



VersaStack with Cisco UCS M5 Servers and IBM SAN Volume Controller

Design Guide for VersaStack with Cisco UCS M5, IBM SVC, FlashSystem 900, and Storwize V5030 with vSphere 6.5U1

Last Updated: April 27, 2018



About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, 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 Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

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

© 2018 Cisco Systems, Inc. All rights reserved.

Table of Contents

Executive Summary	5
Solution Overview	6
Introduction	6
Audience	6
What's New?	6
VersaStack Program Benefits.....	7
Technology Overview	8
Cisco Unified Compute System	9
Cisco UCS Management.....	10
Cisco UCS 6300 Series Fabric Interconnects.....	11
Cisco UCS 5108 Blade Server Chassis	11
Cisco UCS B200 M5 Servers	13
Cisco UCS Network Adapters.....	14
Cisco Nexus 9000 Switches	15
Cisco MDS 9396S Fabric Switch	16
IBM Spectrum Virtualize.....	16
IBM SAN Volume Controller	18
IBM FlashSystem 900	20
IBM Storwize V5000 Second Generation	21
VMware vCenter Server 6.5 update1	22
Optional Components for the Architecture.....	22
Cisco UCS 6200 Fabric Interconnect	22
Cisco UCS 2204XP/2208XP Fabric Extenders	22
Cisco UCS C220 M5 Rack Mount Server	23
Cisco UCS Central.....	23
Cisco Cloud Center	23
Cisco UCS Director.....	23
Solution Design.....	24
Physical Topology.....	25
Cisco Unified Computing System.....	26
Cisco UCS LAN Connectivity	26
Cisco UCS SAN Connectivity	27
Cisco UCS C-Series Server Connectivity.....	28
Cisco UCS Server Configuration for VSphere.....	28

IBM Storage Systems	30
IBM SAN Volume Controller – I/O Groups	30
IBM SAN Volume Controller – iSCSI Connectivity	30
IBM SAN Volume Controller – FC Connectivity	31
IBM SAN Volume Controller – Connectivity to Cisco UCS, FlashSystem 900, and Storwize V5030	32
IBM SVC – Tiered Storage	33
Network Design	34
Virtual Port-Channel Design	34
VLAN Design	35
VSAN Design for Host to Storage Connectivity	37
Design Considerations	39
Cisco Unified Computing System Fabric Interconnects	39
Cisco Unified Computing System I/O Component Selection	40
Cisco Unified Computing System Chassis/FEX Discovery Policy	42
Storage Design and Scalability	43
Management Network Design	44
Virtual Port Channel Configuration	44
Jumbo Frame Configuration	44
Compute and Virtualization High Availability Considerations	45
Deployment Hardware and Software	46
Hardware and Software Revisions	46
Validation	47
Test Plan	47
Cisco UCS Validation	47
Network Validation	47
Storage Validation	47
vSphere Validation	47
Bill of Materials	47
Summary	49
References	50
Products and Solutions	50
Interoperability Matrixes	51
About the Authors	52

Executive Summary

Cisco Validated Designs (CVDs) deliver 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 the customers and to guide them from design to deployment.

Customers looking to deploy applications using shared data center infrastructure face a number of challenges. A recurrent infrastructure challenge is to achieve the levels of IT agility and efficiency that can effectively meet the company business objectives. Addressing these challenges requires having an optimal solution with the following key characteristics:

- Availability: Help ensure applications and services availability at all times with no single point of failure
- Flexibility: Ability to support new services without requiring underlying infrastructure modifications
- Efficiency: Facilitate efficient operation of the infrastructure through re-usable policies
- Manageability: Ease of deployment and ongoing management to minimize operating costs
- Scalability: Ability to expand and grow with significant investment protection
- Compatibility: Minimize risk by ensuring compatibility of integrated components

Cisco and IBM have partnered to deliver a series of VersaStack solutions that enable strategic data center platforms with the above characteristics. VersaStack solution delivers an integrated architecture that incorporates compute, storage and network design best practices thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance and support that can be used in various stages (planning, designing and implementation) of a deployment.

The VersaStack solution, described in this CVD, delivers a converged infrastructure platform specifically designed for Virtual Server Infrastructure. In this deployment, IBM SAN Volume Controller (SVC) standardizes storage functionality across different arrays and provides a single point of control for virtualized storage for greater flexibility and with the addition of Cisco UCS M5 servers and Cisco UCS 6300 series interconnects, the solution provides better compute performance and network throughput with 40G support for ethernet and 16G support for fibre channel connectivity. The design showcases:

- Cisco Nexus 9000 switching architecture running NX-OS mode
- IBM SVC providing single point of management and control for IBM FlashSystem 900 and IBM Storwize V5030
- Cisco Unified Compute System (UCS) servers with Intel Xeon processors
- Storage designs supporting Fibre Channel and iSCSI based storage access
- VMware vSphere 6.5U1 hypervisor
- Cisco MDS Fibre Channel (FC) switches for SAN connectivity

Solution Overview

Introduction

VersaStack solution is a pre-designed, integrated and validated architecture for the data center that combines Cisco UCS servers, Cisco Nexus family of switches, Cisco MDS fabric switches and IBM Storwize and FlashSystem Storage Arrays into a single, flexible architecture. VersaStack is designed for high availability, with no single points of failure, while maintaining cost-effectiveness and flexibility in design to support a wide variety of workloads.

VersaStack design can support different hypervisor options, bare metal servers and can also be sized and optimized based on customer workload requirements. VersaStack design discussed in this document has been validated for resiliency (under fair load) and fault tolerance during system upgrades, component failures, and partial as well as complete loss of power scenarios.

This document discusses the design principles that comprise the VersaStack solution, which is a validated Converged Infrastructure(CI) jointly developed by Cisco and IBM. The solution is a predesigned, best-practice data center architecture with VMware vSphere built on the Cisco Unified Computing System (UCS).The solution architecture presents a robust infrastructure viable for a wide range of application workloads implemented as a Virtual Server Infrastructure (VSI).

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:

- Cisco UCS B200 M5
- Cisco UCS 6300 Fabric Interconnects
- IBM SVC 2145-SV1 release 7.8.1.4
- IBM FlashSystem 900 and IBM Storwize V5030 release 7.8.1.4
- Support for the Cisco UCS release 3.2
- Validation of IP-based storage design with Nexus NX-OS switches supporting iSCSI based storage access

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

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

VersaStack Program Benefits

Cisco and IBM have carefully 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
- Implementation and deployment instructions
- Technical specifications (rules for what is, and what is not, a VersaStack configuration)
- Cisco Validated Designs and IBM Redbooks focused on a variety of use cases

Cisco and IBM have also assembled a robust and experienced support team focused on VersaStack solutions. The team includes customer account and technical sales representatives as well as professional services and technical support engineers. The support alliance between IBM and Cisco provides customers and channel services partners direct access to technical experts who are involved in cross vendor collaboration and have access to shared lab resources to resolve potential multi-vendor issues.

Technology Overview

VersaStack is a data center architecture comprised of the following infrastructure components for compute, network, and storage:

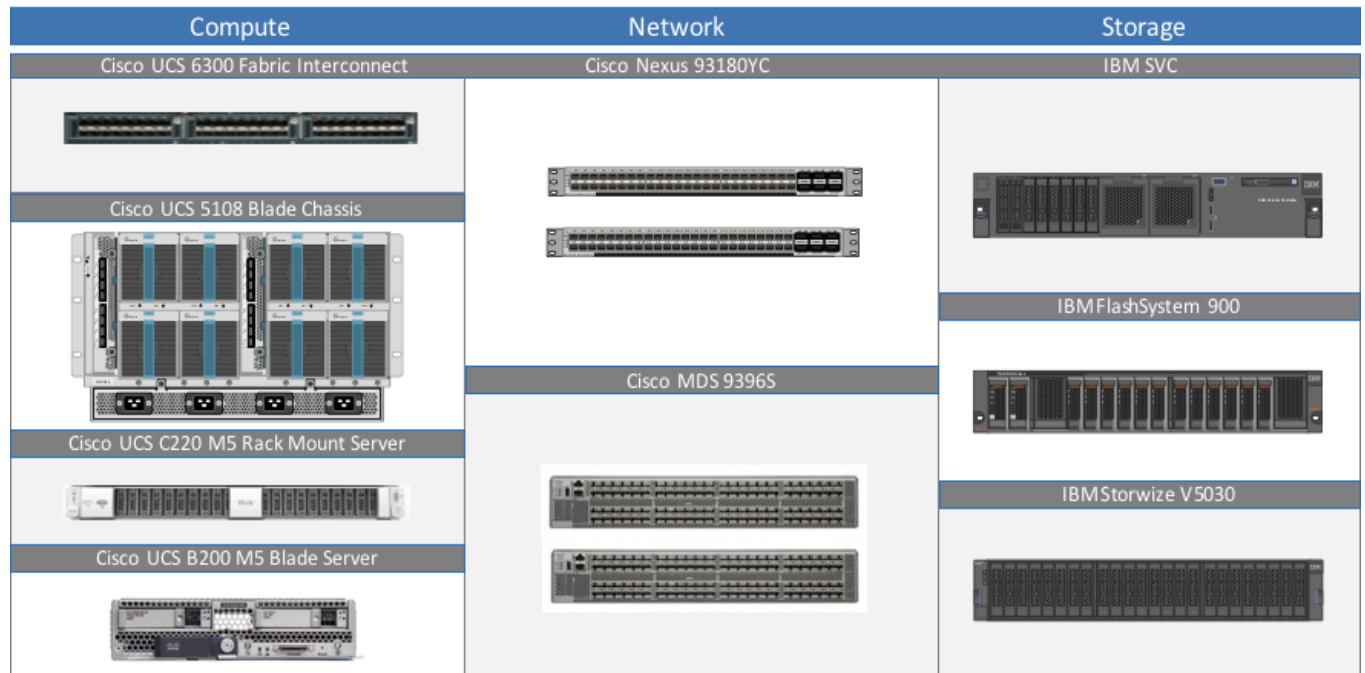
- Cisco Unified Computing System
- Cisco Nexus and Cisco MDS Switches
- IBM SAN Volume Controller, FlashSystem, and IBM Storwize family storage

These components are connected and configured according to best practices of both Cisco and IBM and provide an ideal platform for running a variety of workloads with confidence. The reference architecture covered in this document leverages:

- Cisco UCS 5108 Blade Server chassis
- Cisco UCS B-Series Blade servers
- Cisco UCS C-Series Rack Mount servers
- Cisco UCS 6300 Series Fabric Interconnects (FI)
- Cisco Nexus 93180YC-EX Switches
- Cisco MDS 9396S Fabric Switches
- IBM SAN Volume Controller (SVC) 2145-SV1 nodes *
- IBM FlashSystem V900
- IBM Storwize V5030
- VMware vSphere 6.5 Update 1

* This design guide showcases two 2145-SV1 nodes setup as a two node cluster. This configuration can be customized for customer specific deployments.

Figure 1 VersaStack with Cisco UCS M5 and IBM SVC – Components



One of the key benefits of VersaStack is the ability to maintain consistency in both scale-up and scale-down models. VersaStack can scale-up for greater performance and capacity. You can add compute, network, or storage resources as needed; or it can also scale-out when you need multiple consistent deployments such as rolling out additional VersaStack modules. Each of the component families shown in Figure 1 offer platform and resource options to scale the infrastructure up or down while supporting the same features and functionality.

The following sections provide a technical overview of the compute, network, storage and management components of the VersaStack solution.

Cisco Unified Compute System

The Cisco Unified Computing System (UCS) is a next-generation data center platform that integrates computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership (TCO) and to increase business agility. The system integrates a low-latency; lossless unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform where all resources are managed through a unified management domain.

The Cisco Unified Computing System in the VersaStack architecture utilizes the following components:

- Cisco UCS Manager (UCSM) provides unified management of all software and hardware components in the Cisco UCS to manage servers, networking, and storage configurations. The system uniquely integrates all the system components, enabling the entire solution to be managed as a single entity through Cisco UCSM software. Customers can interface with Cisco UCSM through an intuitive graphical user interface (GUI), a command-line interface (CLI), and a robust application-programming interface (API) to manage all system configuration and operations.
- Cisco UCS 6300 Series Fabric Interconnects are a core part of Cisco UCS, providing both network connectivity and management capabilities for the system. The Cisco UCS 6300 Series offers line-rate,

low-latency, lossless 10 and 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.

- Cisco UCS 5108 Blade Server Chassis supports up to eight blade servers and up to two fabric extenders in a six-rack unit (RU) enclosure.
- Cisco UCS B-Series Blade Servers provide performance, efficiency, versatility and productivity with the latest Intel based processors.
- Cisco UCS C-Series Rack Mount Servers deliver unified computing innovations and benefits to rack servers with performance and density to support a wide range of workloads.
- Cisco UCS Network Adapters provide wire-once architecture and offer a range of options to converge the fabric, optimize virtualization and simplify management.

Cisco UCS Management

Cisco Unified Computing System has revolutionized the way servers are managed in the data center. Some of the unique differentiators of Cisco UCS and Cisco UCS Manager are:

- 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.
- 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.
- Policy Based Resource Classification –Once a compute resource is discovered by Cisco UCS Manager, it can be automatically classified based on defined policies. This capability is useful in multi-tenant cloud computing.
- Combined Rack and Blade Server Management –Cisco UCS Manager can manage B-series blade servers and C-series rack server under the same Cisco UCS domain. This feature, along with stateless computing makes compute resources truly hardware form factor agnostic.
- 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 Cisco UCS Manager with other management systems.
- Service Profiles and Stateless Computing –A service profile is a logical representation of a server, carrying its various identities and policies. Stateless computing enables procurement of a server within minutes compared to days in legacy server management systems.
- Policies, Pools, Templates –The management approach in Cisco UCS Manager is based on defining policies, pools and templates, instead of cluttered configuration, which enables a simple, loosely coupled, data driven approach in managing compute, network and storage resources.
- Built-in Multi-Tenancy Support –The combination of policies, pools and templates, organization hierarchy and a service profiles based approach to compute resources makes Cisco UCS Manager inherently friendly to multi-tenant environment typically observed in private and public clouds.

Cisco UCS 6300 Series Fabric Interconnects

The Cisco UCS 6300 series Fabric interconnects (FI) new to VersaStack architecture provide a single point for connectivity and management for the entire system by integrating all compute components into a single, highly available management domain controlled by Cisco UCS Manager. Cisco UCS 6300 FIs support the **system's unified fabric with low**-latency, lossless, 40Gbps cut-through switching that supports IP, storage, and management traffic using a single set of cables. Cisco UCS FIs are typically deployed in redundant pairs. Because it supports unified fabric, the Cisco UCS 6300 Series Fabric Interconnect provides both LAN and SAN connectivity for all servers within its domain.

The 6332-16UP Fabric Interconnect used in the design is the management and communication backbone for Cisco UCS B-Series Blade Servers, C-Series Rack Servers, and 5100 Series Blade Server Chassis.

The 6332-16UP offers 40 ports in one rack unit (RU), including:

- 24 40-Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)
- 16 1- and 10-Gbps and FCoE or 4-, 8-, and 16-Gbps Fibre Channel unified ports
- Enhanced features and capabilities include:
- Increased bandwidth up to 2.43 Tbps
- Centralized unified management with Cisco UCS Manager
- Efficient cooling and serviceability such as front-to-back cooling, redundant front-plug fans and power supplies, and rear cabling

Figure 2 Cisco UCS 6332-16UP Fabric Interconnect



For more information about various models of the Cisco UCS 6300 Fabric Interconnect, see <https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html>.



Older generations of Cisco UCS Fabric Interconnect are also supported in this design guide.

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5108 Blade Server Chassis is a fundamental building block of the Cisco Unified Computing System, delivering a scalable and flexible blade server architecture. A Cisco UCS 5108 Blade Server chassis is six rack units (6RU) high and can house up to eight half-width or four full-width Cisco UCS B-series blade servers.

For a complete list of blade servers supported, see: <http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>.

The Cisco UCS 5108 chassis contains two I/O bays in the rear that can support Cisco UCS 2200 and 2300 Series Fabric Extenders. The two fabric extenders can be used for both redundancy and bandwidth aggregation. 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 (full width) blades. The chassis is also capable of supporting 40 Gigabit Ethernet. Cisco UCS 5108 blade server chassis uses a unified fabric and fabric-extender technology to simplify and reduce cabling by eliminating the need for dedicated chassis management and blade switches. The unified fabric also reduces TCO by reducing the number of network interface cards (NICs), host bus adapters (HBAs), switches, and cables that need to be managed, cooled, and powered. This architecture enables a single Cisco UCS domain to scale up to 20 chassis with minimal complexity.

For more information, see:

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

Figure 3 Cisco UCS 5108 Blade Server Chassis



Cisco UCS 2304 Fabric Extender

Cisco UCS 2304 Fabric Extender brings the unified fabric into the blade server enclosure, providing multiple 40 Gigabit Ethernet connections between blade servers and the fabric interconnect, simplifying diagnostics, cabling, and management. It is a third-generation I/O Module (IOM) that shares the same form factor as the second-generation Cisco UCS 2200 Series Fabric Extenders and is backward compatible with the shipping Cisco UCS 5108 Blade Server Chassis.

The Cisco UCS 2304 connects the I/O fabric between the Cisco UCS 6300 Series Fabric Interconnects and the Cisco UCS 5100 Series Blade Server Chassis, enabling a lossless and deterministic Fibre Channel over Ethernet (FCoE) fabric to connect all blades and chassis together. Because the fabric extender is similar to a distributed line card, it does not perform any switching and is managed as an extension of the fabric interconnects. This approach removes switching from the chassis, reducing overall infrastructure complexity and enabling Cisco UCS to scale to many chassis without multiplying the number of switches needed, reducing TCO and allowing all chassis to be managed as a single, highly available management domain.

The Cisco UCS 2304 Fabric Extender has four 40 Gigabit Ethernet, FCoE-capable, Quad Small Form-Factor Pluggable (QSFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2304 can provide one 40 Gigabit Ethernet ports connected through the midplane to each half-width slot in the chassis, giving it a total eight 40G interfaces to the compute. Typically configured in pairs for redundancy, two fabric extenders provide up to 320 Gbps of I/O to the chassis.

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

Figure 4 Cisco UCS 2304 Fabric Extender



Cisco UCS B200 M5 Servers

The Cisco UCS B200 M5 Blade Server delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads including Virtual Desktop Infrastructure (VDI), web infrastructure, distributed databases, converged infrastructure, and enterprise applications such as Oracle and SAP HANA. The B200 M5 server can quickly deploy stateless physical and virtual workloads through programmable, easy-to-use Cisco UCS Manager software and simplified server access through Cisco SingleConnect technology. It includes:

- Latest Intel® Xeon® Scalable processors with up to 28 cores per socket
- Up to 24 DDR4 DIMMs for improved performance
- Two GPUs
- Two Small-Form-Factor (SFF) drives
- Two Secure Digital (SD) cards or M.2 SATA drives
- Up to 80 Gbps of I/O throughput

Main Features

The Cisco UCS B200 M5 server is a half-width blade. Up to eight servers can reside in the 6-Rack-Unit (6RU) Cisco UCS 5108 Blade Server Chassis, offering one of the highest densities of servers per rack unit of blade chassis in the industry. You can configure the Cisco UCS B200 M5 to meet your local storage requirements without having to buy, power, and cool components that you do not need.

The Cisco UCS B200 M5 provides the following features:

- Up to two Intel Xeon Scalable CPUs with up to 28 cores per CPU
- 24 DIMM slots for industry-standard DDR4 memory at speeds up to 2666 MHz, with up to 3 TB of total memory when using 128-GB DIMMs
- Modular LAN On Motherboard (mLOM) card with Cisco UCS Virtual Interface Card (VIC) 1340, a 2-port, 40 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable mLOM mezzanine adapter
- Optional rear mezzanine VIC with two 40-Gbps unified I/O ports or two sets of 4 x 10-Gbps unified I/O ports, delivering 80 Gbps to the server; adapts to either 10- or 40-Gbps fabric connections

- Two optional, hot-pluggable, Hard-Disk Drives (HDDs), Solid-State Disks (SSDs), or NVMe 2.5-inch drives with a choice of enterprise-class RAID or passthrough controllers
- Cisco FlexStorage local drive storage subsystem, which provides flexible boot and local storage capabilities and allows you to boot from dual, mirrored SD cards
- Support for up to two optional GPUs
- Support for up to one rear storage mezzanine card

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

Figure 5 Cisco UCS B200 M5 Blade Server



Cisco UCS Network Adapters

The Cisco Unified Computing System supports converged network adapters (CNAs) to provide connectivity to the blade and rack mount servers. CNAs obviate the need for multiple network interface cards (NICs) and host bus adapters (HBAs) by converging LAN and SAN traffic in a single interface. While Cisco UCS supports wide variety of interface cards, this CVD utilizes Cisco Virtual Interface Card (VIC) 1340 model.

Cisco UCS Virtual Interface Card 1340

The Cisco UCS Virtual Interface Card (VIC) 1340 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. The Cisco 1340 VIC supports an optional port-expander which enables the 40-Gbps Ethernet capabilities of the card. The Cisco UCS VIC 1340 supports 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.

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

Figure 6 Cisco VIC 1340



Cisco Nexus 9000 Switches

In the VersaStack design, Cisco UCS Fabric Interconnects and IBM storage systems are connected to the Cisco Nexus 9000 switches. Cisco Nexus 9000 provides Ethernet switching fabric for communications between the Cisco UCS domain, the IBM storage system and the enterprise network.

The Cisco Nexus 9000 family of switches supports two modes of operation: NX-OS standalone mode and Application Centric Infrastructure (ACI) fabric mode. In standalone mode, the switch performs as a typical Cisco Nexus switch with increased port density, low latency and 40G connectivity. In fabric mode, the administrator can take advantage of Cisco Application Centric Infrastructure (ACI).

Cisco Nexus series switches provide an Ethernet switching fabric for communications between the **Cisco UCS and the rest of a customer's network. There are many factors to take into account when** choosing the main data switch in this type of architecture to support both the scale and the protocols required for the resulting applications. All Nexus switch models including the Nexus 5000 and Nexus 7000 are supported in this design, and may provide additional features such as FCoE or OTV. However, be aware that there may be slight differences in setup and configuration based on the switch used. The solution design validation leverages the Cisco Nexus 9000 series switches, which deliver high performance 40GbE ports, density, low latency, and exceptional power efficiency in a broad range of compact form factors.

Many of the most recent VersaStack designs also use this switch due to the advanced feature set and the ability to support Application Centric Infrastructure (ACI) mode. When leveraging ACI fabric mode, the Nexus 9000 series switches are deployed in a spine-leaf architecture. Although the reference architecture covered in this design does not leverage ACI, it lays the foundation for customer migration to ACI in the future, and fully supports ACI today if required.

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

This VersaStack design deploys a single pair of Nexus 9000 top-of-rack switches within each placement, using the traditional standalone mode running NX-OS.

The traditional deployment model using Nexus 9000 switches delivers numerous benefits for this design:

- High performance and scalability with L2 and L3 support per port (Up to 60Tbps of non-blocking performance with less than 5 microsecond latency)
- Layer 2 multipathing with all paths forwarding through the Virtual port-channel (vPC) technology
- VXLAN support at line rate
- Advanced reboot capabilities include hot and cold patching
- Hot-swappable power-supply units (PSUs) and fans with N+1 redundancy

Cisco Nexus 93180YC-EX

The Cisco Nexus 93180YC-EX Switch are 1RU switches that support 48 10/25-Gbps Small Form Pluggable Plus (SFP+) ports and 6 40/100-Gbps Quad SFP+ (QSFP+) uplink ports. All ports are line rate, delivering 3.6

Tbps of throughput. The switch supports Cisco Tetration Analytics Platform with built-in hardware sensors for rich traffic flow telemetry and line-rate data collection.

Figure 7 Cisco Nexus 93180YC-EX Switch



For detailed information about the Cisco Nexus 9000 product line, see

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

Cisco MDS 9396S Fabric Switch

In the VersaStack design, Cisco MDS switches provide SAN connectivity between the IBM storage systems and Cisco UCS domain. The Cisco MDS 9396S 16G Multilayer Fabric Switch used in this design is the next generation of the highly reliable, flexible, and affordable Cisco MDS 9000 Family fabric switches. This powerful, compact, 2-rack-unit switch scales from 48 to 96 line-rate 16-Gbps Fibre Channel ports in 12 port increments. Cisco MDS 9396S is powered by Cisco NX-OS and 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. Cisco MDS 9396S provides up to 4095 buffer credits per group of 4 ports and supports some of the advanced functions such as Virtual SAN (VSAN), Inter-VSAN routing (IVR), port-channels and multipath load balancing and flow-based and zone-based QoS.

Figure 8 Cisco MDS 9396S



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

IBM Spectrum Virtualize

IBM Spectrum Virtualize™ **provides an ideal way to manage and protect the huge volumes of data** organizations use for big-data analytics and new cognitive workloads. It is a proven offering that has been available for years in IBM SAN Volume Controller (SVC), the IBM Storwize® family of storage solutions, IBM FlashSystem® V9000 and IBM VersaStack™—with more than 130,000 systems running IBM Spectrum Virtualize. These systems are delivering better than five nines of availability while managing more than 5.6 exabytes of data.

IBM Spectrum Virtualize Software V7.8.1.4 includes the following key improvements:

- Reliability, availability, and serviceability with NPIV host port fabric virtualization, Distributed RAID support for encryption, and IP Quorum in GUI
- Scalability with support for up to 10,000 volumes, depending on the model; and up to 20 Expansion Enclosures on SV1 models and on FlashSystem V9000
- Manageability with CLI and GUI support for host groups
- Virtualization of iSCSI-attached external storage arrays including XIV(R) Gen 3, Spectrum Accelerate, FlashSystem A9000, FlashSystem A9000R and any Storwize arrays.
- Performance with 64 GB read cache
- Data economics with Comprestimator in the GUI
- Licensing of compression formally added to the External Virtualization software package of Storwize V5030
- Software licensing metrics to better align the value of SVC software with Storage use cases through Differential Licensing

Key existing functions of IBM Spectrum Virtualize include:

- Virtualization of Fibre Channel-attached external storage arrays with support for almost 370 different brands and models of block storage arrays, including arrays from competitors
- Manage virtualized storage as a single storage pool, integrating "islands of storage" and easing the management and optimization of the overall pool
- **Real-time Compression™ for in-line**, real-time compression to improve capacity utilization
- Virtual Disk Mirroring for two redundant copies of LUN and higher data availability
- Stretched Cluster and HyperSwap® for high availability among physically separated data centers
- Easy Tier® for automatic and dynamic data tiering
- Distributed RAID for better availability and faster rebuild times
- Encryption to help improve security of internal and externally virtualized capacities
- FlashCopy® snapshots for local data protection
- Remote Mirror for synchronous or asynchronous remote data replication and disaster recovery through both Fibre Channel and IP ports with offerings that utilize IBM Spectrum Virtualize software
- Clustering for performance and capacity scalability
- Online, transparent data migration to move data among virtualized storage systems without disruptions
- Common look and feel with other offerings that utilize IBM Spectrum Virtualize software
- IP Link compression to improve usage of IP networks for remote-copy data transmission

More information can be found at the IBM Spectrum Virtualize website:

<http://www-03.ibm.com/systems/storage/spectrum/virtualize/index.html>

IBM SAN Volume Controller


IBM SAN Volume Controller (SVC), is a combined hardware and software storage virtualization system with a single point of control for storage resources. SVC includes many functions traditionally deployed separately in disk systems and by including these in a virtualization system, SVC standardizes functions across virtualized storage for greater flexibility and potentially lower costs.

Figure 9 IBM SAN Volume Controller (2145-SV1)



Built with IBM Spectrum Virtualize™ software—part of the IBM Spectrum Storage™ family—SVC helps organizations achieve better data economics by supporting the new workloads that are critical to company's success. SVC systems can handle the massive volumes of data from mobile and social applications, enable rapid and flexible cloud services deployments, and deliver the performance and scalability needed to gain insights from the latest analytics technologies.

Figure 10 IBM SVC – Easy-to-use Graphical User Interface



IBM SVC provides an easy-to-use graphical interface for centralized management. With this single interface, administrators can perform configuration, management and service tasks in a consistent manner over multiple storage systems—even from different vendors—vastly simplifying management and helping reduce the risk of errors. IT staff can also use built-in monitoring capabilities to securely check the health and performance of the system remotely from a mobile device. In addition, plug-ins to support VMware vCenter help enable more efficient, consolidated management in these environments.

With the addition of the latest Storage Engine Model SV1, SVC delivers increased performance and additional internal storage capacity. Model SV1 consists of two Xeon E5 v4 Series eight-core processors and 64 GB of memory. It includes 10 Gb Ethernet ports standard for 10 Gb iSCSI connectivity and service technician use, and supports up to four I/O adapter cards for 16 Gb FC and 10 Gb iSCSI/FCoE connectivity. It also includes two integrated AC power supplies and battery units.

Up to eight **SVC Storage Engines**, or 'nodes', can be clustered to help deliver greater performance, bandwidth, scalability, and availability. An SVC clustered system can contain up to four node pairs or I/O groups. Model SV1 storage engines can be added into existing SVC clustered systems that include previous generation storage engine models such as 2145-DH8.

In summary, SVC combines a variety of IBM technologies to:

- Enhance storage functions, economics and flexibility with sophisticated virtualization
- Leverage hardware-accelerated data compression for efficiency and performance
- Store up to five times more active data using IBM® Real- time Compression™
- Move data among virtualized storage systems without disruptions
- Optimize tiered storage—including flash storage—automatically with IBM Easy Tier®

- Improve network utilization for remote mirroring and help reduce costs
- Implement multi-site configurations for high availability and data mobility

More SVC product information is available on the IBM SAN Volume Controller website: <http://www-03.ibm.com/systems/storage/software/virtualization/svc/>.

IBM FlashSystem 900

IBM® FlashSystem™ 900 is a fully optimized, all flash storage array designed to accelerate the applications that drive business. Featuring IBM FlashCore™ technology, FlashSystem 900 delivers the high performance, ultra-low latency, enterprise reliability and superior operational efficiency required for gaining competitive advantage in today's dynamic marketplace.

Figure 11 IBM FlashSystem 900



FlashSystem 900 is composed of up to 12 MicroLatency modules—massively parallel flash arrays that can provide nearly 40 percent higher storage capacity densities than previous FlashSystem models. FlashSystem 900 can scale usable capacity from as low as 2 TB to as much as 57 TB in a single system. MicroLatency modules also support an offload AES-256 encryption engine, high-speed internal interfaces, and full hot-swap and storage capacity scale-out capabilities, enabling organizations to achieve lower cost per capacity with the same enterprise reliability.

Key features include:

- Accelerate critical applications, support more concurrent users, speed batch processes and lower virtual desktop costs with **the extreme performance of IBM® FlashCore™ technology**
- Harness the power of data with the **ultra-fast response times of IBM MicroLatency™**
- Leverage macro efficiency for high storage density, low power consumption and improved resource utilization
- **Protect critical assets and boost reliability with IBM Variable Stripe RAID™, redundant components and concurrent code loads**
- Power faster insights with hardware accelerated architecture, MicroLatency modules and advanced flash management capabilities of FlashCore technology

To read more about FlashSystem 900 on the web, see ibm.com/storage/flash/900.

IBM Storwize V5000 Second Generation

Designed for software-defined environments and built with IBM Spectrum Virtualize software, the IBM Storwize family is an industry-leading solution for storage virtualization that includes technologies to complement and enhance virtual environments, delivering a simpler, more scalable, and cost-efficient IT infrastructure.

Storwize V5000 is a highly flexible, easy to use, virtualized hybrid storage system designed to enable midsized organizations to overcome their storage challenges with advanced functionality. Storwize V5000 second-generation models offer improved performance and efficiency, enhanced security, and increased scalability with a set of three models to deliver a range of performance, scalability, and functional capabilities.

Storwize V5000 second-generation offer three hybrid models – IBM Storwize V5030, IBM Storwize V5020 and IBM Storwize V5010 – providing the flexibility to start small and keep growing while leveraging existing storage investments. To enable organizations with midsized workloads achieve advanced performance, IBM Storwize V5030F provides an all-flash solution at an affordable price. In this Design Guide, an IBM Storwize V5030 hybrid model was deployed and validated.

Figure 12 IBM Storwize V5030



Storwize V5030 control enclosure models offer:

- Two 6-core processors and up to 64 GB of cache
- Support for 504 drives per system with the attachment of 20 Storwize V5000 expansion enclosures and 1,008 drives when clustered with a second Storwize V5030 control enclosure
- External virtualization to consolidate and provide Storwize V5000 capabilities to existing storage infrastructures
- Real-time Compression for improved storage efficiency
- Encryption of data at rest stored within the Storwize V5000 system and externally virtualized storage systems

All Storwize V5000 second-generation control enclosures include:

- Dual-active, intelligent node canisters with mirrored cache
- Ethernet ports for iSCSI connectivity
- Support for 16 Gb FC, 12 Gb SAS, 10 Gb iSCSI/FCoE, and 1 Gb iSCSI for additional I/O connectivity

- Twelve 3.5-inch (LFF) drive slots or twenty-four 2.5-inch (SFF) drive slots within the 2U, 19-inch rack mount enclosure
- Support for the attachment of second-generation Storwize V5000 LFF and SFF 12 Gb SAS expansion enclosures
- Support for a rich set of IBM Spectrum Virtualize functions including thin provisioning, IBM Easy Tier, FlashCopy, and remote mirroring
- Either 100-240 V AC or -48 V DC power supplies
- Either a one- or three-year warranty with customer replaceable units (CRU) and on-site service

All Storwize V5000 functional capabilities are provided through IBM Spectrum Virtualize Software for Storwize V5000. For additional information about Storwize V5000 functional capabilities and software, go to the **IBM Storwize V5000** website:

http://www03.ibm.com/systems/uk/storage/disk/storwize_v5000/index.html

VMware vCenter Server 6.5 update1

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual 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>.

Optional Components for the Architecture

Cisco UCS 6200 Fabric Interconnect

The Cisco UCS 6248-UP and 6296-UP Fabric Interconnect models present 10 Gbps alternatives to the UCS 6332-16UP Fabric Interconnect supporting 8 Gbps fibre channel within their unified ports. These fabric interconnects can provide the same management and communication backbone to Cisco UCS B-Series and C-Series as the Cisco UCS 6300 series fabric interconnects.



Cisco UCS 2204XP/2208XP Fabric Extenders

The Cisco UCS 2204XP and 2208XP Fabric Extenders are four and eight port 10Gbps converged fabric options to the Cisco UCS 2304 Fabric Extender, that deliver combined respective bandwidths of 80Gbps or 160Gbps when implemented as pairs.



Cisco UCS C220 M5 Rack Mount Server

The Cisco UCS C220 M5 server extends the capabilities of the Cisco UCS portfolio in a 1-Rack-Unit (1RU) form factor. It incorporates the Intel® Xeon® Scalable processors, supporting up to 20 percent more cores per socket, twice the memory capacity, 20 percent greater storage density, and five times more PCIe NVMe Solid-State Disks (SSDs) compared to the previous generation of servers. These improvements deliver significant performance and efficiency gains that will improve your application performance.

Cisco UCS C-Series servers added into a VersaStack are easily managed through Cisco UCS Manager with Single Wire Management, using 10G or 40G uplinks to the FI.

Figure 13 Cisco UCS C220 M5 Rack Server



For more information, see <https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/datasheet-c78-739279.html>.

Cisco UCS Central

Cisco UCS Central scales policy decisions and management of multiple Cisco UCS domains into a single source. Service Profiles and maintenance can be uniform in their implementation across local and/or remotely deployed fabric interconnects that can be part of multiple VersaStack deployments.

Cisco Cloud Center

Cisco Cloud Center can securely extend VersaStack resources between private and public cloud. **Application blueprints rapidly scale the user's needs, while implementing them in a manner that is secure, and provides the controls for proper governance.**

Cisco UCS Director

Cisco UCS Director provides increased efficiency through automation capabilities throughout FlashStack components. Pure Storage has enabled an adapter to include FlashArray support within FlashStack, and there are pre-built automation workflows available, centered around VersaStack.

Solution Design

This VersaStack with Cisco UCS M5 and IBM SVC architecture aligns with the converged infrastructure configurations and best practices as identified in the previous VersaStack releases. The system includes hardware and software compatibility support between all components and aligns to the configuration best practices for each of these components. All the core hardware components and software releases are listed and supported on both the Cisco compatibility list:

http://www.cisco.com/en/US/products/ps10477/prod_technical_reference_list.html

and IBM Interoperability Matrix:

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

The system supports high availability at network, compute and storage layers such that no single point of failure exists in the design. The system utilizes 10 and 40 Gbps Ethernet jumbo-frame based connectivity combined with port aggregation technologies such as virtual port-channels (VPC) for non-blocking LAN traffic forwarding. A dual SAN 16Gbps environment provides redundant storage access from compute devices to the storage controllers.

The VersaStack Datacenter with Cisco UCS M5 and IBM SVC solution, as described in this design guide, conforms to a number of key design and functionality requirements. Some of the key features of the solution are highlighted below.

- The system is able to tolerate the failure of compute, network or storage components without significant loss of functionality or connectivity
- The system is built with a modular approach thereby allowing customers to easily add more network (LAN or SAN) bandwidth, compute power or storage capacity as needed
- The system supports stateless compute design thereby reducing time and effort required to replace or add new compute nodes
- The system provides network automation and orchestration capabilities to the network administrators using NX-OS CLI and restful API
- The systems allow the compute administrators to instantiate and control application Virtual Machines (VMs) from VMware vCenter
- The system provides storage administrators a single point of control to easily provision and manage the storage using IBM management GUI
- The solution supports live VM migration between various compute nodes and protects the VM by utilizing VMware HA and DRS functionality
- The system can be easily integrated with optional Cisco (and third party) orchestration and management platforms such as Cisco UCS Central, Cisco UCS Director and Cisco CloudCenter.

Physical Topology

This VersaStack with Cisco UCS M5 and IBM SVC solution utilizes Cisco UCS platform with Cisco B200 M5 half-width blades connected and managed through Cisco UCS 6332-16UP Fabric Interconnects and the integrated Cisco UCS manager. These high performance servers are configured as stateless compute nodes where ESXi 6.5 U1 hypervisor is loaded using SAN (iSCSI and FC) boot. The boot disks to store ESXi hypervisor image and configuration along with the block and file based datastores to host application Virtual Machines (VMs) are provisioned on the IBM storage devices.

The link aggregation technologies play an important role in VersaStack solution providing improved aggregate bandwidth and link resiliency across the solution stack. Cisco UCS, and Cisco Nexus 9000 platforms support active port channeling using 802.3ad standard Link Aggregation Control Protocol (LACP). In addition, the Cisco Nexus 9000 series features virtual Port Channel (vPC) capability which allows links that are physically connected to two different Cisco Nexus devices to appear as a single "logical" port channel. Each Cisco UCS Fabric Interconnect (FI) is connected to both the Cisco Nexus 93180 switches using virtual port-channel (vPC) enabled 40GbE uplinks for a total aggregate system bandwidth of 80GBps. Additional ports can be easily added to the design for increased throughput. Each Cisco UCS 5108 chassis is connected to the UCS FIs using a pair of 40GbE ports from each IO Module for a combined 80GbE uplink. When Cisco UCS C-Series servers are used, each of the Cisco UCS C-Series servers connects directly into each of the FIs using a 10/40Gbps converged link for an aggregate bandwidth of 20/80Gbps per server.

To provide the compute to storage system connectivity, this design guides highlights two different storage connectivity options:

- Option 1: iSCSI based storage access through Cisco Nexus Fabric
- Option 2: FC based storage access through Cisco MDS 9396S



While storage access from the Cisco UCS compute nodes to IBM SVC storage nodes can be iSCSI or FC based, IBM SVC nodes, FlashSystem 900 and Storwize V5030 communicate with each other on Cisco MDS 9396S based FC SAN fabric only.

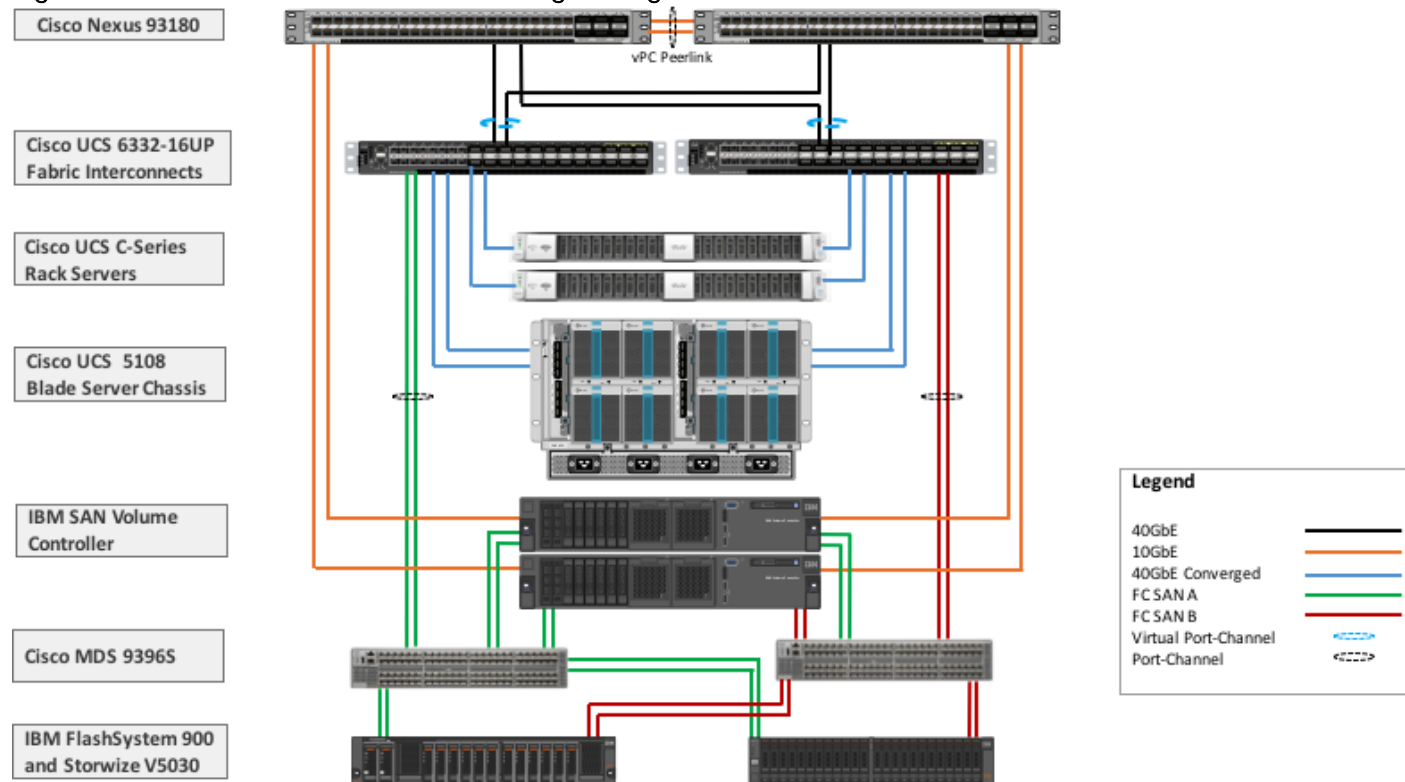
Figure 14 provides a high-level topology of the system connectivity. The VersaStack infrastructure satisfies the high-availability design requirements and is physically redundant across the network, compute and storage stacks. The integrated compute stack design presented in this document can withstand failure of one or more links as well as the failure of one or more devices.

IBM SVC Nodes, IBM FlashSystem 900 and IBM Storwize V5030 are all connected using a Cisco MDS 9396S based redundant FC fabric. To provide FC based storage access to the compute nodes, Cisco UCS Fabric Interconnects are connected to the same Cisco MDS 9396S switches and zoned appropriately. To provide iSCSI based storage access, IBM SVC is connected directly to the Cisco Nexus 93180 switches. One 10GbE port from each IBM SVC Node is connected to each of the two Cisco Nexus 93180 switches providing an aggregate bandwidth of 40Gbps.



Based on the customer requirements, the compute to storage connectivity in the SVC solution can be deployed as an FC-only option, iSCSI-only option or a combination of both. Figure 14 shows connectivity option to support both iSCSI and FC.

Figure 14 VersaStack ISCSI and FC Storage Design with IBM SVC



The following sections cover physical and logical connectivity details across the stack including various design choices at compute, storage, virtualization and network layers.


Cisco Unified Computing System

The VersaStack compute design supports both Cisco UCS B-Series and C-Series deployments. Cisco UCS supports the virtual server environment by providing robust, highly available, and extremely manageable compute resources. In this validation effort, multiple Cisco UCS B-Series and C-Series ESXi servers are booted from SAN using iSCSI or FC (depending on the storage design option).

Cisco UCS LAN Connectivity

Cisco UCS Fabric Interconnects are configured with two port-channels, one from each FI, to the Cisco Nexus 93180 switches. These port-channels carry all the data and IP-based storage traffic originated on the Cisco Unified Computing System. Virtual Port-Channels (vPC) are configured on the Cisco Nexus 93180 to provide device level redundancy. The validated design utilized two uplinks from each FI to the Cisco Nexus switches for an aggregate bandwidth of 160GbE (4 x 40GbE). The number of links can be increased based on customer data throughput requirements.


Figure 15 Cisco UCS - LAN Connectivity



Cisco UCS SAN Connectivity

Cisco UCS Fabric Interconnects (FI) are connected to the redundant Cisco MDS 9396S based SAN fabric to provide Fibre Channel storage connectivity. In addition to the LAN connections covered in Figure 15, two 16Gbps FC ports from each of the FIs are connected to one of the two Cisco MDS fabric switches to support a SAN-A/B design. These ports are configured as FC port-channel to provide 32Gbps effective bandwidth from each FI to each fabric switch. This design is shown in Figure 16.

Figure 16 Cisco UCS – SAN Connectivity




Cisco MDS switches are deployed with N-Port ID Virtualization (NPIV) enabled to support the virtualized environment running on Cisco UCS blade and rack servers. To support NPIV on the Cisco UCS servers, the Cisco UCS Fabric Interconnects that connect the servers to the SAN fabric, are enabled for N-Port Virtualization (NPV) mode by configuring to operate in end-host mode. NPV enables Cisco FIs to proxy fabric login, registration and other messages from the servers to the SAN Fabric without being a part of the SAN fabric. This is important for keeping the limited number of Domain IDs that Fabric switches require to a minimum. FC port-channels are utilized for higher aggregate bandwidth and redundancy. Cisco MDS also provide zoning configuration to enable single initiator (vHBA) to talk to multiple targets.

Cisco UCS C-Series Server Connectivity

In all VersaStack designs, Cisco UCS C-series rack mount servers are always connected via the Cisco UCS FIs and managed through Cisco UCS Manager to provide a common management look and feel. Cisco UCS Manager 2.2 and later versions allow customers to connect Cisco UCS C-Series servers directly to Cisco UCS Fabric Interconnects without requiring a Fabric Extender (FEX). While the Cisco UCS C-Series connectivity using Cisco Nexus FEX is still supported and recommended for large scale Cisco UCS C-Series server deployments, direct attached design allows customers to connect and manage Cisco UCS C-Series servers on a smaller scale without buying additional hardware. In this VersaStack design, Cisco UCS C-Series servers can be directly attached to Cisco UCS FI using two 10/40Gbps converged connections (one connection to each FI) as shown in Figure 17.

Figure 17 Cisco UCS C220 Connectivity



Cisco UCS Server Configuration for VSphere

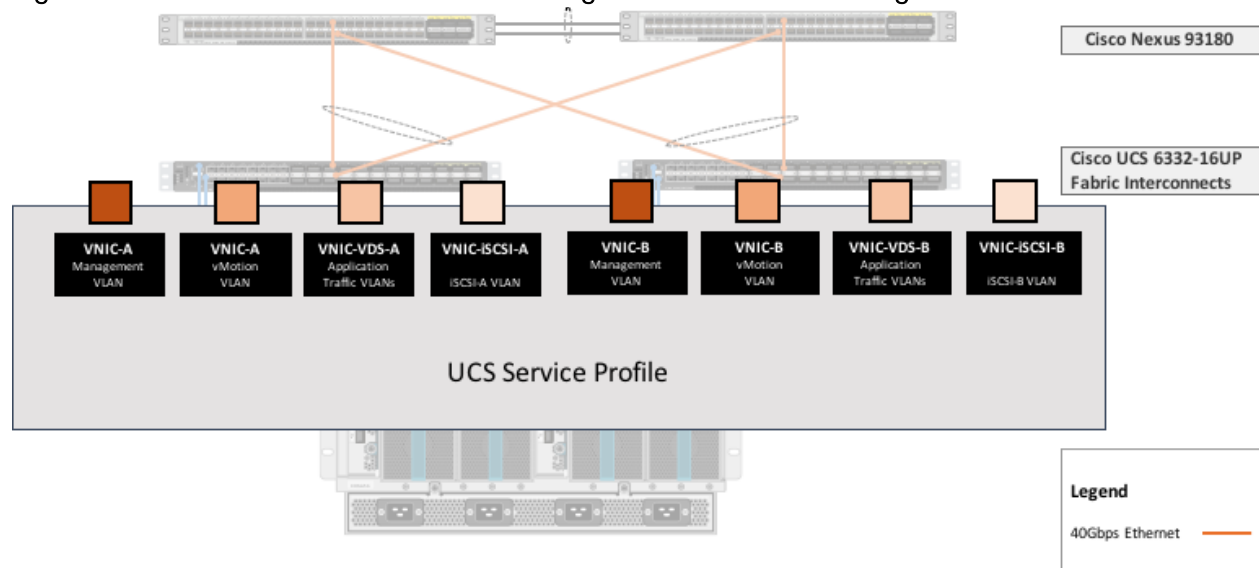
The ESXi nodes consist of Cisco UCS B200-M5 series blades with Cisco 1340 VIC. These nodes are allocated to a VMware High Availability (HA) cluster to support infrastructure services and applications. At the server level, the Cisco 1340 VIC presents multiple virtual PCIe devices to the ESXi node and the vSphere environment identifies these interfaces as vmnics or VMhbas. The ESXi operating system is unaware of the fact that the NICs or HBAs are virtual adapters.

In the VersaStack design for iSCSI storage, eight vNICs are created and utilized as follows (Figure 18):

- One vNIC for iSCSI-A traffic
- One vNIC for iSCSI-B traffic
- Two vNICs for in-band management traffic
- Two vNICs for vMotion traffic
- Two vNICs for application related data including storage access if required. These vNICs are assigned to a distributed switch (vDS)

These vNICs are pinned to different Fabric Interconnect uplink interfaces and are assigned to separate vSwitches and virtual distributed switches based on type of traffic. The vNIC to vSwitch assignment is covered later in the document.

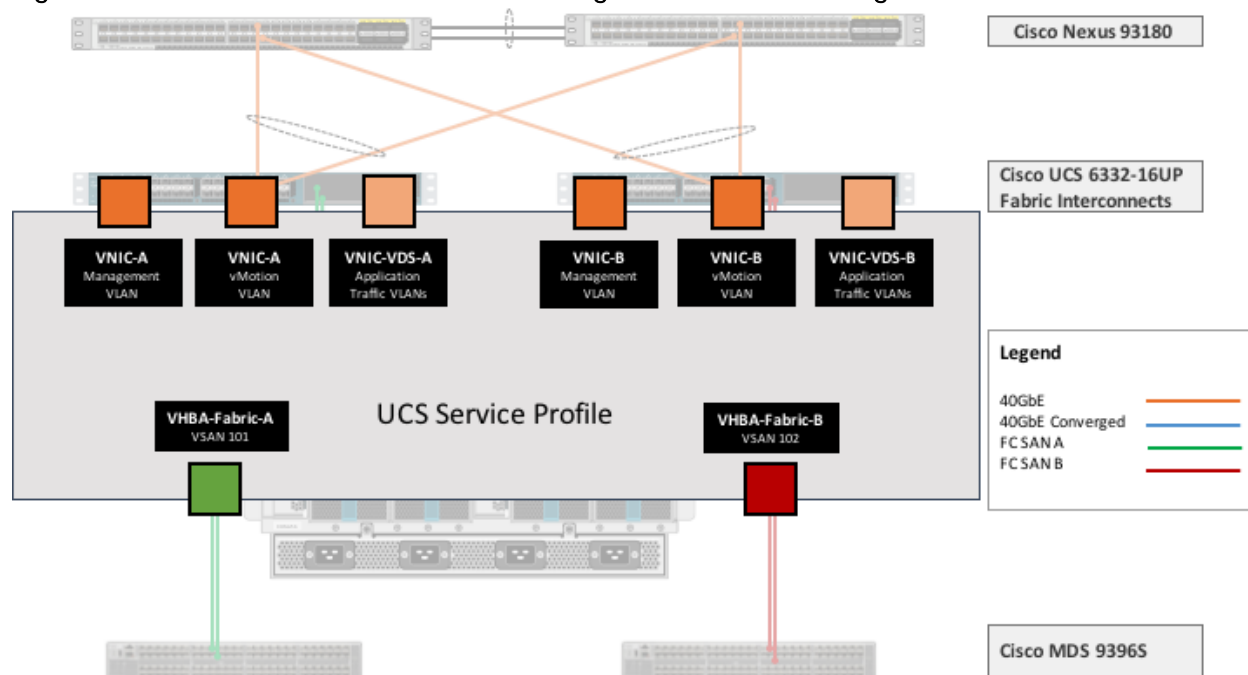
Figure 18 Cisco UCS –Server Interface Design for iSCSI based Storage



In the VersaStack design for FC storage, six vNICs and two vHBAs are created and utilized as follows (Figure 19):

- Two vNICs for in-band management traffic
- Two vNICs for vMotion traffic
- Two vNICs for application related data including application storage access if required. These vNICs are assigned to APIC controlled distributed switch
- One vHBA for VSAN-A FC traffic
- One vHBA for VSAN-B FC traffic

Figure 19 Cisco UCS - Server Interface Design for FC based Storage



IBM Storage Systems

IBM SAN Volume Controller, IBM FlashSystem 900 and IBM Storwize V5030, covered in this VersaStack design, are deployed as high availability storage solutions. IBM storage systems support fully redundant connections for communication between control enclosures, external storage, and host systems.

Each storage system provides redundant controllers and redundant iSCSI and FC paths to each controller to avoid failures at path as well as hardware level. For high availability, the storage systems are attached to two separate fabrics, SAN-A and SAN-B. If a SAN fabric fault disrupts communication or I/O operations, the system recovers and retries the operation through the alternative communication path. Host (ESXi) systems are configured to use ALUA multi-pathing, and in case of SAN fabric fault or node canister failure, the host seamlessly switches over to alternate I/O path.

IBM SAN Volume Controller – I/O Groups

The SAN Volume Controller design is highly scalable and can be expanded up to eight nodes in one clustered system. An I/O group is formed by combining a redundant pair of SAN Volume Controller nodes. In this design document, one pair of SVC IBM 2145-SV1 nodes (I/O Group 0) were deployed:




Based on the specific storage requirements, the number of I/O Groups in customer deployments will vary.

IBM SAN Volume Controller – iSCSI Connectivity

To support iSCSI based IP storage connectivity, each IBM SAN Volume Controller is connected to each of the Cisco Nexus 93180 switch. The physical connectivity is shown in Figure 20. Two 10GbE ports from each IBM SAN Volume Controller are connected to each of the two Cisco Nexus 93180 switches providing an aggregate bandwidth of 40Gbps. In this design, 10Gbps Ethernet ports between the SAN Volume Controllers

I/O Group and the Nexus fabric utilize redundant iSCSI-A and iSCSI-B paths and can tolerate link or device failures. Additional ports can be easily added for additional bandwidth.

Figure 20 IBM SAN Volume Controller - iSCSI based Storage Access



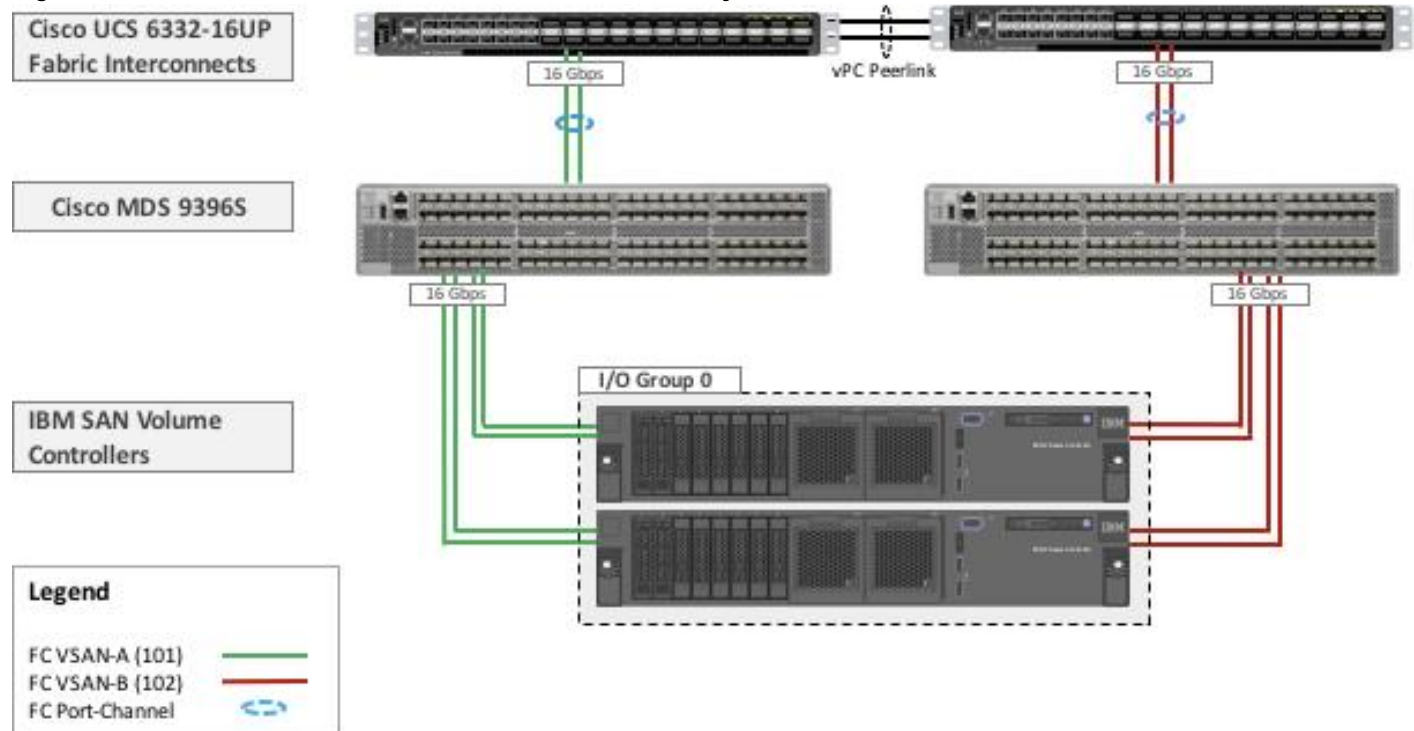
IBM SAN Volume Controller - FC Connectivity

To support FC based storage connectivity, each IBM San Volume Controller is connected to Cisco UCS Fabric Interconnect using redundant SAN Fabric configuration provided by two Cisco MDS 9396S switches. Figure 21 shows the resulting FC switching infrastructure.



This design guide covers both iSCSI and FC based connectivity options to highlight flexibility of the VersaStack storage configuration. A customer can choose an IP-only or an FC-only configuration design when virtualizing storage arrays behind SVC nodes.


Figure 21 IBM SAN Volume Controller - FC Connectivity



IBM SAN Volume Controller – Connectivity to Cisco UCS, FlashSystem 900, and Storwize V5030

Figure 22 illustrates the connectivity between SAN Volume Controllers and the FlashSystem 900 and Storwize V5030 storage arrays using Cisco MDS 9396S switches. All connectivity between the Controllers and the Storage Enclosure is 16 Gbps End-to-End. The connectivity from Cisco UCS to Cisco MDS is 32Gbps FC port-channel) (2 x 16Gbps FC ports). Cisco MDS switches are configured with two VSANs; 101 and 102 to provide FC connectivity between the ESXi hosts, IBM SVC nodes and the storage system enclosures. Zoning configuration on Cisco MDS switches enabled required connectivity between various elements connected to the SAN fabric.

Figure 22 IBM SAN Volume Controller – Connectivity to the Storage Devices



IBM SVC – Tiered Storage

The SAN Volume Controller makes it easy to configure multiple tiers of storage within the same SAN Volume Controller cluster. The system supports single-tiered storage pools, multi-tiered storage pools, or a combination of both. In a multi-tiered storage pool, MDisk's with more than one type of disk tier attribute can be mixed e.g. a storage pool can contain a mix of generic_hdd and generic_ssd MDisk's.








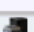

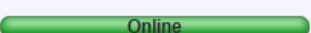
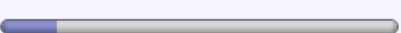


In the VersaStack with UCS M5 and IBM SVC design, both single-tiered and multi-tiered storage pools were utilized as follows:

- **GOLD**: All-Flash storage pool based on IBM FlashSystem 900
- **SILVER**: EasyTier* enabled mixed storage pool based on IBM FlashSystem 900, IBM Storwize V5030's SSD and IBM Storwize15K Enterprise class storage
- **BRONZE**: SAS based storage on IBM Storwize V5030 to support less I/O intensive volumes such as the SAN Boot LUNs for ESXi Hosts, etc.

* Easy Tier monitors the I/O activity and latency of the extents on all volumes in a multi-tier storage pool over a 24-hour period. It then creates an extent migration plan that is based on this activity and dynamically moves high activity (or hot extents) to a higher disk tier within the storage pool. It also moves extents whose activity dropped off (or "cooled") from the high-tier MDisk's back to a lower-tiered MDisk.

Figure 23 outlines a sample configuration to show how various MDisk's are created and assigned to storage pools for creating the multi-tiered storage system discussed above.

Figure 23 IBM SAN Volume Controller - Multi-tiered Storage Configuration

Name	State	Capacity	Storage System - LUN
 Bronze	 Online		71.00 GiB / 10.00 TiB (1%)
 V5030-SVC-Bronze	✓ Online	10.00 TiB	Storwize V5030 Node1 - 00000000000...
 Gold	 Online		3.27 TiB / 10.00 TiB (33%)
 FS900-SVC-Gold	✓ Online	10.00 TiB	Flash System 900 - 0000000000000000
 Silver	 Online		1.98 TiB / 15.00 TiB (13%)
 FS900-SVC-Silver	✓ Online	10.00 TiB	Flash System 900 - 00000000000000001
 V5030-SVC-Silver	✓ Online	5.00 TiB	Storwize V5030 Node1 - 00000000000...

Network Design


In this VersaStack design, a pair of redundant Cisco Nexus 93180 switches provide Ethernet switching fabric for storage and application communication including communication with the existing enterprise networks. Similar to previous versions of VersaStack, the core network constructs such as virtual port channels (vPC) and VLANs plays an important role in providing the necessary Ethernet based IP connectivity.

Virtual Port-Channel Design

In the current VersaStack design, following devices are connected to the Nexus fabric using a vPC (Figure 24):

- Cisco UCS fabric interconnects
- Connection to In-Band management infrastructure switch

Figure 24 Network Design – vPC Enabled Connections



VLAN Design

To enable connectivity between compute and storage layers of the VersaStack and to provide in-band management access to both physical and virtual devices, several VLANs are configured and enabled on various paths. The VLANs configured for the infrastructure services include:

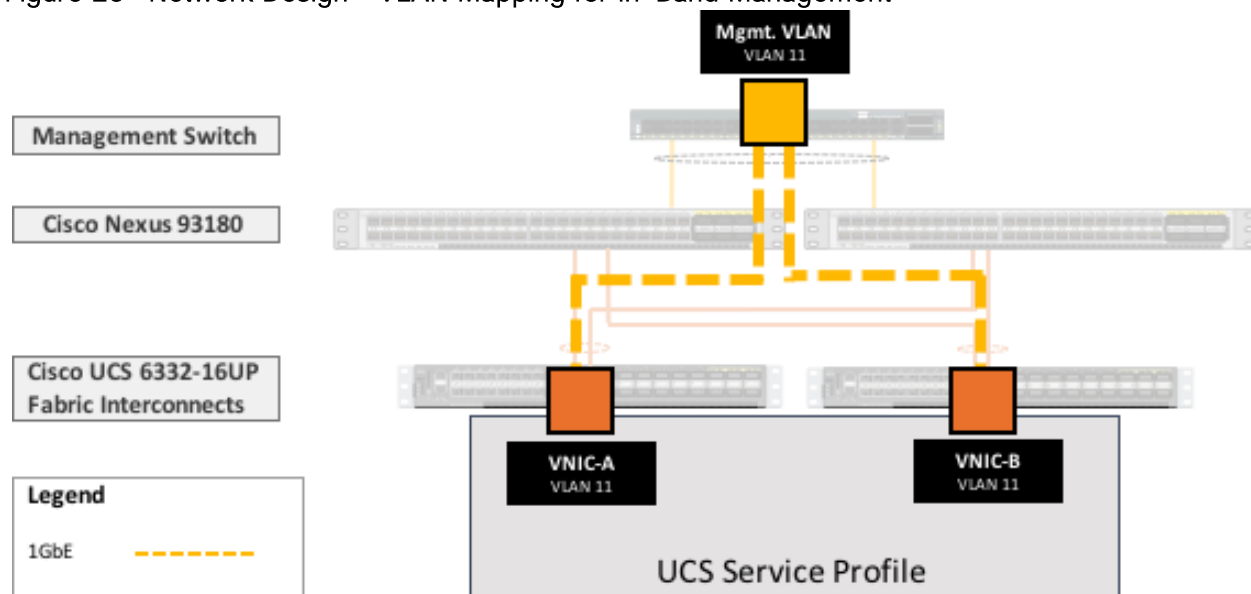
- iSCSI VLANs to provide access to iSCSI datastores including boot LUNs
- Management and vMotion VLANs used by compute and vSphere environment
- Application VLANs used for application and virtual machine communication

These VLAN configurations are described in the following sections.

In-Band Management VLAN Configuration

To provide in-band management access to the ESXi hosts and the infrastructure/management Virtual Machines (VMs), the existing management infrastructure switch is connected to both the Cisco Nexus 93180 switches using a vPC. In this design, VLAN 11 is the pre-existing management VLAN on the infrastructure management switch. Within Cisco UCS service profile, VLAN 11 is enabled on vNIC-A and vNIC-B interfaces as shown in Figure 25.

Figure 25 Network Design – VLAN Mapping for In-Band Management



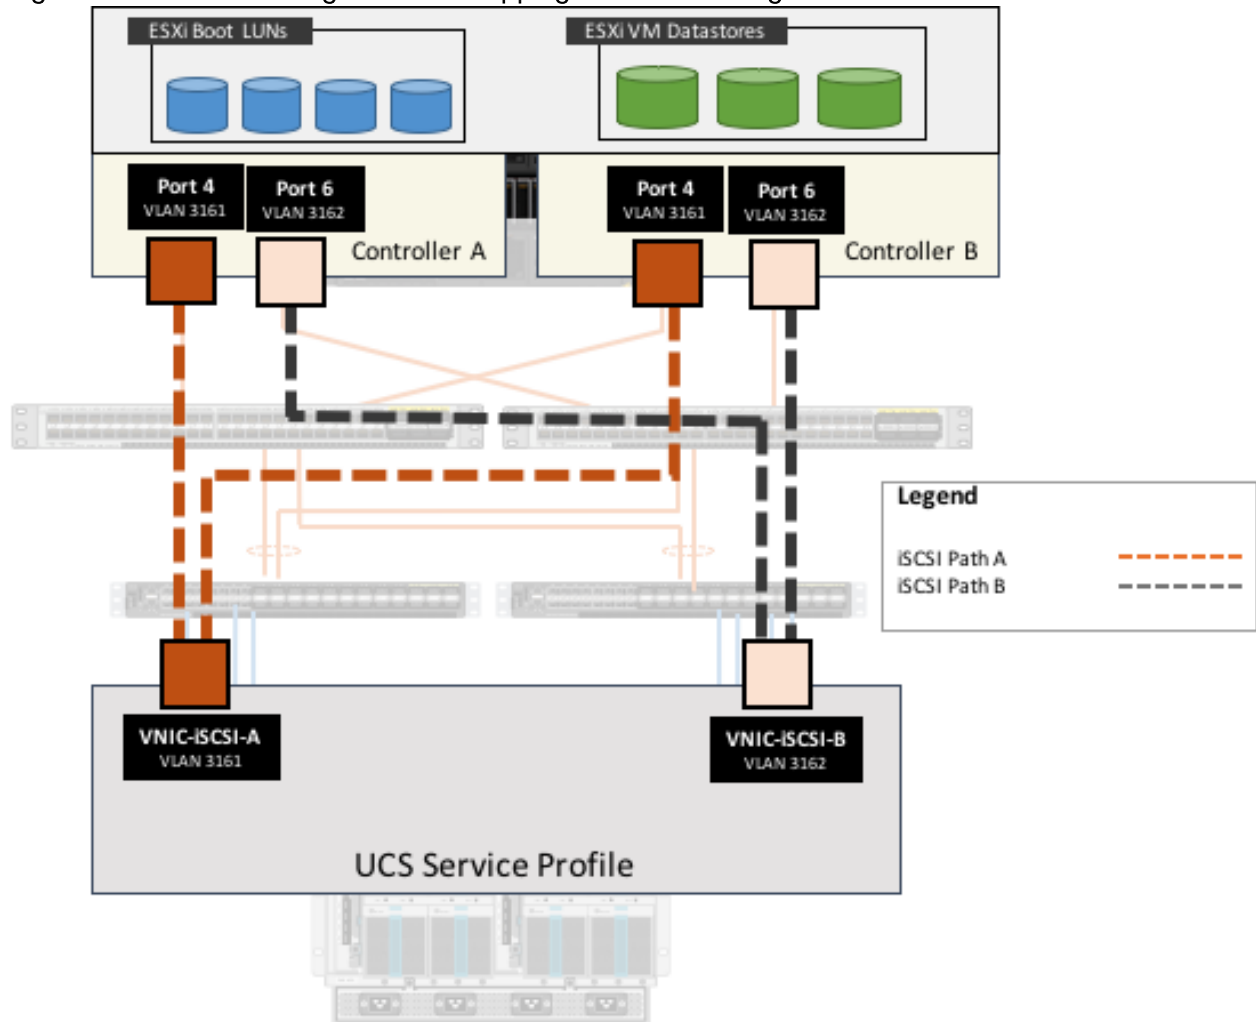
iSCSI VLAN Configuration

To provide redundant iSCSI paths, two VMkernel interfaces are configured to use dedicated NICs for host to storage connectivity. In this configuration, each VMkernel port provided a different path that the iSCSI storage stack and its storage-aware multi-pathing plug-ins can use.

To setup iSCSI-A path between the ESXi hosts and the IBM SVC nodes, VLAN 3161 is configured on the Cisco UCS, Cisco Nexus and on the IBM SVC interfaces. To setup iSCSI-B path between the ESXi hosts and the IBM SVC, VLAN 3162 is configured on the Cisco UCS, Cisco Nexus and on the appropriate IBM SVC node interfaces. Within Cisco UCS service profile, these VLANs are enabled on vNIC-iSCSI-A and vNIC-

iSCSI-B interfaces respectively as shown in Figure 26. The iSCSI VLANs are set as native VLANs on the vNICs to enable boot from SAN functionality.


Figure 26 Network Design – VLAN Mapping for iSCSI Storage Access



Virtual Machine Networking VLANs for VMware vDS

The VLANs needed to enable VM networking has to be created on the Nexus switches and have to be defined on Cisco UCS and enabled on the vNICs; vNIC-VDS-A and vNIC-VDS-B as shown in Figure 27. The picture below shows one VLAN created, multiple VLANs can be created as needed for virtual machine networking.

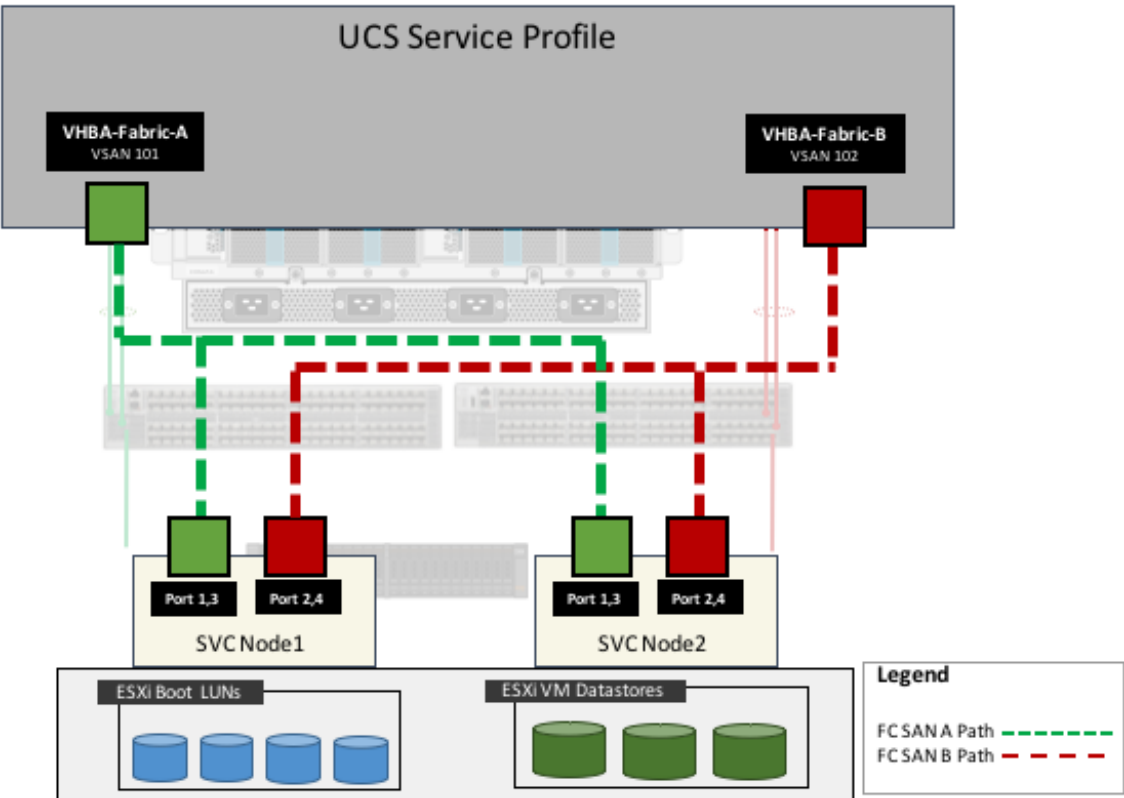
Figure 27 Network Design – VLANs for VM Networks with VMware VDS



VSAN Design for Host to Storage Connectivity

In the VersaStack design, isolated fabric (A and B) are created on the Cisco MDS 9396S switches, VSAN 101 and VSAN 102 provide the dual SAN paths as shown in Figure 28. Use of SAN zoning on Cisco MDS switches allows the isolation of traffic within specific portions of the storage area network. Cisco UCS FIs connect to the Cisco MDS switches using Port Channels while IBM SVC controllers are connected using independent FC ports.

Figure 28 VSAN Design – Cisco UCS Server to IBM SVC Connectivity



Virtual Switching Architecture

Application deployment utilizes port groups on VMware distributed switch (VDS). However, for some of the core connectivity such as out of band management access, vSphere vMotion and storage LUN access using iSCSI, vSphere vSwitches are deployed. To support this multi-vSwitch requirement, multiple vNIC interfaces are setup in Cisco UCS services profile and storage, management and VM data VLANs are then enabled on the appropriate vNIC interfaces. Figure 29 shows the distribution of VMkernel ports and VM port-groups on an iSCSI connected ESXi server. For an ESXi server, supporting iSCSI based storage access, In-band management and vMotion traffic is handled by infrastructure services vSwitches and iSCSI-A and iSCSI-B traffic is handled by two dedicated iSCSI vSwitches. The resulting ESXi host configuration therefore has a combination of four vSwitches and a single distributed switch which handles application specific traffic.

Figure 29 ESXi Host vNIC and vmk Distribution for iSCSI Based Storage Access

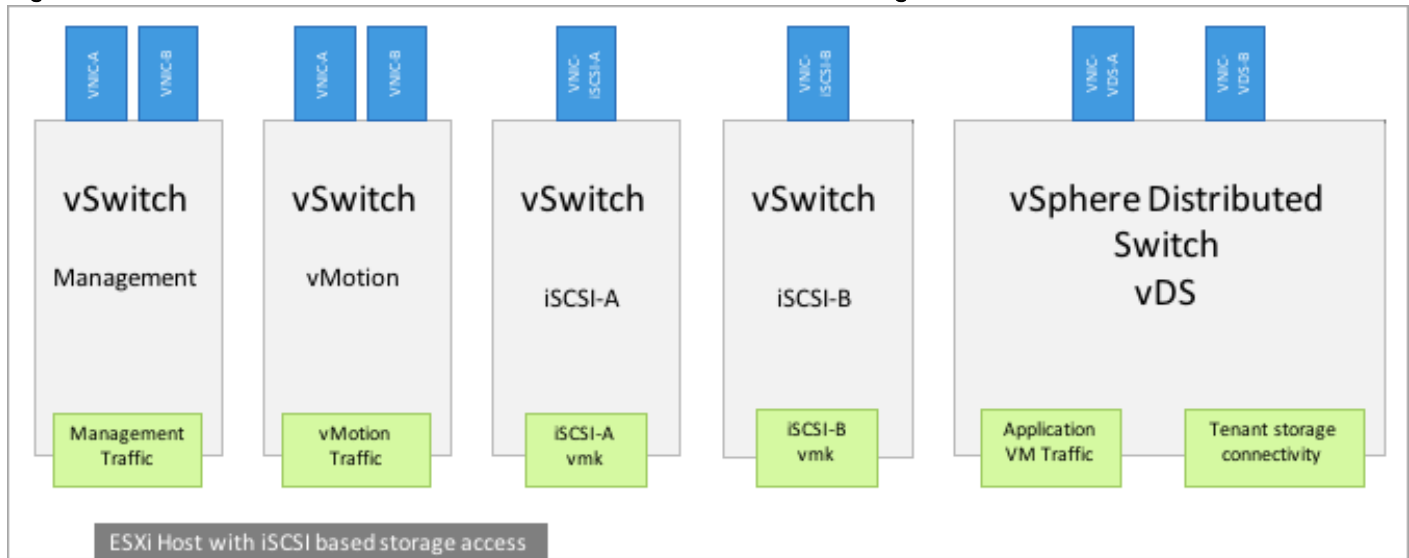
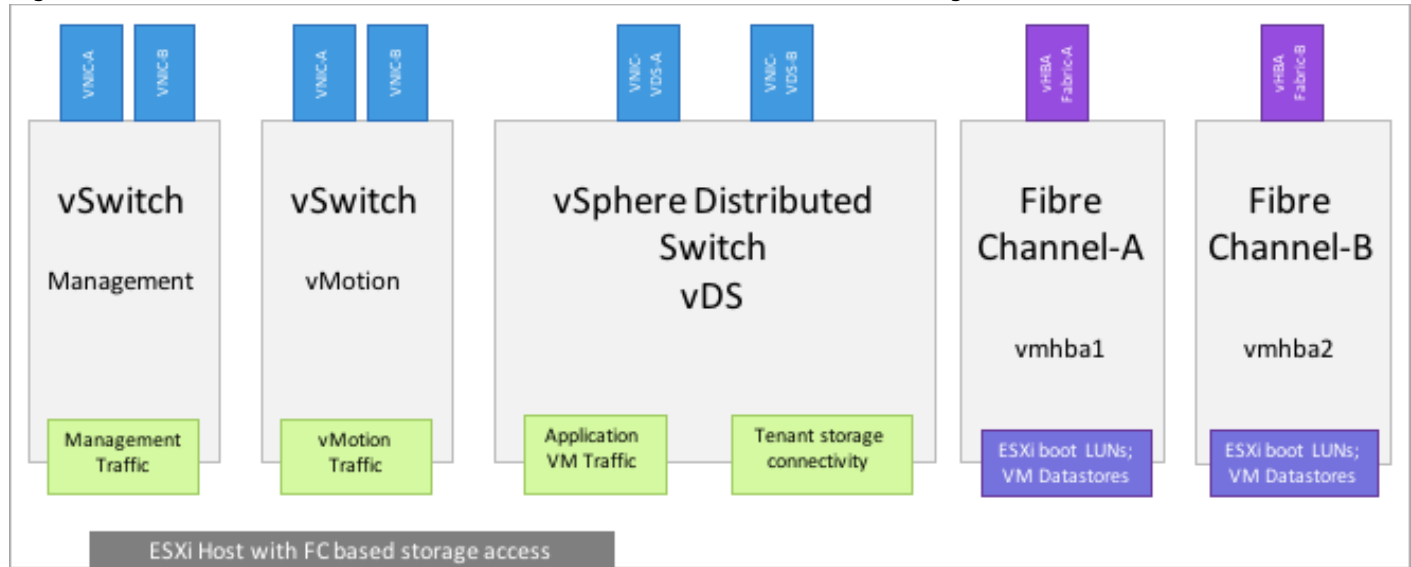


Figure 30 below shows similar configuration for an FC connected ESXi host. In this case, In-band management and vMotion traffic is handled by two dedicated vSwitches but the two vSwitches for iSCSI traffic are not needed. The Fibre Channel SAN-A and SAN-B traffic is handled by two dedicated vHBAs. The resulting ESXi host configuration therefore has a combination of one vSwitch and one (vDS) distributed switch.

Figure 30 ESXi Host vNIC, vHBA and vmk Distribution for FC Based Storage Access



Design Considerations

VersaStack designs incorporate connectivity and configuration best practices at every layer of the stack to provide a highly available best performing integrated system. VersaStack is a modular architecture that allows customers to adjust the individual components of the system to meet their particular scale or performance requirements. This section covers some of the design considerations for the current design and a few additional design selection options available to the customers.

Selection of I/O components has a direct impact on scale and performance characteristics when ordering the Cisco components.

Cisco Unified Computing System Fabric Interconnects

The third generation of Cisco Unified Fabric used in this solution is designed to fit easily into a Cisco UCS environment comprising Cisco UCS fabric extenders, VICs, and Cisco Nexus 9300 platform switches. Cisco UCS 6300 Series Fabric Interconnects have been carefully designed to combine the cost-saving advantages of merchant silicon with custom Cisco innovations. They provide an intelligent progression from previous-generation fabric interconnects. Two Cisco UCS 6332 models; Cisco UCS 6332 and Cisco UCS 6332-16UP are offered to provide deployment flexibility.

This VersaStack solution leveraged Cisco UCS 6332-16UP Fabric Interconnects to enable 40Gbps LAN connectivity along with 16Gbps Fibre channel connectivity to the IBM SVC controllers.

The capabilities of all Fabric Interconnect models that can be part of VersaStack are summarized below.

Table 1 Key Differences of the FI 6200 Series and FI 6300 Series

	FI 6200 Series		FI 6300 Series	
Features	FI 6248	FI 6296	FI 6332	FI 6332-16UP
Height	1RU	2RU	1RU	1RU
Form factor	Modular	Modular	Fixed	Fixed
Physical ports	48 (using GEM)	96 (using GEM)	32	40
Max 10G ports	48	96	$96^* + 2^{**} = 98$	$72^* + 16UP = 88$
Max 40G ports	-	-	32	24
Max FC ports	48 x 2/4/8G FC	96 x 2/4/8G FC	-	16 x 4/8/16G FC
Max unified ports	48	96	-	16
Expansion module	1 x Exp module	3 x Exp modules	-	-

* Using 40G to 4x10G breakout cables ** Requires QSA module

Cisco Unified Computing System I/O Component Selection

For blade chassis, the Cisco UCS 2304 fabric extender used in the design provides 160 Gbps of bandwidth to each of the two fabrics in a Cisco UCS 5108 Blade Server Chassis, for up to 320 Gbps of bandwidth. Each fabric extender provides eight 40-Gbps server links and four 40-Gbps SFP+ uplinks to the fabric interconnects. Server links are supported by multiple 10GBASEKR connections across the blade chassis midplane. These connections are combined into 40-Gbps links when Cisco UCS VIC 1340s are used with the optional Port Expander Card, or left as 10-Gbps links otherwise. In addition to extending the network to blade servers, the fabric extender contains logic that extends the management network from Cisco UCS Manager into each server's Cisco integrated management console.

To preserve any existing investment, the Cisco UCS 6300 Series Fabric Interconnects support the preceding generation of 10Gbps fabric extenders: the Cisco UCS 2204XP (up to 80 Gbps per chassis) and 2208XP (up to 160Gbps per chassis) Fabric Extenders. This support allows you to upgrade your fabric interconnects and your blade server chassis as your connectivity demands require.

Table 2 IOM key differences, IOM 2200 series & IOM 2304


	IOM 2200 Series		IOM 2300
Features	2204XP	2208XP	2304
Form Factor	Modular	Modular	Modular
Total ports	20	40	12
Max 10G ports	20	40	-
Max 40G ports	-	-	12
Uplink ports	4	8	4
Server ports	16	32	8

Cisco UCS VIC 1340 Connectivity to Cisco UCS 2200 Series Fabric Extender


The UCS nodes in this VersaStack design consist of Cisco 1340 VIC's, the Cisco UCS VIC 1340 can be connected to a Cisco UCS 2200 Series Fabric Extender either by itself or with a port expander inserted in the blade server's mezzanine connector. By itself, the Cisco UCS VIC 1340 provides four 10Gbps interfaces, combined by port channels into two 20Gbps interfaces, with one connected to each fabric (Figure 31).

Figure 31 Cisco UCS VIC 1340 to Cisco UCS 2200 Connectivity

Cisco VIC 1340 without Port Expander



Cisco VIC 1340 with Port Expander




If you add a mezzanine-based port expander to the blade server, the Cisco UCS VIC 1340 provides two 40-Gbps interfaces: one to each fabric. However, the interfaces are still four 10Gbps interfaces combined in a port channel, and not true 40-Gbps interfaces.

Cisco UCS VIC 1340 Connectivity to Cisco UCS 2304 Fabric Extender

With the inclusion of the Cisco UCS 2304 Fabric Extender in VersaStack design, the 40-Gbps interfaces to each fabric become true native 40Gbps interfaces, and a hardware port channel is no longer required or used (Figure 32). True 40Gbps interfaces can now be defined for the operating system and implemented directly on the Cisco UCS VIC 1340.

Figure 32 Cisco UCS VIC 1340 to Cisco UCS 2304 Connectivity



Validated I/O Component Configurations

The validated I/O component configuration for this VersaStack with Cisco UCS M5 and IBM SVC design are:

- Cisco 6332-16UP Fabric Interconnects
- Cisco UCS B200 M5 with Cisco VIC 1340
- Cisco UCS FEX 2304

Cisco Unified Computing System Chassis/FEX Discovery Policy

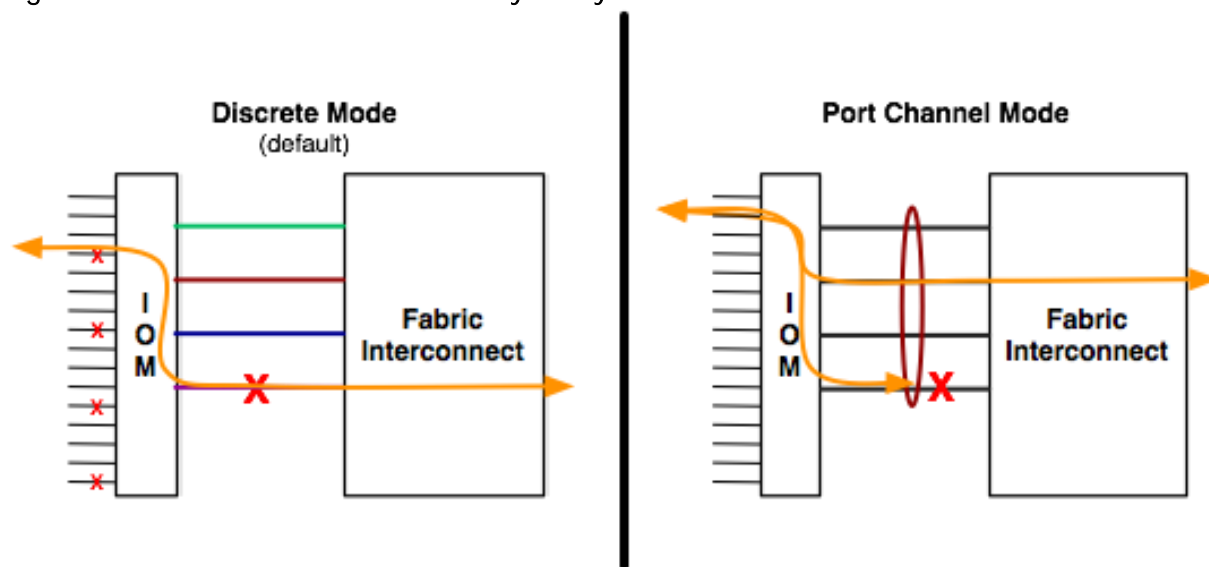
The fabric extenders for the Cisco UCS 5108 blade chassis are deployed with up to 4x40Gb to 6300 series fabric interconnects and up to 8x10Gb to 6200 series fabric interconnects. The links are automatically configured as port channels by the Chassis/FEX Discovery Policy within Cisco UCS Manager, which ensures rerouting of flows to other links in the event of a link failure.

The Cisco UCS chassis/FEX discovery policy determines how the system reacts when you add a new chassis or FEX. Cisco UCS Manager uses the settings in the chassis/FEX discovery policy to determine the minimum threshold for the number of links between the chassis or FEX and the fabric interconnect and whether to group links from the IOM to the fabric interconnect in a fabric port channel. Cisco UCS Manager

cannot discover any chassis that is wired for fewer links than are configured in the chassis/FEX discovery policy.

Cisco Unified Computing System can be configured to discover a chassis using Discrete Mode or the Port-Channel mode (Figure 33). In Discrete Mode, each FEX KR connection and therefore server connection is tied or pinned to a network fabric connection homed to a port on the Fabric Interconnect. In the presence of a failure on the external "link" all KR connections are disabled within the FEX I/O module. In Port-Channel mode, the failure of a network fabric link allows for redistribution of flows across the remaining port channel members. Port-Channel mode therefore is less disruptive to the fabric and hence recommended in the VersaStack designs.

Figure 33 Cisco UCS Chassis Discovery Policy - Discrete Mode vs. Port Channel Mode



Storage Design and Scalability

Virtualizing storage with IBM SVC helps make new and existing storage more effective. Currently, more than 370 storage controllers from various third parties can be virtualized behind SVC. IBM SVC configurations start with a single 'I/O group' comprised of two high performance servers, referred to as 'nodes' or 'storage engines'. Up to four pairs of nodes may be clustered into a single storage system with up to 128 cores, 2 Terabytes (TB) of system memory, 128 FC ports and up to 2,944 drives - which can support a total of 32 Petabytes (PB) of virtualized storage capacity. These nodes are all interconnected on a 16G FC interconnect provided by the redundant Cisco MDS 9396S based fabrics. Multiple 16G links provide high availability for inter-cluster communications between nodes, as well as parallelism when accessing block data from the I/O group.

SVC supports up to twenty 12 Gbps serial attached SCSI (SAS) expansion enclosures per I/O group for up to 736 disk drives. These enclosures provide low-cost storage capacity to complement external virtualized storage. SVC also supports the configuration of flash drives within nodes themselves.

For a current list of supported virtualized controllers, see the Support Hardware Website for SVC:

<http://www-01.ibm.com/support/docview.wss?uid=ssg1S1005828>.

Management Network Design

This VersaStack design uses two different networks to manage the solution:

- An out-of-band management network to configure and manage physical compute, storage and network components in the solution. Management ports on each physical device (Cisco UCS FI, IBM Storage Controllers, Cisco Nexus and Cisco MDS switches) in the solution are connected to a separate, dedicated management switch.
- An in-band management network to manage and access ESXi servers and VMs (including infrastructure VMs hosted on VersaStack). Cisco UCS does allow out-of-band management of servers and VMs by using disjoint layer 2 feature but for the current design, In-Band manageability is deemed sufficient.



If a customer decides to utilize Disjoint Layer-2 feature, additional uplink port(s) on the Cisco UCS FIs are required to connect to the management switches and additional vNICs have to be associated with these uplink ports. The additional vNICs are necessary since a server vNIC cannot be associated with more than one uplink.

In certain data center designs, there is a single management segment which requires combining both out-of-band and in-band networks into one.

Virtual Port Channel Configuration

Virtual Port Channel (vPC) allows Ethernet links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single Port Channel. vPC provides a loop-free topology and enables fast convergence if either one of the physical links or a device fails. In the VersaStack design, when possible, vPC is a preferred mode of Port Channel configuration.

vPC on Nexus switches running in NXOS mode requires a peer-link to be explicitly connected and configured between peer-devices (Nexus 93180 switch pair). In addition to the vPC peer-link, the 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 out-of-band management network.

Jumbo Frame Configuration

Enabling jumbo frames in a VersaStack environment optimizes throughput between devices by enabling larger size frames on the wire while reducing the CPU resources to process these frames. VersaStack supports wide variety of traffic types (vMotion, NFS, iSCSI, control traffic, etc.) that can benefit from a larger frame size. In this validation effort the VersaStack was configured to support jumbo frames with an MTU size of 9000. In VMware vSphere, the jumbo frames are configured by setting MTU sizes at both vSwitches and VMkernel ports. On IBM storage systems, the interface MTUs are modified to enable the jumbo frame. In this validation effort, the VersaStack was configured to support jumbo frames with an MTU size of 9000.



When setting the Jumbo frames, it is important to make sure MTU settings are applied uniformly across the stack to prevent fragmentation and the negative performance.

Compute and Virtualization High Availability Considerations

In the VersaStack solution, each ESXi server is deployed using vNICs and vHBAs that provide redundant connectivity to the unified fabric. All of the server NICs and HBAs are configured to use both the Cisco UCS FIs to avoid traffic disruptions.

VMware vCenter is used to deploy VMware HA clusters to allow VMs to failover in the event of a server failure. VMware vMotion and VMware HA are enabled to auto restart VMs after a failure. Host Monitoring is enabled to monitor heartbeats of all ESXi hosts in the cluster for faster detection. Admission Control is also enabled on the blade servers to ensure the cluster has enough resources to accommodate a single host failure.

VMware vSphere hosts use SAN multi-pathing to access LUNs on the IBM storage devices. If any component (NIC, HBA, FEX, FI, Cisco MDS, IBM controller, cables) along a path fails, all storage traffic will reroute to an alternate path. When both paths are active, traffic is load balanced.

Deployment Hardware and Software

Hardware and Software Revisions

Table 3 below outlines the hardware and software versions used for the 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 visit the following links for more information:

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

Table 3 Validated Hardware and Software Revisions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6300 Series, Cisco UCS B-200 M5	3.2(3a)	Includes the Cisco UCS-IOM 2304, Cisco UCS Manager, and Cisco UCS VIC 1340
	Cisco nenic Driver	1.0.16.0	Ethernet driver for Cisco VIC
	Cisco fnic Driver	1.6.0.37	FCoE driver for Cisco VIC
Network	Cisco Nexus Switches	7.0(3)I4(7)	NXOS
	Cisco MDS 9396S	7.3(0)DY(1)	FC switch firmware version
Storage	IBM SVC	7.8.1.4	Software version
	IBM FlashSystem 900	1.4.5.0	Software version
	IBM Storwize V5030	7.8.1.4	Software version
Software	VMware vSphere ESXi	6.5 update 1	Software version
	VMware vCenter	6.5 update 1	Software version

Validation

Test Plan

The VersaStack solution was validated by deploying multiple VMs running IOMeter tool. The system was validated for resiliency by failing various aspects of the system under the load. The types of tests executed on the system are as follows:

Cisco UCS Validation

- Failure and recovery of links from Cisco UCS Chassis (IOM) to FI-A and FI-B, one at a time
- Rebooting Cisco UCS FI, one at a time
- Removing the physical cables between FI and Cisco Nexus 93180 switches to simulate path failure

Network Validation

- Fail/power off both Cisco 93180 switches, one after other
- Failure and recovery of links with in vPC from UCS FI and Cisco Nexus 93180 switches
- Failure and recovery of physical links from Cisco Nexus 93180 switches and the management switch

Storage Validation

- Failure and recovery of the Cisco MDS switches, one at a time
- Failure and recovery of the links between Cisco MDS and IBM storage controllers
- Failure and recovery of links between Cisco UCS FI and Cisco MDS
- Failure and recovery of the links between Cisco Nexus 93180 switches and IBM SVC Nodes

vSphere Validation

- Failure and recovery of ESXi hosts in a cluster (rebooting of hosts, shutting down of hosts etc.)
- In case of a host failure, verify VM auto restart within the HA cluster
- VM vMotion across ESXi servers
- Storage vMotion between iSCSI and FC datastores (when applicable)
- Restart/Isolate vCenter server

Bill of Materials

To find various components of VersaStack system, complete the following steps:

1. Go to the Main CCW page: <https://apps.cisco.com/Commerce/home>.
2. Under Find Products and Solutions, Click on the Search for solutions.
3. Type VersaStack. System will pull all the VersaStack variations.
4. Select one of the solutions and click View Components.

Summary

VersaStack delivers a platform for enterprise and cloud datacenters using Cisco UCS Blade Servers, Cisco Fabric Interconnects, Cisco Nexus 9000 switches, Cisco MDS switches, and Fibre Channel or iSCSI attached IBM Storage Arrays.

VersaStack with Cisco UCS M5 and IBM SVC architecture aligns with the converged infrastructure configurations and best practices as identified by previous VersaStack releases. The system includes hardware and software compatibility support between all components and aligns to the configuration recommendations for each of these components. The VersaStack design described in this document has been validated for resiliency (under fair load) and fault tolerance during system upgrades, component failures, and partial, as well as, complete loss of power scenarios.

References

Products and Solutions

Cisco Unified Computing System:

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

Cisco UCS 6300 Series Fabric Interconnects:

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

Cisco UCS 5100 Series Blade Server Chassis:

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

Cisco UCS B-Series Blade Servers:

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

Cisco UCS C-Series Rack Servers:

<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/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:

https://www.vmware.com/tryvmware_tpl/vsphere-55_evalcenter.html

IBM SAN Volume Controller

<http://www-03.ibm.com/systems/storage/software/virtualization/svc/>

IBM FlashSystem 900:

<http://www-03.ibm.com/systems/storage/flash/900/>

IBM Storwize V5000:

http://www-03.ibm.com/systems/storage/disk/storwize_v5000/overview.html

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.

Warren Hawkins, Virtualization Test Specialist for IBM Spectrum Virtualize, IBM

Working as part of the development organization within IBM Storage, Warren Hawkins is also a speaker and published author detailing best practices for integrating IBM Storage offerings into virtualized infrastructures. Warren has a background in supporting Windows and VMware environments working in second-line and third-line support in both public and private sector organizations. Since joining IBM in 2013, Warren has played a crucial part in customer engagements and, using his field experience, has established himself as the Test Lead for the IBM Spectrum Virtualize™ product family, focusing on clustered host environments.