



# Cisco Solution for EMC VSPEX Microsoft Hyper-V Architectures

Design for 50 and 100 Virtual Machines


Last Updated: October 3, 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 has 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.



Tim Cerling

### **Tim Cerling, Technical Marketing Engineer, Datacenter Group, Cisco Systems**

Tim's focus is on delivering customer-driven solutions on Microsoft Hyper-V and System Center products. He has been in the IT business since 1979. He started working with Windows NT 3.5 on the DEC Alpha product line during his 19 year tenure with DEC, and he has continued working with Windows Server technologies since then with Compaq, Microsoft, and now Cisco. During his twelve years as a Windows Server specialist at Microsoft, he co-authored a book on Microsoft virtualization technologies – Mastering Microsoft Virtualization. Tim holds a BA in Computer Science from the University of Iowa.

# Acknowledgements

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

- Vadiraja Bhatt-Cisco
- Mehul Bhatt-Cisco
- Vijay Kumar D-Cisco
- Hardik Patel-Cisco
- TJ Singh-Cisco
- Bathu Krishnan-Cisco
- Sindhu Sudhir-Cisco
- Kevin Phillips-EMC
- John Moran-EMC
- Kathy Sharp-EMC

# About Cisco Validated Design (CVD) Program

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

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

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

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

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

© 2012 Cisco Systems, Inc. All rights reserved



# Cisco Solution for EMC VSPEX Microsoft Hyper-V Architectures

---

## Executive Summary

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

## Introduction

As part of an effort to improve and enhance the performance and capabilities of its product line, Cisco and EMC from time to time release revisions of its hardware and software. Therefore, some functions described in this guide may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

## Target Audience

The reader of this document is expected to have the necessary training and background to install and configure Microsoft Hyper-V, EMC VNXe series storage, Cisco Nexus 5548UP and 3048 switches, and Cisco Unified Computing System (UCS) C220 M3 rack servers. External references are provided wherever applicable and it is recommended that the reader be familiar with these documents.

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



---

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

Copyright 2012 Cisco Systems, Inc. All rights reserved.

## Purpose

This document describes the steps required to deploy and configure a Cisco solution for the EMC VSPEX for Microsoft Hyper-V architectures. The document covers two types of Microsoft Hyper-V architectures:

- Microsoft Hyper-V for 50 virtual machines
- Microsoft Hyper-V for 100 virtual machines

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

## Business Needs

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

Business applications are moving into the consolidated compute, network, and storage environment. The Cisco solution for the EMC VSPEX using Microsoft Hyper-V helps to reduce 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. The following are the business needs for the Cisco solution for EMC VSPEX Microsoft Hyper-V architectures:

- Provide an end-to-end virtualization solution to utilize the capability of the unified infrastructure components.
- Provide a Cisco VSPEX for Microsoft Hyper-V Infrastructure as a Service (IaaS) solution for efficiently virtualizing 50 or 100 virtual machines for varied customer use cases.
- Provide a reliable, flexible, and scalable reference design.

## Solution Overview

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

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

- Cisco C-series Unified Computing System servers
- Cisco Nexus 5000 series or 3000 series switches depending on the scale of the solution
- Cisco virtual Port Channels for network load balancing and high availability
- EMC VNXe3150 or VNXe3300 storage components as per the scale needs
- Microsoft Windows Server 2008 R2 SP1 Hyper-V
- Microsoft SQL Server 2008 R2 SP1 database
- Microsoft System Center 2012 Virtual Machine Manager

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

## Technology Overview

This section describes the various technologies used in this solution and their benefits.

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

The main components of 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 5500/5600 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 C220 M3 Rack Mount Servers

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

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



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

## I/O Adapters


The Cisco UCS rack mount server has various Converged Network Adapters (CNA) options. The UCS P81E Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

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

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

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

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



## Cisco Adapter Fabric Extender

Cisco Adapter FEX extends the Cisco FEX technology into traditional rack servers. The Cisco Adapter FEX technology enables the server adapter to be logically partitioned into multiple virtual network interface cards (vNICs). Each vNIC behaves like a physical NIC port and meets the network connectivity needs for each application, so that security and quality of service (QoS) policies can be applied for each vNIC and each application.

**Figure 3** *Network Adapter vNICs as Physical Ports on the Cisco Nexus 5500 Series*

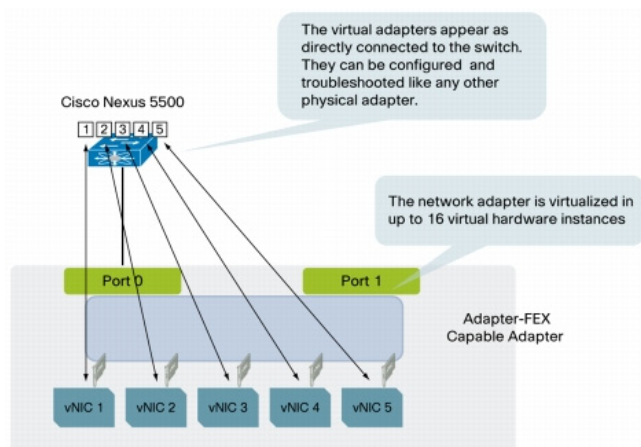


Figure 3 shows that a server with the virtualized adapters (called vNICs) can offer an operating system a number of virtual adapters, and with the A-FEX technology, vNICs are presented as directly connected interfaces to the Cisco Nexus 5500 series switches. All the switching between vNICs occurs on the

upstream Cisco Nexus 5500 series switches, as though they were interfaces of a remote linecard or fabric extenders. In addition to this, all the features right from access control lists (ACLs) to private VLANs, quality of service (QoS), and so on, are available on the remote interfaces.

The redundancy or teaming configuration is not required on the operating system anymore since it is implemented in hardware and controlled by the Cisco Nexus 5500 series switches.

The provisioning model allows the network administrator to define profiles with specific network definitions (mode access or trunk, VLAN, and so on). The server administrator has the choice of defining the number of vNICs and the profile to map them with.

## Cisco Nexus 5548UP Switch

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

**Figure 4** *Cisco Nexus 5548UP Switch*



## Cisco Nexus 3048 Switch

The Cisco Nexus® 3048 Switch is a line-rate Gigabit Ethernet top-of-rack (ToR) switch and is part of the Cisco Nexus 3000 Series Switches portfolio. The Cisco Nexus 3048, with its compact one-rack-unit (1RU) form factor and integrated Layer 2 and 3 switching, complements the existing Cisco Nexus family of switches. This switch runs the industry-leading Cisco® NX-OS Software operating system, providing customers with robust features and functions that are deployed in thousands of data centers worldwide. The Cisco Nexus 3048 switch is shown in [Figure 5](#).

**Figure 5** *Cisco Nexus 3048 Switch*



## Microsoft Windows Server 2008 R2 SP1 Hyper-V

Microsoft Hyper-V is an integral part of Windows Server and provides a foundational virtualization platform that enables you to transition to the cloud. With Windows Server 2008 R2 you get a compelling solution for core virtualization scenarios – production server consolidation, dynamic datacenter, business continuity, VDI, and test & development.

Microsoft Hyper-V provides you better flexibility with features like live migration and cluster shared volumes for storage flexibility.

Microsoft Hyper-V also delivers greater scalability with support for up to 64 logical processors, 2 TB of RAM, NUMA awareness, and improved performance with support for dynamic memory and enhanced networking support.

## Microsoft System Center 2012 Virtual Machine Manager

Microsoft System Center 2012 Virtual Machine Manager (VMM) is a management solution for the virtualized datacenter. This solution enables you to configure and manage your virtualization host, networking, and storage resources in order to create, deploy, and manage virtual machines and services to private clouds that you have created.

## EMC Storage Technologies and Benefits

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

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

The EMC VNXe series storage arrays have the following customer benefits:

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

## Software Suites

The following are the available EMC software suites:

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

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

## Software Packs

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

This is the available EMC protection software pack.

## EMC Avamar

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

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

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

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

# Architectural Overview

This Cisco Validated Design discusses the deployment model for the following two Microsoft Hyper-V server virtualization solutions:

- Microsoft Hyper-V solution for 50 virtual machines
- Microsoft Hyper-V solution for 100 virtual machines

Table 1 lists the mix of hardware components, their quantities and software components used for different solutions:

**Table 1**      *Hardware and Software Components for Various Solutions*

Components	Hyper-V 50 VMS	Hyper-V 100 VMs
Servers	Three Cisco C220 M3 servers	Four Cisco C220 M3 servers
Adapters	2 Cisco GigE I350 LOM 1 Broadcom NetXtreme II 5709 quad-port per server	2 Cisco GigE I350 LOM 1 Cisco UCS P81E VIC per server
Network Switches	Two Cisco Nexus 3048 switches	Two Cisco Nexus 5548UP switches

**Table 1** *Hardware and Software Components for Various Solutions*

Components	Hyper-V 50 VMS	Hyper-V 100 VMs
Storage	EMC VNXe3150	EMC VNXe3300
Network Speed	1 GE	10 GE
Hypervisor	Microsoft Windows Server 2008 R2 SP1 Hyper-V	Microsoft Windows Server 2008 R2 SP1 Hyper-V

Table 2 lists the various hardware and software components which occupies different tiers of the Cisco solution for EMC VSPEX using Microsoft Hyper-V architectures under test.

**Table 2** *Hardware and Software Components of Hyper-V Architectures*

Vendor	Name	Version	Scope of VSPEX solution
Cisco	C220 M3 servers	1.4(4a).1 - CIMC C220M3.1.4.4c.0 - BIOS	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs
Cisco	Cisco Nexus 5548UP Switches	5.1(3)N1(1a)	Only Microsoft Hyper-V 100 VMs
Cisco	Cisco Nexus 3048 Switches	5.0(3)U2(2b)	Only Microsoft Hyper-V 50 VMs
EMC	EMC VNXe3150	2.2.0.16150	Only Microsoft Hyper-V 50 VMs
EMC	EMC VNXe3300	2.2.0.16150	Only Microsoft Hyper-V 100 VMs
EMC	EMC Avamar	6.0.0-592	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs
EMC	Data Domain OS	5.1.0.9-282511	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs
Microsoft	Windows Server 2008 R2	2008 R2 SP1	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs
Microsoft	System Center VMM	SCVMM 2012 with update Rollup1	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs

**Table 2** *Hardware and Software Components of Hyper-V Architectures*

Vendor	Name	Version	Scope of VSPEX solution
Microsoft	Microsoft Windows Server 2008 R2	2008 R2 SP1	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs
Microsoft	Microsoft SQL server	2008 R2 SP1	Both Microsoft Hyper-V 50 VMs and Microsoft Hyper-V 100 VMs

Table 3 outlines the C220 M3 server configuration details (per server basis) across all the Microsoft Hyper-V architectures.

**Table 3** *Cisco UCS C220 M3 Server Hardware Configuration*

Component	Capacity
Memory (RAM)	64 GB (8X8 MB DIMM)
Processor	2 x Intel® Xenon® E5-2650 CPUs, 2 GHz, 8 cores, 16 threads
Network Adapter	2 x Cisco 1GigE 1350 LOM (LAN on Motherboard)
Local Storage	2 x 600 GB SAS 15k RPM hard disk.

## Storage Guidelines

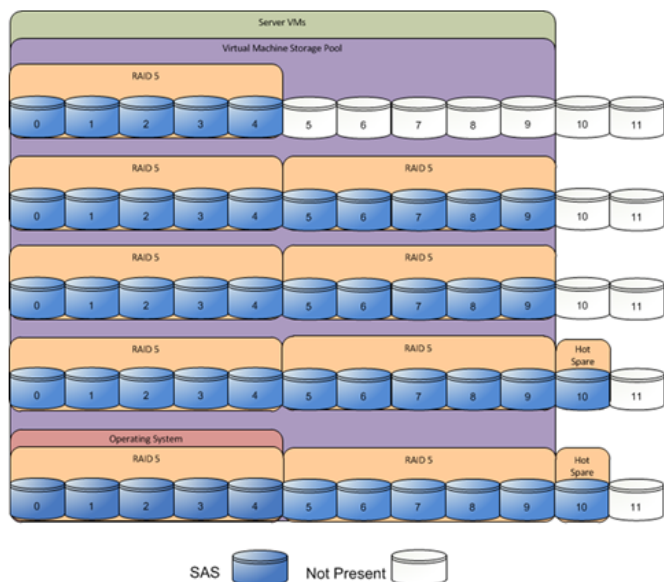
The architecture diagrams in this section show the physical disk layout. Disk provisioning on the EMC VNXe series is simplified through the use of wizards, so that administrators do not choose which disks belong to a given storage pool. The wizard may choose any available disk of the proper type, regardless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

The reference architecture uses the following configuration:

- Disk allocations for different architectures. The following are the different architectures:
  - 50 VMs—Forty-five 600 GB SAS disks are allocated to a single storage pool as nine 4+1 RAID 5 groups (sold as 5-disk packs).
  - 100 VMs—Seventy-seven 600 GB SAS disks are allocated to a single storage pool as eleven 6+1 RAID 5 groups (sold as 7-disk packs) for 100 virtual machines architecture.
- EMC recommends that in addition to the above numbers at least one hot spare disk is allocated for each 30 disks of a given type.

The EMC 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.

**Figure 6** *Storage Architecture for 50 VMs on EMC VNXe3150*



**Figure 7** *Storage Architecture for 100 VMs on EMC VNXe3300*

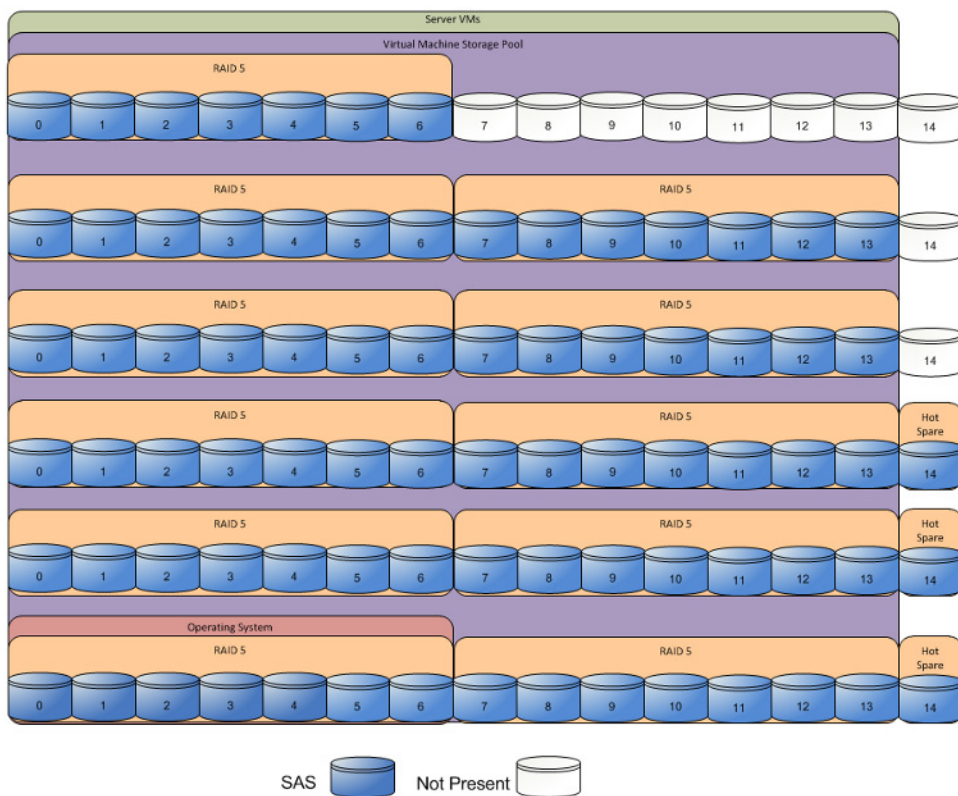


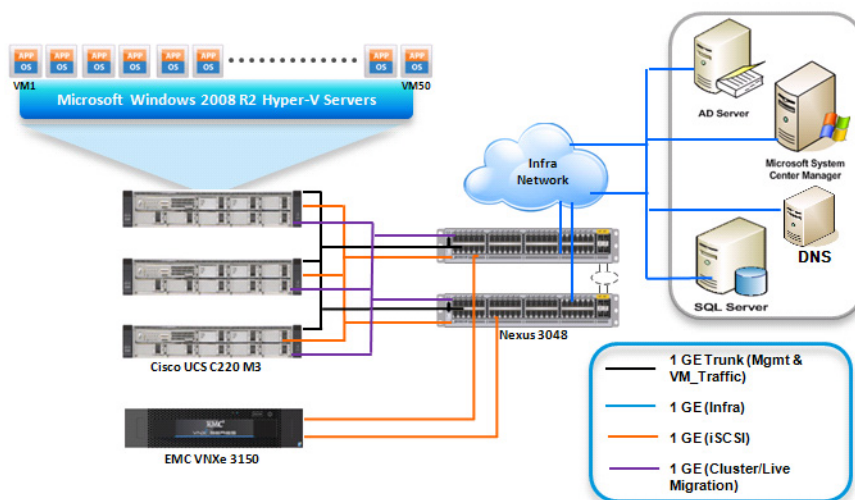
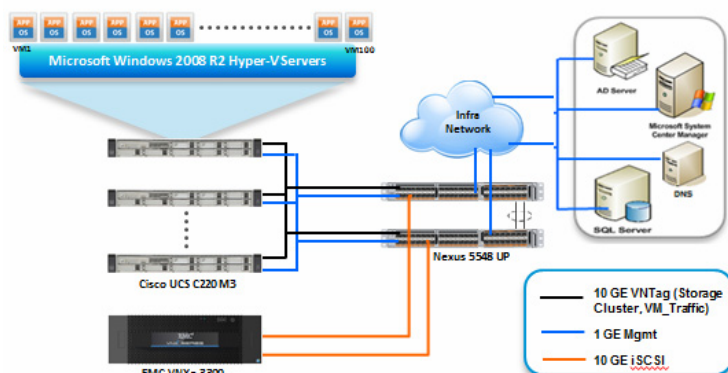
Table 4 provides the datastores size details for the two types of architectures laid out in Figure 6 and Figure 7.



**Table 4**      **Datastores Details for the Microsoft Hyper-V Architectures**

Parameters	50 Virtual Machines	100 Virtual Machines
Disk capacity & type	600 GB SAS	600 GB SAS
Number of disks	45	77
RAID type	4 + 1 RAID 5 groups	6 + 1 RAID 5 groups
Number of RAID Groups	9	11

Both reference architectures assume that there is an existing infrastructure / management network available where a virtual machine or physical machine hosting SCVMM server, Database server, and Microsoft Windows Active Directory / DNS server are present. [Figure 8](#) and [Figure 9](#) show high level solution architecture for up to 50 and up to 100 virtual machines, respectively.

**Figure 8**      **Reference Architecture for 50 Virtual Machines****Figure 9**      **Reference Architecture for 100 Virtual Machines**

As it is evident in the above diagrams, following are the high level design points of Microsoft Hyper-V 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 uplink 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 which hosts the SCVMM, Database, and Active Directory servers. However, design does require accessibility of certain VLANs from the infrastructure network to reach the servers.

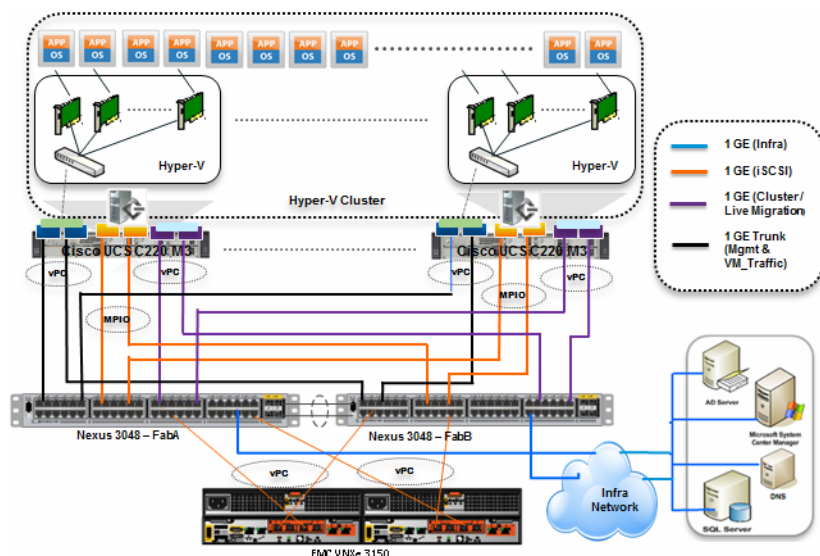
Microsoft Windows Server 2008 R2 SP1 Hyper-V is used as hypervisor operating system on each server and is installed on local hard drives. Typical load is 25 virtual machines per server.

## Architecture for 50 Microsoft Hyper-V Virtual Machines

Figure 10 shows the logical layout of 50 Microsoft Hyper-V virtual machines. Following are the key aspects of this solution:

- Three Cisco C220 M3 servers are used.
- The solution uses two Cisco Nexus 3048 switches, dual-port Cisco 1GigE I350 LOM and 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.
- NIC teaming of the adapters on the host provide load balancing and redundancy as shown in Figure 10. Team 1 has two LoM ports for host management and VM access, separated via VLANs. Team 2 has two Broadcom ports for all cluster traffic.
- EMC VNXe3150 is used as a storage array.

**Figure 10** Cisco Solution for 50 Virtual Machines Using Microsoft Hyper-V

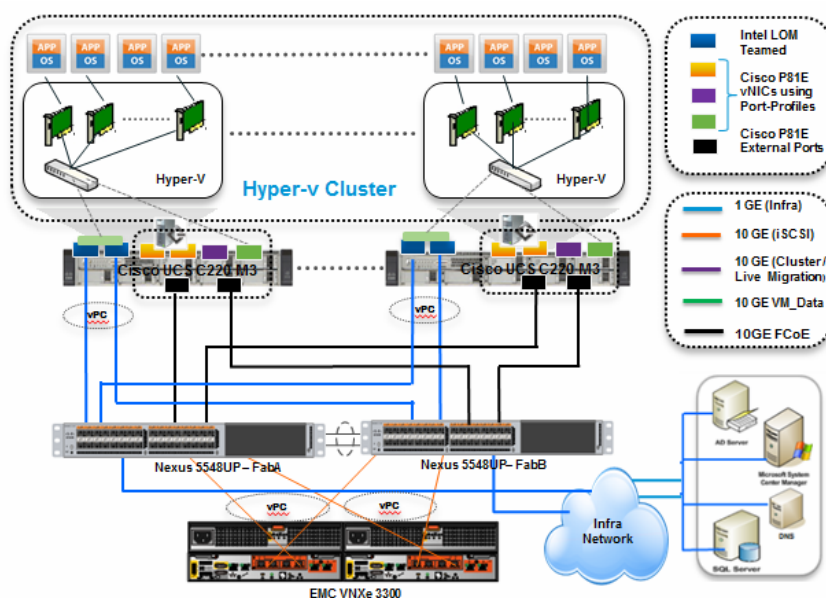


## Architecture for 100 Microsoft Hyper-V Virtual Machines

Figure 11 shows the logical layout of 100 Microsoft Hyper-V virtual machines. Following are the key aspects of this solution:

- Four Cisco C220 M3 servers are used.
- The solution uses two Cisco Nexus 5548UP switches and 10 Gbps Cisco VIC adapters. This results in the 10Gbps solution for the storage access and network and makes live migration and storage access much faster compared to the 1 Gbps solution.
- Virtual port-channels on storage side networking provide high-availability and load balancing.
- Cisco VIC P81E supports Adapter-FEX feature of Cisco Nexus 5500 series switches. It provides NIC level redundancy at the adapter level. On the switch side the ports are set to vntag mode. Each vNIC carved out of NIV enabled Cisco VIC P81E creates a corresponding virtual ethernet interfaces on the switch with unique virtual links or channels. 2 vNICs allow storage traffic, 1 vNIC allows cluster traffic and one vNIC allows VM traffic.
- NIC Teaming of Cisco 1GigE I350 LOM port provides load balancing and redundancy as shown in the Figure 11. This teamed NIC allows management VLAN traffic. The EMC VNXe3300 is used as a storage array.

**Figure 11** Cisco Solution for 100 Virtual Machines Using Microsoft Hyper-V



## Sizing Guidelines

It is important to define a reference workload in virtual infrastructures. 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 a 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. [Table 5](#) provides the characteristics of the virtual machine:

**Table 5** *Virtual Machine Characteristics*

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

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

## Applying the Reference Workload

When considering an existing server that will move into a virtual infrastructure, you have the opportunity to gain efficiency by correctly 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. 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 the 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 Build Application*

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

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

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

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

**Example 2     Point of Sale System**

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

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

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

In this case the one virtual machine uses the resources of eight reference virtual machines. If this was implemented on a resource pool for 50 virtual machines, there are 42 virtual machines of capability remaining in the pool.

**Example 3     Web Server**

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

The following are the requirements to virtualize this application:

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

In this case the virtual machine would use the resources of four reference virtual machines. If this was implemented on a resource pool for 100 virtual machines, there are 96 virtual machines of capability remaining in the pool.

**Summary of Examples**

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

In more advanced cases, there may be trade-offs between memory and I/O or other relationships where in increasing the amount of one resource decreases the need for another. In these cases, the interactions between resource allocations become highly complex, which is out of the scope of this document.

However, once the change in the 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. You can also use the Microsoft Assessment and Planning (MAP) toolkit to assist in the analysis of the current workload. You can download the toolkit from the following Microsoft link:

<http://www.microsoft.com/map>

**Networking Configuration Guidelines**

This document provides details for setting up a redundant, highly-available configuration. As such, references are made as to which component is being configured with each step whether that be A or B. For example, SP A and SP B, are used to identify the two EMC VNXe storage controllers that are provisioned with this document while the Nexus A and Nexus B identify the pair of Cisco Nexus switches that are configured. Additionally, this document details steps for provisioning multiple UCS hosts and these are identified sequentially, M100N1 and M100N2, and so on. Finally, when indicating

that the reader should include information pertinent to their environment in a given step, this is indicated with the inclusion of *<italicized text>* as part of the command structure. See the following example for the VLAN create command on the Cisco Nexus Switch:

```
switchA(config)# vlan {vlan-id | vlan-range}
switchA(config)# vlan <storage VLAN ID>
```

This document is intended to allow the reader to fully configure the customer environment. In order to do so, there are various steps which will require you to insert your own naming conventions, IP addresses, and VLAN schemes, as well as record appropriate iSCSI IQN or MAC addresses. [Table 8](#) details the list of VLANs necessary for deployment as outlined in this guide.

## VSPEX Configuration Guidelines

To configure the Cisco solution for EMC VSPEX Microsoft Hyper-V architectures, follow these steps:

1. [Pre-Deployment Tasks, page 22](#)
2. [Cabling Information, page 23](#)
3. [Prepare and Configure the Cisco Nexus 5548UP Switch, page 26](#)
4. [Infrastructure Servers, page 37](#)
5. [Prepare the Cisco UCS C220 M3 Servers, page 38](#)
6. [Prepare the EMC VNXe3300 Storage, page 73](#)
7. [Microsoft Windows Failover Cluster Setup, page 90](#)
8. [Microsoft System Center-2012 VMM Configuration, page 104](#)
9. [Validating Cisco Solution for EMC VSPEX Microsoft Hyper-V Architectures, page 136](#)

The above steps are described in the following sections.

## Pre-Deployment Tasks

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

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

## Customer Configuration Data

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

The Customer Configuration Data section provides a table to maintain a record of relevant information. This form can be expanded or contracted as required, and information may be added, modified, and recorded as deployment progresses.

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

## VSPEX M100 Configuration Details

### Cabling Information

The following information is provided as a reference for cabling the physical equipment in a VSPEX M100 environment. The tables in this section include 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.

Follow the cabling 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, ensure that the configurations match the cabling details provided in the tables and figures in this section.

[Figure 12](#) shows the VSPEX M100 cabling diagram. The alphabets labeled indicate connections to the end points rather than port numbers on the physical device. For example, connection A is a 10 Gb target port connected from EMC VNXe3300 SP B to Cisco Nexus 5548 A and connection R is a 10 Gb target port connected from Cisco VIC P81E uplink port 1 on Server 2 to Cisco Nexus 5548 B. Connections U and V are 10 Gb vPC peer-links connected from Cisco Nexus 5548 A to Cisco Nexus 5548 B.



**Figure 12**      **Cabling Details for VSPEX Microsoft Hyper-V 100 Virtual Machines**




Table 6 and Table 7 show that there are five major cabling in these architectures:

- Inter switch links
- Data connectivity for servers (trunk links)
- Management connectivity for servers
- Storage connectivity
- Infrastructure connectivity

Table 6 provides the Cisco Nexus 5548 A Ethernet Cabling Information Local Device Local Port Connection Remote.

**Table 6**      **Cabling details for 100 VMs on Cisco Nexus 5548UP A**

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
E	Eth1/1	1	Access	1(D)	2	C220 Server1- 1GE LOM 1
F	Eth1/2	1	Access	1(D)	3	C220 Server2- 1GE LOM 1
G	Eth1/3	1	Access	1(D)	4	C220 Server3- 1GE LOM 1
H	Eth1/4	1	Access	1(D)	5	C220 Server4- 1GE LOM 1



**Table 6** *Cabling details for 100 VMs on Cisco Nexus 5548UP A*

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
U	Eth1/7	1,40,45,46	Trunk	10(D)	7	VPC peer link
V	Eth1/8	1,40,45,46	Trunk	10(D)	7	VPC peer link
M	Eth1/9	1,40,45,46	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
N	Eth1/10	1,40,45,46	vntag	10(D)	-	C220 Server2- P81E VIC Port 0
O	Eth1/11	1,40,45,46	vntag	10(D)	-	C220 Server3- P81E VIC Port 0
P	Eth1/12	1,40,45,46	vntag	10(D)	-	C220 Server4- P81E VIC Port 0
(not shown)	Eth1/15	1,40,45,46	Trunk	10(D)	15	Uplink to Infrastructure network
(not shown)	Eth1/17	1,40,45,46	Trunk	10(D)	17	Uplink to Infrastructure network
A	Eth2/1	40	Access	10(D)	21	EMC VNXe3300 (eth10) - SP B
C	Eth2/2	40	Access	10(D)	22	EMC VNXe3300 (eth10) - SP A

**Table 7** *Cabling details for 100 VMs on Cisco Nexus 5548UP B*

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
I	Eth1/1	1	Access	1(D)	2	C220 Server1- 1GE LOM 1
J	Eth1/2	1	Access	1(D)	3	C220 Server2- 1GE LOM 1
K	Eth1/3	1	Access	1(D)	4	C220 Server3- 1GE LOM 1
L	Eth1/4	1	Access	1(D)	5	C220 Server4- 1GE LOM 1
U	Eth1/7	1,40,45,46	Trunk	10(D)	7	VPC peer link
V	Eth1/8	1,40,45,46	Trunk	10(D)	7	VPC peer link
Q	Eth1/9	1,40,45,46	vntag	10(D)	-	C220 Server1- P81E VIC Port 0
R	Eth1/10	1,40,45,46	vntag	10(D)	-	C220 Server2- P81E VIC Port 0
S	Eth1/11	1,40,45,46	vntag	10(D)	-	C220 Server3- P81E VIC Port 0

**Table 7**      **Cabling details for 100 VMs on Cisco Nexus 5548UP B**

Cable ID	Switch Interface	VLAN	Mode	Speed (Gbps)	Port Channel	Remote Device port
T	Eth1/12	1,40,45,46	vntag	10(D)	-	C220 Server4-P81E VIC Port 0
(not shown)	Eth1/16	1,40,45,46	Trunk	10(D)	15	Uplink to Infrastructure network
(not shown)	Eth1/18	1,40,45,46	Trunk	10(D)	17	Uplink to Infrastructure network
A	Eth2/1	40	Access	10(D)	21	EMC VNXe3300 (eth10) - SP B
C	Eth2/2	40	Access	10(D)	22	EMC VNXe3300 (eth10) - SP A

## Prepare and Configure the Cisco Nexus 5548UP Switch

The following section provides a detailed procedure for configuring the Cisco Nexus 5548 switches for use in EMC VSPEX Microsoft Hyper-V 100 VMs.

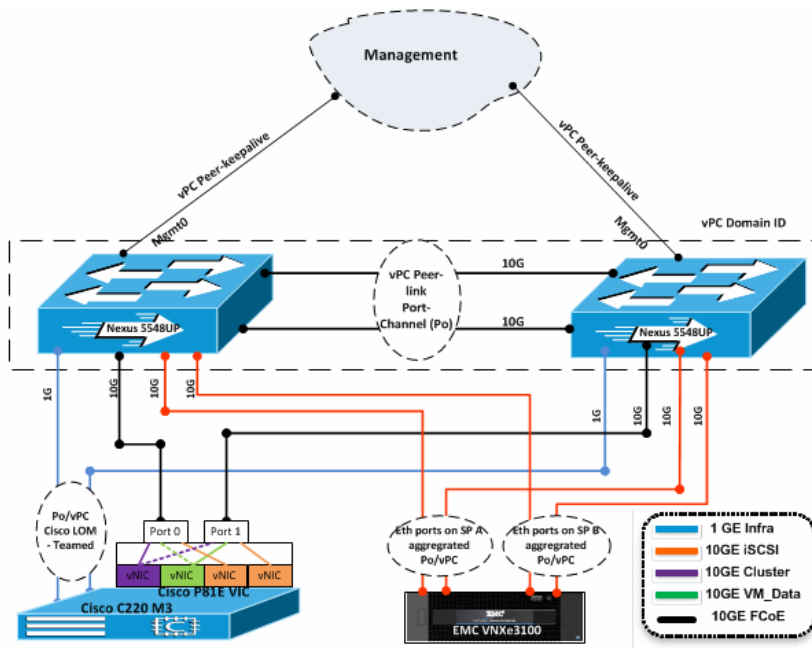
[Figure 13](#) shows two switches configured for vPC. In vPC, a pair of switches acting as vPC peer endpoints looks like a single entity to port-channel-attached devices, although the two devices that act as logical port-channel endpoint are still two separate devices. This provides hardware redundancy 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 5548 UP running release 5.1(3)N1(1a).

**Figure 13**      **Networking Configuration for Microsoft Hyper-V 100 Virtual Machines**



## Initial Setup of Cisco Nexus Switches

These steps provide details for the initial setup on both Cisco Nexus 5548 switches.

### For Cisco Nexus A and Cisco Nexus B

After booting and connecting to the serial or console port of the switch, the NX-OS setup should automatically start.

1. Enter yes to enforce secure password standards.
2. Enter the password for the admin user.
3. Enter the password a second time to commit the password.
4. Enter yes to enter the basic configuration dialog.
5. Create another login account (yes/no) [n]: Enter.
6. Configure read-only SNMP community string (yes/no) [n]: Enter.
7. Configure read-write SNMP community string (yes/no) [n]: Enter.
8. Enter the switch name: <Nexus A Switch name> Enter.
9. Continue with out-of-band (mgmt0) management configuration? (yes/no) [y]: Enter.
10. Mgmt0 IPv4 address: <Nexus A mgmt0 IP> Enter.
11. Mgmt0 IPv4 netmask: <Nexus A mgmt0 netmask> Enter.
12. Configure the default gateway? (yes/no) [y]: Enter.
13. IPv4 address of the default gateway: <Nexus A mgmt0 gateway> Enter.
14. Enable the telnet service? (yes/no) [n]: Enter.
15. Enable the ssh service? (yes/no) [y]: Enter.

16. Type of ssh key you would like to generate (dsa/rsa):rsa.
17. Number of key bits <768–2048> :1024 Enter.
18. Configure the ntp server? (yes/no) [y]: n Enter
19. NTP server IPv4 address: <NTP Server IP> Enter.
20. Enter basic FC configurations (yes/no) [n]: Enter.
21. Would you like to edit the configuration? (yes/no) [n]: Enter.
22. Be sure to review the configuration summary before enabling it.
23. Use this configuration and save it? (yes/no) [y]: Enter.
24. Configuration may be continued from the console or by using SSH. To use SSH, connect to the mgmt0 address of Nexus A or B.
25. Log in as user admin with the password previously entered.

## Enabling Features and Global Configuration

### For Cisco Nexus A and Cisco Nexus B

1. Type config t to enter the global configuration mode.
2. Type feature lacp.
3. Type feature interface-vlan
4. Type feature vpc.

## Set Global Configurations

These steps provide details for setting global configurations.

### For Cisco Nexus A and Cisco Nexus B

1. From the global configuration mode, type spanning-tree port type network default to make sure that, by default, the ports are considered as network ports in regards to spanning-tree.
2. Type spanning-tree port type edge bpduguard default to enable bpduguard on all edge ports by default.
3. Type spanning-tree port type edge bpdufilter default to enable bpdufilter on all edge ports by default.

## Configure VLANs

The steps in this section provide details for creating the VLANs as per the below given reference [Table 8](#).

**Table 8** VLANs for EMC VSPEX Microsoft Hyper-V M100 Setup

VLAN Name	VLAN Purpose	ID used in this Document	Host NICs in VLANs
storage	For iSCSI traffic	40	2 Cisco VNICs
VM_traffic	For VM data	45	1 Cisco VNIC

**Table 8** VLANs for EMC VSPEX Microsoft Hyper-V M100 Setup

VLAN Name	VLAN Purpose	ID used in this Document	Host NICs in VLANs
cluster	For live migration	46	1 Cisco VNIC
default	For management and cluster	1	2 Cisco 1 GigE 1350 LOM in team

**Note**

For details on network addresses, see the section [Customer Configuration Data Sheet, page 138](#). This section provides tabulated record of relevant information (to be filled at the customer's end). This form can be expanded or contracted as required, and information may be added, modified, and recorded as the deployment progresses.

**For Nexus A and Nexus B**

1. Type config-t.
2. Type vlan <storage VLAN ID>.
3. Type name storage
4. Type exit.
5. Type vlan <cluster VLAN ID>.
6. Type name cluster
7. Type exit.
8. Type vlan <vm\_traffic VLAN ID>.
9. Type name VM\_traffic
10. Type exit.

## Configure Port-Channels

### Create Port-Channels

**For Cisco Nexus 5548 A and Cisco 5548 B**

1. From the global configuration mode, type interface Po7.
2. Type description vPC peer-link.
3. Type exit.
4. Type interface Eth1/7-8
5. Type channel-group 7 mode active.
6. Type no shutdown.
7. Type exit.
8. Type interface Po2.
9. Type description <Cisco 1GigE LOM 1 on UCS Server 1 – For Nexus A>/<Cisco 1GigE LOM 2 on UCS Server 1 – For Nexus B>
10. Type exit.

11. Type interface Eth1/1.
12. Type channel-group 2 mode active.
13. Type no shutdown.
14. Type exit
15. Type interface Po3.
16. Type description<Cisco 1GigE LOM 1 on UCS Server 2 – For Nexus A>/<Cisco 1GigE LOM 2 on UCS Server 2 – For Nexus B>.
17. Type exit.
18. Type interface Eth1/2.
19. Type channel-group 3 mode active.
20. Type no shutdown.
21. Type exit.
22. Type interface Po4.
23. Type description <Cisco 1GigE LOM 1on UCS Server 3 – For Nexus A>/<Cisco 1GigE LOM 2 on UCS Server 3 – For Nexus B>.
24. Type exit.
25. Type interface Eth1/3.
26. Type channel-group 4 mode active.
27. Type no shutdown.
28. Type exit
29. Type interface Po5.
30. Type description <Cisco 1GigE LOM 1on UCS Server 4>.
31. Type exit.
32. Type interface Eth1/4.
33. Type channel-group 5 mode active.
34. Type no shutdown.
35. Type exit
36. Type interface Po15.
37. Type description <Infrastructure Network>.
38. Type exit.
39. Type interface Eth1/15.
40. Type channel-group 15 mode active.
41. Type no shutdown.
42. Type exit
43. Type interface Po17.
44. Type description < Infrastructure Network>.
45. Type exit.
46. Type interface Eth1/17.

47. Type channel-group 17 mode active.
48. Type no shutdown.
49. Type exit
50. Type interface Po21.
51. Type description <VNxe Storage Processor B>
52. Type exit.
53. Type interface Eth2/1.
54. Type channel-group 21 mode active.
55. Type no shutdown.
56. Type exit
57. Type interface Po22.
58. Type description <VNxe Storage Processor A>
59. Type exit.
60. Type interface Eth2/2.
61. Type channel-group 22 mode active.
62. Type no shutdown.
63. Type exit

### Add Port Channel Configurations

These steps provide details for adding Port Channel configurations.

#### For Cisco Nexus A and Cisco Nexus B

1. From the global configuration mode, type interface Po7.
2. Type switchport mode trunk.
3. Type switchport trunk allowed vlan <storage VLAN ID, cluster VLAN ID, vm\_traffic VLAN ID >.
4. Type spanning-tree port type network.
5. Type no shutdown.
6. Type exit.
7. Type interface Po15.
8. Type switchport mode trunk.
9. Type switchport trunk allowed vlan < mgmt. VLAN ID vm\_traffic VLAN ID >.
10. Type spanning-tree port type network.
11. Type no shut.
12. Type exit.
13. Type interface Po17.
14. Type switchport mode trunk.
15. Type switchport trunk allowed vlan < mgmt. VLAN ID, vm\_traffic VLAN ID >.
16. Type spanning-tree port type network.
17. Type no shut.

18. Type exit.
19. Type interface Po2.
20. Type switchport mode access.
21. Type spanning-tree port type edge.
22. Type no shut.
23. Type exit.
24. Type interface Po3.
25. Type switchport mode access.
26. Type spanning-tree port type edge.
27. Type no shut.
28. Type exit.
29. Type interface Po4.
30. Type switchport mode access.
31. Type spanning-tree port type edge.
32. Type no shut.
33. Type exit.
34. Type interface Po5.
35. Type switchport mode access.
36. Type spanning-tree port type edge.
37. Type no shut.
38. Type exit.
39. Type interface Po21.
40. Type switchport mode access.
41. Type switchport access vlan <storage VLAN ID>
42. Type spanning-tree port type edge.
43. Type no shut.
44. Type exit.
45. Type interface Po22.
46. Type switchport mode access.
47. Type switchport access vlan <storage VLAN ID>
48. Type spanning-tree port type edge.
49. Type no shut.

## Configure Virtual Port Channels

These steps provide details for configuring virtual Port Channels (vPCs).

### For Cisco Nexus A and Cisco Nexus B

1. From the global configuration mode, type vpc domain <Nexus vPC domain ID>.



2. Type peer-keepalive destination <Nexus B mgmt0 IP> source <Nexus A mgmt0 IP>.
3. Type exit.
4. Type interface Po7.
5. Type vpc peer-link.
6. Type exit.
7. Type interface Po15.
8. Type vpc 15.
9. Type exit.
10. Type interface Po17.
11. Type vpc 17.
12. Type exit.
13. Type interface Po2.
14. Type vpc 2.
15. Type exit.
16. Type interface Po3.
17. Type vpc 3.
18. Type exit.
19. Type interface Po4.
20. Type vpc 4.
21. Type exit.
22. Type interface Po5.
23. Type vpc5.
24. Type exit.
25. Type interface Po21.
26. Type vpc 21.
27. Type exit.
28. Type interface Po22.
29. Type vpc 22.
30. Type exit.
31. Type copy run start

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 port-channel summary” and “show vpc” commands as shown in [Figure 14](#).

**Figure 14**      **Show VLAN Brief Output**

```
EMC-5548B# sh vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Po2, Po3, Po4, Po5, Po7, Po15 Po17, Eth1/5, Eth1/6, Eth1/13 Eth1/14, Eth1/15, Eth1/17 Eth1/19, Eth1/20, Eth1/21 Eth1/22, Eth1/23, Eth1/24 Eth1/25, Eth1/26, Eth1/27 Eth1/28, Eth1/29, Eth1/30 Eth1/31, Eth1/32, Eth2/3, Eth2/4 Eth2/5, Eth2/6, Eth2/7, Eth2/8
40	storage	active	Po7, Po15, Po17, Po21, Po22 Eth2/7, Veth32773, Veth32775 Veth32778, Veth32781
45	vm_traffic	active	Po7, Po15, Po17, Eth2/7 Veth32769, Veth32770, Veth32771 Veth32772
46	cluster	active	Po7, Po15, Po17, Eth2/7 Veth32776, Veth32779, Veth32782 Veth32784

Ensure that on both switches, all required VLANs are in “active” status and right set of ports and port-channels are part of the necessary VLANs.

Port-channel configuration can be verified using “show port-channel summary” command. [Figure 15](#) shows the expected output of this command.

**Figure 15**      **Show Port Channel Summary Output**

```
EMC-5548B# sh port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        S - Suspended    r - Module-removed
        s - Switched      R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
```

Group	Port-Channel	Type	Protocol	Member Ports
2	Po2(SU)	Eth	LACP	Eth1/1(P)
3	Po3(SU)	Eth	LACP	Eth1/2(P)
4	Po4(SU)	Eth	LACP	Eth1/3(P)
5	Po5(SU)	Eth	LACP	Eth1/4(P)
7	Po7(SU)	Eth	LACP	Eth1/7(P)    Eth1/8(P)
15	Po15(SU)	Eth	LACP	Eth1/16(P)
17	Po17(SU)	Eth	LACP	Eth1/18(P)
21	Po21(SU)	Eth	LACP	Eth2/1(P)
22	Po22(SU)	Eth	LACP	Eth2/2(P)

In this example, port-channel 7 is the vPC peer-link port-channel, port-channels 2, 3, 4 and 5 are connected to the Cisco 1GigE I350 LOM on the host, port-channels 15 and 17 are connected to the infrastructure network, and port-channels 21 and 22 are connected to the storage array. Make sure that state of the member ports of each port-channel is “P” (Up in port-channel). Note that port may not come up if the peer ports are not properly configured. Common reasons for port-channel port being down are:

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

vPC status can be verified using “show vpc” command. Example output is shown in [Figure 16](#).

**Figure 16 Show vPC Brief Output**

```

EMC-5548B# sh vpc brief
Legend:
      (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id          : 111
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 : 8
Peer Gateway           : Disabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

vPC Peer-link status
-----
id  Port  Status Active vlans
-----
1   Po7   up    1,20-22,40,42-46

vPC status
-----
id  Port  Status Consistency Reason Active vlans
-----
2   Po2   up    success success 1
3   Po3   up    success success 1
4   Po4   up    success success 1
5   Po5   up    success success 1
15  Po15   up    success success 1,20-22,40,42-46
17  Po17   up    success success 1,20-22,40,42-46
21  Po21   up    success success 40
22  Po22   up    success success 40

```

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

## Configure Adapter FEX

The Cisco NX-OS Adapter-FEX feature combines the advantages of the FEX link architecture with server I/O virtualization to create multiple virtual interfaces over a single Ethernet interface. This allows you to deploy a dual port NIC on the server and to configure more than two virtual interfaces that the server sees as a regular Ethernet interface. The advantage of this approach is that it allows you to reduce power and cooling needs and to reduce the number of network ports.

Adapter-FEX can be thought of as a way to divide a single physical link into multiple virtual links or channels. Each channel is identified by a unique channel number and its scope is limited to the physical link.

The physical link connects a port on a server network adapter with an Ethernet port on the switch. This allows the channel to connect a vNIC on the server with a vEthernet interface on the switch.

Packets on each channel are tagged with a VNTag that has a specific source virtual interface identifier (VIF). The VIF allows the receiver to identify the channel that the source used to transmit the packet.

For more information on Adapter-FEX, check the below URLs:

Cisco Nexus 5000 Series NX-OS Adapter-FEX Software Configuration Guide, Release 5.1(3)N1(1)

[http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/adapter-fex/513\\_n1\\_1/b\\_Configuring\\_Cisco\\_Nexus\\_5000\\_Series\\_Adapter-FEX\\_rel\\_5\\_1\\_3\\_N1.pdf](http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/adapter-fex/513_n1_1/b_Configuring_Cisco_Nexus_5000_Series_Adapter-FEX_rel_5_1_3_N1.pdf)

Cisco Adapter Fabric Extender

[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data\\_sheet\\_c78-657397.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/data_sheet_c78-657397.html)

This section provides information about how to enable and configure the Cisco Nexus 5500 series for Adapter-FEX. After completing the following steps in this section you will carve out vNICs on Cisco VIC P81E adapter (explained in the later section “Creating and configuring vNICs on Cisco P81E VIC”) using the below created port-profiles.

## Enabling Switch for Adapter-FEX

Following steps show enabling the virtualization feature on both the Cisco Nexus switches.

### For Cisco Nexus A and Cisco Nexus B

1. Type Configure terminal
2. Type install feature-set virtualization
3. Type feature-set virtualization
4. Type vethernet auto-create

## Configuring the Switch for Adapter-FEX

Following steps show creation of port-profiles on both the switches and put the FEX interfaces into vntag mode for NIC.

### For Cisco Nexus A and Cisco Nexus B

1. Type Configure terminal
2. Type port-profile type vethernet <port-profile name- storage>
3. Type switchport access vlan <storage VLAN ID>
4. Type state enabled
5. Type exit
6. Type port-profile type vethernet <port-profile name- cluster>
7. Type switchport access vlan <cluster VLAN ID>
8. Type state enabled
9. Type exit
10. Type port-profile type vethernet <port-profile name- vm\_traffic>
11. Type switchport access vlan <vm\_traffic VLAN ID>
12. Type state enabled
13. Type exit
14. Type interface ethernet1/9-12
15. Type switchport mode vntag
16. Type exit
17. Type copy run start

**Figure 17** Port-Profile Brief Output

```
EMC-5548A# sh port-profile brief
```

Port Profile	Profile State	Conf Items	Eval Items	Assigned Intfs	Child Profs
cluster	1	1	1	4	0
storage	1	1	1	4	0
vm_traffic	1	1	1	4	0

## Enable Jumbo Frames

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

The following commands enable jumbo frames on the Cisco Nexus switches.

### For Cisco Nexus A and Cisco Nexus B

```
switch(config)#policy-map type network-qos jumbo
switch(config-pmap-nq)#class type network-qos class-default
switch(config-pmap-c-nq)#mtu 9216
switch(config-pmap-c-nq)#exit
switch(config-pmap-nq)#exit
switch(config)#system qos
switch(config-sys-qos)#service-policy type network-qos jumbo
```

**Figure 18** Validate Jumbo Frames Support in the Storage Processors

```
EMC-5548B# ping 10.10.40.60 packet-size 8972 c 10
PING 10.10.40.60 (10.10.40.60): 8972 data bytes
8980 bytes from 10.10.40.60: icmp_seq=0 ttl=254 time=3.859 ms
8980 bytes from 10.10.40.60: icmp_seq=1 ttl=254 time=2.396 ms
8980 bytes from 10.10.40.60: icmp_seq=2 ttl=254 time=2.462 ms
8980 bytes from 10.10.40.60: icmp_seq=3 ttl=254 time=2.461 ms
8980 bytes from 10.10.40.60: icmp_seq=4 ttl=254 time=2.463 ms
8980 bytes from 10.10.40.60: icmp_seq=5 ttl=254 time=2.461 ms
8980 bytes from 10.10.40.60: icmp_seq=6 ttl=254 time=2.463 ms
8980 bytes from 10.10.40.60: icmp_seq=7 ttl=254 time=2.466 ms
8980 bytes from 10.10.40.60: icmp_seq=8 ttl=254 time=2.459 ms
8980 bytes from 10.10.40.60: icmp_seq=9 ttl=254 time=2.463 ms
--- 10.10.40.60 ping statistics ---
10 packets transmitted, 10 packets received, 0.00% packet loss
round-trip min/avg/max = 2.396/2.595/3.859 ms
```

## Infrastructure Servers

Most environments will already have an Active Directory in their infrastructure either running on a virtual machine or on a physical server. This section will not cover the installation of an Active Directory Domain Controller, however it will cover the brief installation of standalone SQL Server 2008 R2 SP1 and System Center VMM 2012 on Microsoft Windows Server 2008 R2 SP1. The following infrastructure servers were used to validate the VSPEX Microsoft Hyper-V architectures.

**Table 9** *Infrastructure Server Details Used for the VSPEX Solution*

Server Name	Role	OS
M50AD.M50VSPEX.COM	Domain Controller, DNS and DHCP	Microsoft Windows Server 2008 R2 SP1
M50DB.M50VSPEX.COM	SQL Server for SCVMM	Microsoft Windows Server 2008 R2 SP1
M50SCVMM.M50VSPEX.COM	SCVMM - 2012	Microsoft Windows Server 2008 R2 SP1

**Note**

For details on network addresses, see the section [Customer Configuration Data Sheet, page 138](#).

## Prepare the Cisco UCS C220 M3 Servers

Preparing the Cisco C220 M3 servers is a common step for all the Hyper-V architectures. Firstly, you need to install the C220 M3 server in a rack. For more information on mounting the Cisco C220 servers, see the installation guide on details about how to physically mount the server:  
[http://www.cisco.com/en/US/docs/unified\\_computing/ucs/c/hw/C220/install/install.html](http://www.cisco.com/en/US/docs/unified_computing/ucs/c/hw/C220/install/install.html)

To prepare the servers, follow these steps:

1. [Configuring Cisco Integrated Management controller \(CIMC\), page 38](#)
2. [Enable Virtualization Technology in BIOS, page 39](#)
3. [Configuring RAID, page 40](#)

These steps are discussed in detail in the following sections.

### Configuring Cisco Integrated Management controller (CIMC)

This section describes procedures to prepare the Cisco UCS C220 M3 servers.

#### Connecting and Powering on the Server (Standalone Mode)

For connecting and powering on the server (Standalone Mode), follow these steps:

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

**Figure 19** *CIMC Configuration Utility*

```

CIMC Configuration Utility  Version 1.5  Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]          NIC redundancy
Shared LOM:     [ ]          None: [X]
Shared LOM 10G: [ ]          Active-standby:[ ]
Cisco Card:     [ ]          Active-active: [ ]
IPV4 (Basic)
DHCP enabled:   [ ]          Factory Defaults
CIMC IP:        10.29.150.101 CIMC Factory Default:[ ]
Subnetmask:     255.255.255.0 Default User (Basic)
Gateway:        10.29.150.1  Default password:
VLAN (Advanced)
VLAN enabled:   [ ]          Reenter password:
VLAN ID:        1
Priority:        0
*****
<Up/Down arrow> Select items    <F10> Save    <Space bar> Enable/Disable
<F5> Refresh                    <ESC> Exit

```

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.

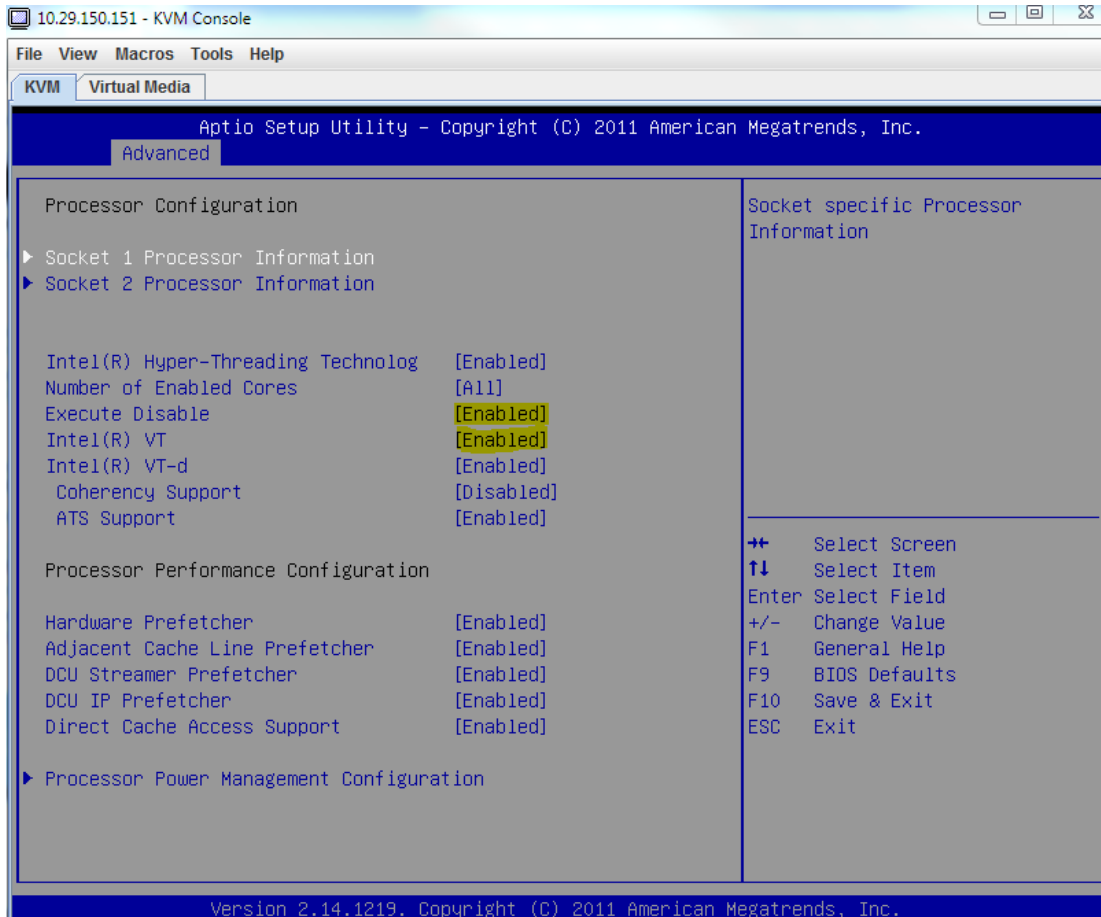
## Enable Virtualization Technology in BIOS

Microsoft Hyper-V requires an x64-based processor, hardware-assisted virtualization (Intel VT enabled), and hardware data execution protection (Execute Disable enabled).

To enable Intel ® VT and Execute Disable in BIOS, follow these steps:

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 **Advanced > Processor Configuration**.

**Figure 20** *Cisco UCS C220 M2 KVM Console*



4. Enable Execute Disable and Intel VT as shown in [Figure 20](#).

## Configuring RAID

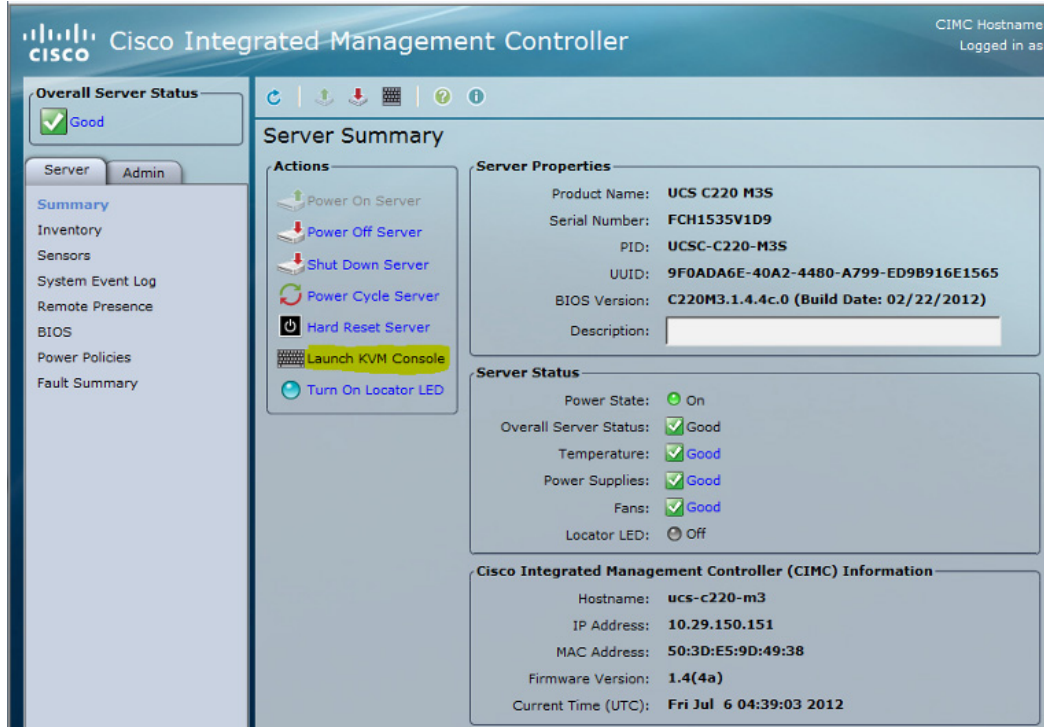
The RAID controller type is Cisco UCSC RAID SAS 2008 and supports 0, 1, 5 RAID levels. We need to configure RAID level 1 for this setup and set the virtual drive as boot drive.

To configure RAID controller, follow these steps:

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.

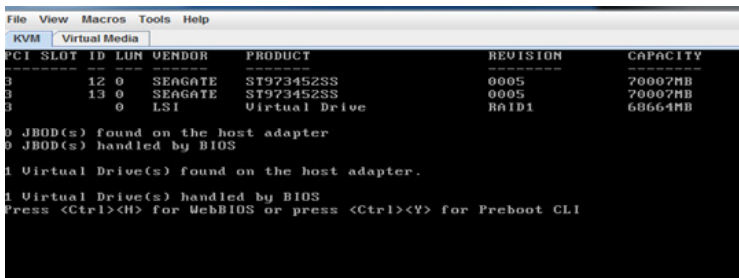


**Figure 21** Cisco UCS C220 M2 CIMC GUI

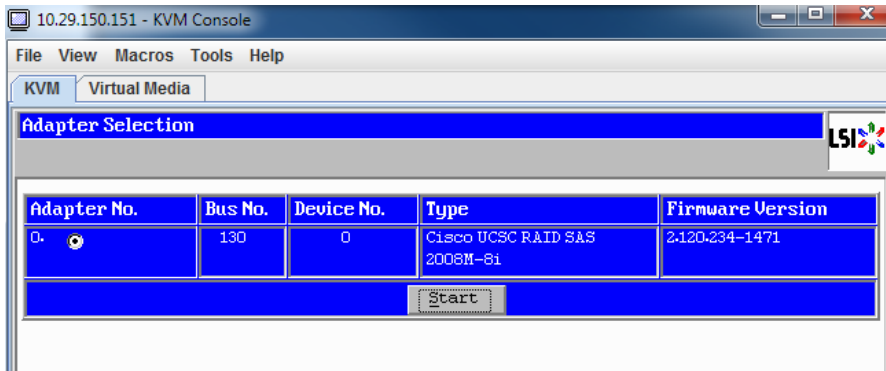


- During bootup, press <Ctrl> <H> when prompted to configure RAID in the WebBIOS.

**Figure 22** KVM Console Showing Cisco UCS C220 M2 Server Booting



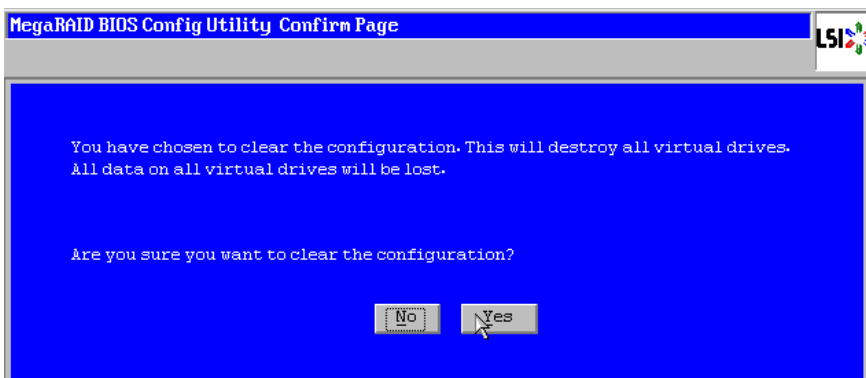
- Choose the adapter and click **Start**.

**Figure 23      Adapter Selection for RAID Configuration**

5. Choose the “New Configuration” radio button and click **Next**.

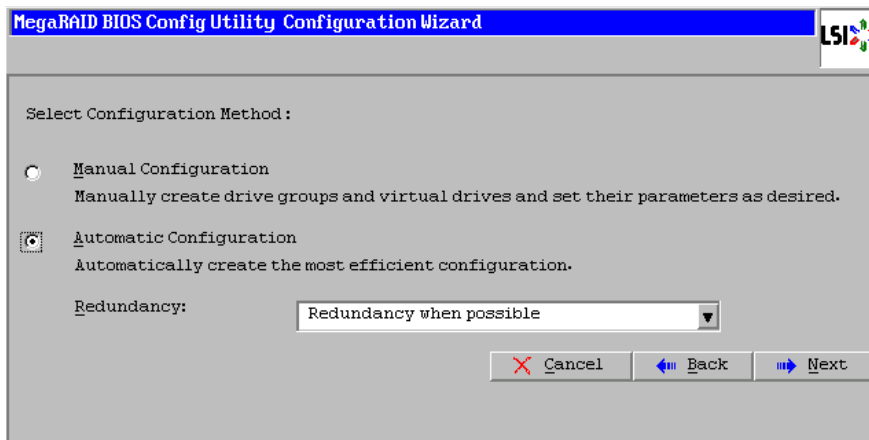
**Figure 24      MegaRAID BIOS Config Utility Configuration**

6. Click **Yes** and then click **Next** to clear the configuration.

**Figure 25      Confirmation Window for Clearing the Configuration**

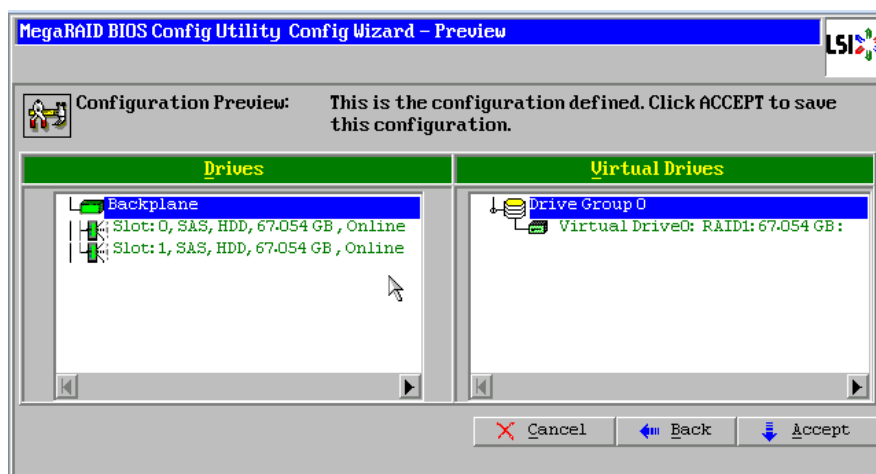
7. If you choose the “Automatic Configuration” radio button and the “Redundancy when possible” option from the “Redundancy” drop-down list and if only two drives are available, the WebBIOS creates a RAID 1 configuration.

**Figure 26**      **Selecting the Configuration Method**



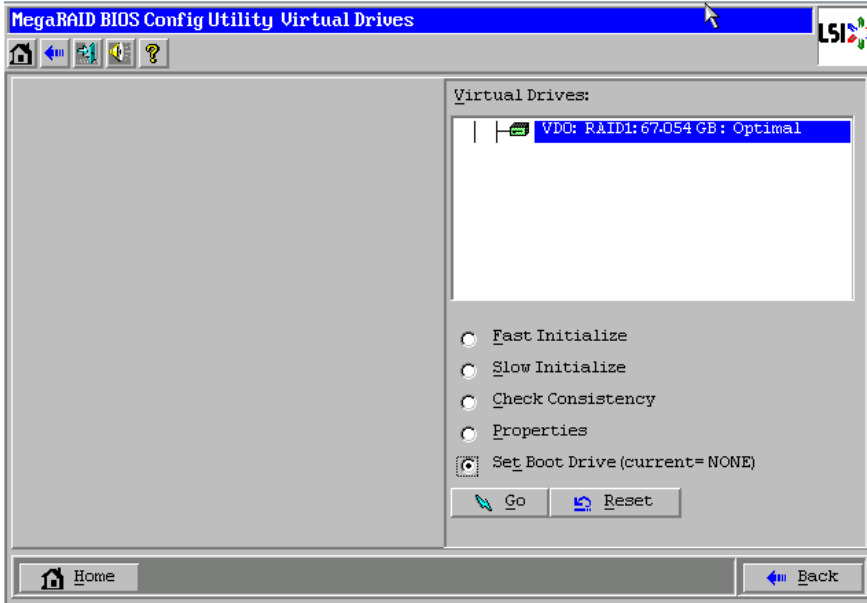
8. Click **Accept** when you are prompted to save the configuration.

**Figure 27**      **Configuration Preview**



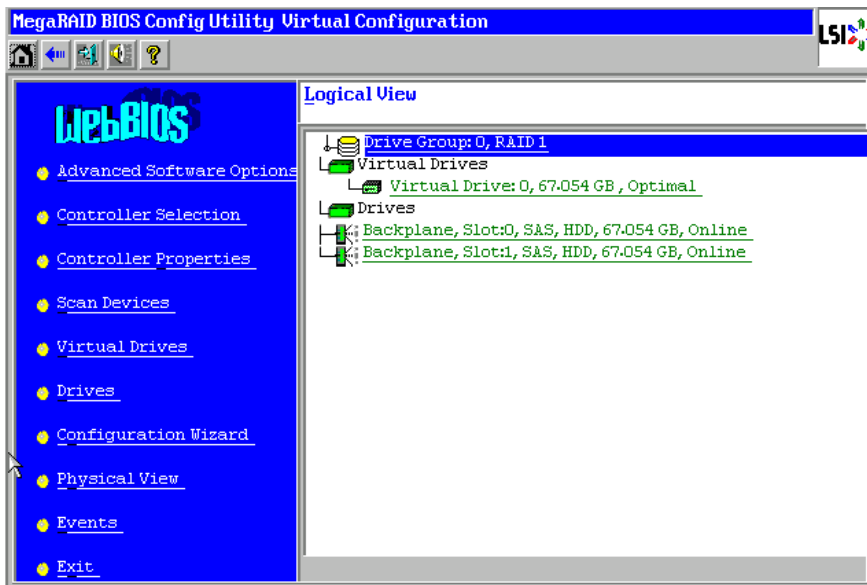
9. Click **Yes** when you are prompted to initialize the new virtual drives.
10. Choose the “Set Boot Drive” radio button for the virtual drive created above and click **Go**.

**Figure 28**      **Setting the Boot Option for Virtual Drives**



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

**Figure 29**      **Virtual Drive Configuration in WebBIOS**



# SQL Server Installation

This document will not go into details and best practices for SQL server installation and configuration. Consult your DBA team to ensure your SQL deployment is configured for best practices according to your corporate standards. For more information on the system requirements for the VMM database, see the Microsoft TechNet link:

<http://technet.microsoft.com/en-us/library/gg610574>

1. Run setup, choose Installation > New Installation...
  1. When prompted for feature selection, install ALL of the following:
  2. Database Engine Services
3. Management Tools – Basic and Complete (for running queries and configuring SQL services)
4. On the Instance configuration, choose a default instance, or a named instance. Default instances are fine for testing and labs. Production clustered instances of SQL will generally be a named instance. For the purposes of the POC, choose default instance to keep things simple.
5. On the Server configuration screen, set SQL Server Agent to Automatic. Click “Use the same account for all SQL Server Services, and input the SQL service account and password (see the section [Customer Configuration Data Sheet, page 138](#)).
6. On the Collation Tab – make sure SQL\_Latin1\_General\_CP1\_CI\_AS is selected, as that is the ONLY collation supported.
7. On the Account provisioning tab – add a domain user account or a group you already have set up for SQL admins.
8. On the Data Directories tab – set your drive letters correctly for your SQL databases, logs, TempDB, and backup.
9. Setup will complete.
10. Apply any service pack or update for SQL 2008 R2 SP1.
11. Once the installation is complete, configure a remote instance for of SQL server for VMM as given in the below URL:
 

<http://technet.microsoft.com/en-us/library/cc764295.aspx>
12. If you are using a domain administrator account or the local system account, SPN (Service Principal Name) for the server is registered in the Active Directory directory service. See the link on Microsoft KB article to register the SPN:
 

<http://support.microsoft.com/kb/909801>

## Microsoft System Center VMM

System Center 2012 - Virtual Machine Manager (VMM) is a management solution for the virtualized datacenter. It enables you to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created. For an overview of System Center 2012 - VMM, see the Microsoft TechNet link:

<http://technet.microsoft.com/en-us/library/gg671827>

This section deals with the installation of the System Center Virtual Machine Manager 2012 on a virtual machine running Microsoft Windows Server 2008 R2 SP1 OS. However, this section does not cover how to create and build a Microsoft Windows Server 2008 R2 with SP1 in a virtual environment.

Before installing a VMM management server, ensure that the computer meets the minimum hardware requirements and that all the prerequisite software is installed. For information about hardware and software requirements for VMM, see the link:

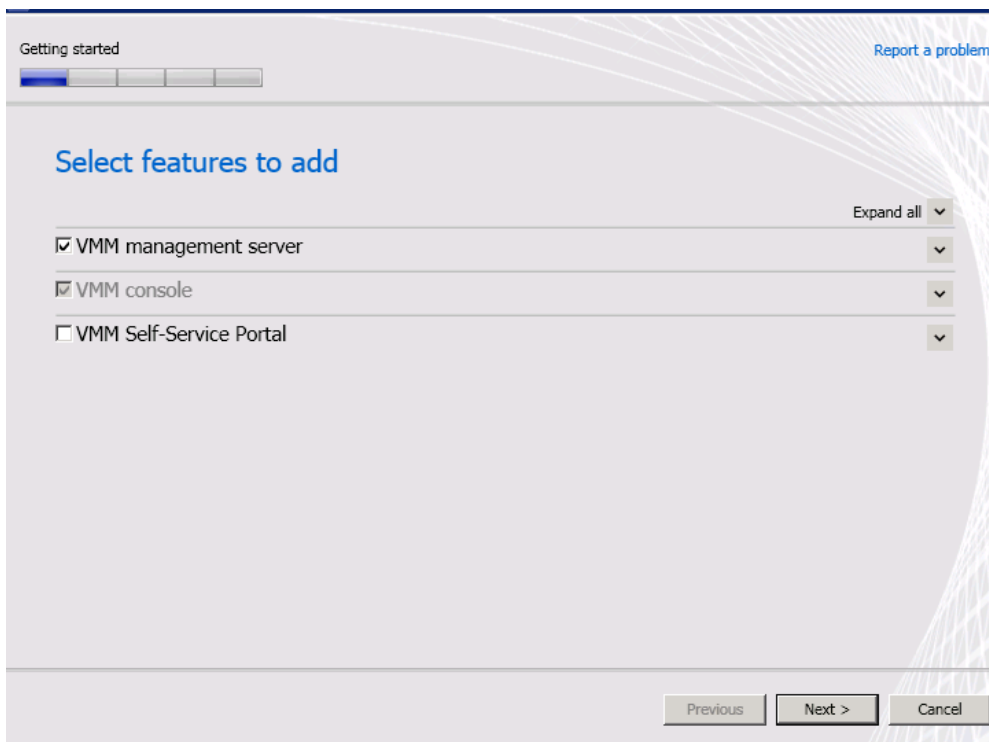
<http://technet.microsoft.com/en-us/library/gg610592>

## Installing the VMM Server and VMM Administrator Console

To install the VMM Server and VMM administrator console, follow these steps:

1. Join the Microsoft Windows Server 2008 R2 virtual machine to the Active Directory domain and login with the domain administrator credentials.
2. To start the Microsoft System Center 2012 Virtual Machine Manager Setup Wizard, on your installation media, right-click **setup.exe**, and click **Run as administrator**.
3. In the main setup page, click **Install**. If you have not installed Microsoft .NET Framework 3.5 SP1, VMM prompts you to install it now.
4. In the “Select features to install” page, choose the **VMM management server** check box and click **Next**.

**Figure 30**      *Selecting Feature in Microsoft SCVMM 2012 Setup Wizard*



5. In the “Installation location” page, use the default path or type a different installation path for the VMM program files and click **Next**.
6. In the “Database configuration” page, enter the database “Server name” and provide the appropriate credentials (See the section [Customer Configuration Data Sheet, page 138](#)). Choose the “New database” radio button and click **Next**.

**Figure 31 Database Configuration in Microsoft SCVMM 2012 Setup Wizard**

**Database configuration**

Provide information about the database that you would like to use for your VMM management server.

Server name:

Port:

☒ Use the following credentials

User name and domain:   
Format: Domain\UserName

Password:

Instance name:

Select an existing database or create a new database.

☒ New database:

☐ Existing database:

7. In the “Configure service account and distributed key management” page, specify the account that will be used by the Virtual Machine Manager Service.

**Figure 32 Configuring Service Account in Microsoft SCVMM 2012 Setup Wizard**

**Microsoft System Center 2012 Virtual Machine Manager Setup Wizard**

Configuration [Report a problem](#)

**Configure service account and distributed key management**

**Virtual Machine Manager Service Account**

Select the account to be used by the VMM service. Highly available VMM installations require the use of a domain account.  
[Which type of account should I use?](#)

☐ Local System account

☒ Domain account

User name and domain:  Password:  

**Distributed Key Management**

Select whether to store encryption keys in Active Directory instead of on the local machine. Highly available VMM installations require the keys be stored in Active Directory.

☐ Store my keys in Active Directory

Provide the location in Active Directory. For example, CN=DKM,DC=contoso,DC=com.

[How do I configure distributed key management?](#)

For more information about which type of account to use, under “Specifying a Service Account for VMM”, see the link:

<http://technet.microsoft.com/library/gg697600.aspx>

8. In the “Port configuration” page, provide unique port numbers for each feature and that are appropriate for your environment and click **Next**.
9. In the “Library configuration” page, choose whether to create a new library share or to use an existing library share on the computer.
10. In the “Installation summary” page, review your selections and click **Install** to install the VMM management server.
11. In the “Setup completed successfully” page, click **Close** to finish the installation.

For more information on installing the System Center 2012 - VMM, see the TechNet article:

<http://technet.microsoft.com/en-us/library/gg610617.aspx>

## Install Microsoft Windows Server on Cisco UCS C220 M3 Servers

This section describes installation of Microsoft Windows Server 2008 R2 SP1 along with driver installation.

### Installation of Microsoft Windows Server 2008 R2 SP1

To install Microsoft Windows Server 2008 R2 Sp1 on all the Cisco UCS C220 M3 bare metal server using the virtual media, follow these steps:

1. Find the drivers for your installed devices on the Cisco UCS C-Series Drivers DVD that came with your C-Series server or download them from: <http://www.cisco.com/cisco/software/navigator.html> and extract them to a local machine such as your laptop.
2. Log in to CIMC Manager using your administrator user ID and password.

**Figure 33** *CIMC Manager Login Page*

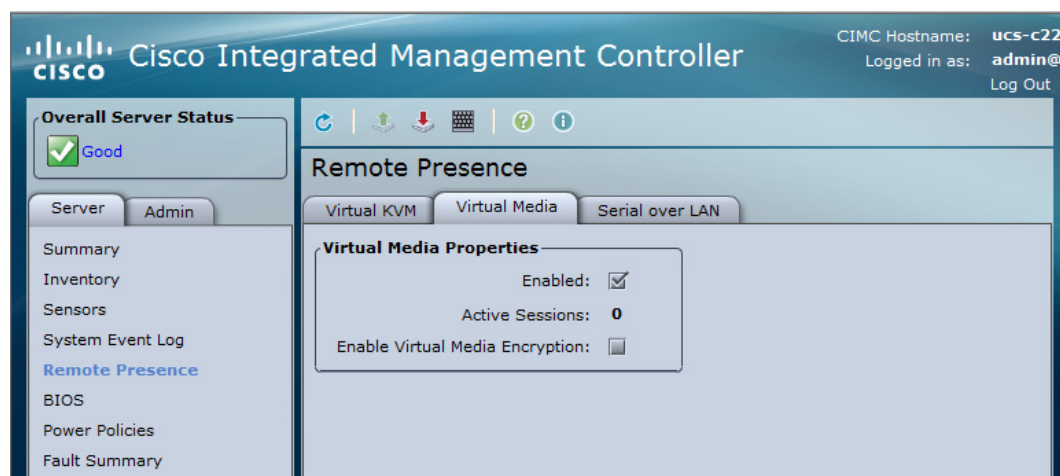


3. Enable the Virtual Media feature, which enables the server to mount virtual drives:
  - a. In the **CIMC Manager Server** tab, click **Remote Presence**.



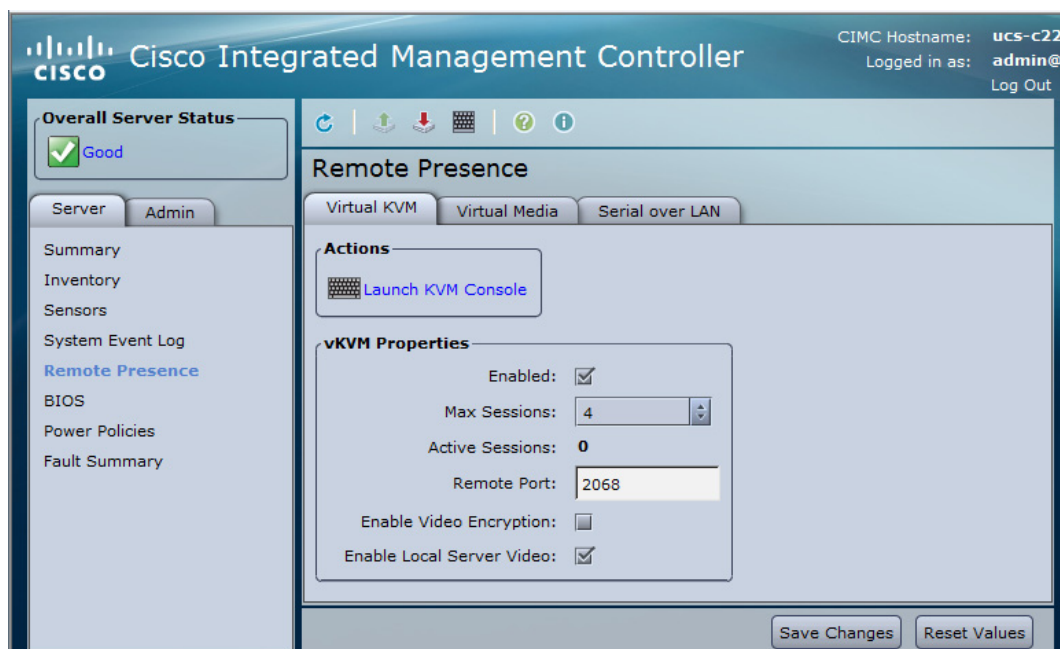
- b. In the “Remote Presence” pane, choose the **Virtual Media** tab and check the **Enable Virtual Media Encryption** check box.
- c. Click **Save Changes**.

**Figure 34**      *Enabling Virtual Media in CIMC*



4. In the “Remote Presence” pane, choose the **Virtual KVM** tab and click **Launch KVM Console**.

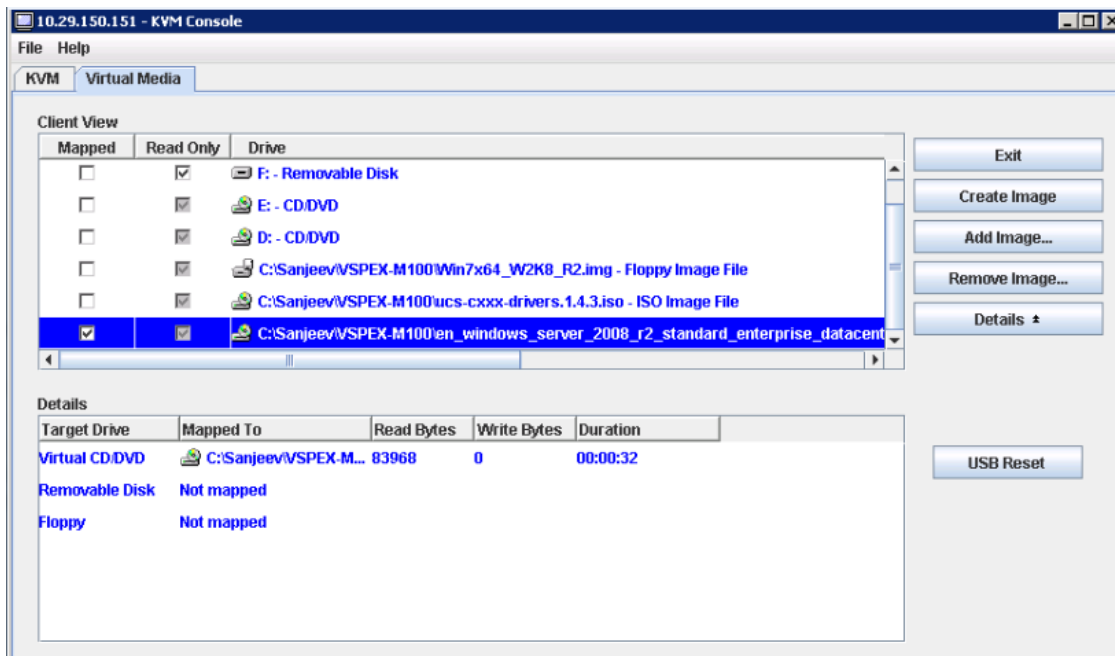
**Figure 35**      *Setting Virtual KVM Properties in CIMC*



5. When the “Virtual KVM Console” window launches, choose the **Virtual Media** tab.
6. In the “Virtual Media” window, provide the path to the Windows installation image by clicking **Add Image**. Use the dialog to navigate to your Microsoft Windows 2008 R2 SP1 ISO file and choose it.

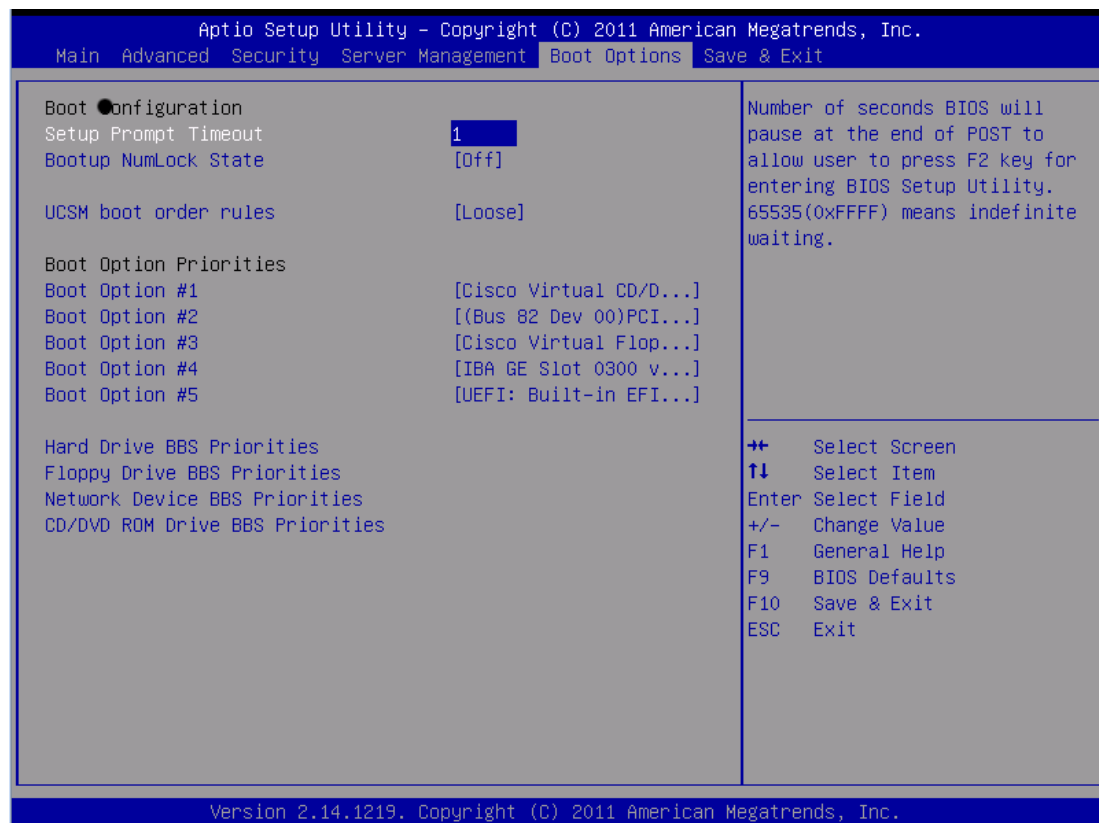
The ISO image is displayed in the “Client View” pane.

**Figure 36** Adding an ISO Image in CIMC



7. When the mapping is complete, power cycle the server so that the BIOS recognizes the media that you just added.
8. In the “Virtual KVM Console” window, watch during bootup for the F2 prompt and then press **F2** to enter the BIOS setup. Wait for the setup utility screen to appear.
9. In the “BIOS Setup utility” screen, choose the **Boot Options** tab and verify that the virtual DVD device that you added in the step 6 is listed as a bootable device.
10. Move the device to the top under “Boot Option Priorities” as shown in [Figure 37](#).

**Figure 37 Cisco UCS C220 M3 BIOS Setup Utility**




- Exit the BIOS Setup utility.

The Microsoft Windows installation begins when the image is booted.

- Press **Enter** when prompted to “boot from CD”.

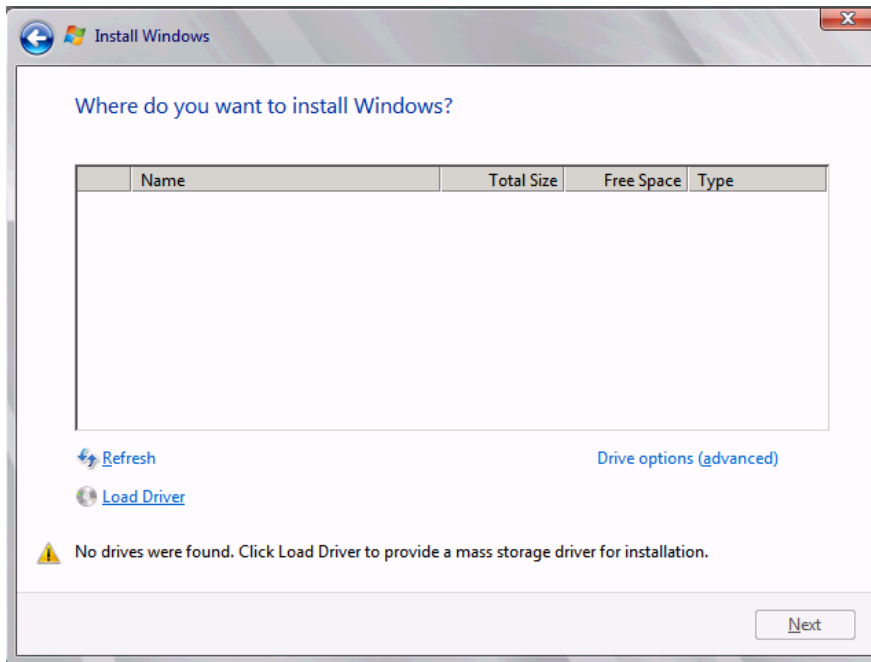
**Figure 38 prompt to Boot from CD**



- Observe the Windows installation process and respond to prompts in the wizard as required for your preferences and company standards.
- When Windows prompts you with “Where do you want to install Windows?”, install the drivers for your mass storage device.

To install the drivers, follow these steps:

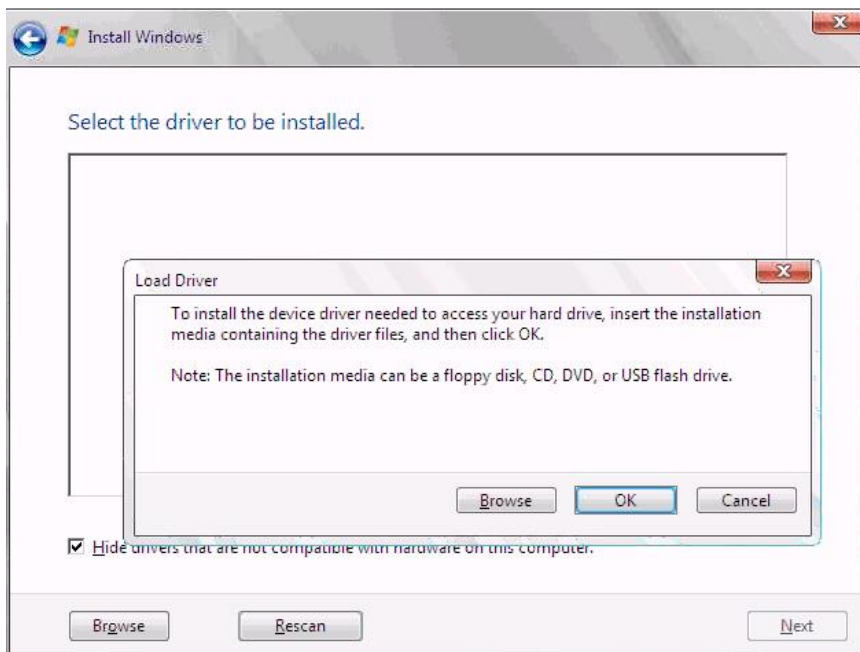
**Figure 39**      **Load Drivers for Installing Microsoft Windows**



- a. In the “Install Windows” window, click **Load Driver**.

You are prompted by a “Load Driver” dialog to choose the driver to be installed. In the next steps, you first define a virtual device with your driver ISO image.

**Figure 40**      **Selecting Drivers for Installing Microsoft Windows**



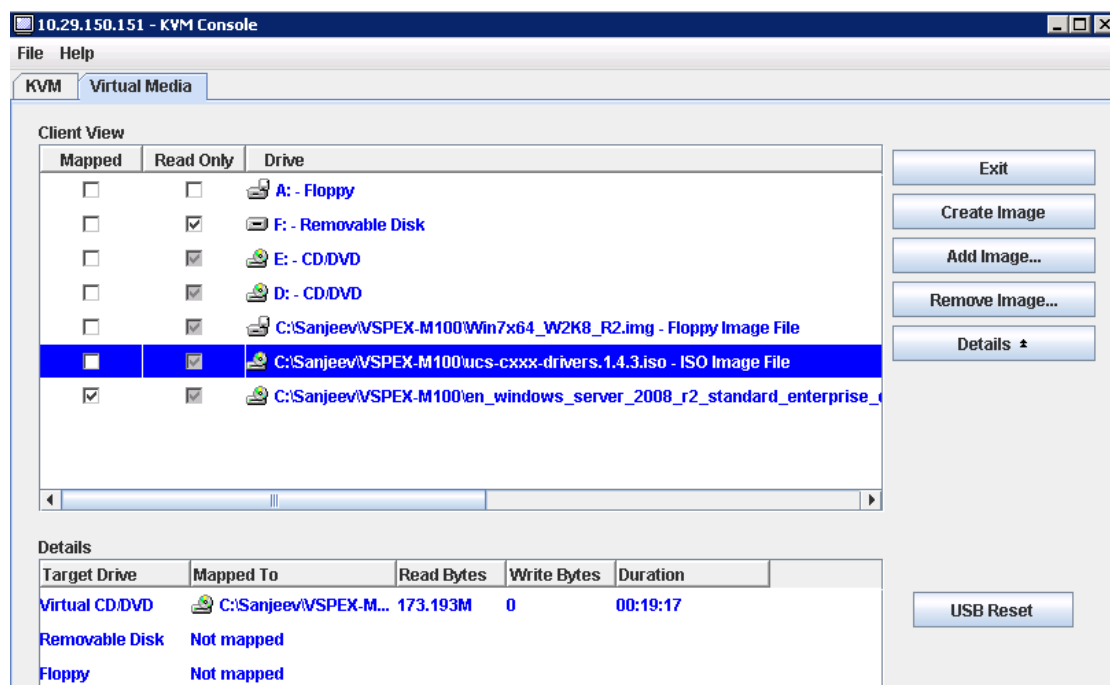
- b. If not already open, open a KVM Virtual Media window as you did in Step 5.

- c. In the “Virtual Media” window, unmount the virtual DVD that you mapped in Step 6 (uncheck the check box under Mapped).
- d. In the “Virtual Media” window, click **Add Image**.
- e. Use the dialog to navigate to the location where you saved the Cisco driver ISO image for your mass storage device in Step 1 and choose it.

The ISO appears in the “Client View” pane.

- f. In the “Virtual Media” window, check the check box under “Mapped” to mount the driver ISO that you just chose. Wait for mapping to complete, as indicated in the “Details” pane. After mapping is complete, you can choose the device for Windows installation.

**Figure 41 Adding an ISO Image in CIMC**



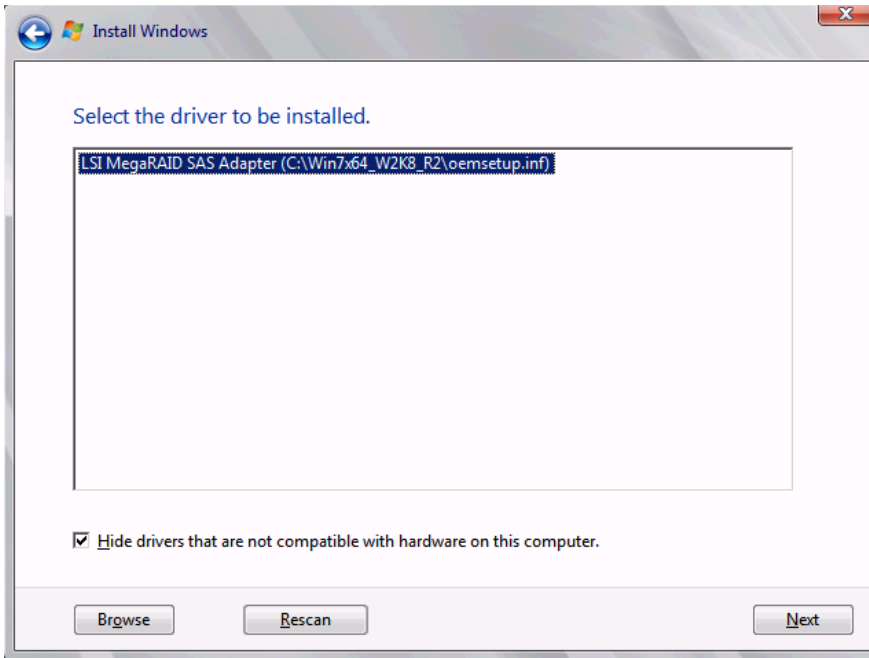
- g. In the “Load Driver” dialog that you opened in Step a, click **Browse**.
- h. Use the dialog to choose the virtual DVD device that you just created.
- i. Navigate to the location of the drivers, choose them, and click **OK**.

Windows loads the drivers and when finished, the driver is listed under the prompt “Select the driver to be installed”.

Driver Path - CDROM Drive:\Windows\Storage\LSI\2008M\W2K8R2\x64

- j. After Windows loads the drivers, choose the driver for your device from the list in the “Install Windows” window and click **Next**. Wait while the drivers for your mass storage device are installed, as indicated by the progress bar.

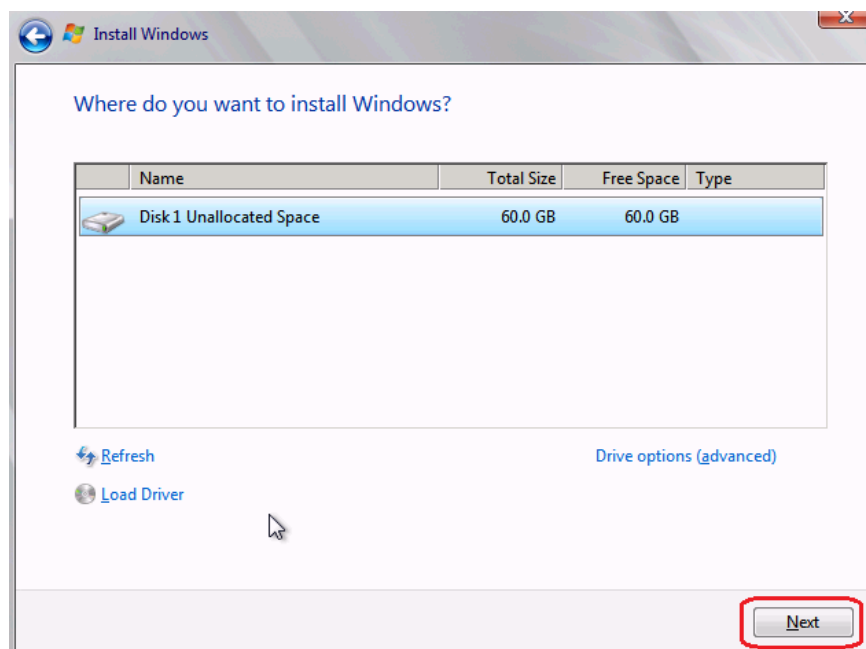
**Figure 42**      **Selecting Drivers for Microsoft Windows Installation**



15. After driver installation finishes, unmap the driver ISO and map the Windows installation image. To unmap the driver ISO and map the Windows installation image, follow these steps:
  - a. In the “Virtual Media” window, uncheck the check box under “Mapped” that corresponds to the driver ISO.
  - b. In the “Virtual Media” window, check the check box under “Mapped” that corresponds to your Windows installation image (the same one that you defined in Step 6).

Wait for the mapping to complete. Observe the progress in the “Details” pane.
  - c. In the “Install Windows” window, choose the disk or partition where you want to install Windows from the list, and then click **Next**.

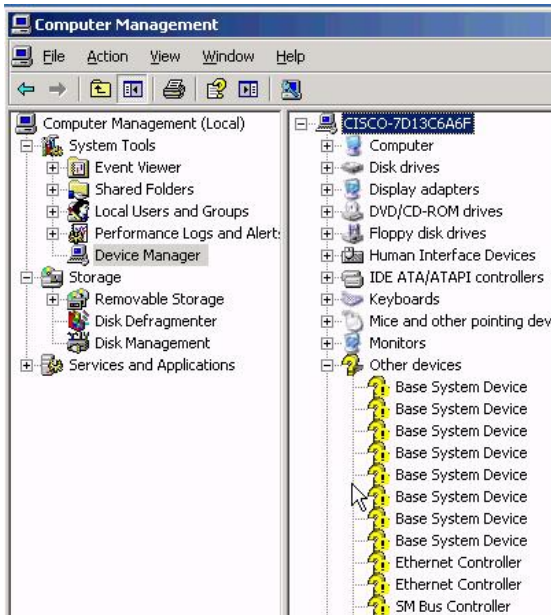
**Figure 43**      **Allocating Disk Space for installing Microsoft Windows**



16. Complete the Windows installation according to the requirements and standards of your company. Continue to observe the Windows installation process and answer prompts as required for your preferences. Verify that Windows lists the drivers that you added.
17. After the Windows installation is complete, Windows reboots the server again and you are prompted to press **Ctrl-Alt-Del** and to log in to access the Windows desktop. Use the login that you supplied during the Windows installation process.


## Device Driver Installation

The Cisco UCS C220 M3 in VSPEX M100 solution contains two Cisco 1GigE I350 LOM (LAN-on-motherboard) and a Cisco VIC P81E adapter for which you need to install the drivers. This section explains how to locate and install the chipset and adapter drivers for Microsoft Windows Server 2008 R2 SP1.

**Figure 44**      **Device Manager**

To locate and install the chipset and adapter drivers for Microsoft Windows Server 2008 R2 SP1, follow these steps:

1. Use a Windows File Manager to navigate to the folder where you extracted the Cisco driver package that you got from the Cisco UCS C-Series Drivers DVD or downloaded from Cisco.com in Step 1 of the [Installation of Microsoft Windows Server 2008 R2 SP1](#) section. Drivers for all of the devices are included in the folders named for each device.
2. Install Intel chipset drivers from the ...\\Windows\\ChipSet\\Intel\\C220\\W2K8R2\\setup.exe and reboot the server.
3. Install the LAN on motherboard (LOM) drivers from ...\\Windows\\Network\\Intel\\I350\\W2K8R2\\x64\\PROWinx64.exe and reboot the server if prompted.

**Figure 45**      **LOM Drive Installation**

4. Install the drivers for Cisco P81E VIC from ...\\D:\\Windows\\Network\\Cisco\\P81E\\W2K8R2\\x64 folder.
5. Repeat the driver installation process for each device that still needs drivers, as indicated by yellow flags in the Microsoft Windows Server 2008 R2 SP1 Device Manager.




## Network Configuration

This section provides steps to configure the NICs and assign IP addresses on all the Windows host servers.

As shown in [Figure 46](#) (for complete diagram, see [Figure 13](#), the two Cisco 1GigE I350 LOM are connected to different switches and teamed at the host side for redundancy and load balancing.

At the Cisco VIC P81E adapter with two 10 GigE uplink ports, four vNICs are created. Cisco Nexus 5500 series switches with Adapter-FEX feature and Cisco VIC P81E in NIV mode provides redundancy at the adapter level. In [Figure 46](#), the vNIC in purple and the vNIC in green shown as solid lines connecting to uplink ports 0 and 1, respectively, indicate active links. The dotted line connection to uplink ports 2 and 1 for these vNICs indicate standby links. It failovers to standby link in the event of active link failing, thus providing redundancy at the adapter level.

**Figure 46** Cisco UCS C220 M3 Network Adapters Configuration



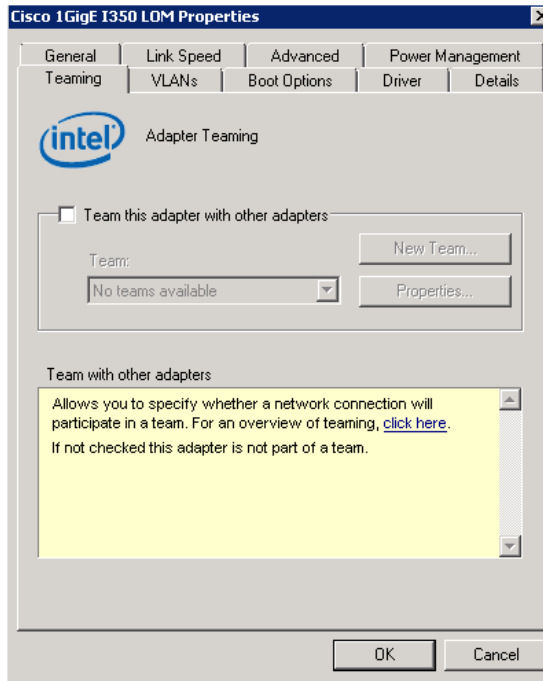
### NIC Teaming of Cisco 1 GigE LOM

Teaming of the Cisco 1 GigE LOMs provides redundancy and doubles the available bandwidth. This teamed adapter is used for Microsoft Hyper-V host management.

For NIC teaming of Cisco 1 GigE LOMs, follow these steps:

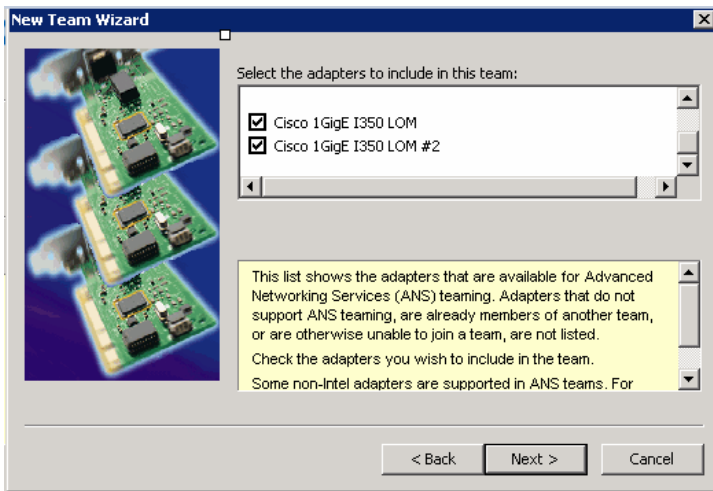
1. Open **Server Manager > Diagnostics > Device Manager**.
2. Right-click the Cisco 1GigE I350 LOM. Choose the **Teaming** tab in the “Cisco 1GigE I350 LOM Properties” window.
3. Check the “Team this adapter with other adapters” check box and click **New Team**.

**Figure 47** *Cisco 1GigE I350 LOM Properties*



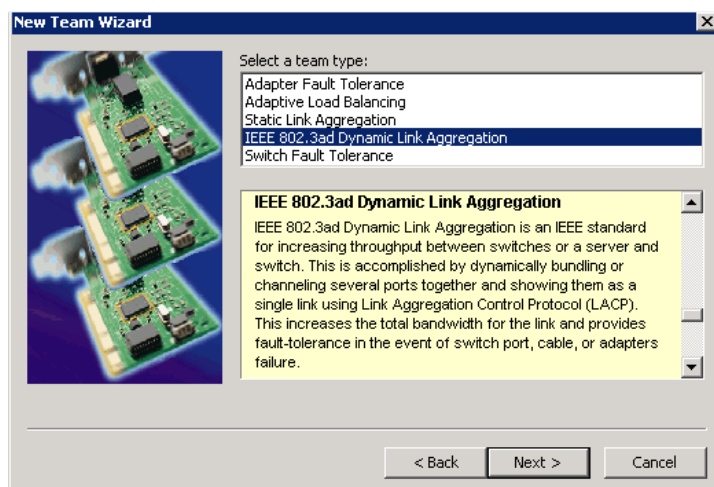
4. Check the “Cisco1GigE I350 LOM # 2” check box in the “New Team Wizard” window.

**Figure 48** *Selecting Adapters for Teaming*



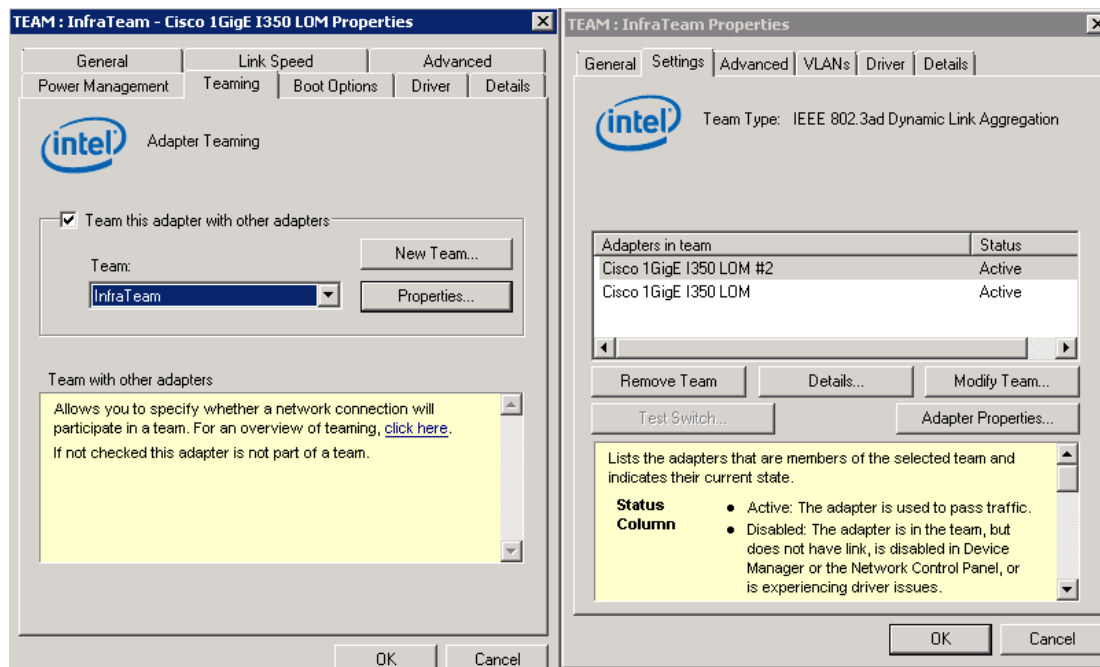
5. Choose “IEEE 802.3ad Dynamic Link Aggregation” from the “Select a team type:” list and click **Next**.

**Figure 49** *Selecting the Type of Adapter Teaming*




A new team will be created as shown in Figure 50.

**Figure 50** *Cisco 1GigE I350 LOM Team Properties*



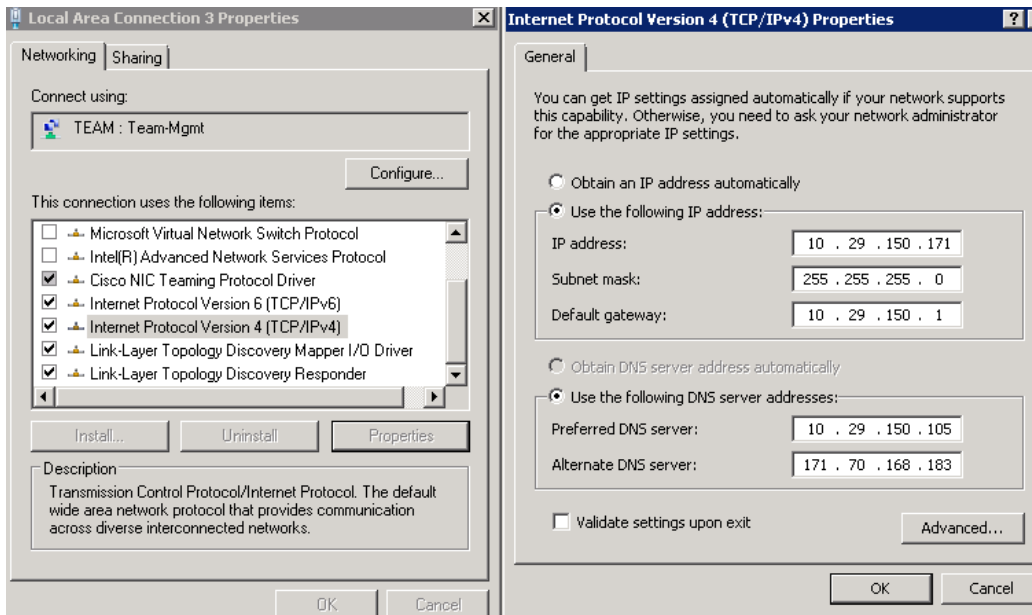
- Choose the teamed adapter, right-click and choose **Properties**.

**Figure 51** *Network Connections page after NIC teaming*



7. Choose **Internet Protocol Version 4 (TCP/IPv4) > Properties** and assign an IP address from the management VLAN (VLAN 1) subnet.

**Figure 52** *Assigning IP Address to the Teamed Adapter*



8. Repeat the above steps to complete the configuration of Cisco 1GigE I350 LOM on all the Cisco UCS C220 M3 servers.

After completion of the above steps successfully, the status of the port-channel summary output on both Cisco Nexus A and B looks like the one shown in [Figure 53](#).

**Figure 53**      **Show Port Channel Summary Output**

```

EMC-5548B# sh port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual    H - Hot-standby (LACP only)
        S - Suspended     R - Module-removed
        s - Switched      r - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
-----
Group  Port-      Type   Protocol  Member Ports
Channel
-----
2      Po2(SU)    Eth    LACP      Eth1/1(P)
3      Po3(SU)    Eth    LACP      Eth1/2(P)
4      Po4(SU)    Eth    LACP      Eth1/3(P)
5      Po5(SU)    Eth    LACP      Eth1/4(P)
7      Po7(SU)    Eth    LACP      Eth1/7(P)      Eth1/8(P)
15     Po15(SU)   Eth    LACP      Eth1/16(P)
17     Po17(SU)   Eth    LACP      Eth1/18(P)
21     Po21(SU)   Eth    LACP      Eth2/1(P)
22     Po22(SU)   Eth    LACP      Eth2/2(P)

```

This Cisco LOM teamed adapter will be used for the host management.

**Note**

Before enabling the Microsoft Hyper-V role, NIC teaming must be completed.

## Creating and Configuring vNICs on Cisco P81E VIC

Cisco P81E VIC supports the Adapter FEX feature and functionality that is enabled on both the Cisco Nexus 5500 series switches. Adapter-FEX can be thought of as a way to divide a single physical link into multiple virtual links or channels. Each channel is identified by a unique channel number and its scope is limited to the physical link. The physical link connects a port on a server network adapter with an Ethernet port on the switch. This allows the channel to connect a vNIC on the server with a Ethernet interface on the switch. Packets on each channel are tagged with a VNTag that has a specific source virtual interface identifier (VIF). The VIF allows the receiver to identify the channel that the source used to transmit the packet.

**Table 10**      **Mapping of Port Profiles on Cisco Nexus Switches with vNICs on Server Adapter**

vNIC	MTU	Uplink	Port-profile Name	Channel Number	Uplink Failover
eth0	9000	Uplink_0	storage	1	Disabled
eth1	9000	Uplink_1	storage	2	Disabled
eth2	1500	Uplink_0	vm_traffic	5	Enabled
eth3	9000	Uplink_1	cluster	6	Enabled

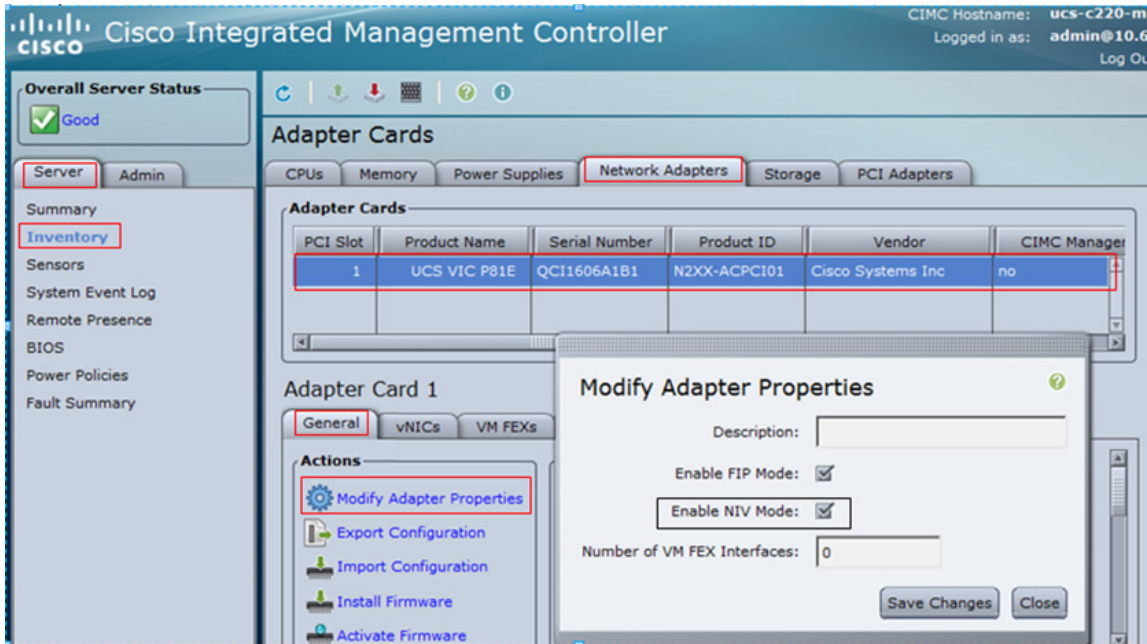
Referring to [Table 10](#), the following steps explain how to modify the existing vNICs properties and create two additional vNICs on Cisco P81E VIC for the VSPEX M100 configuration.

To modify the existing vNICs properties and create the additional vNICs, follow these steps:

1. Using a web browser, connect to the CIMC using the IP address configured in the CIMC Configuration section.
2. Click **Inventory** on the left pane under the **Server** tab and choose the **Network Adapters** tab on the right pane.
3. Choose **UCS VIC P81E** under the “Adapter Cards”.

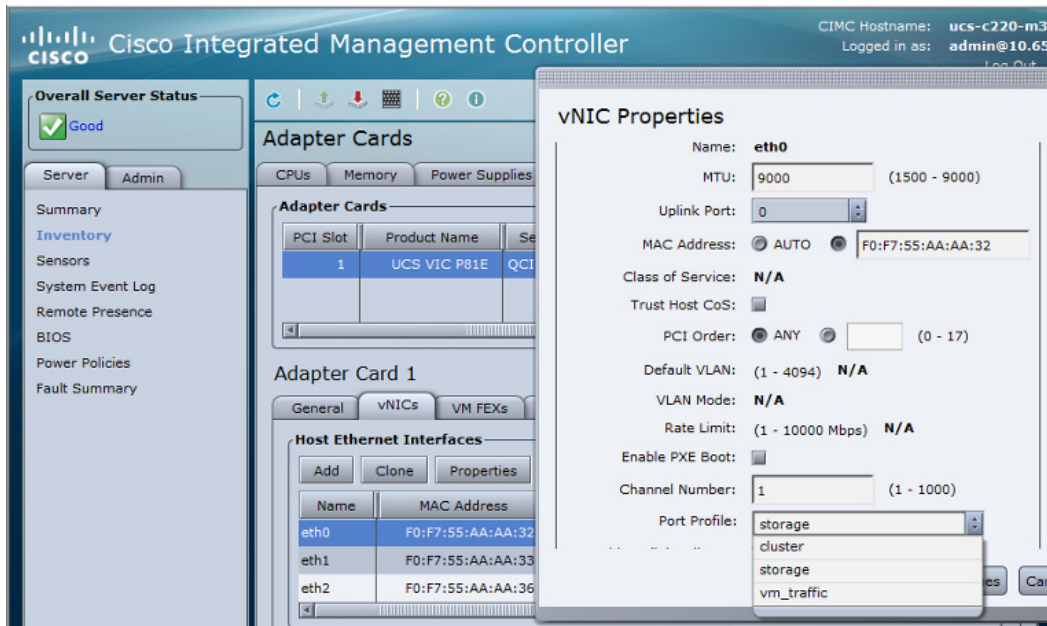
- Click **Modify Adapter Properties** under the **General** tab and enable the Network Interface Virtualization (NIV) mode on the network adapter as shown in Figure 54.

**Figure 54**      *Enabling NIV Mode for Adapter Cards*



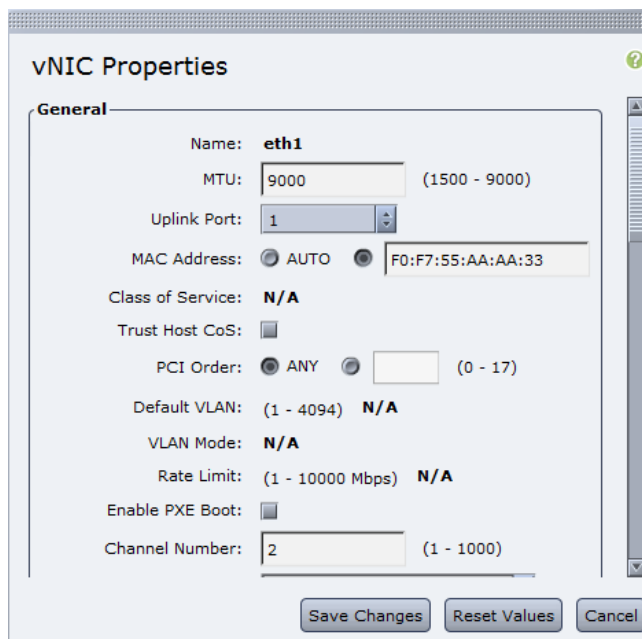
- Reboot the server.
- Once the server is up, repeat the above steps 1 -3. Now choose the **vNICs** tab. Choose **eth0** and click **Properties**.
- Set the MTU size as "9000", uplink port as "0", channel number as "1" in the MTU, Uplink Port, and Channel Number fields respectively. Choose "storage" from the "Port-Profile" drop-down list. Do not "Enable Uplink Failover" for the port-profile name storage.

Figure 55 vNIC Properties



8. Choose **eth1** and click **Properties**.
9. Set the MTU size as “9000”, uplink port as “1”, and channel number as “2” in the MTU, Uplink Port, and Channel Number fields respectively. Choose “storage” from the “Port-Profile” drop-down list. Do not “Enable Uplink Failover” for the port-profile name storage.

Figure 56 CIMC- vNIC Properties 1



10. Create a third vNIC eth2 for the vm\_traffic port-profile by clicking **Add**. Enter a name for the vNIC in the Name field. Set the MTU to “1500”, uplink port to “0” in the MTU and Uplink Port fields respectively. Enter a unique channel number in the Channel Number field.

**Figure 57**      **Adding vNIC**

**Add vNIC**

**General**

Name:

MTU:  (1500 - 9000)

Uplink Port:

MAC Address: ☒ AUTO ☐

Class of Service: **N/A**

Trust Host CoS: ☐

PCI Order: ☒ ANY ☐  (0 - 17)

Default VLAN: (1 - 4094) **N/A**

VLAN Mode: **N/A**

Rate Limit: (1 - 10000 Mbps) **N/A**

Enable PXE Boot: ☐

Channel Number:  (1 - 1000)

11. Scroll down and choose “vm\_traffic” from the “Port-Profile” drop-down list. Check the **Enable Uplink Failover** check box. Enter a value for “Failback Timeout”. Click **Add vNIC**. This vNIC will be used to create virtual switch in Microsoft Hyper-V virtual network manager.



**Figure 58 Adding vNIC 1**

**Add vNIC**

Default VLAN: (1 - 4094) **N/A**

VLAN Mode: **N/A**

Rate Limit: (1 - 10000 Mbps) **N/A**

Enable PXE Boot: ☐

Channel Number: 5 (1 - 1000)

Port Profile: vm\_traffic

Enable Uplink Failover: ☒

Failback Timeout: (0 - 600) 5

**Ethernet Interrupt**

Interrupt Count: 8 (1 - 514)

Coalescing Time: 125 (0 - 65535 us)

Coalescing Type: MIN

Interrupt Mode: MSIX

Add vNIC Reset Values Cancel

12. Similarly, create a fourth vNIC eth3 with the parameters as defined in the [Table 10](#). Eth3 vNIC will be used for live migration in failover cluster setup.
13. Repeat the steps 1 -12 to complete the task on other Cisco UCS C220 M3 Servers.

The switches respond by creating a VevEthernet interface for each vNIC on the server network adapter and associate the port-profile and channel number to the VevEthernet interface. [Figure 59](#) shows the Cisco Nexus switch's partial running-config output showing the veEthernet interfaces created for each vNICs on the four Cisco UCS C220 M3 servers.

**Figure 59** *Running-config Showing veEthernet Interfaces on Cisco Nexus Switch*

```

interface Vethernet32769
  inherit port-profile vm_traffic
  bind interface Ethernet1/9 channel 5

interface Vethernet32770
  inherit port-profile vm_traffic
  bind interface Ethernet1/11 channel 5

interface Vethernet32771
  inherit port-profile vm_traffic
  bind interface Ethernet1/12 channel 5

interface Vethernet32772
  inherit port-profile vm_traffic
  bind interface Ethernet1/10 channel 5

interface Vethernet32773
  inherit port-profile storage
  bind interface Ethernet1/9 channel 2

interface Vethernet32775
  inherit port-profile storage
  bind interface Ethernet1/10 channel 2

interface Vethernet32776
  inherit port-profile cluster
  bind interface Ethernet1/10 channel 6

interface Vethernet32778
  inherit port-profile storage
  bind interface Ethernet1/11 channel 2

interface Vethernet32779
  inherit port-profile cluster
  bind interface Ethernet1/11 channel 6

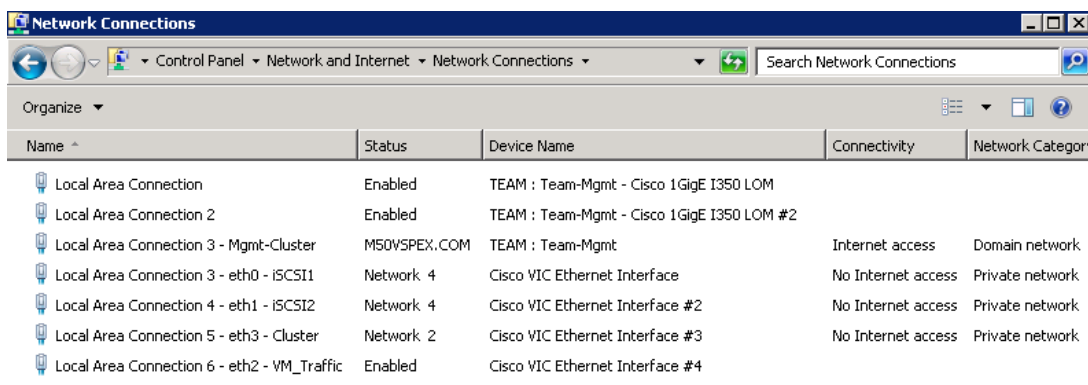
interface Vethernet32781
  inherit port-profile storage
  bind interface Ethernet1/12 channel 2

interface Vethernet32782
  inherit port-profile cluster
  bind interface Ethernet1/12 channel 6

interface Vethernet32784
  inherit port-profile cluster
  bind interface Ethernet1/9 channel 6

```

14. Assign static IP addresses to the NICs as per their VLAN participation. Ping to verify if the assigned IP addresses are working properly. (Optionally, rename the NICs on all the other servers for easy identification and make sure they are consistent to avoid any problems with clustering.

**Figure 60** *Renaming the NICs for Easy Identification*


Name	Status	Device Name	Connectivity	Network Category
Local Area Connection	Enabled	TEAM : Team-Mgmt - Cisco 1GigE I350 LOM		
Local Area Connection 2	Enabled	TEAM : Team-Mgmt - Cisco 1GigE I350 LOM #2		
Local Area Connection 3 - Mgmt-Cluster	M50VSPEX.COM	TEAM : Team-Mgmt	Internet access	Domain network
Local Area Connection 3 - eth0 - iSCSI1	Network 4	Cisco VIC Ethernet Interface	No Internet access	Private network
Local Area Connection 4 - eth1 - iSCSI2	Network 4	Cisco VIC Ethernet Interface #2	No Internet access	Private network
Local Area Connection 5 - eth3 - Cluster	Network 2	Cisco VIC Ethernet Interface #3	No Internet access	Private network
Local Area Connection 6 - eth2 - VM_Traffic	Enabled	Cisco VIC Ethernet Interface #4		

15. Ping to validate if the assigned IPs are working properly. [Figure 61](#) shows ping status from a host to both storage processors SP A (10.10.40.50) and SP B (10.10.40.60) with jumbo frames.

**Figure 61** Jumbo Frame Validation from Source Switch to Destination Storage

```

Administrator: Command Prompt
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\administrator.M50USPEX>hostname
M100N1

C:\Users\administrator.M50USPEX>ping 10.10.40.50 -f -l 8972

Pinging 10.10.40.50 with 8972 bytes of data:
Reply from 10.10.40.50: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.50: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.50: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.50: bytes=8972 time<1ms TTL=255

Ping statistics for 10.10.40.50:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\administrator.M50USPEX>ping 10.10.40.60 -f -l 8972

Pinging 10.10.40.60 with 8972 bytes of data:
Reply from 10.10.40.60: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.60: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.60: bytes=8972 time<1ms TTL=255
Reply from 10.10.40.60: bytes=8972 time<1ms TTL=255

Ping statistics for 10.10.40.60:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\administrator.M50USPEX>

```

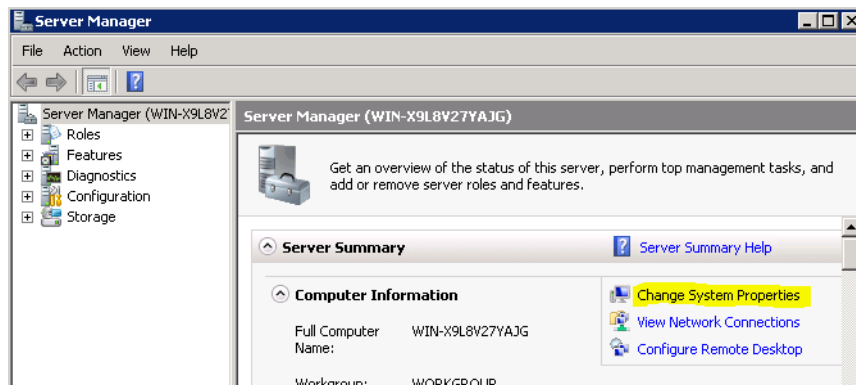
## Host Rename and Domain Join

This section covers step-by-step instructions to rename and join the hosts to a domain.

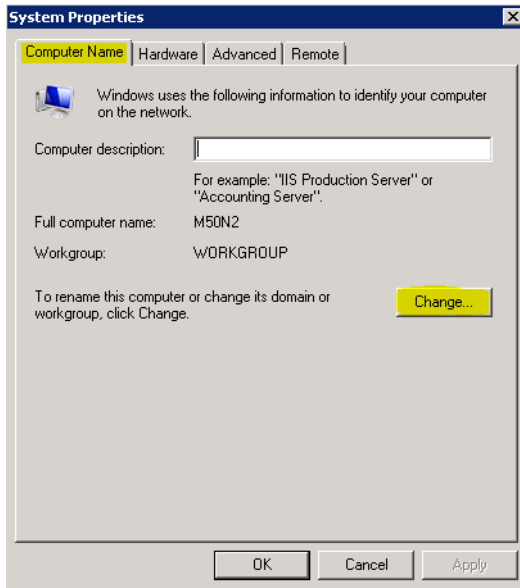
To rename and join the hosts to a domain, follow these steps:

1. Rename the Windows hostname on all the servers as per your naming convention. The four servers' hostnames are in this document are M100N1, M100N2, M100N3, and M100N43.
2. Open the Server Manager and click **Change System Properties**.

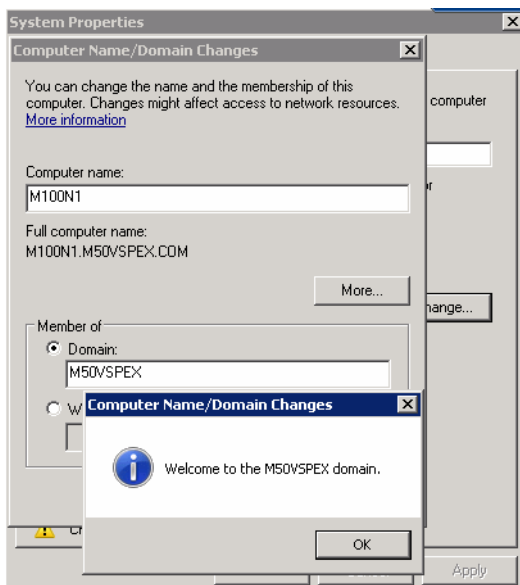
**Figure 62** Changing System Properties in Server Manager Window



3. Click **Change** under the **Computer Name** tab in the “System Properties” window.

**Figure 63**      **System Properties Window**

4. Under “Member of”, choose the “Domain” radio button. Type the name of the domain you want your servers to join and click **OK**.
5. Provide the appropriate credentials in the “Windows Security” pop-up screen and click **OK**.

**Figure 64**      **Specifying the Domain**

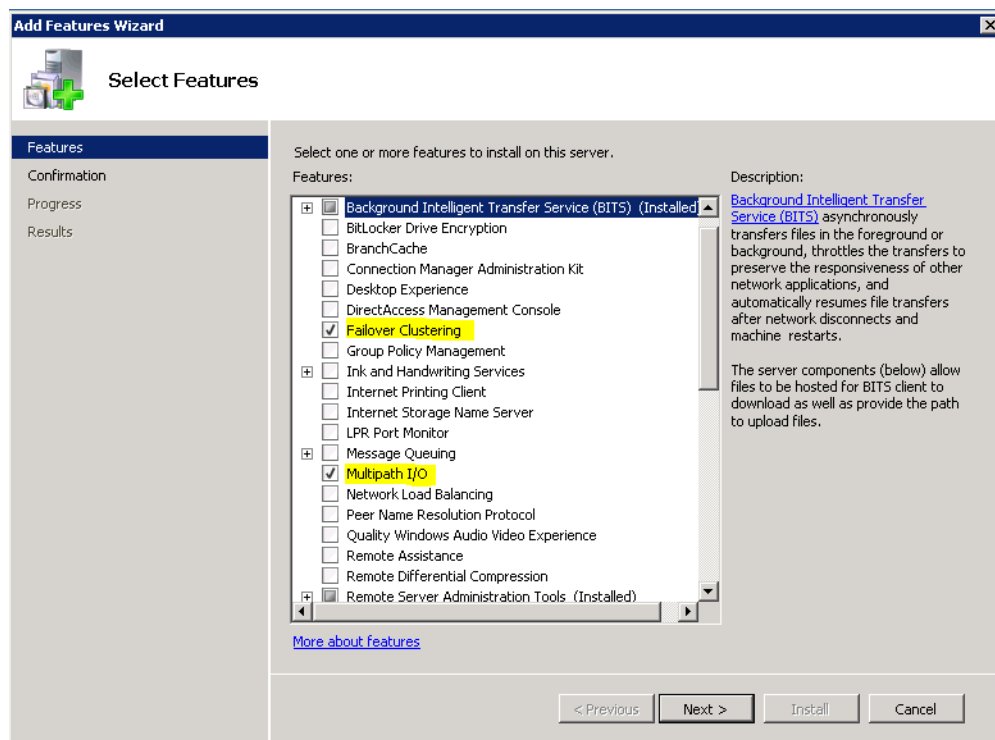
6. Click **OK** in the “Welcome” screen and then restart the server to apply these changes.
7. Login to the server and apply/update any latest hotfixes on all the Windows Server 2008 R2 SP1.
8. Repeat the above steps on all other servers to rename the host and join to the domain.

## Install Roles and Features

To install roles and features, follow these steps:


1. Click **Start > Server Manager**.
2. In the Features selection page of the “Add Features Wizard”, check the “Failover Clustering” and “Multipath I/O” check boxes. Click **Next**.
3. In the “Confirmation selection” page, click **Install**.

**Figure 65** Adding Features to the Server



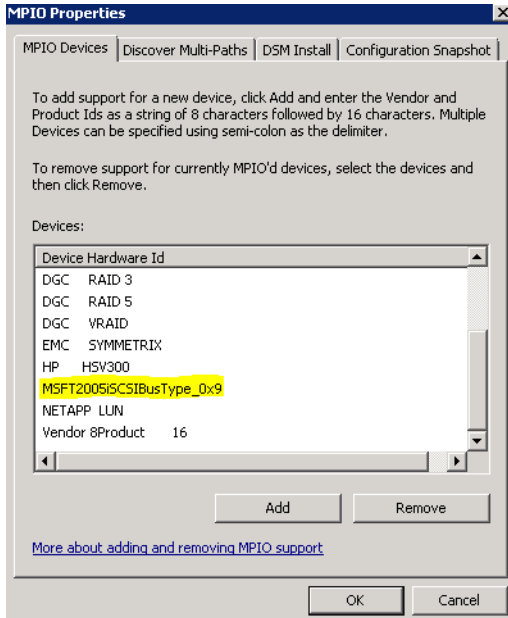
4. Click **Start > Administrative Tools > MPIO**.
5. Choose the **Discover Multi-Paths** tab, check the “Add support for iSCSI devices” check box, and click **Add**. You will be prompted to reboot the server.

**Figure 66** Enabling iSCSI Support in MPIO Properties Window



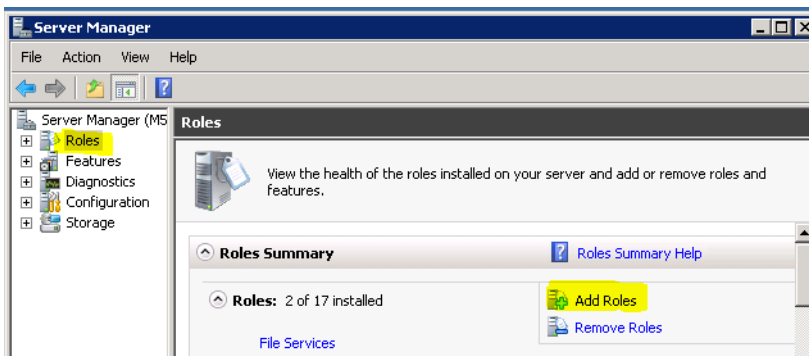
6. After rebooting, if you open the MPIO Control Panel applet, you should see the iSCSI bus listed as a device.

**Figure 67** *MPIO Properties Window*



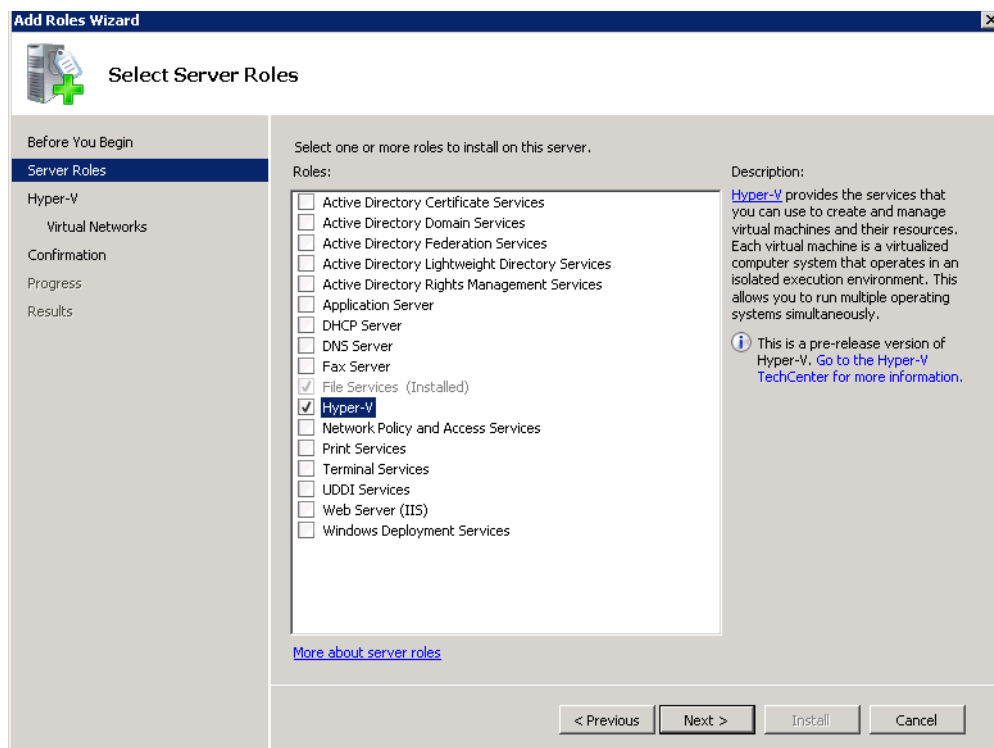
7. Click **Start > Server Manager**.
8. In the “Roles Summary” view of the “Server Manager” window, click **Add Roles**.

**Figure 68** *Adding Roles in the Server Manager Window*



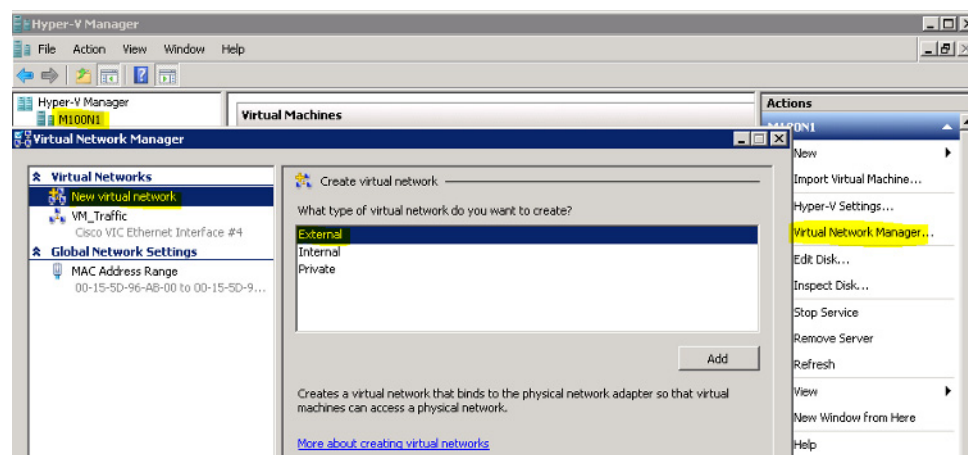
9. In the “Server Roles” selection page, check the “Hyper-V” check box and click **Next**.

**Figure 69**      **Selecting Roles for the Server**



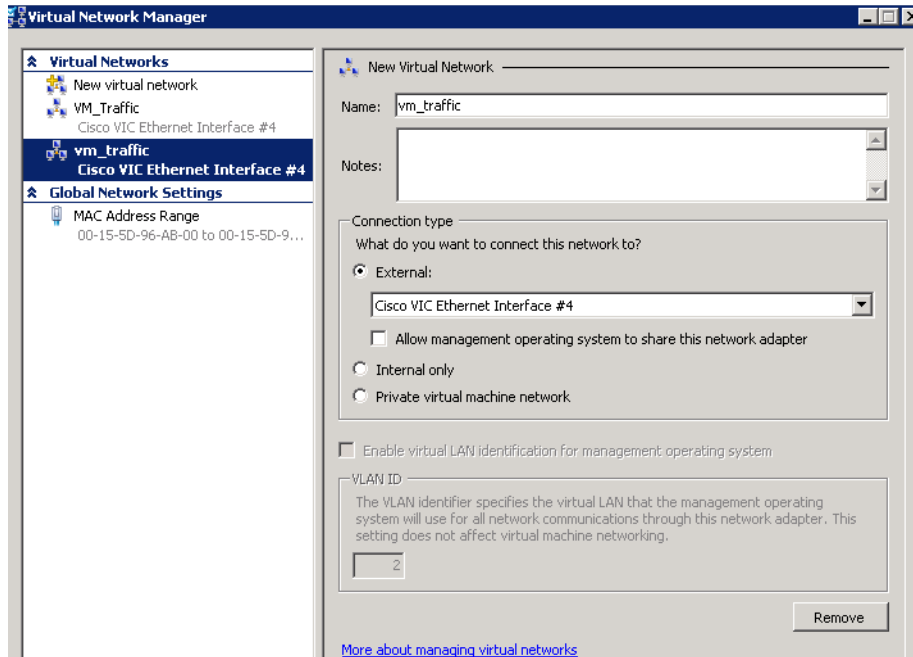
10. In the **Hyper-V > Virtual Networks** page, click **Next** without choosing any Ethernet cards.
11. In the “Confirmation Selection” page, click **Install** and reboot the server when prompted.
12. Login to the server and click **Start > Administrative Tools > Hyper-V Manager**.
13. In the “Hyper-V Manager”, choose the server and click **Virtual Network Manager**.
14. In “New virtual network”, choose **External** from the “What type of virtual network do you want to create” list and click **Add**.

**Figure 70**      **Adding New Virtual Networks**



15. Type a name and choose the “External” radio button. From the drop-down list, choose the NIC assigned to vm\_traffic VLAN and click **Next**.

**Figure 71**      **Add Virtual Networks in Virtual Network Manager 1**



16. Repeat the above steps on all other servers to install the roles and features.

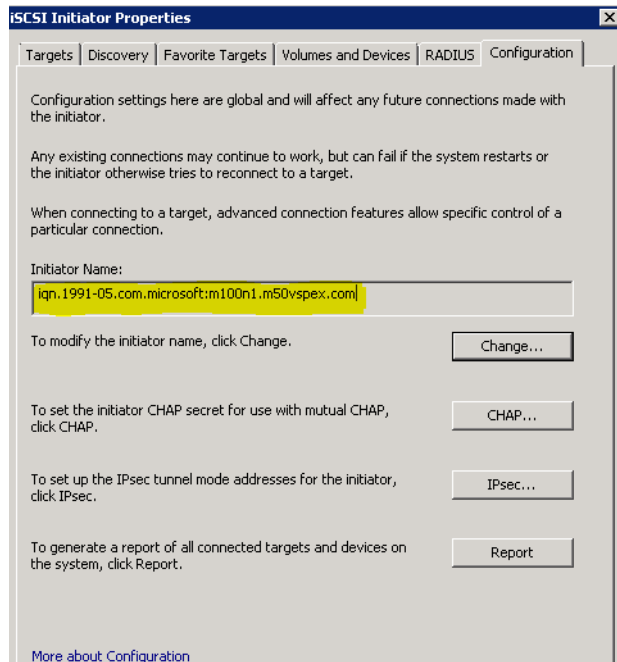
## Enable iSCSI Initiator

To enable iSCSI initiator, follow these steps:

1. Enable the iSCSI Initiator by clicking **Start > Administrative Tools > iSCSI initiator**. Click **Yes** to start the Microsoft iSCSI service.
2. Repeat the above step on all other servers and note down the initiator name. This is required to create the hosts.



**Figure 72**      **Configuring iSCSI Initiator**



## Prepare the EMC VNXe3300 Storage

This section explains the following topics:

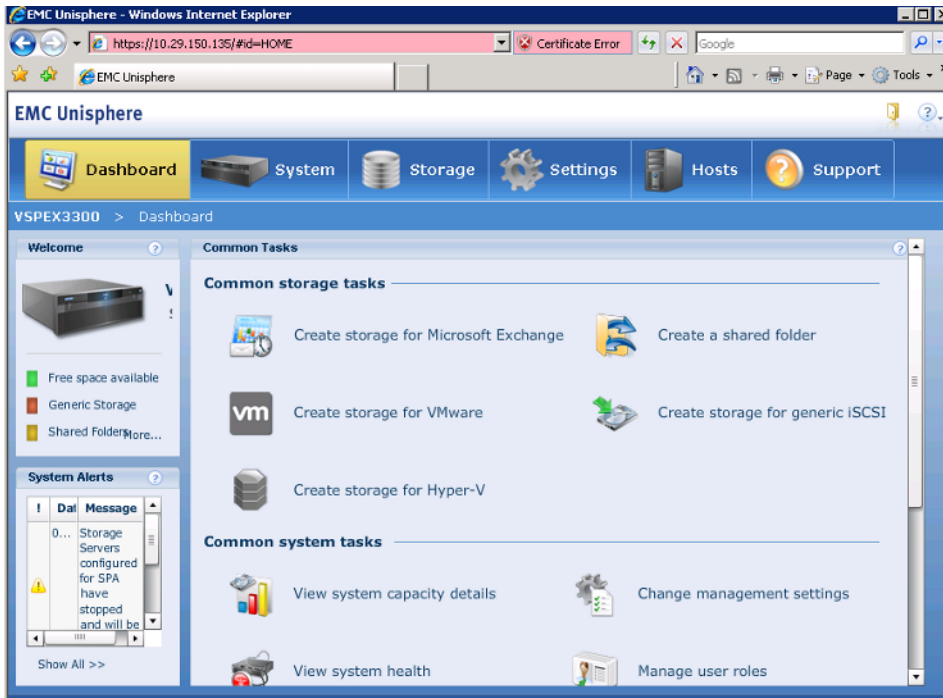
- Preparing the storage array
- Aggregating data ports for high-availability
- Creating storage pools for Hyper-V datastores
- Creating iSCSI server and assigning host access privileges

### Initial Setup of EMC VNXe

To configure the initial setup of the EMC VNXe, follow these steps:

1. Connect the eEthernet cables from the management and data ports to the network as shown in [Figure 72](#).
2. Assign an IP address to the management interface or Download and run the Connection Utility to establish an IP address for managing the EMC 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. Connect to the EMC VNXe system through a web browser using the management IP address.

**Figure 73** *EMC Unisphere GUI*



**Note**

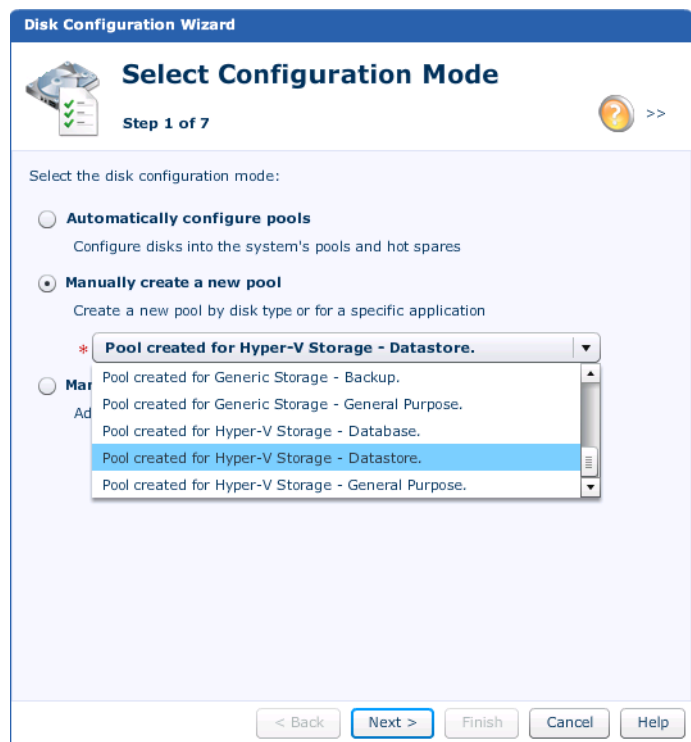
The SP A and SP B network data ports must be connected to the same subnet. In general, both SPs should have mirrored configurations for all front-end cabling (including VLANs) in order to provide Failover.

## Create Storage Pools

Create a pool with the appropriate number of disks as shown in the EMC storage layout in [Figure 74](#).

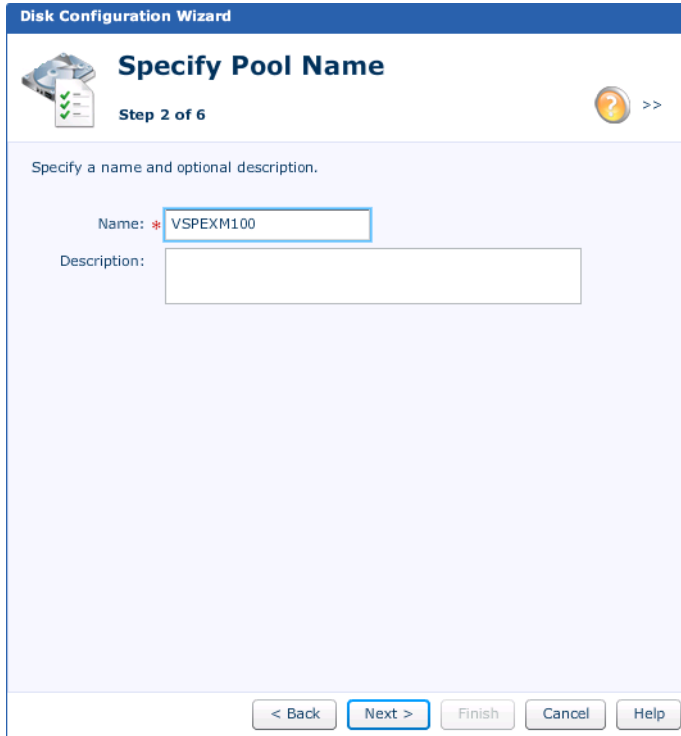
1. In Unisphere, choose **System > Storage Pools**.
2. Choose **Configure Disks**.
3. Manually create a new pool by **Disk Type** for SAS drives in the “Select Configuration Mode” screen of the “Disk Configuration Wizard”.

**Figure 74**      **Selecting the Disk Configuration Mode**



4. Choose “Pool created for Hyper-V Storage - Datastore”.
5. Specify a name for the pool in the “Name” field.

**Figure 75** Specifying Pool Name in the Disk Configuration Wizard



Disk Configuration Wizard

## Specify Pool Name

Step 2 of 6

Specify a name and optional description.

Name: \* VSPEXM100

Description:

< Back Next > Finish Cancel Help

6. Choose “Balanced Perf/Capacity” in the “Select Storage Type” window. The validated configuration uses a single pool with 77 drives.

**Figure 76** *Selecting Storage Type in the Disk Configuration Wizard*

**Disk Configuration Wizard**

**Select Storage Type**

Step 3 of 6

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

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

Rating	Disk Type	Max Capacity	Storage Profile
☆☆☆	SAS	0 GB (None Available)	Balanced Perf/Capac
☆☆	SAS	0 GB (None Available)	High Performance
☆	EFD	0 GB (None Available)	Best Performance
	NL SAS	0 GB (None Available)	High Capacity

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

Hyper-V SAS storage pool using RAID 5(6+1).

< Back   **Next >**   Finish   Cancel   Help

Figure 77 shows the details of the storage pool you have created.

**Figure 77** *Window Showing Storage Pools Details*

**EMC Unisphere**

Dashboard   **System**   Storage   Settings   Hosts   Support

VSPEX3300 > System > Storage Pools > Storage Pools Details

**Storage Pools Details**

**Summary**

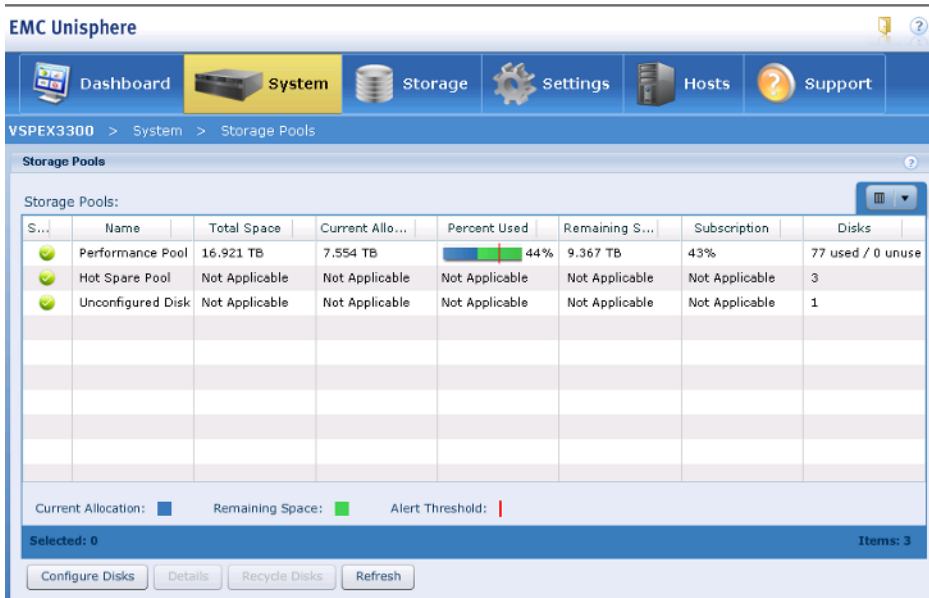
Status: ✔ Ok   Subscription: 43%  
 Name: Performance Pool   Alert Threshold: 70%  
 Remaining Space: 9.367 TB  
 Total Space: 16.921 TB

**General**   Utilization   Disks   Pool History View

Name: Performance Pool  
 Description:   
 Type: SAS  
 Number of Disks: 77  
 Disk Speed: 6 Gbps  
 Pool Type: RAID 5

Apply Changes   Cancel Changes   Add Disks   Recycle Disks

**Figure 78** Window Showing Storage Pools in EMC Unisphere



**Note**

You should also create your Hot Spare disks at this point.



**Note**

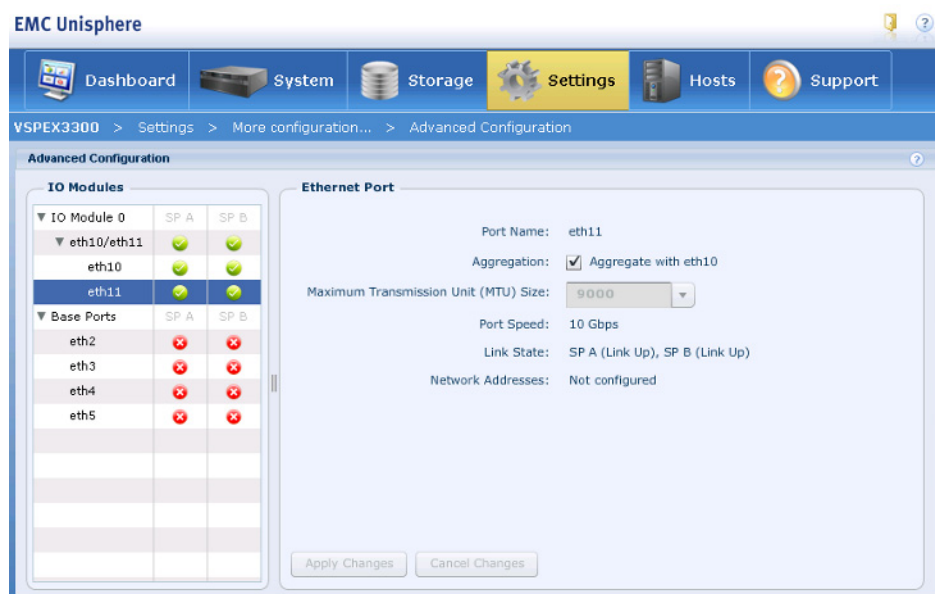
As a performance best practice, all the drives in the pool should be of the same size.

## Configure Advanced Features—Link Aggregation and Jumbo Frames

To configure advanced features – link aggregation and jumbo frames, follow these steps:

1. In Unisphere choose **Settings > More Configuration > Advanced Configuration**.
2. Choose **eth10** and set the MTU size to “9000” in the MTU field.
3. Choose **eth11** and set the MTU size to “9000” in the MTU field.
4. Check the “Aggregate with eth10” check box.

**Figure 79**      **Advanced Configuration Window in EMC Unisphere**



## Create iSCSI Servers

To create iSCSI servers, follow these steps:

1. In Unisphere, choose **Settings > iSCSI Server Settings > Add iSCSI Server**.
2. Type the server name, IP address, subnet mask and default gateway for the new iSCSI server in the Server Name, IP address, Subnet Mask and Default Gateway fields respectively. Choose “SP A” from the “Storage Processor” drop-down list and choose the aggregated ports from the “Ethernet Port” drop-down list as shown in [Figure 80](#).

**Figure 80**      **Adding iSCSI Server for Storage Processor - A**

**iSCSI Server**

Step 1 of 4

Specify the Network Interface for the new iSCSI Server:

Server Name: \* iSCSIServer00

IP Address: \* 10 . 10 . 40 . 50

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

Default Gateway: 10 . 10 . 40 . 1

[Hide advanced](#)

Storage Processor: **SP A**

Ethernet Port: **eth10/eth11 (Link Up)**

VLAN ID: 0 <click to edit>

< Back   Next >   Finish   Cancel   Help

3. Click **Add iSCSI Server** again. Type the server name, IP address, subnet mask and default gateway for the new iSCSI server in the Server Name, IP address, Subnet Mask and Default Gateway fields respectively. Choose “SP B” from the “Storage Processor” drop-down list and choose the aggregated ports from the “Ethernet Port” drop-down list as shown in [Figure 81](#).



**Figure 81** Adding iSCSI Server for Storage Processor - B

**iSCSI Server**

Step 1 of 4

Specify the Network Interface for the new iSCSI Server:

Server Name: \* iSCSIServer01

IP Address: \* 10 . 10 . 40 . 60

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

Default Gateway: 10 . 10 . 40 . 1

[Hide advanced](#)

Storage Processor: SP B

Ethernet Port: eth10/eth11 (Link Up)

VLAN ID: 0 <click to edit>

< Back Next > Finish Cancel Help

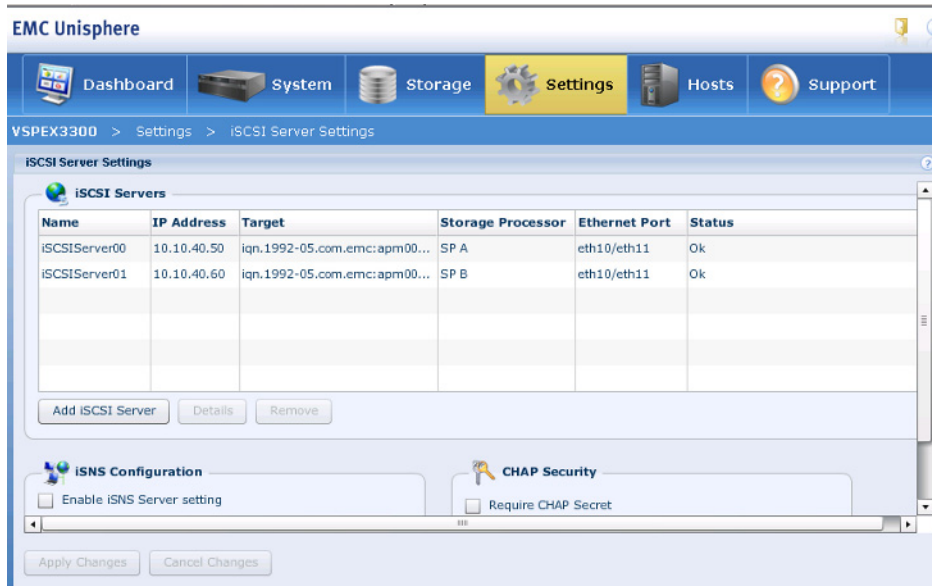


**Note**

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

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

**Figure 82** *iSCSI Server Settings Window in EMC Unisphere*

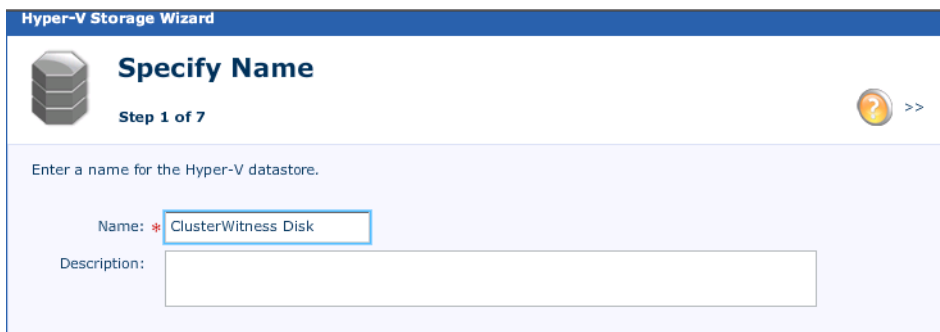


## Create Microsoft Hyper-V Datastores

To create Hyper-V datastores, follow these steps:

1. In Unisphere, choose **Storage > Hyper-V > Create**.
2. Specify a name for the Hyper-V datastore in the Name field and click **Next**.

**Figure 83** *Specifying Hyper-V Datastore Name*



3. Choose the pool and iSCSI server. Enter “10GB” in the Size field. Do not enable Thin Provisioning and click **Next**.



### Note

The default size of the Hyper-V datastore is 100 GB. The maximum size possible is 1.999 TB and the minimum size required is 10 GB.

**Figure 84**      **Configuring Storage Pool and Size for the Datastore**

Hyper-V Storage Wizard

## Configure Storage

Step 2 of 7

Configure the storage pool and size for this Hyper-V datastore:

Type	Pool	Server	Available	Percent Used	Subscription
SAS	Performance Pool	iSCSIServer00	8.973 TB	<div><div></div></div> 45%	43%
SAS	Performance Pool	iSCSIServer01	8.973 TB	<div><div></div></div> 45%	43%

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

Size: \*  GB

Thin: ☐ Enabled

< Back   Next >   Finish   Cancel   Help

4. Change protection by choosing the “Do not configure protection storage for this storage resource” radio button. If you need additional protection, then additional storage would be required. For more information, see the EMC VNXe Storage Configuration Guide.

**Figure 85**      **Setting Configure Protection**

Hyper-V Storage Wizard

## Configure Protection

Step 3 of 7

Configure protection storage for replication and snapshots:

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

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

< Back   Next >   Finish   Cancel   Help

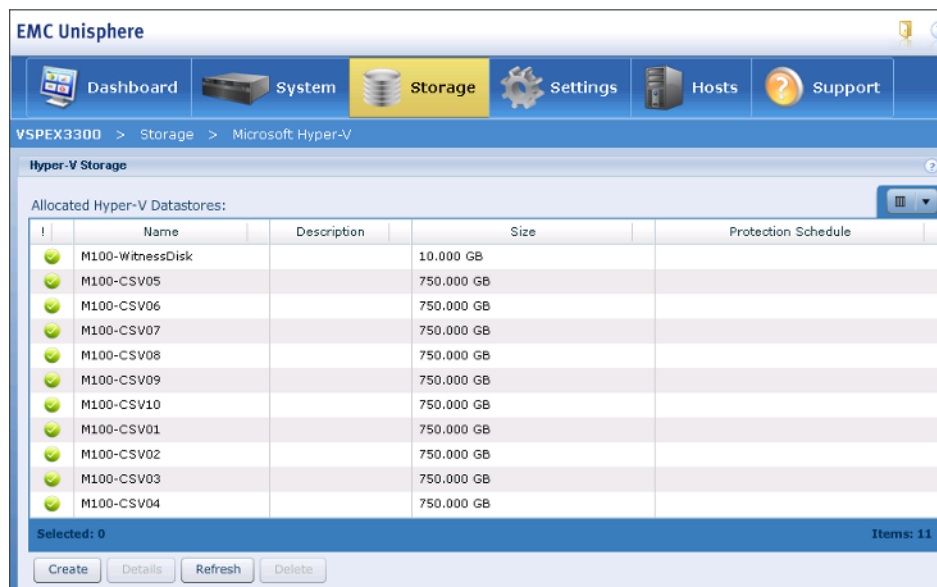
5. Click **Next** in the “Configure Host Access” screen as no host are connected. “Configure Host Access” will be completed in the EMC VNXe3300 Deployment Procedure in the later section.
6. Repeat the above steps and create 10 Hyper-V datastores for Cluster Shared Volumes of 750GB size.



**Note**

We recommend 750 GB size Hyper-V datastore for CSV. You can create bigger size datastores.

**Figure 86** *Hyper-V Datastores for Cluster Shared Volume*

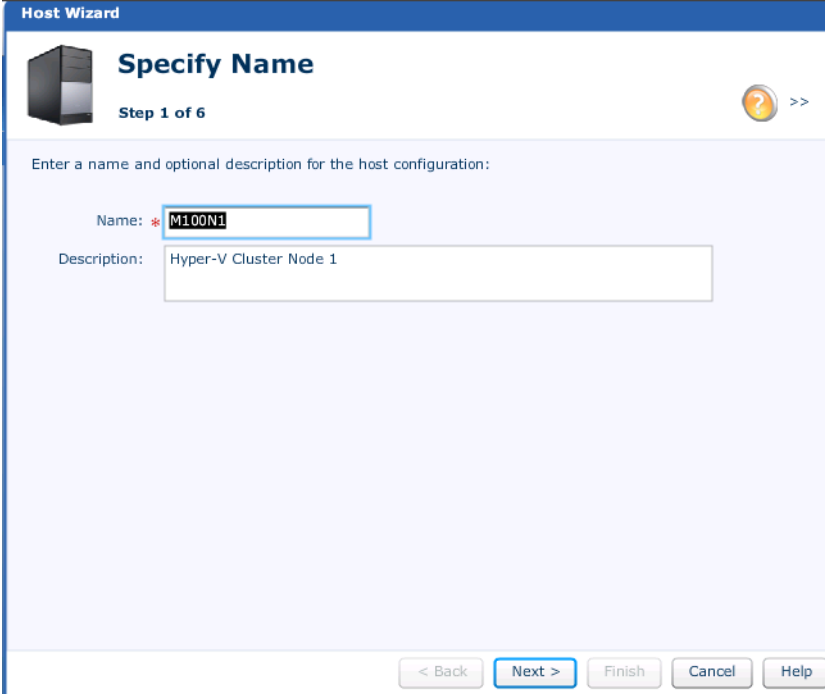


## Create Hosts and Provide Host Access

To create hosts and provide host access, follow these steps:

1. In EMC Unisphere, click the **Hosts** tab and click **Create Host**.  
The “Host Wizard” page appears.
2. In the Name and Description fields, type the name and description of the new host. Click **Next**.  
The “Operating System” page appears.

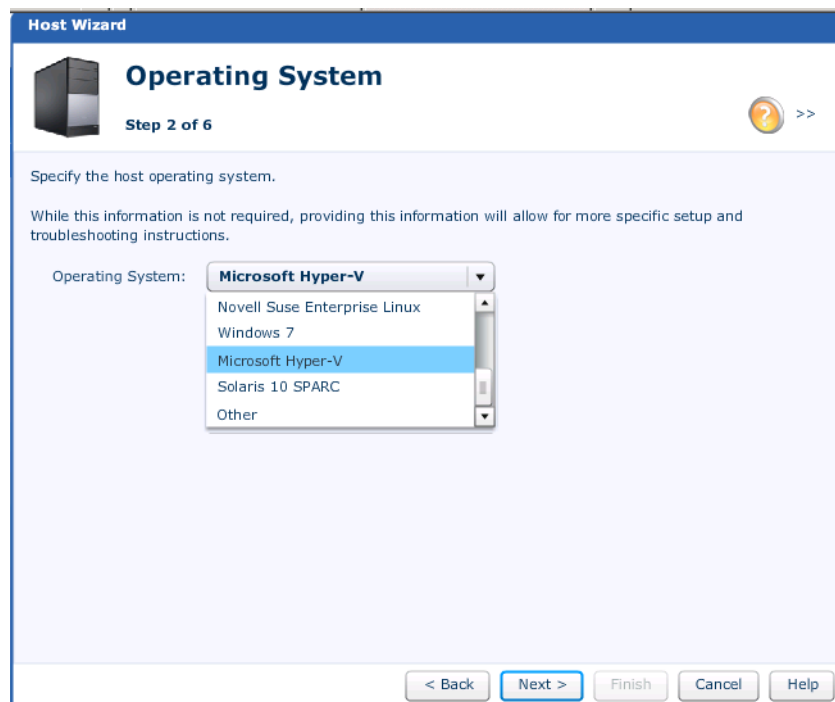
**Figure 87**      *Specifying Name for Host Configuration*



The image shows a screenshot of the 'Host Wizard' window, specifically the 'Specify Name' step. The window has a blue header bar with the text 'Host Wizard'. Below the header, there is a server icon and the title 'Specify Name'. To the right of the title, it says 'Step 1 of 6' and there is a question mark icon with a double arrow. The main area of the window contains the instruction 'Enter a name and optional description for the host configuration:'. Below this, there are two input fields. The first is labeled 'Name: \*' and contains the text 'M100N1'. The second is labeled 'Description:' and contains the text 'Hyper-V Cluster Node 1'. At the bottom of the window, there are five buttons: '< Back', 'Next >', 'Finish', 'Cancel', and 'Help'. The 'Next >' button is highlighted with a blue border.

3. Choose the host OS from the “Operating System” drop-down list. Click **Next**.  
The “Network Address” page appears.

**Figure 88**      **Specifying Host Operating System**



**Host Wizard**

**Operating System**

Step 2 of 6

Specify the host operating system.

While this information is not required, providing this information will allow for more specific setup and troubleshooting instructions.

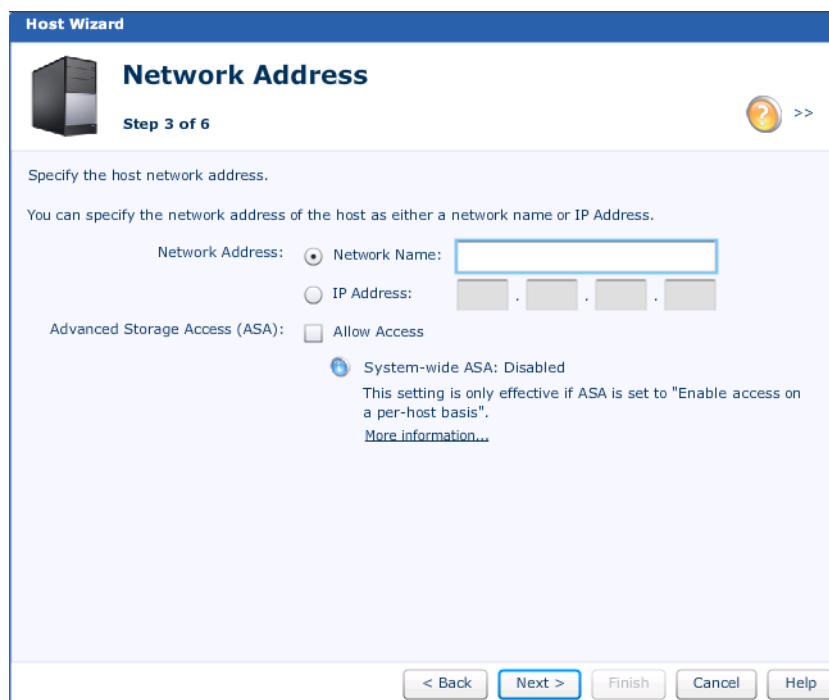
Operating System: **Microsoft Hyper-V**

- Novell Suse Enterprise Linux
- Windows 7
- Microsoft Hyper-V**
- Solaris 10 SPARC
- Other

< Back   Next >   Finish   Cancel   Help

4. Choose the “Network Name” or “IP Address” radio button to enter the details of the host.

**Figure 89**      **Specifying the Host Network Address**



**Host Wizard**

**Network Address**

Step 3 of 6


Specify the host network address.

You can specify the network address of the host as either a network name or IP Address.

Network Address: ☒ Network Name:

☐ IP Address:  .  .  .

Advanced Storage Access (ASA): ☐ Allow Access

 System-wide ASA: Disabled  
This setting is only effective if ASA is set to "Enable access on a per-host basis".  
[More information...](#)

< Back   Next >   Finish   Cancel   Help

5. Enter the IQN of the Hyper-V host in the IQN field. The host IQN can be found in the Configuration tab of the iSCSI initiator. Enter the CHAP secret password if required in your environment in the CHAP Secret field.

The “Summary” page appears.

**Figure 90**      **Specifying iSCSI Unique Number (IQN)**

**Host Wizard**

**iSCSI Access**

**Step 4 of 6**

If this host is connected to iSCSI storage, you must specify a valid iSCSI address (IQN). You can enter up to two IQNs, each with an optional CHAP secret, now:

The iSCSI Initiator Node Name (IQN) is a unique name used to identify a host using the iSCSI protocol. To use CHAP (optional), specify a CHAP Secret when adding an IQN.

IQN:

CHAP Secret:

Confirm CHAP Secret:

[Add Another IQN](#)

< Back   **Next >**   Finish   Cancel   Help

6. Review the host details and click **Finish**.



**Figure 91 Host Configuration Summary**

**Host Wizard**  
**Summary**  
 Step 5 of 6

Confirm the following Host configuration:

Name:	M100N5
Description:	Hyper-V Cluster Node 5
Operating System:	Microsoft Hyper-V
Network Name:	M100N5.M50VSPEX.COM
Advanced Storage Access (ASA):	Not Allowed
IQNs:	iqn.1991-05.com.microsoft:m100n5.m50vspex.com
CHAP Secret:	Not Specified

< Back   Next >   Finish   Cancel   Help

- Repeat the above steps to create Host list for all Hyper-V hosts.

**Figure 92 Hosts Window Listing All the Hyper-V Hosts**

**EMC Unisphere**

Dashboard   System   Storage   Settings   **Hosts**   Support

VSPEX3300 > Hosts > Hosts

**Hosts**

Name	Description	Network Address	Advanced Stor...	IQN	Operating System
M100N1		10.29.150.171	Disabled	iqn.1991-05.com.mi	Microsoft Hyper-V
M100N2		10.29.150.172	Disabled	iqn.1991-05.com.mi	Microsoft Hyper-V
M100N3		10.29.150.173	Disabled	iqn.1991-05.com.mi	Microsoft Hyper-V
M100N4		10.29.150.174	Disabled	iqn.1991-05.com.mi	Microsoft Hyper-V

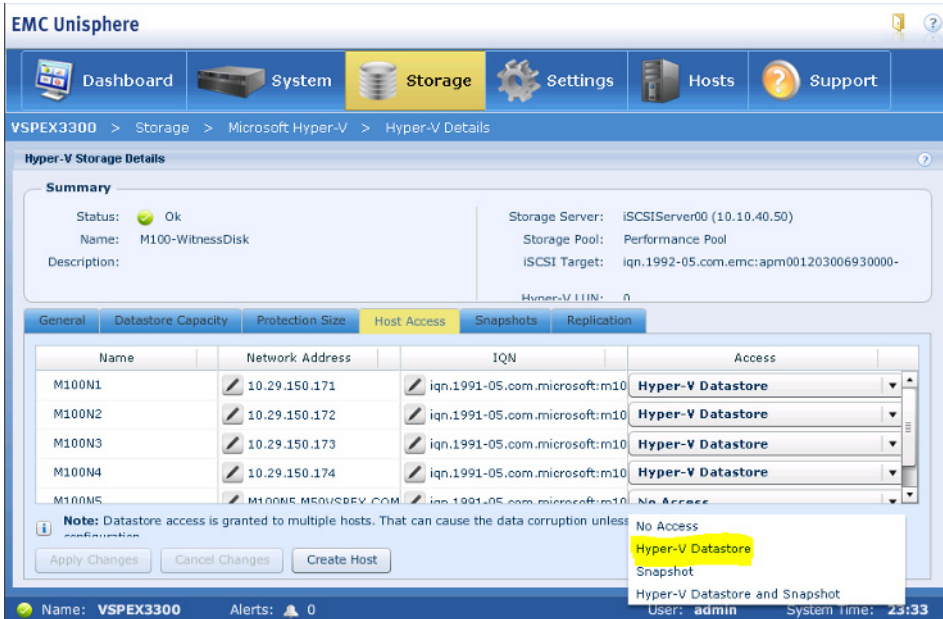
Selected: 0   Items: 4

Create Host   Create Subnet   Create Netgroup   Details   Refresh   Delete

- Click **Storage > Microsoft Hyper-V**.
- Choose an Allocated Hyper-V Datastore and click **Details**.
- Click the **Host Access** tab.

11. From the “Access” column drop-down list, choose “Hyper-V Datastore” for all the hosts participating in a Microsoft Windows Hyper-V cluster.

**Figure 93** *Hyper-V Datastore Details Window*



12. Repeat the above steps to provide Host Access for all the Hyper-V datastores created earlier.

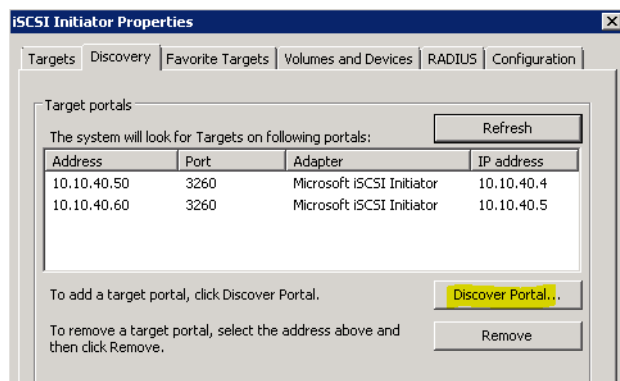
## Microsoft Windows Failover Cluster Setup

### iSCSI Initiator Configuration

To connect iSCSI targets and configure advanced settings, follow these steps:

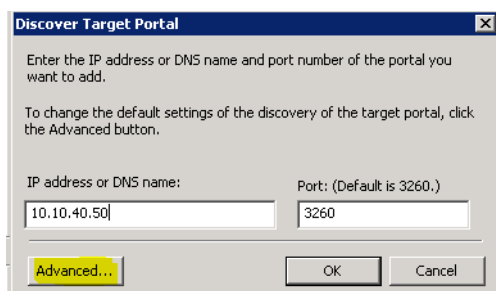
1. Login to the Microsoft Windows Server 2008 R2 Hyper-V host.
2. Click **Start > Administrative Tools > iSCSI initiator**.
3. In the “iSCSI Initiator Properties” dialog box, click **Discovery** and then click **Discover Portal**.

**Figure 94** Adding Target Portal in iSCSI Initiator Properties Window



4. In the “IP address” or “DNS” name field, enter the EMC VNXe iSCSI Server IP for SP A created earlier in the EMC VNXe3300 Deployment Procedure and click **Advanced**.

**Figure 95** Changing the Default Settings of Target Portal



5. The Advanced Settings dialog box appears. Complete the following in the Advanced Setting window:
  - a. Choose “Microsoft iSCSI Initiator” from the Local adapter list box.
  - b. Choose the “IP address” of the first NIC connected to the iSCSI server from the Initiator IP list box.
  - c. If you are required to use the CHAP, enter the details else ignore and click **Ok**.

**Figure 96**      **Advanced Settings Window**

**Advanced Settings**

General | IPsec

Connect using

Local adapter: Microsoft iSCSI Initiator

Initiator IP: 10.10.40.4

Target portal IP:

CRC / Checksum

☐ Data digest ☐ Header digest

☐ Enable CHAP log on

CHAP Log on information

CHAP helps ensure connection security by providing authentication between a target and an initiator.

To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.

Name: iqn.1991-05.com.microsoft:m100n1.m50vspex.com

Target secret:

☐ Perform mutual authentication

To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.

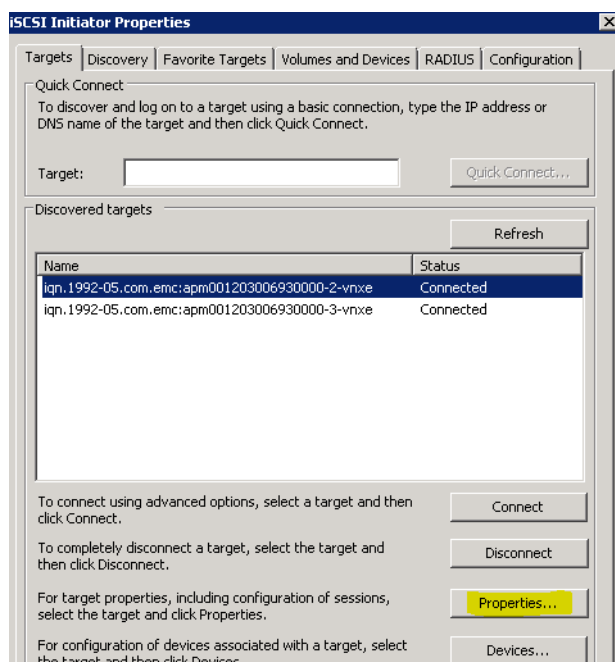
☐ Use RADIUS to generate user authentication credentials

☐ Use RADIUS to authenticate target credentials

OK Cancel Apply

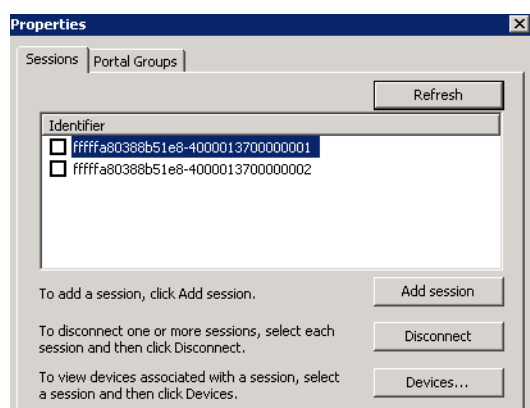
6. In the “iSCSI Initiator Properties” dialog box, verify the “Target Portals” details that are displayed in the “Discovery” window.
7. Choose the **Targets** tab, and then choose the **VNXe** as the target name.
8. Click **Connect**.  
The “Connect To Target dialog box” appears.
9. Complete the following in the Targets Window:
  - a. Choose “Add this connection to the list of Favorite Targets”.
  - b. Clear **Enable multi path** and click **Ok**. In the Targets tab of the iSCSI Initiator Properties dialog box, verify that the status of the target shows connected.

**Figure 97** *Window Displaying Discovered Targets*



10. Repeat steps from 2-7 to add target portal IP of the VNXe iSCSI Server IP for SP B using the second NIC configured for iSCSI.
11. Choose the **Targets** tab, and then choose the first “VNXe” target name.
12. Click **Properties** and in the Properties dialog box click **Add Session**.

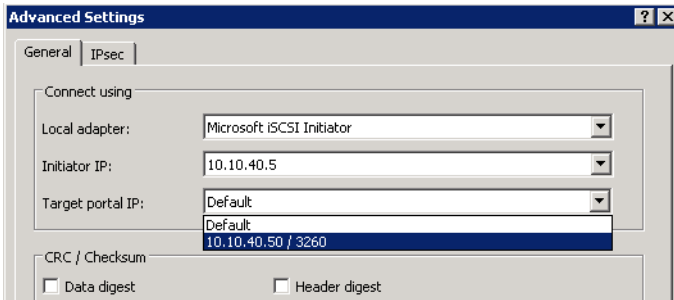
**Figure 98** *iSCSI Initiator Properties Displaying Session Identifiers*



13. Choose “Add this connection” to the list of Favorite Targets and click **Advanced**.  
The Advanced Settings dialog box appears.
14. Complete the following in the Advanced Settings window:
  - a. Choose “Microsoft iSCSI Initiator” from the Local adapter list box.
  - b. Choose the “IP address” of the second NIC connected to the iSCSI server from the Initiator IP list box.

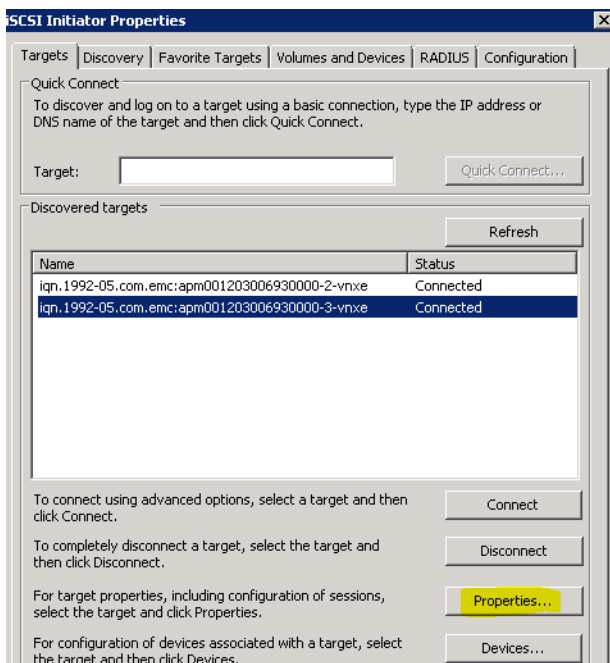
- c. Choose the “Target portal IP” of the VNXe iSCSI Server IP for SP A in the Target portal IP list box. Click **Ok**.

**Figure 99**      **Advanced Settings Window**



15. Choose “Targets” and then enter the second “VNXe” target name.
16. Click **Properties** and in the Properties dialog box click **Add Session**.

**Figure 100**      **Status of the Discovered Targets**



17. Choose “Add this connection” to the list of Favorite Targets and click **Advanced**.  
The Advanced Settings dialog box appears.
18. Complete the following in the Advanced Settings window:
  - a. Choose “Microsoft iSCSI Initiator” from the Local adapter list box.
  - b. Choose the IP address of the first NIC connected to the iSCSI server from the Initiator IP list box.
  - c. Choose the Target portal IP of the VNXe iSCSI Server IP for SP B in the Target portal IP list box. Click **Ok**.

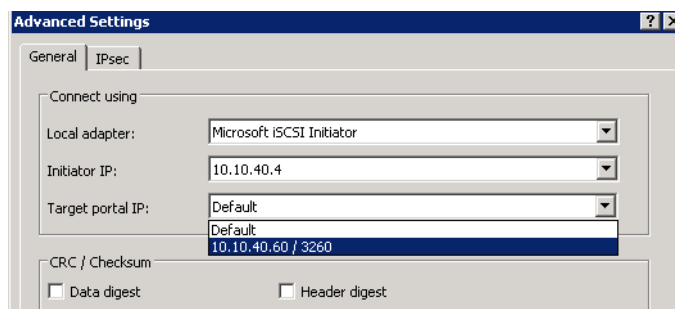
**Figure 101**      **Advanced Settings**

Figure 102 shows the partial output of command **mpclaim.exe -b** captured in a file.

**Figure 102**      **mpclaim.exe Configuration Output**

```

config_b.txt - Notepad
File Edit Format View Help
MPIO Storage Snapshot on Monday, 09 July 2012, at 02:44:37.565

Registered DSMs: 1
=====
+-----+-----+-----+-----+-----+-----+
| DSM Name | Version | PRP | RC | RI | PVP | PVE |
+-----+-----+-----+-----+-----+-----+
| Microsoft DSM | 006.0001.07601.17514 | 0020 | 0003 | 0001 | 030 | False |
+-----+-----+-----+-----+-----+-----+

Microsoft DSM
=====
MPIO Disk10: 02 Paths, Round Robin, ALUA Not Supported

Path ID          State          SCSI Address    Weight
-----
0000000077010003 Active/Optimized 001|000|003|001 0
0000000077010002 Active/Optimized 001|000|002|001 0

MPIO Disk9: 02 Paths, Round Robin, ALUA Not Supported

Path ID          State          SCSI Address    Weight
-----
0000000077010003 Active/Optimized 001|000|003|000 0
0000000077010002 Active/Optimized 001|000|002|000 0

MPIO Disk8: 02 Paths, Round Robin, ALUA Not Supported

Path ID          State          SCSI Address    Weight
-----
0000000077010001 Active/Optimized 001|000|001|008 0
0000000077010000 Active/Optimized 001|000|000|008 0

MPIO Disk7: 02 Paths, Round Robin, ALUA Not Supported

Path ID          State          SCSI Address    Weight
-----
0000000077010001 Active/Optimized 001|000|001|007 0
0000000077010000 Active/Optimized 001|000|000|007 0

MPIO Disk6: 02 Paths, Round Robin, ALUA Not Supported

Path ID          State          SCSI Address    Weight
-----
0000000077010001 Active/Optimized 001|000|001|006 0
0000000077010000 Active/Optimized 001|000|000|006 0

```

19. Login to the “Windows Hyper-V host” and open “Server Manager”.
20. In “Server Manager”, expand “Storage” and click **Disk Management**.
21. In the “Disk Management” right window pane, choose and right-click all the SAN disks to online and initialize them.

Once all the disks are initialized, format with the NTFS file system and assign drive letters to them.

22. Login to other Hyper-V hosts to which the same above SAN disks are provisioned and bring them online.

## Cluster Validation

**Table 11** *Microsoft Windows Failover Cluster Details*

Cluster Node Name	Node IP Address	Cluster IP Address	Cluster Name
M100N1.M50VSPEX.COM	10.29.150.171	10.29.150.175	M100Clus
M100N2.M50VSPEX.COM	10.29.150.172		
M100N3.M50VSPEX.COM	10.29.150.173		
M100N4.M50VSPEX.COM	10.29.150.174		

For cluster validation, follow these steps:

1. Login to M100N1 host using a domain administrative account with local privileges.
2. Open Server Manager and browse to **Features > Failover Cluster Manager**.
3. Validate cluster feasibility:
  - a. Choose “Validate a Configuration”, and click **Next**.
  - b. Add all the nodes one at a time into the Enter server name text field, and click **Next**
  - c. Choose “Run all tests” and click **Next**.
  - d. Click **Next > Next**.
  - e. Review the report and resolve any issues found by the validation wizard before continuing.

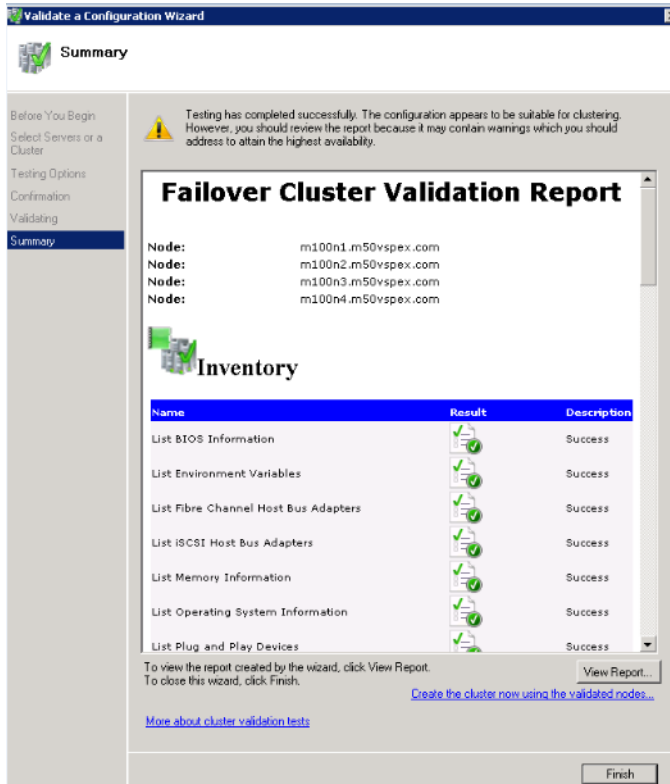


**Note** The warning in [Figure 103](#) is expected because the iSCSI NICs are on the same subnet. For failsafe and HA feature of the EMC VNXe storage to work, the peer ports on both storage processors must be on the same subnet and that is the reason for iSCSI NICs to be on same subnet.

- f. Click **Finish**.



**Figure 103**      **Failover Cluster Validation Report**



## Failover Cluster Setup

To setup a failover cluster, follow these steps:

1. In the Failover Cluster Manager, choose **Create a Cluster**.
2. In the “Welcome” screen, click **Next**.
3. Add all the nodes one at a time into the Enter server name text field and click **Next**.

**Figure 104**      **Selecting Servers for Adding into the Cluster**

**Create Cluster Wizard**

**Select Servers**

Before You Begin  
**Select Servers**  
 Access Point for Administering the Cluster  
 Confirmation  
 Creating New Cluster  
 Summary

Add the names of all the servers that you want to have in the cluster. You must add at least one server.

Enter server name:

Selected servers:

- m100n1.m50vspex.com
- m100n2.m50vspex.com
- m100n3.m50vspex.com
- m100n4.m50vspex.com

< Previous   Next >   Cancel

4. Enter the “Cluster Name”, “Cluster IP”, and click **Next**.

**Figure 105**      **Cluster Details**

**Create Cluster Wizard**

**Access Point for Administering the Cluster**

Before You Begin  
 Select Servers  
**Access Point for Administering the Cluster**  
 Confirmation  
 Creating New Cluster  
 Summary

Type the name you want to use when administering the cluster.

Cluster Name:

One or more IPv4 addresses could not be configured automatically. For each network to be used, make sure the network is selected, and then type an address.

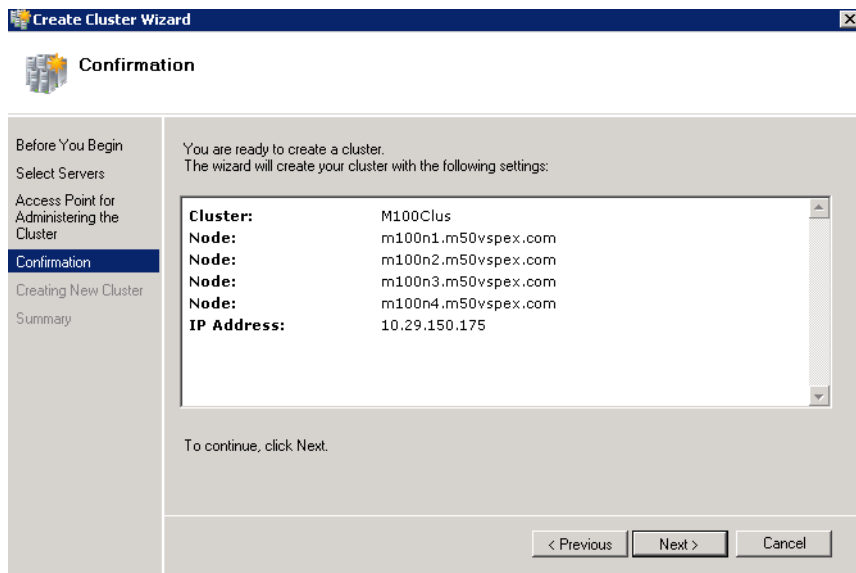
	Networks	Address
<input checked="" type="checkbox"/>	10.29.150.0/24	10.29.150.175
<input type="checkbox"/>	10.10.46.0/24	Click here to type an address
<input type="checkbox"/>	10.10.40.0/24	Click here to type an address

[More about the administrative Access Point for a cluster](#)

< Previous   Next >   Cancel

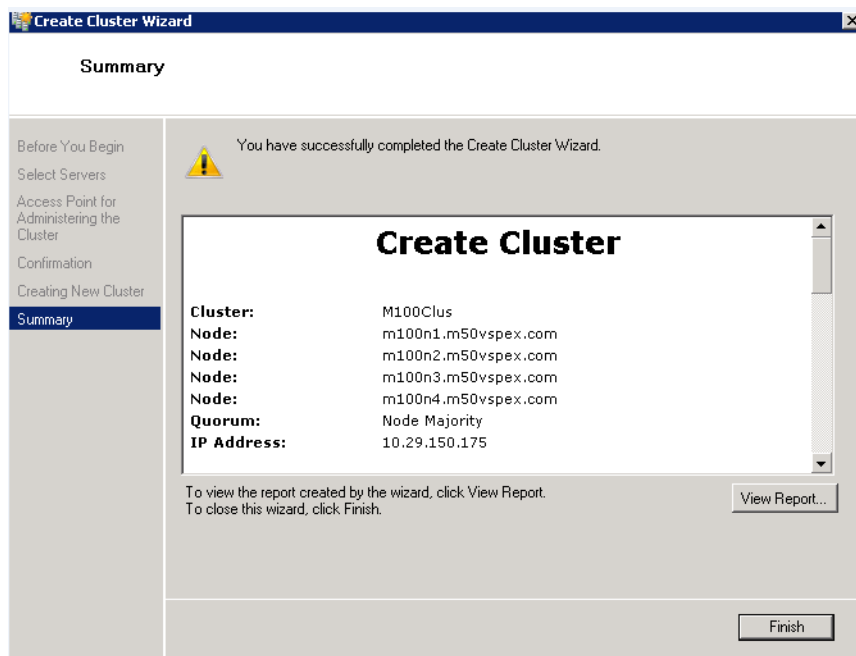
5. In the “Confirmation” page, review and click **Next**.

**Figure 106 Confirmation Window for Creating Cluster**



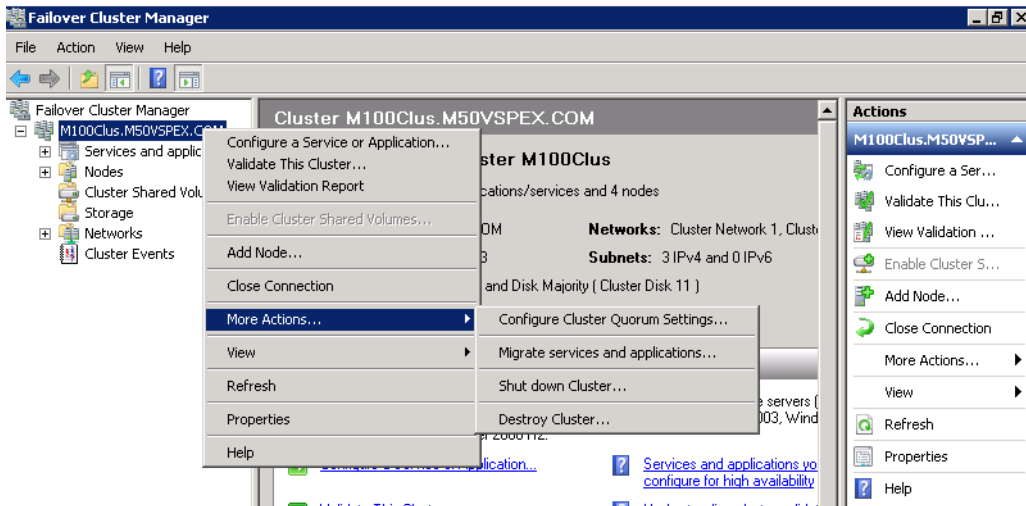
6. In the “Summary” window click **Finish**.

**Figure 107 Window Showing Summary of the Created Cluster**



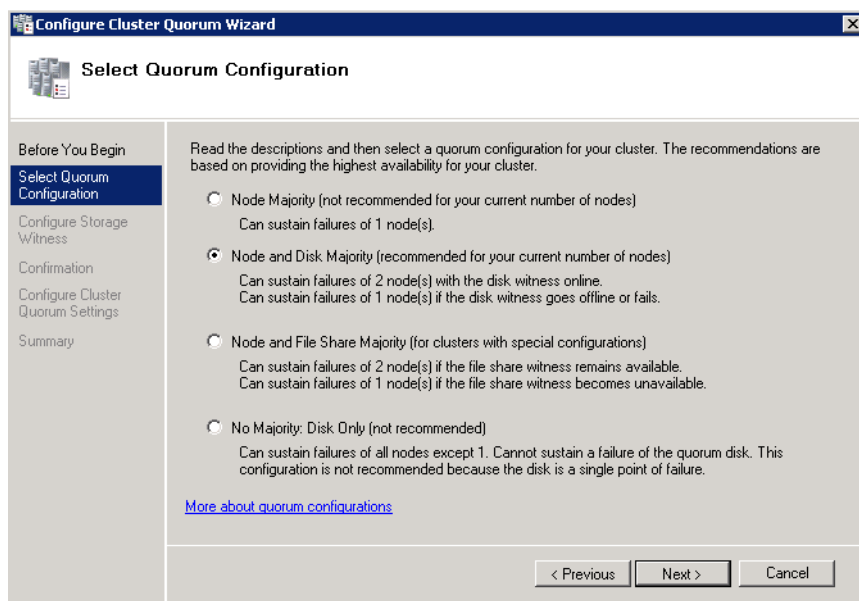
7. In the “Failover Cluster Manager”, right-click the cluster name, choose **More Actions, Configure Cluster Quorum Settings** and click **Next**.

**Figure 108** *Failover Cluster Manager*



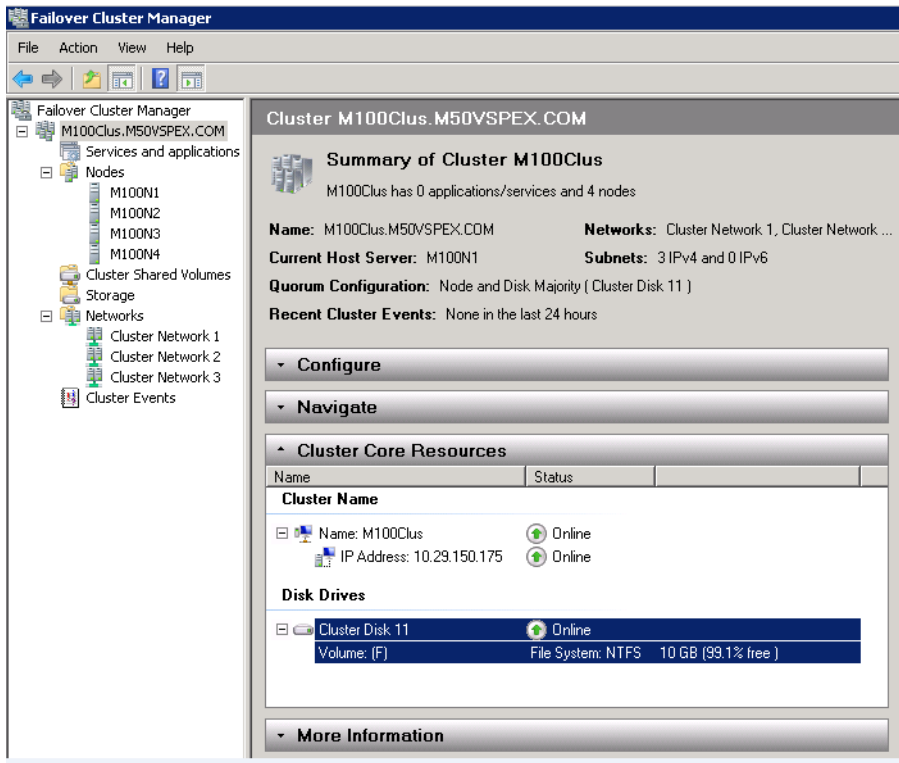
8. In the “Select Quorum Configuration” page, choose the “Node and Disk Majority” radio button and click **Next**.

**Figure 109** *Configure Quorum Configuration*



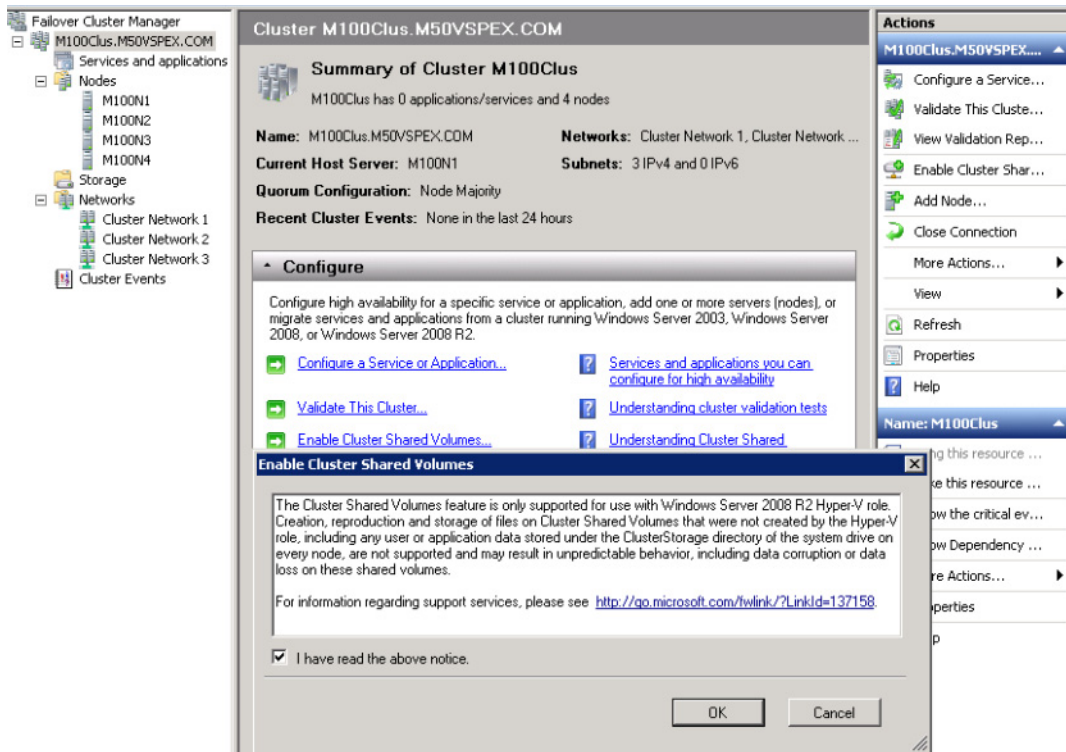
9. In the “Configure Storage Witness” page choose the 10GB disk and click **Next**.
10. In the “Confirmation” page review and click **Finish**.

**Figure 110**      **Failover Cluster Manager Online Status**



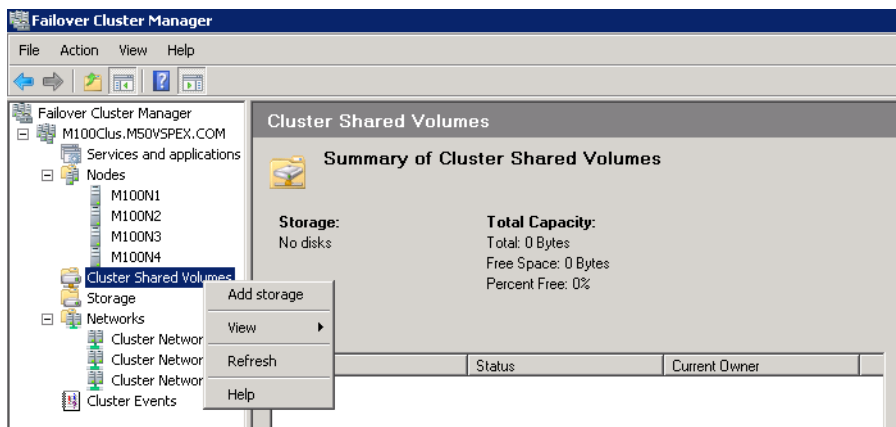
11. To enable the Cluster Shared Volume, open “Failover Cluster Manager” from the node that currently owns the cluster.
12. In the “Configure Section” choose **Enable Cluster Shared Volumes**.
13. Check the check box “I have read the above notice” and click **Ok**.

**Figure 111 Enabling Cluster Shared Volumes (CSV)**



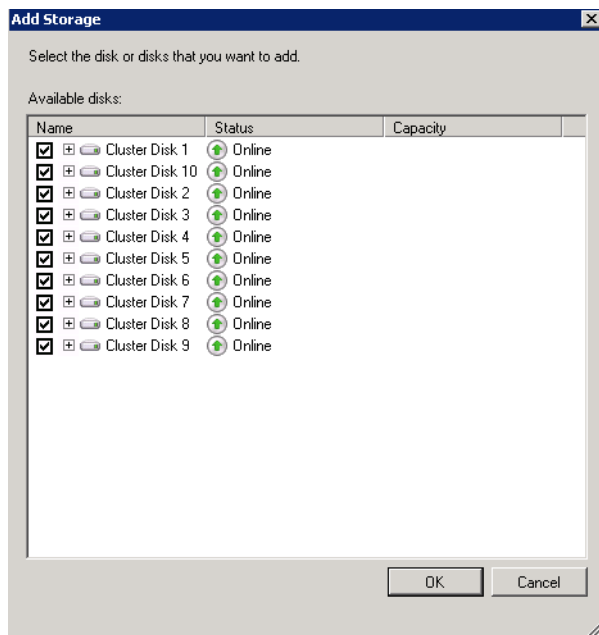
14. Right-click “Cluster Shared Volumes” and choose **Add Storage**.

**Figure 112 Adding Storage to CSV**



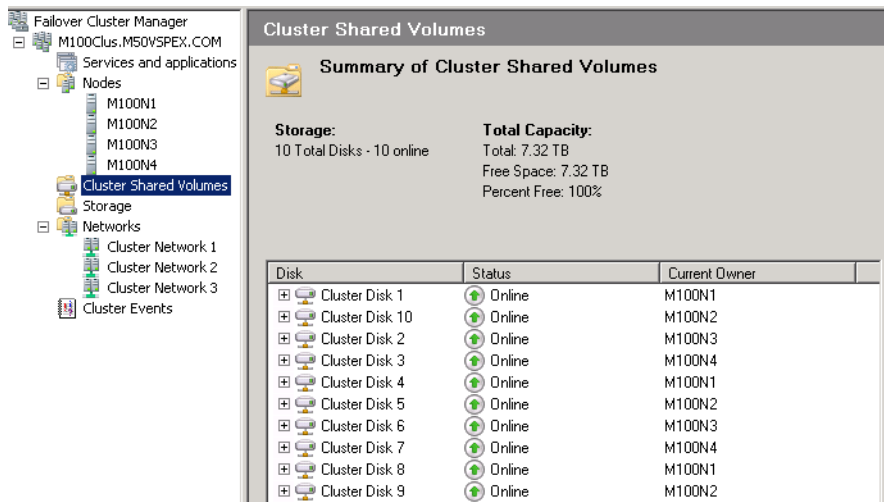
15. Choose all the volumes under “Available disks” and click **Ok**.

**Figure 113**      **Select Available Disks For CSV Storage**



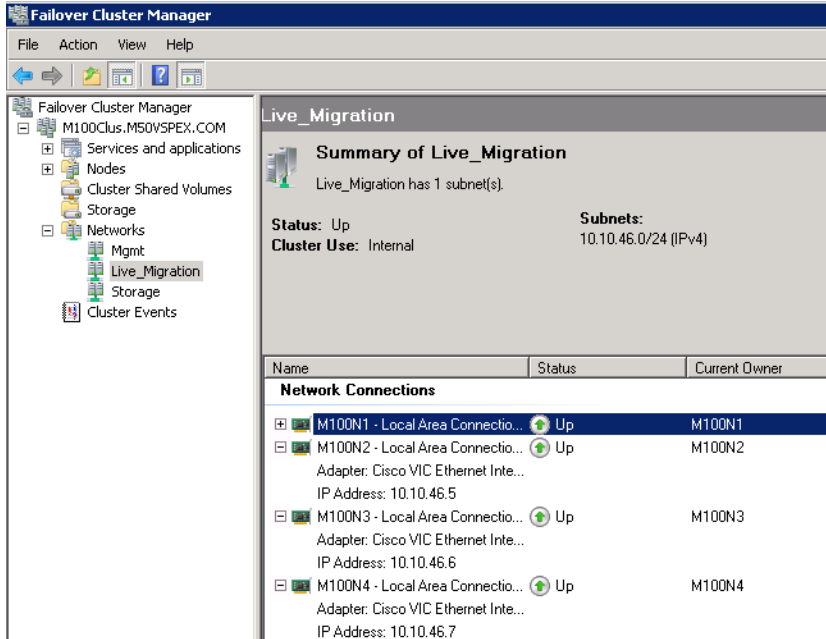
16. Once all the disks are added successfully, the status shows Online.

**Figure 114**      **Window Showing CSV Disk Status**



17. Rename the cluster networks (optional).

**Figure 115** *Live Migration Summary in Failover Cluster Manager*



## Microsoft System Center-2012 VMM Configuration

This section provides configuration details of Microsoft System Center Virtual machine Manager (VMM).

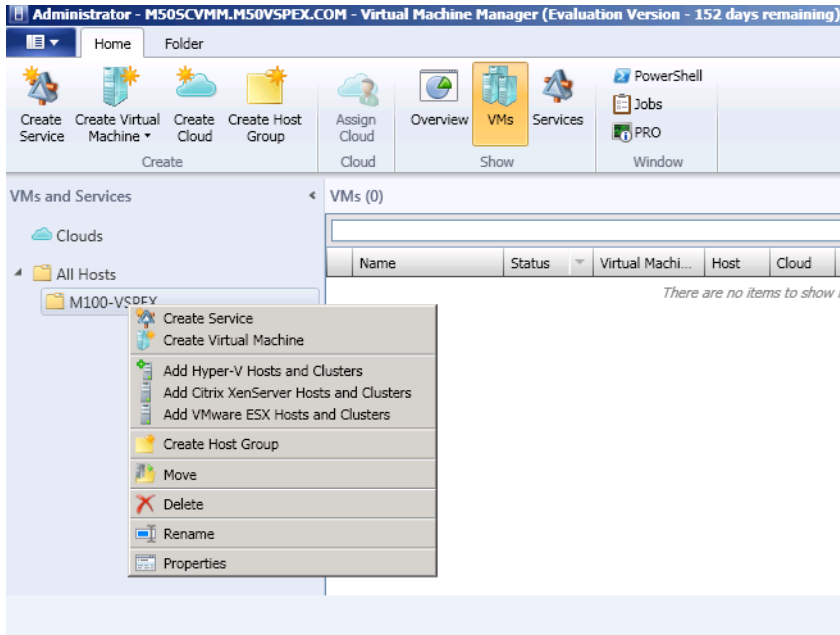
### Add Microsoft Hyper-V Hosts and Cluster

To add Microsoft Hyper-V hosts and cluster, follow these steps:

1. Create a “Host Group” and click **Add Hyper-V Hosts and Clusters**.

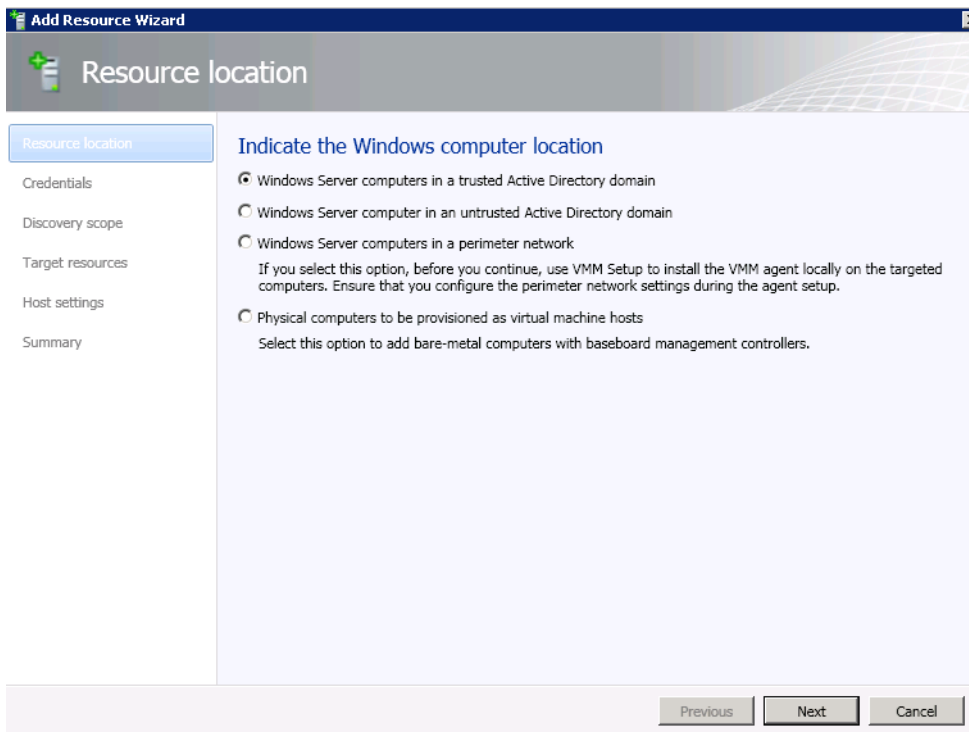


**Figure 116**      **Create a Host Group in SCVMM**



2. In the “Resource location” page of the “Add Resource Wizard”, choose the radio button “Windows server computers in a trusted Active directory domain” and click **Next**.

**Figure 117**      **Selecting Resource Location**

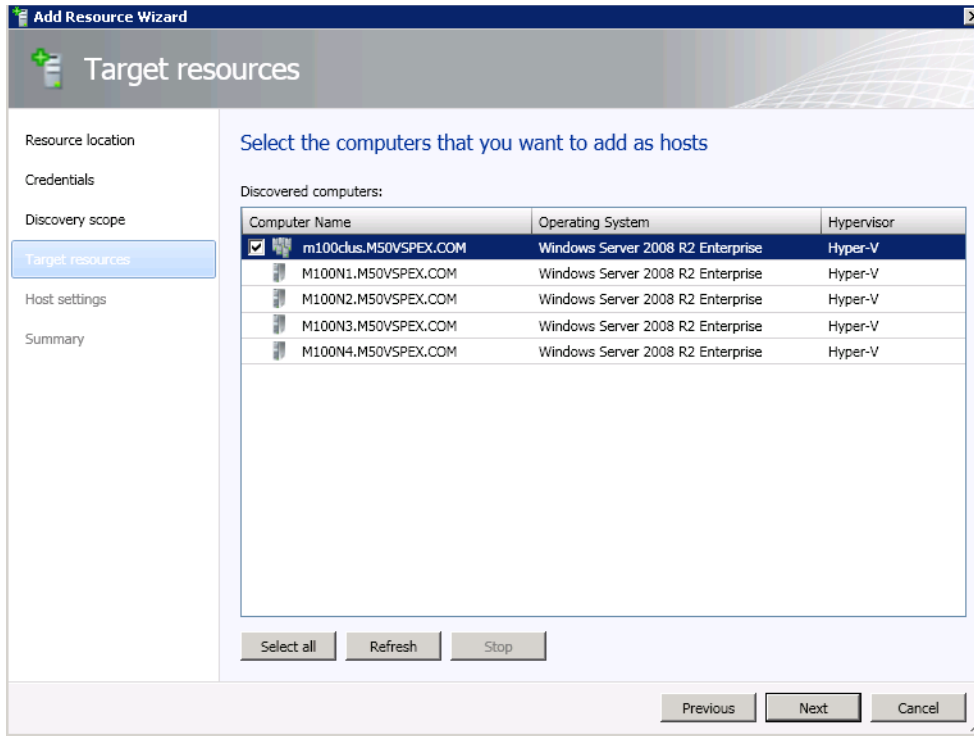


3. In the “Credentials” page of the “Add Resource Wizard”, provide the appropriate credentials and click **Next**.
4. In the “Discovery” page of the “Add Resource Wizard”, choose the radio button “Specify Windows Server Computer by Name” and enter the cluster name created in the previous sections and click **Next**.

**Figure 118**      **Specifying Discovery Scope for VM Hosts**

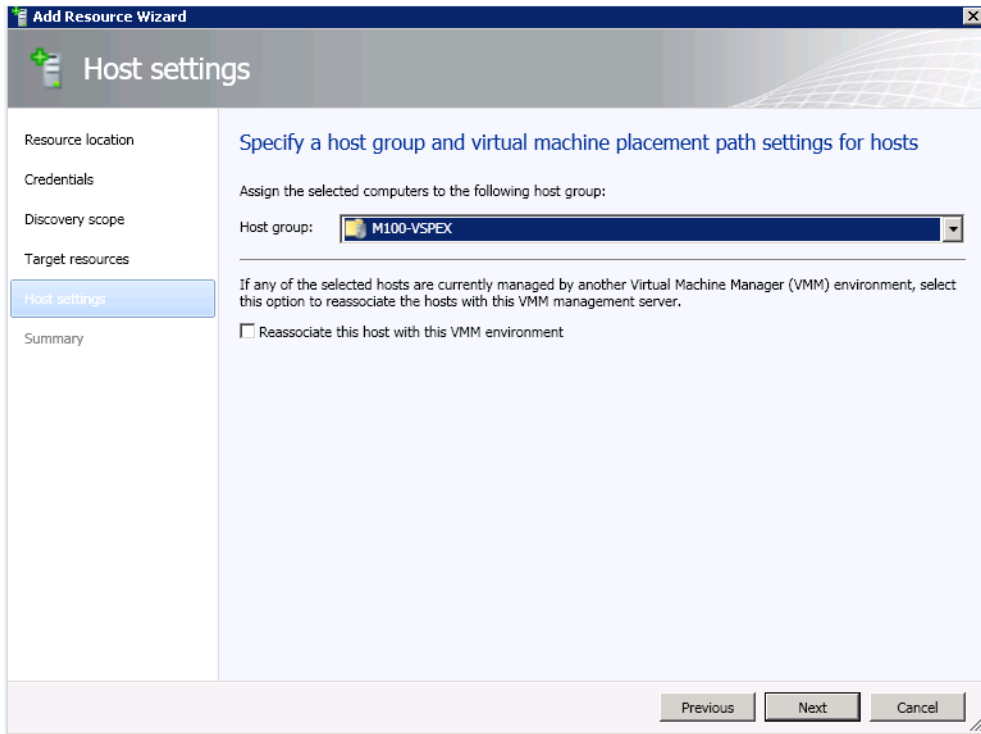
5. In the “Target Resources” page of the “Add Resource Wizard” choose the “cluster name” and click **Next**.

**Figure 119**      **Target Resources to Add as Hosts**



6. In the “Host Settings” page of the “Add Resource Wizard” choose the “Host group” created in step 1 and click **Next**.

**Figure 120**      *Specifying Host Group and VM Path for Hosts*



## Create a Template for Virtual Machine Deployment


This section covers how to create A Virtual Machine, A Virtual Machine Template, and Highly Available Virtual Machines from the Template.

To create a virtual machine from a blank .vhd file, follow these steps:

1. Open the VMs and Services workspace.
2. In the Home tab, in the “Create” group, click the **Create Virtual Machine** drop-down list, and then click **Create Virtual Machine**.

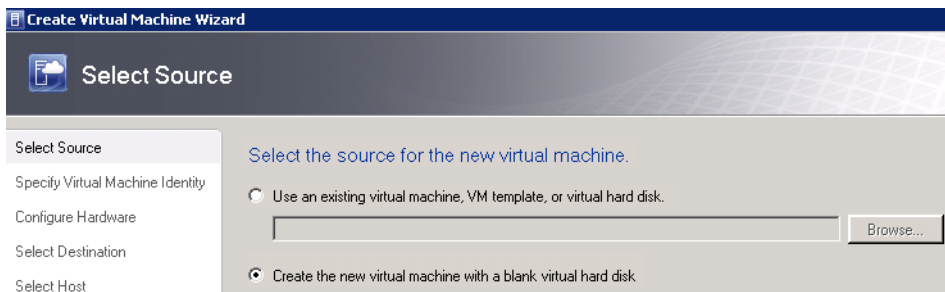
The “Create Virtual Machine” wizard appears.

**Figure 121 Create Virtual Machine**



3. In the “Select Source” page, choose the radio button “Create the new virtual machine with a blank virtual hard disk” and click **Next**.

**Figure 122 Selecting Source for New VM**



4. In the “Specify Virtual Machine Identity” page, enter the virtual machine name and optional description, and then click **Next**.
5. In the “Configure Hardware” page, perform any one of the following, and click **Next**.
  - a. To use an existing hardware profile, in the Hardware profile list, click the desired profile.
  - b. Configure hardware settings manually.
6. In the “Select Destination” page, choose to place the virtual machine on a virtual machine host.
  - a. In the “Select Host” page, review the placement ratings and transfer type, click the **Desired Host**, and then click **Next**.
  - b. In the “Configure Settings” page, under “Locations”, either accept the default virtual machine path on the host for the virtual machine files, or click **Browse** to specify a different location. If desired, choose the “Add this path” to the list of default virtual machine paths on the host check box.  
  
Under Machine Resources, click **Virtual Hard Disk**. You can accept the default values, or choose a different destination path on the host for the .vhd file. To change the .vhd file name, enter a new name in the File name box.
  - c. In the “Select Networks” page (if it appears), optionally choose the desired logical network, the virtual network, and the VLAN ID (if applicable), and click **Next**.

- d. In the “Add Properties” page, configure the action to take when the host starts or stops, and the operating system that you install on the virtual machine. Click **Next**.
  - e. In the “Summary” page, confirm the settings and click **Create**.
7. Install the Operating System on the Virtual Machine created above with latest updates and service pack. Install any Roles and Features, applications and enable any required services to make this a golden image for template creation.

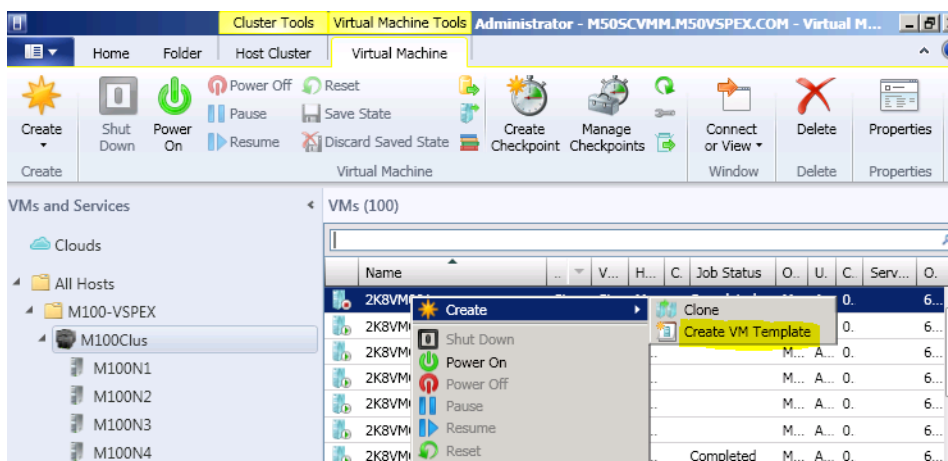
## Create a Virtual Machine Template

This section covers on how-to create a virtual machine template from an existing virtual machine that is deployed on a host.

To create a virtual machine template, follow these steps:

1. Shutdown the VM for template creation and Open the Library workspace.
2. In the Virtual Machine tab, choose and right-click the VM that needs to be converted to a template and click **Create** and then **Create VM Template**.

**Figure 123**      **Creating VM Template**

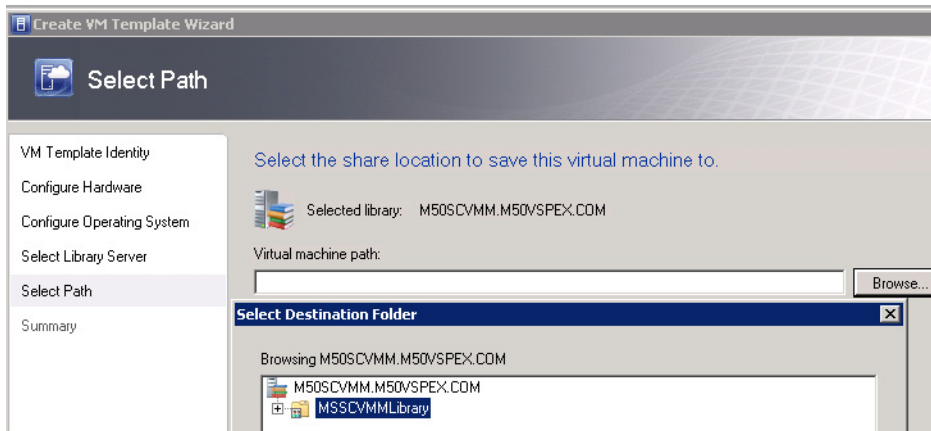


A warning message is displayed “creating a template destroys the source virtual machine, and any user data on the source virtual machine may be lost”.

3. To continue, click **Yes**.
4. In the “VM Template Identity” page, provide a name for the virtual machine template, and click **Next**.
5. In the “Configure Hardware” page click **Next**.
6. In the “Configure Operating System” page, configure the guest operating system settings.  
If you have an existing guest operating system profile that you want to use, in the Guest OS profile list, click the desired guest operating system profile.
7. After you have configured the guest operating system settings, click **Next**.
8. In the “Select Library Server” page, click the “library server” for the virtual machine, and click **Next**.

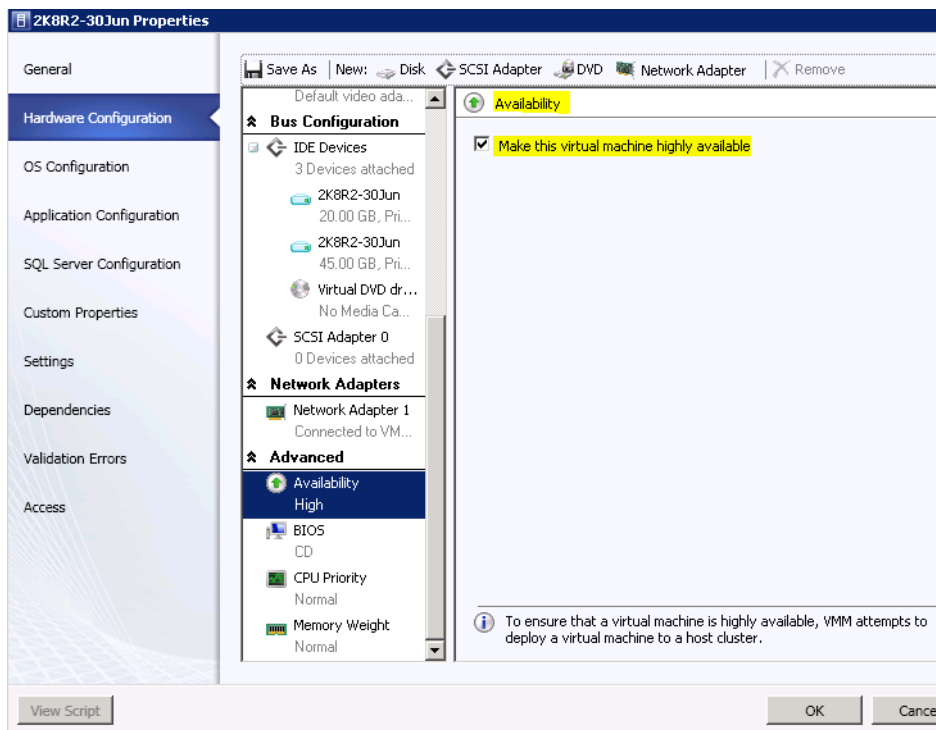
9. In the “Select Path” page, click **Browse**, click a “library share” and optional folder path, click **OK**, and click **Next**.

**Figure 124**      *Selecting Location for Saving the VM*



10. In the “Summary” page, confirm the settings, and click **Create**.
11. In the “Templates” tab, choose and right-click the template created above and click **Properties**. Make the necessary changes like making the VM highly available as shown in [Figure 125](#).

**Figure 125**      *Modifying Template Properties for VM High Availability*



## Create a Highly Available Virtual Machine from a Template

You can use the following procedure to create a virtual machine from a virtual machine template in System Center 2012 - Virtual Machine Manager (VMM).

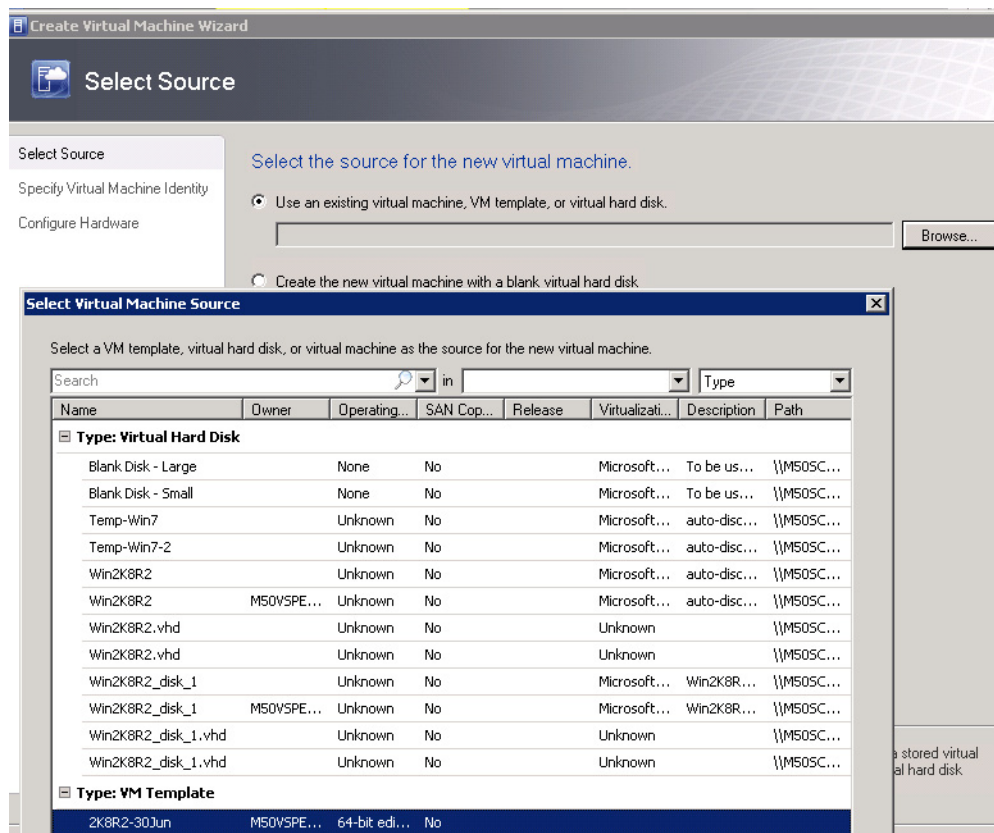
To create a highly available virtual machine from a template, follow these steps:

1. Open the VMs and Services workspace.
2. In the “Home” tab, in the “Create group” click the **Create Virtual Machine** drop-down button, and then click **Create Virtual Machine**.

The “Create Virtual Machine Wizard” opens.

3. In the “Select Source” page, ensure that you choose “Use an existing virtual machine”, “VM template”, or “virtual hard disk”, and then click **Browse**.
4. In the “Select Virtual Machine Source” dialog box, click the appropriate virtual machine template, and then click **Ok**.

**Figure 126**      *Selecting Source for Virtual Machine*



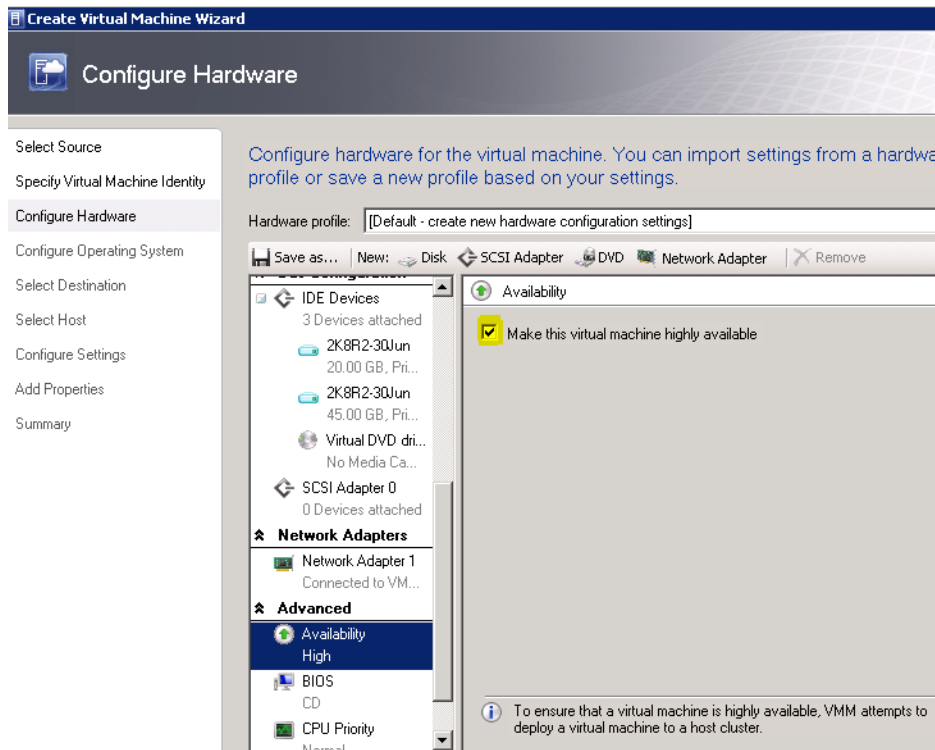
5. In the “Select Source” page, click **Next**.
6. In the “Specify Virtual Machine Identity” page, enter the virtual machine name and optional description, and click **Next**.
7. In the Configure Hardware page, configure the hardware settings.

If you have an existing hardware profile with settings that you want to use, in the Hardware profile list, click the desired hardware profile.



8. After you have configured the hardware settings, click **Next**.

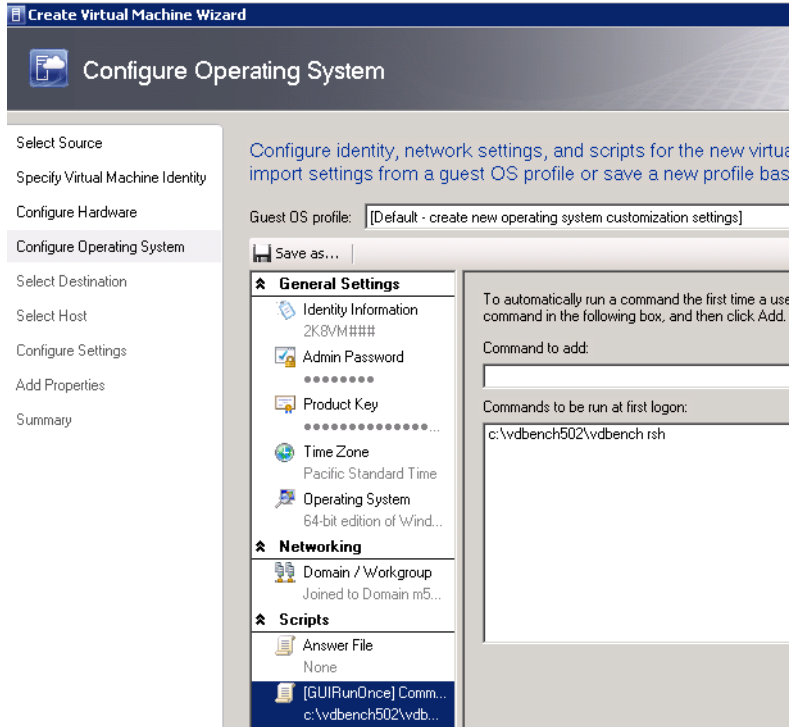
**Figure 127**      *Configuring Hardware for the VM*



In the “Configure Operating System” page, configure the guest operating system settings. If you have an existing guest operating system profile that you want to use, in the Guest OS profile list, click the desired guest operating system profile.

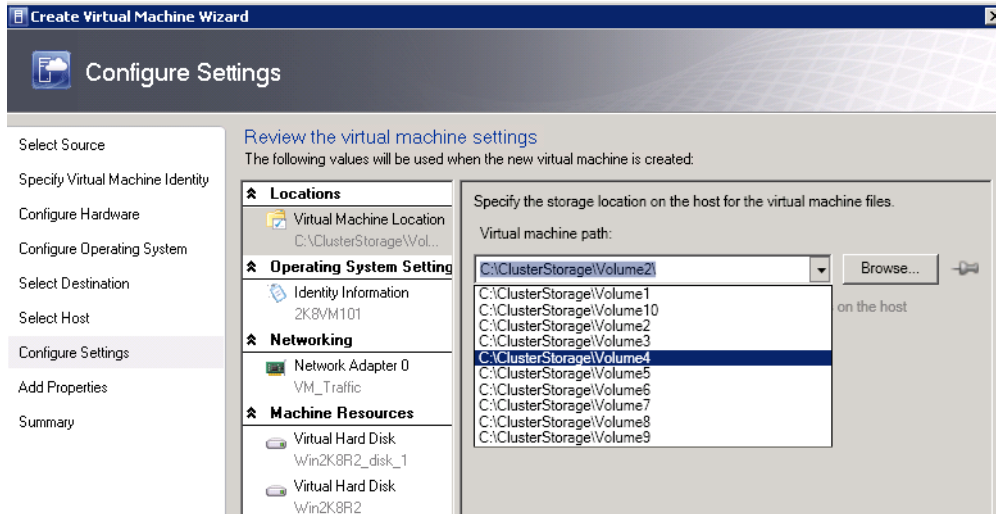
9. After you have configured the guest operating system settings, click **Next**.

**Figure 128**      **Configuring OS for the VM**



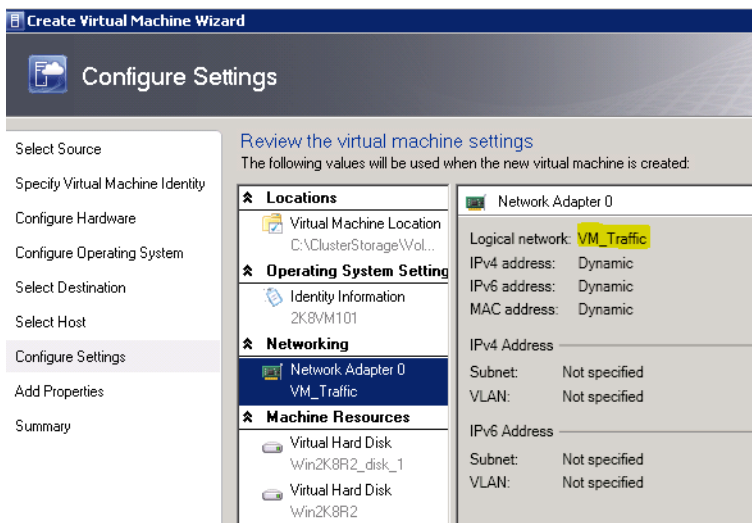
10. In the “Select Destination” page, choose whether to place the virtual machine on a virtual machine host.
11. In the “Select Host” page, review the placement ratings and transfer type, click a desired host that is available for placement, and click **Next**.
12. In the “Configure Settings” page, follow these steps:
  - a. Under Locations, click the **Virtual Machine** Location.
  - b. Either accept the default virtual machine path on the host for the virtual machine files, or click **Browse** to specify a different location.
  - c. If desired, choose “Add this path” to the list of default virtual machine paths on the host check box.

**Figure 129** Window Showing VM Settings

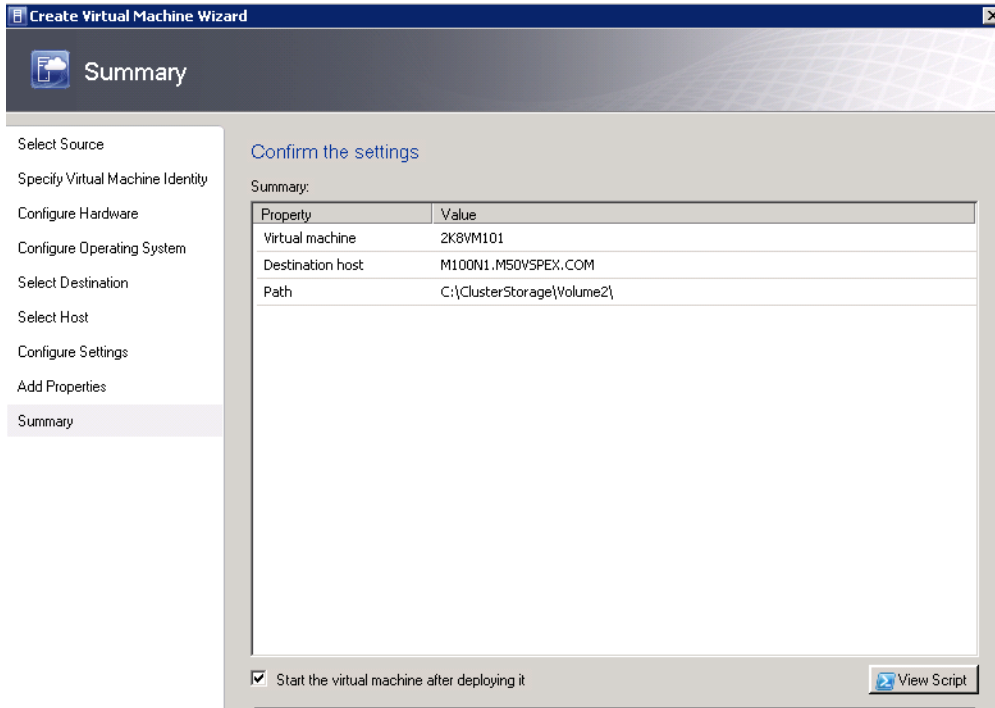


- d. Under Operating System Settings, click **Identity Information**.  
You can either accept or change the computer name.

**Figure 130** Window Showing Virtual Machine Settings



- e. Under “Networking”, click a network adapter to view the configured network settings.  
f. Under “Machine Resources”, click **Virtual Hard Disk**, review and optionally modify the settings, and click **Next**.  
g. In the “Add Properties” page, configure the action to take when the host starts or stops.  
h. In the “Summary” page, confirm the settings and click **Create**.

**Figure 131**      **Virtual Machine Settings Confirmation Window**

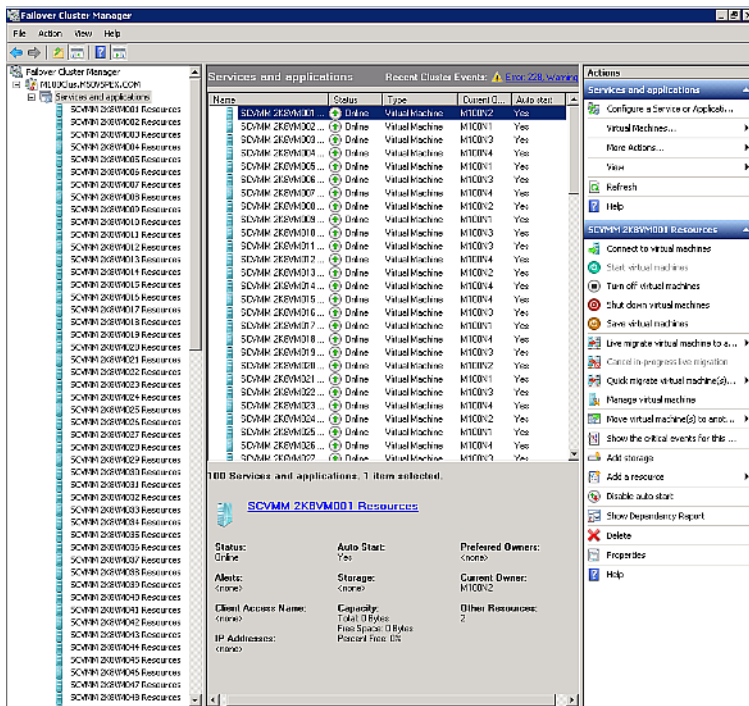
13. [Figure 132](#) shows the 100 highly available VMs deployed from a VM template.

**Figure 132** 100 Highly Available Virtual Machines in the SCVMM Console

Name	Status	Virtual Machine...	Host	Clo...	Job Status	Owner
2K8VM079	Running	Running	M100N3		Completed	M50VSPEx\Administrator
2K8VM080	Running	Running	M100N2		Completed	M50VSPEx\Administrator
2K8VM081	Running	Running	M100N1		Completed	M50VSPEx\Administrator
2K8VM082	Running	Running	M100N4		Completed	M50VSPEx\Administrator
2K8VM083	Running	Running	M100N3		Completed	M50VSPEx\Administrator
2K8VM084	Running	Running	M100N2		Completed	M50VSPEx\Administrator
2K8VM085	Running	Running	M100N1		Completed	M50VSPEx\Administrator
2K8VM086	Running	Running	M100N3		Completed	M50VSPEx\Administrator
2K8VM087	Running	Running	M100N4		Completed	M50VSPEx\Administrator
2K8VM088	Running	Running	M100N3		Completed	M50VSPEx\Administrator
2K8VM089	Running	Running	M100N2		Completed	M50VSPEx\Administrator
2K8VM090	Running	Running	M100N4		Completed	M50VSPEx\Administrator
2K8VM091	Running	Running	M100N1		Completed	M50VSPEx\Administrator
2K8VM092	Running	Running	M100N4		Completed	M50VSPEx\Administrator
2K8VM093	Running	Running	M100N3		Completed	M50VSPEx\Administrator
2K8VM094	Running	Running	M100N2		Completed	M50VSPEx\Administrator
2K8VM095	Running	Running	M100N1		Completed	M50VSPEx\Administrator
2K8VM096	Running	Running	M100N2		Completed w/ Info	M50VSPEx\Administrator
2K8VM097	Running	Running	M100N2		Completed w/ Info	M50VSPEx\Administrator
2K8VM098	Running	Running	M100N2		Completed w/ Info	M50VSPEx\Administrator
2K8VM099	Creating...	Stopped	M100N1		81 %	M50VSPEx\Administrator
2K8VM100	Creating...	Stopped	M100N1		23 %	M50VSPEx\Administrator

14. Figure 133 shows the highly available VMs as seen in the “Failover Cluster Manager”.

**Figure 133** 100 Highly Available Virtual Machines in the Failover Cluster Manager



## VSPEX M50 Configuration Details

### Cabling Information

This information is provided as a reference for cabling the physical equipment in a VSPEX M50 environment. The tables in this section include both the local and remote device and the 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.

Ensure that you follow the cable directions in this section. Failure to do so results in the necessary changes to the deployment procedures that follow because specific port locations are mentioned. Before starting, ensure that the configuration matches what is described in the tables and diagrams in this section.

**Figure 134** shows the VSPEX M50 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 the 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

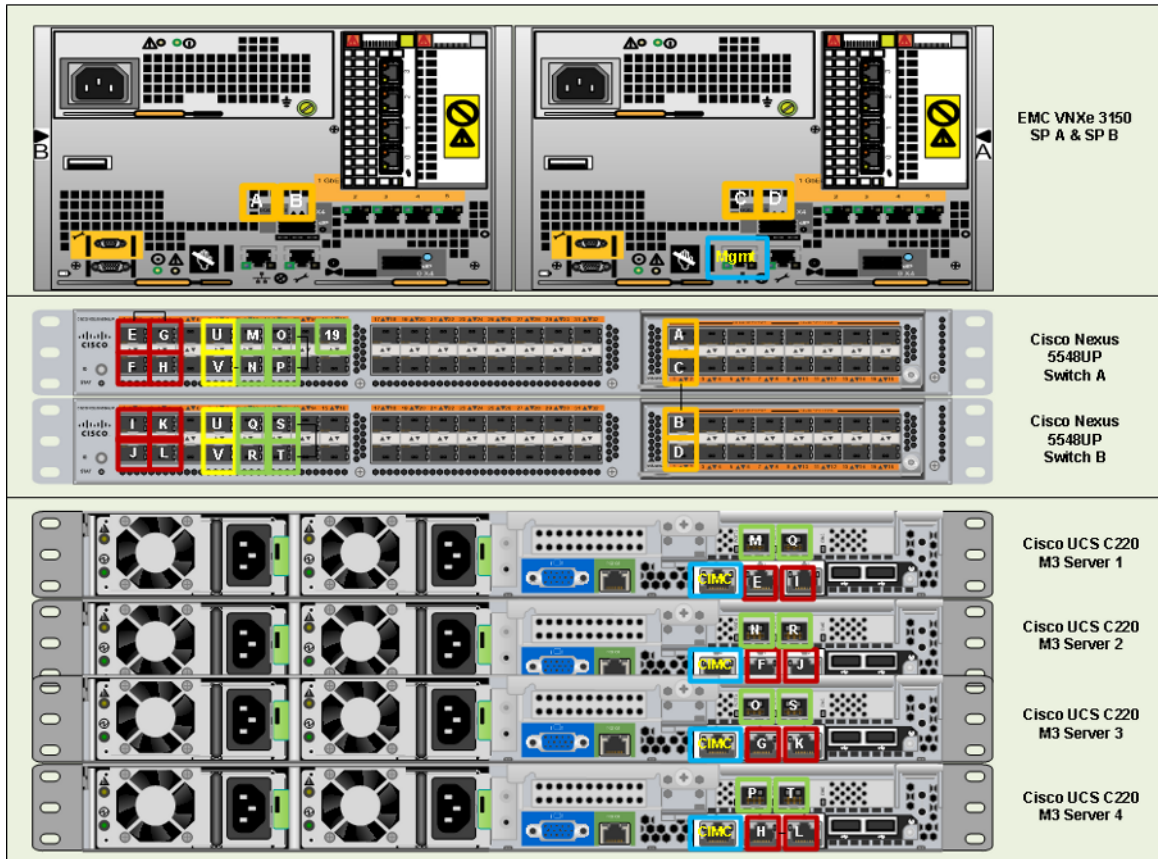
**Figure 134**      **Cabling Details for VSPEX Microsoft Hyper-V 50 Virtual Machines**

Table 12 and Table 13 lists five major cabling sections in these architectures:

1. Inter switch links
2. Data connectivity for servers (trunk links)
3. Management connectivity for servers
4. Storage connectivity
5. Infrastructure connectivity

**Table 12**      **Cisco Nexus 3048 A Ethernet Cabling Information Local Device Local Port Connection Remote**

Cable ID on both ends	Ethernet Interface	VLAN	Mode	Speed	Port Channel	Remote Device Port
E	Eth1/3	1, 23	trunk	1G	3	C220 Server1- 1GE LOM 1
F	Eth1/4	1,23	trunk	1G	4	C220 Server2- 1GE LOM 1
G	Eth1/5	1,23	trunk	1G	5	C220 Server3- 1GE LOM 1
W	Eth1/51	1,20,22,23	trunk	10G	10	VPC peer link
X	Eth1/52	1,20,22,23	trunk	10G	10	VPC peer link
K	Eth1/13	20	access	1G	--	C220 Server1- Broadcom NIC 1
L	Eth1/15	20	access	1G	--	C220 Server2- Broadcom NIC 1



**Table 12** *Cisco Nexus 3048 A Ethernet Cabling Information Local Device Local Port Connection Remote (continued)*

Cable ID on both ends	Ethernet Interface	VLAN	Mode	Speed	Port Channel	Remote Device Port
M	Eth1/17	20	access	1G	--	C220 Server3- Broadcom NIC 1
Q	Eth1/14	22	access	1G	14	C220 Server1- Broadcom NIC 3
R	Eth1/16	22	access	1G	16	C220 Server2- Broadcom NIC 3
S	Eth1/18	22	access	1G	18	C220 Server3- Broadcom NIC 3
Not shown	Eth1/9	1,20,22,23	trunk	10G	15	Uplink to Infrastructure n/w
Not shown	Eth1/10	1,20,22,23	trunk	10G	17	Uplink to Infrastructure n/w
A	Eth1/25	20	access	1G	25	EMC VNXe3150 (eth2) - SPB
C	Eth1/26	20	access	1G	26	EMC VNXe3150 (eth2) - SPA

**Table 13** *Cisco Nexus 3048 B Ethernet Cabling Information Local Device Local Port Connection Remote*

Cable ID on both ends	Ethernet Interface	VLAN	Mode	Speed	Port Channel	Remote Device Port
H	Eth1/3	1, 23	trunk	1G	3	C220 Server1- 1GE LOM 2
I	Eth1/4	1,23	trunk	1G	4	C220 Server2- 1GE LOM 2
J	Eth1/5	1,23	trunk	1G	5	C220 Server3- 1GE LOM 2
Y	Eth1/51	1,20,22,23	trunk	10G	10	VPC peer link
Z	Eth1/52	1,20,22,23	trunk	10G	10	VPC peer link
N	Eth1/13	20	access	1G	--	C220 Server1- Broadcom NIC 2
O	Eth1/15	20	access	1G	--	C220 Server2- Broadcom NIC 2
P	Eth1/17	20	access	1G	--	C220 Server3- Broadcom NIC 2
T	Eth1/14	22	access	1G	14	C220 Server1- Broadcom NIC 4
U	Eth1/16	22	access	1G	16	C220 Server2- Broadcom NIC 4
V	Eth1/18	22	access	1G	18	C220 Server3- Broadcom NIC 4
Not shown	Eth1/9	1,20,22,23	trunk	10G	15	Uplink to Infrastructure n/w
Not shown	Eth1/10	1,20,22,23	trunk	10G	17	Uplink to Infrastructure n/w
B	Eth1/25	20	access	1G	25	EMC VNXe3150 (eth10) - SPB
D	Eth1/26	20	access	1G	26	EMC VNXe3150 (eth10) - SPA

Connect all the cables as outlined in [Figure 134](#) and in [Table 12](#) and [Table 13](#).

## Prepare and Configure the Cisco Nexus 3048 Switch

The following section provides a detailed procedure for configuring the Cisco Nexus 3048 switches for use in EMC VSPEX M50 solution.

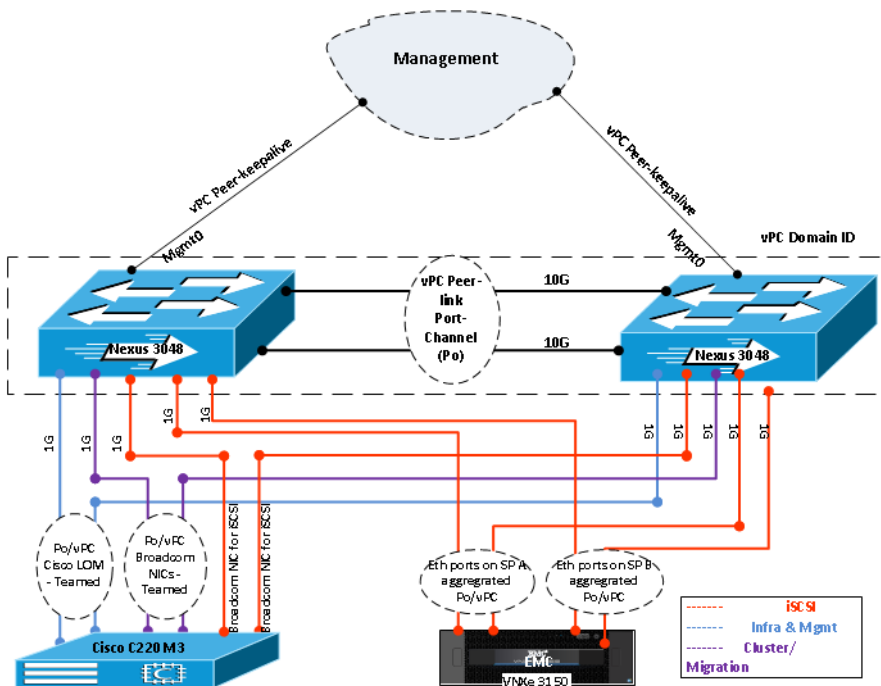


Figure 135 shows two Cisco Nexus switches configured for vPC. In vPC, a pair of switches acting as vPC peer endpoints looks like a single entity to port-channel-attached devices, although the two devices that act as logical port-channel endpoint are still two separate devices. This provides hardware 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 135**      **Networking Configuration for EMC VSPEX M50**



## Initial Setup of Cisco Nexus Switches

See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the initial setup on both Cisco Nexus 3048 switches.

## Enable Features, Jumbo Frames and Global Configuration

See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the global configuration on both Cisco Nexus 3048 switches.

## Configure VLANs

In this VSPEX M50 configuration, create 3 VLANs on both the Cisco Nexus switches using the below table as reference. Storage VLAN is for iSCSI traffic between the host iSCSI NICs and the storage array. Cluster VLAN is for the cluster communication traffic (heartbeat, CSV, and live migration) between the failover cluster nodes. VM\_traffic VLAN is for the virtual machines data traffic. Default VLAN is used by the host for management and infrastructure traffic.

**Table 14** *VLANs for EMC VSPEX Microsoft Hyper-V M50 Setup*

VLAN Name	VLAN Purpose	ID used in this document	Network Address	Host NICs participating in VLAN
Storage	For iSCSI traffic	20	10.10.20.0/24	2 Broadcom NICs
Cluster	For Live Migration	22	10.10.22.0/24	2 Broadcom NICs in team
Vm_traffic	For VM data	23	10.10.23.0/24	2 Cisco 1GigE I350 LOM in team and on trunk link
Default	For Mgmt,& Cluster	1	10.29.150.0/24	

### For Cisco Nexus A and Cisco Nexus B

1. Type `config-t`.
2. Type `vlan <storage VLAN ID>`.
3. Type `name storage`
4. Type `exit`.
5. Type `vlan <cluster VLAN ID>`.
6. Type `name cluster`
7. Type `exit`.
8. Type `vlan <vm_traffic VLAN ID>`.
9. Type `name vm_traffic`
10. Type `exit`.

## Configure Port Channels

This section describes configuring port-channels on both the Cisco Nexus switches.

### Create Port Channels

#### For Cisco Nexus A and Cisco Nexus B

From the global configuration mode, type interface Po10.

1. Type `description vPC peer-link`.
2. Type `exit`.
3. Type `interface Eth1/51-52`.
4. Type `channel-group 10 mode active`.
5. Type `no shutdown`.

6. Type exit.
7. Type interface Po3.
8. Type description <Cisco 1GigE LOM 1 on UCS Server 1 - For Nexus A>/< Cisco 1GigE LOM 2 on UCS Server 1 - For Nexus B>
9. Type exit.
10. Type interface Eth1/3
11. Type channel-group 3 mode active.
12. Type no shutdown.
13. Type exit.
14. Type interface Po4.
15. Type description < Cisco 1GigE LOM 1 on UCS Server 2 - For Nexus A>/< Cisco 1GigE LOM 2 on UCS Server 2 - For Nexus B>.
16. Type exit.
17. Type interface Eth1/4.
18. Type channel-group 4 mode active.
19. Type no shutdown.
20. Type exit.
21. Type interface Po5.
22. Type description <Cisco 1GigE LOM 1 on UCS Server 3 - For Nexus A>/< Cisco 1GigE LOM 2 on UCS Server 3 - For Nexus B >.
23. Type exit.
24. Type interface Eth1/5.
25. Type channel-group 5 mode active.
26. Type no shutdown.
27. Type exit.
28. Type interface Po14.
29. Type description <Broadcom NIC 3 on UCS Server 1- For Nexus A>/< Broadcom NIC 4 on UCS Server 1 - For Nexus B>
30. Type exit.
31. Type interface Eth1/14.
32. Type channel-group 14 mode active.
33. Type no shutdown.
34. Type exit.
35. Type interface Po16.
36. Type description <Broadcom NIC 3 on UCS Server 2- For Nexus A>/< Broadcom NIC 4 on UCS Server 2 - For Nexus B >.
37. Type exit.
38. Type interface Eth1/16.
39. Type channel-group 16 mode active.

40. Type no shutdown.
41. Type exit.
42. Type interface Po18.
43. Type description <Broadcom NIC 3 on UCS Server 3- For Nexus A>/< Broadcom NIC 4 on UCS Server 3 - For Nexus B >.
44. Type exit.
45. Type interface Eth1/18.
46. Type channel-group 18 mode active.
47. Type no shutdown.
48. Type exit.
49. Type interface Po25.
50. Type description <VNXe Storage Processor B>.
51. Type exit.
52. Type interface Eth1/25.
53. Type channel-group 25 mode active.
54. Type no shutdown.
55. Type exit.
56. Type interface Po26.
57. Type description <VNXe Storage Processor A>
58. Type exit.
59. Type interface Eth1/26.
60. Type channel-group 26 mode active.
61. Type no shutdown.
62. Type exit.

## Add Port Channel Configurations

These steps provide details for adding Port Channel configurations.

### For Cisco Nexus A and Cisco Nexus B

From the global configuration mode, type interface Po10.

1. Type switchport mode trunk.
2. Type switchport trunk allowed vlan <default VLAN ID, storage VLAN ID, cluster VLAN ID, vm\_traffic VLAN ID>.
3. Type spanning-tree port type network.
4. Type no shutdown.
5. Type exit.
6. Type interface Po3.
7. Type switchport mode trunk.
8. Type switchport trunk allowed vlan <default VLAN ID, vm\_traffic VLAN ID>.

9. Type spanning-tree port type edge.
10. Type no shut.
11. Type exit.
12. Type interface Po4.
13. Type switchport mode trunk.
14. Type switchport trunk allowed vlan <default VLAN ID, vm\_traffic VLAN ID>.
15. Type spanning-tree port type edge.
16. Type no shut.
17. Type exit.
18. Type interface Po5.
19. Type switchport mode trunk.
20. Type switchport trunk allowed vlan <default VLAN ID, vm\_traffic VLAN ID>.
21. Type spanning-tree port type edge.
22. Type no shut.
23. Type exit.
24. Type interface Po14.
25. Type switchport mode access.
26. Type switchport access vlan <cluster VLAN ID>.
27. Type spanning-tree port type edge.
28. Type no shut.
29. Type exit.
30. Type interface Po16.
31. Type switchport mode access.
32. Type switchport access vlan <cluster VLAN ID>.
33. Type spanning-tree port type edge.
34. Type no shut.
35. Type exit.
36. Type interface Po18.
37. Type switchport mode access.
38. Type switchport access vlan <cluster VLAN ID>.
39. Type spanning-tree port type edge.
40. Type no shut.
41. Type interface Po25.
42. Type switchport mode access.
43. Type switchport access vlan <storage VLAN ID>.
44. Type spanning-tree port type edge.
45. Type no shut.
46. Type interface Po26

47. Type switchport mode access.
48. Type switchport access vlan <storage VLAN ID>.
49. Type spanning-tree port type edge.
50. Type no shut.

## Configure Virtual Port Channels

These steps provide details for configuring virtual Port Channels (vPCs).

### For Cisco Nexus A and Cisco Nexus B

From the global configuration mode, type vpc domain <Nexus vPC domain ID>.

1. Type peer-keepalive destination <Nexus B mgmt0 IP> source <Nexus A mgmt0 IP>.
2. Type exit.
3. Type interface Po10.
4. Type vpc peer-link.
5. Type exit.
6. Type interface Po3.
7. Type vpc 3.
8. Type exit.
9. Type interface Po4.
10. Type vpc 4.
11. Type exit.
12. Type interface Po5.
13. Type vpc 5.
14. Type exit.
15. Type interface Po14.
16. Type vpc 14.
17. Type exit.
18. Type interface Po16.
19. Type vpc 16.
20. Type exit.
21. Type interface Po18.
22. Type vpc18.
23. Type exit.
24. Type interface Po25.
25. Type vpc .
26. Type exit.
27. Type interface Po26
28. Type vpc 26.

29. Type exit.

30. Type copy run start.

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 port-channel summary” and “show vpc” commands as shown in [Figure 137](#) and [Figure 138](#).

**Figure 136 Command for Showing VLAN Details**

```
N3048A# sh vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Po3, Po4, Po5, Po10, 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	storage	active	Po10, Po25, Po26, Eth1/9 Eth1/10, Eth1/12, Eth1/13 Eth1/15, Eth1/17
22	cluster	active	Po10, Po14, Po16, Po18, Eth1/9 Eth1/10,
23	vm_traffic	active	Po3, Po4, Po5, Po10, Eth1/9, Eth1/10

Ensure that on both switches, all required VLANs are in “active” status and right set of ports and port-channels are part of the necessary VLANs.

Port-channel configuration can be verified using “show port-channel summary” command. [Figure 137](#) shows the expected output of this command.

**Figure 137 Command for Showing Port Channel Summary**

```
N3048B(config)# sh port-channel summary
```

Flags: D - Down P - Up in port-channel (members)  
I - Individual H - Hot-standby (LACP only)  
S - Suspended r - Module-removed  
s - Switched R - Routed  
U - up (port-channel)

Group	Port-Channel	Type	Protocol	Member Ports
3	Po3(SU)	Eth	LACP	Eth1/3(P)
4	Po4(SU)	Eth	LACP	Eth1/4(P)
5	Po5(SU)	Eth	LACP	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)
25	Po25(SU)	Eth	LACP	Eth1/25(P)
26	Po26(SU)	Eth	LACP	Eth1/26(P)

```
N3048B(config)#
```

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, port channels 14, 16 and 18 are connected to the Broadcom NICs on the host, and port channels 25 and 26 are connected to the storage array. Make sure that state of the member ports of each port channel is “P” (Up in port-channel).



**Note**

The port may not come up if the peer ports are not properly configured.

Common reasons for port channel port being down are:

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

vPC status can be verified using “show vpc” command. [Figure 138](#) shows an example output.

**Figure 138** Command for Showing vPC Details

```
N3048A# 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               : secondary
Number of vPCs configured : 8
Peer Gateway           : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled

vPC Peer-link status
-----
id  Port  Status Active vlans
--  --  ---  -
1   Po10  up    1,20-23,150

vPC status
-----
id  Port  Status Consistency Reason Active vlans
--  --  ---  -
3   Po3   up    success success 1,23
4   Po4   up    success success 1,23
5   Po5   up    success success 1,23
14  Po14  up    success success 22
16  Po16  up    success success 22
18  Po18  up    success success 22
25  Po25  up    success success 20
26  Po26  up    success success 20
```

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”.

## Infrastructure Servers

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Active Directory Domain Controller

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Microsoft SQL Server

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.



## Microsoft System Center VMM

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Prepare the Cisco UCS C220 M3 Servers

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Configure Cisco Integrated Management controller (CIMC)

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Configure RAID

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Enable Virtualization Technology in BIOS

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Installing Microsoft Windows Server OS on UCS C220 M3 Servers

See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Device Driver Installation

VSPEX M50 solution contains Cisco GigE I350 LOM and quad-port Broadcom BCM5709C NetXtreme II GigE adapter.

To install device drivers follow these steps:

1. See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) and execute the steps 1 to 3 to install the chipset drivers and drivers for Cisco GigE I350 LOM.
2. To install the Broadcom drivers follow these steps:
  - a. Download the Broadcom Management Applications Installer (x64) from the URL given below and install.  
[http://www.broadcom.com/support/ethernet\\_nic/netxtremeii.php](http://www.broadcom.com/support/ethernet_nic/netxtremeii.php)
  - b. Remove any existing drivers and install the Broadcom Management Applications Installer (x64) downloaded in the above step. In addition to the Broadcom device drivers, the installer installs the management applications.

## Network Configuration

This section provides steps to configure the NIC teaming of Cisco 1 GigE I350 LOM adapters and Broadcom BCM5709C NetXtreme II GigE adapters and assign IP addresses on all the Windows host servers.

### NIC Teaming of Cisco 1 GigE LOM

See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the task. Assign an IP address to this teamed adapter from the management VLAN subnet.

### NIC Teaming of Broadcom BCM5709C NetXtreme II GigE adapter

In this section, only the NICs connected to “cluster” VLAN is teamed.



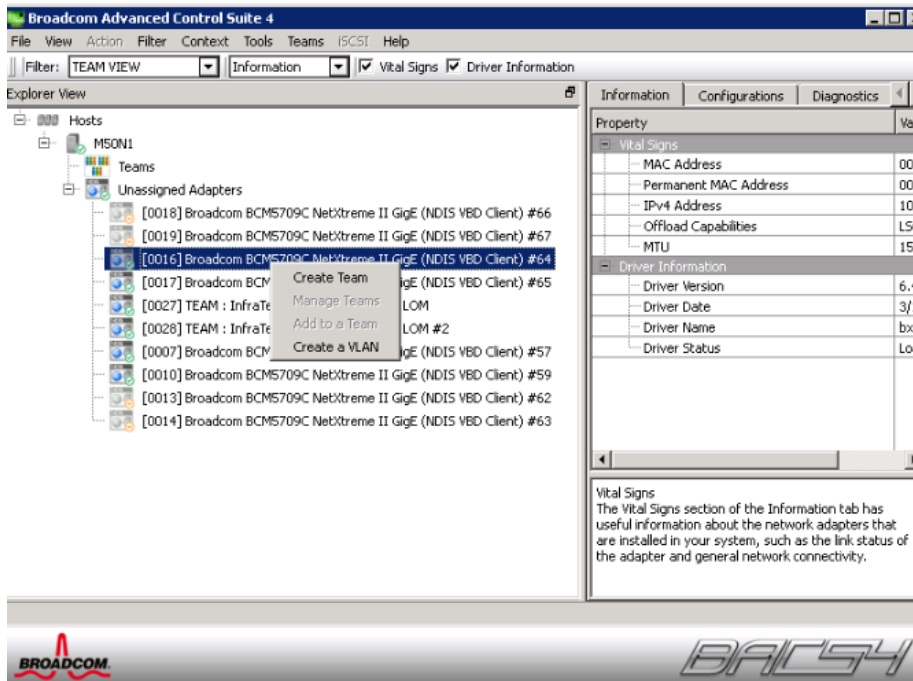
#### Note

NICs connected to the “storage” VLAN are not teamed and instead Microsoft MPIO feature is used for redundancy and load balancing.

To team the Broadcom NICs connected to the “cluster” VLAN follow these steps:

1. Click **Start > All Programs > Broadcom Advanced Control Suite 4**.
2. Choose and right-click a “NIC” and click **Create Team** as shown in [Figure 139](#).

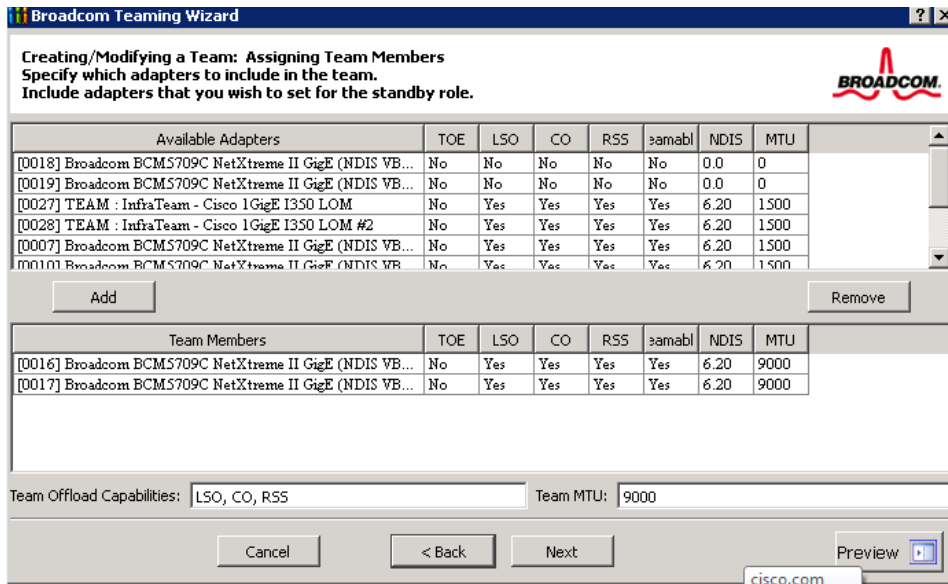
**Figure 139**      *Teaming Broadcom Adapters*



3. In the “Welcome to the Broadcom Teaming Wizard” page click **Next**.
4. Enter a Name for the Team and click **Next**.

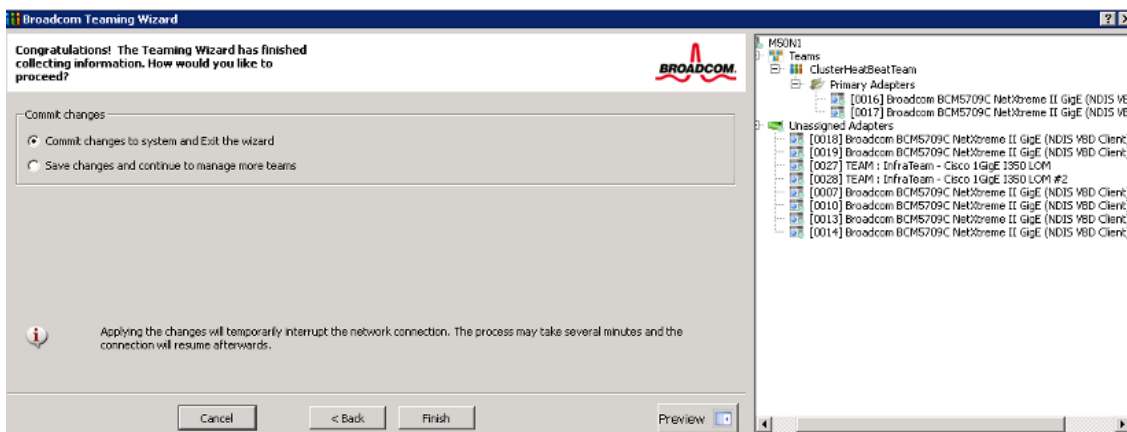
5. In the “Team Type” page, choose FEC/GEC Generic Trunking and click **Next**.
6. In the next screen, choose the second NIC (Connected to "cluster VLAN") and click **Add**. Set the MTU to 9000.

**Figure 140**      **Selecting Broadcom Adapters for Teaming**



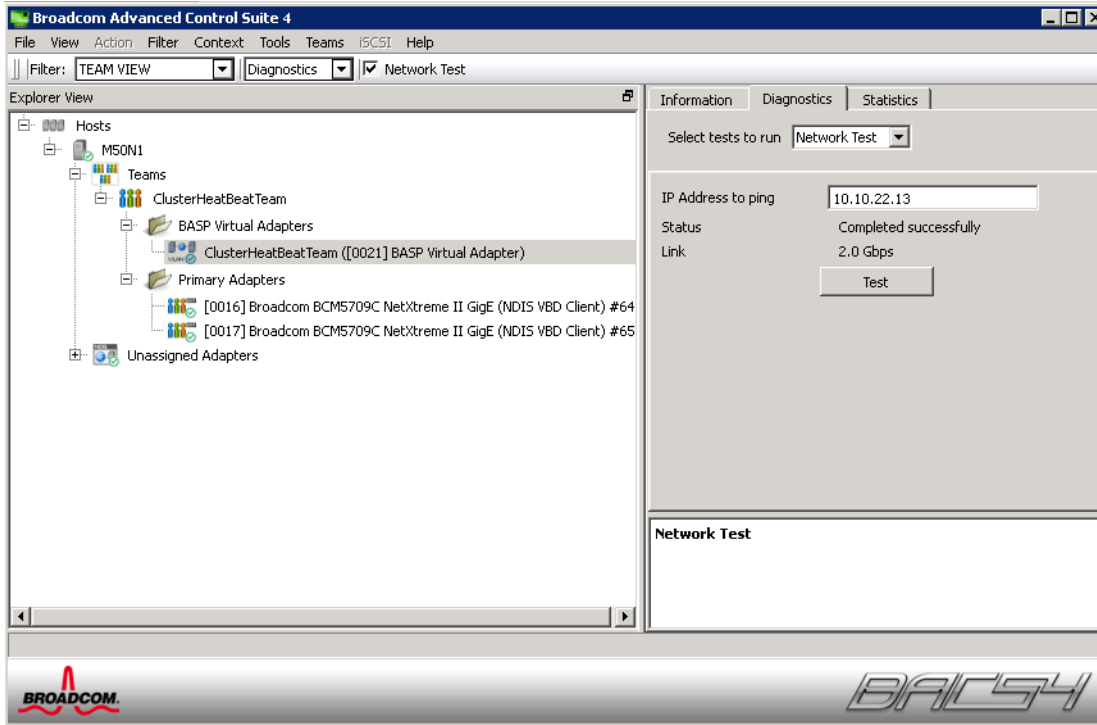
7. In the next screen click **Do not configure a standby member** and click **Next**.
8. Choose “No for Configure LiveLink” and click **Next**.
9. Choose “Skip Manage VLAN in Manage VLAN” and click **Next**.
10. Choose “Commit changes to system” and Exit the wizard.
11. Click **Preview** to validate and then click **Finish**.

**Figure 141**      **Confirmation Window to Commit Changes**



12. Validate the team as shown in [Figure 142](#).

**Figure 142** Window Showing Teamed Broadcom Adapters



13. Assign an IP address to the teamed adapter.

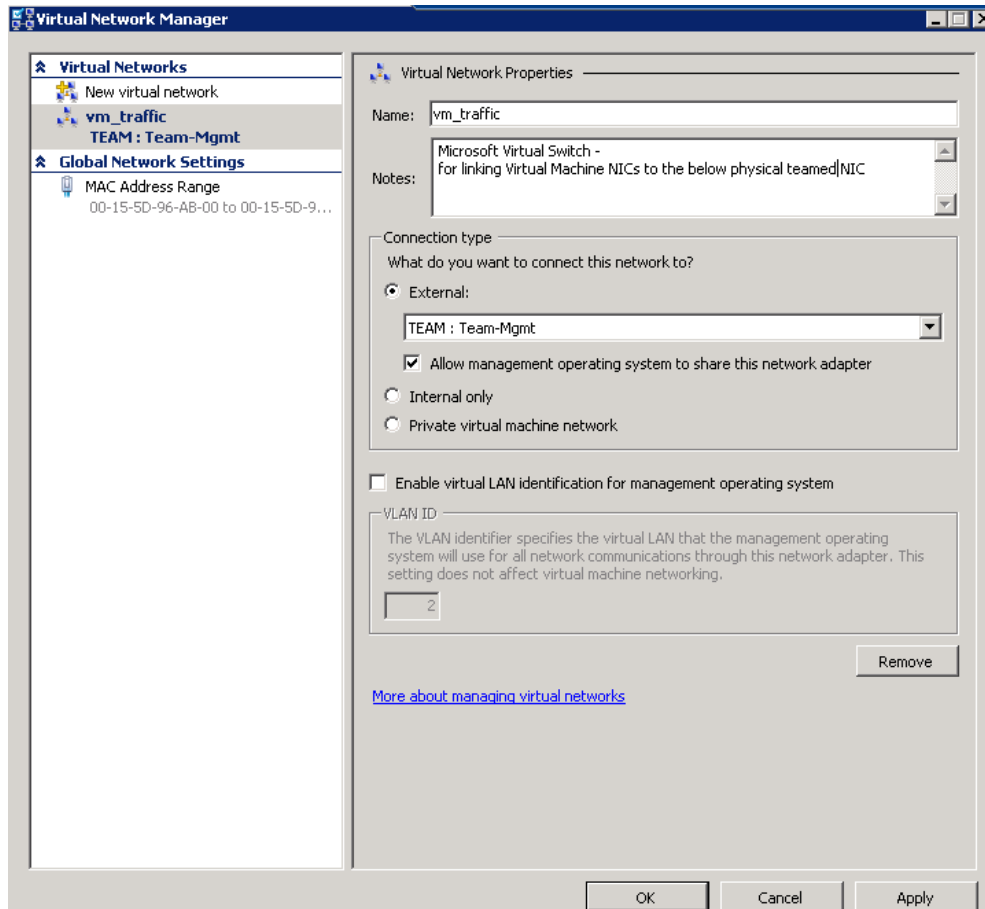
## Host Rename and Domain Join

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task


## Install Roles and Features

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) and complete the steps from 1 to 13.

In step 14, choose the teamed Cisco 1GigE I350 LOM adapter for the External Connection type for creating a virtual switch. Choose “Allow management operating system to share the network adapter”.

**Figure 143** *Microsoft Hyper-V Virtual Network Manager*

The above step also creates a virtual NIC for the host machine and retains the static IP address assigned in the earlier step for host management. Note, this NIC in management VLAN allows host management traffic.

**Figure 144** *Control Panel Network Connections*

## Enable iSCSI initiator

See the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Prepare the EMC VNXe3150 Storage

The interface and configuration of the EMC VNXe3150 is very similar to the EMC VNXe3300, so see the section “Prepare the EMC VNXe3300 storage” in [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Initial Setup of EMC VNXe

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Create Storage Pools

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task. Here you create a single storage pool using 45 disks.

## Configure Advanced features—Link Aggregation and Jumbo Frames

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Create iSCSI Servers

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Create Hosts

See the steps in the corresponding section of the [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Create Hyper-V Datastores

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task. For the 50 Virtual Machines configuration create six 750GB Hyper-V datastores for CSV and another small datastore for the cluster witness disk.

## Microsoft Windows Failover Cluster Setup

See the steps in the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## iSCSI Initiator Configuration

See the steps in the corresponding section of the [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Cluster Validation

See the corresponding section of [VSPEX M100 Configuration Details, page 23](#) to complete the task

## Failover Cluster setup

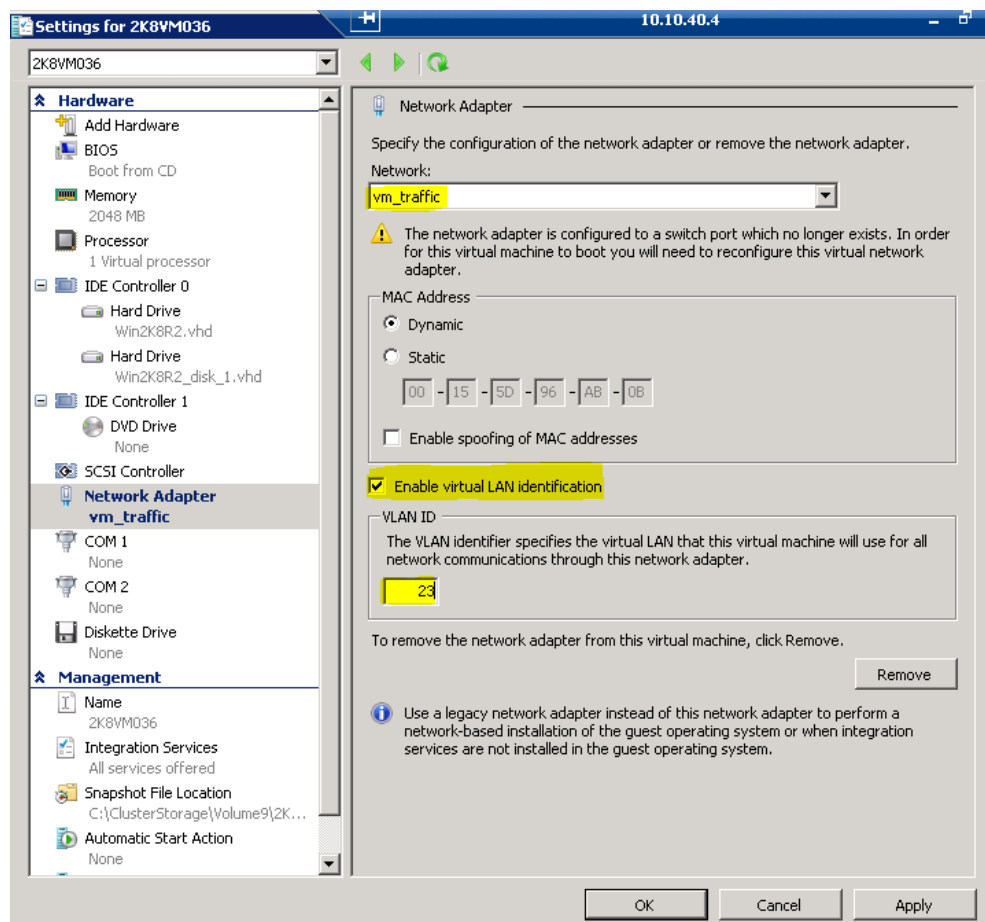
See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the task.

## Microsoft System Center-2012 VMM Configuration

See the corresponding section in the [VSPEX M100 Configuration Details, page 23](#) to complete the task. You need to ensure to connect the VM vNICs to the vm\_traffic VLAN. This can be achieved by following the steps while creating the VM or editing the settings for the VM. This step is to allow traffic from management VLAN and vm\_traffic VLAN to pass through the teamed Cisco 1GigE I350 LOM. On the other end of the Cisco Nexus switch port where the teamed adapters are connected are configured as trunk ports to allow multiple VLAN traffic.

In the settings for VM, choose “Enable Virtual LAN Identification” and enter 23 in the field for VLAN ID.

**Figure 145** Window Showing Network Adapter Settings



## Validating Cisco Solution for EMC VSPEX Microsoft Hyper-V 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.

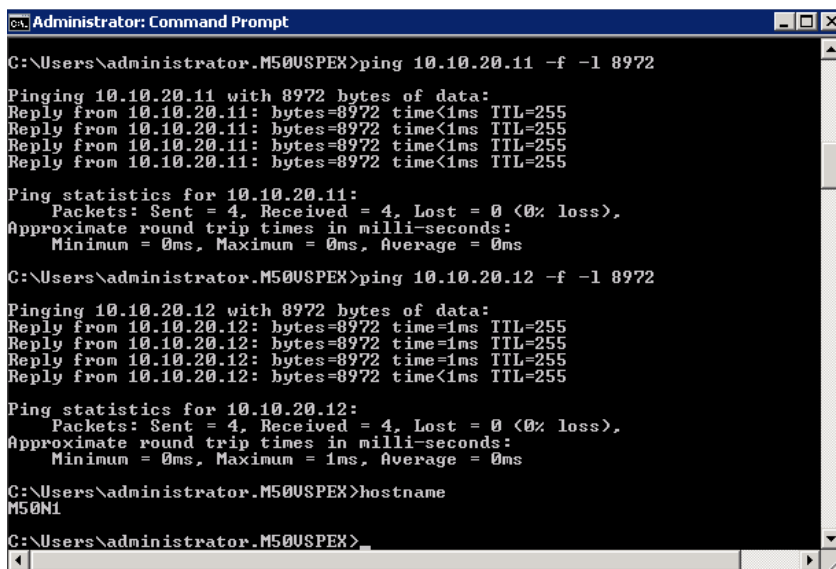
### Post Install Checklist

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

For post install checklist follow these steps:

1. Test Live Migration of VMs from one host to other using SCVMM.
2. Restart hosts and check if VMs migrate to available hosts.
3. Ping with “do not fragment switch” to validate if jumbo frames are supported end-to-end on storage and cluster VLANs.
4. Deploy a single virtual machine using the System Center Virtual Machine Manager (SCVMM) interface.

**Figure 146**      *Validating Jumbo Frames Support*



```

Administrator: Command Prompt
C:\Users\administrator.M50USPEX>ping 10.10.20.11 -f -l 8972
Pinging 10.10.20.11 with 8972 bytes of data:
Reply from 10.10.20.11: bytes=8972 time<1ms TTL=255
Reply from 10.10.20.11: bytes=8972 time<1ms TTL=255
Reply from 10.10.20.11: bytes=8972 time<1ms TTL=255
Reply from 10.10.20.11: bytes=8972 time<1ms TTL=255
Ping statistics for 10.10.20.11:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\Users\administrator.M50USPEX>ping 10.10.20.12 -f -l 8972
Pinging 10.10.20.12 with 8972 bytes of data:
Reply from 10.10.20.12: bytes=8972 time=1ms TTL=255
Reply from 10.10.20.12: bytes=8972 time=1ms TTL=255
Reply from 10.10.20.12: bytes=8972 time=1ms TTL=255
Reply from 10.10.20.12: bytes=8972 time<1ms TTL=255
Ping statistics for 10.10.20.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms
C:\Users\administrator.M50USPEX>hostname
M50N1
C:\Users\administrator.M50USPEX>

```

## Verify the Redundancy of the Solution Components

The 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. Ensure 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 Cisco Nexus switches.



2. Administratively shutdown one of the two data links connected to the storage array. Ensure that storage is still available from all the Microsoft Hyper-V hosts. Upon administratively enabling the shutdown port, the traffic should be rebalanced.
3. Reboot one of the two Cisco 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 EMC 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 Microsoft Hyper-V 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**

In 50 virtual machines architectures, there is enough head room for memory in other servers to accommodate 25 additional virtual machines. However, for 100 virtual machines solution, memory would be oversubscribed when one of the Hyper-V nodes in the cluster goes down. So, for 100 virtual machines solution, dynamic memory features should be used to oversubscribe physical memory on the remaining hosts.

## Cisco Validation Test Profile

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

**Table 15**      **VDBench Details**

Profile characteristic	Value
Number of virtual machines	50 or 100 depending on architecture
Virtual machine OS	Microsoft Windows Server 2008 R2 SP1
Processors per virtual machine	1
Number of virtual processors per physical CPU core	4
RAM per virtual machine	2 GB
Average storage available for each virtual machine	75 GB
Average IOPS per virtual machine	25 IOPS
Number of datastores to store virtual machine disks	10 CSVs
Disk and RAID type for datastores	RAID 5, 600 GB, 15k rpm, 3.5-inch SAS disks

# Bill of Material

Table 16 gives details of the components used in the CVD for 50/100 virtual machines configuration.

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

Description	Part #
UCS C220 M3 rack servers	UCSC-C220-M3S
CPU for C220 M3 rack servers	UCS-CPU-E5-2650
Memory for C220 M3 rack servers	UCS-MR-1X082RY-A
RAID local storage for rack servers	UCSC-RAID-11-C220
Cisco VIC adapter for 100 VMs solutions	N2XX-ACPCI01
Broadcom 1Gbps adapter for 50 VMs solution	N2XX-ABPCI03-M3
Cisco Nexus 5548UP switches for 100 VMs solutions	N5K-C5548UP-FA
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](http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/C220M3_SFF_SpecSheet.pdf)

# Customer Configuration Data Sheet

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

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

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

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

**Table 18**      **Microsoft Hyper-V Server Information**

Server Name	Purpose	Primary IP	Private Net (storage) addresses	
	Microsoft Hyper-V Host 1			
	Microsoft Hyper-V Host 2			
	Microsoft Hyper-V Host 3			
	Microsoft Hyper-V Host 4			

**Table 19**      **Array Information**

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

**Table 20**      **Network Infrastructure Information**

Name	Purpose	IP	Subnet Mask	Default Gateway
	Cisco Nexus 5548UP Switch A			
	Cisco Nexus 5548UP Switch B			

**Table 21**      **VLAN Information**

Name	Network Purpose	VLAN ID	Allowed Subnets
vlan-infra	Management and cluster traffic		
Vlan-vm_traffic	For VM data traffic		
vlan-storage	For iSCSI traffic		
vlan-cluster	For CSV and Live Migration		

**Table 22**      **Service Accounts**

Account	Purpose	Password (optional, secure appropriately)
	Microsoft Windows Server administrator	
	Array administrator	
	SCVMM administrator	
	SQL Server administrator	

# References

Cisco Unified Computing System:

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

Cisco UCS C-Series Servers Documentation Roadmap

<http://www.cisco.com/go/unifiedcomputing/c-series-doc>

Cisco Nexus:

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

Cisco Nexus 5000 Series NX-OS Software Configuration Guide:

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

EMC VNXe3xxx series resources

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

EMC VNX5xxx series resources

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

Network Adapter Virtualization Design (Adapter-FEX) with Cisco Nexus 5500 Switches

[http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/guide\\_c07-690080\\_ns1118\\_Networking\\_Solutions\\_White\\_Paper.html](http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/guide_c07-690080_ns1118_Networking_Solutions_White_Paper.html)

Configuring Port Channels

[http://www.cisco.com/en/US/docs/switches/datacenter/sw/5\\_x/dcnm/interfaces/configuration/guide/if\\_portchannel.html](http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/dcnm/interfaces/configuration/guide/if_portchannel.html)

Configuring Port Profiles

[http://www.cisco.com/en/US/docs/switches/datacenter/sw/5\\_x/dcnm/interfaces/configuration/guide/if\\_portprofile.html](http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/dcnm/interfaces/configuration/guide/if_portprofile.html)

Configuring vPCs

[http://www.cisco.com/en/US/docs/switches/datacenter/sw/5\\_x/dcnm/interfaces/configuration/guide/if\\_vPC.html](http://www.cisco.com/en/US/docs/switches/datacenter/sw/5_x/dcnm/interfaces/configuration/guide/if_vPC.html)

System Center 2012 - Virtual Machine Manager

<http://technet.microsoft.com/en-us/library/gg610610>

Microsoft SQL Server installation guide

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

