

Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.5 for up to 1000 Virtual Machines

Deployment Guide for VSPEX with Cisco UCS B200 M4 and C220 M4 Servers, Cisco Nexus 9000 Series Switches, Cisco MDS Switches, EMC VNX5400 and VMware vSphere 5.5


Last Updated: December 18, 2014



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors

Vijay Durairaj, Technical Marketing Engineer, Cisco Systems, Inc.



Vijay Durairaj is a Technical Marketing Engineer with Cisco UCS Solutions and Performance Group. Vijay has over 10 years of experience in compute, network, storage and server virtualization design. Vijay has delivered solutions for Server and Desktop virtualization and has extensive experience with Cisco Unified Computing System, Cisco Nexus products and Storage technologies. Vijay has worked on performance and benchmarking with Cisco UCS servers. Vijay is a VMware Certified Professional and Cisco Unified Computing systems Design specialist.

Acknowledgments

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

- Bathu Krishnan, Cisco Systems
- Shiva Shastri, Cisco Systems
- Kevin Phillips, EMC

About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit

<http://www.cisco.com/go/designzone>.

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

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

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website 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. (0809R)

© 2014 Cisco Systems, Inc. All rights reserved



Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.5 for up to 1000 Virtual Machines

Summary

The Cisco solution for EMC VSPEX is a validated modular architecture built with proven best-of-breed technologies. The complete end-to-end solutions enable you to make informed decisions while choosing the hypervisor, compute, storage and networking components. VSPEX drastically reduces time and effort around server virtualization planning and configuration. VSPEX infrastructure accelerates your IT transformation, while providing greater flexibility of choice, and efficiency at lower risk. This Cisco Validated Design (CVD) focuses on the VSPEX VMware architecture for mid-market businesses with a required load of less than 1000 Virtual Machines.

Introduction

Virtualization is a key deployment model for achieving better utilization of underlying system resources leading to reduced Total Cost of Ownership (TCO). However, choosing the appropriate platform for virtualization can be a complex task. Platforms should be flexible, reliable and cost effective to facilitate deployment of various enterprise applications while also providing the means to scale and manage with ease. An architecture that allows for resource sharing across points of delivery (PoD) is desirable to address the issue of stranded capacity, both within and external to the integrated stack. In this regard, Cisco's solution implemented under EMC's VSPEX program provides a comprehensive and validated infrastructure platform for virtual machine (VM) deployment to suite any customers' needs.

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, Cisco Unified Computing System (UCS) and Cisco Unified Computing Systems Manager (UCSM). External references are provided where applicable and it is recommended that the reader be familiar with these documents.



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

Purpose of this Document

This document details the steps required to deploy and configure a Cisco solution on EMC VSPEX using VMware 5.5 as hypervisor. Cisco's validation is a confirmation of component compatibility, connectivity and correct operation of the entire integrated stack. This document covers the VMware architectures for Small-to-Medium-sized businesses, typically requiring about 1000 VMs. This document highlights two variants of the solution:

- FC-Variant—EMC VNX series storage array directly attached to the Cisco UCS Fabric Interconnects (FIs) using Fibre-Channel (FC) for storage access.
- NFS-Variant—EMC VNX series storage array using NFS for storage access through a pair of Cisco Nexus 9000 series switches operating in standalone mode.

These two architectures are referred to as "FC Variant" and "NFS Variant" throughout this document. While readers of this document are expected to have sufficient knowledge to install and configure the products used, configuration details that are important to the deployment of this solution are specifically mentioned.

Business Needs

VSPEX solutions are built with proven best-of-breed technologies and this Cisco Validated Design provides guidance to create complete solutions on VSPEX that enable you to make informed decisions about an application ready platform. VSPEX infrastructure accelerates IT transformation through faster deployments, while ensuring greater flexibility of choice and efficiency at a lower risk.

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

- Provide an end-to-end virtualization solution to take full advantage of unified infrastructure components.
- Provide a Cisco VSPEX for VMware ITaaS solution for an efficient virtual environment catering to various customer use cases.
- Show implementation progression of a VMware vCenter 5.5 design with results.
- Provide a reliable, flexible and scalable reference design.

Solution Overview

The Cisco solution for EMC VSPEX with VMware 5.5 provides an end-to-end architecture with Cisco, EMC, VMware, and Microsoft technologies that demonstrate support for up to 1000 generic virtual machines with high availability and server redundancy.

The components used for the design and deployment are as follows:

- Cisco Unified Compute System (UCS) 2.2(3a)

- Cisco Unified Computing System B-series or C-series servers, as per customer choice
- Cisco UCS VIC adapters
- Cisco Nexus 9396PX switches
- Cisco MDS 9148S Switches
- Cisco Nexus 1000v virtual switch
- EMC VNX5400, VNX5600, or VNX5800 storage components based on customer needs
- VMware vCenter 5.5
- Microsoft SQL 2012 database (for vCenter & Update Manager)
- VMware DRS
- VMware HA
- Microsoft Windows 2012 Datacenter (for Infra. Server and VSPEX virtual machines)

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

Technology Overview

Cisco Unified Computing System

The Cisco Unified Computing System is a state-of-the-art data center platform that unites computing, network, storage access, and virtualization into a single cohesive system.

The main components of the Cisco Unified Computing System are:

- **Computing**—The system is based on an entirely new class of computing system that incorporates rack mount and blade servers based on Intel Xeon Processor E5-2600 v3. The Cisco UCS servers offer the patented Cisco Extended Memory Technology to support applications with large datasets and allow more virtual machines per server.
- **Network**—The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Virtualization**—The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage Access**—The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet (NFS or iSCSI), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This provides customers with 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 Nexus 9396PX Switch

The Cisco Nexus 9000 family of switches supports two modes of operation: NxOS standalone mode and Application Centric Infrastructure (ACI) fabric mode. In standalone mode, the switch performs as a typical Cisco Nexus switch with increased port density, low latency and 40G connectivity. In fabric mode, the administrator can take advantage of Cisco ACI

The Cisco Nexus 9396PX (Figure 1) is a two rack-unit switch which delivers a comprehensive line-rate layer 2 and layer 3 features in a two-rack-unit form factor. It supports line rate 1/10/40 GE with 960 Gbps of switching capacity. It is ideal for top-of-rack and middle-of-row deployments in both traditional and Cisco Application Centric Infrastructure (ACI)-enabled enterprise, service provider, and cloud environments.

Specifications At-a-Glance

- Forty-eight 1/10 Gigabit Ethernet Small Form-Factor Pluggable (SFP+) non-blocking ports
- Twelve 40 Gigabit Ethernet Quad SFP+ (QSFP+) non-blocking ports
- Low latency (approximately 2 microseconds)
- 50 MB of shared buffer
- Line rate VXLAN bridging, routing, and gateway support
- Fibre Channel over Ethernet (FCoE) capability
- Front-to-back or back-to-front airflow

Figure 1 Cisco Nexus 9396PX Switch



Cisco MDS 9148S 16G Multilayer Fabric Switch

The Cisco MDS 9148S 16G Multilayer Fabric Switch (Figure 2) is the next generation of the highly reliable Cisco MDS 9100 Series Switches. It includes up to 48 auto-sensing line-rate 16-Gbps Fibre Channel ports in a compact easy to deploy and manage 1-rack-unit (1RU) form factor. In all, the Cisco MDS 9148S is a powerful and flexible switch that delivers high performance and comprehensive Enterprise-class features at an affordable price.

Features and Capabilities

- Flexibility for growth and virtualization
- Optimized bandwidth utilization and reduced downtime
- Enterprise-class features and reliability at low cost

- PowerOn Auto Provisioning and intelligent diagnostics
- In-Service Software Upgrade and dual redundant hot-swappable power supplies for high availability
- High-performance interswitch links with multipath load balancing

Specifications At-a-Glance

- Performance and Port Configuration
 - 2/4/8/16-Gbps auto-sensing with 16 Gbps of dedicated bandwidth per port
 - Up to 256 buffer credits per group of 4 ports (64 per port default, 253 maximum for a single port in the group)
 - Supports configurations of 12, 24, 36, or 48 active ports, with pay-as-you-grow, on-demand licensing.

Figure 2 *Cisco MDS9148S 16G Fabric Switch*



Cisco UCS Manager

Cisco Unified Computing System (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 a unified management domain with centralized management capabilities and controls multiple chassis and thousands of virtual machines.


Fabric Interconnect

These devices provide a single point for connectivity and management for the entire system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency regardless of a server or virtual machine's topological location in the system.

Cisco UCS 6248UP Fabric Interconnect

Cisco UCS 6200 Series Fabric Interconnects ([Figure 3](#)) support the system's 10-Gbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic using a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blade, rack servers, and virtual machines are interconnected using the same mechanisms. The Cisco UCS 6248UP is a 1-RU Fabric Interconnect that features up to 48 universal ports that can support 10 Gigabit Ethernet, Fibre Channel over Ethernet, or native Fibre Channel connectivity. The Cisco UCS 6296UP packs 96 universal ports into only two rack units.

Figure 3 *Cisco UCS 6248UP Fabric Interconnect*



Cisco UCS 2208XP Fabric Extender

The Cisco UCS 2208XP Fabric Extender ([Figure 4](#)) has eight 10 Gigabit Ethernet, FCoE-capable, Enhanced Small Form-Factor Pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.

Figure 4 *Cisco UCS 2208 XP*



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 2208XP Fabric Extenders.

A passive mid-plane supports up to 2x 40 Gbit Ethernet links to each half-width blade slot or up to 4x 40 Gbit links to each full-width slot. It provides 8 blades with 1.2 terabits (Tb) of available Ethernet throughput for future I/O requirements.



Note

The Cisco UCS 6324 FI supports only 512 Gbps.

The chassis is capable of supporting future 80 Gigabit Ethernet standards. The Cisco UCS Blade Server Chassis is shown in [Figure 5](#).

Figure 5 *Cisco Blade Server Chassis (front and back view)*



Cisco UCS B200 M4 Blade Server

Optimized for data center or cloud, the Cisco UCS B200 M4 Blade Server (Figure 6) can quickly deploy stateless physical and virtual workloads, with the programmability of the Cisco UCS Manager and simplified server access of SingleConnect technology. The Cisco UCS B200 M4 is built with the Intel® Xeon® E5-2600 v3 processor family, up to 768 GB of memory (with 32 GB DIMMs), up to two drives, and up to 80 Gbps total bandwidth. It offers exceptional levels of performance, flexibility, and I/O throughput to run the most demanding applications.

In addition, Cisco Unified Computing System has the architectural advantage of not having to power and cool switches in each blade chassis. Having a larger power budget available for blades allows Cisco to design uncompromised expandability and capabilities in its blade servers.

The Cisco UCS B200 M4 Blade Server delivers the following:

- Suitability for a wide range of applications and workload requirements
- Highest-performing CPU and memory options without constraints in configuration, power or cooling
- Half-width form factor offering industry-leading benefits
- Latest features of Cisco UCS Virtual Interface Cards (VICs)

Figure 6 *Cisco UCS B200 M4 Blade Server*



Cisco C220 M4 Rack-Mount Servers

The Cisco UCS C220 M4 Rack-Mount Server (Figure 7) is the most versatile, high-density, general-purpose enterprise infrastructure and application server in the industry today. It delivers world-record performance for a wide range of enterprise workloads, including virtualization, collaboration, and bare-metal applications.

The enterprise-class Cisco UCS C220 M4 rack-mount server extends the capabilities of the Cisco Unified Computing System portfolio in a one rack-unit (1RU) form-factor. It provides the following:

- Dual Intel® Xeon® E5-2600 v3 processors for improved performance suitable for nearly all 2-socket applications
- Next-generation double-data-rate 4 (DDR4) memory and 12 Gbps SAS throughput

- Innovative Cisco UCS virtual interface card (VIC) support in PCIe or modular LAN on motherboard (MLOM) form factor

The Cisco UCS C220 M4 rack-mount server also offers maximum reliability, availability, and serviceability (RAS) features, including the following:

- Tool-free CPU insertion
- Easy-to-use latching lid
- Hot-swappable and hot-pluggable components
- Redundant Cisco Flexible Flash SD cards .

In Cisco UCS managed operations, the Cisco UCS C220 M4 rack-mount server takes advantage of our standards-based unified computing innovations to significantly reduce the customers' TCO and increase business agility.

Figure 7 *Cisco C220 M4 Rack-Mount Server*




Cisco I/O Adapters for Blade and Rack-Mount Servers

Cisco VIC 1340 Virtual Interface Card

The Cisco UCS Blade Server has various Converged Network Adapters (CNA) options. The Cisco UCS VIC 1340 Virtual Interface Card (VIC) option is used in this Cisco Validated Design.


The Cisco UCS Virtual Interface Card (VIC) 1340 ([Figure 8](#)) is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

Figure 8 *Cisco UCS 1340 VIC Card*



The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 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 1340 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 and management.

Figure 9 *Cisco UCS VIC 1340 Architecture*



Cisco VIC 1227 Virtual Interface Card

The Cisco UCS Rack-mount Server has various Converged Network Adapters (CNA) options. The Cisco UCS 1227 Virtual Interface Card (VIC) is option is used in this Cisco Validated Design.

The Cisco UCS Virtual Interface Card (VIC) 1227 is a dual-port Enhanced Small Form-Factor Pluggable (SFP+) 10-Gbps Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCI Express (PCIe) modular LAN-on-motherboard (mLOM) adapter designed exclusively for Cisco UCS C-Series Rack-Mount Servers. New to Cisco rack-mount servers, the mLOM slot can be used to install a Cisco VIC without consuming a PCIe slot, which provides greater I/O expandability. It incorporates next-generation converged network adapter (CNA) technology from Cisco, providing investment protection for future feature releases. The card 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 1227 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 10 *Cisco UCS VIC 1227 Card*



Cisco Nexus 1000v Virtual Switch

The Cisco Nexus 1000V Series Switches are virtual machine access switches for the VMware vSphere environments running the Cisco NX-OS operating system. Operating inside the VMware ESX or ESXi hypervisors, the Cisco Nexus 1000V Series provides:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Non-disruptive operational model for your server virtualization and networking teams

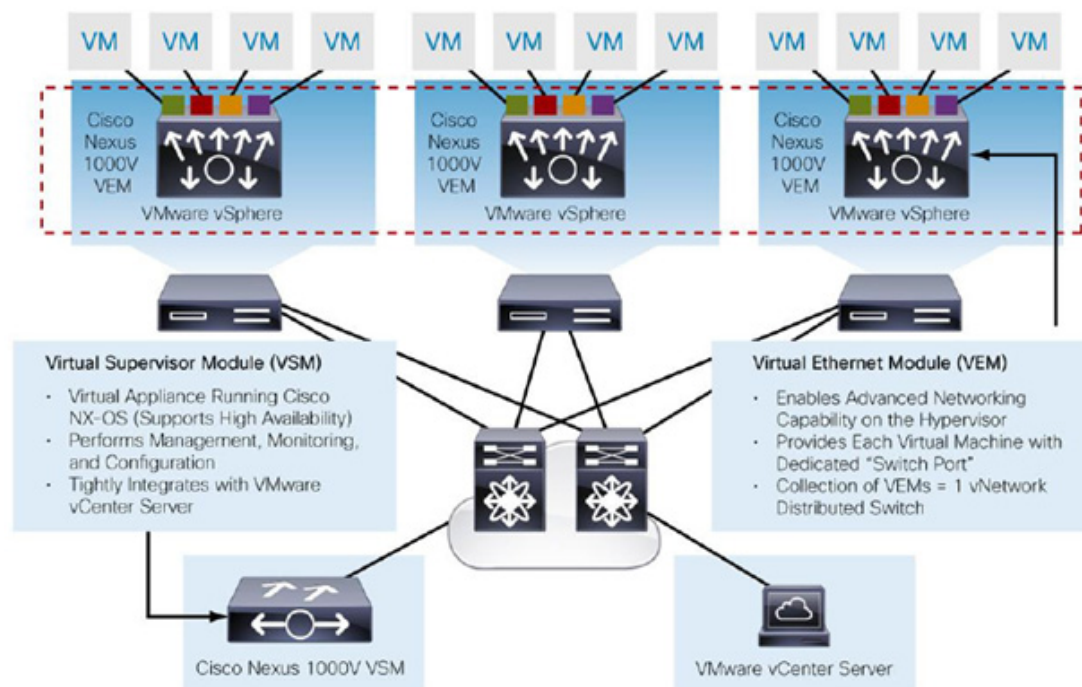
When server virtualization is deployed in the data center, virtual servers typically are not managed the same way as physical servers. Server virtualization is treated as a special deployment, leading to longer deployment time, with a greater degree of coordination among server, network, storage, and security administrators. With the Cisco Nexus 1000V Series, you can have a consistent networking feature set and provisioning process all the way from the virtual machine access layer to the core of the data center network infrastructure. Virtual servers can now use the same network configuration, security policy, diagnostic tools, and operational models as their physical server counterparts attached to dedicated physical network ports. Virtualization administrators can access predefined network policies that follow mobile virtual machines to ensure proper connectivity saving valuable time to focus on virtual machine administration. This comprehensive set of capabilities helps you to deploy server virtualization faster and realize its benefits sooner.

Cisco 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 11 Cisco Nexus 1000v virtual Distributed Switch Architecture



Cisco UCS Differentiators

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

- **Embedded management:** In Cisco Unified Computing System, 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 eight blade servers; thus giving enormous scaling on management plane.
- **Unified Fabric:** In Cisco Unified Computing System, 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.
- **Auto Discovery:** By inserting the blade server in the chassis or connecting the 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 wire-once architecture of Cisco Unified Computing System, where compute capability of Cisco Unified Computing System can extend easily while keeping the existing external connectivity to LAN, SAN and management networks.
- **Policy based resource classification:** When a compute resource is discovered by Cisco 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 presents the policy-based resource classification of Cisco UCS Manager.

- **Combined Rack and Blade Server Management:** Cisco UCS Manager can manage B-series blade servers and C-series rack-mount server under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic. This CVD presents a combination of B and C series servers to demonstrate stateless and form factor independent computing work load.
- **Model-based Management Architecture:** Cisco UCS Manager architecture and management database is model based and data driven. Open, standard-based XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management system, such as VMware vCloud director, Microsoft system center, and Citrix CloudPlatform.
- **Policies, Pools, Templates:** Management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- **Loose Referential Integrity:** In Cisco 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 flexibilities where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.
- **Policy Resolution:** In Cisco 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 other policy by name is resolved in the org hierarchy with closest policy match. If no policy with specific name is found in the hierarchy till root org, then special policy named "default" is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibilities to owners of different orgs.
- **Service Profiles and Stateless Computing:** Service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- **Built-in Multi-tenancy Support:** Combination of policies, pools and templates, loose referential integrity, policy resolution in org hierarchy and service profile based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
- **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 offloads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.
- **Simplified QoS:** Even though fibre-channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

VMware vSphere 5.5

VMware vSphere 5.5 is a next-generation virtualization solution from VMware which builds upon ESXi 5.1 and provides greater levels of scalability, security, and availability to virtualized environments. vSphere 5.5 offers improvements in performance and utilization of CPU, memory, and I/O. It also offers users the option to assign up to thirty two virtual CPU to a virtual machine-giving system administrators more flexibility in their virtual server farms as processor-intensive workloads continue to increase.

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

EMC Storage Technologies and Benefits

This architecture has two variants:

- FC-Variant of the solution, where EMC VNX series storage devices are attached to Cisco UCS Fabric Interconnects directly, accessing storage over Fibre Channel protocol.
- NFS-Variant of the solution, where EMC VNX series storage devices are accessed from Cisco UCS Fabric Interconnects thru a pair of Nexus 9000 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 platforms that can support iSCSI, NFS, Fibre Channel, and CIFS protocols depending on the customer's specific needs. This solution was validated using NFS for data storage of Virtual Machines and Fibre Channel for hypervisor SAN boot. VNX series storage arrays have the following customer benefits:

- Next-generation unified storage, optimized for virtualized applications
- Capacity optimization features including compression, deduplication, thin provisioning, and application-centric copies
- High availability, designed to deliver five 9s availability
- Multiprotocol support for file and block
- Simplified management with EMC Unisphere™ for a single management interface for all network-attached storage (NAS), storage area network (SAN), and replication needs

Software suites available:

- 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 available:

- Total Value Pack - Includes all protection software suites and the Security and Compliance Suite

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 - translating into storage, bandwidth and operational savings.

The following are two of the most common recovery requests made to backup administrators:

- 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 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 Cisco UCS 2.2(3a) with combined B-series and C-series servers with VNX5400 directly attached to Cisco UCS fabric interconnect. NFS-Variant of the architecture uses UCS 2.2(3a) with B-series and C-series servers with VNX5600 or VNX5800 storage array attached to the Cisco Nexus 9396PX switches. In both variants, the Cisco UCS C220 M4 rack-mount servers are connected directly to Cisco UCS Fabric Interconnect with single-wire management feature. VMware vSphere 5.5 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 the VMware Architectures*

Vendor	Name	Version	Description
Cisco	UCSM	2.2(3a)	UCS Manager
Cisco	Nexus 9396PX	6.1(2)I2(2a)	Cisco Nexus 9396x switches (NFS Variant only)
Cisco	MDS 9148S	6.2(9)	Cisco MDS 9148S 16G Multi Fabric switches (NFS Variant only)
Cisco	UCS 6248UP FI	5.2(3)N2(2.23a)	UCS Fabric Interconnects
Cisco	UCS 5108 Chassis	NA	UCS Blade server chassis
Cisco	UCS 2208XP FEX	2.2(3a)	UCS Fabric Extenders for Blade Server chassis
Cisco	UCS B200 M4 servers	2.2(3a)	Cisco B200 M4 blade servers
Cisco	UCS VIC 1340	4.0(1b)	Cisco VIC 1340 adapters
Cisco	UCS C220 M4 servers	2.0.3e – CIMC C220M4.2.0.3c - BIOS	Cisco C220 M4 rack servers
Cisco	UCS VIC 1227	4.0(1b)	Cisco UCS VIC adapter
Cisco	Nexus 1000v	5.2.1.SV3.1.2	Cisco Nexus 1000v virtual switch.
EMC	VNX5400	VNX Block OE 05.33	EMC VNX storage array (FC-variant only)
EMC	VNX5600	VNX File OE 8.1 VNX Block OE 05.33	EMC VNX storage array (NFS-variant only)
EMC	VNX5800	VNX File OE 8.1 VNX Block OE 05.33	EMC VNX storage array (NFS-variant only)
EMC	Avamar	6.1 SP1	EMC data backup software
EMC	Data Domain OS	5.2	EMC Data domain operating system
VMware	ESXi 5.5	5.5.0 Build 1331820	Hypervisor
VMware	vCenter Server	5.5.0 Build 1312299	VMware Management
Microsoft	Windows Server 2012	2012 Datacenter	Operating System to host vCenter server & VSPEX VMs.
Microsoft	SQL Server	2012	Database Server SQL Enterprise Edition for vCenter

Table 2 *Outline of the Cisco UCS B200 M4 or C220 M4 Blade Server Configuration (per server basis) Across the Two Variants of VMware Architectures*

Component	Capacity
Memory (RAM)	128 GB (16 x 16GB DDR4 DIMMs)
Processor	2 x Intel® Xeno® E5-2660 v3 CPUs, 2.6 GHz, 10 cores, 20 threads
Local Storage	Cisco FlexStorage 12G SAS RAID controller

Both variants of the 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 Cisco UCS B-Series or Cisco UCS C-Series servers and storage array type change depending on the number of virtual machines. [Table 3](#) highlights the change in hardware components, as required by a different scale. Typically, 60 reference virtual machines are deployed per server.

Table 3 *Hardware Components for Different Scale*

Component	Up to 300 VSPEX VMs	Up to 600 VSPEX VMs	Up to 1000 VSPEX VMs
Servers	5 x Cisco C220 M4 or B200 M4 servers	10 x Cisco B200 M4 or C220 M4 servers	15 x Cisco B200 M4 or C220 M4 servers
Blade servers 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 12 Reference Architecture for FC Variant

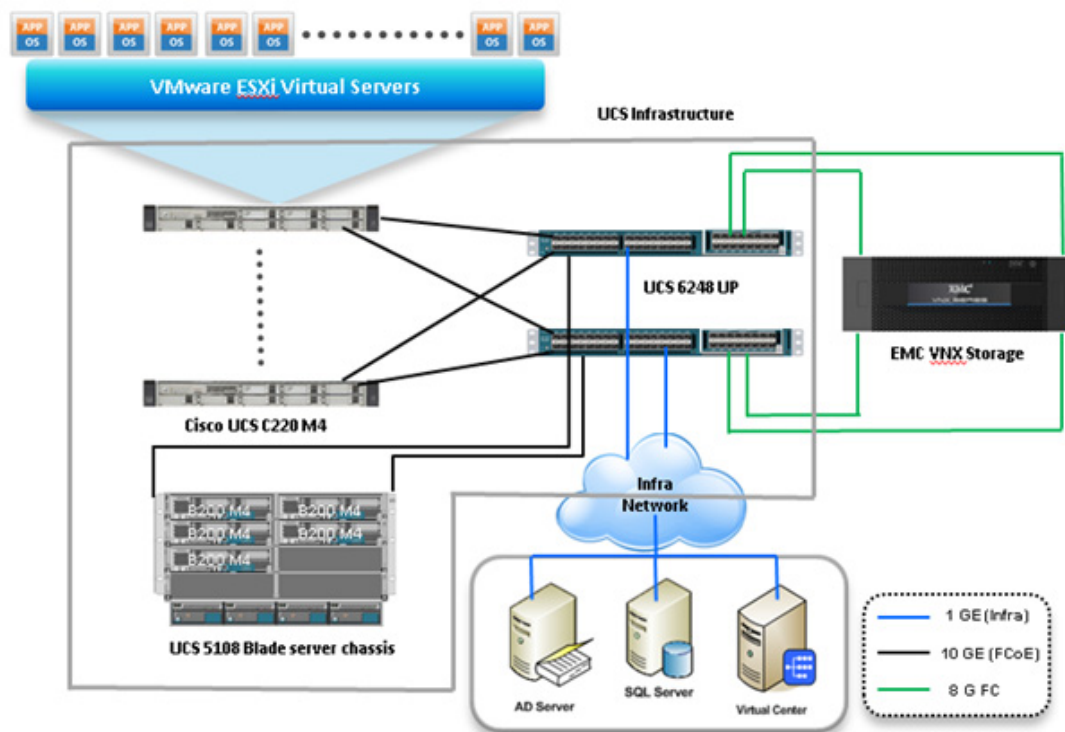
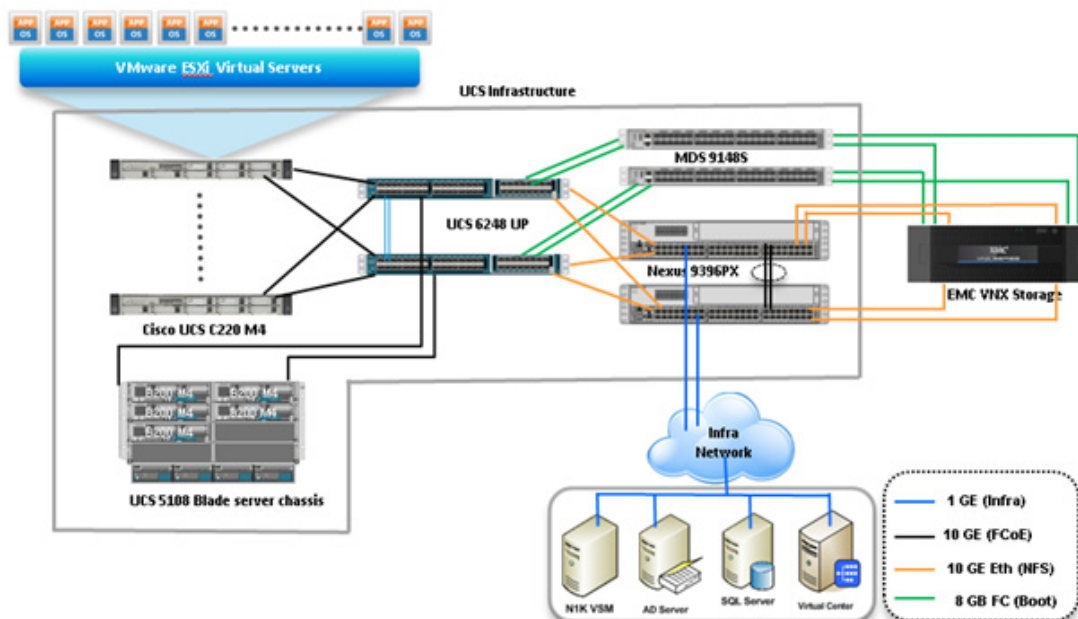


Figure 13 Reference Architecture for NFS Variant



The following are key design points of the VMware architecture for a mid-market segment:

- For smaller scale, storage array dedicated to a given Cisco UCS domain is preferable for simplicity, so VNX5400 storage is directly attached to Cisco UCS FIs. For larger scales, multiple Cisco UCS domains may share a storage, as well as a Cisco UCS domain may want to use multiple storage

arrays. In that case, NFS storage access thru Cisco Nexus 9396 is preferable for VM Datastore and FC storage access thru MDS 9148S is preferable for ESXi SAN boot, as shown in NFS-variant of 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, this design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

ESXi 5.5 is used as hypervisor operating system on each server and is installed on SAN LUNs in both architectures. However the virtual machines storage is accessed thru FC or NFS protocols depending on the architecture. Typical load is 60 reference virtual machines per server.

Memory Configuration Guidelines

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

ESXi/ESXi Memory Management Concepts

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

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

For more information about vSphere memory management concepts, refer to the VMware vSphere Resource Management Guide at

http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf

Virtual Machine Memory Concepts

Figure 14 illustrates the use of memory settings parameters in the virtual machine.

Figure 14 *Virtual Machine Memory Settings*



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

- Configured memory—Memory size of virtual machine assigned at creation.
- Touched memory—Memory actually used by the virtual machine. vSphere allocates only guest operating system memory on demand.
- Swappable—Virtual machine memory that can be reclaimed by the balloon driver or by vSphere swapping. Ballooning occurs before vSphere swapping. If this memory is in use by the virtual machine (that is, touched and in use), the balloon driver causes the guest operating system to swap. Also, this value is the size of the per-virtual machine swap file that is created on the VMware Virtual Machine File System (VMFS) file system (VSWP file). If the balloon driver is unable to reclaim memory quickly enough, or is disabled or not installed, vSphere forcibly reclaims memory from the virtual machine using the VMkernel swap file.

Allocating Memory to Virtual Machines

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

Table 4 *Resources used by a Single Virtual Machine*

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

The 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 gigabytes of memory may consume about 100 megabytes of memory overhead, where a virtual machine with two virtual CPUs and 32 gigabytes of memory may consume approximately 500 megabytes of memory overhead. This memory overhead is in addition to the memory allocated to the virtual machine and must be available on the ESXi host.
- "Right-size" memory allocations—Over-allocating memory to virtual machines can waste memory unnecessarily, but it can also increase the amount of memory overhead required to run the virtual machine, thus reducing the overall memory available for other virtual machines. Fine-tuning the memory for a virtual machine is done easily and quickly by adjusting the virtual machine properties. In most cases, hot-adding of memory is supported and can provide instant access to the additional memory if needed.
- Intelligently overcommit—Memory management features in vSphere allow for overcommitment of physical resources without severely impacting performance. Many workloads can participate in this type of resource sharing while continuing to provide the responsiveness users require of the application. When looking to scale beyond the underlying physical resources, consider the following:

- Establish a baseline before overcommitting—Note the performance characteristics of the application before and after. Some applications are consistent in how they utilize resources and may not perform as expected when vSphere memory management techniques take control. Others, such as Web servers, have periods where resources can be reclaimed and are perfect candidates for higher levels of consolidation.
- Use the default balloon driver settings—The balloon driver is installed as part of the VMware Tools suite and is used by ESXi/ESX if physical memory comes under contention. Performance tests show that the balloon driver allows ESXi/ESX to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESXi/ESX to use host-swapping to make up for the lack of available physical memory which adversely affects performance.
- Set a memory reservation for virtual machines that require dedicated resources—Virtual machines running Search or SQL services consume more memory resources than other application and Web front-end virtual machines. In these cases, memory reservations can guarantee that those services have the resources they require while still allowing high consolidation of other virtual machines.

Storage Guidelines

VSPEX architecture for VMware virtual machines for mid-market segment in this design, 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 more information, refer to the VMware whitepaper Comparison of Storage Protocol Performance in VMware vSphere 5 at http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf.

Storage Best Practices

The 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. This redundancy is in the form of dual adapters connected to separate fabric switches.
- Partition alignment—Partition misalignment can lead to severe performance degradation due to I/O operations having to cross track boundaries. Partition alignment is important both at the VMFS level as well as within the guest operating system. Use the vSphere Client when creating VMFS datastores to be sure they are created aligned. When formatting volumes within the guest, Windows 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 multitier application, you can expect each tier of application to have different I/O requirements. As a general recommendation, pay close attention to the amount of virtual machine disk files hosted on a single VMFS volume. Over-subscription of the I/O resources can go unnoticed at first and slowly begin to degrade performance if not monitored proactively.

Virtual Networking

NFS-variant architecture demonstrates use and benefits of Cisco Nexus 1000v virtual switching technology. Each Cisco UCS B200 M4 Blade Server and Cisco UCS C220 M4 Rack-Mount Server has one physical adapter with two 10 GE links going to fabric A and fabric B for high availability. Cisco UCS VIC 1340 or VIC 1227 presents four virtual Network Interface Cards (vNICs) to the hypervisor, two vNICs per fabric path. In FC-variant, the Cisco UCS VIC 1227 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 Cisco UCS Manager. The vNICs are used in active-active configuration for load-balancing and high-availability. The following are the 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 Cisco Unified Computing System, 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 the Nexus 5548UP switches or infrastructure switches that connect to the Cisco UCS Fabric Interconnect uplinks for faster convergence of STP in the events of FI reboots or FI uplink flaps.
- Jumbo MTU for vMotion and Storage Traffic—this best practice is implemented in the architecture by configuring jumbo MTU end-to-end.

VSPEX VMware Storage Virtualization

Storage Layout

The architecture diagram in this section shows the physical disk layout. Disk provisioning on the 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 15 illustrates storage architecture for 300 virtual machines on VNX5400 for FC-variant of architecture.

Figure 15 Storage layout for up to 300 reference VMs on VNX5400




Figure 16 illustrates storage architecture for 600 virtual machines on VNX5600 for NFS-variant of architecture.

Figure 16 *Storage Layout for up to 600 Reference Virtual Machines on VNX5600*




Figure 17 illustrates storage architecture for 1000 virtual machines on VNX5800 for NFS-variant of architecture.

Figure 17 *Storage Layout for up to 1000 Reference Virtual Machines on VNX5800*




Table 5 provides the size of data stores for various architectures laid out in the above figures:

Table 5 *Datastores for the Different Architectures*

Parameter	300 virtual machines	600 virtual machines	1000 virtual machines
Storage array	VNX5400	VNX5600	VNX5800
Disk capacity & 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

Table 6 *Storage Pools for the Different Architectures*

Configuration	Number of Pools	Number of 15k SAS Drives per Pool	Number of Flash drives per pool	Number of LUNs per Pool	Number of FS per Storage pool for File	LUN Size (GB)	FS Size (TB)
300 Virtual Machines	2	45	2	20	2	800	4
	1	20	2	20	2	400	3
Total	3	110	6	60	6	40 x 800GB LUNs 20 x 400GB LUNs	4 x 7TB FS 2 x 3TB FS
600 Virtual Machines	4	45	2	20	2	800	7
	1	40	2	20	2	700	6
Total	5	220	10	100	10	80 x 800GB LUNs 20 x 700GB LUNs	8 x 7TB FS 2 x 6TB FS
1000 Virtual Machines	8	45	2	20	2	800	7
Total	8	360	16	160	16	16 x 800GB LUNs	16x 7TB FS

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

Storage Virtualization

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

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

Service Profile Design

This architecture implements the following design steps to 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. The boot policy in the 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. The rack-mount servers are automatically put in the pool as and when they are fully discovered by Cisco 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 Cisco Unified Computing System and Cisco 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 this case, if a server has physical fault (for example, faulty memory, or PSU or fan), using the following steps, a new server can be procured within minutes:

- Put the faulty server in maintenance mode using vCenter. This would move virtual machines 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 Cisco UCS Manager.
- Associate the service profile to the newly deployed rack-mount 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.

The architecture achieves the statelessness of the computing in the datacenter. If there are enough identifiers in all the id-pools, and if more servers are attached to Cisco Unified Computing System, in the future, more service profiles can be derived from the service profile template and the private cloud infrastructure can be easily expanded. Blade and rack-mount servers can be added in the same server pool.

Network High Availability Design—FC-Variant

Figure 18 illustrates the logical layout of the FC-Variant of architecture. The following are the key aspects of this solution:

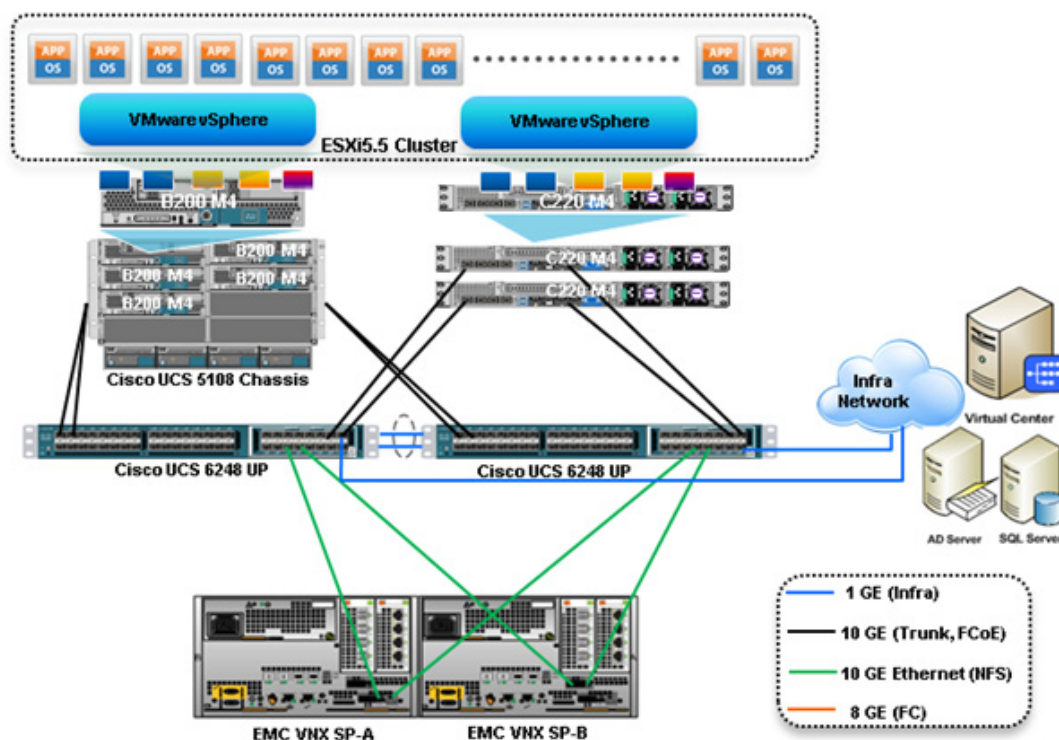
- Mix of Cisco UCS B200 M4 and C220 M4 servers are used, managed by Cisco UCS Manager (UCSM).
- Fabric A and Fabric B are used with host based FC multi-pathing for high availability.

- EMC VNX5400 storage array is directly attached to Cisco UCS Fabric Interconnects.
- Two 10GE links between FI and FEX provides enough bandwidth oversubscription for the SMB segment private cloud. The oversubscription can be reduced by adding more 10GE links between FI and FEX if needed by the virtual machines 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 the 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 thru SAN-B and storage-processor inter-link.
- On hosts, boot order lists vHBA on both fabrics for high-availability.

Figure 18 Logical Layout of the FC-Variant of Architecture



Network High Availability Design—NFS-Variant

Figure 19 illustrates the logical layout of the architecture. The key aspects of this solution are as follows:


- Mix of Cisco UCS B200 M4 and C220 M4 servers are used, managed by Cisco UCS Manager (UCSM)
- 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 9396PX and Cisco UCS Fabric Interconnect's for high availability
- vPC and port-aggregation is used between Cisco Nexus 9396PX 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 virtual machines running on the ESXi hosts.
- VHBAs on Fabric A and Fabric B used for booting ESXi hypervisor images over SAN and it provides FC based storage access high availability.

Storage is made highly available by deploying following practices:

- FC SAN access for booting the hypervisor images in the NFS-Variant of architecture is different from FC-variant. In NFS variant, the Fabric Interconnects (A and B) are connected upstream Fabric Switch MDS9148S (A and B) for Fabric zoning. Other difference is the virtual machines 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.
- MDS 9148S switches (A and B) are connected to both Storage Processors SPA and SPB (over FC) and N9396PX switches (A and B) are connected to both Data Movers (over Ethernet), however, a given FI's connected to both MDS 9148S and N9396PX switches for FC and NFS storage access respectively.
- vPC on Nexus 9396PX 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.
- Similar to FC-variant, In NFS variant Architecture the VNX Storage Processors are always in the active/active mode; if the target cannot be reached on SAN-A, server can access the LUNs thru SAN-B and storage-processor inter-link.
- On hosts, theboot order lists vHBA on both fabrics for high-availability.

Figure 19 Logical Layout of the NFS-Variant of Architecture



Jumbo MTU

Jumbo MTU (size 9000) is used for the 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 help ensure that IP packets are not fragmented by intermediate network nodes. The checklist of end-points where the Jumbo MTU needs to be configured are as follows:

- Ethernet ports on VNX Storage Processors
- System QoS classes in Nexus 9396PX switches
- System QoS classes in Cisco UCS Manager
- vNICS in service profiles
- Nexus 1000v switch or vSwitches on the ESXi hosts
- VM-Kernel ports used for vMotion and storage access on the ESXi hosts

The next section details the sizing guidelines of the Cisco solution for EMC VSPEX VMware architectures.

Sizing Guideline

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

Defining the Reference Workload

To simplify the discussion, an example of a customer reference workload is provided. By comparing your actual customer usage to this reference workload, you can extrapolate which reference architecture to choose.

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

Table 7 *Virtual Machine Characteristics*

Characteristic	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

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

Applying the Reference Workload

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

The reference architectures create a pool of resources sufficient to host a target number of reference virtual machines as described above. It is entirely possible that customer virtual machines may not exactly match the specifications above. In this case, you can say that a single specific customer virtual machine is the equivalent of some number of reference virtual machines, and assume that 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.

Example 1: Custom Built Application

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

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

- CPU resources for one virtual machine
- Memory resources for two virtual machines
- Storage capacity for one virtual machine

- IOPS for one virtual machine

In this example, a single virtual machine uses the resources of two of the reference virtual machines. If the original pool had the capability to provide 300 virtual machines worth of resources, the new capability is 298 virtual machines.

Example 2: Point of Sale System

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

The following are the requirements to virtualize this application:

- CPUs of four reference virtual machines
- Memory of eight reference virtual machines
- Storage of two reference virtual machines
- IOPS of eight reference virtual machines

In this case the one virtual machine uses the resources of eight reference virtual machines. If this was implemented on a resource pool for 300 virtual machines, there are 292 virtual machines of capability remaining in the pool.

Example 3: Web Server

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

The following are the requirements to virtualize this application:

- CPUs of two reference virtual machines
- Memory of four reference virtual machines
- Storage of one reference virtual machines
- IOPS of two reference virtual machines

In this case the virtual machine would use the resources of four reference virtual machines. If this was implemented on a resource pool for 300 virtual machines, there are 296 virtual machines of capability remaining in the pool.

Example 4: Decision Support Database

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

The following are the requirements to virtualize this application:

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

In this case the one virtual machine uses the resources of fifty-two reference virtual machines. If this was implemented on a resource pool for 300 virtual machines, there are 248 virtual machines of capability remaining in the pool.

Summary of Example

The four examples presented illustrate the flexibility of the resource pool model. In all four cases the workloads simply reduce the number of available resources in the pool. If all four examples were implemented on the same virtual infrastructure, with an initial capacity of 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 outside the scope of this document. However, when the change in resource balance has been examined, and the new level of requirements is known; these virtual machines can be added to the infrastructure using the method described in the examples.

The next section details the deployment of the Cisco solution for EMC VSPEC VMware architectures.

VSPEX Configuration Guidelines

The configuration for the Cisco solution for EMC VSPEX VMware architectures is divided into the following steps:

1. Pre-deployment tasks
2. Connect the network cables
3. Configure the Cisco Nexus 9396PX switches (NFS-Variant only)
4. Configure the MDS 9148S switches (NFS-variant only)
5. Prepare the Cisco UCS FIs and configure Cisco UCS Manager
6. Configure the data-stores for ESXi images
7. Install the ESXi servers and vCenter infrastructure
8. Install and configure the vCenter server
9. Install the Cisco Nexus 1000v VMS VM (NFS-Variant only)
10. Configure the storage for virtual machine data stores, install and instantiate the virtual machines thru the vCenter
11. Test the installation

The following sections detail the steps mentioned above.

Pre-deployment Tasks

The pre-deployment tasks include the procedures that do not directly relate to the 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, etc. These tasks should be performed before the customer visit to decrease the time required onsite.

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

Table 8 *Customer Specific Configuration Data*

Requirement	Description	Reference
Hardware	Cisco UCS Fabric Interconnects, Fabric Extenders and Cisco UCS chassis for network and compute infrastructure	See corresponding product documentation
	Cisco Nexus 9396PX switches and Cisco MDS 9148S Fabric Switches for NFS and FC access respectively.	
	Cisco UCS B200 M4 and/or C220 M4 servers to host virtual machines	
	VMware vSphere™ 5.5 server to host virtual infrastructure servers	
	Note: This requirement may be covered in the existing infrastructure EMC VNX storage: Multiprotocol storage array with the required disk layout as per architecture requirements	
Software	Cisco Nexus 1000v VMS and VEM installation media	See corresponding product documentation
	VMware ESXi™ 5.5 installation media	
	VMware vCenter Server 5.5 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 R2 installation media (suggested OS for VMware vCenter & hosting VSPEX VMs)	

Licenses	Microsoft SQL Server 2012 R2 Note: This requirement may be covered in the existing infrastructure	Consult your corresponding vendor to obtain license keys
	Cisco Nexus 1000v license key	
	VMware vCenter 5.5 license key	
	VMware ESXi 5.5 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.

[Appendix—Customer Configuration Data Sheet](#) provides a table to maintain a record of relevant information. This form can be expanded or contracted as required, and information may be added, modified, and recorded as deployment progresses.

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

Connect the Network Cables

See the Cisco Nexus 9396PX, Cisco UCS FI, FEX, Blade Servers Chassis, B-Series and C-Series servers and the EMC VNX Storage Array Configuration guide for detailed information about how to mount the hardware on the rack. [Figure 20](#) illustrates the connectivity details for the VSPEX VMware architecture covered in this document.

Connectivity for FC-Variant

As shown in [Figure 20](#), there are four major cabling sections in this architecture:

- Fabric Interconnect's to storage array Fibre-Channel links (shown in Green)
- Fabric Interconnect's links to Fabric Extenders on Cisco UCS Chassis (shown in Black)
- Fabric Interconnect's links to Cisco UCS C220M4 servers (shown in Black)
- Infrastructure connectivity (not shown)

Figure 20 Detailed Connectivity Diagram of the FC-Variant of Architecture




Table 9 details the cable connectivity for the architecture.

Table 9 Cable Connectivity

Cable ID	Peer 1	Peer 2	VLAN	Mode	Description
A	FI-A, FC 2/9	SP-A,	Storage VSAN	Appliance	Directly attached storage on FI
B	FI-A, FC 2/10	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
C	FI-B, FC 2/9	SP-B,	Storage VSAN	Appliance	Directly attached storage on FI
D	FI-B, FC 2/10	Sp-A,	Storage VSAN	Appliance	Directly attached storage on FI
E	FIA, Eth 1/3	C220-M4 VIC port 1	N/A	Server	Server to fabric A. VLANs are allowed on per vNIC basis
F	FIB, Eth 1/3	C220-M4 VIC port 2	N/A	Server	Server to fabric A. VLANs are allowed on per vNIC basis
G,H	FI-A, Eth 1/1, 1/2	5108 Chassis, FEX 2208 Left	N/A	Server	FI/FEX 20GE port-channel connectivity

I,J	FI-B, Eth 1/1, 1/2	5108 Chassis, FEX 2208 Right	N/A	Server	FI/FEX 20GE port-channel connectivity
(not shown)	Eth1/31 on FI-A and FI-B	Uplink switch	All	Uplink	Uplink to Infrastructure network

Connectivity for NFS-Variant

Divide the connectivity into FC and Ethernet components.

The ethernet cable connectivity can be divided into the following categories in this architecture:

1. Cisco Nexus 9K switches to VNX Data Mover 10G Ethernet links (Blue)
2. Cisco Nexus 9k vPC peer links (Black)
3. Fabric Interconnect's to Cisco Nexus 9k 10G Ethernet links - crisscrossed (Purple)
4. Fabric Interconnect's to Fabric Extenders links - per fabric (Black)
5. Fabric Interconnect's to Cisco UCS C220 M4 Server links - per Fabric (Black)
6. Infrastructure connectivity (not shown)

Fibre Channel connectivity is relatively simple and can be divided into the following categories:

1. MDS 9148S switches to storage array SPs links (Orange)
2. MDS 9148S switches to Fabric Interconnect's (Green)

Figure 21 *Connectivity of the Architecture*

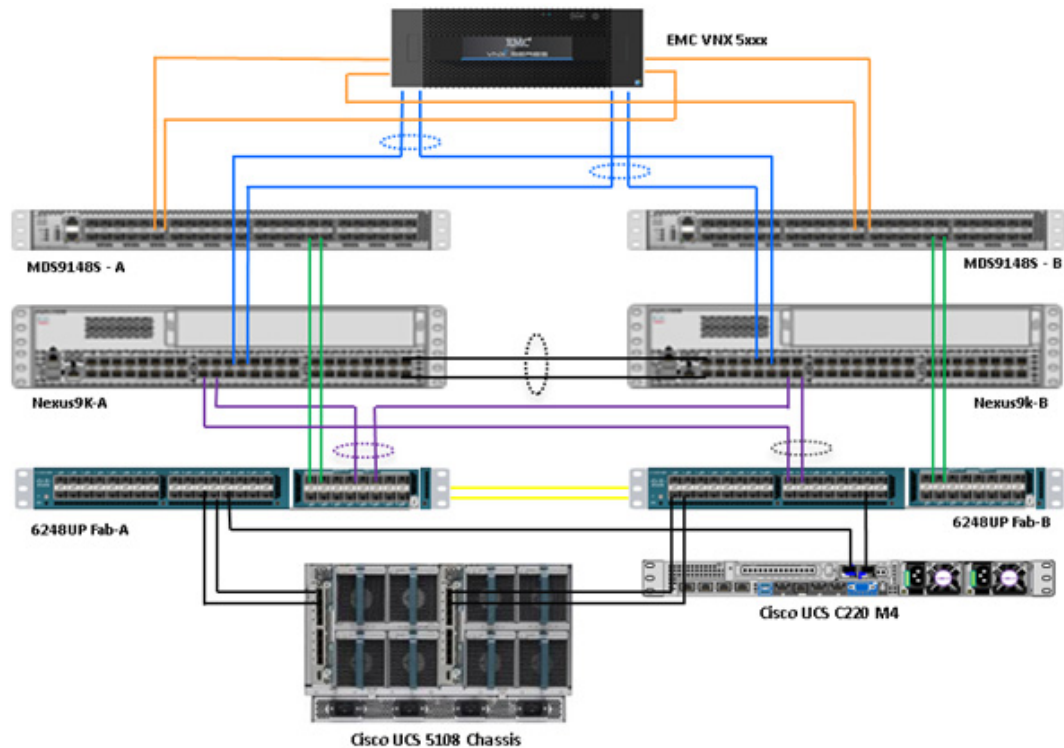


Table 10 details the ethernet cable connectivity for the NFS-Variant of the architecture.

Table 10 *Ethernet Cable Connectivity for the NFS-Variant*

Cable ID	Peer 1	Peer 2	VLAN	Mode	Description
A	N9K-A, Eth 1/33	DM-2, A1/0	Storage	Access (on N9k)	N9k vPC member port to storage Data Mover
B	N9K-A, Eth 1/34	DM-3, B1/0	Storage	Access (on N9k)	N9k vPC member port to storage Data Mover
C	N9K-B, Eth 1/33	DM-2, A1/1	Storage	Access (on N9k)	N9k vPC member port to storage Data Mover
D	N9K-B, Eth 1/34	DM-3, B1/1	Storage	Access (on N9k)	N9k vPC member port to storage Data Mover
E,F	Eth 1/1, Eth 1/2 on N9K-A	Eth 1/1, Eth 1/2 on N9K-B	All	Trunk	vPC peer-links between N9Ks
G	N9K-A, Eth 1/17	FI-A, Eth 1/17	All	Trunk	N9K vPC member port, FI PC member port
H	N9K-A, Eth 1/18	FI-B, Eth 1/17	All	Trunk	N9K vPC member port, FI PC member port
I	N9K-B, Eth 1/17	FI-A, Eth 1/18	All	Trunk	N9K vPC member port, FI PC member port

J	N9K-B, Eth 1/18	FI-B, Eth 1/18	All	Trunk	N9K 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
O	C220M4, VIC Port 1 & Port 2	FI-A, Eth1/5 & FI-B Eth1/5	N/A	Server (on FI)	FI/ IOM links
P (not shown)	Eth 1/31 on FI-A and FI-B	Uplink switch	All	Uplink	Uplink to Infrastructure network

Table 11 details the FC cable connectivity for the NFS-Variant of the architecture.

Table 11 FC Cable Connectivity for NFS-Variant

Cable ID	Peer 1	Peer 2	VSAN	Description
Q	MDS-A, FC 1/3	SP-A, 0/0	Storage VSAN	MDS 9148S to storage SP, crisscrossed Eth cables
R	MDS-A, FC 1/4	SP-B, 0/0	Storage VSAN	MDS 9148S to storage SP, crisscrossed Eth cables
S	MDS-B, FC 1/3	SP-A, 0/1	Storage VSAN	MDS 9148S to storage SP, crisscrossed Eth cables
T	MDS-B, FC 1/4	SP-A, 0/1	Storage VSAN	MDS 9148S to storage SP, crisscrossed Eth cables
U	MDS-A, FC 1/1	FI-A, FC2/9	Storage VSAN	MDS 9148S to FI straight FC cables
V	MDS-A, FC 1/2	FI-A, FC2/10	Storage VSAN	MDS 9148S to FI straight FC cables
W	MDS-B, FC 1/1	FI-B, FC2/9	Storage VSAN	MDS 9148S to FI straight FC cables
X	MDS-B, FC 1/2	FI-B, FC2/9	Storage VSAN	MDS 9148S to FI straight FC cables

Connect all the cables as outlined above; the next step is to configure the Cisco Nexus 9K switches, storage array and Cisco UCS Manager.

Prepare and Configure the Cisco Nexus 9396PX Switches

This section explains the switch configuration needed for the Cisco solution for EMC VSPEX VMware architecture. Details about configuring the password, management connectivity and strengthening the device are not covered here; please refer to the Cisco Nexus 9000 Series Configuration guide for details.

Configure Global VLANs

Figure 22 and Figure 23 details how to configure VLAN on a Cisco Nexus 9K switch A and B.

Figure 22 Cisco Nexus 9K Switch A VLAN Configuration

```

VSPEX-N9K-FAB-A# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# vlan 40
VSPEX-N9K-FAB-A(config-vlan)# name Storage
VSPEX-N9K-FAB-A(config-vlan)# no shut
VSPEX-N9K-FAB-A(config-vlan)# exit
VSPEX-N9K-FAB-A(config)# vlan 41
VSPEX-N9K-FAB-A(config-vlan)# name vMotion
VSPEX-N9K-FAB-A(config-vlan)# no shut
VSPEX-N9K-FAB-A(config-vlan)# exit
VSPEX-N9K-FAB-A(config)# vlan 45
VSPEX-N9K-FAB-A(config-vlan)# name VM-Data
VSPEX-N9K-FAB-A(config-vlan)# no shut
VSPEX-N9K-FAB-A(config-vlan)# exit
VSPEX-N9K-FAB-A(config)#

```

Figure 23 Cisco Nexus 9K Switch B VLAN Configuration

```

VSPEX-N9K-FAB-B# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-B(config)# vlan 40
VSPEX-N9K-FAB-B(config-vlan)# name Storage
VSPEX-N9K-FAB-B(config-vlan)# no shut
VSPEX-N9K-FAB-B(config-vlan)# exit
VSPEX-N9K-FAB-B(config)# vlan 41
VSPEX-N9K-FAB-B(config-vlan)# name vMotion
VSPEX-N9K-FAB-B(config-vlan)# no shut
VSPEX-N9K-FAB-B(config-vlan)# exit
VSPEX-N9K-FAB-B(config)# vlan 45
VSPEX-N9K-FAB-B(config-vlan)# name VM-Data
VSPEX-N9K-FAB-B(config-vlan)# no shut
VSPEX-N9K-FAB-B(config-vlan)# exit
VSPEX-N9K-FAB-B(config)#

```

Table 12 details the VLANs to be configured on both switches A and B in addition to your application specific VLANs.

Table 12 Configured VLANs on Cisco Nexus 9K Switch A and Switch 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)

Refer to the section Customer Configuration Data Sheet for the actual VLAN IDs in your deployment.

Configure Virtual Port-Channel (vPC)

The Virtual port-channel effectively enables two physical switches to behave as a single virtual switch, and the port-channel can be formed across the two physical switches. To configure the vPC, follow these steps:

1. Enable the LACP feature on both switches.
2. Enable the vPC feature on both switches.

3. Configure a unique vPC domain ID, identical on both switches.
4. Configure the mutual management IP addresses on both switches and configure peer-gateway. Refer to the Cisco Nexus 9k Switch VPC configuration in .

```
VSPEX-N9K-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# feature lacp
VSPEX-N9K-FAB-A(config)# feature vpc
VSPEX-N9K-FAB-A(config)# vpc domain 101
VSPEX-N9K-FAB-A(config-vpc-domain)# peer-keepalive destination 10.29.180.12
Note:
-----:: Management VRF will be used as the default VRF ::-----
VSPEX-N9K-FAB-A(config-vpc-domain)# peer-gateway
VSPEX-N9K-FAB-A(config-vpc-domain)# exit
VSPEX-N9K-FAB-A(config)# exit
VSPEX-N9K-FAB-A#
```

5. Configure the port-channel on the inter-switch links. Make sure that "vpc peer-link" is configured on this port-channel.

```
VSPEX-N9K-FAB-A# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# interface port-channel 1
VSPEX-N9K-FAB-A(config-if)# switchport mode trunk
VSPEX-N9K-FAB-A(config-if)# spanning-tree port type network
VSPEX-N9K-FAB-A(config-if)# speed 10000
VSPEX-N9K-FAB-A(config-if)# vpc peer-link
Please note that spanning tree port type is changed to "network" port type on vP
C peer-link.
This will enable spanning tree Bridge Assurance on vPC peer-link provided the ST
P Bridge Assurance
(which is enabled by default) is not disabled.
VSPEX-N9K-FAB-A(config-if)# description VPC-Peerlink
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# exit
VSPEX-N9K-FAB-A#
```

6. Add the ports with the LACP protocol on the port-channel using "channel-group 1 force mode active" command under the interface subcommand as shown below.

```
VSPEX-N9K-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# interface ethernet 1/1-2
VSPEX-N9K-FAB-A(config-if-range)# channel-group 1 force mode active
VSPEX-N9K-FAB-A(config-if-range)# no shut
VSPEX-N9K-FAB-A(config-if-range)# exit
VSPEX-N9K-FAB-A(config)# exit
VSPEX-N9K-FAB-A#
```

7. Repeat the steps from 1 to 6 to create the VPC domain on the Cisco Nexus 9k Fabric B as shown below.

```

VSPEX-N9K-FAB-B# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
VSPEX-N9K-FAB-B(config)# feature lacp
VSPEX-N9K-FAB-B(config)# feature vpc
VSPEX-N9K-FAB-B(config)# vpc domain 101
VSPEX-N9K-FAB-B(config-vpc-domain)# peer
peer-gateway      peer-keepalive      peer-switch
VSPEX-N9K-FAB-B(config-vpc-domain)# peer-keepalive destination 10.29.180.11
Note:
-----:: Management VRF will be used as the default VRF ::-----
VSPEX-N9K-FAB-B(config-vpc-domain)# peer-gateway
VSPEX-N9K-FAB-B(config-vpc-domain)# exit
VSPEX-N9K-FAB-B(config)# exit
VSPEX-N9K-FAB-B#

```

```

VSPEX-N9K-FAB-B(config)# int port-channel 1
VSPEX-N9K-FAB-B(config-if)# switchport mode trunk
VSPEX-N9K-FAB-B(config-if)# spanning-tree port type network
VSPEX-N9K-FAB-B(config-if)# speed 10000
VSPEX-N9K-FAB-B(config-if)# vpc peer-link
Please note that spanning tree port type is changed to "network" port type on vP
C peer-link.
This will enable spanning tree Bridge Assurance on vPC peer-link provided the ST
P Bridge Assurance
(which is enabled by default) is not disabled.
VSPEX-N9K-FAB-B(config-if)# description VPC-Peerlink
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# exit
VSPEX-N9K-FAB-B#

```

```

VSPEX-N9K-FAB-B# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
VSPEX-N9K-FAB-B(config)# interface ethernet 1/1-2
VSPEX-N9K-FAB-B(config-if-range)# channel-group 1 force mode active
VSPEX-N9K-FAB-B(config-if-range)# no shut
VSPEX-N9K-FAB-B(config-if-range)# exit
VSPEX-N9K-FAB-B(config)# exit
VSPEX-N9K-FAB-B#

```

8. Verify vPC status using "show vpc" command.

```

VSPEX-N9K-FAB-A(config)# sh vpc brief
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 inconsistency reason : success
vPC role               : secondary
Number of vPCs configured : 0
Peer Gateway           : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status   : Disabled

vPC Peer-link status
-----
id   Port   Status Active vlans
--   --
1    Po1    up      1
VSPEX-N9K-FAB-A(config)#

```

Configure Port Channels Connected to FabricInterconnects

Interfaces connected to fabric interconnects need to be in the trunk mode and storage, vMotion, infra and application VLANs are allowed on this port. From the switch side, interfaces connected to FI-A and FI-B are in a vPC, and from the FI side the links connected to Nexus 9396 A and B switches are in regular LACP port-channels. It is a good practice to use good description for each port and port-channel on the switch for better diagnosis if any problem arises later. Refer to [Figure 24](#) and [Figure 25](#) for the exact configuration commands for Cisco Nexus 9K switch A and B:

Figure 24 Cisco Nexus 9K Switch A vPC Config. for FIs

```

VSPEX-N9K-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# interface port-channel 17
VSPEX-N9K-FAB-A(config-if)# description to FI-A
VSPEX-N9K-FAB-A(config-if)# switchport mode trunk
VSPEX-N9K-FAB-A(config-if)# vpc 17
VSPEX-N9K-FAB-A(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# interface port-channel 18
VSPEX-N9K-FAB-A(config-if)# description to FI-B
VSPEX-N9K-FAB-A(config-if)# switchport mode trunk
VSPEX-N9K-FAB-A(config-if)# vpc 18
VSPEX-N9K-FAB-A(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# interface ethernet 1/17
VSPEX-N9K-FAB-A(config-if)# description vpc port-channel to FI-A
VSPEX-N9K-FAB-A(config-if)# switchport mode trunk
VSPEX-N9K-FAB-A(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-A(config-if)# channel-group 17 mode active
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# interface ethernet 1/18
VSPEX-N9K-FAB-A(config-if)# description vpc port-channel to FI-B
VSPEX-N9K-FAB-A(config-if)# switchport mode trunk
VSPEX-N9K-FAB-A(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-A(config-if)# channel-group 18 mode active
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)#

```

Figure 25 Cisco Nexus 9K Switch B vPC Config. for FIs

```

VSPEX-N9K-FAB-B# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-B(config)# interface port-channel 17
VSPEX-N9K-FAB-B(config-if)# description to FI-A
VSPEX-N9K-FAB-B(config-if)# switchport mode trunk
VSPEX-N9K-FAB-B(config-if)# vpc 17
VSPEX-N9K-FAB-B(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface port-channel 18
VSPEX-N9K-FAB-B(config-if)# description to FI-B
VSPEX-N9K-FAB-B(config-if)# switchport mode trunk
VSPEX-N9K-FAB-B(config-if)# vpc 18
VSPEX-N9K-FAB-B(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface ethernet 1/17
VSPEX-N9K-FAB-B(config-if)# description vpc port-channel to FI-A
VSPEX-N9K-FAB-B(config-if)# switchport mode trunk
VSPEX-N9K-FAB-B(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-B(config-if)# channel-group 17 mode active
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface ethernet 1/18
VSPEX-N9K-FAB-B(config-if)# description vpc port-channel to FI-B
VSPEX-N9K-FAB-B(config-if)# switchport mode trunk
VSPEX-N9K-FAB-B(config-if)# switchport trunk allowed vlan 1,40-41,45
VSPEX-N9K-FAB-B(config-if)# channel-group 18 mode active
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)#

```

Configure Storage Connectivity

From each Cisco N9K switch 2 link connects to each Data Mover on the VNX storage array. A virtual port-channel is created for the two links connected to a single Data Mover, but connected to two different switches. In this example configuration, links connected to Data mover-2 of VNX storage array are connected to Ethernet port 1/33 on each switch and links connected to Data mover-3 are connected to Ethernet port 1/34 on each switch. A virtual port-channel (id 33) is created for port Ethernet 1/33 on each switch and a different virtual port-channel (id 34) is created for port Ethernet 1/34 on each switch. Note that only storage VLAN is required on this port, and so, the port is in access mode. [Figure 26](#) and [Figure 27](#) detail the configuration on the port-channels and interfaces:

Figure 26 *Cisco Nexus 9K Switch A vPC Config. for VNX Arrays*

```

VSPEX-N9K-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# interface port-channel 33
VSPEX-N9K-FAB-A(config-if)# description to VNX5500-DM-2
VSPEX-N9K-FAB-A(config-if)# vpc 33
VSPEX-N9K-FAB-A(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# interface port-channel 34
VSPEX-N9K-FAB-A(config-if)# description to VNX5500-DM-3
VSPEX-N9K-FAB-A(config-if)# vpc 34
VSPEX-N9K-FAB-A(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-A(config-if)# exit

```

```

VSPEX-N9K-FAB-A(config)# interface ethernet 1/33
VSPEX-N9K-FAB-A(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-A(config-if)# channel-group 33 mode active
VSPEX-N9K-FAB-A(config-if)# no shu
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)# interface ethernet 1/34
VSPEX-N9K-FAB-A(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-A(config-if)# channel-group 34 mode active
VSPEX-N9K-FAB-A(config-if)# no shutdown
VSPEX-N9K-FAB-A(config-if)# exit
VSPEX-N9K-FAB-A(config)#

```

Figure 27 *Cisco Nexus 9K Switch B vPC Config. for VNX Arrays*

```

VSPEX-N9K-FAB-B# conf t
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-B(config)# interface port-channel 33
VSPEX-N9K-FAB-B(config-if)# description to VNX5500-DM-2
VSPEX-N9K-FAB-B(config-if)# vpc 33
VSPEX-N9K-FAB-B(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface port-channel 34
VSPEX-N9K-FAB-B(config-if)# description to VNX5500-DM-3
VSPEX-N9K-FAB-B(config-if)# vpc 34
VSPEX-N9K-FAB-B(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface ethernet 1/33
VSPEX-N9K-FAB-B(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-B(config-if)# channel-group 33 mode active
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)# interface ethernet 1/34
VSPEX-N9K-FAB-B(config-if)# switchport access vlan 40
VSPEX-N9K-FAB-B(config-if)# channel-group 34 mode active
VSPEX-N9K-FAB-B(config-if)# no shutdown
VSPEX-N9K-FAB-B(config-if)# exit
VSPEX-N9K-FAB-B(config)#

```

Configure Ports Connected to Infrastructure Network

The port connected to the infrastructure network needs to be in trunk mode and they require an infrastructure VLAN as part of the allowed VLANs list. Additional VLANs may need to be added as required by your applications. For example, Windows virtual machines may need access to the 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, all the ports and port channels are configured with the 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 the following figures. Note that the ports are "up" only after the peer devices are also configured properly, so revisit this subsection after configuring VNX storage array and fabric interconnects.

```
VSPEX-N9K-FAB-A# show vlan
```

VLAN	Name	Status	Ports
1	default	active	Po1, Po17, Po18, Eth1/1, Eth1/2 Eth1/3, Eth1/4, Eth1/5, Eth1/6 Eth1/7, Eth1/8, Eth1/9, Eth1/10 Eth1/11, Eth1/12, Eth1/13 Eth1/14, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22 Eth1/23, Eth1/24, Eth1/25 Eth1/26, Eth1/27, Eth1/28 Eth1/29, Eth1/30, Eth1/31 Eth1/32, Eth1/35, Eth1/36 Eth1/37, Eth1/38, Eth1/39 Eth1/40, Eth1/41, Eth1/42 Eth1/43, Eth1/44, Eth1/45 Eth1/46, Eth1/47, Eth1/48 Eth2/1, Eth2/2, Eth2/3, Eth2/4 Eth2/5, Eth2/6, Eth2/7, Eth2/8 Eth2/9, Eth2/10, Eth2/11 Eth2/12
40	Storage	active	Po1, Po17, Po18, Eth1/1, Eth1/2 Eth1/17, Eth1/18, Eth1/33 Eth1/34
41	vMotion	active	Po1, Po17, Po18, Eth1/1, Eth1/2 Eth1/17, Eth1/18
45	VM-Data	active	Po1, Po17, Po18, Eth1/1, Eth1/2 Eth1/17, Eth1/18
VLAN	Type	Vlan-mode	
1	enet	CE	
40	enet	CE	
41	enet	CE	
45	enet	CE	

The "show vlan" command can be restricted to a given VLAN or set of VLANs as shown in the above figure. Make sure that on both switches all of the required VLANs are in "active" status and the correct set of ports and port channels are part of the necessary VLANs.

The port channel configuration can be verified using the "show port-channel summary" command. The figure below shows the expected output of this command.

```

VSPEX-N9K-FAB-A(config)# sh 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)
        M - Not in use. Min-links not met
-----
Group  Port-      Type   Protocol  Member Ports
      Channel
-----
1      Po1(SU)     Eth    LACP      Eth1/1(D)  Eth1/2(P)
17     Po17(SU)  Eth    LACP      Eth1/17(P)
18     Po18(SU)  Eth    LACP      Eth1/18(P)
33     Po33(SU)  Eth    LACP      Eth1/33(s)
34     Po34(SD)  Eth    LACP      Eth1/34(s)
VSPEX-N9K-FAB-A(config)#

```

In this example, port channel 1 is the vPC peer-link port channel, port-channels 33 and 34 are connected to the storage arrays and port-channels 17 and 18 are connected to the Cisco UCS FIs. Make sure that the state of the member ports of each port-channel is "P" (Up in port-channel), except one of the two port-channels connected to the storage device. The Data Mover of the VNX storage array runs in active/passive mode, so one of the port-channel will remain down. Apart from that one port-channel, note that the port may not come up if the peer ports are not properly configured. Common reasons for a 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.

The vPC status can be verified using the "show vpc" command. An example of the output is shown in the figure below:

```

VSPEX-N9K-FAB-A(config)# sh vpc brief
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 inconsistency reason : success
vPC role                 : secondary
Number of vPCs configured : 4
Peer Gateway             : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status      : Disabled

vPC Peer-link status
-----
id   Port   Status Active vlans
--   -
1    Po1    up     1,40-41,45

vPC status
-----
id   Port   Status Consistency Reason      Active vlans
--   -
17   Po17   up     success  success  1,40-41,45
18   Po18   up     success  success  1,40-41,45
33   Po33   up     success  success  40
34   Po34   down*  success  success  -
VSPEX-N9K-FAB-A(config)#

```

Make sure that the vPC peer status is "peer adjacency formed ok" and all the port-channels, including the peer-link port-channel, have status "up", except for one of the two port-channels connected to the storage array as previously explained.

Configure QoS

The Cisco solution for EMC VSPEX VMware architectures require MTU is set at 9216 (jumbo frames) for efficient storage and vMotion traffic. Unlike the Cisco Nexus 5000 series switches, MTU configuration on Cisco Nexus 9000 series switches has to be configured on each port-channel/Interface level.

To configure the jumbo MTU on the Cisco Nexus 9000 series switches, use the following steps on both switch A and B; refer to the following figure for the CLI to configure "mtu9216":


```

VSPEX-N9K-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-N9K-FAB-A(config)# system jumbomtu 9216
VSPEX-N9K-FAB-A(config)# interface port-channel 17-18
VSPEX-N9K-FAB-A(config-if-range)# mtu 9216
VSPEX-N9K-FAB-A(config-if-range)# exit
VSPEX-N9K-FAB-A(config)# interface port-channel 33-34
VSPEX-N9K-FAB-A(config-if-range)# mtu 9216
VSPEX-N9K-FAB-A(config-if-range)# exit
VSPEX-N9K-FAB-A(config)# copy running-config startup-config
[#####] 100%
Copy complete.
VSPEX-N9K-FAB-A(config)#

```

Prepare and Configure MDS 9148S Switches

This section explains the Fabric switch configuration needed for the Cisco solution for EMC VSPEX VMware architecture. Details about configuring the password, management connectivity and strengthening the device are not covered here; for detailed instruction, refer to the MDS 9148S Switch Configuration guide.

Configure Global VSANs

Before configuring the global VSAN on the MDS 9148 switches, it is important to enable the NPIV feature on the Cisco MDS 9148S switches. [Figure 28](#) details how to enable the NPIV feature on the MDS 9148S switch.

Figure 28 *Enabling the NPIV feature on the MDS 9148S Switch*

```

VSPEX-MDS-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-MDS-FAB-A(config)# feature npiv
VSPEX-MDS-FAB-A(config)# end
VSPEX-MDS-FAB-A#

```

[Figure 29](#) shows how to configure VSAN on the MDS 9148S switch A and B.

Table 13 *Configured VSAN to Access Storage Array*

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

For the actual VSAN ID of your deployment, refer to the section "Customer Configuration Data Sheet".

Figure 29 *Creating VSAN on MDS 9148S Switch A*

```

VSPEX-MDS-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-MDS-FAB-A(config)# vsan database
VSPEX-MDS-FAB-A(config-vsan-db)# vsan 10
VSPEX-MDS-FAB-A(config-vsan-db)# vsan 10 interface fc 1/1
VSPEX-MDS-FAB-A(config-vsan-db)# vsan 10 interface fc 1/2
VSPEX-MDS-FAB-A(config-vsan-db)# vsan 10 interface fc 1/3
VSPEX-MDS-FAB-A(config-vsan-db)# vsan 10 interface fc 1/4
VSPEX-MDS-FAB-A(config-vsan-db)# end
VSPEX-MDS-FAB-A#

```

After creating the VSAN, the membership is assigned and the peer interfaces on the links need to be configured properly. A healthy fibre channel port is shown in figures below:

```

VSPEX-MDS-FAB-A# sh vsan membership
vsan 1 interfaces:
    fc1/5          fc1/6          fc1/7          fc1/8
    fc1/9          fc1/10         fc1/11         fc1/12
    fc1/13         fc1/14         fc1/15         fc1/16
    fc1/17         fc1/18         fc1/19         fc1/20
    fc1/21         fc1/22         fc1/23         fc1/24
    fc1/25         fc1/26         fc1/27         fc1/28
    fc1/29         fc1/30         fc1/31         fc1/32
    fc1/33         fc1/34         fc1/35         fc1/36
    fc1/37         fc1/38         fc1/39         fc1/40
    fc1/41         fc1/42         fc1/43         fc1/44
    fc1/45         fc1/46         fc1/47         fc1/48

vsan 10 interfaces:
    fc1/1          fc1/2          fc1/3          fc1/4

vsan 4079(evfp_isolated_vsan) interfaces:

vsan 4094(isolated_vsan) interfaces:

```

```

VSPEX-MDS-FAB-A(config)# sh int fc 1/1-4 brief
-----
Interface  Vsan   Admin  Admin  Status   SFP   Oper  Oper  Port
          Mode   Mode                                     Mode  Speed  Channel
          Mode   Mode                                     Mode  (Gbps)
-----
fc1/1      10     auto   on     up        sw1   F     8     --
fc1/2      10     auto   on     up        sw1   F     8     --
fc1/3      10     auto   on     up        sw1   F     8     --
fc1/4      10     auto   on     up        sw1   F     8     --
VSPEX-MDS-FAB-A(config)#

```

Repeat these steps to enable the NPIV feature and to create the VSANs on MDS 9148 switch B.



Note

The figures above demonstrates the use of the NXOS interface range CLI to configure multiple interfaces at the same time.

Prepare the Cisco UCS Fabric Interconnects and Configure Cisco UCS Manager

The next step is to configure the Cisco UCS FIs and Cisco UCS Manager. This task can be subdivided in to following segments:

- [Initial Configuration of Cisco UCS Fabric Interconnects](#)
- [Configuring Server Discovery](#)
- [Upstream Global Network Configuration](#)
- [Configure Identifier Pools](#)
- [Configure Server Pool and Qualifying Policy](#)
- [Configure Service Profile Template](#)
- [Instantiate Service Profiles from the Service Profile Template](#)

The following subsections detail the steps to configure the Cisco UCS Manager.

Initial Configuration of Cisco UCS Fabric Interconnects

At this point, the Cisco UCS Fabric Interconnects, blade servers chassis, FEX and C-series server must be mounted on the rack and the 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 are plugged in, ideally drawing power from two different power strips. Connect the 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 the initial configuration of the FIs, complete the following steps:

1. Attach RJ-45 serial console cable to the first FI, and connect the other end to the serial port of laptop. Configure a password for the "admin" account, fabric ID "A", Cisco UCS system name, management IP address, subnet mask and default gateway and cluster IP address (or Cisco UCS Manager Virtual IP address), as the initial configuration script walks you thru the configuration as shown in the image below. Save the configuration, which will bring up the Cisco UCS Manager CLI login prompt:

```

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. Disconnect the RJ-45 serial console from the recently configured FI and attach it to the other FI. The other FI will detect that its peer has been configured and will prompt to join the cluster. The only information required is the FI specific management IP address, subnet mask and default gateway, as shown in the image below. Save the configuration.

```

Enter the configuration method. (console/gui) ? console
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IP Address: 10.65.121.226
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):

```

- When the initial configurations on both FIs are completed, disconnect the serial console cable. The Cisco UCS Manager is accessible through the web interface (<https://<ucsm-virtual-ip>/>) or SSH. Connect to Cisco UCS Manager using SSH and view the HA status. Since there is a common device connected between the two FIs (a rack server or blade server chassis), the status will say "HA NOT READY", but both FI A and FI B must be visible in the "Up" state as shown in the image below:

```
VSPEX-FI-A# show cluster state
Cluster Id: 0xec91409a491011e2-0xb7a4547feaa1564

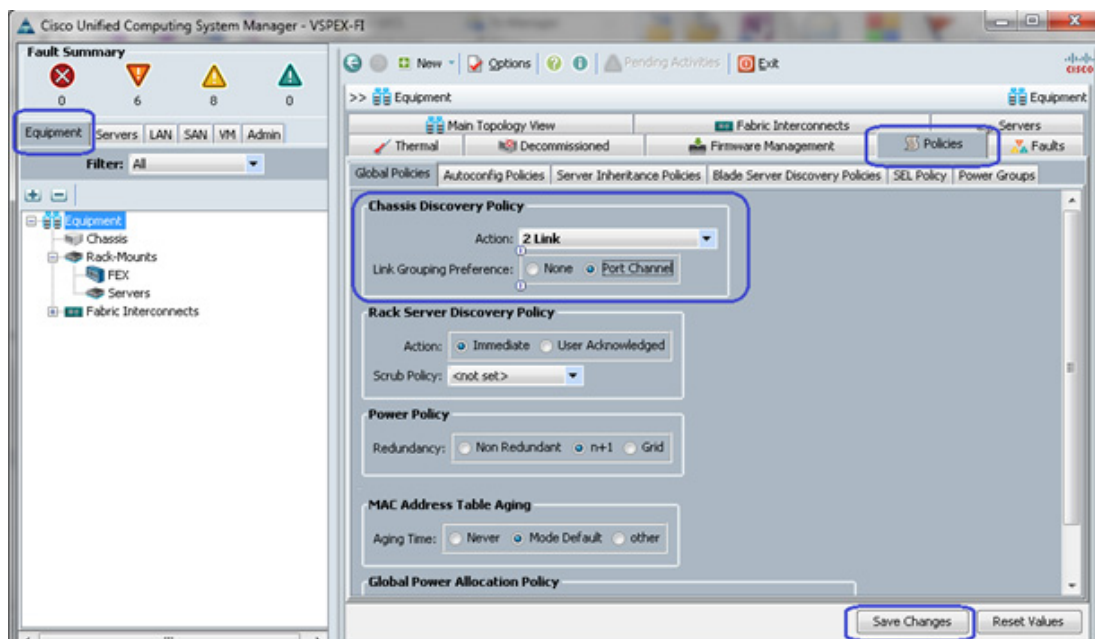
A: UP, PRIMARY
B: UP, SUBORDINATE

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

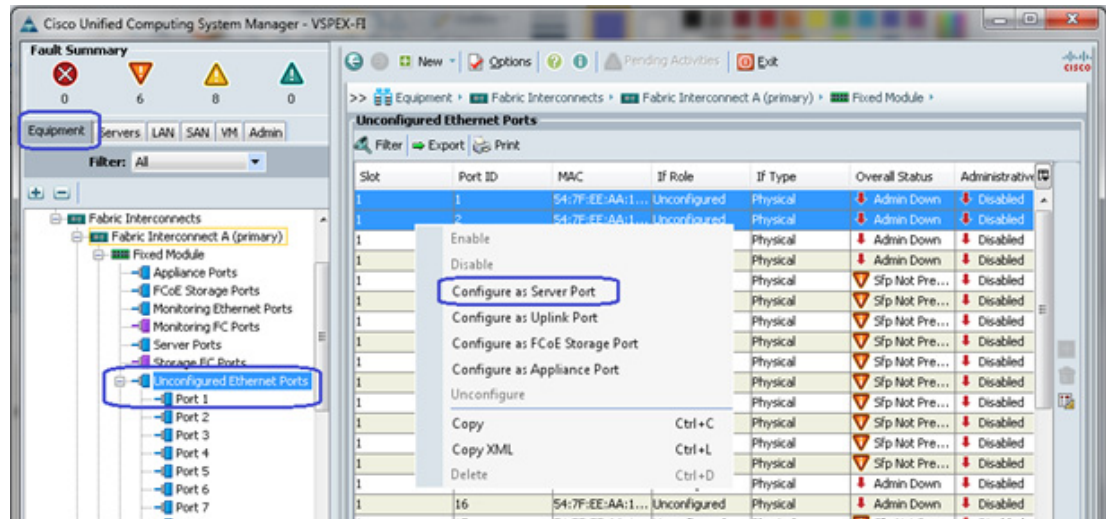
Configuring Server Discovery

All the FI Ethernet ports are unconfigured and shutdown by default. Classify the ports as server facing ports, directly attached storage array facing ports, and uplink ports. The next steps detail how to configure the ports for the proper server auto-discovery.

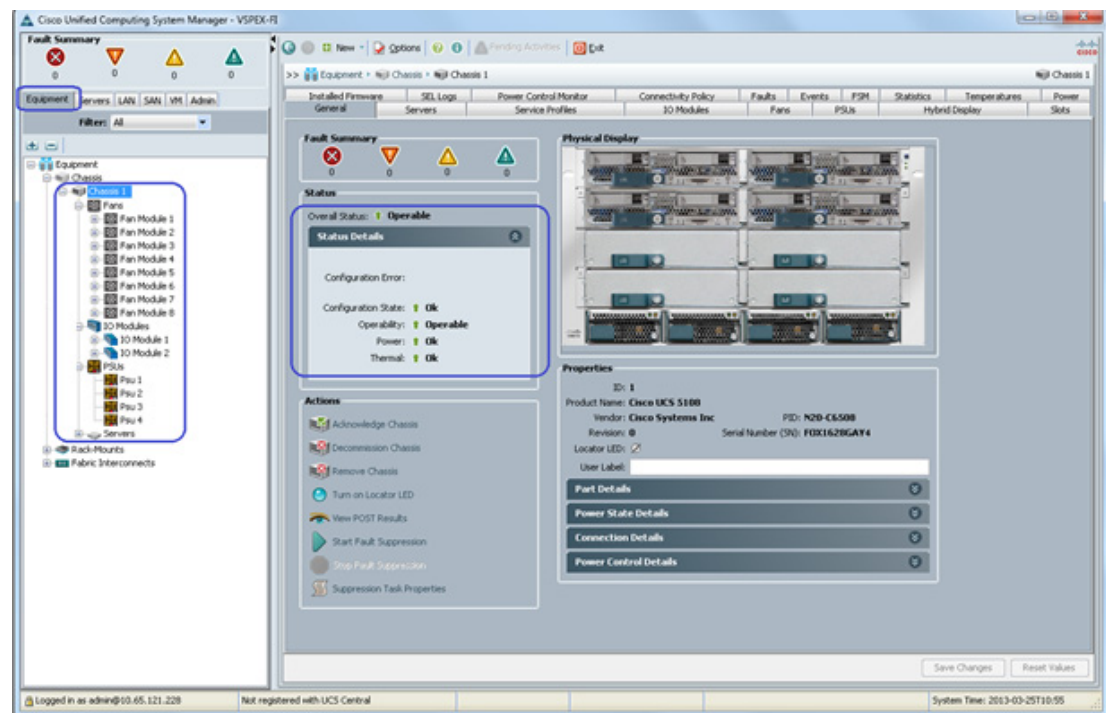
- Configure the chassis discovery policy that specifies the server side connectivity. Using a web browser, access the Cisco UCS Manager using the management virtual IP address and download the Java applet to launch the Cisco UCS Manager GUI. Click the "Equipment" tab and then click the "Policies" tab. From the "Chassis Discovery Policy" section, select "Action" as "2 Link", as two 10 GE links are connected between FI and FEX per fabric. Change "Link Grouping Preference" to "Port Channel" for better utilization of the bandwidth and link the level high-availability as shown in the image below. Save the changes.



- Identify the ports connected to the Chassis on a per FI basis. Click "Equipment", expand "Fabric Interconnects", choose a particular FI, for example "Fabric Interconnect A", click "Unconfigured Ethernet Ports", and select the two ports connected to the Chassis left IOModule1. Right-click, then click "Configure as Server Port" as shown in the image below. Click "Yes" on the confirmation pop-up window. The success pop-up displays when the ports are configured.



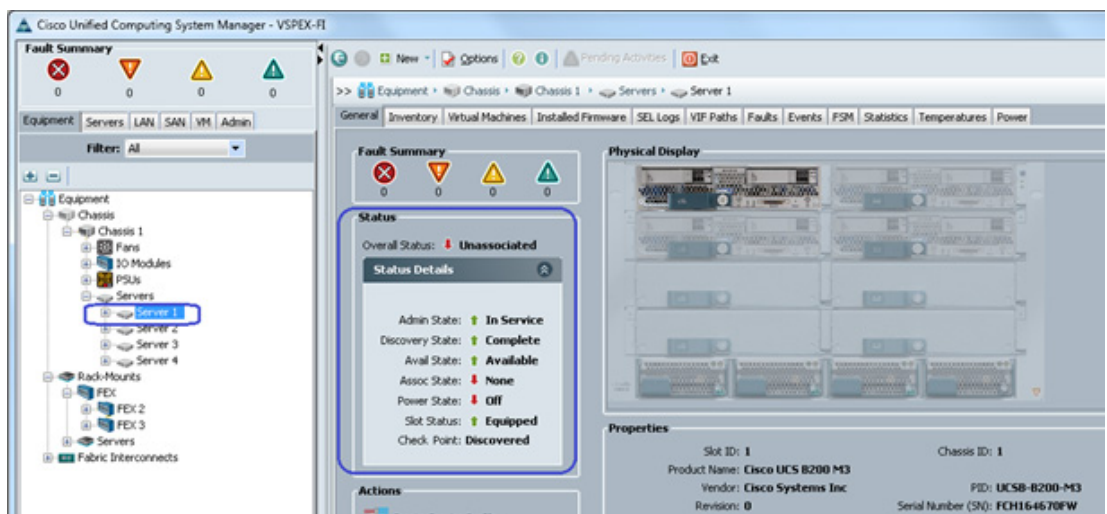
- Repeat step 2 for the Fabric Interconnect-B.
- When the server ports are configured on both FIs, the Cisco UCS Chassis auto-discovery will start and in few minutes, the fully discovered chassis is shown, with all its IOMs, fans, power supplies and so on, as shown in the image below:



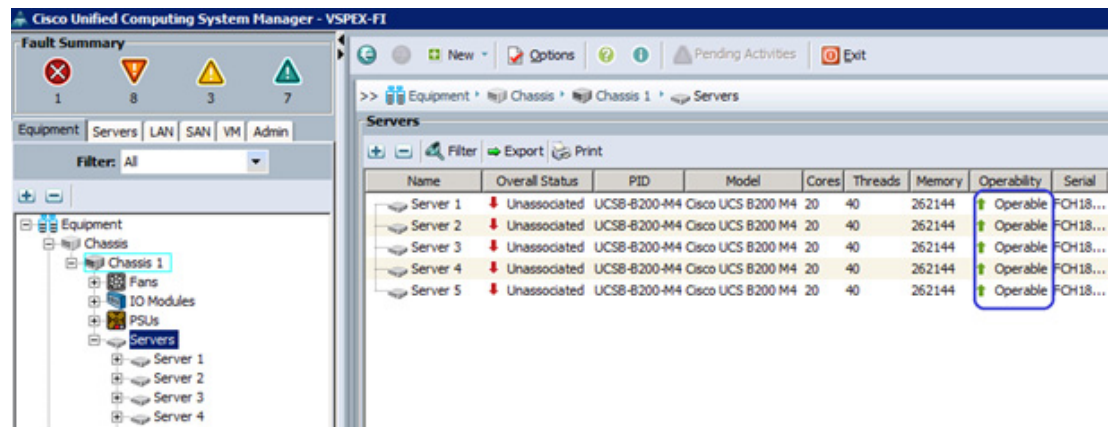
- To discover the Cisco UCS C220 M4 Rack-Mount servers, configure the "Unconfigured" ports on FI-A and FI-B as Server Ports and in few minutes the rack-mount server auto discovery starts. When the servers are discovered, they will display in the "Equipment" tab with an overall status as "Unassociated" and the availability state as "Available", the discovery state as "Complete" as shown in the image below for the rack-mount server:



The blade server's status can be seen as shown in the image below:



- When all the blade servers are discovered, a summary can be viewed by clicking "Equipment" > "Chassis" > "Chassis <id>" > "Servers" as shown in the image below:



In the case of rack-mount servers, from the Equipment tab, under "Equipment" click "Rack-Mounts" then "Servers".

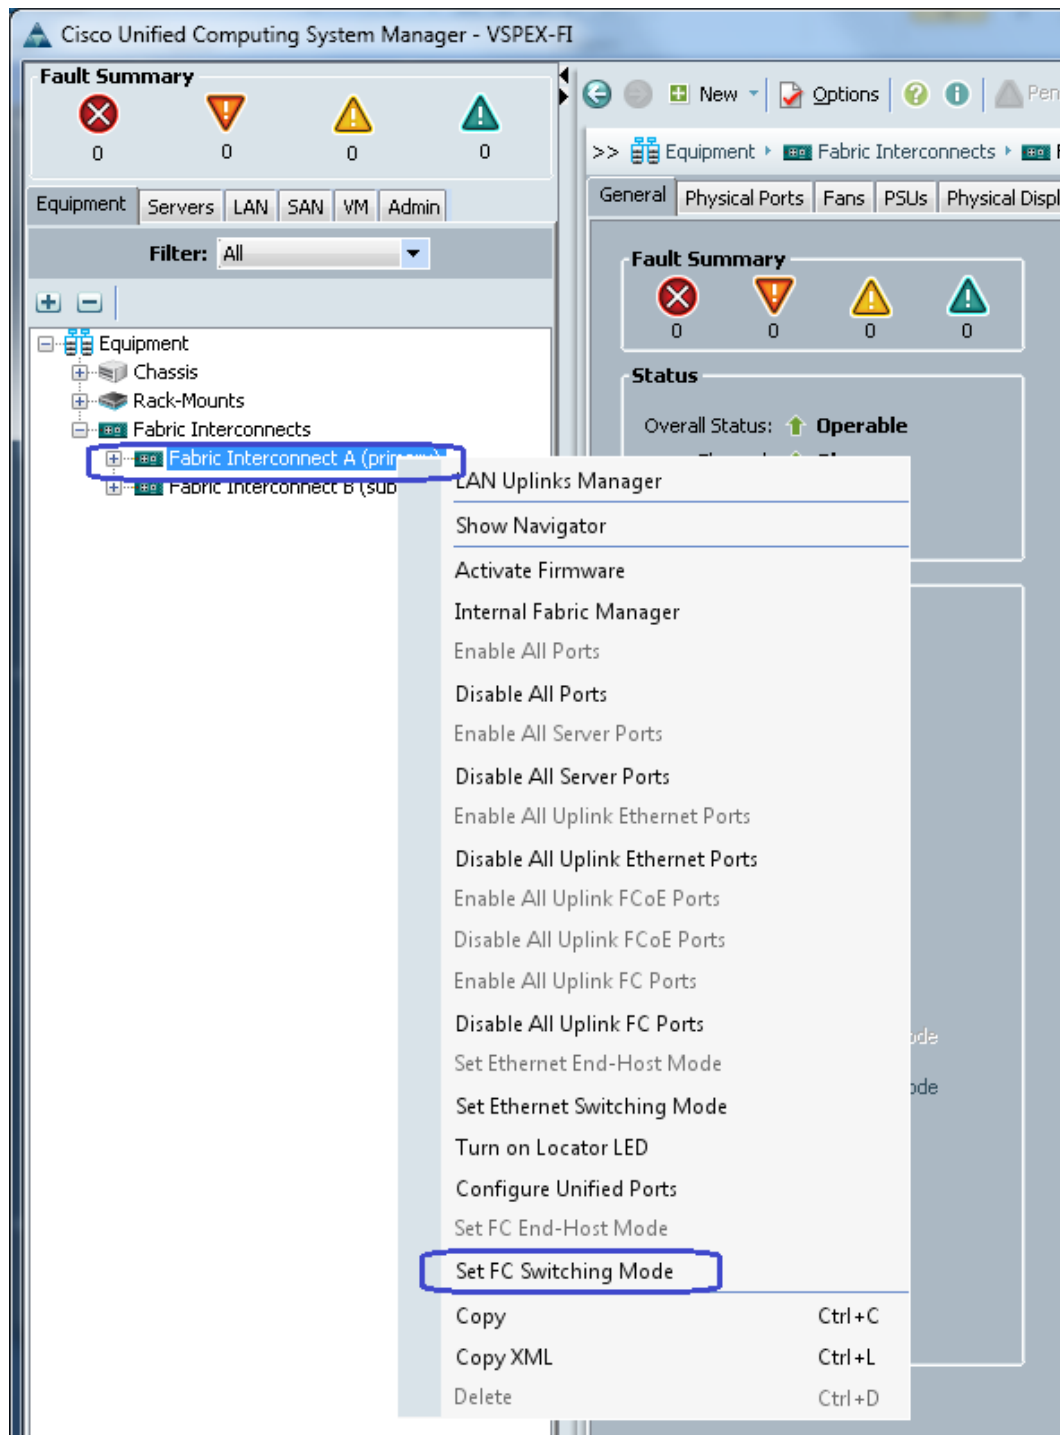
Upstream Global Network Configuration

This subsection includes a few upstream global network configurations listed below:

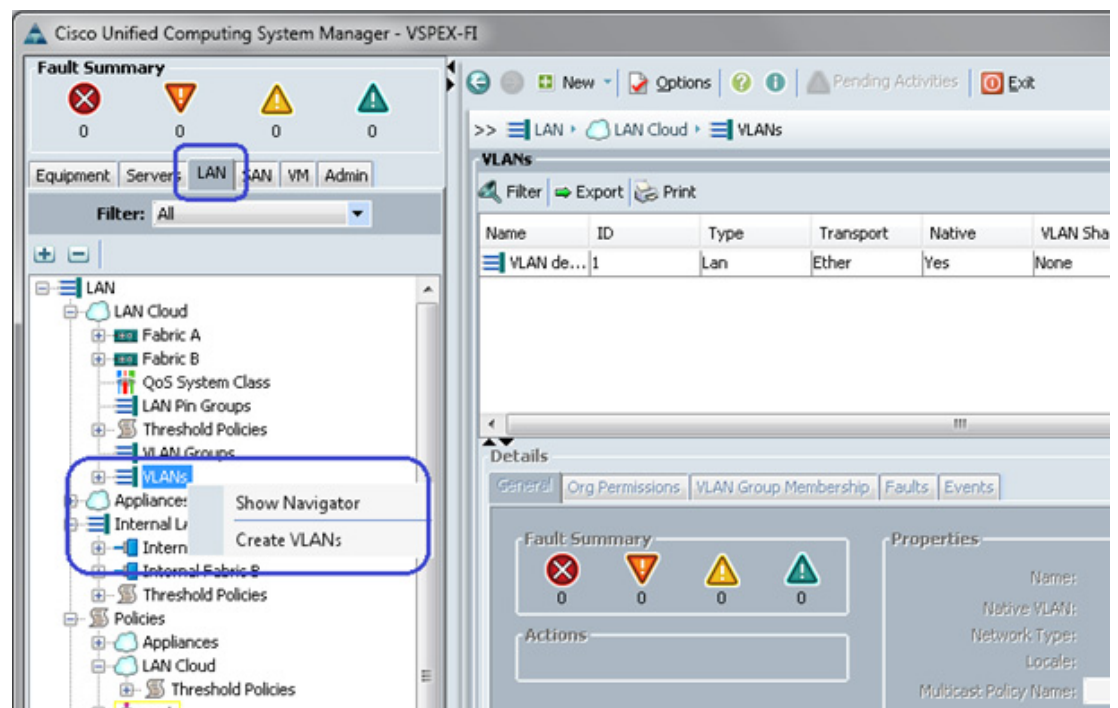
- Move to FC switching mode (FC-variant only)
- Uplink VLAN configuration
- Uplink VSAN configuration (NFS-variant only)
- Appliance VSAN configuration (FC-variant only)
- Configure uplink ports
- Configure universal ports as FC ports
- Configure FC uplink ports (NFS-variant only)
- Configure FC appliance ports (FC-variant only)
- Configure FC Zoning policies (FC-variant only)
- Configure QoS classes and QoS policy for jumbo MTU

To configure the items listed above, complete the following step.

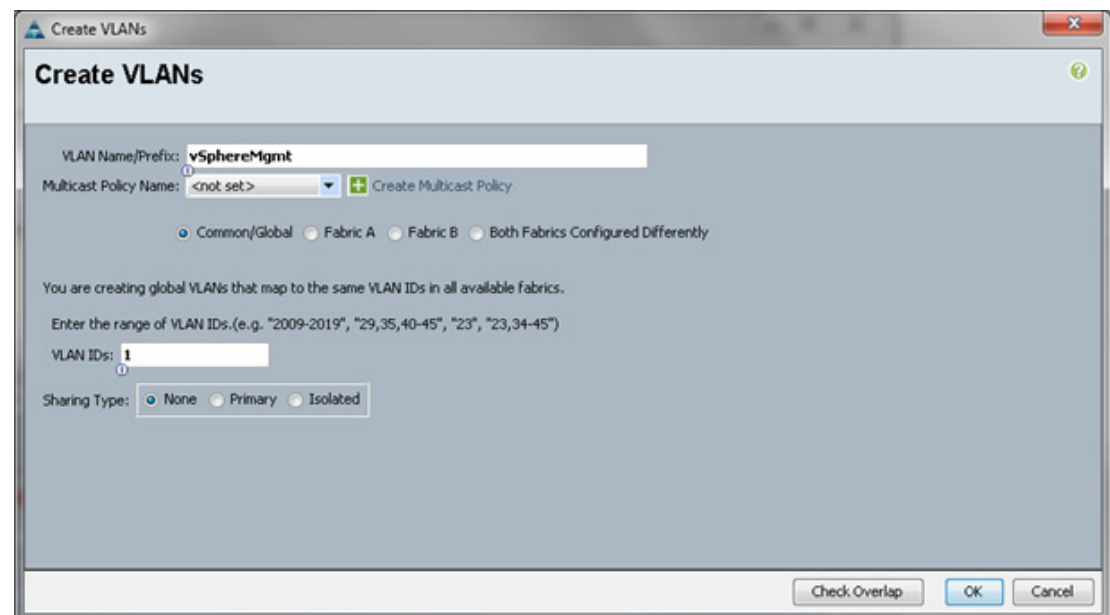
1. (FC-variant only) From the "Equipment" tab, select and right-click "Fabric Interconnect A" and click "Set FC Switching Mode" as shown in the image below:



2. (FC-Variant only) A warning message displays stating that the Fabric Interconnects will be restarted. Click "Yes". Both FIs will reboot (first the secondary FI and then the primary FI). This action is traffic disruptive, so make sure to perform this operation during a maintenance window, if working in a production environment.
3. From the "LAN" tab, expand "LAN" > "LAN Cloud" and right-click "VLANs", then click "Create VLANs" as shown in the image below:

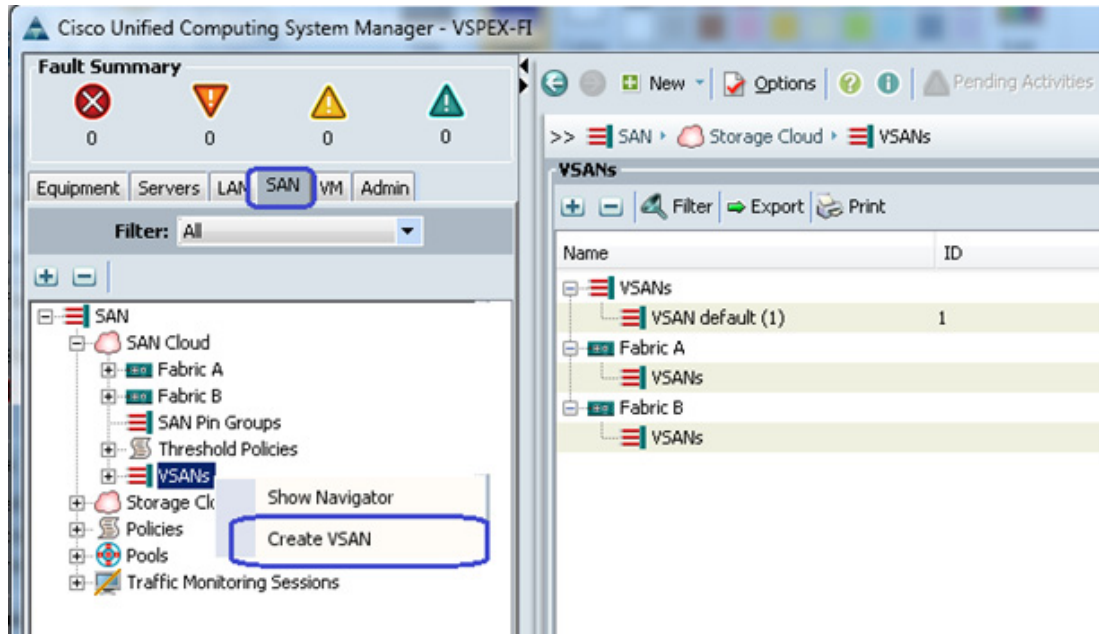


4. Provide a name for the VLAN and assign an ID. Keep the VLAN as default "Common/Global" as shown in the image below:



5. Click "OK" and deploy the VLAN. Repeat these steps for "vSphereMgmt", "VM-Data" and "vMotion" VLANs. For the NFS-variant, also create the "Storage" VLAN. Refer to [Appendix—Customer Configuration Data Sheet](#) for the VLAN values.

- (NFS-variant only) NFS-variant of the architecture uses NFS for VM data access, but still uses FC SAN boot for the hypervisors. Click the "SAN" tab and expand "SAN Cloud", right-click "VSANs". Click "Create VSAN" as shown in the image below:



- (NFS-variant only) Provide a name for the VSAN and provide a VSAN ID and its corresponding FCoE VLAN ID as shown in the image below. The FCoE VLAN ID should not have a conflict with any of the VLANs configured previously. Leave the FC zoning disabled (default).

Create VSAN

Name: **Storage**

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.

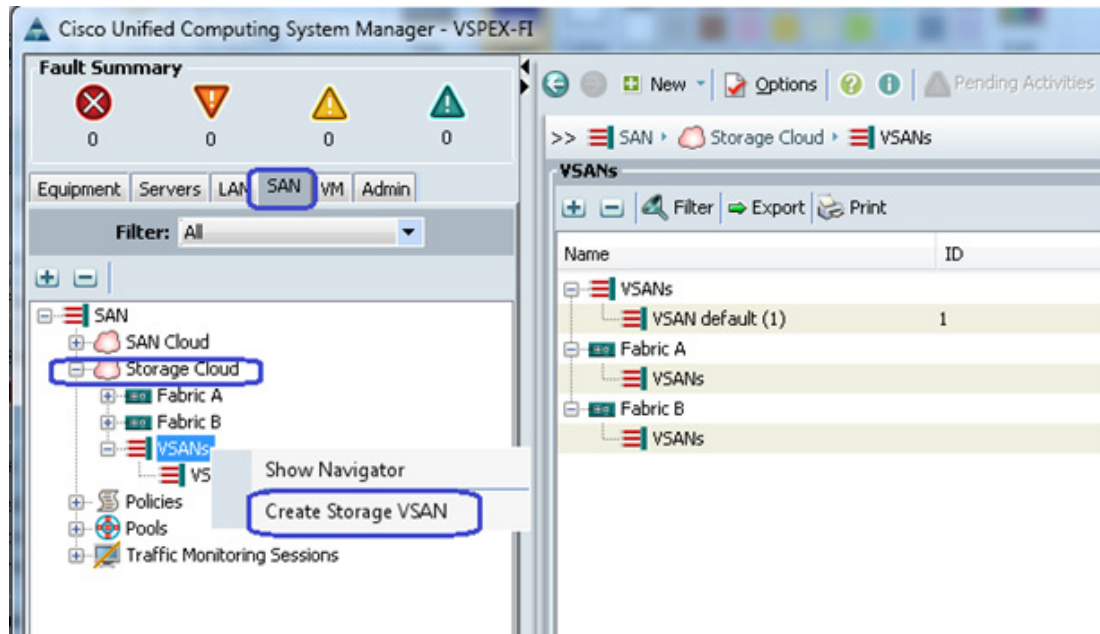
VSAN ID: **10**

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

8. (FC-variant only) Click the "SAN" tab and expand "Storage Cloud", and right-click "VSANs". Click "Create Storage VSAN" as shown in the image below:



9. (FC-variant) Provide a name to the VSAN, enable FC zoning and provide a VSAN ID and its corresponding FCoE VLAN ID as shown below. The FCoE VLAN ID should not conflict with any of the VLANs configured previously.

Create Storage VSAN

Name:

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.

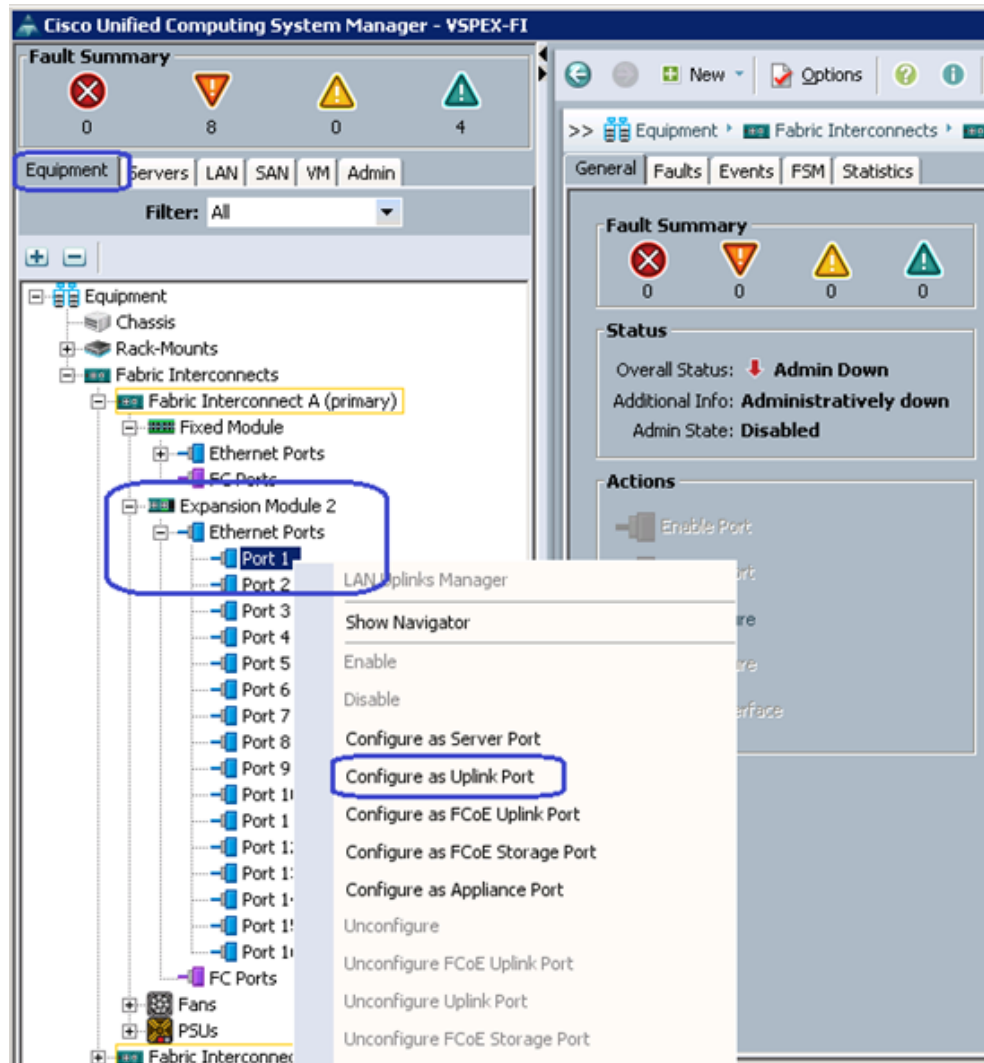
VSAN ID:

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:

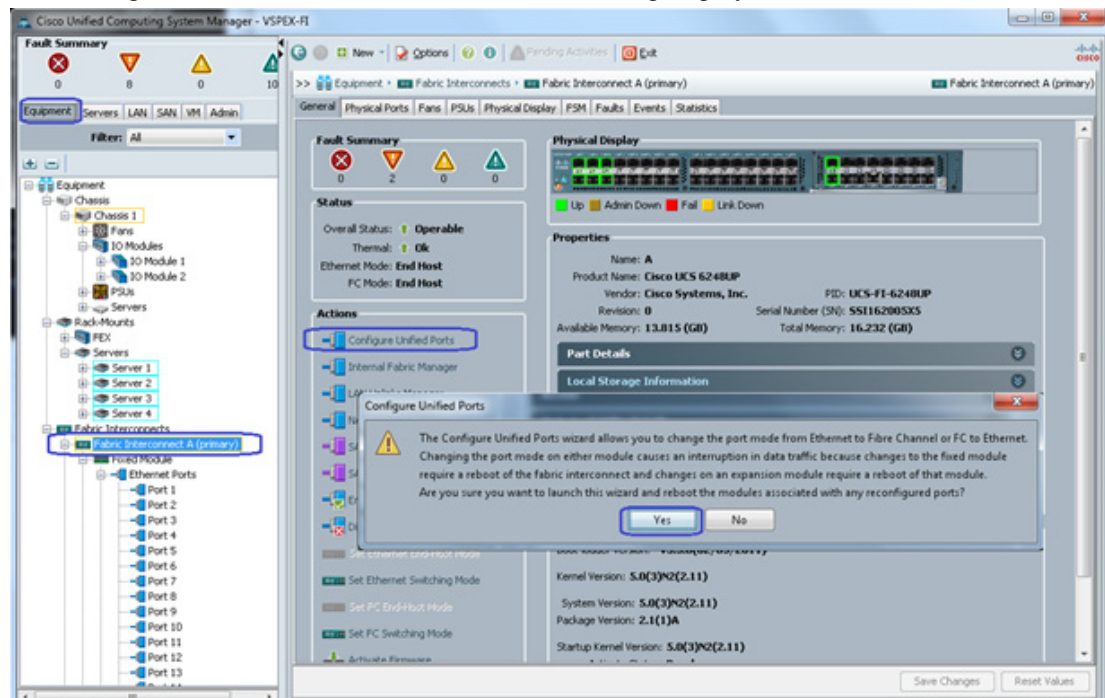
OK Cancel

10. To configure the 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 the uplink port), right-click the Ethernet port, and click "Configure as Uplink Port" as shown below. Repeat this step for all the uplink ports on each FI.

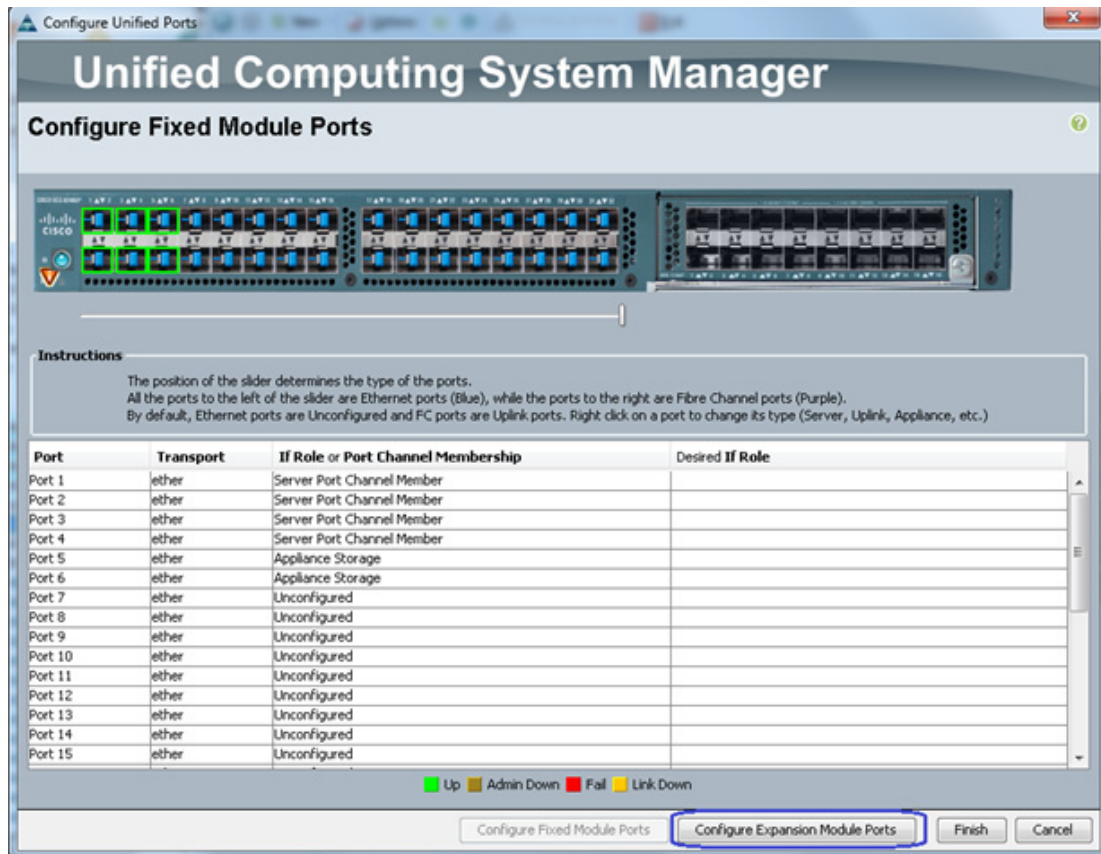


11. The Cisco UCS 6248UP Fabric Interconnects have Universal Ports. The physical ports are 10G Ethernet ports by default, but can be converted into Fibre-Channel ports. FC connectivity to EMC VNX storage array is required for SAN boot; some of the ports need to be converted into FC ports.

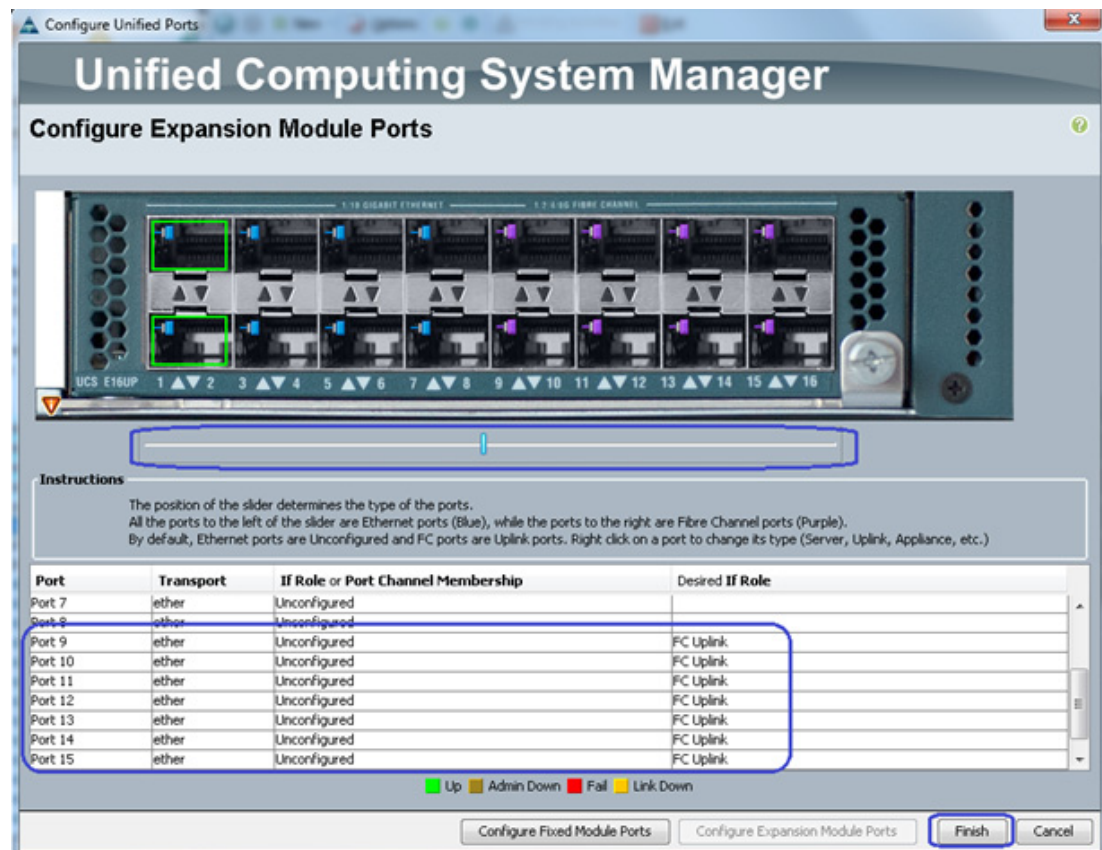
Convert the ports from the expansion module into the FC port. Click the "Equipment" tab, expand "Fabric Interconnects" and click "Fabric Interconnect A". From the links on the right, click the "Configure Unified Ports" as shown below. A warning displays, click "Yes".



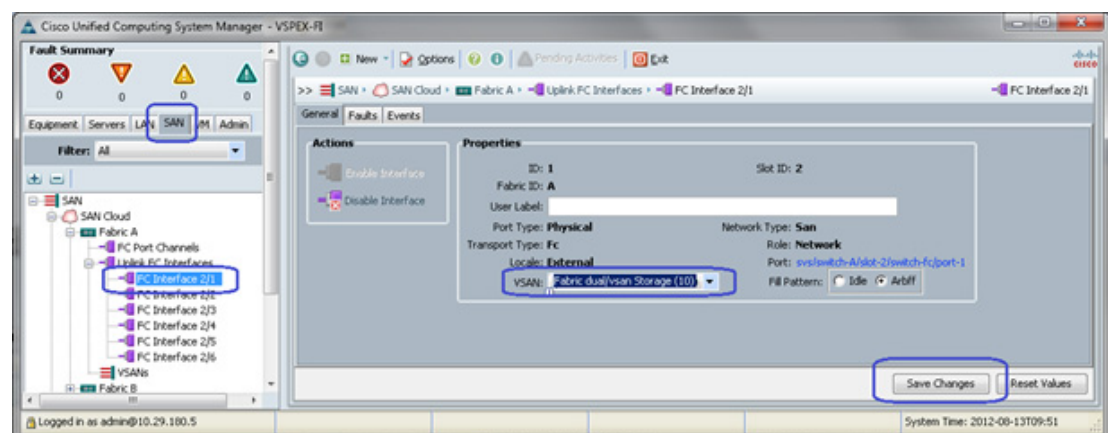
12. From the "Configure Unified Ports" wizard, click "Configure Expansion Module Ports".



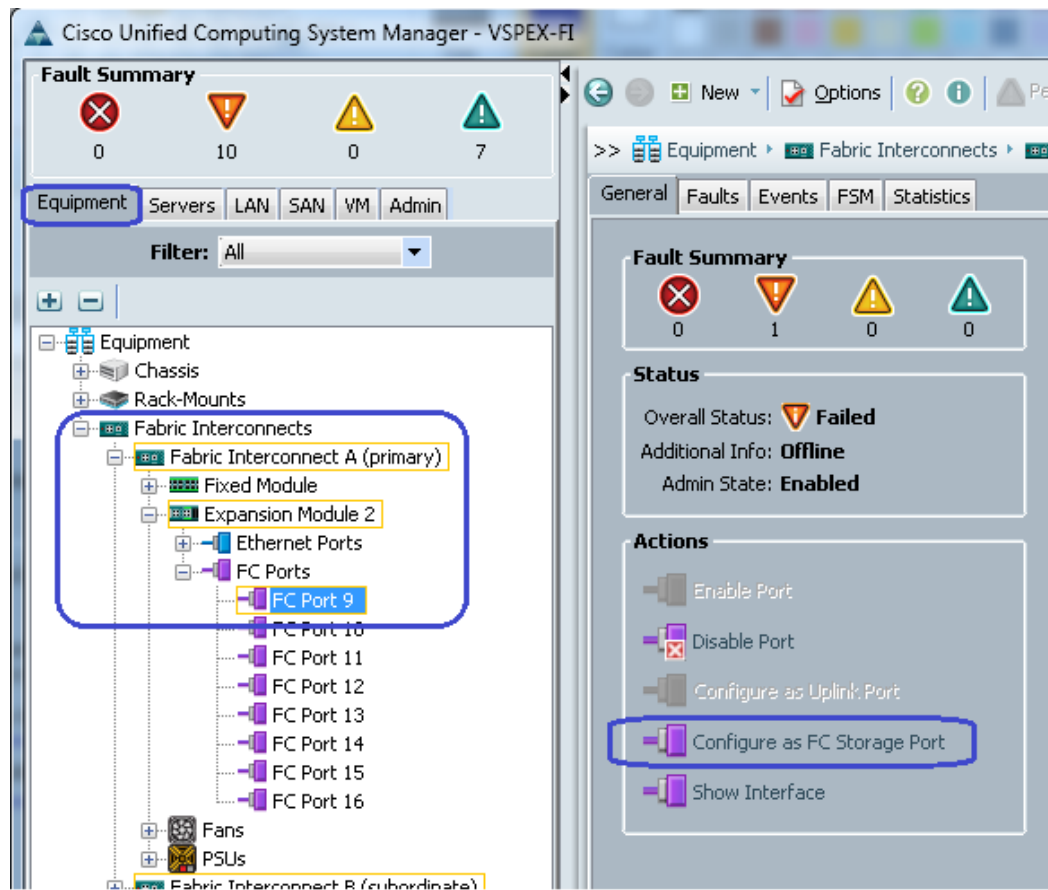
13. Select the slider bar at the top to set in the middle. Make sure that ports 2/9 to 2/15 display "FC Uplink" as shown in the image below. Click "Finish". A pop-up a warning message displays about the FI immediately rebooting. Click "OK" to reboot the FI.



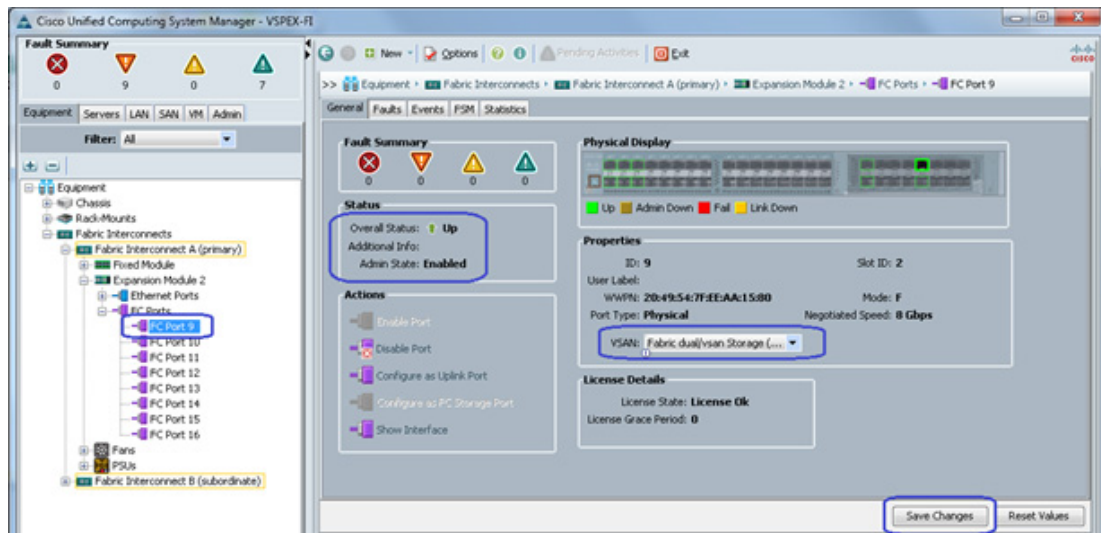
14. When the FI is rebooted, repeat steps 12 and 13 for FI-B.
15. (NFS-Variant only) Click the "SAN" tab, expand "SAN Cloud" > "Fabric A" > "Uplink FC Interfaces", and choose the FC interface. Change the VSAN to the Storage VSAN that was created in steps 6 and 7 as shown below, and click "Save Changes".



16. (FC-Variant only) The Physical FC ports need to be classified as FC storage ports for the attached storage array. Click the "Equipment" tab, expand "Fabric Interconnect" > "Fabric Interconnect A" > "Expansion Module 2" > "FC Ports" and click the individual FC port as shown below. From the right side menu, click "Configure as FC Storage port" link.



17. (FC-Variant only) Make sure that the port comes up as shown below. From the "VSAN" drop-down list, select the Storage VSAN configured in steps 9 and 10, and click "Save Changes".



18. (FC-Variant only) The EMC VNX storage array will Fibre-Channel flogi into the FIs. Using the WWPN of the VNX storage array, specify the zoning policy on the FI. Use the SSH connection to the Cisco UCS Manager Virtual IP address and issue "connect nxos a" command. In the read-only NXOS shell, issue "show flogi database" command and note the WWPN of the storage array as high-lighted in the image below:

```

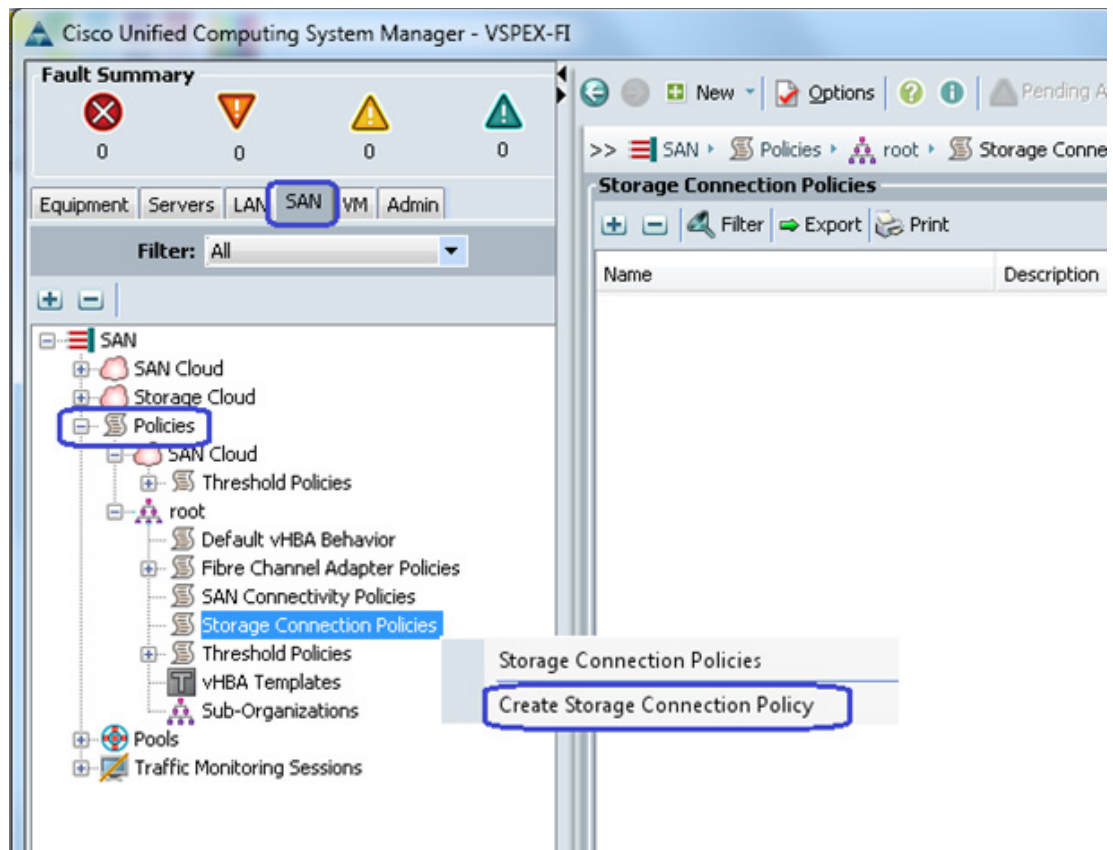
10.65.121.228 - PuTTY
VSPEX-FI-A# connect nxos a
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and 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)#

```

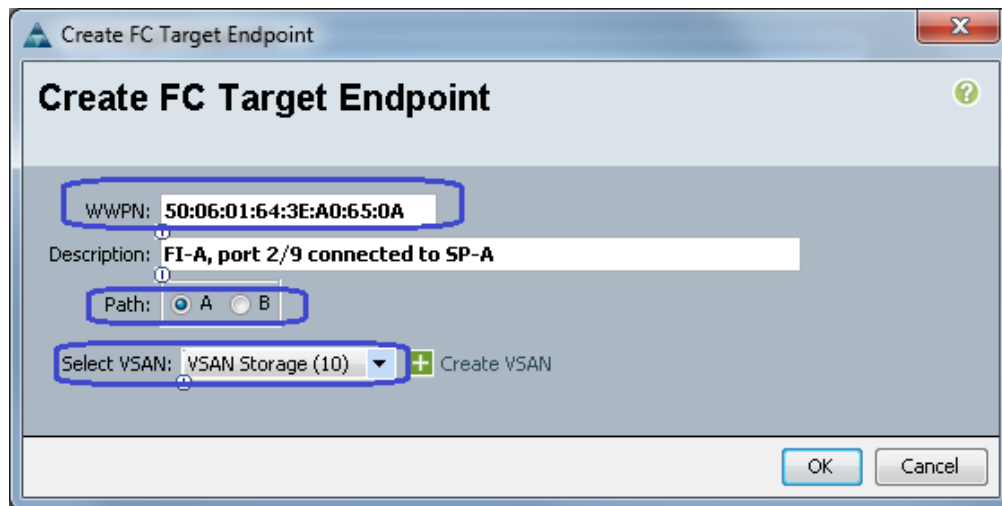
19. (FC-Variant only) From the Cisco UCS Manager GUI, click "SAN", expand "SAN" > "Policies" > "root" and right-click "Storage connection policies" and click "Create Storage Connection Policy" as shown in the image below:



20. (FC-Variant only) Enter the name "Fabric-A" and an optional description. Select "Single Initiator Multiple Targets" as the Zoning Type. Click + to add a new FC Target Endpoint.



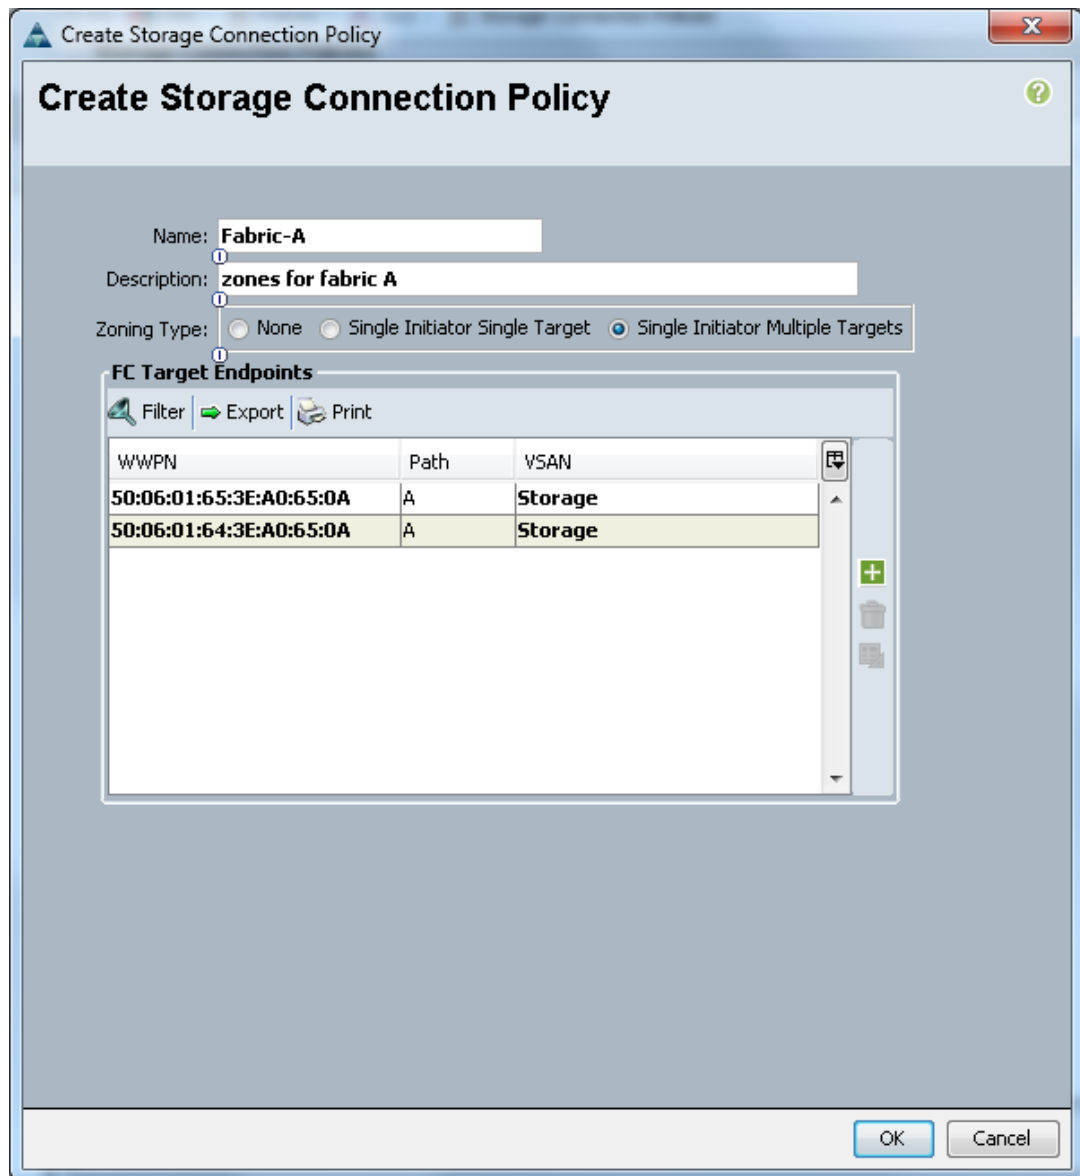
21. (FC-Variant only) Copy the WWPN from the "show flogi database" output from step 18 and paste it into the WWPN field. Provide an optional description, select Path "A" and Storage VSAN from the drop-down menu as shown in the image below:



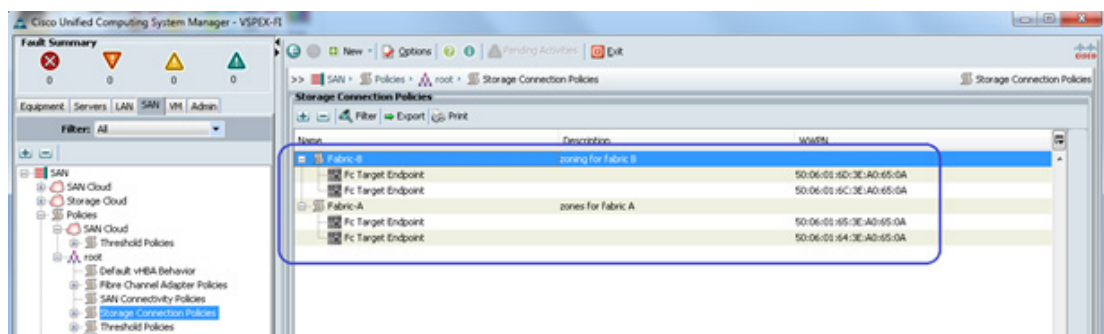
The image shows a 'Create FC Target Endpoint' dialog box. It has a title bar with a close button (X) and a help icon (?). The main area contains the following fields and controls:

- WWPN:** A text field containing '50:06:01:64:3E:A0:65:0A'.
- Description:** A text field containing 'FI-A, port 2/9 connected to SP-A'.
- Path:** Radio buttons for 'A' (selected) and 'B'.
- Select VSAN:** A dropdown menu showing 'VSAN Storage (10)'.
- Create VSAN:** A green button with a plus sign.
- OK** and **Cancel** buttons at the bottom right.

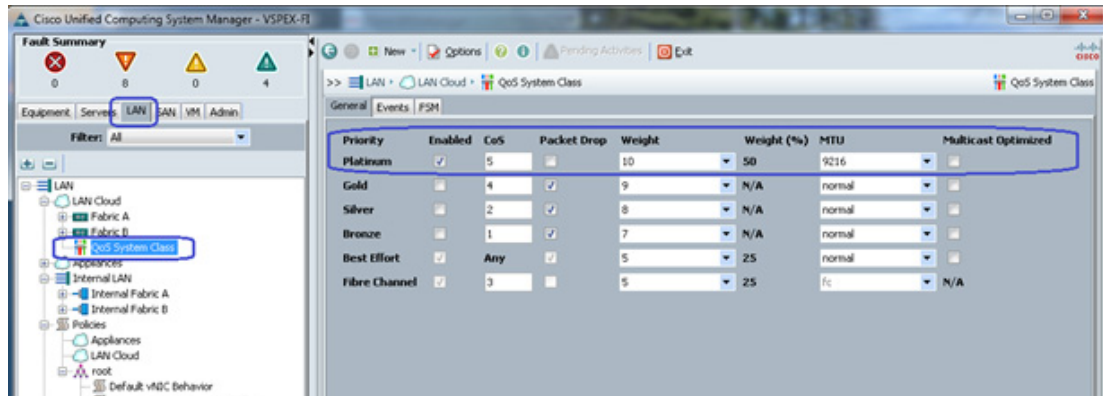
22. (FC-Variant only) Add the second FC target end-point for Fabric A and click OK.



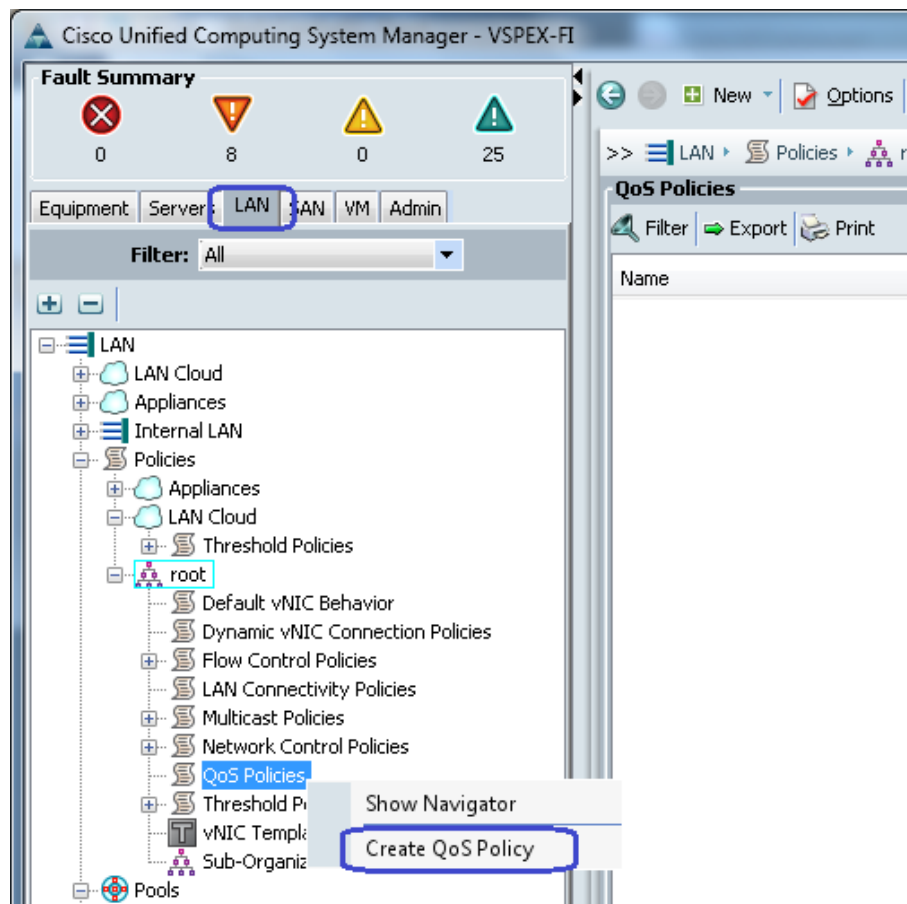
23. (FC-Variant only) Repeat steps 18 to 22 for Fabric B. The result will look similar to the image shown below:



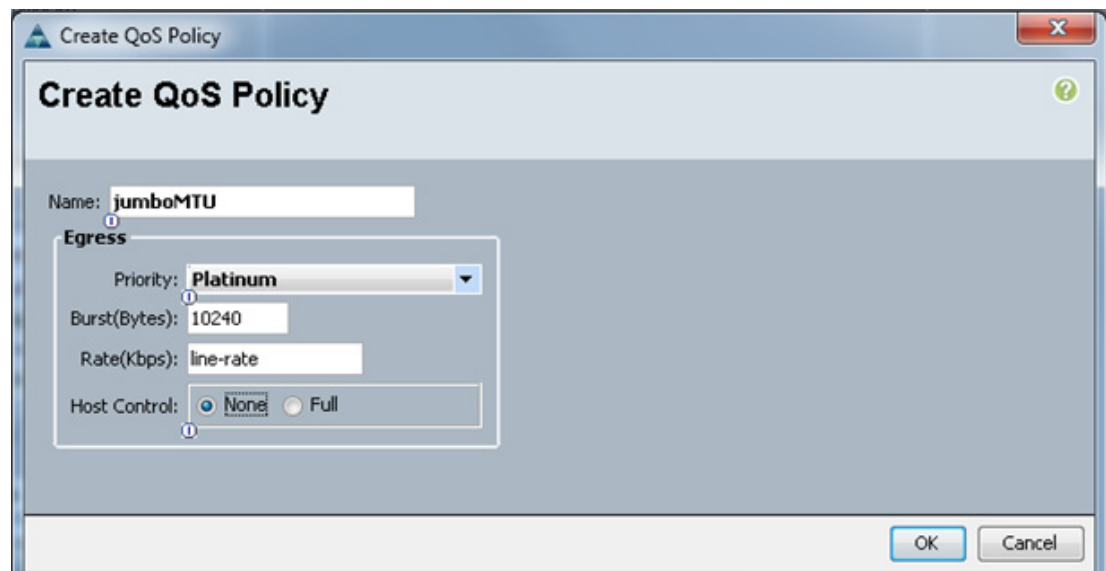
24. Configure QoS. Click the "LAN" tab, expand "LAN" > "LAN Cloud" and click "QoS System Class". Enable "Platinum" priority and set MTU to "9216". Keep the other configuration set as default and save the configuration.



25. From the "LAN" tab, expand "LAN" > "Policies" > "root", and right-click "Create QoS Policy" as shown in the image below:



26. Create a QoS policy named "jumboMTU" and select Priority "Platinum". Click "OK" to save the configuration.



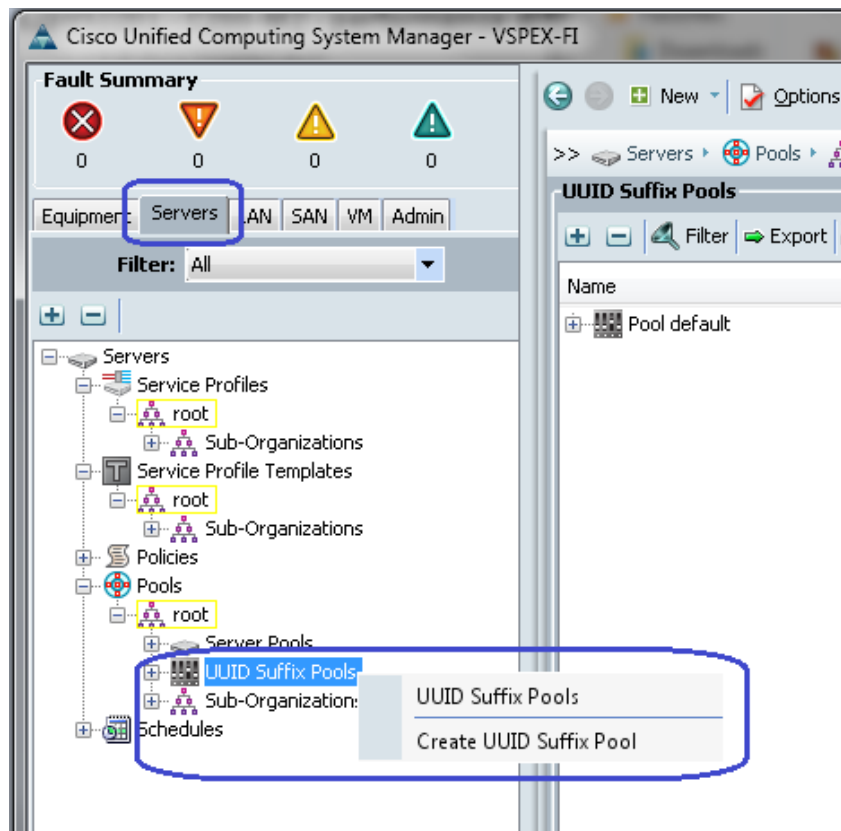
Configure Identifier Pools

In this section, the following identifier pools, used by the service profiles, will be configured:

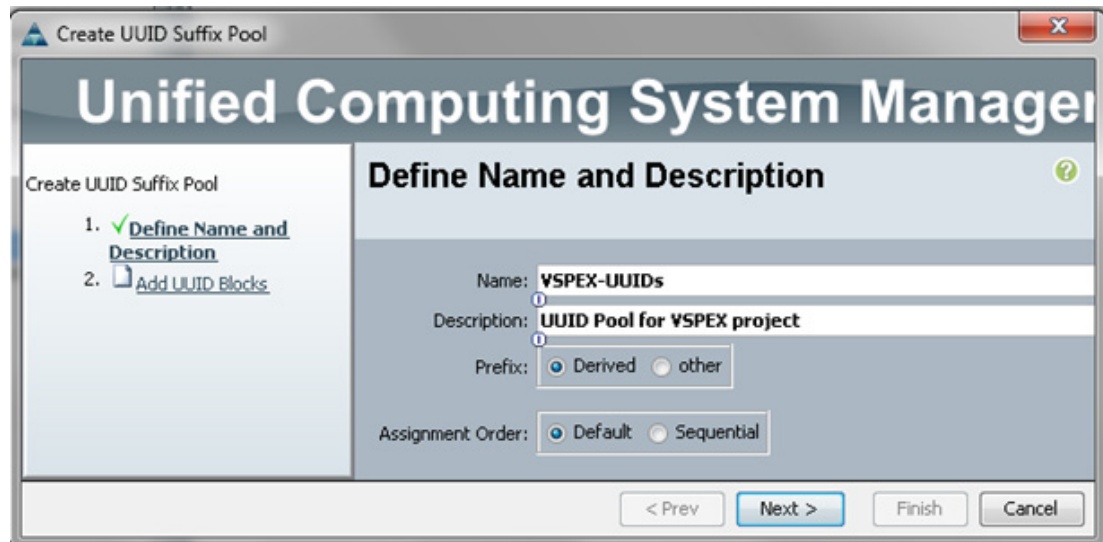
- Server UUID pool
- MAC address pool
- WWN pool
- Management IP address pool

To configure the identifier pools, complete the following steps:

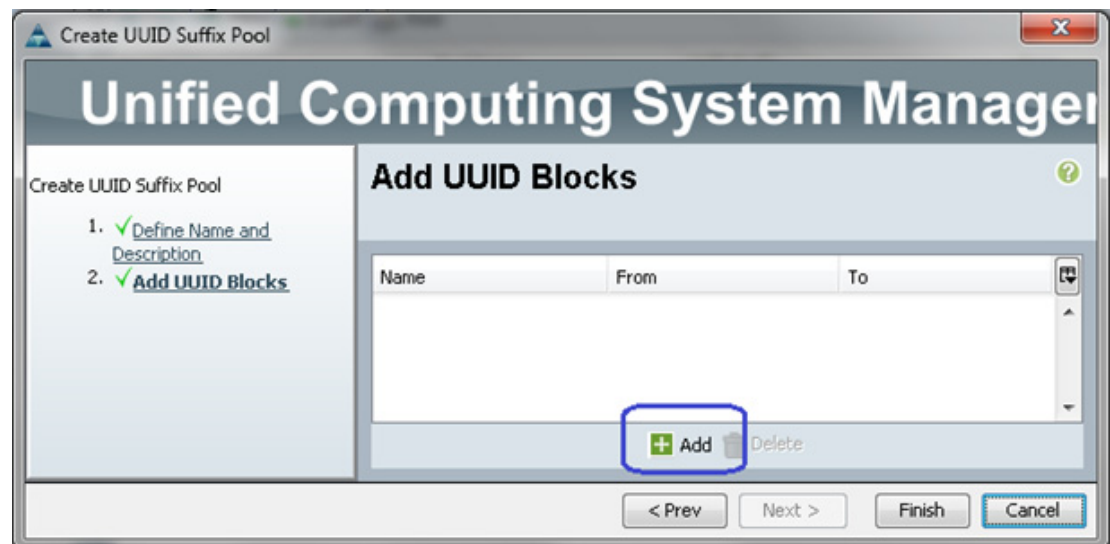
1. From the "Servers" tab, expand "Servers" > "Pools" > "root", and right-click "UUID Suffix pools" and click "Create UUID Suffix Pool" as shown in the image below.



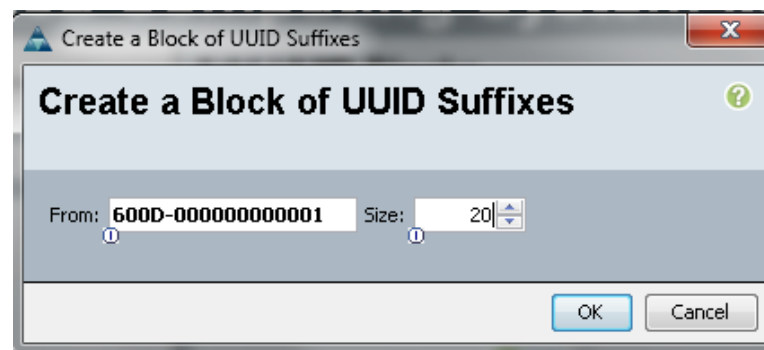
2. Provide a name and description for the UUID suffix pool. Keep the other configuration as default.



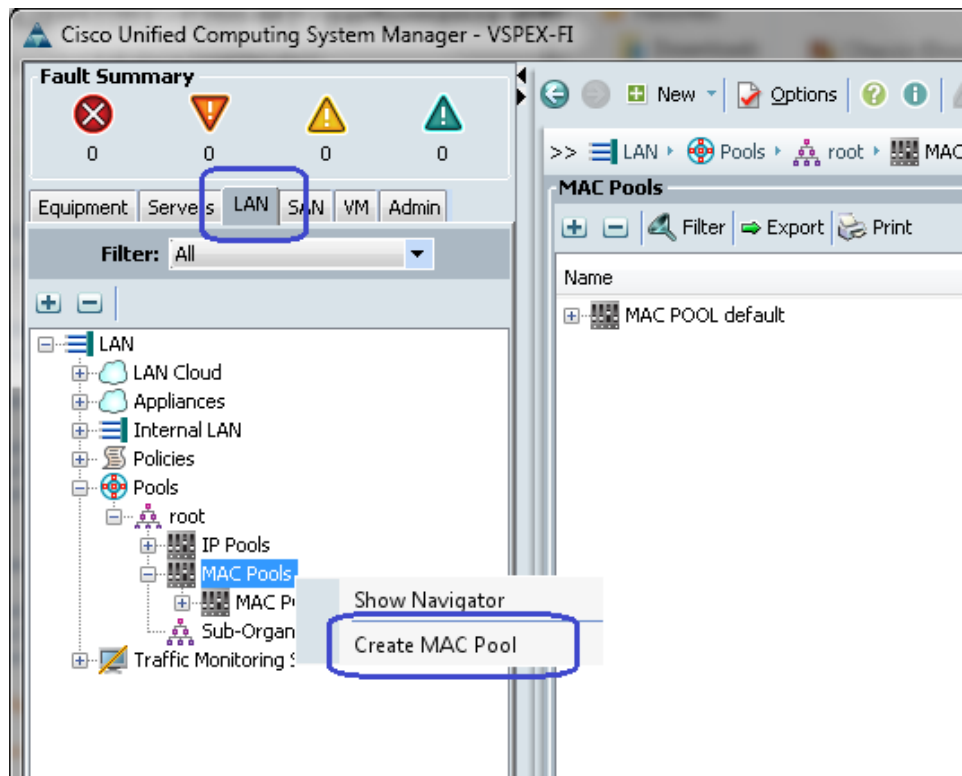
3. Click "Add" to add UUID block.



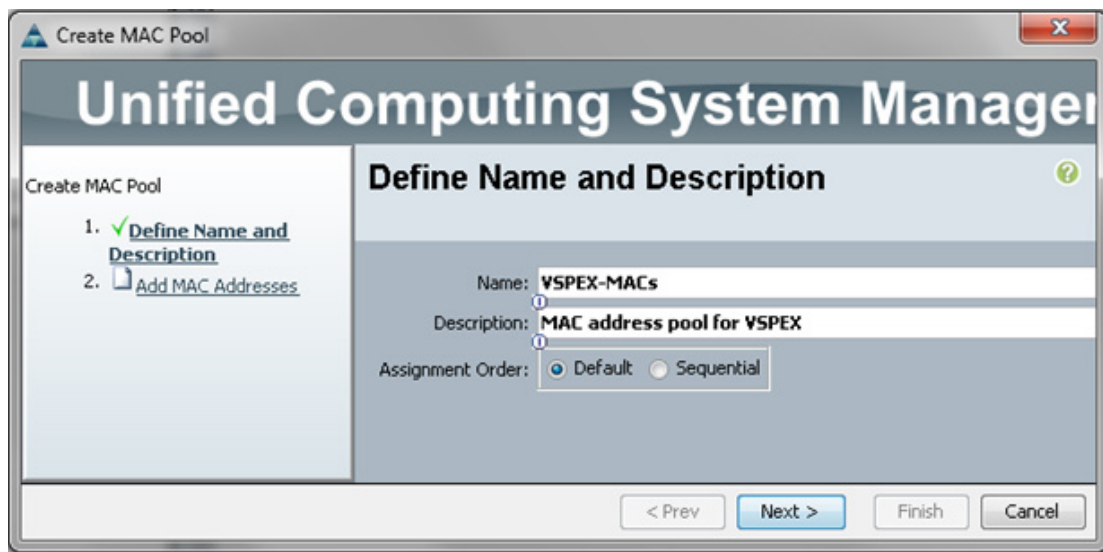
4. Specify the beginning of the UUID Suffixes and select a large size of UUID Suffixes to accommodate future expansion.



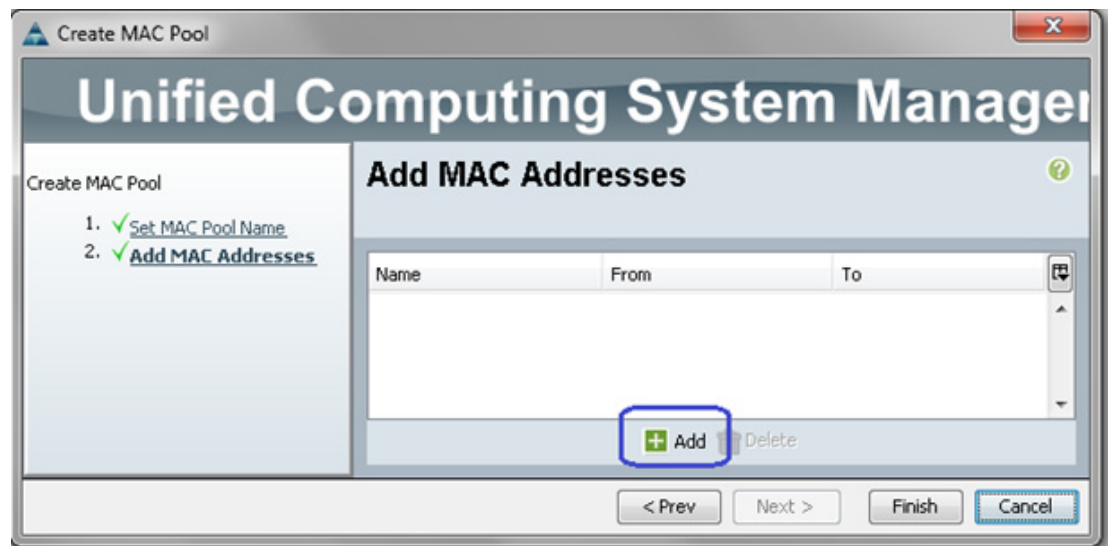
5. Click "OK" and then "Finish" to deploy the UUID pool.
6. Go to the "LAN" tab, expand "LAN" > "Pools" > "root", right-click "MAC Pools" and click "Create MAC Pool" as shown in the image below:



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



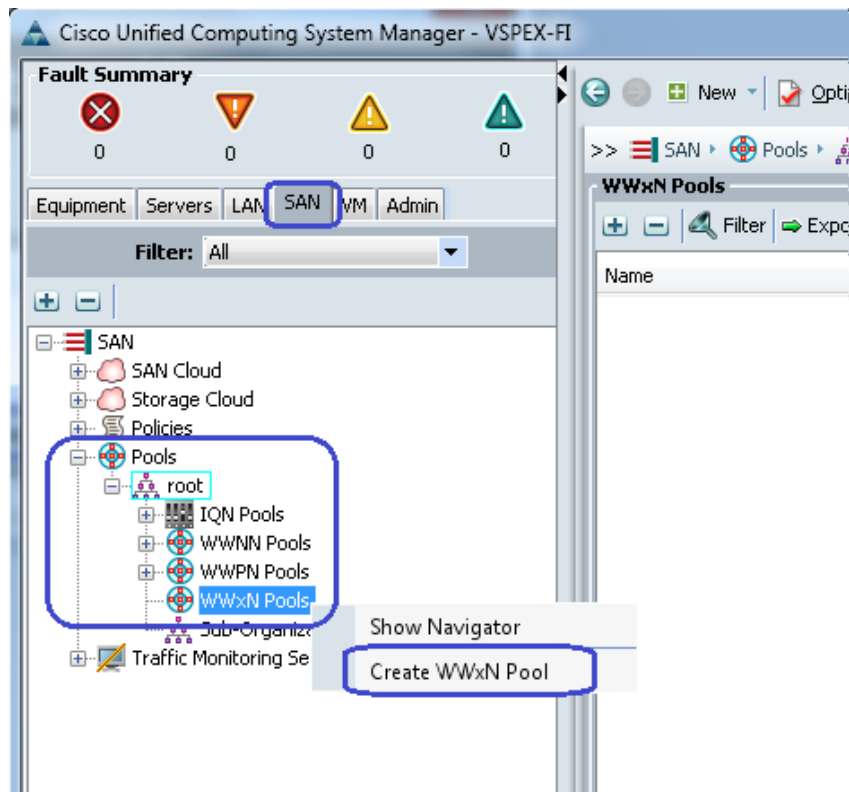
8. Click "Add" to add MAC pool block.



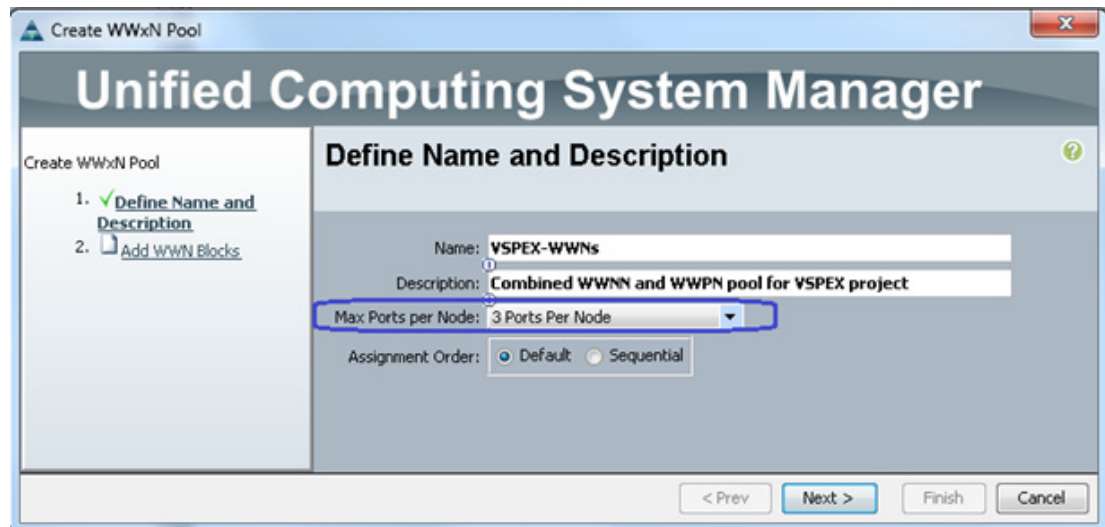
9. Provide the initial MAC address and size of the block. Provide a large number of MAC addresses to accommodate future expansion. It is recommended to have six MAC addresses per server.



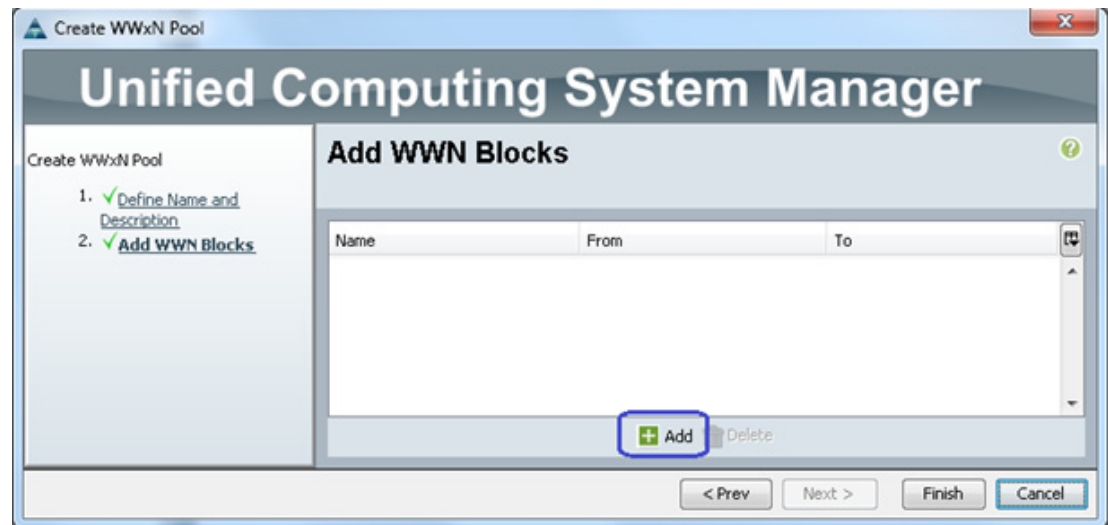
10. Click "OK" and "Finish" to complete the configuration.
11. Go to the "SAN" tab, expand "SAN > Pools" > "root", right-click "WWxN Pools" and click "Create WWxN Pool" as shown in the image below:



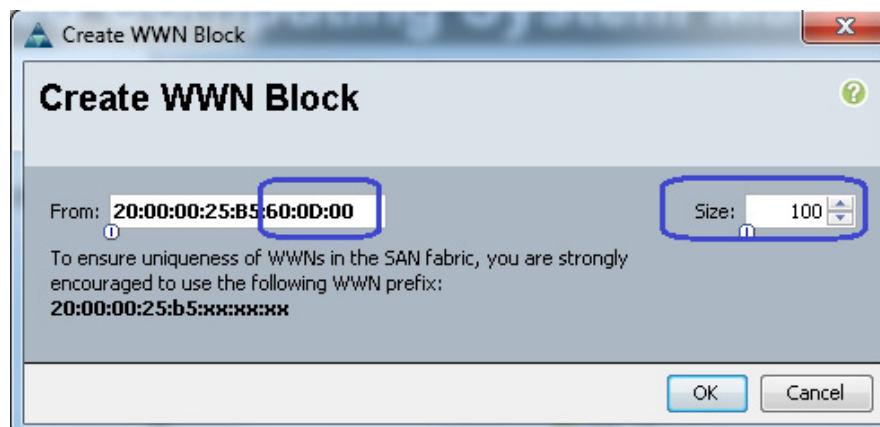
12. Provide a name, description and select "3 Ports per Node" from the drop-down menu as show in the image below:



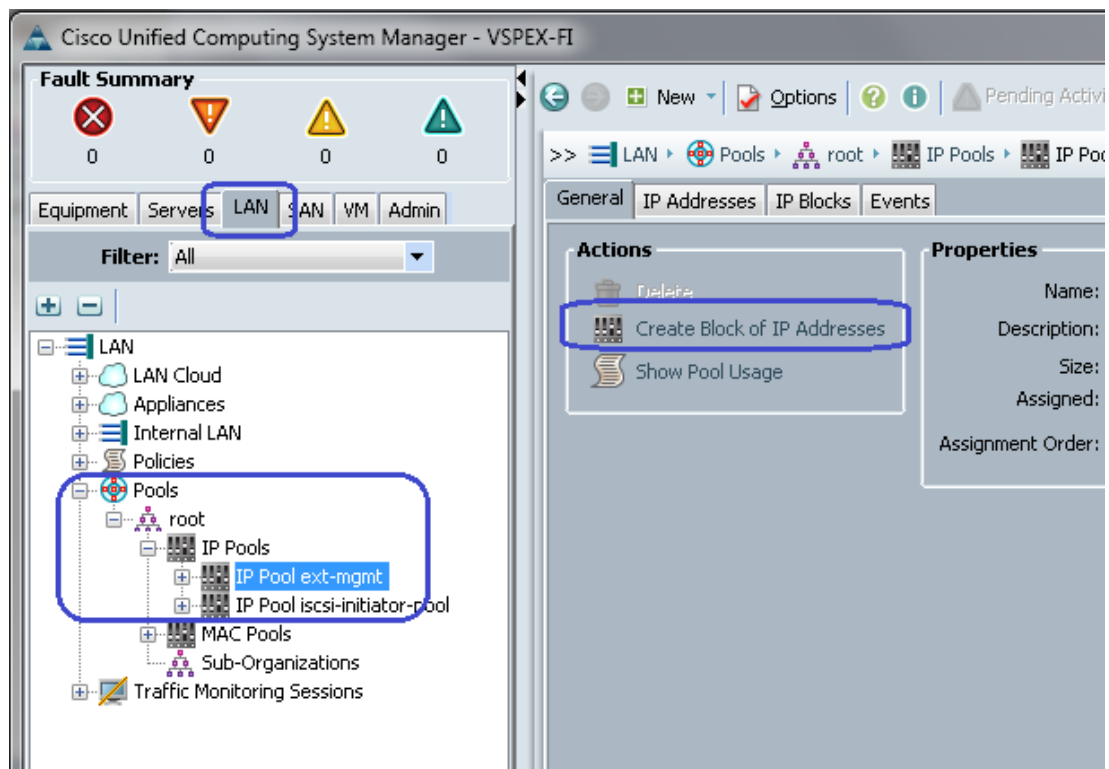
13. Click "Add" to add a block of WWxN IDs.



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



15. Create the management IP address block for KVM to access the servers. The default pool for the server CIMC management IP addresses are created with the name "ext-mgmt". Go to the "LAN" tab, expand "LAN" > "Pools" > "root" > "IP Pools" and select "IP Pool ext-mgmt" and click "Create Block of IP addresses" link as shown in the image below:



16. Provide the initial IP address, size of the pool, default gateway and subnet mask as shown below. Click "OK" to deploy the configuration. The IP addresses will be assigned to various rack-mount server CIMC management access from this block.

This concludes the configuration of all identifier pools and blocks.

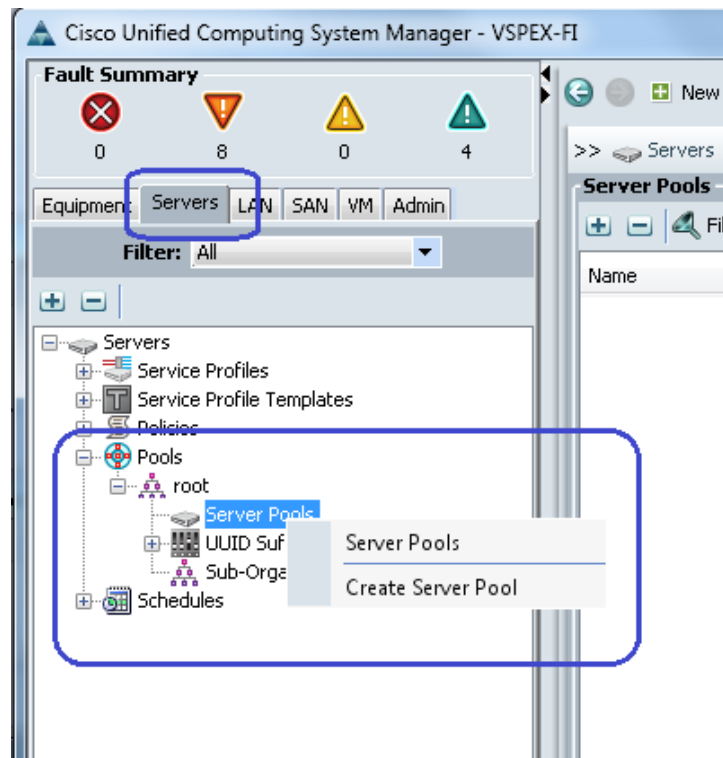
Configure Server Pool and Qualifying Policy

The creation and policy-based auto-population of server pools can be divided in to following tasks:

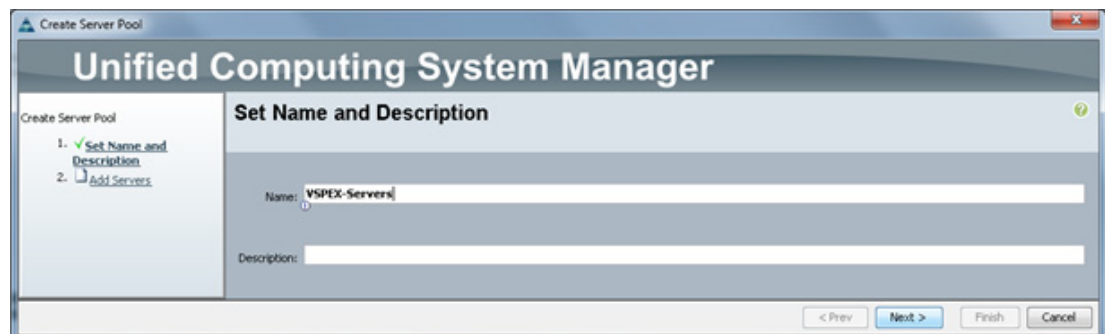
- Creation of server pool
- Creation of server pool policy qualification
- Creation of server pool policy

To configure the server pool and qualifying policy, complete the following steps:

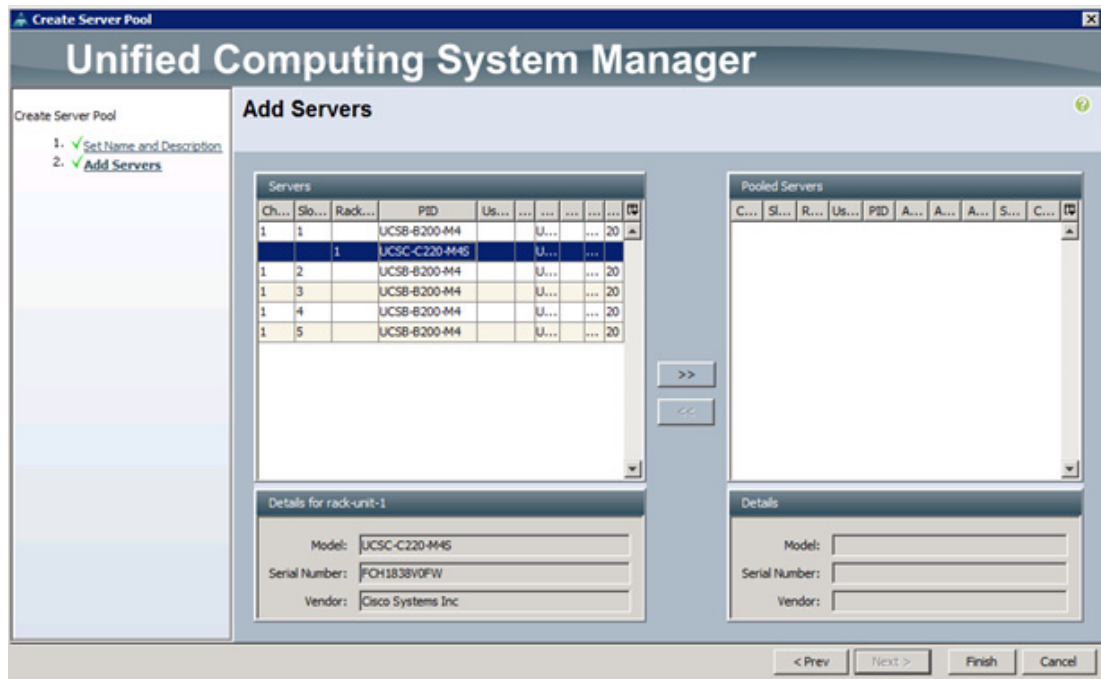
1. Go to the "Servers" tab, expand "Servers" > "Pools" > "root", right-click "Server Pools" and click "Create Server Pool".



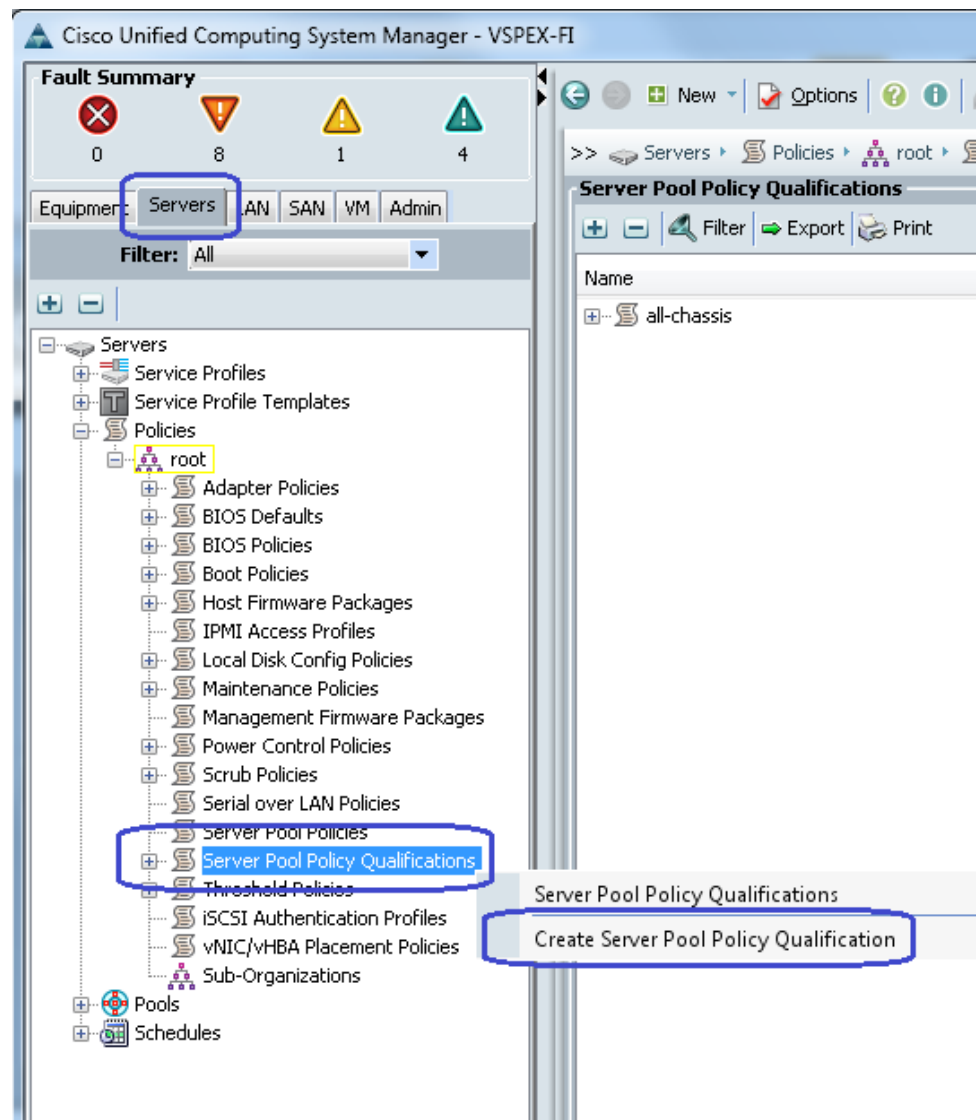
2. Enter a name for the server pool, and click Next.



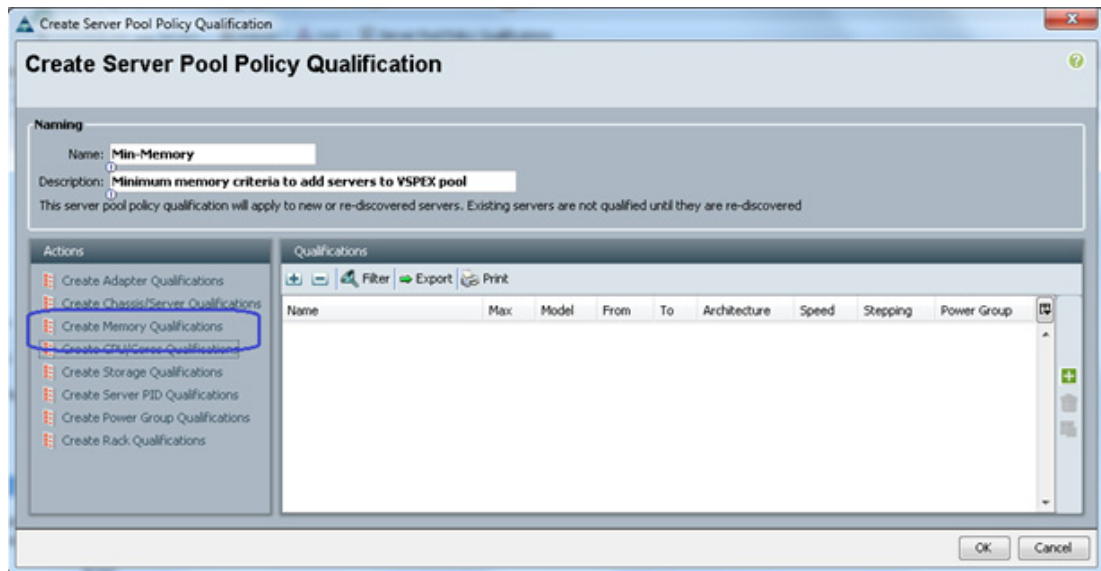
3. Click "Finish" to create the empty server pool. Add the compute resources to this pool dynamically, based on policy.



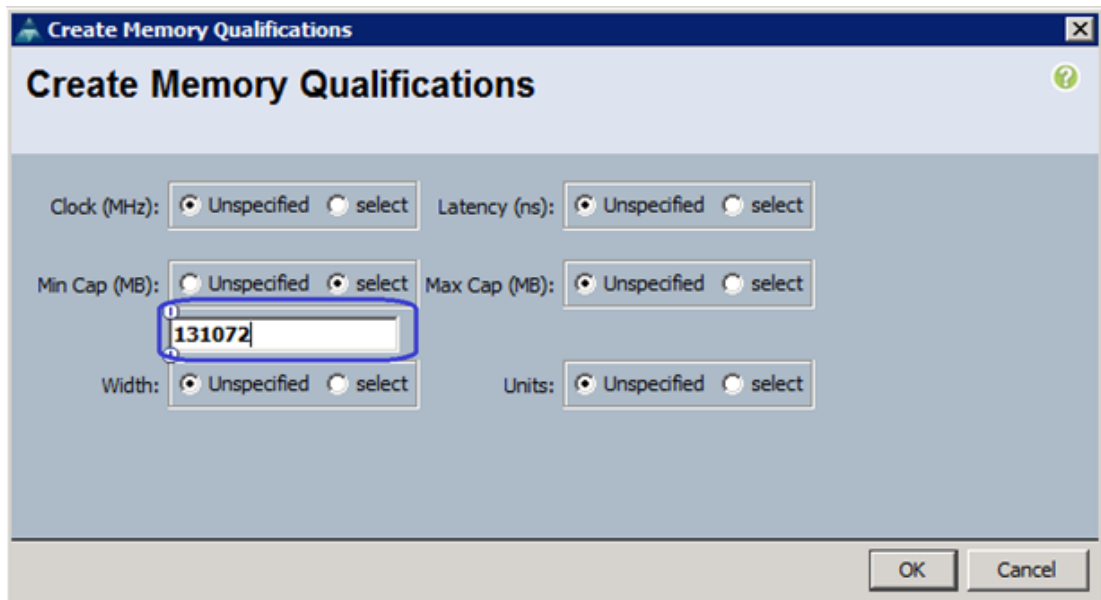
- From the "Servers" tab, expand "Servers" > "Policies" > "root", right-click "Server Pool Policy Qualifications" and click "Create Server Pool Policy Qualification" as shown in the image below:



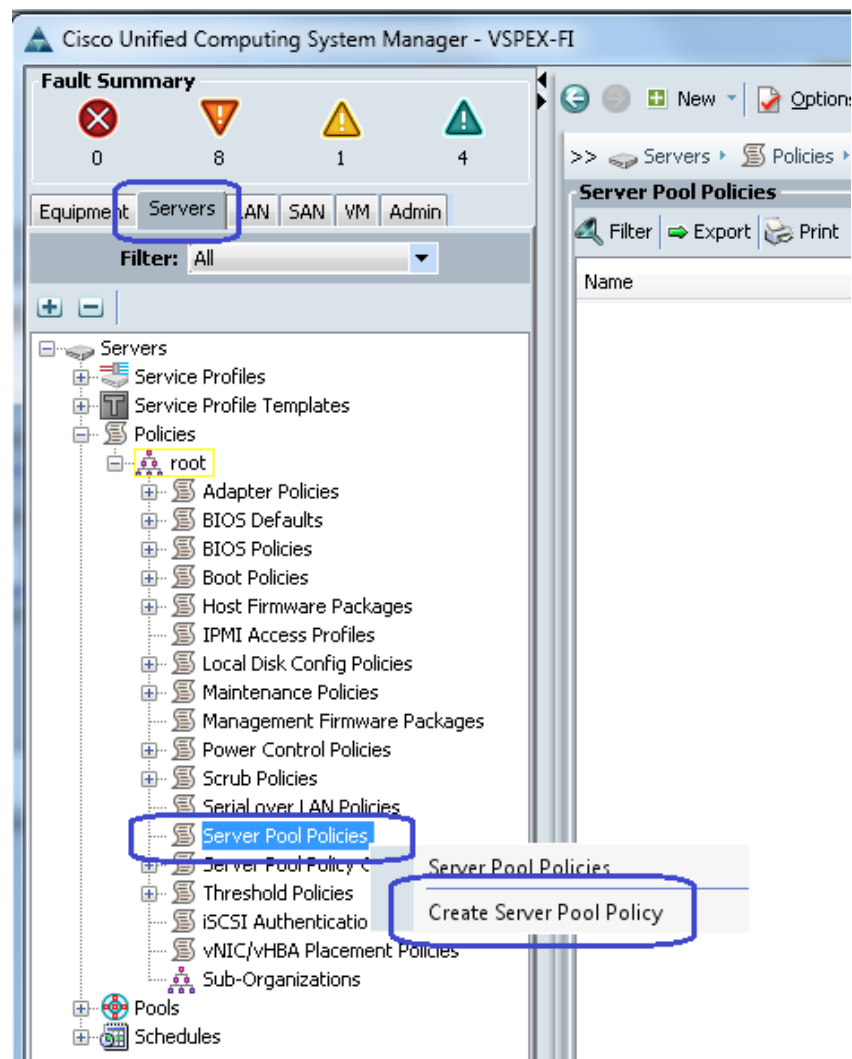
5. Provide a name for the server policy qualification criterion. For example, select memory qualification criterion.




6. Set a minimum of 128 GB RAM as the pool qualification criterion. Note that this is an example criterion; choose a criterion that suites your requirement. Click "OK" twice to create the qualification.



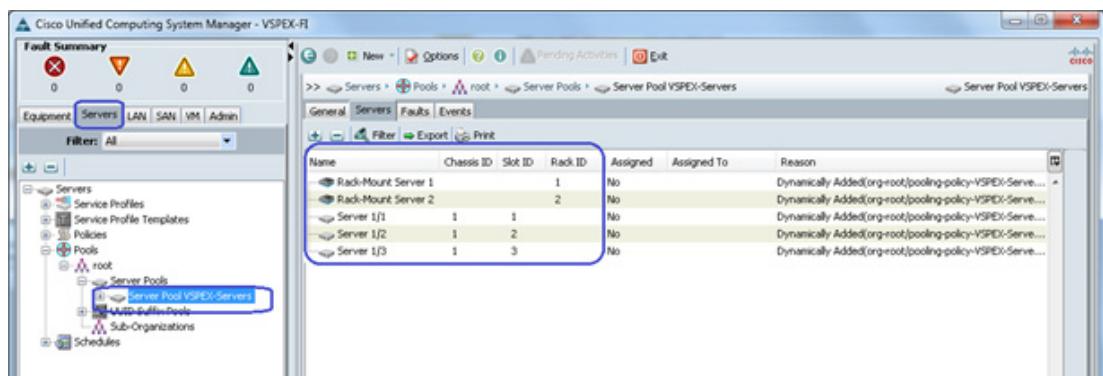
7. Click "Servers" > "Policies" > "root", right-click "Server Pool Policies" and click "Create Server Pool Policy".



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



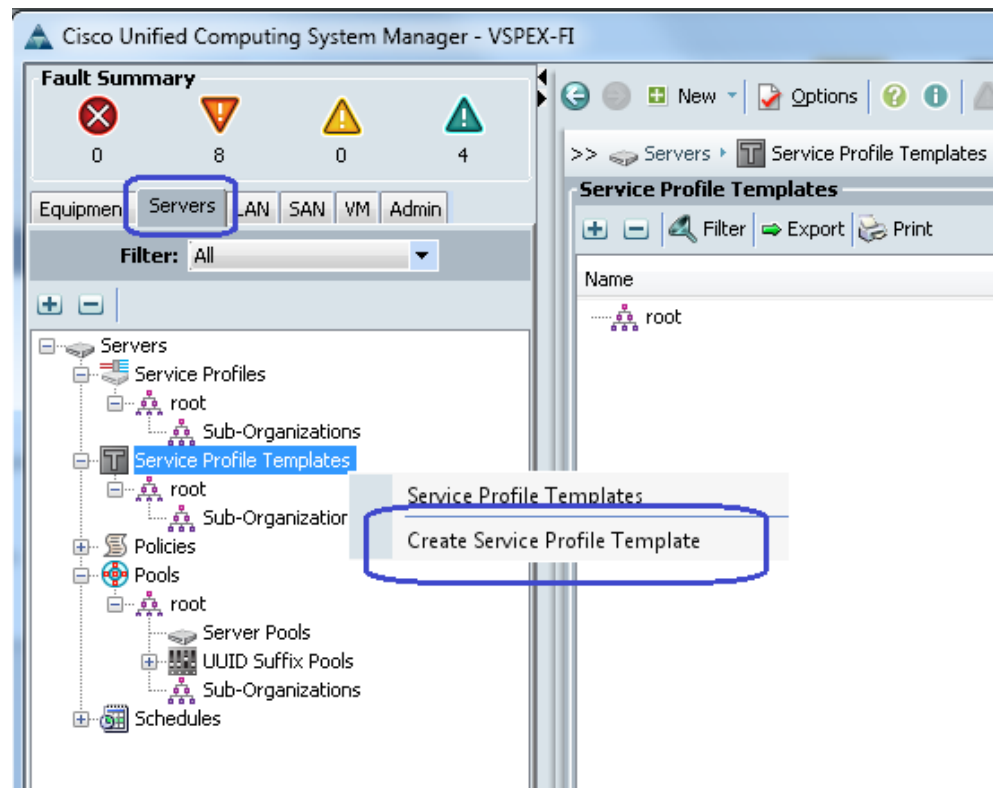
- Return to the server pool created in step 1 above and click the "Servers" tab to view all the compute resources that meet the qualification criteria and are dynamically added to the server pool as shown below. The image below is taken from the FC-Variant of the architecture, where a combination of Cisco UCS B200 M4 blade servers and Cisco UCS C220 M4 rack-mount servers are used to share the workload. This architecture showcases the form-factor independent architecture using Cisco UCS Manager.



Configure Service Profile Template

To instantiate individual service profiles, a service profile template needs to be configured. To create the service profile template, complete the following steps:

- Go to the "Servers" tab. Expand "Servers" > "Service Profile Templates", right-click and click "Create Service Profile Template".



2. Provide a service profile template name, keep the type as "Initial Template", and choose UUID pool for UUID assignment as shown in the image below:

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: VSPEX-Service-Profile

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: VSPEX-UUIDs(20/20)

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. Select configure LAN connectivity as "Expert". Click "Add" to create a vNIC.

Create Service Profile Template

Unified Computing System Manager

Networking

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) Create Dynamic vNIC Connection Policy

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

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

Name	MAC Address	Fabric ID	Native VLAN

Delete Add Modify

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

Name	Overlay vNIC Name	iSCSI Adapter Policy	MAC Address

Add Delete Modify

< Prev Next > Finish Cancel

4. Create a system VNIC for Fabric A. Provide "system-A" VNIC name, choose the MAC pool created previously, choose the fabric ID as "fabric A", select "vMotion" and "vSphereMgmt" VLANs with "vSphereMgmt" as native VLAN. Choose 9000 MTU, "VMware" adapter policy and "jumboMTU" QoS policy as shown in the image below:

Create vNIC

Name:

Use vNIC Template: ☐

[+ Create vNIC Template](#)

MAC Address

MAC Address Assignment:

[+ Create MAC Pool](#)

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"/>

[+ Create VLAN](#)

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: [+ Create LAN Pin Group](#)

Operational Parameters

Adapter Performance Profile

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

Dynamic vNIC Connection Policy: [+ Create Dynamic vNIC Connection Policy](#)

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

Network Control Policy: [+ Create Network Control Policy](#)

OK Cancel

5. Create one more VNIC with the exact same properties on Fabric B.
6. (NFS-variant only) Create two more VNICs similar to steps 3, 4 and 5 for the NFS server access. Provide the name "Storage-A" and "Storage-B" for VNICs on Fabric A and B, choose only "Storage" VLAN and mark it as native VLAN and choose "VMware" and "jumboMTU" for adapter policy and QoS policy.

7. Create a VNIC for VM data traffic. Provide the "data-A" for VNIC name, same MAC address pool name, "fabric A" as fabric ID, "VM-Data" as native VLAN, and "VMware" adapter policy as shown in the image below:

Create vNIC

Name:

Use 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

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	Storage	<input type="radio"/>
<input checked="" type="checkbox"/>	VM-Data	<input checked="" type="radio"/>
<input type="checkbox"/>	vmotion	<input type="radio"/>
<input type="checkbox"/>	Shared-Mgmt	<input type="radio"/>

MTU:

Pin Group:

Operational Parameters

Adapter Performance Profile

Adapter Policy:

Dynamic vNIC Connection Policy:

QoS Policy:

Network Control Policy:

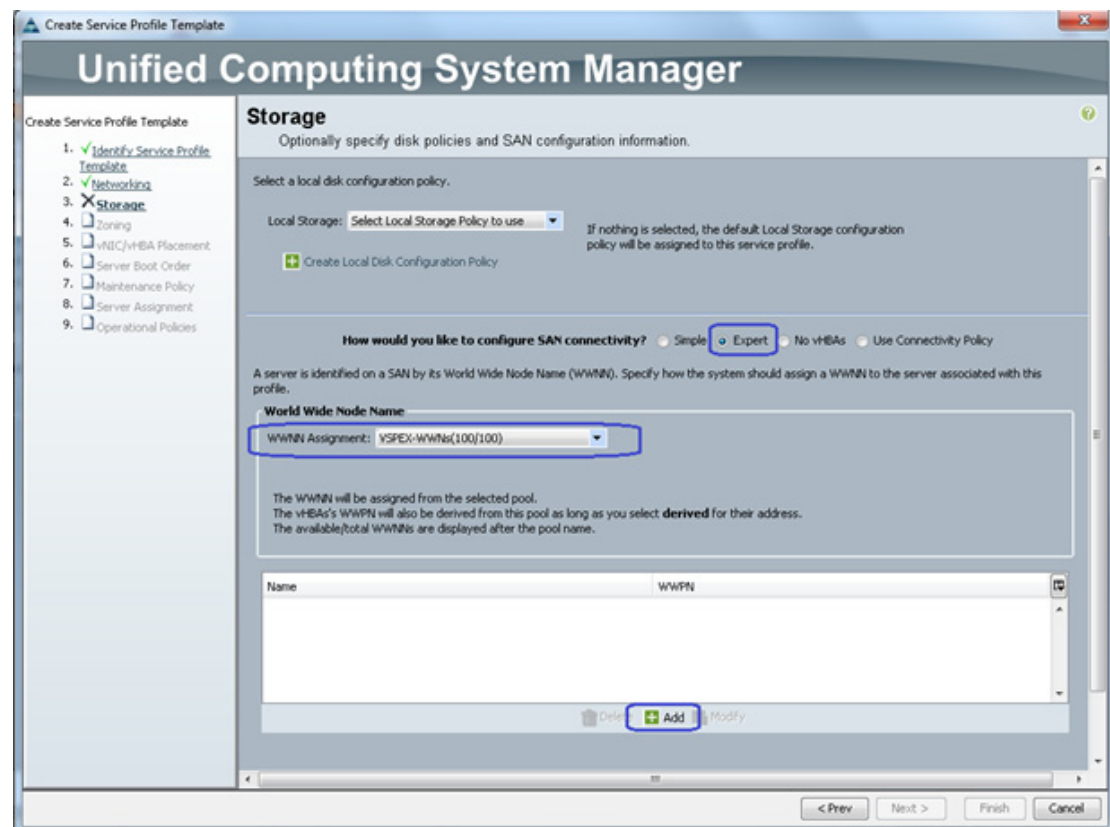
8. Create one more VNIC for fabric B for VM data traffic. [Table 14](#) summarizes all the VNICs created on the service profile:

Table 14 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.

- From the "Storage" page of the wizard, choose the "Expert" radio button for SAN connectivity and select the "VSPEX-WWNs" WWNN pool created from the drop-down menu as shown in the image below. Click "Add" to add vHBA.



- Create a vHBA with vHBA-A name, keep the WWPN assignment as "Derived", select fabric ID as "A", select VSAN as Storage VSAN from drop-down menu, and select Adapter policy as "VMWare" as shown in the image below. Click "OK" to deploy the vHBA.

Create vHBA

Name:

World Wide Port Name:

+ Create vHBA Template

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

Fabric ID: ☐ A ☐ B

Select VSAN: + Create VSAN

Pin Group: + Create SAN Pin Group

Persistent Binding: ☐ Disabled ☐ Enabled

Max Data Field Size:

Operational Parameters

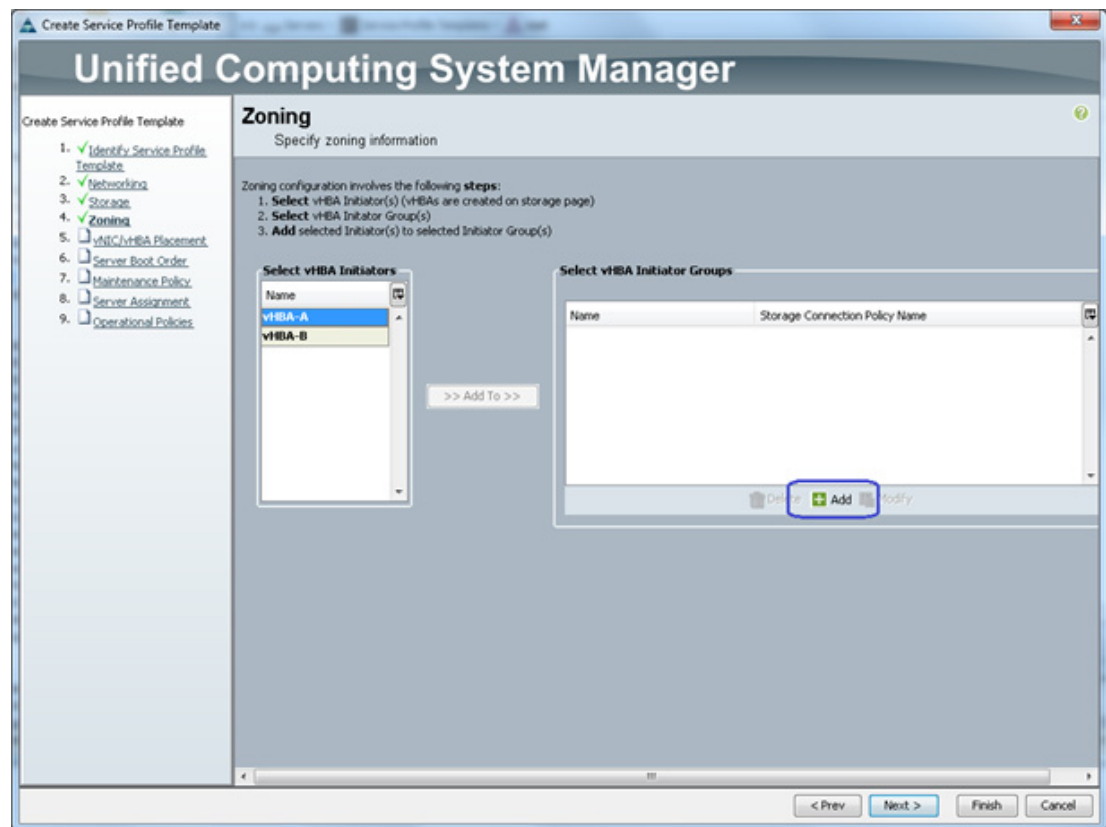
Adapter Performance Profile

Adapter Policy: + Create Fibre Channel Adapter Policy

QoS Policy: + Create QoS Policy

OK Cancel

11. Repeat step 11 for vHBA-B on fabric B, keeping the configuration the same.
12. For FC-Variant of the solution, from the "Zoning" page of the wizard, click "Add" as shown in the image below:



13. (FC-Variant only) Provide "SAN-A" a name to the vHBA initiator group and select the previously configured "Fabric-A" zoning policy from the drop-down menu and click "OK".

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**

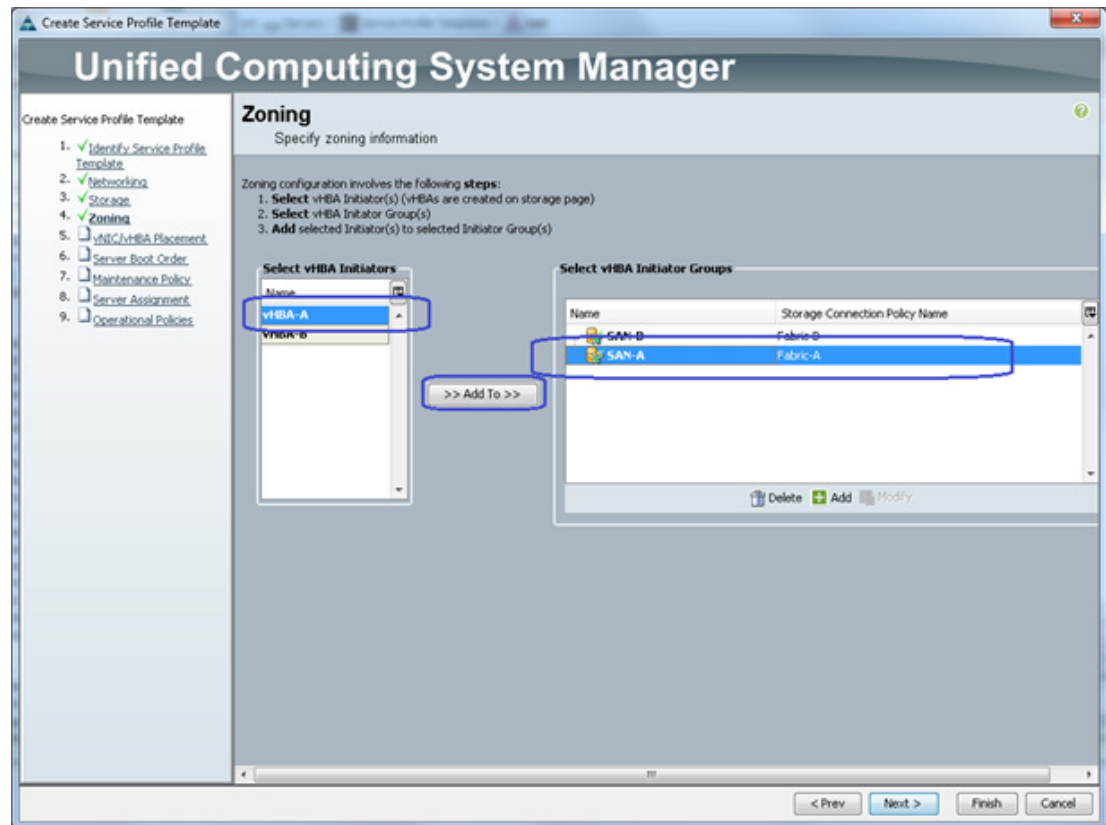
FC Target Endpoints

Filter Export Print

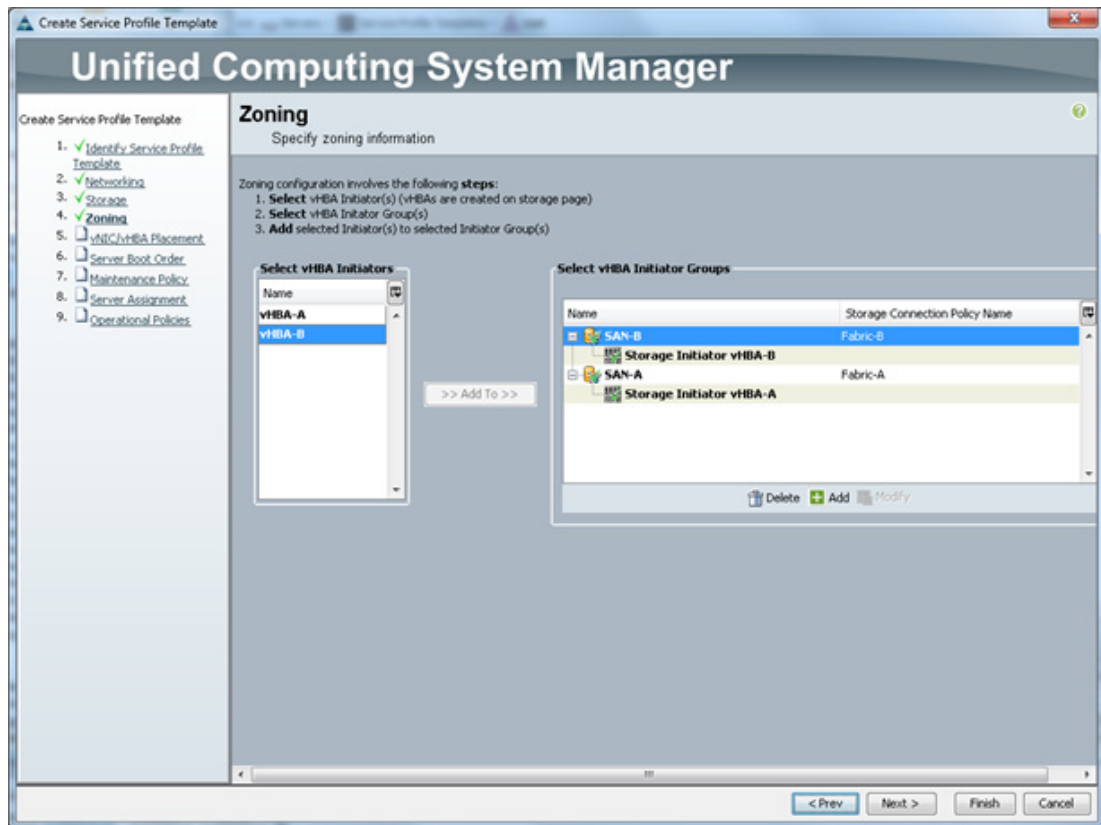
WWPN	Path	VSAN
50:06:01:64:3E:A0:65:0A	A	Storage
50:06:01:65:3E:A0:65:0A	A	Storage

OK Cancel

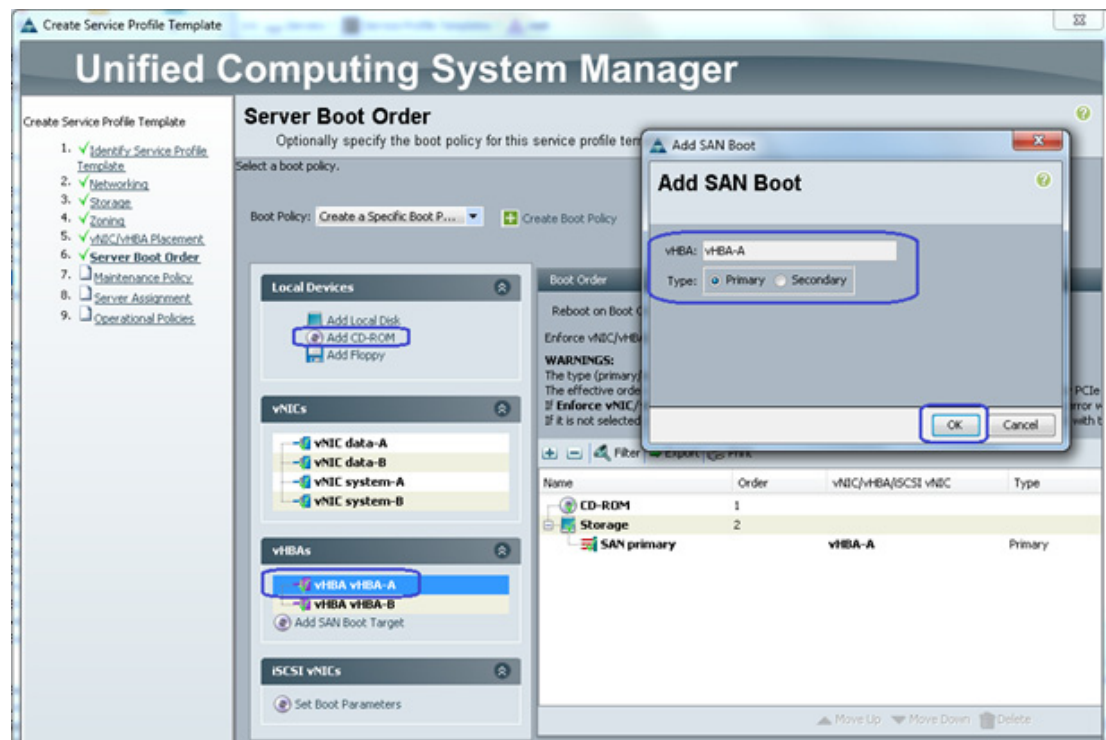
- (FC-Variant) Repeat steps 14 and 15 for the zoning on fabric B. Select "vHBA-A" from the list of initiators and "SAN-A" from the initiator-group, and click "Add To" as shown in the image below.



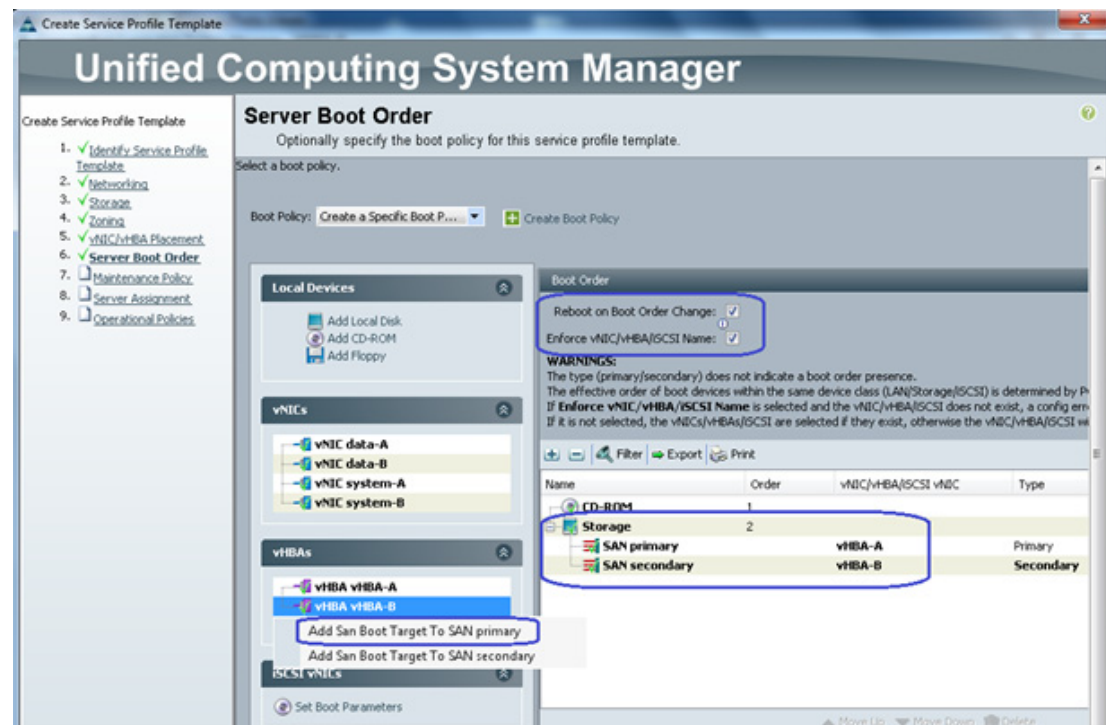
15. Repeat step 16 for fabric B. The result will look like the image below. Click "Next" and select the default configuration on the "vNIC/vHBA Placement" step.



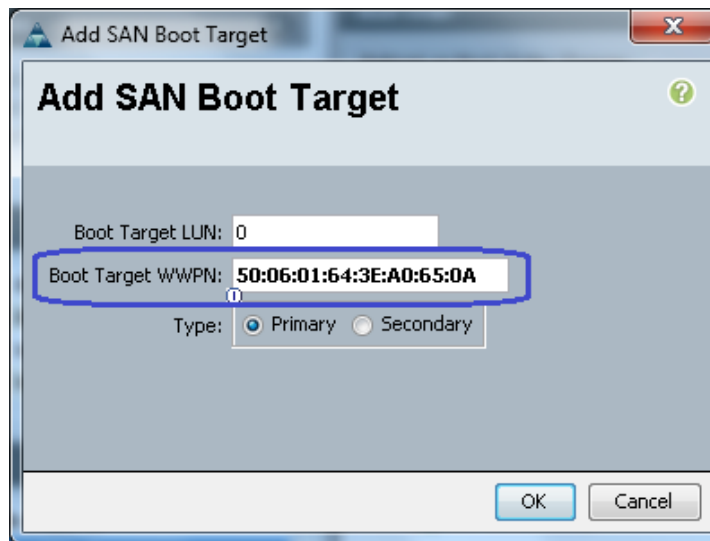
16. For the NFS-variant of the solution, select the default configuration on "Zoning" and "vNIC/vHBA Placement" steps by clicking "Next".
17. For both architectures, from the "Server Boot Order" step, select "Create a Specific Boot Policy" from the drop-down menu. Select "Add CD-ROM" as the first boot order choice. Click "vHBA-A" and provide the name "vHBA-A", keeping the type as primary.



18. Select vHBA-B as the next (secondary) choice to boot from SAN. When both vHBAs are added, make sure that "Reboot on Boot Order Change" and "Enforce vNIC/vHBA/iSCSI name" checkboxes are checked. Click "Add SAN Boot Target" under the vHBAs and click "Add San Boot Target to SAN primary" as shown in the image below:



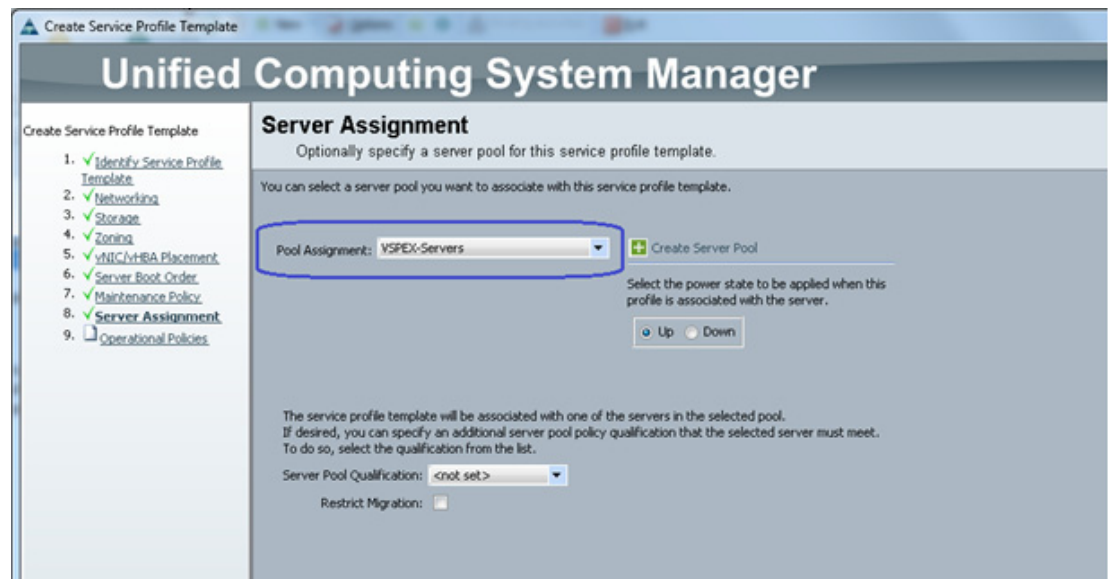
19. Provide a target WWPN for the VNX storage device (which can be obtained using "show flogi database" NXOS CLI command executed under "connect nxos {a|b}" shell as described in the previous subsection). Keep the target as primary.



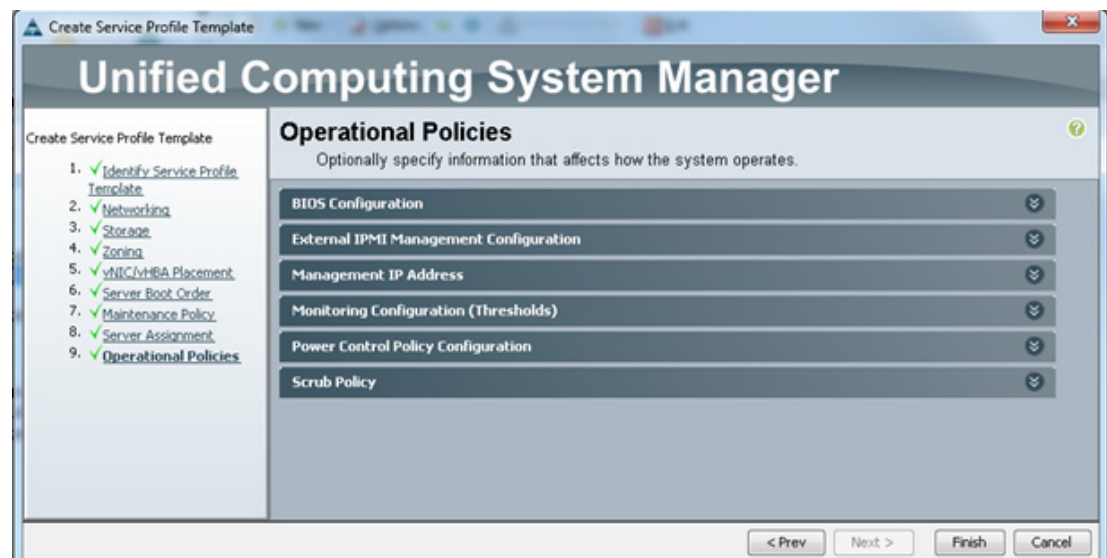
20. Repeat step 19 for the fabric B. The end result will look like the image below:



21. Click "Next". Retain the default and click "Next". Assign the server to "Pool Assignment" and select the Server Pool created in the previously as shown in the image below. Click "Next".



22. From the "Operation Policies" page, retain the defaults and click "Finish" to deploy the Service Profile Template.

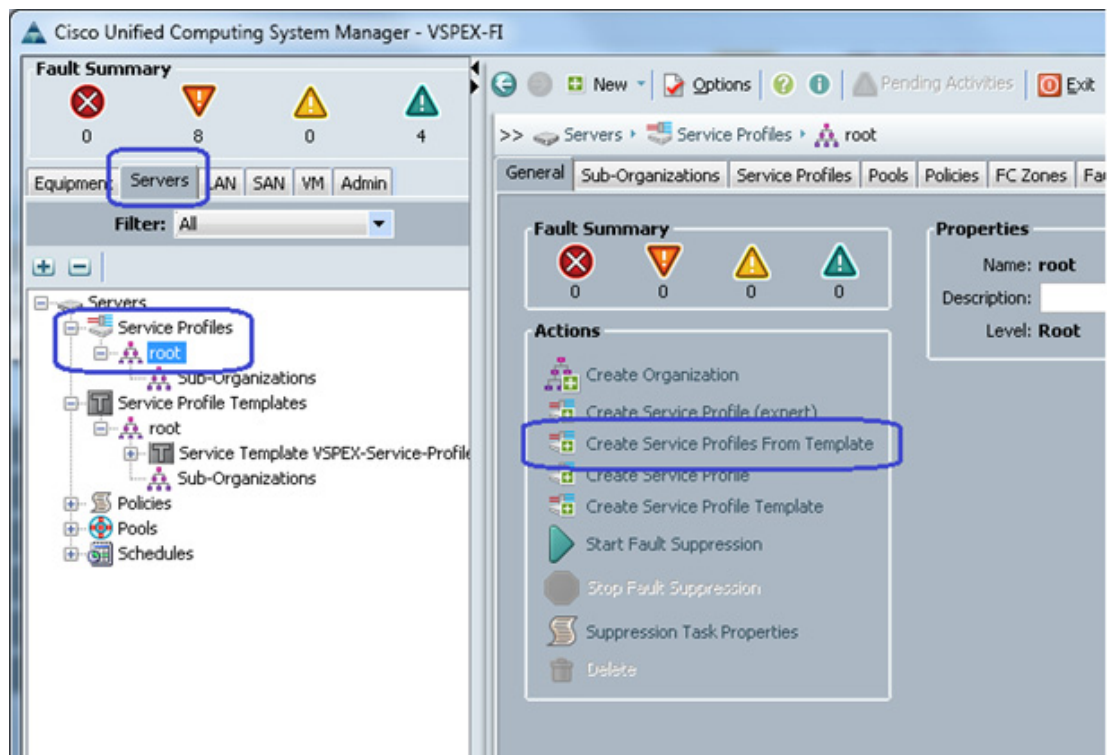


This concludes the service profile template creation.

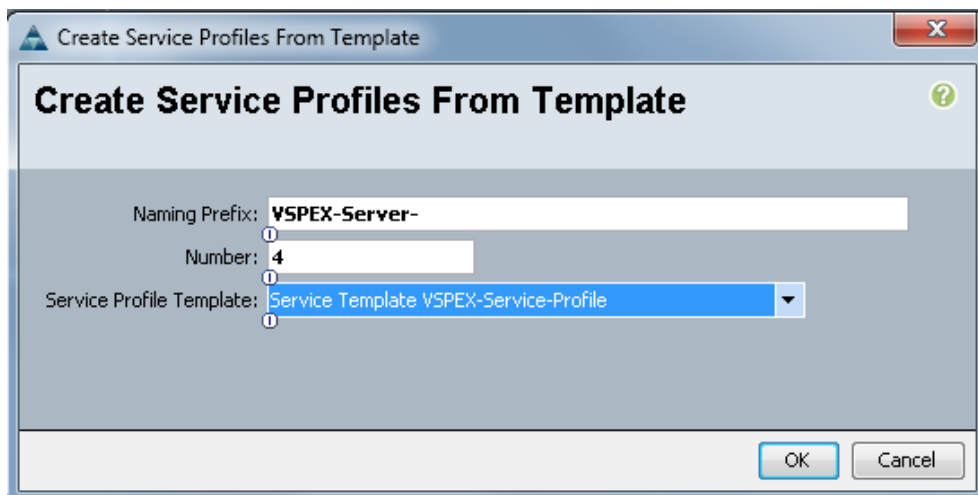
Instantiate Service Profiles from the Service Profile Template

As the final step to configure Cisco UCS Manager, instantiate the service profiles from the service profile template created previously. Complete the following steps:

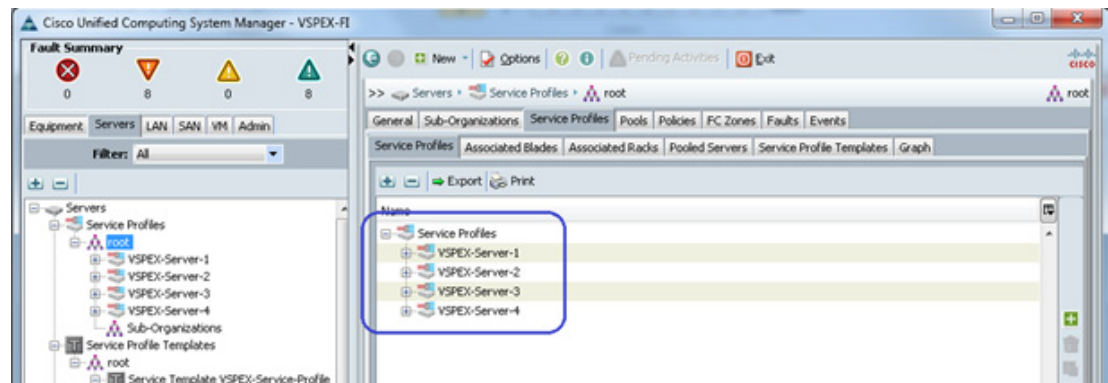
1. From the "Servers" tab, expand "Servers" > "Service profiles" > "root" and click "Create Service Profile from Template" link as shown in the image below:



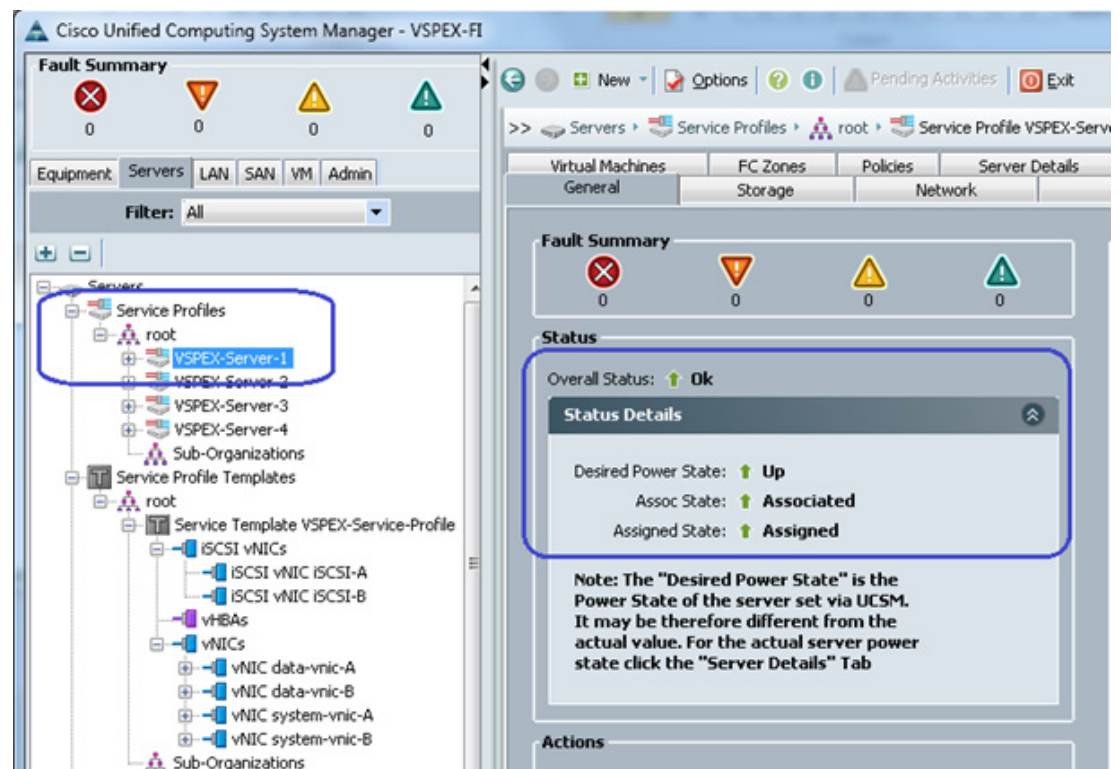
2. Provide a naming prefix, number of service profiles to be instantiated, and choose the service profile template from the drop-down menu. Refer to the sizing guidelines for the number of servers needed for your deployment.



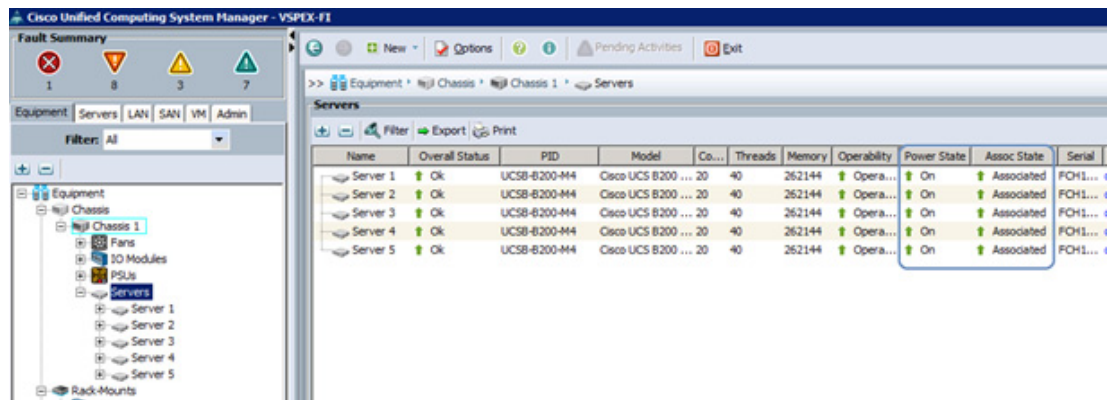
3. Four service profiles are created in this example as shown below:



4. When the service profile template is assigned to a server pool, the service profiles instantiated from the template will be assigned to an individual server resource from the server pool as available. Click a service profile to view its association state and with which server it is associated.



5. All five servers will be associated; view the summary by clicking "Servers" under the "Equipment" tab as shown below:



Configure Data Stores for ESXi Images

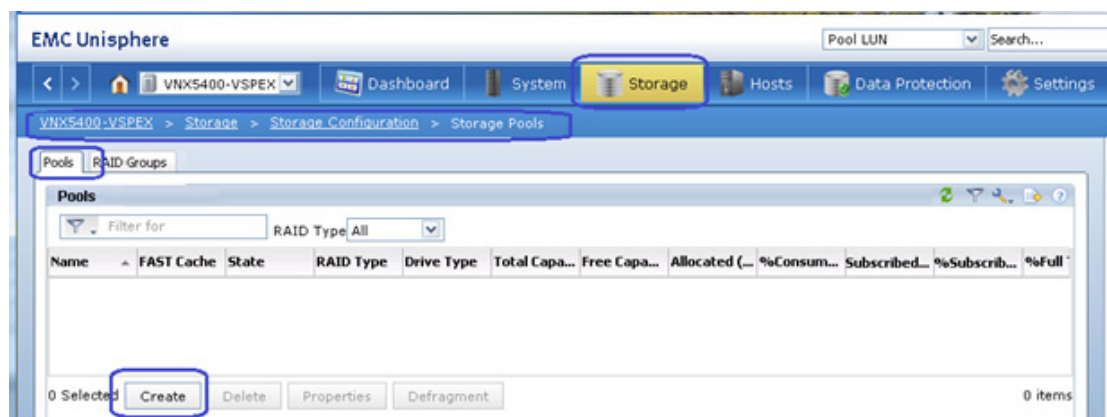
This section details the steps to create FC accessible datastores for the ESXi boot image on a per server basis. This includes the following steps:

- [Configure a Storage Pool](#)
- [Register Hosts](#)
- [Configure Storage Groups](#)

Configure a Storage Pool

To create a storage pool and configure boot LUNs on a per server basis, complete the following steps:

1. Connect to the EMC VNX Unisphere GUI, click the "Storage" tab. Select "Storage Configuration" > "Storage Pools". Click the "Pools" tab and click "Create" as shown below:



2. Choose "RAID5 (4 + 1)" for "Performance". Click "Manual" and click "Select" to manually select 5 SAS disks to create the storage pool.

VNX5400-VSPEX - Create Storage Pool

General Advanced

Storage Pool Parameters

Storage Pool Type: ☒ Pool ☐ RAID Group

☒ Scheduled Auto-Tiering

Storage Pool ID: 3

Storage Pool Name: Pool 3

Extreme Performance

RAID Configuration: RAID5 (4+1) Number of Flash Disks: 0

Performance

RAID Configuration: RAID5 (4+1) Number of SAS Disks: 5 (Recommended)

Distribution

Extreme Performance : 366.906 GB (12.03%)
Performance : 2684.038 GB (87.97%)

Disks

☐ Automatic ☐ Use Power Saving Eligible Disks

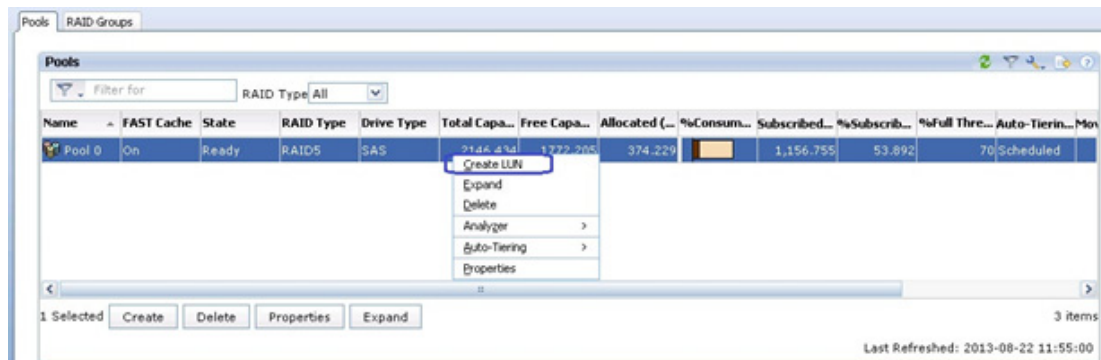
☒ Manual Total Raw Capacity: 3050.944...

Disk	Capacity	Drive Type	Model	State
Bus 0 Enclosure 7 Disk 10	91.727 GB	SATA Flash	SS160510 CL...	Unbound
Bus 0 Enclosure 7 Disk 9	91.727 GB	SATA Flash	SS160510 CL...	Unbound
Bus 0 Enclosure 7 Disk 8	91.727 GB	SATA Flash	SS160510 CL...	Unbound
Bus 0 Enclosure 7 Disk 7	91.727 GB	SATA Flash	SS160510 CL...	Unbound
Bus 1 Enclosure 1 Disk 6	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 1 Enclosure 1 Disk 5	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 1 Enclosure 1 Disk 4	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 1 Enclosure 1 Disk 3	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 1 Enclosure 1 Disk 2	536.808 GB	SAS	STE60005 CL...	Unbound

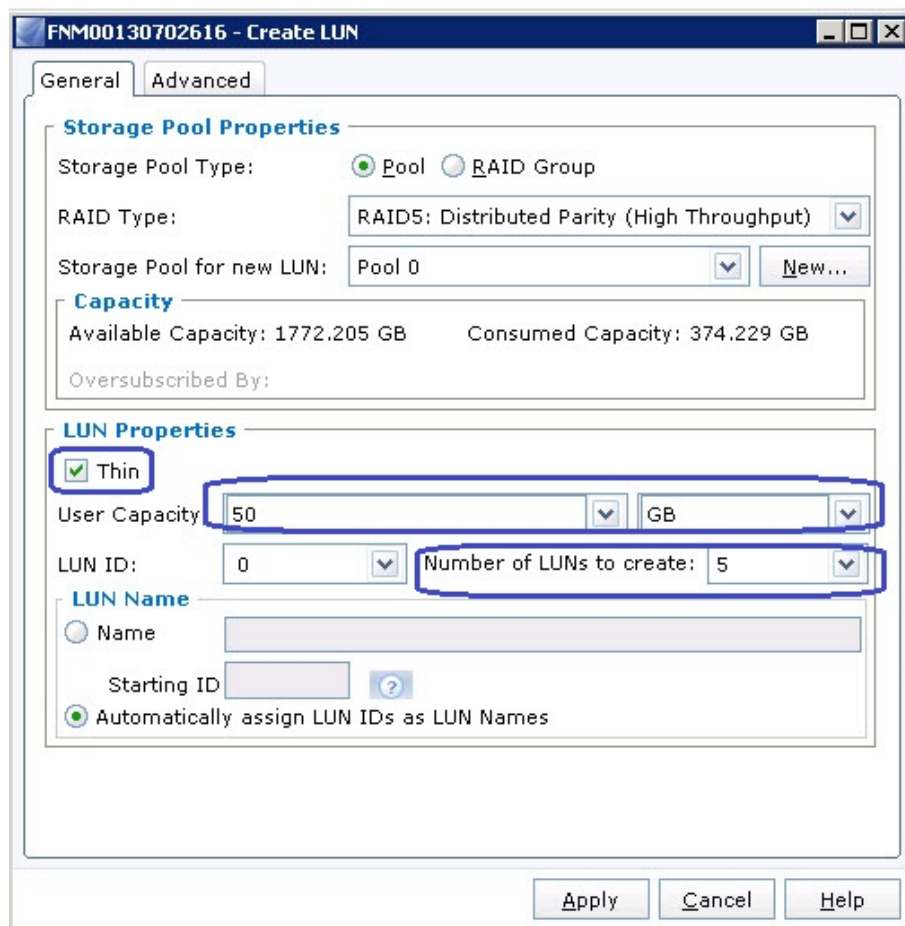
☒ Perform a background verify on the new storage

OK Apply Cancel Help

- From the newly created RAID group, right-click and click "Create LUN" as shown below:



4. Create 5 LUNs for five ESXi hosts, with 50 GB capacity each as shown below. Make sure that the "Thin" provisioning checkbox is checked.



Register Hosts

When the service profiles are associated in Cisco UCS Manager, the vHBAs will perform 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, complete the following steps:

1. (NFS-variant only) For the NFS-variant of the solution, the storage connectivity is thru MDS 9148S switches. The FC zoning must be configured manually on MDS 9148S switches. The FC-variant architecture, where storage is attached to FIs, FC zoning is taken care of by Cisco UCS Manager implicitly. The following steps detail how to configure zoning on MDS 9148S switches.
2. Log in to the MDS 9148S switch A and configure a zoneset for SAN fabric A. 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 Cisco UCS Manager as shown in step 7. The entire zoneset configuration will look like the image below. Activate the zoneset in the storage VSAN after it is configured.

```

VSPEX-MDS-FAB-A# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-MDS-FAB-A(config)# zoneset name VSPEX-Server-Fabric-A vsan 10
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer1-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:0e
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer2-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:1d
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer3-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:1c
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer4-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:1b
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer5-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:1a
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer6-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:09
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer7-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:08
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer8-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:07
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer9-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:06
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# zone name VSPEX-ESXiServer10-fc0
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:05
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:64:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# member pwn 50:06:01:6c:3e:a0:52:02
VSPEX-MDS-FAB-A(config-zoneset-zone)# exit
VSPEX-MDS-FAB-A(config-zoneset)# exit
VSPEX-MDS-FAB-A(config)# zoneset activate name VSPEX-Server-Fabric-A vsan 10
Zoneset activation initiated. check zone status
VSPEX-MDS-FAB-A(config)# copy running-config startup-config
[#####] 100%
Copy complete.
VSPEX-MDS-FAB-A(config)#

```

3. (NFS-variant only) Validate the successful activation of the zoneset by issuing "show zoneset brief" command as shown below:

```
VSPEX-MDS-FAB-A(config)# show zoneset brief
zoneset name VSPEX-Server-Fabric-A vsan 10
  zone VSPEX-ESXiServer1-fc0
  zone VSPEX-ESXiServer2-fc0
  zone VSPEX-ESXiServer3-fc0
  zone VSPEX-ESXiServer4-fc0
  zone VSPEX-ESXiServer5-fc0
  zone VSPEX-ESXiServer6-fc0
  zone VSPEX-ESXiServer7-fc0
  zone VSPEX-ESXiServer8-fc0
  zone VSPEX-ESXiServer9-fc0
  zone VSPEX-ESXiServer10-fc0
VSPEX-MDS-FAB-A(config) #
```

4. (NFS-variant only) On the MDS 9148S switch B, create a zoneset for fabric B as shown below:

```

VSPEX-MDS-FAB-B# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
VSPEX-MDS-FAB-B(config)# zoneset name VSPEX-Server-Fabric-B vsan 10
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer1-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer2-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:0d
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer3-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:0c
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer4-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:0b
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer5-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:0a
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer6-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:19
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer7-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:18
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer8-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:17
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer9-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:16
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# zone name VSPEX-ESXiServer10-fc1
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 20:00:00:25:b5:66:dd:15
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:65:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# member pwn 50:06:01:6d:3e:a0:52:02
VSPEX-MDS-FAB-B(config-zoneset-zone)# exit
VSPEX-MDS-FAB-B(config-zoneset)# exit
VSPEX-MDS-FAB-B(config)# zoneset activate name VSPEX-Server-Fabric-B vsan 10
Zoneset activation initiated. check zone status
VSPEX-MDS-FAB-B(config)# copy running-config startup-config
[#####] 100%
Copy complete.
VSPEX-MDS-FAB-B(config)#

```

The Zoneset on fabric B will look like the image below:

```

VSPEX-MDS-FAB-B(config)# show zoneset brief
zoneset name VSPEX-Server-Fabric-B vsan 10
  zone VSPEX-ESXiServer1-fc1
  zone VSPEX-ESXiServer2-fc1
  zone VSPEX-ESXiServer3-fc1
  zone VSPEX-ESXiServer4-fc1
  zone VSPEX-ESXiServer5-fc1
  zone VSPEX-ESXiServer6-fc1
  zone VSPEX-ESXiServer7-fc1
  zone VSPEX-ESXiServer8-fc1
  zone VSPEX-ESXiServer9-fc1
  zone VSPEX-ESXiServer10-fc1
VSPEX-MDS-FAB-B(config)#

```

5. (NFS-variant only) To validate the zoneset configuration across the entire SAN fabric, SSH to UCS FI-A, issue "connect nxos" command, and run "show npv flogi-table". It will list all ten FLogI sessions; one from each vHBA on fabric A in storage VSAN as shown below.

```

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

```

SERVER	INTERFACE	VSAN	FCID	PORT NAME	NODE NAME	EXTERNAL INTERFACE
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.

```

6. (NFS-variant only) The "show flogi database" command on MDS 9148S switch will show 14 FLogI sessions: 10 from B200 M4 vHBAs, 2 from FI-A's FC ports, and 2 from VNX storage array's SP-A and SP-B FC ports as shown below. Verify the FLogI entries on SAN fabric B.


```

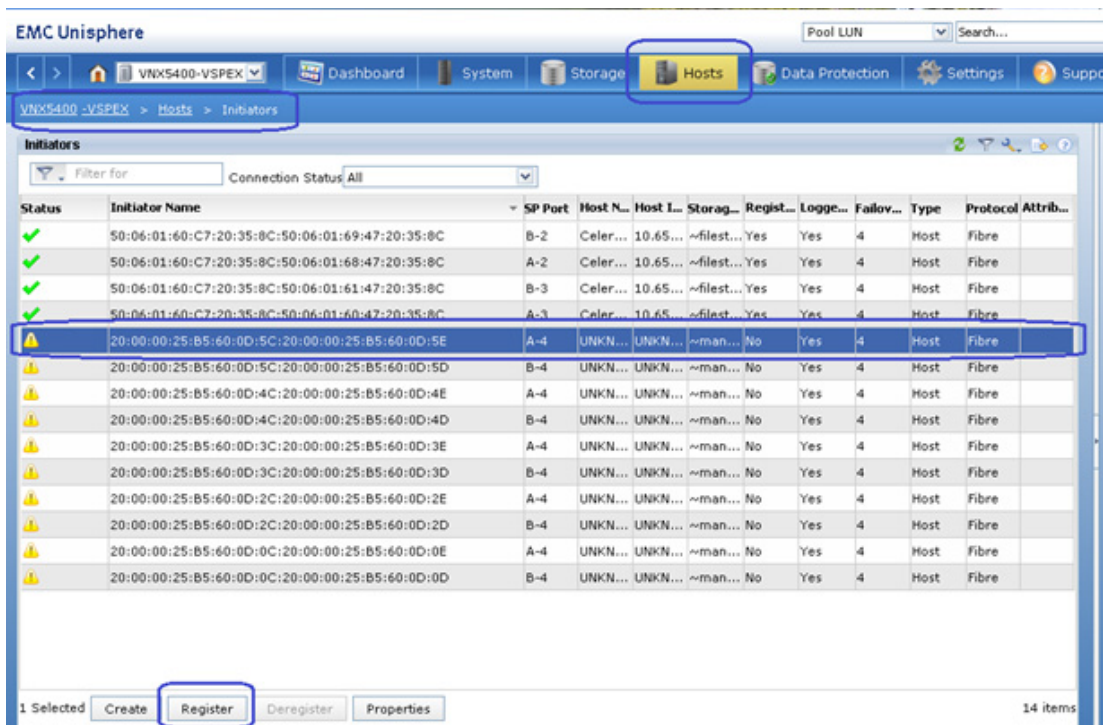
VSPEX-MDS-FAB-A# show flogi database
-----
INTERFACE      VSAN    FCID      PORT NAME      NODE NAME
-----
fc1/29          10      0x5c0000   20:41:00:0d:ec:f7:04:00  20:0a:00:0d:ec:f7:04:01
fc1/29          10      0x5c0002   20:00:00:25:b5:66:dd:0e  20:00:00:25:b5:60:0d:0e
fc1/29          10      0x5c0004   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.

VSPEX-MDS-FAB-A#

```

- On Unisphere GUI, click the "Hosts" tab and click "Initiators". Select the first unregistered initiator and click "Register".



- From Cisco UCS Manager GUI, click the "Servers" tab, expand "Servers" > "Service Profiles" > "root" > <a specific service-profile>, and click "vHBAs". This will list the WWPN identifies of the list. Using these IDs, associate WWPN to the server.
- From the "Register Initiator Record" wizard, select the Initiator Type as "CLARiiON/VNX" and Failover Mode as failovermode 4 from the drop-down menus. Click "New Host", provide the hostname and (future) management IP address of the host. Click "OK".

Register Initiator Record

Initiator Information

WWN/IQN: 20:00:00:25:85:60:0D:4C:20:00:00:25:85:60:0D:4E

SP - port: A-4 (Fibre)

Initiator Type: CLARiiON/VNX Failover Mode: /e-Active mode(ALUA)-failovermode 4

Host Agent Information

☒ New Host ☐ Existing Host ☐ Selected Host

Host Name: VSPEX-Server-1

IP Address: 10.65.121.239

[Advanced Options](#)

OK Cancel Help

10. Select the second vHBA's WWPN from the same server, click "Register" and click "Existing Host". Click "Browse Host" to select the host:

Register Initiator Record

Initiator Information

WWN/IQN: 20:00:00:25:85:60:0D:4C:20:00:00:25:85:60:0D:4D

SP - port: B-4 (Fibre)

Initiator Type: CLARiiON/VNX Failover Mode: /e-Active mode(ALUA)-failovermode 4

Host Agent Information

☐ New Host ☒ Existing Host ☐ Selected Host


Host Name:

IP Address:

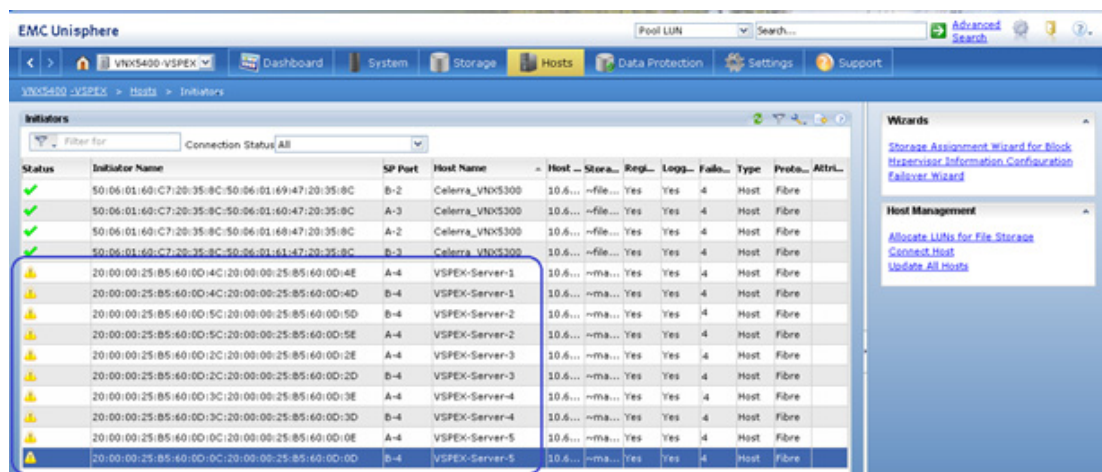
[Advanced Options](#)

OK Cancel Help

11. Select the previously registered host and click OK.




12. Repeat these steps for all the servers in the group. The result will look like the image below:




Configure Storage Groups

When the hosts as well as the LUNs are created on the VNX storage array, storage groups need to be created to assign access to LUNs for various hosts. Boot LUN will be dedicated to a specific server. To configure storage groups, complete the following steps:

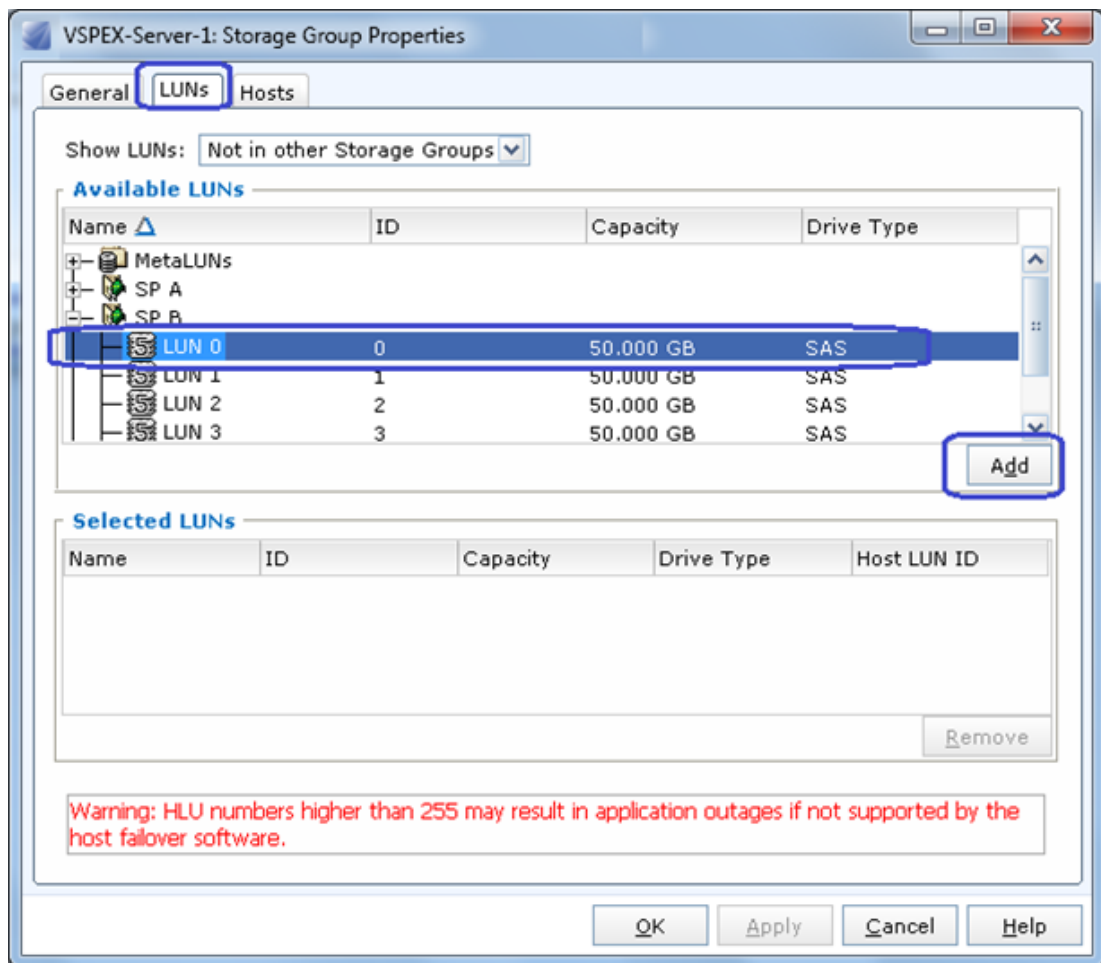
1. Click the "Hosts" tab on the EMC VNX Unisphere GUI and click "Storage Groups". Click "Create" to create a new storage group.



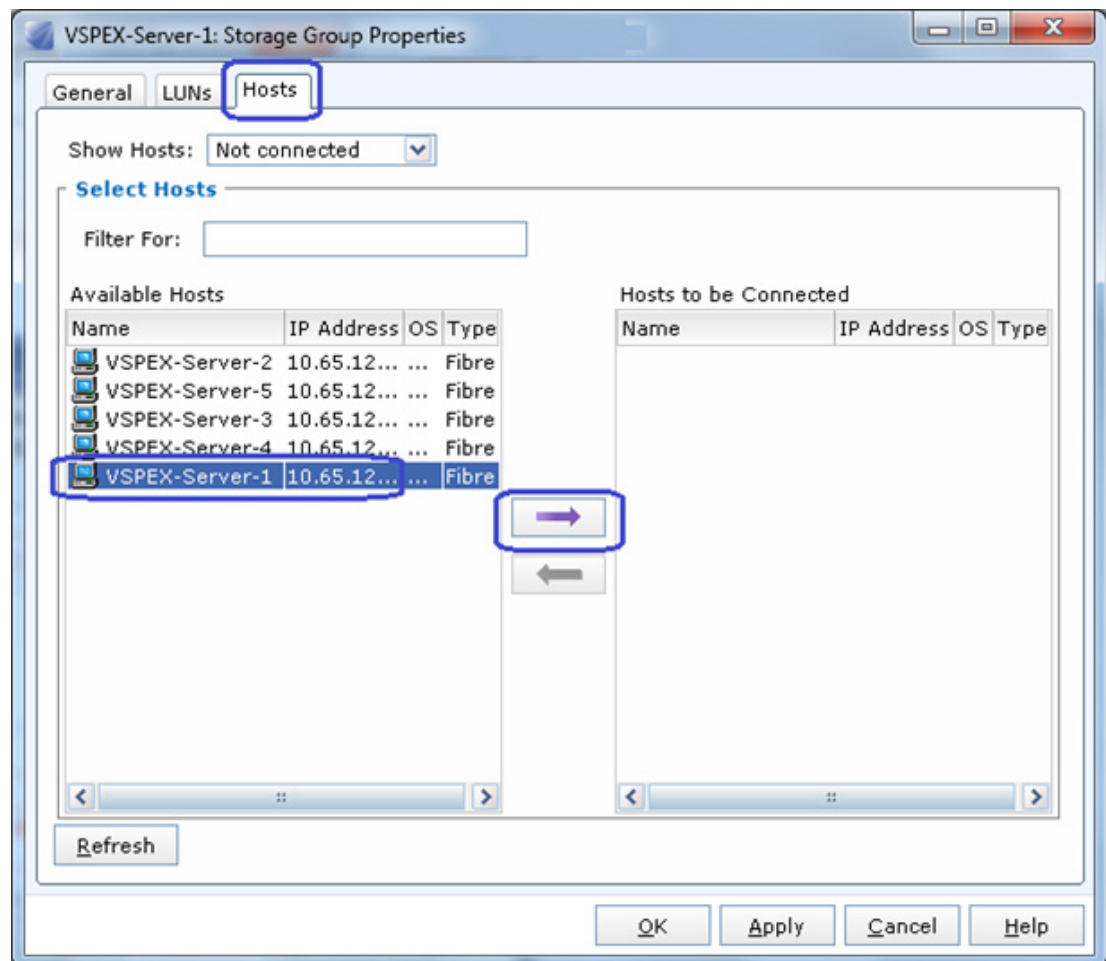
2. Provide a name for the storage group.



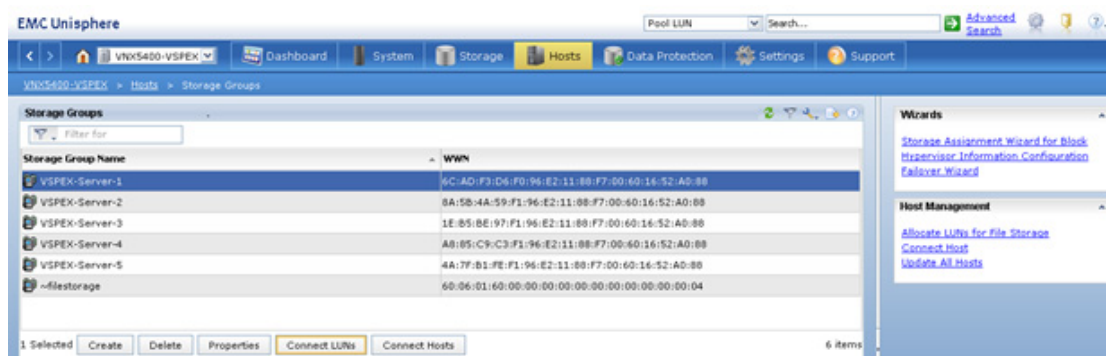
3. A success message displays. The system prompts to create LUNs and connect hosts. Click "Yes".
4. From the "LUNs" tab, select a single LUN and click "Add".



5. Click the "Hosts" tab and select a single server to add to the storage group. Click "OK" to deploy the storage group.



6. Repeat these steps for all five servers. The end result will look like the image below:

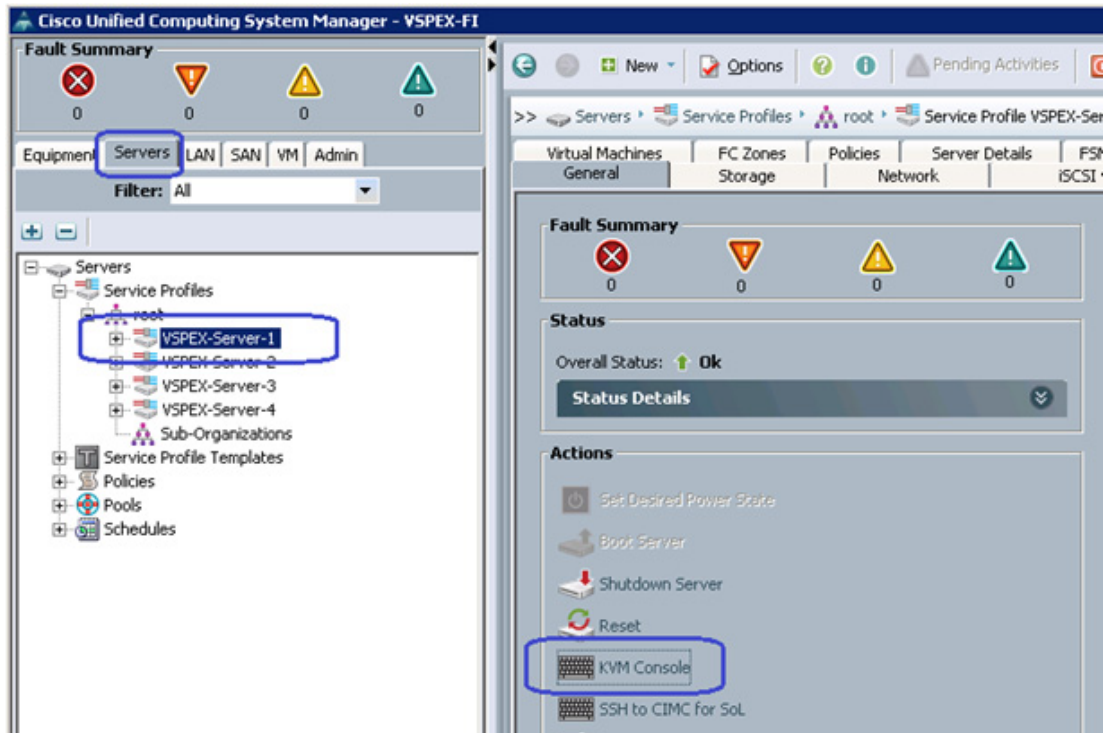


The end-to-end FC storage access from servers in UCS is now available to the specific boot LUN on the VNX storage devices. The next step is to install the ESXi images on the server.


Install ESXi Servers and vCenter Infrastructure

To install the ESXi image on the Cisco UCS servers, complete the following steps:

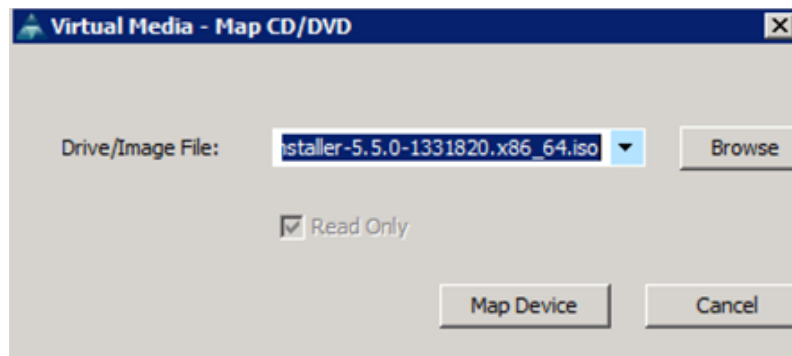
1. From Cisco UCS Manager GUI, select the "Servers" tab, expand "Servers" > "Service Profiles" > "root", and select a particular service profile. Click the "KVM Console" link.



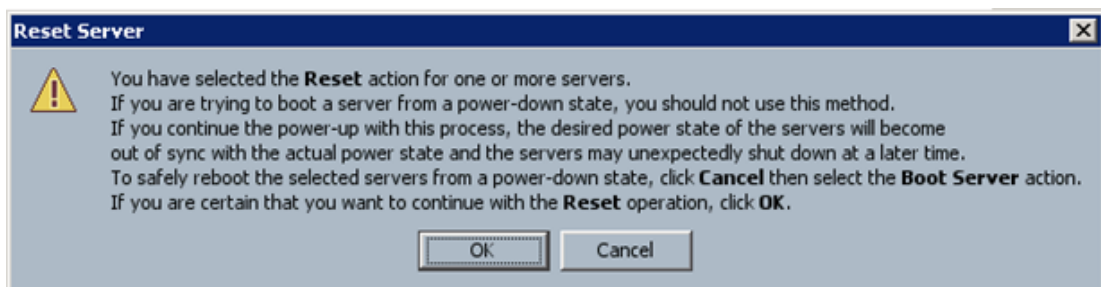
2. When the Java applet of KVM is launched, click the "Virtual Media" tab and click the "Activate Virtual Devices" tab as shown in the image below. Highlight the "Map CD/DVD" to map the ESXi5.5 .iso image. Navigate the local directory structure and select the ISO image of the ESXi 5.5 hypervisor installer media.



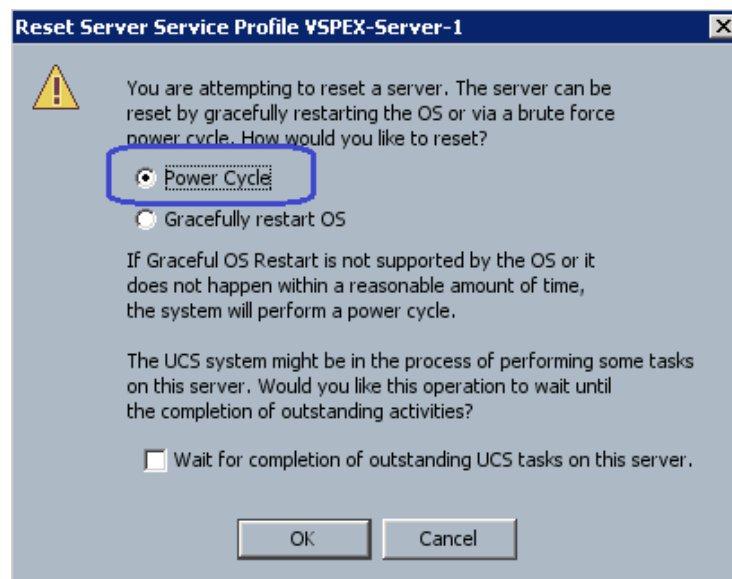
3. When the ISO image appears in the list, click the "Map Device" tab and click "Reset" to reboot the server.



4. Click "OK".



5. Click "Power Cycle" and click "OK"

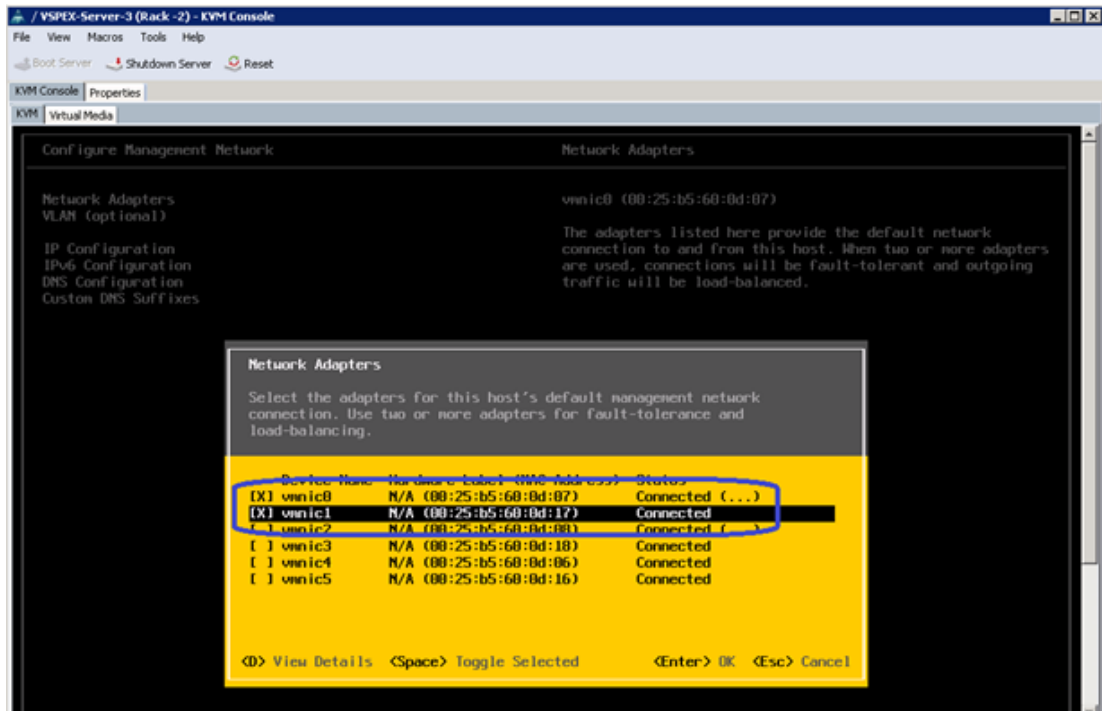


6. Click the "KVM" tab to view the ESXi boot media booted from the virtual CD-ROM drive.
The ESXi installation media will boot from the virtual disk mounted on the KVM. Complete the following steps to install ESXi 5.5 hypervisor on the boot LUN.


Note

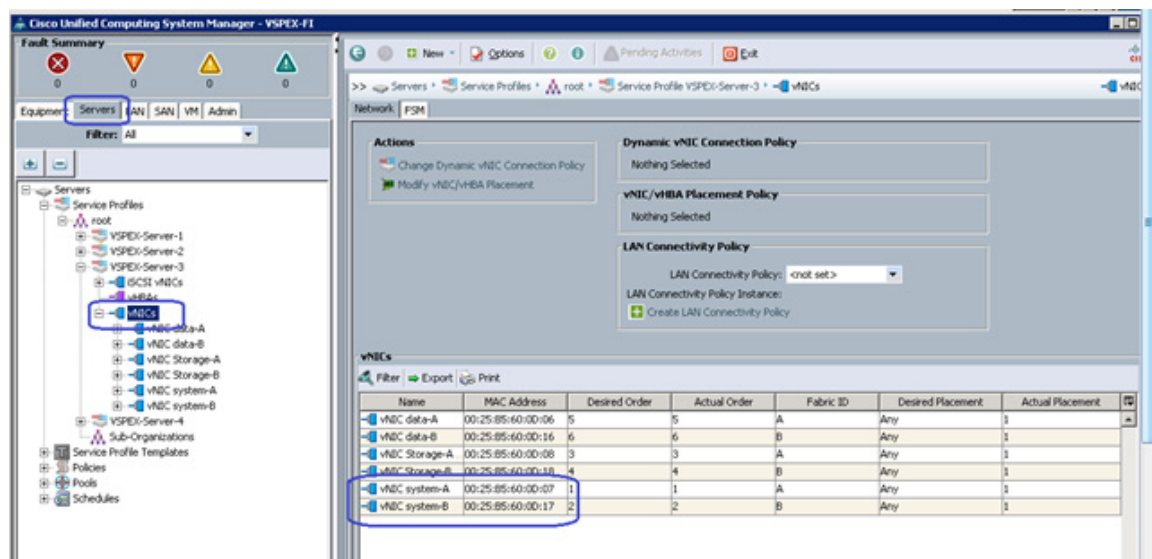
Make sure that you select the boot LUN and not the local disk to install the hypervisor image. You can select all the default parameters or per your requirements.

When the ESXi is installed, log into the system by pressing "F2" from the KVM window. Configure the basic management network for the ESXi host. Make sure to select two system VNICs as shown in the image below:



The easiest way to determine which vmnic adapter should be used for the vSphere management purpose is to identify the vmnic by MAC address. The MAC addresses of the VNICs (vmnic's) are summarized on the following Cisco UCS Manager GUI window.

- 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 as shown below:



- Repeat the ESXi installation steps for all four servers.

VMware vCenter Server Deployment

This section describes the installation of VMware vCenter for VMware environment and will accomplish the following:

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

For detailed information about installing a vCenter Server, go to:

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

To configure the vCenter server, complete the following high-level steps:

1. Create the vCenter host VM.

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

2. Install vCenter guest OS

Install the guest OS on the vCenter host virtual machine. VMware recommends using Windows Server 2012 R2. To make sure that adequate space is available on the vCenter and vSphere Update Manager installation drive, refer to the vSphere Installation and Setup Guide.

3. Install vCenter server

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

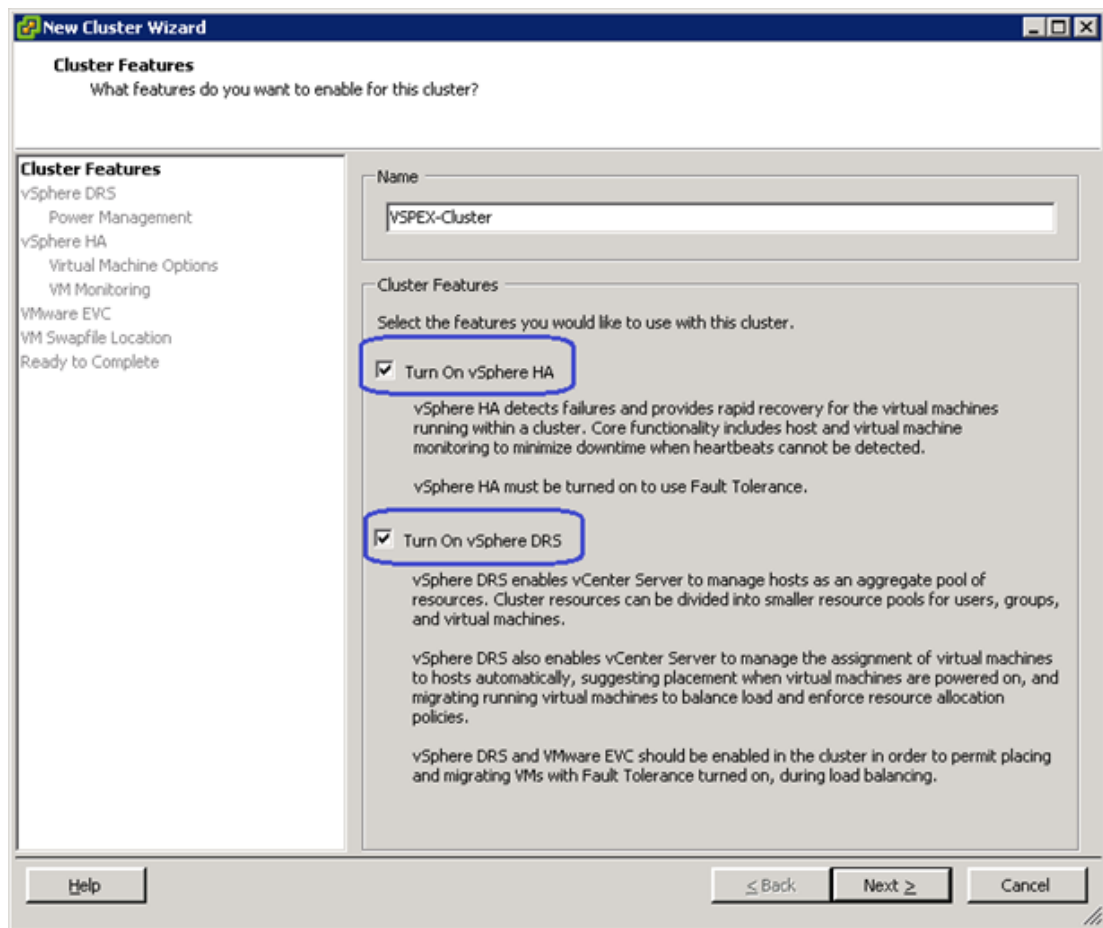
4. Apply vSphere license keys

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

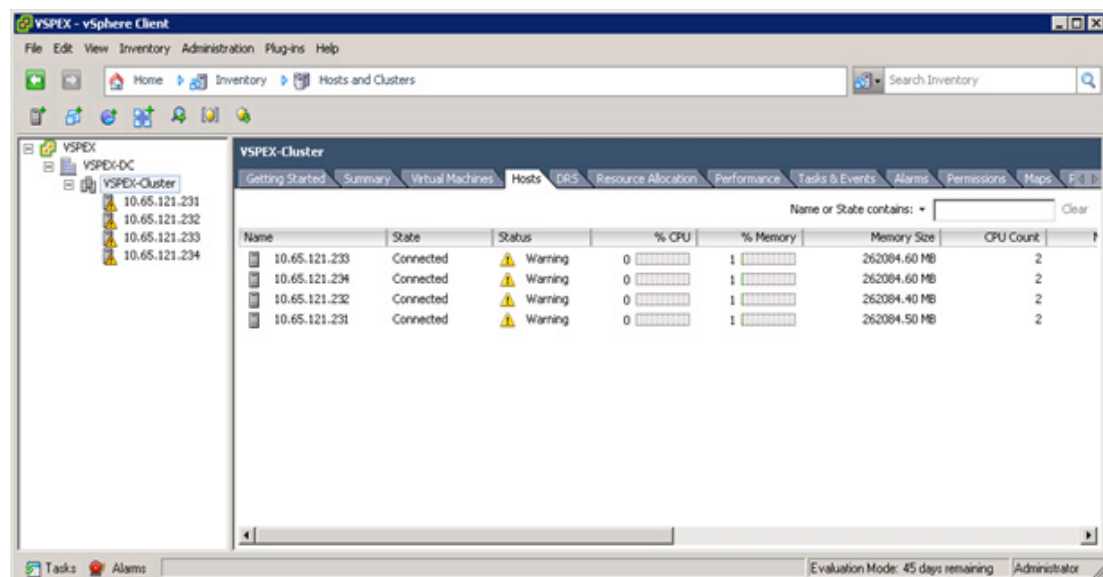
Configuring Cluster, HA and DRS on the vCenter

To add the VMware on a virtual machine vCenter, complete the following steps:

1. Log into VMware ESXi Host using VMware vSphere Client.
2. Create a vCenter Datacenter.
3. Create a new management cluster with DRS and HA enabled.
4. Right-click the cluster and in the corresponding context menu, click Edit Settings.
5. Select the checkboxes Turn On vSphere HA and Turn On vSphere DRS.
6. Click OK, to save changes.



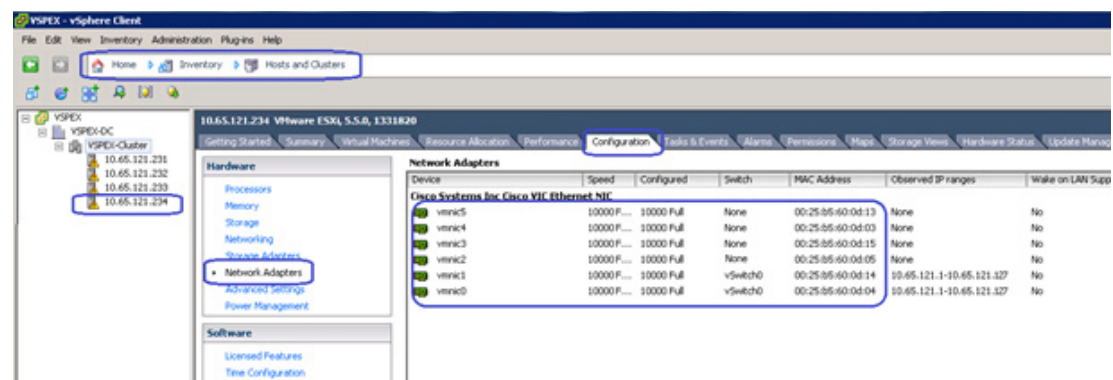
7. Add all ESXi hosts to the cluster by providing servers' management IP addresses and login credentials one by one.



8. Configure EMC Power Path for each hosts using VMware Update Manager (VUM). When all the hosts are added in vCenter, it is very easy to manage software installed on the hosts from a central location. Use the VMware Update Manager to install EMC power path for better FC SAN availability from each host. Failover Mode on the VNX storage array works hand in hand with host side power path for storage high availability.

Virtual Networking Configuration

In the Cisco UCS Manager service profile, created six VNICs per server for NFS-Variant and four VNICs per server for FC-Variant were created. This shows up as six or four network adaptors or vmnics in ESXi server. View these adapters in the vCenter by selecting "Home" > "Inventory" > "Hosts and Clusters" view, select a particular server, click the "Configuration" tab and click "Network Adapters" as shown in the image below.



Create a table like the one shown below:

Table 15 Service Profile VNIC and vSphere vmnic Relations

UCSM VNIC Name	vSphere NIC name	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
VM Data-A	vmnic4		data-uplink
VM Data-B	vmnic5		data-uplink

* Applicable for NFS-variant of the solution only.

Presented are 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

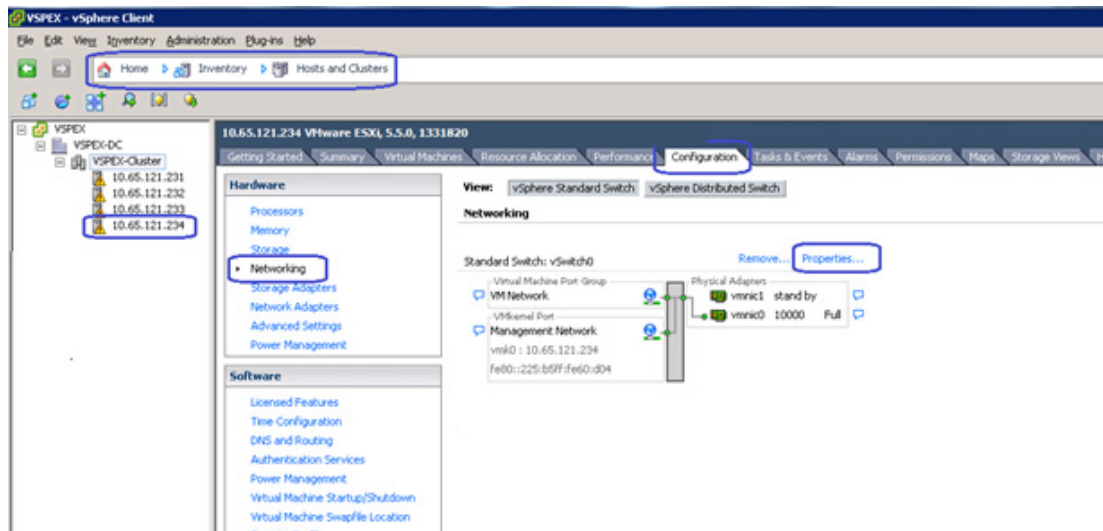
You can use either virtual switching strategy with any variant of architecture. This section focuses on the vSphere native virtual switching.

Create three native virtual switches for the virtual network configuration as follows:

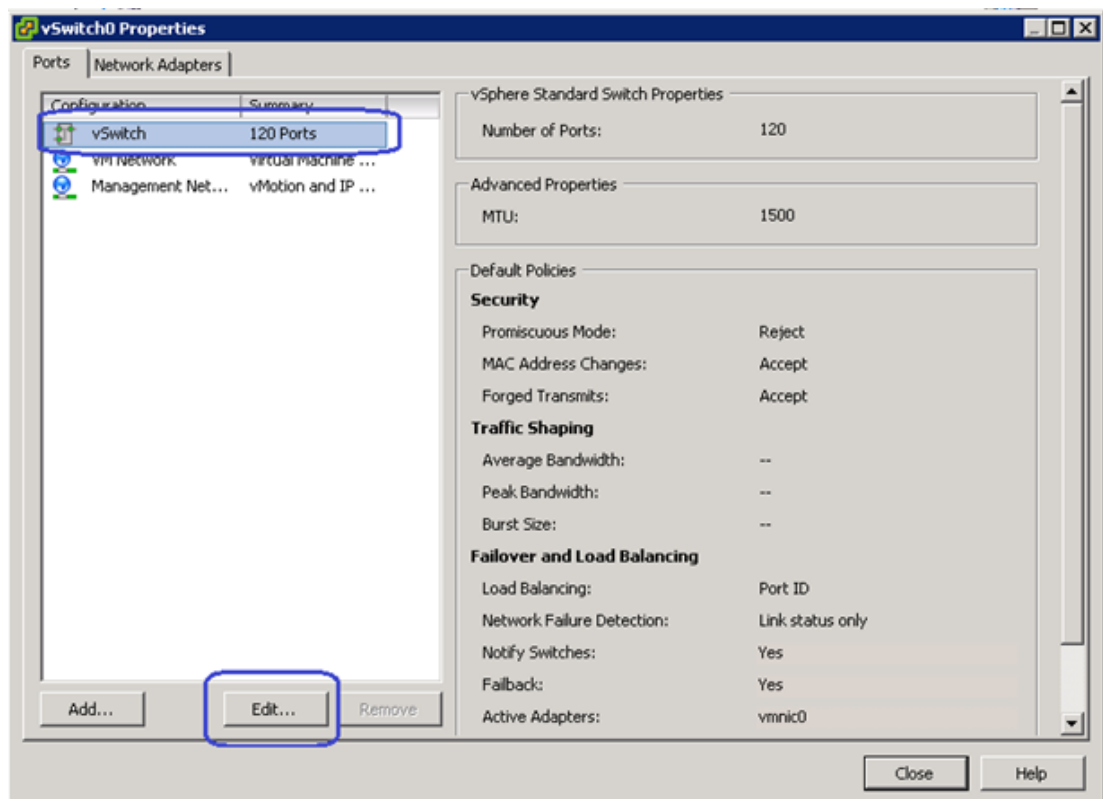
- vSwitch0 - Standard, default vSwitch for management and vMotion traffic
- vSwitch1 - For Storage traffic (NFS-variant)
- vSwitch2 - For VM data traffic

Each vSwitch listed above will have two vmnics, one on each fabric for load balancing and high-availability. For vMotion traffic, jumbo MTU needs to be configured in the virtual network. To configure the two vSwitches, complete the following steps:

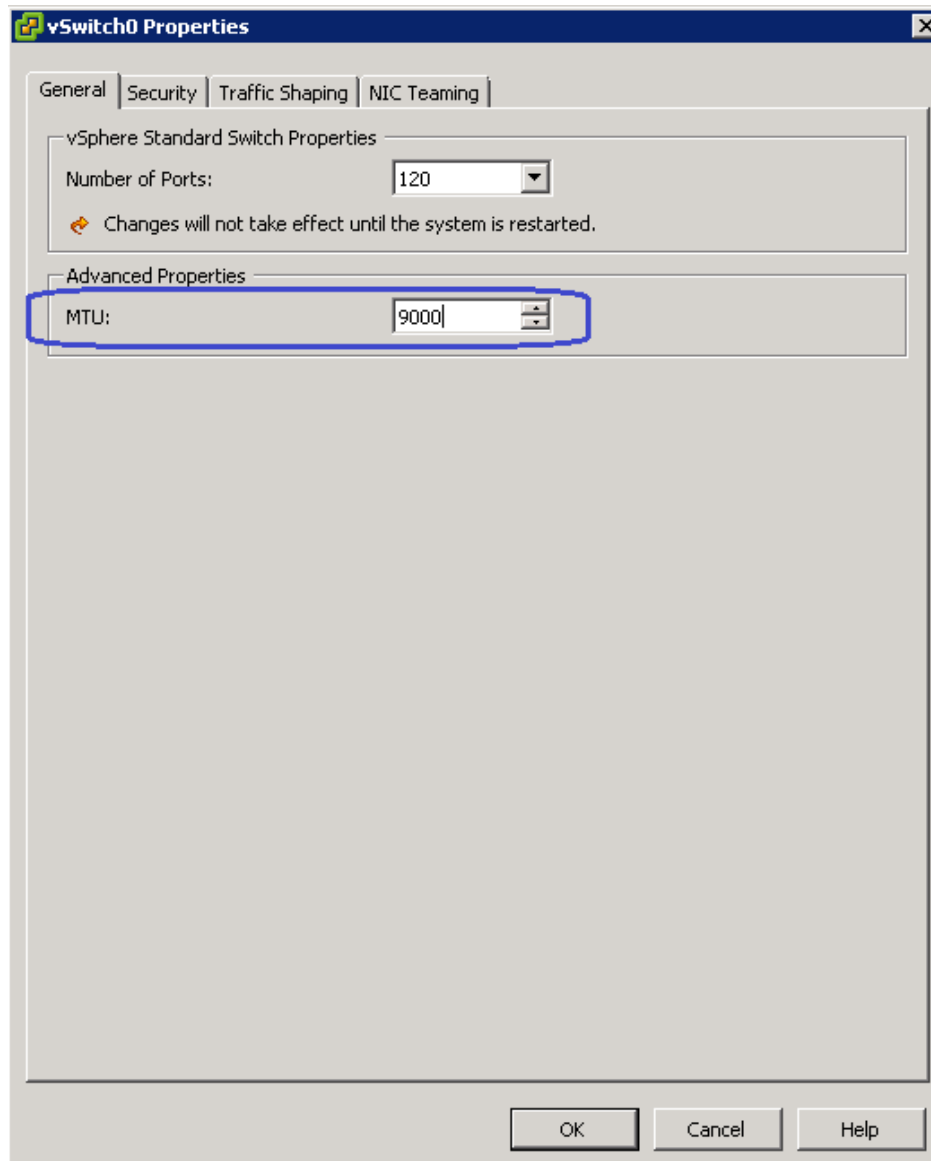
1. Select the "Home" > "Inventory" > "Hosts and Clusters" panel on vCenter, expand the VSPEX cluster and select an ESXi host. Click the "Configuration" tab, select "Networking" and click "Properties" of vSwitch0 as shown below:



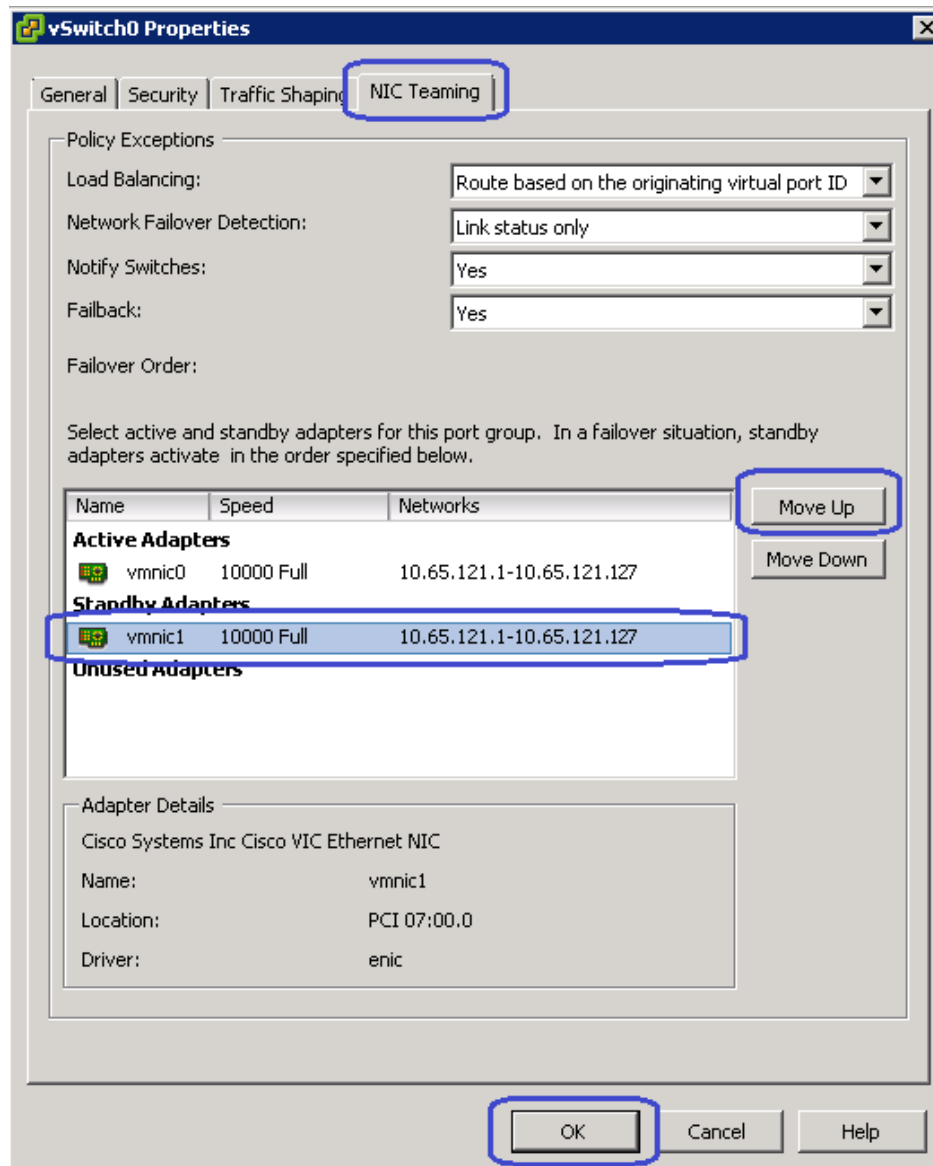
2. Select the "vSwitch" and click "Edit" as shown below:



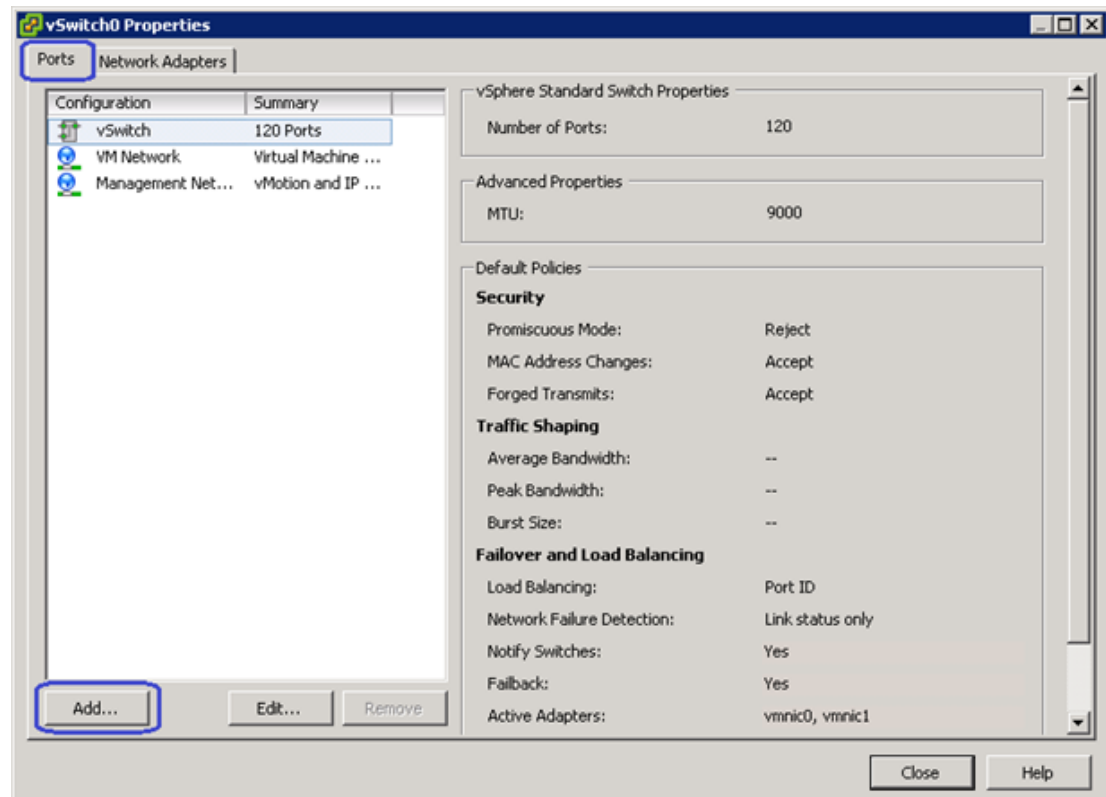
3. Change the MTU to '9000' in the "General" tab.



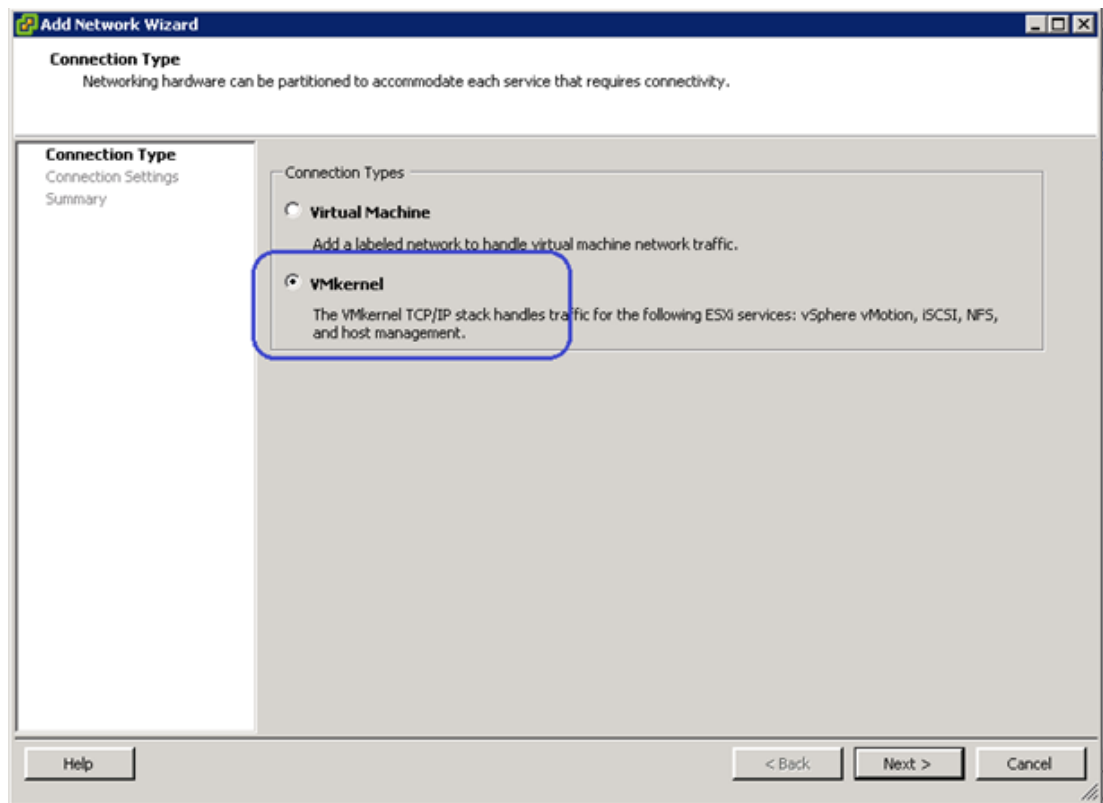
4. Click the "NIC Teaming" tab. Move up the "Standby Adapter" to the "Active Adapters" list, and click "OK".



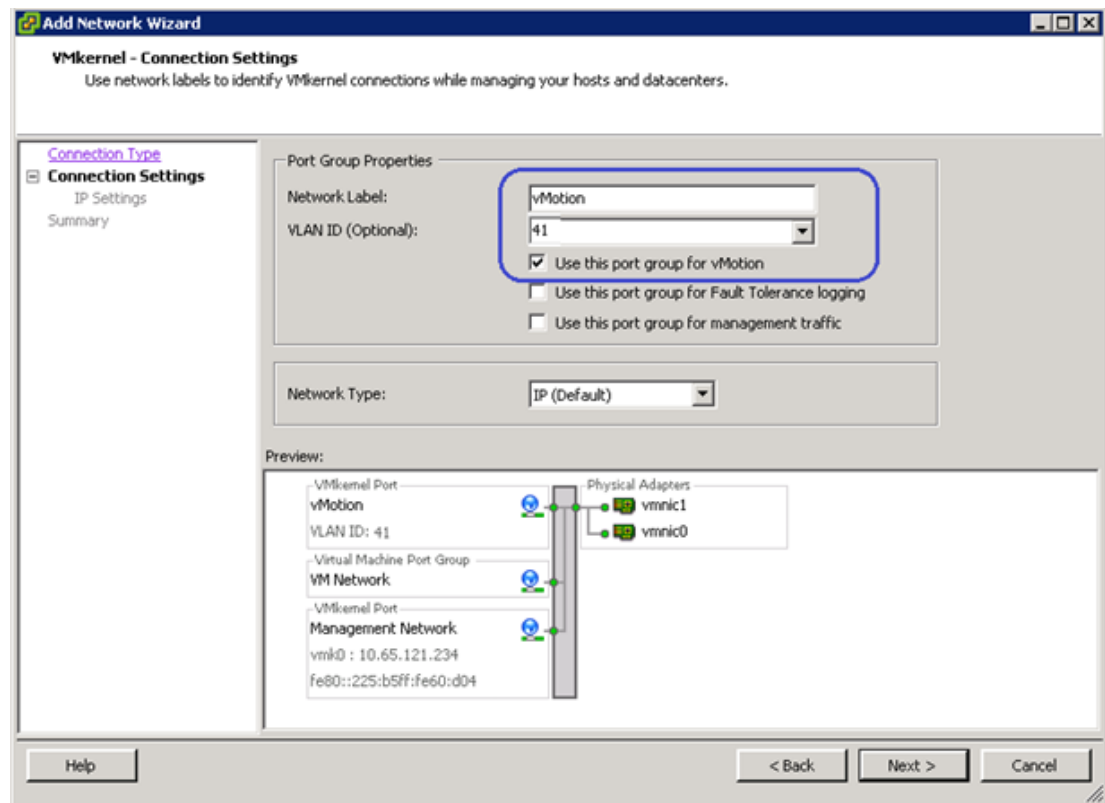
5. The vSwitch0 configuration window displays. Click the "Ports" tab and click "Add".



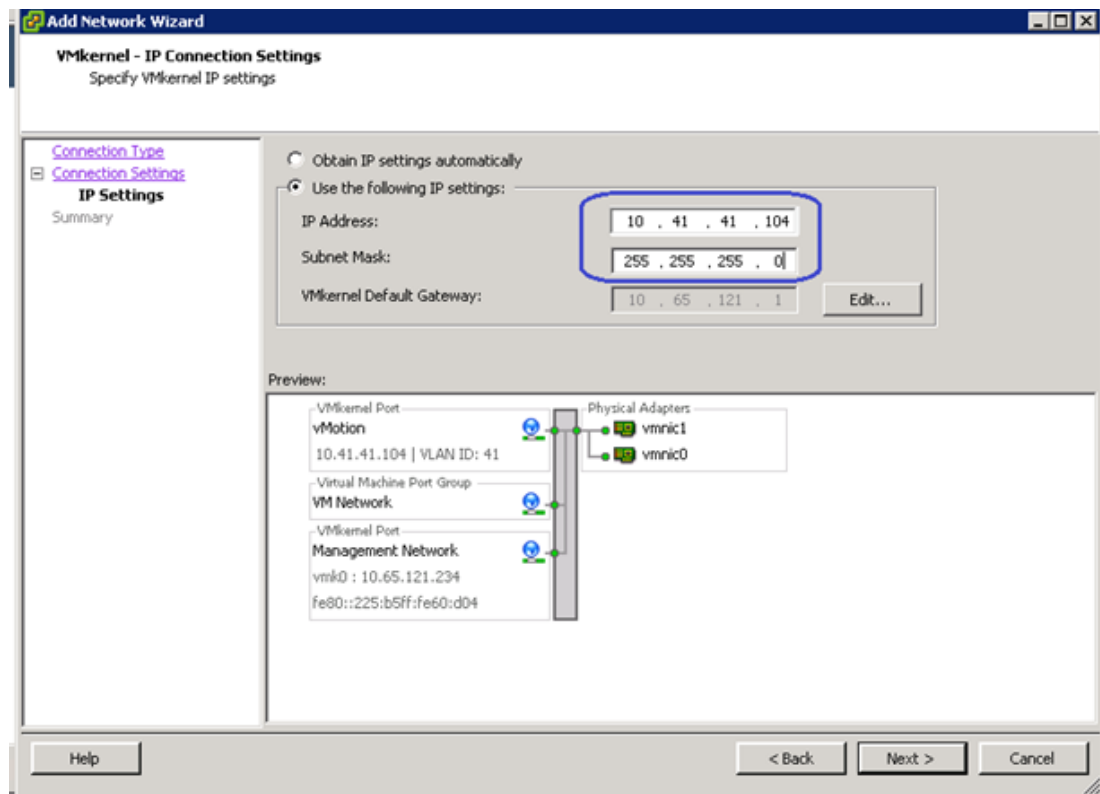
6. Select "VMKernel" and click "Next" on the Add Network Wizard window.



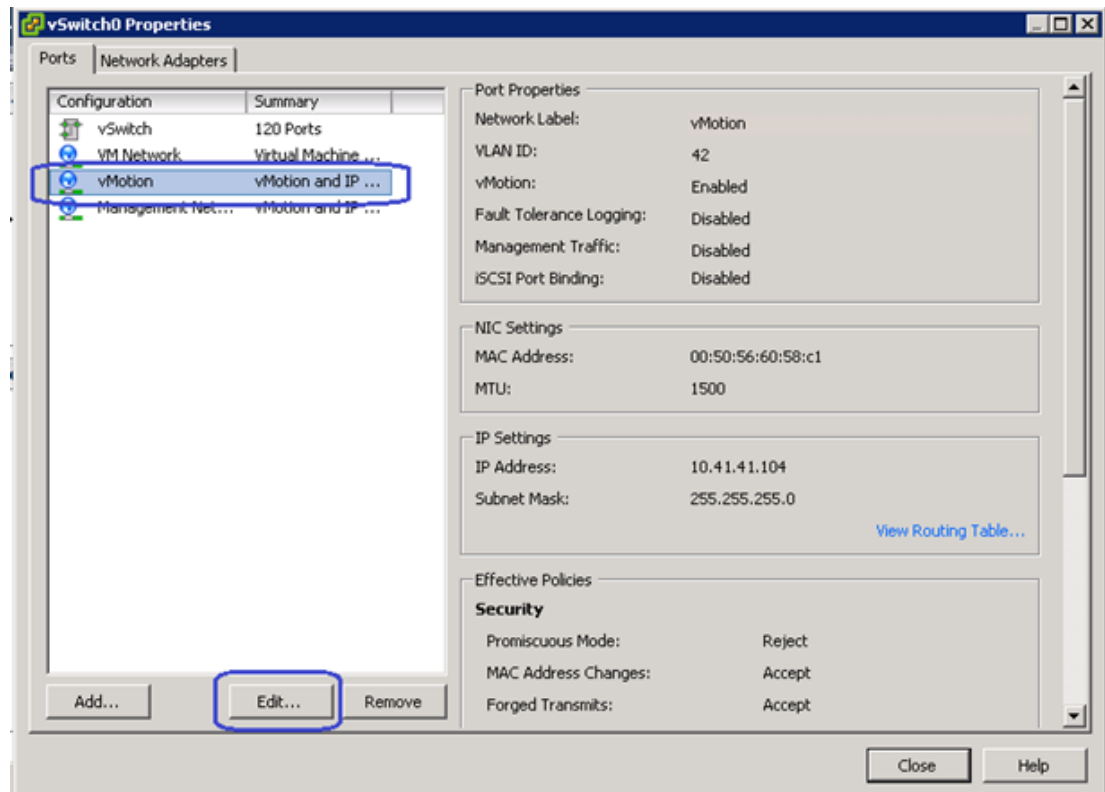
7. Enter "vMotion" as the name for the network label. Select the VLAN ID, as standard vSwitch0 carries both management and vMotion VLANs. Management traffic leaves vSwitch0 untagged, using the native VLAN of the vNIC, but the vMotion traffic must be tagged with appropriate VLAN ID. Select the "Use this port group for vMotion" checkbox.



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



- Click "Next" and deploy the vmkernel. From the vSwitch0 properties window, select the newly created vMotion port group and click "Edit" as shown below:



10. Set the MTU to "9000" and click OK. Click "Close" on the parent window.
11. Repeat these steps for all the ESXi hosts in the cluster. When 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 "do not fragment".

```

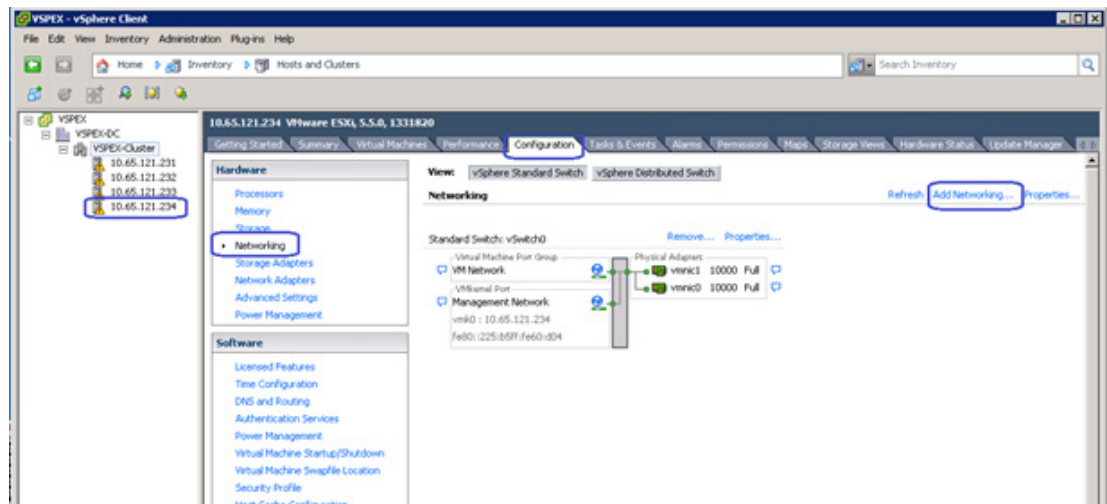
~ # vmkping -d -s 8972 10.41.41.101
PING 10.41.41.101 (10.41.41.101): 8972 data bytes
8980 bytes from 10.41.41.101: icmp_seq=0 ttl=64 time=0.161 ms
8980 bytes from 10.41.41.101: icmp_seq=1 ttl=64 time=0.114 ms
8980 bytes from 10.41.41.101: icmp_seq=2 ttl=64 time=0.154 ms

--- 10.41.41.101 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.114/0.143/0.161 ms
~ # vmkping -d -s 8972 10.41.41.102
PING 10.41.41.102 (10.41.41.102): 8972 data bytes
8980 bytes from 10.41.41.102: icmp_seq=0 ttl=64 time=0.078 ms
8980 bytes from 10.41.41.102: icmp_seq=1 ttl=64 time=0.081 ms
8980 bytes from 10.41.41.102: icmp_seq=2 ttl=64 time=0.114 ms

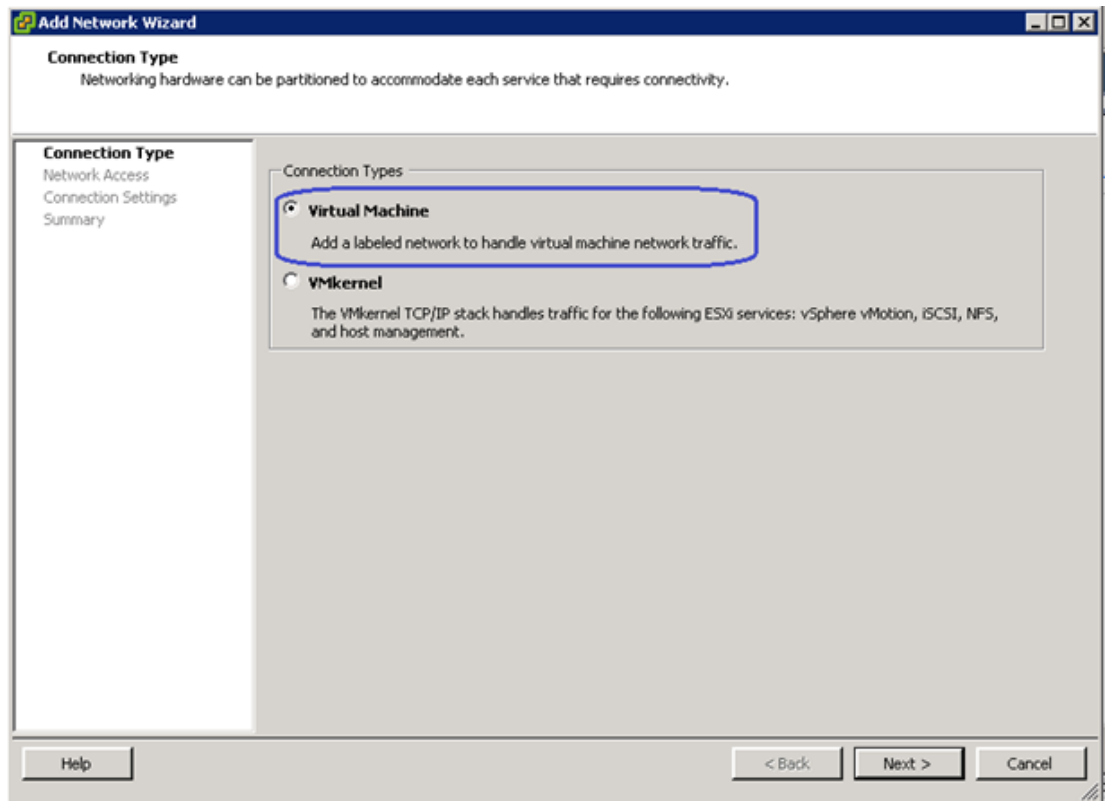
--- 10.41.41.102 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.078/0.091/0.114 ms
~ #

```

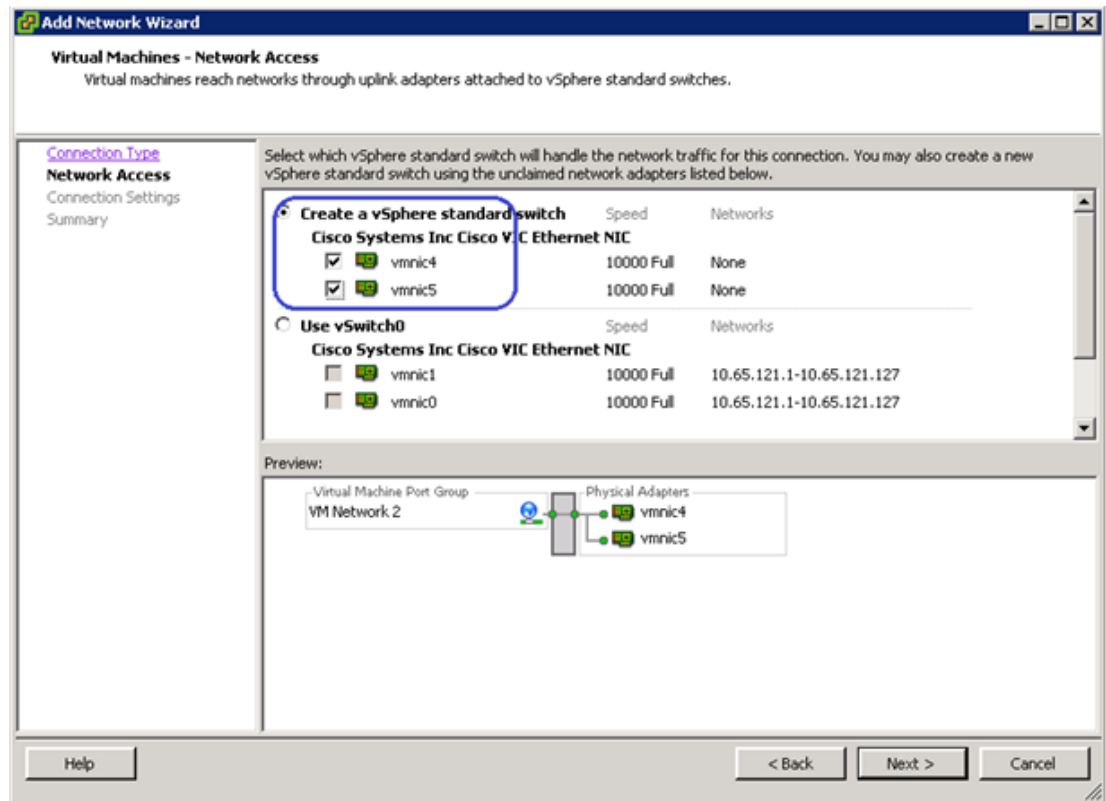
12. From vCenter, select a host from the "Hosts and Clusters" view, click "Configuration" and "Networking". Click "Add Networking" as shown below:



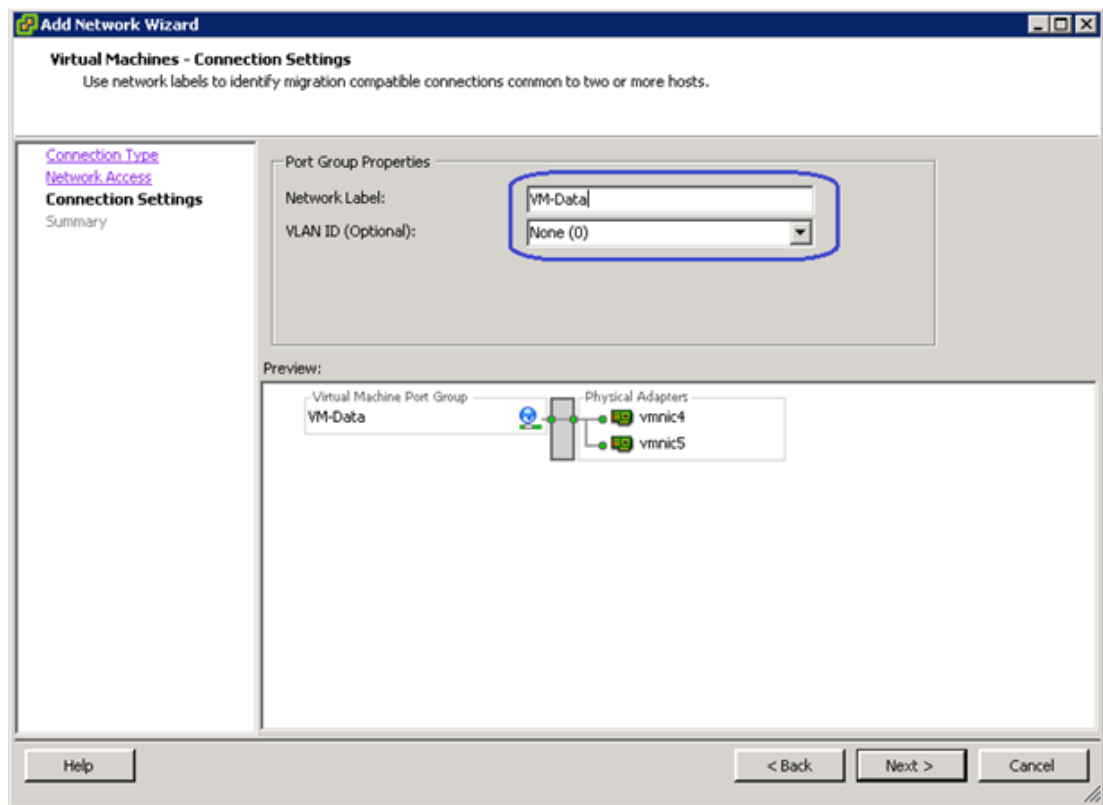
13. Select "Virtual Machine" on the Add Networking Wizard, click "Next".



14. Select the two vmnics corresponding to the VM-Data VNICS and click "Next".



15. Provide a "VM-Data" network label and keep VLAN ID as "0" to signify absence of VLAN tag. Click "Next".



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

Install and Configure Cisco Nexus 1000v

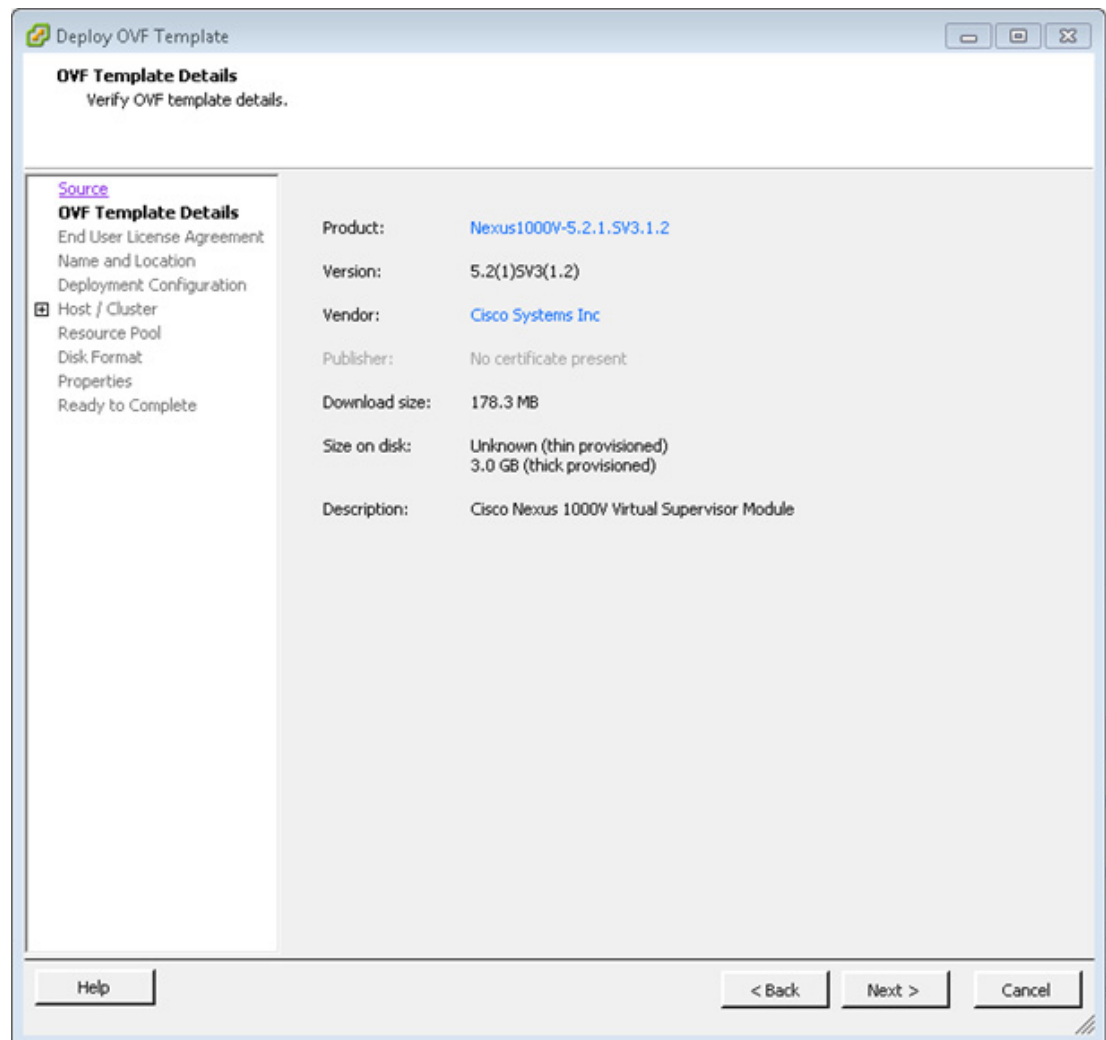
Cisco Nexus 1000v is a Cisco NX-OS based virtual switch that replaces the native vSwitch in the VMware ESXi hosts by a virtual Distributed Switch (vDS). The control plane of the Cisco Nexus 1000v switch is installed in a VMware Virtual Machine and is known as Virtual Switching Module (VSM). VSM virtual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

- Install Cisco Nexus 1000v VSM VM
- Connect Cisco Nexus 1000v VSM to VMware vCenter
- Configure port-profiles in VSM and migrate vCenter networking to vDS

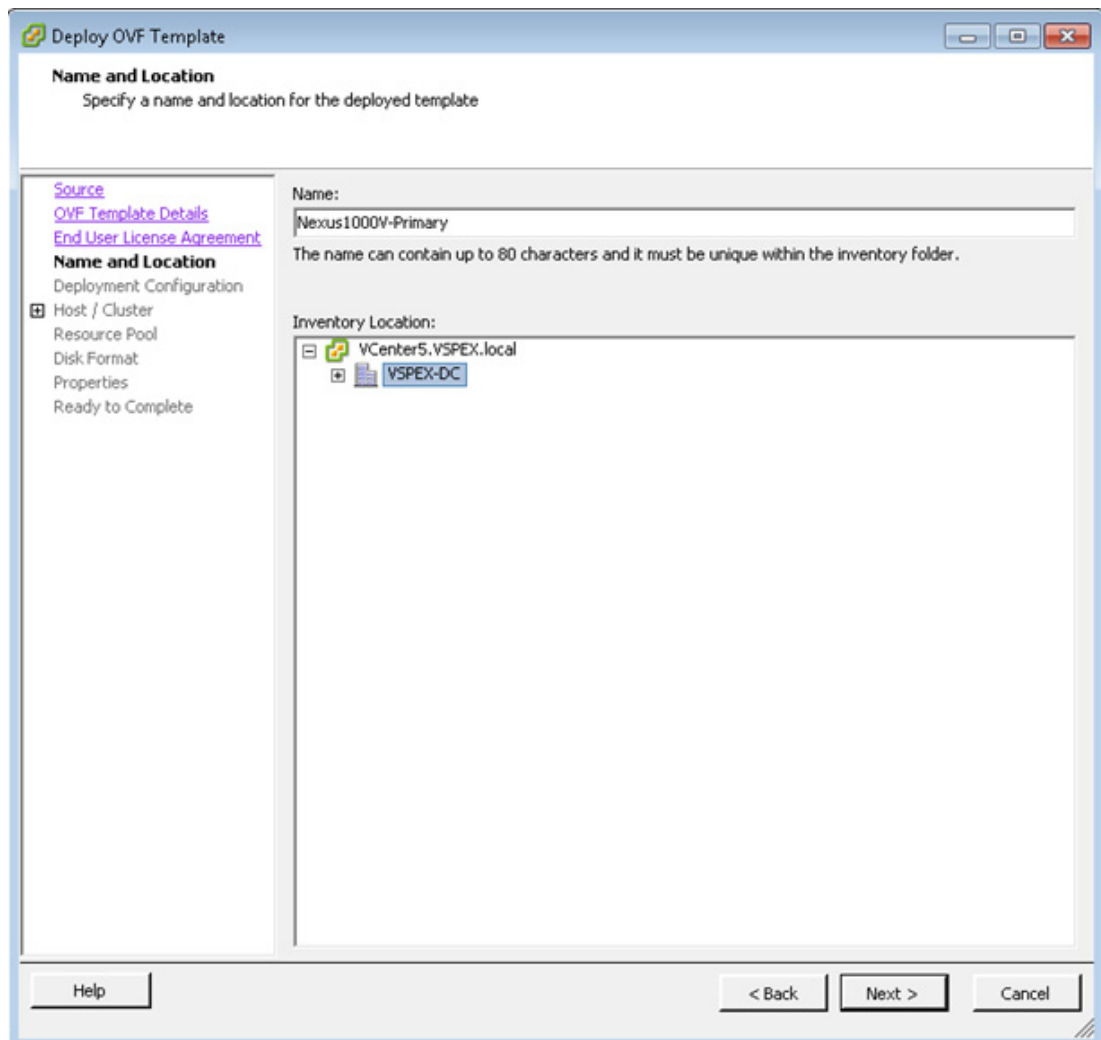
Install Cisco Nexus 1000v VSM VM

The CiscoNexus 1000v VSM VM installation media is available as a VMware virtual machine OVF template. The VSM VM must be deployed on the infrastructure network and not on one of the VSPEX ESXi servers. To install VSM VM, complete the following steps:

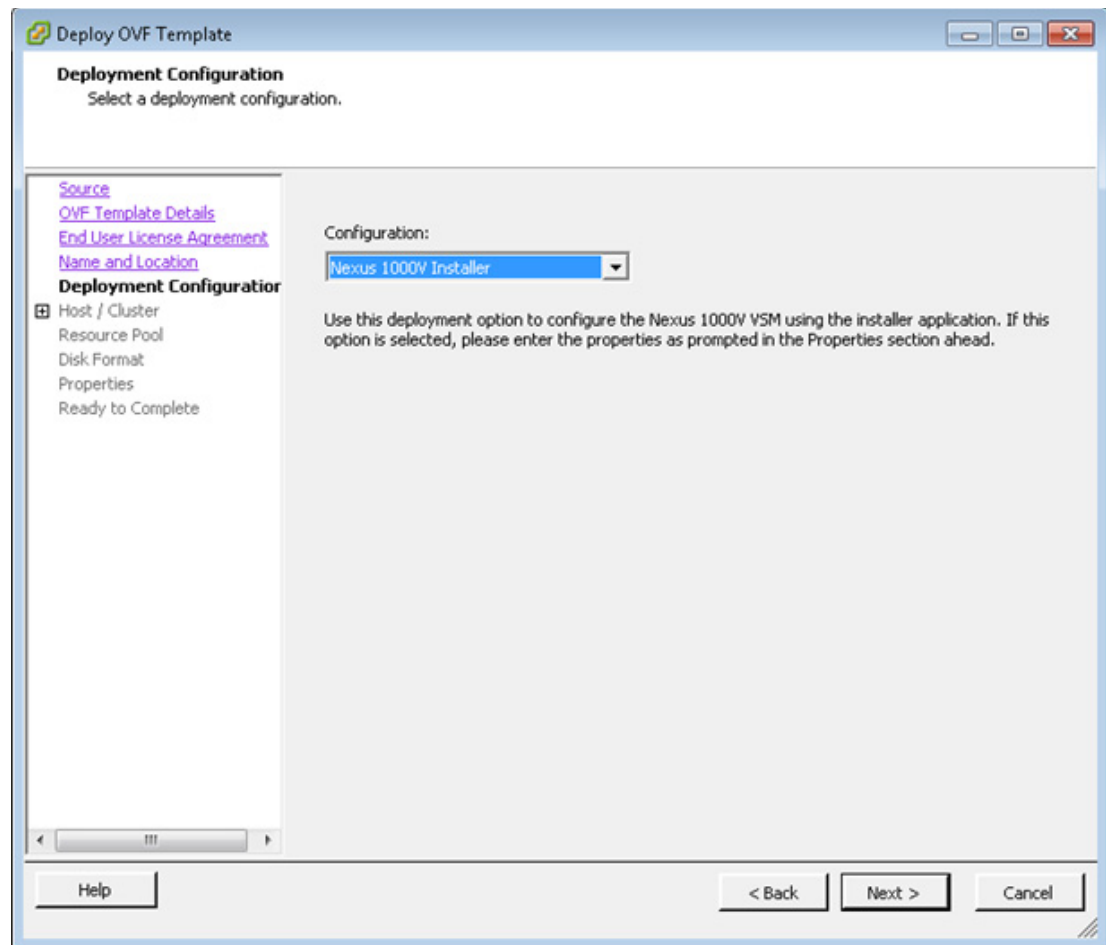
1. From the "Hosts and Cluster" tab in vCenter, select the infrastructure ESX/ESXi host and click "File" > "Deploy new Virtual Machine" thru OVF template. Select Nexus 1000v VSM OVF and click "Next".



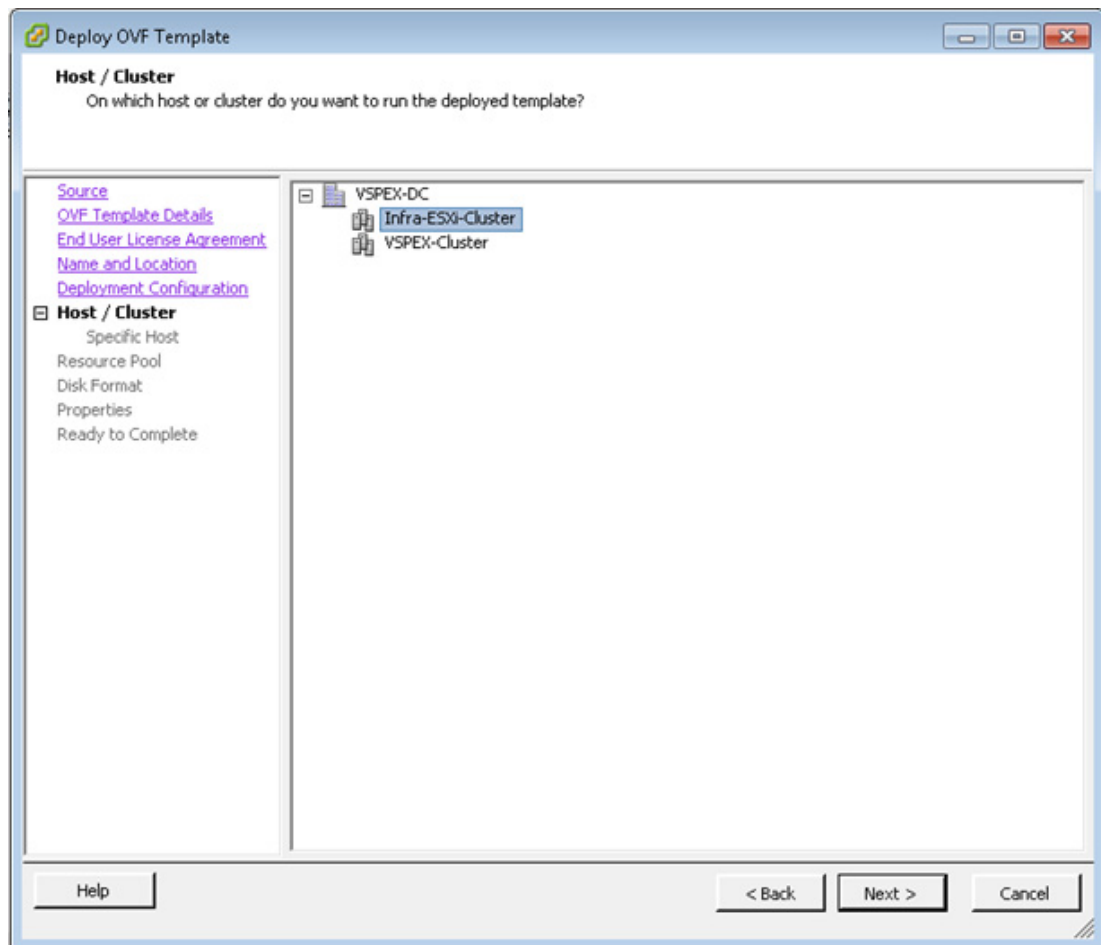
2. Select the datacenter where you want to install the VSM.



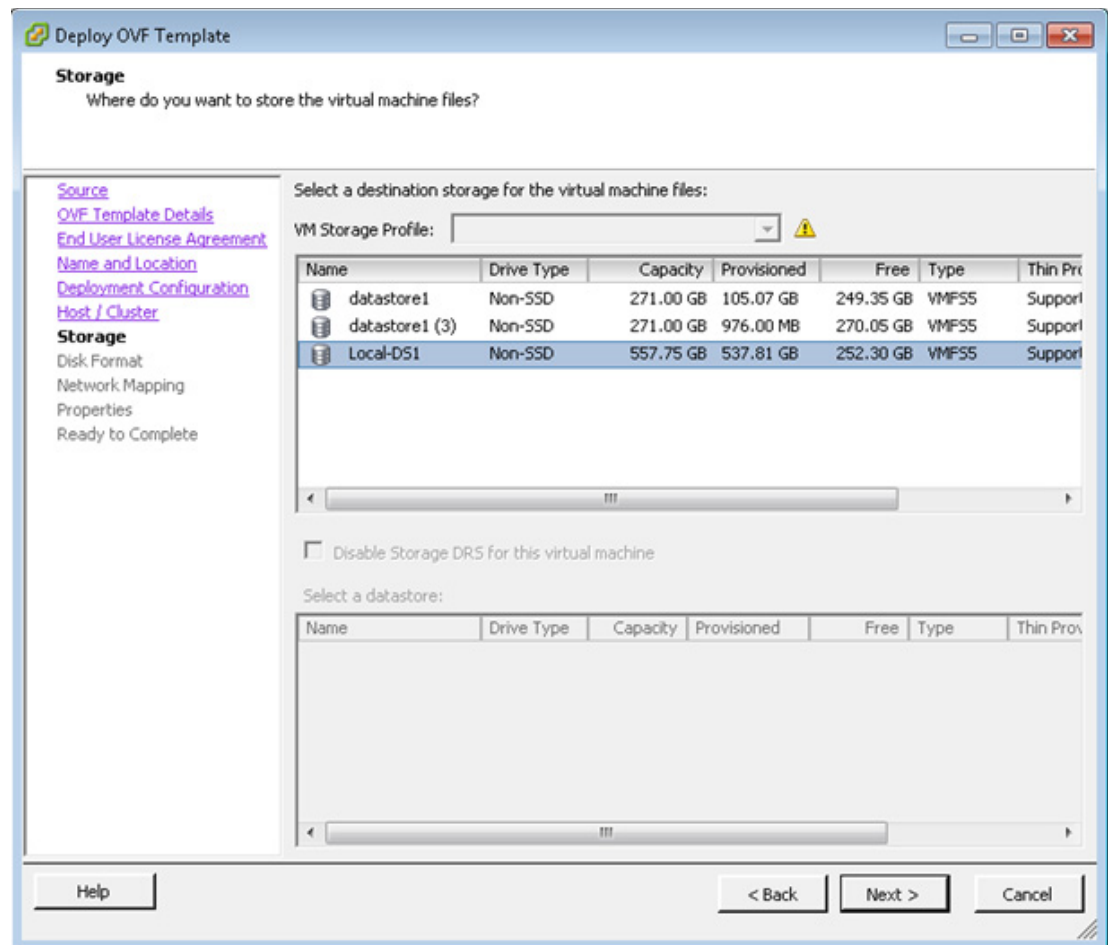
3. Choose "Manually Configure Nexus 1000v".



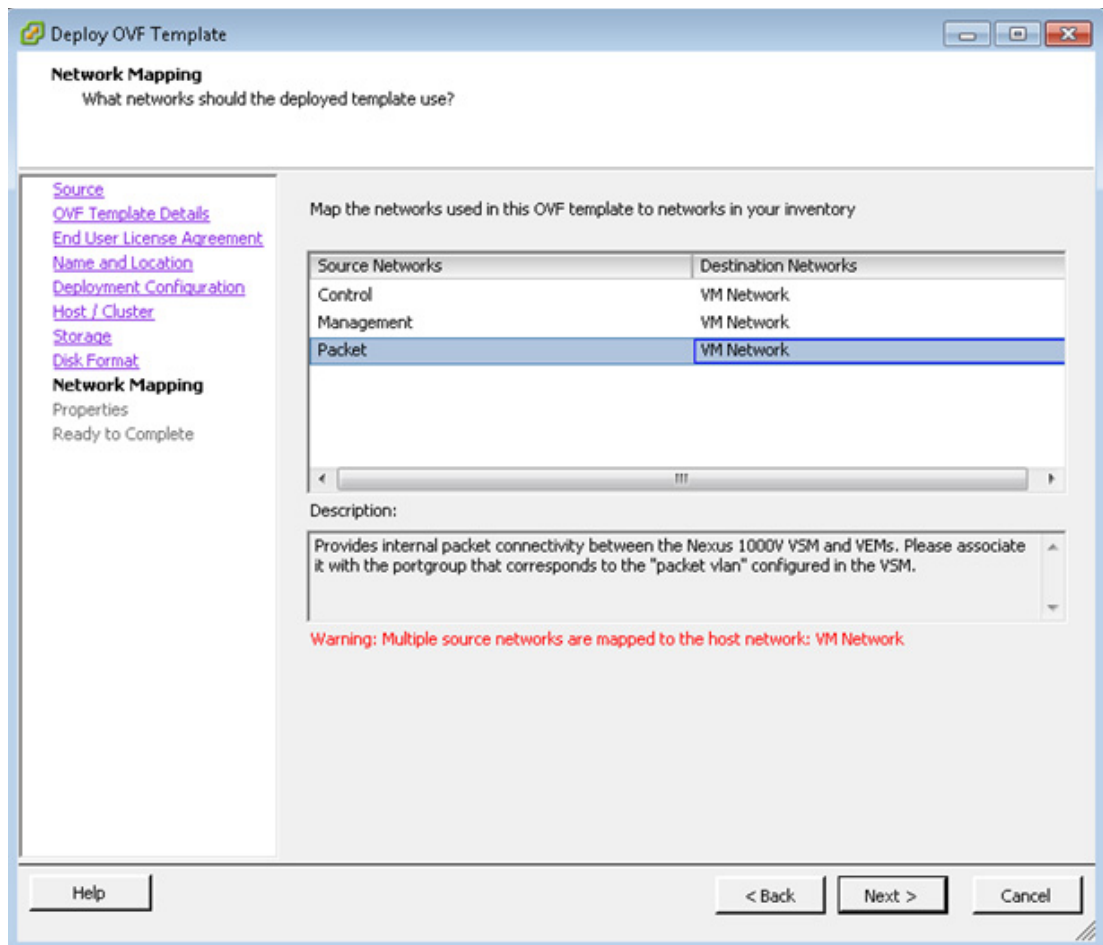
4. Select the host to install the N1k VSM VM.



5. Select the datastore where the VM should be deployed.



6. Select the network mapping for control, management and packet VLANs to the management (infra) VLAN.



- Configure the domain ID (a unique number across multiple N1k VSMs, if there are more than one), administrator password, management IP address and subnet mask.

Deploy OVF Template

Properties
Customize the software solution for this deployment.

Source
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Deployment Configuration](#)
[Host / Cluster](#)
[Storage](#)
[Disk Format](#)
[Network Mapping](#)
Properties
 Ready to Complete

a. VSM Domain ID
DomainId
 Enter the Domain Id (1-1023).

b. Nexus 1000V Admin User Password
Password
 Enter the password. Must contain at least one capital, one lowercase, one number.
 Enter password
 Confirm password

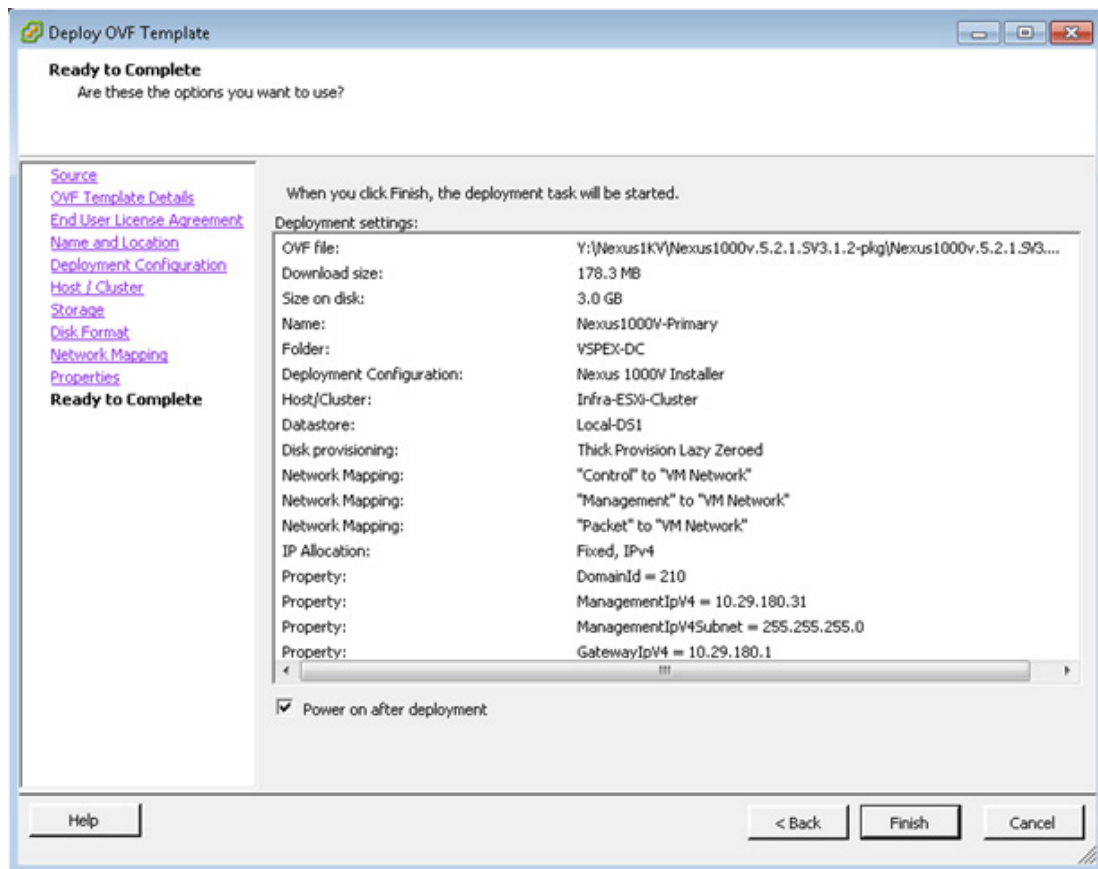
c. Management IP Address
ManagementIpV4
 Enter the VSM IP in the following form: 192.168.0.10
 , , ,

d. Management IP Subnet Mask
ManagementIpV4Subnet
 Enter the Subnet Mask in the following form: 255.255.255.0
 , , ,

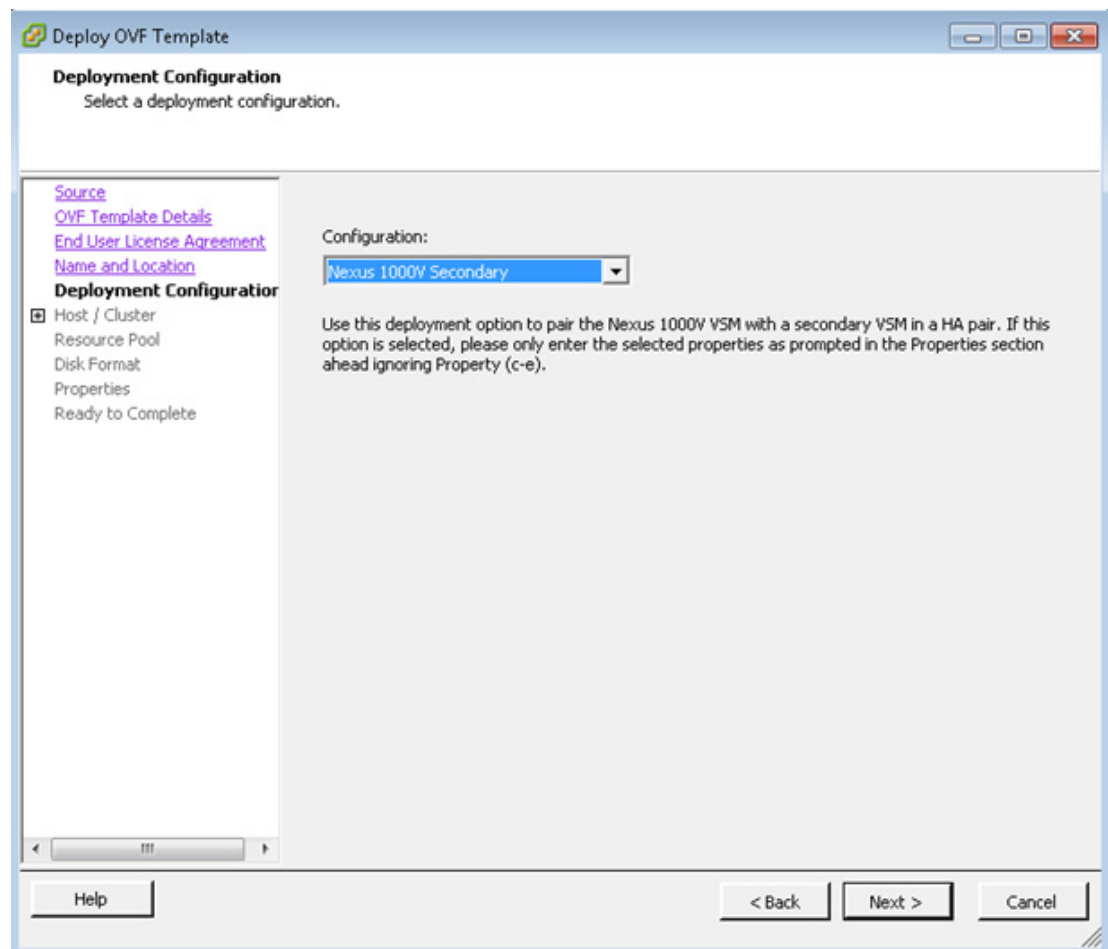
e. Management IP Gateway
GatewayIpV4
 Enter the gateway IP in the following form: 192.168.0.1
 , , ,

Help < Back Next > Cancel

8. Verify the configuration; click "Power on after deployment" and click "Finish" to complete the OVF deployment of VSM virtual machine.



9. When the VSM VM is powered on, click "Console" of the virtual machine in the vCenter and make sure you can logon using user "admin" and the "password" you have provided during initial configuration.
10. 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 "Deployment Configuration", choose "Nexus 1000v Secondary" from the drop-down menu as shown below:



11. Repeat steps 4, 5 and 6 from the original VSM VM deployment. On the "Properties" step, 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.

Deploy OVF Template

Properties
Customize the software solution for this deployment.

[Source](#)
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Deployment Configuration](#)
[Host / Cluster](#)
[Storage](#)
[Disk Format](#)
[Network Mapping](#)
Properties
 Ready to Complete

a. VSM Domain ID
DomainId
 Enter the Domain Id (1-1023).
 210

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
 10 . 29 . 180 . 32

d. Management IP Subnet Mask
ManagementIpV4Subnet
 Enter the Subnet Mask in the following form: 255.255.255.0
 255 . 255 . 255 . 0

e. Management IP Gateway
GatewayIpV4
 Enter the gateway IP in the following form: 192.168.0.1
 10 . 29 . 180 . 1

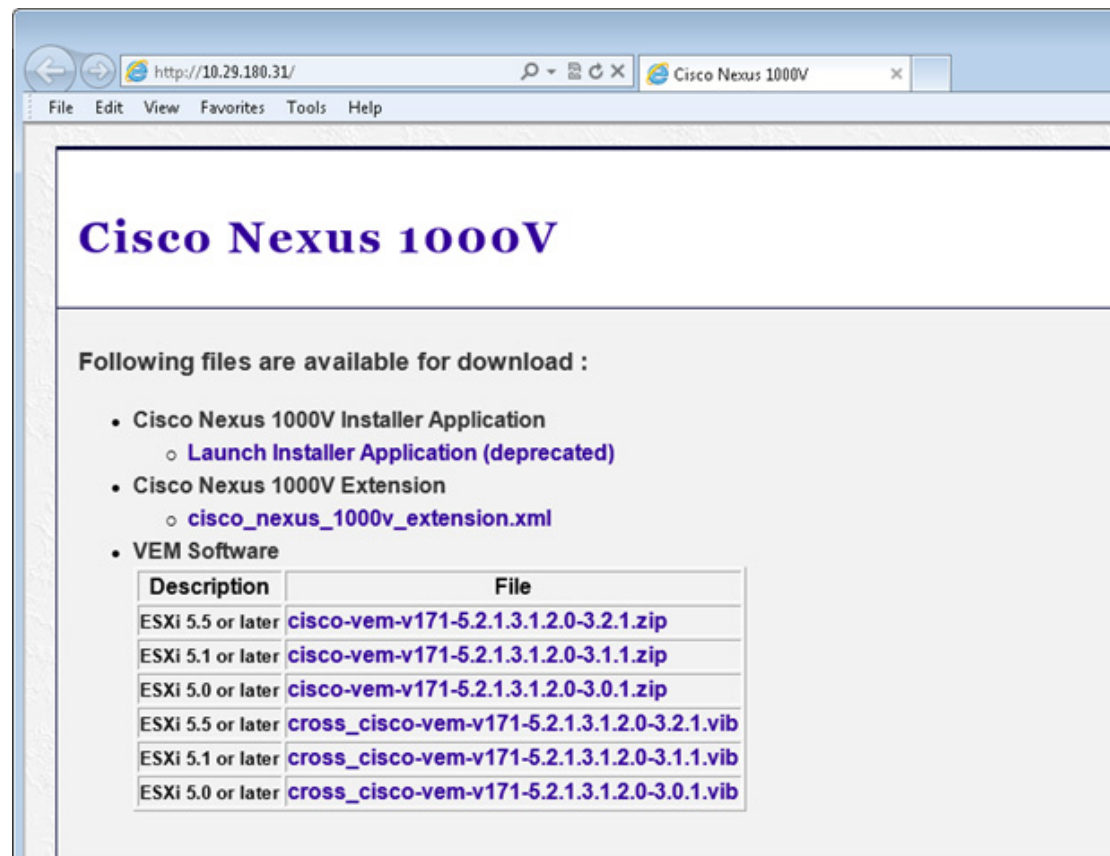
Help < Back Next > Cancel

12. When the VSM IP details are updated, then click Next and check "Power on after deployment" and click Finish to complete the VSM configuration.

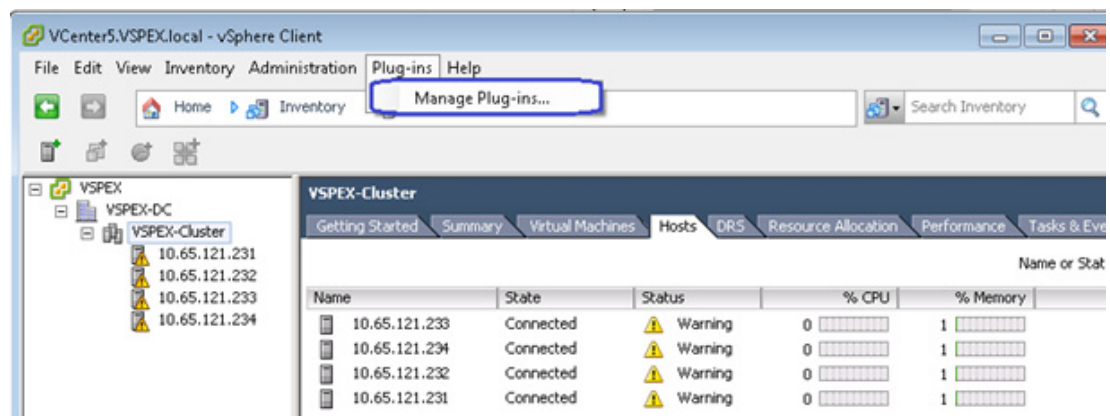
Connecting VSM to vCenter

When the initial setup of VMS VM is finished, the next step is to add it as a plug-in in the vCenter. Complete the following steps:

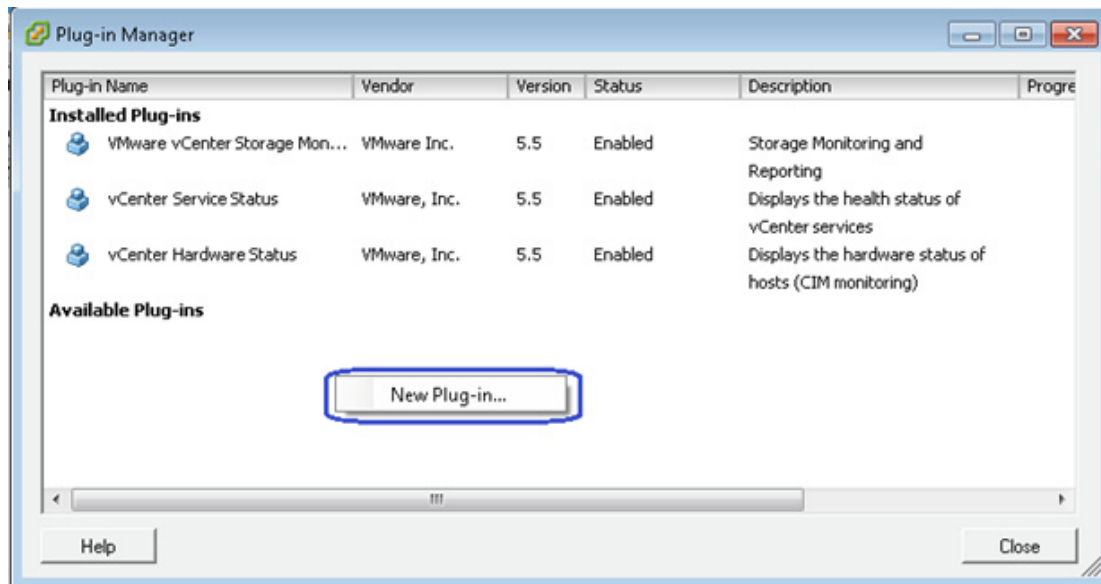
1. Using your browser, access management IP address of the VSM VM.
2. Right-click the `cisco_nexus_1000v_extension.xml` link and save to a location on your local hard drive.



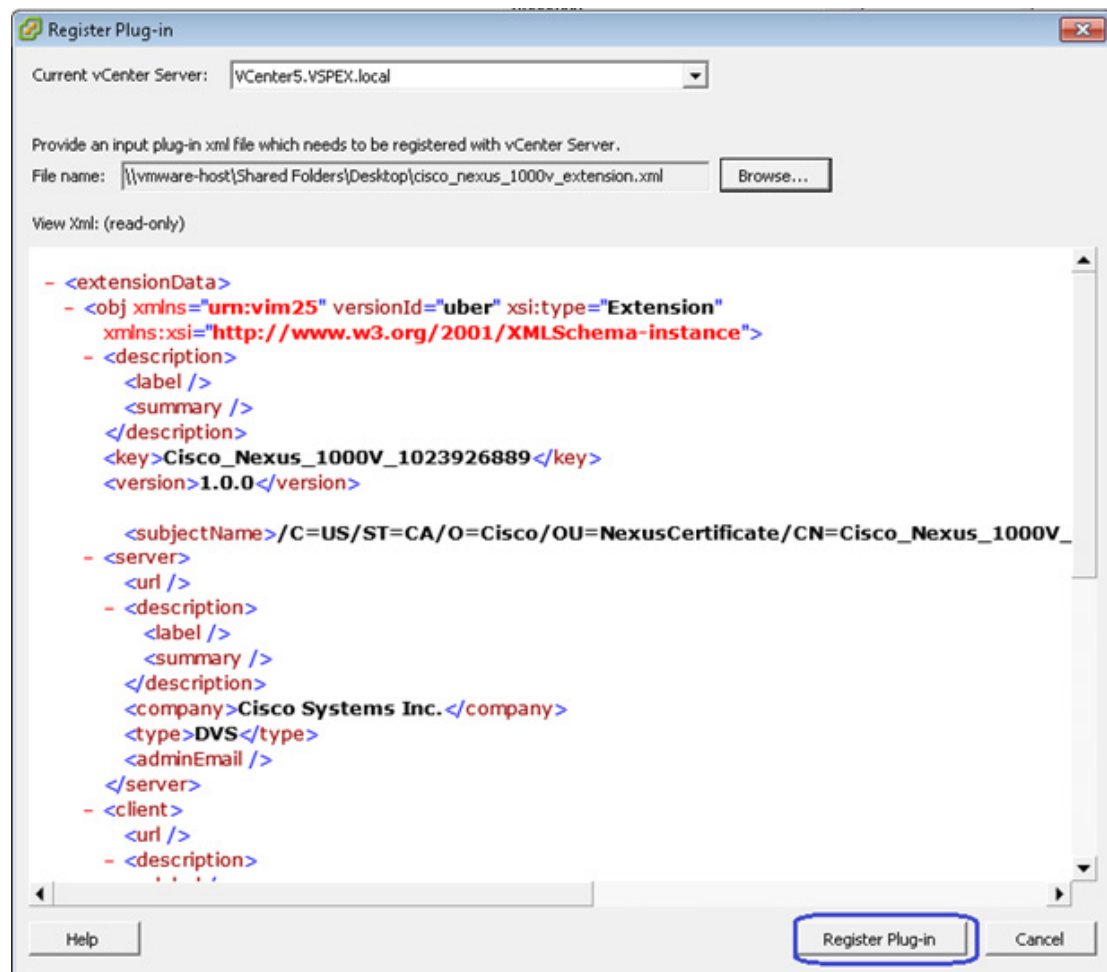
3. From the vCenter, click Plug-ins > Manage Plug-ins.



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



5. Click Browse, select the cisco_nexus-1000v_extension.xml file you just downloaded.
6. Click Register 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. Configure the SVS connection to the vCenter as shown in the image below:

```

10.29.180.31 - PuTTY
login as: admin
Nexus 1000v Switch
Using keyboard-interactive authentication.
Password:
Last login: Sun Nov 23 05:40:38 2014 from 10.21.72.131
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2014, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
switch# config terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# svcs connection vspex
switch(config-svc-conn)# remote ip address 10.29.180.30
switch(config-svc-conn)# vmware dvs datacenter-name vspex-dc
switch(config-svc-conn)# protocol vmware-vim
switch(config-svc-conn)# connect

```

10. Validate the connection using "show svcs connection" and make sure that operational status is "connected" and sync status is "Complete" as shown below:

```

switch(config-svc-conn)# show svcs connections
connection vspex:
  ip address: 10.29.180.30
  remote port: 80
  protocol: vmware-vim https
  certificate: default
  datacenter name: VSPEX-DC
  admin:
  max-ports: 12000
  DVS uuid: ec 82 16 50 40 ff 5e ae-39 03 58 f8 7e 75 dd 83
  config status: Enabled
  operational status: Connected
  sync status: Complete
  version: VMware vCenter Server 5.5.0 build-1312298
  vc-uuid: 4929AD9E-EE59-4DB6-B5A3-6B17238CA684
  ssl-cert: self-signed or not authenticated
switch(config-svc-conn)#

```

11. Create an uplink port-profile 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 as shown below:


```
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 allowed vlan 41
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. Notice the MTU 9000 is configured on a 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 on per vNIC basis. It also provides high-availability by moving the traffic to the alternative adapter when a given fabric is down.

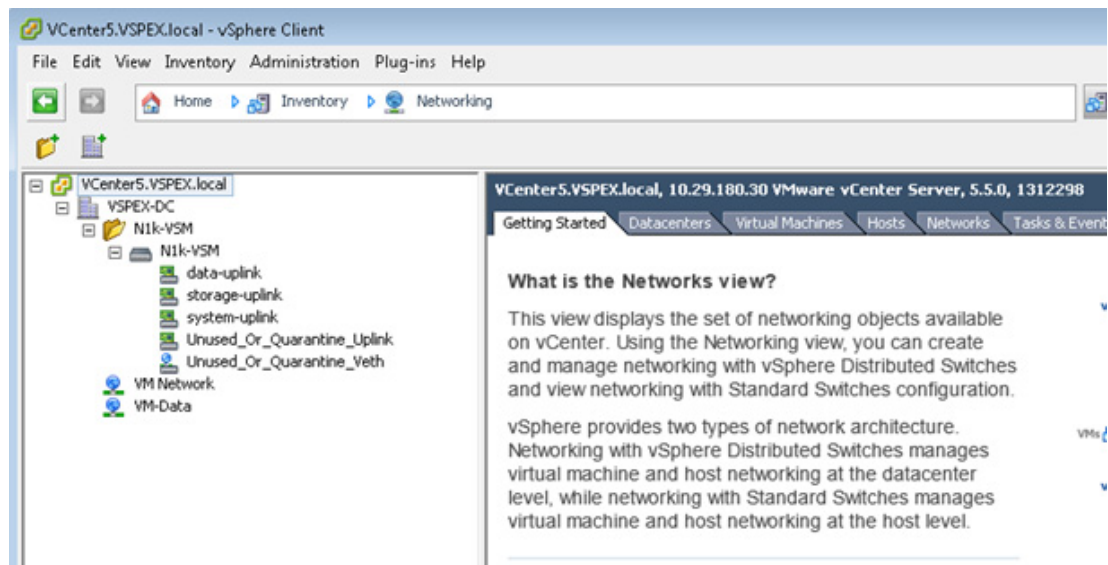
12. Create a storage-uplink port-profile as shown below. The storage-uplink port-profile corresponds to the Storage vNICs of the service profile.

```
N1k-VSM(config)# port-profile type ethernet storage-uplink
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)# 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
N1k-VSM(config)#
```

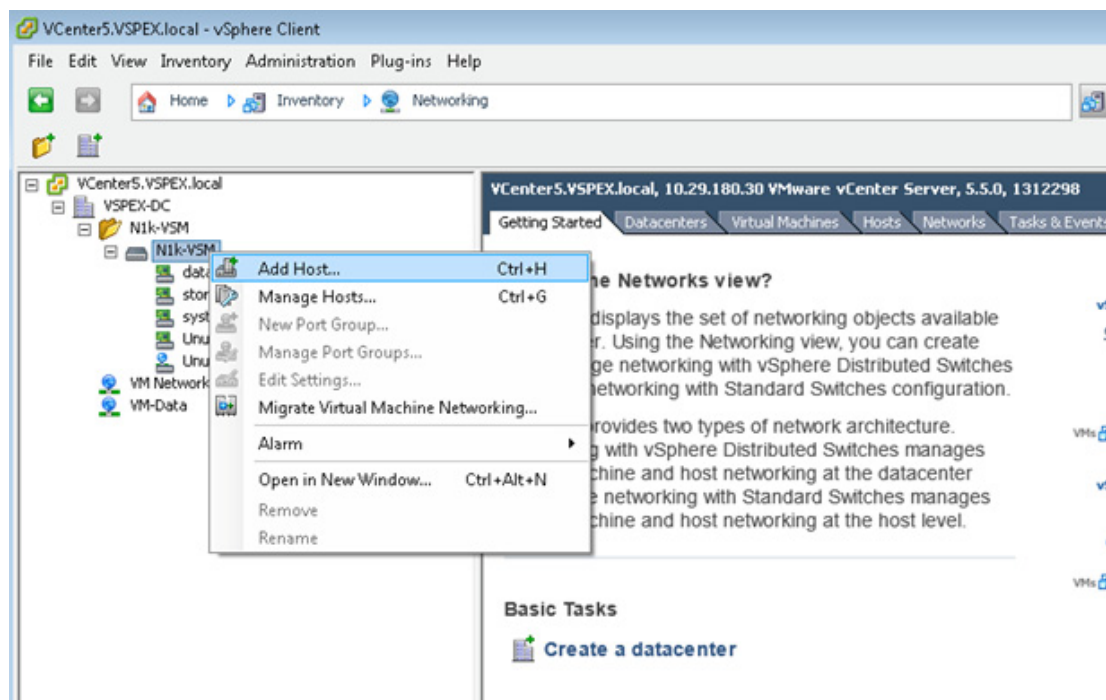
13. Create a data-uplink port-profile as shown below. The data-uplink port-profile corresponds to the Data vNICs of the service profile.

```
N1k-VSM(config)# port-profile type ethernet data-uplink
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)# 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
N1k-VSM(config)#
```

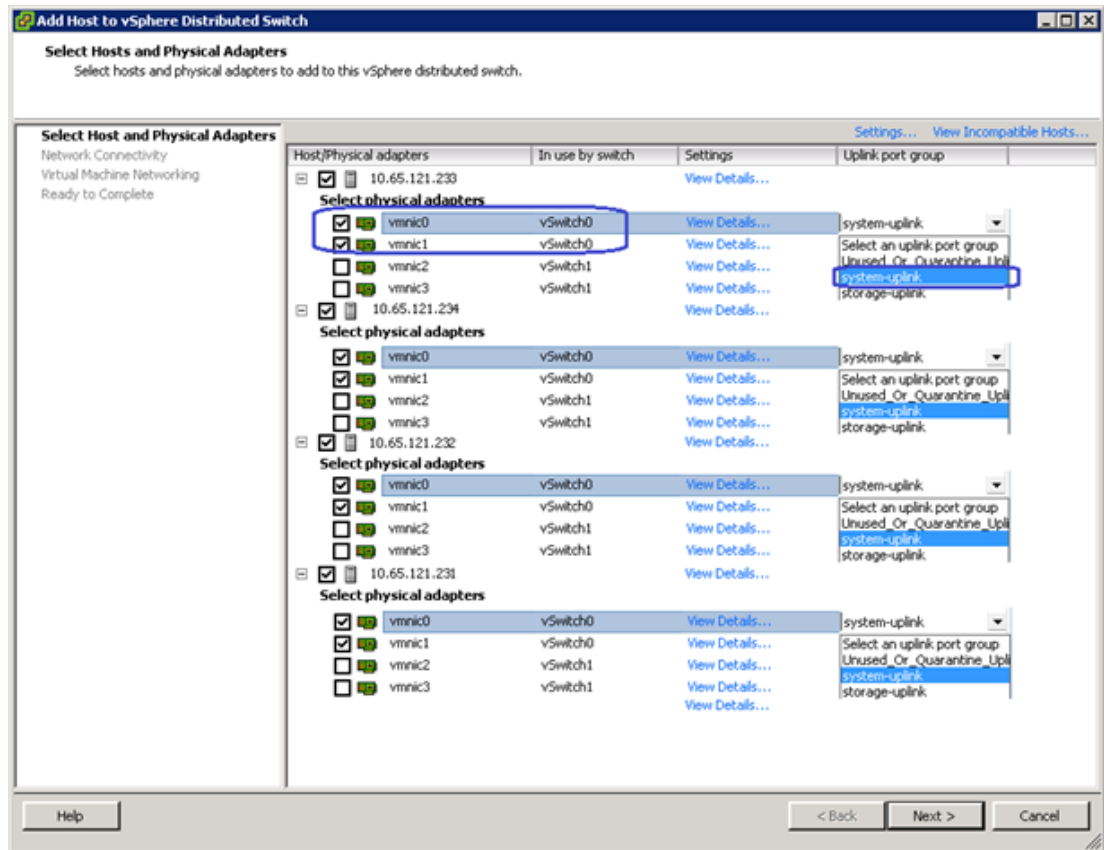
14. When VSM is connected to the vCenter, it shows up as a virtual Distributed Switch in the vCenter's "Network" view as shown below:



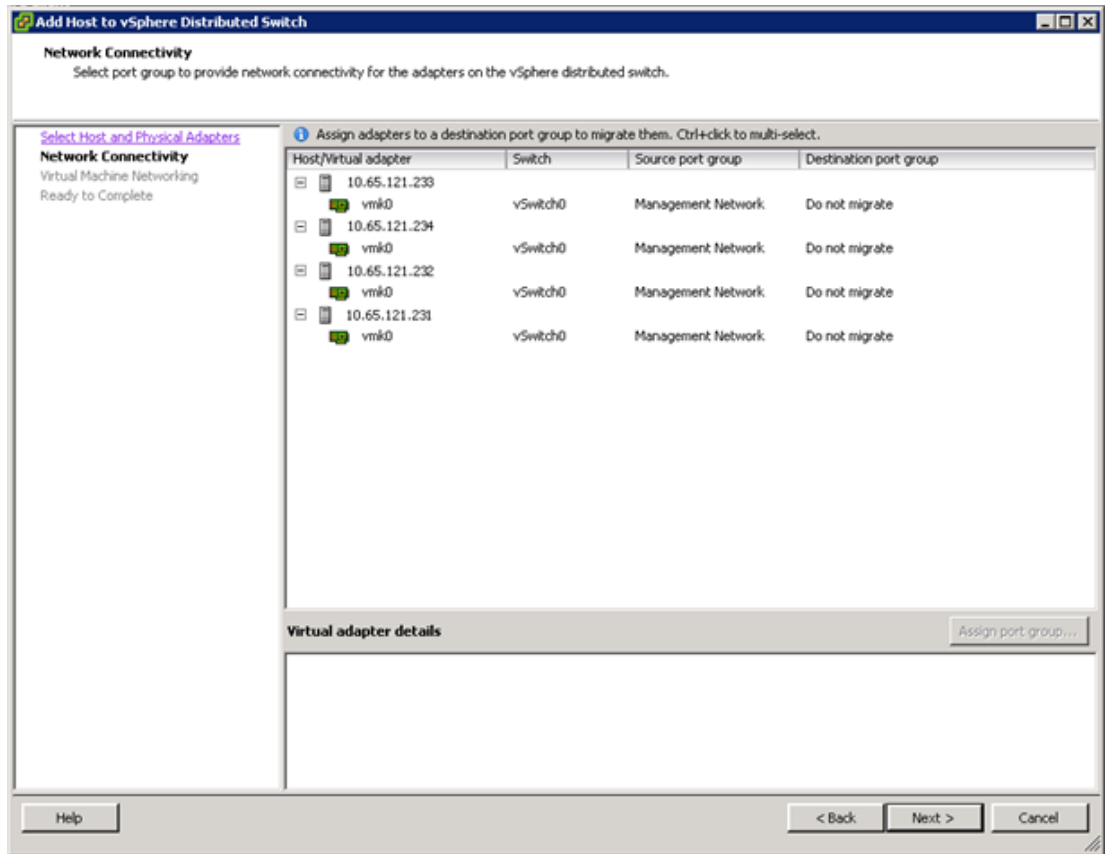
15. Add hosts to the vDS as shown below:



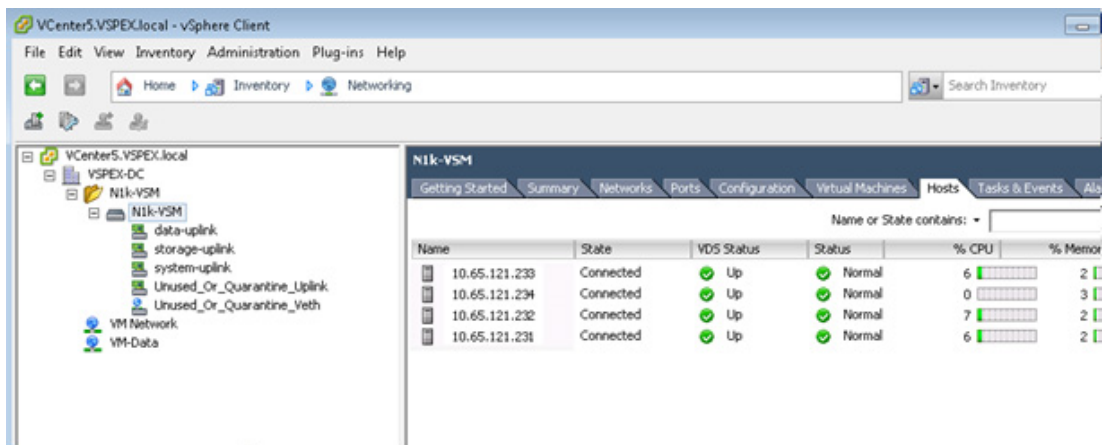
16. On the next dialog box, select all the VSPEX ESXi hosts and add appropriate adapters using the uplink port-profiles created in previous step. Refer to [Table 15](#) for vNICs to uplink port-profile mapping.



- Do not migrate management VM kernel from the native vSwitch to vDS in the last step, click "Next" and finish the add host wizard.



18. Make sure that all the hosts are successfully added to the vDS.



Configure Port Profiles and Add Virtual Machines

The last step of the Cisco Nexus 1000v configuration and its integration with vCenter is creating port profiles and using them in the virtual machines in the vCenter. This is possible to do after making disk space for VMs on the storage array and deploying the VMs. To configure port profiles for VMs, complete the following steps:

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

```
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)#
```

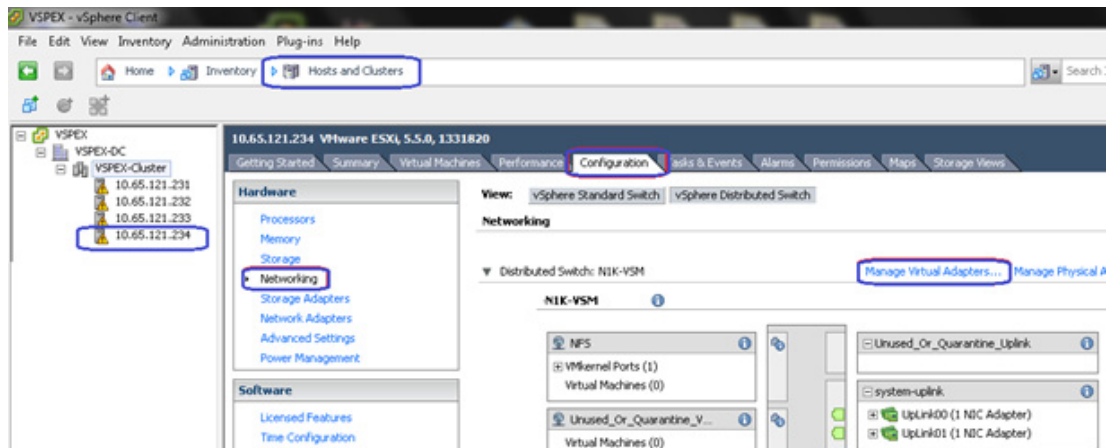
2. Create a port-profile for vMotion traffic. Max-ports can be set to number of hosts you have in the architecture.

```
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 requirements. You can set "max ports" to appropriate values based on the number of VMs being configured. The following is a sample port profile:

```
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)#
```

4. When the port profiles are configured, launch vCenter Center using vSphere client, then select "Hosts and Clusters" tab, select the ESXi host, click the "Configuration" tab, select "Networking", view "vSphere Distributed Switch", and click "Manager Virtual Adapters" as shown in the image below:



- Click "Add" in the wizard, click "New virtual adapter", and click "Next".
- Select "VMKernel".
- 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 "Finish" to deploy the VNIC.
- Add one more VMKNic (VM Kernel NIC) for vMotion. When providing the port-profile name, make sure that you select "vMotion" port-profile and click on the check box "Use this virtual adapter for vMotion".
- Repeat creating the two vmknics virtual adapters for all the ESXi hosts.
- Connectivity between all the vmknics can be tested by enabling SSH access to ESXi host, logging on 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.
- When NFS share is available, the NFS datastore can be discovered and mounted thru vCenter. Virtual machines can be deployed on these NFS datastore using the VM-Data port-profile for the network access. Verify the port-profile usage using "show port-profile brief", "show port-profile usage", or "show port-profile name <name>" command. The following is one sample output.

```

N1k-VSM# show port-profile brief
-----
Port Profile                                Profile Type    Profile State  Conf Items  Eval Items  Assigned Intfs  Child Profs
-----
NFS                                         Vethernet      1              4           4           4              0
Unused_Or_Quarantine_Uplink               Ethernet       1              1           0           0              0
Unused_Or_Quarantine_Veth                  Vethernet      1              1           0           0              0
data-uplink                              Ethernet       1              5           5           12              0
VM-DATA                                   Vethernet      1              3           3           100             0
storage-uplink                            Ethernet       1              5           5           12              0
system-uplink                             Ethernet       1              5           5           12              0
vMotion                                   Vethernet      1              4           4           4              0
-----
Profile Type    Assigned Intfs  Total Prfls  Sys Prfls  Parent Prfls  Child Prfls  UsedBy Prfls
-----
Vethernet      108            4           3           4           0           3
Ethernet       12             3           1           3           0           1
N1k-VSM#

```

- If you issue "show port-profile name <uplink-port-profile-name>" command, you can see the implicit creation of port-channels on the 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 would be also listed as assigned interfaces. Port-channel status can be further viewed / validated using "show port-channel brief" command from VSM VM.

This concludes the Cisco Nexus 1000v configuration.

Configure Storage for Virtual Machine Data Stores

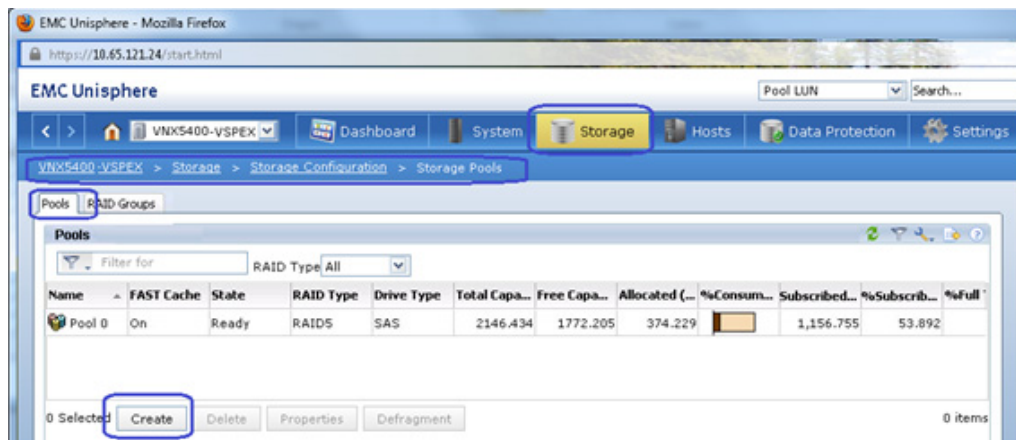
This section is divided in two subsections:

- [Configuring Virtual Machine Data Store for the FC-variant of the Solution](#)
- Configuring VM data store for the NFS-variant of the solution

Configuring Virtual Machine Data Store for the FC-variant of the Solution

To configure the data store, complete the following steps:

- Log into EMC VMX Unisphere and click the "Storage" tab. Click "Storage Configuration" > "Storage Pools" and click the "Pools" tab. Click "Create".



2. Create a new pool. Select "Manual" disk selections and choose 45 SAS disks and 2 Flash disks to create one pool.

VNX5400-VSPEX - Create Storage Pool

General Advanced

Storage Pool Parameters

Storage Pool Type: ☒ Pool ☐ RAID Group

☒ Scheduled Auto-Tiering

Storage Pool ID: 3

Storage Pool Name: Pool 3

Extreme Performance

RAID Configuration: RAID1/0 (4+4) Number of Flash Disks: 2

Performance

RAID Configuration: RAID5 (4+1) Number of SAS Disks: 45 (Recommended)

Distribution

Extreme Performance : 183.453 GB (0.75%)
Performance : 24156.343 GB (99.25%)

Disks

☐ Automatic ☐ Use Power Saving Eligible Disks

☒ Manual

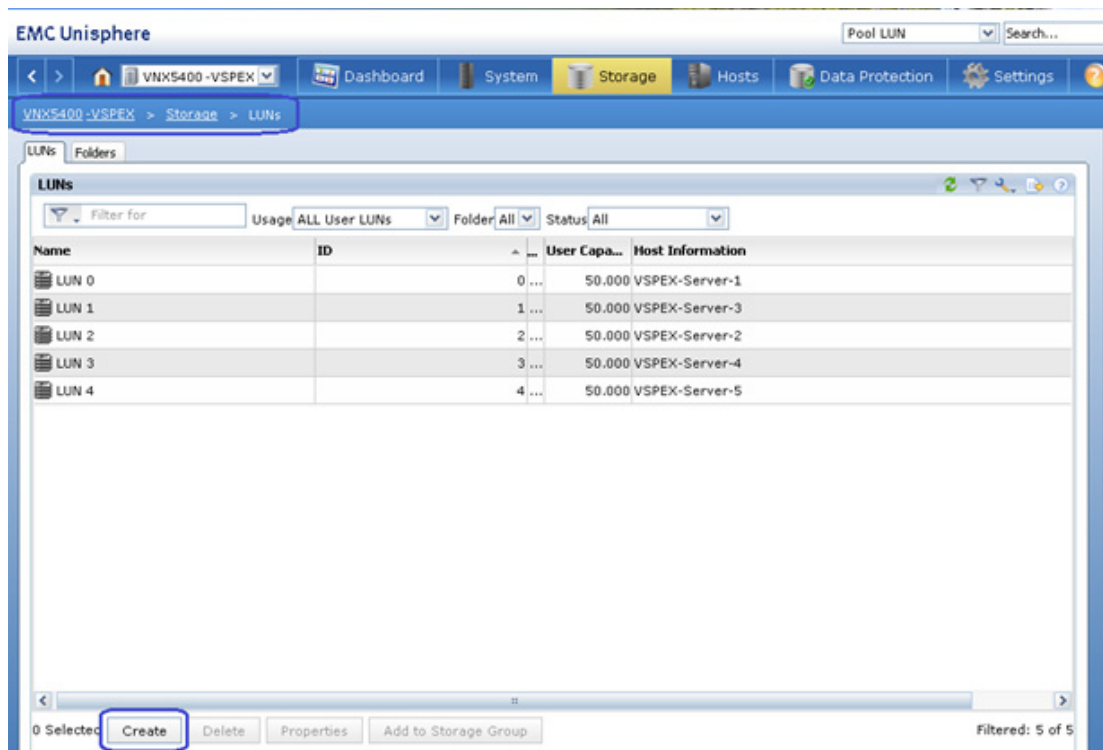
Total Raw Capacity: 24339.79...

Disk	Capacity	Drive Type	Model	State
Bus 0 Enclosure 1 Disk 5	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 6	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 7	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 8	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 9	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 10	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 11	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 12	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 13	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 14	536.808 GB	SAS	STE60005 CL...	Unbound

☒ Perform a background verify on the new storage

OK Apply Cancel Help

- Repeat this step for two more times to create a total of three pools for VM data storage. For the third pool, add 20 SAS drives and two Flash drives.
- Click "Storage" > "LUNs". Viewable are five boot LUNs created for five hosts. Click "Create" to create the LUN for VM data store.

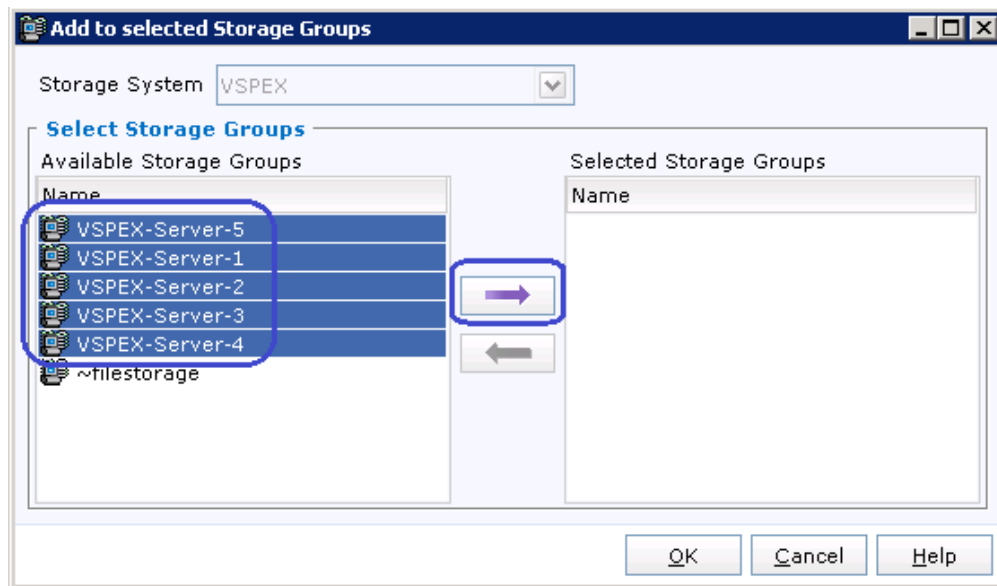


5. Select storage pool type as "Pool", and select the first VM data Pool ID from the drop-down menu, which was created in step 2. Make sure "Thin" provisioning check-box is checked. Select User Capacity to 5 TB and create two LUNs per pool and click "Apply".

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

Name	ID	State	Thin	User Capacity	Current Owner	Host Information	Initial Tier	Additional L...	Tiering Pol...
LUN 10	10	Ready	On	8192.000	SP A	VSPEX-Server-2; VSP...	Highest Av...		Auto-Tier
LUN 13	13	Ready	On	8192.000	SP B	VSPEX-Server-2; VSP...	Highest Av...		Auto-Tier

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



9. On the LUNs screen, you will see the storage group (and host) access for LUNs .

EMC Unisphere

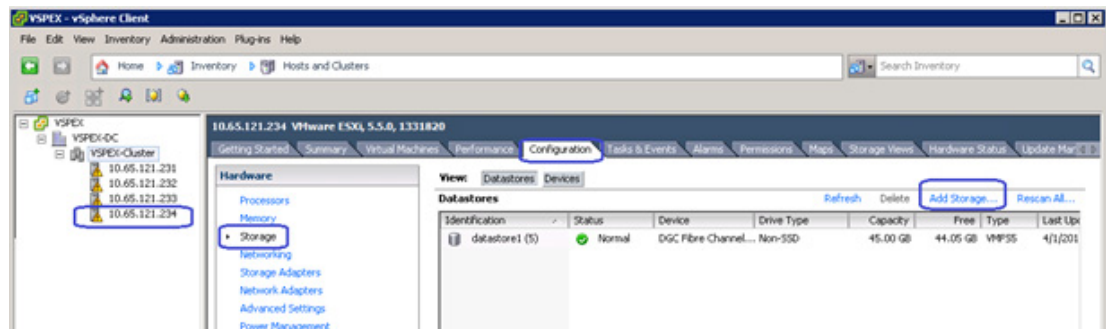
VNX5400-VSPEX > Storage > LUNs

LUNs Folders

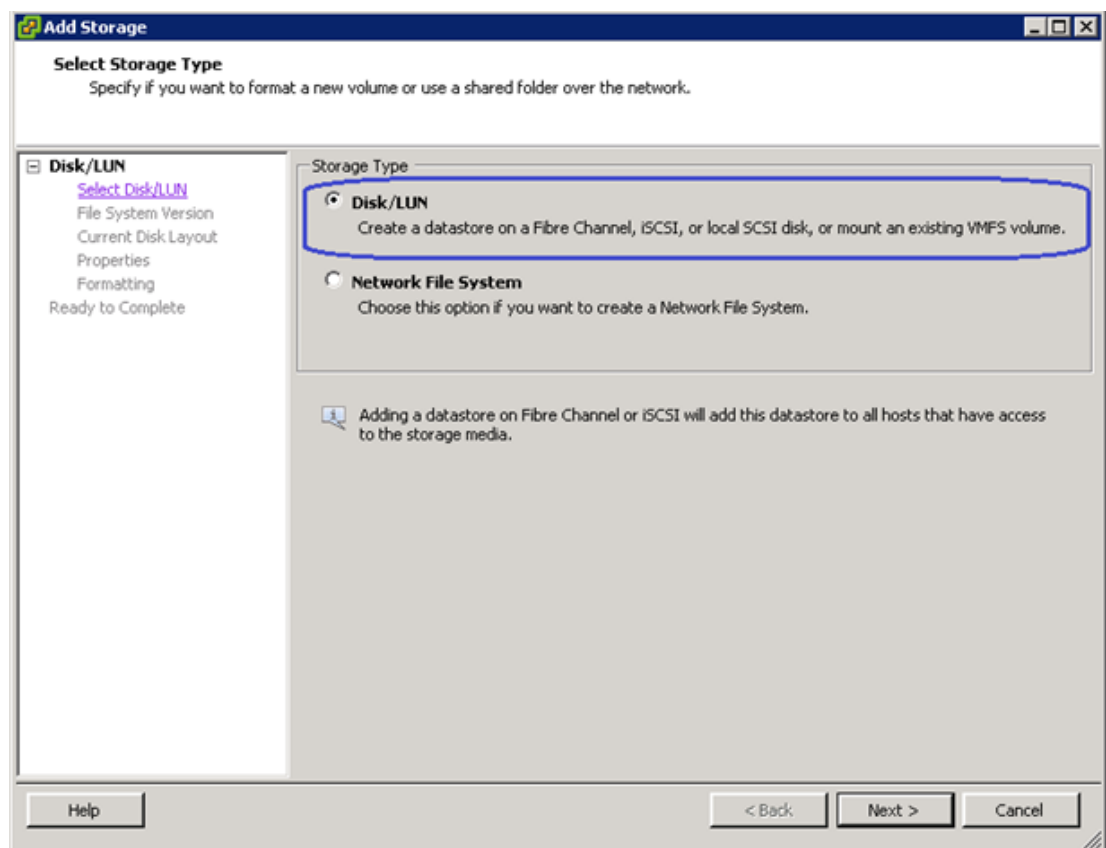
Filter for: Usage: ALL User LUNs Folder: All Status: All

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	2141.336
LUN 6		6 Ready	2141.336	2141.336
LUN 7		7 Ready	2141.336	2141.336
LUN 8		8 Ready	2141.336	2141.336
LUN 9		9 Ready	2141.336	2141.336
LUN 10		10 Ready	2141.336	2141.336
LUN 11		11 Ready	2141.336	2141.336
LUN 12		12 Ready	2141.336	2141.336
LUN 13		13 Ready	2141.336	2141.336
LUN 14		14 Ready	2141.336	2141.336
LUN 15		15 Ready	2141.336	2141.336
LUN 16		16 Ready	2141.336	2141.336
LUN 17		17 Ready	2141.336	2141.336
LUN 18		18 Ready	2141.336	2141.336

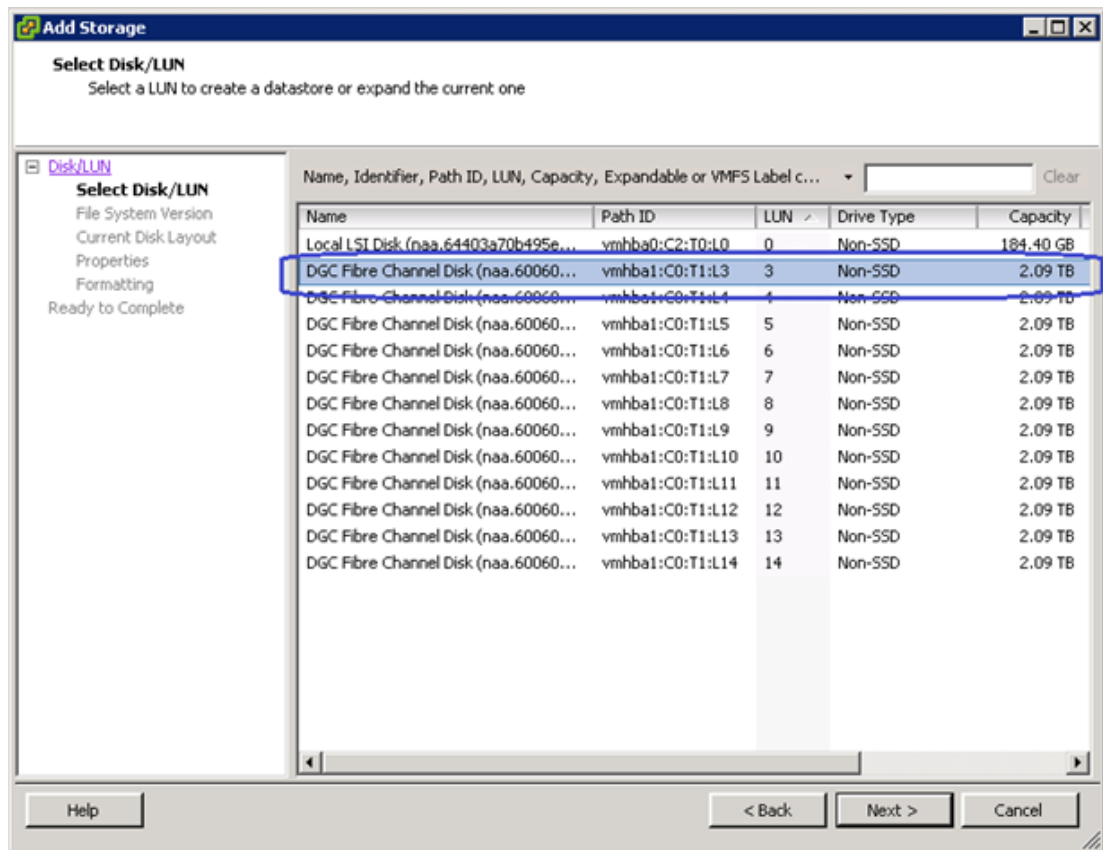
10. Log into vCenter GUI, select a particular host from the "Hosts and Clusters" view, click "Configuration" and "Storage". Click "Add Storage".



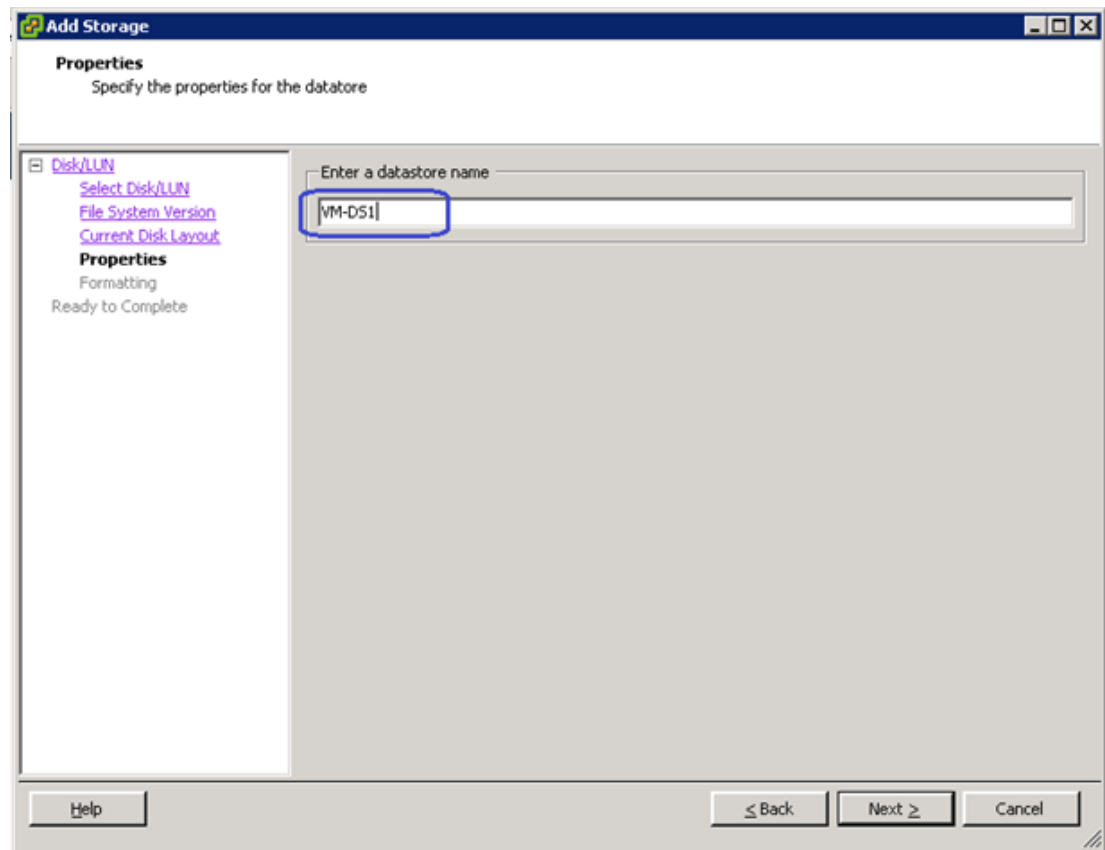
11. Click "Disk/LUN" and click "Next".



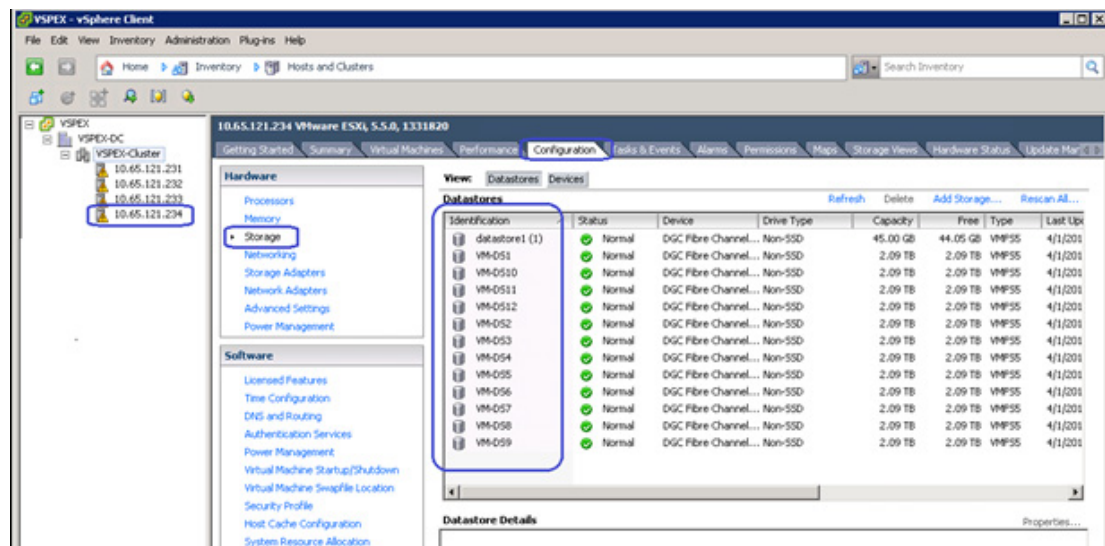
12. Select the first "DGC Fibre Channel Disk (..)" from the list and click "Next".



- Enter "VM-DS1" as the name for the first data-store and click Next and then Finish.



14. Repeat steps 10 to 13 for all the datastores. When data stores are added to one host, it will automatically be available for the other hosts too. The result would look like following (on each host):



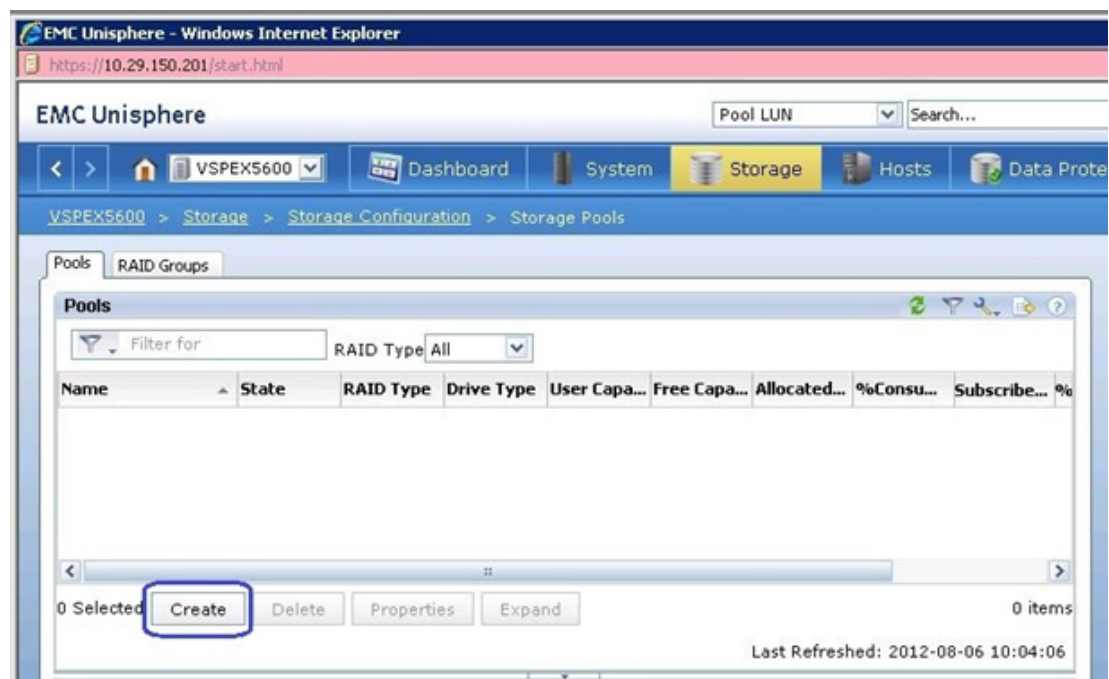
Configuring VM Data Store for the NFS-Variant of the Solution

To Create Storage Pools for NFS Datastore, complete the following steps:

1. Click "Storage" > "Storage Configuration" > "Storage pools".



2. From the Storage Pools > click "Create" .



3. From the Create Storage Pool menu, manually select 45 SAS drives and 2 SATA Flash drives and add to the pool. Under "Extreme Performance", select RAID10 (4 + 4) and under "Performance", select RAID type as RAID5 (4+1) from the drop-down list. Number of Flash / SAS disks would be automatically populated based on number of disks that you just manually added to the pool.

VSPEX - Create Storage Pool

General Advanced

Storage Pool Parameters

Storage Pool Type: ☒ Pool ☐ RAID Group

☒ Scheduled Auto-Tiering

Storage Pool ID: 3

Storage Pool Name: Pool 3

Extreme Performance

RAID Configuration: RAID1/0 (4+4) Number of Flash Disks: 2

Performance

RAID Configuration: RAID5 (4+1) Number of SAS Disks: 45 (Recommended)

Distribution

Extreme Performance : 183.453 GB (0.75%)
Performance : 24156.343 GB (99.25%)

Disks

☐ Automatic ☐ Use Power Saving Eligible Disks

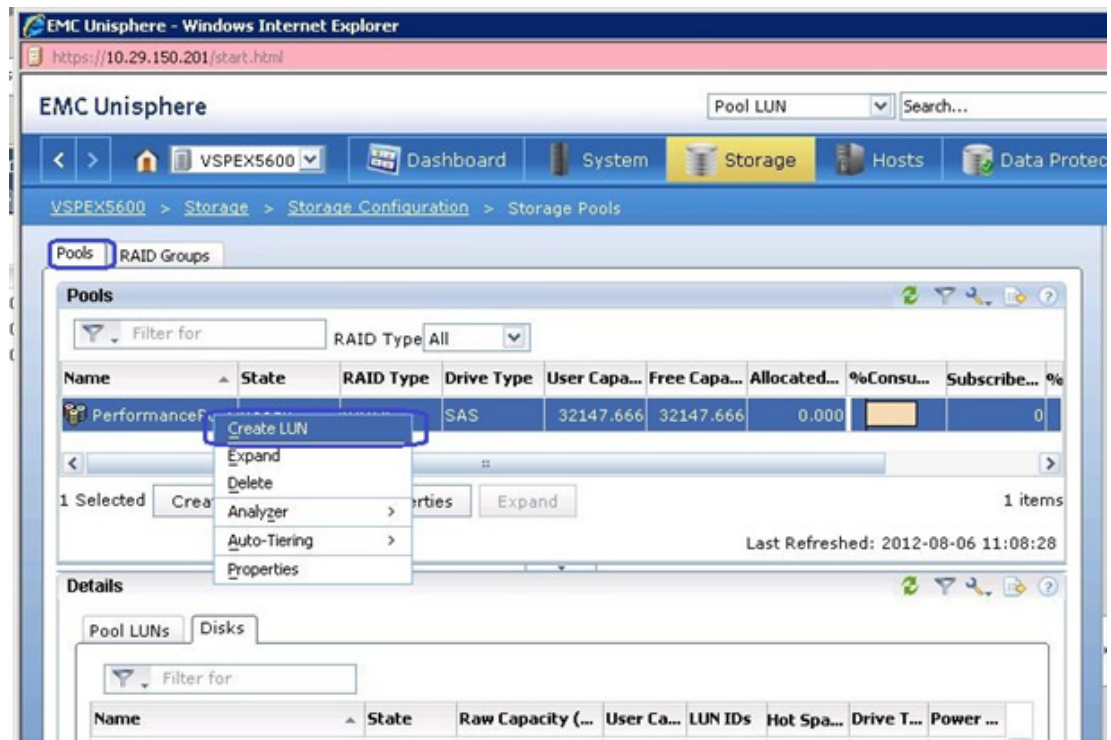
☒ Manual Total Raw Capacity: 24339.79...

Disk	Capacity	Drive Type	Model	State
Bus 0 Enclosure 1 Disk 5	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 6	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 7	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 8	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 9	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 10	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 11	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 12	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 13	536.808 GB	SAS	STE60005 CL...	Unbound
Bus 0 Enclosure 1 Disk 14	536.808 GB	SAS	STE60005 CL...	Unbound

☒ Perform a background verify on the new storage

OK Apply Cancel Help

- Repeat step 3 for required number of pools depending on your architecture.
- To create LUNs from the newly created pools for NFS Datastore; click "Storage" > right-click "PerformancePool" (or "Pool 0", whatever name you have given) > select "Create LUN".



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

VSPEX5600 - Create LUN

General Advanced

Storage Pool Properties

Storage Pool Type: ☒ Pool ☐ RAID Group

RAID Type: Mixed: Multi-tiered with mixed RAID types

Storage Pool for new LUN: Pool 1 New...

Capacity

Available Capacity: 2993.335 GB Consumed Capacity: 16415.750 GB

Oversubscribed By:

LUN Properties

☒ Thin

User Capacity: 800 GB

LUN ID: 49 Number of LUNs to create: 20

LUN Name

☐ Name

Starting ID ?

☒ Automatically assign LUN IDs as LUN Names

Apply Cancel Help

7. Select the pool and Select all the newly created LUNs and click "Add to Storage Group" as shown below. Make sure you select all the LUNs from the pools.

VSPEX5600 > 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		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

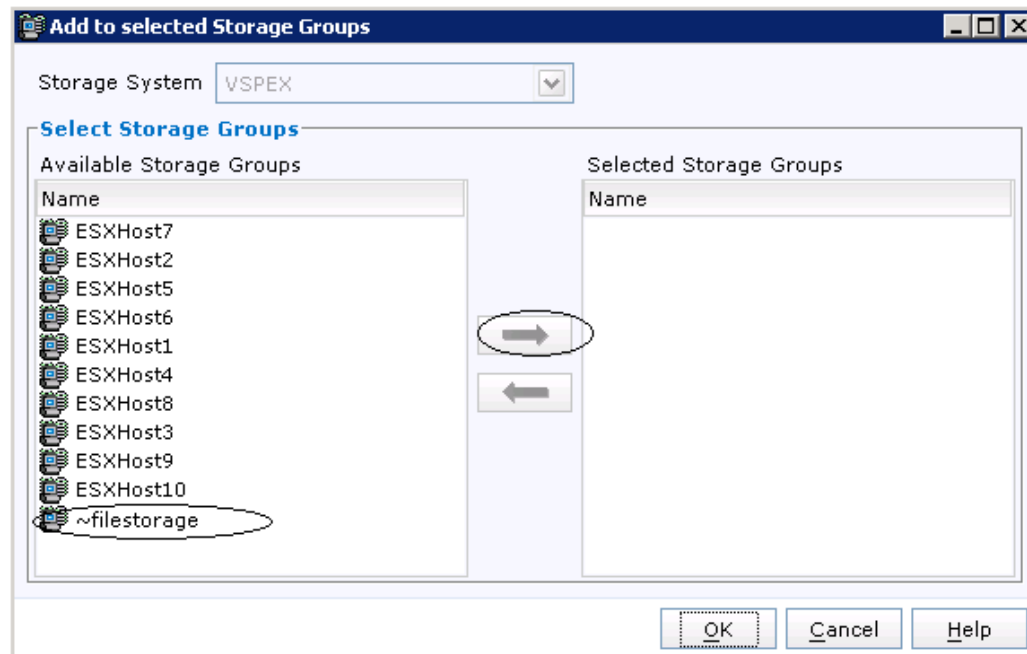
Filter for Usage ALL User LUNs

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	

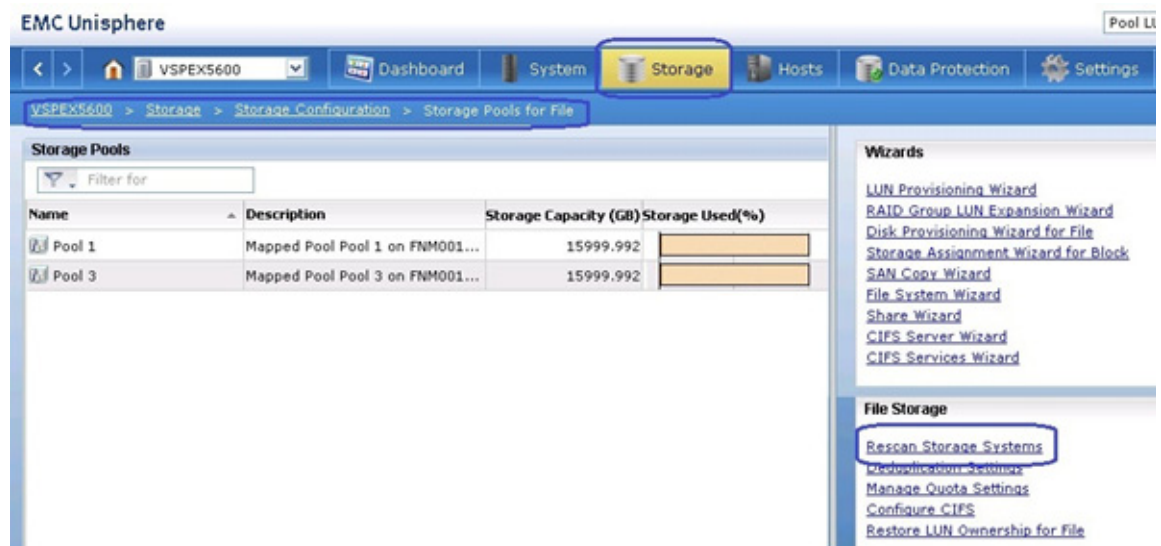
150 Selected Delete Properties Add to Storage Group Filtered: 150 of 150

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

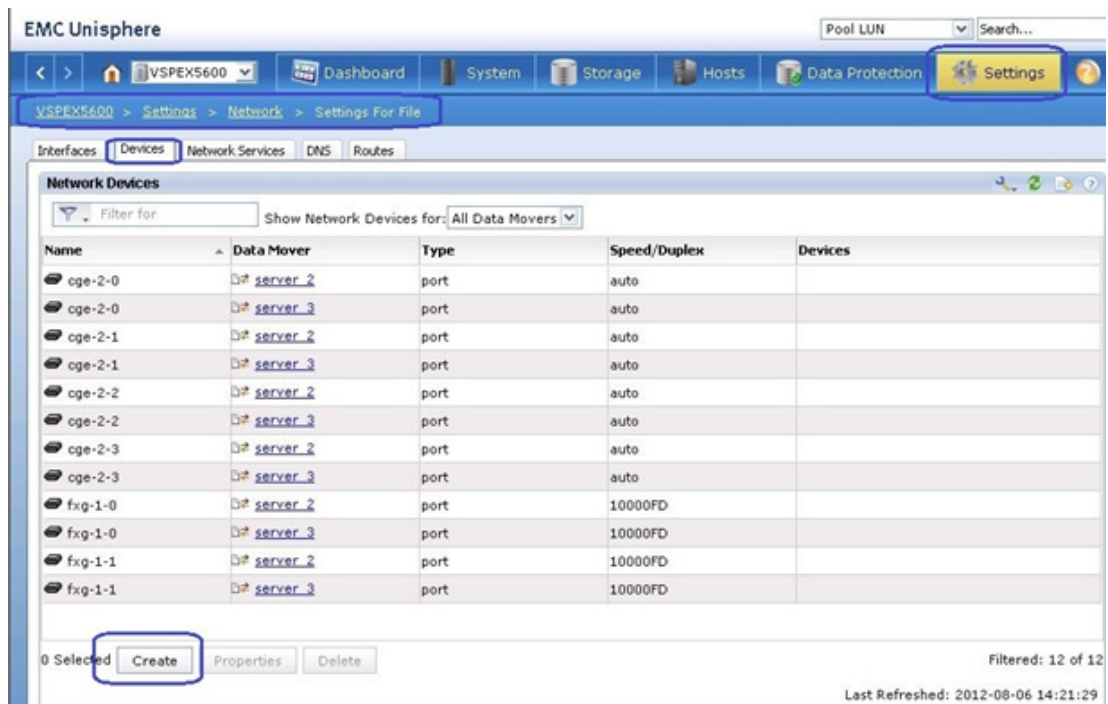
- From the Available Storage Groups, Select "~filestorage" and click the Arrow tab as highlighted below. When "~filestorage" is selected, click "OK".



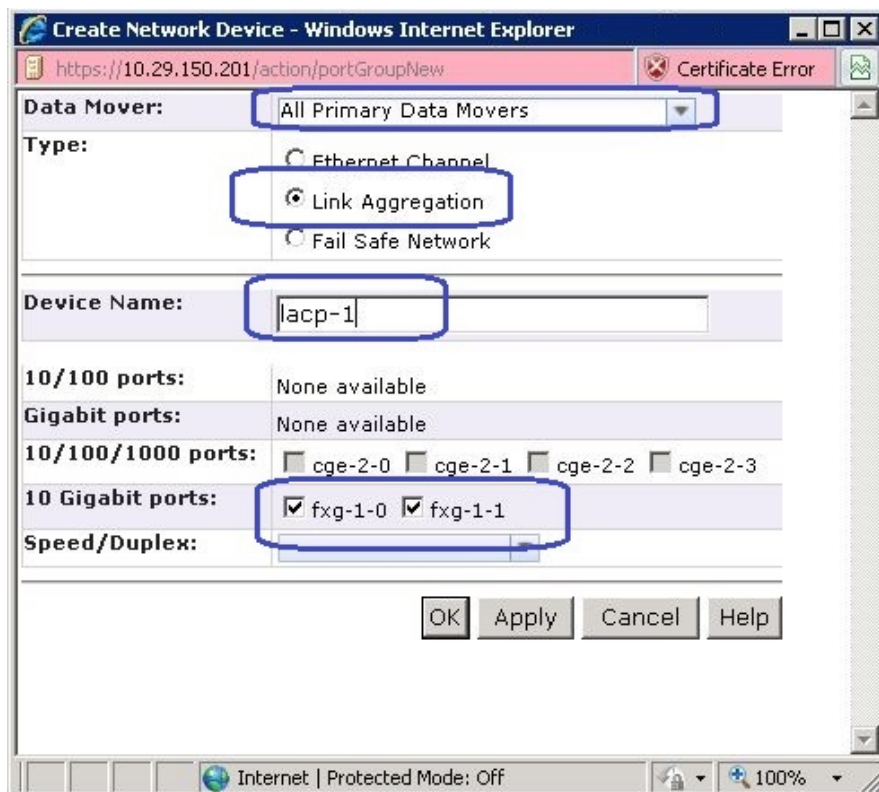
9. Go to "Storage" > "Storage Configuration" > "Storage Pools for Files", and click "Rescan storage systems" as shown below. Rescan will take up to 4 minutes of time. When rescan successfully finishes (track the progress at "Background task for files" page under "System" menu), click "Refresh" and the newly created storage pools will be visible as shown below.



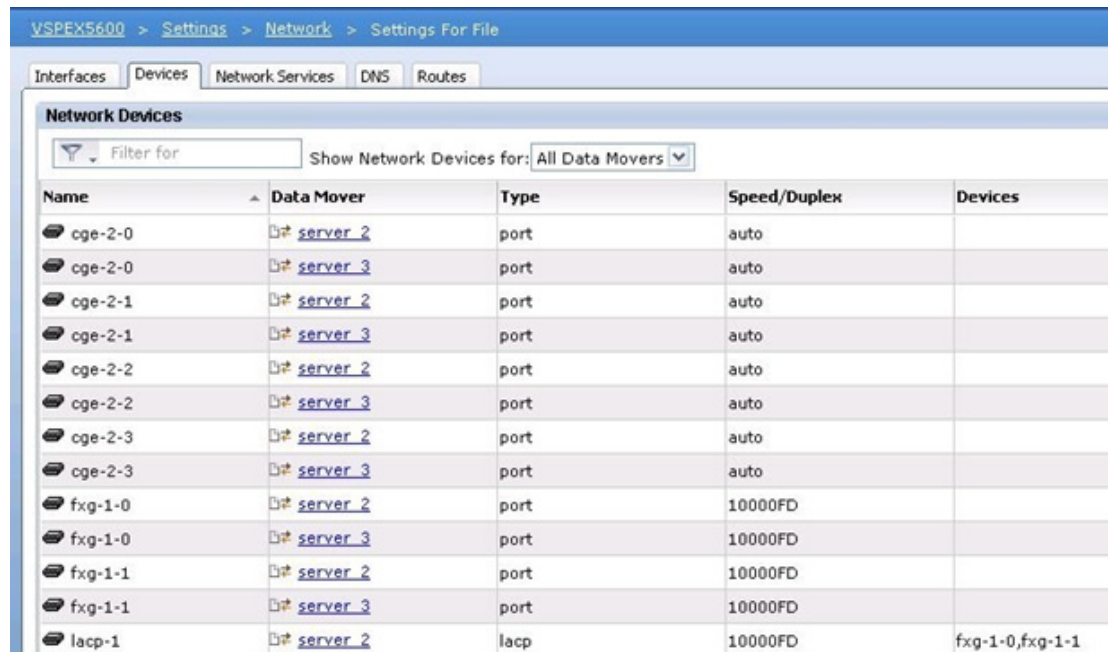
10. The NFS volume is created. Next step is to create highly available network access for the NFS volume. To Create LACP interface. Navigate to "Settings" > "Network" > Select "Settings for File" and click "Devices" tab. Click "Create":



11. In the Data Mover drop-down list, select "All Primary Data Movers", select Type as "Link Aggregation" and Type Device name as "lACP-1". Check the 10 Gigabit ports "fxg-1-0" and "fxg-1-1" as highlighted below. Click "OK" to proceed to the Network Device creation.



The image below shows the creation of LACP Network device name as "lacp-1".



VSPEX5600 > Settings > Network > Settings For File

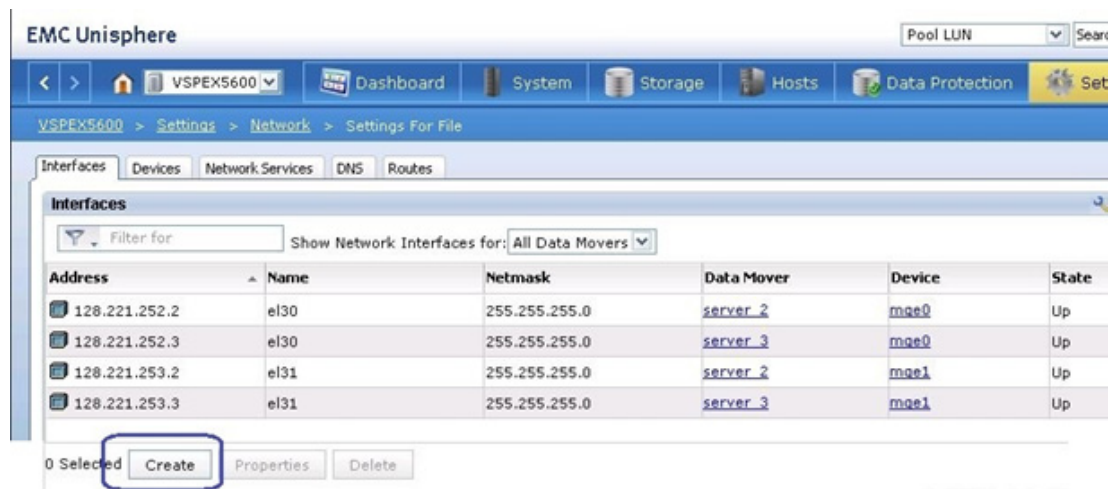
Interfaces Devices Network Services DNS Routes

Network Devices

Filter for Show Network Devices for: All Data Movers

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

12. From the "Settings for File" tab, select "Interfaces" and click "Create".



EMC Unisphere

Pool LUN Search

VSPEX5600 > Settings > Network > Settings For File

Interfaces Devices Network Services DNS Routes

Interfaces

Filter for Show Network Interfaces for: All Data Movers

Address	Name	Netmask	Data Mover	Device	State
128.221.252.2	el30	255.255.255.0	server_2	mge0	Up
128.221.252.3	el30	255.255.255.0	server_3	mge0	Up
128.221.253.2	el31	255.255.255.0	server_2	mge1	Up
128.221.253.3	el31	255.255.255.0	server_3	mge1	Up

0 Selected Create Properties Delete

13. Select 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" and MTU value as "9000" to allow jumbo frames for the lacp interface.

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:	

OK Apply Cancel Help

14. To create file system for NFS data store, navigate to "Storage" > "Storage Configuration" > select "File Systems" and click "Create".
15. From the "Create File System" window, select "Storage Pool" 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", enable "Thin", 7340032 MB (7TB) as Max Capacity, and Select Data Mover as "Server_2" as shown below. Click "OK" to create "NFS-DS-1" File system.

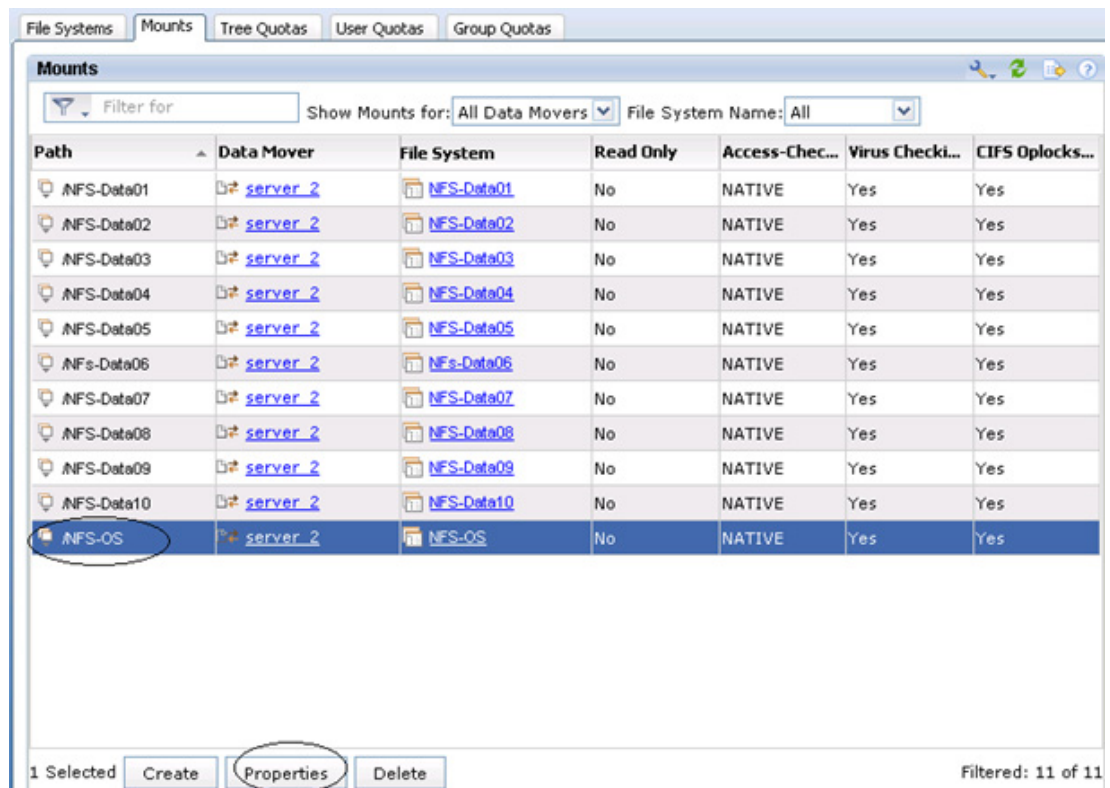
The screenshot shows a web-based configuration interface for creating a file system. The window title is "VSPEX5600 - Create File System - Mozilla Firefox". The URL in the address bar is "https://10.6.117.30/action/filesystemDisplay".

The configuration fields are as follows:

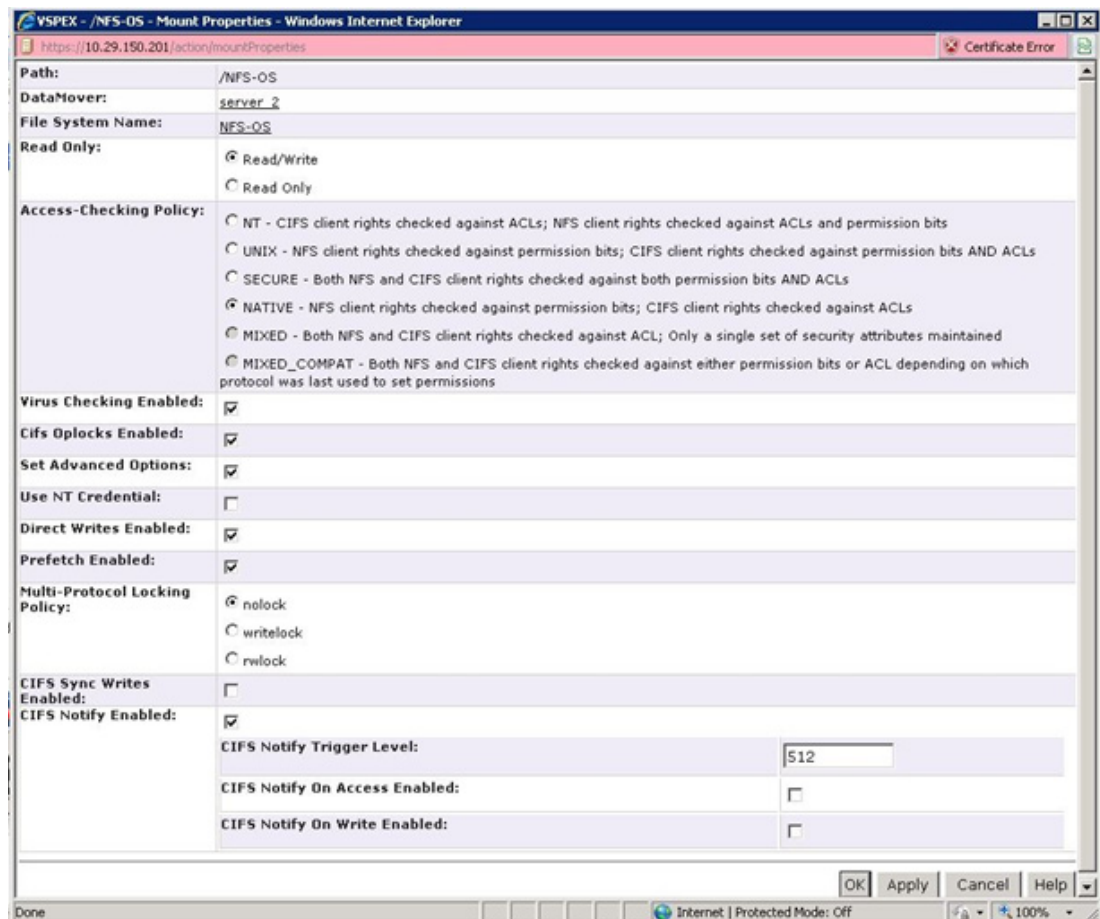
- Create from:** Radio buttons for "Storage Pool" (selected) and "Meta Volume".
- File System Name:** Text input field containing "NFS-DS-1".
- Storage Pool:** Dropdown menu showing "Pool 3 5.6 TB (5898232 MB)".
- Storage Capacity:** Text input field containing "5" and a dropdown menu set to "TB".
- Auto Extend Enabled:** Checkmark is selected.
- Thin Enabled:** Checkmark is selected.
- High Water Mark:** Text input field containing "90". A note indicates: "% (Ranges from 50-99; if left blank defaults to 90)".
- Maximum Capacity (MB):** Text input field containing "7340032". A note indicates: "Required when thin is enabled."
- Slice Volumes:** Checkmark is selected.
- Deduplication Enabled:** Checkmark is not selected.
- VMware VAAI nested clone support:** Checkmark is not selected. A note indicates: "(Must be selected at file system creation time)".
- Data Mover (R/W):** Dropdown menu showing "server_2".
- Mount Point:** Radio buttons for "Default" (selected) and "Custom".

At the bottom right, there are four buttons: "OK", "Apply", "Cancel", and "Help".

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 would 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 "Mounts" tab. Select the path "/NFS-OS" for the file system "NFS-OS" and click "Properties".



18. From the "/NFS-OS" mount properties. Make sure "Read/Write" and "Native" Access policy is selected. Select the "Set Advanced Options" checkbox and check the "Direct Writes Enabled" checkbox as shown below and click "OK".



19. Complete the steps to enable Direct Writes for all the remaining NFS Data file systems.
20. To Create NFS-Exports for all the NFS File systems, click "Storage" > "Shared Folders" > Select "NFS" and click "Create".



21. Select "server-2" in Choose Data Mover drop down list. Select "NFS-OS" in File System drop-down. Specify Path as "/NFS-OS". From the "Root Hosts" and "Access Hosts" field add the IP address of all the ESXi hosts "VMKernel Storage NIC". Separate multiple host vmkernel IP's by : (colon) and Click "OK".

VSPEX - Create NFS Export - Windows Internet Explorer

https://10.29.150.201/action/exportDisplay

Choose Data Mover: server_2

File System: NFS-OS (/NFS-OS)

Path: /NFS-OS

Host Access

Read-only Export: ☐

Read-only Hosts:

Read/Write Hosts:

Root Hosts: 10.10.40.21

Access Hosts: 10.10.40.21

OK Apply Cancel Help

22. The NFS Exports created for the NFS file system are shown below.

EMC Unisphere

Pool LUN Search...

VSPEX5600 > Storage > Shared Folders > NFS

NFS Exports

Filter for Show NFS Exports for: All Data Movers Select a File System: All File Systems


Path	File System	Data Mover
/NFS-OS	NFS-OS	server_2

23. Repeat this step for all the NFS File Systems created on the storage array.
This concludes the NFS storage for VM datastores.

Install and Instantiate VMs from vCenter

This section explains how to deploy virtual machines using vCenter GUI.

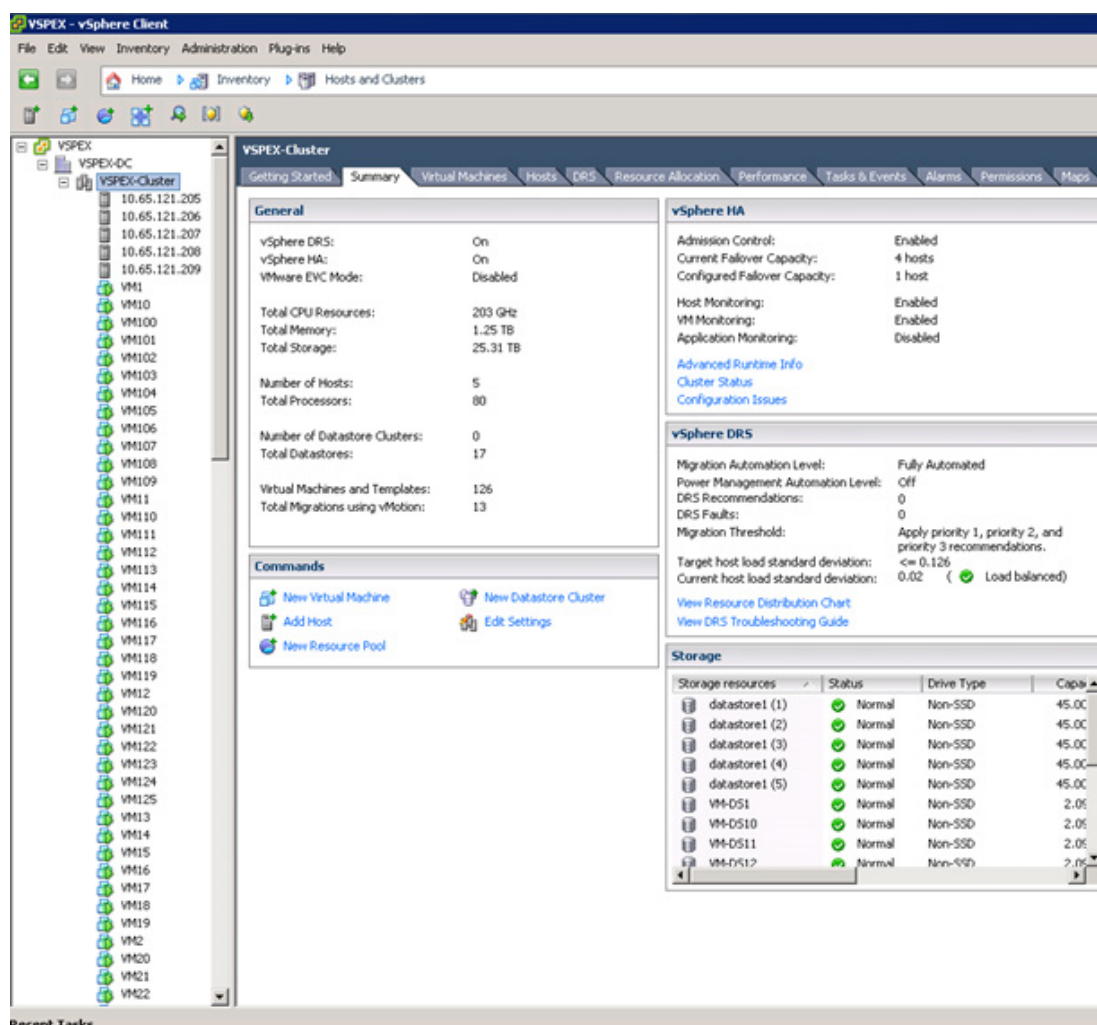
Figure 30 *Template-Based Deployments for Rapid Provisioning*



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

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

Make sure to spread the virtual machines across different VM data-stores to properly load-balance the storage usage. The final snap-shot of VMs in a cluster would look similar to the following image:



Validating the Cisco Solution for EMC VSPEX VMware Architectures

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

Post Install Checklist

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

- Create a test virtual machine that accesses the datastore and is able to do read/write operations. Perform the virtual machine migration (vMotion) to a different host on the cluster.

- Perform storage vMotion from one datastore to another datastore and helps ensure correctness of data.
- During the vMotion of the virtual machine, have a continuous ping to default gateway and make sure that network connectivity is maintained during and after the migration.

Verify the Redundancy of the Solution Components

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

1. Administratively shutdown one of the two server ports connected to the Fabric Extender A. Make sure that connectivity is not affected. Upon administratively enabling the shutdown port, the traffic should be rebalanced. This can be validated by clearing interface counters and showing the counters after forwarding some data from virtual machines on the Nexus switches.
2. Administratively shutdown both server ports connected to Fabric Extender A. ESXi hosts should be able to use fabric B in this case.
3. Administratively shutdown one of the two data links connected to the storage array from FI. Make sure that storage is still available from all the ESXi hosts. Upon administratively enabling the shutdown port, the traffic should be rebalanced. Repeat this step for each link connected to the Storage Processors one after another.
4. Reboot one of the two Fabric Interconnects while storage and network access from the servers are going on. The switch reboot should not affect the operations of storage and network access from the VMs. Upon rebooting the FI, the network access load should be rebalanced across the two fabrics.
5. Reboot the active storage processor of the 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 put it out of the maintenance mode. This should rebalance the VM distribution across the cluster.

Cisco Validation Test Profile

The "VDbench" testing tool was used with Windows 2012 server to test scaling of the solution in Cisco labs. The table below lists the test profile information.

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 16 provides the details of the components used in the CVD for 300 virtual machines configuration.

Table 16 *Component Description*

Description	Part #
5 x UCS B200 M4 blade server / UCS C220 M4 rack server	
CPU for B220 M4/C220M4 Rack servers (2 per server)	
Memory for B220 M4/C220M4 Rack servers (4 per server)	
Cisco UCS 1225 VIC adapter (1 per server)	UCSC-PCIE-CSC-02
UCS 6248UP Fabric Interconnects (2)	UCS-FI-6248UP
10 Gbps SFP+ multifiber mode	SFP-10G-S

For more information about the part numbers and options available for customization, please see Cisco C220 M4 server specsheet at

<http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m4-sff-spec-sheet.pdf>

Appendix—Customer Configuration Data Sheet

Before you start the configuration, gather some customer-specific network and host configuration information. Table 17 provides information on assembling the required network and host address, numbering, and naming information. This worksheet can also be used as a "leave behind" document for future reference.

Table 17 *Common Server Information*

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

Table 18 *ESXi Server Information*

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

Table 19 *Storage Array Information*

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

Table 20 *Network Infrastructure Information*

Description	IP	Subnet Mask	Default Gateway
UCSM Virtual IP address			
UCS Fabric Interconnect A address			
UCS Fabric Interconnect B address			
N9K A management IP address			
N9K B management IP address			
N1kv management IP addresss			
MDS 9148S Switch A management IP address			
MDS 9148S Switch B management IP address			

Table 21 *VLAN Information*

Name	Network Purpose	VLAN ID	Subnet
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 22 *VSAN Information*

Name	Network Purpose	VSAN ID	FCoE VLAN ID
Storage	Storage access		

Table 23 *Service Accounts*

Account	Purpose	Password (optional, secure appropriately)
Admin	UCSM administrator	
Admin	N9K switches administrator	
Admin	N1kv switch administrator	
Admin	MDS9148S 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

Cisco UCSM 2.2(3) configuration guides:

CLI:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/cli/config/guide/2-2/b_UCSM_CLI_Configuration_Guide_2_2.html

GUI:

http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/sw/gui/config/guide/2-2/b_UCSM_GUI_Configuration_Guide_2_2.html

VMware vSphere:

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

VMware vSphere 5.5 documentation:

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

EMC VNX5xxx series resources:

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

Microsoft SQL Server 2012 R2 installation guide

[http://msdn.microsoft.com/en-us/library/bb500469\(v=sql.110\).aspx](http://msdn.microsoft.com/en-us/library/bb500469(v=sql.110).aspx)