

VersaStack Data Center with All-Flash Storage and VMware vSphere 6.0 Design Guide

Design Guide for Cisco Unified Computing System 3.1 and IBM FlashSystem V9000 with VMware vSphere 6.0 Update 1a

Last Updated: May 18, 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

About Cisco Validated Designs	2
Executive Summary	5
Solution Overview	6
Introduction	6
Audience	6
What's New?	6
VersaStack Program Benefits	8
Integrated System	8
Fabric Infrastructure Resilience	9
Network Virtualization	9
Technology Overview	10
Validated System Components	10
Cisco Unified Computing System	11
Cisco UCS 6248UP Fabric Interconnects	12
Cisco UCS 5108 Blade Server Chassis	13
Cisco UCS B200 M4 Servers	13
Cisco I/O Adapters	14
Cisco UCS Fabric Extender Modules (FEX)	15
Cisco UCS Differentiators	15
Cisco Nexus 9000 Series Switch	17
Cisco MDS 9100 Multilayer Fabric Switch	18
Cisco Nexus 1000v	19
IBM FlashSystem V9000	19
VMware vSphere	21
Domain and Element Management	21
Cisco Unified Computing System Manager	22
Cisco UCS Central (Optional)	22
IBM FlashSystem V9000 Easy-to-Use Management GUI	23
VMware vCenter Server	25
Cisco Virtual Switch Update Manager	25
IBM Spectrum Protect Snapshot	26
Solution Design	27
Network Fabric Design	29

Storage Design	30
All-Flash SAN Storage using Fibre Channel Connectivity	30
Scalability	33
Cisco UCS Server Configuration for vSphere	34
Cisco Nexus 1000v	36
Jumbo Frames	39
Management Network	39
Validation	40
Hardware and Software Revisions	40
Hardware and Software Options for VersaStack	41
Summary	43
References	44
Products and Solutions	44
Interoperability Matrixes	45
About the Authors	46
Acknowledgements	46

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 Cisco, IBM® and VMware vSphere 6.0 Update 1a on VersaStack® solution with the Cisco UCS Servers, Cisco Nexus 9000 switches and IBM FlashSystem V9000. A VersaStack solution is a validated approach for deploying Cisco and IBM technologies as a shared cloud infrastructure.

VersaStack™ is pre-designed, 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 Storwize and FlashSystem Enterprise 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 6.0 Update 1a built on the VersaStack with All-Flash 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.

The purpose of this document is to describe the Cisco and IBM VersaStack solution, which is a validated approach for deploying Cisco and IBM technologies. This validated design provides a framework for deploying VMware vSphere, the most popular virtualization platform in enterprise class data centers, on VersaStack.

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:

- Support for the Cisco UCS 3.1(1e) release and Cisco UCS B200-M4 servers
- Support for the latest release of IBM FlashSystem V9000 software 7.6.0.4
- VMware vSphere 6.0 U1a
- Validation of the Cisco Nexus 9000 switches including an IBM FlashSystem V9000 storage array with 16G Host connectivity
- Validation of Cisco MDS 9148S Switches with 16G ports

For more information on previous VersaStack models, please refer the VersaStack guides at:

<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 FlashSystem V9000 All-Flash Storage	Cisco UCS and Cisco Nexus Switches
Real time compression	Unified Fabric
Enhanced IP Replication	Virtualized IO
Form factor scaling capability	Extended Memory
Application agnostic tiering	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 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 and MDS 9000 for switching, Cisco UCS platform for the Compute and IBM FlashSystem V9000 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 Switches, 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:

- Cisco Unified Computing System (UCS)
- Cisco Nexus 9372PX Series Switches
- Cisco MDS 9148S Series Switches
- IBM FlashSystem V9000
- VMware vSphere 6.0 U1a

This section provides a technical overview of the above components

Figure 2 VersaStack All-Flash architecture



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 Fibre Channel based access is validated with iSCSI and FCOE as supported access methods 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 Unified Computing System (Cisco UCS) fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

Cisco UCS 6248UP Fabric Interconnects

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

Cisco UCS 6200 Series Fabric Interconnects support the system's 80 Gbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic using a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blade, rack servers, and virtual machines are interconnected using the same mechanisms. The Cisco UCS 6248UP is a 1-RU Fabric Interconnect that features up to 48 universal ports that can support 80 Gigabit Ethernet, Fibre Channel over Ethernet, or native Fibre Channel connectivity.

Figure 3 Cisco Fabric Interconnect



For more information, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6200-series-fabric-interconnects/index.html>


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. Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92 percent efficient and can be configured to support non-redundant, N+1 redundant and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2204XP or 2208XP Fabric Extenders. A passive mid-plane provides up to 40 Gbps of I/O bandwidth per server slot and up to 80 Gbps of I/O bandwidth for two slots. The chassis is capable of supporting future 80 Gigabit Ethernet standards.

For more information, see:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html>

Figure 4 Cisco UCS 5108 Blade Chassis



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 1536 GB of RAM (using 64 GB DIMMs), two solid-state drives (SSDs) or hard disk drives (HDDs), and up to 80 Gbps throughput connectivity. The UCS B200 M4 Blade Server mounts in a Cisco UCS 5100 Series blade server chassis or 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


VersaStack solution is typically validated using Cisco VIC 1340 or Cisco VIC 1380. Older models of the VIC are also supported with the VersaStack systems and interoperability matrix from IBM and Cisco has to be referred by the customers before solution deployment. 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



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 Fabric Extender Modules (FEX)

Each Cisco UCS chassis is equipped with a pair of Cisco UCS Fabric Extenders. The fabric extenders have two different models, 2208XP and 2204XP. Cisco UCS 2208XP has eight 10 Gigabit Ethernet, FCoE-capable ports that connect the blade chassis to the fabric interconnect. The Cisco UCS 2204 has four external ports with identical characteristics to connect to the fabric interconnect. Each Cisco UCS 2208XP has thirty-two 10 Gigabit Ethernet ports connected through the midplane to the eight half-width slots (4 per slot) in the chassis, while the 2204XP has 16 such ports (2 per slot).

Table 2 Fabric Extenders Models

Cisco UCS Fabric Extenders	Networking Facing Interface	Host Facing Interface
Cisco UCS 2204XP	4	16
Cisco UCS 2208XP	8	32

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 database 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 com-

pute 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. Simplified QoS - Even though Fibre Channel and Ethernet are converged in 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 MDS 9100 Multilayer Fabric Switch

VersaStack design used the MDS 9148S switches for SAN connectivity, the Cisco® MDS 9148S 16G Multilayer Fabric Switch is the next generation of the highly reliable, flexible, and low-cost Cisco MDS 9100 Series switches. It combines high performance with exceptional flexibility and cost effectiveness. This powerful, compact one rack-unit (1RU) switch scales from 12 to 48 line-rate 16 Gbps Fibre Channel ports. The Cisco MDS 9148S delivers advanced storage networking features and functions with ease of management and compatibility with the entire Cisco MDS 9000 Family portfolio for reliable end-to-end connectivity.

Table 3 summarizes the main features and benefits of the Cisco MDS 9148S.

Table 3 Features and Benefits

Feature	Benefit
Common software across all platforms	Reduce total cost of ownership (TCO) by using Cisco NX-OS and Cisco Prime DCNM for consistent provisioning, management, and diagnostic capabilities across the fabric.
PowerOn Auto Provisioning	Automate deployment and upgrade of software images.
Smart zoning	Reduce consumption of hardware resources and administrative time needed to create and manage zones.
Intelligent diagnostics/Hardware based slow port detection	Enhance reliability, speed problem resolution, and reduce service costs by using Fibre Channel ping and traceroute to identify exact path and timing of flows, as well as Cisco Switched Port Analyzer (SPAN) and Remote SPAN (RSPAN) and Cisco Fabric Analyzer to capture and analyze network traffic.
Virtual output queuing	Help ensure line-rate performance on each port by eliminating head-of-line blocking.
High-performance ISLs	Optimize bandwidth utilization by aggregating up to 16 physical ISLs into a single logical Port Channel bundle with multipath load balancing.
In-Service Software Upgrades	Reduce downtime for planned maintenance and software upgrades.

For more information, see: <http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-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/c/en/us/products/switches/nexus-1000v-switch-vmware-vsphere/index.html>

IBM FlashSystem V9000

The storage controller leveraged for this validated design, IBM FlashSystem V9000, is engineered to satisfy the most demanding of workloads. IBM FlashSystem V9000 is a virtualized, flash-optimized, enterprise-class storage system that provides the foundation for implementing an effective storage infrastructure with simplicity, and transforming the economics of data storage. Designed to complement virtual server environments, these modular storage systems deliver the flexibility and responsiveness required for changing business needs.

FlashSystem V9000 uses a fully featured and scalable all-flash architecture that performs at up to 2.5 million input/output operations per second (IOPS) with IBM MicroLatency, is scalable to 19.2 gigabytes per second (GBps), and delivers an effective capacity of up to 2.28 petabytes (PB). Using its flash-optimized design, FlashSystem V9000 can provide response times of 200 microseconds. It delivers better acquisition costs than a high-performance spinning disk for the same effective capacity while achieving five times the performance, making it ideal for environments that demand extreme performance.

Figure 8 IBM V9000 Storage Array



IBM FlashSystem V9000 has the following host interfaces:

- SAN-attached 16 Gbps or 8 Gbps Fibre Channel ports
- 1 Gbps iSCSI
- Optional 10 Gbps iSCSI/FCoE

Each IBM FlashSystem V9000 node canister has up to 128GB internal cache to accelerate and optimize writes, and hardware acceleration to boost the performance of Real-time Compression.

IBM FlashSystem V9000 can deploy the full range of Storwize software features, including:

- FlashCopy for near-instant data backups
- IBM Real-time Compression Accelerators
- IBM EasyTier for automated storage tiering
- Thin provisioning
- Synchronous data replication with Metro Mirror
- Asynchronous data replication with Global Mirror
- Data virtualization
- HyperSwap Split-Clusters
- Highly available configurations
- External storage virtualization and data migration

IBM FlashSystem V9000 with Real-time Compression

A key differentiator in the storage industry is IBM's Real-time Compression. Unlike other approaches to compression, Real-time Compression is designed to work on active primary data, by harnessing dedicated hardware acceleration. This achieves extraordinary efficiency on a wide range of candidate data, such as production databases and email systems, enabling storage of up to five times as much active data in the same physical disk space.

FlashSystem V9000 also provides a built in, non-intrusive compression analysis tool, that measures the yield that compression would have of an existing volume. This allows for accurate planning for capacity for workloads that you may want to move to compressed volumes without having to run any host side tools or schedule time to perform that assessment.

IBM Easy Tier

IBM Easy Tier is a solution that combines functionalities that can add value to other storage, in combination with the highly advanced key points that IBM FlashSystem V9000 can offer, such as IBM MicroLatency and maximum performance.

The great advantage of the tiering approach is the capability to automatically move the most frequently accessed data to the highest performing storage system. In this case, FlashSystem V9000 is the highest performing storage, and the less frequently accessed data can be moved to slower external storage, which can be SSD-based storage or disk-based storage.

FlashSystem V9000 addresses the combination of the lowest latency with the highest functionality and can provide the lowest latency for clients that use traditional disk array storage and need to increase the performance of their critical applications.

For more information, see: <http://www-03.ibm.com/systems/storage/flash/v9000/>

External Storage Virtualization and Data Migration

Virtualizing external storage allows IBM FlashSystem V9000 to use existing SAN storage (from a range of vendors) to act as additional capacity to the FlashSystem V9000 pools. When virtualizing traditional HDD storage, the benefits of Easy Tier are readily noticeable as the hotter access data then remains on V9000 and cooling IO patterns cause movement to the external virtualized storage, automatically.

Even if you are not looking at the long term use of External Virtualization, you can leverage external virtualization for data migration if you have some existing volumes on the SAN from where you want to migrate the data into FlashSystem V9000. This capability allows you to migrate data directly into FlashSystem V9000 without a host operator having to do a backup and restore operation to populate the data, making migrations easy and fast.

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/vsphere>

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 FlashSystem V9000 Integrated 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. UCS Manager offers unified embedded management interface that integrates server, network, and storage. 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.

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


For more information on Cisco UCS Manager, see:

<http://www.cisco.com/en/US/products/ps10281/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 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 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 9 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.

IBM FlashSystem V9000 Easy-to-Use Management GUI

The IBM FlashSystem V9000 built-in user interface (Figure 10) hides complexity and makes it possible for administrators to quickly and easily complete common block storage tasks from the same interface, such as creating and deploying volumes and host mappings. Users can also monitor performance in real-time (Figure 11).

The IBM FlashSystem V9000 management interface has the ability to check for the latest updates, and through 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 FlashSystem V9000 Management GUI Example



Figure 11 Real-time Performance Monitoring on the IBM FlashSystem V9000 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 enterprises and Service Provider environments that demand extreme performance.


The solution offers scalable performance at Compute, Network and Storage layers to accelerate critical business applications and decrease data center costs simultaneously. As a result, organizations can gain a competitive advantage through a more flexible, responsive, and efficient environment with VersaStack deployment.

This VersaStack design uses the Cisco UCS platform with Cisco B200M4 half-width blades and C240M4 UCS C-series rack mount servers running ESXi 6.0 U1a and IBM FlashSystem V9000. This design enables organizations to consolidate and provide new capabilities to their existing infrastructures, with flexible deployment options.

This VersaStack Model includes Fibre Channel based connectivity to the IBM FlashSystem V9000. The IBM FlashSystem V9000 also supports iSCSI and FCOE based connectivity for the Block storage and these could be other options that are supported though not validated in this design.

Figure 12 illustrates the new All-Flash VersaStack design. The infrastructure is physically redundant across the stack, addressing Layer 1 high-availability requirements where the integrated stack can withstand failure of a link or failure of a device. The solution also incorporates additional Cisco and IBM technologies and features that to further increase the design efficiency.

Figure 12 VersaStack All-Flash Architecture



Network Fabric Design

The Network fabric within 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 Port Channels (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 Port Channels 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 Fabric Interconnects in this case. This provides device-level redundancy and connectivity even if one of the 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 UCS domain.

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 has capabilities and performance necessary for Enterprises without having to upgrade the networking infrastructure as the networking needs to 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 Cisco Nexus 9000 Connectivity to Cisco UCS Fabric Interconnects




Figure 13 illustrates the connections between Cisco Nexus 9000 and the Fabric Interconnects. vPC requires a "peer link" which is documented as port channel 10 in this diagram. In addition to the vPC peer-link, vPC peer keepalive link is a required component of a vPC configuration. The peer keepalive link allows each vPC enabled switch to monitor the health of its peer. This link accelerates convergence and reduces the occurrence of split-brain scenarios. In this validated solution, the vPC peer keepalive link uses the management network. This link is not shown in the figure above.

Storage Design

IBM FlashSystem V9000 consists of 'building blocks', and optional additional storage expansion enclosures, all interconnected on a 16G fibre channel interconnect provided by the Cisco MDS 9148S fabrics. Multiple 16G links provide high availability for inter-cluster communications, as well as parallelism when accessing block data. The building block consists of two FlashSystem V9000 Control Enclosures, and one FlashSystem V9000 Storage Enclosure. Within the building block the controllers act as active-active data path to the volumes. Each Storage Enclosure provides RAID5 protection for its capacity.

The current VersaStack design leveraged single scalable building block with potential to scale in two dimensions: performance, and capacity. Each building block provides linear scale of performance; so four building blocks are capable of providing 4x the performance of a single building block. Each building block also carries the choice of its capacity, ranging from as little as 4TB to as much as 57TB per flash enclosure. Once your performance needs are satisfied, up to 4 additional Storage Enclosure can be added, each again ranging in 4-57TB.

Leveraging the MDS switches, additional building blocks as well as additional Storage Enclosures can be added to the IBM FlashSystem V9000 non-disruptively, allowing your VersaStack solution to grow or take on new/additional workload.

With the addition of more Storage Enclosures comes the choice of storage pool configuration. With the 2nd and subsequent Storage Enclosures, you will be presented with a wizard that will prompt a selection of **"Maximum Performance" (which is selected by default) or for that of "Maximum Availability"**. The selection of maximum performance adds the new storage enclosures capacity to the existing storage pool. With a **single pool, all storage volume allocations are 'striped' across the available capacity, which** creates parallelism when accessing the individual flash storage enclosures, and provides the maximum performance. The selection of maximum availability then causes each new flash storage enclosure to generate its own pool of capacity. Although multiple pools means the operator must thoughtfully choose which pool to allocate volumes from, the pools then form what is commonly referred to as failure domains, which go to **isolate failure impact. This form of availability also goes to isolate potential of "noisy neighbors", as to ensure that critical workload is not impacted by the "noise" of lesser workload.**

All-Flash SAN Storage using Fibre Channel Connectivity

IBM FlashSystem V9000 system is deployed as a high availability storage solution with no single point of failure. IBM FlashSystem V9000 supports fully redundant connections for communication between control enclosures, external storage, and host systems.

Each IBM FlashSystem V9000 control enclosure within a system has a pair of active-active I/O paths. If a control enclosure fails, the other control enclosure seamlessly assumes I/O responsibilities of the failed control enclosure. Control enclosures communicate over the FC SAN, for increased redundancy and performance.

Each FlashSystem V9000 control enclosure has up to eight fibre-channel ports that can attach to multiple SAN fabrics. For high availability, node canisters are attached to two fabrics.

If a SAN fabric fault disrupts communication or I/O operations, the system recovers and retries the operation through an alternative communication path. Host systems should be configured to use multi-pathing, so that if a SAN fabric fault or node canister failure occurs, the host can retry I/O operations.

Figure 14 SAN Connectivity – All-Flash Block Storage




Figure 14 above details the SAN connectivity within VersaStack All-Flash system using Cisco MDS 9000 Fibre Channel switches. Isolated fabric topologies are created on the Cisco MDS 9148S switches using VSANs to support Host and Cluster interconnect traffic. The use of VSANs allow the isolation of traffic within specific portions of the storage area network. If a problem occurs in one VSAN, that problem can be handled with a minimum of disruption to the rest of the network. VSANs can also be configured separately and independently.

Redundant host VSANs support host or server traffic to the V9000 Controllers through independent fabrics, with UCS Fabric Interconnects connected to the Cisco MDS switches using Port Channels and another pair of redundant VSANs providing Controllers to Storage enclosure connectivity using 16G ports on the MDS Switches.

Table 4 VSAN Usage

VSANs	Cisco MDS A	Cisco MDS B
Host VSANs	101	102
Cluster Interconnect VSANs	201	202

Creation of separate fabrics on the MDS switches provides logical fabric isolation with:

1. No access for any host or server to the storage enclosure accidentally.
2. No congestion to the host or server-side SAN can cause potential performance implications for both the host or server-side SAN and the FlashSystem V9000

Figure 15 illustrates connectivity between Cisco MDS 9000 and the Fabric Interconnects connected through port channels. Uplinks between the Cisco UCS Fabric Interconnects and Cisco MDS 9000 Family SAN Switch are aggregated in to port channels, which provide uplink bandwidth of 32 Gbps from each Fabric Interconnect to the Cisco MDS switch.

FlashSystem Controllers have 16 Gbps ports connected to the Cisco MDS switches and have two ports per fabric providing redundancy. Dedicated VSANs enable isolated host connectivity on each MDS switch.

Figure 15 Cisco MDS connections to FlashSystem V9000 for Host Connectivity





Figure 16 illustrates connectivity between FlashSystem V9000 Controllers and the Storage Enclosure using Cisco MDS 9000 switches. All connectivity between the Controllers and the Storage Enclosure is 16 Gbps End-to-End.

Figure 16 Cisco MDS connections to FlashSystem V9000 for Cluster Interconnect



Scalability

As mentioned earlier, VersaStack can scale up for greater performance and capacity by adding compute, network, or storage resources individually as needed.

The IBM FlashSystem V9000 supports the ability to grow both the storage capacity and performance after deployment within the VersaStack. IBM FlashSystem V9000 scale up or scale out is achieved by using scalable building blocks and additional storage enclosures.

IBM FlashSystem V9000 supports a maximum configuration of twelve 1.2 TB, 2.9 TB, or 5.7 TB IBM MicroLatency modules per scalable building block. The IBM FlashSystem V9000 can be purchased with 4, 6, 8, 10, or 12 modules of 1.2 TB, 2.9 TB, or 5.7 TB sizes.

FlashSystem V9000 delivers up to 57 TB per building block, scales to four building blocks, and offers up to four more 57 TB V9000 storage enclosure expansion units for large-scale enterprise storage system capability. Building blocks can be either fixed or scalable. You can combine scalable building blocks to create larger clustered systems in such a way that operations are not disrupted.

A scalable building block can be scaled up by adding IBM FlashSystem V9000 AE2 storage enclosures for increased storage capacity. You can add a maximum of four extra storage enclosures, one extra storage enclosure per building block, to any scaled solution. A scalable building block can be scaled out by combining up to four building blocks to provide higher IOPS and bandwidth needs for increased performance.


With IBM Real-time Compression technology, FlashSystem V9000 further extends the economic value of All-Flash systems. FlashSystem V9000 provides up to two times the improvement in Real-time Compression over the model it is replacing, by using dedicated Compression Acceleration Cards. Using the optional Real-time Compression and other design elements, the V9000 provides up to 57 TB usable capacity and up to 285 TB effective capacity in only 6U. This scales to 456 TB usable capacity and up to 2.28 PB effective capacity in only 36U.

The Fibre Channel switch virtual fabric (VSAN) for Cluster Interconnect is dedicated, and is not shared with hosts or server-side virtual storage area networks. After connecting the components in a scalable building block, no physical cable connects any host to any switch in the internal Fibre Channel VSAN fabric.

This private fabric is therefore not affected by traditional host-side SAN traffic, saturation issues, or accidental or intentional zoning issues, therefore providing maximum availability and maximum cluster performance.

This VersaStack design used one pair of MDS 9148S switches with dedicated VSANs separating cluster communication and Host traffic, dedicated MDS switch pairs can be used if customers choose to separate the cluster and host traffic with physically separated fabrics. Dedicated private fabric ensures maximum availability and maximum cluster performance in a fully scaled environment.

Figure 17 IBM FlashSystem V9000 Fully-Scaled HA Communication



The Fully scaled architecture of All Flash VersaStack is not validated as part of the VersaStack program, but it is a supported configuration and the addition of building blocks can be based on the customer's requirements

Cisco UCS Server Configuration for vSphere

The VersaStack design supports both Cisco UCS B-Series and C-Series deployments. The Cisco Unified Computing System supports the virtual server environment by providing robust, highly available, and extremely manageable compute resources. The components of the Cisco Unified Computing System offer

physical redundancy and a set of logical structures to deliver a very resilient VersaStack compute domain. In this validation effort, multiple Cisco UCS B-Series ESXi servers are SAN booted using Fibre Channel transport protocol.

The ESXi nodes consisted of Cisco UCS B200-M4 series blades with Cisco 1340 VIC adapters. These nodes were allocated to a VMware DRS and HA enabled cluster supporting infrastructure services such as vSphere Virtual Center, Microsoft Active Directory and database services.

The VersaStack Solution defines two LAN Virtual port channels (Po11 and Po12) to connect Cisco UCS Fabric Interconnects to the Cisco Nexus 9000 for network access and two SAN port channels to connect Fabric Interconnects to the Cisco MDS 9000 for storage access, storage traffic from servers use the two SAN port-channels to communicate to the IBM FlashSystem V9000.

At the server level, the Cisco 1340 VIC presents four virtual PCIe devices to the ESXi node, two virtual 10 Gb Ethernet NICs (vNIC) and two Fibre Channel vHBAs. The vSphere environment identifies these interfaces as vmnics and vmhbas, the ESXi operating system is unaware these are virtual adapters. The result is a dual-homed ESXi node to the remaining network with redundant SAN connectivity to the IBM FlashSystem V9000.

In the VersaStack design, the following virtual adapters were used:

- 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


Each UCS blade in VersaStack 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 IBM FlashSystem V9000 system ports to maximize performance and redundancy. Each blade had a boot volume created on the FlashSystem V9000 Storage array. The FlashSystem V9000 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 providing HA failover without single point of failure at hypervisor layer and MPIO host software was utilized for path management.

ESXi hosts have a single zone created for each vHBA in both the Fabrics respectively. This zone contains the ESXi host port, and the Front-End ports from each AC2 control enclosure nodes that the host will need to access within the Fabric. This method ensures workload from the hosts is distributed across all the IBM FlashSystem V9000 ports. Alternatively, if the environment consists of a mix of low and high throughput hosts, the Front-End ports on the AC2 controllers can be dedicated to the hosts based on their workload category.

Figure 18 shows logical connectivity from two ESXi Hosts to an IBM FlashSystem V9000.

Figure 18 Logical Connectivity for IBM FlashSystem V9000 with 2 vHBA ports per ESXi host




Cisco Nexus 1000v

The Cisco Nexus 1000v is a virtual distributed switch that fully integrates into a vSphere enabled environment. The Cisco Nexus 1000v operationally emulates a physical modular switch where:

- Virtual Supervisor Module (VSM) provides control and management functionality to multiple modules
- Cisco Virtual Ethernet Module (VEM) is installed on ESXi nodes and each ESXi node acts as a module in the virtual switch

Figure 19 Cisco Nexus 1000v Architecture



The VEM takes configuration information from the VSM and performs Layer 2 switching and advanced networking functions, as follows:


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



VersaStack architecture will fully support other intelligent network services offered through the Cisco Nexus 1000v such as Cisco VSG, ASA1000v, and vNAM.

Figure 20 illustrates a single ESXi node in VersaStack with a VEM registered to the Cisco Nexus 1000v VSM. The ESXi vmnics are presented as Ethernet interfaces in the Cisco Nexus 1000v.

Figure 20 Cisco Nexus 1000v VEM in an ESXi Environment



The configuration required to enable NFS access is included in this document, allowing connectivity to any existing NFS datastores for migration of virtual machines if required. However, NFS is not validated in the solution and is not supported on IBM FlashSystem V9000.

Cisco Nexus 1000v Interfaces and Port Profiles

Port profiles are logical templates that can be applied to the Ethernet and virtual Ethernet interfaces available on the Cisco Nexus 1000v. Cisco Nexus 1000v aggregates the Ethernet uplinks into a single port channel named the "System-Uplink" port profile for fault tolerance and improved throughput (Figure 20).



Since the Cisco Nexus 1000v provides link failure detection, disabling Cisco UCS Fabric Failover within the vNIC template is recommended.

The virtual machine facing virtual Ethernet ports employ port profiles customized for each virtual machines network, security and service level requirements. The VersaStack architecture employs four core VMkernel NICs (vmknics) for the ESXi servers, each with its own port profile:

- vmk0 - ESXi management
- vmk1 - NFS interface
- vmk2 - vMotion interface

The NFS and vMotion interfaces are private subnets supporting data access and VM migration across the VersaStack infrastructure. The management interface support remote vCenter access and if necessary, ESXi shell access.

The Cisco Nexus 1000v also supports Cisco's MQC to assist in uniform operation and enforcement of QoS policies across the infrastructure. The Cisco Nexus 1000v supports marking at the edge and policing traffic from VM-to-VM.

For more information, see the Best Practices in Deploying Cisco Nexus 1000V Series Switches on Cisco UCS B and C Series Cisco UCS Manager Servers:

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

Jumbo Frames

VersaStack accommodates a myriad of traffic types (vMotion, NFS, FCoE, control traffic, etc.) and is capable of absorbing traffic spikes and protect against traffic loss. Cisco UCS and Nexus QoS system classes and policies deliver this functionality. In this validation effort the VersaStack was configured to support jumbo frames with an MTU size of 9000. Enabling jumbo frames allows the VersaStack environment to optimize throughput between devices while simultaneously reducing the consumption of CPU 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 FlashSystem V9000 and Cisco UCS FI's were physically connected to a separate dedicated management switch. Management ports on 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 uplink ports on the 6248 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.

Validation

Hardware and Software Revisions

Table 5 below describes the hardware and 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. Please refer to the following links for more information:

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

Table 5 Hardware and Software Revisions Validated

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series,	3.1(1e)	Embedded management
	Cisco UCS C 220 M3/M4	3.1(1e)	Software bundle release
	Cisco UCS B 200 M3/ M4	3.1(1e)	Software bundle release
	Cisco ESXi eNIC	2.3.0.7	Ethernet driver for Cisco VIC
	Cisco ESXi fnic Driver	1.6.0.25	FCoE driver for Cisco VIC
Network	Cisco Nexus 9372	6.1(2)I3(5)	Operating system version
	Cisco MDS 9148S	6.2(13b)	FC switch firmware version
	IBM FlashSystem V9000	7.6.0.4	Software version
Storage	VMware vSphere ESXi	6.0 update1a	
Software	VMware vCenter	6.0	
	Cisco Nexus 1000v	5.2(1)SV3(1.5a)	Software version

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 Hardware Compatibility List 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. Please refer Table 6 for some 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 6 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
	UCS VIC 12XX series ,13XX series	
	Cisco MDS 91XX, 92XX, 95XX, 97XX series	6.2(9) or later
Storage	IBM Storwize V7000 , IBM FlashSystem V9000, IBM Storwize V5000 LFF Expansion (2076-12F) SFF Expansion (2076-24F)	Version 7.3.0.9 or later

Layer	Hardware	Software
	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

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 FlashSystem V9000:

<http://www-03.ibm.com/systems/storage/flash/v9000/index.html>

IBM Storwize V7000:

http://www-03.ibm.com/systems/storage/disk/storwize_v7000/

IBM 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>

IBM System Storage Interoperation Center:

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

About the Authors

Sreenivasa Edula, Technical Marketing Engineer, 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.

Dave Gimpl, Senior Technical Staff Member, IBM Systems

Dave has over 26 years of engineering experience in IBM's Systems group, and is the Chief Architect of the FlashSystem V9000 and has been involved in the development of the FlashSystem product range from its inception.

Acknowledgements

We would like to acknowledge the following for their contribution to this solution:

- **Chris O'Brien, Manager, Technical Marketing Team, Cisco Systems, Inc.**