




FlexPod Datacenter with VMware vSphere 5.1U1 and Cisco Nexus 9000 Series Switches Design Guide

Last Updated: April 4, 2014



Building Architectures to Solve Business Problems



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>.

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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

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

© 2014 Cisco Systems, Inc. All rights reserved

About the Authors

Chris O'Brien, Technical Marketing Manager, Server Access Virtualization Business Unit, Cisco Systems

Chris O'Brien is currently focused on developing infrastructure best practices and solutions that are designed, tested, and documented to facilitate and improve customer deployments. Previously, O'Brien was an application developer and has worked in the IT industry for more than 15 years.

Haseeb Niazi, Technical Marketing Engineer, Server Access Virtualization Business Unit, Cisco Systems

Haseeb has over 14 years of experience at Cisco dealing in Data Center, Security, and WAN Optimization related technologies. As a member of various solution teams and advanced services, Haseeb has helped many enterprise and service provider customers evaluate and deploy a wide range of Cisco solutions. Haseeb holds a master's degree in Computer Engineering from the University of Southern California.

Chris Reno, Reference Architect, Infrastructure and Cloud Engineering, NetApp

Chris Reno is a reference architect in the NetApp Infrastructure and Cloud Enablement group and is focused on creating, validating, supporting, and evangelizing solutions based on NetApp products. Before being employed in his current role, he worked with NetApp product engineers designing and developing innovative ways to perform Q&A for NetApp products, including enablement of a large grid infrastructure using physical and virtualized compute resources. In these roles, Chris gained expertise in stateless computing, netboot architectures, and virtualization.

FlexPod Datacenter with VMware vSphere 5.1U1 and Cisco Nexus 9000 Series Switches 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 NetApp® VMware vSphere 5.1 Update 1 on FlexPod® solution with the Cisco Nexus 9000. A FlexPod solution is a validated approach for deploying Cisco and NetApp 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.

Changes in FlexPod

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

- Support for the Cisco's latest switch offering, the Nexus 9000, which in future will serve as the foundation for Cisco's Application Centric Infrastructure (ACI)
- Support for the latest release of NetApp Data ONTAP® 8.2
- An IP-based storage design supporting both NAS datastores and iSCSI based SAN LUNs.

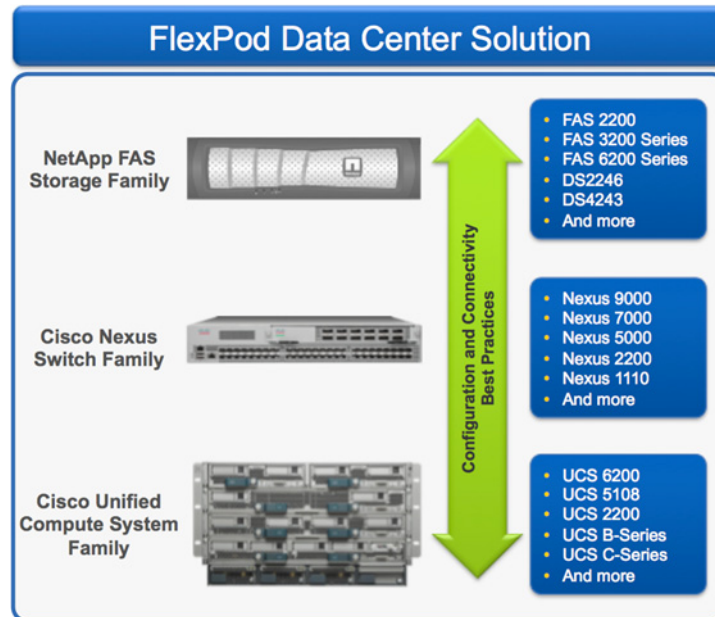
Introduction

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. Enterprise customers are moving away from silos of IT operation toward more cost-effective virtualized environments, leading eventually to cloud computing to increase agility and reduce costs. To accelerate this process and simplify the evolution to a shared cloud infrastructure, Cisco and NetApp have developed and continue to enhance a solution called VMware vSphere® on FlexPod.

FlexPod is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (UCS), the Cisco Nexus® family of switches, and NetApp fabric-attached storage (FAS) or V-Series systems. FlexPod is a suitable platform for running a variety of virtualization hypervisors as well as bare metal operating systems and enterprise workloads. FlexPod delivers a baseline configuration and also has the flexibility to be sized and optimized to accommodate many

different use cases and requirements. This document describes VMware vSphere 5.1 Update 1 built on the FlexPod model from Cisco and NetApp and discusses design choices and deployment of best practices using this shared infrastructure platform.

Figure 1 FlexPod Component Families



Problem Statement

As customers transition toward shared infrastructure or cloud computing they face a number of challenges such as how to begin the transition, how to calculate return on investment (ROI), how to manage the infrastructure and how to plan for growth.

The FlexPod architecture is designed to help with proven guidance and measurable value. By introducing standardization, FlexPod helps customers mitigate the risk and uncertainty involved in planning, designing, and implementing a new data center infrastructure. The result is a more predictive and adaptable architecture capable of meeting and exceeding customers' IT demands.

FlexPod Program Benefits

Cisco and NetApp have thoroughly validated and verified the FlexPod 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 includes, 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 FlexPod configuration)

- Frequently asked questions (FAQs)
- Cisco Validated Designs (CVDs) and NetApp Verified Architectures (NVAs) focused on a variety of use cases

Cisco and NetApp have also built a robust and experienced support team focused on FlexPod solutions, from customer account and technical sales representatives to professional services and technical support engineers. The support alliance provided by NetApp 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.

FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment. The following IT initiatives are addressed by the FlexPod solution.

Integrated System

FlexPod 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

FlexPod is a highly available and scalable infrastructure that IT can evolve over time to support multiple physical and virtual application workloads. FlexPod 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

FlexPod 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 FlexPod while maintaining consistent application and security policy enforcement across the stack even with workload mobility.

FlexPod provides a uniform approach to IT architecture, offering a well-characterized and documented shared pool of resources for application workloads. FlexPod 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

FlexPod

System Overview

FlexPod is a best practice data center architecture that includes three components:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus switches
- NetApp fabric-attached storage (FAS) systems

These components are connected and configured according to best practices of both Cisco and NetApp and provide the ideal platform for running a variety of enterprise workloads with confidence. As previously noted, the reference architecture covered in this document leverages the Nexus 9000 for the switching element. FlexPod 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 FlexPod stacks).

One of the key benefits of FlexPod is the ability to maintain consistency at scale. Each of the component families shown in [Figure 1](#) (Cisco UCS, Cisco Nexus, and NetApp FAS) 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 FlexPod.

Design Principles

FlexPod addresses four primary design principles:

- Application availability. Makes sure that services are accessible and ready to use.
- Scalability. Addresses increasing demands with appropriate resources.
- Flexibility. Provides new services or recovers resources without requiring infrastructure modifications.
- Manageability. Facilitates efficient infrastructure operations through open standards and APIs.

**Note**

Performance and comprehensive security are key design criteria that are not directly addressed in this project but have been addressed in other collateral, benchmarking, and solution testing efforts. This design guide validates the functionality and basic security elements.

FlexPod and Data ONTAP Modes of Operation

NetApp Data ONTAP 8.0 and later support two modes of operation: clustered Data ONTAP and Data ONTAP operating in 7-Mode. The ability to use either mode in a FlexPod configuration has been validated.

It is a fundamental design decision to leverage clustered Data ONTAP or 7-Mode, because these cannot be run simultaneously on the same controller, and the choice will influence hardware requirements, the logical construction of the FlexPod stack, and ultimately the operational practices of the enterprise. Organizations having the following requirements should consider adopting clustered Data ONTAP:

- Large to midsize enterprises that are seeking scalable, shared IT solutions for nondisruptive operations
- New installations

- Existing clustered Data ONTAP 8.x and Data ONTAP GX organizations that are looking to upgrade
- Organizations deploying an enterprise content repository

Organizations with the following characteristics or needs might want to use the Data ONTAP 7-Mode design:

- Existing Data ONTAP 7G and Data ONTAP 8.x 7-Mode customers who are looking to upgrade
- Midsize enterprises; customers who are primarily interested in the FAS2000 series
- Customers who absolutely require synchronous SnapMirror®, MetroCluster™, SnapLock® software, or Data ONTAP Edge


Note

It is advisable to seek counsel from experts. Contact your NetApp account team or partner for further guidance.

This solution leverages clustered Data ONTAP, which enables many advances in the storage feature set, especially for scalability, flexibility, and resiliency.

FlexPod and Cisco Nexus 9000 Modes of Operation

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

Nexus 9000 stand-alone mode FlexPod design consists of a single pair of Nexus 9000 top of rack switches. The integration of ACI fabric in the future will introduce Nexus 9500 and 9300 based 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 leveraging the 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

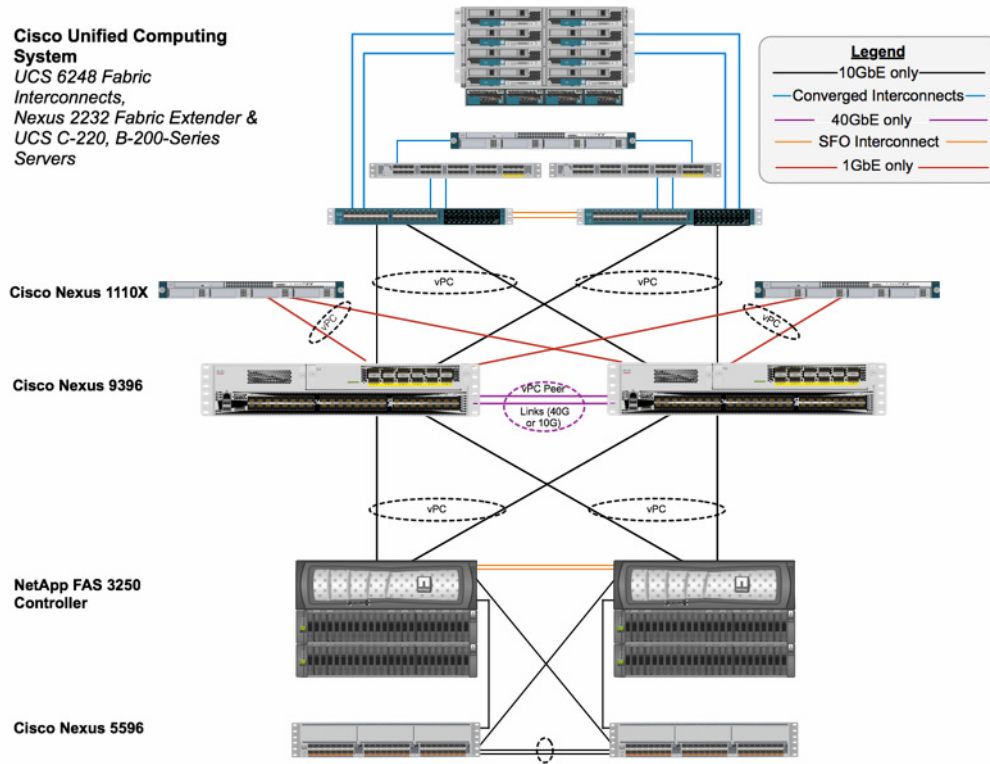
The future of networking with ACI is about providing a network that is deployed, monitored, and managed in a fashion that supports DevOps and rapid application change. ACI does so through the reduction of complexity and a common policy framework that can automate provisioning and managing of resources.

At the time of Cisco ACI launch later this year, the ACI solution will comprise of the [Cisco Nexus 9000 Series Switches](#) running in ACI fabric mode, Cisco Application Policy Infrastructure Controller (APIC), and [Cisco Application Virtual Switch \(AVS\)](#).

FlexPod: Nexus 9000 Standalone Design

[Figure 2](#) details the FlexPod: Cisco Nexus 9000 standalone design with clustered Data ONTAP. As the illustration shows, the design is fully redundant in the compute, network, and storage layers. There is no single point of failure from a device or traffic path perspective.

Figure 2 FlexPod: Cisco Nexus 9000 Standalone Design with Clustered Data ONTAP



The FlexPod: Cisco Nexus 9000 standalone design is an end-to-end IP-Based storage solution that supports SAN access using iSCSI. The solution provides a 10GbE-enabled, 40G capable, fabric defined by Ethernet uplinks from the Cisco UCS Fabric Interconnects and NetApp storage devices connected to the Cisco Nexus switches, as The Nexus 9000 standalone design does not employ a dedicated SAN switching environment and requires no direct Fibre Channel connectivity as iSCSI is the SAN protocol leverage.

As illustrated, link aggregation technologies play an important role, providing improved aggregate bandwidth and link resiliency across the solution stack. The NetApp storage controllers, 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 NetApp FAS controllers 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.

This dedicated uplink design leverages IP-based storage-capable NetApp FAS controllers. From a storage traffic perspective, both standard LACP and the Cisco vPC link aggregation technologies play an important role in the FlexPod distinct uplink design. Figure 2 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 NetApp storage controllers' 10GbE provide a robust connection between host and storage.

Figure 2 shows the initial storage configuration of this solution as a two-node HA pair with clustered Data ONTAP. A storage configuration comprise an HA pair, which consists of like storage nodes such as FAS22xx, 32xx, or 62xx series and storage shelves housing disks. Scalability is achieved by adding storage capacity (disk/shelves) to an existing HA pair, or by adding HA pairs into the cluster or storage domain.



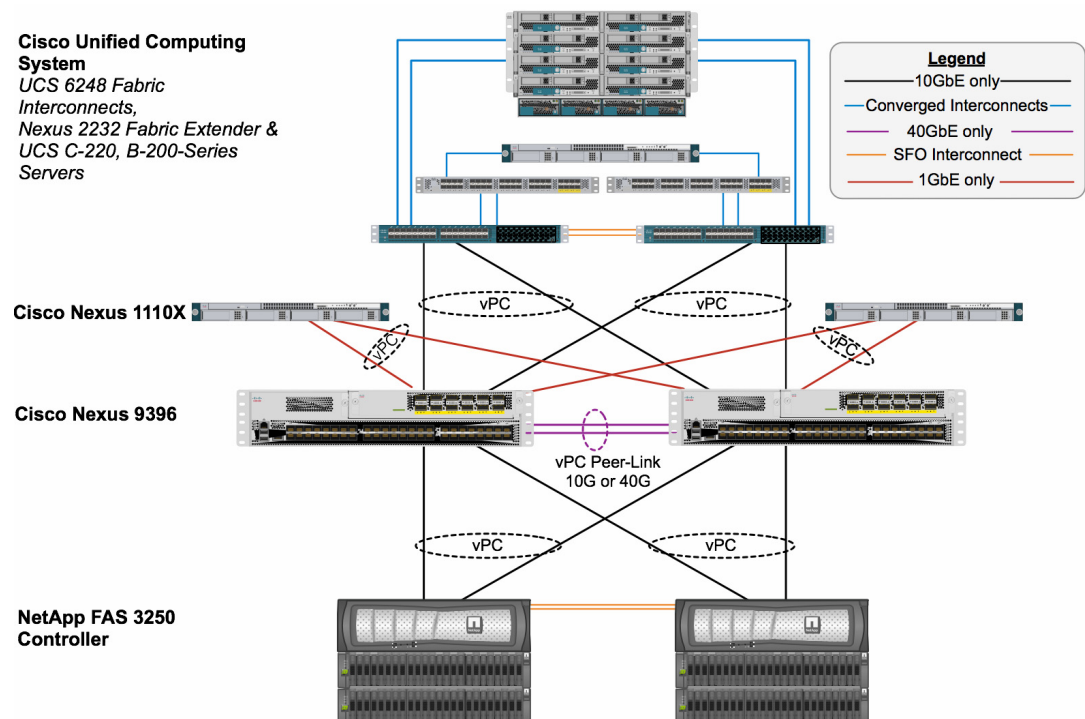
Note

For SAN environments, the NetApp clustered Data ONTAP offering allows up to four HA pairs or 8 nodes and for NAS environments, 12 HA pairs or 24 nodes to form a logical entity.

The HA interconnect allows each HA node pair to assume control of its partner's storage (disk/shelves) directly. The local physical high-availability storage failover capability does not extend beyond the HA pair. Furthermore, a cluster of nodes does not have to include similar hardware. Rather, individual nodes in an HA pair are configured alike, allowing customers to scale as needed, as they bring additional HA pairs into the larger cluster.

Figure 3 details the FlexPod: Cisco Nexus 9000 standalone design with Data ONTAP operating in 7-Mode. As depicted, the FAS devices are configured in an HA pair and scalability is achieved through the addition of storage capacity (disk/shelves), as well as through additional controllers such as FAS2200, 3200, or 6200 series. The controllers are only deployed in HA pairs, meaning more HA pairs can be added for scalability, but each pair is managed separately.

Figure 3 FlexPod: Cisco Nexus 9000 Standalone Design with Data ONTAP in 7-Mode



For more information about the virtual design of the environment consisting of VMware vSphere, Cisco Nexus 1000v virtual distributed switching, and NetApp storage controllers, refer to the section [Logical Build](#).

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
- NetApp Unified Storage
- VMware vSphere

Cisco Unified Computing System

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

The Cisco Unified Computing System consists of the following components:

- Cisco UCS Manager (<http://www.cisco.com/en/US/products/ps10281/index.html>) provides unified, embedded management of all software and hardware components in the Cisco UCS.
- Cisco UCS 6200 Series Fabric Interconnects (<http://www.cisco.com/en/US/products/ps11544/index.html>) 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 Unified Computing System.
- Cisco UCS 5100 Series Blade Server Chassis (<http://www.cisco.com/en/US/products/ps10279/index.html>) 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 (<http://www.cisco.com/en/US/partner/products/ps10280/index.html>) increase performance, efficiency, versatility and productivity with these Intel based blade servers.
- Cisco UCS C-Series Rack Mount Server (<http://www.cisco.com/en/US/products/ps10493/index.html>) deliver unified computing in an industry-standard form factor to reduce total cost of ownership and increase agility.
- Cisco UCS Adapters (http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html) 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 and fixed 10/40/100 Gigabit Ethernet switch configurations with scalability up to 30 Tbps of nonblocking performance with less than five-microsecond latency, 1152 10 Gbps or 288 40 Gbps nonblocking Layer 2 and Layer 3 Ethernet ports and wire speed VXLAN gateway, bridging, and routing support.

For more information, see:

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

Cisco Nexus 2232PP 10GE Fabric Extender

The Cisco Nexus 2232PP 10G provides 32 10 Gb Ethernet and Fibre Channel Over Ethernet (FCoE) Small Form-Factor Pluggable Plus (SFP+) server ports and eight 10 Gb Ethernet and FCoE SFP+ uplink ports in a compact 1 rack unit (1RU) form factor.

When a C-Series Rack-Mount Server is integrated with Cisco UCS Manager, through the Cisco Nexus 2232 platform, the server is managed using the Cisco UCS Manager GUI or Cisco UCS Manager CLI. The Cisco Nexus 2232 provides data and control traffic support for the integrated Cisco UCS C-Series server.

Cisco Nexus 1000v

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

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

For more information, see:

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

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

NetApp FAS and Data ONTAP

NetApp solutions offer increased availability while consuming fewer IT resources. A NetApp solution includes hardware in the form of controllers and disk storage and the NetApp Data ONTAP operating system, the #1 storage OS. The storage efficiency built into Data ONTAP provides substantial space savings, allowing more data to be stored at a lower cost.

NetApp offers the NetApp Unified Storage Architecture. The term "unified" refers to a family of storage systems that simultaneously support storage area network (SAN), network-attached storage (NAS), and iSCSI across many operating environments such as VMware, Windows[®], and UNIX[®]. This single architecture provides access to data by using industry-standard protocols, including NFS, CIFS, iSCSI, FCP, SCSI, FTP, and HTTP. Connectivity options include standard Ethernet (10/100/1000, or 10GbE) and Fibre Channel (1, 2, 4, or 8Gb/sec). In addition, all systems can be configured with high-performance solid state drives (SSDs) or serial ATA (SAS) disks for primary storage applications, low-cost SATA disks for secondary applications (backup, archive, and so on), or a mix of the different disk types.

For more information, see

<http://www.netapp.com/us/products/platform-os/data-ontap-8/index.aspx>

Data ONTAP operating in 7-Mode

As previously mentioned, customers have a choice of deploying their NetApp storage environment operating in 7-Mode or clustered Data ONTAP. Data ONTAP operating in 7-Mode provides customers a broad suite of application integrations, storage efficiencies, and a legacy of customer satisfaction.

Data ONTAP 7-Mode is deployed on an HA pair of controllers that is discrete from any other storage systems in the environment and is managed as such. For this reason, the scalability with clustered Data ONTAP is superior to that of 7-Mode, which is further discussed in the next section on clustered Data ONTAP.

Clustered Data ONTAP

With clustered Data ONTAP, NetApp provides enterprise-ready, unified scale-out storage. Developed from a solid foundation of proven Data ONTAP technology and innovation, clustered Data ONTAP is the basis for large virtualized shared storage infrastructures that are architected for nondisruptive operations over the system lifetime. Controller nodes are deployed in HA pairs, with these HA pairs participating in a single storage domain or cluster.

Data ONTAP scale-out is a way to respond to growth in a storage environment. As the storage environment grows, additional controllers are added seamlessly to the resource pool residing on a shared storage infrastructure. Host and client connections as well as datastores can move seamlessly and nondisruptively anywhere in the resource pool, so that existing workloads can be easily balanced over the available resources, and new workloads can be easily deployed. Technology refreshes (replacing disk shelves, adding or completely replacing storage controllers) are accomplished while the environment remains online and continues serving data.

Data ONTAP is the first product to offer a complete scale-out solution, and it offers an adaptable, always-available storage infrastructure for today's highly virtualized environment.

VMware vSphere

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructure resources—CPUs, storage, networking—as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

The VMware vSphere environment delivers a robust application environment. For example, with VMware vSphere, all applications can be protected from downtime with VMware High Availability (HA) without the complexity of conventional clustering. In addition, applications can be scaled dynamically to meet changing loads with capabilities such as Hot Add and VMware Distributed Resource Scheduler (DRS).

For more information, see:

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

Domain and Element Management

This section of the document provides general descriptions of the domain and element managers used during the validation effort. The following managers were used:

- Cisco UCS Manager
- NetApp OnCommand®

- VMware vCenter™ Server
- Cisco UCS Director

Cisco Unified Computing System Manager

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

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

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 existing virtual machine mobility.

For more Cisco UCS Manager information, see:

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

NetApp OnCommand System Manager

NetApp OnCommand System Manager makes it possible for administrators to manage individual or clusters of NetApp storage systems through an easy-to-use interface simplifying common storage tasks such as creating volumes, LUNs, qtrees, shares, and exports, which saves time and prevents errors. System Manager works across all NetApp storage: FAS2000, FAS3000, and FAS6000 series and V-Series systems.

NetApp OnCommand Unified Manager complements the features of System Manager by enabling the monitoring and management of storage within the NetApp storage infrastructure.

The solution uses both OnCommand System Manager and OnCommand Unified Manager to provide storage provisioning and monitoring capabilities within the infrastructure.

VMware vCenter Server

VMware vCenter Server is the simplest and most efficient way to manage VMware vSphere, irrespective of the number of VMs you have. It 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. A single administrator can manage 100 or more virtualization environment workloads using VMware vCenter Server, more than doubling typical productivity in managing physical infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

For more information, see:

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

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 FlexPod validation process.

NetApp Virtual Storage Console

The NetApp VSC software delivers storage configuration and monitoring, datastore provisioning, VM cloning, and backup and recovery of VMs and datastores. VSC also includes an application-programming interface (API) for automated control. VSC delivers a single VMware plug-in that provides end-to-end VM lifecycle management for VMware environments using NetApp storage. VSC is delivered as a VMware vCenter Server plug-in. It is available to all VMware vSphere Clients that connect to the VMware vCenter Server. This is different from a client-side plug-in that must be installed on every VMware vSphere Client. The VSC software can be installed either on the VMware vCenter Server or on a separate Microsoft Windows Server® instance or VM.

Cisco Nexus 1000v vCenter Plugin

Cisco Nexus 1000V V2.2 (Advanced Edition) supports a plug-in for the vCenter Web Client. 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 administrators can export the necessary networking details from the vCenter server, investigate the root cause of and prevent the networking issues, and deploy the virtual machines with suitable policies. The server administrators can monitor and manage the resources effectively with the network details provided in the vCenter plug-in.

Cisco UCS Director for FlexPod Solution

Cisco UCS Director is an integral companion for FlexPod because it allows holistic management through centralized automation and orchestration from a single, unified view. When FlexPod and Cisco UCS Director are combined, IT can shift time and focus from managing infrastructure to delivering new service innovation. Used together, FlexPod and Cisco UCS Director deliver:

- Enhanced IT agility with a prevalidated, unified architecture that easily scales up or out to large data-center environments without design changes
- Dramatically reduced capital and operating expenses through end-to-end management of the FlexPod platform with real-time reporting of utilization and consumption based on trends set to customer-specific time frames
- Enhanced collaboration between computing, network, storage, and virtualization teams, allowing subject matter experts to define policies and processes that are utilized when resources are consumed
- Support for multiple infrastructure stacks in a single data center, as well as across multiple data centers globally

The extensive Cisco UCS Director task library lets you quickly assemble, configure, and manage workflows for FlexPod, Clustered ONTAP, and FlexPod Express. You can use the workflows immediately or publish them in an infrastructure catalog. Specific workflows can be assigned to the entire organization or specific groups based on your organizational structure, which can be imported from Lightweight Directory Access Protocol (LDAP). The drag-and-drop workflow designer tool eliminates the need for service engagements or the need to bring together multi-product solutions or third-party adapters.

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 within individual Cisco UCS Manager instances. 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.

FlexPod: Cisco Nexus 9000 Standalone Design

Physical Build

Hardware and Software Revisions

Table 1 describes the hardware and software versions used during solution validation. It is important to note that Cisco, NetApp, and VMware have interoperability matrixes that should be referenced to determine support for any specific implementation of FlexPod. Refer to the following links for more information:

- NetApp Interoperability Matrix Tool
<http://support.netapp.com/matrix/>
- Cisco UCS Hardware and Software Interoperability Tool
<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>
- VMware Compatibility Guide
<http://www.vmware.com/resources/compatibility/search.php>

Table 1 Validated Software Versions

Layer	Device	Image	Comments
Compute	Cisco UCS Fabric Interconnects 6200 Series, UCS B-200 M3, UCS C-220 M3	2.1(3a)	Includes the Cisco UCS-IOM 2208XP, Cisco UCS Manager, and UCS VIC 1240
	Cisco eNIC	2.1.2.38	
	Cisco Nexus 1000v	4.2(1)SV2(2.1a)	
	Cisco Nexus 1110-X	4.2(1)SP1(6.2)	
Network	Cisco Nexus 9396 NX-OS	6.1.(2)I2(1)	
Storage	NetApp FAS 3250-AE	Data ONTAP 8.2 P5	
	Nexus 5596 Cluster Switches	5.2(1)N1(1)	
Software	VMware vSphere ESXi	5.1u1	
	VMware vCenter	5.1u1	
	OnCommand Unified Manager for clustered Data ONTAP	6.0	
	NetApp Virtual Storage Console (VSC)	4.2.1	
	Cisco UCS Director	4.1	

Logical Build

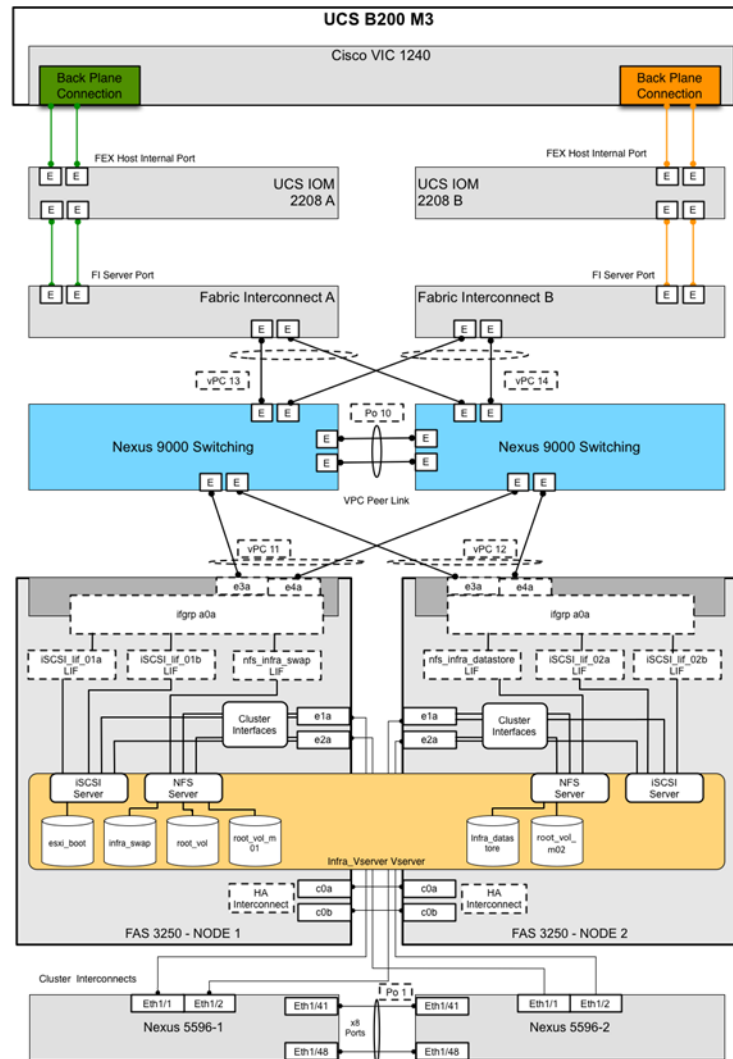
Figure 2 and Figure 3 illustrate the Nexus 9000 standalone design structure. The design is physically redundant across the stack, addressing Layer 1 high-availability requirements. The solution also incorporates additional Cisco and NetApp technologies and features that make for an even more effective design. This section of the document discusses the logical configuration validated for FlexPod. The topics covered include:

- FlexPod: Nexus 9000 standalone design with clustered Data ONTAP
- FlexPod: Nexus 9000 standalone design with Data ONTAP operating in 7-Mode

FlexPod: Cisco Nexus 9000 standalone Design with Clustered Data ONTAP

Figure 4 shows design details of the Nexus 9000 based FlexPod with NetApp clustered Data ONTAP storage. The details of compute, network and storage components within this design are covered in the upcoming sections.

Figure 4 Detailed Cisco Nexus 9000 Design for Clustered Data ONTAP



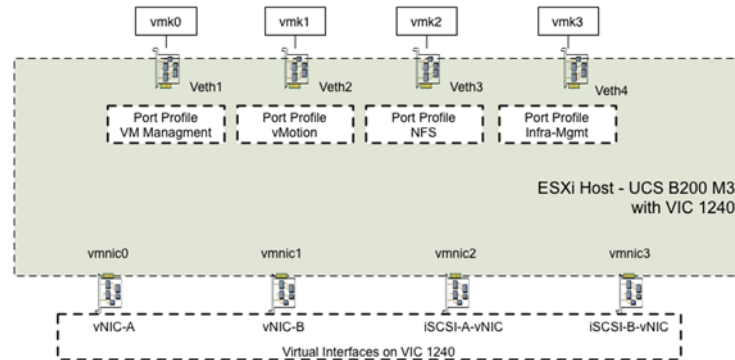
Cisco Unified Computing System

The FlexPod design supports both Cisco UCS B-Series and Cisco UCS C-Series deployments. This section of the document discusses the integration of each deployment into FlexPod. The Cisco Unified Computing System supports the virtual server environment by providing a robust, highly available, and extremely manageable compute resource. The components of the Cisco Unified Computing System offer physical redundancy and a set of logical structures to deliver a very resilient FlexPod compute domain. In this validation effort, multiple Cisco UCS B-Series ESXi servers are booted from SAN using iSCSI. The ESXi nodes consisted of Cisco UCS B200-M3 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.

As shown in Figure 4, the FlexPod defines two LAN port channels (Po11 and Po12) to connect Cisco UCS Fabric Interconnects to the Nexus 9000. Both NFS and SAN (iSCSI) traffic from the servers use these two port-channels to communicate to the storage system. At the server level, the Cisco 1240 VIC presents four virtual PCIe devices to the ESXi node, two virtual 10 Gb Ethernet NICs (vNIC) and two

10 Gb iSCSI vNICs (Figure 5). The vSphere environment identifies these interfaces as vmnics and the ESXi operating system is unaware these are virtual adapters. The result is a dual-homed ESXi node to the remaining network.

Figure 5 ESXi Server Utilizing vNICs



Cisco Unified Computing System I/O Component Selection

FlexPod allows customers to adjust the individual components of the system to meet their particular scale or performance requirements. Selection of I/O components has a direct impact on scale and performance characteristics when ordering the Cisco UCS components. Figure 6 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:

- 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 2 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)

FlexPod 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 6 shows an overview of backplane connectivity for both the I/O Modules and Cisco VICs.

Figure 6 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 the figure, 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 number of KR lanes varies based on selection of VIC 1240, VIC 1240 with Port Expander and VIC 1280.

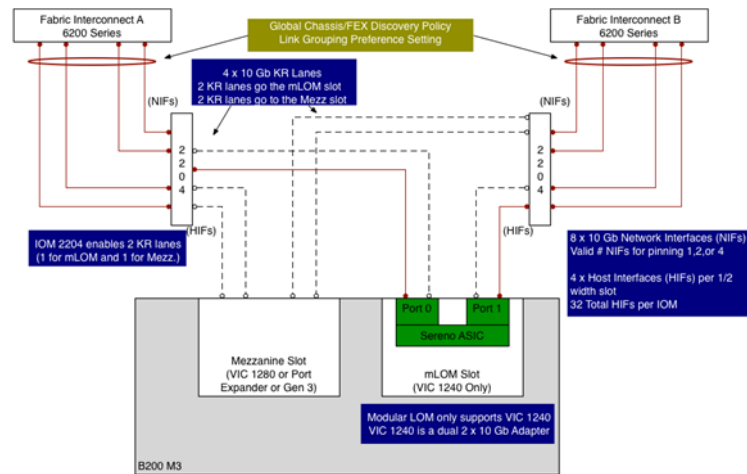
Validated I/O Component Configurations

Two of the most commonly validated I/O component configurations in FlexPod designs are:

- Cisco UCS B200M3 with VIC 1240 and FEX 2204
- Cisco UCS B200M3 with VIC 1240 and FEX 2208

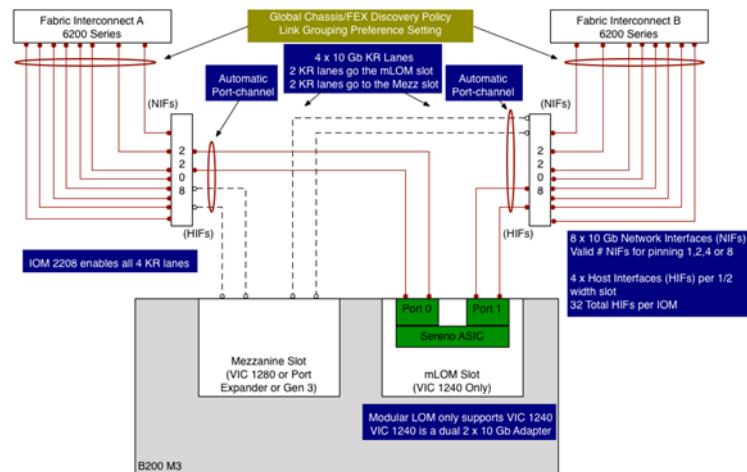
Figure 7 and Figure 8 show the connectivity for these two configurations.

Figure 7 Validated Backplane Configuration—VIC 1240 with FEX 2204



In Figure 7, 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 8 Validated Backplane Configuration—VIC 1240 with FEX 2208




In Figure 8, the FEX 2208XP enables 8 KR lanes to the half-width blade while the global discovery policy dictates the formation of a fabric port channel. Since VIC 1240 is not using a Port Expander module, this configuration results in 40GbE connection 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 9). 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. When there is 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 FlexPod designs.

Figure 9 Chassis Discovery Policy—Discrete Mode vs. Port Channel Mode



Cisco Unified Computing System - QoS and Jumbo Frames

FlexPod accommodates a myriad of traffic types (vMotion, NFS, FCoE, control traffic, etc.) and is capable of absorbing traffic spikes and protect against traffic loss. Cisco UCS and Nexus QoS system classes and policies deliver this functionality. In this validation effort the FlexPod was configured to support jumbo frames with an MTU size of 9000. Enabling jumbo frames allows the FlexPod environment to optimize throughput between devices while simultaneously reducing the consumption of CPU resources.



Note

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.

Cisco Unified Computing System - C-Series Server Design

Cisco UCS Manager 2.1 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 2.1. 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 2.1 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.



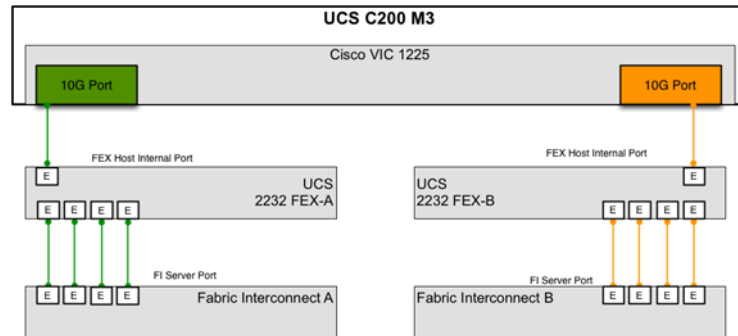
Note

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

Figure 10 illustrates the connectivity of the Cisco UCS C-Series server into the Cisco UCS domain. From a functional perspective the 1 RU 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 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 B and C series deployment models and the virtual constructs formed at the vSphere or Nexus 1000v layer are unaware of the platform in use.

Figure 10 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 NetApp storage system and the enterprise network. In the FlexPod design, Cisco UCS Fabric Interconnects, NetApp storage systems, and Cisco Nexus 1110 devices 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 make sure high availability of the overall FlexPod system

Figure 11 Cisco Nexus 9000 Connections to Cisco UCS Fabric Interconnects and NetApp FAS

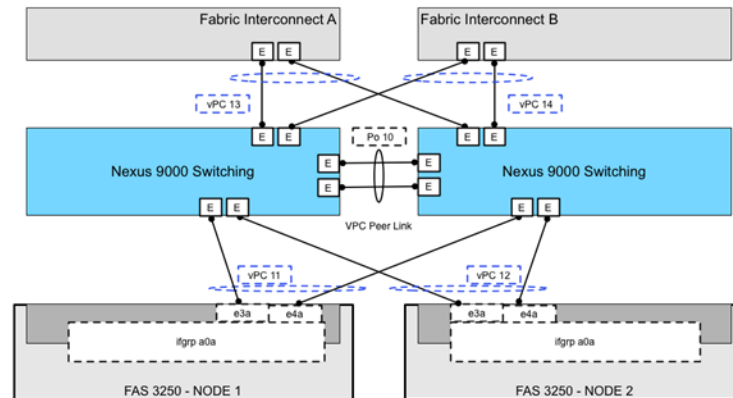


Figure 11 shows the connections between Cisco Nexus 9000, Cisco UCS Fabric Interconnects and NetApp FAS 3250. vPC requires a "peer link" which is documented as port channel 10 in this diagram. In addition to the vPC peer-link, vPC peer keepalive link is a required component of a vPC configuration. The peer keepalive link allows each vPC enabled switch to monitor the health of its peer. This link accelerates convergence and reduces the occurrence of split-brain scenarios. In this validated solution, the vPC peer keepalive link uses the 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 FlexPod 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 (CFS over E) 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 the NetApp storage controller and UCS 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>

VMware vCenter and vSphere

VMware vSphere 5.1 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
- vSphere Virtual Machine File System (VMFS)
- vSphere High Availability (HA)
- vSphere Distributed Resource Scheduler (DRS)

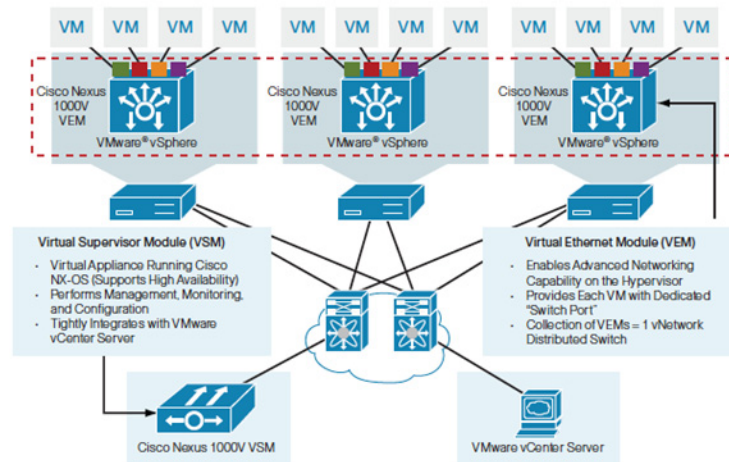
Cisco Nexus 1000v

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) runs on a Nexus 1110X and 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 12 describes the Cisco Nexus 1000v architecture.

Figure 12 Cisco Nexus 1000v Architecture



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

- 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

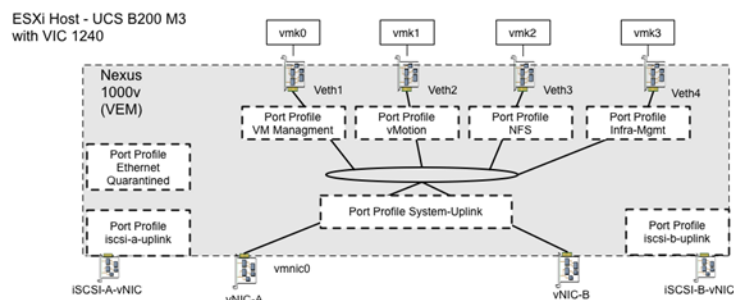


Note

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

Figure 13 shows a single ESXi node with a VEM registered to the Cisco Nexus 1000v VSM. The ESXi vmnics are presented as Ethernet interfaces in the Nexus 1000v.

Figure 13 Cisco Nexus 1000v VEM in an ESXi Environment



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 Nexus 1000v. Nexus 1000v aggregates the Ethernet uplinks into a single port channel named the "System-Uplink" port profile for fault tolerance and improved throughput (Figure 13).



Note

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

The VM facing virtual Ethernet ports employ port profiles customized for each virtual machines network, security and service level requirements. The FlexPod 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 FlexPod 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 ultimately 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 on "Best Practices in Deploying Cisco Nexus 1000V Series Switches on Cisco UCS B and C Series Cisco UCS Manager Servers" see http://www.cisco.com/en/US/prod/collateral/switches/ps9441/ps9902/white_paper_c11-558242.html


NetApp Storage Controllers

Clustered Data ONTAP

Clustered Data ONTAP allows one or more storage HA pairs that are interconnected to be managed as a single system or pool of storage resources.

Figure 14 details the logical configuration of the clustered Data ONTAP environment used during validation. The physical cluster consists of two NetApp storage controllers (nodes) configured in an HA pair and two cluster interconnect switches.

Figure 14 NetApp Storage Controller—Clustered Data ONTAP



The fundamental connections or network types defined for a clustered Data ONTAP solution include:

- HA interconnect: a dedicated interconnect between two nodes permitting the formation of HA pairs. These are also known as storage failover pairs.
- Cluster interconnect: a dedicated high-speed, low-latency, private network used for communication between nodes.
- Management network: a network used for the administration of nodes, cluster, and storage virtual machines (SVMs), also known as Vservers.
- Data network: a network used by clients to access data.

As illustrated, the storage controllers use multiple constructs to abstract the physical resources. These elements include:

- Ports. A physical port such as e0a or e1a or a logical port such as a virtual LAN (VLAN) or an interface group (ifgrp).
- Ifgrps. A collection of physical ports to create one logical port constitutes an interface group. NetApp's interface group is a link aggregation technology and may be deployed in single (active/passive), multiple ("always on"), or dynamic (active LACP) mode, but it is recommended to use only dynamic interface groups to take advantage of LACP-based load distribution and link failure detection.
- LIF. A logical interface that is associated to a physical port, interface group, or VLAN interface. More than one LIF may be associated to a physical port at the same time. There are three types of LIFs:
 - NFS LIF
 - iSCSI LIF
 - FC LIF

LIFs are logical network entities that have the same characteristics as physical network devices but are not tied to physical objects. LIFs used for Ethernet traffic are assigned specific Ethernet-based details such as IP addresses and iSCSI-qualified names and then are associated with a specific physical port capable of supporting Ethernet. NAS LIFs can be nondisruptively migrated to any other physical network port throughout the entire cluster at any time, either manually or automatically (by using policies).

- Storage virtual machines. An SVM is a secure virtual storage server that contains data volumes and one or more LIFs, through which it serves data to the clients. An SVM securely isolates the shared virtualized data storage and network and appears as a single dedicated server to its clients. Each SVM has a separate administrator authentication domain and can be managed independently by an SVM administrator.

Nodes 1 and 2 form a two-node storage failover pair through the HA interconnect direct connection. The FlexPod design uses the following port and interface assignments:

- Ethernet ports e3a and e4a on each node are members of a multimode LACP interface group for Ethernet data. This design leverages an interface group that has LIFs associated with it to support NFS and iSCSI traffic.
- Ports e0M are on each node and support a LIF dedicated to node management. Port e0b is defined as a failover port supporting the "node_mgmt" role.
- Port e0a supports cluster management data traffic through the cluster management LIF. This port and LIF allow for administration of the cluster from the failover port and LIF if necessary.
- Ports c0a and c0b on each node support the HA interconnect processes. These connections do not support any data traffic but only control processes.
- Ports e1a and e2a are cluster interconnect ports for data traffic. These ports connect to each of the Cisco Nexus 5596 cluster interconnect switches.
- The Cisco Nexus Cluster Interconnect switches support a single ISL port channel (Po1).



Note

The cluster interconnect switch configuration is provided by NetApp at https://library.netapp.com/ecm/ecm_get_file/ECMP1115327.

The solution defines a single infrastructure SVM to own and export the data necessary to run the VMware vSphere infrastructure. This SVM specifically owns the following flexible volumes:

- **Root volume.** A flexible volume that contains the root of the SVM namespace.
- **Root volume load-sharing mirrors.** A mirrored volume of the root volume to accelerate read throughput. In this instance, it is labeled root_vol_m01 and root_vol_m02.
- **Boot volume.** A flexible volume that contains ESXi boot LUNs. These ESXi boot LUNs are exported through iSCSI to the Cisco UCS servers.
- **Infrastructure datastore volume.** A flexible volume that is exported through NFS to the ESXi host and is used as the infrastructure NFS datastore to store VM files.
- **Infrastructure swap volume.** A flexible volume that is exported through NFS to each ESXi host and used to store VM swap data.

The NFS datastores are mounted on each VMware ESXi host in the VMware cluster and are provided by NetApp clustered Data ONTAP through NFS over the 10GbE network.

The SVM has a minimum of one LIF per protocol per node to maintain volume availability across the cluster nodes. The LIFs use failover groups, which are network policies defining the ports or interface groups available to support a single LIF migration or a group of LIFs migrating within or across nodes in a cluster. Multiple LIFs may be associated with a network port or interface group. In addition to

failover groups, the clustered Data ONTAP system uses failover policies. Failover policies define the order in which the ports in the failover group are prioritized. Failover policies define migration policy in the event of port failures, port recoveries, or user-initiated requests.

The most basic possible storage failover scenarios in this cluster are as follows:

- Node 1 fails, and Node 2 takes over Node 1's storage.
- Node 2 fails, and Node 1 takes over Node 2's storage.

The remaining node network connectivity failures are addressed through the redundant port, interface groups, and logical interface abstractions afforded by the clustered Data ONTAP system.

NetApp clustered Data ONTAP can be deployed without the cluster interconnect switches when deploying a two-node storage system. This is commonly referred to a switchless cluster. In this scenario, the cluster interfaces e1a and e2a on the storage controllers are connected to their partner nodes instead of the cluster interconnect switches.

This design configuration eliminates the need of cluster interconnect switches for clustered Data ONTAP storage systems composed of only two nodes. Future scale-outs can be easily accommodated by adding in the cluster interconnects using a predefined nondisruptive upgrade procedure, making sure of zero downtime to the storage system.

With the use of a switchless cluster deployment model, Data ONTAP 7-Mode storage systems can also be migrated to clustered Data ONTAP without any additional hardware.

Data ONTAP Operating in 7-Mode

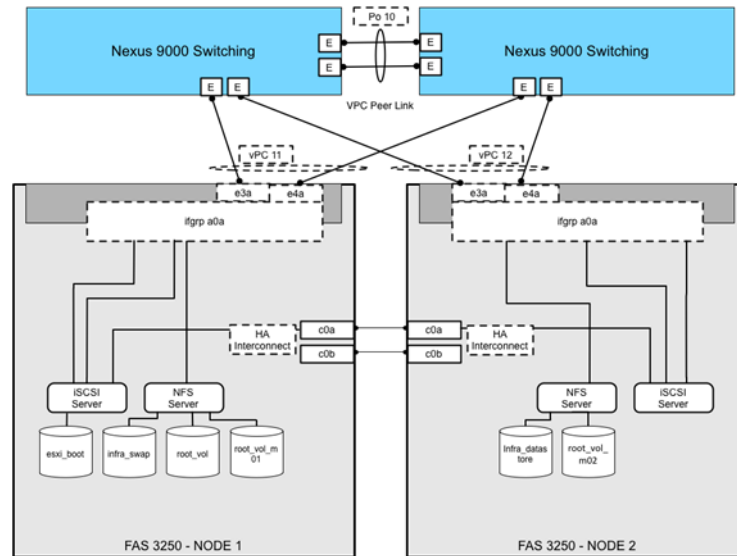
Figure 15 shows a FlexPod unit with Data ONTAP operating in 7-Mode. Data ONTAP 7-Mode consists of only two storage controllers with shared media. From a design perspective, the Cisco Nexus and Cisco UCS component configurations are identical to the previously defined FlexPod configuration with clustered Data ONTAP.

The NetApp FAS controllers use redundant 10GbE converged adapters configured in a two-port interface group (ifgrp). Each port of the ifgrp is connected to one of the upstream switches, allowing multiple active paths by utilizing the Cisco Nexus vPC feature. Ifgrp is a mechanism that allows the aggregation of a network interface into one logical unit. Combining links aids in network availability and bandwidth. NetApp provides three types of ifgrps for network port aggregation and redundancy:

- Single mode
- Static multimode
- Dynamic multimode

Using dynamic multimode ifgrps is recommended due to the increased reliability and error reporting and because they are also compatible with Cisco Virtual Port Channels. A dynamic multimode ifgrp uses Link Aggregation Control Protocol (LACP) to group multiple interfaces together to act as a single logical link. This provides intelligent communication between the storage controller and the Cisco Nexus switch and enables load balancing across physical interfaces as well as failover capabilities.

Figure 15 NetApp Storage Controller—7-Mode



Conclusion

FlexPod is the optimal shared infrastructure foundation to deploy a variety of IT workloads. Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and applications. From virtual desktop infrastructure to SAP®, FlexPod can efficiently and effectively support business-critical applications running simultaneously from the same shared infrastructure. The flexibility and scalability of FlexPod also enable customers to start out with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements.

References

Cisco Unified Computing System:

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

Cisco UCS 6200 Series Fabric Interconnects:

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

Cisco UCS 5100 Series Blade Server Chassis:

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

Cisco UCS B-Series Blade Servers:

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

Cisco UCS Adapters:

http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html

Cisco UCS Manager:

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

Cisco Nexus 9000 Series Switches:

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

Cisco Nexus 5000 Series Switches:

<http://www.cisco.com/en/US/products/ps9670/index.html>

Cisco Nexus 1000v:

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

Cisco Prime Data Center Manager:

<http://www.cisco.com/en/US/products/ps9369/index.html>

VMware vCenter Server:

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

VMware vSphere:

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

Interoperability Matrixes

VMware and Cisco Unified Computing System:

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

© 2014 NetApp, Inc. and Cisco Systems, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. and Cisco Systems, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, FlexPod, MetroCluster, OnCommand, SnapLock, and SnapMirror are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries.