




FlexPod with Cisco UCS Mini Design Guide

Last Updated: February 11, 2015



Building Architectures to Solve Business Problems



About the Author

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

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, Chris was an application developer and has worked in the IT industry for more than 15 years.

Lindsey Street, Solutions Architect, Infrastructure and Cloud Engineering, NetApp

Lindsey Street is a Solutions Architect in the NetApp Infrastructure and Cloud Engineering team. She focuses on the architecture, implementation, compatibility, and security of innovative vendor technologies to develop competitive and high-performance end-to-end cloud solutions for customers. Lindsey started her career in 2006 at Nortel as an interoperability test engineer, testing customer equipment interoperability for certification. Lindsey has her Bachelors of Science degree in Computer Networking and her Masters of Science in Information Security from East Carolina University.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to acknowledge the significant contribution and expertise that resulted in developing this document:

- Haseeb Niazi, Cisco Systems, Inc.
- Gangoor Sridhara, Cisco Systems, Inc.
- John George, NetApp (Contributor)

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)

© 2015 Cisco Systems, Inc. All rights reserved



FlexPod with Cisco UCS Mini 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 use of Cisco Unified Computing System (UCS) Mini platform into the FlexPod Express and Data Center solutions. The FlexPod® solution portfolio combines NetApp® storage systems, Cisco® UCS servers, and Cisco Nexus fabric into a single, flexible architecture. FlexPod solutions are designed and validated to reduce deployment time, project risk, and the cost of IT.

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.

Overview of Market Segments

The consumption of integrated systems; server, shared-storage and network hardware, to provide a shared compute infrastructure is a pervasive technology trend within IT organizations. The initial business drivers behind this movement include

- **Reducing cost:** to reduce the time and money spent on managing servers, storage and applications. This also includes the power and cooling costs associated with the infrastructure.
- **Supporting business requirements:** to create a more flexible and agile infrastructure to respond to the business needs faster and ultimately promote business growth
- **Reducing time-to-market:** to provide the ability to deploy new applications and workloads in minutes (versus days, weeks, months)



On further investigation, the broader integrated infrastructure market can be subdivided into smaller segments each with their own challenges, deployment characteristics and use cases. Key target markets for integrated infrastructure include:

- Remote Office/Branch Office (ROBO)
- Small and Midsize Business (SMB)
- Data Center
- Security or Demilitarized Zone (DMZ)



Note

Each of these segments can be readily addressed by the scale-up and scale-out capabilities of the FlexPod integrated infrastructure.

Remote Office/Branch Office

Remote and Branch office (ROBO) environments are typically large companies with offices and retail locations that contain a relatively large user population or have a significant requirement for compute applications.

These environments have an extremely wide spread of applications and user populations ranging from something like an insurance agent location, which might have one or two users, to a large district office that hosts several hundred employees. Similarly in a retail business, stores can range from a small site with minimal compute requirements to larger stores requiring multiple in-store applications spread across multiple servers for availability or performance reasons. In many of these environments (especially in the retail space), 24x7 availability of local IT staff is not only costly but highly inefficient and unlikely. Therefore, configuration, setup and ongoing management must ideally be performed by the central IT organization.

To this end, IT organizations have attempted to relocate applications centrally in the data center to achieve their goals. Unfortunately, there are a number of applications that cannot be centrally positioned and require locality to achieve appropriate service levels due to the following causes:

- **Performance:** Applications that cannot tolerate the latency and bandwidth limitations of a WAN link because they either require fast response times or generate large amounts of data
- **Availability:** Applications that cannot tolerate the unreliability of a WAN link because they require continuous availability
- **Compliance:** Applications that cannot be hosted outside of the branch office or require a local copy of data because they must meet security, regulatory, or internal policy requirements

Some applications that have defied most centralization efforts and continue to be deployed locally at the edge of the network are as follows:

- **Edge infrastructure services:** Such as Microsoft AD DS, print service, Dynamic Host Configuration Protocol (DHCP) server, Domain Name System (DNS) server, and file server
- **Line-of-business client-server applications:** Point-of-Sale (POS) systems, bank teller in-office control points (IOCP), and electronic-medical-record (EMR) systems
- **Branch-office management software:** Configuration management, operations management, monitoring services, update and patching services, and backup and recovery services

These applications can be found at businesses such as:

- Retail stores
- Oil rigs

- Health care

Managing a large number of geographically dispersed sites utilizing centralized management and control is paramount to address compliance requirements and to provide a consistent IT infrastructure to the ROBO.

Small and Midsize Business

The SMB segment is a broad market sector that covers a very wide variety of firms. These firms often have constrained IT budgets, limited IT staff and diverse IT requirements. The typical method of separating out this segment is to split it up by the number of employees. For example;

- Small (1-99 employees)
- Lower midsize (100-499 employees)
- Upper midsize (500-999 employees)

Simplicity is the primary demand of the SMB segment. In some cases an SMB firm will consume managed services to deliver IT a clear indicator their concern is not with the business of IT but of business outcomes. Typically the SMB segment requires more shared infrastructure resources than a ROBO solution which may translate into a number of distinct domains with the SMB data center.

Use cases include:

- VDI deployments
- Regional HQs
- Large offices in Metros



Note

The FlexPod infrastructure can easily scale to address the requirements of each sub-segment.

Simplicity is the primary demand of the SMB segment. In many cases, SMB firms will choose to deliver IT services by selecting a managed services route, a clear indicator that the firms do not want to invest in sustaining distributed IT operations. Typically the SMB segment requires more shared infrastructure resources than a ROBO solution which may translate into a number of distinct domains with the SMB data center.

Some of the typical SMB use cases include:

- VDI deployments to support distributed workforce
- Regional HQs hosting a number critical applications
- Large offices in Metros with metro specific application requirements

Data Centers

The data center market segment is shifting toward heavily virtualized private, hybrid, and public cloud computing models running on industry-standard systems. These environments require uniform design points that can be repeated for ease of management and scalability leading to the need for predesigned compute, network, and storage building blocks optimized to lower the initial design cost, simplify management, enable horizontal scalability and provide high levels of utilization. This infrastructure trend does not remove the requirement on IT to address security, compliance, and acceptable fault

domains within the data center. These requisites can either be internally defined and delivered or outsourced in a managed service model. Enterprise and Service Provider data centers interested in small failure domains fall into this sector.

Security Zones or Demilitarized Zone (DMZ)

Changing business models, a dynamic threat landscape, and complex, fragmented security solutions have required many organizations to revisit their security posture and their infrastructure. To address today's security challenges, IT organizations are looking to implement simpler, scalable infrastructure solutions. Integrated infrastructures provide a well-known well-defined platform with uniform attack surfaces and vectors that are easily understood, self-contained and therefore readily protected. This use case is particularly evident in the DMZ where public meets private and customers require an isolated domain for security purposes.

FlexPod Solutions

The FlexPod® solution portfolio combines NetApp® storage systems, Cisco® Unified Computing System servers, and Cisco Nexus fabric into a single, flexible architecture. There are three primary FlexPod solutions FlexPod Express, FlexPod Data Center and FlexPod Select.

FlexPod Express

FlexPod Express is a suitable platform for running a variety of virtualization hypervisors as well as bare metal operating systems and enterprise workloads. FlexPod Express delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements.

FlexPod Express:

- Combines all application and data needs into one platform
- Suitable for small-midsize organizations, remote and departmental deployments
- Provides easy infrastructure scaling
- Reduces cost and complexity

FlexPod Data Center

FlexPod Data Center can scale up for greater performance and capacity or scale out for environments that need consistent, multiple deployments; FlexPod also has the flexibility to be sized and optimized to accommodate different use cases including application workloads such as MS SQL Server, Exchange 2010, MS SharePoint 2010, SAP, Red Hat, VDI (VMware View/Citrix XenDesktop), or Secure Multi-tenancy (SMT) environments.

FlexPod Data Center:

- Provides storage, networking and server components from NetApp and Cisco
- Is a validated enterprise-class IT platform
- Speeds deployment of business critical applications

- Reduces cost, minimal risk, increased flexibility
- Scales up or out for future growth

FlexPod Select

FlexPod Select offers pre-configured NetApp and Cisco components for dedicated high-performance analytics workloads. FlexPod® Select solutions are designed for high-performance workloads, such as big data, HPC, databases, and data warehouses. These pre-configured, scalable solutions combine storage, networking, and servers. Validated with analytical applications to reduce risk, improve flexibility, and expedite deployment.

FlexPod Select components include NetApp® E-series, EF Series and FAS storage systems, Cisco® Unified Computing System™ servers and fabric interconnect, and Cisco Nexus® fabric extender.

FlexPod Select:

- Is a pre-configured for dedicated high-performance workloads
- Accelerates time-to-value
- Reduces risk, increases flexibility, speeds deployment
- Provides an open platform

FlexPod with Cisco UCS Mini

FlexPod with Cisco UCS Mini is part of the FlexPod Express and FlexPod Data Center solutions because solution is not customized for any specific workload. The FlexPod with Cisco UCS Mini allows IT departments to address ROBO, SMB or Data Center infrastructure challenges using a streamlined architecture following compute, network and storage best practices. The following sections describe each of the solution components, their structure and functionality.

FlexPod with Cisco UCS Mini Model

Cisco Nexus Switching
(conditional)



Cisco UCS 5108 Chassis with
Cisco UCS 6324 Fabric
Interconnect (Mini)



NetApp FAS2500



Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS®) is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. 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 main components of the Cisco Unified Computing System are as follows:

- **Compute**—The system is based on an entirely new class of computing system that incorporates rack mount and blade servers based on Intel processors.
- **Network**—The system is integrated onto a low-latency, lossless, 10-Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- **Virtualization**—The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access**—The system provides consolidated access to both SAN storage and Network Attached Storage (NAS) over the unified fabric. By unifying the storage access the Cisco Unified Computing System can access storage over Ethernet (NFS or iSCSI), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This provides customers with storage choices and investment protection. In addition, the server administrators can pre-assign storage-access policies to storage resources, for simplified storage connectivity and management leading to increased productivity.

Cisco Unified Computing System (Cisco UCS) fuses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

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

Figure 1 *The Cisco Unified Computing System is a Highly Available Cohesive Architecture*

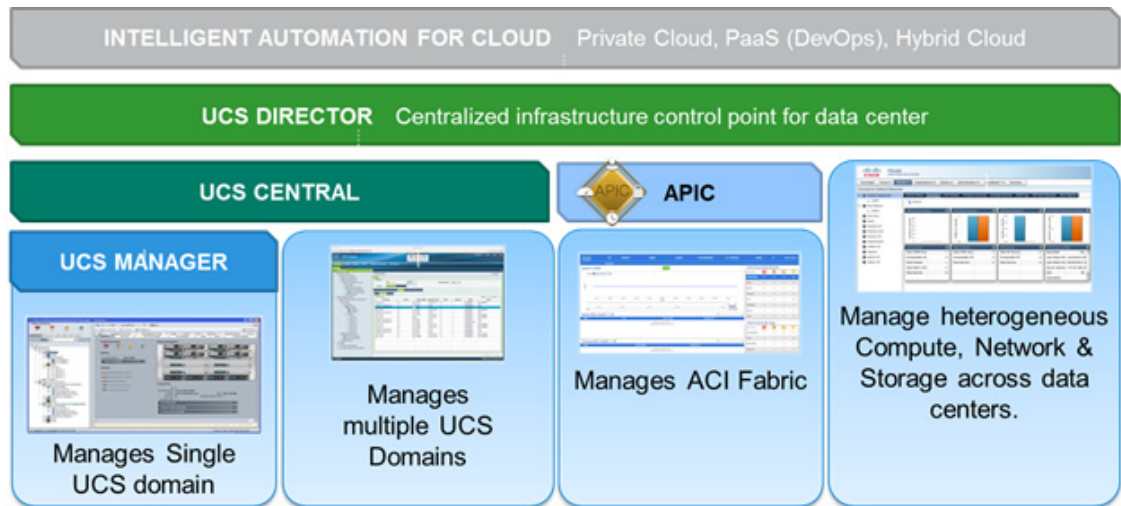


Cisco UCS 6324 Fabric Interconnect Overview

The Cisco UCS 6324 Fabric Interconnect Fabric Interconnect provides the management, LAN, and storage connectivity for the Cisco UCS 5108 Blade Server Chassis and direct-connect rack-mount servers. It provides the same full-featured Cisco UCS management capabilities and XML API as the full-scale Cisco UCS solution in addition to integrating with Cisco UCS Central Software and Cisco UCS Director (Figure 2).

From a networking perspective, the Cisco UCS 6324 Fabric Interconnect uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10 Gigabit Ethernet on all ports, switching capacity of up to 500 Gbps, and 80-Gbps uplink bandwidth for each chassis, independent of packet size and enabled services. Sixteen 10-Gbps links connect to the servers, providing a 20-Gbps link from each Cisco UCS 6324 Fabric Interconnect to each server. The product family supports Cisco® low-latency, lossless 10 Gigabit Ethernet unified network fabric capabilities, which increase the reliability, efficiency, and scalability of Ethernet networks. The fabric interconnect supports multiple traffic classes over a lossless Ethernet fabric from the blade through the fabric interconnect. Significant TCO savings come from an Fibre Channel over Ethernet (FCoE)-optimized server design in which network interface cards (NICs), host bus adapters (HBAs), cables, and switches can be consolidated.

Figure 2 Cisco UCS Management and Fabric Architecture



Unified Fabric with FCoE: I/O Consolidation

The Cisco UCS 6324 Fabric Interconnect is built to consolidate LAN and storage traffic onto a single unified fabric, eliminating the capital expenditures (CapEx) and operating expenses (OpEx) associated with multiple parallel networks, different types of adapter cards, switching infrastructure, and cabling within racks. The unified ports allow the fabric interconnect to support direct connections from Cisco UCS to Fibre Channel, FCoE, and Small Computer System Interface over IP (iSCSI) storage devices.

Cisco UCS Manager

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

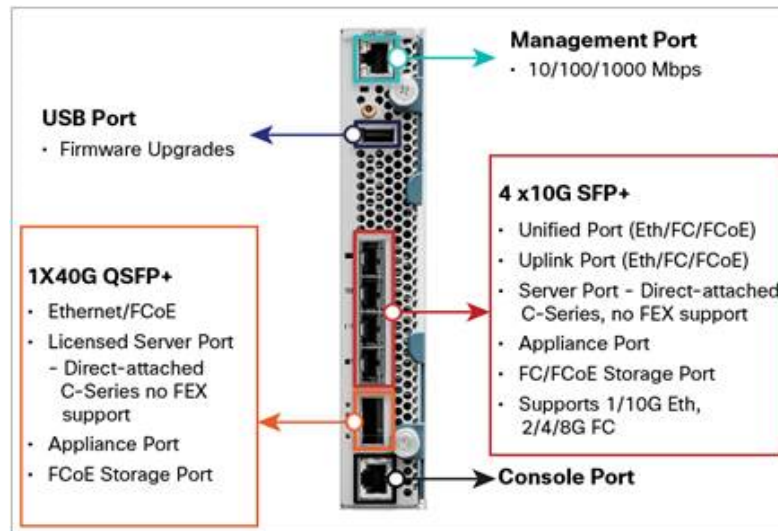
Optimization for Virtualization

For virtualized environments, the Cisco UCS 6324 Fabric Interconnect supports Cisco virtualization-aware networking and Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) architecture. Cisco Data Center VM-FEX allows the interconnects to provide policy-based virtual machine connectivity, with network properties moving with the virtual machine and a consistent operational model for both physical and virtual environments.

Cisco UCS 6324 Fabric Interconnect

The Cisco UCS 6324 Fabric Interconnect Fabric Interconnect (Figure 3) is a 10 Gigabit Ethernet, FCoE, and Fibre Channel switch offering up to 500-Gbps throughput and up to four unified ports and one scalability port.

Figure 3 Cisco UCS 6324 Fabric Interconnect



I/O Adapters

The FlexPod with Cisco UCS Mini uses the Cisco UCS VIC Virtual Interface Card (VIC) for all designs. VIC obviates the need for multiple network interface cards (NICs) and host bus adapters (HBAs) by converging LAN and SAN traffic in a single physical interface. The Cisco Virtual Interface Card (VIC) delivers 256 virtual interfaces and supports Cisco VM-FEX technology. The Cisco VIC provides I/O policy coherency and visibility to enable true workload mobility in virtualized environments. The Cisco VIC is available in a number of form factors for both B-Series blades and C-Series rack servers. This wire-once architectural approach and centralized management through the Cisco UCS Manager helps you:

- Simplify cabling
- Reduce points of management
- Lower costs

Cisco UCS Central (Optional)

Cisco UCS Central Software extends the simplicity and agility of managing a single Cisco UCS domain across multiple Cisco UCS domains. Cisco UCS Central Software allows organizations to easily work on a global scale, putting computing capacity close to users while managing infrastructure with centrally defined policies. As illustrated in [Figure 4](#), Cisco UCS Central supports a centralized policy model across multiple Cisco UCS domains in a given organization simplifying operations, visibility and control.

customers the flexibility to select the right storage controllers to support their business requirements. The storage efficiency built in to Data ONTAP provides substantial space savings, which allows more data to be stored at a lower cost.

NetApp offers the NetApp Unified Storage Architecture that simultaneously supports 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, and FCP. Connectivity options include standard Ethernet (10/100/1000, or 10GbE) and Fibre Channel (2, 4, 8 or 16Gb/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, refer to [NetApp Data ONTAP 8.3 Operating System](#).

The NetApp FAS family of controllers consists of several different sizes of controllers to meet the needs of a wide range of customers. For the FlexPod architecture with the Cisco UCS mini, NetApp recommends the FAS2552, the FAS2554, or the FAS8020 controller to support a remote office solution.

The FAS2552 is a 2-rack unit controller that includes 2 controllers and 24 2.5-inch disks in a single chassis. The FAS2554 is a 4-rack unit controller that includes 2 controllers and 24 3.5-inch disks in a single chassis. The FAS8020 is a 3-rack unit controller that has no disks in the chassis. The FAS2552 and FAS2554 both support a maximum of 5 external disk shelves, and the FAS8020 supports up to 20 external disk shelves.

NetApp Clustered Data ONTAP Overview

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 that participate in a single storage domain or cluster.

Data ONTAP scale-out is a way to respond to growth in a storage environment. Growth can encompass adding new remote office locations, adding individual disk shelves, or adding additional controllers to an existing system. 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 non-disruptively 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 (such as replacing disk shelves or 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.

NetApp Storage Virtual Machines

A cluster serves data through at least one and possibly multiple storage virtual machines (SVMs). An SVM is a logical abstraction that represents a set of physical resources of the cluster. Data volumes and logical network interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved non-disruptively from one node to another. For example, a flexible volume may be non-disruptively moved to a new node and aggregate, or a data LIF could be transparently reassigned to a different physical network port. In this manner, the SVM abstracts the cluster hardware and is not tied to specific physical hardware.

An SVM is capable of supporting multiple data protocols concurrently. Volumes within the SVM can be junctioned together to form a single NAS namespace, which makes all of an SVM's data available through a single share or mount point to NFS and CIFS clients. SVMs also support block-based protocols, and LUNs can be created and exported using iSCSI, Fibre Channel, or Fibre Channel over Ethernet. Any or all of these data protocols may be configured for use within a given SVM.

Because it is a secure entity, an SVM is only aware of the resources that have been assigned to it and has no knowledge of other SVMs and their respective resources. Each SVM operates as a separate and distinct entity with its own security domain. Tenants may manage the resources allocated to them through a delegated SVM administration account. Each SVM may connect to unique authentication zones such as Active Directory®, LDAP, or NIS.

Cisco Nexus Switching Technologies

Customer demands of the network fabric are evolving as the number of use cases continues to expand. To this end Cisco's Nexus switching portfolio continues to advance to meet these new requirements. The Cisco Nexus 9000 product family is designed to meet both current and emerging needs for multiple stages of network transformation.

The Cisco Nexus 9000 provides two modes of operation NX-OS mode and Application Centric Infrastructure (ACI) mode. In NXOS mode the Cisco Nexus 9000 supports traditional 3-tier architectures, Spine-Leaf (CLOS) architectures, or Generation 1 Software Defined Network (SDN) solutions such as OpenFlow. In Application Centric Infrastructure (ACI) mode the Nexus 9000 provides advanced capabilities and an operational model focused on the application to accelerate business. Operating in either Cisco NX-OS Software mode or in Application Centric Infrastructure (ACI) mode, these switches are ideal for traditional or fully automated data center deployments.

The FlexPod with Cisco UCS Mini uses the Cisco NX-OS mode (standalone) of the Cisco Nexus 9300 Series access leaf switch to provide important customer benefits such as:

- 1 Gbps, 10 Gbps and 40-Gbps Ethernet support
- Programmability enabling emerging networking models, including first-generation SDN, automation, and development and operations (DevOps) solutions
- Power efficiency
- Cisco ACI capable; easing migration to the next generation data center infrastructure

The FlexPod with Cisco UCS Mini fully supports other Cisco Nexus access switching platforms. Considerations for the positioning of other Cisco Nexus family switches, Nexus 5000 or Nexus 3000 in an organization may include:

- FC and FCoE Scenarios
- DFA - FabricPath Requirements

Very specific feature needs not present or supported on the Cisco Nexus 9300 series, for example Private VLANs or NetFlow.



Note

The remainder of this document captures the design using Cisco Nexus 9000 access or leaf switch models.

Supported Operating Systems

The FlexPod with UCS Mini is a flexible infrastructure supporting hypervisor based virtualization, single operating system deployments (bare metal) or a combination of each. The operating systems supported on the FlexPod with Cisco UCS Mini infrastructure are dependent on the hardware and software interoperability matrix of the Cisco Unified Computing System and NetApp Fabric-Attached Storage controllers. Please defer to these reference resources for specific hardware and software combinations:

- Cisco UCS Manager Managed UCS Server Compatibility at:
<http://www.cisco.com/c/en/us/support/servers-unified-computing/unified-computing-system/products-technical-reference-list.html>
- NetApp Interoperability Matrix Tool at:
<http://mysupport.netapp.com/NOW/products/interoperability/>

FlexPod Data Center Models with Cisco UCS Mini

The FlexPod Data Center with UCS Mini allows organizations to address their data center needs using a standardized infrastructure with well-known or contained failure domains. This structured approach to the data center has both cost and operational efficiencies associated with it.

From an operational perspective the organization has the ability to create a FlexPod package addressing their specific needs from a compute, network and storage perspective. This structure can then be readily deployed based on FlexPod best practices and repeated as necessary to address increased demand for data center services. Employing Cisco UCS Central allows for centralized compute policy across one or multiple instances of the FlexPod supporting scale and delivering consistency.

The FlexPod Data Center Models with Cisco UCS Mini are characterized by the NetApp FAS connection points within the architecture. Currently, these models include:

- Direct Attached SAN Storage
- IP-based Storage

Direct Attached SAN Storage

Figure 5 illustrates the use of directly attached clustered FAS controllers to the Cisco UCS Mini environment. The Cisco UCS Mini supports directly attaching NetApp storage to the Cisco UCS 6324 Fabric Interconnect. This removes the requirement to have a dedicated FC switching environment as all SAN switching and zoning are performed by the Cisco UCS 6324 Fabric Interconnect and managed through the Cisco UCS Manager. The FAS controllers use a switchless cDOT deployment model through the use of 10 GbE loopback cables. SAN A and B best practices are honored and ALUA enabled host based multi-pathing allows for redundant and optimal 8-Gbps Fibre Channel paths into the NetApp controllers. This model supports only Fibre Channel access to storage. Scalability is achieved by adding storage capacity (disk/shelves) to an existing HA pair.

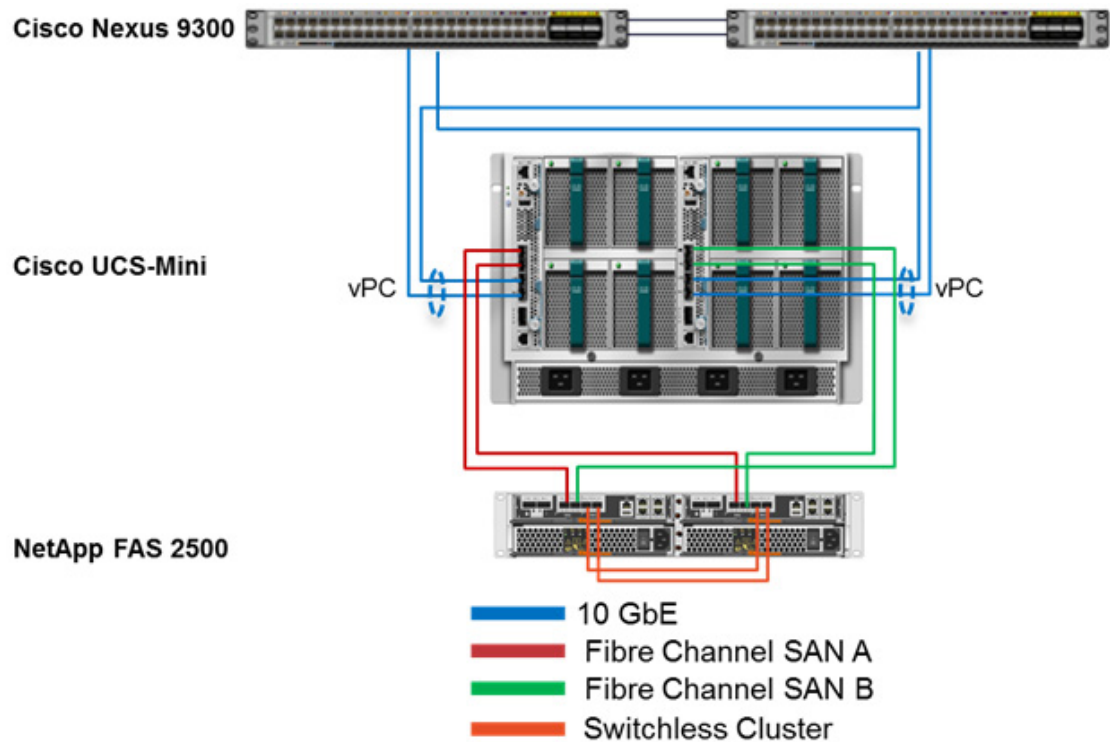
The Cisco Nexus 9300 in standalone mode (NX-OS) provides a non-blocking, redundant 10 GbE-enabled, 40 GbE capable, Ethernet fabric. The Cisco Nexus 9000 does not support any storage traffic with this FlexPod infrastructure design. Storage traffic is contained between the Cisco UCS Fabric Interconnects and FAS controllers. As shown, link aggregation technologies play an important role, providing improved aggregate bandwidth and link resiliency between the network fabric and compute domain. The 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 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.

Figure 5 FlexPod Data Center with Cisco UCS Mini - Direct Attached SAN Storage Design



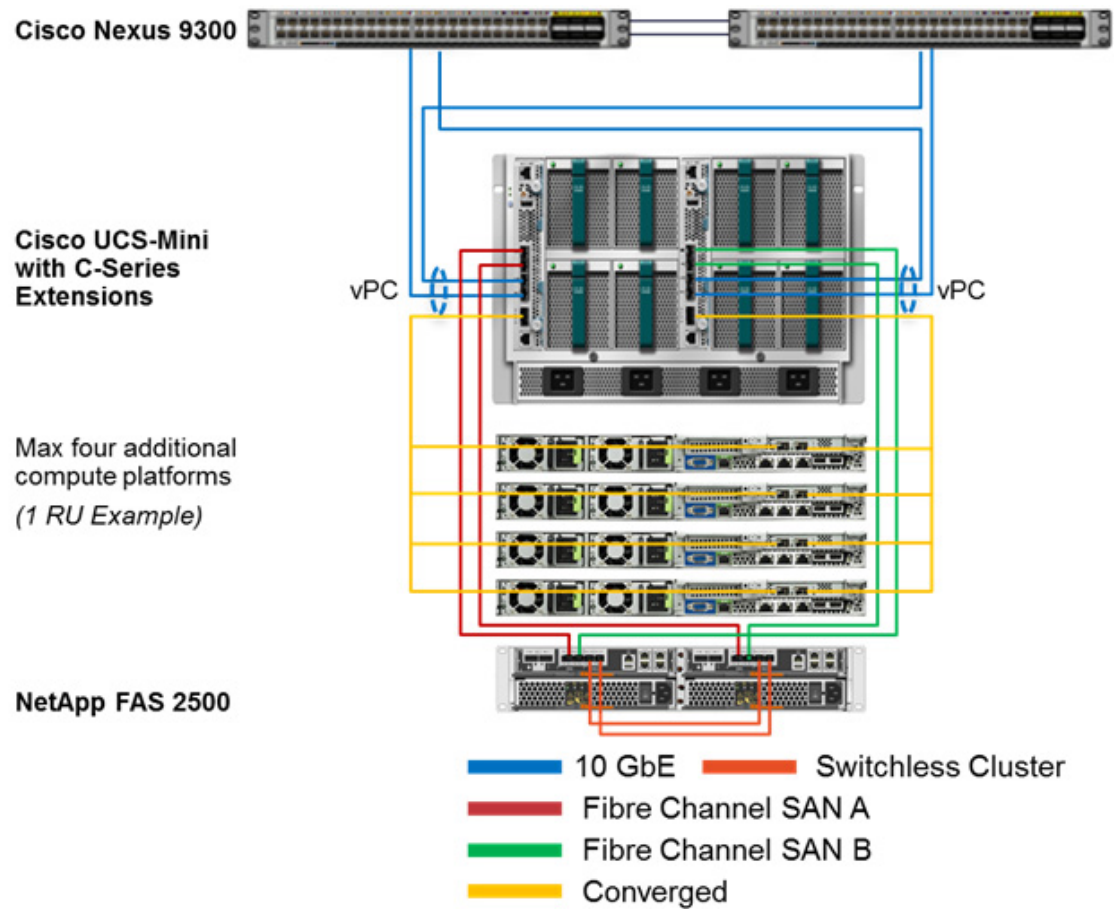
Compute Scaling

The Cisco UCS Mini platform requires the presence of one Cisco UCS 5108 blade chassis. The Cisco UCS 5108 blade chassis can support a maximum of eight half-width blades. In addition, the Cisco UCS 6324 Fabric Interconnect supports a 40 GbE FCoE Enhanced Quad SFP (QSFP+) port that provides connectivity for up to four additional Cisco UCS rack mount servers readily allowing an organization to expand the compute capacity of the FlexPod Data Center with Cisco UCS Mini system. [Figure 6](#) illustrates the maximum number of compute nodes currently supported by this architecture (twelve).

**Note**

At the time of this writing the Cisco UCS Mini supports Cisco UCS B200 M3, C220 M3 and C240 M3 server models. Please check the release notes for hardware internal dependencies.

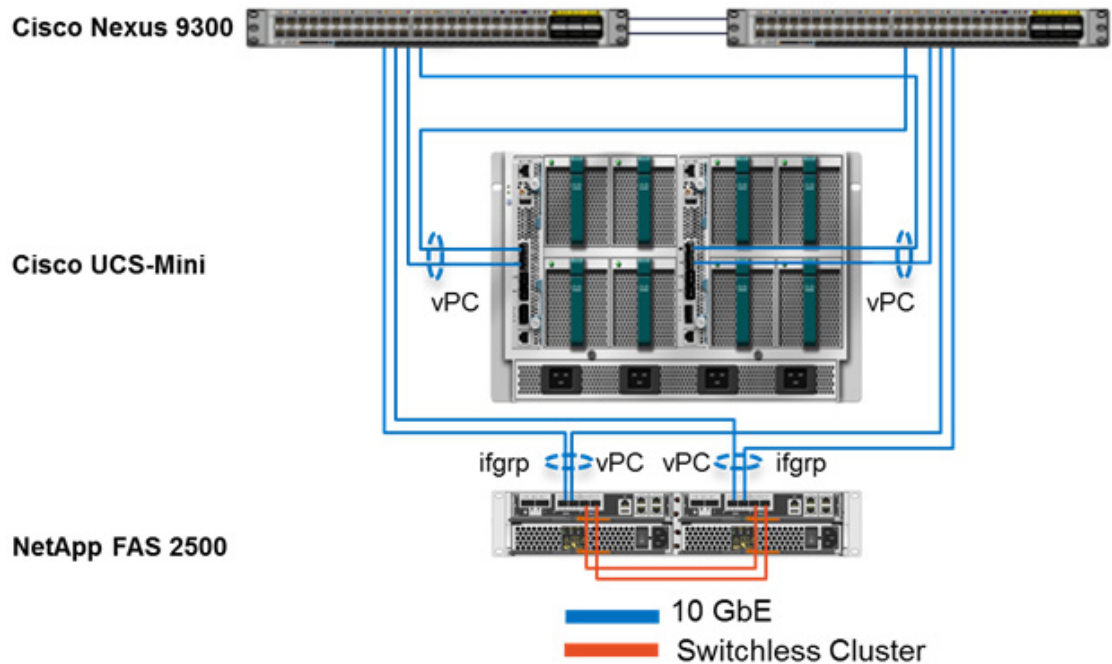
Figure 6 FlexPod Data Center with Cisco UCS Mini - Direct Attached SAN Storage Design and C-Series Extensions



IP-Based Storage

[Figure 7](#) details the FlexPod Data Center with Cisco UCS Mini IP-based storage design. This design uses the Cisco Nexus 9000 in standalone mode 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 7 FlexPod Data Center with Cisco UCS Mini - IP-Based Storage Design



The 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 Cisco 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 detailed, port aggregation provides 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.

The FlexPod Data Center with Cisco UCS Mini 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 uplink design. Figure 7 describes 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 7 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 the FAS25xx. 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.

Compute Scaling

The FlexPod Data Center with Cisco UCS Mini and IP-based storage is extensible from a compute perspective.

As noted in the previous section, the Cisco UCS Mini platform requires the presence of one Cisco UCS 5108 blade chassis. The Cisco UCS 5108 blade chassis can support a maximum of eight half-width blades. With this IP-based model there are two 10 Gbe SFP+ ports available to support two rack mount UCS C-series servers.

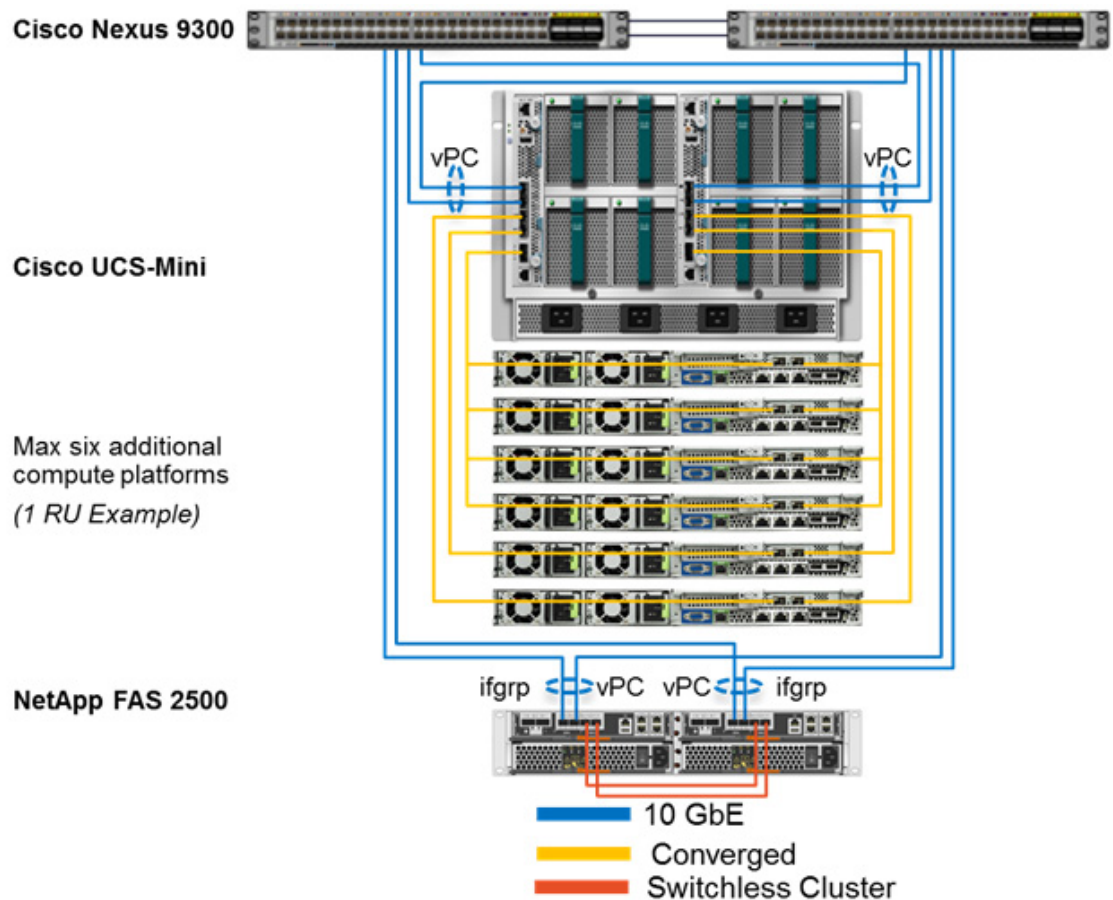
In addition, the Cisco UCS 6324 Fabric Interconnect supports a QSFP+ port that provides connectivity for up to four additional Cisco UCS rack servers; allowing an organization to expand the compute capacity of the FlexPod Data Center with Cisco UCS Mini system. [Figure 8](#) illustrates the maximum number of compute nodes currently supported by this architecture (fourteen).



Note

At the time of this writing, the Cisco UCS Mini supports Cisco UCS B200 M3, C220 M3 and C240 M3 server models. Please check the release notes for hardware internal dependencies.

Figure 8 FlexPod Data Center with Cisco UCS Mini - IP-Based Storage Design and C-Series Extensions



FlexPod Express Models with Cisco UCS Mini

The FlexPod Express with Cisco UCS Mini models were developed to address greenfield and brownfield ROBO settings. In greenfield environments, it is strongly suggested to reference the FlexPod Data Center with Cisco UCS Mini models previously defined. In fact, the organization may consider using a Nexus 3000 series switch as valid option given budgetary constraints. In existing or brown-field

deployments, the FlexPod Express with Cisco UCS Mini models assumes there is an existing switching infrastructure on site and as a result the design limits storage traffic to Cisco Unified Computing System attached SAN. There is no storage traffic traversing the existing switching infrastructure. The existing switching infrastructure support is limited to Ethernet connectivity of the Cisco UCS domain for ingress and egress traffic.



Note The FlexPod Express with Cisco UCS Manager Mini designs are validated using the Nexus 3524 switching platform



Note The FlexPod Express with Cisco UCS Mini does not dictate the Ethernet switching platform in the environment. Non-Cisco switching platforms are not defined in the FlexPod solution portfolio or for that matter supported by Cisco TAC, but when present third-party switches may provide the functionality to deliver file based storage access (NFS) or other traffic such as vMotion. This means the FlexPod Express with UCS Mini though defined with iSCSI can implement NFS based LIFs on the NetApp cDOT SVM if desired provided there is a resilient data path between the two UCS fabric interconnects. Performance is dependent on the available bandwidth of the path and switching platform.

The FlexPod Express with Cisco UCS Mini designs are validated and documented to support local and remote centralized management. These designs are eligible to take advantage of policy consolidation at HQ and delivered to distant sites.

For example, the FlexPod Express model may be deployed as a "standalone" configuration, locally managed as an independent infrastructure. The local administrator would use Cisco UCS Manager and a local virtual machine manager such as vSphere vCenter to oversee the compute domain. An alternate approach, is to deploy the FlexPod Express in a distributed fashion taking advantage of the remote management toolset and capabilities available with the infrastructure. This means centralized policy and significant operational advantages. For instance, utilize Cisco UCS Central to manage multiple Cisco UCS domains across the organizations.

Table 1 captures the tools and possible deployment locale:

Table 1 FlexPod Management Tools and Location

Local	Remote
Cisco UCS Manager	Cisco UCS Central extended from the Data Center
Cisco Nexus 1000v VSM	Cisco Nexus 1000v VSM
vSphere vCenter	vSphere vCenter
Microsoft System Center / VMM	Microsoft System Center /VMM
OnCommand System & Unified Manager, VSC	OnCommand System & Unified Manager, VSC
FAS Command line	Nexus Command Line, Python scripting
Nexus Command Line, Python scripting	NetApp Virtual Storage Console (VSC)
NetApp Virtual Storage Console (VSC)	

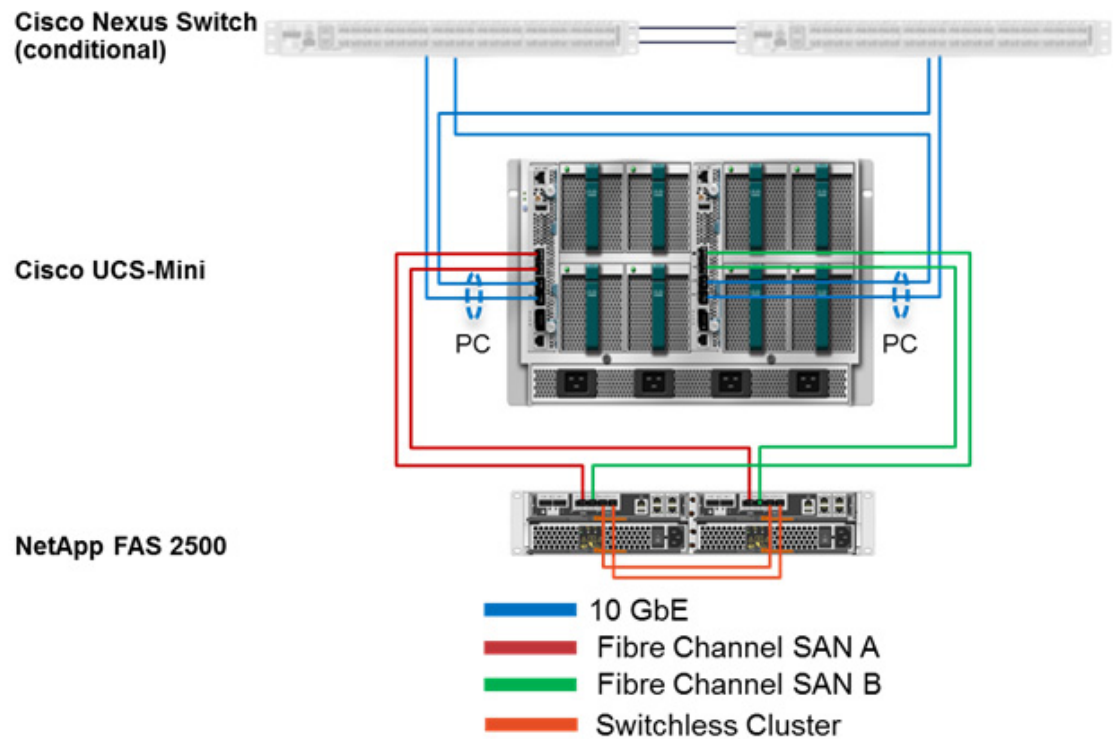
Direct Attached SAN Storage

The FlexPod Express with Direct Attached SAN storage utilizes the Cisco UCS 6324 Fibre Channel zoning and switching capabilities to directly support the FAS 2500 series controllers. In this model, Fibre Channel target and initiator communications are limited on direct connections between the Cisco UCS and NetApp FAS. SAN A and B best practices are maintained within the FlexPod demarcated

domain. [Figure 9](#) captures the architectural simplicity of the FlexPod Express.

The Ethernet switching environment is not clearly defined in this model as caveatted earlier. Regardless the Cisco UCS domain provides multiple 10 GbE links to support any link aggregation technologies available in the switching domain. Note the use of vPC is not available unless a Nexus switch is introduced.

Figure 9 FlexPod Express with Cisco UCS Mini - Direct Attached SAN Design



Compute Scaling

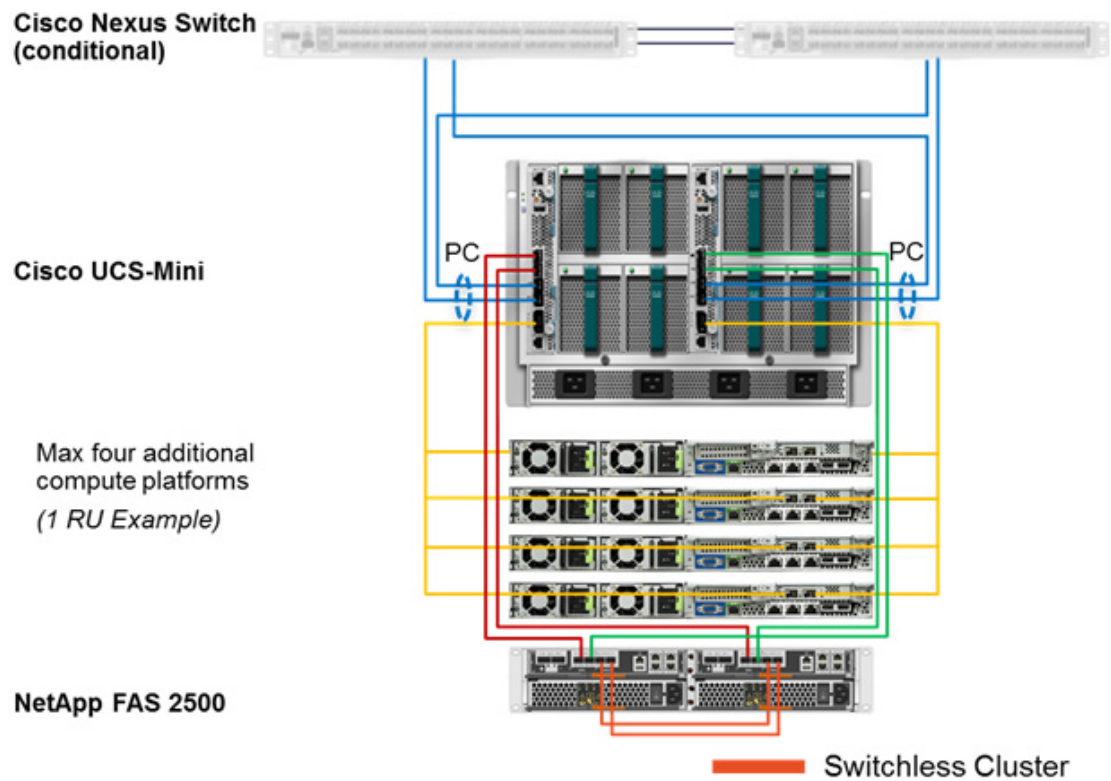
The FlexPod Express with Direct Attached SAN Storage provides compute scale-out within the Cisco UCS 5108 chassis as well as with directly attached C-series servers via the 40 GbE QSFP+. [Figure 10](#) shows an example of the scale assuming four Ethernet connections from the Cisco UCS 6324 Fabric Interconnect are designated for ingress/egress traffic. The maximum number of compute nodes in this configuration would be twelve, eight half-width Cisco UCS blades and four Cisco UCS rack mount models.



Note

At the time of this writing the Cisco UCS Mini supports Cisco UCS B200 M3, C220 M3 and C240 M3 server models. Please check the release notes for hardware internal dependencies.

Figure 10 FlexPod Express with Cisco UCS Mini - Direct Attached SAN Storage Design and C-Series Extensions



IP-Based Storage

The FlexPod Express with Cisco UCS Mini IP-Based Storage design uses iSCSI connections between the NetApp FAS controller and the Cisco UCS 6324 Fabric Interconnect to provide block-based access to data. This is an all 10 GbE design which leverages multiple FAS based iSCSI targets uses iSCSI multipathing to provide storage high availability. This design contains all iSCSI traffic to the "defined" FlexPod links between the Cisco UCS Mini Fabric Interconnect and FAS 2500 series controller as shown in [Figure 11](#). As stated earlier, the FlexPod Express with Cisco UCS Mini does not dictate the Ethernet switching platform in the environment.




Note

The ingress/egress Ethernet traffic to the Cisco UCS domain can utilize the multiple 1/10Gbe links on the Cisco UCS 6324 fabric interconnect.

The use of port channels is dependent on the existing infrastructure at the locale and the capabilities of the resident switching platforms.

Figure 11 FlexPod Express with Cisco UCS Mini - IP-Based Storage Design



Compute Scaling

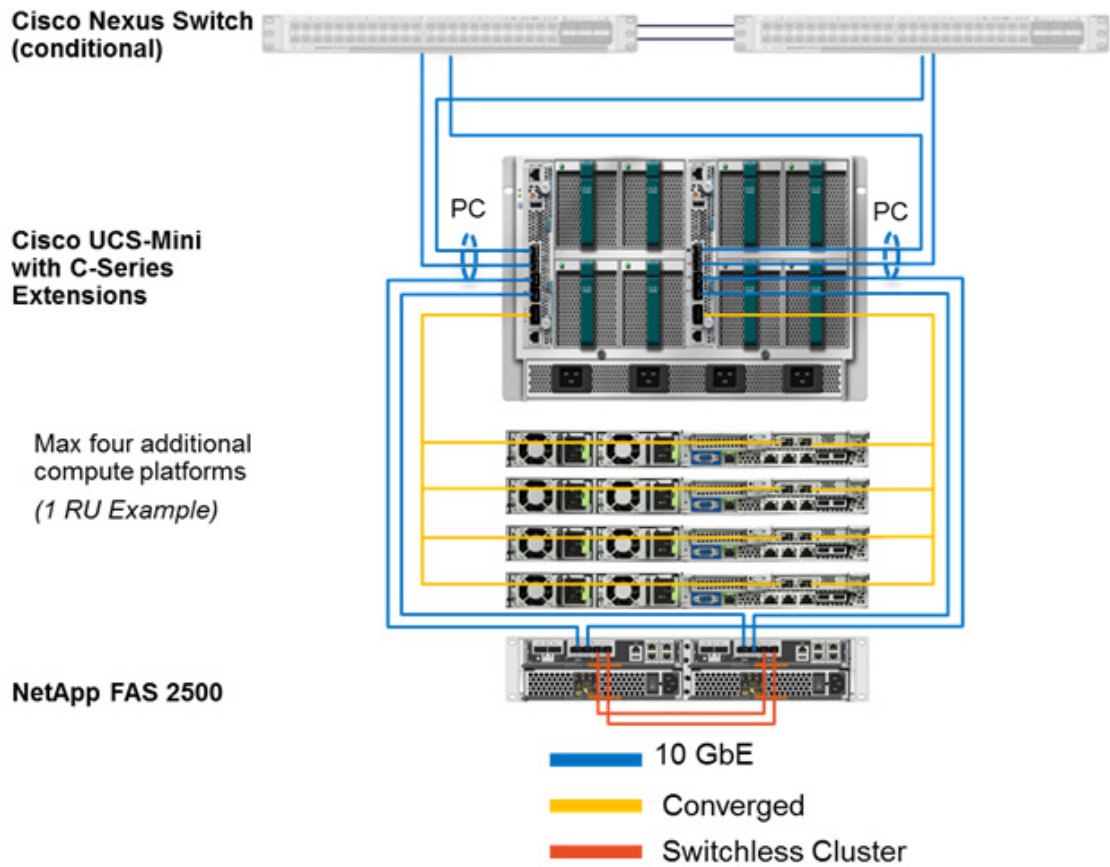
The FlexPod Express with IP-Based Storage design provides compute scale-out within the Cisco UCS 5108 chassis as well as with directly attached C-series servers via the 40 GbE QSFP+. Figure 12 illustrates the maximum compute scale achieved assuming four Ethernet connections from the Cisco UCS 6324 Fabric Interconnect are designated for ingress/egress traffic to the existing network infrastructure. The maximum number of compute nodes in this configuration would be twelve, eight half-width Cisco UCS blades and four Cisco UCS rack mount models.



Note

At the time of this writing the Cisco UCS Mini supports Cisco UCS B200 M3, C220 M3 and C240 M3 server models. Please check the release notes for hardware internal dependencies.

Figure 12 FlexPod Express with Cisco UCS Mini - IP-Based Storage Design with C-Series Extensions



Market Segment to Solution Matrices

The following tables address specific market segments and by default the use cases they encompass. These tables are intended to guide an organization toward the FlexPod with Cisco UCS Mini that may best address their specific requirements.



Note

Each deployment guide lists the hardware and software images used during the validation effort. It is important to remember that FlexPod allows organizations to use supported NetApp and Cisco images as referenced in the Cisco UCS Manager Managed UCS Server Compatibility matrix and the NetApp Interoperability Matrix tool. For more details, go to section [Supported Operating Systems](#).

Table 2 *FlexPod with Cisco UCS Mini for ROBO*

Title	Subtitle (URL)
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with IP-Based Storage (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS-Mini and VMware vSphere 5.5 with IP-Based Storage

Table 3 *FlexPod with Cisco UCS Mini for SMB*

Title	Subtitle (URL)
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5	Deployment Guide for FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with IP-Based Storage (PENDING)	Deployment Guide for FlexPod Datacenter with Cisco UCS-Mini and VMware vSphere 5.5 with IP-Based Storage
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with IP-Based Storage (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS-Mini and VMware vSphere 5.5 with IP-Based Storage

Table 4 *FlexPod with UCS Mini for Data Center*

Title	Subtitle (URL)
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5	Deployment Guide for FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with IP-Based Storage	Deployment Guide for FlexPod Datacenter with Cisco UCS-Mini and VMware vSphere 5.5 with IP-Based Storage

Table 5 *FlexPod with Cisco UCS Mini for DMZ*

Title	Subtitle (URL)
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5	Deployment Guide for FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with IP-based Storage (PENDING)	Deployment Guide for FlexPod Datacenter with Cisco UCS-mini and VMware vSphere 5.5 with IP-based Storage
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with Direct Attached SAN Storage
FlexPod Express with Cisco UCS Mini and VMware vSphere 5.5 with IP-based Storage (PENDING)	Deployment Guide for FlexPod Express with Cisco UCS-mini and VMware vSphere 5.5 with IP-based Storage

Conclusion

FlexPod with Cisco UCS Mini provides a new compelling form factor to address compute, network and storage requirements from the ROBO, the SMB, as well as the Enterprise and service provider data centers. Designed and validated to reduce deployment time, project risk, and the cost of IT the FlexPod with Cisco UCS Mini solution delivers a solid infrastructure to address a multitude of IT initiatives.