Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 50 Virtual Machines

Last Updated: June 26, 2013

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Sanjeev Naldurgkar

Sanjeev Naldurgkar, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems

Sanjeev Naldurgkar is a Technical Marketing Engineer at Cisco Systems with Server Access Virtualization Business Unit (SAVBU). With over 12 years of experience in information technology, his focus areas include UCS, Microsoft product technologies, server virtualization, and storage technologies. Prior to joining Cisco, Sanjeev was Support Engineer at Microsoft Global Technical Support Center. Sanjeev holds a Bachelor's Degree in Electronics and Communication Engineering and Industry certifications from Microsoft, and VMware.

Acknowledgements

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

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

About Cisco Validated Design (CVD) Program

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

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

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

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

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco Virtualization Solution for EMC VSPEX with VMware vSphere 5.1 for 50 Virtual Machines

Executive Summary

Cisco solution for EMC VSPEX proven and modular infrastructures are built with best of-breed technologies to create complete virtualization solutions that enable you to make an informed decision in the hypervisor, compute, and networking layers. VSPEX eases server virtualization planning and configuration burdens. VSPEX 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 virtual machines with Cisco solution for EMC VSPEX.

Introduction

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

Target Audience

The reader of this document is expected to have the necessary training and background to install and configure VMware vSphere, EMC VNXe3150, Cisco Nexus 3048 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

Purpose of this Document

This document describes the steps required to deploy and configure the Cisco solution for EMC VSPEX for VMware architecture to a level that will allow for confirmation that the basic components and connections are working correctly. This CVD covers the VMware vSphere 5.1 for 50 Virtual Machines private cloud architecture. While readers of this CVD are expected to have sufficient knowledge to install and configure the products used, configuration details that are important for deploying this solution are specifically mentioned.

The 50 virtual machine environment discussed is based on a defined reference workload. While not every virtual machine has the same requirement, this document contains methods and guidance to adjust the system to be cost-effective when deployed.

A private cloud architecture is a complex system offering. This document facilitates its setup by providing up-front software and hardware material lists, step-by-step sizing guidance and worksheets, and verified deployment steps. When the last component has been installed, there are validation tests to ensure that your system is operating properly. Following the procedures defined in this document ensures an efficient and painless journey to the cloud.

Business Needs

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

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

- Provide an end-to-end virtualization solution to take full advantage of unified infrastructure components.
- Provide a Cisco VSPEX for VMware ITaaS (IT as a Service) solution for efficiently virtualizing up to 50 virtual machines for varied customer use cases.
- Provide a reliable, flexible and scalable reference design.

Solutions Overview

This section provides a list of components used for deploying the Cisco solution for EMC VSPEX for 50 VMs using VMware vSphere 5.1.

Cisco Solution for EMC VSPEX with VMware Architectures

This solution provides an end-to-end architecture with Cisco, EMC, VMware, and Microsoft technologies that demonstrate support for up to 50 generic virtual machines and provide high availability and server redundancy.

Following are the components used for the design and deployment:

- Cisco C-series Unified Computing System servers
- Cisco Nexus 3000 Series Switch
- Cisco virtual Port Channels (vPC) for network load balancing and high availability
- EMC VNXe3150 storage components
- EMC Next Generation Backup Solutions
- VMware vSphere 5.1
- Microsoft SQL Server Database
- VMware DRS
- VMware HA

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

Technology Overview

Cisco Unified Computing System

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

The main components of the Cisco UCS are:

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

The Cisco Unified Computing System is designed to deliver:

- A reduced Total Cost of Ownership (TCO) and increased business agility.
- Increased IT staff productivity through just-in-time provisioning and mobility support.

- A cohesive, integrated system which unifies the technology in the data center.
- Industry standards supported by a partner ecosystem of industry leaders.

Cisco UCS C220 M3 Rack-Mount Servers

Building on the success of the Cisco UCS C200 M2 Rack Servers, the enterprise-class Cisco UCS C220 M3 server further extends the capabilities of the Cisco Unified Computing System portfolio in a 1-rack-unit (1RU) form factor. And with the addition of the Intel® Xeon® processor E5-2600 product family, it delivers significant performance and efficiency gains.

Figure 1 Cisco UCS C220 M3 Rack Server

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

Cisco 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.

Figure 2 Cisco Nexus 3048 Switch

EMC Storage Technologies and Benefits

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. The EMC VNXe 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 iSCSI 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 Available

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

Software Packs Available

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

This is the available EMC protection software pack.

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. The EMC VNXe storage arrays are multi-protocol platforms that can support the iSCSI, NFS, and CIFS protocols depending on the customer's specific needs. The solution was validated using iSCSI for data storage.

EMC NetWorker and Data Domain

EMC's NetWorker coupled with Data Domain deduplication storage systems seamlessly integrate into virtual environments, providing rapid backup and restoration capabilities. Data Domain deduplication results in vastly less data traversing the network by leveraging the Data Domain Boost technology, which greatly reduces the amount of data being backed up and stored, translating into storage bandwidth and operational savings.

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

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

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

EMC Avamar

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

VMware vSphere 5.1

VMware vSphere 5.1 transforms a computer's physical resources by virtualizing the CPU, memory, storage and network functions. This transformation creates fully functional virtual machines that run isolated and encapsulated operating systems and applications just like physical computers.

The high availability features of VMware vSphere 5.1 such as vMotion and Storage vMotion enable seamless migration of virtual machines and stored files from one vSphere server to another with minimal or no performance impact. Coupled with vSphere DRS and Storage DRS, virtual machines have access to the appropriate resources at any point in time through load balancing of compute and storage resources.

VMware vCenter

VMware vCenter is a centralized management platform for the VMware Virtual Infrastructure. It provides administrators with a single interface for all aspects of monitoring, managing, and maintaining the virtual infrastructure that can be accessed from multiple devices.

VMware vCenter is also responsible for managing some of the more advanced features of the VMware virtual infrastructure like VMware vSphere High Availability (vSphere HA) and Distributed Resource Scheduling (DRS), along with the vMotion and Update Manager.

Architectural overview

This CVD focuses on VMware solution for up to 50 virtual machines.

For the VSPEX solution, the reference workload was defined as a single virtual machine. Characteristics of a virtual machine are defined in Table 1.

Table 1 Virtual Machine Characteristics

Characteristics	Value
Virtual machine operating system	Microsoft Windows Server 2012
Virtual processor per virtual machine (vCPU)	1
RAM per virtual machine	2 GB
Available storage capacity per virtual machine	100 GB
I/O operations per second (IOPS) per VM	25
I/O pattern	Random
I/O read/write ratio	2:1

See "Sizing Guideline" section on page 20 for more detailed information.

Solution architecture overview

Table 2 lists the mix of hardware components, their quantities and software components used in this architecture.

Table 2 Hardware and Software Components

Components	Specifications
Servers	3 Cisco C220 M3 Rack-Mount Servers
Adapters	1 Broadcom NetXtreme II 5709 quad-port per server
Network Switches	2 Cisco Nexus 3048 switches
Storage	EMC VNXe3150
Network Speed	1G Ethernet
Hypervisor	VMware ESXi 5.1

Table 3 lists the various hardware and software versions of the components which occupies different tiers of the Cisco solution for EMC VSPEX with VMware architectures under test.

Table 3 Firmware and Software Versions of Components

Vendor	Name	Version	Description
Cisco	C220 M3 Servers	1.4(7a) – CIMC C220M3.1.4.7b.0 - BIOS	Cisco C220M3 Rack Servers
Cisco	Nexus 3048 Switches	5.0(3)U2(2b)	Nexus 3000 Series Switches running NX-OS
EMC	VNXe3150	2.4.0.20932	VNXe Storage Array
EMC	Avamar	6.1 SP1	EMC Next-Generation Backup
EMC	NetWorker	6.1 SP1	EMC Next-Generation Backup
EMC	Data Domain OS	8.0 SP1	Data Next-Generation Backup
VMware	ESXi 5.1	5.1 build 799733	Hypervisor
VMware	vCenter Server	5.1 build 799731	VMware Management
Micorsoft	Windows Server 2008 R2	2008 R2 SP1	OS to Host vCenter Server
Microsoft	SQL Server	2008 R2	Database Server SQL R2 Enterprise Edition for vCenter

Table 4 outlines the C220 M3 server configuration across all the VMware architectures. Table 4 shows the configuration on per server basis.

Table 4 Cisco UCS C220 M3 Server Hardware Configuration

Components	Capacity
Memory (RAM)	64 GB (8x8MB DIMM)
Processor	2 x Intel® Xeon ® E5-2650 CPUs, 2 GHz, 8 cores, 16 threads
Local Storage	Cisco UCS RAID SAS 2008M-8i Mezzanine Card, With 2 x 67 GB slots for RAID 1 configuration each

The reference architecture assumes that there is an existing infrastructure/ management network available where a virtual machine hosting vCenter server and Windows Active Directory/ DNS server are present. The below diagram illustrates high level solution architecture for 50 virtual machines.

Figure 3 Reference Architecture for 50 Virtual Machines

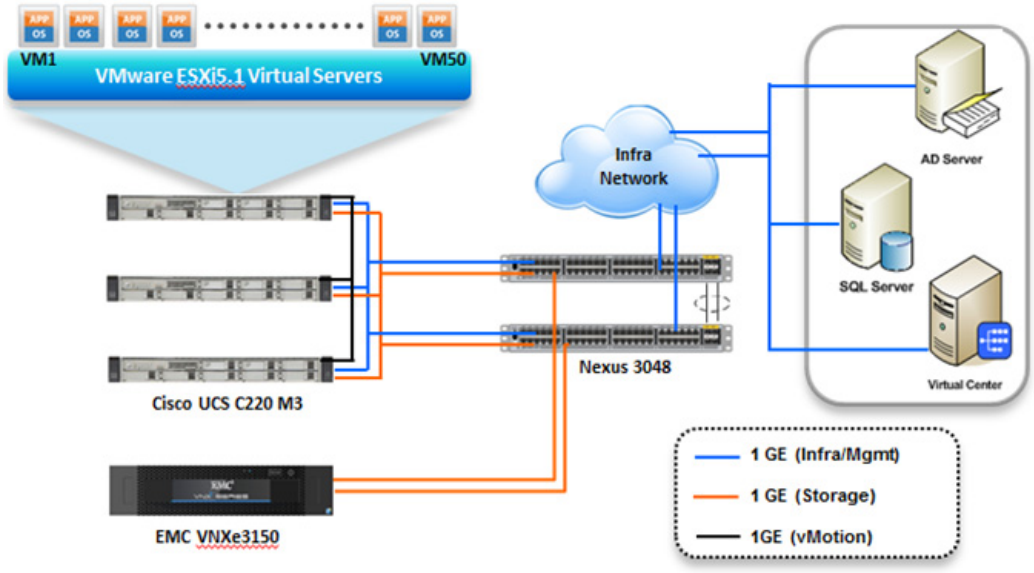

As it is evident in the above diagrams, following are the high level design points of VMware architectures:

- 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 dictate or require any specific layout of infrastructure network. The vCenter server and Microsoft Windows Active Directory are hosted on infrastructure network. However, design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

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

Memory Configuration Guidelines

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

ESX/ESXi Memory Management Concepts

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

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

For more information about vSphere memory management concepts, see the VMware vSphere Resource Management Guide at: http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf

Virtual Machine Memory Concepts

Figure 4 shows the use of memory settings parameters in the virtual machine.

Figure 4 Virtual Machine Memory Settings

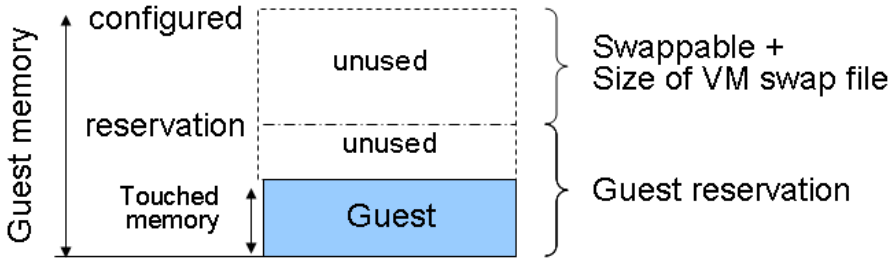

The vSphere memory settings for a virtual machine include the following parameters:

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

Allocating Memory to Virtual Machines

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 1 outlines the resources used by a single virtual machine:

Following are the recommended best practices:

- Account for memory overhead—Virtual machines require memory beyond the amount allocated, and this memory overhead is per-virtual machine. Memory overhead includes space reserved for virtual machine devices, depending on applications and internal data structures. The amount of overhead required depends on the number of vCPUs, configured memory, and whether the guest operating system is 32-bit or 64-bit. As an example, a running virtual machine with one virtual CPU and two GB of memory may consume about 100 MB of memory overhead, where a virtual machine with two virtual CPUs and 32 GB of memory may consume approximately 500 MB of memory overhead. This memory overhead is in addition to the memory allocated to the virtual machine and must be available on the ESXi host.
- "Right-size" memory allocations—Over-allocating memory to virtual machines can waste memory unnecessarily, but it can also increase the amount of memory overhead required to run the virtual machine, thus reducing the overall memory available for other virtual machines. Fine-tuning the memory for a virtual machine is done easily and quickly by adjusting the virtual machine properties. In most cases, hot-adding of memory is supported and can provide instant access to the additional memory if needed.
- Intelligently overcommit—Memory management features in vSphere allow for over commitment of physical resources without severely impacting performance. Many workloads can participate in this type of resource sharing while continuing to provide the responsiveness users require of the application. When looking to scale beyond the underlying physical resources, consider the following:
  - Establish a baseline before over committing. Note the performance characteristics of the application before and after. Some applications are consistent in how they utilize resources and may not perform as expected when vSphere memory management techniques take control. Others, such as Web servers, have periods where resources can be reclaimed and are perfect candidates for higher levels of consolidation.
  - Use the default balloon driver settings. The balloon driver is installed as part of the VMware Tools suite and is used by ESX/ESXi if physical memory comes under contention. Performance tests show that the balloon driver allows ESX/ESXi to reclaim memory, if required, with little to no impact to performance. Disabling the balloon driver forces ESX/ESXi to use host-swapping to make up for the lack of available physical memory which adversely affects performance.
  - Set a memory reservation for virtual machines that require dedicated resources. Virtual machines running Search or SQL services consume more memory resources than other application and Web front-end virtual machines. In these cases, memory reservations can guarantee that the services have the resources they require while still allowing high consolidation of other virtual machines.

Storage Guidelines

VSPEX architecture for VMware virtual machines at 50 VMs scale uses iSCSI to access storage arrays. This simplifies the design and implementation for the small to medium level businesses. vSphere provides many features that take advantage of EMC storage technologies such as VNX VAAI plug-in for

NFS storage and storage replication. Features such as VMware vMotion, VMware HA, and VMware Distributed Resource Scheduler (DRS) use these storage technologies to provide high availability, resource balancing, and uninterrupted workload migration.

Storage Protocol Capabilities

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

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

For more information, see the VMware white paper Comparison of Storage Protocol Performance in VMware vSphere 5.1: http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf

Storage Best Practices

Following are the vSphere storage best practices:

- Host multi-pathing—Having a redundant set of paths to the storage area network is critical to protecting the availability of your environment. In this solution, the redundancy comes from the "Fabric Failover" feature of the dynamic vNICs of Cisco UCS for NFS storage access.
- Partition alignment—Partition misalignment can lead to severe performance degradation due to I/O operations having to cross track boundaries. Partition alignment is important both at the NFS level as well as within the guest operating system. Use the vSphere Client when creating NFS datastores to be sure they are created aligned. When formatting volumes within the guest, Windows 2008 aligns NTFS partitions on a 1024KB offset by default.
- Use shared storage—In a vSphere environment, many of the features that provide the flexibility in management and operational agility come from the use of shared storage. Features such as VMware HA, DRS, and vMotion take advantage of the ability to migrate workloads from one host to another host while reducing or eliminating the downtime required to do so.
- Calculate your total virtual machine size requirements—Each virtual machine requires more space than that used by its virtual disks. Consider a virtual machine with a 20GB OS virtual disk and 16GB of memory allocated. This virtual machine will require 20GB for the virtual disk, 16GB for the virtual machine swap file (size of allocated memory), and 100MB for log files (total virtual disk size + configured memory + 100MB) or 36.1GB total.
- Understand I/O Requirements—Under-provisioned storage can significantly slow responsiveness and performance for applications. In a multi-tier application, you can expect each tier of application to have different I/O requirements. As a general recommendation, pay close attention to the amount of virtual machine disk files hosted on a single NFS volume. Over-subscription of the I/O resources can go unnoticed at first and slowly begin to degrade performance if not monitored proactively.

VMware Memory Virtualization for VSPEX

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

Memory Compression

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

For more details about Vsphere memory management concepts, see the VMware Vsphere Resource Management Guide at: http://www.VMware.com/files/pdf/mem_mgmt_perf_Vsphere5.pdf

Virtual Networking Best Practices

Following are the vSphere networking best practices:

- Separate virtual machine and infrastructure traffic—Keep virtual machine and VMkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation.
- 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 ESX/ESXi hosts to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESX/ESXi hosts from affecting the STP topology. Loops are not possible in virtual switches.
- Jumbo MTU for vMotion and Storage traffic—This best practice is implemented in the architecture by configuring jumbo MTU end-to-end.

VMware Storage Layout for VSPEX

This section explains the EMC storage layout used for this solution.

Storage Layout

The architecture diagram in this section shows the physical disk layout. Disk provisioning on the VNXe series is simplified through the use of wizards, so that administrators do not choose which disks belong to a given storage pool. The wizard may choose any available disk of the proper type, regardless of where the disk physically resides in the array.

Figure 5 shows storage architecture for 50 virtual machines on VNXe3150:

Figure 5 Storage Architecture for 50 VMs on EMC VNXe3150

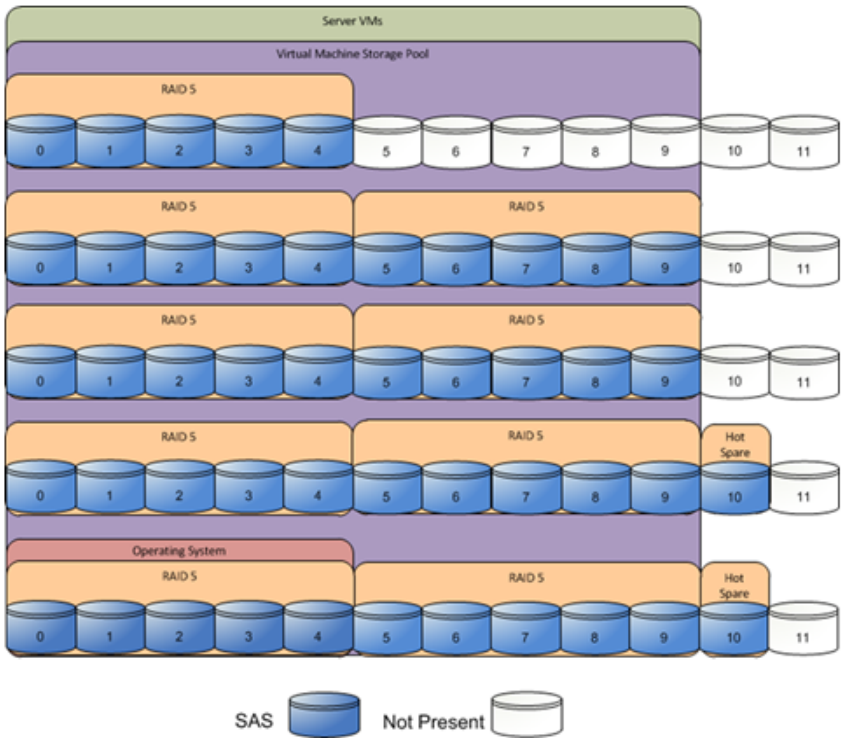

Table 5 provides size of datastores for VMware 50 VMs architecture laid out in Figure 5.

Table 5 Datastore Details for V50 Architecture

Parameters	50 Virtual Machines
Disk capacity and type	300GB SAS
Number of disks	45
RAID type	4 + 1 RAID 5 groups
Number of pools	1
Hot spare disks	2

The reference architecture uses the following configuration:

- Forty-five 300 GB SAS disks are allocated to a single storage pool as nine 4+1 RAID 5 groups (sold as nine packs of five disks)
- At least one hot spare disk is to be allocated for each 30 disks of a given type.
- At least four iSCSI LUNS are allocated to the ESXi cluster from the single storage pool to serve as datastores for the virtual machines.

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

Storage Virtualization

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

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

Architecture for 50 VMware virtual machines

Figure 6 demonstrates 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 Nexus 3048 switches, two Intel mLoM and a quad-port 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.
- On server side, NIC teaming provides simplified load balancing and network high availability.
- Jumbo MTU set to 9000 end-to-end for efficient storage and vMotion traffic.
- EMC VNXe3150 with two storage processors is used as a storage array.

Figure 6 Logical Layout Diagram for VMware 50 VMs

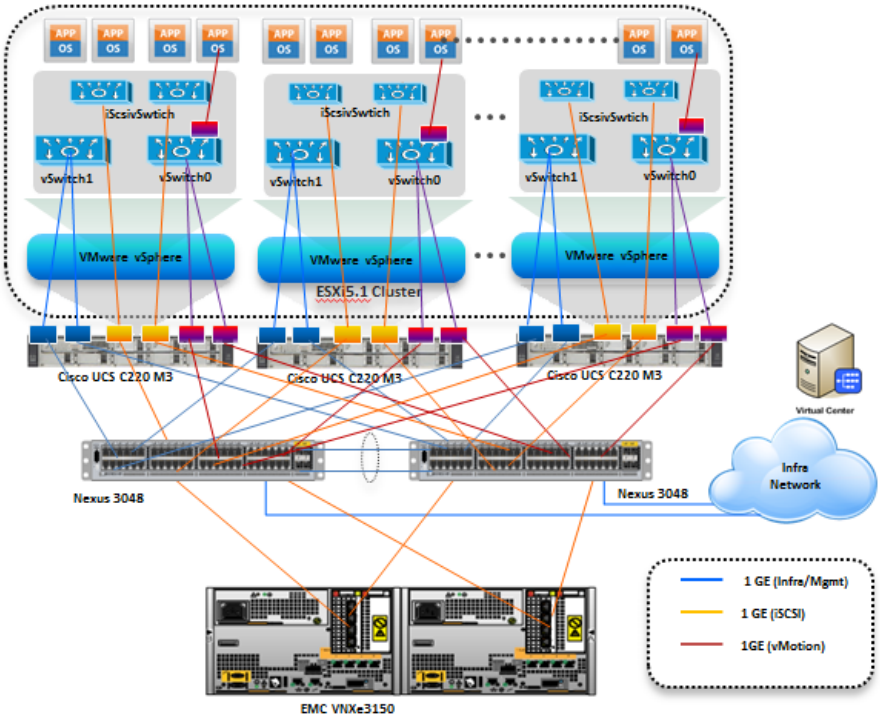

Sizing Guideline

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

Defining the Reference Workload

To simplify the discussion, we have defined a representative reference workload. By comparing your actual 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 characteristics is shown in Table 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 your virtual machines may not exactly match the specifications above. In that case, you can say that a single specific virtual machine is the equivalent of some number of reference virtual machines, and assume that number of virtual machines have been used in the pool. You can continue to provision virtual machines from the pool of resources until it is exhausted. Consider these examples:

Example 1 Custom Built Application

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

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

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

In this example, a single virtual machine uses the resources of two of the reference VMs. If the original pool had the capability to provide 50 VMs worth of resources, the new capability is 48 VMs.

Example 2 Point of Sale System

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

The following are the requirements to virtualize this application:

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

In this case the one virtual machine uses the resources of eight reference virtual machines. If this was implemented on a resource pool for 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 are the requirements to virtualize this application:

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

In this case the virtual machine would use the resources of four reference virtual machines. If this was implemented on a resource pool for 50 virtual machines, there are 46 virtual machines of capability remaining in the pool.

Example 4 Decision Support Database

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

The following are the requirements to virtualize this application:

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

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

Summary of Example

The three examples presented illustrate the flexibility of the resource pool model. In all three cases the workloads simply reduce the number of available resources in the pool. If all three examples were implemented on the same virtual infrastructure, with an initial capacity of 50 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 outside the scope of this document.

However, once the change in resource balance has been examined, and the new level of requirements is known; these virtual machines can be added to the infrastructure using the method described in the examples.

VSPEX Configuration Guidelines

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

1. Pre-deployment tasks
2. Customer configuration data
3. Cabling information
4. Prepare and configure the Cisco Nexus Switches
5. Prepare the Cisco UCS C220 M3 Servers
6. Install ESXi 5.1 on Cisco UCS C220 M3 Servers
7. VMware vCenter server deployment
8. Adding ESXi hosts to vCenter or configuring hosts and vCenter server
9. Configure ESXi networking
10. Prepare the EMC VNXe Series Storage
11. Configure discover address for iSCSI adapters
12. Configuring vSphere HA and DRS
13. Test and validate the installation

Each of these steps are discussed 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 Table 6. These are used throughout the of this document to provide detail on setup procedures and deployment best practices for the various components of the solution.
- Gather tools—Gather the required and optional tools for the deployment. Use Table 2, Table 3 and Table 4 to confirm that all equipment, software, and appropriate licenses are available before the deployment process.
- Gather data—Collect the customer-specific configuration data for networking, naming, and required accounts. Enter this information into the "Customer Configuration Data Sheet" section on page 121 for reference during the deployment process.

Table 6 Deployment prerequisites

Requirement	Description	Reference
Hardware	Cisco UCS C220 M3 servers to host virtual machines	EMC-Cisco Reference Architecture: *VSPEX Server Virtualization with VMware vSphere 5.1 for up to 50 Virtual Machines*.
	Cisco Nexus switches: Two Cisco Nexus 3048 switches for high availability	
	VMware vSphere™ 5.1 server to host virtual infrastructure servers **Note** This requirement may be covered in the existing infrastructure	
	EMC VNXe storage—Multiprotocol storage array with the required disk layout as per architecture requirements	
Software	VMware ESXi™ 5.1 installation media	See the corresponding product documentation
	VMware vCenter Server 5.1 installation media	
	EMC VSI for VMware vSphere: Unified Storage Management – Product Guide	
	EMC VSI for VMware vSphere: Storage Viewer—Product Guide	
	Microsoft Windows Server 2012 installation media (suggested OS for VMware vCenter)	
	Microsoft SQL Server 2008 R2 SP1 **Note** This requirement may be covered in the existing infrastructure	
Licenses	VMware vCenter 5.1 license key	Consult your corresponding vendor to obtain license keys
	VMware ESXi 5.1 license key	
	Microsoft SQL Server license key **Note** This requirement may be covered in the existing infrastructure	
	Microsoft Windows Server 2008 R2 SP1 license key	

Customer Configuration Data

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

"Customer Configuration Data Sheet" section on page 121 provides a set of tables to maintain a record of relevant information. This form can be expanded or contracted as required, and information may be added, modified, and recorded as deployment progresses.

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

Cabling Information

The following information is provided as a reference for cabling the physical equipment in a VSPEX V50 environment. Figure 8 and Figure 9 in this section provide both local and remote device and port locations in order to simplify cabling requirements.

This document assumes that out-of-band management ports are plugged into an existing management infrastructure at the deployment site.

Be sure to follow the cable directions in this section. Failure to do so will result in necessary changes to the deployment procedures that follow because specific port locations are mentioned. Before starting, be sure that the configuration matches what is described in Figure 7, Figure 8, and Figure 9.

Figure 7 shows a VSPEX V50 cabling diagram. The labels indicate connections to end points rather than port numbers on the physical device. For example, connection A is a 1 Gb target port connected from EMC VNXe3150 SP B to Cisco Nexus 3048 A and connection R is a 1 Gb target port connected from Broadcom NIC 3 on Server 2 to Cisco Nexus 3048 B. Connections W and X are 10 Gb vPC peer-links connected from Cisco Nexus 3048 A to Cisco Nexus 3048 B.

Figure 7 VSPEX V50 Cabling Diagram

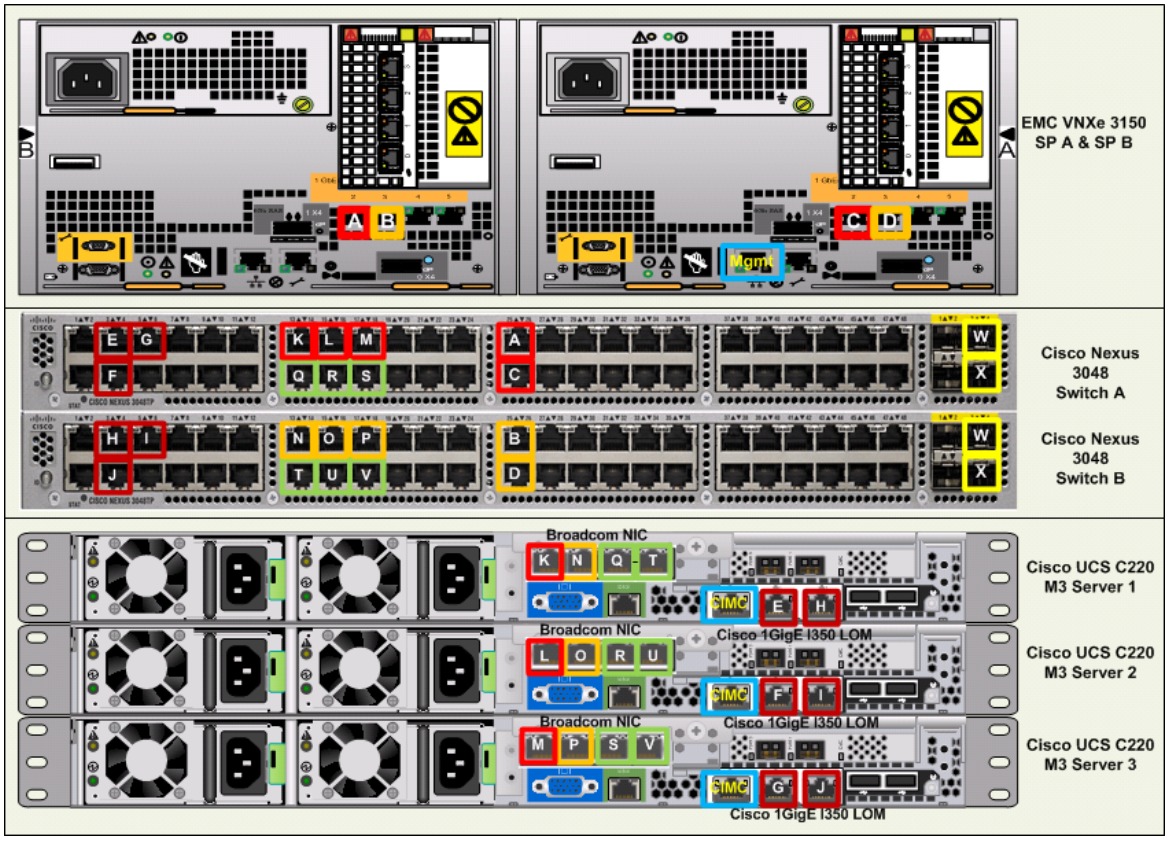

Figure 7, Figure 8, Figure 9 elaborates the detailed cable connectivity for the 50 virtual machines configuration.

Figure 8 Cisco Nexus 3048-A Ethernet Cabling Information

Cable ID on both ends	Ethernet Interface	VLAN ID	Mode	Speed	Port Channel	Remote Device Port
E	Eth1/3	1, 23,613	trunk	1G	3	C220 Srv1- 1GE LOM 1
F	Eth1/4	1,23,613	trunk	1G	4	C220 Srv2- 1GE LOM 1
G	Eth1/5	1,23,613	trunk	1G	5	C220 Srv3- 1GE LOM 1
W	Eth1/51	1,20,22,23,613	trunk	10G	10	vPC peer link
X	Eth1/52	1,20,22,23,613	trunk	10G	10	vPC peer link
K	Eth1/13	20	access	1G	--	C220 Srv1- Broadcom NIC 1
L	Eth1/15	20	access	1G	--	C220 Srv2- Broadcom NIC 1
M	Eth1/17	20	access	1G	--	C220 Srv3- Broadcom NIC 1
Q	Eth1/14	1,22	trunk	1G	14	C220 Srv1- Broadcom NIC 3
R	Eth1/16	1,22	trunk	1G	16	C220 Srv2- Broadcom NIC 3
S	Eth1/18	1,22	trunk	1G	18	C220 Srv3- Broadcom NIC 3
Not shown	Eth1/9	1, 23,613	trunk	10G	--	Uplink to Infra n/w
Not shown	Eth1/10	1,23,613	trunk	10G	--	Uplink to Infra n/w
A	Eth1/25	20	access	1G	--	VNXe3150 (eth10) - SPA
C	Eth1/26	20	access	1G	--	VNXe3150 (eth10) - SPB

Figure 9 Cisco Nexus 3048-B Ethernet Cabling Information

Cable ID on both ends	Ethernet Interface	VLAN ID	Mode	Speed	Port Channel	Remote Device Port
H	Eth1/3	1, 23,613	trunk	1G	3	C220 Srv1- 1GE LOM 2
I	Eth1/4	1,23,613	trunk	1G	4	C220 Srv2- 1GE LOM 2
J	Eth1/5	1,23,613	trunk	1G	5	C220 Srv3- 1GE LOM 2
Y	Eth1/51	1,20,22,23,613	trunk	10G	10	VPC peer link
Z	Eth1/52	1,20,22,23,613	trunk	10G	10	VPC peer link
N	Eth1/13	21	access	1G	--	C220 Srv1- Broadcom NIC 2
O	Eth1/15	21	access	1G	--	C220 Srv2- Broadcom NIC 2
P	Eth1/17	21	access	1G	--	C220 Srv3- Broadcom NIC 2
T	Eth1/14	1,22	trunk	1G	14	C220 Srv1- Broadcom NIC 4
U	Eth1/16	1,22	trunk	1G	16	C220 Srv2- Broadcom NIC 4
V	Eth1/18	1,22	trunk	1G	18	C220 Srv3- Broadcom NIC 4
Not shown	Eth1/9	1, 23,613	trunk	10G	--	Uplink to Infra n/w
Not shown	Eth1/10	1,23,613	trunk	10G	--	Uplink to Infra n/w
B	Eth1/25	21	access	1G	--	VNXe3150 (eth11) – SPA
D	Eth1/26	21	access	1G	--	VNXe3150 (eth11) - SPB

Connect all the cables as outlined in the Figure 7, Figure 8, and Figure 9.

Prepare and Configure the Cisco Nexus 3048 Switch

This section provides a detailed procedure for configuring the Cisco Nexus 3048 switches for use in EMC VSPEX V50 solution.

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

As it is apparent from these figures, there are five major cabling sections in these architectures:

1. Inter switch links
2. vMotion connectivity for servers
3. Infrastructure and Management connectivity for servers
4. Storage connectivity

Figure 10 shows two switches configured for vPC. In vPC, a pair of switches acting as vPC peer endpoints looks like a single entity to port-channel-attached devices, although the two devices that act as logical port-channel endpoint are still two separate devices. This provides hardware redundancy with port-channel benefits. Both switches form a **vPC Domain**, in which one vPC switch is Primary while the other is secondary.

Note The configuration steps detailed in this section provides guidance for configuring the Cisco Nexus 3048 running release 5.0(3)U2(2b).

Figure 10 Network Configuration for EMC VSPEX V50

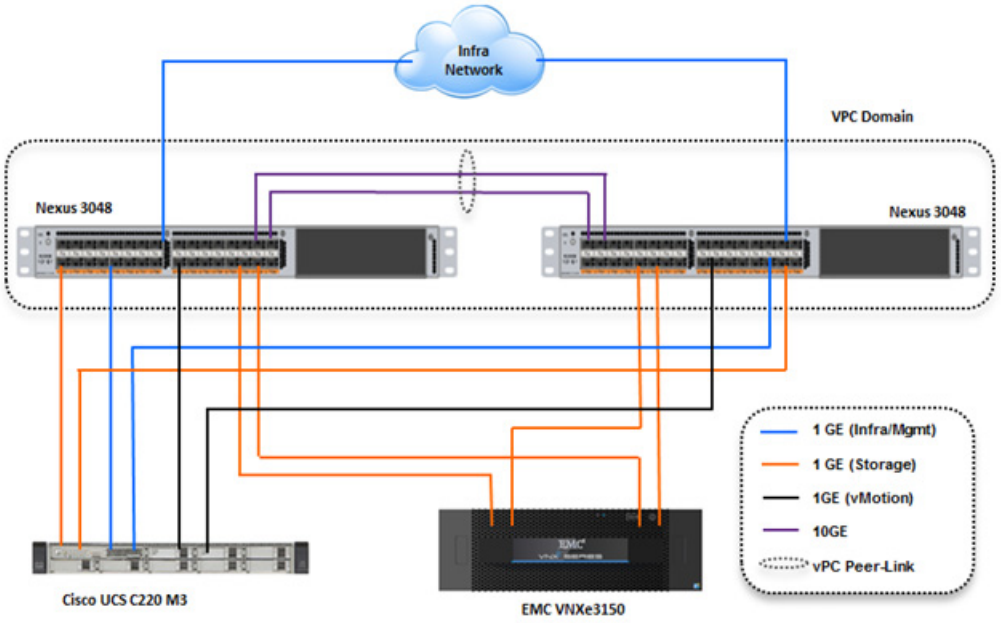

Initial Setup of Nexus Switches

This section details the Cisco Nexus 3048 switch configuration for use in a VSPEX V50 environment.

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

On initial boot and connection to the serial or console port of the switch, the NX-OS setup should automatically start. This initial configuration addresses basic settings such as the switch name, the mgmt0 interface configuration, and SSH setup and defines the control plane policing policy.

Initial Configuration of Cisco Nexus 3048 Switch A and B

Figure 11 Initial Configuration

```
Abort Power On Auto Provisioning and continue with normal setup ?(yes/no)[n]: yes
  ---------- System Admin Account Setup ------------
Do you want to enforce secure password standard (yes/no): yes
  Enter the password for "admin":**********
Confirm the password for "admin":**********
  -------- Basic System Configuration Dialog ---------
 Would you like to enter the basic configuration dialog (yes/no): y
 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 : Switch1
 Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]:
   Mgmt0 IPv4 address : 10.29.150.11
   Mgmt0 IPv4 netmask : 255.255.255.0
 Configure the default gateway for mgmt? (yes/no) [y]:
   IPv4 address of the default gateway : 10.29.150.1
 Enable the telnet service? (yes/no) [n]:
 Enable the ssh service? (yes/no) [y]:
   Type of ssh key you would like to generate (dsa/rsa) : rsa
   Number of  key bits <768-2048> : 1024
 Configure the ntp server? (yes/no) [n]:
 Configure CoPP System Policy Profile ( default / l2 / l3 ) [default]:

The following configuration will be applied:
 switchname Switch1
interface mgmt0
ip address 10.29.150.11 255.255.255.0
no shutdown
exit
vrf context management
ip route 0.0.0.0/0 10.29.150.1
exit
 no telnet server enable
 ssh key rsa 1024 force
 ssh server enable
 policy-map type control-plane copp-system-policy ( default )

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

Software Upgrade (Optional)

It is always recommended to perform any required software upgrades on the switch at this point in the configuration. Download and install the latest available NX-OS software for the Cisco Nexus 3048 switch from the Cisco software download site. There are various methods to transfer both the NX-OS kick-start and system images to the switch. The simplest method is to leverage the USB port on the Switch. Download the NX-OS kick-start and system files to a USB drive and plug the USB drive into the external USB port on the Cisco Nexus 3048 switch.

Copy the files to the local bootflash and update the switch by using the following procedure.

Figure 12 Procedure to update the Switch

```
copy usb1:<<kickstart_image_file>> bootflash:
copy usb1:<<system_image_file>> bootflash:
install all kickstart bootflash:<<kickstart_image_file>> system bootflash:<<system_image_file>>
```

Enable Features

Enable certain advanced features within NX-OS. This is required for configuring some additional options. Enter configuration mode using the (config t) command, and type the following commands to enable the appropriate features on each switch.

Enabling Features in Cisco Nexus 3048 Switch A and B

Figure 13 Command to Enable Features

```
feature interface-vlan
feature lacp
feature vpc
```

Global Port-Channel Configuration

The default port-channel load-balancing hash uses the source and destination IP to determine the load-balancing algorithm across the interfaces in the port channel. Better distribution across the members of the port channels can be achieved by providing more inputs to the hash algorithm beyond the source and destination IP. For this reason, adding the source and destination TCP port to the hash algorithm is highly recommended.

From configuration mode (config t), type the following commands to configure the global port-channel load-balancing configuration on each switch.

Configuring Global Port-Channel Load-Balancing on Cisco Nexus Switch A and B

Figure 14 Commands to Configure Global Port-Channel and Load-Balancing

```
port-channel load-balance ethernet source-dest-port
```

Global Spanning-Tree Configuration

The Cisco Nexus platform leverages a new protection feature called bridge assurance. Bridge assurance helps to protect against a unidirectional link or other software failure and a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of a few states depending on the platform, including **network** and **edge**.

The recommended setting for bridge assurance is to consider all ports as network ports by default. From configuration mode (config t), type the following commands to configure the default spanning-tree options, including the default port type and BPDU guard on each switch.

Configuring Global Spanning-Tree on Cisco Nexus Switch A and B

Figure 15 Configuring Spanning-Tree

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

Enable Jumbo Frames

Cisco solution for EMC VSPEX with VMware architectures require MTU set at 9000 (jumbo frames) for efficient storage and live migration traffic. MTU configuration on Nexus 5000 series switches fall under global QoS configuration. You may need to configure additional QoS parameters as needed by the applications.

From configuration mode (config t), type the following commands to enable jumbo frames on each switch.

Enabling Jumbo Frames on Cisco Nexus 3048 Switch A and B

Figure 16 Enabling Jumbo Frames

```
policy-map type network-qos jumbo
class type network-qos class-default
mtu 9000
system qos
service-policy type network-qos jumbo
```

Configure VLANs

For VSPEX M50 configuration, create the layer 2 VLANs on both the Cisco Nexus 3048 Switches using the Table 7 as reference. Create your own VLAN definition table with the help of "Customer Configuration Data Sheet" section on page 121.

From configuration mode (config t), type the following commands to define and describe the L2 VLANs.

Table 7 Reference VLAN Definitions for EMC VSPEX with VMWare V50 Setup

VLAN Name	VLAN Purpose	ID used in this document	Network Address	Host NICs in VLAN
iSCSI-A	For iSCSI-A traffic	20	10.10.20.0/24	1 Broadcom NIC
iSCSI-B	For iSCSI-B traffic	21	10.10.21.10/24	1 Broadcom NIC
vMotion	For Live Migration	22	10.10.22.0/24	2 Broadcom NICs in a team and on trunk link
VMComm	For VM data	23	10.10.23.0/24	2 Cisco 1GigE I350 LOM in team and on trunk link
Mgmt	For Mgmt	613	10.29.150.0/24	

Defining L2 VLANs on Cisco Nexus 3048 Switch A and B

Figure 17 Commands to Define L2 VLANs

```
vlan 1
vlan <iscsi-a vlan_id>
  name iscsi-a
vlan <iscsi-b vlan_id>
  name iscsi-b
vlan <vmotion vlan_id>
  name vmotion
vlan <vmcomm vlan_id>
  name vmcomm
vlan <mgmt vlan_id>
  name mgmt
```

Virtual Port-Channel (vPC) Global Configuration

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.

The vPC feature requires an initial setup between the two Cisco Nexus switches to function properly. From configuration mode (config t), type the following commands to configure the vPC global configuration for Switch A.

Configuring vPC Global on Cisco Nexus Switch A

Figure 18 Commands to Configure vPC Global Configuration on Switch A

```
vpc domain 101
  role priority 10
  peer-keepalive destination <mgmt0_ip_address of switchB>

int eth1/51-52
  channel-group 10 mode active

int Po10
  description vPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1, <iscsia vlan_id>, <iscsib vlan_id>, <vmotion
vlan_id>,   <vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type network
  vpc peer-link
  no shut
```

From configuration mode (config t), type the following commands to configure the vPC global configuration for Switch B.

Configuring vPC Global on Cisco Nexus Switch B

Figure 19 Commands to Configure vPC Global Configuration on Switch B

```
vpc domain 101
  role priority 10
  peer-keepalive destination <mgmt0_ip_address of switchA>

int eth1/51-52
  channel-group 10 mode active

int Po10
  description vPC peer-link
  switchport mode trunk
  switchport trunk allowed vlan 1, <iscsia vlan_id>, <iscsib vlan_id>, <vmotion
vlan_id>,   <vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type network
  vpc peer-link
  no shut
```

Configuring Storage Connections

Switch interfaces connected to the VNXe storage ports are configured as access ports. Each controller will have two links to each switch.

From the configuration mode (config t), type the following commands on each switch to configure the individual interfaces.

Cisco Nexus 3048 Switch A with VNXe SPA configuration

Figure 20 Commands to Configure VNXe Interface on Switch A

```
interface Ethernet1/25
  description VNXe3150 SPA:eth10
  switchport access vlan 20
  spanning-tree port type edge
  no shut

interface Ethernet1/26
  description VNXe3150 SPB:eth10
  switchport access vlan 20
  spanning-tree port type edge
  no shut
```

Cisco Nexus 3048 Switch B with VNXe SPA configuration

Figure 21 Commands to Configure VNXe Interface on Switch B

```
interface Ethernet1/25
  description VNXe3150 SPA:eth11
  switchport access vlan 21
  spanning-tree port type edge
  no shut

interface Ethernet1/26
  description VNXe3150 SPB:eth11
  switchport access vlan 21
  spanning-tree port type edge
  no shut
```

Configuring Server Connections

Each server has six network adapters (two Intel and four Broadcom ports) connected to both switches for redundancy as shown in Figure 10. This section provides the steps to configure the interfaces on both the switches that are connected to the servers.

Cisco Nexus Switch A with Server 1 configuration

Figure 22 Commands to Configure Interface on Switch A for Server 1 Connectivity

```
interface Ethernet1/3
  description Server1 vmnic0 Intel MLoM1
  switchport mode trunk
  channel-group 3 mode on

interface port-channel3
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 3
  no shut

interface Ethernet1/13
  description Server1 vmnic4 Broadcom 01
  switchport access vlan 20
  spanning-tree port type edge
  no shut

interface Ethernet1/14
  description Server1 vmnic6 Broadcom 03
  switchport mode trunk
  channel-group 14 mode on

interface port-channel14
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 14
  no shut
```

Cisco Nexus Switch B with Server 1 configuration

Figure 23 Commands to Configure Interface on Switch B and Server 1 Connectivity

```
interface Ethernet1/3
  description Server1 vmnic1 Intel MLoM2
  switchport mode trunk
  channel-group 3 mode on

interface port-channel3
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 3
  no shut

interface Ethernet1/13
  description Server1 vmnic5 Broadcom 02
  switchport access vlan 21
  spanning-tree port type edge
  no shut

interface Ethernet1/14
  description Server1 vmnic7 Broadcom 04
  switchport mode trunk
  channel-group 14 mode on

interface port-channel14
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 14
  no shut
```

Cisco Nexus Switch A with Server 2 configuration

Figure 24 Commands to Configure Interface on Switch A and Server 2 Connectivity

```
interface Ethernet1/4
  description Server2 vmnic0 Intel MLoM1
  channel-group 4 mode on

interface port-channel4
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 4
  no shut

interface Ethernet1/15
  description Server2 vmnic4 Broadcom 01
  switchport access vlan 20
  spanning-tree port type edge
  no shut

interface Ethernet1/16
  description Server2 vmnic6 Broadcom 03
  channel-group 16 mode on

interface port-channel16
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 16
  no shut
```

Cisco Nexus Switch B with Server 2 configuration

Figure 25 Commands to Configure Interface on Switch B and Server 2 Connectivity

```
interface Ethernet1/4
  description Server2 vmnic1 Intel MLoM2
  channel-group 4 mode on

interface port-channel4
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 4
  no shut

interface Ethernet1/15
  description Server2 vmnic5 Broadcom 02
  switchport access vlan 21
  spanning-tree port type edge
  no shut

interface Ethernet1/16
  description Server2 vmnic7 Broadcom 04
  channel-group 16 mode on

interface port-channel16
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 16
  no shut
```

Cisco Nexus Switch A with Server 3 configuration

Figure 26 Commands to Configure Interface on Switch A and Server 3 Connectivity

```
interface Ethernet1/5
  description Server3 vmnic0 Intel MLoM1
  channel-group 5 mode on

interface port-channel5
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 5
  no shut

interface Ethernet1/17
  description Server3 vmnic4 Broadcom 01
  switchport access vlan 20
  spanning-tree port type edge
  no shut

interface Ethernet1/18
  description Server3 vmnic2 Broadcom 03
  channel-group 18 mode on

interface port-channel18
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 18
  no shut
```

Cisco Nexus Switch B with Server 3 configuration

Figure 27 Commands to Configure Interface on Switch B and Server 3 Connectivity

```
interface Ethernet1/5
  description Server3 vmnic1 Intel MLoM2
  channel-group 5 mode on

interface port-channel5
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmcomm vlan_id>, <mgmt vlan_id>
  spanning-tree port type edge
  vpc 5
  no shut

interface Ethernet1/17
  description Server3 vmnic5 Broadcom 02
  switchport access vlan 21
  spanning-tree port type edge
  no shut

interface Ethernet1/18
  description Server3 vmnic3 Broadcom 04
  channel-group 18 mode on

interface port-channel18
  switchport mode trunk
  switchport trunk allowed vlan 1,<vmotion vlan_id>
  spanning-tree port type edge
  vpc 18
  no shut
```

Configure ports connected to infrastructure network

Port connected to infrastructure network need to be in trunk mode, and they require at least infrastructure and management 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 the following figures.

Figure 28 Show vlan brief

```
SwitchA#sh vlan brief

VLAN Name                             Status    Ports
---- -------------------------------- --------- -------------------------------
1    default                          active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18, Eth1/1, Eth1/2, Eth1/6
                                                Eth1/7, Eth1/8, Eth1/9, Eth1/10
                                                Eth1/11, Eth1/12, Eth1/19
                                                Eth1/20, Eth1/21, Eth1/22
                                                Eth1/23, Eth1/24, Eth1/27
                                                Eth1/28, Eth1/29, Eth1/30
                                                Eth1/31, Eth1/32, Eth1/33
                                                Eth1/34, Eth1/35, Eth1/36
                                                Eth1/37, Eth1/38, Eth1/39
                                                Eth1/40, Eth1/41, Eth1/42
                                                Eth1/43, Eth1/44, Eth1/45
                                                Eth1/46, Eth1/47, Eth1/48
                                                Eth1/49, Eth1/50
20   iscsi-a                          active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18, Eth1/9, Eth1/13, Eth1/15
                                                Eth1/17, Eth1/25, Eth1/26
21   iscsi-b                          active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18, Eth1/9
22   vmotion                          active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18, Eth1/9
23   vmcomm                           active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18,Eth1/9
613  mgmt                             active    Po3, Po4, Po5, Po10, Po14, Po16
                                                Po18,Eth1/9
```

Figure 29 Show Port-Channel Summary Output

```
SwitchA# sh port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
--------------------------------------------------------------------------------
Group Port-       Type     Protocol     Member Ports
      Channel
--------------------------------------------------------------------------------
3     Po3(SU)     Eth      NONE         Eth1/3(P)
4     Po4(SU)     Eth      NONE         Eth1/4(P)
5     Po5(SU)     Eth      NONE         Eth1/5(P)
10    Po10(SU)    Eth      LACP         Eth1/51(P)   Eth1/52(P)
14    Po14(SU)    Eth      NONE         Eth1/14(P)
16    Po16(SU)    Eth      NONE         Eth1/16(P)
18    Po18(SU)    Eth      NONE         Eth1/18(P)
```

In this example, port-channel 10 is the vPC peer-link port-channel, port-channels 3, 4 and 5 are connected to the Cisco 1GigE I350 LOM on the host and port-channels 14, 16 and 18 are connected to the Broadcom NICs on the host. 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 brief** command. Example output is shown in Figure 30:

Figure 30 Show vpc Brief Output

```
SwitchA# sh vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 101
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary
Number of vPCs configured       : 6
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po10   up     1,20-23,613

vPC status
----------------------------------------------------------------------------
id     Port     Status  Consistency  Reason    Active vlans
------ ---------- ------ ----------- ---------- ----------
3      Po3      up     success      success     1,23,613
4      Po4      up     success      success     1,23,613
5      Po5      up     success      success     1,23,613
14     Po14     up     success      success     1,22
16     Po16     up     success      success     1,22
18     Po18     up     success      success     1,22
```

Make sure that vPC peer status is peer adjacency formed ok and all the port-channels, including the peer-link port-channel, have status up.

Prepare the Cisco UCS C220 M3 Servers

This section provides the detailed procedure for configuring a Cisco Unified Computing System C-Series standalone server for use in VSPEX M50 configurations. Perform all the steps mentioned in this section on all the hosts.

For information on physically mounting the servers, see:

http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/install.html

Configure Cisco Integrated Management Controller (CIMC)

These steps describe the setup of the initial Cisco UCS C-Series standalone server. Follow these steps on all servers:

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 31 CIMC Configuration Utility

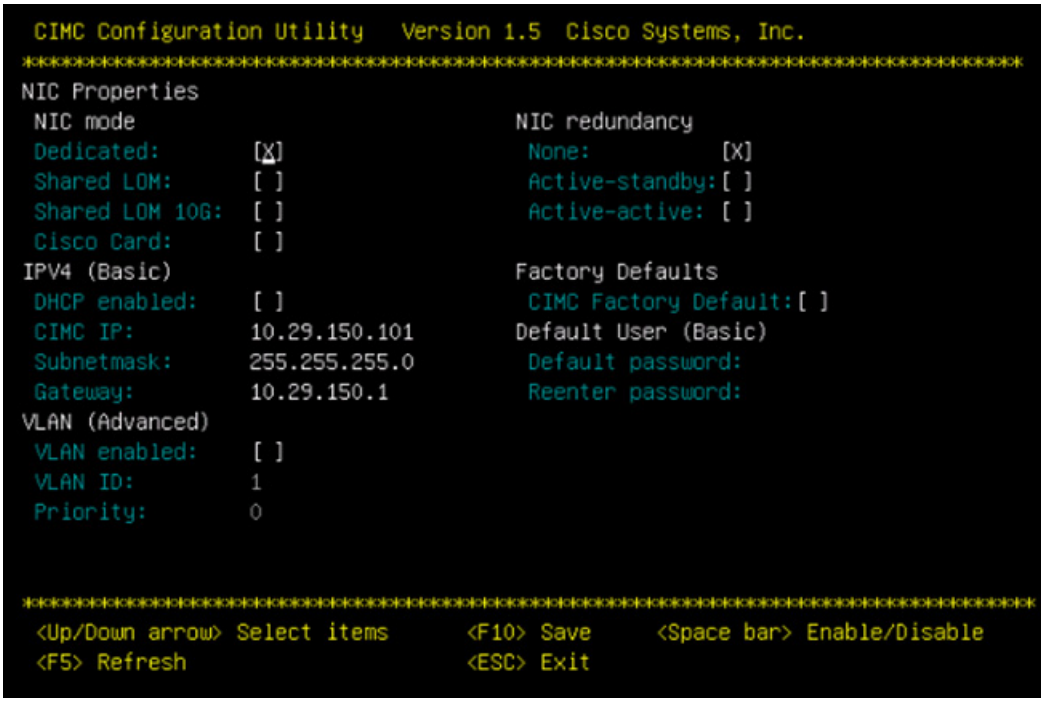

Once the CIMC IP is configured, the server can be managed using the https based Web GUI or CLI.

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 requires an x64-based processor, hardware-assisted virtualization (Intel VT enabled), and hardware data execution protection (Execute Disable enabled). Follow these steps on all the servers to enable Intel ® VT and Execute Disable in BIOS:

1. Press the **Power** button to boot the server. Watch for the prompt to press **F2**.
2. During bootup, press **F2** when prompted to open the BIOS Setup Utility.
3. Choose the **Advanced** tab > **Processor Configuration**.
4. Enable **Execute Disable** and **Intel VT** as shown in Figure 32.

Figure 32 Cisco UCS C220 M3 KVM Console

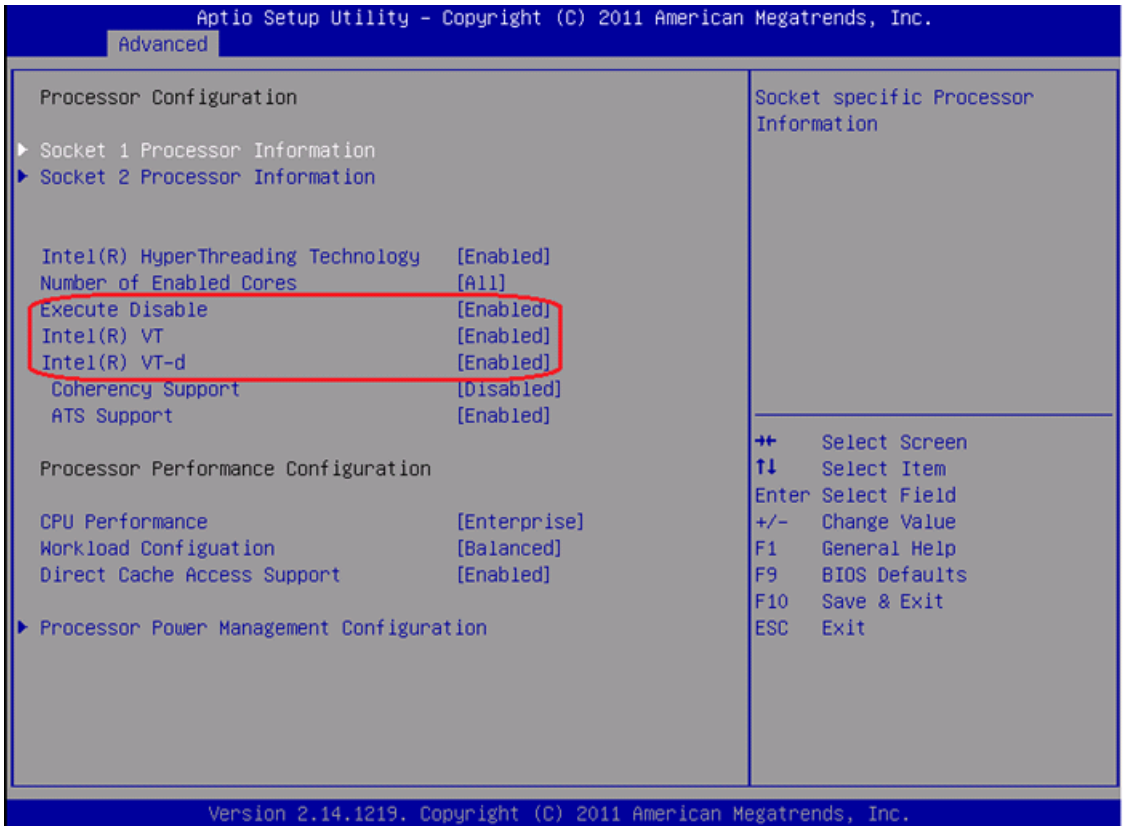

Configuring RAID

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, follow these steps on all the servers:

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 33 Cisco UCS C220 M3 CIMC GUI

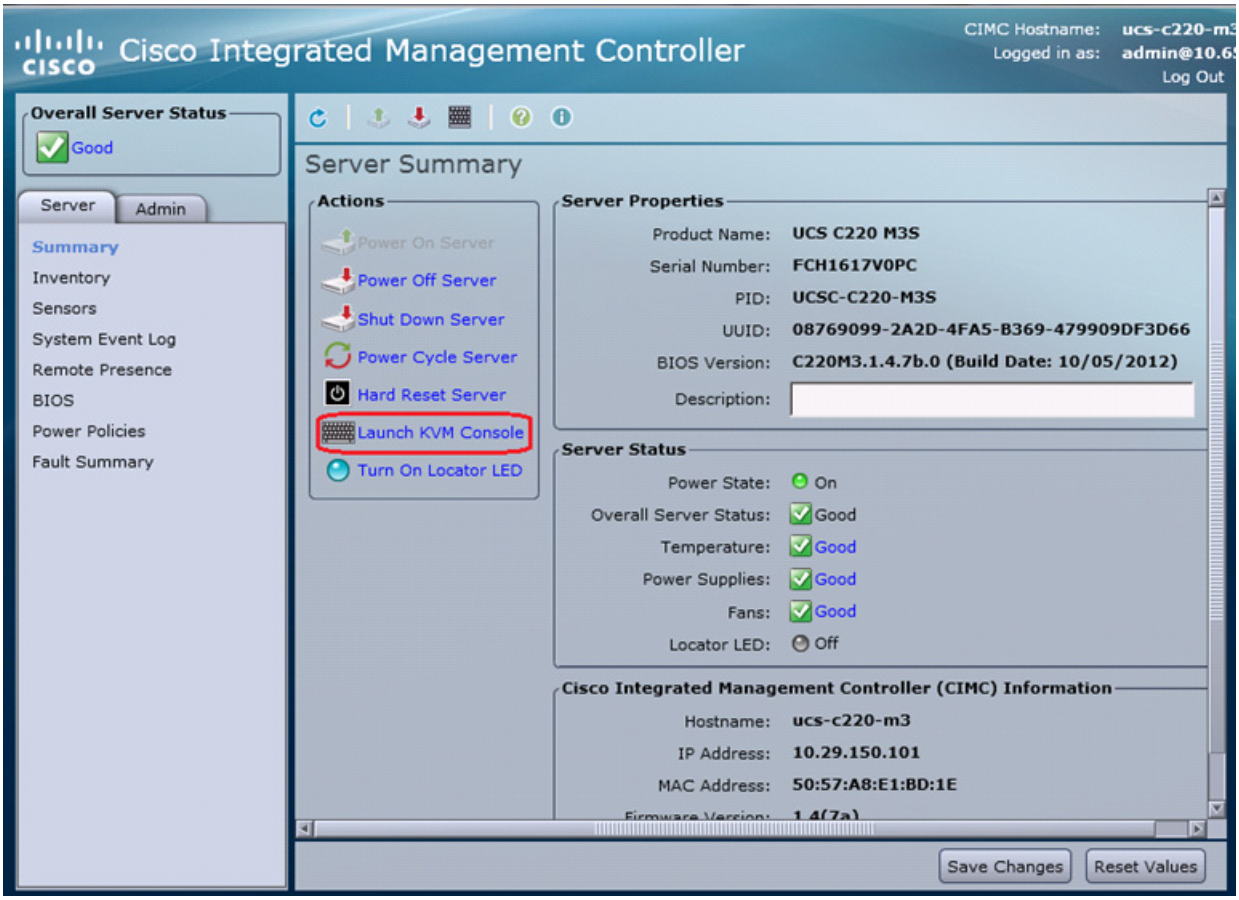

3. During bootup, press **<Ctrl> <H>** when prompted to configure RAID in the WebBIOS.

Figure 34 Cisco UCS C220 M3 KVM Console - Server Booting

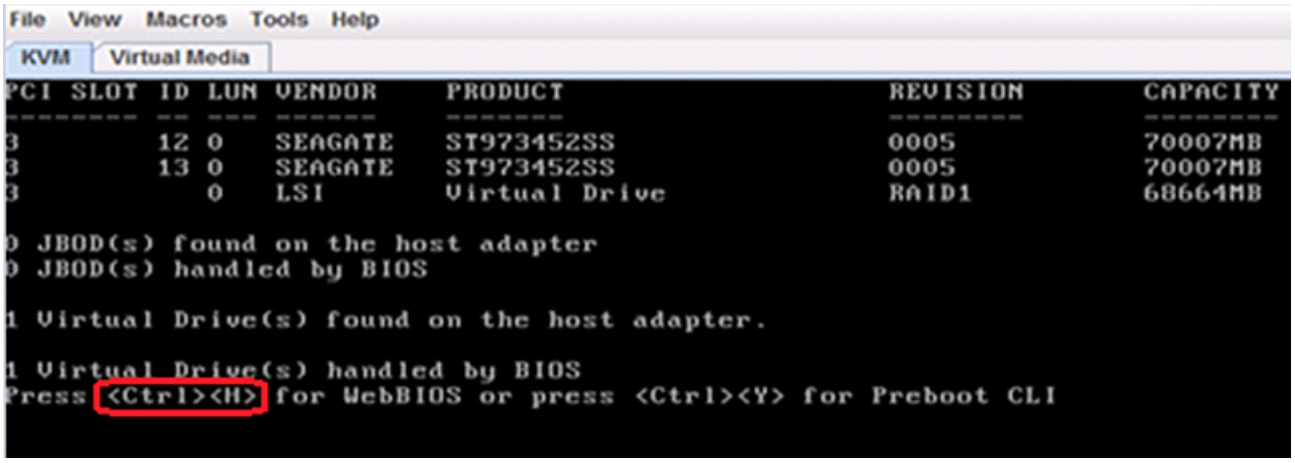

4. Select the adapter and click **Start**.

Figure 35 Adapter Selection for RAID Configuration

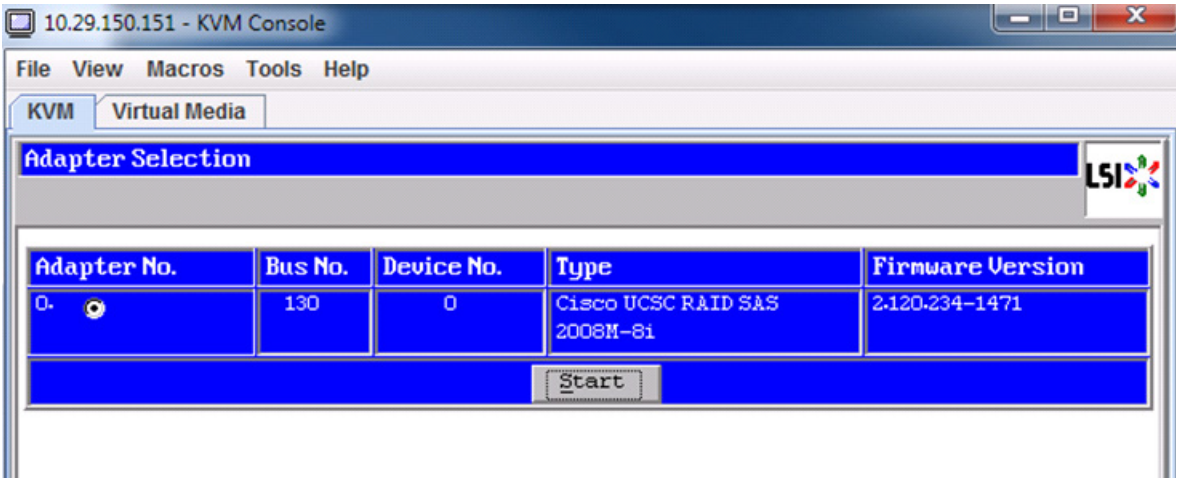

5. Click **New Configuration** and click **Next**.

Figure 36 MegaRAID BIOS Config Utility Configuration Wizard

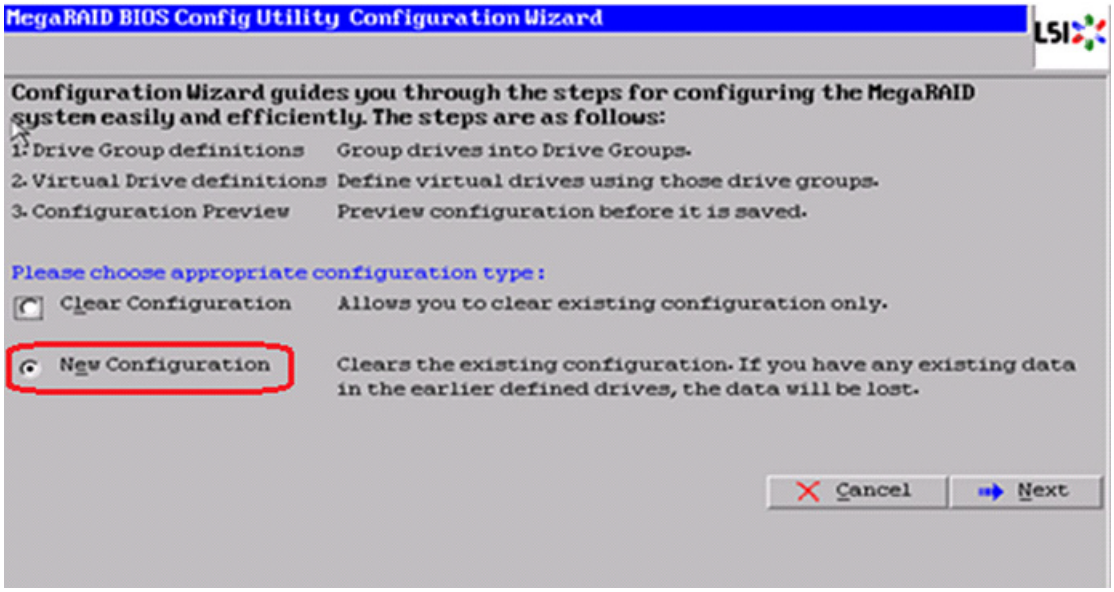

6. Select **Yes** and click **Next** to clear the configuration.

Figure 37 MegaRAID BIOS Config Utility Confirmation Page

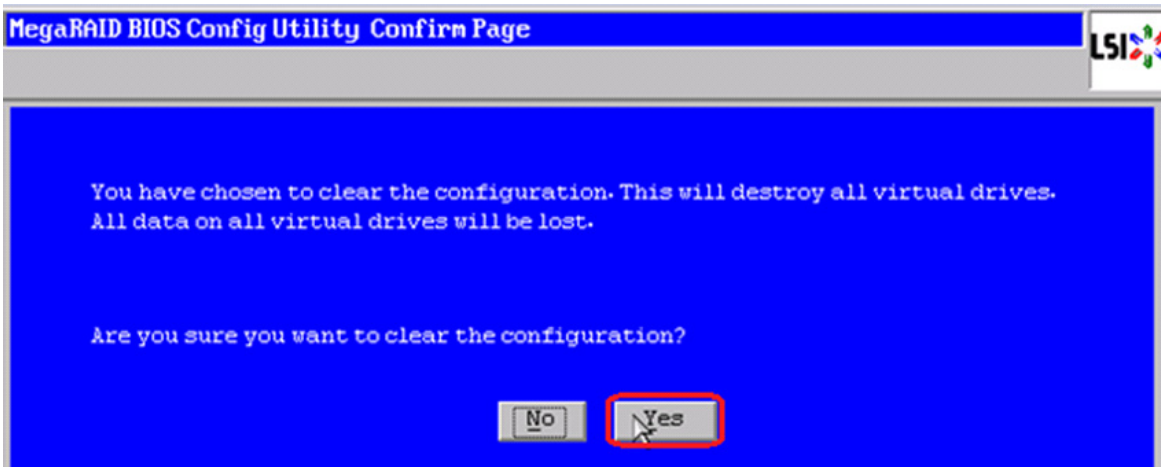

7. If you click **Automatic Configuration** radio button and **Redundancy when possible** for Redundancy and only two drives are available, WebBIOS creates a RAID 1 configuration.

Figure 38 MegaRAID BIOS Config Utility Configuration Wizard - Select Configuration Method

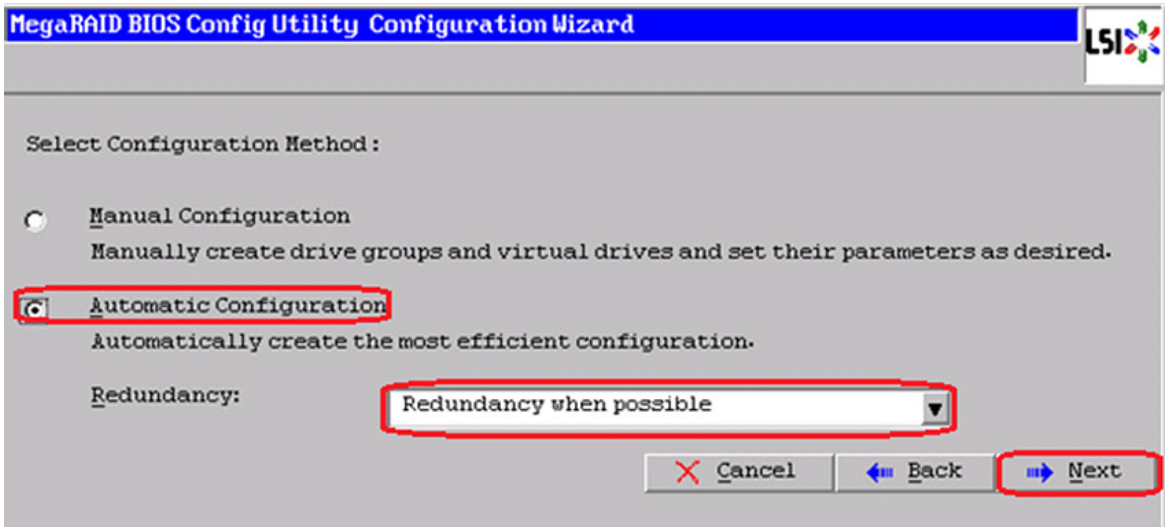

8. Click **Accept** when you are prompted to save the configuration.

Figure 39 MegaRAID BIOS Config Utility Config Wizard - Preview

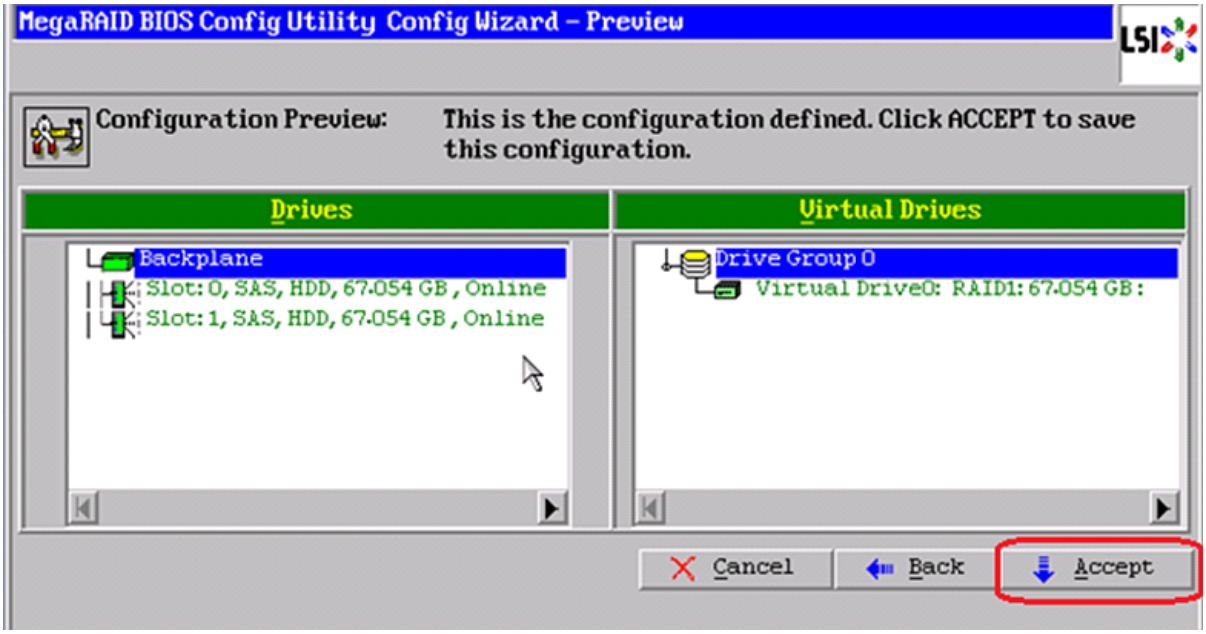

9. Click **Yes** when prompted to initialize the new virtual drives.

Figure 40 MegaRAID BIOS Config Utility Confirmation Page

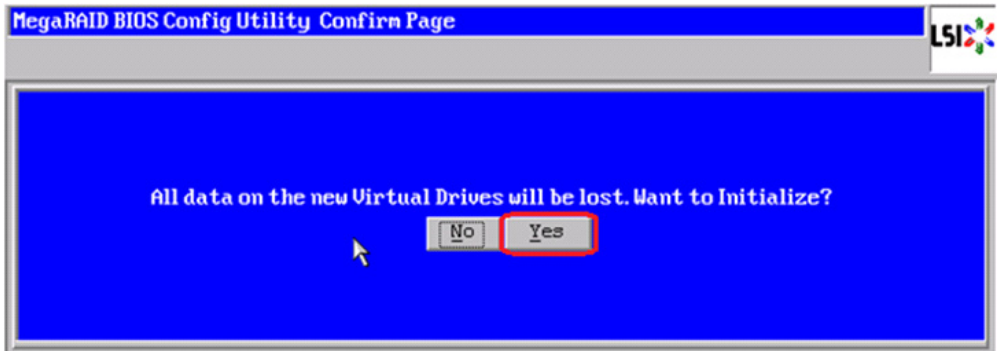

10. Click the **Set Boot Drive** radio button for the virtual drive created above and click **GO**.

Figure 41 MegaRAID BIOS Config Utility Virtual Drives

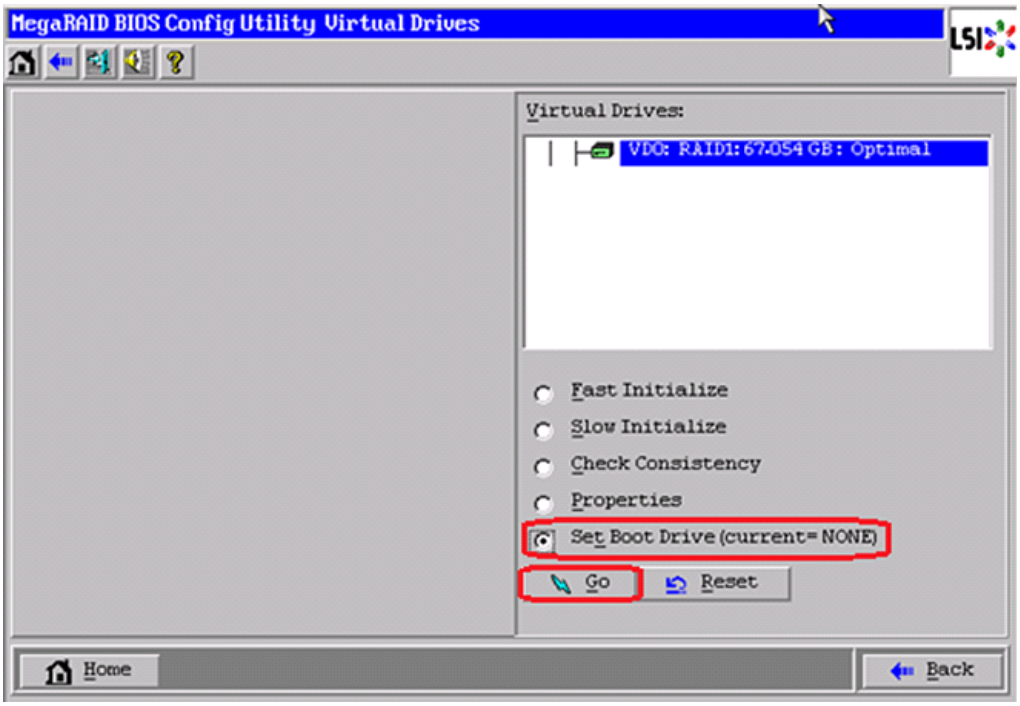

11. Click **Exit** and reboot the system.

Figure 42 MegaRAID BIOS Config Utility Virtual Configuration

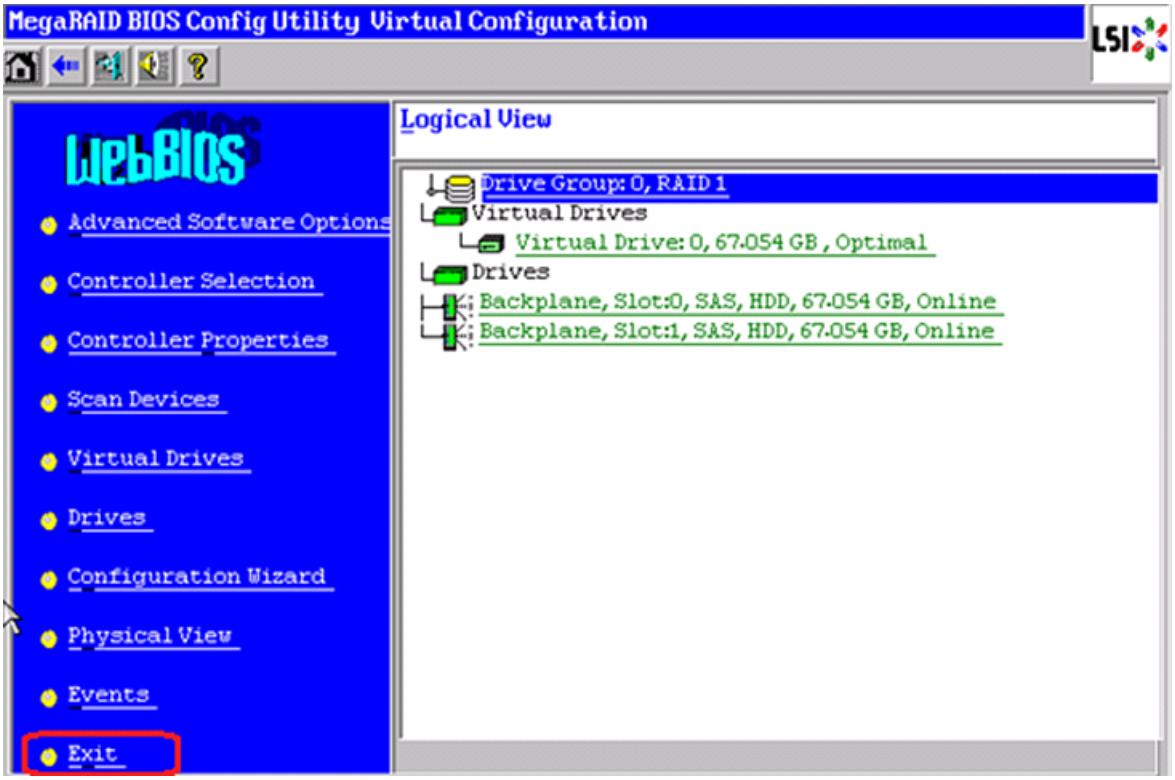

Install ESXi 5.1 on Cisco UCS C220 M3 Servers

This section provides detailed procedures for installing ESXi 5.1 in a V50 VSPEX configuration. Multiple methods exist for installing ESXi in such an environment. This procedure highlights using the virtual KVM console and virtual media features within the Cisco UCS C-Series CIMC interface to map remote installation media to each individual server.

Connect and log into the Cisco UCS C-Series Standalone Server CIMC Interface

1. Open a web browser and enter the IP address for the Cisco UCS C-series CIMC interface. This will launch the CIMC GUI application
2. Log in to CIMC GUI with admin user name and credentials.

Figure 43 CIMC Manager Login Page

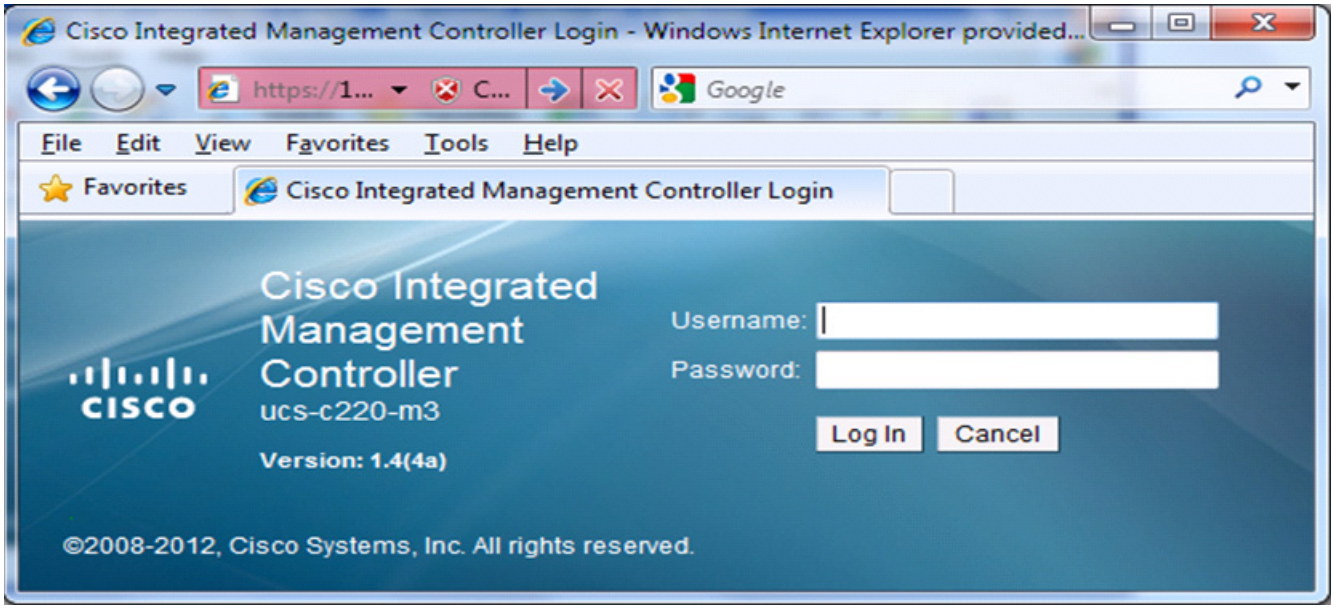

3. In the home page, choose the **Server** tab.
4. Click **launch KVM Console**.
5. Enable the Virtual Media feature, which enables the server to mount virtual drives:
   a. In the CIMC Manager Server tab, click **Remote Presence**.
   b. In the Remote Presence window, click the **Virtual Media** tab and check the check box to enable Virtual Media.
   c. Click **Save Changes**.

Figure 44 CIMC Manager Remote Presence - Virtual Media

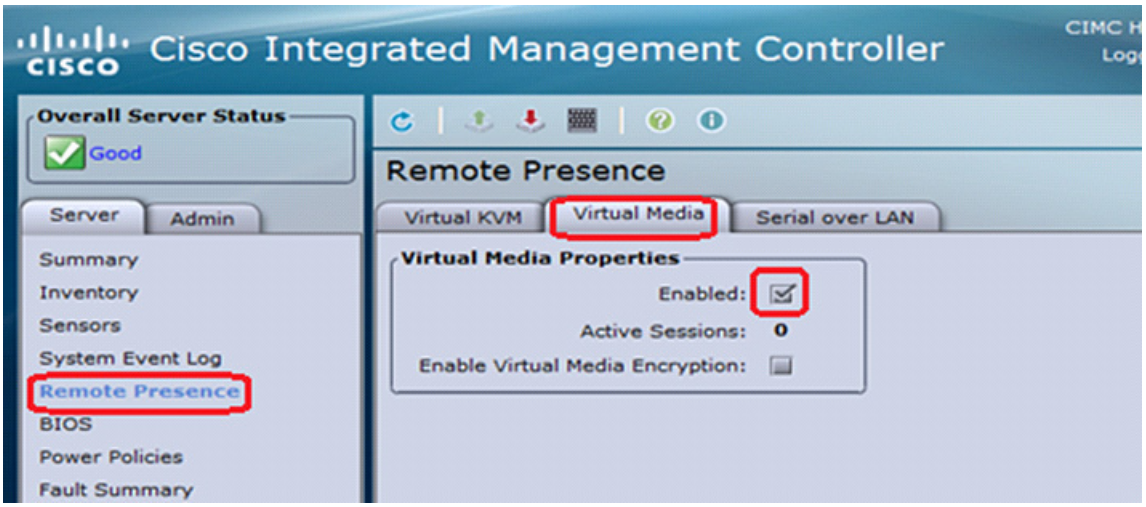

6. On the Remote Presence window, click the **Virtual KVM** tab and then click **Launch KVM Console**.

Figure 45 CIMC Manager Remote Presence - Virtual KVM

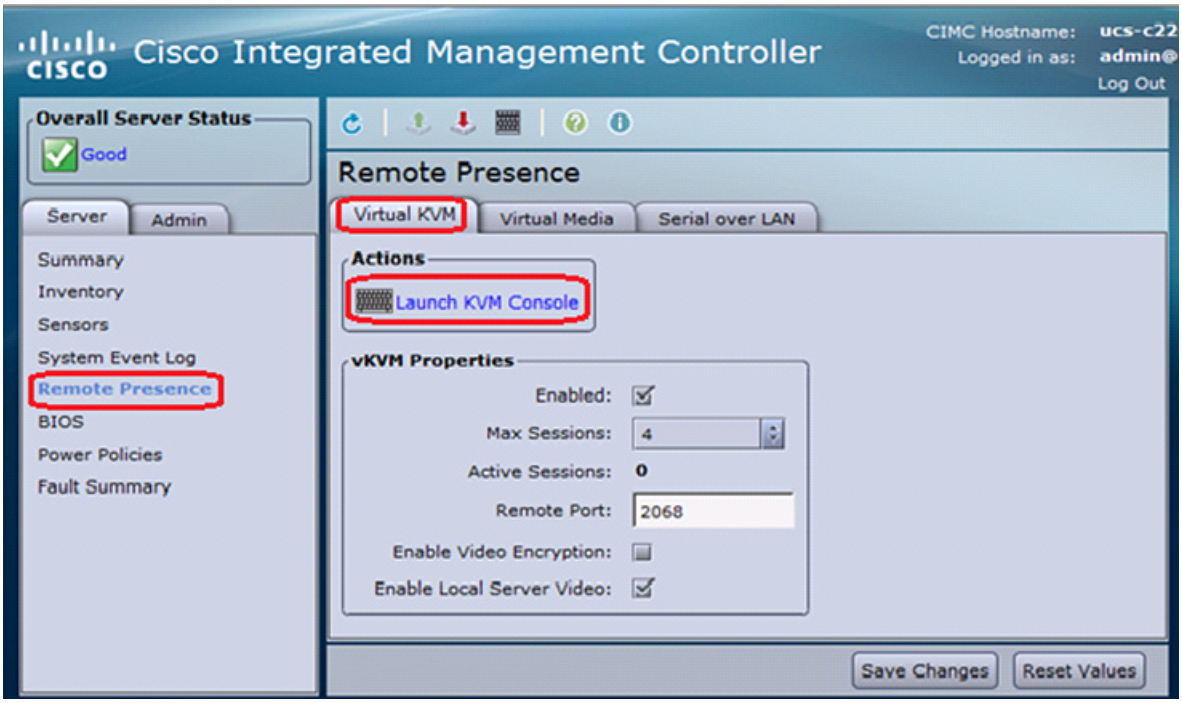

7. When the Virtual KVM Console window launches, click the **Virtual Media** tab.
8. In the Virtual Media Session window, provide the path to the Windows installation image by clicking **Add Image** and then use the dialog to navigate to your VMware ESXi ISO file and select it. The ISO image is displayed in the Client View pane.

Figure 46 CIMC Manager Virtual Media - Add Image

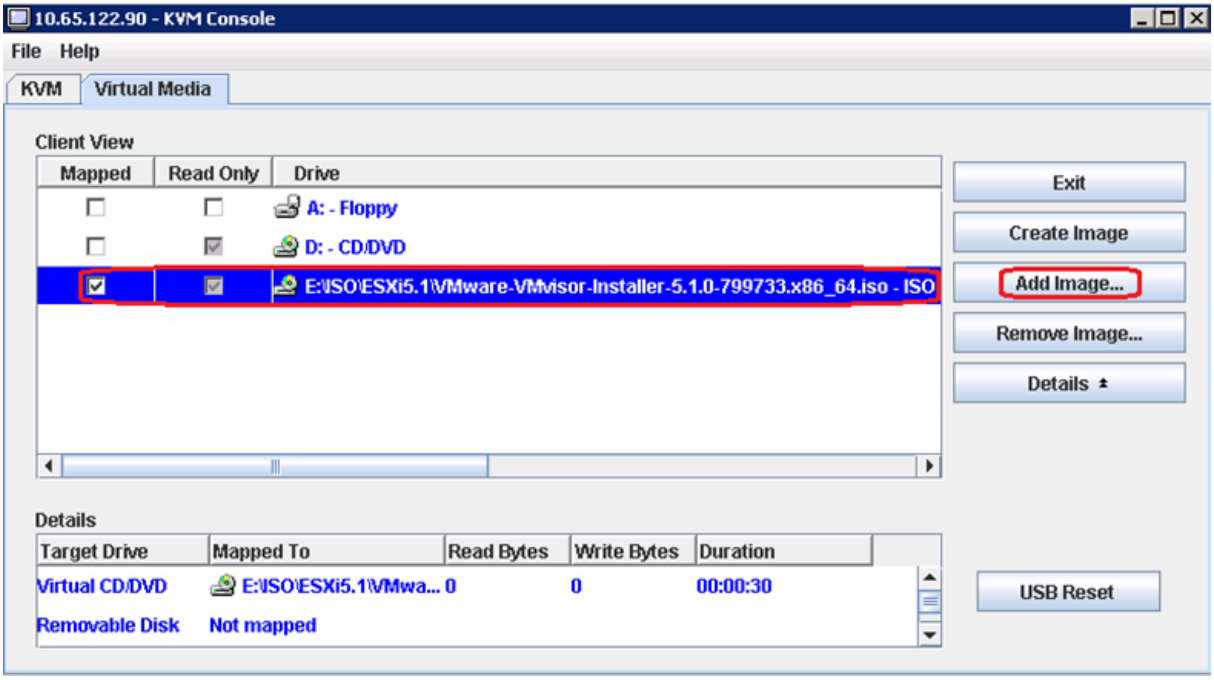

9. When mapping is complete, power cycle the server so that the BIOS recognizes the media that you just added.
10. In the Virtual KVM Console window, watch during bootup for the F2 prompt, and then press **F2** to enter BIOS setup. Wait for the setup utility screen to appear.
11. On the BIOS Setup utility screen, choose the **Boot Options** tab and verify that you see the virtual DVD device that you just added in the above step 8 listed as a bootable device and move it up to the top under Boot Option Priorities as shown in Figure 47.

Figure 47 Cisco UCS C220 M3 BIOS Setup Utility

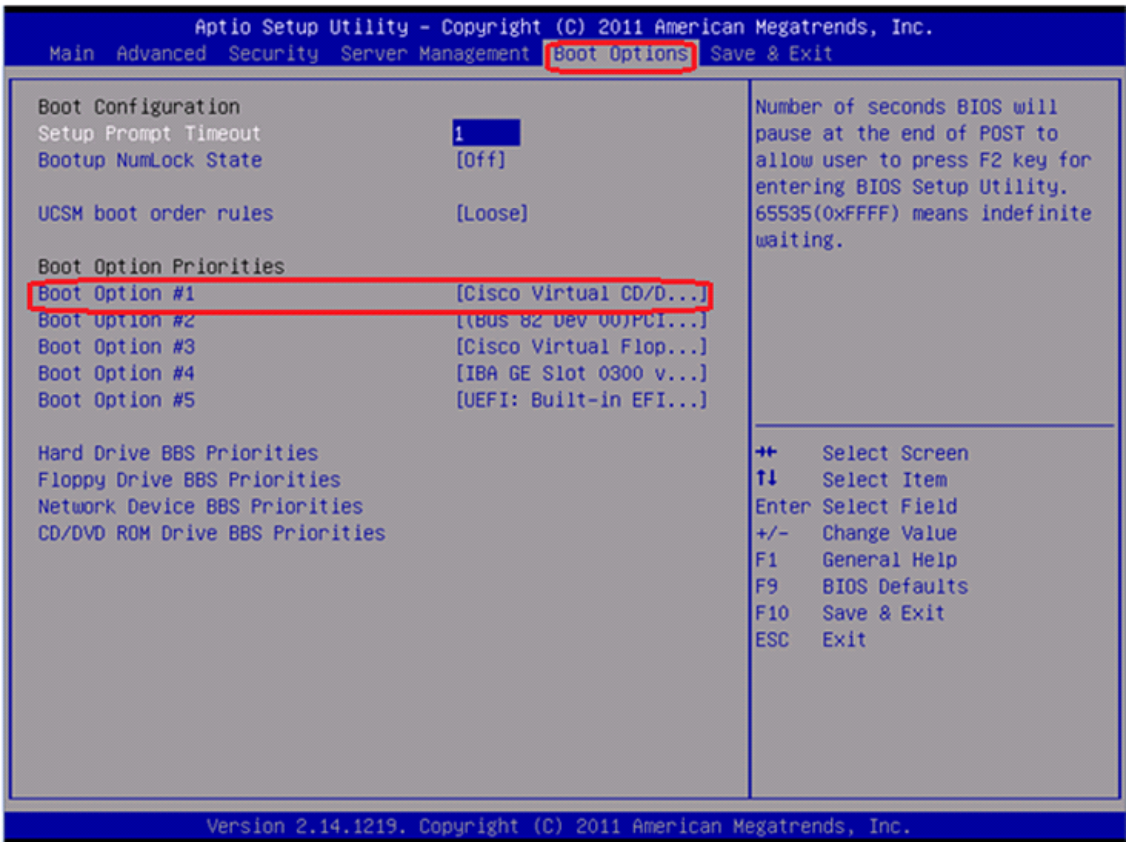

12. Exit the BIOS Setup utility and restart the server.

Installing VMware ESXi

The section explains the steps to complete the installation of ESXi 5.1 on Cisco C220 M3 server and basic configuration of management network for remote access. Follow these steps on all the three Cisco C220 M3 servers:

1. Boot the host from the Virtual CD drive and then select the standard installer option

Figure 48 ESXi installation boot menu

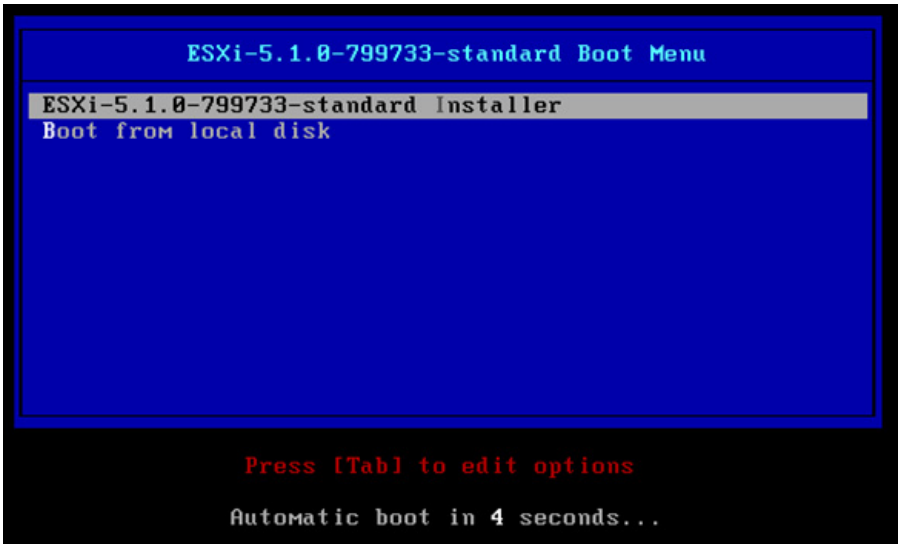

2. Once the loading of ESXi is complete, press **Enter** to continue.

Figure 49 ESXi installation boot welcome screen

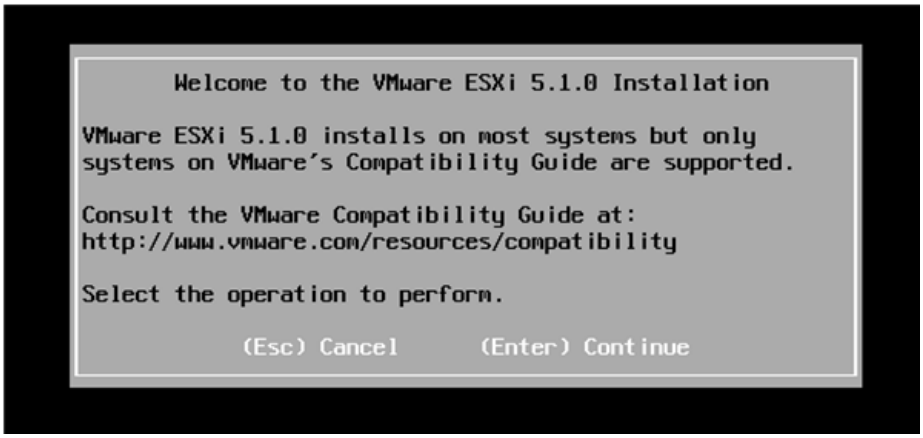

3. In the End User License Agreement (EULA) page, press **F11** to accept and continue.
4. In the Select a Disk to Install or Upgrade window, choose the Local RAID drive that was created in the previous section and press **Enter**.
5. Select the language and press **Enter**.
6. Enter a Password for the ESXi management and press **Enter**.

Figure 50 ESXi installation – Set Password

7. Press **F11** to confirm the installation and wait for the installation to complete.

Figure 51 ESXi installation – Confirm Install

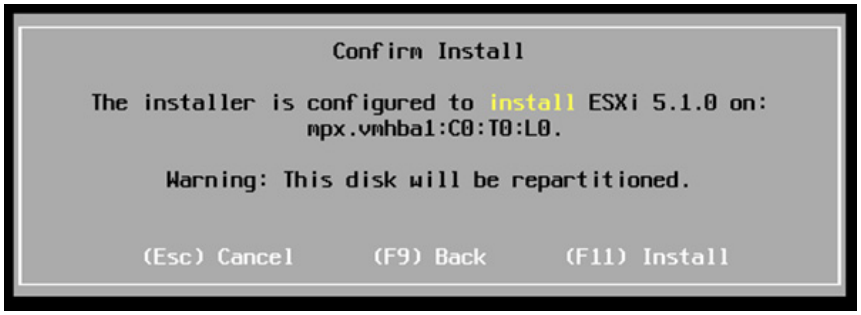

8. After the installation is complete, press **Enter** to reboot.

Figure 52 ESXi installation Complete - Reboot

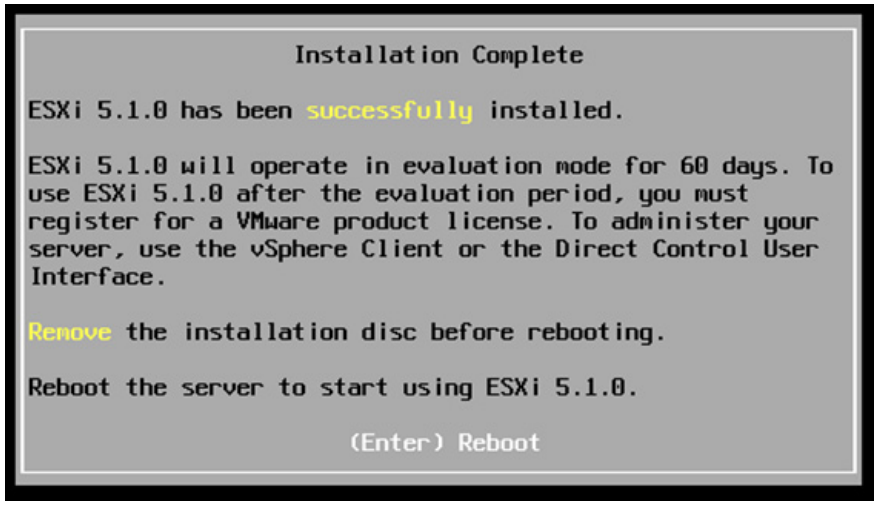

9. Once the server boots up, press **F2** to customize your system.

10. When prompted for a username/password, enter root as the username and the password that was set during the installation.
11. In the system Customization window, select Configure Management Network and press **Enter**.

Figure 53 ESXi System Customization

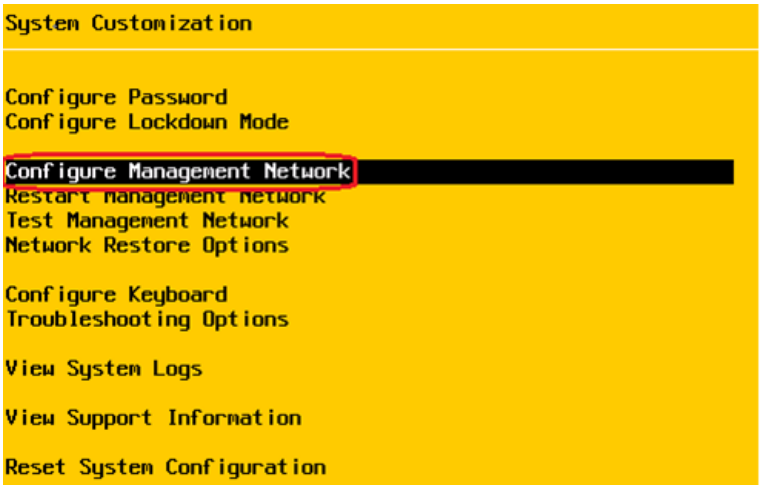

12. In the Configure Management Network, select Network Adapters and press **Enter** to change.
13. Select the vmic that is configured on the switch end for management traffic and press **Enter**.

Figure 54 ESXi Configure Management Network – Network Adapters

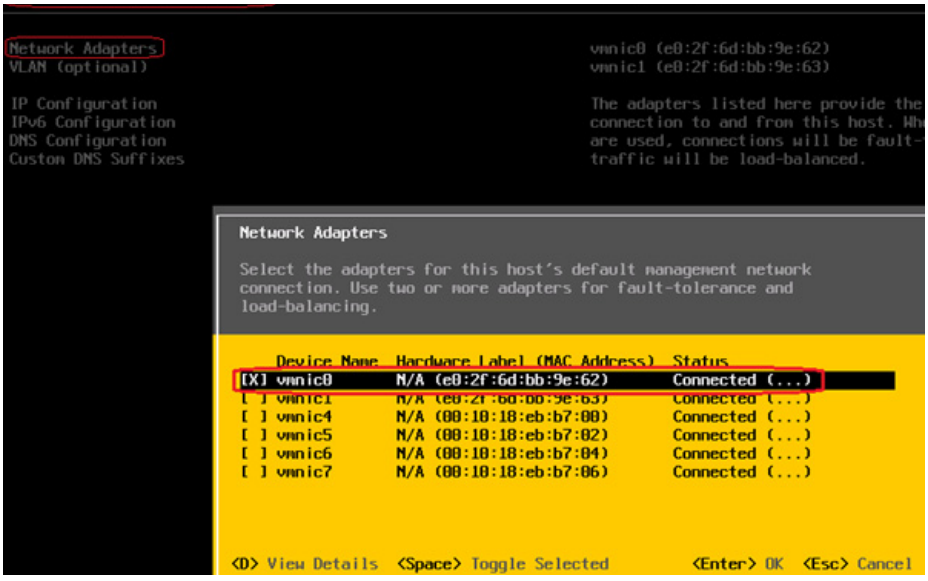

14. Go to IP Configuration under Configure Management Network and press **Enter**.
15. Select Set static IP address and network configuration and enter in your IP Address, Subnet Mask and Default Gateway and press **Enter**.

Figure 55 ESXi Configure Management Network – IP Configuration

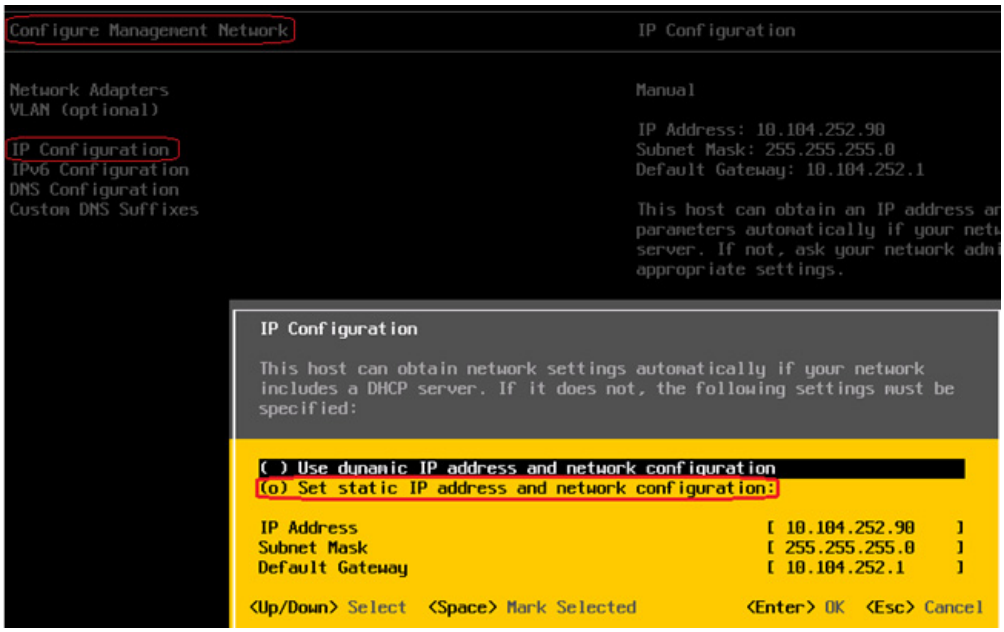

16. This step is optional. If the management vmic is not in the default/native vlan then select VLAN (Optional) in the Configure Management Network and press **Enter** to change the settings.
17. Enter a VLAN ID for the management network.

Figure 56 ESXi Configure Management Network – VLAN

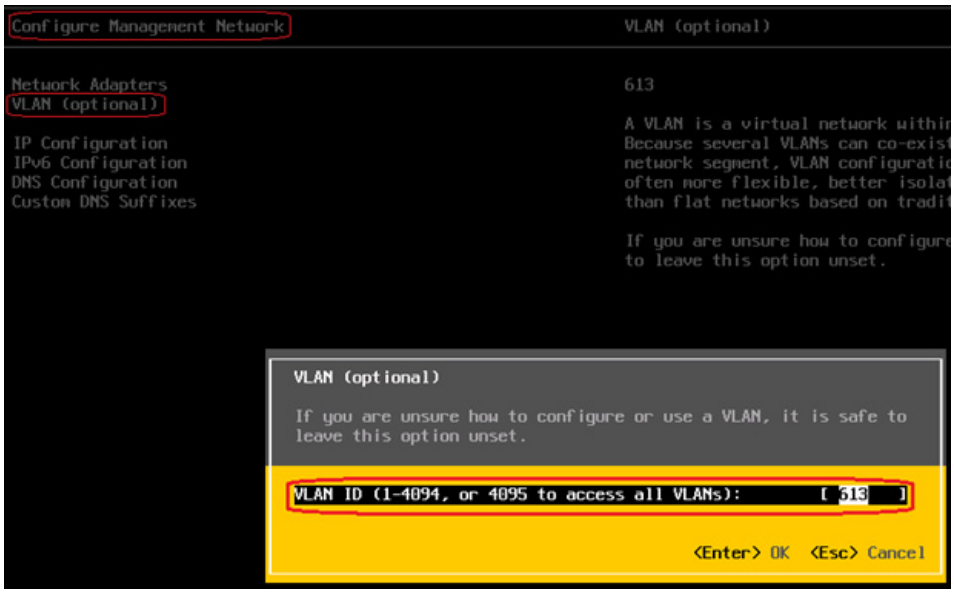

18. Press **ESC** to exit the Configure Management Network window and when prompted press **Y** to apply these changes and restart the management network.

Once the above steps are completed you can further configure and manage the ESXi remotely using the vSphere Client, the vSphere Web client and vCenter Server. To manage the host with the vSphere Client, the vSphere Web Client, and vCenter Server, the applications must be installed on a computer that serves

as a management station with network access to the ESXi host. To download and install vSphere client, open a web browser and enter the IP address of the ESXi management network that you configured in the above step 15. Use the below URL for detailed information on installing the vSphere client:

http://pubs.vmware.com/vsphere-51/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html#com.vmware.vsphere.install.doc/GUID-B15F4221-ABDB-4CA7-A6C4-6C96E72F04A5.html

VMware vCenter Server Deployment

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 more information on installing a vCenter Server, see:

http://pubs.vmware.com/vsphere-51/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html#com.vmware.vsphere.install.doc/GUID-BC044F6C-4733-4413-87E6-A00D3BDEDE58.html

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

Following steps provides high level configuration steps to configure vCenter server:

1. Create the vCenter host VM

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

2. Install vCenter guest OS

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

3. Create vCenter ODBC connection

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

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

4. Install vCenter server

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

5. Apply vSphere license keys

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

Adding ESXi hosts to vCenter or Configuring Hosts and vCenter Server

This section describes on how to populate and organize your inventory by creating a virtual datacenter object in vCenter server and adding the ESXi hosts to it. A virtual datacenter is a container for all the inventory objects required to complete a fully functional environment for operating virtual machines.

1. Open a vSphere client session to a vCenter server by providing IP address and credentials as shown in Figure 57.

Figure 57 vSphere Client

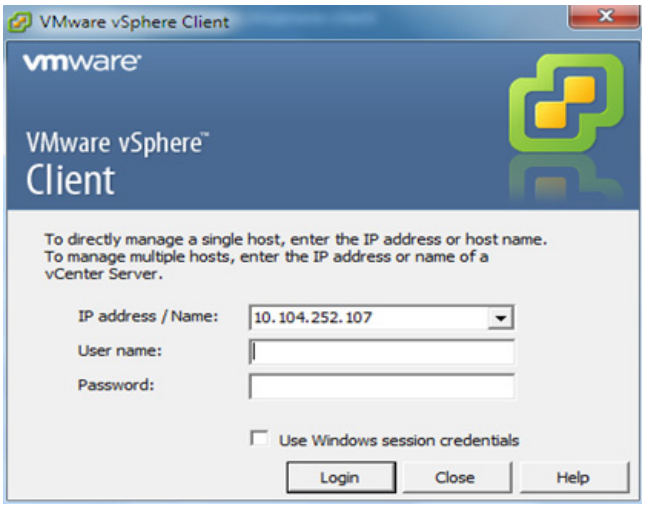

2. Go to **Home > Inventory > Hosts and Clusters.**

Figure 58 VMware vCenter – Host and Clusters

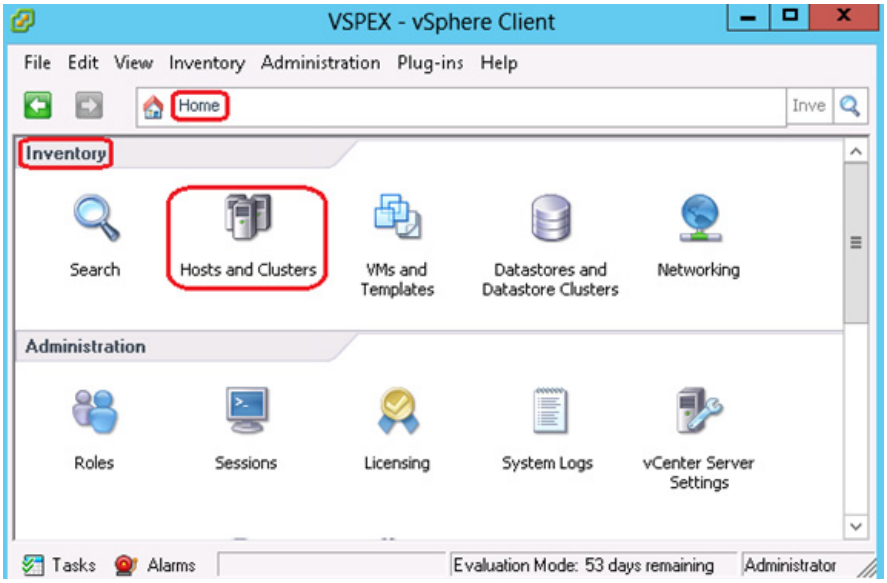

3. Choose **File** > **New** > **Datacenter**. Rename the datacenter.

Figure 59 VMware vCenter – Create Datacenter

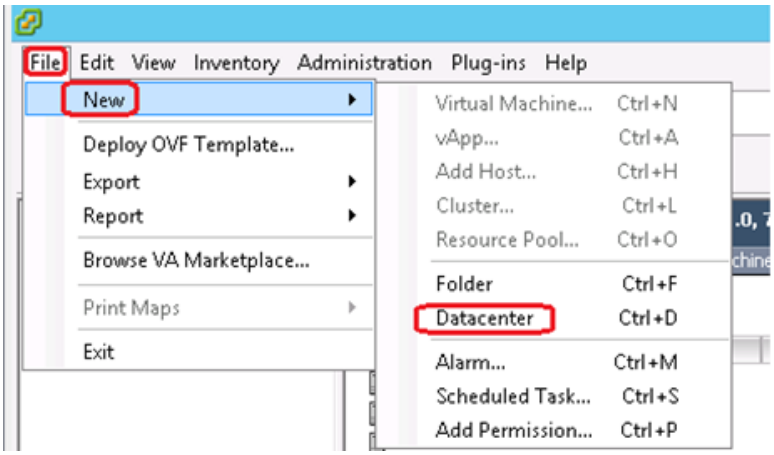

4. In vSphere client, choose **Home** > **Inventory** > **Hosts and Clusters**, select the datacenter you created in step 3 and click **Add a host**.

Figure 60 VMware vCenter – Add a Host

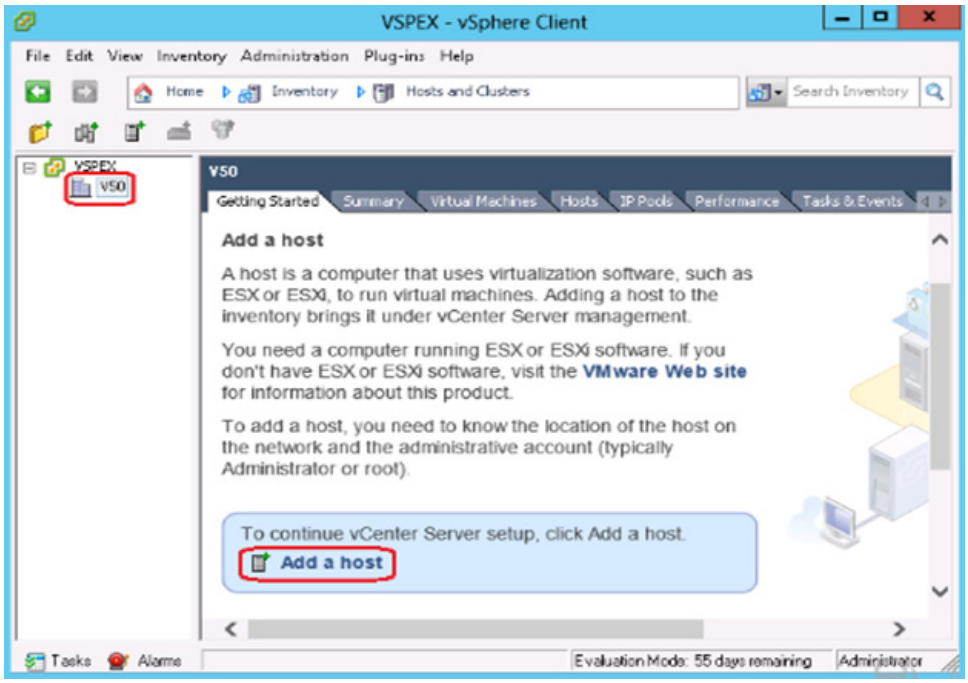

5. In the Add Host Wizard window, enter host name or IP address and the credentials. Click **Next**.

Figure 61 Add Host Wizard –Connection Settings

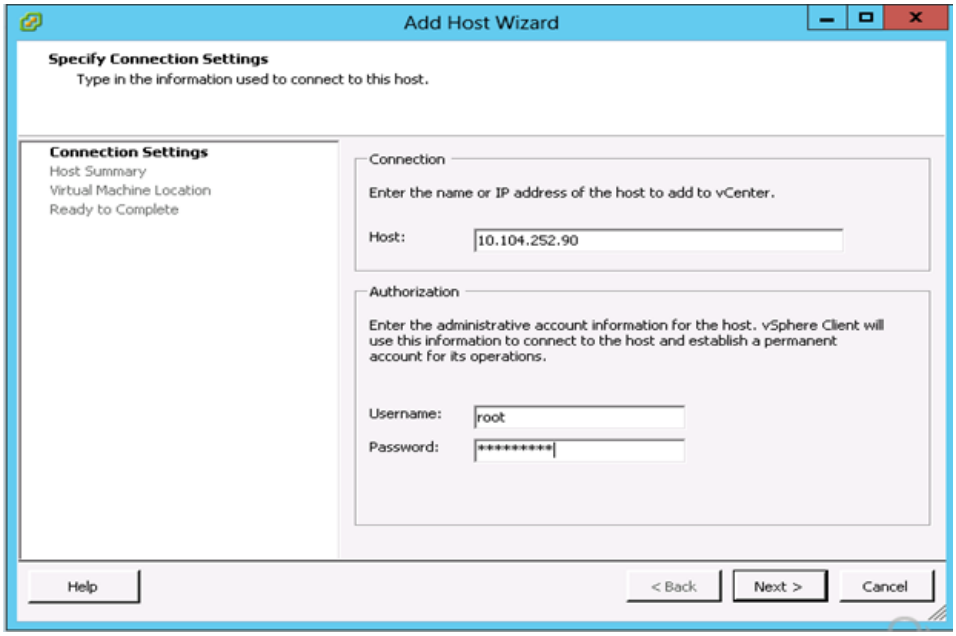

6. A Security Alert window pops up. Click **Yes** to proceed.
7. Click **Next** after reviewing the host information in the Host Summary.

Figure 62 Add Host Wizard – Host Summary

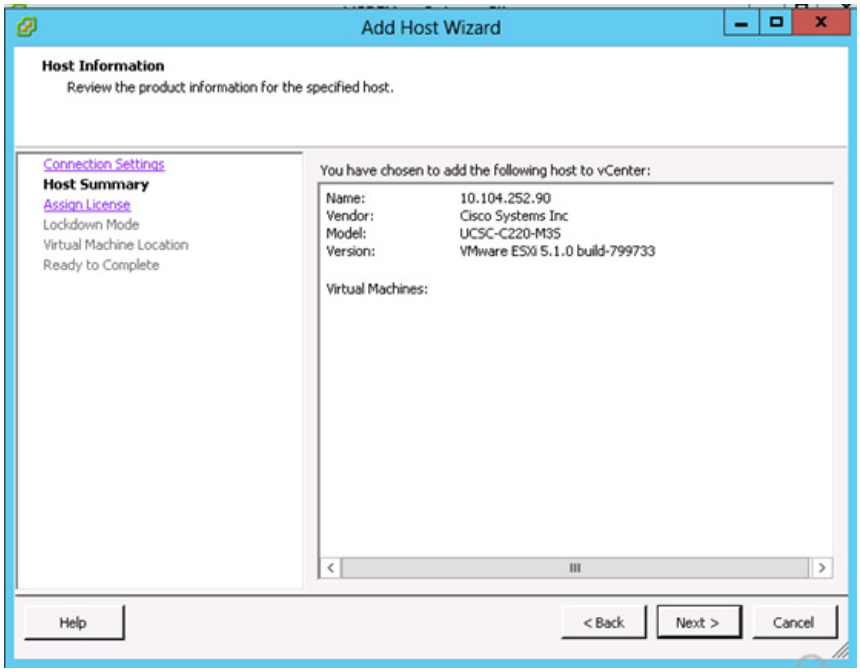

8. Enter the license key for Assign License and click **Next**. You can choose Evaluation Mode to assign the license key later also.

Figure 63 Add Host Wizard –Assign License

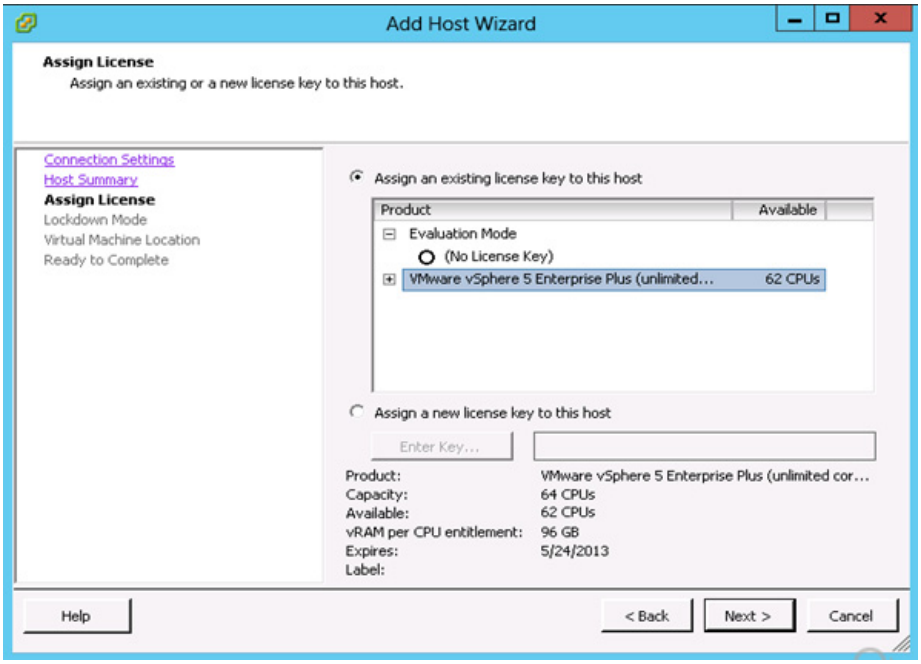

9. In the next screen, per your security policies, you can check the check box Enable Lockdown Mode or keep it unchecked.
10. Click **Finish** in the Ready to Complete Page after reviewing the options you have selected.

Figure 64 Add Host Wizard –Ready to Complete

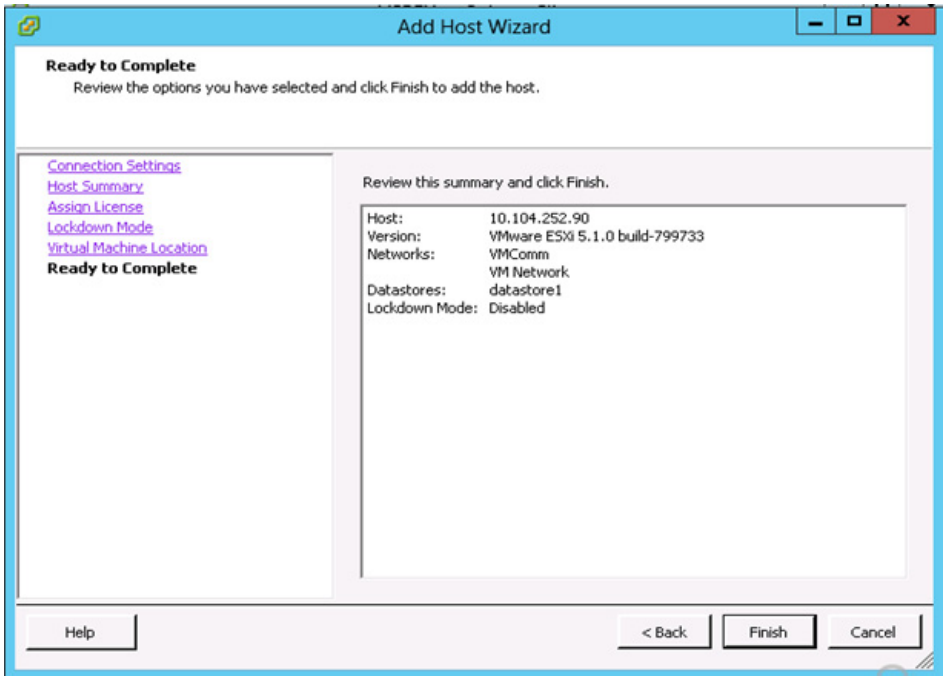

11. Repeat the steps 5 to10 to add the other two ESXi hosts to the vCenter. Once they are added successfully you should see them in the vCenter under the datacenter created as shown in the below figure.

Figure 65 VMware vCenter with hosts added

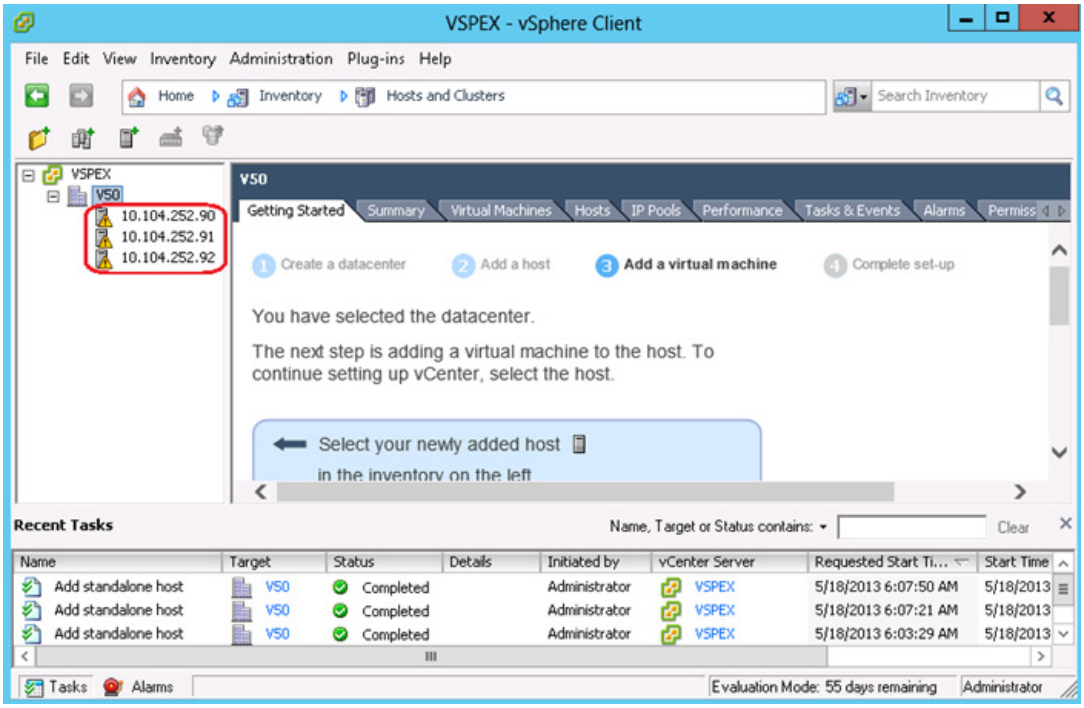

Configure ESXi Networking

This section instructs how to configure the networking for ESXi hosts using the vCenter server. In this document we will be using four virtual standard switches (vSS). vSwitches used for management and vMotion traffic would have two vmnics, one on each fabric for load balancing and high-availability. The other two vSwitches used for iSCSI storage will have each with a single vmnic uplink and a single vmk port, bound to the iSCSI adapter. The VSPEX solution recommends an MTU set at 9000 (jumbo frames) for efficient storage and migration traffic. The below table can be used as a reference for creating and configuring virtual switch on ESXi hosts.

Table 8 ESXi vSwitch reference table for VSPEX V50 solution

vSwitch Name	Traffic Type	NIC Teaming Policy	Load Balancing Algorithm	Port Groups	Uplink Adapters (NICs)
vSwitch0 (Default)	Mgmt and Virtual Machine	Active/Active	IP Hash	1 VMkernel and 1 Virtual Machine	2 Intel mLoM
vSwitch1	vMotion	Active/Active	IP Hash	1 VMkernel	2 Broadcom NICs
vSwitch2	iSCSI	-	-	1 VMkernel	1 Broadcom NIC
vSwitch3	iSCSI	-	-	1 VMkernel	1 Broadcom NIC

Configuring vSwitch for Management and VM traffic

During the installation of VMware ESXi, a virtual standard switch (vSS) will be created. By default, ESXi chooses only one physical NIC as a virtual switch uplink. This section describes the steps to add a redundant link and configure the load balancing for the default vSwitch0. Policies set at the standard switch apply to all of the port groups on the standard switch. The exceptions are the configuration options that are overwritten at the standard port group. The vSwitch for management and VM port group uses the same policy settings configured on the standard vSwitch properties.

1. Connect to the vCenter server using vSphere client.
2. Select an ESXi host on the left pane of Hosts and Clusters window. Click **Configuration > Networking > vSwitch0 Properties** on the right pane of the window.

Figure 66 VMware vCenter - Networking

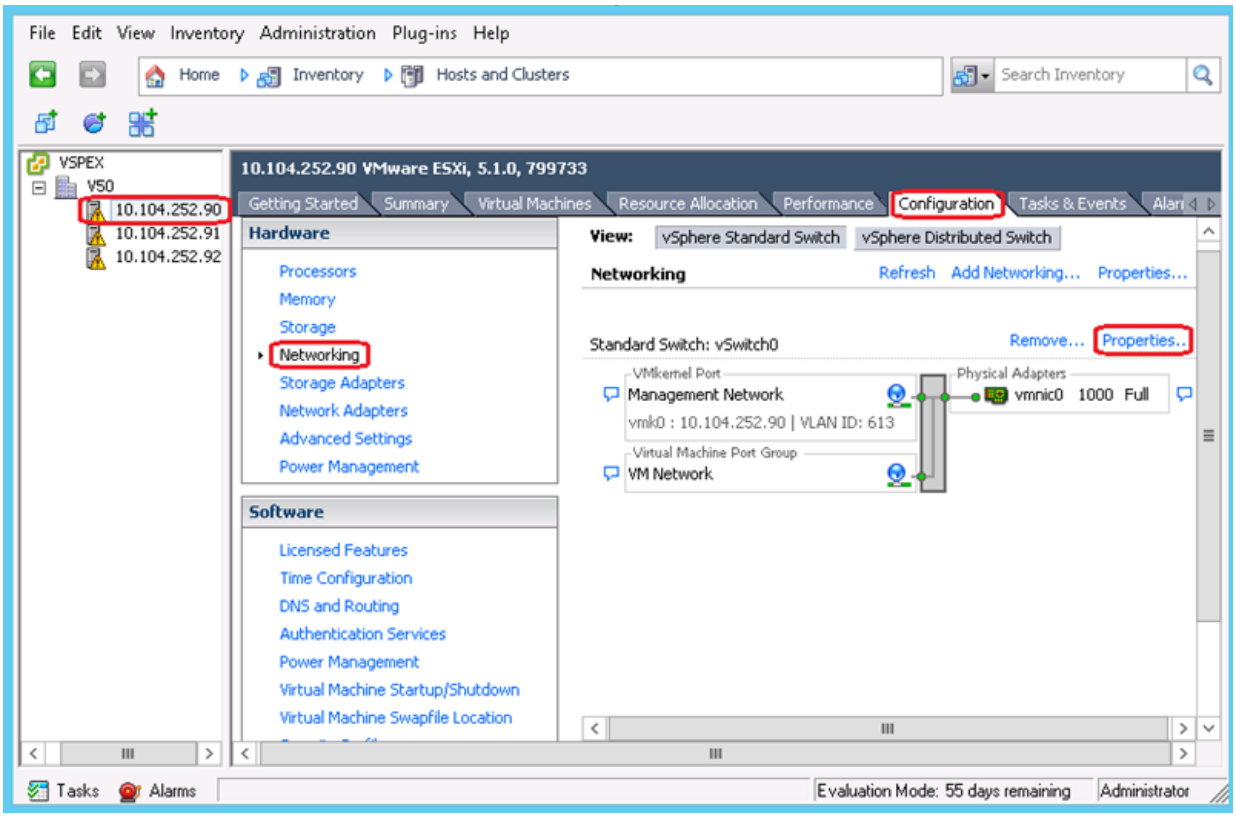

3. Click the **Network Adapters** tab and click **Add** in the vSwitch0 Properties window.

Figure 67 Management vSwitch Properties

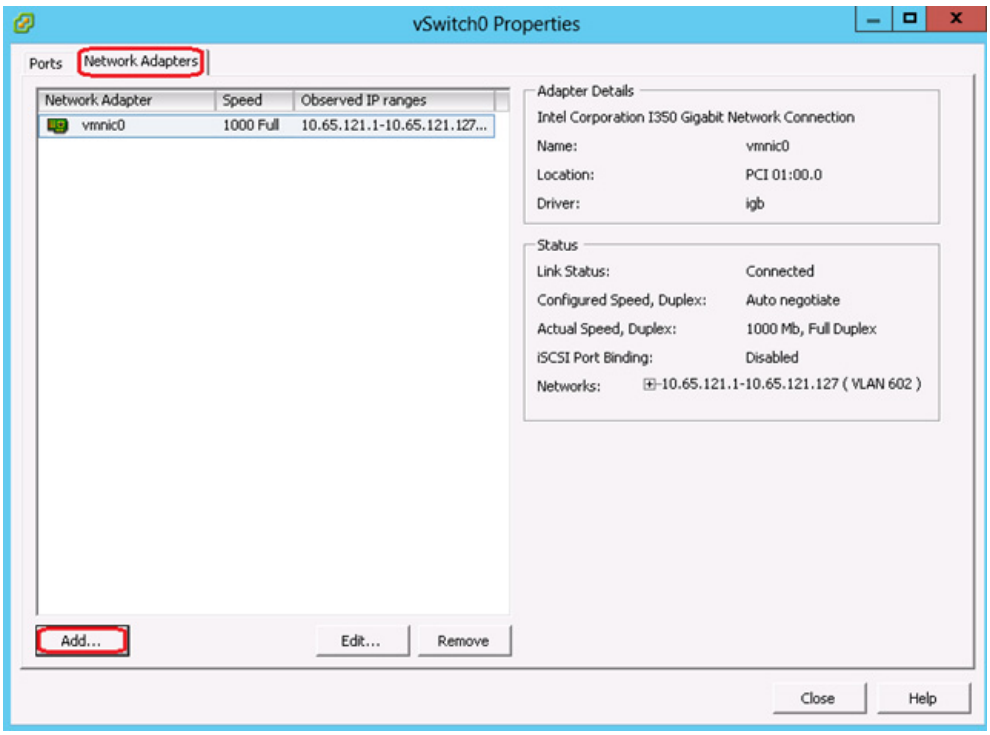

4. Select the appropriate vmnic and click **Next**. In this configuration, vmnic0 and vmnic1 connect to Cisco Nexus Switch A and Switch B respectively, which were configured as virtual port channels.

Figure 68 Management vSwitch – Add Adapter Wizard

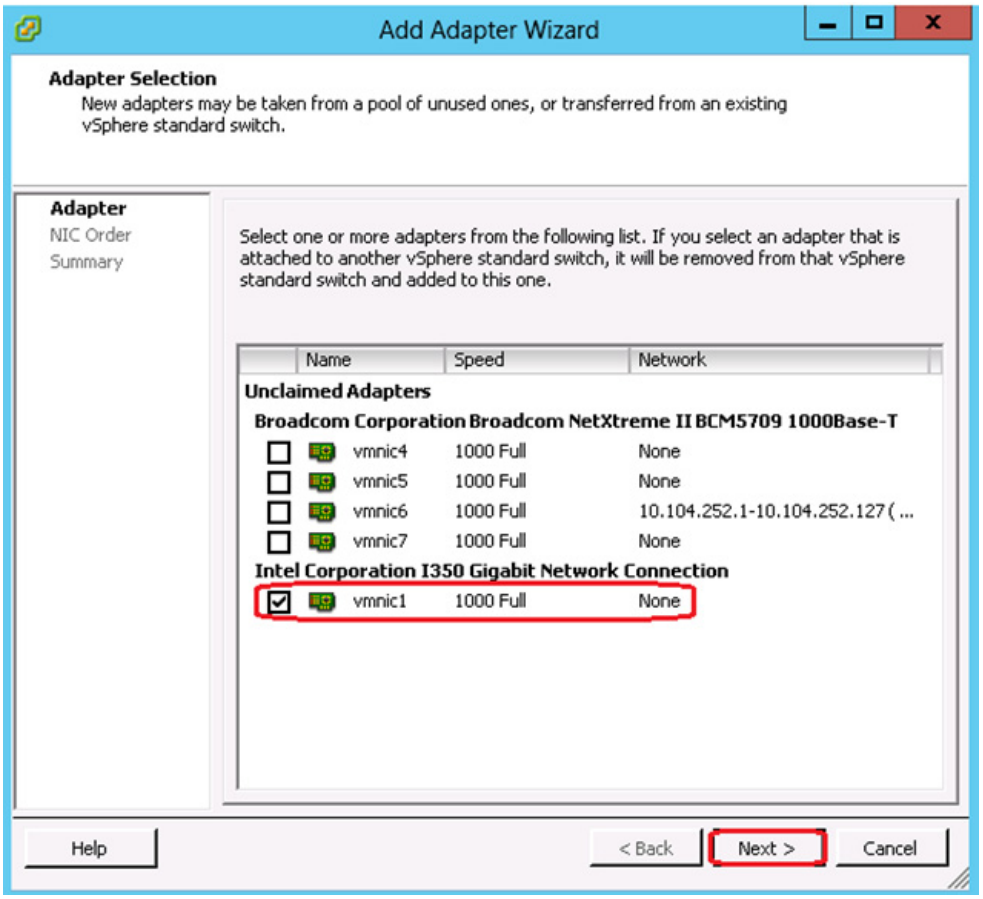

5. Set the NIC Order in the Add Adapter Wizard as per your requirement and click **Next**.
6. In the Summary page review and click **Finish**.
7. In the vSwitch0 Properties, click **Ports** tab and select **vSwitch**. Click **Edit** to make the required changes.

Figure 69 Management vSwitch – Ports Edit

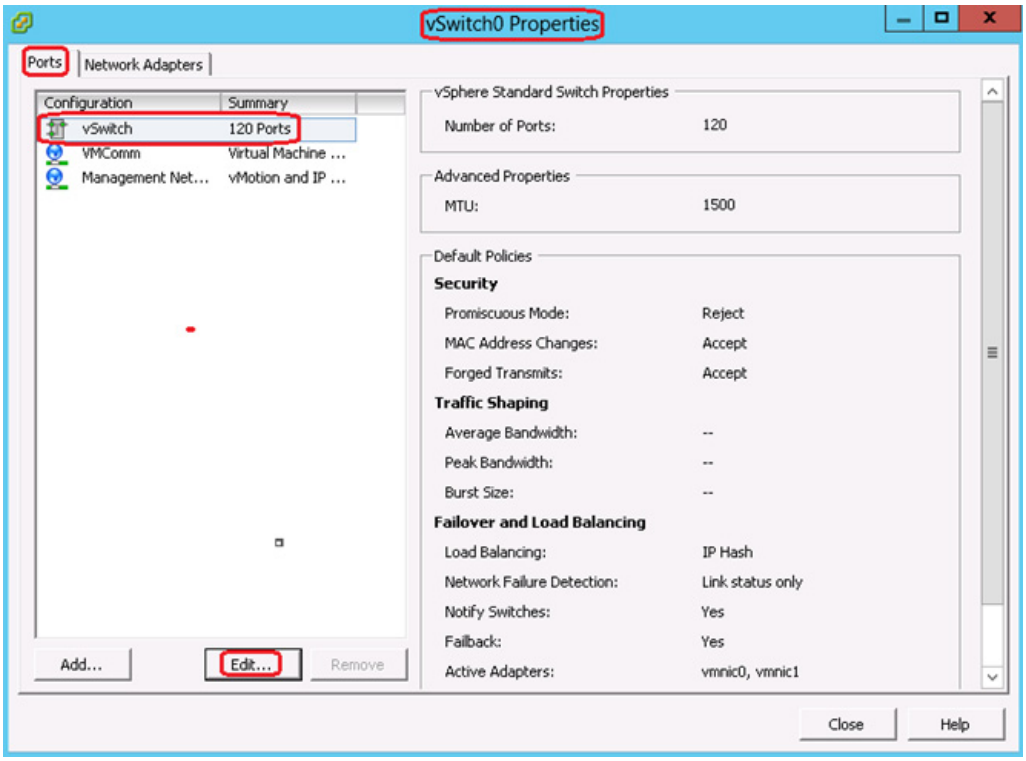

8. Click **NIC Teaming** tab, for Load Balancing choose Route based on IP hash from the drop-down list. Both vmnics should be listed under Active Adapters.

Figure 70 Management vSwitch – NIC Teaming

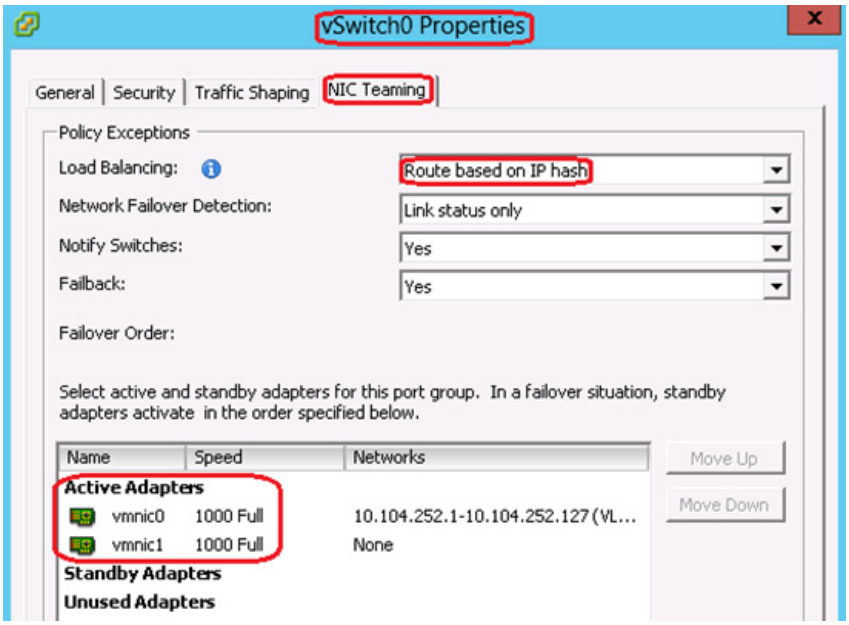

Note When IP-hash load balancing is used, do not use beacon probing and do not configure standby uplinks.

Since we are using Ether channel for link aggregation on the Cisco Nexus switches, the supported vSwitch NIC teaming mode is IP hash. For more details, see the VMware KB articles at:

http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1004048

http://kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1001938

9. Edit VM Network port group by providing a Network Label and VLAN ID. Management Network and VM Network will use the same standard vSwitch NIC Teaming Load Balancing policy configured in step 8.

Figure 71 Management vSwitch – VM Network VLAN Settings

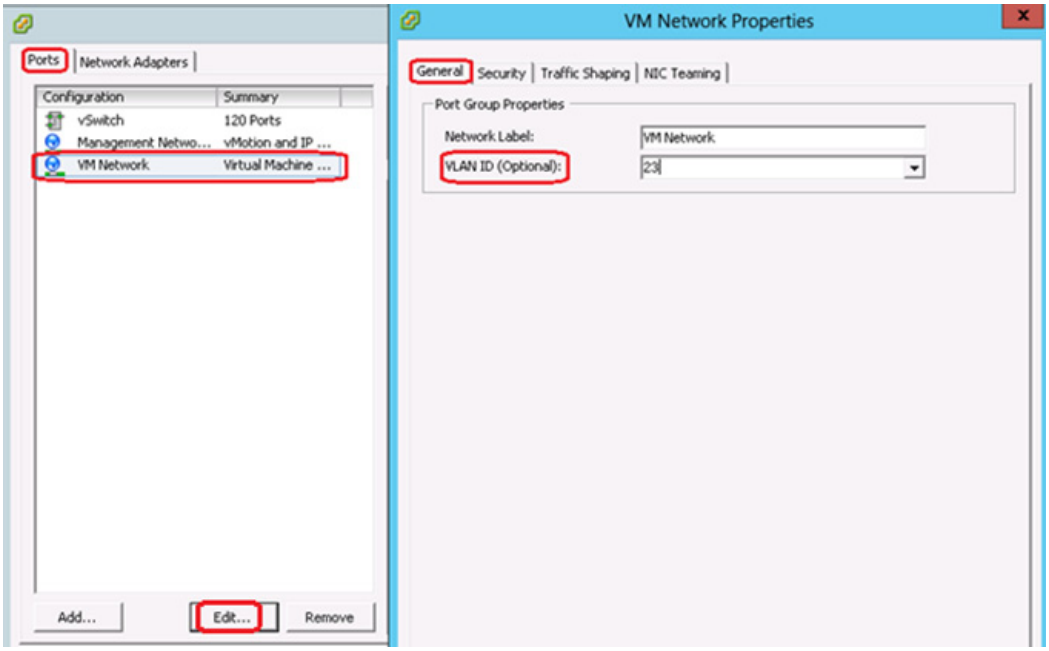

Figure 72 Management VMkernel Port-Group

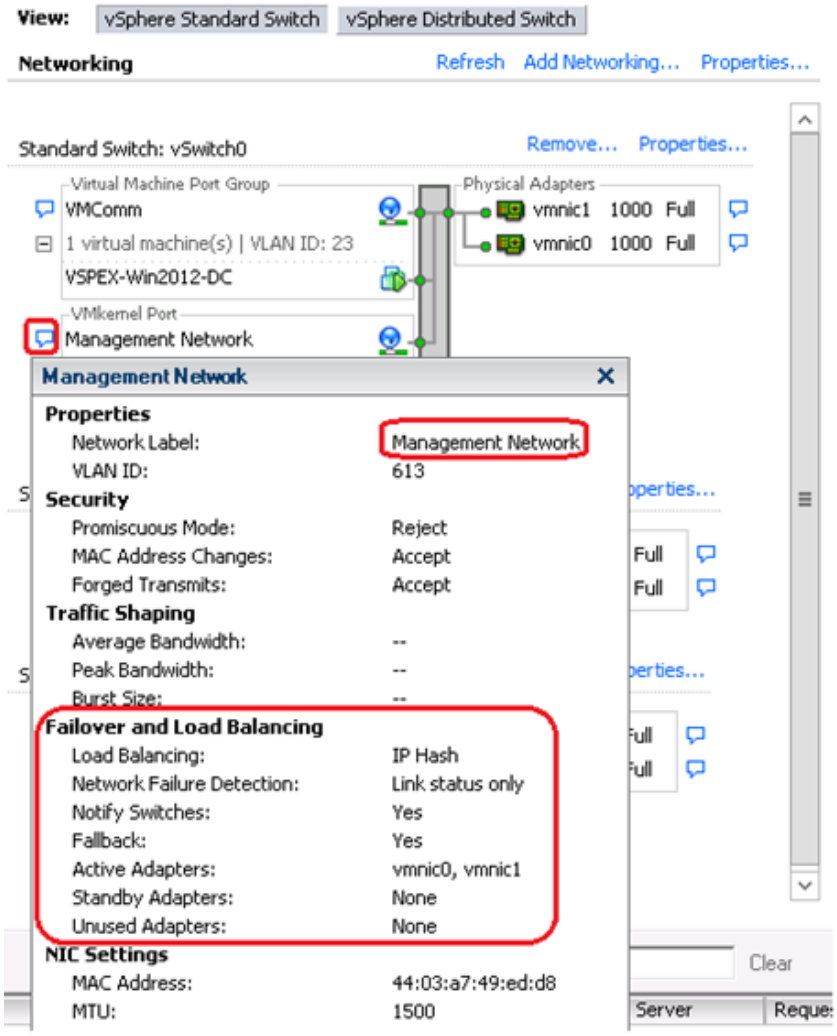

10. Repeat all the steps in this section on the other two ESXi hosts.

Create and Configure vSwitch for vMotion

In this section we will create a new vSwitch for vMotion, add two uplink active adapters and set the load balancing policy to route based on IP hash. The MTU on this vSwitch will be set to 9000 (jumbo frames)

1. Connect to the vCenter server using vSphere client.
2. Select an ESXi host on the left pane of Hosts and Clusters window. Click **Configuration > Networking > vSwitch0 Properties** on the right pane of the window.

Figure 73 VMware vCenter – Add Networking

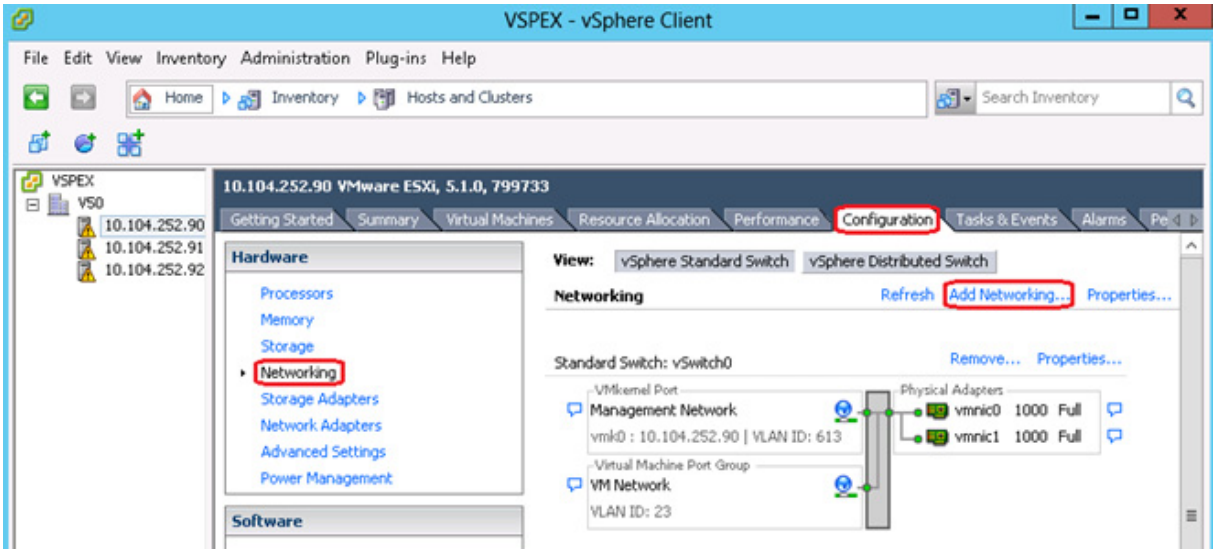

3. Select VMkernel under Connection Types and click **Next**.

Figure 74 Add Network Wizard – Connection Type

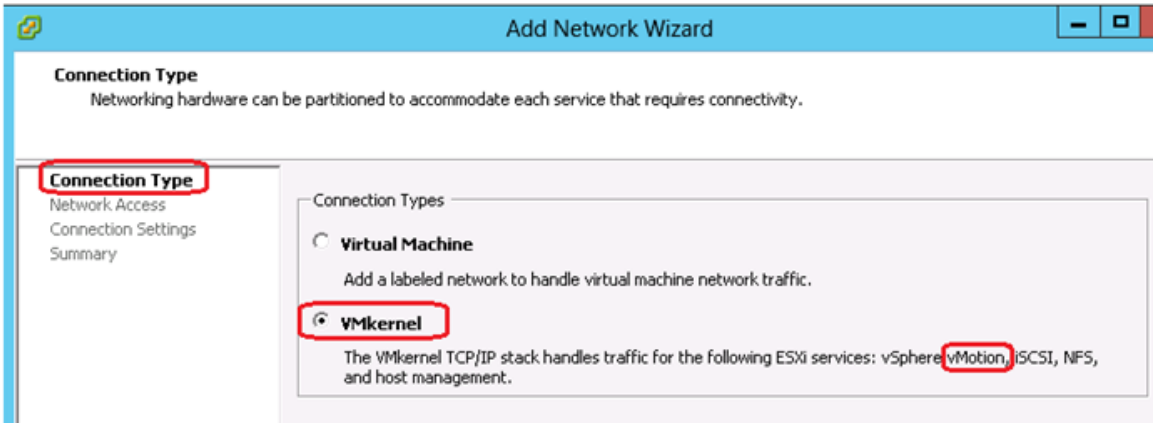

4. In the Network Access section select Create a vSphere standard switch and select the two vmnics configured for vMotion VLAN.

Figure 75 Add Network Wizard – Network Access

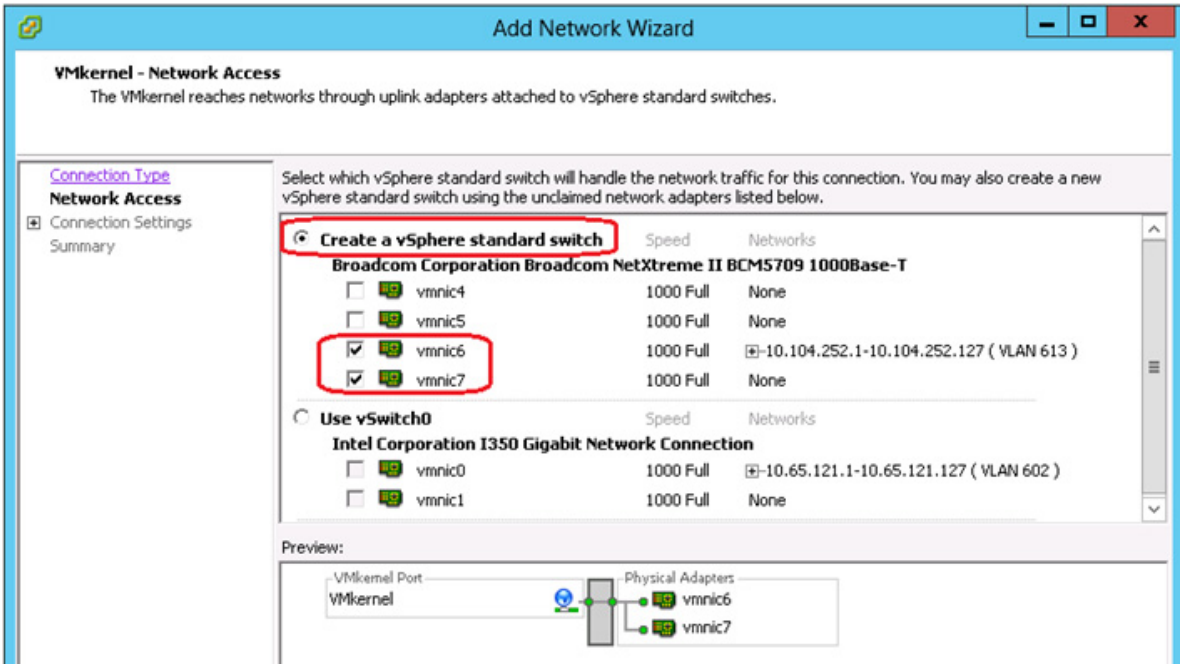

5. In the Connection Settings page, provide a Name for Network Label, a VLAN ID and select Use this port group for vMotion.

Figure 76 Add Network Wizard – Connection Settings

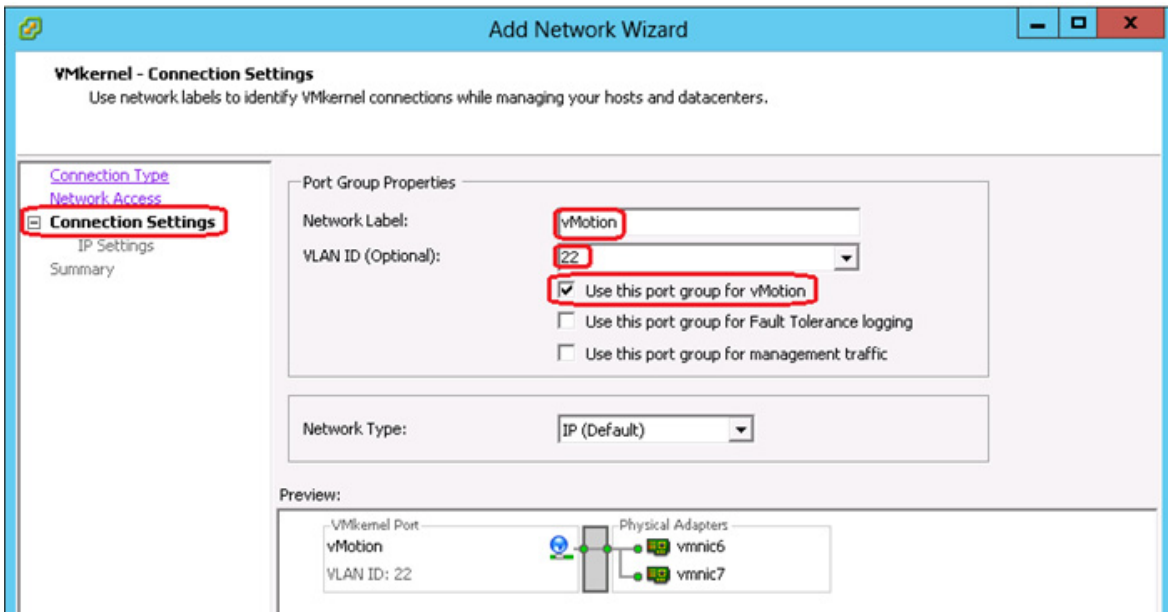

6. In the IP Settings screen enter an IP Address, Subnet Mask and click **Next**.

Figure 77 Add Network Wizard – VMkernel IP Settings

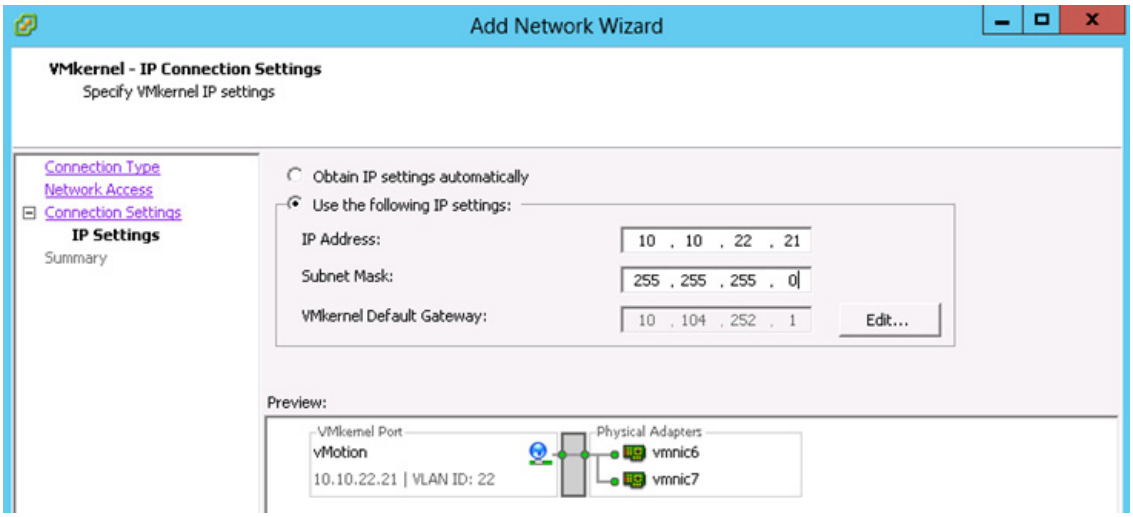

7. In the Summary page click **Finish**. The created vSwitch for vMotion gets listed in the vSwitch1area as shown in Figure 78.

Figure 78 VMware vCenter – Network Configuration

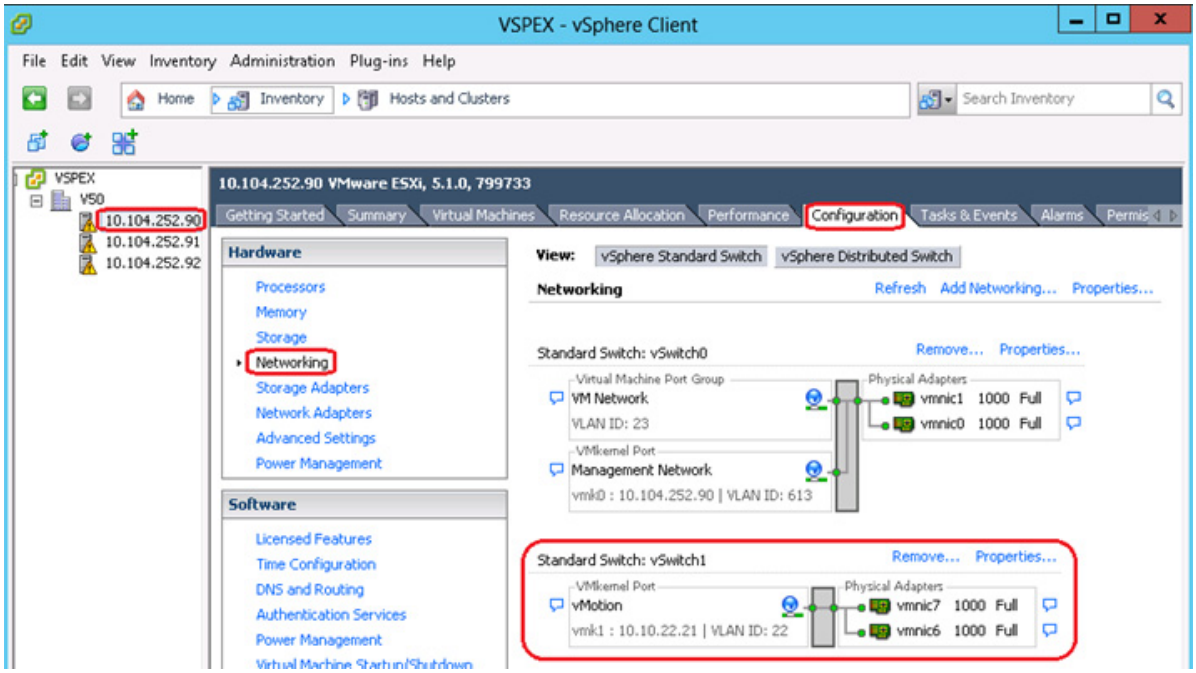

8. In the vSwitch1 area, click **Properties** and under the **General** tab set the MTU size to 9000.

Figure 79 vSwitch Properties - MTU

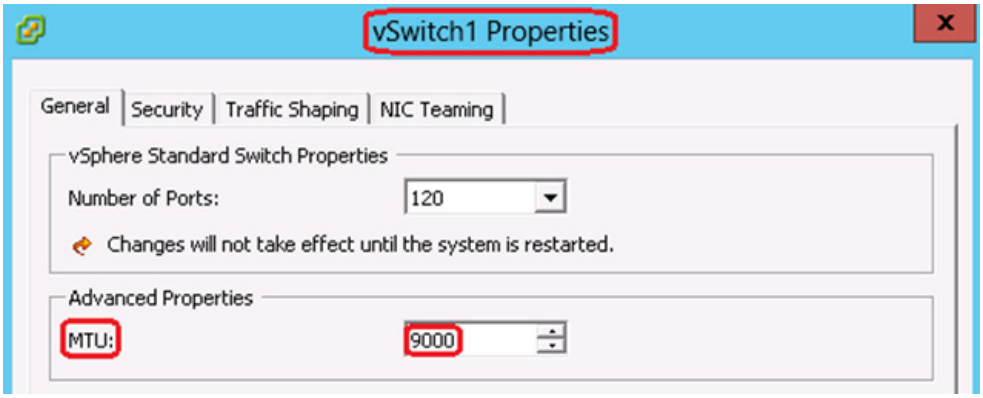

9. Click the **NIC Teaming** tab and for Load Balancing policy, choose Route based on IP hash option from the drop-down list. Both vmnics should be listed under Active Adapters.

Figure 80 vSwitch Properties – NIC Teaming Policy

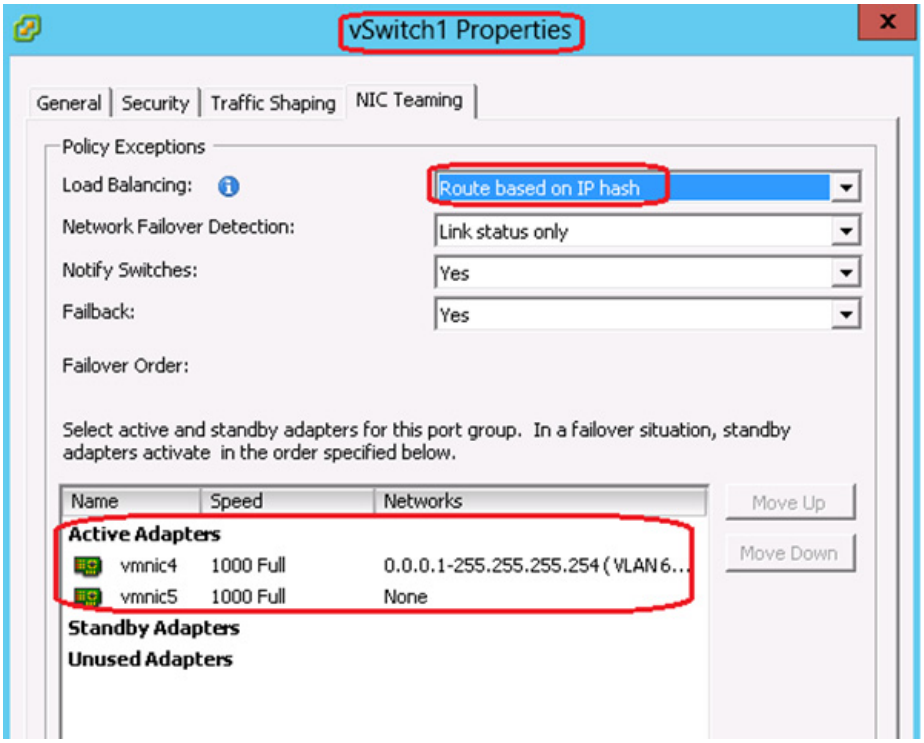

10. Repeat all the steps in this section on the other two ESXi hosts.

Create and Configure vSwitch for iSCSI

The iSCSI adapter and physical NIC connect through a virtual VMkernel adapter, also called virtual network adapter or VMkernel port. You create a VMkernel adapter (vmk) on a vSphere switch (vSwitch) using 1:1 mapping between each virtual and physical network adapter.

One way to achieve the 1:1 mapping when you have multiple NICs, is to designate a separate vSphere switch for each virtual-to-physical adapter pair as shown in the below figure.

Figure 81 1:1 adapter mapping on separate vSphere standard switches

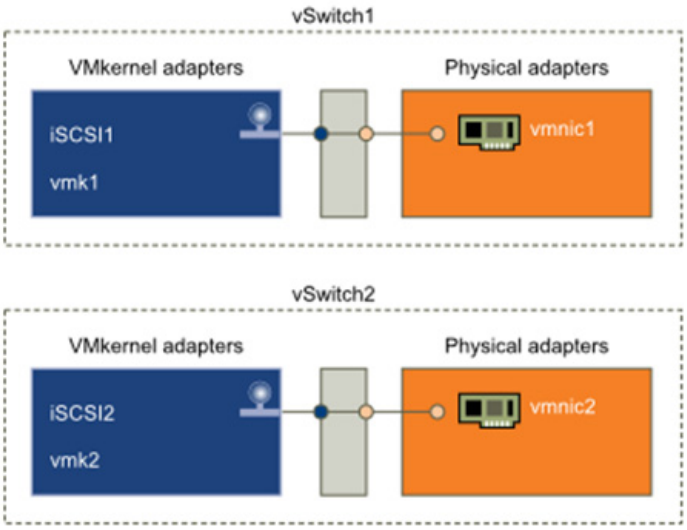

Note If you use separate vSphere switches, you must connect them to different IP subnets. Otherwise, VMkernel adapters might experience connectivity problems and the host will fail to discover iSCSI LUNs.

In this section we will be creating two vSwitches, each with a single vmnic uplink and a single VMkernel port, bound to the iSCSI adapter. A MTU size of 9000 will also be set on the vSwitches and VMkernal ports to enable jumbo frames. Follow these steps to create vSwitch:

1. Connect to the vCenter server using vSphere client.
2. Select an ESXi host on the left pane of Hosts and Clusters window. Click **Configuration** > **Networking** > **Add Networking** on the right pane of the window.

Figure 82 VMware vCenter – Add Networking

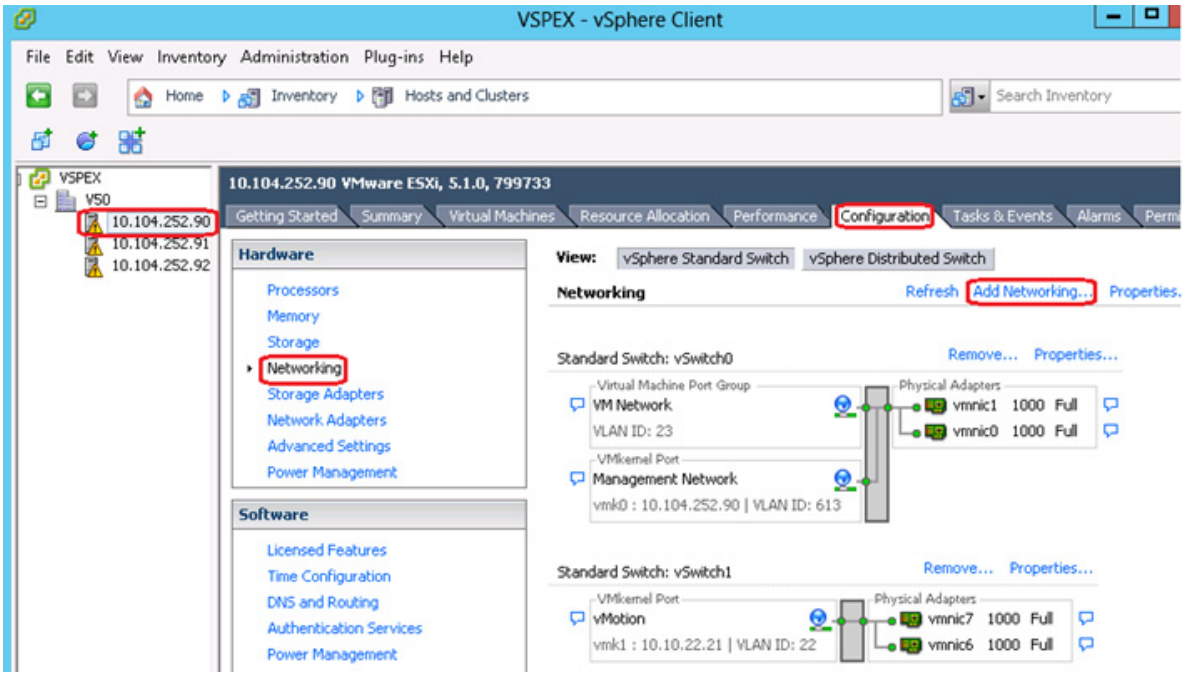

3. Select VMkernel under Connection Types and click **Next**.

Figure 83 Add Network Wizard – Connection Type

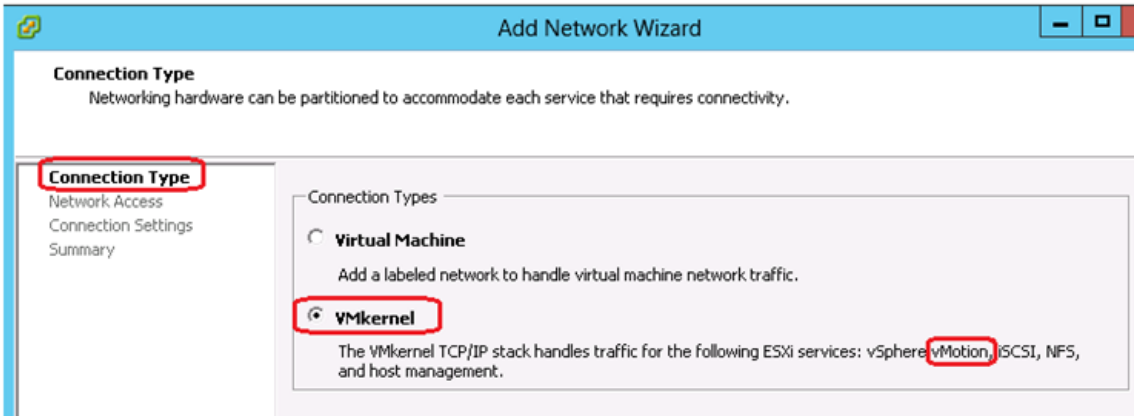

4. In the Network Access section select Create a vSphere standard switch and select the appropriate vmnic.

Figure 84 Add Network Wizard – Network Access

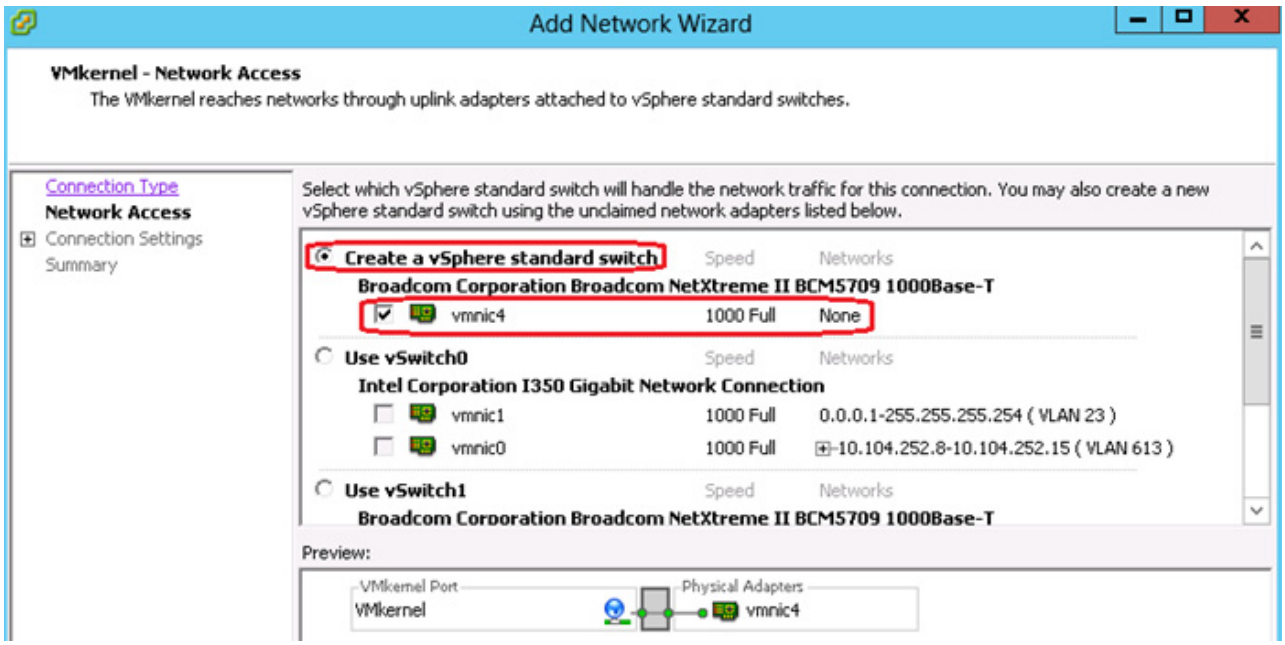

5. In the Connection Settings window, provide a Name for the Network Label. Leave the VLAN ID to None because the corresponding switch port at the other end is configured as vlan access port.

Figure 85 Add Network Wizard – Connection Settings

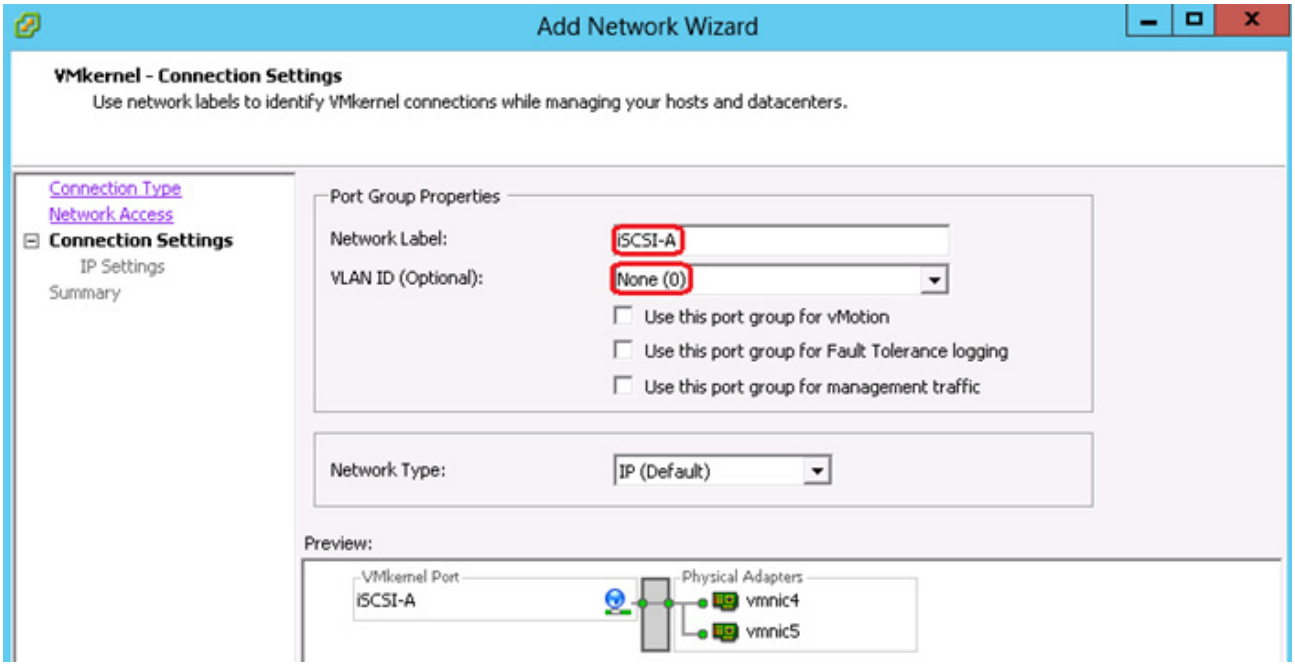

6. In the IP Settings screen enter an IP Address, Subnet Mask and click **Next**. This IP address range is from the iSCSI-A vlan.

Figure 86 Add Network Wizard – VMkernel IP Settings

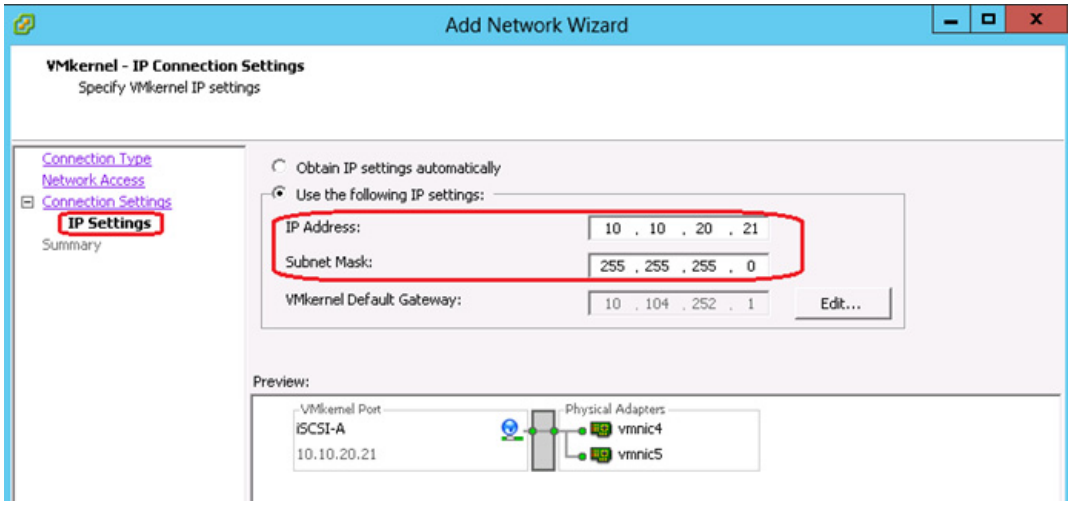

7. In the Summary page click **Finish**. You should now be able to see the vSwitch for iSCSI-A created.

Figure 87 iSCSI vSwitch - Properties

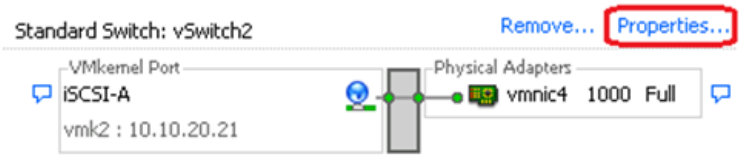

8. In the vSwitch2 Properties, select the vSwitch under **Ports** tab and click **Edit**. Set the MTU size to 9000 under the **General** tab to enable jumbo frames.

Figure 88 iSCSI vSwitch Properties - Edit

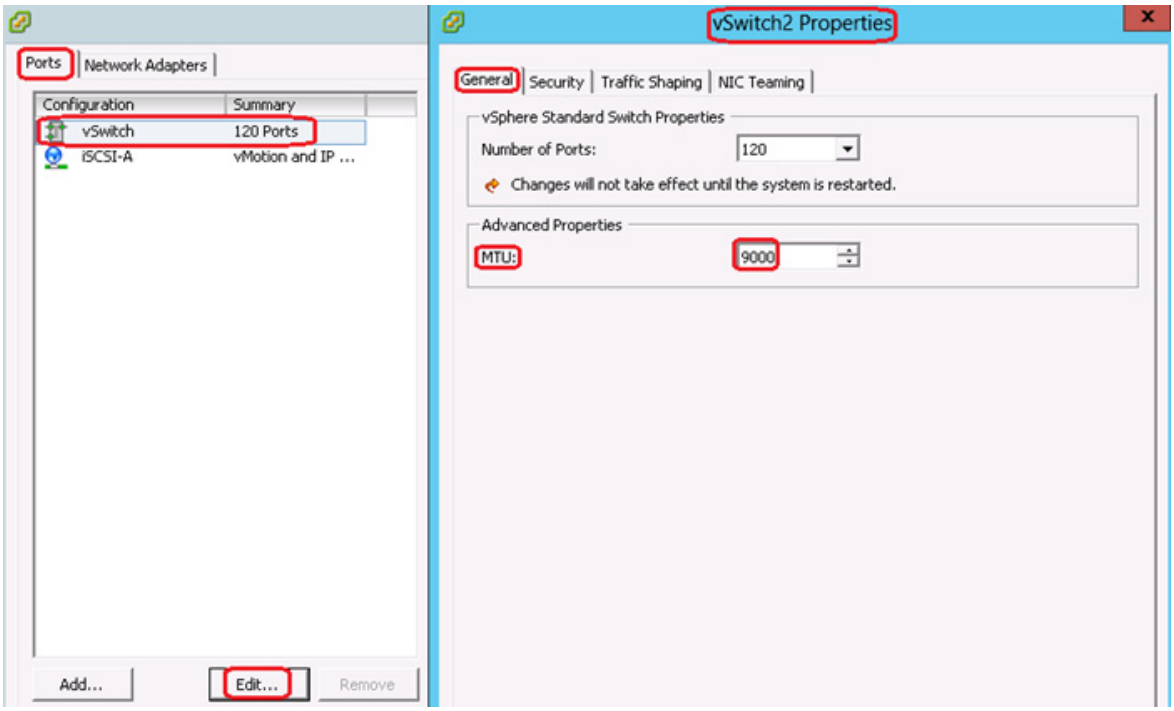

9. Set the MTU size to 9000 for the iSCSI-A VMkernel.

Figure 89 iSCSI vSwitch Properties - MTU

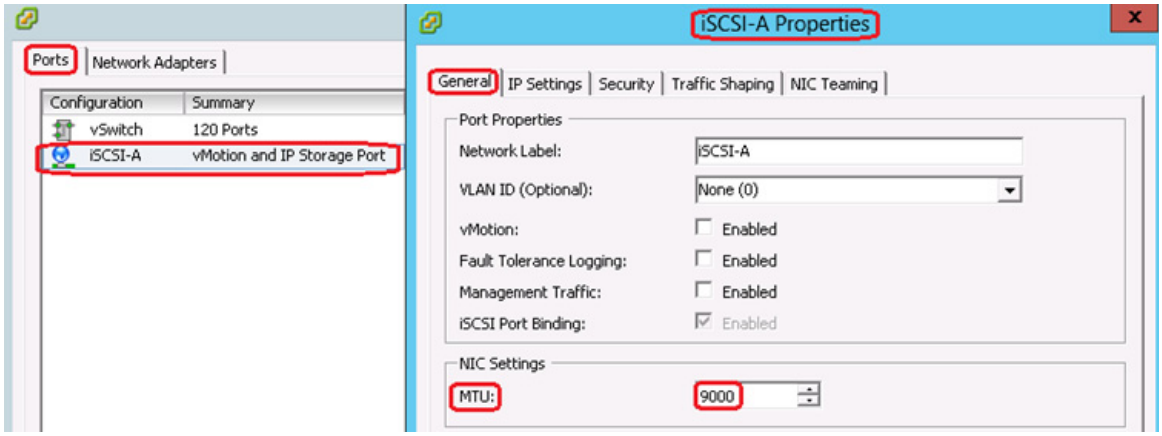

10. Repeat the steps 2 to 9 to create another vSwitch with Network label as iSCSI-B. Select the appropriate vmnic uplink adapter and assign IP address from the iSCSI-B VLAN range as shown in Figure 90.

Figure 90 iSCSI vSwitches vCenter view

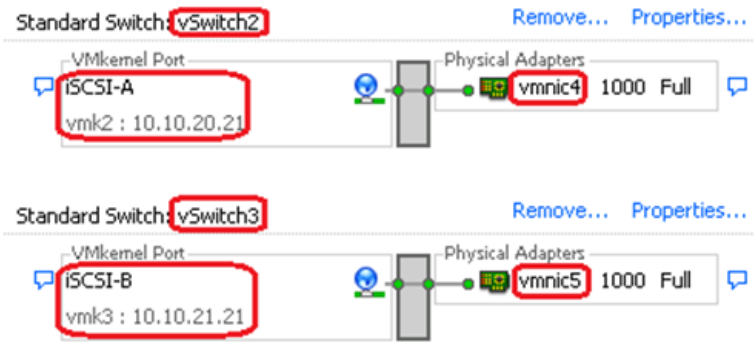

Configuring Software iSCSI Adapter (vmhba)

The software iSCSI adapter (vmbha) uses standard NICs to connect your host to a remote iSCSI target on the IP network. The software iSCSI adapter that is built into ESXi facilitates this connection by communicating with the physical NICs through the network stack. Before you can use the software iSCSI adapter, you must set up networking (this is completed in the previous section - Create and Configure vSwitch2 for iSCSI), activate the adapter, and configure parameters such as discovery addresses and CHAP. This section describes these steps to configure the software iSCSI storage adapter to access the remote storage configured on the VNXe3150 storage array.

1. Login to the vSphere Client, and select a host from the inventory panel.
2. Click the **Configuration** tab and click **Storage Adapters** in the Hardware panel.
3. Select and right click vmhba under the iSCSI Software Adapter and click **Properties**. If you do not see the software iSCSI adapter, then click **Add** and select Software iSCSI Adapter.

Figure 91 iSCSI Software Adapter

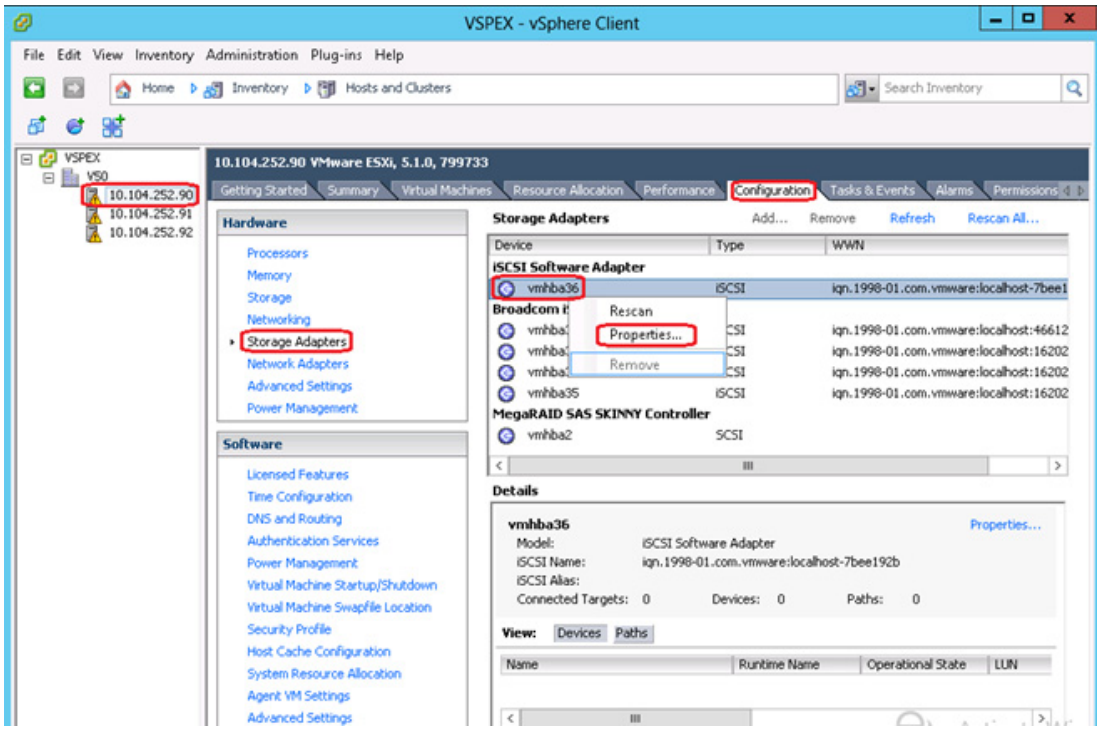

4. In the **General** tab, click **Configure**.

Figure 92 iSCSI Software Adapter Properties

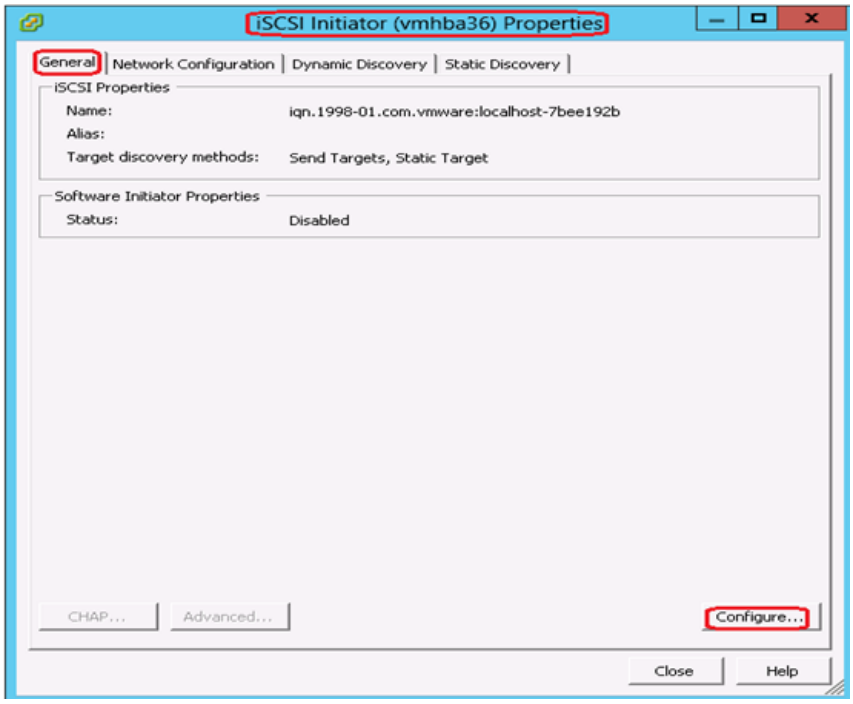

5. Make sure that the adapter is enabled and click **OK**. After enabling the adapter, the host assigns the default iSCSI name to it. If you change the default name, follow iSCSI naming conventions.

Figure 93 iSCSI Software Adapter - Enable

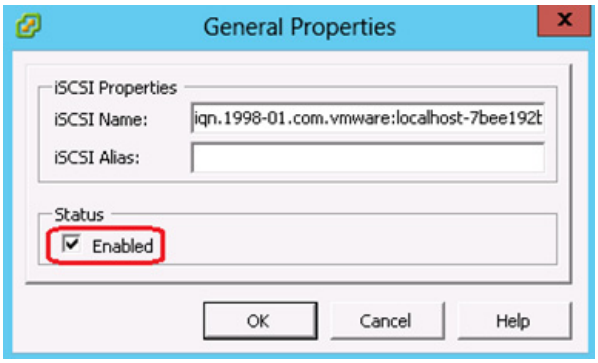

5. Click the **Network Configuration** tab and click **Add**.

Figure 94 iSCSI Software Adapter Properties – Network Configuration

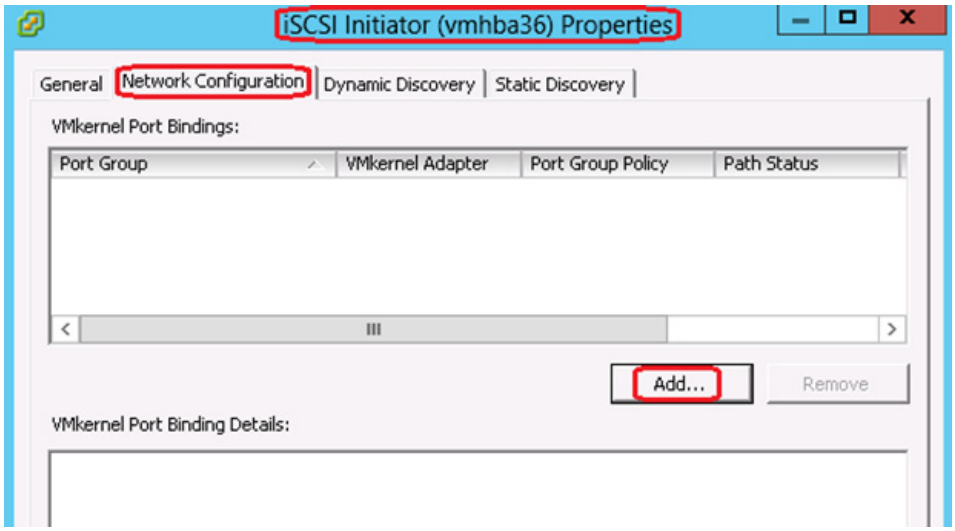

6. Select both the VMkernel adapters (vmk2 and vmk3) one at a time and click **OK** to bind it with the iSCSI adapter.

Figure 95 iSCSI Software Adapter Properties – Bind VMkernel

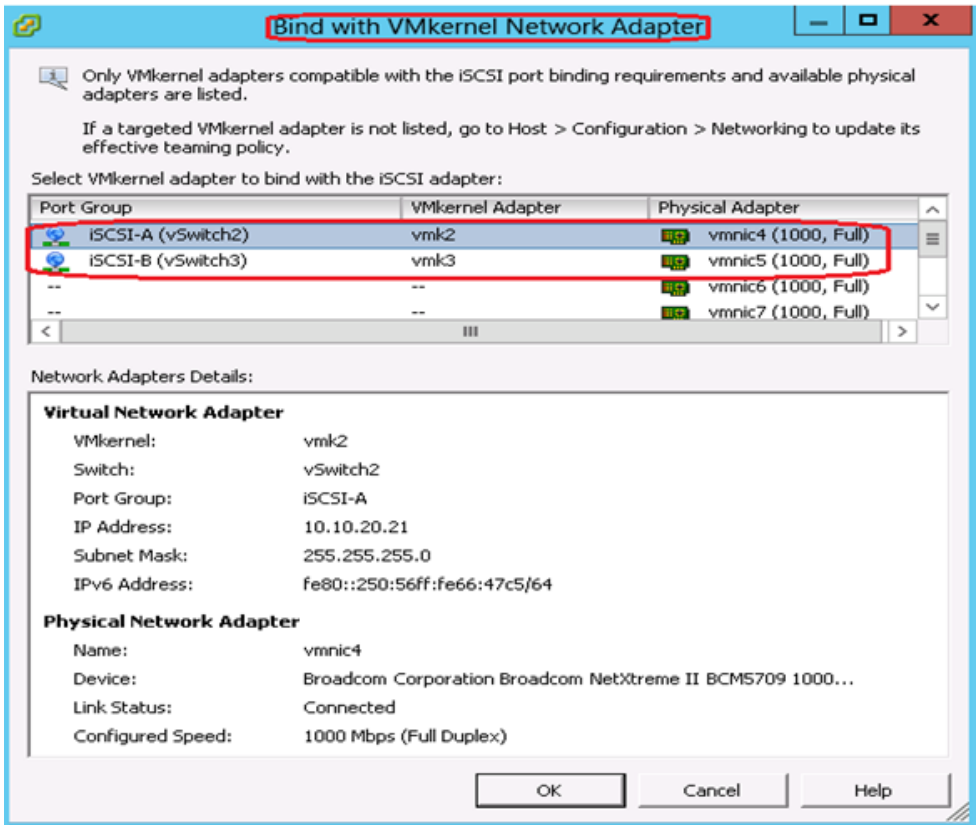

The network connection appears on the list of VMkernel port bindings for the iSCSI adapter as shown in Figure 96.

Figure 96 iSCSI Software Adapter Properties – Network Configuration

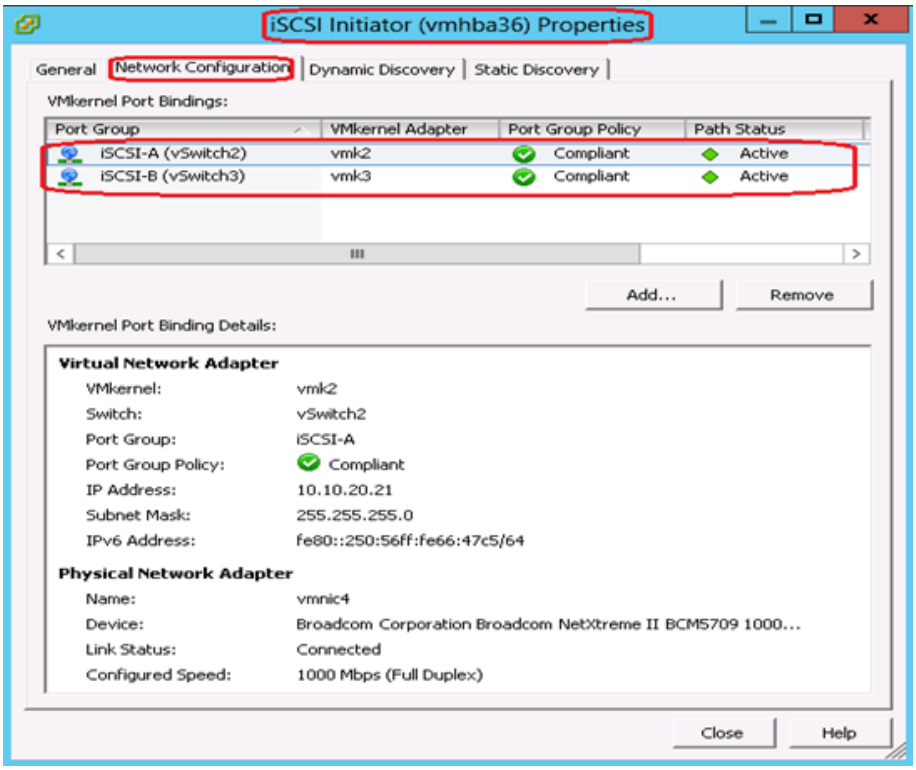

Prepare the EMC VNXe3150 Storage

Figure 97 shows the steps involved in deploying VMware environments by using VNXe storage.

Figure 97 VNXe VMware vSphere storage provisioning flowchart

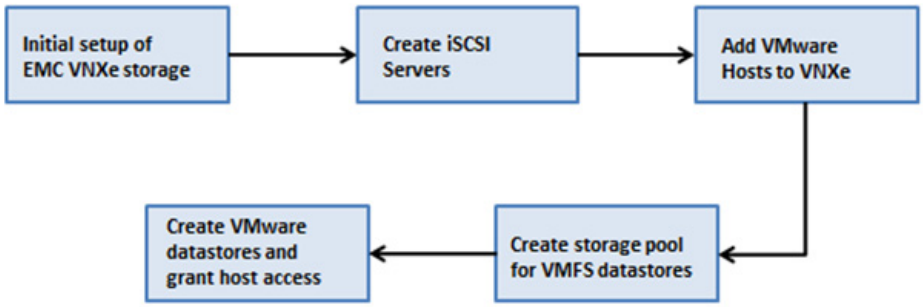

Initial Setup of VNXe

1. Connect the Ethernet cables from the management and data ports to the network as shown in the cabling guide.

2. Assign an IP address to the management interface or Download and run the Connection Utility to establish an IP address for managing the VNXe storage system. The Connection Utility can be downloaded directly from the product support page.

   http://www.emc.com/support-training/support/emc-powerlink.htm

3. Complete the installation/upgradation of the software and activate the licenses.
4. Connect to the VNXe system from a web browser using the management IP address.

Figure 98 EMC Unisphere - Dashboard Page

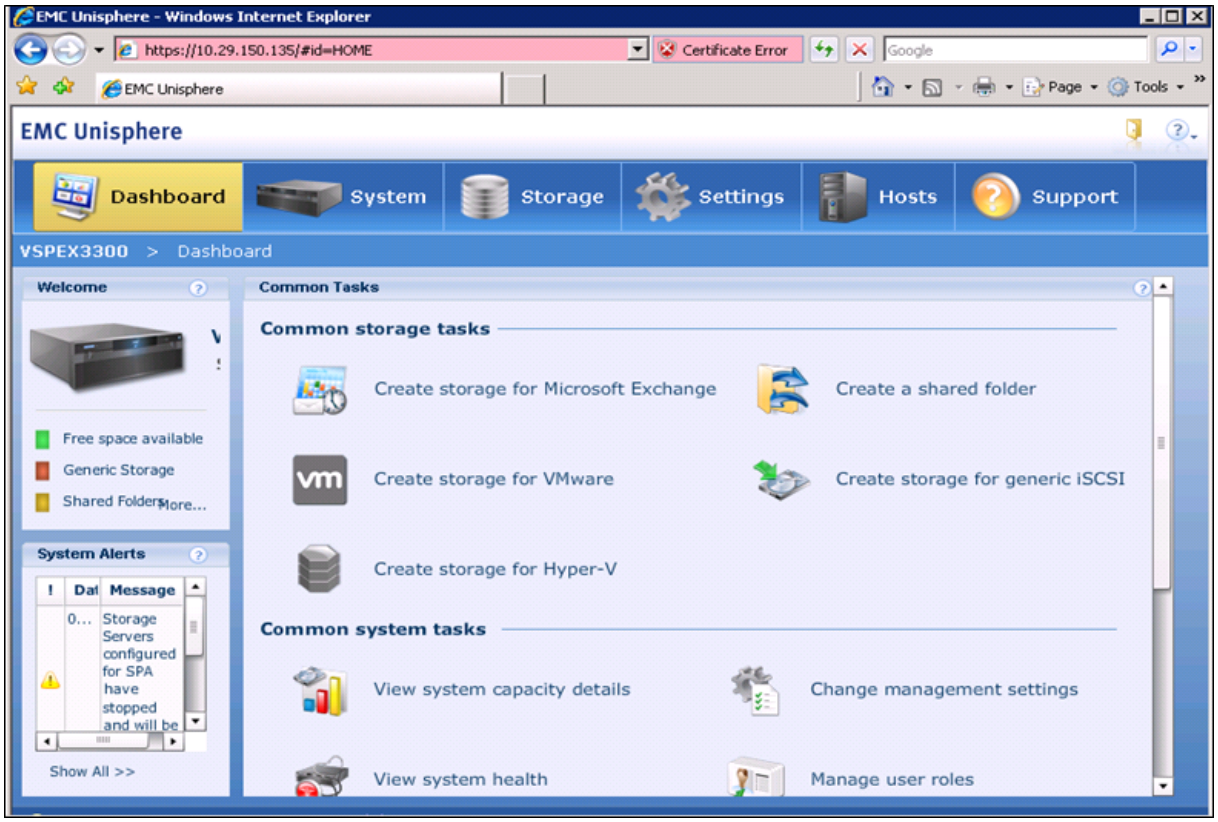

Note The SP A and SP B network data ports must be connected to the same subnet. In general, both SPs should have mirrored configurations for all front-end cabling (including VLANs) in order to provide Failover.

iSCSI Server Configuration

The iSCSI Storage Server is the portal through which storage will be accessed by the hosts within the Fast Track configuration. The goal of the proposed iSCSI server configuration is to provide redundancy, multi-pathing and balanced access across all 1 GigE connections and both storage processors. Each 1 GigE module will have 2 ports, referred to as eth10 and eth11. Considering there is an I/O module for each storage processor, both SPA and SPB will have eth10 and eth11 connections.

iSCSI servers will run on either SPA or SPB. This means storage assigned to a given iSCSI server will only be available to one SP at a given time. To utilize both SPA and SPB concurrently, two iSCSI servers will be created.

With respect to iSCSI server high availability, the eth10 and eth11 connections are paired across the storage processors. If an iSCSI server running with an IP address dedicated to eth10 on SP A needs to move to SP B, for maintenance as an example, the IP address will move to the corresponding eth10 port on SPB. Therefore subnet connectivity will need to be the same for the associated eth10 and eth11 connections across the storage processors. Figure 99 shows a logical example of the connections.

Figure 99 VNXe Array Logical Network Connections

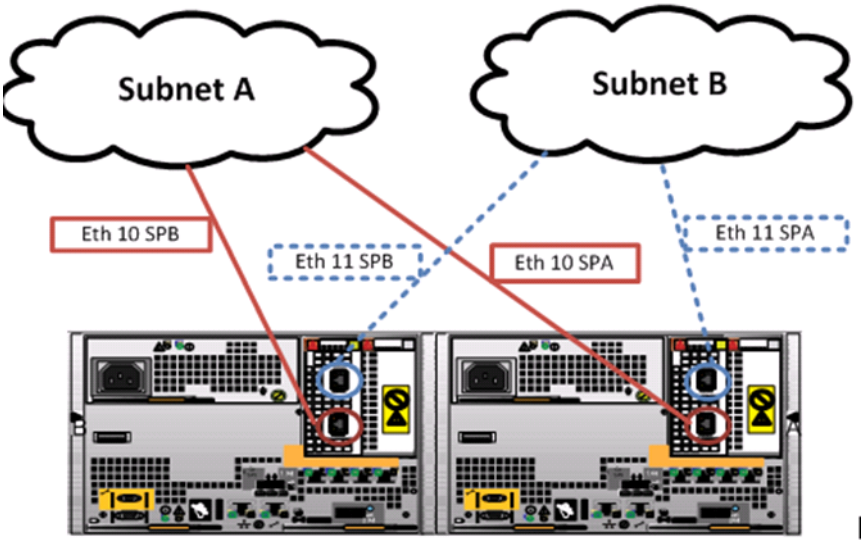

Table 9 Sample IP Configuration

iSCSI Server A	iSCSI Server B
IP Address Eth10 Subnet A (10.10.20.10/24)	IP Address Eth10 Subnet A (10.10.20.11/24)
IP Address Eth11 Subnet B (10.10.20.10/24)	IP Address Eth11 Subnet B (10.10.21.11/24)

Configure iSCSI Storage Servers

1. In the EMC Unisphere home page, choose **Settings > iSCSI Server Settings > Add iSCSI Server**.
2. Enter the desired Server Name, IP address, and Subnet Mask (default gateway is not required as the connection from the servers is on the same subnet). Click **Show Advanced** and select the appropriate storage processor (SPA) and Ethernet Port (eth10) as shown in Figure 100.

Figure 100 EMC Unisphere - iSCSI Server SP-A eth10

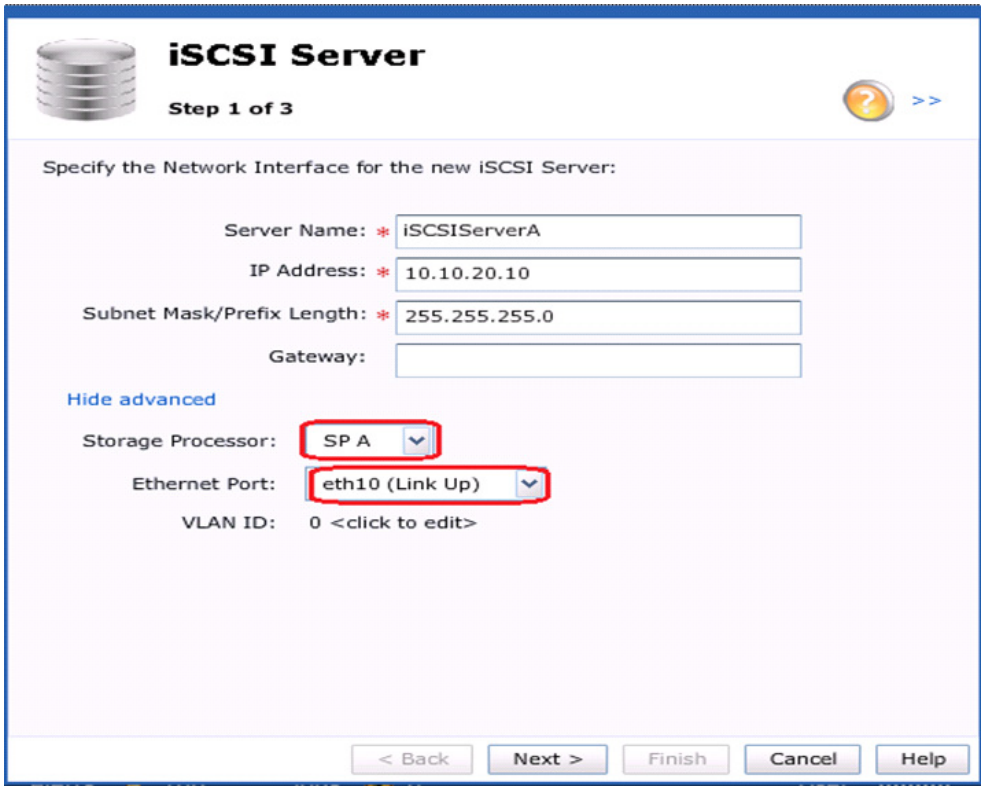

3. Repeat the steps 1 and 2 to create a second iSCSI server on SP-B and eth10.

Figure 101 EMC Unisphere - iSCSI Server SP-B eth10

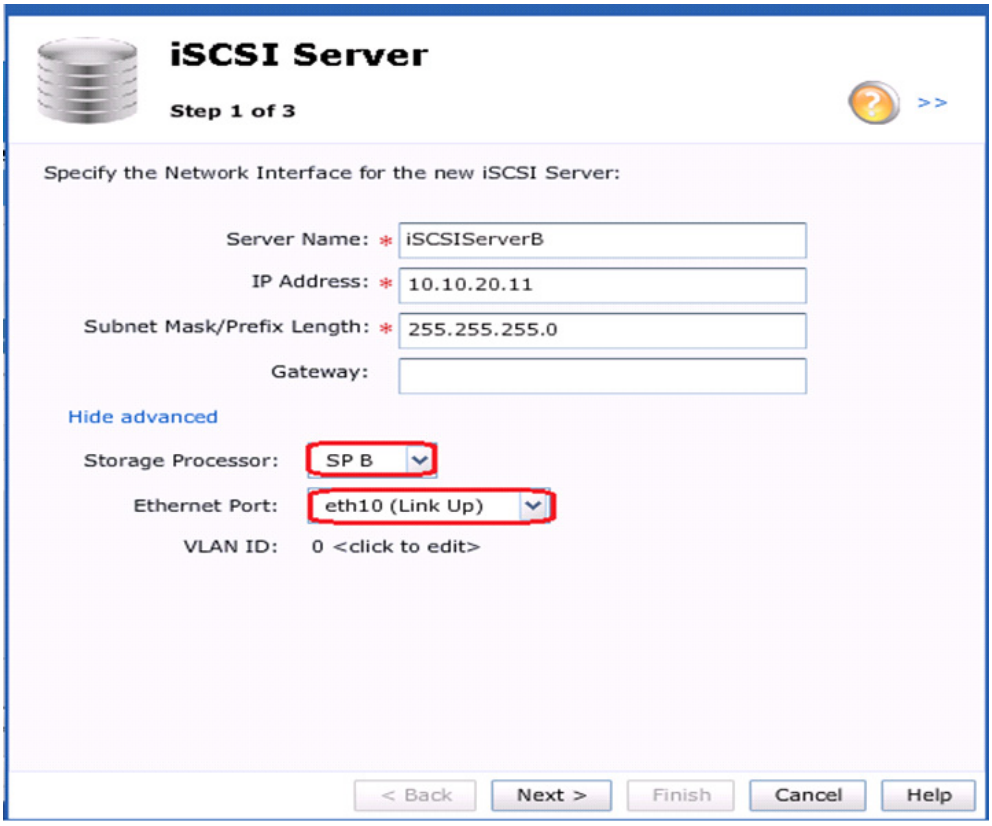

4. Select the previously created iSCSI server and select **Details**.

Figure 102 EMC Unisphere - iSCSI Server Settings

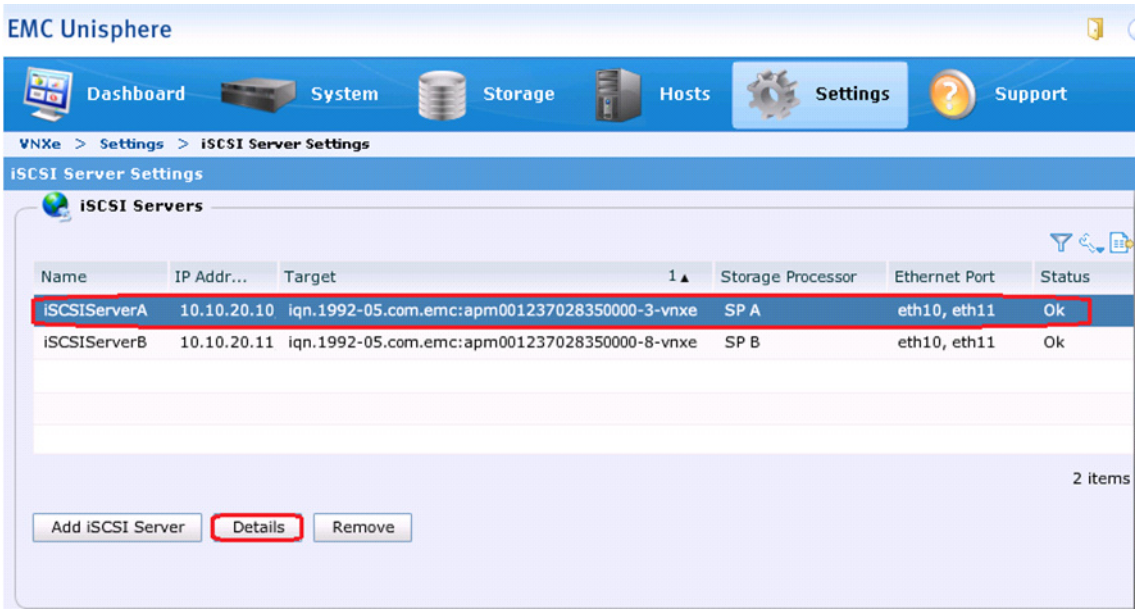

5. In the **iSCSI Server Details** window, click **Add Network Interface.**
6. Enter the appropriate IP Address, Subnet Mask and Gateway information.

Figure 103 EMC Unisphere - iSCSI Server SP-A eth11

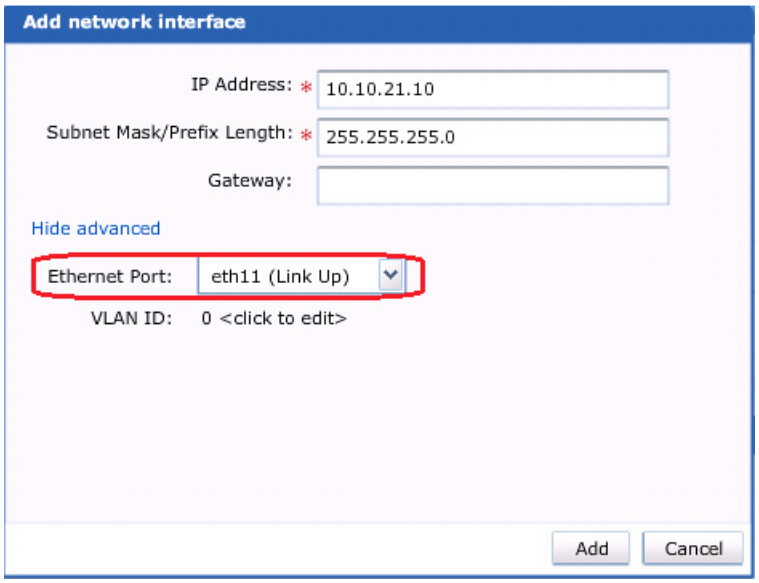

7. Repeat the steps 4,5, and 6 for the iSCSI Server instance assigned to the other storage processor, SP-B.

Figure 104 EMC Unisphere - iSCSI Server SP-B eth11

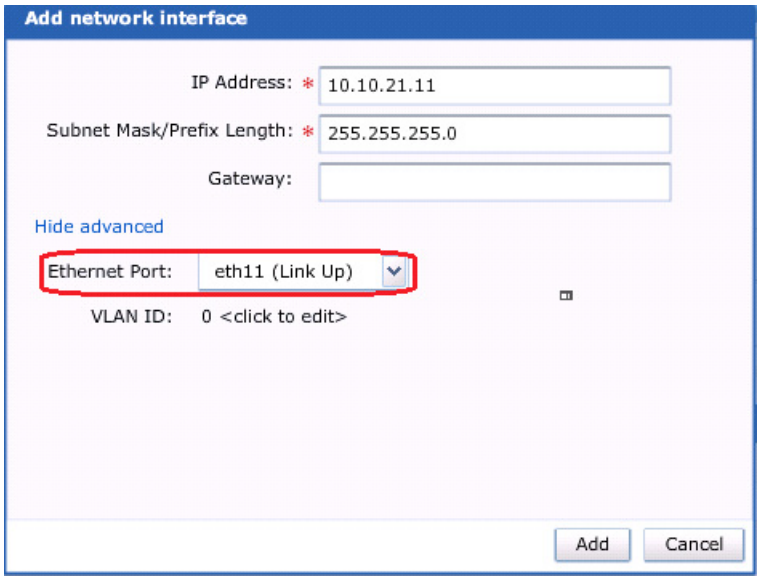

Note In VNXe storage systems, for fail safe networking (FSN) and high availability features to work, the peer ports on both storage processors must belong to the same subnet. For more information about high availability in VNXe storage systems is available in the below URL.

http://www.emc.com/collateral/hardware/white-papers/h8178-vnxe-storage-systems-wp.pdf

Figure 105 EMC Unisphere - iSCSI Server Settings

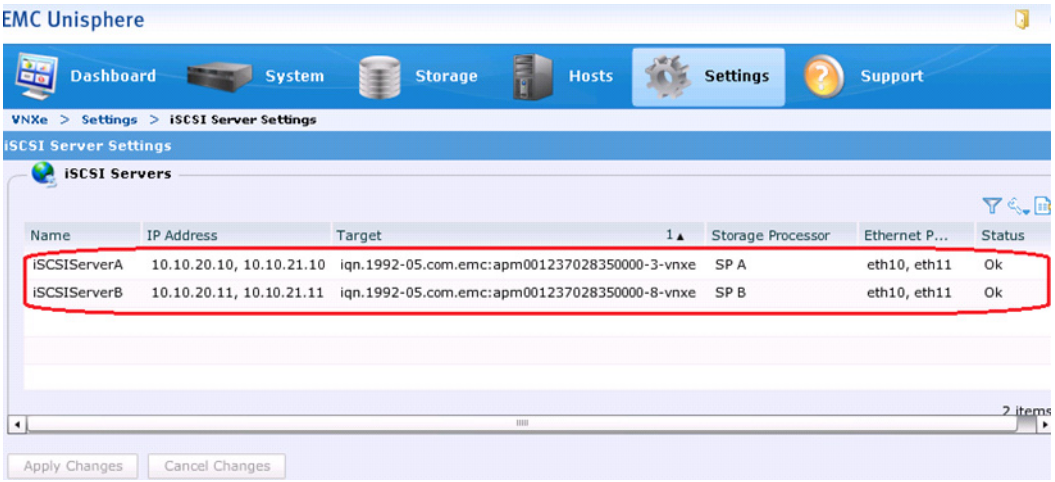

Add VMware Hosts to EMC VNXe Storage

To integrate virtual infrastructure with VNXe for storage provisioning, follow these steps:

1. In Unisphere, choose **Hosts** > **VMware**. The VMware page appears.

Figure 106 EMC Unisphere – Hosts

2. Click **Find ESX Hosts**.

Figure 107 EMC Unisphere – Find ESX Hosts

3. In the Find ESX Hosts window, click the radio button **Other (vCenter or ESX Host)** and enter the IP address of the VMware vCenter server.

Figure 108 EMC Unisphere – Find ESX Hosts

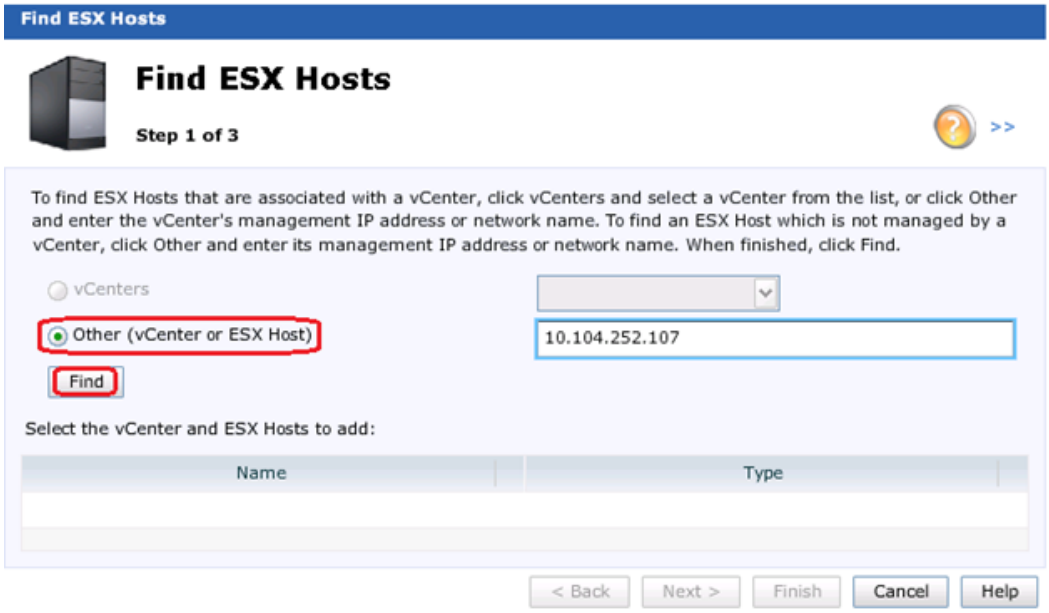

4. Click **Find** and in the Enter Credentials for vCenter/ESX Host window, enter the Username and Password for the vCenter server. Click **OK**.

Figure 109 EMC Unisphere – vCenter Credentials

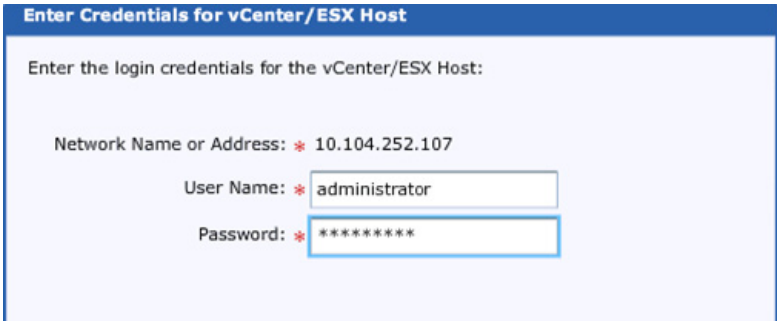

5. Select the vCenter and ESX hosts to add by checking the check boxes against each of the ESX hosts and click **Next**.

Figure 110 EMC Unisphere – Find ESX Hosts Selection

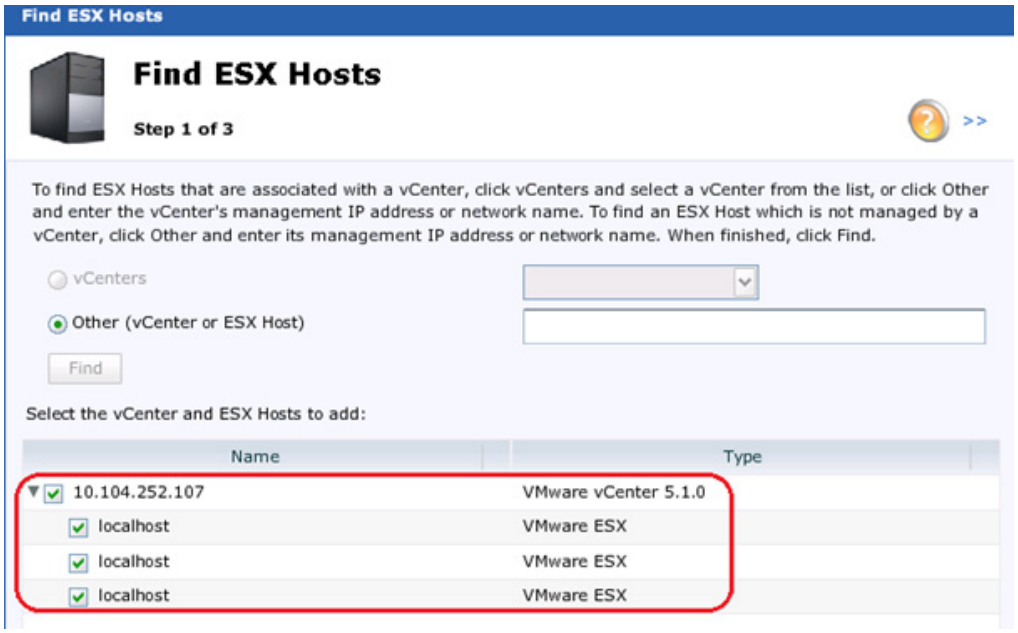

6. In the Summary page, confirm the selected ESX hosts and click **Finish**.

Create Storage Pools for VMware

A storage pool is an aggregation of storage disk resources that are configured with a particular storage profile. The storage profile defines the type of disks that are used to provide storage and the RAID configuration of the component disks. The storage pool configuration defines the number of disks and the quantity of storage associated with the pool.

Create a pool with the appropriate number of disks as suggested in the storage layout figure.

1. In the EMC Unisphere home page, choose **System > Storage Pools**.
2. In the Storage Pools window, click **Configure Disks**.

Figure 111 EMC Unisphere – Storage Pools Configure Disks

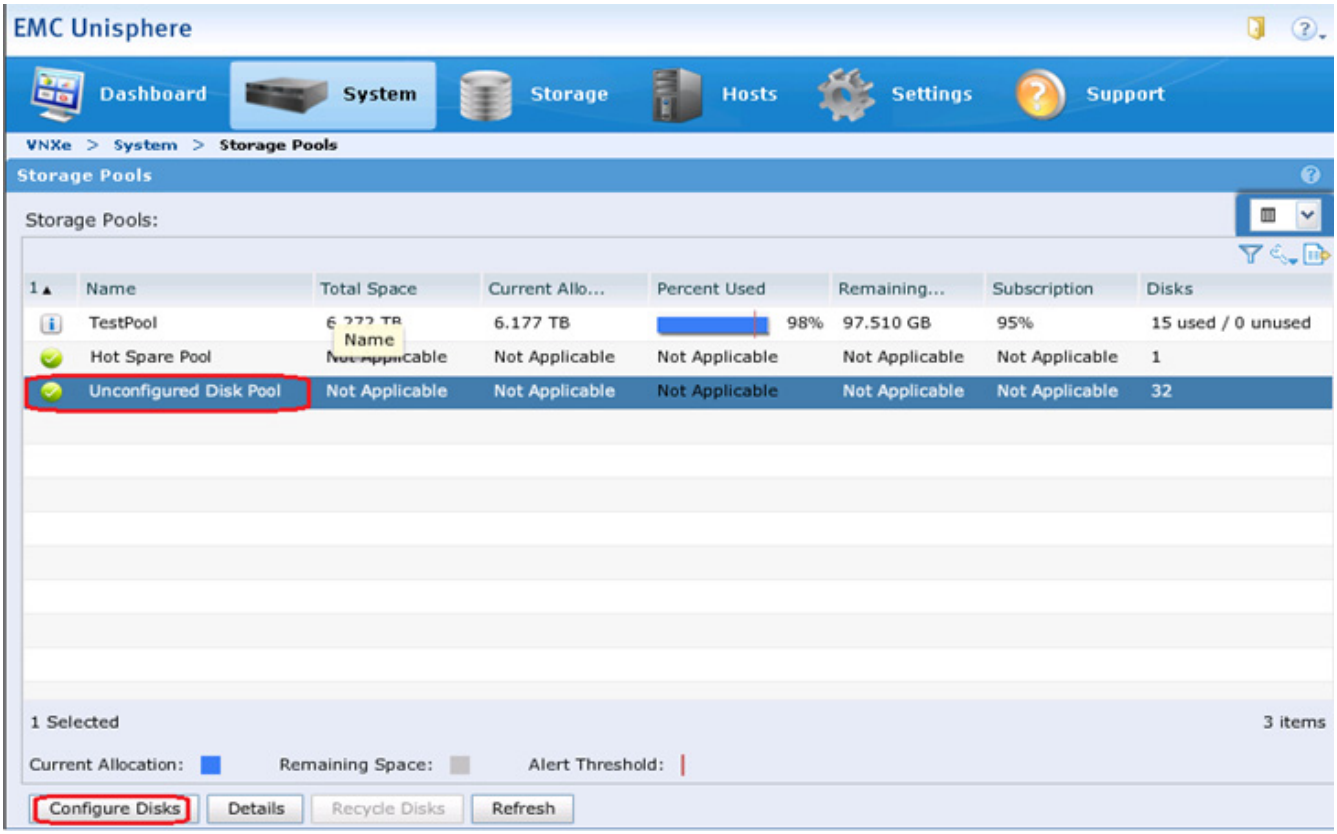

3. For Disk Configuration Mode, click the radio button **Manually create a new pool**.
4. For creating a new pool by disk type, choose Pool created for VMFS VMware Storage – Datastore from the drop-down list.

Figure 112 EMC Unisphere – Disk Configuration Wizard, Select Mode

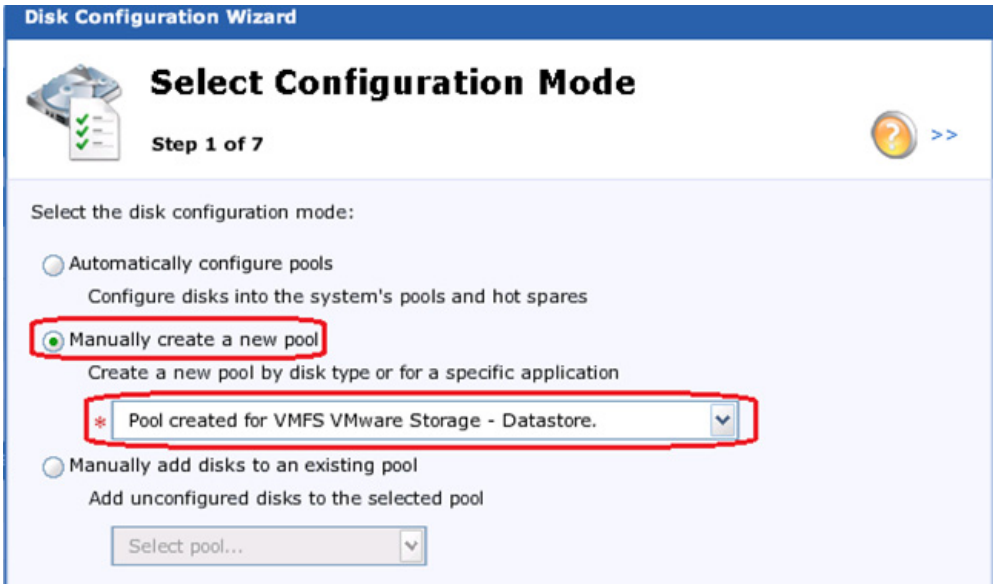

5. Specify Pool Name.

Figure 113 EMC Unisphere – Disk Configuration Wizard, Specify Pool Name

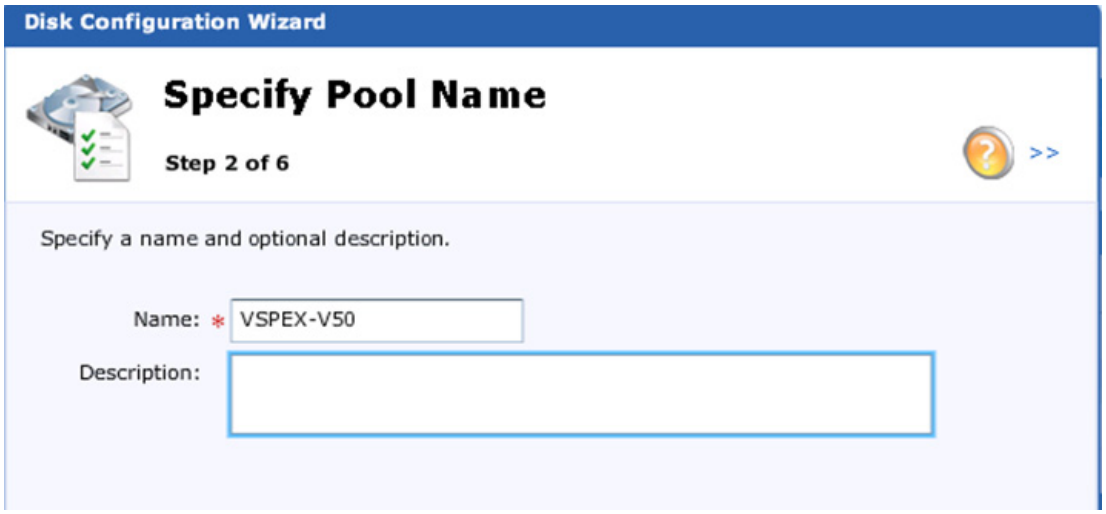

6. Select Balanced Perf/Capacity in the Select Storage Type window. The validated configuration uses a single pool with 45 drives.

Figure 114 EMC Unisphere – Disk Configuration Wizard, Select Storage Type

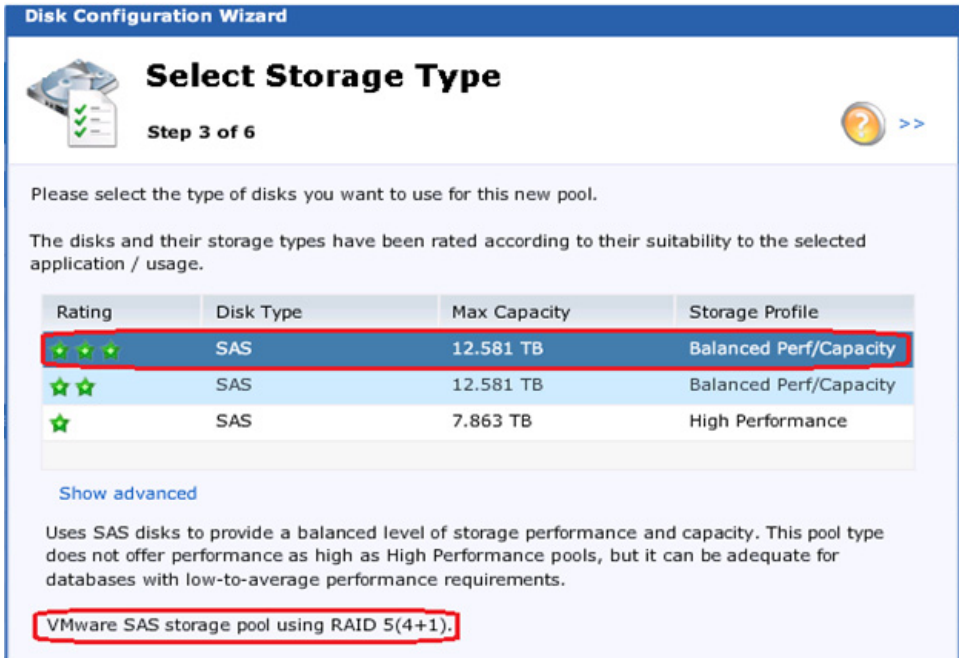

7. In the Select Amount of Storage window, for the field 600GB SAS (15000 RPM) Disks, choose Use 45 disks of 47 disks from the drop-down list and click **Next**.

Figure 115 EMC Unisphere – Disk Configuration Wizard, Select Amount of Storage

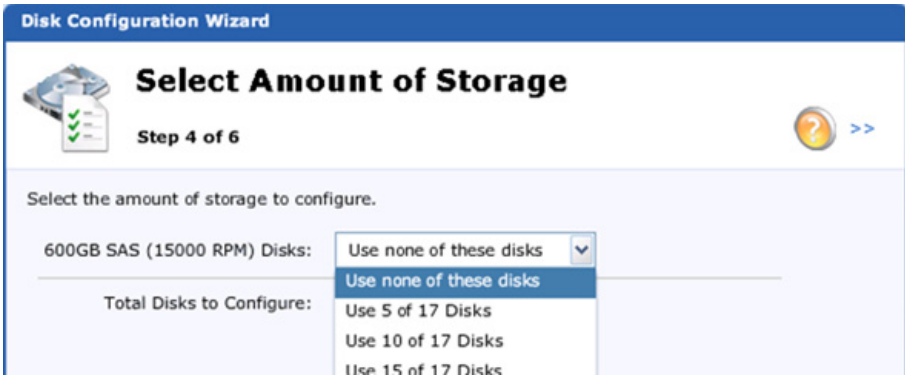

8. Review the summary and click **Finish**.

Note You should also create your Hot Spare disks at this point. As a performance best practice, all of the drives in the pool should be of the same size.

Configure Jumbo Frames

The Cisco networking environment will have a Maximum Transmission Unit (MTU) size of 9000 for the iSCSI connections to the VNXe. To match the configured MTU size via EMC Unisphere, follow these steps:

1. In the EMC Unisphere home page, choose **Settings > More Configuration > Advanced Configuration**.
2. Select eth10 and set the MTU size to 9000.
3. Select eth11 and set the MTU size to 9000.

Figure 116 EMC Unisphere – Advanced Configuration Jumbo Frames

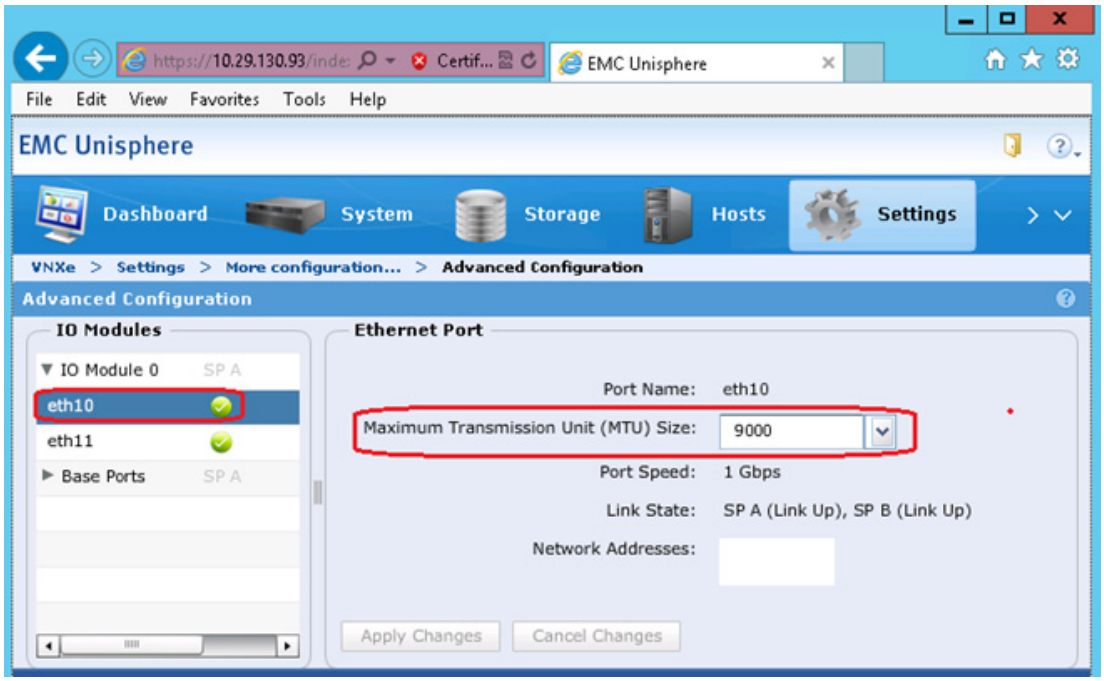

Create VMware Datastores

There are two options to create datastores in the EMC VNXe storage array. We have used generic iSCSI storage option to create the virtual disks and present them to the ESXi hosts. However, you can create VMware/VMFS datastores using VMware option in the storage section of EMC Unisphere.

Note Depending on the VNXe code revision, the datastores created by the VNXe vSphere engine may create the datastores as VMFS version 3.x, instead of VMFS version 5.x.

To create an iSCSI Virtual disk on iSCSI storage on a VNXe platform, follow these steps:

1. In the EMC Unisphere home page, choose **Storage > Generic iSCSI Storage**.

Figure 117 EMC Unisphere – Generic iSCSI Storage

2. Click **Create** in the Generic iSCSI Storage window.

Figure 118 EMC Unisphere – Generic iSCSI Storage, Create

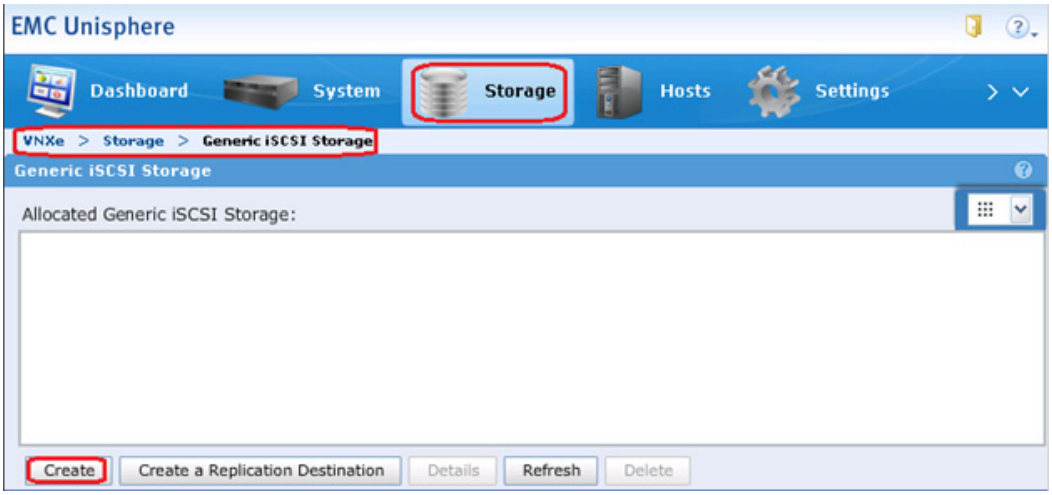

3. In the Generic iSCSI Storage, enter a Name for the generic storage resource and click **Next**.

Figure 119 EMC Unisphere – Generic iSCSI Storage Wizard, Specify Name

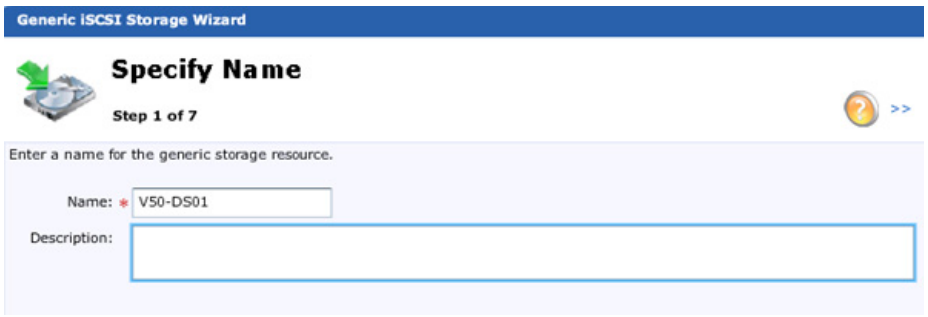

4. For Storage Server field, choose iSCSIServerA (SPA) and enter the required storage size in the Size field.

Figure 120 EMC Unisphere – Generic iSCSI Storage Wizard, Configure Storage

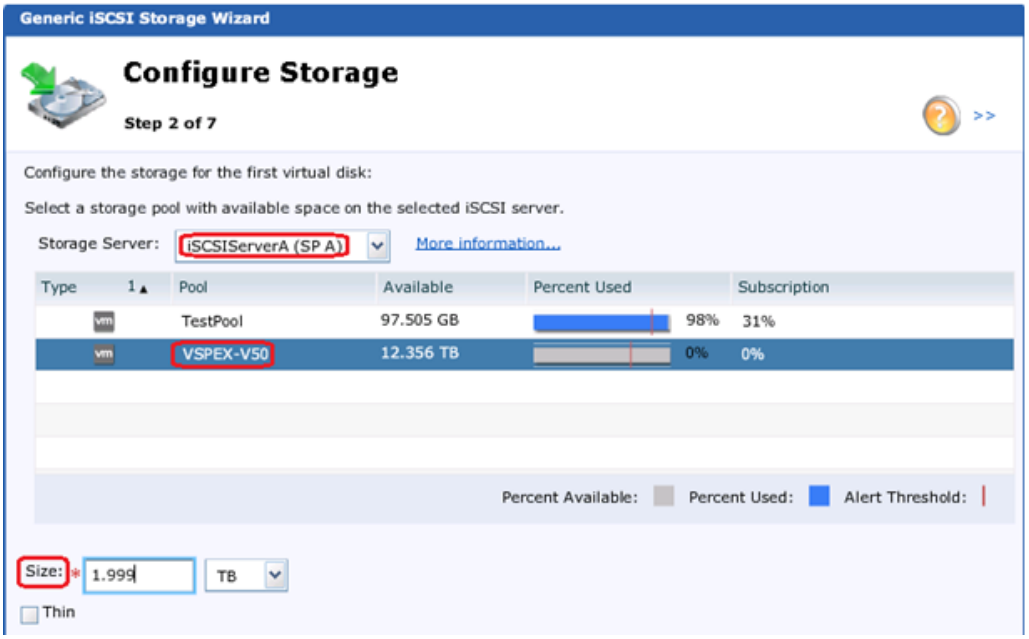

5. Change protection by clicking the radio button **Do not configure protection storage for this storage resource**. If you need additional protection, then additional storage is required. For more details on opting additional storage, see *EMC VNXe Storage Configuration Guide*.

Figure 121 EMC Unisphere – Generic iSCSI Storage Wizard, Configure Protection

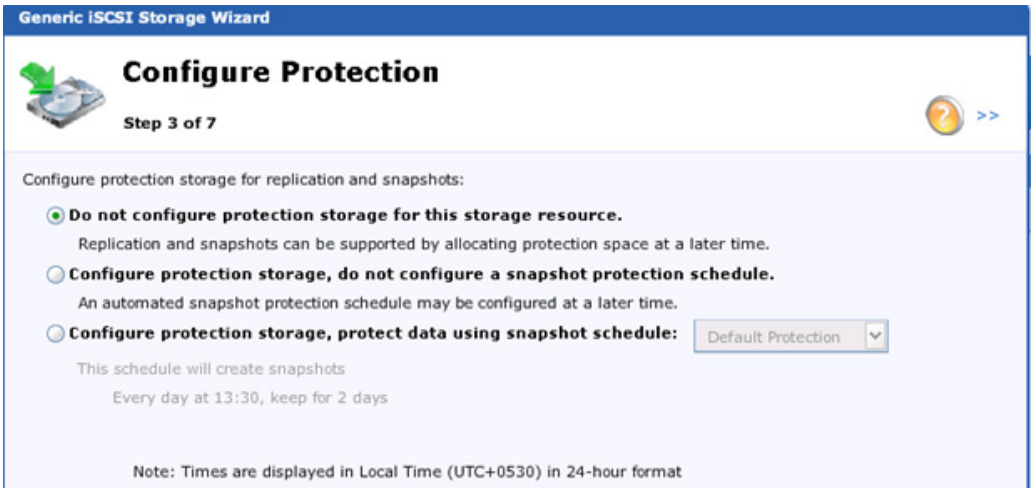

6. In the Configure Host Access window, for Access, choose Virtual Disk from the drop-down list. This is required because only the access level can permit the host to access the storage and not the snapshots. VNXe provides four levels of access for generic iSCSI storage as shown in Figure 122.

Figure 122 EMC Unisphere – Generic iSCSI Storage Wizard, Configure Host Access

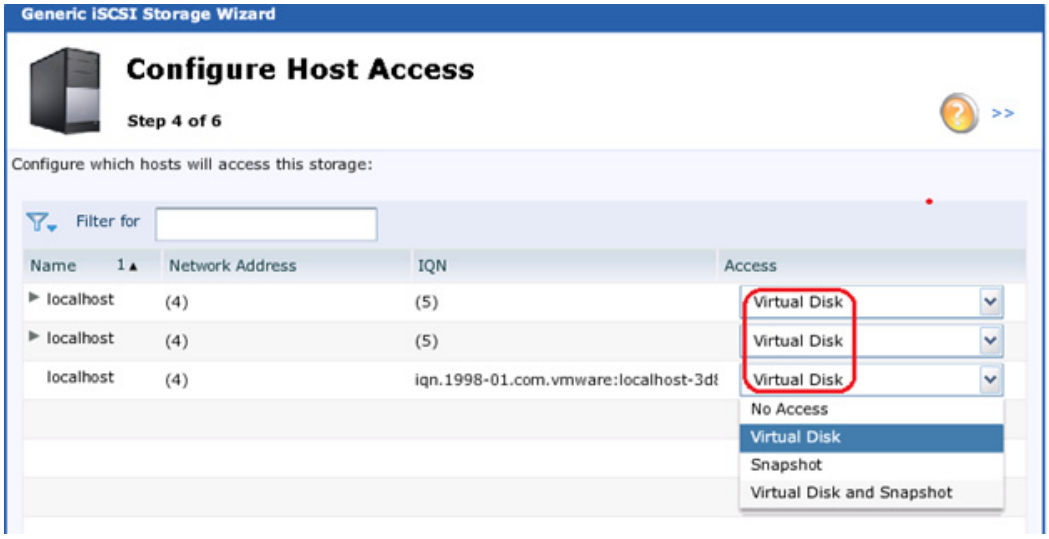

7. Click **Next** and Click **Finish** in the Summary window after verifying the details.
8. In the Results page, click **Add virtual Disk** radio button to create more iSCSI virtual disks or click **None** radio button to exit.

Figure 123 EMC Unisphere – Generic iSCSI Storage Wizard, Results

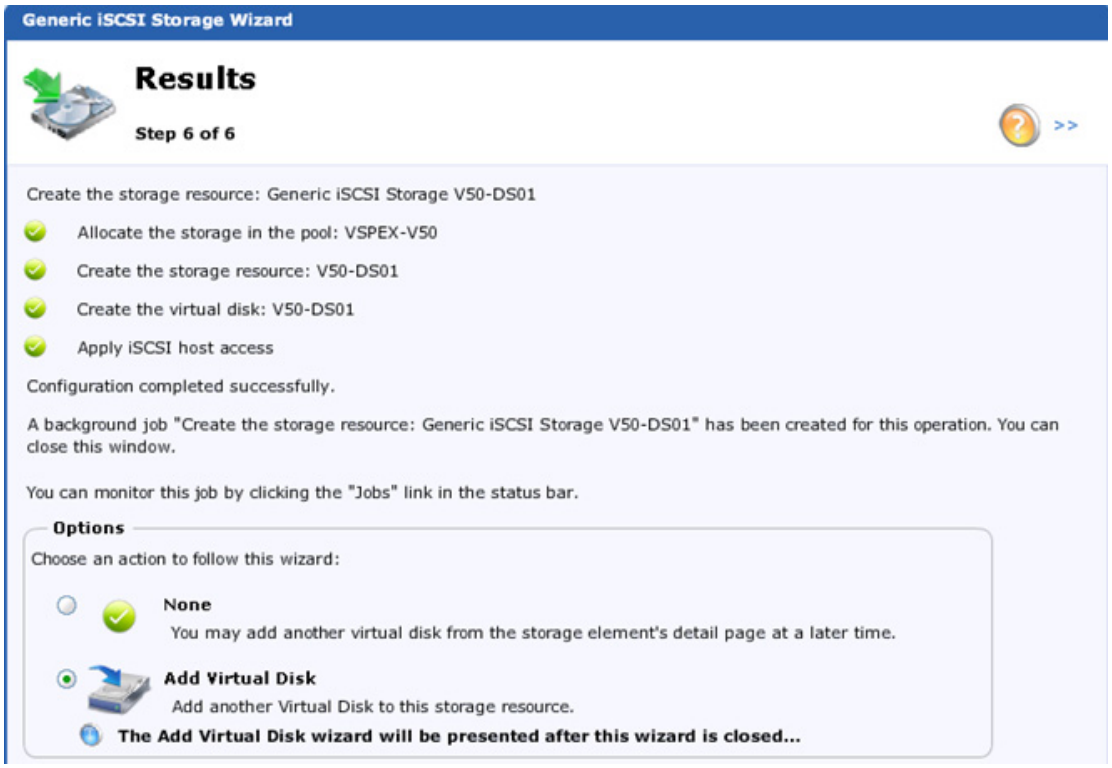

Configuring Discover Addresses for iSCSI Adapters

We now have to setup target discovery addresses so that the iSCSI adapter can determine the shared storage resources configured and made available for access to the ESXi hosts on the VNXe3150 storage array in the above section. The ESXi system supports Dynamic and Static Discovery methods. With dynamic discovery, all targets associated with an IP address or host name and the iSCSI name are discovered. With static discovery, you must specify the IP address or host name and the iSCSI name of the target you want to access. The iSCSI HBA must be in the same VLAN as both ports of the iSCSI array.

1. Login to the vSphere Client, and select a host from the inventory panel.
2. Click the **Configuration** tab and click **Storage Adapters** in the Hardware panel.
3. Select and right-click vmhba under the iSCSI Software Adapter and click **Properties**.

Figure 124 VMware vCenter – Storage Adapter Configuration

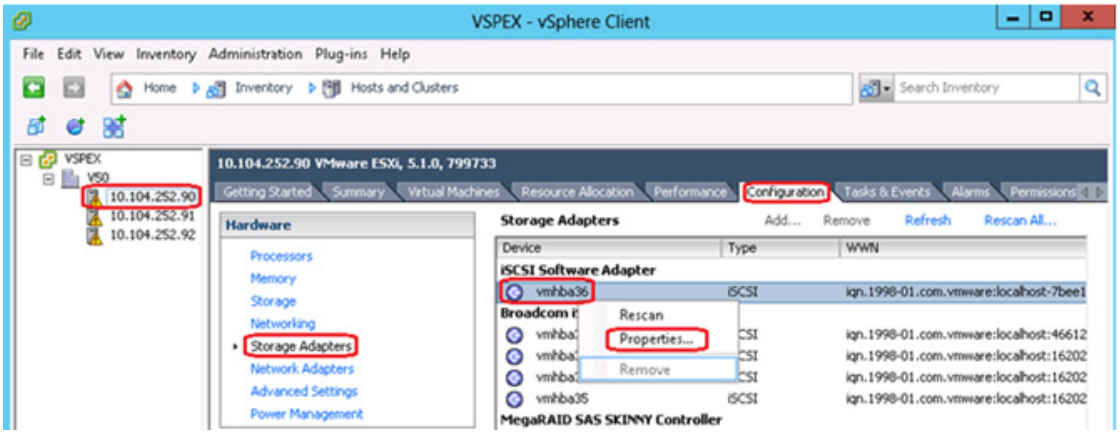

4. Click the **Dynamic Discovery** tab.
5. To add an address for the Send Targets discovery, click **Add**.

Figure 125 iSCSI Initiator Properties

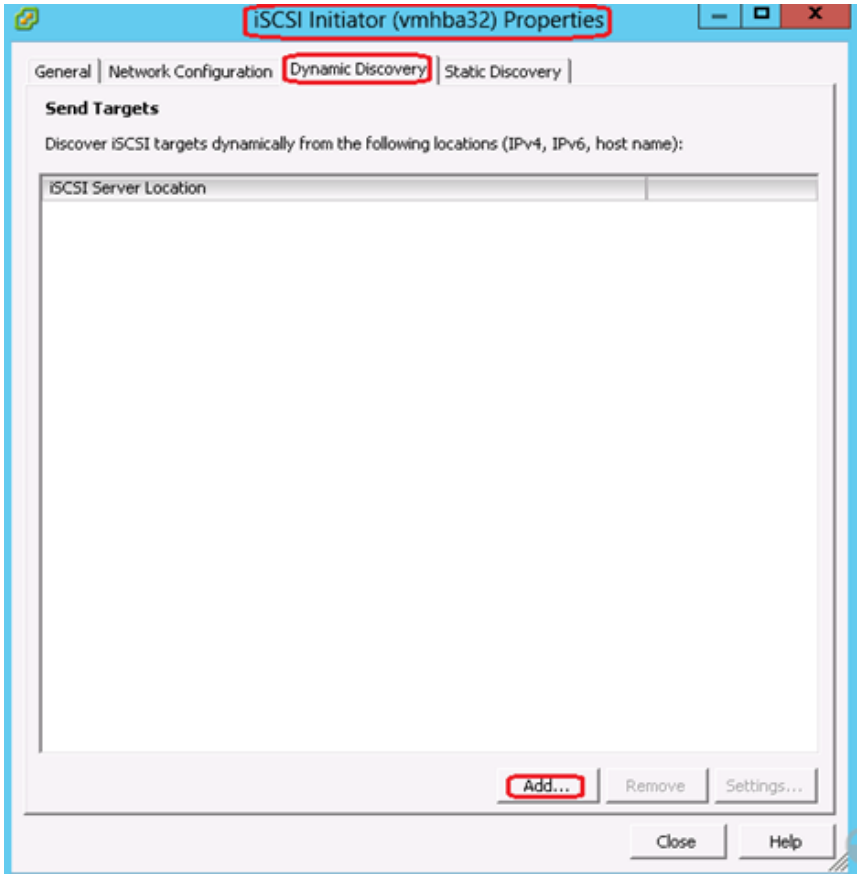

6. Enter the IP Address of the iSCSI Server (VNXe SP-A eth10) in the Add Send Target Server window. Leave the port number to default 3260 and click **OK**.

Figure 126 Add Send Target Server

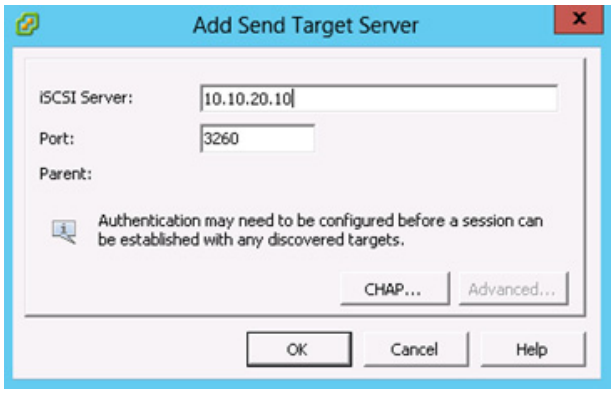

7. Repeat the steps 4 and 5 to add all the other Send Target portal IP addresses of the VNXe storage array.

Figure 127 iSCSI Initiator Properties – Dynamic Discovery

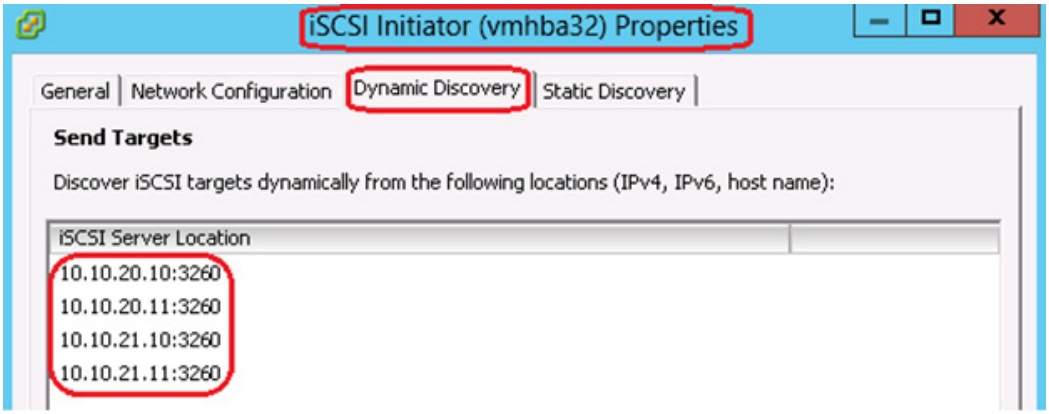

8. After your host establishes the Send Targets session with this system, any newly discovered targets appear in the Static Discovery list.

Figure 128 iSCSI Initiator Properties – Static Discovery

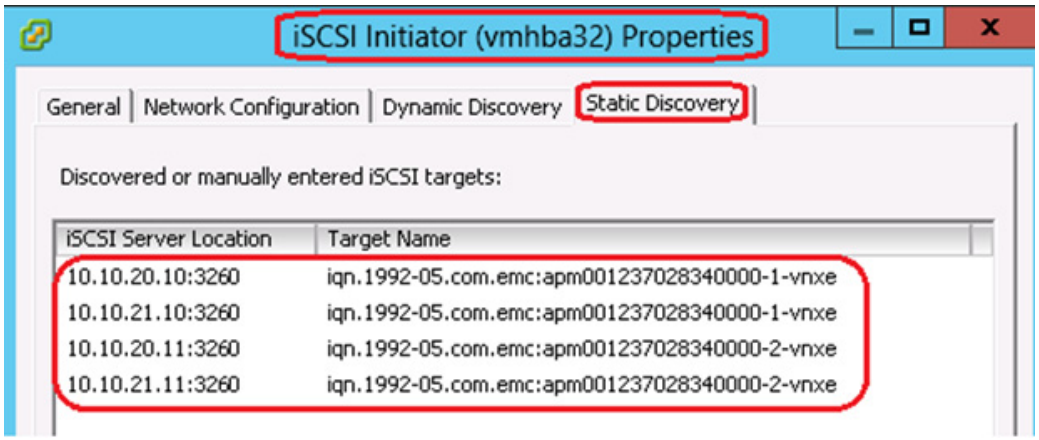

9. After configuring Dynamic Discovery for your iSCSI adapter, rescan the adapter.
10. Select a host in the inventory panel and click **Configuration > Storage > Add Storage**.

Figure 129 VMware vCenter – Add Storage

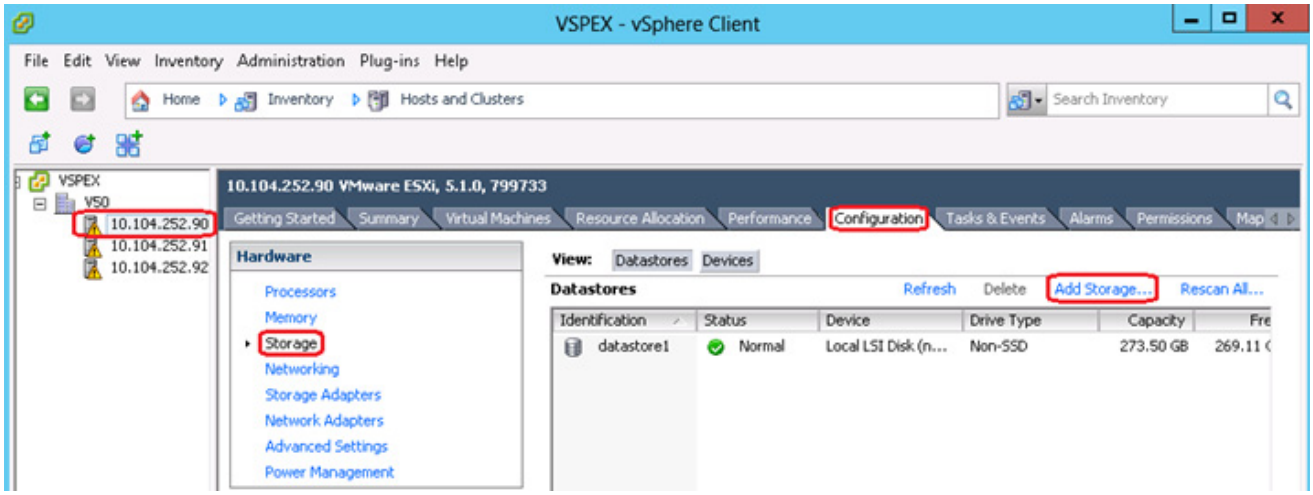

11. In the Add Storage wizard, click the radio button **Disk/LUN** in the Storage Type area and click **Next**.

Figure 130 Add Storage – Disk/LUN

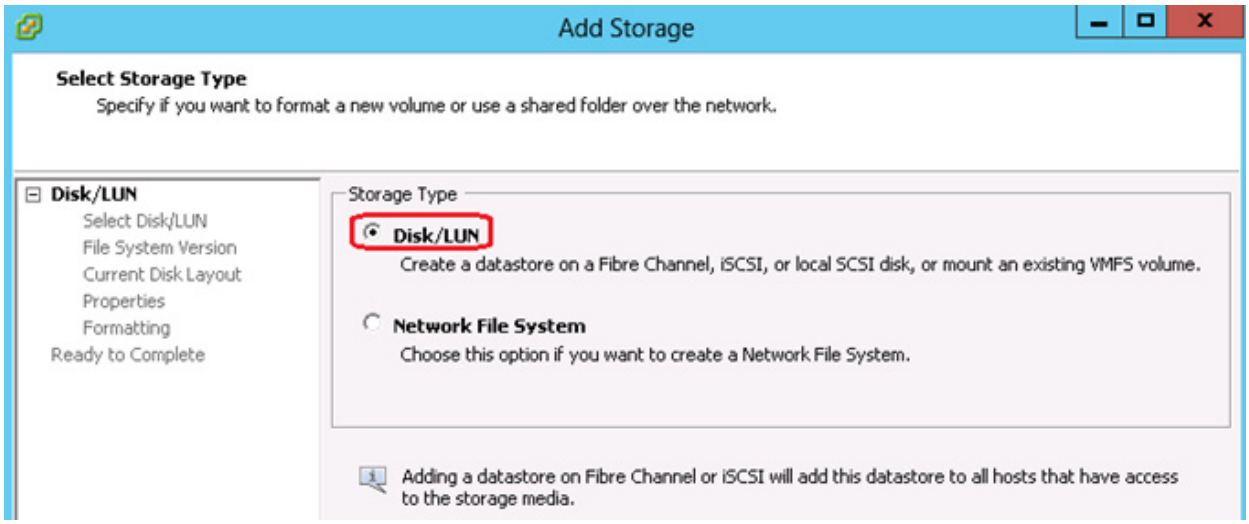

12. Select a LUN from the list to create a datastore and click **Next**.

Figure 131 Add Storage – Select Disk/LUN

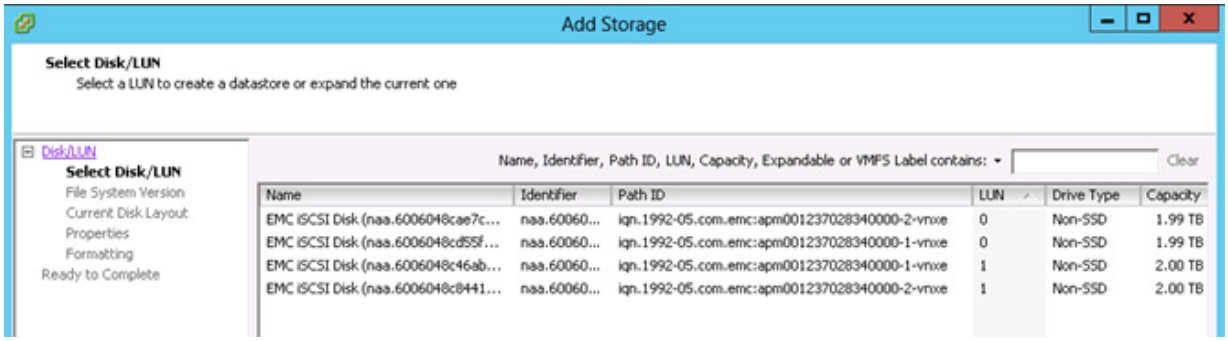

13. Next select VMFS-5 as the File System Version in the Add Storage wizard and click **Next**.

Figure 132 Add Storage – File System Version

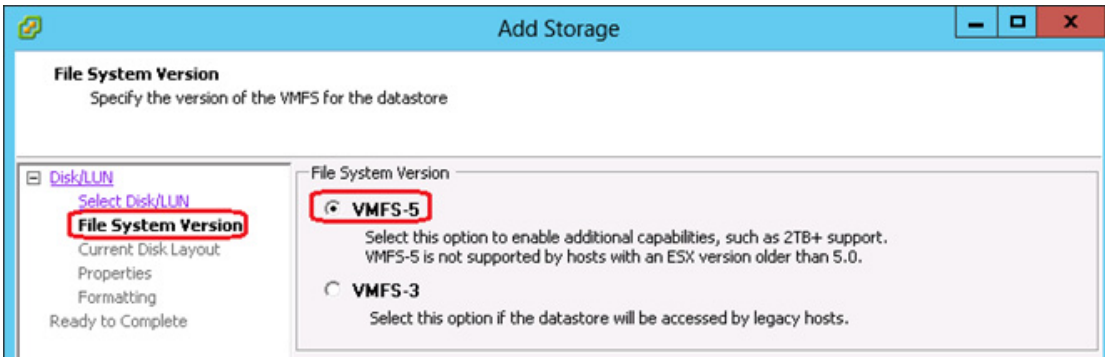

14. Review the current disk layout and click **Next** to create a partition.

Figure 133 Add Storage – Current Disk Layout

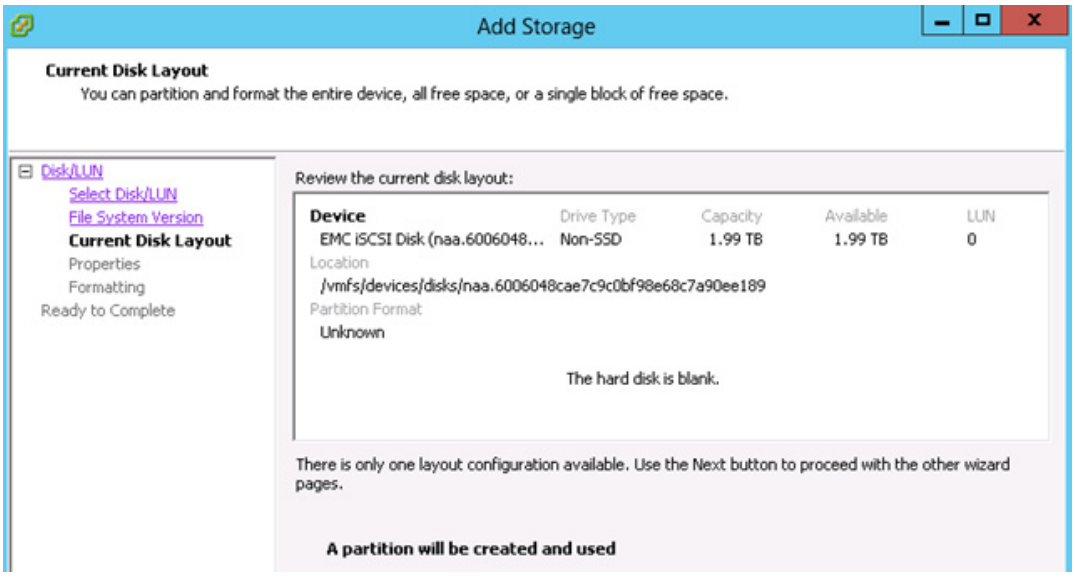

15. In the Properties window of the Add Storage wizard, enter a name for the datastore and click **Next**.

Figure 134 Add Storage – Properties

16. In the Formatting window of Add Storage wizard, specify the file size and capacity of the datastore for formatting and click **Next**.

Figure 135 Add Storage – Formatting

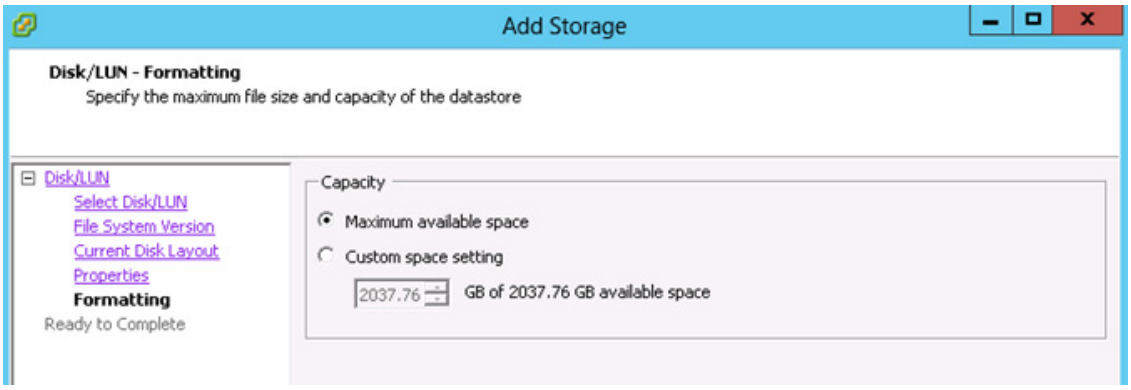

17. In the Ready to Complete window of Add Storage wizard review the disk layout and click **Finish** to add storage.

Figure 136 Add Storage – Ready to Complete

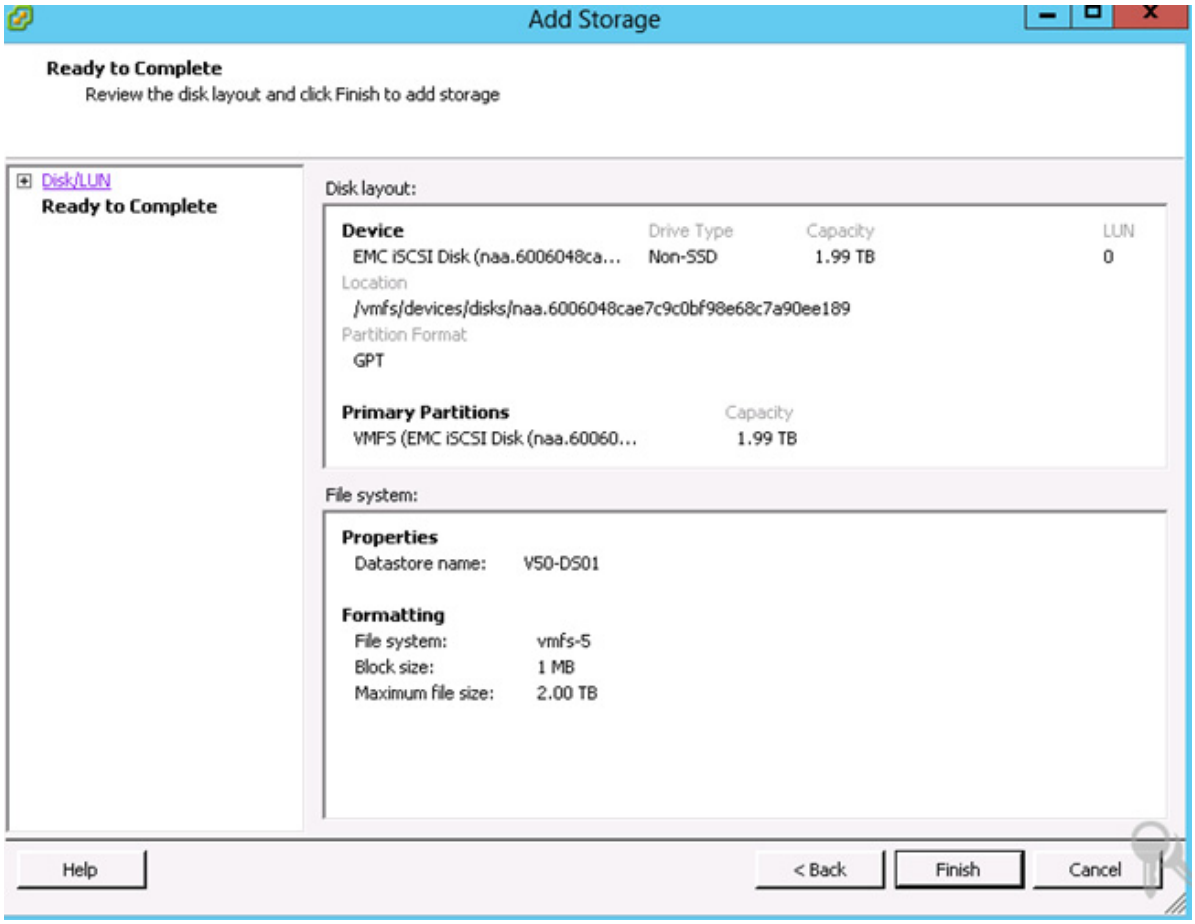

18. Repeat the steps 8 to 15 to add other Disk/LUNs.

Figure 137 VMware vCenter – Storage View

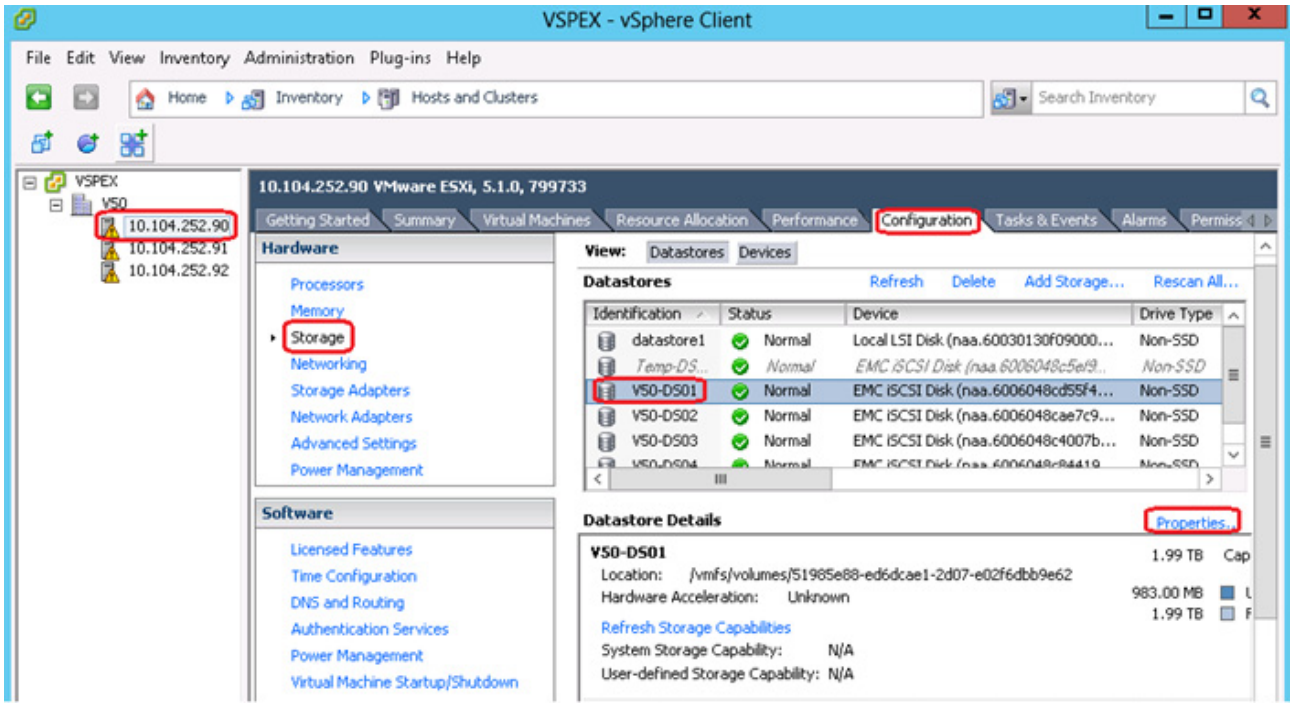

Configuring vSphere HA and DRS

vSphere High Availability (HA) and Distributed Resource Scheduler (DRS) provide protection for virtual machines (VMs) running on ESXi hosts, as well as optimize resources and performance and simplify VM management. The HA feature continuously monitors all ESX Server hosts in a cluster and detects failures, and will automatically restart VMs on other host servers in an ESX cluster in case of a host failure. DRS aggregates vSphere host resources into clusters and automatically distributes these resources to virtual machines by monitoring utilization and continuously optimizing virtual machine distribution across vSphere hosts. vCenter Server is required for the use of this feature.

The section explains steps to create cluster and configure HA and DRS for the VSPEX V50 solution:

1. Login to the vCenter using vSphere Client, and choose **Datacenter** from the inventory panel.
2. 2. Right-click and choose **New Cluster** to launch the New Cluster Wizard.

Figure 138 VMware vCenter – New Cluster

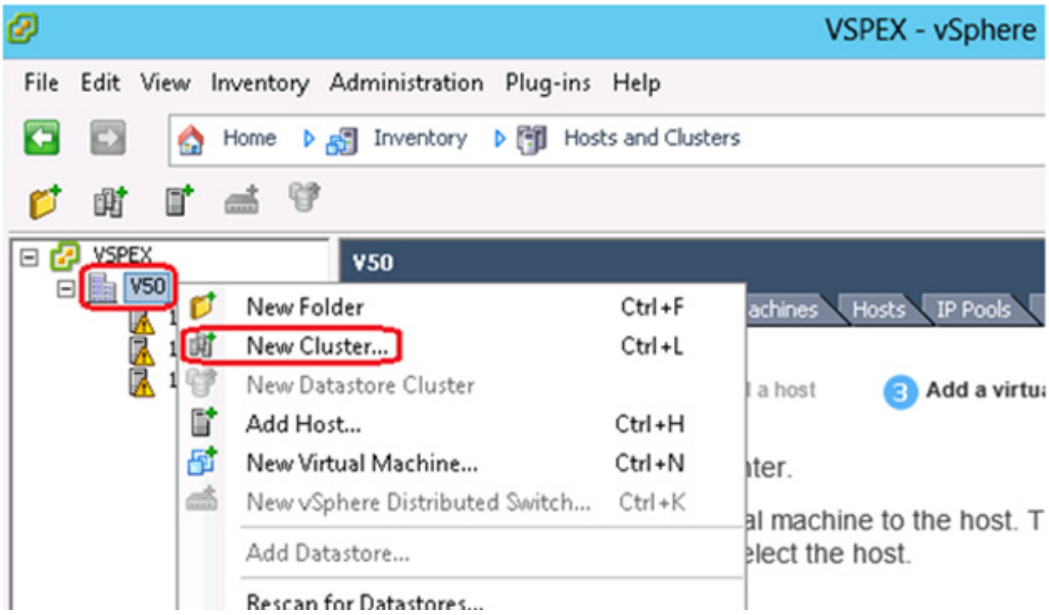

3. In the Cluster Features window, select Turn ON vSphere HA and Turn On vSphere DRS.

Figure 139 New Cluster Wizard – Cluster Features

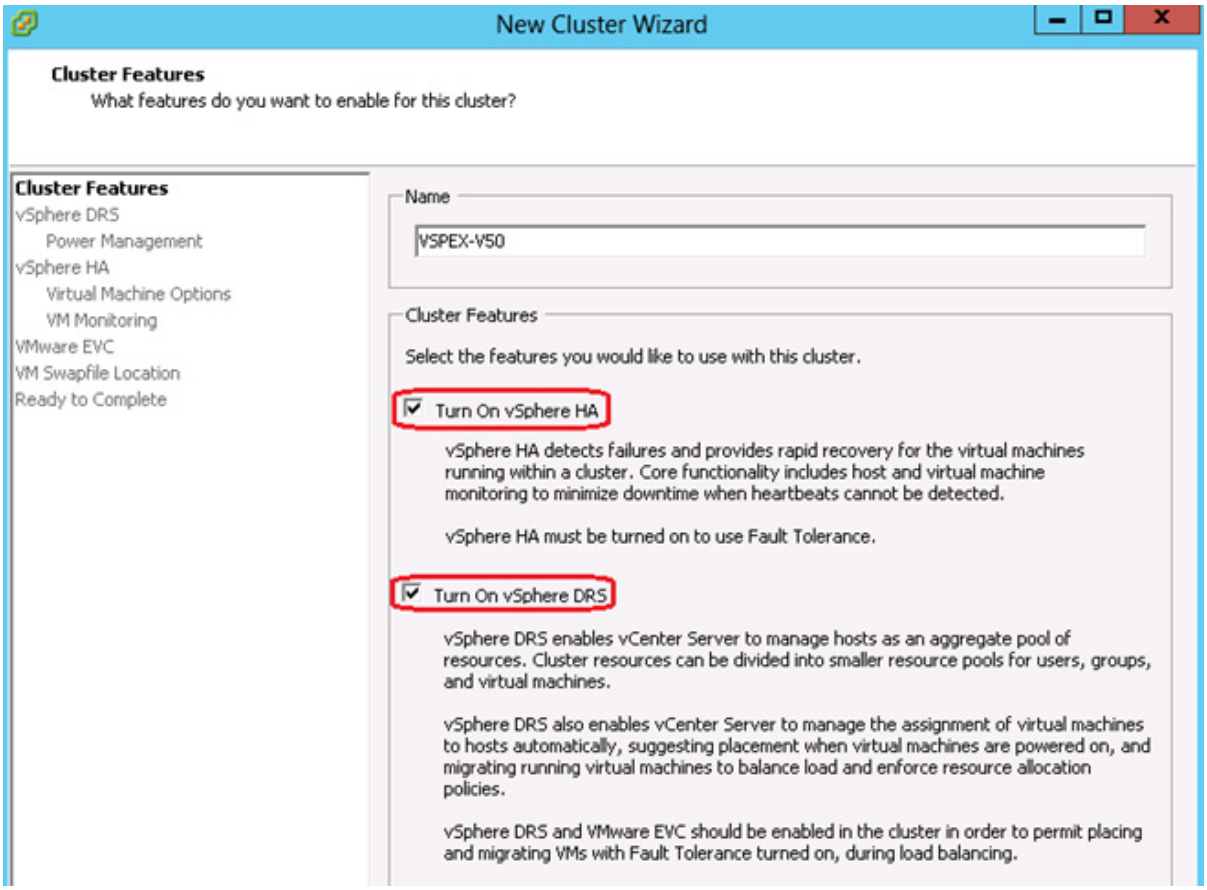

4. Click **Next** and in the subsequent steps choose the options suitable for your environment to complete the configuration of HA and DRS.

Figure 140 New Cluster Wizard – Ready to Complete

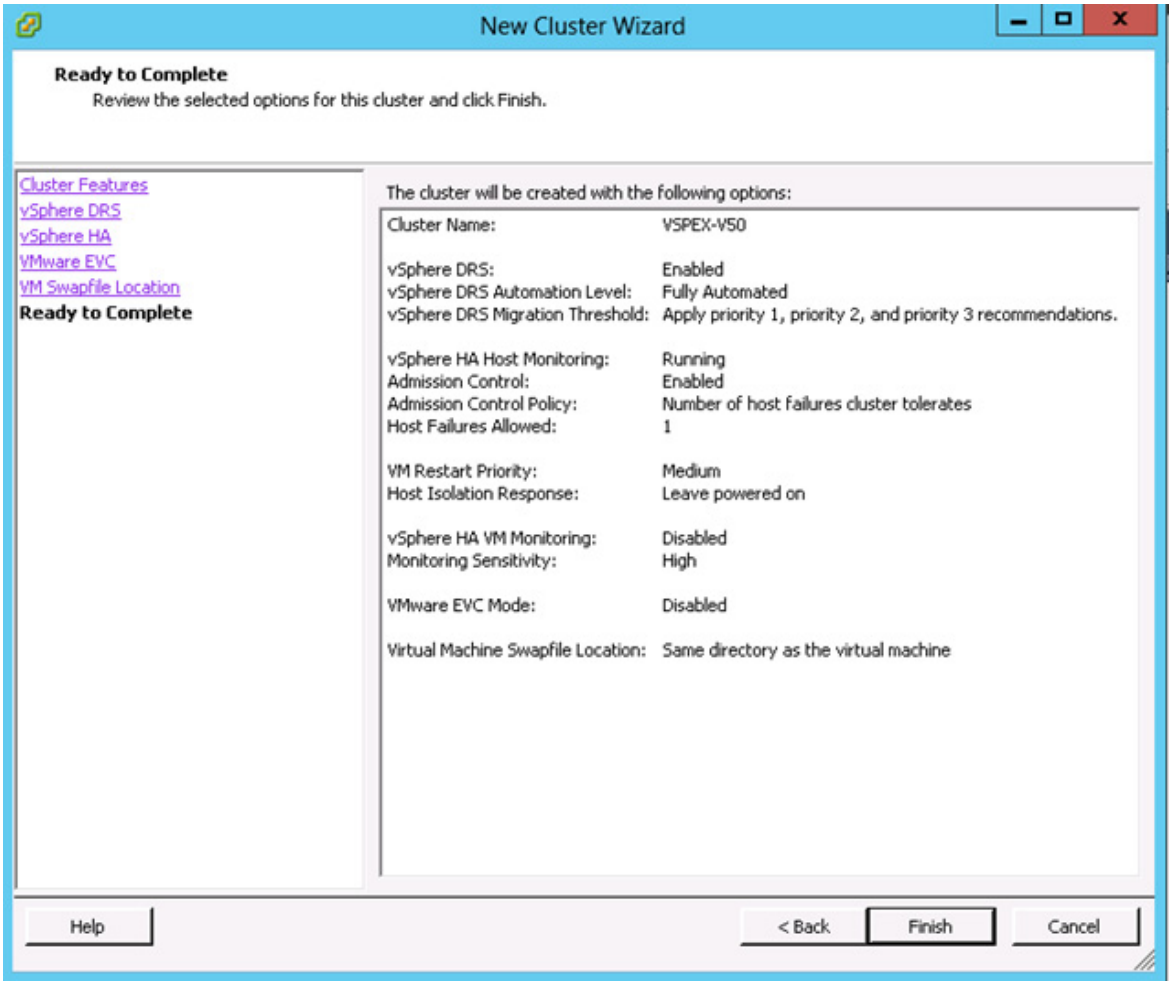

5. Right-click the Cluster and choose **Add Hosts**.

Figure 141 VMware vCenter – Add Host to Cluster

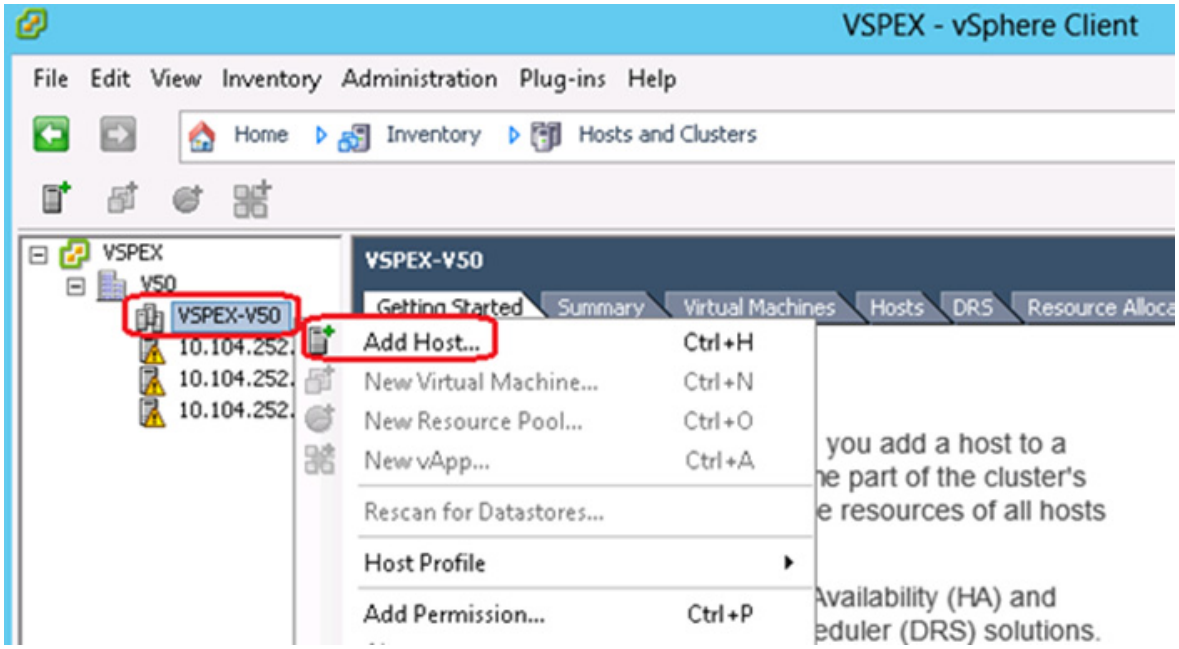

6. In Specify Connection Settings enter the name or IP address of the host to add to the cluster.

Figure 142 Add Host Wizard – Connection Settings

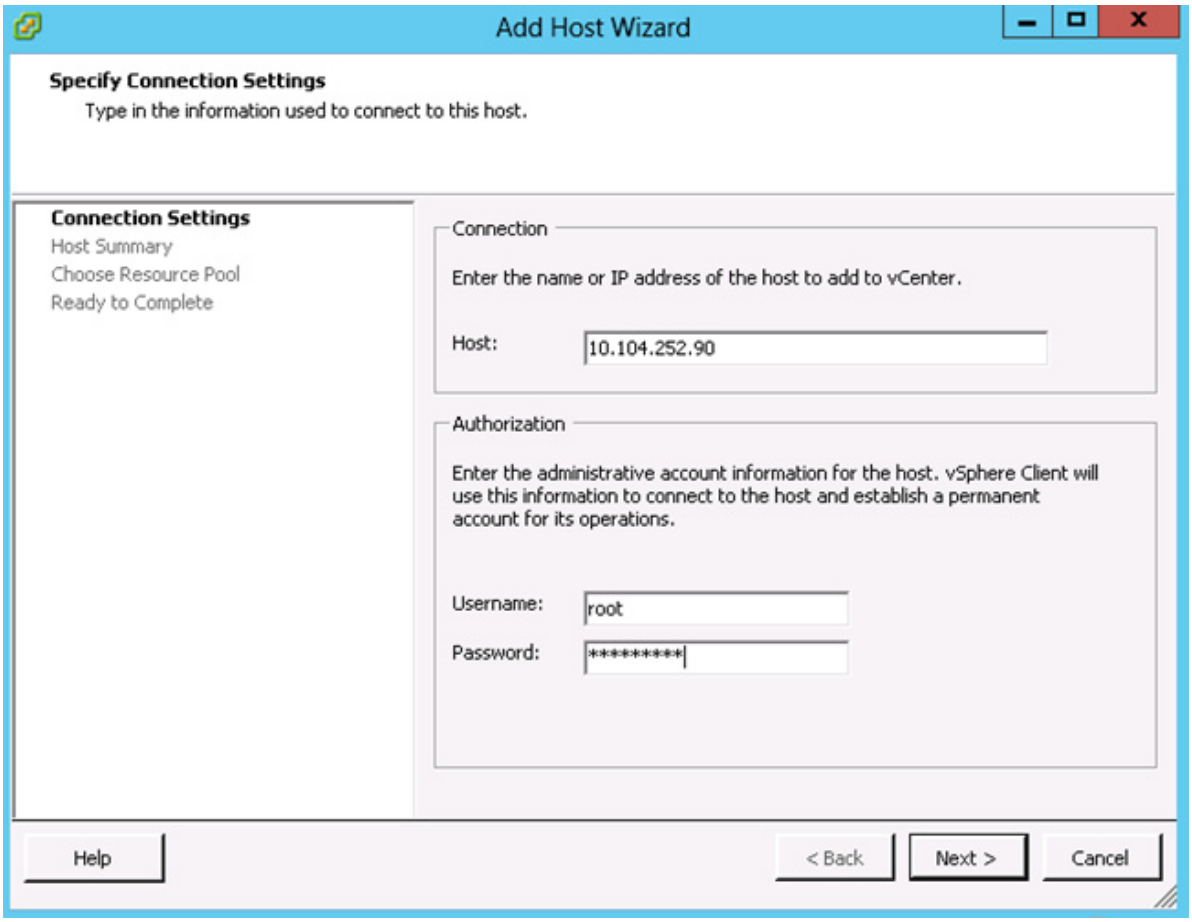

7. In the Security Alert dialog box click Yes to continue.
8. In the Host Summary page review the information and click Next.

Figure 143 Add Host Wizard – Host Summary

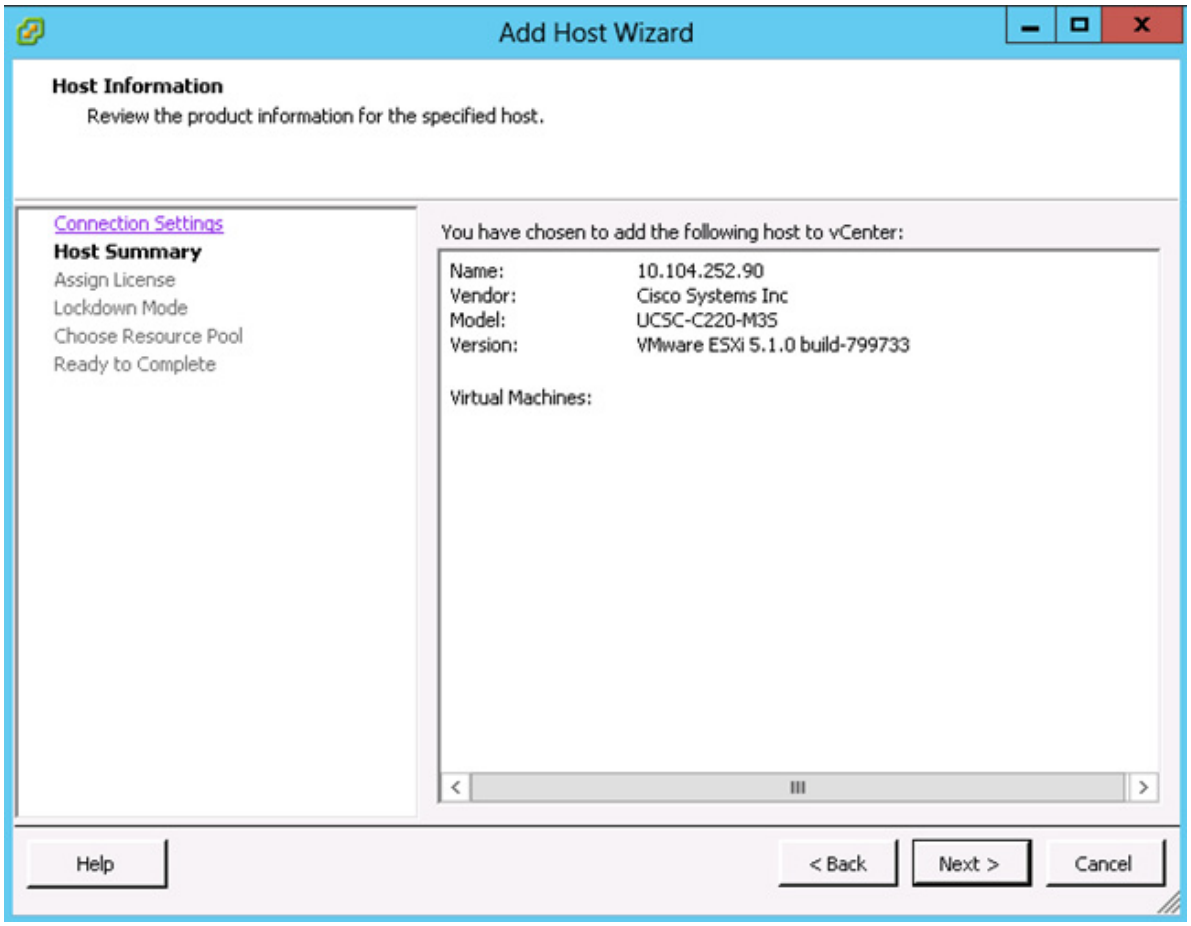

9. Assign an existing or a new key to the host in the Assign License page. You can choose evaluation and assign the key at a later stage.

Figure 144 Add Host Wizard – Assign License

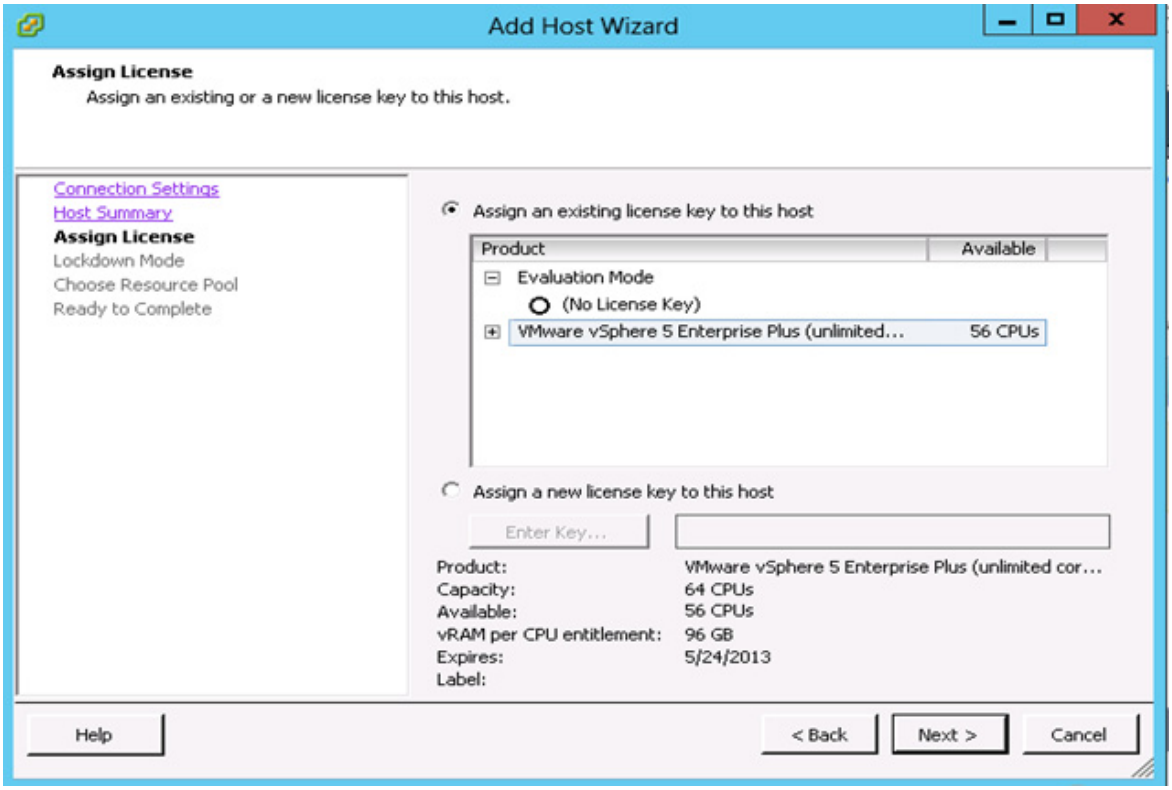

10. In the Lockdown Mode section you can specify whether to Enable Lockdown Mode for this host.

Figure 145 Add Host Wizard – Lockdown Mode

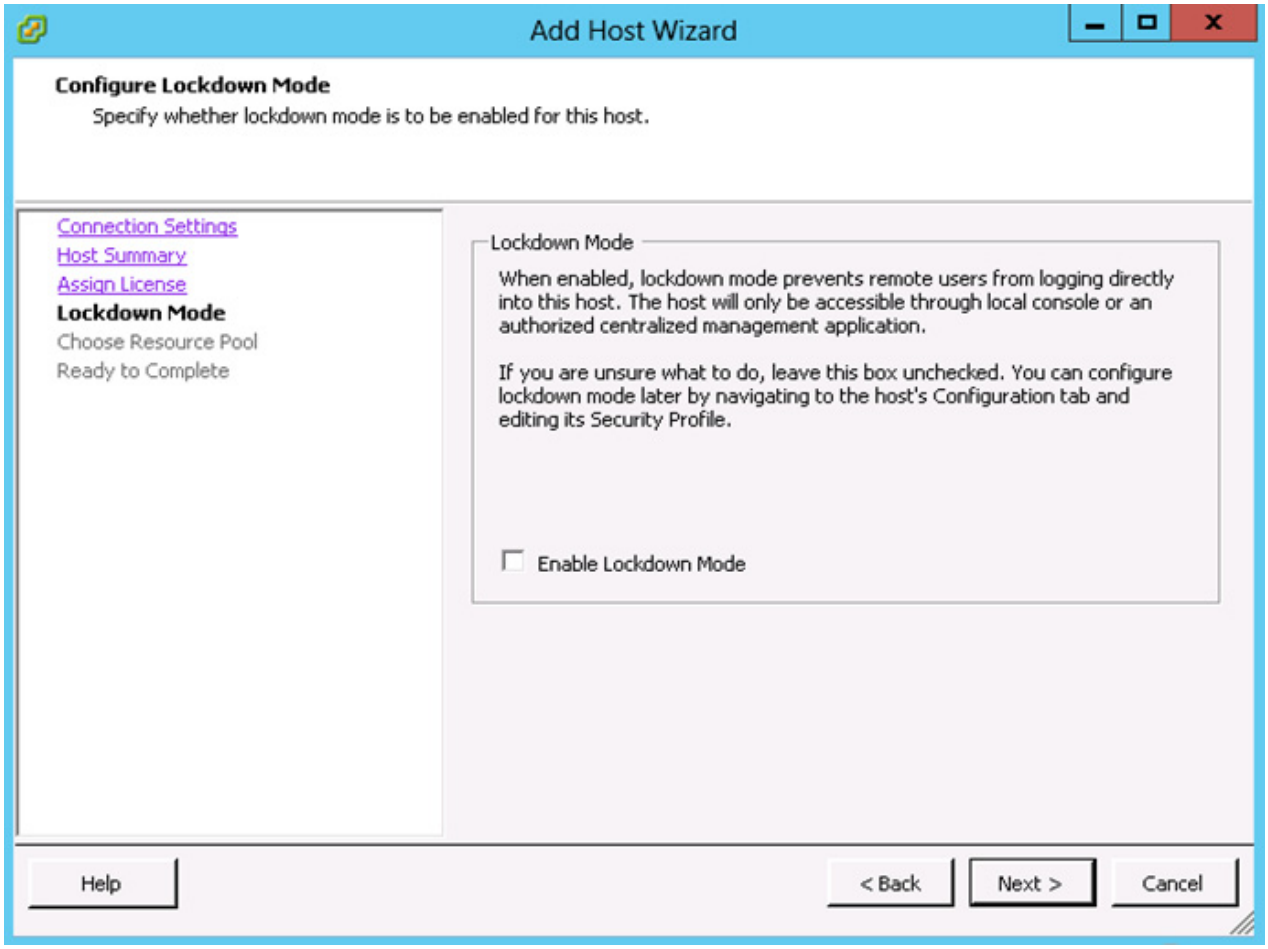

11. In the next window choose where to place this host's virtual machines in the resource pool hierarchy and click **Next**.

Figure 146 Add Host Wizard – Choose Resource Pool

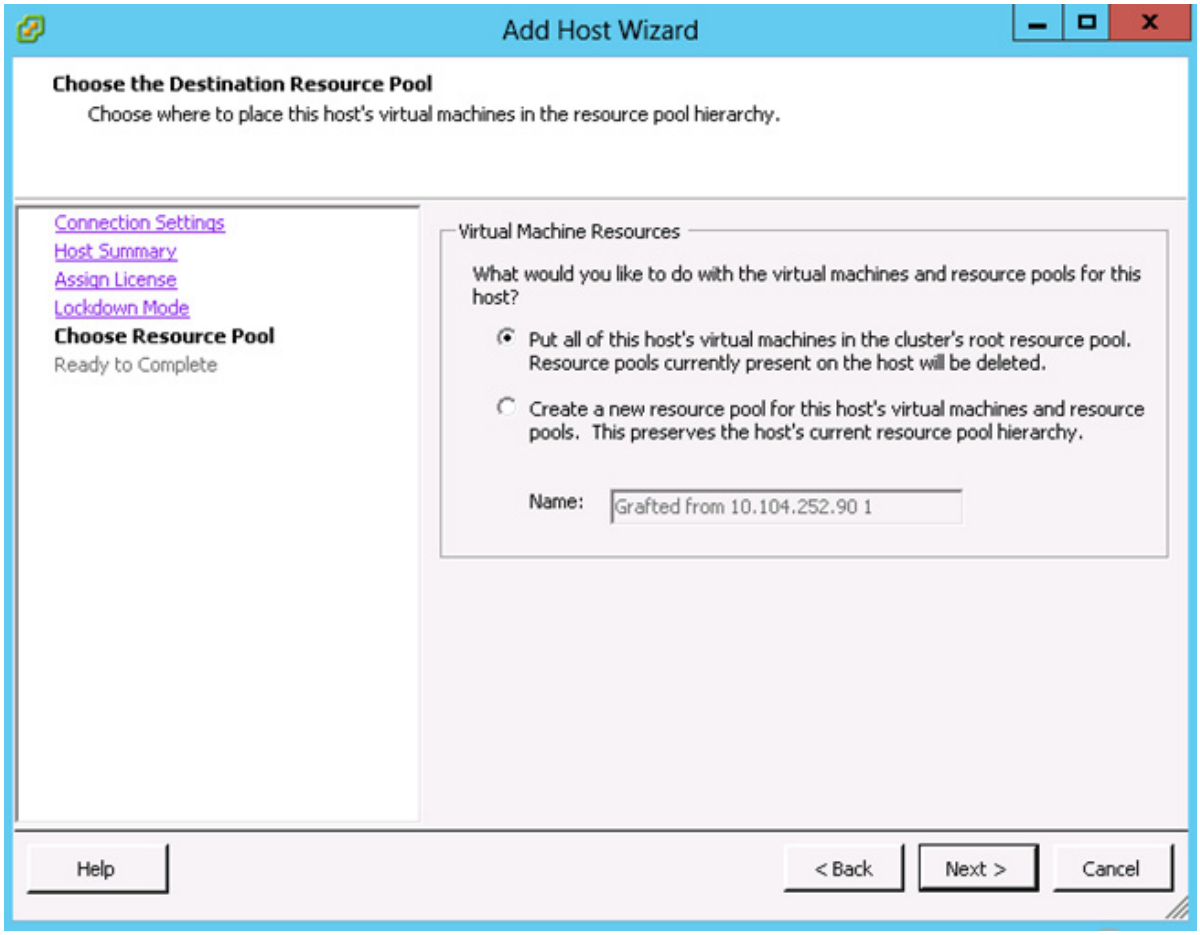

12. In the Ready to Complete window review the summary and click **Finish**.

Figure 147 Add Host Wizard – Ready to Complete

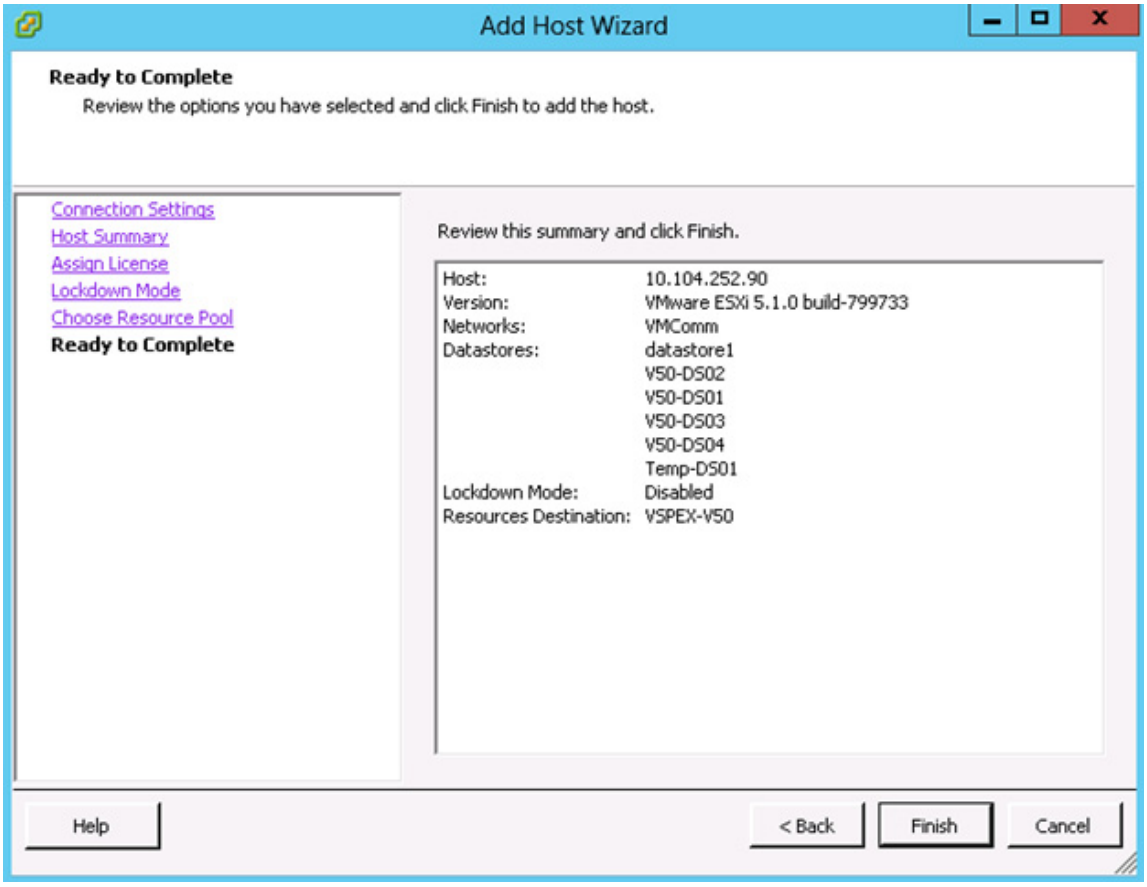

13. Repeat the steps 5 to 8 to add other ESXi hosts to the cluster.

Figure 148 VMware vCenter – Cluster View

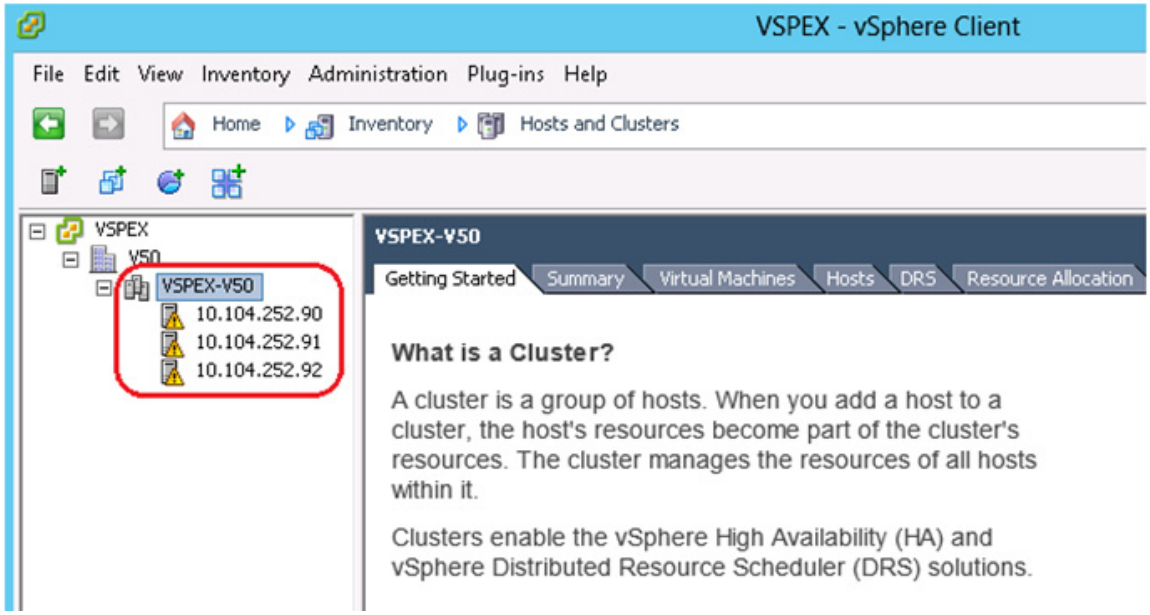

Template-Based Deployments for Rapid Provisioning

Figure 149 Rapid Provisioning Using VM Templates

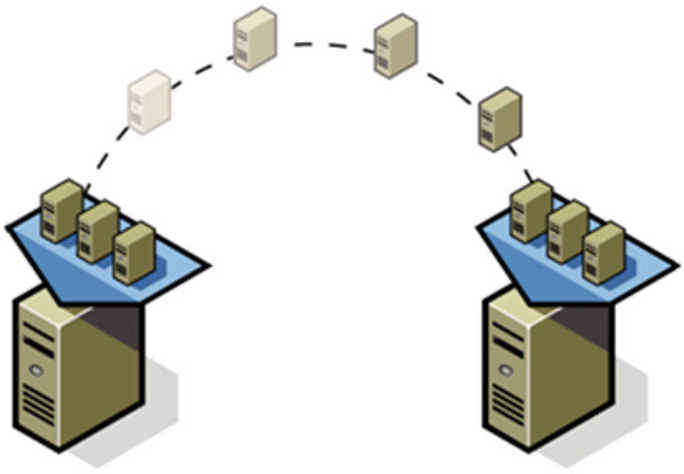

In an environment with established procedures, deploying new application servers can be streamlined, but can still take many hours or days to complete. Not only must you complete an OS installation, but downloading and installing service packs and security updates can add a significant amount of time.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

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

Make sure to spread VMs across different VM datastores to properly load-balance the storage usage.

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 the Configuration tab.

After connecting to the ESXi host through SSH, initiate ping to the iSCSI storage server and to vMotion VMkernel port of other ESXi hosts with large MTU size and set "Do Not Fragment" bit of IP packet to 1. Use the vmkping command as shown in Figure 150.

Figure 150 Validation of Jumbo Frames support

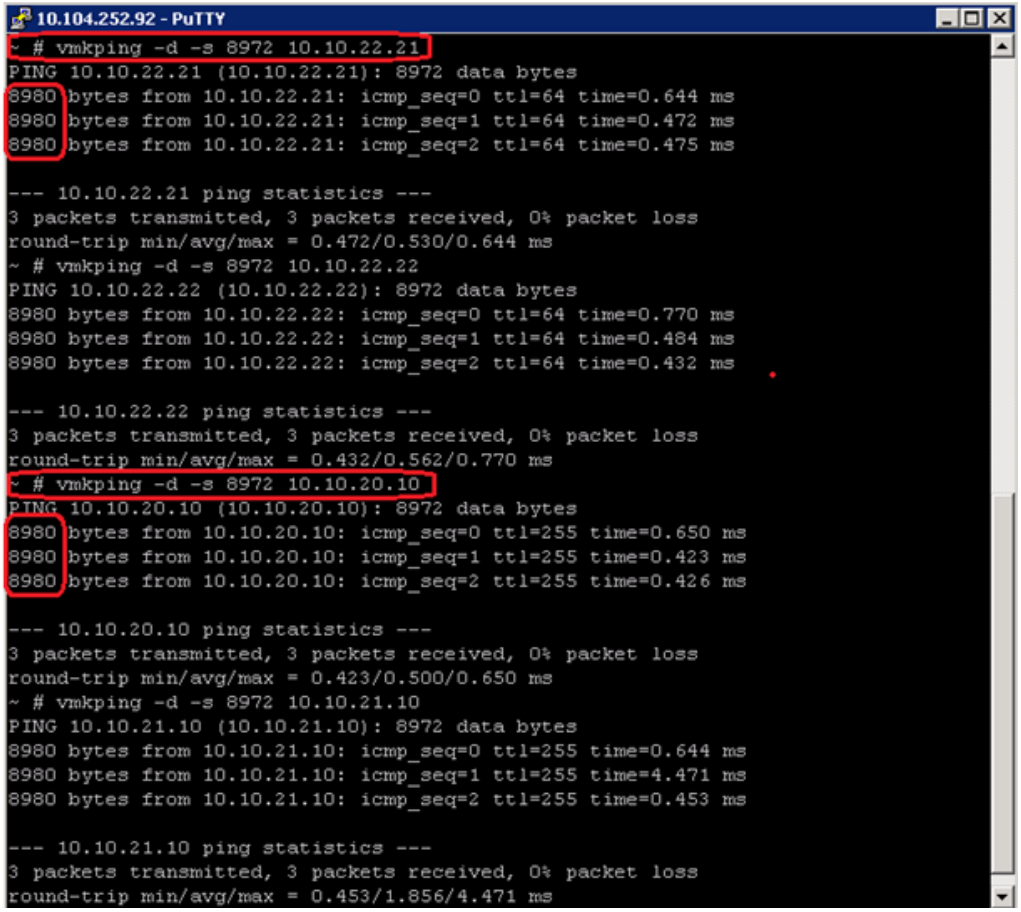

Ensure that the packet size is 8972 due to various L2/L3 overhead. Ping must be successful. If ping is not successful verify that 9000 MTU configured at each of the following steps:

1. 9000 MTU on all the Ethernet interfaces configured for iSCSI server on the VNXe storage device.

2. Make sure that a “jumbo-mtu” policy map is created at Nexus 3000 series switches with default class having MTU 9000. Make sure that “jumbo-mtu” policy is applied to system classes on the ingress traffic.
3. Make sure that 9000 MTU is set on the vSwitches as well as the VMkernel ports on them used configured for vMotion and iSCSI networking.

Validating Cisco Solution for EMC VSPEX with VMware Architectures

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

Post Install Checklist

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

- On each vSphere server, verify that the vSwitch that hosts the client VLANs has been configured with sufficient ports to accommodate the maximum number of virtual machines it may host.
- On each vSphere server used as part of this solution, verify that all required virtual machine port-groups have been configured and that each server has access to the required VMware datastores.
- On each vSphere server used in the solution, verify that an interface is configured correctly for vMotion/iSCSI 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 VMs. Upon rebooting the switch, the network access load should be rebalanced across the two switches.

4. Reboot the active storage processor of the VNXe storage array and make sure that all the iSCSI targets are still accessible during and after the reboot of the storage processor.
5. Fully load all the virtual machines of the solution. Shutdown one of the ESXi nodes in the cluster. All the VMs running on that host should be migrated to other active hosts. No VM should lose any network or storage accessibility during or after the migration. Note that in 50 virtual machines architectures, there is enough head room for memory in other servers to accommodate 25 additional virtual machines.

Cisco validation test profile

"vdbench" testing tool was used with Windows Server 2012 to test scaling of the solution in Cisco labs. Figure 150 provides details on the test profile used.

Table 10 VDBench Details

Profile Characteristics	Value
Number of virtual machines	50
Virtual machine OS	Windows Server 2010
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 machines	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 Material

Table 11 gives details of the components used in the CVD for 50 virtual machines configuration.

Table 11 Component Description

Description	Part #
Cisco UCS C220 M3 rack servers	UCSC-C220-M3S
CPU for Cisco UCS C220 M3 rack servers	UCS-CPU-E5-2650
Memory for Cisco UCS C220 M3 rack servers	UCS-MR-1X082RY-A
RAID local storage for rack servers	UCSC-RAID-11-C220
Broadcom 1Gbps adapter for 50 VMs solution	N2XX-ABPCI03-M3
Cisco 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 Cisco C220 M3 server specsheet at:

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

Customer Configuration Data Sheet

Before you start the configuration, gather some customer-specific network and host configuration information. Table 12, Table 13, Table 14, Table 15, Table 16, Table 17 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.

Table 12 Common Server Information

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

Table 13 VMware Server Information

Server Name	Purpose	Primary IP	Private Net (storage) addresses		VMkernel IP	vMotion IP
	ESXi Host 1					
	ESXi Host 2					
	ESXi Host 3					

Table 14 Array Information

Array name	
Admin account	
Management IP	
Storage pool name	

Table 14 Array Information (continued)

Array name	
Datastore name	
iSCSI Server IP (SP-A and SP-B)	

Table 15 Network Infrastructure Information

Name	Purpose	IP	Subnet Mask	Default Gateway
	Cisco Nexus 3048 Switch A			
	Cisco Nexus 3048 Switch B			

Table 16 VLAN Information

Name	Network Purpose	VLAN ID	Allowed Subnets
vlan-mgmt	Virtual Machine Networking ESXi Management		
vlan-iscsi-a	iSCSI Server		
vlan-iscsi-b	iSCSI Server		
vlan-vMotion	vMotion traffic network		
vlan-vmcomm (multiple)	Data VLAN of customer VMs as needed		

Table 17 Service Accounts

Account	Purpose	Password (optional, secure appropriately)
	Microsoft Windows Server administrator	
Root	ESXi root	
	Array administrator	
	vCenter administrator	
	SQL server administrator	
	Nexus 3048 administrator	

References

Cisco UCS:

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

Cisco Nexus:

http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html

EMC VNXe3xxx series resources:

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

VMware vSphere:

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

VMware vSphere 5.1 documentation:

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

Microsoft SQL Server installation guide

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