

Cisco and Hitachi Adaptive Solutions for Converged Infrastructure Design Guide

Last Updated: June 2019



About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to:

<http://www.cisco.com/go/designzone>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT 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 Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, 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)

© 2021 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction.....	6
Audience	6
Purpose of this Document.....	6
Solution Summary	6
Technology Overview	8
Cisco Unified Computing System	8
Cisco UCS Differentiators	8
Cisco Intersight	9
Cisco UCS Manager.....	10
Cisco UCS Fabric Interconnects	10
Cisco UCS Fabric Extenders.....	11
Cisco UCS 5108 Blade Server Chassis	12
Cisco UCS Virtual Interface Card	12
Cisco UCS B-Series Blade Servers	12
Cisco UCS C-Series Rack Servers	13
Cisco Nexus 9000 Series Switch.....	13
Cisco MDS 9000 Series Switch	14
Cisco Data Center Network Manager.....	15
Hitachi Virtual Storage Platform	15
Hitachi Virtual Storage Platform F1500 and G1500	15
Hitachi Virtual Storage Platform Fx00 Models and Gx00 Models.....	17
Storage Virtualization Operating System RF	17
VMware Support	18
Solution Design	20
Requirements	20
Physical Topology.....	20
Compute Connectivity.....	22
Network Connectivity	24
Storage Connectivity.....	25
End-to-End Data Path.....	29
Compute Design.....	31
Cisco UCS	31
VMware vSphere.....	32

Network Design	34
Nexus Configuration.....	34
Storage Design.....	34
Capacity Saving with Deduplication and Compression Options.....	35
LUN Multiplicity Per HBA and Different Pathing Options.....	36
Data Center Network Manager.....	36
Design Considerations.....	39
Cisco Best Practices	39
Hitachi Best Practices	39
Solution Validation	43
Validated Hardware	43
Summary	45
Appendix: Solution Design for Cisco and Hitachi Adaptive Solutions for CI as Direct Attached Storage (DAS)	46
Requirements	46
Physical Topology.....	46
Connectivity.....	48
End-to-End Data Path.....	49
Validated Hardware for Direct Attached Storage.....	50
Direct Attached Storage Deployment Guide.....	51
Appendix: Solution References.....	52
Compute.....	52
Network and Management	52
Storage.....	53
Virtualization Layer.....	53
Compatibility Matrixes.....	53
About the Authors.....	55
Acknowledgements	55

Executive Summary

Cisco Validated Designs consist 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.

Cisco and Hitachi are working together to deliver a converged infrastructure solution that helps enterprise businesses meet the challenges of today and position themselves for the future. Leveraging decades of industry expertise and superior technology, this Cisco CVD offers a resilient, agile, and flexible foundation for today's businesses. In addition, the Cisco and Hitachi partnership extends beyond a single solution, enabling businesses to benefit from their ambitious roadmap of evolving technologies such as advanced analytics, IoT, cloud, and edge capabilities. With Cisco and Hitachi, organizations can confidently take the next step in their modernization journey and prepare themselves to take advantage of new business opportunities enabled by innovative technology.

This document describes the Cisco and Hitachi Adaptive Solutions for Converged Infrastructure as a Virtual Server Infrastructure (VSI), which is a validated approach for deploying Cisco and Hitachi technologies as private cloud infrastructure. The recommended solution architecture consists of Cisco Unified Computing System (Cisco UCS), Cisco Nexus 9000 Series switches, Cisco MDS Fibre channel switches, and Hitachi Virtual Storage Platform (VSP). In addition, it includes validations of both VMware vSphere 6.5 and VMware vSphere 6.7 to support a larger range of customer deployments within vSphere and delivers a number of new features for optimizing storage utilization and facilitating private cloud common to these releases.

Solution Overview

Introduction

Modernizing your data center can be overwhelming, and it's vital to select a trusted technology partner with proven expertise. With Cisco and Hitachi as partners, companies can build for the future by enhancing systems of record, supporting systems of innovation, and growing their business. Organizations need an agile solution, free from operational inefficiencies, to deliver continuous data availability, meet SLAs, and prioritize innovation.

The Adaptive Solutions for CI as a Virtual Server Infrastructure (VSI) is a best practice datacenter architecture built on the collaboration of Hitachi and Cisco to meet the needs of enterprise customers utilizing virtual server workloads. This architecture is composed of the Hitachi VSP connecting through the Cisco MDS multilayer switches to Cisco UCS, and further enabled with the Cisco Nexus family of switches. Additionally, this design has been updated to include a Direct Attached Storage (DAS) model to support the direct VSP to FI configuration option.

This design is presented as a validated reference architecture, that covers specifics of products utilized within the Cisco validation lab, but the solution is considered relevant for equivalent supported components listed within Cisco and Hitachi published compatibility matrixes.

Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, storage administrators, technical architects, and customers who want to modernize their infrastructure to meet SLAs and their business needs at any scale.

Purpose of this Document

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 Adaptive Solutions for CI as a VSI for VMware vSphere utilizing Cisco Unified Computing and Hitachi VSP, along with Cisco Nexus and MDS switches as a validated approach for deploying Cisco and Hitachi technologies as an integrated compute stack.

Solution Summary

Adaptive Solutions for CI is a powerful and scalable architecture, leveraging the strengths of both Cisco and Hitachi. The Adaptive Solutions for CI data center implementation uses the following components:

- Cisco Unified Computing System
- Cisco Nexus Family of Switches
- Cisco MDS Family of Multilayer SAN Switches
- Hitachi Virtual Storage Platform

These products have been brought together as a validated reference architecture. The components are configured using the configuration and connectivity best practices from both companies to implement a reliable

and scalable VSI, validated for both vSphere 6.5 U2 and vSphere 6.7 U1. The specific products listed in this design guide and the accompanying deployment guide have gone through a battery of validation tests confirming functionality and resilience for the components as listed. Adjustments to the architecture are supported, provided they comply with the respective compatibility lists of both companies, and relevant product specific requirements of those changes are followed.

The documented example of the implementation of this design is here: [Deployment Guide for Cisco and Hitachi Converged Infrastructure with Cisco UCS Blade Servers, Cisco Nexus 9336C-FX2 Switches, Cisco MDS 9706 Fabric Switches, and Hitachi VSP G1500 and VSP G370 Storage Systems with vSphere 6.5 and vSphere 6.7](#)

For the Direct Attached Storage implementation, detailed in the appendix, the documented example can be found here: [Deployment Guide for Cisco and Hitachi Converged Infra-structure with Cisco UCS Blade Servers, Cisco Nexus 9336C-FX2 Switches, and Hitachi VSP G370 with vSphere 6.5 and vSphere 6.7 Connected as Direct Attached Storage](#)

Technology Overview

This section provides a technical overview of the compute, network, storage and management components in this solution. For additional information about any of the components covered in this section, see [Solution References](#).

Cisco Unified Computing System

Cisco UCS is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10–100 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform with a unified management domain for managing all resources.

Cisco Unified Computing System consists of the following subsystems:

- Compute – The compute piece of the system incorporates servers based on latest Intel’s x86 processors. Servers are available in blade and rack form factor, managed by Cisco UCS Manager.
- Network – The integrated network fabric in the system provides a low-latency, lossless, 10/25/40/100 Gbps Ethernet fabric. Networks for LAN, SAN and management access are consolidated within the fabric. The unified fabric uses the innovative Single Connect technology to lower costs by reducing the number of network adapters, switches, and cables. This in turn lowers the power and cooling needs of the system.
- Virtualization – The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtual environments to support evolving business needs.
- Storage access – Cisco UCS system provides consolidated access to both SAN storage and Network Attached Storage over the unified fabric. This provides customers with storage choices and investment protection. Also, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.
- Management: The system uniquely integrates compute, network and storage access subsystems, enabling it to be managed as a single entity through Cisco UCS Manager software. Cisco UCS Manager increases IT staff productivity by enabling storage, network, and server administrators to collaborate on Service Profiles that define the desired physical configurations and infrastructure policies for applications. Service Profiles increase business agility by enabling IT to automate and provision resources in minutes instead of days.

Cisco UCS Differentiators

Cisco UCS has revolutionized the way servers are managed in data center and provides a number of unique differentiators that are listed below:

- Embedded Management – Servers in Cisco UCS are managed by embedded software in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
- Unified Fabric – Cisco UCS uses a wire-once architecture, where a single Ethernet cable is used from the FI from the server chassis for LAN, SAN and management traffic. Adding compute capacity does not require additional connections. This converged I/O reduces overall capital and operational expenses.
- Auto Discovery – By simply inserting a blade server into the chassis or a rack server to the fabric interconnect, discovery of the compute resource occurs automatically without any intervention.

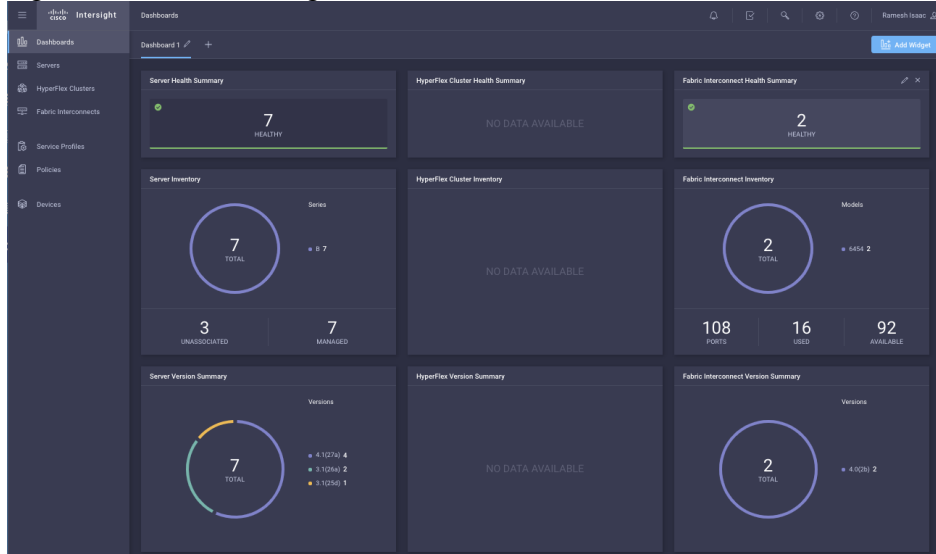
- Policy Based Resource Classification – Once a compute resource is discovered, it can be automatically classified to a resource pool based on policies defined which is particularly useful in cloud computing.
- Combined Rack and Blade Server Management – Cisco UCS Manager is hardware form factor agnostic and can manage both blade and rack servers under the same management domain.
- Model based Management Architecture – Cisco UCS Manager architecture and management database is model based, and data driven. An open XML API is provided to operate on the management model which enables easy and scalable integration of Cisco UCS Manager with other management systems.
- Policies, Pools, and Templates – The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Policy Resolution – In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools and templates can be defined at different levels of organization hierarchy.
- Service Profiles and Stateless Computing – A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- Built-in Multi-Tenancy Support – The combination of a profiles-based approach using policies, pools and templates and policy resolution with organizational hierarchy to manage compute resources makes Cisco UCS Manager inherently suitable for multi-tenant environments, in both private and public clouds.

Cisco Intersight

Cisco Intersight gives IT operations management to claimed devices across differing sites, presenting these devices within a unified dashboard. The adaptive management of Intersight provides visibility and alerts to firmware management, showing compliance across managed UCS domains, as well as proactive alerts for upgrade recommendations. Integration with Cisco TAC allows the automated generation and upload of tech support files from the customer.

Each Cisco UCS server or Cisco HyperFlex system automatically includes a Cisco Intersight Base edition at no additional cost when the customer accesses the Cisco Intersight portal and claims the device. In addition, customers can purchase the Cisco Intersight Essentials edition using the Cisco ordering tool.

A view of the unified dashboard provided by Intersight can be seen in [Figure 1](#).

Figure 1 Cisco Intersight Dashboard View

For more information about Cisco Intersight, see: <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/intersight/datasheet-c78-739433.html>

Cisco UCS Manager

Cisco UCS Manager (UCSM) provides unified, integrated management for all software and hardware components in Cisco UCS. Using [Cisco Single Connect](#) technology, it manages, controls, and administers multiple chassis for thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity through an intuitive GUI, a CLI, or a through a robust API.

Cisco UCS Manager is embedded into the Cisco UCS Fabric Interconnects and provides a unified management interface that integrates server, network, and storage. Cisco UCS Manager performs auto-discovery to detect inventory, manage, and provision system components that are added or changed. It offers comprehensive set of XML API for third party integration, exposes thousands of integration points and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

For more information about Cisco UCS Manager, see: <https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

Cisco UCS Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide a single point for connectivity and management for the entire Cisco UCS system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly-available management domain controlled by the Cisco UCS Manager. Cisco UCS FIs provide a single unified fabric for the system, with low-latency, lossless, cut-through switching that supports LAN, SAN and management traffic using a single set of cables.

The 4th generation (6400) Fabric Interconnect and the 3rd generation (6300) Fabric Interconnect deliver options for both high workload density, as well as high port count, with both supporting either Cisco UCS B-Series blade servers, or Cisco UCS C-Series rack-mount servers (Cisco UCS C-Series are not part of this validated design).

Figure 2 Cisco UCS 6332-16UP Fabric Interconnect Front Side View

Figure 3 Cisco UCS 6454 Fabric Interconnect Front Side View

The Fabric Interconnect models featured in this design are:

- Cisco UCS 6332-16UP Fabric Interconnect is a 1RU 40GbE/FCoE switch and 1/10 Gigabit Ethernet, FCoE and FC switch offering up to 2.24 Tbps throughput. The switch has 24x40Gbps fixed Ethernet/FCoE ports with unified ports providing 16x1/10Gbps Ethernet/FCoE or 4/8/16Gbps FC ports. This model is aimed at FC storage deployments requiring high performance 16Gbps FC connectivity to Cisco MDS switches or FC direct attached storage.
- Cisco UCS 6454 Fabric Interconnect is a 54 port 1/10/25/40/100GbE/FCoE switch, supporting 8/16/32Gbps FC ports and up to 3.82Tbps throughput. This model is aimed at higher port count environments that can be configured with 32Gbps FC connectivity to Cisco MDS switches or FC direct attached storage.

Table 1 provides a comparison of the port capabilities of the different Fabric Interconnect models.

Table 1 Cisco UCS 6200 and 6300 Series Fabric Interconnects

Features	6248	6296	6332	6332-16UP	6454
Max 10G ports	48	96	96* + 2**	72* + 16	48
Max 25G ports	-	-	-	-	48
Max 40G ports	-	-	32	24	6
Max 100G ports	-	-	-	-	6
Max unified ports	48	96	-	16	8
Max FC ports	48x 2/4/8G	96x 2/4/8G	-	16x 4/8/16G	8x 8/16/32G

* Using 40G to 4x10G breakout cables

** Requires QSA module

For more information about Cisco UCS Fabric Interconnects, see the following data sheets:

[Cisco UCS 6300 - https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-736682.html](https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-736682.html)

[Cisco UCS 6454 - https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html](https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html)

Cisco UCS Fabric Extenders

The Cisco UCS Fabric extenders (FEX) or I/O Modules (IOMs) multiplexes and forwards all traffic from servers in a blade server chassis to a pair of Cisco UCS Fabric Interconnects over a 10Gbps or 40Gbps unified fabric links. All traffic, including traffic between servers on the same chassis, or between virtual machines on the same server, is forwarded to the parent fabric interconnect where Cisco UCS Manager runs, managing the profiles and polices for the servers. FEX technology was developed by Cisco. Up to two FEXs can be deployed in a chassis.

For more information about the benefits of FEX, see: <http://www.cisco.com/c/en/us/solutions/data-center-virtualization/fabric-extender-technology-fex-technology/index.html>

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a fundamental building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server architecture. The Cisco UCS blade server chassis uses an innovative unified fabric with fabric-extender technology to lower TCO by reducing the number of NICs, HBAs, switches, and cables that need to be managed, cooled, and powered. It is a 6-RU chassis that can house up to 8 x half-width or 4 x full-width Cisco UCS B-series blade servers. A passive mid-plane provides up to 80Gbps of I/O bandwidth per server slot and up to 160Gbps for two slots (full-width). The rear of the chassis contains two I/O bays to house a pair of Cisco UCS 2000 Series Fabric Extenders to enable uplink connectivity to FIs for both redundancy and bandwidth aggregation.

Figure 4 Cisco UCS 5108 Blade Server Chassis



For more information about the Cisco UCS Blade Chassis, see:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-5100-series-blade-server-chassis/data_sheet_c78-526830.html.

Cisco UCS Virtual Interface Card

The Cisco UCS Virtual Interface Card (VIC) converges LAN and SAN traffic with one adapter using blade or rack servers from Cisco. The Cisco UCS VIC 1400 Series extends the network fabric directly to both servers and virtual machines so that a single connectivity mechanism can be used to connect both physical and virtual servers with the same level of visibility and control. Cisco VICs provide complete programmability of the Cisco UCS I/O infrastructure, with the number and type of I/O interfaces configurable on demand with a zero-touch model. The VIC presents virtual NICs (vNICs) as well as virtual HBAs (vHBAs) from the same adapter, provisioning from 1 to 256 Express PCIe devices within UCSM.

For more information about the Cisco UCS VIC 1440 see:

<https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.html>

Cisco UCS B-Series Blade Servers

Cisco UCS B-Series Blade Servers are based on Intel Xeon processors. They work with virtualized and non-virtualized applications to increase performance, energy efficiency, flexibility, and administrator productivity. The latest M5 B-Series blade server models come in two form factors; the half-width Cisco UCS B200 Blade Server and the full-width Cisco UCS B480 Blade Server. These M5 server uses the latest Intel Xeon Scalable processors with up to 28 cores per processor. The Cisco UCS B200 M5 blade server supports 2 sockets, 3TB of RAM (using 24 x128GB DIMMs), 2 drives (SSD, HDD or NVMe), 2 GPUs and 80Gbps of total I/O to each server. The Cisco UCS B480 blade is a 4 socket system offering 6TB of memory, 4 drives, 4 GPUs and 160 Gb aggregate I/O bandwidth.

Figure 5 Cisco UCS B200 M5**Figure 6 Cisco UCS B480 M5**

Each supports the Cisco VIC 1400 series adapters to provide connectivity to the unified fabric.

For more information about Cisco UCS B-series servers, see:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/datasheet-c78-739296.html>

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack Servers deliver unified computing in an industry-standard form factor to reduce TCO and increase agility. Each server addresses varying workload challenges through a balance of processing, memory, I/O, and internal storage resources. The most recent M5 based C-Series rack mount models come in three main models; the Cisco UCS C220 1RU, the Cisco UCS C240 2RU, and the Cisco UCS C480 4RU chassis, with options within these models to allow for differing local drive types and GPUs.

For more information about Cisco UCS C-series servers, see:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/datasheet-listing.html>

Cisco Nexus 9000 Series Switch

The Cisco Nexus 9000 Series Switches offer both modular and fixed 10/40/100 Gigabit Ethernet switch configurations with scalability up to 60 Tbps of non-blocking performance with less than five-microsecond latency, wire speed VXLAN gateway, bridging, and routing support.

The Nexus featured in this design is the Nexus 9336C-FX2 implemented in NX-OS standalone mode (Figure 7).

Figure 7 Cisco Nexus 9336C-FX2 Front Side View

The Nexus 9336C-FX2 implements Cisco Cloud Scale ASICs, giving flexible, and high port density, intelligent buffering, along with in-built analytics and telemetry. Supporting either Cisco ACI or NX-OS, the Nexus delivers a powerful 40/100Gbps platform offering up to 7.2 Tbps of bandwidth in a compact 1RU TOR switch.

For more information about the Cisco Nexus 9000 product family, see:

<https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/datasheet-listing.html>

Cisco MDS 9000 Series Switch

The Cisco MDS 9000 family of multilayer switches give a diverse range of storage networking platforms, allowing you to build a highly scalable storage network with multiple layers of network and storage management intelligence. Fixed and modular models implement 2/4/8/10/16/32-32Gbps FC, 1/10/40*Gbps FCIP, 10/40GE FCoE, and up to 48 Tbps of switching bandwidth.



* Refers to 40GE FCIP which is a future update.

The Cisco MDS 9132T Fibre Channel Fabric Switch is a 32-Gbps 32-Port Fibre Channel Switch provides Fibre Channel connectivity from the server rack to the SAN core. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the dual benefits of greater bandwidth and consolidation. Small-scale SAN architectures can be built from the foundation using this low-cost, low-power, non-blocking, line-rate, and low-latency, bi-directional airflow capable, fixed standalone SAN switch connecting both storage and host ports. This switch also offers state-of-the-art SAN analytics and telemetry capabilities that have been built into this next-generation hardware platform using the built-in dedicated Network Processing Unit designed to complete analytics calculations in real time. The switch offers fully automated zoning feature called AutoZone. AutoZone can help create zones between each Initiator and Target automatically. Fully populated MDS 9132T can also provide higher number of buffer to buffer credits, 4/8/16/32G FC connectivity, reversible airflow options, along with redundant PSUs and Fan Trays.

Figure 8 Cisco MDS 9132T with 16-Port Fibre Channel Expansion Module



The Cisco MDS 9706 Multilayer Director is featured in this design as one of the options to use within an Adaptive Solutions for CI placement ([Figure 9](#)).

Figure 9 Cisco MDS 9706 Multilayer Director Front Side View



This six-slot switch presents a modular, redundant supervisor design, giving FC and FCoE line card modules, SAN extension capabilities, as well as NVMe over FC support on all ports. The MDS 9706 offers a lower TCO through SAN consolidation, high availability, traffic management, and packet level visibility using SAN analytics, along with management and monitoring capabilities available through Cisco Data Center Network Manager (DCNM).

For more information about the MDS 9000 product family, see: <https://www.cisco.com/c/en/us/products/storage-networking/mds-9700-series-multilayer-directors/datasheet-listing.html>

Cisco Data Center Network Manager

Cisco Data Center Network Manager (DCNM) provides configuration and operations management for both the MDS, UCS and Nexus lines of switches. Cisco Data Center Network Manager (Cisco DCNM) automates the infrastructure of Cisco Nexus 5000, 6000, 7000, and 9000 Series Switches and Cisco MDS 9000 Series switches for data center management. Cisco DCNM enables you to manage many devices while providing ready-to-use control, management, and automation capabilities, along with SAN Analytics and automation. It is optimized for large deployments with little overhead, but traditional deployments are supported as well for implementations. Fabric deployments can be customized by the user to meet business needs, delivering a range of features including SAN Telemetry, end-to-end topology views, and simplified zoning.

For more information about DCNM, see <https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/prime-data-center-network-manager/datasheet-c78-740978.html>

Hitachi Virtual Storage Platform

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. With the industry's only 100% uptime warranty, Virtual Storage Platform delivers the highest uptime and flexibility for your block-level storage needs.

Hitachi Virtual Storage Platform F1500 and G1500

The all-flash Hitachi Virtual Storage Platform F1500 (VSP F1500) and the Hitachi Virtual Storage Platform G1500 (VSP G1500) storage systems provide high performance, high availability, and reliability for always-on, enterprise-class data centers. VSP F1500 and VSP G1500 deliver guaranteed data availability and feature the industry's most comprehensive suite of local and remote data protection capabilities, including true active-active metro-clustering. The storage systems scale to meet the demands of IT organizations' ever-increasing workloads. When combined with server virtualization, the mission-critical storage virtualization supports a new breed of applications at cloud scale while reducing complexity.


Key features of the Hitachi Virtual Storage Platform F1500 and G1500 are:

- High-speed VSDs - VSP F1500 and G1500 are equipped with new virtual storage directors (VSDs) that use the latest generation of Intel Xeon 2.3-GHz 8-core microprocessor to efficiently manage the front-end and back-end directors, PCI-Express interface, shared memory, and service processor.
- All-flash VSP F1500 - The VSP F1500 all-flash storage system is configured exclusively with the latest generation of flash module drives (FMDs) to provide performance optimized for intense I/O operations. Designed for flash-first, high-performance workloads and leveraging Hitachi SVOS RF-based deduplication and compression, VSP F1500 offers up to five times greater ROI. Accelerated flash architecture delivers consistent, low-latency IOPS at scale. Adaptive flash management distributes writes and rebalances load over time, and FMDs from Hitachi deliver enterprise performance with superior functionality and greater cost value.

- Global Storage Virtualization - Enables an always-on infrastructure with enterprise-wide scalability that provides a complete separation between host and storage. The scalability is independent of connectivity, location, storage system, and vendor. Data center replication support allows provisioning and management of virtual storage machines up to 100 meters apart.
- Integrated Active Mirroring - Enables volume extensibility between systems and across sites through the provisioning and management of active-active volumes up to 100 km apart. Combined with remote data center replication, this mirroring is an ideal solution for critical applications that require zero recovery point and recovery time objectives. Active mirroring is enabled by the global-active device feature.
- Capacity Saving Functions - Data deduplication and data compression enable you to reduce the bitcost for stored data by deduplicating and compressing the data stored on internal flash drives. Data deduplication and compression are performed by the controllers of the storage system.
- Accelerated Compression - Enables you to reduce your bitcost for the stored data by utilizing the compression function in the flash module compression (FMC) drives. Accelerated compression allows you to assign FMC capacity to a pool that is larger than the physical capacity of the FMC parity groups. The data access performance of the storage system is maintained when the accelerated compression function is used, as the compression engine is offloaded to the FMC drives.
- Unified Storage with Enterprise Scalability - Allows you to centrally manage multi-vendor storage resources across all virtualized internal and external storage pools, whether deployed for SAN or object storage.
- Centralized Storage Management - Software such as Hitachi Storage Advisor simplifies administrative operations and streamlines basic management tasks.
- Hitachi Accelerated Flash - Offers a patented data center-class design and rack-optimized form factor that delivers up to 8.1 PB of HAF per system. It supports a sustained performance of 100,000 8K I/O operations per second per device, with fast and consistent response time.
- Server Virtualization - Integration with leading virtual server platforms provides end-to-end visibility, from an individual virtual machine to the storage logical unit, protecting large-scale multivendor environments.
- Customer-driven Nondisruptive Migration Capability - Enables movement, copy, and migration of data between storage systems, including non-Hitachi storage systems, without interrupting application access and local and remote copy relationships.

The VSP F1500 all-flash storage system is equipped with advanced high-density FMDs, providing up to 4M IOPS and 40 PB of capacity for multi-workload consolidation. The following illustration shows a fully configured VSP G1500 with two controllers containing 8 virtual storage director (VSD) pairs (128 CPU cores), 2 TB cache, and 12 16U drive chassis containing 2,304 SFF drives with a physical storage capacity of approximately 2.7 PB.

Figure 10 Hitachi VSP G1500 Front Side View



Hitachi Virtual Storage Platform Fx00 Models and Gx00 Models

Based on Hitachi's comprehensive suite of storage technology, the all-flash VSP F350, F370, F700, F900 and the VSP G350, G370, G700, and G900 include a range of versatile, high-performance storage systems that deliver flash-accelerated scalability, simplified management, and advanced data protection.

Hitachi Accelerated Flash (HAF) storage delivers best-in-class performance and efficiency in the Hitachi VSP F700, F900, G700, and G900 storage systems. HAF features patented flash module drives (FMDs) that are rack-optimized with a highly dense design that delivers greater than 338 TB effective capacity per 2U tray based on a typical 2:1 compression ratio. IOPS performance yields up to five times better results than that of enterprise solid-state drives (SSDs), resulting in leading performance, lowest bit cost, highest capacity, and extended endurance.

HAF integrated with Storage Virtualization Operating System (SVOS) RF running on VSP F700, F900, G700, and G900 enables transactions executed within sub-millisecond response even at petabyte scale.

For more information about Hitachi Virtual Storage Platform F Series, see: <https://www.hitachivantara.com/en-us/products/storage/virtual-storage-platform-f-series.html>

For more information about Hitachi Virtual Storage Platform G Series, see: <https://www.hitachivantara.com/en-us/products/storage/virtual-storage-platform-g-series.html>

Storage Virtualization Operating System RF

The Hitachi Storage Virtualization Operating System (SVOS) RF abstracts information from storage systems, virtualizes and pools available storage resources, and automates key data management functions such as configuration, mobility, optimization, and protection. This unified virtual environment enables you to maximize the utilization and capabilities of your storage resources while at the same time reducing operations overhead and risk. Standards-compatible for easy integration into IT environments, storage virtualization and management capabilities provide the utmost agility and control, helping you build infrastructures that are continuously available, automated,

and agile. Hitachi VSP F350, F370, F700, F900, F1500, G350, G370, G700, G900, and G1500 all utilize SVOS RF as the operating system. This provides the same features and functionality across storage systems regardless of the size of the workloads or storage system.

SVOS RF is the latest version of SVOS. Flash performance is optimized with a patented flash-aware I/O stack, which accelerates data access. Adaptive inline data reduction increases storage efficiency while enabling a balance of data efficiency and application performance. Industry-leading storage virtualization allows SVOS RF to use third-party all-flash and hybrid systems as storage capacity, consolidating resources for a higher ROI and providing a high-speed front-end to slower, less predictable systems.

SVOS RF provides the foundation for superior storage performance, high availability, and IT efficiency. The enterprise-grade capabilities in SVOS RF include centralized management across storage systems and advanced storage features, such as active-active data centers and online migration between storage systems without user or workload disruption. Features of SVOS RF include:

- Advanced efficiency providing user-selectable data reduction
- External storage virtualization
- Thin provisioning and automated tiering
- Flash performance acceleration
- Deduplication and compression of data stored on internal flash drives
- Storage service-level controls
- Data-at-rest encryption
- Performance instrumentation across multiple storage platforms
- Centralized storage management
- Standards-based application program interfaces (REST APIs)

For more information on Hitachi Storage Virtualization Operating System RF, see:

<https://www.hitachivantara.com/en-us/products/storage/storage-virtualization-operating-system.html>

VMware Support

Hitachi VSP F series and G series are aligned with the VMware software-defined storage vision, providing the following support:

- vSphere Metro Storage Cluster (vMSC): Using Hitachi Global-active device (GAD), you can create and maintain synchronous, remote copies of data volumes. A virtual storage machine is configured in the primary and secondary storage systems using the actual information of the primary storage system, and the global-active device primary and secondary volumes are assigned the same virtual LDEV number in the virtual storage machine. This enables the host to see the pair volumes as a single volume on a single storage system, and both volumes receive the same data from the host. Configuring GAD as the backend storage for a vSphere Metro Storage Cluster provides an ideal solution for maximizing availability and uptime by clustering physical data centers that reside within metro distances of each other.
- Hitachi Storage Provider for VMware vCenter: Hitachi Storage Provider works with VMware vSphere API for Storage Awareness (VASA) to provide access to Hitachi VSP F series and G series. Storage Provider enables policy-based storage management using VMware Storage Policy-based Management (SPBM) and

VMware Virtual Volumes (VVol). Storage Provider provides a simplified method for VMware admins and storage admins to deliver effective storage that meets advanced VM requirements.

- vSphere Storage APIs - Array Integration (VAAI): VAAI enables multiple storage functions (primitives) within vSphere to be offloaded to certified storage hardware. This reduces ESXi processor utilization by allowing certified storage hardware to perform these functions on the storage systems themselves. In many cases, VAAI accelerates these functions allowing them to complete in less time as compared to performing the functions within software on the ESXi host.
- Hitachi Storage Replication Adapter (SRA): Hitachi Storage Replication Adapter (SRA) for VMware Site Recovery Manager provides a disaster recovery (DR) solution that works with both your storage environment and your VMware environment. Arrays at both sites are "paired" during Site Recovery Manager configuration, and VMware administrators use the Site Recovery Manager application to manage the configuration and definition of the DR plan.
- vStorage API for Multipathing (VAMP): Hitachi VSP F series and G series support VAMP to provide enhanced control of I/O path selection and failover.
- vStorage API for Data Protection (VADP): Hitachi VSP F series and G series support VADP to enable backup applications to perform file-level or VM-level backup of running virtual machines.

Solution Design

The Adaptive Solutions for CI Solution Design implements a Virtual Server Infrastructure built to be powerful, scalable, and reliable, using the best practices of both Cisco and Hitachi. This section explains the architecture about how it was built, as well as the design options used within the solution.

Requirements

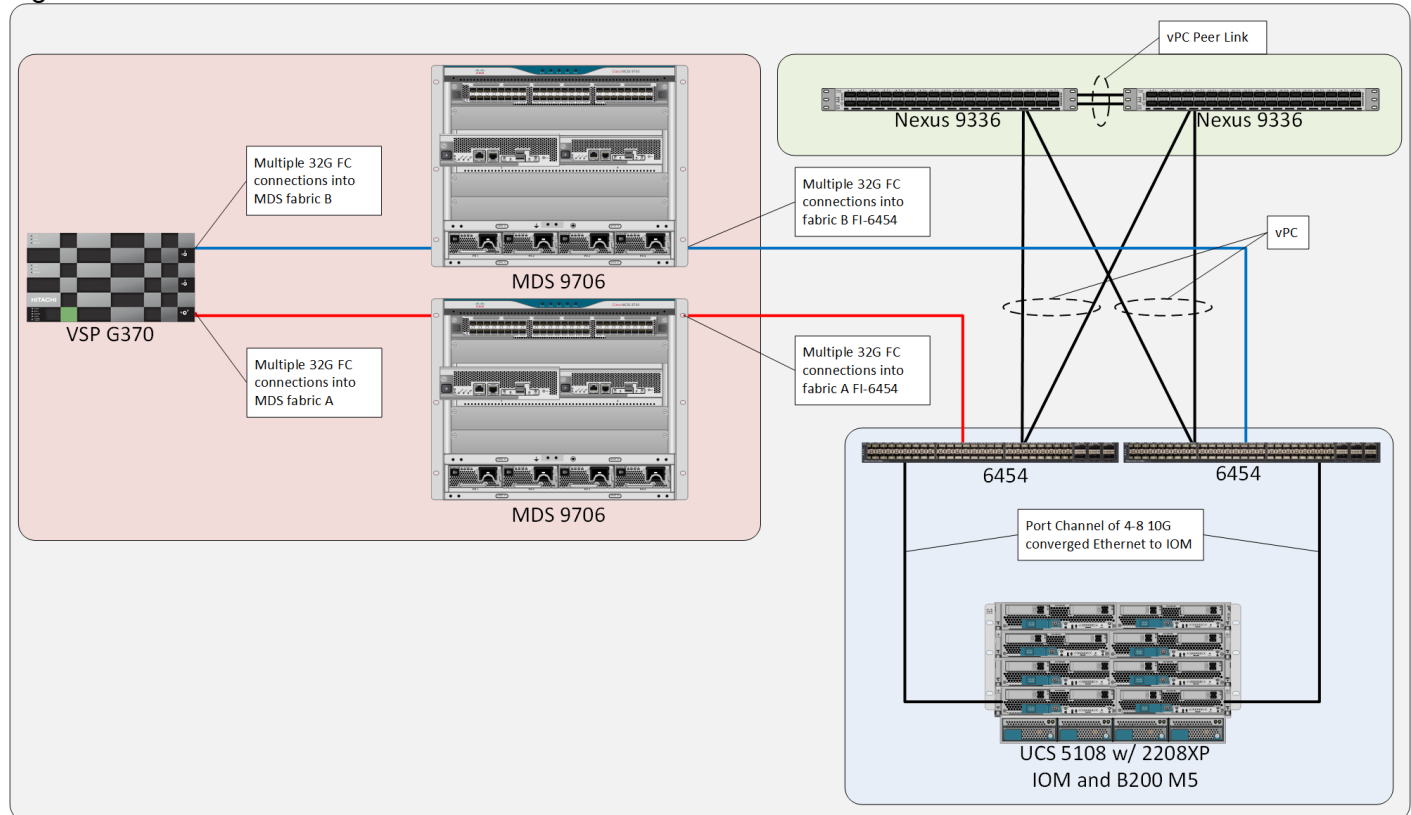
The Adaptive Solutions for CI datacenter is intended to provide a Virtual Server Infrastructure that addresses the primary needs of hosting virtual machines. This design assumes existing management infrastructure and routing have been pre-configured. These existing items include, but may not be limited to:

- An Out of Band management network
- A terminal server for console access
- An Active Directory/DNS Server
- Layer 3 connectivity to the Internet and any other adjacent enterprise networks
- Additional management components used for deployment

Physical Topology

The design is presented in two similar architectures, validated by the same Cisco Nexus and MDS switching. The first features the Cisco UCS fourth generation 6454 Fabric Interconnect paired with the Hitachi VSP G370. This pairing gives the option of 32G connectivity at each port coming from the FI going through to the VSP controller. The Cisco UCS 6454 and the VSP G370 present a good combination for a branch office type placement and is shown in [Figure 11](#).

Figure 11 Cisco and Hitachi Data Center with VSP G370 and Cisco UCS 6454



The components of this integrated architecture shown in [Figure 11](#) are:

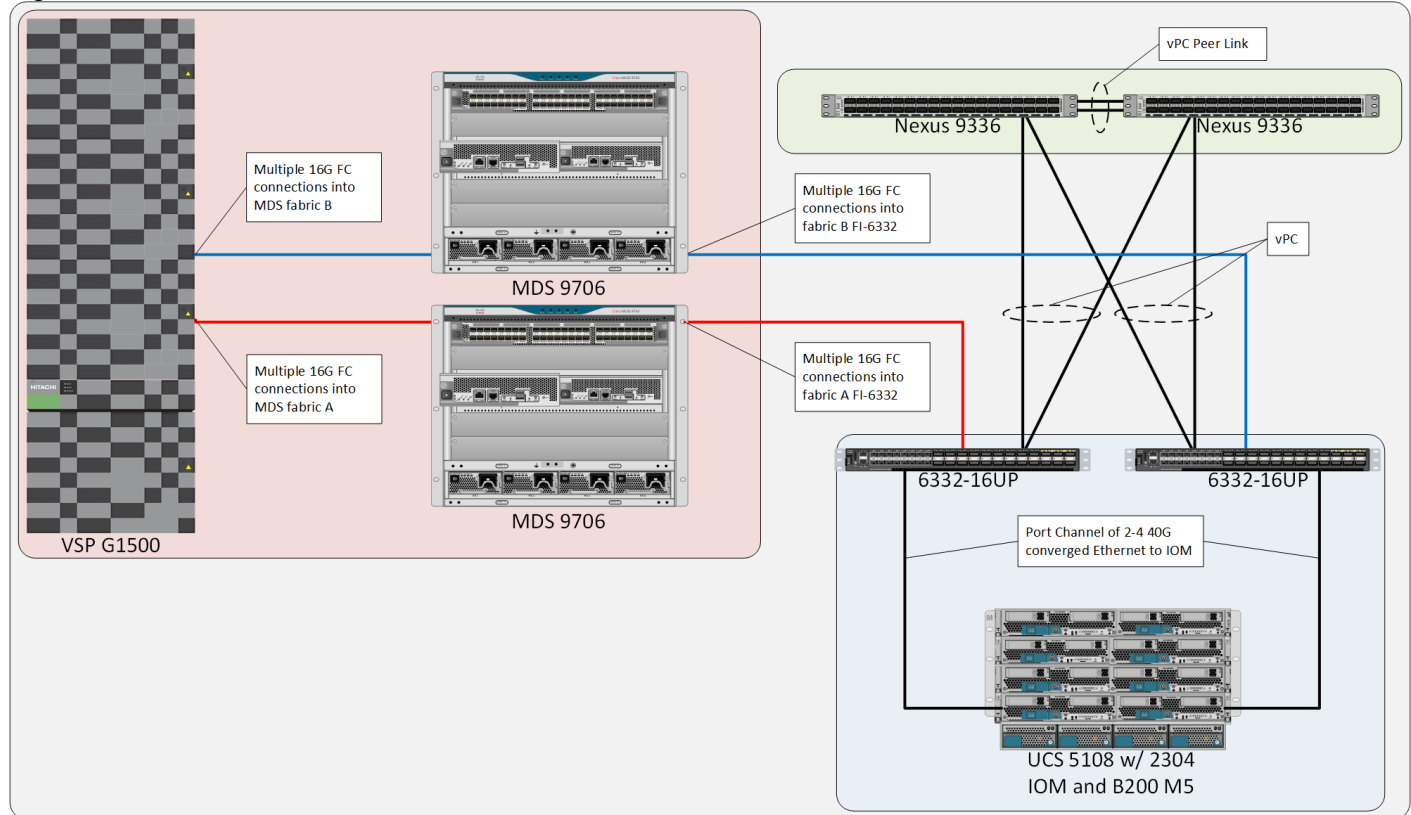
- Cisco Nexus 9336C-FX2 – 100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6454 Fabric Interconnect – Unified management of UCS compute, and the compute’s access to storage and networks.
- Cisco UCS B200 M5 – High powered, versatile blade server, conceived for virtual computing.
- Cisco MDS 9706 – 32Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- Hitachi VSP G370 – Mid-range, high performance storage system with optional all-flash configuration

Management components of the architecture additionally include:

- Cisco UCS Manager – Management delivered through the Fabric Interconnect, providing stateless compute, and policy driven implementation of the servers managed by it.
- Cisco Intersight (optional) – Comprehensive unified visibility across UCS domains, along with proactive alerts and enablement of expedited Cisco TAC communications.
- Cisco Data Center Network Manager – Multi-layer network configuration and monitoring.

The second architecture features the Cisco UCS third generation 6332-16UP Fabric Interconnect paired with the Hitachi VSP G1500. This pairing gives massively scalable storage capacity, along with high bandwidth 40G converged Ethernet connections to the Cisco UCS chassis. Putting this together produces an ideal converged infrastructure for a large main office placement, as shown in [Figure 12](#).

Figure 12 Cisco and Hitachi Data Center with VSP G1500 and Cisco UCS 6332-16UP



The components of this integrated architecture shown in Figure 12 are:

- Cisco Nexus 9336C-FX2 – 100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6332-16UP Fabric Interconnect – Unified management of UCS compute, and the compute’s access to storage and networks.
- Cisco UCS B200 M5 – High powered, versatile blade server, conceived for virtual computing.
- Cisco MDS 9706 – 32Gbps Fibre Channel connectivity within the architecture, as well as interfacing to resources present in an existing data center.
- Hitachi VSP G1500 – Enterprise-level, high performance storage system with optional all-flash configuration

Management components of the architecture include:

- Cisco UCS Manager – Management delivered through the Fabric Interconnect, providing stateless compute, and policy driven implementation of the servers managed by it.
- Cisco Intersight (optional) – Comprehensive unified visibility across UCS domains, along with proactive alerts and enablement of expedited Cisco TAC communications.
- Cisco Data Center Network Manager (optional) – Multi-layer network configuration and monitoring.


Compute Connectivity

Each compute chassis in the design is connected to the managing fabric interconnect with at least two ports per IOM. Ethernet traffic from the upstream network and Fibre Channel frames coming from the VSP are converged

within the fabric interconnect to be both Ethernet and Fibre Channel over Ethernet transmitted to the Cisco UCS servers through the IOM. These IOM connections from the Cisco UCS Fabric Interconnects to the IOMs are automatically configured as port channels with the specification of a Chassis/FEX Discovery Policy within UCSM.

These connections from the Cisco UCS 6454 Fabric Interconnect to the 2208XP IOM hosted within the chassis are shown in [Figure 13](#).


Figure 13 Compute Connectivity for Cisco UCS 6454 FI to Cisco UCS 2208XP IOM



The 2208XP IOM are shown with 4x10Gbps ports to deliver an aggregate of 80Gbps to the chassis, full population of the 2208XP IOM can support 8x10Gbps ports, allowing for an aggregate of 160Gbps to the chassis.

The equivalent connections for the Cisco UCS 6332-16UP Fabric Interconnect to the 2304 IOM are shown in [Figure 14](#).

Figure 14 Compute Connectivity for Cisco UCS 6332-16UP FI to Cisco UCS 5108 with 2304 IOM



The 2304 IOM are shown with 2x40Gbps ports to deliver an aggregate of 160Gbps to the chassis, full population of the 2304 IOM can support 4x40Gbps ports, allowing for an aggregate of 320Gbps to the chassis.


Network Connectivity

The network connection to each of the fabric interconnects is implemented as vPC from the upstream Nexus switches. In the switching environment, the vPC provides the following benefits:

- Allows a single device to use a Port Channel across two upstream devices
- Eliminates Spanning Tree Protocol blocked ports and use all available uplink bandwidth
- Provides a loop-free topology
- Provides fast convergence if either one of the physical links or a device fails
- Helps ensure high availability of the network


The upstream network switches which connect to the Cisco UCS 6454 Fabric Interconnects can utilize 10/25/40/100G port speeds. In this design, the 100G ports available from the 40/100G configurable ports (49-54) were used to construct the port channels on the fabric interconnect talking to the vPCs presented by the pair of upstream switches (Figure 15).

Figure 15 Network Connectivity from the Nexus 9336C-FX2 to the Cisco UCS 6454 FI



The network connections used by the Cisco UCS 6332-16UP Fabric Interconnects are configured in a similar manner but have options of 10/40G ports to talk to the upstream Nexus switch. In this design, the 40G ports were used for the construction of the port channels on the fabric interconnect that connected to the vPCs presented by the pair of upstream switches (Figure 16).

Figure 16 Network Connectivity from the Nexus 9336C-FX2 to the Cisco UCS 6332-16UP FI



Storage Connectivity

Both VSP platforms are connected through the Cisco MDS 9706 to the respective fabric interconnects they are associated with. For the fabric interconnects, these are configured as SAN Port Channels, with N_Port ID Virtualization (NPIV) enabled on the MDS. This configuration allows:

- Increased aggregate bandwidth between the fabric interconnects and the MDS
- Load balancing between the links using SID, DID, OXID
- Higher bandwidth and link redundancy in the event of a link failure

Figure 17 illustrates the connectivity of the VSP G370 storage system to the Cisco UCS 6454 Fabric Interconnects.

Figure 17 Storage Connectivity for the VSP G370 and Cisco UCS 6454 Connecting through the MDS 9706


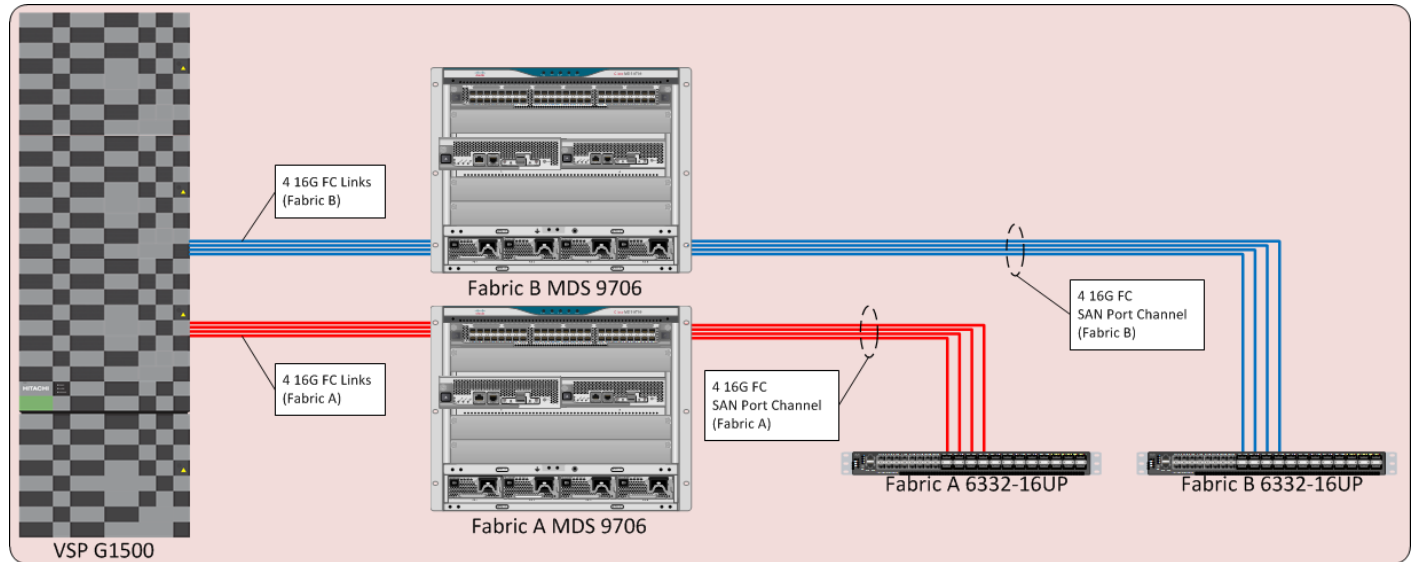


Figure 18 illustrates the connectivity of the VSP G1500 storage system to the Cisco UCS 6332-16UP Fabric Interconnects.

Figure 18 Storage Connectivity for the VSP G1500 and Cisco UCS 6332-16UP Connecting through the MDS 9706




Hitachi VSP FC Port to Fabric Assignments

Each VSP storage system is comprised of multiple controllers and fibre channel adapters that control connectivity to the fibre channel fabrics. Channel board adapters (CHAs) are used within the VSP F1500 and G1500, and channel boards (CHBs) are used within the VSP Fx00 models and Gx00 models. The multiple CHA/CHBs within each storage system allow for designing multiple layers of redundancy within the storage architecture, increasing availability and maintaining performance during a failure event. The VSP F350, F370, F700, F900, F1500, G350, G370, G700, G900, and G1500 CHA/CHBs each contain up to four individual fibre channel ports, allowing for redundant connections to each fabric in the Cisco UCS infrastructure.

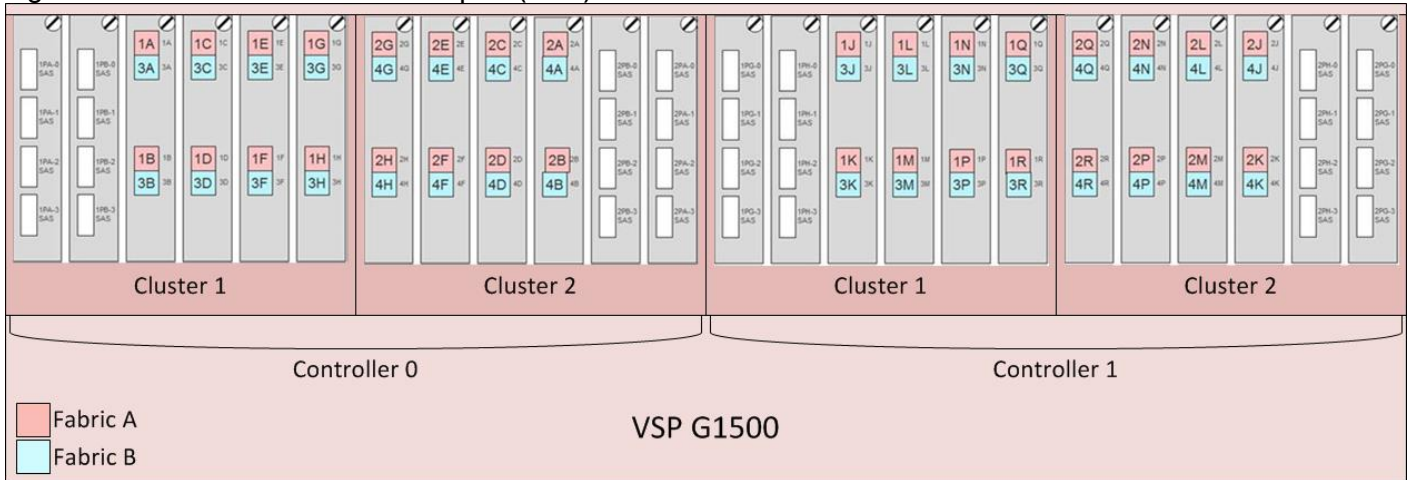
Hitachi VSP Fx00 models and Gx00 models have two controllers contained within the storage system. The port to fabric assignments for the VSP G370 used in this design are shown in Figure 19, illustrating multiple connections to each fabric and split evenly between VSP controllers and 32Gb CHBs:

Figure 19 VSP G370 Channel Boards (CHB) Ports



Hitachi VSP F1500 and G1500 have two controllers contained within the storage system, each of which contains two cluster controllers. The port to fabric assignments for the VSP G1500 used in this design are shown in Figure 20, illustrating multiple connections to each fabric and split evenly between VSP controllers, clusters, and 16Gb CHAs.

Figure 20 VSP G1500 Channel Adapter (CHA) Ports



Hitachi VSP LUN Presentation and Path Assignments

Due to Cisco UCS' ability to provide alternate paths for the boot LUN on each fibre channel fabric, four paths to each boot LUN were assigned, comprised of two paths on each fabric. For LUNs used as VMFS volumes, redundant paths for each LUN considering controller and cluster failure were assigned, comprised of two to four paths per fabric depending on the VSP model used. Figure 21 illustrates the boot LUN and VMFS LUN pathing configuration for the VSP G370.

Figure 21 Logical View of LUN to VSP G370 Port Association

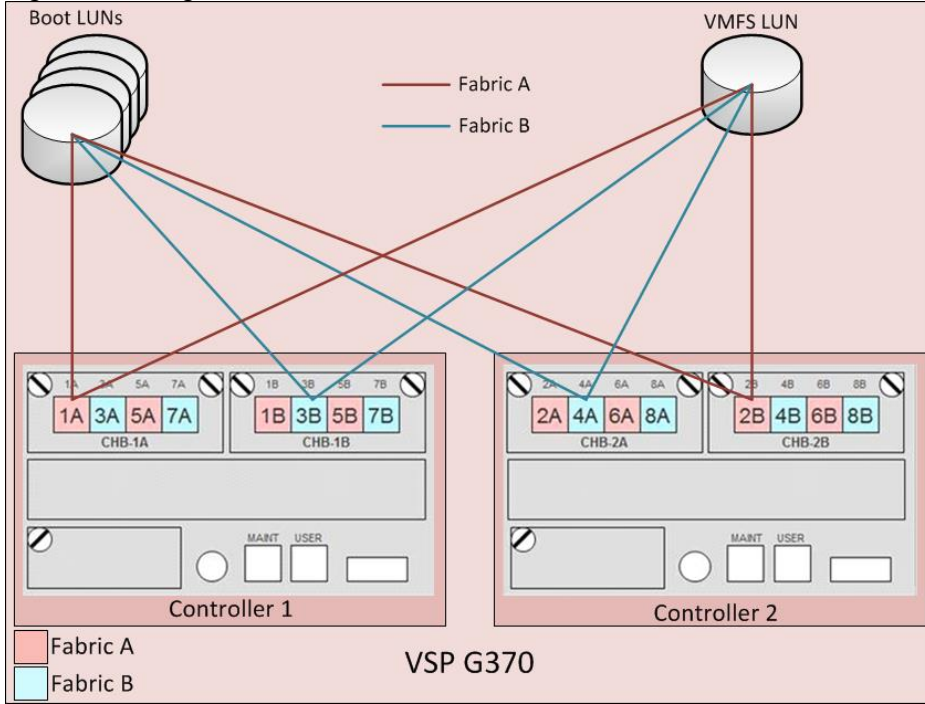
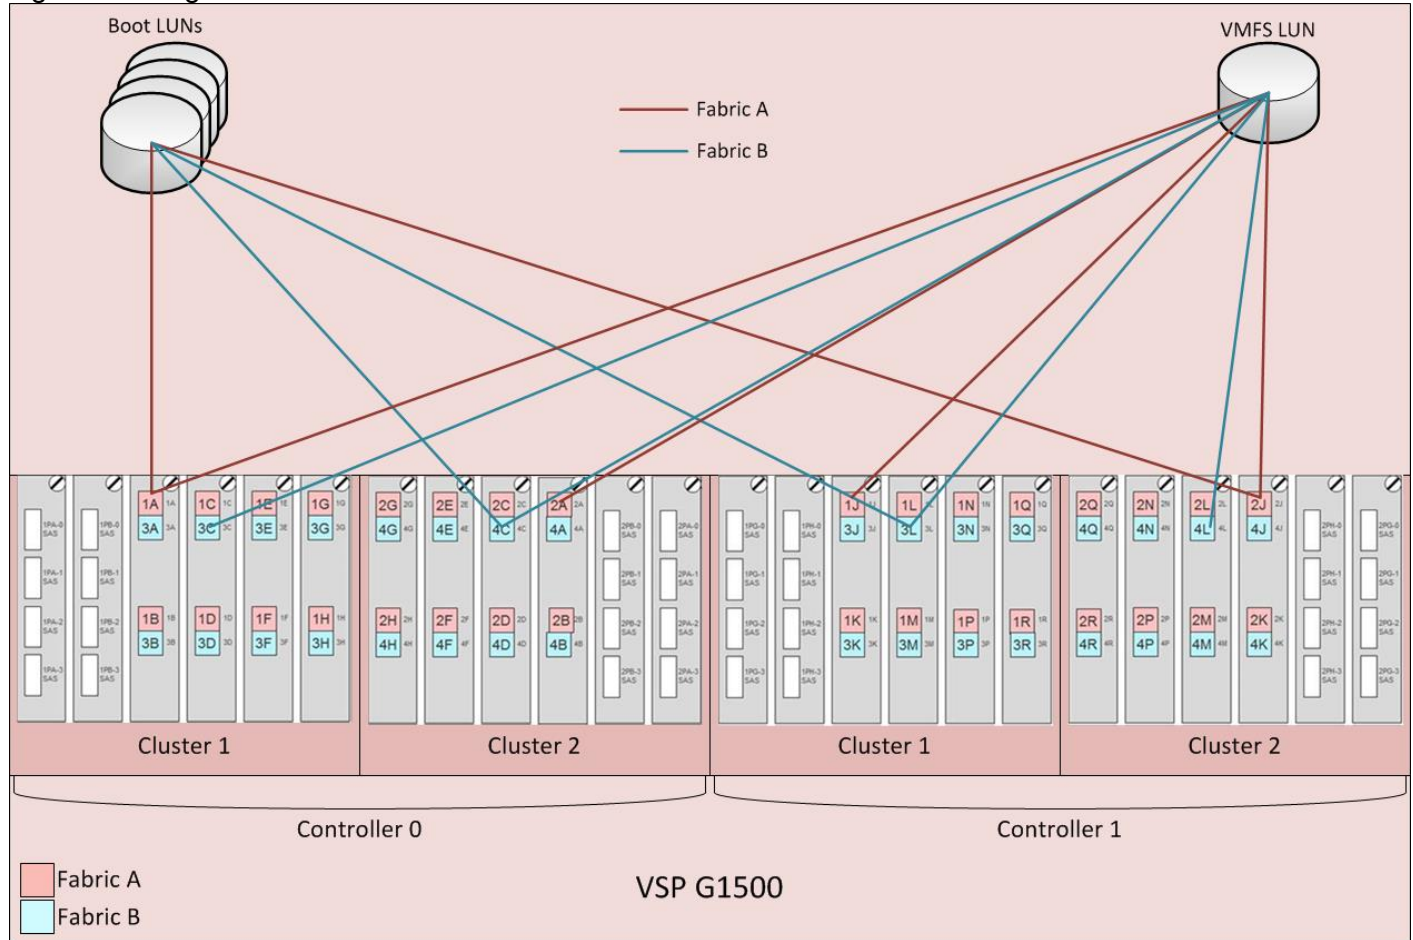


Figure 22 illustrates the boot LUN and VMFS LUN pathing configuration for the VSP G1500.

Figure 22 Logical View of LUN to VSP G1500 Port Association



MDS Zoning

Zoning within the MDS is configured for each host with single initiator multiple target zones, leveraging the Smart Zoning feature of the MDS for greater efficiency. The design implements a simple, single VSAN layout per fabric within the MDS, but configuration of differing VSANs for greater security and tenancy are supported.

Initiators (UCS hosts) and targets (VSP controller ports) are set up with device aliases within the MDS for easier identification within zoning and flogi connectivity. Configuration of zoning and the zonesets containing them can be managed with CLI but is also available for creation and editing with DCNM for a simpler administrative experience.

More information about zoning and the Smart Zoning feature is described in the [Storage Design](#) section.

End-to-End Data Path

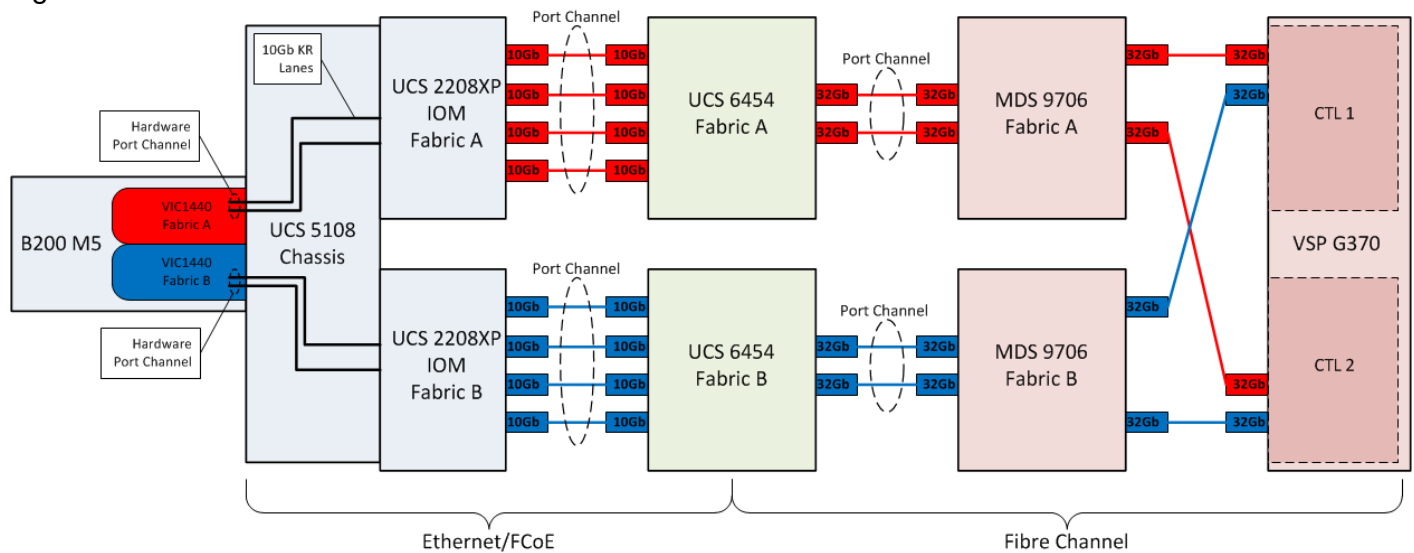
Both architectures in this design are built around the implementation of fibre channel storage, this is a high bandwidth solution within both options, but includes 32G end-to-end FC for the Cisco UCS 6454 to the VSP G370. Storage traffic flows from a Cisco UCS B200 blade in a UCS environment to Hitachi VSP G370 is as follows:

- The Cisco UCS B200 M5 server, equipped with a VIC 1440 adapter₍₁₎, connects to each fabric at a link speed of 20Gbps.

- Pathing through 10Gb KR lanes of the Cisco UCS 5108 Chassis backplane into the Cisco UCS 2208XP IOM (Fabric Extender).
- Connecting from each IOM to the Fabric Interconnect with pairs of 10Gb uplinks automatically configured as port channels during chassis association, that carry the FC frames as FCoE along with the Ethernet traffic coming from the chassis blades.
- Continuing from the Cisco UCS 6454 Fabric Interconnects into the Cisco MDS 9706 with multiple 32G FC ports configured as a port channel for increased aggregate bandwidth and link loss resiliency.
- Ending at the Hitachi VSP G370 fibre channel controller ports with dedicated F_Ports on the Cisco MDS 9706 for each N_Port WWPN of the VSP controller, with each fabric evenly split between the controllers and CHBs.

(1) The VIC 1440 will work with the UCS 6454 to provide 40G/40G when equipped with the 4th generation IOM which is not available at the time of this validation.

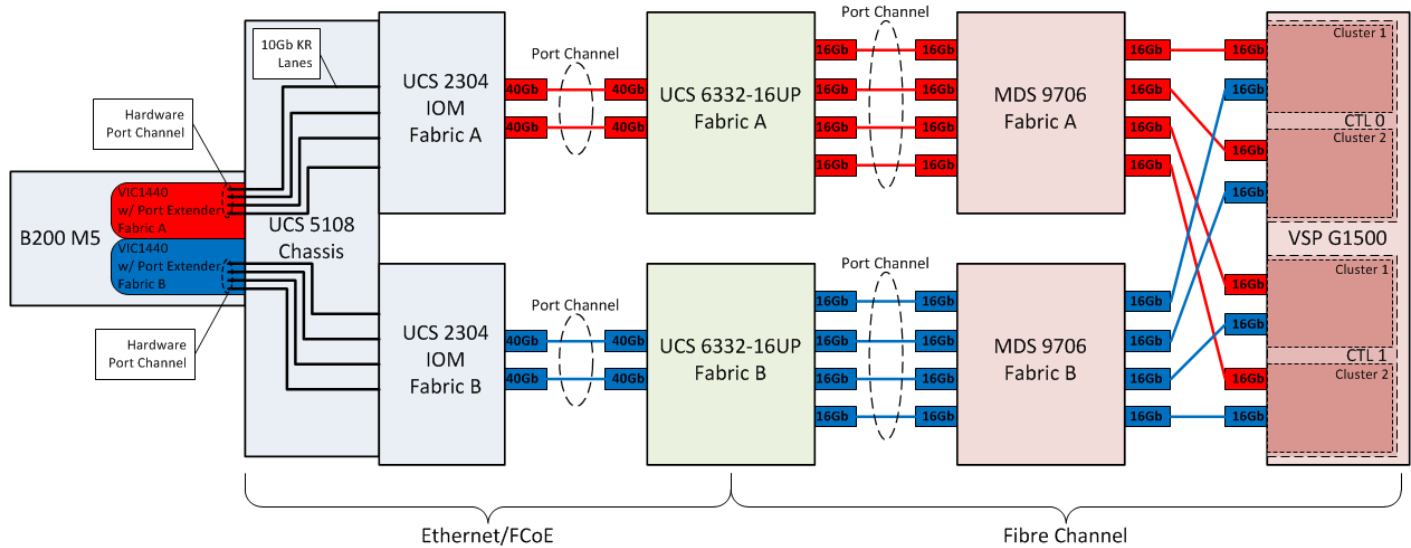
Figure 23 End-to-End Data Path from Cisco UCS B200 M5 to VSP G370



The equivalent Storage traffic flow from a Cisco UCS B200 blade in a UCS environment to Hitachi VSP G1500 is as follows:

- The Cisco UCS B200 M5 server, equipped with a VIC 1440 adapter and a Port Expander card, connects to each fabric at a link speed of 40Gbps.
- Pathing through 10Gb KR lanes of the Cisco UCS 5108 Chassis backplane into the Cisco UCS 2304 IOM (Fabric Extender).
- Connecting from each IOM to the Fabric Interconnect with pairs of 40Gb uplinks automatically configured as port channels during chassis association, that carry the FC frames as FCoE along with the Ethernet traffic coming from the chassis blades.
- Continuing from the Cisco UCS 6332-16UP Fabric Interconnects into the Cisco MDS 9706 with multiple 16G FC ports configured as a port channel for increased aggregate bandwidth and link loss resiliency.
- Ending at the Hitachi VSP G1500 fibre channel controller ports with dedicated F_Ports on the Cisco MDS 9706 for each N_Port WWPN of the VSP controller, with each fabric evenly split between the controllers, clusters, and CHAs.

Figure 24 End-to-End Data Path from Cisco UCS B200 M5 to VSP G1500



Compute Design

Cisco UCS

This section explains the design decisions used within the Cisco UCS compute layer for both resiliency and ease of implementation.

Cisco UCS B-Series

The Cisco UCS B-Series servers were selected for this converged infrastructure. Supporting up to 3TB of memory in a half width blade format, these Cisco UCS servers are ideal virtualization hosts. These servers are configured in the design with:

- Diskless SAN boot – Persistent operating system installation, independent of the physical blade for true stateless computing.
- VIC 1440 – Dual-port 40Gbps capable of up to 256 Express (PCIe) virtual adapters
- VIC Port Expander⁽¹⁾ – Enablement of the full 40G bandwidth for the adapter

(1) The Port Expander is not supported when used with the 2208XP IOM connecting to the Cisco UCS 6454 FI, hosts configured with a Port Expander using this combination of IOM and FI will only see the 20Gbps that a VIC 1440 will show when used without a Port Expander.

Cisco UCS Service Profiles and Cisco UCS Service Profile Templates

Cisco UCS Service Profiles (SP) were configured with identity information pulled from pools (WWPN, MAC, UUID, etc) as well as policies covering firmware to power control options. These SP are provisioned from Cisco UCS Service Profile Templates that allow rapid creation, as well as guaranteed consistency of the hosts at the UCS hardware layer.

Cisco UCS vNIC Templates

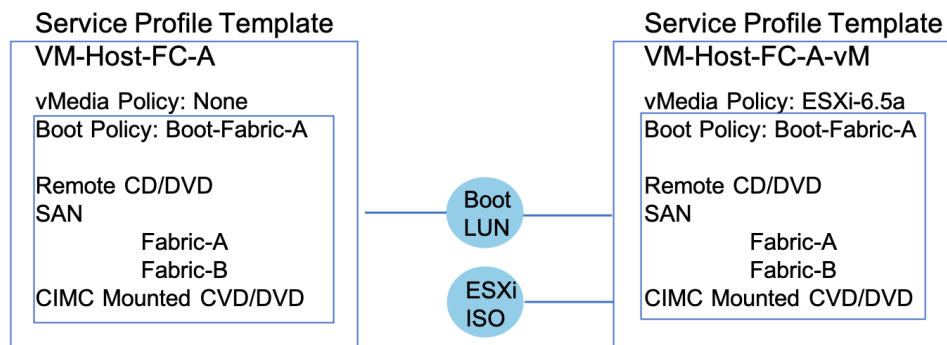
Cisco UCS virtual Network Interface Cards (vNICs) are created as virtual adapters from the Cisco UCS VICs in the host, and vNIC Templates provide a repeatable, reusable, and adjustable sub-component of the SP template for handling these vNICs. These vNICs templates were adjusted within the options for:

- Fabric association or failover between fabrics
- VLANs that should be carried
- Native VLAN specification
- VLAN and setting consistency with another vNIC template
- vNIC MTU
- Consistent Device Naming (CDN) used to guarantee the expected order of interfaces
- Specification of a MAC Pool

Cisco UCS vMedia

The installation of ESXi is simplified at scale through the use of a Cisco UCS vMedia policy to present the installation ISO through the network. The HTTP service was used to validate this from an existing resource in the environment, but HTTPS, NFS, and CIFS are additional options for presenting the ISO.

Figure 25 Logical View of vMedia Policy within Cisco UCS Service Profile



During the initial setup, the template for the ESXi hosts were created, as shown on the left (Figure 25). This template was cloned and adjusted to add a vMedia Policy to allow for a boot from ISO. Hosts are provisioned from this vMedia enabled template and after installation, the provisioned Service Profiles are unbound from the template enabled for vMedia and bound to the template without the vMedia Policy.

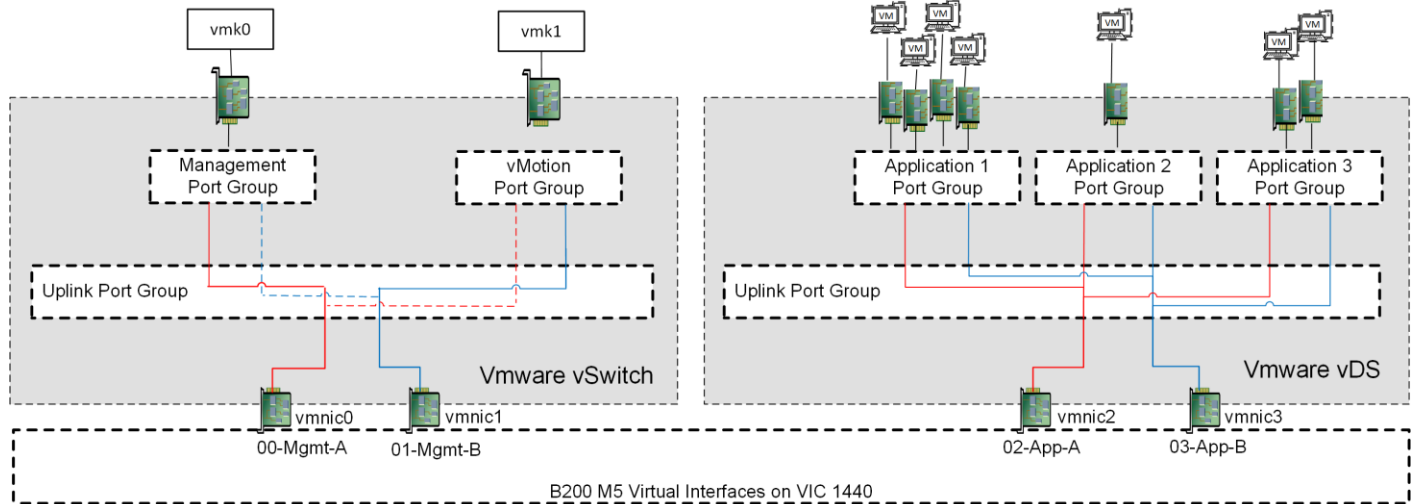
VMware vSphere

VMware vSphere is the hypervisor in this design, with validation covering both vSphere 6.5 U2 and vSphere 6.7 U1 to support a larger customer install base. Design implementations between the two vSphere releases did not differ in this architecture and were both managed by the same vCenter 6.7 instance for consistency and simplification of the environment.

Virtual Networking Configuration

The virtual networking configuration on the Cisco UCS B200s takes advantage of the Cisco UCS VIC adapter to create multiple virtual adapters to present to the ESXi installation as shown in Figure 26.

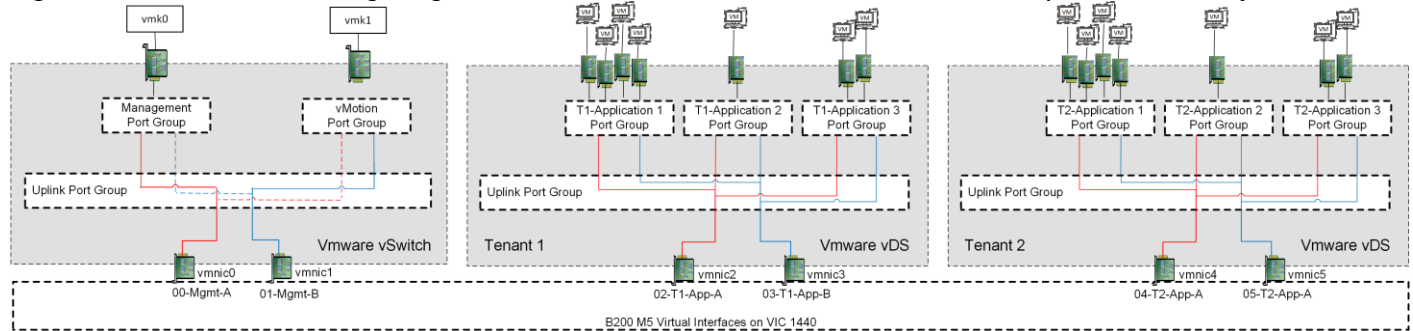
Figure 26 Virtual Networking Diagram for a Cisco UCS B200 M5 ESXi Host



The ESXi management interfaces are carried on a single pair of vNICs delivering dedicated VLANs for In-Band management and for vMotion used by the host. These vNICs are connected to a VMware vSwitch for simplification of configuration, and to ensure portability if the vCenter was somehow compromised. The VMkernel configuration of both management and vMotion are pinned in an active/standby configuration setting on opposing links, to keep these types of traffic contained within a particular Cisco UCS fabric interconnect when switching this traffic between ESXi hosts, thus preventing the need to send it up into the Nexus switch to pass between fabrics.

All application traffic is set off of another pair of vNICs coming from the VIC adapter that are associated with a vDS to allow of quick expansion of additional application port groups, and assurance of consistency between ESXi hosts. As an option, additional vNICs can be provisioned for differing tenants, to allow a level of multi-tenancy with each application tenant getting their own vDS.

Figure 27 Virtual Networking Diagram for a Cisco UCS B200 M5 ESXi Host with Expanded Tenancy



Virtual Storage Configuration

No modifications within vSphere or at the compute layer are necessary from a base install of ESXi to take advantage of Hitachi storage hardware acceleration within vSphere. The entire line of Hitachi storage systems are certified for VMware vSphere Storage APIs Array Integration (VAAI) operations within vSphere. Certain modifications to the host groups connecting to the Cisco UCS blades running ESXi are necessary to enable full VAAI functionality for the environment and are described in the [Design Considerations](#) section.

Network Design

Nexus Configuration

The Nexus configuration covers the basic networking needs within the stack for Layer 2 to Layer 3 communication.

The following NX-OS features are implemented within the Nexus switches for the design:

- `feature interface-vlan` – Allows for VLAN IP interfaces to be configured within the switch.
- `feature hsrp` – Allows for Hot Standby Routing Protocol is set for basic routing between the tenant networks, using the VLAN interfaces of each switch for redundancy. These VLAN interfaces on each switch are configured in an active standby relationship and share a virtual IP that is pointed to by the tenant networks as their respective default gateways.
- `feature lacp` – Allows for the utilization of Link Aggregation Control Protocol (802.3ad) by the port channels configured on the switch. Port channeling is a link aggregation technique offering link fault tolerance and traffic distribution (load balancing) for improved aggregate bandwidth across member ports.
- `feature vpc` – Allows for two Nexus switches to provide a cooperative port channel called a virtual Port Channel (vPC). vPCs present the Nexus switches as a single “logical” port channel to the connecting upstream or downstream device. Configured vPCs present the same fault tolerance and traffic distribution as a port-channel but allow these benefits to occur across links distributed between the two switches. Enablement of vPCs will require a connecting vPC Peer-Link between the two switches, as well as an out of band vPC keep alive to handle switch isolation scenarios.

Management Connectivity

Out-of-band management is handled by an independent switch that could be one currently in place in the customer’s environment. Each physical device had its management interface carried through this out-of-band switch, with in-band management carried as a differing VLAN within the solution for ESXi, vCenter and other virtual management components.

Jumbo Frames

Jumbo frames are a standard recommendation across Cisco designs to help leverage the increased bandwidth availability of modern networks. To take advantage of the bandwidth optimization and reduced consumption of CPU resources gained through jumbo frames, they were configured at each network level to include the virtual switch and virtual NIC.

Setting Jumbo frames in any environment requires end to MTU modification. To enable jumbo frames from this solution to any other datacenter segments will require the MTU modification across all the links. For management traffic as well as traffic going out WAN devices, it is best to leave the MTU settings to standard 1500 bytes.

Storage Design

Several design options are available with Hitachi VSP storage systems in order to service differing vSphere workloads and environments. Choose from smaller, mid-range storage which can service 600,000 IOPS and 2.4PB of capacity to enterprise-class storage which can service up to 4.8 million IOPS and 34.6PB of capacity. [Table 2](#) provides a comparison of the different models of VSP available within the families tested in this design.

Table 2 Comparison of VSP Family Models

VSP Model	F350, F370, G350, G370	F700, G700	F900, G900	F1500, G1500
-----------	------------------------	------------	------------	--------------

VSP Model	F350, F370, G350, G370	F700, G700	F900, G900	F1500, G1500
Storage Class	Mid-Range			Enterprise
Maximum IOPS	600K to 1.2M IOPS	1.4M IOPS	2.4M IOPS	4.8M IOPS
	9 to 12GB/s bandwidth	24GB/s bandwidth	41GB/s bandwidth	48GB/s bandwidth
Maximum Capacity	2.8 to 4.3PB (SSD)	6PB (FMD)	8.1PB (FMD)	8.1PB (FMD)
	2.4 to 3.6PB (HDD)	13PB (SSD)	17.3PB (SSD)	34.6PB (SSD)
		11.7PB (HDD)	14PB (HDD)	6.7PB (HDD)
Drive Types	480GB, 1.9/3.8/7.6/15TB SSD 600GB, 1.2/2.4TB 10K HDD 6/10TB 7.2K HDD	3.5/7/14TB FMD 480GB, 1.9/3.8/7.6/15TB SSD 600GB, 1.2/2.4TB 10K HDD 6/10TB 7.2K HDD	3.5/7/14TB FMD 1.9/3.8/7.6/15TB SSD 600GB, 1.2/2.4TB 10K HDD 6/10TB 7.2K HDD	1.75/3.5/7/14TB FMD 7/14TB FMD HDE 960GB, 1.9/3.8/7.6/15TB SSD 600GB 15K HDD 600GB, 1.2/1.8/2.4TB 10K HDD 4/6TB 7.2K HDD
Maximum FC Interfaces	16x (16/32Gb FC)	64x (16/32Gb FC)	80x (16/32Gb FC)	192x (8/16Gb FC)

Capacity Saving with Deduplication and Compression Options

Hitachi Virtual Storage Platform delivers superior all-flash performance for business-critical applications. It guarantees continuous data availability with a combination of high-density flash module drives (FMD). These drives use patented flash I/O management and specialized offload engines to maximize flash utilization. The key factor affecting accommodation on a flash device is not performance, but capacity. This makes key factors the high raw capacity which the flash device has and the saving ratio which comes from deduplication and compression functionalities.

Regarding deduplication and compression, the Hitachi Virtual Storage Platform family has two main types;

- Hitachi Storage Virtualization Operation System (SVOS) provides software-based deduplication and/or compression with post processing
- FMD hardware-based compression with Inline processing

When you use FMD hardware-based compression, enabling the accelerated compression option on all parity groups of FMD drives is required. You can use either software-based or hardware-based deduplication and compression, or a combination of both. With a combination of both options, software-based deduplication and hardware-based compression are used.

LUN Multiplicity Per HBA and Different Pathing Options

This design implements Single Initiator-Multi Target (SI-MT) zoning in conjunction with single vHBAs per fabric on the Cisco UCS infrastructure. This means that each vHBA within Cisco UCS will see multiple paths on their respective fabric to each LUN. Using this design requires the use of Cisco Smart Zoning within the MDS switches.

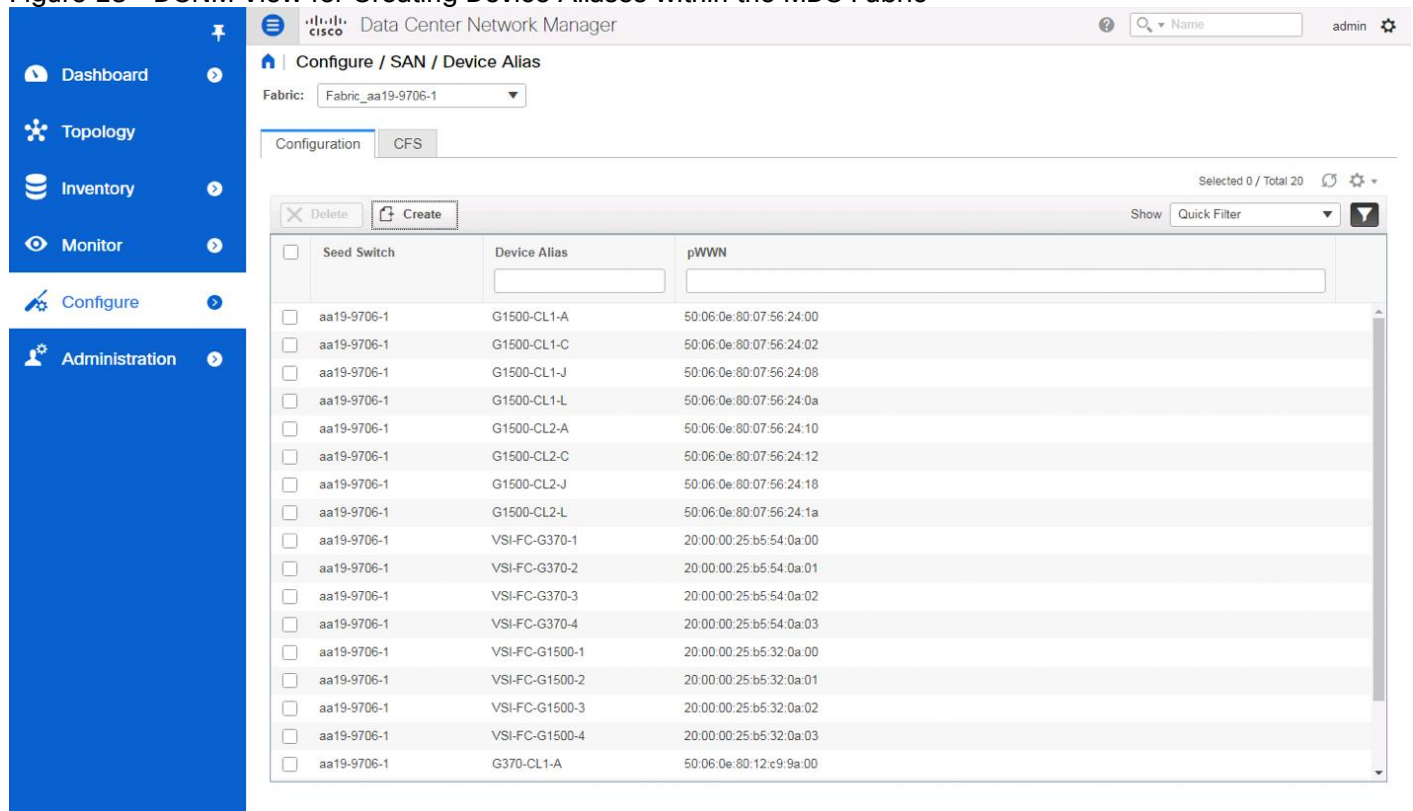
Different pathing options including Single Initiator-Single Target (SI-ST) are supported but may reduce availability and performance especially during a component failure or upgrade scenario within the overall data path.

Balance your bandwidth and application needs, vSphere cluster utilization, and availability requirements when evaluating alternative pathing options to deploy this solution.

Data Center Network Manager

Data Center Network Manager (DCNM) was used for device alias creation through an easy-to-use GUI interface (Figure 28).

Figure 28 DCNM View for Creating Device Aliases within the MDS Fabric



Device aliases created in this manner are valid for the seed switch specified, along with all others configured for that fabric.

The configuration of the fabric zoning and zoneset is also available within DCNM. Zones are easily created from selectable device aliases or end ports, with Smart Zoning specifiers of host, initiator, or both supported (Figure 29).

Figure 29 DCNM View for Creating Zones within a Zoneset

The screenshot displays the Cisco Data Center Network Manager (DCNM) interface for configuring zones within a zoneset. The interface is divided into several sections:

- Navigation Sidebar:** Includes Dashboard, Topology, Inventory, Monitor, Configure, and Administration.
- Top Navigation:** Shows the current configuration path: **Configure / SAN / Zoning**. It also displays the selected Fabric (**Fabric_aa19-9706-2**), VSAN (**VSAN-B(102)**), and Switches (**aa19-9706-2**).
- Zonesets Panel:**
 - Buttons: **Activate**, **Deactivate**
 - Table:

Zonesets	Status	Modified
ucp-vsi-zoneset	Active	No
- Zones Panel:**
 - Buttons: **Show**, **Quick Filter**
 - Table:

In Zoneset	Zone Name	Smart Zone
<input checked="" type="checkbox"/>	VSI-FC-G1500-1	true
<input type="checkbox"/>	VSI-FC-G1500-2	true
<input type="checkbox"/>	VSI-FC-G1500-3	true
<input type="checkbox"/>	VSI-FC-G1500-4	true
<input type="checkbox"/>	VSI-FC-G370-1	true
<input type="checkbox"/>	VSI-FC-G370-2	true
<input type="checkbox"/>	VSI-FC-G370-3	true
<input type="checkbox"/>	VSI-FC-G370-4	true
- Zone Members Panel:**
 - Buttons: **Show**, **Quick Filter**
 - Table:

Zone	Zoned By	Device Type	Name	Switch Interf..	
<input type="checkbox"/>	VSI-FC-G1500-1	WWN	Storage	CL3-C	aa19-9706-2 f..
<input type="checkbox"/>	VSI-FC-G1500-1	WWN	Storage	CL4-L	aa19-9706-2 f..
<input type="checkbox"/>	VSI-FC-G1500-1	WWN	Storage	CL4-A	aa19-9706-2 f..
<input type="checkbox"/>	VSI-FC-G1500-1	WWN	Storage	CL3-L	aa19-9706-2 f..
<input type="checkbox"/>	VSI-FC-G1500-1	WWN	Host	VSI-FC-G1500-1	AA19-6332-16.
- Available to Add Panel:**
 - Buttons: **Show**, **Quick Filter**
 - Radio buttons: **Zone by:** End Ports, Device Alias
 - Table:

Type	Name	Switch Interface	Fcid	PWWN
<input type="checkbox"/>	CL4-L	aa19-9706-2 fc1/13	0x280020	50:06:0e:f..
<input type="checkbox"/>	CL4-C	aa19-9706-2 fc1/9	0x2800a0	50:06:0e:f..
<input type="checkbox"/>	G370-CL4-B	aa19-9706-2 fc1/16	0x2801a0	50:06:0e:f..

Zones created and added to a zoneset can be activated or deactivated from within DCNM for the fabric.

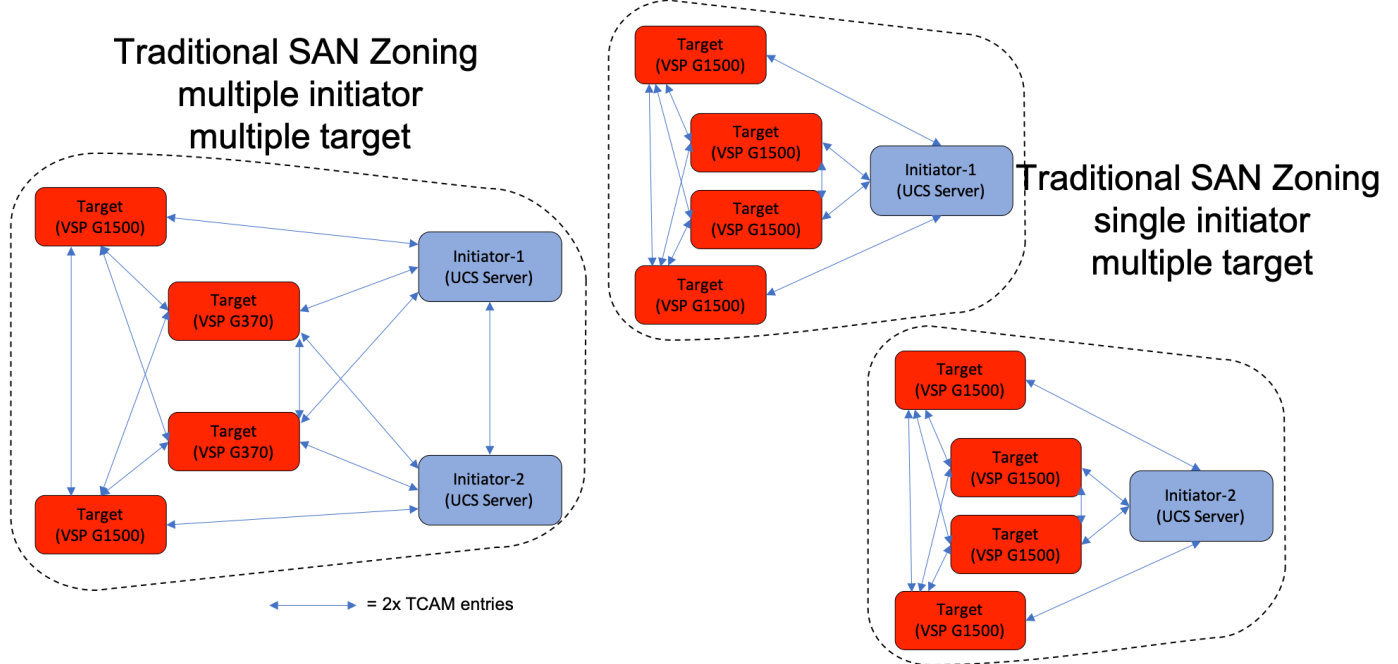
DCNM was used for the deployment of the validated architecture, using a pre-existing resource sitting outside of the Adaptive Solutions for CI converged infrastructure. Details of the installation of DCNM are not covered in the Deployment Guide, but pointers to the Cisco installation and configuration documents for DCNM are provided.

Smart Zoning

Zoning is configured as Single Initiator/Multiple Target (SI-MT) to optimize traffic intended to be specific to the initiator (UCS host vHBA) and the targets (VSP controller ports). Using SI-MT zoning provides reduced administrative overhead versus configuring single initiator/single target zoning, and results in the same SAN switching efficiency when configured with Smart Zoning.

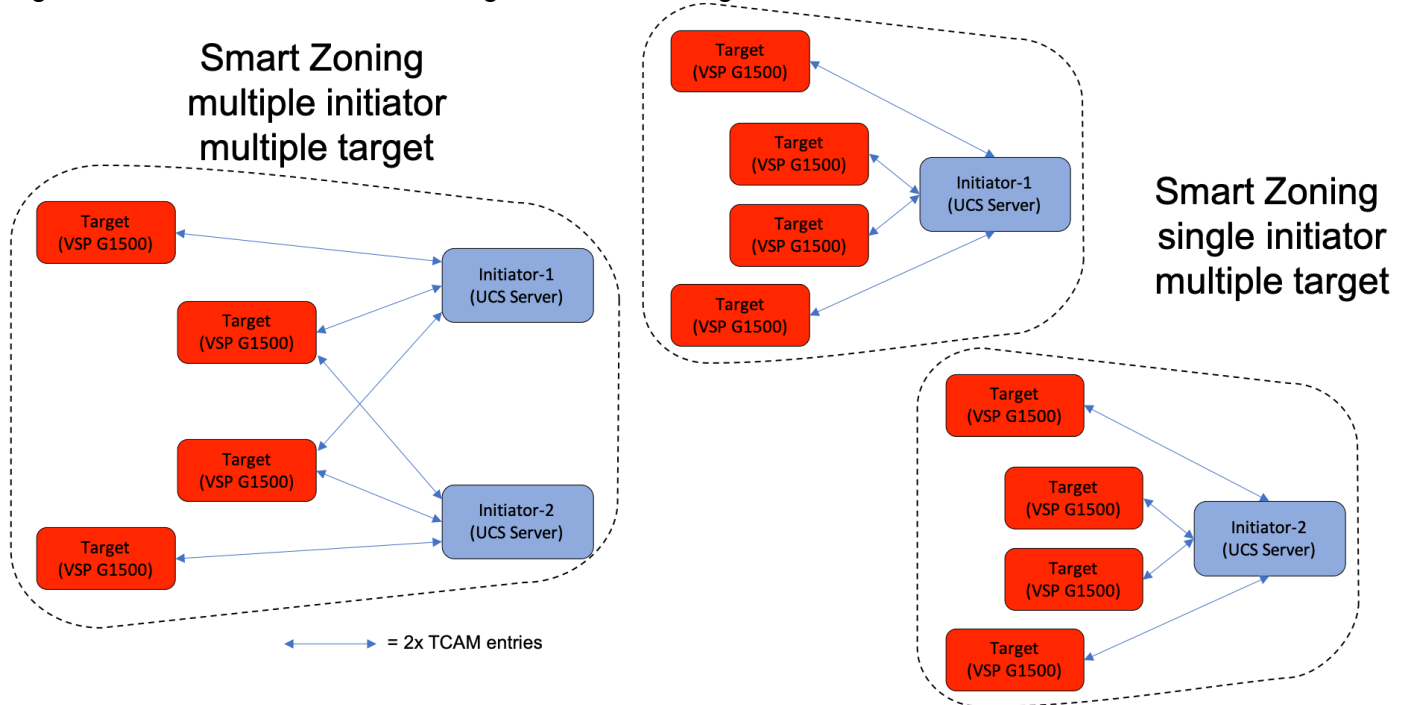
Smart Zoning is configured on the MDS to allow for reduced TCAM (ternary content addressable memory) entries, which are fabric ACL entries of the MDS allowing traffic between targets and initiators. When calculating TCAMs used, two TCAM entries will be created for each connection of devices within the zone. Without Smart Zoning enabled for a zone, targets will have a pair of TCAMs established between each other, and all initiators will additionally have a pair of TCAMs established to other initiators in the zone as shown in Figure 30.

Figure 30 TCAM View for SAN Zoning without Smart Zoning Enabled



Using Smart Zoning, Targets and Initiators are identified, reducing TCAMs needed to only occur Target to Initiator within the zone as shown in Figure 31.

Figure 31 TCAM View for SAN Zoning with Smart Zoning Enabled



Large multiple initiator to multiple target zones can grow exponentially without smart zoning enabled. Single initiator/single target zoning will produce the same amount of TCAM entries with or without Smart Zoning but will match the TCAM entries used for any multiple target zoning method that is done with Smart Zoning.

Design Considerations

Cisco Best Practices

The following Cisco design best practices and recommendations were used as references in this design.

Nexus vPC Best Practices

These best practices originated from Nexus 7000 vPC implementation but are valid for vPC configuration within Nexus 9000 switches:

https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf

MDS Best Practices

The MDS design followed basic Cisco SAN concepts for the functionality of the SAN that was deployed. These concepts do not take advantage of many of the more advanced features that optimize much more complex SAN environments that the MDS 9000 and Director Class MDS 9700 series switches can offer. Some of these more advanced feature recommendations can be found here:

<https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9700-series-multilayer-directors/white-paper-c11-738426.html>

Cisco UCS Configuration

Cisco UCS common practices as well as a thorough background on the value of the concepts utilized within Cisco UCS are presented in this white paper:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-manager/whitepaper_c11-697337.html

Performance and Tuning for Cisco UCS M5 Servers

The BIOS within Cisco UCS servers present a large number of options for optimizing the servers for differing workloads. The following white paper was referenced for adjusting the BIOS selections to optimize for virtualization workloads:

https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper_c11-740098.pdf

Hitachi Best Practices

The following Hitachi VSP storage design best practices and recommendations were used in this design.

VMFS Provisioning Best Practices

- Fibre Channel Port Options – These settings are required to be set on each fibre channel port used in the solution.
 - Port Security – Set the port security to Enable. This allows multiple host groups on the fibre channel port.
 - Fabric – Set fabric to ON. This allows connection to a fibre channel switch.

- Connection Type – Set the connection type to P-to-P. This allows a point-to-point connection to a fibre channel switch.

One Host Group Per VMware ESXi Host Configuration

If you plan to deploy VMware ESXi hosts, each host's WWN should be in its own host group. This approach provides granular control over LUN presentation to ESXi hosts. This is the best practice for SAN boot environments such as Cisco UCS, because ESXi hosts do not have access to other ESXi hosts' boot LUNs.

- Host Group Configuration and Host Mode Options – On the Hitachi Virtual Storage Platform family storage, create host groups using Hitachi Storage Navigator. Change the following host mode and host mode options to enable VMware vSphere Storage APIs – Array Integration (VAAI):
 - Host Mode – 21[VMware Extension]
 - Host Mode Options:
 - Enable 54-(VAAI) Support Option for the EXTENDED COPY command
 - Enable 63-(VAAI) Support Option for vStorage APIs based on T10 standards
 - Enable 114-(VAAI) Support Option for Auto UNMAP

VMware Round Robin PSP Rotation IOPS Limit

The VMware ESXi Round Robin Path Selection Plug-in (PSP) balances the load across all active storage paths. A path is selected and used until a specific quantity of I/O has been transferred. The I/O quantity at which a path change is triggered is known as the limit. After reaching that I/O limit, the PSP selects the next path in the list. The default I/O limit is 1000 but can be adjusted if needed to improve performance. Specifically, it can be adjusted to reduce latency seen by the ESXi host when the storage system is not seeing latency. The recommended PSP limit for most environments is 1000. Based on testing in Hitachi labs, in certain circumstances setting the value to 20 can provide a potential 3-5% reduction in latency as well as a 3-5% increase in IOPS.

Auto UNMAP

- In VMware vSphere 6.5, ESXi supports manual and automatic asynchronous reclamation of free space on VMFS5 and VMFS6 datastores. It automatically issues the UNMAP command to release free storage space in background on thin-provisioned storage arrays that support UNMAP operation.
- In VMware vSphere 6.7, more granular UNMAP rates are available and are supported by Hitachi storage systems.
- Be aware that the UNMAP operations consume processor utilization on the VSP storage arrays, so test and plan ahead before increasing UNMAP rates from their default settings.

Hitachi Dynamic Tiering and Active Flash

- Using Hitachi Dynamic Tiering, you can configure a storage system with multiple storage tiers using different types of data drives. active flash monitors a page's access frequency level in real time to promote pages that suddenly became busy from a slower media to high-performance flash media.
- In a VMware environment, many workloads tend to be highly random with smaller block size. This may not be suitable for deduplication and compression, even with all flash configuration. Hitachi Dynamic Tiering with active flash may a good option to improve capacity and cost by efficiently using the flash tier minimally.

Storage DRS Recommendations

- Storage DRS generates recommendations or performs Storage vMotion migrations to balance space use across the datastore cluster. It also distributes I/O within the datastore cluster and helps alleviate high I/O load on certain datastores. VMware recommends not mixing SSD and hard disks in the same datastore cluster. However, this does not apply to the datastores provisioned from a Hitachi Dynamic Tiering pool.

The following are recommendations for VMware vSphere Storage DRS with Hitachi storage:

- Enable only Space metrics when a datastore cluster contains multiple datastores that are provisioned from the same dynamic provisioning pool with or without Hitachi Dynamic Tiering. Moving a noisy neighbor within the same dynamic provisioning pool does not improve the performance.
- Enable Space and I/O metrics when a datastore cluster contains multiple datastores that are provisioned from different dynamic provisioning pools. Moving a noisy neighbor to the other dynamic provisioning pool balances out the performance.

Storage I/O Control (SIOC) and HBA Queue Depth Recommendations

VMware vSphere Storage I/O Control (SIOC) extends the constructs of shares and limits to handle storage I/O resources. You can control the amount of storage I/O that is allocated to virtual machines during periods of I/O congestion, which ensures that more important virtual machines get preference over less important virtual machines for I/O resource allocation. In a mixed VMware environment, increasing the HBA LUN queue depth will not solve a storage I/O performance issue. It may overload the storage processors on your storage systems. The best practice from Hitachi Data Systems is to set the default HBA LUN queue depth to 32.

Compression Recommendation and Considerations

Regarding compression, using FMD hardware-based compression is recommended for the following reasons:

- No performance degradation appears due to the truly hardware-offloaded in-line or real-time accelerated compression.
- Regarding the compression saving ratio, the differences between software-based and hardware-based are insignificant.
- Inline processing-based compression provides you with reduction of initial capacity and cost. You can estimate the required FMD capacity with the [Hitachi Data Reduction Estimator](#). ([Read more about using this tool](#), including how to get access to it.)
- Software-based compression consumes extra storage compute resources. This post processing-based compression requires full allocated capacity to temporarily store for the initial phase as well.

Deduplication Recommendation and Considerations

Deduplication is highly effective in the virtualization environment, which tends to have duplicated data. This includes data such as the same operating system images, templates, and backups. From lab validation results at Hitachi, enabling deduplication achieved a 60–70 percent capacity saving for a datastore where 8 virtual machines with an operating system VMDK resides (Microsoft Windows Server 2012 R2).

Enabling FMD hardware accelerated compression enhances deduplication with more than a 20 percent capacity saving. This combination of deduplication and compression achieved more than 80–90 percent capacity savings in total. You can also estimate saving ratio and deduped capacity with the Hitachi Data Reduction Estimator.

A main concern related to deduplication is performance degradation. This comes from mainly the following two factors:

- It consumes extra storage compute resources to perform deduplication and metadata management.
- The garbage collection running as a background task also require processing overhead. This task may increase storage CPU (MP) usage from 2 percent to 15 percent.

The following are some of the considerations regarding software-based deduplication:

- It may impact I/O performance. Verify the performance by utilizing best practices or the cache optimization tool (COT) tool before using the capacity saving function.
- Because approximately 10 percent of the capacity is used for metadata and garbage data, capacity saving function should be applied only when the saving is expected to be 20 percent or higher.
- In deduplication and compression, processing is performed per 8 KB. Therefore, if the block size of the file system is an integral multiple of 8 KB, then the capacity saving is likely to be effective.
- The capacity saving function is not a good fit for high-write workloads. If the write workload rate is higher than garbage collection throughput, then the storage cache write-pending increases, causing performance degradation.
- The capacity saving effect vary depends on your application and workload. You need to know your application workload and suitability before enabling a capacity saving feature.

Flash Module Drive Configurations and Recommendations

The key factor affecting accommodation on a flash device is not performance, but capacity. The required flash memory drive (FMD) capacity can vary, whether there is dedupeable or compressible data. You can estimate it by using the Hitachi Data Reduction Estimator. The following are some recommendations for FMD:

- If your application requires high IOPS and low latency, and if your data is compressible, FMD accelerated compression (without dedupe) might be an option.
- RAID-6 is the recommended RAID level for pool-VOLs, especially for a pool where recovery time from a pool failure due to a drive failure is not acceptable.
- Configure a parity group across the drive-boxes to maximize the performance by increasing the number of back-end paths.

Solution Validation

The solution was validated by deploying virtual machines running the IOMeter tool. The system was validated for resiliency by failing various aspects of the system under load. Examples of the types of tests executed include:

- Failure and recovery of fibre channel booted ESXi hosts in a cluster
- Rebooting of fibre channel booted hosts
- Service Profile migration between blades
- Failure of partial and complete IOM links to Fabric Interconnects
- Failure and recovery of redundant links to VSP controllers from MDS switches
- Disk removal to trigger a parity group rebuild on VSP storage

Validated Hardware

[Table 3](#) lists the hardware and software versions used during solution validation. It is important to note that Cisco, Hitachi, and VMware have compatibility matrixes that should be referenced to determine support and are available in the [Appendix](#).

Table 3 Validated Hardware and Software

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	7.0(3)I7(5a)
Compute	Cisco UCS Fabric Interconnect 6332	4.0(1b)
	Cisco UCS 2304 IOM	4.0(1b)
	Cisco UCS Fabric Interconnect 6454	4.0(1b)
	Cisco UCS 2208XP IOM	4.0(1b)
	Cisco UCS B200 M5	4.0(1b)
	VMware vSphere	6.7 U1 VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
	ESXi 6.7 U1 nenic	1.0.25.0
	ESXi 6.7 U1 nfnic	4.0.0.14

Component		Software Version/Firmware Version
	VMware vSphere	6.5 U2 VMware-ESXi-6.5.0-9298722-Custom-Cisco-6.5.2.2.iso
	ESXi 6.5 U2 nenic	1.0.25.0
	ESXi 6.5 U2 fnic	1.6.0.44
	VM Virtual Hardware Version	13 ⁽¹⁾
Storage	Hitachi VSP G1500	80-06-42-00/00
	Hitachi VSP G370	88-02-03-60/00
	Cisco MDS 9706 (DS-X97-SF1-K9 & DS-X9648-1536K9)	8.3(1)
	Cisco Data Center Network Manager	11.0(1)

(1) Virtual Hardware Version 13 was used for the convenience of migrating virtual machines between vSphere 6.5 and vSphere 6.7 hosts during validation and is not a requirement within the solution for environments that will utilize vSphere 6.7.

Summary

The Adaptive Solutions for CI is a Virtual Server Infrastructure, built by a partnership between Cisco and Hitachi to support virtual server workloads within VMware vSphere 6.5 and VMware vSphere 6.7. Adaptive Solutions for CI is a best practice data center architecture to meet the growing needs of enterprise customers utilizing virtual server workloads. The solution is built utilizing Cisco UCS Blade Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches and fibre channel-attached Hitachi VSP storage. It is designed and validated using compute, network and storage best practices for high-performance, scalability, and resiliency throughout the architecture.

Appendix: Solution Design for Cisco and Hitachi Adaptive Solutions for CI as Direct Attached Storage (DAS)

The Hitachi VSP and the Cisco UCS Fabric Interconnect can be configured in a compact architecture with the Fabric Interconnect handling the FC switching between initiators and targets. The same 16G/32G FC speeds can be achieved, but scalability is reduced and the performance analysis and monitoring capabilities of the MDS platform become unavailable.

Requirements

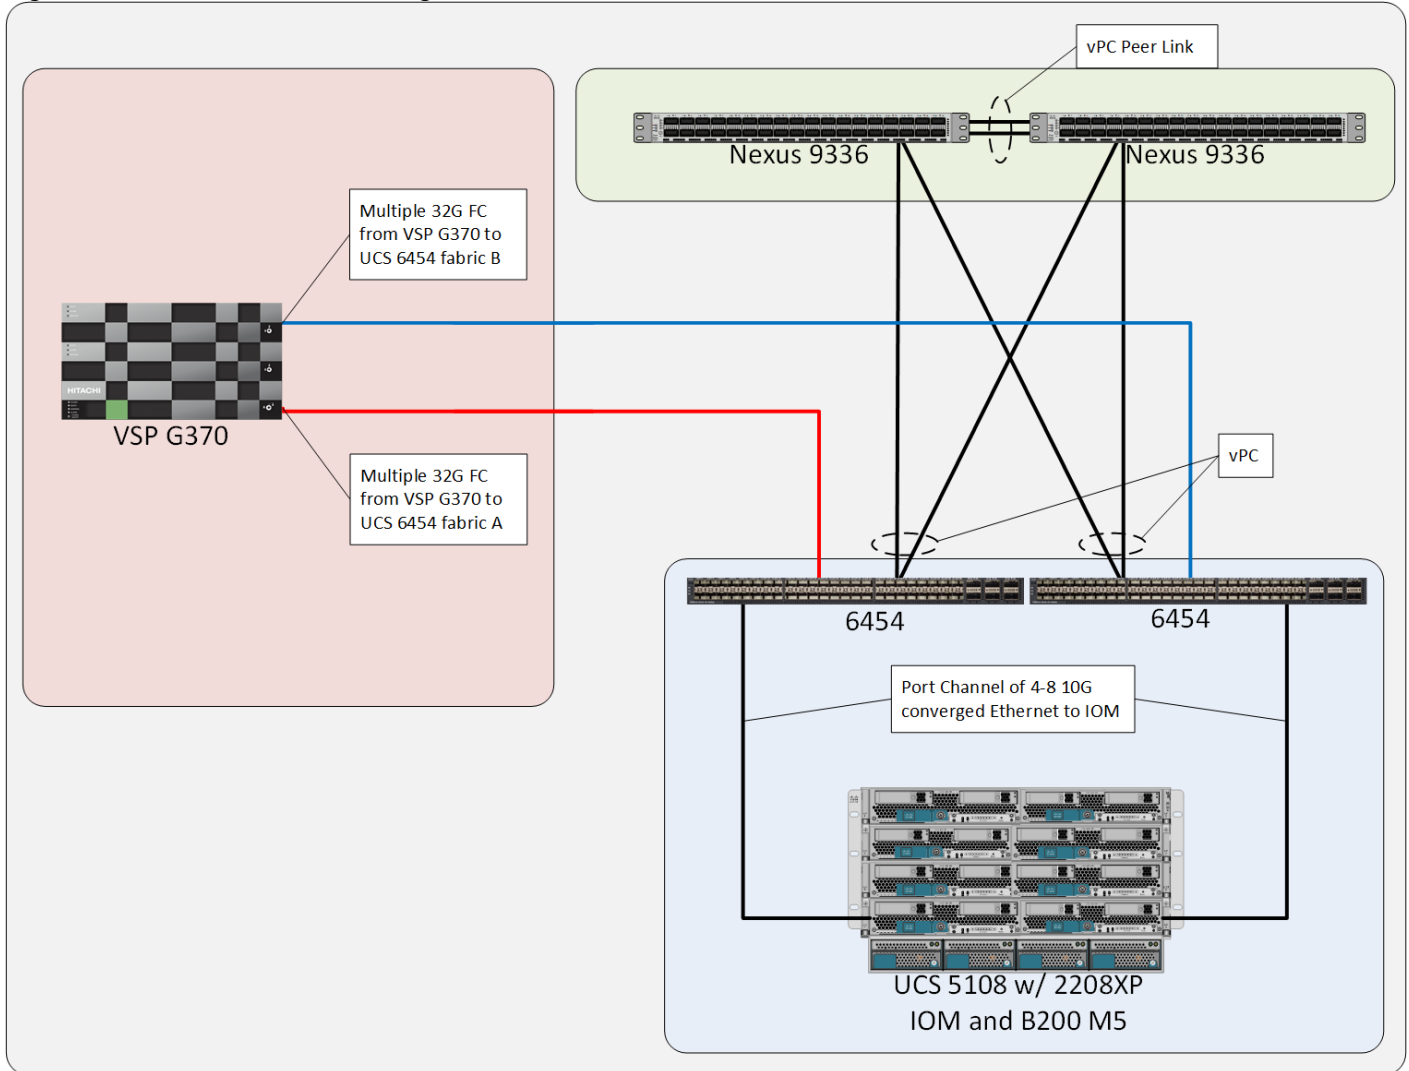
The requirements for the Direct Attached Storage design is the same for components used in the primary design. These existing items include, but may not be limited to:

- An Out of Band management network
- A terminal server for console access
- An Active Directory/DNS Server
- Layer 3 connectivity to the Internet and any other adjacent enterprise networks
- Additional management components used for deployment

Physical Topology

The design described in this appendix details the connectivity between the Cisco UCS fourth generation 6454 Fabric Interconnect and the Hitachi VSP G370. The third generation 6332-16UP Fabric Interconnect can similarly be configured for fibre channel switching if desired. The adjusted topology for the DAS solution is shown in [Figure 32](#).

Figure 32 Direct Attached Storage with Hitachi VSP G370 and Cisco UCS 6454



This architectural pairing involves the following:

- Cisco Nexus 9336C-FX2 – 100Gb capable, LAN connectivity to the UCS compute resources.
- Cisco UCS 6454 Fabric Interconnect – Unified management of UCS compute, and the compute’s access to storage and networks.
- Cisco UCS B200 M5 – High powered, versatile blade server, conceived for virtual computing.
- Hitachi VSP G370 – Mid-range, high performance storage system with optional all-flash configuration

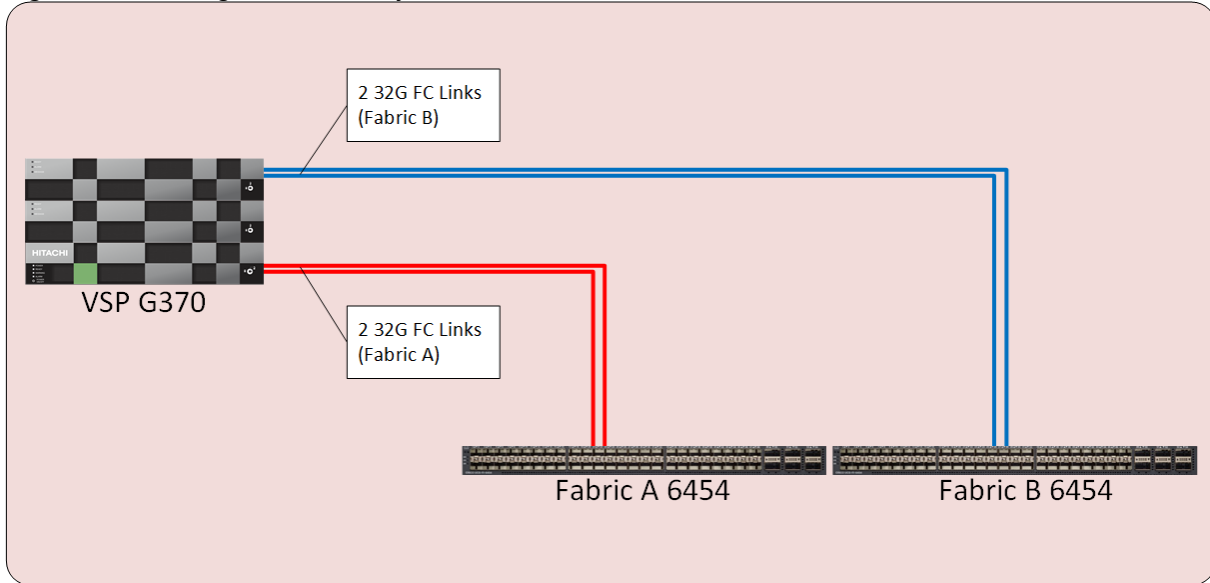
Management components of the architecture additionally include:

- Cisco UCS Manager – Management delivered through the Fabric Interconnect, providing stateless compute, and policy driven implementation of the servers it manages.
- Cisco Intersight (optional) – Comprehensive unified visibility across UCS domains, along with proactive alerts and enablement of expedited Cisco TAC communications.

Connectivity

Network and compute connectivity will have no differences from the previously shown primary topologies. The storage connectivity does change as the MDS is removed, and the Fabric Interconnect takes over FC switching as shown in [Figure 33](#).


Figure 33 Storage Connectivity for the Hitachi VSP G370 and Cisco UCS 6454



Hitachi VSP G370 FC Port to Fabric Assignments and LUN Presentation

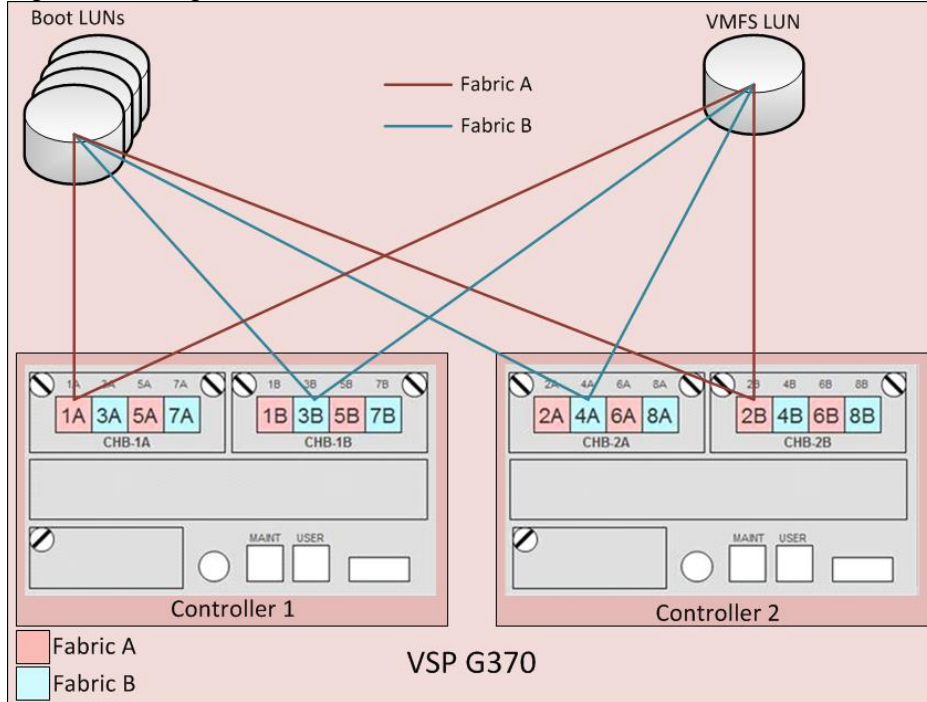
The directly connected VSP G370 storage system continues to have redundant connections to each fabric in the Cisco UCS infrastructure. The port to fabric assignments for the VSP G370 used in this design are shown in [Figure 34](#) illustrating multiple connections to each fabric and split evenly between VSP G370 controllers and 32Gb CHBs:

Figure 34 VSP G370 Channel Boards (CHB) Ports



LUN presentation remains unchanged compared to the MDS-connected design, with associations of boot LUNs and VMFS LUNs sharing the same connections as defined by the UCS boot paths for up to four connections. Similar to the MDS-connected designs, additional connections configurable to expand VMFS FC bandwidth as needed as long as additional unified ports on the Fabric Interconnects are available. [Figure 35](#) illustrates the boot LUN and VMFS LUN pathing configuration for the VSP G370.

Figure 35 Logical View of LUN to VSP G370 Port Association



Fabric Interconnect FC Switching

The UCS Fabric Interconnects take over FC switching from the MDS for the connectivity between the VSP G370 and the host connected to the FIs. The default mode for the FIs is to be in FC End-Host Mode where the FI presents attached hosts as server ports (N-ports) to a connected SAN switch. With FC switching mode enabled, the zoning and FC traffic between the server vHBAs and the VSP connected storage ports are switched within the FIs.

Fabric Interconnect Zoning

Zoning is configured automatically for the boot initiators and targets of server vHBAs that are configured with VSANs that are enabled for FC zoning. Zones created are Single Initiator Single Target as shown in the command output from an example FI zoneset below, and are created upon association of a Service Profile to a UCS blade:

```
UCS-6454-A(nx-os)# sh zoneset active vsan 101
zoneset name ucs-UCS-6454-vsan-101-zoneset vsan 101
 zone name ucs_UCS-6454_A_2_VSI-G370-01_Fabric-A vsan 101
 * fcid 0x840060 [pwwn 20:00:00:25:b5:54:0a:00] <<<-- VSI-G370-01 vHBA
 * fcid 0x840020 [pwwn 50:06:0e:80:12:c9:9a:11] <<<-- VSP G370 CL2-B

 zone name ucs_UCS-6454_A_1_VSI-G370-01_Fabric-A vsan 101
 * fcid 0x840060 [pwwn 20:00:00:25:b5:54:0a:00] <<<-- VSI-G370-01 vHBA
 * fcid 0x840000 [pwwn 50:06:0e:80:12:c9:9a:00] <<<-- VSP G370 CL1-A
```

The boot initiators defined by the UCS Boot Policy primary and secondary SAN Targets are configured for these example zones above. Additional initiators that are not defined as primary or secondary SAN Targets can be included in the zoning but will need to be specified in a UCS Storage Connection Policy and have the Storage Connection Policy configured within the UCS Service Policy zoning section.

End-to-End Data Path

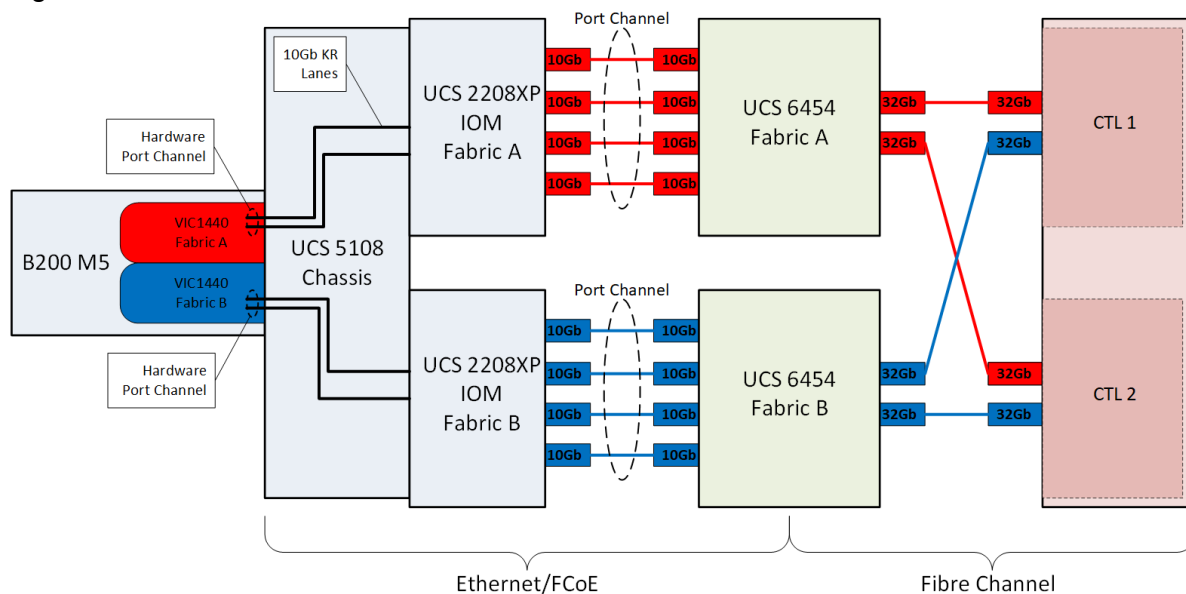
The DAS architecture becomes slightly shorter than the primary architectures discussed in the main design. This shortened implementation is still a high bandwidth fibre channel storage solution, with 32G end-to-end FC for the

Cisco UCS 6454 to the VSP G370. The storage traffic for DAS flows from a Cisco UCS B200 blade in a UCS environment to Hitachi VSP G370 is as follows:

- The Cisco UCS B200 M5 server, equipped with a VIC 1440 adapter⁽¹⁾, connects to each fabric at a link speed of 20Gbps.
- Pathing through 10Gb KR lanes of the Cisco UCS 5108 Chassis backplane into the Cisco UCS 2208XP IOM (Fabric Extender).
- Connecting from each IOM to the Fabric Interconnect with pairs of 10Gb uplinks automatically configured as port channels during chassis association, that carry the FC frames as FCoE along with the Ethernet traffic coming from the chassis blades.
- Within the Cisco UCS 6454 Fabric Interconnects, the FCoE from the blades is converted back to FC and is switched to the appropriate 32G FC ports configured for the connected VSP G370 controller.
- Ending at the Hitachi VSP G370 fibre channel controller ports with dedicated F_Ports on the Cisco UCS 6454 for each N_Port WWPN of the VSP G370 controller, with each fabric evenly split between the controllers and CHBs.

(1) The VIC 1440 will work with the UCS 6454 to provide 40G/40G when equipped with the 4th generation IOM which is not available at the time of this validation.

Figure 34 End-to-End Data Path from Cisco UCS B200 M5 to VSP G370



Validated Hardware for Direct Attached Storage

The UCS firmware differed for the DAS validation as the initial 4.0(1) UCSM release did not support FC switching for the 4th Gen 6454 FI (the 4.0(1) and earlier releases of UCSM do support FC switching in the 2nd and 3rd Gen FIs). The FC switching was enabled in 4.0(2) for the 6454 and was installed along with some accompanying changes in nfnic and nenic drivers within vSphere. The full list of what was used in the DAS architecture is shown below:

Table 4 Validated Hardware and Software

Component	Software Version/Firmware Version

Component		Software Version/Firmware Version
Network	Cisco Nexus 9336C-FX2	7.0(3)I7(5a)
Compute	Cisco UCS Fabric Interconnect 6454	4.0(2b)
	Cisco UCS 2208XP IOM	4.0(2b)
	Cisco UCS B200 M5	4.0(2b)
	VMware vSphere	6.7 U1 VMware_ESXi_6.7.0_10302608_Custom_Cisco_6.7.1.1.iso
	ESXi 6.7 U1 nenic	1.0.27.0
	ESXi 6.7 U1 nfnic	4.0.0.33
	VM Virtual Hardware Version	13 ⁽¹⁾
Storage	Hitachi VSP G370	88-02-03-60/00

(1) Hardware Version 13 was kept for initial transfer support between vSphere 6.5 and vSphere 6.7 of the VM test harness used.

Direct Attached Storage Deployment Guide

A full deployment guide for this DAS solution option has been put together, and can be found here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_hitachi_adaptivesolutions_ci_da.html

Appendix: Solution References

Compute

Cisco Unified Computing System:

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

Cisco UCS 6300 Series Fabric Interconnects:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-736682.html>

Cisco UCS 6400 Series Fabric Interconnects:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html>

Cisco UCS 5100 Series Blade Server Chassis:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-5100-series-blade-server-chassis/data_sheet_c78-526830.html

Cisco UCS 2300 Series Fabric Extenders:

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/datasheet-c78-675243.html>

Cisco UCS 2200 Series Fabric Extenders:

https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-6300-series-fabric-interconnects/data_sheet_c78-675243.html

Cisco UCS B-Series Blade Servers:

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

Cisco UCS VIC 1440 Adapter:

<https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.html>

Cisco UCS Manager:

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

Network and Management

Cisco Nexus 9000 Series Switches:

<https://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/datasheet-listing.html>

Cisco MDS 9000 Series Multilayer Switches:

<http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html>

Cisco Data Center Network Manager 11:

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/prime-data-center-network-manager/datasheet-c78-740978.html>

Storage

Cisco and Hitachi Adaptive Solutions CI with Red Hat OpenShift Platform:

<https://www.hitachivantara.com/en-us/pdf/architecture-guide/cisco-hitachi-adaptive-solutions-ci-with-red-hat-openshift-platform-implementation-guide.pdf>

Data Protection with Hitachi Ops Center Protector on Cisco and Hitachi Adaptive Solutions:

https://knowledge.hitachivantara.com/Documents/Application_Optimized_Solutions/VMWare/Data_Protection_with_Hitachi_Ops_Center_Protector_on_Cisco_and_Hitachi_Adaptive_Solutions_for_Converged_Infrastructure_Implementation_Guide

Hitachi Virtual Storage Platform F Series:

<https://www.hitachivantara.com/en-us/pdf/datasheet/vsp-f-series-all-flash-enterprise-cloud-solutions-datasheet.pdf>

Hitachi Virtual Storage Platform G Series:

<https://www.hitachivantara.com/en-us/pdf/datasheet/vsp-g-series-hybrid-flash-midrange-cloud-solutions-datasheet.pdf>

Virtualization Layer

VMware vCenter Server:

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

VMware vSphere:

<https://www.vmware.com/products/vsphere>

Compatibility Matrixes

Cisco UCS Hardware Compatibility Matrix:

<https://ucshcltool.cloudapps.cisco.com/public/>

Cisco Nexus Recommended Releases for Nexus 9K:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/recommended_release/b_Minimum_and_Recommended_Cisco_NX-OS_Releases_for_Cisco_Nexus_9000_Series_Switches.html

Cisco MDS Recommended Releases:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/b_MDS_NX-OS_Recommended_Releases.html

Cisco Nexus and MDS Interoperability Matrix:

<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx/Matrix1.html>

Cisco MDS 9000 Family Pluggable Transceivers Data Sheet:

https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9000-series-multilayer-switches/product_data_sheet09186a00801bc698.html

Hitachi Interoperability:

https://support.hitachivantara.com/en_us/interoperability.html sub-page -> (VSP G1X00, F1500, Gxx0, Fxx0, VSP, HUS VM VMWare Support Matrix)

VMware and Cisco Unified Computing System:

<http://www.vmware.com/resources/compatibility>

About the Authors

Ramesh Isaac, Technical Marketing Engineer, Cisco Systems, Inc.

Ramesh Isaac is a Technical Marketing Engineer in the Cisco UCS Data Center Solutions Group. Ramesh has worked in the data center and mixed-use lab settings since 1995. He started in information technology supporting UNIX environments and focused on designing and implementing multi-tenant virtualization solutions in Cisco labs before entering Technical Marketing where he has supported converged infrastructure and virtual services as part of solution offerings as Cisco. Ramesh has certifications from Cisco, VMware, and Red Hat.

Tim Darnell, Master Solutions Architect and Product Owner, Hitachi Vantara

Tim Darnell is a Master Solutions Architect and Product Owner in the Hitachi Vantara Converged Product Engineering Group. Tim has worked on data center and virtualization technologies since 1997. He started his career in systems administration and has worked in a multitude of roles since, from technical reference authoring to consulting in large, multi-national corporations as a technical advisor. He is currently a Product Owner at Hitachi Vantara, responsible for the Unified Compute Platform Converged Infrastructure line of products that focus on VMware vSphere product line integrations. Tim holds multiple VCAP and VCP certifications from VMware and is a RedHat Certified Engineer.

Acknowledgements

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

- Haseeb Niazi, Technical Marketing Engineer, Cisco Systems, Inc.
- Archana Sharma, Technical Marketing Engineer, Cisco Systems, Inc.
- Bhavin Yadav, Technical Marketing Engineer, Cisco Systems, Inc.
- Michael Nakamura, Director Virtualization Engineering, Hitachi Vantara