



Microsoft SharePoint 2010 With VMware vSphere 5.0 on FlexPod

A Cisco Validated Design for 100,000 Microsoft Sharepoint Users on
Cisco UCS B-Series Servers


Last Updated: April 25, 2012



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors



SY Abrar

SY Abrar

SY is a Technical Marketing Engineer with the Server Access Virtualization Business Unit (SAVBU) at Cisco. SY has over 10 years of experience in information technology; his focus area includes Microsoft product technologies, server virtualization and storage design. Prior to joining Cisco, Abrar was a Technical Architect at NetApp. Abrar holds a Bachelors degree in Computer Science and he is storage certified and Microsoft certified Technology specialist.



Vadiraja Bhatt

Vadiraja Bhatt

Vadi is a Performance Architect at Cisco, managing the solutions and benchmarking effort on the Cisco Unified Computing System Platform. Vadi has over 17 years of experience in performance and benchmarking the large enterprise systems, deploying mission critical applications. Vadi specializes in optimizing and fine tuning complex hardware and software systems and has delivered many benchmark results on TPC and other industry standard benchmarks. Vadi has six patents in the Database (OLTP and DSS) optimization area.



Rob Barker

Rob Barker

Rob is a Reference Architect with NetApp and focuses on Microsoft SharePoint Server and NetApp's supporting technologies, SnapManager for SharePoint and Data ONTAP. With over 10 years of experience working with SharePoint as both a software developer and administrator, Rob provides a unique perspective and experience for how storage technology and software extensibility work together to build scalable and flexible SharePoint solutions. Prior to joining NetApp, Rob was a Senior Technical Evangelist at Microsoft focused on Microsoft Office and SharePoint and also worked in the consulting field focused on custom SharePoint application development. He is the author of several MS Press books that discuss how to build business applications with Microsoft SharePoint Server.



Alex Fontana

Alex Fontana

Alex is a Senior Solutions Architect at VMware, with a focus on virtualizing Microsoft business critical applications. Alex has worked in the information technology industry for over 12 years during with the last seven years spent designing and deploying Microsoft solutions on VMware technologies. Alex specializes in Microsoft operating systems and applications with a focus on Active Directory, Exchange, and VMware vSphere. Alex is co-author of Virtualizing Microsoft Tier 1 Applications with VMware vSphere 4 and holds VMware, Microsoft and ITIL certifications.



Microsoft SharePoint 2010 With VMware vSphere 5.0 on FlexPod

Introduction

FlexPod is a predesigned, base configuration that is built on the Cisco Unified Computing System (UCS), Cisco Nexus data center switches, NetApp FAS storage components, and a range of software partners. FlexPod serves as an integrated infrastructure stack for all virtualization solutions.

FlexPod configuration may vary depending on customer's needs. The FlexPod architecture is highly modular and a FlexPod unit can be scaled up easily as customer's requirements and demand change. FlexPod can scale up for greater performance and capacity or it can scale out for environments that need consistent, multiple deployments. FlexPod is a baseline configuration, but also has the flexibility to be sized and optimized to accommodate many different use cases.

FlexPod for VMware includes NetApp storage, Cisco networking, the Cisco Unified Computing System, and VMware virtualization software in a single package. This solution is deployed and tested on the defined set of hardware and software. For more information about FlexPod for VMware architecture, go to: <http://www.netapp.com/us/technology/flexpod>.

This Cisco Validated Design demonstrates how enterprises can apply best practices for VMware vSphere, VMware vCenter, Cisco Unified Computing System, Cisco Nexus family switches, NetApp FAS.

Audience

This document is intended to assist solution architects, sales engineers, field engineers and consultants in planning, design, and deployment of Microsoft SharePoint server 2010 hosted on VMware virtualization solutions on the Cisco Unified Computing System. This document assumes that the reader has an architectural understanding of the Cisco Unified Computing System, VMware, Microsoft Office SharePoint 2010 Server, NetApp storage system, and related software.



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

Copyright © 2010 Cisco Systems, Inc. All rights reserved.

Solution Overview

This solution provides an end-to-end architecture with Cisco, VMware, Microsoft, and NetApp technologies. It demonstrates that the Microsoft SharePoint 2010 servers can be virtualized to support 100,000 users with 10 percent concurrency, provide high availability and server redundancy.

The following are the components used for the design and deployment of Microsoft SharePoint:

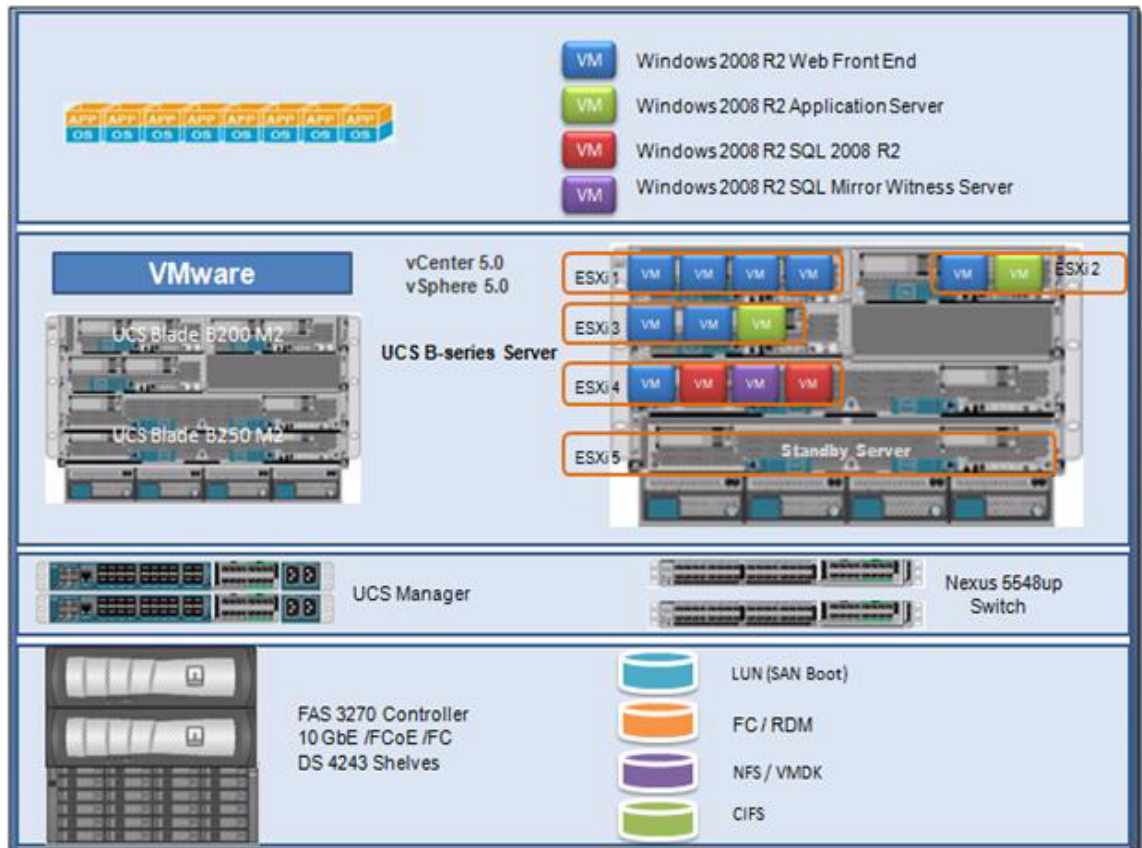
- Microsoft SharePoint 2010 SP1 application
- Cisco Unified Computing System server platform
- VMware vSphere 5 virtualization platform
- Pass Through switching for the virtual servers
- Data Center Business Advantage Architecture
- LAN and SAN architectures
- NetApp storage components
- NetApp OnCommand System Manager 2.0
- Microsoft Network Load Balancer
- Microsoft SQL Mirroring
- VMware DRS
- VMware HA

This solution is designed to host scalable, mixed application workloads. The scope of this Cisco Validated Design is limited to the Microsoft SharePoint 2010 deployment only. The Microsoft SharePoint farm consists of the following virtual machines:

- Eight Load Balanced Web Front End Servers (WFE)
- Two Application servers with redundant services
- Two SQL Servers Mirrored with witness server

Figure 1 illustrates the assignment of application servers to Cisco UCS server blades in a large datacenter.

Figure 1 *FlexPod for VMware SharePoint 2010 Application and Operating System Legend*



Cisco Unified Computing System

Cisco Unified Computing System™ is the first converged data center platform that combines industry-standard, x86-architecture servers with networking and storage access into a single converged system. The system is entirely programmable using unified, model-based management to simplify and speed deployment of enterprise-class applications and services running in bare-metal, virtualized, and cloud-computing environments. The system's unified I/O infrastructure uses a unified fabric to support both network and storage I/O, while Cisco® Fabric Extender technology extends the fabric directly to servers and virtual machines for increased performance, security, and manageability.

The main components of the Cisco Unified Computing System are:

- **Computing**—The system is based on an entirely new class of computing system that incorporates blade servers based on Intel Xeon 5500/5600 Series Processors. The Cisco UCS blade servers offer the patented Cisco Extended Memory Technology to support applications with large datasets and allow more virtual machines per server.
- **Network**—The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

- **Virtualization**—The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access**—The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), and iSCSI. This provides customers with choice for storage access and investment protection. In addition, the server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity, and management for increased productivity.
- **Management**—The system uniquely integrates all system components which enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application programming interface (API) to manage all system configuration and operations.

The Cisco Unified Computing System is designed to deliver the following:


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

Fabric Interconnect

These devices provide a single point for connectivity and management for the entire system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by the Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency regardless of a server or virtual machine's topological location in the system.

Cisco UCS 6248UP Fabric Interconnect


The Cisco UCS 6200 Series Fabric Interconnects support the system's 10-Gbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic using a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blade, rack servers, and virtual machines are interconnected using the same mechanisms. The Cisco UCS 6248UP is a 1-RU fabric interconnect that features up to 48 universal ports that can support 10 Gigabit Ethernet, Fibre Channel over Ethernet, or native Fibre Channel connectivity. The Cisco UCS 6296UP packs 96 universal ports into only two rack units (Figure 2).

Figure 2 *Cisco UCS 6248UP 20-Port Fabric Interconnect*

Cisco UCS 2100 Series Fabric Extenders

These zero-management, low-cost, low-power consuming devices distribute the system's connectivity and management planes into rack and blade chassis to scale the system without complexity. Designed never to lose a packet, Cisco fabric extenders eliminate the need for top-of-rack switches and blade-server-resident Ethernet and Fibre Channel switches and management modules, dramatically reducing infrastructure cost per server.

The Cisco UCS 2208XP Fabric Extenders bring the unified fabric and management planes into Cisco UCS 5108 Blade Server Chassis. Typically deployed in pairs, each device brings up to 80 Gbps of bandwidth to the blade server chassis, for a total of up to 160 Gbps across up to eight servers. Each half-width blade has access to up to 80 Gbps of bandwidth (Figure 3).

Figure 3 *Cisco UCS 2208XP*

Cisco UCS Blade Chassis


The Cisco UCS 5100 Series Blade Server Chassis (Figure 4) is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis.

The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2104XP Fabric Extenders.

A passive mid-plane provides up to 20 Gbps of I/O bandwidth per server slot and up to 40 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

Figure 4 *Cisco UCS Blade Server Chassis (front and back view)*



Cisco UCS Manager

Cisco Unified Computing System Manager provides unified, embedded management of all software and hardware components of the Cisco UCS through an intuitive GUI, a command line interface (CLI), or an XML API. The Cisco UCS Manager provides unified management domain with centralized management capabilities and controls multiple chassis and thousands of virtual machines.

Cisco UCS B200 M2 Blade Server

The Cisco UCS B200 M2 Blade Server (Figure 5) is a half-width, two-socket blade server. The system uses two Intel Xeon 5600 Series Processors, up to 96 GB of DDR3 memory, two optional hot-swappable small form factor (SFF) serial attached SCSI (SAS) disk drives, and a single mezzanine connector for up to 20 Gbps of I/O throughput. The server balances simplicity, performance, and density for production-level virtualization and other mainstream data center workloads.

Figure 5 *Cisco UCS B200 M2 Blade Server*



Cisco UCS B250 M2 Extended Memory Blade Server

The Cisco UCS B250 M2 Extended Memory Blade Server (Figure 6) is a full-width, two-socket blade server featuring Cisco Extended Memory Technology. The system supports two Intel Xeon 5600 Series Processors, up to 384 GB of DDR3 memory, two optional SFF SAS disk drives, and two mezzanine connections for up to 40 Gbps of I/O throughput. The server increases performance and capacity for demanding virtualization and large-data-set workloads with greater memory capacity and throughput.

Figure 6 *Cisco UCS B250 M2 Extended Memory Blade Server*



Cisco UCS Service Profiles

Programmatically Deploying Server Resources

Cisco UCS Manager provides centralized management capabilities, creates a unified management domain, and serves as the central nervous system of the Cisco UCS. Cisco UCS Manager is embedded device management software that manages the system from end-to-end as a single logical entity through an intuitive GUI, CLI, or XML API. Cisco UCS Manager implements role- and policy-based management using service profiles and templates. This construct improves IT productivity and business agility. Now infrastructure can be provisioned in minutes instead of days, shifting IT's focus from maintenance to strategic initiatives.

Dynamic Provisioning with Service Profiles

Cisco Unified Computing System resources are abstract in the sense that their identity, I/O configuration, MAC addresses and WWNs, firmware versions, BIOS boot order, and network attributes (including QoS settings, ACLs, pin groups, and threshold policies) all are programmable using a just-in-time deployment model. The manager stores this identity, connectivity, and configuration information in service profiles that reside on the Cisco UCS 6200 Series Fabric Interconnect. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, enabling flexibility in the use of system resources. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.


Service Profiles and Templates

A service profile contains configuration information about the server hardware, interfaces, fabric connectivity, and server and network identity. The Cisco UCS Manager provisions servers utilizing service profiles. The UCS Manager implements a role-based and policy-based management focused on service profiles and templates. A service profile can be applied to any blade server to provision it with the characteristics required to support a specific software stack. A service profile allows server and network definitions to move within the management domain, enabling flexibility in the use of system resources.

Service profile templates are stored in the Cisco UCS 6200 Series Fabric Interconnects for reuse by server, network, and storage administrators. Service profile templates consist of server requirements and the associated LAN and SAN connectivity. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.

The Cisco UCS Manager can deploy the service profile on any physical server at any time. When a service profile is deployed to a server, the Cisco UCS Manager automatically configures the server, adapters, Fabric Extenders, and Fabric Interconnects to match the configuration specified in the service profile. A service profile template parameterizes the UUIDs that differentiate between server instances.

This automation of device configuration reduces the number of manual steps required to configure servers, Network Interface Cards (NICs), Host Bus Adapters (HBAs), and LAN and SAN switches. Figure 7 shows the Service profile which contains abstracted server state information, creating an environment to store unique information about a server.

Figure 7 **Service Profile**

Cisco Nexus 5548UP Switch

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

Figure 8 **Cisco Nexus 5548UP Switch**

I/O Adapters


The Cisco UCS blade server has various Converged Network Adapters (CNA) options. The UCS M81KR Virtual Interface Card (VIC) option is used in this Cisco Validated Design.

This Cisco UCS M81KR VIC (Figure 9) is unique to the Cisco UCS blade system. This mezzanine card adapter is designed around a custom ASIC that is specifically intended for VMware-based virtualized systems. It uses custom drivers for the virtualized HBA and 10-GE network interface card. As is the case with the other Cisco CNAs, the Cisco UCS M81KR VIC encapsulates fibre channel traffic within the 10-GE packets for delivery to the Fabric Extender and the Fabric Interconnect.

Cisco UCS M81KR VIC (Figure 9) provides the capability to create multiple VNICs (up to 128 in version 1.4) on the CNA. This allows complete I/O configurations to be provisioned in virtualized or non-virtualized environments using just-in-time provisioning, providing tremendous system flexibility and allowing consolidation of multiple physical adapters.

System security and manageability is improved by providing visibility and portability of network policies and security all the way to the virtual machines. Additional M81KR features like VN-Link technology and pass-through switching, minimize implementation overhead and complexity.

Figure 9 *Cisco UCS M81KR VIC*



Cisco VM-FEX Technology

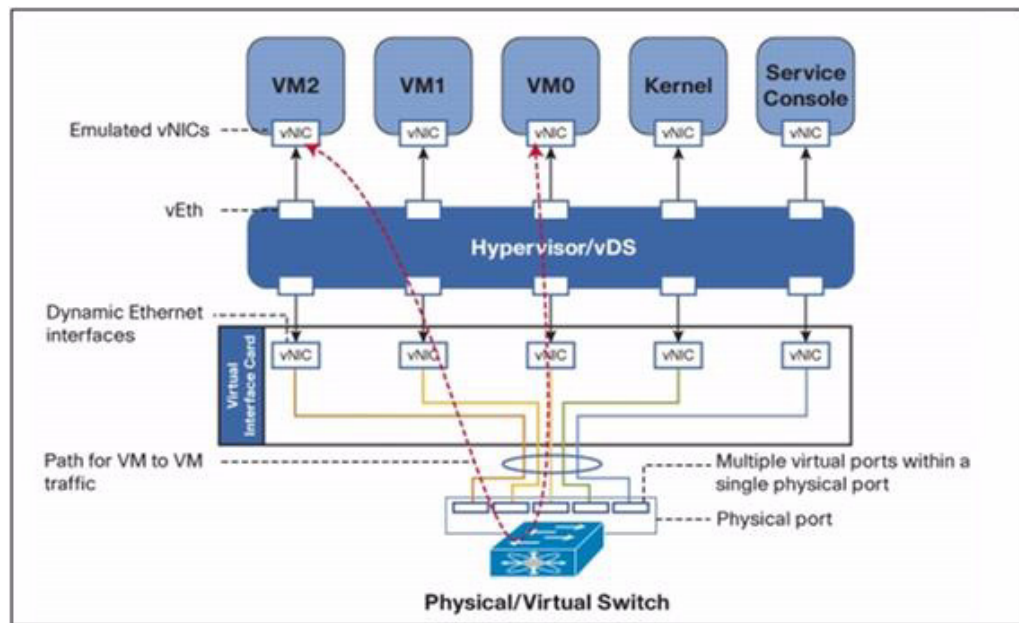
The Virtual Interface Card provides the first implementation of the Cisco VM-FEX technology. The VM-FEX technology eliminates the virtual switch within the hypervisor by providing individual virtual machine virtual ports on the physical network switch. Virtual machine I/O is sent directly to the upstream physical network switch, in this case, the Cisco UCS 6200 Series Fabric Interconnect, which takes full responsibility for virtual machine switching and policy enforcement.

In a VMware environment, the VIC presents itself as three distinct device types: a Fibre Channel interface, a standard Ethernet interface, and a special Dynamic Ethernet interface. The Fibre Channel and Ethernet interfaces are consumed by standard VMware vmkernel components and provide standard capabilities. The Dynamic Ethernet interfaces are not visible to vmkernel layers and are preserved as raw PCIe devices.

Using the Cisco vDS VMware plug-in and Cisco VM-FEX technology, the VIC provides a solution that is capable of discovering the Dynamic Ethernet interfaces and registering all of them as uplink interfaces for internal consumption of the vDS. The vDS component on each host discovers the number of uplink interfaces that it has and presents a switch to the virtual machines running on the host as shown in Figure 10. All traffic from an interface on a virtual machine is sent to the corresponding port of the vDS switch. The traffic is mapped immediately to a unique Dynamic Ethernet interface presented by the VIC. This vDS implementation guarantees the 1:1 relationship with a virtual machine interface and an uplink port. The Dynamic Ethernet interface selected is a precise proxy for the virtual machine's interface.

The Dynamic Ethernet interface presented by the VIC has a corresponding virtual port on the upstream network switch.

Figure 10 Virtual Machine Interface Displaying Its Own Virtual Port on the Physical Switch



Cisco UCS Manager running on the Cisco UCS Fabric Interconnect works in conjunction with the VMware vCenter software to coordinate the creation and movement of virtual machines. Port profiles are used to describe the virtual machine interface attributes such as VLAN, port security, rate limiting, and QoS marking. Port profiles are managed and configured by network administrators using Cisco UCS Manager. To facilitate integration with the VMware vCenter, Cisco UCS Manager pushes the catalog of port profiles into VMware vCenter, where they are represented as distinct port groups. This integration allows the virtual machine administrators to simply select from a menu of port profiles as they create virtual machines. When a virtual machine is created or moved to a different host, it communicates its port group to the Virtual Interface Card. The VIC gets the port profile corresponding to the requested profile from the Cisco UCS Manager and the virtual port on the Fabric Interconnect switch is configured according to the attributes defined in the port profile.

Cisco VM-FEX technology addresses the concerns raised by server virtualization and virtual networking and provides the following benefits:

- **Unified virtual and physical networking**—Cisco VM-FEX technology consolidates the virtual network and physical network into a single switching point that has a single management point. Using Cisco VM-FEX technology, the number of network management points can be reduced by an order of magnitude.
- **Consistent performance and feature availability**—All traffic is controlled at the physical switch, leading to consistent treatment for all network traffic, virtual or physical. Each virtual machine interface is coupled with a unique interface on the physical switch, which allows precise decisions to be made related to the scheduling of and operations on traffic flow from and to a virtual machine.
- **Reduced broadcast domains**—The virtual machine's identity and positioning information is known to the physical switch, so the network configuration can be precise to the specified port.

Modes of Operations for VM-FEX Technology

Cisco VM-FEX technology supports virtual machine interfaces that run in the following modes:

- Emulated mode

The hypervisor emulates a NIC (also referred to as a back-end emulated device) to replicate the hardware it virtualizes for the guest virtual machine. The emulated device presents descriptors, for read and write, and interrupts to the guest virtual machine just as a real hardware NIC device would. One such NIC device that VMware ESXi emulates is the vmxnet3 device. The guest OS in turn instantiates a device driver for the emulated NIC. All the resources of the emulated devices' host interface are mapped to the address space of the guest OS.

- PCIe Pass-Through or VMDirectPath mode

Virtual Interface Card uses PCIe standards-compliant IOMMU technology from Intel and VMware's VMDirectPath technology to implement PCIe Pass-Through across the hypervisor layer and eliminate the associated I/O overhead. The Pass-Through mode can be requested in the port profile associated with the interface using the "high-performance" attribute.

VMware vSphere 5.0

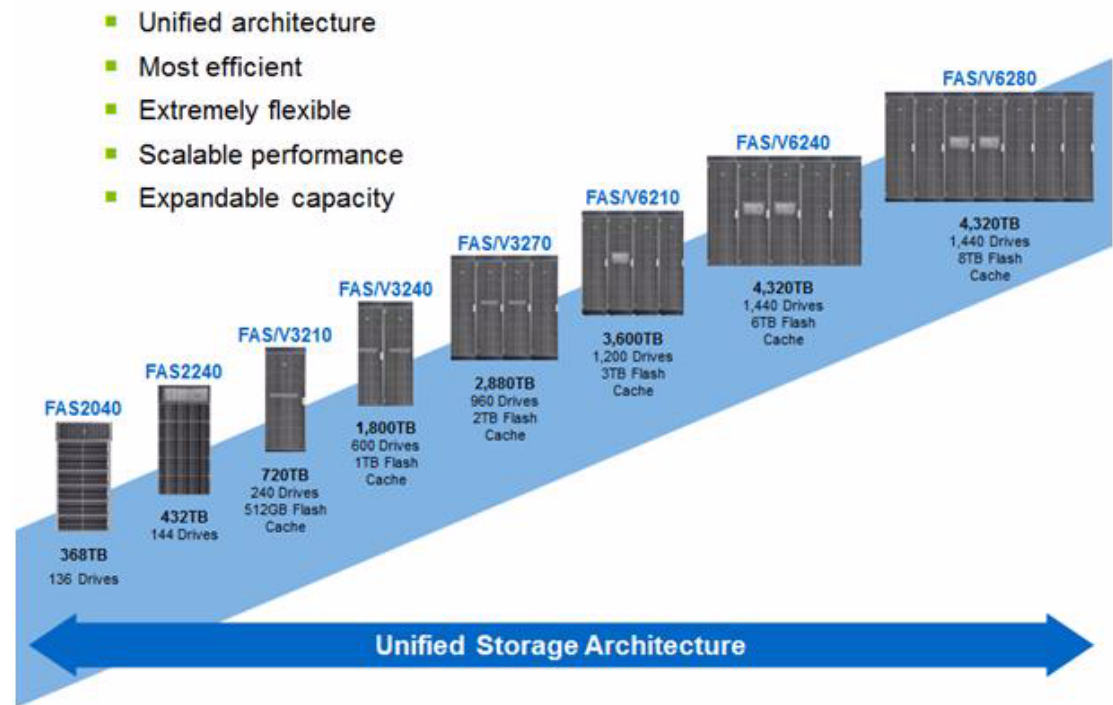
VMware vSphere 5.0 is a next-generation virtualization solution from VMware which builds upon ESXi 4 and provides greater levels of scalability, security, and availability to virtualized environments. vSphere 5.0 offers improvements in performance and utilization of CPU, memory, and I/O. It also offers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

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

NetApp Storage Technology and Benefits

Data ONTAP® is the fundamental NetApp software platform that runs on all NetApp storage systems. Data ONTAP is a highly optimized, scalable operating system that supports mixed NAS and SAN environments and a range of protocols, including Fiber Channel, iSCSI, FCoE, NFS, and CIFS. The platform includes a patented file system, Write Anywhere File Layout (WAFL), and storage virtualization capabilities. By leveraging the Data ONTAP platform, the NetApp Unified Storage Architecture offers the flexibility to manage, support, and scale to different business environments by using a common knowledge base and tools. This architecture allows users to collect, distribute, and manage data from all locations and applications at the same time. This allows the investment to scale by standardizing processes, cutting management time, and increasing availability. Figure 11 shows the different NetApp Unified Storage Architecture platforms.

Figure 11 *NetApp Unified Storage Architecture Platforms*



The NetApp storage hardware platform used in this solution is the FAS3270. The FAS3200 series is an ideal platform for primary and secondary storage for a Microsoft SharePoint 2010 SP1 server deployment.

An array 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 makes use of a subset of these tools and enhancements.

RAID-DP

RAID-DP® is NetApp's implementation of double-parity RAID 6, which is an extension of NetApp's original Data ONTAP WAFL® RAID 4 design. Unlike other RAID technologies, RAID-DP provides the ability to achieve a higher level of data protection without any performance impact while consuming a minimal amount of storage. For more information on RAID-DP® go to:

<http://www.netapp.com/us/products/platform-os/raid-dp.html>.

Thin Provisioning and FlexVol

Thin provisioning is a function of NetApp FlexVol® which allows storage to be provisioned just like traditional storage. However, the storage is not consumed until the data is written (just-in-time storage). Use NetApp ApplianceWatch in Microsoft's SCOM to monitor thin provisioned LUNs to increase disk efficiency. Microsoft recommends a 20 percent growth factor above the database size and a 20 percent free disk space in the database LUN which on disk is over 45 percent free disk space. For more information on ApplianceWatch go to:

http://blogs.netapp.com/storage_nuts_n_bolts/2010/06/netapp-appliancewatch-for-microsofts-scom-and-scvmhyper-v.html.

NetApp OnCommand System Manager 2.0

System Manager is a powerful management tool for NetApp storage. The System Manager tool helps administrators to manage single NetApp® storage systems as well as clusters quickly and easily.

Some of the benefits of the System Manager Tool are the following:

- Easy to install
- Easy to manage from a browser
- Does not require storage expertise
- Increases storage productivity and response time
- Cost effective
- Leverages storage efficiency features such as:
 - Thin provisioning
 - Compression
 - Deduplication

NetApp Deduplication

NetApp Deduplication technology leverages NetApp WAFL block sharing to perform protocol-agnostic data-in-place deduplication as a property of the storage itself. With legacy versions of Exchange Server the deduplication rate was one to five percent, while with Exchange Server 2010 the deduplication rate is 10 to 35 percent. In a virtualized environment, it is common to see 90 percent deduplication rate on the application and operating system data.

NetApp Snapshot

NetApp Snapshot™ technology provides zero-cost, near-instantaneous backup, point-in-time copies of the volume or LUN by preserving the Data ONTAP WAFL consistency points (CPs).

Creating Snapshot copies incurs minimal performance effect because data is never moved, as it is with other copy-out technologies. The cost for Snapshot copies is at the rate of block-level changes and not 100 percent for each backup as it is with mirror copies. Using Snapshot can result in savings in storage cost for backup and restore purposes and opens up a number of efficient data management possibilities.

SnapManager for Microsoft SharePoint Server

SnapManager for Microsoft SharePoint Server reduces storage costs and manages the growth of the SharePoint environment efficiently.

SnapManager provides the ability to easily migrate and store SharePoint files on Common Internet File Systems (CIFS)/Server Message Block (SMB) shares outside of Microsoft SQL Server® to improve scalability of large SharePoint 2007 or 2010 deployments.

With its browser-based user interface, SnapManager enables you to automate your data backup processes and other administrative functions. You can also plan and implement a highly reliable disaster recovery strategy. Together with SnapMirror®, you can simplify remote replication of SharePoint Server data so it can be recovered rapidly in the event of a disaster.

NetApp Strategy for Storage Efficiency

As seen in the previous section on technologies for storage efficiency, NetApp's strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core Data ONTAP operating system and the WAFL file system. Unlike its competitors' technologies, NetApp's technologies surrounding its FAS and V-Series product line have storage efficiency built into their core. Customers who already have other vendors' storage systems and disk shelves can still leverage all the storage saving features that come with the NetApp FAS system simply by using the NetApp V-Series product line. This is again in alignment with NetApp's philosophy of storage efficiency because customers can continue to use their existing third-party storage infrastructure and disk shelves, yet save more by leveraging NetApp's storage-efficient technologies.

Microsoft SharePoint 2010 SP1

Microsoft SharePoint 2010 is an extensible and scalable web-based platform consisting of tools and technologies that support the collaboration and sharing of information within teams, throughout the enterprise and on the web. The total package is a platform on which one can build business applications to help better store, share, and manage information within an organization. Microsoft SharePoint turns users into participants, allowing users to easily create, share and connect with information, applications and people. SharePoint 2010 provides all the good features present in the earlier versions of the product along with several new features and important architectural changes to improve the product.

Three-Tier Role Based Architecture


The three-tier role based architecture of Microsoft SharePoint 2010 Farm includes Web Server Role, Application Server Role and Database Server Role (Figure 12).

Web Server Role—SharePoint web server is responsible to host web pages, web services, and web Parts that are necessary to process requests served by the farm. Also, the server is responsible to direct requests to the appropriate application servers.

Application Server Role—SharePoint Application Server is associated with Services, where each service represents a separate application service that can potentially reside on a dedicated application server. Services with similar usage and performance characteristics can be grouped on a server. The grouped services can be then be scaled out onto multiple servers.

Database Server Role—SharePoint Database can be categorized broadly by their roles as Search Database, Content Database and Service Database. In larger environments, SharePoint databases are grouped by roles and deployed onto multiple database servers.

Figure 12 Three-Tier Architecture



Advantages of Three-Tier Architecture

Among the many major benefits three-tier architecture provides, a few are described below:

- **Maintainability**—Three-tier architecture follows a modular approach; each tier is independent of each other. Thus each tier can be updated without affecting application as a whole.
- **Scalability**—Scalability is a major benefit to incorporate three-tier architecture, we can scale each tier as and when required without dependence on the other tiers. I.e. independently scaling each tier. For instance in scaling of web servers, provisioning of servers at multiple geographical locations enables faster end user response time, avoiding high network latency. Another aspect is scaling of application servers which require high computing resources, thus a farm of clustered application servers can provide on demand application performance.
- **Availability and Reliability**—Applications using three tier approach can exploit three-tier modular architecture to scale components and servers at each tier, thus providing redundancy, avoiding single point of failure and in-turn improving availability of overall system.

Microsoft SharePoint 2010 SP1 Sizing Considerations

In the context of SharePoint, the term “farm” is used to describe a collection of one or more SharePoint servers and one or more SQL servers that together provide a set of basic SharePoint services bound together by a single Configuration Database in SQL. Farm marks the highest level of SharePoint administrative boundary.

SharePoint 2010 SP1 can be configured as a small, medium, or large farm deployment. Remember, the topology service provides you with an almost limitless amount of flexibility, so you can tailor the topology of your farm to meet your specific needs.

A key concept of any SharePoint design is “right sizing” the SharePoint implementation to meet the needs of the organization.

Classic errors in design include both over building and under building the SharePoint environments. Over building can result in an overly complex SharePoint environment, with upwards of a dozen servers (or even more), often with an overwhelming number of features enabled (such as managed metadata, workflows, forms, business intelligence and other SharePoint Enterprise features) that exceed IT's abilities to support them. The feature set can also exceed the user community's skills, especially when

sufficient training is not offered, resulting in the impression that SharePoint is “too complicated” or “never works right.” Under building can yield equally painful results that include slow page loads, time outs during uploads or downloads, system outages, or overly simplistic feature sets that frustrate users.

Some of the common questions related to sizing such as the number of servers required to create a farm, the size of a farm, the capacity the farm should possess and so on can be answered by proper sizing. Analyzing the demand characteristics that the solution is expected to handle i.e. understanding both the workload characteristics such as number of users, concurrent users at peak time and most frequently used operations and the dataset characteristics such as content size and distribution is necessary to do a proper sizing.

The farm used in this solution has eight web front end servers, two application server and a mirrored database server. The farm serves various tiers fulfilling the realistic enterprise needs while being flexible, scalable, and maintainable. The farm is failure proof which makes it a reliable farm.

The following sections briefly describe the sizing aspects of each tier of the three-tier role based SharePoint farm.

Web Front-End Server

WFE servers form the SharePoint connection point for clients that request content or services. All client requests results in some load on the WFE servers as the WFE servers render pages before returning requested pages to a browser. WFE servers do not require large quantities of disk storage, but rely heavily on processor and memory for performance.

Table 1 describes the processor and memory load characteristics for WFE servers.

Table 1 Processor and Memory Load Characteristics for the WFE Server

Service Application	CPU Load	Memory Load
SharePoint Foundation Service	High	High
Timer Service	Very High	Very High
Logging Service	Very High	Very High
Sand Boxed Solution	Low	Low
Workflow Capabilities	High	High

Application Server

Different service applications have different workload profiles, but specific servers can be dedicated to specific service applications. Scale out can be achieved by assigning multiple servers for a specific service application.

Most application services do not require local storage on the application server. The main hardware requirements for application servers are processor and memory.

Table 2 describes the processor and memory load characteristics of application servers.

Table 2 Processor and Memory Load Characteristics for the Application

Server

Service Application	CPU Load	Memory Load
SharePoint Foundation Service	None	None
Central Administration Service	Medium	Medium
Timer Service	Very High	Very High
Logging Service	None	None
User Profile Service	Very High	Very High
Word Viewing Service	High	High
SharePoint Search Service	Very High	Very High

The following are the best practices for Web Front-End and Application Servers:

- Have more than one Web server in the farm that hosts the Microsoft SharePoint Foundation Web Application which allows end users to access the SharePoint sites and data.
- Make the choice of service applications which go on the Web server considering that these applications may impact the overall performance of the Web servers and impact the performance perception of end users.
- Evaluate the need of the Application server based service applications for the SharePoint solution. Install and enable only those service applications which will enable the organization to meet specific business and technology goals and keeping in mind the skill sets of IT and the user population to both support and use these features.
- Enabling some SharePoint 2010 Enterprise features, such as Excel Services or PerformancePoint can put strain on the application and SQL servers and require additional configuration such as the installation of SQL Server Analysis Services for PerformancePoint.
- Installing “companion products” such as Project Server 2010 or Office Web Applications may justify the addition of other application servers to the farm.

Microsoft SharePoint 2010 SP1 Search

Microsoft SharePoint 2010 SP1 Search service offers significant benefits for users, but places a large workload burden on the farm. When considering the farm performance, search performance must be considered specifically in the context of the farm.

The following are the components of the search servers:

- **Crawl component**—The crawl component crawls and indexes content in the SharePoint content databases and in other types of storage repositories. The crawl role builds the index and submits index updates to the search query role. Crawl components aggressively use CPU bandwidth. Memory is not critical for the crawl component.
- **Query component**—The query component responds to the user's search requests. When users enter a search in a SharePoint site, SharePoint 2010 SP1 submits the query to a server that hosts the query role and returns a result set. All servers that host the query role have a copy of the index that the crawl role generates.

In the scope of this performance study, the Microsoft SharePoint 2010 Web Front-End and application virtual servers hosted ESXi 5 serves on Cisco UCS B200 M2. The Cisco UCS B200 M2 Blade Server balances simplicity, performance, and density for production-level virtualization and other mainstream

data-center workloads. The server is a half-width, 2-socket blade server with substantial throughput and scalability. The UCS B200 M2 server extends the capabilities of the Cisco Unified Computing System. It uses Intel's latest Xeon 5600 series multi-core processors to deliver enhanced performance and efficiency.

Database Server

All the data including content, configuration and metadata are stored in the SQL server. Not all service applications affect database servers, because only some of them require databases. However, storage access times and storage capacity are a key requirement for this role (Table 3).

Table 3 *Processor and Memory Load Characteristics for the Database*

Service Application	CPU Load	I/O Load	Storage
SharePoint Foundation Service	Very High	High	High
Timer Service	None	None	None
Logging Service	Very High	High	High
User Profile Service	High	High	Very High
Word Viewing Service	None	None	None
Sandboxed Solution	None	None	None
Workflow Capabilities	None	None	None
SharePoint Search Service	High	High	High

In the default configuration, SharePoint 2010 stores data by uploading it to a SharePoint site in a SQL Server database, with SQL Server 2008 R2 being the recommended version. Since the process of uploading a document to the SQL database is not as efficient as simply storing a file on a file share, optimizing the I/O on the SQL server is very important.

When using a Cisco UCS-based environment, an organization can choose to create physical SQL 2008 servers or to create a VMware virtual SQL 2008 environment and implement a server cluster or high-availability mirror ("HA mirror").

In the scope of this performance study, Database Virtual servers hosted ESXi 5 serves on Cisco UCS B250 M2. The Cisco UCS 250 Blade servers have two Intel Xeon 5600 Series processors which adjust server performance according to application needs, and have DDR3 memory technology with memory scalable up to 384 GB for demanding virtualization and large dataset applications, or have a more cost-effective memory footprint for less demanding workloads. Two dual-port mezzanine cards for up to 40 Gbps of I/O per blade are required. The options for the Mezzanine card include either the Cisco UCS VIC M81KR Virtual Interface Card, a converged network adapter (Emulex or QLogic compatible), or a single 10GB Ethernet Adapter.

Microsoft SharePoint 2010 SP1 Design Considerations

The design best practices for designing SharePoint environments are arrived at to show the many advantages to organizations choosing the Cisco platform and are applicable for organization of all sizes. There are many options that need to be considered in the areas of SQL Server design, SharePoint "front

end" server configurations, configuration of the software and the hardware that will run the SharePoint 2010 farm(s), when designing an enterprise class SharePoint 2010 environment. Organizations must look ahead when planning for the management and governance of the SharePoint environment. This ensures that the environment continues to offer acceptable levels of performance and reliability.

An obvious immediate benefit with Cisco is a single trusted vendor providing all the components needed for a SharePoint farm.

The additional capabilities offered by Cisco Unified Computing System include the following:

- **Dynamic provisioning and service profiles**—Cisco UCS Manager supports service profiles, which contain abstracted server states, creating a stateless environment. It implements role-based and policy-based management focused on service profiles and templates. These mechanisms fully provision one or many servers and their network connectivity in minutes, rather than hours or days. This can be very valuable in SharePoint environments, where new servers may need to be provisioned on short notice, or even whole new farms for specific development activities.
- **Embedded multi-role management**—Management is embedded in the Fabric Interconnects, with all attached systems handled as a single, redundant management domain. Cisco UCS Manager controls all aspects of system configuration and operation, eliminating the need to use multiple, separate element managers for each system component. The result is a reduction in management modules and consoles while harmonizing data center roles for high productivity.
- **Cisco VN-Link virtualization support and virtualization adapter**—Virtual machines get virtual links that allow virtual machines to be managed in the same manner as physical links. Now virtual links can be centrally configured and managed without the complexity of traditional systems that interpose multiple switching layers in virtualized environments. I/O configurations and network profiles move along with virtual machines, helping increase security and efficiency while reducing complexity. The adapter also helps improve performance and reduces network interface card (NIC) infrastructure.
- **Cisco UCS Virtual Interface Cards**—Cisco offers a variety of adapter cards designed for use with Cisco UCS B-Series Blade Servers. All Cisco UCS B-Series Network Adapters allow for the reduction of the number of required NICs and HBAs, are managed via Cisco UCS Manager software, feature dual 10 Gigabit connections to the chassis mid-plane, and can be used in a redundant configuration with two Fabric Extenders and two Fabric Interconnects.
- **Cisco Unified Fabric and Fabric Interconnects**—The Cisco Unified Fabric leads to a dramatic reduction in network adapters, blade-server switches, and cabling by passing all network and storage traffic over one cable to the parent Fabric Interconnects, where it can be processed and managed centrally. This improves performance and reduces the number of devices that need to be powered, cooled, secured, and managed. The 6200 series offer key features and benefits, including:
 - High performance Unified Fabric with line-rate, low-latency, lossless 10 Gigabit Ethernet, and Fibre Channel over Ethernet (FCoE).
 - Centralized unified management with Cisco UCS Manager software.
 - Virtual machine optimized services with the support for VN-Link technologies.

The next critical components of the SharePoint 2010 environment are the servers, virtual or physical, that run the Windows Server 2008, the SQL Server 2008 and the SharePoint 2010. A key part of the design process involves complying with organizational standards while meeting the anticipated needs of the organization for the foreseeable life cycle of the technology. For example, some organizations have standards that require all servers to be virtualized, and designs that will meet the anticipated end user requirements for the next three to five years.

The growth of SharePoint in an organization demands more hardware and the implementation of this new hardware either physical or virtual are often the biggest bottlenecks. A lot of work is involved in implementation of the server configurations and in controlling the overall complexity and the number of servers.

The above mentioned design considerations and Microsoft best practices were followed to design the different SharePoint 2010 SP1 server roles and for the servers that were deployed virtually on vSphere as well as for the servers installed natively on the Cisco UCS blade servers.

Microsoft SharePoint 2010 Farm Architecture on FlexPod

The enterprise deployment design was determined using results from the evaluation deployment based on concurrent users, request per second, and page response times for different features. The final design incorporated additional Cisco UCS, VMware, and NetApp end-to-end solution components. The environment was comprised of eight Web front end servers, two application servers, and a mirrored SQL database with a witness server (Figure 13).

Figure 13 Large Microsoft SharePoint 2010 Farm Scenario

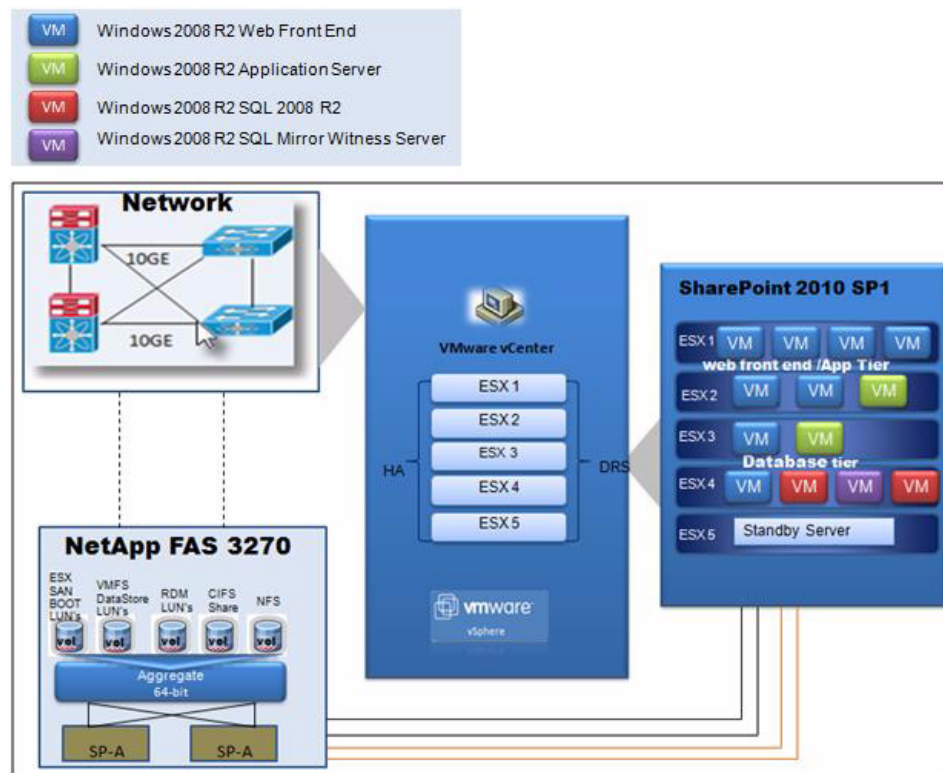


Table 4 lists the various hardware and software components that occupy different tiers of the SharePoint 2010 SP1 Large farm under test.

Table 4 Hardware and Software Components of the Microsoft SharePoint Large Farm


Vendor	Name	Version	Description
Cisco	Cisco 6120 XP	Cisco UCS Manager 2.0(1q)	Cisco UCS 6100 Series Fabric Interconnects
Cisco	Cisco UCS Chassis	5108	Chassis
Cisco	Nexus 5548UP Switch	NX-OS	Nexus Switch 5000 series
Cisco	Cisco UCS Blade Server	B200 M2	Half width Blade server
Cisco	Cisco UCS Blade Server	B250 M2	Full Width Blade Server
VMware	ESXi 5.0	5.0	Hypervisor
VMware	vCenter Server	5.0	VMware Management
Microsoft	Windows Server 2008 R2	2008 R2	Operating System
Microsoft	Microsoft SharePoint server	SharePoint 2010 SP1	Web Front End SharePoint 2010 Enterprise Edition
Microsoft	Microsoft SharePoint Server	SharePoint 2010 SP1	Application Server SharePoint 2010 Enterprise Edition
Microsoft	SQL Server	2008 R2	Database Server SQL R2 Enterprise Edition
Microsoft	SQL Server	2008 R2	Witness server SQL2008 R2 Standard Edition
Microsoft	VSTS 2010	SP1	Load test Controller VSTS 2010 Ultimate Edition
Microsoft	VSTS 2010	SP1	Load test Agents VSTS 2010 Ultimate Edition
NetApp	FAS 3270 Controller	8.0.1	NetApp storage Controller FC /FCOE /Ethernet
NetApp	DS 4243		Shelf's SAS Drives

Table 5 lists the details of all the hardware components used in the deployment. Figure 14 illustrates the FlexPod components for the large Microsoft SharePoint 2010 SP1 Farm.

Table 5 Deployment Hardware Components

Servers	Quantity	Role	Hosted Server Role	Processor	Max Memory Supported	Storage	Network	Protocol
Cisco UCS Blade Server B200 M2	3	ESXi Server	Virtual Servers WFE/ Application Tier	Intel Xeon 5680	96 GB	Internal Drive SAS 2x2.5" & NetApp FAS 3270"	4 Gbps	FC
Cisco UCS Blade Server B250 M2	2	ESXi Server	Virtual Servers SQL Server Database Tier	Intel Xeon 5680	384 GB	Internal Drive SAS 2x2.5" & NetApp FAS 3270"	4 Gbps	FC
Cisco UCS Blade Server B200 M2	1	VSTS Controller	Load Test Controller	Intel Xeon 5680	96 GB	Internal Drive SAS 2x2.5"	10 GB	FCoE
Cisco UCS Blade Server B200 M2	12	VSTS Agent	Load Test Agents	Intel Xeon 5680	96 GB	Internal Drive SAS 2x2.5"	10 GB	FCoE

Figure 14 *FlexPod Components for a Large Microsoft SharePoint 2010 SP1 Farm*



FlexPod Configuration Guidelines

Physical and Virtual CPUs

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

Virtual SMP

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

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

- Simplistically, if multiple, idle physical CPUs are not available when the virtual machine wants to run, the virtual machine remains in a special wait state. The time a virtual machine spends in this wait state is called ready time.

- Even idle processors perform a limited amount of work in an operating system. In addition to this minimal amount, the ESXi host manages these "idle" processors, resulting in some additional work by the hypervisor. These low-utilization vCPUs compete with other vCPUs for system resources.

In VMware ESXi 5 and ESXi, the CPU scheduler underwent several improvements to provide better performance and scalability. For details, see the paper *VMware vSphere 5: The CPU Scheduler in VMware ESXi 5* at http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf. For example, in ESXi 5, the relaxed co-scheduling algorithm was refined so that scheduling constraints due to co-scheduling requirements are further reduced. These improvements resulted in better linear scalability and performance of SMP virtual machines.

VMware recommends the following practices when considering the allocation of vCPUs for SharePoint 2010:

- Allocate the minimum requirement for production virtual machines based on Microsoft guidelines, the role of the virtual machine, and the size of the environment. Additional vCPUs can be added later if necessary.
- Test, development, and proof-of-concept environments can get along with fewer vCPUs allocated to virtual machines. These environments typically require a fraction of the resources needed to satisfy user demand in production.
- When over-committing CPU resources (number of vCPUs allocated to running virtual machines is greater than the number of physical cores on a host), monitor the responsiveness of SharePoint to understand the level of over-commitment which can be provided while still performing at an acceptable level.

Microsoft SharePoint 2010 minimum processor requirements recommended by Microsoft may be excessive in some environments. For this reason, VMware recommends reducing the number of virtual CPUs if monitoring of the actual workload shows that the virtual machine is not benefitting from the increased virtual CPUs. Having virtual CPUs allocated but sitting idle reduces the consolidation level and efficiency of the ESXi/ESXi host.

Over-Commitment

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

When choosing to over-commit CPU resources, monitor vSphere and SharePoint to be sure responsiveness is maintained at an acceptable level. Table 6 lists counters that can be monitored to help drive consolidation numbers higher while maintaining performance.

Table 6 ESXitop CPU Performance Metrics

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

Hyper-Threading

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

Non-Uniform Memory Access (NUMA)

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

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

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

Memory Configuration Guidelines

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

ESXi/ESX Memory Management Concepts

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


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

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

Virtual Machine Memory Concepts

Figure 15 illustrates the use of memory settings parameters in the virtual machine.

Figure 15 Virtual Machine Memory Settings



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

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

Allocating Memory to Microsoft SharePoint 2010 Virtual Machines

The proper sizing of memory for a Microsoft SharePoint 2010 virtual machine is based on many factors. With the number of application services and use cases available determining a suitable configuration for an environment requires creating a baseline configuration, testing, and making adjustments, as discussed later in this paper. Regardless of how much memory virtual machines require, there are best practices to consider when planning for the underlying virtual infrastructure to support SharePoint. The following are descriptions of the recommended best practices:

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

As with over-committing CPU resources, proactive monitoring is a requirement. Table 7 lists counters that can be monitored to avoid performance issues resulting from memory over-commitment.

Table 7 ESXitop Memory Counters

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


Storage Guidelines

VMware vSphere provides many features that take advantage of commonly used storage technologies such as storage area networks and storage replication. Features such as VMware vMotion, VMware HA, and VMware Distributed Resource Scheduler (DRS) use these storage technologies to provide high availability, resource balancing, and uninterrupted workload migration.

Virtual Server Configuration

Figure 16 shows that VMware storage virtualization can be categorized into three layers of storage technology:

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

Figure 16 VMware Storage Virtualization Stack

VMFS File System

The VMFS file system was created by VMware to allow multiple vSphere hosts to read and write to the same storage concurrently. VMFS is a clustered file system that allows you to simplify virtual machine provisioning and administration by consolidating virtual machines into smaller units of storage. Unique virtualization-based capabilities provided by the VMFS file system include live migration using vMotion, and increased availability using VMware HA.

Virtual machines are stored on VMFS datastores as a unique set of encapsulated files, including configuration files and virtual disks (VMDK files). VMFS is supported on both iSCSI and Fibre Channel attached storage.

Raw Device Mapping (RDM)

For instances where isolation or direct access to the underlying storage subsystem is required a raw device mapping can be used in place of virtual disks. Raw device mappings use a mapping file that is located on a VMFS datastore to point to a physical LUN. The physical LUN is accessible to the virtual machine in its raw form and must be formatted from within the virtual machine. Unlike VMFS, a raw device mapping is typically only assigned to a single virtual machine; however, RDMs can be shared, for example, in a Microsoft Cluster configuration where multiple nodes use SCSI reservations to handle arbitration. RDMs cannot provide all of the features available with VMFS and should be limited to use only when technically required.

For a more information on SAN system design, see the *VMware SAN System Design and Deployment Guide* at <http://www.vmware.com/resources/techresources/772>.

Storage Protocol Capabilities

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

For SharePoint 2010 on vSphere there is no single option that is considered superior to another. It is recommended that this decision be made based on your established storage management practices within the virtualized environment.

For more information, see the VMware white paper *Comparison of Storage Protocol Performance in VMware vSphere 5* at http://www.vmware.com/files/pdf/perf_vsphere_storage_protocols.pdf.

Storage Best Practices

The following are vSphere storage best practices:

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

Network Configuration

Table 8 lists the network components used for the deployment.

Table 8 Network Configuration Details

Network Components	Number of Components Used
Cisco Fabric Interconnect 6248UP	2
Nexus 5548up Series Switch	2

Storage Configuration

Table 9 lists the storage system and controller types used in the deployment.

- ESXi 5.0 is SAN Booted on Cisco UCS blade B200 M2 and B250 M2 servers. For more information, refer to the section [“Boot from SAN.”](#)
- Virtual servers are configured with pass through switching. For more information, refer to the section [“Network Configuration.”](#)

Table 9 Storage Components

NetApp Storage Subsystem	Number of Controllers Used	Protocol	Disks
FAS 3270 Controller	2	FC/FCoE/Ethernet	
DS 4243	3 x Shelves http://www.netapp.com/us/products/storage-systems/disk-shelves-and-storage-media/disk-shelves-tech-specs.html		24 x 3 disk drives SAS Drives With 15k RPM

Microsoft SharePoint 2010 VMware Memory Virtualization

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

Memory Compression

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

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

Virtual Networking

The Cisco Virtual Machine Fabric Extender (VM-FEX) collapses virtual and physical networking into a single infrastructure. The VM-FEX allows data center administrators to provision, configure, manage, monitor, and diagnose virtual machine network traffic and bare metal network traffic within a unified infrastructure.

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

- Each virtual machine includes a dedicated interface on the parent switch
- All virtual machine traffic is sent directly to the dedicated interface on the switch
- The software-based switch in the hypervisor is eliminated
- VM-FEX is supported on Red Hat Kernel-based Virtual Machine (KVM) and VMware ESXi hypervisors. Live migration and vMotion are also supported with VM-FEX.

Benefits

- Simplified operations—Eliminates the need for a separate, virtual networking infrastructure
- Improved network security—Contains VLAN proliferation
- Optimized network utilization—Reduces broadcast domains
- Enhanced application performance—Offloads virtual machine switching from host CPU to parent switch application-specific integrated circuits (ASICs)

Virtual Networking Best Practices

The following are the vSphere networking best practices:

- Separate virtual machine and infrastructure traffic—Keep virtual machine and VMkernel or service console traffic separate. This can be accomplished physically using separate virtual switches that uplink to separate physical NICs, or virtually using VLAN segmentation.
- Use NIC Teaming—Use two physical NICs per vSwitch, and if possible, uplink the physical NICs to separate physical switches. Teaming provides redundancy against NIC failure and, if connected to separate physical switches, against switch failures. NIC teaming does not necessarily provide higher throughput.
- Enable PortFast on ESXi/ESX host uplinks—Failover events can cause spanning tree protocol recalculations that can set switch ports into a forwarding or blocked state to prevent a network loop. This process can cause temporary network disconnects. To prevent this situation, set the switch ports connected to ESXi/ESX hosts to PortFast, which immediately sets the port back to the forwarding state and prevents link state changes on ESXi/ESX hosts from affecting the STP topology. Loops are not possible in virtual switches.
- Converged Network and Storage I/O with 10Gbps Ethernet —When possible consolidating storage and network traffic can provide simplified cabling and management over having to maintain separate switching infrastructures.

VMware vSphere Performance

With every release of VMware vSphere, the overhead of running an application on the VMware vSphere virtualized platform is reduced by way of new performance improving features. Typical virtualization overhead for applications, such as SharePoint 2010, is less than 10 percent. Many of these features not

only improve performance of the virtualized application itself, but also allow for higher consolidation ratios. Understanding these features and taking advantage of them in your SharePoint 2010 environment helps guarantee the highest level of success in your virtualized deployment (Table 10).

Table 10 VMware vSphere Performance

ESXitop Metric	Description	Implication
NUMA Support	ESXi/ESXi uses a NUMA load-balancer to assign a home node to a virtual machine. Because memory for the virtual machine is allocated from the home node, memory access is local and provides the best performance possible. Even applications that do not directly support NUMA benefit from this feature.	See VMware vSphere 4: The CPU Scheduler in VMware ESXi 5 at http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf
Transparent Page Sharing	Virtual machines running similar operating systems and applications typically have identical sets of memory content. Page sharing allows the hypervisor to reclaim the redundant copies and keep only one copy, which frees up the total host memory consumption. Because most SharePoint virtual machines run the same operating system and application binaries, total memory usage can be reduced to increase consolidation ratios.	See Understanding Memory Resource Management in VMware ESXi 5.0 at http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf
Memory Ballooning	By using a balloon driver loaded in the guest operating system, the hypervisor can reclaim host physical memory if memory resources are under contention. This is done with little to no impact to the performance of the application.	See Understanding Memory Resource Management in VMware ESXi 5.0 at Http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf
Memory Compression	Before a virtual machine resorts to host swapping, due to memory over-commitment the pages elected to be swapped attempt to be compressed. If the pages can be compressed and stored in a compression cache, located in main memory, the next access to the page causes a page decompression as opposed to a disk swap out operation, which can be an order of magnitude faster.	See Understanding Memory Resource Management in VMware ESXi 5.0 at Http://www.vmware.com/files/pdf/perf-vsphere-memory_management.pdf
Large Memory Page Support	An application that can benefit from large pages on native systems, such as MS SQL, can potentially achieve a similar performance improvement on a virtual machine backed with large memory pages. Enabling large pages increases the memory page size from 4KB to 2MB.	See Performance Best Practices for VMware vSphere 5.0 at http://www.vmware.com/pdf/Perf_Best_Practices_vSphere5.0.pdf and Performance and Scalability of Microsoft SQL Server on VMware vSphere 4 at http://www.vmware.com/files/pdf/perf_vsphere_sql_scalability.pdf

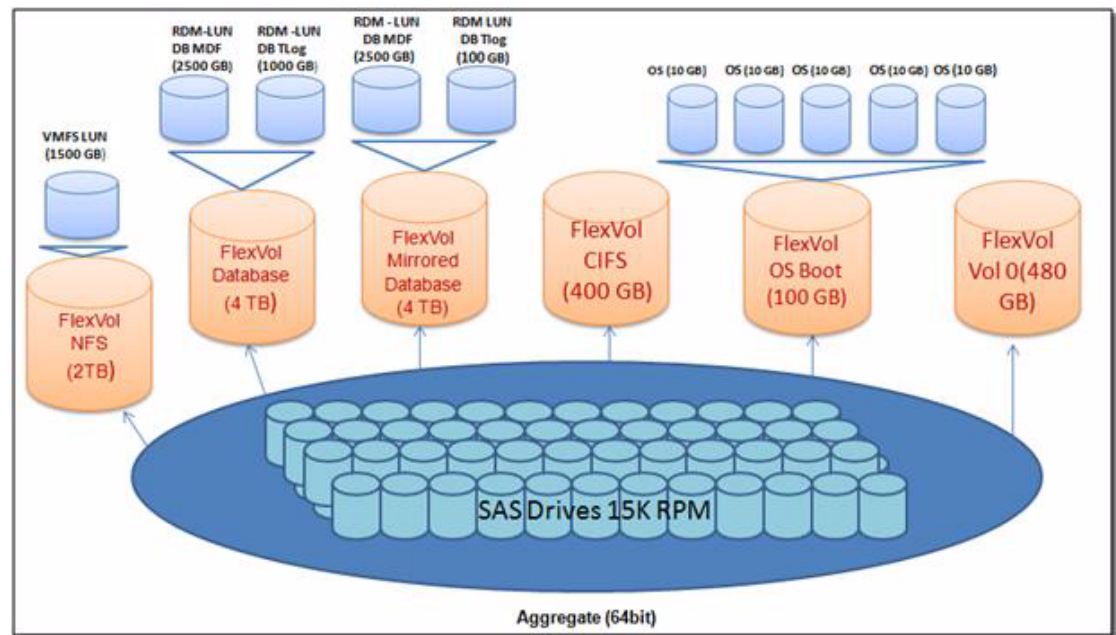
Microsoft SharePoint 2010 VMware Storage Virtualization

Storage Layout

A single large 64 bit aggregate is created for the Microsoft SharePoint 2010. Single large aggregate may maximize performance for the SharePoint data. The Microsoft SharePoint 2010 volumes are allocated based on load distribution across the 3270 controllers.

Figure 17 illustrates the Microsoft SharePoint 2010 Storage.

Figure 17 *Microsoft SharePoint 2010 Storage*



Aggregate, Volume, and LUN Sizing

The aggregate contains all the physical disks for a workload and for SharePoint 2010 SP1.

The larger size of 64-bit aggregates makes it possible to add a lot more disks to an aggregate than is feasible with 32-bit aggregates. Therefore, for scenarios where the disks are the bottleneck in improving performance, 64-bit aggregates with a higher spindle count can give a performance boost. All the FlexVol volumes that are created inside a 64-bit aggregate span across all the data drives in the aggregate, thus providing more disk spindles for the I/O activity on the FlexVol volumes.



Note

NetApp recommends having at least 10% free space available in an aggregate that is hosting SharePoint data. This allows optimal performance of the storage system.

The volume is generally sized at 90 percent of the aggregate size, housing both the actual LUNs and the snapshots of those LUNs. This sizing takes into account the database, the content database, the transaction logs, the growth factor, and 20 percent of free disk space.

For more information about 64-bit Aggregate, go to <http://media.netapp.com/documents/tr-3786.pdf>.

In this design, the user database log has been separated onto its own volume. This provides more granular control of which disks it resides on and how many LUNs it is composed of. Spreading the log across multiple LUNs can help improve its ability to handle higher I/O requirements. This also allows the log file to be managed independently of the data file.

Storage Considerations

When planning for content storage on SharePoint 2010, you must choose suitable storage architecture. SharePoint 2010 content storage has a significant dependency on the underlying database; therefore, database and SQL Server requirements will drive the storage choices.

Storage Virtualization

VMFS is a cluster file system that provides storage virtualization optimized for virtual machines. Each virtual machine is encapsulated in a small set of files and VMFS is the default storage system for these files on physical SCSI disks and partitions. VMware supports Fibre-Channel, iSCSI, and NAS shared storage protocols.

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

Setting up NetApp Storage for Microsoft SharePoint 2010 SP1

In this solution, the Microsoft SharePoint 2010 farm has a pair of FAS 3270 storage arrays equipped with three disk shelves, each disk shelf consists of 24 drives for a total of 72 SAS 15k RPM drives with a storage space total of 43 TB.

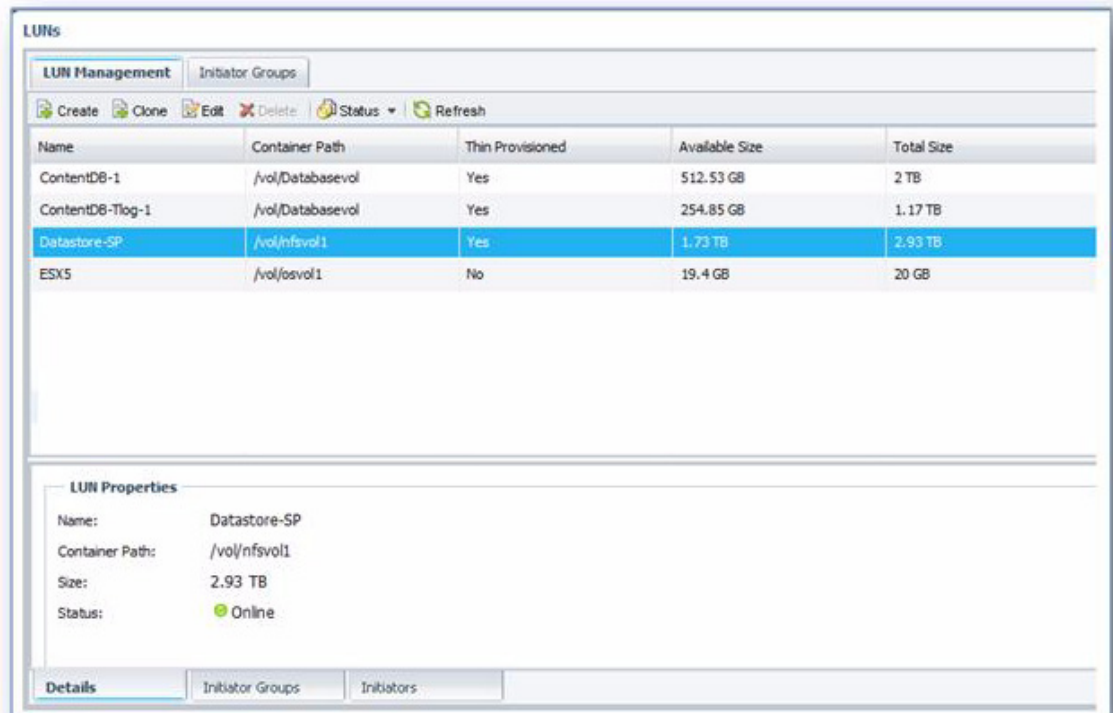
The storage provisioning for this solution uses some of the NetApp best practices for Microsoft SharePoint 2010 SP1 and VMware best practices for SharePoint 2010 SP1 virtualization. The solution also uses the storage flexibility available in SharePoint 2010 SP1.

For information about the NetApp best practices, go to <http://www.netapp.com/us/library/technical-reports/tr-3887.html>.

VMFS Datastore for VMotion

In this solution for data center, a shared Virtual machine file system (VMFS) Datastore is provisioned, which stores the VM images is required to support VMware DRS/VMotion feature for this solution.

The ESXi hosts that are participating in VMotion are given access to this shared datastore on the NetApp array by including them in the Initiator Group for that LUN. The ESXi servers in the large farm are mapped to the /vol/VM/VM LUN (VMFS datastore) in the LUN Initiator Group ESXi Servers as shown in Figure 18. The mapping involves specifying the WWPNs of the VHBAs defined in the Cisco UCS Service Profile for each ESXi server in the Initiator Group “ESXi Servers.”

Figure 18 **Managing LUNs**

Soft Zoning and VSAN

Before the ESXi configuration is done on the vCenter, you need to make sure the soft zoning and VSAN configuration on the SAN fabric switches permit the Cisco UCS blade servers access to the storage array. This implies that the WWPNs of the NetApp FAS 3270 controllers and the WWPNs of the vHBAs on the Cisco UCS blade servers must all be in VSAN and part of the same zone and zoneset. On each of the FC switches in the large farm, the VSAN and soft zoning configuration can be seen using the command “show zoneset active” as follows:

```
rk2-n5548-a(config-zoneset)# zoneset activate name zoneset-SP2010 vsan 3
WARNING: Specified zoneset already active and unchanged
rk2-n5548-a(config)# sh zoneset active
zoneset name zoneset-SP2010 vsan 3
zone name ESX5-SP2010-Blade2 vsan 3
* fcid 0x2b0000 [pwn 20:42:00:05:9b:77:4b:00]
* fcid 0x2b0002 [pwn 20:41:00:05:9b:77:4b:00]
* fcid 0x2b0005 [pwn 20:00:00:25:b5:04:01:0e]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]

zone name ESX5-SP2010-Blade1 vsan 3
* fcid 0x2b0001 [pwn 20:00:00:25:b5:04:01:0f]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]

zone name ESX5-SP2010-Blade3 vsan 3
* fcid 0x2b0006 [pwn 20:00:00:25:b5:04:01:0d]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]

zone name ESX5-SP2010-Blade4 vsan 3
* fcid 0x2b0007 [pwn 20:00:00:25:b5:04:01:0c]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]

zone name ESX5-SP2010-Blade5 vsan 3
* fcid 0x2b0008 [pwn 20:00:00:25:b5:04:01:0b]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]

zone name ESX5-SP2010-Blade6 vsan 3
* fcid 0x2b0009 [pwn 20:00:00:25:b5:04:01:0a]
* fcid 0x2b0003 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x2b0004 [pwn 50:0a:09:84:8d:93:40:7f]
```

```
rk2-n5548-b(config)# sh zoneset active
zoneset name zoneset-SP2010 vsan 3
zone name ESX5-SP2010-Blade2 vsan 3
* fcid 0x150005 [pwn 20:00:00:25:b5:04:02:0e]
* fcid 0x150002 [pwn 20:41:00:05:73:a2:66:80]
* fcid 0x150000 [pwn 20:42:00:05:73:a2:66:80]
* fcid 0x150004 [pwn 50:0a:09:83:9d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]

zone name ESX5-SP2010-Blade5 vsan 3
* fcid 0x150008 [pwn 20:00:00:25:b5:04:02:0b]
* fcid 0x150004 [pwn 50:0a:09:83:8d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]

zone name ESX5-SP2010-Blade1 vsan 3
* fcid 0x150001 [pwn 20:00:00:25:b5:04:02:0f]
* fcid 0x150004 [pwn 50:0a:09:83:8d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]

zone name ESX5-SP2010-Blade3 vsan 3
* fcid 0x150006 [pwn 20:00:00:25:b5:04:02:0d]
* fcid 0x150004 [pwn 50:0a:09:83:8d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]

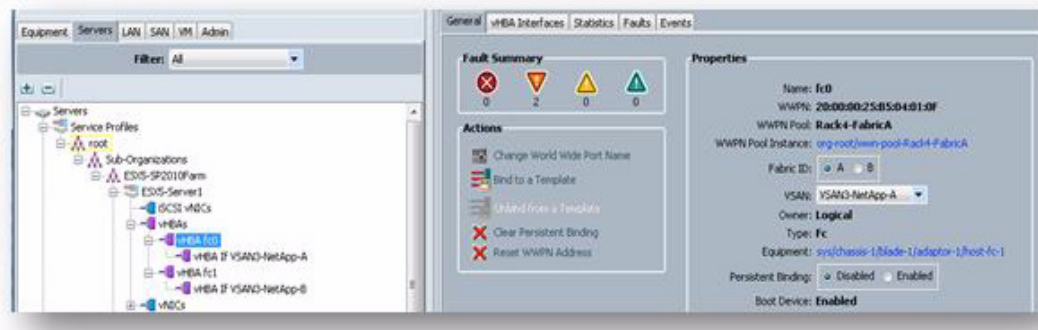
zone name ESX5-SP2010-Blade4 vsan 3
* fcid 0x150007 [pwn 20:00:00:25:b5:04:02:0c]
* fcid 0x150004 [pwn 50:0a:09:83:8d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]

zone name ESX5-SP2010-Blade6 vsan 3
* fcid 0x150009 [pwn 20:00:00:25:b5:04:02:0a]
* fcid 0x150004 [pwn 50:0a:09:83:8d:93:40:7f]
* fcid 0x150003 [pwn 50:0a:09:84:9d:93:40:7f]
```

Cisco UCS Service Profile

For high availability, the ESXi hosts need a dual fibre-channel connections to the FC SAN (Figure 19). VMware VSphere 5 provides a native MPIO which is leveraged in this solution to provide that redundancy. Initially, the Cisco UCS Service Profile for each ESXi host is configured to have two vHBA going through Fabric Interconnects; Fabric Interconnect A and the other going through Fabric Interconnect B. Since there is only a single VSAN in this solution, both are configured on the same VSAN 3.

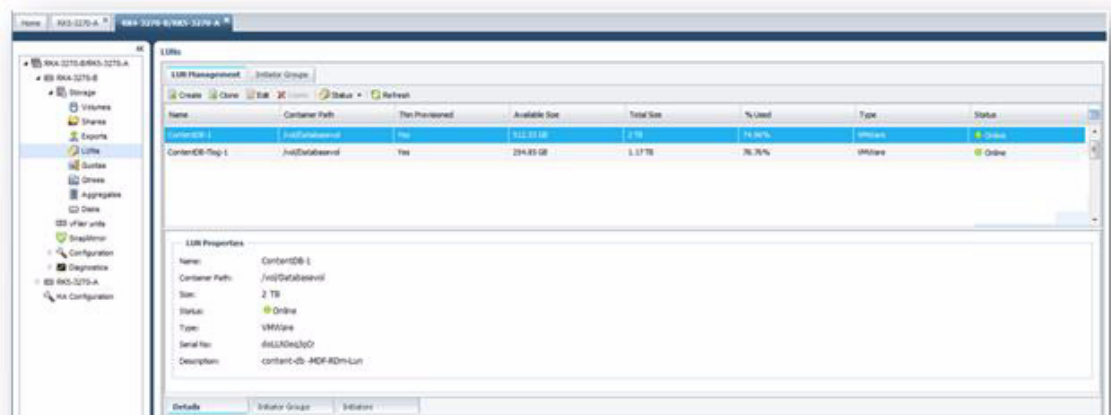
Figure 19 VSAN 3 for Fabric Interconnect A and B



Microsoft SharePoint 2010 SP1 Database and Log LUNs

The Microsoft SharePoint 2010 database storage is provisioned on separate RDM LUN's for databases and logs. Disks are configured with RAID DP (Figure 20). Databases (.mdf) and (.ldf) files are on separate RDM LUNs. FC is used as the transport protocol to storage subsystem.

Figure 20 LUN Mapping



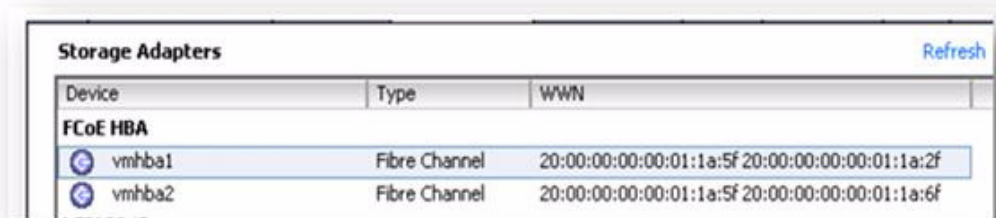
ESXi Native MPIO

Each ESXi host minimally needs FC connectivity to the VMFS datastore containing the virtual machine images of the virtualized Microsoft SharePoint 2010 servers and virtual machine for Database servers in the Microsoft SharePoint farm are mapped with RDM LUNs on the SAN for storing the Content database and transaction log files.

Through the vCenter, you can see the list of LUNs accessible by each ESXi host and the dual paths available through the SAN fabric.

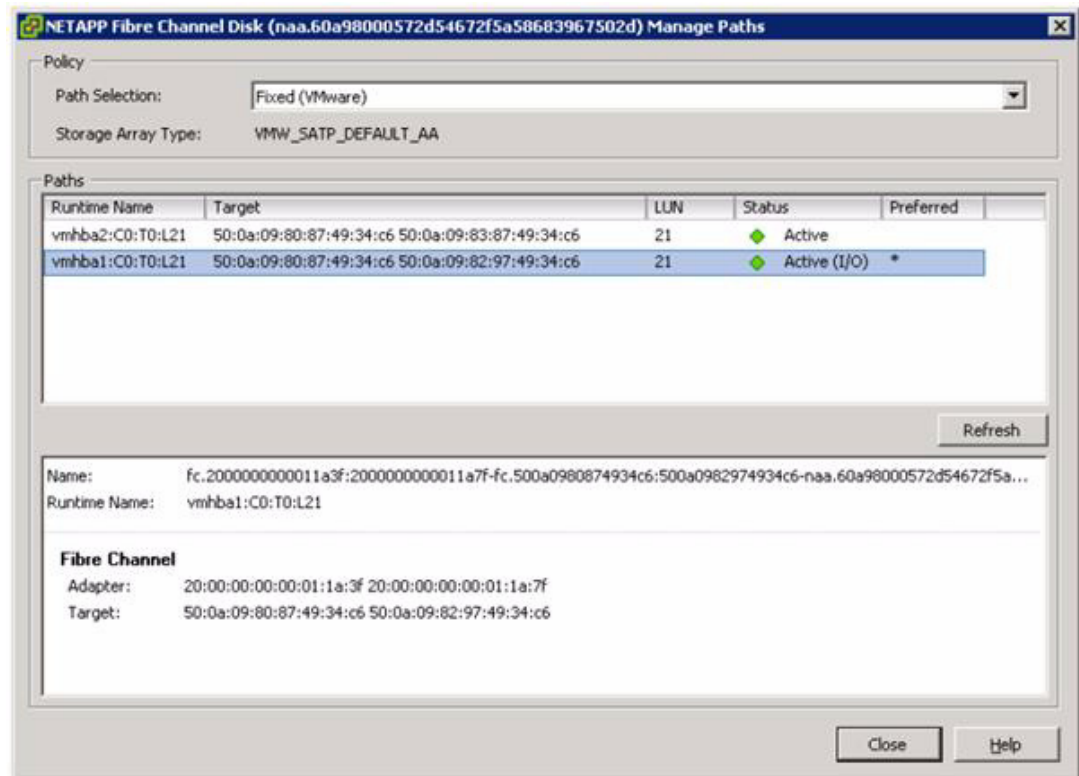
Figure 21 shows the Dual FC Paths to VMFS Datastore and RDM LUNs.

Figure 21 *Dual FC Paths to VMFS Datastore and RDM LUNs*



Device	Type	WWN
FCoE HBA		
vmhba1	Fibre Channel	20:00:00:00:01:1a:5f 20:00:00:00:01:1a:2f
vmhba2	Fibre Channel	20:00:00:00:01:1a:5f 20:00:00:00:01:1a:6f

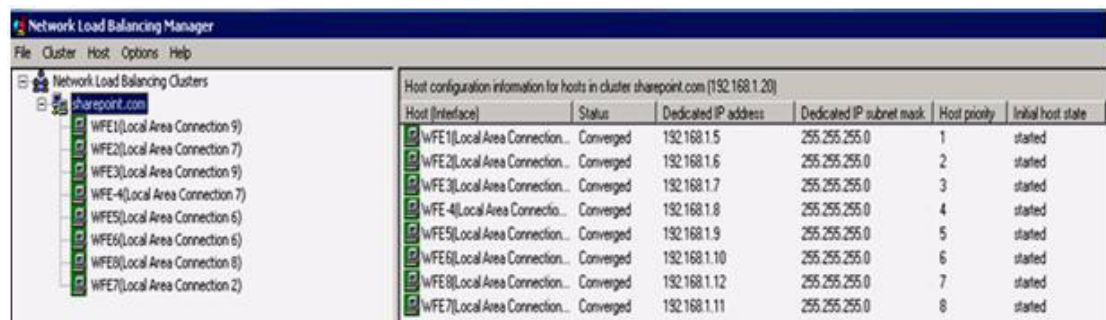
ESXi native MPIO allows you to specify the preferred FC path to each LUN. This can be done by selecting “Edit Settings” for the virtual machine and selecting “Manage Paths” for the LUN. The pop-up window, as shown in Figure 22, allows the administrator to select the path selection method based on whether a round-robin scheme or a single path from the drop-down options as desired. For this solution, the path selection is set to “Fixed (VMware)” and a preferred path (marked as “Active (I/O)”) is constantly used as long as it is available. When this path is not available, the secondary path will become the new fixed path.

Figure 22 *ESXi Native MPIO Path Selection*

Microsoft SharePoint 2010 High Availability

Large farm topology discussed in this document, is implemented with Microsoft SharePoint High Availability. High Availability for a three tier SharePoint Farm under test is implemented as described:

To implement high availability for front-end web servers, eight front-end web servers are used. These servers host the same web application, with Windows network Load balancing (NLB) feature (Figure 23).

Figure 23 *Network Load Balancing Manager for WFE Servers*

NLB is available in Windows Server 2008 Web, Standard, Enterprise, and Datacenter editions.

For information on configuring NLB, go to:

[http://technet.microsoft.com/en-us/library/cc754833\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc754833(WS.10).aspx)

- To implement high availability for service applications, two application servers are chosen that host the same services.
- To implement high availability for databases, two machines are used to run the SQL Server. You can configure these for database mirroring (which requires duplicate storage), or you can configure them as part of a SQL Server failover cluster (which requires shared storage).

For the purpose for this study Database tier is implemented with SQL mirroring along with a witness server. SQL Mirroring is integrated with SharePoint 2010 Application. Figure 24 shows the Database Mirroring.

Figure 24 **Database Mirroring**

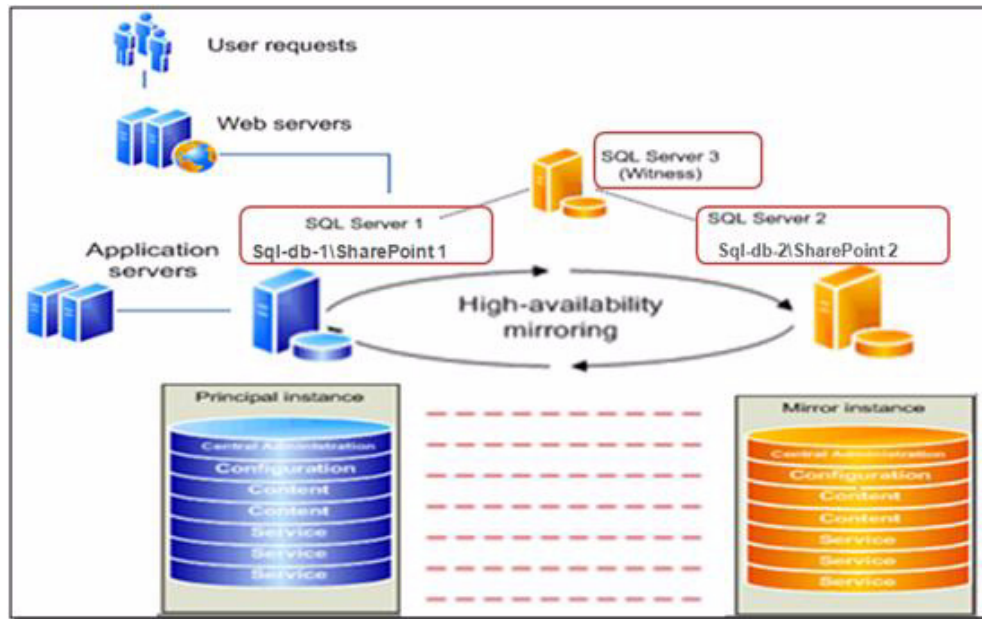
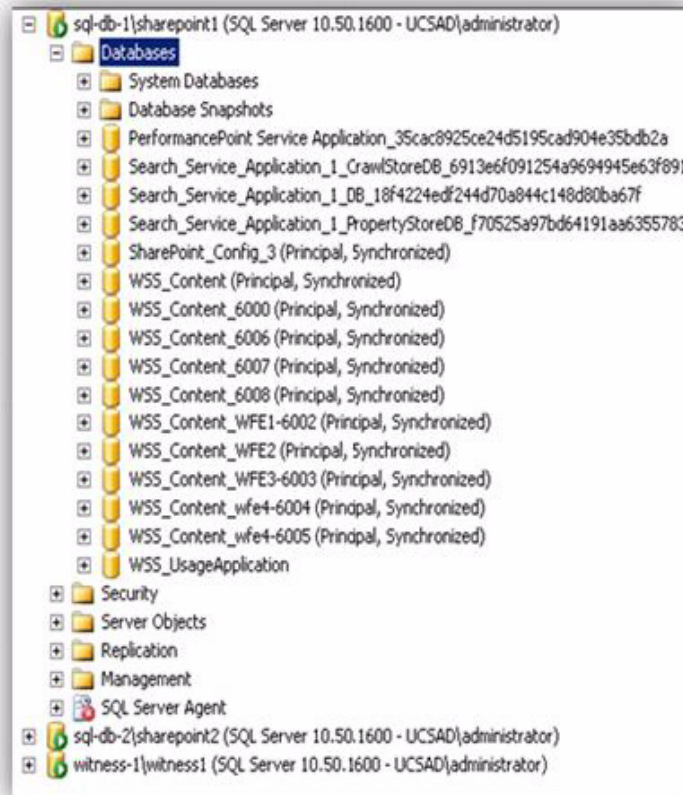



Figure 25 *SQL Mirroring Configuration for Microsoft SharePoint Databases***Figure 26** *Microsoft SharePoint 2010 Configuration of a Failover Server*

For more information on configuring mirroring, go to:

<http://technet.microsoft.com/en-us/library/ms188712.aspx>.

Service Profile Configuration

You can create service profiles in the following ways:

- Manually—You can create a service profile manually by using the Cisco UCS Manager GUI.
- From a template—You can create a service profile from a service profile template.
- By cloning—You can create a service profile by cloning an existing service profile. Cloning creates a replica of the existing service profile. Cloning is equivalent to creating a template from the service policy and then creating a service profile from the template.

In this CVD, a service profile initial template is created and a service profile is then instantiated from the template.

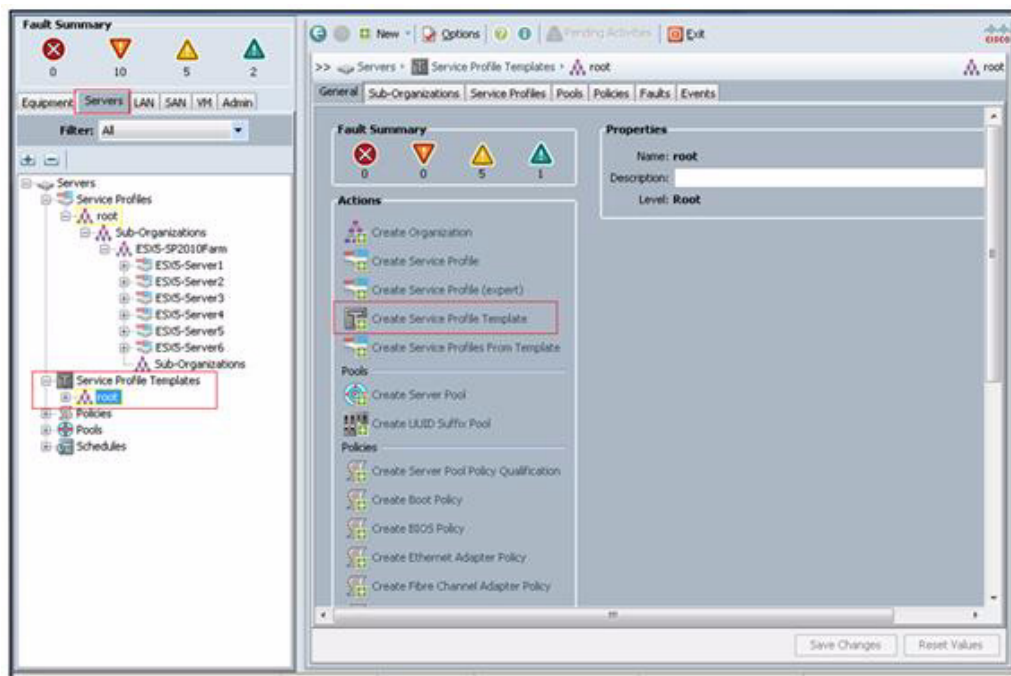
A service profile template parameterizes the UUIDs that differentiate one instance of a server from another instance, which otherwise is identical. The following are the different types of templates:

- **Initial Template**—An initial template is used to create a new server with UUIDs from a service profile. The changes made to the template do not reflect in the server because there is no link between the server and the template after deployment of the server. All changes to items defined by the template must be made individually on each deployed server.
- **Updating Template**—An updating template maintains a link between the template and the deployed servers, and changes to the template, mostly firmware revisions are reflected in the servers deployed with the template on a schedule determined by the administrator.

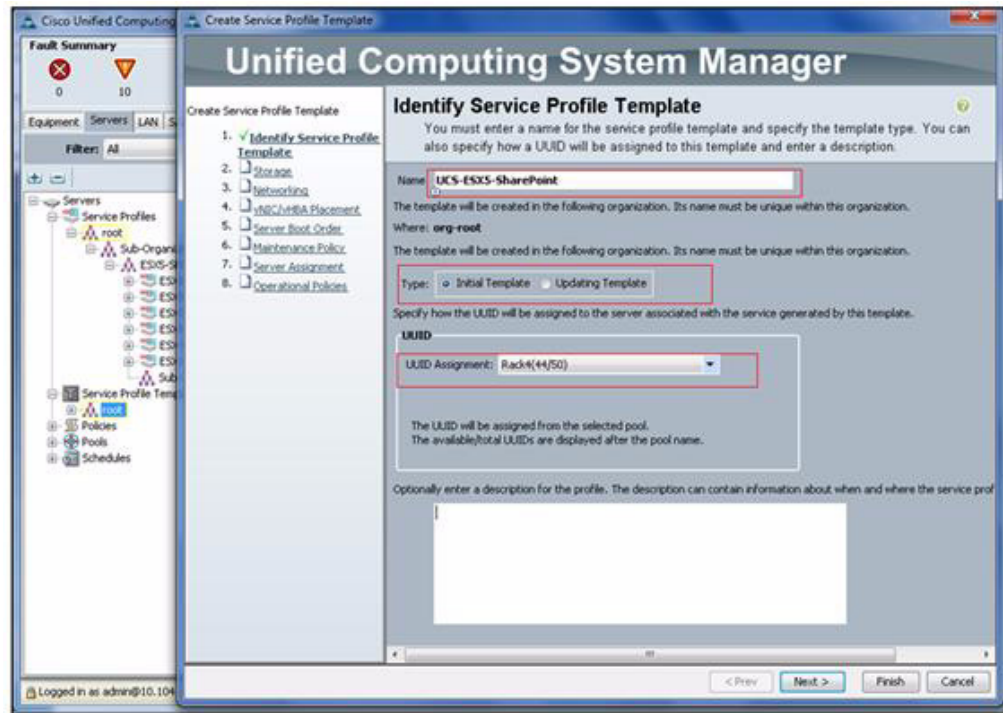
Service profiles, templates, and other management data are stored in a high-speed persistent storage on the Cisco Unified Computing System Fabric Interconnects, with mirroring between fault-tolerant pairs of Fabric Interconnects.

Perform the following steps to configure a service profile template for ESXi host on each of the Blade servers

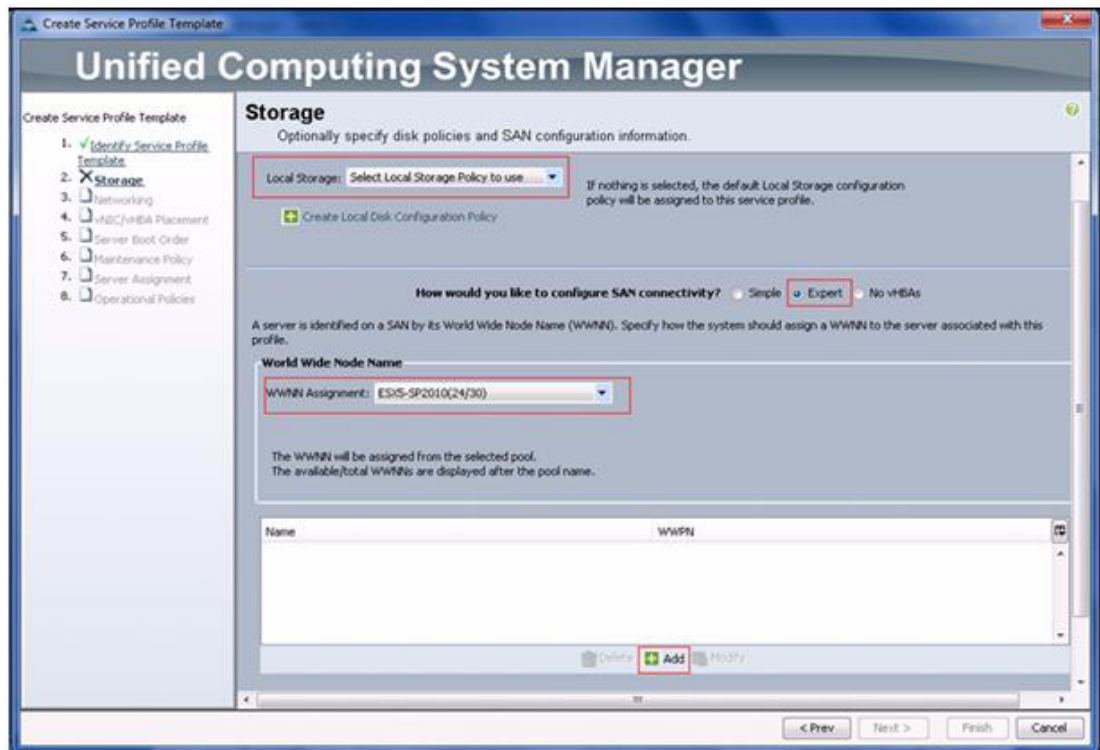
1. Click **Create Service Profile Template** link on the right pane of the window.



2. Enter a valid name in the Name field in the service profile template.



3. Select the type of the template and click the **Initial Template** radio button.
4. Click the **UUID Assignment** drop-down arrow to select the pool name and the total available UUIDs are displayed after the pool name.
5. Click **Next**.
6. For Storage Configuration, perform the following steps:
 - a. Click the **Local Storage** drop-down arrow and do not select any option to apply the default storage configuration policy. When doing a SAN boot for the B200/B250 server, the RAID policy configured in the storage LUN is used.
 - b. Click **Expert** to configure SAN Connectivity.
 - c. Select the World Wide Node Name (WWNN) from the WWNN Assignment drop-down list.
 - d. Click **Add** to assign WWPN for VHBA



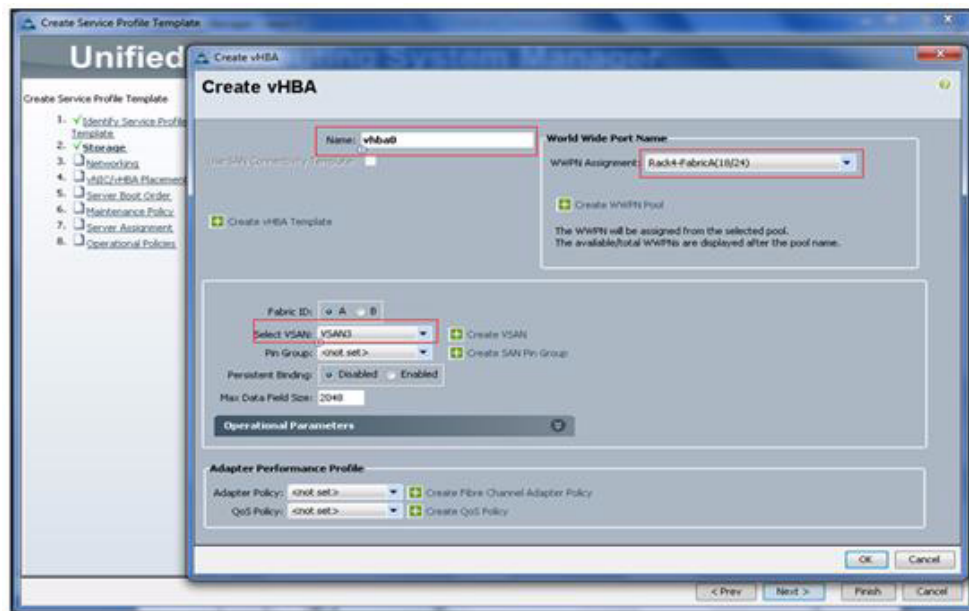
7. In the current setup you need to configure four VHBAs:

- Two VHBAs are used for SAN boot with SAN LUN configured with RAIDDP
- Two VHBAs are used for installation configured with SAN LUN configured with RAIDDP.

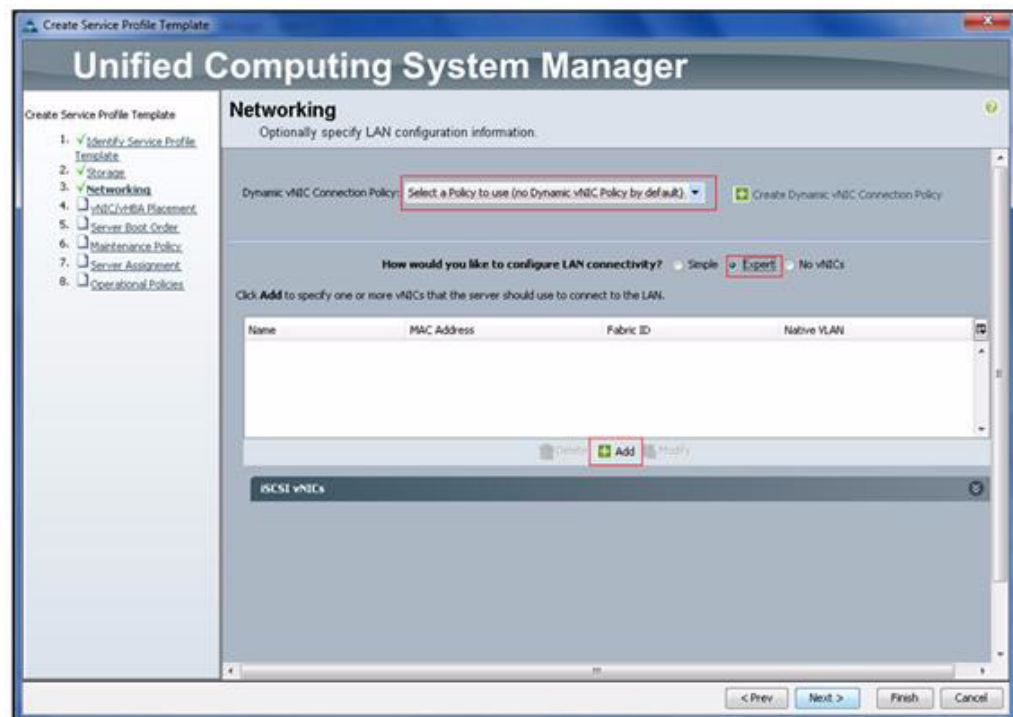
Each of the VHBAs for SAN Boot and installation is mapped to Fabric A and Fabric B respectively. This allows redundancy at the Fabric interconnect level.

8. To configure the VHBA, perform the following steps:

- a. Select the WWPN pool as configured in the previous section. Different WWPN pools for Fabric A and Fabric B have been configured.
- b. Click Fabric ID radio button A for vhba0 and Fabric ID radio button B for vhba1.
- c. Click the Select VSAN drop-down arrow to select VSAN configured in the previous section.
- d. Follow the same steps for the other three VHBAs. Click **OK** to save the configuration settings.
- e. Click **Next**.



9. In Networking window, perform the following steps to configure LAN:
 - a. Click **Expert** to configure LAN connectivity.
 - b. Click the plus button to create the Dynamic VNIC Connection Policy.
 - c. Click Add to add one or more VNICs used by the server to connect to the LAN.



10. The Create VNIC Window allows you to add two VNICs; eth0 and eth1, each configured with Fabric A and Fabric B. Click **OK** to save the VNIC configuration.

Create vNIC

Name:

Use LAN Connectivity Template: ☐

MAC Address

MAC Address Assignment:

The MAC address will be automatically assigned from the selected pool.

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ Enable Failover

Select	Name	Native VLAN
<input type="checkbox"/>	default	<input type="radio"/>
<input type="checkbox"/>	809	<input type="radio"/>
<input checked="" type="checkbox"/>	810810	<input checked="" type="radio"/>
<input type="checkbox"/>	810811	<input type="radio"/>

MTU:

Pin Group:

Operational Parameters

Adapter Performance Profile

Adapter Policy:

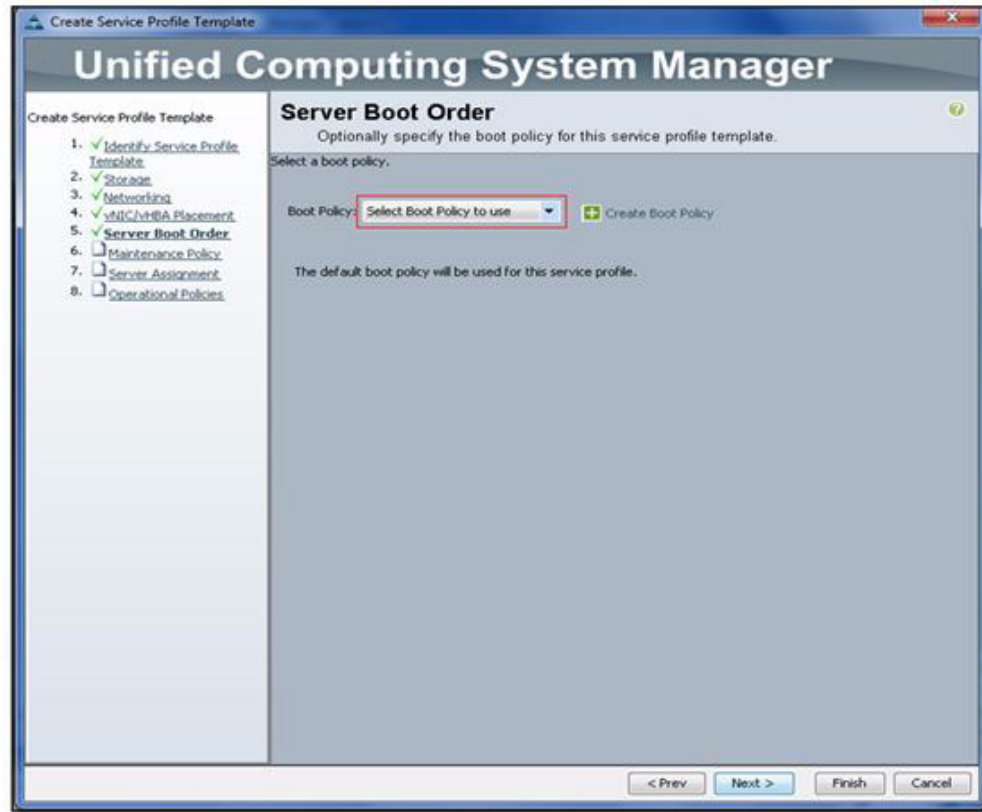
QoS Policy:

Network Control Policy:

11. In the VNIC/VHBA placement window, from the Select Placement drop-down list, select the default setting, **Let System Performance Placement**.
12. Click **Next**.



13. In the Server Boot order window, do not select any boot policy. You need to configure a new boot policy in SAN Configuration section, as you will be doing a SAN Boot. Click **Next**.



14. Keep the default settings for the Maintenance Policy as shown below.

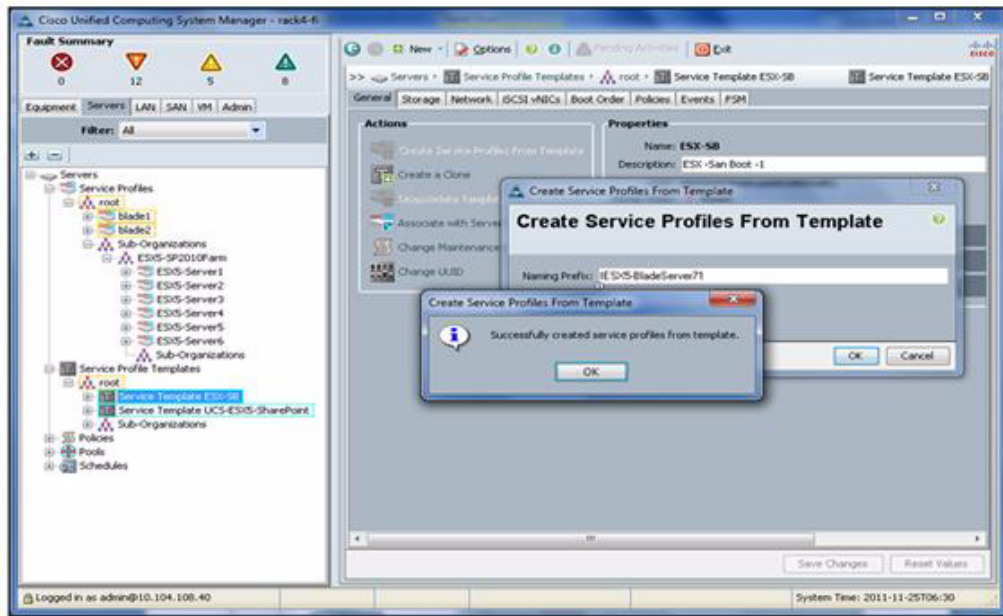


You can define your custom policy for each of the three screens, for instance, in Operational Policy window you can define a BIOS Policy which will be assigned to server as per the requirement.

15. Click Finish.

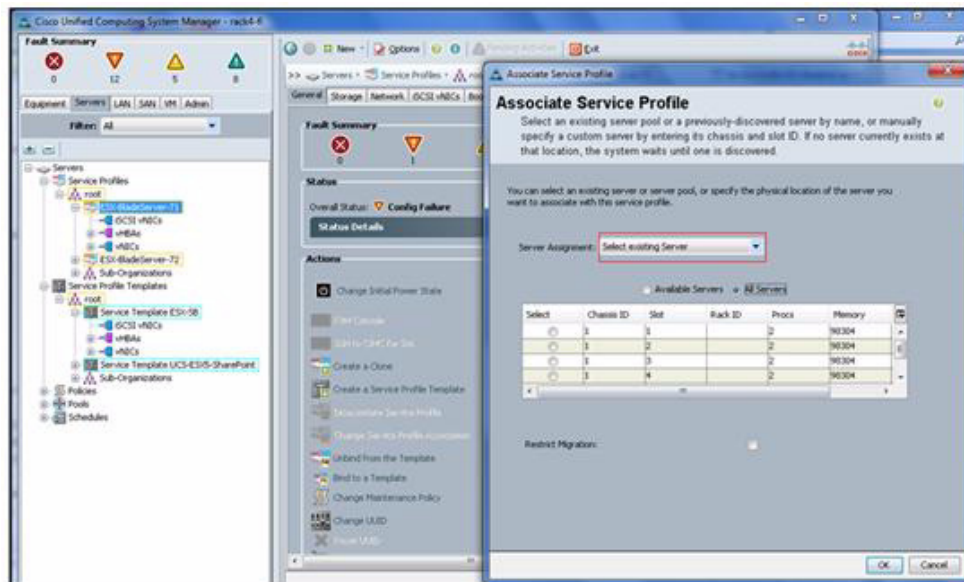
You can view the Service profile template on the Servers tab option in the Service Profile Templates window.

16. Create a Service Profile from Service profile template and associate it with a B200/B250 M2 server placed in the Cisco UCS Chassis 5108.




17. When the service profile is created, you need to associate the service profile with the available server slot in Chassis 5108.
 - a. Select the created Service profile and go to Change Service Profile Association window.
 - b. Click the Server Assignment drop-down arrow and select **Existing Server** from the list.

The workflow of the Service Profile Association displays.



The screen shot below shows the progress in the Service profile under the FSM status tab, once you start associating the Service Profile to the available server.



The above steps demonstrate the stateless nature of the Cisco Unified Computing System enabling data center servers. The server's identity (using MAC or WWN addressing or UUIDs) as well as the build and the operational policy information such as firmware, BIOS revisions, network, and storage connectivity profiles can be dynamically provisioned or migrated to any physical server in the system.

The next section details the SAN Configuration followed by a definition of the Boot policy in a Cisco UCS Service profile and enabling SAN Boot.

Boot from SAN

Boot from SAN is a critical feature which helps to achieve stateless computing in which there is no static binding between a physical server and the OS / applications the server is supposed to run. The OS is installed on a SAN LUN and boot from the service profile. When a service profile is moved to another server, the PWWN of the HBAs and the server policy will also move along with the server profile. The new server looks the same as the old server.

The following are the benefits of boot from SAN:

- **Reduce Server Footprints**—Boot from SAN eliminates the need alleviates the necessity for each server to have its own direct-attached disk (internal disk) which is a potential point of failure. The following are the advantages of thin diskless servers:
 - Require less physical space
 - Require less power
 - Require fewer hardware components
 - Less expensive
- **Disaster and Server Failure Recovery**—All the boot information and production data stored on a local SAN can be replicated to a SAN at a remote disaster recovery site. When server functionality at the primary site goes down in the event of a disaster, the remote site can take over with minimal downtime.
- **Recovery from server failures**—Recovery from server failures is simplified in a SAN environment. With the help of server snapshots, mirrors of a failed server can be quickly recovered. This greatly reduces the time required for server recovery.

- **High Availability**—A typical data center is highly redundant in nature with redundant paths, redundant disks and redundant storage controllers. The operating system images are stored on SAN disks which eliminates potential problems caused due to mechanical failure of a local disk.
- **Rapid Redeployment**—Businesses that experience temporary high production workloads can take advantage of SAN technologies to clone the boot image and distribute the image to multiple servers for rapid deployment. Such servers may only need to be in production for hours or days and can be readily removed when the production need has been met. Highly efficient deployment of boot images makes temporary server usage highly cost effective.
- **Centralized Image Management**—When operating system images are stored on SAN disks, all upgrades and fixes can be managed at a centralized location. Servers can readily access changes made to disks in a storage array.

With boot from SAN, the server image resides on the SAN and the server communicates with the SAN through a Host Bus Adapter (HBA). The HBA BIOS contain instructions that enable the server to find the boot disk. After Power On Self Test (POST), the server hardware component fetches the designated boot device in the hardware BOIS settings. Once the hardware detects the boot device, it follows the regular boot process.


There are four distinct portions of the SAN procedure:

1. Storage array configuration
2. Cisco UCS configuration of service profile
3. SAN zone configuration
4. Host Registration on Storage

Storage Array Configuration

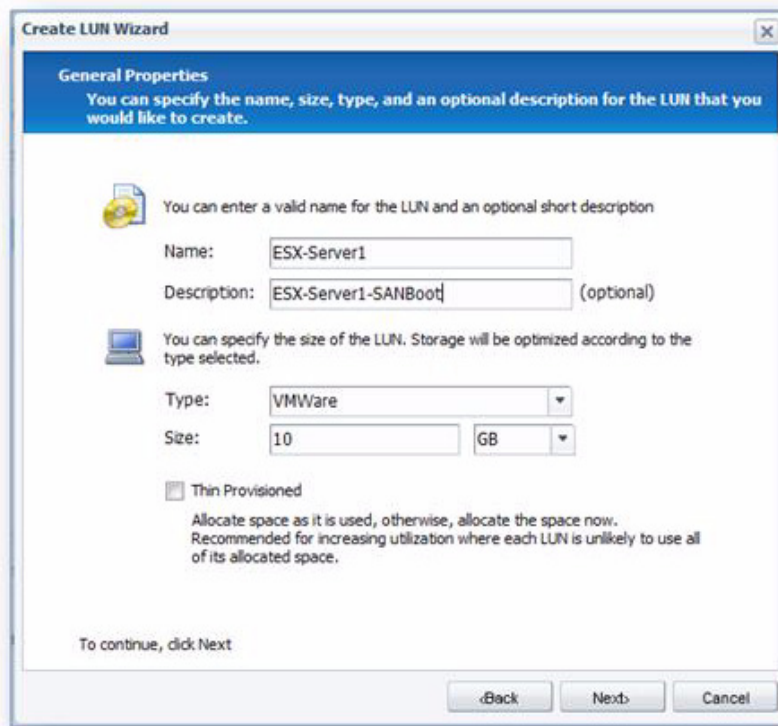
For this deployment, a NetApp 3270 Controller is used as a Storage device. Figure 39 gives an overview of SAN connectivity for the SharePoint deployment over Cisco UCS Blade servers.

Figure 27 **Storage Connectivity**



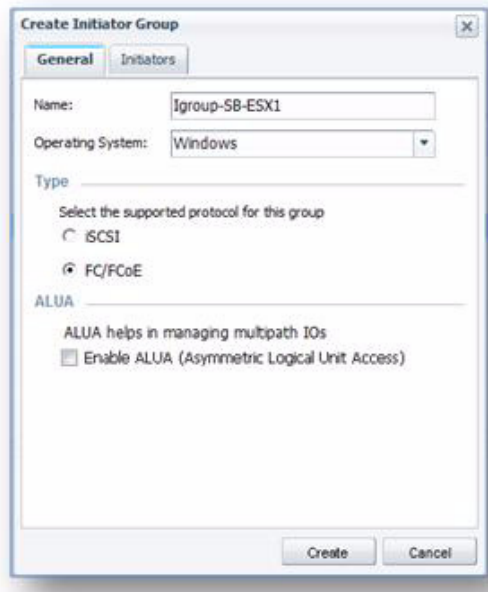
To configure the LUN parameters, perform the following steps:


1. Create a LUN on the storage using the NetApp Storage GUI.

**Note**

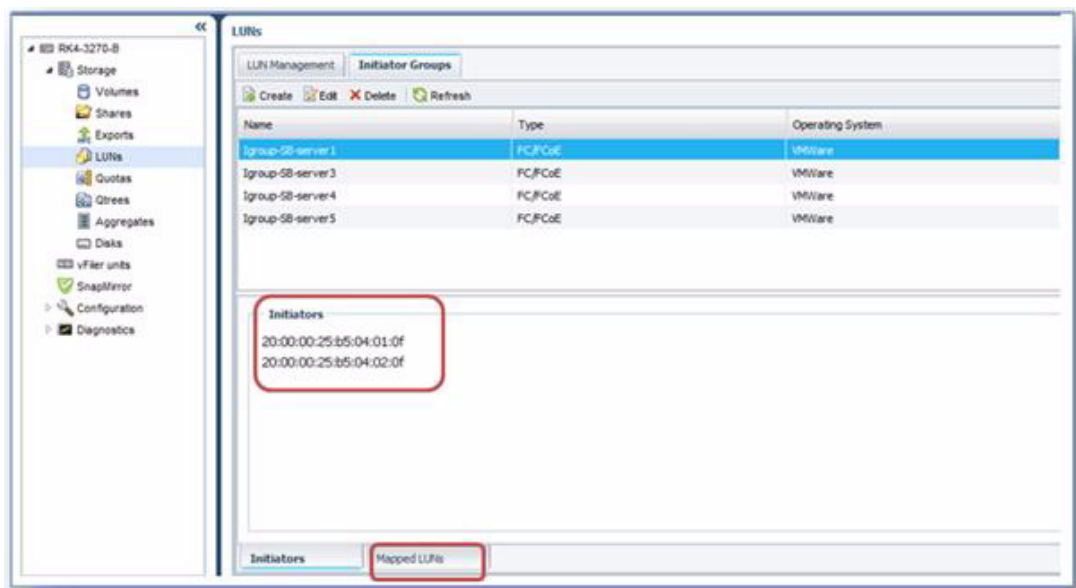
NetApp storage GUI is shown for your reference, but it is recommended to use NetApp OnCommand System Manager 2.0 when provisioning the storage or configuring a storage controller.

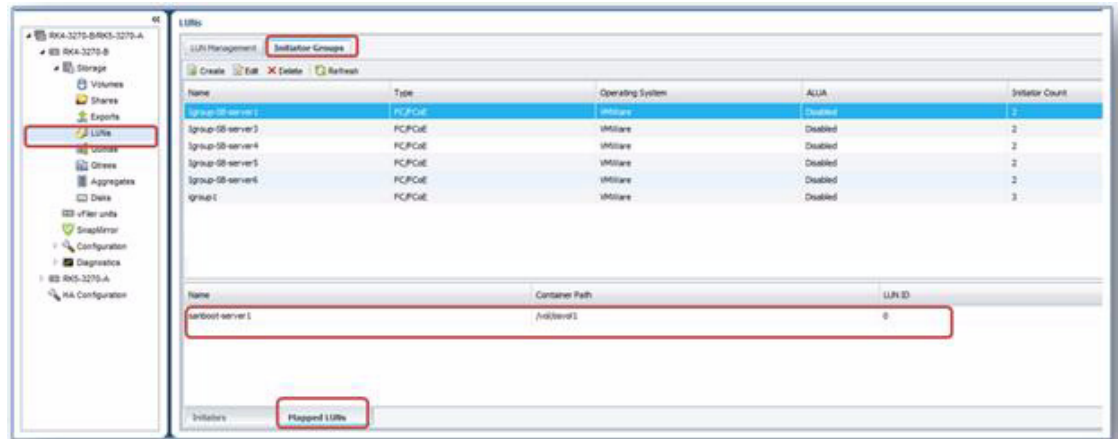
2. Create an igroup using the storage using the NetApp Snap Manager 2.0 R1.





3. Map the LUN to the initiator group using the storage using the NetApp Storage GUI.





You have now configured a Storage Group which has a 10G LUN.

In the next section, configuring a Boot policy in a Cisco UCS Service Profile is discussed. You need to have followed the steps for [Storage Array Configuration](#) to register the host.

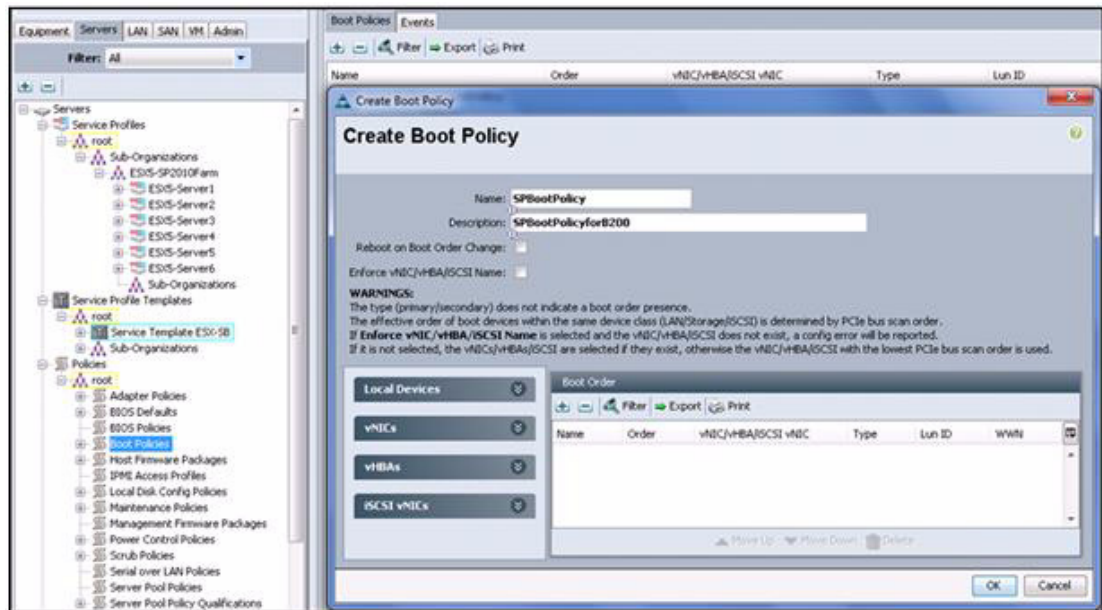
Now you need to identify the Host ID from the configured OS installation LUN. This Host ID will be mapped to the LUN ID during addition of SAN Boot Target in the Boot Policy of the Cisco UCS Service Profile.

Cisco UCS Manager Configuration

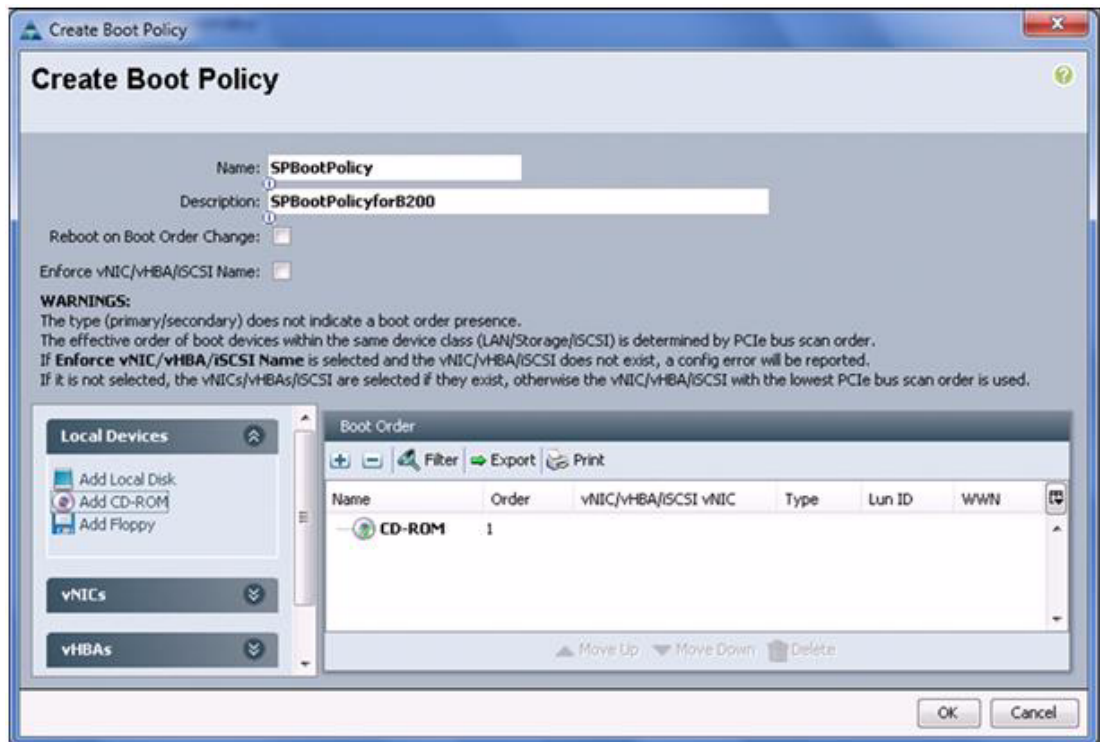
Perform the following steps to enable boot from SAN from the Cisco UCS Manager window:

1. Create a boot policy:
 - a. In the left pane of the Cisco UCS Manager window, click the **Servers** tab and select **Policies**.
 - b. From the right pane, click the **Boot Policies** tab. Enter the Name and the Description, check the **Reboot On Boot Order Change** check box and uncheck **Enforce VNIC/VHBA/iSCSI Name** check box.

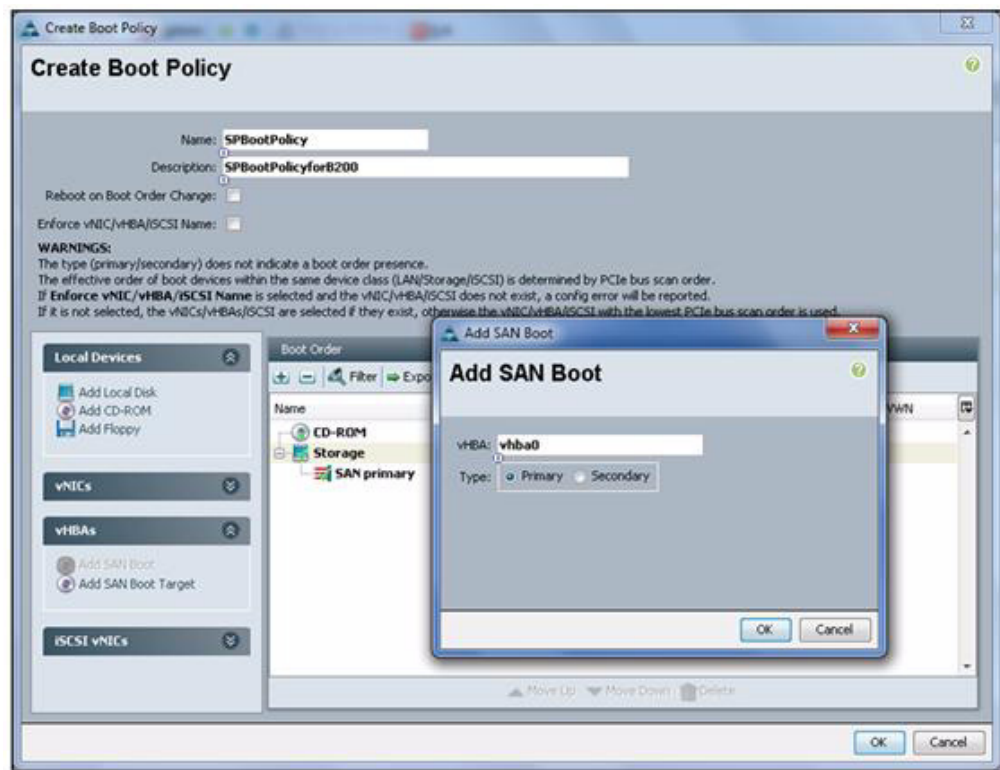
Create Boot Policy displays.



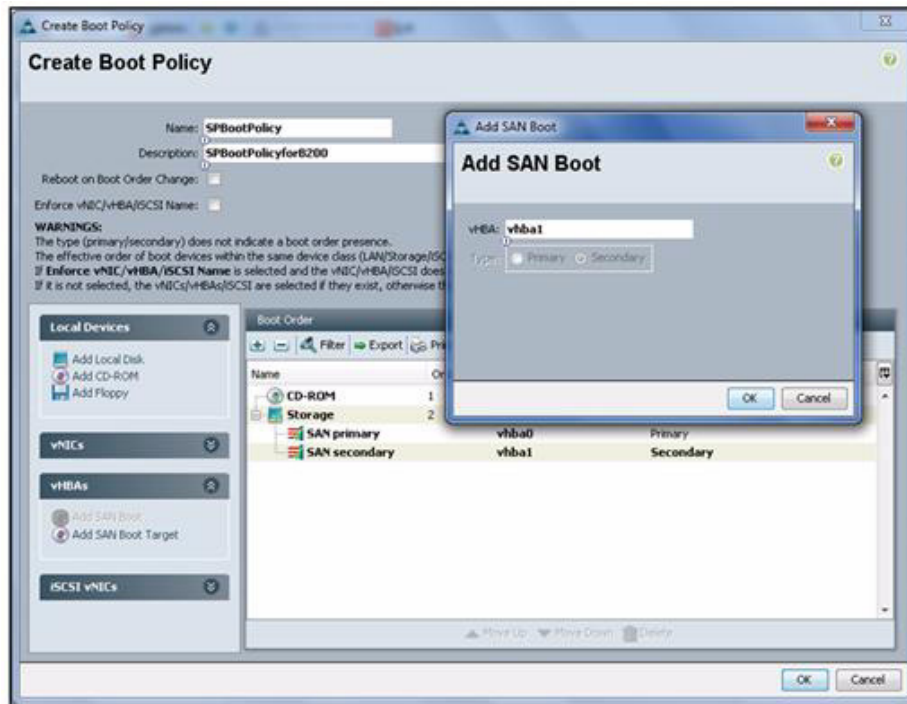
2. Add the first target as CD-ROM. This enables you to install the Operating System through the KVM Console.



3. Add SAN Boot for SAN Primary. Select the radio button **Primary** for the field Type in the ADD SAN Boot window.



4. Add SAN Boot for SAN Secondary. Select the radio button **Secondary** for the field Type in the ADD SAN Boot window.



- In this CVD, vhb0 and vhb1 are used for SAN Boot; for the SharePoint application server installation vhb2 and vhb3 are used. You will need to identify the SAN WWPN ports and add SAN Boot Targets.

The SAN Boot Targets that need to be added are:

vhba0

- Storage Port SP -1a - Primary Target - 50:0a:09:81:9d:93:40:7f
- Storage Port SP -1b - Secondary Target - 50:0a:09:84:8d:93:40:7f

rk3-N5548-A# sh flogi database

```
rk2-n5548-a# sh flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/25	3	0x2b000c	20:00:00:25:b5:04:01:0b	20:00:00:25:b5:00:00:0c
fc1/29	3	0x2b0001	20:00:00:25:b5:04:01:0f	20:00:00:25:b5:00:00:0f
fc1/29	3	0x2b0002	20:41:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/29	3	0x2b0007	20:00:00:25:b5:04:01:0c	20:00:00:25:b5:00:00:1d
fc1/29	3	0x2b0009	20:00:00:25:b5:04:01:0a	20:00:00:25:b5:00:00:1c
fc1/30	3	0x2b0000	20:42:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/30	3	0x2b0005	20:00:00:25:b5:04:01:0e	20:00:00:25:b5:00:00:0e
fc1/30	3	0x2b0006	20:00:00:25:b5:04:01:0d	20:00:00:25:b5:00:00:0d
fc1/30	3	0x2b0008	50:0a:09:81:9d:93:40:7f	50:0a:09:81:9d:93:40:7f
fc1/32	3	0x2b0004	50:0a:09:84:8d:93:40:7f	50:0a:09:80:8d:93:40:7f

Total number of flogi = 10.

rk3-N5548-B# sh flogi database

```
rk2-n5548-b# sh flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/25	3	0x15000c	20:42:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/29	3	0x150002	20:41:00:05:73:a2:66:80	20:03:00:05:73:a2:66:81
fc1/29	3	0x150005	20:00:00:25:b5:04:02:0e	20:00:00:25:b5:00:00:0e
fc1/29	3	0x150007	20:00:00:25:b5:04:02:0c	20:00:00:25:b5:00:00:1d
fc1/29	3	0x150009	20:00:00:25:b5:04:02:0a	20:00:00:25:b5:00:00:1c
fc1/30	3	0x150000	20:42:00:05:73:a2:66:80	20:03:00:05:73:a2:66:81
fc1/30	3	0x150006	20:00:00:25:b5:04:02:0d	20:00:00:25:b5:00:00:0d
fc1/30	3	0x150008	20:00:00:25:b5:04:02:0b	20:00:00:25:b5:00:00:0c
fc1/31	3	0x150004	50:0a:09:83:8d:93:40:7f	50:0a:09:80:8d:93:40:7f
fc1/32	3	0x150003	50:0a:09:84:9d:93:40:7f	50:0a:09:80:8d:93:40:7f

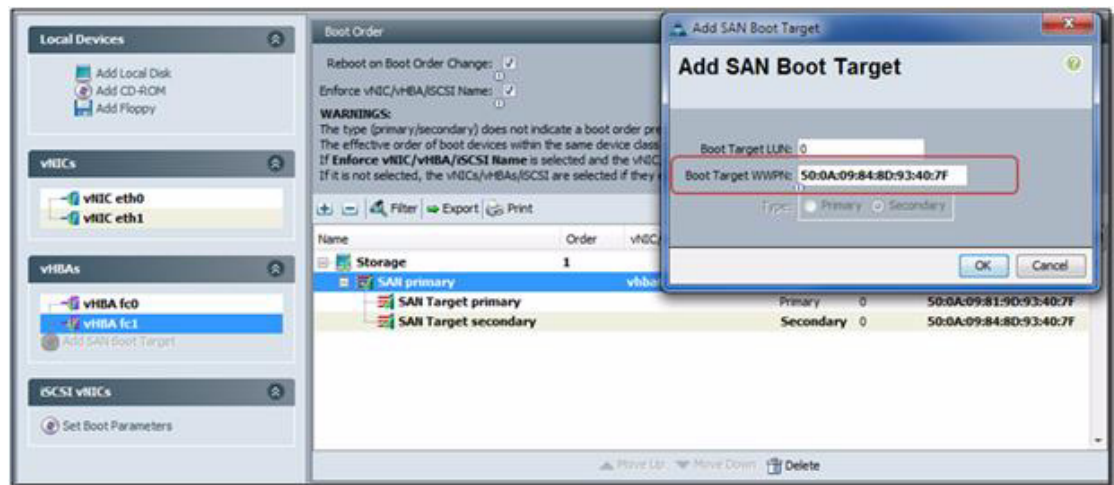
Total number of flogi = 10.

The final mapping is shown in the above screen shots.

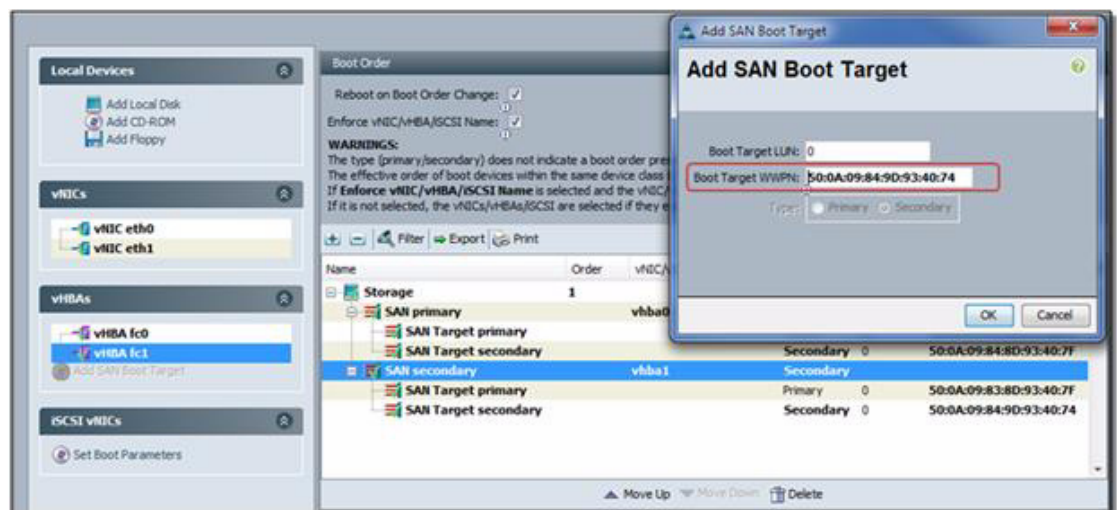
vhba1

- Storage Port SP-0c - Primary Target - 50:0a:09:83:8d:93:40:7f
- Storage Port SP-0d - Secondary Target - 50:0a:09:84:9d:93:40:7f

- Add the San Boot Target for SAN Primary.



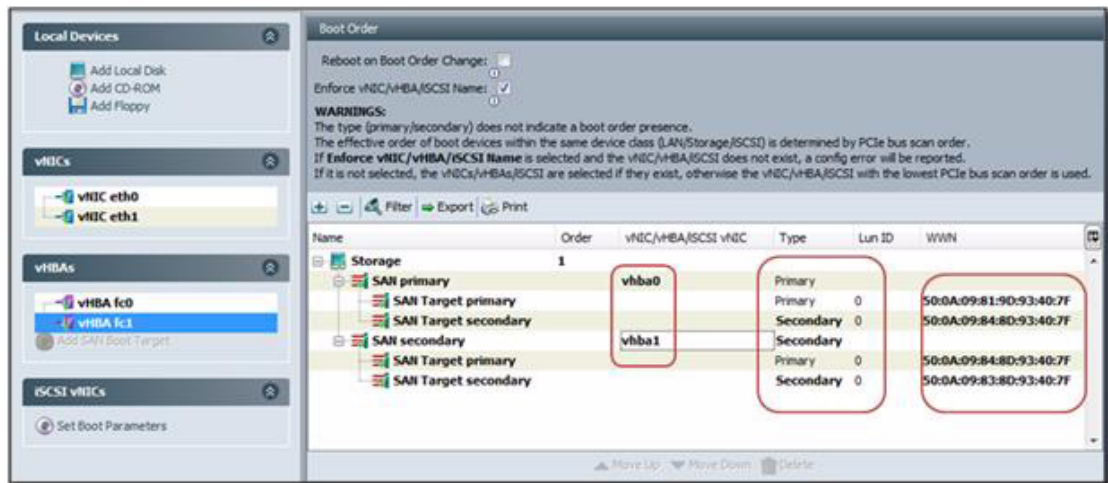
7. Add a secondary to the created Primary SAN Boot Target. You need to create the secondary under vhba0.




The Boot Target LUN ID is identified from the Host ID of the LUN created for SAN Boot in the Storage Group of the corresponding host.

8. Repeat step 6. and step 7. to add the Secondary SAN Boot Target for vhba1.

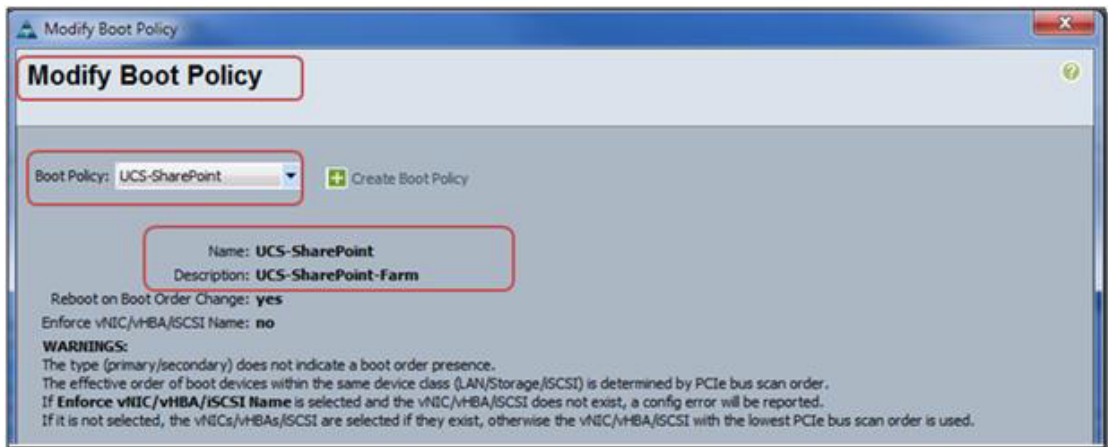
The SAN Boot Target Summary displays.




9. Associate the created Boot Policy to the Service Profile for SharePoint servers.



10. Select the SAN Boot Policy that you have created and Add to the Boot policy.



11. Click the **FSM** tab to view the Service Profile Modification.



You can see the WWPN of vhba0 and vhba1 in N5548-A and N5548-B after the server is rebooted as shown in the following screen shots.

```
rk2-n5548-a# sh flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/25	3	0x2b000c	20:00:00:25:b5:04:01:0b	20:00:00:25:b5:00:00:0c
fc1/29	3	0x2b0001	20:00:00:25:b5:04:01:0f	20:00:00:25:b5:00:00:0f
fc1/29	3	0x2b0002	20:41:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/29	3	0x2b0007	20:00:00:25:b5:04:01:0c	20:00:00:25:b5:00:00:1d
fc1/29	3	0x2b0009	20:00:00:25:b5:04:01:0a	20:00:00:25:b5:00:00:1e
fc1/30	3	0x2b0000	20:42:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/30	3	0x2b0005	20:00:00:25:b5:04:01:0e	20:00:00:25:b5:00:00:0e
fc1/30	3	0x2b0006	20:00:00:25:b5:04:01:0d	20:00:00:25:b5:00:00:0d
fc1/30	3	0x2b0008	50:0a:09:81:9d:93:40:7f	50:0a:09:81:9d:93:40:7f
fc1/32	3	0x2b0004	50:0a:09:84:8d:93:40:7f	50:0a:09:80:8d:93:40:7f

Total number of flogi = 10. WWPN of vhba0

```
rk2-n5548-b# sh flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/25	3	0x15000c	20:42:00:05:9b:77:4b:00	20:03:00:05:9b:77:4b:01
fc1/29	3	0x150002	20:41:00:05:73:a2:66:80	20:03:00:05:73:a2:66:81
fc1/29	3	0x150005	20:00:00:25:b5:04:02:0e	20:00:00:25:b5:00:00:0e
fc1/29	3	0x150007	20:00:00:25:b5:04:02:0c	20:00:00:25:b5:00:00:1d
fc1/29	3	0x150009	20:00:00:25:b5:04:02:0a	20:00:00:25:b5:00:00:1c
fc1/30	3	0x150000	20:42:00:05:73:a2:66:80	20:03:00:05:73:a2:66:81
fc1/30	3	0x150006	20:00:00:25:b5:04:02:0d	20:00:00:25:b5:00:00:0d
fc1/30	3	0x150008	20:00:00:25:b5:04:02:0b	20:00:00:25:b5:00:00:0c
fc1/31	3	0x150004	50:0a:09:83:8d:93:40:7f	50:0a:09:80:8d:93:40:7f
fc1/32	3	0x150003	50:0a:09:84:9d:93:40:7f	50:0a:09:80:8d:93:40:7f

Total number of flogi = 10. WWPN of vhb1

Zone Configuration

After creating the boot policy, you need to configure the zone on the Nexus 5548 and register the host on storage array. The following commands are used to configure VSAN and add zones in Nexus 5548 for VHBA configured in service profile of the Cisco UCS B200 blade server.

The following command configures the Nexus 5548 unified switch ports to FC ports.

```
rk2-n5548-b(config)# slot 1
rk2-n5548-b(config)# slot 1
rk2-n5548-b(config-if)# interface fc 1/25
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b(config-if)# interface fc 1/26
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b(config-if)# interface fc 1/27
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b(config-if)# interface fc 1/28
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b(config-if)# interface fc 1/29
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b(config-if)# interface fc 1/30
rk2-n5548-b(config-if)# switchport mode F
rk2-n5548-b(config-if)# no shut
rk2-n5548-b# config t
```

The following command executed on the Nexus 5548up switch creates zones for each ESXI Server and adds a corresponding PWWN number of storage and vHBA ports.

```

Enter configuration commands, one per line. End with CNTL/Z.
zone name ESX5-SP2010-Blade4 vsan 3
rk2-n5548-b(config)# zone name ESX5-SP2010-Blade4 vsan 3
rk2-n5548-b(config-zone)# member pwwn 20:00:00:25:b5:04:02:0c
rk2-n5548-b(config-zone)# member pwwn 50:0a:09:83:8d:93:40:7f
rk2-n5548-b(config-zone)# member pwwn 50:0a:09:84:9d:93:40:7f
rk2-n5548-b(config-zone)# copy r s
[#####] 100%
rk2-n5548-b(config)# zoneset name zoneset-SP2010 vsan 3
rk2-n5548-b(config-zoneset)# member ESX5-SP2010-Blade4
rk2-n5548-b(config-zoneset)# copy r s
[#####] 100%
rk2-n5548-b(config-zoneset)# zoneset activate name zoneset-SP2010 vsan 3
Zoneset activation initiated. check zone status
rk2-n5548-b(config)# sh zoneset active
zoneset name zoneset-SP2010 vsan 3

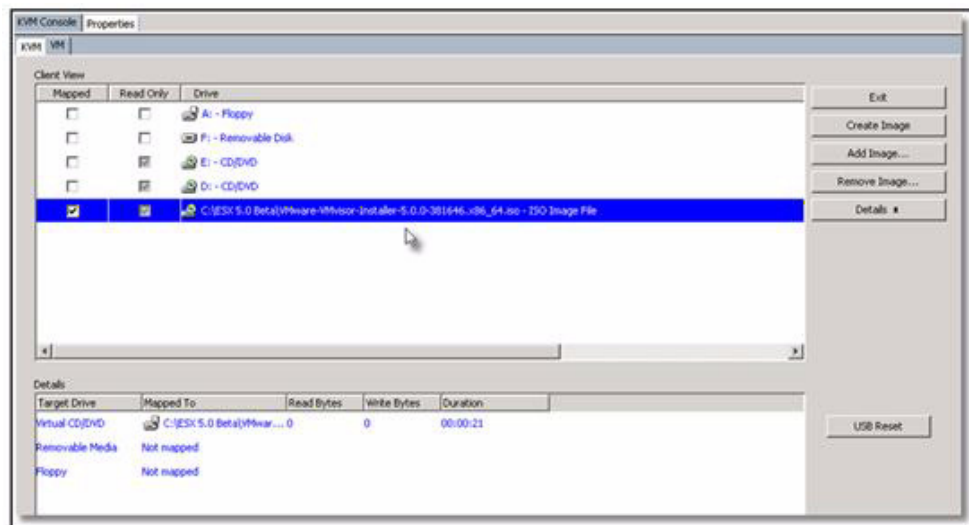
```

Now the Cisco UCS server blade is ready for OS installation using the SAN Boot LUN.


OS Installation

After the SAN and Service Profile configuration for Boot from SAN is completed, you need to start the ESXi Server installation process. Perform the following steps to start the OLE installation process:


1. From the Cisco UCS Manager, select the Service Profile and connect to the server through the KVM Console.
2. Select the VMware ESXi 5.0 ISO image through Launch Virtual Media. After selecting the ISO image, reboot the server and start the OS installation.




3. After the OS image loads, select the installer at the boot prompt .



4. Select **Enter** to continue with the installation.




5. Press the F11 key to accept the license agreement.
6. Select the NetApp LUN.



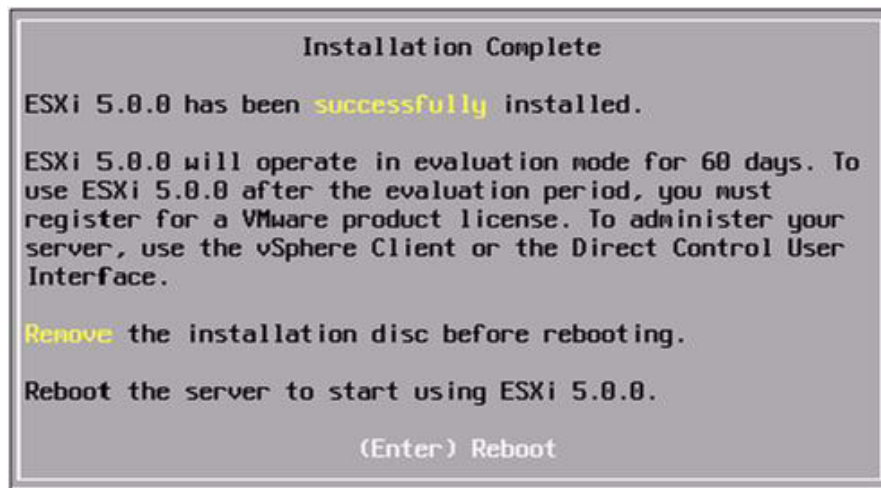
7. Select the Keyboard layout.




8. ESXi 5.0 is configured and needs to be installed on the NetApp LUN.



9. After the installation is complete, type Enter to reboot.



ESXi 5.0.0 is ready to use after the reboot.



VMware vCenter Server Deployment

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

- A running VMware vCenter virtual machine
- A running SQL virtual machine acting as the vCenter database server
- A vCenter DataCenter with associated ESXi hosts
- VMware DRS and HA functionality enabled

For detailed information about installing a vCenter Server, go to:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html.

For detailed information about vSphere Virtual Machine Administration, go to:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.install.doc_50/GUID-A71D7F56-6F47-43AB-9C4E-BAA89310F295.html.

For detailed information on creating a virtual machine in the vSphere 5 client, go to:

http://pubs.vmware.com/vsphere-50/index.jsp?topic=/com.vmware.vsphere.vm_admin.doc_50/GUID-0433C0DC-63F7-4966-9B53-0BECDD6420.html.

For detailed information about installing and configuring Microsoft SharePoint 2010 servers, go to:


<http://technet.microsoft.com/en-us/library/cc262957.aspx>

Template-Based Deployments for Rapid Provisioning

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

Virtual machine templates can help speed up this process by eliminating most of these monotonous tasks. By completing the core installation requirements, typically to the point where the application is ready to be installed, you can create a golden image which can be sealed and used as a template for all of your virtual machines. Depending on how granular you want to make a specific template, the time to deployment can be as little as the time it takes to install, configure, and validate the application. You can use PowerShell tools for SharePoint and VMware vSphere PowerCLI to bring the time and manual effort down dramatically (Figure 64).

Figure 28 **Rapid Provisioning**



For more information about Microsoft SharePoint 2010 servers, refer to the *Installation and Configuration of Microsoft SharePoint 2010 Servers* available at:

<http://technet.microsoft.com/en-us/library/cc262957.aspx>

Perform the following steps to deploy VMware on virtual machine vCenter:

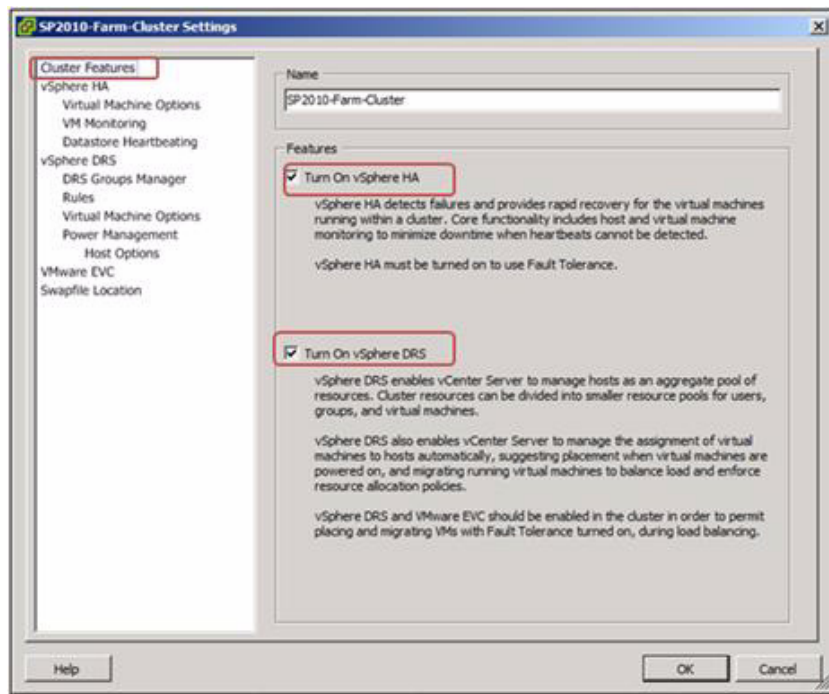
1. Log into VMware ESXi Host using VMware vSphere Client.
2. Build a SQL Server VM using Windows Server 2008 R2 x 64 image.
3. Create the required databases and database users using the script provided in the vCenter installation directory.



Note

VMware vCenter can use one of a number of vendor databases. This deployment assumes Microsoft SQL Server 2008 R2. If a database server already exists and it is compatible with vCenter you can create the required database instance for vCenter and skip this step.

4. Build a vCenter virtual machine on another Windows Server 2008 R2 virtual machine instance.
5. Install SQL Server 2008 R2 Native Client on the vCenter virtual machine.
6. Create Data Source Name referencing the SQL instance on the vCenter machine.
7. Install VMware vCenter Server referencing the SQL server data source previously established.
8. Create a vCenter Datacenter.
9. Create a new management cluster with DRS and HA enabled;
 - a. Right-click on the cluster and in the corresponding context menu, click **Edit Settings**.
 - b. Select the checkboxes **Turn On vSphere HA** and **Turn On vSphere DRS**.
 - c. Click **OK** to save changes.



DRS Affinity Rules

In this experiment, five ESXi hosts in VMware DRS Affinity Rules have been configured. These rules allow assignment of individual VMs or Groups of VMs to individual ESXi Servers or Groups of ESXi Servers by configuring VMware DRS Affinity Rules. These rules also govern the placement of virtual servers on ESXi host.

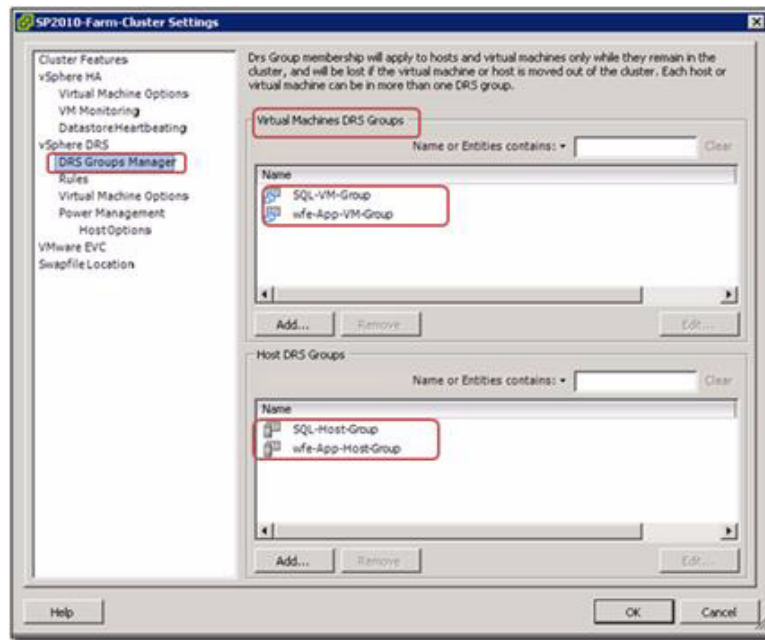
In the SharePoint 2010 farm web front-end tier, application tier virtual servers are pinned to a group of Cisco UCS B200 M2 blade servers, and database servers are pinned to another group of Cisco UCS B250 Servers.

Cisco B250 M2 servers are well suited for the database tier as these servers have two-socket blade servers featuring Cisco Extended Memory Technology. This system supports two Intel Xeon 5600 Series Processors, up to 384 GB of DDR3 memory and two mezzanine connections for up to 40 Gbps of I/O throughput. The server provides increased performance and capacity for demanding virtualization and large-data-set workloads with greater memory capacity and throughput.

DRS is running in fully automated mode. In this mode, vSphere independently decides the VM to be assigned to the ESXi Server.

Perform the following steps to set the DRS to run in fully automatic mode:

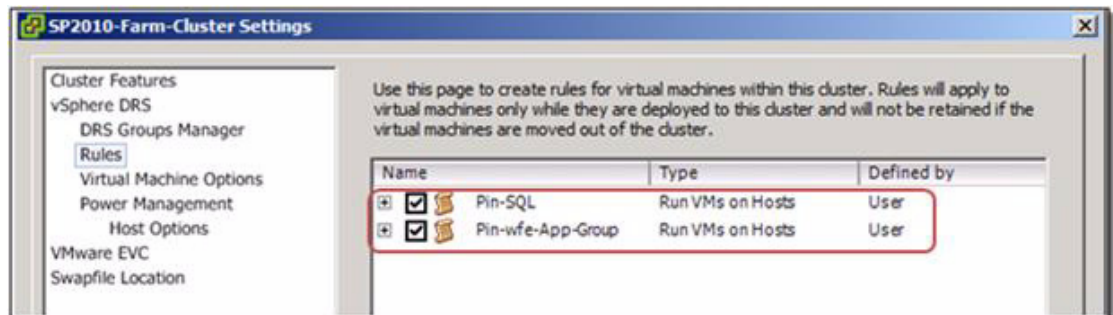
1. Open the vSphere Client and find the cluster on which the rules will be enforced.
2. Right-click on the cluster and in the corresponding context menu, select **Edit Settings**.
3. In the left pane of SP2010-Farm-Cluster Settings window, choose **DRS Groups Manager**.
4. In the right pane of the window, under Virtual Machines DRS Groups click **Add**.



5. Select the virtual machines to add to the virtual machine group.
6. Add the selected virtual machines to the VM group. When you are done, provide a group name, then click **OK**.

In the DRS Groups Manager, you can create a Host DRS Group. This is the host DRS group where the newly created DRS groups for virtual machine will be assigned.

7. In the Host DRS Groups and click **Add**.
8. Select the hosts that you want to add to the host DRS group and select the **DRS Groups Manager**.



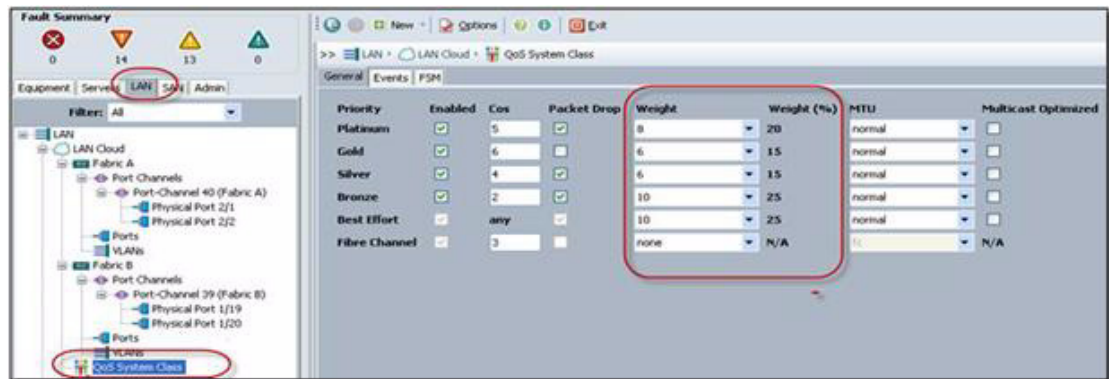
VMotion Configuration

In order to successfully use vMotion, you must first configure your hosts correctly.


Ensure that you have correctly configured the hosts in each of the following areas:

- Each host must be correctly licensed for vMotion.
- Each host must meet shared storage requirements for vMotion.
- Each host must meet the networking requirements for vMotion.

In our scenario, a dedicated VLAN for ESXi host is configured with platinum QoS Policy.



VMotion port groups are defined as VMkernel port groups since they are access ports for the ESXi hypervisors.



Network Configuration

The advantages of Universal Pass Through Switching (UPT) are as follows:

- VN-Link in Hardware with VMDirectPath Hypervisor Bypass
- Pointers passed between adapter DMA engine, and VM memory space
- Direct reads/writes from guest's OS buffers to adapter
- Bypass hypervisor for CPU savings
- Each interface handled in parallel with full FI exposure
- Uses Nexus 1000V kernel bits (cannot co-exist on a server)
- VMotion supported

VM-FEX Configuration for VMware Environment

This section includes the following topics:

- Configuration steps to integrate VM-FEX with vCenter Manager,
- Define VM-FEX Cluster VMware Datacenter with Port profile / client
- Install / configure ESXi host based VM-FEX VEM (Virtual Ethernet Module) Module
- Apply port profiles to VM's in VMware vCenter Manager

Prerequisites

The following are the hardware and software components that need to be installed and configured in your environment for VM-FEX configuration:

- Hardware Component
 - Cisco UCS System
 - Full or Half Width Blade
 - Dual or Single Cisco VIC adapters
- Software Component
 - vCenter Application Installed and loaded on system (Physical machine / Virtual machine)
 - Cisco UCS Manager
 - VM-FEX VEM Modules for ESXi Host

VM-FEX UCS Configuration

vCenter Extension Download

vCenter requires a “plug-in” to be installed that contains information about the Cisco UCS Manager that will connect to it for registration as a VDS. The plug-in (“extension file” or “extension XML file”) typically contains the following authentication information:

- Extension key
- Public SSL certification

Downloading Extension Keys

To download the extension keys, perform the following steps:

1. Launch the Cisco UCS Manager.
2. From the left navigation pane, select the **VM** tab and click the VMware link.




3. Click **Export vCenter Extension** located under Actions in the right pane of the Cisco Unified Computing System Manager as.
 - If the vCenter version is 4.0 U1 and above, you need to download a single extension XML file. To download a single extension XML file, click **Export vCenter Extension**.
 - If the vCenter version is less than 4.0, you will need to download multiple vCenter Extensions. To download multiple vCenter Extensions XML file, click **Export Multiple vCenter Extensions**.



4. Download the relevant extension keys, verifying with the vCenter version in your environment and save them to your local disk or to a shared network folder path.


Note

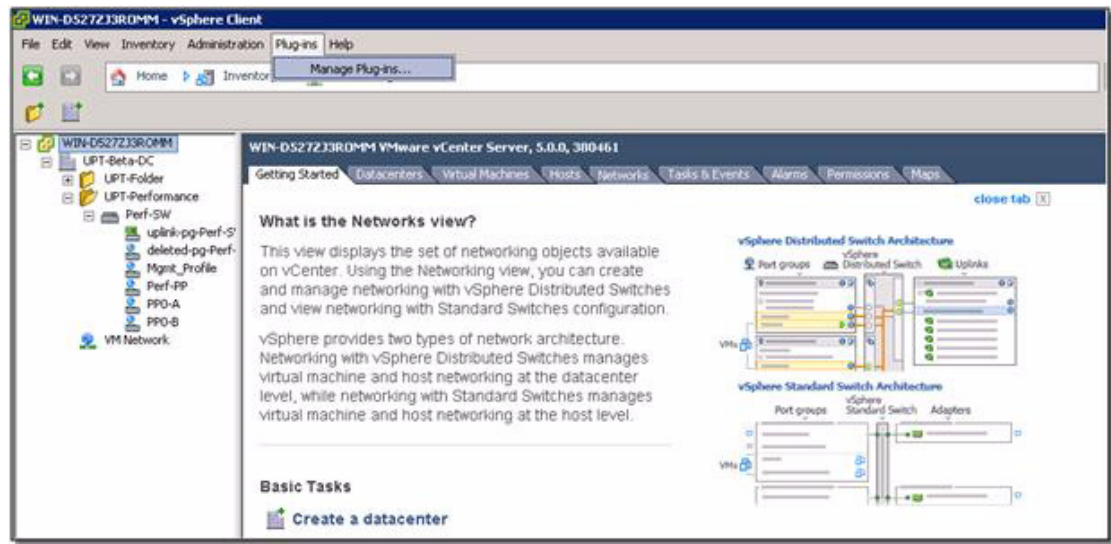
The folder specified for downloading the extension key must be created manually. Cisco UCS Manager does not create this folder.



vCenter Plug-in Registration

After downloading the Cisco UCS VM-FEX extension key, you need to register the key with the vCenter Manager using the plug-in provided by vCenter. This is required for vCenter to establish a connection with the Cisco UCS Manager for downloading the port profiles defined by Cisco UCS VM-FEX.

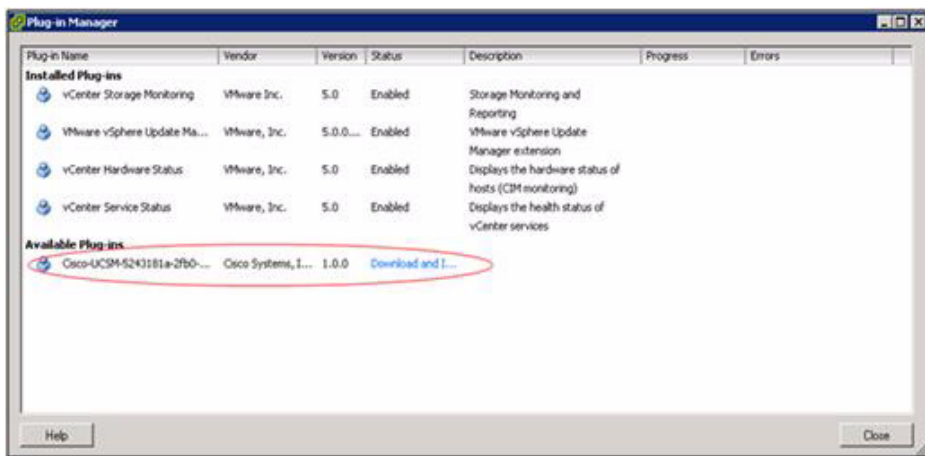
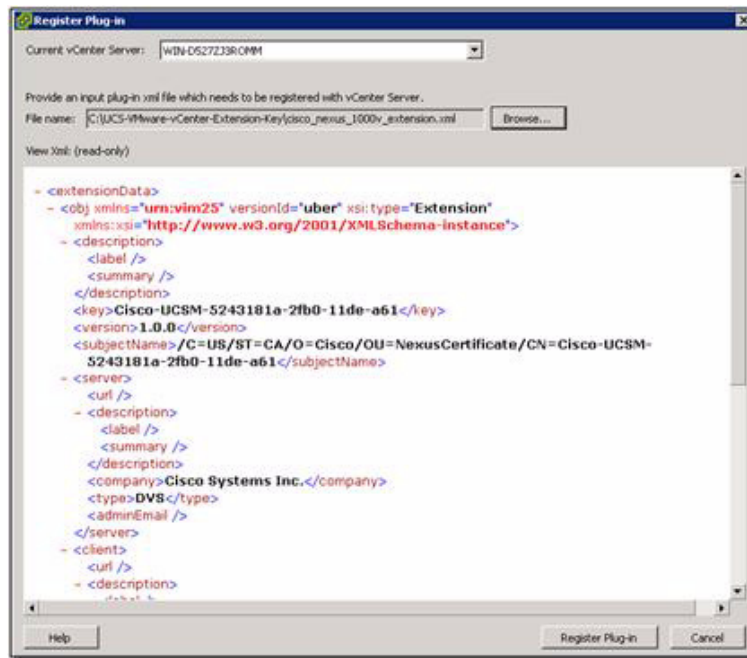
1. Launch the vCenter Manager.
2. Click **Plug-ins** and select **Manage Plug-ins** from the drop-down menu.



3. Right-click on the Plug-in Manager window and click **New Plug-in** to register it.



4. In the Register plug-in window, click **Browse** and navigate to the path of the Cisco UCS Extension Key that is already downloaded. Click **Register Plug-in**.
5. When vCenter registers the Cisco UCS VM-FEX Extension Key in the Plug-in, you can verify this under the Available Plug-ins in the Plug-in Manager window.

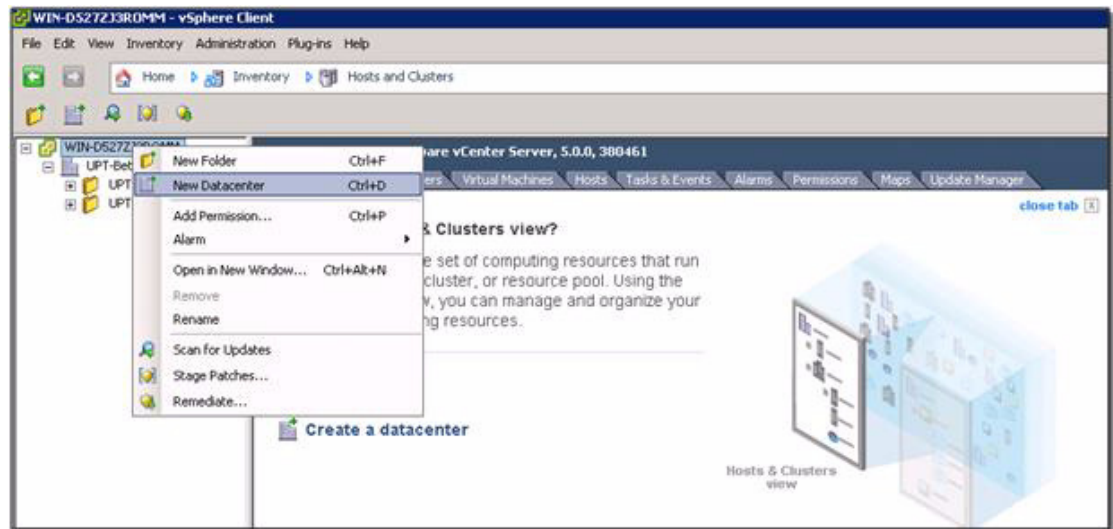


vCenter Datacenter Creation

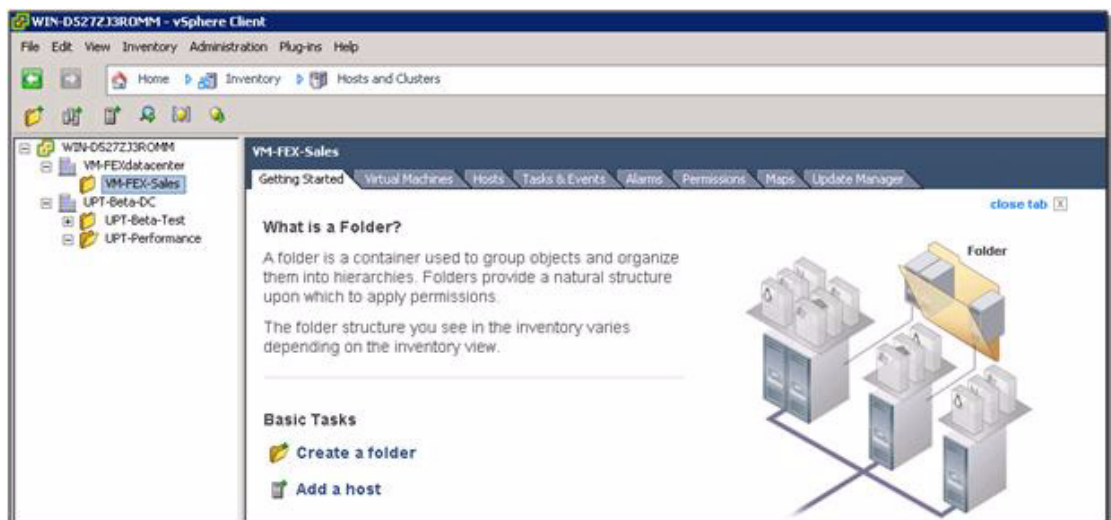
After installing the vCenter Plug-in, you need to create the vCenter Datacenter hierarchy to add host, folders and clusters to a datacenter and add datacenters in vCenter.

A single VM-FEX switch can belong to one datacenter under the vCenter and span across hosts which belong to the same datacenter. Therefore, you need create similar hierarchy that you defined in vCenter datacenter when defining Cisco UCS VM-FEX in Cisco UCS Manager.

1. Launch the vCenter Manager and right-click on the vCenter Name.
2. From the drop-down list select **New Datacenter**.



3. Provide a valid name for the new Datacenter.
4. After the datacenter is created, create a folder under the new datacenter.



VM-FEX DVS Switch Configuration

After creating a new datacenter with its folder on the vCenter, you need to create a VM-FEX DVS switch by following the same hierarchy as vCenter datacenter.

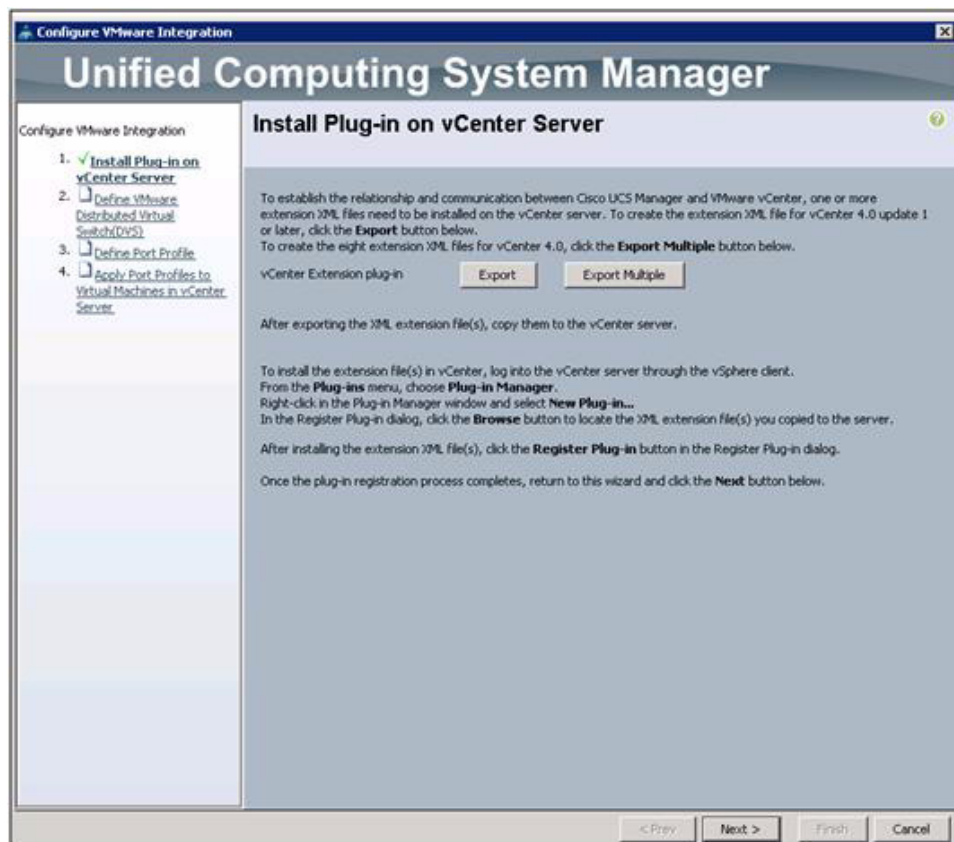
You can configure VM-FEX either by choosing VMware Integration in Cisco UCS Manager or by using Configure vCenter. In the VMware Integration using the Cisco UCS Manager method, you can complete the entire configuration in a customized manner. This provides end to end configuration without human error. In the Configure vCenter method, you can configure settings like port profile, client, etc. at a later time.

In this configuration, choose VMware Integration option for configuration and perform the following steps to configure the switch:

1. Launch the Cisco UCS Manager and click the **VM** tab in the left pane of the Cisco UCS Manager.
2. Click **Configure VMware Integration** in the right pane of the Cisco UCS Manager. Ignore the extension key since the Plug-in is already registered in the vCenter.
3. Click **Save Changes**.

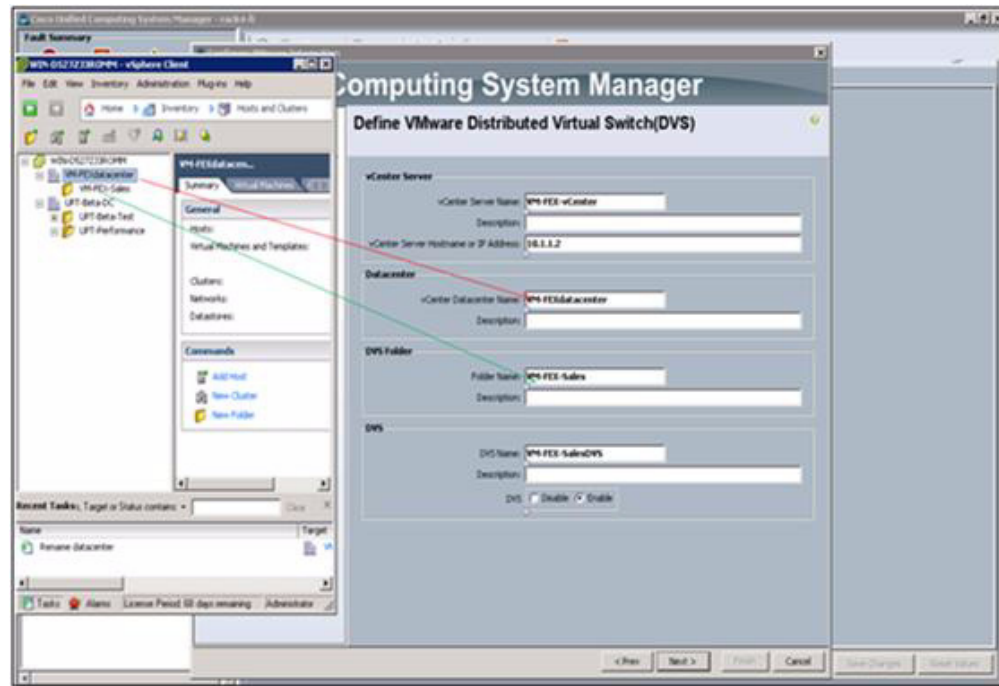


4. Click **Next**.

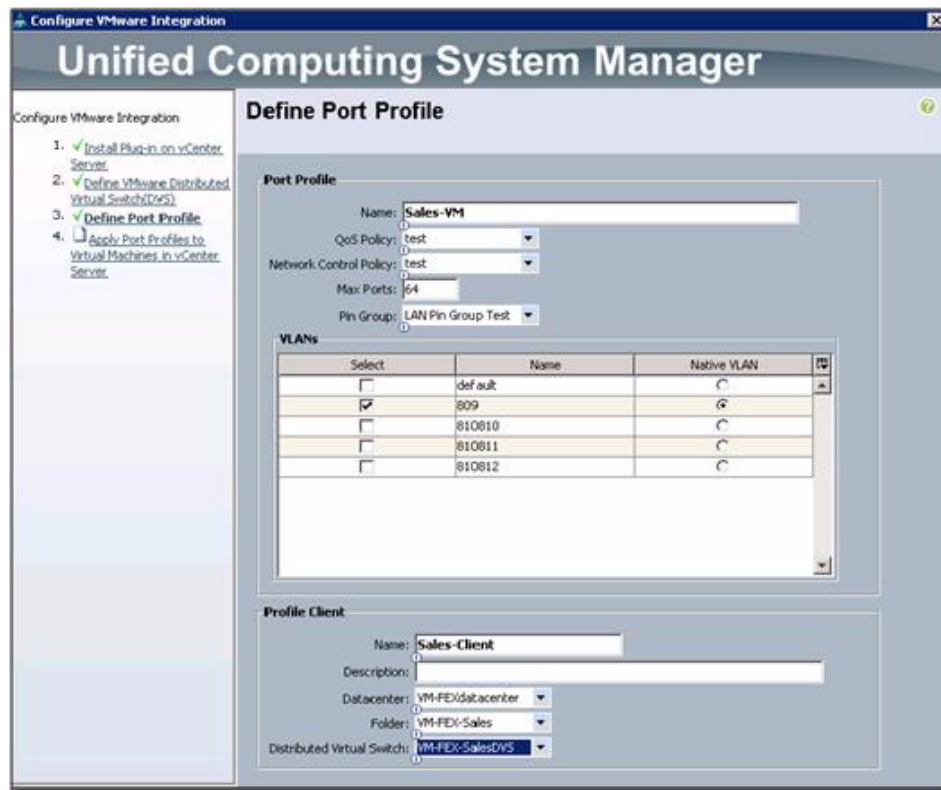


5. In defining VMware Distributed Virtual Switch configuration, you need to have the same naming convention as defined during vCenter new datacenter and have the same folder.
 - a. Under Datacenter, the vCenter Datacenter Name for VMware DVS configuration should match the vCenter Datacenter Name in the vCenter Manager.

- b. The folder name under VMware DVS should match with the vCenter Datacenter folder name created under datacenter.
- c. Click **Enable** to automatically turn on the DVS switch to operational mode.
- d. Click **Next**.



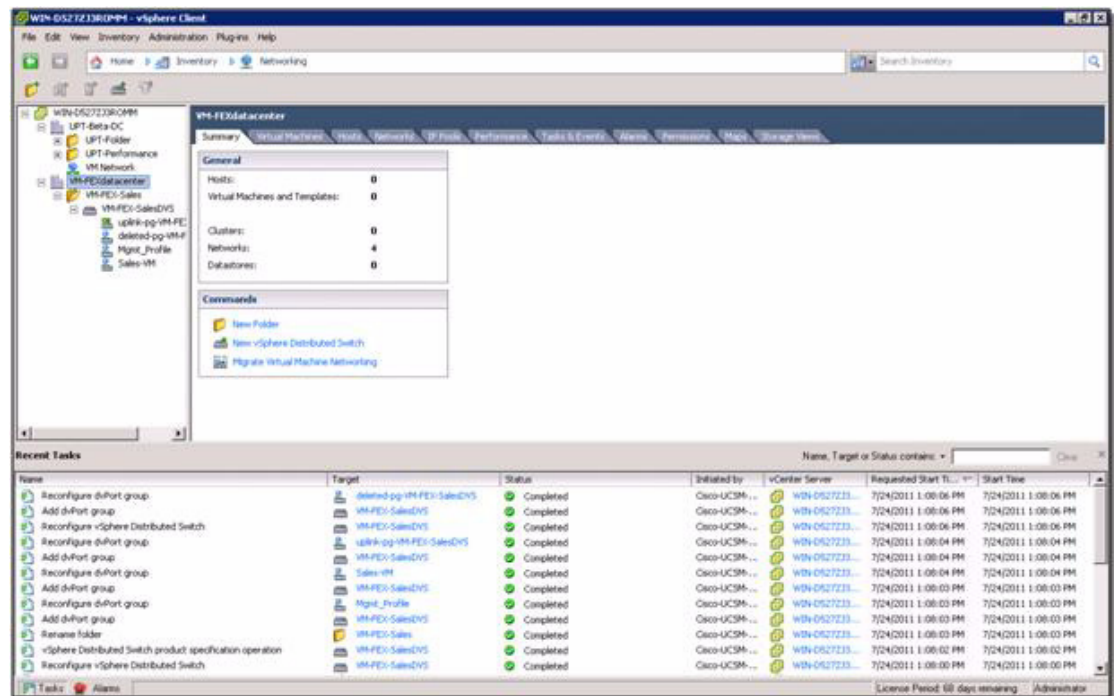
6. The define Port Profile/Client window provides options to define network properties and then define Port Client (vCenter Datacenter). The Profile Client displays all the available Datacenter that you had created in the Cisco UCS VM tab. You need to choose the Port Profile to be applied to a specific Datacenter.
7. Click **Next**.



8. Click **Finish** in the Apply Port profiles to Virtual machines in vCenter Server window to create the VM-FEX-DVS switch in the vCenter Manager under the VM-FEX datacenter.

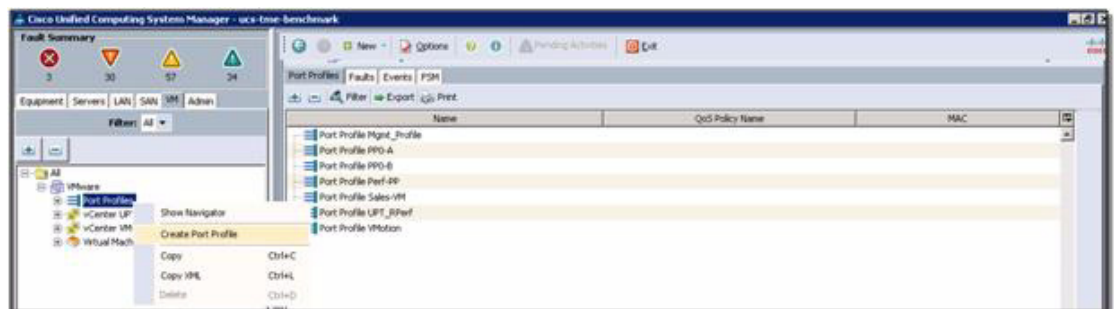
After the creation of VM-FEX-DVS through the Cisco UCS Manager you can verify the completion status of the VM-FEX-DVS in the vCenter.

9.

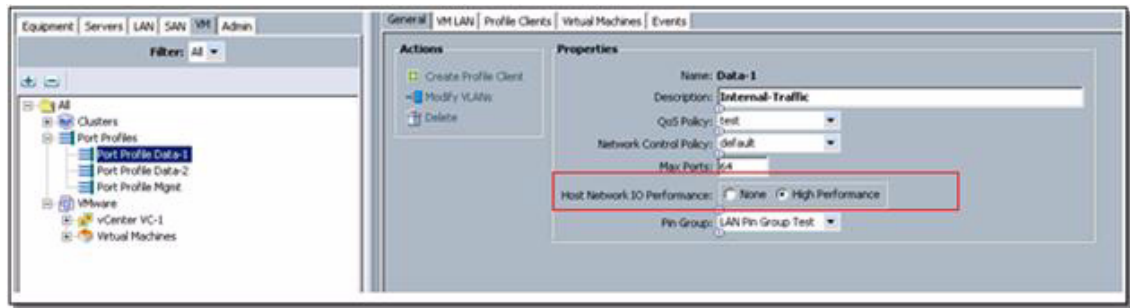


You can manually create Port Profiles and apply them to the relevant Port Client of the desired vCenter Datacenter. You can create different Port Profiles with different network policies for Vmotion, Fault Tolerant (FT) or to virtual machine's data port groups in vCenter.

10. Launch Cisco UCS Manager and expand the VMware link located in the left pane of the Cisco UCS Manager window.
11. Right-click on **Port Profiles** and choose **Create Port Profiles** from the drop-down list.



12. In **Create Port Profile**, enter information for all of the fields.



13. Select **High Performance** for Host Network I/O Performance.

ESXi Host Preparation

This section describes the steps to prepare ESXi host for installing the VM-FEX VEM module. To prepare the ESXi host you need to perform the following:

- Install a Cisco VIC
- On the Host Enabling Virtualization Technology in BIOS Settings
- Assign a UCS Service Profile
- Define Dynamic Ethernet Policy for VM
- Attach VM to VM-FEX DVS switch running inside the ESXi host

Assign a Cisco UCS Service Profile

To assign the Cisco UCS Service Profile, perform the following steps:

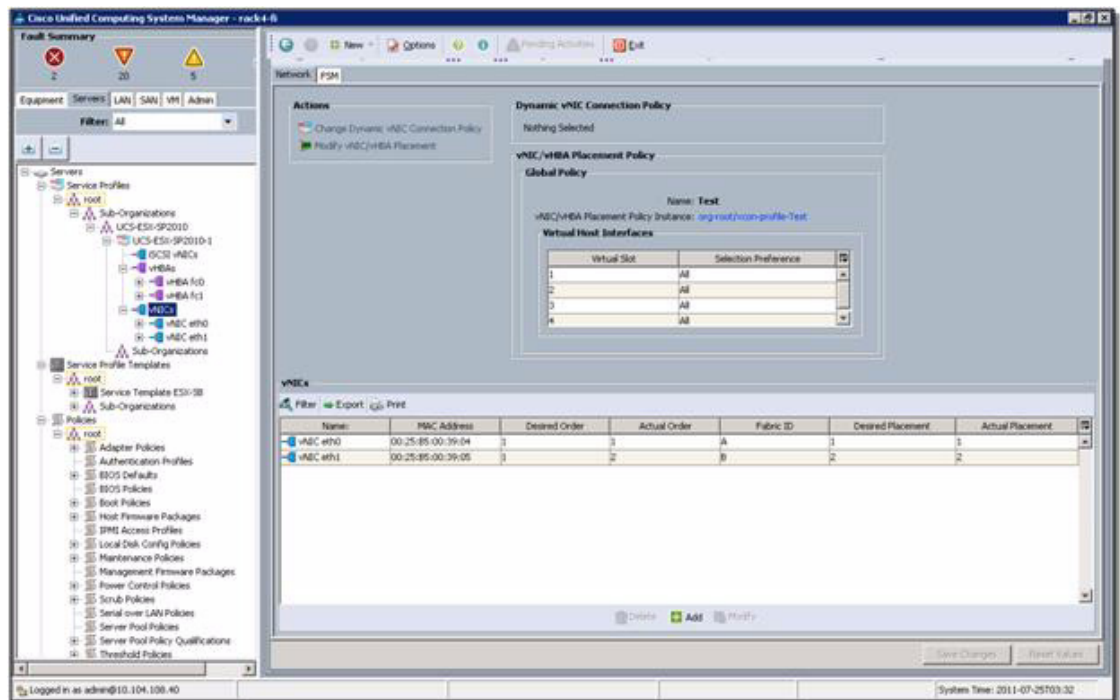
1. Create and associate the service profile for the ESXi blade that you have inserted on the Cisco VIC adapter. Within the service profile, create two ethernet static VNICs on both sides of fabric A and B (No Fabric Failure) and if applicable, create FC VHBAs for SAN Connectivity.

When the VM-FEX is connected to virtual machines, the VM-FEX does not use the static VNICs to apply them on virtual machines. The static VNICs are primarily used for Service Console connectivity (before ESXi host is added to vCenter) and migration of VM-FEX in the hardware.

Define Dynamic Ethernet Policy for Virtual Machines

To define the dynamic ethernet policy, perform the following steps:

1. After the blade is associated with the service profile, create the Dynamic Ethernet policy and apply on the VNICs.

**Note**

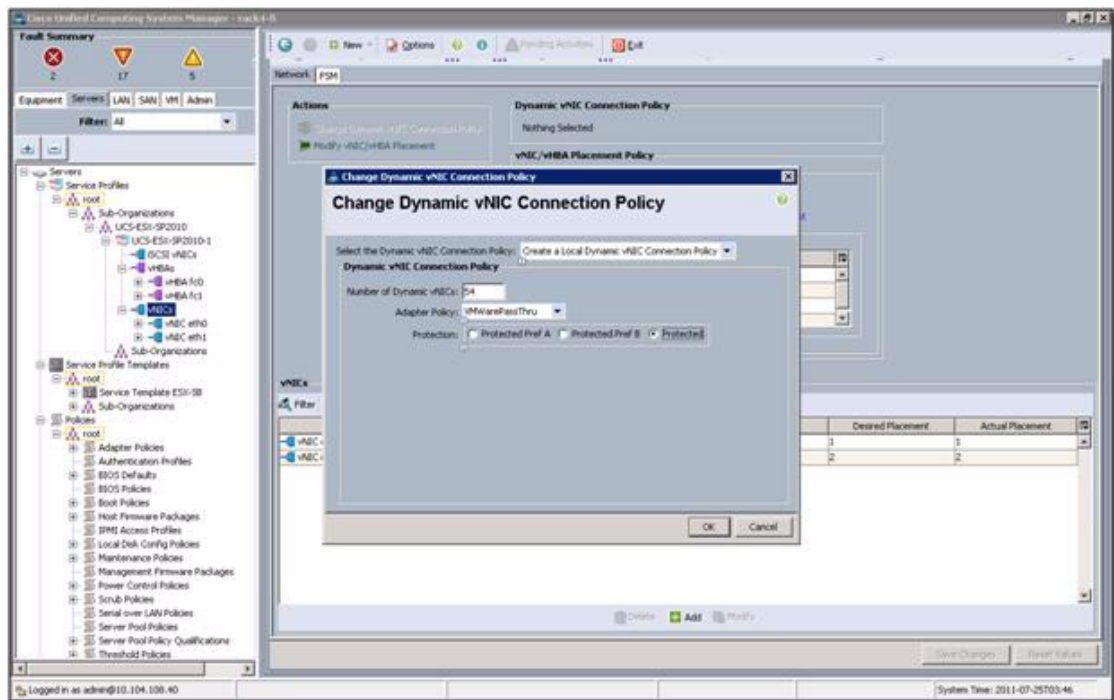
The Dynamic Ethernet VNICs have one-to-one mapping with the virtual machine VNICs that are created on the ESXi host during the service profile association, but do not contain any network properties like VLAN, QOS, etc.

The number of Dynamic VNICs that can be created depends on the number of acknowledged links between IOM and Fabric Interconnect.

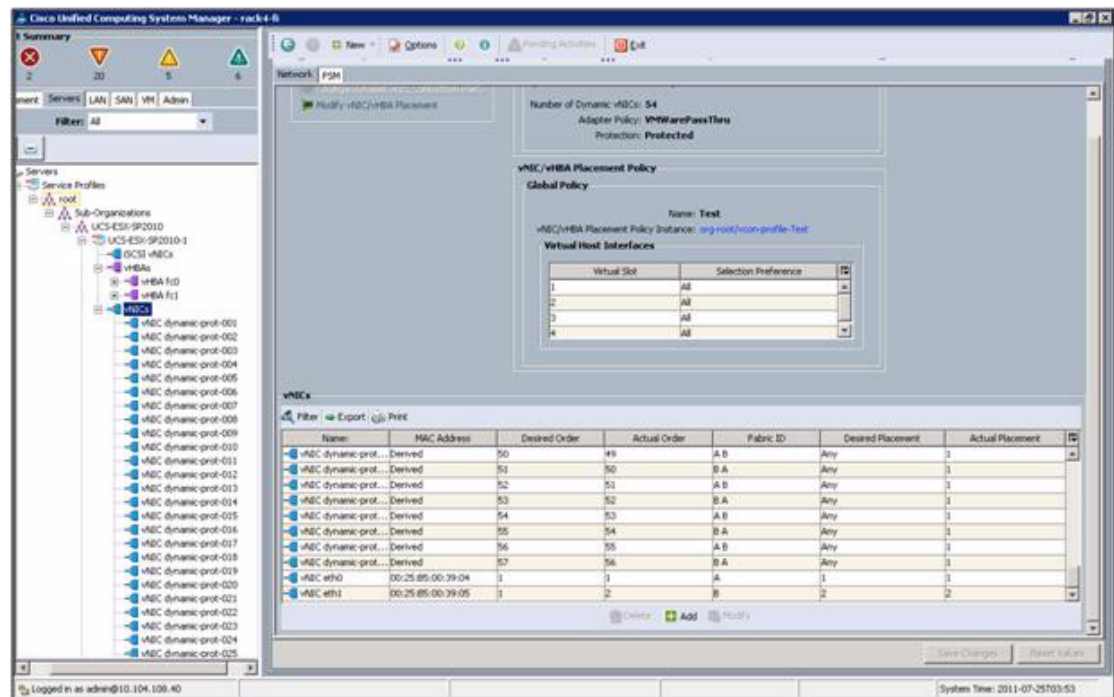
For more information about Dynamic VNICS, refer to:

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/vm_fex_best_practices_deployment_guide_ns1124_Networking_Solutions_White_Paper.html

2. Choose the VNIC option from the **Servers** tab on the left pane of the Service Profile.
3. Click **Change Dynamic VNIC Connection Policy** under Actions located in the right pane of the Service Profile.
4. Define the number of Dynamic VNICs you need, based on the formula explained with the adapter policy. Set the Adapter Policy to VMWarePassThru and set the fabric failure protection option to **Protected**.



After applying the Dynamic VNIC Policy, the service profile will go into the reconfigure phase.



- Attach VM to VM-FEX DVS Switch Running Inside the ESXi Host.
- Open the KVM console on the Service Profile and install ESXi OS.

- When the installation is complete you can run `lspci` command on the ESXi console and see if all the Dynamic Ethernet devices are listed.

```
# lspci | grep Dynamic
000:010:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic0]
000:011:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic1]
000:012:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic2]
000:013:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic3]
000:014:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic4]
000:015:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic5]
000:016:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic6]
000:017:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic7]
000:018:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic8]
000:019:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic9]
000:020:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic10]
000:021:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic11]
000:022:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic12]
000:023:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic13]
000:024:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic14]
000:025:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic15]
000:026:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic16]
000:027:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic17]
000:028:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic18]
000:029:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic19]
000:030:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic20]
000:031:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic21]
000:032:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic22]
000:033:00.0 Network controller: Cisco Systems Inc Cisco VIC Ethernet NIC Dynamic [vf_vmmnic23]
```

VM-FEX VEM Installation on ESXi Host

VM-FEX VEM module installed on the ESXi host is responsible for path control. The implementation maintains and conveys the control information about the emulated interface to the virtual switching layer. Information such as the MAC address of the vNIC as provisioned by VMware vSphere, the MAC address filters that will be implemented, and the connected or disconnected state of the vNIC to the virtual switch.

Perform the following steps to install VM-FEX on the ESXi Host:

- After installing the ESXi host, enable the SSH access to the root user by default SSH is not enabled.

Perform the following steps to enable the SSH access:

- On the ESXi Host start screen press **F2** for “Customize System.”
- Login with the local password.
- Select and enter the troubleshooting options.
- Click **Enable SSH** to activate SSH on your VMware ESXi 5.0 host.

- Download the corresponding VM-FEX VEM modules for the ESXi versions to a shared path and copy (SCP) these VEM agents to ESXi host. Using the following command you can install VM-FEX VEM:

```
ESXiupdate -b <VEM> --maintenancemode update -nosigcheck
```

- After installing the VM-FEX VEM, check to make sure the VEM modules are loaded into the kernel using the following command:

```
# vmkload_mod -l | grep vem
```

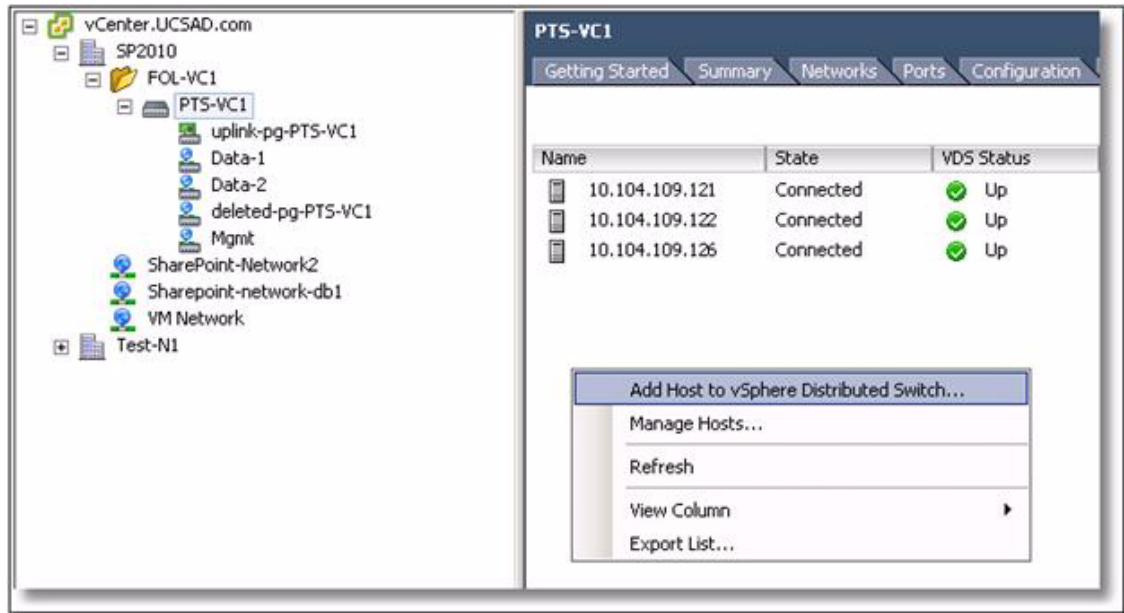
When you reboot the ESXi host (optional), you may see a warning message that the VEM modules are loaded, but has no signature attached on the ESXi console. You can ignore this message.

Register ESXi Host in the vCenter

After installing the VEM modules on the ESXi host, you need to attach the VEM modules to the appropriate datacenter in vCenter.

Perform the following steps to attach the VEM modules to the datacenter:

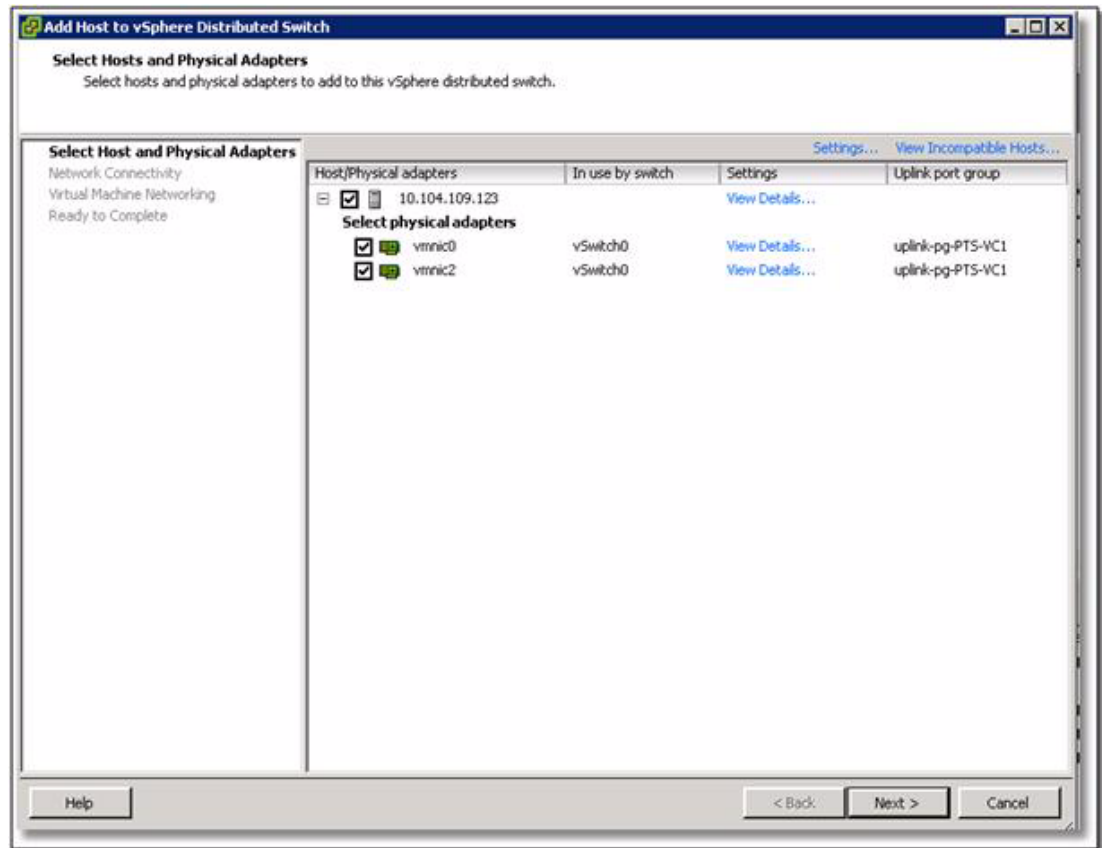
1. Select the ESXi host to be added to the vCenter datacenter.
2. Right-click the ESXi host and select **Add Host to Vsphere Distributed Switch**.



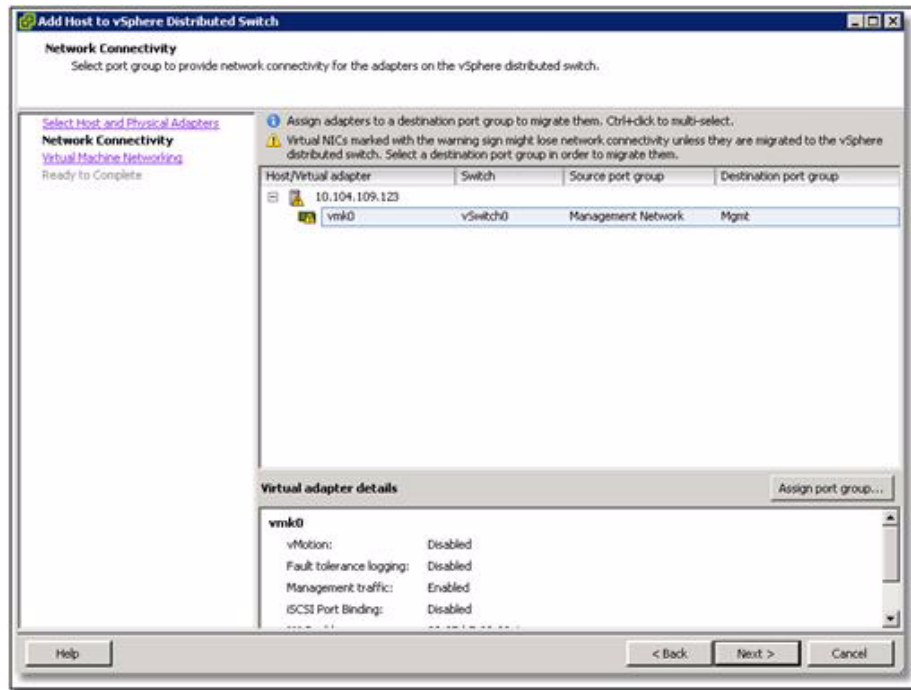
Note

A Service Console port group will be defined in the default vSwitch which is created as part of the ESXi installation.


3. Select the hosts and physical adapters to add to the VMware vSphere Distributed Switch.
4. Click **Next**.



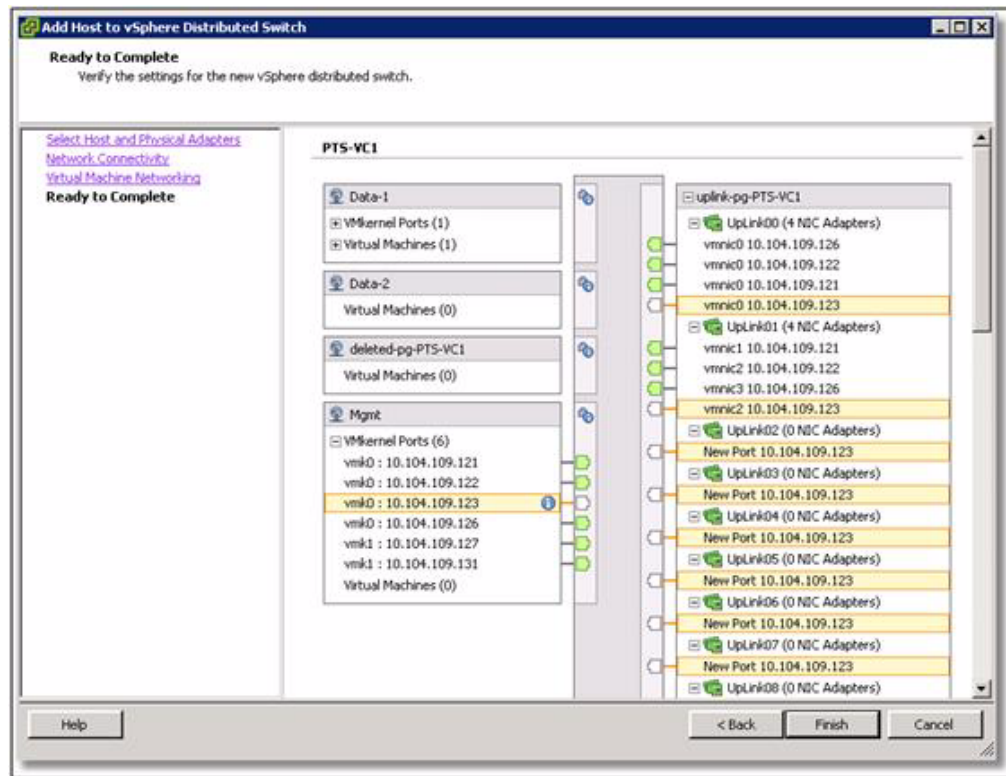
5. Restore the settings of the Network Connectivity window.
6. Click **Next**.
7. In the Network Connectivity window, select **Assign Port Group** to provide network connectivity for the adapters on the VMware vSphere Distributed Switch.



8. Select the Service Console port group on vSwitch0 and Port Profile as defined in the Cisco Unified Computing System to migrate to VM-FEX DVS switch.
9. In the Ready to Configure window, verify the settings for the new VMware vSphere Distributed Switch. Select the management network properties already defined in static VNIC.



10. Click **Finish**.

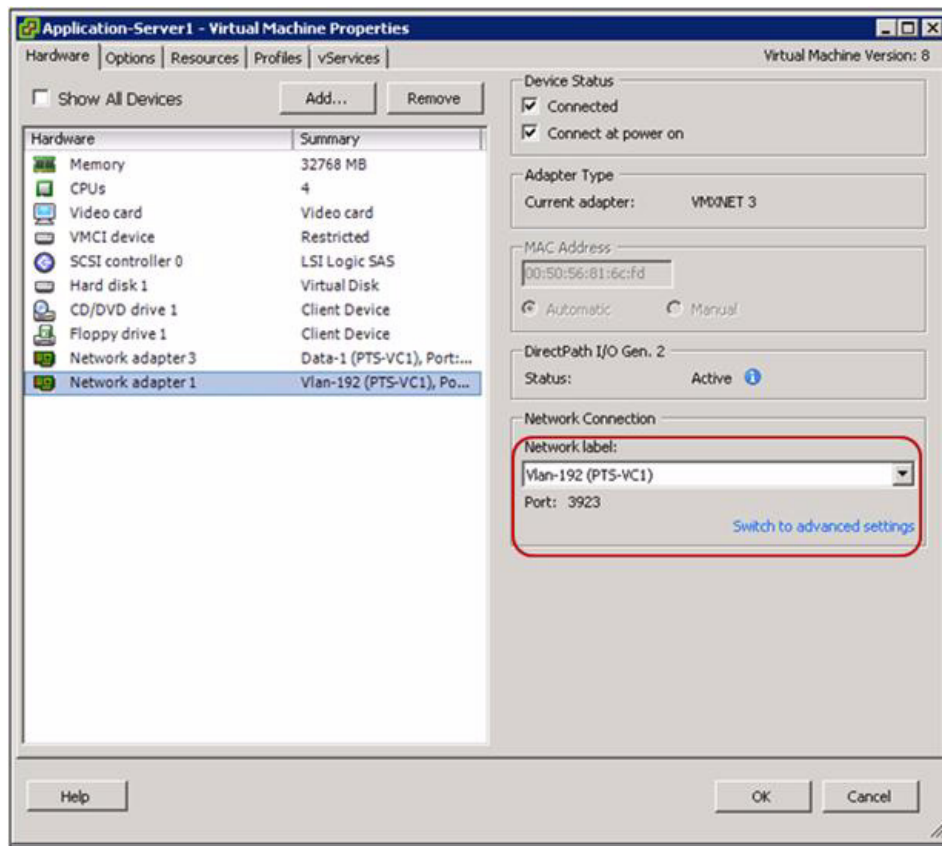


The migration of the Service Console port group on vSwitch0 to VN-FEX DVS switch is complete.

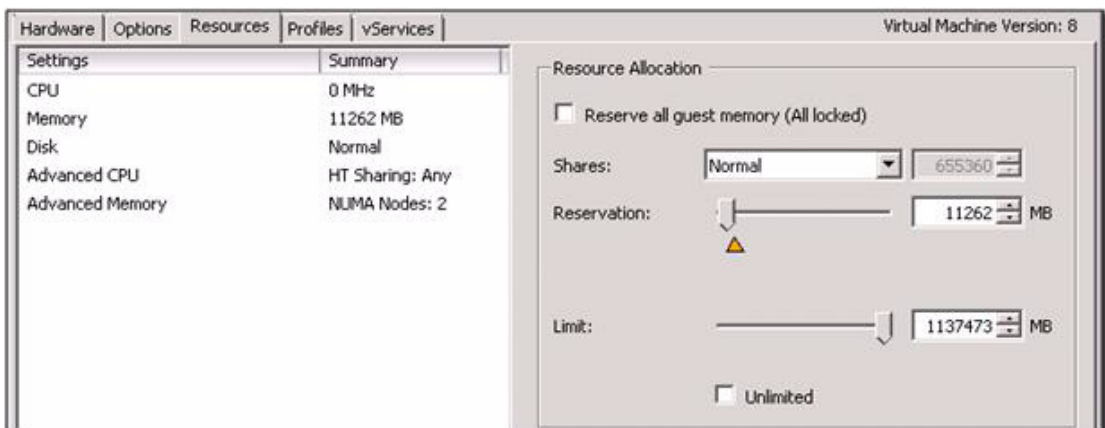
Applying the Port Profile to a Virtual Machine

After migrating the ESXi host Service Console and the corresponding Static VNIC to the VM-FEX DVS switch, perform the following steps to apply the Port Profile to the virtual machine:

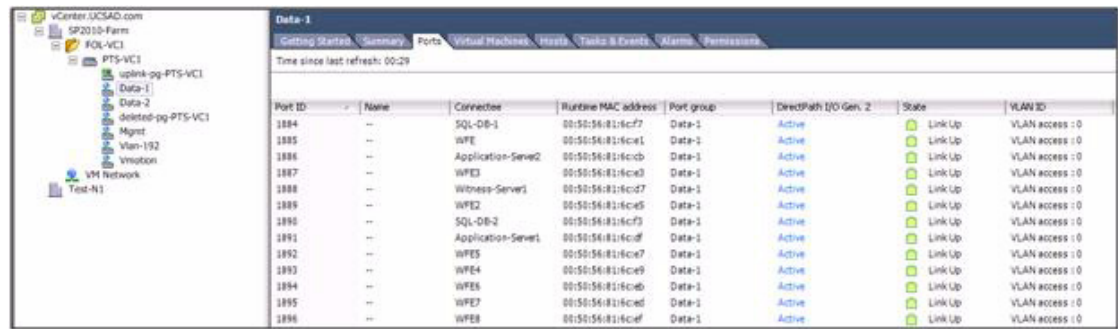
1. Login to the vCenter Manager .
2. Apply the Cisco UCS Port Profile to the virtual machine the under Network Adapter setting.



3. Allocate memory to the virtual machine in the Cisco UCS Port profile.



4. Make sure that the Directpath I/O Gen. 2 is active in the vCenter manager. After the port profile is applied Directpath I/O Gen2 status automatically goes to an active state.



Port ID	Name	Connector	Runtime MAC address	Port group	DirectPath I/O Gen. 2	State	VLAN ID
1884	--	SQL-DB-1	00:50:56:81:6cf7	Data-1	Active	Link Up	VLAN access : 0
1885	--	WFE	00:50:56:81:6ce1	Data-1	Active	Link Up	VLAN access : 0
1886	--	Application-Server2	00:50:56:81:6ceb	Data-1	Active	Link Up	VLAN access : 0
1887	--	WFE3	00:50:56:81:6cea	Data-1	Active	Link Up	VLAN access : 0
1888	--	Witness-Server1	00:50:56:81:6cd7	Data-1	Active	Link Up	VLAN access : 0
1889	--	WFE2	00:50:56:81:6ce5	Data-1	Active	Link Up	VLAN access : 0
1890	--	SQL-DB-2	00:50:56:81:6cf3	Data-1	Active	Link Up	VLAN access : 0
1891	--	Application-Server1	00:50:56:81:6cdf	Data-1	Active	Link Up	VLAN access : 0
1892	--	WFE5	00:50:56:81:6ce7	Data-1	Active	Link Up	VLAN access : 0
1893	--	WFE4	00:50:56:81:6ce9	Data-1	Active	Link Up	VLAN access : 0
1894	--	WFE6	00:50:56:81:6ceb	Data-1	Active	Link Up	VLAN access : 0
1895	--	WFE7	00:50:56:81:6ced	Data-1	Active	Link Up	VLAN access : 0
1896	--	WFE8	00:50:56:81:6cef	Data-1	Active	Link Up	VLAN access : 0

Windows Networking

In the SharePoint environment, there is a possibility of high network traffic between clients and WFEs, and between WFEs and the database, or WFEs and Application servers.

The new features to SharePoint 2010 such as Office Web Applications, digital asset storage, and playback introduce additional network traffic handling over previous versions of Office SharePoint.

The best practice is to separate the Client-WFE HTTP traffic from the WFE-database traffic to improve the network efficiency (Figure 29).

Figure 29 VLAN Inbound and Outbound Network Connections

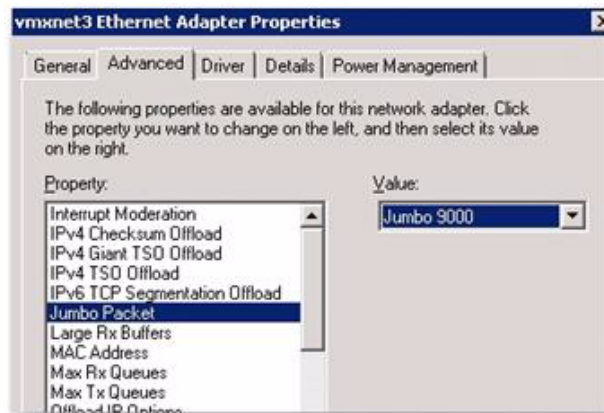


Guest Operating System Networking

Using the VMXNET3 virtualized network adapter yields the best performance for virtualized SharePoint 2010 deployments.

- Three requirements exist to take advantage of the VMXNET3:
- VMware Virtual hardware version 7
- VMware tools must be installed on the guest
- The MTU size must be set to 9000 end to end and all intermediate switches and routers must have jumbo frames enabled. This must also be set in the guest.

Figure 98 shows the VMXNET Guest Adapter Settings.

Figure 30 VMXNET Guest Adapter Settings

Validating Microsoft SharePoint 2010 Server Performance

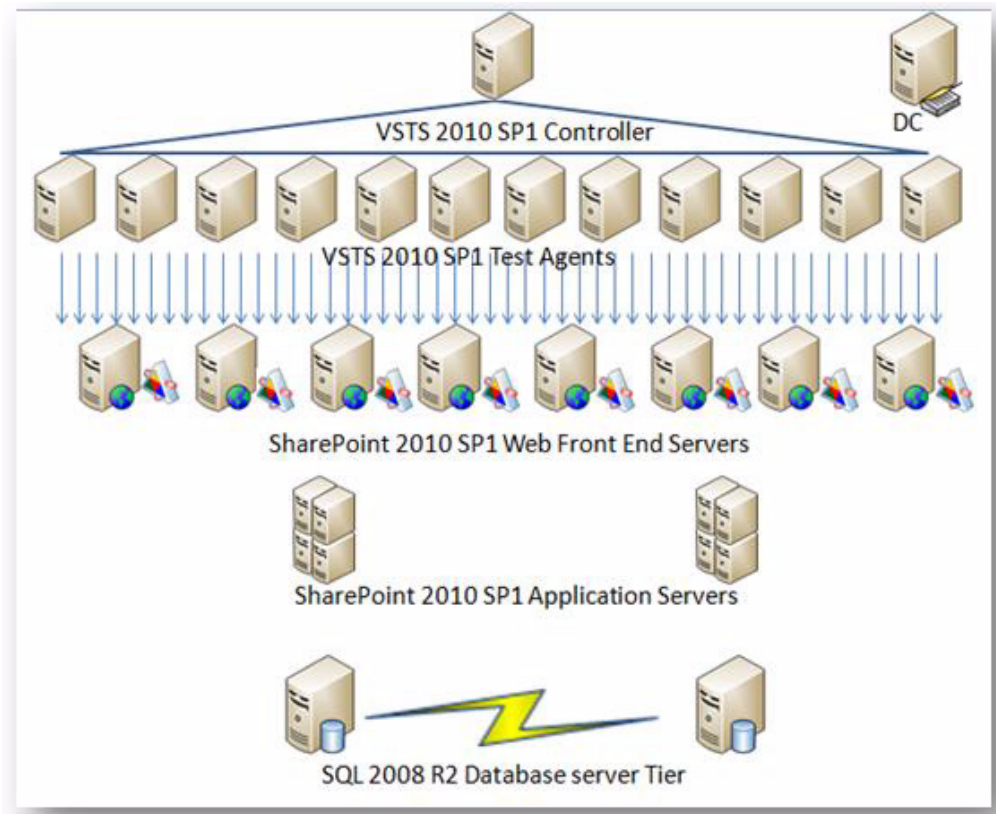
The physical architecture of the test farm consists of one VSTS 2010 SP1 controller and 12 VSTS Agents. Servers and the network infrastructures, which together creates the SharePoint Environment is tailored with respect to the size and topology as explained in the subsequent sections.

Modeling a Microsoft SharePoint Server 2010 SP1 environment begins with analyzing current requirements and estimating the expected future demand and targets for the deployment. Decisions are made on the key solution that an environment must support, and establish all the important metrics and parameters. The Microsoft SharePoint Server 2010 SP1 environment is modeled considering enterprise workload characteristics like number of users, most frequently used operations, datasets like content size, and content distribution demands, and is tailored in accordance with Microsoft recommendations.

Microsoft SharePoint Farm Under Test

Figure 99 shows the specific architecture.

Figure 31 **Microsoft SharePoint Farm Under Test Architecture**



Workload Characteristics

Sizing a SharePoint environment workload is one of the key factors of the solution. The system under test should sustain the described work load demands, user base, and usage characteristics.

Table 11 details the workload characteristics needed to size the farm under test.

Table 11 **Workload Characteristics**

Workload Characteristics	Value
Number Of Users (of unique users in a 24 hour period)	100000
Concurrent users at Peak Time (Distinct users generating requests in a given time frame)	10,000

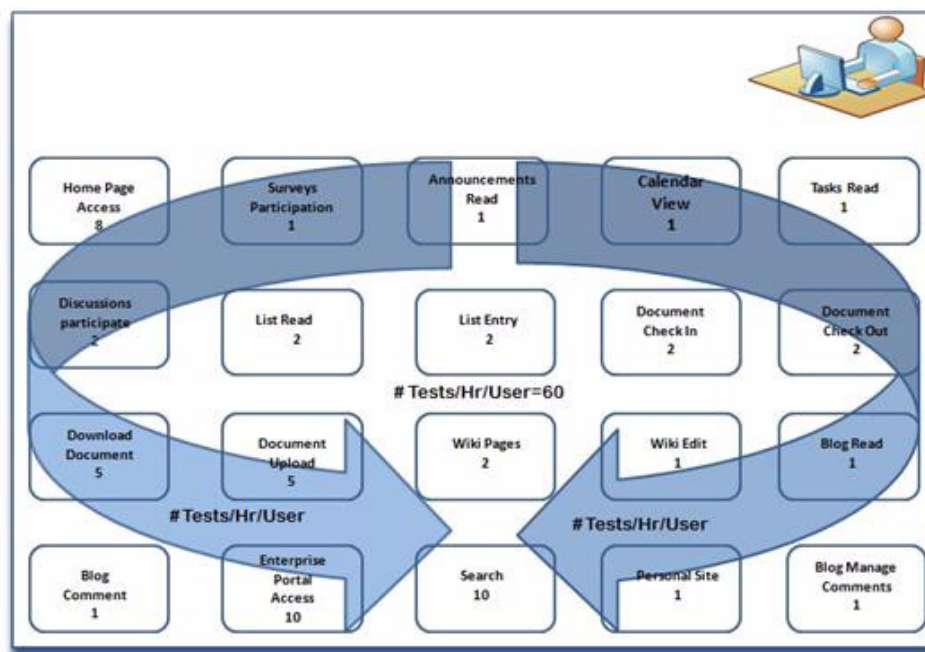
Workload Mix (60 RPH)

The requirements for the farm includes the number of users and their usage characteristics. This performance test considers a heavy profile, where a single active user requests 60 different requests per hour to the Microsoft SharePoint 2010 farm under test, for example 60 requests/ hour/ user.

User activities or usage characteristics are modeled based on an enterprise business environment needs such as organizations like marketing and engineering. In general the environment hosts central team sites, publishing portals for internal teams as well as enterprise collaboration of organizations, teams and projects. Sites created in these environments are used as communication portals host applications for business solutions and provides channel for general collaboration. Searching, editing and uploading documents, participate in discussions, blog post, comment on blogs, etc. are among the most common activities. Considering these activities to be part of a typical enterprise user, the following set of activities is included in the workload used for the performance test.

Figure 100 shows various requests, which a user should make over a period of one hour depicting a workload generated during peak hour. These workloads are loaded onto the Microsoft SharePoint farm with a brief warm up time.

Figure 32 *Requests Over a One-Hour Period*



Dataset Capacity

Defining Dataset Capacity of the Farm Under Test

A dataset holds the Microsoft SharePoint 2010 content of the defined workload.

Table 12 provides a few key metrics which were used to determine the capacity of the dataset for the test.

Table 12 Dataset Characteristics

Dataset Characteristics	Value
Content Database size each	3 TB
Number of sites	4000
Search index size (number of items)	41,972,821
Total number of documents	57,590,002
Number of content databases	20
Total number of databases	31
Number of site collections	500
Number of Web applications	20

Performance Test Framework

This performance test provides the responsiveness, throughput, reliability, and scalability of a Microsoft SharePoint 2010 SP1 farm under a given workload. The results from this performance test and analysis can help you estimate the hardware configuration required for a Microsoft SharePoint 2010 SP1 farm to support up to 100000 users with 10 percent concurrency under production operation on FlexPod.

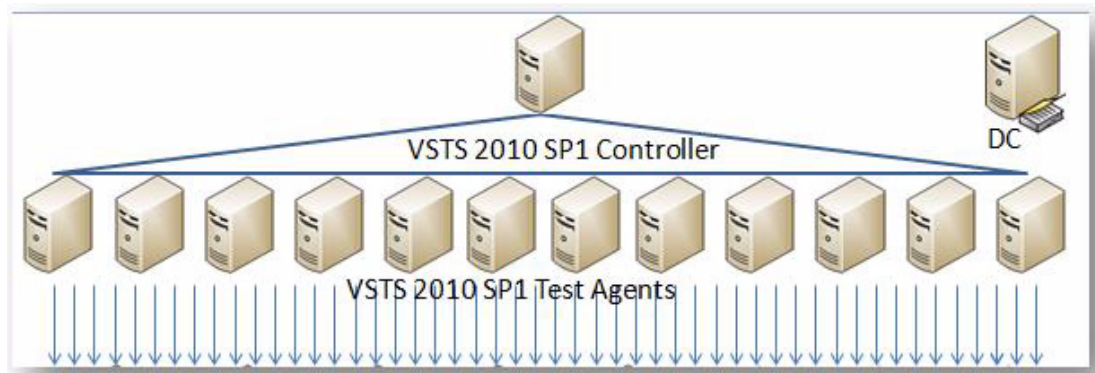
Test Methodology

For this performance test the load on the farm is applied using Microsoft LTK (Load Testing Kit). The kit is modified to make it more flexible and enhanced to generate the desired load. The Load Testing Kit (LTK) generates a Visual Studio Team System 2010 (VSTS); SP1 load test is based on the Windows SharePoint Services 3.0 Internet Information Services logs. The content database of size 3 TB is created which has sites and site collections and other important features which constitute the dataset. For more information on the created dataset, refer to the section [Defining Dataset Capacity of the Farm Under Test](#).

VSTS Test Rig

A group of servers are used to generate a simulated load for testing. A single controller machine is used to run tests remotely and simultaneously on several servers with the help of one or more agents, collectively called a as rig. The rig is employed to generate more load than a single computer can generate. A controller is used to coordinate with the agents i.e. send the load test to the agents and the agents perform the test received from the controller. The controller is also responsible to collect the test results.

Figure 101 shows the rig created for this performance test. In the test scenario rig consists of one controller and twelve agents with a Domain Controller.

Figure 33 **Controller and Test Agents****Note**

The agent takes a set of tests and a set of simulation parameters as input. A key concept in Test Edition is that tests are independent of the computer on which they are run.

Performance Tuning

Caching Of Microsoft SharePoint 2010 Farm

Microsoft SharePoint 2010 SP1 has several methods of caching data and objects to help improve the Performance for the end user. When a client requests a page, a copy of the page is stored temporarily in the output cache. Although the duration of the cache is typically small (the default is 60 seconds), this can boost the performance of WFE servers and reduce latency.

Cache profiles are available so that the site administrators can control the output cache behavior. Administrators can also determine whether different users like the content editors can receive cached pages. Administrators can also adjust the output cache at the site, site collection, or Web application level.

Microsoft SharePoint 2010 uses the object cache to temporarily store objects such as lists, libraries, or page layouts on the WFE server. Caching enables the WFE server to render pages more quickly, reducing the amount of data that is required from the SQL Server databases.

Microsoft SharePoint 2010 has a BLOB cache that temporarily stores digital assets, such as image or media files, but you can use it with any file type. Using the BLOB cache in conjunction with the Bit Rate Throttling feature in IIS 7.0 also enables the progressive download feature for digital assets. Progressive download feature enables to download media files in chunks of data, and playback feature starts after downloading the first chunk rather than the whole file. You can enable and control the size of the BLOB cache at the Web application level. The default size is 10 GB and the default setting is disabled.

In the performance test we have enabled the Microsoft SharePoint 2010 cache to improve the overall response time of the farm.

For more information about Cache Setting Operations, go to:

<http://technet.microsoft.com/en-us/library/cc261797.aspx>.

Environment Configuration Tuning

Table 13 shows the settings required at the time of configuring the environment to enhance its performance and capacity.

Table 13 Environment Configuration Tuning

Settings	Value	Notes
Site collection:		
Object Caching (On Off)	ON	Enabling the output cache improves server efficiency by reducing calls to the database for data that is frequently requested.
Anonymous Cache	Enabled	
Profile (select)		
Object Cache (Off n MB)	ON – 10GB	
Cross List Query Cache Changes (Every Time Every n seconds)	60 seconds	
Site collection cache profile (select)	Intranet (Collaboration Site)	“Allow writers to view cached content” is checked, bypassing the ordinary behavior of not letting people with edit permissions to have their pages cached.
Object Cache (Off n MB)	ON – 1GB	The default is 100 MB. Increasing this setting enables additional data to be stored in the front-end Web server memory.
Usage Service:		
Trace Log – days to store log files (default: 14 days)	5 days	The default is 14 days. Lowering this setting can save disk space on the server where the log files are stored.
Database Server – Default Instance: Max degree of parallelism	1	<p>The default is 0. To ensure optimal performance, Microsoft strongly recommend that you set max degree of parallelism to 1 for database servers that host SharePoint Server 2010 SP1 databases.</p> <p>For more information about how to set max degree of parallelism, see max degree of parallelism Option (http://go.microsoft.com/fwlink/?LinkId=189030).</p>

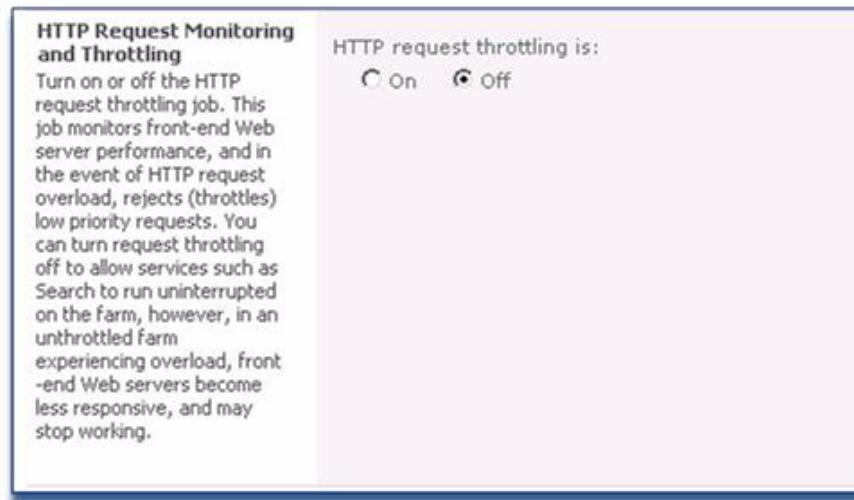
HTTP Throttling

HTTP throttling is a new feature in Microsoft SharePoint 2010 that allows the server to discard the server requests when it is too busy. Every five seconds a dedicated job runs which will check the server resources to compare with the resource levels configured for the server. By default the Server CPU,

Memory, Request in Queue and Request wait time are being monitored. The server enters a throttling period after three consecutive unsuccessful HTTP GET checks. The server remains in this period till a successful check is done. Requests generated prior to the server entering into the throttling mode will be taken and completed. Any new HTTP GET and Search Robot request will generate a 503 error message and will be logged in the event viewer. Also while the server is in a throttling period no new timer jobs will get started.

To monitor the server performance the HTTP throttling is turned off in the performance test, as shown in Figure 102.

Figure 34 *HTTP Request Throttling*



Performance Results and Analysis

The Microsoft SharePoint 2010 server performance in general varies from environment to environment, depending on the complexity of the deployment and the components involved in the architecture.

Performance of the SharePoint architecture is determined by the user experience.

The following are the major performance counters that measure user experience:

- Requests per second—Number of requests per second taken by the SharePoint Server
- Average response time—Amount of time the SharePoint Server takes to return the results of a request to the user
- Average page response time—Average time to download the page and all of its dependent requests like images, css, js etc.


The following sections detail the results as we applied the described workload (60 RPH) on the created SharePoint Farm.

Requests Per Second

Figure 35 shows the highest number of received requests per second for several user loads (60 RPH). The graph shows the smooth performance of the Virtual servers on Cisco UCS blade servers B200 M2 as the request per second scale ascending linearly with the user load without causing any or much stress

on the server. The following graph also shows the possibility to further scale up the user load with an ensured stable server performance. However the decline in request per second is seen which is due to ASP.Net and IIS limitations due to the default configuration setting at the front-end web servers.


Figure 35 Request Per Second



Average Page Time

Figure 36 shows the SharePoint 2010 average page time well below 1 sec. The SharePoint achieved sub-second response time in the performance test for a concurrent user load 10000 users, the response time varies with increased load on the Microsoft SharePoint 2010.


Figure 36 Average Page Time



Average Response Time


Figure 37 shows average response time metrics for several user loads (60 RPH) of the SharePoint farm. The designed SharePoint farm can support more than 100,000 with the 10 percent user concurrency and achieve sub-second response time. The spikes in the graph are the result of the web front end cache flush. The duration of the cache is typically small (the default is 60 seconds), but this can greatly boost the performance of WFE servers and reduce latency. The graph shows the average response time to be well below one second proving the efficiency and potential of FlexPod.

Figure 37 **Average Response Time**



Pages Per Second


Figure 38 shows the pages per second metrics for several user loads (60 RPH) of the SharePoint farm. SharePoint 2010 farm served on an average 332 pages per second.

Figure 38 *Pages Per Second*

Virtual Server - Processor Utilization

Web Front-End Server

Figure 39 shows the CPU utilization of the eight Microsoft SharePoint Web Front End servers. Under heavy workloads of 10000 users and with Network Load Balancer to balance the load at the web front end tier, Virtual servers are hosted on Cisco UCS B200 M2 blade servers. CPU utilization remained at around 40% on an average. The graph shows the linear growth in the CPU utilization for the increase in the user load and their workloads. The graph shows virtual servers on Cisco UCS B200 M2 blade server's capacity to accommodate much more workloads with ease without causing any stress.

Figure 39 *Processor Utilization*

Application Server

Figure 40 and Figure 41 show the Application Server CPU Utilization. The application server hosted the Central Administration and search services. The spikes shown on the graph are due to the search crawl at the time of the performance test. At Application Tier both the virtual servers spiked to 100 percent CPU utilization.




Figure 40 *Application Server 1 CPU Utilization*

Figure 41 **Application Server 2 CPU Utilization**

Database Server

Figure 42 shows the SQL 2008 R2 database server CPU utilization. The high availability scenario requires one server to be active and another mirrored. The CPU spikes in the graph shows the search crawl services while updating the search databases. On an average, the overall CPU utilization at the database tier with virtual servers remained at around 60 percent. This is when the server is hosted on the Cisco UCS B250 M2. FlexPod accommodates more CPU utilization without causing stress on its performance.

Figure 42 **Database Server CPU Utilization**


Network Utilization

Figure 43 shows the network utilization at Web front-end, application and database tiers of the SharePoint 2010 SP1 large farm. The graph also shows the aggregated performance numbers of the network utilization on all servers in the farm.


Note

The network is enabled with QoS policies and universal pass through switching.

Figure 43 **Network Utilization**




SharePoint 2010 Server Memory Utilization

Figure 44 shows the memory utilization of the SharePoint 2010 server farm under heavy workloads of 100000 users.

Maximum memory utilization on the Web front-end servers, Application server and SQL Server at the maximum user load is within 50 percent of the available physical memory. This indicates the CPU availability for further expansion, while providing high availability for all Microsoft SharePoint roles hosted on various Microsoft SharePoint tiers.

Figure 44 *Microsoft SharePoint 2010 Server Memory Utilization*




VMware Physical Host CPU Utilization

VMware Physical Host 1

Figure 45 shows the CPU utilization of the ESXi 5 host on Cisco UCS B200 M2 Blade Server. Where ESXi 5 physical server hosts four web front end servers where each virtual server is configured to 4 vCPU. At the maximum user load, CPU usage is within 50 percent of the available physical CPU, indicating CPU availability for further expansion.

Figure 45 *Physical CPU Core Utilization Time*



Note

The physical ESXi 5 servers mentioned above are part of VMware vSphere's DRS (Distributed Resource Scheduler) which is mainly used for load balancing virtual machines on a cluster and have configured with automated mode. Mobility of virtual servers is managed by DRS.


VMware Physical Host- 2

Figure 46 shows the CPU utilization of the ESXi 5 host on Cisco UCS B250 M2 Server. ESXi 5 physical server hosts four virtual servers i.e. SQL 2008 R2 Server 1, SQL 2008 R2 Server 2, WFE2 and Witness-Server. Each virtual server is configured to 4 vCPU.

We have configured two ESXi 5.0 physical servers considering high availability at physical server level. These are pinned to two B250 servers by configuring VMware DRS Affinity Rules that governs the placement of virtual servers on physical servers.

At the maximum user load, CPU utilization is within 60 percent of the available physical CPU, indicating CPU availability for further expansion.

Figure 46 *Physical CPU Core Utilization Time*



Note


The physical ESXi 5 servers shown in Figure 46 are part of VMware vSphere's DRS (Distributed Resource Scheduler) which is mainly used for load balancing the virtual machines on a cluster and are configured with automated mode. Mobility of virtual servers is managed by DRS.

VMware Physical Host

The following graph shows the CPU utilization of the ESXi 5 host on Cisco UCS B200 M2 Blade Server. ESXi 5 server hosted three virtual servers, for example one application server and two Web front-end servers, Where each virtual server is configured with four vCPU.

At the maximum user load, CPU utilization is within 60 percent of the available physical CPU, indicating CPU availability for further expansion.

Figure 47 *Physical CPU Core Utilization Time*




Note

The physical ESXi 5 servers mentioned above are part of VMware vSphere's DRS (Distributed Resource Scheduler) which is mainly used for load balancing the virtual machines on a cluster and have configured with automated mode. Mobility of Virtual servers is managed by DRS.

VMware Physical Host 4

Figure 48 shows the CPU utilization of the ESXi 5 host on Cisco UCS B200 M2 Server. ESXi 5 server hosted two virtual servers, for example an application server and a Web front-end server. With each virtual server configured with 4 vCPU. At the maximum user load, CPU utilization is within 60 percent of the available physical CPU, indicating CPU availability for further expansion.

Figure 48 *Physical CPU Core Utilization Time*




Note

The physical ESXi 5 servers mentioned above are part of VMware vSphere's DRS (Distributed Resource Scheduler) which is mainly used for load balancing virtual machines on a cluster and have configured with automated mode. Mobility of Virtual servers is managed by DRS.

VMware vSphere DRS Failover Time

The VMware vSphere is configured with DRS automated mode. The vSphere DRS collect resource usage information from servers and virtual machines, and then generate recommendations to optimize virtual machine allocation. These recommendations were automatically executed. Figure 49 shows the time taken for the mobility of the virtual servers between the ESXi 5.0 hosts.

Figure 49 **DRS Failover Time**

Observations done under the test scenario shows that on average the virtual servers mobility has taken 27 seconds to move virtual servers between the ESXi 5.0 hosts.


Migrate virtual machine	WFE	Completed	System	vCenter-UCSAD.com	12/1/2011 6:29:00 AM	12/1/2011 6:29:00 AM	12/1/2011 6:29:30 AM
Migrate virtual machine	WFE4	Completed	System	vCenter-UCSAD.com	12/1/2011 6:29:01 AM	12/1/2011 6:29:01 AM	12/1/2011 6:29:28 AM
Migrate virtual machine	WFE6	Completed	System	vCenter-UCSAD.com	12/1/2011 5:49:03 AM	12/1/2011 5:49:03 AM	12/1/2011 5:49:33 AM

NetApp FAS 3270—Read Write Throughput from Storage

For each of the load test run, the LUN statistics are monitored using the command `LUN stats -i -c -a -o` that captures the LUN stats every second. The following is a sample of the “`lun stats -i -c -a -o`” output that are analyzed.

Figure 50 shows the performance graph of VMware VMFS data store read/write at the time of the performance test.


Figure 50 VMware DataStore—Read/Write Performance



RDM LUNs—Read/Write Throughput

Figure 51 shows the Read/write throughput of RDM LUN. RDM LUNs are mapped to the ESXi 5.0 physical host, hosting SQL 2008 R2 database server in the performance tests.

Figure 51 Database—Read/Write Performance



Performance Results and Analysis

The test is functionally successful, meeting the criteria set to achieve a 100000 user workload with approximately 10 percent concurrency. Table 14 provides the summary of the performance results and the results of the most important realistic enterprise concerns.

Table 14 Performance Results

Important Concerns	Performance Results
Response Time (User Concern)	Less than One Second
Throughput (Business Concern)	60 Tests Per User Per Hour
Request Per Second	561 RPS
Resource Utilization (System Concern)	Under 50% CPU utilization of ESXi Hosts hosting SharePoint 2010 SP1 Large Farm.

Summary

This Cisco Validated Design introduces FlexPod which constitutes various technologies mainly, Cisco Unified Computing System, VMware vSphere 5.0 and NetApp Storage Technologies together to form a highly reliable, robust and a virtualized solution for Microsoft SharePoint 2010.

The Cisco Unified Computing System meets server virtualization challenges with a next-generation data center platform that unites and computes, network, storage access, and virtualization support in a cohesive system that is managed centrally and coordinated with virtualization software such as VMware ESXi Server. The system integrates enterprise-class servers in a 10 Gigabit Ethernet unified network fabric that provides the I/O bandwidth and functions that virtual machines and the virtualization software both need. Cisco Extended Memory Technology offers a highly economical approach for establishing the large memory footprints that high virtualization density requires. Finally, the Cisco Unified Computing System integrates the network access layer into a single easily managed entity in which, links to virtual machines can be configured, managed, and moved as readily as physical links. The Cisco Unified Computing System continues Cisco's long history of innovation and delivers innovation in architecture, technology, partnerships, and services.

VMware vSphere uses virtualization to transform datacenters into scalable, aggregated computing infrastructures. VMware vSphere manages large collections of infrastructure, such as CPUs, storage, and networking, as a seamless and dynamic operating environment, and also manages the complexity of a datacenter. The VMware vSphere software stack is composed of the virtualization, management, and interface layers. VMware vSphere provides greater levels of scalability, security, and availability to virtualized environments providing flexibility in their virtual server farms as processor-intensive workloads continue to increase.

NetApp's strategy for storage efficiency is based on the built-in foundation of storage virtualization and unified storage provided by its core Data ONTAP operating system and the WAFL file system. NetApp's technologies surrounding its FAS and V-Series product line have storage efficiency built into their core. NetApp's highly optimized, scalable operating system supports mixed NAS and SAN environments and a range of protocols, including Fibre Channel, iSCSI, FCoE, NFS, and CIFS; also, includes a patented file system and storage virtualization capabilities. The various NetApp's tools and technologies for storage like RAID-DP®, Thin Provisioning and FlexVol, Deduplication, NetApp Snapshot™ and SnapDrive® for Windows® assist in deployment, backup, recovery, replication, management, data protection and helps in provisioning storage resources.

Microsoft SharePoint 2010 an extensible and scalable web-based platform consists of tools and technologies that support the collaboration and sharing of information within teams, throughout the enterprise and on the web. SharePoint 2010 finds its implementation to fulfill various enterprise demands. Implementation of SharePoint requires an in depth study of requirements. Demand Characteristics such as expected user load, workloads, concurrent users at peak time, request per second and content size define the size of a SharePoint implementation. Three-tier architecture provisions an ideal SharePoint topology. Several servers at individual tier render various SharePoint components; together to make up a SharePoint 2010 Large farm. Servers at web tier renders web and search query functions, servers on the application tier are responsible for search indexing and various service application functions and server at the database tier hosts SQL Server databases for the farm.

This performance study Cisco Validated Design is intended to understand the performance capacity of Large SharePoint farm on FlexPod. Where, Cisco UCS blade servers providing improved application performance and operational efficiency, NetApp providing its tools and technologies to improve storage efficiency and VMware with its innovative stack of virtualization.

The performance study showed that SharePoint Large farm could easily support 100000 users and with a minimum 10 percent concurrency. The Microsoft SharePoint farm comprising of Cisco UCS blade servers, NetApp storage, VMware virtualization with a Network access of 10 Gbps connectivity between the tiers provided an average response time well below one second. FlexPod with its various innovative technological performance benefits essentially yielded the efficient performance results.

Bill of Material

Table 15 provides the details of the components used in this Cisco Validated Design.

Table 15 Component Description

Description	Part #
Cisco UCS 5108 Blade Server Chassis	N20-C6508
Cisco UCS 2104XP Fabric Extender/4 external 10Gb ports	N20-I16584
Cisco UCS B200 M2 Blade Server; dual Intel Xeon X5680 CPUs (3.33 GHz & 6 Cores), 96GB RAM (DDR3 1067 MHz)	N20-B6620-1
Cisco UCS B250 M2 Blade Server, dual Intel Xeon X5680 CPUs (3.33 GHz & 6 Cores), 384 GB RAM (DDR3 1067 MHz)	N20-B6625-2
Cisco UCS 6120XP 20-port Fabric Interconnect	N10-S6100
4-port 10 GE/4-port 4Gb FC/Expansion Module/UCS 6100 Series	N10-E0440
FAS3270 single enclosure HA (single 3U chassis)	FAS3270A
Dual-port 10-GbE Unified Target Adapter with fiber	X1139A-R6
Disk Shelf with 24 450-GB SAS Drives, 15k RPM, 4 PSU, 2 IOM3 modules	DS4243-1511-24S-QS-R54
NFS Software License	SW-T7C_NFS-C
Nexus 5548up	N5K-C5548UP-FA
Nexus 5548up Storage Protocols Services License	N5548P-SSK9
10GBASE-SR SFP Module	SFP-10G-SR
8Gbps Fibre Channel SW SFP+, LC	DS-SFP-FC8G-LW

Table 16 Software Details

Platform	Software Type
Cisco UCS 6120	Management
Cisco UCS 6120	OS
Nexus 5548up	OS
NetApp 3270	OS
Blade servers	OS
Database	SQL
Application	SharePoint 2010 Enterprise
Application	Management
Application	Management

References

Microsoft SharePoint 2010:

<http://technet.microsoft.com/en-us/sharepoint/ee263917>

Cisco UCS:

<http://www.cisco.com/en/US/netsol/ns944/index.html>

VMware vSphere:

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

NetApp Storage Systems:

<http://www.netapp.com/us/products/storage-systems/>

Cisco Nexus:

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

Cisco Validated Design-FlexPod for VMware:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html

Cisco Nexus 5000 Series NX-OS Software Configuration Guide:

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

NetApp TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection:

<http://www.netapp.com/us/library/technical-reports/tr-3298.html>

Microsoft Visual Studio Ultimate 2010:

<http://www.microsoft.com/visualstudio/en-us/products/2010-editions/ultimate>

Microsoft TechNet Articles

- Capacity management and sizing overview for SharePoint Server 2010

<http://technet.microsoft.com/en-us/library/ff758647.aspx>

- Cache settings for a Web application (SharePoint Server 2010)

<http://msdn.microsoft.com/en-us/library/ms189910.aspx>

- Microsoft SharePoint 2010

<http://technet.microsoft.com/en-us/sharepoint/ee263917>

Acknowledgements

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

- Mike Mankovsky, Cisco Systems, Inc.
- Nick DeRose, NetApp
- Christopher Reno, NetApp
- Wande He, VMware
- Frank Cicalese, Cisco Systems, Inc.

About Cisco Validated Design (CVD) Program

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

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

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

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

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

© 2010 Cisco Systems, Inc. All rights reserved

