMotion traffic network		
vlan-data (multiple)	Data VLAN of customer VMs as needed		

Table 21 VSAN Information

Name	Network Purpose	VSAN ID	Allowed Subnets
vsan-storage	FC connectivity from server to storage		

Table 22 Service Accounts

Account	Purpose	Password (optional, secure appropriately)
	Windows Server administrator	
root	ESXi root	
	Array administrator	
	vCenter administrator	
	SQL Server administrator	
admin	Nexus 5548UP administrator	
admin	Cisco UCS Manager administrator	

References

Cisco UCS:

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

VMware vSphere:

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

Cisco Nexus 5000 Series NX-OS Software Configuration Guide:

http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide.html

EMC VNX 5xxx series resources:

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

Microsoft SQL Server installation guide:

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