

# VersaStack with Cisco UCS Mini and IBM V5000 Design Guide

**Last Updated:** January 27, 2016



## About Cisco Validated Designs

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

<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 NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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

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

© 2016 Cisco Systems, Inc. All rights reserved.

# Table of Contents

Executive Summary .....	5
Solution Overview .....	6
Introduction .....	6
Audience .....	6
<b>What's New?</b> .....	6
VersaStack Program Benefits.....	7
Integrated System .....	7
Fabric Infrastructure Resilience.....	8
Network Virtualization .....	8
Technology Overview .....	9
Validated System Components .....	10
Cisco Unified Computing System.....	10
Cisco UCS 5108 Blade Server Chassis .....	12
Cisco UCS 6324 Fabric Interconnect .....	13
Cisco UCS B200 M4 Servers .....	13
Cisco I/O Adapters .....	14
Cisco UCS Differentiators .....	15
Cisco Nexus 9000 Series Switch .....	17
Cisco Nexus 1000v.....	18
IBM Storwize V5000 .....	18
VMware vSphere .....	19
Domain and Element Management.....	19
Cisco Unified Computing System Manager .....	19
Cisco UCS Central (Optional).....	20
Cisco UCS Manager (Mini) Management with Cisco UCS Central .....	21
IBM Storwize® V5000 Easy-to-Use Management GUI .....	22
VMware vCenter Server.....	24
Cisco Virtual Switch Update Manager .....	24
IBM Spectrum Protect Snapshot .....	25
Solution Design.....	26
Storage Design.....	28
IP-Based SAN Storage using iSCSI .....	29
Direct Attached SAN Storage using Fibre Channel.....	29

Physical Topology.....	30
Design Considerations.....	32
Storage .....	32
Cisco UCS Server Configuration for vSphere.....	33
Jumbo Frames.....	34
Management Network .....	34
High Availability .....	34
Compute and Virtualization .....	34
Network.....	35
Scale and Performance .....	35
Compute .....	35
Storage .....	37
Validation.....	39
Hardware and Software Revisions .....	39
Hardware and Software Options for VersaStack .....	40
Summary .....	42
References .....	43
Products and Solutions .....	43
Interoperability Matrixes.....	43
About the Authors.....	45
Acknowledgements .....	45

## Executive Summary

---

Cisco® Validated Designs include 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 customers.

This document describes the VersaStack® solution with the Cisco UCS Mini and IBM V5000 into a single, flexible architecture. A VersaStack solution is a validated approach for deploying Cisco and IBM technologies as a shared cloud infrastructure.

VersaStack™ is predesigned, integrated platform architecture for the data center that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® family of switches, and IBM storage arrays. VersaStack is designed with no single point of failure and a focus of simplicity, efficiency, and versatility. VersaStack is a suitable platform for running a variety of virtualization hypervisors as well as bare metal operating systems to support enterprise workloads.

VersaStack delivers a baseline configuration and also has the flexibility to be sized and optimized to accommodate many different use cases and requirements. System designs discussed in this document have been validated for resiliency by subjecting to multiple failure conditions while under load. Fault tolerance to operational tasks such as firmware and operating system upgrades, switch, cable and hardware failures, and loss of power has also been ascertained. This document describes a solution with VMware vSphere 5.5 Update 2 built on the VersaStack with block storage. The document from Cisco and IBM discusses design choices made and best practices followed in deploying the shared infrastructure platform.

## Solution Overview

---

### Introduction

Industry trends indicate a vast data center transformation toward converged solutions and cloud computing. Enterprise customers are moving away from disparate layers of compute, network and storage to integrated stacks providing the basis for a more cost-effective virtualized environment that can lead to cloud computing for increased agility and reduced cost.

To accelerate this process and simplify the evolution to a shared cloud infrastructure, Cisco and IBM have developed the solution on VersaStack™ for VMware vSphere®. Enhancement of this solution with automation and self-service functionality and development of other solutions on VersaStack™ are envisioned under this partnership.

By integrating standards based components that are compatible, scalable and easy to use, VersaStack addresses customer issues during the planning, design and implementation stages. When deployed, the efficient and intuitive front-end tools provide the means to manage the platform in an easy and agile manner. The VersaStack architecture thus mitigates customer risk and eliminates critical pain points while providing necessary guidance and measurable value. The result is a consistent platform with characteristics to meet changing workloads of any customer.

### Audience

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

### What's New?

The following design elements distinguish this version of VersaStack from previous models:

- Validation of the Cisco UCS Mini with Cisco Nexus 9000 switches and IBM Storwize V5000 storage array
- Support for the Cisco UCS 3.0(2d) and 3.0(2e) release and Cisco UCS B200-M4 servers
- Support for the release of IBM Storwize software 7.4.0.6
- Fibre Channel and IP based storage designs with block storage based datastores
- Support for iSCSI and FC access to LUNs

For more information on earlier VersaStack models, see:

<http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html>

## VersaStack Program Benefits

Cisco and IBM have thoroughly validated and verified the VersaStack solution architecture and its many use cases while creating a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model. This portfolio will include, but is not limited to the following items:

- Best practice architectural design
- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for what is, and what is not, a VersaStack configuration)
- Frequently asked questions (FAQs)
- Cisco Validated Designs (CVDs) and IBM Redbooks focused on a variety of use cases

Cisco and IBM have also built a robust and experienced support team focused on VersaStack solutions, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance provided by IBM and Cisco provides customers and channel services partners with direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

VersaStack supports tight integration with hypervisors leading to virtualized environments and cloud infrastructures, making it the logical choice for long-term investment. Table 1 lists the features in VersaStack:

**Table 1 VersaStack Component Features**

IBM Storwize V5000 Storage	Cisco UCS and Cisco Nexus 9000 Switches
Easy Tier technology	Unified Fabric
Metro Mirror and Global Mirror	Virtualized IO
Enhanced IP replication	Extended Memory
Form factor scaling capability	Stateless Servers through policy based management
Flash optimization	Centralized Management
Big data and analytics enablement	Investment Protection
External virtualization	Scalability
IBM FlashCopy	Automation

## Integrated System

VersaStack is a pre-validated infrastructure that brings together compute, storage, and network to simplify, accelerate, and minimize the risk associated with data center builds and application rollouts. These integrated systems provide a standardized approach in the data center that facilitates staff expertise, application onboarding, and automation as well as operational efficiencies relating to compliance and certification.

## Fabric Infrastructure Resilience

VersaStack is a highly available and scalable infrastructure that IT can evolve over time to support multiple physical and virtual application workloads. VersaStack has no single point of failure at any level, from the server through the network, to the storage. The fabric is fully redundant and scalable and provides seamless traffic failover should any individual component fail at the physical or virtual layer.

## Network Virtualization

VersaStack delivers the capability to securely connect virtual machines into the network. This solution allows network policies and services to be uniformly applied within the integrated compute stack using technologies such as virtual LANs (VLANs), quality of service (QoS), and the Cisco Nexus 1000v virtual distributed switch. This capability enables the full utilization of VersaStack while maintaining consistent application and security policy enforcement across the stack even with workload mobility.

VersaStack provides a uniform approach to IT architecture, offering a well-characterized and documented shared pool of resources for application workloads. VersaStack delivers operational efficiency and consistency with the versatility to meet a variety of SLAs and IT initiatives, including:

- Application rollouts or application migrations
- Business continuity/disaster recovery
- Desktop virtualization
- Cloud delivery models (public, private, hybrid) and service models (IaaS, PaaS, SaaS)
- Asset consolidation and virtualization

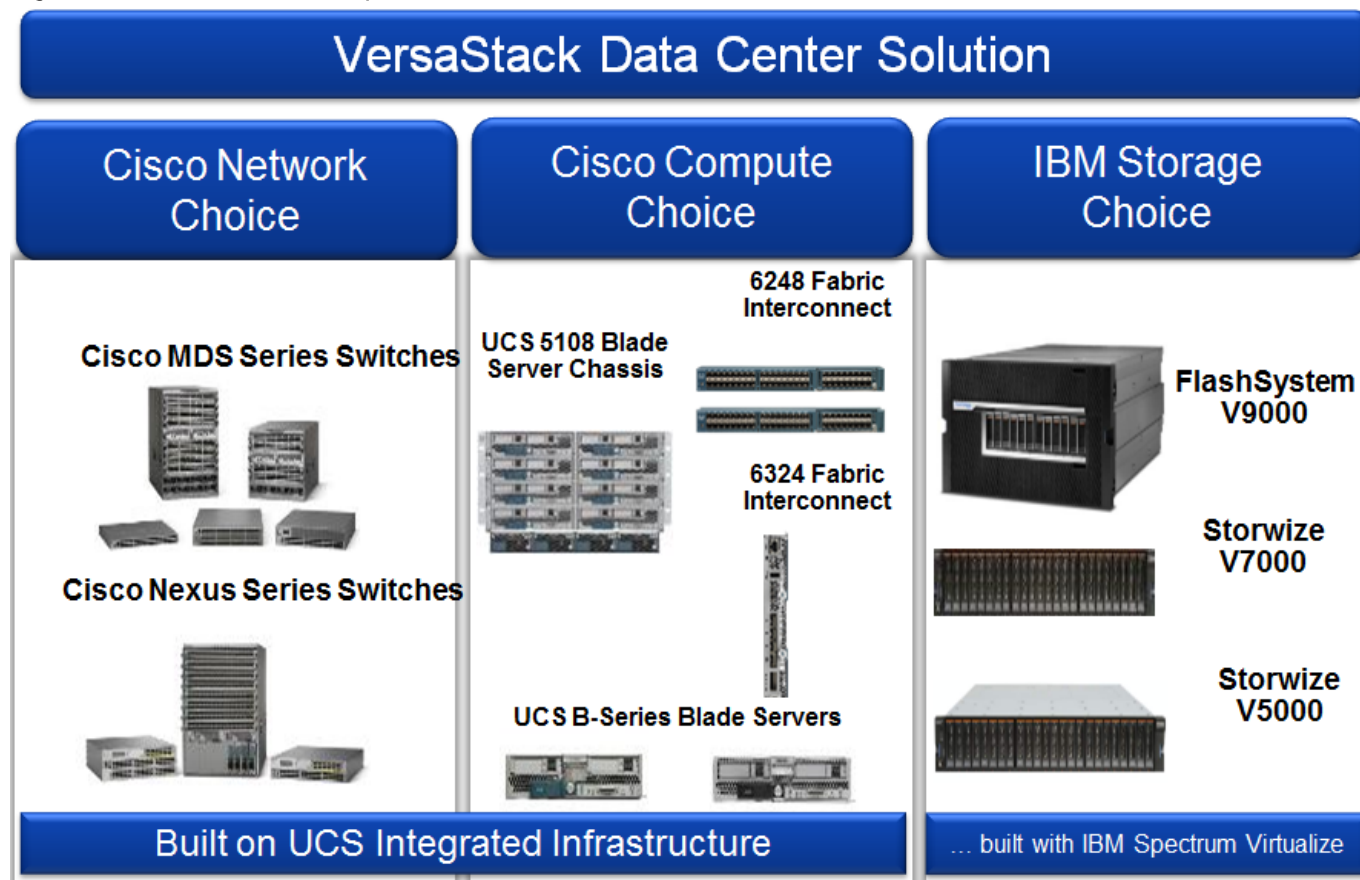


## Technology Overview

VersaStack is a best practice data center architecture that includes the following components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus and MDS switches
- IBM FlashSystem and IBM Storwize family storage

Figure 1 VersaStack Components



These components are connected and configured according to best practices of both Cisco and IBM and provide the ideal platform for running a variety of enterprise workloads with confidence.

The reference architecture covered in this document leverages the Cisco Nexus 9000 for switching, Cisco UCS Mini platform for the Compute and IBM Storwize V5000 Storage. VersaStack can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed), or it can scale out for environments that need multiple consistent deployments (rolling out additional VersaStack stacks).

One of the key benefits of VersaStack is the ability to maintain consistency at scale. Each of the component families shown in Figure 1 (Cisco Unified Computing System, Cisco Nexus, and IBM storage arrays) offers platform and resource options to scale the infrastructure up or down, while supporting the same features and functionality that are required under the configuration and connectivity best practices of VersaStack.

## Validated System Components

The following components are required to deploy this VersaStack design with Cisco UCS Mini:

- Cisco UCS Mini
- Cisco Nexus 9372PX Series Switch
- IBM Storwize V5000
- VMware vSphere 5.5 U2

This section provides a technical overview of the above components

### Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. The Cisco Unified Computing System accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and non-virtualized systems.

The main components of the Cisco UCS are:


- Compute —The system is based on the industry leading data center computing system that incorporates rack mount and blade servers based on Intel processors.
- Network —The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- Virtualization —The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access —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. Only iSCSI and Fibre Channel based access is supported in this version of VersaStack solution.
- Management—The system uniquely integrates all system components to enable the entire solution to be managed as a single entity by the Cisco UCS Manager. The Cisco UCS Manager has an intuitive graphical user interface (GUI), a command-line interface (CLI), and a powerful scripting library module for Microsoft PowerShell built on a robust application programming interface (API) to manage all system configuration and operations.

Cisco UCS Mini used in this design delivers all of the above capabilities in an easy-to-deploy compact form factor. Cisco UCS Mini is for smaller deployments with less server needs but with the same Enterprise-class features and management as that of a Cisco UCS system. Cisco UCS Mini allows IT departments to address ROBO, SMB or Data Center infrastructure challenges using a streamlined architecture following compute, network and storage best practices.

The Cisco UCS 6324 Fabric Interconnect extends the Cisco UCS architecture into environments with lesser resource requirements. Providing the same unified server and networking capabilities as the full-scale Cisco UCS solution, the Cisco UCS 6324 Fabric Interconnect embeds the connectivity within the Cisco UCS 5108 Blade Server Chassis to provide a smaller domain of up to 15 servers (8 blade servers and up to 7 direct-connect rack servers).

Figure 2 shows the Cisco UCS Mini in a highly available cohesive architecture with integrated unified fabric interconnects.

**Figure 2 The Cisco Unified Computing System in a Highly Available Cohesive Architecture**



Cisco UCS Mini consists of the following components.

- [Cisco UCS 5108 Blade Server Chassis](#) – Cisco UCS chassis can accommodate up to eight half-width Cisco UCS B200 M4 Blade Servers.
- [Cisco UCS 6324 Fabric Interconnect](#) – Cisco UCS 6324 is embedded within the Cisco UCS 5108 Blade Server Chassis and provides the same unified management capabilities as the standalone Cisco UCS 6200 Series Fabric Interconnects.

- [Cisco UCS Manager](#) – UCS Manager provides unified, embedded management of all software and hardware components in a Cisco UCS Mini solution.
- [Cisco UCS B200 M4 Blade Server](#) – Cisco UCS B200 M4 Blade Server addresses a broadest set of workloads, delivering performance, versatility, and density without compromise.

### Optional Solution Components

- [Cisco UCS C220 M4 Rack Server](#) – This one-rack-unit (1RU) server offers superior performance and density over a wide range of business workloads.
- [Cisco UCS C240 M4 Rack Server](#) – This 2RU server is designed for both performance and expandability over a wide range of storage-intensive infrastructure workloads.
- [Cisco UCS Central](#) – Cisco UCS Central manages multiple Cisco UCS Mini and UCS domains.

### Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server chassis. The Cisco UCS 5108 Blade Server Chassis is six rack units (6RU) high and can mount in an industry-standard 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series Blade Servers and can accommodate both half-width and full-width blade form factors. The UCS Mini chassis supports the B22 M3, B200 M3, B420 M3, and B200 M4 blade servers today. Cisco UCS C-series rack mount servers can also be connected into the UCS Mini chassis – see link below for a complete list of supported servers. Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+ 1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays. On the Cisco UCS Mini chassis, the I/O bays are used to accommodate the UCS 6324 Fabric Interconnect modules. A passive mid-plane provides up to 80 Gbps of I/O bandwidth per server slot and up to 160 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 40 Gigabit Ethernet standards.

For more information, see:

<http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/ucsmini-specsheet.pdf>

Figure 3 Cisco UCS 5108 Blade Chassis



## Cisco UCS 6324 Fabric Interconnect

The key to delivering the power of Cisco Unified Computing System in a smaller form factor on the Cisco UCS Mini is the Cisco UCS 6324 Fabric Interconnect. The Fabric Interconnect modules (up to two) plug into the back of the Cisco UCS Mini Series blade server chassis. A mid-plane connects the blade servers to the Fabric Interconnects. The Cisco UCS 6324 Fabric Interconnect combines the Fabric Extender and Fabric Interconnect functions into one plug-in module, and allows direct connection to an external switch.

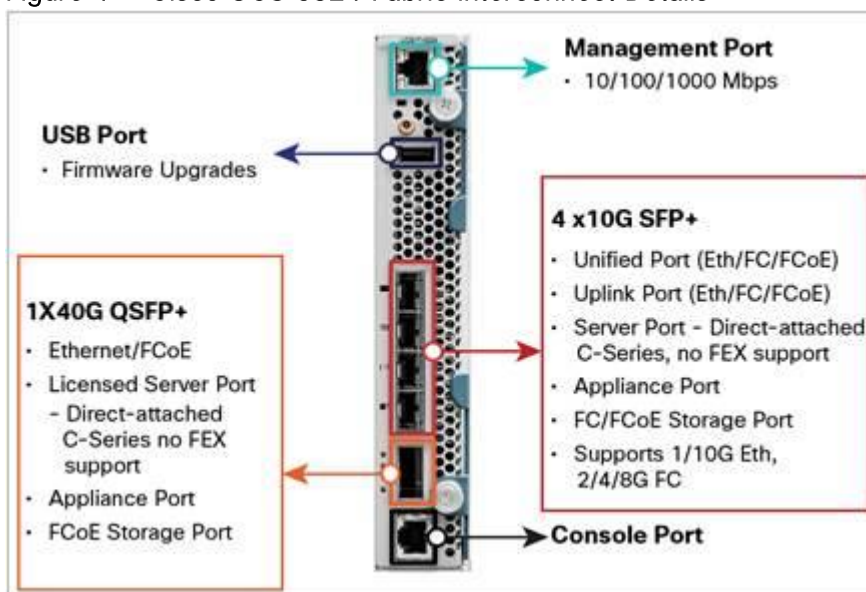
The Cisco UCS 6324 Fabric Interconnect supports the integrated Cisco UCS Management software (UCS Manager) and allows direct LAN and storage connectivity for the blade servers and directly-connected rack-mount servers in one plug-in module.

From a networking perspective, the Cisco UCS 6324 Fabric Interconnect supporting deterministic, low-latency, line-rate traffic with a maximum switching capacity of up to 500Gbps, independent of packet size and enabled services. Sixteen 10Gbps links connect to the servers, providing a 20Gbps link from each Cisco UCS 6324 Fabric Interconnect to each server.

The fabric interconnects supports multiple traffic classes over a lossless Ethernet fabric from the blade through the fabric interconnect. Significant TCO savings come from an optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

The Cisco UCS 6324 Fabric Interconnect is a 10 Gigabit Ethernet and Fibre Channel switch offering up to 500-Gbps throughputs and up to four unified ports and one scalability port.

**Figure 4 Cisco UCS 6324 Fabric Interconnect Details**



For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html>

## Cisco UCS B200 M4 Servers

The enterprise-class **Cisco UCS B200 M4 Blade Server** extends the capabilities of Cisco's Unified Computing System portfolio in a half-width blade form factor. The Cisco UCS B200 M4 uses the power of the latest Intel® Xeon® E5-2600 v3 Series processor family CPUs with up to 768 GB of RAM (using 32 GB DIMMs), two solid-state drives (SSDs) or hard disk drives (HDDs), and up to 80 Gbps throughput connectivity. The Cisco UCS B200 M4 Blade Server mounts in a Cisco UCS 5100 Series blade server chassis

or Cisco UCS Mini blade server chassis. It has 24 total slots for registered ECC DIMMs (RDIMMs) or load-reduced DIMMs (LR DIMMs) for up to 768 GB total memory capacity (B200 M4 configured with two CPUs using 32 GB DIMMs). It supports one connector for Cisco's VIC 1340 or 1240 adapter, which provides Ethernet and FCoE.

For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b200-m4-blade-server/index.html>

Figure 5 Cisco UCS B200 M4 Blade Server



## Cisco I/O Adapters

The Cisco UCS blade server has various Converged Network Adapters (CNA) options. The Cisco UCS Virtual Interface Card (VIC) 1340 used in the solution is a 2-port 40-Gbps Ethernet or dual 4 x 10-Gbps Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional port expander, the Cisco UCS VIC 1340 capabilities is enabled for two ports of 40-Gbps Ethernet.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host that can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco UCS VIC 1340 supports Cisco® Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS Fabric interconnect ports to virtual machines, simplifying server virtualization deployment and management.


Figure 6 Cisco UCS Virtual Interface Card (VIC) 1340



The personality of the card is determined dynamically at boot time using the service profile associated with the server. The number, type (NIC or HBA), identity (MAC address and World Wide Name [WWN]), failover policy, bandwidth, and quality-of-service (QoS) policies of the PCIe interfaces are all determined using the service profile. The capability to define, create, and use interfaces on demand provides a stateless and agile server infrastructure.



Figure 7 Cisco UCS Virtual Interface Card (VIC) 1340 Architecture



Each PCIe interface created on the VIC is associated with an interface on the Cisco UCS fabric interconnect, providing complete network separation for each virtual cable between a PCIe device on the VIC and the interface on the fabric interconnect.

For more information, see: <http://www.cisco.com/c/en/us/products/interfaces-modules/ucs-virtual-interface-card-1340/index.html>

## Cisco UCS Differentiators

**Cisco's Unified Compute System is revolutionizing the way servers are managed in data-center.** The following are the unique differentiators of Cisco UCS and Cisco UCS Manager.

1. **Embedded Management** –In Cisco UCS, the servers are managed by the embedded firmware in the Fabric Interconnects, eliminating need for any external physical or virtual devices to manage the servers.
2. **Unified Fabric** –In Cisco UCS, from blade server chassis or rack servers to FI, there is a single Ethernet cable used for LAN, SAN and management traffic. This converged I/O results in reduced cables, SFPs and adapters - reducing capital and operational expenses of overall solution.
3. **Auto Discovery** –By simply inserting the blade server in the chassis or connecting rack server to the fabric interconnect, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto-discovery enables the wire-once architecture of UCS, where compute capability of UCS can be extended easily while keeping the existing external connectivity to LAN, SAN and management networks.
4. **Policy Based Resource Classification** –Once a compute resource is discovered by UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is

useful in multi-tenant cloud computing. This CVD showcases the policy based resource classification of UCS Manager.

5. Combined Rack and Blade Server Management –Cisco UCS Manager can manage B-series blade servers and C-series rack server under the same UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
6. Model based Management Architecture –Cisco UCS Manager architecture and management data-base is model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of UCS Manager with other management systems.
7. Policies, Pools, Templates –The management approach in 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.
8. Loose Referential Integrity –In Cisco UCS Manager, a service profile, port profile or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy or a referred policy can be deleted even though other policies are referring to it. This provides different subject matter experts to work independently from each-other. This provides great flexibility where different experts from different domains, such as network, storage, security, server and virtualization work together to accomplish a complex task.
9. 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. A policy referring to another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with **specific name is found in the hierarchy of the root organization, then special policy named “default”** is searched. This policy resolution practice enables automation friendly management APIs and provides great flexibility to owners of different organizations.
10. 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.
11. Built-in Multi-Tenancy Support –The combination of policies, pools and templates, loose referential integrity, policy resolution in organization hierarchy and a service profiles based approach to compute resources makes UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.
12. Extended Memory – The enterprise-class Cisco UCS B200 M4 blade server extends the capabilities **of Cisco's Unified Computing System portfolio in a half-width blade form factor**. The Cisco UCS B200 M4 harnesses the power of the latest Intel® Xeon® E5-2600 v3 Series processor family CPUs with up to 1536 GB of RAM (using 64 GB DIMMs) – allowing huge VM to physical server ratio required in many deployments, or allowing large memory operations required by certain architectures like big data.
13. Virtualization Aware Network –Cisco VM-FEX technology makes the access network layer aware about host virtualization. This prevents domain pollution of compute and network domains with virtualization when virtual network is managed by port-**profiles defined by the network administrators’**



team. VM-FEX also off-loads hypervisor CPU by performing switching in the hardware, thus allowing hypervisor CPU to do more virtualization related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM and Hyper-V SR-IOV to simplify cloud management.

14. Simplified QoS –Even though Fibre Channel and Ethernet are converged in Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network Quality of Service (QoS) is simplified in UCS Manager by representing all system classes in one GUI panel.

## Cisco Nexus 9000 Series Switch

The Cisco Nexus 9000 Series Switches offer both modular (9500 switches) and fixed (9300 switches) 1/10/40/100 Gigabit Ethernet switch configurations designed to operate in one of two modes:

- Cisco NX-OS mode for traditional architectures
- ACI mode to take full advantage of the policy-driven services and infrastructure automation features of ACI

### Architectural Flexibility

- Delivers high performance and density, and energy-efficient traditional 3-tier or leaf-spine architectures
- Provides a foundation for Cisco ACI, automating application deployment and delivering simplicity, agility, and flexibility

### Scalability

- Up to 60 Tbps of non-blocking performance with less than 5-microsecond latency
- Up to 2304 10-Gbps or 576 40-Gbps non-blocking layer 2 and layer 3 Ethernet ports
- Wire-speed virtual extensible LAN (VXLAN) gateway, bridging, and routing support

### High Availability

- Full Cisco In-Service Software Upgrade (ISSU) and patching without any interruption in operation
- Fully redundant and hot-swappable components
- A mix of third-party and Cisco ASICs provide for improved reliability and performance

### Energy Efficiency

- The chassis is designed without a midplane to optimize airflow and reduce energy requirements
- The optimized design runs with fewer ASICs, resulting in lower energy use
- Efficient power supplies included in the switches are rated at 80 Plus Platinum

### Investment Protection

Cisco 40-Gb bidirectional transceiver allows for reuse of an existing 10 Gigabit Ethernet cabling plant for 40 Gigabit Ethernet

- Designed to support future ASIC generations

- Support for Cisco Nexus 2000 Series Fabric Extenders in both NX-OS and ACI mode
- Easy migration from NX-OS mode to ACI mode

The VersaStack design covered in this document uses NX-OS mode of operation using a pair of Cisco Nexus 9300 Series (Cisco Nexus 9372PX) switches. By using Cisco Nexus 9300 Series switches, it lays the foundation for migrating to ACI at a future time.

For more information, see:

<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>

## Cisco Nexus 1000v

Cisco Nexus 1000V Series Switches provide a comprehensive and extensible architectural platform for virtual machine (VM) and cloud networking. Integrated into the VMware vSphere hypervisor and fully compatible with VMware vCloud® Director, the Cisco Nexus 1000V Series provides:

- Advanced virtual machine networking based on Cisco NX-OS operating system and IEEE 802.1Q switching technology
- Cisco vPath technology for efficient and optimized integration of virtual network services
- Virtual Extensible Local Area Network (VXLAN), supporting cloud networking
- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Non-disruptive operational model for your server virtualization and networking teams
- Virtualized network services with Cisco vPath providing a single architecture for L4 -L7 network services such as load balancing, firewalling and WAN acceleration. For more information, see:

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

— <http://www.cisco.com/en/US/products/ps10785/index.html>

## IBM Storwize V5000

The storage controller leveraged for this validated design, The IBM Storwize® V5000 is a highly flexible, easy to use, virtualized storage system that enables midsized organizations to meet the challenges of rapid data growth and limited IT budgets. As an intermediate Storwize family offering, Storwize V5000 enables organizations to consolidate and provide new capabilities to their existing infrastructures. With flexible licensing options, Storwize V5000 makes it easy to deploy complementary solutions, such as external virtualization and replication. Designed to complement virtual server environments, these modular storage systems deliver the flexibility for changing business needs.

IBM Storwize® V5000 has the following host interfaces:

- SAN-attached 8 Gbps Fibre Channel or 10 Gbps iSCSI/FCoE
- 1 Gbps iSCSI
- 4-port 6 Gbps SAS host interface card

- Each IBM Storwize® V5000 node canister has up to 4GB internal cache to accelerate and optimize.

IBM Storwize® V5000 can deploy a range of Storwize software features, including:

- IBM Easy Tier for automated storage tiering
- External storage virtualization and data migration
- Synchronous data replication with Metro Mirror
- Asynchronous data replication with Global Mirror
- FlashCopy for near-instant data backups

### IBM Easy Tier

IBM Easy Tier further improves performance while increasing efficiency, by automatically identifying and moving active data to faster storage, such as flash. This means flash storage is used for the data that will benefit the most, to deliver better performance from even small amounts of flash capacity. Even in systems without flash, hot spots are automatically detected and data is redistributed, to optimize performance.

For more information, see: [http://www.ibm.com/systems/storage/disk/storwize\\_v5000/](http://www.ibm.com/systems/storage/disk/storwize_v5000/)

### VMware vSphere

VMware vSphere is the leading virtualization platform for managing pools of IT resources consisting of processing, memory, network and storage. Virtualization allows for the creation of multiple virtual machines to run in isolation, side-by-side and on the same physical host. Unlike traditional operating systems that dedicate all server resources to one instance of an application, vSphere provides a means to manage server hardware resources with greater granularity and in a dynamic manner to support multiple instances.

For more information, see:

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

## Domain and Element Management

This section of the document provides general descriptions of the domain and element managers relevant to the VersaStack:

- Cisco UCS Manager
- Cisco UCS Central
- IBM Storwize V5000 Unified management GUI
- **VMware vCenter™ Server**

### Cisco Unified Computing System Manager

Cisco Unified Computing System (UCS) Manager provides unified, embedded management for all software and hardware components in the 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 command-line interface (CLI), or an XML API. Cisco UCS Manager offers unified embedded 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 a comprehensive XML API for third part integration, exposes 9000 points of integration and facilitates custom development for automation, orchestration, and to achieve new levels of system visibility and control.

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

Within the Cisco UCS Mini, Cisco UCS 6324 Fabric Interconnect hosts and runs Cisco UCS Manager in a highly available configuration, enabling the fabric interconnects to fully manage all Cisco UCS elements. The Cisco UCS 6324 Fabric Interconnects support out-of-band management through dedicated 10/100/1000-Mbps Ethernet management ports. Cisco UCS Manager typically is deployed in a clustered active-passive configuration with two Cisco UCS 6324 Fabric Interconnects connected through the cluster interconnect built into the chassis.

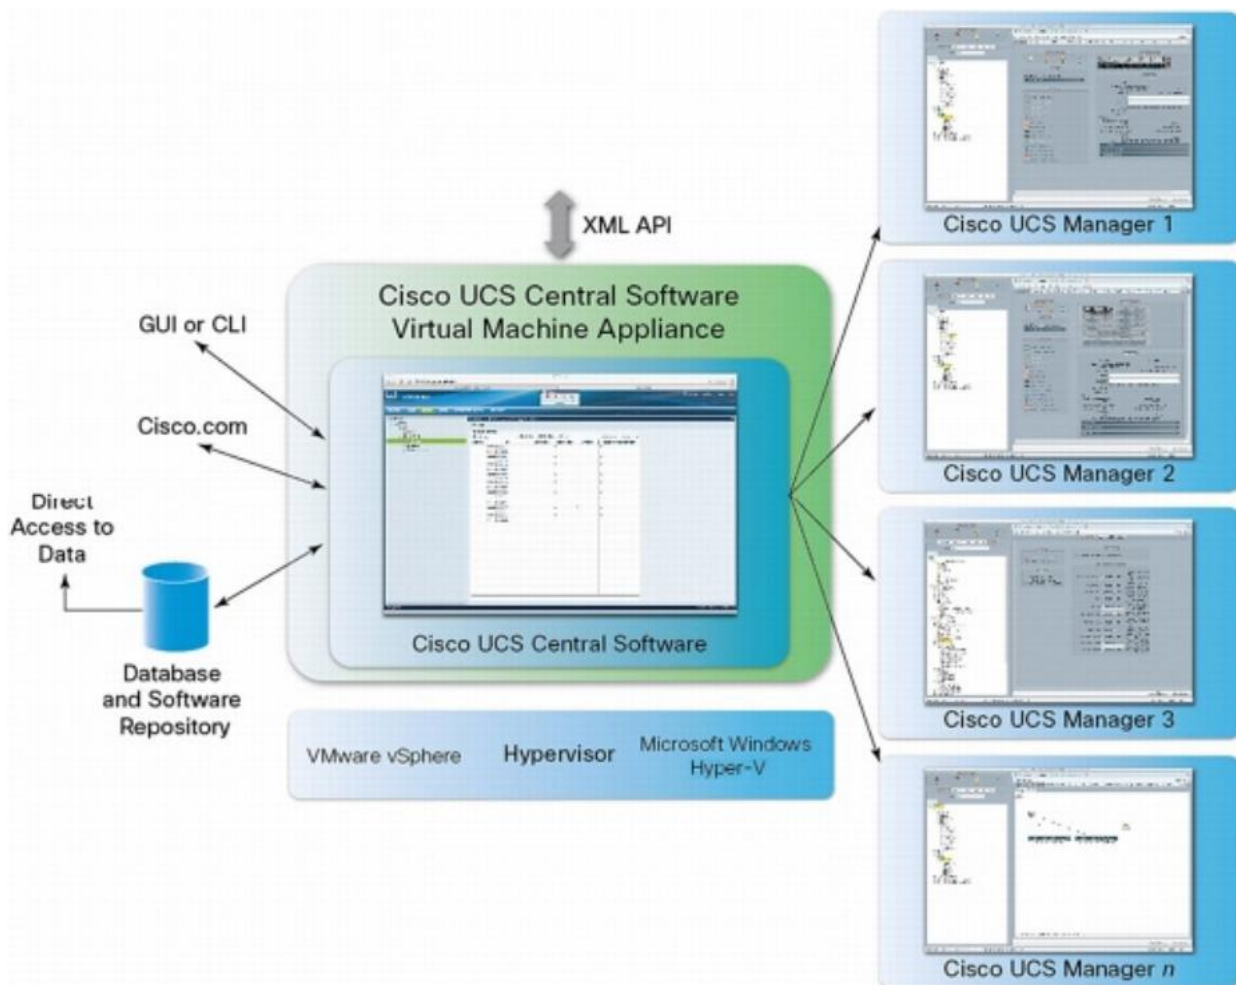
Cisco UCS Manager supports the 6324 Fabric Interconnect that integrates the FI into the Cisco UCS Chassis and provides an integrated solution for a smaller deployment environment. Cisco UCS Mini simplifies the system management and saves cost for smaller scale deployments. The hardware and software components support Cisco unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

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

## Cisco UCS Central (Optional)

Cisco UCS Central is an optional component in the VersaStack design which provides Cisco UCS customers managing growth within a single data center, growth across multiple sites, or both, Cisco UCS Central Software centrally manages multiple Cisco UCS domains using the same concepts that Cisco UCS Manager uses to support a single Cisco UCS domain. Cisco UCS Central Software manages global resources (including identifiers and policies) that can be consumed within individual Cisco UCS Manager instances. It can delegate the application of policies (embodied in global service profiles) to individual Cisco UCS domains, where Cisco UCS Manager puts the policies into effect. Cisco UCS Central software manages multiple, globally distributed Cisco UCS domains from a single pane. Every instance of Cisco UCS Manager and all of the components managed by it form a domain. Cisco UCS Central integrates with Cisco UCS Manager, and utilizes it to provide global configuration capabilities for pools, policies, and firmware.

Figure 8 Cisco UCS Central Software Architecture



Cisco UCS Central Software makes global policy and compliance easier. When Cisco UCS domains are registered with Cisco UCS Central Software, they can be configured to automatically inherit global identifiers and policies that are centrally defined and managed. Making identifiers such as universal unique identifiers (UUIDs), MAC addresses, and worldwide names (WWNs) global resources allows every server worldwide to be configured uniquely so that identifier conflicts are automatically avoided. Globally defined policies take this concept significantly further: defining and enforcing server identity, configuration, and connectivity policies centrally essentially ensure standards compliance. The system simply will not configure a server in a way that is inconsistent with standards, so configuration drift and an entire class of errors that can cause downtime are avoided.

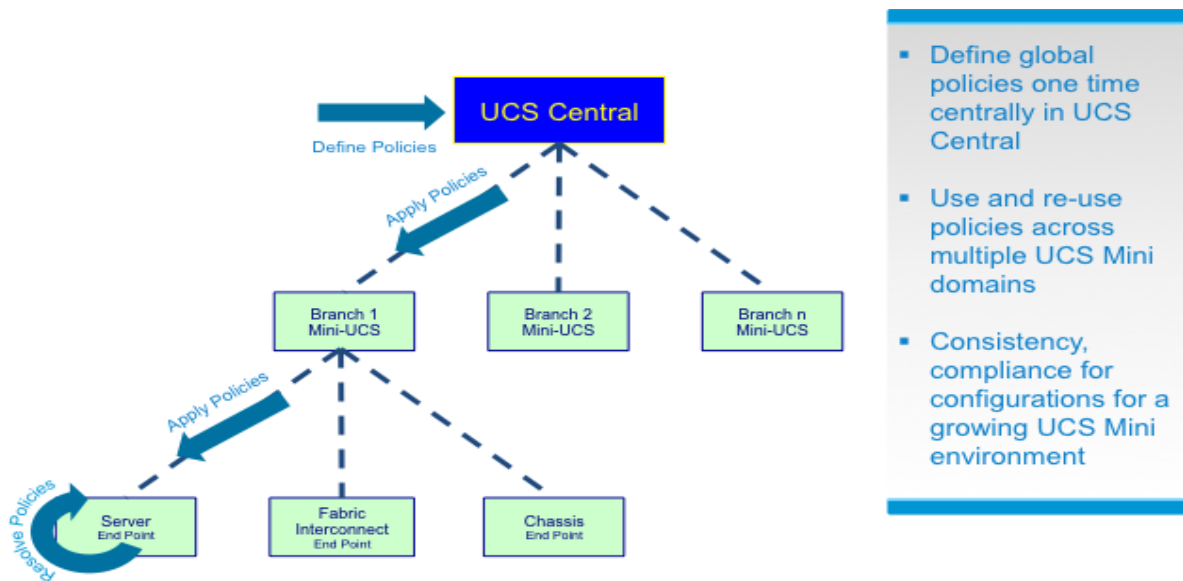
### Cisco UCS Manager (Mini) Management with Cisco UCS Central

Cisco UCS Central Software is designed and operates comparable to Cisco UCS Manager in that policies and configuration definitions, which make up a Cisco UCS service profile, can be created at a central location and then applied to the endpoint recipient, where they are resolved. With Cisco UCS Manager, the endpoint recipients are the Cisco UCS infrastructure (servers, network, etc.).

For Cisco UCS Central Software, the recipients are individual Cisco UCS Manager Instances that have been registered with Cisco UCS Central Software. With Cisco UCS Central Software, global Cisco UCS service

profiles are defined centrally and are passed to Cisco UCS Manager instances according to the way they are registered with Cisco UCS Central Software (Figure 9).

Figure 9 Cisco UCS Manager (Mini) Management with Cisco UCS Central



33

### IBM Storwize® V5000 Easy-to-Use Management GUI

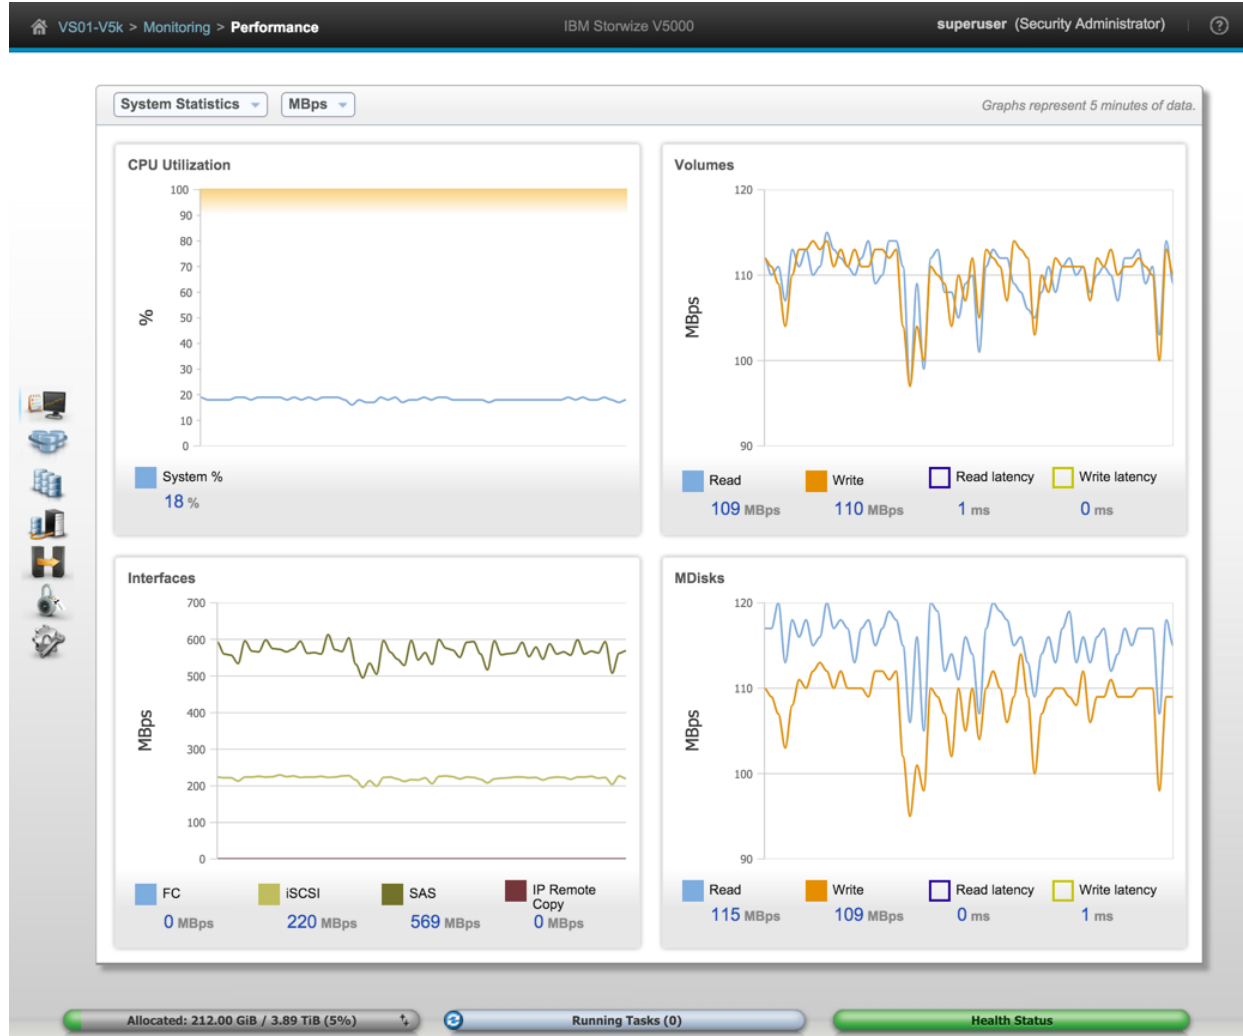
The IBM Storwize V5000 Unified built-in user interface (Figure 10) hides complexity and makes it possible for administrators to quickly and easily complete common block and file storage tasks from the same interface, such as creating and deploying volumes, file sets and host mappings. Users can also monitor performance in real-time (Figure 11).

The V5000 Unified management interface has the ability to check for the latest updates, and via an upgrade wizard, keep you running the latest software release with just a few mouse clicks. The interface provides auto-discovery and presets that help the admin greatly reduce setup time as well as help them easily implement a successful deployment. The interface is web-accessible and built into the product, removing the need for the administrator to download and update management software.

Figure 10 IBM Storwize® V5000 Management GUI Example



Figure 11 Real-time Performance Monitoring on the IBM Storwize® V5000 Management GUI



## VMware vCenter Server

VMware vCenter is a virtualization management application for managing large collections of IT infrastructure resources such as processing, storage and networking in a seamless, versatile and dynamic manner. It is the simplest and most efficient way to manage VMware vSphere hosts at scale. It provides unified management of all hosts and virtual machines from a single console and aggregates performance monitoring of clusters, hosts, and virtual machines. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, virtual machines, storage, the guest OS, and other critical components of a virtual infrastructure. A single administrator can manage 100 or more virtualization environment workloads using VMware vCenter Server, more than doubling typical productivity in managing physical infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

For more information, see: <http://www.vmware.com/products/vcenter-server/overview.html>

## Cisco Virtual Switch Update Manager

The Cisco Nexus 1000V virtual switch is a software-based Layer 2 switch for VMware ESX virtualized server environments. The Cisco Nexus 1000V provides a consistent networking experience across both physical



and the virtual environments. It consists of two components: the Virtual Ethernet Module (VEM), a software switch that is embedded in the hypervisor, and a Virtual Supervisor Module (VSM), a module that manages the networking policies and the quality of service for the virtual machines.

Cisco Virtual Switch Update Manager (Cisco VSUM) enables you to install, upgrade, and monitor the Cisco Nexus 1000V for VMware vSphere and also migrate hosts to the Cisco Nexus 1000V, using the VMware vSphere Web Client.

Cisco VSUM enables you to do the following:

- Install the Cisco Nexus 1000V switch.
- Migrate the VMware vSwitch and VMware vSphere Distributed Switch (VDS) to the Cisco Nexus 1000V.
- Monitor the Cisco Nexus 1000V.
- Upgrade the Cisco Nexus 1000V and added hosts from an earlier version to the latest version.
- Install the Cisco Nexus 1000V license.
- View the health of the virtual machines in your data center using the Dashboard - Cisco Nexus 1000V.

## IBM Spectrum Protect Snapshot

IBM Spectrum Protect Snapshot delivers high levels of protection for key applications and databases using advanced integrated application snapshot backup and restore capabilities.

It lets you perform and manage frequent, near-instant, non-disruptive, application-aware backups and restores using integrated application and VM-aware snapshot technologies.

For more information, see:

<http://www-03.ibm.com/software/products/en/spectrum-protect-snapshot>

## Solution Design

This VersaStack Solution was designed to address the IT infrastructure requirements of mid-sized enterprises, smaller IT and ROBO (remote office and branch office) environments. These deployments support a diverse set of applications and typically have access to a very limited IT staff and budget. As such, simplicity and ease of deployment and management are critical factors when selecting an infrastructure platform to address IT needs.

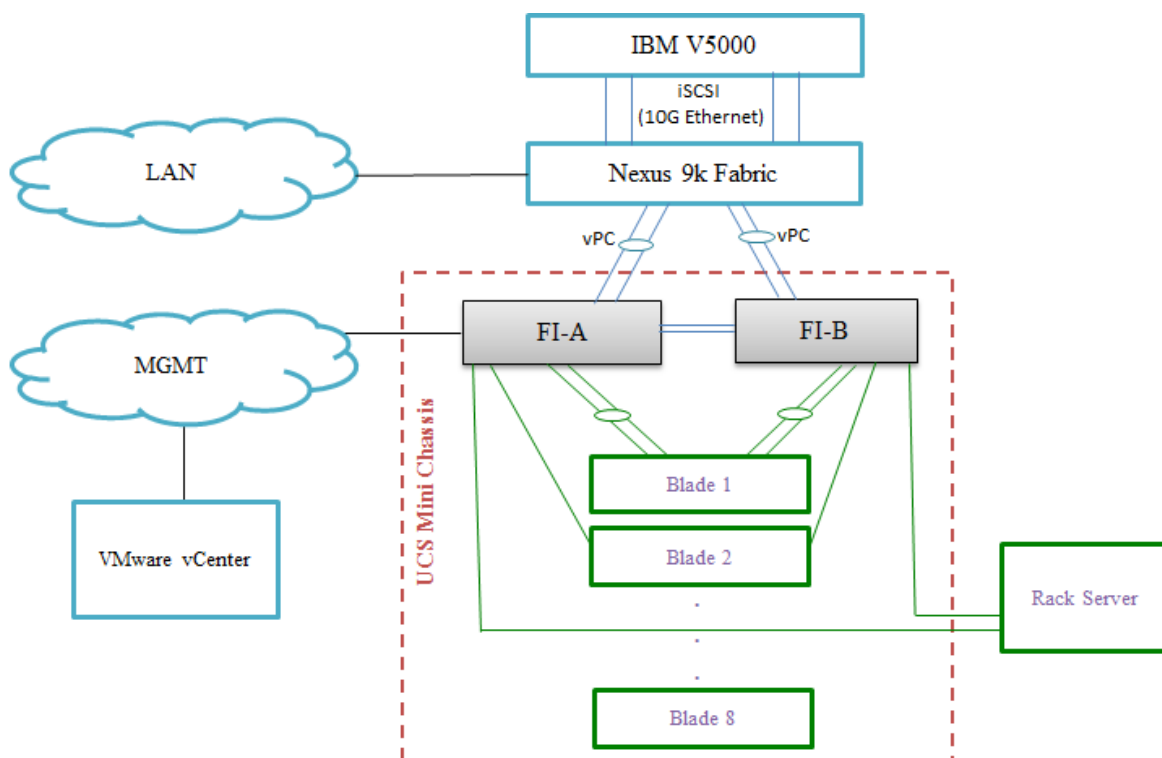
The VersaStack design uses the compact version of Cisco UCS (Cisco UCS Mini) with Cisco B200M4 half-width blades running ESXi 5.5 U2 and IBM Storwize V5000 to enable organizations meet the challenges of rapid growth and limited IT budgets. This design enables organizations to consolidate and provide new capabilities to their existing infrastructures, with flexible deployment options.

The VersaStack Models with Cisco UCS Mini are characterized by the IBM Storwize V5000 connection points within the architecture. Currently, these models include:

- IP-based Storage
- Direct Attached SAN Storage

Figure 12 illustrates the fully redundant VersaStack IP-Based Storage logical topology in the compute, network, and storage layers.

Figure 12 VersaStack IP-Based Storage Logical Topology

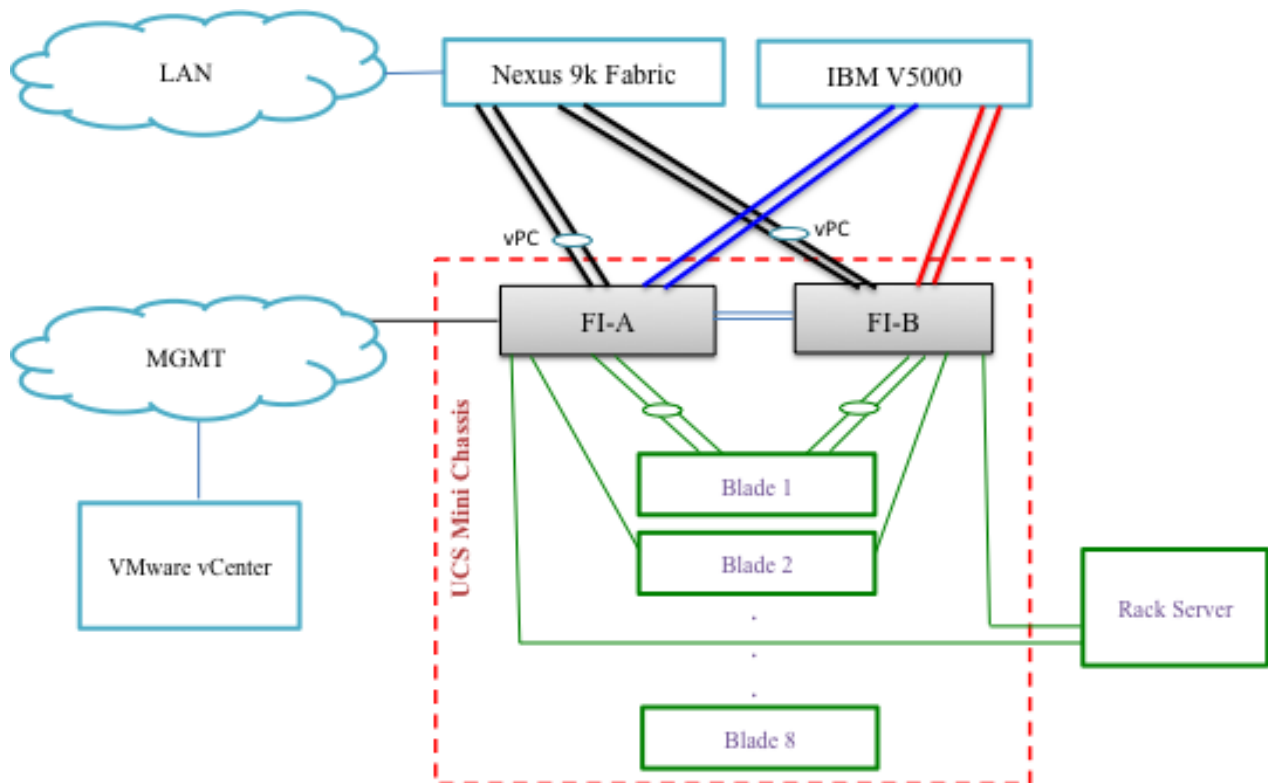


The Network fabric with in the solution consists of two Cisco Nexus 9372PX switches deployment for high availability and provides 10G enabled, 40G capable network fabric. Link aggregation using virtual PortChannels (vPC) is used in this design to provide higher aggregate bandwidth and fault tolerance. Cisco Nexus 9000 platforms support link aggregation using 802.3ad standard Link Aggregation Control Protocol (LACP). Virtual PortChannels allow links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single logical link to a third device - Cisco UCS Mini in this case. This provides device-level redundancy and connectivity even if one of the Cisco Nexus switches fail. It also provides a loop-free topology without blocked ports that typically occurs with spanning tree, enabling all available uplink bandwidth to be used, thereby increasing the aggregate bandwidth into the Cisco UCS domain. **These links are used to provide connectivity to the rest of the customer’s network and carry storage traffic.**

Cisco Nexus 9000 family of switches supports two modes of operation: NX-OS standalone mode and ACI mode. NX-OS standalone mode is used in this VersaStack design. Cisco Nexus 9000 switches have capabilities and performance necessary for medium-size businesses and Enterprises without having to upgrade the networking infrastructure as the networking needs grow. Cisco Nexus switches provides 40G connectivity at low latency and high port-density. Cisco Nexus 9300 series switches used in this design also provides investment protection by providing the foundation for migrating to ACI with centralized policy based management.

Figure 13 illustrates the fully redundant VersaStack Direct Attached Storage topology in the compute, network, and storage layers.

Figure 13 VersaStack Direct Attached Storage Logical Topology

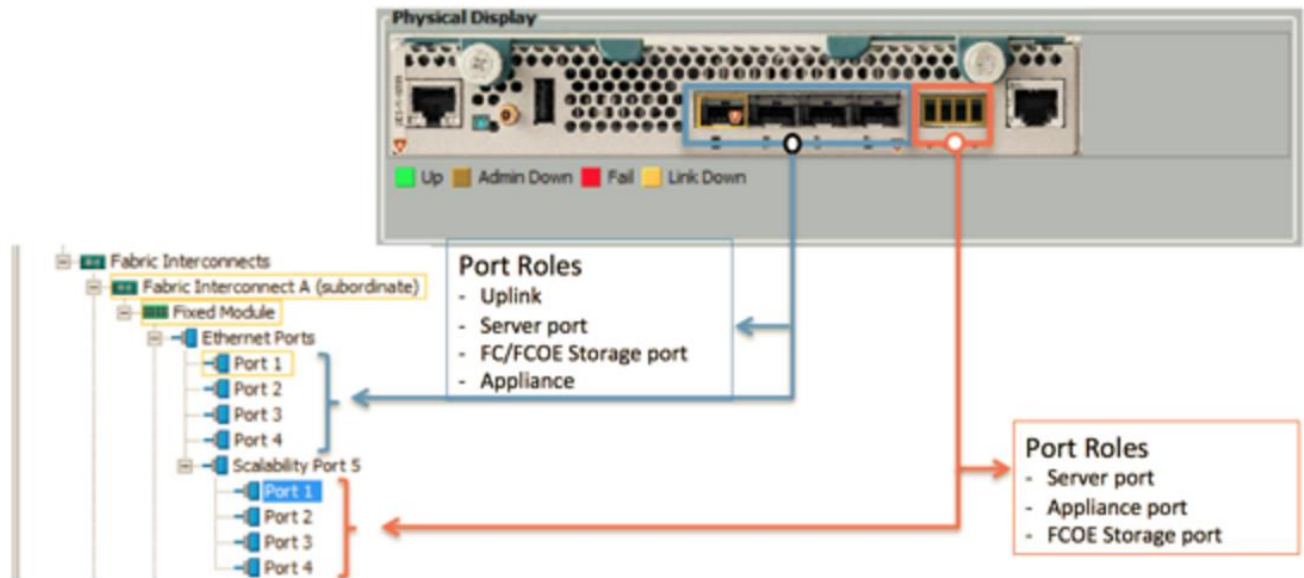


As stated earlier, Cisco UCS Mini has built-in redundant fabric interconnects (FI) providing unified network and storage access. Each Cisco UCS fabric interconnect has 4x10Gbps ports and a 40 Gbps scalability port.

2x10Gbps ports will be used as Uplink ports and connect to the switching infrastructure supporting applications. The other 10Gbps pair is used for the user and storage traffic in the IP-Based storage design as illustrated in Figure 12. The remaining ports can be configured as Appliance ports for direct storage connectivity as an alternate design option with Direct Attached storage as illustrated in Figure 13. The number of rack servers supported will be affected if storage is directly attached to the fabric interconnects.

The 40 Gbps (Scalability Port) can be used with a breakout cable that transforms the port to 4x10Gbps cables. These cables are connected to Cisco UCS C-Series Rack servers within the architecture; however, they can be used for server/storage connections.

Figure 14 Ethernet and Scalability Ports in Cisco UCS Built-in FI



## Storage Design

Within IBM Storwize V5000, there are one or two pairs of node canisters, which are known as I/O groups. The IBM Storwize V5000 therefore supports four node canisters in a clustered system, which provides two I/O groups.

When a host server performs I/O to one of its volumes, all the I/O for that volume is directed to the I/O group where the volume has been defined. Under normal conditions, the same node within that I/O GROUP also always processes these I/Os.

Both nodes of the I/O group act as preferred nodes for their own specific subset of the total number of volumes that the I/O group presents to the host servers (a maximum of 2048 volumes per hosts). However, both nodes also act as a failover node for the partner node within the I/O group. Therefore, a node takes **over the I/O workload from its partner node (if required) without affecting the server's application.**

The active/active architecture provides availability to process I/Os for both controller nodes and allows the application to continue running smoothly, even if the server has only one access route or path to the storage controller. This type of architecture eliminates the path/LUN thrashing that is typical of active/passive architecture.

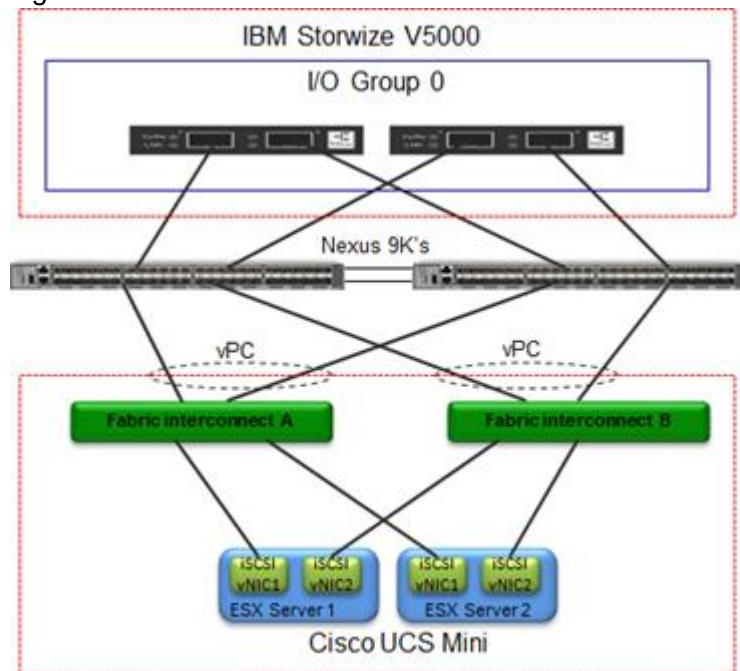
ESXi hosts in the Cisco UCS Mini use iSCSI LUNs in the IP-Based storage design and FC LUNs in the Direct Attached SAN storage design to SAN boot and do not require local disks. A volume in the Storwize V5000 I/O group appears as a single device to the applications on hosts, even though the multipathing driver can detect four separate devices for each volume. The multipathing driver selects one or more of these devices during I/O. If the connection between the host and one Storwize V5000 node port is lost, the multipathing driver can select an alternative path to the Storwize V5000 I/O group. The I/O between the host and Storwize V5000 continue without error.

## IP-Based SAN Storage using iSCSI

This design is an end-to-end IP-Based storage solution that supports SAN access using iSCSI. The design **doesn't employ a dedicated SAN switching environment and requires no direct Fibre Channel connectivity**, as iSCSI is the SAN protocol leveraged. The storage design uses IBM Storwize V5000 array to provide block storage using iSCSI. Two-node IBM Storwize V5000 system is redundantly connected to the Cisco Nexus fabric using four 10 Gbps links, with each node having connectivity to both the switches. Each node has two 10G Ethernet ports, each of which is used for iSCSI data transfers. The design uses two distinct IP domains with distinct iSCSI targets and uses multipathing to provide storage high availability.

The design is highly available with no single point of failure between the compute and storage subsystems. The redundancy between Cisco UCS Mini and Storwize V5000 array are shown in Figure 15. The redundant controllers on the Storwize are connected to Cisco Nexus 9000 redundant switches using the cabling configuration show below. Virtual PortChannels are not used in the connectivity between Storwize V5000 and Cisco Nexus 9000 series.

Figure 15 Storwize V5000 Active-Active Architecture with IP-Based Storage

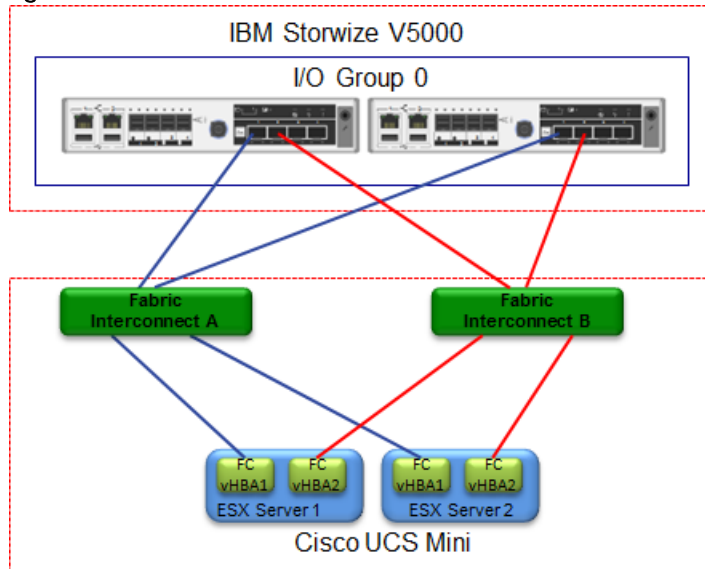


## Direct Attached SAN Storage using Fibre Channel

For smaller FC based deployments VersaStack can be deployed in a direct attached configuration to eliminate the requirement for separate Fibre Channel switches.

In this type of deployment, the flexibility of the Fabric Interconnects is leveraged and they are changed to FC Switching Mode. With the introduction of firmware version 2.1(1a), Cisco has introduced Cisco UCS Manager based FC zoning for direct connected topologies. This eliminates the need to have an upstream switch host the FC zoning database as the zoning can be handled within the Fabric Interconnects. The FC zoning is automated through Cisco UCS SAN Connectivity policies and applied on the Fabric Interconnects to simplify deployment of new servers.

**Figure 16 Storwize V5000 Active-Active Architecture with Direct Attached SAN Storage**



There are two SAN fabrics created; one each on the Cisco UCS fabric interconnects. The IBM Storwize V5000 system node canisters have four 8Gbps fibre channel ports, two ports on each canister are used and are connected to Fabric-A and Fabric-B respectively. Each ESXi host from the Cisco UCS mini is also connected to both the fabrics. This provides High availability if a SAN fabric fault or a node failure occurs with the help of multipathing software as discussed earlier.

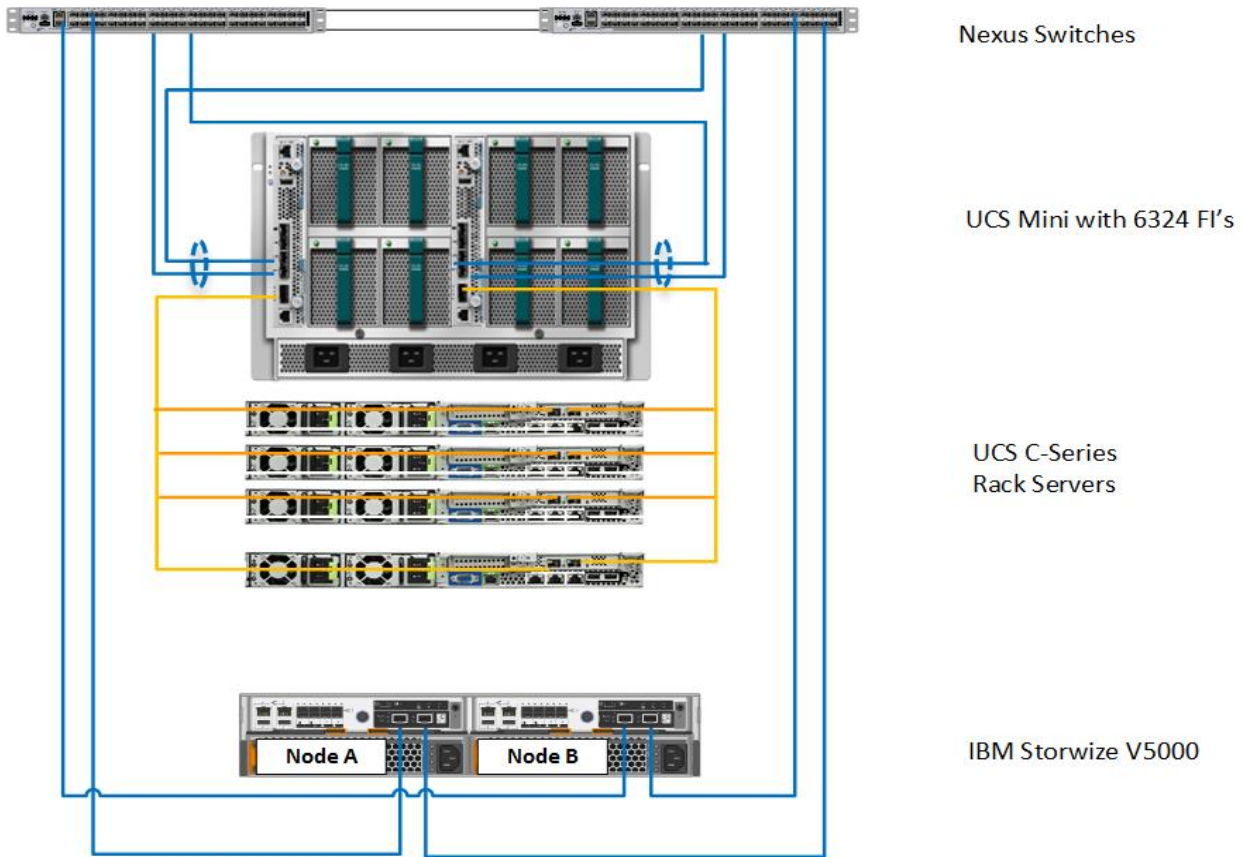
## Physical Topology

A detailed topology of the VersaStack architecture used in the validation is shown in the figures below.

For more information, see the VersaStack guides at:

<http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/versastack-designs.html>

Figure 17 VersaStack IP-Based Storage Physical Topology



Nexus Switches

UCS Mini with 6324 FI's

UCS C-Series Rack Servers

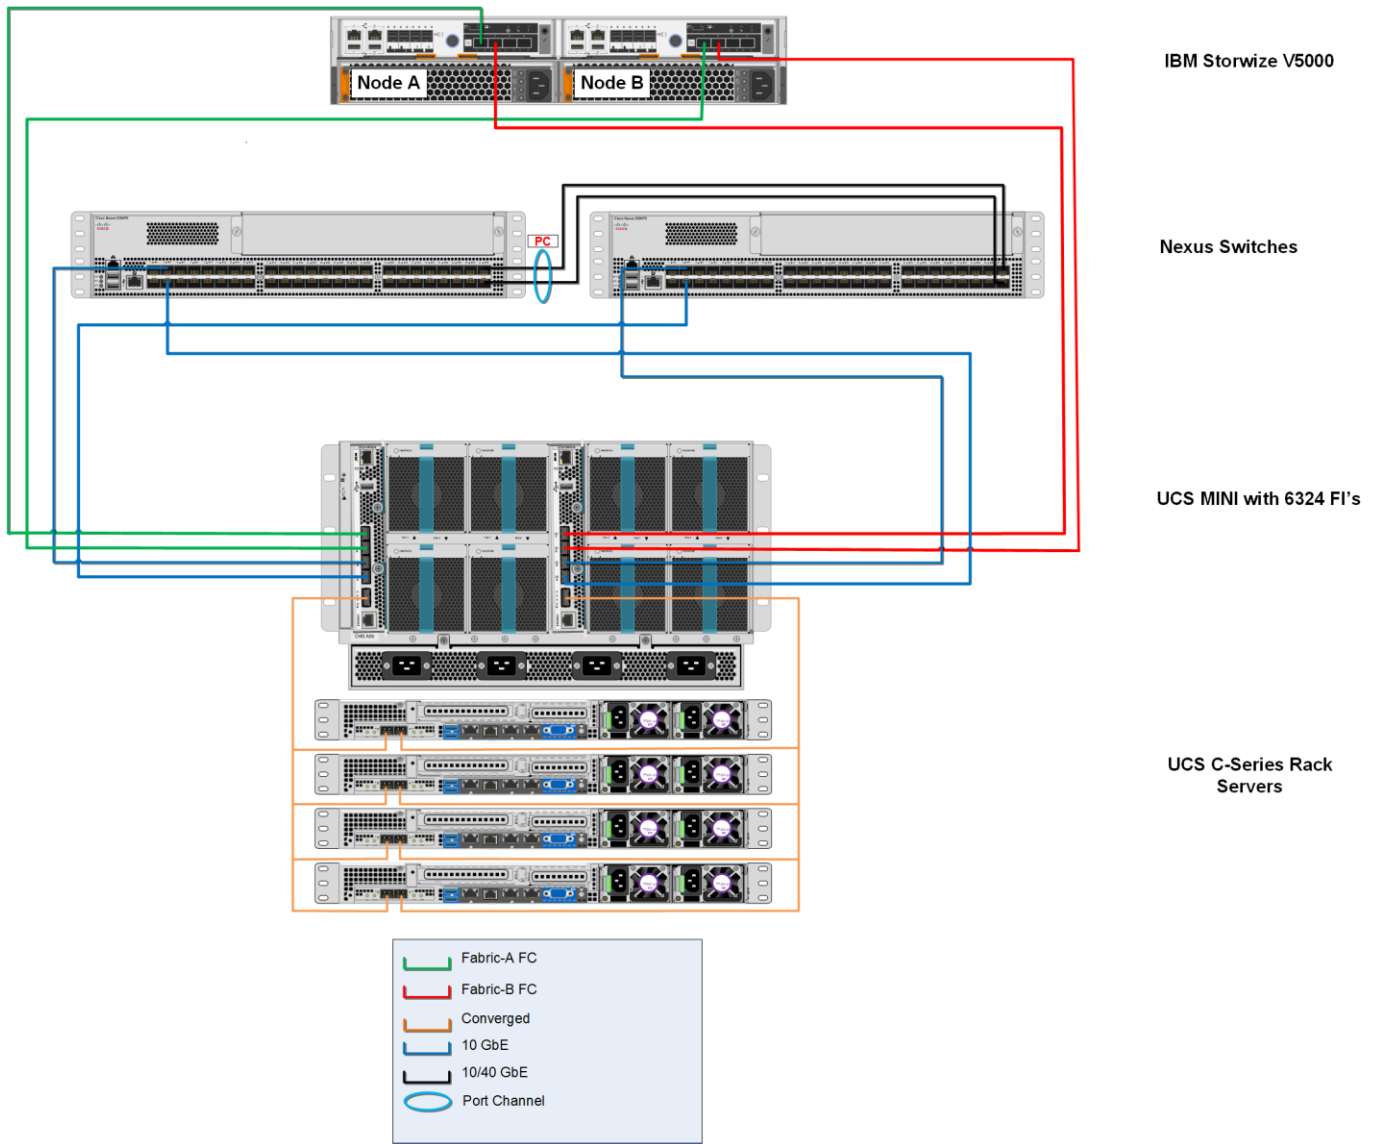
IBM Storwize V5000

*Legend*

— 10GbE    - - - vPC    — Converged



Figure 18 VersaStack Direct Attached SAN Storage Physical Topology



## Design Considerations

### Storage

- Storwize V5000 was deployed with redundant storage controllers, power supplies, cabling and paths to each of the node canisters within the controller
- Host system multipathing and failover support
- The RAID capabilities of the underlying disk subsystem.
- Dual Subnet configuration for storage connectivity to Cisco UCS Mini Fabric Interconnects
- Ten Gigabit Ethernet data path connectivity from each node canister to each Fabric Interconnect and 8GB FC connectivity in the direct attached design.
- Backup and Recovery capability via snapshots



- Easy Tier to seamlessly migrate extents to the most appropriate tier

### Storage – Pools

With the use of IBM EasyTier technology, we can create a hybrid pool and use EasyTier to move commonly accessed data to the fastest storage. Alternatively, we could create single-tier type pools.

### Storage – iSCSI boot

Each Cisco UCS blade was deployed using iSCSI SAN boot. Each Cisco UCS blade used a single Service Profile level IQN for all connections. Each blade had a boot volume created on the Storwize V5000 Storage array. The Storwize V5000 Storage array provides an initiator group to only honor connections from this single service profile. During iSCSI SAN boot connectivity, the blade connects to both a primary and secondary target. This provides for normal boot operations even when the primary path is offline. The host software utilized MPIO and the Storwize software assisted with iSCSI session and path management. Also the VMware hosts in question were deployed in a cluster to allow for HA failover and to avoid a single point of failure at the hypervisor layer.

### Storage – FC SAN boot

Each Cisco UCS blade used a single Service Profile **with two vHBA's having connectivity to SAN fabric-A** and SAN Fabric-B. Each vHBA is zoned to a different set of Storwize V5000 system ports to maximize performance and redundancy. Each blade had a boot volume created on the Storwize V5000 Storage array. The Storwize V5000 Storage array provides LUN masking to only honor connections from the assigned Host. During SAN boot connectivity; the blade connects to both a primary and secondary target. This provides for normal boot operations even when the primary path is offline. The ESX hosts were deployed in a cluster similar to iSCSI design providing HA failover without single point of failure at hypervisor layer and MPIO host software was utilized for path management.

## Cisco UCS Server Configuration for vSphere

Each Cisco UCS Server running ESXi 5.5 U2 was deployed using Cisco UCS 1340 VIC network adapter. At the server level, each Cisco VIC presents multiple vPCIe devices to ESXi node which vSphere identifies as vmnics. ESXi is unaware that these NICs are virtual adapters. In the VersaStack design, the following virtual NICs were used:

### IP-Based Storage

- One vNIC carries isolated iSCSI-A traffic to FI-A
- One vNIC carries isolated iSCSI-B traffic to FI-B
- One vNIC carry data and management traffic to FI-A
- One vNIC carry data and management traffic to FI-B
- Direct Attached SAN Storage
- One vHBA carries isolated SAN-A traffic to FI-A
- One vHBA carries isolated SAN-B traffic to FI-B
- One vNIC carry data and management traffic to FI-A
- One vNIC carry data and management traffic to FI-B

## Jumbo Frames

VersaStack solution was configured for jumbo frames with an MTU size of 9000 Bytes on the storage network links between the Cisco UCS Mini and Storwize V5000 array and also for the NFS Storage access and vMotion traffic. Though NFS storage is not supported in this VersaStack design, NFS links have been created to allow access to existing NFS shares in the environment. The jumbo MTU size is important so that larger frames can be sent and received on the wire, which reduces the CPU load and a more efficient use of the available resources. Jumbo frames were enabled at the NIC and virtual switch level.

## Management Network

A separate out-of-band management network was used for configuring and managing compute, storage and network infrastructure components in the solution. Management ports on each Storwize V5000 and Cisco UCS Mini FI were physically connected to a separate dedicated management switch. Management ports on Cisco Nexus 9372PXs were also connected into the same management switch.

Access to vCenter and ESXi hosts were done in-band. Out-of-band access to these components can be enabled, but would require additional ports on the 6324 Fabric Interconnects. A disjoint layer-2 configuration can then be used to keep the management and data plane networks completely separate. This would require 2 additional vNICs (e.g. OOB-Mgmt-A, OOB-Mgmt-B) on each server, which are then associated with the management uplink ports.

## High Availability

VersaStack platform was designed for maximum availability of the complete infrastructure (compute, network, storage, and virtualization) with no single points of failure.

## Compute and Virtualization

- Cisco UCS mini server is highly redundant with redundant power supplies and redundant 6324 fabric interconnects.
- NIC failover between Cisco UCS Fabric Interconnects enabled through Cisco UCS Manager. This is done for all management and virtual machine vNICs.
- VMware vCenter deployed with VMware HA and DRS enabled clusters
- Cisco UCS B200 M4 servers deployed in an N+1 configuration in all management and application VM clusters to provide backup in the event of an ESXi host failure – up to one host failure per cluster is supported.
- VMware vMotion was enabled and VMware HA was enabled to auto restart VMs in the event of a host/server failure in the cluster
- Host Monitoring enabled to monitor heartbeats of all ESXi hosts in the cluster to ensure quick detection in the event of an ESXi host failure
- Admission Control was enabled to ensure the cluster has enough resources to accommodate a single host failure



NIC Failover should be disabled for iSCSI NICs and for the NICs connected to Cisco Nexus 1000v.

---

## Network

- Virtual PortChannel (vPC) or link-aggregation capabilities of the Cisco Nexus 9000 family of switches were used for the network connectivity between Cisco UCS Mini and the customer's network. vPC provides Layer 2 multipathing by allowing multiple parallel paths between nodes with load balancing that result in increased bandwidth and redundancy. A vPC based architecture is therefore highly resilient and robust that scales the available Layer 2 bandwidth by using all available links. Technical benefits of vPC include:
  - Allows a single device to use a PortChannel across two upstream devices
  - Eliminates Spanning Tree Protocol blocked ports
  - Provides a loop-free topology
  - Uses all available uplink bandwidth
  - Provides fast convergence if either the link or a device fails
  - Provides link-level resiliency
  - Helps ensure high availability
- The following vPC and Cisco Nexus best practices were implemented:
  - Spanning tree port type **'edge trunk'** used on ports connected to hosts on Cisco UCS Mini
  - **Spanning tree Bridge Protocol Data Unit (BPDU) Guard and filter** enabled on **'edge trunk'** ports
  - All criteria for vPC consistency checks implemented
  - Link Aggregation Control Protocol (LACP) used on all vPC PortChannels
  - **Unique vPC Domain ID used (same on both peers) with a lower 'role priority'** on the primary switch
  - Same vPC ID and PortChannel ID used for ease of configuration and troubleshooting
  - vPC IP Address Resolution Protocol (ARP) synchronization enabled on all vPC peers
  - vPC auto-recovery feature enabled
  - Bridge Assurance only on vPC Peer links - enabled by default
  - Loopguard - disabled by default
  - CDP enabled for infrastructure visibility and ease of troubleshooting

## Scale and Performance

### Compute

The VersaStack design with Cisco UCS Mini and IP-based storage can support a maximum of eight half-width blade servers within the Cisco UCS 5108 chassis. With this IP-based model there are two 10 Gbps SFP+ ports available to support two rack mount Cisco UCS C-series servers. Additionally, the Cisco UCS

Mini can support up to four Cisco UCS rack mount servers using the 40 GbE Enhanced Quad SFP (QSFP+) ports on the Cisco UCS 6324 Fabric Interconnects.

Figure 19 shows the maximum compute scale possible in a VersaStack design with IP-based storage. Assuming that two ports on each Cisco UCS 6324 Fabric Interconnect are used for Ethernet traffic, the maximum compute scale possible with this design is a fourteen; eight half-width Cisco UCS B-series blade servers and six Cisco UCS C-series rack mount servers.

Figure 19 Compute Scalability with IP-Based Storage





Figure 20 shows the maximum compute scale possible in a VersaStack design with Direct Attached SAN storage. Assuming that two ports on each Cisco UCS 6324 Fabric Interconnect are used for Ethernet traffic and two more ports used for SAN storage connectivity, the maximum compute scale possible with this design is a twelve; eight half-width Cisco UCS B-series blade servers and four Cisco UCS C-series rack mount servers.

Figure 20 Compute Scalability with Direct Attached SAN Storage




## Storage

The VersaStack design provides flexible Storage scaling with the Storwize V5000 array. The Storwize V5000 control enclosures provide 6 GB SAS internal dual-port drive connectivity, and each control enclosure also has four 6 GB SAS x4 (Mini-SAS HD SFF-8644) ports (2 ports per node canister) for 6 GB SAS expansion enclosure connectivity.

The Storwize V5000 Control Enclosure supports attachment of up to 19 Storwize V5000 Large Form Factor (LFF) or 19 Small Form Factor (SFF) expansion enclosures with a maximum of 480 drives providing up to 1.92 PB of storage. Intermix of LFF and SFF enclosures are supported. The expansion enclosures can be added to the system non-disruptively.

Each Storwize V5000 expansion unit ships with two expansion canisters. Each expansion canister provides 6 Gb SAS connectivity to the internal drives and two external 6 Gb SAS x4 ports (Mini-SAS HD SFF-8644 connectors labelled Port 1 and Port 2) that are used for connections to the Storwize V5000 node canisters and for connecting the expansion units between each other. One of the expansion ports (Port 3 or Port 4) on the Storwize V5000 node canister is connected to the Port 1 on the expansion canister. The Port 2 on the expansion canister is connected to the Port 1 on the expansion canister in the adjacent enclosure, and so on.

Figure 21 Storage Scalability



Up to 10 expansion enclosures can be connected to Port 3 in a daisy-chained manner, and up to nine expansion enclosures can be connected to Port 4 in a daisy-chained manner, for a total of 19 expansion enclosures per one control enclosure.

## Validation

### Hardware and Software Revisions

Error! Reference source not found. lists the software versions used during solution validation.

It is important to note that Cisco, IBM, and VMware have interoperability matrices that should be referenced to determine support for any specific implementation of VersaStack. See the following links for more information:

- [IBM System Storage Interoperation Center](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)
- [VMware Compatibility Guide](#)

**Table 2 Validated Software Versions**

Layer	Device	Version or Release	Details
Compute	Cisco UCS fabric interconnect	3.0(2d) & 3.0(2e)	Embedded management
	Cisco UCS C 220 M3/M4	3.0(2d) & 3.0(2e)	Software bundle release
	Cisco UCS B 200 M3/ M4	3.0(2d) & 3.0(2e)	Software bundle release
	Cisco eNIC	2.1.2.69	Ethernet driver for Cisco VIC
	Cisco fNIC	1.6.0.16	FCoE driver for Cisco VIC
Network	Cisco Nexus 9372PX	6.1(2)I3(4b)	Operating system version
Storage	IBM Storwize V5000	7.4.0.6	Software version
Software	Cisco UCS hosts	VMware vSphere ESXi™ 5.5u2	Operating system version
	VMware vCenter™	5.5u2	VM (1 each): VMware vCenter
	Cisco Nexus 1000v	5.2(1)SV3(1.4)	Software version
	Virtual Switch	1.5	Virtual Switch



	Update Manager (VSUM)		Deployment Software
--	--------------------------	--	------------------------

## Hardware and Software Options for VersaStack

While VersaStack Deployment CVDs are configured with specific hardware and software, the components used to deploy a VersaStack can be customized to suit the specific needs of the environment as long as all the components and operating systems are on the HCL lists referenced in this document. VersaStack can be deployed with all advanced software features such as replication and storage virtualization on any component running support levels of code. It is recommended to use the software versions specified in the deployment CVD when possible. Other operating systems such as Linux and Windows are also supported on VersaStack either as a Hypervisor, guest OS within the hypervisor environment, or directly installed onto bare metal servers. Note that basic networking components such as IP only switches are typically not on the IBM HCL. See the Table 3 for examples of additional hardware options.

Interoperability links:

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>

**Table 3 Examples of Other VersaStack Hardware and Software Options**

Layer	Hardware	Software
Compute	Rack C220,C240,C460 (M3/M4)	2.2(3b) or later
	5108 chassis	2.2(3b) or later
	Blade: B200,B420,B260,460 (M3/M4)	2.2(3b) or later
Network	93XX, 55XX, 56XX, 35XX series	6.1(2)I3(1) or later
	Cisco Nexus 1000v	5.2(1)SV3(1.1) or later
	Cisco UCS VIC 12XX series ,13XX series	
	Cisco MDS 91XX, 92XX, 95XX, 97XX series	6.2(9) or later
Storage	IBM Storwize V7000 , IBM FlashSystem	Version 7.3.0.9



Layer	Hardware	Software
	V9000	or later
	Large Form Factor Expansion (2076-12F) Small Form Factor Expansion (2076-24F)	
	IBM Storwize V7000 Unified File Modules (2073-720)	Version 1.5.0.5-1 or later
Software	VMware vSphere ESXi	5.5u1 or later
	VMware vCenter	5.5u1 or later
	Windows	2008R2, 2012R2

## Summary

---

The VersaStack solution combines the innovation of Cisco UCS Integrated Infrastructure with the efficiency of the IBM storage systems. The Cisco UCS Integrated Infrastructure includes the Cisco Unified Computing System (Cisco UCS), Cisco Nexus, Cisco MDS switches, and Cisco UCS Director.

The IBM Storage Systems enhances virtual environments with Data Virtualization, Real-time Compression and Easy Tier features.

VersaStack is the optimal integrated infrastructure platform to host a variety of IT workloads. Cisco and IBM have created the foundation for a flexible and scalable platform for multiple use cases and applications. From virtual desktop infrastructure to SAP®, VersaStack can efficiently and effectively support business-critical applications running simultaneously on the same shared infrastructure. The modularity of components and architectural flexibility provide a level of scalability that will enable customers to start with a right-sized infrastructure that can continue to grow with and adapt to any customer business need.

The VersaStack solution is backed by Cisco Validated Designs to provide faster delivery of applications, greater IT efficiency, and less risk. Cisco is offering its Data Center Solution Support for Critical Infrastructure Service Delivery, which provides entitled customers and partners a single contact to resolve all support issues.

## References

---

### Products and Solutions

Cisco Unified Computing System:

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

Cisco UCS 6200 Series Fabric Interconnects:

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

Cisco UCS 5100 Series Blade Server Chassis:

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

Cisco UCS B-Series Blade Servers:

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

Cisco UCS Adapters:

[http://www.cisco.com/en/US/products/ps10277/prod\\_module\\_series\\_home.html](http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html)

Cisco UCS Manager:

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

Cisco Nexus 9000 Series Switches:

<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/tsd-products-support-series-home.html>

Cisco Application Centric Infrastructure:

<http://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/index.html>

VMware vCenter Server:

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

VMware vSphere:

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

IBM Storwize V5000:

[http://www.ibm.com/systems/storage/disk/storwize\\_v5000](http://www.ibm.com/systems/storage/disk/storwize_v5000)

### Interoperability Matrixes

Cisco UCS Hardware Compatibility Matrix:

<http://www.cisco.com/c/en/us/support/servers-unified-computing/unified-computing-system/products-technical-reference-list.html>

VMware and Cisco Unified Computing System:

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

IBM System Storage Interoperation Center:

<http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss>

## About the Authors

---

Sreenivasa Edula, Technical Marketing Engineer, Cisco UCS Data Center Solutions Engineering, Cisco Systems, Inc.

Sreeni has over 17 years of experience in Information Systems with expertise across Cisco Data Center technology portfolio, including DC architecture design, virtualization, compute, network, storage and cloud computing.

Adam Reid - Test Specialist, Systems & Technology Group, IBM

Adam has more than 15 years of Computer Engineering experience. Focused more recently on IBM's **Storwize Storage Systems, he's been deeply involved with VMware and the testing and configuration of** virtualized environments pivotal to the future of software defined storage. Adam has designed and tested validated systems to meet the demands of a wide range of mid-range and enterprise environments.

## Acknowledgements

For their support and contribution to the design, validation and creation of this Cisco Validated Design (CVD), the authors acknowledge:

- **Chris O'Brien** - Manager, Technical Marketing Team, Cisco Systems, Inc.
- Sally Neate, Test Architect, Systems and Technology Group, IBM