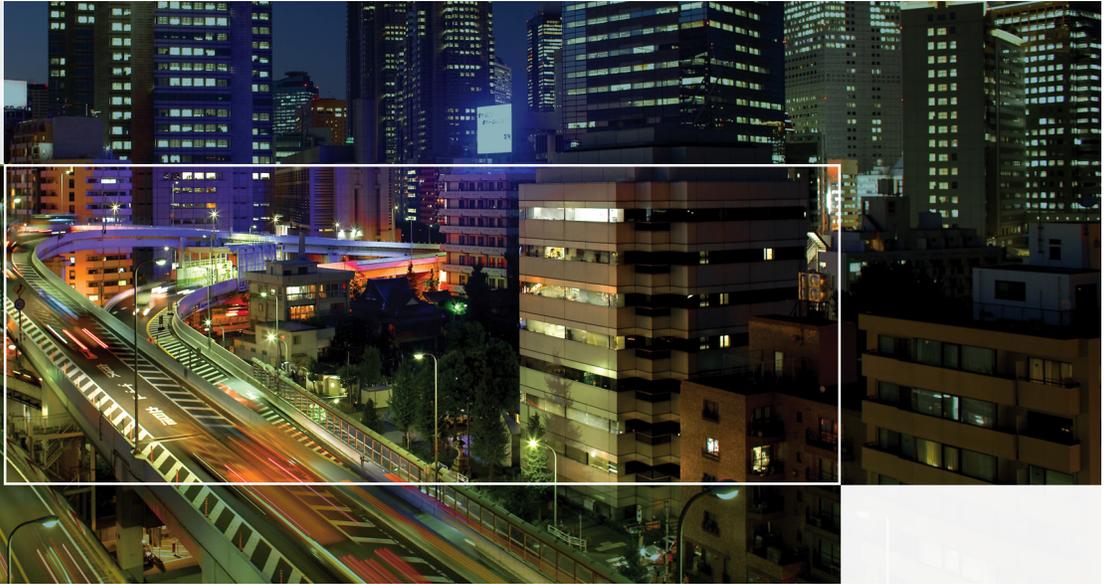


Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 Solution for up to 1000 Virtual Machines

Last Updated: October 18, 2013



Building Architectures to Solve Business Problems





About the Authors



Mehul Bhatt

Mehul Bhatt, 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.



Prashanto Kochavara

Prashanto, EMC Corporation

Prashanto has been working for the EMC solutions group for over 3 years. Prashanto is a SME on EMC Storage and Virtualization technologies including VMware and Hyper-V. He has vast amount of experience in end-to-end solution planning and deployments of VSPEX architectures. Prior to joining the solutions group at EMC, Prashanto has interned as a Systems Engineer (EMC) and Software Developer (MBMS Inc.). Prashanto holds a Bachelors degree in Computer Engineering from SUNY Buffalo and will be graduating with a Masters Degree in Computer Science from North Carolina State University in December 2013.

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
- Tim Cerling-Cisco
- Vijaykumar D-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 Solution for up to 1000 Virtual Machines

Executive Summary

Cisco solution for EMC VSPEX is a pre-validated and modular architecture built with proven best of-breed technologies to create complete end-to-end virtualization solutions that enable you to make an informed decision while choosing the hypervisor, compute, storage and networking layers. VSPEX drastically reduces server virtualization planning and configuration burdens. VSPEX infrastructures accelerate your IT Transformation by enabling faster deployments, greater flexibility of choice, efficiency, and lower risk. This Cisco Validate Design document focuses on the VSPEX VMware architecture for mid-market business segments with less than 1000 typical Virtual Machines load.

Introduction

Virtualization is a key and critical strategic deployment model for reducing the Total Cost of Ownership (TCO) and achieving better utilization of the platform components like hardware, software, network and storage. However, choosing 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, EMC VNX series storage arrays, and Cisco Unified Computing System (UCS) and UCS Manager. External references are provided where ever applicable and it is recommended that the reader be familiar with these documents.



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

Copyright 2012 Cisco Systems, Inc. All rights reserved.

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

Purpose of this Document

This document describes the steps required to deploy and configure a Cisco solution for EMC VSPEX for VMware architectures to a level that will allow for confirmation that the basic components and connections are working correctly. The document covers VMware architectures for SMBs, typically for 1000 VMs or less. This document showcases two variants of the solution:

- **FC-variant**—EMC VNX series storage array directly attached to UCS FIs using FC for storage access.
- **NFS-variant**—EMC VNX series storage array using NFS for storage access through pair of Cisco Nexus 5000-Series switches.

While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are required for deploying these solutions are specifically mentioned.

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 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 1000 generic virtual machines and provide high availability and server redundancy.

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

- Cisco B-Series or C-Series Unified Computing System servers (as per customer's choice)
- Cisco UCS 5108 Chassis

- Cisco UCS 2204XP Fabric Extenders
- Cisco UCS 6248UP Fabric Interconnects
- Cisco UCS Manager 2.1(3a)
- Cisco VIC adapters
- Cisco Nexus 5548UP Switches
- Cisco Nexus 1000v Virtual Switch
- EMC VNX5400, VNX5600, VNX5800 storage array as per the scalability needs
- VMware vCenter 5.1
- Microsoft SQL database
- VMware DRS
- VMware HA

The solution is designed to host scalable, and mixed application workloads for up to 1000 reference virtual machines.

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 V2 Series Processors.
- **Network**—The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Virtualization**—The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access**—The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet (NFS), Fibre Channel, Fibre Channel over Ethernet (FCoE). 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.

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership (TCO) 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.
- Industry standards supported by a partner ecosystem of industry leaders.

Cisco UCS Manager

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

Cisco UCS Fabric Interconnect

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

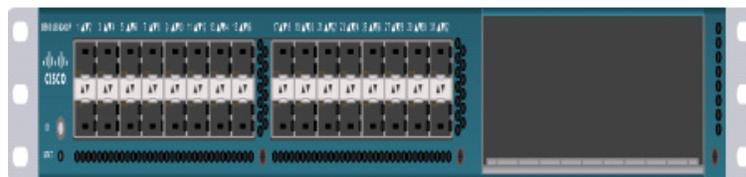
The Cisco UCS 6200 Series provides the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS 5100 Series Blade Server Chassis. All chassis, and therefore all blades, attached to the Cisco UCS 6200 Series Fabric Interconnects become part of a single, highly available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

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

Cisco UCS 6248UP Fabric Interconnect

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

Figure 1 *Cisco UCS 6248UP Fabric Interconnect*



Cisco UCS Fabric Extenders

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

Cisco UCS 2232PP Fabric Extender

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

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

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

Figure 2 Cisco UCS 2232PP Fabric Extender



Cisco C220 M3 Rack Mount Servers

Building on the success of the Cisco UCS C220 M3 Rack Servers, the enterprise-class Cisco UCS C220 M3 server further extends the capabilities of the Cisco Unified Computing System portfolio in a 1-rack-unit (1RU) form factor. And with the addition of the Intel® Xeon® processor E5-2600 V2.

Figure 3 Cisco UCS C220 M3 Rack Mount Server



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

Cisco UCS Blade Chassis

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

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

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

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

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

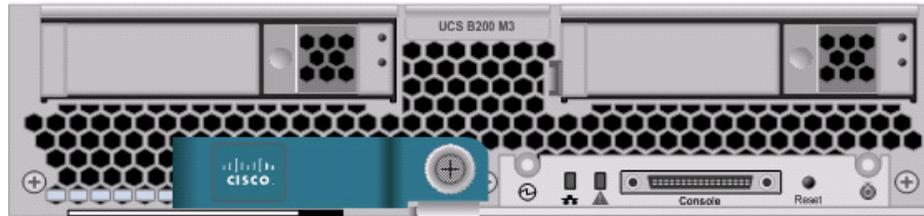


Cisco UCS Blade Servers

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

Building on the success of the Cisco UCS B200 M2 blade servers, the enterprise-class Cisco UCS B200 M3 server, further extends the capabilities of Cisco's Unified Computing System portfolio in a half blade form factor. The Cisco UCS B200 M3 server harnesses the power and efficiency of the Intel Xeon E5-2600 V2 processor product family, up to 768 GB of RAM, 2 drives or SSDs and up to 2 x 20 GbE to deliver exceptional levels of performance, memory expandability and I/O throughput for nearly all applications. In addition, the Cisco UCS B200 M3 blade server offers a modern design that removes the need for redundant switching components in every chassis in favor of a simplified top of rack design, allowing more space for server resources, providing a density, power and performance advantage over previous generation servers. The Cisco UCS B200M3 Server is shown in [Figure 5](#).

Figure 5 Cisco UCS B200 M3 Blade Server



Cisco 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 6](#).

Figure 6 Cisco Nexus 5548UP switch



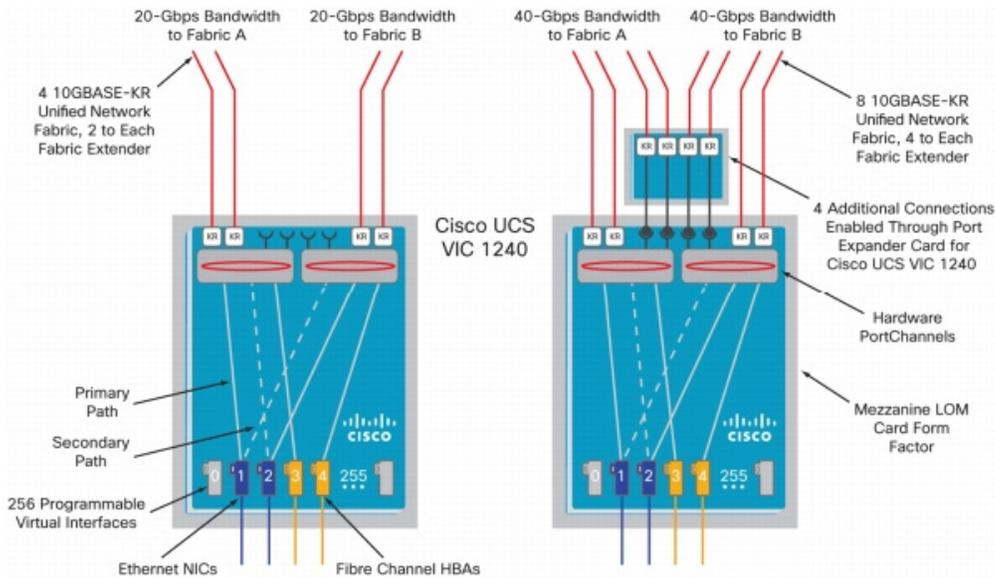
Cisco I/O Adapters

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 7 Cisco UCS VIC 1240



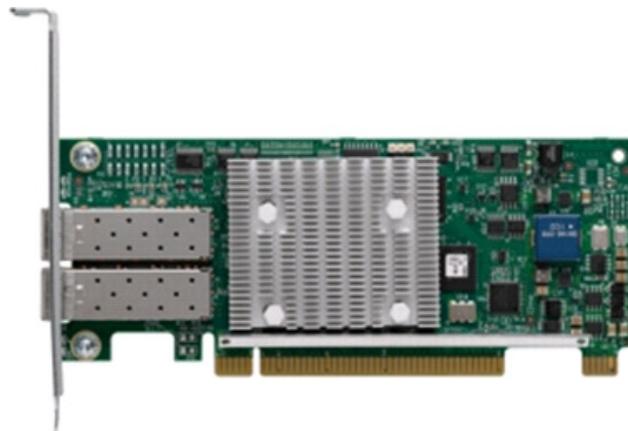
The Cisco UCS rack mount server has various Converged Network Adapters (CNA) options. The UCS 1225 Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

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

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

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

Figure 8 Cisco UCS 1225 VIC



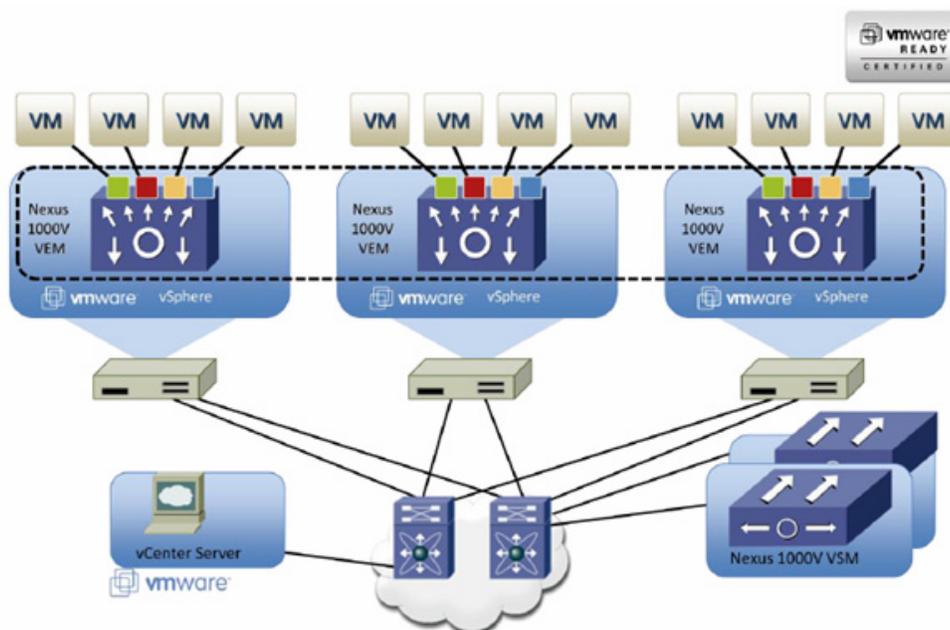
Nexus 1000v Virtual Switch

Nexus 1000v is a virtual Ethernet switch with two components:

- Virtual Supervisor Module (VSM)—The control plane of the virtual switch that runs NX OS.
- Virtual Ethernet Module (VEM)—A virtual line card embedded into each VMware vSphere hypervisor host (ESXi).

Virtual Ethernet Modules across multiple ESXi hosts form a virtual Distributed Switch (vDS). Using the Cisco vDS VMware plug-in, 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 the host. All traffic from an interface on a virtual machine is sent to the corresponding port of the vDS switch. The traffic is then sent out to physical link of the host using the special uplink port-profile. This vDS implementation guarantees consistency of features and better integration of host virtualization with rest of the Ethernet fabric in the data center.

Figure 9 Nexus 1000v Virtual Distributed Switch Architecture



UCS 2.1 Single Wire Management

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

UCS Differentiators

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

- 1. Embedded management**—In UCS, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers. Also, a pair of FIs can manage up to 40 chassis, each containing 8 blade servers. This gives enormous scaling on the management plane.
- 2. Unified fabric**—In UCS, from blade server chassis or rack server fabric-extender to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters – reducing capital and operational expenses of overall solution.
- 3. Auto Discovery**—By simply inserting the blade server in the chassis or connecting rack server to the fabric extender, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of UCS, where compute capability of UCS can be extended easily while keeping the existing external connectivity to LAN, SAN and management networks.
- 4. Policy based resource classification**—Once a compute resource is discovered by UCS Manager, 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 UCS Manager.
- 5. Combined Rack and Blade server management**—UCS Manager 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**—UCS Manager 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 UCS Manager with other management system, such as VMware vCloud director, Microsoft System Center, and Citrix Cloud Platform.
- 7. Policies, Pools, Templates**—The management approach in UCS Manager 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 UCS Manager, 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 UCS Manager, 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 UCS Manager 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 UCS Manager 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 vSphere 5.0 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 Nexus 1000v switch integrated into the vCenter Server, deploying and administering virtual machines is similar to deploying and administering physical servers. Network administrators can continue to own the responsibility for configuring and monitoring network resources for virtualized servers as they did with physical servers. System administrators can continue to “plug-in” their virtual machines into the network ports that have Layer 2 configurations, port access and security policies, monitoring features, and so on, that have been pre-defined by the network administrators; in the same way they need to plug in their physical servers to a previously-configured access switch. In this virtualized environment, the network port configuration/policies move with the virtual machines when the virtual machines are migrated to different server hardware.

EMC Storage Technologies and Benefits

This document describes the steps required to deploy and configure a Cisco solution for EMC VSPEX with VMware architectures. This architecture has two variants:

- FC-variant of the solution, where EMC VNX series storage devices are attached to UCS FIs directly, accessing storage over Fibre Channel protocol.
- NFS-variant of the solution, where EMC VNX series storage devices are accessed from UCS through a pair of Nexus 5000 Series switches, accessing storage over NFS protocol.

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 and FC 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 CVD focuses on the architecture for EMC VSPEX for VMware private cloud, targeted for mid-market segment, using VNX storage arrays. There are two variants of the architecture: FC-variant and NFS-variant. FC-variant of the architecture uses UCS 2.1(3a) with combined B-series and C-series servers with VNX5400 directly attached to UCS fabric interconnect. NFS-variant of the architecture uses UCS 2.1(3a) with B-series and C-series servers with VNX5600 or VNX5800 storage array attached to the Nexus 5548 switches. In both variants, the C220 M3 servers are connected with single-wire management feature. VMware vSphere 5.1 is used as server virtualization architecture. FC-variant of architecture show cases VMware's native virtual switching, while NFS-variant of architecture show cases Cisco Nexus 1000v based virtual switches. However, either architecture can use any of the virtual switching components.

[Table 1](#) lists the various hardware and software components which occupies different tiers of the Cisco solution for EMC VSPEX VMware architectures under test.

Table 1 *Hardware and Software Components of VMware Architectures*

Vendor	Name	Version	Description
Cisco	Cisco UCS B200M3 servers	2.1(3a)	Cisco UCS B200M3 Blade Servers (FC-variant)
Cisco	Cisco UCS 5108 Chassis	-	Cisco UCS Blade Server Chassis
Cisco	Cisco VIC 1240	2.1(3a)	Cisco Virtual Interface Card (adapter) firmware
Cisco	Cisco 2204XP Fabric Extender	2.1(3a)	Cisco UCS fabric extender firmware
Cisco	Cisco UCS 6248UP Fabric Interconnect	5.0(3)N2(2.11)	Cisco UCS fabric interconnect firmware
Cisco	Cisco 2232PP Fabric Extender	5.0(3)N2(2.11.2)	Cisco UCS Fabric Extender
Cisco	Cisco UCS C220M3 Servers	1.5(2) or later – CIMC C220M3.1.5.2.23 - BIOS Cisco C220M3 rack servers	Cisco UCS C220M3 Rack Servers
Cisco	Cisco UCS 1240 VIC	2.1(3a)	Cisco UCS VIC adapter (FC-variant)

Table 1 Hardware and Software Components of VMware Architectures

Vendor	Name	Version	Description
Cisco	Cisco UCS Manager	2.1(3a)	Cisco UCS Manager software
Cisco	Cisco Nexus 5548UP Switches	5.1(3)N1(1a)	Cisco Nexus 5000 series switches running NX-OS
Cisco	Cisco nexus 1000v	4.2(1).SV2(2.1)	Cisco Nexus 1000v Virtual Switch (NFS-variant)
EMC	EMC VNX5400	VNX Block OE 05.33	EMC VNX storage array (FC-variant)
EMC	EMC VNX5600	VNX File OE 8.1 VNX Block OE 05.33	EMC VNX storage array (NFS-variant)
EMC	EMC VNX5800	VNX File OE 8.1 VNX Block OE 05.33	EMC VNX storage array (NFS-variant)
EMC	EMC Avamar	6.1 SP1	EMC data backup software
EMC	Data Domain OS	5.2	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 2012	Windows Server 2012 Datacenter	Operating system to host vCenter server
Microsoft	Microsoft SQL server	2008 R2	Database server SQL R2 Enterprise edition for vCenter

Table 2 outlines the B200M3 or C220M3 server configuration (per server basis) across the two variants of VMware architectures.

Table 2 Server Configuration Details

Component	Capacity
Memory (RAM)	128GB (8 X 16GB DIMM)
Processor	2 x Intel® Xenon® E5-2600 V2, CPUs 2.0 GHz, 8cores, 16 threads
Local storage	Cisco UCS RAID SAS 2008M-8i Mezzanine Card, With 2 x 67 GB slots for RAID 1 configuration each

Both the variants of architecture assume that there is an existing infrastructure/ management network available where a virtual machine hosting vCenter server and Windows Active Directory/ DNS server are present.

The required number of C or B series servers and storage array type change depending on number of virtual machines. Table 3 highlights the change in hardware components, as required by different scale. Typically, 50 reference Virtual Machines are deployed per server.

Table 3 Hardware Components for Different Scale

Components	VMware 300 VMs	VMware 600 VMs	VMware 1000 VMs
Servers	6 x Cisco C220M3 or B200M3 servers	12 x Cisco B200M3 or C220M3 servers	18 x Cisco B200M3 or C220M3 servers
Blade Server Chassis	1 x Cisco 5108 Blade Server Chassis	2 x Cisco 5108 Blade Server Chassis	3 x Cisco 5108 Blade Server Chassis
Storage	EMC VNX5400	EMC VNX5600	EMC VNX5800

Figure 10 and Figure 11 show a high level Cisco solution for EMC VSPEX VMware FC-variant and NFS-variant architectures respectively.

Figure 10 Reference Architecture for FC-Variant

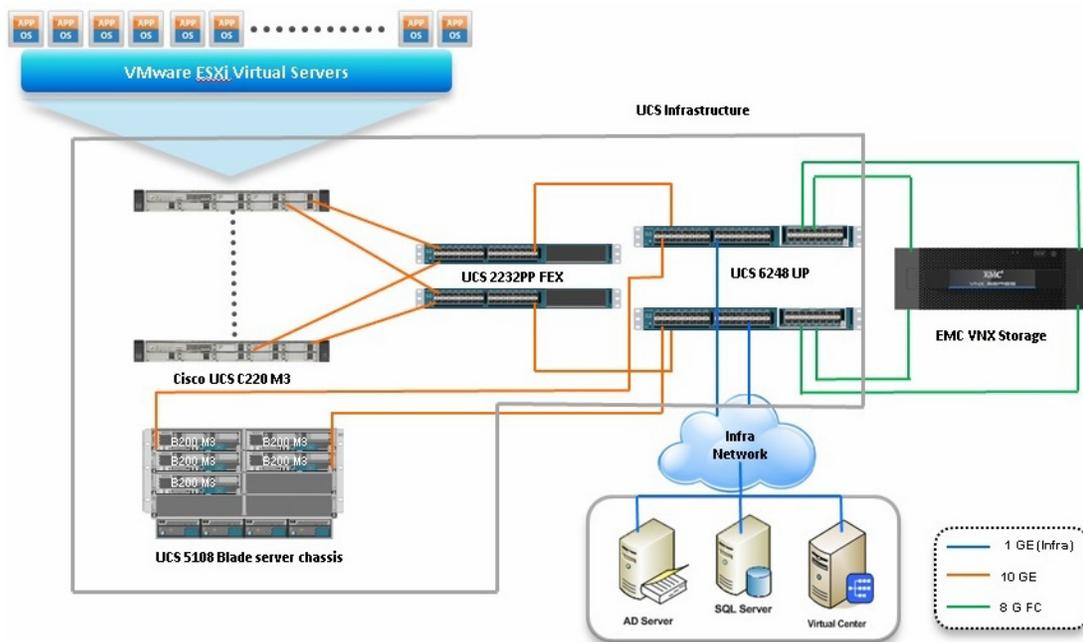
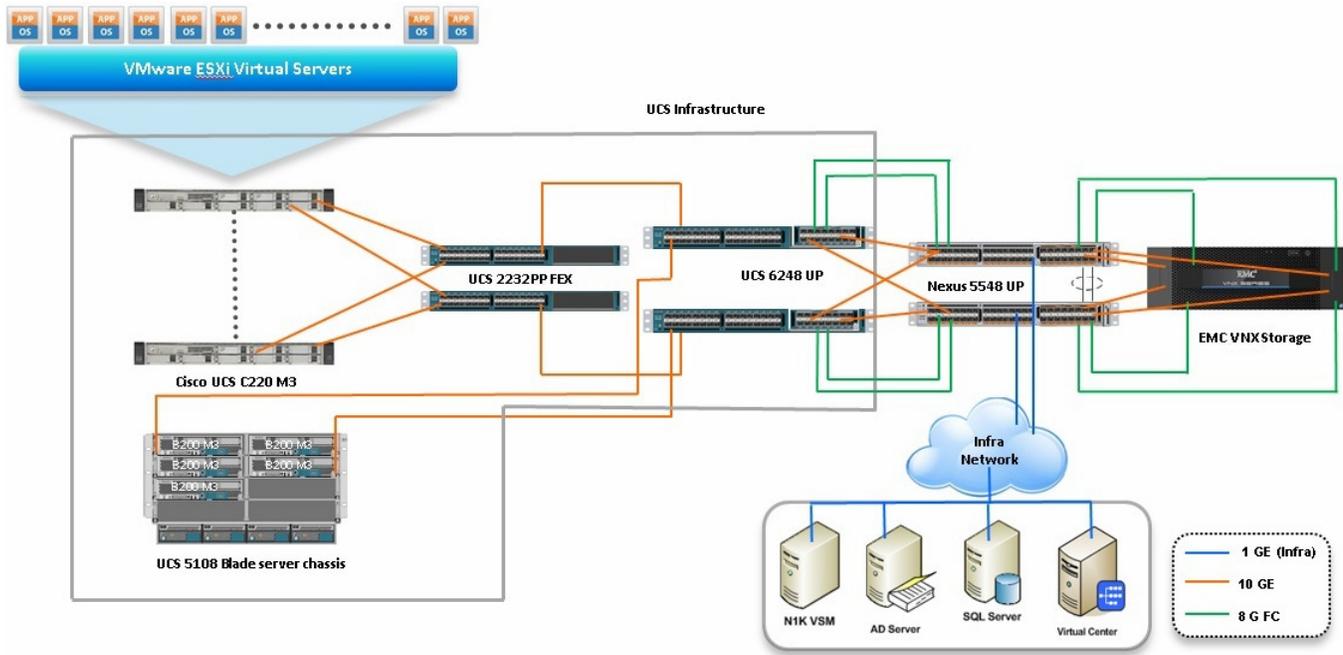


Figure 11 Reference Architecture for NFS-Variant



Following are the key design points of the architectures for mid-market segment:

- For smaller scale, storage array dedicated to a given UCS domain is preferable for simplicity, so VNX5400 storage is directly attached to the Cisco UCS FIs. For larger scales, multiple UCS domains may share a storage or multiple storage arrays; In this case, storage access through Cisco Nexus 5548UP is preferable, as shown in the NFS-variant architecture.
- Infrastructure network is on a separate 1GE network.
- Network redundancy is built-in by providing two switches, two storage controllers and redundant connectivity for data, storage, and infrastructure networking.

This design does not dictate or require any specific layout of infrastructure network. The vCenter server, Microsoft AD server and Microsoft SQL server are hosted on infrastructure network. However, design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

ESXi 5.1 is used as hypervisor operating system on each server and is installed on SAN LUNs in both the architectures. However, virtual machines' storage is accessed through FC or NFS protocols depending on the architecture. Typical load is 50 reference 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:

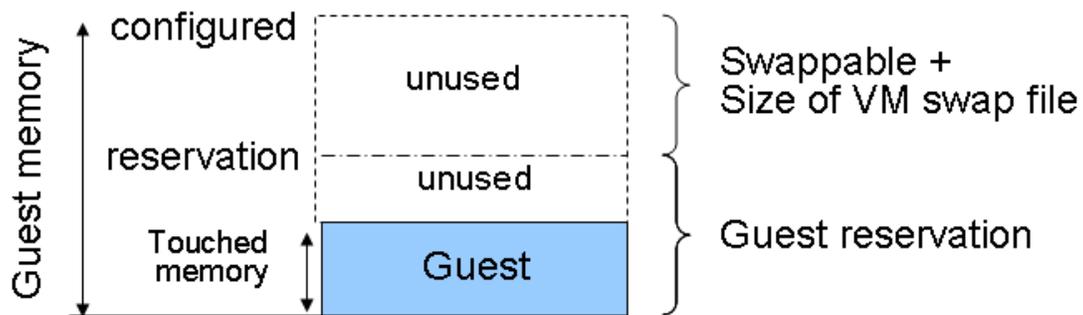
- Memory sharing across virtual machines that have similar data (that is, same guest operating systems).
- Memory over commitment, which means allocating more memory to virtual machines than is physically available on the ESX/ESXi host.
- A memory balloon technique whereby virtual machines that do not 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 12 shows the use of memory settings parameters in the virtual machine.

Figure 12 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. Also, this value is the size of per-virtual machine swap file that is created on the VMware Virtual Machine File System (VMFS) VSWP file. If the balloon driver is unable to reclaim memory quickly enough, or is disabled, or not installed, vSphere forcibly reclaims memory from the virtual machine using the VMkernel swap file.

Allocating Memory to Virtual Machines

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 machine operating system	Microsoft Windows Server 2012
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

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 32GB 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.

Storage Guidelines

VSPEX architecture for VMware VMs for mid-market segment uses FC or NFS to access storage arrays. FC is used with smaller scale with VNX5400 storage array, while NFS is used with VNX5600 or VNX5800 storage array. vSphere provides many features that take advantage of EMC storage technologies such as auto discovery of storage resources and ESXi hosts in vCenter and VNX respectively. Features such as VMware vMotion, VMware HA, and VMware Distributed Resource Scheduler (DRS) use these storage technologies to provide high availability, resource balancing, and uninterrupted workload migration.

Storage Protocol Capabilities

VMware vSphere provides vSphere and storage administrators with the flexibility to use the storage protocol that meets the requirements of the business. This can be a single protocol datacenter wide, such as NFS, or multiple protocols for tiered scenarios such as using Fibre Channel for high-throughput storage pools and NFS for high-capacity storage pools.

For 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: 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 2012 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

NFS-variant architecture demonstrates use and benefits of Nexus 1000v virtual switching technology. Each B200 M3 blade server and C220 M3 rack server has one physical adapter with two 10 GE links going to fabric A and fabric B for high availability. Cisco UCS VIC 1225 or 1240 presents four virtual Network Interface Cards (vNICs) to the hypervisor, two vNICs per fabric path. In FC-variant, the Cisco UCS VIC 1225 adapter also presents two virtual Host Bus Adapters (vHBAs) to the hypervisor, one per fabric path. The MAC addresses to these vNICs are assigned using MAC address pool defined on the UCS Manager. The vNICs are used in active-active configuration for load-balancing and high-availability. Following are vSphere networking best practices implemented in this architecture:

- Separate virtual machine and infrastructure traffic—Keep virtual machine and VMkernel or service console traffic separate. This is achieved by having two vSwitches per hypervisor:
 - vSwitch (default)—Used for management and vMotion traffic
 - vSwitch1—Used for Virtual Machine data traffic
- Use NIC Teaming—Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches. This is achieved by using two vNICs per vSwitch, each going to different Fabric Interconnects. Teaming provides redundancy against NIC failure, switch (FI or FEX) failures, and in case of UCS, upstream switch failure (due to “End-Host-Mode” architecture).
- Enable PortFast on ESX/ESXi host uplinks—Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. Cisco UCS Fabric Extenders are not really Ethernet switches, they are line cards to the Fabric Interconnect, and Cisco UCS Fabric Interconnects run in end-host-mode and avoid running Spanning Tree Protocol. Given this, there is no need to enable port-fast on the ESXi host uplinks. However, it is recommended that you enable

portfast on Cisco Nexus 5548UP switches or infrastructure switches that connect to Cisco UCS Fabric Interconnect uplinks for faster convergence of STP in the events of FI reboot or FI uplink flap.

- Jumbo MTU for vMotion and storage traffic—This practice is implemented in the architecture by configuring jumbo MTU end-to-end.

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 VNX series is simplified through the use of wizards, so that administrators do not choose which disks belong to a given storage pool. The wizard may choose any available disk of the proper type, regardless of where the disk physically resides in the array.

Figure 13 shows storage architecture for 300 virtual machines on VNX5400 for FC-variant of architecture.

Figure 13 Storage Architecture for 300 VMs on EMC VNX5400

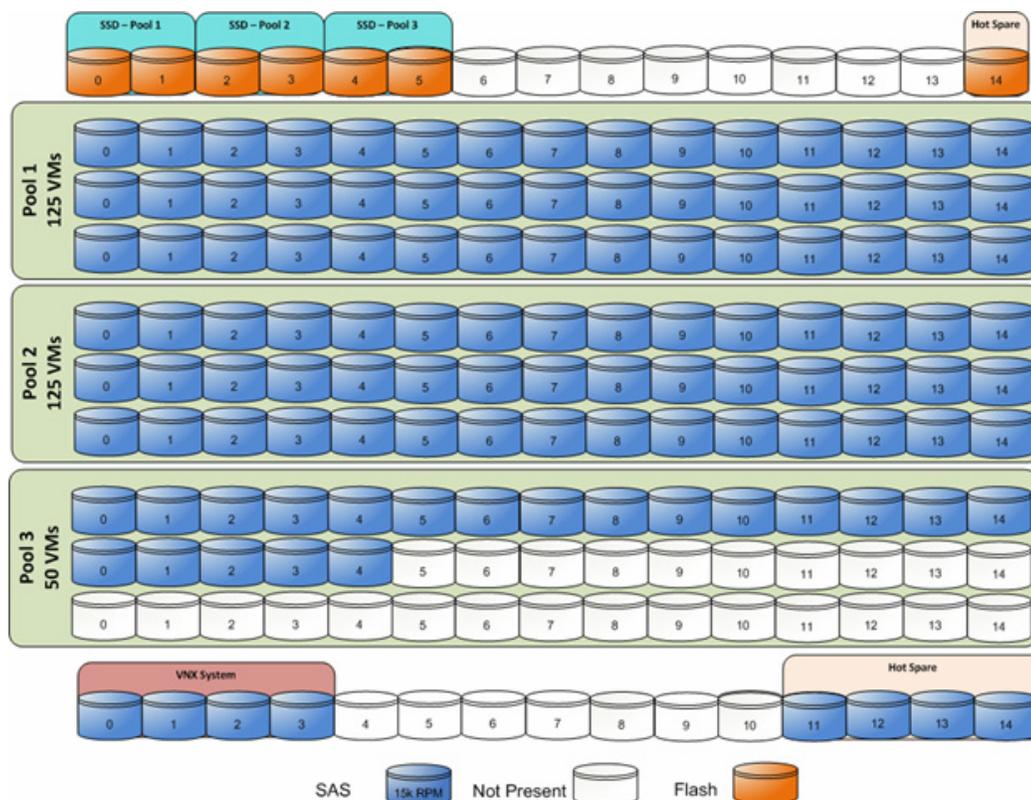


Figure 14 shows storage architecture for 600 virtual machines on VNX5600 for NFS-variant of architecture.

Figure 14 Storage Architecture for 600 VMs on EMC VNX5600

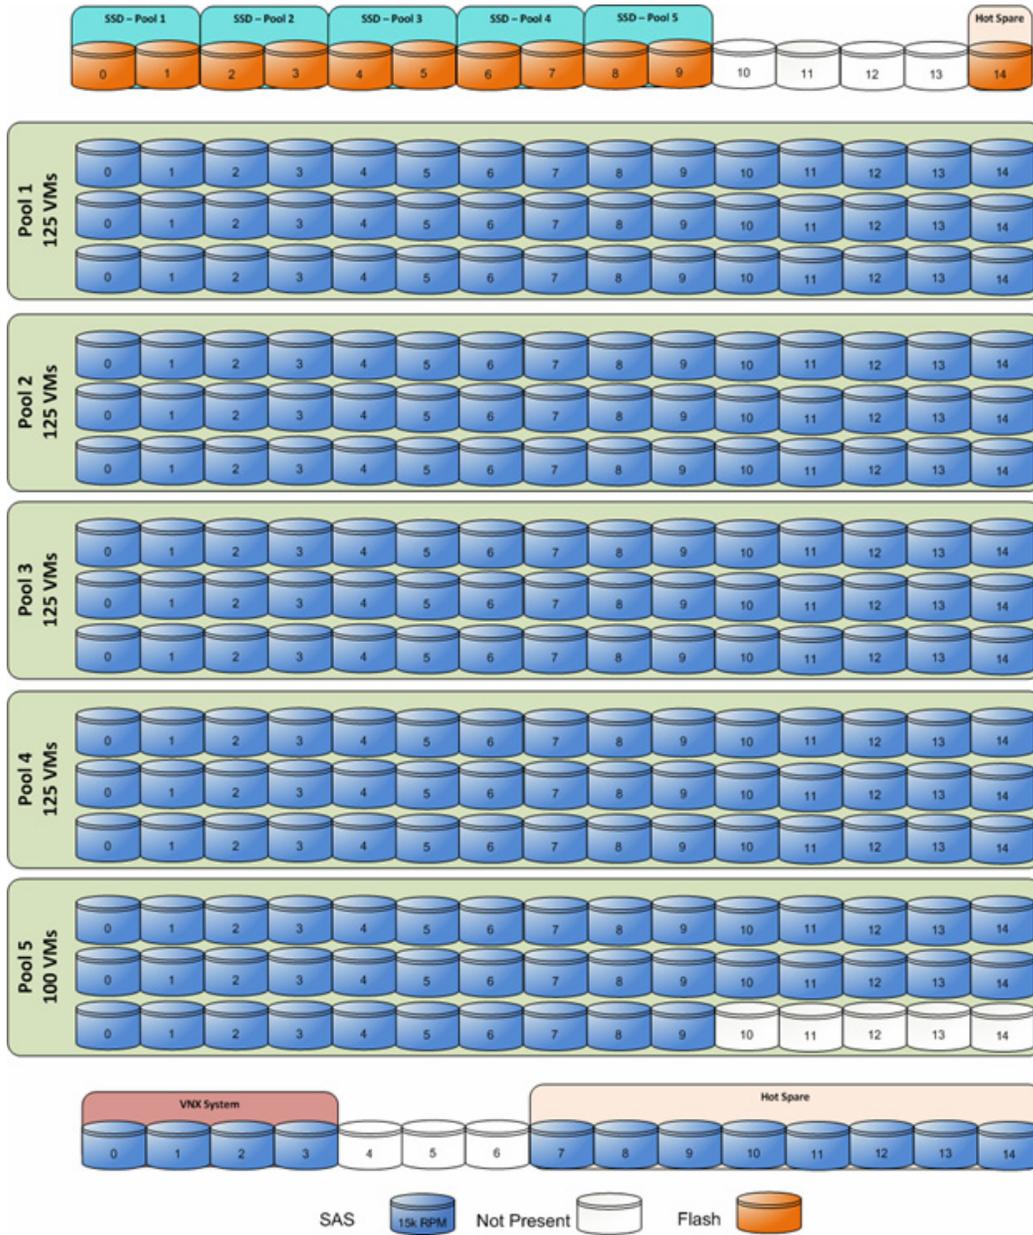


Figure 15 shows storage architecture for 1000 virtual machines on VNX5800 for NFS-variant of architecture.

Figure 15 Storage Architecture for 1000 VMs on EMC VNX5800

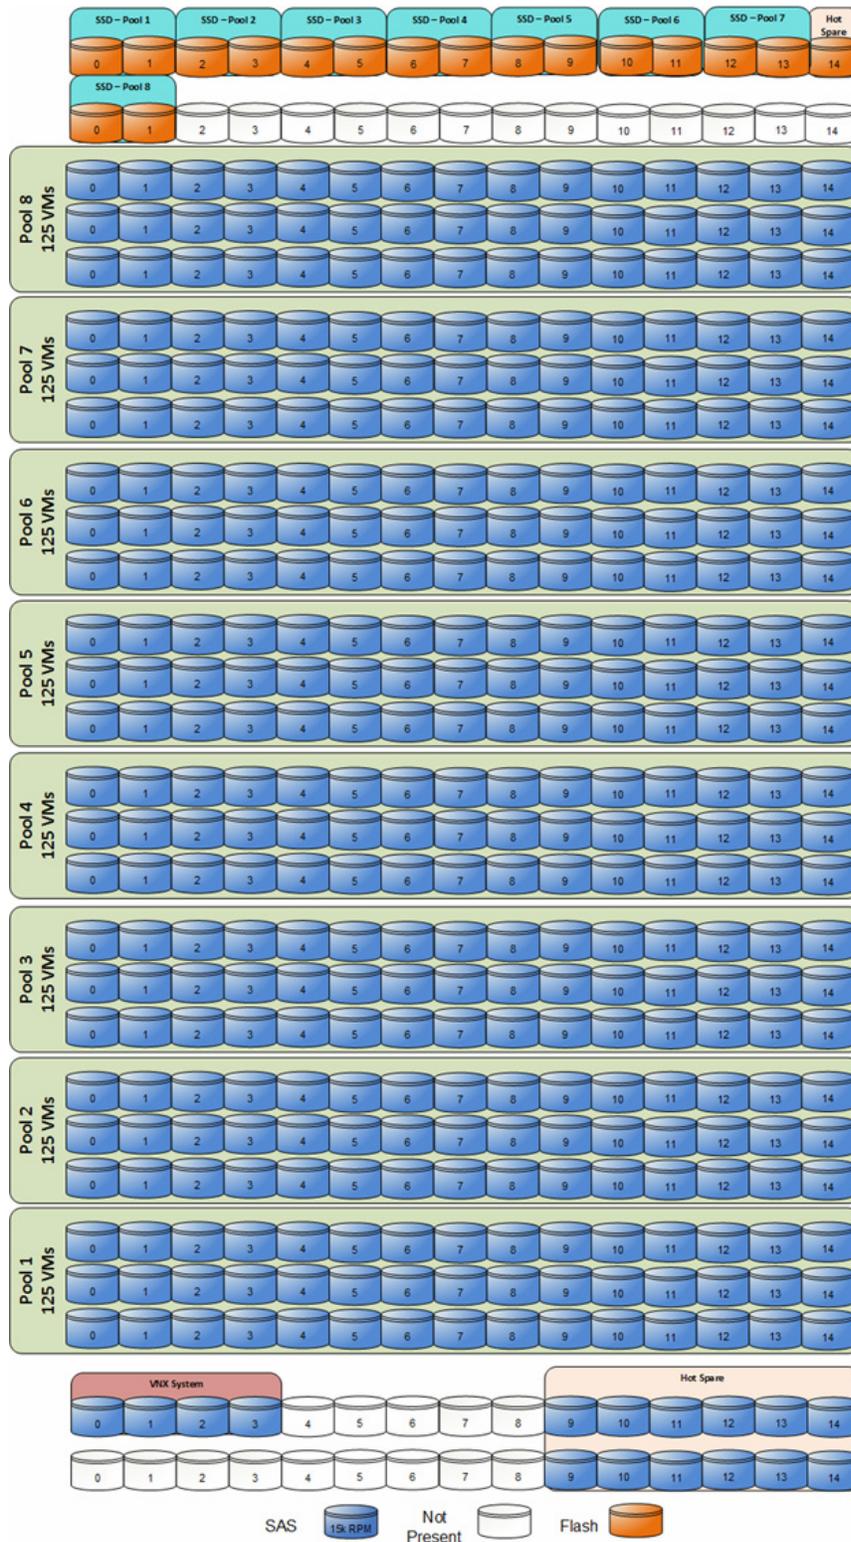


Table provides the data stores size for various architectures shown in [Figure 13](#), [Figure 14](#), and [Figure 15](#).

Table 5 Datastores for Different Scale

Parameter	300 VMs	600 VMs	1000 VMs
Storage Array	VNX5400	VNX5600	VNX5800
Disk capacity and type	600 GB SAS	600 GB SAS	600 GB SAS
Number of disks	110	220	360
RAID type	RAID 5 groups	RAID 5 groups	RAID 5 groups
Fast VP Config	6 x 200 GB Flash Drives	10 x 200 GB Flash Drives	16 x 200 GB Flash Drives
Hot spares	4 x 600 GB SAS 1 x 200 GB Flash	8 x 600 GB SAS 1 x 200 GB Flash	12 x 600 GB SAS 1 x 200 GB Flash

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.

Service Profile Design

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

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

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

- Local disks are NOT used for booting. Boot policy in service profile template suggests host to boot from the storage devices using FC protocol for both architectures.

- Server pool is defined with automatic qualification policy and criteria. Rack servers are automatically put in the pool as and when they are fully discovered by UCS Manager. This eliminates the need to manually assign servers to server pool.
- Service profile template is associated to the server pool. This eliminates the need to individually associating service profiles to physical servers.

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

- Put the faulty server in maintenance mode using vCenter. This would move VMs running on fault server to other healthy servers on the cluster.
- Disassociate the service profile from the faulty server and physically remove the server for replacement of faulty hardware (or to completely remove the faulty server).
- Physically install the new server and connect it to the Fabric Extenders. Let the new server be discovered by UCS Manager.
- Associate the service profile to the newly deployed rack server. This would boot the same ESXi server image from the storage array as what the faulty server was running.
- The new server would assume the role of the old server with all the identifiers intact. You can now end the maintenance mode of the ESXi server in vCenter.

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

Network Availability Design - FC-Variant

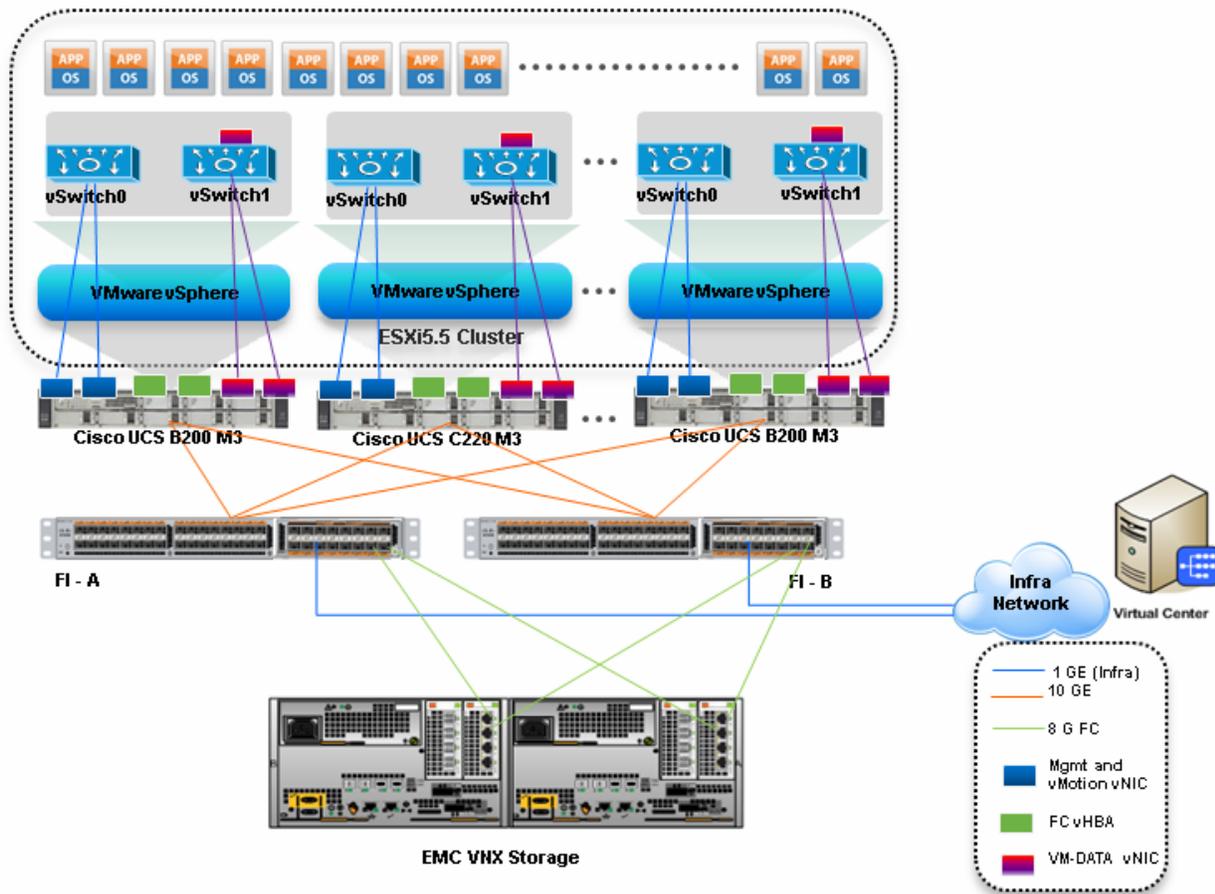
Figure 16 shows the logical layout of FC-variant architecture. The following are the key aspects of this solution:

- Mix of Cisco UCS B200 M3 and C220 M3 servers are used, managed by UCS Manager.
- Fabric A and Fabric B are used with host based FC multi-pathing for high availability
- EMC VNX5400 storage array is directly attached to UCS Fabric Interconnects
- Two 10GE links between FI and FEX provides enough bandwidth oversubscription for the SMB segment private cloud. The oversubscription can be reduced by adding more 10GE links between FI and FEX if needed by the VMs running on the ESXi hosts.
- Two vSwitches are used per host, as discussed in the Virtual Networking design section.

Storage is made highly available by deploying following practices:

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

Figure 16 Logical Layout of FC-Variant Architecture



Network Availability Design - NFS-Variant

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

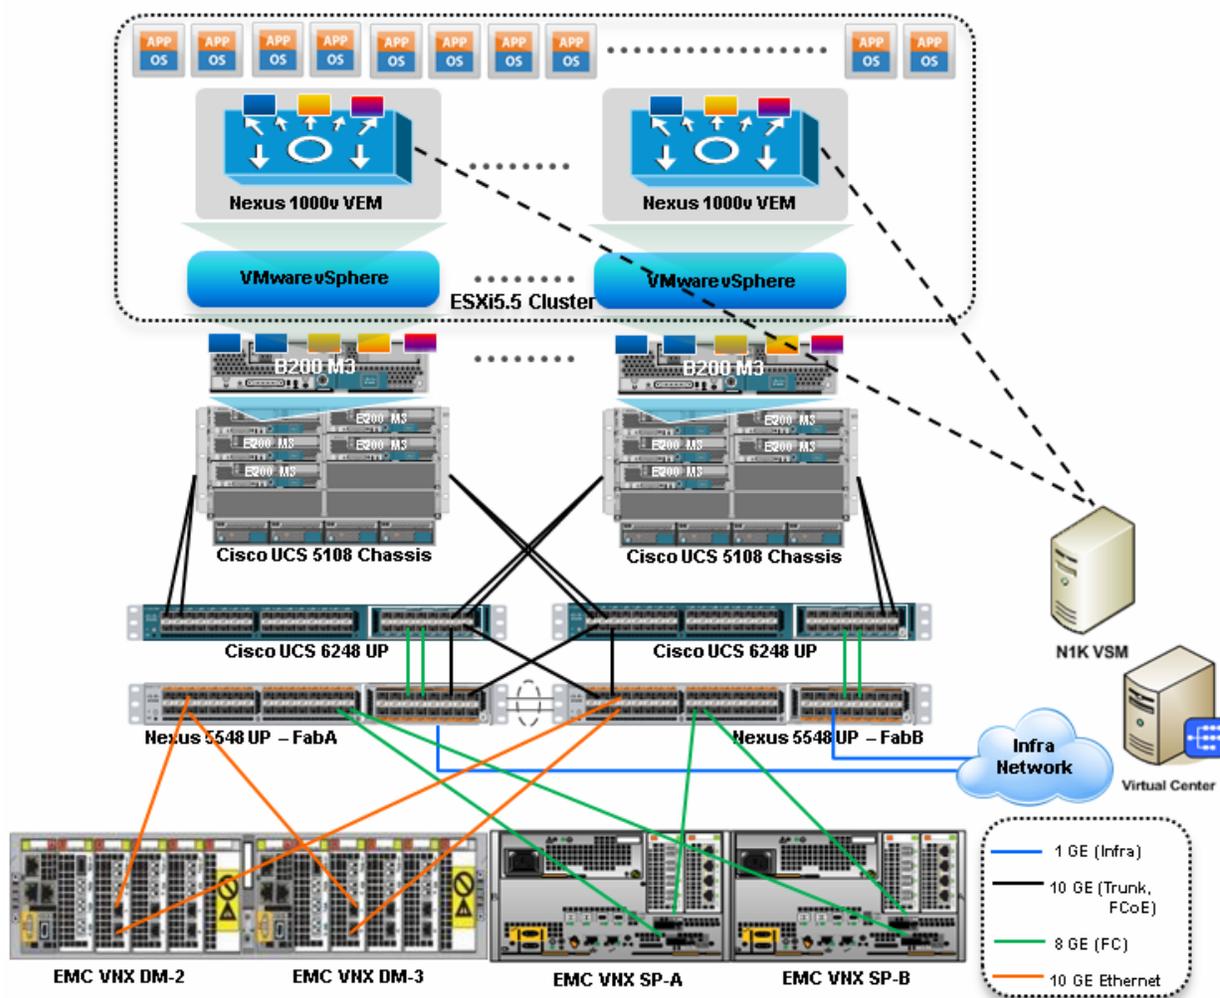
- Mix of Cisco UCS B200 M3 and C220 M3 servers are used, managed by UCS Manager
- Cisco Nexus 1000v distributed virtual switch is used for virtual switching
- vNICs on fabric A and fabric B are used for NFS based access high-availability
- vPC is used between Nexus 5548UP and UCS FIs for high availability
- vPC and port-aggregation is used between Nexus 5548UP and VNX storage for high availability
- Two 10GE links between FI and FEX provides enough bandwidth oversubscription for the private cloud. The oversubscription can be reduced by adding more 10GE links between FI and FEX if needed by the VMs running on the ESXi hosts.

Storage is made highly available by deploying following practices:

- FC access for booting the hypervisor images in the NFS-variant of architecture is exactly same as the FC-variant. Only difference is – the VMs are stored on the NFS mount servers in the NFS-variant of architecture.

- VNX storage arrays provide two Storage Processors (SPs): SP-A and SP-B for FC and two Data Movers (DMs): DM-2 and DM-3 for NFS.
- Both Nexus 5548UP switches (A and B) are connected to both Storage Processors (over FC) and Data Movers (over Ethernet), however, a given FI connects to the same port on each SP. FI-A connects to “eth10” port of SP-A and SP-B, while FI-B connects to “eth11” port of SP-A and SP-B.
- vPC on Nexus 5548UP switches and port-aggregation on VNX storage arrays for high availability of NFS servers
- Data Movers are always in the active/stand-by mode; the L2 links are up on both DMs, LACP would be down on the stand-by DM.
- If a single link on port-channel fails, the other link would bear all the load. If the whole Data Mover fails, then the standby DM takes over the role of active DM.

Figure 17 Logical Layout of NFS-Variant Architecture



Jumbo MTU

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

- NFS Storage access
- vMotion traffic

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

1. Ethernet ports on VNX Storage Processors
2. System QoS classes in Nexus 5548UP switches
3. System QoS classes in UCS Manager
4. vNICs in service profiles
5. Nexus 1000v switch or vSwitches on the ESXi hosts
6. VM-Kernel ports used for vMotion and storage access on the ESXi hosts

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 characteristics shown in [Table 4](#).

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 3GB 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 30GB on local hard drive storage.

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

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

In this example, a single virtual machine uses the resources of two of the reference VMs. Once this VM is deployed, the solution's new capability would be 298 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 16GB of memory. It uses 200GB 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 four reference VMs
- Memory of eight reference VMs
- Storage of two reference VMs
- IOPS of eight reference VMs

In this case the one virtual machine uses the resources of eight reference virtual machines. Once this VM is deployed, the solution's new capability would be 292 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 25GB 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 two reference VMs
- Memory of four reference VMs
- Storage of one reference VMs
- IOPS of two reference VMs

In this case the virtual machine would use the resources of four reference virtual machines. Once this VM is deployed, the solution's new capability would be 296 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 48GB of memory. It uses 5TB 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 52 reference virtual machines. If this was implemented on a resource pool for 300 virtual machines, there are 248 virtual machines of capability remaining in the pool.

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 300 virtual machines they can all be implemented, leaving the capacity of 236 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 section 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. Connect network cables
3. Configure Cisco Nexus 5548UP switches (NFS-variant only)
4. Prepare Cisco UCS FIs and configure Cisco UCS using UCS Manager
5. Configure data stores for ESXi images
6. Install VMware ESXi servers and vCenter infrastructure.
7. Install and configure vCenter server
8. Install Cisco Nexus 1000v VMS VM (NFS-variant only)
9. Configure storage for VM data stores, install and instantiate VMs through vCenter
10. 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 6](#) to confirm that all equipment, software, and appropriate licenses are available before the deployment process.

- Gather data—Collect the customer-specific configuration data for networking, naming, and required accounts. Enter this information into the [Customer Configuration Data Sheet, page 200](#) for reference during the deployment process.

Table 6 **Customer Specific Configuration Data**

Requirement	Description	Reference
Hardware	Cisco UCS B200M3 or C220M3 servers to host virtual machines	See the corresponding product documentation
	Cisco UCS 5108 Blade Server Chassis	
	Cisco Nexus 2232PP Fabric Extender	
	Cisco UCS 6248UP Fabric Interconnect	
	VMware vSphere™ 5.1 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: Multiprotocol storage array with the required disk layout as per architecture requirements	

Table 6 Customer Specific Configuration Data

Requirement	Description	Reference
Software	Cisco Nexus 1000v VSM and VEM installation media	See the corresponding product documentation
	VMware ESXi™ 5.1 installation media	
	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 2012 installation media (suggested OS for VMware vCenter)	
	Microsoft SQL Server 2008 R2 SP1 Note This requirement may be covered in the existing infrastructure	
Licenses	VMware vCenter 5.1 license key	Consult your corresponding vendor 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 200](#) 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.

Connect Network Cables

See the Cisco Nexus 5548UP, UCS FI, FEX, Blade server chassis, B-series and C-series server and EMC VNX storage array configuration guide for detailed information about how to mount the hardware on the rack. Following diagrams show connectivity details for the VSPEX VMware architecture covered in this document.

Connectivity for FC-Variant:

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

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

Figure 18 Detailed Connectivity Diagram of the FC-Variant Architecture

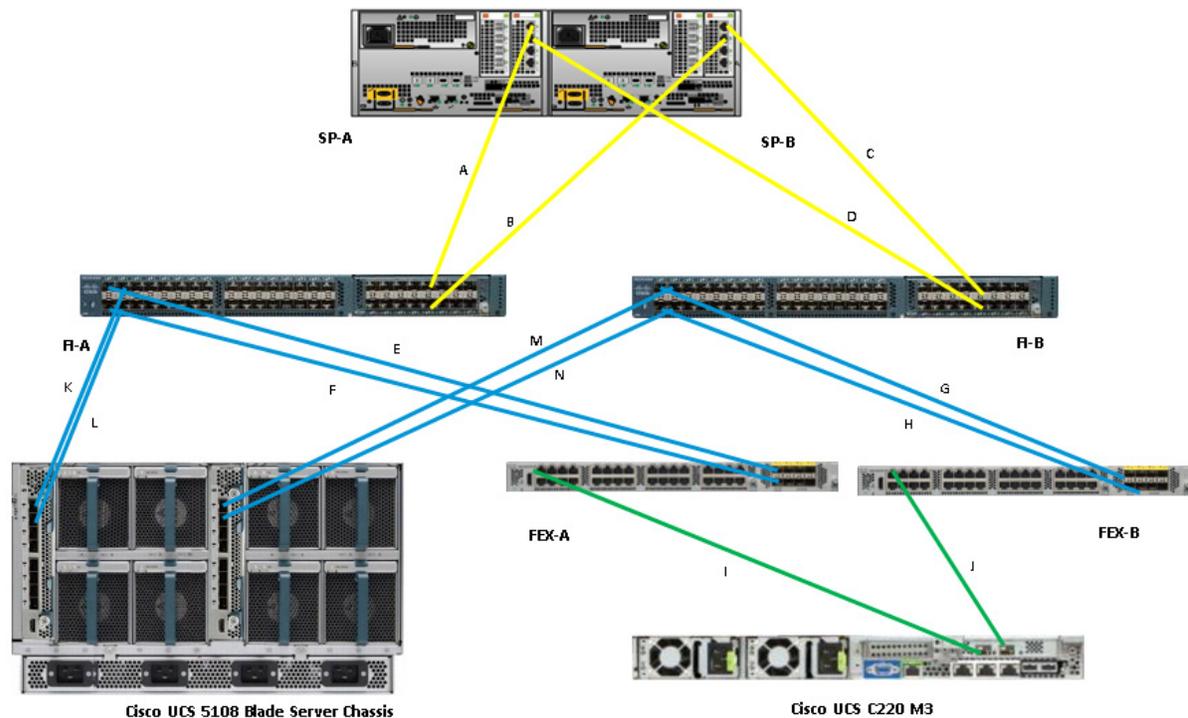


Figure 19 elaborates the detailed cable connectivity for the architecture.

Figure 19 Connectivity Details of FC-Variant Architecture

Cable ID	Peer1	Peer2	VLAN	Mode	Description
A	FI-A, FC 2/9	SP-A,	Storage VSAN	Appliance	Directly attached storage on FI
B	FI-A, FC 2/10	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
C	FI-B, FC 2/9	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
D	FI-B, FC 2/10	Sp-A,	Storage VSAN	Appliance	Directly attached storage on FI
E,F	FI-A, Eth 1/1, 1/2	FEX-A uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
G,H	FI-A, Eth 1/1, 1/2	FEX-B uplinks	N/A	Server	FI/FEX 20GE port-channel connectivity
I	FEX-A, port 1	C220-M3 VIC port 1	N/A	VNTag (internal)	Server to fabric A. VLANs are allowed on per vNIC basis
J	FEX-B, port 1	C220-M3 VIC port 2	N/A	VNTag (internal)	Server to fabric B. VLANs are allowed on per vNIC basis
K,L	FI-A, Eth 1/3, 1/4	5108 Chassis, FEX 2208 Left	N/A	Server	FI/FEX 20GE port-channel connectivity
M,N	FI-B, Eth 1/3, 1/4	5108 Chassis, FEX 2208 Right	N/A	Server	FI/FEX 20GE port-channel connectivity
(not shown)	Eth 2/1, 2/2 on FI-A and FI-B	Uplink switch	All	Uplink	Uplink to Infrastructure network

Connectivity for NFS-Variant

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

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

Figure 20 Detailed Connectivity Diagram of the NFS-Variant Architecture

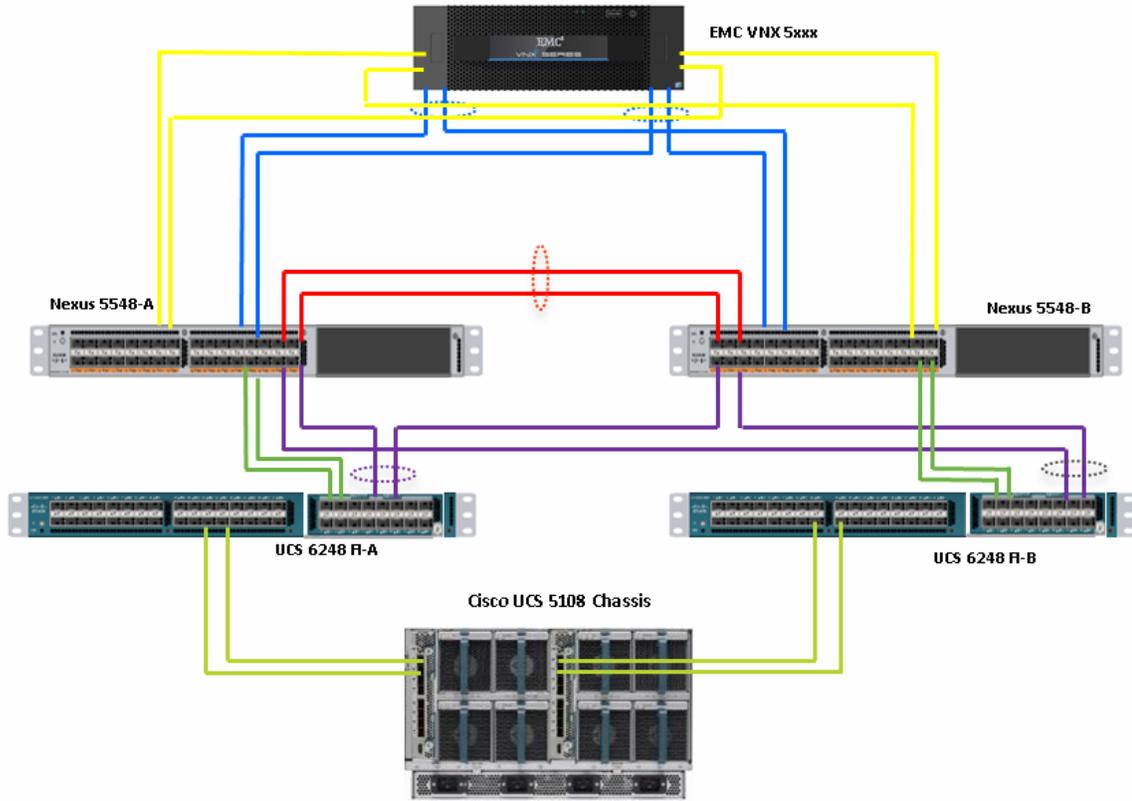


Figure 21 elaborates the detailed cable connectivity for the architecture.

Figure 21 Ethernet Cable Connectivity Details of NFS-Variant Architecture

Cable ID	Peer1	Peer2	VLAN	Mode	Description
A	FI-A, Eth 1/5	DM-2, A1/0	Storage	Access (on N5k)	N5k to storage Data Mover
B	FI-A, Eth 1/6	DM-3, B1/0	Storage	Access (on N5k)	N5k to storage Data Mover
C	FI-B, Eth 1/5	DM-2, A1/1	Storage	Access (on N5k)	N5k to storage Data Mover
D	FI-B, Eth 1/6	DM-3, B1/1	Storage	Access (on N5k)	N5k to storage Data Mover
E,F	Eth 1/15, Eth 1/16 on N5k-A	Eth 1/15, Eth 1/16 on N5k-B	All	Trunk	vPC peer-links between N5ks
G	N5k-A, Eth 1/3	FI-A, Eth 1/13	All	Trunk	N5k vPC member port, FI PC member port
H	N5k-A, Eth 1/4	FI-B, Eth 1/13	All	Trunk	N5k vPC member port, FI PC member port
I	N5k-B, Eth 1/3	FI-A, Eth 1/14	All	Trunk	N5k vPC member port, FI PC member port
J	N5k-B, Eth 1/4	FI-B, Eth 1/14	All	Trunk	N5k vPC member port, FI PC member port
K,L	FI-A, Eth 1/1, 1/2	FEX-A uplinks	N/A	Server (on FI)	FI / IOM links
M,N	FI-B, Eth 1/1, 1/2	FEX-B uplinks	N/A	Server (on FI)	FI / IOM links

Figure 22 FC Cable Connectivity Details for NFS-Variant Architecture

Cable ID	Peer1	Peer2	VSAN	Description
O	N5k-A, FC 1/31	SP-A, 0/0	Storage VSAN	N5k to storage SP, crisscrossed
P	N5k-A, FC 1/32	SP-B, 0/0	Storage VSAN	N5k to storage SP, crisscrossed
Q	N5k-B, FC 1/31	SP-A, 0/1	Storage VSAN	N5k to storage SP, crisscrossed
R	N5k-B, FC 1/32	SP-A, 0/1	Storage VSAN	N5k to storage SP, crisscrossed
S	N5k-A, FC1/29	FI-A, FC1/29	Storage VSAN	N5k to FI straight cables
T	N5k-A, FC1/30	FI-A, FC1/30	Storage VSAN	N5k to FI straight cables
U	N5k-B, FC1/29	FI-B, FC1/29	Storage VSAN	N5k to FI straight cables
V	N5k-B, FC1/30	FI-B, FC1/30	Storage VSAN	N5k to FI straight cables

By connecting all the cables as outlined above, and you would be ready to configure Cisco Nexus 5548UP Switch, EMC VNX Series Storage Array and Cisco UCS Manager.

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 *Cisco Nexus 5000 Series Configuration Guide*.

Configure Global VLANs and VSANs

Figure 23 shows how to configure VLAN on a switch.

Figure 23 **Creating VLAN**

```
UCS-N5k-FabA# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
UCS-N5k-FabA(config)# vlan 40
UCS-N5k-FabA(config-vlan)# name Storage
UCS-N5k-FabA(config-vlan)# exit
UCS-N5k-FabA(config)# exit
UCS-N5k-FabA#
```

Following VLANs in Table 7 need to be configured on both switches A and B in addition to your application specific VLANs:

Table 7 **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 200](#).

We have used one VSAN in this solution. Table 8 gives the VSAN name and the description.

Table 8 **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 200](#).

Figure 24 and Figure 25 show the creation of VSAN and assigning VSAN to the fibre channel interface.

Figure 24 *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 25](#).

Figure 25 *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                                     Mode  Speed  Channel
          (Gbps)
-----
fc1/29     10     F      on     up          sw1    F     8     --
fc1/30     10     F      on     up          sw1    F     8     --
fc1/31     10     F      on     up          sw1    F     8     --
fc1/32     10     F      on     up          sw1    F     8     --
UCS-N5k-FabA#
```

It is also crucial to enable NPIV feature on the Cisco Nexus 5548UP switches. [Figure 26](#) show how to enable NPIV feature on Nexus 5548UP switches.

Figure 26 *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 27](#).

Figure 27 **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 28](#). Ensure that “vpc peer-link” is configured on this port-channel.

Figure 28 **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 vPC peer-link.
This will enable spanning tree Bridge Assurance on vPC peer-link provided the STP 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 29](#).

Figure 29 Window Showing Successful vPC Configuration

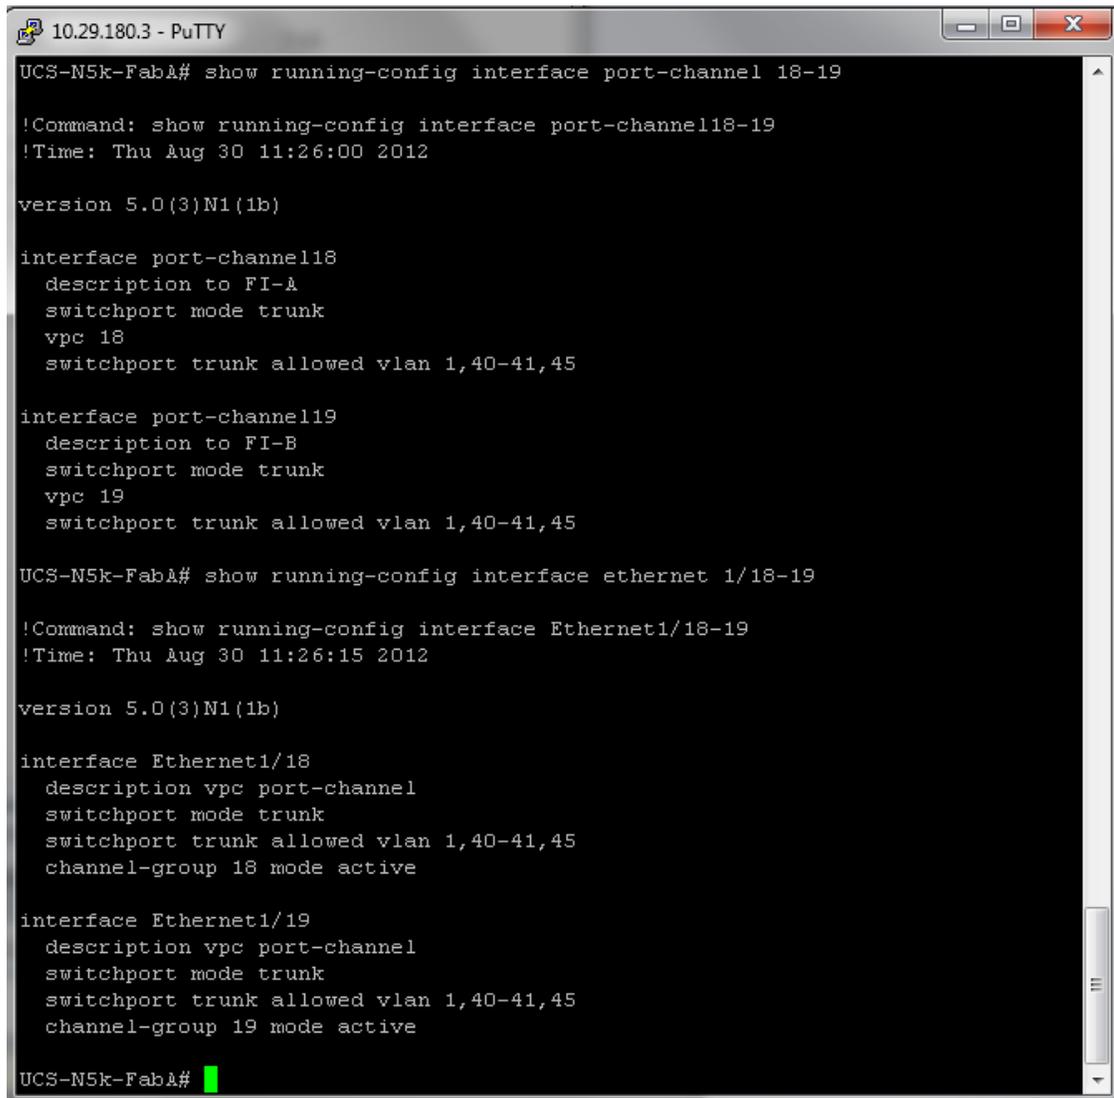
```
UCS-N5k-FabA# show vpc
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

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 30](#) shows the configuration commands.

Figure 30 Port-channel Configuration


```

10.29.180.3 - PuTTY
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

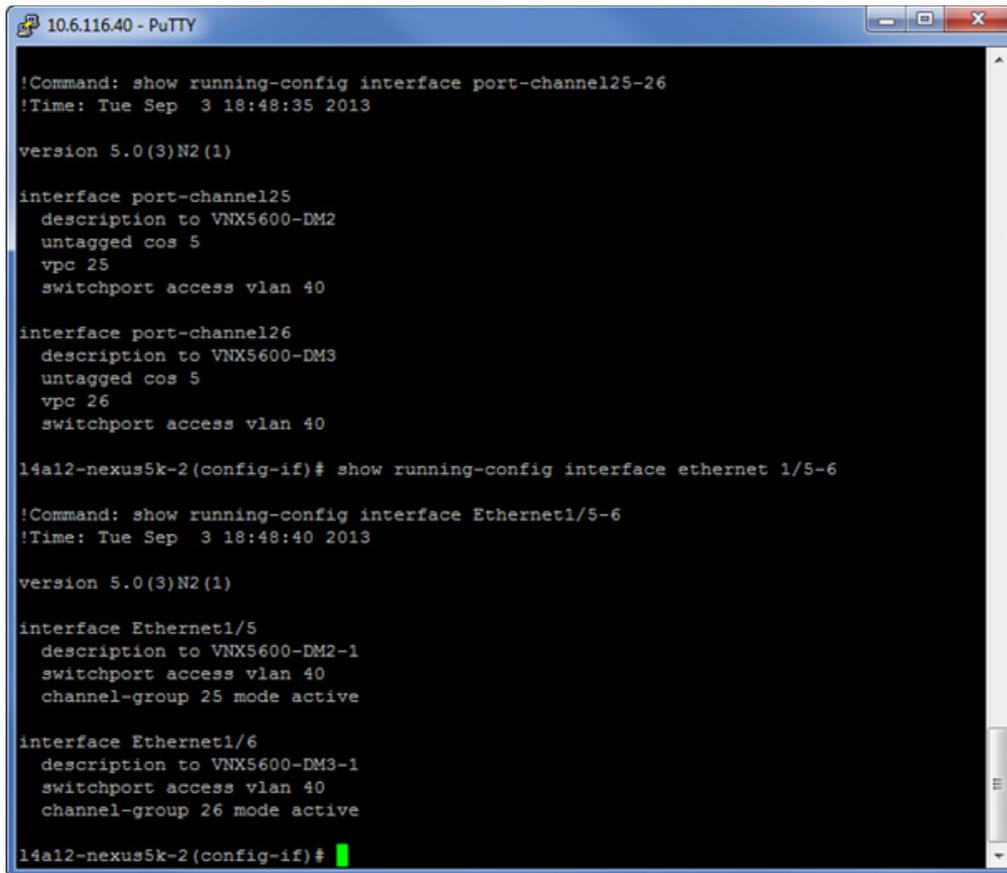
From each switch one link connects to each storage processor on the VNX storage array. A virtual port-channel is created for the two links connected to a single storage processor, but connected to two different switches. In this example configuration, links connected to the storage processor A (SP-A) of VNX 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 31 shows the configuration on the port-channels and interfaces.

Figure 31 Configuration of Port-channel and Interfaces



```

10.6.116.40 - PuTTY
!Command: show running-config interface port-channel25-26
!Time: Tue Sep 3 18:48:35 2013

version 5.0(3)N2(1)

interface port-channel25
  description to VNX5600-DM2
  untagged cos 5
  vpc 25
  switchport access vlan 40

interface port-channel26
  description to VNX5600-DM3
  untagged cos 5
  vpc 26
  switchport access vlan 40

l4a12-nexus5k-2(config-if)# show running-config interface ethernet 1/5-6

!Command: show running-config interface Ethernet1/5-6
!Time: Tue Sep 3 18:48:40 2013

version 5.0(3)N2(1)

interface Ethernet1/5
  description to VNX5600-DM2-1
  switchport access vlan 40
  channel-group 25 mode active

interface Ethernet1/6
  description to VNX5600-DM3-1
  switchport access vlan 40
  channel-group 26 mode active

l4a12-nexus5k-2(config-if)#

```

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 32](#).



Note

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

Figure 32 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

```

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 32](#). 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 33](#) shows the expected output of this command.

Figure 33 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-Channel	Type	Protocol	Member Ports
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 34.

Figure 34 Verifying VPC Status

```
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 : 4
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

vPC status
-----
id  Port  Status Consistency Reason Active vlans
--  --
18  Po18  up     success success 1,40-41
19  Po19  up     success success 1,40-41
25  Po25  up     success success 40
26  Po26  down*  success success -

UCS-N5k-FabA#
```

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”, except one of the two port-channels connected to the storage array as explained before.

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.
4. NFS traffic flowing from storage array to fabric interconnect need to be marked with Ethernet Class of Service (CoS) 5 for proper classification at the fabric interconnect. Use the “untagged cos 5” command at the port-channels connected to the storage arrays.

Figure 35 shows the exact Cisco Nexus CLI for the steps mentioned above.

Figure 35 **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 35 shows the NX-OS interface range CLI to configure multiple interfaces at the same time.

Prepare UCS FIs and configure UCS Manager

Configure UCS FIs and UCS Manager can be subdivided in to following segments:

1. [Initial Configuration of Cisco UCS FIs, page 51](#)
2. [Configuration for Server Discovery, page 53](#)
3. [Upstream/ Global Network Configuration, page 58](#)
4. [Configure Identifier Pools, page 78](#)
5. [Configure Server Pool and Qualifying Policy, page 85](#)
6. [Configure Service Profile Template, page 91](#)
7. [Instantiate Service Profiles from the Service Profile Template, page 108](#)

Following subsections provided details on each of the steps mentioned above.

Initial Configuration of Cisco UCS FIs

At this point of time, the Cisco UCS FIs, FEX, and Blade Servers or Rack Servers must be mounted on the rack and appropriate cables must be connected. Two 100 Mbps Ethernet cables must be connected between two FIs for management pairing. Two redundant power supplies are provided per FI, it is highly

recommended that both the power supplies are plugged in, ideally drawing power from two different power strips. Connect mgmt0 interfaces of each FI to the infrastructure network, and put the switch port connected to FI in access mode with access VLAN as management VLAN.

To perform initial FI configuration, follow these steps:

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

Figure 36 Initial Configurations of Cisco UCS Fabric Interconnect

```

10.65.121.10 - PuTTY
Enter the configuration method. (console/gui) ? console
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
Enforce strong password? (y/n) [y]:
Enter the password for "admin":
Confirm the password for "admin":
Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes
Enter the switch fabric (A/B) []: A
Enter the system name: VSPEX-FI
Physical Switch Mgmt0 IPv4 address : 10.65.121.226
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 10.65.121.1
Cluster IPv4 address : 10.65.121.228
Configure the DNS Server IPv4 address? (yes/no) [n]:
Configure the default domain name? (yes/no) [n]:
Following configurations will be applied:
Switch Fabric=A
System Name=VSPEX-FI
Enforced Strong Password=yes
Physical Switch Mgmt0 IP Address=10.65.121.226
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.65.121.1
Cluster Enabled=yes
Cluster IP Address=10.65.121.228
NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): █

```

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

Figure 37 **Configuring Peer a Fabric Interconnect**

```

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

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

Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IP Address: 10.65.121.226
Peer Fabric interconnect Mgmt0 IP Netmask: 255.255.255.0
Cluster IP address      : 10.65.121.228

Physical Switch Mgmt0 IPv4 address : 10.65.121.227

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

```

3. Once initial configurations on both FIs are completed, you can disconnect the serial console cable. Now, UCS Manager will be accessible through web interface (<https://<ucsm-virtual-ip>/>) or SSH. Connect to UCS Manager using SSH, and see HA status. As there is common device connected between two FIs (a rack server or blade server chassis), the status shows as “HA NOT READY”, but you must see both FI A and FI B in “Up” state as shown [Figure 38](#).

Figure 38 **Cisco UCS Fabric Interconnect - Cluster State**

```

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

A: UP, PRIMARY
B: UP, SUBORDINATE

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

```

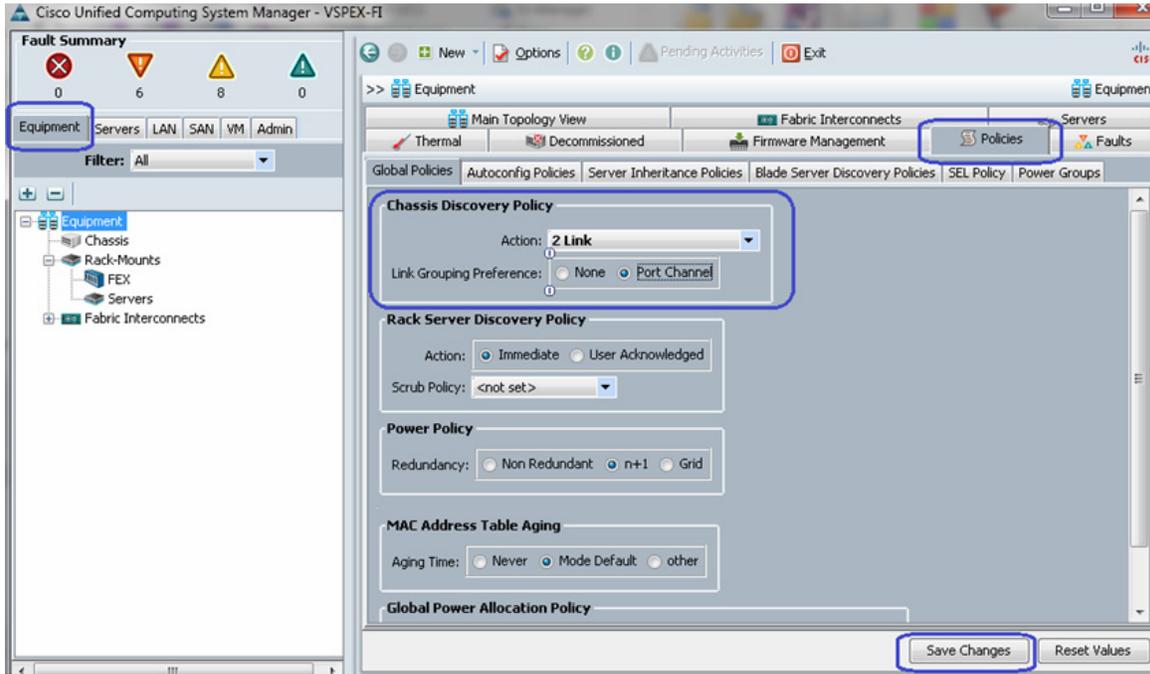
Configuration for Server Discovery

All the Ethernet ports of FIs are unconfigured and shutdown by default. You need to classify these ports as server facing ports, directly attached storage array facing ports, and uplink ports.

To configure the ports for proper server auto-discovery, follow these steps:

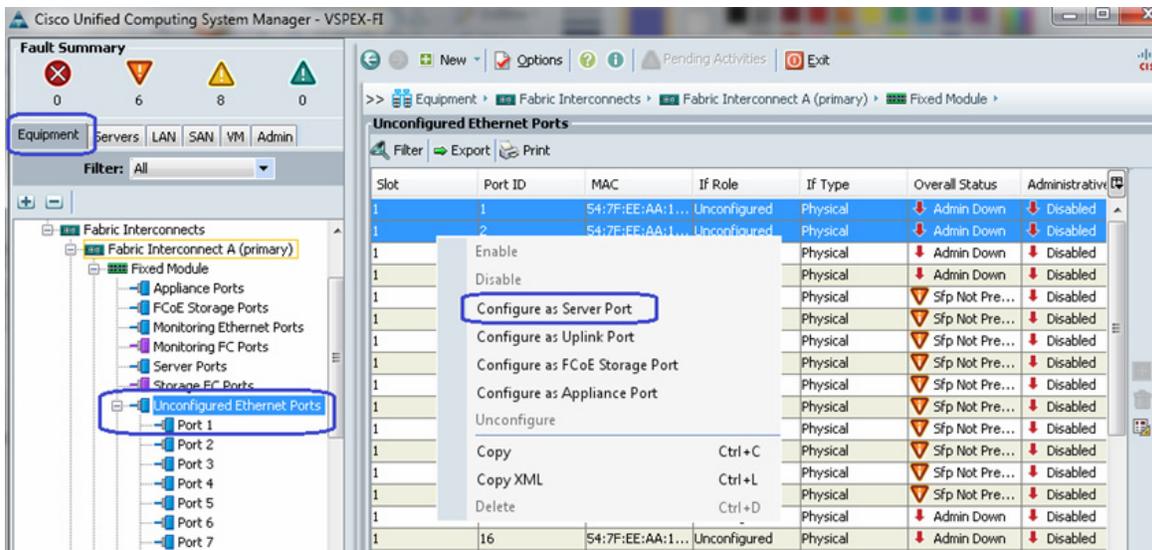
1. To configure chassis discovery policy that specifies server side connectivity, using a web browser, access the UCS Manager from the management virtual IP address and download the Java applet to launch UCS Manager GUI. Click **Equipment** tab in the left pane, and then **Policies** tab in the right pane. In Chassis Discovery Policy, For Actions field choose **2 Link**. Two links represent the two 10 GE links that are connected between FI and FEX per fabric. Also, change Link Grouping Preference to **Port Channel** for better bandwidth utilization and link level high-availability as shown in [Figure 39](#). Save the changes.

Figure 39 Configuring Chassis Discovery Policy



- Next, identify ports connected to the Chassis or FEX per FI basis. Click the **Equipment** tab, expand **Fabric Interconnects**, choose an FI, for example, Fabric Interconnect A, click **Unconfigured Ethernet Ports**, and select the two ports connected to the FEX-A. Right-click, and choose **Configure as Server Port**. Click **Yes** on the confirmation pop-up window.

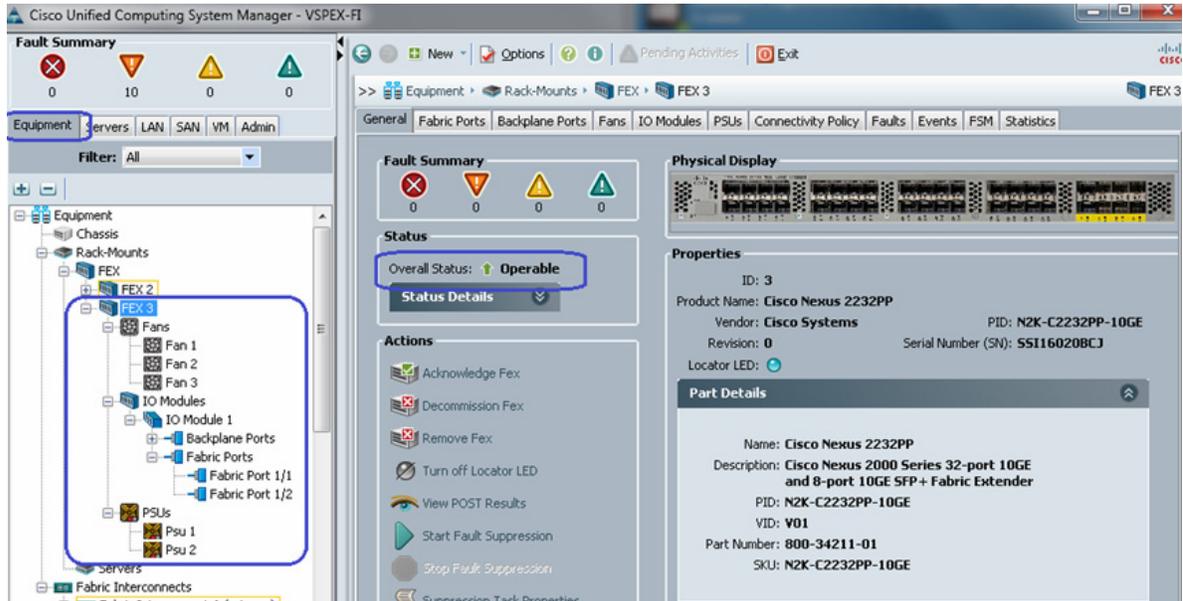
Figure 40 Configuring Ethernet Ports as Server Ports



- Repeat step 2 for the other FI as well.

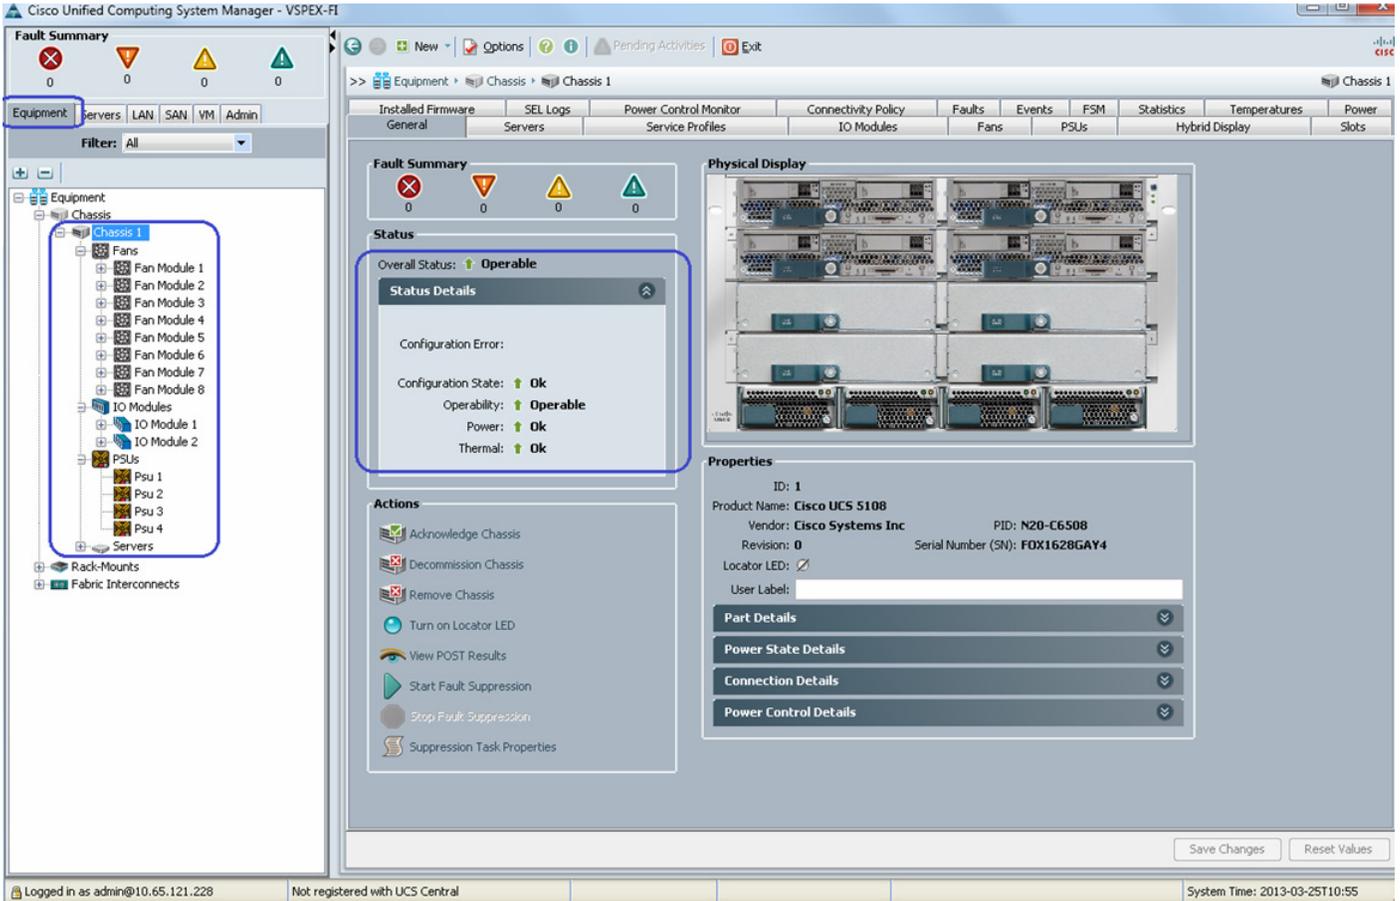
- Once server ports are configured on both FIs, the Chassis or FEX auto-discovery gets started. In case of FEX, after the deep discovery of FEX is complete, you will see two Fabric Extenders in the **Equipment** tab with overall status shown as Operable.

Figure 41 Overall Status of FEX After Auto-Discovery



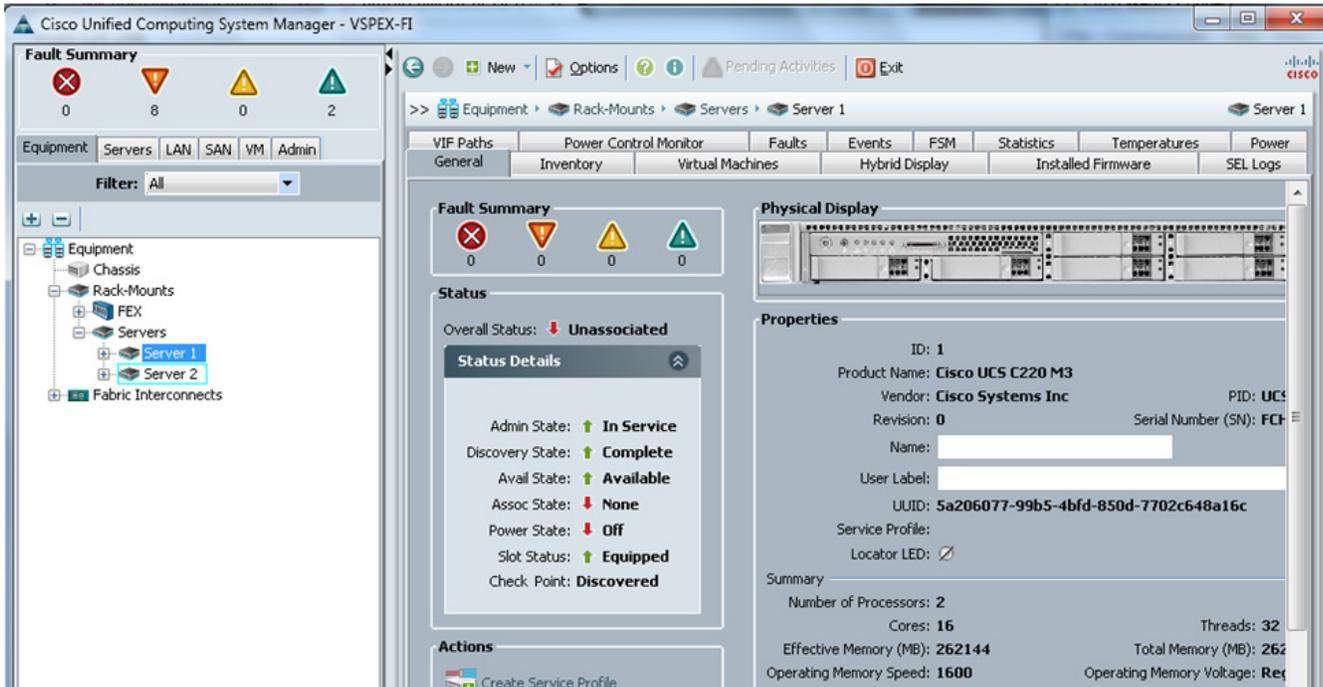
Similarly, if server ports are connected to the chassis, you will see that the chassis is fully discovered, with all its IOMs, fans, power supplies and so on.

Figure 42 Overall Status of Chassis After Auto-Discovery



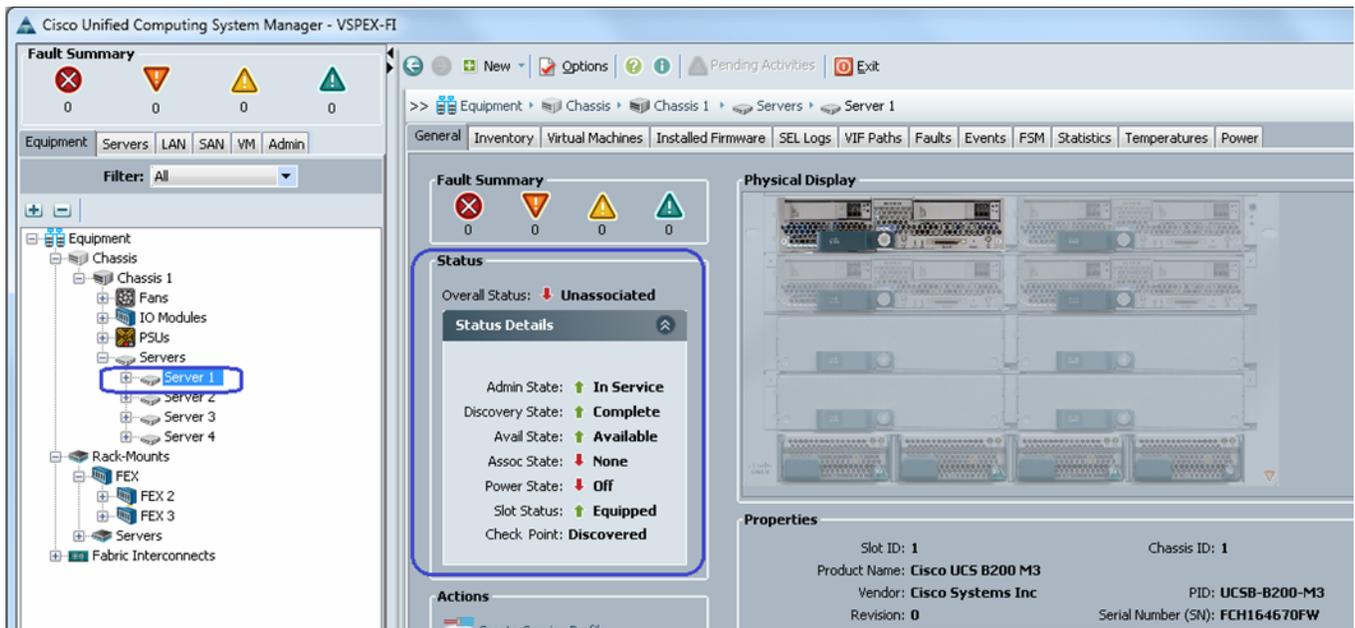
- After the Chassis and FEX auto-discovery, the Blade Server and Rack Server auto-discovery will get started respectively. As and when the servers are discovered, you will see them getting added in the **Equipment** tab with overall status shown as Unassociated and availability state as Available, and discovery state as Complete.

Figure 43 Overall Status of Rack Servers After Discovery



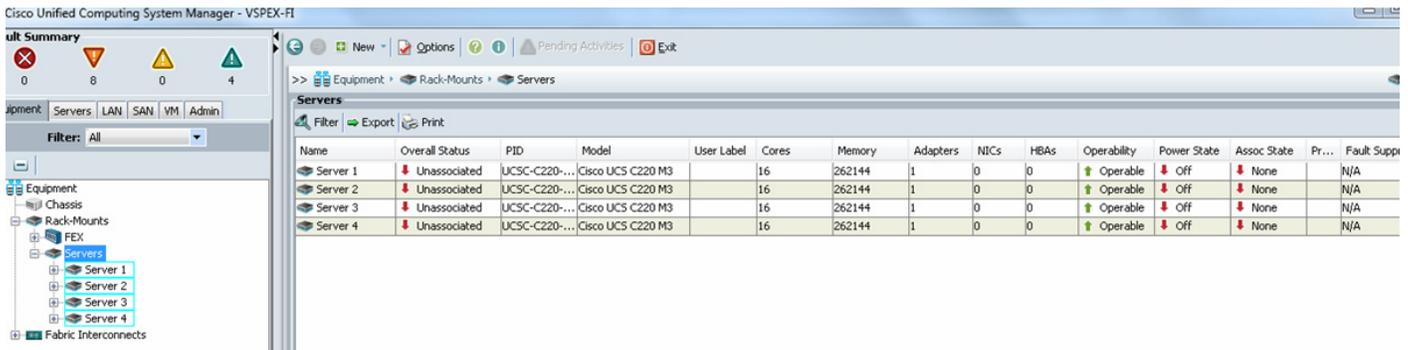
Similarly, a Blade Server’s status is shown in Figure 44.

Figure 44 Overall Status of Blade Servers After Discovery



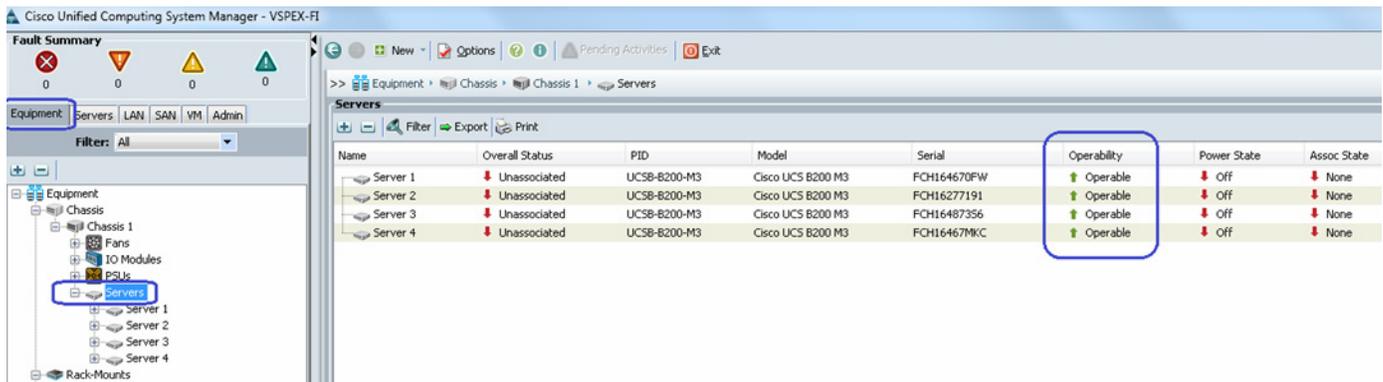
- Once all the servers are discovered, you can see the summary of all of them by choosing **Equipment** tab > **Rack-Mounts** > **Servers** as shown below.

Figure 45 Summary of Rack Servers After the Discovery



Or, in case of Blade Servers, you can see the summary by choosing **Equipment** tab > **Chassis** > **Chassis <id>** > **Servers**.

Figure 46 Summary of Blade Servers After the Discovery



Upstream/ Global Network Configuration

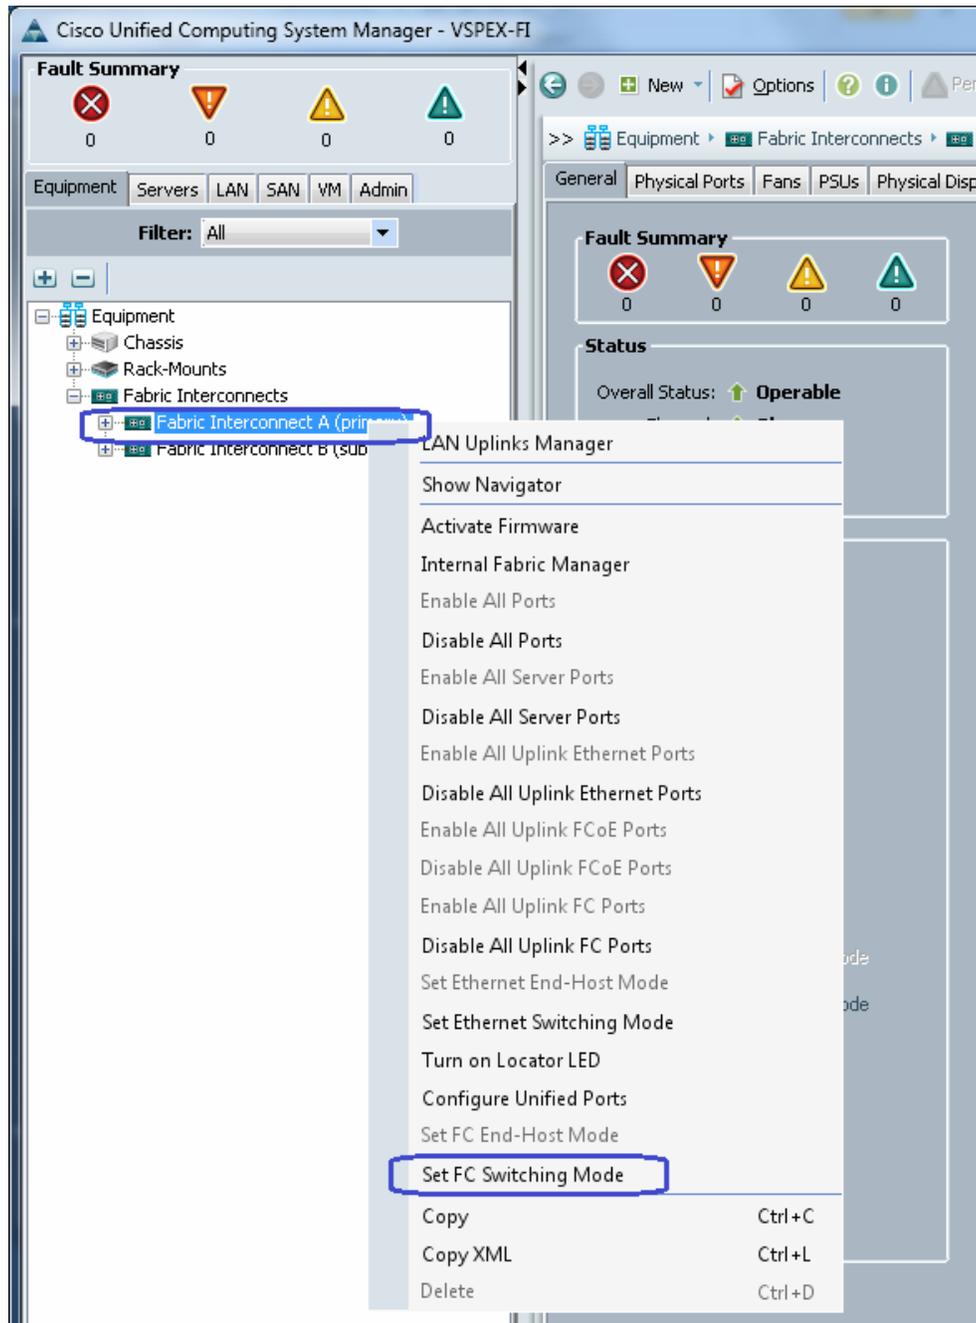
This subsection lists a few upstream/ global network configuration:

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

To configure upstream/ global network, follow these steps:

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

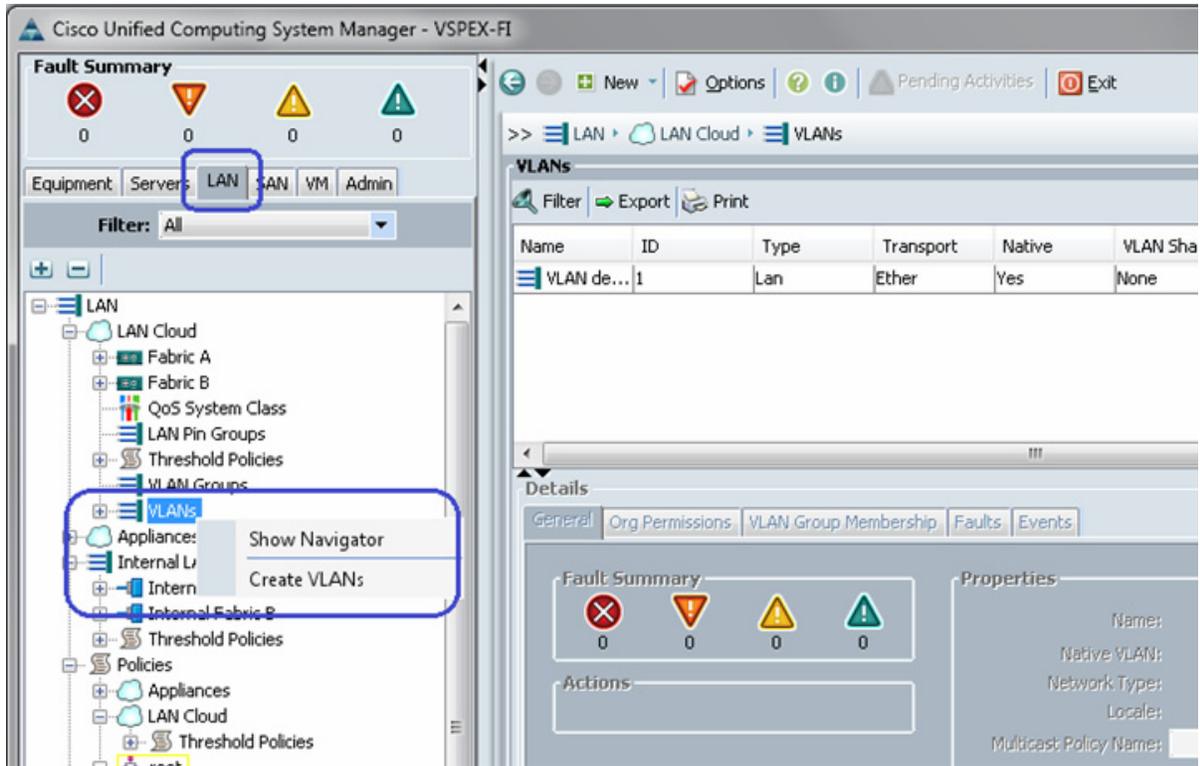
Figure 47 Setting FC Switching Mode on Fabric A for FC-Variant



2. (FC-variant only) You might see a warning message that Fabric Interconnects need to be restarted as a result of this action. Click **Yes**. Both the FIs will reboot (first the secondary FI and then the primary FI). This action is traffic disruptive, so make sure that you perform this operation during maintenance window, if you are working in a production environment.

- From the LAN tab, expand LAN > LAN Cloud, and right-click on VLANs, and choose **Create VLANs**.

Figure 48 *Creating VLANs*



- Enter the name of the VLAN and assign the VLAN ID. Keep the VLAN as default with the option Common/Global.

Figure 49 VLAN Details for Creating VLAN

Create VLANs

VLAN Name/Prefix:

Multicast Policy Name:

Common/Global
 Fabric A
 Fabric B
 Both Fabrics Configured Differently

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

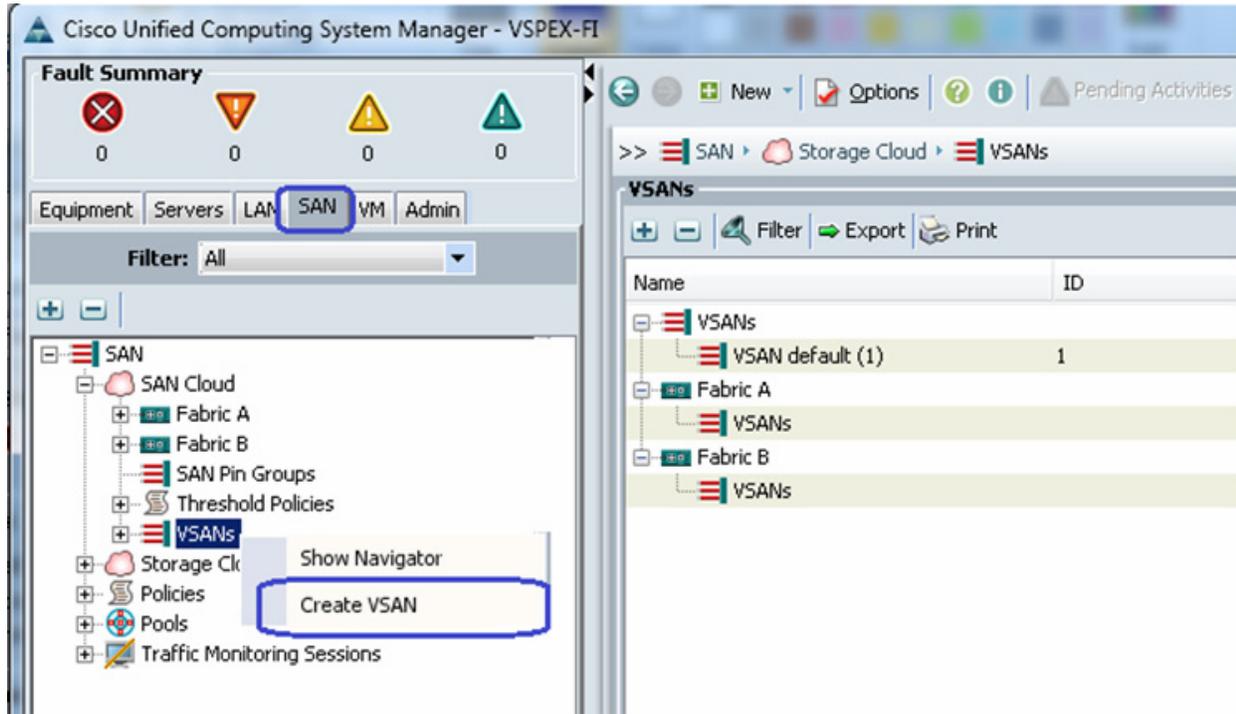
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs:

Sharing Type: None Primary Isolated

5. Click **Ok** and deploy the VLAN. Repeat this steps for vSphereMgmt, VM-Data and vMotion VLANs. For NFS-variant, create the Storage VLAN.
6. (NFS-variant only) NFS-variant of the architecture uses NFS for VM data access, but still uses FC SAN boot for the hypervisors. In the **SAN** tab, expand **SAN Cloud**, and right-click on VSANs. Choose **Create VSAN**.

Figure 50 Creating VSAN for NFS-Variant



7. (NFS-variant only) Enter a VSAN name in the Name field and provide VSAN ID and its corresponding FCoE VLAN ID. FCoE VLAN ID should not have conflict with any of the VLANs configured before. Keep the FC zoning disabled (default setting).

Figure 51 VSAN Details for Creating VSAN

Create VSAN

Name:

FC Zoning Settings

FC Zoning: Disabled Enabled

Do **NOT** enable local zoning if fabric interconnect is connected to an upstream FC/FCoE switch.

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

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

VSAN ID:

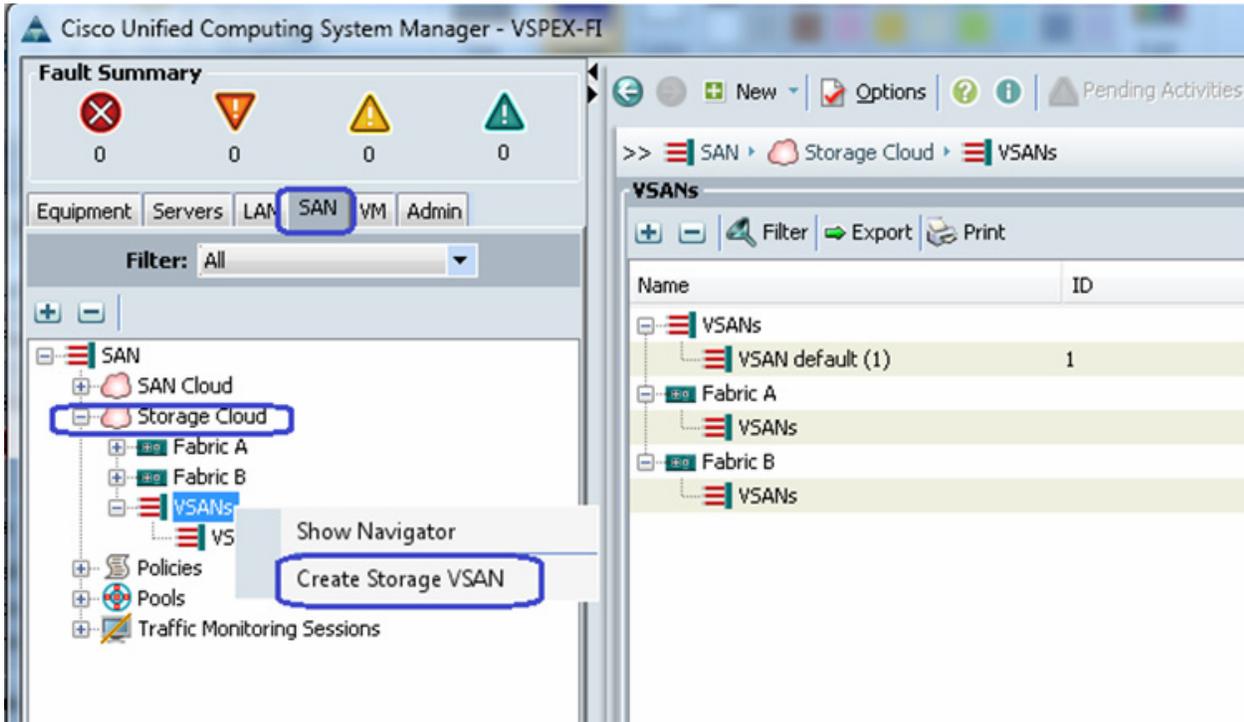
A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN:

OK Cancel

8. (FC-variant only) Click the **SAN** tab, and expand **Storage Cloud**, and right-click on VSANs. Choose **Create Storage VSAN**.

Figure 52 Creating Storage VSAN for FC-Variant



- (FC-variant) Enter the VSAN name in the Name field, enable FC zoning and provide VSAN ID and its corresponding FCoE VLAN ID. FCoE VLAN ID should not have conflict with any of the VLANs configured before.

Figure 53 VSAN Details for Creating Storage VSAN

Create Storage VSAN

Name: **Storage**

FC Zoning Settings

FC Zoning: Disabled Enabled

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

Common/Global Fabric A Fabric B Both Fabrics Configured Differently

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

VSAN ID: **10**

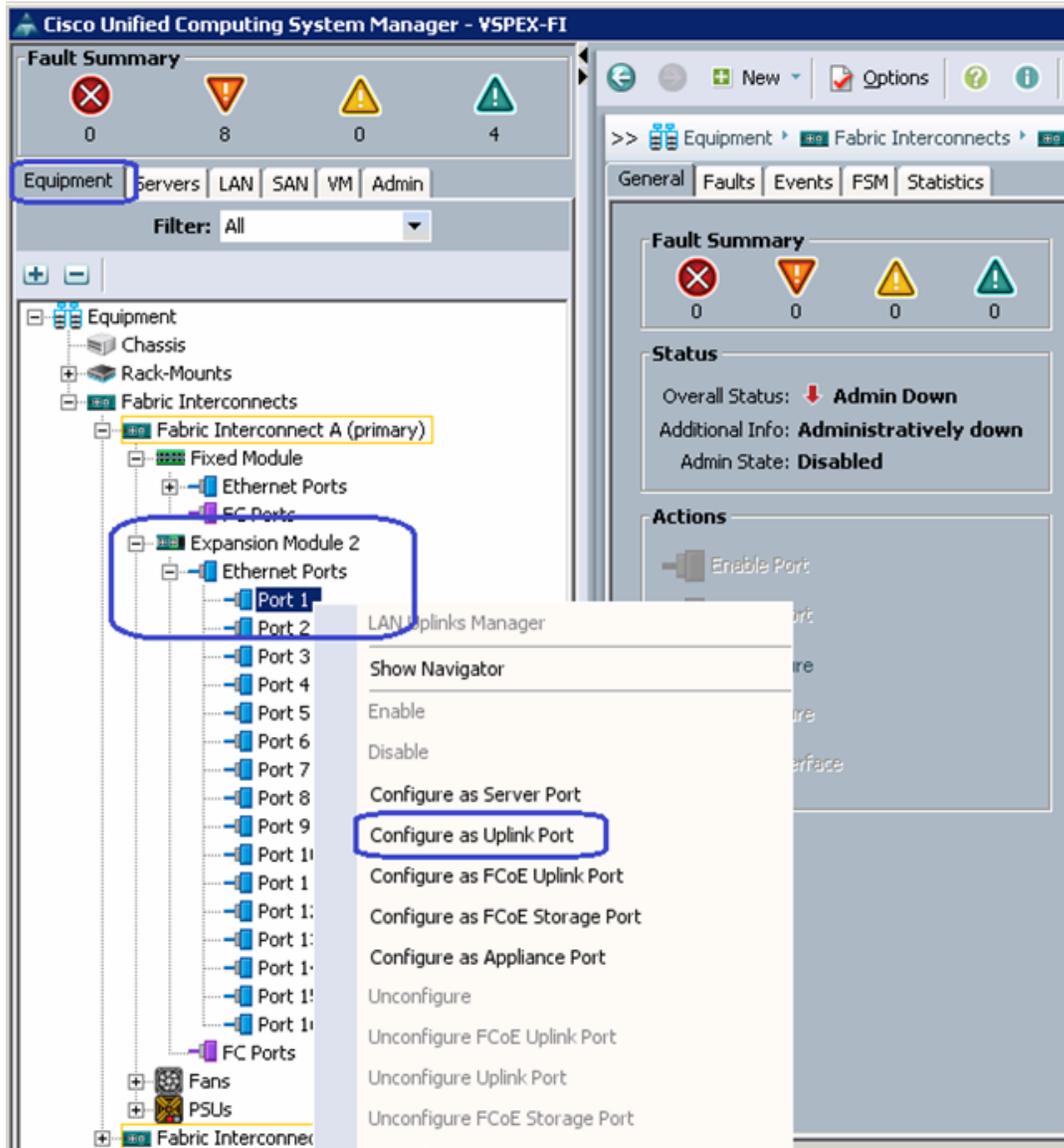
A VLAN can be used to carry FCoE traffic and can be mapped to this VSAN.
Enter the VLAN ID that maps to this VSAN.

FCoE VLAN: **10**

OK Cancel

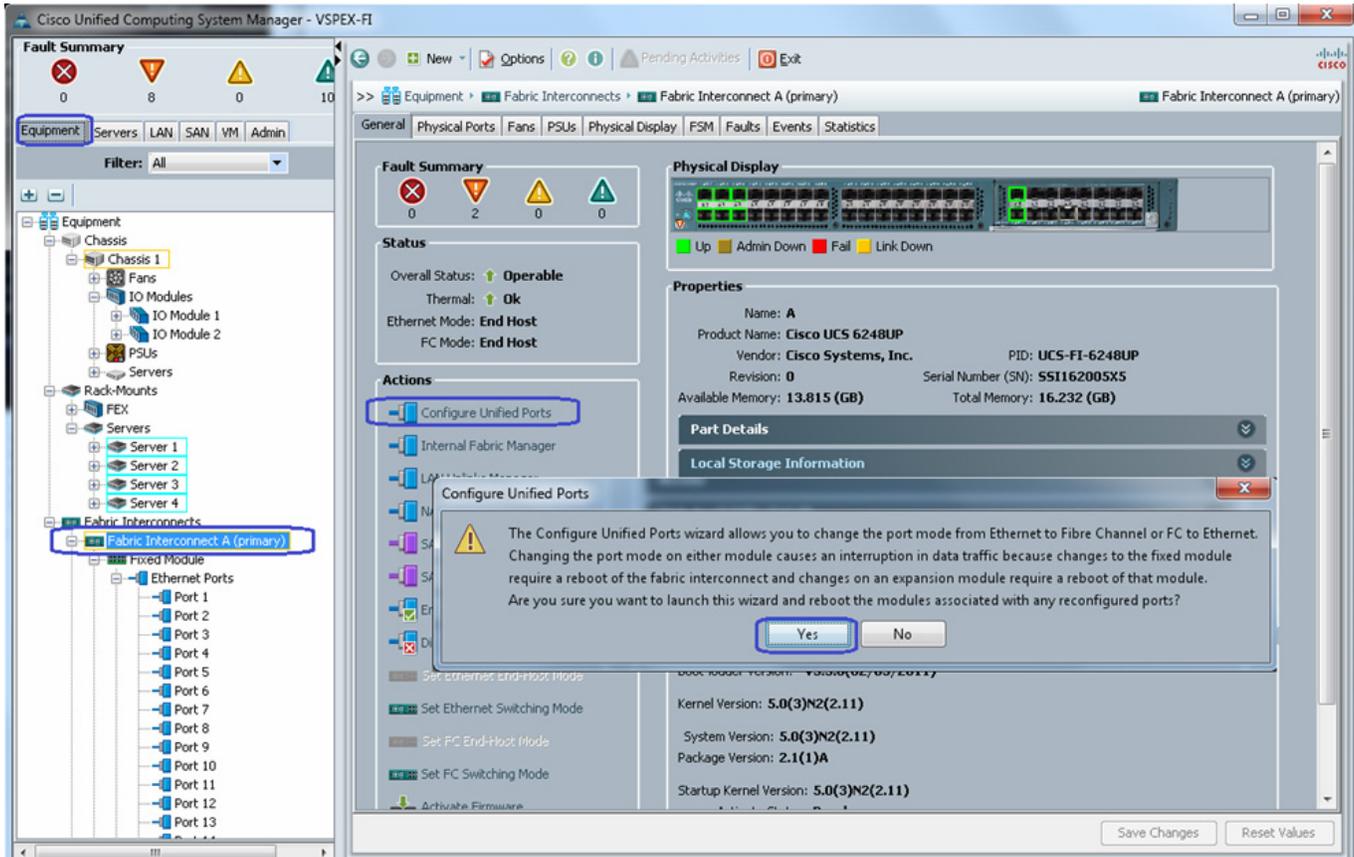
- To configure Uplink ports connected to the infrastructure network, click the **Equipment** tab, expand **Fabric Interconnects**, choose a particular FI, expand **Expansion Module 2** (this may vary depending on which port you have chosen as uplink port), right-click on the Ethernet port, and choose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

Figure 54 Configuring Ethernet Ports as Uplink Ports



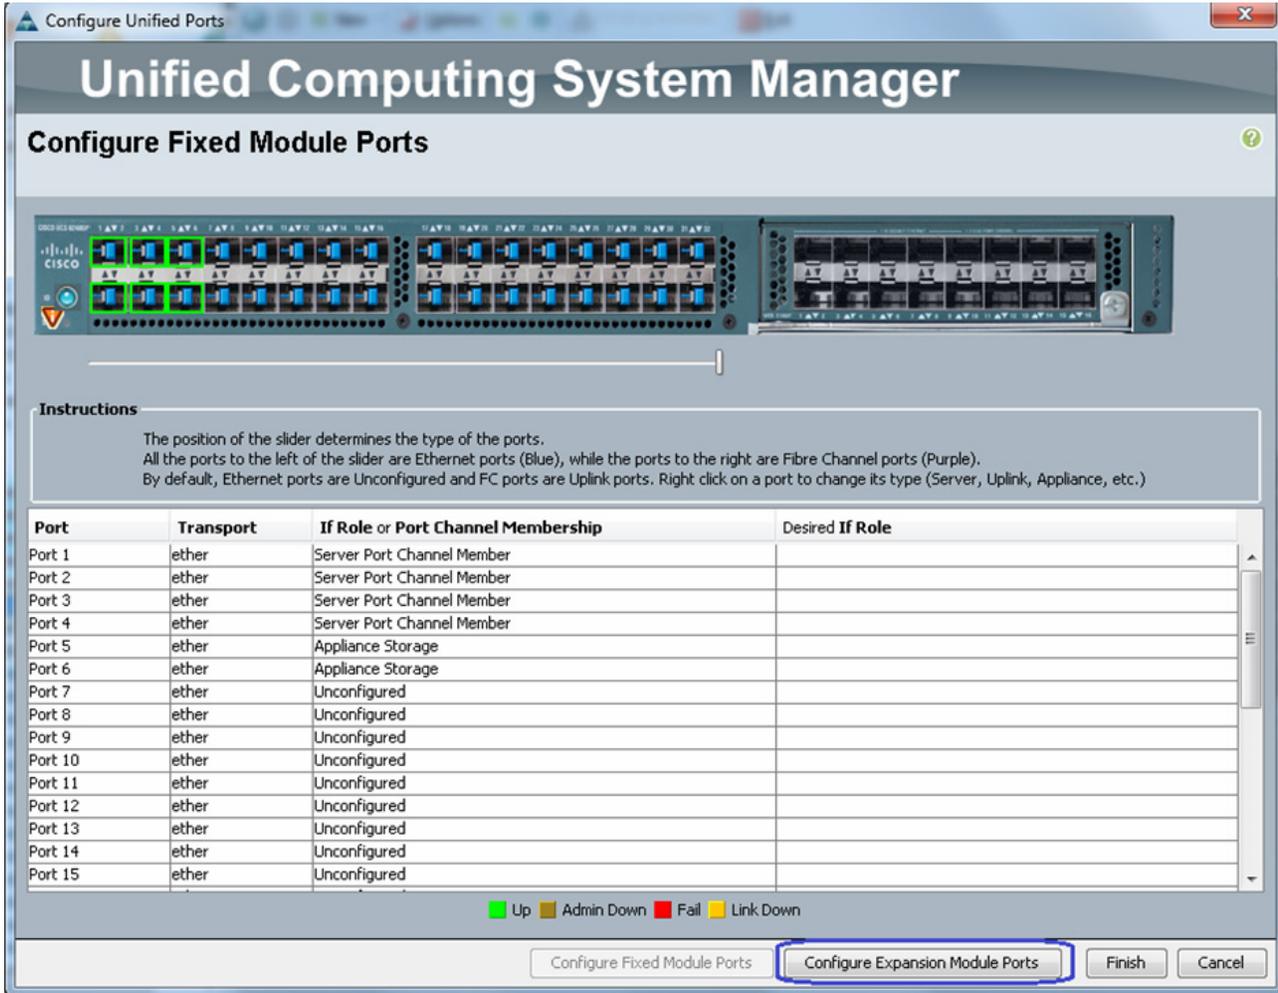
11. Cisco UCS 6248UP Fabric Interconnects have Universal Ports. The physical ports are 10G Ethernet ports by default, but can be converted into Fibre-Channel ports as well. We need the FC connectivity to EMC VNX storage array at least for SAN boot. For that, some of the ports need to be converted to FC ports. We can convert ports from expansion module into FC port. For that, click **Equipment** tab, expand Fabric Interconnects and click **Fabric Interconnect A**. In the right pane, click **Configure Unified Ports**. Click **Yes**, in the warning window.

Figure 55 Configuring Ethernet Ports as FC Ports for Storage Connectivity



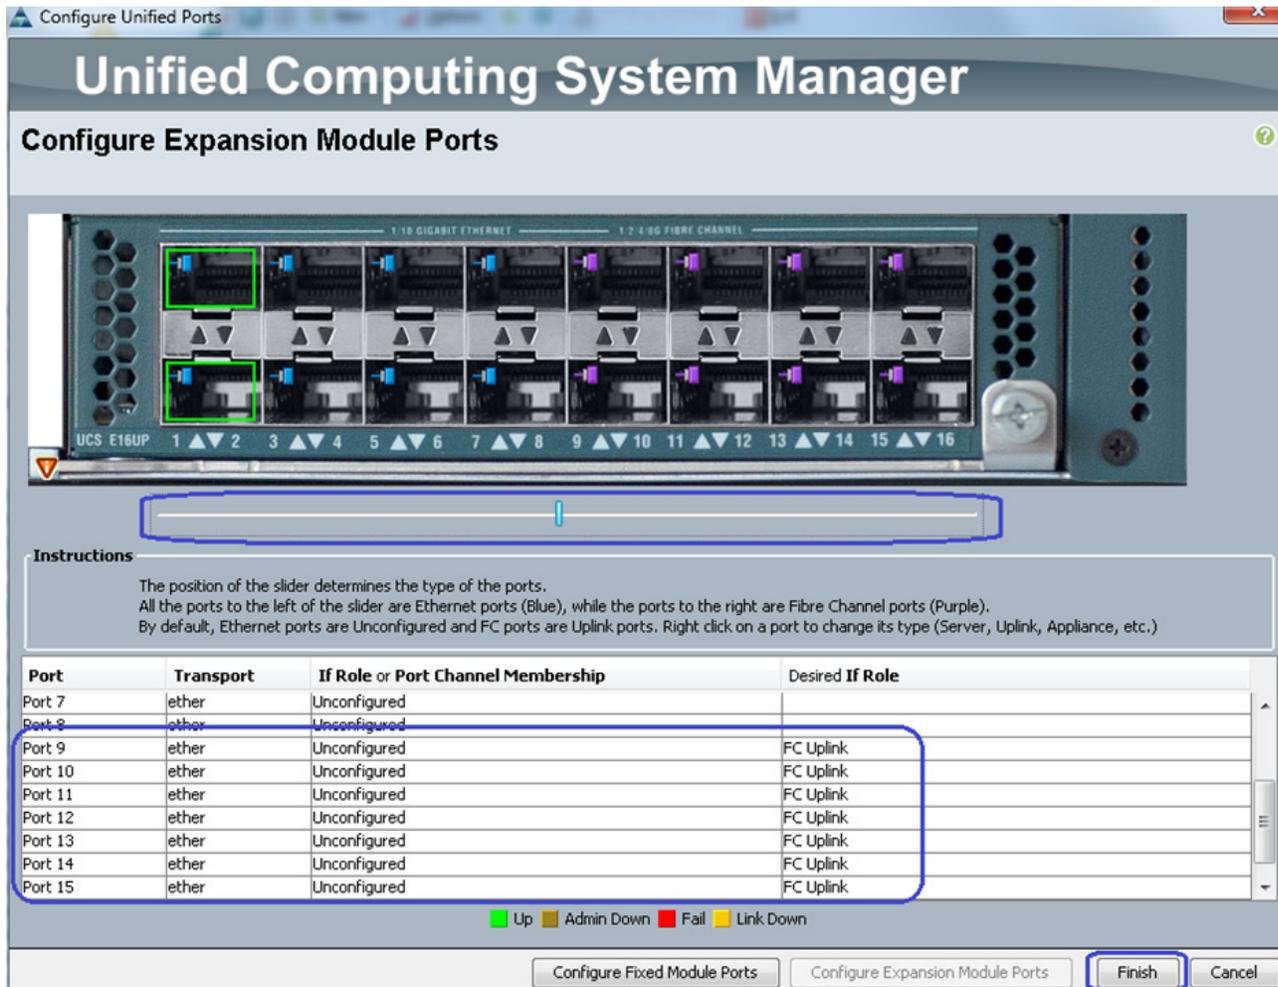
12. In the Configure Unified Ports wizard, click **Configure Expansion Module Ports**.

Figure 56 Configuring Expansion Module Ports



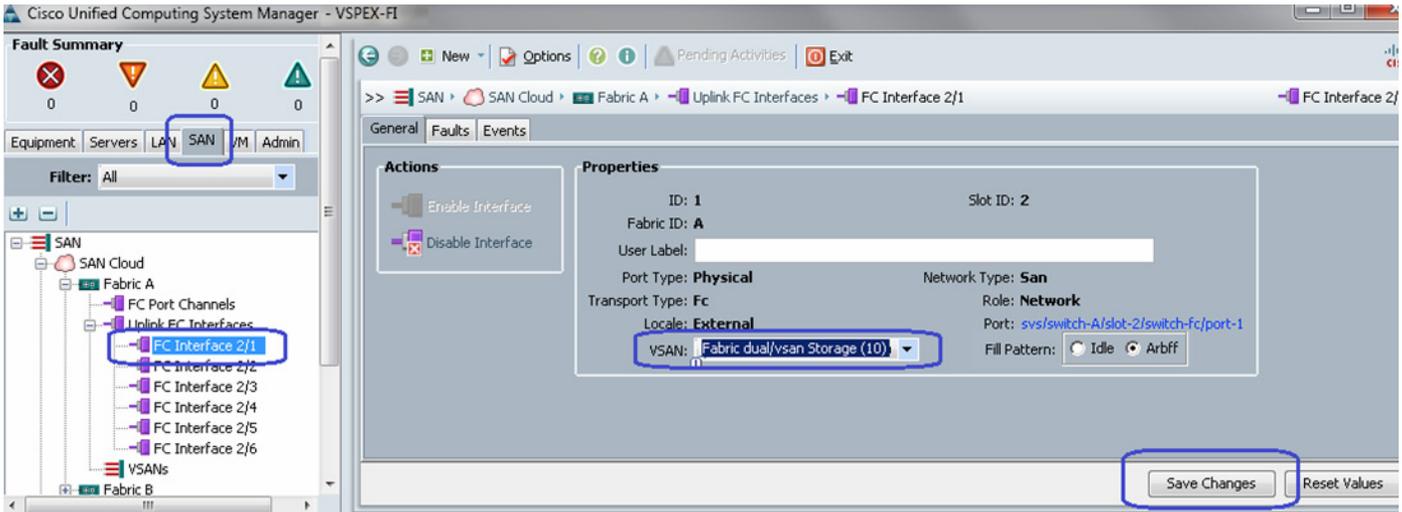
13. Move the slider to the middle of the bar as shown in Figure 57. Make sure that the ports 2/9 to 2/15 are showing as FC Uplink. Click **Finish**. A warning message window pops up to restart FIs. Click **OK**.

Figure 57 Verifying FC Uplink on Ports 9 to 15



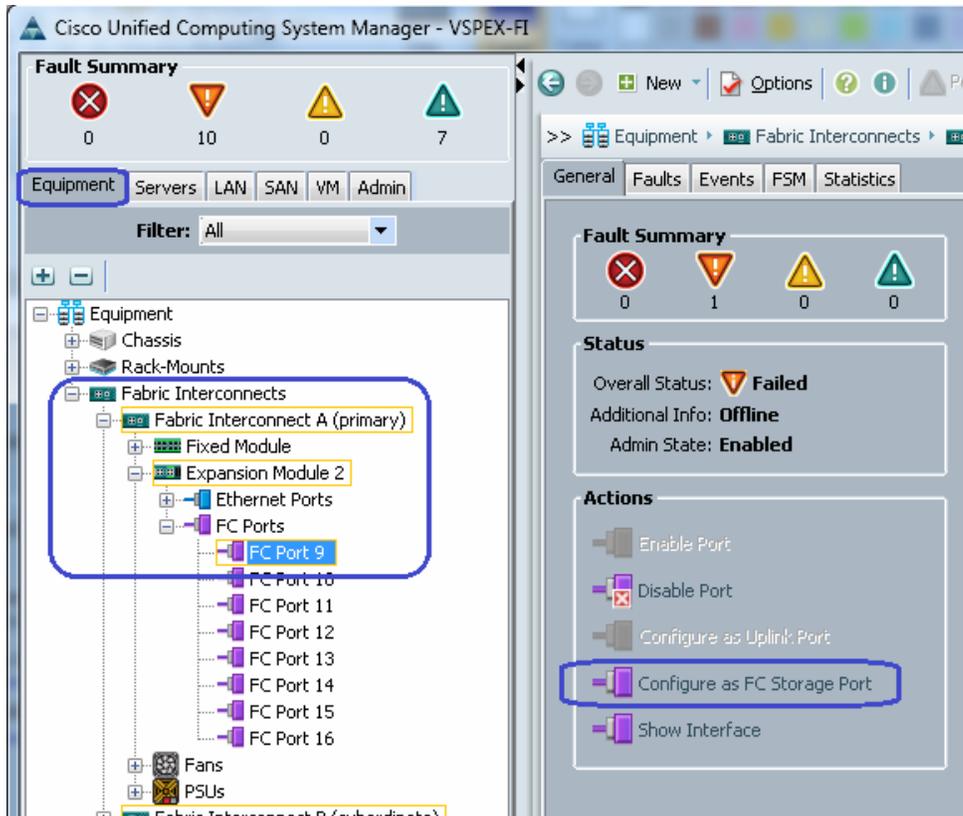
14. Once the FI is rebooted, repeat steps 12, 13 and 14 on FI-B.
15. (NFS-variant only) Click the **SAN** tab, expand **SAN Cloud > Fabric A > Uplink FC Interfaces**, and choose the FC interface. Change the VSAN to the Storage VSAN that was created in steps 6 and 7 and click **Save Changes**.

Figure 58 Selecting Storage VSAN for NFS-Variant



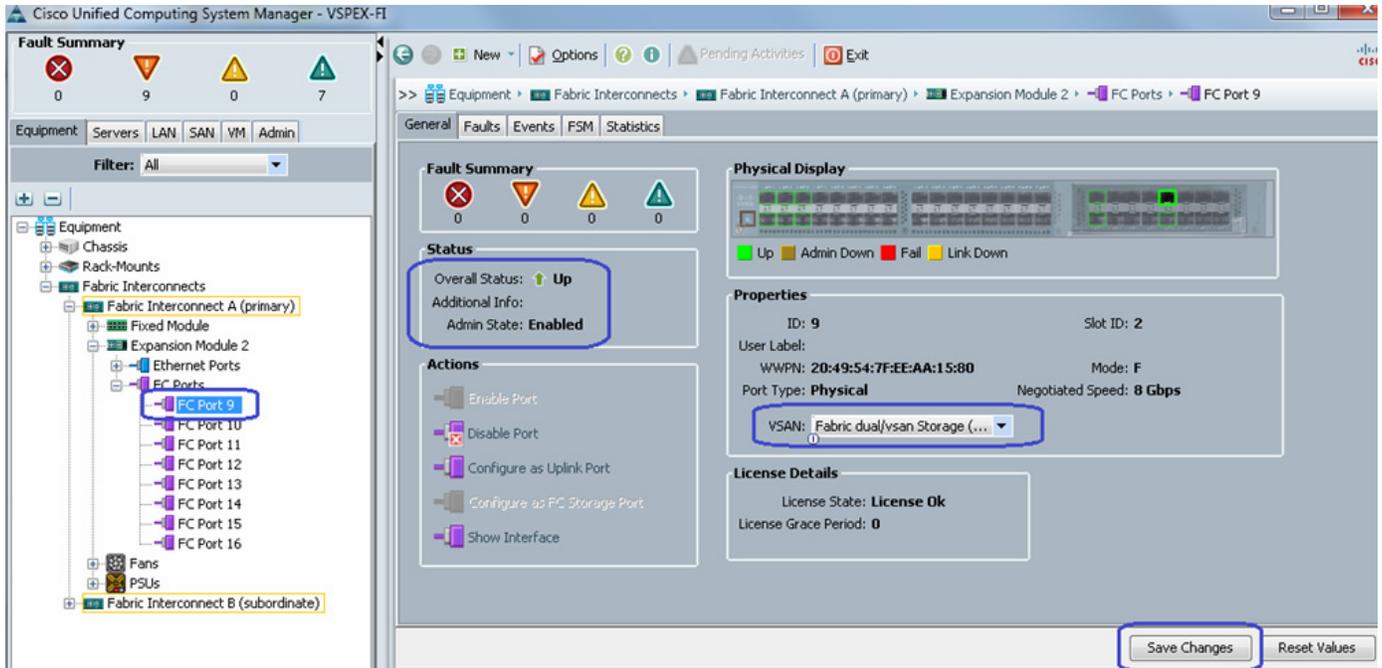
- (FC-variant only) Physical FC ports further need to be classified as FC storage ports for directly attached storage array. Click the **Equipment** tab, expand **Fabric Interconnect > Fabric Interconnect A > Expansion Module 2 > FC Ports** and select each of the FC port and on the right plane, click **Configure as FC Storage port**.

Figure 59 Configuring FC Port as FC Storage Port for FC-Variant



17. (FC-variant only) Make sure that the port are up. From the VSAN drop down list, select the Storage VSAN configured in steps 9 and 10, and click **Save Changes**.

Figure 60 **Selecting Storage VSAN for FC-Variant**



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

Figure 61 Port Name (WWPN) of the Storage Array

```

10.65.121.228 - PuTTY
VSPEX-FI-A# connect nxos a
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
VSPEX-FI-A(nxos)# show flogi database
-----
INTERFACE          VSAN    FCID    PORT NAME          NODE NAME
-----
fc2/9              10      0x2003ef 50:06:01:64:3e:a0:65:0a 50:06:01:60:be:a0:65:0a
fc2/10             10      0x2002ef 50:06:01:65:3e:a0:65:0a 50:06:01:60:be:a0:65:0a

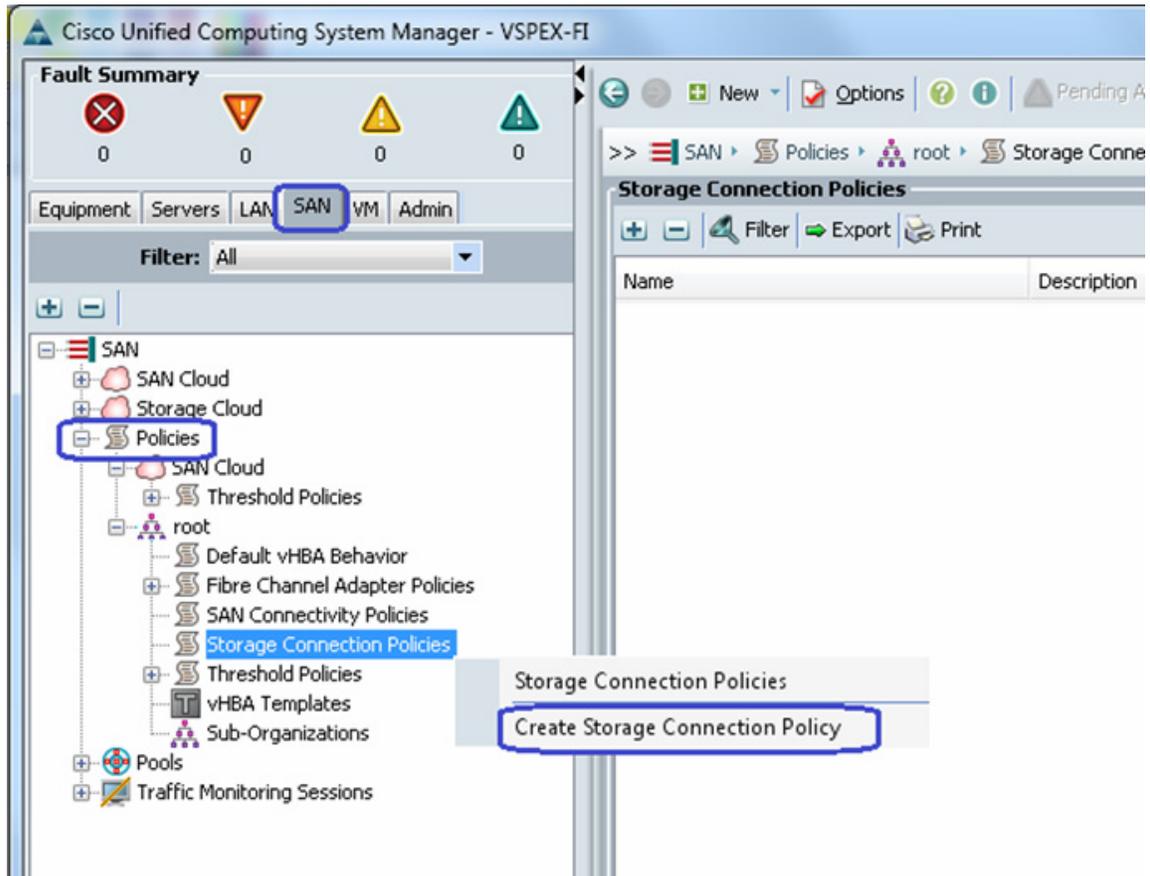
Total number of flogi = 2.

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

```

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

Figure 62 **Creating Storage Connection Policy**



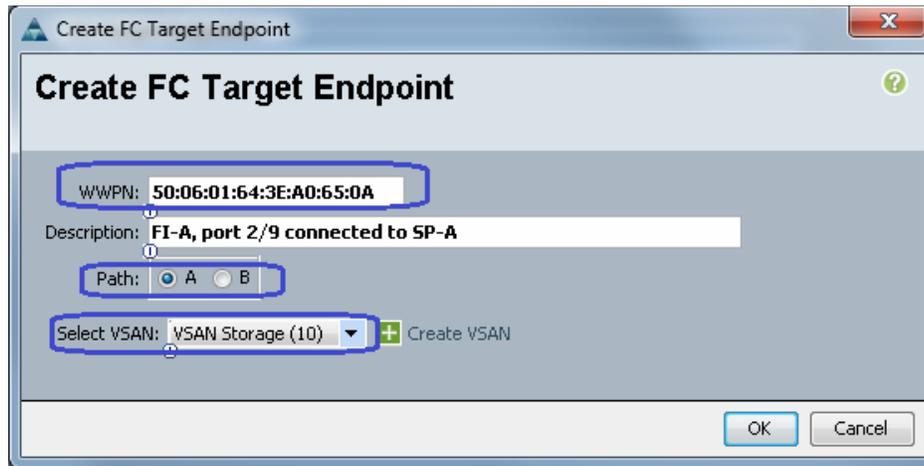
20. (FC-variant only) Enter the name Fabric-A in the name field and optional description. Choose Single Initiator Multiple Targets as the Zoning Type. Click  to add a new FC Target Endpoint.

Figure 63 Details for Creating Storage Connection Policy for FC-Variant



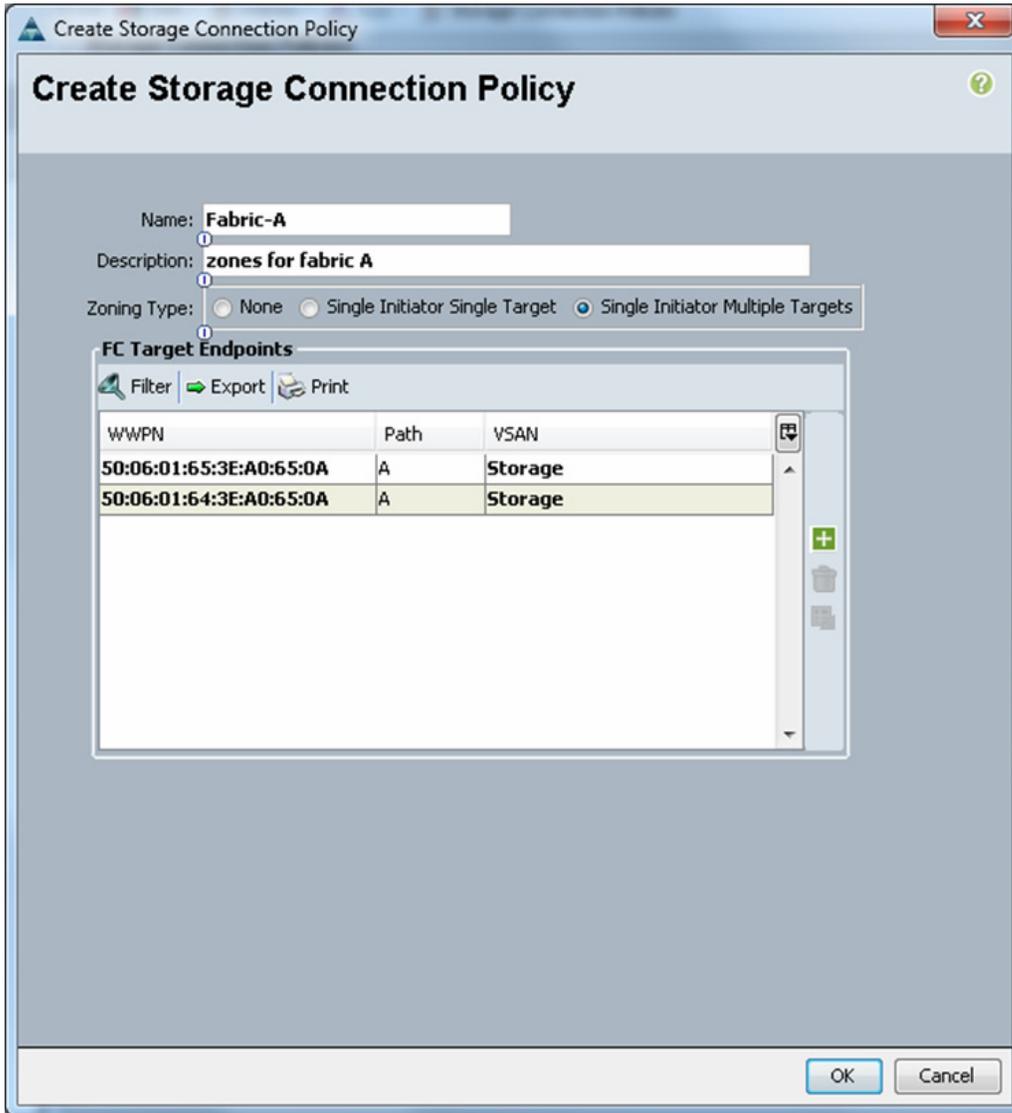
21. (FC-variant only) Copy the WWPN from the **show flogi database** output from step 18 and paste it in the WWPN field. Provide optional description, click the radio button **Path A** and for VSAN choose **Storage VSAN** from the drop-down list.

Figure 64 **Creating FC Target Endpoint for Fabric A**



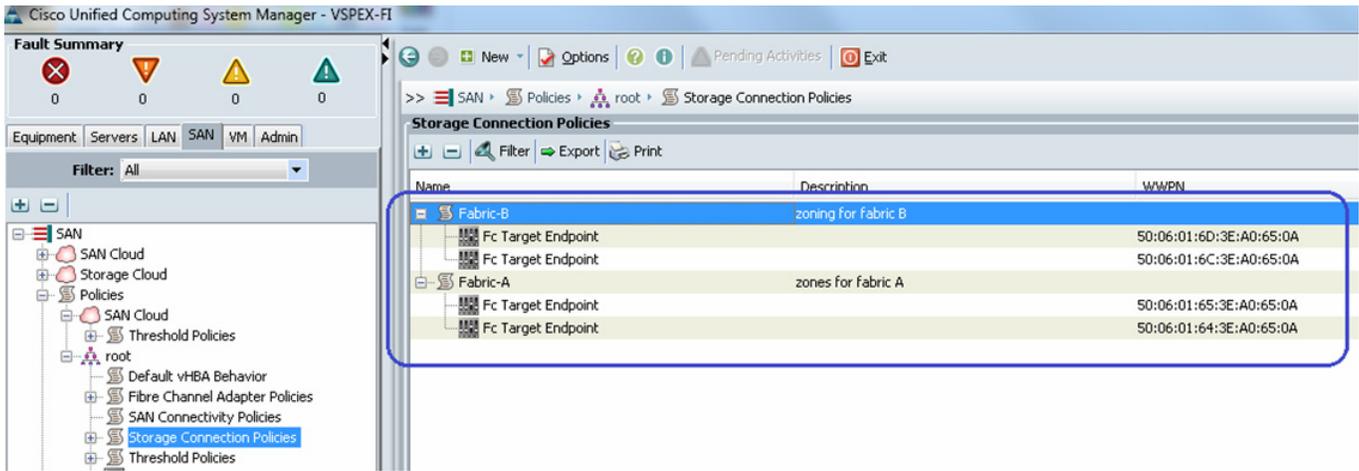
22. (FC-variant only) Similarly, add the second FC target endpoint for fabric A and click **OK**.

Figure 65 Adding Second FC Target Endpoint for Fabric A



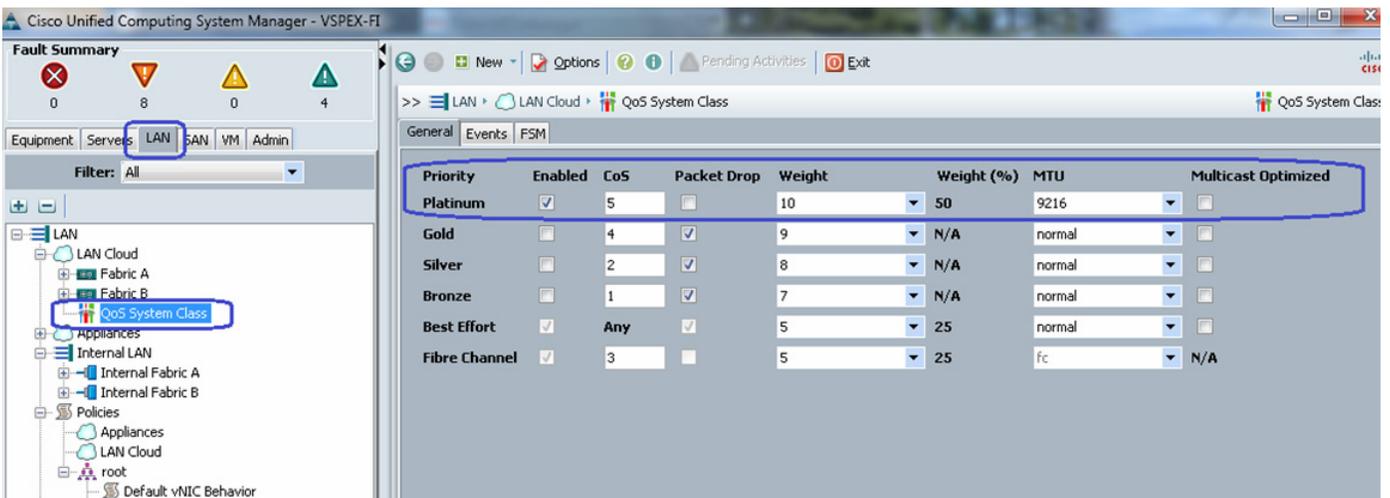
- (FC-variant only) Repeat steps 18 to 22 for Fabric B as well. The end result should look similar to [Figure 66](#).

Figure 66 FC Target Endpoints on Fabric A and Fabric B



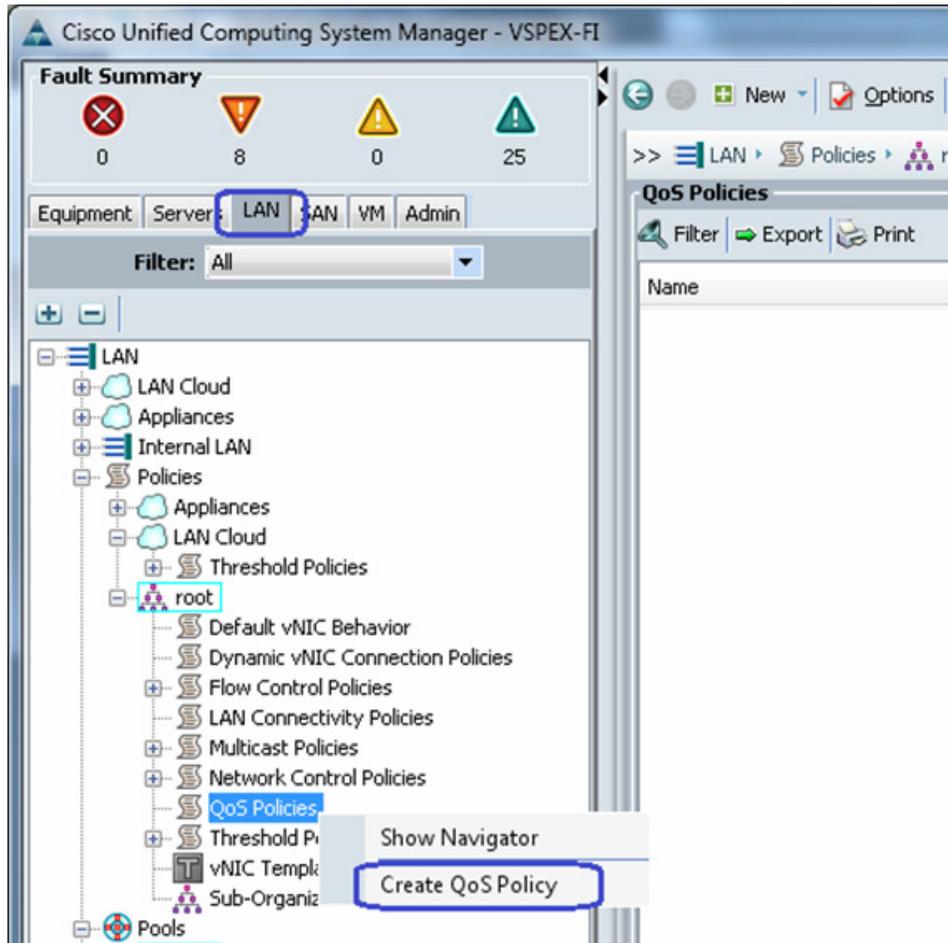
- The next global configuration task is QoS configuration. Click the **LAN** tab, expand **LAN > LAN Cloud**, and choose **QoS System Class**. Check the check box next to **Platinum** for setting the priority, and set MTU to **9216**. Keep other configuration as default and save the configuration.

Figure 67 Configuring QoS



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

Figure 68 Creating QoS Policy



26. Create a QoS policy with name jumboMTU and for the Priority field choose **Platinum** from the drop-down list. Click **OK** to save the configuration.

Configure Identifier Pools

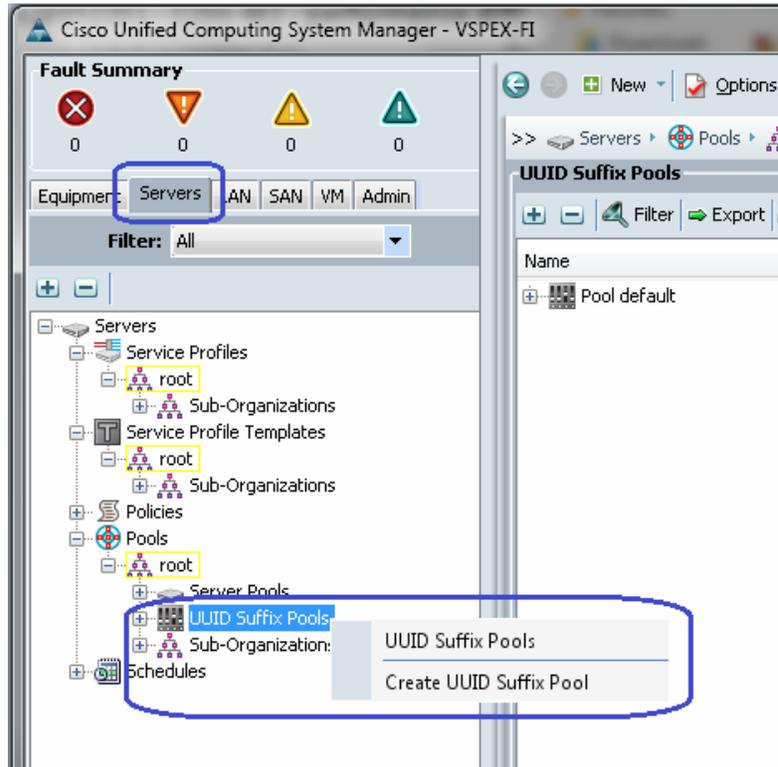
In this section, we would configure following identifier pools used by service profile:

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

To configure pools mentioned above, follow these steps:

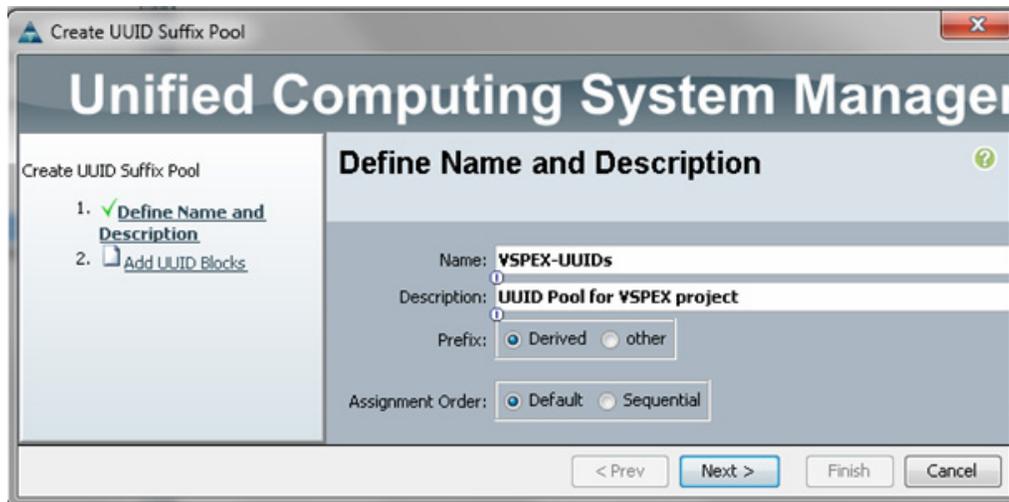
1. From the **Servers** tab, expand **Servers > Pools > root**, and right-click on UUID Suffix pools and click **Create UUID Suffix Pool**.

Figure 69 *Creating UUID Suffix Pool*



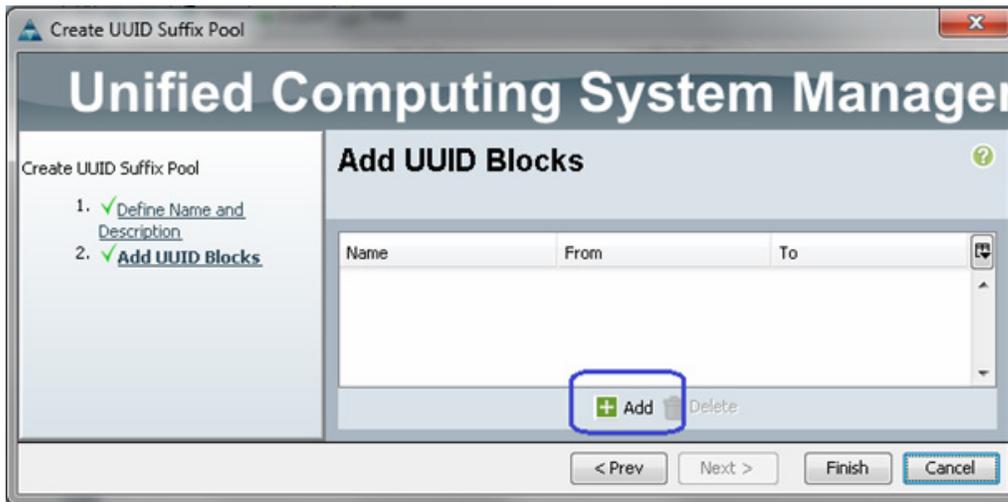
2. Enter the name and description to the UUID suffix pool. Keep other configuration as default.

Figure 70 *Details for Creating UUID Suffix Pool*



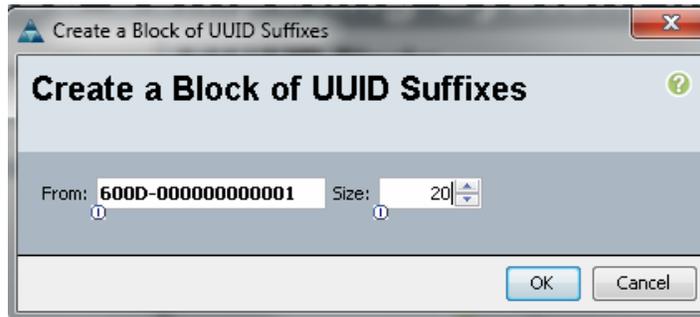
3. Click to  add UUID block.

Figure 71 Adding UUID Block



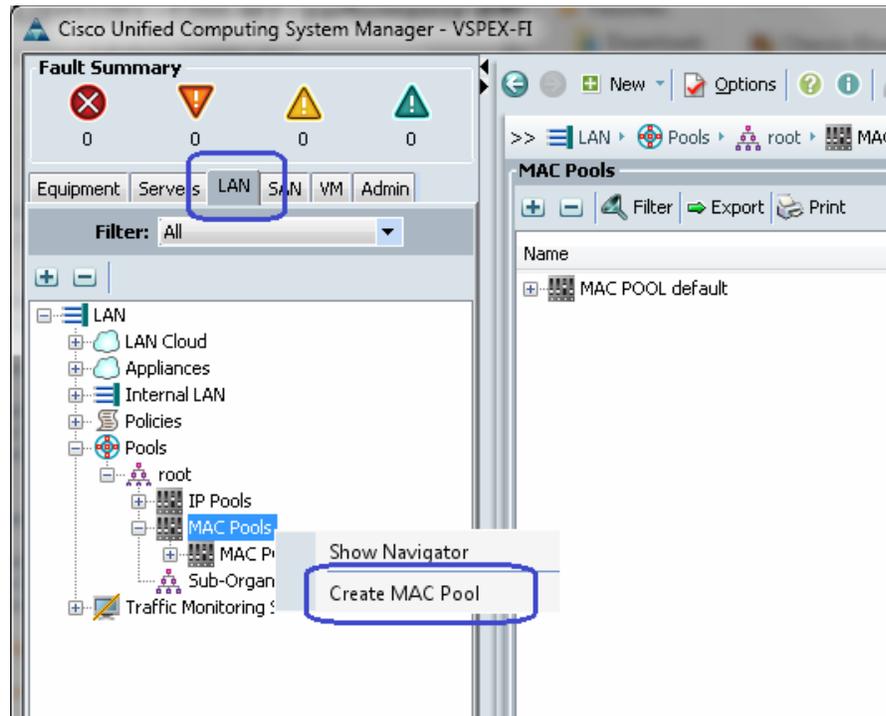
- Specify the beginning of the UUIDs, and have a large size of UUID block to accommodate future expansion.

Figure 72 Specifying Block Size



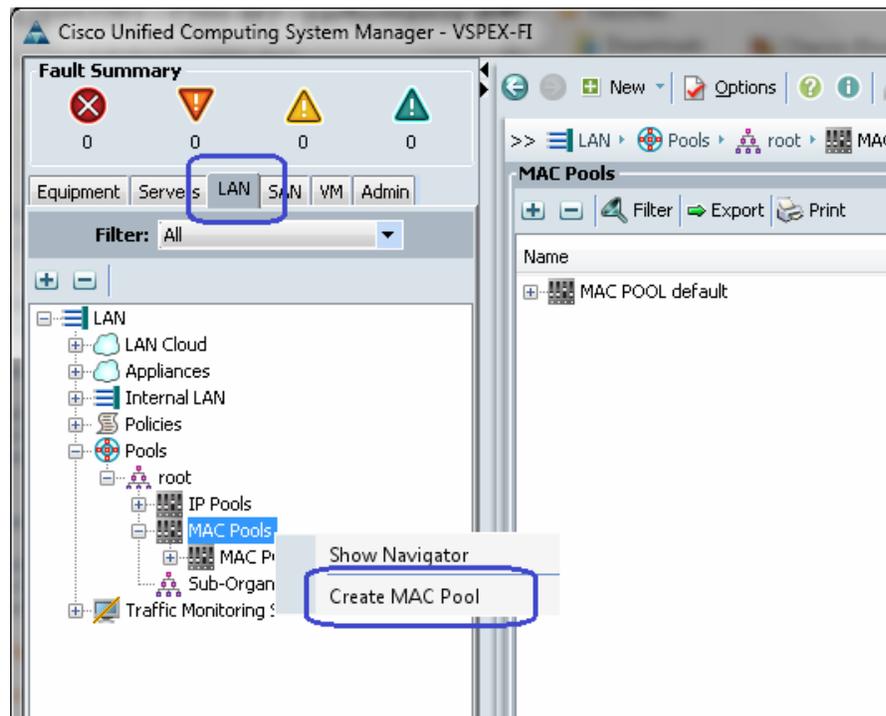
- Click **OK** and then **Finish** to deploy UUID pool.
- Click the **LAN** tab, expand **LAN > Pools > root**, right-click on MAC Pools and select **Create MAC Pool**.

Figure 73 **Creating MAC Pool**



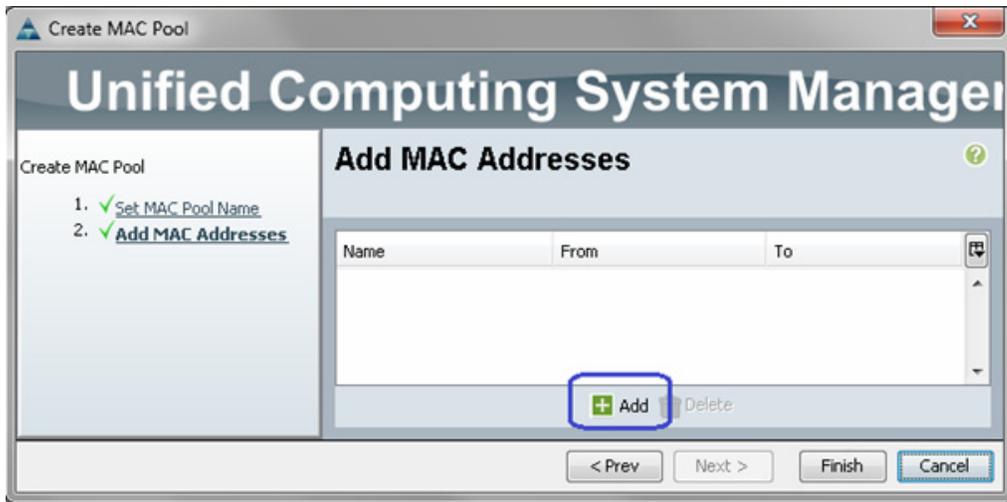
7. Enter the name and description for MAC pool and click **Next**.

Figure 74 **Details for Creating MAC Pool**



- Click **+ Add** to add MAC pool block.

Figure 75 Adding MAC Address



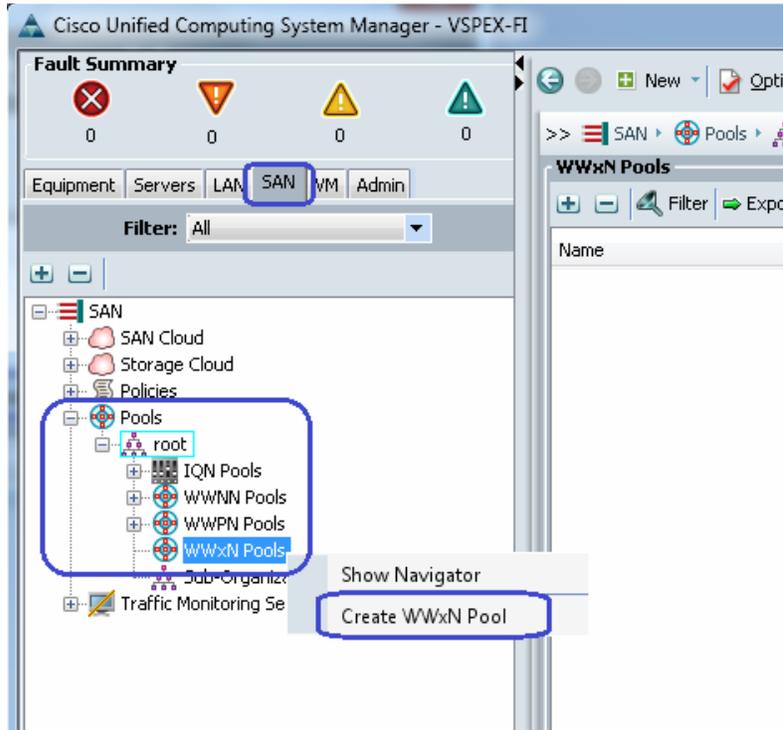
- Enter the initial MAC address and size of the block. As always, provide large number of MAC addresses to accommodate future expansion. We will require 6 MAC addresses per server.

Figure 76 Specifying MAC Address Block



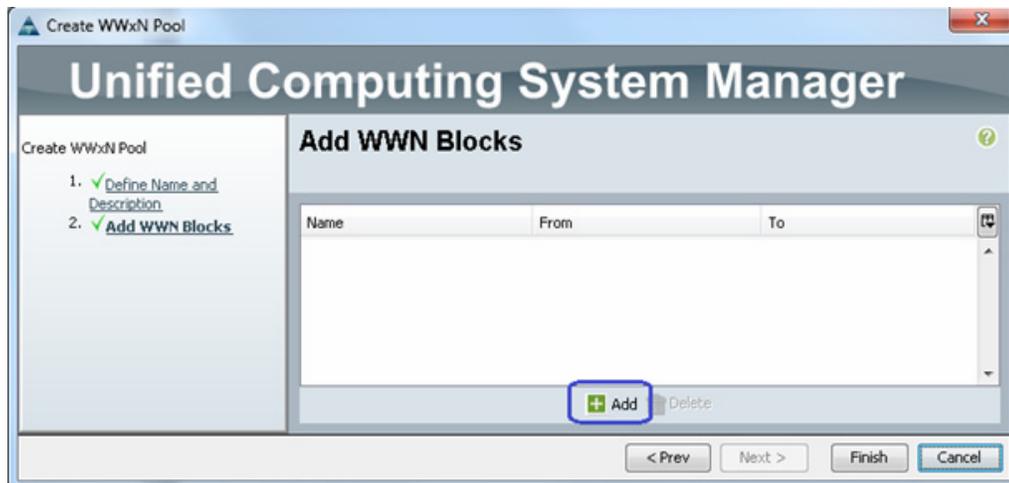
- Click **OK** and **Finish** to complete configuration.
- From the **SAN** tab, expand **SAN > Pools > root**, right-click on WWxN Pools, and choose **Create WWxN Pool**.

Figure 77 *Creating WWxN Pool*



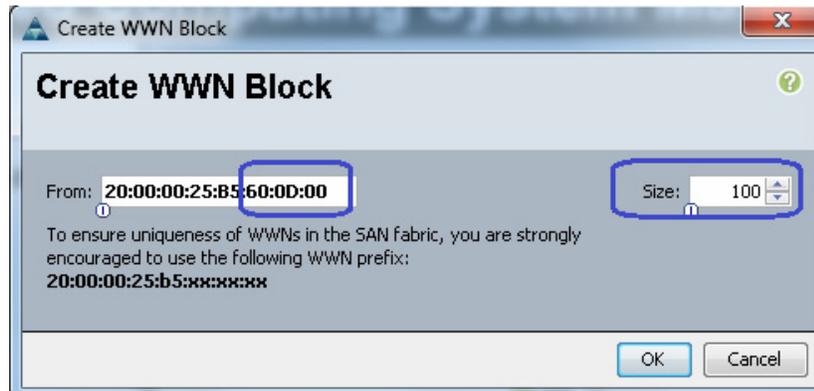
12. Enter name, and description for WWxN and choose **3 Ports per Node** from the drop-down list for max ports.
13. Click  to add a block of WWxN IDs.

Figure 78 *Adding WWxN Block*



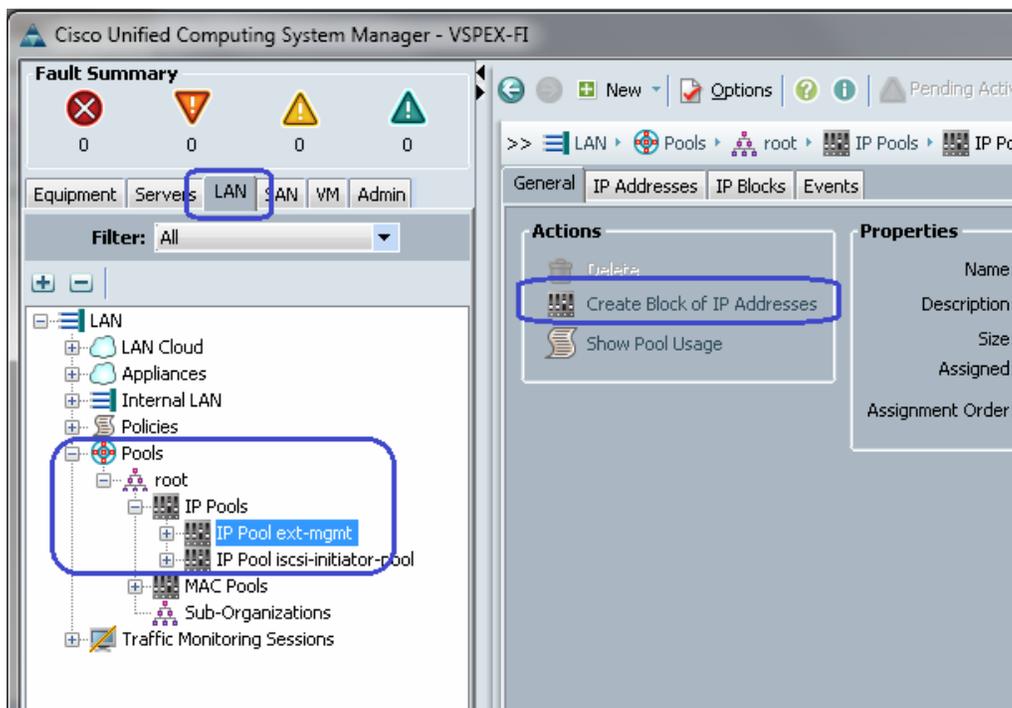
14. Enter first WWN IDs and sufficiently large number of block size. Click **OK** and **Finish**.

Figure 79 Specifying WWxN Block Size



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

Figure 80 Creating IP Address Block



- Enter the initial IP address, size of the pool, default gateway and subnet mask. Click **OK** to deploy the configuration. IP addresses will be assigned to various Rack-Mount server CIMC management access from this block.

Figure 81 Specifying the IP address Block Size

The screenshot shows a dialog box titled "Create Block of IP Addresses". The dialog contains the following fields and values:

- From:** 10.65.121.231
- Size:** 8
- Subnet Mask:** 255.255.255.0
- Default Gateway:** 10.65.121.1
- Primary DNS:** 0.0.0.0
- Secondary DNS:** 0.0.0.0

Buttons for "OK" and "Cancel" are located at the bottom right of the dialog.

Configure Server Pool and Qualifying Policy

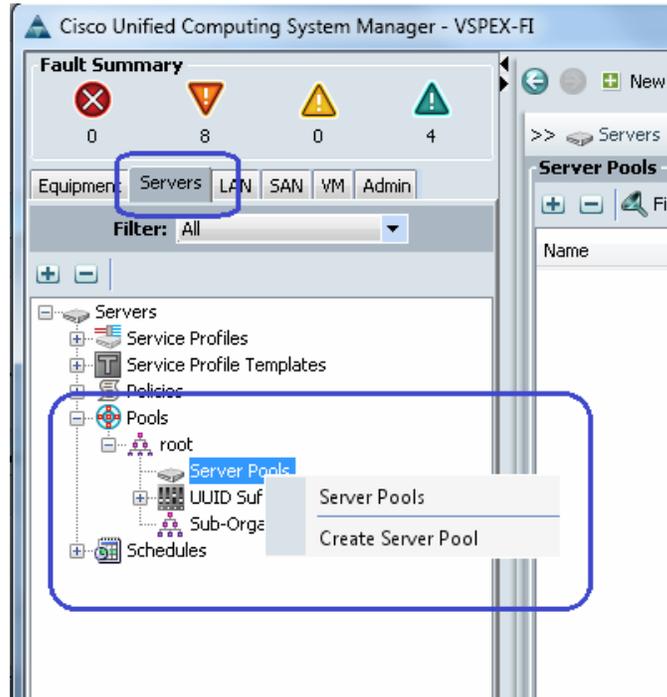
Creation and policy based auto-population of server pool can be sub-divided into the following tasks:

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

Follow these steps to complete the three tasks mentioned above:

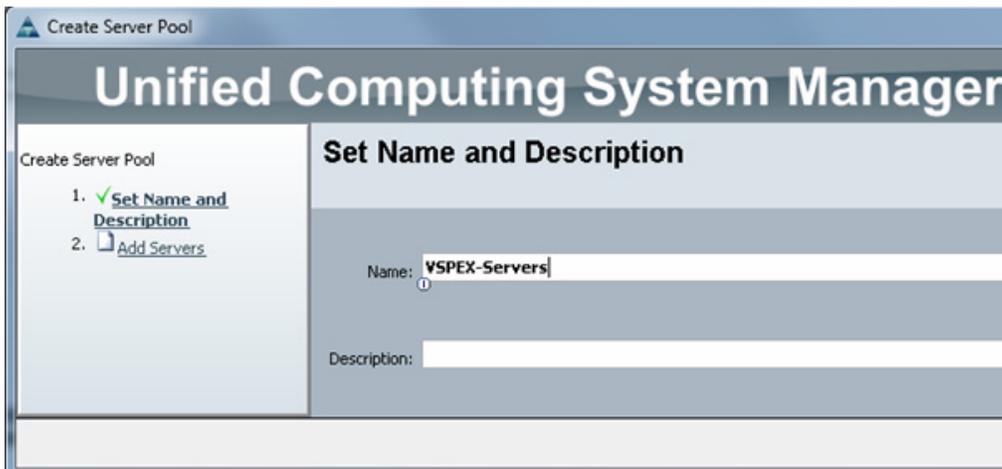
1. From the **Servers** tab, expand **Servers > Pools > root**, right-click on Server Pools and choose **Create Server Pool**.

Figure 82 Creating Server Pools



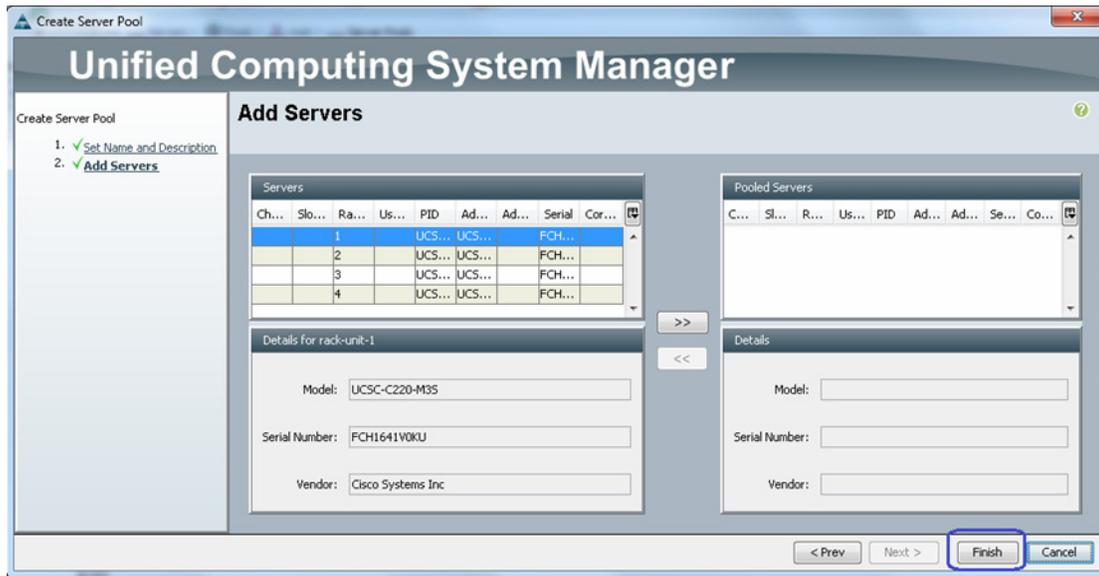
2. Enter the name of the server pool in the Name field, and click **Next**.

Figure 83 Entering Details in the Create Server Pool Wizard



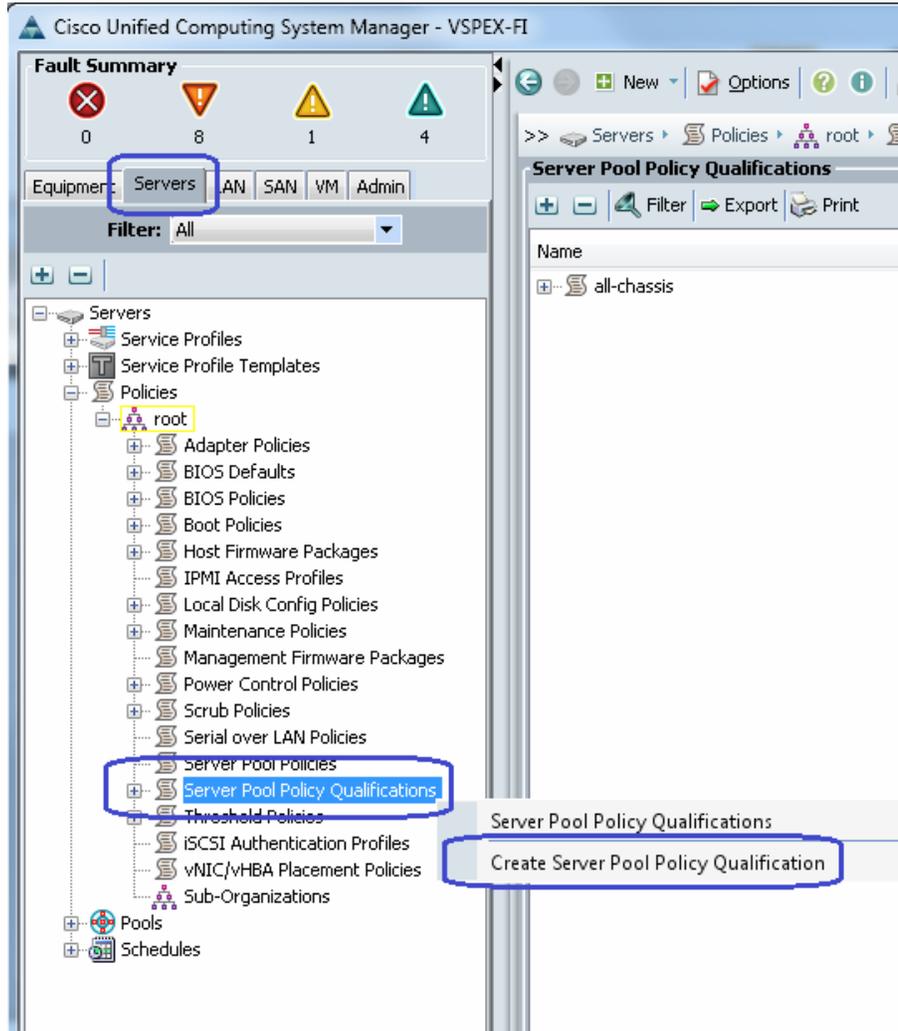
3. Click **Finish** to create the empty server pool. We would add the compute resources to this pool dynamically, based on policy.

Figure 84 Adding Servers in the Create Server Pool Wizard



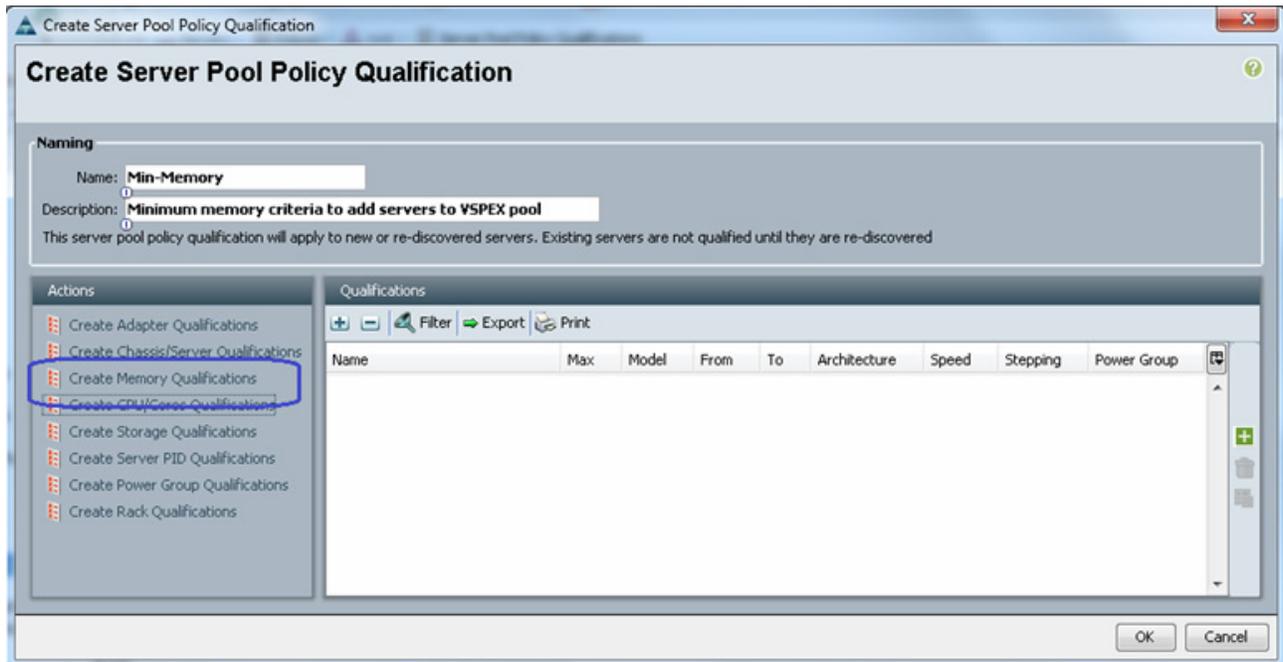
4. From the **Servers** tab, expand **Servers > Policies > root**, right-click on **Server Pool Policy Qualifications** and choose **Create Server Pool Policy Qualification**.

Figure 85 Creating Server Pool Policy Qualification



5. Enter a name for the server policy qualification criterion in the Name field. In the left pane under Actions choose **Create Memory Qualifications**.

Figure 86 **Creating Memory Qualification**

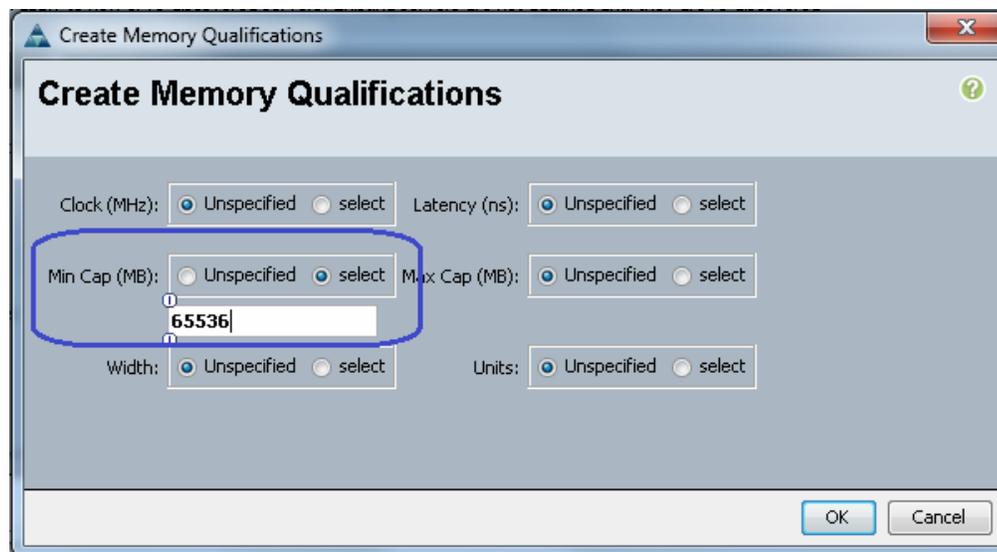


- Set the minimum RAM capacity as 64GB RAM for the pool qualification criterion. Click **OK** twice to create the qualification.



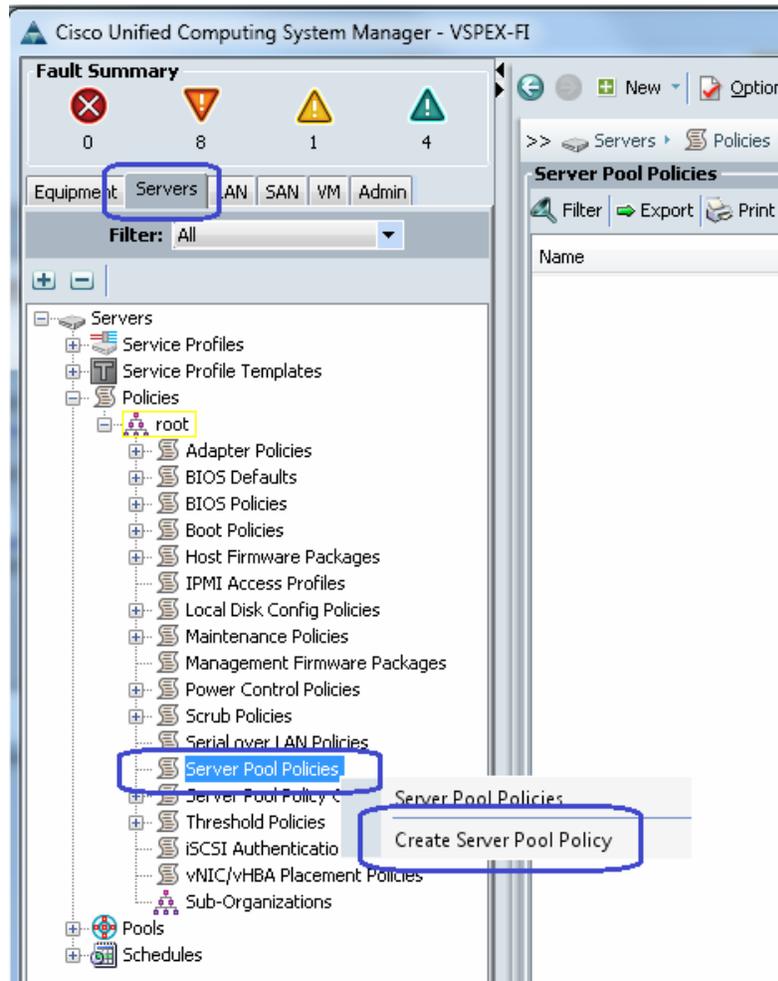
Note This is just an example criterion, you can choose a criterion that suites your requirement.

Figure 87 **Specifying Minimum RAM Capacity for Memory Qualification**



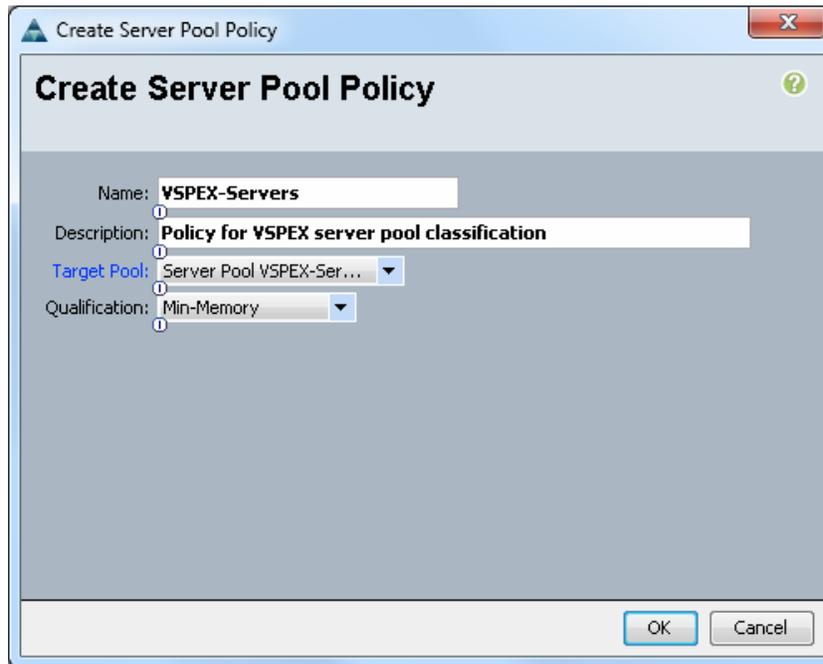
- From the **Servers** tab, expand **Servers > Policies > root**, right-click on Server Pool Policies and choose **Create Server Pool Policy**.

Figure 88 Creating Server Pool Policy



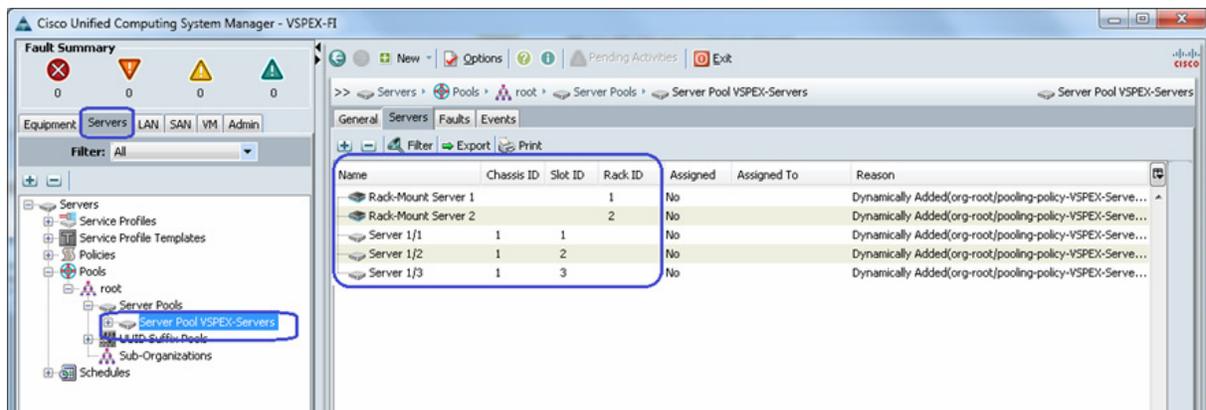
8. Enter a name and description to the server pool policy. Choose recently created Target Pool and Qualification. Click **OK** to deploy the configuration.

Figure 89 Details for Creating Server Pool Policy



- If you go back to the server pool created in step 1 above and click the **Servers** tab on right pane, you will see that all the compute resources that meet the qualification criteria are dynamically added to the server pool. [Figure 90](#) shows the screen capture taken from the FC-variant of the architecture, where combination of Cisco UCS B200M3 Blade Servers and Cisco UCS C220M3 Rack Servers are used to share the workload. This architecture showcases the form-factor independent architecture with unified managed using UCS Manager.

Figure 90 Qualified Compute Resources Automatically Added to the Server Pool



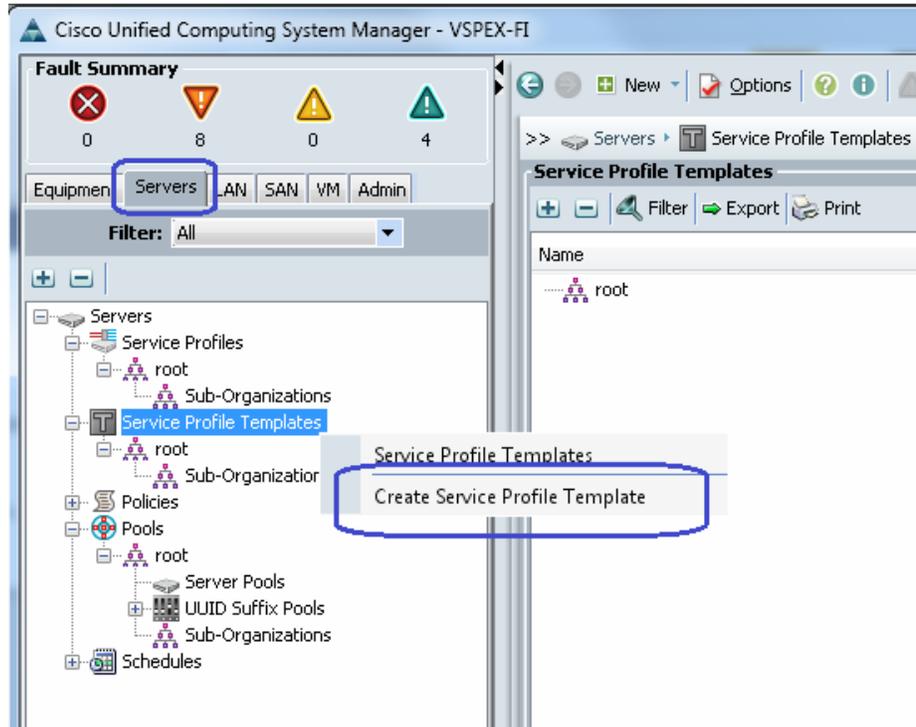
Configure Service Profile Template

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

To create service profile template, follow these steps:

1. From the **Servers** tab. Expand **Servers > Service Profile Templates**, right-click on service profile templates and choose **Create Service Profile Template**.

Figure 91 *Creating Service Profile Template*



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

Figure 92 **Creating Service Profile Template - Entering Details**

Create Service Profile Template

Unified Computing System Manager

Identify Service Profile Template

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

Name:

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

Where: **org-root**

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

Type: Initial Template Updating Template

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

UUID

UUID Assignment:

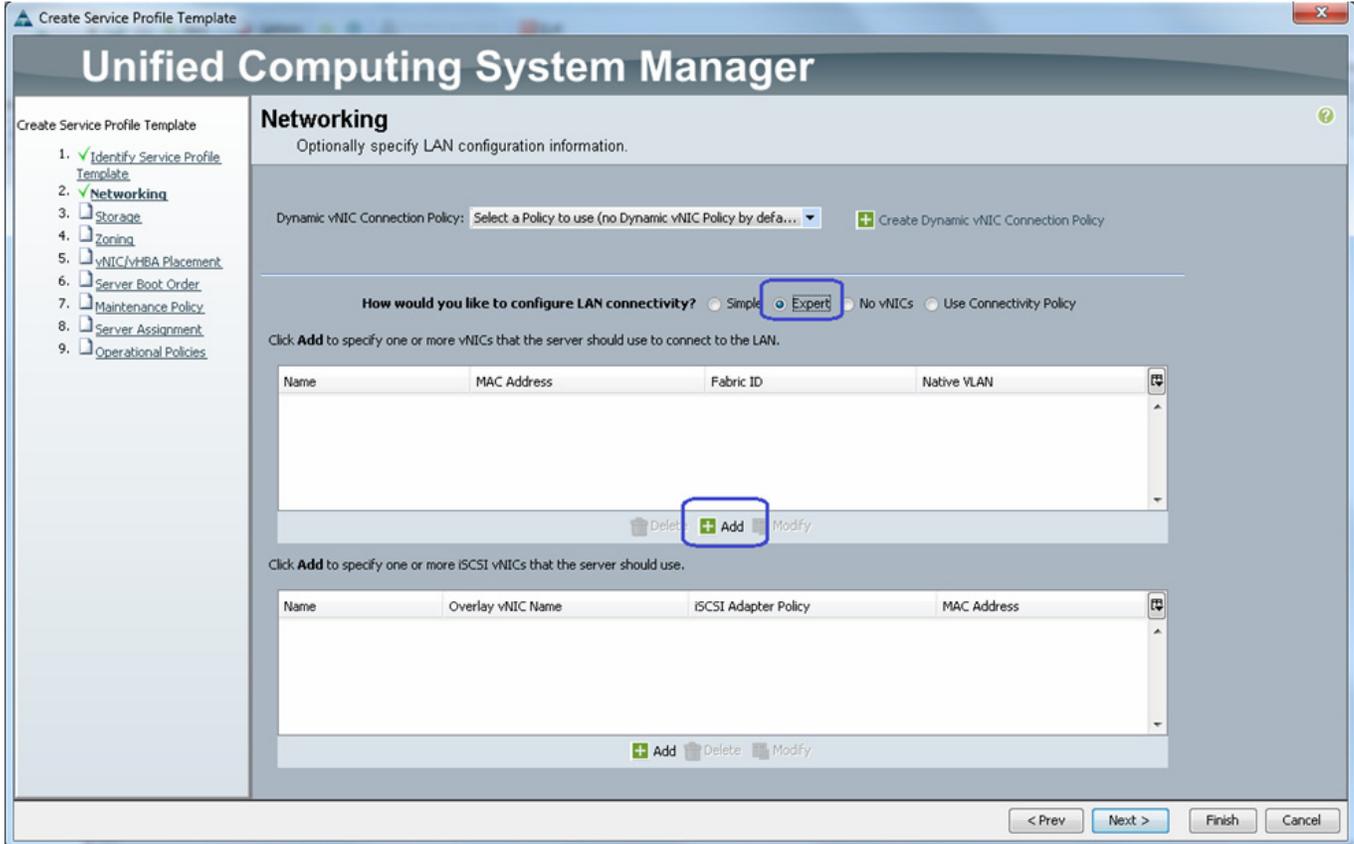
The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

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

< Prev Next > Finish Cancel

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

Figure 93 Creating Service Profile Template - LAN Configuration Details



4. Create a system vNIC for fabric A. Enter system-A as the vNIC name, choose the MAC pool created in section D, click the radio button **fabric A** for fabric ID, check the check boxes **vMotion** and **vSphereMgmt** VLANs with vSphereMgmt as native VLAN. For MTU enter **9000**, for Adapter Policy field, choose **VMware** and choose **jumboMTU** for the QoS Policy field.

Figure 94 Creating a System vNIC

Create vNIC

Name:

Use vNIC Template:

Create vNIC Template

MAC Address

MAC Address Assignment:

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

Fabric ID: Fabric A Fabric B Enable Failover

VLANs

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

MTU:

Warning

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

Pin Group:

Operational Parameters

Adapter Performance Profile

Adapter Policy:

Dynamic vNIC Connection Policy:

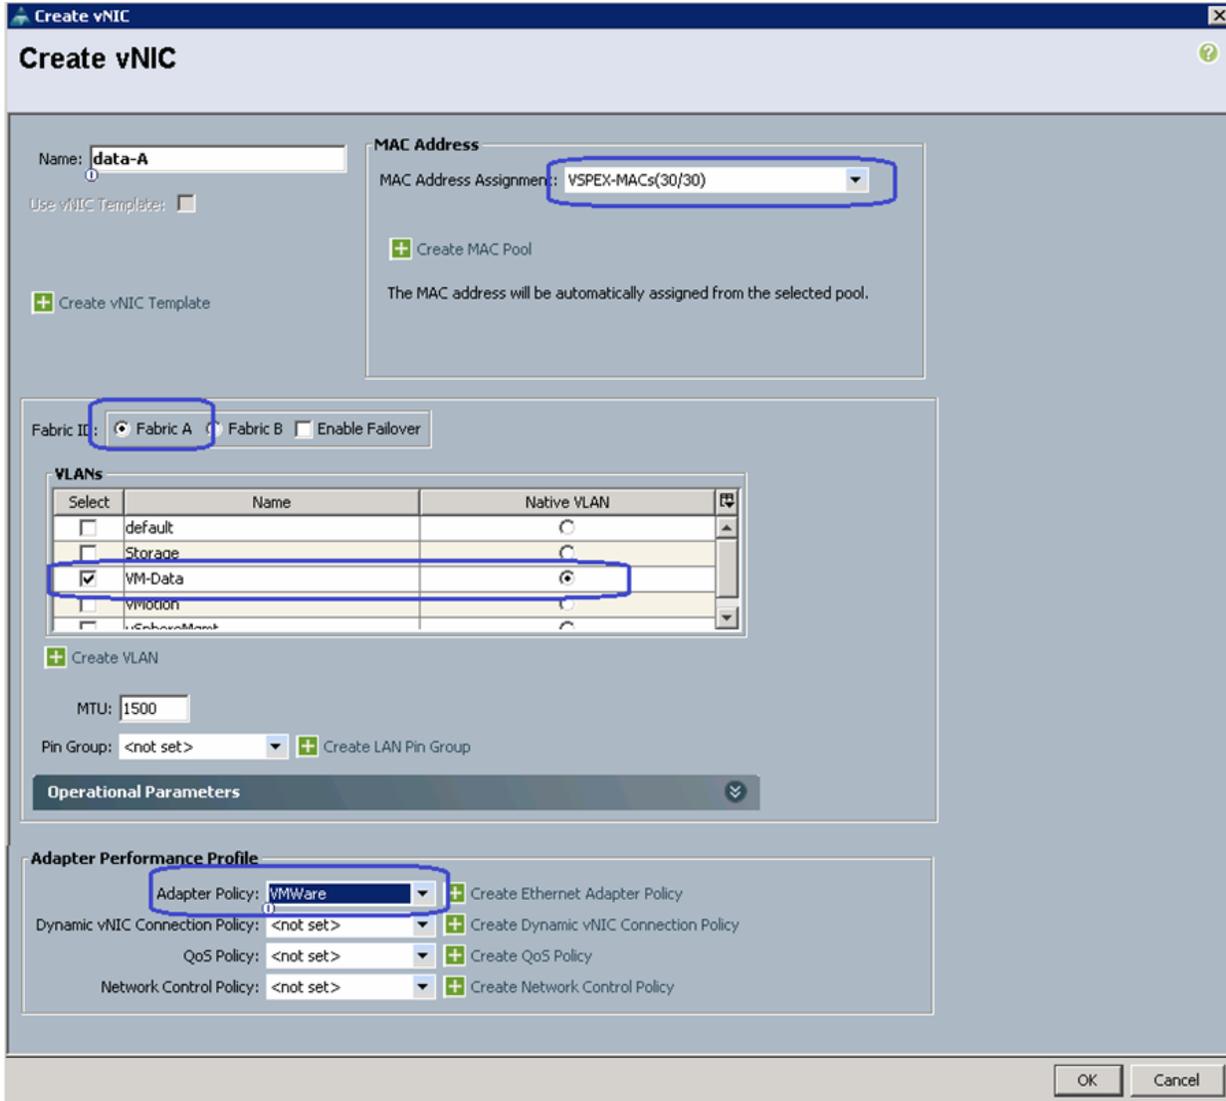
QoS Policy:

Network Control Policy:

- Similarly, create one more vNIC with exact same properties on fabric B.
- (NFS-variant only) Create two more vNICs similar to steps 3, 4 and 5 for NFS server access. Enter the names Storage-A and Storage-B for vNICs on fabric A and B respectively, choose only Storage VLAN and mark it as native VLAN and choose **VMware** and **jumboMTU** for adapter policy and QoS policy respectively.

- Finally, create a vNIC for VM data traffic. Enter **data-A** for vNIC name, same MAC address pool name, **Fabric A** for Fabric ID, **VM-Data** as native VLAN, and **VMware** as adapter policy.

Figure 95 Creating vNIC for VM Data Traffic



- Similarly, create one more vNIC for Fabric B for VM data traffic. [Table 9](#) summarizes all the vNICs created on the service profile.

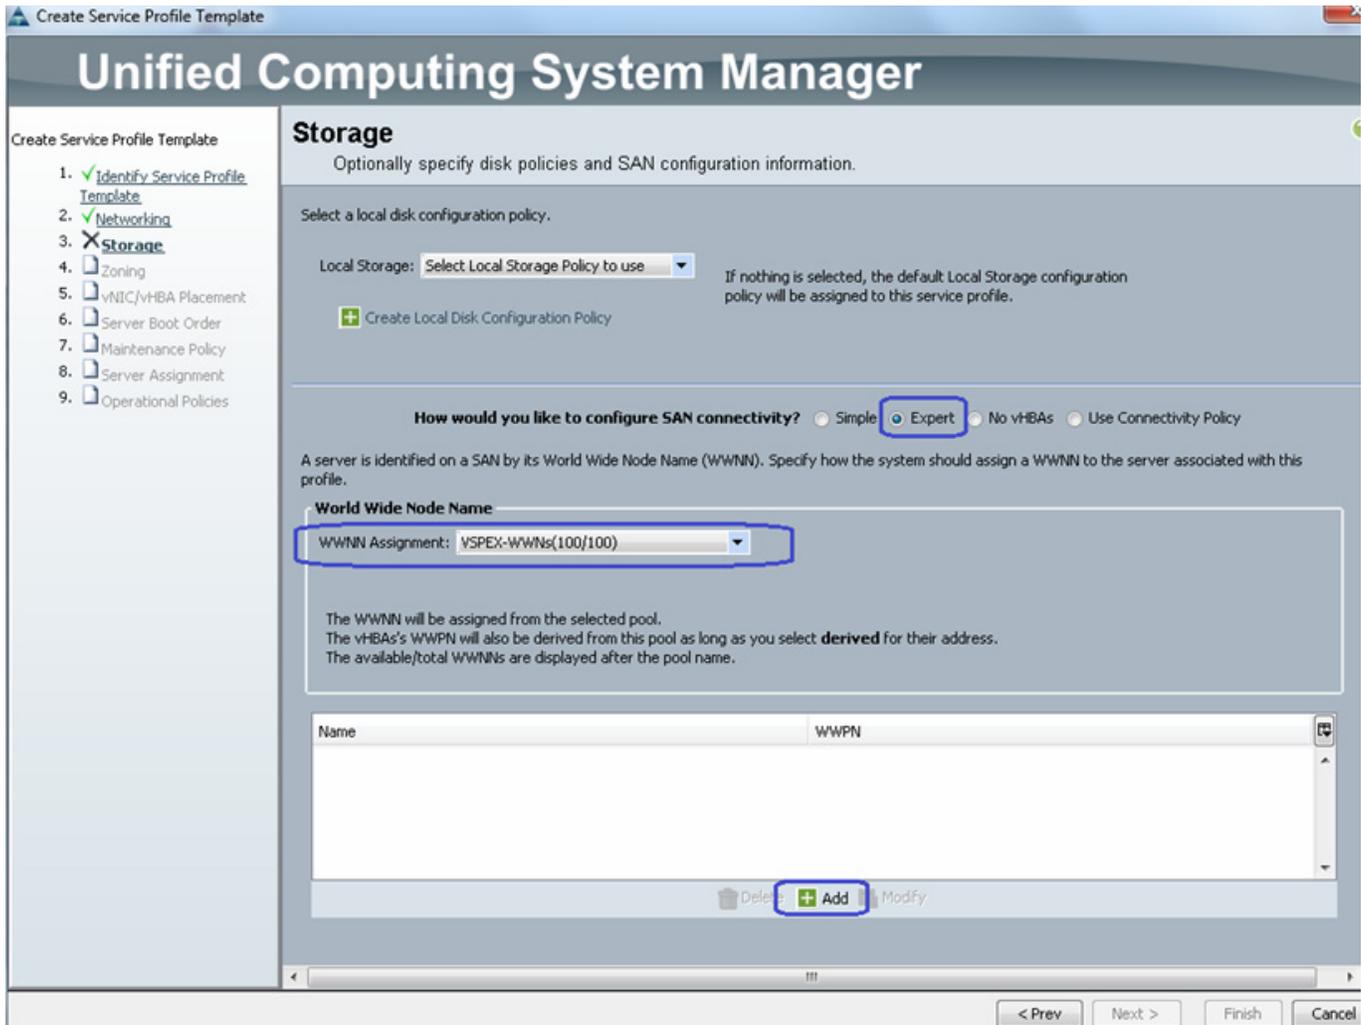
Table 9 Summary of all the vNICs created on the service profile

vNIC Name	MAC address assignment	VLANs	Native VLAN	Fabric	MTU	Adapter Policy	QoS Policy
System-A	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	A	9000	VMware	jumboMTU
System-B	MAC pool	vSphereMgmt, vMotion	vSphereMgmt	B	9000	VMware	jumboMTU
Storage-A*	MAC pool	Storage	Storage	A	9000	VMware	jumboMTU
Storage-B*	MAC pool	Storage	Storage	B	9000	VMware	jumboMTU
Data-A	MAC pool	VM-Data	VM-Data	A	1500	VMware	-
Data-B	MAC pool	VM-Data	VM-Data	B	1500	VMware	-

*Storage vNICs are created for NFS-variant only

- In the Storage window, click the **Expert** radio button for SAN connectivity and choose the option VSPEX-WWNs for WWNN pool from the drop-down list. Click to add vHBA.

Figure 96 Creating Service Profile Template - Storage Configuration Details



10. Enter the vHBA name as vHBA-A in the Name field, Choose the WWPN assignment as **Derived** from the drop-down list, click the **A** radio button for Fabric ID, choose VSAN as **Storage VSAN** from the drop-down list, and choose the Adapter policy as **VMWare**. Click **OK** to deploy the vHBA.

Figure 97 Creating vHBA on Fabric A

Create vHBA

Name:

Use vHBA Template:

Create vHBA Template

World Wide Port Name

WWPN Assignment:

Create WWPN Pool

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

Fabric ID: A B

Select VSAN: Create VSAN

Pin Group: Create SAN Pin Group

Persistent Binding: Disabled Enabled

Max Data Field Size:

Operational Parameters

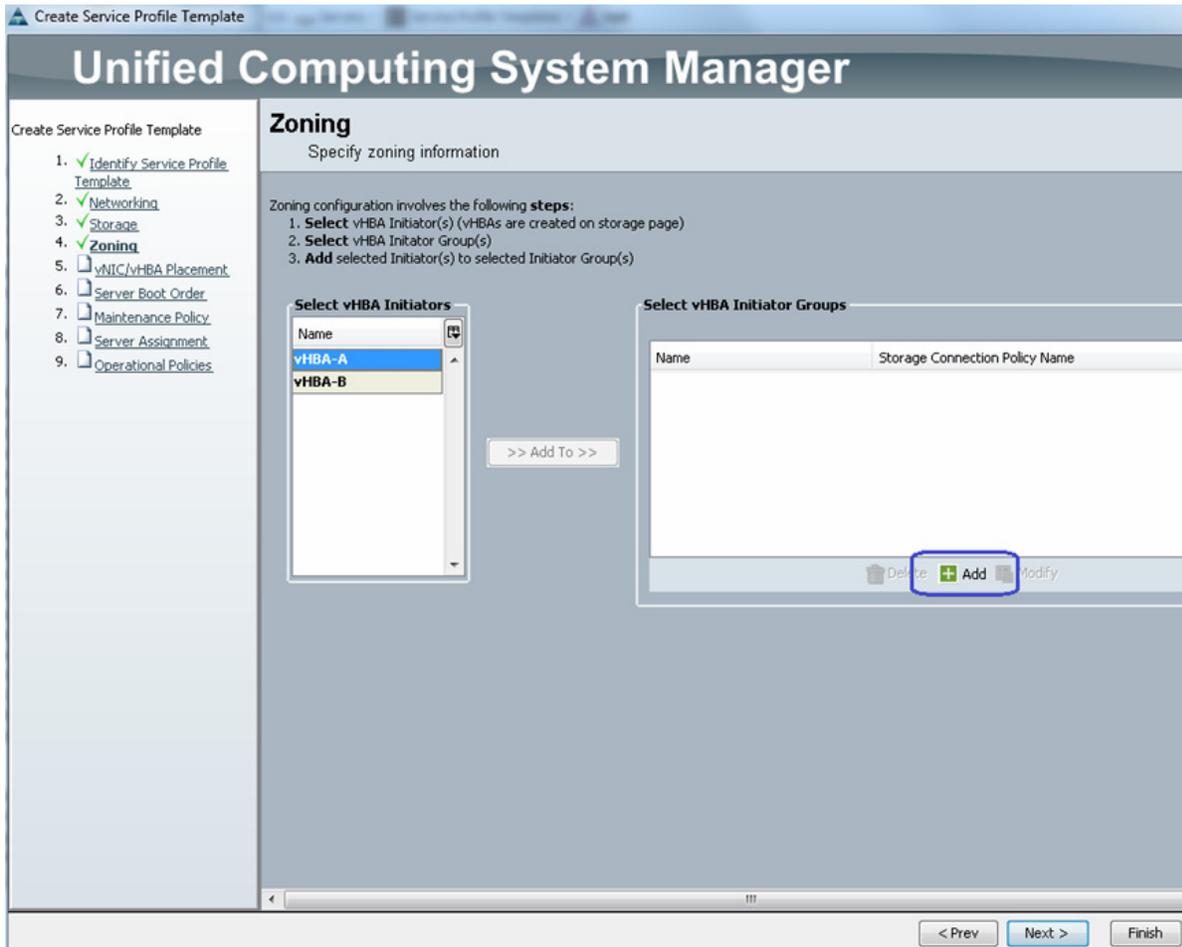
Adapter Performance Profile

Adapter Policy: Create Fibre Channel Adapter Policy

QoS Policy: Create QoS Policy

11. Repeat step 11 for vHBA-B on fabric B, keep all the configuration same as fabric A.
12. For FC-variant of the solution, in the **Zoning** window, click to create a vHBA initiator group.

Figure 98 Creating Service Profile Template - Zoning Details



13. (FC-variant only) Enter SAN-A in the Name field of Create vHBA Initiator Group window, and choose **Fabric-A** zoning policy from the drop-down list and click **OK**.

Figure 99 Creating vHBA Initiator Group for FC-Variant

Create vHBA Initiator Group

vHBA Initiator Group

Name: **SAN-A**

Description:

Storage Connection Policy: **Fabric-A** + Create Storage Connection Policy

Global Storage Connection Policy

Global storage connection policy **defined under org** is assigned to this vHBA initiator group.

Properties

Storage Connection Policy: **Fabric-A**
 Description: **zones for fabric A**
 Zoning Type: **Single Initiator Multiple Targets**

FC Target Endpoints

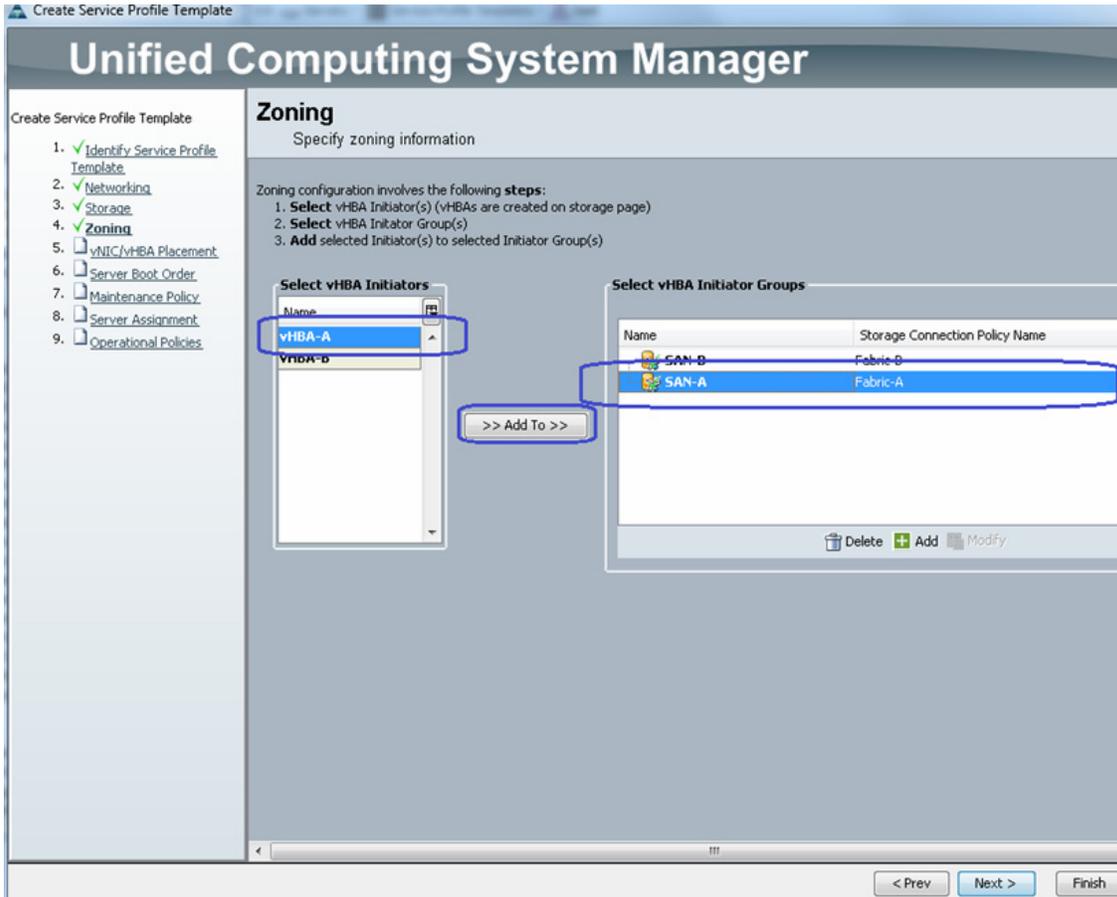
Filter Export Print

WWPN	Path	VSAN
50:06:01:64:3E:A0:65:0A	A	Storage
50:06:01:65:3E:A0:65:0A	A	Storage

OK Cancel

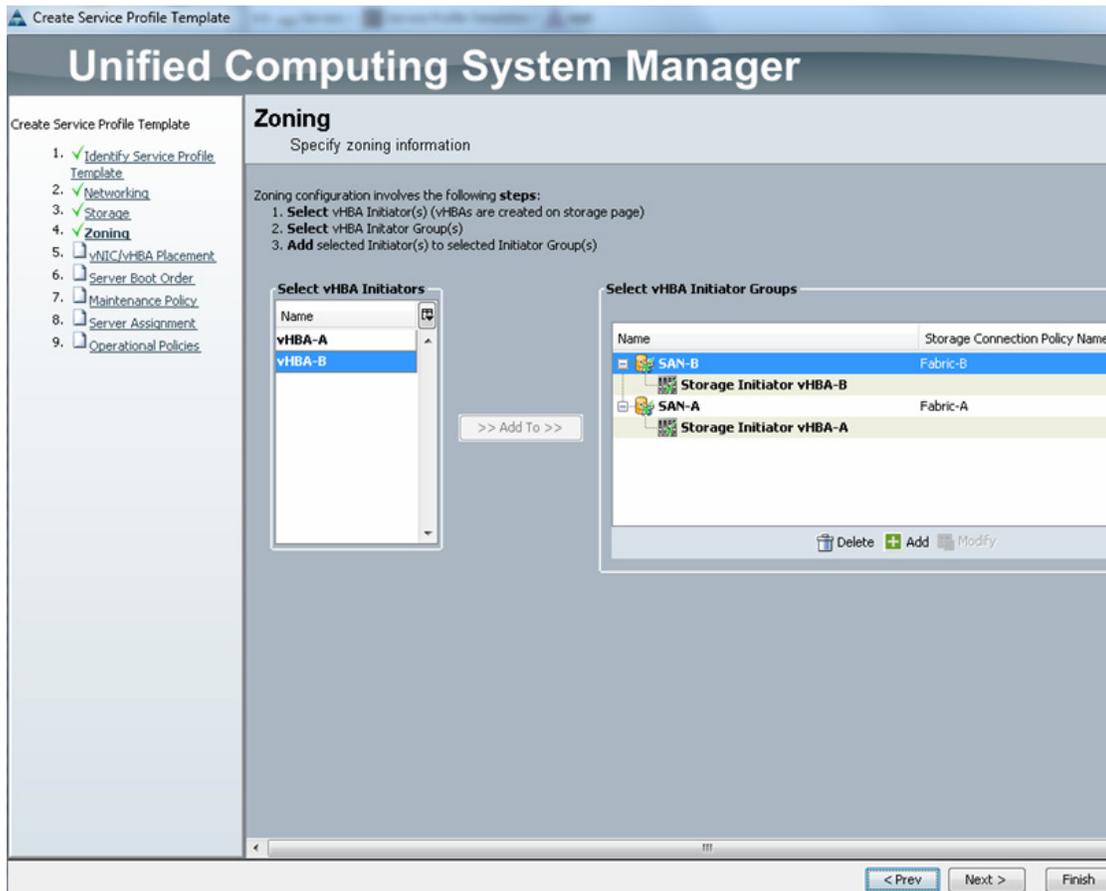
14. (FC-variant) Repeat steps 14 and 15 for the zoning on fabric B. Now choose **vHBA-A** from the list of initiators and **SAN-A** from the initiator-group, and click to add initiator to the selected initiator group.

Figure 100 Adding vHBA to vHBA Initiator Group



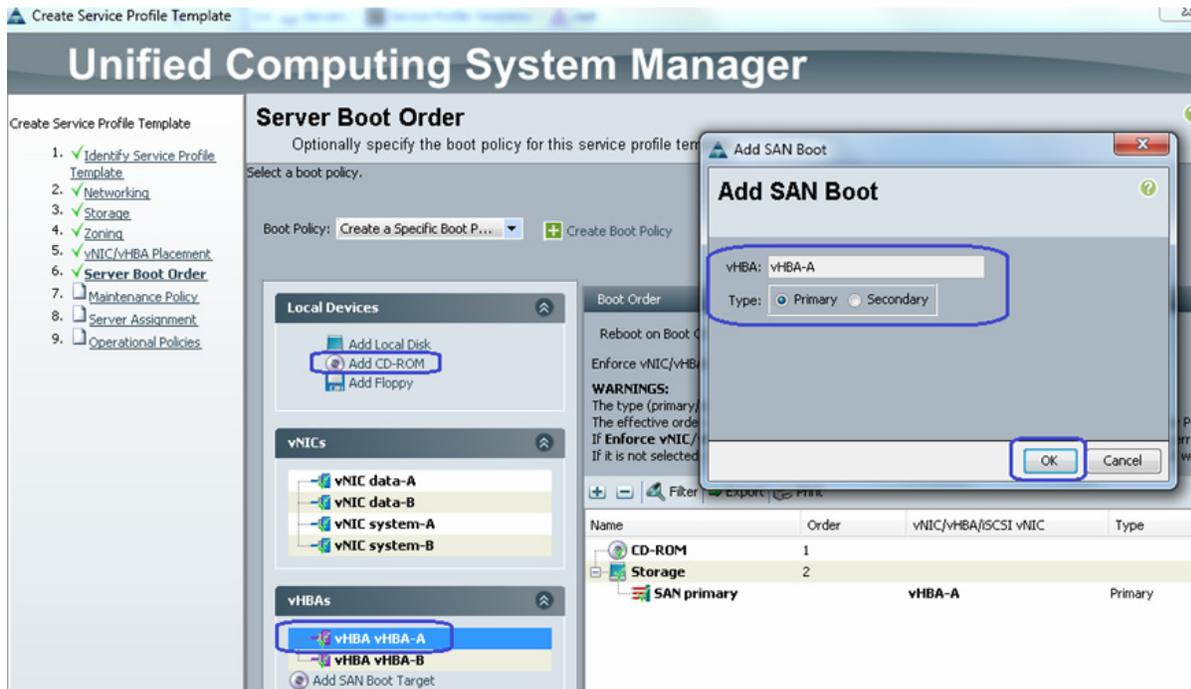
- Repeat step 16 for fabric B as well. The end result should look like Figure 101. Click **Next** and choose the default configuration on the vNIC/vHBA Placement window.

Figure 101 Window Showing vHBAs Added to the Initiator Group



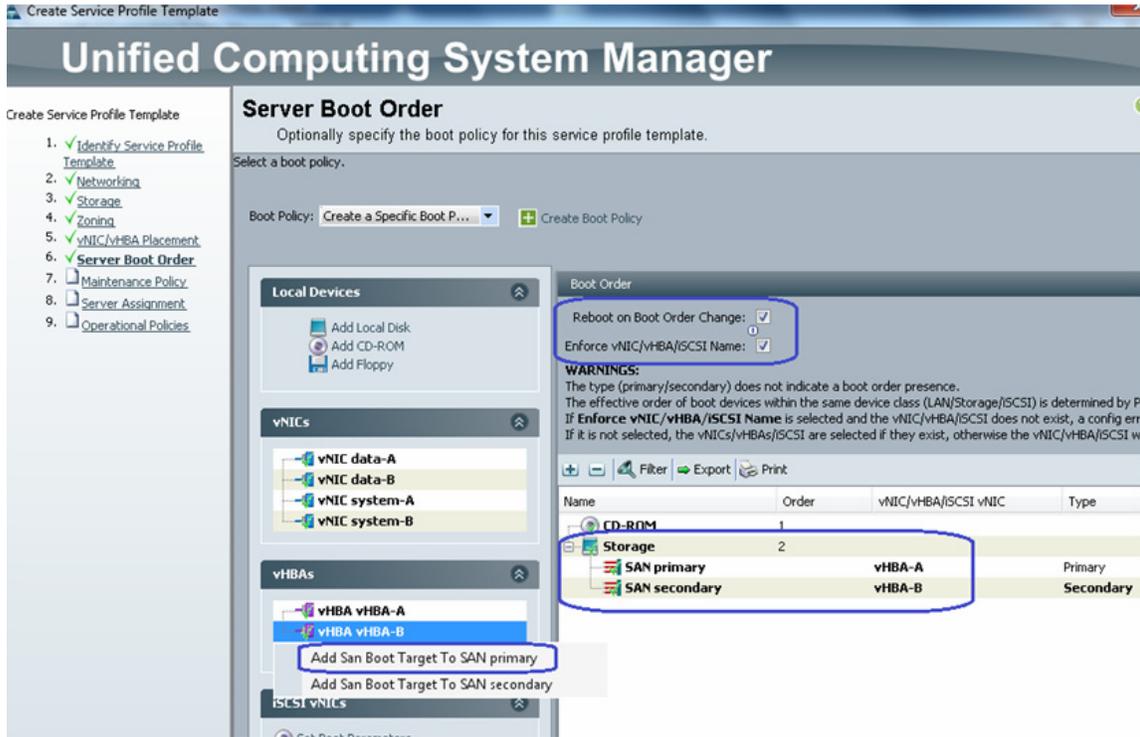
16. For the NFS-variant of the solution, choose the default configuration in Zoning and vNIC/vHBA Placement window. Click **Next**.
17. For both the architectures, in the Server Boot Order window, choose Create a Specific Boot Policy from the drop-down list. Choose the option **Add CD-ROM** as the first boot order choice. Choose **vHBA-A** as the next choice, and provide a name vHBA-A. Click the radio button **Primary** for Type and click **OK**.

Figure 102 Creating Service Profile Template - Configuring Boot Order



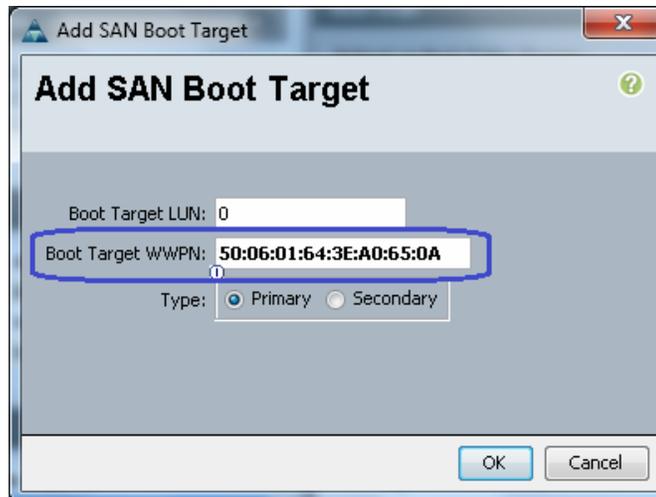
- Similarly, choose **vHBA-B** as the next (secondary) choice to boot from SAN. Once both the vHBAs are added, make sure that the check boxes **Reboot on Boot Order Change** and **Enforce vNIC/vHBA name** are checked. Click **Add SAN Boot Target** under the vHBAs, and click **Add San Boot Target** to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

Figure 103 Adding SAN Boot Target to SAN Primary



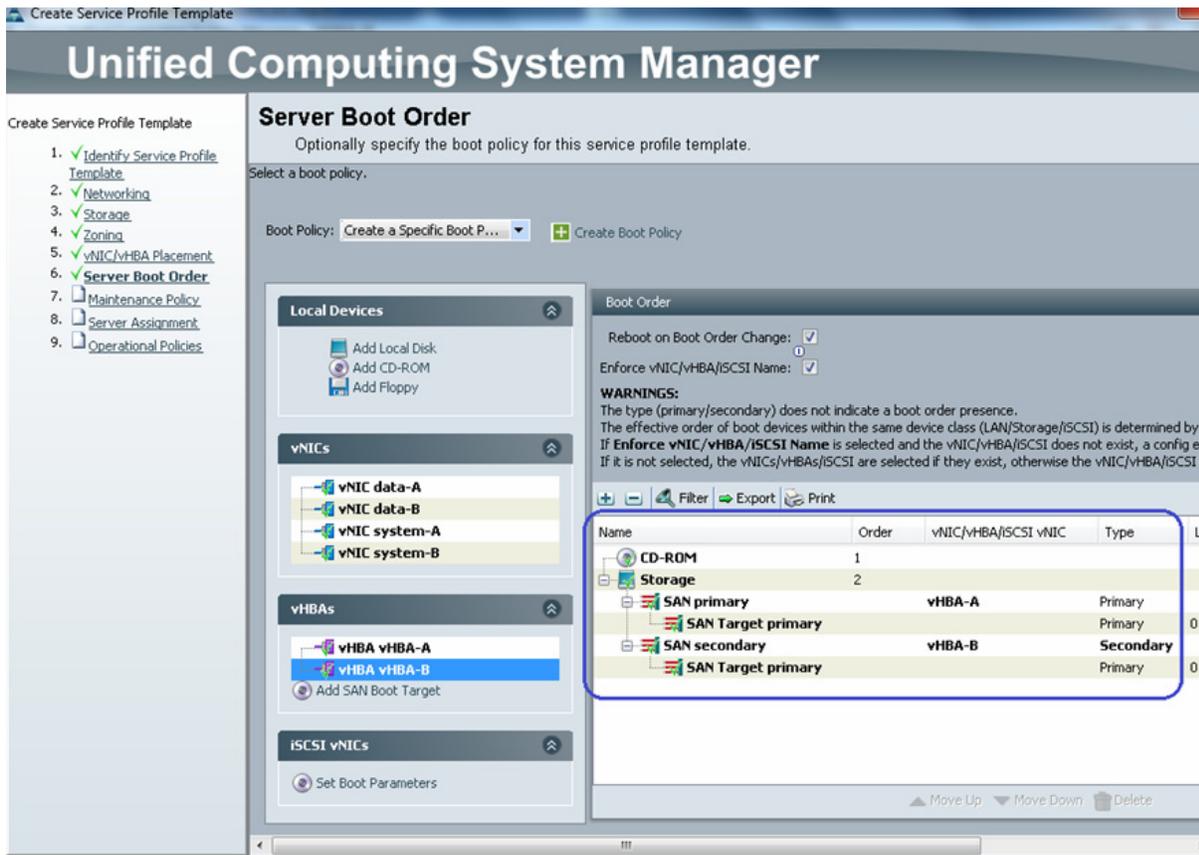
19. Provide target WWPN of the VNX storage device (which can be obtained using show **logi database NX-OS** CLI command executed under **connect nxos {alb}** shell as described in the previous subsection). Keep the target as **Primary**.

Figure 104 Adding SAN Boot Target as Primary



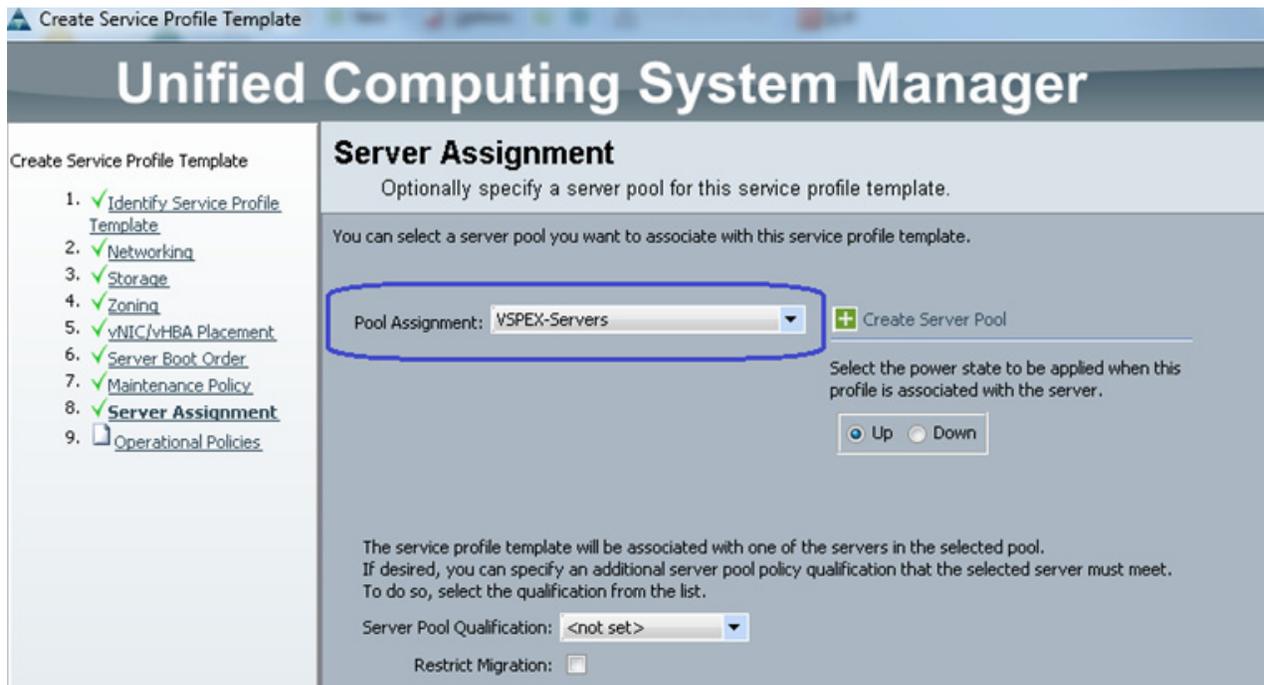
20. Repeat step 21 for the fabric B as well. The end result will look like [Figure 105](#).

Figure 105 Server Boot Order After the Configuration is Complete



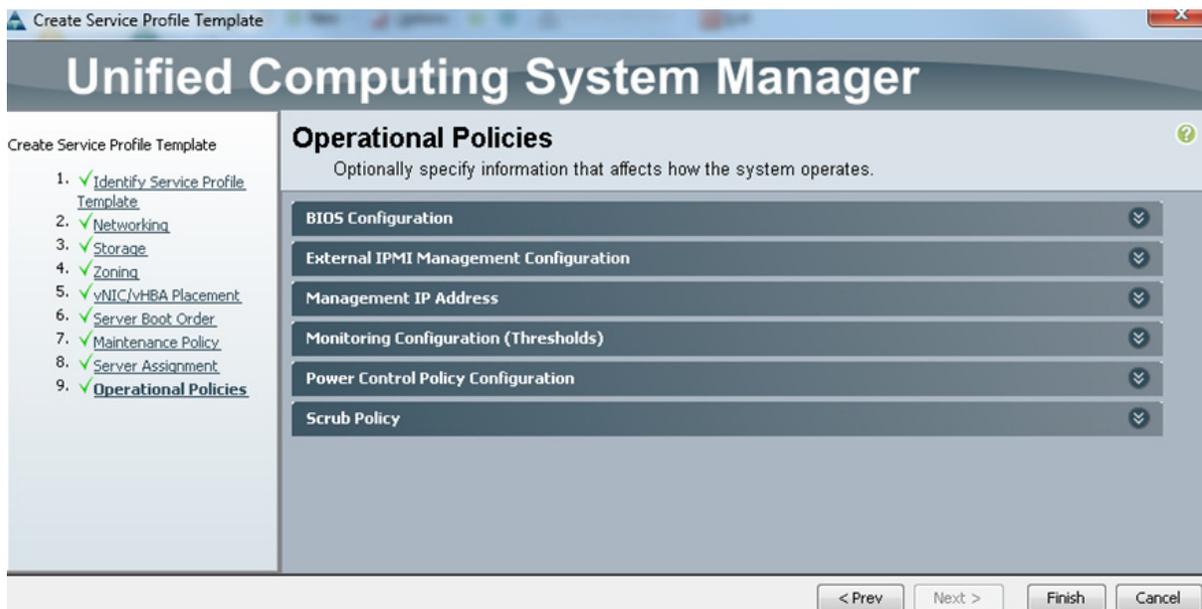
- Click **Next** to go to the Maintenance Policy window. Keep all the fields at default and click **Next** to continue to Server Assignment window. For Pool Assignment, choose the Server Pool created in the previous sub-section. Click **Next**.

Figure 106 *Creating Service Profile Template - Configuring Server Assignment*



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

Figure 107 *Creating Service Profile Template - Restore Default Settings for Operational Policy*

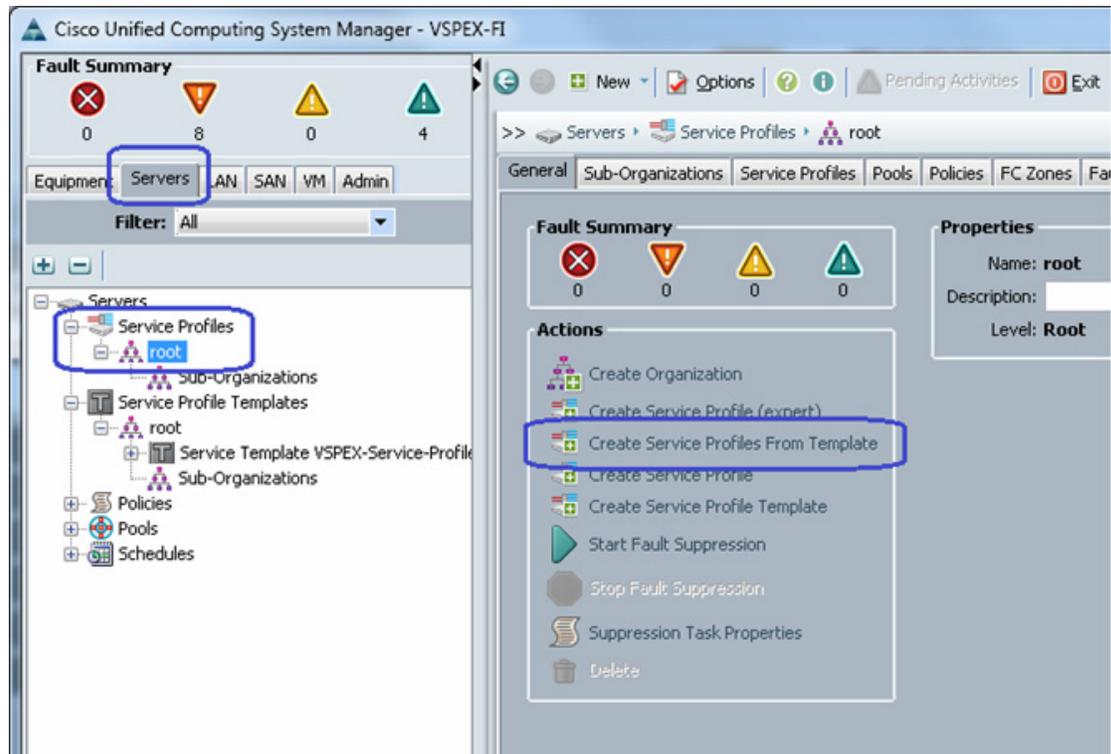


Instantiate Service Profiles from the Service Profile Template

As a final step to configure UCS Manager, we need to instantiate service profiles from the service profile template created in “[Configure Service Profile Template](#)” section on page 91. Follow these steps to instantiate service profiles from the service profile template:

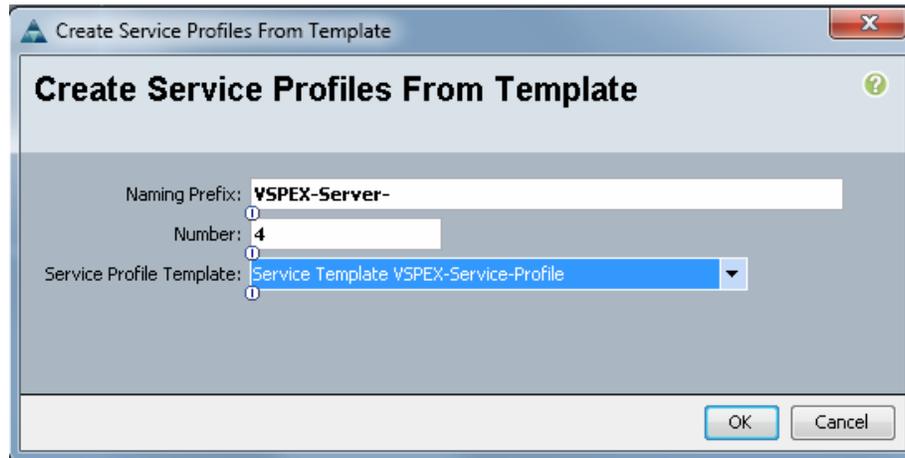
1. From the **Servers** tab, expand **Servers > Service profiles > root**, and click the **Create Service Profile from Template** link in the right pane.

Figure 108 *Creating Service Profile from Template*



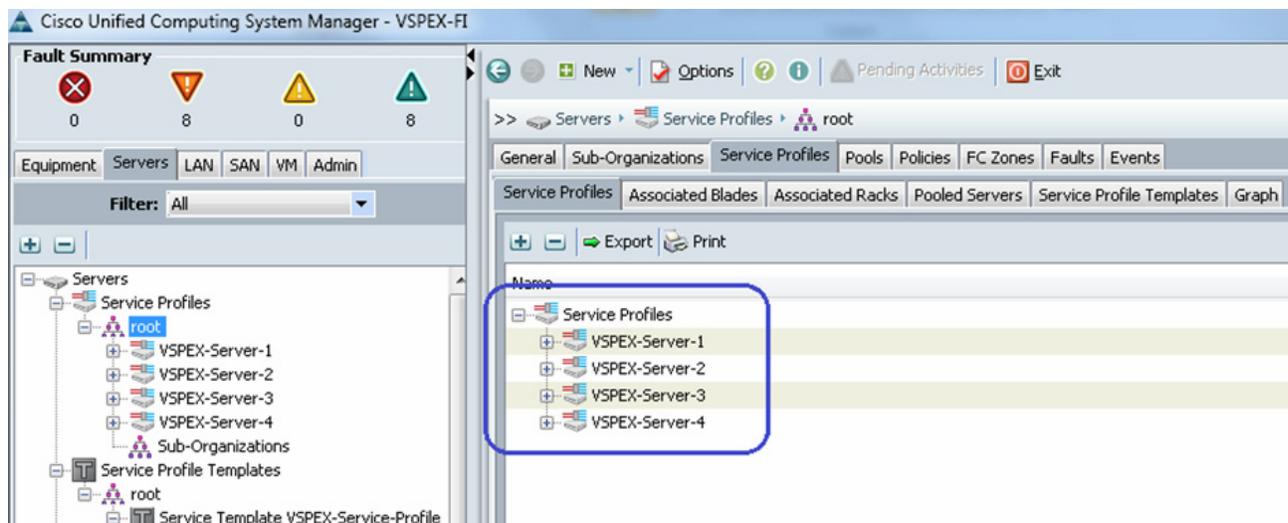
2. Providing the naming prefix, number of service profiles to be instantiated and choose the service profile template from the drop-down list. Refer to the sizing guidelines for the number of servers needed for your deployment.

Figure 109 Details for Creating Service Profiles



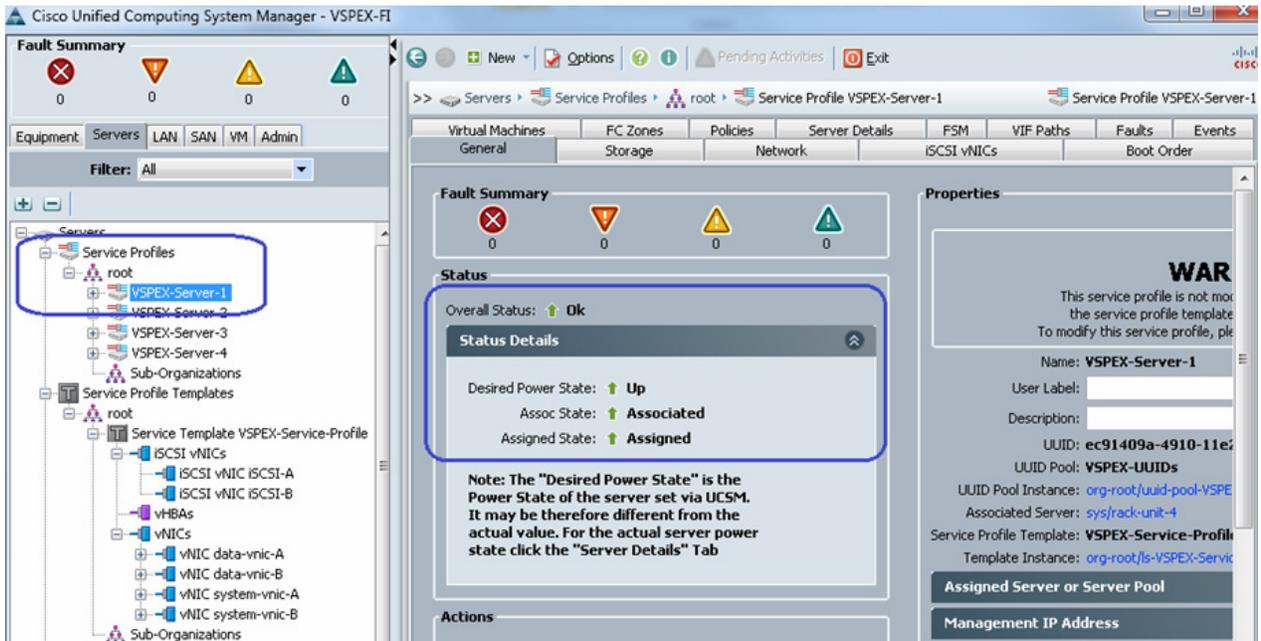
3. There are four service profiles are created in this example.

Figure 110 Window Showing All the Service Profiles Created from the Template



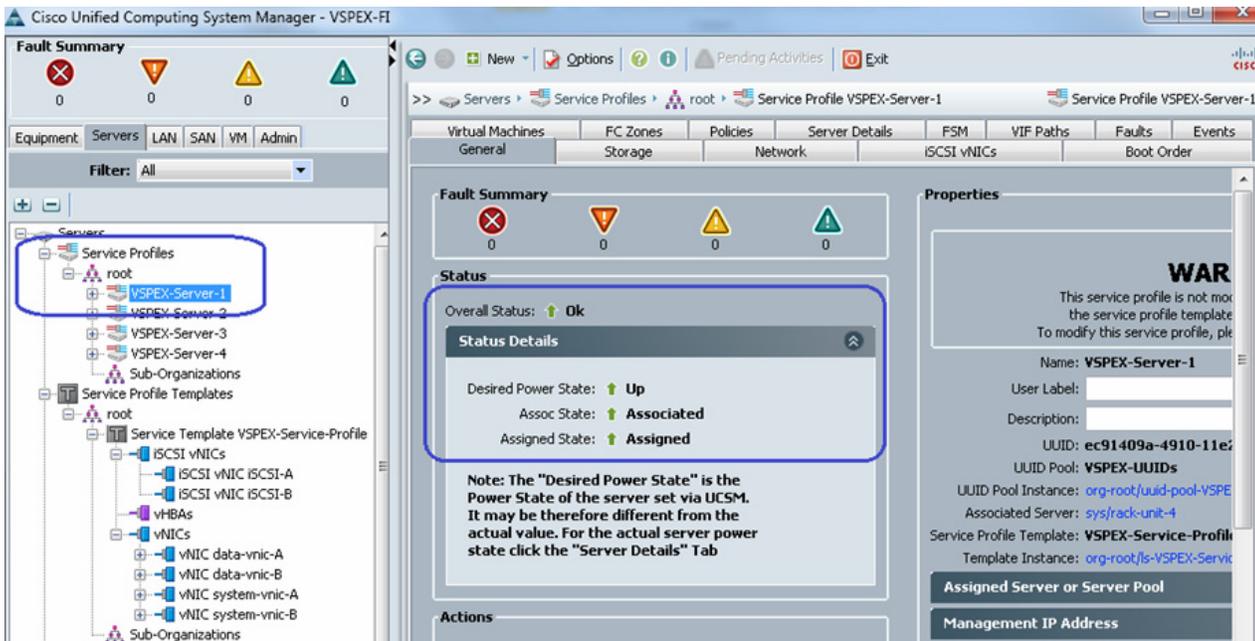
4. As the service profile template is assigned to a server pool, the service profiles instantiated from the template would be assigned to individual server resource from the server pool as far as they are available. You can click on a given service profile to see its association state, and with which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

Figure 111 Status Details Of Service Profiles



- Eventually, all the four servers will be associated – you can see the summary by clicking **Servers** in the **Equipment** tab.

Figure 112 Summary of Service Profiles Showing Assigned State as Associated



**Note**

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

Configure Data-Stores for ESXi images

This section provides necessary steps to create FC accessible data-stores for the ESXi boot image per server basis. This can be done in three steps:

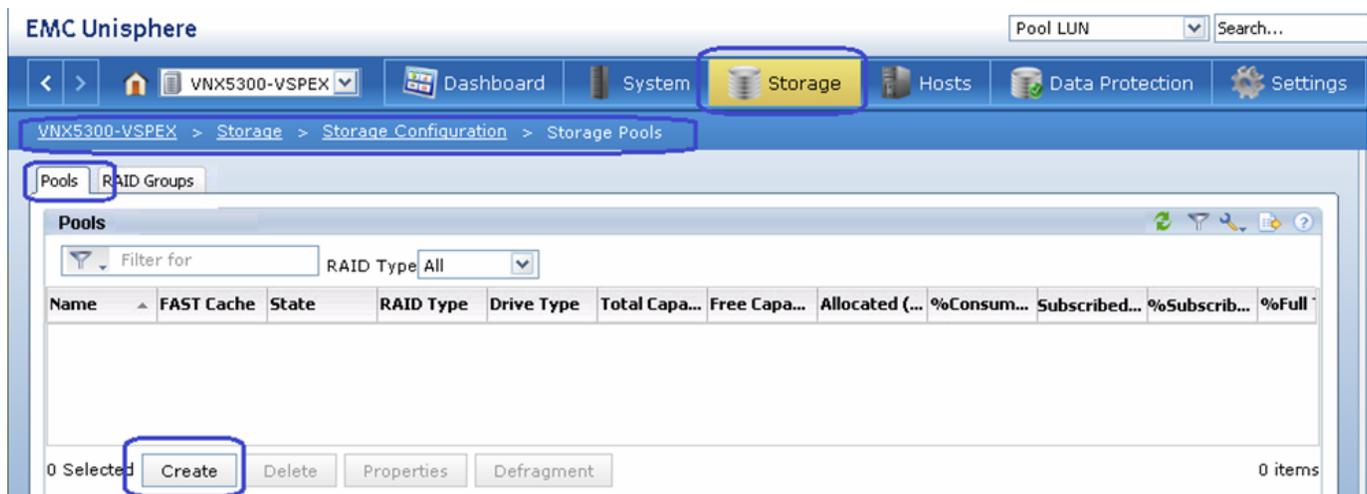
1. [Configure Storage Pool, page 111](#)
2. [Register Hosts, page 114](#)
3. [Configure Storage Groups, page 122](#)

Configure Storage Pool

To create storage pool and carve boot LUNs per server basis, follow these steps:

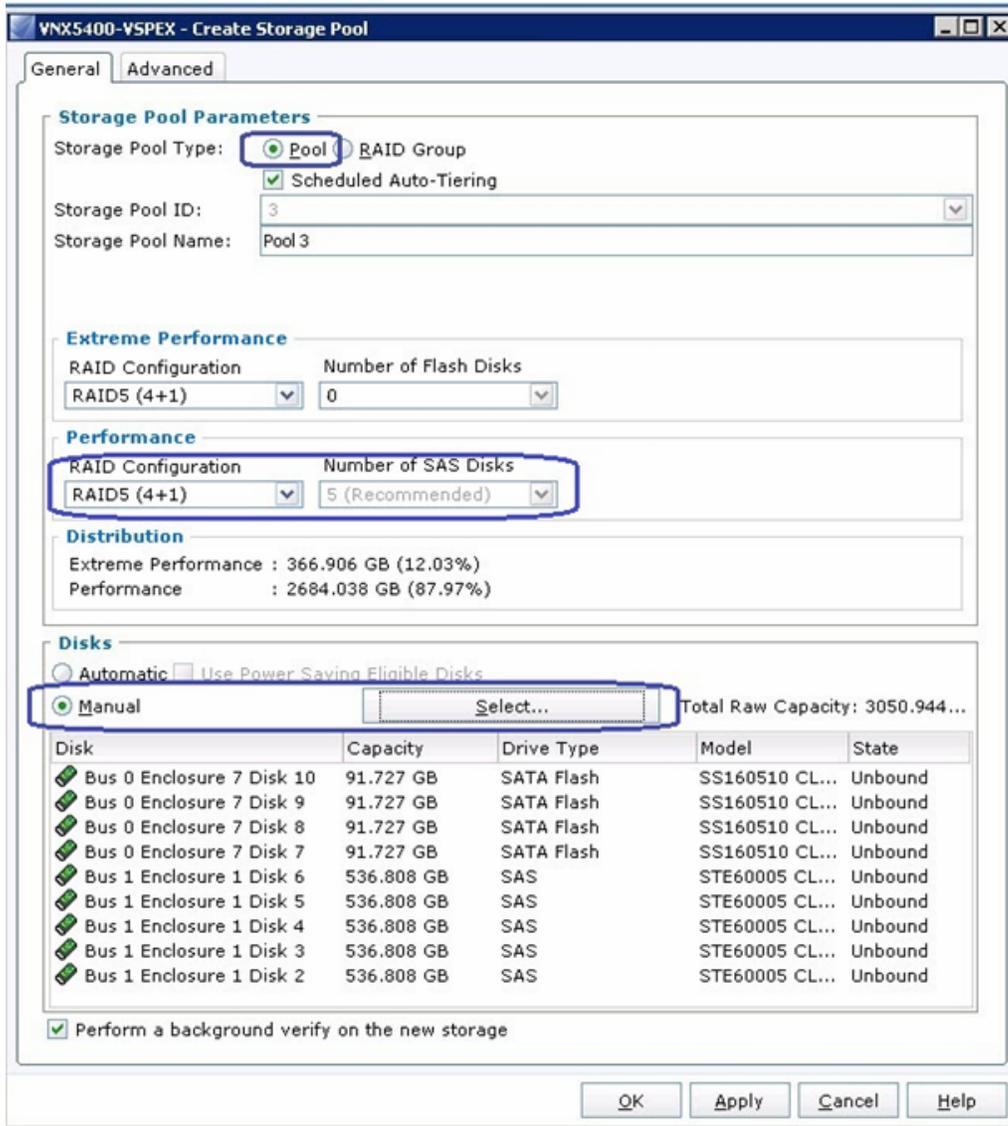
1. Connect to EMC VNX Unisphere GUI, click the **Storage** tab. Choose **Storage Configuration > Storage Pools**. Click the **Pools** tab, in the Pools window, choose **Create**.

Figure 113 *Creating Storage Pools in EMC Unisphere*



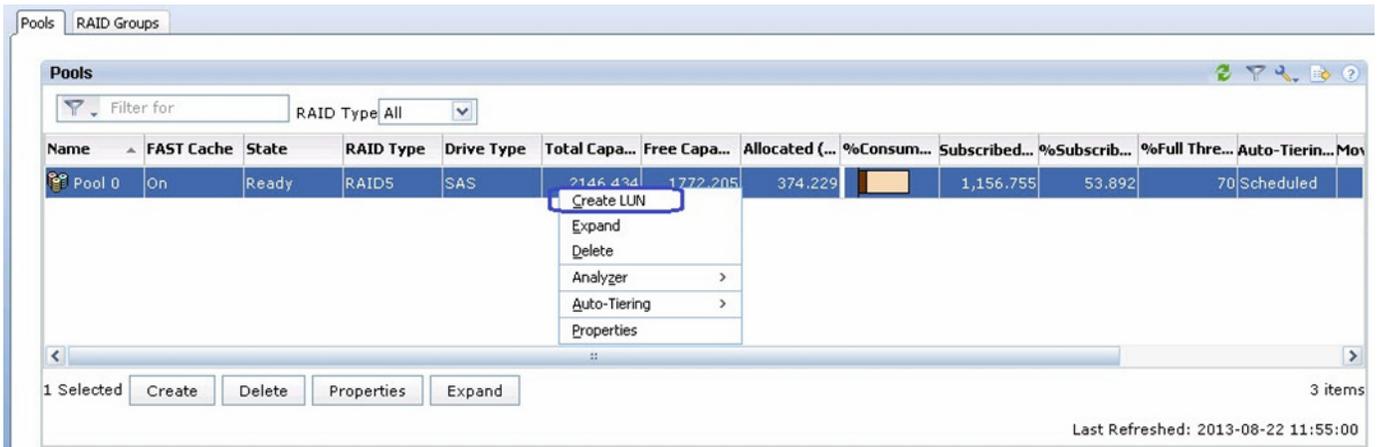
2. Choose the RAID Configuration as RAID5 (4 + 1) from the drop-down list for performance. Click the **Manual** radio button, and click **Select** to manually select 5 SAS disks to create the storage pool.

Figure 114 Details for Creating Storage Pools



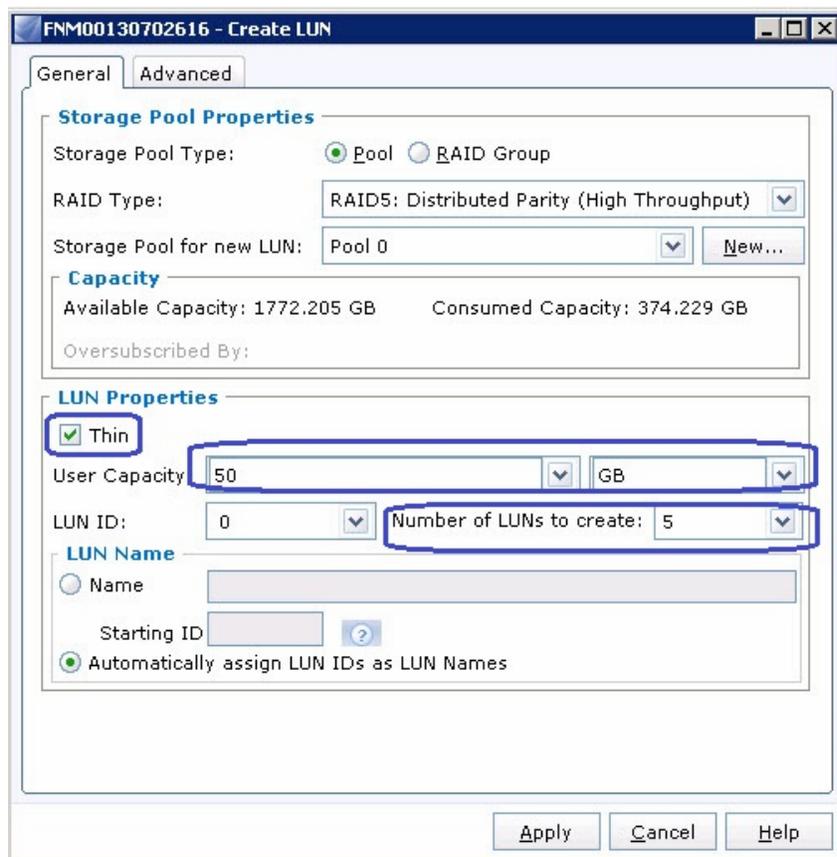
- From the newly created RAID group, right-click and choose **Create LUN**.

Figure 115 Creating LUN from Created RAID Group



4. Choose 5 for Number of LUNs to create, with 50 GB User Capacity each. Make sure the **thin provisioning** check box is checked.

Figure 116 Details for Creating LUN



Register Hosts

After the service profiles are associated in UCS Manager, the vHBAs will do flogi in the network and the SAN initiators will be identified by the VNX storage array. To register the hosts identified by the WWPN of the server, follow these steps:

(NFS-variant only) For NFS-variant of the solution, the storage connectivity is through Nexus 5000 switches. In that case, the FC zoning must be configured manually on Nexus 5548UP switches. In contrast to this, in case of FC-variant architecture, where storage is attached to FIs, FC zoning is taken care by UCS Manager implicitly.

To configure zoning on Nexus 5548 UP switches, follow these steps:

1. Login to the Nexus 5548UP switch A and configure a zoneset for SAN fabric A. You need to create one zone for each ESXi host, containing WWPN of SP-A and SP-B of VNX storage and WWPN of the vHBA on fabric A of the ESXi server. WWPN list is available from UCS Manager as shown in step 7. Entire zoneset configuration looks like [Figure 117](#). Activate the zoneset in the storage VSAN after the zoneset is configured completely.

Figure 117 Zoneset Configuration on Cisco Nexus 5548UP Switch

```

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 pwnn 20:00:00:25:b5:66:dd:0e
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:1d
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:1c
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:1b
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:1a
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:09
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:08
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:07
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:06
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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 pwnn 20:00:00:25:b5:66:dd:05
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 50:06:01:64:3e:a0:52:02
UCS-N5k-FabA(config-zoneset-zone)# member pwnn 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)#

```

- (NFS-variant only) Validate the successful activation of zoneset by issuing **show zoneset brief** command.

Figure 118 Running `show zoneset brief` Command for 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#
```

3. (NFS-variant only) Similarly, on Nexus 5548UP switch B, create a zoneset for fabric B and include vHBAs on fabric B on the servers. Zoneset on fabric B looks like [Figure 119](#).

Figure 119 Running `show zoneset brief` Command for 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#
```

4. (NFS-variant only) To further validate zoneset configuration across entire SAN fabric, SSH to UCS FI-A, issue `connect nxos` command, and run `show npv flogi-table`. It should list all ten FLogI sessions, one from each vHBA on fabric A in storage VSAN.

Figure 120 Running `show npv flogi-table` to List All the FLogI Sessions

```
V250-UCS-A# connect nxos
Cisco Nexus Operating System (NX-OS) Software
T&C 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)#
```

5. (NFS-variant only) 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, and 2 from VNX storage array's SP-A and SP-B FC ports. Similarly, verify the FLogI entries on SAN fabric B.

Figure 121 Running show flogi database Command on Cisco Nexus 5548UP

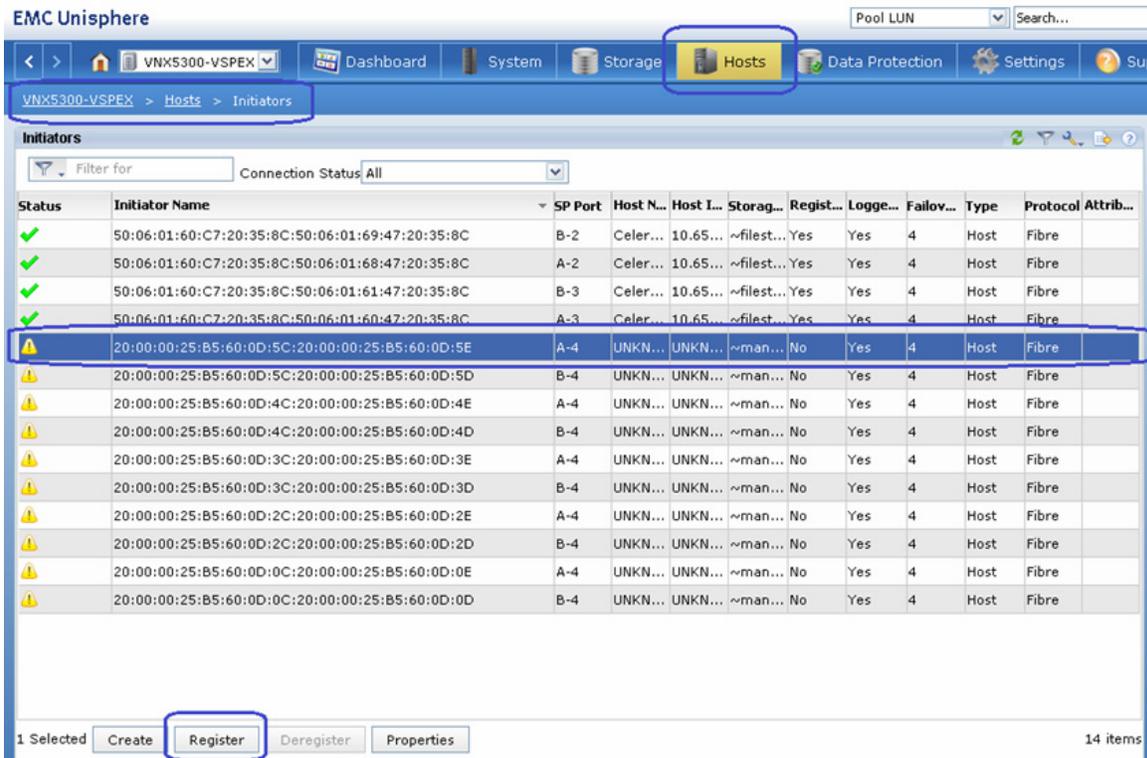
```
UCS-N5k-FabA# show flogi database
-----
INTERFACE          VSAN    FCID          PORT NAME      NODE NAME
-----
fc1/29              10      0x5c0000     20:41:00:0d:ec:f7:04:00  20:0a:00:0d:ec:f7:04:01
fc1/29              10      0x5c0002     20:00:00:25:b5:66:dd:0e  20:00:00:25:b5:60:0d:0e
fc1/29              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#
```

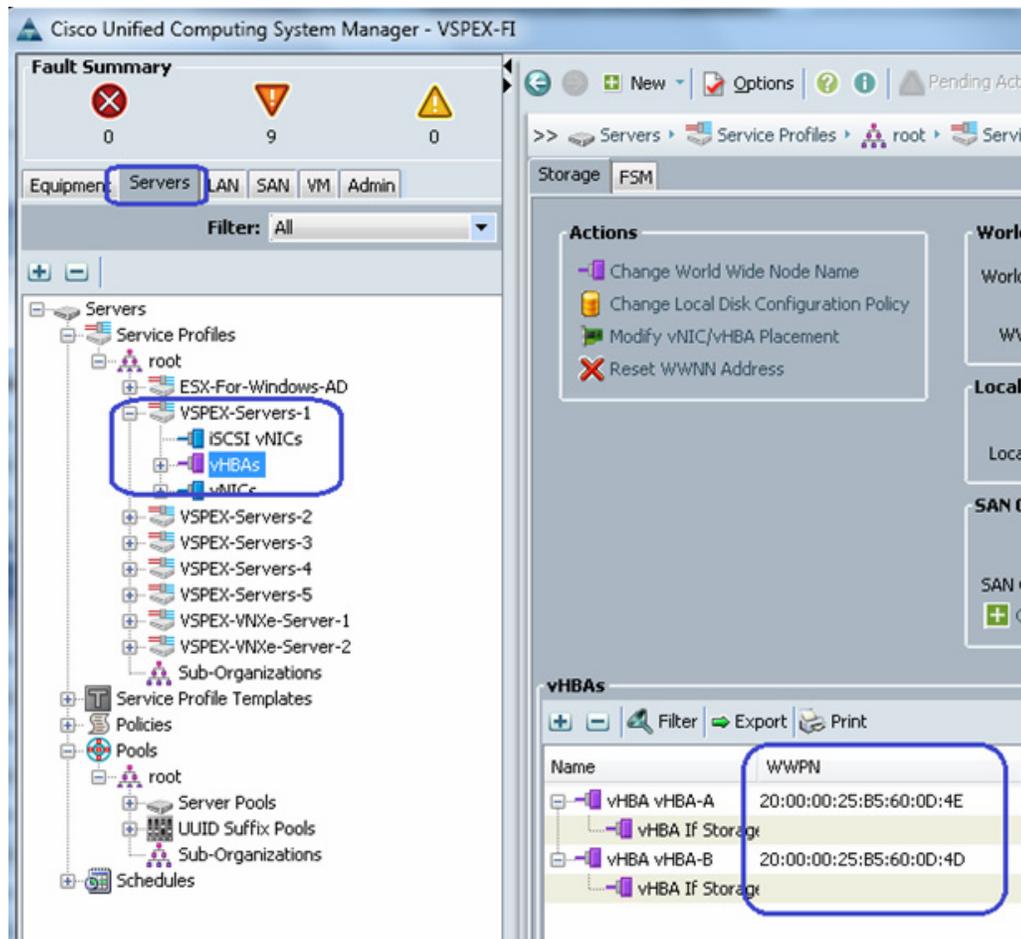
- In the Unisphere GUI, click the **Hosts** tab, and click **Initiators**. Choose the first unregistered initiator and click **Register**.

Figure 122 Registering Unregistered Initiator in the EMC Unisphere



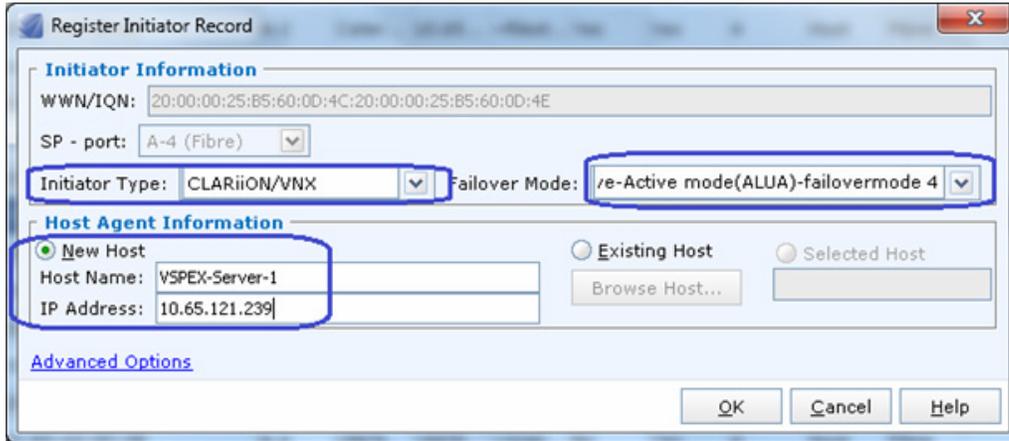
- From the UCS Manager GUI, click the **Servers** tab, expand **Servers > Service Profiles > root > Service profile** (a specific service-profile), and click **vHBAs**. This will list the WWPN identities. Using this IDs, you can associate WWPN to the server.

Figure 123 *WWPN Identities of vHBAs in Cisco UCS Manager GUI*



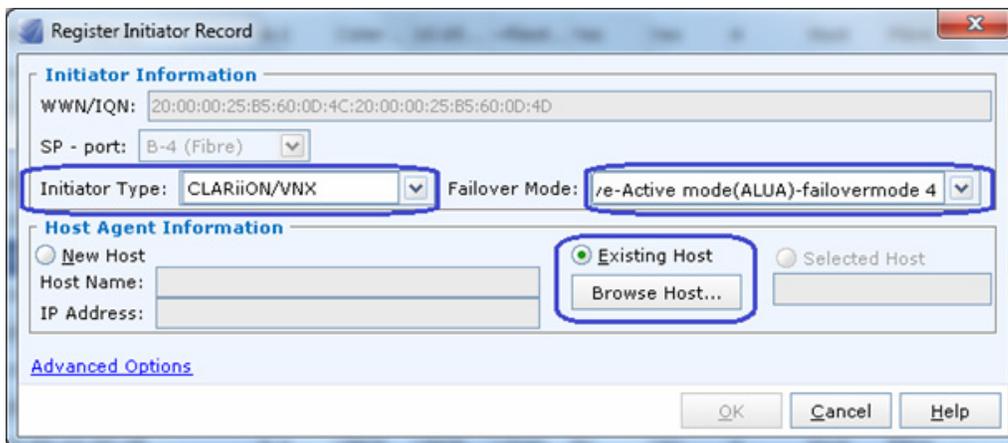
- In the Register Initiator Record wizard, choose the Initiator Type as **CLARiiON/VNX** and Failover Mode as **failovermode 4** from the drop-down lists respectively. Click the **New Host** radio button, enter the hostname and (future) management IP address of the host. Click **OK**.

Figure 124 Registering Initiator Record - Part 1



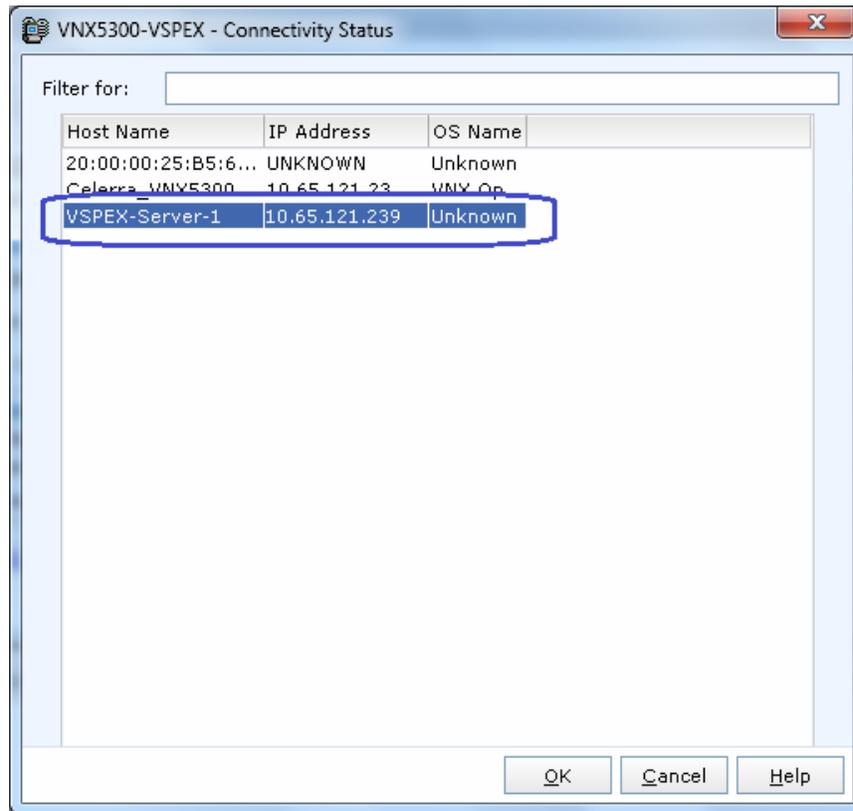
9. Choose the second vHBA's WWPN from the same server, click **Register** and now, click the **Existing Host** radio button. Click **Browse Host** to manually select the host.

Figure 125 Registering Initiator Record - Part 2



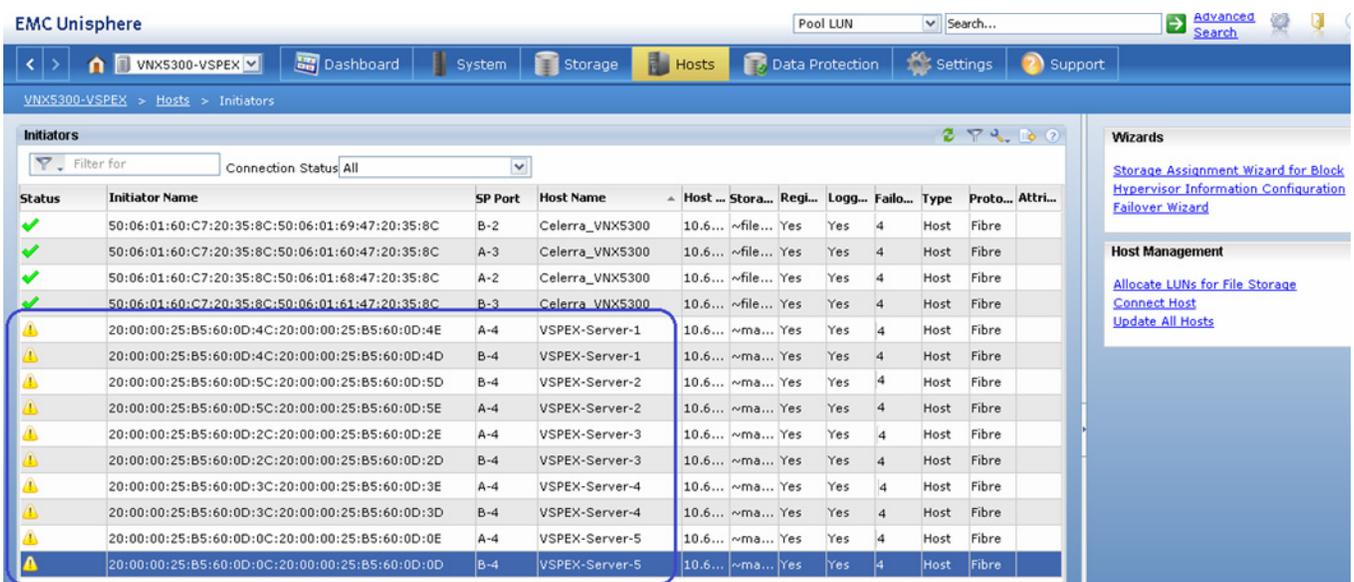
10. Choose the previously registered host, and click **OK**.

Figure 126 Choosing Registered Host for EMC VNX Connectivity



11. Repeat these steps for all the servers in the group. End result looks like Figure 127.

Figure 127 Initiator Window Showing All the Hosts with Initiator Names

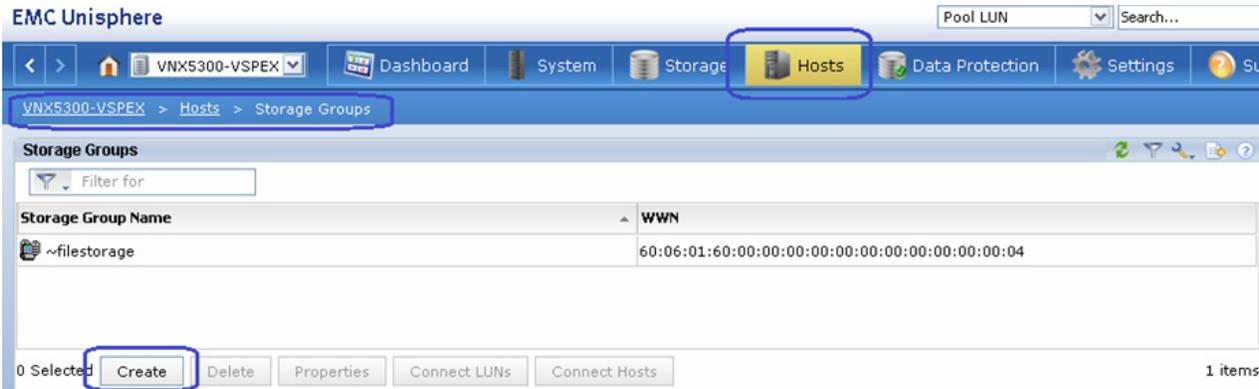


Configure Storage Groups

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

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

Figure 128 *Creating a Storage Group*



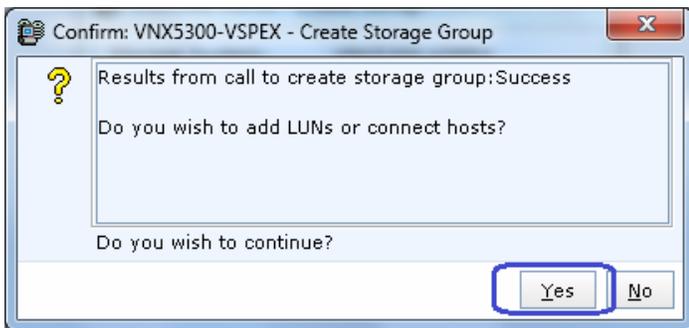
2. Enter a name to the storage group name in the Name field.

Figure 129 *Name the Storage Group*



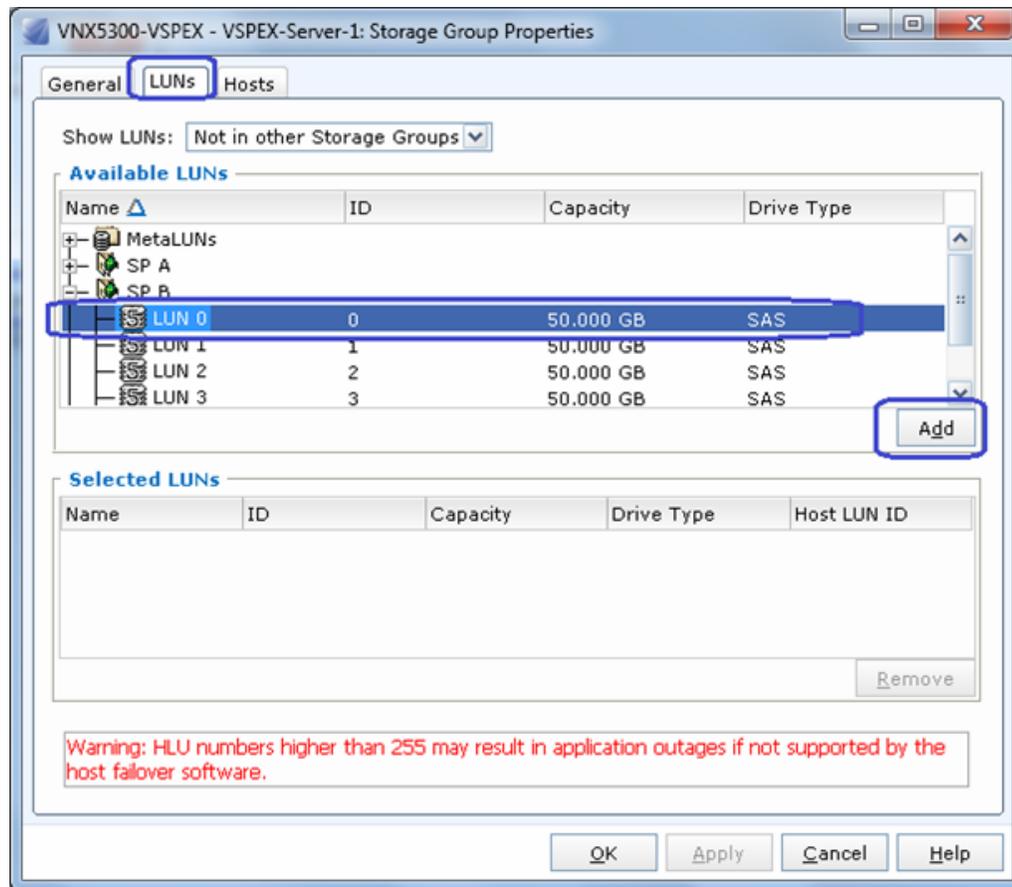
3. You will see the success message after the creation of new storage group. The system prompts to create LUNs and connect hosts. Click **Yes**.

Figure 130 *Successful Completion of Storage group Creation*



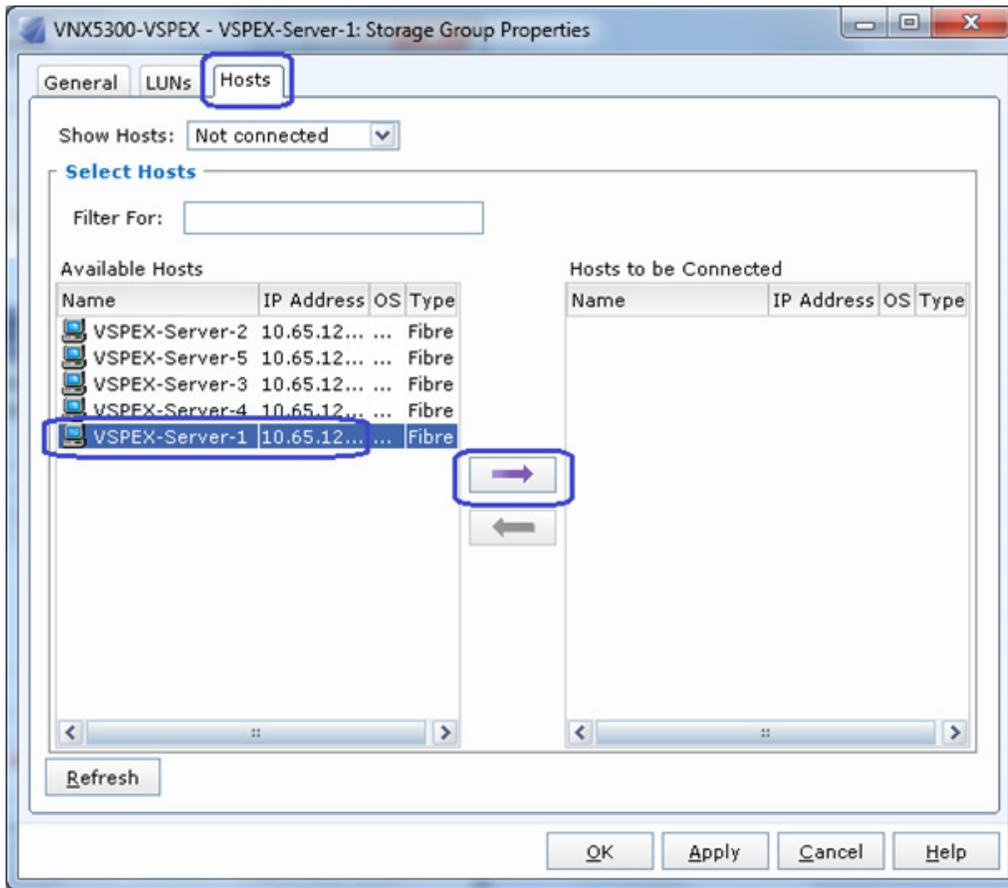
- From the **LUNs** tab, select a LUN and click **Add**.

Figure 131 Adding LUNs to the Storage Group



- Click the **Hosts** tab, and choose a server to add on the storage group. Click **OK** to deploy the storage group.

Figure 132 Adding Host to the Storage Group



6.Repeat these steps for all the five servers. The end result looks like [Figure 133](#).

Figure 133 Storage Group Showing All the Servers Added

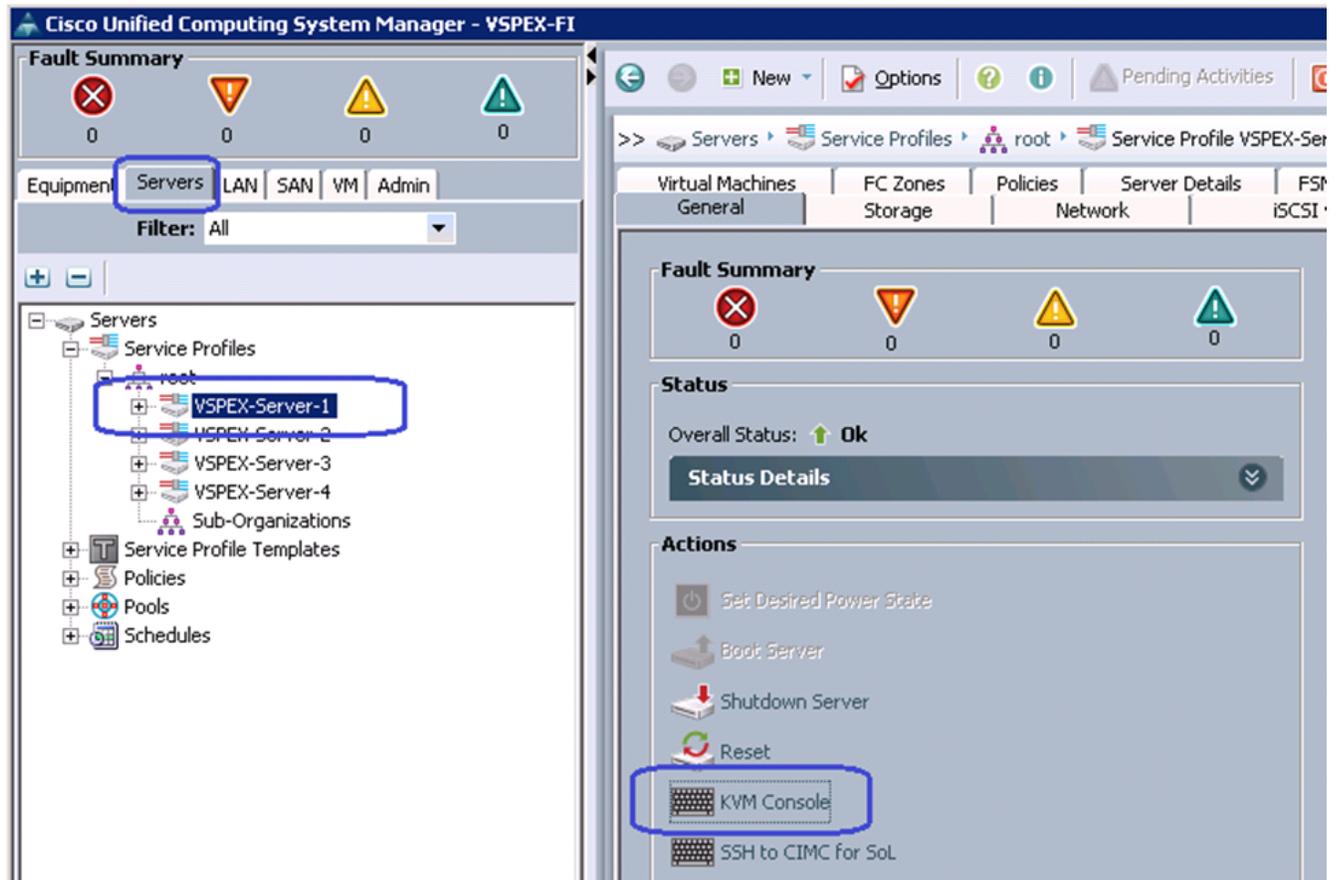
Now, we have end-to-end FC storage access from servers in UCS to the specific boot LUN on the VNX storage devices. We are ready to install ESXi images on the server.

Install ESXi Servers and vCenter Infrastructure

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

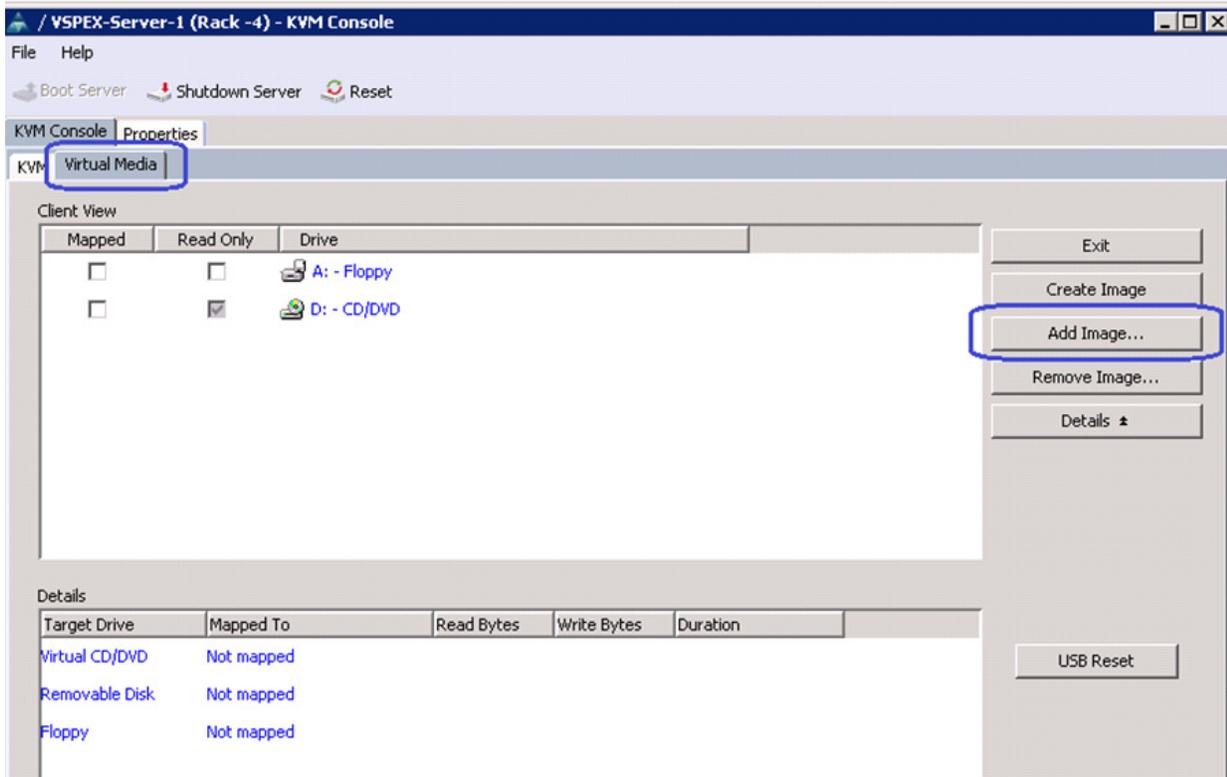
1. From UCS Manager GUI, click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a particular service profile. Click KVM Console in the right pane of the window.

Figure 134 Launching KVM Console



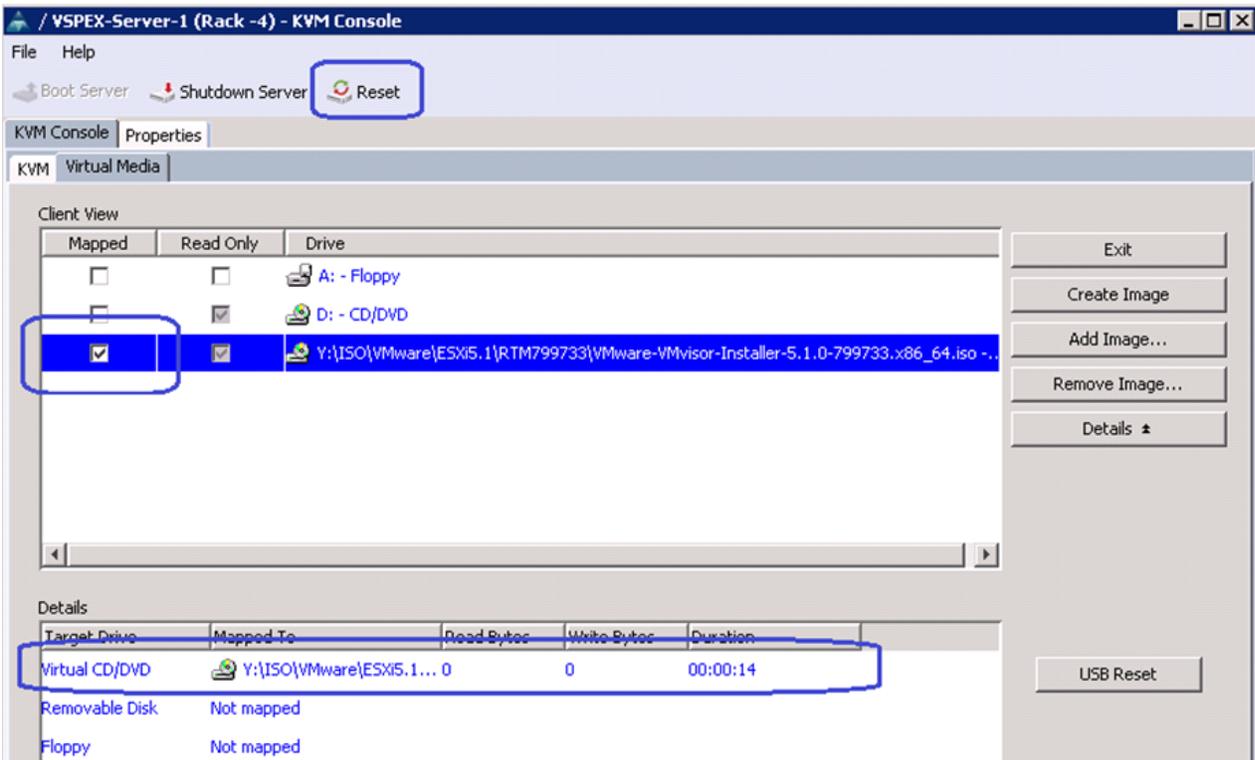
2. Once the Java pallet of KVM is launched, click the **Virtual Media** tab and click **Add Image**. A window appears to select an ISO image. Browse through the local directory structure and select ISO image of the ESXi 5.1 hypervisor installer media.

Figure 135 Adding an ISO Image of the ESXi 5.1 Hypervisor Installer Media



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

Figure 136 Mapping the ISO Image and Resetting the Server



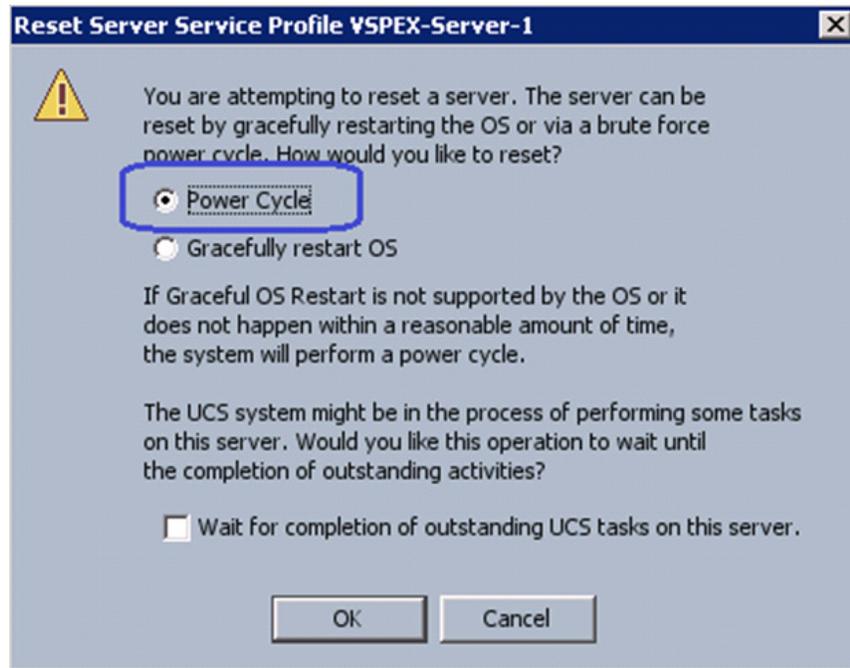
4. Click **OK** in the Reset Server warning message window.

Figure 137 Warning Message for Resetting the Server



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

Figure 138 *Selecting Resetting Option*

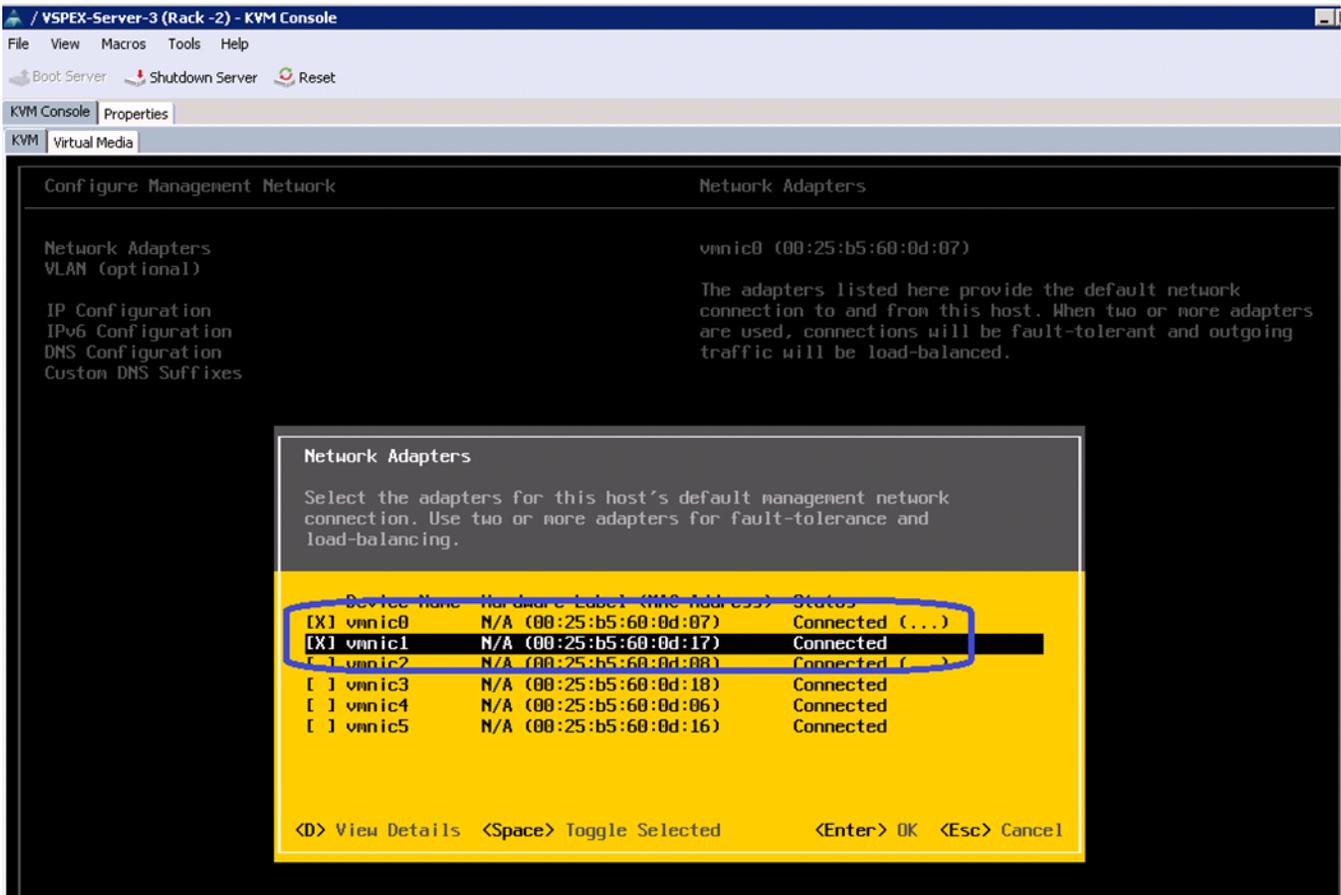


6. Click the **KVM** tab to see how ESXi boot media is booted from the virtual CD-ROM drive.

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

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

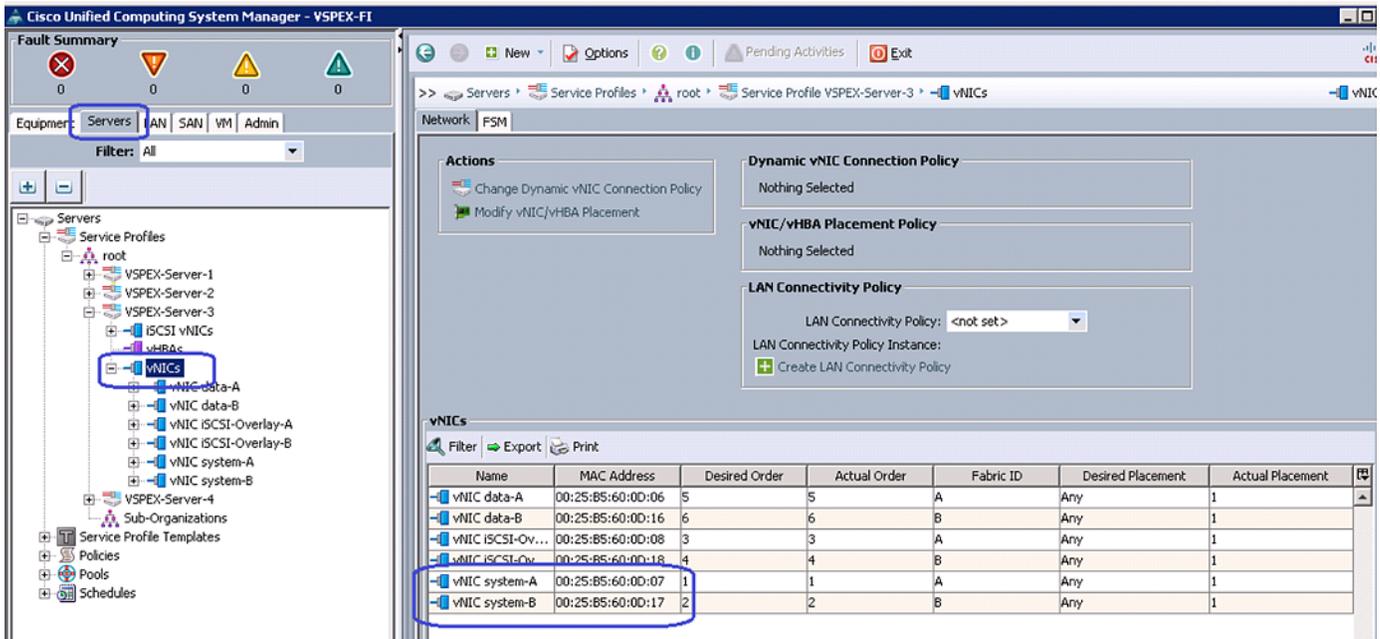
Figure 139 Selecting Adapters for Default Management Network Connection



Note

Easiest way to figure out which vmnic adapter should be used for the vSphere management purpose, you can identify the vmnic by MAC address. The MAC addresses of the vNICs (vmnic's) are summarized in UCS Manager GUI as show in [Figure 140](#). Click the **Servers** tab, expand **Servers > Service Profiles > root**, and select a particular service profile and click **vNICs**. The vNIC names and MAC addresses are listed in the right pane of the window.

Figure 140 vNIC Names and Their MAC Addresses are Shown in the vNICs Area



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

VMware vCenter Server Deployment

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

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

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

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

The following steps provide high level configuration procedure to configure vCenter server:

1. Create the vCenter host VM

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

2. Install vCenter guest OS

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

3. Install vCenter server

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

4. Apply vSphere license keys

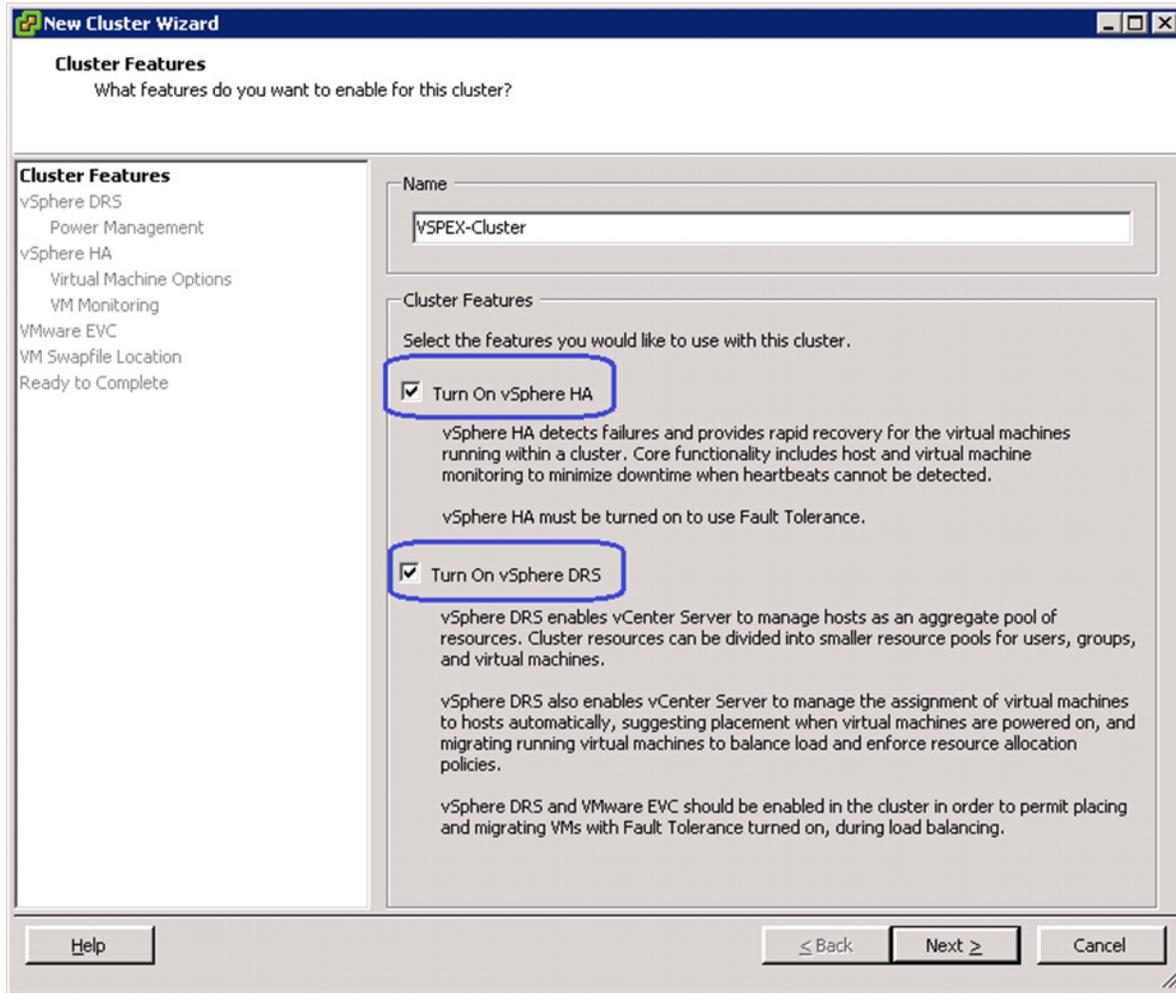
To perform license maintenance, log into the vCenter Server and select the Administration - Licensing menu from the vSphere client. Use the vCenter License console to enter the license keys for the ESXi hosts. After this, they can be applied to the ESXi hosts as they are imported into vCenter.

Configuring Cluster, HA and DRS on VMware vCenter

Follow these steps to configure cluster, HA, and DRS on vCenter:

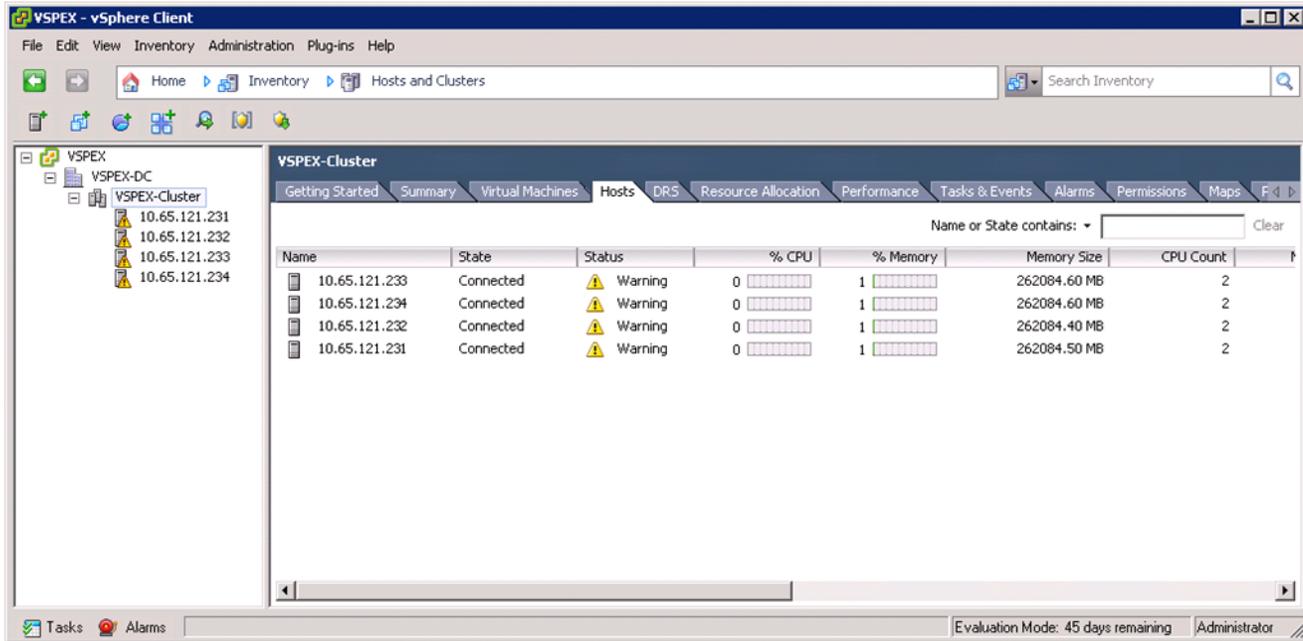
1. Log into VMware ESXi Host using VMware vSphere Client.
2. Create a vCenter Datacenter.
3. Create a new management cluster with DRS and HA enabled.
 - a. Right-click on the cluster and, in the corresponding context menu, click **Edit Settings**.
 - b. Check the check boxes Turn On vSphere HA and Turn On vSphere DRS, as shown in [Figure 141](#).
 - c. Click **OK**, to save changes.

Figure 141 Configuring HA and DRS on Cluster



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

Figure 142 Adding ESXi Hosts to the Cluster



Virtual Networking Configuration

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

Figure 143 Network Adapter Showing vmnics in ESXi Server

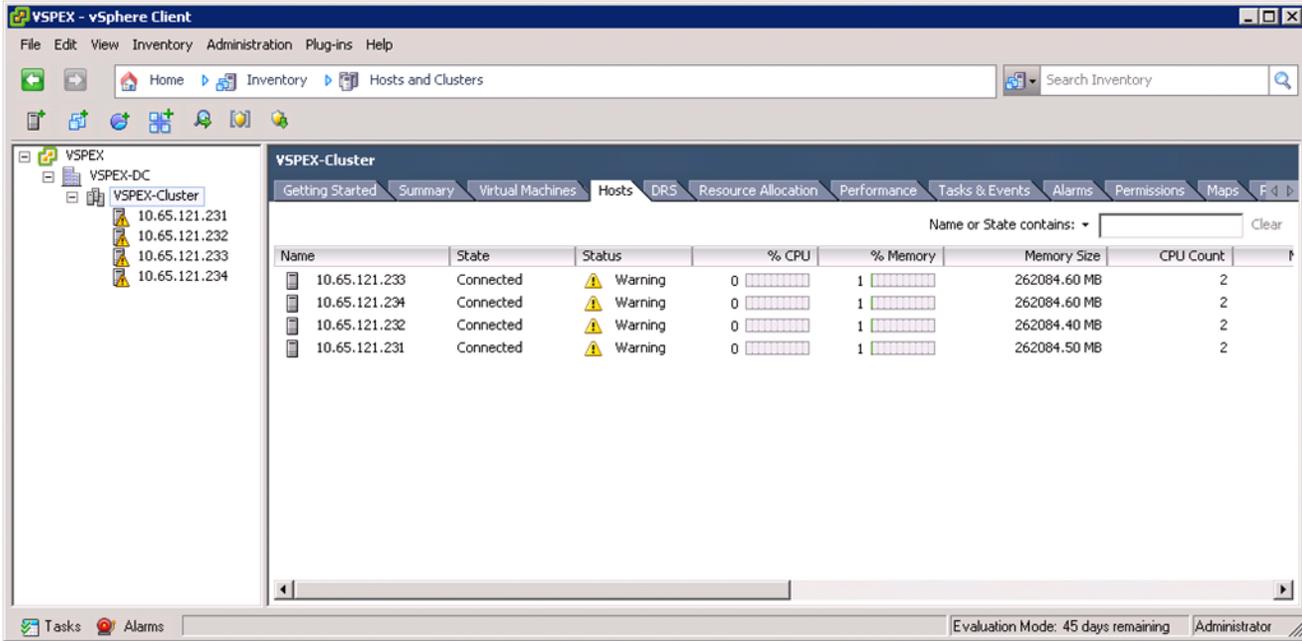


Table 10 shows UCS Manager service profile and vSphere vmnic per ESXi host basis:

Table 10 Service Profile vNIC and vSphere vmnic Relations

UCS Manager vNIC Names	vSphere VM NIC Names	MAC Address	Uplink Port-Profile*
System-A	vmnic0		system-uplink
System-B	vmnic1		system-uplink
Storage-A*	vmnic2		storage-uplink
Storage-B*	vmnic3		storage-uplink
Data-A	vmnic4		data-uplink
Data-B	vmnic5		data-uplink

*Applicable for the NFS-variant of the solution only.

We show casing two different approaches for the virtual networking layer of this architecture:

1. VMware vSphere native virtual switching in FC-variant of architecture
2. Cisco Nexus 1000v virtual switching in NFS-variant of architecture

**Note**

You can use either virtual switching strategy with any variant of architecture. This section focuses on the vSphere native virtual switching. See [“Install and configure Nexus 1000v” section on page 149](#).

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

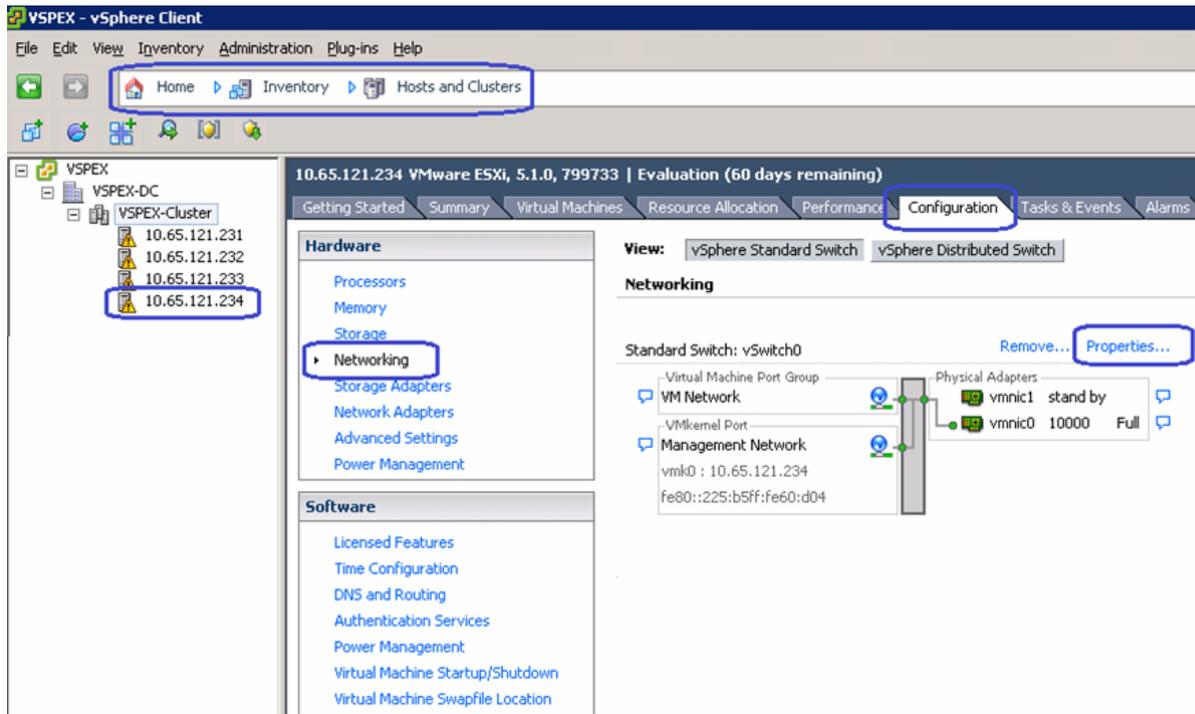
1. vSwitch0—Standard, default vSwitch for management and vMotion traffic.
2. vSwitch1—For VM data traffic.

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

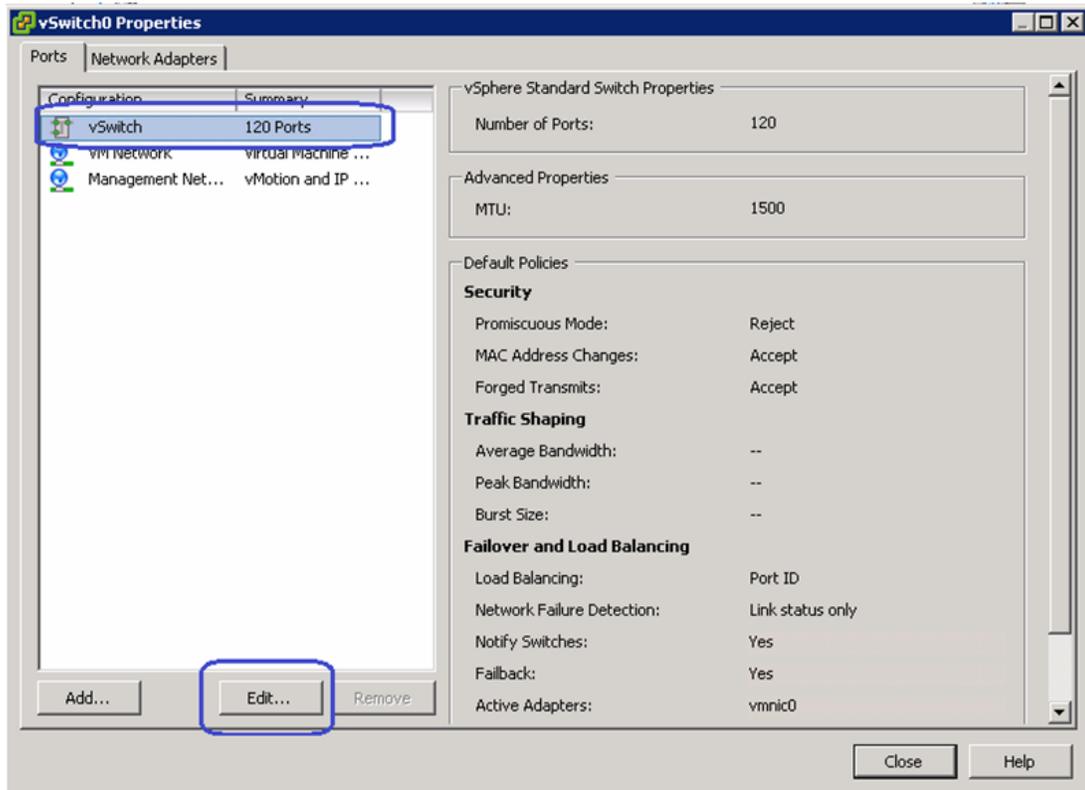
Follow these steps to configure the two vSwitches:

1. In the vSphere client, choose **Home > Inventory > Hosts and Clusters**. In the Hosts and Clusters window, click the **Configuration** tab on the right pane of the window. Click **Networking** in the Hardware area. Click **Properties** to see the details of vSwitch0.

Figure 144 Viewing Details of vSwitch0

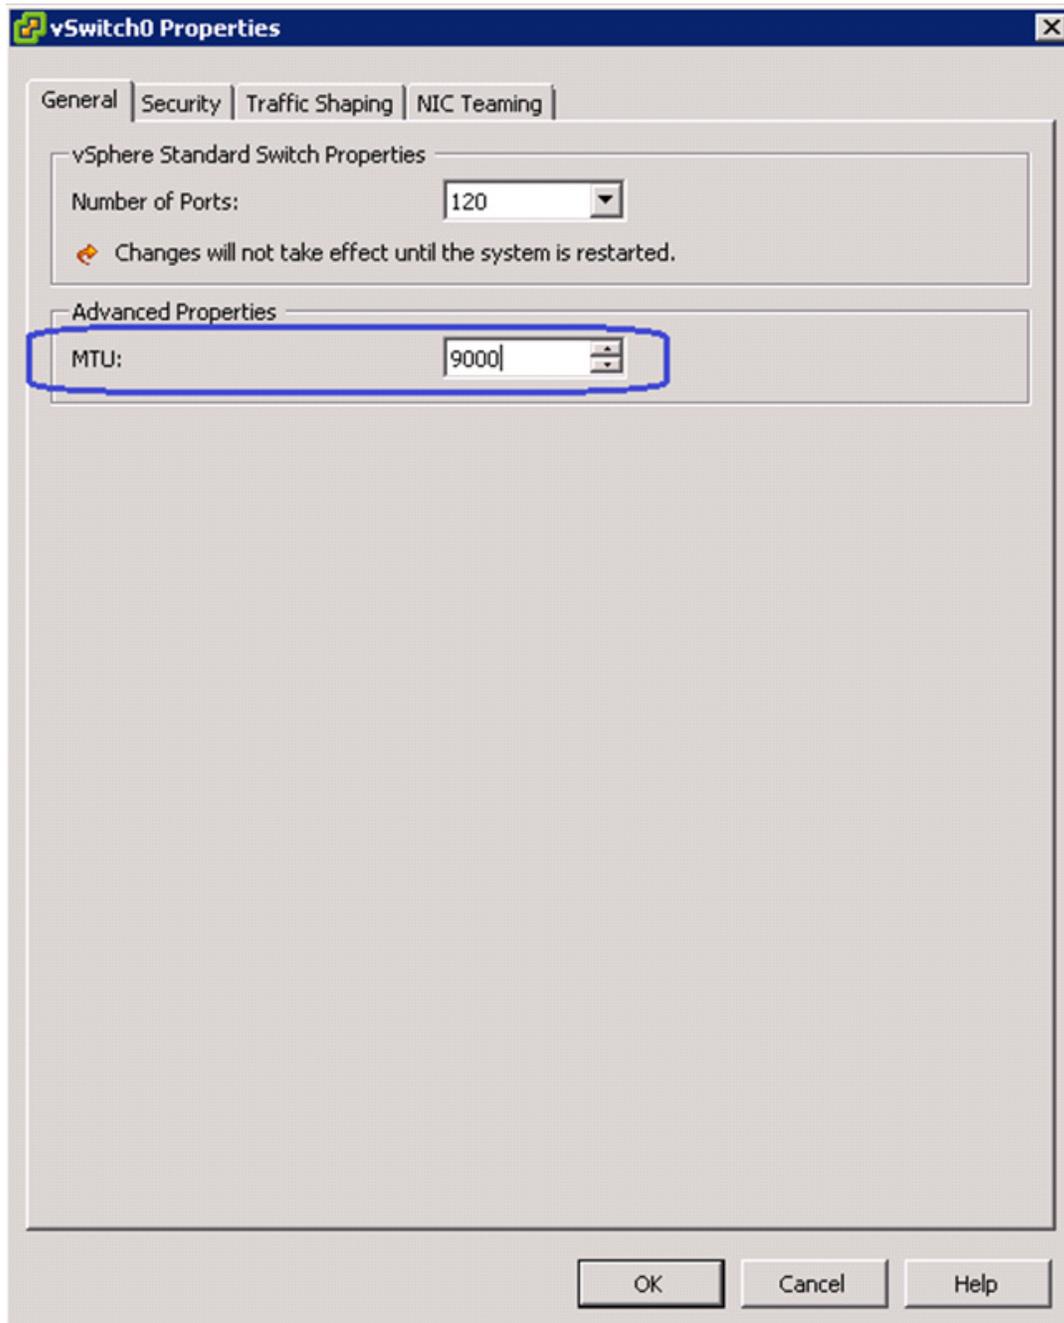


2. Select vSwitch in the vSwitch0 Properties window, and click **Edit**.

Figure 145 *Editing vSwitch0 Properties*

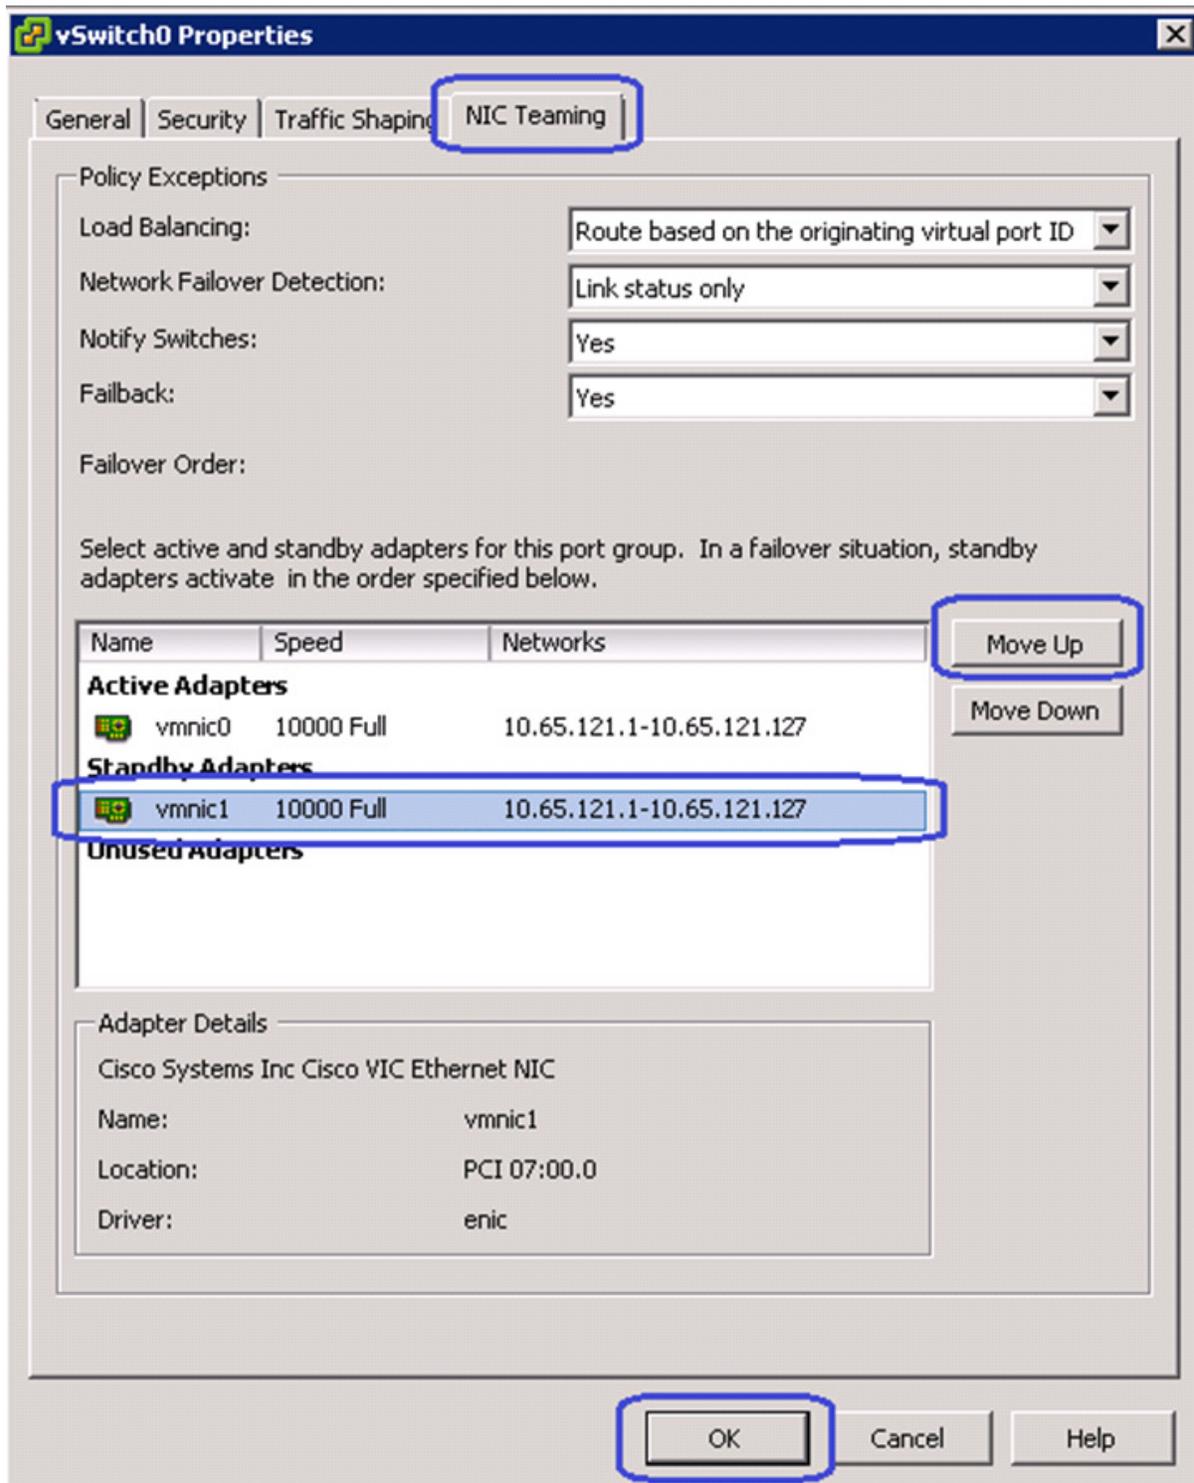
3. Click the **General** tab of the vSwitch0 Properties window, change the MTU in the Advanced Properties area to 9000.

Figure 146 Setting Jumbo MTU



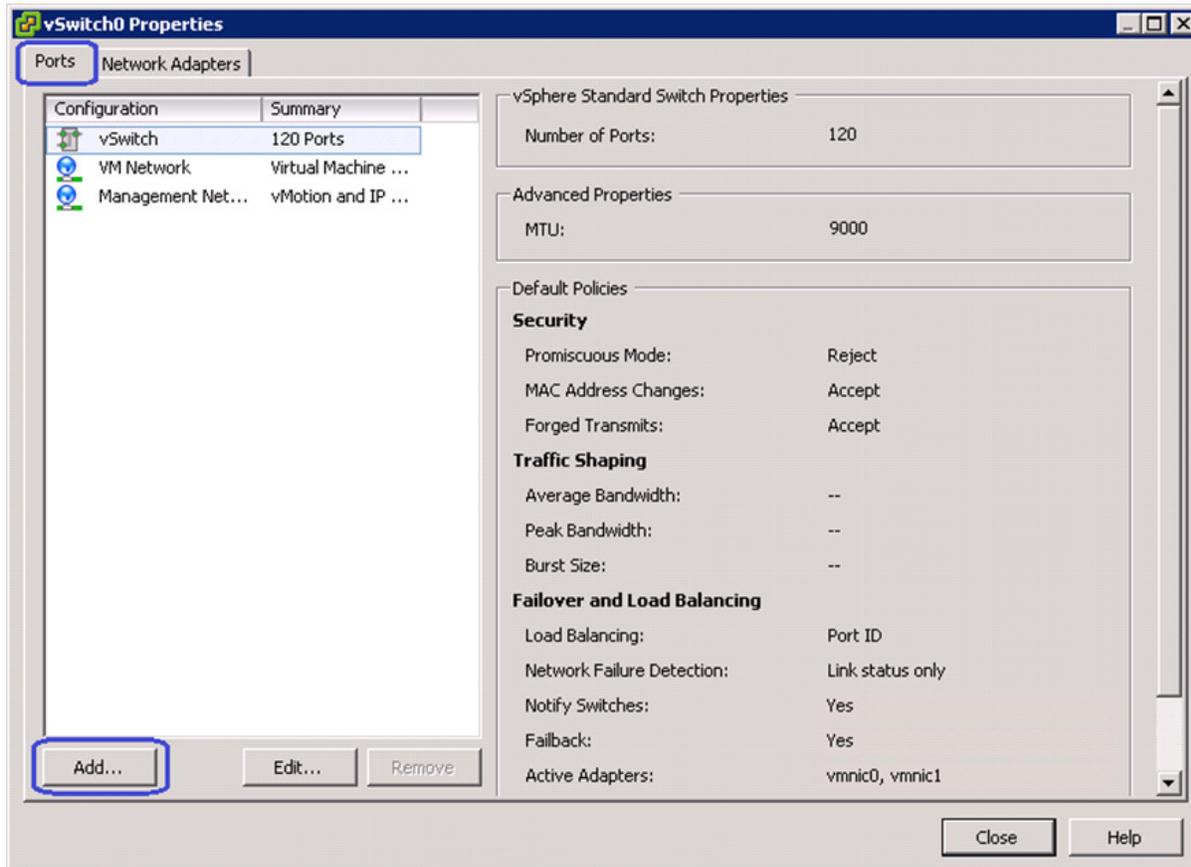
4. Click the **NIC Teaming** tab in the vSwitch0 Properties window. Select the adapter under Standby Adapters and click **Move up** to get it under Active Adapters, and click **OK**.

Figure 147 Moving Standby Adapter to Active Adapter



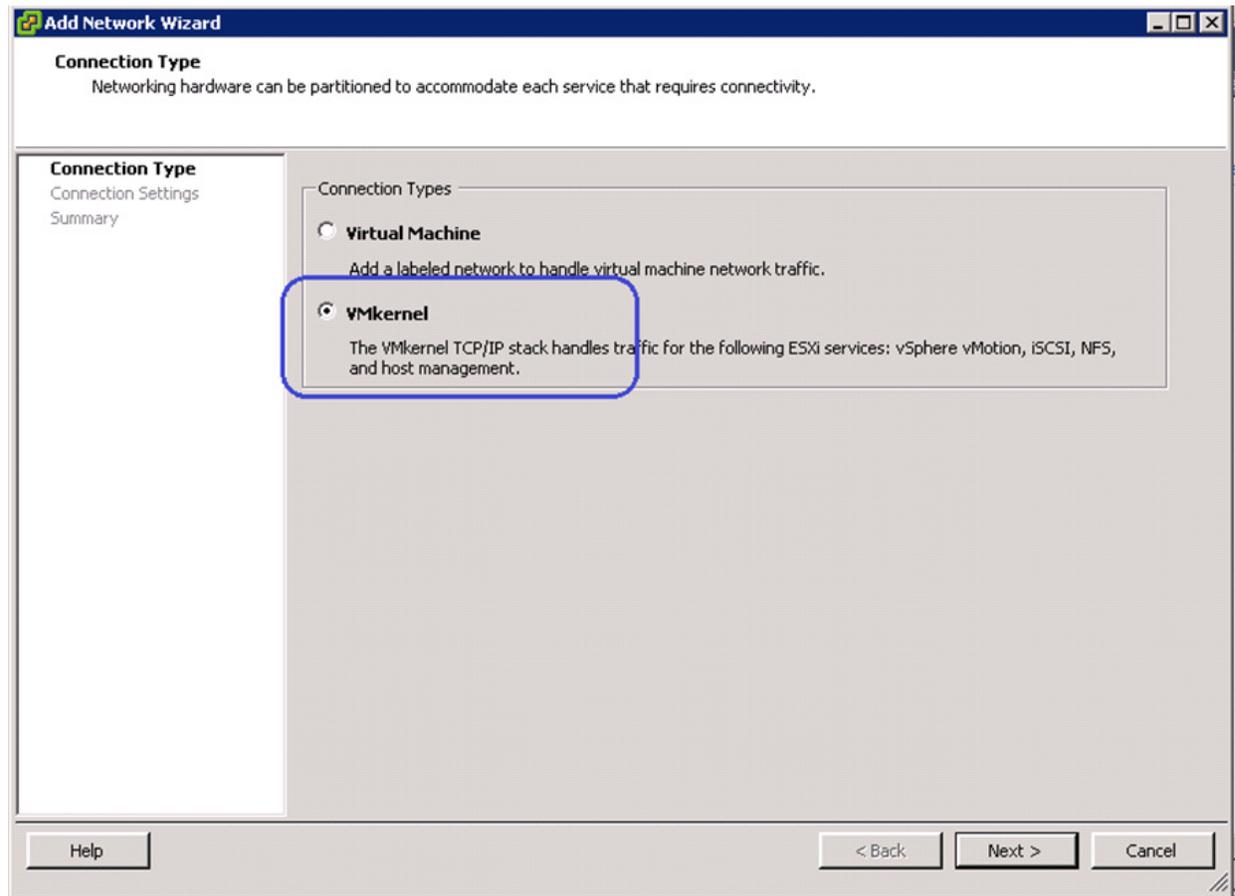
- In the vSwitch0 configuration window, click the **Ports** tab, and click **Add**.

Figure 148 Adding Ports to vSwitch0



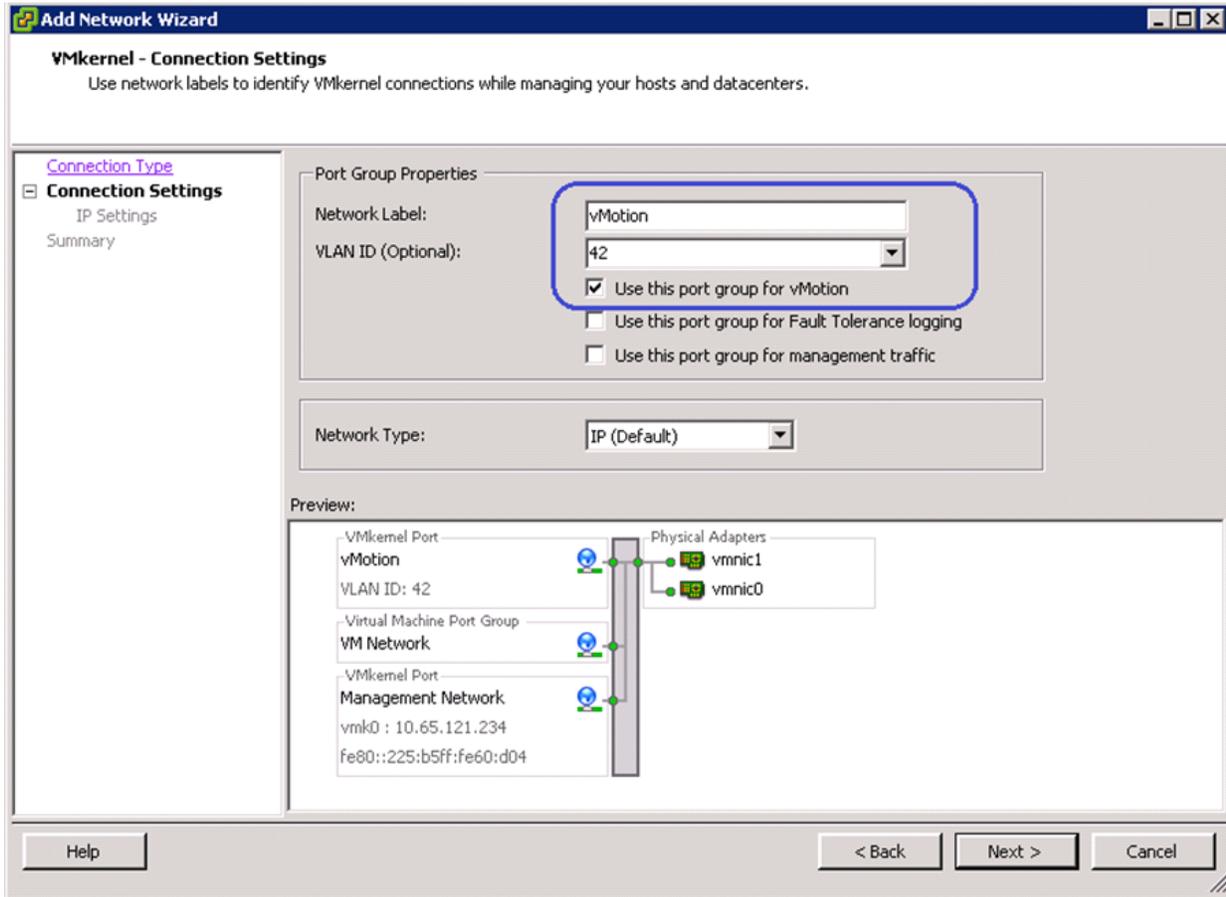
6. Click the **VMKernel** radio button in the Connection Types area and click **Next** in the Add Network Wizard window.

Figure 149 Specifying Connection Type



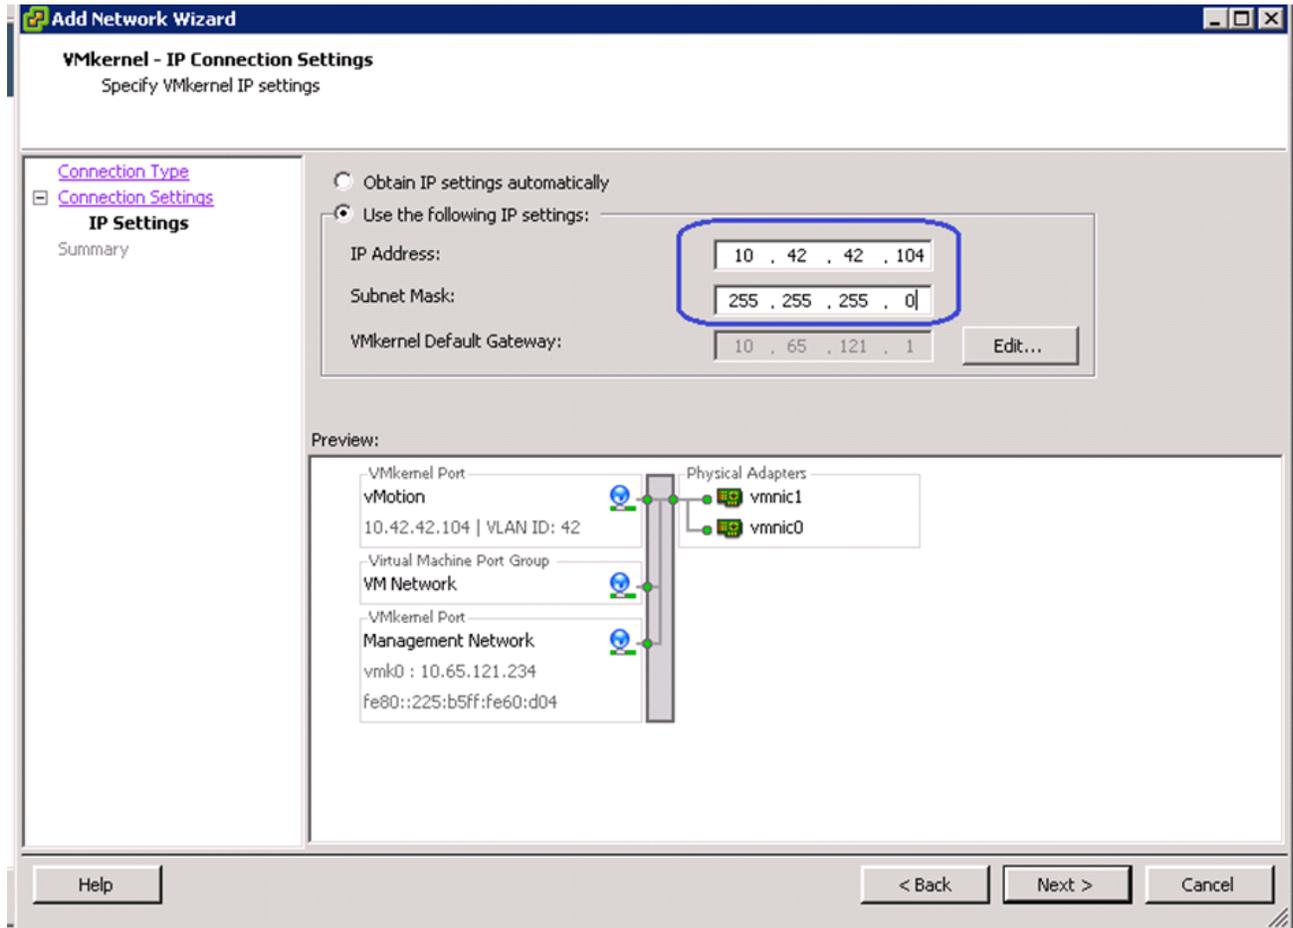
7. Enter vMotion in the Network Label field. Choose the VLAN ID from the drop-down list. The standard vSwitch0 carries both management and vMotion VLANs. Management traffic leaves vSwitch0 untagged using the native VLAN of the vNIC. The vMotion traffic must be tagged with the appropriate VLAN ID. Check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

Figure 150 Specifying Port Group Properties for Setting VMkernel Connection



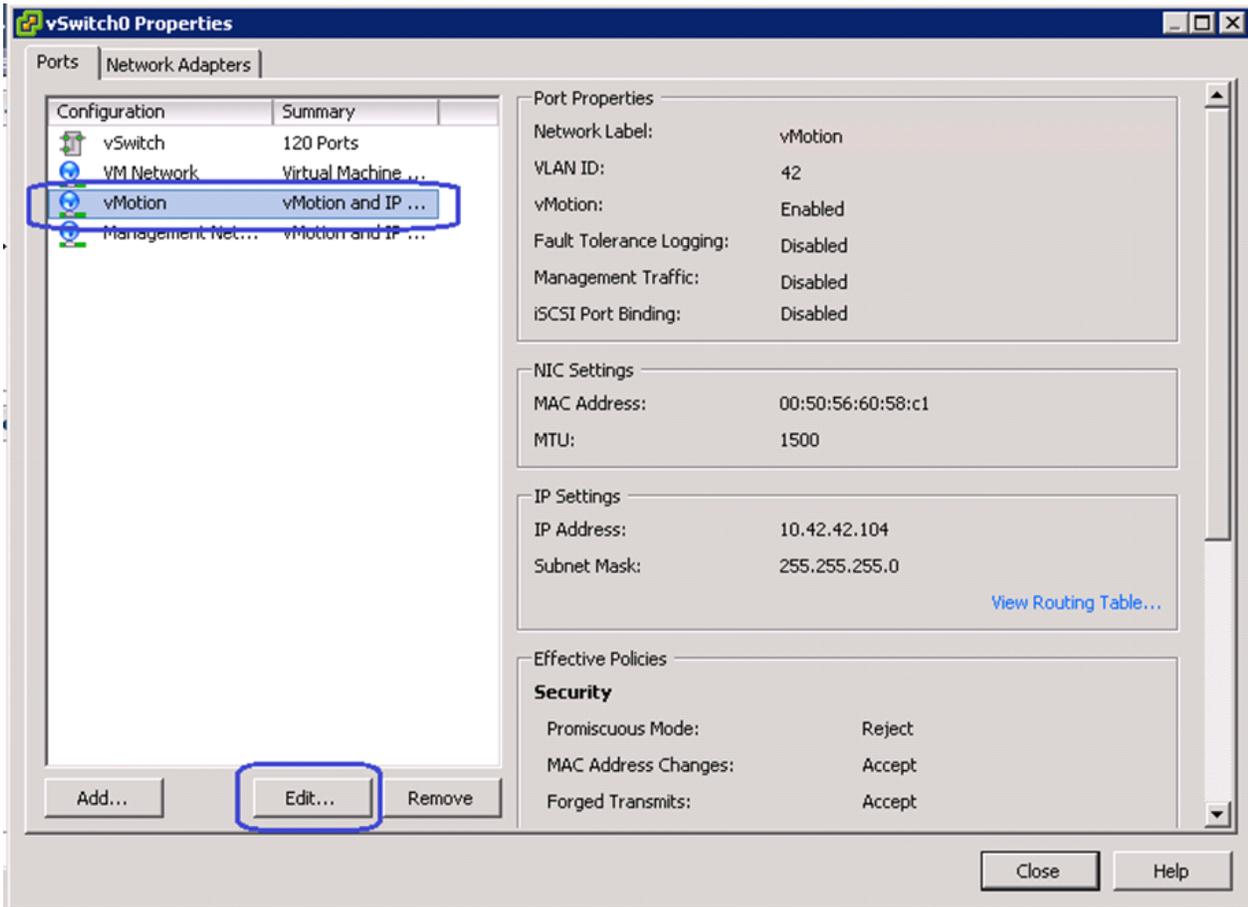
- Configure IP address and subnet mask for the vmkernel interface. Click **Next** and deploy the vmkernel.

Figure 151 Specifying IP Address and Subnet Mask for Setting VMkernel Connection



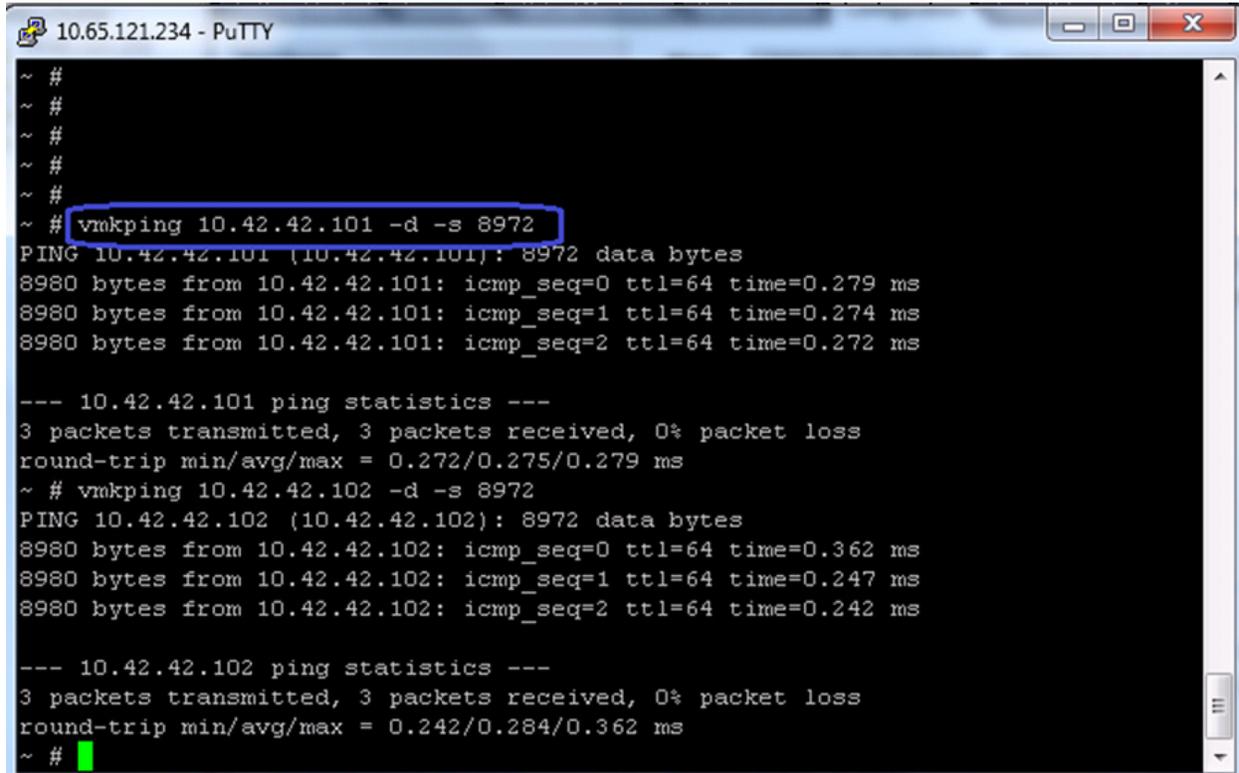
- In the vSwitch0 properties window, select the newly created vMotion port group and click **Edit**.

Figure 152 Editing vMotion vSwitch0 to Set Jumbo MTU



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

Figure 153 Pinging the VMKernel Port with Jumbo MTU



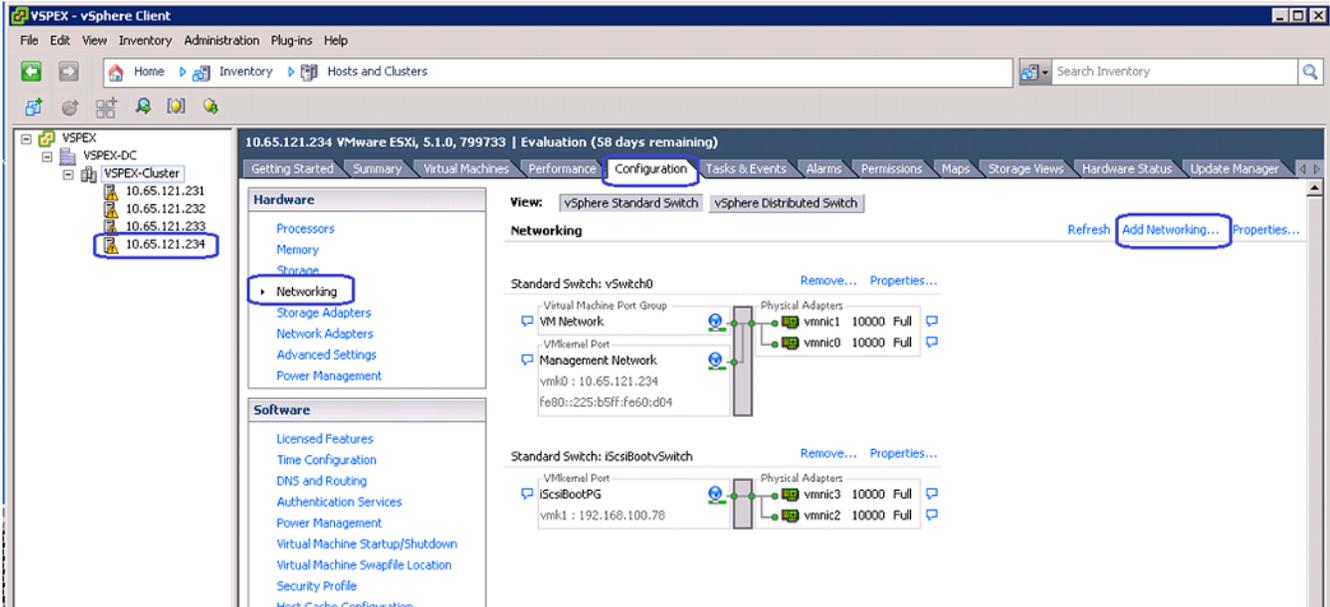
```
10.65.121.234 - PuTTY
~ #
~ #
~ #
~ #
~ #
~ # vmkping 10.42.42.101 -d -s 8972
PING 10.42.42.101 (10.42.42.101): 8972 data bytes
8980 bytes from 10.42.42.101: icmp_seq=0 ttl=64 time=0.279 ms
8980 bytes from 10.42.42.101: icmp_seq=1 ttl=64 time=0.274 ms
8980 bytes from 10.42.42.101: icmp_seq=2 ttl=64 time=0.272 ms

--- 10.42.42.101 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.272/0.275/0.279 ms
~ # vmkping 10.42.42.102 -d -s 8972
PING 10.42.42.102 (10.42.42.102): 8972 data bytes
8980 bytes from 10.42.42.102: icmp_seq=0 ttl=64 time=0.362 ms
8980 bytes from 10.42.42.102: icmp_seq=1 ttl=64 time=0.247 ms
8980 bytes from 10.42.42.102: icmp_seq=2 ttl=64 time=0.242 ms

--- 10.42.42.102 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.242/0.284/0.362 ms
~ #
```

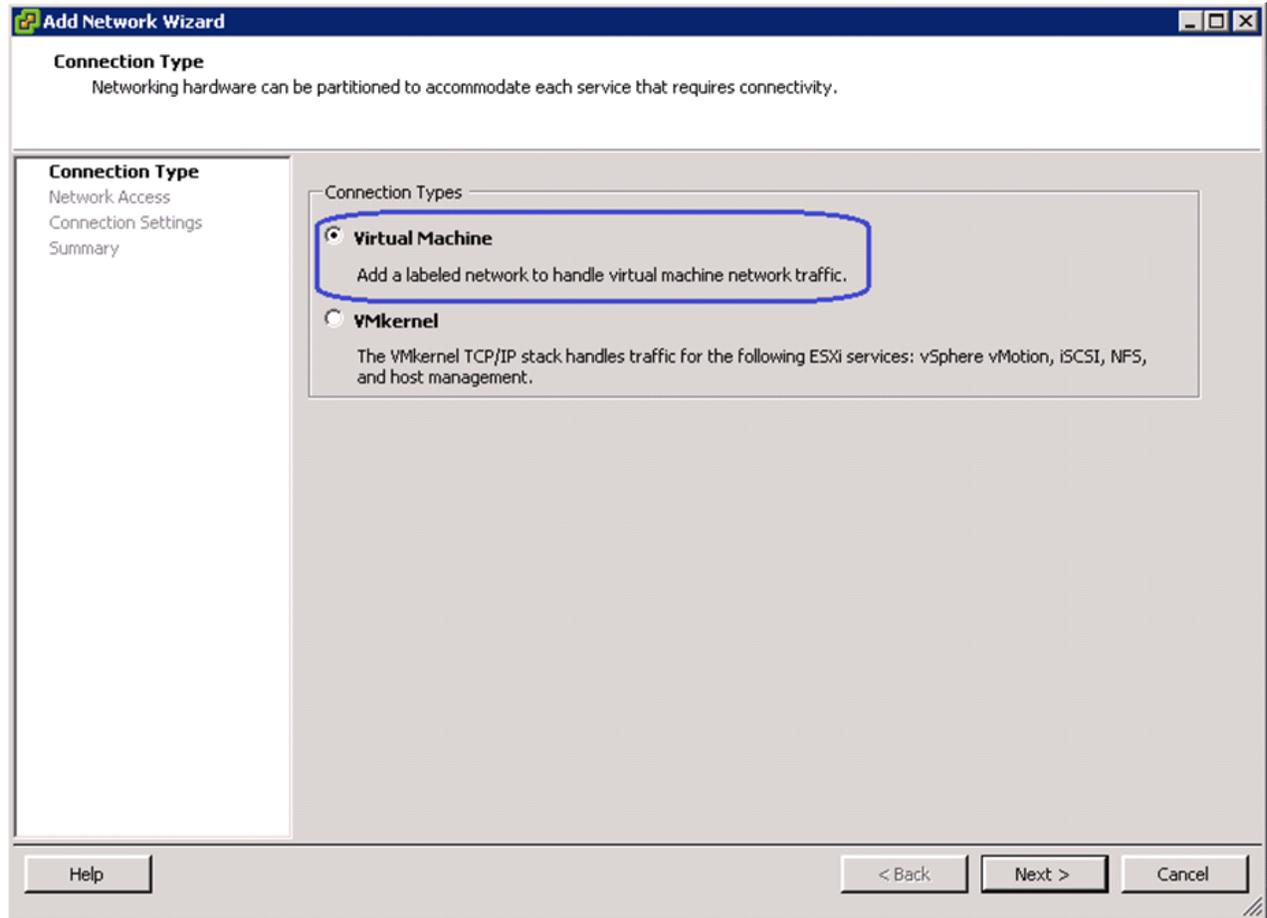
12. In the vCenter GUI, choose **Home > Inventory > Hosts and Clusters**. In the Hosts and Clusters window, click the **Configuration** tab on the right pane of the window. Click **Networking** in the Hardware area. Click **Add Networking**.

Figure 154 Add Networking in VMware vSphere Client



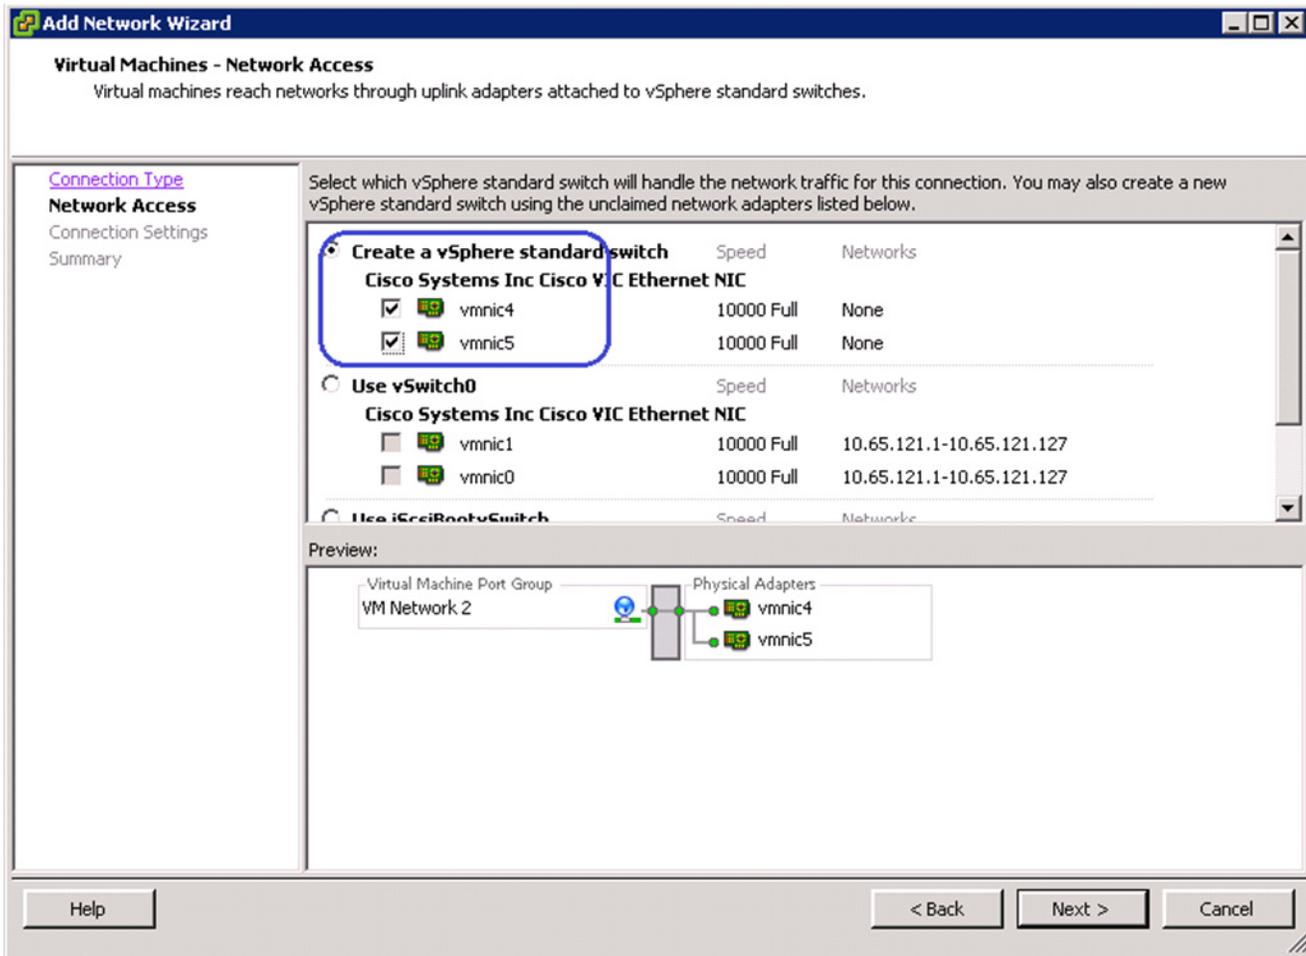
13. Click the **Virtual Machine** radio button in the Add Networking Wizard, click **Next**.

Figure 155 **Specifying the Connection Type**



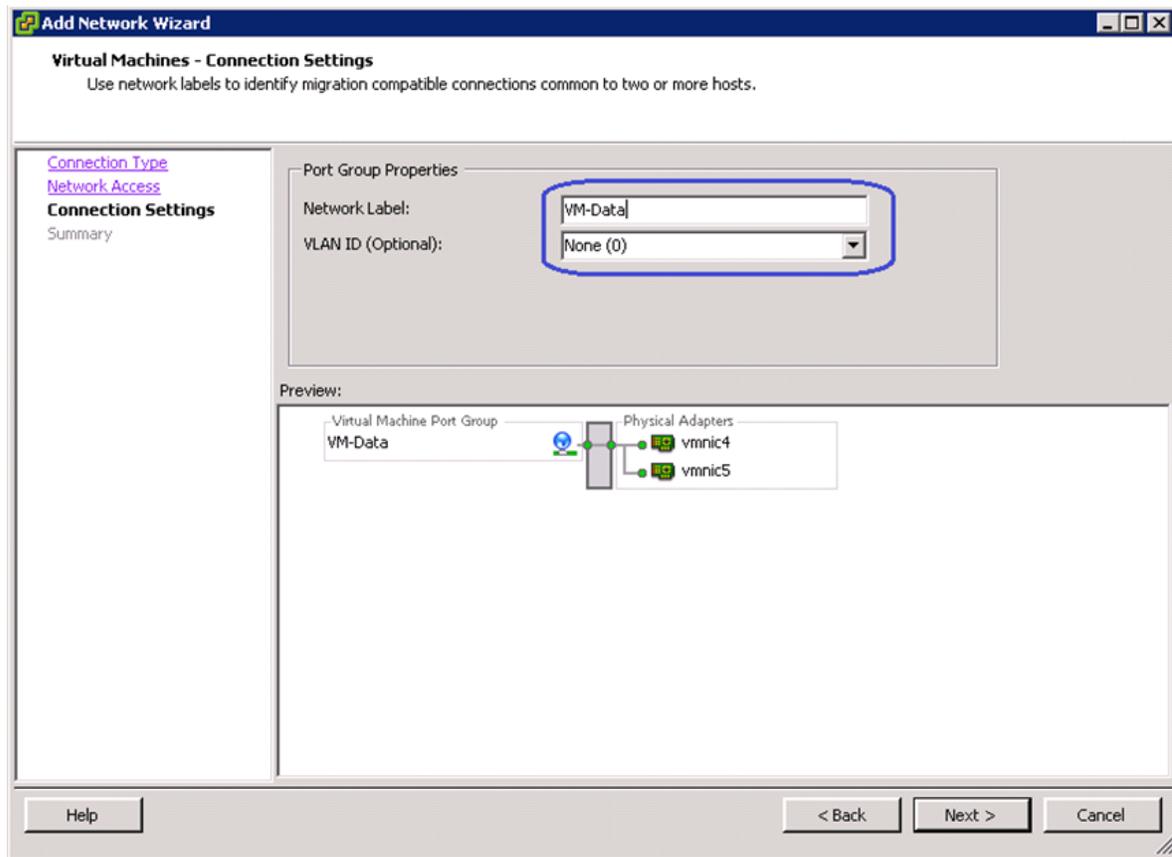
14. Select the two vmnics corresponding to the VM-Data vNICs by checking the check boxes and click **Next**.

Figure 156 Selecting the vSphere Standard Switch for Handling Network Traffic



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

Figure 157 **Specifying Port Group Properties**



16. Repeat steps 15 to 18 for all the ESXi hosts in the cluster.

Install and configure Nexus 1000v

Cisco Nexus 1000v is a Cisco's NX-OS based virtual switch that replaces the native vSwitch in the VMware ESXi hosts by a virtual Distributed Switch (vDS). The control plane of Nexus 1000v switch is installed in a VMware Virtual Machine, and is known as Virtual Switching Module (VSM). VSM virtual machine (VSM VM) is available as a VMware OVF template. Following are the major steps to deploy Nexus 1000v architecture:

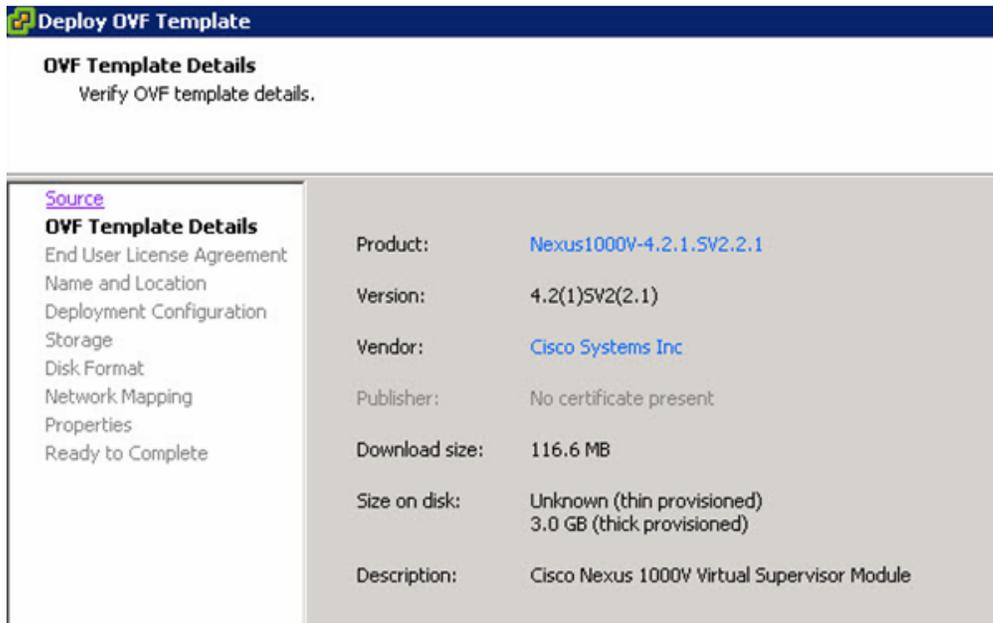
1. Install Nexus 1000v VSM VM
2. Connect Nexus 1000v VSM to VMware vCenter
3. Configure port-profiles in VSM and migrate vCenter networking to vDS

Install Nexus 1000v VSM VM

As mentioned before, the Nexus 1000v VSM VM installation media is available as VMware virtual machine OVF template. The VSM VM must be deployed on the infrastructure network, and not on one of the VSPEX ESXi servers. Follow these steps to install VSM VM:

1. From the **Hosts and Cluster** tab in vCenter, choose the infrastructure ESX/ESXi host and click **File > Deploy new Virtual Machine** through OVF template. Choose Nexus 1000v VSM OVF, and click **Next**.

Figure 158 Deploy OVF Template - Verifying OVF Template Details



2. Select the datacenter where you want to install the VSM in next page

Figure 159 Deploy OVF Template - Specifying Inventory Location

Deploy OVF Template

Name and Location
Specify a name and location for the deployed template

[Source](#)
[OVF Template Details](#)
[End User License Agreement](#)

Name and Location
Deployment Configuration

- Host / Cluster
- Resource Pool
- Disk Format
- Properties
- Ready to Complete

Name:
Nexus1000V-1780-Primary
The name can contain up to 80 characters and it must be unique within the inventory folder.

Inventory Location:

- VSPEX-Infra1.vspex.com
 - VSPEX-FC
 - VSPEX-Infra**
 - VSPEX-NFS

- For Configuration field, choose **Manually Configure Nexus 1000v** from the drop-down list.

Figure 160 Deploy OVF Template - Choosing Deployment Configuration

Deploy OVF Template

Deployment Configuration
Select a deployment configuration.

[Source](#)
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)

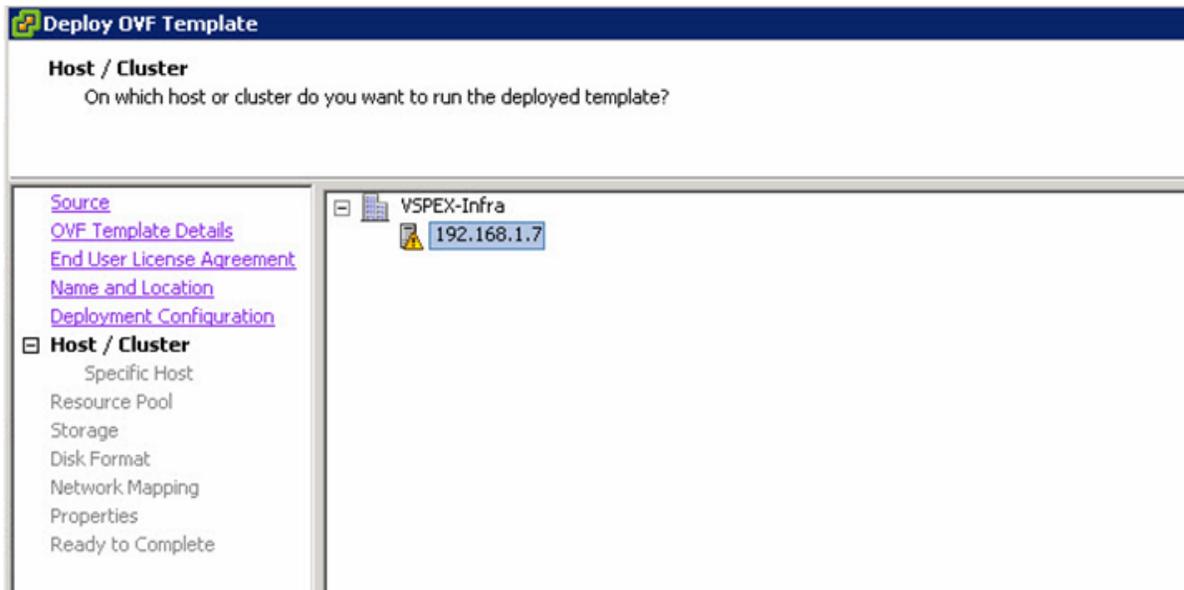
Deployment Configuration
Storage
Disk Format
Network Mapping
Properties
Ready to Complete

Configuration:
Manually Configure Nexus 1000V

Use this deployment option to manually configure the Nexus 1000V VSM without the use of the installer application or OVF Properties. If this option is selected, please ignore the properties section ahead.

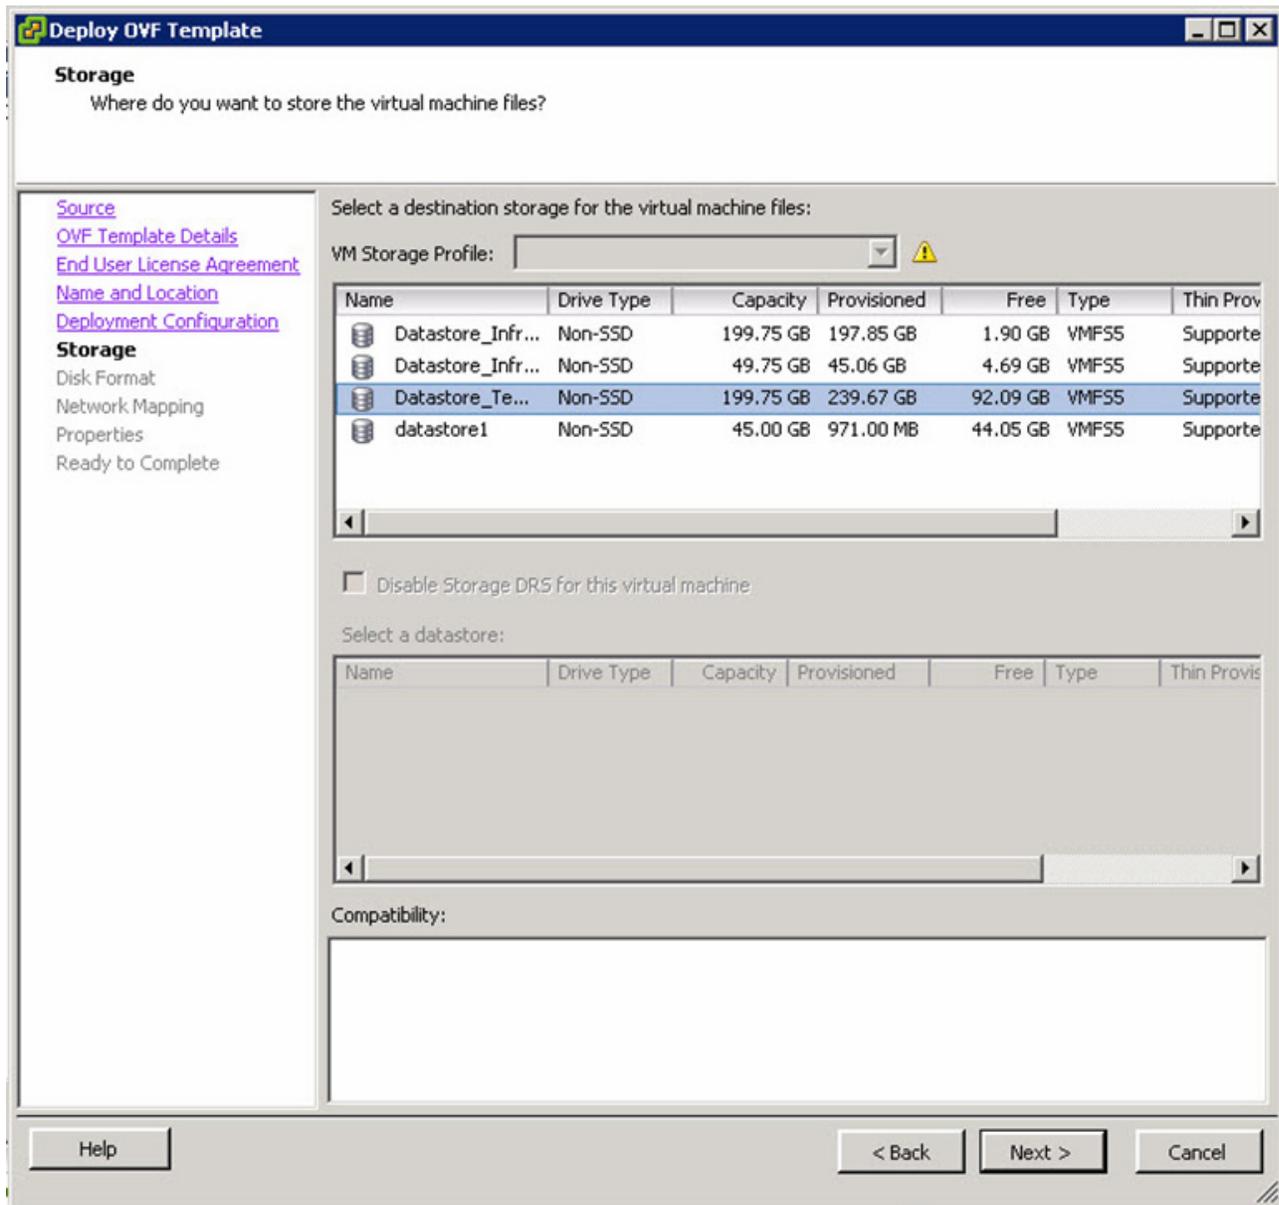
- Choose the host on which you want to install the N1k VSM VM.

Figure 161 Deploy OVF Template - Choosing the Host



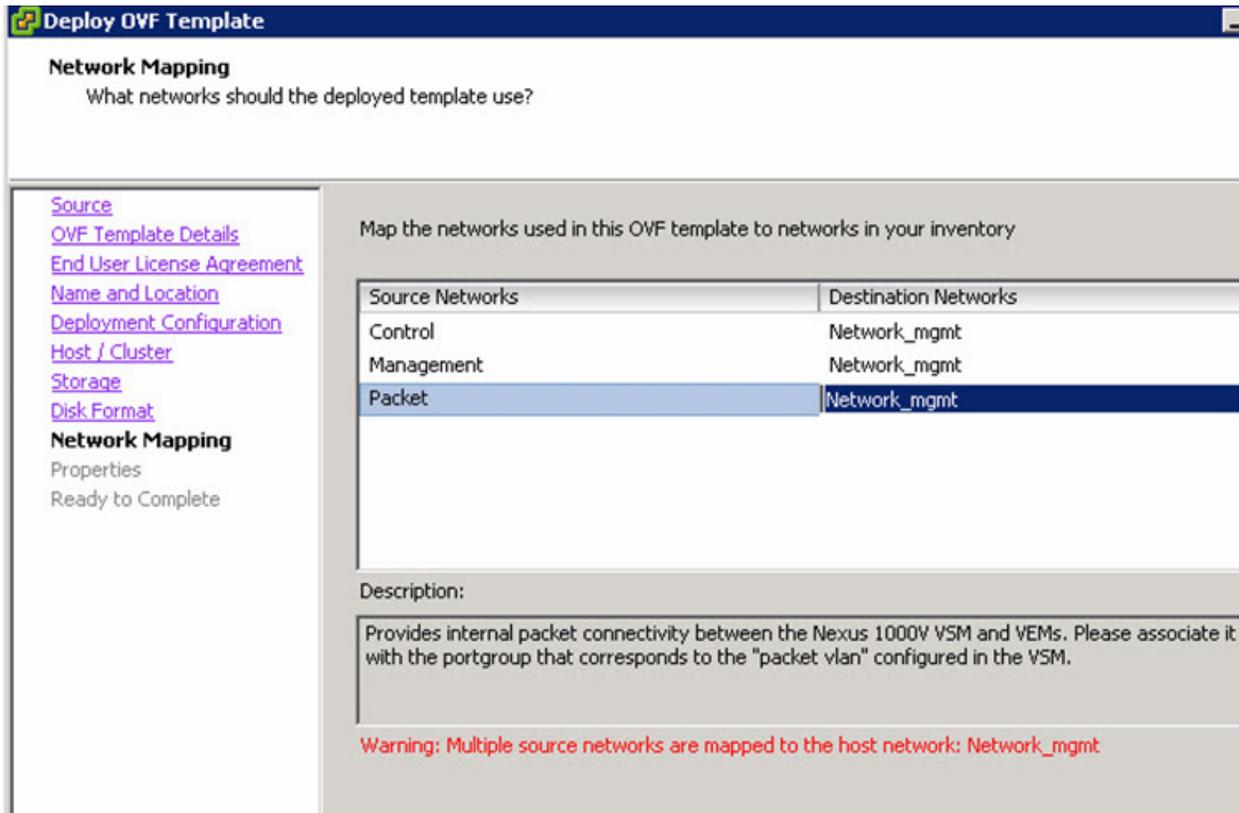
5. Choose the datastore where the VM should be deployed. Click **Next**.

Figure 162 Deploy OVF Template - Choosing the Datastore



- Choose the destination Network for mapping the Source Network's Control, Management, and Packet VLANs to the management (infra) VLAN. Click **Next**.

Figure 163 Deploy OVF Template - Choosing the Destination Network for Mapping



7. In the Properties window, configure Domain Id (a unique number across multiple N1k VSMs, if there are more than one), Administrator User Password, Management IP address and Management Subnet Mask as shown in Figure 164. Click **Next**.

Figure 164 Deploy OVF Template - Entering Details for the Deployment

🏠
Deploy OVF Template

Properties
Customize the software solution for this deployment.

[Source](#)

[OVF Template Details](#)

[End User License Agreement](#)

[Name and Location](#)

[Deployment Configuration](#)

[Storage](#)

[Disk Format](#)

[Network Mapping](#)

Properties

[Ready to Complete](#)

a. VSM Domain ID

DomainId
Enter the Domain Id (1-4095).

b. Nexus 1000V Admin User Password

Password
Enter the password. Must contain at least one capital, one lowercase, one number.

Enter password

Confirm password

c. Management IP Address

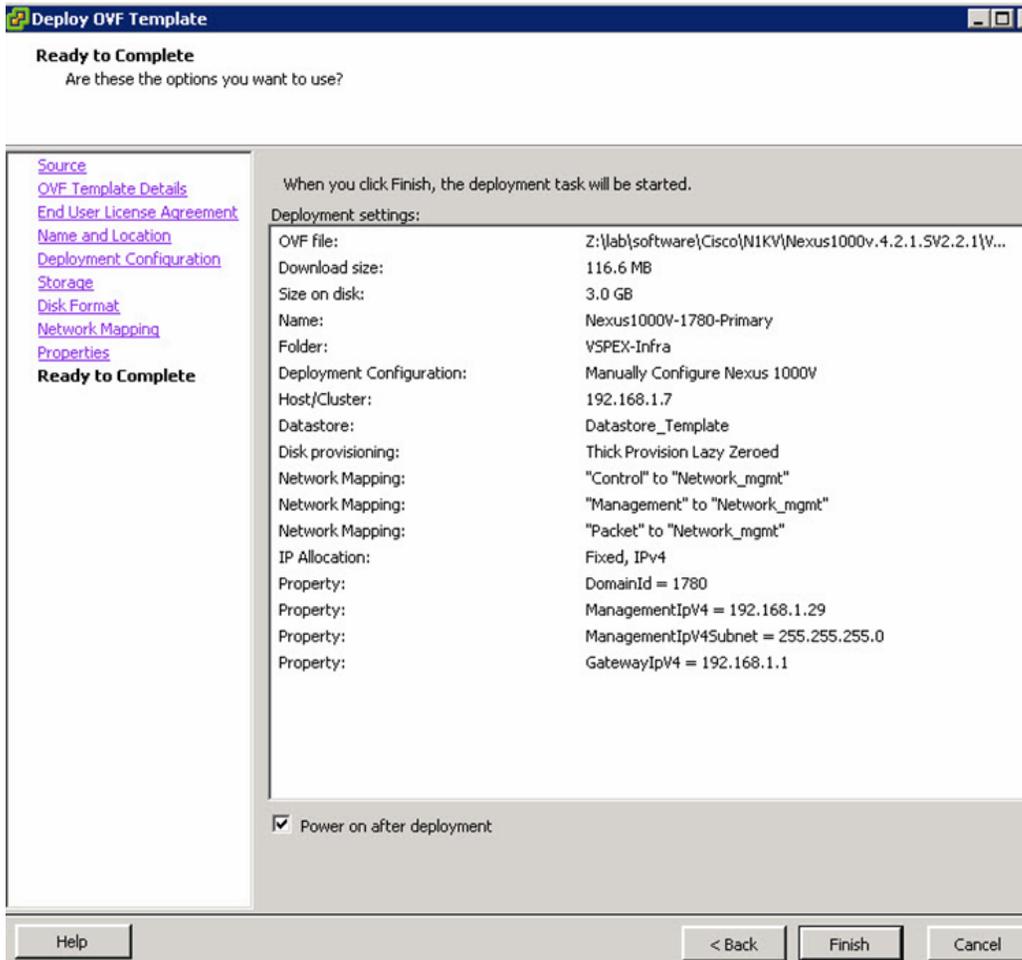
ManagementIPv4
Enter the VSM IP in the following form: 192.168.0.10

d. Management IP Subnet Mask

ManagementIPv4Subnet
Enter the Subnet Mask in the following form: 255.255.255.0

8. Verify the configuration, check the check box **Power on after deployment** and click **Finish** to complete the OVF deployment of VSM virtual machine.

Figure 165 Deploy OVF Template - Verify the Options



9. When VSM VM is powering up for the first time, click **Console** of the virtual machine in the vCenter and perform initial configuration of the VSM VM. During the OVF deployment, basic configuration information was already provided. Verify the details and press **no** if you don't want to change anything.

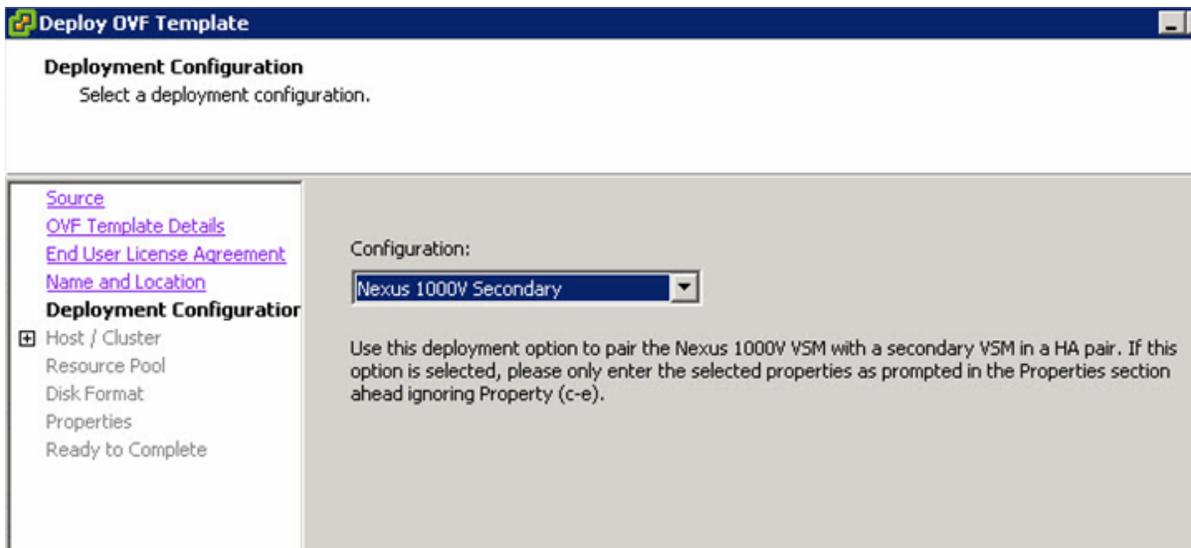
Figure 166 Verifying the Details in the Virtual Machine Console in vCenter

```
The following configuration will be applied:
switchname USM-1020
interface mgmt0
ip address 192.168.1.28 255.255.255.0
no shutdown
vrf context management
ip route 0.0.0.0/0 192.168.1.1
ssh key rsa 2048 force
ssh server enable
feature http-server
svs-domain
no control vlan
no packet vlan
svs mode L3 interface mgmt0
domain id 1020

Would you like to edit the configuration? (yes/no) [n]: no_
```

10. Once initial configuration is complete, you would see the exec mode CLI and you can SSH to the VM.
11. It is highly recommended that you deploy two VMs in HA mode for VSM. To deploy secondary VM, repeat steps 1 and 2. On step 3, for the configuration field, choose **Nexus 1000v Secondary** from the drop-down list.

Figure 167 Deploy OVF Template - Choosing Deployment Configuration for Secondary VM



12. Repeat steps 4, 5 and 6 from the original VSM VM deployment. In the Properties window, give the same domain ID as the primary VSM VM domain ID and password. No need to provide IP address/subnet mask, as secondary VM will take over the operations if primary fails.

Figure 168 Deploy OVF Template - Entering Details for the Secondary VM Deployment

The screenshot shows a web-based wizard titled "Deploy OVF Template". On the left, there is a navigation pane with links for "Source", "OVF Template Details", "End User License Agreement", "Name and Location", "Deployment Configuration", "Host / Cluster", "Storage", "Disk Format", "Network Mapping", "Properties", and "Ready to Complete". The "Properties" section is currently active, displaying the instruction: "Customize the software solution for this deployment." The main content area is divided into four sections:

- a. VSM Domain ID**: Labeled "DomainId", it asks to "Enter the Domain Id (1-4095)." and has a text input field containing "1780".
- b. Nexus 1000V Admin User Password**: Labeled "Password", it asks to "Enter the password. Must contain at least one capital, one lowercase, one number." It includes two input fields: "Enter password" and "Confirm password", both containing "*****".
- c. Management IP Address**: Labeled "ManagementIpV4", it asks to "Enter the VSM IP in the following form: 192.168.0.10" and has a text input field containing "0 . 0 . 0 . 0".
- d. Management IP Subnet Mask**: Labeled "ManagementIpV4Subnet", it asks to "Enter the Subnet Mask in the following form: 255.255.255.0" and has a text input field containing "0 . 0 . 0 . 0".

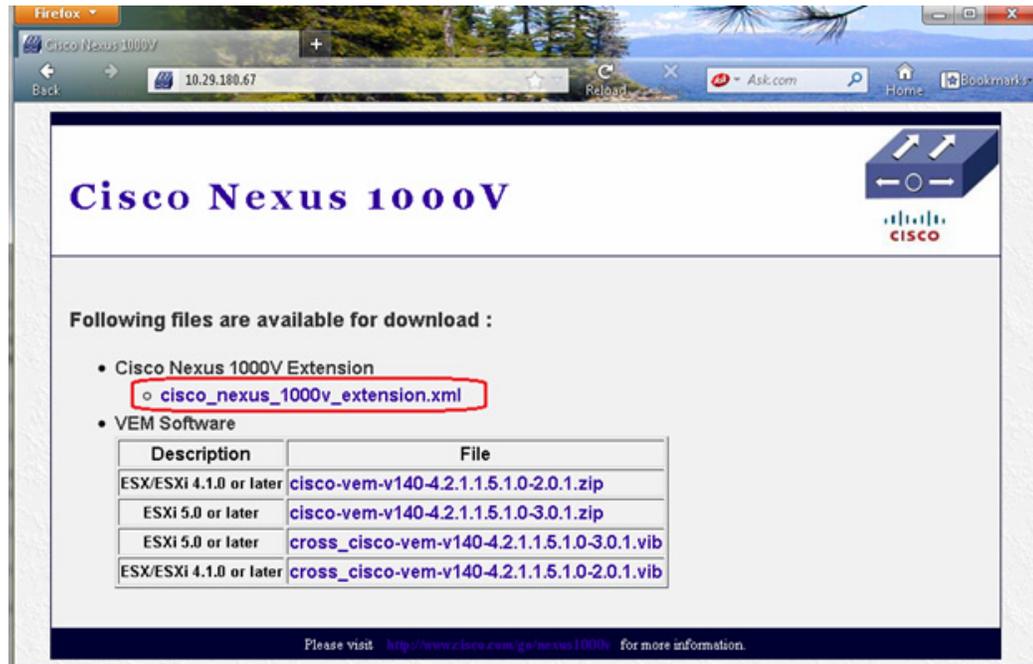
13. For the other configurations in the wizard, follow the steps same as primary VSM VM configuration.

Connecting VSM to vCenter

Once initial setup of VMS VM is completed, we need to add it as a plug-in/ extension in the vCenter:

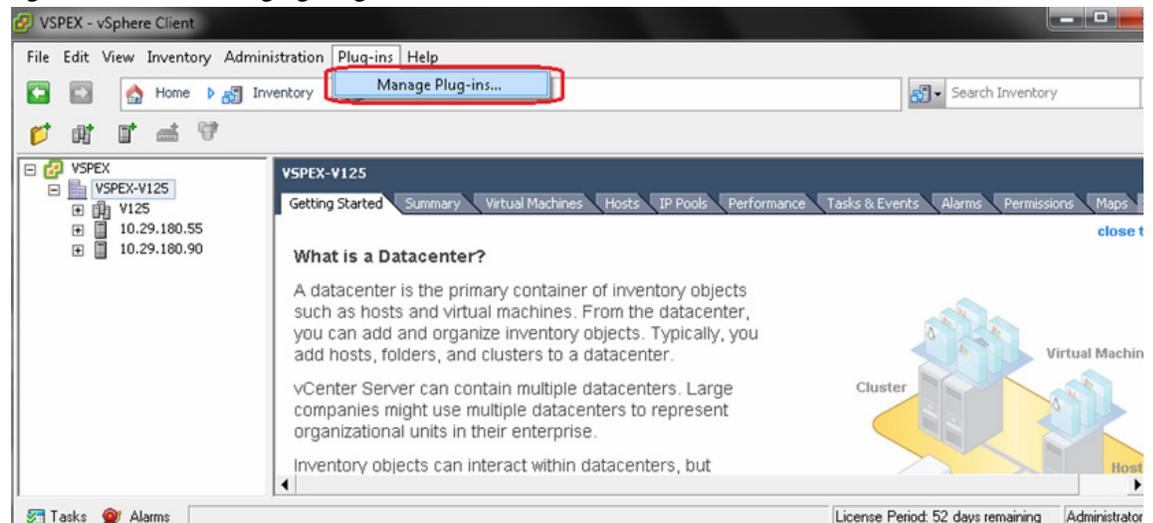
1. Using the browser, access management IP address of the VSM VM.
2. Right-click the `cisco_nexus_1000v_extension.xml` link and save it to a location on your local hard drive.

Figure 169 Save the Cisco Nexus Extension XML on your Local Drive



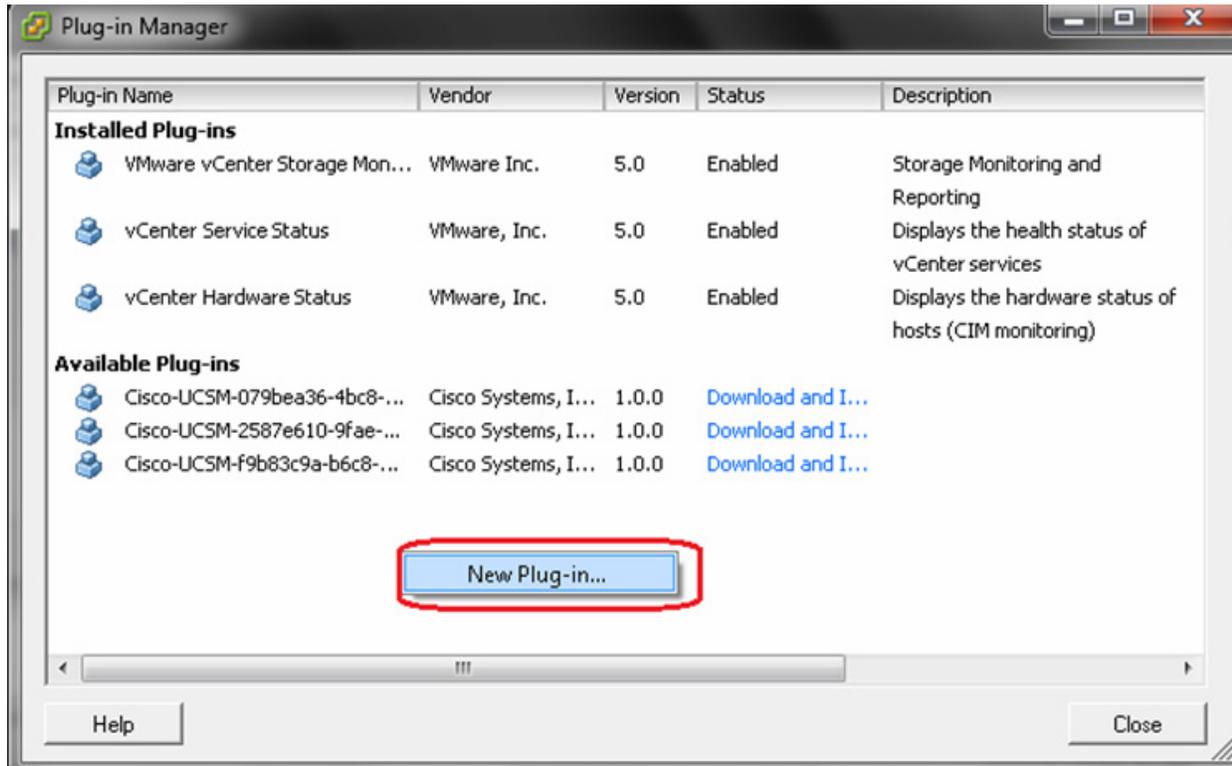
- From the vCenter, click **Plug-ins > Manage Plug-ins**.

Figure 170 Managing Plug-ins in vCenter



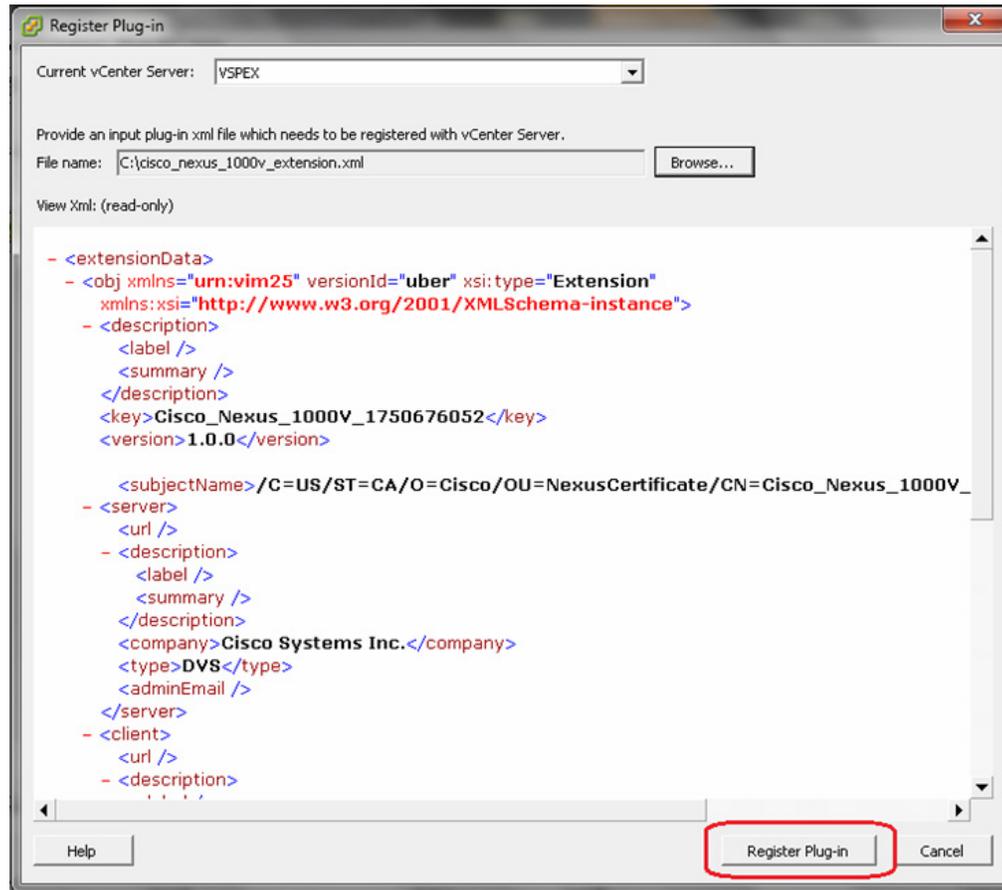
- Scroll to the bottom of Available Plug-ins, right-click in the empty space and choose **New Plug-in**.

Figure 171 Creating a New Plug-in



5. Click **Browse**, choose the cisco_nexus-1000v_extension.xml file you just downloaded.
6. Click **Register Plug-in**.

Figure 172 Registering the Downloaded XML File as the New Plug-in



7. If you receive a certificate warning, click **Ignore**.
8. Click **OK**. The Plug-in Manager page appears showing the plug-in that was just added.
9. Now configure the SVS connection to the vCenter as shown in [Figure 174](#).

Figure 173 *Configuring SVS Connection*

```

10.29.150.167 - PuTTY
login as: admin
Nexus 1000v Switch
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
N1k-VSM# conf t
Enter configuration commands, one per line. End with CNTL/Z.
N1k-VSM(config)# svcs connection vspex
N1k-VSM(config-svs-conn)# remote ip address 10.29.150.166
N1k-VSM(config-svs-conn)# vmware dvs datacenter-name vspex
N1k-VSM(config-svs-conn)# protocol vmware-vim
N1k-VSM(config-svs-conn)# connect
N1k-VSM(config-svs-conn)#

```

10. Validate the connection using **show svcs connections** and make sure that the operational status is **connected** and sync status is **Complete** as shown in [Figure 175](#).

Figure 174 *Running show svcs connections to Verify the Status*

```

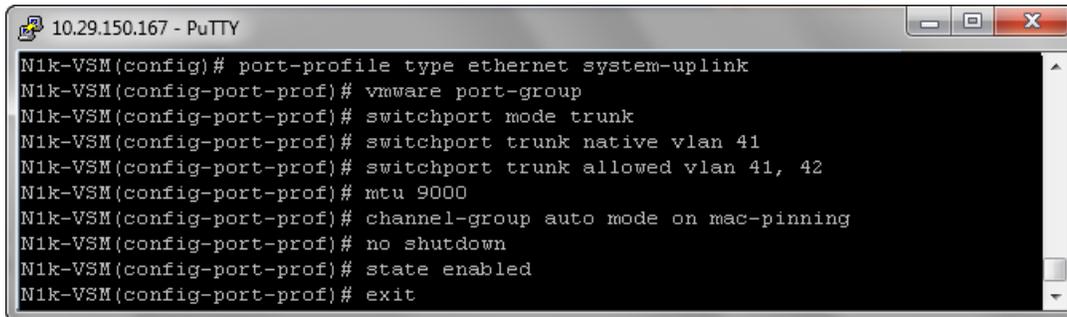
N1k-VSM(config)# svcs connection vspex
N1k-VSM(config-svs-conn)# remote ip address 10.29.150.166
N1k-VSM(config-svs-conn)# protocol vmware-vim
N1k-VSM(config-svs-conn)# vmware dvs datacenter-name vspex
N1k-VSM(config-svs-conn)# connect
N1k-VSM(config-svs-conn)# exit
N1k-VSM(config)# exit
N1k-VSM# show svcs connections

connection vspex:
  ip address: 10.29.150.166
  remote port: 80
  protocol: vmware-vim https
  certificate: default
  datacenter name: vspex
  admin:
  max-ports: 8192
  DVS uuid: 5c 03 21 50 ba e4 17 23-ad 02 05 ab 6e cb af 20
  config status: Enabled
  operational status: Connected
  sync status: Complete
  version: VMware vCenter Server 5.0.0 build-455964
  vc-uuid: 9D562753-5C78-4C08-973A-D67EA6023CB8
N1k-VSM#

```

11. You must create uplink port-profiles for the static vNICs of the service profile. Uplink port-profile is used to apply configuration on the uplink of the vDS, effectively the physical adapter of the ESXi server. Configure the system-uplink port-profile.

Figure 175 **Configuring System-Uplink Port-profile**



```

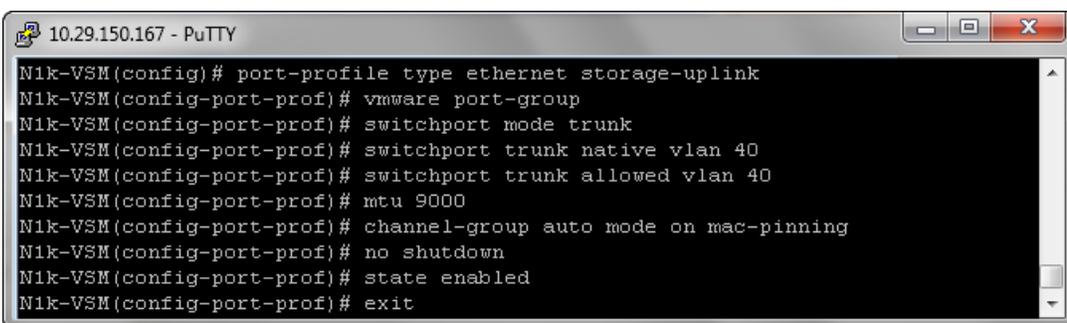
10.29.150.167 - PuTTY
N1k-VSM(config)# port-profile type ethernet system-uplink
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode trunk
N1k-VSM(config-port-prof)# switchport trunk native vlan 41
N1k-VSM(config-port-prof)# switchport trunk allowed vlan 41, 42
N1k-VSM(config-port-prof)# mtu 9000
N1k-VSM(config-port-prof)# channel-group auto mode on mac-pinning
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# exit

```

The system-uplink port-profile is applied to the system vNICs of the service profile. MTU 9000 is configured on uplink port-profile to enable jumbo frames. “channel-group auto mode on mac-pinning” is a very important configuration which ‘pins’ the VM vNICs to uplinks on the vDS. MAC pinning feature does static load balancing per vNIC basis. It also provides high-availability by moving the traffic to the alternate adapter when a given fabric is down.

12. Create storage-uplink port-profile as shown in [Figure 176](#). The storage-uplink port-profile corresponds to the storage vNICs of the service profile.

Figure 176 **Configuring Storage-Uplink Port-Profile**



```

10.29.150.167 - PuTTY
N1k-VSM(config)# port-profile type ethernet storage-uplink
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode trunk
N1k-VSM(config-port-prof)# switchport trunk native vlan 40
N1k-VSM(config-port-prof)# switchport trunk allowed vlan 40
N1k-VSM(config-port-prof)# mtu 9000
N1k-VSM(config-port-prof)# channel-group auto mode on mac-pinning
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# exit

```

13. Create data-uplink port-profile as shown in [Figure 177](#). The data-uplink port-profile corresponds to the data vNICs of the service profile.

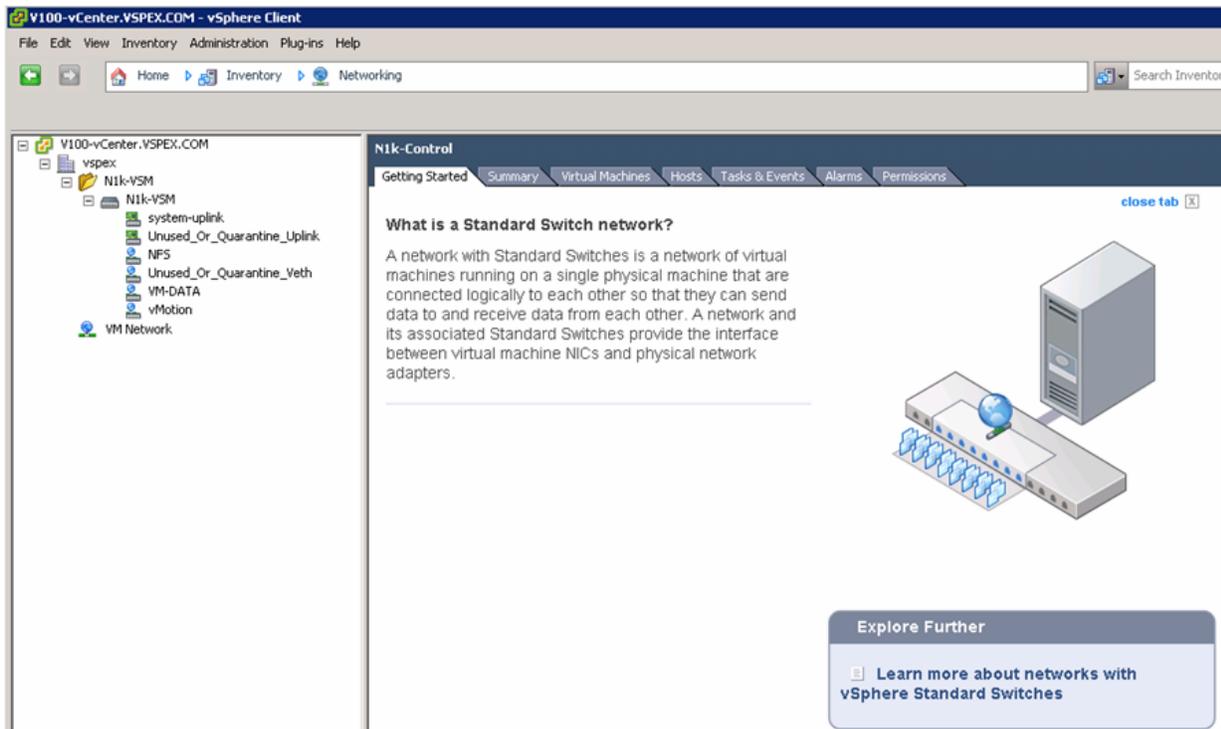
Figure 177 *Configuring Data-Uplink Port-Profile*

```

10.29.150.167 - PuTTY
N1k-VSM(config)# port-profile type ethernet data-uplink
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode trunk
N1k-VSM(config-port-prof)# switchport trunk native vlan 45
N1k-VSM(config-port-prof)# switchport trunk allowed vlan 45
N1k-VSM(config-port-prof)# mtu 9000
N1k-VSM(config-port-prof)# channel-group auto mode on mac-pinning
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# exit
    
```

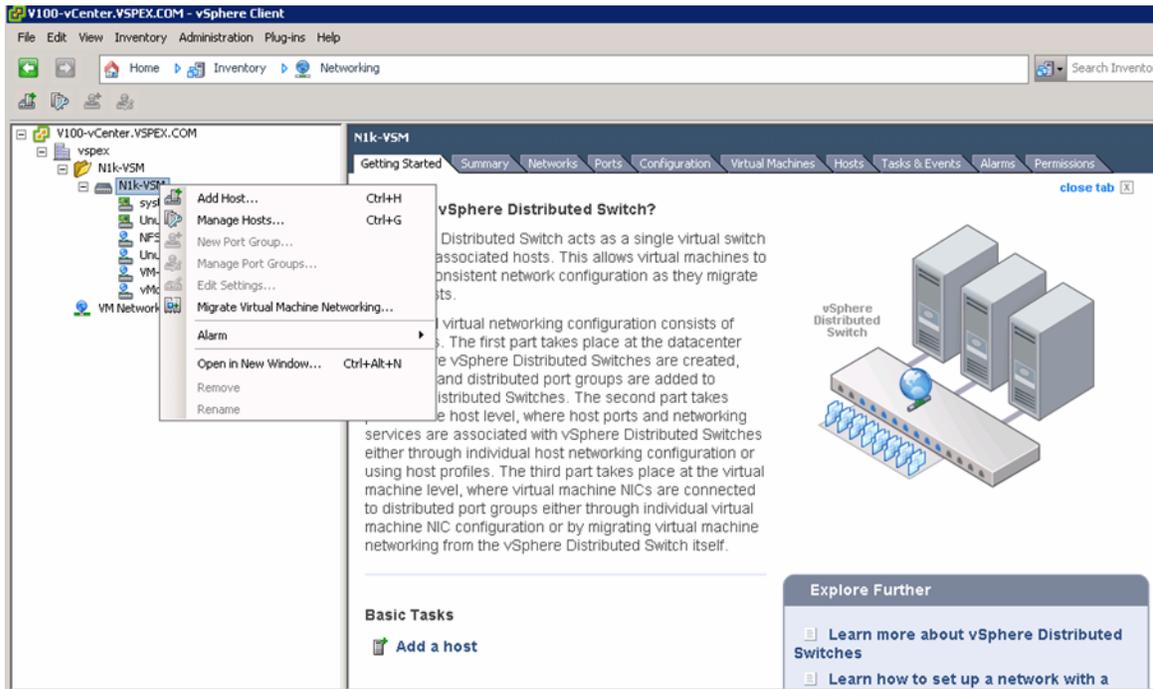
- Once VSM is connected to the vCenter, it shows up as a virtual Distributed Switch in the vCenter’s “Network” view as shown in [Figure 179](#).

Figure 178 *VSM Shown as Virtual Distributed Switch in vCenter*



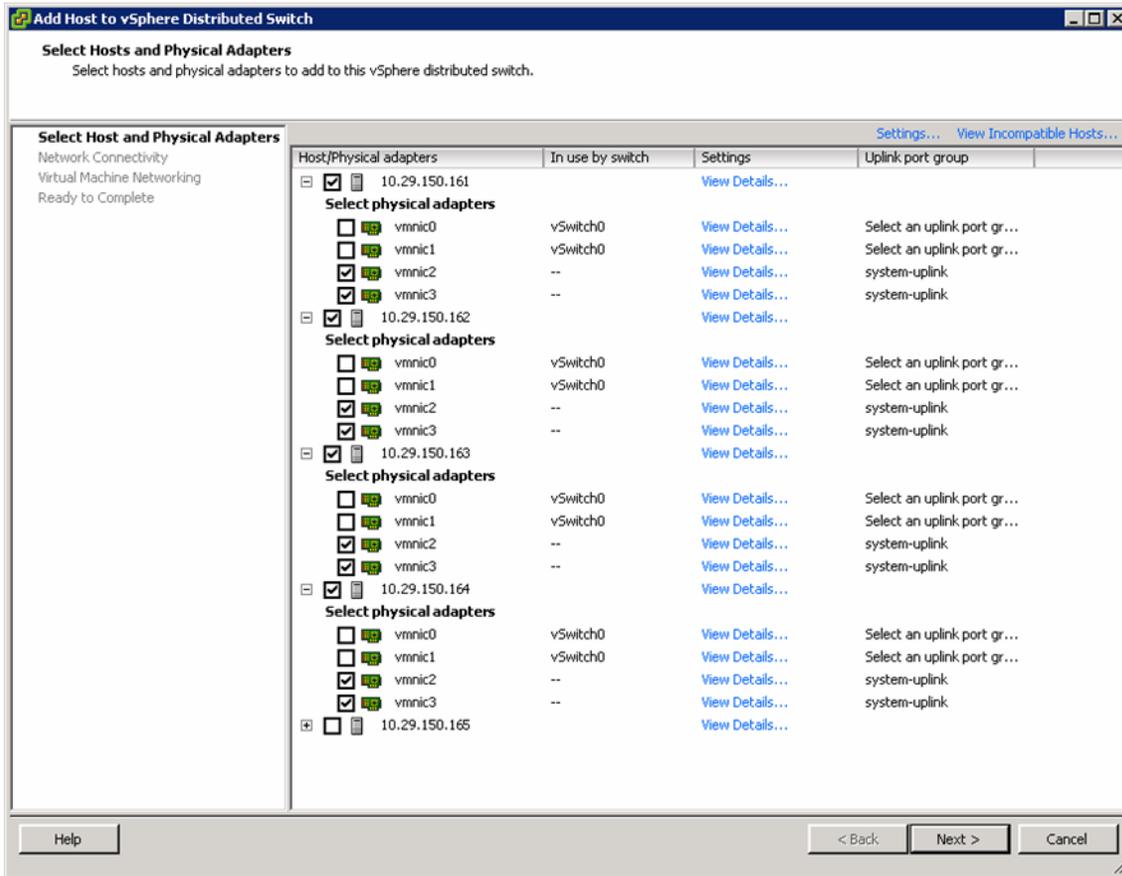
- Add hosts to the vDS.

Figure 179 Adding Host to vDS



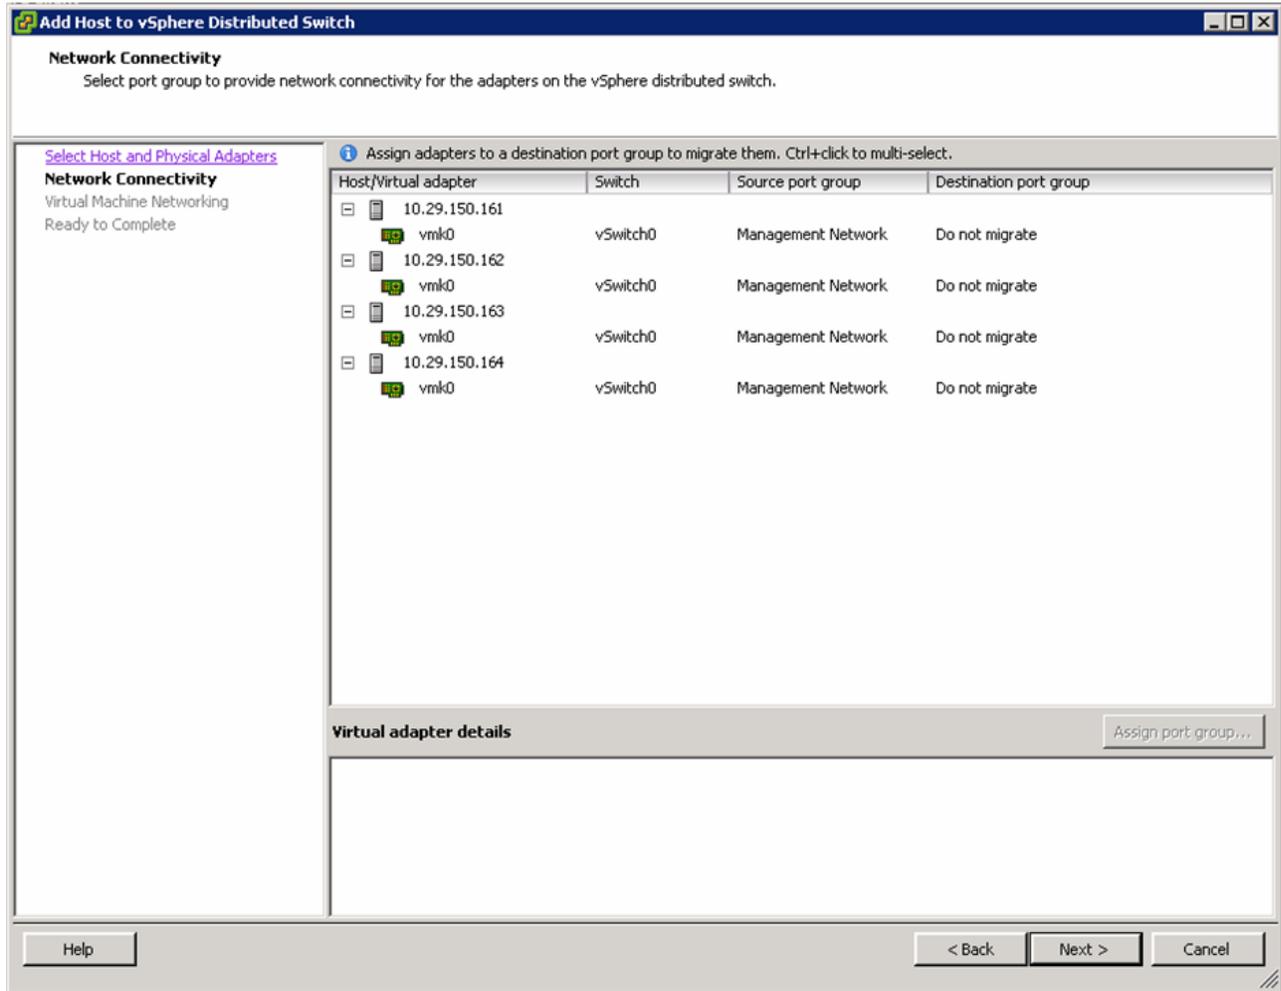
16. Select all the VSPEX ESXi hosts and add appropriate adapters using the uplink port-profiles created in the previous step. See [Table 10](#) for vNICs to uplink port-profile mapping.

Figure 180 Selecting Host and Adapters to Add to vDS



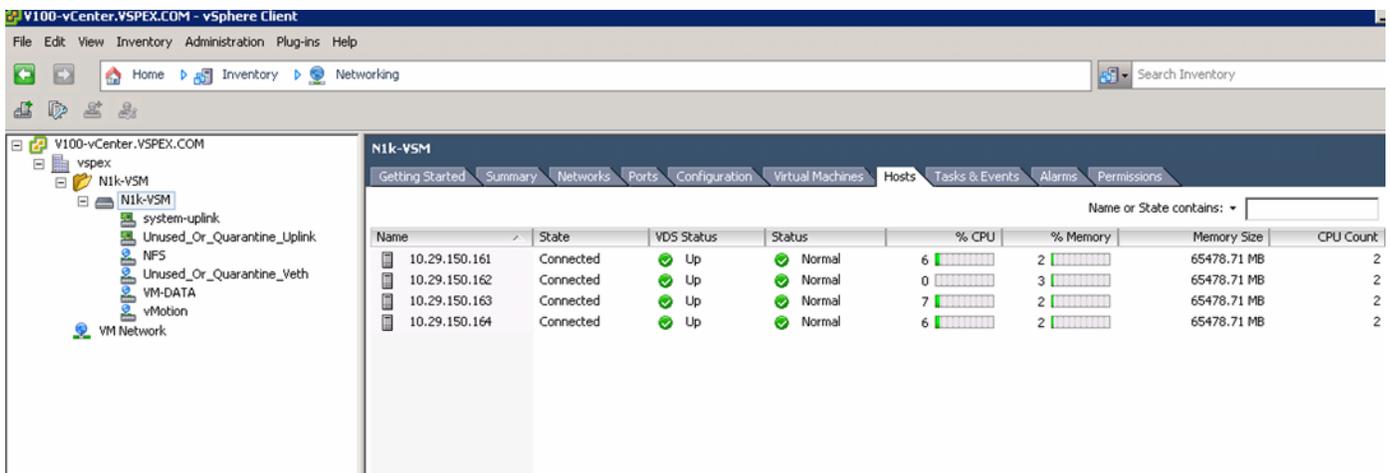
- Do not migrate management VM kernel from the native vSwitch to vDS, click **Next** and then click **Finish**.

Figure 181 Selecting Port Groups for Network Connectivity



18. Verify that all the hosts are successfully added to the vDS.

Figure 182 Verifying the Hosts Added to the vDS



Configure Port-Profiles and Add Virtual Machines

The last step of Nexus 1000v configuration and its integration with vCenter is creation of port-profiles and using them in the virtual machines in the vCenter. Note that this is possible only after carving out disk space for VMs on the storage array and deploying the VMs.

To configure port-profiles for VMs, follow these steps:

1. Create a port-profile for storage (NFS) access. Max-ports can be set to number of hosts you have in the architecture.

Figure 183 *Creating a Port-Profile for NFS Storage Access*

```
N1k-VSM(config)# port-profile type vethernet NFS
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode access
N1k-VSM(config-port-prof)# switchport access vlan 40
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# max-ports 5
N1k-VSM(config-port-prof)# description port-profile for NFS share access
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# exit
N1k-VSM(config)#
```

2. Create a port-profile for vMotion traffic. Max-ports can be set to number of hosts you have in the architecture.

Figure 184 *Creating a Port-Profile for vMotion Traffic*

```
N1k-VSM(config)# port-profile type vethernet vMotion
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode access
N1k-VSM(config-port-prof)# switchport access vlan 41
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# max-ports 5
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# description port-profile for vMotion traffic
N1k-VSM(config-port-prof)# exit
N1k-VSM(config)#
```

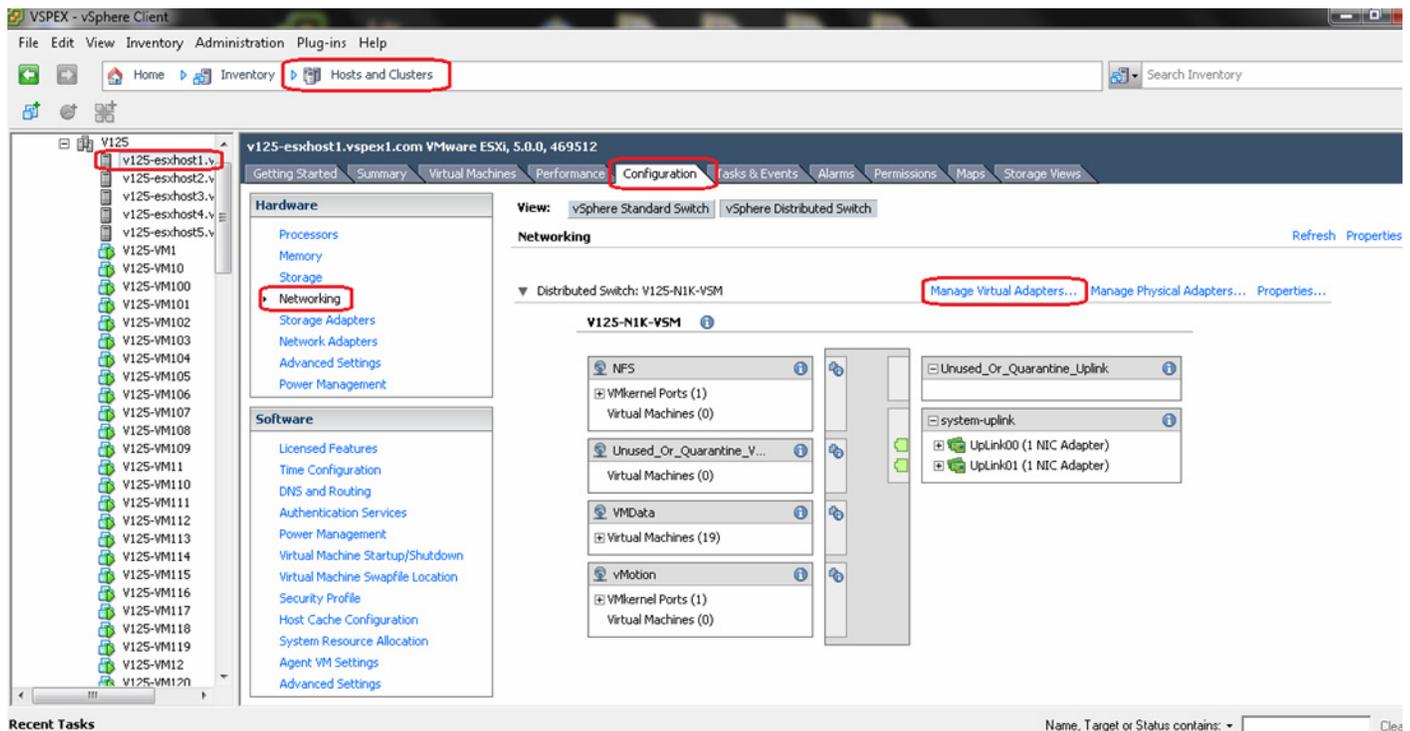
3. Create port-profiles for the virtual machine data traffic used by various applications as per your needs. You can set max ports to appropriate values based on the number of VMs being configured. Following is a sample port-profile:

Figure 185 Creating a Port-Profile for VM Data Traffic

```
N1k-VSM(config)# port-profile type vethernet VM-DATA
N1k-VSM(config-port-prof)# vmware port-group
N1k-VSM(config-port-prof)# switchport mode access
N1k-VSM(config-port-prof)# switchport access vlan 45
N1k-VSM(config-port-prof)# no shutdown
N1k-VSM(config-port-prof)# max-ports 101
N1k-VSM(config-port-prof)# description port-profile for virtual machine Ethernet
data traffic
N1k-VSM(config-port-prof)# state enabled
N1k-VSM(config-port-prof)# exit
N1k-VSM(config)#
```

- Once the port-profiles are configured, choose the **Hosts and Clusters** tab, choose the **ESXi Host**, click the **Configuration** tab, choose **Networking**, view **vSphere Distributed Switch**, and click the **Manage Virtual Adapters...** link.

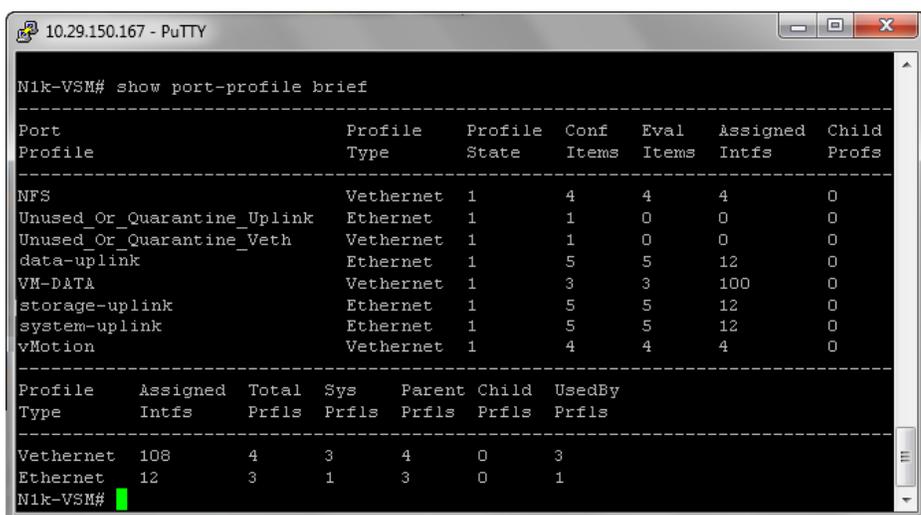
Figure 186 Manage Virtual Adapters in vCenter



- Click **Add** in the wizard, click **New virtual adapter**, and click **Next**.
- Select **VMKernel** on the next dialog box.
- Select port-profile NFS for the storage access, and click **Next**.
- Configure IP address from the NFS subnet and configure subnet mask. Click **Next** and the **Finish** to deploy the vNIC.
- Similarly, add one more vmknics (VM Kernel NIC) for vMotion. When providing the port-profile name, make sure that you choose vMotion port-profile and check the check box **Use this virtual adapter for vMotion**.

10. Repeat creation of the two vmknic virtual adapters for all the ESXi hosts.
11. Connectivity between all the vmknics can be tested by enabling SSH access to ESXi host, logging in to ESXi host using SSH and using **vmkping** command and ping to all vMotion IP addresses from each of the hosts. Similarly, all the hosts must be able to ping NFS share IP address.
12. Once NFS share is reachable, the NFS datastore can be discovered and mounted through vCenter. Virtual machines can be deployed on these NFS datastore using the VM-Data port-profile for the network access.
13. Verify the port-profile usage using **show port-profile brief**, **show port-profile usage**, or **show port-profile name <name>** command. Sample output is as shown in [Figure 187](#).

Figure 187 Running **show port-profile brief** to Verify Port-Profile Usage



By running the command **show port-profile name <uplink-port-profile-name>**, you can see the implicit creation of port-channels per ESXi host basis due to the “channel-group auto mode on mac-pinning” CLI configured under the port-profile. In addition to the Ethernet uplink ports, port-channels are also listed as assigned interfaces. Port-channel status can be further viewed/ validated using **show port-channel brief** command from VSM VM.

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

This subsection is divided in two subsections:

1. [Configure VM Datastores for FC-Variant of the Solution, page 170](#)
2. [Configuring VM Datastore for the NFS-Variant of the Solution, page 180](#)

Configure VM Datastores for FC-Variant of the Solution

See [Figure 13](#) for Storage Architecture for 300 VMs on VNX 5400 to have a high-level overview of storage architecture. Follow these steps to configure the data store:

1. Login to the EMC VMX Unisphere, and click the **Storage** tab. Click **Storage Configuration > Storage Pools** and click the **Pools** tab. Click **Create**.

Figure 188 **Creating Storage Pools**

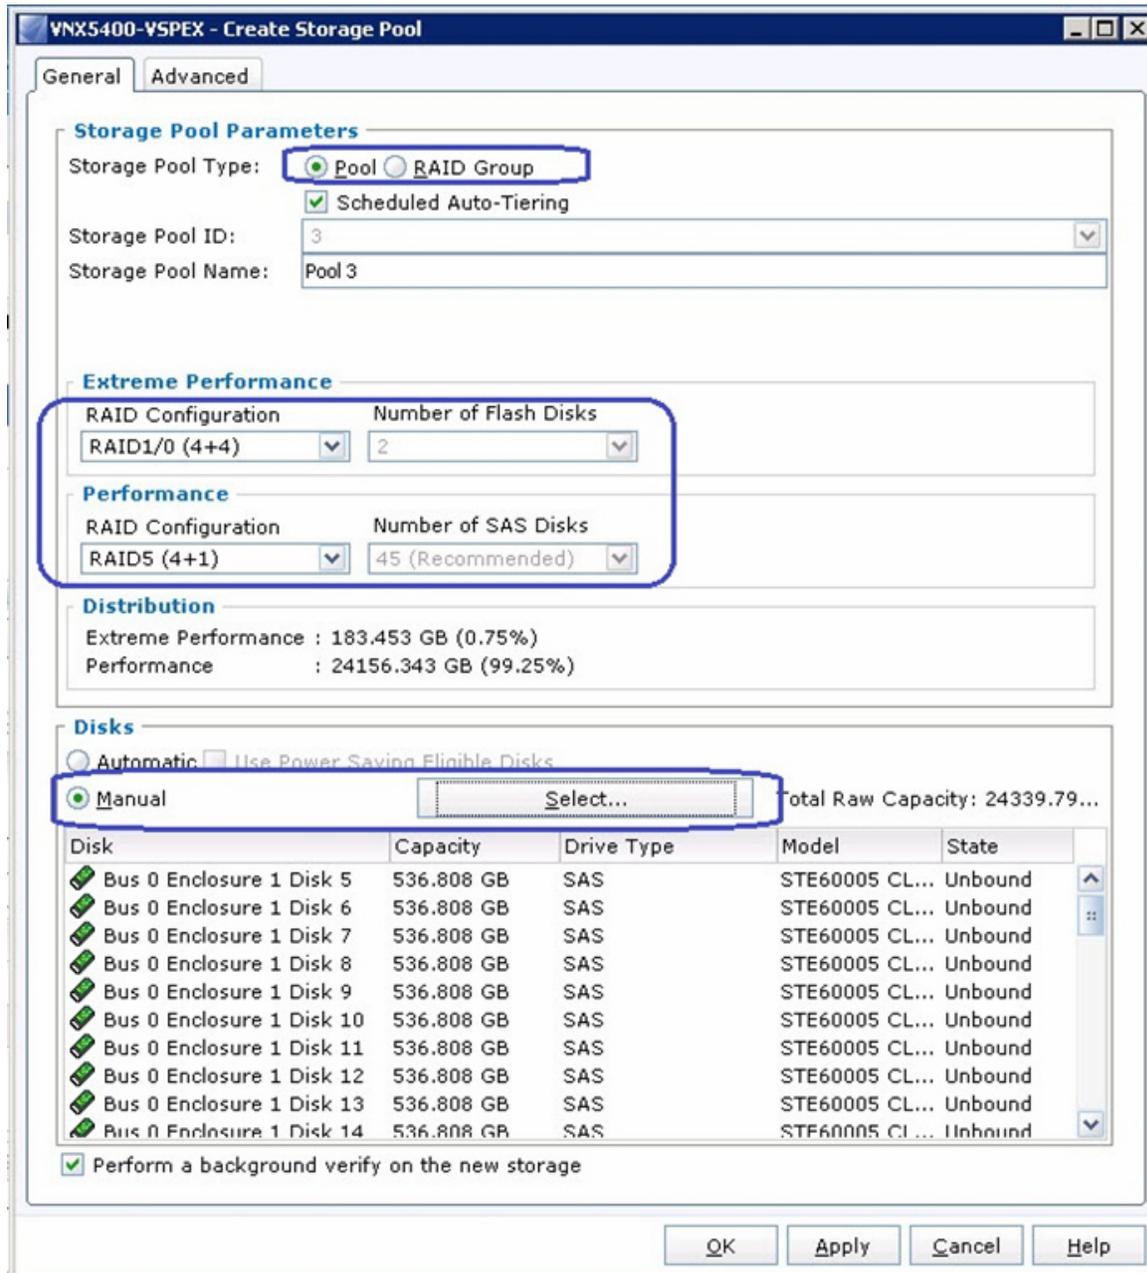
The screenshot shows the EMC Unisphere interface for creating storage pools. The top navigation bar includes 'Storage', 'Hosts', 'Data Protection', and 'Settings'. The breadcrumb trail is 'VNX5300-VSPEX > Storage > Storage Configuration > Storage Pools'. The 'Pools' tab is active, and a table displays the following data:

Name	FAST Cache	State	RAID Type	Drive Type	Total Capa...	Free Capa...	Allocated (...)	%Consum...	Subscribed...	%Subscrib...	%Full
Pool 0	On	Ready	RAID5	SAS	2146.434	1772.205	374.229		1,156.755	53.892	

At the bottom of the interface, the 'Create' button is highlighted, along with 'Delete', 'Properties', and 'Defragment' buttons. The status indicates '0 Selected' and '0 items'.

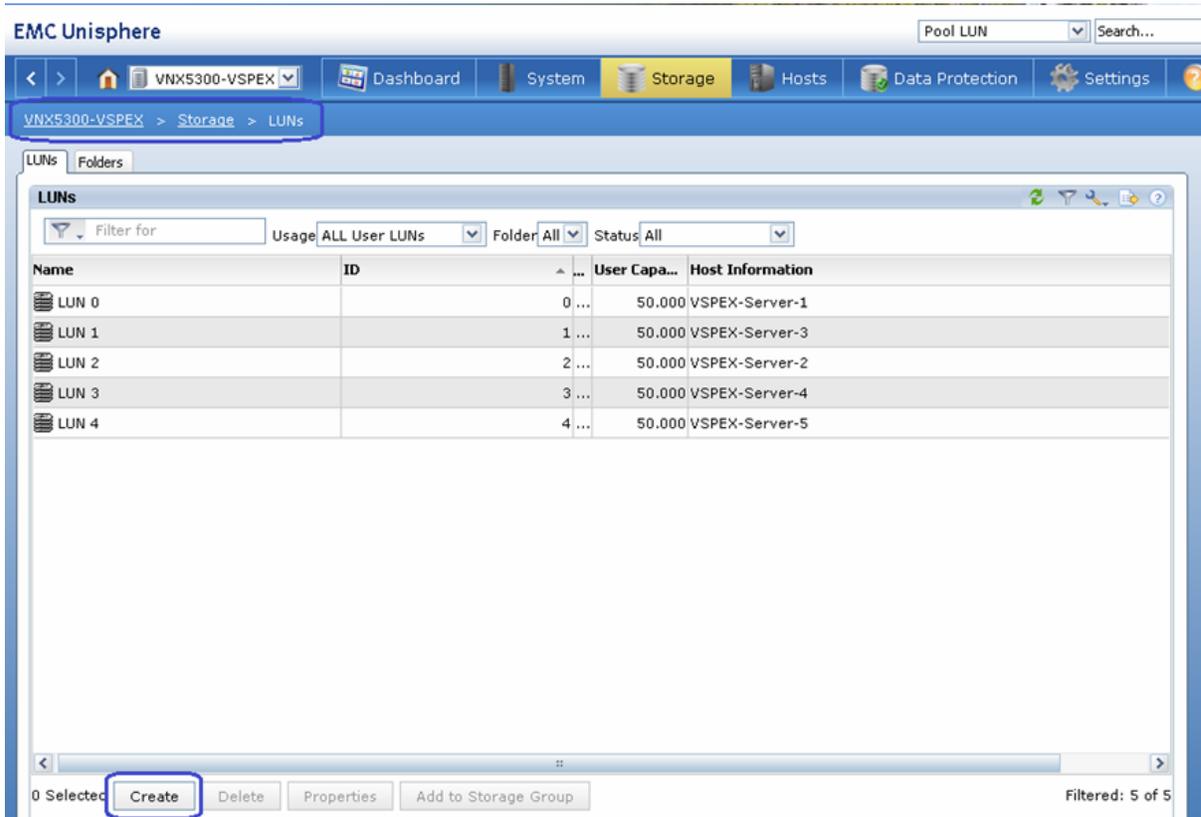
2. Create a new pool. Choose **Manual** for disk selections and choose 45 for the number of SAS disks and 2 for the number of Flash disks to create one pool.

Figure 189 Details for Creating Storage Pool



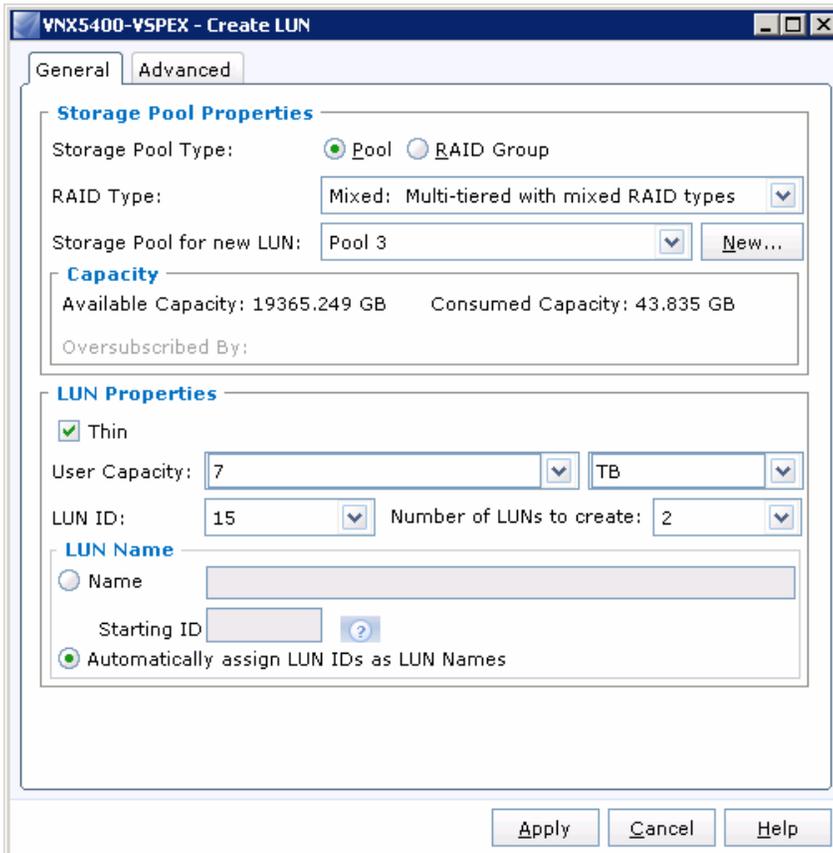
- Repeat this step for 2 more times to create total of three pools for VM datastorage. For the third pool, you would add 20 SAS drives and 2 Flash drives (See Figure 13 for Storage Architecture for 300 VMs on VNX 5400).
- Click **Storage > LUNs**. You will see 5 boot LUNs created for 5 hosts. Click **Create** to create the LUN for VM datastore.

Figure 190 **Creating LUNs for VM Datastore**



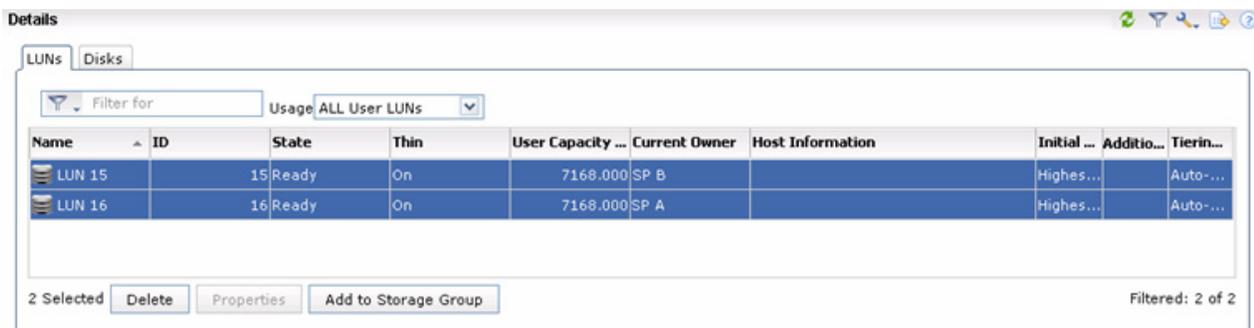
5. Click the **Pool** radio button for storage pool type, and choose the first VM data Pool ID from the drop-down list, which was created in step 2. Make sure to check the Thin check box for provisioning. Select User Capacity as **7 TB** from the drop-down list and Number of LUNs to Create per pool as 2, and click **Apply**.

Figure 191 Details for Creating LUN



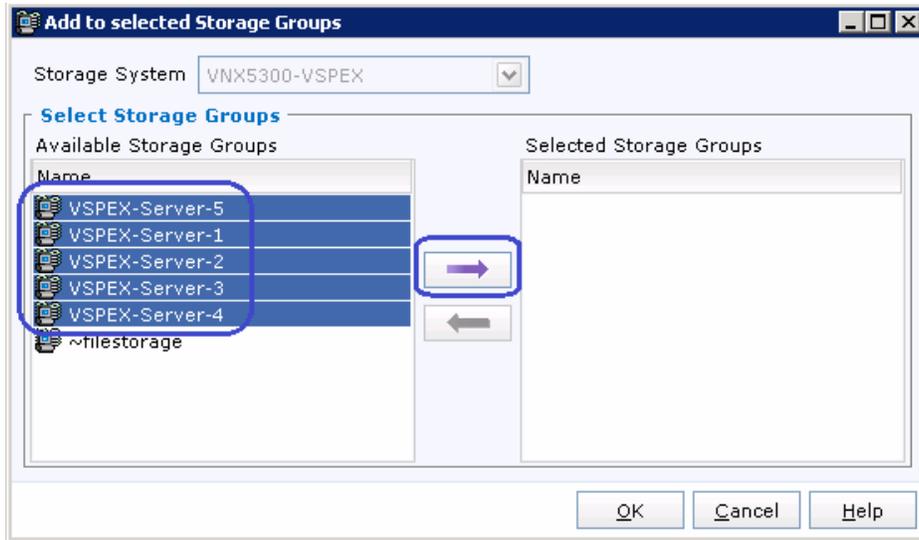
6. Repeat step 5 for all the 3 pools in the system. Note that 3rd pool will have reduced LUN size of 2.2 TB.
7. Select all the newly created LUNs, and click **Add to Storage Group**.

Figure 192 Adding the Created LUN to Storage Group



8. Select all ESXi servers, and move them to the right side. This will allow all ESXi hosts to see the datastore, which is essential for the vMotion of VMs across the cluster.

Figure 193 Moving Storage Groups to Selected Storage Groups



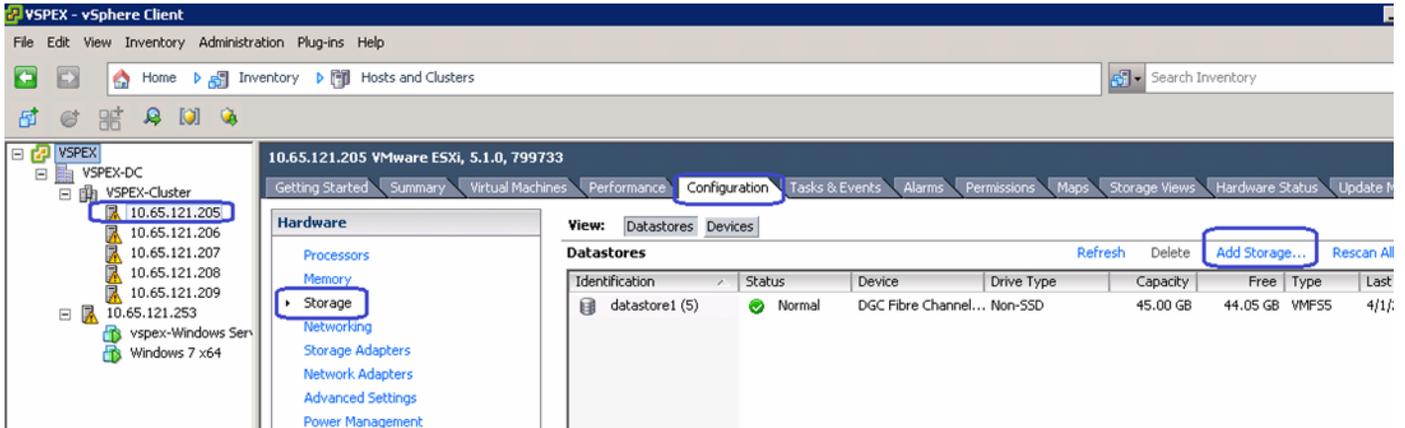
- On the LUNs window, you will see the storage group (and hence host) access for LUNs as shown in Figure 194.

Figure 194 Storage Group and Host Information

Name	ID	State	User Capacity (GB)	Host Information
LUN 0	0	Ready	50.000	VSPEX-Server-1
LUN 1	1	Ready	50.000	VSPEX-Server-3
LUN 2	2	Ready	50.000	VSPEX-Server-2
LUN 3	3	Ready	50.000	VSPEX-Server-4
LUN 4	4	Ready	50.000	VSPEX-Server-5
LUN 5	5	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 6	6	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 7	7	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 8	8	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 9	9	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 10	10	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 11	11	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5
LUN 12	12	Ready	2141.336	VSPEX-Server-4; VSPEX-Server-3; VSPEX-Server-2; VSPEX-Server-1; VSPEX-Server-5

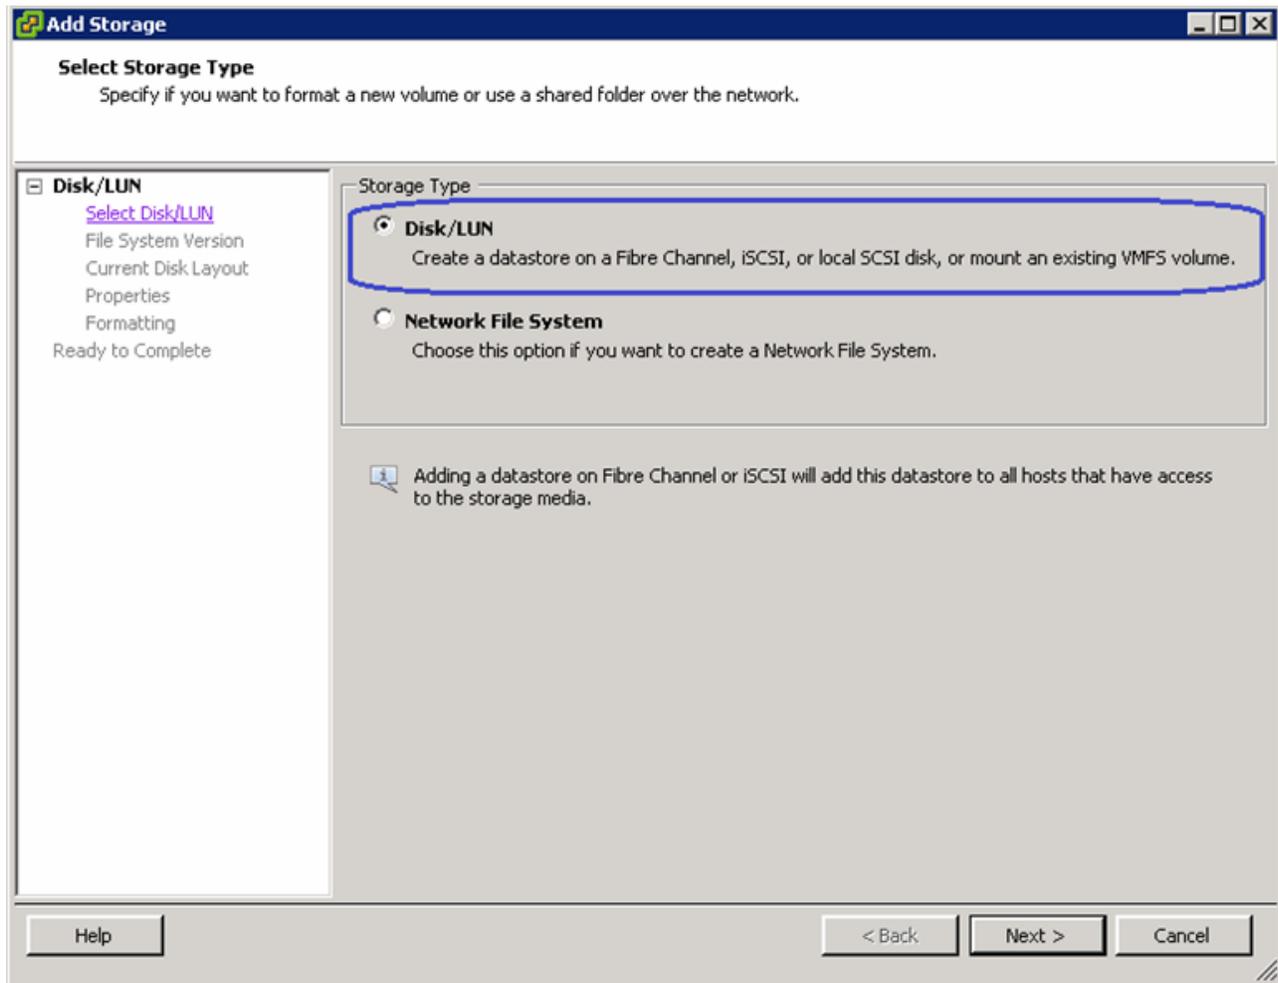
- login to vCenter GUI, select a particular host from the **Hosts and Clusters** view, click **Configuration** and then **Storage**. Click the **Add Storage** link.

Figure 195 Adding Storage



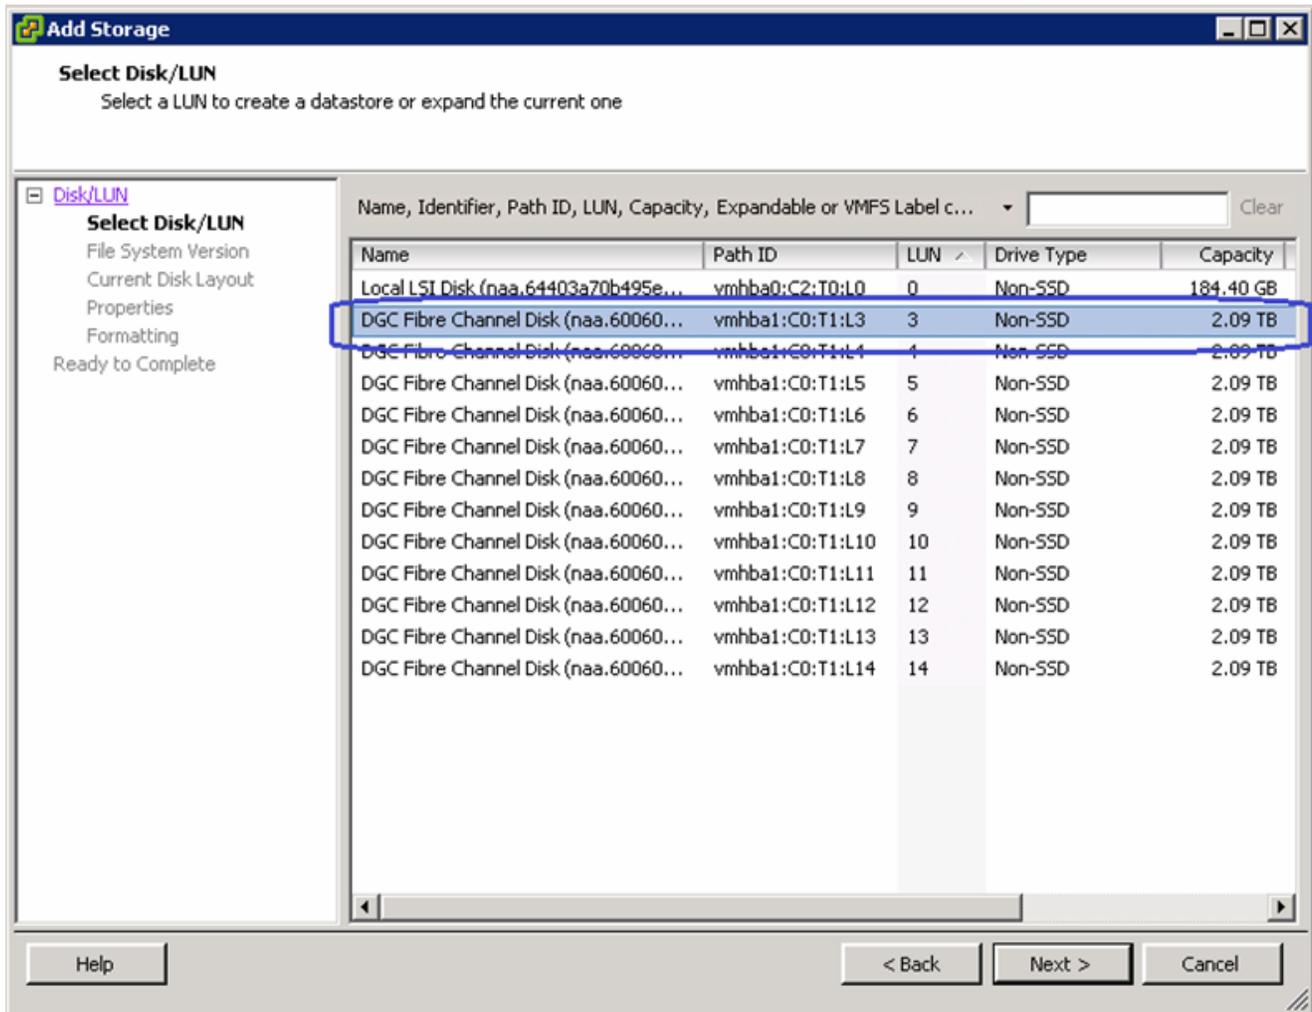
11. Click the **Disk/LUN** radio button in the wizard, click **Next**.

Figure 196 Choose the Storage Type



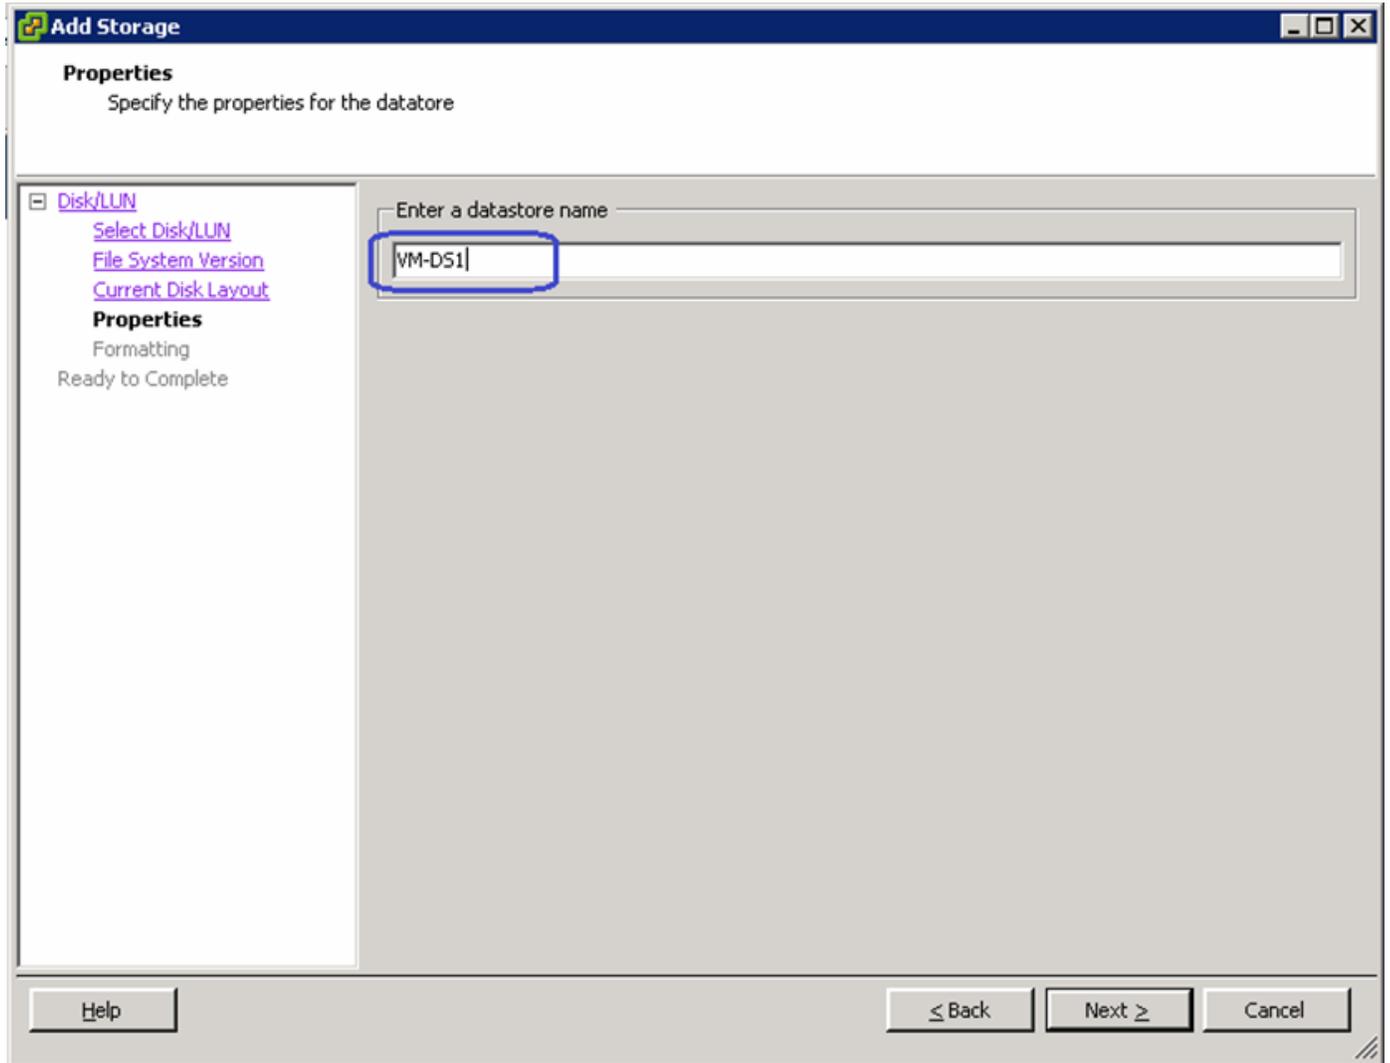
12. Select the **DGC Fibre Channel Disk (..)** from the list and click **Next**.

Figure 197 Selecting a LUN



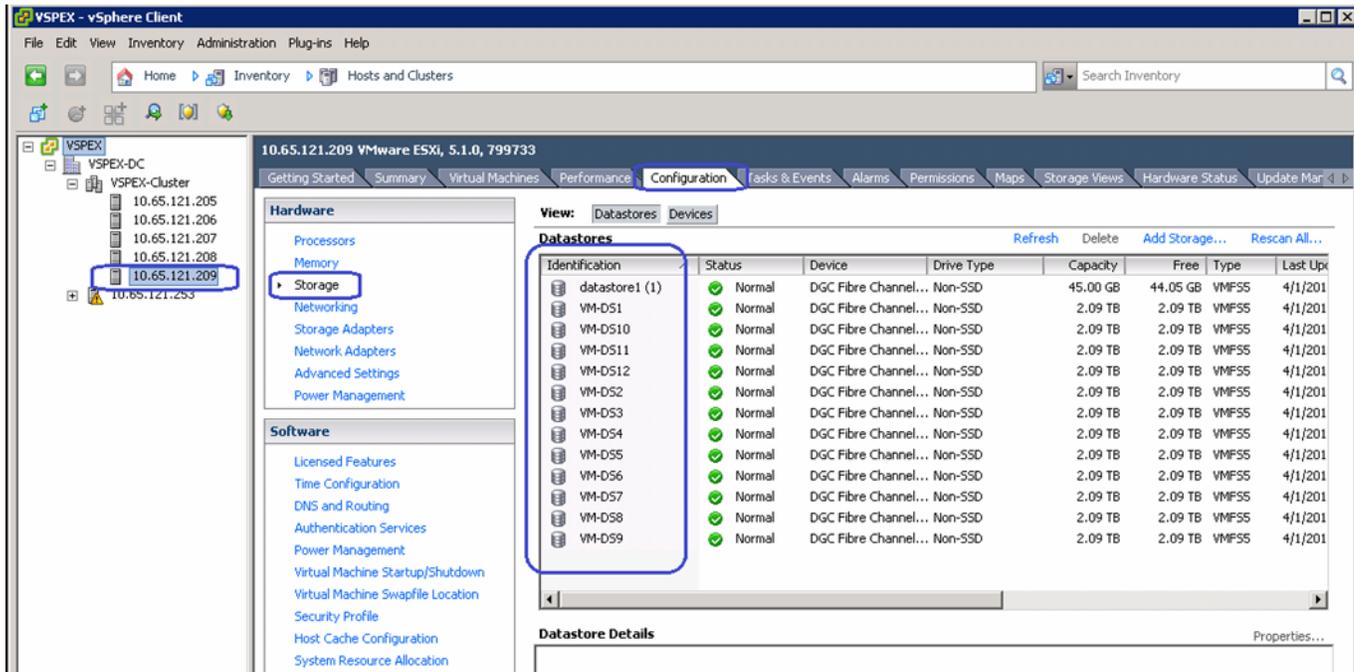
13. Enter the name of the first datastore as VM-DS1 and click **Next** and then **Finish**.

Figure 198 Specifying a Name for Datastore



14. Repeat steps 10 to 13 for all the datastores. Once data stores are added to one host, it will automatically show up for the other hosts too. The end result looks like [Figure 199](#) (on each host):

Figure 199 Datastores Added to the Hosts



Configuring VM Datastore for the NFS-Variant of the Solution

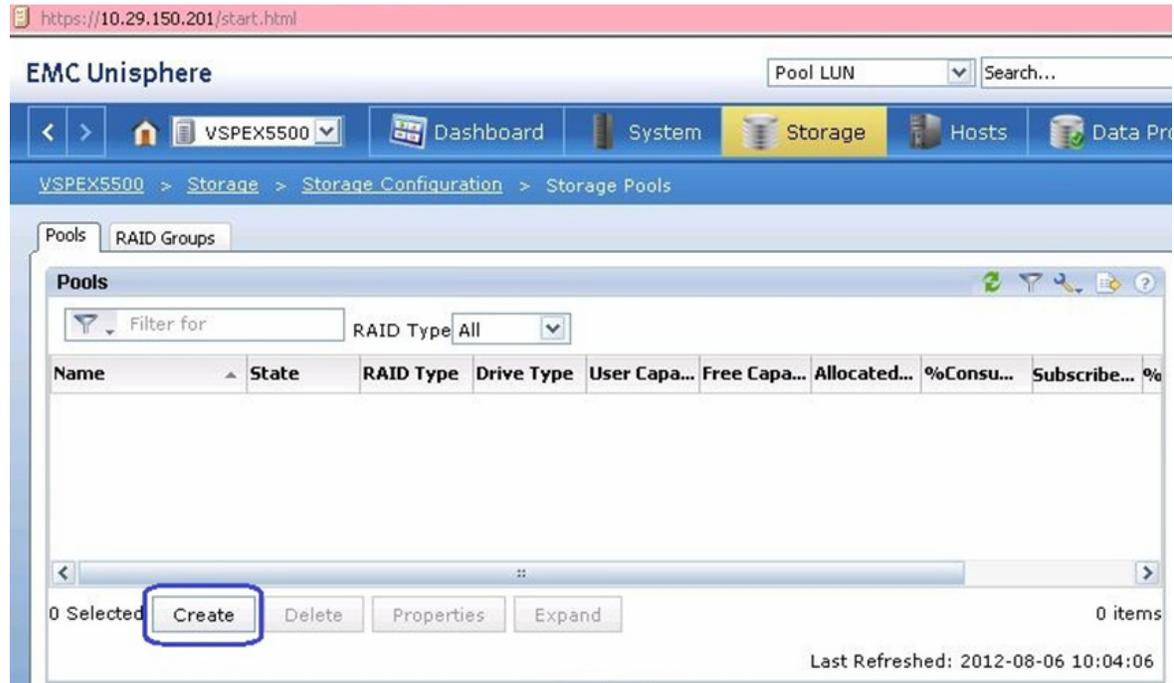
1. To Create Storage Pools for NFS Datastore. Click **Storage > Storage Configuration > Storage pools**.

Figure 200 Selecting Storage Pools in EMC Unisphere



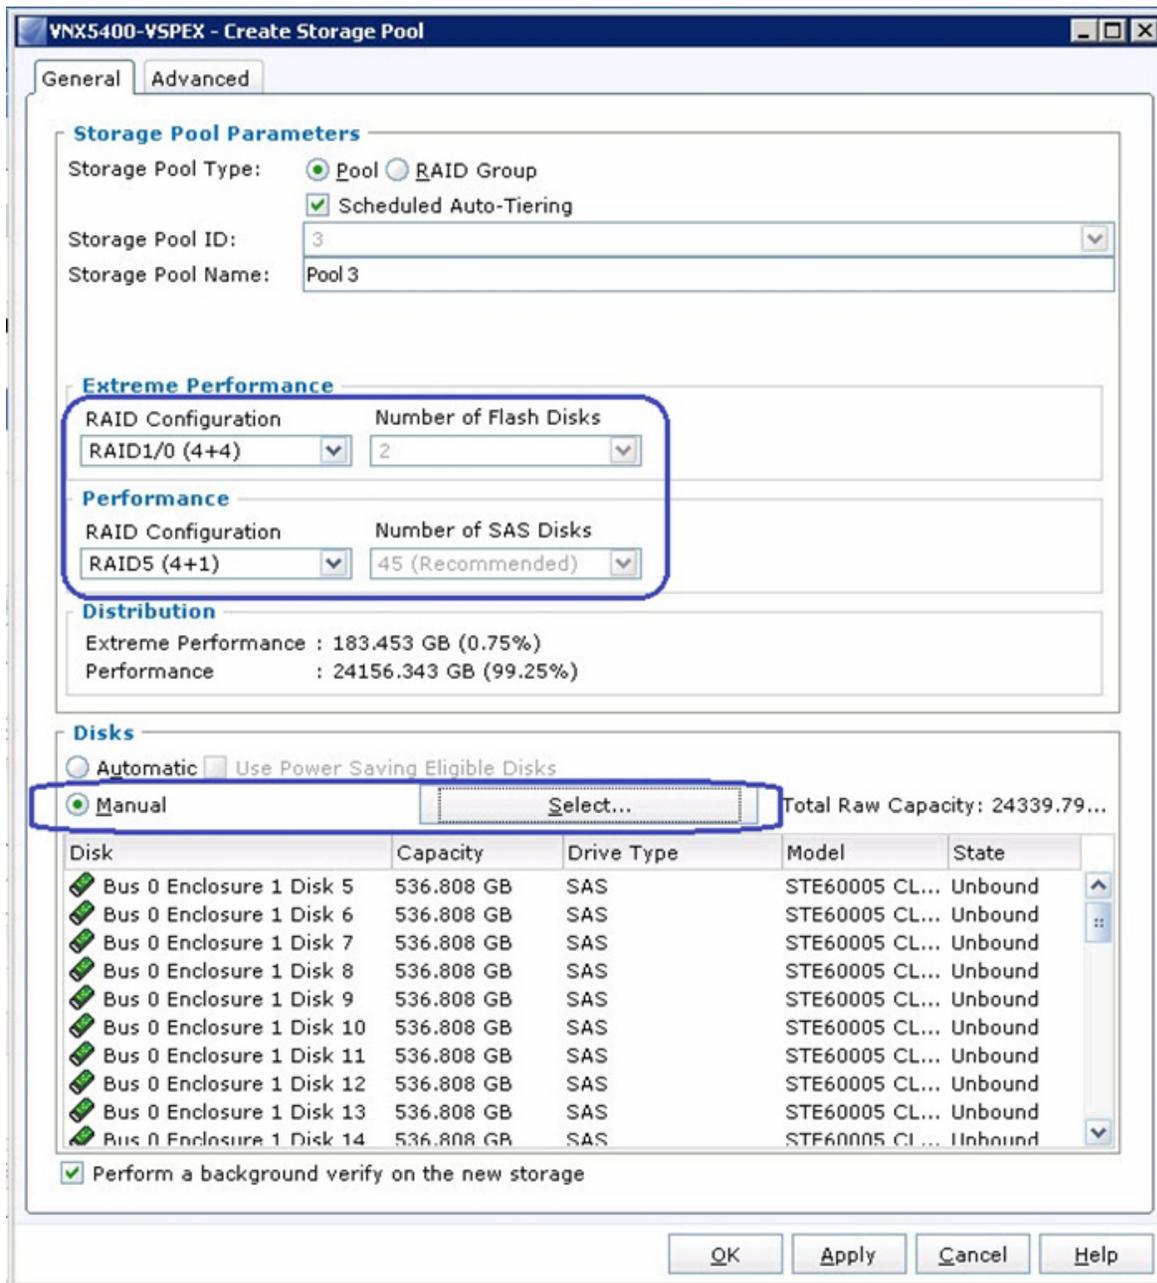
2. In the EMC Unisphere GUI, choose **Storage Pools > Pools**. Click **Create**.

Figure 201 **Creating Storage Pools**



3. In the Create Storage Pool window, manually select 45 SAS drives and 2 SATA Flash drives and add to the pool. Under Extreme Performance, choose RAID10 (4 + 4) and under Performance, choose RAID configuration as RAID5 (4+1) from the drop-down list. Number of Flash/ SAS disks will be automatically populated based on the number of disks that is manually added to the pool.

Figure 202 Details for Creating Storage Pools



- Repeat step 3 for required number of pools depending on your architecture. See the Storage Layout diagram for 600 and 1000 VMs, Figure 14 and Figure 15 respectively.
- To create LUNs from the newly created pools for NFS Datastore; choose **Storage**, right-click on PerformancePool (or “Pool 0”, whatever name you have given). Choose **Create LUN**.

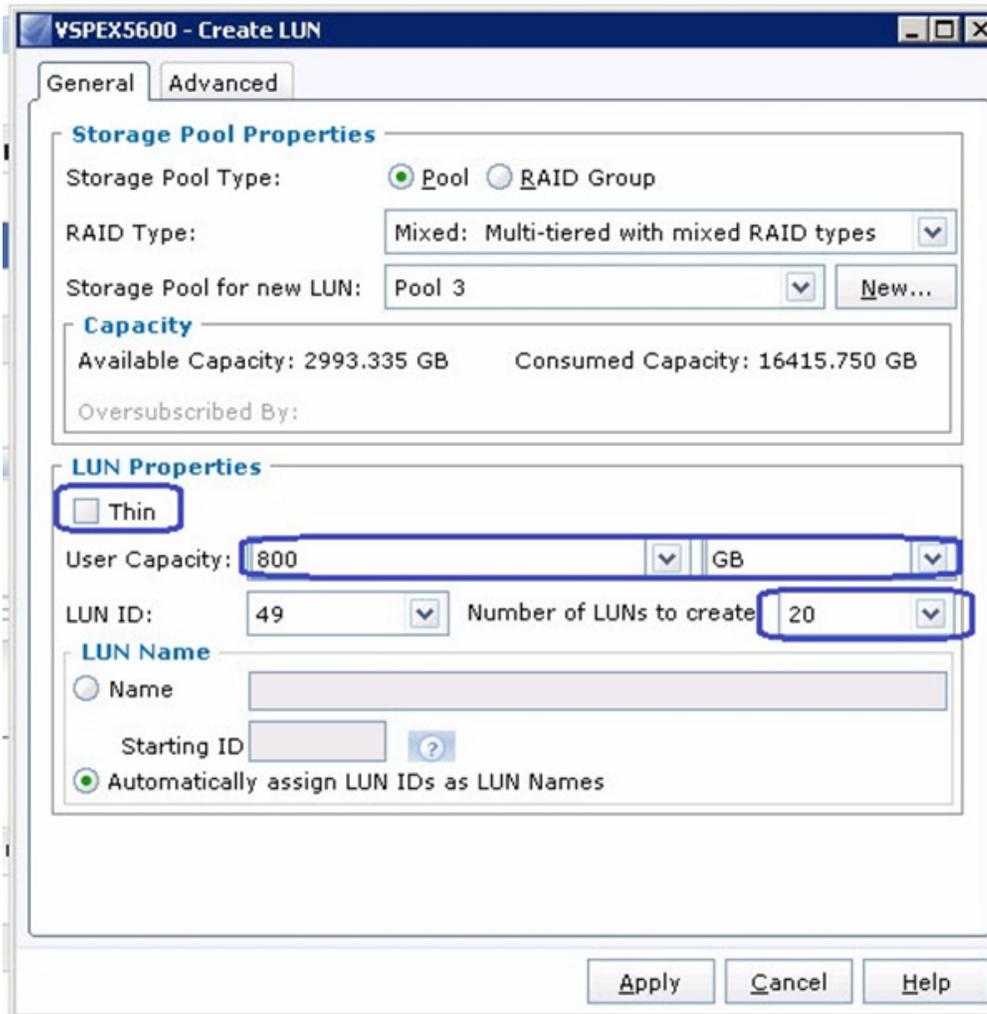
Figure 203 Creating LUNs form Newly Created Pool

The screenshot shows the EMC Unisphere interface for VSPEX5500. The breadcrumb trail is: VSPEX5500 > Storage > Storage Configuration > Storage Pools. The 'Pools' tab is active, and a context menu is open over a pool entry, with 'Create LUN' selected. Below the table, the 'Details' section shows 'Pool LUNs' and 'Disks' tabs.

Name	State	RAID Type	Drive Type	User Capa...	Free Capa...	Allocated...	%Consu...	Subscribe...	%
PerformanceP...	...	RAID5	SAS	32147.666	32147.666	0.000			0

6. Make sure, the check box **Thin** is unchecked, User Capacity is 800 G, and number of LUNs to create is 20.

Figure 204 Details for Creating LUN



7. Select the pool and Select all the newly created LUNs and Click **Add to Storage Group** as shown [Figure 205](#). Make sure you select all the LUNs from the pools.

Figure 205 Adding Pools and LUNs to Storage Group

VSPEX5500 > Storage > Storage Configuration > Storage Pools

Pools RAID Groups

Pools

Filter for RAID Type All

Name	State	RAID Type	Drive Type	User Capa...	Free Capa...	Allocated...	%Consu...	Subscribe...	%Subscri...	Auto-Tieri...
PerformancePool	Ready	RAID5	Mixed	95307.785	48716.965	46590.820	<div style="width: 50%;"></div>	46,440.527	48.727	Scheduled

1 Selected Create Delete Properties Expand 1 items

Last Refreshed: 2012-08-06 14:14:35

Details

Pool LUNs Disks

Filter for Usage ALL User LUNs

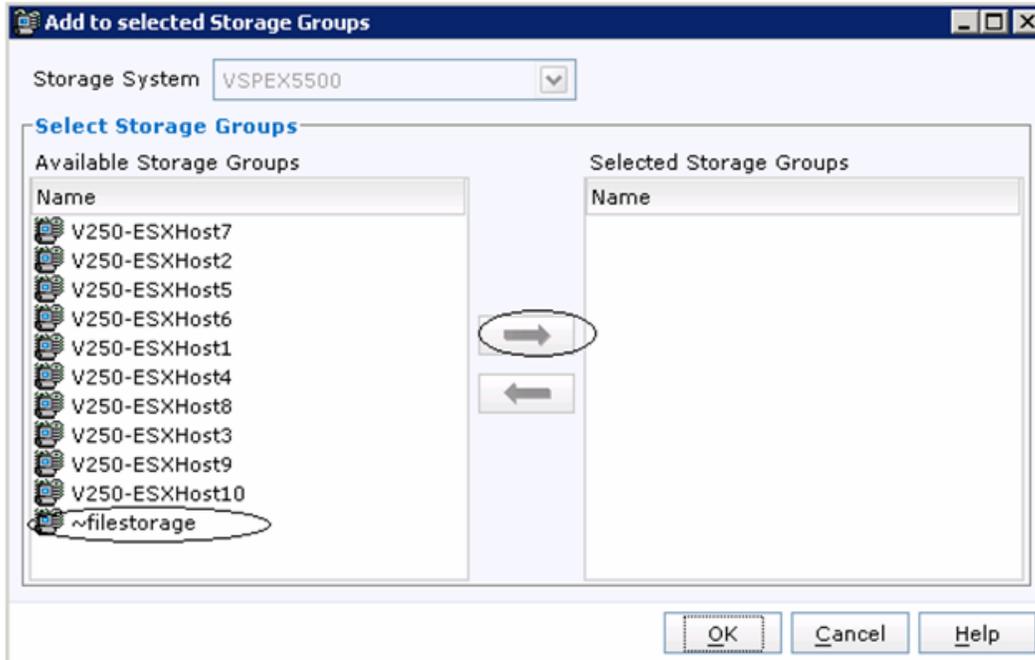
Name	ID	State	User Capacity (GB)	Current Owner	Host Information
LUN 11	11	Ready	800.000	SP A	
LUN 12	12	Ready	800.000	SP B	
LUN 13	13	Ready	800.000	SP A	
LUN 14	14	Ready	800.000	SP B	
LUN 15	15	Ready	800.000	SP A	
LUN 16	16	Ready	800.000	SP B	
LUN 17	17	Ready	800.000	SP A	
LUN 18	18	Ready	800.000	SP B	
LUN 19	19	Ready	800.000	SP A	
LUN 20	20	Ready	800.000	SP B	

150 Selected Delete Properties Add to Storage Group Filtered: 150 of 150

Last Refreshed: 2012-08-06 14:14:37

- From the Available Storage Groups, Select “~filestorage” and click the **Arrow** tab. Once “~filestorage” is selected on right pane, click **OK**.

Figure 206 Moving Storage Groups to Selected Storage Groups



9. In the EMC Unisphere GUI, choose **Storage > Storage Configuration > Storage Pools for Files**, and click **Rescan storage systems** on the right pane. Rescan will take up to 4 minutes of time. Once rescan successfully finishes (track the progress at Background task for files page under System menu), click **Refresh** and newly created storage pools must be visible in the left pane of the window.

Figure 207 Created Storage Pools are Shown in Storage Pools for File Window of EMC Unisphere

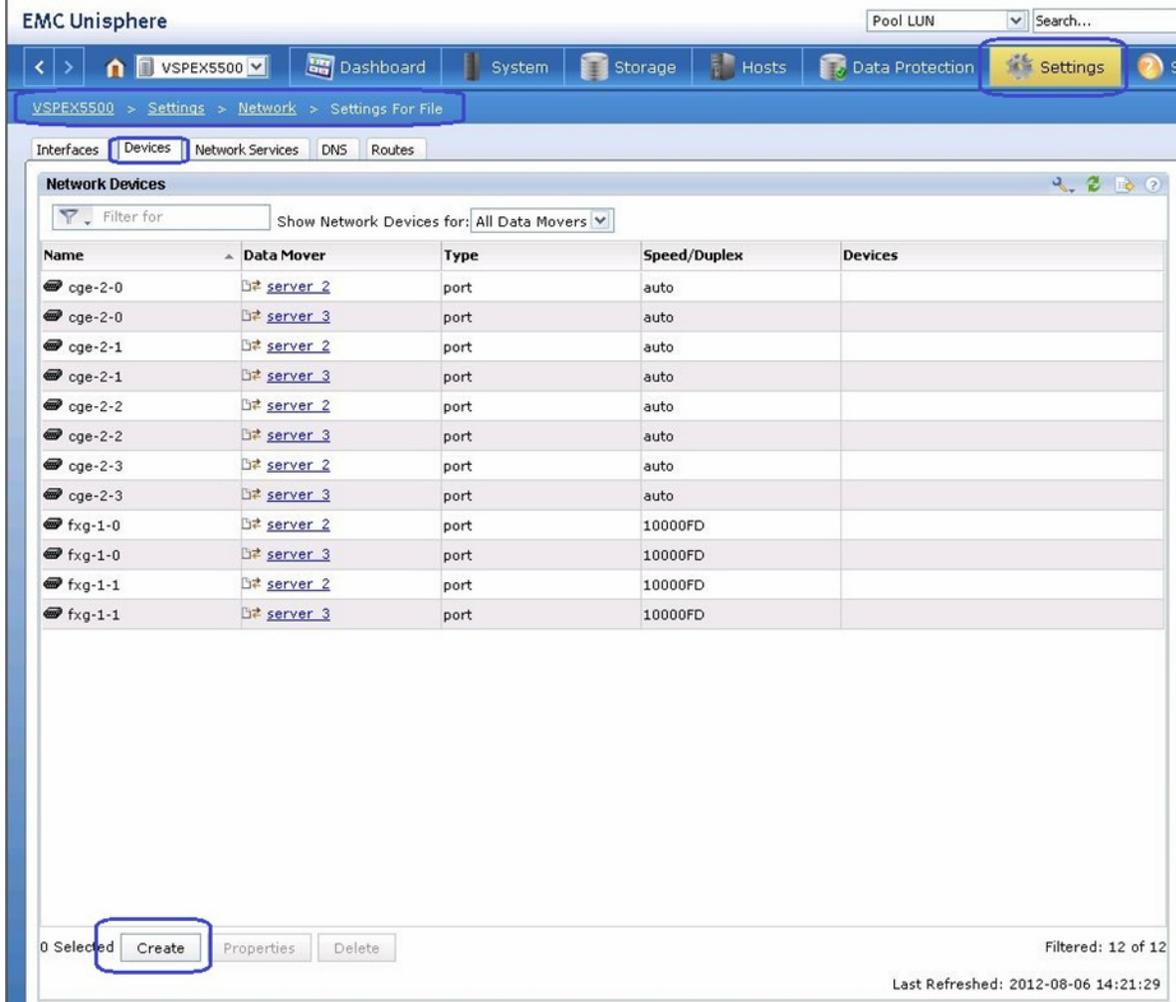
The screenshot shows the EMC Unisphere interface for the 'Storage Pools for File' window. The main content area displays a table with the following data:

Name	Description	Storage Capacity (GB)	Storage Used(%)	Type	Disk Type	Automatic Extension
Pool 1	Mapped Pool Pool 1 on FNM001...	15999.992	<div style="width: 100%; height: 10px; background-color: #f0f0f0;"></div>	Mapped pool	Mixed	Enabled
Pool 3	Mapped Pool Pool 3 on FNM001...	15999.992	<div style="width: 100%; height: 10px; background-color: #f0f0f0;"></div>	Mapped pool	Mixed	Enabled

The interface also includes a sidebar with various management options such as 'Wizards', 'Tiering', 'Data Migration', 'Block Storage', 'File Storage', and 'Atmos Management'. The 'File Storage' section is highlighted, and the 'Rescan Storage Systems' link is circled in red.

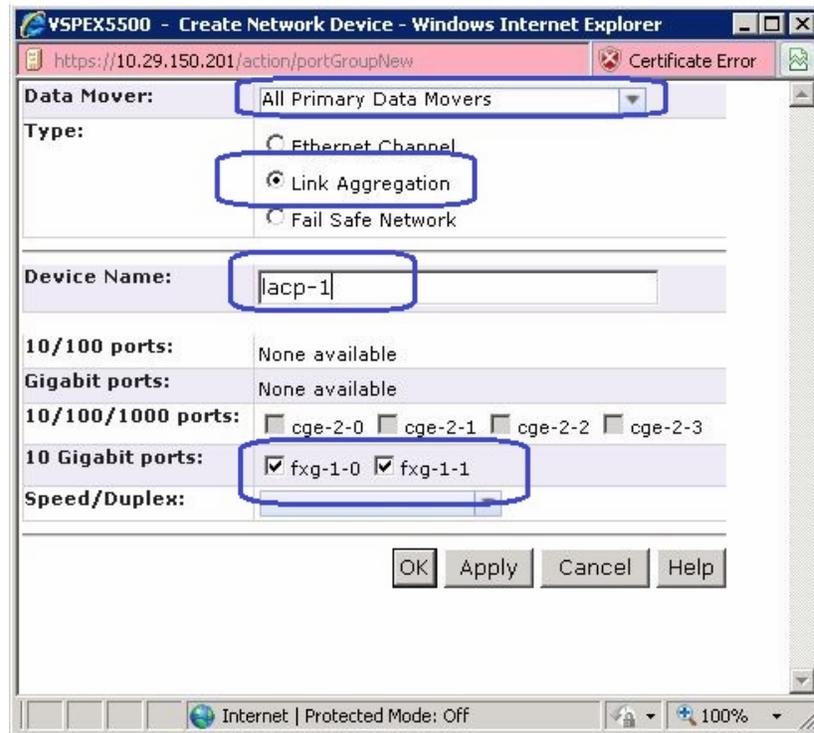
- At this point, the NFS volume is created. Next step is to create highly available network access for the NFS volume. To Create LACP interface, navigate to **Settings > Network > Settings for File**, and click the **Devices** tab. Click **Create**.

Figure 208 Creating LACP Interface



11. Choose Date Mover as **All Primary Data Movers** from the drop-down list, click the **Link Aggregation** for Type and enter the Device name as lacp-1. Check the two check boxes for 10 Gigabit ports **fxg-1-0** and **fxg-1-1**. Click **OK** to proceed to the Network Device Creation.

Figure 209 *Creating Network Device*



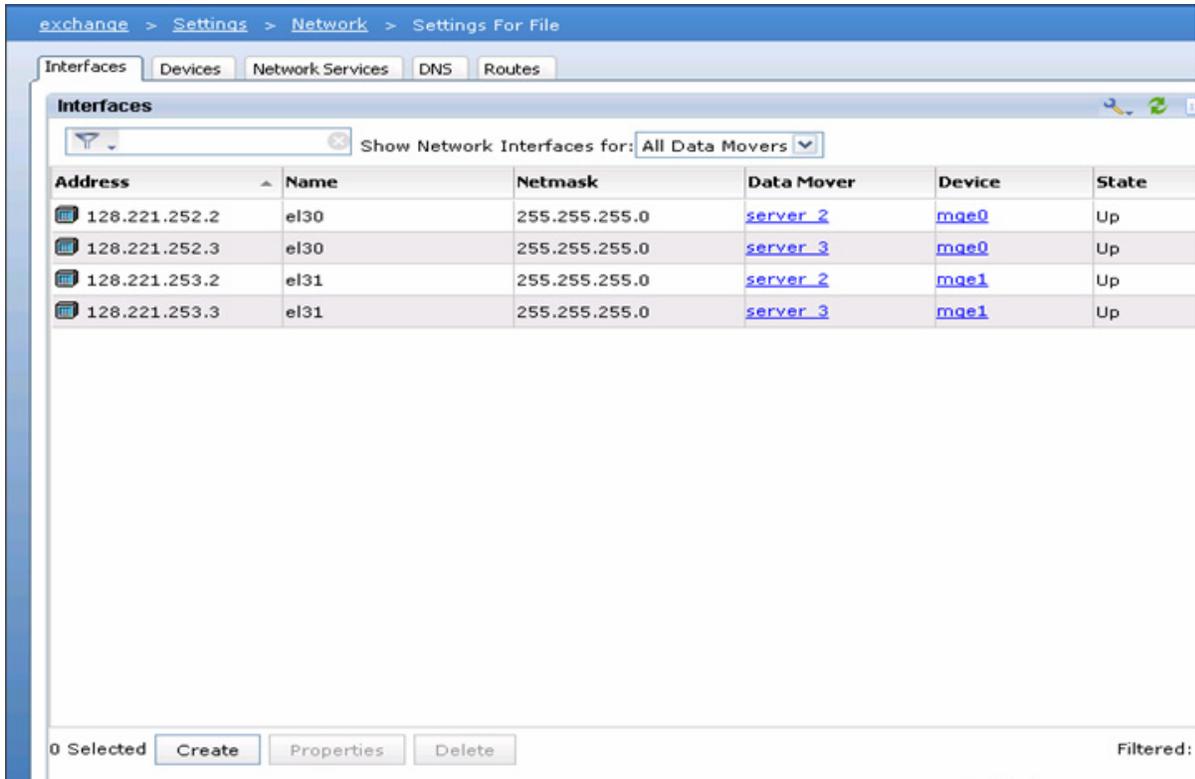
12. Figure 210 shows the creation of LACP Network device name as “lACP-1”

Figure 210 *Created LACP Network Device is Shown Under Network Devices*

Name	Data Mover	Type	Speed/Duplex	Devices
cge-2-0	server 2	port	auto	
cge-2-0	server 3	port	auto	
cge-2-1	server 2	port	auto	
cge-2-1	server 3	port	auto	
cge-2-2	server 2	port	auto	
cge-2-2	server 3	port	auto	
cge-2-3	server 2	port	auto	
cge-2-3	server 3	port	auto	
fxg-1-0	server 2	port	10000FD	
fxg-1-0	server 3	port	10000FD	
fxg-1-1	server 2	port	10000FD	
fxg-1-1	server 3	port	10000FD	
lACP-1	server 2	lACP	10000FD	fxg-1-0,fxg-1-1

- From the **Settings for File** tab. Click **Interfaces** and then click **Create**.

Figure 211 **Creating Network Interface**



- Choose the Data Mover as **server_2** and Choose Device name as **lACP-1** from the drop-down list. Specify the valid IP address, Netmask and Interface name as “fs01” & MTU value as “9000” to allow jumbo frames for the lACP interface.

Figure 212 **Details for Creating Network Interface**

The screenshot shows a web browser window titled "VSPEX5500 - Create Network Interface - Windows Internet Explorer". The address bar shows "https://10.29.150.201/action/interfaceDisplay" and a "Certificate Error" warning. The form contains the following fields and values:

Data Mover:	server_2
Device Name:	larp-1
Address:	10.10.40.11
Name:	fs01
Netmask:	255.255.255.0
Broadcast Address:	10.10.40.255
MTU:	9000
VLAN ID:	

At the bottom of the form are four buttons: OK, Apply, Cancel, and Help. The browser status bar at the bottom shows "Internet | Protected Mode: Off" and a zoom level of "100%".

- To Create File system for NFS data store, navigate to **Storage > Storage Configuration**. Click **File Systems** and then click **Create**.

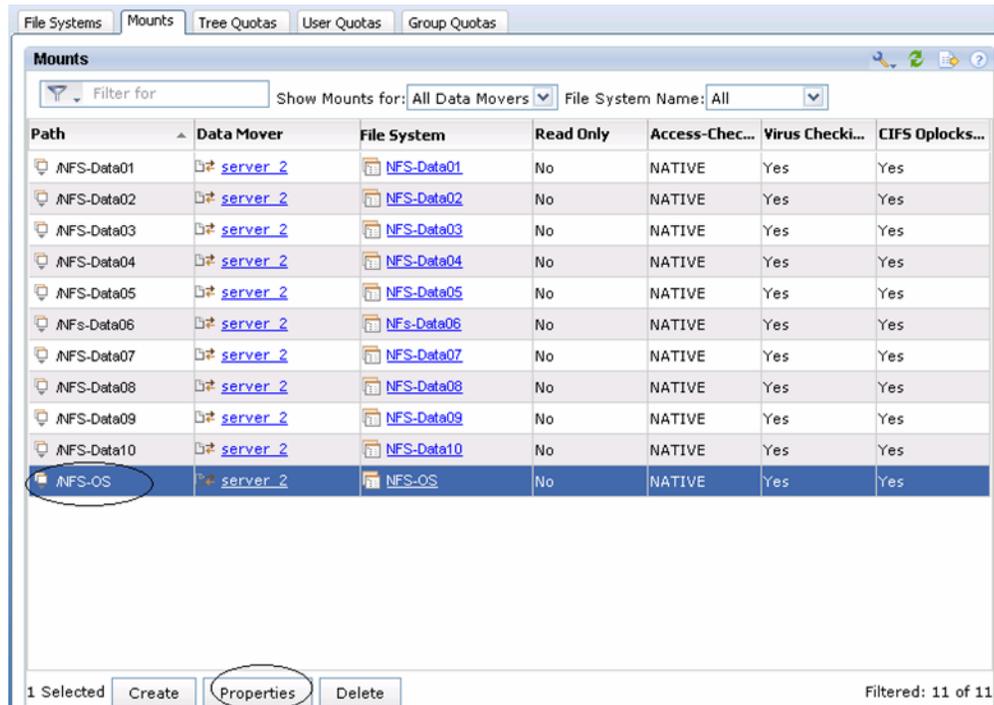
From the Create File System window, click the **Storage Pool** radio button and Specify File System Name as **NFS-DS-1** for Virtual machine datastore. Then, Select Storage Pool from the drop-down list. Specify Storage Capacity as **5 TB**, check the check box **Thin Enabled**, **7340032 MB (7TB)** as Max Capacity, and choose Data Mover as **Server_2**. Click **OK** to create NFS-DS-1 File system.

Figure 213 Details for Creating Storage Pool

Create from		<input checked="" type="radio"/> Storage Pool <input type="radio"/> Meta Volume
File System Name:		NFS-DS-1
Storage Pool:		Pool 3 15.6 TB (16383980 MB)
Storage Capacity:		7 TB
Auto Extend Enabled:		<input checked="" type="checkbox"/>
Thin Enabled:		<input checked="" type="checkbox"/>
High Water Mark:	90	% (Ranges from 50-99; if left blank defaults to 90)
Maximum Capacity (MB):	7340032	Required when thin is enabled.
Slice Volumes:		<input checked="" type="checkbox"/>
Deduplication Enabled:		<input type="checkbox"/>
VMware VAAI nested clone support:		<input type="checkbox"/> (Must be selected at file system creation time)
Data Mover (R/W):		server_2
Mount Point:		<input checked="" type="radio"/> Default <input type="radio"/> Custom

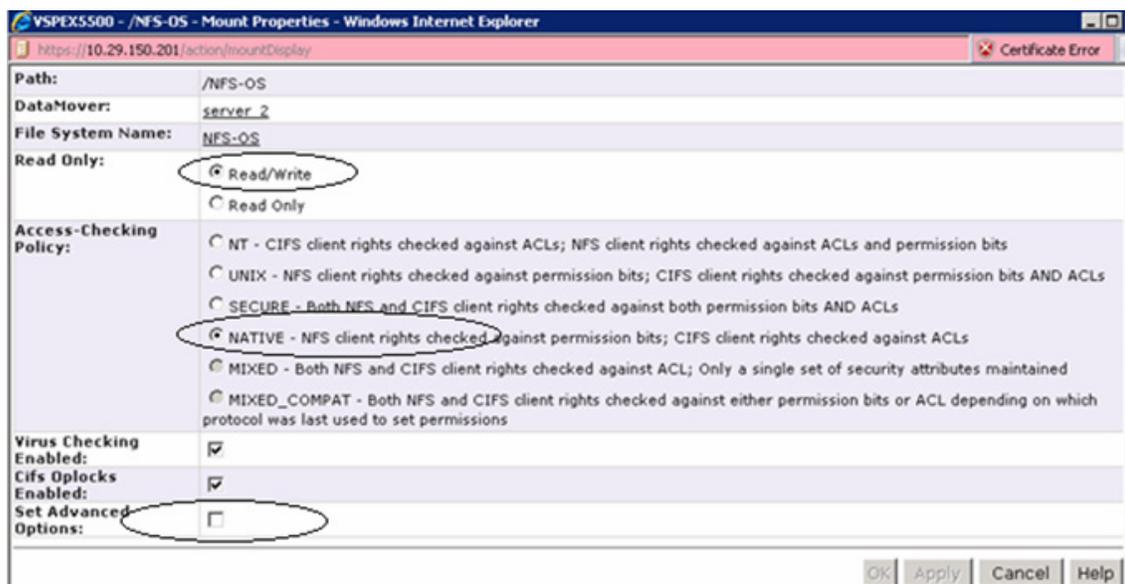
16. Wait until the NFS-DS-1 File system creation process to complete. Verify the process using Background Tasks for File under System menu. Once the NFS-DS-1 is successfully created, repeat steps 15 and 16 for one more NFS file system NFS-DS-2 for the given pool. You will need to create total 10 File systems for 600 VM setup and 16 File Systems for 1000 VM setup.
17. To enable Direct Writes for all the NFS File system. Select **Storage > Storage Configuration > File Systems**. Click the **Mounts** tab, select the path /NFS-OS for the file system NFS-OS and click Properties.

Figure 214 Select /NFS-OS to Enable Direct Writes



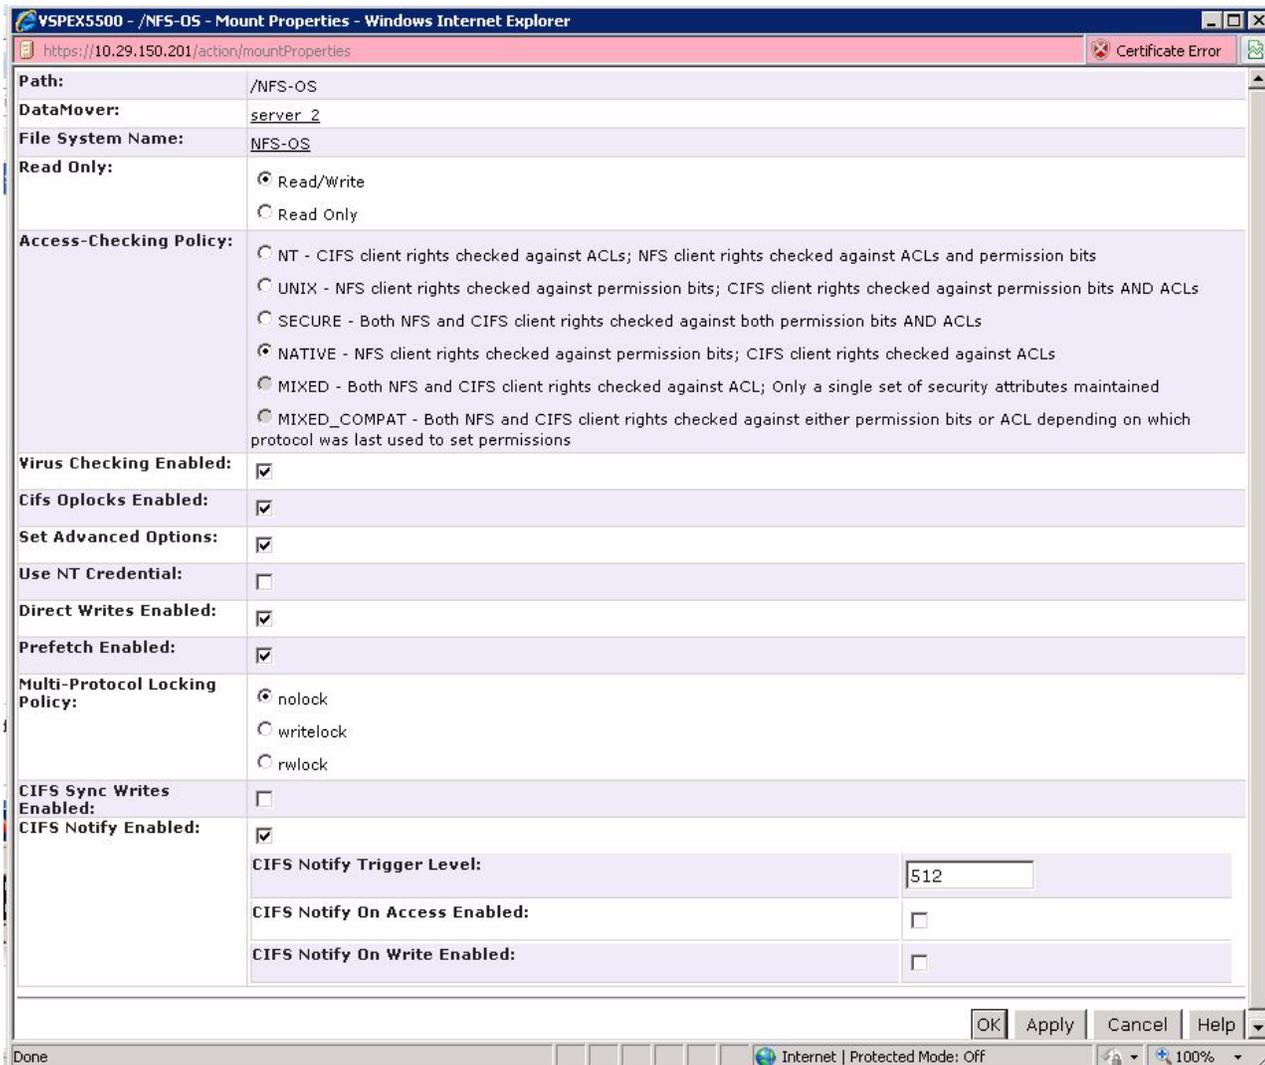
18. From the /NFS-OS mount properties. Make sure the radio buttons **Read/Write** and **Native** for Access policy are selected. Then, check the **Set Advanced Options** check box.

Figure 215 /NFS-OS Mount Properties - Part 1



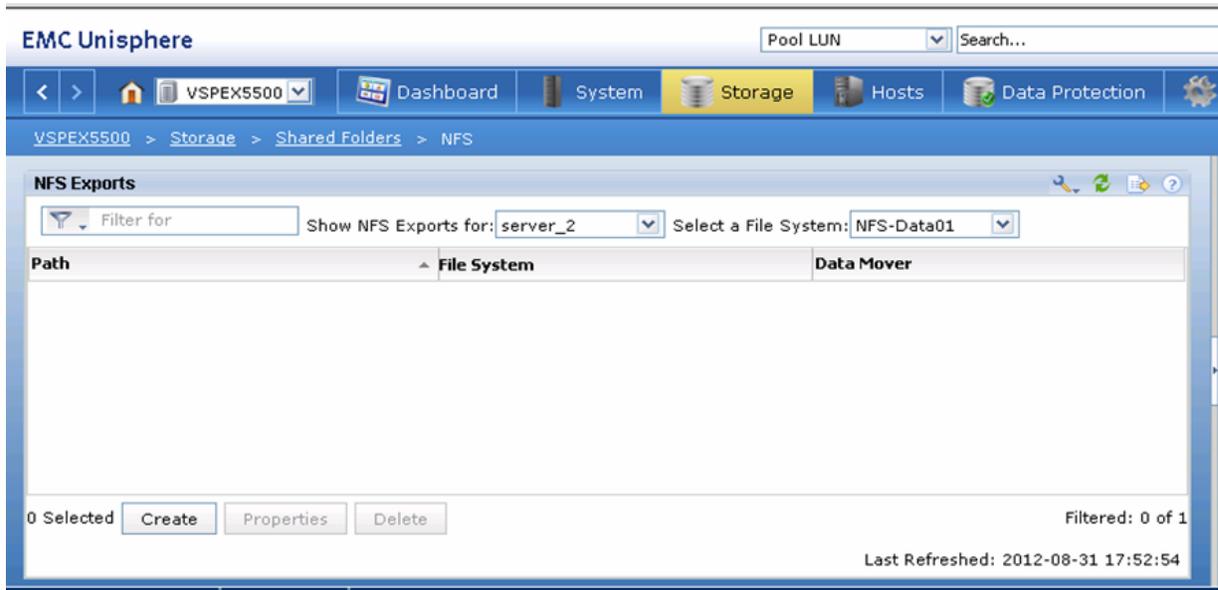
19. Check the **Set Advanced Options** and the **Direct Writes Enabled** check box and click **OK**.

Figure 216 /NFS-OS Mount Properties - Part 2



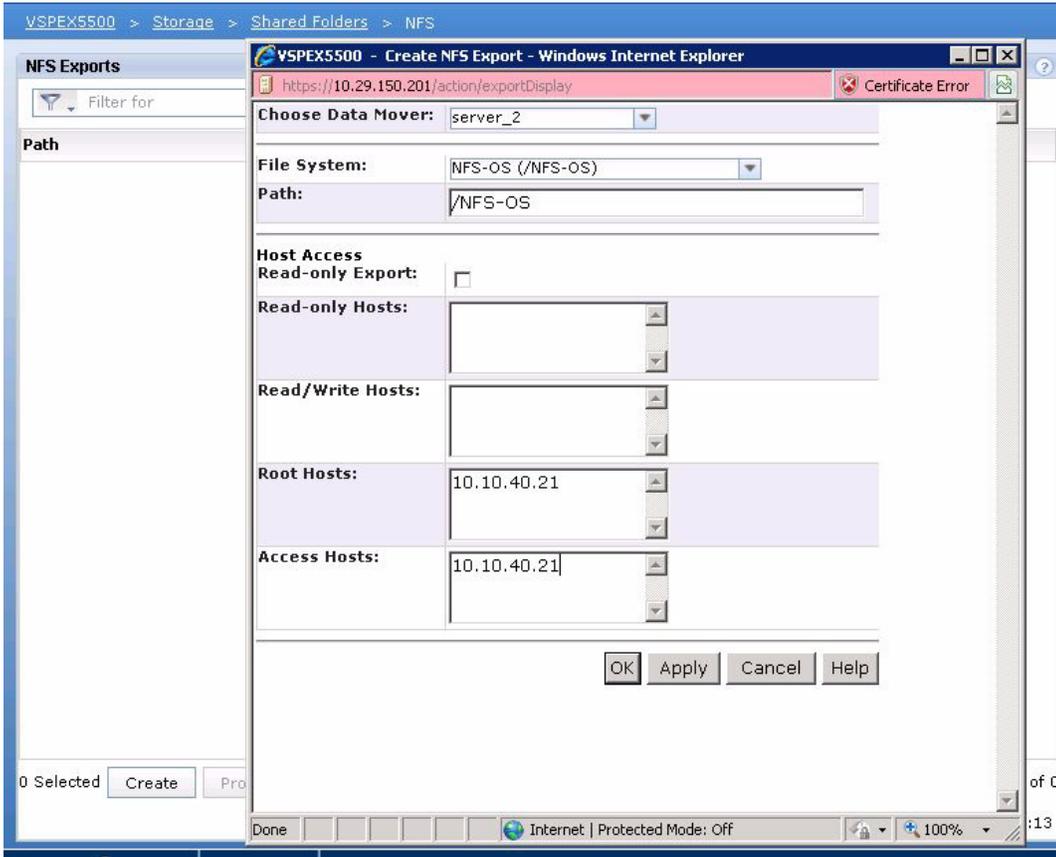
20. Follow the Steps to enable Direct Writes for all the remaining NFS Data file systems.
21. To Create NFS-Exports for all the NFS File systems. Click **Storage > Shared Folders > NFS** and click **Create**.

Figure 217 **Creating NFS Exports**



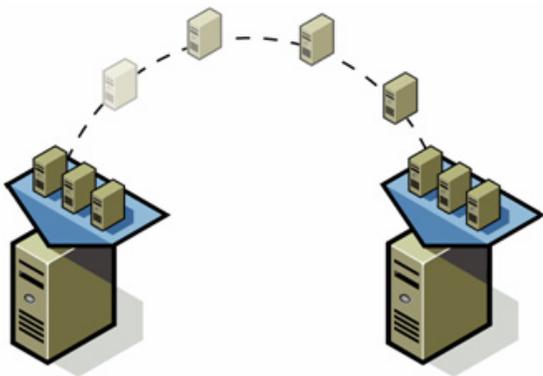
22. Select Data mover as **server-2** and Choose File system as **NFS-OS** and specify Path as **/NFS-OS**. In the Root Hosts and Access Hosts fields, enter the IP address of all the ESXi hosts VMKernel Storage NIC. Separate multiple host vmkernel IP's by : (colon) and click **OK**. Repeat this step for all the NFS File Systems created on the storage array.

Figure 218 Details for Creating NFS Export



Template-Based Deployments for Rapid Provisioning

Figure 219 Rapid Provisioning



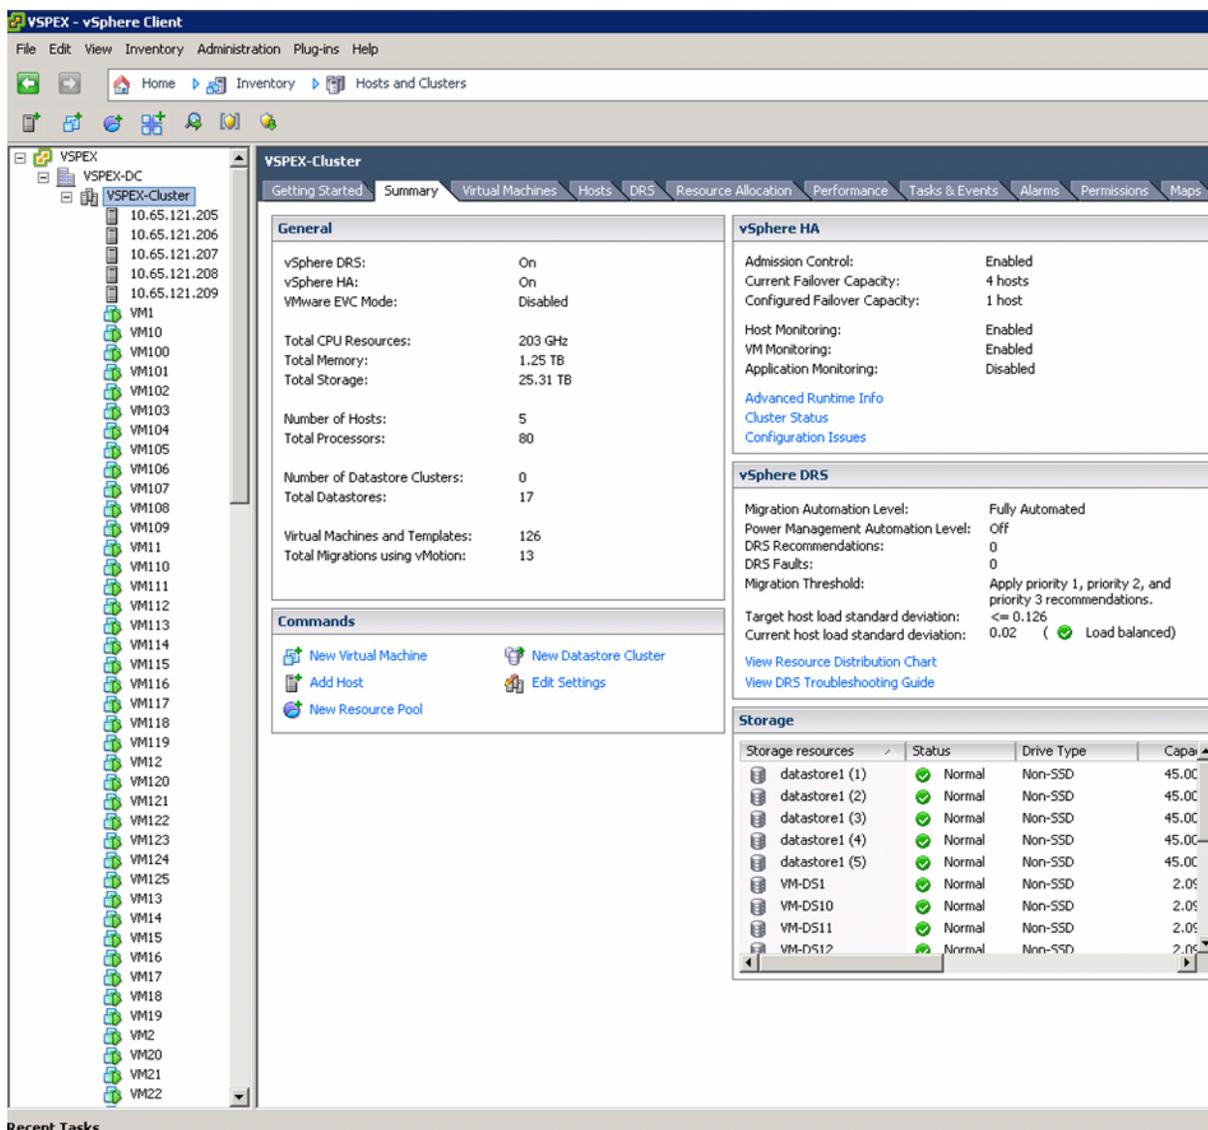
In an environment with established procedures, deploying new application servers can be streamlined, but can still take many hours or days to complete. Not only must you complete an OS installation, but downloading and installing service packs and security updates can add a significant amount of time. Many applications require features that are not installed with Windows by default and must be installed

prior to installing the applications. Inevitably, those features require more security updates and patches. By the time all deployment aspects are considered, more time is spent waiting for downloads and installs than is spent configuring the application.

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

Make sure that the VMs are spread across different VM datastores to properly load-balance the storage usage. The final snapshot of VMs in a cluster looks like [Figure 220](#).

Figure 220 Summary Window Showing VMs in the Cluster in vCenter



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.

- Create a test virtual machine that accesses the datastore and is able to do read/ write operations. Perform the virtual machine migration (vMotion) to a different host on the cluster.
- Perform storage vMotion from one datastore to another datastore and ensure correctness of data.
- During the vMotion of the virtual machine, make sure to have a continuous ping to default gateway and to check if 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. A continuous ping from VM to VM, and vCenter to ESXi hosts should not show significant failures (one or two ping drops might be observed at times, during FI reboot). Also, all the datastores must be visible and accessible from all the hosts at all the time.

1. Administratively shutdown one of the two server ports connected to the Fabric Extender A. Make sure that the connectivity is not affected. Upon administratively enabling the shutdown port, the traffic should be rebalanced. This can be validated by clearing interface counters and showing the counters after forwarding some data from virtual machines on the Nexus switches.
2. Administratively shutdown both server ports connected to Fabric Extender A. ESXi hosts should be able to use fabric B in this case.
3. Administratively shutdown one of the two data links connected to the storage array from FI. Make sure that the storage is still available from all the ESXi hosts. Upon administratively enabling the shutdown of port, the traffic should be rebalanced. Repeat this step for each link connected to the Storage Processors one after another.
4. Reboot one of the two Cisco UCS Fabric Interconnects while storage and network access from the servers are up. The switch reboot should not affect the operations of storage and network access from the VMs. Upon rebooting the FI, the network access load should be rebalanced across the two fabrics.
5. Reboot the active storage processor of the VNX storage array and make sure that all the datastores are still accessible during and after the reboot of the storage processor.
6. Fully load all the virtual machines of the solution. Put one of the ESXi host in maintenance mode. All the VMs running on that host should be migrated to other active hosts. No VM should lose any network or storage accessibility during or after the migration. This test assumes that enough RAM is available on active ESXi hosts to accommodate VMs from the host put in maintenance mode.
7. Reboot the host in maintenance mode, and remove it from the maintenance mode; this should re-balance the VM distribution across the cluster.

Cisco Validation Test Profile

“vdbench” testing tool was used with Windows Server 2012 to test scaling of the solution in Cisco labs. [Table 11](#) details on the test profile used.

Table 11 *Test Profile Details*

Profile characteristic	Value
Number of virtual machines	300, 600, or 1000
Virtual machine OS	Windows Server 2012
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

Bill of Material

[Table 12](#) gives the list of the components used in this CVD. The number of actual Cisco UCS servers will vary depending on the size of the architecture. For actual number of servers and chassis, see [Table 3](#).

Table 12 *List of Hardware Components Used in the CVD*

Description	Part #
Cisco UCS C220M3 Rack Servers	UCSC-C220-M3S
Cisco UCS B200M3 Blade Server	UCSB-B200-M3
CPU for Cisco UCS Servers (2 per server)	UCS-CPU-E52650B
Memory for Cisco UCS Servers (8 per server)	UCS-MR-1X162RY-A
Cisco UCS 1225 VIC Adapter (1 per rack server)	UCSC-PCIE-CSC-02
Cisco UCS 1240 VIC Adapter (1 per blade server)	UCSB-MLOM-40G-01
Cisco UCS 2232PP Fabric Extenders (2)	N2K-C2232PP-10GE
Cisco UCS 2208XP Fabric Extenders (2)	UCS-IOM-2208XP
Cisco UCS 6248UP Fabric Interconnects (2)	UCS-FI-6248UP
Cisco UCS Nexus 5548UP Switches (2)	N5K-C5548UP-FA
10 Gbps SFP+ multifiber mode	SFP-10G-SR

Customer Configuration Data Sheet

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

Table 13 **Common Server Information**

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

Table 14 **ESXi Server Information**

Server Name	Purpose	Management	Private Net (storage) addresses	vMotion IP
	ESXi Host 1			
	ESXi Host 2			

Table 15 **Array Information**

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

Table 16 Network Infrastructure Information

Description	IP	Subnet Mask	Default Gateway
UCS Manager Virtual IP address			
UCS Fabric Interconnect A address			
UCS Fabric Interconnect B address			
N5k A management IP address			
N5k B management IP address			
N1kv management IP address			

Table 17 VLAN Information

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

Table 18 VSAN Information

Name	Network Purpose	VSAN ID	Allowed Subnets
Storage	Storage access		

Table 19 Service Accounts

Account	Purpose	Password (optional, secure appropriately)
	UCS Manager administrator	
	N5k switches administrator	
	N1kv switch administrator	
	Windows Server administrator	
Root	ESXi root	
	Array administrator	
	vCenter administrator	
	SQL Server administrator	

References

Cisco UCS:

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

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>