



# Cisco Solution for EMC VSPEX VMware vSphere 5.0 Architectures

Design for 50, 100 and 125 Virtual Machines


Last Updated: October 24, 2013



Cisco  
Validated  
Design



Building Architectures to Solve Business Problems



## About the Authors



Mehul Bhatt

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

Mehul Bhatt has over 12 years of Experience in virtually all layers of computer networking. His focus area includes Unified Compute Systems, network and server virtualization design. Prior to joining Cisco Technical Marketing team, Mehul was Technical Lead at Cisco, Nuova systems and Bluecoat systems. Mehul holds a Masters degree in computer systems engineering and holds various Cisco career certifications.



Hardik Patel

### **Hardik Patel, Virtualization System Engineer, Server Access Virtualization Business Unit, Cisco Systems**

Hardik Patel has over 9 years of experience with server virtualization and core application in the virtual environment with area of focus in design and implementation of systems and virtualization, manage and administration, UCS, storage and network configurations. Hardik holds Masters degree in Computer Science with various career oriented certification in virtualization, network and Microsoft.



VijayKumar D

### **VijayKumar D, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems**

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

# Acknowledgements

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

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



# About Cisco Validated Design (CVD) Program

The CVD program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information visit [www.cisco.com/go/designzone](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)

© 2012 Cisco Systems, Inc. All rights reserved



# Cisco Solution for EMC VSPEX VMware vSphere 5.0 Architectures

---

## Executive Summary

Cisco solution for the EMC VSPEX is a pre-validated and modular architecture built with proven best-of-breed technologies to create and complete an end-to-end virtualization solution. The end-to-end solutions enable you to make an informed decision while choosing the hypervisor, compute, storage and networking layers. VSPEX eliminates the server virtualization planning and configuration burdens. The VSPEX infrastructures accelerate your IT Transformation by enabling faster deployments, greater flexibility of choice, efficiency, and lower risk. This Cisco Validated Design document focuses on the VMware architecture for 50, 100 and 125 virtual machines with Cisco solution for the EMC VSPEX.

## Introduction

Virtualization is a key and critical strategic deployment model for reducing the Total Cost of Ownership (TCO) and achieving better utilization of the platform components like hardware, software, network and storage. However, choosing an appropriate platform for virtualization can be challenging. Virtualization platforms should be flexible, reliable, and cost effective to facilitate the deployment of various enterprise applications. In a virtualization platform to utilize compute, network, and storage resources effectively, the ability to slice and dice the underlying platform is essential to size to the application requirements. The Cisco solution for the EMC VSPEX provides a very simplistic yet fully integrated and validated infrastructure to deploy Virtual Machines in various sizes to suit various application needs.

## Target Audience

The reader of this document is expected to have the necessary training and background to install and configure VMware vSphere 5.0, EMC VNXe series, EMC VNX5300, Cisco Nexus 3048 switch, Cisco Nexus 5548UP switch, Cisco Nexus 1000v switch, and Cisco Unified Computing (UCS) C220 M3 rack servers. External references are provided wherever applicable and it is recommended that the reader be familiar with these documents.

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



---

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

Copyright 2012 Cisco Systems, Inc. All rights reserved.

## Purpose of this Guide

This document describes the steps required to deploy and configure the Cisco solution for the EMC VSPEX for VMware architecture. The document covers three types of VMware architectures:

- VMware vSphere 5.0 for 50 Virtual Machines
- VMware vSphere 5.0 for 100 virtual machines
- VMware vSphere 5.0 for 125 virtual machines

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

## Business Needs

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

For more detailed information on server capacity, network interface, and storage configuration, see the EMC VSPEX Server Virtualization Solution VMware vSphere 5 for 125 Virtual Machines Enabled by VMware vSphere 5.0 and EMC VNX5300—Reference Architecture and associated documentation.

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

- Provide an end-to-end virtualization solution to utilize the capabilities of the unified infrastructure components.
- Provide a Cisco VSPEX for VMware ITaaS solution for efficiently virtualizing 50, 100 or 125 virtual machines for varied customer use cases.
- Show implementation progression of VMware vCenter 5.0 design and the results.
- Provide a reliable, flexible and scalable reference design.

## Solution Overview

The Cisco solution for EMC VSPEX using VMware vSphere 5.0 provides an end-to-end architecture with Cisco, EMC, VMware, and Microsoft technologies that demonstrate support for up to 50, 100 and 125 generic virtual machines and provide high availability and server redundancy.

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

- Cisco C-series Unified Computing System servers
- Cisco Nexus 5000 series or 3000 series switches depending on the scale of the solution
- Cisco Nexus 1000v virtual switch
- Cisco virtual Distributed Switch across multiple VMware ESXi hypervisors
- Cisco virtual Port Channels for network load balancing and high availability

- EMC VNXe3150, VNXe3300 or VNX5300 storage components as per the scale needs
- VMware vCenter 5
- Microsoft SQL database
- VMware DRS
- VMware HA

The solution is designed to host scalable, and mixed application workloads. The scope of this CVD is limited to the Cisco solution for EMC VSPEX VMware solutions for 50, 100 and 125 virtual machines only.

## Technology Overview

### Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation data center platform that unites compute, network, and storage access. The platform, optimized for virtual environments, is designed using open industry-standard technologies and aims to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. It is an integrated, scalable, multi chassis platform in which all resources participate in a unified management domain.

The main components of Cisco Unified Computing System are:

- **Computing**—The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500/5600 Series Processors. Selected Cisco UCS blade 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, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.
- **Management**—The system uniquely integrates all system components which enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver:

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

## Cisco C220 M3 Rack-Mount Servers

Building on the success of the Cisco UCS C220 M3 Rack-Mount Servers, the enterprise-class Cisco UCS C220 M3 server further extends the capabilities of the Cisco Unified Computing System portfolio in a 1-rack-unit (1RU) form factor. And with the addition of the Intel® Xeon® processor E5-2600 product family, it delivers significant performance and efficiency gains. [Figure 1](#) shows the Cisco UCS C220 M3 rack server.

**Figure 1** *Cisco UCS C220 M3 Rack Server*



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

## I/O Adapters


The Cisco UCS Rack-Mount Server has various Converged Network Adapters (CNA) options. The Cisco UCS P81E Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

This Cisco UCS P81E VIC is unique to the Cisco UCS Rack-Mount Server system. This mezzanine card adapter is designed around a custom ASIC that is specifically intended for virtualized systems. As is the case with the other Cisco CNAs, the Cisco UCS P81E VIC encapsulates fibre channel traffic within the 10-GE packets for delivery to the Ethernet network.

UCS P81E VIC provides the capability to create multiple VNICs (up to 128) on the CNA. This allows complete I/O configurations to be provisioned in virtualized or non-virtualized environments using just-in-time provisioning, providing tremendous system flexibility and allowing consolidation of multiple physical adapters.

System security and manageability is improved by providing visibility and portability of network policies and security all the way to the virtual machines. Additional P81E features like VN-Link technology and pass-through switching, minimize implementation overhead and complexity. [Figure 2](#) shows the Cisco UCS P81E VIC.

**Figure 2** *Cisco UCS P81e VIC*



## Cisco Nexus 5548UP Switch

The Cisco Nexus 5548UP is a 1RU 1 Gigabit and 10 Gigabit Ethernet switch offering up to 960 gigabits per second throughput and scaling up to 48 ports. It offers 32 1/10 Gigabit Ethernet fixed enhanced Small Form-Factor Pluggable (SFP+) Ethernet/FCoE or 1/2/4/8-Gbps native FC unified ports and three expansion slots. These slots have a combination of Ethernet/FCoE and native FC ports. The Cisco Nexus 5548UP switch is shown in [Figure 3](#).

**Figure 3** *Cisco Nexus 5548UP switch*



## Cisco Nexus 3048 Switch

The Cisco Nexus® 3048 Switch is a line-rate Gigabit Ethernet top-of-rack (ToR) switch and is part of the Cisco Nexus 3000 Series Switches portfolio. The Cisco Nexus 3048, with its compact one-rack-unit (1RU) form factor and integrated Layer 2 and 3 switching, complements the existing Cisco Nexus family of switches. This switch runs the industry-leading Cisco® NX-OS Software operating system, providing customers with robust features and functions that are deployed in thousands of data centers worldwide. The Cisco Nexus 3048 switch is shown in [Figure 4](#).

**Figure 4** *Cisco Nexus 3048 Switch*



## Cisco Nexus 1000v Virtual Switch

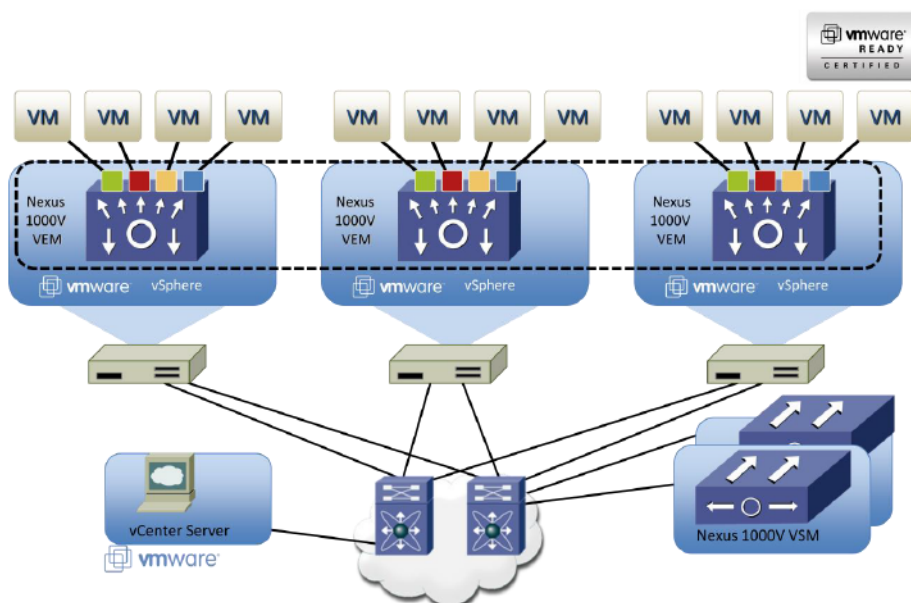
Nexus 1000v is a virtual Ethernet switch with two components:

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

Virtual Ethernet Modules across multiple ESXi hosts form a virtual Distributed Switch (vDS). Using the Cisco vDS VMware plug-in, the Virtual Interface Card (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 the 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 the rest of the Ethernet fabric in the Data Center.

The Cisco Nexus 1000v vDS architecture is shown in [Figure 5](#).

**Figure 5** *Cisco Nexus 1000v Switch*



## VMware vSphere 5.0

VMware vSphere 5.0 is a next-generation virtualization solution from VMware which builds upon ESXi 4 and provides greater levels of scalability, security, and availability to virtualized environments. vSphere 5.0 offers improvements in performance and utilization of CPU, memory, and I/O. It also offers users the option to assign up to thirty two virtual CPU to a virtual machine—giving system administrators more flexibility in their virtual server farms as processor-intensive workloads continue to increase.

The vSphere 5.0 provides the VMware vCenter Server that allows system administrators to manage their ESXi hosts and virtual machines on a centralized management platform. With the Cisco Fabric Interconnects Switch integrated into the vCenter Server, deploying and administering virtual machines is similar to deploying and administering physical servers. Network administrators can continue to own the responsibility for configuring and monitoring network resources for virtualized servers as they did with physical servers. System administrators can continue to “plug-in” their virtual machines into the network ports that have Layer 2 configurations, port access and security policies, monitoring features, and so on, that have been pre-defined by the network administrators; in the same way they need to plug in their physical servers to a previously-configured access switch. In this virtualized environment, the network port configuration/policies move with the virtual machines when the virtual machines are migrated to different server hardware.

## EMC Storage Technologies and Benefits

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

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

The VNXe series is purpose-built for the IT manager in smaller environments and the VNX series is designed to meet the high-performance, high-scalability requirements of midsize and large enterprises. The EMC VNXe and VNX storage arrays are multi-protocol platform that can support the iSCSI, NFS, and CIFS protocols depending on the customer’s specific needs. The solution was validated using NFS for data storage.

VNXe series storage arrays have following customer benefits:

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

## Software Suites

The following are the available EMC software suites:

- Remote Protection Suite—Protects data against localized failures, outages, and disasters.



- Application Protection Suite—Automates application copies and proves compliance.
- Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

## Software Packs

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

This is the available EMC protection software pack.

## EMC Avamar

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

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

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

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

# Architectural Overview

This CVD discusses the deployment model for the following three VMware virtualization solutions:

- VMware solution for 50 virtual machines
- VMware solution for 100 virtual machines
- VMware solution for 125 virtual machines

[Table 1](#) lists the mix of hardware components, their quantities and software components used for different VMware solutions:

**Table 1** *Hardware and software components for various solutions*

Components	VMware 50 Virtual Machines	VMware 100 Virtual Machines	VMware 125 Virtual Machines
Servers	Three Cisco C220 M3 servers	Four Cisco C220 M3 servers	Five Cisco C220 M3 servers
Adapters	One Broadcom NetXtreme II 5706 per server	One Cisco UCS P81E VIC per server	One Cisco UCS P81E VIC per server

**Table 1** Hardware and software components for various solutions

Components	VMware 50 Virtual Machines	VMware 100 Virtual Machines	VMware 125 Virtual Machines
Network Switches	Two Cisco Nexus 3048 switches	Two Cisco Nexus 5548UP switches	Two Cisco Nexus 5548UP switches
Virtual Switch	One Cisco Nexus 1000v	One Cisco Nexus 1000v	One Cisco Nexus 1000v
Storage	EMC VNXe3150	EMC VNXe3300	EMC VNX5300
Network Speed	1 GE	10 GE	10 GE
Hypervisor	VMware ESXi 5.0	VMware ESXi 5.0	VMware ESXi 5.0

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

**Table 2** Hardware and software components of VMware architectures

Vendor	Name	Version	Description
Cisco	C220 M3 servers	1.4(4a).1 - CIMC C220M3.1.4.4c.0 - BIOS	Cisco C220 M3 rack servers
Cisco	Cisco Nexus 5548UP Switches	5.1(3)N1(1a)	Cisco Nexus 5000 series switches running NX-OS
Cisco	Cisco Nexus 3048 Switches	5.0(3)U2(2b)	Nexus 3000 series switches running NX-OS
Cisco	Cisco Nexus 1000v switch	4.2(1)SV1(5.1a)	Cisco Nexus 1000 virtual switch
EMC	EMC VNXe3150	2.2.0.16150	EMC VNXe storage array
EMC	EMC VNXe3300	2.2.0.16150	EMC VNXe storage array
EMC	EMC VNX5300	7.0.50-2	EMC VNX storage array
EMC	EMC Avamar	6.0.0-592	EMC data backup software
EMC	Data Domain OS	5.1.0.9-282511	EMC data domain operating system
VMware	ESXi 5.0	5.0 build 623860	VMware Hypervisor
VMware	vCenter Server	5.0 build 455964	VMware management

**Table 2** *Hardware and software components of VMware architectures*

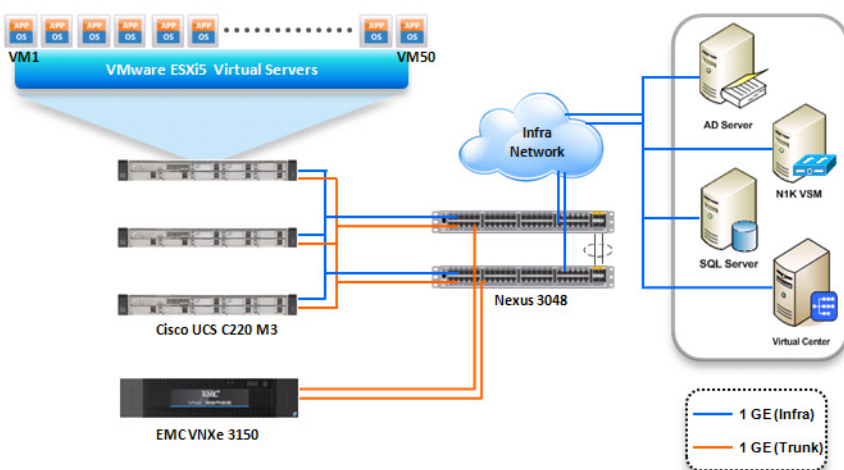
Vendor	Name	Version	Description
Microsoft	Microsoft Windows Server 2008 R2	2008 R2 SP1	Operating system to host vCenter server
Microsoft	Microsoft SQL server	2008 R2	Database server SQL R2 Enterprise edition for vCenter

Table 3 outlines the C220 M3 server configuration details (per server basis) across all the VMware architectures.

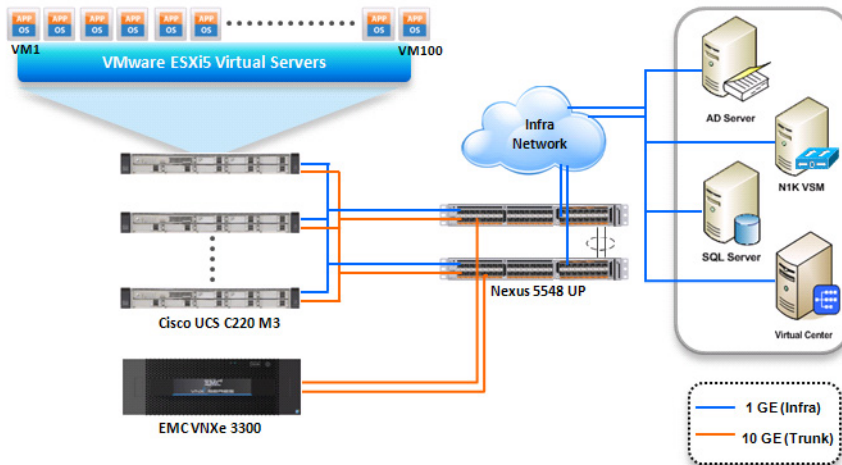
**Table 3** *Server configuration details*

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

All the three reference architectures 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. A new VM hosting the Nexus 1000v VMS service would be deployed as part of the Cisco solution for the EMC VSPEX architecture. Figure 6, Figure 7, Figure 8 illustrate high-level solution architecture for 50, 100 and 125 virtual machines.

**Figure 6** *Reference Architecture for 50 Virtual Machines*

**Figure 7** *Reference Architecture for 100 Virtual Machines*



**Figure 8** *Reference Architecture for 125 Virtual Machines*

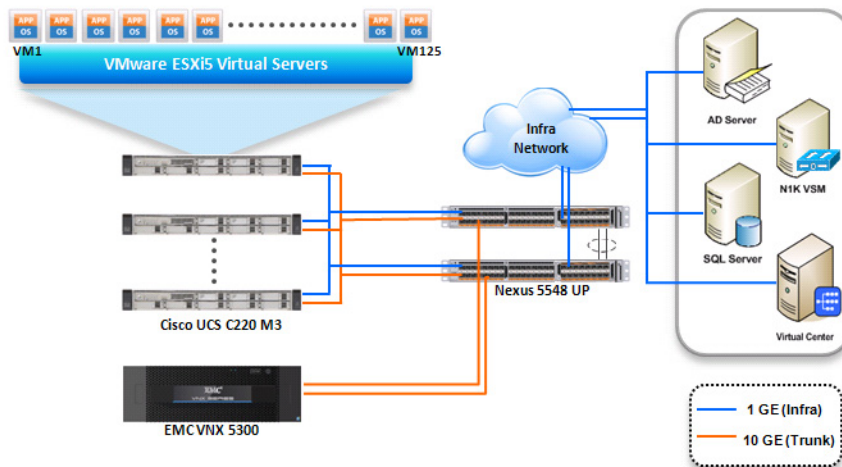


Figure 6, Figure 7, Figure 8 illustrate that the high-level design points of VMware architectures are as follows:

- Only Ethernet is used as network layer 2 media to access storage as well as TCP/IP network
- Infrastructure network is on a separate 1GE network
- Network redundancy is built in by providing two switches, two storage controllers and redundant connectivity for data, storage and infrastructure networking.

This design does not recommend or require any specific layout of infrastructure network. The VMware vCenter server and the Cisco Nexus 1000v VSM virtual machines are hosted on infrastructure network. However, design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

ESXi 5.0 is used as hypervisor operating system on each server and is installed on local hard drives. Typical load is 25 virtual machines per server.

## Memory Configuration Guidelines

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

### ESXi/ESXi Memory Management Concepts

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


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

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

### Virtual Machine Memory Concepts

The [Figure 9](#) illustrates the use of memory settings parameters in the virtual machine.

**Figure 9** Virtual Machine Memory Settings



The VMware 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. VMware vSphere allocates only guest operating system memory on demand.

- **Swappable**—Virtual machine memory can be reclaimed by the balloon driver or by VMware vSphere swapping. Ballooning occurs before VMware vSphere swapping. If this memory is in use by the virtual machine (that is, touched and in use), the balloon driver causes the guest operating system to swap.

## Allocating Memory to Virtual Machines

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

**Table 4** *Resources for a single virtual machine*

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

Following are the recommended best practices:

- **Account for memory overhead**—Virtual machines require memory beyond the amount allocated, and this memory overhead is per-virtual machine. Memory overhead includes space reserved for virtual machine devices, depending on applications and internal data structures. The amount of overhead required depends on the number of vCPUs, configured memory, and whether the guest operating system is 32-bit or 64-bit. As an example, a running virtual machine with one virtual CPU and two GB of memory may consume about 100 MB of memory overhead, where a virtual machine with two virtual CPUs and 32 GB of memory may consume approximately 500 MB of memory overhead. This memory overhead is in addition to the memory allocated to the virtual machine and must be available on the ESXi host.
- **“Right-size” memory allocations**—Over-allocating memory to virtual machines can waste memory unnecessarily, but it can also increase the amount of memory overhead required to run the virtual machine, thus reducing the overall memory available for other virtual machines. Fine-tuning the memory for a virtual machine is done easily and quickly by adjusting the virtual machine properties. In most cases, hot-adding of memory is supported and can provide instant access to the additional memory if needed.
- **Intelligently overcommit**—Memory management features in VMware vSphere allow for over commitment of physical resources without severely impacting performance. Many workloads can participate in this type of resource sharing while continuing to provide the responsiveness users require of the application. When looking to scale beyond the underlying physical resources, consider the following:

- Establish a baseline before overcommitted. 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 VMware 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 the services have the resources they require while still allowing high consolidation of other virtual machines.

As with overcommitted CPU resources, proactive monitoring is a requirement. [Table 5](#) lists counters that can be monitored to avoid performance issues resulting from overcommitted memory.

**Table 5** *ESXitop Memory Counters*

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

## Storage Guidelines


VSPEX architecture for VMware 50, 100, and 125 virtual machine scale uses NFS to access storage arrays. This simplifies the design and implementation for the small to medium level businesses. VMware vSphere provides many features that take advantage of EMC storage technologies such as VNX VAAI plugin for NFS storage and storage replication. Features such as VMware vMotion, VMware HA, and VMware Distributed Resource Scheduler (DRS) use these storage technologies to provide high availability, resource balancing, and uninterrupted workload migration.

## Virtual Server Configuration

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

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

**Figure 10** *VMware Storage Virtualization Stack*



## Storage Protocol Capabilities

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

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

For more information, see the VMware white paper Comparison of Storage Protocol Performance in VMware vSphere 5: [http://www.vmware.com/files/pdf/perf\\_vsphere\\_storage\\_protocols.pdf](http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf)

## Storage Best Practices

Following are the VMware 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 can be in the form of dual adapters connected to separate fabric switches, or a set of teamed network interface cards for NFS.



- **Partition alignment**—Partition misalignment can lead to severe performance degradation due to I/O operations having to cross track boundaries. Partition alignment is important both at the NFS level as well as within the guest operating system. Use the VMware vSphere Client when creating NFS datastores to be sure they are created aligned. When formatting volumes within the guest, Windows 2008 aligns NTFS partitions on a 1024KB offset by default.
- **Use shared storage**—In a VMware 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 NFS volume. Over-subscription of the I/O resources can go unnoticed at first and slowly begin to degrade performance if not monitored proactively.

## VSPEX VMware Memory Virtualization

VMware vSphere 5.0 has a number of advanced features that help to maximize performance and overall resources utilization. This section describes the performance benefits of some of these features for the VSPEX deployment.

### Memory Compression

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

For more details about VMware vSphere memory management concepts, see the VMware vSphere Resource Management Guide at: [http://www.VMware.com/files/pdf/mem\\_mgmt\\_perf\\_Vsphere5.pdf](http://www.VMware.com/files/pdf/mem_mgmt_perf_Vsphere5.pdf)

### Virtual Networking

The Cisco Nexus 1000v collapses virtual and physical networking into a single infrastructure. The Nexus 1000v allows data center administrators to provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare metal network traffic within a unified infrastructure.

The Nexus 1000v software extends Cisco data-center networking technology to the virtual machine with the following capabilities:

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

## Benefits

- Simplified operations—Seamless virtual networking infrastructure similar to Cisco Nexus 5000 / 7000 series CLI interface
- Improved network security—Contains VLAN proliferation
- Optimized network utilization—Reduces broadcast domains
- Reduced network complexity—Separation of network and server administrator's domain by providing port-profiles by name

## Virtual Networking Best Practices

Following are the VMware vSphere networking best practices:

- Separate virtual machine and infrastructure traffic—Keep virtual machine and VMkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation.
- Use NIC Teaming—Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches. Teaming provides redundancy against NIC failure and, if connected to separate physical switches, against switch failures. NIC teaming does not necessarily provide higher throughput.
- Enable PortFast on ESX/ESXi host uplinks—Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. To prevent this situation, set the switch ports connected to ESXi/ESXi hosts to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESXi/ESXi hosts from affecting the STP topology. Loops are not possible in virtual switches.
- MAC pinning—MAC pinning based load balancing and high availability is recommended over the virtual Port-Channel (vPC) based load balancing because of the simplicity in the MAC pinning approach. MAC pinning provides more static allocation of virtual machines' vNICs on the physical uplink, however, given 25 virtual machines per server, there will be a fair distribution of network load across the Virtual Machines.
- Converged Network and Storage I/O with 10Gbps Ethernet—Consolidating storage and network traffic can provide simplified cabling and management over maintaining separate switching infrastructures.

## VMware vSphere Performance

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

**Table 6** *VMware vSphere Performance*

ESXitop Metric	Description	Implication
NUMA Support	ESX/ESXi uses a NUMA load-balancer to assign a home node to a virtual machine. Because memory for the virtual machine is allocated from the home node, memory access is local and provides the best performance possible. Even applications that do not directly support NUMA benefit from this feature.	See The CPU Scheduler in VMware ESXi 5: <a href="http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf">http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf</a>
Transparent page sharing	Virtual machines running similar operating systems and applications typically have identical sets of memory content. Page sharing allows the hypervisor to reclaim the redundant copies and keep only one copy, which frees up the total host memory consumption. If most of your application virtual machines run the same operating system and application binaries then total memory usage can be reduced to increase consolidation ratios.	See Understanding Memory Resource Management in VMware ESXi 5.0: <a href="http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf">http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf</a>
Memory ballooning	By using a balloon driver loaded in the guest operating system, the hypervisor can reclaim host physical memory if memory resources are under contention. This is done with little to no impact to the performance of the application.	See Understanding Memory Resource Management in VMware ESXi 5.0: <a href="http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf">http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf</a>

**Table 6 VMware vSphere Performance**

ESXitop Metric	Description	Implication
Memory compression	Before a virtual machine resorts to host swapping, due to memory over commitment the pages elected to be swapped attempt to be compressed. If the pages can be compressed and stored in a compression cache, located in main memory, the next access to the page causes a page decompression as opposed to a disk swap out operation, which can be an order of magnitude faster.	See Understanding Memory Resource Management in VMware ESXi 5.0: <a href="http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf">http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf</a>
Large memory page support	An application that can benefit from large pages on native systems, such as MS SQL, can potentially achieve a similar performance improvement on a virtual machine backed with large memory pages. Enabling large pages increases the memory page size from 4KB to 2MB.	See Performance Best Practices for VMware vSphere 5.0: <a href="http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf">http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf</a> and see Performance and Scalability of Microsoft SQL Server on VMware vSphere 4: <a href="http://www.vmware.com/files/pdf/perf_vsphere_sql_scalability.pdf">http://www.vmware.com/files/pdf/perf_vsphere_sql_scalability.pdf</a>

## Physical and Virtual CPUs

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

## Virtual SMP

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

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

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

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

[http://www.vmware.com/pdf/Perf\\_Best\\_Practices\\_vSphere5.0.pdf](http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf). For example, in VMware ESXi 5, the relaxed co-scheduling algorithm was refined so that scheduling constraints due to co-scheduling requirements are further reduced. These improvements resulted in better linear scalability and performance of the SMP virtual machines.

## Overcommitment

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

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

**Table 7**      **List of Counters**

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

## Hyper-Threading

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

## Non-Uniform Memory Access (NUMA)

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

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

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

## VSPEX VMware Storage Virtualization

Disk provisioning on the EMC VNXe series is simplified through the use of wizards, so that administrators need not choose the disks that belong to the given storage pool. The wizard will automatically choose the available disk, regardless of where the disk physically resides in the array. On the other hand, disk provisioning on the EMC VNX series requires administrators to choose disks for each of the storage pools.

### Storage Layout

This section illustrates the physical disk layouts on the EMC VNXe and VNX storage arrays.

Figure 11 shows storage architecture for 50 virtual machines on VNXe3150:

**Figure 11**      **Storage Architecture for 50 Virtual Machines on EMC VNXe3150**




Figure 12 shows storage architecture for 100 virtual machines on VNXe3300:

**Figure 12 Storage Architecture for 100 Virtual Machines on EMC VNXe3300**

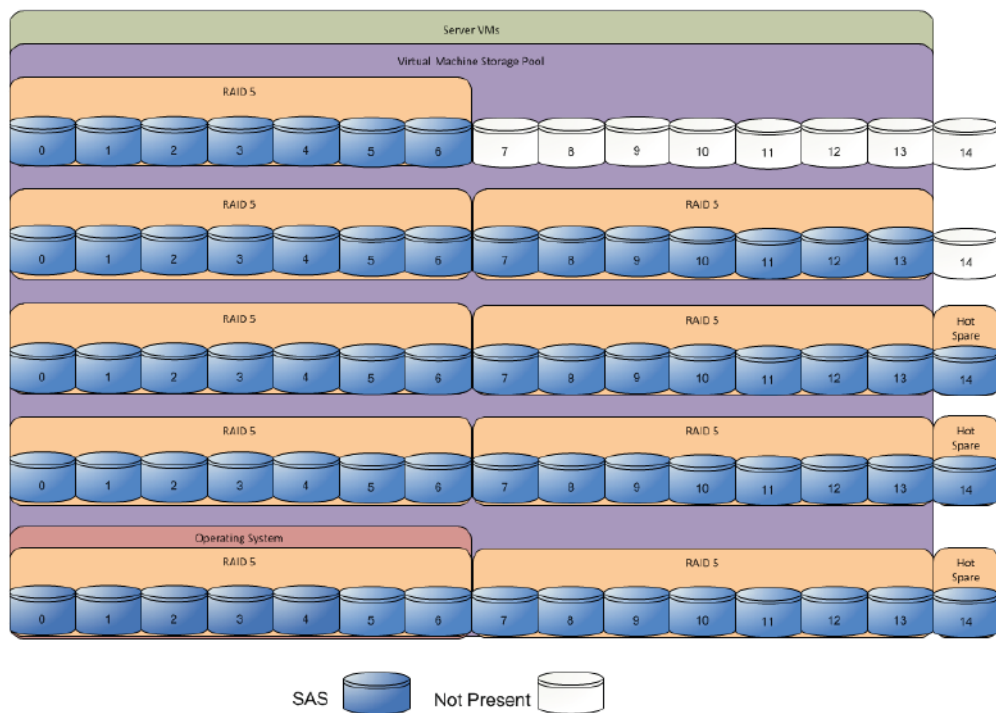


Figure 13 shows storage architecture for 125 virtual machines on VNX5300:



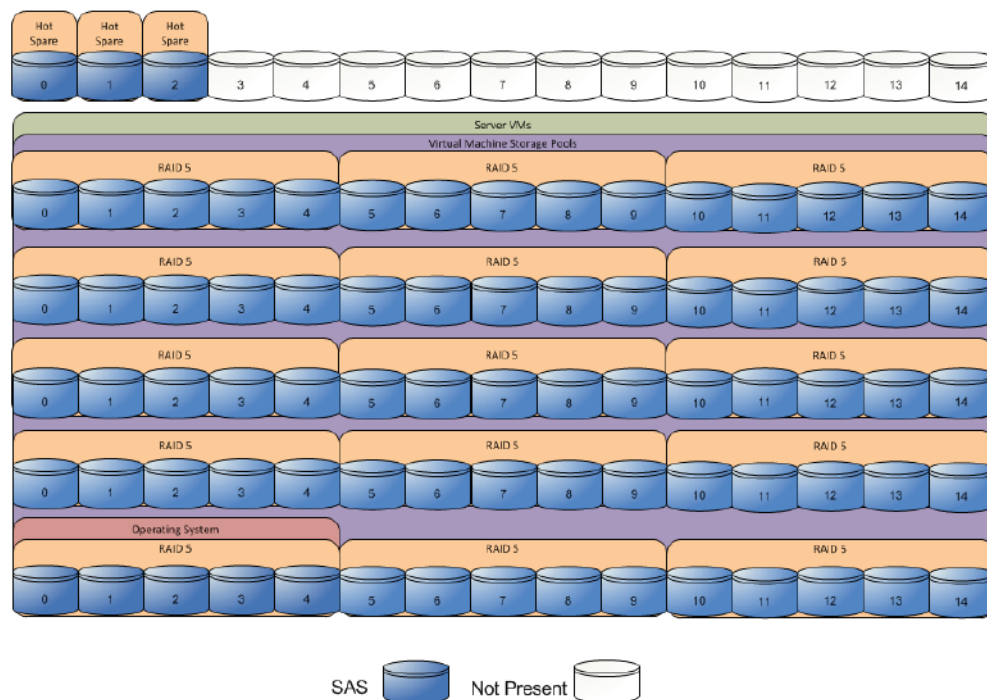
**Figure 13 Storage Architecture for 125 Virtual Machines on EMC VNX5300**

Table 8 provides the data store sizes for various architectures shown in Figure 11, Figure 12, Figure 13:

**Table 8 Data store sizes**

Parameters	50 virtual machines	100 virtual machines	125 virtual machines
Disk capacity & type	600 GB SAS	600 GB SAS	600 & 300 GB SAS
Number of disks	30	63	60 (600 GB) 15 (300 GB)
RAID type	4 + 1 RAID 5 groups	6 + 1 RAID 5 groups	4 + 1 RAID 5 groups
Number of pools	6	9	15

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

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

## Storage Virtualization

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

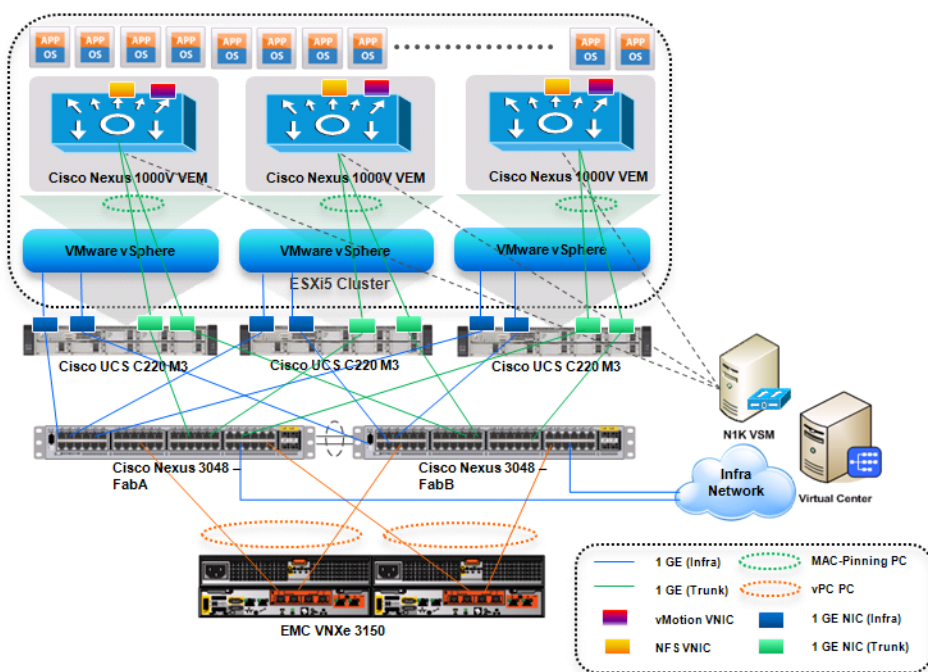
It is preferable to deploy virtual machine files on shared storage to take advantage of VMware VMotion, VMware High Availability™ (HA), and VMware Distributed Resource Scheduler™ (DRS). This is considered a best practice for mission-critical deployments, which are often installed on third-party, shared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

## Architecture for 50 VMware Virtual Machines

Figure 14 shows the logical layout of 50 VMware virtual machines. Following are the key aspects of this solution:

- Three Cisco C220 M3 servers are used.
- The solution uses Cisco Nexus 3048 switches and Broadcom 1Gbps NIC. This results in the 1Gbps solution for the storage access.
- Virtual port-channels on storage side networking provide high-availability and load balancing.
- Cisco Nexus 1000v distributed Virtual Switch provides port-profiles based virtual networking solution.
- On server side, port-profile based MAC pinning feature provides simplified load balancing and network high availability.
- EMC VNXe3150 is used as a storage array.

**Figure 14** Cisco Solution VMware Architecture for 50 Virtual Machines

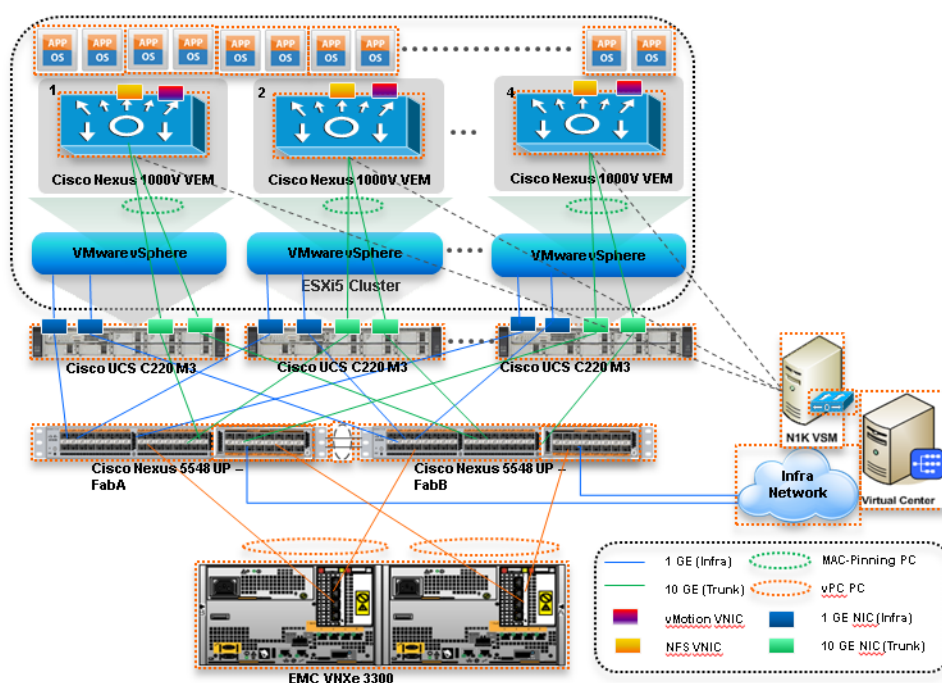


## Architecture for 100 VMware Virtual Machines

Figure 15 shows the logical layout of 100 VMware virtual machines. Following are the key aspects of this solution:

- Four Cisco C220 M3 servers are used.
- The solution uses Cisco Nexus 5548UP switches and 10 Gbps Cisco VIC adapters. This results in the 10Gbps solution for the storage access and network and makes vMotion 9 times faster compared to the 1 Gbps solution.
- Virtual port-channels on storage side networking provide high-availability and load balancing.
- Cisco Nexus 1000v distributed Virtual Switch provides port-profiles based virtual networking solution.
- On server side, port-profile based MAC pinning feature provides simplified load balancing and network high availability.
- EMC VNXe3300 is used as a storage array.

**Figure 15** Cisco Solution VMware Architecture for 100 Virtual Machines



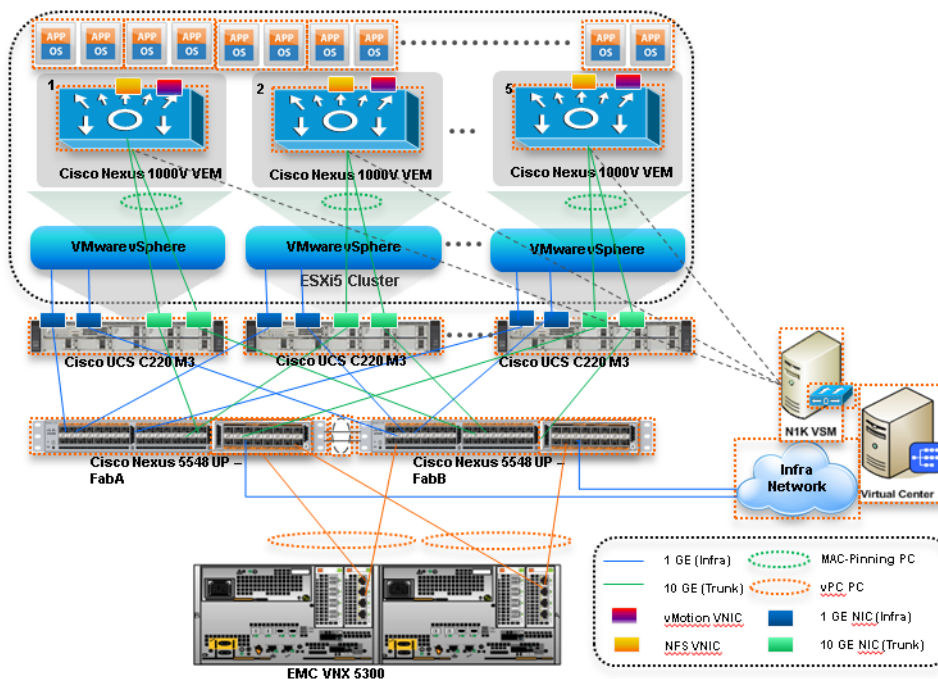
## Architecture for 125 VMware Virtual Machines

Figure 16 shows the logical layout of 125 VMware virtual machines. Following are the key aspects of this solution:

- Five Cisco C220 M3 servers are used.
- The solution uses Cisco Nexus 5548UP switches and 10 Gbps Cisco VIC adapters. This results in the 10Gbps solution for the storage access and network and makes vMotion 9 times faster compared to the 1 Gbps solution.

- Virtual port-channels on storage side networking provide high-availability and load balancing.
- Cisco Nexus 1000v distributed Virtual Switch provides port-profiles based virtual networking solution.
- On server side, port-profile based MAC pinning feature provides simplified load balancing and network high availability.
- EMC VNX5300 is used as a storage array.

**Figure 16** Cisco Solution VMware Architecture for 125 Virtual Machines



## Sizing Guidelines

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

## Defining the Reference Workload

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

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

**Table 9**      **Virtual machine characteristics**

Characteristic	Value
Virtual machine operating system	Microsoft Windows Server 2008 R1 SP1
Virtual processor per virtual machine (vCPU)	1
RAM per virtual machine	2 GB
Available storage capacity per virtual machine	100 GB
I/O operations per second (IOPS) per VM	25
I/O pattern	Random
I/O read/write ratio	2:1

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

## Applying the Reference Workload

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

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

### **Example 1**      **Custom Built Application**

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

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

- CPU resources for 1 VM
- Memory resources for 2 VMs
- Storage capacity for 1 VM
- IOPS for 1 VM

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

**Example 2     Point of Sale System**

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

The following resources that are needed from the resource pool to virtualize this application:

- CPUs of 4 reference VMs
- Memory of 8 reference VMs
- Storage of 2 reference VMs
- IOPS of 8 reference VMs

In this case the one virtual machine uses the resources of eight reference virtual machines. If this was implemented on a resource pool for 50 virtual machines, there are 42 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 resources that are needed from the resource pool to virtualize this application:

- CPUs of 2 reference VMs
- Memory of 4 reference VMs
- Storage of 1 reference VMs
- IOPS of 2 reference VMs

In this case the virtual machine would use the resources of four reference virtual machines. If this was implemented on a resource pool for 125 virtual machines, there are 121 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 10 CPUs and 48 GB of memory. It uses 5 TB of storage and generates 700 IOPS during an average busy cycle.

The following resources that are needed from the resource pool to virtualize this application:

- CPUs of ten reference VMs
- Memory of 24 reference VMs
- Storage of 52 reference VMs
- IOPS of 28 reference VMs

In this case the one virtual machine uses the resources of 52 reference virtual machines. If this was implemented on a resource pool for 100 virtual machines, there are 48 virtual machines of capability remaining in the pool.

**Summary of Example**

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

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

# VSPEX Configuration Guidelines

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

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

1. Pre-deployment tasks.
2. Prepare servers.
3. Prepare switches, connect network and configure switches.
4. Prepare and configure storage array.
5. Install ESXi servers and vCenter infrastructure.
6. Install and configure SQL server database.
7. Install and configure vCenter server.
8. Install and configure Nexus 1000v.
9. Test the installation.

These steps are described in detail in the following sections.

## Pre-Deployment Tasks

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

- **Gather documents**—Gather the related documents listed in the Preface. These are used throughout the text of this document to provide detail on setup procedures and deployment best practices for the various components of the solution.
- **Gather tools**—Gather the required and optional tools for the deployment. Use 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 10**      **Customer- Specific Configuration Data**

Requirement	Description	Reference
Hardware	Cisco UCS C220 M3 servers to host virtual machines	EMC-Cisco Reference Architecture: <i>VSPEX Server Virtualization with VMware vSphere 5 for up to 50, 100 or 125 Virtual Machines.</i>
	VMware vSphere™ 5 server to host virtual infrastructure servers Note: This requirement may be covered in the existing infrastructure	
	Cisco Nexus switches: Two Cisco Nexus 5548UP or 3048 switches for high availability	
	EMC VNX/VNXe storage: Multiprotocol storage array with the required disk layout as per architecture requirements	
Software	VMware ESXi™ 5.0 installation media	See the corresponding product documentation
	VMware vCenter Server 5.0 installation media	
	Cisco Nexus 1000v virtual switch installation media	
	EMC VSI for VMware vSphere: Unified Storage Management – Product Guide	
	EMC VSI for VMware vSphere: Storage Viewer—Product Guide	
	Microsoft Windows Server 2008 R2 SP1 installation media (suggested OS for VMware vCenter)	
	Microsoft SQL Server 2008 R2 SP1 Note: This requirement may be covered in the existing infrastructure	
Licenses	VMware vCenter 5.0 license key	Consult your corresponding vendor obtain license keys
	VMware ESXi 5.0 license keys	
	Microsoft SQL Server license key	
	Note: This requirement may be covered in the existing infrastructure	



## Customer Configuration Data

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

The section [Customer Configuration Data Sheet, page 150](#) provides tabulated record of relevant information (to be filled at the customer's end). This form can be expanded or contracted as required, and information may be added, modified, and recorded as the deployment progresses.

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

## Preparing Servers

Preparing the Cisco C220 M3 servers is a common step for all the VMware architectures. Firstly, you need to install the C220 M3 server in a rack. For more information on mounting the Cisco C220 servers, see the installation guide on details about how to physically mount the server:

[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/c/hw/C220/install/install.html](http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/install.html)

To prepare the servers, follow these steps:

1. [Configure Management IP Address for CIMC Connectivity, page 37](#)
2. [Enabling Virtualization Technology in BIOS, page 38](#)
3. [Configuring RAID, page 40](#)

These steps are discussed in detail in the following sections.

### Configure Management IP Address for CIMC Connectivity

To power-on the server and configure the management IP address, follow these steps:

1. Attach a supplied power cord to each power supply in your server, and then attach the power cord to a grounded AC power outlet.
2. Connect a USB keyboard and VGA monitor by using the supplied KVM cable connected to the KVM connector on the front panel.
3. Press the Power button to boot the server. Watch for the prompt to press F8.
4. During bootup, press F8 when prompted to open the BIOS CIMC Configuration Utility.
5. Set the NIC mode to Dedicated and NIC redundancy to None.
6. Choose whether to enable DHCP for dynamic network settings, or to enter static network settings.
7. Press F10 to save your settings and reboot the server.

**Figure 17** *Configuring CIMC IP in CIMC Configuration Utility*

```

CIMC Configuration Utility  Version 1.5  Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]          NIC redundancy
Shared LOM:     [ ]          None: [X]
Shared LOM 10G: [ ]          Active-standby: [ ]
Cisco Card:     [ ]          Active-active: [ ]

IPv4 (Basic)
DHCP enabled:   [ ]          Factory Defaults
CIMC IP:        10.29.150.101 CIMC Factory Default:[ ]
Subnetmask:     255.255.255.0 Default User (Basic)
Gateway:        10.29.150.1  Default password:
                                   Reenter password:

VLAN (Advanced)
VLAN enabled:   [ ]
VLAN ID:        1
Priority:        0

*****
<Up/Down arrow> Select items  <F10> Save  <Space bar> Enable/Disable
<F5> Refresh                  <ESC> Exit

```

When the CIMC IP is configured, the server can be managed using the https based Web GUI or CLI.



**Note**

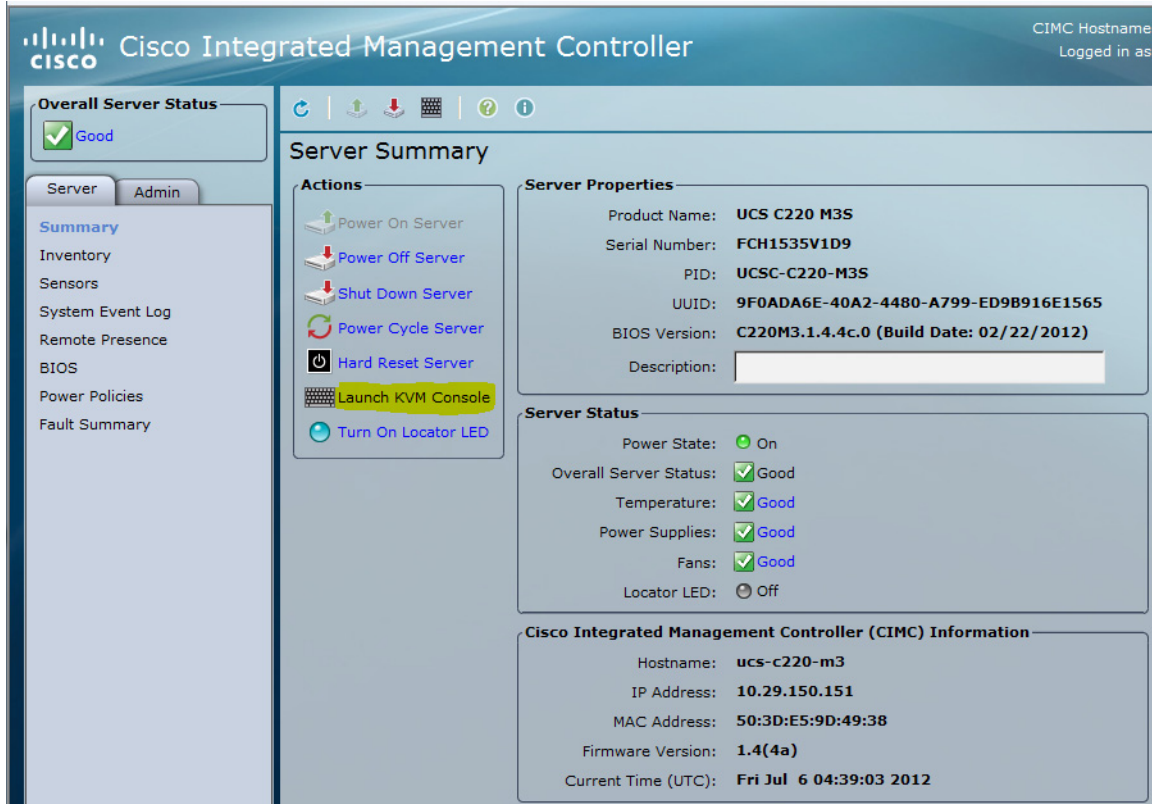
The default username for the server is “admin” and the default password is “password”. Cisco strongly recommends changing the default password.

## Enabling Virtualization Technology in BIOS

VMware vCenter requires an x64-based processor, hardware-assisted virtualization (Intel VT enabled), and hardware data execution protection (Execute Disable enabled). Perform the following steps to enable Intel ® VT and Execute Disable in BIOS.

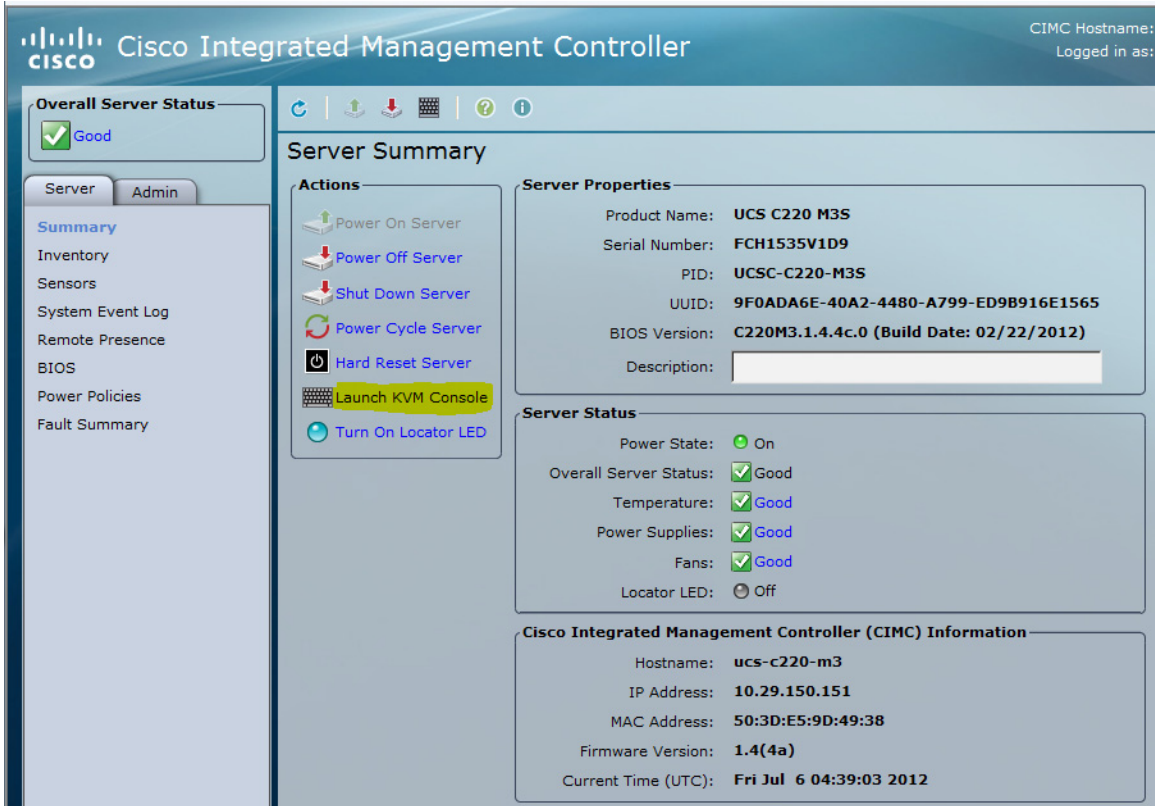
1. Using a web browser, connect to the CIMC using the IP address configured in the CIMC Configuration section.
2. Launch the KVM from the CIMC GUI.

**Figure 18**      **Launching KVM Console Through CIMC GUI**



3. Press the Power button to boot the server. Watch for the prompt to press **F2**.
4. During bootup, press **F2** when prompted to open the BIOS Setup Utility.
5. Choose **Advanced tab > Processor Configuration**.

**Figure 19**      *Enabling Virtualization Technology in KVM Console*



6. Enable Execute Disable and Intel VT as shown in [Figure 19](#).

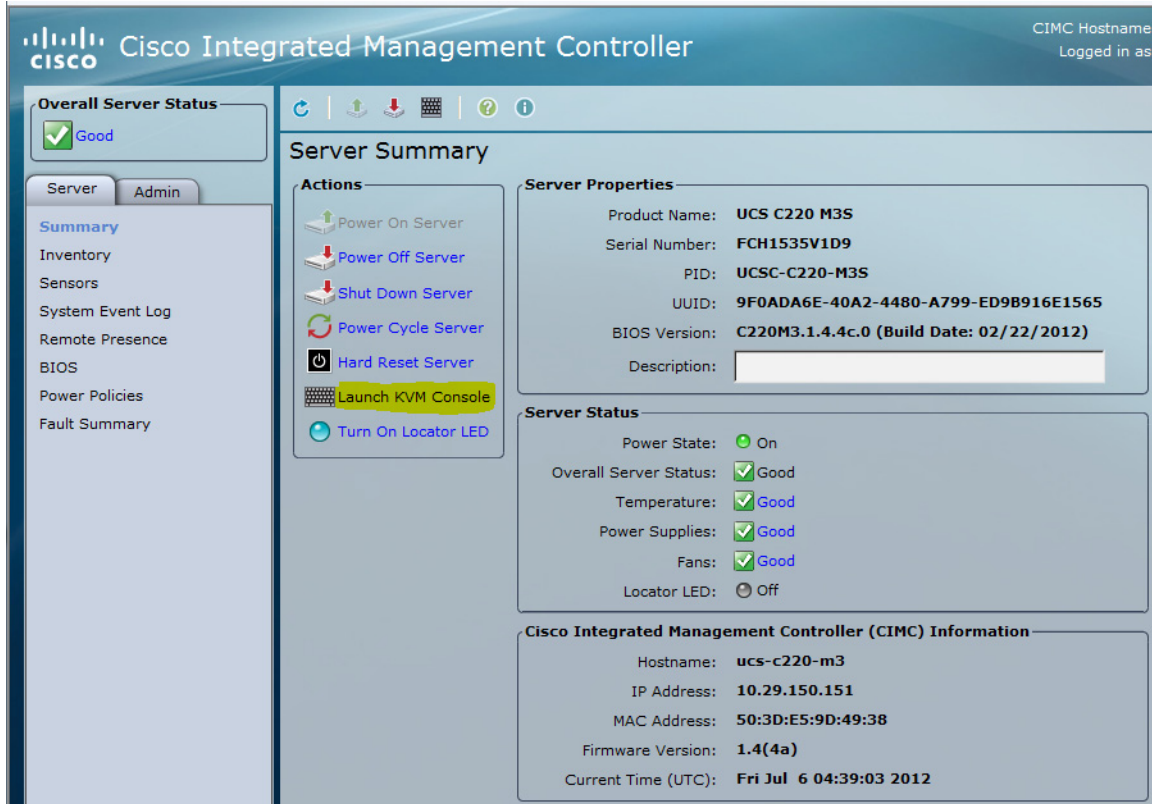
## Configuring RAID

The RAID controller type is Cisco UCSC RAID SAS 2008 and supports 0, 1, 5 RAID levels. We need to configure RAID level 1 for this setup and set the virtual drive as boot drive.

To configure RAID controller, perform the following steps:

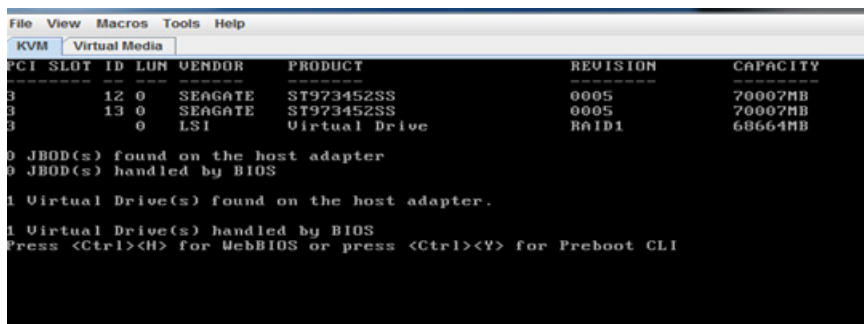
1. Using a web browser, connect to the CIMC using the IP address configured in the CIMC Configuration section.
2. Launch the KVM from the CIMC GUI.

**Figure 20**      **Launching KVM Console Through CIMC GUI**

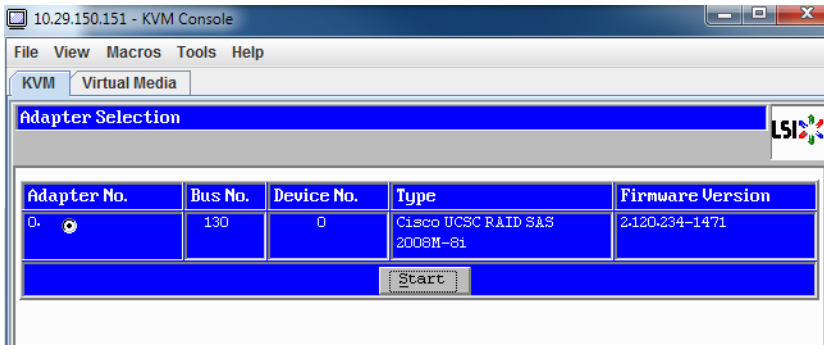


- During bootup, press <Ctrl> <H> when prompted to configure RAID in the WebBIOS.

**Figure 21**      **Opening WebBIOS Window**



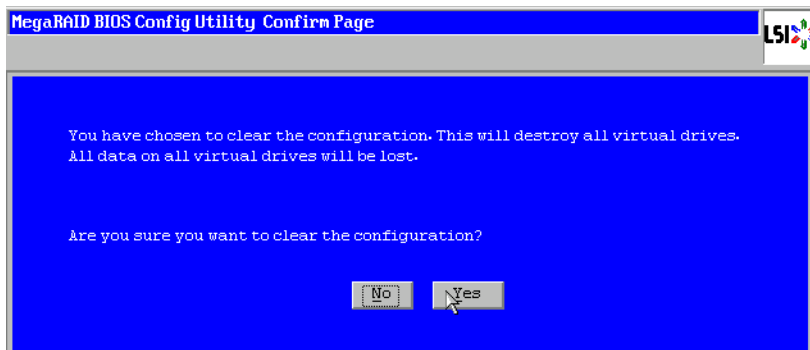
- Choose the adapter and click the **Start** button.

**Figure 22 Adapter Selection Window**

5. Choose **New Configuration** and click **Next**.

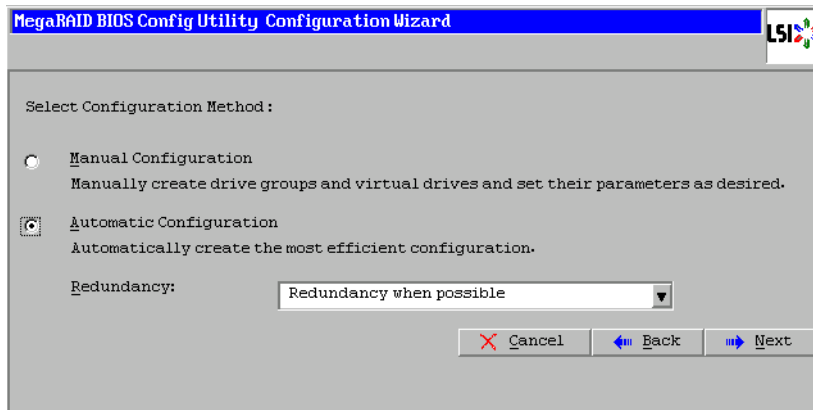
**Figure 23 MegaRAID Configuration Wizard**

6. Choose **Yes** and click **Next** to clear the configuration.

**Figure 24 MegaRAID Confirmation Window**

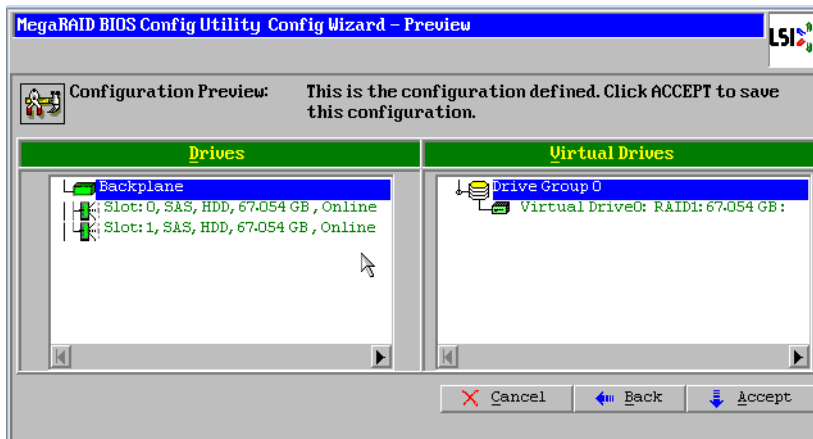
7. If you choose “Automatic Configuration” radio button and “Redundancy when possible” from the drop-down list for “Redundancy”, and if only two disks are available, then WebBIOS creates a RAID 1 configuration.

**Figure 25**      **Selecting Configuration**




8. Click **Accept** when you are prompted to save the configuration.

**Figure 26**      **MegaRAID Configuration Preview**



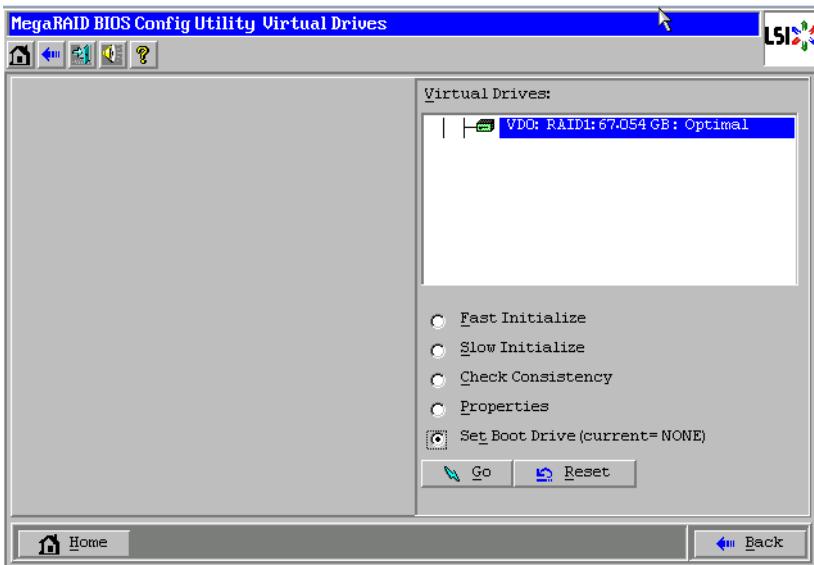
9. Click **Yes** when prompted to initialize the new virtual drives.

**Figure 27**      **Initializing New Virtual Drives**



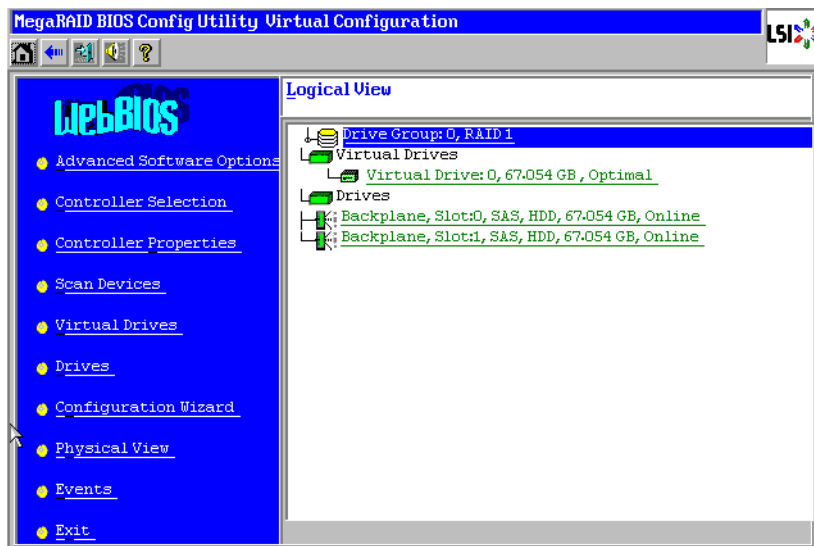
10. Choose the **Set Boot Drive** for the virtual drive created above and click the **GO** button.

**Figure 28**      **Setting Virtual Drive as Boot Drive**



11. Click **Exit** and reboot the system.

**Figure 29**      **Logical View of Virtual Configuration in WebBIOS**



## Preparing Switches, Connecting Network, and Configuring Switches

See the Nexus 3048 or Nexus 5548UP configuration guide for detailed information about how to mount the switches on the rack. Following diagrams show connectivity details for the three VMware architectures covered in this document.

Figure 30, Figure 32, Figure 34 show there are five major cabling sections in these architectures:

1. Inter switch links.
2. Data connectivity for servers (trunk links).
3. Management connectivity for servers.



4. Storage connectivity.
5. Infrastructure connectivity.

## Topology Diagram for 50 Virtual Machines

**Figure 30**      *Topology Diagram for 50 Virtual Machines*





Table 11 and Figure 31 provide the detailed cable connectivity for the 50 virtual machines configuration.

**Table 11**      *Cabling details for 50 Virtual Machines*

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
A,C	Eth1/7	All	Trunk	10(D)	7	VPC peer link
B,D	Eth1/8	All	Trunk	10(D)	7	VPC peer link
E,H	Eth1/1	1	Access	1(D)	2	C220 Server1- 1GE LOM 1
F,I	Eth1/2	1	Access	1(D)	3	C220 Server2- 1GE LOM 1
G,J	Eth1/3	1	Access	1(D)	4	C220 Server3- 1GE LOM 1
K,N	Eth1/9	40-45	vntag	10(D)	-	C220 Server1- Broadcom adapter
L,O	Eth1/10	40-45	vntag	10(D)	-	C220 Server1- Broadcom adapter
M,P	Eth1/11	40-45	vntag	10(D)	-	C220 Server1- Broadcom adapter
Q,S	Eth2/1	40	Access	10(D)	21	VNXe3150 - SPA

**Table 11**      **Cabling details for 50 Virtual Machines**

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
R,T	Eth2/2	40	Access	10(D)	22	VNXe3150 - SPB
(not shown)	Eth1/15	1,41-45	Trunk	10(D)	15	Uplink to Infrastructure network
(not shown)	Eth1/17	1,41-45	Trunk	10(D)	17	Uplink to Infrastructure network

**Figure 31**      **Detailed Backplane Connectivity for 50 Virtual Machines**

After connecting all the cables as per [Table 11](#), you can configure the switch.

## Topology Diagram for 100 Virtual Machines

**Figure 32**      *Topology Diagram for 100 Virtual Machines*




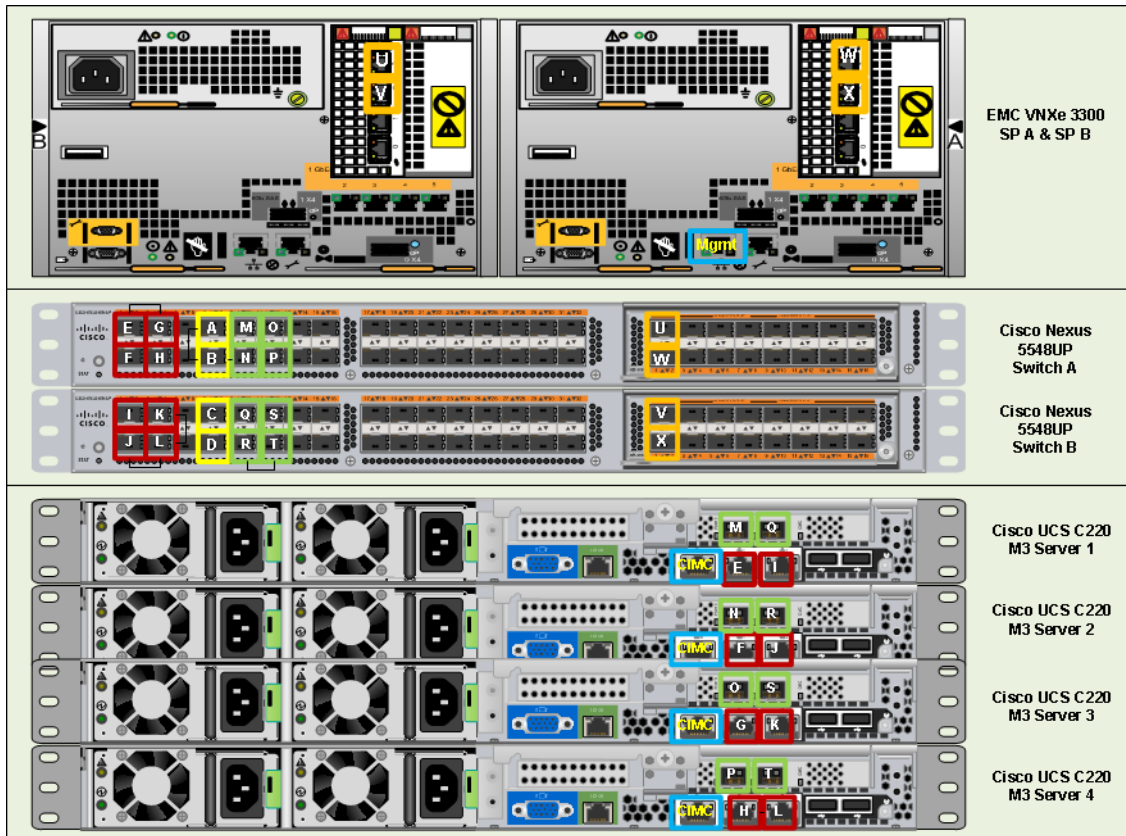
Table 12 and Figure 33 provides the detailed cable connectivity for the 100 virtual machines configuration.

**Table 12**      *Cabling Details for 100 Virtual Machines*

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device Port
A,C	Eth1/7	All	Trunk	10(D)	7	VPC peer link
B,D	Eth1/8	All	Trunk	10(D)	7	VPC peer link
E,I	Eth1/1	1	Access	1(D)	2	C220 Server1- 1GE LOM 1
F,J	Eth1/2	1	Access	1(D)	3	C220 Server2- 1GE LOM 1
G,K	Eth1/3	1	Access	1(D)	4	C220 Server3- 1GE LOM 1
H,L	Eth1/4	1	Access	1(D)	5	C220 Server3- 1GE LOM 1
M,Q	Eth1/9	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
N,R	Eth1/10	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
O,S	Eth1/11	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0

**Table 12**      **Cabling Details for 100 Virtual Machines**

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device Port
P,T	Eth1/12	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
U,W	Eth2/1	40	Access	10(D)	21	VNXe3300 (Eth 10)- SPA
V,X	Eth2/2	40	Access	10(D)	22	VNXe3300 (Eth 10)- SPB
(not shown)	Eth1/15	1,41-45	Trunk	10(D)	15	Uplink to Infrastructure network
(not shown)	Eth1/17	1,41-45	Trunk	10(D)	17	Uplink to Infrastructure network

**Figure 33**      **Detailed Backplane Connectivity for 100 Virtual Machines**

After connecting all the cables as per [Table 12](#), you can configure the switch.

## Topology Diagram for Hundred and Twenty Five Virtual Machines

**Figure 34**      *Topology Diagram for 125 Virtual Machines*




Table 13 and Figure 35 provides the detailed cable connectivity for the 125 virtual machines configuration.

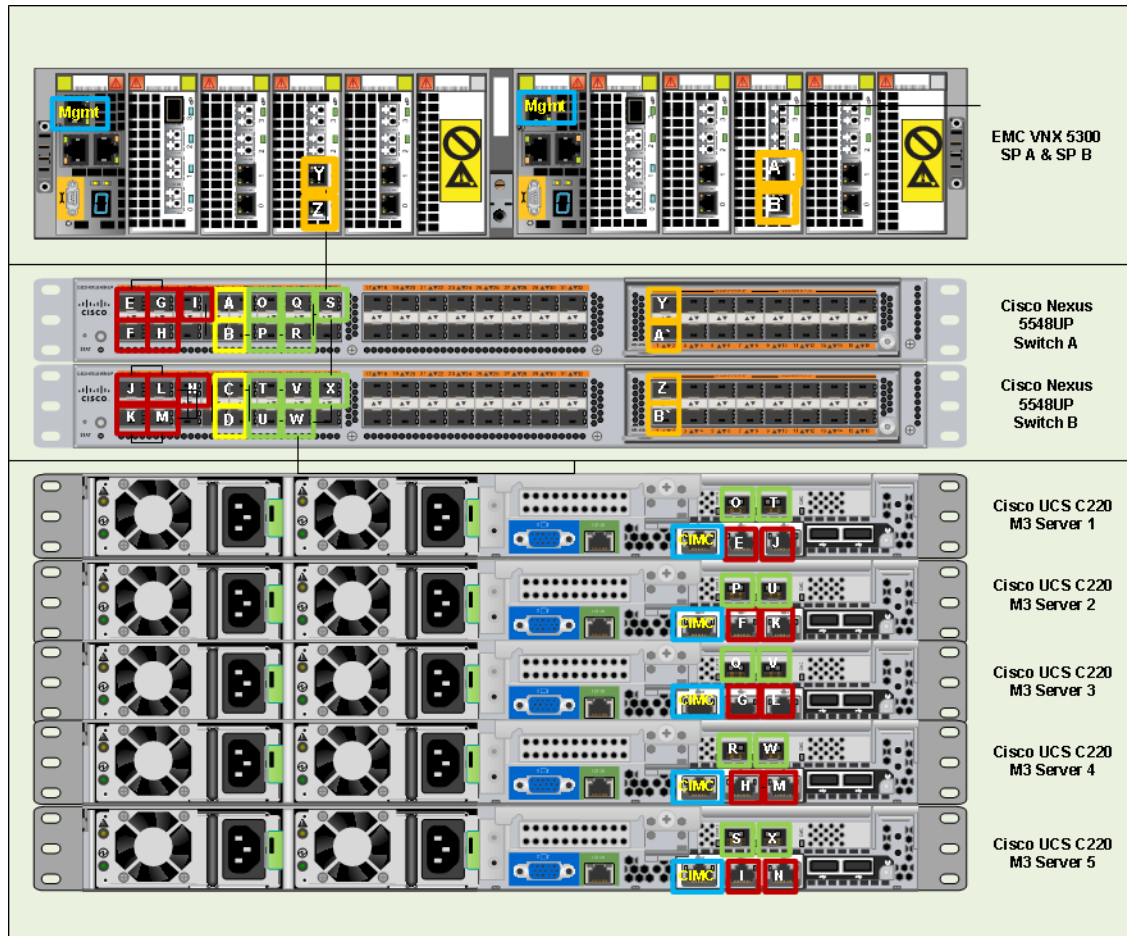
**Table 13**      *Cabling details for 100 Virtual Machines*

Cable ID	Switch Interface	VLAN	Mode	Speed (Gpbs)	Port Channel	Remote Device Port
A,C	Eth1/7	All	Trunk	10(D)	7	VPC peer link
B,D	Eth1/8	All	Trunk	10(D)	7	VPC peer link
E,J	Eth1/1	1	Access	1(D)	2	C220 Server1- 1GE LOM 1
F,K	Eth1/2	1	Access	1(D)	3	C220 Server2- 1GE LOM 1
G,L	Eth1/3	1	Access	1(D)	4	C220 Server3- 1GE LOM 1
H,M	Eth1/4	1	Access	1(D)	5	C220 Server3- 1GE LOM 1
I,N	Eth1/5	1	vntag	1(D)	5	C220 Server3- 1GE LOM 1
O,T	Eth1/9	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
P,U	Eth1/10	40-45	vntag	10(D)	-	C220 Server1- P81E VIC Port 0

**Table 13**      **Cabling details for 100 Virtual Machines**

<b>Cable ID</b>	<b>Switch Interface</b>	<b>VLAN</b>	<b>Mode</b>	<b>Speed (Gbps)</b>	<b>Port Channel</b>	<b>Remote Device Port</b>
Q,V	Eth1/11	40-45	vntag	10(D)	-	C220 Server1-P81E VIC Port 0
R,W	Eth1/12	40-45	vntag	10(D)	-	C220 Server1-P81E VIC Port 0
S,X	Eth1/13	40-45	vntag	10(D)	-	C220 Server1-P81E VIC Port 0
Y,A'	Eth2/1	40	Access	10(D)	21	VNXe3300 (Eth 10)- SPA
Z,B'	Eth2/2	40	Access	10(D)	22	VNXe3300 (Eth 10)- SPB
(not shown)	Eth1/15	1,41-45	Trunk	10(D)	15	Uplink to Infrastructure network
(not shown)	Eth1/17	1,41-45	Trunk	10(D)	17	Uplink to Infrastructure network

**Figure 35 Detailed Backplane Connectivity for 100 Virtual Machines**



After connecting all the cables as per [Table 13](#), you can configure the switch.

## Configuring Cisco Nexus Switches

This section explains switch configuration needed for the Cisco solution for EMC VSPEX VMware architectures. Details about configuring password, management connectivity and strengthening the device are not covered here, please refer to the Nexus 3000 and 5000 series configuration guide for that.

### Configure Global VLANs

Following is an example to configure VLAN on a switch:

```
switch# configure terminal
switch(config)# vlan 40
switch(config-vlan)# name Storage
switch(config-vlan)# exit
```

Following VLANs in [Table 14](#) need to be configured on both switches A and B in addition to your application specific VLANs:

**Table 14**      **Configured VLANs on switch A and B**

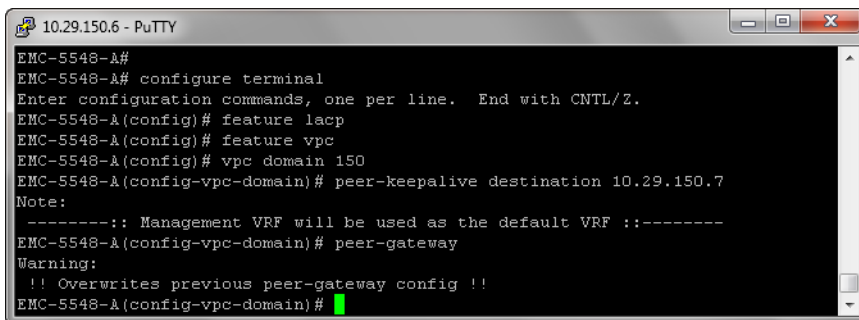
VLAN Name	Description
Storage	VLAN to access storage array from the servers
vMotion	VLAN for virtual machine vMotion
N1K-Mgmt	Management VLAN for Nexus 1000v virtual switch
N1K-Control	VLAN for control traffic between Nexus 1000v VSM and ESXi servers
N1K-Packet	VLAN for packet traffic between Nexus 1000v VSM and ESXi servers
VM-Data	VLAN for the virtual machine (application) traffic (can be multiple VLANs)

For actual VLAN IDs of your deployment, see the section Customer Configuration Data Sheet.

### Configuring Virtual Port Channel (VPC)

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

1. Enable LACP feature on both switches.
2. Enable vPC feature on both switches.
3. Configure a unique vPC domain ID, identical on both switches.
4. Configure mutual management IP addresses on both the switches and configure peer-gateway as shown in the [Figure 36](#).

**Figure 36**      **Configuring Peer-Gateway**


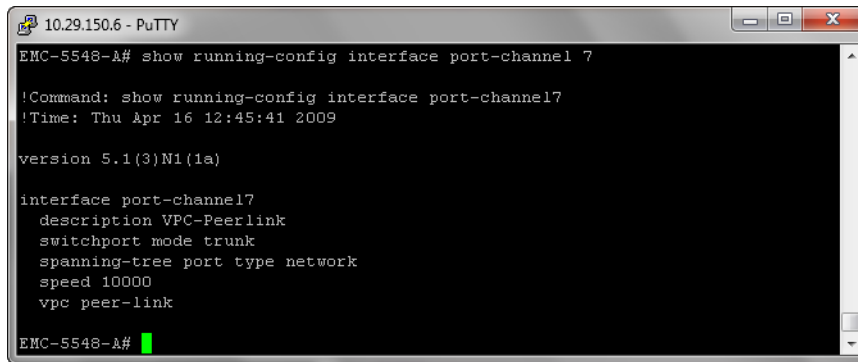
```

10.29.150.6 - PuTTY
EMC-5548-A#
EMC-5548-A# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
EMC-5548-A(config)# feature lacp
EMC-5548-A(config)# feature vpc
EMC-5548-A(config)# vpc domain 150
EMC-5548-A(config-vpc-domain)# peer-keepalive destination 10.29.150.7
Note:
-----:: Management VRF will be used as the default VRF ::-----
EMC-5548-A(config-vpc-domain)# peer-gateway
Warning:
!! Overwrites previous peer-gateway config !!
EMC-5548-A(config-vpc-domain)#

```

5. Configure port-channel on the inter-switch links. Configuration for these ports is shown in [Figure 37](#). Ensure that “vpc peer-link” is configured on this port-channel.



**Figure 37** Configured VPC Peer-link on Port-Channel


```

10.29.150.6 - PuTTY
EMC-5548-A# show running-config interface port-channel 7

!Command: show running-config interface port-channel7
!Time: Thu Apr 16 12:45:41 2009

version 5.1(3)N1(1a)

interface port-channel7
 description VPC-Peerlink
 switchport mode trunk
 spanning-tree port type network
 speed 10000
 vpc peer-link

EMC-5548-A#

```

6. Add ports with LACP protocol on the port-channel using “channel-group 7 mode active” command under the interface subcommand.

### Configuring Infrastructure Ports Connected to Servers

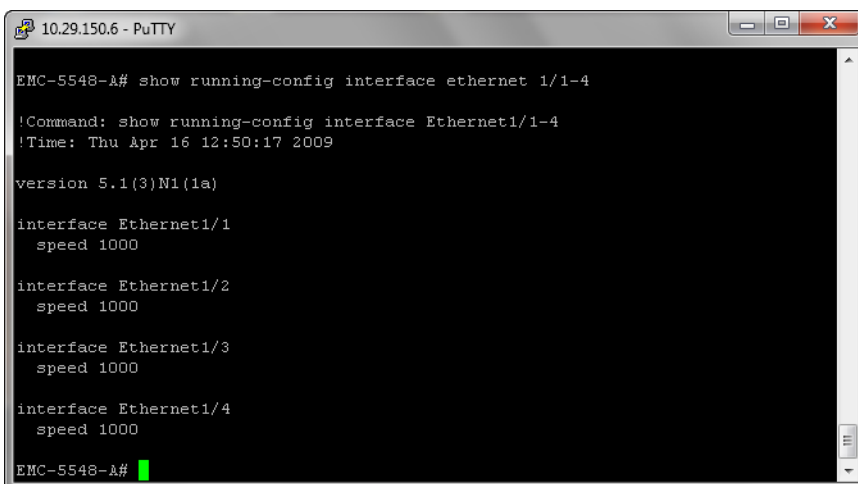
Infrastructure links connected to the LOM ports on the servers always operate at the speed of 1Gbps and carry only the infrastructure VLAN. [Figure 38](#) shows the configuration for infrastructure ports on the switches.



#### Note

In the test environment, VLAN 1 was used as infrastructure VLAN.

In your deployment, if the infrastructure VLAN is not VLAN 1, then you need to explicitly configure as an access VLAN using “switchport access vlan <id>” under the interface subcommand.

**Figure 38** Configuring Infrastructure Ports on Switches


```

10.29.150.6 - PuTTY
EMC-5548-A# show running-config interface ethernet 1/1-4

!Command: show running-config interface Ethernet1/1-4
!Time: Thu Apr 16 12:50:17 2009

version 5.1(3)N1(1a)

interface Ethernet1/1
 speed 1000

interface Ethernet1/2
 speed 1000

interface Ethernet1/3
 speed 1000

interface Ethernet1/4
 speed 1000

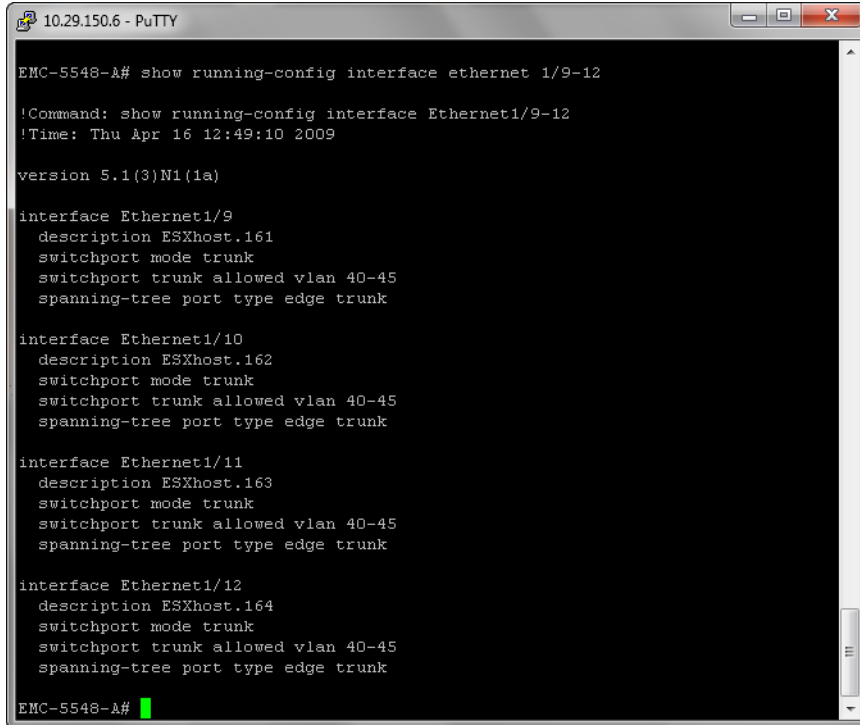
EMC-5548-A#

```

## Configuring Trunk Ports Connected to Servers

Data ports connected to the ESXi servers need to be in the trunk mode. Storage, vMotion, N1k-control, N1k-packet and application VLANs are allowed on this port. For 100 and 125 virtual machines architectures, these trunk ports operate at the speed of 10 Gbps. It is recommended to provide good description for each port and port-channel on the switch, to ease troubleshooting incase of any issues later. Exact configuration commands are shown in [Figure 39](#).

**Figure 39** Port-Channel Configuration Commands



```

10.29.150.6 - PuTTY
EMC-S548-A# show running-config interface ethernet 1/9-12

!Command: show running-config interface Ethernet1/9-12
!Time: Thu Apr 16 12:49:10 2009

version 5.1(3)N1(1a)

interface Ethernet1/9
  description ESXhost.161
  switchport mode trunk
  switchport trunk allowed vlan 40-45
  spanning-tree port type edge trunk

interface Ethernet1/10
  description ESXhost.162
  switchport mode trunk
  switchport trunk allowed vlan 40-45
  spanning-tree port type edge trunk

interface Ethernet1/11
  description ESXhost.163
  switchport mode trunk
  switchport trunk allowed vlan 40-45
  spanning-tree port type edge trunk

interface Ethernet1/12
  description ESXhost.164
  switchport mode trunk
  switchport trunk allowed vlan 40-45
  spanning-tree port type edge trunk

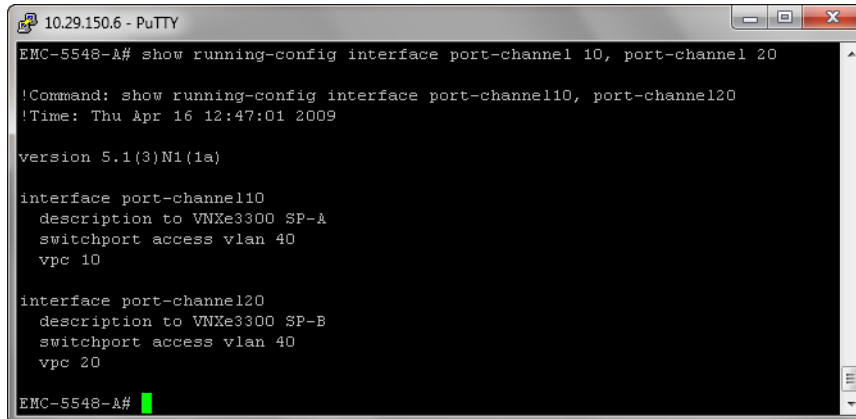
EMC-S548-A#

```

In this test environment, we are not using virtual port-channels across two switches for load-balancing and high-availability, as we have used port-profile based load-balancing and high-availability defined in the Cisco Nexus 1000v virtual distributed switch on the host to reduce complexity of the architecture. As mentioned in the next configuration step, the port-channel and vPC are used on storage side connectivity for load-balancing and high-availability.

## Configuring Storage Connectivity

From each switch one link connects to each storage processor on the VNX/VNXe storage array. A virtual port-channel is created for the two links connected to a single storage processor, but connected to two different switches. In this example configuration, links connected to storage processor A (SP-A) of VNXe3300 storage array are connected to Ethernet port 2/1 on each switch and links connected to storage processor B (SP-B) are connected to Ethernet port 2/2 on each switch. A virtual port-channel (id 10) is created for port Ethernet 2/1 on each switch and a different virtual port-channel (id 20) is created for port Ethernet 2/2 on each switch. [Figure 40](#) shows the configuration on the port-channels.

**Figure 40** Configured Virtual Port-Channels


```

10.29.150.6 - PuTTY
EMC-5548-A# show running-config interface port-channel 10, port-channel 20
!Command: show running-config interface port-channel10, port-channel20
!Time: Thu Apr 16 12:47:01 2009

version 5.1(3)N1(1a)

interface port-channel10
  description to VNXe3300 SP-A
  switchport access vlan 40
  vpc 10

interface port-channel20
  description to VNXe3300 SP-B
  switchport access vlan 40
  vpc 20

EMC-5548-A#

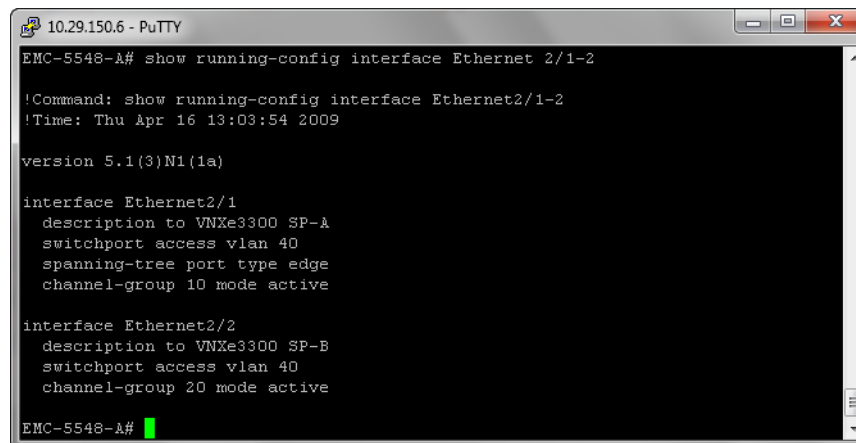
```

Figure 41 shows the exact configuration on each port connected to the storage array.

**Note**

Only storage VLAN is required on this port, and so, the port will be in the access mode.

Port is part of the port-channel configured in Figure 40.

**Figure 41** Port Connections to Storage Arrays


```

10.29.150.6 - PuTTY
EMC-5548-A# show running-config interface Ethernet 2/1-2
!Command: show running-config interface Ethernet2/1-2
!Time: Thu Apr 16 13:03:54 2009

version 5.1(3)N1(1a)

interface Ethernet2/1
  description to VNXe3300 SP-A
  switchport access vlan 40
  spanning-tree port type edge
  channel-group 10 mode active

interface Ethernet2/2
  description to VNXe3300 SP-B
  switchport access vlan 40
  channel-group 20 mode active

EMC-5548-A#

```

**Configure ports connected to infrastructure network**

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

**Verify VLAN and port-channel configuration**

At this point of time, all ports and port-channels are configured with necessary VLANs, switchport mode and vPC configuration. Validate this configuration using the “show vlan”, “show port-channel summary” and “show vpc” commands as shown in Figure 42.

**Figure 42** Validating Created Port-Channels with VLANs

```

10.29.150.6 - PuTTY
EMC-5548-A# show vlan id 40-45

VLAN Name                Status    Ports
-----
40   Storage                active    Po7, Po10, Po20, Eth1/9
                           Eth1/10, Eth1/11, Eth1/12
41   vMotion                active    Po7, Eth1/9, Eth1/10, Eth1/11
                           Eth1/12
42   N1K-Mgmt               active    Po7, Po15, Po17, Eth1/9
                           Eth1/10, Eth1/11, Eth1/12
43   N1K-Control            active    Po7, Po15, Po17, Eth1/9
                           Eth1/10, Eth1/11, Eth1/12
44   N1K-Packet             active    Po7, Po15, Po17, Eth1/9
                           Eth1/10, Eth1/11, Eth1/12
45   VM-Traffic             active    Po7, Po15, Po17, Eth1/9
                           Eth1/10, Eth1/11, Eth1/12

VLAN Type  Vlan-mode
-----
40   enet    CE
41   enet    CE
42   enet    CE
43   enet    CE
44   enet    CE
45   enet    CE

Remote SPAN VLANs
-----

Primary  Secondary  Type      Ports
-----
EMC-5548-A#

```

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

Port-channel configuration can be verified using “show port-channel summary” command. [Figure 43](#) shows the expected output of this command.

**Figure 43** Verifying Port-Channel configuration

```

10.29.150.6 - PuTTY
EMC-5548-A# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        S - Suspended     R - Module-removed
        U - Up            R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met

Group  Port-Channel  Type      Protocol  Member Ports
-----
7      Po7(SU)       Eth       LACP      Eth1/7(P)  Eth1/8(P)
10     Po10(SU)      Eth       LACP      Eth2/1(P)
15     Po15(SU)      Eth       LACP      Eth1/15(P)
17     Po17(SU)      Eth       LACP      Eth1/17(P)
20     Po20(SU)      Eth       LACP      Eth2/2(P)
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#

```

In this example, port-channel 7 is the vPC peer-link port-channel, port-channels 10 and 20 are connected to the storage arrays and port-channels 15 and 17 are connected to the infrastructure network. Make sure that state of the member ports of each port-channel is “P” (Up in port-channel). Note that port may not come up if the peer ports are not properly configured. Common reasons for port-channel port being down are:

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

vPC status can be verified using “show vpc” command. Example output is shown in [Figure 44](#).

**Figure 44** Verifying VPC Status

```

10.29.150.6 - PuTTY
EMC-5548-A# show vpc
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 150
Peer status            : peer adjacency formed ok
vPC keep-alive status  : peer is alive
Configuration consistency status: success
Per-vlan consistency status : success
Type-2 consistency status : success
vPC role               : secondary
Number of vPCs configured : 4
Peer Gateway           : Enabled
Peer gateway excluded VLANs : -
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  --  ---
1   Po7   up    1,40-46

vPC status
-----
id  Port  Status Consistency Reason Active vlans
--  --  ---
10  Po10  up    success success 40
15  Po15  up    success success 1,42-45
17  Po17  up    success success 1,42-45
20  Po20  up    success success 40

EMC-5548-A#
EMC-5548-A#
EMC-5548-A#
EMC-5548-A#

```

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

## Configure QoS

The Cisco solution for the EMC VSPEX VMware architectures require MTU to be set at 9000 (jumbo frames) for efficient storage and vMotion traffic. MTU configuration on Cisco Nexus 5000 and 3000 series switches fall under global QoS configuration. You may need to configure additional QoS parameters as needed by the applications. For more information on the QoS configuration, see Nexus 3000 and 5000 series configuration guide.

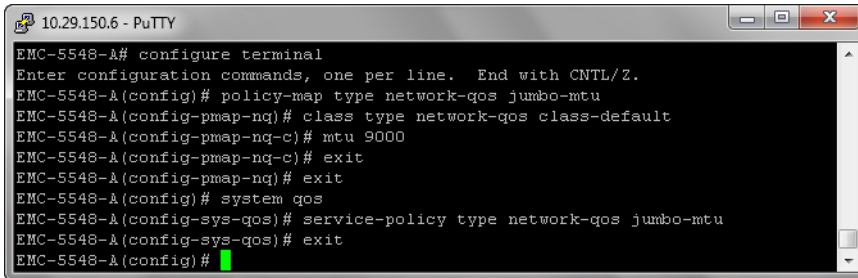
To configure 9000 MTU on the Nexus 5000 and 3000 series switches, follow these steps on both switch A and B (refer to the following figure for CLI):

1. Create a policy map named “jumbo-mtu”.

2. As we are not creating any specific QoS classification, set 9000 MTU on the default class.
3. Configure the system level service policy to the “jumbo-mtu” under the global “system qos” subcommand.

Figure 45 shows the exact Cisco Nexus CLI for the steps mentioned above.

**Figure 45** *Configuring MTU on Cisco Nexus Switches*



```

10.29.150.6 - PuTTY
EMC-5548-A# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
EMC-5548-A(config)# policy-map type network-qos jumbo-mtu
EMC-5548-A(config-pmap-nq)# class type network-qos class-default
EMC-5548-A(config-pmap-nq-c)# mtu 9000
EMC-5548-A(config-pmap-nq-c)# exit
EMC-5548-A(config-pmap-nq)# exit
EMC-5548-A(config)# system qos
EMC-5548-A(config-sys-qos)# service-policy type network-qos jumbo-mtu
EMC-5548-A(config-sys-qos)# exit
EMC-5548-A(config)#

```

## Preparing and Configuring EMC Storage Array

Storage array configuration for VMware architecture varies depending on the scale. For 50 and 100 virtual machines, VNXe3150 and VNXe3300 arrays are used respectively. For 125 virtual machines, VNX5300 array is used. GUI configuration for VNXe and VNX arrays differ significantly, so they are described separately in the following sections.

Note that at a high level, following steps are taken:

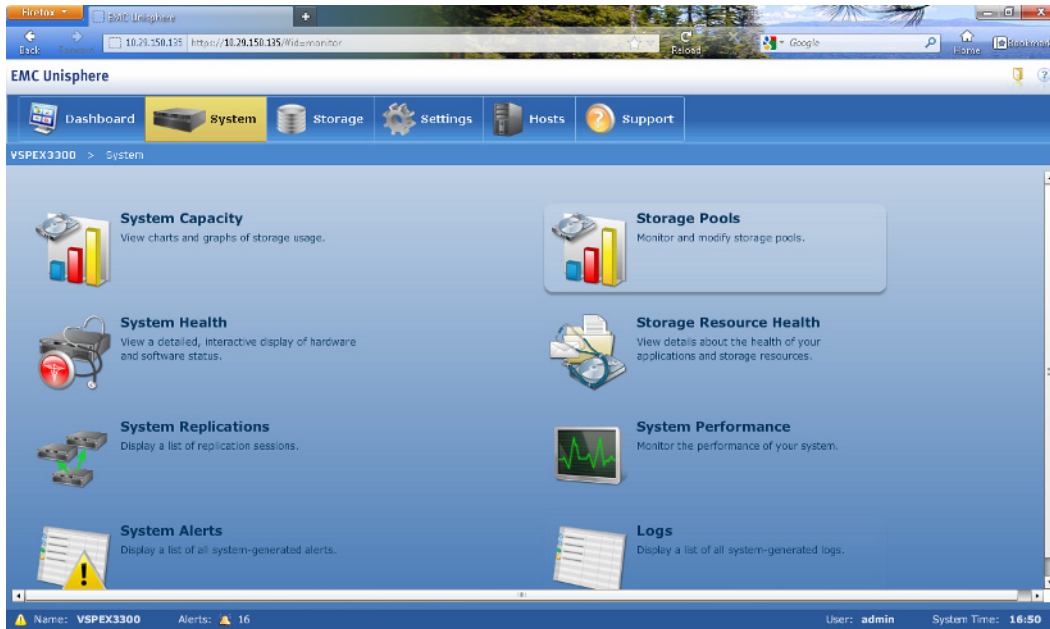
1. Create a single data store for virtual machines operating systems, and at least two data stores for the virtual machines data.
2. Configure NFS share and assign host access privileges.
3. Configure port-channel (aggregation) and jumboframe.

### Configuring VNXe3000 series storage arrays

This section covers configuring storage array for 50 and 100 virtual machines. Follow these steps to configure storage arrays:


1. Using the browser, launch the Unisphere Console with the management IP address of the storage array. Provide user name and password.
2. Click **System > Storage Pools** in the EMC Unisphere window.

**Figure 46 Storage Pools in EMC Unisphere Window**



3. Click **Configure Disks**.

**Figure 47 Configuring Disks in EMC Unisphere**



4. Click the “Manually create a new pool” radio button and choose “Pool created for VMware Storage – Datastores” from drop-down list.

**Figure 48**      **Selecting Configuration Mode**

**Disk Configuration Wizard**

**Select Configuration Mode**

Step 1 of 3

Select the disk configuration mode:

- ☐ **Automatically configure pools**  
Configure disks into the system's pools and hot spares
- ☒ **Manually create a new pool**  
Create a new pool by disk type or for a specific application  
\* **Pool created for VMware Storage - Datastore.**
- ☐ **Manually add disks to an existing pool**  
Add unconfigured disks to the selected pool  
Select pool...

< Back   Next >   Finish   Cancel   Help

- Specify the pool name and provide description in the “Name” and “Description” fields respectively. Click **Next**.

**Figure 49**      **Specifying Pool Names**

**Disk Configuration Wizard**

**Specify Pool Name**

Step 2 of 6

Specify a name and optional description.

Name: \* OS-Datastore

Description: Datastore for Virtual Machine operating systems

< Back   Next >   Finish   Cancel   Help

- Select SAS for balanced performance storage profile.



**Figure 50**      **Selecting Required Storage Space**

**Disk Configuration Wizard**

**Select Amount of Storage**

Step 4 of 6

Select the amount of storage to configure.

300GB SAS Disks: **Use 7 of 78 Disks**

Total Disks to Configure: 7

< Back   Next >   Finish   Cancel   Help

7. See [Table 8](#) for the storage configuration for the given number of virtual machines architecture. Choose storage amount from 300GB SAS Disks drop-down list.

**Figure 51**      **Selecting Storage Disk Type**

**Disk Configuration Wizard**

**Select Storage Type**

Step 3 of 6

Please select the type of disks you want to use for this new pool.

The disks and their storage types have been rated according to their suitability to the selected application / usage.

Rating	Disk Type	Max Capacity	Storage Profile
☆☆☆	SAS	17.299 TB	Balanced Perf/Cap.
☆☆	SAS	10.222 TB	High Performance
☆	EFD	0 GB (None Availab	Best Performance
	NL SAS	0 GB (None Availab	High Capacity

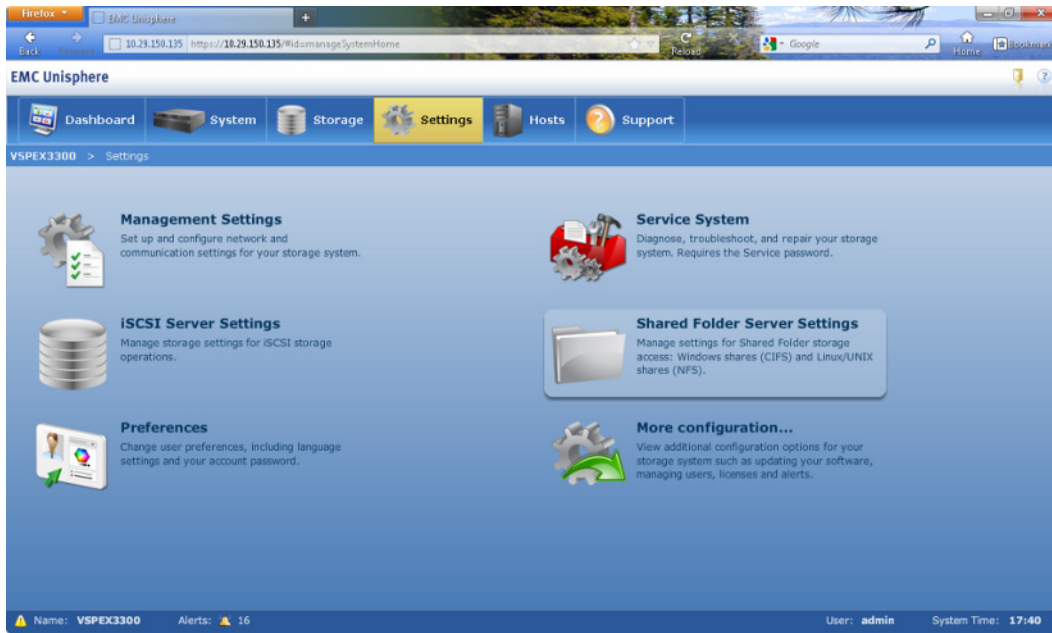
Uses SAS disks to provide a balanced level of storage performance and capacity. This pool type does not offer performance as high as High Performance pools, but it can be adequate for databases with low-to-average performance requirements.

VMware SAS storage pool using RAID 5(6+1).

< Back   Next >   Finish   Cancel   Help

8. Click **Next**. Verify the configuration and click **Finish**, to deploy disk storage. Repeat these steps for two more data stores for the VM data storage.
9. To configure NFS share point, click **Settings** and choose **Shared Folder Server Settings**. Click **Add Shared Folder Server**.

**Figure 52**      **Configuring NFS Share**



10. Provide NFS server name, IP address and related configuration details in the mandatory fields.

**Figure 53**      **Network Interface Details for Shared NFS Server**

**Shared Folder Server**

Step 1 of 4

Specify the Network Interface for the new Shared Folder Server:

Server Name: \* NFS-Share

IP Address: \* 10 . 10 . 40 . 51

Subnet Mask: \* 255 . 255 . 255 . 0

Gateway: \* 10 . 10 . 40 . 1

[Show advanced](#)

< Back   Next >   Finish   Cancel   Help

11. Check the check box **Linux/Unix shares (NFS)** for the NFS storage. Click **Next** to continue.

**Figure 54**      **Selecting Shared Folder Types**

**Shared Folder Server**

Step 2 of 4

Choose the type of shares the Shared Folder Server supports:

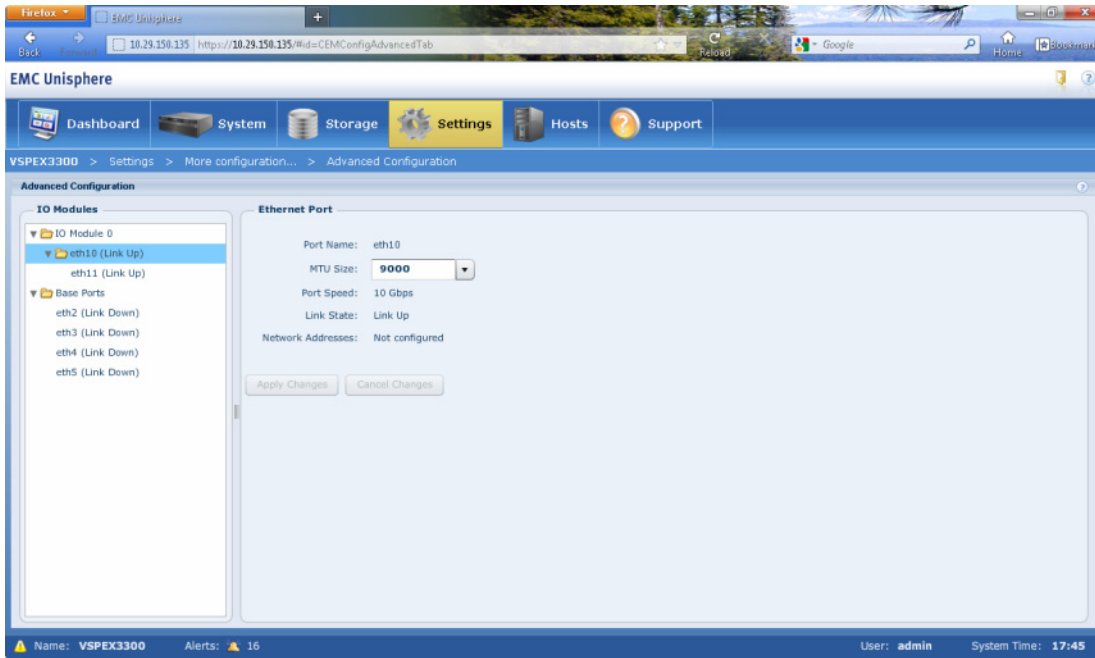
☒ Linux/Unix shares (NFS)

☐ Windows shares (CIFS)

< Back   Next >   Finish   Cancel   Help

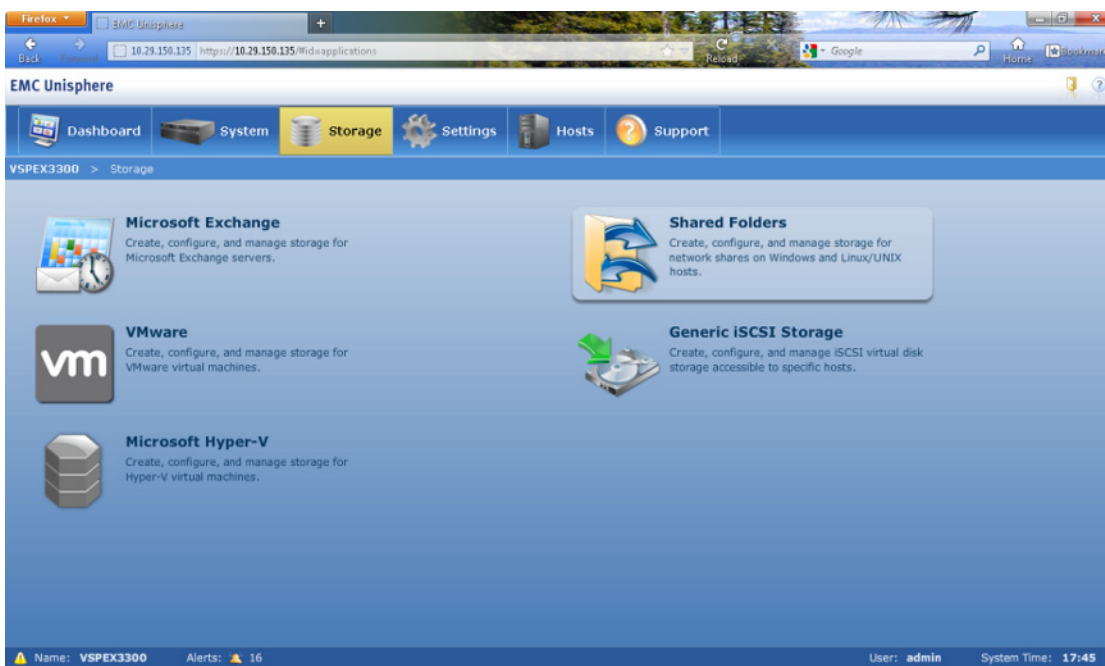
12. For eth10 interface, where the IP address for NFS share is configured, set MTU to 9000 and apply changes. On eth11 interface, aggregate with eth10 interface. This forms the other end of the storage vPC on the Cisco Nexus switches.

**Figure 55** *Setting MTU Size in EMC Unisphere*



13. Click **Storage** in the EMC Unisphere window and then choose “Shared Folders”. Click **Create**.

**Figure 56** *Shared Folder to Configure Shared Storage*



14. Specify the shared folder name and provide description in the “Name” and “Description” fields respectively. This name is accessible through NFS share on the NFS IP address provided in the previous step.

**Figure 57**      **Shared Folder Details in Shared Folder Wizard**

Shared Folder Wizard

**Specify Shared Folder Name**

Step 1 of 8

Specify a name and optional description for the shared folder:

Name: \* OS-Storage

Description: NFS share for Virtual Machine OS datastore

< Back   Next >   Finish   Cancel   Help

15. Configure the shared folder storage on the disks configured in the steps 3 to 8. Same as disks, create three shared folders: one for VM OS and two for VM data.

**Figure 58**      **Configuring Shared Folder Storage**

**Shared Folder Wizard**

**Configure Shared Folder Storage**

Step 2 of 8

Configure the storage pool and size for this shared folder:

Type	Pool	Server	Available	Percent Used	Subscription
vm	OS-Datastore	NFS-Share	1.518 TB	<div><div></div></div> 0%	0%
vm	VM-Datastore	NFS-Share	11.873 TB	<div><div></div></div> 0%	0%

Percent Used: ■ Percent Available: ■ Alert Threshold: |

Size: \*  TB ▼

Thin: ☐ Enabled

< Back   **Next >**   Finish   Cancel   Help

16. Click “Linux/Unix shares (NFS)” radio button for the NFS store. Click **Next** to continue.

**Figure 59**      **Configuring Shared Folder Attributes**

**Shared Folder Wizard**

**Configure Shared Folder Attributes**

Step 3 of 8

Configure the type of shares which will be exported from this shared folder:

Share Type: ☐ Windows shares (CIFS)

The selected storage server does not have CIFS support enabled. You can enable this feature from the Shared Folder Server Settings page.

☒ **Linux/Unix shares (NFS)**

NFS shares are used to share content in Linux/Unix environments.

[Show advanced](#)

< Back   **Next >**   Finish   Cancel   Help

17. To Configure the NFS share, check the check box Create an NFS share. Specify the name and description.

**Figure 60**      **Configuring NFS Share**

Shared Folder Wizard

### Configure NFS Share

Step 4 of 9

Configure the share to be created for this shared folder:

Local Path: /OS-Storage/

☒ Create an NFS share:

Name: OS-Share

Export Path: NFS-Share:/OS-Share

Description: Datastore for Virtual Machines OS datastore

< Back   Next >   Finish   Cancel   Help

**18.** Click **Add ESX Host** to add servers.



**Figure 61**      **Configuring Host Access**

Shared Folder Wizard

## Configure Host Access

Step 5 of 9

Configure host access for the NFS share:

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

Name	Network Address	Access

Create New Host   Add ESX Host   Create New Subnet   Create New Netgroup

< Back   **Next >**   Finish   Cancel   Help

19. Click **Next** to continue.

**Figure 62**      **Adding ESX Host**

**Find ESX Hosts**

**Welcome to the Add ESX Hosts Wizard**

Step 1 of 4

This wizard helps you integrate Virtual Machines from ESX Hosts that use storage served from this System. The steps are:

1. List ESX Hosts (if any) managed by vCenter servers, list ESX Hosts (if any) that are not managed by vCenter servers.

When you click Finish, these ESX Hosts will appear either in your list of hosts, in their virtualization page, or, if they are managed by vCenter servers, in both locations.

Click Next to continue.

< Back   **Next >**   Finish   Cancel   Help

20. Provide IP address of vCenter server and click **Find**. Check all the ESXi hosts check boxes and click **Next**.

**Figure 63**      **Finding vCenters and Manage ESX Hosts**

**Find ESX Hosts**

**Discover vCenters/ESX Hosts**

Step 2 of 4

Click Find to discover vCenters, managed/un-managed ESX Hosts.

☐ vCenters  

☒ Other (vCenter or ESX Host)  

**Find**

Discovered vCenter servers/ESX Hosts:

Name	Type
10.29.150.166	VMware vCenter 5.0.0
<input checked="" type="checkbox"/> V100-ESXHost3	VMware ESX
<input checked="" type="checkbox"/> V100-ESXHost4	VMware ESX
<input checked="" type="checkbox"/> V100-ESXHost1	VMware ESX
<input checked="" type="checkbox"/> V100-ESXHost2	VMware ESX
<input type="checkbox"/> V100-InfraESX	VMware ESX

< Back   **Next >**   Finish   Cancel   Help

21. Default access for all hosts is Read/Write, allow Root and set the same for all the hosts.

**Figure 64**      **Setting Default Access to Hosts**

**Shared Folder Wizard**

**Configure Host Access**

Step 5 of 9

Configure host access for the NFS share:

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

Name	Network Address	Access
V100-ESXHost3	10.29.150.163	Use Default Access
V100-ESXHost4	10.29.150.164	Use Default Access
V100-ESXHost1	10.29.150.161	Use Default Access
V100-ESXHost2	10.29.150.162	Use Default Access

22. Click **Next**, verify the configuration and click **Finish**. This will take you to the previous wizard, where default protection is “Configure protection storage”.

**Figure 65**      **Choosing Protection Storage Type**

Shared Folder Wizard

## Configure Protection

Step 6 of 9

Configure protection storage for replication and snapshots:

- ☐ **Do not configure protection storage for this storage resource.**  
Replication and snapshots can be supported by allocating protection space at a later time.
- ☒ **Configure protection storage, do not configure a snapshot protection schedule.**  
An automated snapshot protection schedule may be configured at a later time.
- ☐ **Configure protection storage, protect data using snapshot schedule:** Default Protection ▾  
This schedule will create snapshots:  
Every day at 01:00, keep for 2 days

Note: Times are displayed in Local Time (UTC-0700) in 24-hour format

< Back   Next >   Finish   Cancel   Help

23. Change protection by clicking the radio button “Do not configure protection storage for this storage resource”. If you need additional protection, then additional storage would be required. For more information, see the *EMC VNXe storage configuration guide*.

**Figure 66**      **Removing Additional Protection**

**Shared Folder Wizard**

**Configure Protection**

Step 6 of 9

Configure protection storage for replication and snapshots:

- ☒ **Do not configure protection storage for this storage resource.**  
Replication and snapshots can be supported by allocating protection space at a later time.
- ☐ **Configure protection storage, do not configure a snapshot protection schedule.**  
An automated snapshot protection schedule may be configured at a later time.
- ☐ **Configure protection storage, protect data using snapshot schedule:** Default Protection ▼  
This schedule will create snapshots:  
Every day at 01:00, keep for 2 days

Note: Times are displayed in Local Time (UTC-0700) in 24-hour format

< Back   **Next >**   Finish   Cancel   Help

24. Click **Next**, choose all the defaults in next window and click **Finish**. Repeat these steps for VM data store shared folders. [Figure 67](#) shows that the NFS data stores are successfully configured.

**Figure 67**      **Shared Folder Storage**

**EMC Unisphere**

VSPEX3300 > Storage > Shared Folders

**Shared Folder Storage**

Allocated Shared Folders:

	Name	Replication	Description	Protocol	Storage Server	Size	Protection Sch...	Deduplication
✓	OS-Storage	Not Replicated	NFS share for Virtual	NFS	NFS-Share	1.517 TB		Disabled
✓	VM-Share1	Not Replicated	NFS share (1) for VM	NFS	NFS-Share	5.500 TB		Disabled
✓	VM-Share2	Not Replicated	NFS share (2) for VM	NFS	NFS-Share	5.500 TB		Disabled

Selected: 1      Items: 3

Create   Add Share   Create a Replication Destination   Details   Refresh   Delete

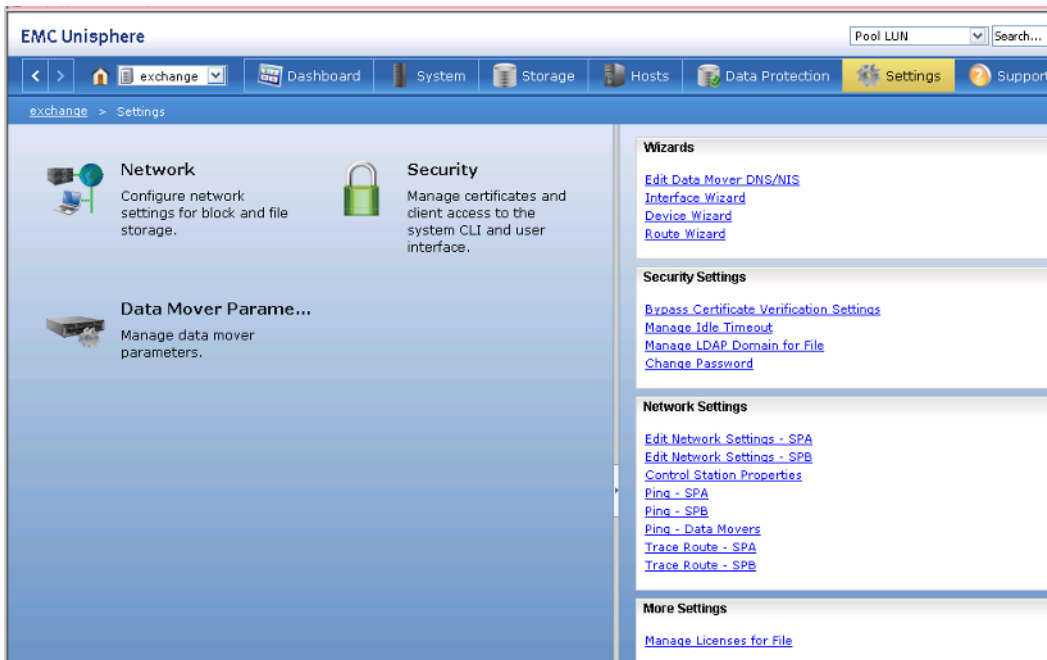
Name: VSPEX3300   Alerts: 20   User: admin   System Time: 11:35

## VNX Configuration

For 125 virtual machines architecture, EMC VNX5300 storage array is used. Next steps explain how to configure VNX5300 for 125 virtual machines architecture.

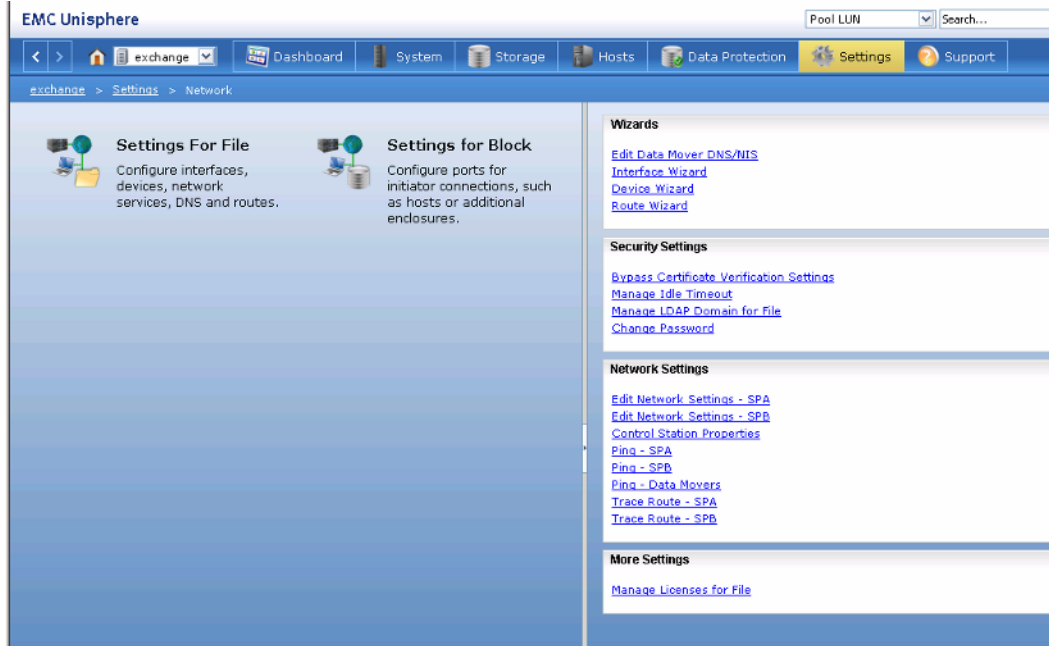
1. Using the browser, launch the Unisphere Console with the management IP address of the storage array. Provide user name and password.
2. To create LACP (Link Aggregation Control Protocol) Network interface, click **Settings** menu and then choose **Network**.

**Figure 68** *Settings Window in EMC Unisphere*



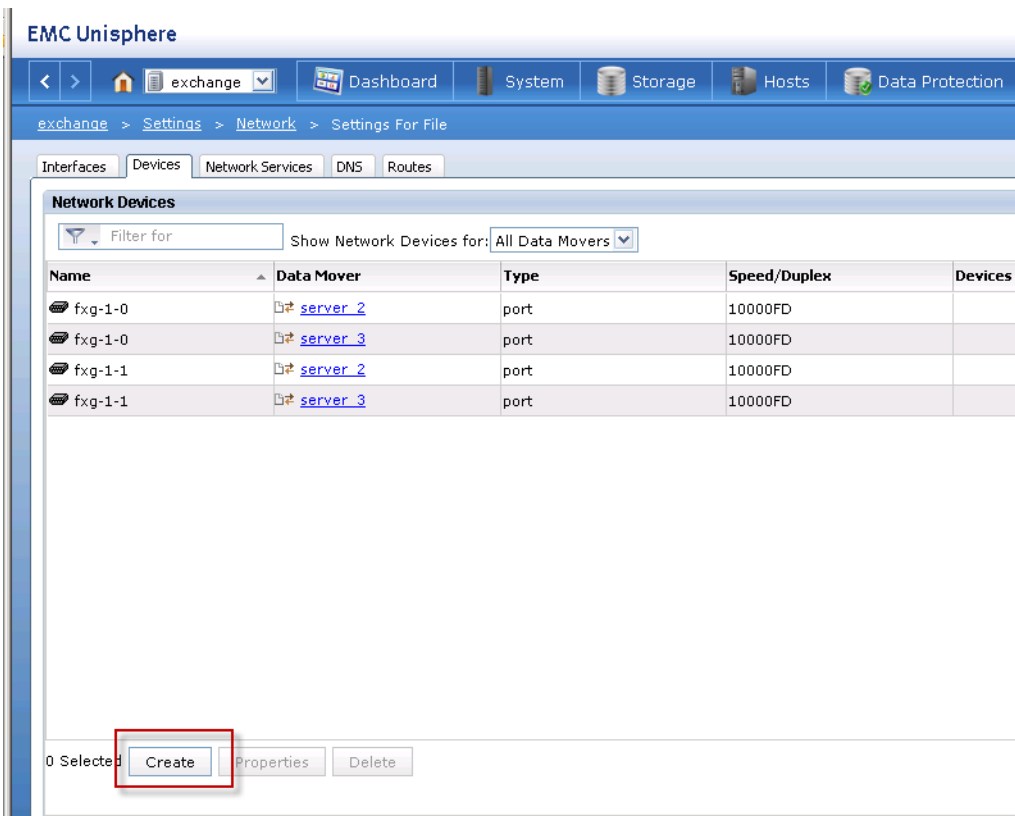
3. Choose **Settings for File** option.

**Figure 69 Settings for Files**



4. Choose **Devices** tab and click **Create** in the Network Devices window.

**Figure 70 Creating Network Devices**



5. Choose **All Primary Data Movers** option from the Data Movers drop-down list and click “Link Aggregation” radio button. In the Device Name field specify the device name as “LACP-1” and check both the 10 Gigabit ports check boxes as highlighted in the [Figure 71](#) and click **OK**.

**Figure 71**      **Device Details**

**Data Mover:** All Primary Data Movers

**Type:**

- ☐ Ethernet Channel
- ☒ Link Aggregation
- ☐ Fail Safe Network

**Device Name:** LACP-1

**10/100 ports:** None available

**Gigabit ports:** None available

**10/100/1000 ports:** None available

**10 Gigabit ports:** ☒ fxg-1-0 ☒ fxg-1-1

**Speed/Duplex:**

**Buttons:** OK Apply Cancel Help

6. [Figure 72](#) shows the creation of LACP Network device name as “LACP-1”.



**Figure 72**      **Created LACP Network Device Details**

EMC Unisphere

exchange > 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
fxg-1-0	<a href="#">server_2</a>	port	10000FD	
fxg-1-0	<a href="#">server_3</a>	port	10000FD	
fxg-1-1	<a href="#">server_2</a>	port	10000FD	
fxg-1-1	<a href="#">server_3</a>	port	10000FD	
LACP-1	<a href="#">server_2</a>	lacp	10000FD	fxg-1-0,fxg-1-1

1 Selected Create Properties Delete

7. Choose **Interfaces** tab and click **Create**.

**Figure 73**      **Network Interface Details**

Address	Name	Netmask	Data Mover	Device	State
128.221.252.2	el30	255.255.255.0	<a href="#">server_2</a>	<a href="#">mge0</a>	Up
128.221.252.3	el30	255.255.255.0	<a href="#">server_3</a>	<a href="#">mge0</a>	Up
128.221.253.2	el31	255.255.255.0	<a href="#">server_2</a>	<a href="#">mge1</a>	Up
128.221.253.3	el31	255.255.255.0	<a href="#">server_3</a>	<a href="#">mge1</a>	Up

0 Selected   [Create](#)   [Properties](#)   [Delete](#)   Filtered:

- Choose `server_2` from the Data Mover drop-down list “`server_2`” and choose “`LACP-1`” as the device name from the Device Name drop-down list.

**Figure 74**      **Creating Network Interface**

Data Mover: [server\\_2](#)

Device Name: [LACP-1](#)

Address:

Name:

Netmask:

Broadcast Address: 0.0.0.0

MTU:

VLAN ID:

[OK](#)   [Apply](#)   [Cancel](#)   [Help](#)

- Specify the IP address and subnet mask in the Address and Netmask fields respectively for the Network Interface and Enter the name of the Network interface as “`fs01`” in the Name field. Specify MTU value as “`9000`”. Click **OK**.

**Figure 75**      **Creating Network Interface**

**Data Mover:** server\_2

**Device Name:** LACP-1

**Address:** 10.10.40.111

**Name:** fs01

**Netmask:** 255.255.255.0

**Broadcast Address:** 10.10.40.255

**MTU:** 9000

**VLAN ID:**

OK Apply Cancel Help

10. Figure 76 shows the creation of Network Interface “fs01” for the LACP device “LACP-1”

**Figure 76**      **Creating Network Interface for LACP Device**

EMC Unisphere

Pool LUN Search...

exchange > Settings > Network > Settings For File

Interfases

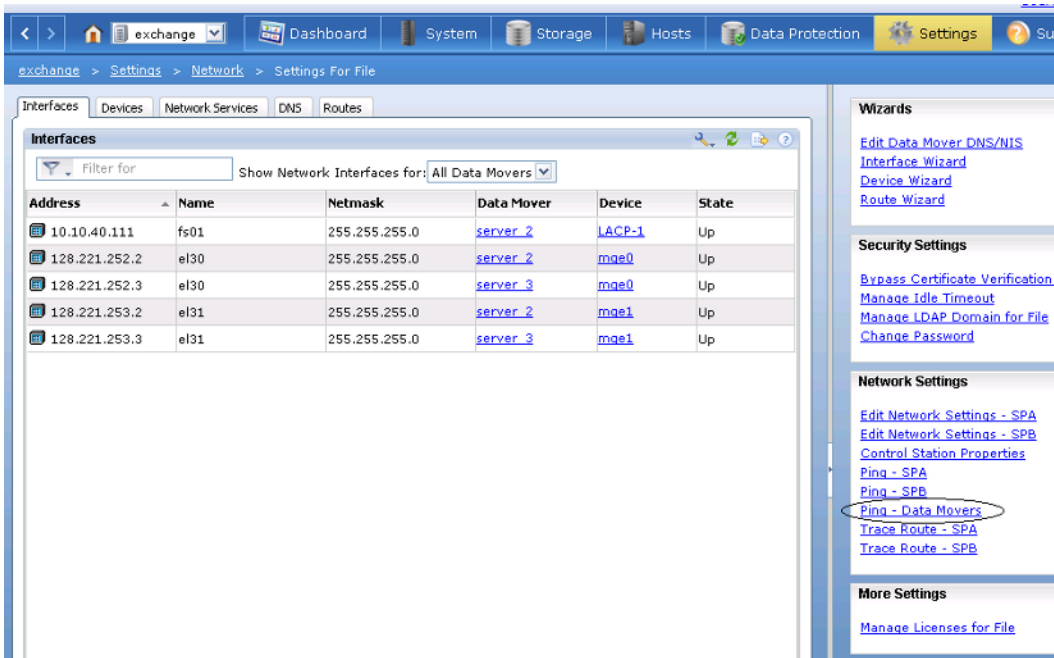
Filter for Show Network Interfaces for: All Data Movers

Address	Name	Netmask	Data Mover	Device	State
10.10.40.111	fs01	255.255.255.0	server_2	LACP-1	Up
128.221.252.2	el30	255.255.255.0	server_2	mae0	Up
128.221.252.3	el30	255.255.255.0	server_3	mae0	Up
128.221.253.2	el31	255.255.255.0	server_2	mae1	Up
128.221.253.3	el31	255.255.255.0	server_3	mae1	Up

1 Selected Create Properties Delete Filtered: 5 of


11. To verify the network connectivity between LACP Network Interface of Data mover & VMkernel IP of ESXi host. Select “Ping – Data Movers” as highlighted in Figure 77.

**Figure 77** Showing Network Interfaces for All Data Movers



- Choose "server\_2" from the Data mover drop-down list.

**Figure 78** Choose and Check Each Data Mover



- Choose the newly created network Interface "10.10.40.111" from the drop-down menu and enter the VMkernel IP of ESXi host "10.10.40.101" on the Destination text box. Click **OK**.

**Figure 79**      **Entering Network Details to Verify Connection**

14. Ensure that the destination VMkernel IP is alive.
15. To create Storage Pools for NFS Datastore, Choose **Storage > Storage Configuration > Storage pools**. In the Pools window click **Create**.

**Figure 80**      **Creating Storage Pools**

16. Specify Storage Pool Name as “PerformancePool” and choose RAID type as RAID5 from the drop-down list. Then, Select the required SAS disks from the drop-down list as shown in [Figure 81](#).

**Figure 81**      **Storage Pool Parameters**

exchange - Create Storage Pool

General   Advanced

**Storage Pool Parameters**

Storage Pool Type: ☒ Pool ☐ RAID Group

☒ Scheduled Auto-Tiering

Storage Pool ID: 0

Storage Pool Name: PerformancePool

RAID Type: RAID5

Number of Disks:

**Performance**

SAS Disks

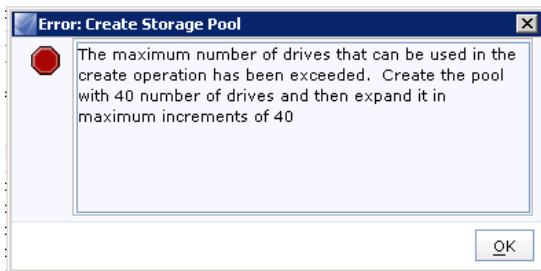
75 (Recommend...)

**Distribution**

Performance : 40260.571 GB (100.00%)

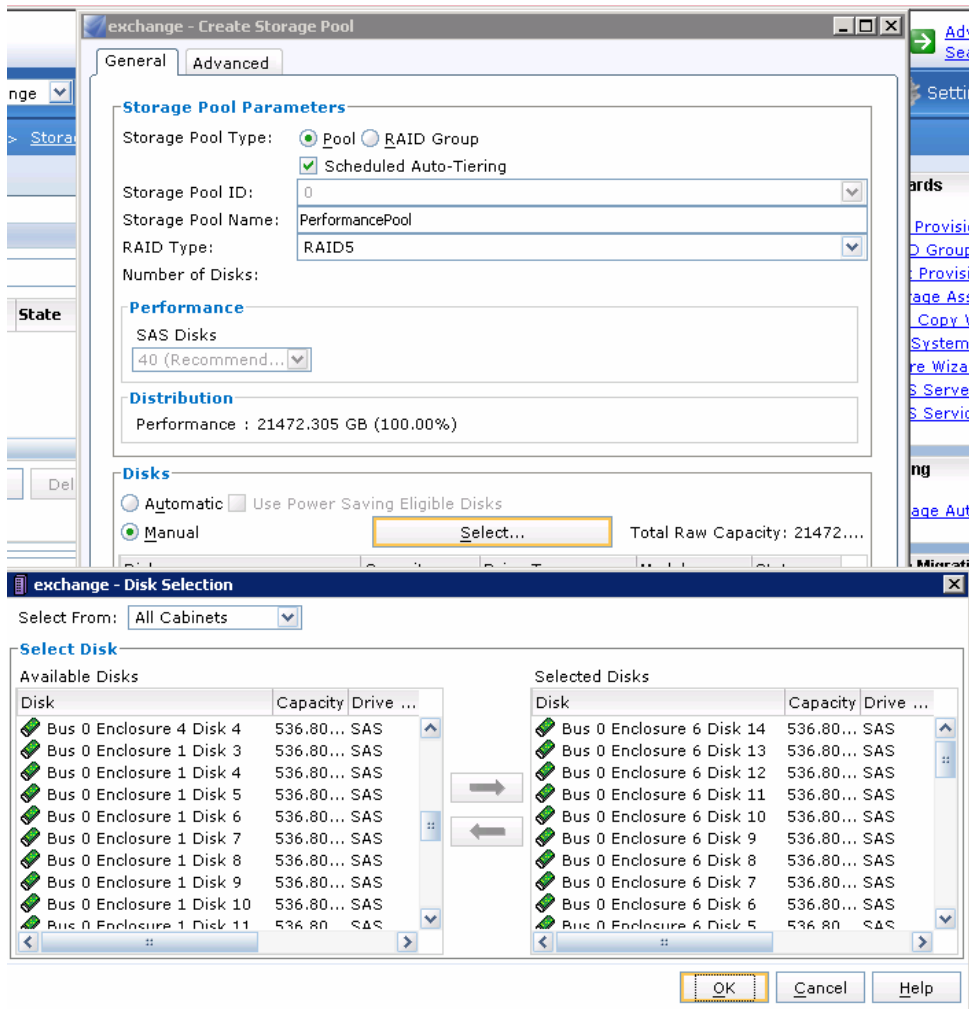
**Note**

VNX5300 does not support more than 40 drives during storage pool creation. In order to choose 75 disks for the given storage pool, create the pool with 40 drives and then expand it with 35 drives.

**Figure 82**      **Error in Creating Storage Pool**

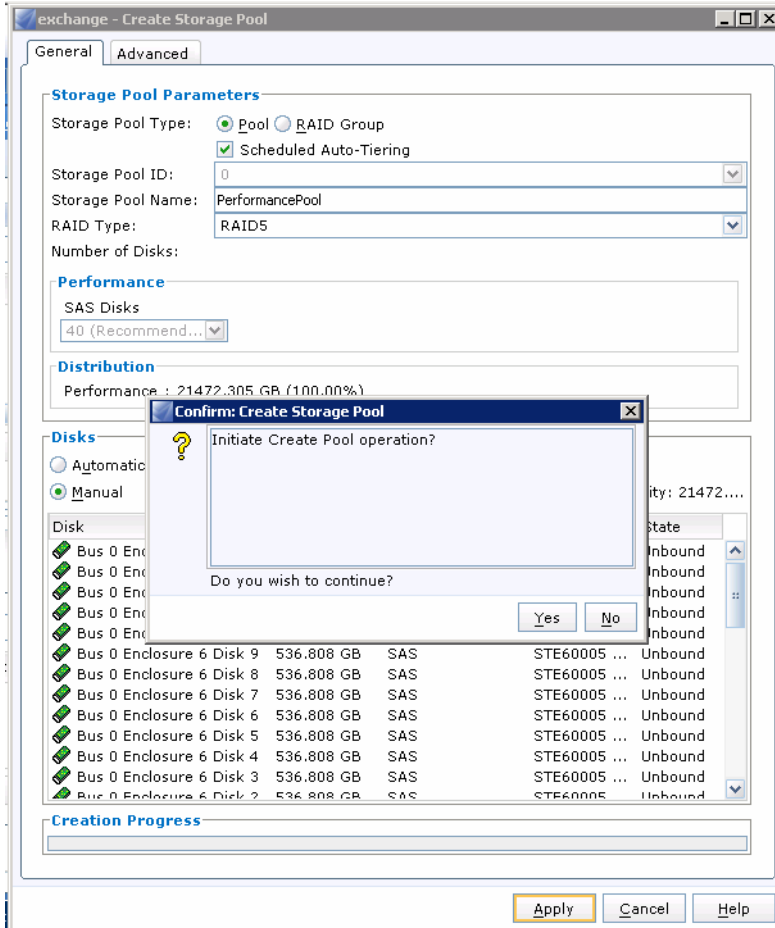
17. Manually select 40 disks from the SAS Disks drop-down list and click **OK**.

**Figure 83 Choose SAS Disks Manually**



18. Click Yes to confirm the storage pool operation.

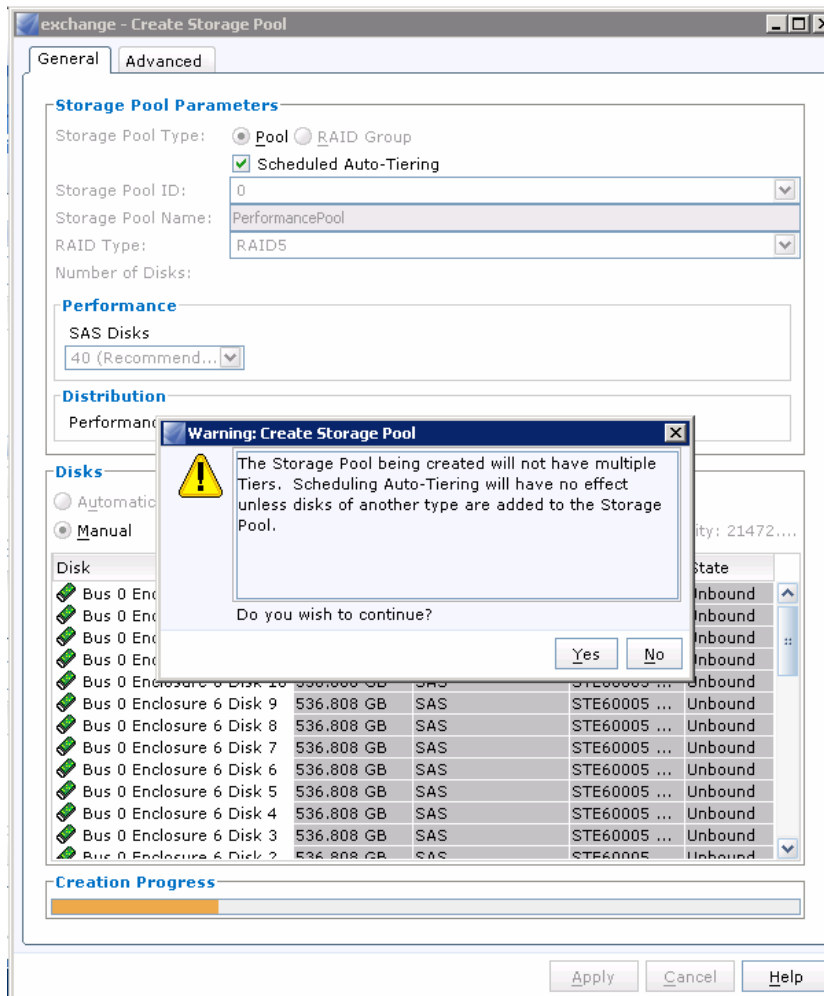
**Figure 84**      **Confirming Storage Pool Creation**



19. Click **Yes** to continue the storage pool creation.

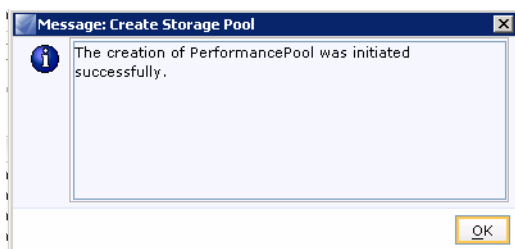


**Figure 85**      **Warning Message on Scheduled Auto Tiering**

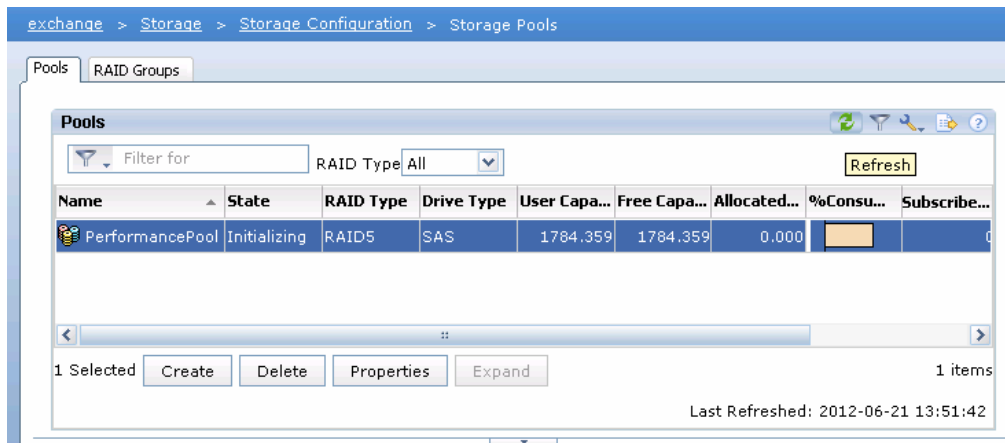


20. Click **OK** on the pop-up window on successful creation of the Performance Pool.

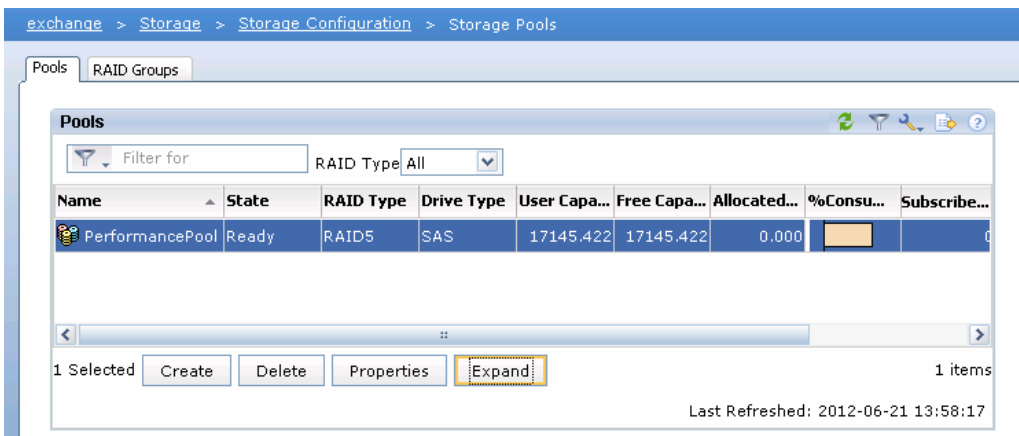
**Figure 86**      **Completion of Storage Pool Creation**



21. Select "Performance Pool" and click **Refresh** until the state shows "Ready".


**Figure 87 State of the Created Storage Pool**

22. Ensure that the performance pool state is changed from Initializing to “Ready”. Click **Expand** button.

**Figure 88 Storage Pool State Showing Ready**

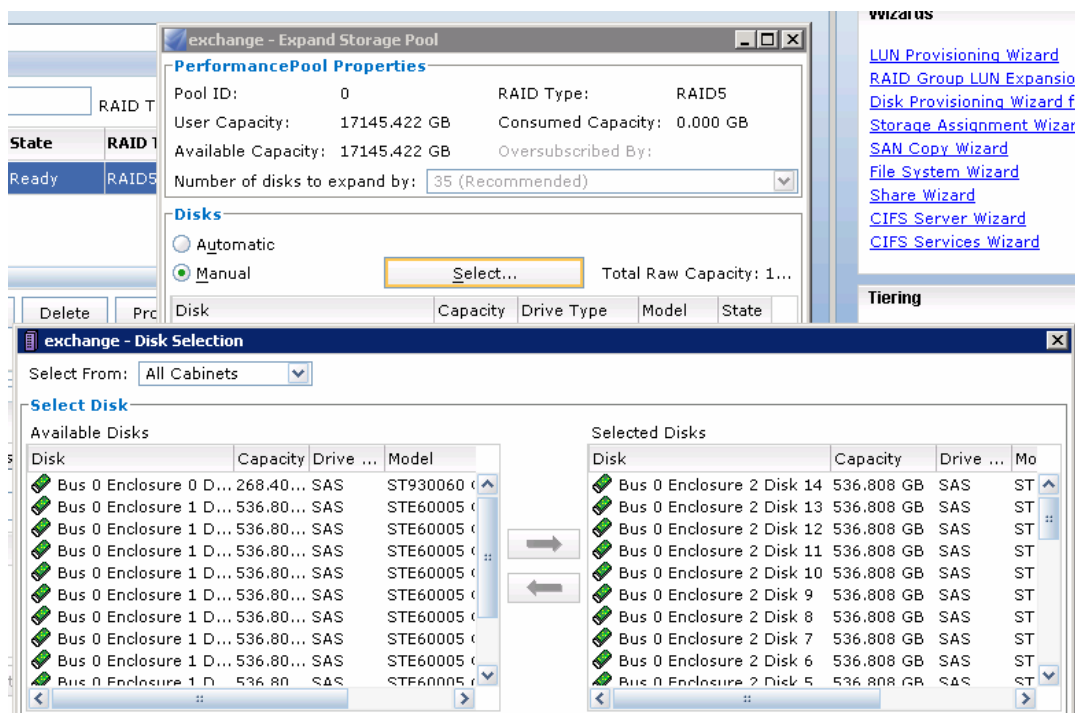
23. In the “Expand Storage Pool window”, you can add remaining set of disk to the pool.  
 24. In the “Expand Storage Pool Window”, click the radio button “Manual” and click **Select** button.

**Figure 89** *Selecting Disks Manually in Expand Storage Pool Window*



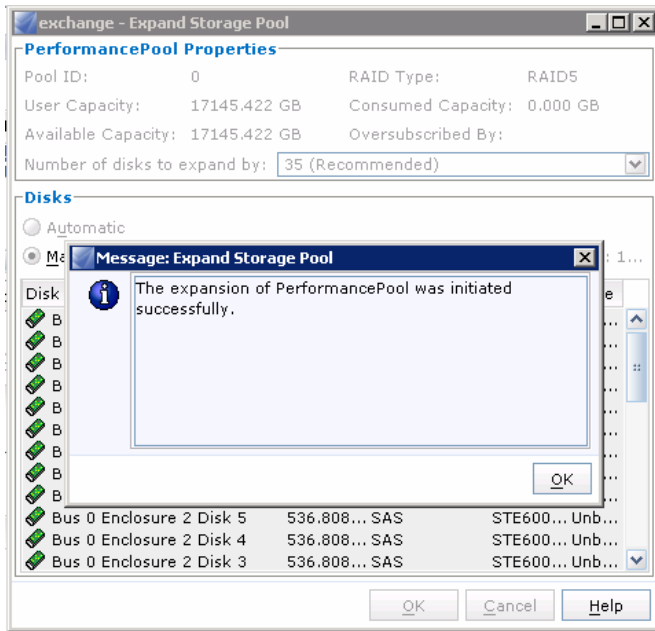
25. Select remaining available drives and drag them left to right pane for selected drives.

**Figure 90** *Adding Available Drives to Storage Pool*



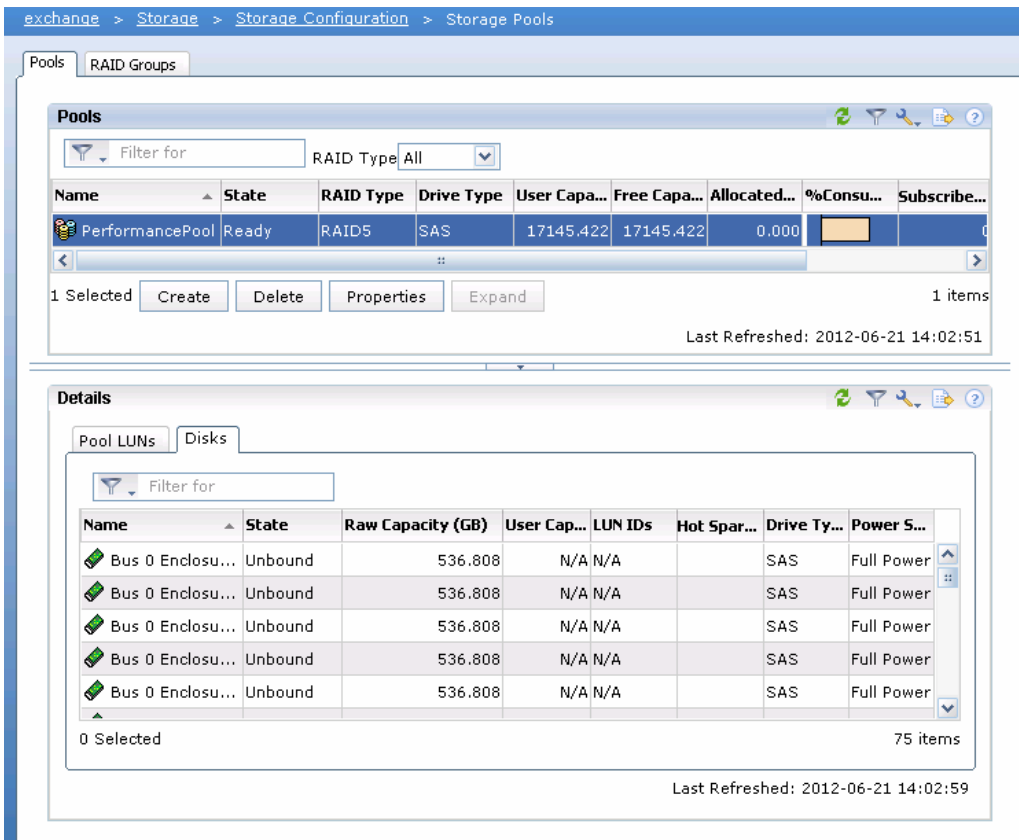
26. Then click **Yes** on the popup, **Yes** on the No Multi-Tier available, and **OK** on the success popup for expansion of the pool initiation.

**Figure 91** Window Showing Successful Addition of Disks



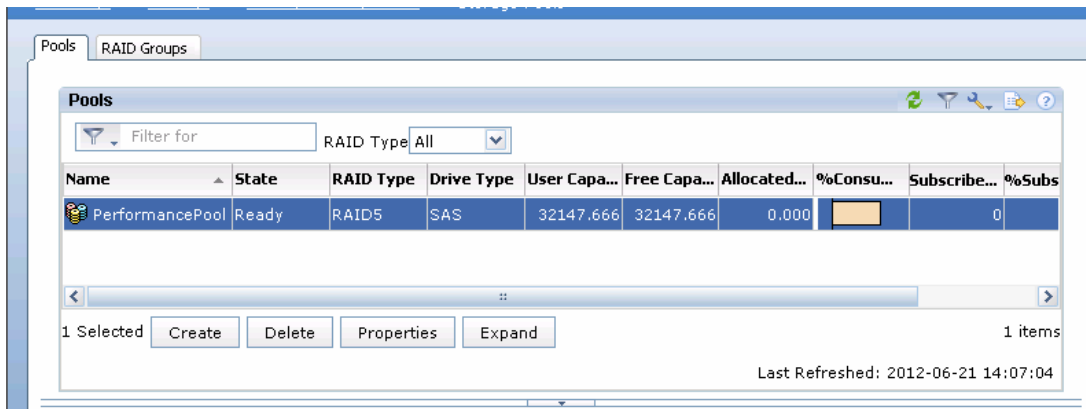
27. Wait for the expansion of the pool to be completed and state as Ready.

**Figure 92** Storage Pool Details



28. The performance pool is ready to use after the completion of the pool expansion.

**Figure 93** Performance Pool State Showing Ready



29. To create Hot Spares for the system. Choose **System > Hardware > Hot Spares** in the EMC Unisphere window. Click **Create**.

**Figure 94** Configuring Hot Spares



30. In the create Spare window, click “RAID Group” radio button for the Storage Pool Type. Choose storage Pool ID as 0, Storage Pool name as RAID Group 0, RAID type as Hot Spare, Number of Disks as 1. Click “Automatic” radio button in the “Disks” pane and click **Apply**.

**Figure 95** Storage Pool Parameters in Create Hot Spare Window

The screenshot shows the 'exchange - Create Hot Spare' window with the 'General' tab selected. The 'Storage Pool Parameters' section includes the following fields:

- Storage Pool Type: ☐ Pool ☒ RAID Group
- Storage Pool ID: 0
- Storage Pool Name: RAID Group 0
- RAID Type: Hot Spare
- Number of Disks: 1

The 'Disks' section includes the following options and a table:

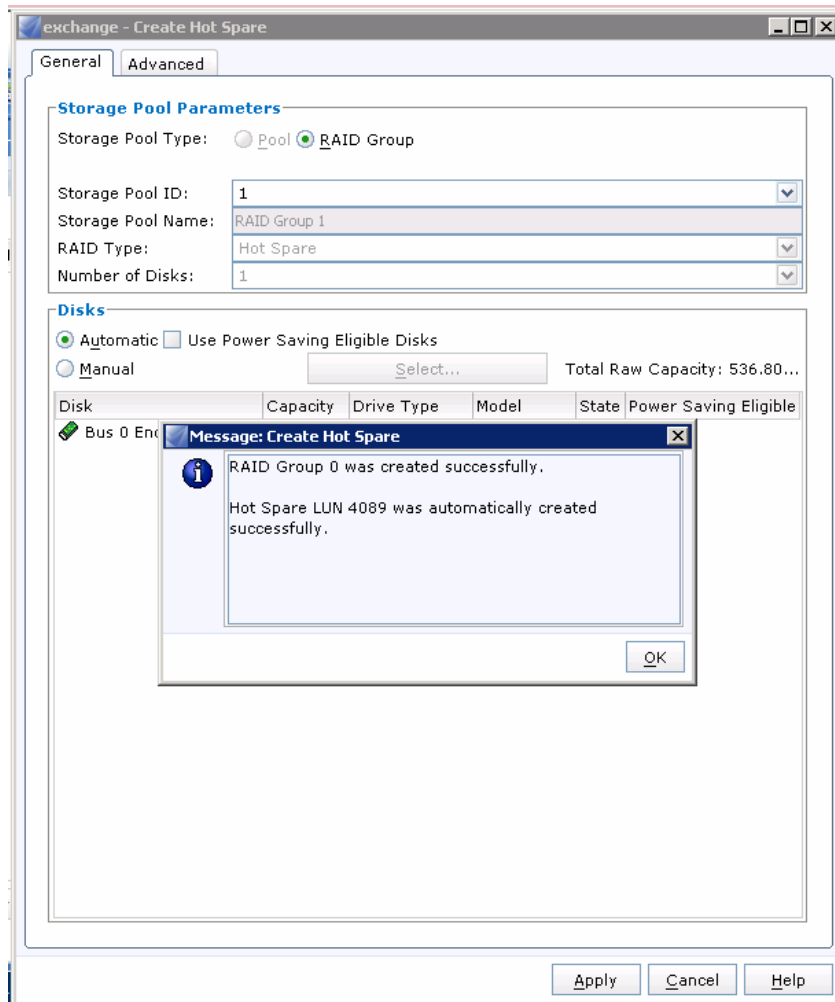
- ☒ Automatic ☐ Use Power Saving Eligible Disks
- ☐ Manual  Total Raw Capacity: 536.80...

Disk	Capacity	Drive Type	Model	State	Power Saving Eligible
Bus 0 Enclosure 1 D...	536.808...	SAS	STE60005...	Un...	No

At the bottom of the window are buttons for 'Apply', 'Cancel', and 'Help'.

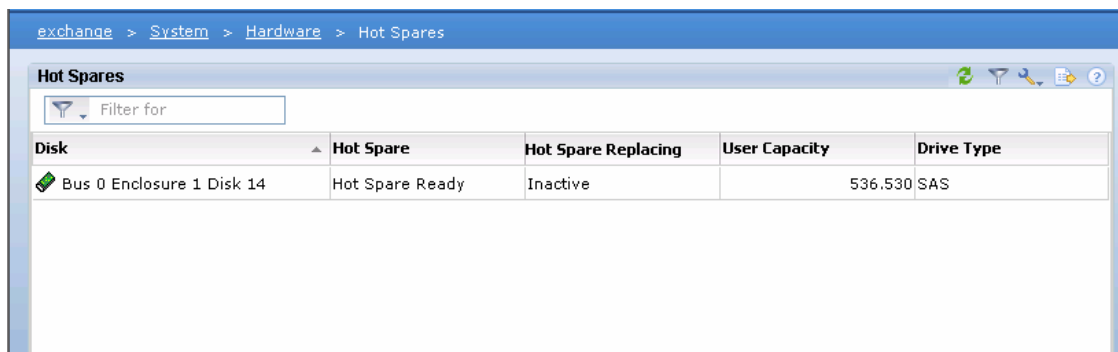
31. Follow procedure in the step 30 to create hot spares as needed. [Figure 96](#) shows the RAID Group 0 has been created successfully. And initiation of RAID Group 1.

**Figure 96** Window Showing Successful Creation of RAID Group



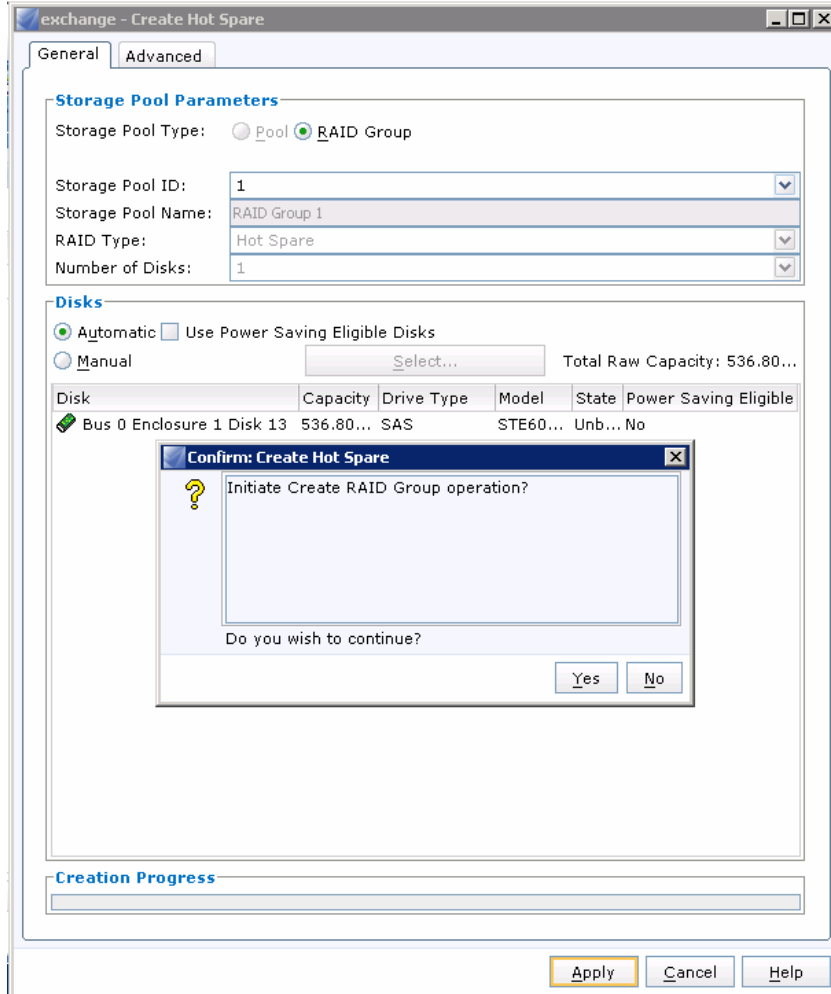
32. Figure 97 shows drive created successfully for Hot Spare window.

**Figure 97** Drive Created for Hot Spare



33. Click **Yes** to Initiate creation of RAID Group operation to create Hot Spare by following procedure in the step 30.

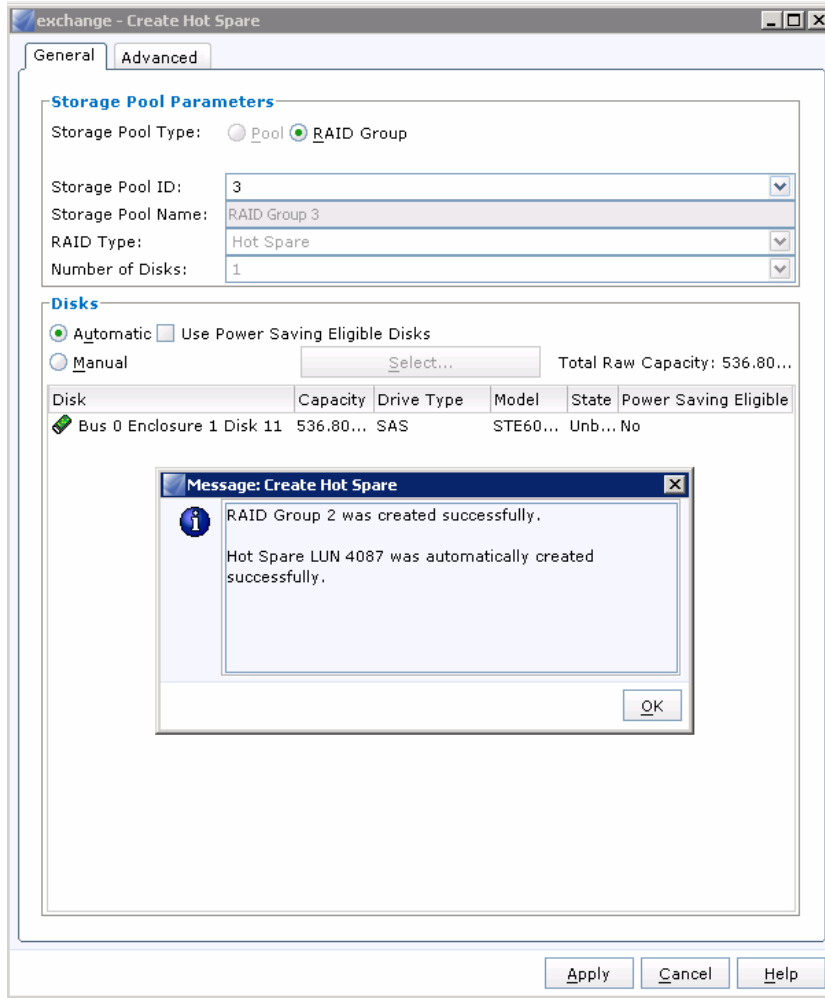
**Figure 98** Confirmation for Creating RAID Group Operation



34. Repeat the step 30 to create Hot Spares.

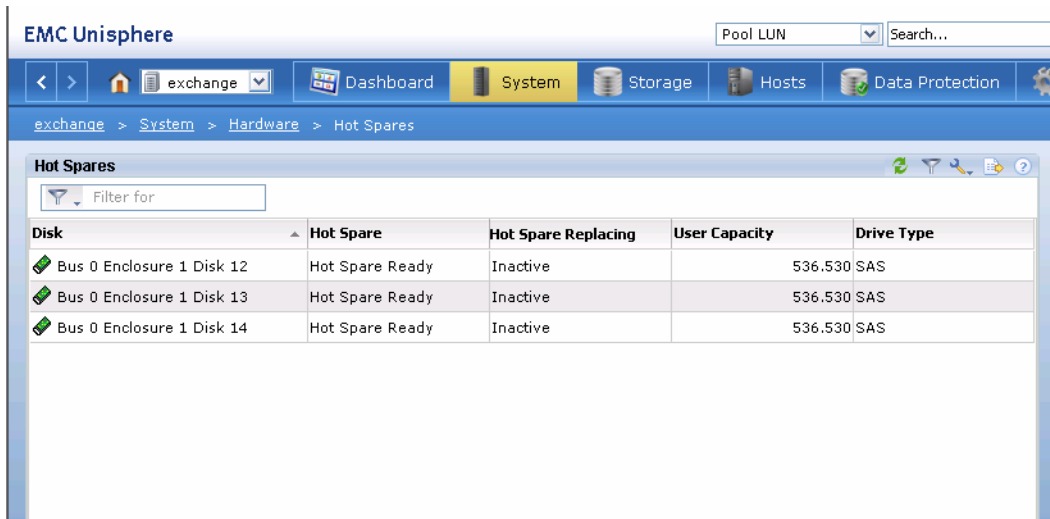


**Figure 99** Window Showing Successful Creation of RAID Group



35. When the Hot Spares are created successfully and ensure that the Hot Spare state is Ready.

**Figure 100** Window Showing Hot Spare Status



36. To create LUNs for storage pools; choose Storage. Right click on new pool created and click **Create LUN**.

**Figure 101**      *Creating LUN in Storage Pools*

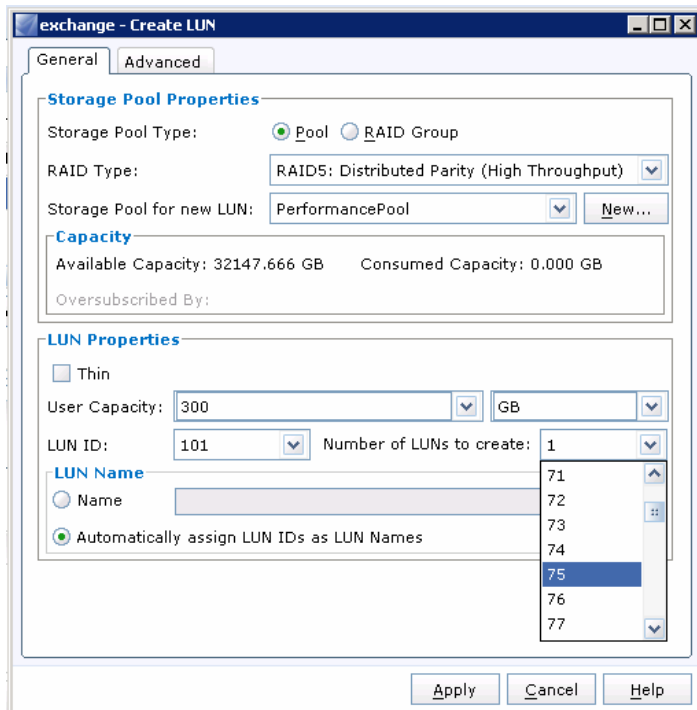
The screenshot shows the EMC Unisphere web interface. At the top, the 'Storage' tab is selected in the navigation bar. The breadcrumb trail is 'exchange > Storage > Storage Configuration > Storage Pools'. The 'Pools' section is active, showing a table of storage pools. The 'PerformancePool' is in a 'Ready' state. A right-click context menu is open over the pool, with 'Create LUN' selected. Below the table, there are buttons for 'Create', 'Delete', 'Analyzer', 'Auto-Tiering', and 'Properties'. The 'Details' pane at the bottom shows the 'Pool LUNs' tab, which is currently empty. The 'Usage' dropdown is set to 'ALL User LUNs'. The table has columns for Name, ID, State, User Capacity (GB), Current Owner, and Host Information. At the bottom of the details pane, there are buttons for 'Delete', 'Properties', and 'Add to Storage Group'. The status '0 Selected' and 'Filtered: 0 of 0' are displayed.

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

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
------	----	-------	--------------------	---------------	------------------

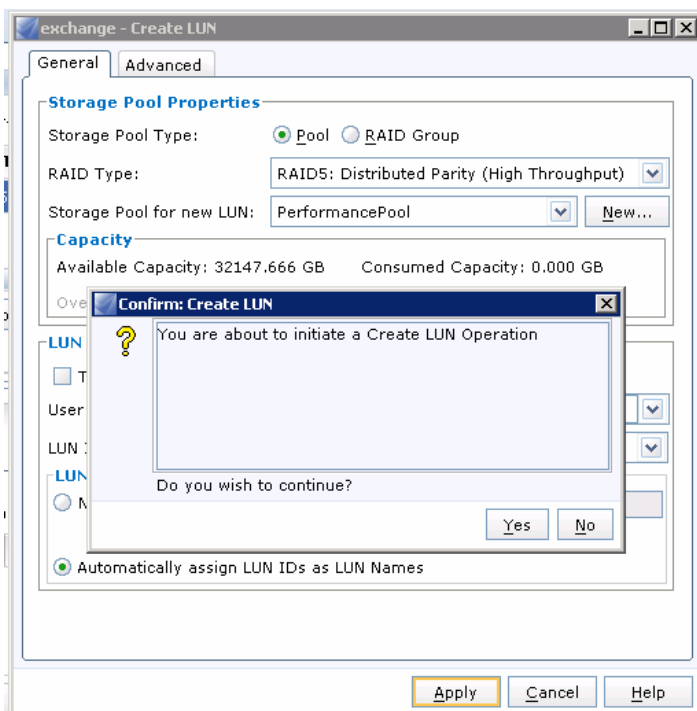
37. In the Storage Pool Properties pane, click “Pool” radio button for Storage Pool Type, choose Type of RAID and Storage Pool for New LUN as shown in [Figure 102](#). In the LUN Properties pane, choose User Capacity as max capacity of drive; in this case it is 300GB. Choose Number of LUNs to Create as 75; as there are 75 drives. In the LUN Name pane, click “automatically assign LUN IDs as LUN Names” radio button. Click **Apply** to initiate process of creating LUNs.

**Figure 102** *Choosing Properties for Storage Pool and LUN*

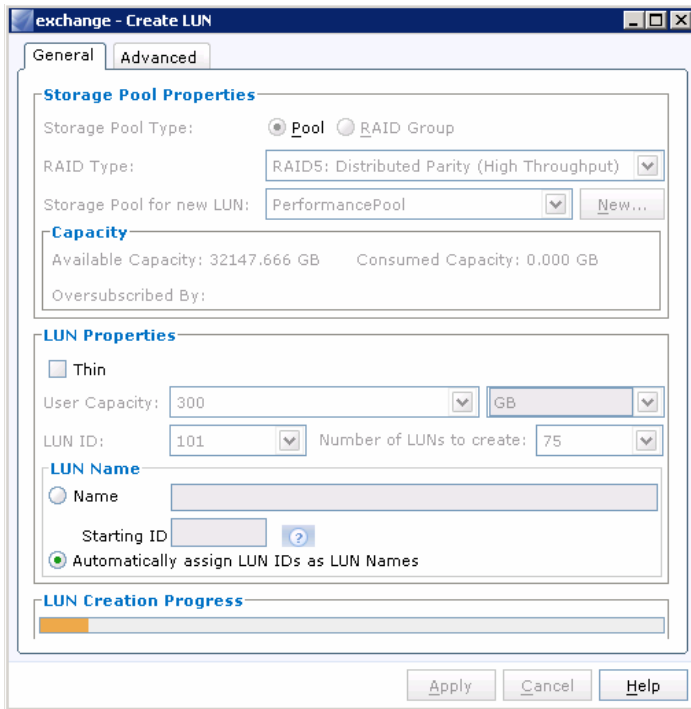


38. Click **Yes** to continue the initiation operation to create LUNs.

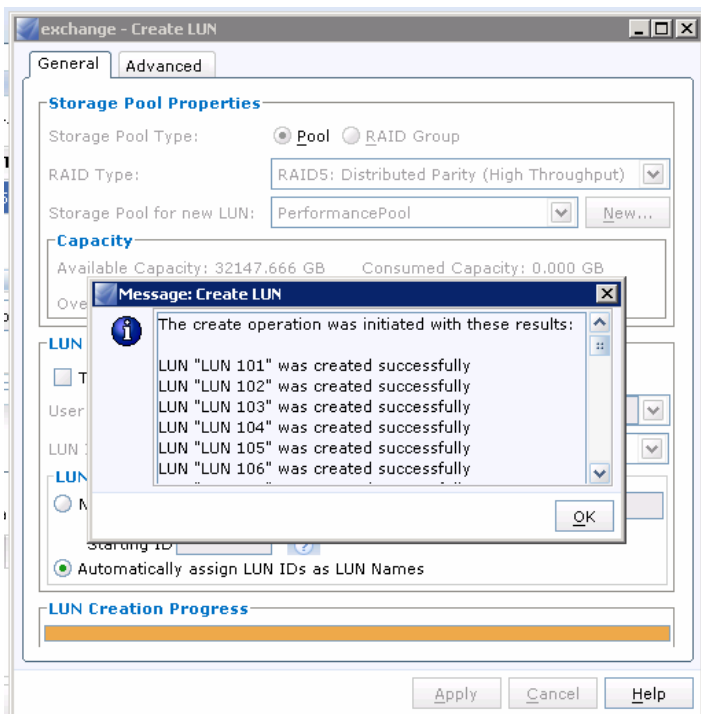
**Figure 103** *Confirmation for Creating LUN Operation*



39. Progress in the LUN creation process is shown by the progress indicator.

**Figure 104** Window Showing LUN Creation in Progress

40. Wait for the acknowledge of the task to be complete. Click **OK**.

**Figure 105** Window Showing LUNs Created Successfully

41. Select the Pool used to create LUNs. Choose **Pool LUNs** tab for Pool LUNs tab in “Details” pane.

**Figure 106** Pool LUNs for Selected Storage Pool

exchange > 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...	%St
PerformancePool	Ready	RAID5	SAS	32147.666	8852.256	23295.410		23,220,264	

1 Selected Create Delete Properties Expand 1 items

Last Refreshed: 2012-06-21 14:27:26

---

**Details**

Pool LUNs Disks

Filter for Usage ALL User LUNs

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
LUN 101	101	Ready	300.000	SP A	
LUN 102	102	Ready	300.000	SP B	
LUN 103	103	Ready	300.000	SP A	
LUN 104	104	Ready	300.000	SP B	
LUN 105	105	Ready	300.000	SP A	

0 Selected Delete Properties Add to Storage Group Filtered: 75 of 75

Last Refreshed: 2012-06-21 14:27:27

42. Select all LUNs created and click **Add to Storage Group**.

**Figure 107 Adding LUNs to Storage Group**

The screenshot displays the EMC Unisphere interface for managing storage pools and LUNs. The top navigation bar includes links for Dashboard, System, Storage, Hosts, and Data Protection. The main content area is divided into two sections: 'Pools' and 'Details'.

**Pools Section:**

- Filter for: [Empty]
- RAID Type: All
- Table with columns: Name, State, RAID Type, Drive Type, User Capa..., Free Capa..., Allocated..., %Consu..., Subscribe..., %St.
- Table content:
 

Name	State	RAID Type	Drive Type	User Capa...	Free Capa...	Allocated...	%Consu...	Subscribe...	%St
PerformancePool	Ready	RAID5	SAS	32147.666	8852.256	23295.410		23,220.264	
- Buttons: Create, Delete, Properties, Expand
- 1 items
- Last Refreshed: 2012-06-21 14:27:26

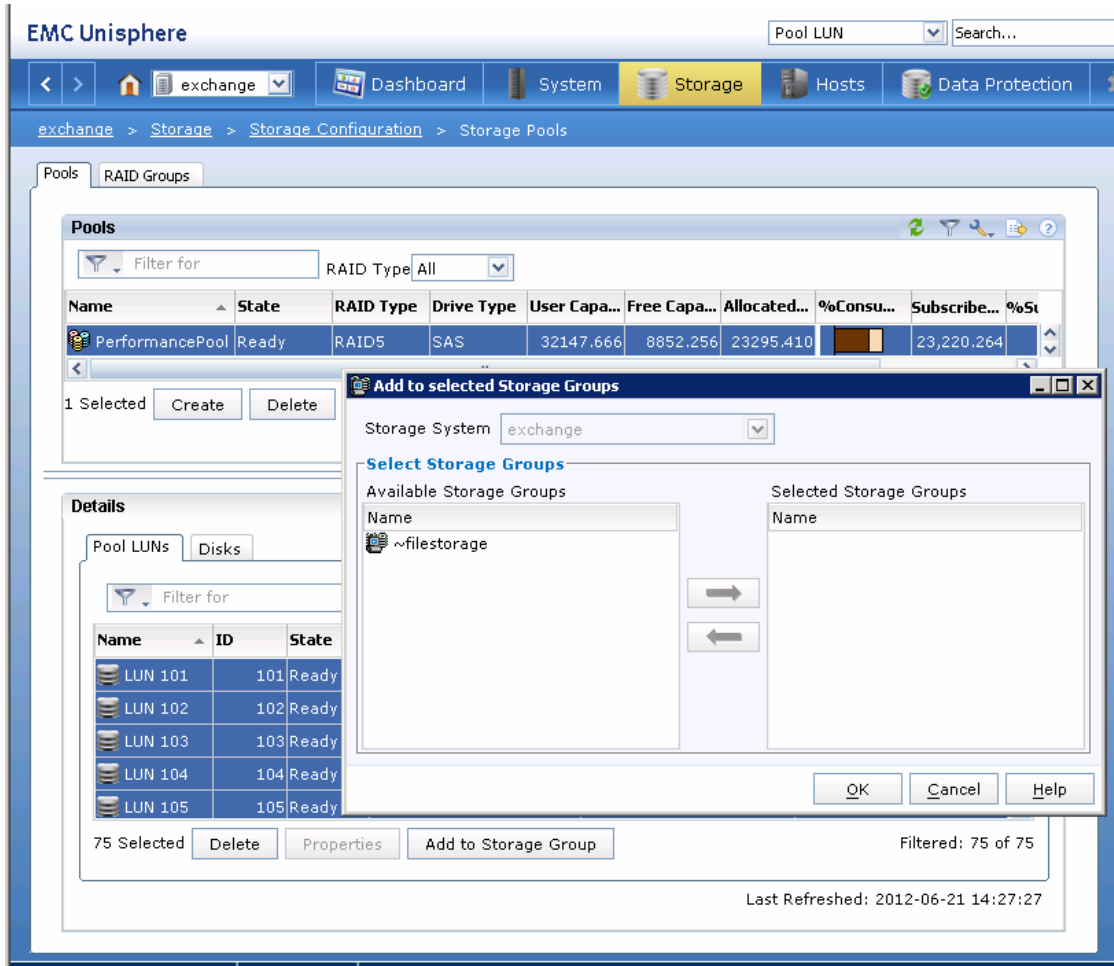
**Details Section:**

- Filter for: [Empty]
- Usage: ALL User LUNs
- Table with columns: Name, ID, State, User Capacity (GB), Current Owner, Host Information
- Table content:
 

Name	ID	State	User Capacity (GB)	Current Owner	Host Information
LUN 101	101	Ready	300.000	SP A	
LUN 102	102	Ready	300.000	SP B	
LUN 103	103	Ready	300.000	SP A	
LUN 104	104	Ready	300.000	SP B	
LUN 105	105	Ready	300.000	SP A	
- Buttons: Delete, Properties, Add to Storage Group
- 75 Selected
- Filtered: 75 of 75
- Last Refreshed: 2012-06-21 14:27:27

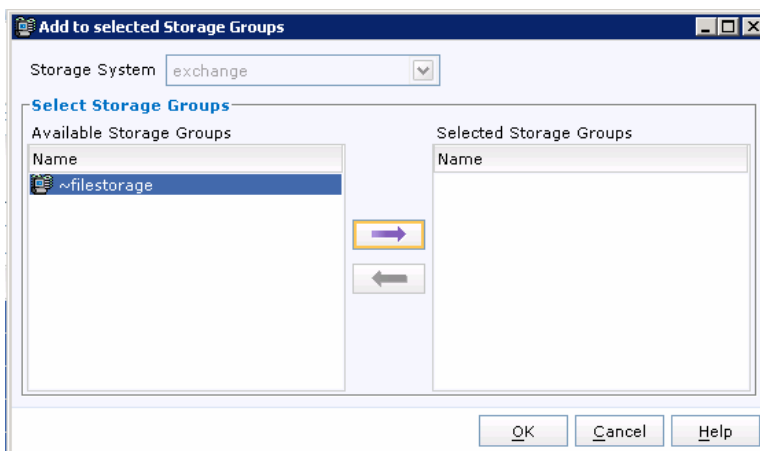
43. In the Select Storage Groups pane, select “~filestorage” as the available storage in the Storage Groups.

**Figure 108**      **Selecting Storage Groups from Available Storage Groups**



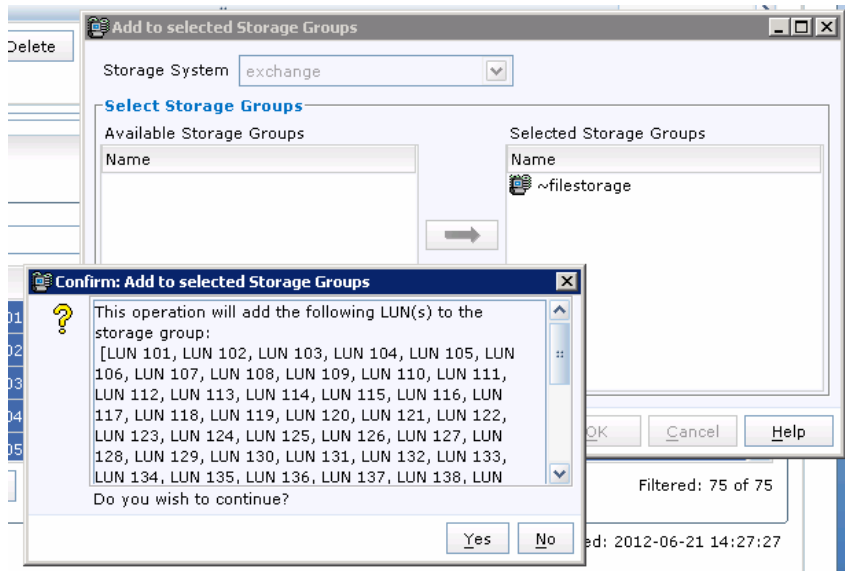
44. In the Select Storage Groups pane, select the available storage “filestorage” and add it to Selected Storage Groups as shown in the [Figure 109](#).

**Figure 109**      **Adding Storage Groups**



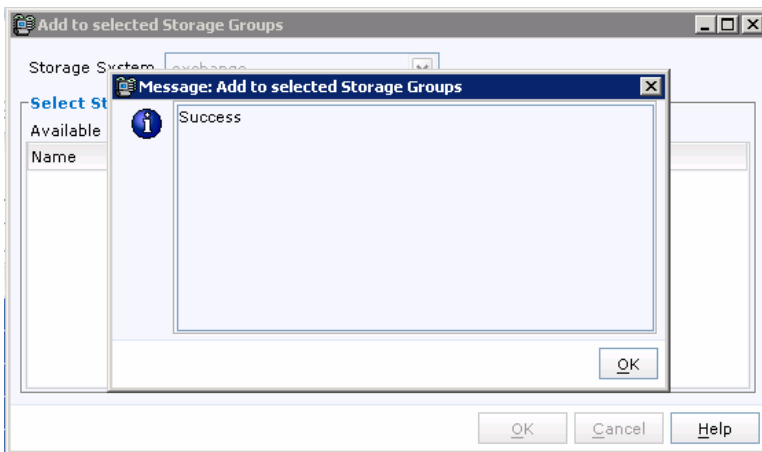
45. Click **Ok**. Click **Yes** in the pop-up window to confirm the operation to add LUNs to the storage group.

**Figure 110** Confirmation for Adding Selected Storage Groups



46. Click **Ok** in the pop-up window showing successful completion of the task.

**Figure 111** Window Showing Successful Addition of Storage Groups



47. Ensure that the LUNs are added to the storage group in the Details pane.



**Figure 112 Adding LUNs to Storage Group**

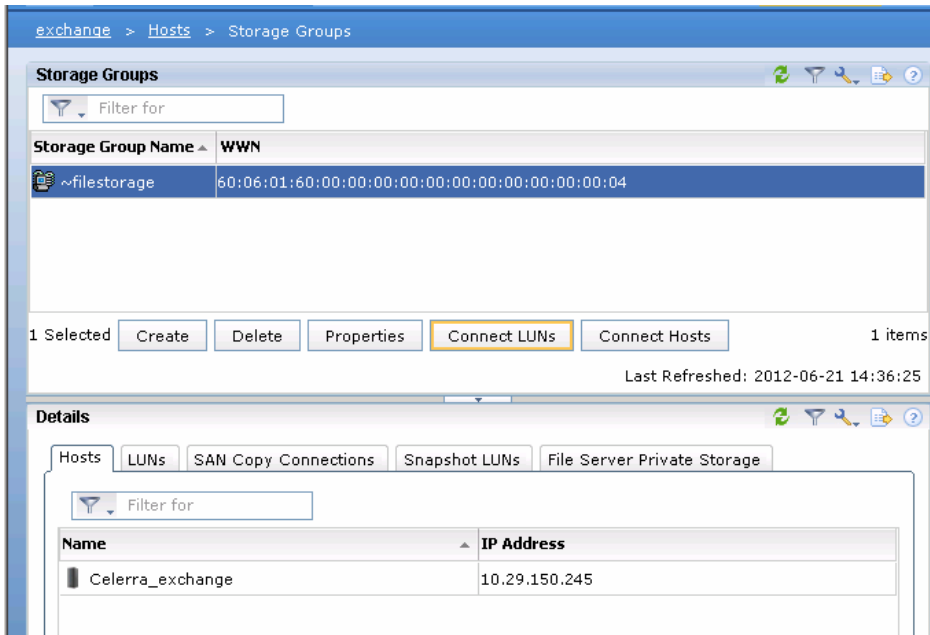
The screenshot shows the EMC Unisphere interface. The top navigation bar includes 'exchange > Storage > Storage Configuration > Storage Pools'. The 'Pools' section is active, showing a table with columns: Name, State, RAID Type, Drive Type, User Capacity, Free Capacity, Allocated Capacity, % Consumed, Subscribed Capacity, and % Subscribed. The table contains one entry: 'PerformancePool' with State 'Ready', RAID Type 'RAID5', Drive Type 'SAS', User Capacity '32147.666', Free Capacity '8852.256', Allocated Capacity '23295.410', % Consumed '23,220.264', and % Subscribed '23,220.264'. Below the table are buttons: 'Create', 'Delete', 'Properties', and 'Expand'. The 'Details' section is also visible, showing a table with columns: Name, ID, State, User Capacity (GB), Current Owner, and Host Information. The table contains four entries: 'LUN 101', 'LUN 102', 'LUN 103', and 'LUN 104'. Below the table are buttons: 'Delete', 'Properties', and 'Add to Storage Group'. The 'Add to Storage Group' button is highlighted.

48. To assign Host ID to LUNs created; choose **Hosts** tab > **Storage Groups**.

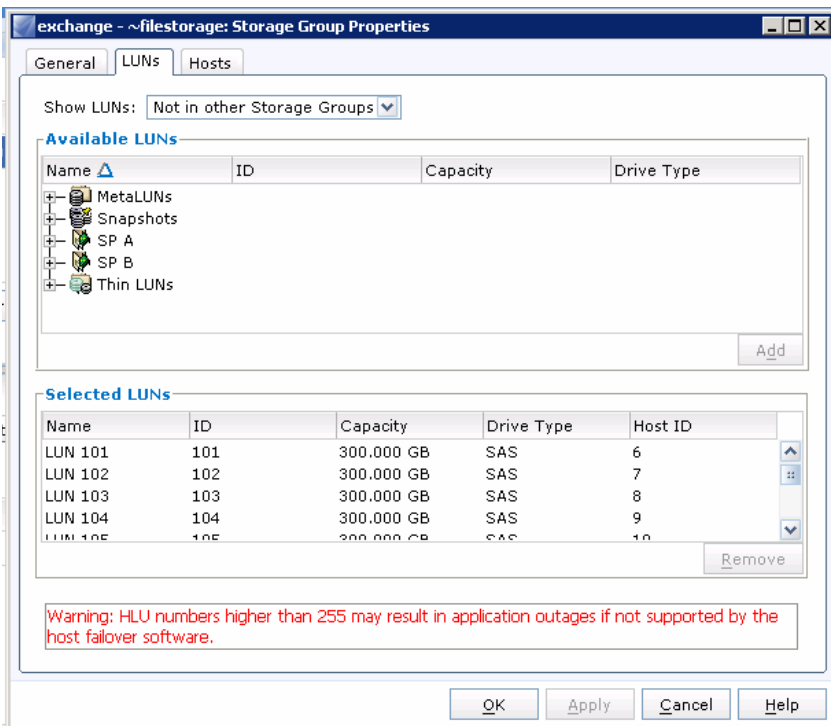
**Figure 113 Selecting Storage Groups in EMC Unisphere**

The screenshot shows the EMC Unisphere interface. The top navigation bar includes 'exchange > Hosts'. The 'Hosts' section is active, showing a table with columns: Name, ID, State, User Capacity (GB), Current Owner, and Host Information. The table contains four entries: 'LUN 101', 'LUN 102', 'LUN 103', and 'LUN 104'. Below the table are buttons: 'Delete', 'Properties', and 'Add to Storage Group'. The 'Add to Storage Group' button is highlighted.

49. Select Storage Group name where all the LUNs were added, “~filestorage” in this case. Click **Connect LUNs**.

**Figure 114**      **Connecting LUNs to Storage Group**

50. Ensure all the LUNs that are part of the filestorage group are shown in the “Selected LUNs” pane.

**Figure 115**      **Verify LUNs in the Storage Group**

51. Choose **Storage** tab > **Storage Configuration** > **Volumes** in the EMC Unisphere window. Verify that the volumes are created.

**Figure 116**      **Verifying the Created Volumes**

Volume	Type	Uses Volumes	Used by	Storage Capacity (MB)	Storage Used
d3	disk		md3	1,990	
d4	disk		md4	1,990	
d5	disk		md5	1,996	
d6	disk		md6	63,990	
md3	meta	d3	root fs d3	1,990	
md4	meta	d4	root fs d4	1,990	
md5	meta	d5	root fs d5	1,996	
md6	meta	d6	root fs d6	63,990	

52. To verify the IP address go to VNX cmd line through ssh with the IP used for configuration of VNX. Alternatively, click **Rescan Storage Systems** under “File Storage” in the right pane of Unisphere GUI.

**Figure 117**      **Verifying IP Address**

```
[root@localhost ~]# ssh nasadmin@10.29.150.245
A customized version of the Linux operating system is used on the
EMC(R) VNX(TM) Control Station. The operating system is
copyrighted and licensed pursuant to the GNU General Public License
("GPL"), a copy of which can be found in the accompanying
documentation. Please read the GPL carefully, because by using the
Linux operating system on the EMC Celerra you agree to the terms
and conditions listed therein.

EXCEPT FOR ANY WARRANTIES WHICH MAY BE PROVIDED UNDER THE TERMS AND
CONDITIONS OF THE APPLICABLE WRITTEN AGREEMENTS BETWEEN YOU AND EMC,
THE SOFTWARE PROGRAMS ARE PROVIDED AND LICENSED "AS IS" WITHOUT
WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT
NOT LIMITED TO, THE IMPLIED MERCHANTABILITY AND FITNESS FOR A
PARTICULAR PURPOSE. In no event will EMC Corporation be liable to
you or any other person or entity for (a) incidental, indirect,
special, exemplary or consequential damages or (b) any damages
whatsoever resulting from the loss of use, data or profits,
arising out of or in connection with the agreements between you
and EMC, the GPL, or your use of this software, even if advised
of the possibility of such damages.

EMC, VNX, Celerra, and CLARIion are registered trademarks or trademarks of
EMC Corporation in the United States and/or other countries. All
other trademarks used herein are the property of their respective
owners.

EMC VNX Control Station Linux release 3.0 (NAS 7.0.50)
nasadmin@10.29.150.245's password:
Last login: Sat Jun 23 10:00:44 2012 from 10.29.180.52
EMC VNX Control Station Linux   wed Jul 27 12:25:40 EDT 2011

    *** slot_0 primary control station ***

[nasadmin@exchange ~]$
```

53. Type the command **nas\_disk -list** in the command line to see the existing dvols.

**Figure 118** List of Existing dvols

```
[nasadmin@exchange ~]$ nas_disk -list
id  inuse  sizeMB  storageID-devID  type  name  servers
1   y      11260  APM00112001158-2007  CLSTD  root_disk  1,2
2   y      11260  APM00112001158-2008  CLSTD  root_ldisk  1,2
3   y       2038  APM00112001158-2009  CLSTD  d3  1,2
4   y       2038  APM00112001158-200A  CLSTD  d4  1,2
5   y       2044  APM00112001158-200B  CLSTD  d5  1,2
6   y      65526  APM00112001158-200C  CLSTD  d6  1,2
```

54. Verify that same volumes are seen in the GUI interface.

**Figure 119** Verifying the Created Volumes

name	Type	Uses Volumes	Used by	Storage Capacity (...)	Storage Used
d3	disk		md3	1,990	
d4	disk		md4	1,990	
d5	disk		md5	1,996	
d6	disk		md6	63,990	
md3	meta	d3	root fs d3	1,990	
md4	meta	d4	root fs d4	1,990	
md5	meta	d5	root fs d5	1,996	
md6	meta	d6	root fs d6	63,990	

55. Type the command **nas\_diskmark -mark -all** to get info messages only and not error messages.

**Figure 120** Storage Configuration Information

```
[nasadmin@exchange ~]$ nas_disk -list
id  inuse  sizeMB  storageID-devID  type  name  servers
1   y      11260  APM00112001158-2007  CLSTD  root_disk  1,2
2   y      11260  APM00112001158-2008  CLSTD  root_ldisk  1,2
3   y       2038  APM00112001158-2009  CLSTD  d3  1,2
4   y       2038  APM00112001158-200A  CLSTD  d4  1,2
5   y       2044  APM00112001158-200B  CLSTD  d5  1,2
6   y      65526  APM00112001158-200C  CLSTD  d6  1,2

[nasadmin@exchange ~]$ nas_diskmark -mark -all
Discovering storage on exchange (may take several minutes)
```

56. Type the command **nas\_disk -list** to see new dvols that are not in use at this point.

**Figure 121** Verifying the Created Volumes with New dvols

Name	Type	Uses Volumes	Used by	Storage Capacity (...)	Storage Us
d3	disk		md3	1,990	
d4	disk		md4	1,990	
d5	disk		md5	1,996	
d6	disk		md6	63,990	
d117	disk			299,999	
d118	disk			299,999	
d119	disk			299,999	
d120	disk			299,999	
d121	disk			299,999	
d122	disk			299,999	
d123	disk			299,999	
d124	disk			299,999	
d125	disk			299,999	
d126	disk			299,999	

57. To create different file systems choose **Storage** tab > **Storage Configuration** > **File Systems** and click **Create**.

**Figure 122** Creating File System

Name	Storage Capacity (GB)	Storage Used	Data Movers
------	-----------------------	--------------	-------------

58. Click the “Storage Pool” radio button for “Create from”. Specify File System Name, specify Storage Pool, Storage Capacity, Data Mover. Click **OK**.

**Figure 123** *Entering Details for Creating File System*

Create from	<input checked="" type="radio"/> Storage Pool <input type="radio"/> Meta Volume
File System Name:	NFS-OS
Storage Pool:	PerformancePool 22.0 TB (23039925 MB)
Storage Capacity:	1 TB
Auto Extend Enabled:	<input type="checkbox"/>
Thin Enabled:	<input type="checkbox"/>
Slice Volumes:	<input checked="" type="checkbox"/>
Deduplication Enabled:	<input type="checkbox"/>
Data Mover (R/W):	server_2
Mount Point:	<input checked="" type="radio"/> Default <input type="radio"/> Custom

OK Apply Cancel Help

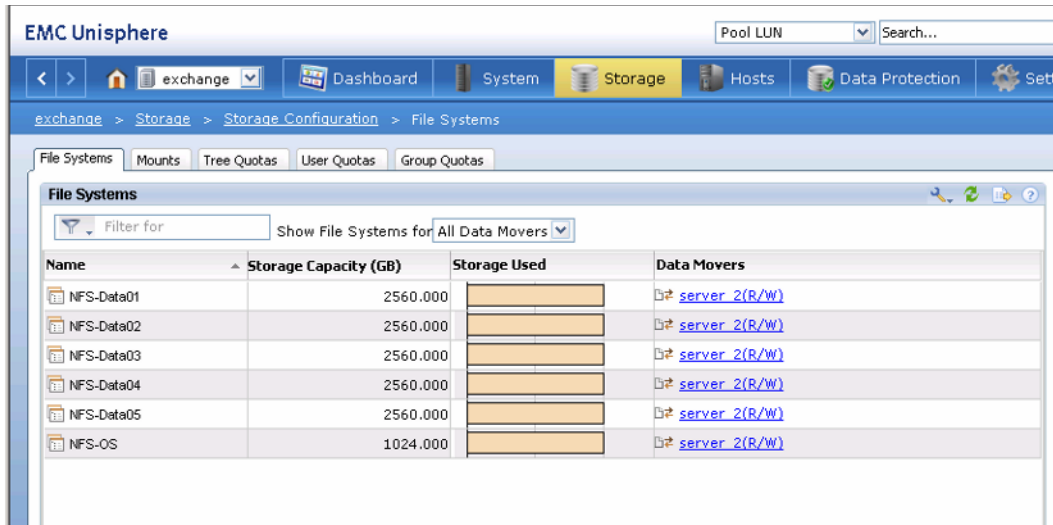
59. As per the above configuration NFS-OS share is created with 1024GB capacity.

**Figure 124** *Window Showing NFS-OS Storage Capacity*

Name	Storage Capacity (GB)	Storage Used	Data Movers
NFS-OS	1024.000		server_2(R/W)

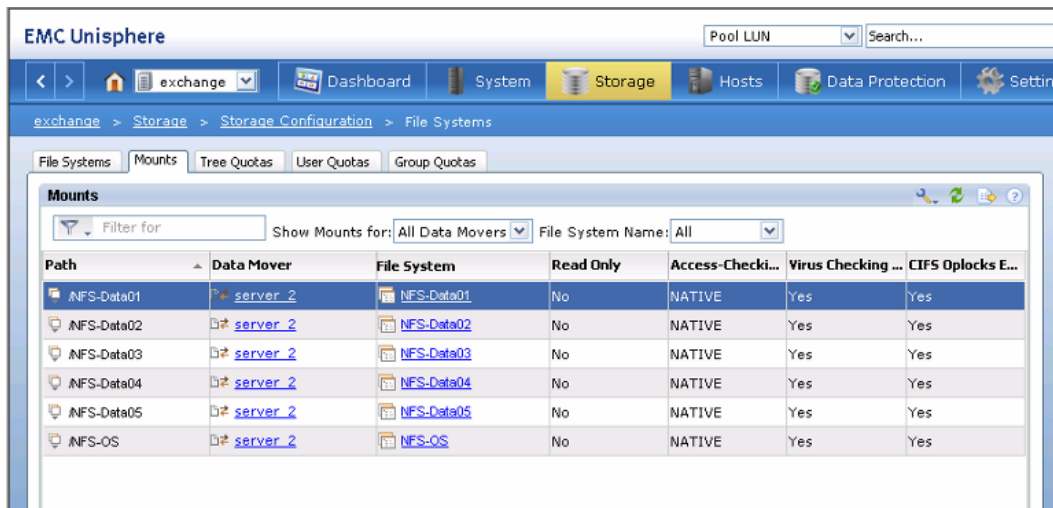
60. Follow the steps 57 to 59 to create as many File systems as needed. After creating, verify them under **File Systems** tab.

**Figure 125 Verifying Created File Systems**

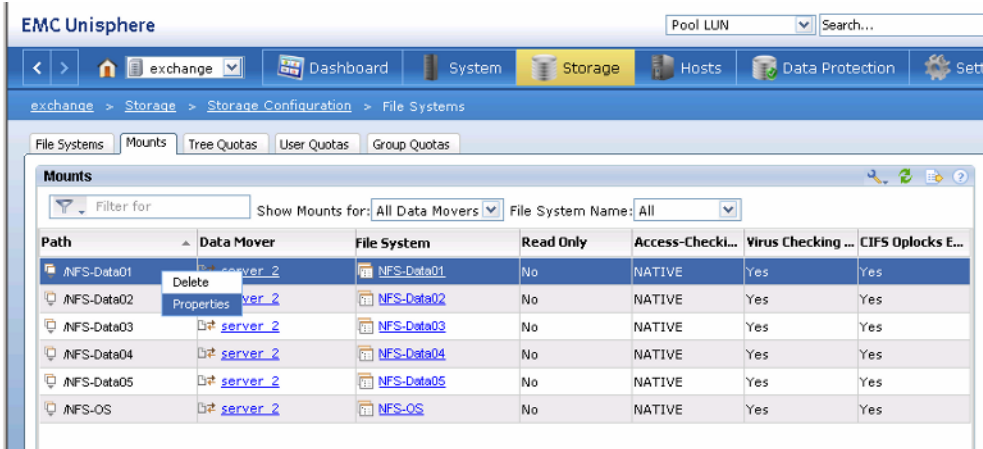


61. Choose **Mounts** tab and select the File Systems just created.

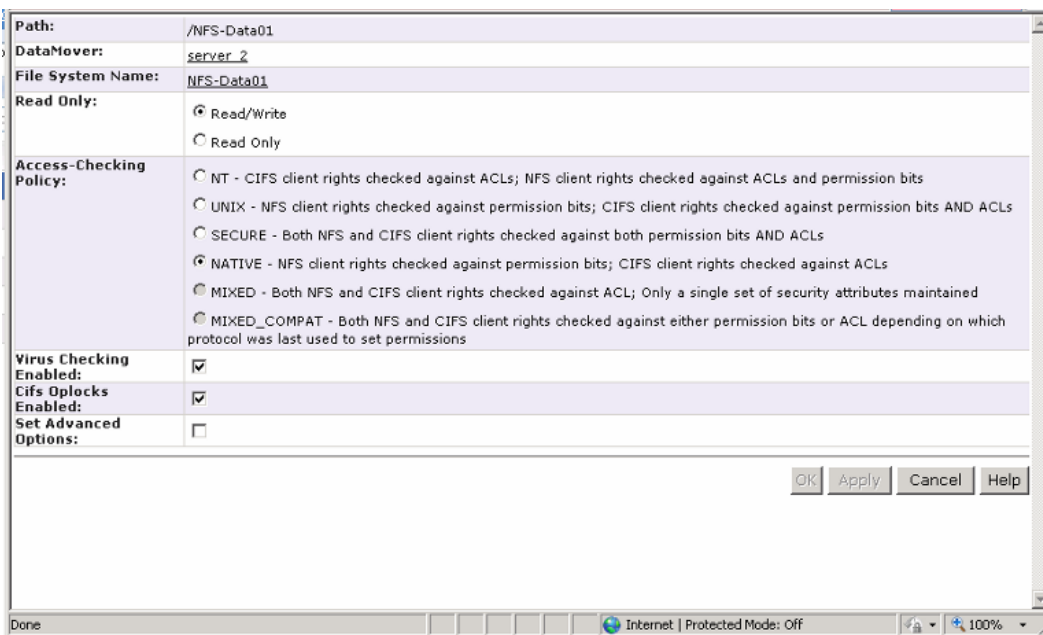
**Figure 126 Selecting the File System**



62. Right click on the selected file system and click **Properties**.

**Figure 127** *File System Properties*

63. In the Properties window, check the check box “Set Advance Options”.

**Figure 128** *Setting Advanced Options in File System*

64. Check the check box “Direct Writes Enabled” and click **OK**. Repeat steps 61 to 64 for all file systems that will be used as NFS datastores.



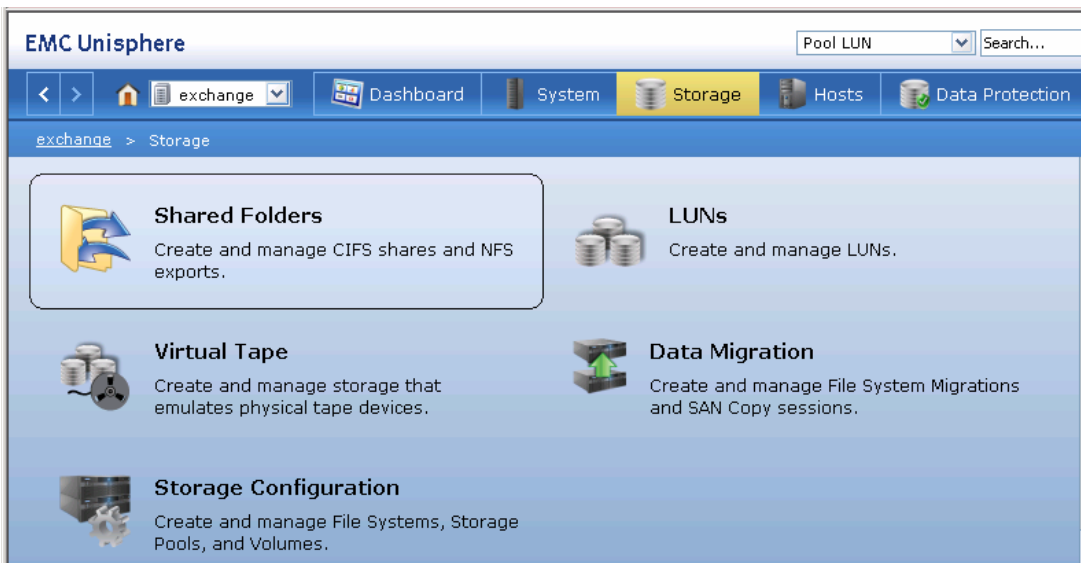
**Figure 129**      **Enabling Direct Writes**

The screenshot shows a configuration window for an NFS export. The settings are as follows:

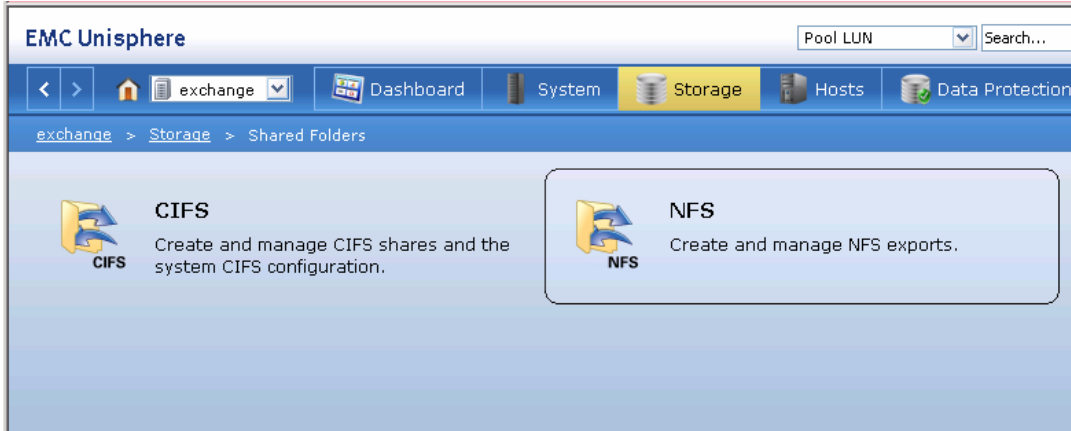
- Path:** /NFS-Data01
- DataMover:** server\_2
- File System Name:** NFS-Data01
- Read Only:**
  - ☒ Read/Write
  - ☐ Read Only
- Access-Checking Policy:**
  - ☐ NT - CIFS client rights checked against ACLs; NFS client rights checked against ACLs and permission bits
  - ☐ UNIX - NFS client rights checked against permission bits; CIFS client rights checked against permission bits AND ACLs
  - ☐ SECURE - Both NFS and CIFS client rights checked against both permission bits AND ACLs
  - ☐ NATIVE - NFS client rights checked against permission bits; CIFS client rights checked against ACLs
  - ☐ MIXED - Both NFS and CIFS client rights checked against ACL; Only a single set of security attributes maintained
  - ☐ MIXED\_COMPAT - Both NFS and CIFS client rights checked against either permission bits or ACL depending on which protocol was last used to set permissions
- Virus Checking Enabled:** ☒
- Cifs Oplocks Enabled:** ☒
- Set Advanced Options:** ☒
- Use NT Credential:** ☐
- Direct Writes Enabled:** ☐
- Prefetch Enabled:** ☒
- Multi-Protocol Locking Policy:**
  - ☒ nolock
  - ☐ writelock
  - ☐ rwlock
- CIFS Sync Writes Enabled:** ☐
- CIFS Notify Enabled:** ☒
  - CIFS Notify Trigger Level:**
  - CIFS Notify On Access Enabled:** ☐
  - CIFS Notify On Write Enabled:** ☐

Buttons at the bottom: OK, Apply, Cancel, Help.

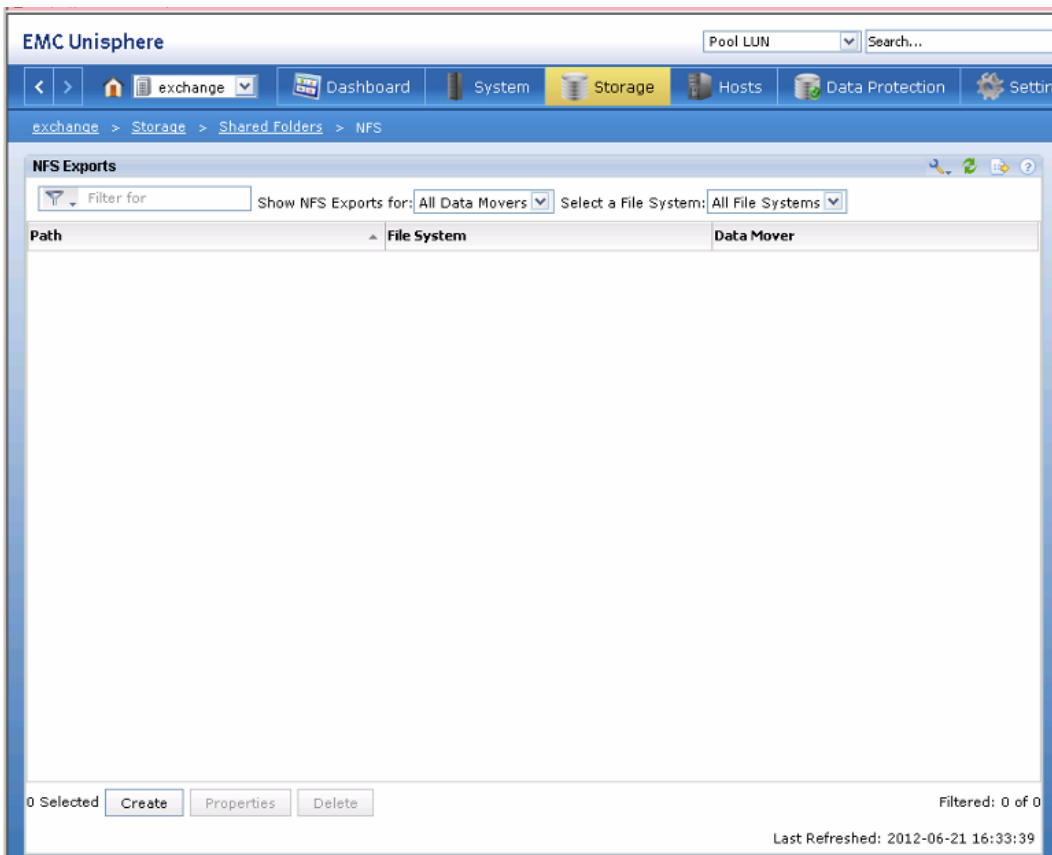
65. To map NFS export and assign to storage group, click **Storage** tab > **Shared Folders**.

**Figure 130**      **Window Showing Shared Folder in NFS Exports**

66. Choose NFS option.

**Figure 131** *Choosing NFS Option to Map NFS Export*

67. In the NFS Exports window, click **Create**.

**Figure 132** *NFS Exports Window*

68. From the drop-down list select the File System created above; verify the path. Add Root Hosts and Access Hosts (These are the IP addresses of the ESX hosts' vmkernel Storage NIC). Click **OK**.



**Note** Multiple hosts are separate by “:”

**Figure 133**      **Create NFS Export**

Choose Data Mover: server\_2

File System: NFS-Data01 (/NFS-Data01)

Path: /NFS-Data01

Host Access

Read-only Export: ☐

Read-only Hosts:

Read/Write Hosts:

Root Hosts:

Access Hosts:

10.10.46.101:10.10.  
46.102:10.10.46.103  
:10.10.46.104:10.10.

OK Apply Cancel Help

Done Internet | Protected Mode: Off 100%

69. NFS export is now available to add to VMware ESXi hosts.

**Figure 134**      **Window Showing Availability of File System for Data Mover**

EMC Unisphere

Pool LUN Search...

exchange > 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-Data01	NFS-Data01	server_2

70. Repeat the above steps to create NFS export for each file system.

**Figure 135** *Creating NFS Export for Each File System*

Choose Data Mover: server\_2

File System: NFS-Data02 (/NFS-Data02)

Path: /NFS-Data02

Host Access

Read-only Export: ☐

Read-only Hosts:

Read/Write Hosts:

Root Hosts: 10.10.40.101:10.10.40.102:10.10.40.103:10.10.40.104:10.10.40.105

Access Hosts: 40.102:10.10.40.103:10.10.40.104:10.10.40.105

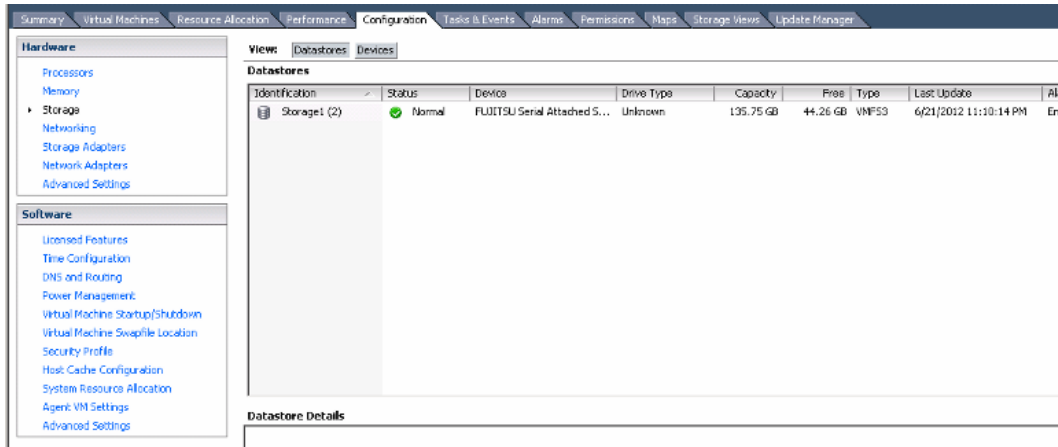
OK Apply Cancel Help

71. All the File Systems are now available to add to the VMware ESXi hosts.

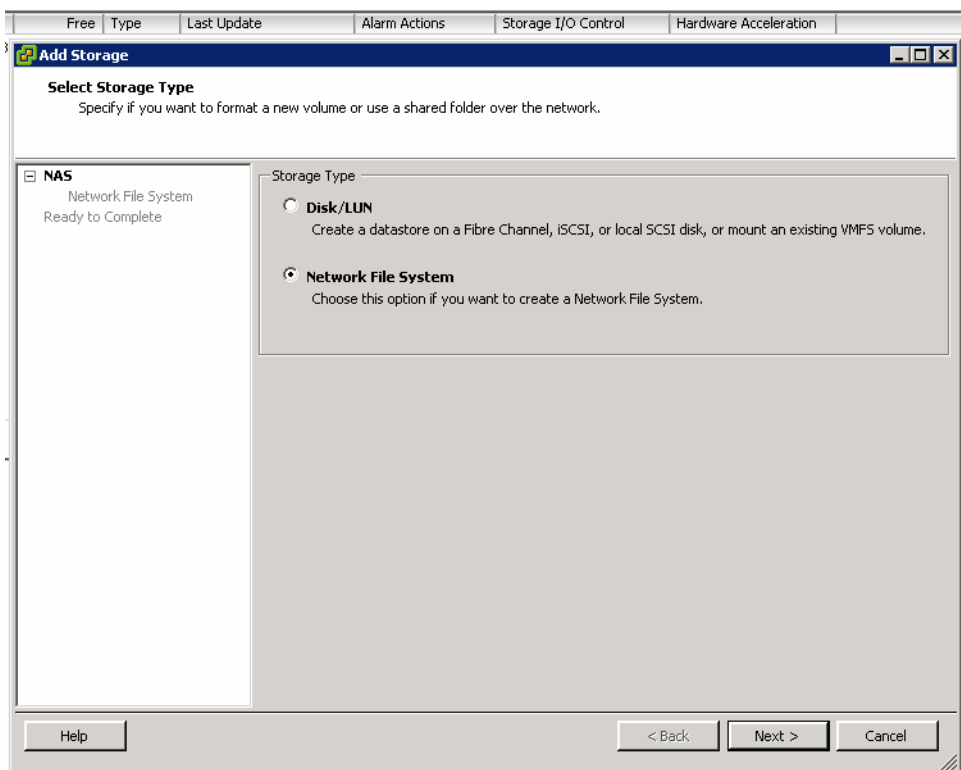
**Figure 136** *Window Showing Availability of File System for Data Mover*

Path	File System	Data Mover
/NFS-Data01	NFS-Data01	server_2
/NFS-Data02	NFS-Data02	server_2
/NFS-Data03	NFS-Data03	server_2
/NFS-Data04	NFS-Data04	server_2
/NFS-Data05	NFS-Data05	server_2
/NFS-OS	NFS-OS	server_2

72. To add NFS export created traditionally in the VMware vCenter, choose **ESXi Hosts > Configuration > Storage**. Click **Add Storage**.

**Figure 137 Adding Storage In VMware ESXi Host**

73. In the Add Storage window, click “Network File System” radio button in the “Storage Type” pane. Click **Next**.

**Figure 138 Selecting Storage Type**

74. In the Network File System window, in the Properties pane enter the IP for NFS server in the Server field, export path for NFS in the Folder field, and in the Datastore pane enter the datastore name. Click **Next**.

**Figure 139**      **Locating Network File System**

The screenshot shows the 'Add Storage' wizard window. The title bar says 'Add Storage'. The main heading is 'Locate Network File System' with the question 'Which shared folder will be used as a vSphere datastore?'. On the left, there is a tree view with 'NAS' expanded, showing 'Network File System' with the status 'Ready to Complete'. The main area is divided into 'Properties' and 'Datastore Name' sections. The 'Properties' section contains: 'Server:' with the value '10.10.40.111' and examples 'nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2'; 'Folder:' with the value '/NFS-Data01' and example '/vols/vol0/datastore-001'; and a checkbox 'Mount NFS read only' which is unchecked. Below the checkbox is a warning icon and text: 'If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.' The 'Datastore Name' section has a text box with the value 'NFS-share01'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', and 'Cancel'.

**Add Storage**

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

**NAS**  
**Network File System**  
Ready to Complete

**Properties**

Server: 10.10.40.111  
Examples: nas, nas.it.com, 192.168.0.1 or FE80:0:0:0:2AA:FF:FE9A:4CA2

Folder: /NFS-Data01  
Example: /vols/vol0/datastore-001

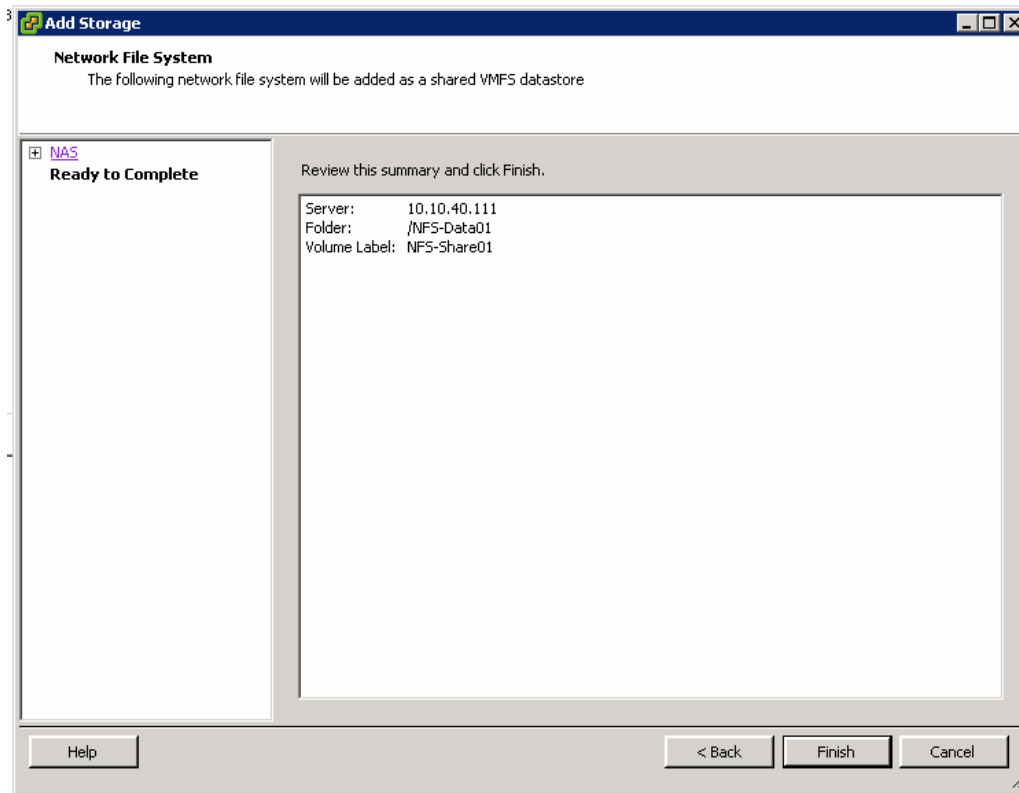
☐ Mount NFS read only

If a datastore already exists in the datacenter for this NFS share and you intend to configure the same datastore on new hosts, make sure that you enter the same input data (Server and Folder) that you used for the original datastore. Different input data would mean different datastores even if the underlying NFS storage is the same.

**Datastore Name**  
NFS-share01

Help < Back Next > Cancel

75. Review the settings and click **Finish**.

**Figure 140 NFS Summary**

76. Repeat the steps 73 to 75, to add all the NFS exports to the VMware ESXi hosts.

**Figure 141 Window Showing all NFS Exports and VMware ESXi Hosts**

Configuration   Tasks & Events   Alarms   Permissions   Maps   Storage Views   Update Manager										
View: Datasources   Devices										
Datasources										
Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update	Alarm Actions	Storage I/O Control	
datastore1 (2)	Normal	Local LSI Disk (naa.60030...	Non-SSD	70.00 GB	69.05 GB	VMFS5	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-OS	Normal	10.10.40.111:/NFS-OS	Unknown	1,008.38 G	1,008.37 G	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-Share01	Normal	10.10.40.111:/NFS-Data01	Unknown	2.46 TB	2.46 TB	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-Share02	Normal	10.10.40.111:/NFS-Data02	Unknown	2.46 TB	2.46 TB	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-Share03	Normal	10.10.40.111:/NFS-Data03	Unknown	2.46 TB	2.46 TB	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-Share04	Normal	10.10.40.111:/NFS-Data04	Unknown	2.46 TB	2.46 TB	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	
NFS-Share05	Normal	10.10.40.111:/NFS-Data05	Unknown	2.46 TB	2.46 TB	NFS	6/21/2012 11:14:40 PM	Enabled	Disabled	

## Installing VMware ESXi Servers and vCenter Infrastructure

To install the VMware ESXi servers, follow these steps:

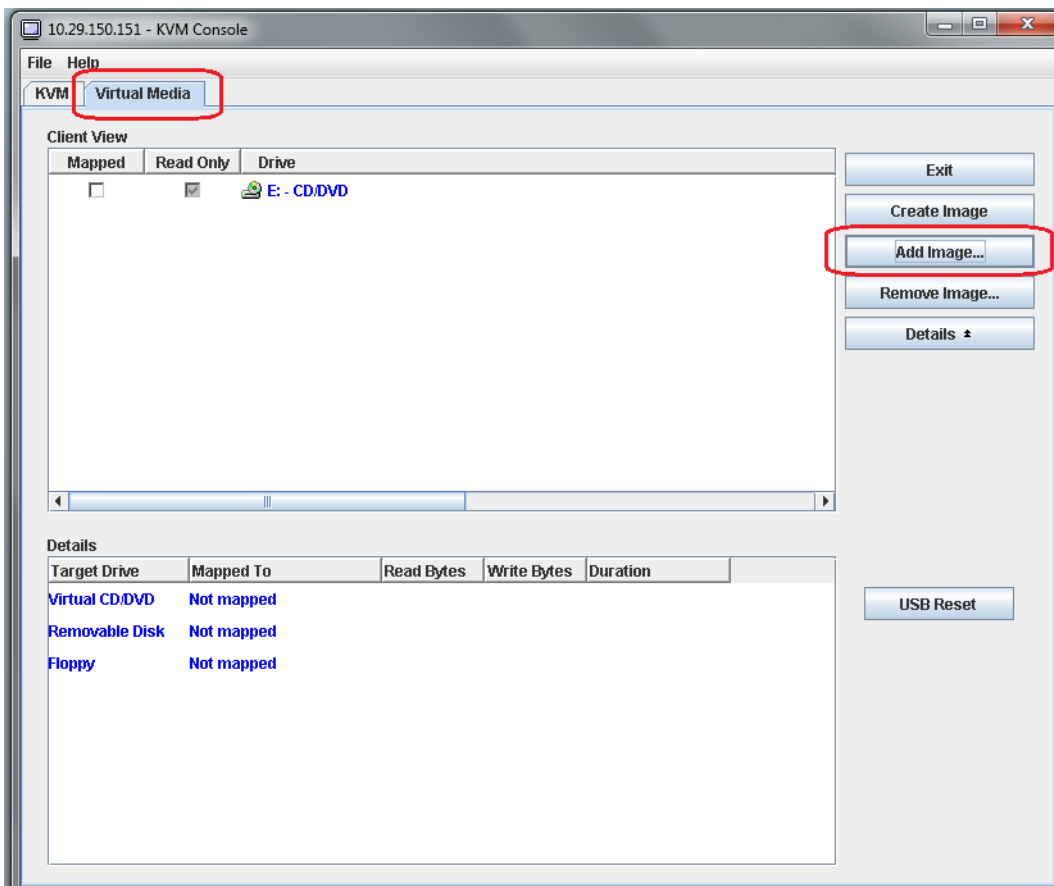
1. Access the CIMC of Cisco C220 M3 servers using the management IP address and launch the KVM for the server as shown in [Figure 142](#).

Figure 142 Server Summary Window in CIMC



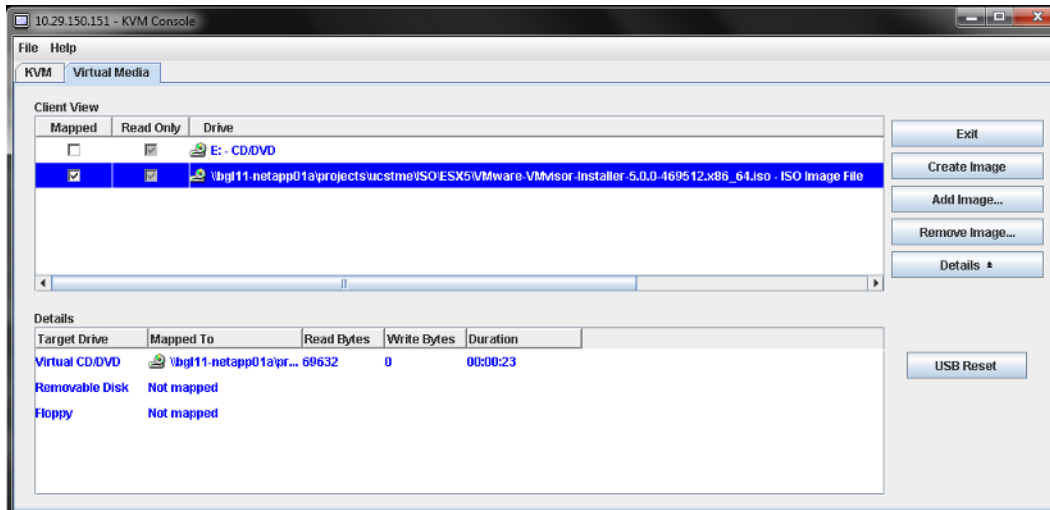
- When the Java applet of the KVM is launched, click **Virtual Media** tab > **Add Image** tab as shown in Figure 143. A new window is displayed to select an ISO image. Navigate in the local directory structure and select the ISO image of the VMware ESXi 5.0 hypervisor installer media.

Figure 143 Adding Image in Virtual Media



- When the ISO image shows up in the list, check the “Mapped” check box and reset the server.



**Figure 144**      **Selecting the ISO Image**

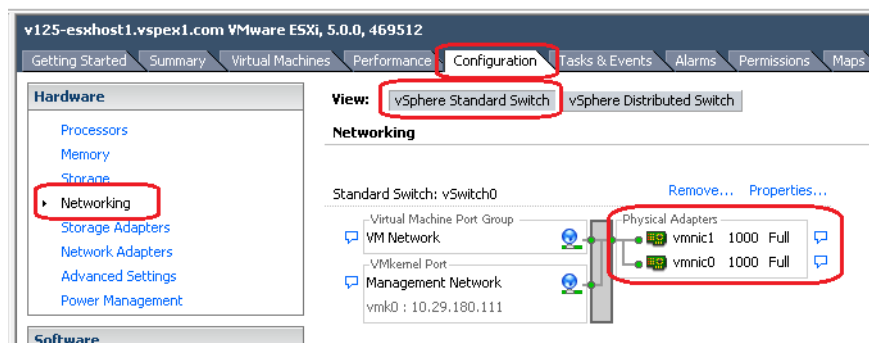
4. On restarting the server, VMware ESXi 5.0 install media will boot. Follow the above mentioned steps to install the hypervisor on each of the servers. ESXi hostnames, IP addresses, and a root password are required for the installation.

The Appendix A Customer Configuration Data provides appropriate values.

5. The VMware ESXi OS should be installed on the local disk of the C220 M3 servers. When the ESXi is installed, verify the network connectivity and accessibility of each server from each other.

## Configure ESXi Networking

During the installation of VMware ESXi, a standard virtual switch (vSwitch) will be created. By default, ESXi chooses only one physical NIC as a virtual switch uplink. This is the 1 GigE mLOM port on the C220 M3 server. To maintain redundancy and bandwidth requirements, the second 1 GigE mLOM port (vmkn1c1) must be added either by using the ESXi console or by connecting to the ESXi host from the VMware vSphere Client. When the two 1 GigE mLOM NICs are added to the native vSwitch, the UI looks like in [Figure 145](#).

**Figure 145**      **ESXi Console Showing Added Physical Adapters**

For 100 and 125 virtual machines architecture, each VMware ESXi server has two 10 GigE interfaces connected to each of the Cisco Nexus 5548UP switch to ensure redundancy which is used for network load balancing, link aggregation, and network adapter failover. Similarly, for 50 virtual machines

architecture, two additional 1 GigE interfaces are connected through the Broadcom adapter to each of the Cisco Nexus 3048 switches. These ports are used for storage access, vMotion and VM data traffic through Cisco Nexus 1000v virtual Distributed Switch.

## Installing and Configuring Microsoft SQL Server Database

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

1. Create a VM for Microsoft® SQL server



### Note

The customer environment may already contain an SQL Server that is designated for this role. In that case, refer to Configure database for VMware vCenter.

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

2. Install Microsoft® Windows on the VM

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

3. Install SQL server

Install SQL Server on the virtual machine from the SQL Server installation media. The Microsoft TechNet website provides information on how to install SQL Server.

4. Configure database for VMware vCenter

To use VMware vCenter in this solution, you will need to create a database for the service to use. The requirements and steps to configure the vCenter Server database correctly are covered in Preparing vCenter Server Databases.



### Note

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

It is a best practice to create individual login accounts for each service accessing a database on SQL Server.

5. Configure database for VMware Update Manager

To use VMware Update Manager in this solution you will need to create a database for the service to use. The requirements and steps to configure the Update Manager database correctly are covered in Preparing the Update Manager Database. It is a best practice to create individual login accounts for each service accessing a database on SQL Server. Consult your database administrator for your organization's policy.

6. Deploy the VNX VAAI for NFS plug-in

The VAAI for NFS plug-in enables support for the VMware vSphere 5 NFS primitives. These primitives reduce the load on the hypervisor from specific storage-related tasks to free resources for other operations. Additional information about the VAAI for NFS plug-in is available in the plug-in download VMware vSphere Storage APIs for Array Integration (VAAI) Plug-in.

**Note**

The same version of the plug-in supports both the VNX and VNXe platforms.

The VAAI for NFS plug-in is installed using VMware vSphere Update Manager. Refer process for distributing the plug demonstrated in the EMC VNX VAAI NFS plug-in – installation HOWTO video available on the [www.youtube.com](http://www.youtube.com) web site. To enable the plug-in after installation, you must reboot the ESXi server.

## Deploying VMware vCenter Server

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

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

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

[http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc\\_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html](http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html).

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

[http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc\\_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html](http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html).

For detailed information on creating a Virtual Machine in the VMware vSphere 5 client, see the link:

[http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.vm\\_admin.doc\\_50/GUID-0433C0DC-63F7-4966-9B53-0BECDDDEB6420.html](http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.vm_admin.doc_50/GUID-0433C0DC-63F7-4966-9B53-0BECDDDEB6420.html).

To configure vCenter server, follow these steps:

**Note**

These steps provide high level configuration procedure to configure vCenter server.

1. Create the vCenter host VM

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

2. Install vCenter guest OS

Install the guest OS on the vCenter host virtual machine. VMware recommends using Windows Server 2008 R2 SP1. To ensure that adequate space is available on the vCenter and vSphere Update Manager installation drive, see VMware vSphere Installation and Setup Guide.

3. Create vCenter ODBC connection

Before installing vCenter Server and vCenter Update Manager, you must create the ODBC connections required for database communication. These ODBC connections will use SQL Server authentication for database authentication. Appendix A Customer Configuration Data provides SQL login information.

For instructions on how to create the necessary ODBC connections see, VMware vSphere Installation and Setup and Installing and Administering VMware vSphere Update Manager.

**4. Install vCenter server**

Install vCenter by using the VMware VIMSetup installation media. Use the customer-provided username, organization, and vCenter license key when installing vCenter.

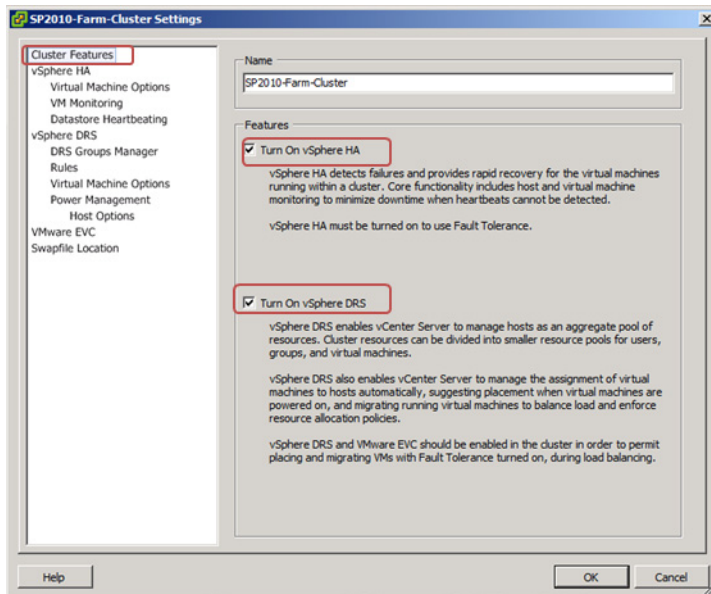
**5. Apply VMware vSphere license keys**

To perform license maintenance, log into the vCenter Server and select the Administration - Licensing menu from the VMware 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 VMware vCenter


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

1. Log into VMware ESXi Host using VMware vSphere Client.
2. Create a vCenter Data Center.
3. Create a new management cluster with DRS and HA enabled.
  1. Right-click on the cluster and in the corresponding Context menu, click **Edit Settings**.
  2. Check the check boxes “Turn On vSphere HA”, and “Turn On vSphere DRS”, as shown in [Figure 146](#).
  3. Click **Ok**, to save the changes.
4. Add all ESXi hosts to the cluster by providing servers’ management IP addresses and login credentials one by one.

**Figure 146** *Farm Cluster Settings*

## Template-Based Deployments for Rapid Provisioning

This section provides information on how to deploy virtual machines using vCenter GUI.

**Figure 147** *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.

## Installing and Configuring Cisco Nexus 1000v

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

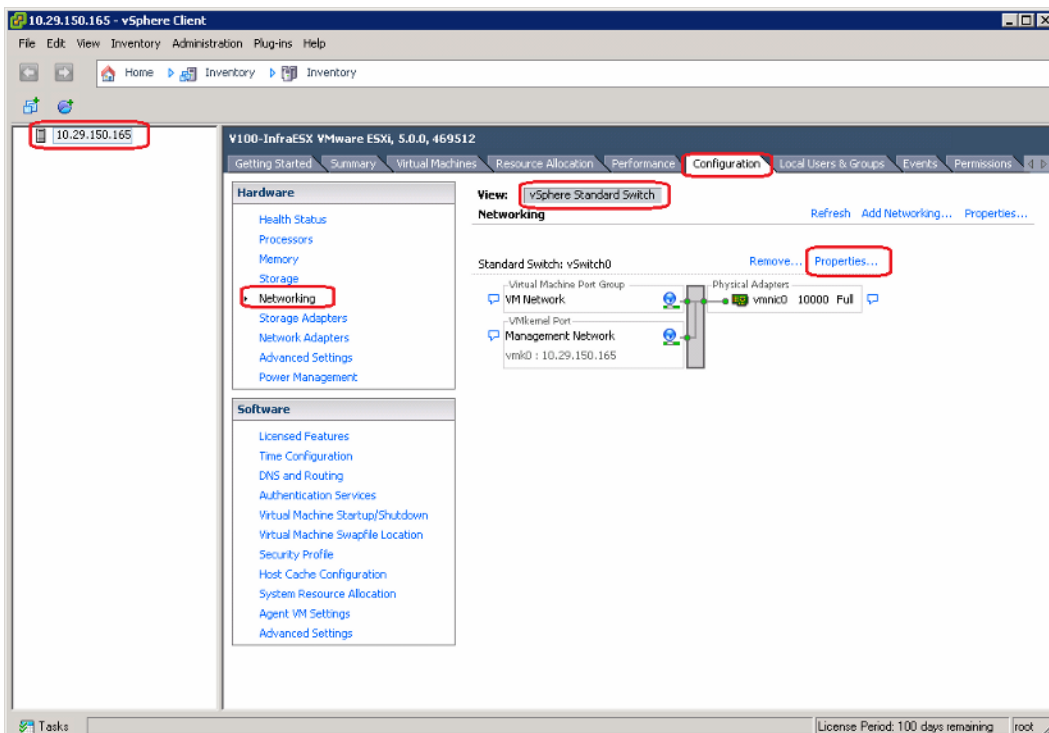
1. [Installing Cisco Nexus 1000v VSM VM, page 122](#)
2. [Connecting Cisco Nexus 1000v VSM to VMware vCenter, page 136](#)
3. [Configuring Port-Profile in VSM and Migrate vCenter Networking to vDS, page 143](#)

### Installing Cisco Nexus 1000v VSM VM

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

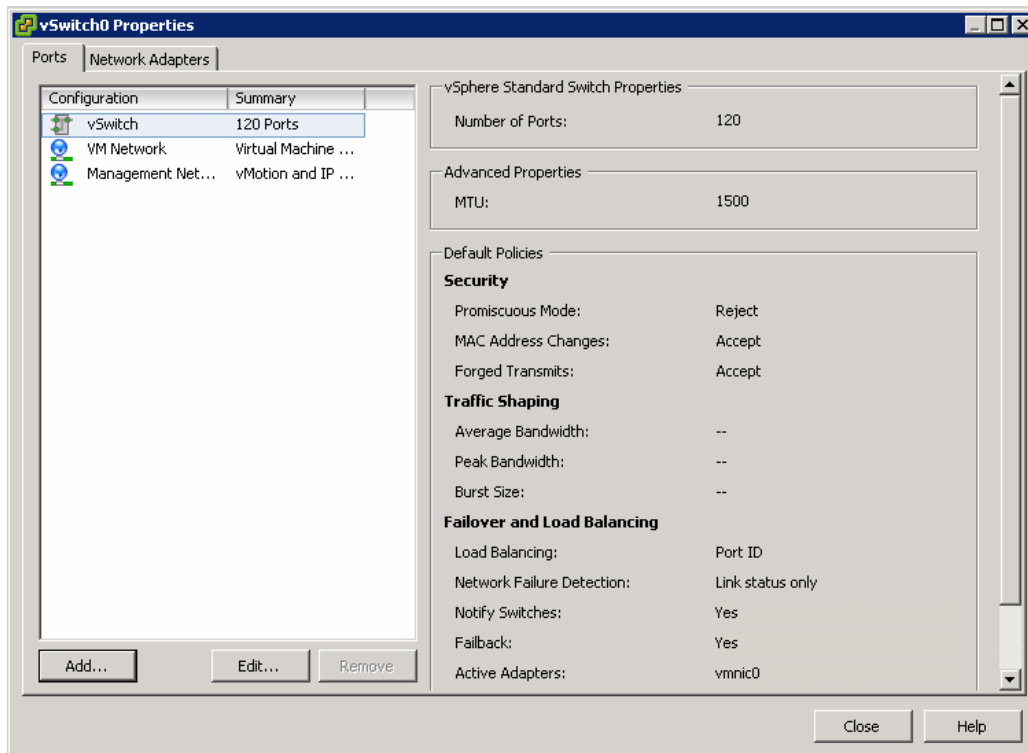
1. Click the infrastructure infraESX VMware ESXi, 5.0 on which the VSM VM is to be deployed. Choose **Configuration > Networking > vSphere Standard Switch**, and click **Properties** in the Networking pane as shown in [Figure 148](#).

**Figure 148**      *Selecting VMware vSphere Standard Switch View*



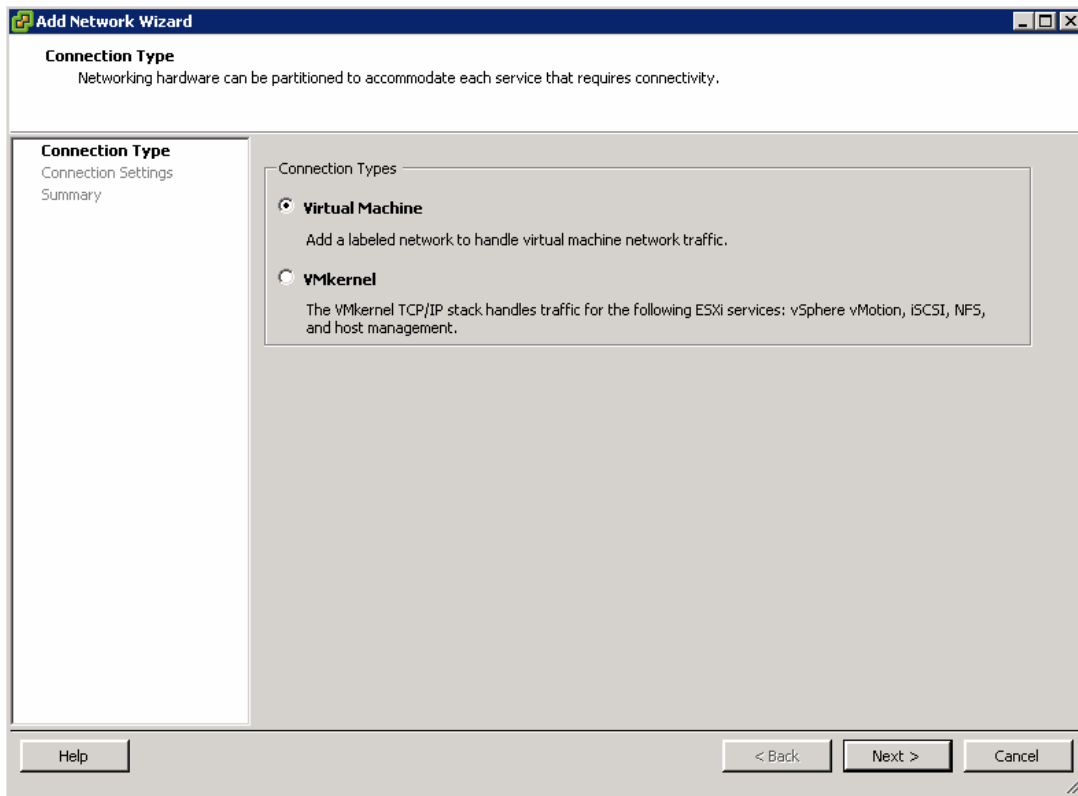
2. Click **Add...** on the “vSwitch0 Properties” window.

**Figure 149** Window Showing vSwitch0 Properties



3. Click the “Virtual Machine” radio button to add new VLANs (networks) and click **Next**.

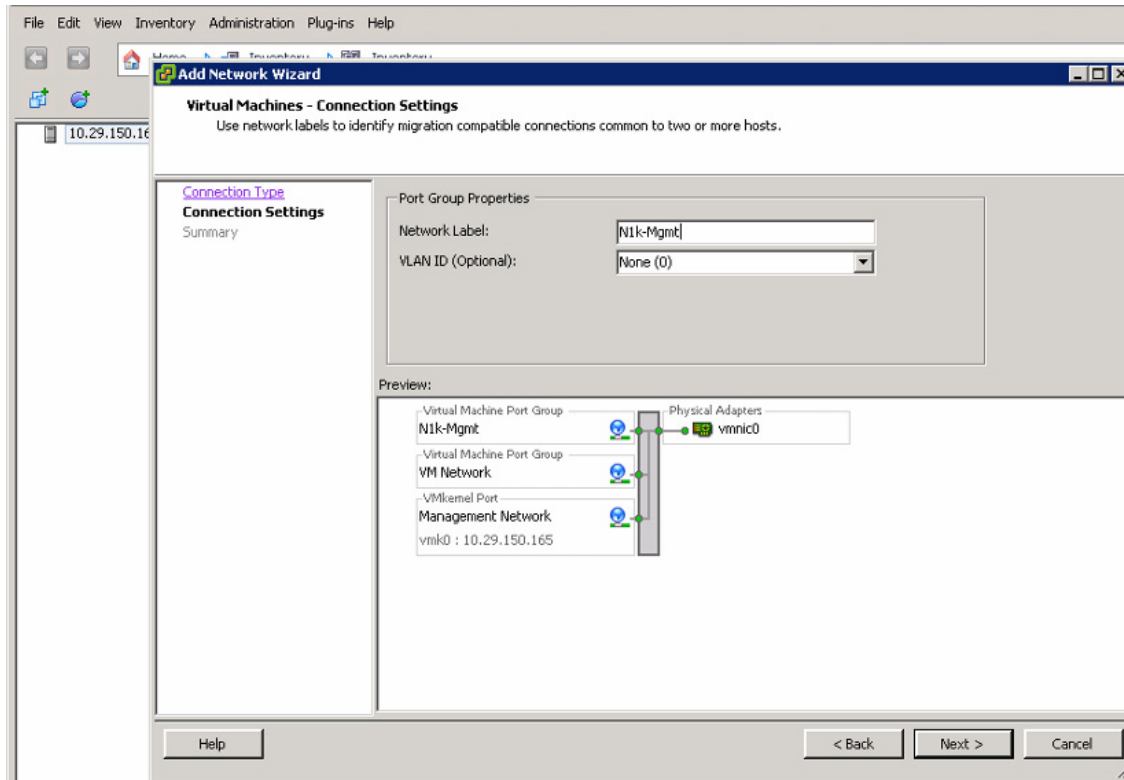
**Figure 150**      **Connection Type Window in Add Network Wizard**



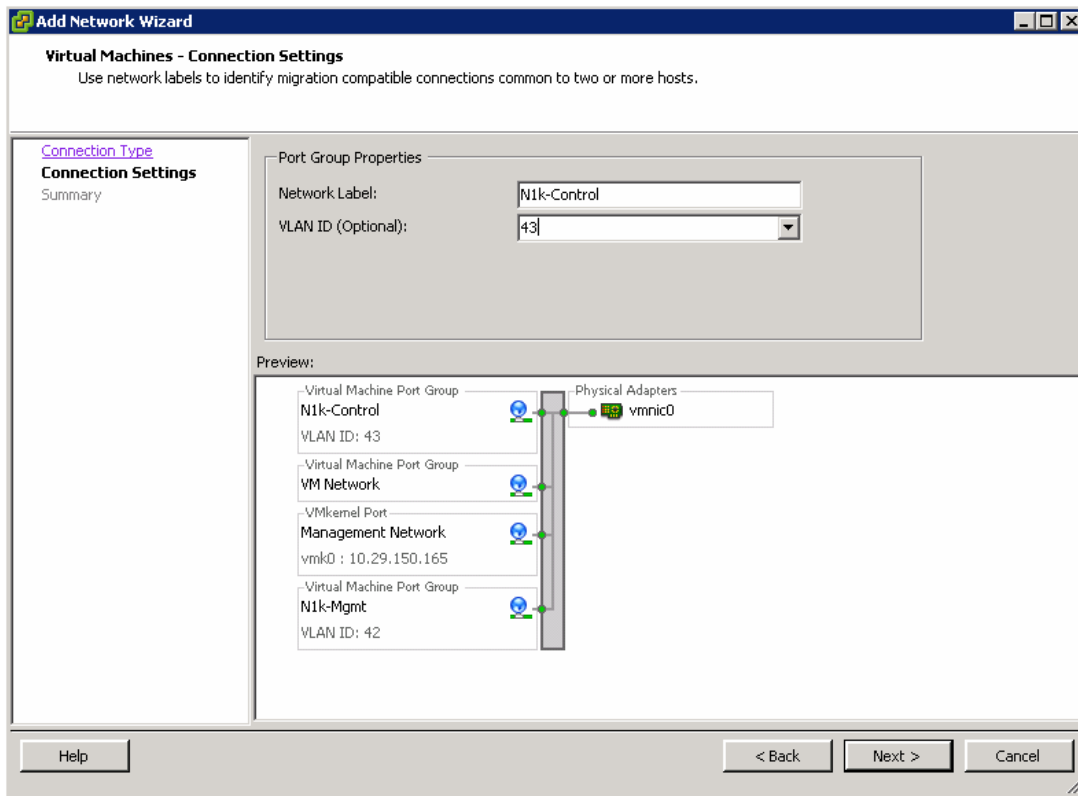
4. In the Add Network wizard, specify the Network Label as “N1k-Mgmt” and provide the VLAN ID as “None (0)” to use the default (native) VLAN. Click **Next** and in the next window click **Finish**.



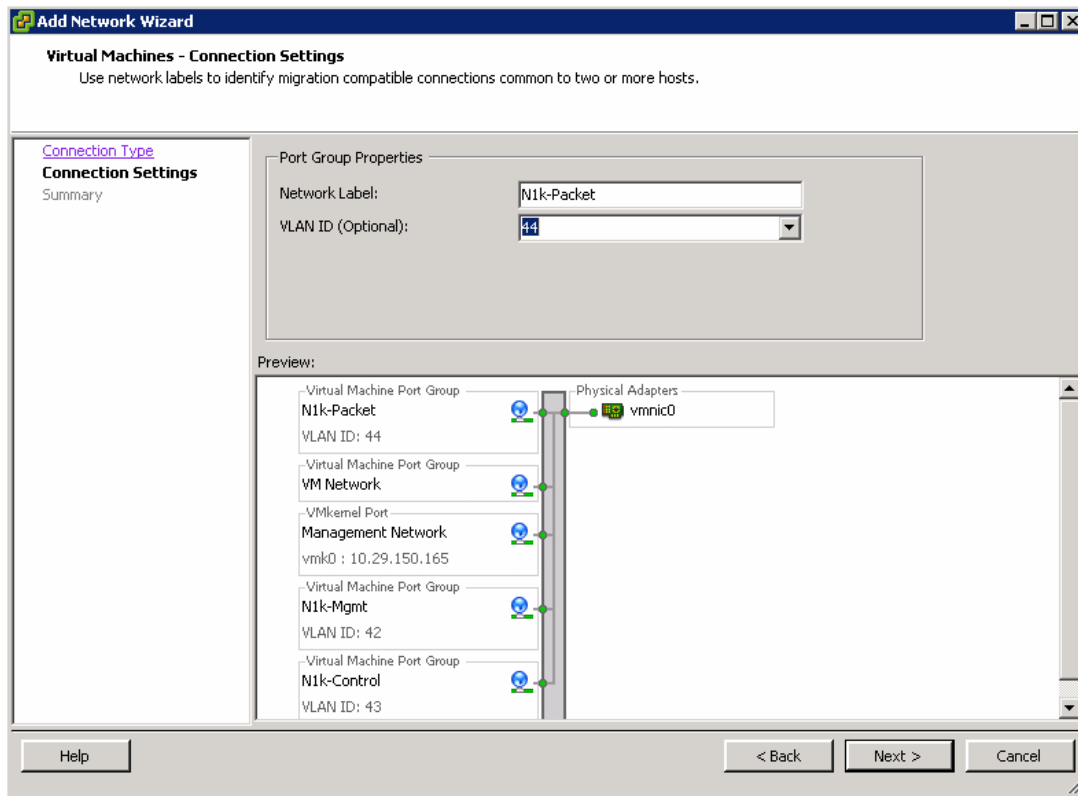
**Figure 151**      **Connection Settings Window in Add Network Wizard**



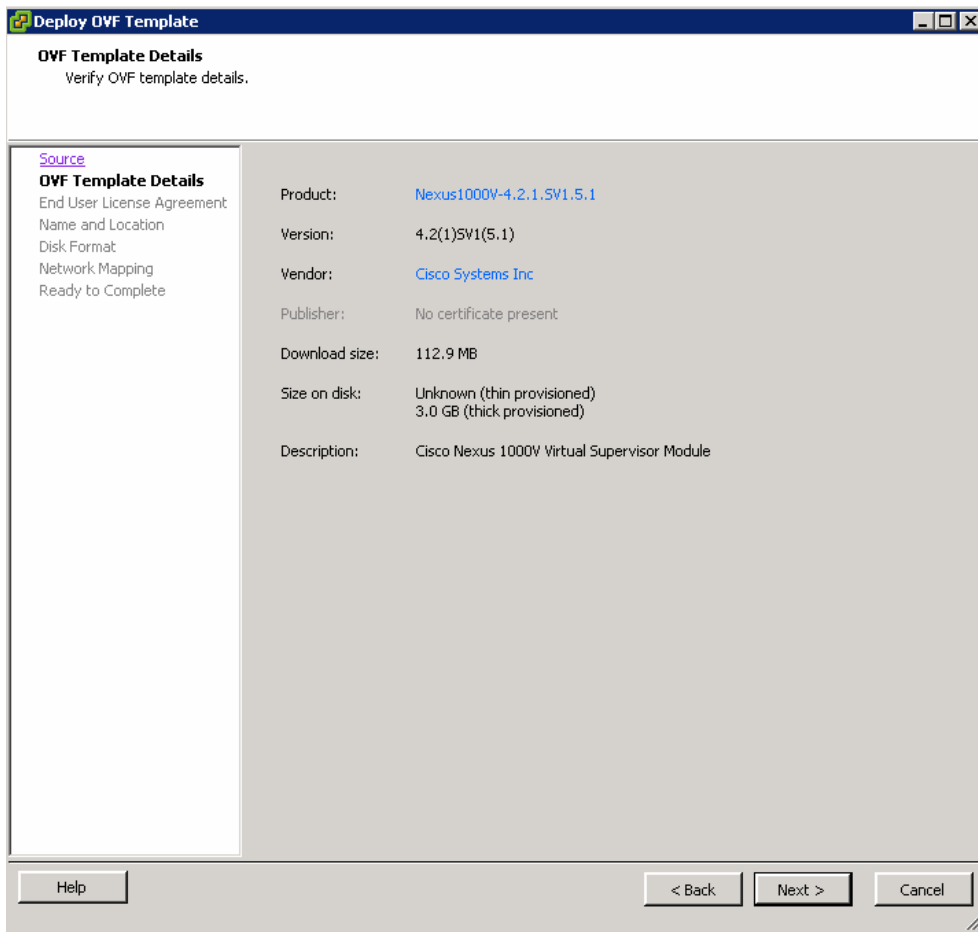
5. Similarly, add two more VLANs “N1k-Control” and “N1k-Packet” with appropriate VLAN IDs as shown in [Figure 152](#), [Figure 153](#).

**Figure 152**      **Connection Settings Window in Add Network Wizard**

**Figure 153** Connection Settings Window in Add Network Wizard



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

**Figure 154**      **Verifying OVF Template Details**

7. Specify VSM virtual machine name in the next window of the Deploy OVF Template wizard.

**Figure 155**      **Specifying Virtual Machine Name**

The screenshot shows a window titled "Deploy OVF Template" with a standard Windows-style title bar. The main area is divided into two panes. The left pane contains a list of steps: "Source", "OVF Template Details", "End User License Agreement", "Name and Location" (which is highlighted), "Disk Format", "Network Mapping", and "Ready to Complete". The right pane is for the "Name and Location" step. It has a sub-header "Name and Location" and a sub-instruction "Specify a name and location for the deployed template". Below this is a text input field labeled "Name:" containing the text "V100-N1k-VSM". Under the input field, a note states: "The name can contain up to 80 characters and it must be unique within the inventory folder." At the bottom of the window, there are three buttons: "Help", "< Back", and "Next >", and a "Cancel" button on the far right.

8. Specify the correct DataStore and click “Flat Disk” radio button. Click **Next**.

**Figure 156**      **Specifying Disk Format**

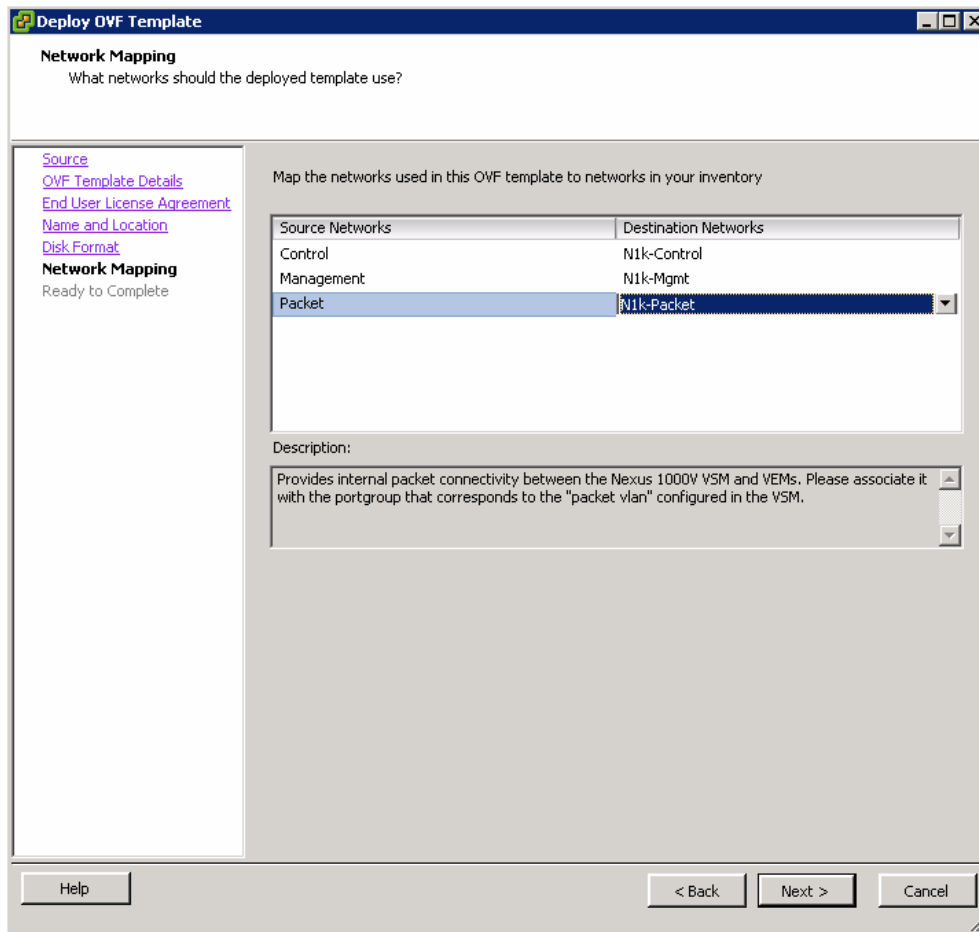
The screenshot shows the 'Deploy OVF Template' wizard window. The title bar reads 'Deploy OVF Template'. The main heading is 'Disk Format' with the subtext 'In which format do you want to store the virtual disks?'. On the left, a navigation pane lists: 'Source', 'OVF Template Details', 'End User License Agreement', 'Name and Location', 'Disk Format' (highlighted), 'Network Mapping', and 'Ready to Complete'. The main area contains the following fields and options:

- Datastore:** A text box containing 'datastore1'.
- Available space (GB):** A text box containing '130.1'.
- Format Options:** Three radio buttons are listed:
  - ☒ Flat Disk
  - ☐ Thick Provision
  - ☐ Thin Provision

At the bottom, there are three buttons: 'Help', '< Back', and 'Next >', followed by a 'Cancel' button.

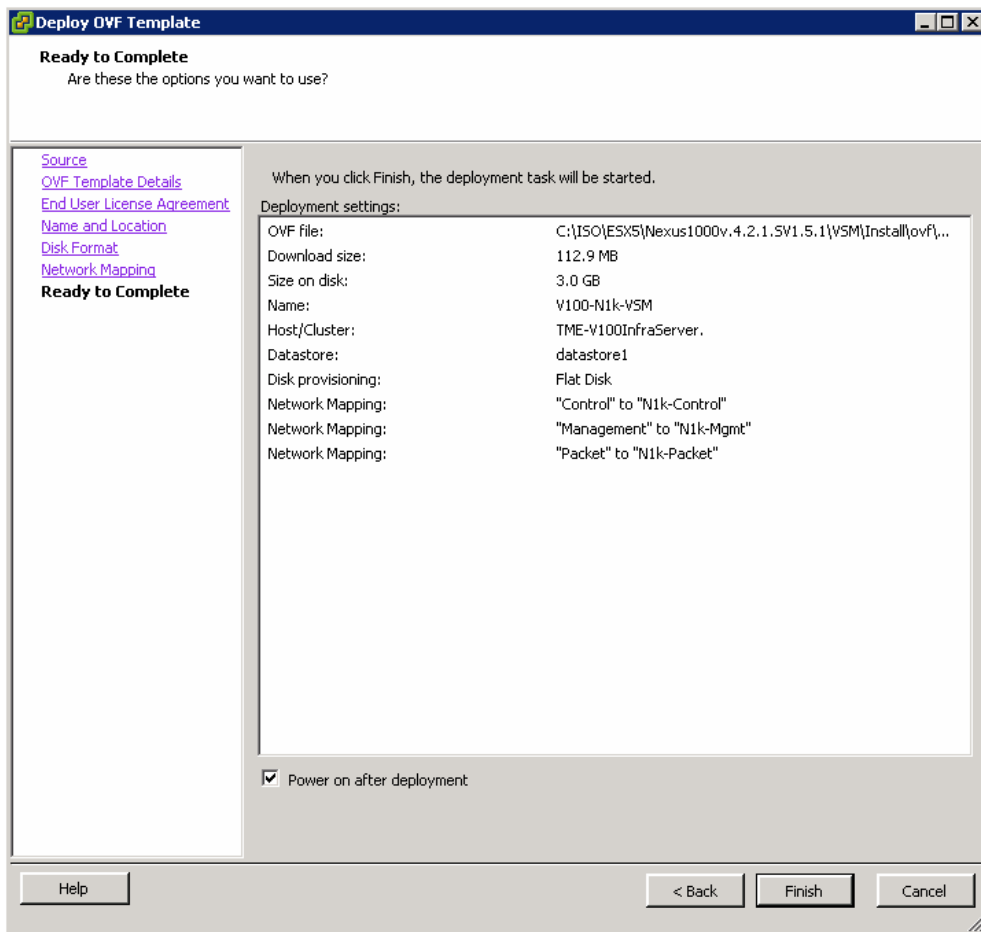
9. Choose the Network Mapping for mapping the networks to the deployed OVF template.

**Figure 157**      **Selecting Destination Packet to Map Networks to the Inventory**



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

**Figure 158**      **Verifying Deployment Settings**



### Initial configuration of VSM VM

When VSM VM is powering up for the first time, click **Console** of the virtual machine in the vCenter and follow these steps to perform initial configuration of the VSM VM.

1. Type “Yes” for the question “Would you like to enter the basic configuration dialog?”



**Figure 159 Basic System Configuration Dialog**

```

----- System Admin Account Setup -----

Enter the password for "admin":
Confirm the password for "admin":
Enter HA role[standalone/primary/secondary]: standalone

Enter the domain id<1-4095>: 10
[#####] 100%

----- Basic System Configuration Dialog -----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): _

```

2. Provide switch name, management IP address, subnet mask and default gateway as shown in [Figure 160](#).

**Figure 160 Entering Configuration Details**

```

Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : U100-USM
Continue with Out-of-band (Mgmt0) management configuration? (yes/no) [y]:
  Mgmt0 IPv4 address : 10.29.150.167
  Mgmt0 IPv4 netmask : 255.255.255.0
Configure the default gateway? (yes/no) [y]:
  IPv4 address of the default gateway : 10.29.150.1
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]: _

```

License Period: 100 days remaining | To release cursor, press CTRL+ALT | root

3. Configure SVS domain parameters, and appropriate VLANs for control and packet VLANs.

**Figure 161**      **Configuring SVS Domain and VLANs**

```

V100-N1k-VSM
Getting Started Summary Resource Allocation Performance Events Console Permissions

Configure the default gateway? (yes/no) [y]:
  IPv4 address of the default gateway [10.29.150.1]:
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Disable the ssh service? (yes/no) [n]:
Enable the http-server? (yes/no) [y]:
Configure the ntp server? (yes/no) [n]:
Vem feature level will be set to 4.2(1)SU1(5.1), Do you want to reconfigure? (
yes/no) [n]:
Configure svcs domain parameters? (yes/no) [y]:
  Enter SVS Control mode (L2 / L3) [L2]: L2
  Enter control vlan <1-3967, 4048-4093> [43]: 43
  Enter packet vlan <1-3967, 4048-4093> [44]: 44_

License Period: 100 days remaining  To release cursor, press CTRL+ALT  root

```

4. Verify the configuration and save the configuration.

**Figure 162**      **Verifying the Configuration Details**

The screenshot shows a terminal window titled "V100-N1k-VSM" with tabs for "Getting Started", "Summary", "Resource Allocation", "Performance", "Events", "Console", and "Permissions". The "Console" tab is active, displaying the following text:

```

Enter control vlan <1-3967, 4048-4093> [43]: 43
Enter packet vlan <1-3967, 4048-4093> [44]: 44

The following configuration will be applied:
switchname V100-USM
interface mgmt0
ip address 10.29.150.167 255.255.255.0
no shutdown
vrf context management
ip route 0.0.0.0/0 10.29.150.1
no telnet server enable
ssh key rsa 1024 force
ssh server enable
feature http-server
svs-domain
  svs mode L2
  control vlan 43
  packet vlan 44
  domain id 10
vlan 43
vlan 44

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

```

At the bottom of the window, there is a status bar with the text "License Period: 100 days remaining", "To release cursor, press CTRL+ALT", and "root".

5. When the initial configuration is complete, you will see the exec mode CLI as shown in [Figure 163](#).

**Figure 163**      **Confirming the Configuration Details**

The screenshot shows a console window for V100-N1k-VSM. The window has tabs for Getting Started, Summary, Resource Allocation, Performance, Events, Console (selected), and Permissions. The console output shows the following text:

```

vlan 44

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

Use this configuration and save it? (yes/no) [y]:

[#####] 100%

Nexus 1000v Switch
V100-USM login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
V100-USM# _

```

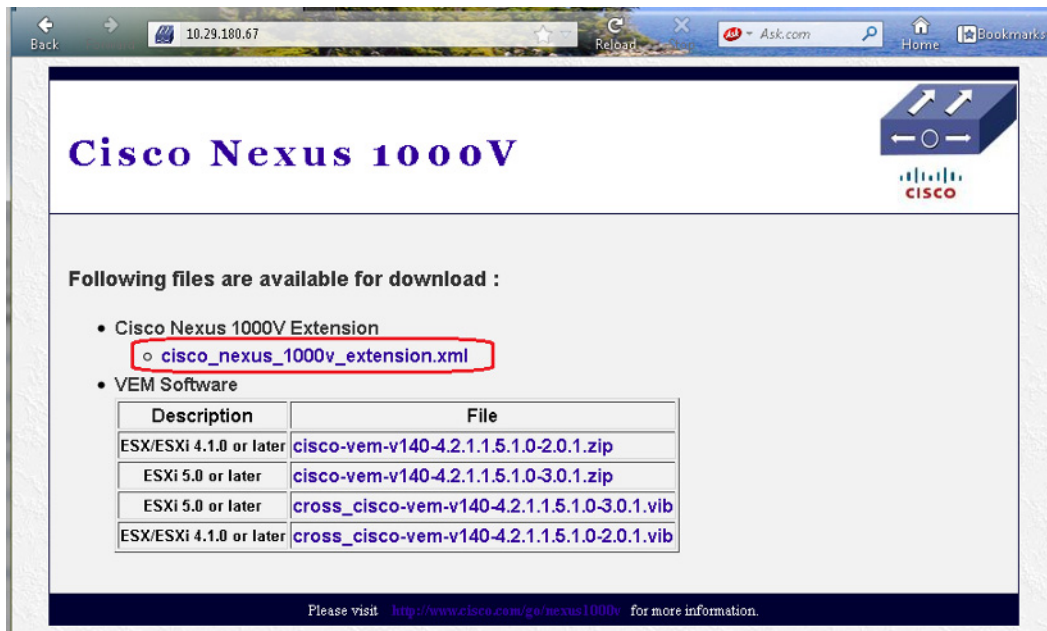
At the bottom of the window, there is a status bar with the text: "License Period: 100 days remaining", "To release cursor, press CTRL+ALT", and "root".

## Connecting Cisco Nexus 1000v VSM to VMware vCenter

When the initial setup of VMS VM is complete, you need to add it as a plugin / extension in the vCenter. To add it as a plug-in, follow these steps:

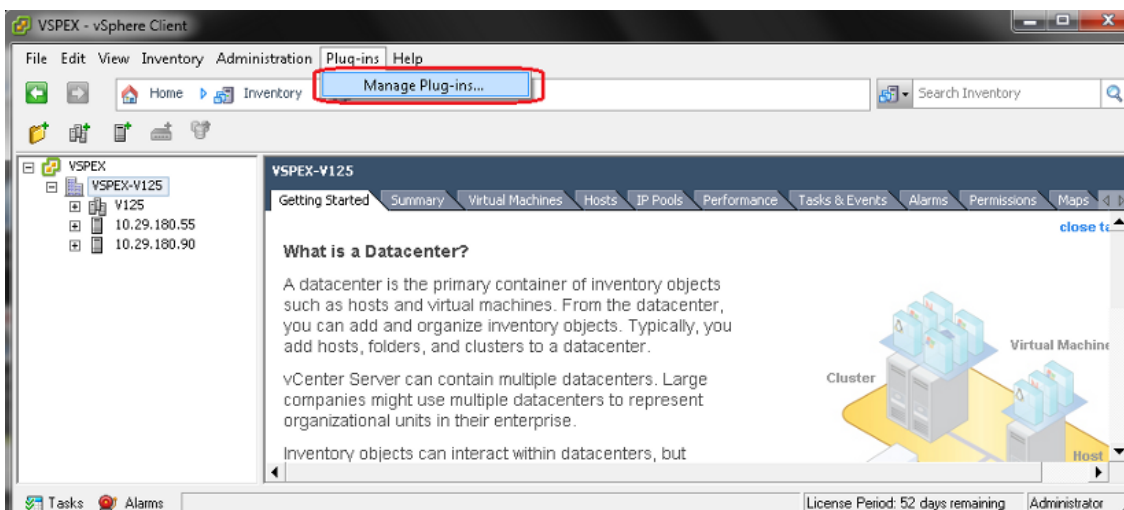
1. Using your browser, access management IP address of the VSM VM.
2. Right-click the `cisco_nexus_1000v_extension.xml` link and save to a location on your local hard drive.

**Figure 164** Downloading the Cisco Nexus 1000v Extension




- From the VMware vSphere client, click **Plug-ins > Manage Plug-ins**.

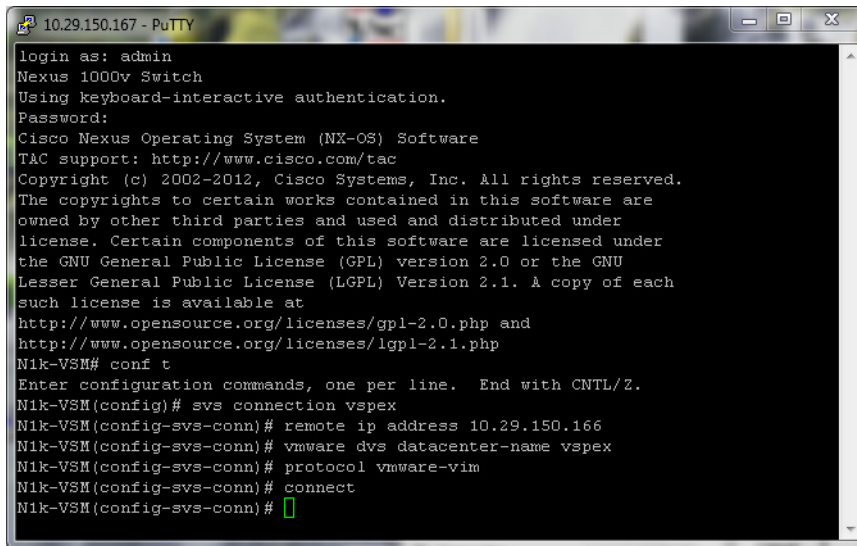
**Figure 165** Adding Plug-In Manage Plug-in Window



- Scroll to the bottom of Available Plug-ins, right-click in the empty space and choose **New Plug-in**.
- Click **Browse**, choose the `cisco_nexus-1000v_extension.xml` file you have downloaded.
- Click **Register Plug-in**.

**Figure 166**     **Registering Plug-in**

7. If you receive a certificate warning, click **Ignore**.
8. Click **OK**. The Plug-in Manager window appears showing the plug-in that was just added.
9. Configure the SVS connection to the vCenter as shown in [Figure 167](#).

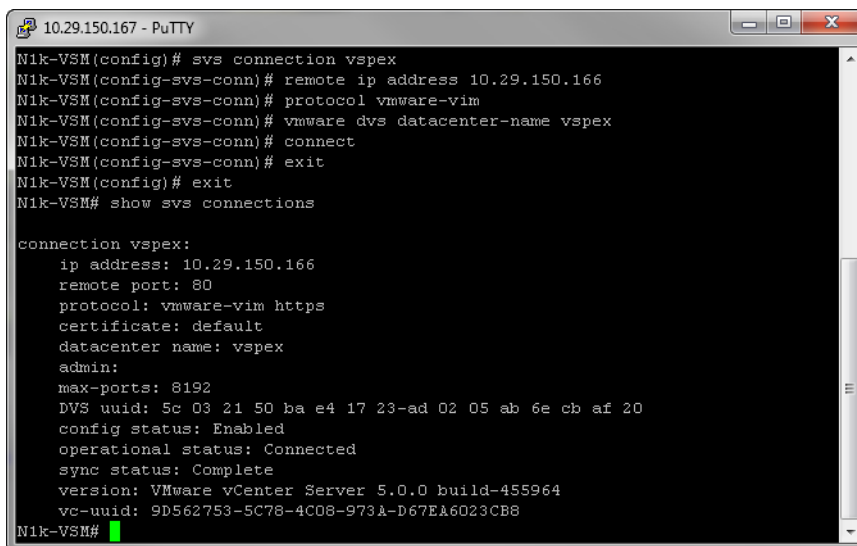
**Figure 167**      **Configuring SVS Connection to vCenter**


```

10.29.150.167 - PuTTY
login as: admin
Nexus 1000v Switch
Using keyboard-interactive authentication.
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2012, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
N1k-VSM# conf t
Enter configuration commands, one per line. End with CNTL/Z.
N1k-VSM(config)# svcs connection vspex
N1k-VSM(config-svs-conn)# remote ip address 10.29.150.166
N1k-VSM(config-svs-conn)# vmware dvs datacenter-name vspex
N1k-VSM(config-svs-conn)# protocol vmware-vim
N1k-VSM(config-svs-conn)# connect
N1k-VSM(config-svs-conn)#

```

10. Validate the connection using “show svcs connection” and ensure that the operational status is “connected” and sync status is “Complete” as shown in [Figure 168](#).

**Figure 168**      **Verifying SVS Connection**


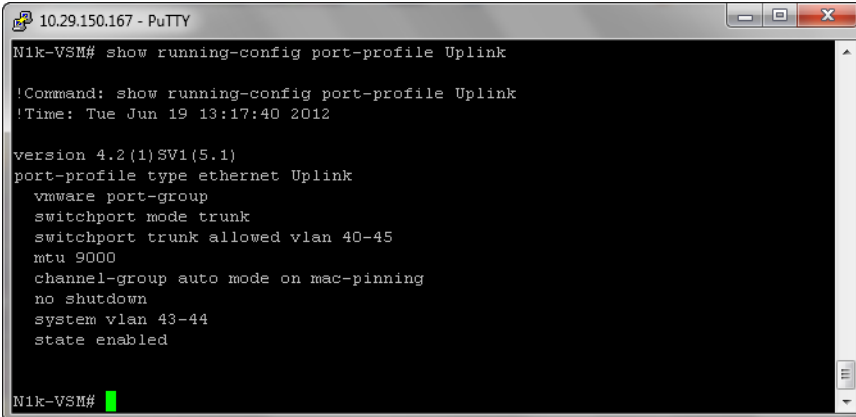
```

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

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

```

11. You must configure at least one uplink port-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 uplink port-profile as shown in [Figure 169](#).

**Figure 169**      **Configuring Uplink Port-Profile**


```

10.29.150.167 - PuTTY
N1k-VSM# show running-config port-profile Uplink

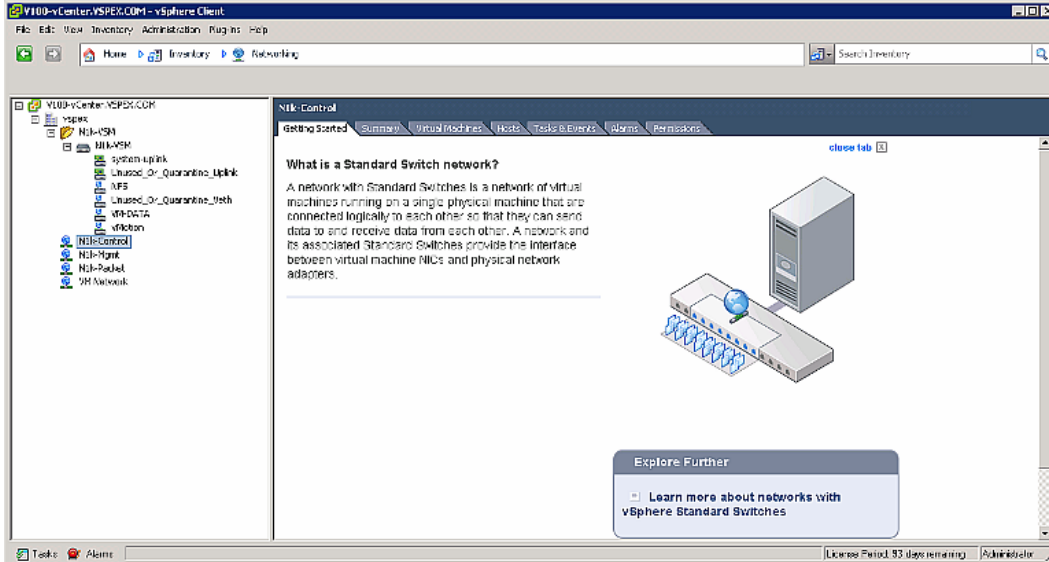
!Command: show running-config port-profile Uplink
!Time: Tue Jun 19 13:17:40 2012

version 4.2(1)SV1(5.1)
port-profile type ethernet Uplink
  vmware port-group
    switchport mode trunk
    switchport trunk allowed vlan 40-45
    mtu 9000
    channel-group auto mode on mac-pinning
    no shutdown
    system vlan 43-44
    state enabled
N1k-VSM#

```

Notice that uplink port-profile uses trunk with storage, vMotion, N1k control, N1k packet and all the necessary virtual machine data VLANs. MTU 9000 is configured on uplink port-profile to enable jumbo frames. “channel-group auto mode on mac-pinning” is a very important configuration which ‘pins’ the VM VNICs to uplinks on the vDS. MAC pinning feature does static load balancing on per VNIC basis. It also provides high-availability by moving the traffic to the alternative adapter when a given fabric is down

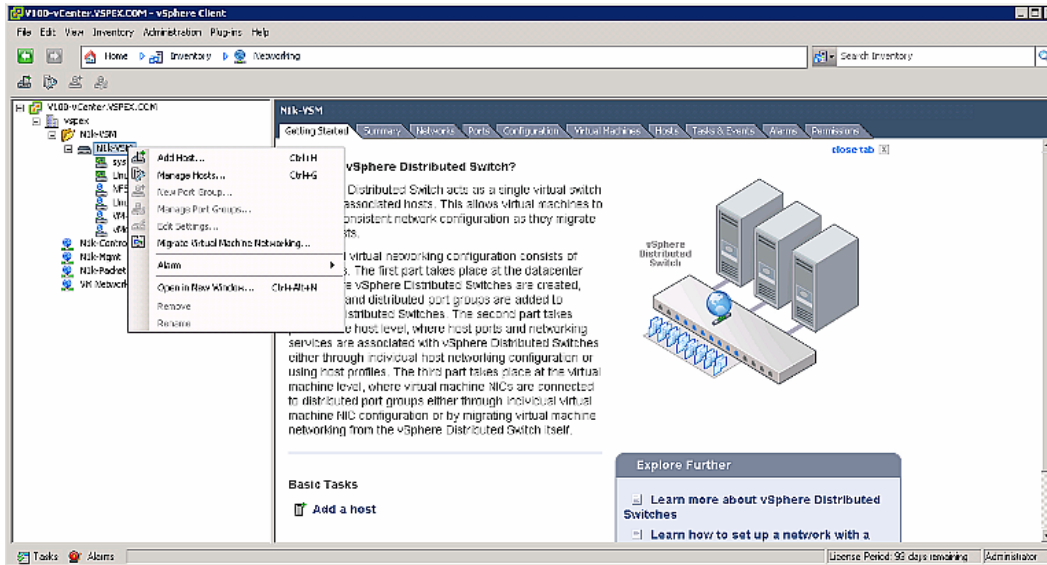
12. When VSM is connected to the vCenter, it shows up as a virtual Distributed Switch in the vCenter’s “Network” view as shown in [Figure 170](#).

**Figure 170**      **Window Showing N1K-Control in vCenter**

13. Add hosts to the vDS as shown in [Figure 171](#).

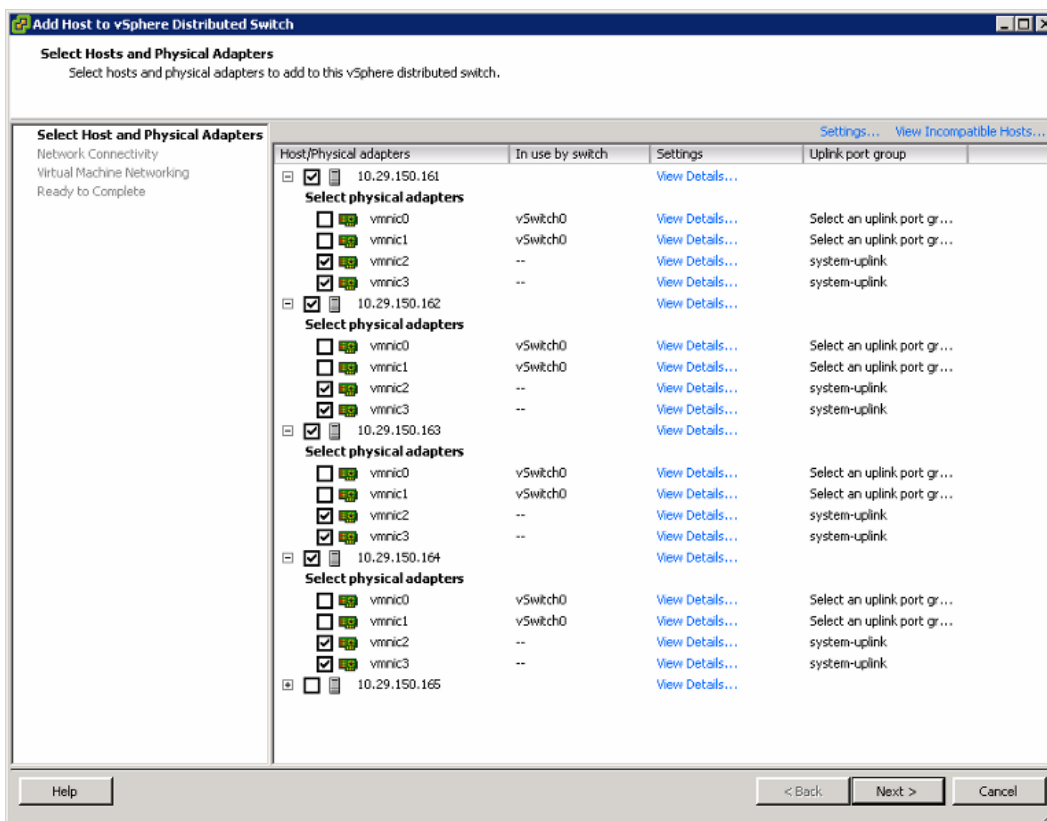


Figure 171 Adding Host in vCenter



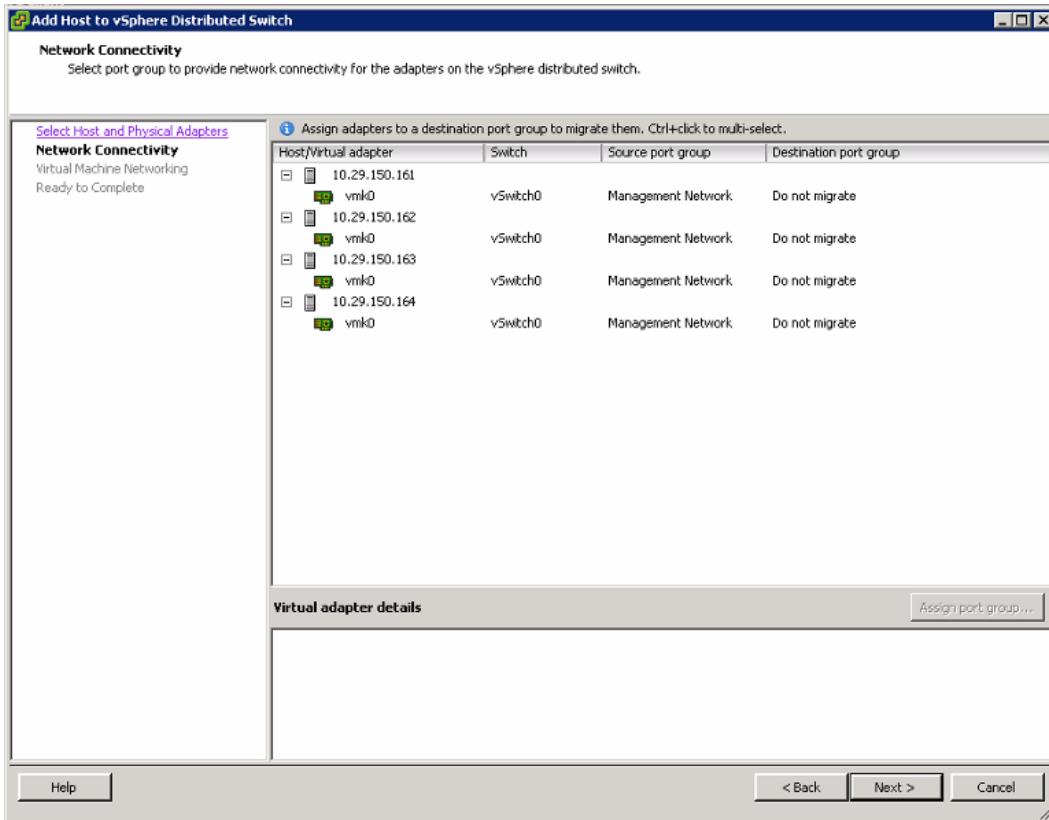
14. On the next dialog box, select all the VSPEX ESXi hosts and add unclaimed adapters (10 GigE links in case of 100 and 125 Virtual Machine architectures) using the uplink port-profile created in the previous step.

Figure 172 Selecting Host and Physical Adapters

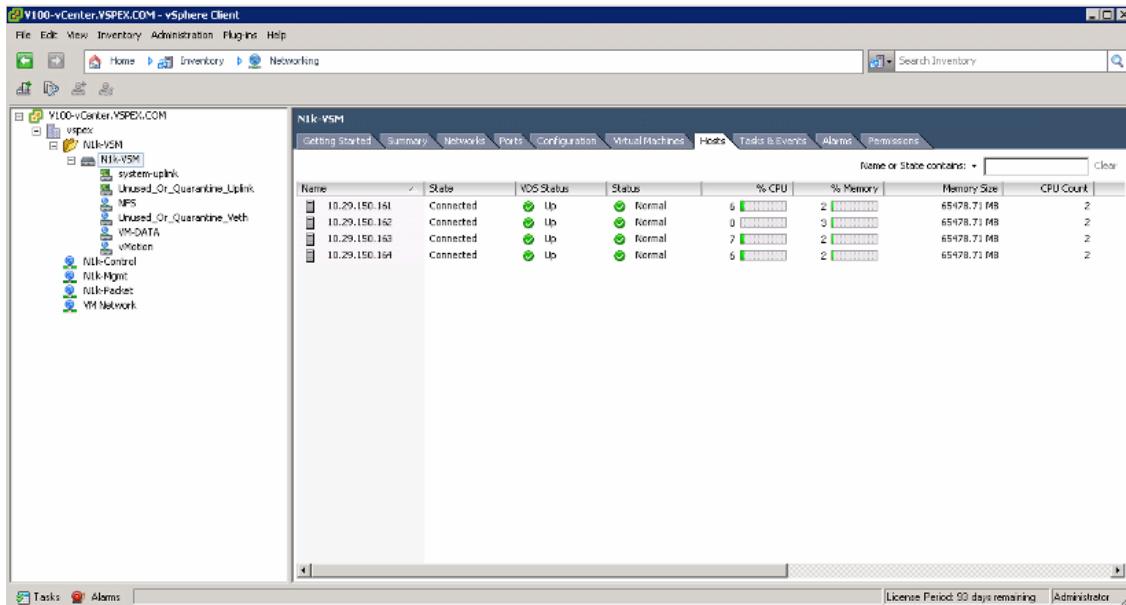


15. Do not migrate management VM kernel from the native vSwitch to vDS in the last step, click **Next** and click **Finish** to exit the add host wizard.

**Figure 173**      *Selecting Port Group for Network Connectivity*



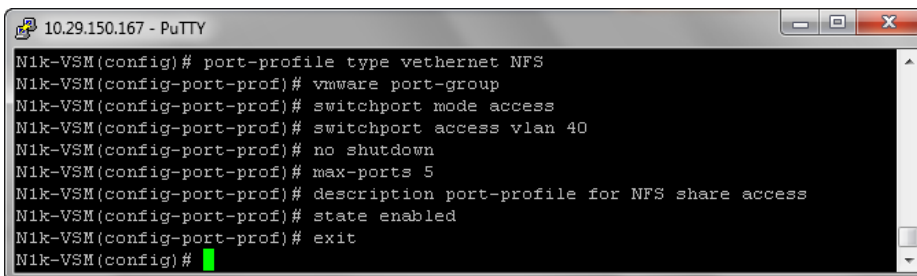
16. Finally, ensure that all the hosts are successfully added to the vDS.

**Figure 174**      **Verifying Successful Addition of Hosts**

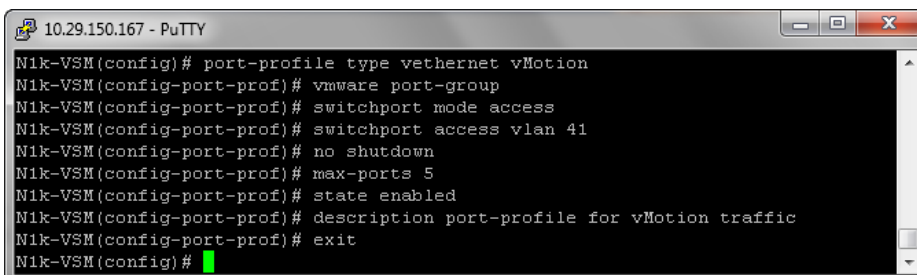
## Configuring Port-Profile in VSM and Migrate vCenter Networking to vDS

The last step of Nexus 1000v configuration and its integration with vCenter is creation of port-profiles and using them in the virtual machines in the vCenter. Use following steps to configure these steps:

1. Create a port-profile for storage (NFS) access.

**Figure 175**      **Creating Port-profiles for NFS**

2. Create a port-profile for vMotion traffic.

**Figure 176**      **Creating Port-profiles for vMotion Traffic**

3. Create port-profiles for the virtual machine data traffic used by various applications as per your needs. You can set “max ports” to appropriate values based on the number of Virtual Machines being configured. [Figure 177](#) shows a sample port-profile.

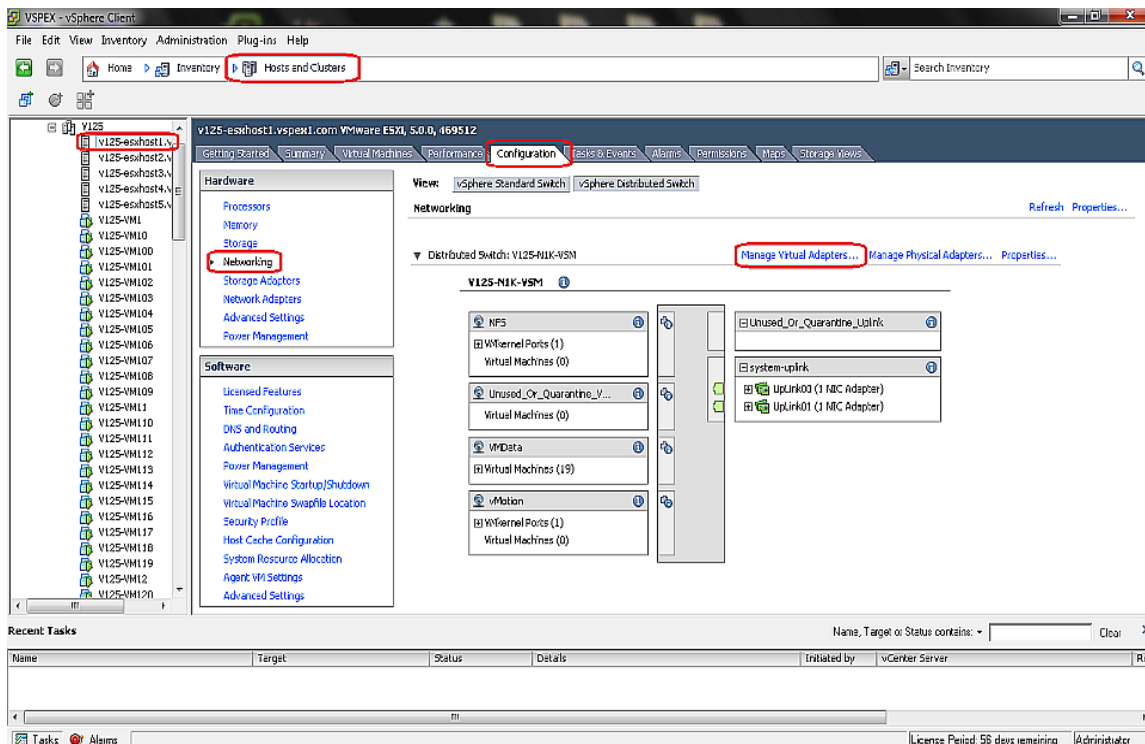
**Figure 177**      **Creating Port-profiles for VM Data Traffic**

```

10.29.150.167 - PuTTY
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, choose “Hosts and Clusters” tab, choose the ESXi host, click **Configuration** tab > **Networking**. In the View pane, choose **vSphere Distributed Switch**, and click “Manager Virtual Adapters...” link as shown in [Figure 178](#).

**Figure 178**      **Managing Virtual Adapters in vCenter**



5. Click **Add** on the wizard, click **New virtual adapter**, and click **Next**.
6. Choose “VMKernel” on the next dialog box.
7. Choose port-profile “NFS” for the storage access, and click **Next**.
8. Configure IP address from the NFS subnet and configure subnet mask. Click “Next” and “Finish” to deploy the VNIC.

9. Similarly, add one more VMKNic (VM Kernel NIC) for vMotion. When providing the port-profile name, ensure that you choose “vMotion” port-profile and check the check box “Use this virtual adapter for vMotion”.
10. Repeat creation of the two vmknics virtual adapters for all the ESXi hosts.
11. 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.
12. When the NFS share is reachable, the NFS datastore can be discovered and mounted through the vCenter. Virtual machines can be deployed on these NFS datastore using the VM-Data port-profile for the network access.
13. Verify the port-profile usage using “show port-profile brief”, “show port-profile usage”, or “show port-profile name <name>” command. [Figure 179](#) shows two sample outputs.

**Figure 179**      **Verifying Port-Profiles**

```
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
Uplink	Ethernet	1	5	5	12	0
VM-DATA	Vethernet	1	3	3	100	0
uplink	Ethernet	0	2	2	0	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

[Figure 180](#) shows the output of uplink port-profile usage, notice 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 are also listed as assigned interfaces. Port-channel status can be further viewed / validated using “show port-channel brief” command from VSM VM.

**Figure 180**      **Verifying Post-Profile Uplinks**

```

10.29.150.167 - PuTTY
Nik-VSM# show port-profile name Uplink

port-profile Uplink
  type: Ethernet
  description:
  status: enabled
  max-ports: 32
  min-ports: 1
  inherit:
  config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 40-45
    mtu 9000
    channel-group auto mode on mac-pinning
    no shutdown
  evaluated config attributes:
    switchport mode trunk
    switchport trunk allowed vlan 40-45
    mtu 9000
    channel-group auto mode on mac-pinning
    no shutdown
  assigned interfaces:
    port-channel4
    port-channel8
    port-channel9
    port-channel10
    Ethernet3/3
    Ethernet3/4
    Ethernet4/3
    Ethernet4/4
    Ethernet5/3
    Ethernet5/4
    Ethernet6/3
    Ethernet6/4
  port-group: Uplink
  system vlans: 43-44
  capability l3control: no
  capability iscsi-multipath: no
  capability vxlan: no
  capability l3-vn-service: no
  port-profile role: none
  port-binding: static

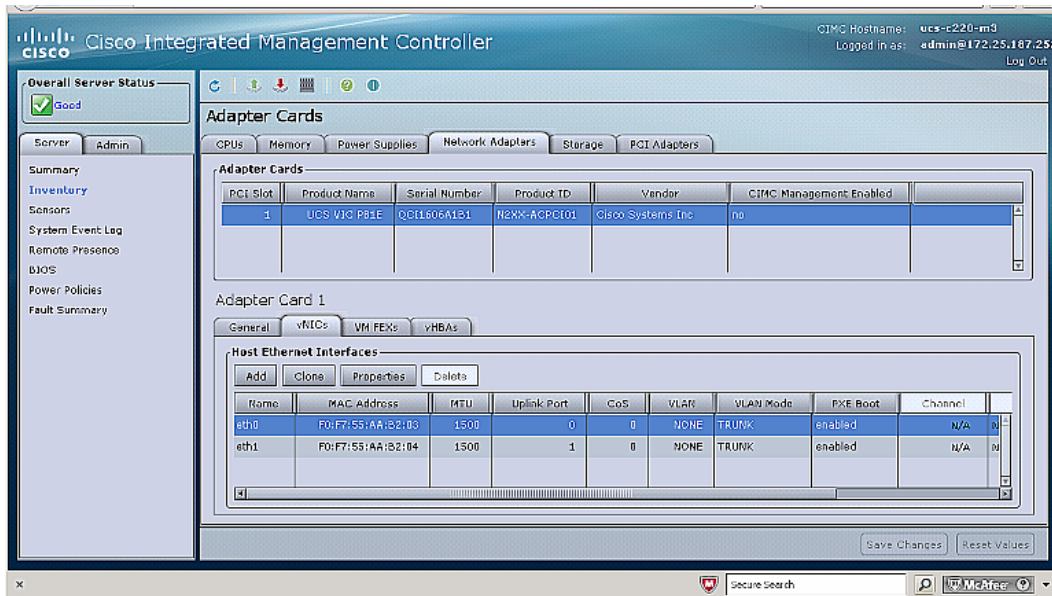
Nik-VSM#

```

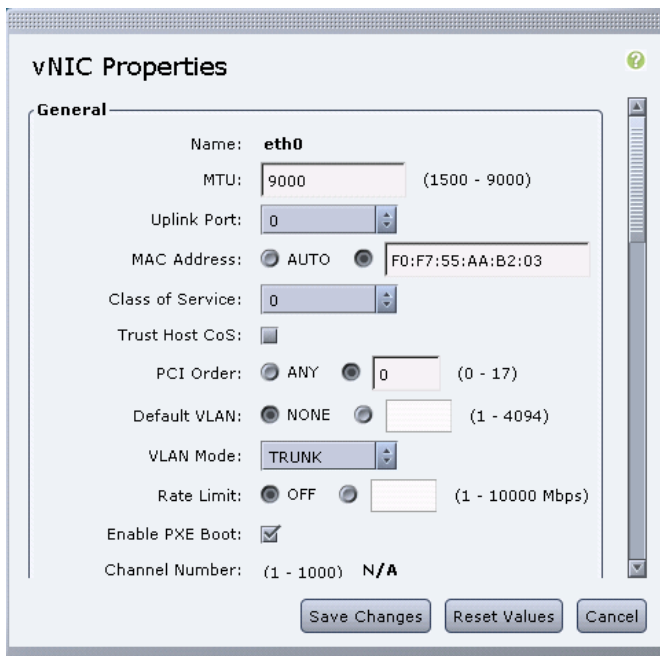
## Configuring Jumbo Frame at the CIMC Interface

To make 9000 MTU work end-to-end, the last piece of the jumbo frame puzzle need to be solved at the CIMC adapter level. Use following steps to configure 9000 MTU on physical adapter:

1. Using the web browser, connect to each of the servers' CIMC management IP address and provide username/password.
2. Click **Inventory** on the left side, and **Network Adapters** tab on the right side. Choose the vNICs tab and click **eth0** vNIC. Click **Properties**. The sample GUI is shown in [Figure 181](#).

**Figure 181** Viewing Details of Adapter Cards in CIMC Inventory

3. On the “vNIC Properties” window, change the MTU to 9000 value and click **Save Changes** button.

**Figure 182** Window Showing vNIC Properties

## Jumbo MTU Validation and Diagnostics

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

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

```

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

```

```

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

```

Ensure that the packet size is 8972 due to various L2/L3 overheads. Ping need to be successful. If ping is not successful, verify the 9000 MTU configured at each of these steps:

1. 9000 MTU on the NFS share IP address on the VNX/VNXe storage device(s).
2. Ensure that a “jumbo-mtu” policy map is created at the Cisco Nexus 5000 / 3000 series servers with default class having MTU 9000. Ensure that the “jumbo-mtu” policy is applied to system classes on the ingress traffic.
3. Ensure that the “mtu 9000” is set on uplink port-profile in the Cisco Nexus 1000v.
4. Ensure that the 9000 MTU is set for vmkernel ports, used for vMotion as well as storage access VNICS.
5. Ensure that the CIMC configuration is validated.

## Validating Cisco Solution for EMC VSPEX VMware Architectures

This section provides a list of items that needs to be reviewed after the solution is 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.

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

Verify the redundancy of the solution components



Following redundancy checks were performed at the Cisco lab to verify solution robustness:

1. Administratively shutdown one of the two data links connected to the server. 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 one of the two data links connected to the storage array. Make sure that storage is still available from all the ESXi hosts. Upon administratively enabling the shutdown port, the traffic should be rebalanced.
3. Reboot one of the two Nexus switches while storage and network access from the servers are going on. The switch reboot should not affect the operations of storage and network access from the Virtual Machines. Upon rebooting the switch, the network access load should be rebalanced across the two switches.
4. Reboot the active storage processor of the VNX/VNXe storage array and make sure that all the NFS shares are still accessible during and after the reboot of the storage processor.
5. Fully load all the virtual machines of the solution. Put one of the ESXi host in maintenance mode. All the Virtual Machines 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. Note that in 50 and 125 virtual machines architectures, there is enough head room for memory in other servers to accommodate 25 additional virtual machines. However, for 100 virtual machines solution, memory would be oversubscribed when one of the ESXi host goes down. So, for 100 virtual machines solution, vCenter memory compression or dynamic memory commitment features should be used to oversubscribe physical memory on the remaining hosts.

## Cisco Validation Test Profile

“vdbench” testing tool was used with Microsoft Windows 2008 R2 SP1 server to test scaling of the solution in Cisco labs. [Table 15](#) provides information on the test profile used.

**Table 15**      **Test Profile Details**

Profile Characteristics	Value
Number of virtual machines	50, 100 or 125 depending on architecture
Virtual machine OS	Windows Server 2008 R2 SP1
Processors per virtual machine	1
Number of virtual processors per physical CPU core	4
RAM per virtual machine	2 GB
Average storage available for each virtual machine	100 GB
Average IOPS per virtual machine	25 IOPS
Number of datastores to store virtual machine disks	2
Disk and RAID type for datastores	RAID 5, 600 GB, 15k rpm, 3.5-inch SAS disks

# Bill of Materials

Table 16 provides the details of the components used in the CVD for 50 virtual machines configuration.

**Table 16**      **Component Description**

Description	Part Number
UCS C220 M3 rack servers	UCSC-C220-M3S
CPU for C220 M3 rack servers	UCS-CPU-E5-2650
Memory for C220 M3 rack servers	UCS-MR-1X082RY-A
RAID local storage for rack servers	UCSC-RAID-11-C220
Cisco VIC adapter for 100 and 125 VMs solutions	N2XX-ACPCI01
Broadcom 1Gbps adapter for 50 VMs solution	N2XX-ABPCI03-M3
Nexus 5548UP switches for 100 and 125 VMs solutions	N5K-C5548UP-FA
Nexus 3048 switches for 50 VMs solution	N3K-C3048TP-1GE
10 Gbps SFP+ multifiber mode	SFP-10G-SR

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

[http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/C220M3\\_SFF\\_SpecSheet.pdf](http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/C220M3_SFF_SpecSheet.pdf).

## Customer Configuration Data Sheet

Before you start the configuration, gather some customer-specific network and host configuration information. Table 17, Table 18, Table 19, Table 20, Table 21, Table 22 provide information on assembling the required network and host address, numbering, and naming information. This worksheet can also be used as a “leave behind” document for future reference.

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

**Table 17**      **Common Server Information**

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

**Table 17 Common Server Information**

Server Name	Purpose	Primary IP
	vCenter Console	
	SQL Server	

**Table 18 ESXi server Information**

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

**Table 19 Array Information**

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

**Table 20 Network Infrastructure Information**

Name	Purpose	IP	Subnet Mask	Default Gateway
	Cisco Nexus 5548UP A / Cisco Nexus 3048 A			
	Cisco Nexus 5548UP B / Cisco Nexus 3048 B			
	Cisco Nexus 1000v VSM			

**Table 21** *VLAN Information*

<b>Name</b>	<b>Network Purpose</b>	<b>VLAN ID</b>	<b>Allowed Subnets</b>
vlan-infra	Virtual Machine Networking VMware ESXi Management		
vlan-nfs	NFS storage network		
vlan-vMotion	VMware vMotion traffic network		
vlan-control	Control VLAN for Cisco Nexus 1000v switch	N/A	
vlan-packet	Packet VLAN for Cisco Nexus 1000v switch	N/A	
vlan-data (multiple)	Data VLAN of customer VMs as needed		

**Table 22** *Service Accounts*

<b>Account</b>	<b>Purpose</b>	<b>Password (option, secure)</b>
	Microsoft Windows server Administrator	
Root	VMware ESXi root	
	Array administrator	
	VMware vCenter administrator	
	Microsoft SQL server administrator	
	Cisco Nexus 5548UP administrator	
	Cisco Nexus 1000v administrator	

## References

Cisco UCS:

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

VMware vSphere:

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

Cisco Nexus:

[http://www.cisco.com/en/US/products/ps9441/Products\\_Sub\\_Category\\_Home.html](http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html)

Cisco Nexus 5000 Series NX-OS Software Configuration Guide:

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

Cisco Nexus 1000v virtual switch Software Configuration Guide

[http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4\\_2\\_1\\_s\\_p\\_1\\_2/software/configuration/guide/n1010\\_vsvcs\\_cfg.html](http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_2_1_s_p_1_2/software/configuration/guide/n1010_vsvcs_cfg.html)

EMC VNXe3xxx series resources:

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

EMC VNX5xxx series resources:

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

Microsoft SQL Server installation guide:

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