VersaStack for Data Center Design Guide

With Cisco Unified Computing System and Cisco Nexus 9000 Series Switches, IBM Storwize V7000 Unified Storage and VMware vSphere 5.5 Update 1

Last Updated: April 10, 2015

Cisco Validated Design

Building Architectures to Solve Business Problems

About the Authors

Jeff Fultz, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems

Jeff has over 20 years of experience in Information Systems and Application Development focusing on data center management, backup, and virtualization optimization related technologies. As a member of multiple engineering solution teams, Jeff has deployed, designed and tested a wide variety of enterprise solutions encompassing Cisco, VMware, HyperV, SQL, and Microsoft Exchange to name a few. Jeff is a Microsoft Certified System Engineer with multiple patents filed in the datacenter solutions space.

Shivakumar Shastri, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems

Shiva has over 18 years of experience in various aspects of IT Infrastructure Engineering. Recent areas of focus include cloud automation, hyper-convergence, security and systems engineering of integrated stacks.

Sally Neate, Test Architect, Systems and Technology Group, IBM

Sally Neate has over 10 years of engineering experience in IBM's Systems and Technology group, and has been involved in the development of the Storwize product range from its inception. As system test lead for the Storwize and SAN Volume Controller 7.2 and 7.3 releases, Sally has designed and tested systems to meet the demands of a wide range of mid-range and enterprise environments.

About Cisco Validated Design (CVD) Program

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

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

VersaStack for Data Center Design Guide

About this Document

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 and IBM® VMware vSphere 5.5 Update 1 on VersaStack® solution with the Cisco Nexus 9000 switches. A VersaStack solution is a validated approach for deploying Cisco and IBM technologies as a shared cloud infrastructure.

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.

Introduction

Industry trends point to a 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 a 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.

VersaStack™ is a predesigned, integrated platform architecture for the data center that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus® family of switches, and IBM Storwize V7000 Unified storage. 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 e such as firmware and operating system upgrades, switch, cable and hardware failures, and loss of power has also been ascertained.. This document describes a solution with VMware vSphere 5.5 Update 1 built on the VersaStack. The document from Cisco and IBM discusses design choices made and best practices followed in deploying the shared infrastructure platform.

Problem Statement

Customers looking to transition toward shared infrastructure, with or without cloud functionality, face a number of challenges. Most of these challenges are due to not knowing where to start and how to proceed in order to achieve expected levels of efficiency and reach business objectives. Even with customers experienced in hosting multi-tenant integrated platforms, there is a need to address critical areas leading up to an optimal solution. These areas include compatibility of selected components, scalability with investment protection and efficient and easy management of the integrated stack. All these factors are major considerations in realizing expected return on investment (ROI) as well as achieving objectives of recovery time (RTO) and recovery point (RPO).

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.

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 shows a list of features in focus on VersaStack:

Table 1 VersaStack Component Features

IBM Storwize V7000 Unified Storage	Cisco UCS and Nexus 9000 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
GTS Private modular cloud	Application Centric Infrastructure (ACI)
Integrated GTS support and delivery services	Investment Protection
Big data and analytics enablement	Scalability
	Automation

Integrated System

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

Fabric Infrastructure Resilience

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

Network Virtualization

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

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

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

System 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 Storwize family storage

These components are connected and configured according to best practices of both Cisco and IBM and provide the ideal platform for running a variety of enterprise workloads with confidence. The reference architecture covered in this document leverages the Cisco Nexus 9000 for the switching element. 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 below (Cisco Unified Computing System, Cisco Nexus, and IBM Storwize) 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.

Design Principles and Objectives

Design principles are guidelines to help ensure the architectural implementation fulfills its mission. It provides a means to tie components and methods to the business objectives.

VersaStack addresses five primary design objectives:

- Availability: Helps ensure applications and services are accessible and ready to use.
- Scalability: Provide for capacity needs while staying within the architectural framework
- Versatility: Ability to quickly support new services without requiring infrastructure modifications.
- Efficiency. Facilitate efficient operation of the infrastructure through policy and API management.
- Simplicity: A strong focus on ease of use to help reduce deployment and operating costs.

The VersaStack architecture priorities scalability, availability and simplicity by using modular, redundant components, managed by centralized easy to use interfaces.

Some architectural principles and rationale for inclusion are described below:

- Least common mechanism: This principle highlights the need to globalize common/shared modules as it has the effect of reducing duplicates which can lead to higher efficiency and provide ease of maintenance. VersaStack implements centralized element managers such as Cisco UCS Manager and IBM Storwize V7000 software for fully scaled out deployments of compute and storage.
- Efficient Mediated Access: Application centric infrastructure (ACI) functionality, in VersaStack, consists of hardware and software components to provide performance with flexibility.

Inclusion of Cisco UCS Director has the potential for further consolidation at the management layer, leading to an even more agile and manageable integrated stack.

Note Performance and comprehensive security are key design criteria that are not directly addressed in this solution but will be addressed in other collateral, benchmarking, and solution testing efforts. This Design Guide validates the functionality and basic security elements.

VersaStack Design Benefits

The following are the VersaStack design benefits:

- VersaStack leverages hardware and software features of IBM Storwize V7000 Unified to transform data center efficiency while simplifying administrative tasks.
- VersaStack harnesses IBM Storwize V7000 Generation 2 hardware to support up to 1.5 PB of internal storage attached to a single V7000 control enclosure, and up to 4PB for a fully-scaled system. Hardware-assisted Real-time compression and flash optimization with Easy Tier further boost storage efficiency and performance.
- A key benefit of VersaStack is the virtualizing capabilities of Storwize software, which enables Administrators to centrally manage external storage volumes from a single point. Virtualization simplifies management tasks with the easy-to-use Storwize GUI, while leveraging all Storwize functional features to boost storage utilization and capability. Administrators can migrate data without disruption to applications, avoiding downtime for backups, maintenance and upgrades.
- VersaStack uses IBM Storwize V7000 Unified to integrate support for file and block data through one interface.
- File support features include Network File System version 4 (NFSv4) for high-level security, and greater performance with Server Message Block (SMB) 2.1. IBM Active Cloud Engine provides automated, policy-based and tiering of file data, with multi-writer capability that allows multiple users at different locations to be able to write to the same file set at a particular site.

VersaStack and Cisco Nexus 9000 Modes of Operation

The Cisco Nexus 9000 family of switches supports two modes of operation: NxOS 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 ACI.

The Cisco Nexus 9000 stand-alone mode VersaStack design consists of a single pair of Cisco Nexus 9000 top of rack switches. When in ACI fabric mode, the Cisco Nexus 9500 and 9300 switches are deployed in a spine-leaf architecture. Although the reference architecture covered in this document does not leverage ACI, it lays the foundation for customer migration to ACI by integrating the Cisco Nexus 9000 switches.

Application Centric Infrastructure (ACI) is a holistic architecture with centralized automation and policy-driven application profiles. ACI delivers software flexibility with the scalability of hardware performance. Key characteristics of ACI include:

- Simplified automation by an application-driven policy model
- Centralized visibility with real-time, application health monitoring
- Open software flexibility for DevOps teams and ecosystem partner integration
- Scalable performance and multi-tenancy in hardware

The future of networking with ACI is about providing a network that is deployed, monitored, and managed in a fashion that supports DevOps and rapid application change. ACI does so through the reduction of complexity and a common policy framework that can automate provisioning and managing of resources.

VersaStack: Cisco Nexus 9000 Design Options

The compute, network and storage components within VersaStack provide interfaces and protocol options for block (iSCSI and Fiber Channel) and file (NFS/CIFS) leading to flexibility and multiple design options. Some of the relevant options that utilize available interfaces and supported protocols are discussed below. The first design depicts a unified infrastructure consisting of both NFS access for VMware datastores and FC for SAN boot as well as additional VMFS datastores. An alternate configuration option in Figure 2 displays a block only deployment provided to point out the flexibility of the components and architecture.

Design Focus: VersaStack Unified design

Figure 1 details the VersaStack: Cisco Nexus 9000 design can be deployed to leverage both file and block storage with IBM Storwize V7000 Unified. As the illustrations show, the designs are fully redundant in the compute, network, and storage layers. There is no single point of failure from a device or traffic path perspective. The Storwize V7000 Controller depicted shows 1 enclosure that contains 2 canisters, 1 for each of the 2 cluster nodes, as well as 24 disk drives.

Figure 1 Scalable Multi-Protocol VersaStack Design

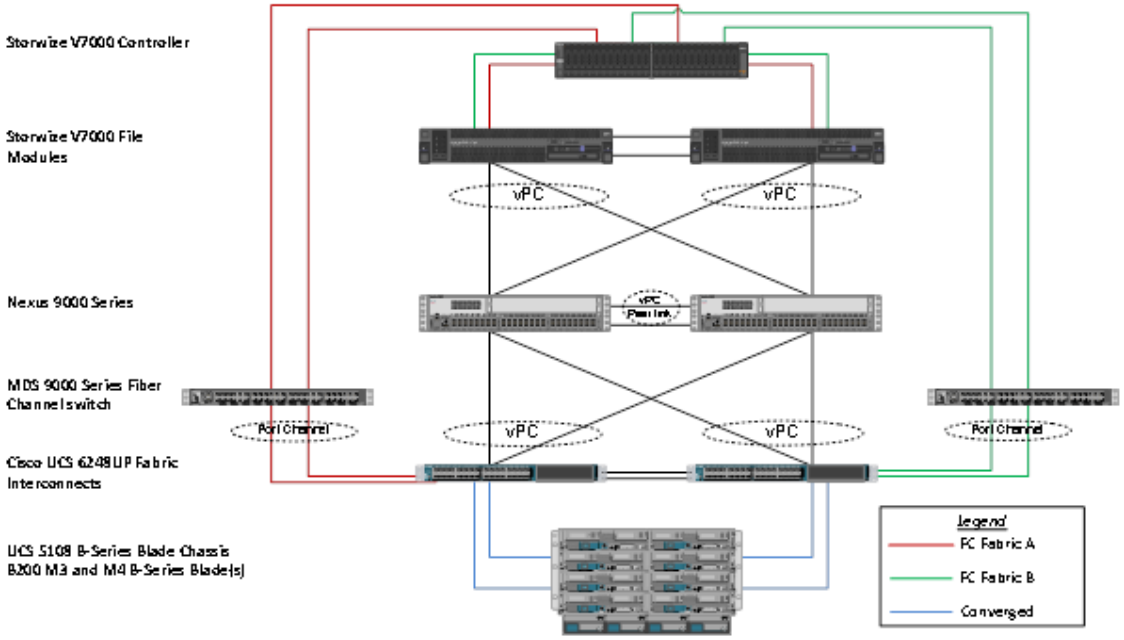

The VersaStack design incorporates an IP and FC-based storage solution that supports file access using NFS, and block access using FC. The solution provides a 10GbE-enabled, 40GbE-capable ethernet fabric defined by ethernet uplinks from the Cisco UCS Fabric Interconnects and from IBM Storwize V7000 storage connected to the Cisco Nexus switches through the V7000 file modules.

The Cisco Nexus design also incorporates an 8Gb-enabled, 16Gb-capable FC fabric for block access, defined by FC uplinks from the IBM Storwize V7000 control enclosure to the Cisco Nexus switches through Cisco MDS switches.

The Cisco Nexus 9000 design can be deployed with a dedicated SAN switching environment, or requiring no direct Fibre Channel connectivity using iSCSI for the SAN protocol.

As illustrated, link aggregation technologies play an important role, providing improved aggregate bandwidth and link resiliency across the solution stack. IBM Storwize V7000 Unified, Cisco Unified Computing System, and Cisco Nexus 9000 platforms support active port channeling using 802.3ad standard Link Aggregation Control Protocol (LACP). Port channeling is a link aggregation technique offering link fault tolerance and traffic distribution (load balancing) for improved aggregate bandwidth across member ports. In addition, the Cisco Nexus 9000 series features virtual PortChannel (vPC) capabilities. vPC allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single "logical" port channel to a third device, essentially offering device fault tolerance. vPC addresses aggregate bandwidth, link, and device resiliency. The Cisco UCS Fabric Interconnects and IBM Storwize V7000 Unified file modules both benefit from the Cisco Nexus vPC abstraction, gaining link and device resiliency as well as full utilization of a non-blocking Ethernet fabric.

Note The Spanning Tree protocol does not actively block redundant physical links in a properly configured vPC-enabled environment, so all ports are forwarding on vPC member ports.

The dedicated ethernet uplink design leverages the IP-based storage-capable IBM Storwize V7000 Unified system. From a storage traffic perspective, both standard LACP and the Cisco vPC link aggregation technologies play an important role in the VersaStack design. Figure 1 illustrates the use of dedicated 10GbE uplinks between the Cisco UCS fabric interconnects and the Cisco Nexus 9000 unified switches. vPC links between the Cisco Nexus 9000 and the IBM storage controllers' 10GbE provide a robust connection between host and storage.

Figure 1 shows the initial storage configuration of this solution with IBM Storwize V7000 Unified. The block storage configuration comprises a Storwize V7000 control enclosure, containing an HA pair of node canisters, and expansion enclosures housing additional disks. For file access, dual redundant V7000 file modules are connected to the V7000 control enclosure. Scalability is achieved by adding storage capacity (expansion enclosures) to an existing V7000 control enclosure, and by adding up to 3 more control enclosures or IO groups that can, in-turn, manage more expansion enclosures. The fully expanded Storwize V7000 storage system can scale up to 1056 disks.

Note IBM Storwize V7000 allows up to four control enclosures, for a total of eight control node canisters in one system. A single pair of V7000 file modules enables file access to a fully-scaled 8-node canister V7000 system.

For more information about the virtual design of the environment consisting of VMware vSphere, Cisco Nexus 1000v virtual distributed switching, and IBM Storwize V7000 Unified, refer to section Logical Build.

VersaStack Block Design

This design option displays FC/ISCSI only block based storage deployment. In this architecture example, the Cisco blade servers would SAN boot from the VersaStack just as in the Unified VersaStack model, thereby allowing the server profiles to easily migrate between blades. The ESXi Datastores would be provisioned on VersaStack as a VMFS as opposed to NFS. Both architectures are designed to easily scale by adding additional compute or storage. Should this design model meet the customer's objectives, a block only model can help reduce Capex and Opex through the simplification of the design and reduced equipment requirements.

Figure 2 VersaStack FC Design

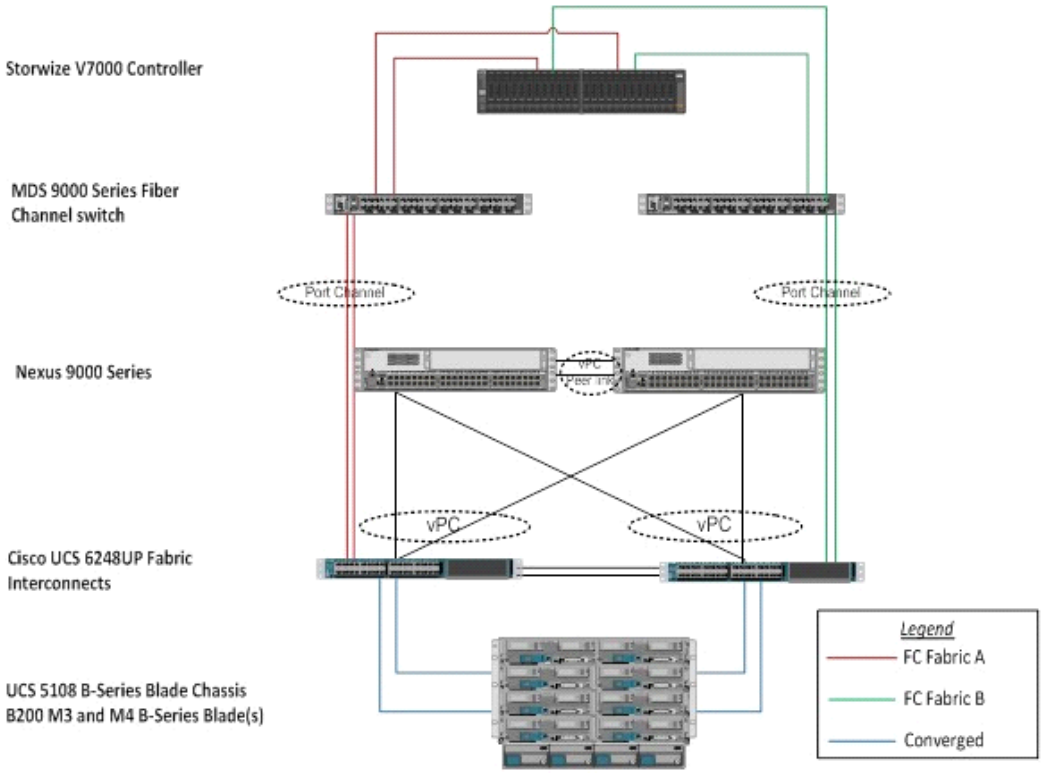

VersaStack Block Design Direct Attached

For smaller deployments VersaStack can be deployed in a direct attached configuration to eliminate the requirement for separate Fibre Channel switches. In this type of deployment the flexibility of the Fabric Interconnects is leveraged and they are changed to FC Switching Mode. The FC zoning is automated via UCS SAN Connectivity Polies and applied on the Fabric Interconnects to simplify deployment of new servers.

Figure 3 VersaStack Direct Attached Design

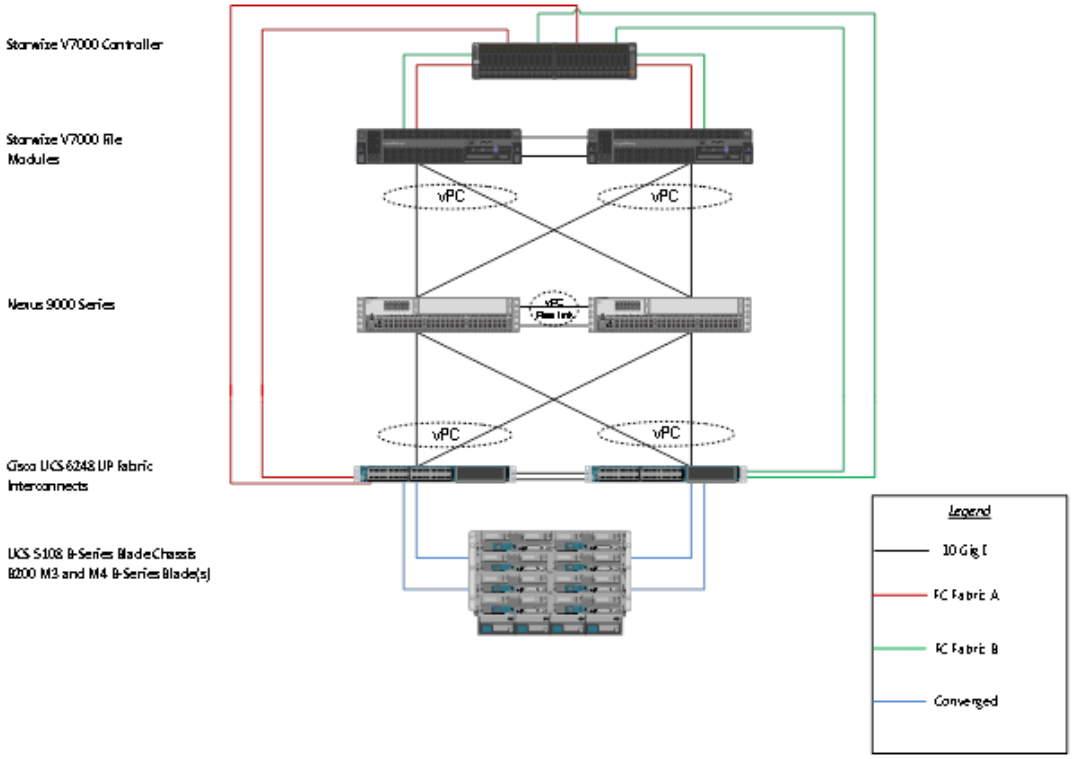

Integrated System Components

The following components are required to deploy this Cisco Nexus 9000 standalone design:

- Cisco Unified Compute System
- Cisco Nexus 9396 Series Switch
- Cisco MDS 9148 Series Switch
- IBM Storwize V7000 Unified with 8GB fiber channel adapters
- VMware vSphere

Cisco Unified Computing System

The Cisco Unified Computing System is a platform 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 compute platform in which all resources participate as 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 Cisco Unified Computing System consists of the following components:

- Cisco UCS Manager provides unified, embedded management of all software and hardware components in the Cisco Unified Computing System.

- Cisco UCS 6200 Series Fabric Interconnects is a family of line-rate, low-latency, lossless, 10-Gbps Ethernet and Fibre Channel over Ethernet interconnect switches providing the management and communication backbone for the Cisco Unified Computing System.
- Cisco UCS 5100 Series 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 increase performance, efficiency, versatility and productivity with these Intel based blade servers.
- Cisco UCS C-Series Rack Mount Server deliver unified computing in an industry-standard form factor to reduce total cost of ownership and increase agility.
- Cisco UCS Adapters wire-once architecture offers a range of options to converge the fabric, optimize virtualization and simplify management.

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

When leveraging ACI fabric mode, the Cisco Nexus 9500 and 9300 switches are deployed in a spine-leaf architecture. Although the reference architecture covered in this document does not leverage ACI, it lays the foundation for customer migration to ACI by introducing the Cisco Nexus 9000 switches. Application Centric Infrastructure (ACI) is a comprehensive architecture with centralized automation and policy-driven application profiles. ACI delivers software flexibility with the scalability of hardware performance.

For more information, refer to:

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

Cisco Nexus 1000v

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

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

  http://www.cisco.com/en/US/products/ps9902/index.html
  http://www.cisco.com/en/US/products/ps10785/index.html

Cisco MDS 9100 Multilayer Fabric Switch

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

Table 2 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 PortChannel bundle with multipath load balancing.
In-Service Software Upgrades	Reduce downtime for planned maintenance and software upgrades.

For more information, refer to:
http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/interoperability/matrix/intmatrx.pdf

IBM Storwize V7000 Unified

The storage controller leveraged for this validated design, IBM Storwize V7000 Unified, is engineered to satisfy the most demanding of workloads. IBM Storwize V7000 Unified 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.

IBM Storwize V7000 Unified has the following host interfaces:

- SAN-attached 8 Gbps Fibre Channel
- 1 Gbps iSCSI
- Optional 10 Gbps iSCSI/FCoE
- NAS-attached 1 Gbps and optional 10 Gbps

Each IBM Storwize V7000 node canister has up to 64GB internal cache to accelerate and optimize writes, and hardware acceleration to boost the performance of Real-time Compression.

IBM Storwize V7000 Unified can deploy the full range of Storwize software features, including:

- IBM Real-time compression
- IBM Easy Tier for automated storage tiering

- External storage virtualization and data migration
- IBM Active Cloud Engine for policy-based file tiering
- Synchronous data replication with Metro Mirror
- Asynchronous data replication with Global Mirror
- FlashCopy for near-instant data backups

IBM Storwize V7000 Unified 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.

IBM Easy Tier

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

For more information, refer to:

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

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, refer to:

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

Domain and Element Management

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

- Cisco UCS Manager
- Cisco UCS Central
- IBM Storwize V7000 Unified management GUI
- VMware vCenter™ Server

Cisco Unified Computing System Manager

Cisco UCS Manager provides unified, centralized, embedded management of Cisco Unified Computing System (UCS) software and hardware components across multiple chassis for up to 160 servers supporting thousands of virtual machines. Administrators use the software to manage the entire Cisco Unified Computing System as a single logical entity, called a domain, through an intuitive GUI, a command-line interface (CLI) or an XML API.

The Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series Fabric Interconnects using a clustered, active-standby configuration for high availability. The software, along with programmability of the Cisco VIC, provides statelessness to servers and gives administrators a single interface for performing server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, auditing, and statistics collection. Statelessness refers to the capability to separate server identity from the underlying hardware, thus setting the stage for assuming new server identity as needed. Cisco UCS servers provide for this functionality by storing server identifiers such as the MAC, UUID, WWN, firmware and BIOS versions in pools within Cisco UCS Manager which is resident in the redundant fabric interconnects. 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.

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

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

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

For more Cisco UCS Manager information, refer to:

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

Cisco UCS Central

For 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 domain. Cisco UCS Central Software manages global resources (including identifiers and policies) that can be consumed by instances of Cisco UCS Manager (Figure 4). It can delegate the application of policies (embodied in global service profiles) to individual domains, where Cisco UCS Manager puts the policies into effect. In its first release, Cisco UCS Central Software can support up to 10,000 servers in a single data center or distributed around the world in as many domains as are used for the servers.

Figure 4 Cisco UCS Management Stack

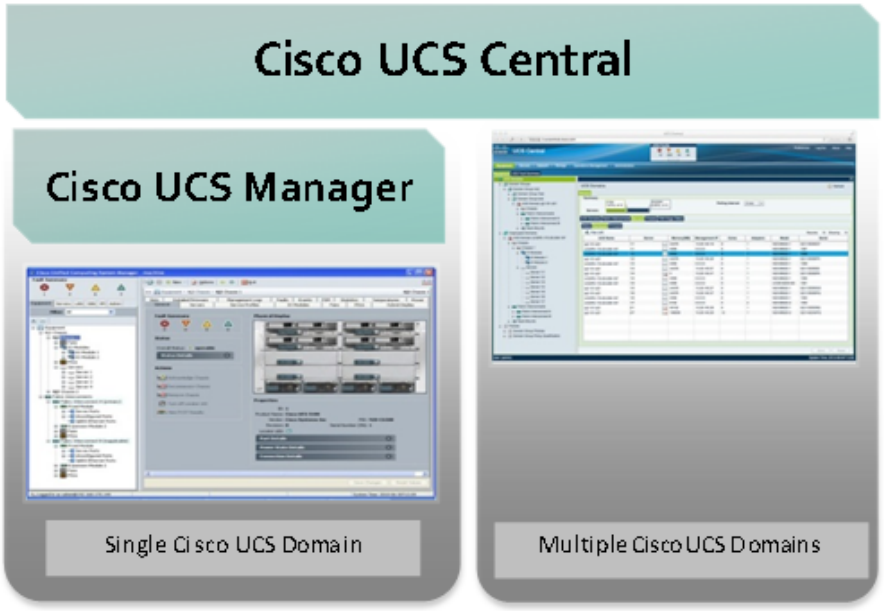

IBM Storwize V7000 Unified Easy-to-Use Management GUI

The IBM Storwize V7000 Unified built-in user interface (Figure 5) hides complexity and makes it possible for administrators to quickly and easily complete common block and file storage tasks from the same interface, such as creating and deploying volumes, file sets and host mappings. Users can also monitor performance in real-time (Figure 6).

The IBM Storwize V7000 Unified 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 5 IBM Storwize V7000 Unified Management GUI Example

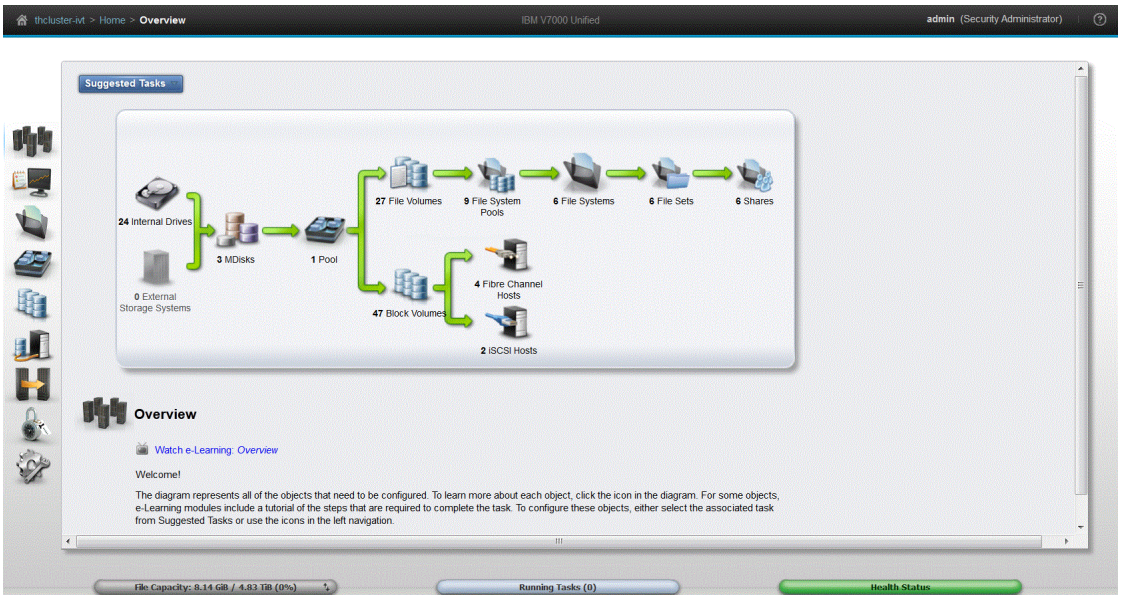

Figure 6 Real-time Performance Monitoring on the IBM Storwize V7000 Unified Management GUI

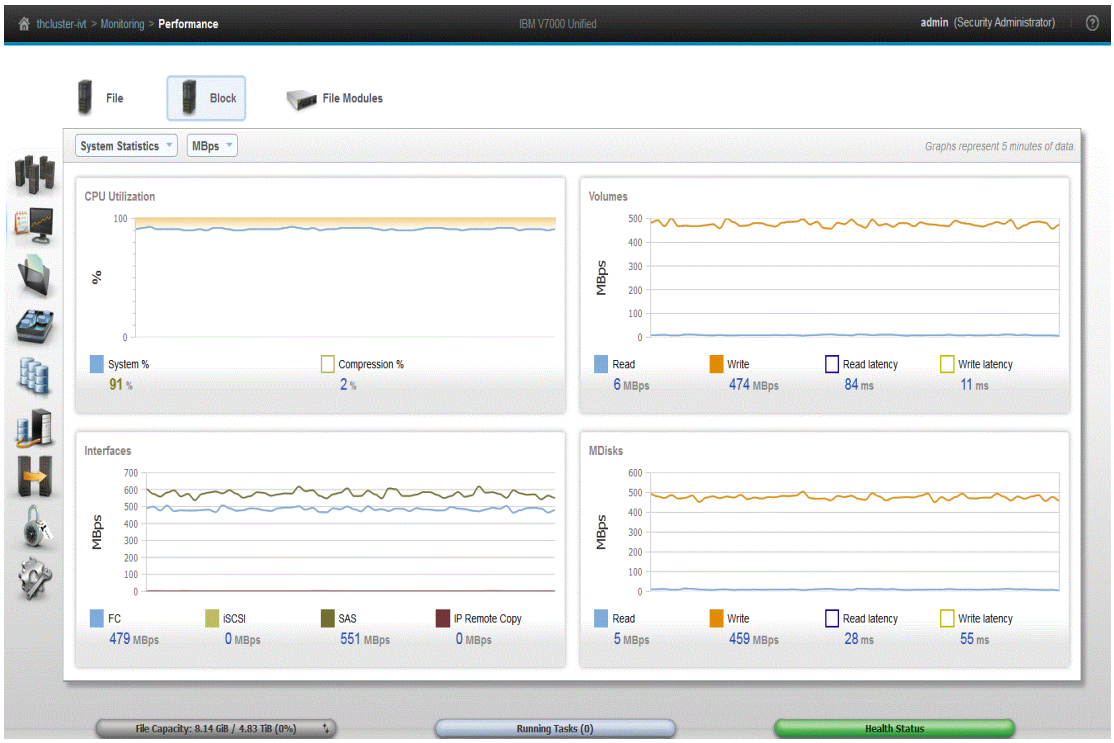

VMware vCenter Server

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, refer to: http://www.vmware.com/products/vcenter-server/overview.html

VMware vCenter Server Plug-Ins

vCenter Server plug-ins extend the capabilities of vCenter Server by providing more features and functionality. Some plug-ins are installed as part of the base vCenter Server product, for example, vCenter Hardware Status and vCenter Service Status, while other plug-ins are packaged separately from the base product and require separate installation. These are some of the plug-ins used during the VersaStack validation process.

Cisco Nexus 1000v vCenter Plug-In

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.

Starting with Cisco Nexus 1000V Release 4.2(1)SV2(1.1), a plug-in for the VMware vCenter Server, known as vCenter plug-in (VC plug-in) is supported on the Cisco Nexus 1000V virtual switch. It provides the server administrators a view of the virtual network and a visibility into the networking aspects of the Cisco Nexus 1000V virtual switch. The vCenter plug-in is supported on VMware vSphere Web Clients only. VMware vSphere Web Client enables you to connect to a VMware vCenter Server system to manage a Cisco Nexus 1000V through a browser. The vCenter plug-in is installed as a new tab in the Cisco Nexus 1000V as part of the user interface in vSphere Web Client.

The vCenter plug-in allows the administrators to view the configuration aspects of the VSM. With the vCenter plug-in, the server administrator can export necessary networking details for further analysis. The server administrator can thus monitor and manage networking resources effectively with the details provided by the vCenter plug-in.

IBM Tivoli Storage FlashCopy Manager

IBM Tivoli Storage FlashCopy Manager 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, refer to:
http://www-03.ibm.com/software/products/en/tivoli-storage-flashcopy-manager

Scalable Multi-Protocol VersaStack Design

Physical Build

Hardware and Software Revisions

Table 3describes 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 3 Validated Software Versions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M3,M4 UCS C-220 M3	2.2(3b)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, and UCS VIC 1240
	Cisco ESXi eNIC	2.1.2.42	
	Cisco Nexus 1000v	5.2(1)SV3(1.1)	
	Cisco ESXi fnic Driver	1.6.0.5	
Network	Cisco Nexus 9396 NX-OS	6.1(2)I3(1)	
	Cisco MDS 9148S	6.2(9)	
Storage	IBM Storwize V7000	Version 7.3.0.8	
	IBM Storwize V7000 Unified	Version 1.5.0.5-1	
Software	VMware vSphere ESXi	5.5u1	
	VMware vCenter	5.5u1	

Hardware and Software options for VersaStack

While VersaStack Deployment CVD's are configured with specific hardware and software, the components used to deploy a VersaStack can be customized to suit the specific needs of the environment as long as all the components and operating systems are on the HCL lists referenced in this document. VersaStack can be deployed with all advanced software features such as replication and storage virtualization on any component running support levels of code. It is recommended to use the software versions specified in the deployment CVD when possible. Other operating systems such as Linux and Windows are also supported on VersaStack either as a Hypervisor, guest OS within the hypervisor

environment, or directly installed onto bare metal servers. Other protocols, iSCSI for example, are supported as well. Note that basic networking components such as IP only switches are typically not on the IBM HCL. Please refer to Table 4 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 4 Examples of other VersaStack Hardware and Software Options

LayercomLayer	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	9372PX, 9396PX, 5548P, 5548UP, 5596UP	6.1(2)I3(1) or later
	Cisco Nexus 1000v	5.2(1)SV3(1.1) or later
	UCS VIC 12XX series ,13XX series	
Storage	Cisco MDS 9148S, 9250i	6.2(9) or later
	IBM Storwize V7000	Version 7.3.0.9 or later
	LFF Expansion (2076-12F) SFF Expansion (2076-24F)	
	IBM Storwize V7000 Unified File Modules (2073-720)	Version 1.5.0.5-1 or later
Software	VMware vSphere ESXi	5.5u1 or later
	VMware vCenter	5.5u1 or later
	Windows	2008R2, 2012R2

Logical Build

Figure 1 above illustrates the VersaStack with Cisco Nexus 9000 in standalone mode. The design is physically redundant at every level and connected for high-availability to provide a reliable platform... The solution also incorporates additional Cisco and IBM technologies and features for a differentiated and effective design. This section of the document discusses the logical configuration validated for VersaStack.

Cisco Unified Computing System

The VersaStack design supports both Cisco UCS B-Series and C-Series deployments. This section of the document discusses the integration of each deployment into VersaStack. 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 booted from SAN using Fibre Channel. The ESXi nodes consisted of Cisco UCS B200-M3 and B200-M4 series blades with Cisco 1240 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.

Cisco Unified Computing System I/O Component Selection

VersaStack allows customers to adjust the individual components of the system to meet their particular scale or performance requirements. Selection of I/O components, at the time of ordering, has a direct impact on scale and performance characteristics of the Cisco UCS components. Figure 7 illustrates the available backplane connections in the Cisco UCS 5100 series chassis. As shown, each of the two Fabric Extenders (I/O module) has four 10GBASE KR (802.3ap) standardized Ethernet backplane paths available for connection to the half-width blade slot. This means that each half-width slot has the potential to support up to 80Gb of aggregate traffic depending on selection of the following:

- Fabric Extender model (2204XP or 2208XP)
- Modular LAN on Motherboard (mLOM) card
- Mezzanine Slot card

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 5 Fabric Extender Model Comparison

	Network Facing Interface	Host Facing Interface
Cisco UCS 2204XP	4	16
Cisco UCS 2208XP	8	32

MLOM Virtual Interface Card (VIC)

VersaStack solution is typically validated using Cisco VIC 1240 or Cisco VIC 1280. Cisco VIC 1240 is a 4-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable modular LAN on motherboard (mLOM) designed exclusively for the M3 generation of Cisco UCS B-Series Blade Servers. When used in combination with an optional Port Expander, the Cisco UCS VIC 1240 capabilities can be expanded to eight ports of 10 Gigabit Ethernet with the use of Cisco UCS 2208 fabric extender.

Mezzanine Slot Card

A Cisco VIC 1280 is an eight-port 10 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE)-capable mezzanine card designed exclusively for Cisco UCS B-Series Blade Servers.

Traffic Aggregation

Selection of the FEX, VIC and Mezzanine cards plays a major role in determining the aggregate traffic throughput to and from a server. Figure 7 provides an overview of backplane connectivity for both the I/O Modules and Cisco VICs.

Figure 7 Cisco UCS B-Series M3 Blade Server Chassis Backplane Connections

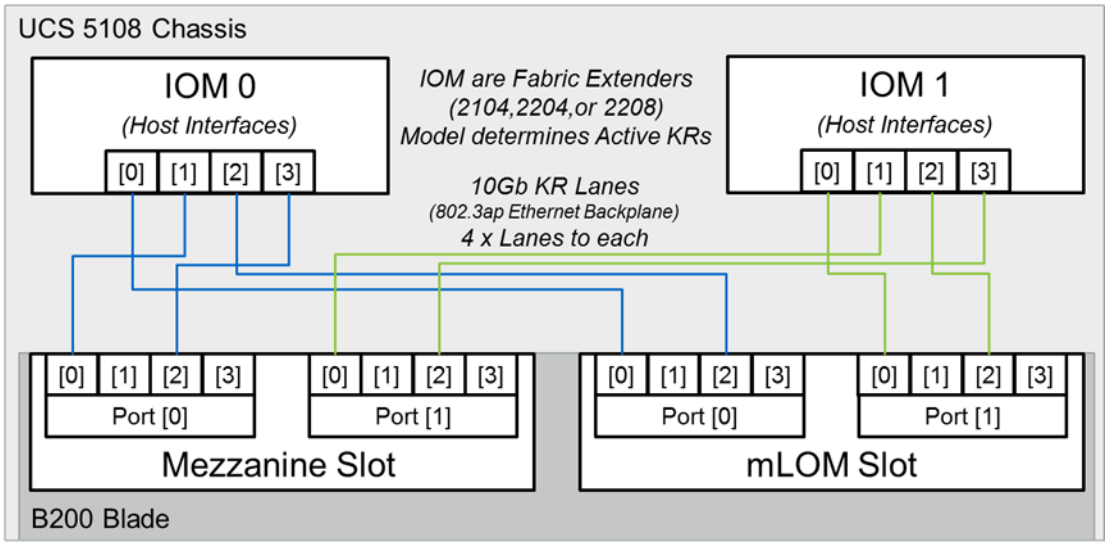

The number of KR lanes indicates the 10GbE paths available to the chassis and therefore blades. As shown in Figure 7, depending on the models of I/O modules and VICs, traffic aggregation differs. 2204XP enables 2 KR lanes per half-width blade slot while the 2208XP enables all four. Similarly, the number of KR lanes varies based on selection of VIC 1240, VIC 1240 with Port Expander and VIC 1280.

Validated I/O Component Configurations

Two of the most commonly validated I/O component configurations in VersaStack designs are as follows:

- Cisco UCS B200M3 with VIC 1240 and FEX 2204
- Cisco UCS B200M3 with VIC 1240 and port expander card

Figure 8 and Figure 9 illustrates the connectivity for these two configurations.

Figure 8 Validated Backplane Configuration—VIC 1240 with FEX 2204

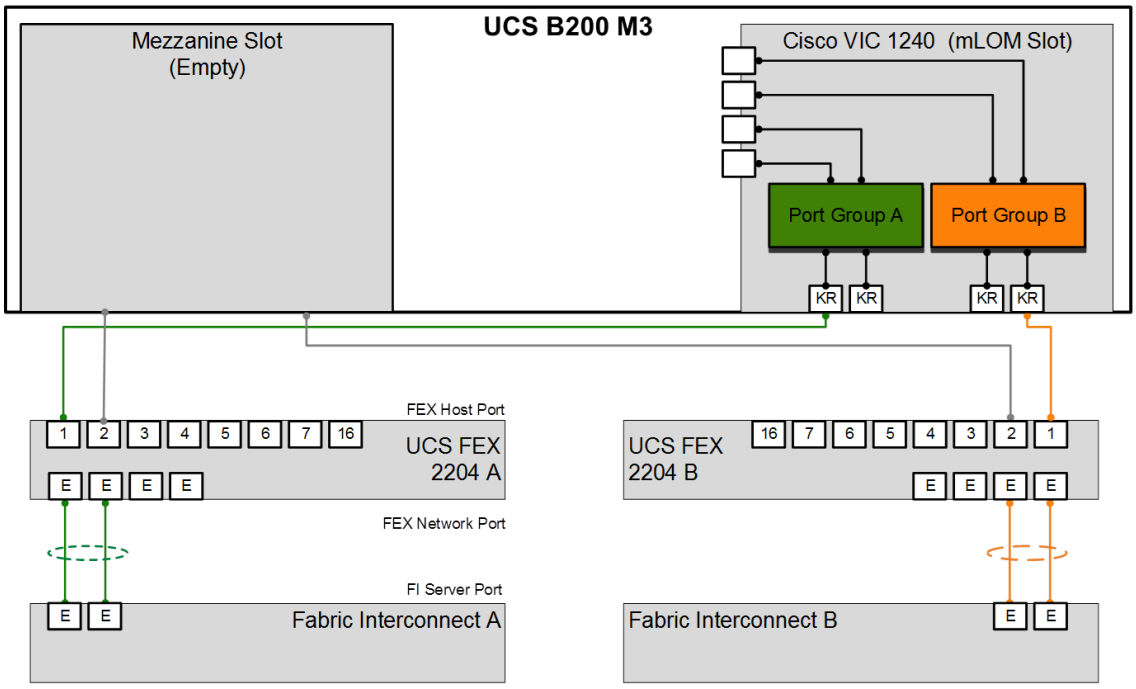

In Figure 8, the FEX 2204XP enables 2 KR lanes to the half-width blade while the global discovery policy dictates the formation of a fabric port channel. This results in 20GbE connection to the blade server.

Figure 9 Validated Backplane Configuration—VIC 1240 with VIC Port Expander Card

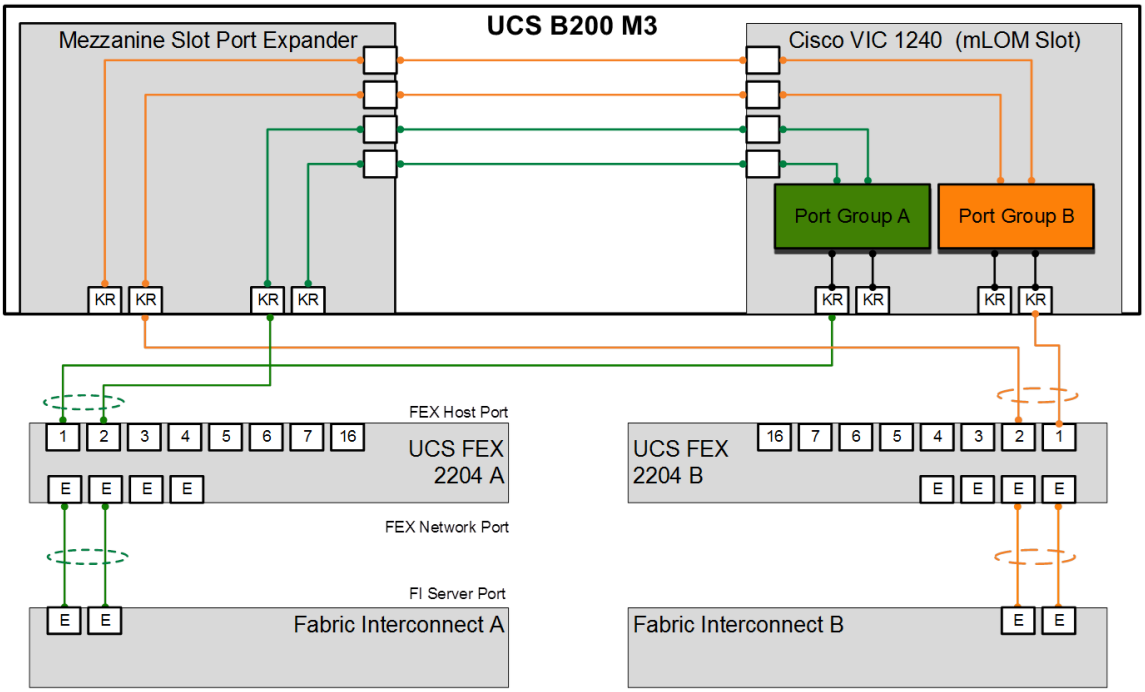

In Figure 9, the FEX 2204XP enables two more 10 Gbps KR lanes to the half-width blade while the global chassis discovery policy dictates the formation of a fabric port channel. This configuration results in two 20 Gbps port channel connections to the blade server.

Cisco Unified Computing System Chassis/FEX Discovery Policy

Cisco Unified Computing System can be configured to discover a chassis using Discrete Mode or the Port-Channel mode (Figure 10). 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 is less disruptive to the fabric and hence recommended in the VersaStack designs.

Figure 10 Chassis Discovery Policy Discrete—Mode vs. Port Channel Mode

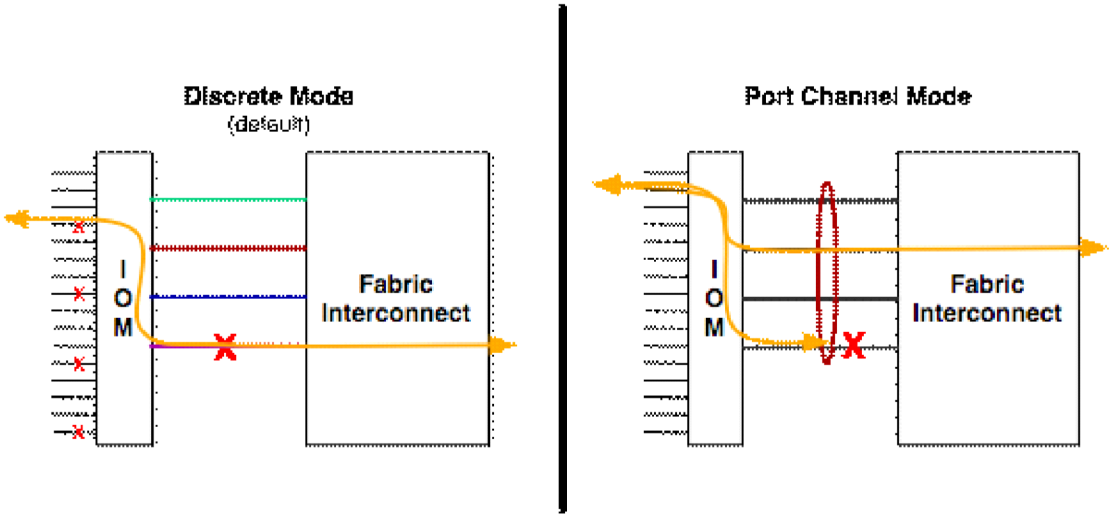

Cisco Unified Computing System—QoS and Jumbo Frames

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 Unified Computing System 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.

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

Cisco Unified Computing System—C-Series Server Design

Cisco UCS Manager 2.2 provides two connectivity modes for Cisco UCS C-Series Rack-Mount Server management:

- Dual-wire Management (Shared LOM): This management mode is supported in the Cisco UCS Manager releases earlier than 2.2. In this mode, shared LAN on Motherboard (LOM) ports on the rack server are used exclusively for carrying management traffic. A separate cable connected to one of the ports on the PCIe card carries the data traffic.
- Single-wire Management (Sideband): Cisco UCS Manager release version 2.2 introduces an additional rack server management mode using Network Controller Sideband Interface (NC-SI). Cisco UCS VIC1225 Virtual Interface Card (VIC) uses the NC-SI, which can carry both data traffic and management traffic on the same cable. Single-wire management allows for denser server to FEX deployments.

Note The VersaStack Nexus 9000 design is capable of supporting both single and dual wire management but the validation was limited to single-wire management designs.

Figure 11 illustrates the connectivity of the Cisco UCS C-Series server into the Cisco UCS domain.

From a functional perspective the 1 RU Cisco Nexus FEX 2232PP replaces the Cisco UCS 2204 or 2208 IOM (located within the Cisco UCS 5108 blade chassis). Each 10GbE VIC port connects to Fabric A or B through the FEX. The FEX and Fabric Interconnects form port channels automatically based on the chassis discovery policy providing a link resiliency to the Cisco UCS C-Series server. This is identical to the behavior of the IOM to Fabric Interconnect connectivity. From a logical perspective the virtual circuits formed within the Cisco UCS domain are consistent between Cisco UCS B and C Series deployment models and the virtual constructs formed at the vSphere or Cisco Nexus 1000v layer are unaware of the platform in use.

Figure 11 Cisco UCS C-Series with VIC 1225

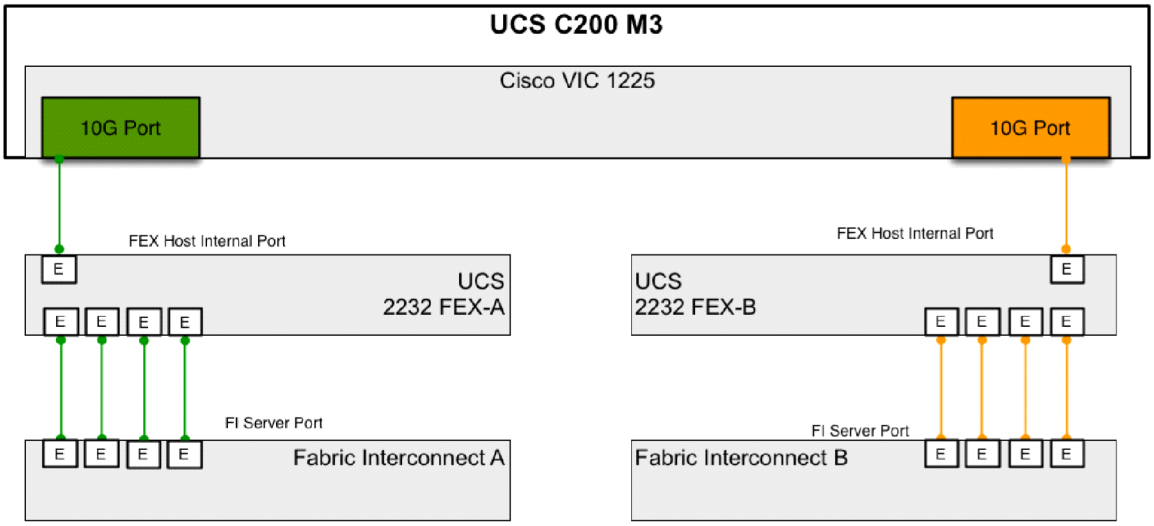

Cisco Nexus 9000

Cisco Nexus 9000 provides Ethernet switching fabric for communications between the Cisco UCS domain, the IBM Storwize V7000 Unified storage system and the enterprise network. In the VersaStack design, Cisco UCS Fabric Interconnects, IBM Storwize V7000 file modules and are connected to the Cisco Nexus 9000 switches using virtual PortChannels (vPC).

Virtual Port Channel (vPC)

A virtual PortChannel (vPC) allows links that are physically connected to two different Cisco Nexus 9000 Series devices to appear as a single PortChannel. In a switching environment, vPC provides the following benefits:

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

Figure 12 Cisco Nexus 9000 Connections to IBM Storwize V7000 Unified

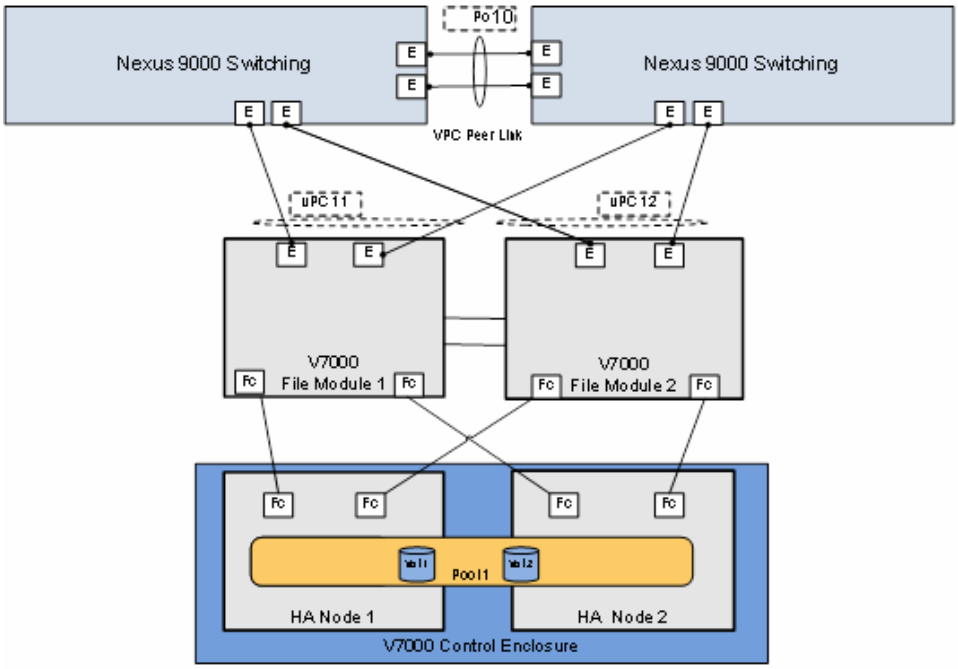

Figure 12 illustrates the connections between Cisco Nexus 9000 and file module for NFS access to the IBM Storwize V7000 file modules. V7000 file modules connected through a vPC require a "peer link" which is documented as port channel 10 in this diagram. In addition to the vPC peer-link, vPC peer keepalive 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 out-of-band management network. This link is not shown in Figure 11.

Cisco Nexus 9000 Best Practices

Cisco Nexus 9000 related best practices used in the validation of the VersaStack architecture are summarized below:

- Cisco Nexus 9000 features enabled
  - Link Aggregation Control Protocol (LACP part of 802.3ad)

  - Cisco Virtual Port Channeling (vPC) for link and device resiliency
  - Enable Cisco Discovery Protocol (CDP) for infrastructure visibility and troubleshooting
- vPC considerations
  - Define a unique domain ID
  - Set the priority of the intended vPC primary switch lower than the secondary (default priority is 32768)
  - Establish peer keepalive connectivity. It is recommended to use the out-of-band management network (mgmt0) or a dedicated switched virtual interface (SVI)
  - Enable vPC auto-recovery feature
  - Enable peer-gateway. Peer-gateway allows a vPC switch to act as the active gateway for packets that are addressed to the router MAC address of the vPC peer allowing vPC peers to forward traffic
  - Enable IP arp synchronization to optimize convergence across the vPC peer link.

Note Cisco Fabric Services over Ethernet (CFSoE) is responsible for synchronization of configuration, Spanning Tree, MAC and VLAN information, which removes the requirement for explicit configuration. The service is enabled by default.

  - A minimum of two 10 Gigabit Ethernet connections are required for vPC
  - All port channels should be configured in LACP active mode
- Spanning tree considerations
  - The spanning tree priority was not modified. Peer-switch (part of vPC configuration) is enabled which allows both switches to act as root for the VLANs
  - Loopguard is disabled by default
  - BPDU guard and filtering are enabled by default
  - Bridge assurance is only enabled on the vPC Peer Link.
  - Ports facing IBM Storwize Unified and Cisco Unified Computing System are defined as "edge" trunk ports For configuration details, refer to the Cisco Nexus 9000 series switches configuration guides:

    http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-installation-and-configuration-guides-list.html

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 13 illustrates the Cisco Nexus 1000v architecture.

Figure 13 Cisco Nexus 1000v Architecture

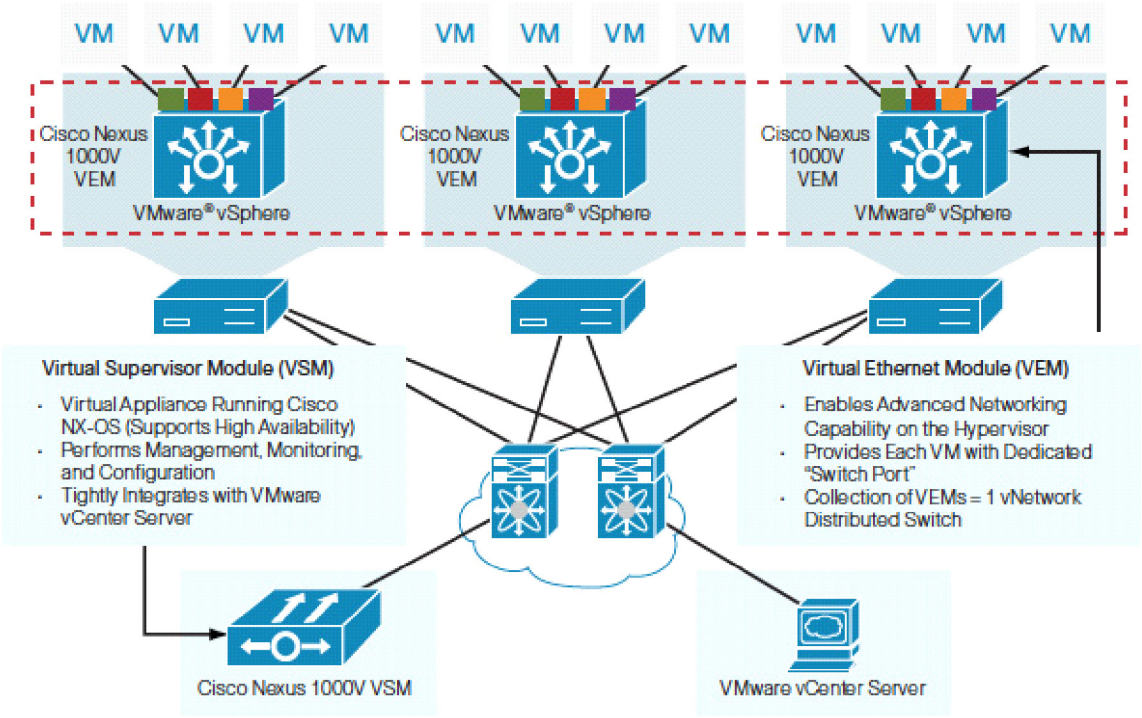

The VEM takes configuration information from the VSM and performs Layer 2 switching and advanced networking functions, as follows:

- 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

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

Figure 14 illustrates a single ESXi node with a VEM registered to the Cisco Nexus 1000v VSM. The ESXi vmnics are presented as Ethernet interfaces in the Cisco Nexus 1000v.

Figure 14 Cisco Nexus 1000v VEM in an ESXi Environment

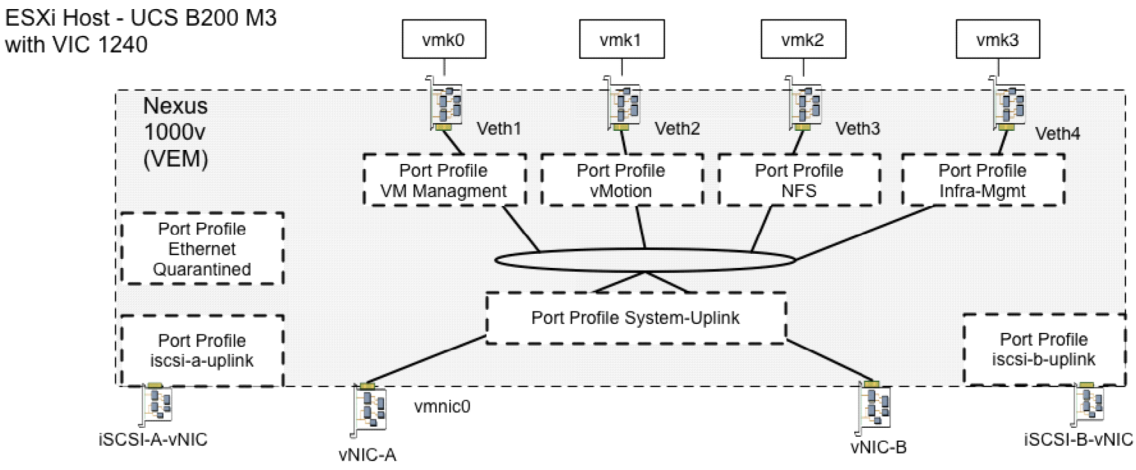

Cisco Nexus 1000v Interfaces and Port Profiles

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 14).

Note 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), each with its own port profile:

- vmk0 - ESXi management
- vmk1 - vMotion interface
- vmk2 - NFS interface
- vmk3 - Infrastructure management

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 about "Best Practices in Deploying Cisco Nexus 1000V Series Switches on Cisco UCS B and C Series Cisco UCS Manager Servers" refer to:

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

IBM Storwize V7000 Unified High Availability

Each Storwize V7000 Unified system can be used to deploy a high availability storage solution with no single point of failure.(Figure 15)

Each IBM Storwize V7000 control enclosure within a system contains dual redundant control node canisters that manage I/O on an active-active basis. If a node canister fails, the other node canister in the control enclosure seamlessly assumes I/O responsibilities of the failed node canister. Node canisters in a control enclosure can communicate both through an internal PCIe link, and over the FC SAN, for increased redundancy and performance.

Each Storwize V7000 node canister has up to eight fibre-channel ports that can attach to multiple SAN fabrics. For high availability, attach node canisters to at least two fabrics.

Figure 15 IBM Storwize V7000 2-node Block HA Cabling

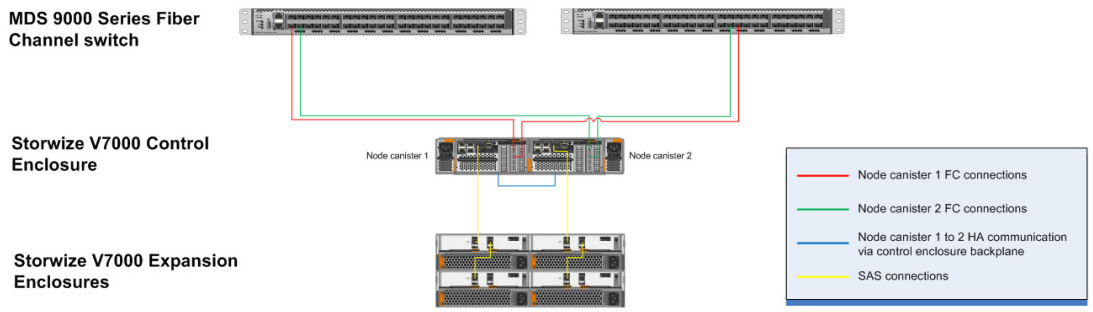

IBM Storwize V7000 supports fully redundant connections for communication between control enclosures, external storage, and host systems. If a SAN fabric fault disrupts communication or I/O operations, the system recovers and retries the operation through an alternative communication path. Host systems should be configured to use multipathing, so that if a SAN fabric fault or node canister failure occurs, the host can retry I/O operations.

IBM Storwize V7000 file modules are connected to a V7000 control enclosure using direct redundant fiber channel connections. The V7000 file modules and V7000 control enclosure also need to be able to communicate through ethernet for management.

The V7000 file modules are connected to each other with redundant 1GbE cables, with a further six 1GbE ports and two 10GbE ports for public file access. One of the 1GbE ports is used for system management, and one of the 10GbE ports is optionally available for system management (Figure 16).

Figure 16 IBM Storwize V7000 Unified 2-Node Block and 10GbE File Traffic Cabling

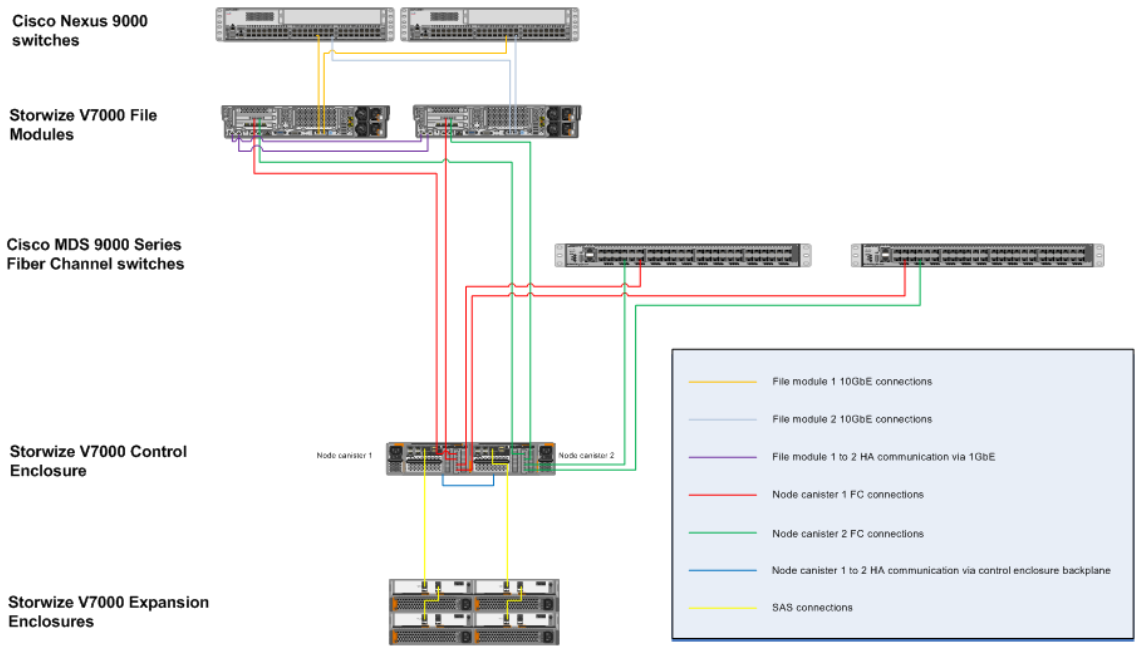

IBM Storwize V7000 supports up to four control enclosures, for a maximum 8 node canisters. Systems containing multiple control enclosures must connect node canisters through FC SAN to allow communication between them. A single pair of V7000 file modules is needed to provide file access to a fully-scaled V7000 system. I/O requests to control enclosures that the file modules are not direct-attached are routed over the FC SAN (Figure 17 and Figure 18).

Figure 17 Storwize V7000 Fully-scaled HA Communication

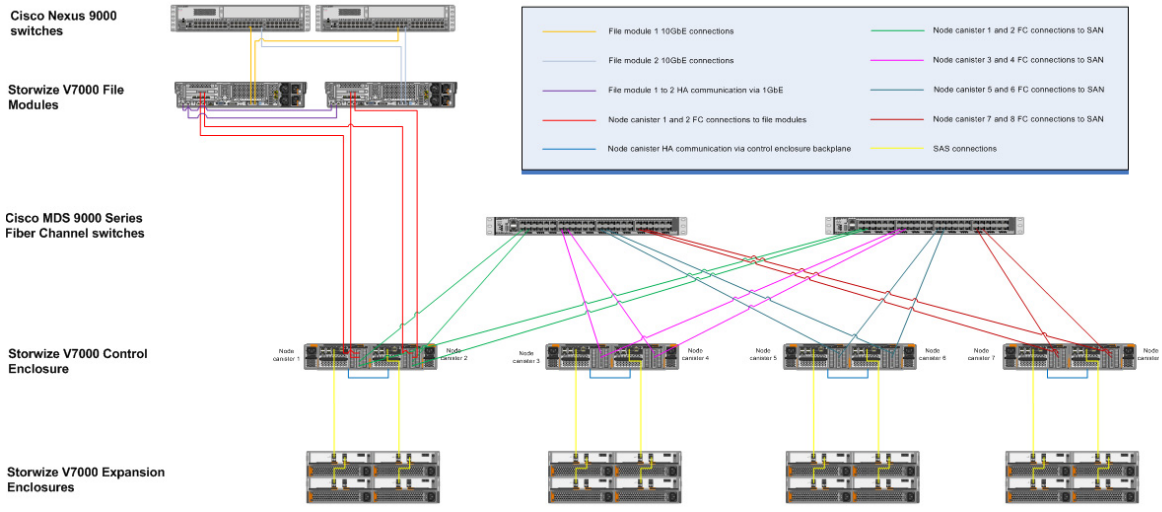

Figure 18 Fibre Channel Connectivity Details

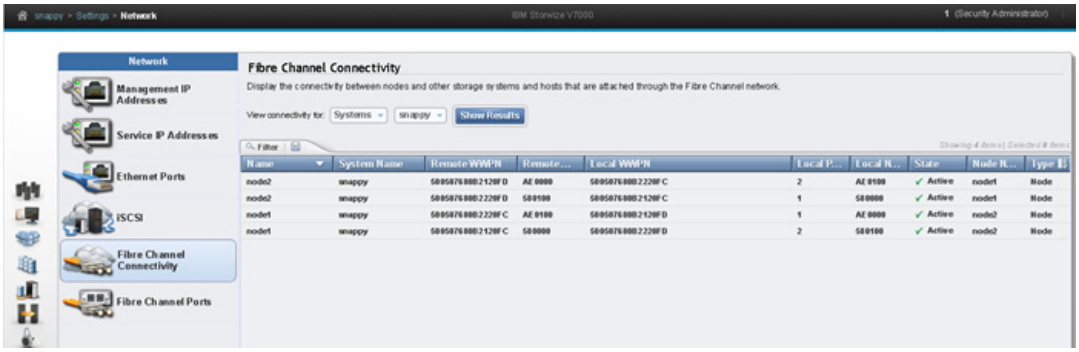

VMware vCenter and vSphere

VMware vSphere 5.5 Update 1 provides a platform for virtualization that includes multiple components and features. In this validation effort, the following key components and features were utilized:

- VMware ESXi
- VMware vCenter Server
- VMware vSphere SDKs
- VMware vSphere Virtual Machine File System (VMFS)
- VMware vSphere High Availability (HA)
- VMware vSphere Distributed Resource Scheduler (DRS)

Conclusion

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.

References

- Cisco Unified Computing System
- Cisco UCS 6200 Series Fabric Interconnects
- Cisco UCS 5100 Series Blade Server Chassis
- Cisco UCS B-Series Blade Servers
- Cisco UCS Adapters
- Cisco UCS Manager
- Cisco Nexus 9000 Series Switches

- Cisco Nexus 1000v
- Cisco Prime Data Center Manager
- IBM Storwize V7000 Unified
- VMware vCenter Server
- VMware vSphere

Interoperability Matrixes

- VMware and Cisco Unified Computing System
- IBM System Storage Interoperation Center
- Design Principles for Security, by Terry V. Benzel, Cynthia E. Irvine, Timothy E. Levin, Ganesha Bhaskara, Thuy D. Nguyen, and Paul C. Clark