# Microsoft SQL Server 2012 Failover Cluster on Bare-Metal Microsoft Windows 2008 with Cisco UCS iSCSI-Based Storage Access: Deployment Guide

# Contents

## Executive Summary

The document describes Microsoft SQL Server 2012 failover cluster deployment in the Cisco Unified Computing System™ (Cisco UCS®) using the Small Computer System Interface over IP (iSCSI) protocol to communicate with storage devices. The document describes how to deploy a Microsoft SQL Server 2012 failover cluster on bare-metal Microsoft Windows 2008 Release 2 (R2) with Service Pack 1 (SP1) with Cisco UCS iSCSI-based storage access. The deployment scenarios discussed in this document follow the 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 easily adapt to existing networking infrastructure for storage enclosure access. 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.

High availability is one of the primary requirements for enterprise-level database platforms because mission-critical applications cannot afford 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 the database applications. Coupled with iSCSI storage at the system level, a clustering-enabled Microsoft SQL Server deployed on the Cisco UCS platform provides a complete back-end solution with optimal total cost of ownership (TCO) and high return on investment (ROI).

### Small Computer Systems Interface over IP

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 (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. The storage resources are referred to as targets. The SCSI block commands are encapsulated into Ethernet packets for transmission over IP networks at both ends of the network by the iSCSI stack.

#### Advantages of iSCSI

Following 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.

## Audience

The target audience for this guide consists of sales engineers, field consultants, professional services staff, IT managers, partner engineering staff, and customers who want to deploy Microsoft SQL Server on iSCSI.

## Hardware and Software Requirements

This section provides information about hardware and software products used in this deployment model.

### Cisco Unified Computing System Overview

Cisco UCS consists of a set of preintegrated data center components, including blade servers, adapters, fabric interconnects, and fabric extenders, that are integrated into a common embedded management system. This approach results in far fewer system components and improved manageability, greater operation efficiency, and more flexibility than other data center platforms.

#### Main Differentiating Technologies

Following are the main differentiating technologies that make Cisco UCS unique and advantageous compared to competing offerings. These technologies are discussed at a high level only, and discussion of other supporting technologies such as Fibre Channel over Ethernet (FCoE) is beyond the scope of this document.

#### Unified Fabric

Unified fabric can dramatically reduce the number of network adapters, blade-server switches, cables, and management touch points, bypassing all the network traffic to the 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 in the fabric switch. Cisco UCS Manager controls all the 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. Significant reduction in the number of management modules and consoles and in the number of agents resident on all the hardware (which must be separately managed and updated) is a critical feature of Cisco UCS. Cisco UCS Manager, using role-based access and visibility, helps enable cross-function communication efficiency, promoting collaboration among data center roles for increased productivity.

#### Cisco Extended Memory Technology

Significantly enhancing the available memory capacity of 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 and 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.

#### 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-based and policy-based management

using service profiles and templates. These mechanisms fully provision one or many servers and their network connectivity in just a few 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 the products of several major hypervisor vendors. It provides IT departments with the flexibility to allow administrators 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 IDs (UUIDs) 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 Virtual Machine Fabric Extender (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

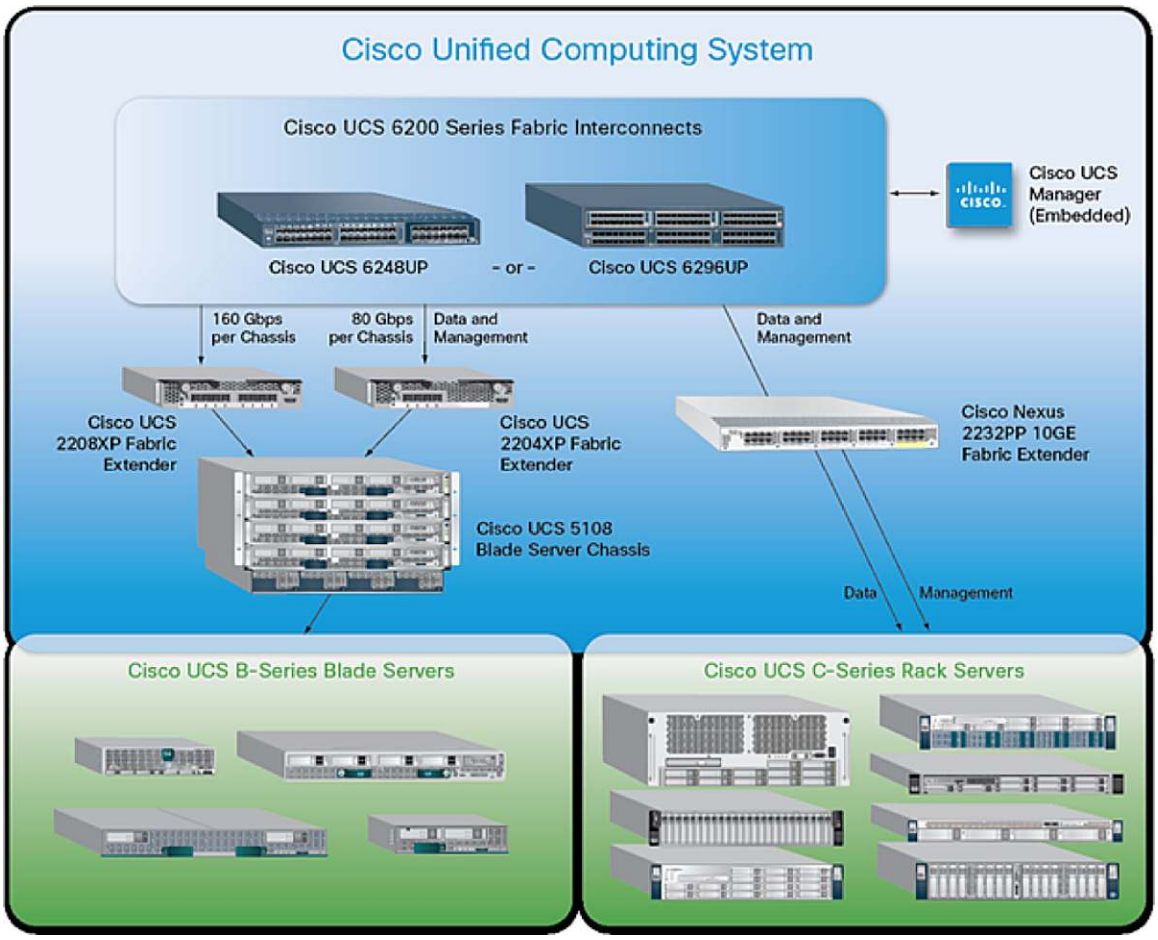


Cisco UCS is designed from the foundation to be programmable and self-integrating. A server's entire hardware stack, ranging from server firmware and settings to network profiles, is configured through model-based management. With Cisco virtual interface cards (VICs), even the number and type of I/O 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 failure due to inconsistent configurations.

Cisco Fabric Extender Technology (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 6200 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 6200 Series expands the Cisco UCS networking portfolio and offers 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 6200 Series provides the flexibility to support LAN and SAN connectivity for all blades in 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.

### 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 throughput of more than 1 terabit per second (Tbps) with low latency. It has 32 fixed 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 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 2200 Series Fabric Extenders

The Cisco UCS 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 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. It uses custom drivers for the virtualized host bus adapter (HBA) and the 10 Gigabit Ethernet network interface card (NIC). As is the case with the other Cisco CNAs, the Cisco UCS M81KR VIC encapsulates Fibre Channel traffic within 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 4-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 8-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.

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 processors p5600 series, up to 192 GB of double-data-rate-3 (DDR3) memory, two optional Small Form Factor

(SFF) SAS solid-state drives (SSDs, 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 SSDs, 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 SSDs 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 SSDs 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 four ports of dual 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.

### Microsoft Windows 2008 R2 SP1 Overview

Microsoft Windows 2008 R2 is Microsoft's multipurpose next-generation operating system designed to increase reliability and flexibility. Microsoft Windows 2008 R2 SP1 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 2008 R2 SP1 are:

- **Improved scalability and reliability:** Microsoft Windows 2008 R2 SP1 is specifically designed to support increased workloads while using fewer resources.

- **Technology improvements:** Microsoft Windows 2008 R2 SP1 includes technology improvements designed with Microsoft Windows 7 enterprise users in mind, augmenting the network experience, security, and manageability.
- **Improved management:** Microsoft Windows 2008 R2 SP1 provides enhanced management consoles and automation for repetitive day-to-day administrative tasks.
- **Improved web application platform:** Microsoft Windows 2008 R2 SP1 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 an average of 10 times the query performance of prior versions of Microsoft SQL Server
- Support for the Microsoft Windows server core to enable greater 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 Models on Cisco UCS

This document describes two Microsoft SQL Server deployment models:

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

### Microsoft SQL Server Single-Host 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 with the running application can potentially occur because system resources (mainly CPU, memory, and I/O) are shared, although tools such as the CPU affinity mask and maximum 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.

The single-host deployment model is shown in Figure 2.

**Figure 2.** Microsoft SQL Server Single-Host Deployment Model

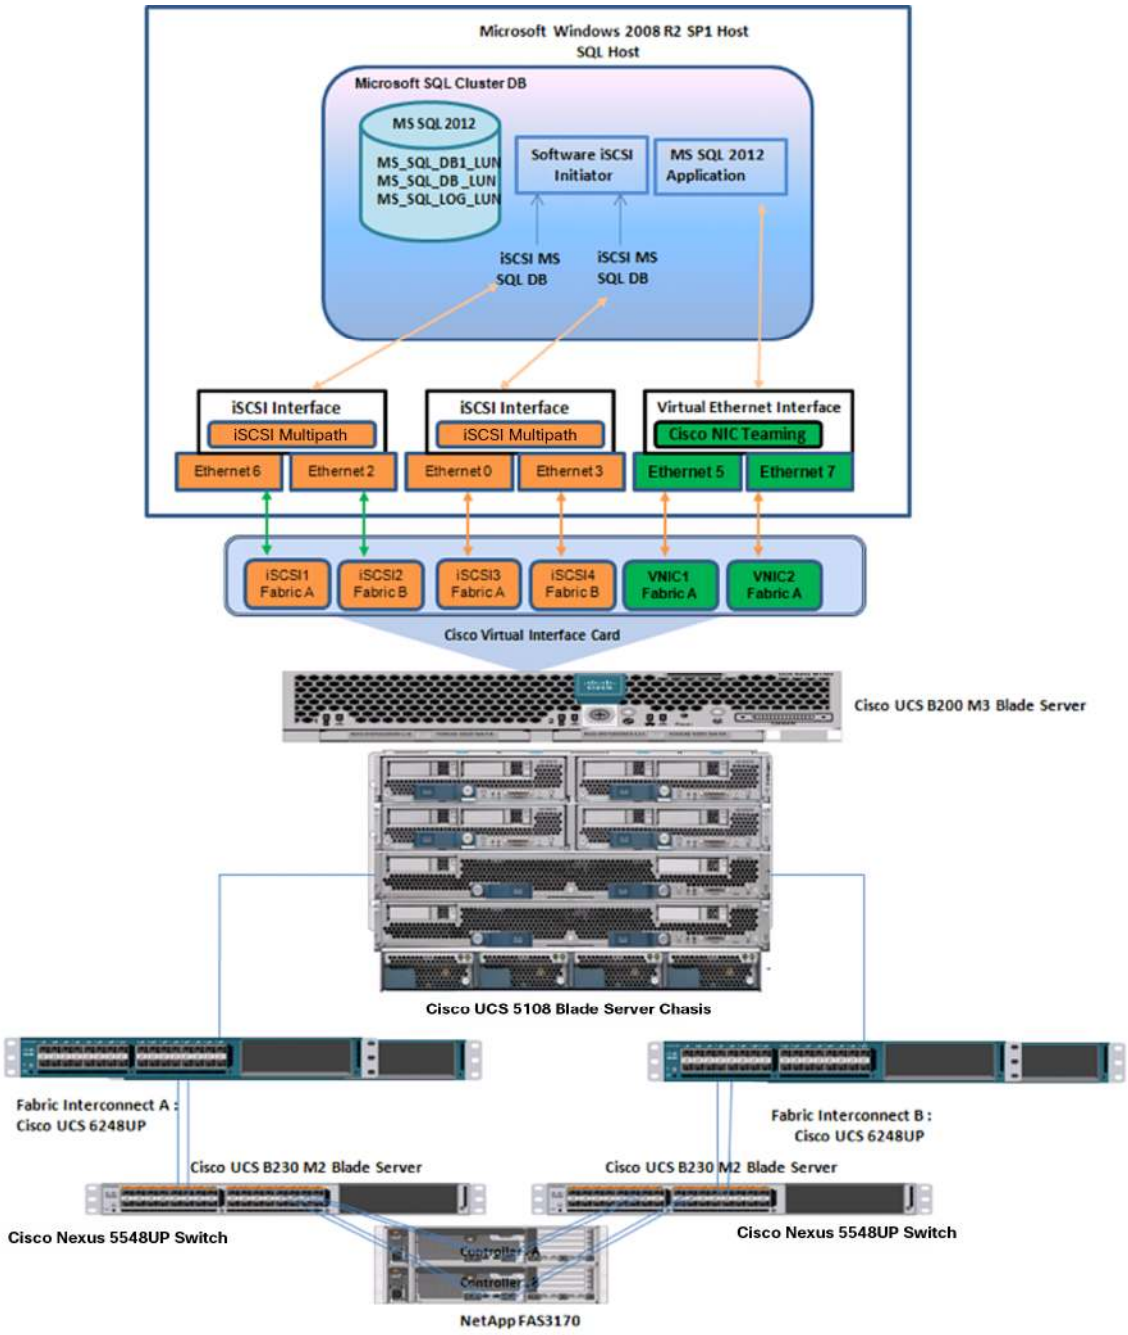


Microsoft SQL Server Failover Cluster Deployment Model

The Microsoft SQL Server failover 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 two or more nodes (or servers), which have one or more shared disks grouped into logical units called resource groups. 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 failover cluster deployment model is shown in Figure 3. Two nodes that are members of the Microsoft Windows 2008 R2 SP1 failover cluster service are deployed on Microsoft Windows 2008 R2 SP1 server hosts on two separate Cisco UCS blades. Both Microsoft Windows 2008 R2 SP1 hosts 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 the NetApp FAS3270 storage. These LUNs are accessed through the iSCSI initiator.

This design demonstrates the flexibility of storage access through the iSCSI protocol with a host-based iSCSI initiator. For high availability, at the physical level the clusters nodes are deployed on two different chassis, and at the software level they are configured with Microsoft Windows SQL Server failover clustering.

**Figure 3.** Microsoft SQL Server Failover Cluster Deployment Model

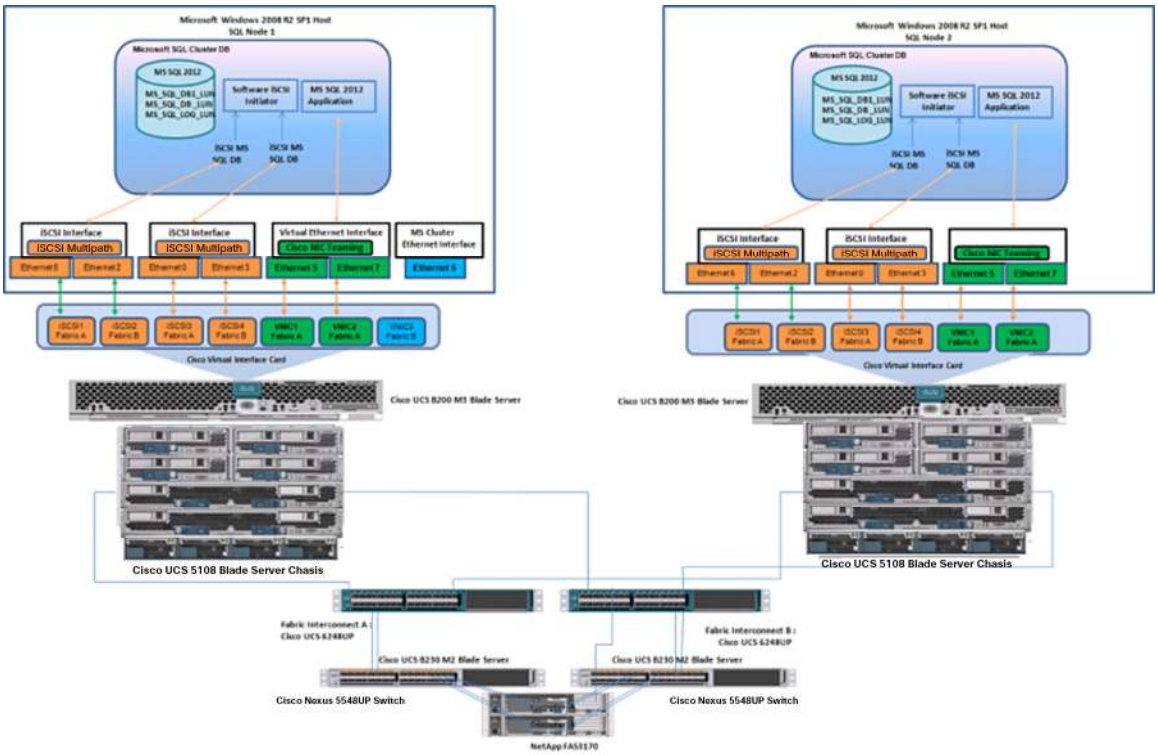


## Storage Requirements for Microsoft SQL Server Database Deployment

Storage configuration is critical to any successful database deployment. As in any physical Microsoft SQL Server deployment, the storage 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 with the maximum values to increase storage performance. RAID 1+0 provides 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 Microsoft Windows 2008 R2 SP1 Server 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). 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 base of knowledge and a single set of 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 4 shows the NetApp unified storage architecture platform.

**Figure 4.** NetApp Unified Storage Architecture

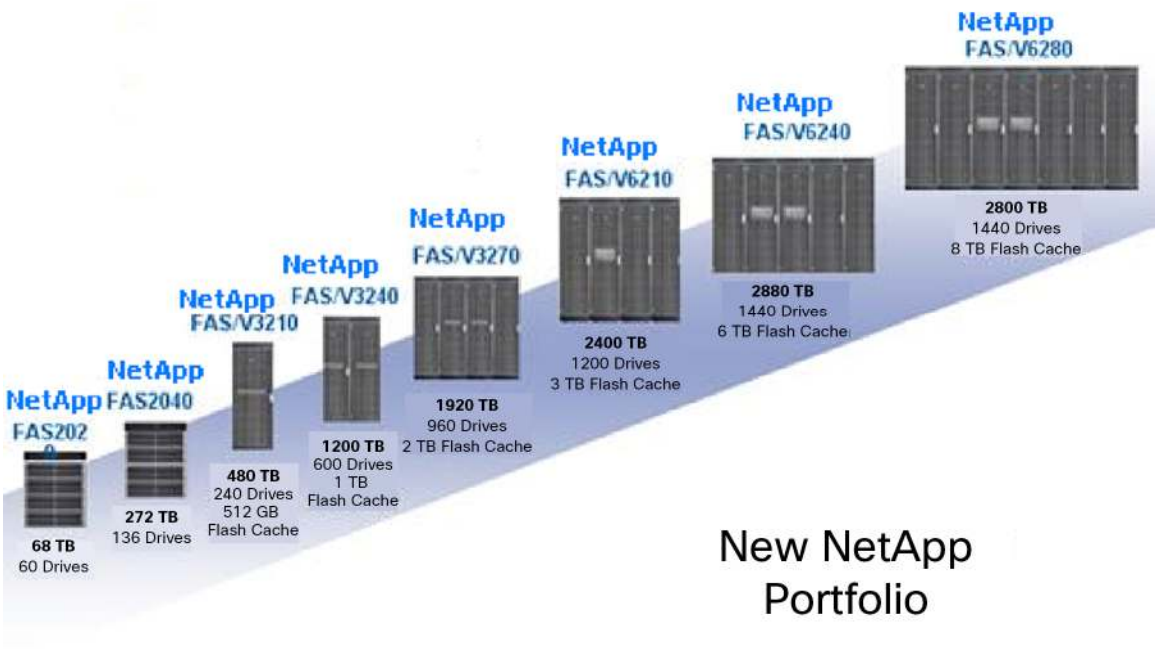


The NetApp storage hardware platform used in this solution is the NetApp FAS3270. The NetApp FAS3200 series is an excellent platform for Microsoft SQL Server 2012 deployments. NetApp storage is:

- Truly unified
- Highly efficient
- Extremely flexible

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 Microsoft Windows 2008 R2 SP1 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 Microsoft SQL Server is deployed in a Microsoft Windows 2008 R2 SP1 server 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 5 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 5.** Cisco UCS Components in iSCSI Network Design

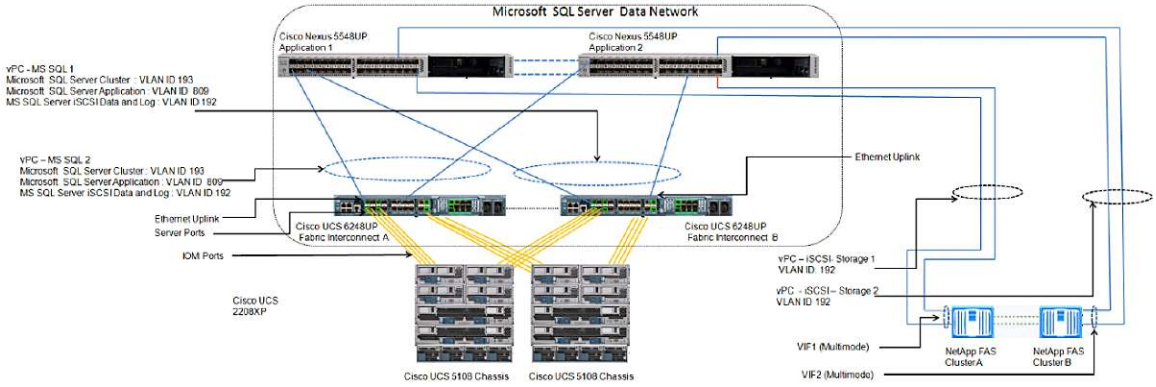


As shown in Figure 5 a pair of Cisco UCS 6248UP fabric interconnects carries both storage and network traffic from the blades with the help of a Cisco Nexus 5548UP Switch. Both the fabric interconnect and the Cisco Nexus 5548UP are clustered with the peer link between them to provide high availability. 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.

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 to Cisco UCS in the event of failure, 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 6 from the CLI on the Cisco Nexus 5548UP Switch.

**Figure 6.** PortChannel Status on Cisco Nexus 5548UP

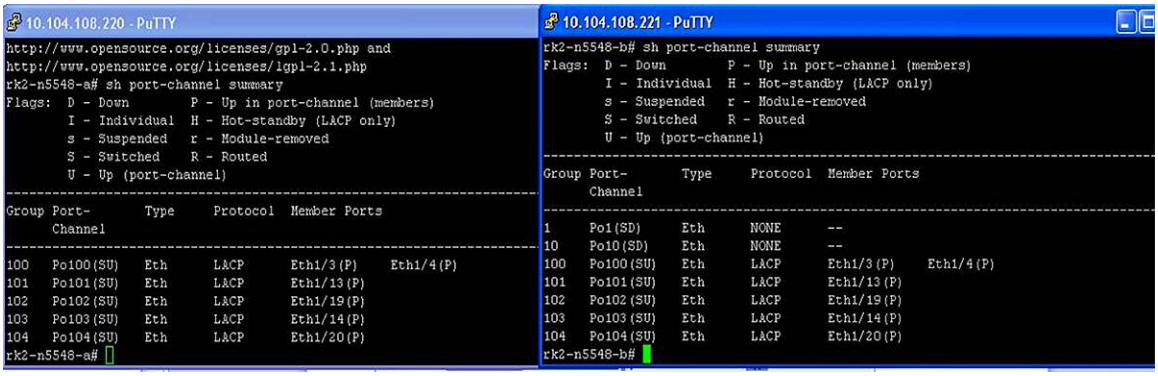


Table 2 shows the vPC configuration details for Cisco UCS 6248UP Fabric Interconnects A and B with the required vPC IDs, VLAN IDs, and Ethernet uplink ports for a Microsoft SQL Server data network design.

**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)<br>109 (SQL network)<br>192 (iSCSI storage) |
| **vPC-MS SQL 2** | 102 | Fabric Interconnect B (Eth 1/15 and 1/16) | 108 (management)<br>109 (SQL network)<br>192 (iSCSI storage) |

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 on the Cisco UCS Manager screen in Figure 7.

**Figure 7.** Uplink Interfaces and PortChannel Status

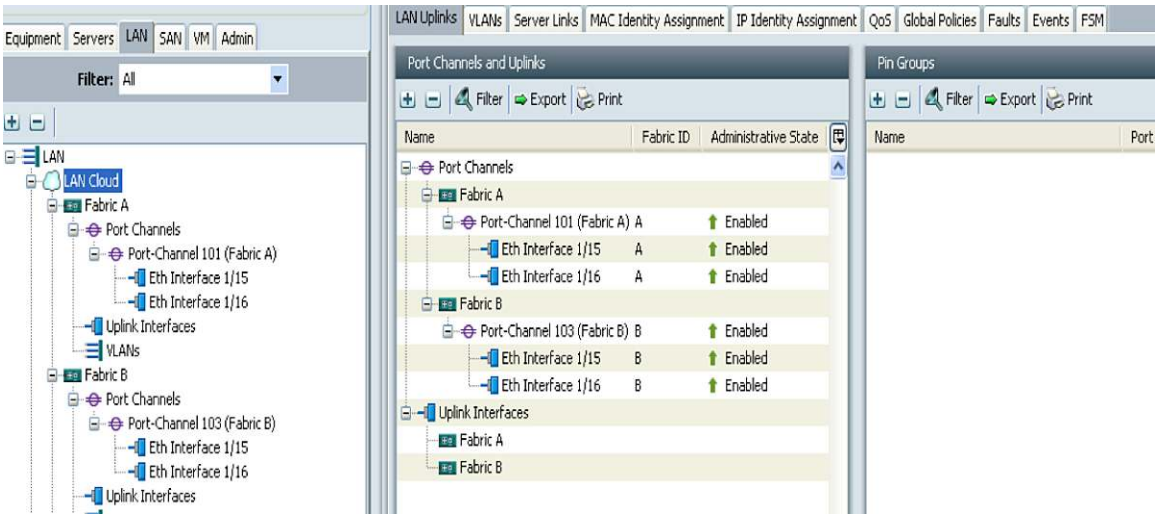


On the Cisco Nexus 5548UP Switch, a separate vPC is created to access NetApp shared iSCSI storage. The vPC is created with the vPC name and corresponding vPC ID and required VLAN IDs, as shown in Table 3.

**Table 3.** NetApp Storage

| vPC Name | iSCSI Ports (Controllers A and B) | vPC ID | VLAN ID |
|---|---|---|---|
| **vPC-iSCSI Storage 1** | e1b and e1c<br>(Controller A) | 103 | 192 |
| **vPC-iSCSI Storage 2** | e1b and e1c<br>(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 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**<br>**Gold priority**<br>**Silver priority**<br>**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 virtual NIC (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 with Microsoft Windows 2008 R2 SP1 over an iSCSI network, QoS system classes and corresponding policies are defined for network traffic generated by iSCSI storage, Microsoft failover clusters, and the Microsoft SQL Server application 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, an **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 Microsoft Windows cluster control data traffic across Microsoft SQL Server failover cluster nodes, a **MSSQLCluster** QoS class is created and defined with the third highest weight (bandwidth) on Cisco UCS.
- To handle Microsoft Windows bare-metal management and 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, an **MSSQLAPP** QoS class is created and defined with the fourth highest weight (bandwidth) 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 classes and policy types with class-of-service (CoS) values that match the Cisco UCS QoS classes.

For more information, see the Cisco Nexus QoS guide 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 |
| **MSSQLCluster** | Silver | 8 | 9000 |
| **MSSQLAPP** | Bronze | 7 | 9000 |

Figure 8 shows Cisco UCS QoS system class and QoS policy configurations defined for static and dynamic vNICs for accessing a Microsoft SQL Server iSCSI network.

**Figure 8.** Cisco UCS QoS System Class and QoS Policy Configuration Window

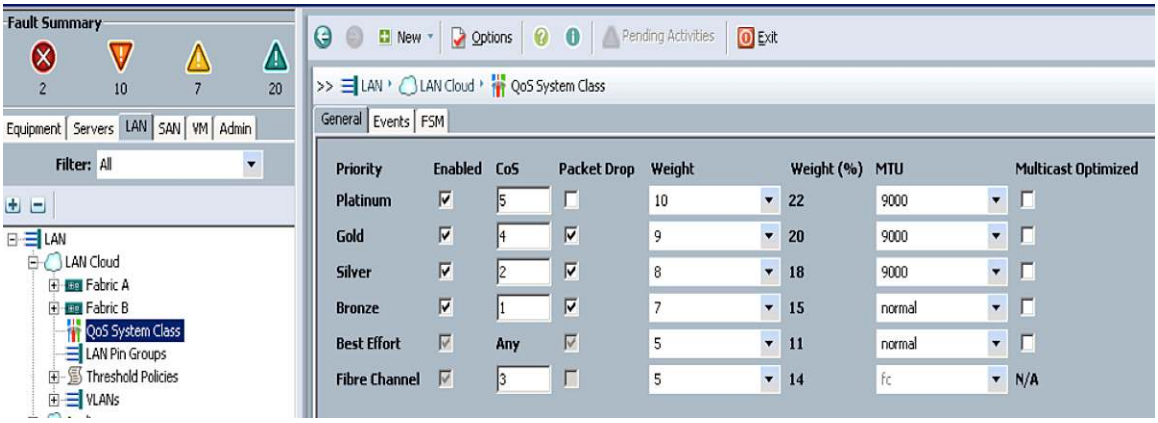


Figure 9 shows how the class priorities are applied to the named QoS policies in Cisco UCS Manager.

**Figure 9.** Applying Priority Classes to QoS Policy in Cisco UCS Manager

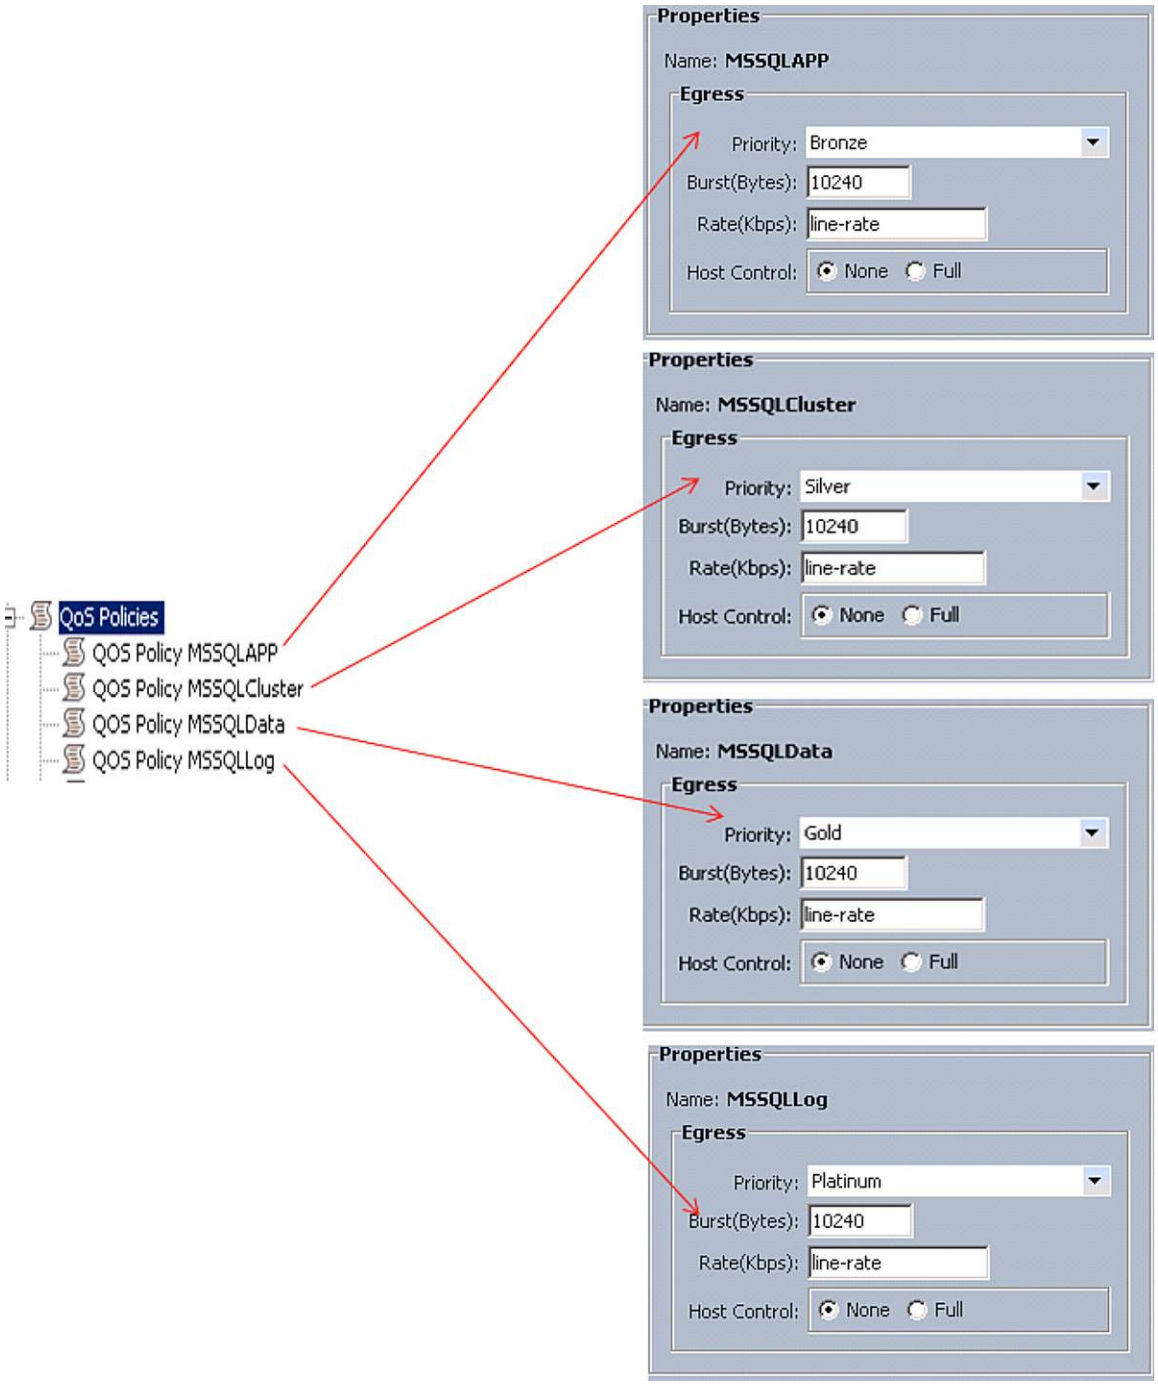


Table 6 shows Cisco UCS and Cisco Nexus 5548UP QoS mapping, with Cisco UCS QoS policy configuration values matched with Cisco Nexus 5548UP QoS policy values to achieve end-to-end QoS.

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 at the global system level.

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

| Cisco UCS QoS | | | | Weight (Percentage) | |
|---|---|---|---|---|---|
| **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<br>QoS: QoS group 5 | Cisco UCS Nexus 5548UP QoS |
| **MSSQLData** | Gold | 9000 | 4 | Network QoS: MTU 9000 and CoS 4<br>QoS: QoS group 4 | Cisco UCS Nexus 5548UP QoS |
| **MSSQLCluster** | Silver | 9000 | 2 | Network QoS: MTU 9000 and CoS 2<br>QoS: QoS group 2 | Cisco UCS Nexus 5548UP QoS |
| **MSSQLAPP** | Bronze | 9000 | 1 | Network QoS: MTU 9000 and CoS 1<br>QoS: QoS group 1 | Cisco UCS Nexus 5548UP QoS |

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 details you should consider when you deploy a Microsoft SQL Server 2012 database using Microsoft Windows 2008 SP1 R2 on Cisco UCS in an iSCSI network environment.

Figure 10 shows a high-level storage design with a NetApp FAS3270 cluster storage system.

**Figure 10.** Design Overview for a NetApp Storage Cluster

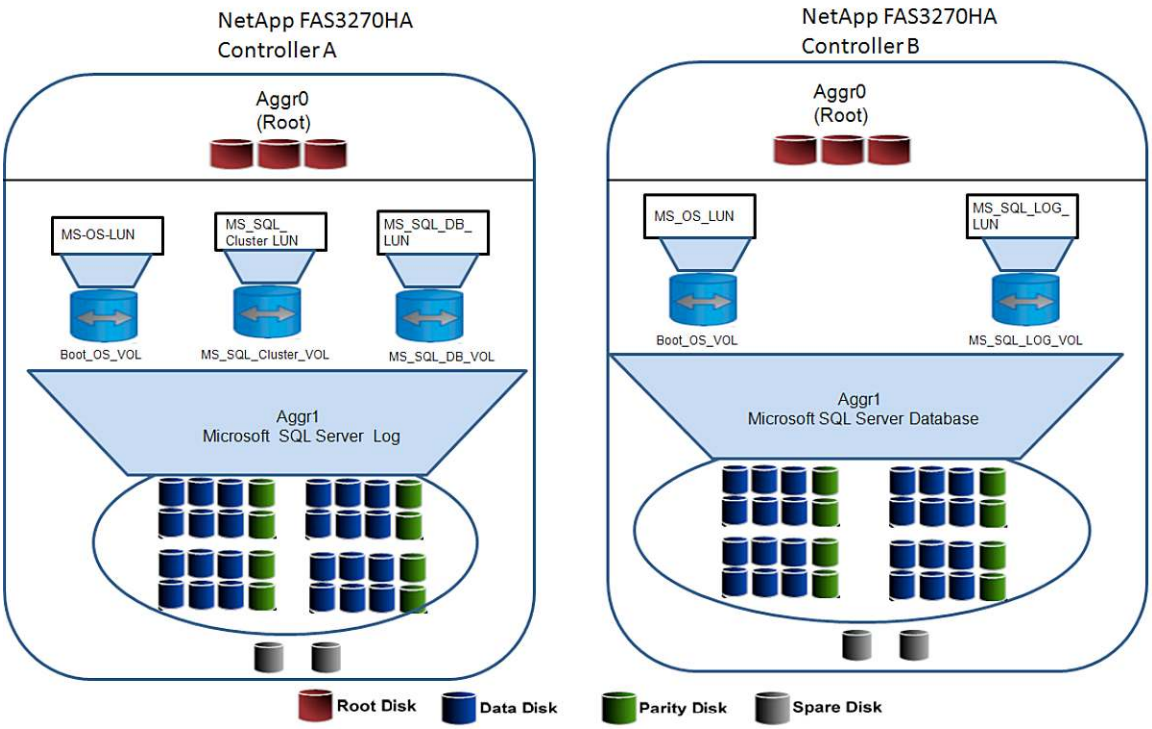

The NetApp aggregation layer provides a large virtualized pool of storage capacity and disk IOPS to be used on demand by Microsoft Windows 2008 R2 SP1 on the aggregation layer. The aggregation-layer sizing is based on the storage requirements for Microsoft SQL Server data and log files to meet the storage 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 needed 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 to handle 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 | MS_OS_LUN | iSCSI boot LUN for the Microsoft Windows host for node 1 of the failover cluster with a Cisco UCS B200 M3 blade server |
| **Aggr1 and Controller A** | MS_SQL_Cluster_VOL | MS_SQL_Cluster_LUN | LUN with Microsoft Windows–based software iSCSI initiator for storing failover cluster quorum data |
| **Aggr1 and Controller A** | MS_SQL_DB_VOL | MS_SQL_DB_LUN | LUN with Microsoft Windows–based software iSCSI initiator for storing the Microsoft SQL Server 2012 database file LUN |
| **Aggr1 and Controller B** | Boot_OS_VOL | MS_OS_LUN | iSCSI boot LUN for the Microsoft Windows host for node 2 of the failover cluster with a Cisco UCS B200 M3 blade server |
| **Aggr1 and Controller B** | MS_SQL_LOG_VOL | MS_SQL_LOG_LUN | LUN with Microsoft Windows–based software iSCSI initiator for storing the Microsoft SQL Server 2012 database log file LUN |

Use the commands in this section to configure NetApp cluster storage systems on the storage controllers to implement the storage layout design described here. To run the following commands, from the CLI log into the storage controller using Secure Shell (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 shown in Table 7. These volumes are exposed to Microsoft Windows host and Microsoft SQL Server operations.

```
FAS3270HA-Controller A> vol create Boot_OS_VOL aggr1 50g
FAS3270HA-Controller A> vol create MS_SQL_DB_VOL aggr1 150g
FAS3270HA-Controller A> vol create MS_SQL_Cluster_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 windows /vol/Boot_OS_VOL/MS_OS_LUN
```

```
FAS3270HA-Controller A> lun create –s 100g –t windows
/vol/MS_SQL_DB_VOL/MS_SQL_DB_LUN
FAS3270HA-Controller A> lun create –s 100g –t windows
/vol/MS_SQL_Cluster_VOL/MS_SQL_Cluster_LUN
```

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

```
FAS3270HA-Controller A> igroup create –I –t windows SQLNODE1-Primary iqn.2012-
01.com.windows:sqlnode1-Primary, iqn.2012-01.com.windows:sqlnode1-Secondary
```

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

```
FAS3270HA-Controller A>
lun map /vol/Boot_OS_VOL/MS_OS_LUN SQLNODE1-Primary
FAS3270HA-Controller A>
lun map /vol/MS_SQL_DB_VOL/MS_SQL_DB_LUN SQLNODE1-Primary
FAS3270HA-Controller A>
lun map /vol/MS_SQL_Cluster_VOL/MS_SQL_Cluster_LUN SQLNODE1-Primary
```

After successfully running these commands, you can verify the storage configuration using the NetApp OnCommand System Manager, as shown in Figure 11.

**Figure 11.** Verification of Storage Configuration

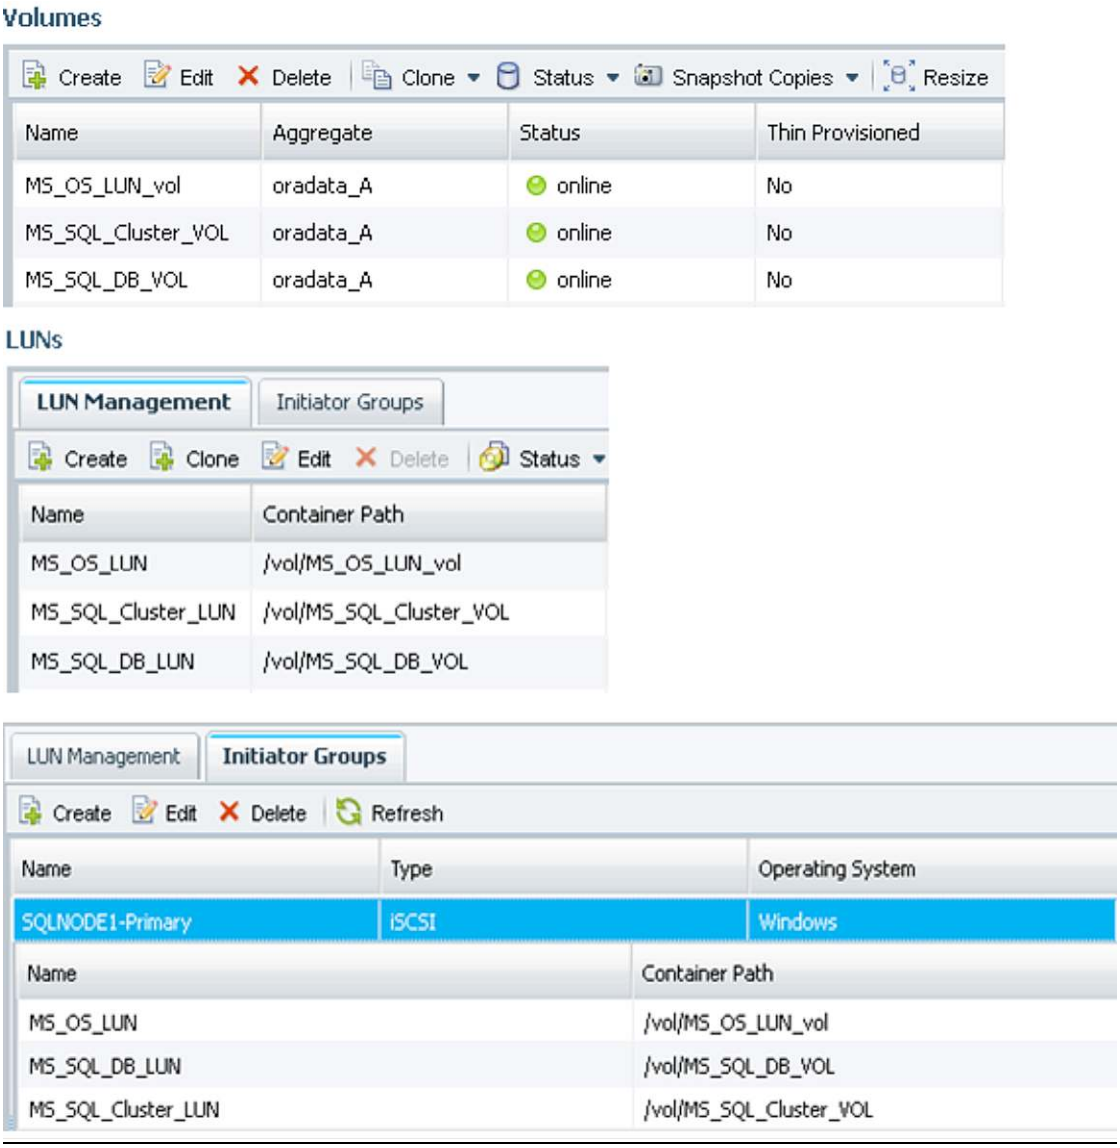


**NetApp FAS3270HA (Controller B)**

- 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 shown in Table 7. These volumes are exposed to Microsoft Windows host and Microsoft SQL Server operations.

```
FAS3270HA-Controller B> vol create Boot_OS_VOL aggr1 50g
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 windows /vol/Boot_OS_VOL/MS_OS_LUN
FAS3270HA-Controller B>
 lun create –s 5g –t v windows /vol/MS_SQL_LOG_VOL/MS_SQL_LOG_LUN
```

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

```
FAS3270HA-Controller B> igroup create –I –t windows SQLNODE2-Primary iqn.2012-
01.com.windows:sqlnode2-Primary,iqn.2012-01.com.windows:sqlnode2-Secondary
```

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

```
FAS3270HA-Controller B>
lun map /vol/Boot_OS_VOL/MS_OS_LUN SQLNODE2-Primary
FAS3270HA-Controller B>
FAS3270HA-Controller B>
lun map vol/ MS_SQL_LOG_VOL/MS_SQL_LOG_LUN SQLNODE2-Primary
```

After successfully running these commands, you can verify the storage configuration using the NetApp Filter view, as shown in Figure 12.

**Figure 12.** Verification of Storage Configuration

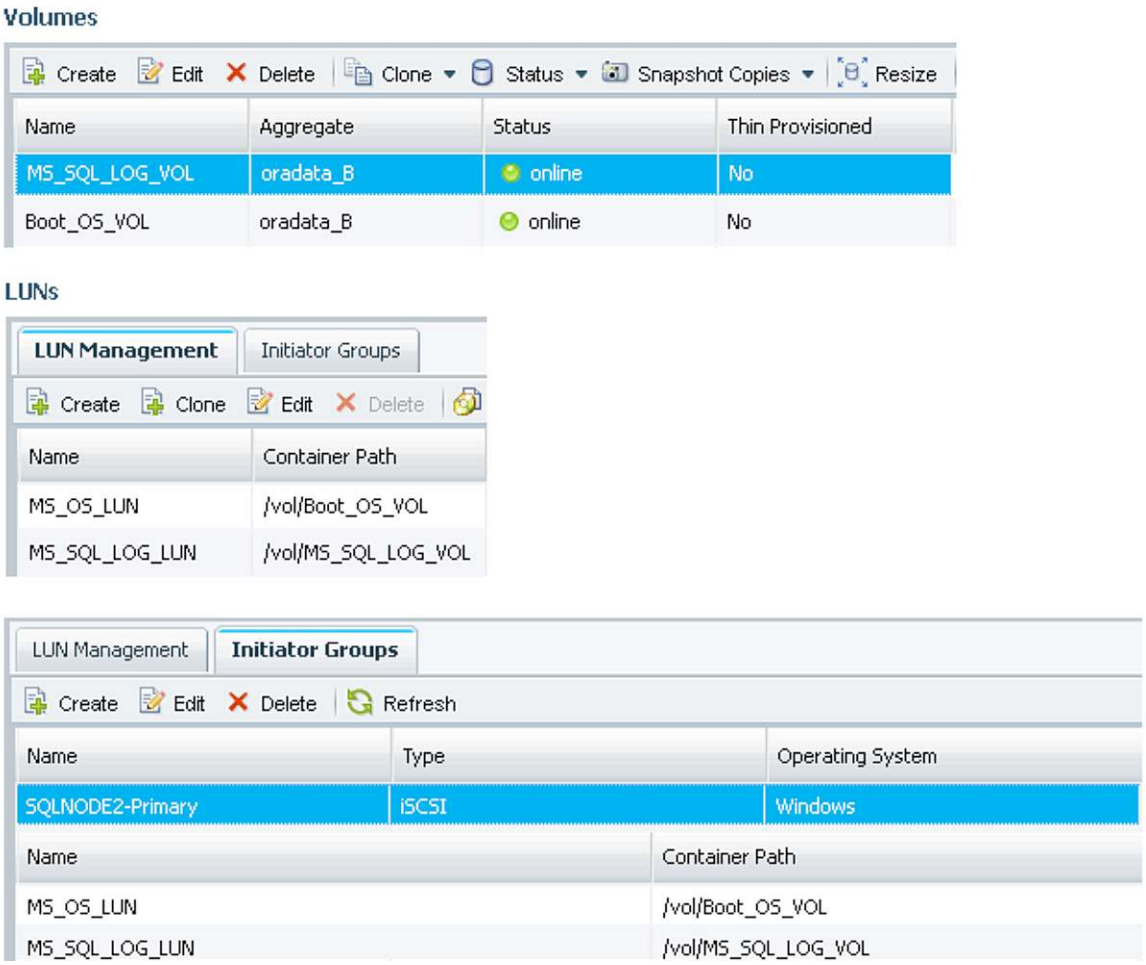


NetApp Multimode Virtual Interfaces

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 (Microsoft Windows hosts).

Figure 13 shows an iSCSI vPC-enabled network design on Cisco Nexus 5548UP and NetApp FAS3270HA Controllers A **a**nd B for accessing a Microsoft SQL Server data network.

**Figure 13.** iSCSI vPC Enabled on Cisco Nexus 5548UP

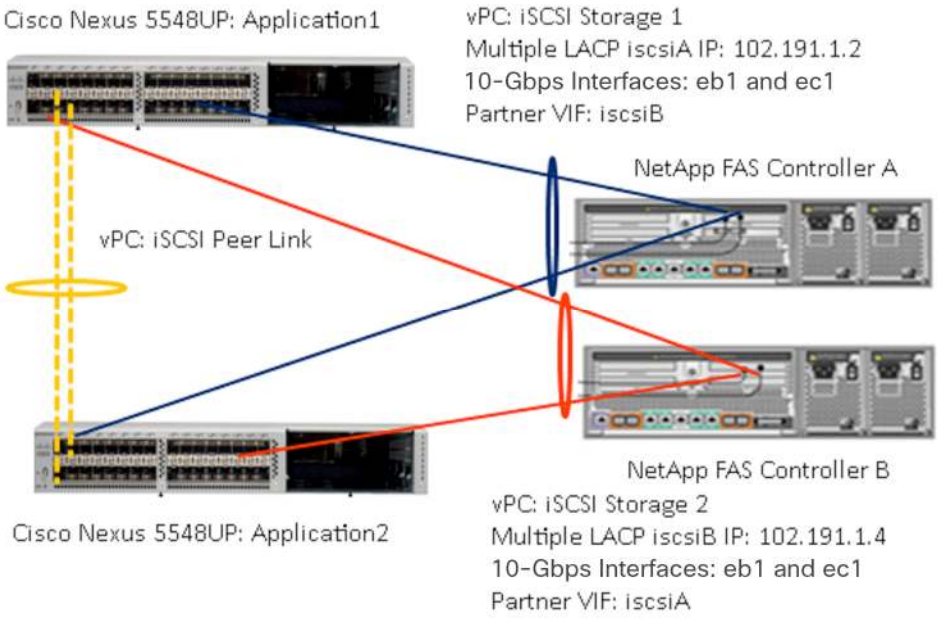


The vPC design layout for Cisco Nexus 5548UP Switches and corresponding NetApp cluster storage system multimode VIFs is as follows:

- 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 Design Topology section earlier in this document.
- 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 frames enabled for accessing storage using the iSCSI protocol. VIF iSCSI B is 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 the e1b and e1c 10 Gigabit Ethernet interfaces for accessing storage through the iSCSI protocol from the Microsoft Windows 2008 R2 host-based software initiator.

**Note:** On the Cisco Nexus 5548UP upstream switch, make sure that the correct QoS class and MTU value and policy types are applied to the PortChannel ports (eth19 and eth 20). PortChannels are connected to the NetApp FAS3270HA (Controllers A and B), on the 10 Gigabit Ethernet interfaces (e1b and e1c), to allow network packets to be tagged from the Cisco Nexus 5548UP fabric. This step is needed because NetApp storage will not tag any network packets with MTU and QoS values.

The commands that follow show how to configure the CoS on the Cisco Nexus 5548UP for untagged packets originating from storage on the PortChannels.

**Cisco Nexus 5548UP Application 1**

```
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
```

**Cisco Nexus 5548UP Application 2**

```
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_NXOS_Quality_of_Service_Configuration_Guide_chapter3.html#con_1150612.

NetApp VIF Configuration Details

The commands shown here 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 192.191.1.4 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.

## Microsoft Windows iSCSI Boot

This section describes the Cisco UCS service profile design for deploying the Microsoft Windows iSCSI boot OS from the NetApp shared iSCSI target on the Cisco UCS B-Series server. In this deployment, the Cisco UCS 1240 or 1280 VIC is used for iSCSI SAN bootup of the Microsoft Windows OS from the NetApp iSCSI target.

The following steps show the basic configuration on the service profile to enable Microsoft Windows 2008 R2 SP1 iSCSI SAN bootup on the Cisco UCS B200 M3 Blade Server from the NetApp iSCSI target. For more information about the configuration steps for deploying Microsoft Windows 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 **Chassis_Slot1_B200M3_SQLNode1** and **Chassis_Slot6_B200M3_SQLNode2** and associate them with Cisco B200 M3 blades using the Cisco UCS 1240 or 1280 VIC to install Microsoft Windows 2008 R2 SP1 from the iSCSI target on NetApp FAS3270. Figure 14 shows the creation of these service profiles.

**Figure 14.** Creating Service Profiles

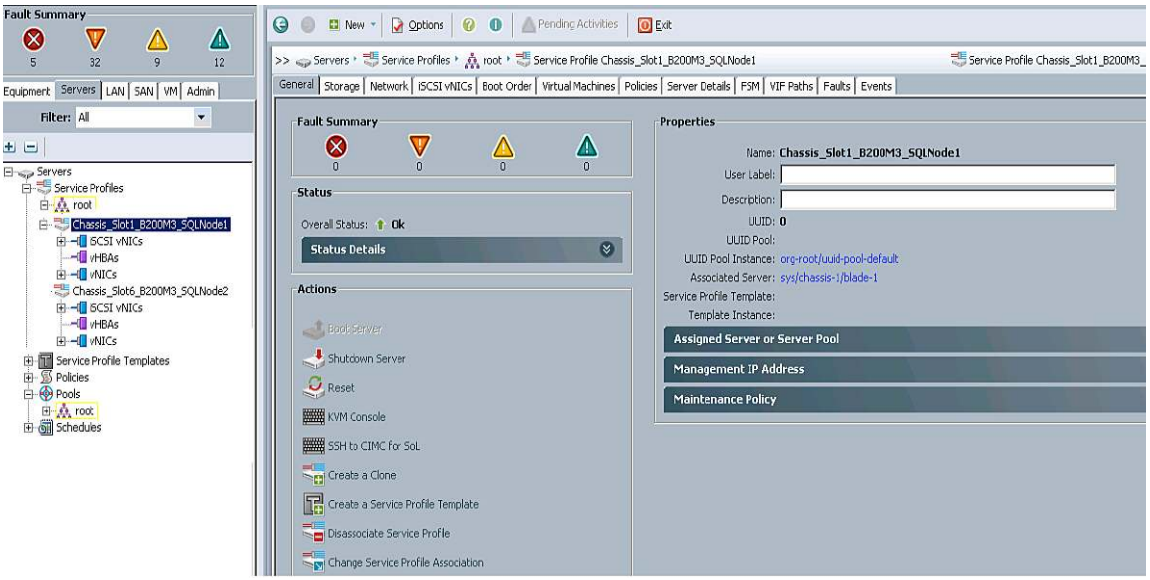

2. On the Service Profiles tab for the newly created service profiles, create four static vNICs, **iSCSI1**, **iSCSI2**, **iSCSI3**, and **iSCSI4**, on Fabric A and Fabric B, respectively, with the MTU value set to 9000, without fabric failure and with network VLAN access to VLAN ID 192 (iSCSI storage), as shown in Figure 15.
3. To access Microsoft Windows for management and SQL applications by clients, create two separate static vNICs (**vNIC1** and **vNIC2**) with VLAN ID 811, as shown in Figure 15.
4. For Microsoft Windows cluster control heartbeats across Microsoft SQL Server failover cluster nodes, create a separate static vNIC (**vNIC3**) with VLAN ID 193, as shown in Figure 15.

**Figure 15.** Creating Static vNICs on Fabric Interconnects

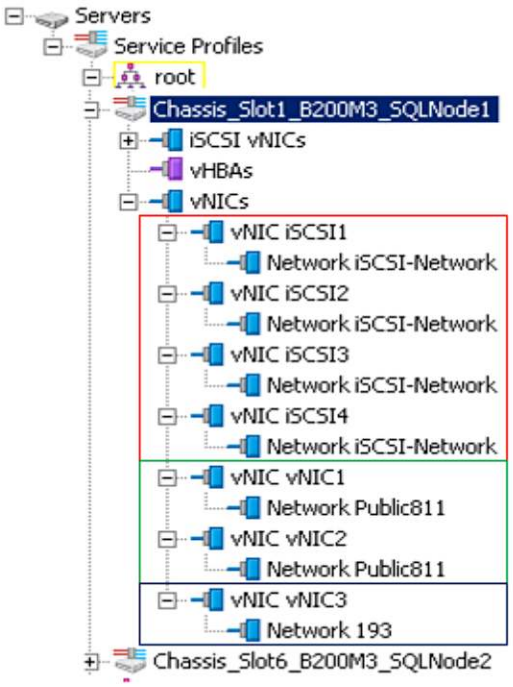


5. On the desired service profile, initially create one iSCSI vNIC, **iSCSI-Primary**, which is required to access the NetApp storage iSCSI target during iSCSI bootup to load the Microsoft Windows operating system over the iSCSI network. Make sure that the iSCSI vNIC **iSCSI-Primary** is overlaid on static vNIC **iSCSI1**, as shown in Figure 16.

   A single iSCSI vNIC on the service profile is required only during the iSCSI installation period, because the Microsoft Windows OS requires the iSCSI NetApp target to be accessed over a single iSCSI path to detect the iSCSI LUN during installation.

**Figure 16.** iSCSI vNICs Overlaid on Static vNICs

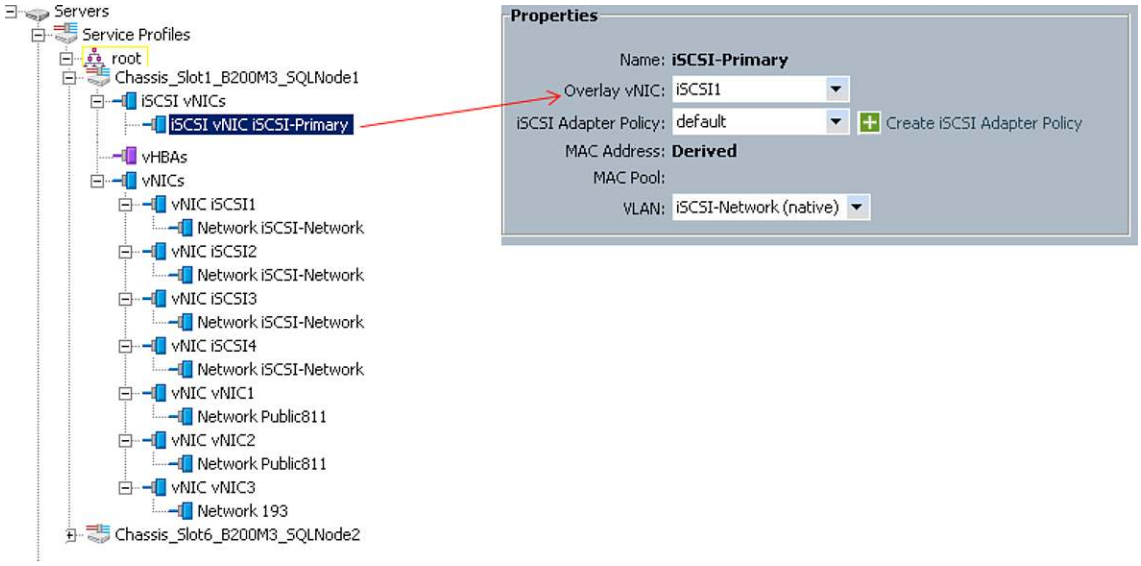


For the Cisco UCS 1280 or 1240 VIC, make sure that the MAC address is marked “Derived” and that the correct VLAN ID (192) is chosen to access the NetApp iSCSI target during Microsoft Windows 2008 R2 SP1 server iSCSI bootup.

6. In Cisco UCS Manager, create a new iSCSI boot policy, **iSCSI-Boot**, with iSCSI vNIC **iSCSI-Primary** as a primary path during Microsoft Windows host iSCSI bootup. Figure 17 shows the boot policy configuration.

   A single iSCSI vNIC defined in the **iSCSI-Boot** policy is required only during the iSCSI installation period, because the Microsoft Windows OS requires the iSCSI NetApp target to be accessed over a single iSCSI path to detect the iSCSI LUN during installation.

**Figure 17.** New iSCSI Boot Policy in Cisco UCS Manager

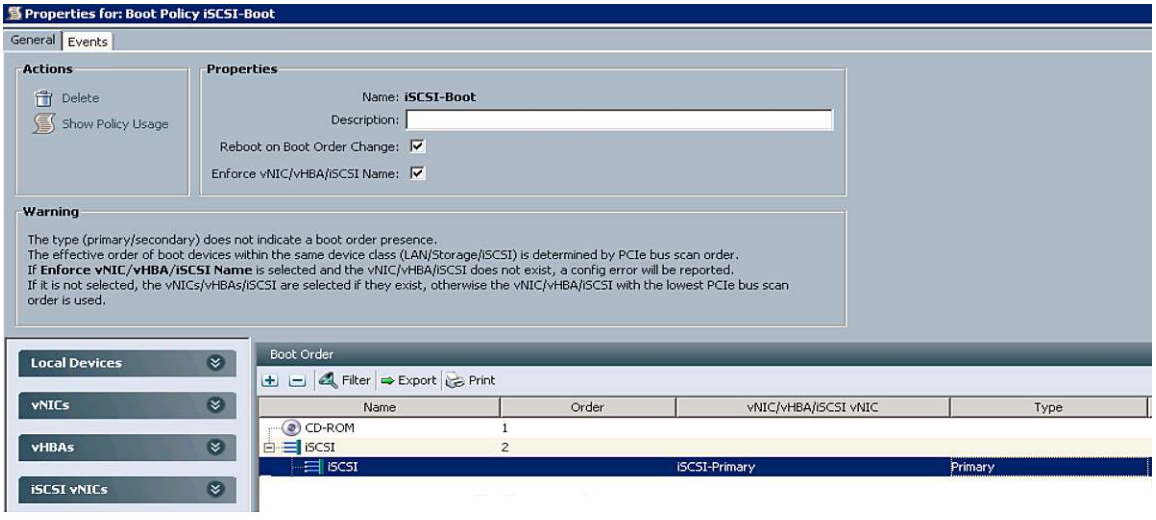

7. After the iSCSI boot policy is created, choose a newly created boot order policy for the desired service profile. For the chosen service profile on the Cisco UCS Manager Boot Order tab, assign **iSCSI-Primary** as the primary iSCSI vNIC with Microsoft Windows iSCSI boot parameters as shown in Table 8 and Figure 18.

   If you have defined a single network subnet without any routing (Layer 2) iSCSI network for the Microsoft Windows OS iSCSI boot or to access the Microsoft SQL Server database over the iSCSI protocol, make sure to set the iSCSI target IP address as the default gateway in the iSCSI boot parameters.

**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-Primary | iqn.2012-01.com.windows:sqlnode1-primary | Static | 192.191.1.2 | 992-08.com.netapp:sn.1574125695 | 3260 | 192.191.1.2 | 0 |

**Figure 18.** Setting iSCSI Boot Parameters

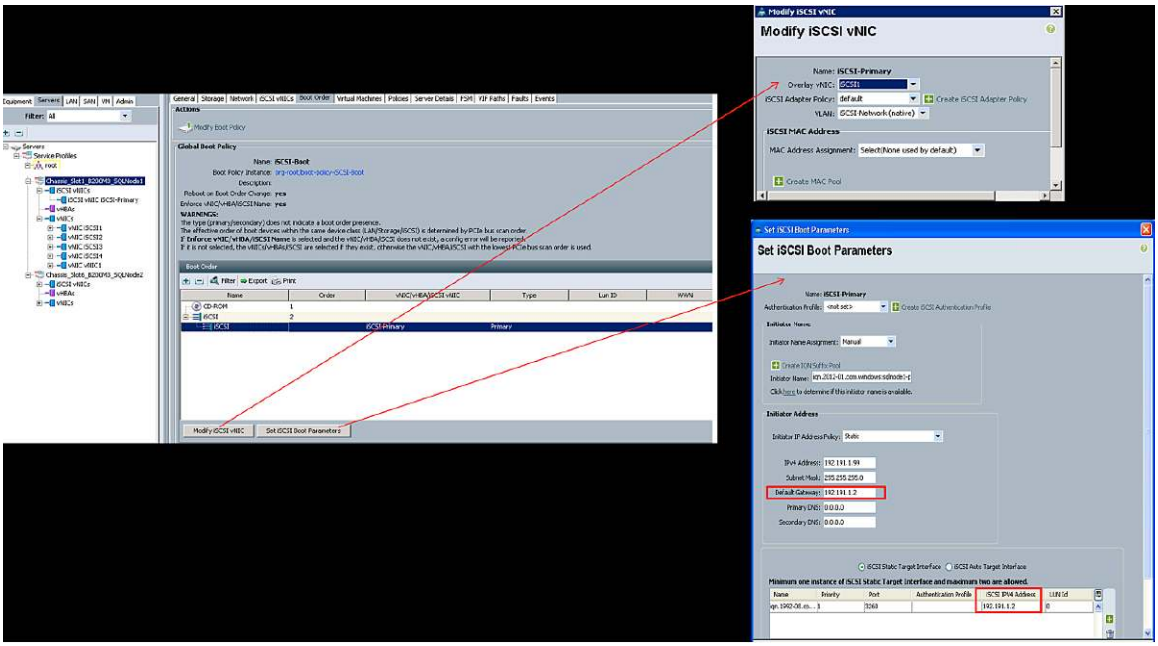


8. Associate the service profile with the desired blade (Cisco UCS B200 M3 in this case). On Cisco UCS in the associated service profile, launch the keyboard, video, and mouse (KVM) console. Through the virtual media interface, map the Microsoft Windows 2008 R2 SP1 ISO image and install the operating 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.

9. After the completion of the Microsoft Windows 2008 R2 SP1 iSCSI OS boot installation and bootup, access the KVM console for verification of the iSCSI-booted Microsoft Windows 2008 R2 SP1 OS. After the system is booted, define the secondary iSCSI vNIC **iSCSI-Secondary** in the service profile and iSCSI boot parameters to achieve redundancy in the event of failure during the iSCSI boot process. Figure 19 shows the final iSCSI

vNIC and boot parameters settings needed for the Microsoft Windows 2008 R2 SP1 iSCSI OS boot. In Figure 19 you can see in the iSCSI boot properties window that the secondary iSCSI vNIC (**iSCSI-Secondary**) is configured and can boot if the primary iSCSI vNIC (**iSCSI-Primary**) fails.

**Figure 19.** iSCSI vNIC and Boot Parameter Settings

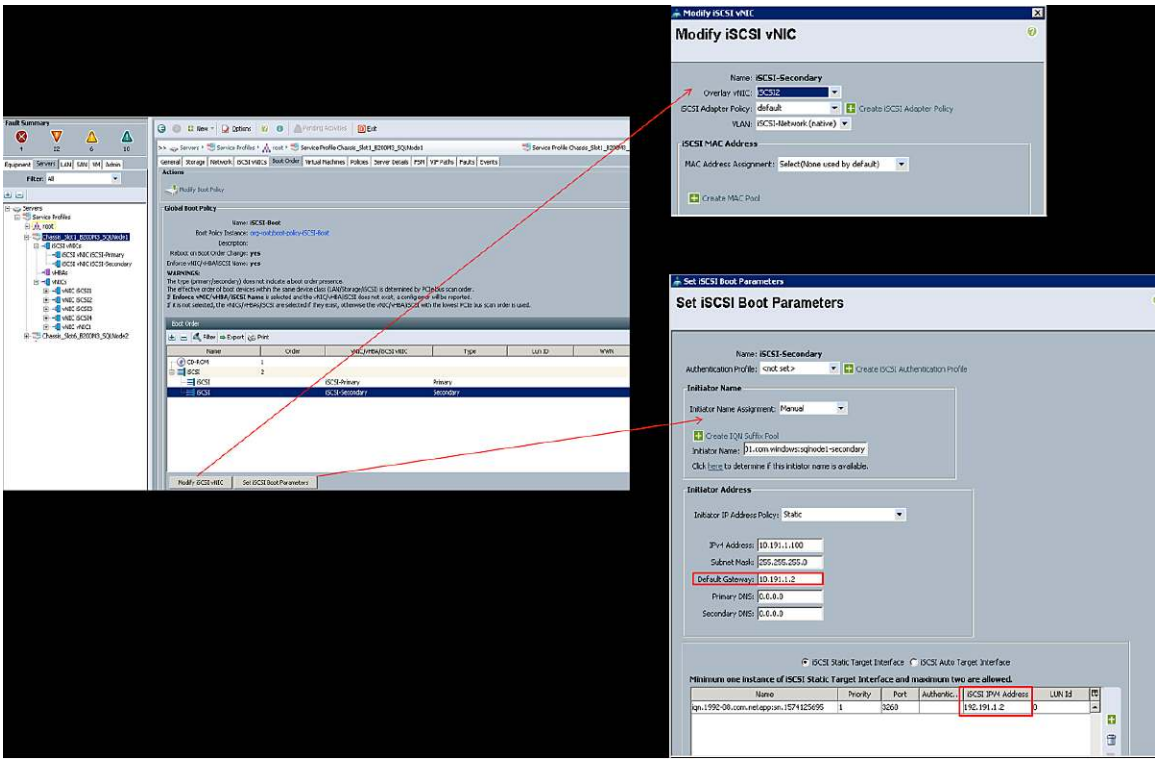


10. To configure the Microsoft Windows 2008 R2 SP1 iSCSI boot management IP address, access the Cisco UCS B200 M3 KVM console. Log in to the Microsoft Windows server, and to identify the Ethernet adapter connection interface that has been configured with VLAN ID 811 (management or application) access based on the MAC address, run the **ipconfig/all** command from the Microsoft Windows CLI. Figure 20 shows Ethernet adapters Local Area Connection 07 and Local Area Connection 05 selected for configuration.

**Figure 20.** Configuring Network Management

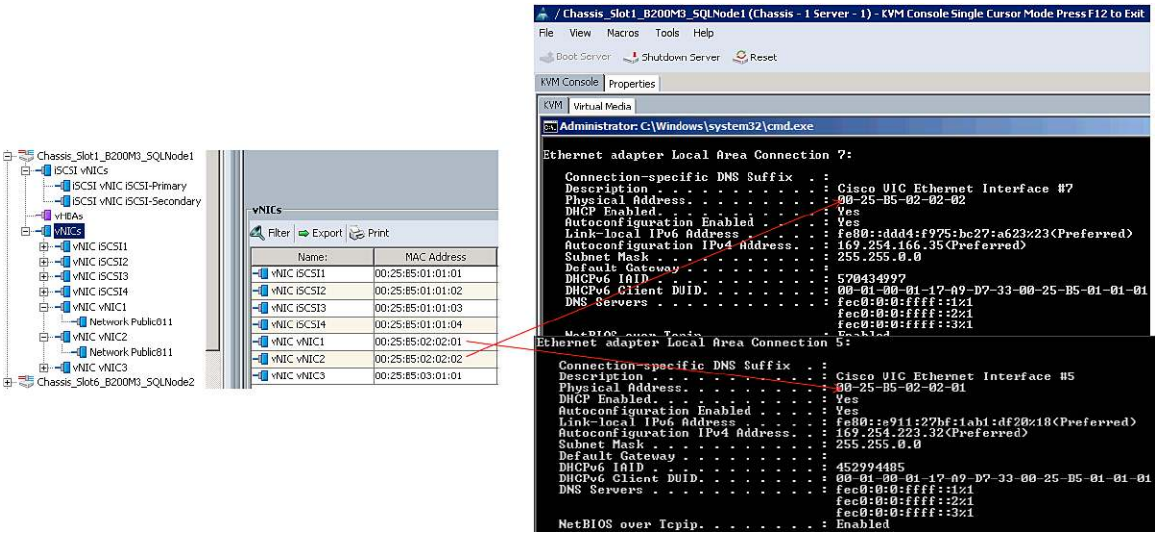


11. After identifying the Ethernet adapter interfaces, configure network redundancy using Cisco VIC network teaming software to create a virtual Ethernet adapter that is configured with the Microsoft Windows 2008 R2 SP1 management IP address.

    Follow these steps to configure the Cisco VIC network teaming software:

    a. Copy the Cisco VIC network teaming software to the iSCSI-booted Microsoft Windows 2008 R2 SP1 host.

    b. Load the Cisco NIC teaming protocol driver.

    c. Create NIC teaming with active-active load balancing on the identified Ethernet interfaces to create a vNIC-teamed Ethernet interface.

    d. Configure the management IP address on the vNIC-teamed Ethernet interface.

    Figure 21 shows the Cisco VIC network teaming and management IP address configuration details.

**Figure 21.** Management IP Configuration Details

```
C:\NicTeaming\W2K8R2\x64>enictool.exe -p "c:\NicTeaming\W2K8R2\x64"
Cisco NIC Teaming protocol driver installed successfully.

C:\NicTeaming\W2K8R2\x64>enictool.exe -c "Local Area Connection 5" "Local Area C
onnection 7" -m 3
The team creation is successfully completed.
```

```
C:\NicTeaming\W2K8R2\x64>enictool.exe -l

Ethernet Adapters available to Team:
Available Teams:
1. Local Area Connection 10 (Cisco NIC Teaming Miniport Driver #2)
```


With these steps, Microsoft Windows 2008 R2 SP1 installation is complete, with the iSCSI boot LUN configured on NetApp FAS3270.

## Microsoft Windows iSCSI Solution Overview

This section provides an overview of iSCSI solution deployment with Cisco UCS. The solution builds a Microsoft Windows iSCSI infrastructure to deploy a Microsoft SQL Server 2012 failover cluster solution on a Microsoft Windows 2008 R2 SP1 bare-metal OS on Cisco UCS B-Series Blade Servers (Cisco UCS B200 M3) connected to NetApp iSCSI storage over an iSCSI network, as described in the section Cisco UCS and Storage iSCSI Network.

### Physical and Logical Architecture

This section provides a physical and logical overview of iSCSI network architecture on Cisco UCS B-Series Blade Servers (Cisco UCS B200 M3) with Cisco VIC 1240 or 1280 adapters for deploying Microsoft Windows SQL Server with Microsoft Windows 2008 R2 SP1 using a native iSCSI software-based initiator to access shared NetApp storage over the iSCSI protocol.

Figure 22 shows the physical and logical architecture of the Cisco UCS configuration and provides an overview of the Microsoft Windows software-based iSCSI initiator configuration for deploying a Microsoft SQL Server 2012 failover cluster using Microsoft Windows 2008 R2 SP1.

**Figure 22.** Physical and Logical Architecture of Cisco UCS Configuration

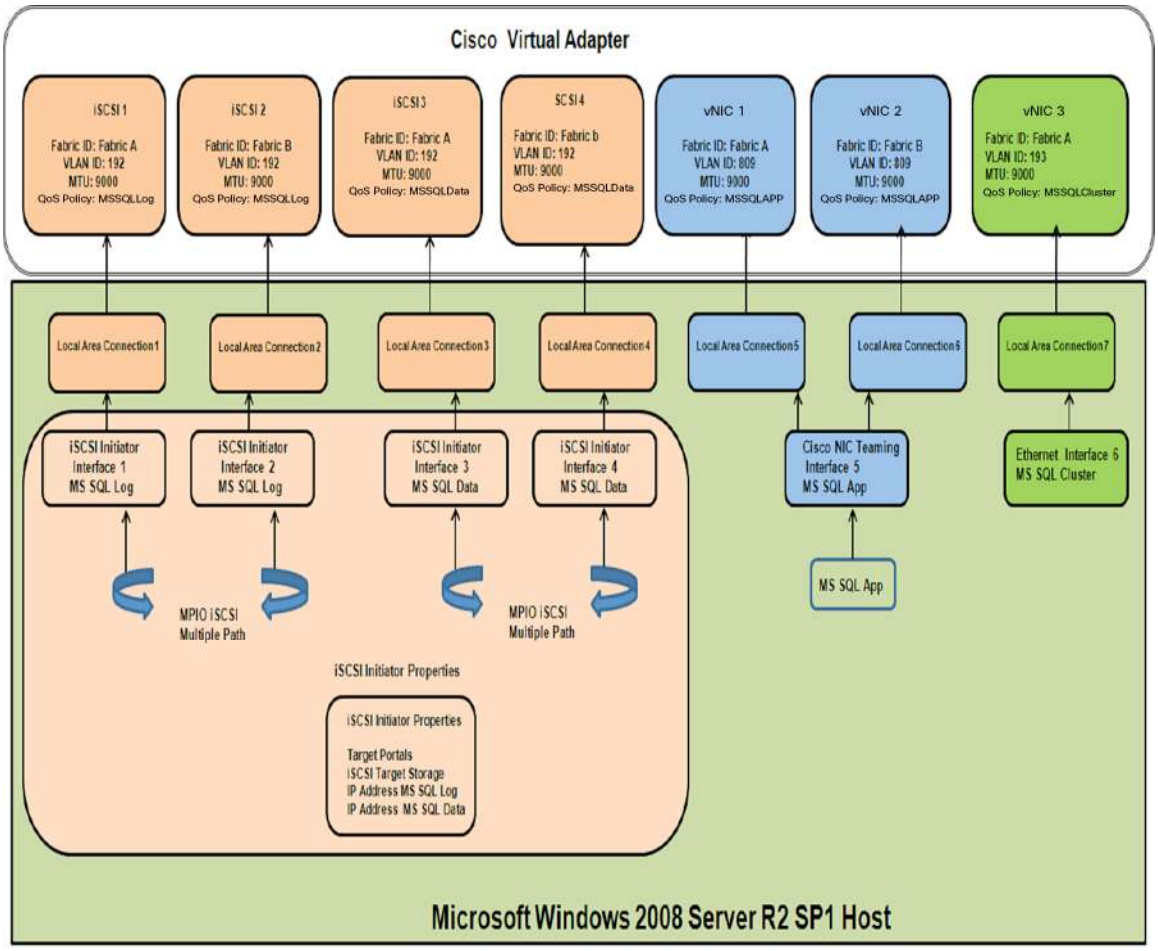


Cisco UCS Service Profile Design

This section explains the network considerations for the Cisco UCS service profile for creating static vNICs to enable a software-based iSCSI software initiator in a Microsoft Windows 2008 R2 SP1 environment.

The following vNICs need to be created in the service profile for the Microsoft Windows 2008 R2 SP1 host to access the iSCSI storage target:

- Create two iSCSI vNICs, **iSCSI-Primary** and **iSCSI-Secondary**, which are overlaid on two static vNICs: **iSCSI1** on Fabric A and **iSCSI2** on Fabric B, respectively. See the Microsoft Windows iSCSI Boot for the implementation steps.
- Create two static vNICs, **iSCSI1** on Fabric A and **iSCSI3** on Fabric B, with VLAN ID 192, no fabric failover, and **MSSQLLog** QoS policy defined to handle Microsoft SQL Server 2012 database log iSCSI storage data network traffic.
- Create two static vNICs, **iSCSI3** on Fabric A and **iSCSI4** on Fabric B, with VLAN ID 192, no fabric failover, and **MSSQLData** QoS policy defined to handle Microsoft SQL Server 2012 database data iSCSI storage data network traffic.

- Create two static vNICs, **vNIC1** on Fabric A and **vNIC2** on Fabric B, with VLAN ID 809, fabric failover, and **MSSQLAPP** QoS policy defined to handle Microsoft Windows 2008 R2 SP1 management and Microsoft SQL Server 2012 application data network traffic.
- Create one static vNIC, **vNIC3**, on Fabric A, with VLAN ID 193, fabric failover, an MTU of 9000, and **MSSQLCluster** QoS policy defined to handle Microsoft cluster data network traffic from the Microsoft Windows 2008 R2 SP1 cluster nodes.

Table 9 lists the static vNICs with network properties created in the service profile.

**Table 9.** Cisco UCS Service Profile Configuration

| vNIC | MAC Address | Fabric ID | Fabric Failover | VLAN ID | MTU | Adapter Policy | QoS Policy |
|---|---|---|---|---|---|---|---|
| **iSCSI1** | 0025:b501:0101 | Fabric A | No | 192 | 9000 | Windows | MSSQLLog |
| **iSCSI2** | 0025:b501:0102 | Fabric B | No | 192 | 9000 | Windows | MSSQLLog |
| **iSCSI3** | 0025:b501:0103 | Fabric A | No | 192 | 9000 | Windows | MSSQLData |
| **iSCSI4** | 0025:b501:0104 | Fabric B | No | 192 | 9000 | Windows | MSSQLData |
| **vNIC1** | 0025:b502:0201 | Fabric A | Yes | 809 | 9000 | Windows | MSSQLApp |
| **vNIC2** | 0025:b502:0202 | Fabric B | Yes | 809 | 9000 | Windows | MSSQLApp |
| **vNIC3** | 0025:b503:0101 | Fabric A | Yes | 193 | 1500 | Windows | MSSQLCluster |

### Microsoft Windows Host iSCSI Design

After booting the Microsoft Windows 2008 R2 SP1 OS through the iSCSI LUN on the Cisco UCS B200 M3 Blade Server, you need to assign the correct IP address to the corresponding Ethernet interface. You identify this address using the MAC address mapping from the Cisco UCS service profile vNIC setting for accessing Microsoft SQL Server 2012 database data and log iSCSI traffic, Microsoft Windows host management traffic, Microsoft SQL Server 2012 application traffic, and Microsoft Windows cluster traffic, as shown in Figure 23.

**Figure 23.** Mapping of Microsoft Windows NIC Adapters to Cisco UCS Static vNICs


Note the following Microsoft Windows 2008 R2 SP1 host logical iSCSI design considerations when deploying a Microsoft SQL Server 2012 failover cluster solution:

- Microsoft Windows 2008 R2 SP1 host management network and Microsoft SQL Server 2012 application network traffic from clients is accessed through the Cisco vNIC-teamed Ethernet interface **(Local Area Connection 10)** with the MTU value set to 9000. The NIC-teamed Ethernet interface is enabled with active-active load balancing and is teamed on physical Ethernet interfaces **Local Area Connection 5** and **Local Area Connection 7**, with Cisco UCS fabric failover enabled to achieve high availability and better network throughput.
- To access the Microsoft SQL Server 2012 database network traffic from the Microsoft Windows 2008 R2 SP1 host, a software-based iSCSI initiator is enabled with multipath on physical interfaces **Local Area**

**Connection 0** and **Local Area Connection 3** with the MTU value set to 9000. iSCSI multipath is configured with active-active load-balancing policy to achieve high availability and better network throughput.

- To access the Microsoft SQL Server 2012 database log network traffic from the Microsoft Windows 2008 R2 SP1 host, a software-based iSCSI initiator is enabled with multipath on physical interfaces **Local Area Connection 6** and **Local Area Connection 2** with the MTU value set to 9000. iSCSI multipath is configured with active-active load-balancing policy to achieve high availability and better network throughput.
- To handle Microsoft 2008 Cluster network traffic across nodes, enable Ethernet interface **Local Area Connection 8** with the MTU value set to 1500 and enable Cisco UCS fabric failure, to achieve high availability.

Also note that in scenarios in which multiple IP gateways are specified, you can route network packets to a specific gateway IP address by defining a lower metric value for a specific Ethernet interface that is set with the designated gateway. This setup provides additional control over the metric that is used for local routes.

In the configuration shown here, Microsoft Windows 2008 R2 SP1 host management network and Microsoft SQL Server 2012 application network traffic network packets are routed to the Cisco vNIC-teamed Ethernet interface (**Local Area Connection 10**) that is set with the appropriate gateway. To achieve this setup manually, set a lower metric value on this Ethernet interface (**Local Area Connection 10**), and on other Ethernet interfaces use the default settings, as shown in Figure 24.

**Figure 24.** Lower Metric Value Setting on Ethernet Interface (Local Area Connection 10)

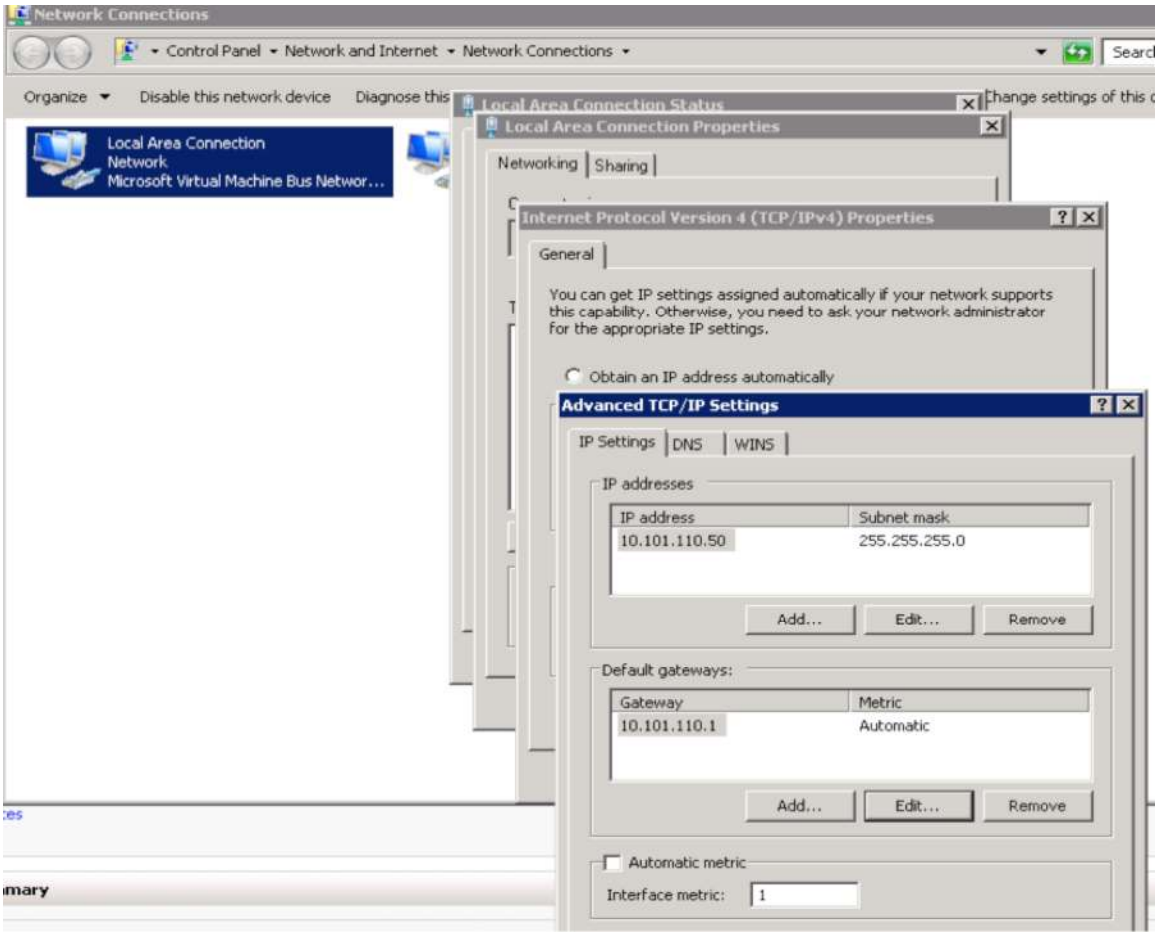


For more information, see http://support.microsoft.com/kb/299540.

To support end-to-end jumbo frames (MTU 9000) to carry the Microsoft SQL Server client, iSCSI, and cluster traffic from the Microsoft Windows 2008 R2 SP1 host, Cisco UCS, and NetApp storage, perform the following steps:

1. Configure MTU 9000 in the Cisco UCS QoS system with Platinum, Gold, Silver, and Bronze classes as shown in the section Cisco UCS Quality-of-Service System and Policy.
2. Configure MTU 9000 in the Jumbo field on the appropriate Ethernet interfaces (0, 10, 8, 6, 3, and 2 in this design) on the Microsoft Windows 2008 R2 SP1 host.
3. Configure NetApp iSCSI VIFs to enable the MTU 9000 value, as shown in the section Microsoft SQL Data Network and Storage Network vPC Mapping.
4. On the Microsoft Windows 2008 R2 SP1 host, 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://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 SP1 host iSCSI software initiator to access the NetApp iSCSI target:

1. Discover the storage controller NetApp FAS3270HA Controller A VIF **iscsiA** (192.191.1.2) with the Microsoft Windows 2008 R2 SP1 host iSCSI initiator Interfaces (192.191.1.99 and 192.191.1.100), and discover the storage controller NetApp FAS3270HA Controller B VIF **iscsiB** (192.191.1.101 and 192.191.1.102) with the Microsoft Windows 2008 R2 SP1 host iSCSI software initiator, as shown in Figure 25.

**Figure 25.** iSCSI Initiator Properties Showing Target IP Addresses

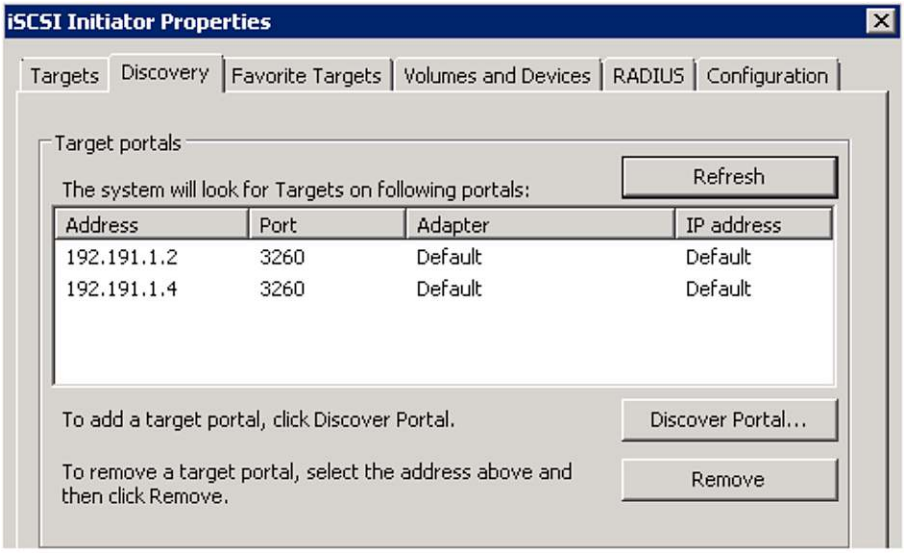


2. To enable iSCSI multipath, select the **Enable multi-path** check box in the target connection configuration for both controllers. Figure 26 shows multipath enabled.

**Figure 26.** Multipath Enabled for Target Connection

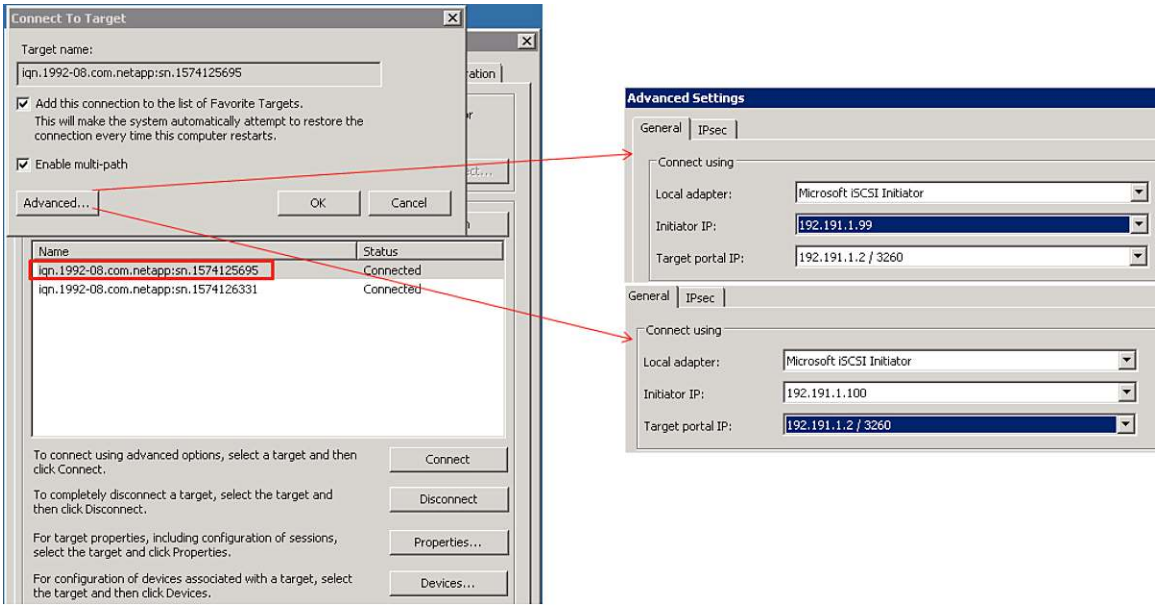

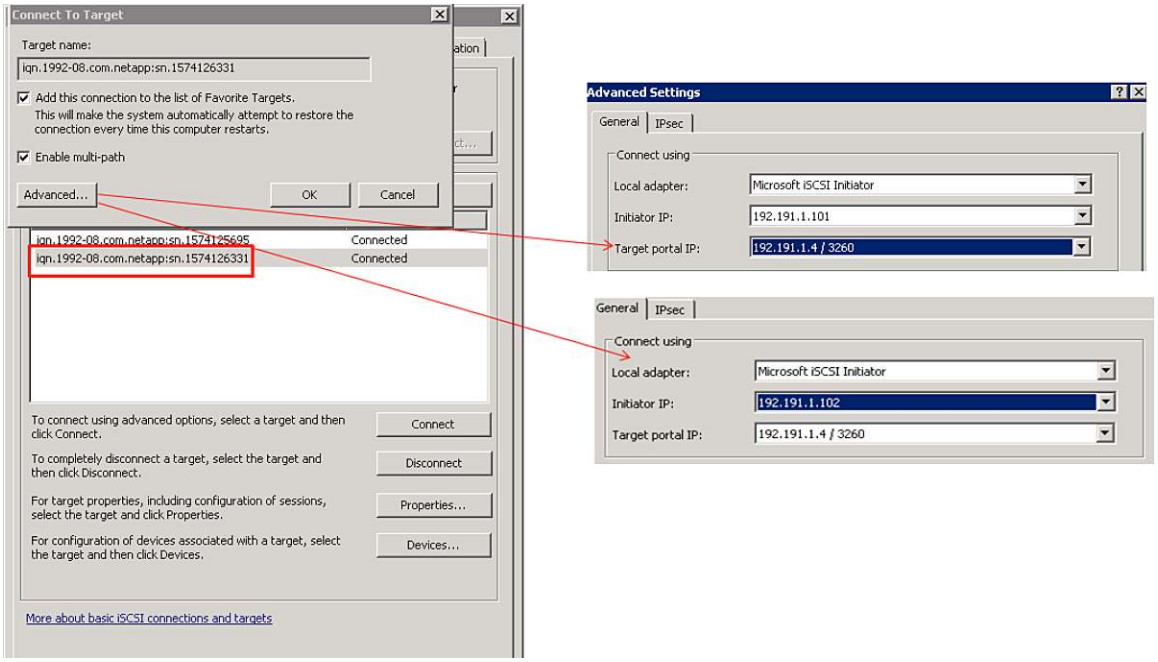


3. In Microsoft Windows 2008 R2, log in to the iSCSI initiators to access NetApp targets; LUNS are automatically exposed for configuration. For a Microsoft SQL Server 2012 single-host installation, you use Disk 1 as the database file and Disk 2 as the log file, as shown in Figure 27. See the section NetApp Storage Configuration Overview for LUN creation and access details.

**Figure 27.** Disk 1 and Disk 2 Details for Microsoft SQL Server 2012 Installation

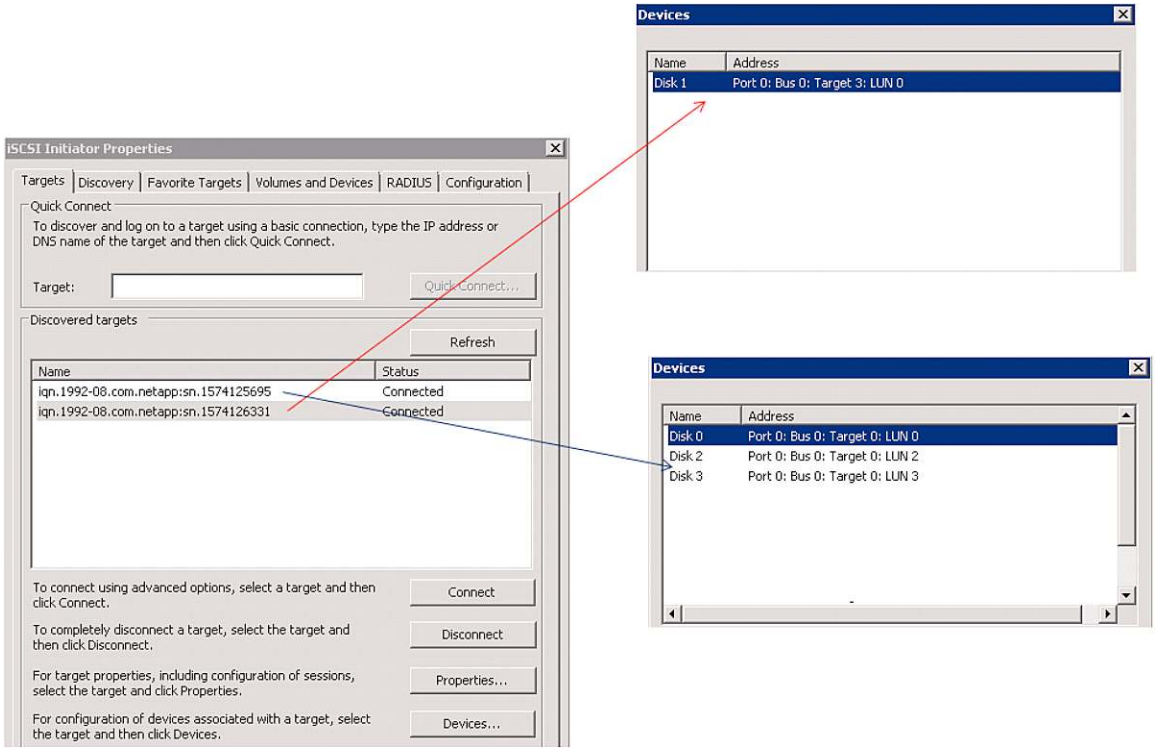

a. Under Disk Management, in Microsoft Windows 2008 R2, scan for new disks and format them. Create the NTFS file for storing Microsoft SQL Server 2012 database data and log files, as shown in Figure 28.

**Figure 28.** Scanning for New Devices on Microsoft Windows 2008 R2

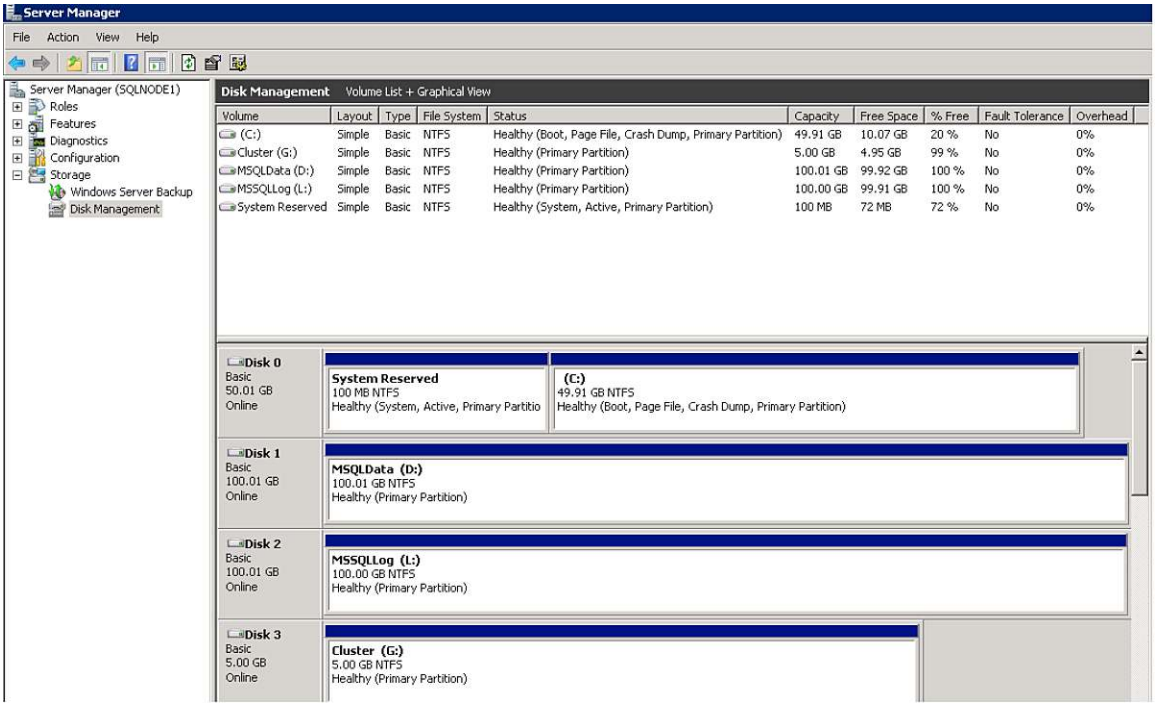


4. Install Microsoft SQL Server on the Microsoft Windows 2008 R2 OS and create the database with the data and log files residing on the designated storage volumes. For more information about installing Microsoft SQL Server 2012, see http://msdn.microsoft.com/en-us/library/bb500469%28v=sql.110%29.aspx.

## Microsoft SQL Server Failover Cluster 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 SP1 host with a software-based iSCSI initiator on Cisco UCS to access NetApp shared iSCSI storage in an iSCSI network environment.

Failover clustering provides very good protection in the event of a hardware failure. Failover to an active 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 such as 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 also can be deployed regionally.

### 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 SP1 host based on an iSCSI software initiator on a Cisco UCS B200 M3 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 B200 M3 Blade Servers mounted on separate Cisco UCS chassis for high availability and managed by a single Cisco UCS platform.

Figure 29 shows the physical and logical design of the Microsoft SQL Server 2012 failover cluster solution on a Microsoft Windows 2008 R2 SP1 host with an iSCSI software initiator.

**Figure 29.** Physical and Logical Design of Microsoft SQL Server 2012 Failover Cluster Solution

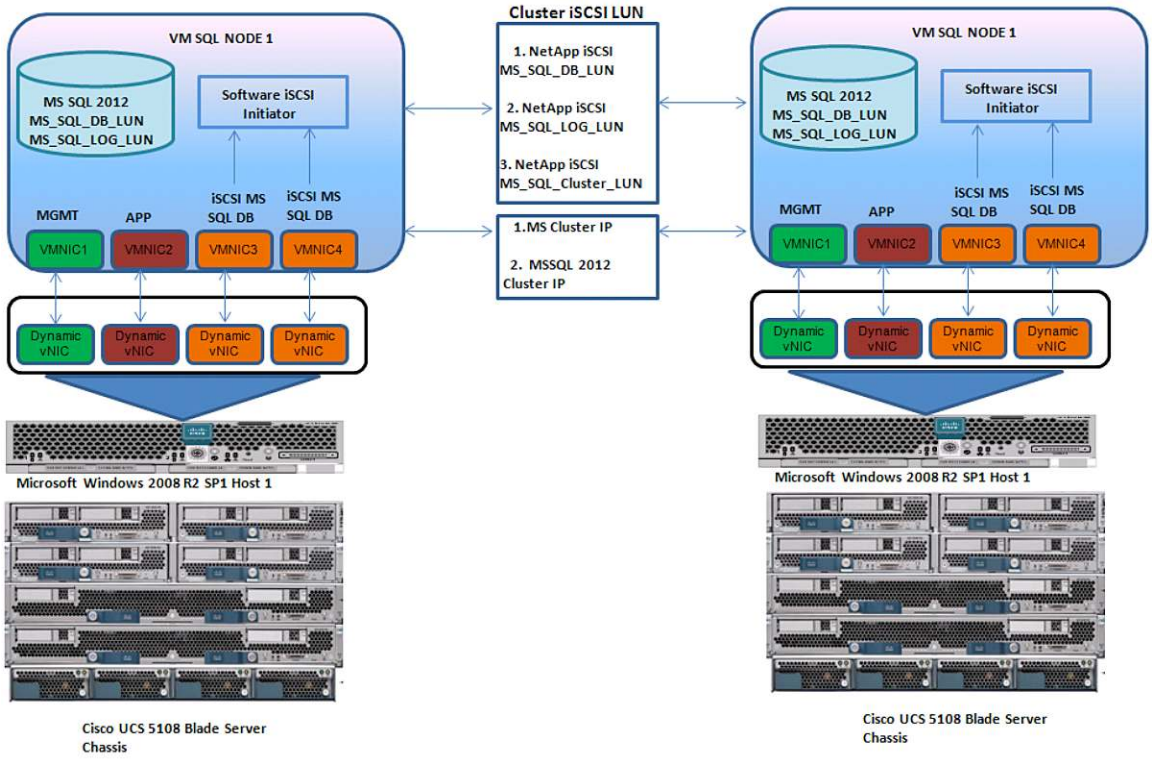


Perform the following steps to implement failover clustering on two Microsoft Windows 2008 R2 SP1 hosts **SQL Node1 and SQL Node2** using software-based iSCSI systems. These Microsoft Windows 2008 R2 SP1 nodes are part of a Microsoft Windows failover cluster, booted through the iSCSI target.

1. To implement Microsoft server clustering and a Microsoft SQL Server 2012 failover cluster server on Microsoft Windows 2008 R2 SP1 hosts, you need to use two Cisco UCS B200 M3 blades (**SQL Node1** and **SQL Node2**) in a dual chassis, as shown in Figure 29.

   Define two service profiles with the required network infrastructure to install Microsoft Windows 2008 R2 SP1 iSCSI boot operating systems, which are clustered to host failover clustering. The design of the Microsoft Windows 2008 R2 SP1 iSCSI boot OS and iSCSI initiator for individual hosts is explained in the sections Microsoft Windows iSCSI Boot and Microsoft Windows iSCSI Solution Overview, respectively.

2. To deploy Microsoft Windows 2008 failover cluster mode on Microsoft Windows 2008 R2 cluster nodes **SQLNODE1** and **SQLNODE2**, attach eight static vNICs as shown in Figure 29.

   a. Configure iSCSI interfaces with multipath on **Ethernet 6** and **Ethernet 2** local area connections Ethernet interfaces to access the iSCSI storage NetApp FAS3270HA Controller A VIF target, which hosts cluster LUNs

for Microsoft SQL Server database data files. The LUN is accessed through the Microsoft Windows 2008 R2 SP1 iSCSI software initiator.

b. Configure iSCSI interfaces with multipath on **Ethernet 0** and **Ethernet 3** local area connections Ethernet interfaces to access the iSCSI storage NetApp FAS3270HA Controller B VIF target, which hosts cluster LUNs for Microsoft SQL Server database Log files. The LUN is accessed through the Microsoft Windows 2008 R2 SP1 iSCSI software initiator.

c. Configure the Cisco NIC teaming virtual Ethernet interface on **Ethernet 5** and **Ethernet 7** local area connection Ethernet interfaces for internal and external clients to access the Microsoft SQL Server 2012 failover cluster server.

d. Configure the Microsoft cluster Ethernet interface on the **Ethernet 5** local area connection Ethernet interface for access cluster network communication across Microsoft SQL Server 2012 failover cluster servers.

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 SP1 cluster nodes **SQLNODE1** and **SQLNODE2**.

   a. Microsoft Windows 2008 R2 SP1 cluster nodes **SQLNODE1** and **SQLNODE2** use the iSCSI software initiator configured with the **Ethernet 2** and **Ethernet 6** local area connection interfaces to access the NetApp cluster storage iSCSI target VIF (NetApp FAS3270HA Controller A **iscsiA** and **Ethernet 0** and **Ethernet 3** local area connection interfaces to access NetApp FAS3270HA Controller B **iscsiB**), with multipath enabled to access the Microsoft SQL Server 2012 database data and log LUNs. On NetApp storage systems (NetApp FAS3270HA Controllers 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 SP1 host, which is part of the Microsoft Windows server.

   b. Make sure that you create igroups on both NetApp storage controllers, NetApp FAS3270HA Controller A and NetApp FAS3270HA Controller B, with both the **SQLNODE1** and **SQLNODE2** 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 NetApp Storage Configuration Overview.

   c. After exposing NetApp storage LUNs to the Microsoft Windows 2008 R2 SP1 host **SQLNODE1** and **SQLNODE2** 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 30.

**Figure 30.** Scanning for New Disks and Assigning New Disks to Cluster Nodes

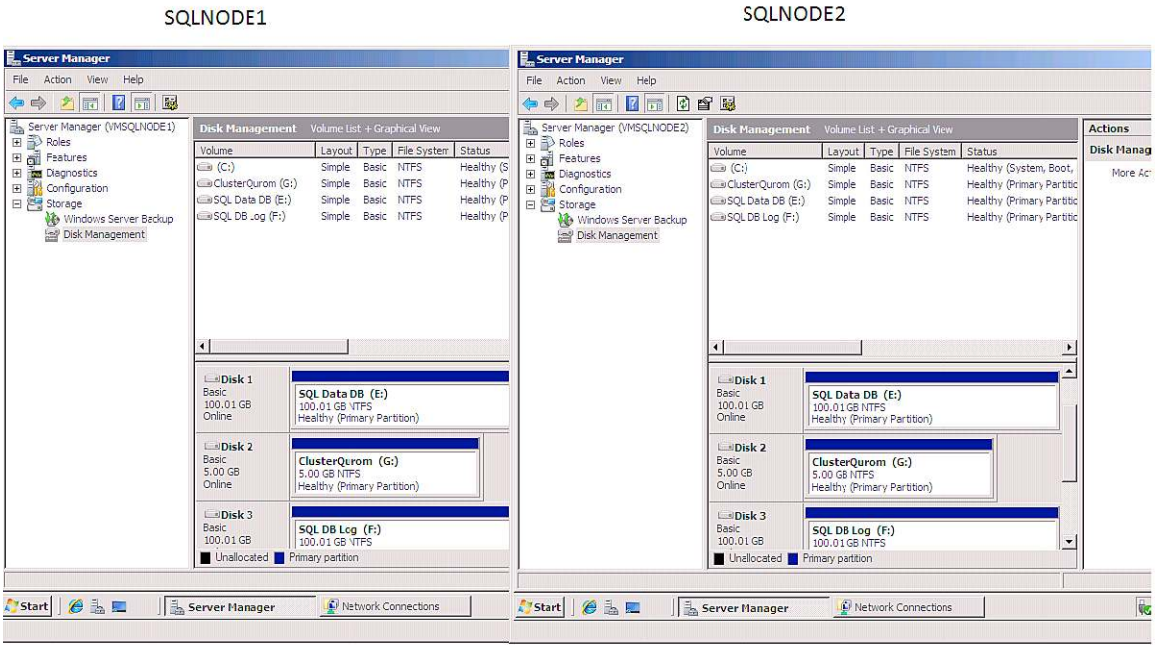


## Installation of Microsoft Windows 2008 Failover Cluster Feature with iSCSI Software Initiator

After configuration of the two Microsoft Windows 2008 R2 SP1 hosts, **SQLNODE1** and **SQLNODE2**, 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 SP1 hosts **SQLNODE1** and **SQLNODE2**, verify that they are part of a Microsoft Active Directory domain.
2. Log in to individual Microsoft Windows 2008 R2 SP1 hosts **SQLNODE1** and **SQLNODE2** by selecting the domain controller with Admin credentials.
3. Add the failover clustering feature on both Microsoft Windows 2008 R2 SP1 hosts: **SQLNODE1** and **SQLNODE2**.
4. After installing the failover cluster feature on both Microsoft Windows 2008 R2 SP1 hosts, log in to either host (Microsoft Windows 2008 R2 SP1 **SQLNODE1** or **SQLNODE2**) and launch the Failover Cluster Management console to validate clustering. Add Microsoft SQL Server hosts **SQLNODE1** and **SQLNODE2** with the fully qualified domain name in the Select Servers or a Cluster window, as shown in Figure 31.

**Figure 31.** Adding Microsoft SQL Server Guest

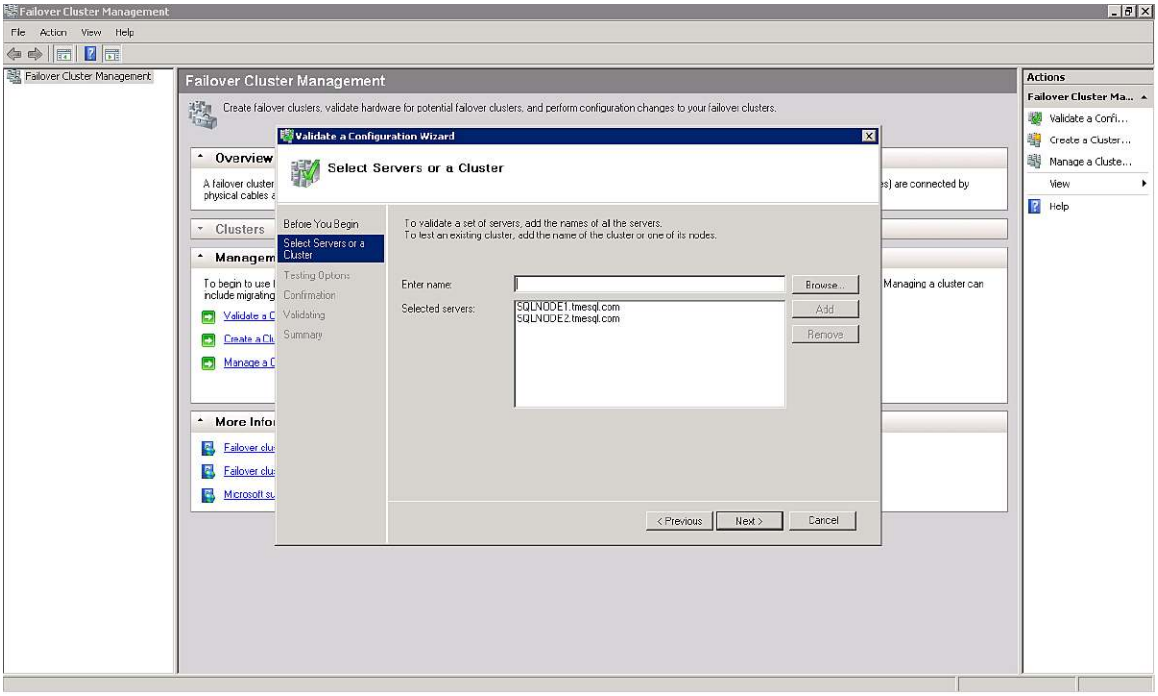


5. Select and add Microsoft Windows 2008 R2 SP1 hosts **SQLNODE1** and **SQLNODE2** and be sure to run all the tests to validate the cluster requirements, as shown in Figure 32.

**Figure 32.** Running Tests to Validate Cluster Requirements

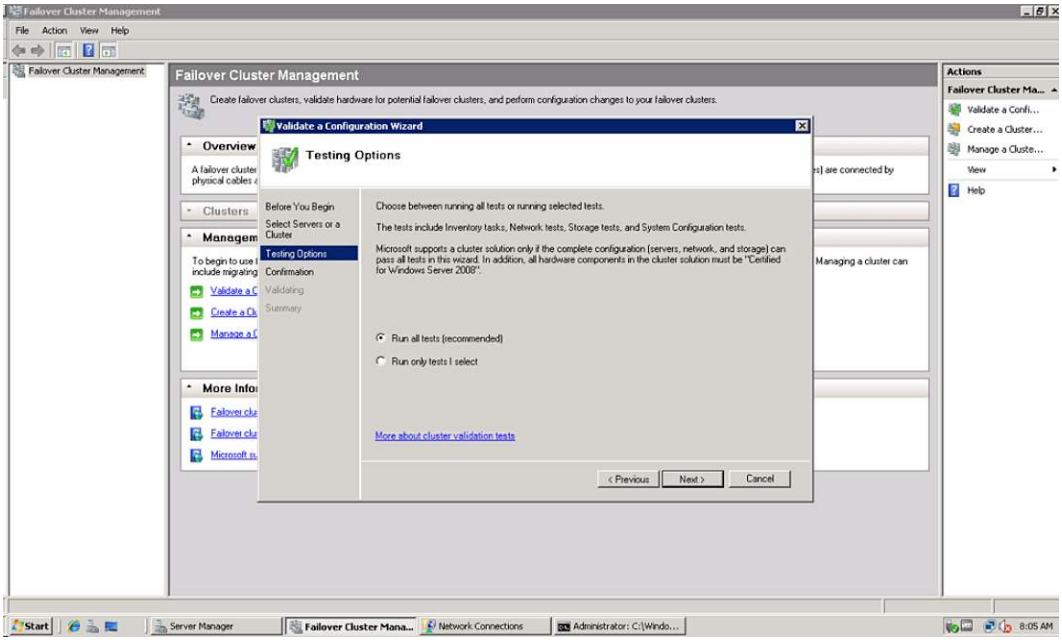

Successful validation of the addition of Microsoft Windows hosts **SQLNODE1** and **SQLNODE2** to the cluster is shown in Figure 33.

If warning messages appear, they need to be documented, and associated risks and action plans need to be identified.

**Figure 33.** Failover Cluster Validation Report

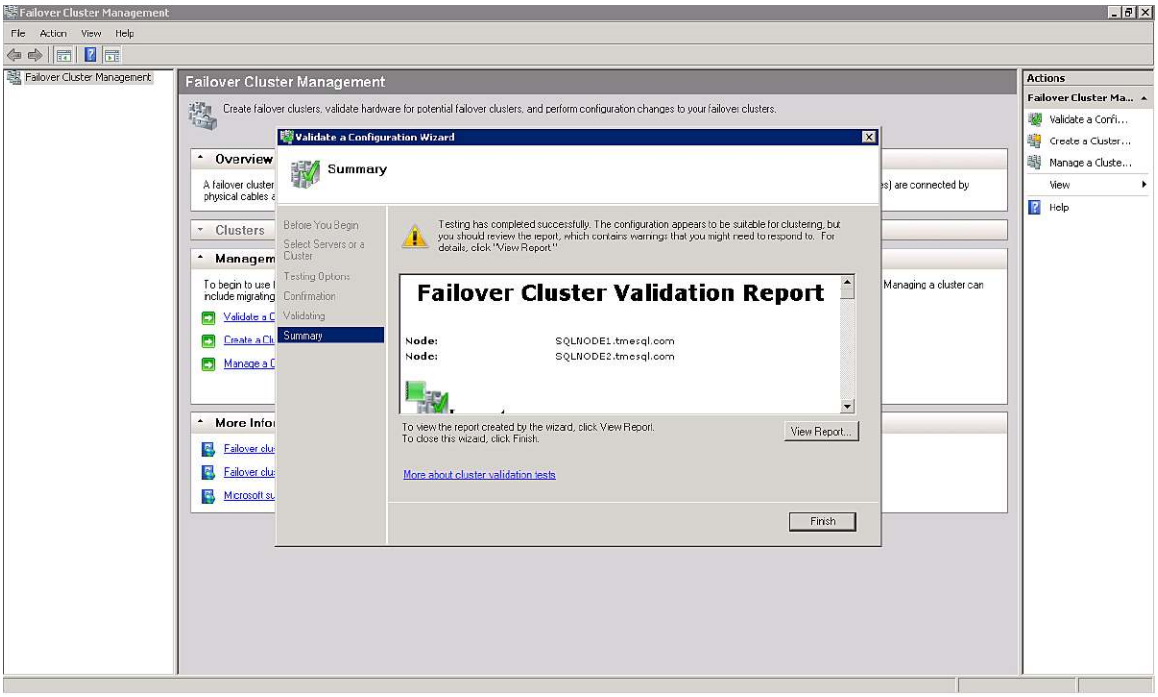


6. Create a cluster and name it **MSSQL** with the cluster and IP address on the management VLAN (809) as shown in Figure 34.

**Figure 34.** Administering the Cluster

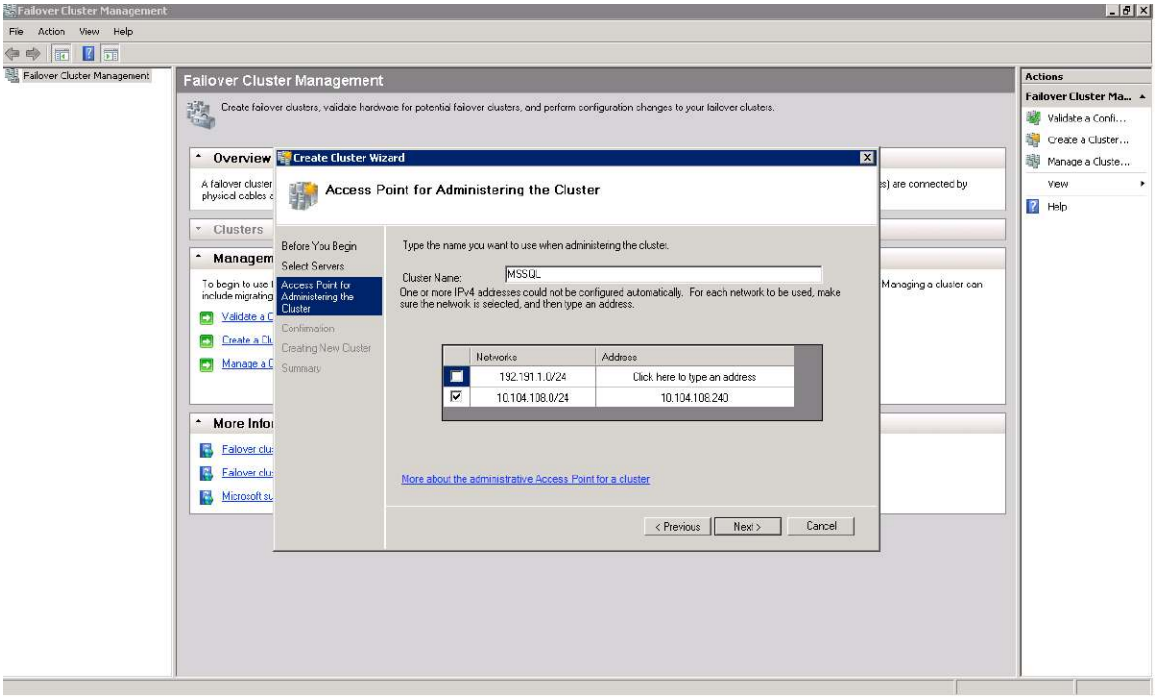


Figure 35 shows cluster summary information prior to creation of the cluster for the **SQLNODE1** and **SQLNODE2** nodes.

**Figure 35.** Summary of the Created Cluster

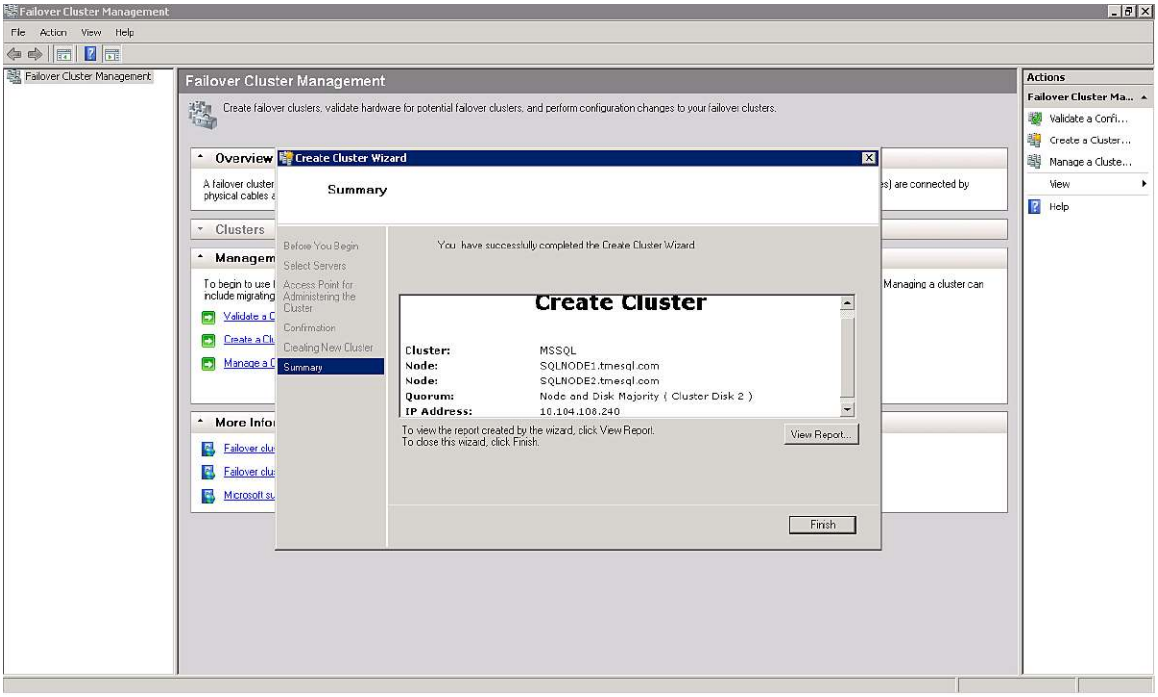

7. To validate the Microsoft cluster installation, log in to either Microsoft Windows 2008 R2 SP1 host (**SQLNODE1** or **SQLNODE2**) cluster node and launch the Failover Cluster Management console as shown in Figure 36.

**Figure 36.** Validating Microsoft Cluster Installation

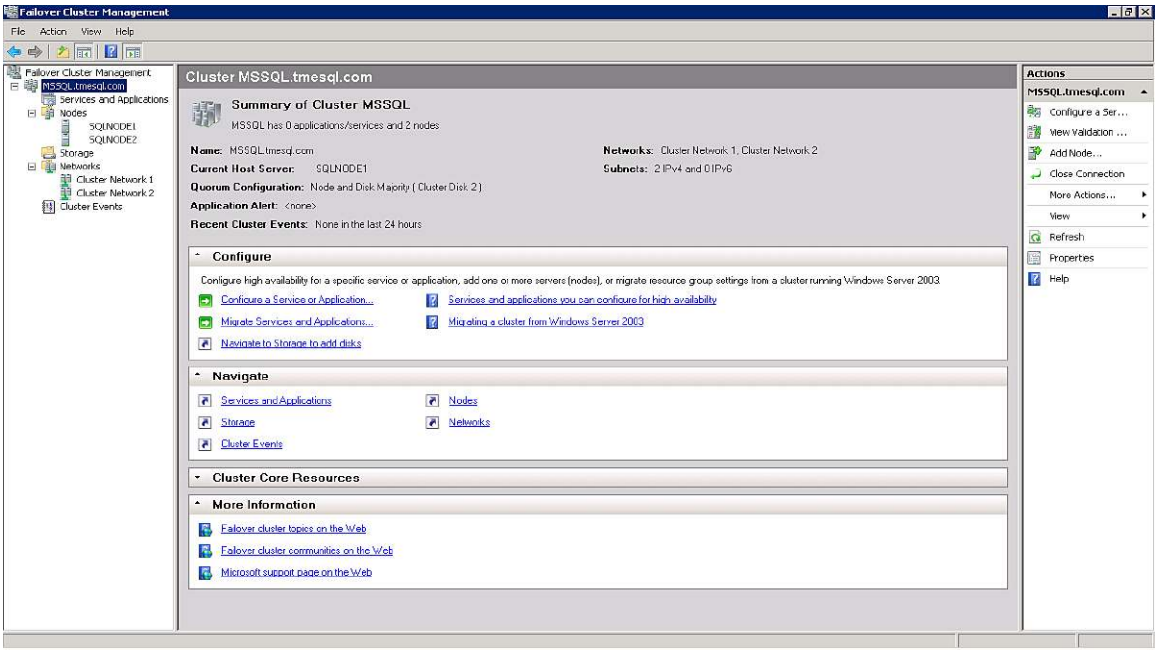


8. Log in to either Microsoft Windows 2008 R2 SP1 host (**SQLNODE1** or **SQLNODE2**) cluster node and choose the cluster witness disk for storing cluster information.
9. On the Failover Cluster Management console, select More Actions and then select the cluster name. In the Configure Cluster Quorum 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 witness disk if it is not been automatically identified, as shown in Figure 37.

**Figure 37.** Configuring MS SQL Cluster Quorum

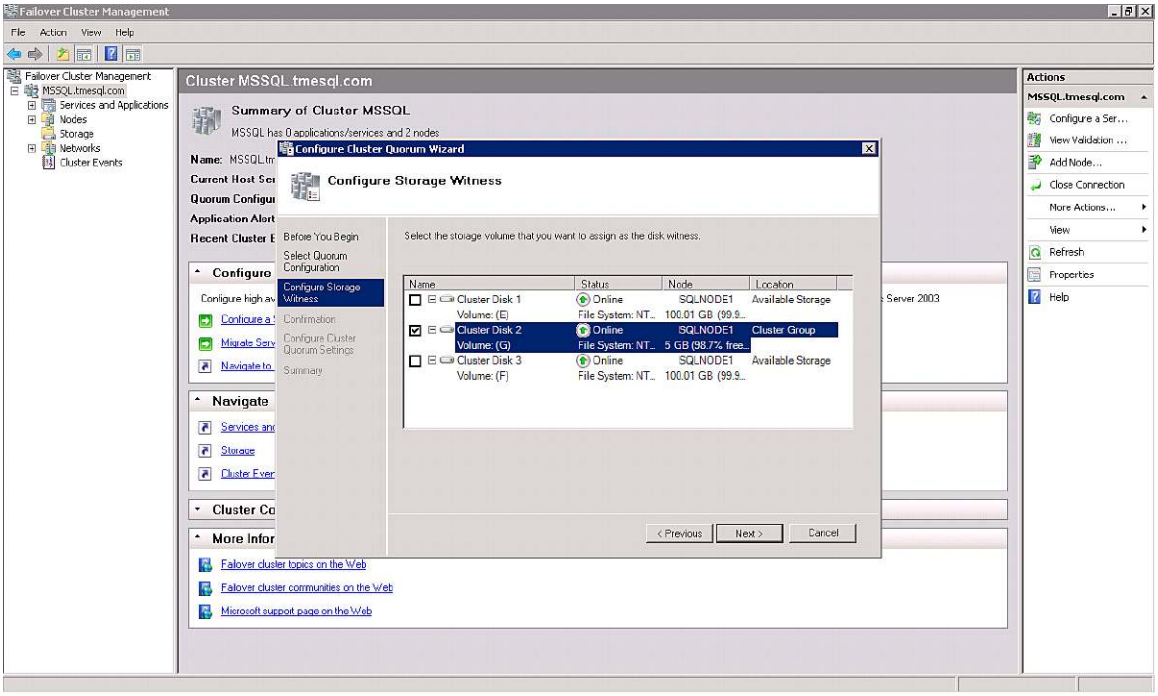


10. On the Failover Cluster Management console, select More Actions and then select Networks. Under Networks, choose Cluster Network 2 and select the radio button **Allow cluster network communication on this network**. Figure 37 shows the configuration of cluster networks.

**Figure 38.** Configuring Cluster Network

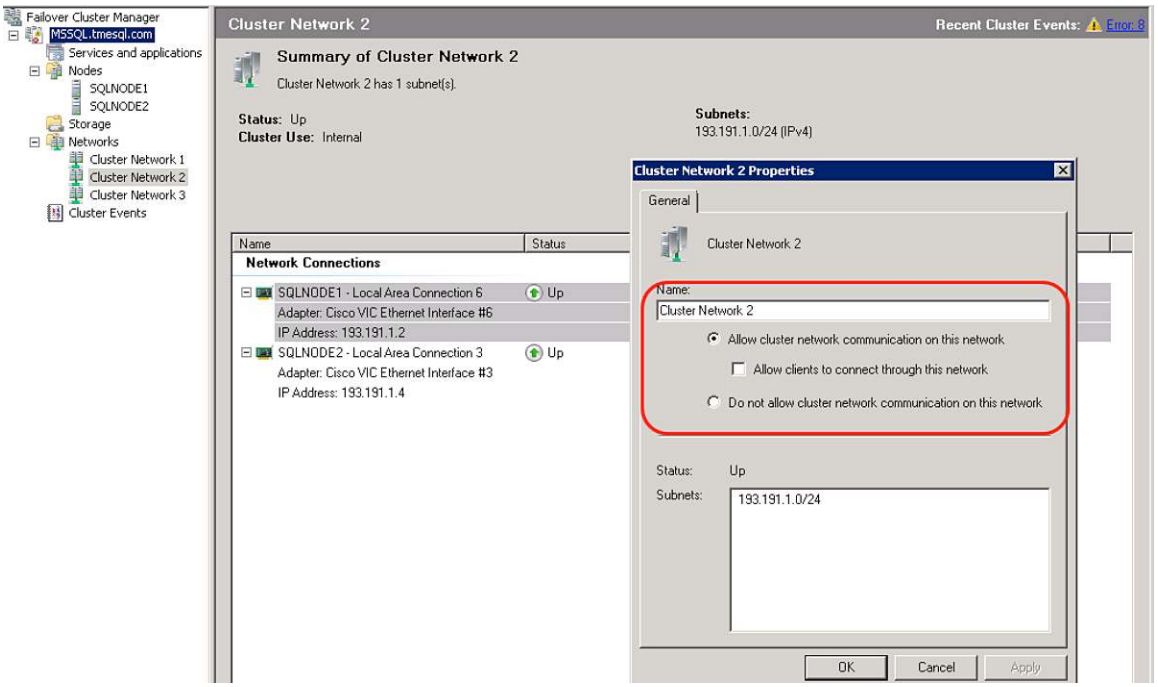

## Installation of Microsoft SQL Server 2012 Failover Cluster Feature with iSCSI Storage

The following steps describe deployment of Microsoft SQL Server 2012 failure clustering on Microsoft Windows 2008 R2 SP1 hosts.

1. Log in to either Microsoft Windows 2008 R2 SP1 cluster node, **SQLNODE1** or **SQLNODE2**, to perform installation on Microsoft SQL Server 2012. This document uses the **SQLNODE1** cluster node.

   a. Copy the Microsoft SQL Server 2012 binaries on the Microsoft Windows 2008 R2 SP1 host **SQLNODE1** cluster node for installation.

   b. Log in to **SQLNODE1** 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 39.

**Figure 39.** Launch Microsoft SQL Server in Microsoft SQL Server Installation Center

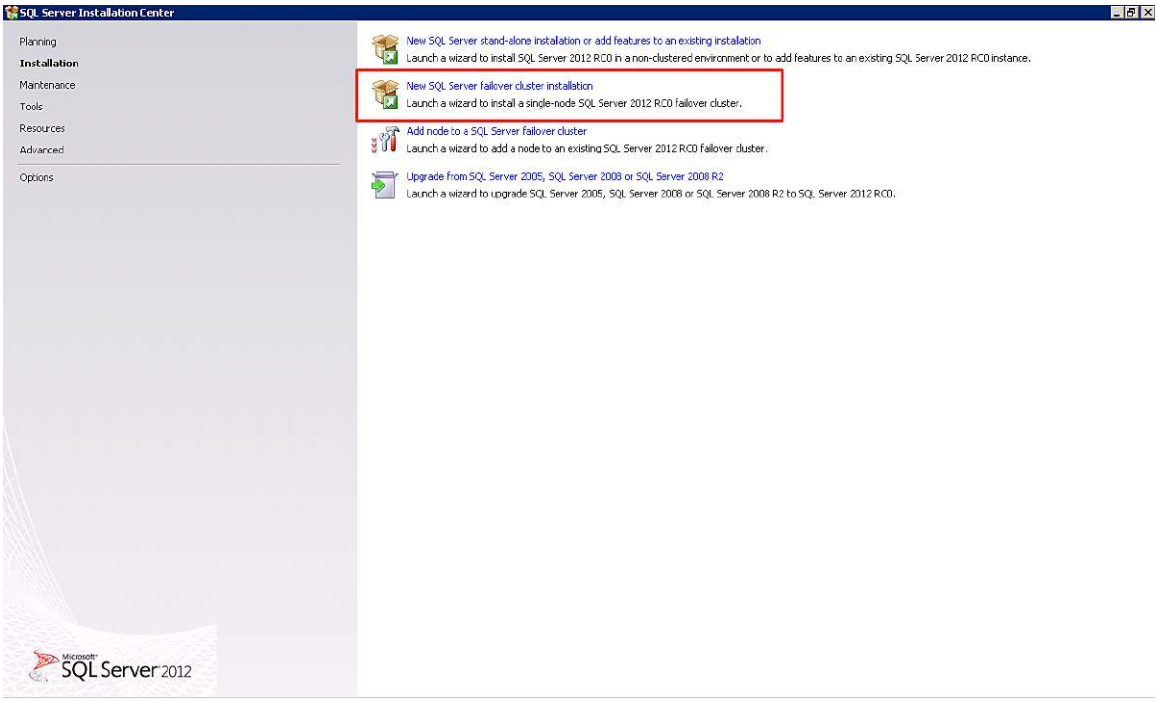


   c. Follow the installation steps in the wizard and provide license keys. Choose the Feature Selection option in the wizard and select appropriate features based on your requirements. For the Microsoft SQL Server binaries shared feature directory, provide the appropriate installation directory on the cluster node, **SQLNODE1**, as shown in Figure 40.

**Figure 40.** Selecting Evaluation Features to Install

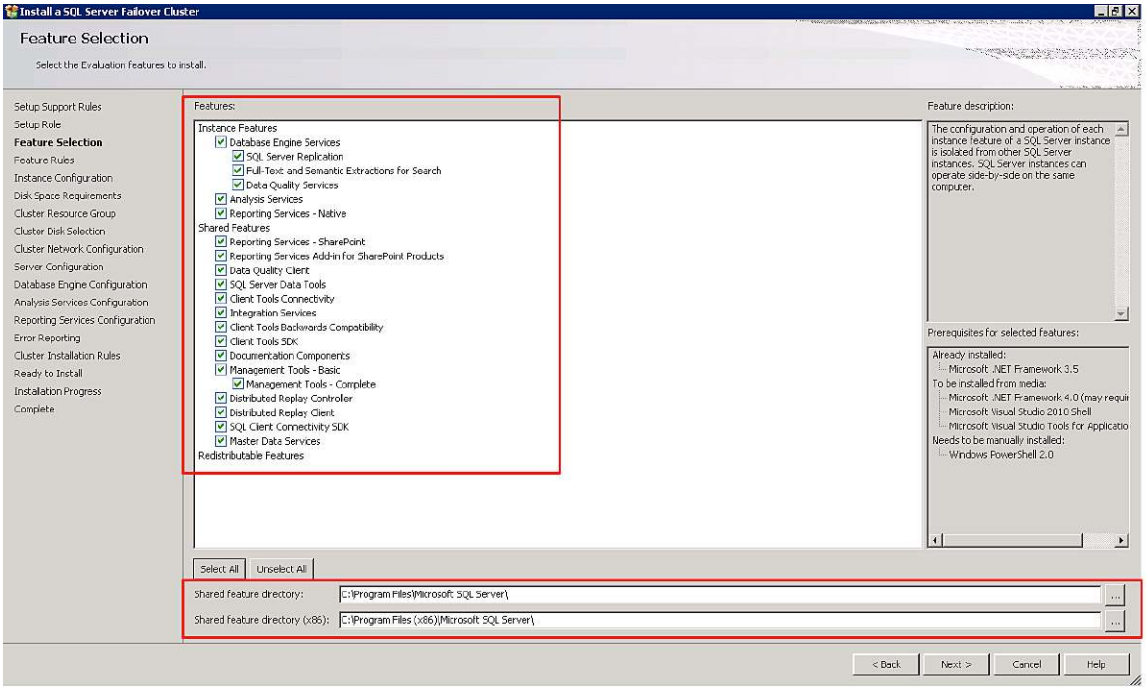


d. In the Instance Configuration window of the wizard, enter **SQL2012** as the name of Microsoft SQL Server 2012 failover cluster, as shown in Figure 41.

**Figure 41.** Instance Configuration in Server Failover Cluster Wizard

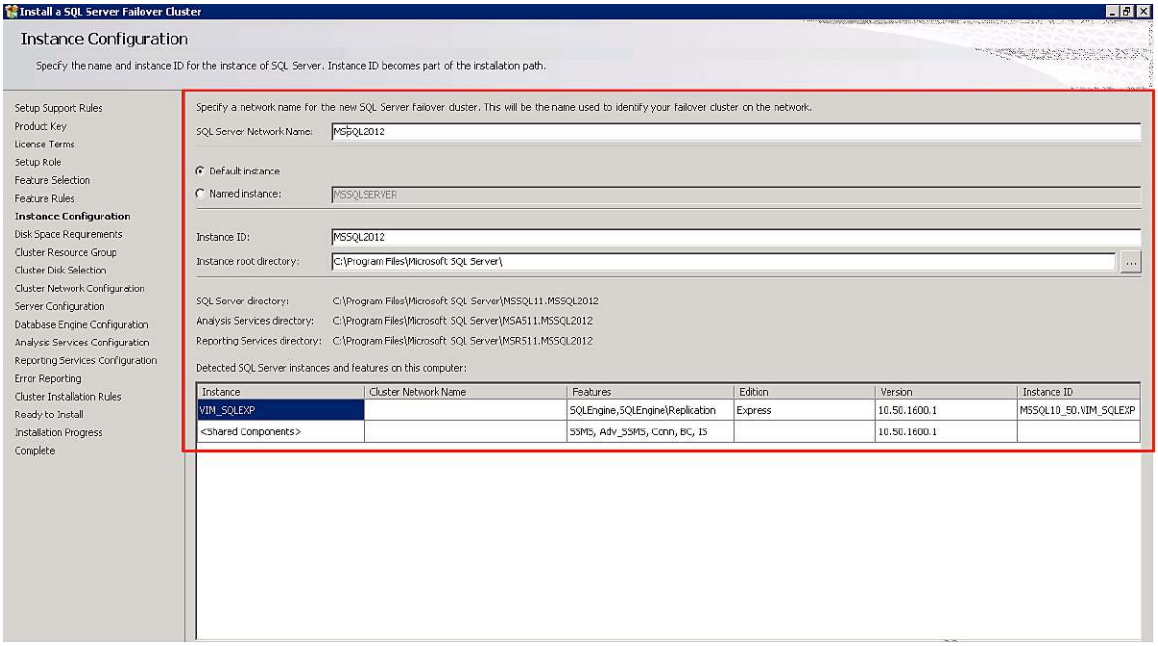

e. In the Cluster Disk Selection window, select **Cluster Disk 1** and **Cluster Disk 3** to store database data and log files. The Cluster Disk 2 resource is already reserved for cluster quorum storage, as shown in Figure 42.

**Figure 42.** Cluster Disk Selection to Store Database Data and Log Files

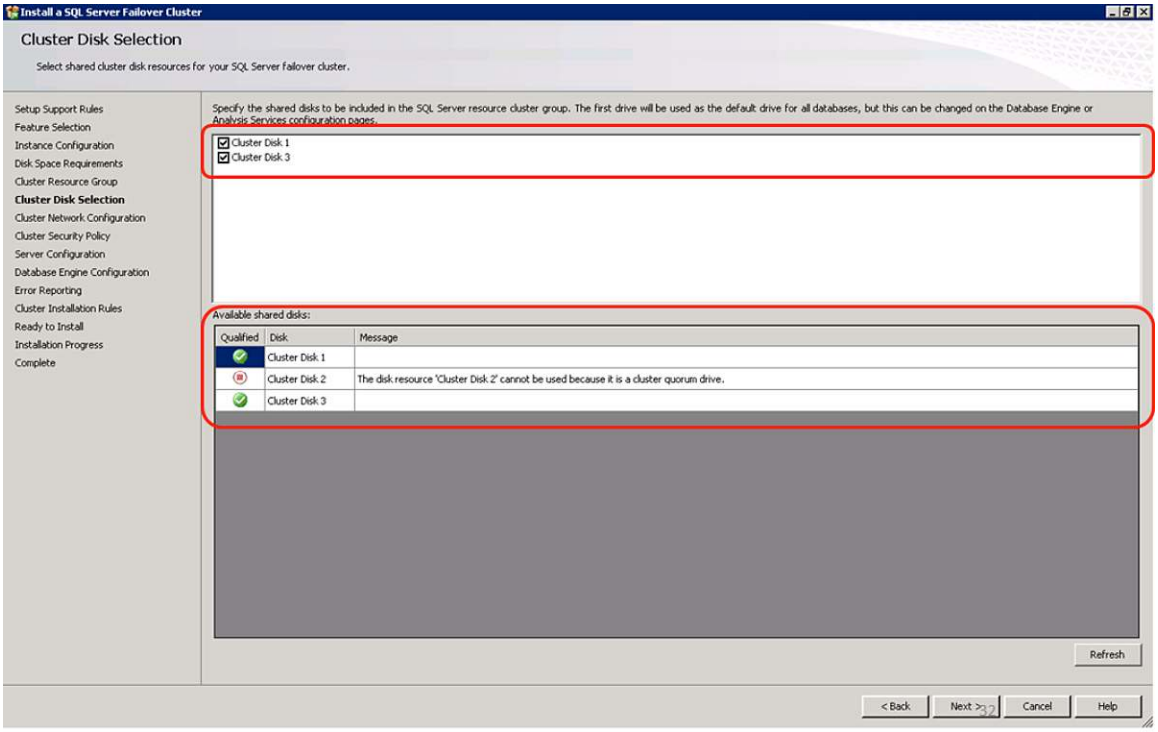


f. In the Cluster Network Configuration window, select the appropriate cluster network subnet from which the Microsoft SQL Server 2012 cluster IP address can be accessed by internal and external clients. This document uses Cluster Network 2, which is configured with the VLAN 809 management IP address, as shown in Figure 43.

**Figure 43.** Cluster Network Configuration in Server Failover Cluster Wizard

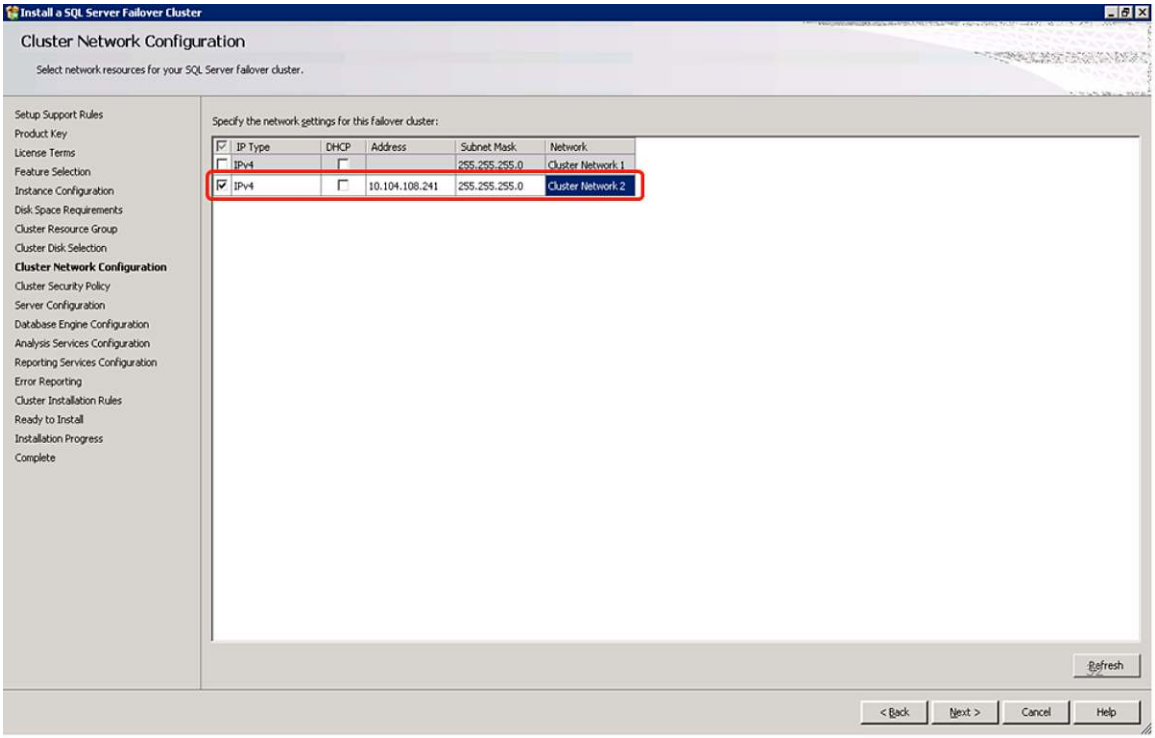


Cluster Network 1 is used to access Microsoft SQL Server 2012 database data and log file storage over the iSCSI network on the Microsoft Windows 2008 R2 **SQLNODE1** and **SQLNODE2** cluster nodes, as shown in Figure 44.

**Figure 44.** Summary of Cluster Network 1 and Cluster Network 2

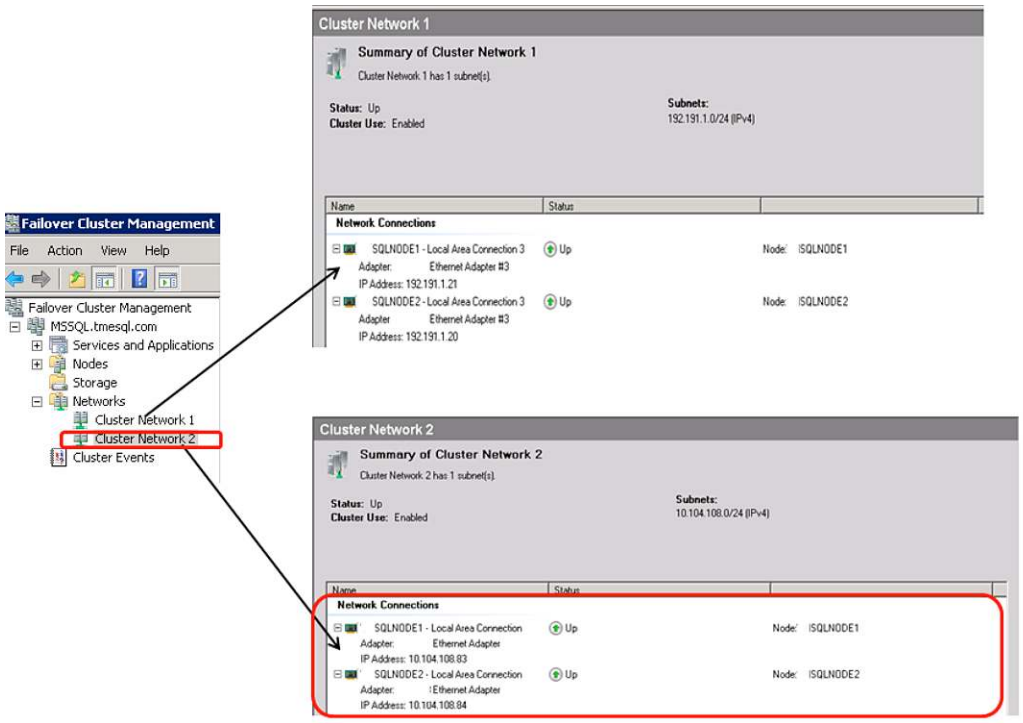


g. In the Database Engine Configuration window, select the appropriate storage drive for storing Microsoft SQL Server 2012 user database data and log files. In this setup, the database data directory is set to E:\MSSQL_DATA, which is mapped to the disk created on NetApp iSCSI storage LUN **MS_SQL_DB_LUN**. 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 45.

**Figure 45.** Database Engine Configuration Showing Database Directory and Log Directory

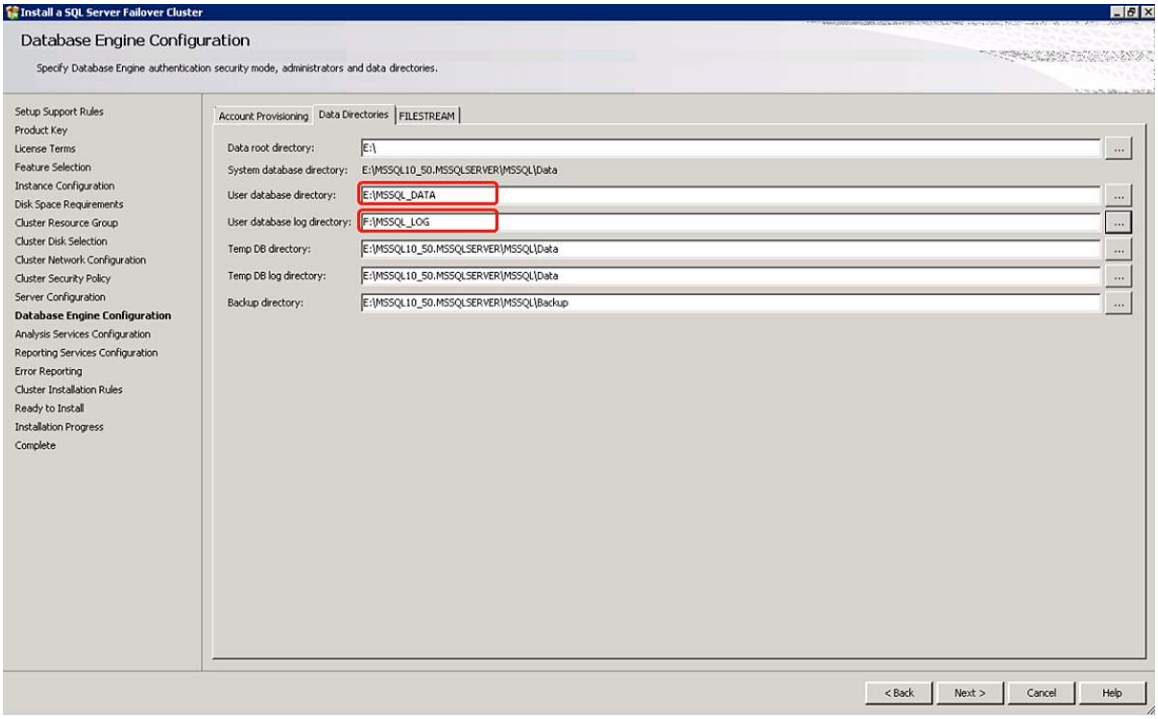


Figure 46 shows the two disks used for database data (E:) and log (F:) files, which are mapped to NetApp iSCSI storage LUNs **MS_SQL_DB_LUN** and **MS_SQL_LOG_LUN**, respectively.

**Figure 46.** Disks for Database Data and Log Files Mapped to NetApp iSCSI Storage

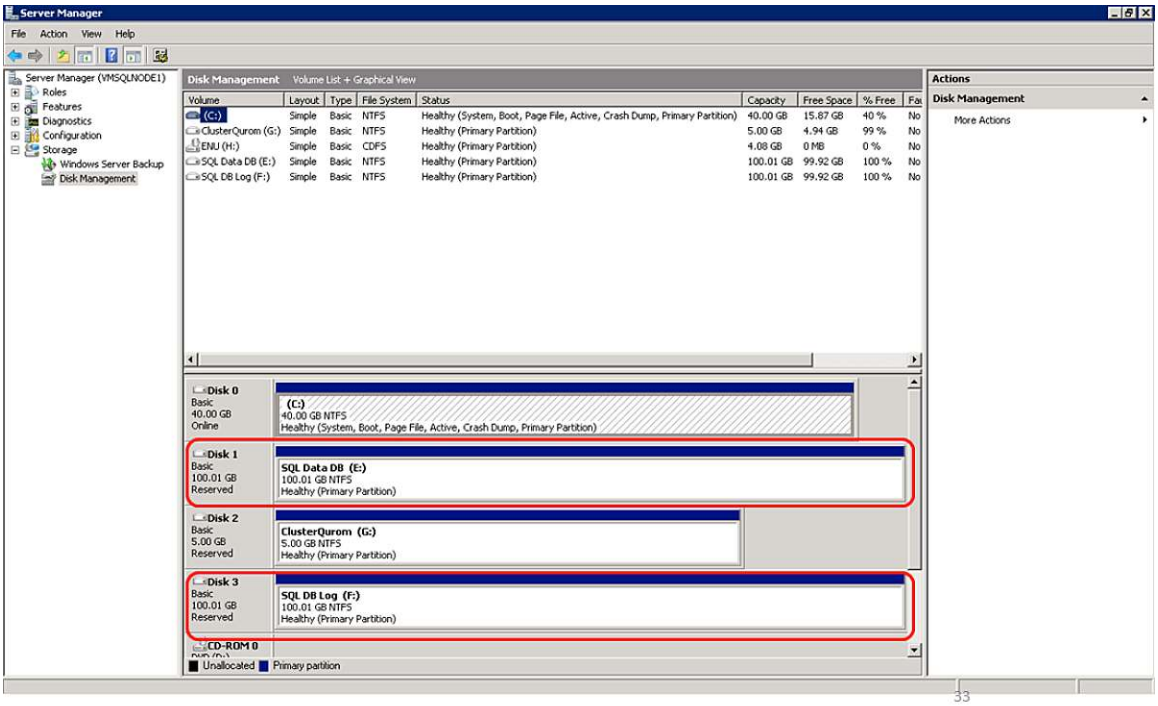


h. Figure 47 shows a summary of the configuration at the end of the Microsoft SQL Server 2012 failover cluster installation setup.

**Figure 47.** Summary of Configuration Details at the End of Failover Cluster Installation Setup

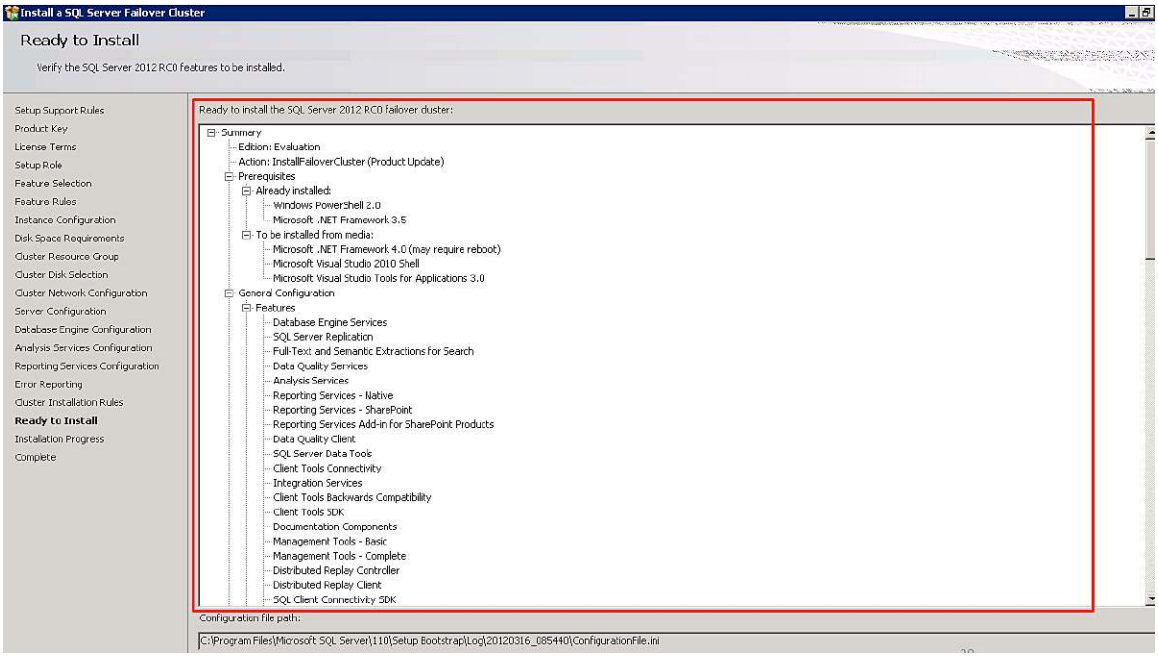


i. Figure 48 shows completion of Microsoft SQL Server 2012 failover cluster installation on the cluster node **SQLNODE1**.

**Figure 48.** Completion of Microsoft SQL Server 2012 Failover Cluster Installation on SQLNODE1

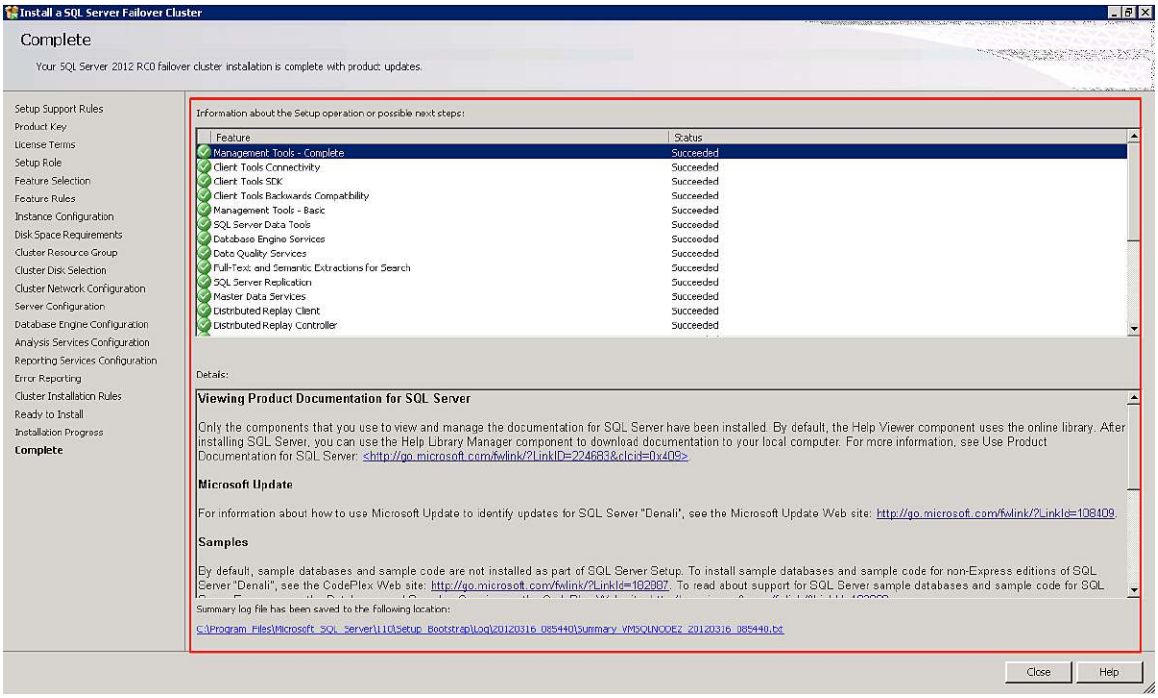


j. To verify that Microsoft SQL Server 2012 failover cluster installation succeeded on the **SQLNODE1** cluster node, launch the Failover Cluster Management console and choose Services and Applications; verify that the **MSSQL2012** instance is added and that the cluster IP address and storage cluster disk status are listed as Online, as shown in Figure 49.

**Figure 49.** Verifying the Storage Cluster Disk Status

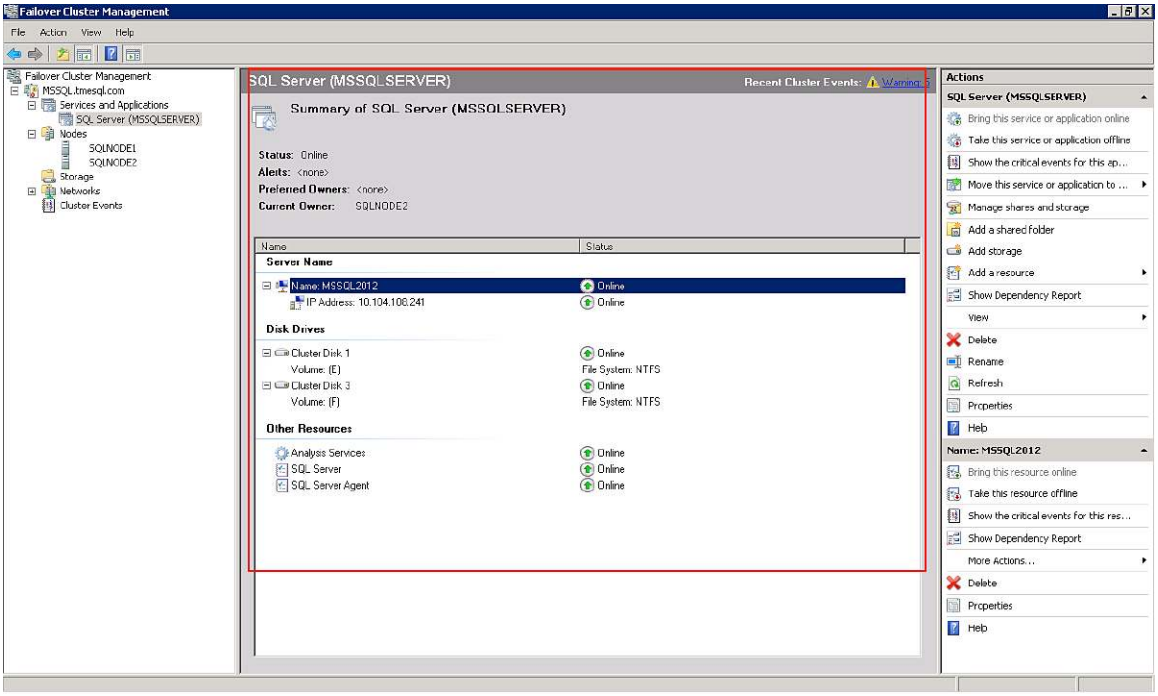


k. To verify that the Microsoft SQL Server 2012 failover cluster instance is accessible on the **SQLNODE1** node, launch Microsoft SQL Server Management Studio and connect to the **MSSQL2012** instance and check the start status, as shown in Figure 50.

**Figure 50.** Verification of Microsoft SQL Server 2012 Failover Cluster Accessibility on SQLNODE1

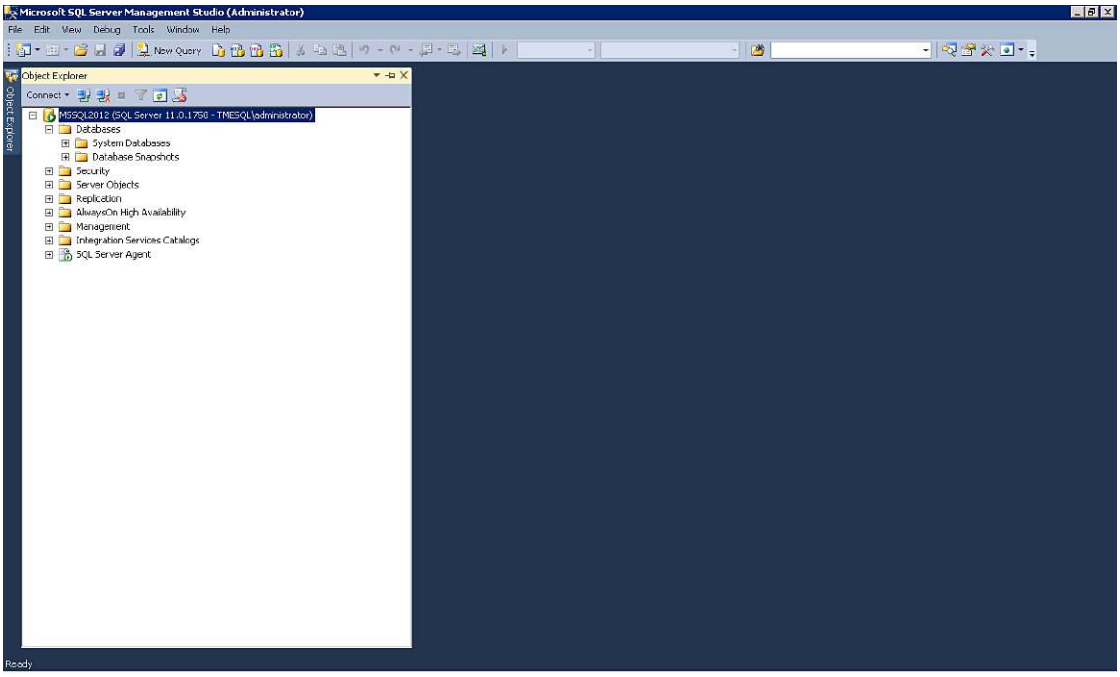

2. Perform the following steps to install Microsoft SQL Server 2012 failover clustering on Microsoft Windows 2008 R2 SP1 cluster host **SQLNODE2** as a secondary node to join the primary failover cluster node installed on **SQLNODE1**.

   a. Copy the Microsoft SQL Server 2012 binaries on the Microsoft Windows 2008 R2 **SQLNODE2** cluster node for installation.

   b. Log in to **SQLNODE2** 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 51.

**Figure 51.** Add Node to Microsoft SQL Server Failover Cluster in SQL Server Installation Center Wizard

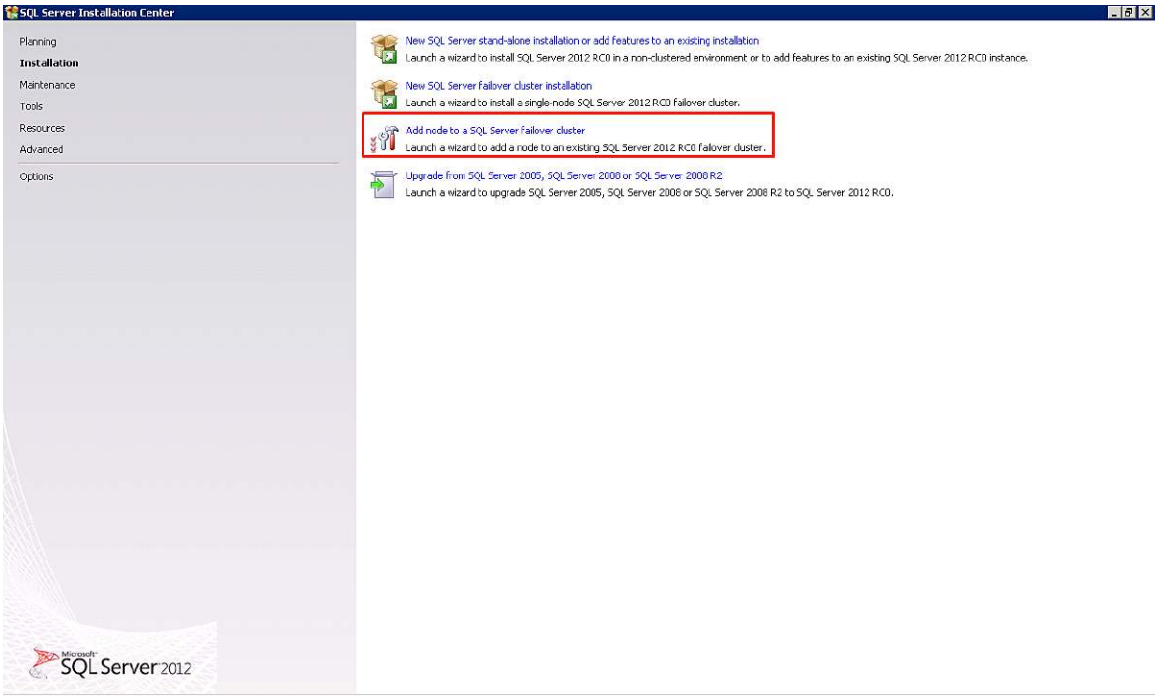


   c. In the Add a Failover Cluster Node window, below the **SQL Server instance name** field, select the **MSSQL2012** instance, which was created during the first step of Microsoft SQL Server 2012 failover cluster deployment, as shown in Figure 52.

**Figure 52.** Cluster Node Configuration in Failover Cluster Node Wizard

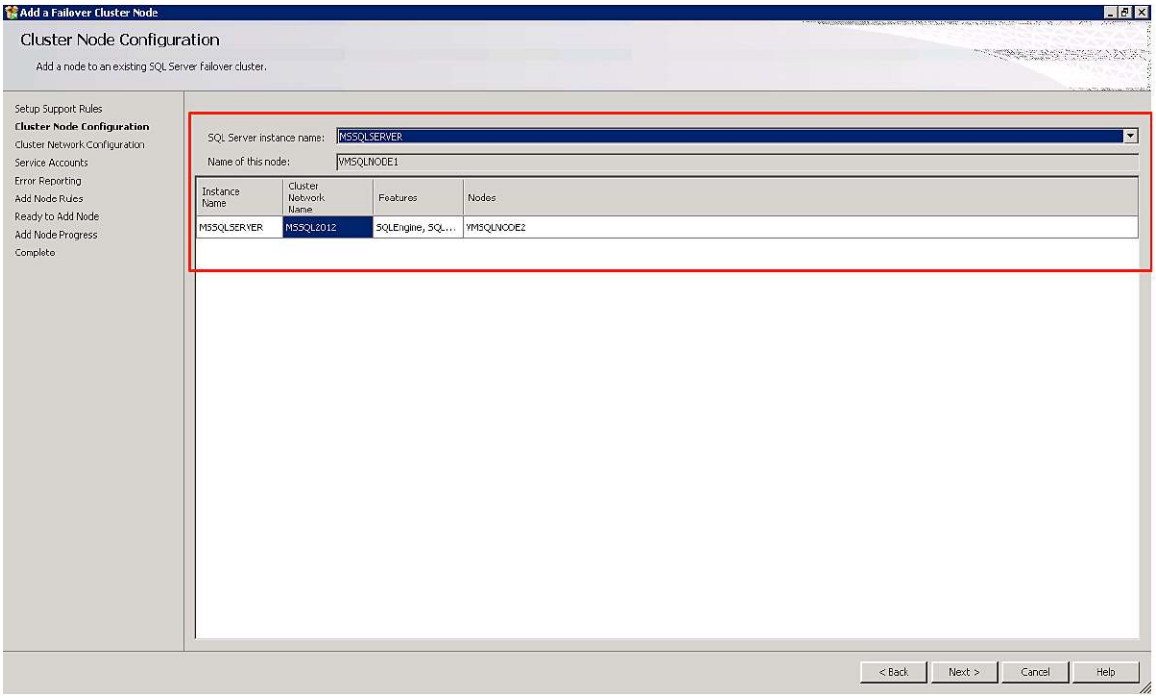


d. In the Cluster Network Configuration window of the wizard, **Cluster Network 2** is automatically selected as the configured IP address during Microsoft SQL Server 2012 failover cluster installation on **SQLNODE1**, as shown in Figure 53.

**Figure 53.** IP Address of Cluster Network 2

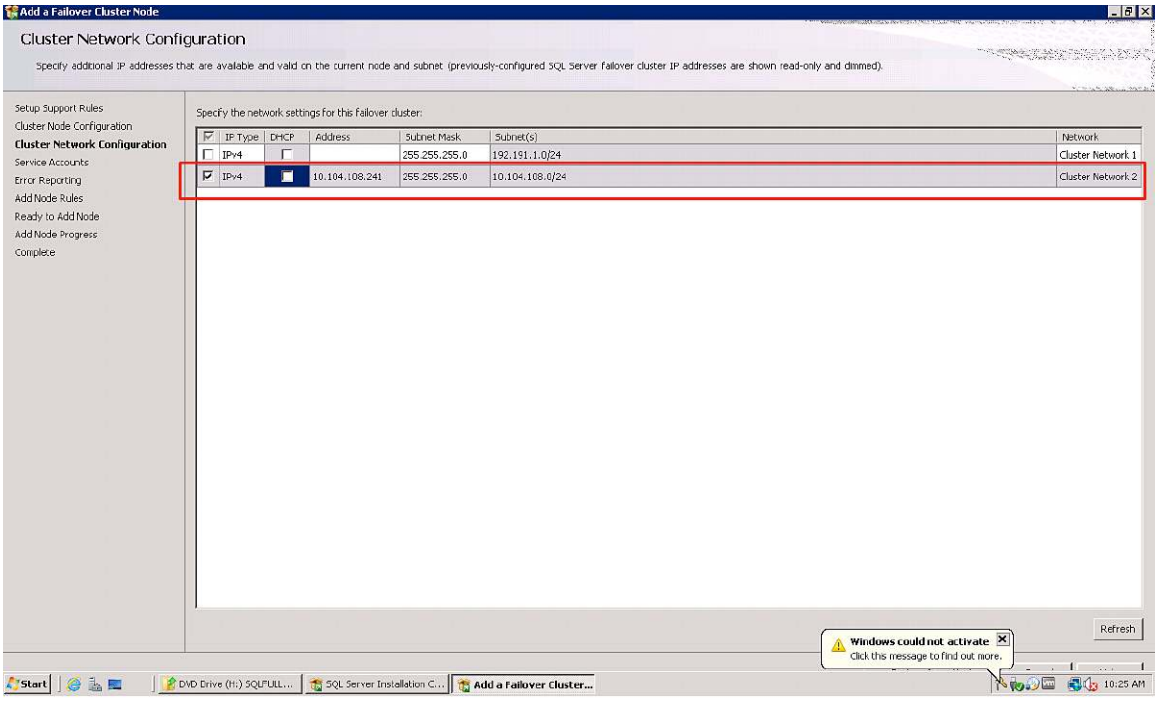

e. Figure 54 shows a summary of the configuration at the end of the installation process.

**Figure 54.** Summary of Cluster Network Configuration

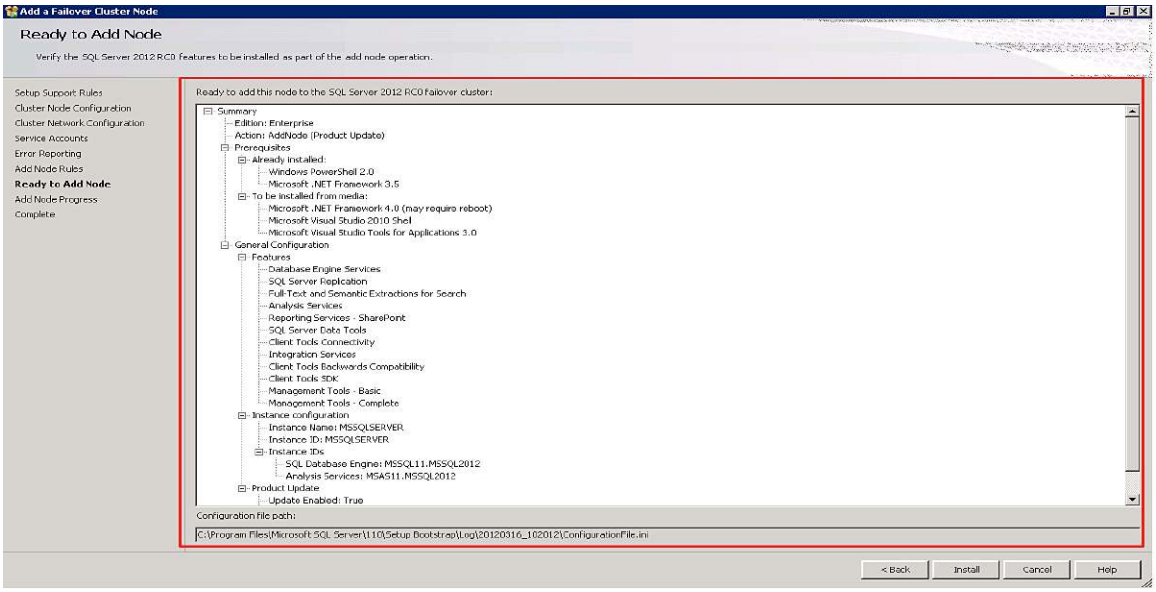


f. Figure 55 shows the completion of Microsoft SQL Server 2012 failover cluster installation on cluster node **SQLNODE2**.

**Figure 55.** Completion of Microsoft SQL Server 2012 Failover Cluster Installation on SQLNODE2

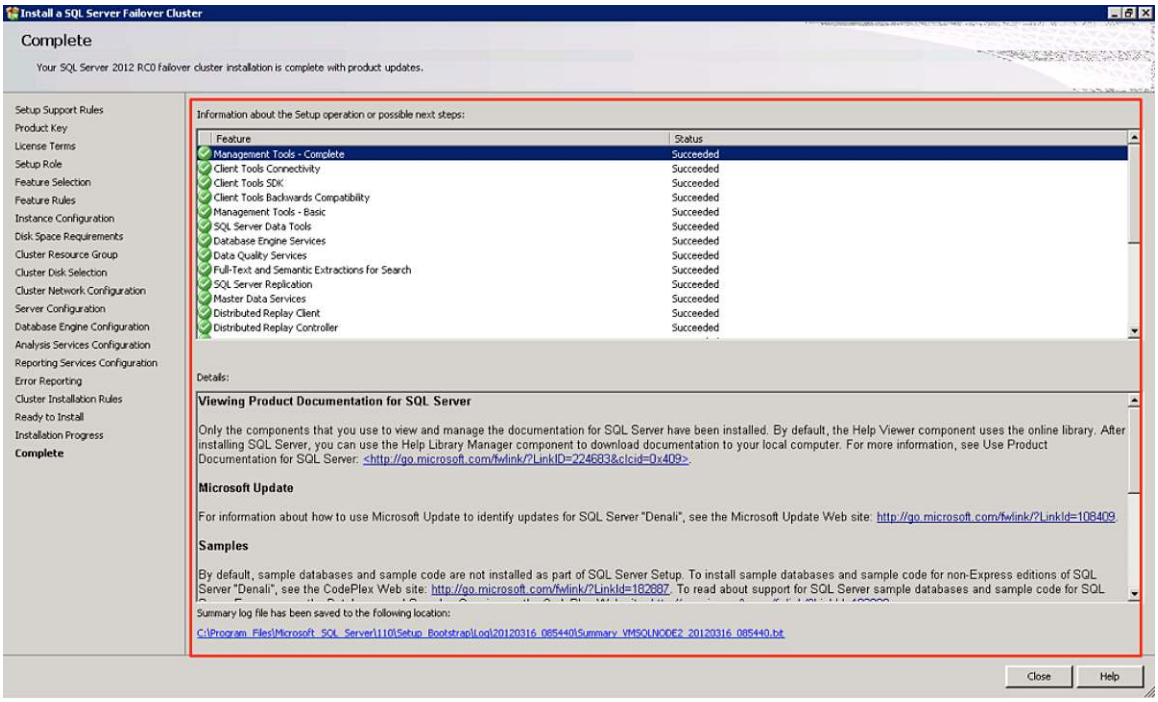

g. Verify that nodes are added to the Microsoft SQL Server 2012 failover cluster on the Microsoft Windows 2008 R2 SP1 host **SQLNODE2** node. Launch the Failover Cluster Management console, and under Services and Applications verify that the **MSSQL2012** instance has been added, that the cluster IP address and storage cluster disks status are listed as “Online”, and that under Nodes, both **SQLNODE1** and **SQLNODE2** are listed, as shown in Figure 56.

**Figure 56.** Verify the Storage Cluster Disk Status and That Both Nodes Are Listed

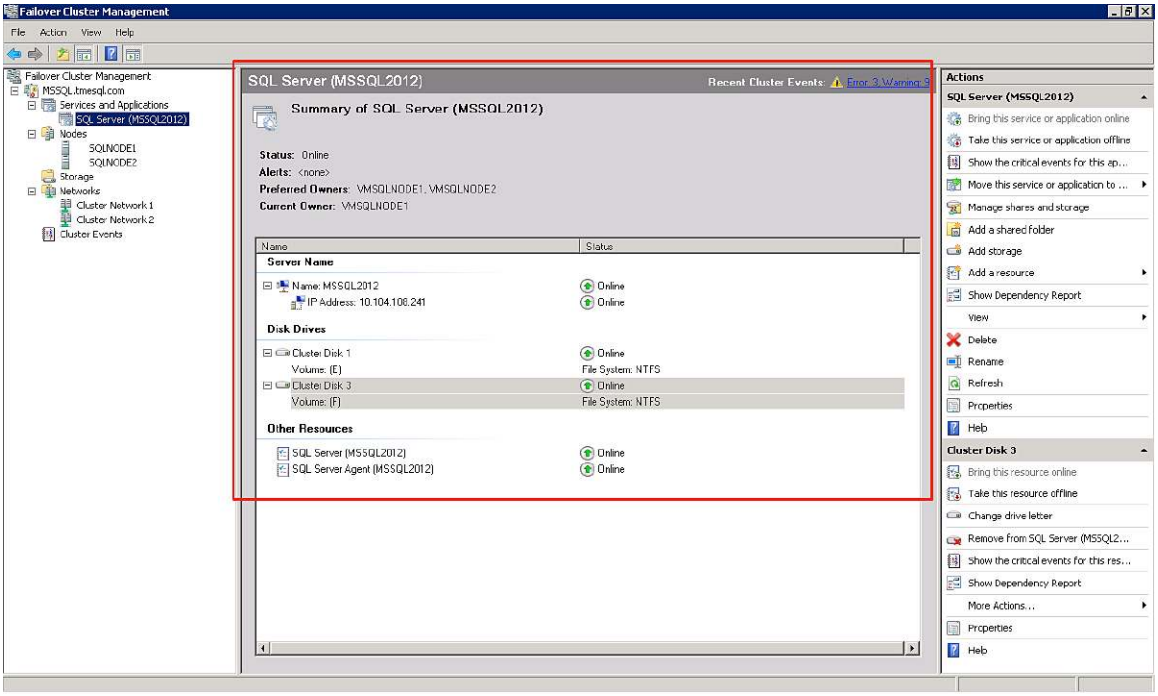


3. Perform the following steps to test failover of Microsoft SQL Server 2012 with the **SQLNODE1** and **SQLNODE2** cluster nodes.

   a. Log in to the **SQLNODE1** cluster node, which currently manages and owns the **MSSQL2012 SQL** instance, as shown in Figure 57. Then shut down the node.

**Figure 57.** SQLNODE1 Managing and Owning the Microsoft SQL Server 2012 Instance

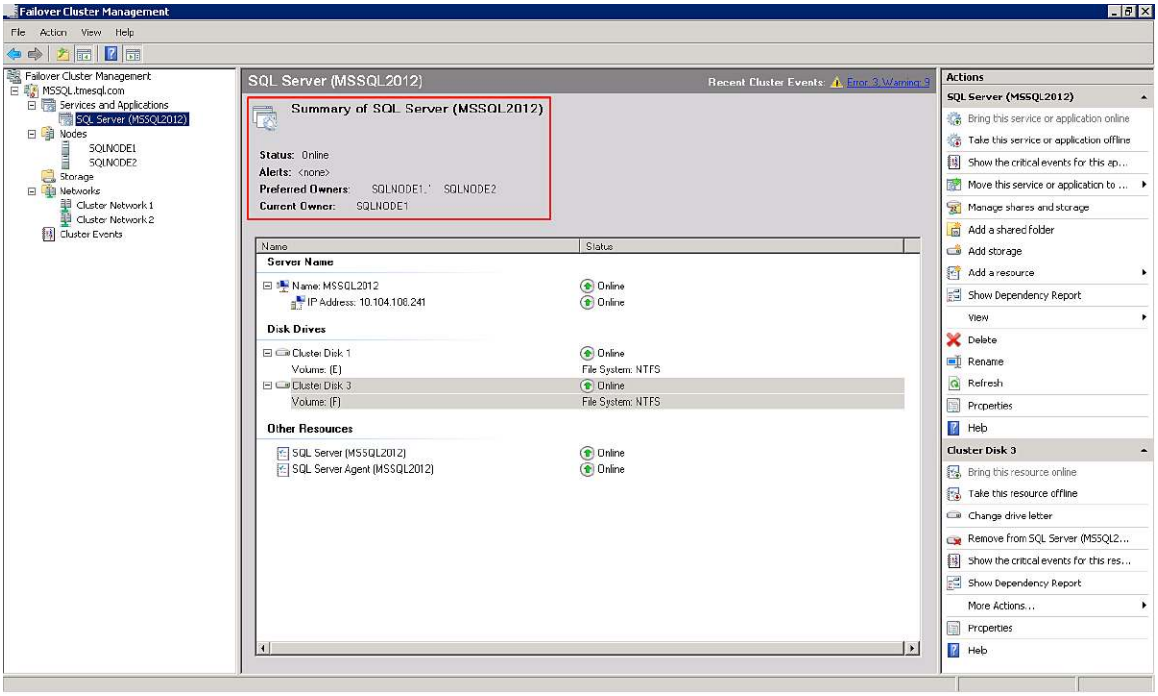


b. After the shutdown of Microsoft Windows 2008 R2 node **SQLNODE1**, the Microsoft Windows failover cluster should be triggered automatically. Subsequently, the Microsoft SQL Server 2012 cluster resource will fail over to the secondary node, **SQLNODE2**. After the failover, **SQLNODE2** will become the current owner of the Microsoft SQL Server 2012 cluster, as shown in Figure 58.

**Figure 58.** Owner of Microsoft SQL Server 2012 Cluster Is SQLNODE2 After Failover

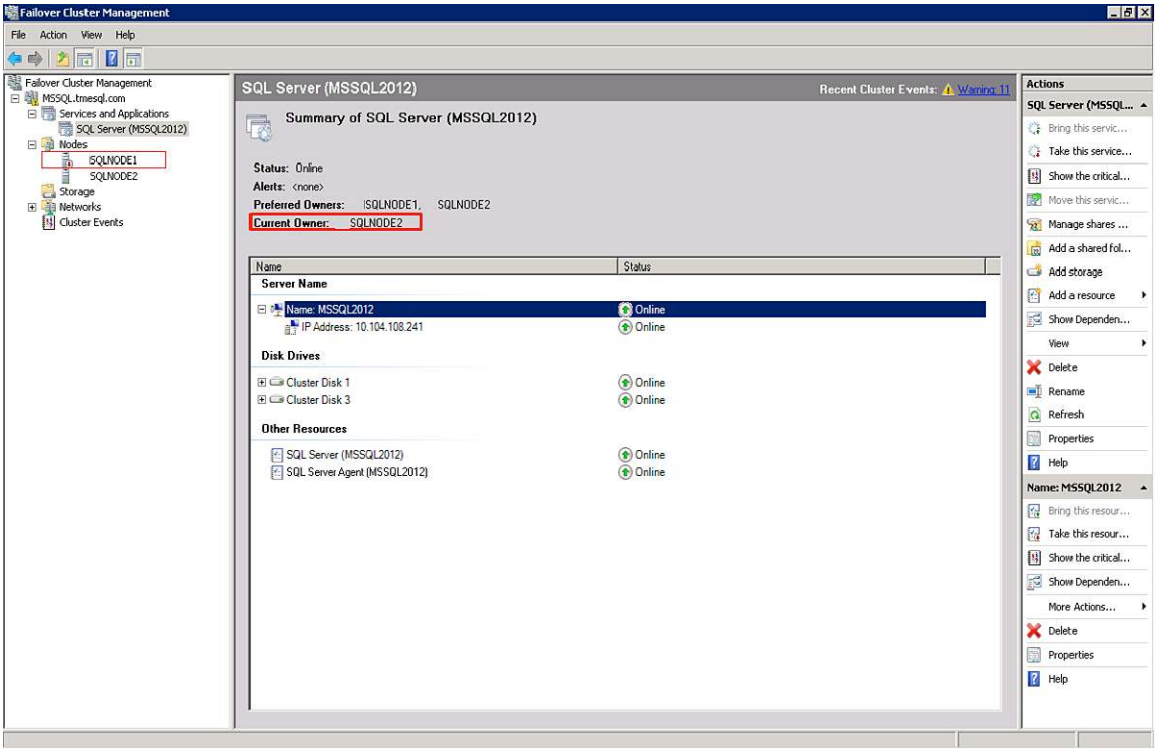


## Conclusion

Microsoft Windows 2008 R2 Enterprise x64 and Microsoft SQL Server 2012 x64 introduce many new features and enhancements. This guide presents the best practices to follow to get the best performance and reliability when deploying Microsoft SQL Server 2012 single-host and failover cluster database servers on bare metal with Microsoft Windows 2008 R2 SP1 using a host-based iSCSI initiator on NetApp iSCSI storage on Cisco UCS blade servers.

## For More Information

See the following documents for additional information about implementation of Microsoft SQL Server 2012 on Microsoft Windows 2008 R2 SP1 with a NetApp iSCSI storage system on Cisco UCS B-Series Blade Servers:

- Microsoft SQL Server 2012 installation 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 UCS hardware and software interoperability matrix:
  http://www.cisco.com/en/US/docs/unified_computing/ucs/interoperability/matrix/r_hcl_B_rel2_0.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 C07-715950-01 04/13