

Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere

Cisco Unified Computing System and Hitachi Virtual Storage Platform
with VMware vSphere 5.0


Last Updated: February 5, 2014



Cisco
Validated
Design



Building Architectures to Solve Business Problems



About the Authors

Chris O'Brien, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems, Inc.

Chris O'Brien is currently focused on developing infrastructure best practices and solutions that are designed, tested, and documented to facilitate and improve customer deployments. Previously, O'Brien was an application developer and has worked in the IT industry for more than 15 years.

Hardik Patel, Virtualization System Engineer, Server Access Virtualization Business Unit, Cisco Systems, Inc.

Hardik is focused in design and implementation of systems and virtualization, manage and administration, Cisco Unified Computing System, storage and network configurations in data center solutions with Cisco Business Unit and partners. Hardik has a Master's degree in Computer Science with various career oriented certification in virtualization including VMware Certified Professional, network and Microsoft.

Tushar Patel, Principle Engineer, Data Center Group, Cisco Systems, Inc.

Tushar Patel is a Principle Engineer for the Cisco Systems Data Center group. Tushar has nearly fifteen years of experience in database architecture, design and performance. Tushar also has strong background in Intel x86 architecture, storage and virtualization. He has helped a large number of enterprise customers evaluate and deploy various database solutions. Tushar holds a Master's degree in Computer Science and has presented to both internal and external audiences at various conferences and customer events.

Jay Chellappan, Solutions Consultant, ISV Solutions and Engineering, Hitachi Data Systems

Jay Chellappan has over seven years of experience architecting solutions using the entire line of Hitachi storage products in Proof of Concepts, Solutions and Reference Architectures. Jay started his career with WIPRO Infotech as a Systems Engineer working in pre- and post-sales. After this, Jay co-founded WISE, a technology services firm in India eventually selling off the company. Following this Jay worked at several large corporations in the San Francisco Bay Area as a Systems, Storage and Oracle Database Administrator for several years. Jay has a Bachelor's degree in Physics and a Master's in Computer Applications, both from the M.S. University of Baroda, India.

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)

© 2012 Cisco Systems, Inc. All rights reserved



Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere

Introduction

The Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere offers a blueprint for creating a cloud-ready infrastructure integrating computing, networking, and storage resources. As a validated reference architecture, the Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere is a safe, versatile, and cost-effective alternative to static-integrated approaches. The Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere provides a uniform platform allowing organizations to standardize their data center environments.

The Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere speeds application and infrastructure deployment and reduces the risk of introducing new technologies to the data center. Administrators select hardware options that meet their needs for today, with the confidence that the architecture can adapt as business needs change.

This document outlines the Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System, and discusses design choices and deployment best practices.

Audience

The intended audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

Purpose

A Cisco Validated Design consists of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of our customers.



The purpose of this document is to describe the Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere by Hitachi Unified Compute Platform (UCP) Select for VMware vSphere with Cisco Unified Computing System which is a validated approach for deploying Cisco and Hitachi technologies as an integrated compute stack.

Problem Statement

Ideally modern IT departments can rapidly stand up virtualized compute, storage, and network in a standardized manner with ease of management, little risk and a roadmap to support cloud deployments in addition to other IT initiatives. Unfortunately, today's data centers are typically comprised of legacy heterogeneous compute, network and storage stacks, which prevent the full realization of these goals. In order to address these issues Cisco and Hitachi Data Systems have collaborated to deliver a pre-validated Cisco Solution for Hitachi Unified Compute Platform Select for VMware vSphere.

Hitachi Unified Compute Platform Select for VMware® vSphere with Cisco® Unified Computing System

The architectural advantage of Hitachi Unified Compute Platform Select for VMware vSphere with Cisco UCS is delivered through a pre-validated infrastructure that integrates the Cisco Unified Computing System (UCS), Hitachi Virtual Storage Platform (VSP), VMware vSphere, and VMware vCenter Server using a Cisco-enabled data center fabric. Each of these platforms deliver advantages from a technology perspective but the overall benefits of the UCP Select for VMware vSphere with Cisco UCS solution are greater than the sum of its individual parts.

The following advantages are delivered within UCP Select for VMware vSphere with Cisco Unified Computing System:

- Fabric Infrastructure Resilience

UCP Select for VMware vSphere with Cisco Unified Computing System is a highly available and scalable infrastructure that IT can evolve over time to support multiple physical and virtual application workloads. UCP Select for VMware vSphere with Cisco UCS contains no single point of failure at any level, from the server through the network to the storage. Both the LAN and SAN fabrics are fully redundant and provide seamless traffic failover should any individual component fail.

- Fabric Convergence

The Cisco Unified Fabric is a data center network that supports both traditional LAN traffic and all types of storage traffic, including the lossless requirements for block-level storage transport through Fibre Channel. The Cisco Unified Fabric creates high-performance, low-latency, and highly available networks serving a diverse set of data center needs. Cisco NX-OS is the software operating the Cisco Unified Fabric. Cisco NX-OS enabled devices are managed from a single dashboard through Cisco Prime Data Center Network Manager (DCNM) allowing network and storage administrators to manage the Unified Fabric, both Ethernet and non-Ethernet portions. Hitachi Unified Compute Platform Select for VMware vSphere with Cisco® Unified Computing System uses the Cisco Unified Fabric, offering wire-once application deployment acceleration as well as efficiencies and cost savings associated with infrastructure consolidation.

- Network Virtualization

UCP Select for VMware vSphere with Cisco Unified Computing System delivers the capability to securely separate and connect virtual machines into the network. Using technologies such as VLANs, QoS and port profiles, UCP Select for VMware vSphere with Cisco Unified Computing System allows network policies and services to be uniformly applied within the integrated compute stack. This capability enables the full utilization of the UCP Select for VMware vSphere with Cisco Unified Computing System while maintaining consistent application and security policy enforcement across the stack even with workload mobility.

UCP Select for VMware vSphere with Cisco Unified Computing System provides a uniform approach to IT architecture, offering a well characterized and documented shared pool of resources for application workloads providing operational efficiency and consistency while being versatile enough to meet a variety of SLAs.

As a result, UCP Select for VMware vSphere with Cisco Unified Computing System can support a number of IT initiatives including:

- New applications or application migrations
- Business Continuity/Disaster Recovery
- Desktop Virtualization
- Cloud delivery models (public, private, hybrid) and service models (IaaS, PaaS, SaaS)
- Asset Consolidation and Virtualization

UCP Select for VMware vSphere with Cisco Unified Computing System

System Overview

UCP Select for VMware vSphere with Cisco Unified Computing System from Cisco and Hitachi Data Systems is a pre-validated integrate compute stack offering organizations the opportunity to deploy and scale compute, network and storage resources efficiently into the data center.

Hitachi Unified Compute Platform Select for VMware vSphere with Cisco® Unified Computing System is a platform designed to support multiple application workloads providing a resilient reliable infrastructure.

[Table 1](#) a high-level view of the components used to develop UCP Select for VMware vSphere with Cisco Unified Computing System. Each of these products brings unique capabilities to the data center; however, when properly integrated these product make an even more compelling and powerful solution.

Table 1 *Products Used in UCP Select for VMware vSphere with Cisco UCS from Cisco and Hitachi*

Role	Product
Compute	Cisco UCS with B200 M3 and B230 M3 blade servers with Cisco UCS Virtual Interface Card 1240
Fabric Interconnect	Cisco UCS 6200UP Series Fabric Interconnect
Unified Access Layer Switches	Cisco Nexus 5500UP Switch
Storage	Hitachi Virtual Storage Platform (VSP) with Hitachi Dynamic Provisioning
Virtualization Software	VMware ESXi vSphere 5.0 VMware vCenter 5.0
Distributed Virtual Switch	Cisco Nexus 1000v Series Switch Or VM-FEX technology
Management	Cisco UCS Manager 2.0 Cisco Prime Data Center Network Manager 6.1 Hitachi Command Suite v7.3 Hitachi Storage Navigator (VSP) Hitachi Storage Manager vCenter Plugin v2.3 Hitachi Storage Provider for VMware vCenter (VASA)
Storage Switch (<i>optional</i>)	Cisco MDS 9500 Series Multilayer Directors with Gen4 Modules

UCP Select for VMware vSphere with Cisco Unified Computing System does not mandate dedicated SAN switching. Instead, it has been validated using both a Unified Access Layer Design, as well as, an alternate Dedicated SAN Design (Alternate Design). Either UCP Select for VMware vSphere with Cisco UCS model supports the Cisco Unified Fabric using NX-OS software on the Nexus and MDS product lines.

Design Principles

Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System addresses four primary design principles scalability, elasticity, availability and manageability. These architecture goals as follows:

- Availability makes application services accessible and ready for use with no single point of failure
- Scalability will match increasing demands with increasing resources
- Elasticity provides new services or recovers used resources without requiring infrastructure modification
- Manageability enables administrator's to easily operate the infrastructure



Note

Performance and Security are key design criteria that were not directly addressed in this project but have been addressed by each vendor in other collateral, benchmarking and solution testing efforts.

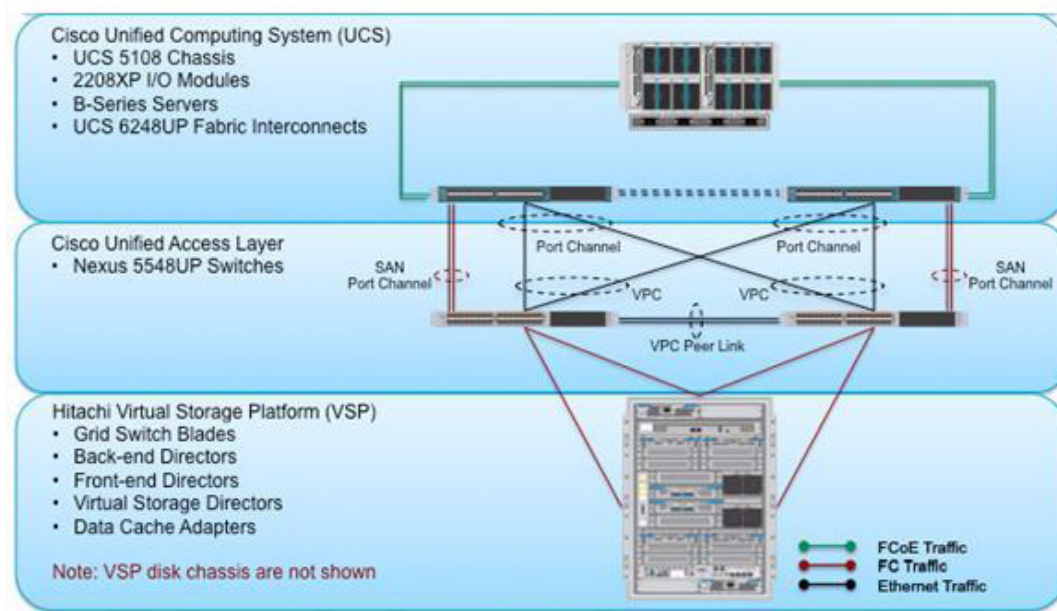
The design pillars are targeted at providing a robust, reliable, manageable and predictable architecture. The results of designing with these goals in mind are the two UCP Select for VMware vSphere with Cisco Unified Computing System models developed and documented in this paper.

Unified Access Layer Design

Figure 1 details the Unified Access Layer Design for this Reference Architecture. As the illustration shows the design is fully redundant in the compute and network layers, there exists no single point of failure from a device or traffic path perspective. At the storage layer, the Hitachi Virtual Storage Platform

(VSP) frame contains a single Control and Disk Unit (Disk chassis not shown) with redundant data and control paths, as well as, redundant hot-swappable components to provide a resilient and highly available storage solution in a single frame, see [Appendix A—Bill of Materials](#) for more details.

Figure 1 *UCP Select for VMware vSphere with Cisco Unified Computing System—Unified Access Layer Design*



The Unified Access Layer design does not employ a dedicated SAN switching environment. The Nexus 5500 series switches support both Ethernet and Fibre Channel traffic. The Hitachi VSP and Cisco UCS 6200 series Fabric Interconnects are directly connected to the Nexus 5548UP unified ports that can be configured to support either Ethernet (including FCoE) or native Fibre Channel traffic. The figure illustrates both Fibre Channel, and Ethernet frames are traversing the Nexus access layer in this design.

The Cisco UCS Fabric Interconnects are deployed in Ethernet end-host mode meaning the Fabric Interconnects are not learning MAC addresses outside of their domain and will only retain the MAC addresses of servers or virtual machines within the Cisco UCS system. End-host mode allows for optimal traffic flow of frames within the Cisco UCS domain as “interior” MACs are learned and the Fabric Interconnects will forward accordingly. End host mode does not require a protocol to contend with physical path redundancy such as Spanning Tree so all ports are unblocked.

For load balancing and fault tolerance, the Cisco Unified Computing System and Nexus 5500 platforms support port channeling of Ethernet and Fibre Channel traffic. The Nexus 5500UP platform features virtual PortChannel (vPC) capabilities. vPC allows links that are physically connected to two different Cisco Nexus 5500 Series devices to appear as a single PortChannel to a third device, in this case the Cisco UCS Fabric Interconnects. In addition to the link and device resiliency provided through vPC, the environment is non-blocking allowing full utilization of the Ethernet channel links.

From a Fibre Channel perspective the Cisco UCS Fabric Interconnects and the Nexus 5500 platforms support Fibre Channel port channeling or link aggregation allowing one logical Fibre Channel link to provide fault-tolerance. Each UCS Fabric Interconnect pair or domain supports a maximum of four Fibre Channel port channels. Each Fibre Channel port channel can support a maximum of sixteen FC uplinks. The Cisco UCS Fabric Interconnects operate in the N-Port Virtualization (NPV) mode meaning the servers are either manually or automatically pinned to specific FC uplinks, in this case the port channel.

The Fibre Channel links from the Nexus 5500 platforms are dual-homed into either Cluster on the Hitachi VSP controller unit to provide link resiliency and traffic load balancing. VMware native multi-pathing using the round-robin hashing algorithm balances the I/O across the connections. Fibre Channel zoning is performed on the Cisco Nexus 5500 access layer and masking on the Hitachi VSP for control and added security within the storage domain.

The [Logical Build](#) section provides more details regarding the design and implementation of the virtual environment consisting of VMware vSphere, Cisco Nexus 1000v virtual distributed switching, Cisco VM-FEX and Hitachi VSP.

Components

The following components are required to deploy the Unified Access Layer design:

- Cisco Unified Compute System
- Cisco Nexus 5500 Series Switch
- Hitachi Virtual Storage Platform
- VMware vSphere

Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation blade and rack server computing. Cisco Unified Computing System is an innovative data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Managed as a single system whether it has one server or 160 servers with thousands of virtual machines, Cisco Unified Computing System decouples scale from complexity. Cisco Unified Computing System accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non-virtualized systems.

The Cisco Unified Computing System consists of the following components:

- Cisco UCS 6200 Series Fabric Interconnects
(<http://www.cisco.com/en/US/products/ps11544/index.html>) is a family of line-rate, low-latency, lossless, 10-Gbps Ethernet and Fibre Channel over Ethernet interconnect switches providing the management and communication backbone for the Unified Computing System. Cisco UCS supports VM-FEX technology, see Cisco VM-FEX section for details.
- Cisco UCS 5100 Series Blade Server Chassis
(<http://www.cisco.com/en/US/products/ps10279/index.html>) supports up to eight blade servers and up to two fabric extenders in a six rack unit (RU) enclosure.
- Cisco UCS B-Series Blade Servers
(<http://www.cisco.com/en/US/partner/products/ps10280/index.html>) increase performance, efficiency, versatility and productivity with these Intel based blade servers.
- Cisco UCS Adapters
(http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html) wire-once architecture offers a range of options to converge the fabric, optimize virtualization and simplify management. Cisco adapters support VM-FEX technology, see Cisco VM-FEX section for details.
- Cisco UCS Manager

(<http://www.cisco.com/en/US/products/ps10281/index.html>) provides unified, embedded management of all software and hardware components in the Cisco UCS.

For more information, see: <http://www.cisco.com/en/US/products/ps10265/index.html>

Cisco Nexus 5500 Series Switch

The Cisco Nexus 5000 Series is designed for data center environments with cut-through technology that enables consistent low-latency Ethernet solutions, with front-to-back or back-to-front cooling, and with data ports in the rear, bringing switching into close proximity with servers and making cable runs short and simple. The switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data center-class Cisco® NX-OS Software for high reliability and ease of management.

The Cisco Nexus 5500 platform extends the industry-leading versatility of the Cisco Nexus 5000 Series purpose-built 10 Gigabit Ethernet data center-class switches and provides innovative advances toward higher density, lower latency, and multilayer services. The Cisco Nexus 5500 platform is well suited for enterprise-class data center server access-layer deployments across a diverse set of physical, virtual, storage-access, and high-performance computing (HPC) data center environments.

The Cisco Nexus 5548UP is a 1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports and one expansion slot supporting modules with 10 Gigabit Ethernet and FCoE ports or to connect to Fibre Channel SANs with 8/4/2/1-Gbps Fibre Channel switch ports, or both.

For more information, see: <http://www.cisco.com/en/US/products/ps9670/index.html>

Cisco Nexus 1000v

Cisco Nexus® 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking. The switches are designed to accelerate server virtualization and multitenant cloud deployments in a secure and operationally transparent manner. Integrated into the VMware vSphere hypervisor and fully compatible with VMware vCloud Director, the Cisco® Nexus 1000V Series provides:

- Advanced virtual machine networking based on Cisco NX-OS operating system and IEEE 802.1Q switching technology
- Cisco vPath technology for efficient and optimized integration of virtual network services
- Virtual Extensible Local Area Network (VXLAN), supporting cloud networking

These capabilities help ensure that the virtual machine is a basic building block of the data center, with full switching capabilities and a variety of Layer 4 through 7 services in both dedicated and multitenant cloud environments. With the introduction of VXLAN on the Nexus 1000V Series, network isolation among virtual machines can scale beyond traditional VLANs for cloud-scale networking.

The Cisco Nexus 1000V Series Switches are virtual machine access switches for the VMware vSphere environments running the Cisco NX-OS operating system. Operating inside the VMware ESX or ESXi hypervisors, the Cisco Nexus 1000V Series provides:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Nondisruptive operational model for your server virtualization and networking teams
- Virtualized network services with Cisco vPath providing a single architecture for L4 -L7 network services such as load balancing , firewalling and WAN acceleration

For more information, see: <http://www.cisco.com/en/US/products/ps9902/index.html>

Cisco VM-FEX

Cisco VM-FEX technology collapses virtual switching infrastructure and physical switching infrastructure into a single, easy-to-manage environment. Benefits include:

- 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)

VM-FEX is supported on Red Hat Kernel-based Virtual Machine (KVM) and VMware ESX hypervisors. Live migration and vMotion are also supported with VM-FEX.

VM-FEX eliminates the virtual switch within the hypervisor by providing individual Virtual Machines (VMs) virtual ports on the physical network switch. VM I/O is sent directly to the upstream physical network switch that takes full responsibility for VM switching and policy enforcement. This leads to consistent treatment for all network traffic, virtual or physical. VM-FEX collapses virtual and physical switching layers into one and reduces the number of network management points by an order of magnitude.

The VIC uses VMware's Direct Path I/O technology to significantly improve throughput and latency of VM I/O. Direct Path allows direct assignment of PCIe devices to VMs. VM I/O bypasses the hypervisor layer and is placed directly on the PCIe device associated with the VM. VM-FEX unifies the virtual and physical networking infrastructure by allowing a switch ASIC to perform switching in hardware not on a software based virtual switch. VM-FEX is offloading the ESXi hypervisor that may improve the performance of any hosted VM applications.

Hitachi Virtual Storage Platform (VSP)

Hitachi Virtual Storage Platform is a highly scalable, true enterprise-class storage system that can virtualize external storage and provide virtual partitioning and quality of service for diverse workload consolidation. The abilities to securely partition port, cache and disk resources, and to mask the complexity of a multivendor storage infrastructure, make Virtual Storage Platform the ideal complement to VMware environments.

As your VMware environment grows to encompass mission-critical and Tier 1 business applications, VSP delivers the highest uptime and flexibility for your block-level storage needs, providing much-needed flexibility for virtual environments.

With the addition of flash acceleration technology a single VSP is now able to service more than 1 million random read IOPS. This extreme scale allows you to increase systems consolidation and virtual machine density by up to 100%, defer capital expenses and operating costs, and improve quality of service for open systems and virtualized applications.

VSP capabilities ensure that businesses can meet service level agreements (SLAs) and stay on budget.

- The industry's only 100% uptime warranty
- 3D scalable design
- 40 percent higher density
- 40 percent less power required
- Industry's leading virtualization

- Nondisruptive data migration
- Fewer administrative resources
- Resilient performance, less risk


3D Scaling of Storage Resources

VSP is the only enterprise storage platform in the industry that offers 3D scaling of resources (see [Figure 2](#)), full native VMware integration and storage virtualization that extends features to more than 100 heterogeneous storage systems, allowing you to leverage existing investments.

3D scaling supports growing VMware environments:


- Scale up for increased performance, adding just the processing and cache you need.
- Scale out for increased capacity, with up to 2048 drives internally.
- Scale deep to leverage the industry's leading external virtualization technology and manage your external storage through common interface and management tools

Figure 2 *Hitachi Virtual Storage Platform*



VSP offers a flexible choice of 2.5-inch and 3.5-inch form factors, drives and energy efficiencies of up to 40 percent savings compared to other storage systems. With a much smaller data center footprint, VSP excels in performance and cost savings while addressing a multitude of enterprise business requirements. VSP also extends the life of storage assets and increases ROA by enabling the connection of legacy storage systems to VSP and allowing them to inherit vStorage API for Array Integration (VAAI) capability.

Figure 3 *Hitachi Virtual Storage Platform Supports 3D Scaling*



One of the unique features of VSP is its ability to use storage virtualization as a use case for VMware environments, which allows administrators to:

- Nondisruptively move VMware Virtual Machine File System (VMFS) data stores within a single frame.
- Nondisruptively migrate to new storage.
- Extend the useful life of existing storage assets that are used for VMware, enabling VAAI.
- Enable data classification for tiered storage management.

VSP capabilities ensure that businesses can meet service level agreements (SLAs) and stay on budget.

- The industry's only 100% uptime warranty
- 3D scalable design
- 40 percent higher density
- 40 percent less power required
- Industry's leading virtualization
- Nondisruptive data migration
- Fewer administrative resources
- Resilient performance, less risk

The flexible, scalable, secure VSP design provides:

- Point-to-point SAS-2 6Gb/sec scalable and reliable architecture
- Integrated data at rest encryption
- Virtualization of up to 255PB of external storage
- Up to 1TB of shared global cache
- Up to 192 Fibre Channel 8Gb/sec host ports
- Up to 96 Fibre Channel over Ethernet (FCoE) 10Gb/sec host ports
- Built-in hardware-based data-at-rest encryption

With full VAAI integration, VSP supports:

- Increased I/O performance and scalability, gained by offloading block-locking mechanism
- Reduction or elimination of SCSI reservation conflicts
- Increased virtual machine density per volume
- Full copy offload to the storage system; block zeroing offload to the storage system

VSP further enables wide striping and thin provisioning with Hitachi Dynamic Provisioning (HDP) software:

- Wide-striping of data across all drives in array groups assigned to the HDP pool to increase performance by striping across spindles in multiple array groups
- Thin provisioning and management of Dynamic Provisioning pools in Hitachi Command Suite management tools

Hitachi Dynamic Provisioning


Hitachi Dynamic Provisioning (HDP) is a microcode function of Virtual Storage Platform that provides thin provisioning and wide striping services for simplified provisioning operations, automatic performance optimization and storage space savings. Provisioning storage from a virtual pool

dramatically simplifies capacity planning and reduces administration costs by cutting the time to provision new storage. It also improves application availability by reducing the downtime needed for storage provisioning.

Capacity is allocated to an application without it being physically mapped until it is used. In this manner it is possible to achieve overall higher rates of storage utilization with just-in-time provisioning. Dynamic provisioning separates an application's storage provisioning from the addition of physical storage capacity to the storage system. When physical storage is installed within the VSP or a virtualized external array it is mapped to one or more Dynamic Provisioning Pools. Physical storage can be non-disruptively added to or removed from the pool as needed and existing allocated capacity is automatically restriped. Additional physical storage is dynamically added to thin provisioned Host LUNs (DP-VOLS) whenever an application requires additional capacity. Physical storage that is no longer required can also be dynamically returned to the pool using Zero Page Reclaim and vSphere's unmap functionality.

Dynamic Provisioning also simplifies Storage Virtualization as it provides a layer of abstraction which masks the difference between internal and external storage resources. There is no difference in the way that Host LUNs are carved out of pools for internal and external storage. The only difference is the way that physical storage capacity is replenished. This is illustrated in [Figure 4](#).

Figure 4 *Hitachi Dynamic Provisioning Example*



Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. This distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots. Hitachi Dynamic Provisioning in a VMware vSphere environment allows a leveling of the workload across a large number of disks. This allows sizing the disks for an average workload instead of a peak burst workload.

The advantages of using Hitachi Dynamic Provisioning in a VMware environment are:


- Performance Balance-Hitachi Dynamic Provisioning distributes data across all disks in the pool.
- Over Provisioning-Hitachi Dynamic Provisioning allows over provisioning of the storage. ?

For more information, see the Hitachi Dynamic Provisioning datasheet and Hitachi Dynamic Provisioning on the Hitachi Data Systems website

Hitachi Dynamic Tiering

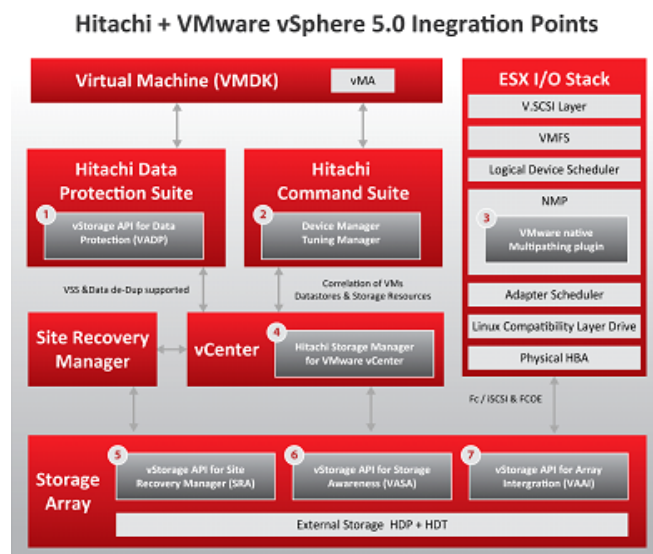
Hitachi Dynamic Tiering is an enhancement to Dynamic Provisioning that increases performance and lowers operating cost using automated data placement. Data is automatically moved according to simple rules. Two or three tiers of storage can be defined using any of the storage media types available for Hitachi Virtual Storage Platform. Tier creation is automatic based on media type and speed. Using ongoing embedded performance monitoring and periodic analysis, the data is moved at a fine grain, sub-LUN level to the most appropriate tier. The most active data moves to the highest tier. During the process the system automatically maximizes the use of storage, keeping the higher tiers fully utilized.

Figure 5 *Hitachi Dynamic Tiering*



Hitachi Virtual Storage Platform is the only 3D scaling storage platform designed for all data types. It is the only storage architecture that flexibly adapts for performance and capacity, and it extends to multivendor storage. With the unique management capabilities of Hitachi Command Suite software, it transforms the data center. You can start small and scale up, out and deep to meet your unique requirements, now and as you grow.

Hitachi Virtual Storage Platform provides numerous native VMware integrations, as shown in Figure 6.

Figure 6 Hitachi Virtual Storage Platform and VMware vSphere Integration Points**Note**

VMFS = VMware Virtual Machine File System, vMA = VMware vSphere Management Assistant, NMP = network management protocol, HBA = host bus adapter, VM = virtual machine, VSS = Microsoft Volume Shadow Copy Service, HDP = Hitachi Dynamic Provisioning, HDT = Hitachi Dynamic Tiering

VMware Native Multipathing Plug-In

One of the key advantages of the VSP is its global cache and fifth-generation switch architecture (see [Figure 6](#)). Together these capabilities provide uniform access and latency from each port to each storage resource and dramatically simplify multipathing.. Hitachi recommends using VMware's native multipathing software, as the controller has hardware-based algorithms for load balancing. This removes complexities for manually managing load within the controller and ESXi Host, as dual controller systems are susceptible to ESXi Host workload imbalances that require administrators to spend time manually diagnosing and mitigating.

vStorage API for Storage Awareness (VASA)

VMware vStorage API for Storage Awareness is a VMware vCenter 5.0 plug-in that provides integrated information of physical storage resources and information based on, topology, capability and state. This information is then used by vSphere for various features, including VMware Storage Distributed Resource Scheduler (SDRS) and profile-based storage.

VMware vSphere Storage APIs for Storage Awareness enables unprecedented coordination between vSphere/vCenter and the storage system. It provides built-in storage insight in vCenter to support intelligent VM storage provisioning, bolster storage troubleshooting and enable new SDRS-related use cases for storage.

Hitachi Data Systems supports vSphere Storage APIs for Storage Awareness through the plug-in, or “provider,” which is available through your HDS sales representative. This provider is compatible with vSphere and vCenter, and supports storage system makes and models as described on the VMware Compatibility Guide.

VMware Array Awareness Integration (VAAI)

VAAI enables the hypervisor to offload storage-specific tasks to the storage system. The benefit is directly proportional to the capabilities and performance of the storage system. Hitachi supports VAAI features natively on its storage systems, without the need for 3rd-party plug-ins or special software, reducing complexity, support and management overhead (see [Table 2](#)).

Hitachi is the only vendor to support all VAAI functions on externally virtualized storage systems. Storage virtualized behind Hitachi VSP inherits the features of VAAI from VSP. This has a substantial impact on the amount of IO workload that can be offloaded from the ESXi hosts into the storage array: Today, if you decide to move devices from one Pod to another using storage vMotion, the ESX host will be required to do the actual storage vMotion work, which will consume resources on the host and network. The Hitachi ability to extend VAAI functionality to external storage makes performance management much easier and reduces host overhead; this is because the VAAI features would offload the Storage vMotion between any VSP externally virtualized storage systems.

There are three primary functions for VAAI:

- **Xcopy:** vSphere can now offload cloning and storage vMotions to the storage system. VSP is the only storage virtualization platform where all of the storage vMotions between internal and external volumes will be handled by the storage system. Other architectures will require the host to do the vMotion as they are unable to offload IO operations that span multiple storage systems.
- **Write Same (Block Zero):** When provisioning a new VMDK, the storage system can now zero out the disk. There are 2 benefits here: First, this is offloaded to the storage system. This is useful for provisioning hundreds and thousands of VMs. Second, when used in conjunction with Hitachi Dynamic Tiering or Hitachi Dynamic Provisioning, hardware thin provisioning can be used for eagerzeroedthick VMDK file formats. This format is required to use VMware Fault Tolerance and clustering software like Microsoft Cluster Server (MSCS). This also simplifies provisioning; all VMs should now be deployed using the eagerzeroedthick format.
- **Hardware-assisted Locking (also known as Atomic Test and Set):** This primitive enables sub-lun locking which eliminates SCSI reservation conflicts that were seen when multiple virtual machines tried to lock the same physical device. This eliminates complexity associated with larger data stores and allows volumes and data stores to be sized based on I/O and capacity requirements. Hitachi offers sizing tools to help administrators with this function.


Table 2 *VAAI Performance Benefits on Hitachi Storage*

VMware Array Awareness Integration (VAAI) Feature	Performance Improvement	Why It Matters
Full Copy	18% faster	<ul style="list-style-type: none"> ▪ Speeds up cloning and storage vMotion ▪ Allows for faster copies of virtual machines (VMs)
Write Same (Block Zeroing)	85% faster performance	<ul style="list-style-type: none"> ▪ Speeds provisioning of new VMs ▪ Key to supporting large-scale VMware or virtual desktop interface (VDI) deployments
Hardware-assisted Locking (SCSI Reservation Lock)	25% to 35% performance improvement	<ul style="list-style-type: none"> ▪ Removes SCSI reservation conflicts ▪ Enables faster locking ▪ Improves VM density and performance <ul style="list-style-type: none"> ▪ Simplifies storage configuration ▪ Enables 2TB volumes ▪ Supports faster VM deployment ▪ Eliminates “boot storms” ▪ Enables scaling

Example: VMDK Provisioning + Block Zeroing + Hitachi Dynamic Provisioning

- A new 30GB “eagerzeroedthick” VMDK file was created on a Dynamic Provisioning volume (see [Figure 8](#)).
- With VAAI disabled, the virtual disk consumed 31GB of storage capacity on the Dynamic Provisioning pool.
- With VAAI enabled, it consumed only 1GB of storage capacity on the Dynamic Provisioning pool.

Figure 7 *A New eagerzeroedthick VMDK File is Created*

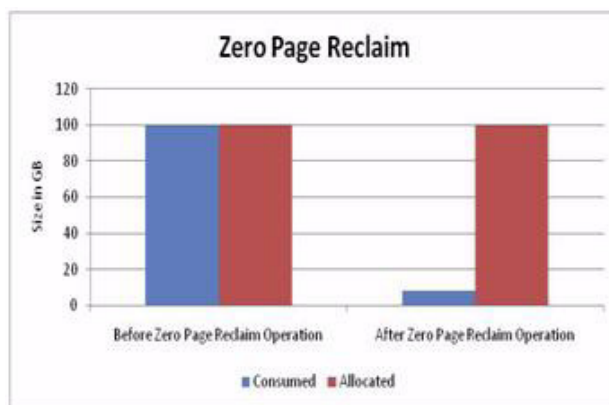


Block zeroing enables the formatting and provisioning of both “zeroedthick” and “eagerzeroedthick” VMDKs and in VMFS data stores to be handled by the storage system rather than the ESX hosts.

Example: VMDK Provisioning + Block Zeroing + Hitachi Dynamic Provisioning + Zero Page Reclaim

- A 100GB “zeroedthick” VMDK file was converted to an “eagerzeroedthick” VMDK file on a Dynamic Provisioning volume (see [A zeroedthick VMDK File is Converted to an eagerzeroedthick VMDCK File](#)).
- Reclaim space.
- With VAAI enabled, zero page reclaim is run and the “eagerzeroedthick” VMDK consumed 8GB of storage capacity on the Dynamic Provisioning pool.

Figure 8 *A zeroedthick VMDK File is Converted to an eagerzeroedthick VMDCK File*



With zero page reclaim, any pages that are zeroed within an “eagerzeroedthick” VMDK file are freed to be reallocated as needed.

For more information, see the Hitachi Virtual Storage Platform on the Hitachi Data Systems website.

VMware vSphere

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructure resources—CPUs, storage, networking—as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

VMware vSphere provides revolutionary benefits, but with a practical, non-disruptive evolutionary process for legacy applications. Existing applications can be deployed on VMware vSphere with no changes to the application or the OS on which they are running.

VMware vSphere provides a set of application services that enable applications to achieve unparalleled levels of availability, and scalability. For example, with VMware vSphere, all applications can be protected from downtime with VMware High Availability (HA) and VMware Fault Tolerance (FT), without the complexity of conventional clustering. In addition, applications can be scaled dynamically to meet changing loads with capabilities such as Hot Add and VMware Distributed Resource Scheduler (DRS).

For more information, see:

<http://www.vmware.com/products/vsphere/mid-size-and-enterprise-business/overview.html>

Dedicated SAN Design (Alternate Design)

Figure 9 illustrates the Dedicated SAN Design for the reference architecture. This design uses Fibre Channel switches to provide access to the Hitachi VSP storage controllers. Each NPV enabled Cisco UCS Fabric Interconnect attaches to one of the MDS 9500 series SAN switches with a Fibre Channel port channel creating two distinct data paths further enforced with independent VSAN definitions. The Fibre Channel port channel offers link fault tolerance and increased aggregate throughput between devices.

The Fibre Channel links from the MDS platforms are dual-homed into either Cluster on the Hitachi Virtual Storage Platform Controller Unit to provide link resiliency and traffic load balancing. VMware native multi-pathing using the round-robin hashing algorithm balances the I/O across the connections. Fibre Channel zoning is performed on the MDS 9500 directors and Fibre Channel masking on the Hitachi VSP for access control and added security within the storage domain.

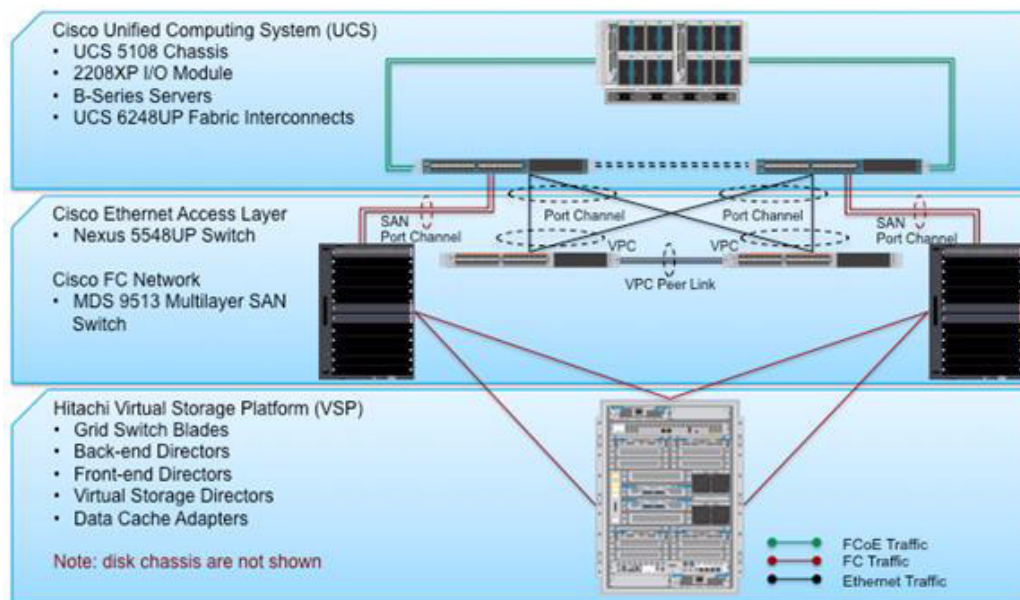


Note

The MDS offers many advanced storage services that were not tested as part of this effort. For more information, see: <http://www.cisco.com/en/US/products/ps5990/index.html>

The Ethernet forwarding configuration is identical to the Unified Access Layer Design. The Cisco UCS Fabric Interconnects are dual-homed to the Nexus 5500 supported access layer through port channels. The Nexus 5500 uses a distinct vPC for each Fabric Interconnect. Again, the use of port channels and vPCs improve the resiliency and capacity of the solution.

Figure 9 *Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System—Dedicated SAN Design*



Additional Components

The Dedicated SAN Design requires a SAN network to support a distinct Fibre Channel transport environment. The Cisco MDS 9500 Series Multilayer Director is highly recommended to implement a greenfield deployment or address current storage consolidation initiatives in brownfield environments. The MDS offers both SAN extension and other intelligent fabric services to all deployments.

Cisco MDS 9500 Series Multilayer Directors

Cisco MDS 9500 Series Multilayer Directors are high-performance, protocol-independent director-class SAN switches. These directors are designed to meet stringent requirements of enterprise data center storage environments. They provide:

- High availability
- Storage infrastructure connectivity for converged LAN and SAN fabrics
- Scalability
- Security
- Transparent integration of intelligent features for storage management

Cisco® MDS 9500 Series Multilayer Directors provide a comprehensive set of intelligent features onto a high-performance, protocol-independent switch fabric for extremely versatile data center SAN solutions. The MDS 9500 platform shares the same operating system (NX-OS) and management interface with other Cisco data center switches.

For more information, see:

<http://www.cisco.com/en/US/products/ps5990/index.html>

Domain and Element Management

This section provides general descriptions of the domain and element managers used during the validation effort. The following managers were used:

- Cisco UCS Manager
- Cisco Prime Data Center Manager
- Hitachi Dynamic Provisioning
- VMware vCenter Server

Cisco UCS Manager

Cisco UCS Manager provides unified, centralized, embedded management of all Cisco Unified Computing System software and hardware components across multiple chassis and thousands of virtual machines. Administrators use the software to manage the entire Cisco UCS as a single logical entity through an intuitive GUI, a command-line interface (CLI), or an XML API.

The Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The software gives administrators a single interface for performing server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection. Cisco UCS Manager service profiles and templates support versatile role- and policy-based management, and system configuration information can be exported to configuration management databases (CMDBs) to facilitate processes based on IT Infrastructure Library (ITIL) concepts.

Service profiles let server, network, and storage administrators treat Cisco UCS servers as raw computing capacity to be allocated and reallocated as needed. The profiles define server I/O properties and are stored in the Cisco UCS 6200 Series Fabric Interconnects. Using service profiles, administrators can provision infrastructure resources in minutes instead of days, creating a more dynamic environment and more efficient use of server capacity.

Each service profile consists of a server software definition and the server's LAN and SAN connectivity requirements. When a service profile is deployed to a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the profile. The automatic configuration of servers, network interface cards (NICs), host bus adapters (HBAs), and LAN and SAN switches lowers the risk of human error, improves consistency, and decreases server deployment times.

Service profiles benefit both virtualized and non-virtualized environments. The profiles increase the mobility of non virtualized servers, such as when moving workloads from server to server or taking a server offline for service or upgrade. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing virtual machine mobility.

For more Cisco UCS Manager information, visit:

<http://www.cisco.com/en/US/products/ps10281/index.html>

Cisco Prime Data Center Manager

Cisco Prime Data Center Network Manager (DCNM) is designed to help you efficiently implement and manage virtualized data centers. It includes a feature-rich, customizable dashboard that provides visibility and control through a single pane of glass to Cisco Nexus and MDS products.

Cisco Prime DCNM optimizes the overall uptime and reliability of your data center infrastructure and helps improve business continuity. This advanced management product:

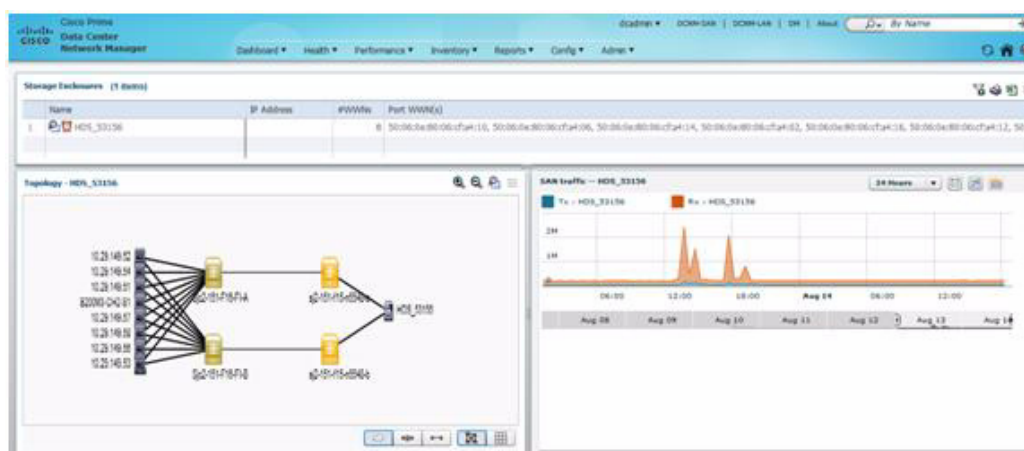
- Automates provisioning of data center LAN and SAN elements

- Proactively monitors the SAN and LAN, and detects performance degradation
- Helps secure the data center network
- Eases diagnosis and troubleshooting of data center outages
- Simplifies operational management of virtualized data centers

The primary benefits of Cisco Data Center Network Manager:

- Faster problem resolution
- Intuitive domain views that provide a contextual dashboard of host, switch, and storage infrastructures
- Real-time and historical performance and capacity management for SANs and LANs
- Virtual-machine-aware path analytics and performance monitoring
- Easy-to-use provisioning of Cisco NX-OS features with preconfigured, customized templates
- Customized reports which can be scheduled at certain intervals

Figure 10 *Cisco Prime Data Center Manager GUI Example*



For more information, see:

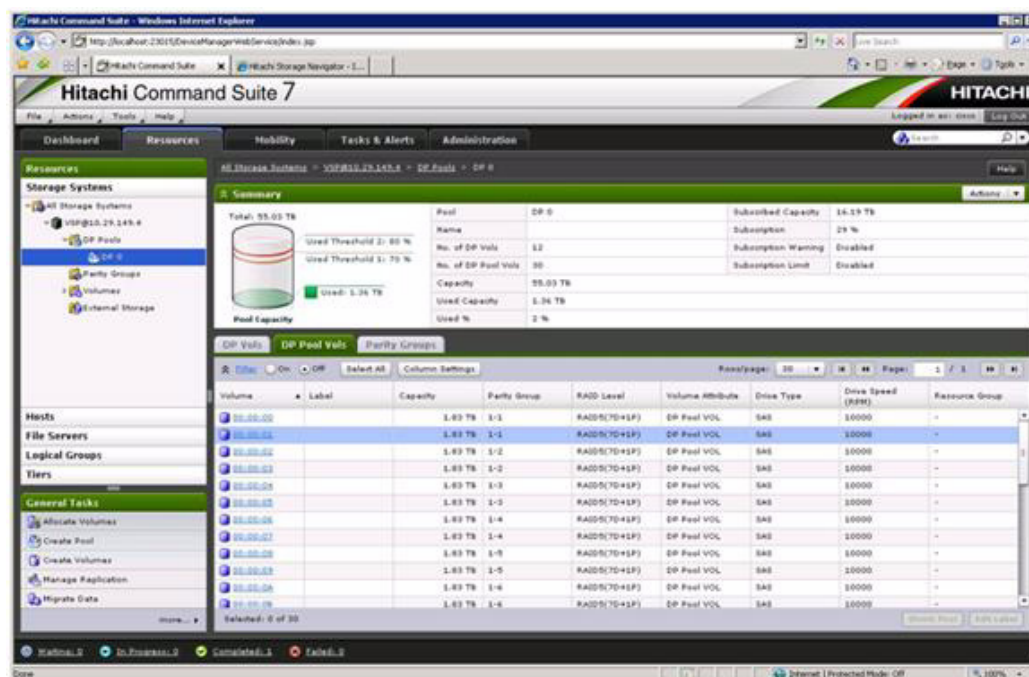
<http://www.cisco.com/en/US/products/ps9369/index.html>

Hitachi Command Suite v7

Hitachi Command Suite is a centralized software management framework that incorporates multiple IT and storage management disciplines, including:

- Storage resource management
- Tiered storage management
- Service level management

Figure 11 *Hitachi Command Suite v7 GUI Example*



Hitachi Command Suite provides advanced data management that improves storage operations, provisioning, optimization and resilience for Hitachi storage environments. It complements Hitachi storage systems, creating the most reliable and easiest-to-manage enterprise storage solution available.

The latest release of Hitachi Command Suite incorporates a number of key new enhancements, including:

- Agentless host discovery, which can be leveraged for basic storage management practices while only requiring host agent deployments for more advanced management functions
- Enhanced integration across HCS products with a new GUI, shared data repositories and task management, so tasks can either be executed immediately or scheduled for execution at a later time
- Improved usability with integrated use case wizards and best practices defaults for common administrative tasks
- Improved scalability to manage larger data centers with a single Hitachi Command Suite management server and improved performance to reduce the time required to execute tasks
- Streamlining of common administrative practices, which can unify management across all Hitachi storage systems and data types

VMware vCenter Server

VMware vCenter Server is the simplest and most efficient way to manage VMware vSphere, no matter whether you have ten VMs or tens of thousands of VMs. It Hitachi Command Suite provides advanced data management that improves storage operations, provisioning, optimization and resilience for Hitachi storage environments .It provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts and VMs. VMware vCenter Server gives administrators deep insight into the status and configuration of clusters, hosts, VMs, storage, the guest

OS and other critical components of a virtual infrastructure. Using VMware vCenter Server, a single administrator can manage 100 or more virtualization environment workloads, more than doubling typical productivity in managing physical infrastructure.

For more information, see:

<http://www.vmware.com/products/vcenter-server/overview.html>

vCenter Plugin - Hitachi Storage Manager for VMWare vCenter

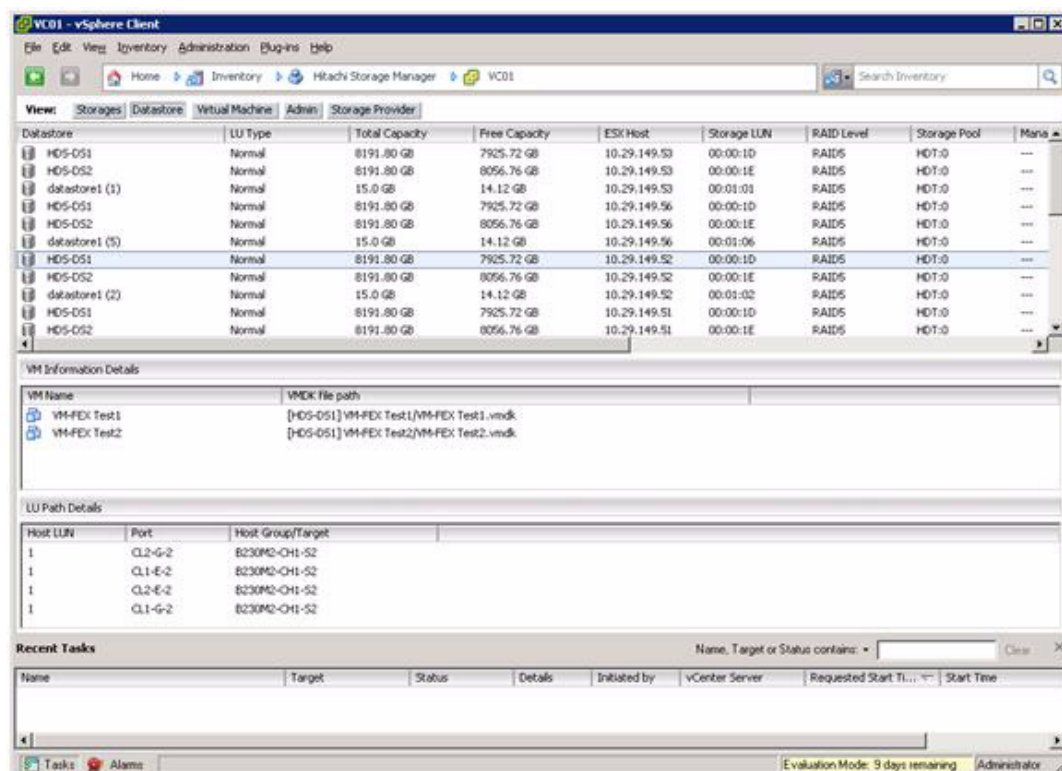
Storage Manager for VMware vCenter provides a scalable and extensible platform that forms the foundation for virtualization management. It centrally manages VMware vSphere environments, allowing IT administrators dramatically improved control over the virtual environment compared, to other management platforms.

Storage Manager for VMware vCenter is composed of the following main components:

- **VI Client**—VMware Infrastructure Client allows administrators and users to connect remotely to the vCenter Server or individual VMware ESX hosts from any Windows PC.
- **vCenter Server**—VMware vCenter server provides a scalable and extensible platform that forms the foundation for virtualization management.
- **ESX Host**—VMware virtual machine software for consolidating and partitioning servers in high-performance environment.
- **Storage Manager for VMware vCenter**—Storage Manager for VMware vCenter connects to VMware vCenter Server and associates the Hitachi Storage system information with VMware ESX Datastore/Virtual Machine information.

Figure 12 is a sample screen capture of Hitachi Storage Manager for VMware vCenter. In this example, VSP data stores and the associated virtual machines can be easily correlated as well as the transport path between LUN and ESXi node.

Figure 12 Hitachi Storage Manager vCenter Plugin Datastore View Example



Reference Architecture Design and Configuration

Physical Build

Hardware and Software Revisions

Table 3 describes the hardware and software versions used during validation. It is important to note that Cisco, Hitachi and VMware have interoperability matrixes that should be referenced to determine support for any one specific implementation of Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System. Please reference the following links:

- Hitachi Data Systems and Cisco UCS Nexus:
<http://www.hds.com/assets/pdf/all-storage-models-for-cisco-ucs.pdf>
- VMware, Cisco Unified Computing System, and Hitachi Data Systems:
<http://www.vmware.com/resources/compatibility>

Table 3 **Validated Software Version**

Device	Image	Comments
Cisco UCS Fabric Interconnects	2.0(3a)	Includes the Cisco UCS-IOM 2208XP and Cisco UCS Manager
Cisco UCS B Series Servers	2.0(3a)	B200 M3 using Cisco UCS VIC 1240 and B230 M2 using Cisco Cisco UCS VIC M81 KR
Cisco Nexus 5000 NX-OS	5.1(3)N2(1a)	
Hitachi Virtual Storage Platform	Microcode 70-04-31-00/ 00	
VMware vSphere ESXi	5.0.0, 469512	
VMware vCenter	5.0.0, 455964	
Cisco MDS 9500	5.2(2d)	
Cisco Nexus 1000v	4.2(1)SV1(5.1a)	
Cisco Prime Data Center Manager	6.1(1a)	

Cabling Information

This section details the physical cabling of UCP Select for VMware vSphere with Cisco Unified Computing System designs consisting of the following sub-sections:


- Common Connectivity
- Unified Access Layer Design Fibre Channel Connectivity
- Dedicated SAN Design Fibre Channel Connectivity



Note

This cabling section does not include integration of the management infrastructure.

Figure 13 *Colorcoded—Unified Access Layer Design Cabling Guide*



Common Connectivity

Table 4 *Cisco Nexus 5548UP—Ethernet Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 A	Eth 1/9-10	10Gbe	Cisco Nexus 5548 B	Eth 1/9 -10
	Eth 1/13-14	10GbE	Cisco UCS Fabric Interconnect A	Eth 1/13-14
	Eth 1/15-16	10GbE	Cisco UCS Fabric Interconnect B	Eth 1/13-14

Table 5 *Cisco Nexus 5548UP—B Ethernet Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 B	Eth 1/9-10	10Gbe	Cisco Nexus 5548 A	Eth 1/9 -10
	Eth 1/13-14	10GbE	Cisco UCS Fabric Interconnect A	Eth 1/13-14
	Eth 1/15-16	10GbE	Cisco UCS Fabric Interconnect B	Eth 1/13-14

Table 6 *Cisco UCS Fabric Interconnect A Ethernet Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	Eth1/ 13-14	10GbE	Cisco Nexus 5548 A	Eth1/ 13-14
	Eth1/ 15-16	10GbE	Cisco Nexus 5548 B	Eth1/13-14
	Eth1/ 1-4	10GbE	Cisco UCS B-Series FEXA	Port 1-4
	Eth 2/ 1-4	8Gbps (FC)	Cisco Nexus 5548UP A	Eth 1/ 29-32

Table 7 *Cisco UCS Fabric Interconnect B Ethernet Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	Eth1/ 13-14	10GbE	Cisco Nexus 5548 A	Eth1/ 13-14
	Eth1/ 15-16	10GbE	Cisco Nexus 5548 B	Eth1/ 15-16
	Eth1/ 1-4	10GbE	Cisco UCS B-Series chassis FEX B	Port 1-4
	Eth/ 2/ 1-4	10GbE	Cisco Nexus 5548UP B	Eth 29-32

Table 8 *Cisco UCS IOM 2208XP—FEX A*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS-IOM FEXA	Eth1/ 1	10GbE	Cisco UCS Fabric Interconnect A	Eth1/ 1
	Eth1/ 2	10GbE	Cisco UCS Fabric Interconnect A	Eth1/ 2
	Eth1/ 3	10GbE	Cisco UCS Fabric Interconnect A	Eth1/ 3
	Eth 1/ 4	10GbE	Cisco UCS fabric interconnect A	Eth1/ 4

Table 9 *Cisco UCS-IOM 2208XP—FEX B*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS-IOM FEXB	Eth1/ 1	10GbE	Cisco UCS Fabric Interconnect B	Eth1/ 1
	Eth1/ 2	10GbE	Cisco UCS Fabric Interconnect B	Eth1/ 2
	Eth1/ 3	10GbE	Cisco UCS Fabric Interconnect B	Eth1/ 3
	Eth 1/ 4	10GbE	Cisco UCS fabric interconnect B	Eth1/ 4

Unified Access Layer Design Fibre Channel Connectivity

Table 10 *Cisco Nexus 5548UP-A Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 A	fc1/25	8Gbps (FC)	HDS-VSP	CL1 -1A
	fc1/26	8Gbps (FC)	HDS-VSP	CL2 -2C
	fc1/27	8Gbps (FC)	HDS-VSP	CL1 – 1E
	fc1/28	8Gbps (FC)	HDS-VSP	CL2 – 2G
	fc1/29-32	8Gbps (FC)	Cisco UCS Fabric Interconnect A	Eth 2/ 1-4

Table 11 *Cisco Nexus 5548UP-B Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco Nexus 5548 B	fc1/25	8Gbps (FC)	HDS-VSP	CL2 -2A
	fc1/26	8Gbps (FC)	HDS-VSP	CL1 -1C
	fc1/27	8Gbps (FC)	HDS-VSP	CL2 – 2E
	fc1/28	8Gbps (FC)	HDS-VSP	CL1 – 1G
	fc1/29-32	8Gbps (FC)	Cisco UCS Fabric Interconnect B	Eth2/ 1-4

Figure 14 *Cisco UCS Fabric Interconnect A Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	FC2/ 1-4	8Gbps (FC)	Cisco Nexus 5548 A	fc1/ 29-32

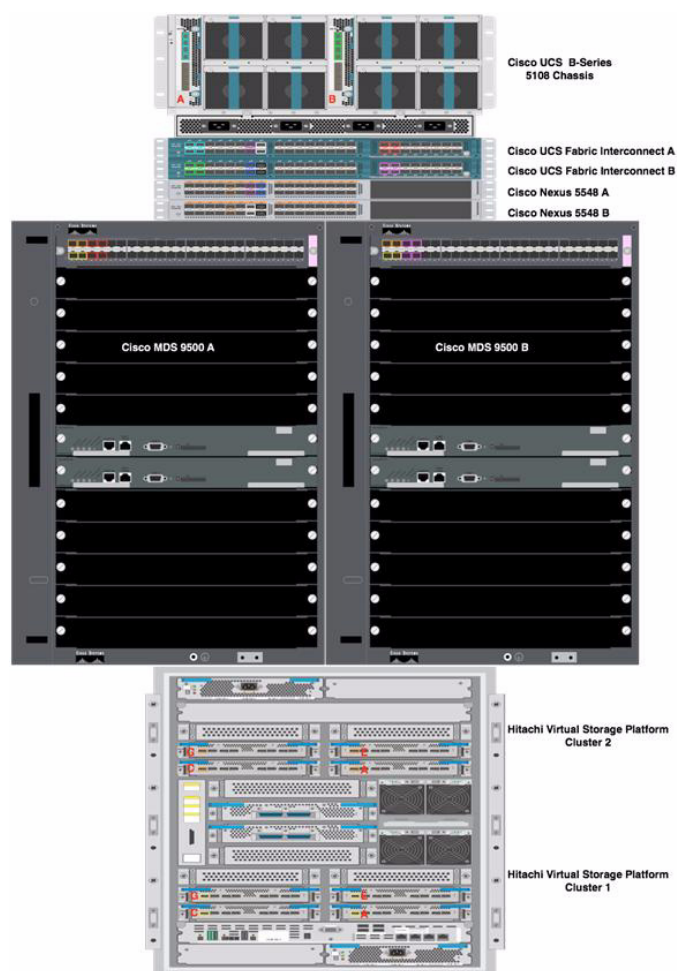
Table 12 *Cisco UCS Fabric Interconnect B Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	FC2/ 1-4	8Gbps (FC)	Cisco Nexus 5548 B	fc1/ 29-32

Table 13 ***Hitachi VSP Fibre Channel Cabling***

Local Device	Local Port	Connection	Remote Device	Remote Port
HDS-VSP	CL1 -1 A	8Gbps (FC)	Cisco Nexus 5548 A	fc1/ 25
	CL1 -1 C	8Gbps (FC)	Cisco Nexus 5548 B	fc1/ 26
	CL1 -1 E	8Gbps (FC)	Cisco Nexus 5548 A	fc1/ 27
	CL1 -1 G	8Gbps (FC)	Cisco Nexus 5548 B	fc1/ 28
	CL2 -2 A	8Gbps (FC)	Cisco Nexus 5548 B	fc1/ 25
	CL2 -2 C	8Gbps (FC)	Cisco Nexus 5548 A	fc1/ 26
	CL2 -2 E	8Gbps (FC)	Cisco Nexus 5548 B	fc1/ 27
	CL2 -2 G	8Gbps (FC)	Cisco Nexus 5548 A	fc1/ 28

Figure 15 **Dedicated SAN Design Fibre Channel Connectivity**



Note

The Dedicated SAN Hitachi Unified Compute Platform Select for VMware® vSphere with Cisco Unified Computing System design uses the same Ethernet cabling scheme documented in the Common Connectivity section.

Table 14 **Cisco UCS Fabric Interconnect A Fibre Channel Cabling**

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect A	2/ 1	8Gbps (FC)	Cisco MDS 9500 A	Fc1 / 5
	2/ 2	8Gbps (FC)	Cisco MDS 9500 A	Fc1 / 6
	2/ 3	8Gbps (FC)	Cisco MDS 9500 A	Fc1 / 7
	2/ 4	8Gbps (FC)	Cisco MDS 9500 A	Fc1 / 8

Table 15 *Cisco UCS Fabric Interconnect B Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco UCS Fabric Interconnect B	2/ 1	8Gbps (FC)	Cisco MDS 9500 B	Fc1/ 5
	2/ 2	8Gbps (FC)	Cisco MDS 9500 B	Fc1/ 6
	2/ 3	8Gbps (FC)	Cisco MDS 9500 B	Fc1/ 7
	2/ 4	8Gbps (FC)	Cisco MDS 9500 B	Fc1/ 8

Table 16 *Cisco MDS 9500 A Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9500 A	Fc1/ 1	8Gbps (FC)	HDS-VSP	CL2-2A
	Fc1/ 2	8Gbps (FC)	HDS-VSP	CL1-1C
	Fc1/ 3	8Gbps (FC)	HDS-VSP	CL2-2E
	Fc1/ 4	8Gbps (FC)	HDS-VSP	CL1-1G
	Fc1/ 5	8Gbps (FC)	Cisco UCS Fabric Interconnect A	fc2/ 1
	Fc1/ 6	8Gbps (FC)	Cisco UCS Fabric Interconnect A	fc2/ 2
	Fc1/ 7	8Gbps (FC)	Cisco UCS Fabric Interconnect A	fc2/ 3
	Fc1/ 8	8Gbps (FC)	Cisco UCS Fabric Interconnect A	fc2/ 4

Table 17 *Cisco MDS 9500 B Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
Cisco MDS 9500 B	Fc1/ 1	8Gbps (FC)	HDS-VSP	CL1-1A
	Fc1/ 2	8Gbps (FC)	HDS-VSP	CL2-2C
	Fc1/ 3	8Gbps (FC)	HDS-VSP	CL2-2G
	Fc1/ 4	8Gbps (FC)	HDS-VSP	CL1-1E
	Fc1/ 5	8Gbps (FC)	Cisco UCS Fabric Interconnect B	fc2/ 1
	Fc1/ 6	8Gbps (FC)	Cisco UCS Fabric Interconnect B	fc2/ 2
	Fc1/ 7	8Gbps (FC)	Cisco UCS Fabric Interconnect B	fc2/ 3
	Fc1/ 8	8Gbps (FC)	Cisco UCS Fabric Interconnect B	fc2/ 4

Table 18 *Hitachi VSP Fibre Channel Cabling*

Local Device	Local Port	Connection	Remote Device	Remote Port
HDS-VSP	CL1 -1A	8Gbps (FC)	Cisco MDS 9500 B	fc1/ 1
	CL1 -1C	8Gbps (FC)	Cisco MDS 9500 A	fc1/ 2
	CL1 -1E	8Gbps (FC)	Cisco MDS 9500 B	fc1/ 4
	CL1 -1G	8Gbps (FC)	Cisco MDS 9500 A	fc1/ 4
	CL2 -2A	8Gbps (FC)	Cisco MDS 9500 A	fc1/ 1
	CL2 -2C	8Gbps (FC)	Cisco MDS 9500 B	fc1/ 2
	CL2 -2E	8Gbps (FC)	Cisco MDS 9500 A	fc1/ 3
	CL2 -2G	8Gbps (FC)	Cisco MDS 9500 B	fc1/ 3

Logical Build

This section details the logical configuration validated for Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System. The topics covered include:

- UCP Select for VMware vSphere with Cisco Unified Computing System —Unified Access Layer Design using the Nexus 1000v
- UCP Select for VMware vSphere with Cisco Unified Computing System —Unified Access Layer Design using VM-FEX
- UCP Select for VMware vSphere with Cisco Unified Computing System —Dedicated SAN Design



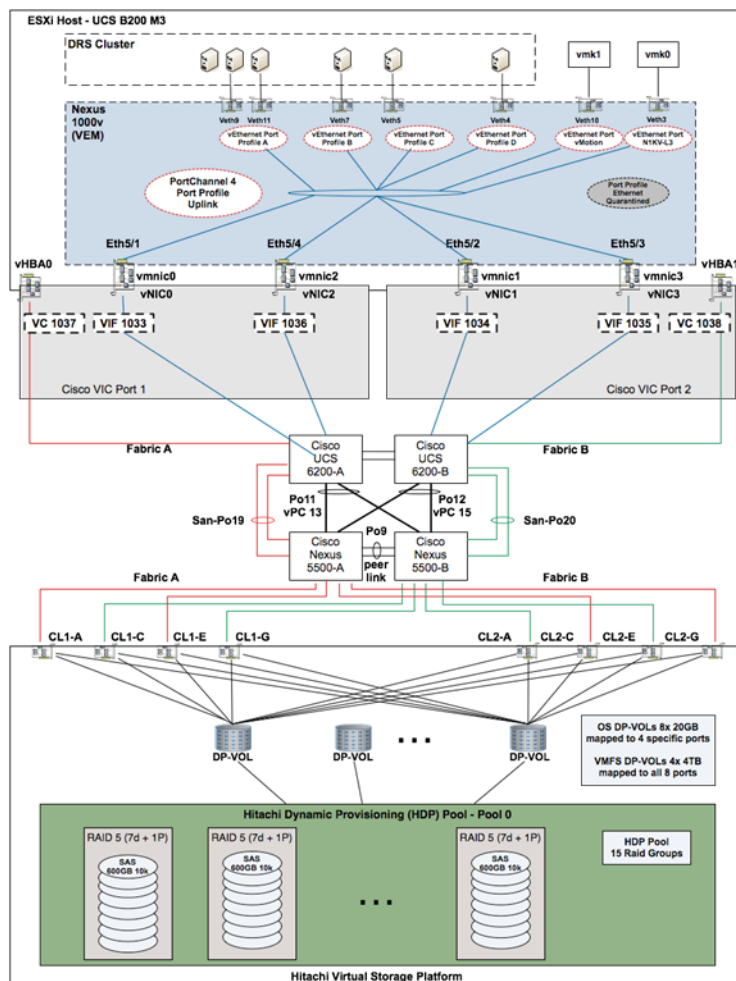
Note

UCP Select for VMware vSphere with Cisco Unified Computing System using the Nexus 1000v or VM-FEX employs the same physical infrastructure and can be deployed simultaneously within the same VMware vCenter server environment.

Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System—Unified Access Layer Design using the Nexus 1000v

Figure 16 illustrates the UCP Select for VMware vSphere with Cisco Unified Computing System Unified Access Layer Design using the Nexus 1000v virtual distributed switch. The validated design is physically redundant across the stack addressing Layer 1 high availability data center requirements, but there are additional Cisco and Hitachi technologies and features that make for an even more robust solution.

Figure 16 UCP Select for VMware vSphere with Cisco Unified Computing System – Unified Access Layer Design Using the Nexus 1000v



Cisco Nexus 5500

As Figure 16 shows, the Nexus 5500 provides Ethernet and Fibre Channel connectivity for the Cisco UCS domain as well as Fibre Channel services to the Hitachi VSP. From an Ethernet perspective, the Nexus 5500 uses virtual PortChannel (vPC) allowing links that are physically connected to two different Cisco Nexus 5000 Series devices to appear as a single PortChannel to a third device in this case the Cisco UCS Fabric Interconnects. vPC provides the following benefits:

- Allows a single device to use a PortChannel across two upstream devices
- Eliminates Spanning Tree Protocol blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth
- Provides fast convergence if either the link or a device fails
- Provides link-level resiliency
- Helps ensure high availability

vPC requires a “peer link” which is documented as port channel 9 in [Figure 16](#).

The Nexus 5500 in the Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System design provides Fibre Channel services to the Cisco Unified Computing System and VSP platforms. Internally the Nexus 5500 platforms are performing FC zoning to enforce access policy between Cisco UCS-based initiators and VSP-based targets. Multiple FC dedicated Nexus 5500 unified ports are assigned to support the FC multi-pathing capabilities and port density requirements of the Hitachi VSP platform.

The Nexus 5500 employs two Fibre Channel port channels; these were configured between the Nexus 5500s and the UCS Fabric Interconnects. Each port channel, SanPo19 and SanPo20 carries a distinct VSAN to maintain SAN “A” , SAN “B” best practice separation. In addition to SAN segmentation, the port channels provide the additional benefits of link aggregation that include link resiliency and improved aggregate available throughput.

For configuration details refer to the Cisco Nexus 5000 series switches configuration guides:

http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html .

Cisco Unified Computing System

The Cisco Unified Computing System supports the virtual server environment by providing a robust, highly available, and extremely manageable compute resource. As [Figure 16](#) illustrates the components of the Cisco Unified Computing System offer physical redundancy and a set of logical structures to deliver the Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System . The Cisco 1240 Virtual Interface Card (VIC) presents multiple virtual PCIe devices to the ESXi host as vNICs. Each vNIC has a virtual identity and circuit supported by the Cisco Unified Computing System hardware and managed by the Cisco UCS Manager. The VMware vSphere environment identifies the vNICs as vmnics available for network connectivity.

This effort referenced the detailed configuration information for the Cisco UCS manager available at http://www.cisco.com/en/US/products/ps10281/prod_configuration_examples_list.html

VMware vCenter and vSphere

VMware vSphere provides a platform for virtualization comprised of multiple components and features. In this validation effort the following were used:

- VMware ESXi—A virtualization layer run on physical servers that abstracts processor, memory, storage, and resources into multiple virtual machines.
- VMware vCenter Server—The central point for configuring, provisioning, and managing virtualized IT environments. It provides essential datacenter services such as access control, performance monitoring, and alarm management.
- VMware vSphere SDKs—Feature that provides standard interfaces for VMware and third-party solutions to access VMware vSphere.
- vSphere Virtual Machine File System (VMFS)—A high performance cluster file system for ESXi virtual machines.
- vSphere High Availability (HA)—A feature that provides high availability for virtual machines. If a server fails, affected virtual machines are restarted on other available servers that have spare capacity.

- vSphere Distributed Resource Scheduler (DRS)—Allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines. This feature includes distributed power management (DPM) capabilities that enable a datacenter to significantly reduce its power consumption.

In this validation effort, multiple Cisco UCS B-Series servers are SAN booted as VMware ESXi nodes. The ESXi nodes consisted of Cisco UCS B200-M3 and B230-M2 series blades with a mix of processor and memory configurations with Cisco VIC adapter models 1240 and M81KR. These nodes were allocated to two distinct VMware DRS and HA enabled clusters. One cluster employed the Nexus 1000v as the virtual distributed switch while the other cluster used the Cisco UCS VM-FEX based virtual distributed switch (see the section Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System —Unified Access Layer Design using Cisco VM-FEX). The Cisco UCS Manager and Nexus 1000v integrate through the vCenter Server extensions as plugins.

**Note**

A single ESXi node cannot support both networking models.

As illustrated in [Figure 16](#), the Cisco 1240 VIC presents four 10 GE VNICs to the ESXi node. VMware identifies these as vmnics that can be assigned to a vSwitch or virtual distributed switch through vCenter. The ESXi operating system is unaware these are virtual network interface adapters. The result is a dual-homed ESXi node to the remaining network.

In addition to VNICs, the ESXi host is presented two virtual host bus adapters (vHBA). These adapters allow the ESXi host to boot remotely from the VSP and the node to access LUNs and VMFS datastores on the VSP. Again, the host is unaware these are virtualized adapters. The ESXi node has connections to two independent Fabrics A and B. The details of the Hitachi storage configuration can be found below in the Hitachi Virtual Storage Platform section.

Cisco Nexus 1000v

The Cisco Nexus 1000v is a virtual distributed switch that fully integrates into a vSphere enabled environment. The Cisco Nexus 1000v operationally emulates a physical modular switch, with a Virtual Supervisor Module (VSM) providing control and management functionality to multiple line cards. In the case of the Nexus 1000v, the ESXi nodes become modules in the virtual switch when the Cisco Virtual Ethernet Module (VEM) is installed. [Figure 17](#) details the Cisco Nexus 1000v architecture.

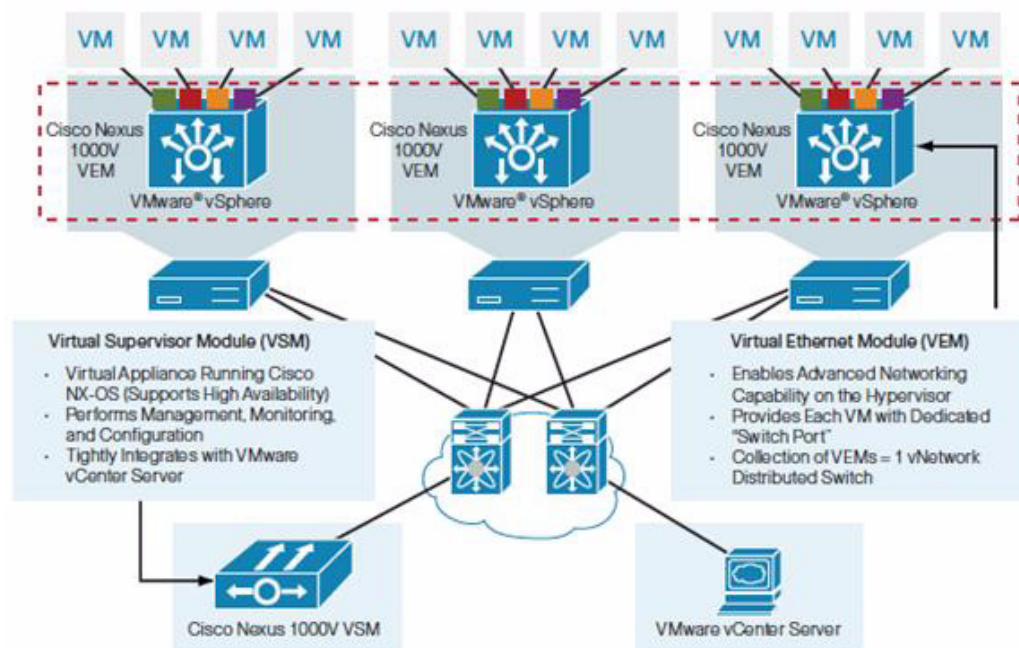
Figure 17 Cisco Nexus 1000v Architecture

Figure 17 shows a single ESXi node with a VEM registered to the Cisco Nexus 1000v VSM. The ESXi vmnics are presented as Ethernet interfaces in the Nexus 1000v. In this example, the ESXi node is the fifth module in the virtual distributed switch as the Ethernet interfaces are labeled as module/interface #. The VEM takes configuration information from the VSM and performs Layer 2 switching and advanced networking functions such as:

- PortChannels
- Quality of service (QoS)
- Security: Private VLAN, access control lists (ACLs), and port security
- Monitoring: NetFlow, Switch Port Analyzer (SPAN), and Encapsulated Remote SPAN (ERSPAN)
- vPath providing efficient traffic redirection to one or more chained services such as the Cisco Virtual Security Gateway and Cisco ASA 1000v

**Note**

The Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System architecture will fully support other intelligent network services offered through the Nexus 1000v such as Cisco VSG, ASA1000v, and vNAM.

The Nexus 1000v supports port profiles. Port profiles are logical templates that can be applied to the Ethernet and virtual Ethernet interfaces available on the Nexus 1000v. In the UCP Select for VMware vSphere with Cisco Unified Computing System architecture, the Nexus 1000v aggregates the Ethernet uplinks into a single port channel for fault tolerance and improved throughput. The VM facing virtual Ethernet ports employ port profiles customized for each virtual machines network, security and service level requirements.

For more information about “Best Practices in Deploying Cisco Nexus 1000V Series Switches on Cisco UCS B and C Series Cisco UCS Manager Servers” see

http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-558242.html

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform for this reference architecture consisted of a single frame that housed a Control Unit and a Disk Unit. [Table 19](#) details the physical characteristics of Virtual Storage Platform that was used for testing.

Table 19 *Hitachi Virtual Storage Platform Characteristics*

Number of Frames	1
Number of Control Units	1
•Number of Front End Directors	5
•Number of FC Ports	80
•Cache	96GB
•Number of Backend Directors	1
Number of Disk Units	1
•Number of Disks	120
•Type of Disks	2.5" SAS, 10000 RPM
•Capacity of Disks	600GB
•Number of Spare Disks	1

Hitachi Dynamic Provisioning

For this reference architecture, Hitachi Dynamic Provisioning was used on Virtual Storage Platform to define the storage block. This simplifies the management of the storage through the use of over provisioning, wide striping and on-line expansion of dynamic provisioning pools.

A storage block is comprised of a single dynamic provisioning pool. The sizing of the pool was based on assumptions made about the disks used and the VMware environment hosted on Cisco UCS servers

- Disk—Disk speed, RAID level, and the required IOPS
- VMware environment—Number of UCS blades and number of VM's to be hosted.

This reference architecture used dynamic provisioning volumes of 4TB since ESX5 (vs. ESX4) fully supports a maximum LUN size over 2TB. Four 4TB VMFS LUNs were provisioned to the four blade servers with each LUN capable of hosting 256 virtual machines. [Table 20](#) details the configuration for the dynamic provisioning pool.

Table 20 **Dynamic Provisioning Pool Configuration**

Pool ID	0
Number of Physical Disks in Pool	120
Number of RAID Groups	15
Raid Type	5
Raid Configuration	7D+1P
Capacity of Each RAID Group	3.66 TB
Total Pool Capacity	55.04 TB
Number of 20GB DP Volumes	10 for OS
Number of 4TB DP Volumes	4 for VMFS

You can add more RAID-5 (7D+1P) groups to the dynamic provisioning pool to max out the pool capacity to more than 5PB. With dynamic provisioning, the data automatically re-distributes across the newly added RAID groups.

**Note**

For 100 percent Data Availability Guarantee HDS requires a RAID6 configuration of 6D + 2P. 4TB is the maximum LUN size in environments requiring replication.

VMware VAAI

VAAI was enabled on the Host Groups. This was done by setting Host Mode to 01[VMware] and enabling Host Mode Option 63 (Support option for vStorage APIs base on T10 standards).

DP Volume Allocation

Each Cisco UCS Blade server was allocated a 20GB Dynamic Provisioning Volume (DP-VOL) for ESX OS Installation for SAN Boot. Each UCS Blade server was also allocated 4 shared 4TB DP-VOL for VMFS.

Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System—Unified Access Layer Design using Cisco VM-FEX

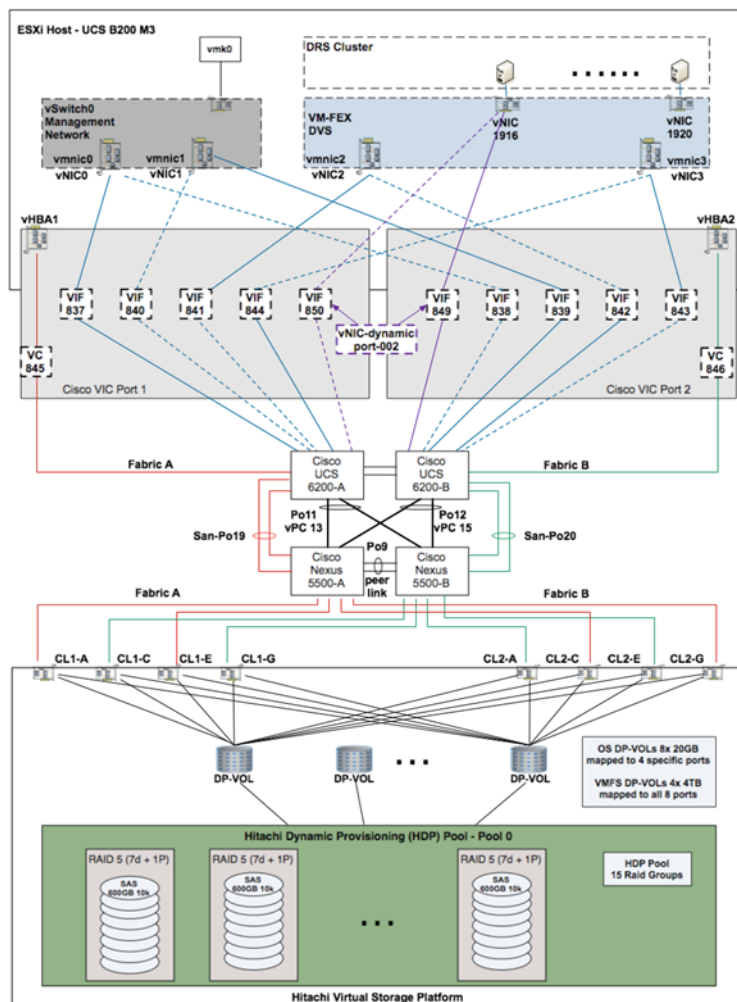
Cisco Virtual Machine Fabric Extender (VM-FEX) is a Cisco technology that addresses management and performance concerns in a data center by unifying physical and virtual switch management. The use of Cisco VM-FEX collapses both virtual and physical networking into a single infrastructure reducing the number of network management points and enabling consistent provisioning, configuration and management policy within the enterprise. The Cisco UCS and VMware vCenter are connected for a fully integrated management solution.

Figure 18 illustrates the use of Cisco VM-FEX in the Hitachi Unified Compute Platform Select for VMware vSphere with Cisco UCS design. On initial review the UCP Select for VMware vSphere with Cisco UCS design using VM-FEX technology appears identical to the previously described Nexus 1000v-based UCP Select for VMware vSphere with Cisco UCS model. The solution offers the same fully redundant physical infrastructure, as well as, identical Hitachi VSP and Cisco Nexus 5500 configurations. It is at the virtual access layer within the Cisco UCS environment that the VM-FEX design differs. The decision to use VM-FEX is typically driven by application requirements such as performance and the operational preferences of the IT organization.

**Note**

During testing, the Cisco VM-FEX and Cisco Nexus 1000v were deployed simultaneously within the same UCP Select for VMware vSphere with Cisco Unified Computing System system.

Figure 18 *UCP Select for VMware vSphere with Cisco Unified Computing System—Unified Access Layer Design Using VM-FEX*



The Cisco UCS Virtual Interface Card (VIC) offers each VM a virtual Ethernet interface or vNIC. This vNIC provides direct access to the Fabric Interconnects and Nexus 5500 series switches where forwarding decision can be made for each VM using a VM-FEX interface.

As detailed in [Figure 18](#), the path for a single VM is fully redundant across the Cisco fabric. The VM has an active virtual interface (VIF) and standby (VIF) defined on the adapter, an adapter that is dual-homed to Fabric A and B. Combined with the Cisco UCS Fabric Failover feature the VM-FEX solution provides fault tolerance and removes the need for software based HA teaming mechanisms. If the active uplink fails the vNIC will automatically fail over to the standby uplink and simultaneously update the network through gratuitous ARP. In this example, the active links are solid and the standby links are dashed.

The Cisco Fabric Extender technology provides both static and dynamic vNICs. As illustrated in this example, vNICs 0 through 4 are static adapters presented to the VMware vSphere environment. Static vNIC0 and vNIC1 are assigned to the local vSwitch supporting vSphere management traffic, while static vNICs 2 and 3 are associated with the VM-FEX distributed virtual switch (DVS). From a VMware vSphere perspective the vNICs are PCIe devices and do not require any special consideration or configuration. As shown, the Cisco UCS vNIC construct equates to a VMware virtual network interface card (vmnic) and is identified as such.

Dynamic vNICs are allocated to virtual machines and removed as the VM reaches the end of its lifecycle. [Figure 18](#) details a dynamic vNIC associated with a particular VM. From a vSphere perspective; the VM is assigned to the VM-FEX DVS on port 1916. This port maps to two VIFs, 849 and 850, which are essentially an active/standby HA pair. The maximum number of Virtual Interfaces (VIF) that can be defined on a Cisco VIC Adapter depends on the following criteria and must be considered in any VM-FEX design:

- The presence of jumbo frames
- The combination of Fabric Interconnects (6100 / 6200) and Fabric Extenders (2104 / 2208)
- The maximum number of port links available to the UCS IOM Fabric Extender
- The number of supported static and dynamic vNICs and vHBAs on the Cisco VIC Adapters
- The version of vSphere version


Note

VM-FEX requires that the ESXi host must have the Cisco Virtual Ethernet Module (VEM) software bundled installed.

Cisco Unified Computing System and VMware vSphere Integration

The Cisco Unified Computing System offers a VMware vCenter plugin for easy integration with the vSphere environment. As seen in [Figure 19](#), the Cisco UCS Manager indicates the vCenter instance it is associated with, VC01, as well as the virtual distributed switch it is allied, namely VM-FEX-DVS. This same figure indicates there are five port profiles created on the UCS and available to the VMware administrator for VM assignment as port groups.

[Figure 19](#) also identifies that VM “VM-FEX Test1” is using vNIC 1916. This is the same VM illustrated in [Figure 18](#) above. 1916 is the VDS port number that aligns with the vCenter view of the networking environment captured in [Figure 20](#). Notice the port group name as it is defined within Cisco UCS Manager and used within vCenter.

Figure 19 Cisco UCS Manager VM-FEX DVS Example

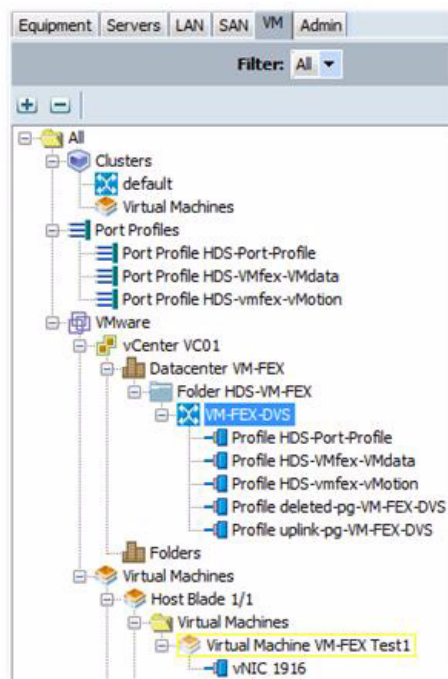
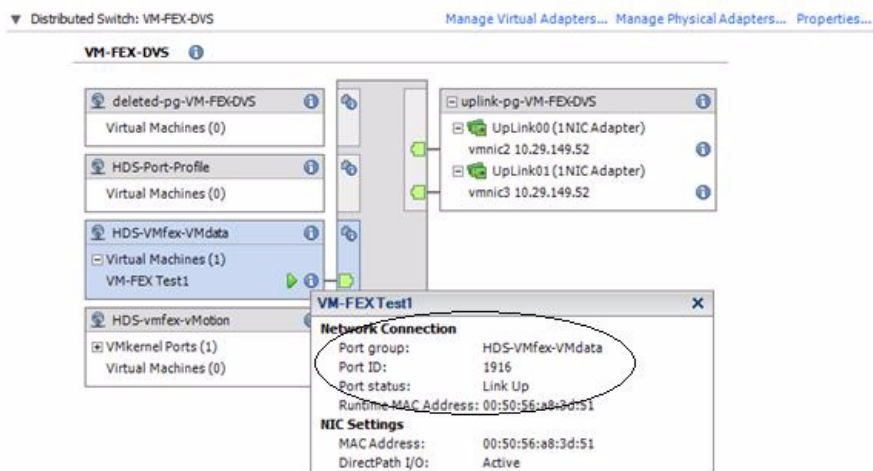


Figure 20 VMware Virtual Distributed Switch View of Cisco VM-FEX DVS



VM-FEX is configurable in standard or high performance mode from the Cisco UCS Manager port profile tab. In standard mode some of the ESXi nodes virtualization stack is used for VM network I/O. In high performance mode the VM completely bypasses the hypervisor and DVS accessing the Cisco VIC adapter directly. The high performance model takes advantage of VMware DirectPath I/O. DirectPath offloads the host CPU and memory resources that are normally consumed managing VM networks. This is a design choice primarily driven by performance requirements and VMware feature availability.

**Note**

The following VMware vSphere features are only available for virtual machines configured with DirectPath I/O on the Cisco Unified Computing Systems through Cisco Virtual Machine Fabric Extender distributed switches.

- vMotion
- Hot adding and removing of virtual devices
- Suspend and resume
- High availability
- DRS
- Snapshots

The following features are unavailable for virtual machines configured with DirectPath on any server platform:

- Record and replay
- Fault tolerance

For more information about “Cisco VM-FEX Best Practices for VMware ESX Environment Deployment Guide” see

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/vm_fex_best_practices_deployment_guide.html#wp9001031

Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System—Dedicated SAN Design (Alternate Design Only)

Figure 21 illustrates the UCP Select for VMware vSphere with Cisco UCS dedicated SAN design using the Nexus 1000v virtual distributed switch. It is important to clarify that the UCP Select for VMware vSphere with Cisco UCS- dedicated SAN design readily supports the use of Cisco's VM-FEX technology. From a compute and Ethernet Networking perspective the UCP Select for VMware vSphere with Cisco UCS dedicated SAN design is identical to the previously defined designs. The only modification is the transport path for block storage access. ?

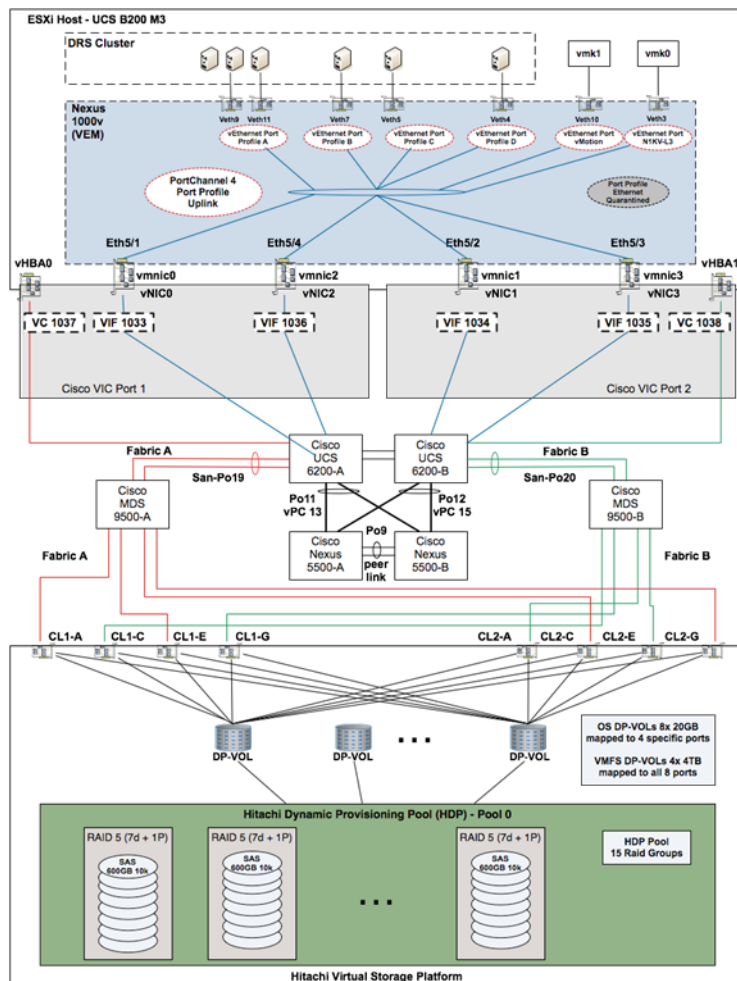
As shown in this example, the Cisco UCS Fabric Interconnects in end host mode (NPV mode) using Fibre Channel F-port channels to connect to the Cisco MDS 9500 series switches with the NPIV feature enabled. Each Cisco UCS Fabric Interconnect and MDS switch is supporting a distinct fabric with aggregated links for path redundancy and resiliency. There is no single point of failure from the Cisco UCS vHBA to the VSP cluster interfaces. In this design, the VMware Native Multipathing Plugin (NMP) provides MPIO support by selecting the physical path for FC transport using a round-robin based hash for load balancing. ?

The VSP clusters are dual-homed to the MDS switches providing redundant access to the storage environment. Essentially there are redundant fabrics with redundant links to redundant switches. In addition, the port density of both the Cisco MDS and VSP will readily allow storage administrators to adjust the fan-in/fan-out oversubscription ratio of the environment to address any potential throughput capacity issues.

**Note**

For additional redundancy it is best practice to divide connections across multiple line cards to eliminate any single point of failure.

Figure 21 *Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System—Dedicated SAN Design Using Nexus 1000v*



To restrict access across the Fibre Channel fabric SAN networks use zoning to associate or group initiators with targets. This fabric-based service helps enforce data ownership and reduce data corruption because the SAN fabric is carved into subsets with access control policies preventing unauthorized access to data and communication between hosts across the SAN.

The Cisco MDS provides Fibre Channel zoning using three zoning modes:

- Enhanced Zoning
- Basic Zoning
- Basic zoning for active-only zoneset

Enhanced zoning is the recommended method for an all Cisco VSAN environment. This mode of zoning allows for full-distribution of the zoneset information across the fabric. This mode is operationally the most efficient.

The remaining two zone modes do not apply to the Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System design. To summarize Basic Zoning mode provides interoperability with non-Cisco SAN switches, while Basic zoning for the active-only zoneset mode is strictly for lab environments and is not recommended for a production SAN.

The best practice recommendation is to default zone deny. This configuration will deny communication between any devices on the SAN fabric. The SAN administrator can then explicitly permit all initiator to target communication. It is also recommended to backup the zones on a periodic basis.

On the MDS platform, zonesets are maintained as access control entries (ACEs) in the Ternary Content Addressable Memory (TCAM) hardware of the switch. This is a finite resource and the following best practices should be utilized to reduce ACE utilization and TCAM consumption:

- Limit large zones
- Multiple-Initiator zones should be avoided unless absolutely required.
- Single-initiator, multiple-target (SMT) zoning is typically recommended



Note

The Cisco MDS NX-OS 5.2(6) has introduced Smart Zoning. Smart Zoning offers operational simplicity and TCAM efficiencies by recognizing that most SAN zoning follows the SMT model. Smart Zoning allows one-to-many, many-to-one, or many-to-many relationships on a per VSAN basis and can also be disabled for certain zones within a VSAN. Cisco smart zoning simplifies zoning to reduce the burden on administrators and enable automation of SAN provisioning without affecting SAN hardware resource consumption. Smart Zoning is supported by DCNM and ideal for large SAN deployments.

For more information about Smart Zoning see:

http://www.cisco.com/en/US/docs/switches/datacenter/mds9000/sw/5_2/configuration/guides/fabric/nx-os/zone_ps5989_TSD_Products_Configuration_Guide_Chapter.html

VSP implements up to 1024 virtual ports for each physical storage port. This improves the security of the data on the VSP array by enabling LUNs to be assigned to virtual JBODs known as Host Storage Domains rather than directly to physical ports. As hosts connect to Virtual Ports rather than Physical Ports on the VSP they only detect the devices associated with their Host Storage Domain. This allows LUN numbers to be preserved when consolidating workloads and offers smaller management domains to improve operational efficiency. The [Appendix B—Hitachi Virtual Storage Platform Configurations](#) has a complete listing of the host storage domains used on VSP in the reference architecture.

Conclusion

Innovation is core to both Cisco and Hitachi. The introduction of Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System is a clear indicator of this and the level of collaboration between the two companies. This reference architecture offers a blueprint for creating a cloud-ready infrastructure integrating computing, networking, and storage resources. The platform is a safe, versatile, and cost-effective alternative to static-integrated approaches. Hitachi Unified Compute Platform Select for VMware vSphere with Cisco Unified Computing System provides a uniform platform allowing organizations to standardize their data center environments and deliver business value through better IT.

Appendix A—Bill of Materials

The following product information is provided for reference and will require modification depending on specific customer environments. Considerations include optics, cabling preferences, application workload and performance expectations.


Note

The Cisco UCS B-Series configurations below address performance-based customer requirements. The UCP Select for VMware vSphere with Cisco Unified Computing System architecture can support numerous Cisco UCS B-Series configurations that are not covered in this Cisco UCS Bill of Materials.

Table 21 Virtual Storage Platform Bill of Materials

Product	Product Description	Qty
VSP-SOLUTION.S	Hitachi Virtual Storage Platform	1
VSP-A0001.S	VSP Hardware Product	1
041-100065-01.P	Virtual Storage Platform Microcode Kit	1
041-100066-01.P	Virtual Storage Platform Product Documentation Library	1
043-991826-01.P	SVC VSP Installation Planning	1
043-991828-01.P	SVC VSP Installation - Control Frame	1
9960623.P	PDU Kit, 30A, 250V, 8XC13R,NEMA L6-30P	4
DKC-F710I-16UFC.P	Fibre 16-Port HOST Adapter(8Gbps)	5
DKC-F710I-APC.P	Additional Controller PS	1
DKC-F710I-BCH.P	Bezel Kit - DKC	1
DKC-F710I-BM64.P	Cache Flash MemoryModule (64GB)	1
DKC-F710I-BUC.P	DEV Cable from Controller to 1st Drive Chassis	1
DKC-F710I-BUH.P	Drive Chassis Bezel - HDS	1
DKC-F710I-C32G.P	Cache Memory Module (32GB)	3
DKC-F710I-DECO.P	Decoration Panel - HDS	1
DKC-F710I-ESW.P	PCI-Express Switch Adapter	1
DKC-F710I-FIHT.P	Filler Panel	15
DKC-F710I-PBR1.P	PDU Bracket	1
DKC-F710I-PHUC.P	DKU Power Cord Kit (USA)	1
DKC-F710I-PLUC.P	DKC Power Cord Kit (USA)	1
DKC-F710I-RK42.P	Rack - 42U	1
DKC-F710I-SBX.P	SFF Drive Chassis	1
043-991846-01.P	SVC VSP SBX 1Mo Yr-A Warranty Service	36
DKC-F710I-SCAP	Disk Adapter	1
DKC-F710I-SCOV.P	Side Cover	1
DKC710I-CBXA.P	Primary Controller Chassis	1
043-991837-01.P	SVC VSP CC 1Mo Yr-A Warranty Service	36
DTI4GL.P	4GB USB memory stick with lanyard	1
IP0662-14.P	LAN Cable 14ft	1
IP0665-11.P	RJ-11 Modular In-Line Coupler 4 Conductor	1
IP0665-45.P	RJ-45 Modular In-Line Coupler 6 Conductor	1
DKC-F710I-600JCM.P	SFF 600GB Disk Drive 2.5inch	120
VSP-SPARES.S	VSP Spare Drives	1
DKC-F710I-600JCM.P	SFF 600GB Disk Drive 2.5inch	1

Table 22 Cisco Unified Computing System Bill of Materials

Custom Name	SKU	Description	Qty
Blade Server UCSB-B200-M3	UCSB-B200-M3	UCS B200 M3 Blade Server w/ o CPU, memory, HDD, mLOM/ mezz	4
	UCS-CPU-E5-2690	2.90 GHz E5-2690/ 135 W 8C/ 20MB Cache/ DDR3 1600MHz	8
	UCS-MR-1X162RY-A	16GB DDR3-1600-MHz RDIMM/ PC3-12800/ dual rank/ 1.35v	64
	UCS-VIC-M82-8P	Cisco UCS VIC 1280 dual 40Gb capable Virtual Interface Card	4
	UCSB-MLOM-40G-01	Cisco UCS VIC 1240 modular LOM for M3 blade servers	4
	N20-BBLKD	UCS 2.5 inch HDD blanking panel	8
	UCSB-HS-01-EP	CPU Heat Sink for UCS B200 M3 and B420 M3	8
Blade Server B230-BASE-M2	B230-BASE-M2	UCS B230 M2 Blade Server w/ o CPU, memory, SSD, mezzanine	4
	UCS-CPU-E72870	2.4 GHz E7-2870 130W 10C/ 30M Cache	8
	UCS-MR-2X164RX-D	2X16GB NBS DDR3-1333-MHz RDIMM/ PC3-10600/ quad rank/ x4/ 1.35v	32
	UCS-VIC-M82-8P	Cisco UCS VIC 1280 dual 40Gb capable Virtual Interface Card	4
	N20-BBLKD-7MM	UCS B230 M2 and M1 7mm SSD blanking panel	8
	N20-BHTS6	CPU heat sink for UCS B230 Blade Server	8
	UCS-MKIT-164RX-D	Mem kit for UCS-MR-2X164RX-D	64
Chassis N20-C6508	N20-C6508	UCS 5108 Blade Svr AC Chassis/ 0 PSU/ 8 fans/ 0 fabric extender	2
	CAB-C19-CBN	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	8
	UCS-IOM-2208XP	UCS 2208XP I/O Module (8 External, 32 Internal 10Gb Ports)	4
	N01-UAC1	Single phase AC power module for UCS 5108	2
	N20-CAK	Access kit for 5108 Blade Chassis including Railkit, KVM dongle	2
	N20-FAN5	Fan module for UCS 5108	16
	N20-PAC5-2500W	2500W AC power supply unit for UCS 5108	8
	N20-FW010	UCS 5108 Blade Server Chassis FW package	2
Fabric Interconnect UCS-FI-6248UP	UCS-FI-6248UP	UCS 6248UP 1RU Fabric Int/ No PSU/ 32 UP/ 12p LIC	2
	UCS-ACC-6248UP	UCS 6248UP Chassis Accessory Kit	2
	UCS-PSU-6248UP-AC	UCS 6248UP Power Supply/ 100-240VAC	4

Table 23 Cisco Unified Computing System Bill of Materials

Name	Catalog Num	Description	Qty
MDS 9513	DS-C9513-4AK9.P	MDS 9513 Base Config: Chassis, 2 Sup-2A, 2 Fabric3, 2 6K PS	2
	DS-13SLT-FAB3.P	MDS 9513 Crossbar Switching Fabric3 Module	4
	DS-13SLT-FAN-F.P	MDS 9513 front installed fan tray	2
	DS-13SLT-FAN-R.P	MDS 9513 rear installed fan tray	2
	DS-CAC-6000W.P	MDS 9513 6000W AC power supply	4
	DS-X9530-SF2AK9.P	MDS 9500 Series Supervisor-2A	4
	CAB-C19-CBN.P	Cabinet Jumper Power Cord, 250 VAC 16A, C20-C19 Connectors	8
	DS-9513-KIT-HDS.P	MDS 9513 Accessory Kit for HDS	2
	DS-X9232-256K9.P	32-Port 8-Gbps Advanced Fibre Channel Switching Module	4
	DS-X9248-256K9.P	48-Port 8-Gbps Advanced Fibre Channel Switching Module	8
	M95S2K9-5.2.1.P	MDS 9500 Supervisor/ Fabric-2, NX-OS Software Release 5.2(1)	2
	DCNM-SAN-M95-K9.P	DCNM for SAN License for MDS 9500	2
	M9500ENT1K9.P	Enterprise package license for 1 MDS9500 switch	2
Nexus 5548UP	N5K-C5548UP-OSM.P	Nexus 5548UP Storage Solutions Bundle, Full Stor Serv Lic, OSM	2
	N55-48PO-SSK9.P	Nexus 5500 Storage License, 48 Ports, OSM	2
	N55-DL2.P	Nexus 5548 Layer 2 Daughter Card	2
	N55-M-BLNK.P	Nexus 5500 Module Blank Cover	2
	N55-PAC-750W.P	Nexus 5500 PS, 750W, Front to Back Airflow(Port-Side Outlet)	4
	N5548P-FAN.P	Nexus 5548P and 5548UP Fan Module, Front to Back Airflow	4
	CAB-C13-C14-2M.P	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	4
	N5548-ACC-KIT.P	Nexus 5548 Chassis Accessory Kit	2
	N5KUK9-521N1.1.P	Nexus 5000 Base OS Software Rel 5.2(1)N1(1)	2
	DCNM-SAN-N5K-K9=P	DCNM for SAN License for Nexus 5000	2

Table 24 Cisco Unified Computing System Bill of Materials

Name	SKU	Description	Qty
Nexus 1000v	L-N1K-VLCPU-16=	Nexus 1000V eDelivery CPU License Qty 16	1
	L-N1K-VLCPU-01	Nexus 1000V eDelivery CPU License 1-Pack	16

Appendix B—Hitachi Virtual Storage Platform Configurations

Table 25 *Hitachi Virtual Storage Platform Port Settings*

Fabric	ON
Security	Enabled
Speed	Auto(8Gbps)
Attribute	Target
Connection Type	Point-to-Point

Table 26 *Hitachi Virtual Storage Platform Zone Sets*

Chassis	Server	Server FC Port	Hitachi VSP Port
CH 1	S 1	FC 1	CL2-2A
			CL1-1C
		FC 0	CL2-2C
CH 1	S 2		CL1-1A
		FC 1	CL2-2E
			CL1-1G
CH 1	S 3	FC 0	CL2-2G
			CL1-1E
			CL1-1A
CH 1	S 4		CL2-2C
		FC 1	CL1-1C
			CL2-2A
CH 1	S 4	FC 0	CL1-1E
			CL2-2G
		FC 1	CL1-1G
CH 2	S 1		CL2-2E
		FC 1	CL2-2A
			CL1-1C
CH 2	S 2	FC 0	CL2-2C
			CL1-1A
		FC 1	CL2-2E
CH 2	S 3		CL1-1G
		FC 0	CL2-2G
			CL1-1E
CH 2	S 4	FC 1	CL2-2A
			CL1-1A
		FC 0	CL2-2C
CH 2	S 4	FC 1	CL1-1C
			CL2-2A
		FC 0	CL1-1E
CH 2	S 4		CL2-2G
		FC 1	CL1-1G
			CL2-2E

References

Deploy Hitachi Unified Compute Platform Select for VMware vSphere 5.1 with Cisco Unified Computing System Using Hitachi Unified Storage VM in a Scalable Environment:

<http://www.hds.com/assets/pdf/hitachi-ucp-select-for-vmware-vsphere-with-cisco-ucs.pdf>

Cisco Unified Computing System:

<http://www.cisco.com/en/US/products/ps10265/index.html>

Cisco UCS 6200 Series Fabric Interconnects:

<http://www.cisco.com/en/US/products/ps11544/index.html>

Cisco UCS 5100 Series Blade Server Chassis:

<http://www.cisco.com/en/US/products/ps10279/index.html>

Cisco UCS B-Series Blade Servers:

<http://www.cisco.com/en/US/partner/products/ps10280/index.html>

Cisco UCS Adapters:

http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

<http://www.cisco.com/en/US/products/ps10281/index.html>

Cisco Nexus 5000 Series Switches:

<http://www.cisco.com/en/US/products/ps9670/index.html>

Cisco Nexus 1000v:

<http://www.cisco.com/en/US/products/ps9902/index.html>

Cisco MDS 9500 Series Multilayer Directors:

<http://www.cisco.com/en/US/products/ps5990/index.html>

Cisco Prime Data Center Manager:

<http://www.cisco.com/en/US/products/ps9369/index.html>

Hitachi Virtual Storage Platform:

http://www.hds.com/products/storage-systems/hitachi-virtual-storage-platform.html?WT.ac=us_mg_pro_hvsp

VMware vCenter Server

<http://www.vmware.com/products/vcenter-server/overview.html>

VMware vSphere:

<http://www.vmware.com/products/datacenter-virtualization/vsphere/index.html>

Interoperability Matrixes:

- HDS and UCS/Nexus:
<http://www.hds.com/assets/pdf/all-storage-models-for-cisco-ucs.pdf>
- VMware and UCS/HDS:
<http://www.vmware.com/resources/compatibility>

Acknowledgments

Mike Meditz, Consulting System Engineer for HDS, Cisco Systems, Inc.

Mike Meditz is globally responsible for the technical relationship with Hitachi Data Systems. His primary functions include joint solution discovery, development, execution, and field enablement. Prior to Cisco, Mike was a Systems Engineer at SBC DataComm (now AT&T) focusing on Cisco data center and security solutions in the San Francisco Bay Area. Following this, Mike worked for a large government agency and was responsible for enterprise architectures, network security, and content delivery. Mike has a Bachelor of Science degree in Finance from Sonoma State University in Rohnert Park, California.

TJ Singh, Network Engineer Cisco Lab Services, Cisco Systems, Inc.

TJ Singh has over 7 years of experience with compute, network, storage physical and virtual topologies. Design, implement and manage Cisco Unified Computing System, Network, Storage infrastructure. TJ holds a bachelors degree with various career oriented certification network and virtualization.

HITACHI is a trademark or registered trademark of Hitachi, Ltd.