Guide

Microsoft SQL Server 2012 Failover Cluster on Cisco UCS with iSCSI-Based Storage Access Deployment Guide

June 2012

Contents

Executive Summary

The document describes the Microsoft SQL Server 2012 failover cluster deployment in a virtual computing environment using the Small Computer System Interface over IP (iSCSI) protocol to communicate with storage devices. The document describes how to deploy Microsoft SQL Server on iSCSI using Cisco® Data Center Virtual Machine Fabric Extender (VM-FEX) technology in the Cisco Unified Computing System™ (Cisco UCS™). The deployment scenarios discussed in this document follow Cisco UCS best practices and recommendations to help ensure that the systems are highly available and scalable and can be efficiently consolidated and centrally managed.

Introduction

A Microsoft SQL Server 2012 database on iSCSI storage offers a cost-effective solution for enterprise-level database deployments. An inexpensive yet reliable and robust storage solution, iSCSI-based storage appliances provide an easy adaption of existing networking infrastructure to access the storage enclosures. Cisco UCS can exploit the bandwidth available to provide scalable, enterprise-class storage access through the iSCSI protocol. Cisco UCS provides up to 80 Gbps of unified bandwidth for disk and network access for a single Cisco UCS 5108 Blade Server Chassis.

To reduce the system infrastructure cost, IT departments are trying to virtualize their computing, storage, and networking infrastructure. Database server consolidation enables many companies to achieve considerable cost savings, reducing the total cost of ownership (TCO). Database server consolidation can also help companies achieve the infrastructure agility they need to stay competitive and to market their solutions. A Microsoft SQL Server database on iSCSI storage can be easily consolidated on a virtualized platform such as VMware, and with the help of Cisco Data Center VM-FEX technology, each guest virtual machine can have direct access to the iSCSI device. Cisco Data Center VM-FEX technology eliminates the software switch in the hypervisor layer. Such a deployment exactly mimics the bare-metal deployment and provides an easy migration path for Microsoft SQL Server from bare metal to a VMware virtual machine deployment.

High availability is one of the primary requirements for enterprise-level database platforms because mission-critical applications cannot afford to any downtime caused by unavailable databases at the network back end. Microsoft SQL Server 2012 integrates with the new Microsoft Windows 2008 failover cluster service to offer failover clustering, providing high availability for database applications. Coupled with iSCSI storage at the virtual machine level, a clustering-enabled Microsoft SQL Server deployed on the Cisco UCS platform provides a complete back-end solution with optimal TCO and high return on investment (ROI).

iSCSI

Small Computer Systems Interface (SCSI) is a standard client-server protocol that is used to enable computers to communicate with storage devices. The iSCSI protocol transfers the SCSI packets over a TCP/IP (Ethernet) network. The most common implementation of iSCSI is over 1 or 10 Gigabit Ethernet. The iSCSI protocol provides an interoperable solution that uses the existing TCP/IP infrastructure to transport block-level storage requests. Using the iSCSI protocol, systems can connect to remote storage and use it as a physical disk even if the remote storage provider or target actually uses virtual physical disks.

An iSCSI SAN typically consists of software or hardware initiators on the host connected to an isolated Ethernet network and storage resources. Storage resources are referred to as targets. The SCSI block commands are encapsulated into Ethernet packets for transmission over IP networks at both the ends of the network by the iSCSI stack.

Advantages of iSCSI

Here are some of the main benefits of the iSCSI protocol compared to the SCSI protocol:

- iSCSI uses the existing TCP/IP network.
- iSCSI reduces total storage costs.
- iSCSI eliminates the distance limitation.
- iSCSI reduces complexity.
- iSCSI uses 10 Gigabit Ethernet.

Cisco Data Center Virtual Machine Fabric Extender Technology

Cisco Data Center VM-FEX is a Cisco technology that addresses management and performance concerns in a data center by unifying physical and virtual switch management. Cisco Data Center VM-FEX collapses virtual and physical networking into a single infrastructure. This unified infrastructure enables data center administrators to provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare-metal network traffic.

Cisco Data Center VM-FEX significantly reduces the number of network management points, enabling physical and virtual network traffic to be treated in a consistent policy-based way. Cisco Data Center VM-FEX technology helps enable a consistent operating model and visibility between physical and virtual environments, and it simplifies enforcement of security and network policy when virtual machines are moved across hosts.

Cisco Data Center VM-FEX Capabilities

The Cisco Data Center VM-FEX software extends Cisco Fabric Extender Technology (FEX Technology) to the virtual machine with the following capabilities:

- Each virtual machine includes a dedicated interface on the parent switch.
- All virtual machine traffic is sent directly to the dedicated interface on the switch.
- The software-based switch in the hypervisor is eliminated.

Advantages Cisco Data Center VM-FEX

- Simplicity
  - One infrastructure for virtual and physical resource provisioning, management, monitoring, and troubleshooting
  - Consistent features, performance, and management for virtual and physical infrastructure
- Robustness
  - Programmable, with capability to renumber VLANs without disruptive changes
  - Capability to troubleshoot and perform traffic engineering for virtual machine traffic from the physical network
- Performance
  - Near-bare-metal I/O performance with VMDirectPath with VMware vMotion
  - Delivery of the required line-rate traffic to the virtual machine

Audience

The target audience for this guide includes sales engineers, field consultants, professional services staff, IT managers, partner engineering staff, and customers who want to deploy Microsoft SQL Server on iSCSI using Cisco Data Center VM-FEX.

Hardware and Software Requirements

Cisco Unified Computing System Overview

Cisco UCS is a set of preintegrated data center components, including blade servers, adapters, fabric interconnects, and fabric extenders, that are integrated within a common embedded management system. This approach results in far fewer system components and much better manageability, operation efficiencies, and more flexibility than comparable data center platforms.

Main Differentiating Technologies

The main differentiating technologies described here are what make Cisco UCS unique and give it advantages over competing offerings. The technologies presented here are high level, and the discussions do not include the technologies (such as Fibre Channel over Ethernet [FCoE]) that support these high-level elements.

Unified Fabric

Unified fabric can dramatically reduce the number of network adapters, blade-server switches, cables, and management touch points by passing all network traffic to parent fabric interconnects, where it can be prioritized, processed, and managed centrally. This approach improves performance, agility, and efficiency and dramatically reduces the number of devices that need to be powered, cooled, secured, and managed.

Embedded Multirole Management

Cisco UCS Manager is a centralized management application that is embedded on the fabric switch. Cisco UCS Manager controls all Cisco UCS elements within a single redundant management domain. These elements include all aspects of system configuration and operation, eliminating the need to use multiple, separate element managers for each system component. Massive reduction in the number of management modules and consoles and in the proliferation of agents resident on all the hardware (which must be separately managed and updated) are important deliverables of Cisco UCS. Cisco UCS Manager, using role-based access and visibility, helps enable cross-function communication efficiency, promoting collaboration between data center roles for increased productivity.

Cisco Extended Memory Technology

Significantly enhancing the available memory capacity of some Cisco UCS servers, Cisco Extended Memory Technology helps increase performance for demanding virtualization and large-data-set workloads. Data centers can now deploy very high virtual machine densities on individual servers as well as provide resident memory capacity for databases that need only two processors but can dramatically benefit from more memory. The high-memory dual in-line memory module (DIMM) slot count also lets users more cost-effectively scale this capacity using smaller, less costly DIMMs.

Cisco Data Center VM-FEX Virtualization Support and Virtualization Adapter

With Cisco Data Center VM-FEX, virtual machines have virtual links that allow them to be managed in the same way as physical links. Virtual links can be centrally configured and managed without the complexity of traditional systems, which interpose multiple switching layers in virtualized environments. I/O configurations and network

profiles move along with virtual machines, helping increase security and efficiency while reducing complexity. Cisco Data Center VM-FEX helps improve performance and reduce network interface card (NIC) infrastructure.

Dynamic Provisioning with Service Profiles

Cisco UCS Manager delivers service profiles, which contain abstracted server-state information, creating an environment in which everything unique about a server is stored in the fabric, and the physical server is simply another resource to be assigned. Cisco UCS Manager implements role- and policy-based management focused on service profiles and templates. These mechanisms fully provision one or many servers and their network connectivity in minutes, rather than hours or days.

Cisco UCS Manager

Cisco UCS Manager is an embedded, unified manager that provides a single point of management for Cisco UCS. Cisco UCS Manager can be accessed through an intuitive GUI, a command-line interface (CLI), or the comprehensive open XML API. It manages the physical assets of the server and storage and LAN connectivity, and it is designed to simplify the management of virtual network connections through integration with several major hypervisor vendors. It provides IT departments with the flexibility to allow people to manage the system as a whole, or to assign specific management functions to individuals based on their roles as managers of server, storage, or network hardware assets. It simplifies operations by automatically discovering all the components available on the system and enabling a stateless model for resource use.

The elements managed by Cisco UCS Manager include:

- Cisco UCS Integrated Management Controller (IMC) firmware
- RAID controller firmware and settings
- BIOS firmware and settings, including server universal user ID (UUID) and boot order
- Converged network adapter (CNA) firmware and settings, including MAC addresses and worldwide names (WWNs) and SAN boot settings
- Virtual port groups used by virtual machines, using Cisco Data Center VM-FEX technology
- Interconnect configuration, including uplink and downlink definitions, MAC address and WWN pinning, VLANs, VSANs, quality of service (QoS), bandwidth allocations, Cisco Data Center VM-FEX settings, and EtherChannels to upstream LAN switches

Cisco Unified Computing System Components

Figure 1 shows the Cisco UCS components.

Figure 1. Cisco UCS Components

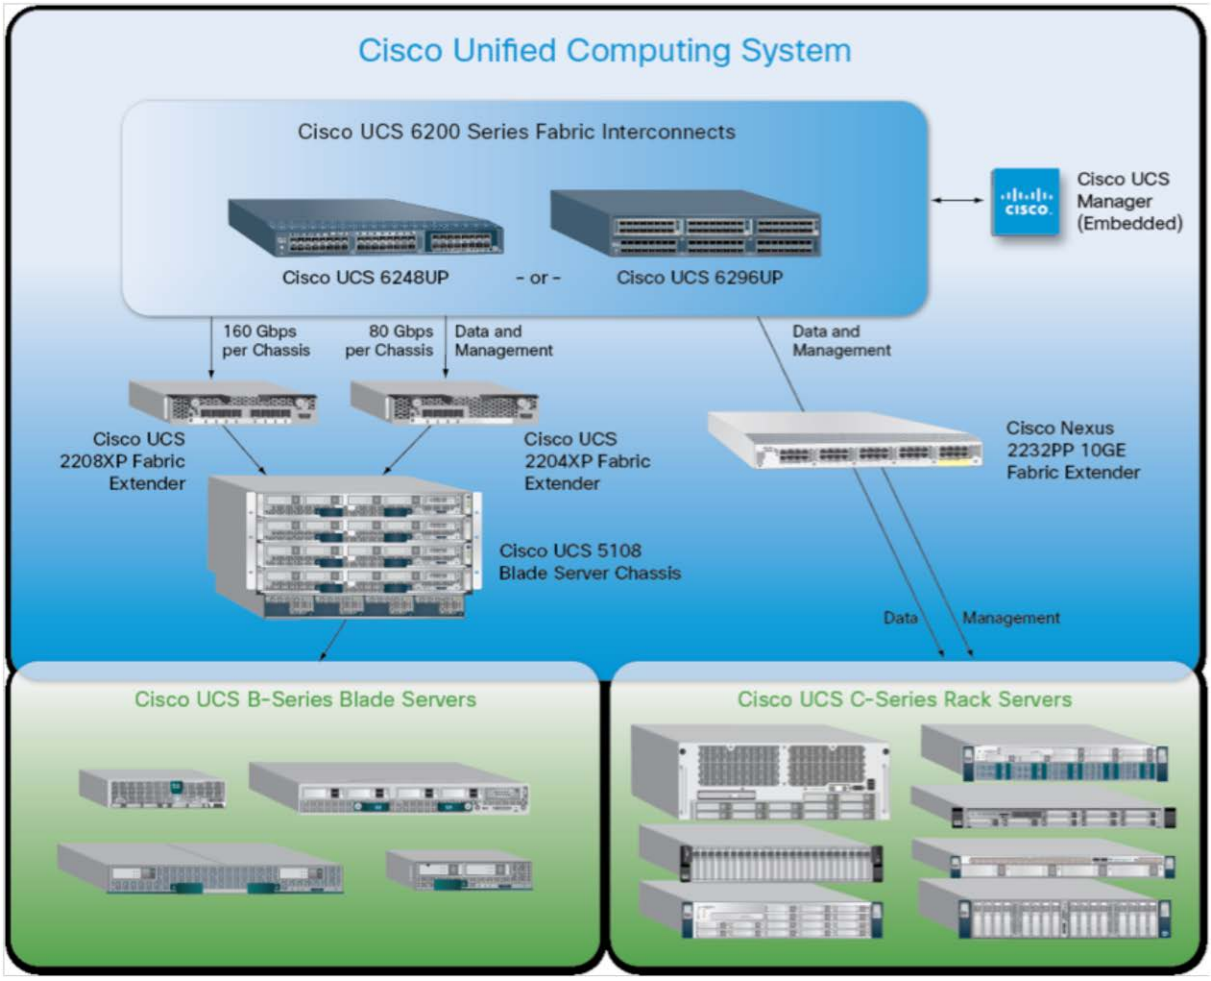

Cisco UCS is designed from the start to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards (VICs), even the number and type of I/O interfaces is programmed dynamically, making every server ready to power any workload at any time.

With model-based management, administrators manipulate a model of a desired system configuration and associate a model's service profile with hardware resources, and the system configures itself to match the model. This automation accelerates provisioning and workload migration with accurate and rapid scalability. The result is increased IT staff productivity, improved compliance, and reduced risk of failures due to inconsistent configurations.

Cisco FEX Technology reduces the number of system components that need to be purchased, configured, managed, and maintained by condensing three network layers into one. It eliminates both blade server and hypervisor-based switches by connecting fabric interconnect ports directly to individual blade servers and virtual machines. Virtual networks are now managed exactly the same way that physical networks are, but enable massive scalability. This approach represents a radical simplification compared to traditional systems, reducing

capital expenditures (CapEx) and operating expenses (OpEx) while increasing business agility, simplifying and accelerating deployment, and improving performance.

Cisco UCS Fabric Interconnects

Cisco UCS fabric interconnects create a unified network fabric throughout Cisco UCS. They provide uniform access to both networks and storage, eliminating the barriers to deployment of a fully virtualized environment based on a flexible, programmable pool of resources. Cisco fabric interconnects comprise a family of line-rate, low-latency, lossless 10 Gigabit Ethernet, IEEE Data Center Bridging (DCB), and FCoE interconnect switches. Based on the same switching technology as the Cisco Nexus® 5000 Series Switches, Cisco UCS 6100 Series Fabric Interconnects provide additional features and management capabilities that make them the central nervous system of Cisco UCS. The Cisco UCS Manager software runs inside the Cisco UCS fabric interconnects. The Cisco UCS 6100 Series Fabric Interconnects expand the Cisco UCS networking portfolio and offer higher capacity, higher port density, and lower power consumption. These interconnects provide the management and communication backbone for the Cisco UCS B-Series Blade Servers and Cisco UCS blade server chassis. All chassis and all blades that are attached to interconnects are part of a single, highly available management domain. By supporting unified fabric, the Cisco UCS 6100 Series provides the flexibility to support LAN and SAN connectivity for all blades within its domain at configuration time. Typically deployed in redundant pairs, Cisco UCS fabric interconnects provide uniform access to both networks and storage, facilitating a fully virtualized environment.

The Cisco UCS fabric interconnect portfolio currently consists of the Cisco 6100 and 6200 Series Fabric Interconnects.

Cisco UCS 6248UP 48-Port Fabric Interconnect

The Cisco UCS 6248UP 48-Port Fabric Interconnect is a one-rack-unit (1RU) 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect providing more than 1-terabit-per-second (Tbps) throughput with low latency. It has 32 fixed ports of Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE Enhanced Small Form-Factor Pluggable (SFP+) ports.

One expansion module slot can provide up to 16 additional Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS 6120XP 20-Port Fabric Interconnect

The Cisco UCS 6120XP 20-Port Fabric Interconnect is a 1RU 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect providing more than 500-Gbps throughput with very low latency. It has 20 fixed 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

One expansion module slot can be configured to support up to six additional 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS 6140XP 40-Port Fabric Interconnect

The Cisco UCS 6140XP 40-Port Fabric Interconnect is a 2RU 10 Gigabit Ethernet, IEEE DCB, and FCoE interconnect built to provide 1.04-Tbps throughput with very low latency. It has 40 fixed 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Two expansion module slots can be configured to support up to 12 additional 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS 6296UP 96-Port Fabric Interconnect

The Cisco UCS 6296UP 96-Port Fabric Interconnect is a 2RU 10 Gigabit Ethernet, FCoE, and native Fibre Channel switch offering up to 1920-Gbps throughput and up to 96 ports. The switch has 48 1/10-Gbps fixed Ethernet, FCoE, and Fibre Channel ports and three expansion slots.

One expansion module slot can provide up to 16 additional Fibre Channel, 10 Gigabit Ethernet, IEEE DCB, and FCoE SFP+ ports.

Cisco UCS 2100 and 2200 Series Fabric Extenders

The Cisco UCS 2100 and 2200 Series Fabric Extenders multiplex and forward all traffic from blade servers in a chassis to a parent Cisco UCS fabric interconnect over 10-Gbps unified fabric links. All traffic, even traffic between blades on the same chassis or virtual machines on the same blade, is forwarded to the parent interconnect, where network profiles are managed efficiently and effectively by the fabric interconnect. At the core of the Cisco UCS fabric extender are application-specific integrated circuit (ASIC) processors developed by Cisco that multiplex all traffic.

Up to two fabric extenders can be placed in a blade chassis.

- The Cisco UCS 2104XP Fabric Extender has eight 10GBASE-KR connections to the blade chassis midplane, with one connection per fabric extender for each of the chassis' eight half slots. This configuration gives each half-slot blade server access to each of two 10-Gbps unified fabric-based networks through SFP+ sockets for both throughput and redundancy. It has four ports connecting the fabric interconnect.
- The Cisco UCS 2204XP Fabric Extender has four 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2204XP has sixteen 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 80 Gbps of I/O to the chassis.
- The Cisco UCS 2208XP Fabric Extender has eight 10 Gigabit Ethernet, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis. Typically configured in pairs for redundancy, two fabric extenders provide up to 160 Gbps of I/O to the chassis.

Cisco UCS M81KR Virtual Interface Card

The Cisco UCS M81KR VIC is unique to the Cisco UCS blade system. This mezzanine adapter is designed based on a custom ASIC that is specifically intended for virtualized systems based on VMware. It uses custom drivers for the virtualized host bus adapter (HBA) and the 10 Gigabit Ethernet NIC. As is the case with the other Cisco CNAs, the Cisco UCS M81KR VIC encapsulates Fibre Channel traffic within the 10 Gigabit Ethernet packets for delivery to the fabric extender and the fabric interconnect.

Cisco UCS Virtual Interface Card 1240

A Cisco innovation, the Cisco UCS VIC 1240 is a four-port 10 Gigabit Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet.

Cisco UCS Virtual Interface Card 1280

A Cisco innovation, the Cisco UCS VIC 1280 is an eight-port 10 Gigabit Ethernet, FCoE-capable mezzanine card designed exclusively for Cisco UCS B-Series Blade Servers.

The Cisco UCS VIC 1240 and 1280 enable a policy-based, stateless, agile server infrastructure that can present up to 256 PCI Express (PCIe) standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs. In addition, the Cisco UCS VIC 1280 supports Cisco Data Center VM-FEX technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, simplifying server virtualization deployment.

Cisco UCS 5100 Series Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a 6RU blade chassis that accepts up to eight half-width Cisco UCS B-Series Blade Servers or up to four full-width Cisco UCS B-Series Blade Servers, or a combination of the two. The Cisco UCS 5108 can accept four redundant power supplies with automatic load sharing and failover and two Cisco UCS 2100 or 2200 Series Fabric Extenders. The chassis is managed by Cisco UCS chassis management controllers, which are mounted in the Cisco UCS fabric extenders and work in conjunction with Cisco UCS Manager to control the chassis and its components.

A single Cisco UCS managed domain can theoretically scale to up to 40 individual chassis and 320 blade servers. At this time, Cisco UCS supports up to 20 individual chassis and 160 blade servers.

Basing the I/O infrastructure on a 10-Gbps unified network fabric allows Cisco UCS to have a streamlined chassis with a simple yet comprehensive set of I/O options. The result is a chassis that has only five basic components:

- The physical chassis with passive midplane and active environmental monitoring circuitry
- Four power supply bays with power entry in the rear and hot-swappable power supply units accessible from the front panel
- Eight hot-swappable fan trays, each with two fans
- Two fabric extender slots accessible from the back panel
- Eight blade server slots accessible from the front panel

Cisco UCS B200 M2 Blade Servers

The Cisco UCS B200 M2 Blade Server is a half-slot, 2-socket blade server. The system uses two Intel Xeon p5600 series processors, up to 192 GB of double-data-rate-3 (DDR3) memory, two optional Small Form Factor (SFF) SAS/SSD disk drives, and a single CNA mezzanine slot for up to 20 Gbps of I/O throughput. The Cisco UCS B200 M2 Blade Server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

Cisco UCS B250 M2 Extended Memory Blade Servers

The Cisco UCS B250 M2 Extended-Memory Blade Server is a full-slot, 2-socket blade server using Cisco Extended Memory Technology. The system supports two Intel Xeon processors 5600 series, up to 384 GB of DDR3 memory, two optional SFF SAS/SSD disk drives, and two CNA mezzanine slots for up to 40 Gbps of I/O throughput. The Cisco UCS B250 M2 blade server provides increased performance and capacity for demanding virtualization and large-data-set workloads, with greater memory capacity and throughput.

Cisco UCS B230 M2 Blade Servers

The Cisco UCS B230 M2 Blade Server is a full-slot, 2-socket blade server offering the performance and reliability of the Intel Xeon processor E7-2800 product family and up to 32 DIMM slots, which support up to 512 GB of

memory. The Cisco UCS B230 M2 supports two SSD drives and one CNA mezzanine slot for up to 20 Gbps of I/O throughput. The Cisco UCS B230 M2 Blade Server platform delivers outstanding performance, memory, and I/O capacity to meet the diverse needs of virtualized environments with advanced reliability and exceptional scalability for the most demanding applications.

Cisco UCS B440 M2 High-Performance Blade Servers

The Cisco UCS B440 M2 High-Performance Blade Server is a full-slot, 2-socket blade server offering the performance and reliability of the Intel Xeon processor E7-4800 product family and up to 512 GB of memory. The Cisco UCS B440 M2 supports four SFF SAS/SSD drives and two CNA mezzanine slots for up to 40 Gbps of I/O throughput. The Cisco UCS B440 M2 blade server extends Cisco UCS by offering increased levels of performance, scalability, and reliability for mission-critical workloads.

Cisco UCS B200 M3 Blade Servers

The Cisco UCS B200 M3 Blade Server delivers performance, versatility, and density without compromise. It addresses the broadest set of workloads, from IT and web infrastructure to distributed databases. Building on the success of the Cisco UCS B200 M2 Blade Server, the enterprise-class Cisco UCS B200 M3 Blade Server further extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M3 harnesses the power of the latest Intel Xeon processor E5-2600 product family, with up to 384 GB of RAM (using 16-GB DIMMs), two disk drives, and up to dual 4x 10 Gigabit Ethernet throughput. In addition, Cisco UCS has the architectural advantage of not having to power and cool excess switches in each blade chassis. With a larger power budget per blade server, Cisco can design uncompromised expandability and capabilities in its blade servers, as evidenced by the new Cisco UCS B200 M3, with its leading memory slot and drive capacity.

VMware ESX 5.0 Architecture Overview

VMware ESX is an enterprise-level computer virtualization solution. VMware ESX is a production-proven virtualization layer that runs on physical servers that abstract processor, memory, storage, and networking resources to be provisioned to multiple virtual machines.

In the VMware ESX architecture, shown in Figure 2, the VMware Virtualization Kernel (VMkernel) is augmented by a management partition known as the console operating system or service console. The primary purpose of the console operating system is to provide a management interface with the host. Various VMware management agents are deployed in the console operating system, along with other infrastructure service agents (for example, name service, time service, and logging agents). Furthermore, individual administrative users can log in to the console operating system to run configuration and diagnostic commands and scripts.

Figure 2. VMware ESX 5.0 Architecture

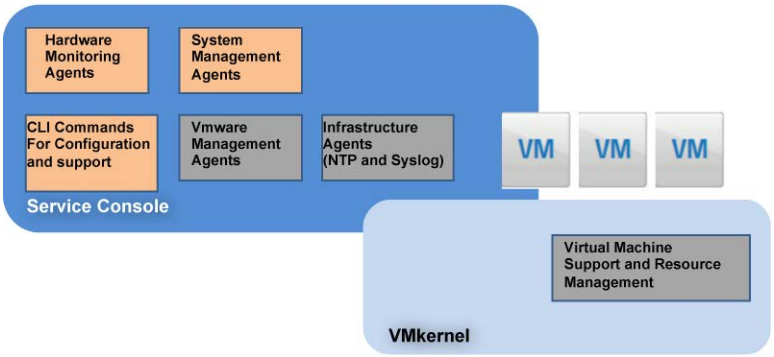

Virtualization using VMware ESX provides an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility. Virtualization allows multiple virtual machines with heterogeneous operating systems (for example, Microsoft Windows 2008 Server and Linux) and applications to run in isolation side by side on the same physical machine. A virtual machine is the representation of a physical machine by software. It has its own set of virtual hardware (RAM, CPU, NICs, hard disks, etc.) on which an operating system and applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components. VMware virtual machines contain advanced hardware features such as 64-bit computing and virtual symmetric multiprocessing. Figure 3 shows server virtualization with VMware ESX in which virtual machines directly access the network through Cisco Data Center VM-FEX.

Figure 3. VMware ESX 5.0 with Cisco Data Center VM-FEX

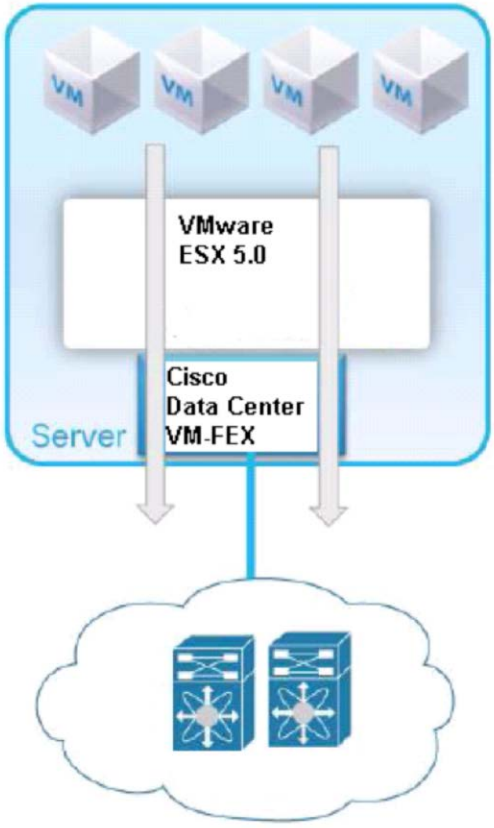

Microsoft Windows 2008 Release 2 Overview

Microsoft Windows Server 2008 Release 2 (R2) is Microsoft's multipurpose next-generation operating system designed to increase reliability and flexibility. Microsoft Windows Server 2008 R2 introduces powerful next-generation tools, built-in virtualization technology, and security and server management enhancements to efficiently manage IT operations, reduce costs, and improve performance of business-critical systems. The main improvements offered in Microsoft Windows Server 2008 R2 are:

- Improved scalability and reliability: Microsoft Windows Server 2008 R2 is specifically designed to support increased workloads while using fewer resources.
- Technology improvements: Microsoft Windows Server 2008 R2 includes technology improvements designed with Microsoft Windows 7 enterprise users in mind, augmenting the network experience, security, and manageability.
- Improved management: Microsoft Windows Server 2008 R2 provides enhanced management consoles and automation for repetitive day-to-day administrative tasks.
- Improved web application platform: Microsoft Windows Server 2008 R2 provides the capability to deliver web-based multimedia experiences efficiently and effectively, with improved administration, diagnostic, development, and application tools and lower infrastructure costs.
- Microsoft Remote Desktop Services (RDS): Microsoft RDS enables users to access applications, data, and even an entire desktop running in the data center over the network. This capability provides both the features and the robustness of a proven solution, giving users flexible access to their data and applications.

Microsoft SQL Server 2012 Overview

Microsoft SQL Server is an enterprise-class relational database management system (RDBMS) that runs on the Microsoft Windows platform and provides a wide range of data management, data integration (including data quality), and business intelligence capabilities.

Some of the main features of Microsoft SQL Server 2012 are:

- High availability, including support for active multiple secondary databases, faster failover performance, fast setup, and integrated management
- ColumnStore Index, enabling the caching of query-critical data from the data warehouse in memory-based columnar format and delivering on average 10 times the query performance of prior versions of Microsoft SQL Server
- Support for Microsoft Windows Server Core to enable better reliability and thorough cross-system security through a reduced surface area
- The new Microsoft Power View browser–based tool, along with enhancements to the Microsoft PowerPivot feature, providing rapid insight through self-service data exploration, visualization, and data mashup capabilities (users can collaborate and share these insights through Microsoft SharePoint)
- A new single business intelligence semantic model and data quality services that help provide credible, consistent data
- Support for big data through bidirectional connectors for Hadoop along with enhancements for creation of massively scalable analytics and data warehouse solutions
- Cloud-ready connectivity built with features that support hybrid IT (integrating on-premises systems with public and private clouds)
- Expanded support for unstructured data and greater interoperability with PHP, Java, and Linux

Overview of Microsoft SQL Server 2012 Deployment Model on Cisco UCS

This document describes two Microsoft SQL Server deployment models:

- Microsoft SQL Server single-instance deployment model
- Microsoft SQL Server failover cluster deployment model

Microsoft SQL Server Single-Instance Deployment Model

In the single-instance model, multiple applications are moved onto a single physical server with multiple Microsoft SQL Server instances. Each application is contained within its own Microsoft SQL Server instance. This model provides isolation of the Microsoft SQL Server instance binaries, allowing each application to be at a different patch level (major or minor version level). However, conflicts can potentially occur with the running application because system resources (mainly CPU, memory, and I/O) are shared, although tools such as the CPU affinity mask and max server memory settings can help provide resource isolation. Database system administration is isolated, but Microsoft Windows system administration shares the same host server. Each Microsoft SQL Server instance on the device can be enrolled within a Microsoft SQL Server control point for management. Another possible implementation is consolidation of several databases under a single Microsoft SQL Server instance to serve various applications. In this model, a single Microsoft SQL Server instance is shared across multiple applications, with each application having its own database.

With the single-instance approach, applications migrated from their physical server to a virtual machine environment can continue to have similar isolation with the Microsoft SQL Server database running on its own virtual machine. A single physical machine hosts multiple virtual machines, and each virtual machine hosts a single Microsoft SQL Server instance. Because a virtual machine can act as a dedicated physical machine, this approach provides an easier migration of the source environment to the consolidation environment. The single-instance deployment model is shown in Figure 4.

Figure 4. Microsoft SQL Server Single-Host Deployment Model

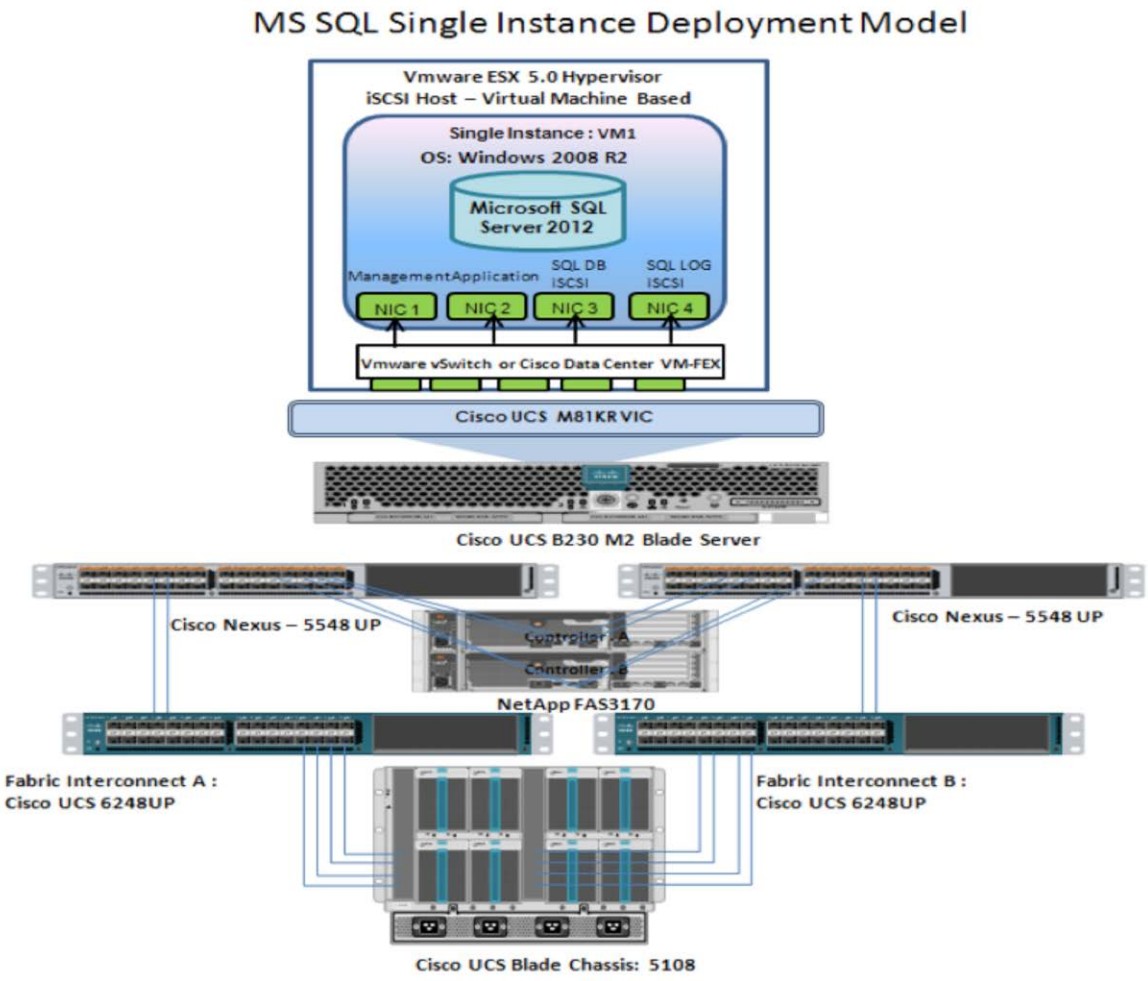

Microsoft SQL Failover Cluster Deployment Model

The Microsoft SQL cluster deployment model allows one Microsoft SQL Server to take over the tasks and responsibilities of another Microsoft SQL Server that has failed. This model helps ensure that users running mission-critical applications experience little or no downtime when such a failure occurs. Downtime can be very expensive, and the database administrator can help reduce it as much as possible. Microsoft SQL Server clustering is a high-availability technology for Microsoft SQL Server instances. It involves the sharing of server resources between one or more nodes (or servers), which have one or more shared disks grouped into logical units called resource groups. A resource group that contains at least one IP address, network name, and disk resource is called a virtual server. The cluster service arbitrates ownership of the resource groups. A single node can own a resource group and its associated resources at any given time.

The Microsoft SQL Server cluster deployment model is shown in Figure 5. Two nodes that are members of the Microsoft Windows 2008 R2 failover cluster service are deployed on VMware ESX virtual machines on two separate Cisco UCS blades. Both VMware ESX and the guest virtual machine (Microsoft Windows 2008 R2) are booted from a logical unit number (LUN) hosted on a NetApp FAS3270 with access through the iSCSI protocol.

The quorum disk for the failover cluster is also accessed through the iSCSI protocol. The database data and log files are stored on separate LUNs carved out of NetApp FAS3270. These LUNs are accessed through the iSCCI initiator originating in both the host and guest virtual machines.

This design demonstrates the flexibility of accessing storage through the iSCSI protocol with either the host-based iSCSI initiator or guest virtual machine–based iSCSI initiator. With universal passthrough (UPT) enabled on the virtual NICs (vNICs), guest virtual machines can access LUNs directly without having to go through the hypervisor layer, eliminating the additional overhead incurred while accessing critical storage resources. With UPT enabled for the iSCSI initiator, you get better response times and higher bandwidth with less CPU use on the VMware ESX host.

Figure 5. Microsoft SQL Server Failover Cluster Deployment Model

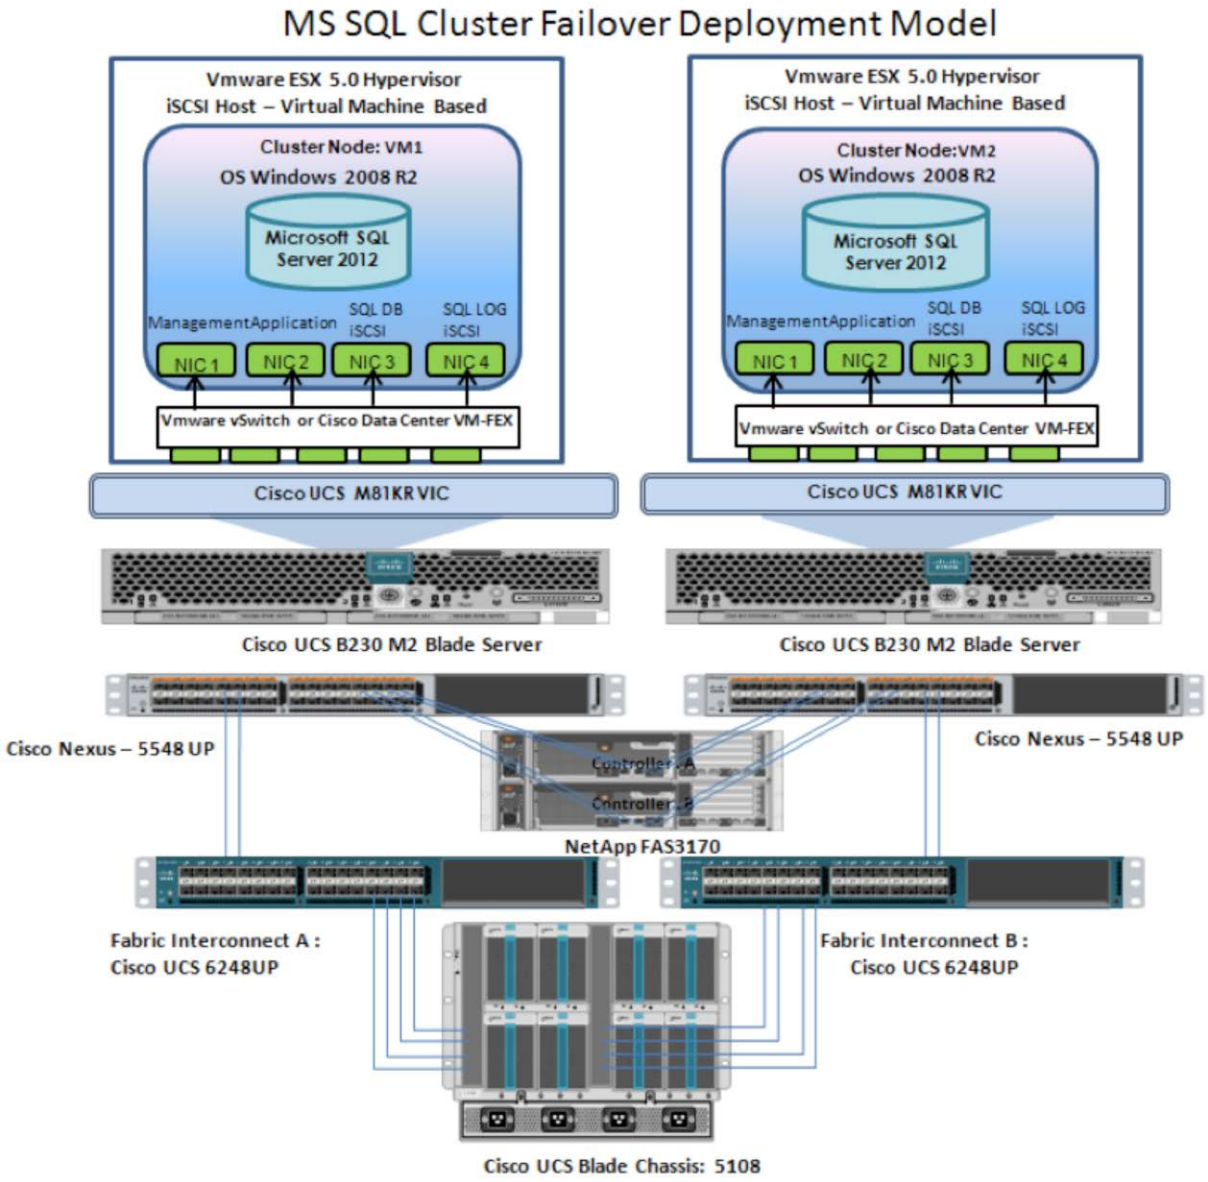

Storage Requirements for Microsoft SQL Server Database Deployment in Virtualized Environments

Storage configuration is critical to any successful database deployment. As with any physical Microsoft SQL Server deployment, the storage in virtualized environments should be sized properly to meet the database I/O requirements. The two important considerations for sizing the storage requirements are:

- Database size measured in GB
- Performance capacity measured by the number of I/O operations per second (IOPS) needed for the database to operate efficiently

To successfully design and deploy storage for a Microsoft SQL Server application, you need to understand the application's I/O characteristics and the Microsoft SQL Server I/O patterns. You need to consider parameters such as the read-to-write ratio of the application and typical I/O rates to configure the I/O characteristics of the application. The number of spindles and the speed should be configured to the maximum possible to increase storage performance. RAID 1+0 provides a better throughput for write-intensive applications. Place log files on RAID 1+0 (or RAID 1) disks for better performance and protection from hardware failures.

This validated solution uses the iSCSI protocol to access the primary database application storage.

Advantages of iSCSI Storage Implementation on the Guest Virtual Machine and VMware ESX Host

The iSCSI protocol allows SCSI commands to be sent over a TCP/IP network. iSCSI uses standard IP network equipment such as Ethernet switches and standard NICs to send SCSI block commands encapsulated in IP packets.

iSCSI offers the following advantages:

- iSCSI uses the existing IP networks and components (NICs, switches, cables, etc.), and therefore a separate network is not required to create the SAN.
- An iSCSI SAN is cost effective compared to a Fibre Channel SAN.
- An iSCSI-based SAN can coexist with the current Fibre Channel–based SAN. This feature gives customers using Fibre Channel the flexibility to scale up their SANs by adding storage capacity using an iSCSI SAN.
- An iSCSI SAN does not have any distance limitation.
- iSCSI is easy to learn, deploy, and maintain because it uses common IP-based network components.
- iSCSI is well suited for implementation of SANs in virtual server environments because it supports software initiators that make such integration easier.
- iSCSI supports the same amount of bandwidth as IP networks and therefore can provide the high bandwidth required for virtual server environments.
- iSCSI supports direct backup to tapes or disks even from virtual servers.

NetApp Storage Technologies and Benefits

NetApp solutions begin with NetApp Data ONTAP 8.0.1, the fundamental software platform that runs on all NetApp storage systems. NetApp Data ONTAP 8.0.1 is a highly optimized, scalable operating system that supports mixed network-attached storage (NAS) and SAN environments and a range of protocols, including Fibre Channel, iSCSI, FCoE, Network File System (NFS), and Common Internet File System (CIFS). It also includes a patented file system and storage virtualization capabilities. Using the NetApp Data ONTAP 8.0.1 platform, the NetApp unified storage architecture offers the flexibility to manage, support, and scale business environments by using a single set of knowledge and tools. From the remote office to the data center, customers collect, distribute, and manage data

from all locations and applications at the same time, scaling their investment by standardizing processes, reducing management time, and increasing availability. Figure 6 shows the NetApp unified storage architecture platforms.

Figure 6. NetApp Unified Storage Architecture

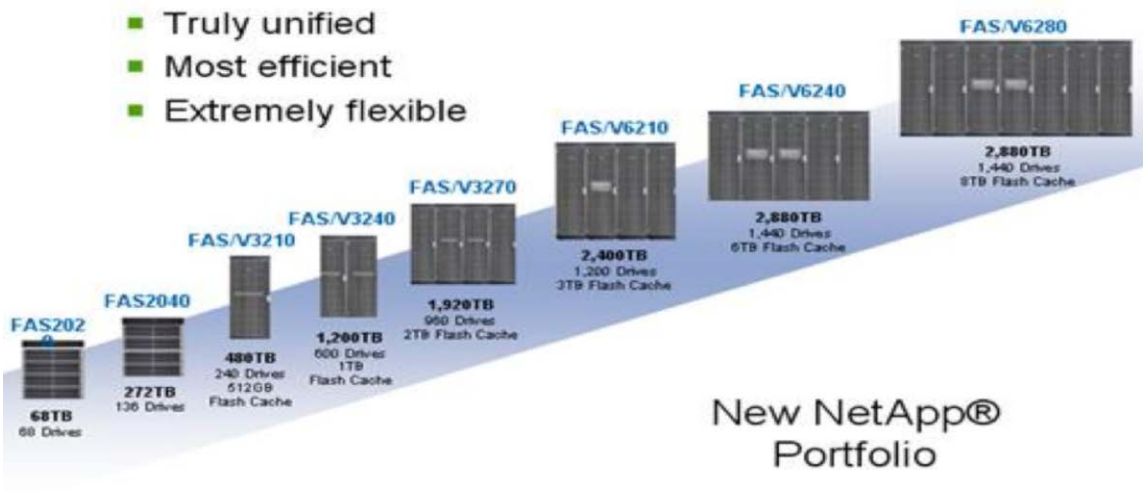

The NetApp storage hardware platform used in this solution is the NetApp FAS3270. The NetApp FAS3200 series is an excellent platform for Microsoft SQL Server 2012 deployments.

A variety of NetApp tools and enhancements are available to augment the storage platform. These tools assist in deployment, backup, recovery, replication, management, and data protection. This solution uses a subset of these tools and enhancements.

Design Topology

This section presents physical and logical high-level design considerations for Cisco UCS networking and computing with VMware ESX virtualization on NetApp storage for Microsoft SQL Server 2012 failover cluster deployments.

Cisco UCS and iSCSI Storage Network

This section explains Cisco UCS iSCSI networking and computing design considerations when deploying Microsoft SQL Server in a VMware ESX environment. In this design, the iSCSI traffic is isolated from the regular management and application data network using the same Cisco UCS infrastructure by defining logical VLAN networks to provide better data security. This design also reduces OpEx and CapEx compared to a topology in which a separate dedicated physical switch is deployed to handle iSCSI traffic.

Figure 7 presents a detailed view of the physical topology, identifying the various levels of the architecture and some of the main components of Cisco UCS in an iSCSI network design.

Figure 7. Cisco UCS Component in iSCSI Network Design

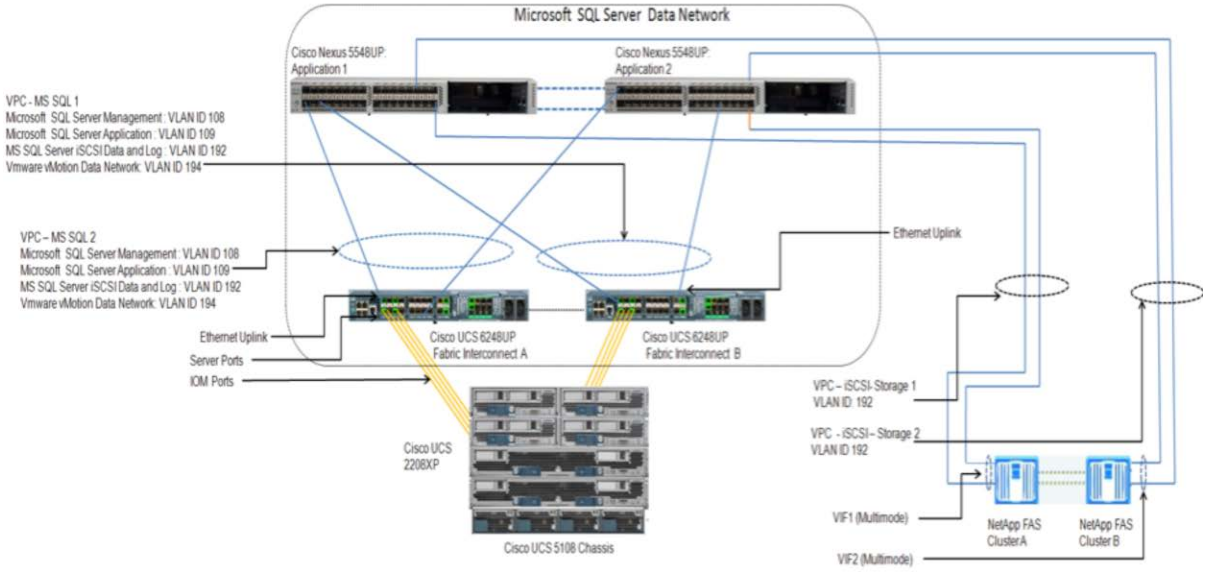

As shown in Figure 7, a pair of Cisco UCS 6248UP fabric interconnects carries both storage and network traffic from the blades with the help Cisco Nexus 5548UP. Both the fabric interconnect and the Cisco Nexus 5548UP are clustered with the peer link between them to provide high availability. Two virtual PortChannels (vPCs) are configured to provide network and storage access paths for the blades to northbound switches. Each vPC has VLANs created for application network data, iSCSI storage data, and management data paths. There is also a dedicated VLAN for VMware vMotion data traffic for VMware ESX Server.

For more information about vPC configuration on the Cisco Nexus 5548UP Switch, see http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9670/configuration_guide_c07-543563.html.

Microsoft SQL Data Network and Storage Network vPC Mapping

Table 1 shows the Cisco Nexus 5548UP vPC configurations with the vPC domains and corresponding vPC names and IDs for Microsoft SQL Servers. To provide Layer 2 and 3 switching, a pair of Cisco Nexus 5548UP Switches with upstream switching are deployed, providing high availability in the event of failure to Cisco UCS to handle management, application, and iSCSI storage data traffic. In the Cisco Nexus 5548UP topology, a single vPC feature is enabled to provide high availability, faster convergence in the event of a failure, and greater throughput.

Table 1. vPC Mapping

vPC Domain	vPC Name	vPC ID
100	vPC-MS SQL 1	101
100	vPC-MS SQL 2	102
100	vPC-iSCSI Storage 1	103
100	vPC-iSCSI Storage 2	104

In the vPC design table, a single vPC domain, Domain 100, is created across Cisco Nexus 5548UP member switches to define vPCs to carry specific network traffic. This topology defines four vPCs with IDs 101 through 104.

vPC IDs 101 and 102 are defined for traffic from Cisco UCS fabric interconnects, and vPC IDs 103 and 104 are defined for traffic to NetApp storage. These vPCs are managed within the Cisco Nexus 5548UP, which connects Cisco UCS fabric interconnects and the NetApp storage system.

When configuring the Cisco Nexus 5548UP with vPCs, be sure that the status for all vPCs is “Up” for connected Ethernet ports by running the commands shown in Figure 8 from the CLI on the Cisco Nexus 5548UP Switch.

Figure 8. PortChannel Status on Cisco Nexus 5548UP

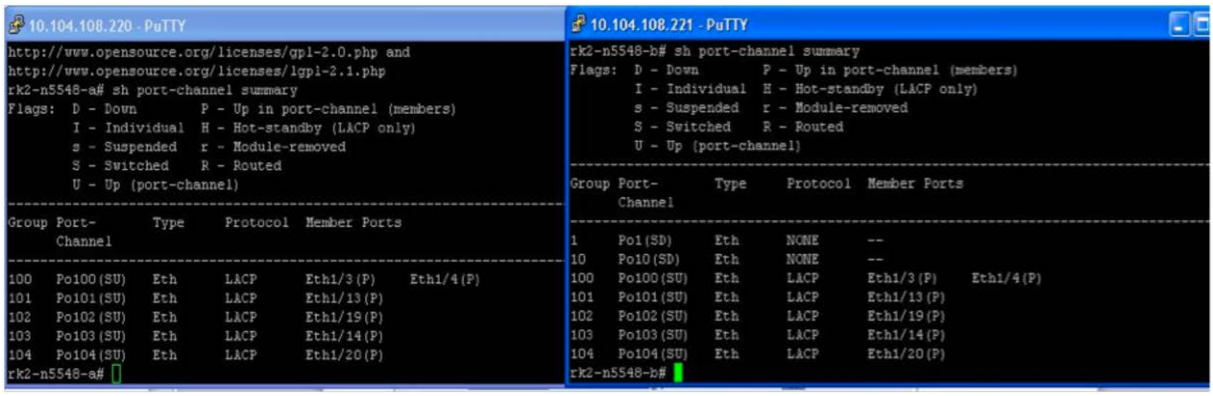

Table 2 shows the vPC configuration details for Cisco UCS 6248UP Fabric Interconnects A and B with the required vPC IDs, VLAN IDs, and Ethernet uplink ports for a Microsoft SQL Server data network design.

Table 2. Fabric Interconnects A and B (Microsoft SQL Server Data Network)

vPC Name	vPC ID	LAN Uplink Ports	VLAN ID
vPC-MS SQL 1	101	Fabric Interconnect A (Eth 1/15 and 1/16)	108 (management) 109 (SQL network) 192 (iSCSI storage) 194 (VMware vMotion)
vPC-MS SQL 2	102	Fabric Interconnect B (Eth 1/15 and 1/16)	108 (management) 109 (SQL network) 192 (iSCSI storage) 194 (VMware vMotion)

On Cisco UCS Fabric Interconnect A, Ethernet uplink ports 15 and 16 are connected to Cisco Nexus 5548UP Application 1 (port 13) and Cisco Nexus 5548UP Application 2 (port 13), which are part of vPC ID 101 and have access to VLAN IDs 108, 109, 192, and 194. The same configuration is replicated for vPC ID 102 on Fabric interconnect B, with ports 15 and 16 connected to port 14 of Cisco Nexus 5548UP Application 1 and Cisco Nexus 5548UP Application 2.

After configuring Cisco UCS 6248UP Fabric Interconnects A and B with vPCs, make sure that the status of all the PortChannels is “Enabled,” as shown in the Cisco UCS Manager screen in Figure 9.

Figure 9. Uplink Interfaces and PortChannel Status

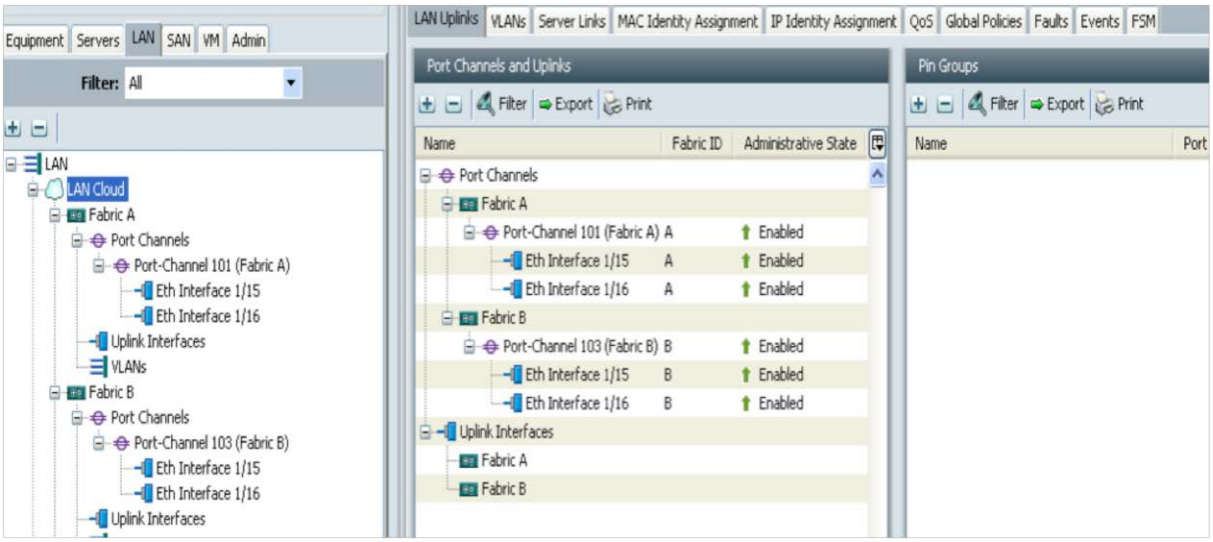

On the Cisco Nexus 5548UP Switch, a separate vPC is created to access NetApp shared iSCSI storage. The vPC is created with the vPC name and corresponding vPC ID and required VLAN IDs, as shown in Table 3.

Table 3. NetApp Storage

vPC Name	iSCSI Ports (Controllers A and B)	vPC ID	VLAN ID
vPC- iSCSI Storage 1	e1b and e1c (Controller A)	103	192
vPC- iSCSI Storage 2	e1b and e1c (Controller B)	104	192

On NetApp Storage Controller A, Ethernet 10-Gbps port e1b is connected to Cisco Nexus 5548UP Application 1 (port 19), and Ethernet port e1c is connected to Cisco Nexus 5548UP Application 2 (port 19), which are part of vPC-iSCSI Storage 1 with vPC ID 103 that allows traffic from VLAN ID 192. On NetApp Storage Controller B, Ethernet 10-Gbps port e1b is connected to Cisco Nexus 5548UP Application 1 (port 20), and Ethernet port e1c is connected to Cisco Nexus 5548UP Application 2 (port 20), which are part of vPC-iSCSI Storage 2 with vPC ID 104 that allows traffic from VLAN ID 192.

Cisco UCS Quality-of-Service System and Policy

Cisco UCS uses IEEE Data Center Bridging (DCB) to handle all traffic within Cisco UCS. This industry-standard enhancement to Ethernet divides the bandwidth of the Ethernet pipe into eight virtual lanes. System classes determine how the DCB bandwidth in these virtual lanes is allocated across the entire Cisco UCS platform.

Each system class reserves a specific segment of the bandwidth for a specific type of traffic, providing an assured level of traffic management even in an oversubscribed system. For example, you can configure the Fibre Channel Priority system class to determine the percentage of DCB bandwidth allocated to FCoE traffic.

Table 4 describes the system classes.

Table 4. System Classes

System Class	Description
• Platinum Priority • Gold Priority • Silver Priority • Bronze Priority	These classes set the quality of service (QoS) for all servers that include one of these system classes in the QoS definition in the service profile associated with the server. Each of these system classes manages one lane of traffic. All properties of these system classes are available for you to assign custom settings and policies.
Best-Effort Priority	This class sets the QoS for the lane that is reserved for basic Ethernet traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a drop policy to allow it to drop data packets if required.
Fibre Channel Priority	This class sets the QoS for the lane that is reserved for FCoE traffic. Some properties of this system class are preset and cannot be modified. For example, this class has a no-drop policy to help ensure that it never drops data packets.

QoS policies assign a system class to the outgoing traffic for a vNIC or virtual HBA (vHBA). You must include a QoS policy in a vNIC policy and then include that policy in a service profile to configure the vNIC.

To provide efficient network utilization and bandwidth control in a Microsoft SQL Server environment on VMware ESX over an iSCSI network, QoS system classes and corresponding policies are defined for network traffic generated by iSCSI storage, VMware vMotion, and the Microsoft SQL Server application and guest virtual machine management network in Cisco UCS as explained here:

- iSCSI storage traffic requires high bandwidth and a fast response time to access Microsoft SQL Server log data in the shared storage. To meet this requirement, a **SQLLog** QoS policy is created and defined with the Platinum class with the highest weight (bandwidth) and a maximum transmission unit (MTU) of 9000 for handling Microsoft SQL Server log transactions, which have a sequential I/O access pattern.
- To handle Microsoft SQL Server database data traffic, which have a more random I/O pattern and are less I/O intensive than log traffic, a **SQLDB** QoS policy is created with the Gold class with the second highest weight (bandwidth) and an MTU of 9000 to handle iSCSI packets.
- To handle VMware vMotion kernel traffic across a VMware ESX cluster during dynamic resource scheduler or manual intervention, VMware ESX requires dedicated network bandwidth for copying virtual machine active memory data. To meet this requirement, **SQLVMotion** QoS policy is created and is defined with the Silver class and with the third highest weight (bandwidth) and an MTU of 9000 to handle jumbo VMkernel packets from vNICs (static) in the service profiles in which the VMware ESX host is installed, which is a part of the VMware ESX host-based iSCSI environment.
- To handle Microsoft SQL Server application data traffic from clients on the network that are not I/O intensive compared to Microsoft SQL Server database data and log traffic and VMware vMotion traffic, a Bronze QoS class with the fourth highest weight (bandwidth) is defined on Cisco UCS.
- To handle VMware ESX host and guest virtual machine network traffic for management and operations that have lower bandwidth requirements, the Best-Effort QoS class with the least weight (bandwidth) is defined on Cisco UCS.

Note: To apply QoS across the entire system, from Cisco UCS to the upstream switches (Cisco Nexus 5548UP Switches), you need to configure similar QoS class and policy types with the right class-of-service (CoS) values that match the Cisco UCS QoS classes.

For more information, refer the Cisco Nexus QoS guide available at http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_chapter3.html - con_1150612.

Table 5 shows each QoS policy name with the corresponding priority, weight, and MTU value. These values are applied to static and dynamic vNICs in the Microsoft SQL Server deployment environment.

Table 5. Cisco UCS QoS Policy

Policy Name	Priority	Weight (Percentage)	MTU
MSSQLLog	Platinum	10	9000
MSSQLData	Gold	9	9000
VMotion	Silver	8	9000
SQLAPP	Bronze	7	9000
Management	Best Effort	5	1500

Figure 10 shows Cisco UCS QoS system class and QoS policy configurations defined for application on static and dynamic vNICs for accessing a Microsoft SQL Server iSCSI network.

Figure 10. Cisco UCS QoS System Class and QoS Policy Configuration Window

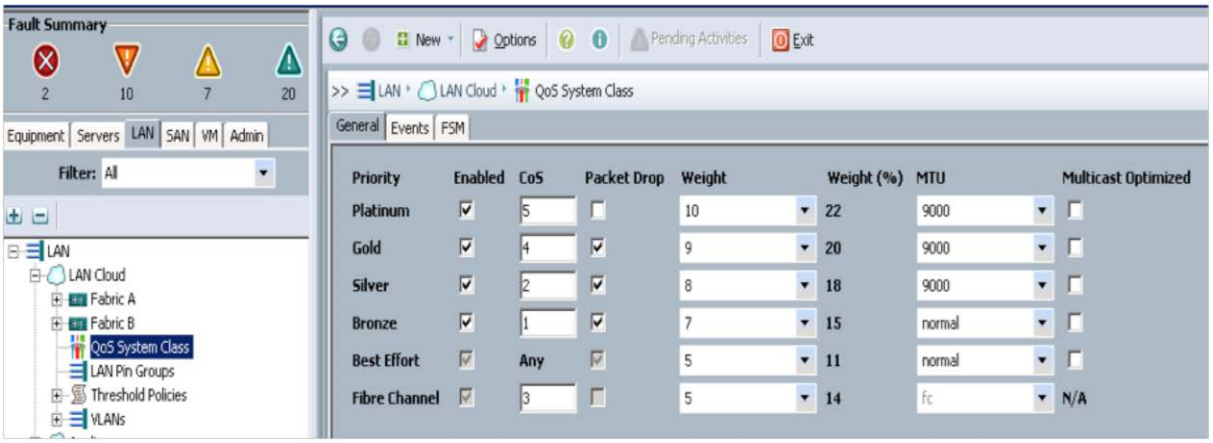

Figure 11 shows how the class priorities are applied to the named QoS policies in Cisco UCS Manager.

Figure 11. Applying Priority Classes to QoS Policy in Cisco UCS Manager

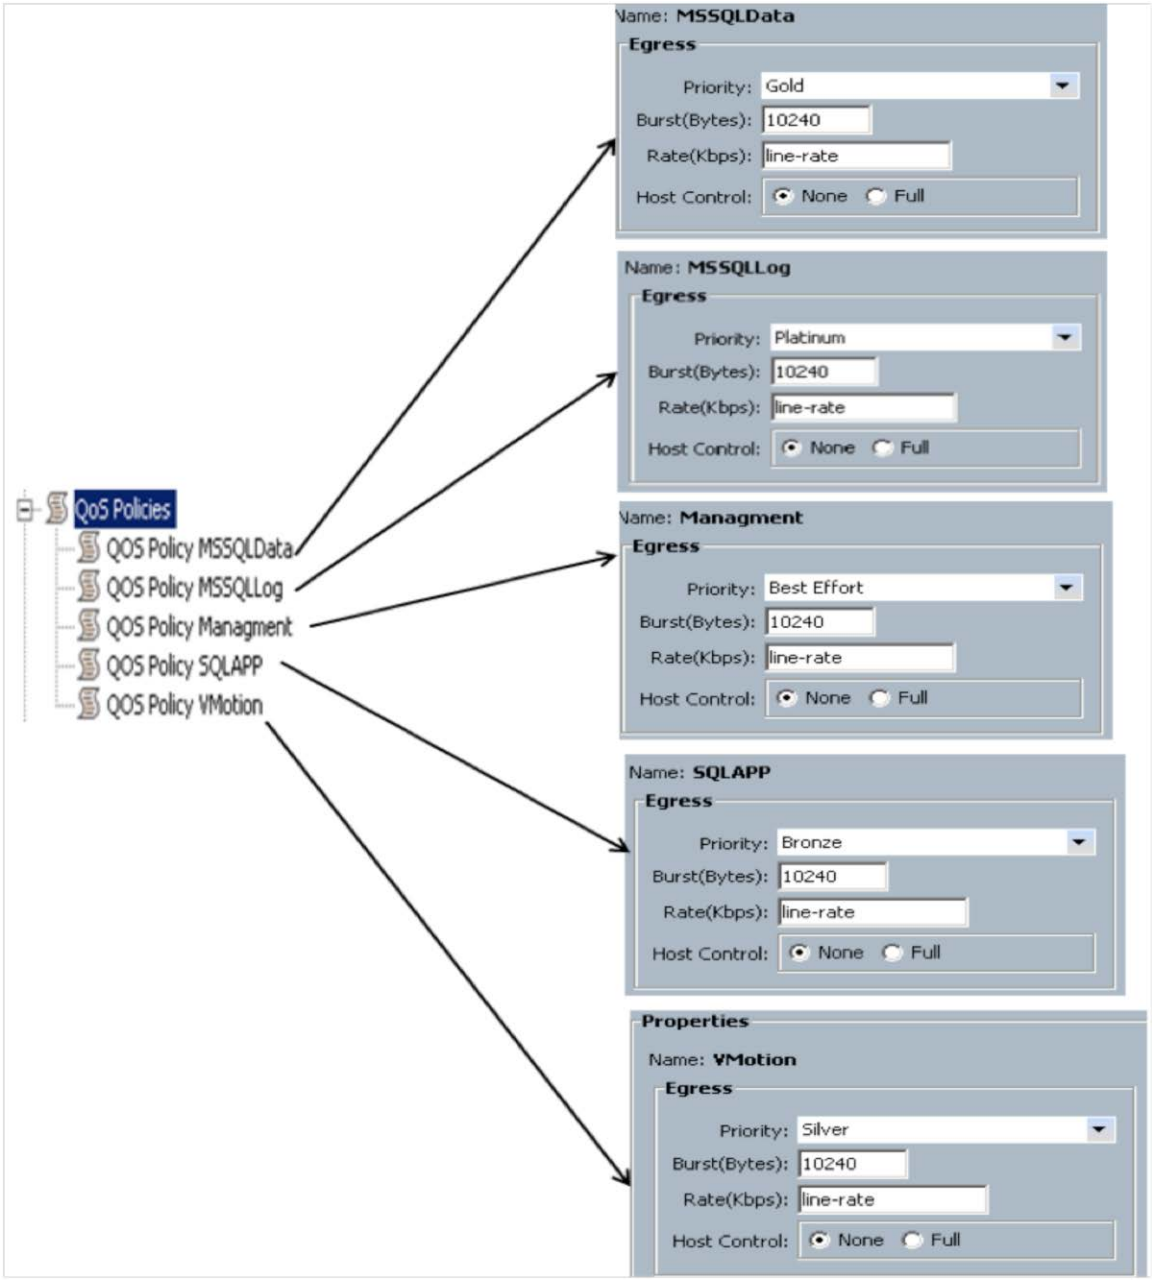

Table 6 shows Cisco UCS and Cisco Nexus 5548UP QoS mapping, with Cisco UCS QoS policy configuration values matched with Cisco Nexus 5548UP QoS policy values to achieve end-to-end QoS.

On the Cisco Nexus 5548UP, a single policy type map is defined with multiple class types, with Cisco UCS QoS matching configuration values that are applied on the global system level.

Table 6. Cisco UCS and Cisco Nexus 5548UP QoS Mapping

Cisco UCS QoS				Cisco Nexus 5548UP QoS	
Policy Name	Priority	MTU	CoS	Class Type: Network QoS and QoS	Policy Type: Network QoS and QoS
MSSQLLog	Platinum	9000	5	Network QoS: MTU 9000 and CoS 5 QoS: QoS group 5	Cisco UCS Nexus 5548UP QoS
MSSQLData	Gold	9000	4	Network QoS: MTU 9000 and CoS 4 QoS: QoS group 4	
VMotion	Silver	9000	2	Network QoS: MTU 9000 and CoS 2 QoS: QoS group 2	
SQLAPP	Bronze	9000	1	Network QoS: MTU 9000 and CoS 1 QoS: QoS group 1	
Management	Best Effort	1500	Any	Network QoS: MTU 1500	

For more information about configuration details, see http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_chapter3.html - con_1150612.

NetApp Storage Configuration Overview

This section discusses NetApp storage layout design considerations required when deploying a Microsoft SQL Server 2012 database on a VMware ESX hypervisor on Cisco UCS in an iSCSI network environment.

Figure 12 shows a high-level storage design overview on a NetApp FAS3270 cluster storage system.

Figure 12. Design Overview on a NetApp Storage Cluster

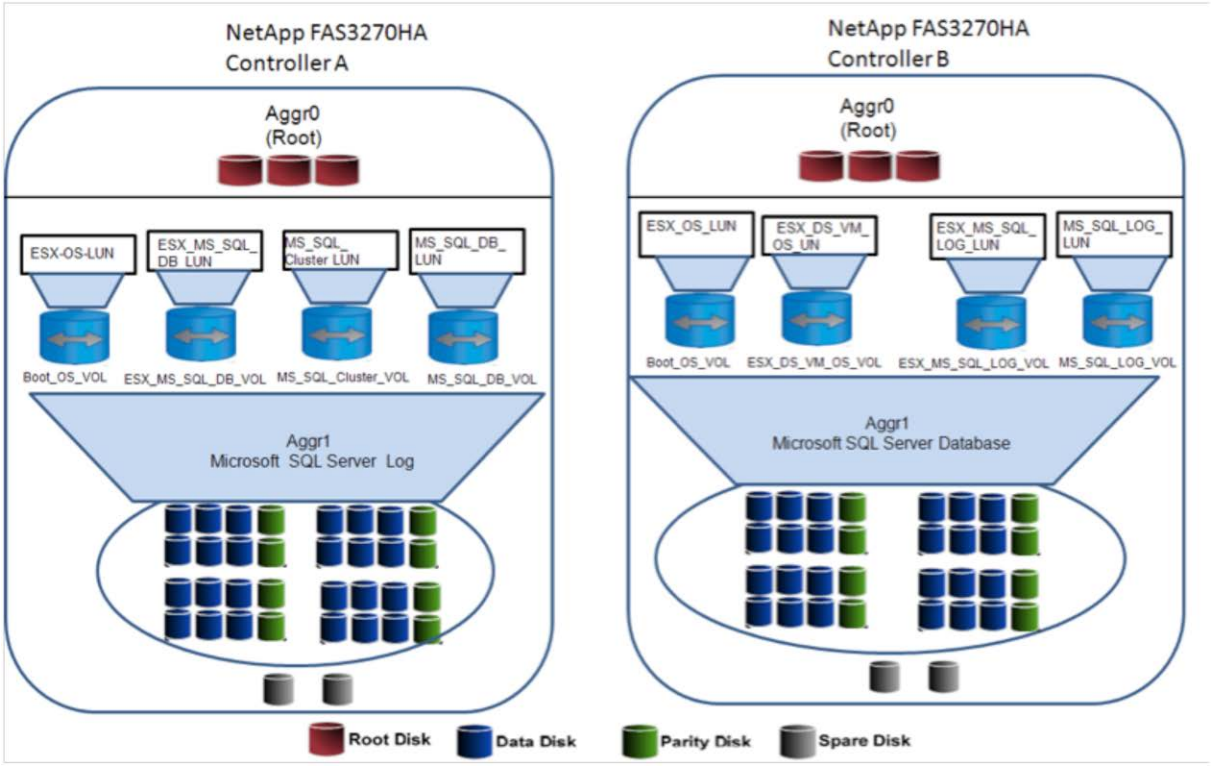

The NetApp aggregation layer provides a large virtualized pool of storage capacity and disk IOPS to be used on demand by all the virtual machines hosted on the aggregation layer. The aggregation-layer sizing is based on the storage requirements for Microsoft SQL Server data and log files to meet the storage capacity, performance, and snapshot backup requirements of an assumed workload. When sizing your environment, you need to perform the necessary planning to determine the exact storage configuration to meet your individual requirements. Aggregation layer 0 (Aggr0) is defined for hosting root NetApp Flexible Volumes (FlexVols), which use the NetApp ONTAP operating system for handling NetApp storage configurations. For detailed NetApp storage command options, see http://now.netapp.com/NOW/public/knowledge/docs/ontap/rel732/pdfs/ontap/210-04499.pdf.

Table 7 shows the NetApp storage layout with volumes and LUNs created for various purposes.

Table 7. NetApp Storage Layout with Volumes and LUNs

NetApp Storage Layout			
Aggregation and NetApp Controller	**NetApp FlexVol**	**Flexible LUN**	**Comments**
Aggr1 and Controller A	Boot_OS_VOL	ESX_OS_LUN	iSCSI boot LUN for VMware ESX host for node 1 of failover cluster with Cisco UCS B230 blade server
Aggr1 and Controller A	ESX_MS_SQL_DB_VOL	ESX_MS_SQL_DB_LUN	LUN with VMware ESX host-based iSCSI initiator for storing Microsoft SQL Server 2012 database file; VMware Virtual Machine Disk Format (VMDK) files are created for the SQL host to store the SQL data on the VMware ESX virtual machine file system
Aggr1 and Controller A	MS_SQL_Cluster_VOL	MS_SQL_Cluster_LUN	LUN with VMware ESX guest-based iSCSI initiator on Cisco Data Center VM-FEX distributed virtual switch (DVS) for storing failover cluster quorum data
Aggr1 and Controller A	MS_SQL_DB_VOL	MS_SQL_DB_LUN	LUN with VMware ESX guest-based iSCSI initiator with vSwitch or Cisco Data Center VM-FEX DVS for storing Microsoft SQL Server 2012 database file LUN
Aggr1 and Controller B	Boot_OS_VOL	ESX_OS_LUN	iSCSI boot LUN for VMware ESX host for node 2 of failover cluster with Cisco UCS B230 blade server
Aggr1 and Controller B	ESX_DS_VM_OS_VOL	ESX_DS_VM_OS_LUN	LUN with VMware ESX host-based initiator for storing guest virtual machine VMDK files
Aggr1 and Controller B	ESX_MS_SQL_LOG_VOL	ESX_MS_SQL_LOG_LUN	LUN with VMware ESX guest- or host-based iSCSI initiator for storing Microsoft SQL Server 2012 database log file
Aggr1 and Controller B	MS_SQL_LOG_VOL	MS_SQL_LOG_LUN	LUN with VMware ESX guest-based iSCSI initiator on vSwitch or Cisco Data Center VM-FEX DVS for storing Microsoft SQL Server 2012 database file LUN

Use the following commands to configure NetApp cluster storage systems on the storage controllers to implement the storage layout design described here. To run these commands, log into the storage controller through the CLI using SSH.

NetApp FAS3270HA (Controller A)

- The following command creates Aggr1 with a RAID group size of 10, 50 disks, and RAID_DP redundancy for hosting NetApp FlexVols and LUNs as shown in Table 7.

```
FAS3270HA-Controller A> aggr create aggr1 -t raid_dp -r 10 50
```

- The following commands create NetApp FlexVols on Aggr1 for hosting iSCSI LUNs as described in Table 7. These volumes are exposed to VMware ESX host and guest virtual machines for Microsoft SQL Server operations.

```
FAS3270HA-Controller A> vol create Boot_OS_VOL aggr1 50g
FAS3270HA-Controller A> vol create ESX_MS_SQL_DB_VOL aggr1 150g
```

```
FAS3270HA-Controller A> vol create MS_SQL_Cluster_VOL aggr1 150g
FAS3270HA-Controller A> vol create MS_SQL_DB_VOL aggr1 150g
```

- The following commands create LUNs on NetApp FlexVols for hosting Microsoft SQL Server database and log files.

```
FAS3270HA-Controller A> lun create -s 40g -t vmware /vol/Boot_OS_VOL/ESX_OS_LUN
FAS3270HA-Controller A> lun create -s 100g -t vmware /vol/
ESX_MS_SQL_DB_VOL/ESX_MS_SQL_DB_LUN
FAS3270HA-Controller A> lun create -s 100g -t vmware
/vol/MS_SQL_Cluster_VOL/MS_SQL_Cluster_LUN
FAS3270HA-Controller A> lun create -s 100g -t vmware
/vol/MS_SQL_DB_VOL/MS_SQL_DB_LUN
```

- The following commands create an initiator group (igroup) for mapping the VMware ESX host boot LUN and Microsoft SQL Server database data and log LUNs.

```
FAS3270HA-Controller A> igroup create -I -t vmware iSCSI-ESX -Boot iqn.2012-
01.com.vmware:ESX
FAS3270HA-Controller A> igroup create -I -t vmware ESX-MS-SQL-Node iqn.1991-
05.com.microsoft:VM
```

- The following commands map LUNs to specific igroups to access the VMware ESX host boot LUN and Microsoft SQL Server database data and log LUNs.

```
FAS3270HA-Controller A>
lun map /vol/Boot_OS_VOL/ESX_OS_LUN  iSCSI-ESX-Boot
FAS3270HA-Controller A>
lun map /vol/ESX_MS_SQL_DB_VOL/ESX_MS_SQL_DB_LUN  ESX-MS-SQL-Node
FAS3270HA-Controller A>
lun map /vol/MS_SQL_Cluster_VOL/MS_SQL_Cluster_LUN ESX-MS-SQL-Node
FAS3270HA-Controller B>
lun map /vol/ MS_SQ_DB_VOL/MS_SQL_DB_LUN  ESX-MS-SQL-Node
```

After successfully running these commands, you can verify the storage configuration using the NetApp Filter view, as shown in Figure 13.

Figure 13. Verification of Storage Configuration

NetApp FAS3270HA (Controller B)

- The following command creates Aggr1 with a RAID group size of 10, 50 disks, and RAID_DP redundancy for hosting NetApp FlexVols and LUNs as shown in Table 7.

```
FAS3270HA-Controller B> aggr create aggr1 -t raid_dp -r 10 50
```

- The following commands create NetApp FlexVols on Aggr1 for hosting iSCSI LUNs as described in Table 7. These volumes are exposed to VMware ESX host and guest virtual machines for Microsoft SQL Server operations.

```
FAS3270HA-Controller B> vol create Boot_OS_VOL aggr1 50g
FAS3270HA-Controller B> vol create ESX_DS_VM_OS_VOL aggr1 150g
FAS3270HA-Controller B> vol create ESX_MS_SQL_LOG_VOL aggr1 150g
FAS3270HA-Controller B> vol create MS_SQL_LOG_VOL aggr1 50g
```

- The following commands create LUNs on NetApp FlexVols for hosting Microsoft SQL Server database and log files.

```
FAS3270HA-Controller B>
lun create -s 30g -t vmware /vol/Boot_OS_VOL/ESX_OS_LUN
FAS3270HA-Controller B>
lun create -s 100g -t vmware /vol/ESX_DS_VM_OS_VOL/ESX_DS_VM_OS_LUN
FAS3270HA-Controller B>
lun create -s 100g -t vmware  /vol/ESX_MS_SQL_LOG_VOL/ESX_MS_SQL_LOG_LUN
FAS3270HA-Controller B>
```

```
lun create –s 5g –t vmware /vol/MS_SQL_LOG_VOL/MS_SQL_LOG_LUN
```

- The following commands create an igroup for mapping the VMware ESX host boot LUN and Microsoft SQL Server database data and log LUNs.

```
FAS3270HA-Controller B> igroup create –I –t vmware iSCSI-ESX –Boot iqn.2012-01.com.vmware:ESX
FAS3270HA-Controller B> igroup create –I –t vmware ESX-MS-SQL-Node iqn.1991-05.com.microsoft:VM
```

- The following commands map LUNs to specific igroups to access the VMware ESX host boot LUN and Microsoft SQL Server database data and log LUNs.

```
FAS3270HA-Controller B>
lun map /vol/Boot_OS_VOL/ESX_OS_LUN  iSCSI-ESX-Boot
FAS3270HA-Controller B>
lun map /vol/ESX_DS_VM_OS_VOL/ESX_DS_VM_OS_LUN  iSCSI-ESX-Boot
FAS3270HA-Controller B>
lun map /vol/ESX_MS_SQL_LOG_VOL/ESX_MS_SQL_LOG_LUN ESX-MS-SQL-Node
FAS3270HA-Controller B>
lun map vol/ MS_SQL_LOG_VOL/MS_SQL_LOG_LUN ESX-MS-SQL-Node
```

After successfully running these commands, you can verify the storage configuration using the NetApp Filter view, as shown in Figure 14.

Figure 14. Verification of Storage Configuration

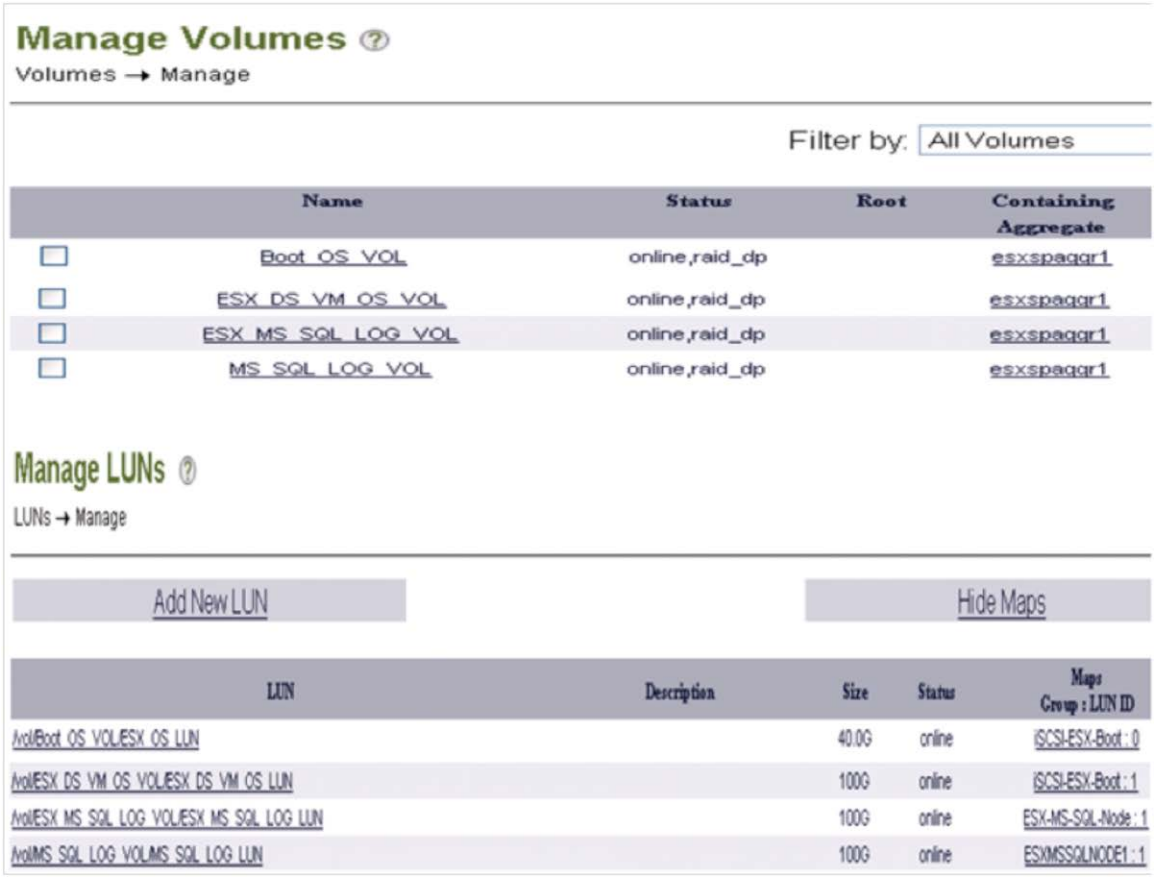

NetApp Multimode Virtual Interfaces

The NetApp multimode virtual interface (VIF) feature is enabled on NetApp storage systems on 10 Gigabit Ethernet ports for configuring the iSCSI target through which LUNs are exposed over the iSCSI protocol to host iSCSI initiators (VMware ESX host and guest virtual machines).

Figure 15 shows an iSCSI vPC-enabled network design on Cisco Nexus 5548UP and NetApp FAS3270HA Controllers A and B to access a Microsoft SQL Server data network.

Figure 15. iSCSI vPC Enabled on Cisco Nexus 5548UP

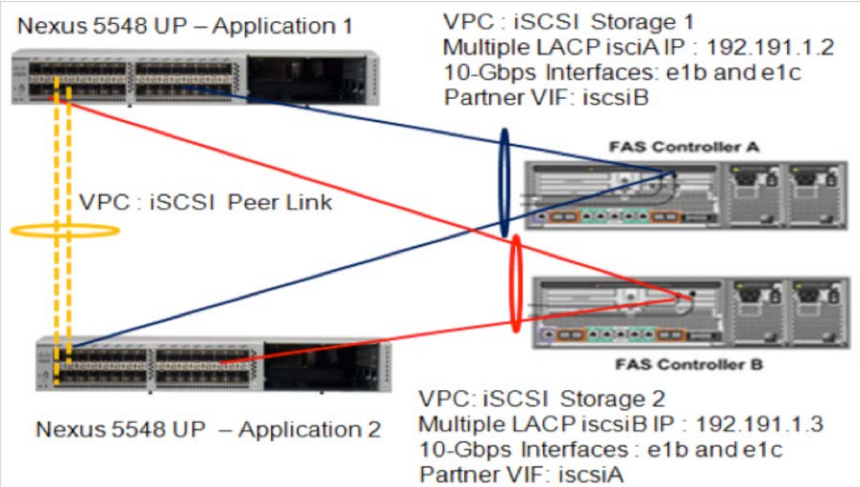

The vPC design layout for Cisco Nexus 5548UP Switches and corresponding NetApp cluster storage system multimode VIFs is as follows:

- Cisco Nexus 5548UP Application 1 and Cisco Nexus 5548UP Application 2 are part of the vPC domain and have two vPCs: vPC iSCSI Storage 1 and vPC iSCSI Storage 2 as described in the above.
- vPC iSCSI Storage 1 has NetApp FAS3270HA (Controller A) 10 Gigabit Ethernet Interfaces e1b and e1c as member ports and is connected to Cisco Nexus 5548UP Application 1 and Cisco Nexus 5548UP Application 2 switches.
- vPC iSCSI Storage 2 has NetApp FAS3270HA (Controller B) 10 Gigabit Ethernet Interfaces e1b and e1c as member ports and is connected to Cisco Nexus 5548UP Application 1 and Cisco Nexus 5548UP Application 2 vPC switches.
- On NetApp FAS3270HA (Controller A) multilevel dynamic VIF, iSCSI A is created on 10 Gigabit Ethernet Interfaces e1b and e1c and has the MTU set to 9000 with jumbo frames enabled for accessing storage using the iSCSI protocol. VIF iSCSI A is configured with cluster failover enabled on the VIF, and the iSCSI B VIF IP address is set on NetApp FAS3270HA (Controller B).
- On NetApp FAS3270HA (Controller B), multilevel dynamic VIF iSCSI B is created on 10 Gigabit Ethernet Interfaces e1b and e1c and has the MTU set to 9000 with jumbo frame enabled for accessing storage using the iSCSI protocol. VIF iSCSI B ifys configured with cluster failover enabled on the VIF, and the iSCSI A VIF IP address is set on NetApp FAS3270HA (Controller A),
- On NetApp FAS3270HA (Controllers A and B), iSCSI is enabled on e1b and e1c 10 Gigabit Ethernet interfaces for accessing storage through the iSCSI protocol from the VMware ESX host or guest virtual machine–level software initiator.

Note: Note: On the Cisco Nexus 5548UP upstream switch, ensure that the correct QoS class and MTU value with policy types are applied to the Port Channel Ports (eth19 and eth 20). Port channels are connected to the NetApp FAS3270HA (Controllers A and B), 10 Gigabit Ethernet interfaces (e1b and e1c), to allow network packets to be tagged from Nexus 5548 fabric. This is done because NetApp Storage will not tag any network packets with MTU and QoS values.

Following commands shows how to configure the COS on Nexus 5548 for untagged packets originating from storage on the Port Channels.

CiscoNexus5548UPApplication1

```
Switch# Configure Terminal
Switch(Conf)# Interface port channel 103
Switch(Conf-if)#untagged cos 5
Switch# sh policy-map type qos

Switch# Configure Terminal
Switch(Conf)# Interface port channel 104
Switch(Conf-if)#untagged cos 4
Switch# sh policy-map type qos
```

CiscoNexus5548UPApplication2

```
Switch# Configure Terminal
Switch(Conf)# Interface port channel 103
Switch(Conf-if)#untagged cos 5
Switch# sh policy-map type qos

Switch# Configure Terminal
Switch(Conf)# Interface port channel 104
Switch(Conf-if)#untagged cos 4
Switch# sh policy-map type qos
```

For more information, see http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_chapter3.html.

NetApp VIF Configuration Details

The following are the NetApp CLI commands for configuring the multilevel dynamic VIF on NetApp FAS3270HA (Controllers A and B) cluster storage systems.

NetApp FAS3270HA (Controller A)

```
FAS3270HA-Controller A> iscsi start
FAS3270HA-Controller A > ifgrp create lacp iscsiA
FAS3270HA-Controller A > ifgrp add iscsiA e1a e1b
FAS3270HA-Controller A > ifconfig iscsiA mtusize 9000 192.191.1.2 netmask
255.255.255.0 partner iscsiB up
```

NetApp FAS3270HA (Controller B)

```
FAS3270HA-Controller B> iscsi start
FAS3270HA-Controller B > ifgrp create lacp iscsiA
FAS3270HA-Controller B > ifgrp add iscsiA e1a e1b
FAS3270HA-Controller B > ifconfig iscsiB mtusize 9000 192.191.1.3 netmask
255.255.255.0 partner iscsiA up
```

Make sure that the MTU is set to 9000 and that jumbo frames are enabled on the Cisco UCS static and dynamic vNICs and on the upstream Cisco Nexus 5548UP Switches.

VMware ESX iSCSI Boot

This section describes the Cisco UCS service profile design for deploying the VMware ESX iSCSI boot OS from the NetApp shared iSCSI target on the Cisco UCS B-Series server. In this deployment, the Cisco UCS M81KR VIC is used for iSCSI SAN bootup of the VMware ESX OS from the NetApp iSCSI target.

The following steps show the basic configuration on the service profile to enable VMware ESX 5.0 iSCSI SAN bootup on the Cisco UCS B230 blade server from the NetApp iSCSI target. For more information about the configuration steps for deploying VMware ESX iSCSI bootup, see the Cisco UCS CLI and GUI detailed configuration steps at http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/cli/config/guide/2.0/b_UCSM_CLI_Configuration_Guide_2_0.pdf.

1. Create service profiles **ESX-VMDirectPath-SQL** and **ESX-VMDirectPath-SQL1** and associate them with Cisco B230 blades using the Cisco UCS M81KR VIC to install VMware ESX 5.0 from the iSCSI target on the NetApp FAS3270. Figure 16 shows the creation of these service profiles.

Figure 16. Creating Service Profiles

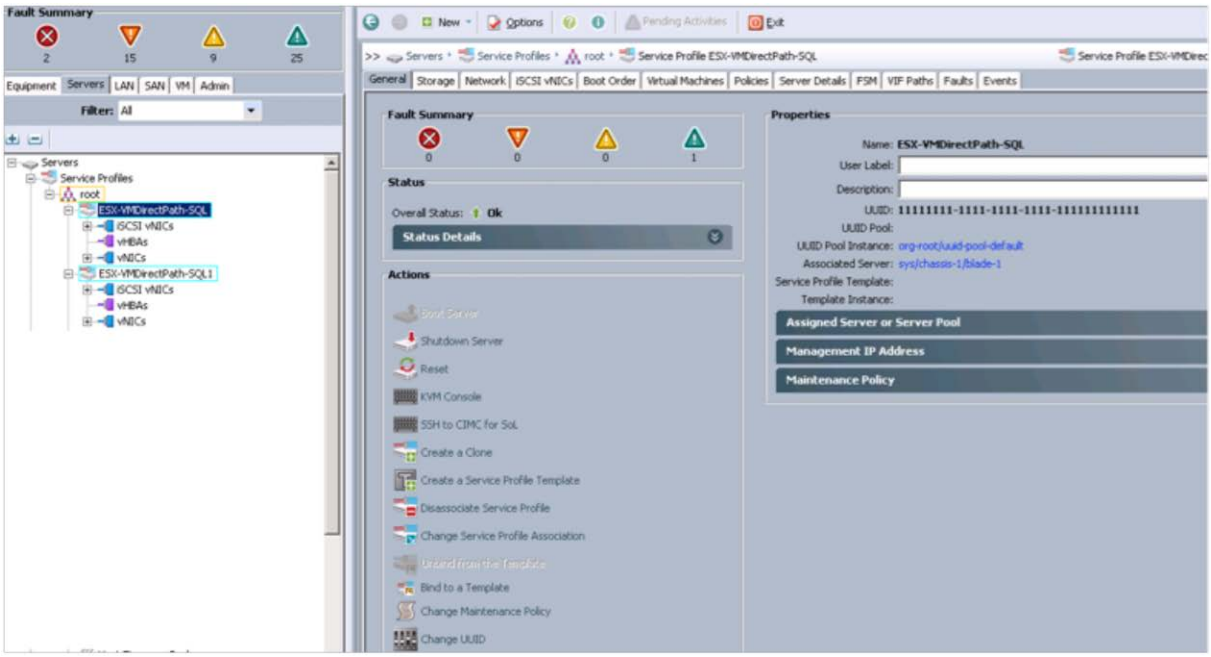

2. On the Service Profiles tab for the newly created service profiles, create two static vNICs, **vNIC2** and **vNIC3**, on Fabric A and Fabric B, respectively, with the MTU value set to 9000, without fabric failure and network VLAN access to VLAN ID 192 (iSCSI storage), as shown in Figure 17.

Figure 17. Creating Static vNICs on Fabric Interconnects

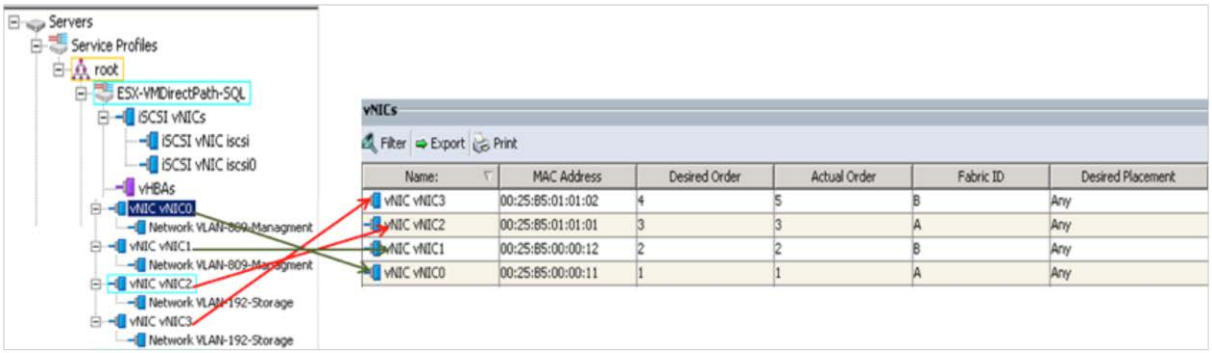

3. To access VMware ESX for management purpose, create two separate static vNICs (**vNIC0** and **vNIC1**) with the appropriate VLAN ID. These vNICs will provide uplinks to the VMware ESX vSwitch and Cisco Data Center VM-FEX DVS, explained in the section below.
4. On the desired service profile, create two iSCSI vNICs, **iscsi** and **iscsi0**, which are required to access the NetApp storage iSCSI target during iSCSI bootup to load the VMware ESX operating system over the iSCSI network. Make sure that the iSCSI vNIC **iscsi** is overlaid on static vNIC **vNIC2**, and that **iscsi0** is overlaid on static vNIC **vNIC3**, as shown in Figure 18.

Figure 18. iSCSI vNICs Overlaid on Static vNICs

For the Cisco UCS M81KR VIC, make sure that the MAC address is marked “Derived” and that the correct VLAN ID (192) is chosen to access the NetApp iSCSI target during VMware ESX iSCSI bootup.

5. In Cisco UCS Manager, create a new iSCSI boot policy, **MSSQL-iSCSI-Boot**, with iSCSI vNIC **iscsi** as a primary path and **iscsi0** as a secondary path to provide redundancy during VMware ESX host iSCSI bootup in case of software or hardware faults. Figure 19 shows the boot policy configuration.

Figure 19. New iSCSI Boot Policy in Cisco UCS Manager

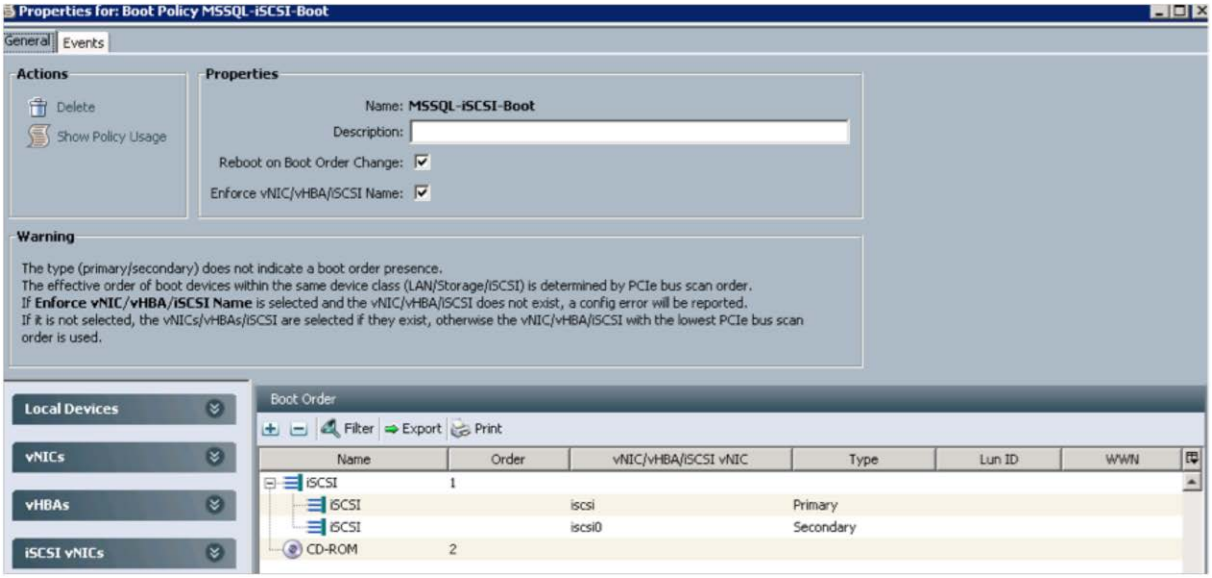

6. After the iSCSI Boot policy is created, choose a newly created boot order policy for the desired service profile. For the chosen service profile on the Cisco UCS Manager Boot Order tab, assign **iscsi** as the primary iSCSI vNIC and **iscsi0** as the secondary iSCSI vNIC with VMware ESX iSCSI boot parameters as shown in Table 8 and Figure 20.

Table 8. iSCSI Boot Parameters

iSCSI vNIC Name	iSCSI Initiator iSCSI Qualified Name (IQN)	Initiator IP Address Policy	Initiator IP Address	iSCSI Target IQN	iSCSI Port	iSCSI Target IP Address	LUN ID
iscsi	iqn.2012-01.com.vmware.ESX5i	Static	192.191.1.5	iqn.1992-08.com.netapp.sn:1574126331	3260	192.191.1.2	0
iscsi0	iqn.2012-01.com.vmware.ESX5i	Static	192.191.1.6	iqn.1992-08.com.netapp.sn:1574126331	3260	192.191.1.2	0

Figure 20. Setting iSCSI Boot Parameters

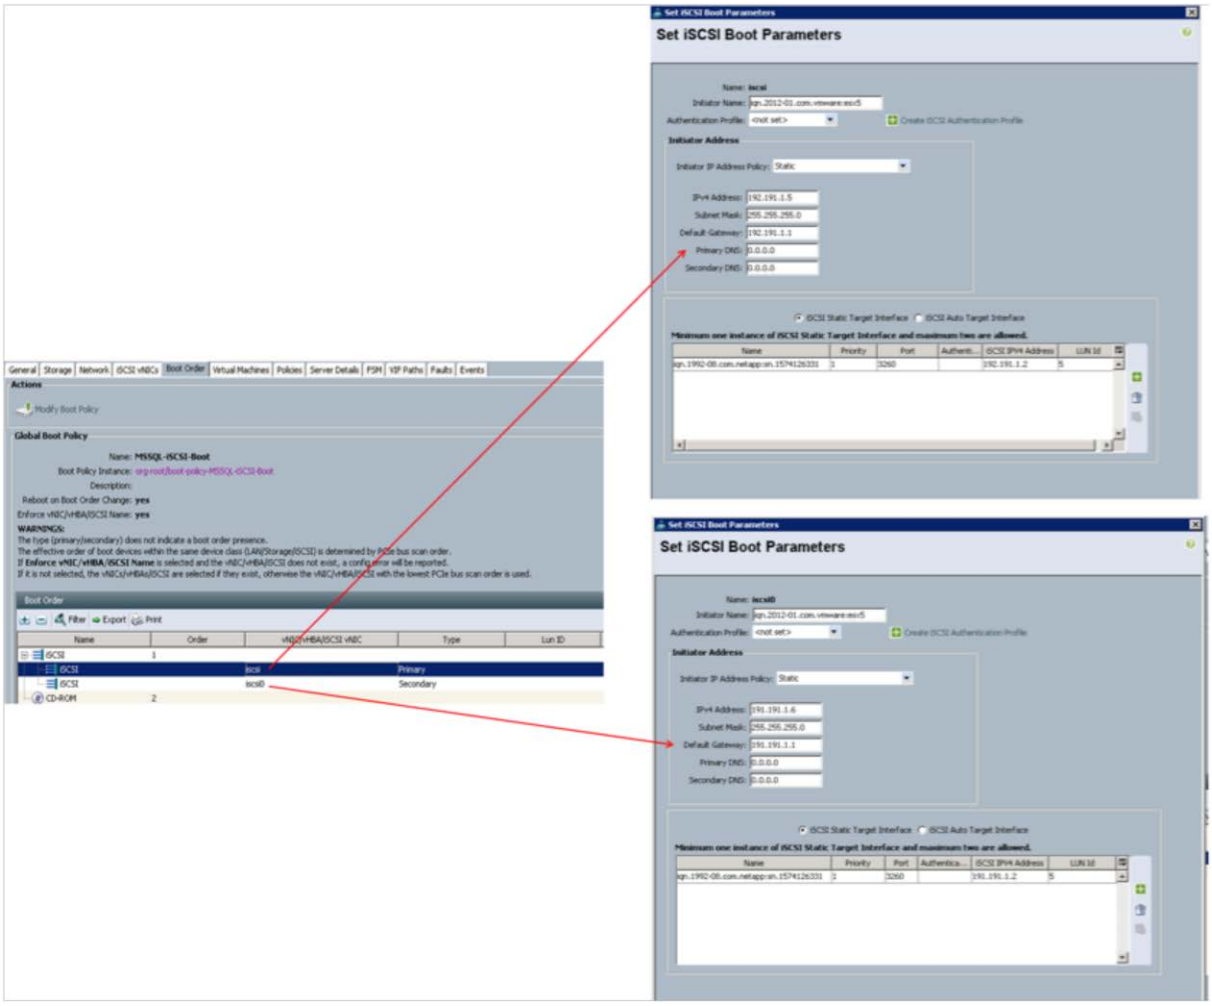

7. Associate the service profile with the desired blade (Cisco UCS B230 in this case). On Cisco UCS in the associated service profile, launch the keyboard, video, and mouse (KVM) console. Through the virtual media interface, map the VMware ESX 5.0 ISO image and install the operating system on the iSCSI boot LUN exposed over the iSCSI network.

   For more information about installing the OS in the iSCSI boot LUN, see http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html.

8. After the completion of the VMware ESX 5.0 iSCSI OS boot installation and VMware ESX 5.0 iSCSI bootup, on the VMware ESX console press the F2 key to configure the management network. Under the Network Adapters option, select static vNICs **vNIC1** and **vNIC2** as uplinks for the default VMware ESX vSwitch named **iSCSIBoot vSwitch**, as shown in Figure 21.

Figure 21. Configuring Network Management

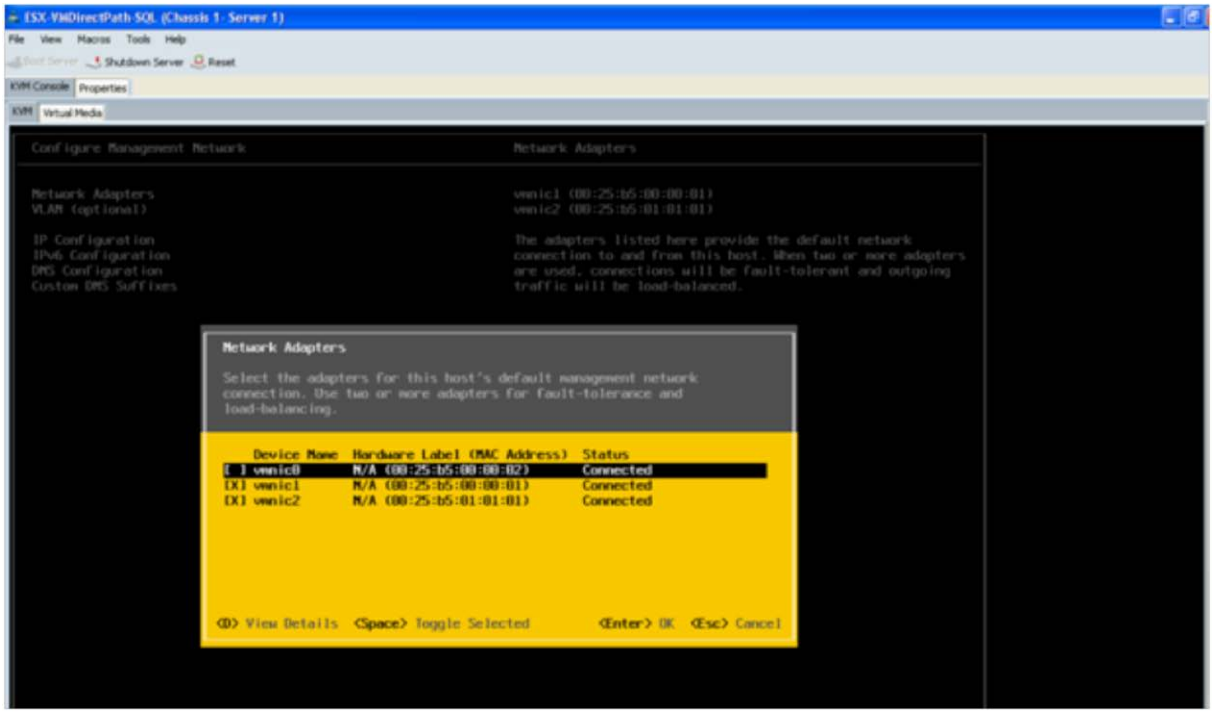

9. Under the IP Configuration option, configure the management IP address on VLAN 108 on the VMkernel management port group. Note that by default the IP address is set to the iSCSI vNIC IP address (VLAN ID 192). Figure 22 shows the management IP address configuration details.

Figure 22. Management IP Configuration Details

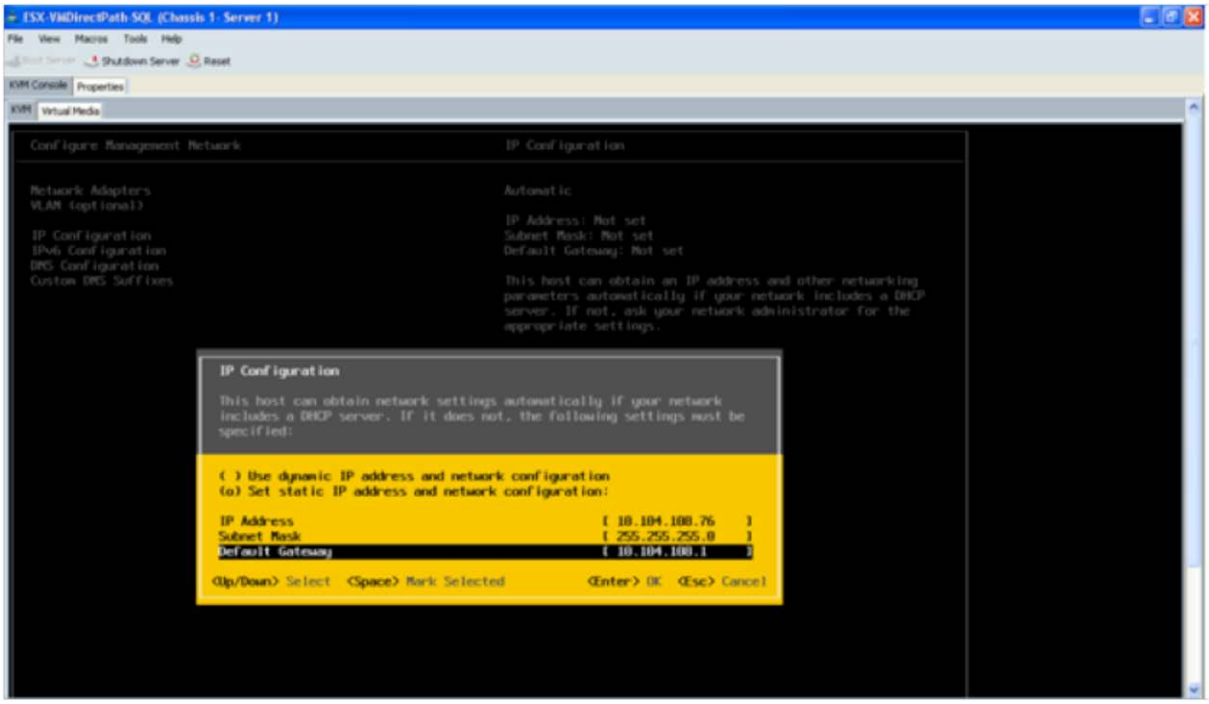

With these steps, VMware ESX 5.0 installation is completed with the iSCSI boot LUN configured on the NetApp FAS3270.

Microsoft SQL Deployment Overview

This section describes various iSCSI network topologies available to deploy a Microsoft SQL Server 2012 single-host database installed on a Microsoft Windows 2008 R2 guest virtual machine on the VMware ESX 5.0 hypervisor on a Cisco UCS B-Series server connected to the NetApp iSCSI storage over an iSCSI network as described in the section Cisco UCS and Storage iSCSI Network.

This section also discusses three scenarios for accessing storage through the iSCSI protocol for Microsoft SQL Server 2012 on a Microsoft Windows 2008 R2 guest virtual machine hosted by the VMware ESX 5.0 hypervisor:

- Guest-based iSCSI initiator on Cisco Data Center VM-FEX DVS
- VMware ESX host-based iSCSI initiator on VMware ESX vSwitch
- Guest-based iSCSI initiator on VMware ESX vSwitch

Guest-Based iSCSI initiator on Cisco Data Center VM-FEX DVS

This section describes a Microsoft SQL Server single-host deployment in a VMware ESX environment using a Cisco Data Center VM-FEX DVS for accessing shared NetApp storage over the iSCSI protocol.

Cisco Data Center VM-FEX is a software DVS that can be used in the VMware ESX environment to provide better visibility and manageability and allow hypervisor VMware VMDirectPath I/O, which provides wire-speed 10-Gbps capability to the guest virtual machine while running I/O-intensive applications.

Cisco Data Center VM-FEX significantly reduces the number of network management points, enabling both physical and virtual network traffic to be treated in a consistent policy-based way.

The Cisco Data Center VM-FEX software extends the Cisco fabric extender technology to the virtual machine with the following capabilities:

- Each virtual machine has a dedicated interface on the parent switch.
- All virtual machine traffic is sent directly to the dedicated interface on the switch.
- The software-based switch in the hypervisor is bypassed.

The following section provides a high-level overview of iSCSI network infrastructure configuration for a Microsoft SQL Server 2012 single-host installation on Cisco UCS, the VMware ESX 5.0 hypervisor, and a Microsoft Windows 2008 R2 guest virtual machine.

Cisco Data Center VM-FEX in Cisco UCS is integrated with VMware ESX 5.0 through the VMware vCenter plug-in. It is assumed that the Cisco Data Center VM-FEX plug-in is integrated with VMware vCenter. For more information, see http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide.pdf.

Figure 23 shows the physical and logical architecture of the Cisco UCS Manager virtual machine–specific configuration and VMware ESX and Cisco VM-FEX DVS configuration to deploy Microsoft SQL Server 2012 single-host database on the Microsoft Windows 2008 R2 guest virtual machine–based iSCSI software initiator.

Figure 23. Physical and Logical Architecture of Cisco UCS Manager

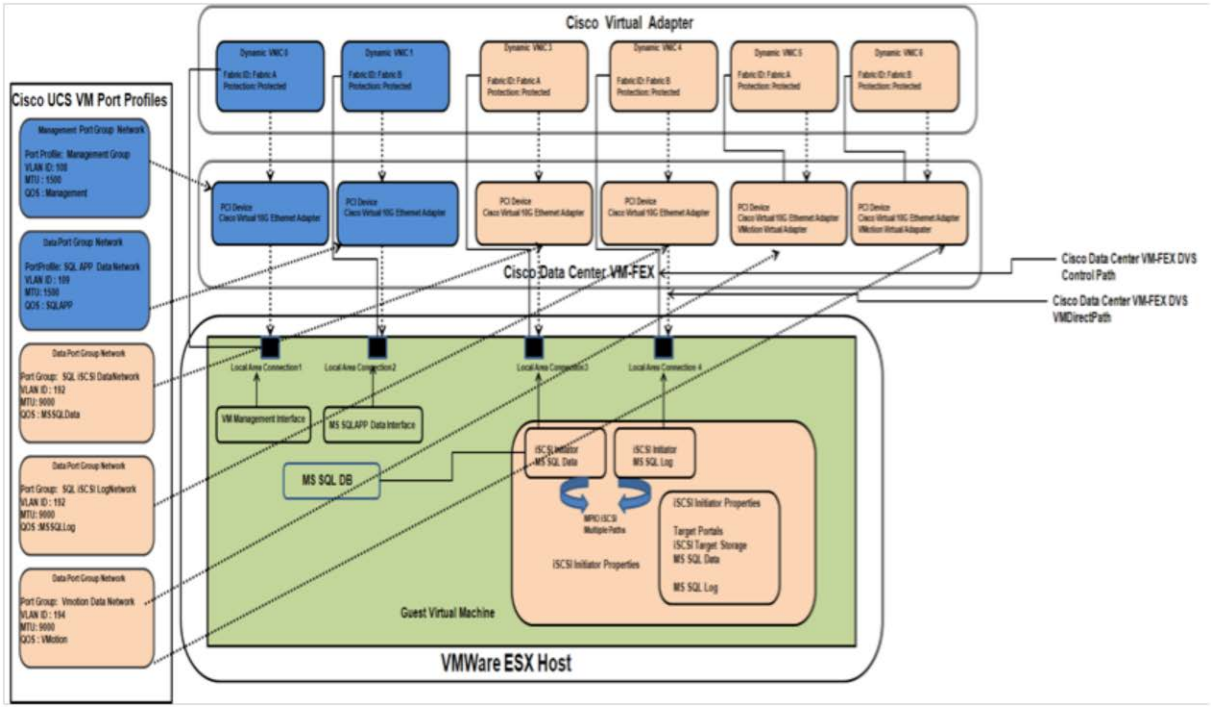

Cisco UCS Manager Virtual Machine Port Profile Design

This section describes Cisco Data Center VM-FEX port profile design considerations in Cisco UCS required to deploy a Microsoft SQL Server 2012 single-host network layout on the Microsoft Windows 2008 R2 guest virtual machine running the native iSCSI software initiator to access the NetApp iSCSI target to store database and log files.

The Cisco UCS Manager port profile for Cisco Data Center VM-FEX provides network properties and settings (VLAN ID, QoS, VMware VMDirectPath, and so on) to apply on the Cisco UCS dynamic vNIC VIFs that are exposed to the VMware ESX hypervisor through the VMware vCenter server. These dynamic vNICs are attached to the guest virtual machine (Microsoft Windows 2008 R2) running the Microsoft SQL Server 2012 single-host database to access the NetApp iSCSI storage and the database and log files for operations.

The following steps describe the Cisco Data Center VM-FEX port profile design process in Cisco UCS Manager on the VM tab:

1. To manage and configure the VMware ESX host and guest virtual machines for administration purposes, define the Cisco UCS virtual machine port profile **Managementgroup** with VLAN ID 108, a 64-port maximum, and a QoS policy with the Best Effort class for management traffic on dynamic vNICs assigned to the guest virtual machine.
2. Define the port profile **SQLAPPDataNetwork** for dynamic vNICs through which internal or external clients can access the Microsoft SQL Server database hosted on the guest virtual machine. This port profile is configured with VLAN ID 109, a 64-port maximum, and a QoS policy of **SQLAPP** with the Bronze class. Also, the **VMDirectPath High Performance** option is enabled on these dynamic vNICs assigned to the guest virtual machine.
3. The Microsoft SQL Server database log file is accessed by the iSCSI software initiator running in the guest virtual machine. To provide traffic isolation for better security and better bandwidth, define the port profile **SQLiSCSILogNetwork** with VLAN ID 192 and a QoS policy of **MSSQLData** with the Platinum class. The **VMDirectPath High Performance** option is enabled for the dynamic vNIC assigned to the guest virtual machine for accessing log LUNs.
4. The Microsoft SQL Server database data file is accessed by the iSCSI software initiator running in the guest virtual machine. To provide traffic isolation for security and better bandwidth, define the Cisco UCS virtual machine port profile **SQLiSCSIDataNetwork** with VLAN ID 192 and a QoS policy of **MSSQLLog** with the Gold class. The option **VMDirectPath High Performance** is enabled for the dynamic vNIC assigned to the guest virtual machine for accessing database LUNs.
5. To handle VMware ESX vMotion traffic for performing guest virtual machine migration, for a failure scenario or for better load balancing of hardware resources, you must use secured and dedicated network bandwidth. To achieve this, define the Cisco UCS virtual machine port profile **VMotion** with VLAN ID 194 and a QoS policy of **VMotion** with the Silver class, which will be assigned to the VMware ESX host VMkernel network Interfaces.

Table 9 provides the Cisco Data Center VM-FEX port profile Cisco UCS design VLAN ID, QoS policy, maximum port count, and high-performance configuration settings for VMware VMDirectPath I/O.

Table 9. Cisco Data Center VM-FEX Port Profile Properties in Cisco UCS

Cisco UCS: Cisco Data Center VM-FEX Port Profile	Port-Profile Properties
Managmentgroup	QoS policy: Management Network control policy: Default Maximum ports: 64 VLAN ID: 108
SQLAPPDataNetwork	QoS policy: SQLAPP Network control policy: Default Maximum ports: 64 Host network I/O performance: High Performance VLAN ID: 109
SQLiSCSIDataNetwork	QoS policy: MSSQLData Network control policy: Default Maximum ports: 64 Host network I/O performance: High Performance VLAN ID: 192
SQLiSCSILogNetwork	QoS policy: MSSQLLog Network control policy: Default Maximum ports: 64 Host network I/O performance: High Performance VLAN ID: 192
VMotion	QoS policy: VMotion Network control policy: Default Maximum ports: 64 VLAN ID: 194

Figure 24 verifies the QoS policies mapping to newly created Cisco Data Center VM-FEX port profiles on the Cisco UCS Manager VM tab.

Figure 24. QoS Policy Mapping with Cisco Data Center VM-FEX Port Profiles

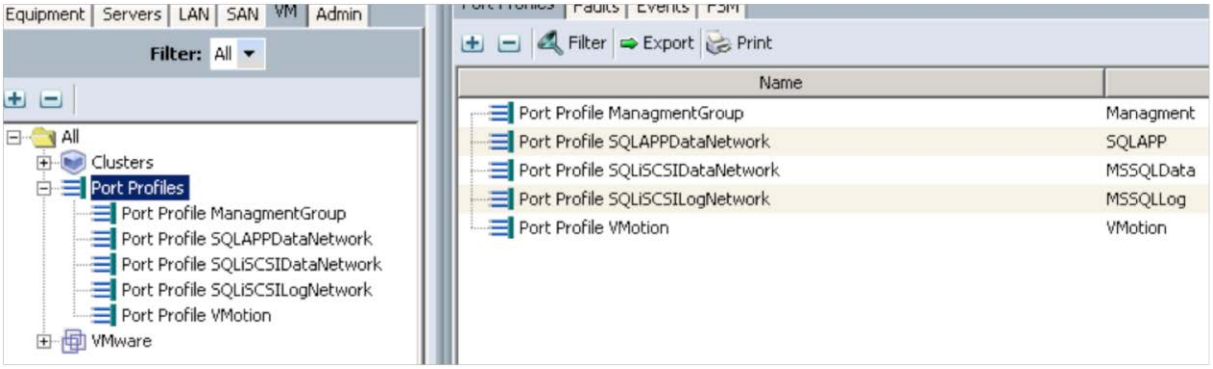

Details of the properties assigned to each of the newly created port profiles can be verified by selecting Port Profiles on the Cisco UCS Manger VM tab, as shown in Figure 25.

Figure 25. Port Profile Properties in Cisco UCS Manager

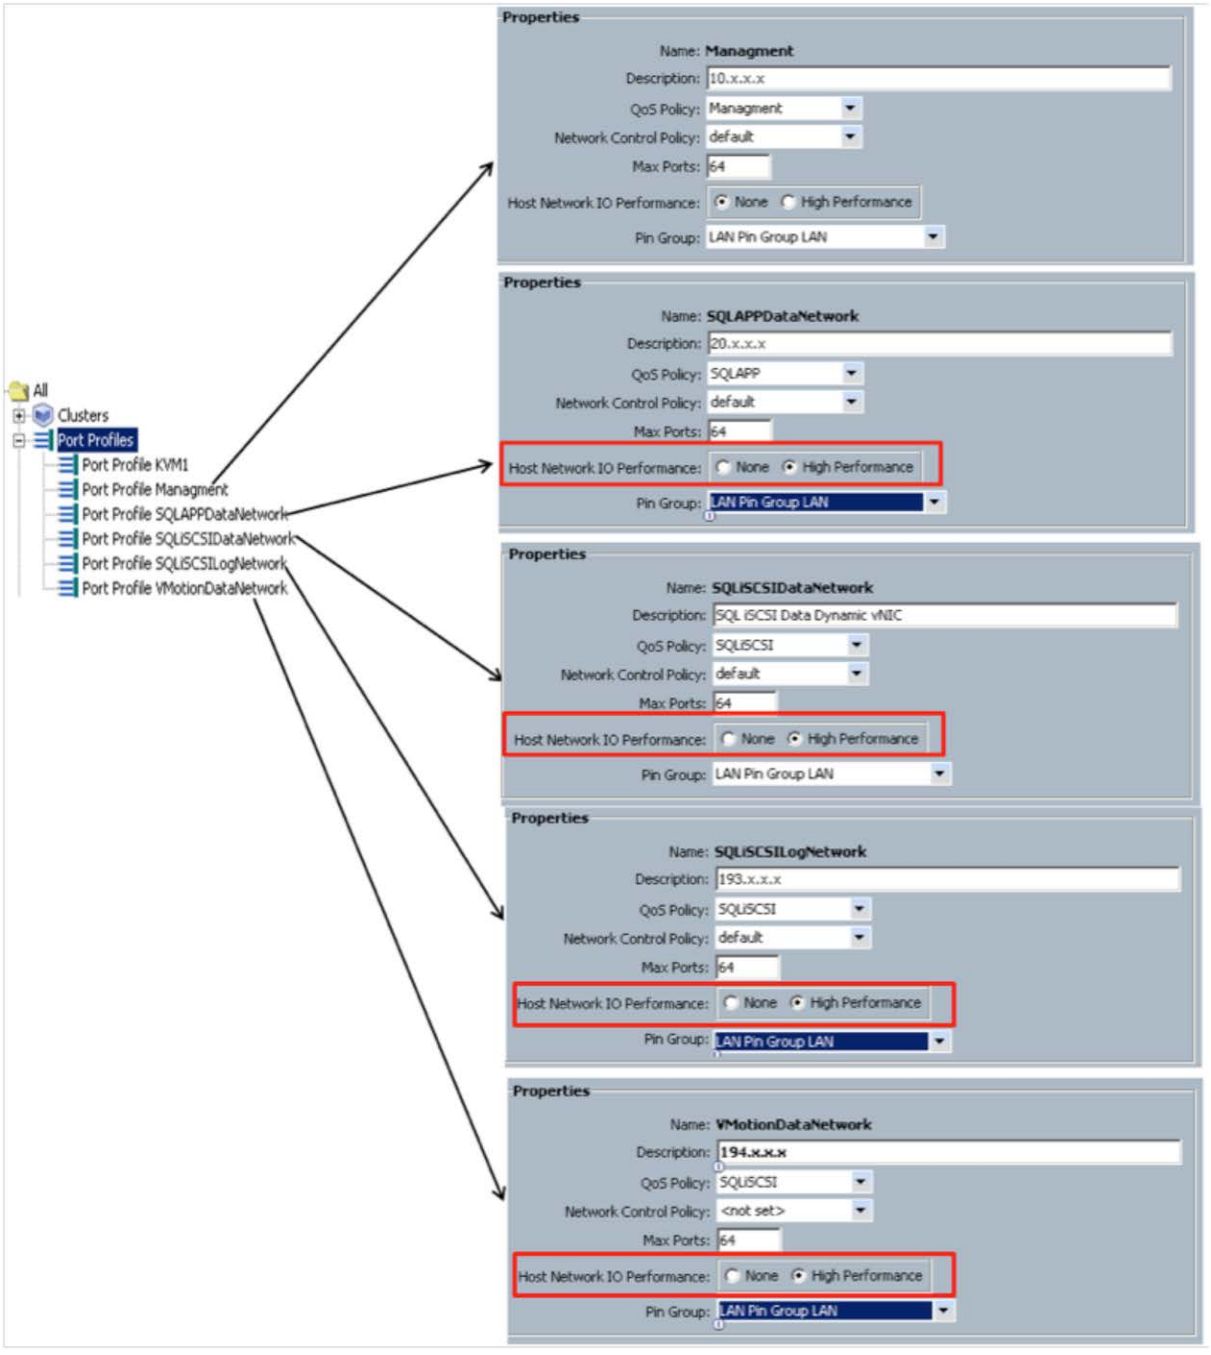

Cisco UCS Service Profile Design

This section explains static and dynamic vNIC network design with the Cisco UCS service profile to deploy a Microsoft SQL Server single-host database on a Microsoft Windows 2008 R2 guest virtual machine with VMware ESX 5.0 using a Cisco Data Center VM-FEX DVS. A goal of this design is to achieve high I/O throughput and high availability.

The following procedure shows the configuration steps to be performed for each service profile to create vNICs to access the iSCSI storage target.

6. In the service profile, create two static vNICs, **iscsi** and **isci0**, which are overlaid on two static vNICs, **vNIC2** on Fabric A and **vNIC3** on Fabric B, respectively, as explained in the VMware ESX iSCSI Boot section.
7. The service profile also has two static vNICs, **vNIC0** on Fabric A and **vNIC1** on Fabric B, with VLAN ID 108, without fabric failover, and a **Management** QoS policy definition to handle VMware ESX host and guest virtual machine management data network traffic.
8. Configure the service profile with dynamic vNIC connection policy with a predefined number of vNICs, which are exposed to the VMware ESX host to connect management VMware VMkernel network adapters (vmnic0 and vmnic1 are part of the VMware ESX vSwitch iSCSI boot). However, these vNICs will be migrated later to the Cisco Data Center VM-FEX DVS.

You need at least six dynamic vNICs for the current design. To derive the number of dynamic vNICs, see http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/vm_fex_best_practices_deployment_guide_ps10277_Products_White_Paper.html.

Figure 26 shows the configuration details for static and dynamic vNICs created for the specific service profiles.

Figure 26. Configuration Details for Static and Dynamic vNICs

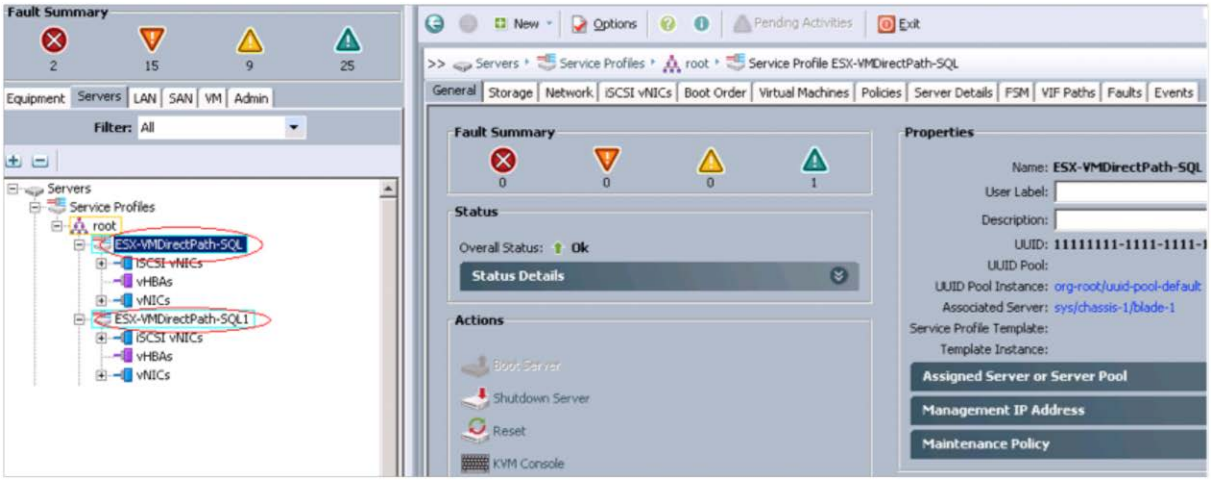

Table 10 shows the properties of static and dynamic vNICs created for the service profile for deploying Microsoft SQL Server 2012 on the Microsoft Windows 2008 R2 guest virtual machine on VMware ESX 5.0 in a Cisco Data Center VM-FEX environment.

Table 10. Cisco UCS Service Profile Network Design

vNIC Name	vNIC Type	Fabric ID	Failover	Adapter Policy	VLAN	MAC Address	QoS
vnic0	Static	Fabric A	No	VMware	108	00:25:B5:00:00:01	Management
vnic1	Static	Fabric B	No	VMware	108	00:25:B5:00:00:02	Management
vnic2	Static	Fabric A	No	VMware	192	00:25:B5:01:01:01	MSSQLLOG
vnic3	Static	Fabric B	No	VMware	192	00:25:B5:01:01:02	MSSQLLOG
PCI device	Dynamic	Fabric A	Yes	VMware PassThrough	108 (Cisco UCS virtual machine port profile Managementgroup)	Derived (Cisco UCS virtual machine port profile Managementgroup)	Management (Cisco UCS virtual machine port profile Managementgroup)
PCI device	Dynamic	Fabric B	Yes	VMware PassThrough	109 (Cisco UCS virtual machine port profile SQLAPPDataNetwork)	Derived (Cisco UCS virtual machine port profile SQLAPPDataNetwork)	MSQLAPP (Cisco UCS virtual machine port profile SQLAPPDataNetwork)
v PCI device	Dynamic	Fabric A	Yes	VMware PassThrough	192 (Cisco UCS virtual machine port profile SQLiSCSIDataNetwork)	Derived (Cisco UCS virtual machine port profile SQLiSCSIDataNetwork)	MSSQLData (Cisco UCS virtual machine port profile SQLiSCSIDataNetwork
PCI device	Dynamic	Fabric B	Yes	VMware PassThrough	192 (Cisco UCS virtual machine port profile SQLiSCSILogNetwork)	Derived (Cisco UCS virtual machine port profile SQLiSCSILogNetwork)	MSSQLLOG (Cisco UCS virtual machine port profile SQLiSCSILogNetwork)

VMware ESX Host and Guest Virtual Machine Network Design

This section discusses the network design layout for a VMware ESX host and Microsoft Windows 2008 R2 guest virtual machine with a standalone Microsoft SQL Server 2012 installed with an iSCSI software initiator connected to NetApp shared storage access.

This section describes the VMware ESX and Microsoft Windows 2008 R2 guest virtual machine physical and logical iSCSI network design to deploy the Microsoft SQL Server 2012 database and log file.

- When VMware ESX 5.0 is booted through the iSCSI LUN on a Cisco UCS B230 blade server, VMware ESX 5.0 host VMNIC network adapters are mapped with Cisco UCS static vNICs on the VMware vCenter server as shown in Figure 27.

Figure 27. Mapping of Network Adapters with Cisco UCS Static vNICs in VMware vCenter

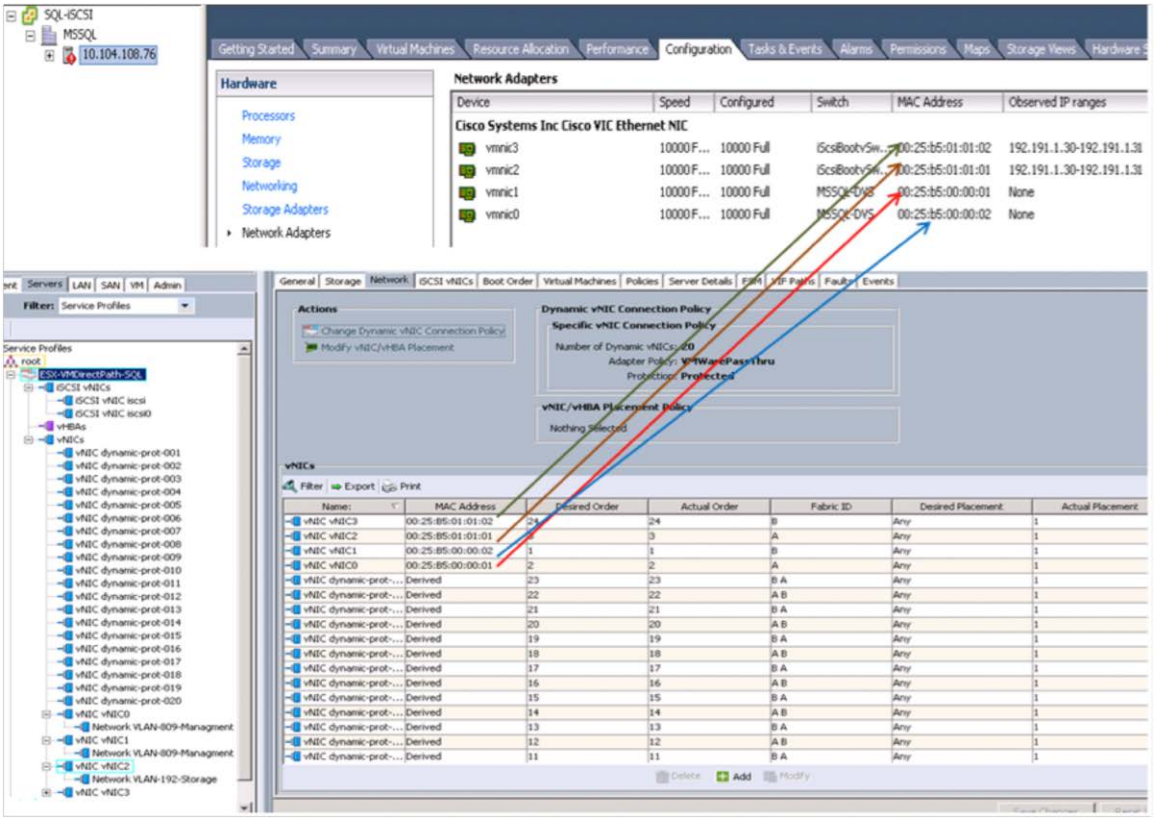

- ● The VMware VMkernel (vmk0) management port and its associated physical VMNIC adapters, **vmnic0** and **vmnic1**, with uplinks on the default **Management Network1** port group on **vSwitch0** defined during installation of VMware ESX 5.0 need to be migrated to the Cisco Data Center VM-FEX DVS. For more information, see Cisco Data Center VM-FEX Administration guide: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide.pdf. Figure 28 shows VMware ESX vSwitch configuration after the migration is complete.

The two uplink ports, **vmnic4** and **vminc2**, of the **MS SQL Server iSCSI Boot** port group of **vSwitch4** should be left undisturbed. Altering these settings can affect VMware ESX bootup through the iSCSI LUNs.

Figure 28. VMware ESX 5.0 vSwitch Configuration Details

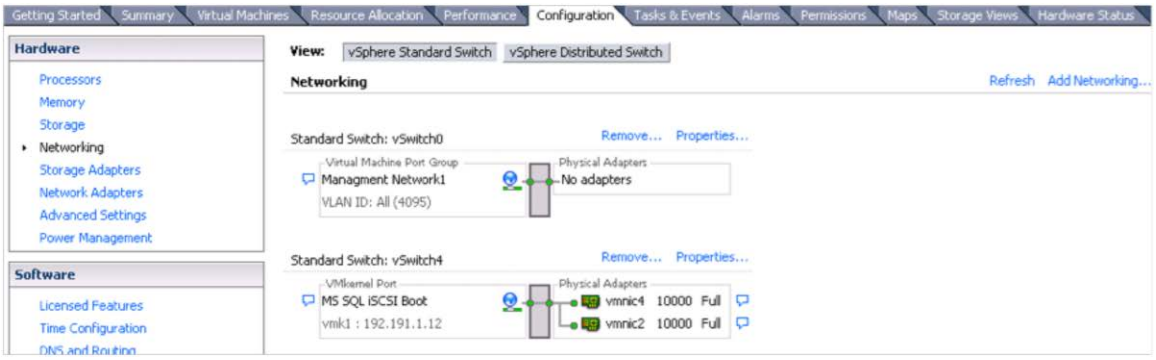

Perform the following steps to deploy a standalone Microsoft SQL Server 2012:

1. On the Microsoft Windows 2008 R2 guest virtual machine; create four virtual adapters to access management, the Microsoft SQL Server application, and the iSCSI storage network.
2. Attach Virtual Network Adapter 1 to the Cisco Data Center VM-FEX port profile **ManagementGroup**, which is defined to access the guest virtual machine management network.
3. Attach Virtual Network Adapter 2 to the Cisco Data Center VM-FEX port profile **SQLiSCSIDataNetwork**, which is defined to access the Microsoft SQL Server database, and attach Virtual Network Adapter 3 to the Cisco Data Center VM-FEX port profile **SQLiSCSILogNetwork** to access Microsoft SQL Server database log files on shared NetApp storage over the iSCSI network.
4. Attach Virtual Network Adapter 4 to the Cisco Data Center VM-FEX port profile **SQLAppDataNetwork**, which is defined to access the Microsoft SQL Server database from the client network.
5. Enable VMware VMDirectPath I/O on Virtual Network Adapters 2, 3, and 4 on the Cisco Data Center VM-FEX port profiles to bypass the VMware ESX kernel stack to achieve high performance and reduce CPU cycles for handling I/O operations on the VMware ESX host.

Figure 29 shows the Microsoft Windows 2008 guest virtual machine with the network adapter settings as explained in the preceding steps.

Figure 29. Microsoft Windows 2008 Guest Virtual Machine Showing Adapter Settings

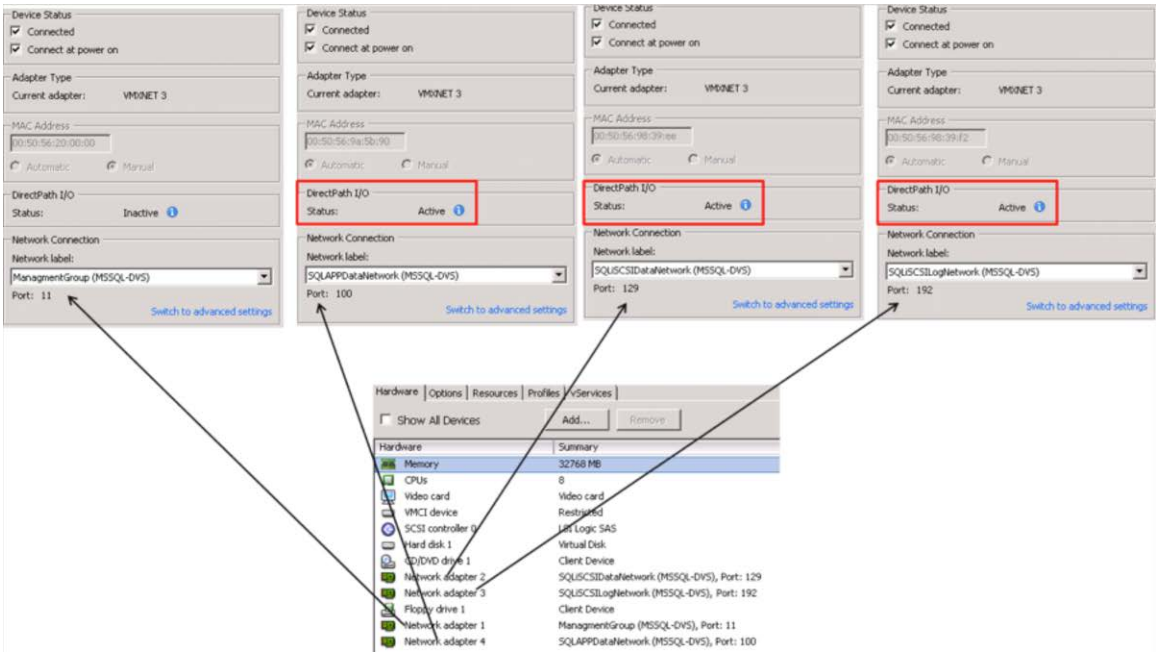

6. Configure each network interface on the Microsoft Windows 2008 R2 guest virtual machine with the required IP address to access management, the Microsoft SQL Server 2012 client network, database data, and the database log through the iSCSI network.

   Make sure that the VMXNET3 driver is selected and that memory reservation on the VM tab in VMware vCenter is performed. Set the correct IP address in the guest virtual machine as shown in Figure 30.

Figure 30. Setting the Correct IP Address in the Guest Virtual Machine

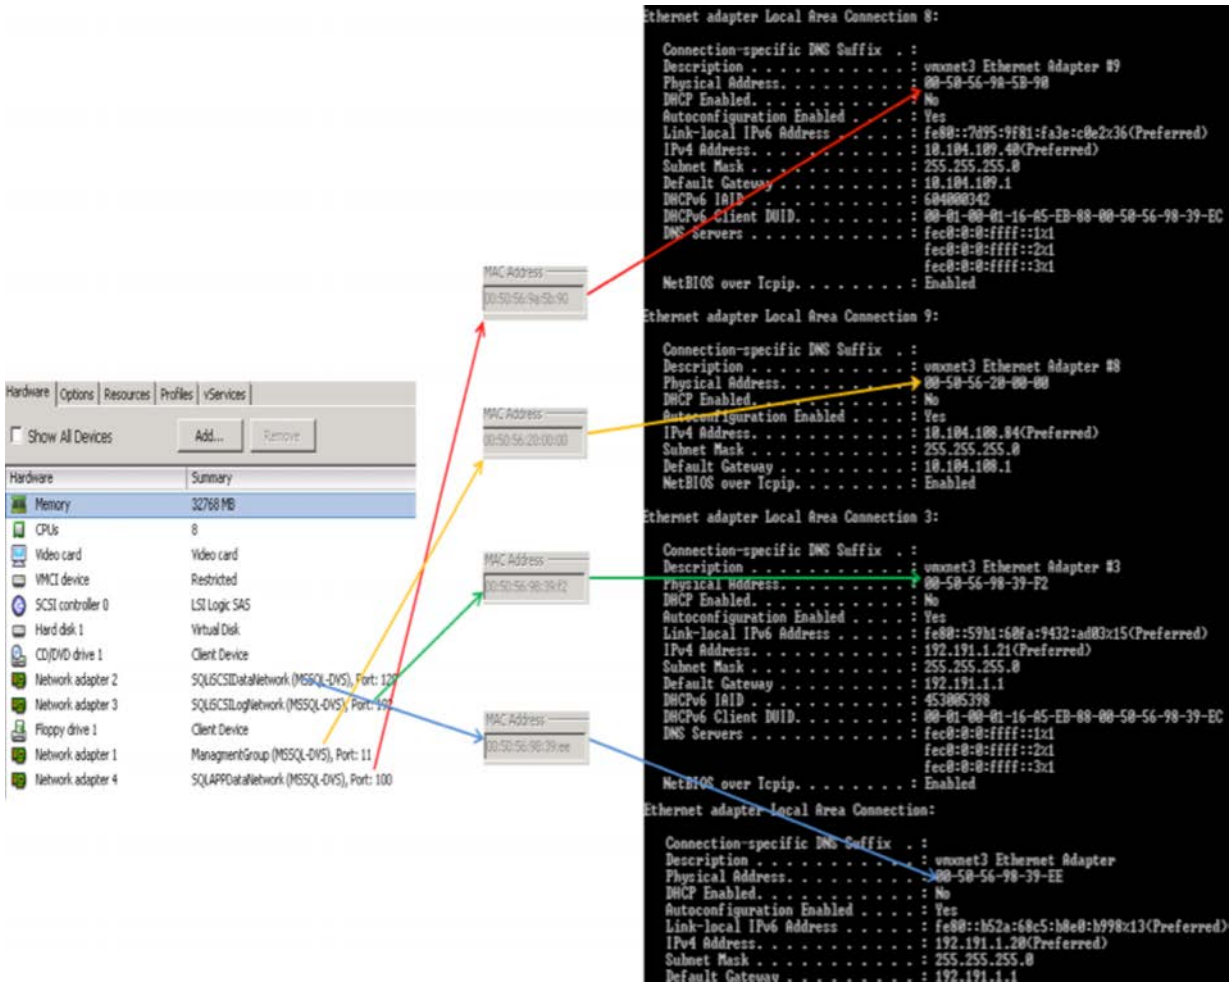

7. Perform the following steps to support end-to-end jumbo frames (MTU 9000) to carry Microsoft SQL Server client and iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine, Cisco UCS, Cisco Data Center VM-FEX, and NetApp storage:

   a. Specify MTU 9000 in the Cisco UCS QoS system class for Platinum, Gold, Silver, and Bronze classes as discussed previously.

   b. Specify MTU 9000 in the Jumbo field on the appropriate network interfaces (0, 8, and 3 in this design) in the guest virtual machine.

   c. Choose the correct QoS policy in Cisco Data Center VM-FEX port profiles **SQLAPPDataNetwork**, **SQLiSCSIDataNetwork**, and **SQLiSCSILogNetwork**.

   d. Configure the NetApp iSCSI VIF to specify MTU 9000.

8. On the Microsoft Windows 2008 R2 guest virtual machine, enable and configure the iSCSI software initiator and multipath I/O (MPIO) to access NetApp iSCSI targets.

For more information about configuring the iSCSI initiator in Microsoft Windows, see http://www.microsoft.com/download/en/details.aspx?id=18986.

The following steps present the high-level procedure for configuring the Microsoft Windows 2008 R2 guest virtual machine iSCSI software initiator for deploying a Microsoft SQL Server 2012 single-host database on a NetApp iSCSI target:

a. Discover the NetApp FAS3270HA Controller A VIF iscsiA (192.191.1.2) with Microsoft Windows 2008 R2 guest virtual machine iSCSI Initiator Adapter 3 (192.191.1.21) and NetApp FAS3270HA Controller B VIF iscsiB (192.191.1.3) with Microsoft Windows 2008 R2 guest virtual machine iSCSI Initiator Adapter 0 (192.191.1.20), as shown in Figure 31.

Figure 31. iSCSI Initiator Adapters and Target IP Addresses

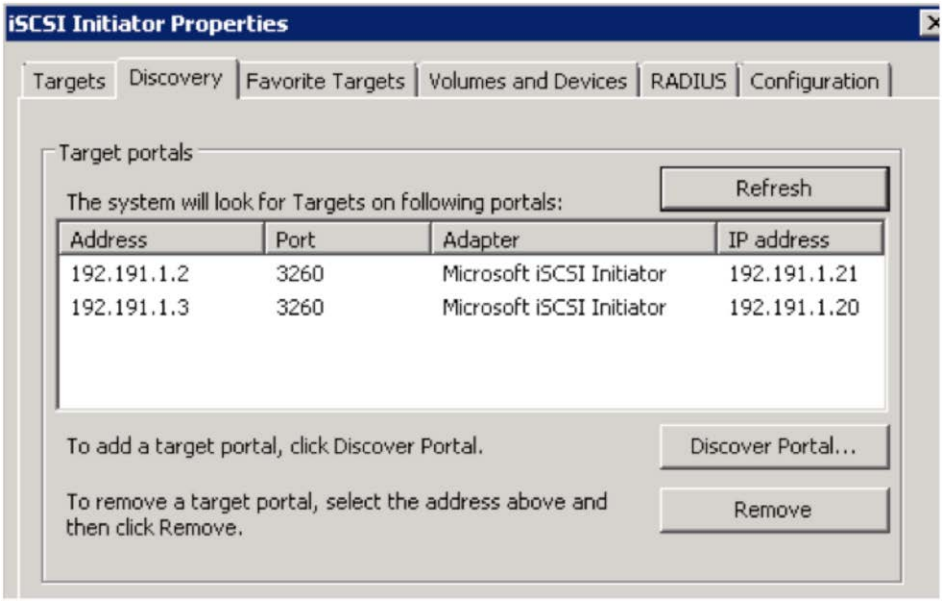

b. To enable iSCSI multipathing make sure that the Enable multi-path check box on the NetApp storage target connection is checked for both controllers. Figure 32 shows multipath enabled.

Figure 32. Multipath Enabled on NetApp Storage Target

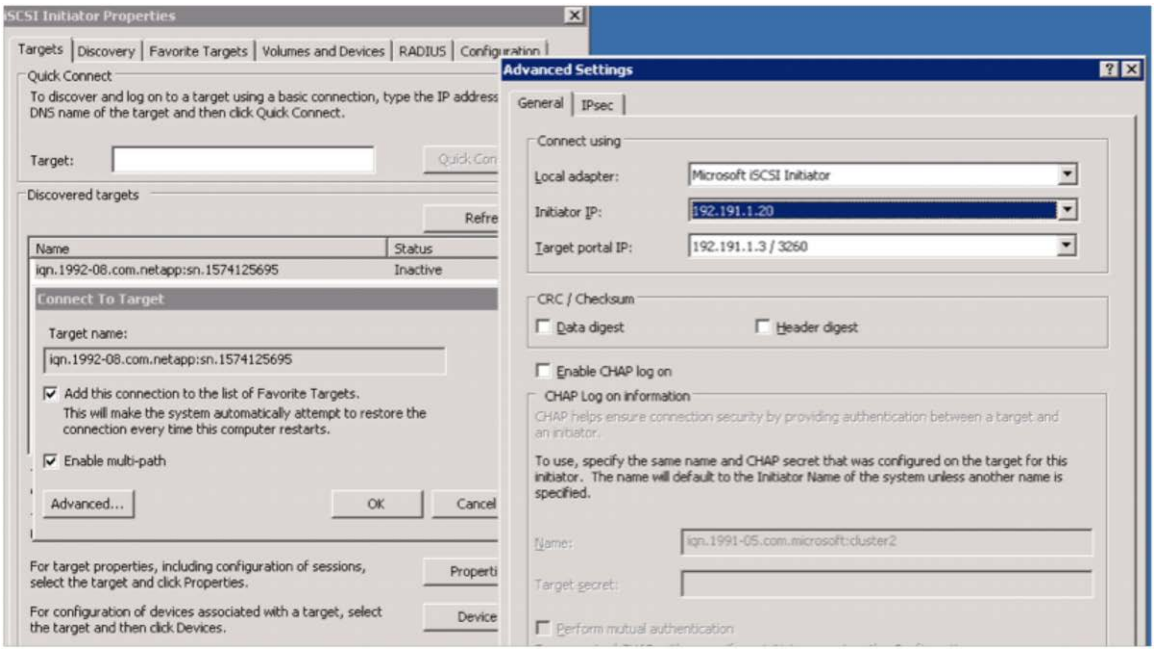

c. Log in to the iSCSI initiators on the Microsoft Windows 2008 R2 guest virtual machine. After successful login, the NetApp targets and LUNS are automatically exposed for configuration. For the Microsoft SQL Server 2012 single-host installation, you use Disk 1 as the database file and Disk 2 as the log file as shown in Figure 33. Refer to the section above for LUN creation and access information details.

Figure 33. Disk 1 and Disk 2 for Microsoft SQL Server 2012 Single-Host Installation

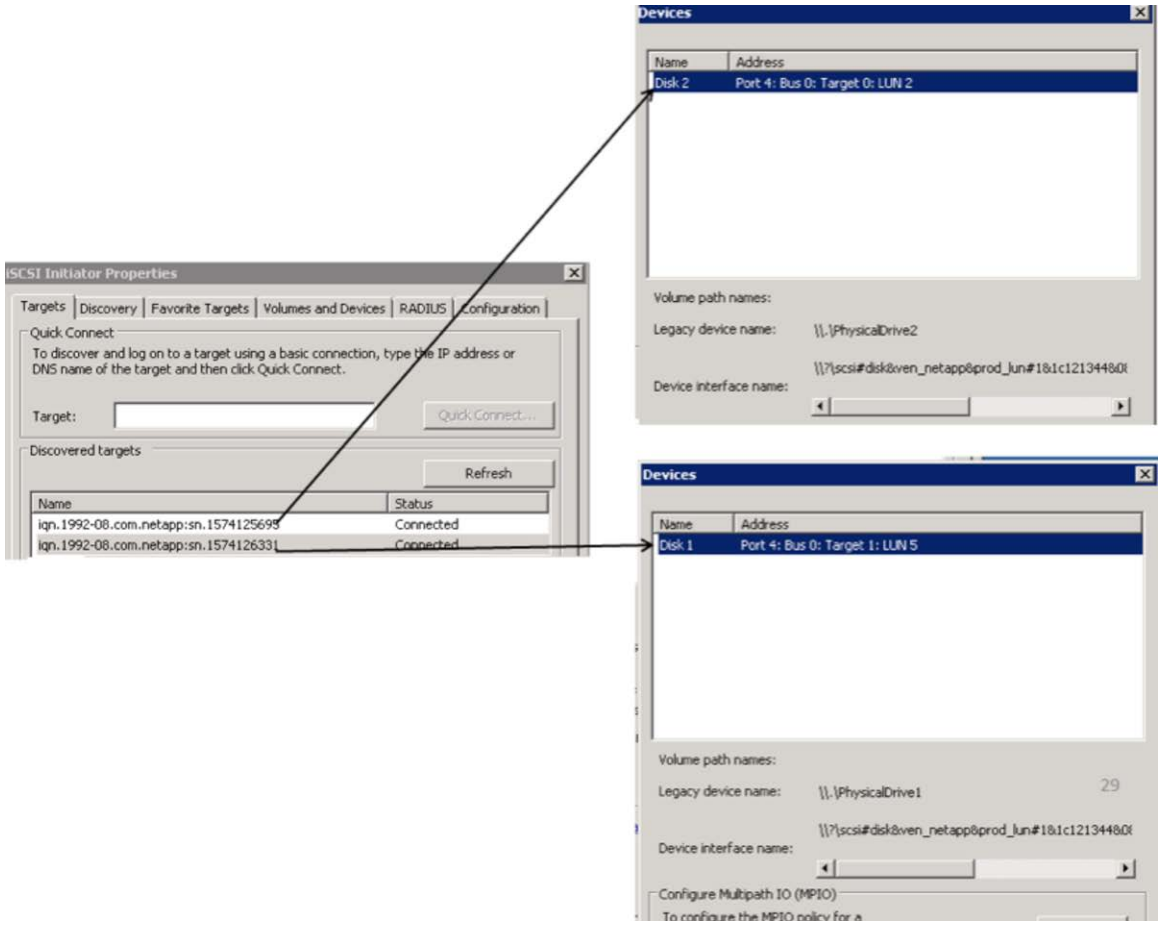

d. Under Disk Management, scan for new disks on the Microsoft Windows 2008 R2 guest virtual machine and format the disks to install the Microsoft SQL Server 2012 single-host database and log file on separate iSCSI LUNs as shown in Figure 34.

Figure 34. Scanning for New Disks Under Disk Management

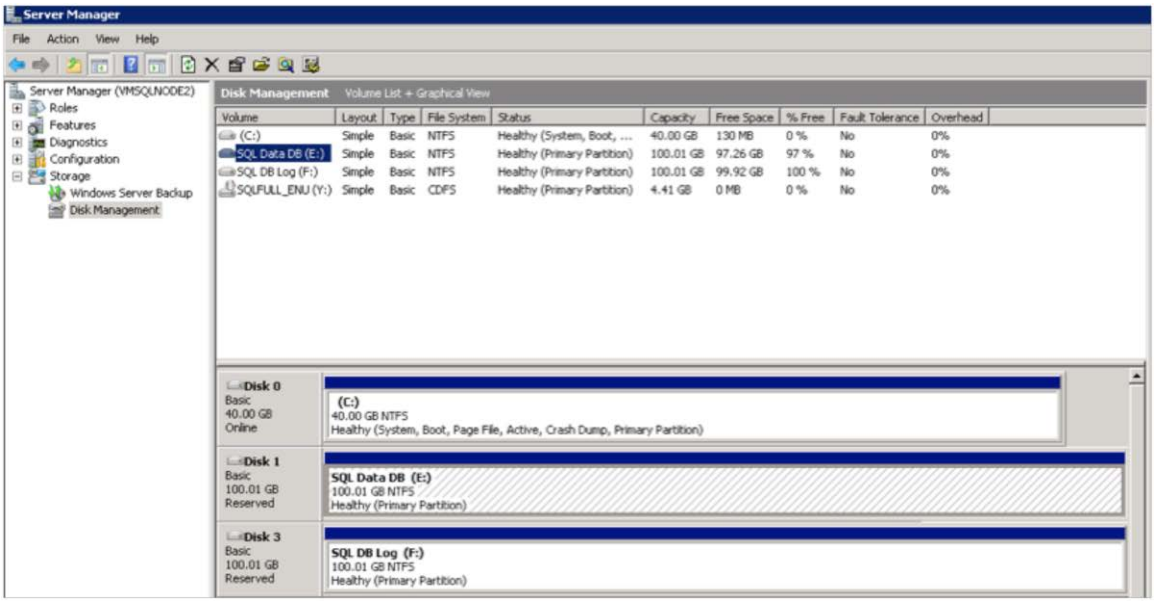

9. On the Microsoft Windows 2008 R2 guest virtual machine, install the standalone Microsoft SQL Server 2012 and using NTFS create a database labeled **SQL_DATA_DB** for data and create a log on **SQL_DB_Log**.

   For more information about installing single-host Microsoft SQL Server 2012, see http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx.

Now Microsoft SQL Server 2012 is successfully installed with database data and a log file accessing an iSCSI target LUN on NetApp FAS3270 storage.

The Cisco UCS Manager virtual machine Admin tab provides a single management interface for managing Cisco Data Center VM-FEX DVSs, configure port profiles with network settings across Cisco Data Center VM-FEX DVSs, and troubleshoot network problems on Microsoft Windows 2008 R2 virtual machines that are part of Cisco Data Center VM-FEX DVSs. Figure 35 shows all these functions in Cisco UCS Manager.

Figure 35. Management and Configuration Details of DVS Switches in Cisco UCS Manager

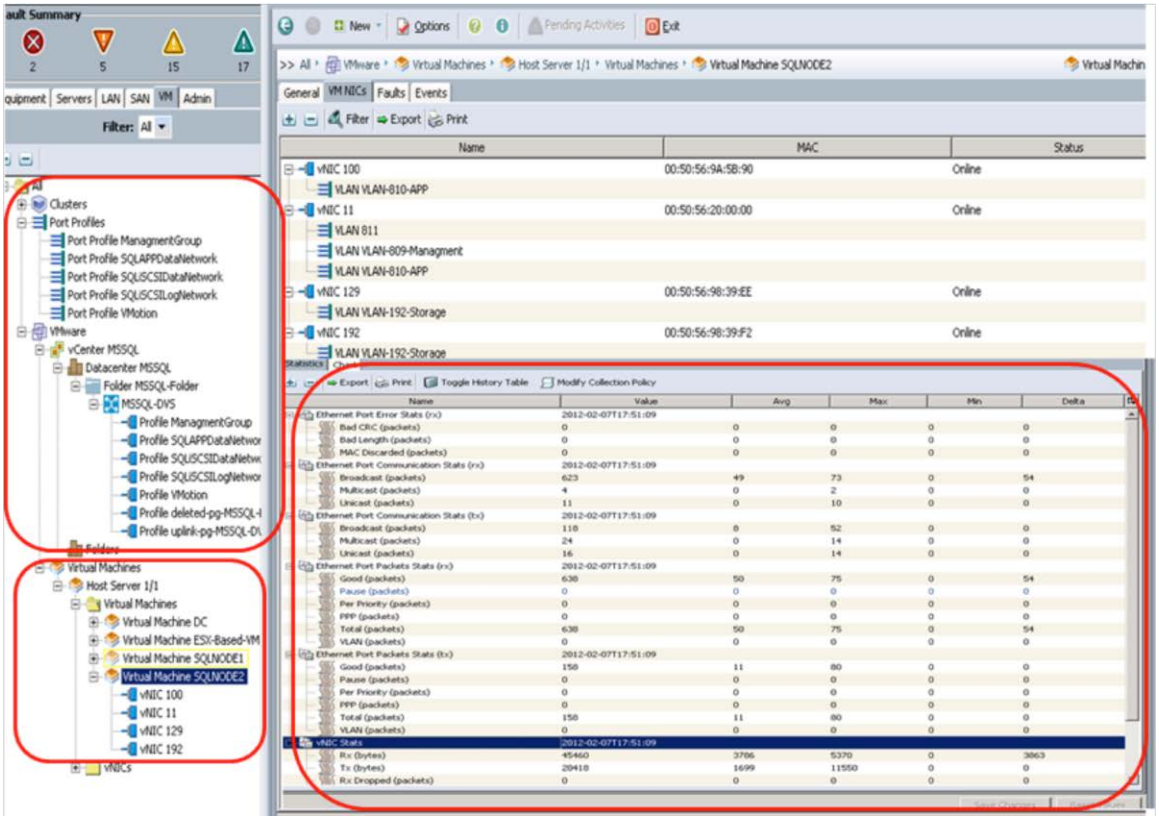

VMware ESX vCenter provides a single pane to view all the guest virtual machine dynamic vNICs that are part of the Cisco Data Center VM-FEX DVS. The pane also shows the port ID, link connectivity status, and Cisco UCS port profile information applied, with the MAC address, MTU, and IP address configured in the Microsoft Windows 2008 R2 guest virtual machine, as shown in Figure 36.

Figure 36. Link Connectivity Status and Cisco UCS Port Profile Information with MAC Address

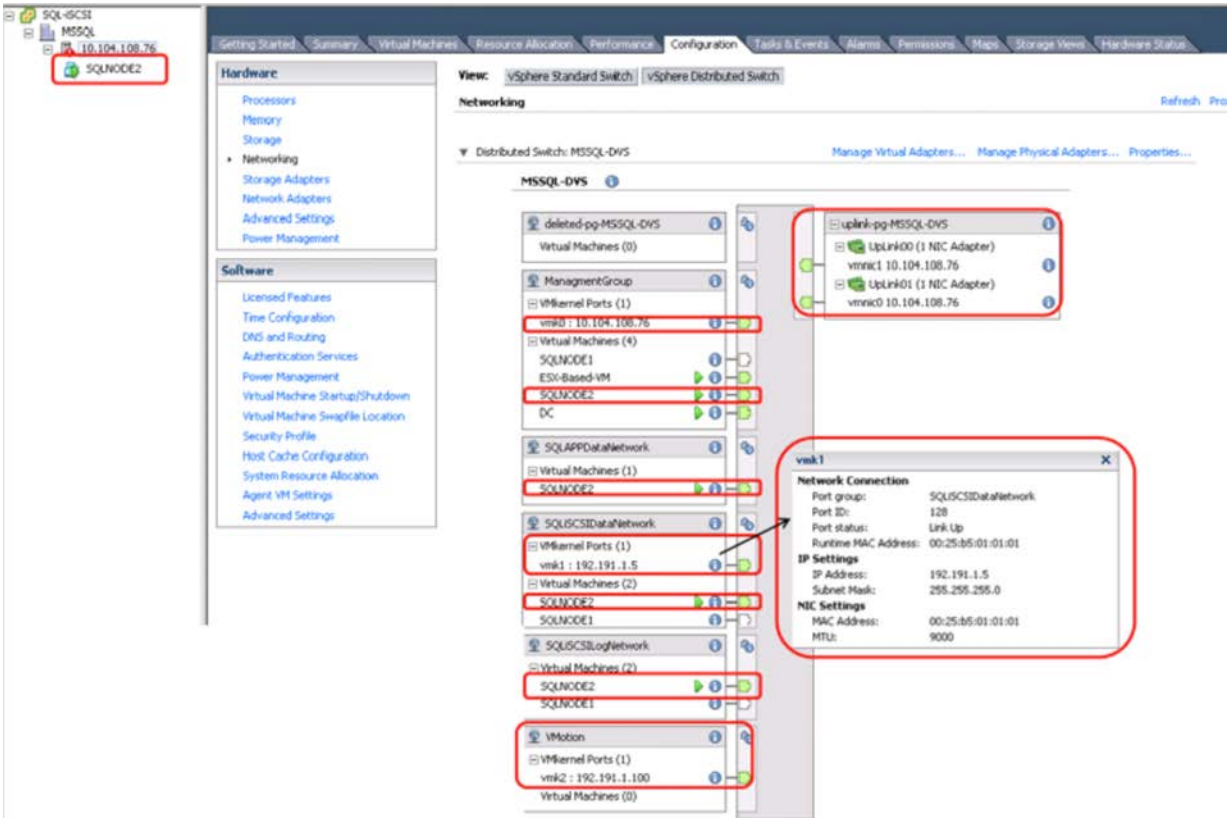

VMware ESX Host-Based iSCSI Initiator on VMware ESX vSwitch

This section provides an overview of a Microsoft Windows 2008 R2 Microsoft SQL Server 2012 single-host deployment in a VMware ESX environment using a VMware ESX host-based iSCSI software initiator to access shared NetApp iSCSI storage. In this scenario, the VMware ESX native vSwitch is configured to access VMware ESX management, Microsoft Windows 2008 R2 guest virtual machine management, Microsoft SQL Server 2012 client, VMware vMotion, and iSCSI-based storage access points.

The following sections describe the high-level design and deployment of Microsoft SQL Server 2012 single-host iSCSI network infrastructure on Cisco UCS, the VMware ESX 5.0 hypervisor, and a Microsoft Windows 2008 R2 guest virtual machine.

Figure 37 shows the logical network architecture of the Cisco UCS and VMware ESX host-based iSCSI software initiator with the VMware ESX vSwitch to deploy a Microsoft SQL Server 2012 single-host system on shared NetApp iSCSI storage.

Figure 37. VMware ESX Host-Based iSCSI Network Separated with Logical VLAN

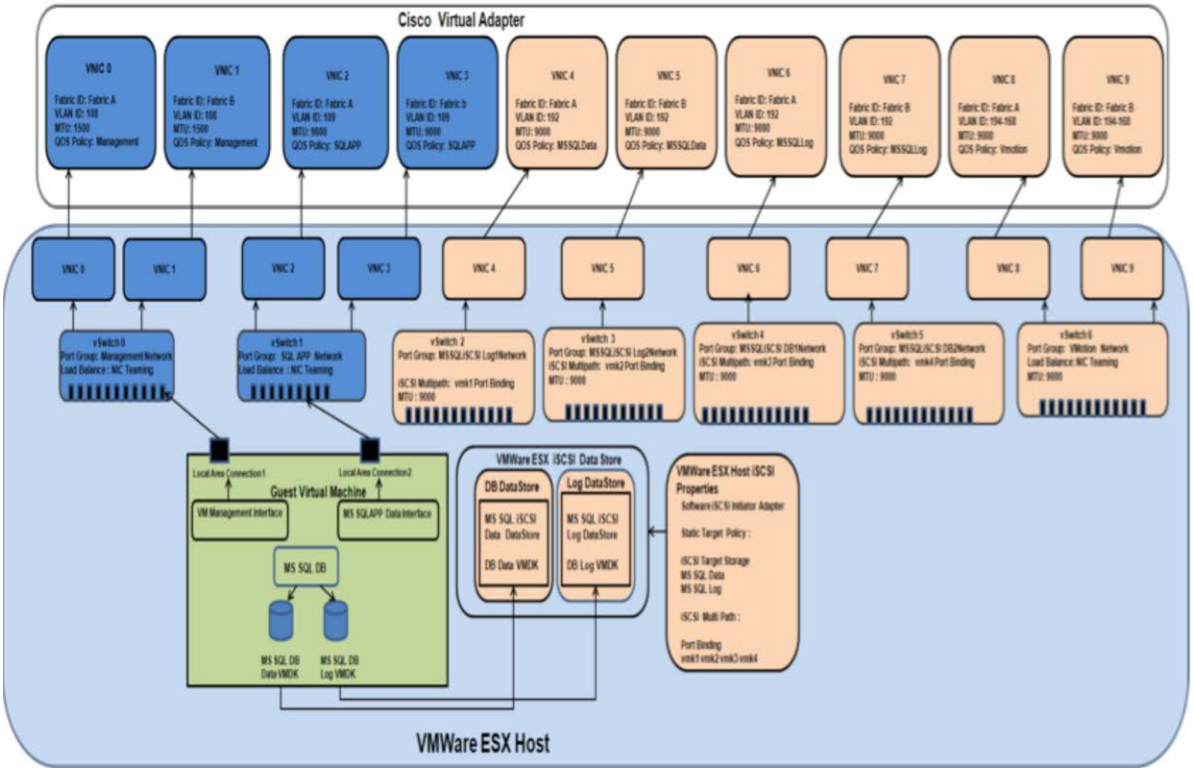

Cisco UCS Service Profile Design

This section describes how to configure static vNICs on the desired service profile to allow storage access through the iSCSI initiator at the VMware ESX host using the native vSwitch.

1. Configure iSCSI vNICs **iscsi** and **isci0**, which are overlaid on two static vNICs, **vNIC4** on Fabric A and **vNIC5** on Fabric B, respectively, as explained in the section VMware ESX iSCSI Boot.
2. Configure static vNICs **vNIC0** on Fabric A and **vNIC1** on Fabric B with VLAN ID 108, with no fabric failover, with **Management** QoS policy defined to handle VMware ESX host and Microsoft Windows 2008 R2 guest virtual machine management data network traffic.
3. Configure static vNICs **vNIC2** on Fabric A and **vNIC3** on Fabric B with VLAN ID 109, with no fabric failover, an MTU of 9000, and **SQLAPP** QoS policy defined to handle Microsoft SQL Server 2012 client data network traffic from the Microsoft Windows 2008 R2 guest virtual machine.
4. Configure static vNICs **vNIC4** on Fabric A and **vNIC5** on Fabric B with VLAN ID 192, with no fabric failover, an MTU of 9000, and **MSSQLData** QoS policy defined to handle Microsoft SQL Server 2012 database data iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine.
5. Configure static vNICs **vNIC6** on Fabric A and **vNIC7** with VLAN ID 192, with no fabric failover, an MTU of 9000, and **MSSQLLog** QoS policy defined to handle Microsoft SQL Server 2012 data log iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine.
6. Configure static vNICs **vNIC8** on Fabric A and **vNIC9** on Fabric B with VLAN ID 193, with no fabric failover, an MTU of 9000, and **VMotion** QoS Policy definition to handle VMware vMotion traffic.

On static vNICs, fabric failover is not enabled; vNIC failover and load-balancing policy is enabled on individual vSwitches.

Table 11 shows the static vNICs with the network properties created on the service profile.

Table 11. Cisco UCS Service Profile Configuration

vNIC	MAC Address	Fabric ID	Fabric Failover	VLAN ID	MTU	Adapter Policy	QoS Policy
vnic0	0025:b500:0011	Fabric A	No	108	1500	VMware	Management
vnic1	0025:b500:0012	Fabric B	No	108	1500	Vmware	Management
vnic2	0025:b501:0101	Fabric A	No	109	9000	VMware	SQLAPP
vnic3	0025:b501:0102	Fabric B	No	109	9000	VMware	SQLAPP
vnic4	0025:b501:0103	Fabric A	No	192	9000	VMware	MSSQLData
vnic5	0025:b501:0104	Fabric B	No	192	9000	VMware	MSSQLData
vnic6	0025:b501:0105	Fabric A	No	192	9000	VMware	MSSQLLog
vnic7	0025:b599:0007	Fabric B	No	192	9000	VMware	MSSQLLog
vnic8	0025:b598:0008	Fabric A	No	194	9000	VMware	Vmotion
vnic9	0025:b599:0010	Fabric B	No	194	9000	VMware	Vmotion

VMware ESX Host and Guest Virtual Machine Network Design

After booting VMware ESX 5.0 through the iSCSI LUN on the Cisco UCS B230 blade server, you can see the mapping of VMware ESX 5.0 host VMNIC network adapters to Cisco UCS static vNICs on the VMware vCenter server, as shown in Figure 38.

Figure 38. Mapping of VMware ESX 5.0 Host VMNIC Adapters to Cisco UCS Static vNICs

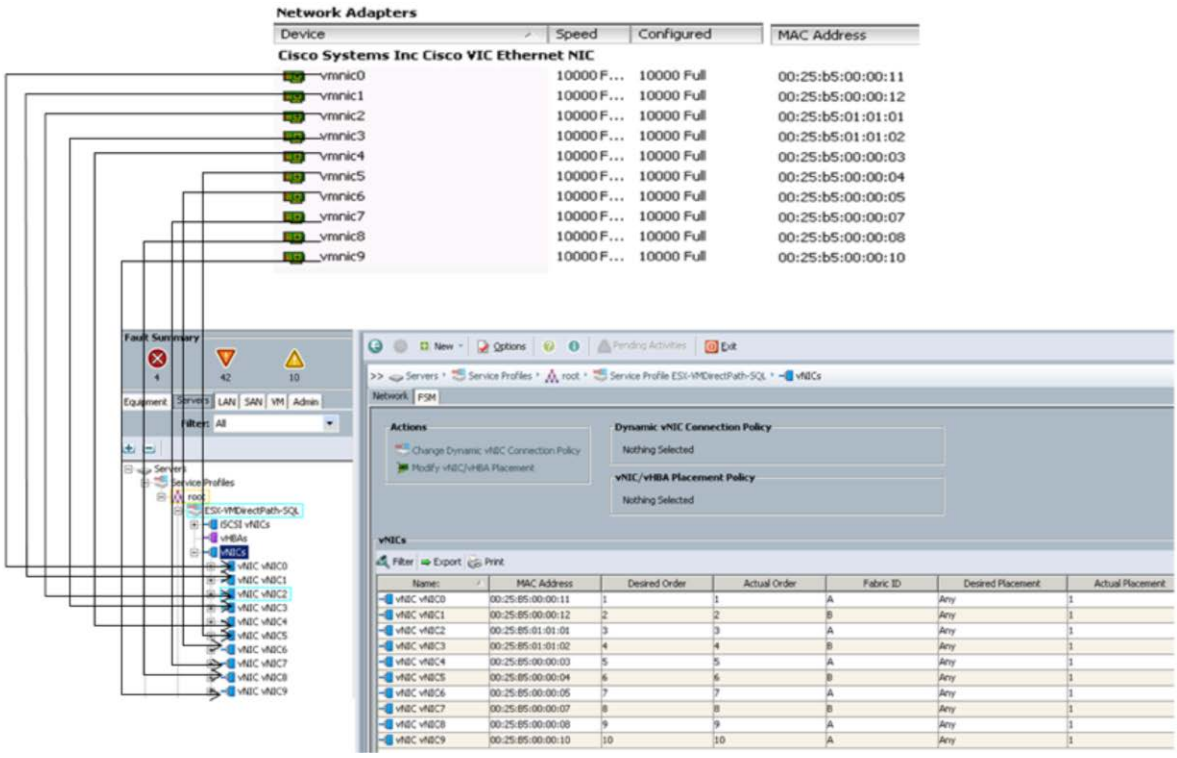

The following are VMware ESX host and Microsoft Windows 2008 R2 guest logical iSCSI design for deploying Microsoft SQL Server 2012:

- Microsoft Windows 2008 R2 guest virtual machine and VMware ESX host management network traffic is accessed through **vSwitch0** in the VMware VMkernel port group **Management Network** on which **vmnic0** and **vmnic1** VMware ESX network adapters are uplinked, with NIC teaming enabled. These VMware ESX network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput.
- To access the Microsoft SQL Server client network traffic from the Microsoft SQL Server guest virtual machine, **vSwitch1** is created with VMware VMkernel port group **MS SQL App Network** on which **vmnic4** and **vmnic5** VMware ESX network adapters are uplinked, with the MTU value set to 9000 and NIC teaming enabled. These network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput.
- To handle VMware vMotion network traffic, **vSwitch6** is created with VMware VMkernel port group **VMware vMotion Network** on which **vmnic8** and **vmnic9** VMware ESX network adapters are uplinked, with VMware vMotion enabled, the MTU value set to 9000, and NIC teaming enabled, and with active-active failover and load-balancing policy to achieve high availability and better network throughput.

Figure 39 shows the VMware ESX 5.0 host Networking tab with the VMware ESX vSwitch configuration on the VMware vCenter server as discussed here.

Figure 39. VMware ESX 5.0 vSwitch Configuration in VMware vCenter Server

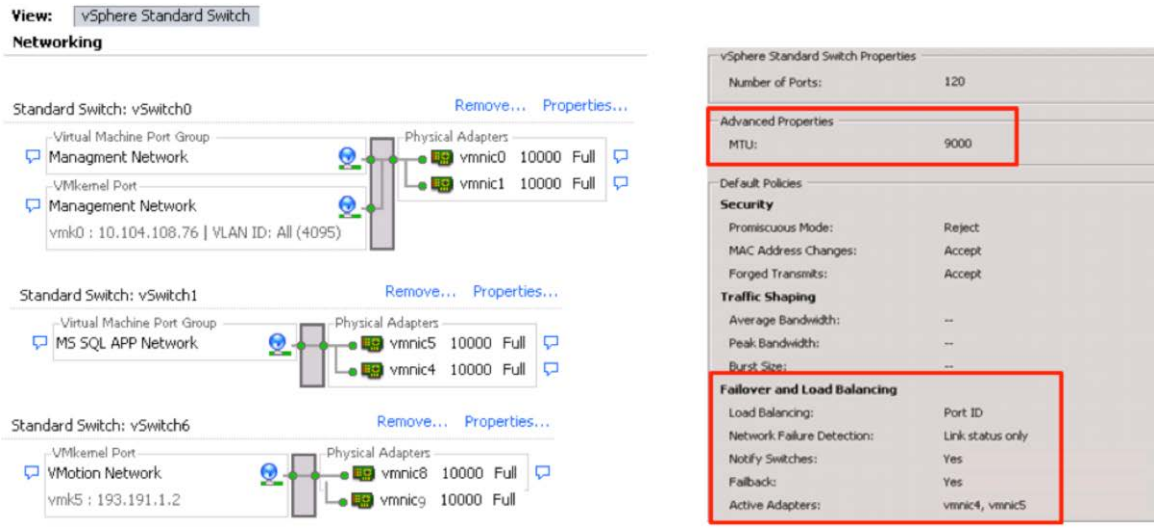

The VMware ESX 5.0 host-based iSCSI software initiator is used to access NetApp iSCSI target LUNs to create the VMware ESX data store for storing the Microsoft SQL Server 2012 database data and log files on the Microsoft Windows 2008 R2 guest virtual machine. The following steps present the high-level configuration process:

1. Access the NetApp storage system iSCSI target VIFs **iscsiA** (192.191.1.2) (NetApp FAS3270 Controller A) and **iscsiB** (192.191.1.3) (NetApp FAS3270 Controller B) using the VMware ESX 5.0 host-based iSCSI software initiator through VMware ESX vSwitches **vSwitch2**, **vSwitch3**, **vSwitch4**, and **vSwitch5**. Table 12 provides the details of the created vSwitches.

Table 12. Details of Created vSwitches

vSwitch	VMware VMkernel interface	Uplink Port	iSCSI Access Target
vSwitch2	vmk2	vmnic6	Log LUN
vSwitch3	vmk3	vmnic7	Log LUN
vSwitch4	vmk1	vmnic2	Data LUN
vSwitch5	vmk4	vmnic3	Data LUN

By default, the **iScsiBootvSwitch** VMware ESX 5.0 vSwitch has two uplink ports (**vmnic2** and **vmnic3**) defined as part of the VMware ESX iSCSI boot installation; these can be deleted for reuse. In this document, these ports are assigned to **vSwitch4** and **vSwitch5** to access the iSCSI storage network as discussed previously.

2. To configure the VMware ESX 5.0 host-based iSCSI software initiator, you need to add the VMkernel interfaces (**vmk1**, **vmk2**, **vmk3**, and **vmk4**) to create the iSCSI binding; refer to the configuration guide at http://www.microsoft.com/download/en/details.aspx?id=18986.

Figure 40 shows the VMware ESX 5.0 host vSwitch and iSCSI initiator adapter configuration settings for accessing the NetApp iSCSI target.

Figure 40. VMware ESX 5.0 Host vSwitch and iSCSI Initiator Adapter Configuration Settings

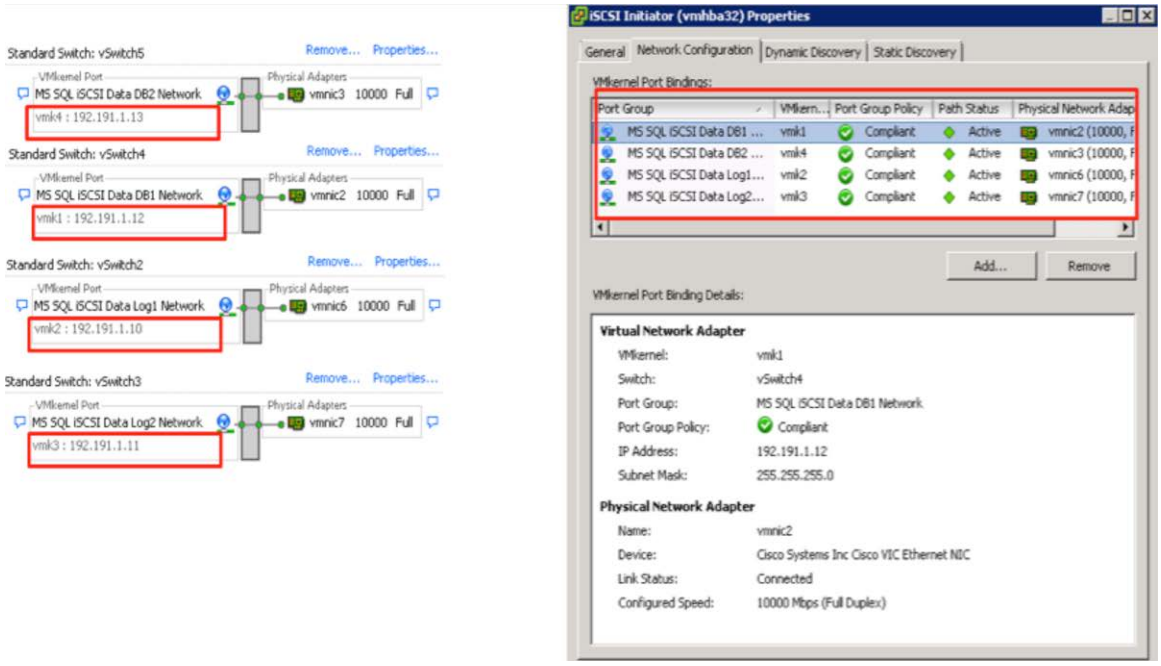

3. To access NetApp storage system iSCSI targets using the VMware ESX 5.0 host-based iSCSI software initiator, manually add iSCSI target IQNs and IP addresses for static discovery, as shown in Figure 41.

Figure 41. iSCSI Initiator Properties in iSCSI Software Adapter

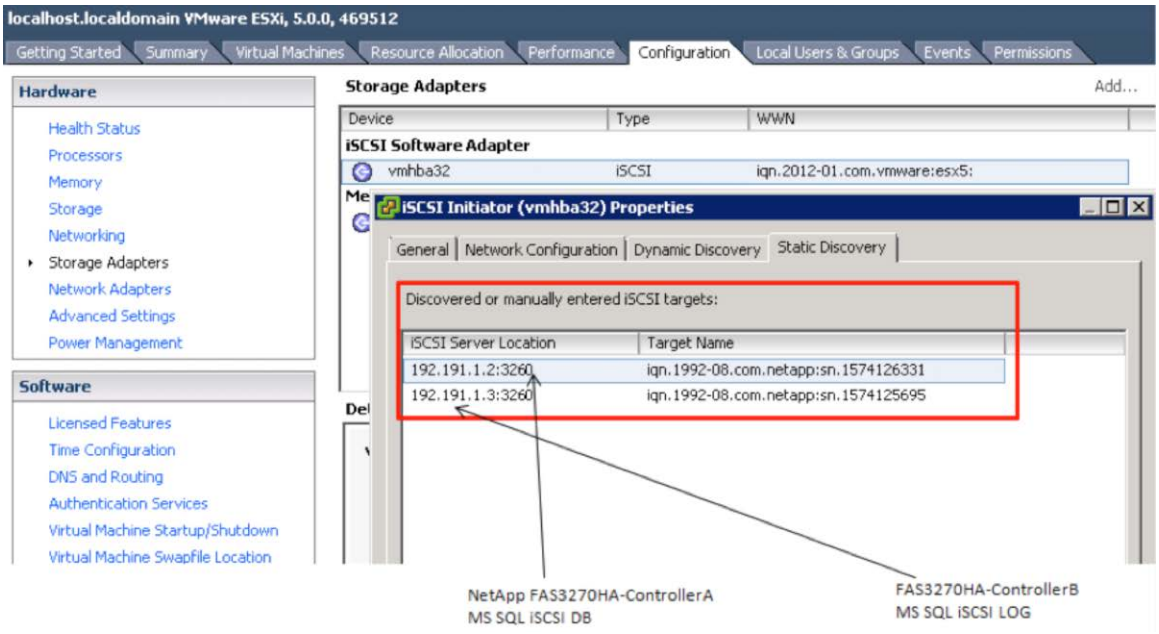

4. After the completion of the iSCSI target configuration, the iSCSI LUNs will be exposed to the VMware ESX 5.0 host in the VMware vCenter storage adapter configuration window, as shown in Figure 42.

Figure 42. iSCSI LUNs Exposed to VMware ESX 5.0 Host in VMware vCenter Server

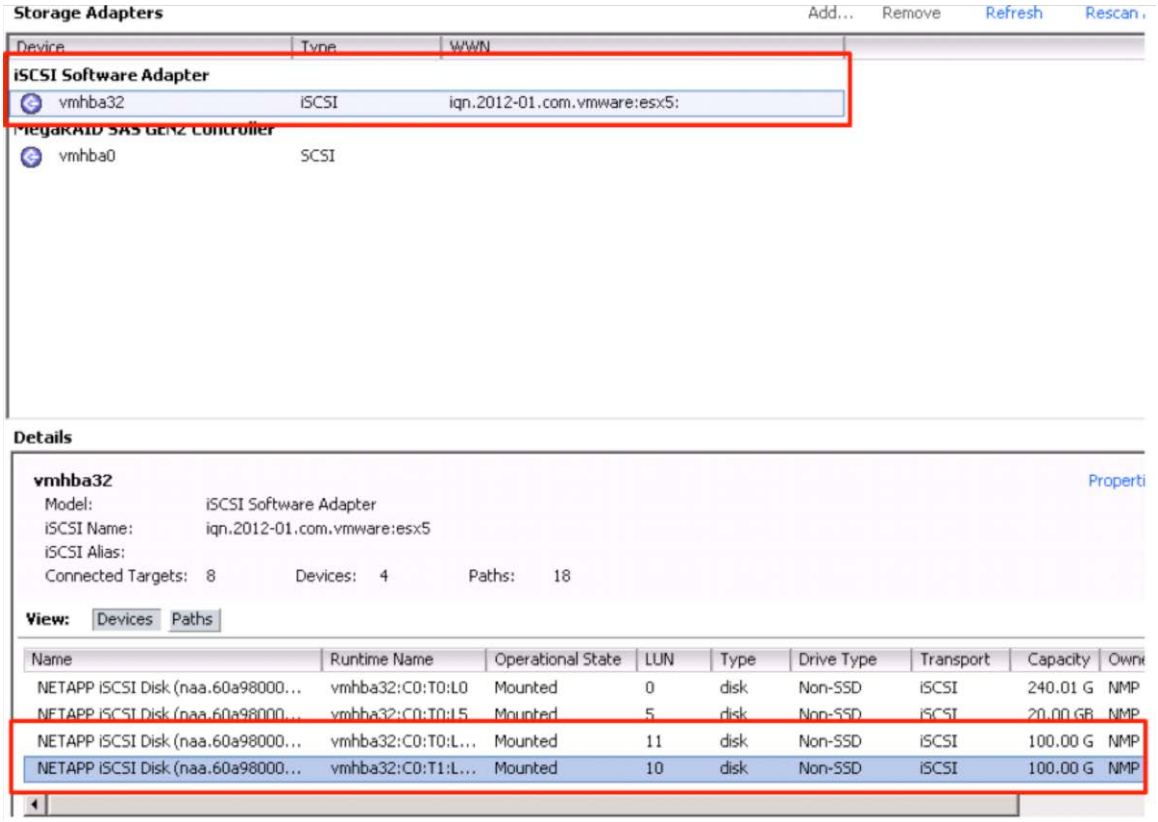

5. On the VMware ESX 5.0 host, create VMFS data stores **MSSQLDatabaseDB** and **MSSQLDatabaseLog** on NetApp exposed to the iSCSI LUN to store Microsoft SQL Server 2012 database and log files, as shown in Figure 43.

Figure 43. Microsoft SQL Server Database Data and Log Exposed to iSCSI LUN

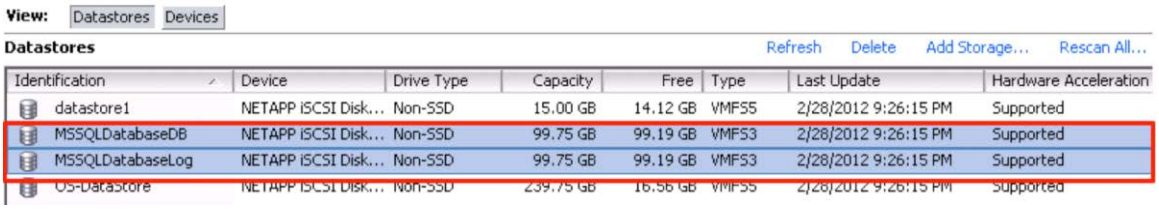

6. On VMware vCenter for the guest virtual machine, configure virtual hard disk (VMDK) **disk2** for storing the Microsoft SQL Server 2012 database file in VMware ESX 5.0 data store **MSSQLDatabaseDB**, and configure VMDK **disk3** for storing the Microsoft SQL Server 2012 database log file in VMware ESX 5.0 data store **MSSQLDatabaseLog**, as shown in Figure 44.

Figure 44. Configuring the Virtual Hard Disk

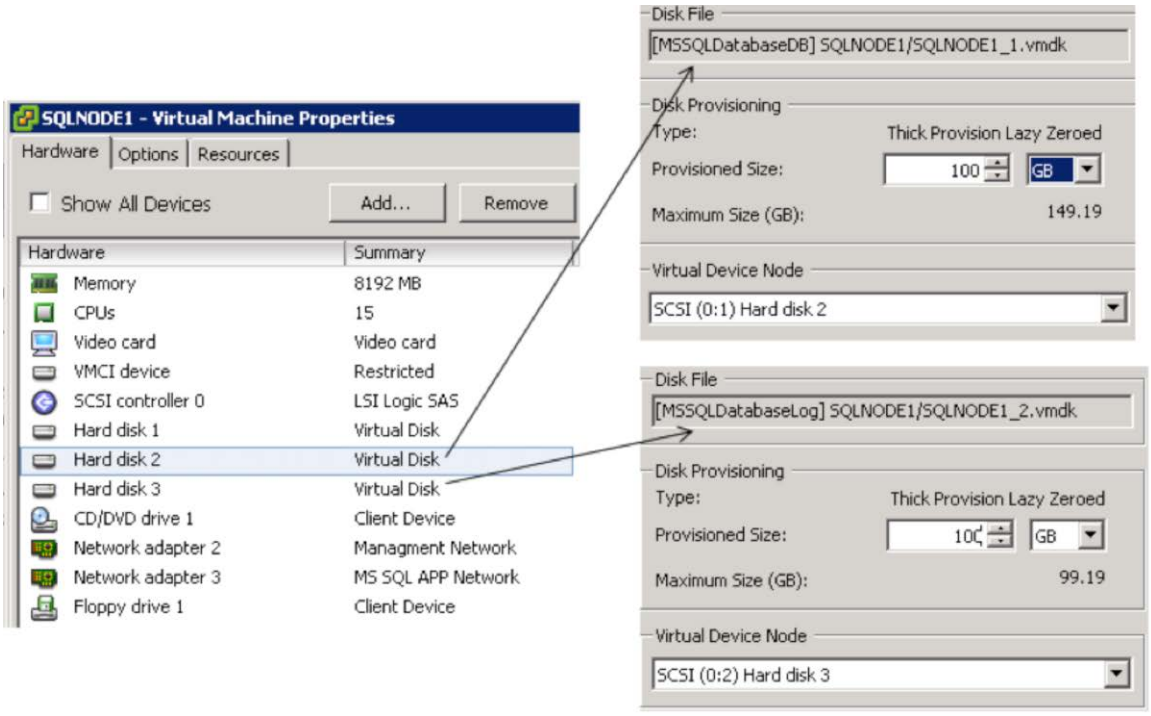

7. On the Microsoft Windows 2008 R2 guest virtual machine, under Disk Management, scan for new disks and then format them and create the NTFS for storing the Microsoft SQL Server 2012 database data and log files, as shown in Figure 45.

Figure 45. Scanning for New Disks in Microsoft Windows 2008 R2 Guest Virtual Machine

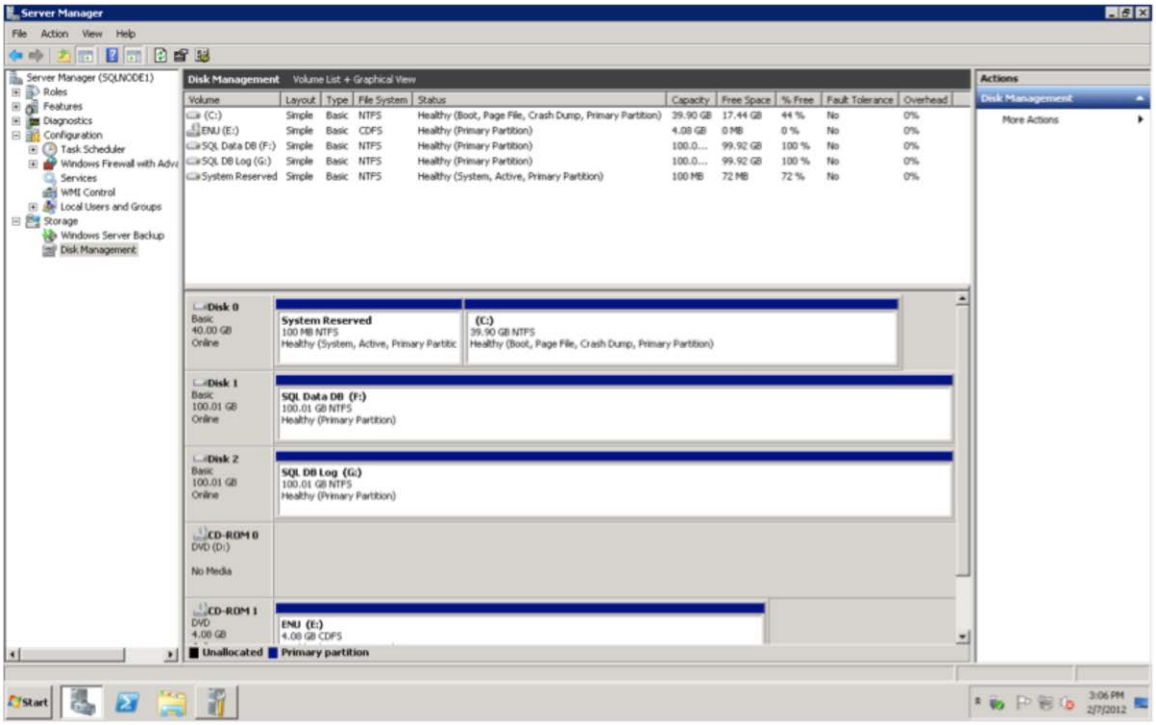

8. Install Microsoft SQL Server on the guest virtual machine and create the database with the data and log files residing on the designated storage volumes created on the guest virtual machine.

   For more information about installing Microsoft SQL Server 2012, see http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx.

Virtual Machine–Based iSCSI Initiator on VMware ESX vSwitch

This section provides an overview of the Microsoft SQL Server single-host deployment in a VMware ESX environment using a Microsoft Windows 2008 R2 guest virtual machine iSCSI software initiator with a VMware ESX native vSwitch to access shared NetApp storage over the iSCSI protocol.

The following section provides a high-level overview of the configuration of iSCSI network infrastructure for a Microsoft SQL Server 2012 single-host installation on Cisco UCS, the VMware ESX 5.0 hypervisor, and a Microsoft Windows 2008 R2 guest virtual machine.

Figure 46 shows the physical and logical architecture of the Cisco UCS configuration and an overview of the VMware ESX vSwitch configuration for deploying Microsoft SQL Server 2012 on a Microsoft Windows 2008 R2 virtual machine with a guest-based iSCSI software initiator.

Figure 46. Physical and Logical Architecture of Cisco UCS Configuration

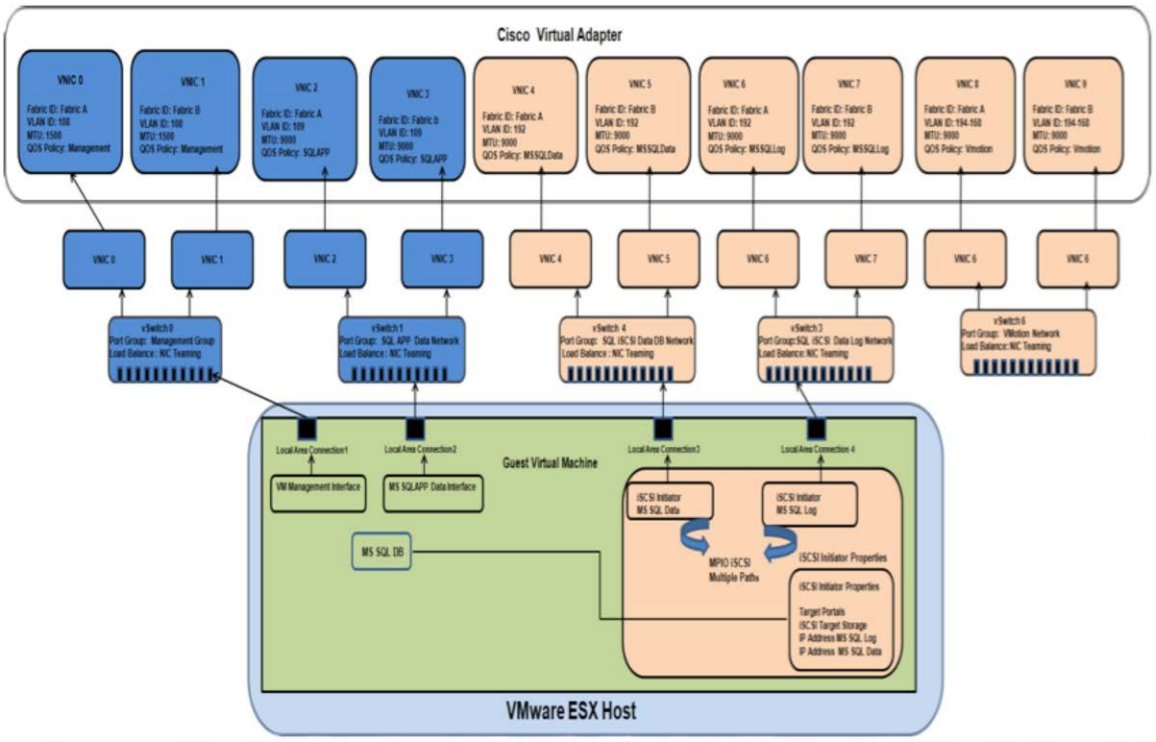

Cisco UCS Service Profile Design

This section explains the network considerations for the Cisco UCS service profile for creating static vNICs to enable the guest virtual machine iSCSI software initiator on VMware ESX 5.0 in a vSwitch environment.

The following vNICs need to be created in the service profile for the virtual machine to access the iSCSI storage target:

- Create two iSCSI vNICs, **iscsi** and **isci0**, which are overlaid on two static vNICs, **vNIC2** on Fabric A and **vNIC3** on Fabric B, respectively. Refer to the VMware ESX iSCSI Boot section for the implementation steps.
- Create two static vNICs, **vNIC0** on Fabric A and **vNIC1** on Fabric B, with VLAN ID 108, no fabric failover, and **Management** QoS policy defined to handle VMware ESX host and Microsoft Windows 2008 R2 guest virtual machine management data network traffic.
- Create two static vNICs, **vNIC4** on Fabric A and **vNIC5** on Fabric B, with VLAN ID 109, no fabric failover, an MTU of 9000, and **SQLAPP** QoS policy defined to handle Microsoft SQL Server 2012 client data network traffic from the Microsoft Windows 2008 R2 guest virtual machine.
- Create two static vNICs, **vNIC2** on Fabric A and **vNIC3** on Fabric B, with VLAN ID 192, no fabric failover, an MTU of 9000, and **MSSQLData** QoS policy defined to handle Microsoft SQL Server 2012 database iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine.
- Create two static vNICs, **vNIC6** on Fabric A and **vNIC7** on Fabric B, with VLAN ID 192, no fabric failover, an MTU of 9000, and **MSSQLLog** QoS policy defined to handle Microsoft SQL Server 2012 log iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine.

- Create two static vNICs, **vNIC8** on Fabric A and **vNIC9** on Fabric B, with VLAN ID 193, no fabric failover, an MTU of 9000, and **VMotion** QoS policy defined to handle VMware vMotion traffic.

Fabric failover is not enabled on static vNICs at the Cisco UCS level. You need to enable vNIC failover and load-balancing policy on individual VMware ESX vSwitches.

Table 13 lists the static vNICs with network properties created in the service profile.

Table 13. Cisco UCS Service Profile Configuration

vNIC	MAC Address	Fabric ID	Fabric Failover	VLAN ID	MTU	Adapter Policy	QoS Policy
vnic0	0025:b500:0011	Fabric A	No	108	1500	VMware	Management
vnic1	0025:b500:0012	Fabric B	No	108	1500	Vmware	Management
vnic2	0025:b501:0101	Fabric A	No	109	9000	VMware	SQLAPP
vnic3	0025:b501:0102	Fabric B	No	109	9000	VMware	SQLAPP
vnic4	0025:b501:0103	Fabric A	No	192	9000	VMware	MSSQLData
vnic5	0025:b501:0104	Fabric B	No	192	9000	VMware	MSSQLData
vnic6	0025:b501:0105	Fabric A	No	192	9000	VMware	MSSQLLog
vnic7	0025:b599:0007	Fabric B	No	192	9000	VMware	MSSQLLog
vnic8	0025:b598:0008	Fabric A	No	194	9000	VMware	VMotion
vnic9	0025:b599:0010	Fabric B	No	194	9000	VMware	VMotion

VMware ESX Host and Guest Virtual Machine iSCSI Design

After booting VMware ESX 5.0 through the iSCSI LUN on the Cisco UCS B230 blade server, you can see the mapping of VMware ESX 5.0 VMNIC network adapters to Cisco UCS static vNICs on the VMware vCenter server, as shown in Figure 47.

Figure 47. Mapping of VMware ESX 5.0 VMNIC Adapters with Cisco UCS Static vNICs

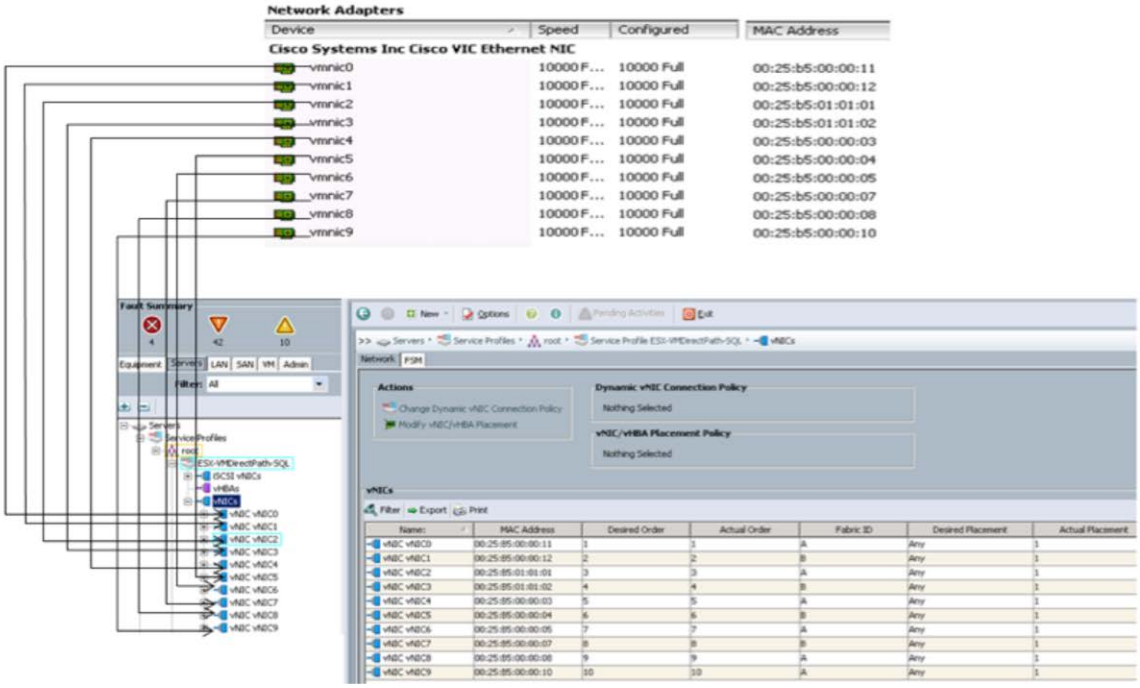

The following are VMware ESX host and Microsoft Windows 2008 R2 guest logical iSCSI design considerations for Microsoft SQL Server 2012 deployment:

- Microsoft Windows 2008 R2 guest virtual machine and VMware ESX host management network traffic is accessed through **vSwitch0** with VMware VMkernel port group **Management Network**, where the vmnic0 and vmnic1 VMware ESX network adapters are uplinked, with NIC teaming enabled. These VMware ESX network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput.
- To access the Microsoft SQL Server 2012 application network traffic from the Microsoft Windows 2008 guest virtual machine on Microsoft SQL Server, create **vSwitch1** with virtual machine port group **MS SQL App Network**, where the **vmnic2** and **vmnic3** VMware ESX network adapters are uplinked, with the MTU value set to 9000 and NIC teaming enabled. These VMware ESX network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput with load balancing.
- To access the Microsoft SQL Server 2012 database network traffic from the Microsoft Windows 2008 R2 guest virtual machine, create **vSwitch4** with virtual machine port group **MS SQL iSCSI Data DB Network**, where the **vmnic2** and **vmnic3** VMware ESX network adapters are uplinked, with the MTU value set to 9000 and NIC teaming enabled. These VMware ESX network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput with load balancing.
- To access the Microsoft SQL Server 2012 database log network traffic from the Microsoft Windows 2008 R2 guest virtual machine, create **vSwitch3** with virtual machine port group **MS SQL iSCSI Data Log Network**, where the **vmnic6** and **vmnic7** VMware ESX network adapters are uplinked, with the MTU value set to

9000 and NIC teaming enabled. These VMware ESX network adapters are configured with active-active failover and load-balancing policy to achieve high availability and better network throughput with load balancing.

- To handle VMware vMotion network traffic, create **vSwitch6** with VMware VMkernel port group **VMotion Network**, on which the **vmnic8** and **vmnic9** VMware ESX network adapters are uplinked, with VMware vMotion enabled, the MTU value set to 9000, and NIC teaming enabled, and with active-active failover and load-balancing policy to achieve high availability and better network throughput with load balancing.

Figure 48 shows the VMware ESX vSwitch configuration on the VMware ESX host Networking tab of the VMware vCenter server.

Figure 48. VMware ESX vSwitch Configuration on VMware ESX Host in VMware vCenter

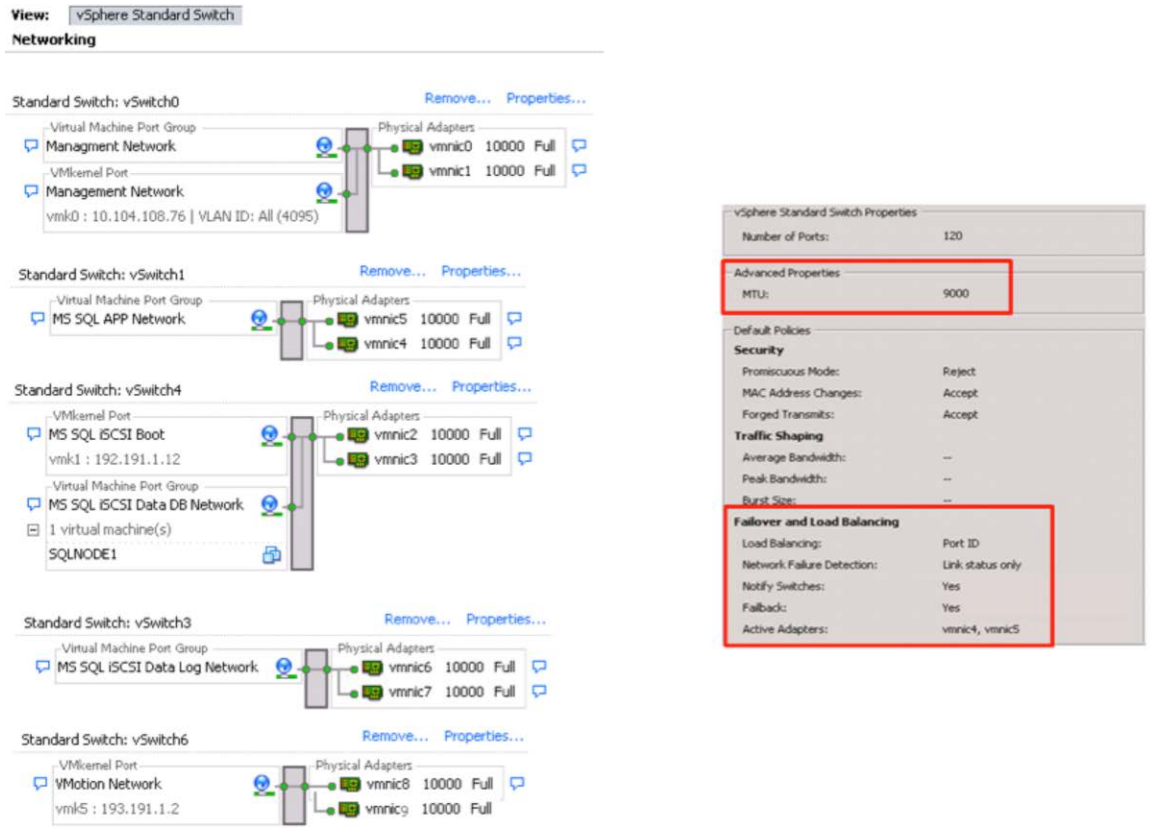

Perform the following steps on the Microsoft Windows 2008 R2 guest virtual machine to configure the software-based iSCSI environment.

1. On VMware ESX for the Microsoft Windows 2008 R2 guest virtual machine, create four virtual adapters to access management, the Microsoft SQL Server client, and the iSCSI storage network.
2. Attach Virtual Network Adapter 1 to the VMware ESX 5.0 **Management Network** port group on **vSwitch0**, which is defined for accessing the guest virtual machine management network.
3. Attach Virtual Network adapter 2 to the VMware ESX 5.0 **MS SQL APP Network** port group on **vSwitch1**, which is defined for accessing the Microsoft SQL Server database from the client network.

4. Attach Virtual Network Adapter 3 to the VMware ESX 5.0 **MS SQL iSCSI Data DB Network** port group on **vSwitch4**, which is defined for accessing the Microsoft SQL Server database network for storing the Microsoft SQL Server 2012 database file on NetApp iSCSI storage.
5. Attach Virtual Network Adapter 4 to the VMware ESX 5.0 **MS SQL iSCSI Data Log Network** port group on **vSwitch3**, which is defined for accessing the Microsoft SQL Server database log network for storing the Microsoft SQL Server 2012 database log file on NetApp iSCSI storage.

Figure 49 shows Microsoft Windows 2008 guest virtual machine virtual network adapter settings on the VMware vCenter server.

Figure 49. Guest Virtual Machine Virtual Network Adapter Settings in VMware vCenter

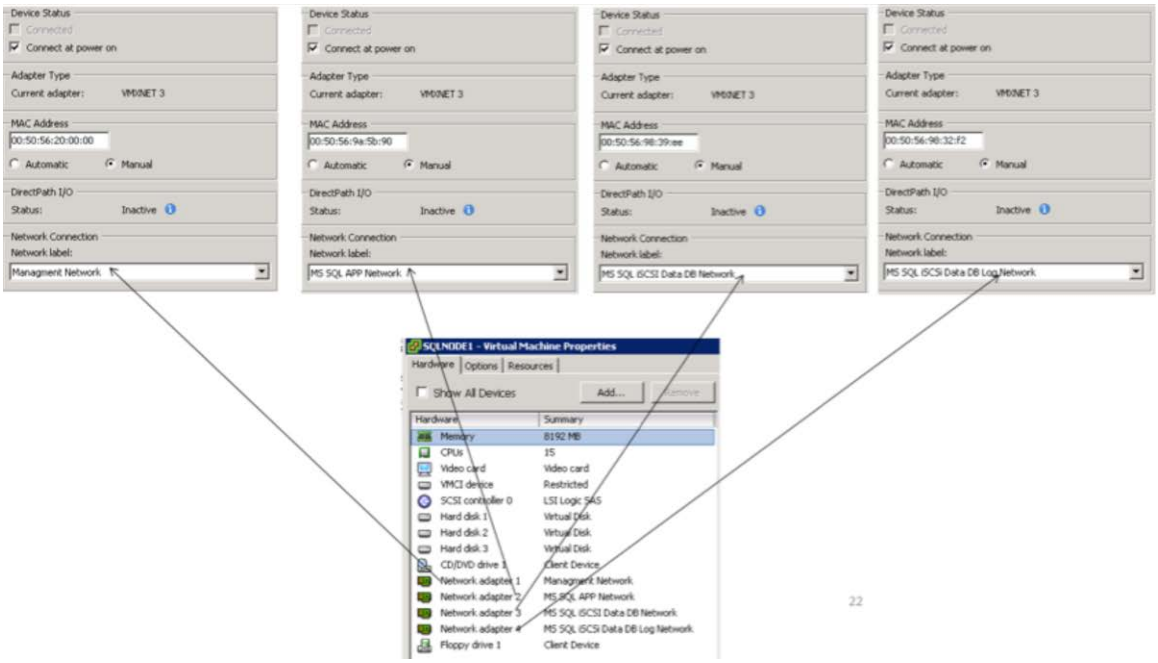

6. Configure each network interface on the Microsoft Windows 2008 R2 guest virtual machine with the required IP address for accessing management, the Microsoft SQL Server 2012 client network, the database data, and the database log through the iSCSI network, as shown in Figure 50.

Figure 50. Configuring Each Network Interface on the Guest Virtual Machine

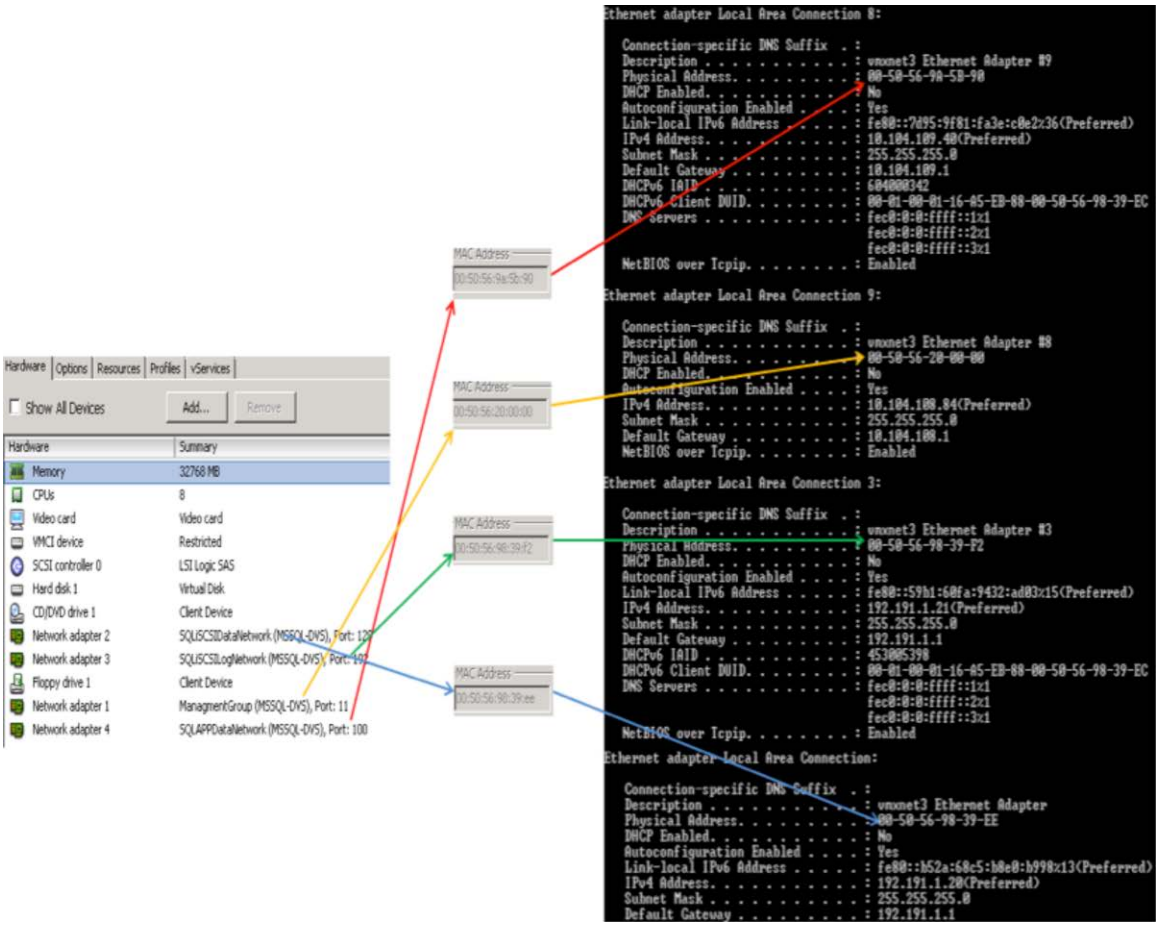

7. To support end-to-end jumbo frame (MTU 9000) to carry Microsoft SQL Server client and iSCSI traffic from the Microsoft Windows 2008 R2 guest virtual machine, Cisco UCS, Cisco Data Center VM-FEX, and NetApp storage, perform the following steps:

   a. Configure MTU 9000 in the Cisco UCS QoS system with Platinum, Gold, Silver, and Bronze classes as shown in the section above

   b. Configure MTU 9000 in the Jumbo field on the appropriate network interfaces (0, 8, and 3 in this design) in the guest virtual machine.

   c. Configure NetApp iSCSI VIFs to enable the MTU 9000 value, as shown in the section above.

8. On the Microsoft Windows 2008 R2 guest virtual machine, enable and configure the iSCSI software initiator and MPIO to access NetApp iSCSI targets.

   For more information about configuring the iSCSI initiator in Microsoft Windows, see http://technet.microsoft.com/en-us/library/ee338476%28v=ws.10%29.

   The following steps provide a high-level overview of the configuration of the Microsoft Windows 2008 R2 guest virtual machine iSCSI software initiator for the NetApp iSCSI target:

   a. Discover the storage controller NetApp FAS3270HA Controller A VIF **iscsiA** (192.191.1.2) with Microsoft Windows 2008 R2 guest virtual machine iSCSI Initiator Adapter 3 (192.191.1.21) and the NetApp FAS3270HA

Controller B VIF **iscsiB** (192.191.1.3) with Microsoft Windows 2008 R2 guest virtual machine iSCSI Initiator Adapter 0 (192.191.1.20), as shown in Figure 51.

Figure 51. iSCSI Initiator Properties Showing Target IP Addresses

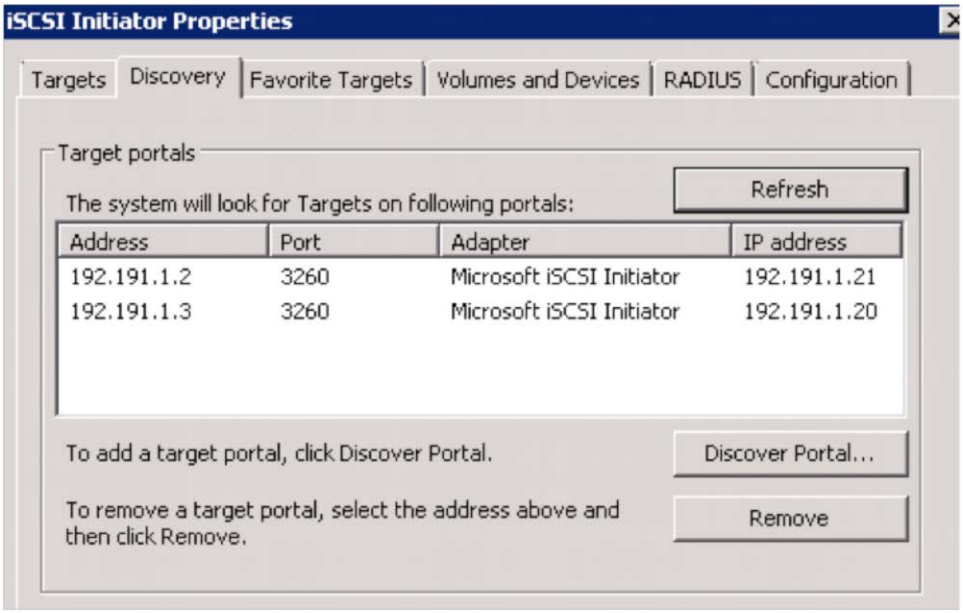

b. To enable iSCSI multipathing, check Enable multi-path in the target connection configuration step for both controllers. Figure 52 shows multipath enabled.

Figure 52. Multipath Enabled on Target Connection

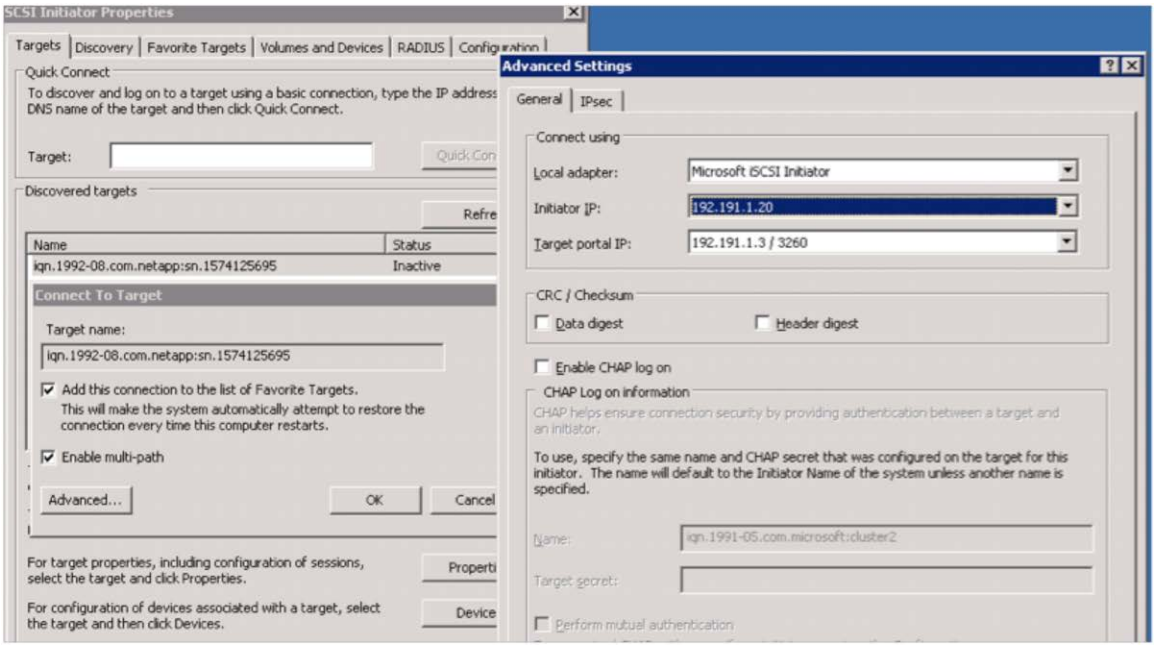

c. On the Microsoft Windows 2008 R2 guest virtual machine, log into the iSCSI initiators to access NetApp targets, and LUNS are automatically exposed for configuration. For a Microsoft SQL Server 2012 single-host installation, you use Disk 1 as the database file and Disk 2 as the log file, as shown in Figure 53. Refer to the section above for LUN creation and access details.

Figure 53. Disk 1 and Disk 2 Details for Microsoft SQL Server 2012 Single-Host Installation

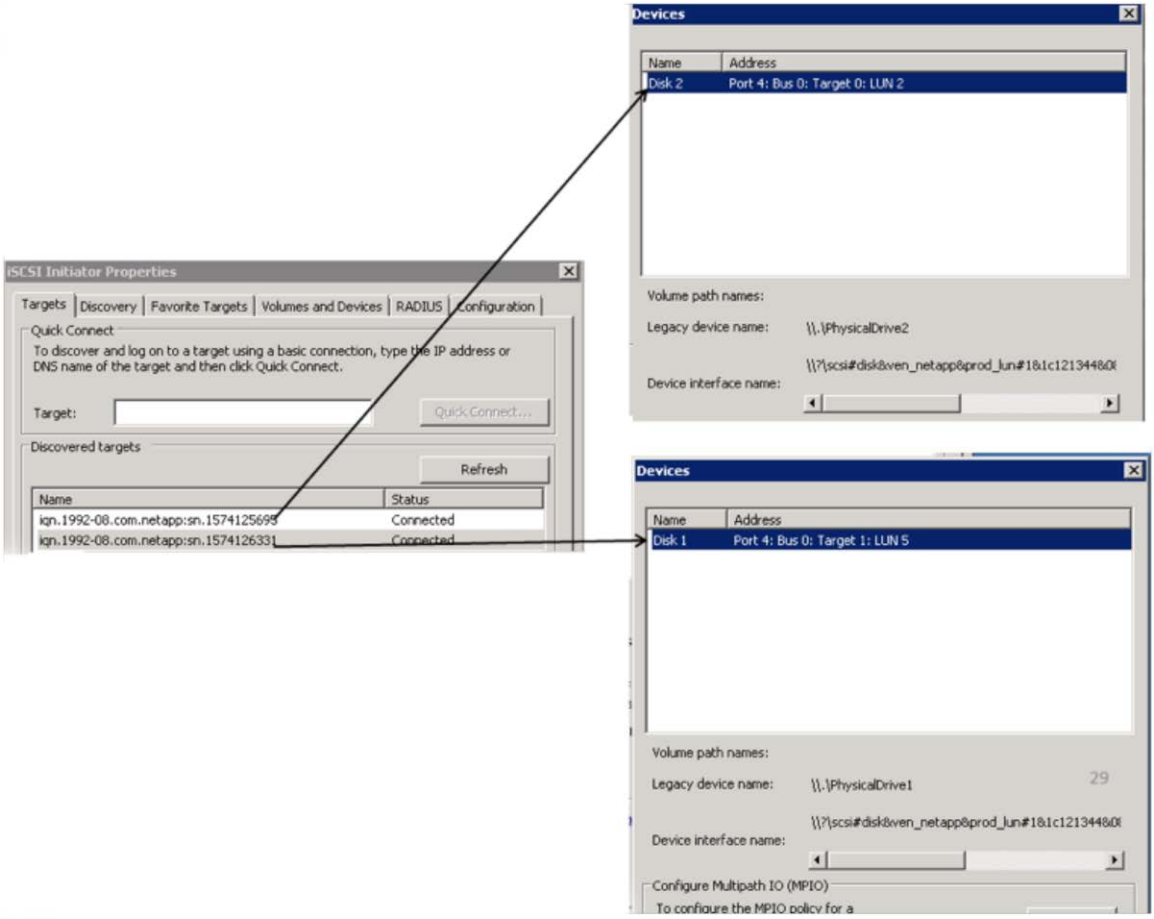

d. Under Disk Management, on the Microsoft Windows 2008 R2 guest virtual machine, scan for new disks and format them. Create the NTFS for storing Microsoft SQL Server 2012 database data and log files, as shown in Figure 54.

Figure 54. Scanning for New Devices on Microsoft Windows 2008 R2 Guest Virtual Machine

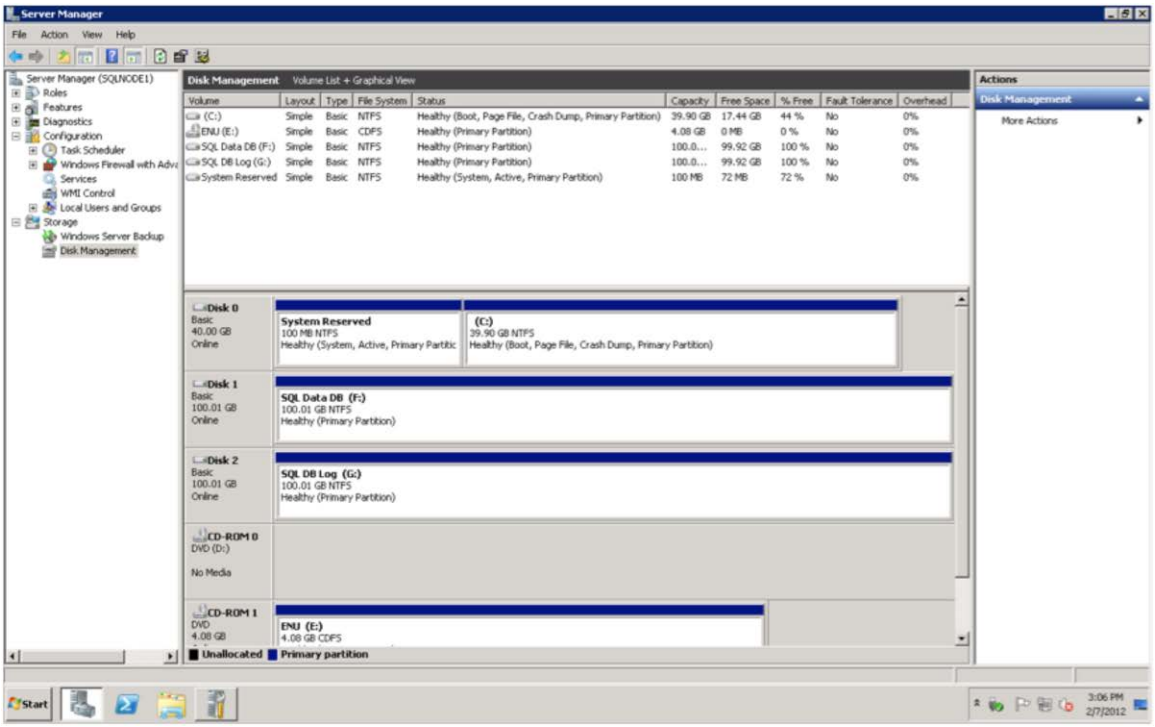

9. Install Microsoft SQL Server on the guest virtual machine and create the database with the data and log files residing on the designated storage volumes created on the guest virtual machine.

   For more information about installing Microsoft SQL Server 2012, see http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx.

Microsoft SQL Server Failover Cluster with Cisco Data Center VM-FEX DVS Solution

This section provides a high-level physical and logical procedure for setting up a Microsoft SQL Server 2012 cluster failover deployment on a Microsoft Windows 2008 R2 guest virtual machine–based software iSCSI initiator on a VMware ESX 5.0 host with a Cisco Data Center VM-FEX DVS in Cisco UCS to access NetApp shared iSCSI storage in an iSCSI network environment.

Failover clustering provides very good protection in the event of hardware failure. Failover to a passive node is fairly quick (between one and five minutes depending on the state of the cluster and database). Failover clustering provides service availability but does not provide data redundancy like database mirroring and log shipping. Data protection has to be provided at the storage level or in combination with other solutions.

Failover clustering provides host-level protection built on Microsoft Windows failover clustering. Cluster nodes typically are co-located within the same site or data center to provide local availability, but they can also be deployed regionally.

VMware ESX vSphere vMotion is not supported on virtual machines that are part of a Microsoft failover cluster domain. Virtual machines that are part of the Microsoft failover cluster will not be able take advantage of the automatic VMware Distributed Resource Scheduler (DRS) feature. For more information about Microsoft clustering

in the VMware vSphere knowledge base, see http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1037959.

Physical and Logical Design

This section provides a high-level overview of the physical and logical infrastructure design considerations required to deploy a Microsoft SQL Server 2012 failover cluster on a Microsoft Windows 2008 R2 guest virtual machine iSCSI software initiator hosted on VMware ESX 5.0 with Cisco UCS Cisco Data Center VM-FEX DVS on a Cisco UCS B230 blade server.

This document describes Microsoft SQL Server 2012 failover clustering within the same site on a single Cisco UCS platform across two Cisco UCS B230 blade servers. However, for high availability you can have blades on different chassis managed by a single Cisco UCS platform.

Figure 55 shows the physical and logical design of the Microsoft SQL Server 2012 failover cluster solution on a Microsoft Windows 2008 R2 guest virtual machine iSCSI software initiator on a VMware ESX 5.0 host with Cisco Data Center VM-FEX DVS.

Figure 55. Physical and Logical Design of Microsoft SQL Server 2012 Failover Cluster Solution

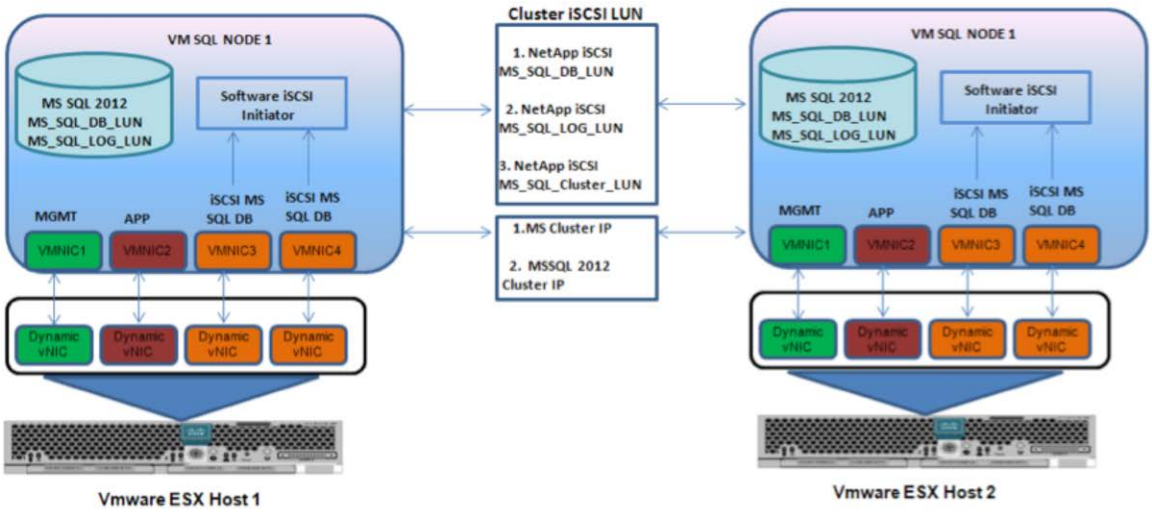

Perform the following steps to implement failover clustering on two Microsoft Windows 2008 R2 guest virtual machine software-based iSCSI systems. These Microsoft Windows 2008 R2 nodes are part of a Microsoft Windows failover cluster, booted through the iSCSI target. Access to the iSCSI target is faciliated through an instance of the Cisco Data Center VM-FEX DVS with direct I/O path enabled.

1. To implement Microsoft server clustering and a Microsoft SQL Server 2012 failover cluster server on Microsoft Windows 2008 R2 guest Virtual machines, you need to use two Cisco UCS B230 blades (**ESX Host 1** and **ESX Host 2**) in a single chassis, as shown in Figure 55.

   Define two service profiles with the required network infrastructure to install VMware ESX 5.0 iSCSI boot operating systems, which are clustered to host two failover cluster Microsoft Windows 2008 R2 guest virtual machines using the VMware vCenter server. The design of the VMware ESX iSCSI boot OS and guest-based iSCSI initiator for individual hosts is explained in the sections VMware ESX iSCSI Boot and above respectively.

2. To deploy Microsoft Windows 2008 failover cluster mode on Microsoft Windows 2008 R2 guest virtual machine cluster nodes **VMSQLNODE1** and **VMSQLNODE2**, attach four dynamic vNICs as shown in Figure 55.

   a. Attach **VMNIC1** with the Cisco Data Center VM-FEX port profile **Managementgroup** to access the Microsoft Windows 2008 R2 guest virtual machine for operations.

   b. Attach **VMNIC2** with the Cisco Data Center VM-FEX port profile **SQLAPPDataNetwork** for internal and external clients to access the Microsoft SQL Server 2012 failover cluster server.

   c. Attach **VMNIC3** with the Cisco Data Center VM-FEX port profile **SQLiSCSIDataNetwork** to access the iSCSI storage NetApp FAS3270HA Controller B VIF target, which hosts cluster LUNs for Microsoft SQL Server database data files. The LUN is accessed through the Microsoft Windows 2008 guest virtual machine iSCSI software initiator.

   d. Attach **VMNIC4** with the Cisco Data Center VM-FEX port profile **SQLiSCSILogNetwork** to access the iSCSI storage NetApp FAS3270HA Controller A VIF target, which hosts cluster LUNs for Microsoft SQL Server database log files. It is accessed through the Microsoft Windows 2008 guest virtual machine iSCSI software initiator.

   e. Attach the Cisco UCS service profile with two Cisco UCS B230 blades (**ESX Host 1** and **ESX Host 2**), which have dynamic vNICs **VMINC1**, **MNIC2**, **VMNIC3**, and **VMNIC4** to Cisco Data Center VM-FEX port profiles with **VMDirectPath** enabled to provide higher performance. Enabling VMDirectPath results in better VMware ESX utilization by reducing the number of VMware ESX host CPU cycles for handling I/O. For more information, refer to the section above.

3. Perform the following steps to design the iSCSI NetApp storage target for deploying a Microsoft SQL Server 2012 failure cluster instance on Microsoft Windows 2008 R2 guest virtual machine cluster nodes **VMSQLNODE1** and **VMSQLNODE2**.

   a. Microsoft Windows 2008 R2 guest virtual machine cluster nodes **VMSQLNODE1** and **VMSQLNODE2** use the iSCSI software initiator configured with the **VMNIC3** and **VMNIC4** network interfaces to access the NetApp cluster storage iSCSI target VIF (NetApp FAS3270HA Controller B **iscsiB** and NetApp FAS3270HA Controller A **iscsiA**) with multipath enabled to access the Microsoft SQL Server 2012 database data and log LUNs, as explained in the section above.

   b. On NetApp storage systems (NetApp FAS3270HA Controller A and B), provision cluster LUNs **MS_SQL_Cluster_LUN** for storing Microsoft cluster quorum data, and **MS_SQL_DB_LUN** and **MS_SQL_LOG_LUN** for storing shared Microsoft SQL Server 2012 failover cluster database data and log files. These LUNs are exposed through the iSCSI network on the Microsoft Windows 2008 R2 guest virtual machine, which is part of the Microsoft Windows server.

   c. Make sure that you create igroups on both NetApp storage controllers, NetApp FAS3270HA Controller A and NetApp FAS3270HA Controller B, with both the **VMSQLNODE1** and **VMSQLNODE2** iSCSI initiator IQN names and map **MS_SQL_Cluster_LUN**, **MS_SQL_DB_LUN** and **MS_SQL_LOG_LUN** to those IQNs as explained in the section above.

   d. On exposing NetApp storage LUNs to the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1** and **VMSQLNODE2** cluster nodes, scan for new disks in the disk manager and format the disk. Assign the same drive letter to both cluster nodes, as shown in Figure 56.

Figure 56. Scanning for New Disks and Assigning New Disks to Cluster Nodes

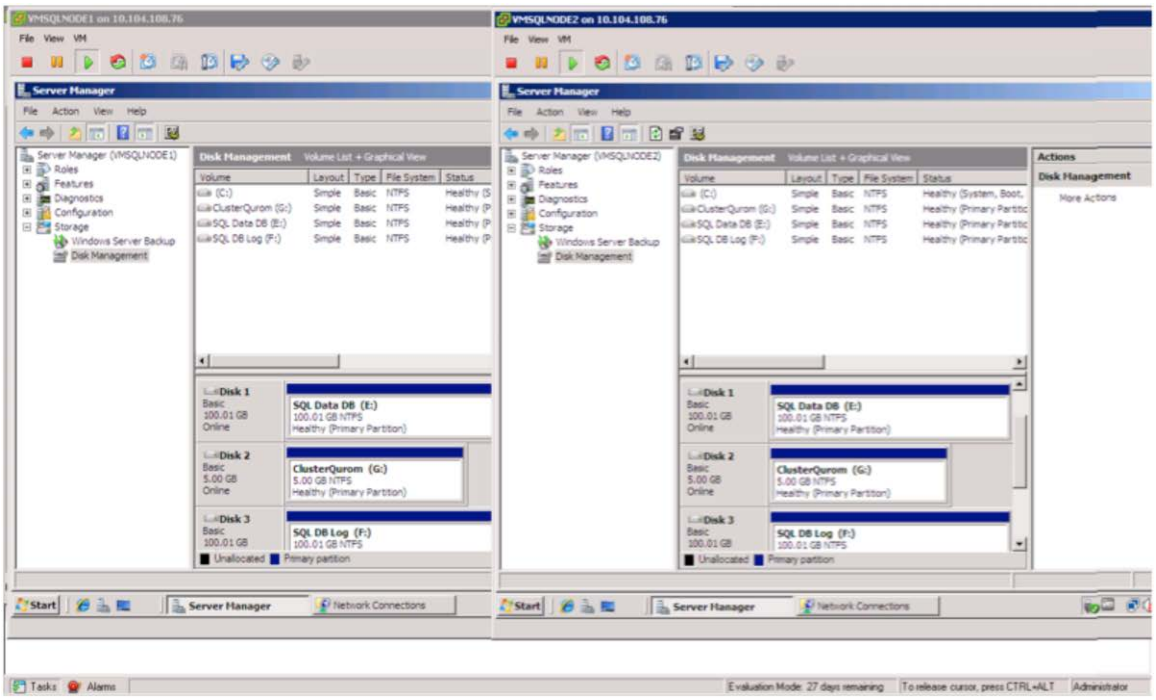

Installation of Microsoft Windows 2008 Failover Cluster Feature with Guest Virtual Machine–Based iSCSI Software Initiator

After the configuration of the two Microsoft Windows 2008 guest virtual machines, **VMSQLNODE1** and **VMSQLNODE2**, is complete, perform the following steps to deploy Microsoft Windows 2008 failover clustering:

1. Before installing the Microsoft failover clustering feature on Microsoft Windows 2008 R2 guest virtual machines **VMSQLNODE1** and **VMSQLNODE2**, make sure that they are part of a domain controller.
2. Log in to individual Microsoft Windows 2008 guest virtual machines **VMSQLNODE1** and **VMSQLNODE2** by selecting the domain controller with Admin credentials.
3. Add the failover clustering feature on both Microsoft Windows 2008 R2 guest virtual machines, **VMSQLNODE1** and **VMSQLNODE2**.
4. After installation of the failover cluster feature on both Microsoft Windows 2008 R2 guest virtual machines, log in to either virtual machine (Microsoft Windows 2008 R2 guest virtual machines **VMSQLNODE1** or **VMSQLNODE2**) and launch the Failover Cluster Management console to validate clustering. Add Microsoft SQL Server guest virtual machines **VMSQLNODE1** and **VMSQLNODE2** with the fully qualified domain name in the Select Servers or a Cluster window, as shown in Figure 57.

Figure 57. Adding Microsoft SQL Server Guest Virtual Machines

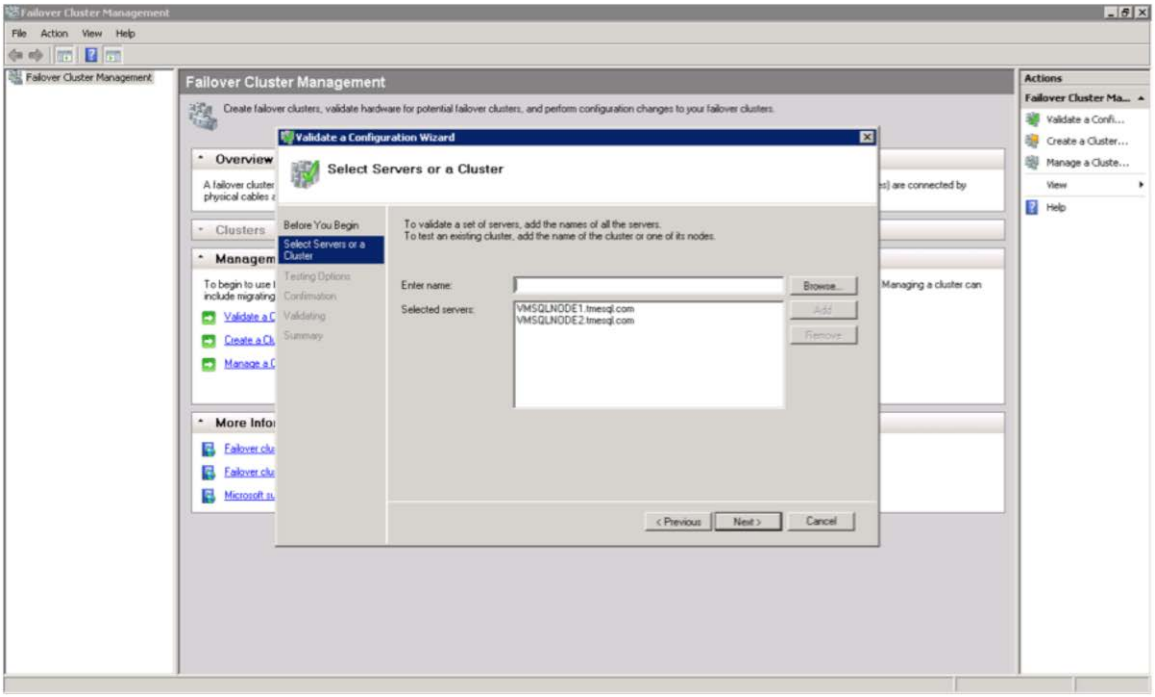

5. Select and add Microsoft Windows 2008 R2 guest virtual machines **VMSQLNODE1** and **VMSQLNODE2** and be sure to run all tests to validate the cluster requirements, as shown in Figure 58.

Figure 58. Running Tests to Validate Cluster Requirements

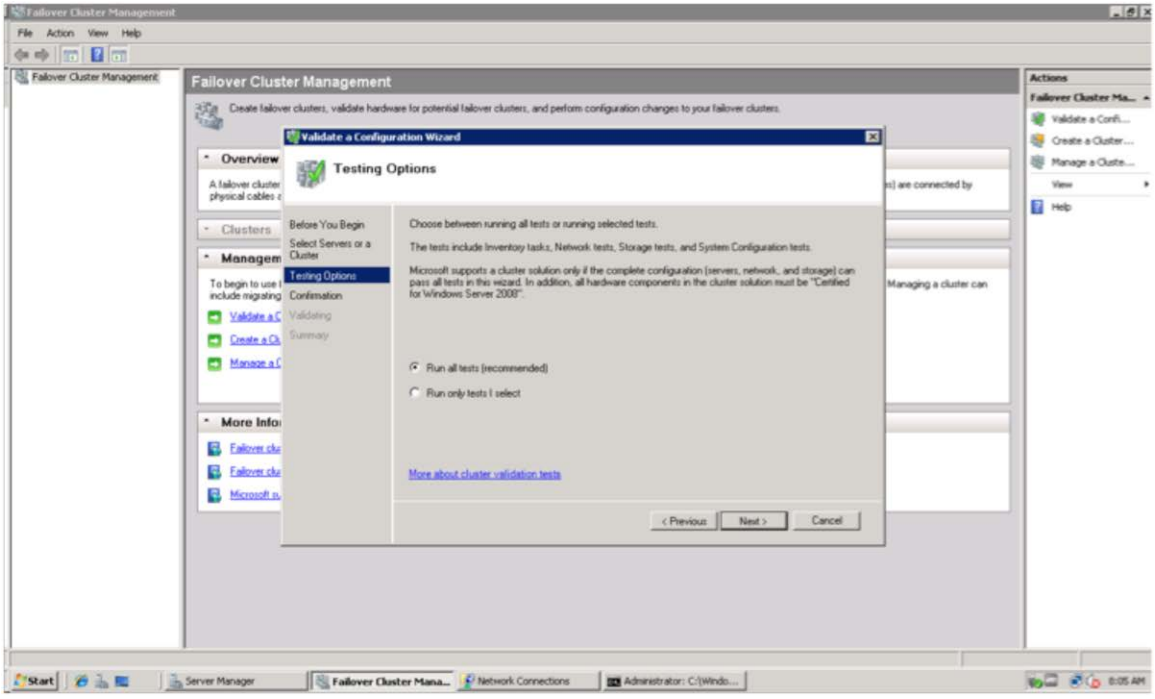

After successful validation of the addition of guest virtual machines **VMSQLNODE1** and **VMSQLNODE2** to the cluster as shown in Figure 59, the Microsoft Windows failover cluster is ready.

Figure 59. Failover Cluster Validation Report

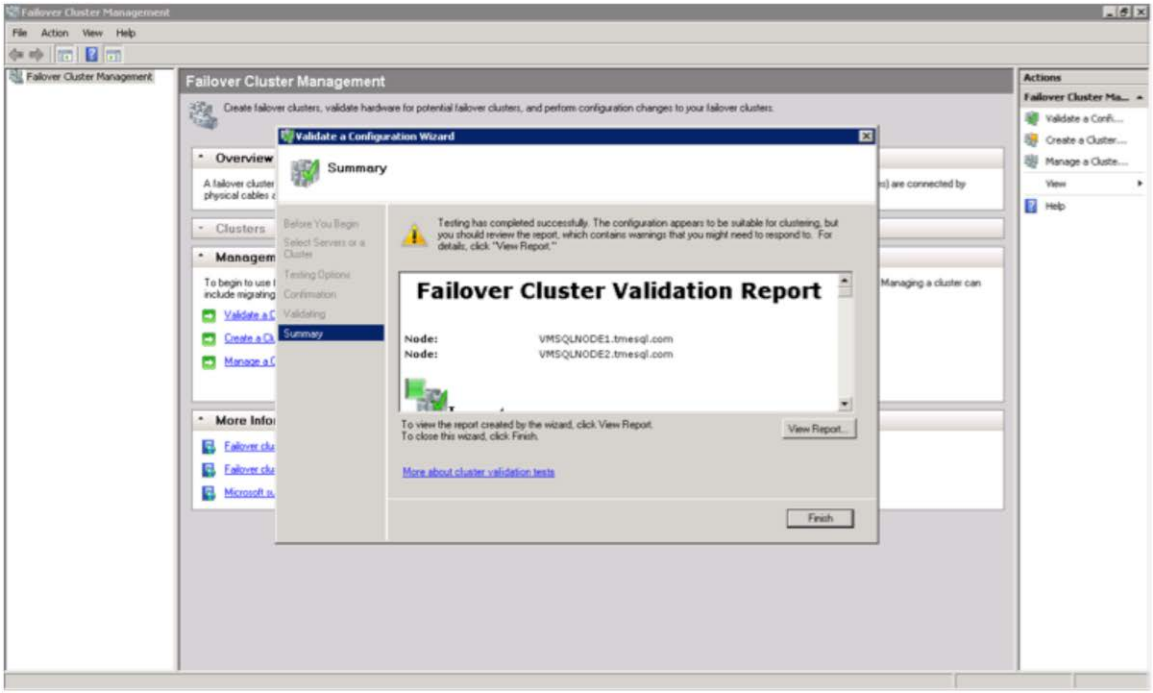

6. Create cluster virtual name **MSSQL** with the cluster and IP address on the management VLAN (108) as shown in Figure 60.

Figure 60. Administering the Cluster

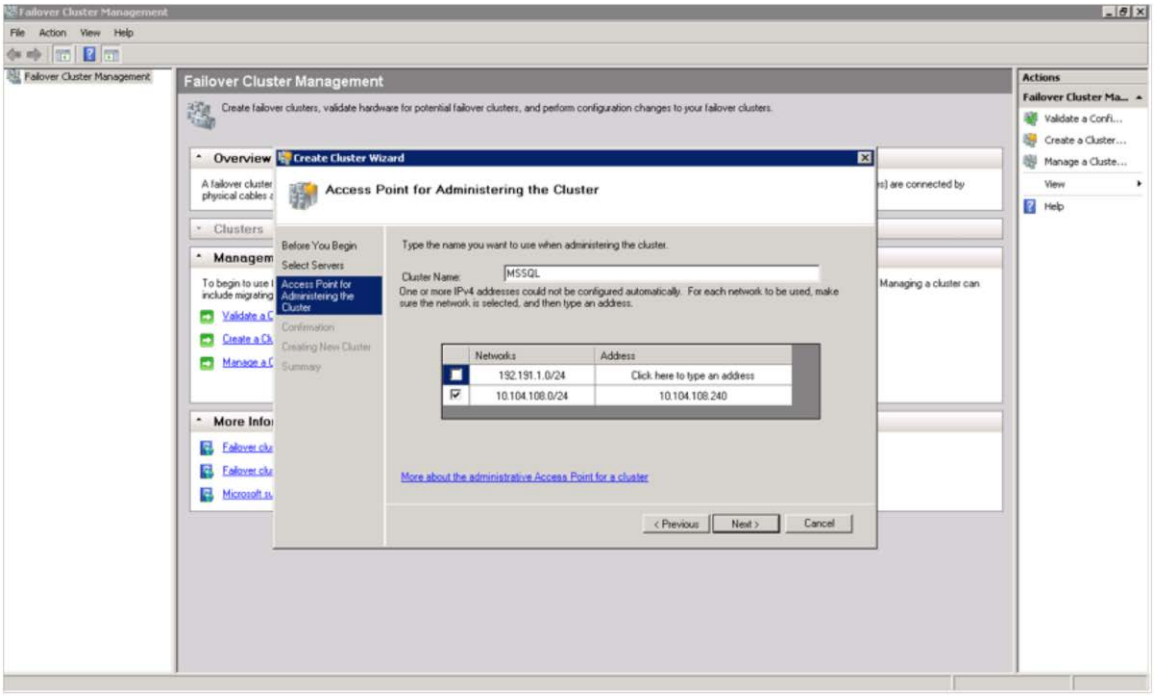

Figure 61 shows cluster summary information prior to creation of the cluster with the **VMSQLNODE1** and **VMSQLNODE2** nodes.

Figure 61. Summary of the Created Cluster

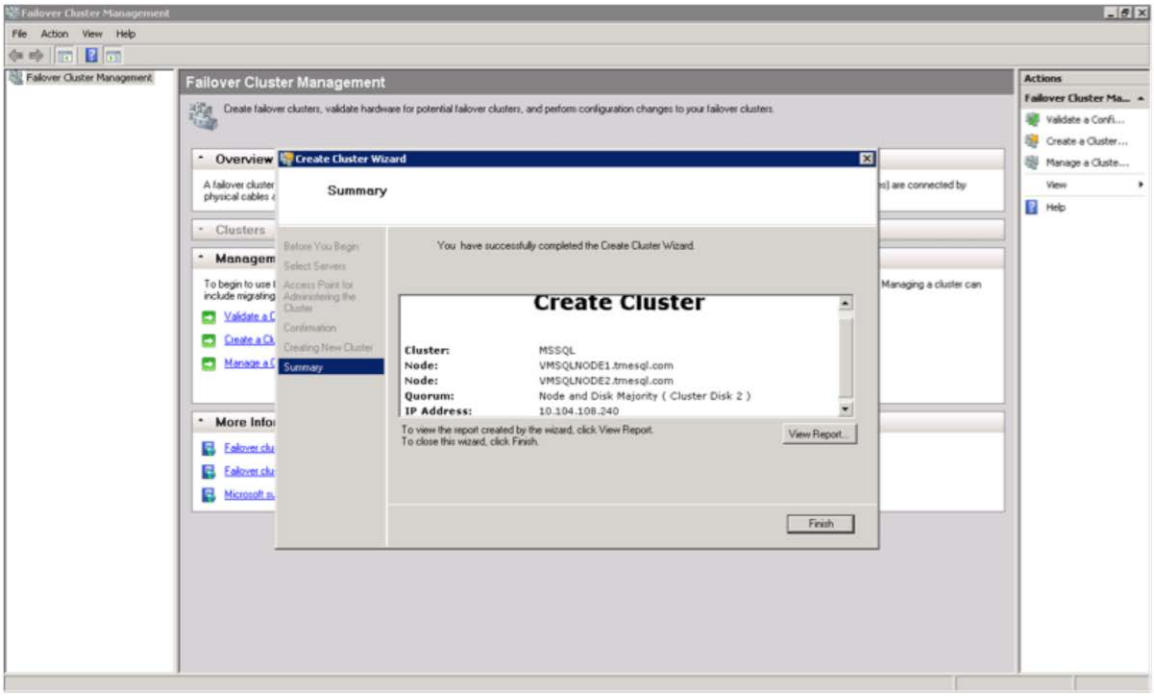

7. To validate Microsoft Cluster installation, log in to either the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1** or **VMSQLNODE2** cluster node and launch the Failover Cluster Management console as shown in Figure 62.

Figure 62. Validating Microsoft Cluster Installation

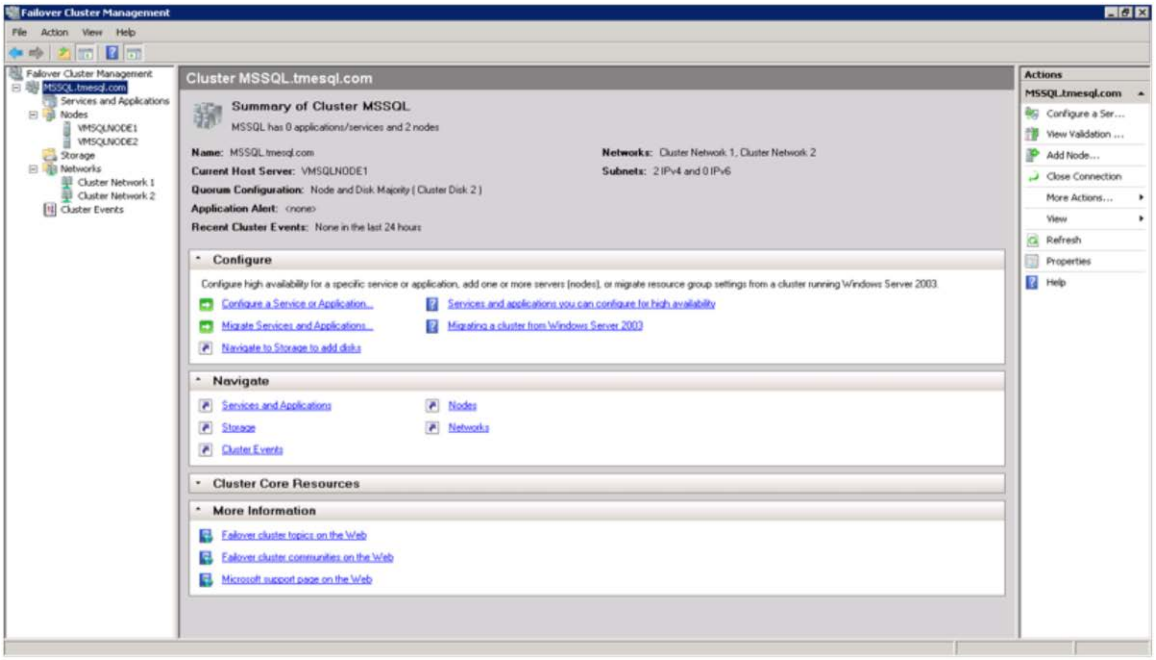

8. Log in to either the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1** or **VMSQLNODE2** cluster node and choose the cluster quorum disk for storing cluster information.
9. On the Failover Cluster Management console, select More Actions and then select the cluster name. In the Configure Cluster Quorum Wizard, under Quorum Configuration, choose Node and Disk Majority (recommended for the number of nodes used here) and select the MS_SQL_Cluster LUN disk drive as the quorum disk as shown in Figure 63. Configuring Storage Witness.

Figure 63. Configuring MS SQL Cluster Quorum

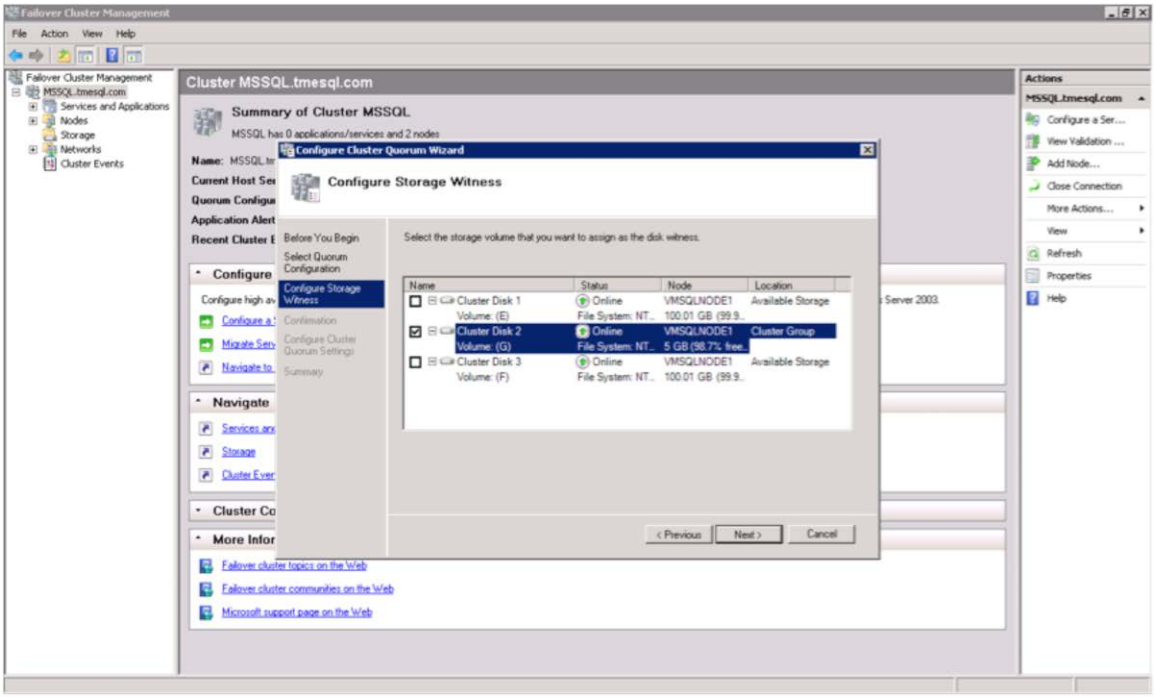

Installation of Microsoft SQL Server 2012 Failover Cluster Feature with Guest Virtual Machine–Based iSCSI Storage

The following steps describe deployment of Microsoft SQL Server 2012 failure clustering on Microsoft Windows 2008 R2 guest virtual machines.

1. Log in to either of the Microsoft Windows 2008 R2 guest virtual machine cluster nodes, **VMSQLNODE1** or **VMSQLNODE2**, to perform installation on Microsoft SQL Server 2012. This document uses the **VMSQLNODE1** cluster node.

   a. Copy the Microsoft SQL Server 2012 binaries on the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1** cluster node for installation.

   b. Log in to VMSQLNODE1 with Admin credentials for installing Microsoft SQL Server 2012 software, launch the Microsoft SQL Server installation .exe file, and choose **New SQL Server failover cluster installation** as shown in Figure 64.

Figure 64. Launch Microsoft SQL Server in Microsoft SQL Server Installation Center

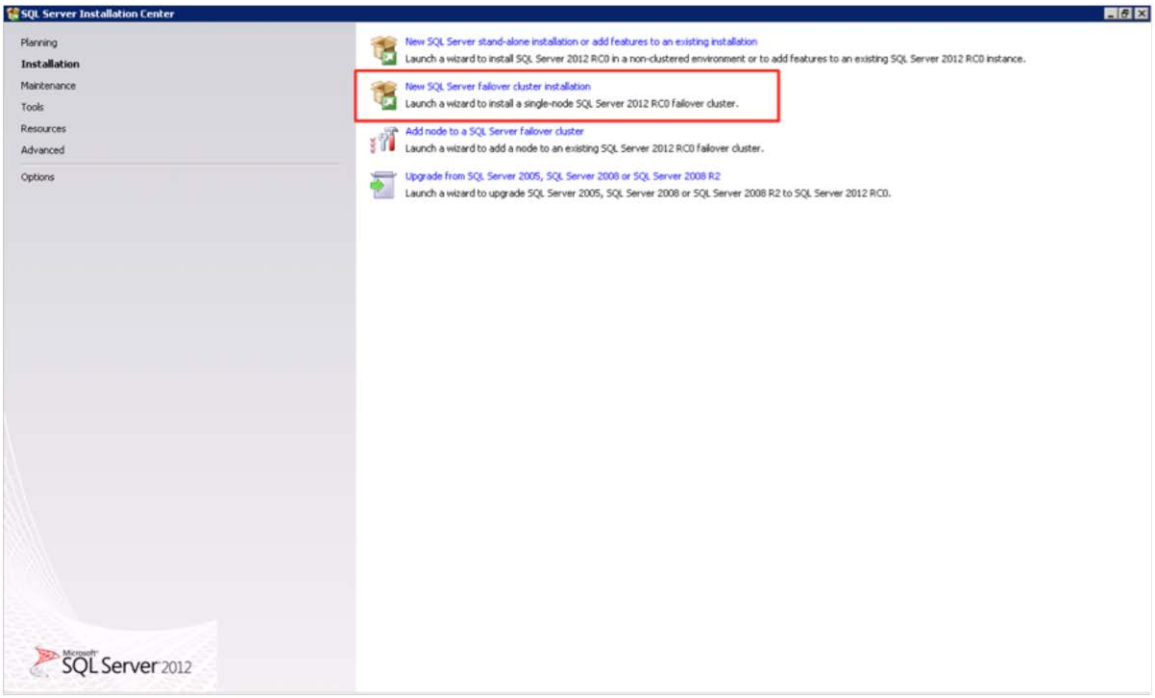

c. Follow the installation steps in the wizard and provide license keys. Choose the Feature Selection option in the wizard and select appropriate features based on your requirements. For the Microsoft SQL Server binaries shared feature directory, provide the appropriate installation directory on the **VMSQLNODE1** cluster virtual machine, as shown in Figure 65.

Figure 65. Selecting Evaluation Features to Install

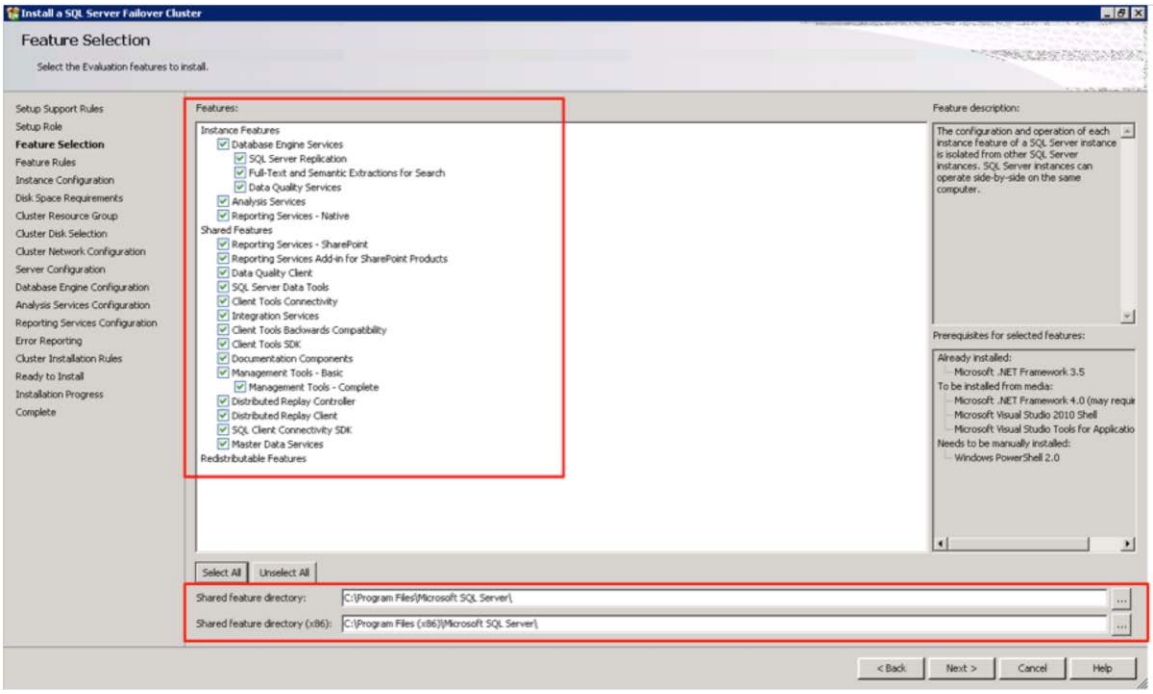

d. In the Instance Configuration window of the wizard, enter **MSSQL2012** as the name of Microsoft SQL Server 2012 failover cluster, as shown in Figure 66.

Figure 66. Instance Configuration in Server Failover Cluster Wizard

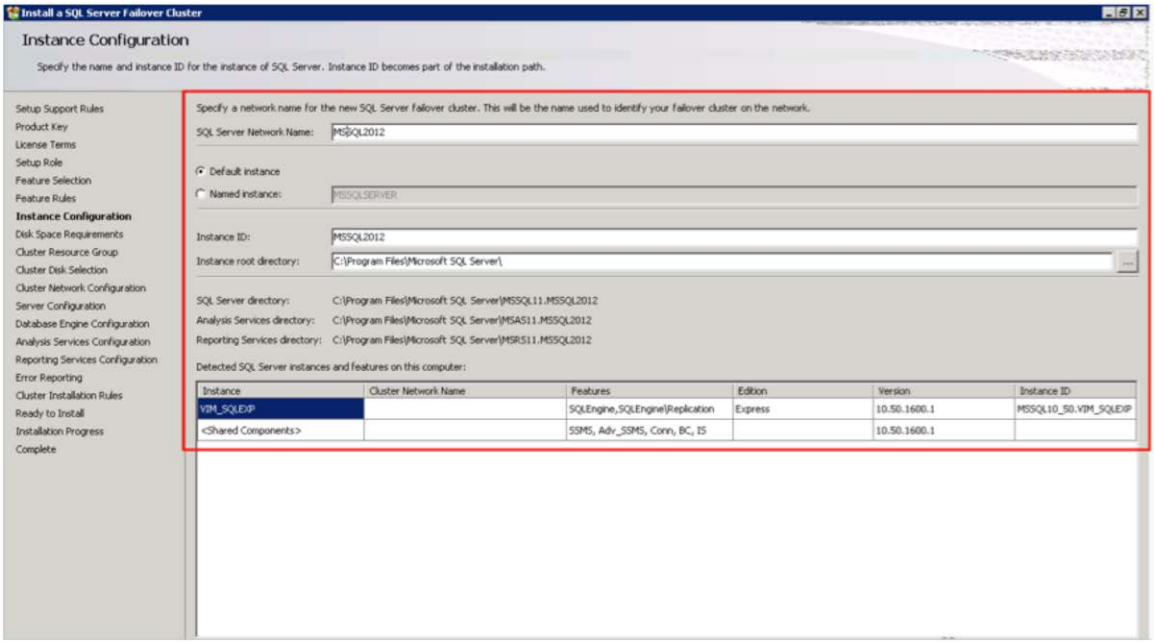

e. In the Cluster Disk Selection window, select Cluster Disk 1 and Cluster Disk 3 to store database data and log files. The Cluster Disk 2 resource is already reserved for cluster quorum storage, as shown in Figure 67.

Figure 67. Cluster Disk Selection to Store Database Data and Log Files

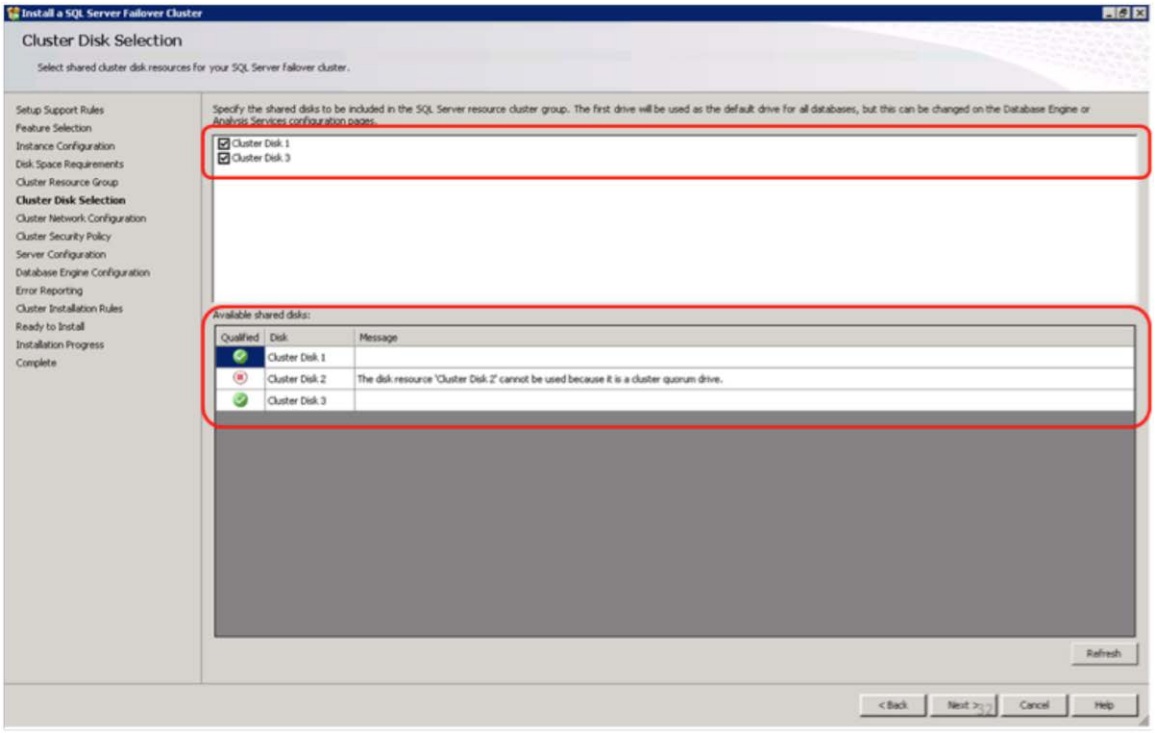

f. In the Cluster Network Configuration window, select the appropriate cluster network subnet from which the Microsoft SQL Server 2012 cluster IP address can be accessed by internal and external clients. This document uses Cluster Network 2, which is configured with the VLAN 108 management IP address, as shown in Figure 68.

Figure 68. Cluster Network Configuration in Server Failover Cluster Wizard

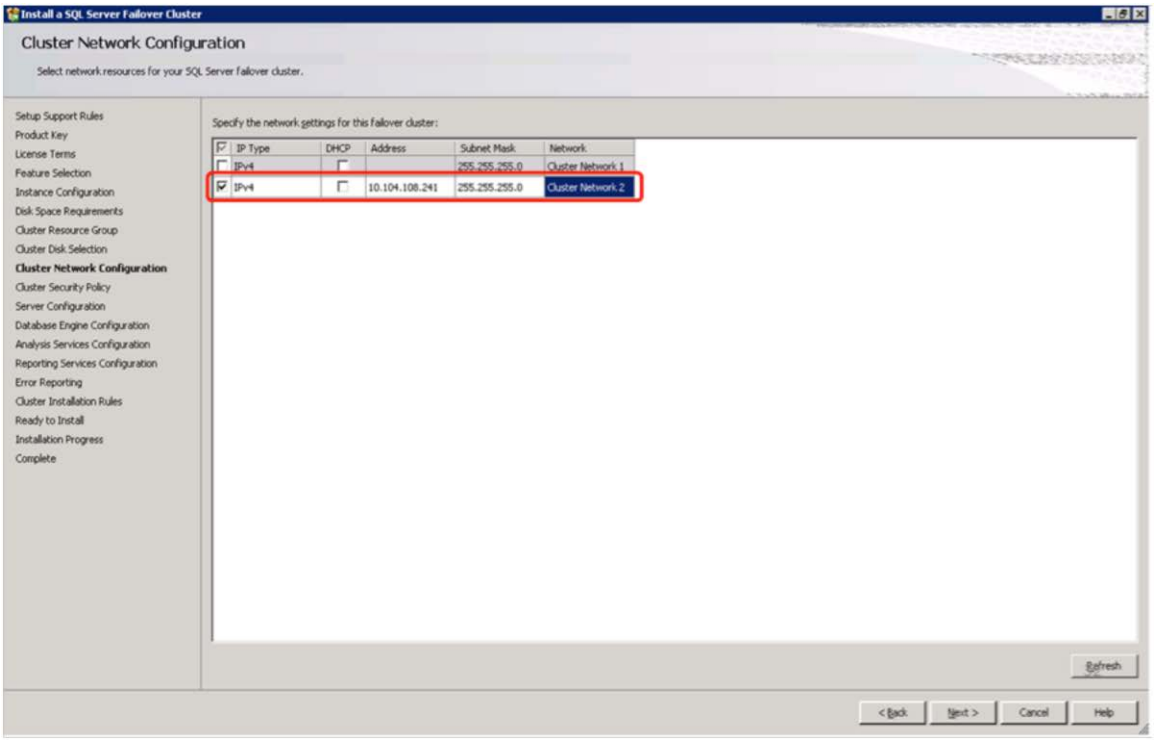

Cluster Network 1 is used to access Microsoft SQL Server 2012 database data and log file storage over the iSCSI network on the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1** and **VMSQLNODE2** cluster nodes, as shown in Figure 69.

Figure 69. Summary of Cluster Network 1 and Cluster Network 2

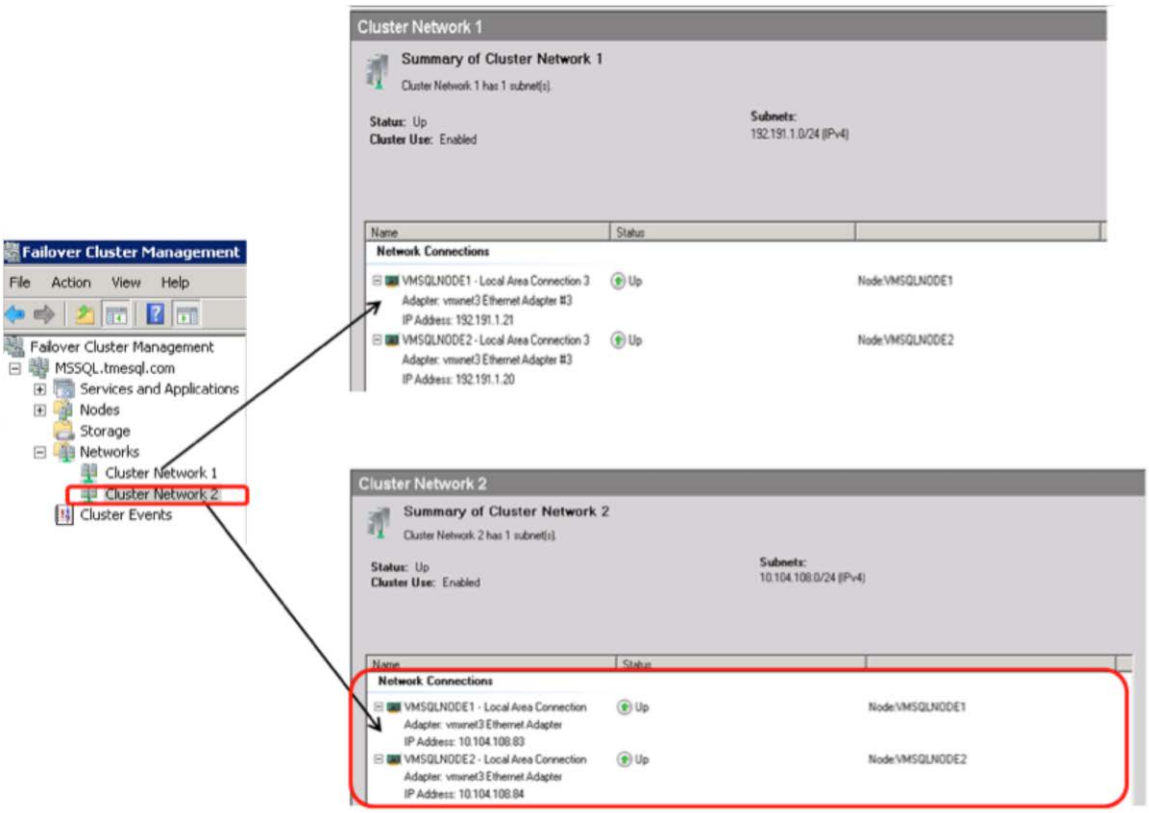

g. In the Database Engine Configuration window, select the appropriate storage drive for storing Microsoft SQL Server 2012 user database data and log files. In this setup, the database data directory is set to E:\MSSQL_DATA, which is mapped to the disk created on NetApp iSCSI storage LUN **MS_SQL_DB_LUN**. The user database log directory is set to F:\MSSQL_LOG, which is mapped to the disk created on NetApp iSCSI storage LUN **MS_SQL_LOG_LUN**, as shown in Figure 70.

Figure 70. Database Engine Configuration Showing Database Directory and Log Directory

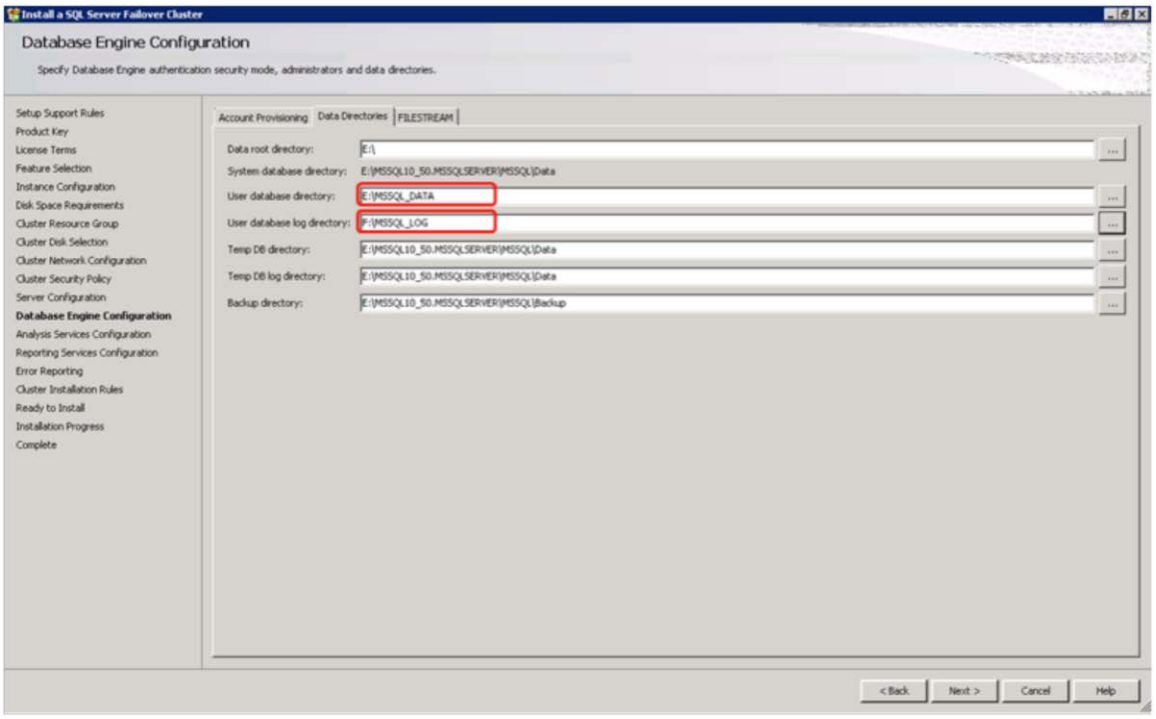

Figure 71 shows the two disks used for database data (E:) and log (F:) files, which are mapped to NetApp iSCSI storage LUNs **MS_SQL_DB_LUN** and **MS_SQL_LOG_LUN**, respectively.

Figure 71. Disks for Database Data and Log Files Mapped to NetApp iSCSI Storage

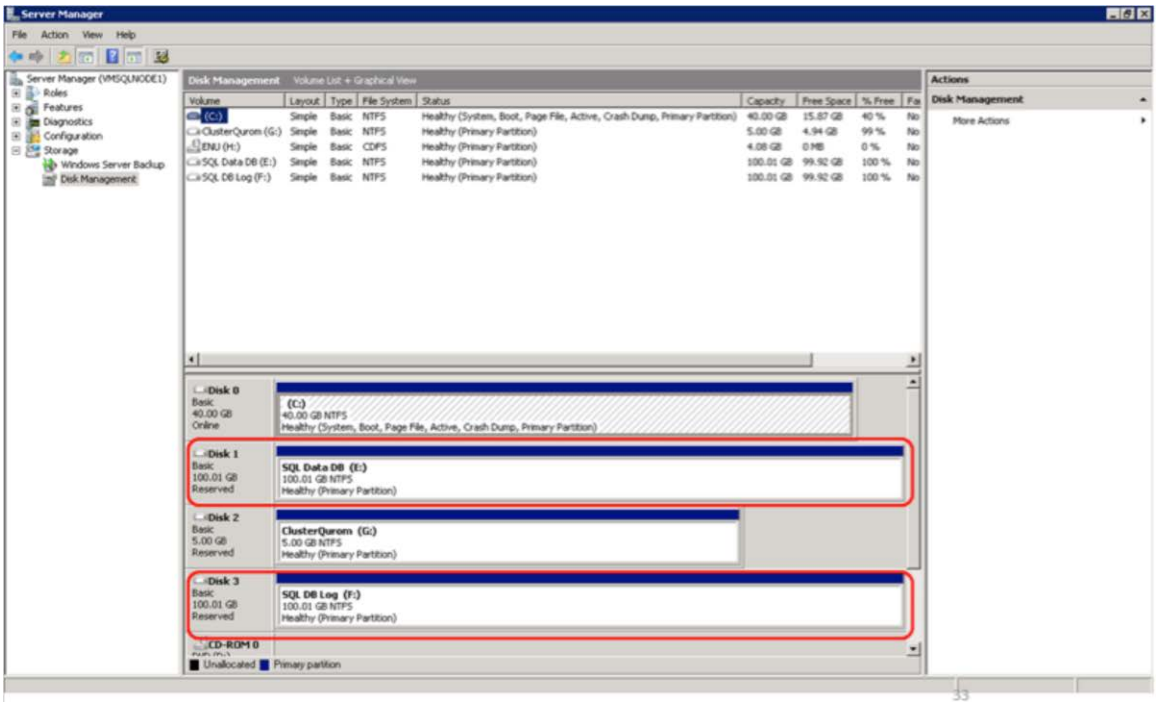

h. Figure 72 shows the summary of the configuration at the end of Microsoft SQL Server 2012 failover cluster installation setup.

Figure 72. Summary of Configuration Details at the End of Failover Cluster Installation Setup

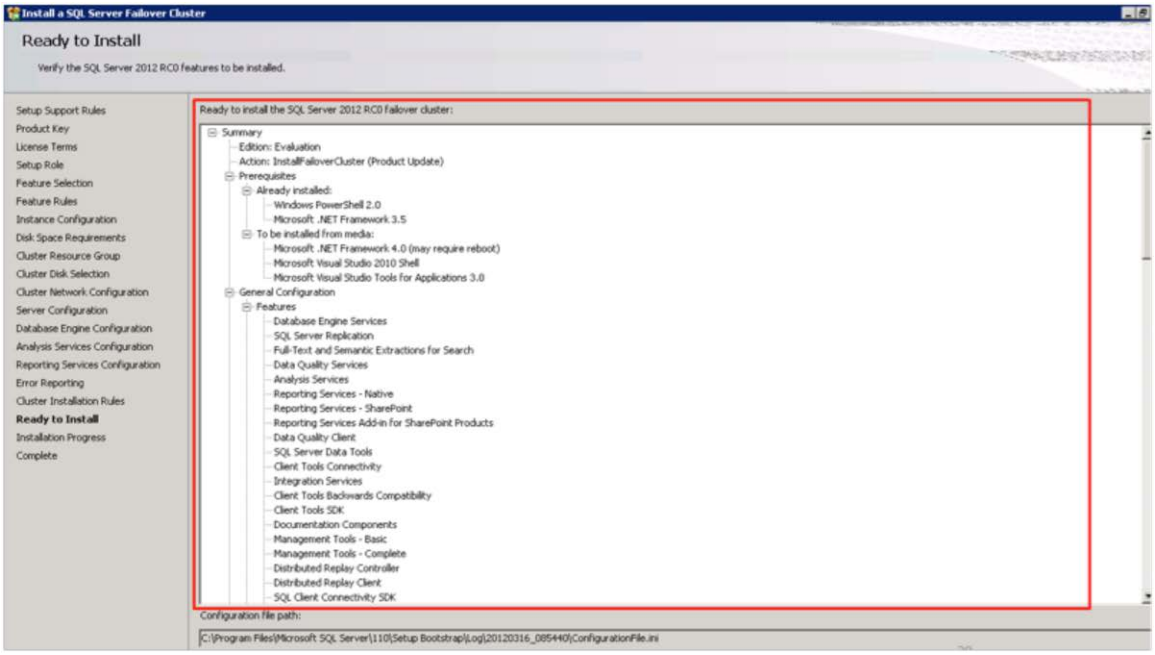

i. Figure 73 shows completion of Microsoft SQL Server 2012 failover cluster installation on the virtual machine **VMSQLNODE1** cluster node.

Figure 73. Completion of Microsoft SQL Server 2012 Failover Cluster Installation on VMSQLNODE1

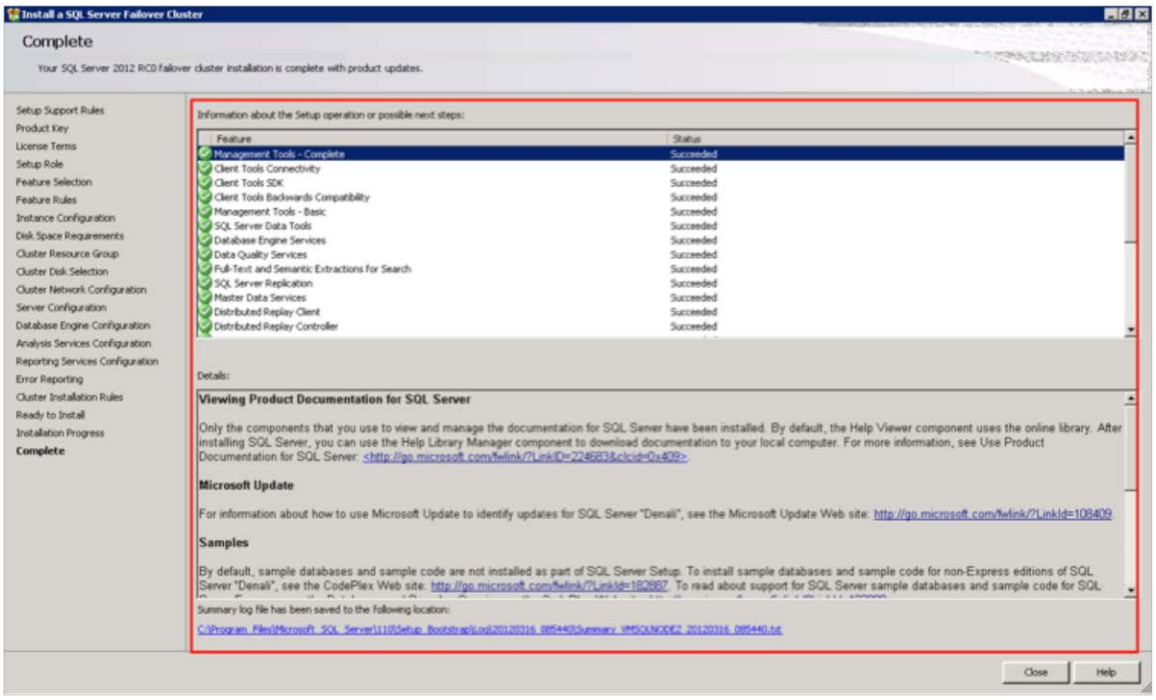

j. To verify that Microsoft SQL Server 2012 failover cluster installation succeeded on the **VMSQLNODE1** cluster node, launch the Failover Cluster Management console and choose Services and Applications and verify that the **MSSQL2012** instance is added and that the cluster IP address and storage cluster disk status are listed as Online, as shown in Figure 74.

Figure 74. Verifying the Storage Cluster Disk Status

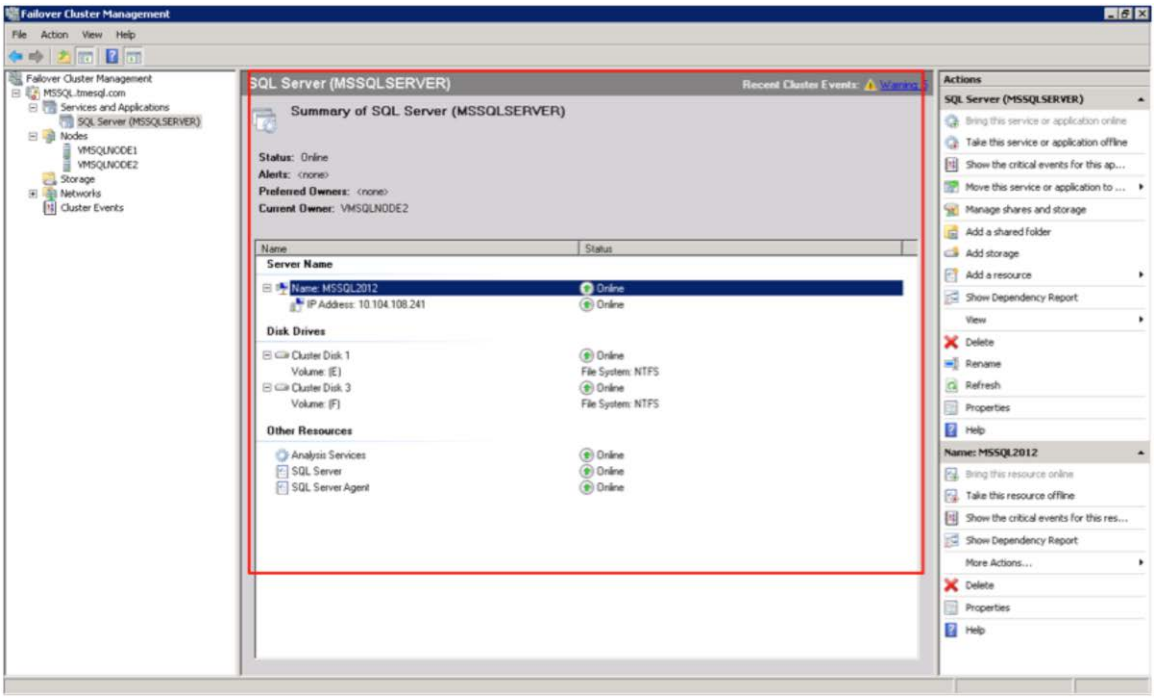

k. To verify that the Microsoft SQL Server 2012 failover cluster instance is accessible on the **VMSQLNODE1** node, launch Microsoft SQL Server Management Studio and connect to the **MSSQL2012** instance and check the start status, as shown in Figure 75.

Figure 75. Verification of Microsoft SQL Server 2012 Failover Cluster Accessibility on VMSQLNODE1

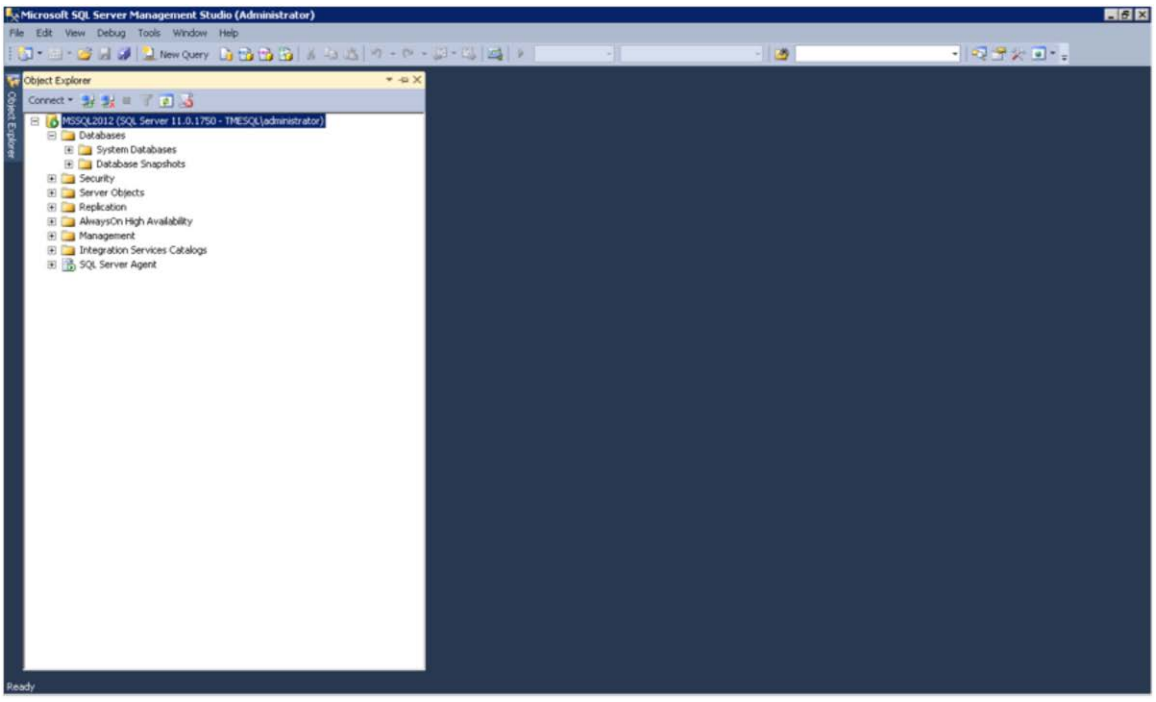

2. Perform the following steps to install Microsoft SQL Server 2012 failover clustering on Microsoft Windows 2008 R2 guest virtual machine cluster **VMSQLNODE2** as a secondary node to join the primary failover cluster node installed on **VMSQLNODE1**.

   a. Copy the Microsoft SQL Server 2012 binaries on the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE2** cluster node for installation.

   b. Log in to **VMSQLNODE2** with Admin credentials for installing Microsoft SQL Server 2012 software. Launch the Microsoft SQL Server installation .exe file and choose the option **Add node to a Microsoft SQL Server failover cluster**, as shown in Figure 76.

Figure 76. Add Node to Microsoft SQL Server Failover Cluster in SQL Server Installation Center Wizard

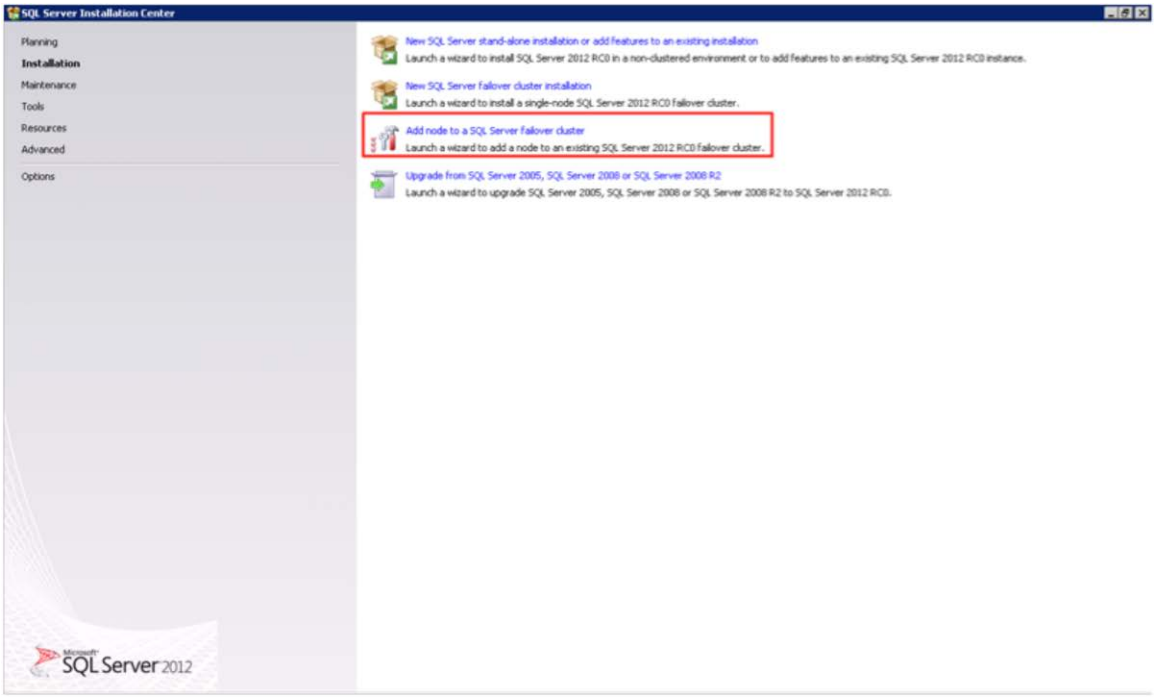

   c. In the Add a Failover Cluster Node window, below the **SQL Server instance name** field, select the **MSSQL2012** instance, which was created during first step of Microsoft SQL Server 2012 failover cluster deployment, as shown in Figure 77.

Figure 77. Cluster Node Configuration in Failover Cluster Node Wizard

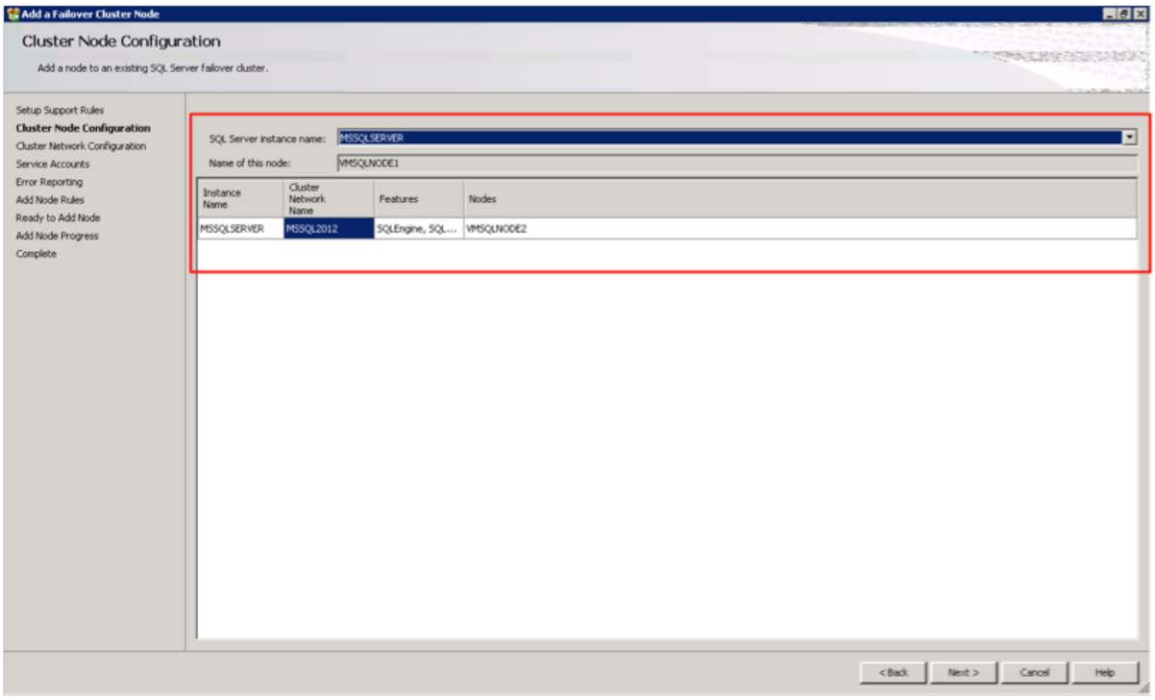

d. In the Cluster Network Configuration window of the wizard, **Cluster Network 2** is automatically selected as the IP address configured during Microsoft SQL Server 2012 failover cluster installation on **VMSQLNODE1**, as shown in Figure 78.

Figure 78. IP Address of Cluster Network 2

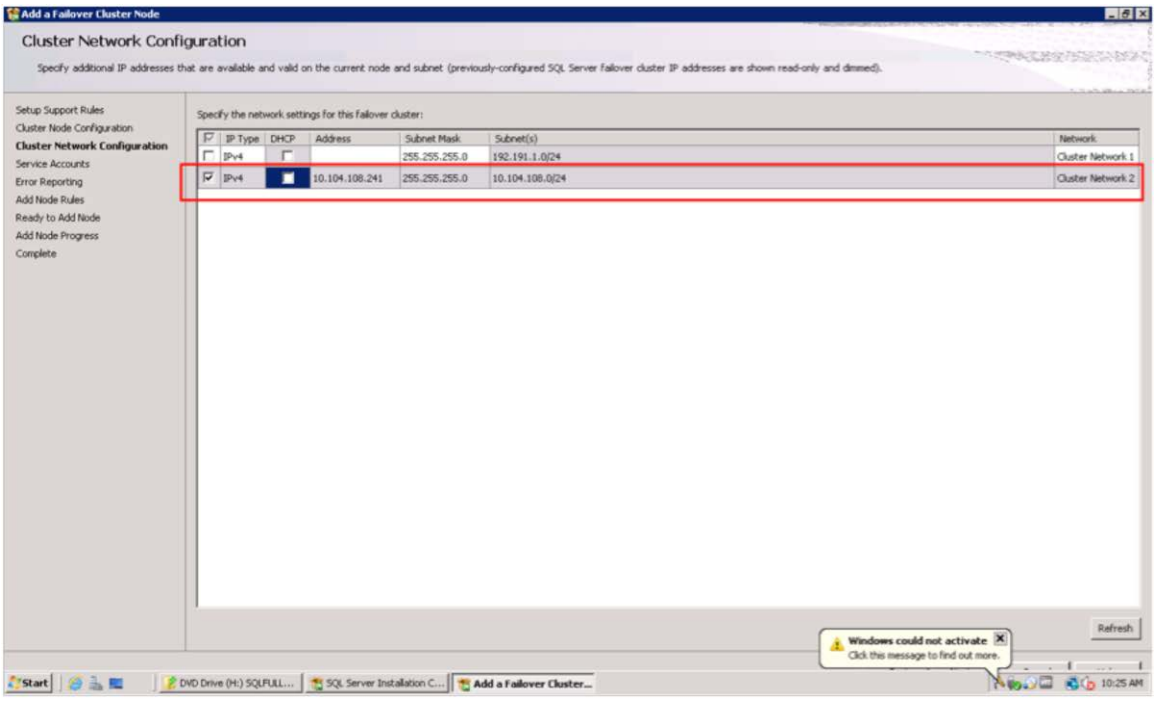

e. Figure 79 shows the summary of the configuration at the end of the installation process.

Figure 79. Summary of Cluster Network Configuration

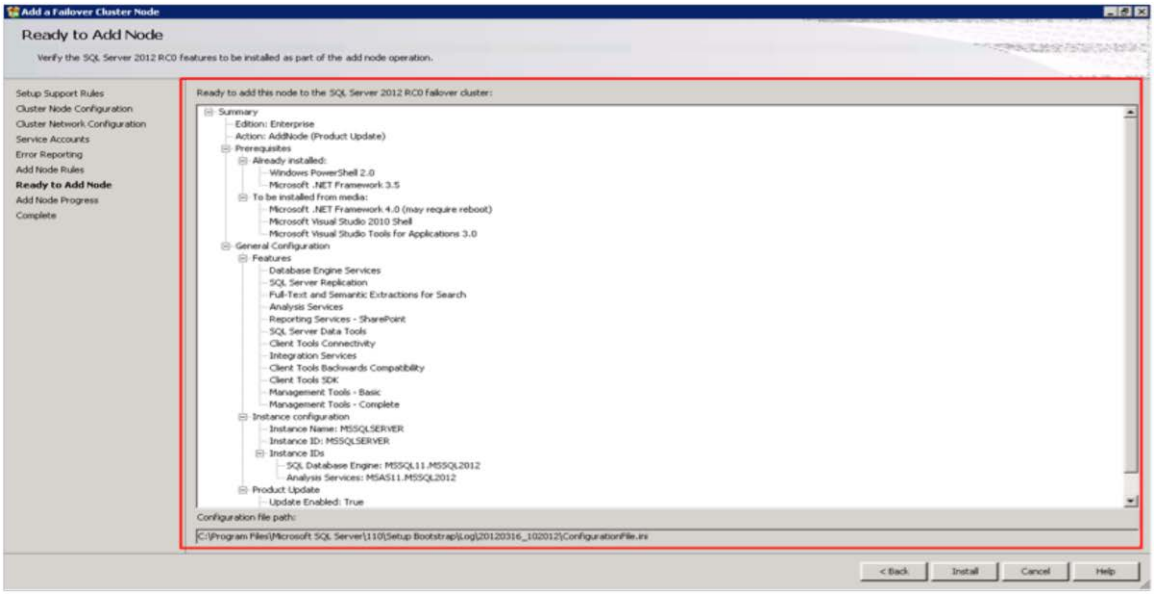

f. Figure 80 shows the completion of Microsoft SQL Server 2012 failover cluster installation on the virtual machine **VMSQLNODE2** cluster node.

Figure 80. Completion of Microsoft SQL Server 2012 Failover Cluster Installation on VMSQLNODE2

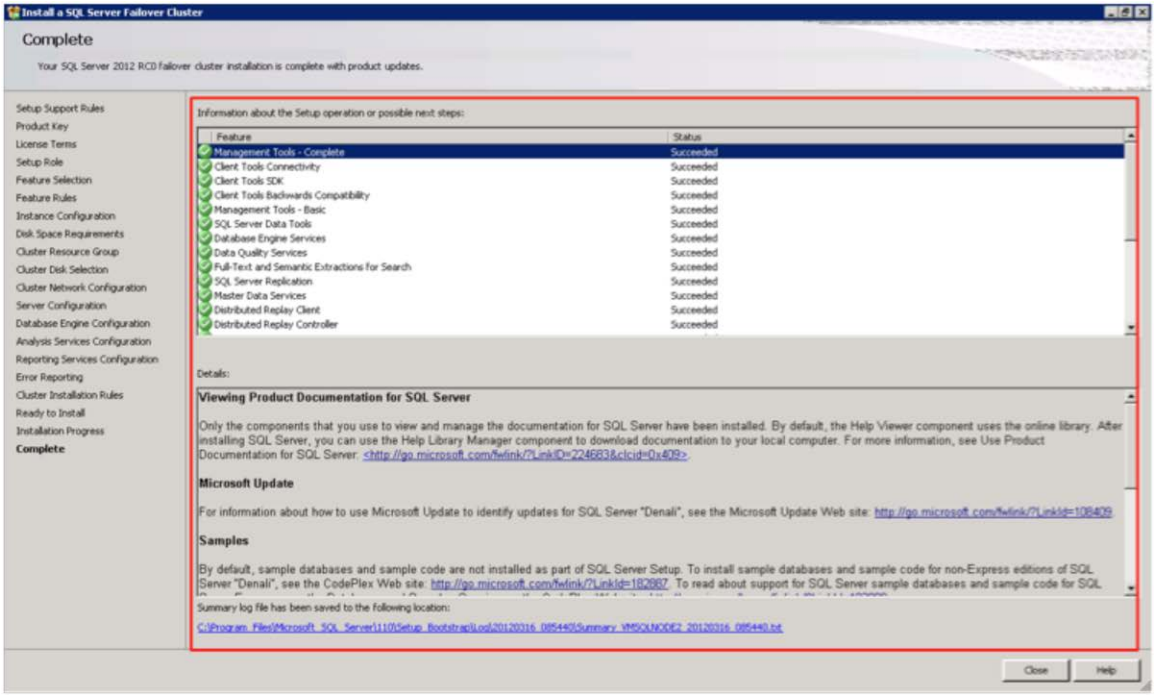

g. Verify that nodes are added to the Microsoft SQL Server 2012 failover cluster on the Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE2** node. Launch the Failover Cluster Management console, and under Services and Applications verify that the **MSSQL2012** instance has been added, that the cluster IP address and storage cluster disks status are listed as Online, and that under Nodes, both **VMSQLNODE1** and **VMSQLNODE2** are listed, as shown in Figure 81.

Figure 81. Verify the Storage Cluster Disk Status and That Both Nodes Are Listed

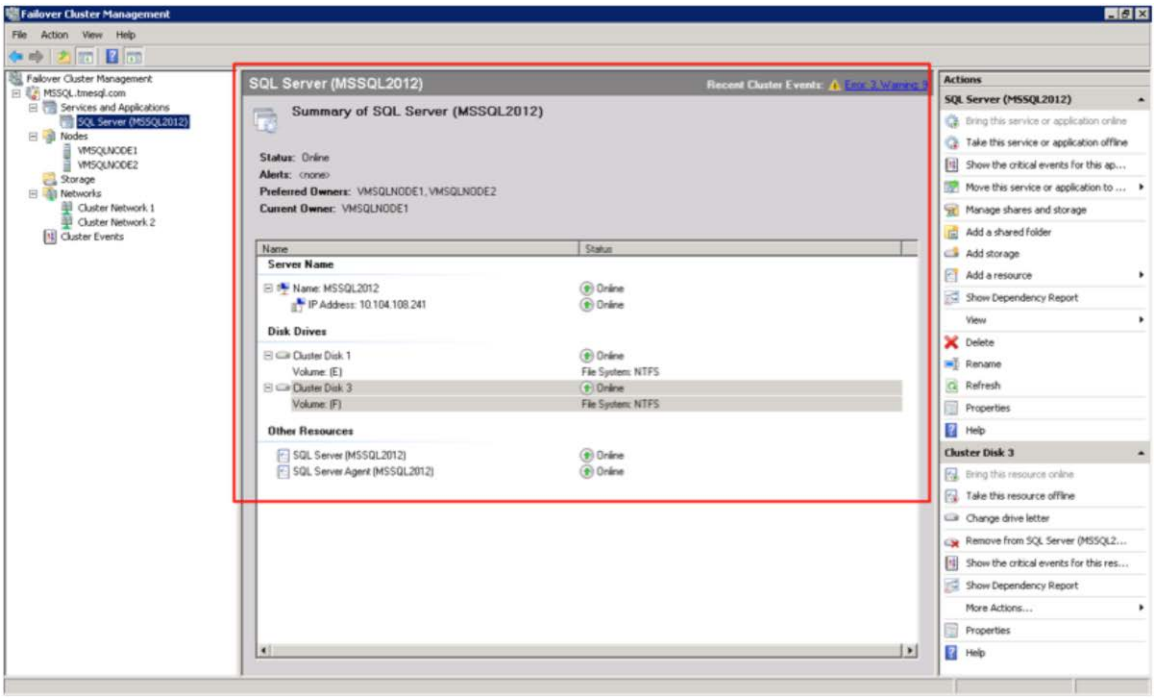

3. Perform the following steps to test failover of Microsoft SQL Server 2012 with the **VMSQLNODE1** and **VMSQLNODE2** cluster nodes.

   a. Log in to the **VMSQLNODE1** cluster node, which currently manages and owns the **MSSQL2012 SQL** instance, as shown in Figure 82, and then shut down the node.

Figure 82. VMSQLNODE1 Managing and Owning the Microsoft SQL Server 2012 Instance

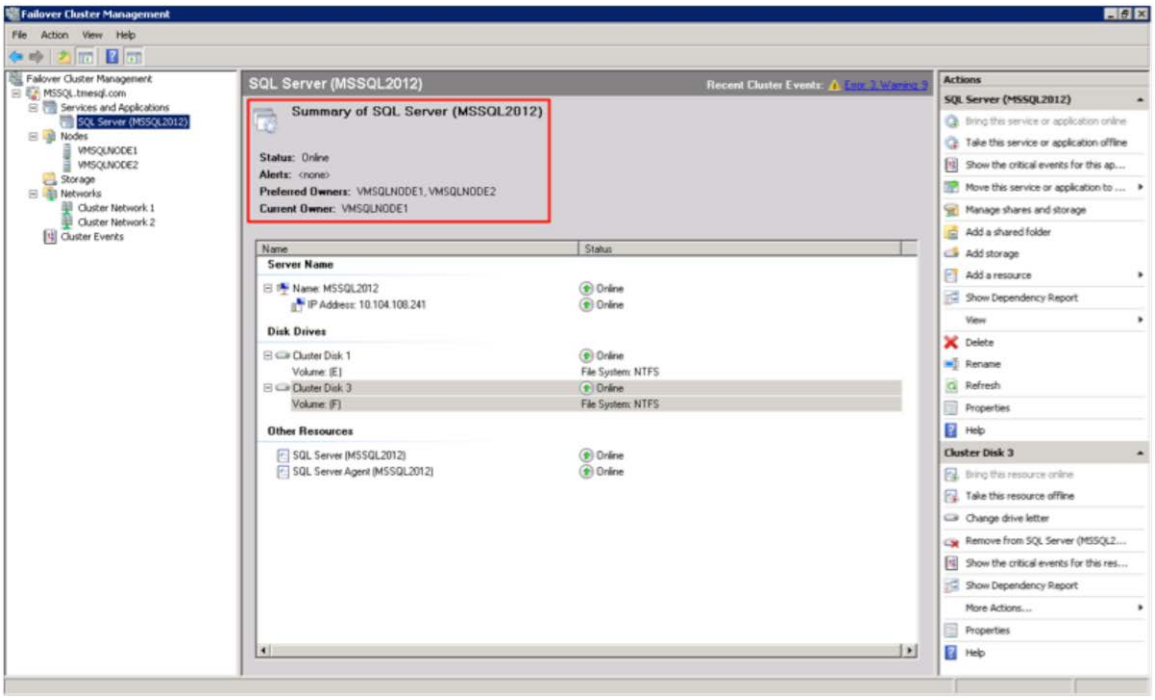

b. After the shutdown of Microsoft Windows 2008 R2 guest virtual machine **VMSQLNODE1**, the Microsoft Windows failover cluster will be triggered automatically. Subsequently, the Microsoft SQL Server 2012 cluster resource will failover to secondary node **VMSQLNODE2**. After the failover, **VMSQLNODE2** will become the current owner of the Microsoft SQL Server 2012 cluster, as shown in Figure 83.

Figure 83. Owner of Microsoft SQL Server 2012 Cluster Is VMSQLNODE2 After Failover

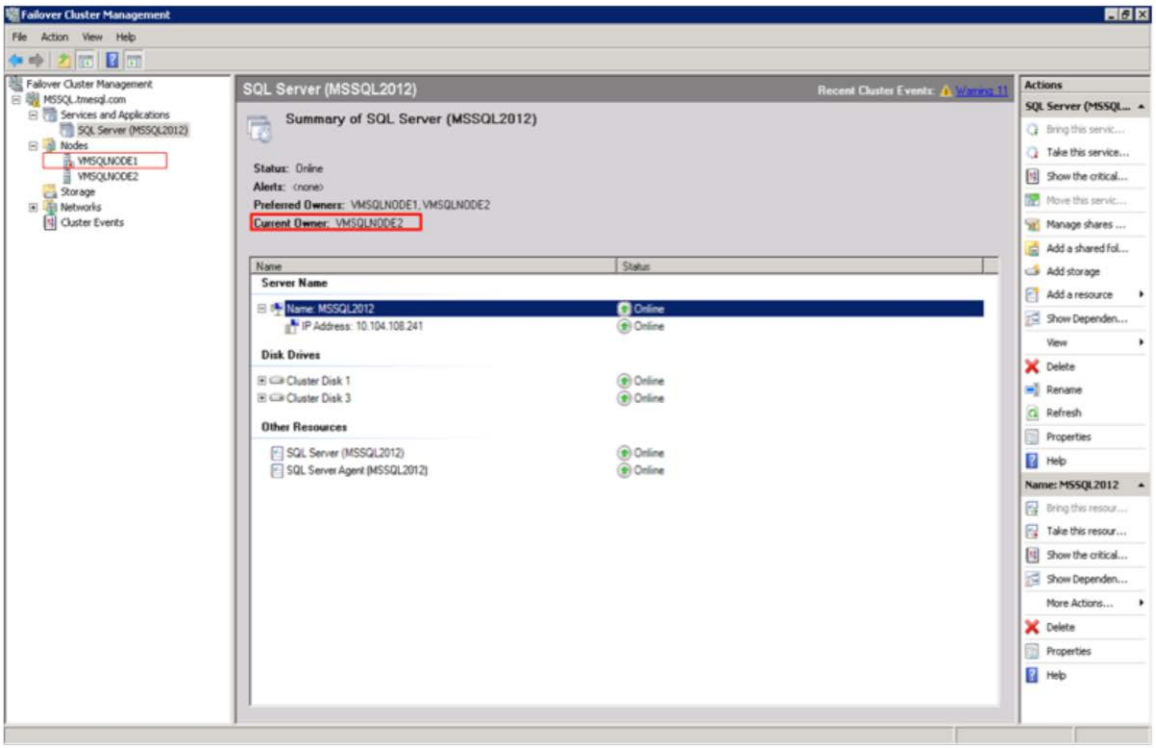

Conclusion

Windows Server 2008 R2 Enterprise x64, SQL Server 2012 x64 and, Vmware ESX 5.0 introduce many new features and enhancements. This Guide has documented best practices to get best performance and reliability for deploying MS SQL 2012 single and failover cluster database server on VMWare ESX 5.0 Hypervisior using Virtual Machince Guest based and Host based iSCSI Initiator on NetApp iSCSI Storage with VM-FEX and native Virtual Switch on UCS B-Series Sytem products.

References

Documents listed here provide additional information relevant to implementing Microsoft SQLServer 2012 on Vmware ESX 5.0 Hypervisior with NetApp iSCSI Storage System on Cisco UCS B-Series System.

- Microsoft SQL Server 2012 Installtion Guide: http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx
- Cisco Nexus QoS Switch Configuration Guide: http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/qos/Cisco_Nexus_5000_Series_NX-OS_Quality_of_Service_Configuration_Guide_chapter3.html - con_1150612
- Cisco VM-FEX Configuration Guide: http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/vm_fex/vmware/gui/config_guide/b_GUI_VMware_VM-FEX_UCSM_Configuration_Guide.pdf
- Cisco VM-FEX Best Practice and Troubleshooting Guide: http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/vm_fex_best_practices_deployment_guide_ps10277_Products_White_Paper.html

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

- Cisco UCS System Hardware and Software Interoperability Matrix: http://www.cisco.com/en/US/docs/unified_computing/ucs/interoperability/matrix/r_hcl_B_rel2_0.pdf
- VMware vSphere Networking ESXi 5.0: http://pubs.vmware.com/vsphere-50/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-50-networking-guide.pdf

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C07-707705-01 06/12