



Technical Report

FlexPod Datacenter Solutions Guide

John George, Chris Reno, Lindsey Street, NetApp
October 2015 | TR-3884

In Collaboration with Cisco

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information toward shared infrastructures and virtualized environments and eventually will move toward cloud computing to increase agility and reduce costs. This transformation appears daunting and complex because companies must address both organizational and technical resistance to this new IT model. Cisco and NetApp developed the FlexPod[®] platform to address these virtualization needs and to simplify the evolution toward shared, virtualized, and cloud infrastructures.

TABLE OF CONTENTS

1	Introduction.....	3
1.1	FlexPod Design.....	3
1.2	Target Audience.....	5
2	Business Benefits.....	5
2.1	Repeatable.....	5
2.2	Flexible.....	5
2.3	Right-Sized and Scalable.....	5
2.4	Lower Total Cost of Ownership.....	6
3	FlexPod Design Elements.....	6
3.1	Cisco Unified Computing System.....	6
3.2	Cisco Nexus 5000, 6000, 7000, and 9000 Series Switches.....	9
3.3	NetApp Storage Infrastructure	12
3.4	Management.....	15
4	Data Center Solutions	16
4.1	FlexPod Infrastructure.....	16
4.2	FlexPod Application Solutions.....	23
4.3	FlexPod Datacenter Integration	26
5	Conclusion	30
	References.....	31

LIST OF FIGURES

Figure 1)	FlexPod layout.	4
Figure 2)	Example of scaling a FlexPod configuration.	5
Figure 3)	VASA provider collects and delivers data to vCenter Server.	16
Figure 4)	FlexPod solutions and applications.	17
Figure 5)	Base FCoE FlexPod infrastructure.....	18
Figure 6)	iSCSI boot FlexPod infrastructure.....	19
Figure 7)	Secure separation architecture overview.	21
Figure 8)	SMT architectural components (graphic supplied by Cisco).	27
Figure 9)	VMware SRM with NetApp storage.....	30

1 Introduction

Industry trends indicate a vast data center transformation toward shared infrastructures. Enterprise customers are moving away from silos of information toward shared infrastructures and virtualized environments and eventually will move toward cloud computing to increase agility and reduce costs. This transformation appears daunting and complex because companies must address both organizational and technical resistance to this new IT model. Cisco and NetApp developed FlexPod to address these virtualization needs and to simplify the evolution toward shared, virtualized, and cloud infrastructures.

FlexPod is a validated infrastructure solution built on the Cisco Unified Computing System (Cisco UCS), Cisco Nexus data center switches, NetApp® FAS storage components, and software from a range of partners. FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. A FlexPod design provides a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different business solutions. This document describes how to build several different solutions on top of FlexPod.

Cisco and NetApp developed FlexPod as a platform that can address current data center needs and simplify the evolution toward an IT-as-a-service (ITaaS) infrastructure.

FlexPod serves as a base infrastructure layer for a variety of IT solutions. You can find a number of recently published FlexPod validated designs at

www.netapp.com/us/solutions/flexpod/datacenter/validated-designs.aspx.

Authorized partners can access the “FlexPod Implementation Guide” at the NetApp [Field Portal](#).


1.1 FlexPod Design

Cisco and NetApp have provided documentation for best practices and the deployment collateral necessary to build the FlexPod shared infrastructure stack. As part of the FlexPod offering, Cisco and NetApp designed a reference architecture with a technical specifications sheet that is highly modular or podlike. Although each customer’s FlexPod unit can vary in its exact configuration to meet the customer’s initial use case, after a FlexPod unit is built, it can easily be scaled as requirements and demand change. This capability includes scaling both up (adding resources within a FlexPod unit) and out (adding FlexPod units).

Specifically, FlexPod is a defined set of hardware and software products that serves as a foundation for data center deployments. FlexPod includes NetApp storage, Cisco networking, and the Cisco Unified Computing System (UCS) in a single deployment. The solution can be scaled while maintaining its integrity by augmenting a single FlexPod instance to provide the appropriate network, compute, or storage capacity needed within a single pod. It can also be scaled using the port density of the Cisco Nexus networking platform to readily accommodate multiple FlexPod instances.

In either case, the flexibility of the pod construct enables numerous solutions to be built on top of one or more FlexPod configurations. This capability provides enterprise flexibility, supportability, and manageability for both virtualized and nonvirtualized environments. Furthermore, you can leverage existing equipment to complete a given FlexPod instance. Figure 1 shows a base FlexPod configuration with hardware families and solution benefits.

Figure 1) FlexPod layout.



The base hardware, which is detailed in the FlexPod technical specifications, includes:

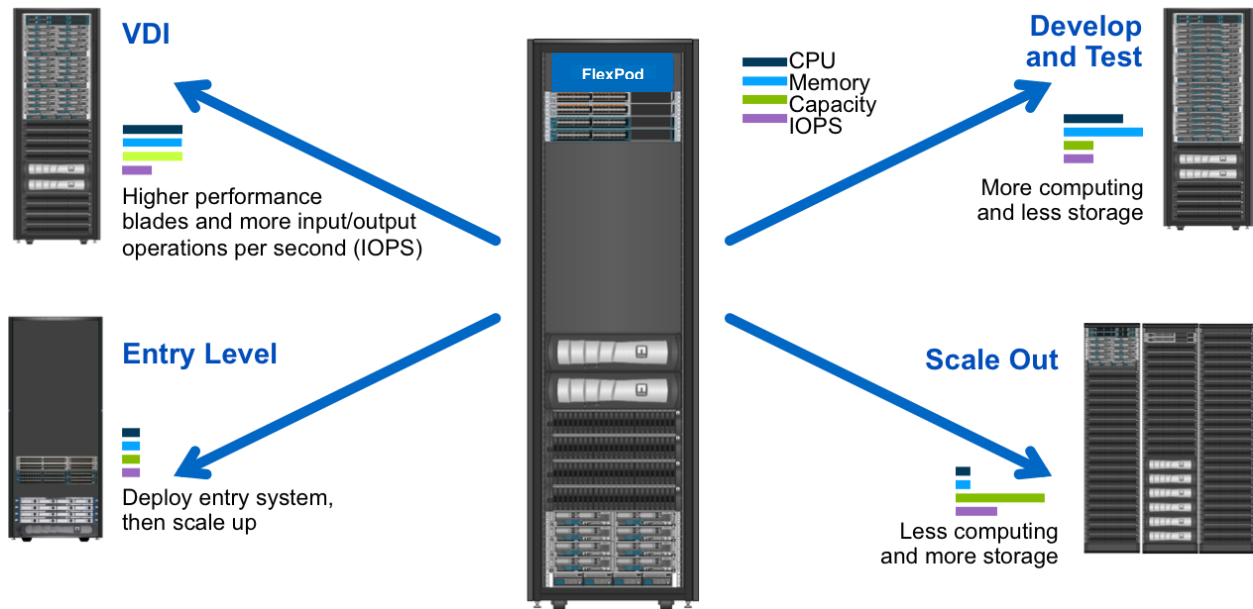
- Two Cisco Nexus switches
- Two Cisco UCS 6248UP fabric interconnects
- Cisco UCS B-Series blades with two fabric extenders per chassis and/or Cisco UCS C-Series rack servers
- NetApp FAS80xx (highly available configuration with dual controllers in a single enclosure)

In Figure 1, storage is provided by a NetApp FAS3250 with accompanying disk shelves. All systems and fabric links feature redundancy, providing end-to-end high availability (HA). Although this is the default base design, each of the components can be scaled flexibly to support the customer's specific business requirements. For example, more (or different) UCS blades and chassis can be deployed to increase compute capacity. Additional disk shelves can be deployed to increase capacity and improve input/output (I/O) throughput. And special hardware or software features can be added to introduce new features, such as NetApp Flash Cache™ intelligent caching for deduplication-aware caching.

Note: FlexPod infrastructures operating with the NetApp clustered Data ONTAP® operating system require a cluster interconnect. This interconnect is one that can be deployed as a switchless cluster for two nodes or with a pair of cluster interconnects for clusters with more than two nodes.

FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments. Figure 2 shows a few FlexPod scaling options. FlexPod is a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different use cases.

Figure 2) Example of scaling a FlexPod configuration.



1.2 Target Audience

This document describes the basic architecture of FlexPod as well as several solutions that can be built on FlexPod. The target audience for this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineering personnel, and customers who want to deploy solutions on top of a FlexPod configuration.

2 Business Benefits

FlexPod solutions offer a variety of business benefits to the consumer. This section describes some of these benefits at a high level.

2.1 Repeatable

The FlexPod architecture and accompanying collateral, delivered by Cisco and NetApp, aid customers, partners, and field personnel with new FlexPod deployments. This configuration provides documentation about sample workloads, required port count, power consumption, and required cooling. For customers who want standard scalable configurations, FlexPod is an excellent choice for their infrastructures because environmental requirements are defined.

2.2 Flexible

One of the founding principles of the FlexPod program is flexibility. The reference architecture serves as a starting point for a variety of workloads but it is not mandated as the only way to bundle FlexPod products. Customers can modify server type and count, storage type and scale, and network switches, for example, based on their specific needs.

2.3 Right-Sized and Scalable

As previously mentioned, the reference architecture can be scaled and right-sized based on a customer's use case. You can achieve proper sizing by using appropriate Cisco and NetApp tools, such as:

- Storage Performance Modeler (SPM)

- Cisco UCS Application Sizer
- Cisco RESCOMP
- FlexPod Guided Solution Sizing (GSS)

Depending on their application requirements, customers can choose the appropriate hardware for their deployments by working with their subject matter experts. Doing so enables customers to avoid purchasing an infrastructure stack that is underused.

Note: The sizing data found in FlexPod GSS covers typical enterprise workloads. However, individual customers might have lighter or heavier workloads related to the greater headroom required for growth or during system failure. Take these factors into consideration, as well as the application profiles, when sizing for individual customers.

2.4 Lower Total Cost of Ownership

Because the FlexPod architecture is repeatable, flexible, and right-sized to application requirements, the purchase cost, deployment time, and provisioning time are all lowered. This standardized approach to the data center also decreases training time for staff and increases their ability to work with tenants' or different business units' equipment. These benefits are substantial for the customer and, when combined with NetApp storage efficiency and Cisco UCS stateless computing, customers can lower the total cost of ownership (TCO) of the infrastructure.

3 FlexPod Design Elements

This section describes the elements that compose a FlexPod architecture. Because FlexPod is not a rigid configuration, the customer can build an infrastructure that includes best-in-class technologies from Cisco and NetApp.

3.1 Cisco Unified Computing System

The Cisco Unified Computing System (UCS) is a next-generation data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce TCO and increase business agility. This system integrates a low-latency, lossless, 10-Gigabit Ethernet (10GbE), unified network fabric with enterprise-class, x86-architecture servers. This system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain. A Cisco UCS deployment consists of Cisco UCS fabric interconnects, blade server chassis, B-Series blades, C-Series rack-mount servers, and adapter connectivity.

For more information, refer to the following documentation:

- Cisco UCS white papers:
www.cisco.com/en/US/netsol/ns944/networking_solutions_white_papers_list.html
- Unified computing and servers:
www.cisco.com/en/US/products/ps10265/index.html

Cisco UCS Director

Cisco UCS Director can be 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, enabling 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 Data 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 the Lightweight Directory Access Protocol. The drag-and-drop workflow designer tool eliminates the need for service engagements and the need to bring together multiproduct solutions or third-party adapters.

For more information about Cisco UCS Director, refer to www.cisco.com/c/en/us/products/servers-unified-computing/ucs-director/index.html.

Cisco UCS Manager

Cisco UCS Manager (Cisco UCSM) is robust device management software that is embedded in every Cisco UCS deployment and can support up to 160 servers. Cisco UCSM allows management of the entire compute environment from a highly available single pane of glass. Cisco UCS Manager provides flexible role-based and policy-based management of resources by using service profiles and templates. The use of service profiles and templates abstracts those elements that typically denote server personality (such as MAC addresses, WWPN, system UUID, and firmware revisions) from the role that the compute resource serves in the data center.

Through this methodology, firmware, boot order, NIC and HBA settings, boot targets, and so on are no longer tied to a specific piece of server hardware. This approach enables IT infrastructure to be deployed in minutes rather than days and allows organizations to focus on strategy rather than on monotonous maintenance tasks.

With respect to element management, Cisco UCSM provides:

- Device discovery
- Firmware management
- Inventory
- Kernel-based virtual machine (VM) console access
- Diagnostics
- Quality of service (QoS)
- Monitoring
- Statistics collection
- Auditing

Cisco UCS Fabric Interconnect

Cisco UCS 6100 and 6200 Series fabric interconnects are a core part of the Cisco Unified Computing System, providing network connectivity and management capabilities for the system, because this device is where the Cisco UCSM resides. Deployed as a clustered pair, Cisco UCS 6100 or 6200 Series fabric interconnects offer line-rate, low-latency, lossless 10GbE and Fibre Channel over Ethernet (FCoE) functions.

Cisco UCS 6100/6200 Series fabric interconnects are the management and communication backbone for Cisco UCS B-Series blade servers, Cisco UCS 5100 Series blade server chassis, and C-Series rack-mount servers. All chassis, blades, and rack mounts attached to Cisco UCS 6100/6200 Series fabric interconnects become part of a single, highly available management domain. In addition, by supporting a unified fabric, Cisco UCS 6100/6200 Series fabric interconnects provide both LAN and storage area network (SAN) connectivity for all blades within their domain.

Note: The Cisco UCS 6200UP fabric interconnect supports unified ports. Therefore, ports can be configured to support 10GbE, Fibre Channel (FC), or FCoE.

Cisco UCS B-Series Blades and Blade Server Chassis

The Cisco UCS 5100 Series blade server chassis is a crucial building block for the Cisco Unified Computing System (UCS), delivering a scalable and flexible blade server chassis for today's and tomorrow's data center while helping reduce TCO.

Cisco's first blade server chassis offering, the Cisco UCS 5108 blade server chassis, is six rack units high and can be mounted in an industry-standard 19" rack. A chassis can house up to eight half-width Cisco UCS B-Series blade servers and can accommodate both half-width and full-width blade form factors. There are currently six blade server options from which to choose:

- Cisco UCS B200 M4 blade server
- Cisco UCS B260 M4 blade server
- Cisco UCS B420 M4 blade server
- Cisco UCS B22 M3 blade server
- Cisco UCS B200 M3 blade server
- Cisco UCS B420 M3 blade server

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92% efficient and can be configured to support nonredundant, N+1-redundant, and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two input/output (I/O) bays for either the Cisco UCS 2104XP or 2208XP fabric extenders.

A passive midplane provides up to 20GBps of I/O bandwidth per half-width server slot and up to 40GBps of I/O bandwidth per full-width server slot. The chassis can support future 40GbE standards.

Cisco UCS C-Series Rack-Mount Servers

Cisco UCS C-Series rack-mount servers extend unified computing innovations to an industry-standard form factor to help reduce TCO and increase business agility. The innovations embodied in this series include a standards-based unified network fabric, Cisco VN-Link virtualization support, and Cisco Extended Memory Technology. The servers support an incremental deployment model and protect customer investment with a migration path to unified computing.

There are currently 10 rack-mount server options from which to choose:

- Cisco UCS C220 M4 rack server
- Cisco UCS C240 M4 rack server
- Cisco UCS C460 M4 rack server
- Cisco UCS C22 M3 rack server
- Cisco UCS C220 M3 rack server
- Cisco UCS C24 M3 rack server
- Cisco UCS C240 M3 rack server
- Cisco UCS C260 M2 rack server

- Cisco UCS C420 M3 rack server
- Cisco UCS C460 M2 rack server

Each server model addresses varying workload challenges through a balance of processing, memory, I/O, and internal storage resources.

For more information, refer to the Cisco UCS C-Series rack servers documentation at www.cisco.com/en/US/products/ps10493/index.html.

Cisco UCS Virtual Interface Card

Cisco virtual interface cards (VICs) were developed from the ground up to accelerate the various new operational modes introduced by server virtualization. VICs are highly configurable, self-virtualized adapters that can create up to 128 PCI Express (PCIe) endpoints per adapter. These PCIe endpoints are created in the adapter firmware and present fully compliant standard PCIe topology to the host OS or hypervisor. Versions of the VIC are available for both B-Series blades and C-Series rack-mount servers.

Each of the PCIe endpoints created by a VIC can be configured individually for the following attributes:

- **Interface type.** FCoE, Ethernet, or Dynamic Ethernet interface device
- **Resource maps that are presented to the host.** PCIe base address registers, interrupt arrays
- **Network presence and attributes.** Maximum transmission unit, VLAN membership
- **Quality-of-service parameters.** 802.1p class, enhanced transmission selection attributes, rate limiting, and shaping

3.2 Cisco Nexus 5000, 6000, 7000, and 9000 Series Switches

The networking foundation for any FlexPod deployment is the Cisco Nexus family of switches. The Cisco Nexus switch family enables a variety of transport protocols and leverages the NX-OS operating system.

Cisco Nexus 5000 Series Switches

Cisco Nexus 5000 Series switches are ideal for enterprise-class data-center server access-layer and smaller-scale, midmarket data-center aggregation-layer deployments. You can deploy these multipurpose, multilayer switches across a diverse set of traditional, virtualized, unified, and high-performance computing (HPC) environments.

As part of the FlexPod foundation, these switches help address the business, application, and operational requirements of evolving data centers. The switches provide:

- Architectural flexibility to support diverse business and application needs
- Infrastructure simplicity to decrease TCO
- Agility for traditional deployments, with easy migration to virtualized, unified, or HPC environments
- Enhanced business resilience with greater operational continuity based on the modular Cisco NX-OS operating system
- Ability to use existing operational models and administrative domains for easy deployment
- Nonblocking, resilient architecture based on Virtual PortChannel, which removes dependency on the spanning-tree protocol

The Cisco Nexus 5000 Series also expands Cisco Fabric Extender (FEX) technology into servers and VMs with Cisco Adapter FEX and Cisco VM-FEX.

This switch family's rich feature set makes the series ideal for rack-level, access-layer applications. The switches protect investments in data center racks with standards-based Ethernet and FCoE features that allow IT departments to consolidate networks based on their own requirements and timing.

The combination of high port density, wire-speed performance, and extremely low latency makes the switch an ideal product to meet the growing demand for 10GbE at the rack level. The switch family has sufficient port density to support single or multiple racks that are fully populated with blade and rack-mount servers.

Built for today's data centers, the switches are designed just like the servers they support. Ports and power connections are at the rear, closer to server ports, helping keep cable lengths as short and efficient as possible. Hot-swappable power and cooling modules can be accessed from the front panel, where status lights offer an at-a-glance view of switch operation. Front-to-back cooling is consistent with server designs, supporting efficient data center hot-aisle and cold-aisle designs. Serviceability is enhanced with all customer-replaceable units accessible from the front panel. The use of Small-Form-Factor Pluggable Plus ports offers increased flexibility to use a range of interconnect solutions, including copper for short runs and fiber for long runs.

FCoE and IEEE Data Center Bridging features support I/O consolidation, simplify management of multiple traffic flows, and optimize performance. Although implementing SAN consolidation requires only the lossless fabric provided by the Ethernet pause mechanism, the Cisco Nexus 5500 Series offers additional features that create an even more easily managed, high-performance unified network fabric.

For more information, refer to the Cisco Nexus 5500 Series switches documentation at Cisco [Nexus 5000 Series switches](#).

Cisco Nexus 6000 Series Switches

The new Cisco Nexus 6000 Series brings high-density 10GbE and 40GbE in energy-efficient, compact-form-factor switches. With a robust, integrated layer 2 and layer 3 feature set, the Cisco Nexus 6000 Series provides a versatile platform that can be deployed in various scenarios. These scenarios include direct-attached 10GbE and 40GbE access and high-density Cisco Fabric Extender aggregation deployments, spine-and-leaf architectures, and compact aggregation deployments to build scalable Cisco Unified Fabric in the data center. Cisco Nexus 6000 Series architectures can adapt to increasing bandwidth demands with low power and a compact space profile, providing capital expenditure (capex) and operating expense (opex) savings.

Cisco Nexus 6000 Series products use the same set of Cisco application-specific integrated circuits (ASICs) and a single software image across all of the products in the family, offering feature consistency and operational simplicity. Cisco Nexus 6000 Series switches support the following features:

- Robust layer 2 and layer 3 functions
- Industry-leading fabric extender architecture with Cisco Nexus 2200 platform fabric extenders
- The Cisco Nexus B22 blade fabric extender
- The in-service software upgrade
- Cisco FabricPath

Operation efficiency and programmability are enhanced on the Cisco Nexus 6000 Series through advanced analytics, PowerOn Auto Provisioning, and Python and Tool Command Language scripts.

The Cisco Nexus 6001 is an important component of the Cisco Unified Data Center architecture, complementing existing Cisco Nexus Family switches. The Cisco Nexus 6001 extends the industry-leading innovations and versatility of the Cisco Nexus 5000 Series switches, which are purpose-built 10GbE data center-class switches. The Cisco Nexus 6001 is an energy-efficient switch that offers high port density in one rack unit (1RU). The switch supports 10GbE, 40GbE, and FCoE and provides integrated layer 2 and layer 3 features at wire speed. The switch also offers low latency of approximately 1ms for any packet size. With a choice of front-to-back (port-side exhaust) and back-to-front (port-side intake) airflow options, the switch is designed for a broad range of traditional data center and large-scale virtualized cloud deployments.

For more information, refer to Cisco Nexus 6000 [Series](#) switches.

Cisco Nexus 7000 Series Switches

The Cisco Nexus 7000 Series offers a comprehensive, one-platform solution for the data center core network. The series also provides aggregation, high density, and end-of-row and top-of-rack server connectivity. For campus core deployments, the switches offer a scalable, highly resilient, high-performance solution.

The Cisco Nexus 7000 Series platform runs on [Cisco NX-OS software](#). It was designed for the most mission-critical deployments in the [data center](#) and on campus.

The Cisco Nexus 7000 Series was designed around three principles:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, high density, and performance all support efficient data center infrastructure growth.
- **Operational continuity.** The Cisco Nexus design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** You can incrementally and cost-effectively adopt new networking innovations and technologies.

New supported technologies include:

- [Cisco Overlay Transport Virtualization \(OTV\)](#)
- [Cisco FabricPath](#)
- [Fibre Channel over Ethernet \(FCoE\)](#)
- [Cisco Locator/ID Separation Protocol \(LISP\)](#)

Cisco Nexus 9000 Series Switches

The Cisco Nexus 9000 Series switches offer both modular and fixed 10-, 40-, and 100-Gigabit Ethernet switch configurations with scalability up to 30TBps of nonblocking performance with less than 5ms latency. The switches also offer 1,152 10GBps or 288 40GBps nonblocking layer 2 and layer 3 Ethernet ports and wire-speed VXLAN gateway, bridging, and routing support.

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

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

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. When Cisco ACI launches later this year, the ACI solution will be composed of Cisco Nexus 9000 Series

switches running in ACI fabric mode, a Cisco Application Policy Infrastructure Controller (APIC), and a Cisco Application Virtual Switch (AVS).

For more information, refer to the Cisco Nexus 5500 Series switches documentation at Cisco [Nexus 9000 Series switches](#).

3.3 NetApp Storage Infrastructure

The NetApp Unified Storage Architecture provides customers with an agile and a scalable storage platform. All NetApp storage systems use the NetApp Data ONTAP operating system. Both Data ONTAP 7-Mode and clustered Data ONTAP are supported in FlexPod. Both operating systems provide SAN (FCoE, FC, iSCSI), NAS (CIFS, NFS), and primary and secondary storage in a single unified platform. In that way, all virtual desktop data components can be hosted on the same storage array.

A single process for activities such as installation, provisioning, mirroring, backup, and upgrading is used throughout the product line, from entry-level to enterprise-class controllers. Having a single set of software and processes makes it simple to handle even the most complex enterprise data-management challenges. Unifying storage and data-management software and processes reduces the complexity of data ownership, enables companies to adapt to their changing business needs nondisruptively, and results in reduced TCO.

In a shared infrastructure, the availability and performance of the storage infrastructure are critical because storage outages and performance issues can affect thousands of users. The storage architecture must provide a high level of availability and performance. For detailed documentation about best practices, NetApp and its technology partners developed best-practice documents.

For more information, refer to the following documentation:

- [NetApp Data Storage Systems](#)
- [TR-3437: Storage Subsystem Resiliency Guide](#)
- [TR-3450: High-Availability Pair Controller Configuration Overview and Best Practices](#)
- [TR-3633: Best Practices for Oracle Databases on NetApp Storage](#)
- [TR-3702: NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V](#)
- [TR-3705: NetApp and VMware View Solution Guide](#)
- [TR-3749: NetApp Storage Best Practices for VMware vSphere](#)
- [TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010](#)
- [TR-4068: VMware vSphere 5 on NetApp Clustered Data ONTAP 8.1](#)
- [TR-4080: Best Practices for Scalable SAN Data ONTAP 8.3.1](#)
- [NVA-0005: VMware vSphere on FlexPod - Validated with Clustered Data ONTAP and Data ONTAP Operating in 7-Mode](#)

NetApp Unified Storage Architecture

When it comes to storage platforms, NetApp views simplicity as the key to delivering value to customers. That is why every NetApp storage controller is available with the Data ONTAP operating system. Data ONTAP is the foundation for providing unified storage through multiprotocol support on every platform. The NetApp family of controllers consists of platforms that can serve many business needs. The FlexPod program leverages the FAS3xxx and FAS6xxx product lines to build efficient infrastructures.

RAID DP

With any shared infrastructure deployment, data protection is critical, because a RAID failure could disconnect hundreds to thousands of end users from their virtual desktops or applications and result in lost productivity. NetApp RAID DP[®] technology offers performance that is comparable to RAID 10, but it

requires fewer disks to achieve equivalent protection. RAID-DP protects against double-disk failure, as opposed to RAID 5, which protects against only one disk failure per RAID group. For more information about RAID DP, refer to [TR-3298: RAID DP: NetApp Implementation of Double-Parity RAID for Data Protection](#).

Storage Efficiency

One of the critical barriers to adopting a shared infrastructure is the increased cost of using shared storage to obtain highly available enterprise-quality infrastructure. Virtual desktop and other enterprise deployments can create a high level of data redundancy, especially for VM OS data. With traditional storage, the total storage required equals the sum of the storage required by each VM. For example, if each VM is 20GB in size and there are 1,000 VMs in the solution, at least 20TB of usable data would be required on the shared storage.

Thin provisioning, data deduplication, and NetApp FlexClone® thin-cloning technology are the critical components of the NetApp solution. These technologies offer multiple levels of storage efficiency across the virtual desktop OS data, installed applications, and user data. This capability helps customers save 50% to 90% of the cost associated with shared storage (based on existing customer deployments and NetApp solutions lab validation).

Flexible Volumes and Aggregates

Aggregates are the NetApp virtualization layer that abstracts physical disks from logical datasets, also known as flexible volumes (also known as NetApp FlexVol® volumes). Aggregates offer storage administrators pooled disk resources, including input/output operations per second (IOPS). The FlexVol volumes contained in each aggregate have their own unique logical capacity. FlexVol volumes can be thin provisioned and the logical capacity resized as needed by the storage administrator.

FlexVol volumes are shared as file-level (NFS or CIFS) mount points or are further allocated as logical unit numbers (LUNs) for block-level (iSCSI or FCP) access. FlexVol volumes can be readily offered to the VMware environment as datastores; NetApp recommends a one-to-one alignment between the FlexVol volume and VMware ESX datastore assets. Using VMware as an example, alignment makes mapping easy between the VMware server administrator's view and the storage administrator's view of the virtual environment.

NetApp also suggests large aggregates to support most VMware environments. VMware environments often have random I/O requirements. Large aggregates provide maximum flexibility to VMware administrators because a large pool of I/O resources is made readily available. With NetApp's inherent storage virtualization techniques that apply to all hypervisors, all datasets or VMs housed in a shared storage infrastructure take advantage of RAID DP from a performance and protection standpoint.

Thin Provisioning

Thin provisioning is a way of logically presenting more storage to hosts than is physically available. With thin provisioning, the storage administrator can access a pool of physical disks (known as an aggregate) to create logical volumes for different applications to use. The administrator can do this while not preallocating space to those volumes. The space is allocated only when the host needs it. The unused aggregate space is available for the existing thin-provisioned volumes to expand or for use in the creation of new volumes. For more information about thin provisioning, refer to [TR-3563: NetApp Thin Provisioning Increases Storage Utilization with on-Demand Allocation](#).

NetApp Deduplication

NetApp deduplication saves space on primary storage by removing redundant copies of blocks in a volume that hosts hundreds of virtual desktops. This process is transparent to the application and user and can be enabled and disabled dynamically. In a virtual desktop infrastructure (VDI) environment, deduplication provides significant space savings, given that each VM has an identical copy of the OS,

applications, and patches. The savings are also achieved for user data hosted on CIFS home directories. For more information about NetApp deduplication, refer to [TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide](#).

Using NetApp deduplication and file FlexClone technology can reduce the overall storage footprint of VDI desktops and improve performance by leveraging transparent storage cache sharing. Data that is deduplicated (or nonduplicated, in the case of file FlexClone data) on disk exists in the storage array cache only once per volume. All subsequent reads from any of the VM disks of a block that is already in cache are read from cache and not from disk, thus improving performance tenfold. Any nondeduplicated data that is not in cache must be read from disk. Data that is deduplicated but does not have as many block references as a heavily deduplicated VM disk (a VMDK) appears in cache only once. However, based on the frequency of access, that data might be evicted earlier than data that has many references or that is heavily used.

NetApp recommends the following deduplication guidelines:

- Deduplication is configured and operates on the FlexVol volumes only.
- Data can be deduplicated up to 255:1 without consuming additional space.
- Each storage platform has different deduplication limits.
- Each volume has dense and nondense size limits.

Transparent Storage Cache Sharing

Transparent storage cache sharing (TSCS) enables customers to benefit from the storage efficiency of NetApp and at the same time significantly increase I/O performance. TSCS is built natively into the Data ONTAP operating system. It works by leveraging block-sharing technologies, such as NetApp primary storage deduplication and FlexClone technology, to reduce the amount of cache required and eliminate duplicate disk reads. Only one instance of any duplicate block is read into cache, thus requiring less cache than that of traditional storage solutions.

VDI implementations can see up to 99% initial space savings (validated in the NetApp solutions lab) by using space-efficient NetApp cloning technologies. This capability translates into higher cache deduplication and high cache hit rates. TSCS is especially effective in addressing the simultaneous system boot, or boot storm, of hundreds to thousands of virtual desktop systems that can overload a traditional legacy storage system.

Some important issues to consider for transparent storage cache sharing include:

- Data ONTAP 7.2.5.1, 7.3P1, or later is required.
- Both `a_sis` and a NetApp NearStore[®] appliance must be licensed for deduplication to work.
- Deduplication should be run before creating NetApp Snapshot[®] copies or running NetApp SnapMirror[®] or NetApp SnapVault[®] updates.

For more information about deduplication, refer to [NetApp TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide](#).

All Flash FAS, Flash Cache, and FlexScale

The NetApp All Flash FAS storage system brings together the exceptional performance characteristics of an all-flash array and the rich feature set and proven enterprise-grade reliability of the NetApp Data ONTAP operating system. This storage system enables customers to run workloads that require high performance at low latencies.

NetApp Flash Cache[™] intelligent caching is a hardware device that extends native Data ONTAP TSCS capabilities. Flash Cache increases the amount of available cache and helps to reduce virtual desktop storm activities. For more information about Flash Cache technology, refer to the [NetApp Flash Cache 2 Technical Specifications](#).

NetApp FlexScale™ software is the tunable software component of Flash Cache. It is a licensed feature of Data ONTAP 7.3 and later. FlexScale enables three different caching modes to be used, based on the type of workload:

- Metadata only
- Normal user data
- Low-priority blocks

For example, FlexScale enables system administrators to tune the VMware View environments of their NetApp controllers.

For more information about All Flash FAS, refer to the following documentation:

- [NetApp All Flash FAS Datasheet](#)
- [FlexPod Datacenter with VMware vSphere 5.5 Update 1 and NetApp All Flash FAS Design Guide](#)
- [FlexPod Datacenter with VMware vSphere 5.5 Update 1 and NetApp All Flash FAS Deployment Guide](#)

NetApp Virtual Storage Console and Operations Manager

Implementation and management complexities associated with deploying a VDI solution are another potential barrier to VDI adoption. To provide operationally agile management of storage on the VMware vSphere platform, the NetApp Virtual Storage Console integrates directly into VMware vCenter to rapidly provision, manage, configure, and back up a VDI implementation.

NetApp Operations Manager offers a comprehensive monitoring and management solution for storage infrastructure. The solution provides comprehensive reports of utilization and trends for capacity planning and space usage. It also monitors system performance, storage capacity, and health to resolve potential problems.

3.4 Management

Although the element managers discussed in the previous sections provide thorough control of their respective components, more integrated management is required to handle the daily operation of the environment.

FlexPod Management Solutions

To give customers choices for this role, NetApp and Cisco joined with trusted partners to form the FlexPod Management Solutions Program, which includes two classes of solutions:

- **FlexPod management solutions.** These solutions, provided by partners, address specific management needs on each of the FlexPod components, including Cisco Nexus, the Cisco Unified Computing System (UCS), and NetApp storage controllers. These packages can provide a variety of services, including configuration management, health and performance monitoring, VM lifecycle management, and more.
- **Validated FlexPod management solutions.** These partner solutions were verified by NetApp and Cisco to fulfill a rigorous set of use case requirements and to provide unified management of the FlexPod environment. Validated solutions must meet a high bar of integration with all of the FlexPod components as well as usability and feature requirements that include:
 - A unified hierarchy of stack resources
 - A service catalog
 - A customer-facing application programming interface (API)
 - Fine-grained resource metering
 - The ability to integrate into an existing management environment

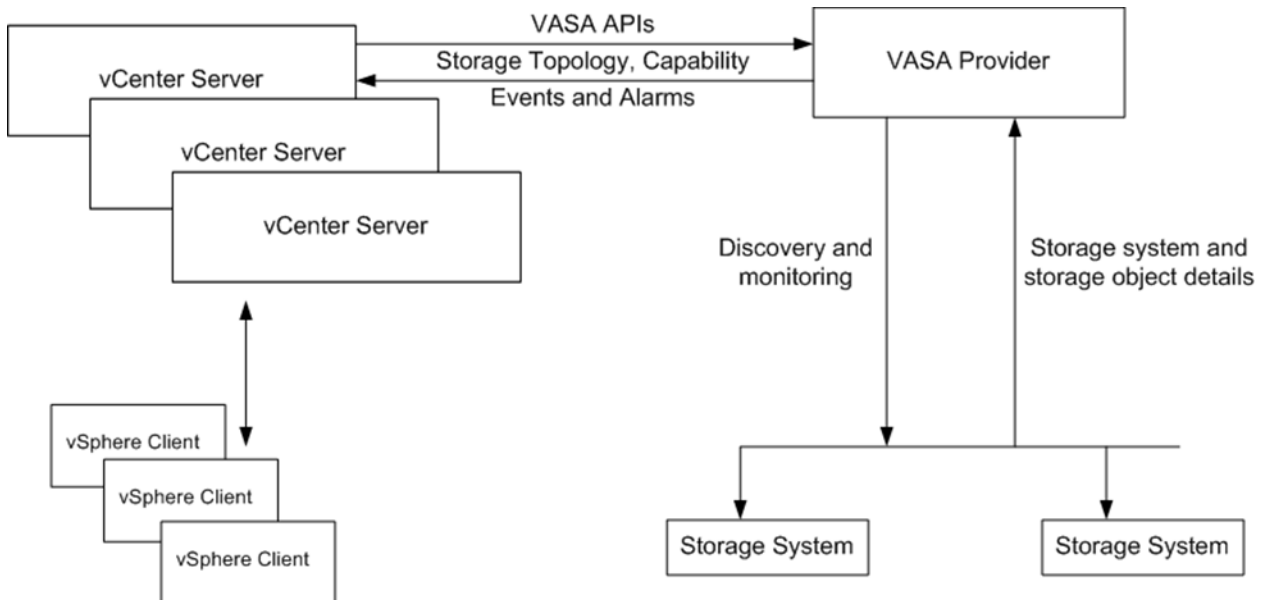
All validated solutions share this core set of capabilities, but our partners also go beyond this prescribed baseline to innovate and differentiate their solutions. This work enables customers to compare with confidence and to find the validated solution that matches the needs of their enterprises.

This program enables customers to easily identify FlexPod management solutions from qualified developers, and these validated solutions offer the additional assurance of FlexPod unified management and verification by Cisco and NetApp. For more information about the FlexPod management software ecosystem, refer to FlexPod [Management Solutions](#).

NetApp VASA Provider

NetApp is a VMware vStorage APIs for Storage Awareness (VASA) provider. Providers collect information about storage and present that information to vSphere. The FAS/V-Series VASA provider architecture includes storage systems running Data ONTAP, a host system running a FAS/V-Series VASA provider, and a vCenter Server that requests storage information from NetApp storage systems. Figure 3 shows how a FAS/V-Series VASA provider collects data and delivers that information to the vCenter Server.

Figure 3) VASA provider collects and delivers data to vCenter Server.



4 Data Center Solutions

When deploying a shared infrastructure, customers often have both virtualized and nonvirtualized workloads deployed in their data centers. This section details available solutions and describes how to deploy FlexPod as the basis for any data-center workload.

4.1 FlexPod Infrastructure

This section describes infrastructure architectures that support both virtualized and nonvirtualized platforms that might be deployed.

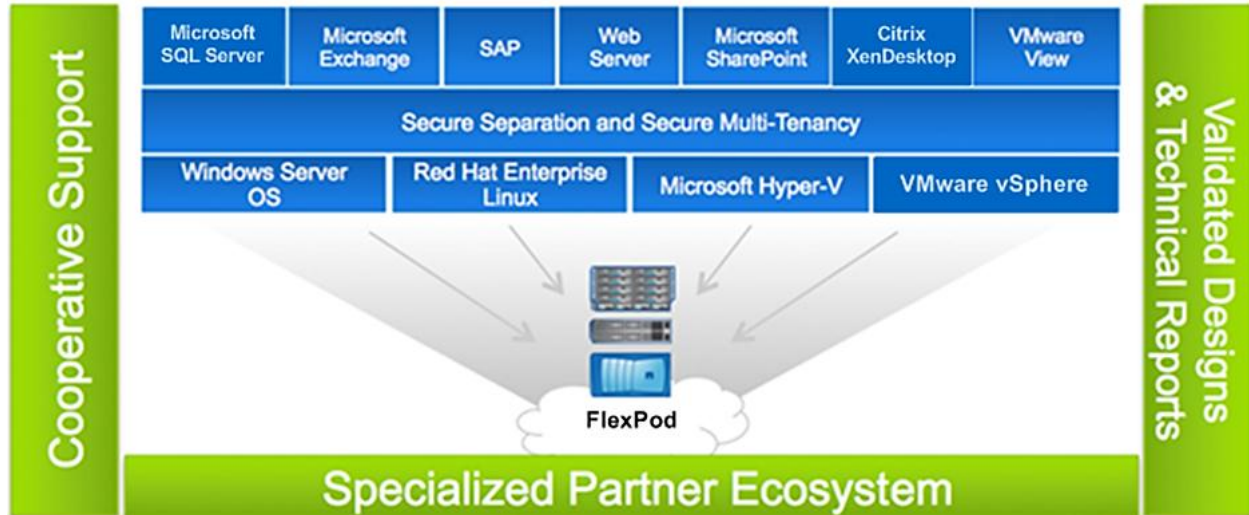
Base FlexPod

FlexPod is a predesigned and validated base configuration that leverages the Cisco Unified Computing System (UCS), Cisco Nexus data center switches, NetApp FAS storage components, and software from a

range of partners. FlexPod can scale up for greater performance and capacity, or it can scale out for environments that need consistent, multiple deployments.

FlexPod includes all of the infrastructure elements that serve as the foundation for solution layering. Many customers today require the ability to support a variety of operating systems to satisfy this demand. FlexPod can be deployed as a virtualized, nonvirtualized, or hybrid environment. Various documented architectures, including FlexPod Datacenter with VMware vSphere, Red Hat Enterprise Linux built on FlexPod, and SAP applications built on FlexPod, provide fully virtualized, nonvirtualized, and hybrid environments, respectively.

Figure 4) FlexPod solutions and applications.



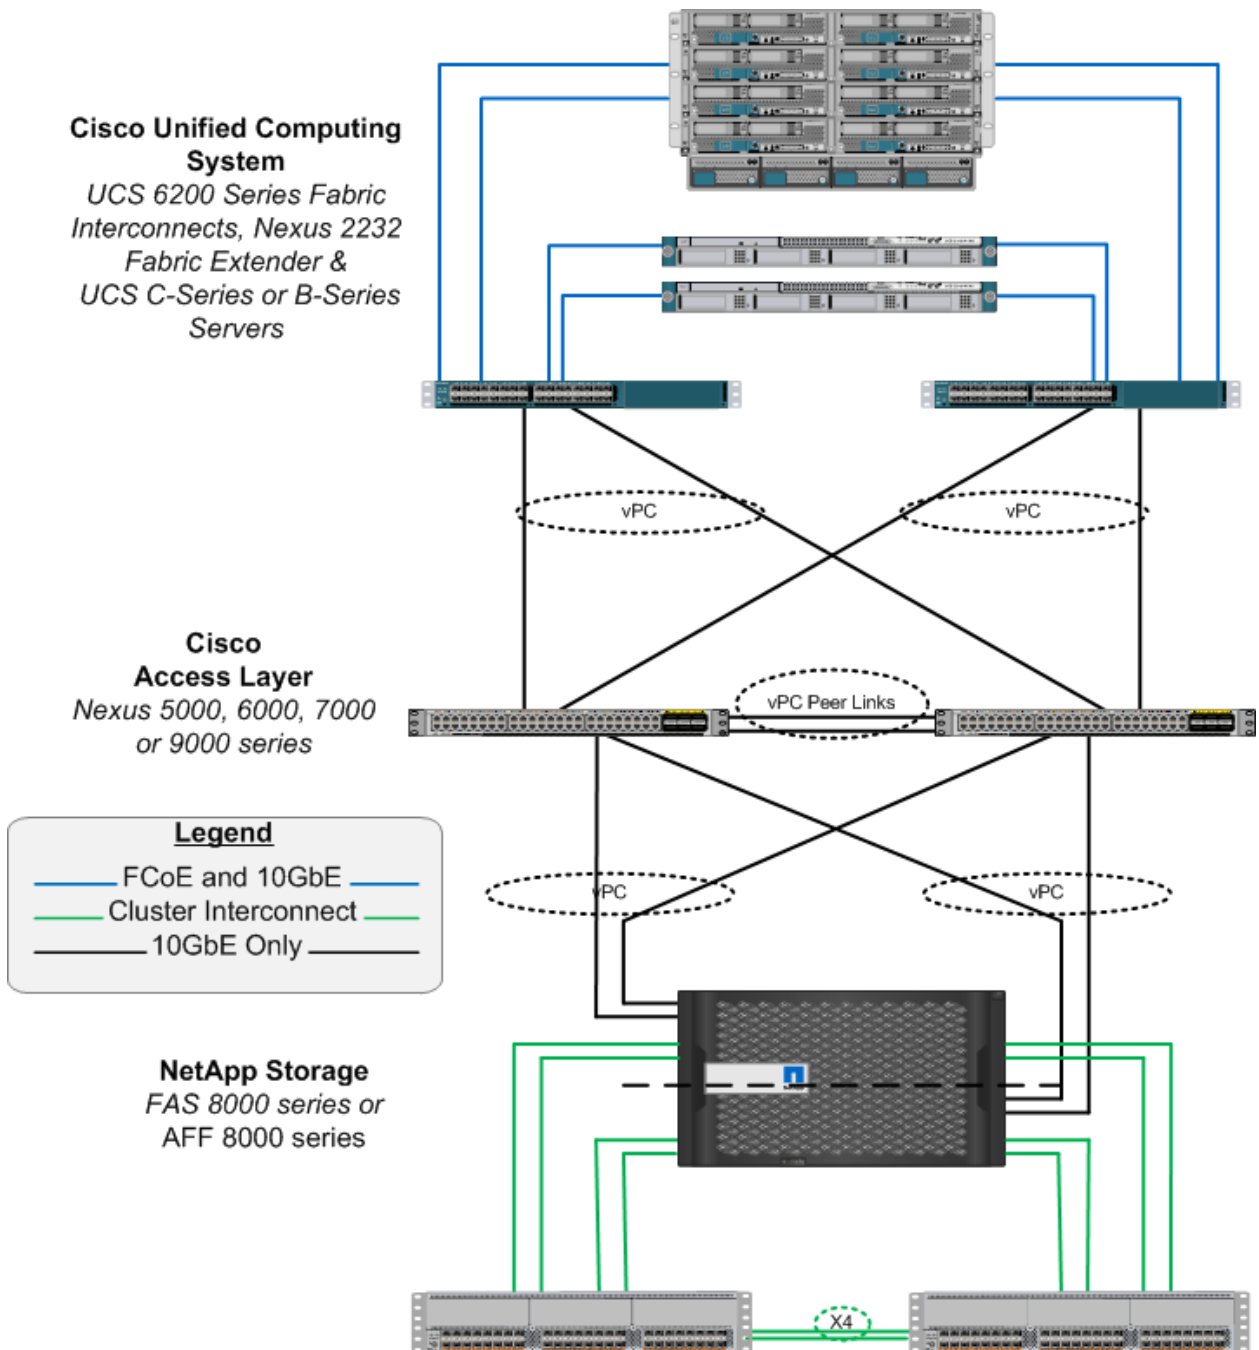
FlexPod also has the flexibility to be sized and optimized to accommodate many different use cases. For example, a customer who wants a FlexPod solution to satisfy VDI requirements would require higher-performance blades and optimized NetApp Flash Cache technologies. For a development and test environment, a customer would require more compute resources and less storage because of inherent NetApp storage efficiencies and the capability to eliminate redundant data.

In contrast, for a data protection and backup environment, a customer would require less computing and more storage capacity. Customers can also start with a small, simple FlexPod configuration and scale up or out, depending on the needs of the business. Figure 5 shows the structure of a single FlexPod configuration using FCoE and a Cisco Nexus 5000 Series switch. Figure 6 shows the structure of a single FlexPod configuration using an iSCSI boot and a Cisco Nexus 7000 Series switch. Native FC is also an option, but requires a switch capable of supporting native FC.

18



Figure 6) iSCSI boot FlexPod infrastructure.



For more information about this solution, refer to the following documentation:

- [CVD: FlexPod Datacenter with NetApp All Flash FAS, Cisco Nexus 9000 ACI, and VMware vSphere](#)
- [CVD: FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5 with IP-Based Storage](#)
- [CVD: FlexPod Datacenter Data ONTAP with VMware vSphere 5.1 Design Guide](#)
- [CVD: FlexPod Datacenter with VMware vSphere 5.1](#)
- [CVD: FlexPod Datacenter with VMware vSphere 5.1 and Nexus 7000 with 7-Mode](#)
- [FlexPod Solutions](#)

- [NVA-0003: FlexPod Datacenter Solution: VMware vSphere Built on FlexPod – NetApp Verified Architecture](#)
- [VMware vSphere Built on FlexPod with IP-Based Storage](#)

Cisco Secure Enclaves Architecture Built on FlexPod

The Cisco Secure Enclaves Architecture built on FlexPod is a threat-centric approach to security that enables customers to consistently address the full attack continuum before, during, and after an attack on a standard platform. The solution is based on a FlexPod Datacenter integrated system and is augmented with services to address business, compliance, and application requirements. The FlexPod Datacenter with Cisco Secure Enclaves solution provides a standardized method for delivering a flexible, functional, and secure application environment that can be readily automated.

In addition to the FlexPod Datacenter infrastructure components, this solution includes:

- VMware vSphere
- Cisco Adaptive Security Appliance
- Cisco NetFlow Generation Appliance
- Cisco Virtual Security Gateway
- Cisco Identity Services Engine
- Cisco Network Analysis Module
- Cisco UCS Director
- Lancope StealthWatch System

For more information about this solution, refer to the following documentation:

- [Cisco Secure Enclaves Architecture Design Guide](#)
- [Cisco Secure Enclaves Architecture Built on FlexPod](#)

FlexPod Datacenter with VMware vSphere

The VMware vSphere built on FlexPod solution architecture layers the VMware vSphere hypervisor on top of the FlexPod infrastructure stack. Cisco, NetApp, and VMware developed VMware vSphere built on FlexPod to address industry virtualization needs and simplify the evolution to shared infrastructure and cloud architectures. VMware vSphere built on FlexPod is the foundation for various virtualized application workloads and solutions, including the secure multi-tenancy and disaster recovery configurations described in section 4.3.

In addition to the FlexPod base infrastructure components, this solution includes:

- VMware vSphere ESXi
- VMware vSphere vCenter Server
- VMware vSphere Client
- NetApp Virtual Storage Console
- Cisco Nexus 1000v Virtual Ethernet Switch

For more information about this solution, refer to the following documentation:

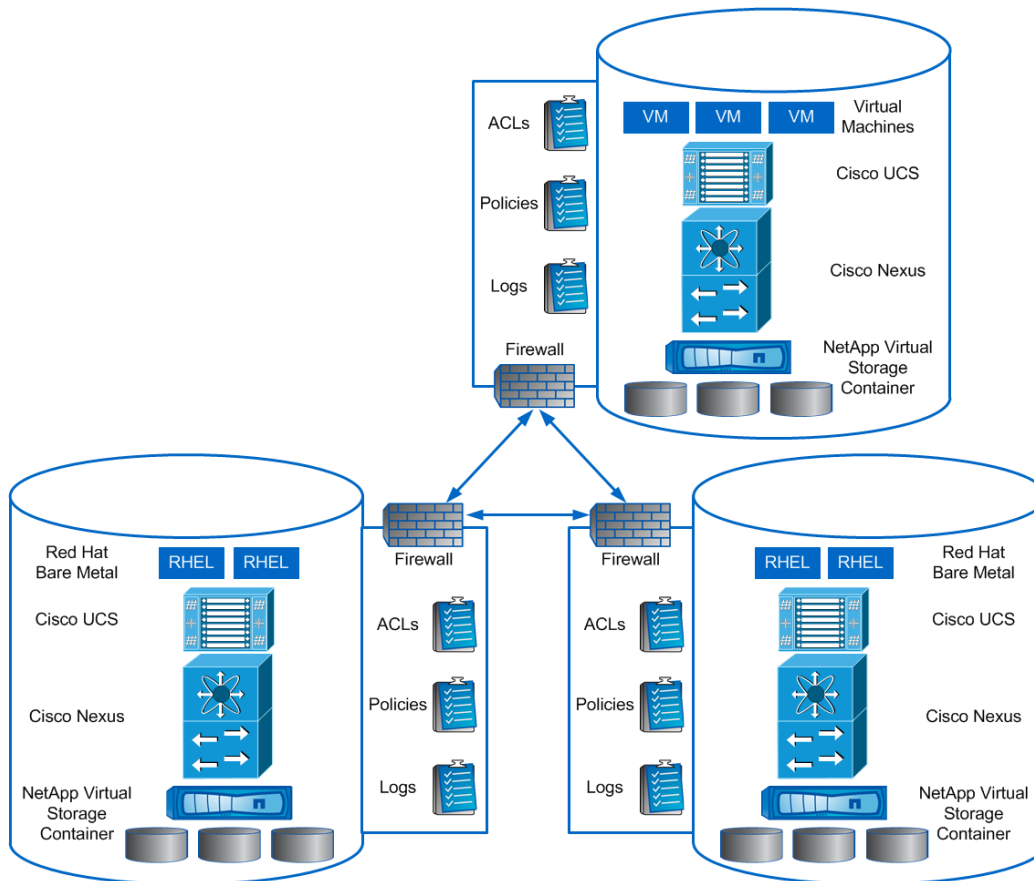
- [CVD: FlexPod Datacenter with Cisco UCS Mini and VMware vSphere 5.5](#)
- [NVA-0005: FlexPod Datacenter Solution: VMware vSphere on FlexPod – NetApp Verified Architecture](#)
- [TR-4113: VMware vSphere 5.0 on FlexPod Data ONTAP 7-Mode Deployment Guide](#)
- [TR-4114: VMware vSphere 5.0 on FlexPod Clustered Data ONTAP Deployment Guide](#)
- [VMware vSphere Built on FlexPod – Data Center Solution Site](#)

- [TR-4119: Oracle RAC Built on FlexPod with VMware Solutions Guide](#)

Secure Separation Built on FlexPod Design Guide

Cisco and NetApp along with Red Hat jointly designed a best-in-class nonvirtualized FlexPod Datacenter architecture with secure separation and they validated the design in the lab environment. The nonvirtualized environment uses a native operating system (Red Hat Enterprise Linux 6) on Cisco Unified Computing System (UCS) blades rather than a thin hypervisor. This section provides guidance for implementing securely separated tenants in a nonvirtualized (also known as bare metal) data center built on FlexPod. Tenants can be a group of users, a group of applications, or any group that needs committed resources from the data center.

Figure 7) Secure separation architecture overview.



For more information about this solution, refer to the [Secure Separation Built on FlexPod Design Guide](#).

Red Hat Enterprise Linux Built on FlexPod

The Red Hat Enterprise Linux built on FlexPod solution architecture layers the Red Hat Enterprise Linux operating system on top of the FlexPod infrastructure stack. You can use FlexPod and the Red Hat Enterprise Linux operating system and its add-on options, such as scalable file system, high availability, and resilient storage, as a trusted platform for mission-critical applications. Red Hat Enterprise Linux built on FlexPod can be used in standalone or cluster configurations on physical Cisco UCS servers or as a guest operating system in virtualized environments.

In addition to the FlexPod base infrastructure components, this solution includes:

- Red Hat Enterprise Linux
- Red Hat Cluster Suite

For more information about this solution, refer to the [Red Hat Enterprise Linux Built on FlexPod Deployment Guide](#).

NetApp Hyper-V Cloud Fast Track with Cisco

NetApp Hyper-V Cloud Fast Track with Cisco, a FlexPod Datacenter solution, includes NetApp storage, Cisco networking, Cisco UCS, and Hyper-V virtualization software in a single package. In this solution, computing and storage fit in one data center rack and networking resides in a separate rack. Depending on port density, the networking components can accommodate multiple Fast Track units for Hyper-V configurations.

In addition to the FlexPod base infrastructure components, this solution includes:

- Microsoft Windows Server 2008 R2 Datacenter Edition
- Microsoft System Center 2007
- Microsoft System Center Virtual Machine Manager
- Microsoft System Center Operations Manager
- Opalis Integration Server
- NetApp OnCommand® Plug-in for Hyper-V
- NetApp SnapDrive® for Windows®
- NetApp SnapManager® for Microsoft Hyper-V
- NetApp Data ONTAP DSM

For more information, refer to the following documentation:

- [FlexPod Validated with Microsoft Private Cloud Reference Architecture and Deployment Guide](#)
- [TR-4094i: FlexPod Validated with Microsoft Private Cloud Fast Track 2.0 Deployment Guide](#)

FlexPod Datacenter with NetApp MetroCluster Software

A growing segment of the industry wants to reduce its investment risk by introducing geographical separation between data centers while logically maintaining a single data center. These virtualized data centers are enabled by the proliferation of dark fiber between cities at ever-extending geographical distances. As customers continue to outgrow their existing data centers, they look for ways to effectively extend their service offerings without forklift data center replacements. Being able to host two data centers in different geographical areas, operating as a single entity, is becoming increasingly popular.

The primary benefit of a stretched-cluster model is to enable on-demand and nonintrusive mobility of workloads across two geographically distant sites. By using VMware Metro Storage Cluster (vMSC) design guidance, workload mobility is combined with active resource balancing across the two data centers to optimally utilize data center resources. Some of the key benefits of a multisite data center solution are:

- Workload mobility
- Automated load balancing across sites
- Ease of maintenance
- Avoidance of disaster and downtime

For more information, refer to the following documentation:

- [CVD: FlexPod Datacenter with Cisco Nexus 7000 and NetApp MetroCluster for Multisite Deployment](#)
- [TR-3548: Best Practices for MetroCluster Design and Implementation](#)
- [TR-4128: vSphere 5 on NetApp MetroCluster Solution](#)

4.2 FlexPod Application Solutions

The previous section explained use cases for deploying FlexPod as the architecture of choice. This section introduces specific application environments with validated best practices that can be built on FlexPod.

SAP Built on FlexPod

Corporations deploying SAP software today are under pressure to reduce cost, minimize risk, and control change by accelerating deployments and increasing the availability of their SAP landscapes. Deployment of business solutions usually exceeds a single production instance of SAP. Business process owners and project managers must coordinate with IT management to optimize the scheduling and availability of systems to support rapid prototyping and development, frequent parallel testing or troubleshooting, and appropriate end-user training levels. The ability to access these systems as project schedules dictate, with current datasets and without affecting production operations, often determines whether SAP projects are delivered on time and within budget.

To meet this challenge, NetApp and Cisco collaborated to create SAP built on FlexPod. This solution differs from other virtualized infrastructure offerings by providing:

- Validated technologies from industry leaders in computing, storage, networking, and server virtualization
- A single platform, built from unified computing, fabric, and storage technologies, that scales to meet the largest data center requirements without disruption or architectural changes in the future
- Integrated components that enable administrators to centrally manage all infrastructure pools
- An open design management framework that integrates with existing third-party infrastructure management solutions
- Support of bare-metal-server VMs and VMs that are based on VMware
- Virtualization on all layers of the solution stack
- Secure multi-tenancy for operating fenced SAP systems or landscapes
- Application and data mobility
- Integrated storage-based backup
- Provisioning of infrastructure components, such as tenants and operating systems
- Automated SAP system copies
- Provisioning of fenced SAP systems based on clones of production systems

SAP built on FlexPod introduces an infrastructure that is based on virtualization technologies on all layers of the solution stack within a pool of shared resources. The SAP systems can run on VMware VMs or on bare metal servers. All management applications are installed in the infrastructure tenant. These management applications run on VMware VMs.

For more information, refer to the following documentation:

- [FlexPod for SAP Applications](#)
- [TR-3921i: SAP Applications Built on FlexPod Implementation Guide](#)

Microsoft Solutions Built on FlexPod

Many customers want to virtualize traditional Microsoft server instances as they move toward a shared infrastructure with lower costs. Many solutions leverage FlexPod Datacenter with VMware vSphere as the foundation for these virtualized workloads. As mentioned previously, the VMware portion of FlexPod Datacenter with VMware vSphere offers the reduced costs of virtualization and high availability.

The Cisco Nexus platform's unified switching fabric enables connections between VMs and storage by using FC, FCoE, iSCSI, NFS, or CIFS over a single network fabric. By purchasing SnapDrive for Windows, SnapManager for Microsoft Exchange, SnapManager for SQL Server, SnapManager for Microsoft Office SharePoint Server, the backup portion of NetApp VSC, and SnapMirror, you can set up advanced backup and recovery solutions.

[TR-3151: FlexPod Datacenter Solutions Sizing Guide](#), described in section 2.3, specifically addresses sizing of Microsoft mixed-workload applications with VMware built on FlexPod. Customer workload requirements can be used to customize the compute and storage provided in VMware built on FlexPod to meet these requirements.

Additional collateral has been developed describing how to deploy Microsoft SharePoint 2010 for a real-world workload of 50,000 users. This design guide demonstrates how enterprises can leverage best practices for FlexPod technologies and software providers Microsoft and VMware. Design validation was completed using Microsoft Visual Studio 2010 Ultimate Workloads and Visual Studio Load Test agents to simulate a realistic 50,000-user load across SharePoint 2010 Enterprise features. Virtual user connections from branch offices, remote offices, and mobile users were optimized using Cisco Wide Area Application Services. Cisco Application Control Engine was implemented for web and application server load balancing in the SharePoint 2010 architecture.

For more information, refer to the following documentation:

- [CVD: Microsoft Exchange 2010 with VMware vSphere on Cisco Unified Computing System with NetApp Storage](#)
- [CVD: Microsoft SharePoint 2010 on FlexPod for VMware](#)
- [TR-3785: Running Microsoft Enterprise Applications on VMware vSphere, NetApp Unified Storage, and Cisco Unified Fabric](#)
- [CVD: FlexPod for Windows Server 2012 Hyper-V Design Guide](#)

VMware View 4.5 Built on FlexPod

The VMware View 4.5 desktop virtualization platform enables administrators to run virtual desktops in the data center and deliver desktops to employees as a secure, managed service. End users gain a familiar, personalized environment that they can access from any number of devices anywhere throughout the enterprise or from home. Administrators gain centralized control, efficiency, and security by having desktop data in the data center.

There are many reasons to consider a virtual desktop solution, such as an ever-growing and diverse base of user devices and the complexity of managing traditional desktops, security, and user-owned/non-IT-supported devices. To design and deploy a successful virtual desktop environment, it is important to understand the requirements of the user community. These are some typical types of users:

- Knowledge workers today are not tied to their office all day; they attend meetings, visit branch offices, and work from home and even coffee shops. These "anywhere workers" expect access to all of their applications and data wherever they are.
- External contractors are increasingly part of everyday business. They need access to many applications and data, yet administrators have little control over the devices they use or the locations

from which they work. Consequently, IT must trade off the cost of providing these workers with a device versus the security risk of allowing them access from their own devices.

- Task workers perform a set of well-defined tasks. These workers access a small set of applications and have limited requirements for their PCs. Because these workers interact with customers, partners, and employees, they often have access to critical data.
- “Road warriors” need access to their virtual desktop from everywhere, regardless of how they connect to a network. These workers expect to be able to personalize their PCs by installing their own applications and storing their own data, such as photos and music, on these devices.
- Shared workstation users are typically found in university and business computer labs, in conference rooms, and in training centers. Shared workstation environments require desktops to be reprovisioned with the latest operating systems or applications as the needs of the organization change.

The virtual desktop user community requirements drive system design decisions. For more information about this solution, refer to the [VMware View 4.5 on FlexPod for VMware Design Guide](#).

Citrix XenDesktop on FlexPod Datacenter with VMware vSphere

Citrix XenDesktop is a desktop virtualization solution that delivers Windows desktops as an on-demand service to any user, anywhere. With FlexCast delivery technology, XenDesktop can quickly and securely deliver individual applications or complete desktops to everyone in the entire enterprise, whether they are task workers, knowledge workers, or mobile workers. Users now have the flexibility to access their desktops on any device, at any time, with a high-definition user experience.

With XenDesktop, IT can manage single instances of each OS, application, and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management. The XenDesktop open architecture enables customers to easily adopt desktop virtualization by using any hypervisor, storage, or management infrastructure.

In this solution, VMware vSphere is the chosen hypervisor. VMware vSphere consists of the management infrastructure or virtual center server software and the hypervisor software that virtualizes the hardware resources on the servers. vSphere offers features such as its Distributed Resource Scheduler, VMware vMotion, high availability, Storage vMotion, VMware Virtual Machine File System, and a multipathing storage layer. For more information about vSphere, refer to the [VMware](#) website.

For more information about this solution, refer to the following documentation:

- [CVD: Citrix XenDesktop on FlexPod with Microsoft Private Cloud](#)
- [Reference Architecture Based Design for Implementation of Citrix XenDesktop Using Cisco Unified Computing System, Microsoft Hyper-V, and NetApp Storage](#)

Citrix XenDesktop with XenServer Built on FlexPod

Citrix XenDesktop is hypervisor agnostic and is another example of a hypervisor that can be used to host VDI-based desktops. XenServer is a complete, managed server virtualization platform built on the powerful Xen hypervisor. Xen technology is widely acknowledged as the fastest and most secure virtualization software in the industry. XenServer is designed for efficient management of Windows and Linux virtual servers and delivers cost-effective server consolidation and business continuity. For more information about XenDesktop, refer to the [Citrix](#) website.

For more information about this solution, refer to [CVD: Reference Architecture Based Design for Implementation of Citrix XenDesktop on Cisco Unified Computing System, Citrix XenServer, and NetApp Storage](#).

4.3 FlexPod Datacenter Integration

The solutions outlined in sections 4.1 and 4.2 are important because they are integrated into existing data center standards, best practices, and workflows. Common themes in data centers are multi-tenancy, development and testing, and backup and recovery. FlexPod solutions integrate seamlessly with established best practices and procedures.

Secure Multi-Tenancy

IT infrastructure too often suffers from oversegmentation of server, networking, and storage resources. Each organization or customer purchases and uses its own equipment, which leads to low utilization, inefficiency, and the inability to scale properly and respond quickly to changing business needs. Virtualization of server and network environments has been effective in increasing the use of storage and network resources. However, the adoption of cloud computing to deliver ITaaS in data centers promises to complete the vision of a fully virtualized environment.

The biggest obstacle to the adoption of ITaaS has been a lack of confidence that data and applications are securely isolated in a cloud-based infrastructure in which servers, networks, and storage are shared resources. To address this need, Cisco, NetApp, and VMware joined together to develop secure multi-tenancy (SMT) in a virtualized data center, which is a carefully designed and lab-validated solution for the next-generation data center.

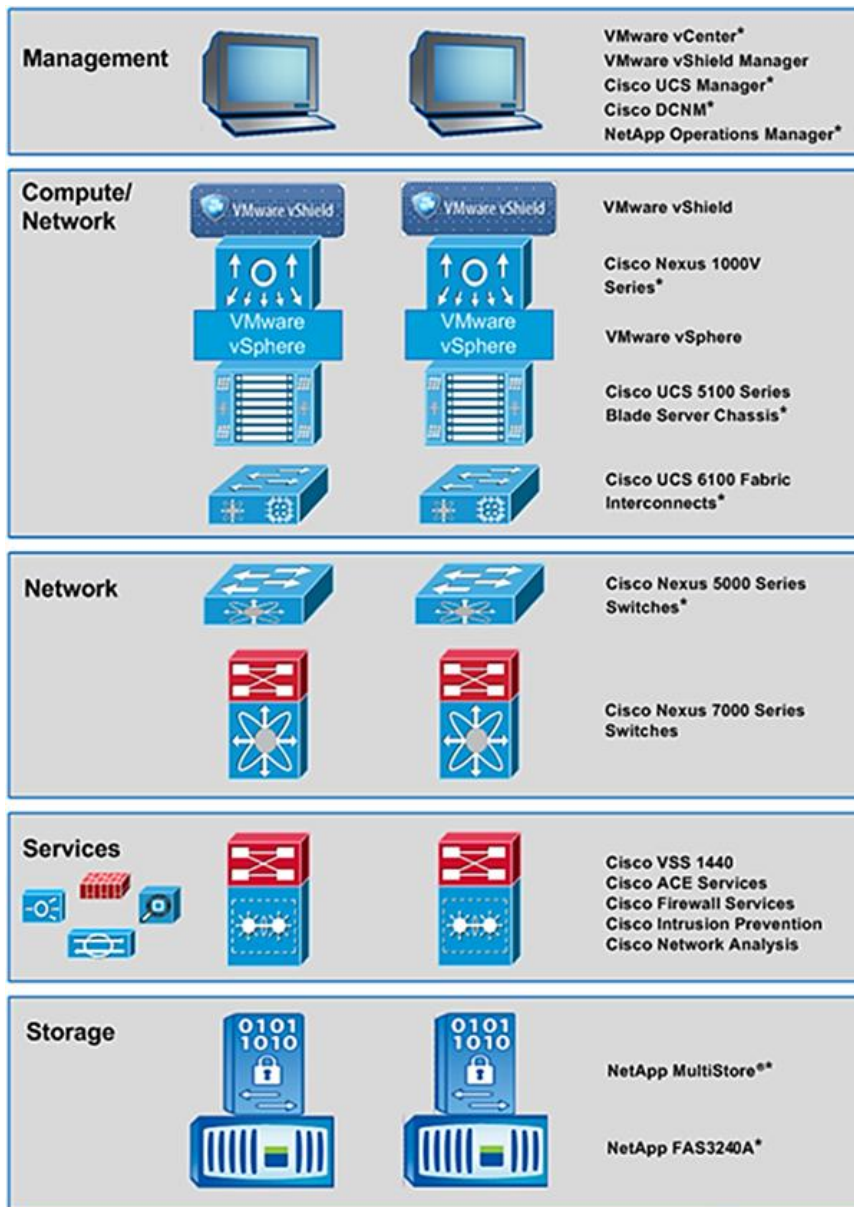
To develop a robust design, organizations must clearly define their requirements and apply proven methodology and design principles. The following four requirements are the critical elements, or pillars, of the SMT architecture:

- **Availability.** This pillar allows the infrastructure to meet the expectation that compute, network, and storage resources will be available even in the event of failure. Each layer has its own method for providing an HA configuration that works cohesively with adjacent layers, similar to the approach used by the secure separation pillar. Security and availability are best deployed in a layered approach.
- **Secure separation.** This pillar makes sure that one tenant does not have access to another tenant's resources, such as VMs, network bandwidth, and storage. Each tenant must be securely separated by using techniques such as access control, virtual LAN segmentation, and virtual storage controllers. Also, each layer has its own means of enforcing policies that help reinforce the policies of the adjacent layers.
- **Service assurance.** This pillar provides isolated compute, network, and storage performance during both steady-state and nonsteady-state operations. For example, the network can provide each tenant with a certain bandwidth assurance by using quality-of-service measures. Resource pools in VMware help balance and provide CPU and memory resources, and the NetApp FlexShare® quality-of-service tool can balance resource contention across storage volumes.
- **Management.** This pillar is required to rapidly provision and manage resources and to view resource availability. Domain and element management provide comprehensive administration of the shared resources that compose the secure cloud architecture. The demarcation point for managing this design is defined by the interactive and programmable interfaces delivered by NetApp, VMware, and Cisco. The administrative interfaces and APIs in this portfolio address infrastructure components such as VMware vCenter and vCloud Director, Cisco UCS Manager and Data Center Network Manager, and the NetApp Manageability Suite. These element managers and their associated open APIs form the foundation for delivering cohesive service lifecycle orchestration with solution partners.

Figure 8 shows the architectural components of SMT.

Note: The components included in FlexPod Datacenter with VMware vSphere are marked with an asterisk in Figure 8.

Figure 8) SMT architectural components (graphic supplied by Cisco).



VMware vSphere 4.1 (or later) built on FlexPod serves as the base layer of this architecture. The additional SMT components can be layered on top of the infrastructure. The additional components are:

- Cisco Nexus 7000 Series switches
- VMware vShield App and Edge virtual firewalls
- VMware Site Recovery Manager (required when implementing DR with SMT)
- NetApp SnapDrive and SnapManager products
- Cisco Services layer

The [Enhanced Secure Multi-Tenancy Design Guide](#) specifies the setup of Microsoft applications such as Microsoft Exchange, SQL Server, and SharePoint as securely separated tenants. The [Enhanced Secure Multi-Tenancy Design Guide](#) also specifies the setup of a disaster recovery (DR) site for recovery of the SMT infrastructure and tenants. To use DR site equipment during normal operations, follow the design

guide specification to set up a development and testing environment running VMware vCloud at the secondary DR site. During DR failover, the production environment preempts the development and testing environment.

For more information, refer to the following documentation:

- [CVD: Enhanced Secure Multi-Tenancy Design Guide](#)
- [WP-7083: Delivering IT as a Service for Virtualized Data Centers](#)
- [WP-7113: An Overview of the Enhanced Secure Multi-Tenancy Architecture](#)

Development and Testing

For businesses of all sizes, the demand for database application development continues to grow at a dramatic rate. Full or partial production database copies are required for many activities, including application development and testing. Often, many copies of the production database are required to support multiple code streams as part of a modular development process.

NetApp solutions can dramatically reduce the time, resources, and effort required for rapid and cost-effective database cloning and deployment. NetApp FlexClone clones require less storage than traditional clones, and they can be created without the storage administrator's support. In addition, the ability to use NetApp storage efficiency technologies from the VMware vSphere layer (through the invocation of storage layer services by the VMware vCenter Server) enables rapid and efficient provisioning of server and storage resources for development and testing.

The server virtualization technologies offered by VMware vSphere enable the creation of a base VM that runs a database application that is properly installed and tested. This base VM, which can actually simulate a real production database system, is stored as a set of files in the shared NetApp storage system. It is self-contained and includes a guest operating system, a database application, and all the data and log partitions. Cloning of the VM using NetApp storage-efficient thin cloning technologies, provided by the NetApp Virtual Storage Console (VSC), makes available everything needed for a functional development and testing environment.

Another approach to this solution is to first create cloned development and testing VMs that contain just the guest operating system and the database binaries. You then mount the database clones created from the production database instance by using NetApp thin cloning technologies on the cloned VMs.

Database and VMware administrators can use powerful, automated NetApp tools, such as SnapManager and VSC provisioning, to easily and quickly create database and VM clones using negligible additional storage and without requiring extensive storage skills or storage administrator support.

The development and testing solution fits on top of VMware built on FlexPod. This solution has the required storage system features licensed, including NetApp FlexClone, iSCSI, and NFS. The NetApp VSC provisioning module is also included in VMware built on FlexPod and you can use it to create space-efficient development and testing clones of production VMs.

Note: You must purchase NetApp SnapDrive and SnapManager separately for the database application.

When NetApp SnapDrive and SnapManager are used, LUNs for the development and testing database clones can be mapped directly to the VM by using iSCSI. They can also be mapped as VMware iSCSI raw device mapping (RDM) LUNs. With VMware built on FlexPod, space-efficient development and testing clones of production environments can be set up quickly using very little additional storage. Doing so accelerates the application development and testing cycle.

Also, the server blades in the Cisco UCS can be grouped. Pools of blades can be defined in development and testing and in the production organizations. These organizations physically isolate development and testing blade resources from production blade resources and thus isolate the development and testing environment from the production environment. In this setup, development and testing blades and

production blades are in two separate VMware ESXi clusters, further separating the two environments. Another way to accomplish development and testing separation is to put all blades into a single organization and a single VMware ESXi cluster and to use resource pools to control resource use in the cluster. Doing so provides virtual isolation of development and testing.

For more information, refer to the following documentation:

- [RA-0005: Accelerating Development of Microsoft SQL Applications in Heterogeneous Environments](#)
- [TR-3861: Oracle Dev/Test on VMware vSphere and NetApp Storage Solutions Overview](#)
- [TR-3742: Using FlexClone to Clone Files and LUNs](#)

Disaster Recovery


Setting up a disaster recovery (DR) environment involves setting up a secondary or recovery data center site with hardware (compute, network, and storage) that closely mirrors the hardware in the primary or protected data center. If the protected data center fails, the VMs that ran in the protected data center can be brought up and into production on the hardware at the recovery data center. VMs and VM application data are periodically replicated from protected data center storage to recovery data center storage. In addition, a test environment can be brought up at the recovery data center to test the disaster recovery plan. Finally, one recovery data center can serve as the DR site for multiple protected data centers.

A FlexPod configuration can function as a recovery DR site. An ideal situation is to have one or more such configurations at the protected site with one or more FlexPod for VMware configurations at the recovery DR site. The configuration for this solution is simplified because the sites can be set up as mirror images of each other.

To implement DR on VMware built on FlexPod, the VMware Site Recovery Manager (SRM) must be purchased. Also, the NetApp adapter for the VMware SRM must be added to the environment to automatically manage the DR plan and to implement offline tests of this plan at the recovery site. Additionally, the backup portion of NetApp VSC (previously NetApp SnapManager for Virtual Infrastructure) must be purchased. Also, NetApp SnapMirror, SnapDrive, and the appropriate SnapManager products must be purchased to back up and replicate VMs and VM data to the DR site. The recovery site networking and VMware environment must be set up manually, similarly to the protected site setup.

Figure 9 shows a logical representation of a DR setup. VMware built on FlexPod configurations can be used for both the protected site and the recovery site with the purchase of the software listed in the previous paragraph.

Figure 9) VMware SRM with NetApp storage.



For more information, refer to the following documentation:

- [TR-3671: Deploying VMware vCenter Site Recovery Manager with NetApp FAS/V-Series Storage Systems](#)
- [TR-3822: Disaster Recovery of Microsoft Exchange, SQL Server, and SharePoint Server Using VMware vCenter Site Recovery Manager, NetApp SnapManager and SnapMirror, and Cisco Nexus Unified Fabric](#)

5 Conclusion

NetApp and Cisco combined their technologies to develop FlexPod, a new solution for virtual computing. FlexPod is a defined set of hardware and software products that serves as an integrated infrastructure stack for all virtualization solutions. Combining best-in-class compute, network, and storage elements, FlexPod offers an uncompromising infrastructure solution for customers who want to deploy virtualized, nonvirtualized, and hybrid infrastructure solutions for a variety of enterprise applications.

References

This document uses the following references:

- TR-3298: RAID-DP: NetApp Implementation of RAID Double Parity for Data Protection
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60325-16&m=tr-3298.pdf
- TR-3437: Storage Subsystem Resiliency Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60064-16&m=tr-3437.pdf
- TR-3450: High-Availability Pair Controller Configuration Overview and Best Practices
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61088-16&m=tr-3450.pdf
- TR-3505: NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60084-16&m=tr-3505.pdf
- TR-3563: NetApp Thin Provisioning Increases Storage Utilization with on-Demand Allocation
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61045-16&m=tr-3563.pdf
- TR-3633: Best Practices for Oracle Databases on NetApp Storage
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60340-16&m=tr-3633.pdf
- TR-3671: Deploying VMware vCenter Site Recovery Manager 4 with NetApp FAS/V-Series Storage Systems
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61583-16&m=tr-3671.pdf
- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-60173-16&m=tr-3702.pdf
- TR-3705: NetApp and VMware View Solution Guide
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61632-16&m=tr-3705.pdf
- TR-3749: NetApp Storage Best Practices for VMware vSphere
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-104655-16&m=tr-3749.pdf
- TR-3824: Storage Efficiency and Best Practices for Microsoft Exchange Server 2010
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61277-16&m=tr-3824.pdf
- TR-4068: VMware vSphere 5 on NetApp Clustered Data ONTAP Best Practices
www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-61602-16&m=tr-4068.pdf

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. TR-3884-1015