



NetApp Verified Architecture

FlexPod Datacenter with NetApp All Flash FAS and VMware Horizon

NVA Design

David Arnette, NetApp
June 2016 | NVA-1110-FP-DESIGN | Version 2.0

Reviewed by



In collaboration with



TABLE OF CONTENTS

1	Executive Summary.....	4
1.1	New in This Release	4
2	Program Summary.....	6
2.1	FlexPod Program Benefits	7
3	Solution Overview	9
3.1	Target Audience.....	9
3.2	Solution Technology	9
3.3	Use Case Summary	24
4	Technology Components	25
4.1	Hardware Components	25
4.2	Software Components	25
5	Solution Design	26
5.1	Cisco UCS Design	26
5.2	Cisco Nexus Network Design.....	28
5.3	NetApp AFF Storage Design.....	31
5.4	VMware vSphere Design	36
5.5	VMware Horizon View Design.....	40
6	Design Considerations	46
7	Best Practices	52
8	Solution Verification.....	53
	Conclusion	54
	Acknowledgements	54
	Appendix.....	54
	References.....	55
	Cisco UCS.....	55
	Cisco Nexus Networking	56
	NetApp AFF Storage	56
	VMware vSphere	57
	VMware Horizon View	57
	Interoperability Matrixes	57

Version History 57

Trademark Information 58

Copyright Information 58

LIST OF TABLES

Table 1) Changes between NVA releases 5

Table 2) Hardware components. 25

Table 3) Solution software components. 25

Table 4) NetApp AFF8000 storage system technical specifications. 34

Table 5) NetApp VSC VM configuration. 35

Table 6) VMware vCenter Server Appliance VM configuration. 39

Table 7) VMware Horizon View Connection VM example configuration. 42

Table 8) Horizon View Composer VM example configuration. 43

Table 9) Microsoft SQL Server database VM example configuration. 44

Table 10) Virtual desktop configuration. 45

Table 11) Design considerations. 46

Table 12) VDI best practices. 52

LIST OF FIGURES

Figure 1) FlexPod component families. 7

Figure 2) Cisco UCS Manager: Java client. 12

Figure 3) Cisco UCS Manager: HTML5 client. 13

Figure 4) ONTAP. 16

Figure 5) NetApp VSC example. 20

Figure 6) VMware vSphere feature overview (graphic supplied by VMware). 22

Figure 7) VMware Horizon View deployment (graphic supplied by VMware). 23

Figure 8) VMware Horizon View linked clone using View Composer. 24

Figure 9) Example of discrete and port channel modes. 28

Figure 10) Scale-out storage with ONTAP on AFF. 31

Figure 11) Distributed storage clusters. 32

Figure 12) ONTAP in a SAN-only environment. 32

Figure 13) ONTAP in a NAS-only environment. 33

Figure 14) Multipath HA to DS2246 shelves of SSD. 35

Figure 15) Cisco Nexus 1000v architecture (graphic provided by Cisco). 38

Figure 16) VMware vCenter Server Appliance size options. 40

Figure 17) Horizon View architecture. 41

Figure 18) FCoE direct-connect topology. 55

1 Executive Summary

Industry trends indicate a transformation of the data center toward shared infrastructure and cloud computing, including a growing use of desktop virtualization. To increase agility and reduce costs, enterprise customers are moving away from IT operation silos toward more cost-effective virtualized environments and ultimately toward cloud computing. This transformation might appear daunting and complex because companies must address resistance to change in both their organizational and technical IT models. In addition, correctly architecting, deploying, and managing a virtual desktop infrastructure (VDI) can be challenging because of the large number of solution components in the architecture.

NetApp, Cisco, and VMware want to accelerate this process and provide the numerous benefits available from a VDI solution. Therefore, they have developed a solution for VMware Horizon on FlexPod® Datacenter.

A successful VDI implementation must provide you with a positive end-user experience. Indeed, the end-user experience must be as good as or better than any previous experiences you have had on a physical PC or virtual desktop. Typically, storage is the leading cause of end-user performance problems. The NetApp® All Flash FAS (AFF) solution with the AFF8000 platform solves the performance problems commonly found in VDI deployments. Deployed as a part of the FlexPod integrated infrastructure, the AFF solution allows you to scale as needed, prevents interruptions for users, and reduces risk to your business.

An infrastructure failure prevents users from working, which causes lost revenue and productivity. That is why what used to be considered a tier 3 or tier 4 application can now be more critical to business operations. An integrated infrastructure with a robust set of data management and availability features is key to system stability and reducing business risk.

FlexPod has multiple built-in features to help improve availability:

- Active-active high availability (HA) across the entire stack of compute, network, and storage
- Network and storage quality of service (QoS)
- Multiple, redundant connections from the servers through the network to the back-end connectivity between the storage controllers and disks
- Nondisruptive operations to seamlessly move virtual machines (VMs) between hosts in the compute cluster or to move either VM or user data within the storage cluster without affecting the user

FlexPod also allows you to increase compute or storage system capacity by simply adding servers, chassis, disks, or shelves as business needs dictate. There is no need to purchase additional compute or storage controllers to add users when additional capacity is required. When the platform requires expansion, additional nodes can be added in a scale-out fashion and managed within the same management framework and interface. Workloads can then be nondisruptively and seamlessly migrated or balanced to the new nodes in the cluster (compute or storage).

1.1 New in This Release

Since the initial version of this design was released in 2014, there have been multiple releases of all of its software components and changes to many of its hardware components. The key differences in version 2.0 include:

- NetApp AFF storage systems are now distinct products from the traditional FAS storage systems with different model names and SSD-specific optimizations not available on non-AFF systems.
- NetApp Data ONTAP® 8.3 now supports only the clustered version of the operating system (OS) with numerous new features and flash optimizations, including inline compression and inline data deduplication as of version 8.3.2.

- Cisco Nexus switching has expanded to include new hardware series models, and the 5500 series that was used in the initial version has been replaced with 9000 series switches.
- The Cisco Unified Computing System (Cisco UCS) boot from SAN configuration now uses iSCSI rather than FCoE for the storage protocol.
- The datastores for virtual machines now use NFS rather than FCoE for the storage protocol.
- The scale for the reference architecture has increased from 2,000 desktops to 2,500 desktops with commensurate changes to the configuration of the underlying hardware.
- All software components, including the OSs, have been upgraded to more recent versions.

Table 1 lists the detailed differences between version 1.0 and version 2.0 of this NVA.

Table 1) Changes between NVA releases.

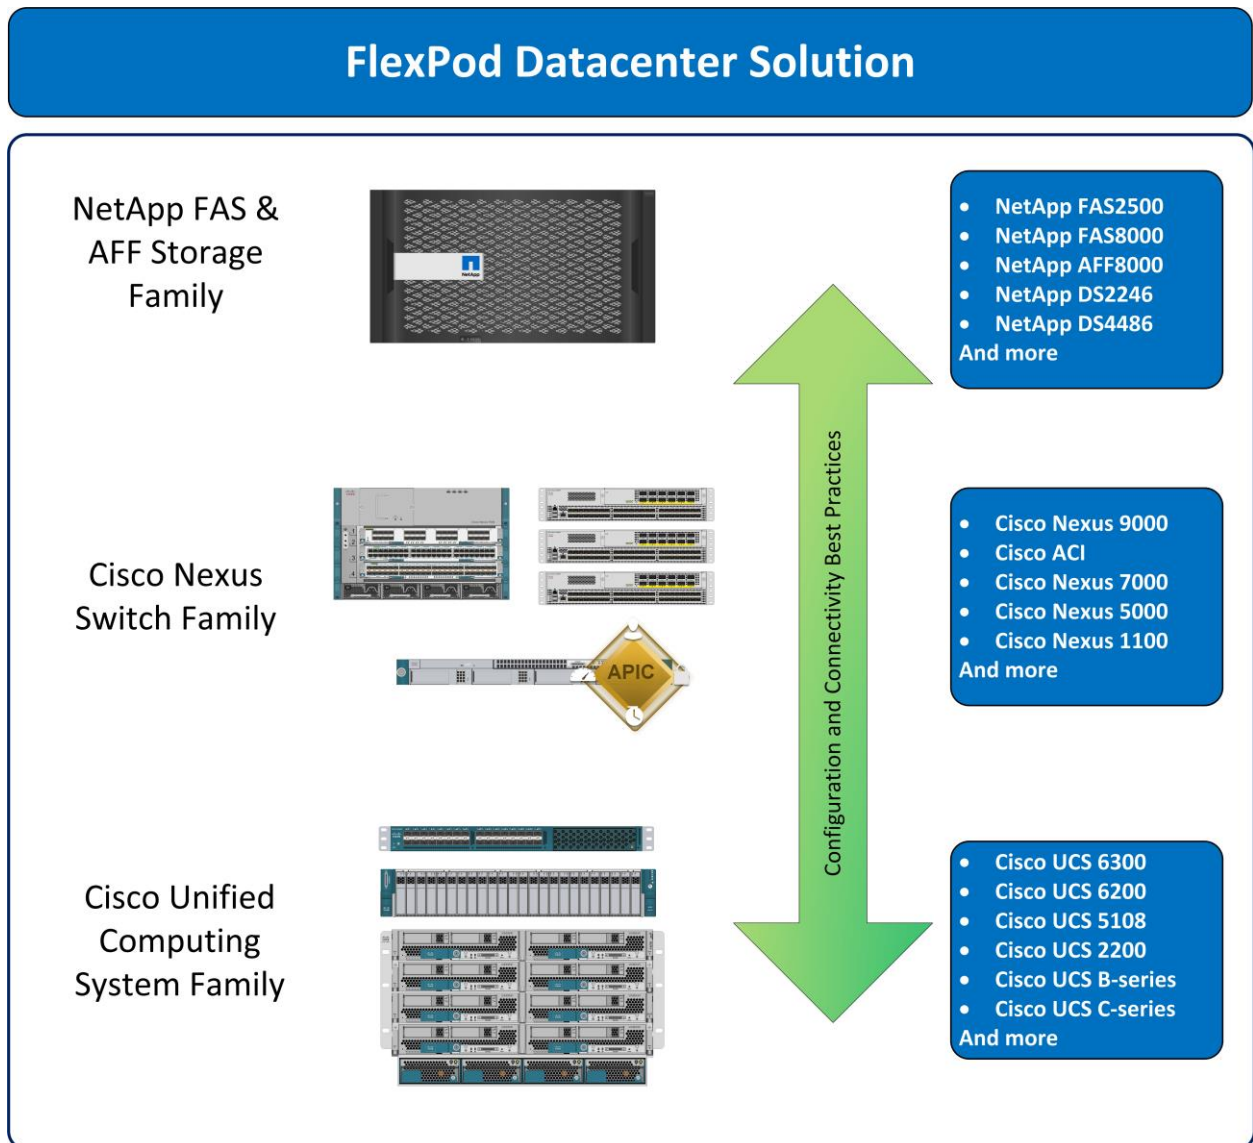
Solution Component	NVA Version 1.0	NVA Version 2.0 (This Document)
Storage		
NetApp controller model	FAS8060	AFF8080cc
NetApp disk size	400GB SSD	800GB SSD
NetApp clustered Data ONTAP	8.2.1	8.3.2
NetApp Virtual Storage Console (VSC)	5.0	6.2
Boot from SAN protocol	FCoE	iSCSI
Datastore access protocol	FCoE	NFS
Networking		
Cisco network switches	Cisco Nexus 5500 series	Cisco Nexus 9000 series
Cisco NX-OS	7.0(0)N1(1)	7.0(3)I2(2a)
Compute		
Cisco UCS Manager	2.2.1c	3.1(1e)
Cisco compute blade model	Cisco UCS B200 M3	Cisco UCS B200 M4
Number of desktop blades	16	16
Virtualization		
VMware vSphere	5.5	6.0 Update 1
VMware vCenter installation type	VMware vCenter Server installed on Windows Server VM	VMware vCenter Server Appliance VM
VMware Horizon View	5.3.1	7.0

Solution Component	NVA Version 1.0	NVA Version 2.0 (This Document)
Operating Systems		
Microsoft Windows client	Windows 7 Update 1	Windows 10
Microsoft Windows Server	Windows 2008 R2	Windows 2012 R2
Solution		
Infrastructure separation	Separate compute clusters and separate storage clusters	Separate compute clusters and the same storage cluster
Desktop persistence type	100% nonpersistent	50% persistent, 50% nonpersistent
Validated solution scale	2,000 desktops	2,500 desktops

2 Program Summary

FlexPod is a predesigned, best practice data center architecture that is built on the Cisco UCS, the Cisco Nexus family of switches, and NetApp AFF series systems. FlexPod can run a variety of virtualization hypervisors as well as bare-metal OSs and enterprise workloads. FlexPod delivers a baseline configuration and can also be sized and optimized to accommodate many different use cases and requirements. Figure 1 lists the component families that make up the FlexPod Datacenter solution.

Figure 1) FlexPod component families.



2.1 FlexPod Program Benefits

NetApp and Cisco have thoroughly validated and verified the FlexPod solution architecture and its many use cases. They have also creating a portfolio of detailed documentation, information, and references to assist you in transforming your data center to this shared infrastructure model. This portfolio includes the following items:

- Best practice architectural design
- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for what is and what is not a FlexPod configuration)
- Frequently asked questions (FAQ)
- Cisco Validated Designs (CVDs) and NetApp Validated Architectures (NVAs) focused on a variety of use cases

NetApp and Cisco have also built a robust and experienced support team focused on FlexPod solutions, from customer account and technical sales representatives to professional services and technical support engineers. This support alliance provides customers and channel services partners with direct access to technical experts who collaborate with cross vendors and have access to shared lab resources to resolve potential issues.

FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment. As a key FlexPod cooperative support partner, VMware provides the virtualization hypervisor and virtual desktop management solution for this verified design with VMware vSphere, VMware vCenter, and VMware Horizon 7.

Integrated System

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

Fabric Infrastructure Resilience

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

Fabric Convergence

FlexPod components are interconnected through the Cisco Unified Fabric network architecture. This architecture supports both traditional LAN traffic and all types of storage traffic, including the lossless requirements for block-level storage transport using Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE). The Cisco Unified Fabric provides high-performance, low-latency, and highly available networks, serving a diverse set of data center needs.

FlexPod uses the Cisco Unified Fabric to offer a wire-once environment that accelerates application deployment. FlexPod also offers efficiencies associated with infrastructure consolidation, including the following:

- Cost savings from the reduction in switches (LAN/SAN switch ports), associated cabling, rack space (capex), and associated power and cooling (opex)
- Migration to faster 10GbE or 40GbE networks and to 100GbE networks in the future
- Evolution to a converged network with little disruption and preservation of investments in the existing infrastructure, management tools, and staff training (expertise)
- Simplified cabling, provisioning, and network maintenance to improve productivity and operational models

Flash-Accelerated Storage

The adoption of flash-accelerated storage is a growing trend in the industry. The benefits gained from flash technologies are well aligned with the needs of shared infrastructures. With shared infrastructures, the benefits of rapid time to market, the ability to scale in consumable increments, reduced risk, and so on are all derived from a standardized approach to infrastructure design. When deploying flash technologies, you should consider the portfolio breadth of the storage vendor. NetApp offers proven technologies such as NetApp Flash Cache™ and NetApp Flash Pool™ intelligent caching and now AFF, so you can be confident that your solution yields reliable performance characteristics for even the most demanding workloads.

Crucially, an all-flash architecture provides predictable and consistent low latency to applications and end users in addition to significantly greater performance and higher performance ceilings. The NetApp AFF8000 series provides the integration and feature richness of the FAS8000 series with even more advanced storage efficiencies, consistent high IOPS, and low-latency performance. Therefore, the AFF8000 series meets or exceeds the capabilities of other options in the all-flash array market.

3 Solution Overview

As you begin your journey toward VDIs, you face a number of questions:

- How do I start the transition?
- What return on investment (ROI) can I expect?
- How do I build a future-proof infrastructure?
- How do I cost-effectively transition from my current infrastructure?
- Will my applications run properly in a virtual desktop environment?
- How do I manage the infrastructure?
- What flash options provide equivalent or better performance than a physical desktop?

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

This document describes VMware vSphere 6.0, VMware Horizon 7, and NetApp AFF with clustered Data ONTAP 8.3.2 built on the FlexPod model from Cisco and NetApp. This document also discusses design choices and best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific releases of the components described in this document, but are also applicable to other versions.

3.1 Target Audience

The intended audience for this document includes sales engineers, field consultants, professional services personnel, IT managers, and partner engineering personnel. This document is also intended for customers who want to take advantage of virtual desktop efficiency to enhance employee productivity and innovation.

3.2 Solution Technology

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

- The Cisco UCS
- Cisco Nexus switches
- NetApp AFF or FAS systems

These components are connected and configured according to the best practices of both Cisco and NetApp and provide the ideal platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed). It can also scale out for environments that need multiple consistent deployments (rolling out additional FlexPod stacks). FlexPod delivers a baseline configuration, and it can also be sized and optimized to accommodate many different use cases.

Typically, the more scalable and flexible a solution is, the more difficult it becomes to maintain a single unified architecture capable of offering the same features and functionality across implementations. This is one of the key benefits of FlexPod. Each of the component families shown in Figure 1 offers platform

and resource options to scale the infrastructure up or down while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

FlexPod addresses four primary design principles: availability, scalability, flexibility, and manageability, as follows:

- **Application availability.** Services are accessible and ready to use.
- **Scalability.** Increasing demands are addressed with appropriate resources.
- **Flexibility.** New services are provided and resources are recovered without infrastructure modification requirements.
- **Manageability.** Efficient infrastructure operations are facilitated through open standards and application programming interfaces (APIs).

FlexPod Networking

Link aggregation technologies play an important role in a FlexPod design, providing improved aggregate bandwidth and link resiliency across the solution stack. NetApp storage controllers, the Cisco UCS fabric interconnects, and Cisco Nexus 9000 switches support active port channeling using the 802.3ad standard Link Aggregation Control Protocol (LACP). Port channeling is a link aggregation technique that provides link fault tolerance and traffic distribution (load balancing) for improved aggregate bandwidth across member ports.

In addition, the Cisco Nexus 9000 series features virtual port channel (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, thereby creating device fault tolerance. vPC addresses aggregate bandwidth, link, and device resiliency. The Spanning Tree protocol does not actively block redundant physical links in a properly configured vPC-enabled environment, so all ports forward the vPC member ports. Cisco UCS fabric interconnects and NetApp AFF controllers benefit from Cisco Nexus vPC abstraction, gaining link and device resiliency and full utilization of a nonblocking Ethernet fabric.

The initial storage configuration of this solution is a two-node HA pair with ONTAP. An HA pair consists of related storage nodes such as the AFF8040 or AFF8080. Scalability is achieved by adding storage capacity (disk and shelves) to an existing HA pair or by adding additional HA pairs into the cluster or storage domain. In both scenarios, the HA interconnect allows each HA node pair to assume control of its partner's storage (disk and shelves) directly.

The local physical HA storage failover capability does not extend beyond the HA pair. Furthermore, a cluster of nodes does not have to include similar hardware. Rather, individual nodes in an HA pair are configured alike, allowing customers to scale as needed as they bring additional HA pairs into the larger cluster.

Network failover is independent of the HA pair construct. Network failover of each node in the cluster is supported by both the back-end cluster interconnect and the front-end switching fabric. This configuration permits cluster, data, and management network interfaces to fail over to different nodes in the cluster, extending beyond the HA pair.

Using clustered Data ONTAP 8.2 or later, NetApp storage systems can be configured to operate without a cluster interconnect switch in a two-node storage system. This configuration is referred to as a two-node switchless cluster, and the verification effort for this design used this configuration.

Cisco Unified Computing System

The Cisco UCS is a next-generation solution for blade and rack server computing. The Cisco UCS is an innovative data center platform that unites compute, network, storage access, and virtualization into a cohesive system designed to reduce total cost of ownership (TCO) and increase business agility. The system integrates a low-latency, lossless 10GbE unified network fabric with enterprise-class, x86-architecture servers.

The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain. Managed as a single system, whether it has two servers or 160 servers with thousands of VMs, the Cisco UCS decouples scale from complexity. The Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

The Cisco UCS consists of the following components:

- [Cisco UCS Manager](#). Provides unified, embedded management of all software and hardware components in the Cisco UCS.
- [Cisco UCS 6200 Series fabric interconnects](#). Are a family of line-rate, low-latency, lossless, 10Gbps Ethernet and FCoE interconnect switches that provide the management and communication backbone for the Cisco UCS.
- [Cisco UCS 5100 Series blade server chassis](#). Support up to eight half-width blade servers and two fabric extenders in a six rack-unit (RU) enclosure.
- [Cisco UCS B-Series](#) Intel-based [blade servers](#). Increase performance, efficiency, versatility, and productivity.
- [Cisco UCS C-Series rack mount servers](#). Deliver unified computing in an industry-standard form factor to reduce TCO and increase agility.
- [Cisco UCS adapters](#). Provide a wire-once architecture that offers a range of options to converge the fabric, optimize virtualization, and simplify management.

For more information about the Cisco UCS, see the Cisco [Servers - Unified Computing](#) site.

Cisco UCS Manager

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

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

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

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

Service profiles benefit both virtualized and nonvirtualized environments. These profiles increase the mobility of nonvirtualized servers, when moving workloads from server to server or taking a server offline

for service or upgrade, for example. Profiles can also be used in conjunction with virtualization clusters to bring new resources online easily, complementing existing VM mobility.

In the latest Cisco UCS Manager release (3.1), Cisco now provides both the Java GUI and a new HTML5 GUI for systems management, as can be seen Figure 3 and Figure 3.

Note: Remote console access still requires Java under Cisco UCS Manager 3.1.

Figure 2) Cisco UCS Manager: Java client.

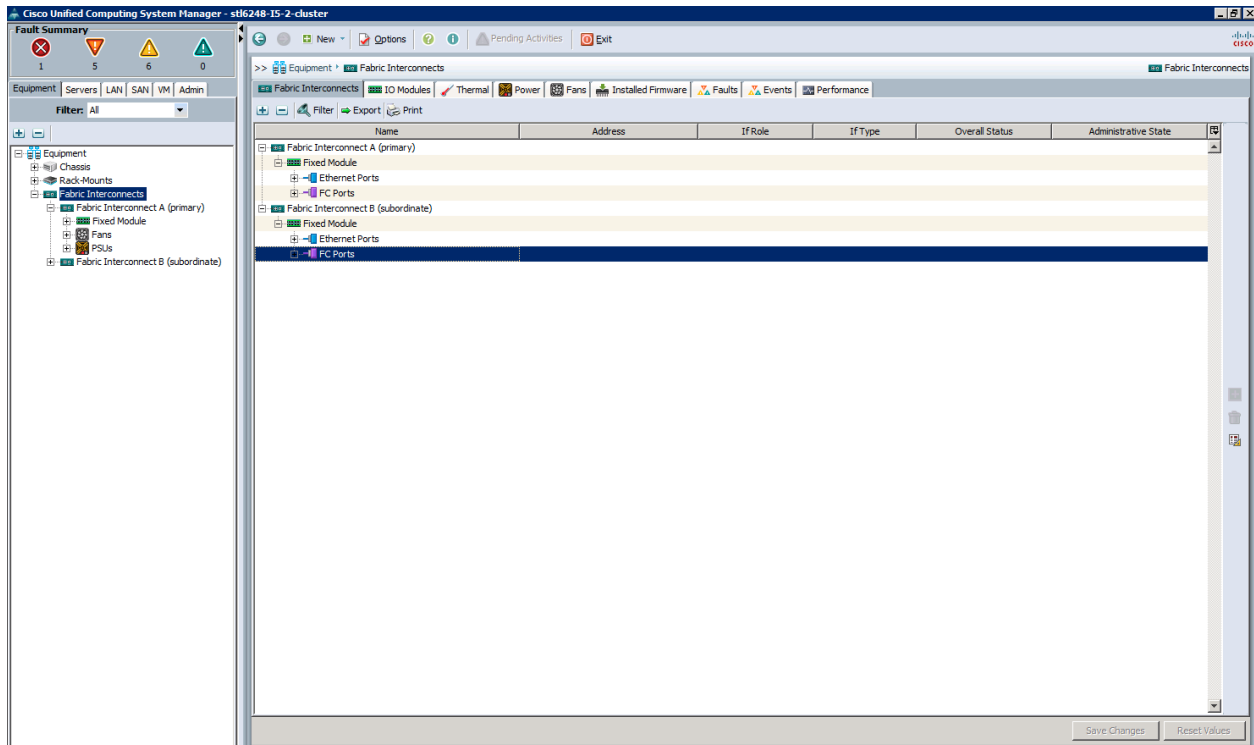
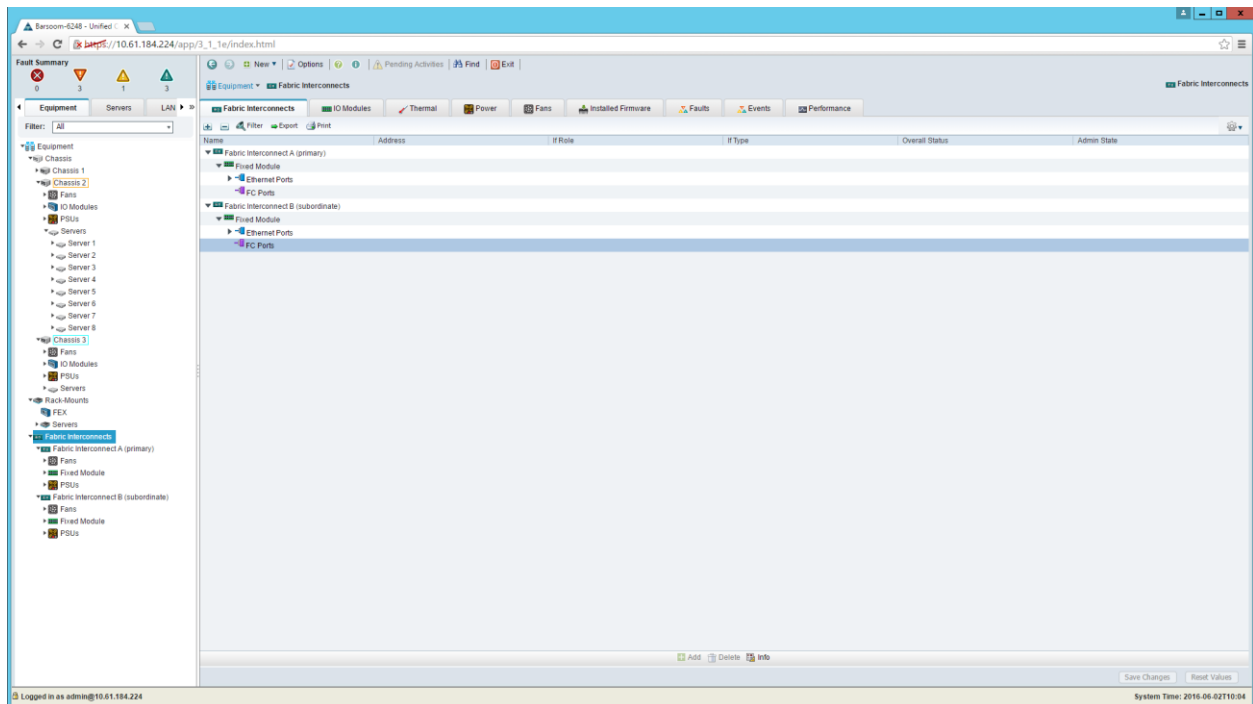


Figure 3) Cisco UCS Manager: HTML5 client.



For more information about Cisco UCS Manager, refer to the [Cisco UCS Manager](#) site.

Cisco Nexus 9000 Series Switches

The Cisco Nexus 9000 Series is designed for data center environments with cut-through switching technology that enables consistent low-latency Ethernet solutions. With front-to-back or back-to-front cooling, the Cisco Nexus 9000 Series possesses data ports in the rear, which brings switching into close proximity with servers and makes cable runs short and simple. This switch series is highly serviceable, with redundant, hot-pluggable power supplies and fan modules. It uses data center-class Cisco NX-OS software for high reliability and ease of management.

The Cisco Nexus 9000 platform extends the industry-leading versatility of the Cisco Nexus Series purpose-built, 10GbE data center-class switches and provides innovative advances toward higher density, lower latency, and multilayer services. The Cisco Nexus 9000 platform is well suited for enterprise data center deployments across a diverse set of physical, virtual, and high-performance computing environments. Cisco Nexus 9000 switches provide 40Gb switching capability and can participate in Cisco Application Centric Infrastructure. However, they do not support FC or FCoE storage protocols. To support these protocols, FlexPod supports FC and FCoE connections directly between the Cisco UCS fabric interconnects and the NetApp AFF storage system, as described in the appendix.

The switch used in the validation of this FlexPod architecture is the Cisco Nexus 9396PX. This switch has the following specifications:

- A two-rack unit, 1/10/40GbE switch
- Forty-eight fixed 1/10GbE ports on the base chassis and one expansion slot supporting up to 12 fixed 40GbE ports
- Throughput of up to 1.92Tbps

Other Cisco Nexus 9000 switches, such as the Cisco Nexus 9372, are also suitable for this architecture. Cisco Nexus 9396PX switches were used for this validation due to inventory availability. However, its use is not a specific requirement for this solution.

For more information, see the [Cisco Nexus 9000 Series Switches](#) site.

NetApp All Flash FAS

A product of more than 20 years of innovation, ONTAP has evolved to meet the changing needs of customers and help drive their success. NetApp ONTAP provides a rich set of data management features and clustering for scale-out, operational efficiency, and nondisruptive operations. This storage OS offers customers one of the most compelling value propositions in the industry. The IT landscape is undergoing a fundamental shift to IT as a service, a model that requires a pool of compute, network, and storage resources that serve a wide range of applications and deliver a wide range of services. Innovations such as ONTAP are fueling this revolution.

NetApp solutions are user friendly, easy to manage, and quick to deploy and offer increased availability while consuming fewer IT resources. This means that they dramatically lower lifetime TCO. Whereas others manage complexity, NetApp eliminates it. A NetApp solution includes hardware in the form of controllers and disk storage and the ONTAP software.

NetApp offers the NetApp Unified Storage Architecture. The term “unified” refers to a family of storage systems that simultaneously support SAN and NAS across many operating environments such as VMware, Windows, and UNIX. This single architecture provides access to data by using industry-standard protocols, including NFS, CIFS, iSCSI, FCP, and FCoE. Connectivity options include standard Ethernet (100/1000 or 10GbE) and FC (2, 4, 8, or 16Gbps).

This FlexPod Datacenter solution includes the NetApp AFF8000 series unified scale-out storage systems. Powered by ONTAP, the AFF8000 series unifies SAN and NAS storage infrastructures. The AFF8000 features a multiprocessor Intel chipset and leverages high-performance memory modules, NVRAM to accelerate and optimize writes, and an I/O-tuned PCIe gen3 architecture that maximizes application throughput. The AFF8000 series comes with integrated UTA2 ports that support 16Gb FC, 10GbE, or FCoE.

If storage requirements change over time, NetApp storage offers the flexibility to change quickly, as needed, and without expensive and disruptive forklift upgrades. For example, a LUN can be changed from FC access to iSCSI access without moving or copying the data. Only a simple dismount of the FC LUN and a mount of the same LUN using iSCSI is required. In addition, a single copy of data can be shared between Windows and UNIX systems while allowing each environment to access the data through native protocols and applications.

NetApp storage solutions provide redundancy and fault tolerance through clustered storage controllers, hot-swappable redundant components (such as cooling fans, power supplies, disk drives, and shelves), and multiple network interfaces. This highly available and flexible architecture enables customers to manage all data under one common infrastructure while achieving mission requirements. The NetApp Unified Storage Architecture allows data storage with higher availability and performance, easier dynamic expansion, and greater ease of management than any other solution.

Outstanding Performance

The NetApp AFF solution shares the same unified storage architecture, ONTAP software, management interface, rich data services, and advanced feature set as the rest of the FAS product families. This unique combination of all-flash media with ONTAP delivers the consistent low latency and high IOPS of all-flash storage, with the industry-leading ONTAP software. In addition, it offers proven enterprise availability, reliability, and scalability; storage efficiency proven in thousands of VDI deployments; unified storage with multiprotocol access; and advanced data services. ONTAP also provides operational agility through tight application integrations.

Optimized Writes

The NetApp WAFL® (Write Anywhere File Layout) file system enables NetApp to process writes efficiently. When the ONTAP software receives I/O, it stores the I/O in battery-backed nonvolatile RAM (NVRAM) in both controllers of the HA pair and sends back an acknowledgement (or ACK), notifying the sender that the write is committed. Acknowledging the write before committing to disk provides a very low response time for write I/O and thus low write latency. This architecture also allows ONTAP to perform many functions to optimize the data layout for optimal write/write coalescing. Before being written to disk, I/Os are coalesced into larger blocks because larger sequential blocks require less CPU for each operation.

With the AFF8000 series, starting with Data ONTAP 8.3.1 and continuing with 8.3.2, additional write optimizations have been introduced to take specific advantage of an all-SSD storage back end. These optimizations reduce the number of internal steps required for write operations to produce substantial performance increases relative to older releases of ONTAP on the same physical equipment.

Enhancing Flash

NetApp ONTAP has leveraged flash technologies since 2009 and has supported SSDs since 2010. This relatively long experience with SSDs has allowed NetApp to tune ONTAP features to optimize SSD performance and enhance flash media endurance.

As discussed in previous sections, SSDs are not in the critical write path because ONTAP acknowledges writes after they are in dynamic random-access memory (DRAM) and logged to NVRAM. Therefore, write latencies are very low. ONTAP also enables efficient use of SSDs when destaging cache by coalescing writes into a single sequential stripe across all SSDs at once. ONTAP writes to free space whenever possible, minimizing overwrites for every dataset, not only for deduped or compressed data.

This wear-leveling feature of ONTAP is native to the architecture, and it also leverages the wear-leveling and garbage-collection algorithms built into SSDs to extend device life. Therefore, NetApp provides up to a five-year warranty with all SSDs (three-year standard warranty, plus the offer of an additional two-year extended warranty, with no restrictions on the number of drive writes).

The parallelism built into ONTAP, combined with the multicore CPUs and large system memories in the NetApp AFF8000 storage controllers, takes full advantage of SSD performance.

NetApp ONTAP

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

NetApp ONTAP scale-out is one way to respond to growth in a storage environment. All storage controllers have physical limits to their expandability. The number of CPUs, memory slots, and space for disk shelves dictates the maximum capacity and controller performance. If more storage or performance capacity is needed, it might be possible to add CPUs and memory or install additional disk shelves. However, ultimately the controller becomes completely populated, with no further expansion possible. At this stage, the only option is to acquire another controller.

If the original controller must be completely replaced by a newer and larger controller, data migration is required to transfer the data from the old controller to the new one. This process is time consuming and potentially disruptive and most likely requires configuration changes on all of the attached host systems.

If the newer controller can coexist with the original controller, you now have two storage controllers that must be individually managed. However, there are no native tools that can balance or reassign workloads across them. The situation becomes even more difficult as the number of controllers increases. If the scale-up approach is used, the operational burden increases consistently as the environment grows, and

the end result is a very unbalanced and difficult-to-manage environment. Technology refresh cycles require substantial planning in advance, lengthy outages, and configuration changes, which introduce risk into the system.

Scale-Out

In contrast, the use of scale-out means that as the storage environment grows, additional controllers are added seamlessly to the resource pool residing on a shared storage infrastructure. Host and client connections as well as datastores can move seamlessly and nondisruptively anywhere in the resource pool. Therefore, existing workloads can be easily balanced over the available resources, and new workloads can be easily deployed. Technology refreshes (replacing disk shelves or adding or completely replacing storage controllers) are accomplished in an environment that remains online and continues serving data.

The benefits of scale-out include the following:

- Nondisruptive operations
- The ability to add additional workloads with no effect on existing services
- Operational simplicity and flexibility

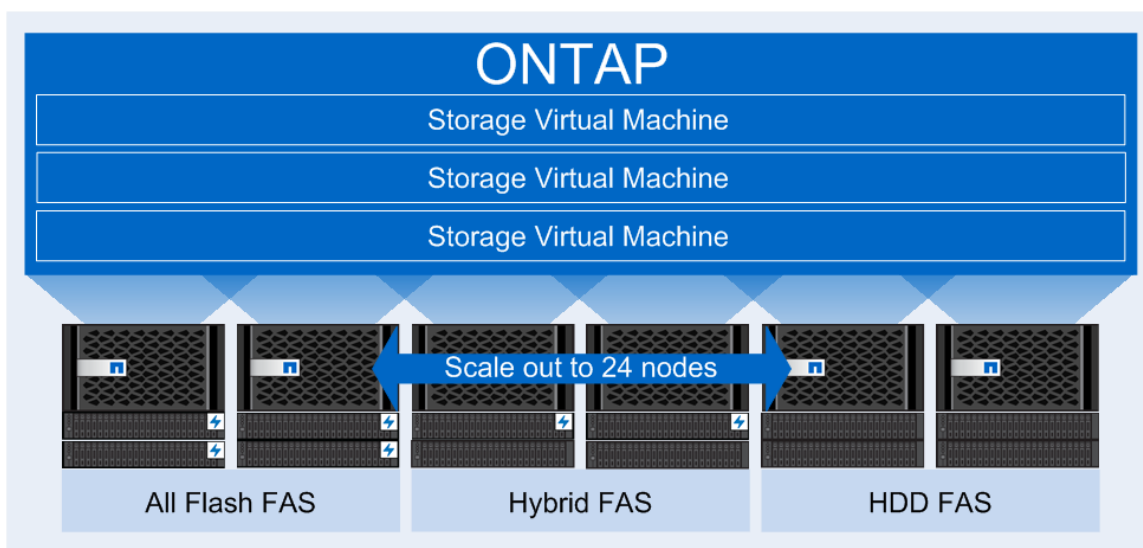
Although scale-out products have been available for some time, these products were typically subject to one or more of the following shortcomings:

- Limited protocol support (NAS only)
- Limited hardware support (supported only a particular type of storage controller or a very limited set)
- Little or no storage efficiency (thin provisioning, deduplication, or compression)
- Little or no data replication capability

Therefore, although these products are well positioned for certain specialized workloads, they are less flexible, less capable, and not robust enough for broad deployment throughout the enterprise.

As depicted in Figure 4, NetApp ONTAP is the first product to offer a complete scale-out solution with an adaptable, always-available storage infrastructure for today's highly virtualized environment. An ONTAP system can scale up to 24 nodes, depending on platform and protocol, and can contain different disk types and controller models in the same storage cluster.

Figure 4) ONTAP.



Note: Storage virtual machines (SVMs) were formerly known as Vservers.

Nondisruptive Operations

The move to a shared infrastructure has made it nearly impossible to schedule downtime to accomplish routine maintenance. NetApp ONTAP is designed to eliminate the planned downtime needed for maintenance operations and lifecycle operations as well as the unplanned downtime caused by hardware and software failures.

Three standard tools make this elimination of downtime possible:

- **NetApp DataMotion™ for Volumes (vol move).** Allows data volumes to be moved from one aggregate to another on the same or a different cluster node.
- **Logical interface (LIF) migrate.** Allows the physical Ethernet interfaces in ONTAP to be virtualized. LIF migrate also allows LIFs to be moved from one network port to another on the same or a different cluster node.
- **Aggregate relocate (ARL).** Allows complete aggregates to be transferred from one controller in an HA pair to the other without data movement.

Used individually and in combination, these tools enable you to nondisruptively perform a full range of operations, from moving a volume from a faster to a slower disk all the way up to a complete controller and storage technology refresh.

As storage nodes are added to the system, all physical resources—CPUs, cache memory, network I/O bandwidth, and disk I/O bandwidth—can easily be kept in balance. NetApp ONTAP systems enable users to:

- Add or remove storage shelves (over 23PB in an 8-node cluster and up to 69PB in a 24-node cluster).
- Move data between storage controllers and tiers of storage without disrupting users and applications.
- Dynamically assign, promote, and retire storage while providing continuous access to data as administrators upgrade or replace storage.

These capabilities allow administrators to increase capacity while balancing workloads and can reduce or eliminate storage I/O hot spots without the need to remount shares, modify client settings, or stop running applications.

Availability

Shared storage infrastructure provides services to thousands of virtual desktops. In such environments, downtime is not an option. The NetApp AFF solution eliminates sources of downtime and protects critical data against disaster through two key features:

- **High availability.** A NetApp HA pair provides seamless failover to its partner in case of hardware failure. Each of the two identical storage controllers in the HA pair configuration serves data independently during normal operation. During an individual storage controller failure, the data service process is transferred from the failed storage controller to the surviving partner.
- **NetApp RAID DP® data protection technology.** During any virtualized desktop deployment, data protection is critical because any RAID failure might disconnect hundreds to thousands of end users from their desktops, resulting in lost productivity. RAID DP provides performance comparable to that of RAID 10, and yet it requires fewer disks to achieve equivalent protection. In contrast to RAID 5, RAID DP provides protection against double disk failure, which can protect against only one disk failure per RAID group. Therefore, RAID DP provides RAID 10 performance and protection at a RAID 5 price point.

NetApp Advanced Data Management Capabilities

This section describes the storage efficiencies, multiprotocol support, VMware integrations, and replication capabilities of the NetApp AFF solution.

Storage Efficiencies

Most desktop virtualization implementations deploy thousands of desktops from a small number of golden VM images, resulting in large amounts of duplicate data. This is especially the case with the VM OS.

The NetApp AFF solution includes built-in thin provisioning, inline and postprocess data deduplication, inline and postprocess data compression, and zero-cost cloning with NetApp FlexClone® data replication technology. These features offer multilevel storage efficiency across virtual desktop data, installed applications, and user data. This comprehensive storage efficiency package enables a significantly reduced storage footprint for virtualized desktop implementations, with a capacity reduction of up to 10:1, or 90%. This analysis is based on existing customer deployments and NetApp solutions lab verification.

Several features make this level of storage efficiency possible:

- **Thin provisioning.** Allows multiple applications to share a single pool of on-demand storage. This feature eliminates the need to provision more storage for a particular application if another application still has plenty of allocated but unused storage.
- **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 the user, and it can be enabled and disabled on the fly. With Data ONTAP 8.3.2, inline deduplication of in-memory data is enabled by default, and postprocess deduplication is also available. To eliminate any potential concerns about postprocess deduplication causing additional wear on the SSDs, NetApp provides up to a five-year warranty for all SSDs (three-year standard plus an additional two-year extended warranty), with no restrictions on the number of drive writes.
- **Inline compression.** Data compression reduces the disk space required, regardless of storage protocol, application, or storage tier. Inline compression also reduces the data that must be moved to SSDs, thereby reducing the wear on SSDs.
- **FlexClone.** Offers hardware-assisted rapid creation of space-efficient, writable, point-in-time images of individual VM files, LUNs, or flexible volumes. It is fully integrated with VMware vSphere vStorage APIs for Array Integration (VAAI). The use of FlexClone technology in VDI deployments provides high levels of scalability and significant cost, space, and time savings. Both file-level cloning and volume-level cloning are tightly integrated with the VMware vCenter Server through the NetApp Virtual Storage Console Provisioning and Cloning vCenter plug-in and native VM cloning offload with VMware VAAI. The NetApp VSC provides the flexibility to rapidly provision and redeploy thousands of VMs with hundreds of VMs in each datastore.

Advanced Storage Features

NetApp ONTAP provides a number of additional features that can be leveraged in a virtual desktop environment, whether for the infrastructure supporting the desktops or the desktops themselves. Some of these features are:

- **NetApp Snapshot® copies.** Manual or automatically scheduled point-in-time copies that write only changed blocks, with no performance penalty. Snapshot copies consume minimal storage space because only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- **Compression.** Compression of data blocks on disk to provide space savings instead of or in addition to those obtained with deduplication.

- **LIFs.** A LIF is a logical interface that is associated with a physical port, interface group (ifgrp), or VLAN interface. More than one LIF can be associated with a physical port at the same time. There are three types of LIFs:
 - NFS LIF
 - iSCSI LIF
 - FC LIF

LIFs are logical network entities that have the same characteristics as physical network devices but are not tied to physical objects. LIFs used for Ethernet traffic are assigned specific Ethernet-based details such as IP addresses and iSCSI-qualified names and are then associated with a specific physical port capable of supporting Ethernet. LIFs used for FC-based traffic are assigned specific FC-based details such as worldwide port names and are then associated with a specific physical port capable of supporting FC or FCoE. NAS LIFs can be nondisruptively migrated to any other physical network port throughout the entire cluster at any time, either manually or automatically (by using policies). SAN LIFs rely on multipath input/output (MPIO) and asymmetric logical unit access (ALUA) to notify clients of any change in the network topology.

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

Multiprotocol Support

By supporting all common NAS and SAN protocols on a single platform, NetApp Unified Storage enables the following functions:

- Direct access to storage by each client
- Network file sharing across different platforms without the need for protocol-emulation products such as SAMBA, NFS Maestro, or PC-NFS
- Simple and fast data storage and data access for all client systems
- Fewer storage systems
- Greater efficiency from each system deployed

ONTAP can support several protocols concurrently in the same storage system. Data ONTAP 7G and Data ONTAP operating in 7-Mode also include support for multiple protocols. Unified storage is important to VMware Horizon View solutions, such as CIFS/SMB for user data, NFS or SAN for the VM datastores, and guest-connect iSCSI LUNs for Windows applications.

The following protocols are supported:

- NFS v3, v4, and v4.1 (including pNFS)
- iSCSI
- FC
- FCoE
- CIFS

VMware vSphere Integrations

The complexity of deploying and managing thousands of virtual desktops can be daunting without the right tools. NetApp Virtual Storage Console for VMware vSphere is tightly integrated with VMware vCenter for rapidly provisioning, managing, configuring, and backing up a VMware Horizon View implementation. NetApp VSC is a VMware vCenter Server plug-in that provides end-to-end VM lifecycle management for VMware vSphere environments using NetApp storage. This integration significantly

increases operational efficiency and agility by simplifying the deployment and management process for thousands of virtual desktops.

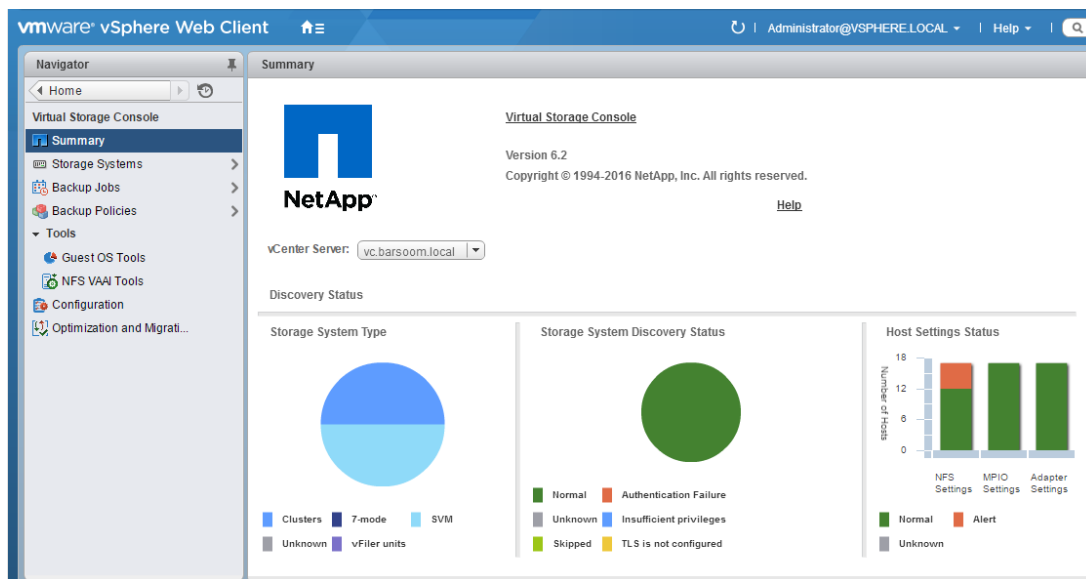
By using VSC for VMware vSphere, the following key NetApp capabilities can be executed from the vSphere client or as part of VMware Horizon View Composer maintenance operations:

- **Monitoring and host configuration.** Provides a view of the NetApp storage environment from a VMware administrator's perspective and optimizes storage and host configurations.
- **Backup and restore.** Automates data protection processes with policy-based backup and restore using NetApp Snapshot and FlexClone technologies.
- **Provisioning and cloning.** Delivers end-to-end datastore provisioning, rapid server and desktop VM cloning, and flexible redeployment services.
- **Optimization and migration.** Detects VM misalignments and optimizes performance by enabling online VM I/O optimization and VM migration.

NetApp VSC also includes an API for automated control. The NetApp VSC software can be installed on a separate Microsoft Windows Server instance or VM or as a standalone vApp.

Figure 5 depicts an example of the NetApp VSC.

Figure 5) NetApp VSC example.



Replication

The backup and recovery capability of NetApp VSC is a unique, scalable, integrated data protection solution for persistent desktop VMware Horizon View environments. The backup and recovery plug-in allows customers to leverage VMware snapshot functionality with NetApp array-based, block-level Snapshot copies to provide consistent backups for the virtual desktops. The backup and recovery plug-in is integrated with NetApp SnapMirror® replication technology, which preserves the deduplicated storage savings from the source on the destination storage array. Therefore, you do not need to rerun deduplication on the destination.

SnapMirror can be used to replicate between any disk type or tier between NetApp AFF and FAS systems, including SSD to SAS, SAS to SATA, SATA to SSD, or any other combination. SnapMirror also covers cascading mirrors on different tiers of storage. An AFF can replicate to another AFF, to a hybrid FAS, or to a hard drive-only FAS, providing customers with cost-effective performance options for business requirements around data protection and DR.

When a VMware Horizon View environment is replicated with SnapMirror, the replicated data can quickly be brought online to provide production access during a site or data center outage. In addition, SnapMirror is fully integrated with VMware Site Recovery Manager and NetApp FlexClone technology. This integrated system can instantly create zero-cost writable copies of the replicated virtual desktops at the remote site that can be used for DR testing or for test and development work.

VMware vSphere

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

VMware vSphere provides revolutionary benefits, but with a practical, nondisruptive evolutionary process for legacy applications. Existing applications can be deployed on VMware vSphere with no changes to the application or the OS on which they run.

VMware vSphere provides a set of application services that enable applications to achieve unparalleled levels of availability and scalability. VMware vSphere delivers the following core capabilities to meet numerous application and enterprise demands:

- **Availability.** Workload mobility is provided through vMotion. HA is provided through vSphere fault domain manager technology, offering VM resiliency in the event of physical server or guest OS failures.
- **Automation.** The VMware Distributed Resource Scheduler (DRS) offers dynamic workload distribution to align resource utilization with business priorities and compute capacity. DRS provides efficient use of compute resources and thus power consumption.
- **Compute.** The VMware vSphere ESXi hypervisor provides efficient memory, storage, and compute abstraction through the use of VMs.
- **Network.** VMware vSphere supports third-party virtual distributed switches such as the Cisco Nexus 1000v, providing a resilient and fully integrated virtualized network access layer.
- **Storage.** Thin provisioning enables overprovisioning of storage resources to improve storage utilization and improve capacity planning. The Virtual Machine File System (VMFS) is a clustered file system allowing multiple hosts simultaneous read and write access to a single volume located on a SCSI-based device through FC, FCoE, or iSCSI. VMFS supports the connection of a maximum of 64 hosts to a single volume of up to 64TB in size.

Figure 6 provides an overview of VMware vSphere capabilities.

Figure 6) VMware vSphere feature overview (graphic supplied by VMware).



VMware vSphere delivers a robust application environment. For example, with VMware vSphere, all applications can be protected from downtime with VMware HA without the complexity of conventional clustering. In addition, applications can be scaled dynamically to meet changing loads with capabilities such as hot add and VMware DRS.

For more information, see the [VMware vSphere](#) product site.

VMware Horizon 7

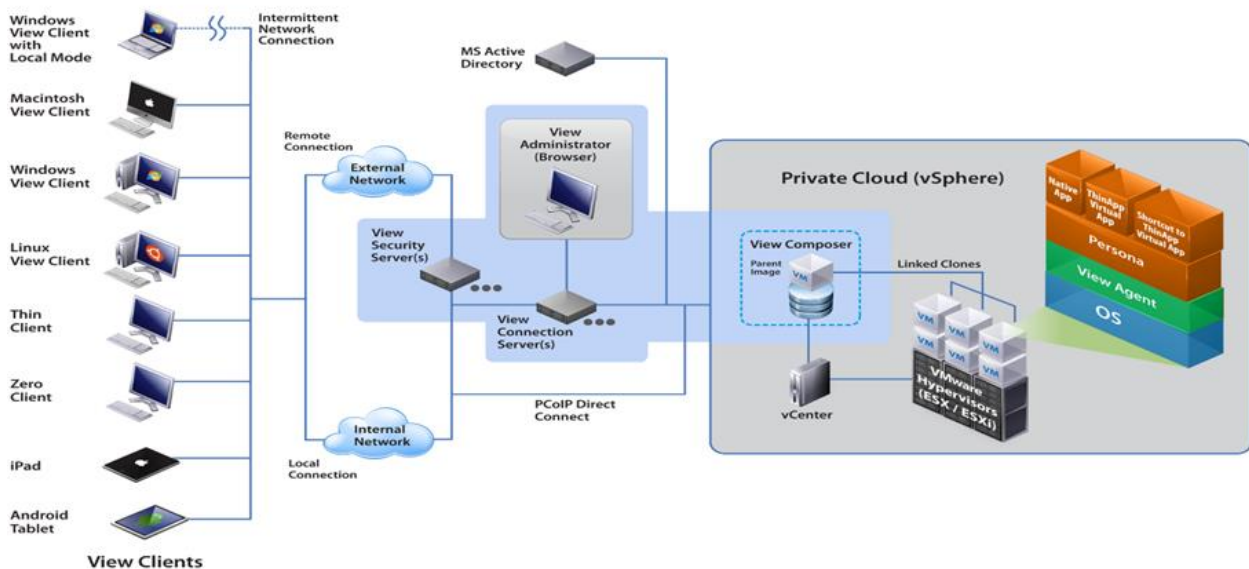
VMware Horizon 7 is an enterprise-class desktop virtualization solution that delivers virtualized or remote desktops and applications to end users through a single platform. VMware Horizon 7 allows IT to manage desktops, applications, and data centrally while increasing flexibility and customization at the endpoint for the user. It enables levels of availability and agility to desktop services that are unmatched by traditional PCs at about half the TCO per desktop.

VMware Horizon 7 is a tightly integrated, end-to-end solution built on the industry-leading virtualization platform VMware vSphere. Figure 7 provides an architectural overview of a VMware Horizon 7 deployment that includes seven main components:

- **View Connection Server.** Streamlines the management, provisioning, and deployment of virtual desktops by acting as a broker for client connections, authenticating and directing incoming user desktop requests. Administrators can centrally manage thousands of virtual desktops from a single console. End users connect through the View Connection Server to securely and easily access their personalized virtual desktops.
- **View Security Server.** An instance of the View Connection Server that adds an additional layer of security between the Internet and the internal network.

- **View Composer Server.** An optional feature that allows you to manage pools of linked-clone desktops by creating master images that share a common virtual disk.
- **View Agent Service.** Communicates between VMs and the Horizon client. View Agent is installed on all VMs managed by vCenter Server so that the View Connection Server can communicate with them. View Agent also provides features such as connection monitoring, virtual printing, persona management, and access to locally connected USB devices. View Agent is installed in the guest OS.
- **Horizon clients.** Can be installed on each endpoint device. End users can access their virtual desktops from devices such as zero clients, thin clients, Windows PCs, Macs, and iOS-based and Android-based mobile devices. Horizon Clients are available for Windows, Mac, Ubuntu, Linux, iOS, and Android to provide the connection to remote desktops from the device of choice.
- **User environment management and personalization.** VMware User Environment Manager offers personalization and dynamic policy configuration across any virtual, physical, and cloud-based environment:
 - Simplifies end-user profile management by providing organizations with a single and scalable solution that leverages existing infrastructure
 - Provides end users with quick access to a Windows workspace and applications with a personalized and consistent experience across devices and locations
- **Real-time application delivery and management.** Supports the following functionality:
 - Easily package applications to avoid compatibility issues
 - Instantly provision applications at scale
 - Dynamically attach applications to users, groups, or devices, even when users are logged on to their desktop
 - Provision, deliver, update, and retire applications in real time

Figure 7) VMware Horizon View deployment (graphic supplied by VMware).



VMware View Connection Server

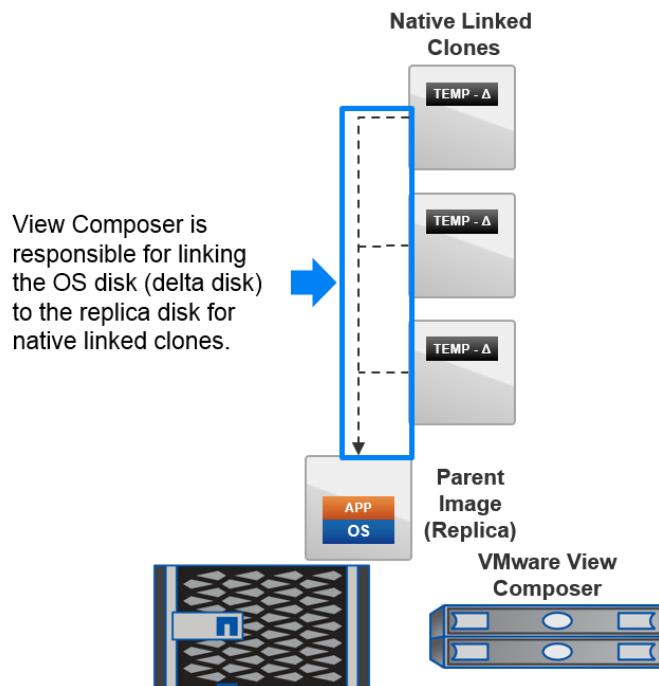
The VMware View Connection Server is responsible for provisioning and managing virtual desktops and for brokering the connections between clients and the virtual desktop machines.

VMware View Composer

The VMware View Composer server is a critical component of solutions that use VMware Horizon linked clones. This server is responsible for the creation and maintenance operations of VMware Horizon linked clones. It works with the View Connection Server to rapidly provision storage-efficient virtual desktops for use in the VMware Horizon desktop environment. These linked-clone desktops created by the Composer can be either dedicated or floating virtual desktops in an automated pool. For this reference architecture, dedicated desktops in an automated pool were created.

The Composer server is also involved during maintenance operations, such as refresh, recompose, and rebalance. These operations improve the storage efficiency, performance, security, and compliance of the virtual desktop environment. Figure 8 shows a VMware Horizon linked clone using VMware View Composer.

Figure 8) VMware Horizon View linked clone using View Composer.



The Composer server can be installed on a VMware vCenter Server or as a standalone server, excluding any servers participating in the VMware Horizon environment, such as the Connection Server, the transfer server, the security server, and so on. For this reference architecture, the Composer server was installed on a separate VM.

3.3 Use Case Summary

The FlexPod Datacenter with VMware Horizon and NetApp AFF solution architecture provides a flexible framework on which to deploy enterprise-class virtual desktop infrastructures. This solution applies to the following use cases:

- Deploying VDI for up to 2,500 users within a single FlexPod platform
- Deploying VDI for an even larger number of users based on the virtual desktop pool of desktops (POD) design

4 Technology Components

This section covers the technology components for the FlexPod Datacenter with VMware Horizon and NetApp AFF solution.

4.1 Hardware Components

During solution testing, Cisco UCS blade servers were used to host the infrastructure and the desktop VMs. The desktops and infrastructure servers were hosted on discrete resources so that the workload to the NetApp AFF system could be precisely measured. It is a NetApp and industry best practice to use separate compute resources for the desktop VMs and the infrastructure VMs because noisy neighbors or bully virtual desktops can affect the infrastructure. Such problems can negatively affect all users, applications, and performance results. Although performance of the AFF platform is sufficient to host the infrastructure storage volumes, a separate NetApp FAS storage system (not shown) was used to host the launcher VMs and LoginVSI infrastructure. This is a typical configuration for a customer environment.

Table 2 lists the hardware components used to implement the solution for validation in the NetApp labs. The hardware components used in any particular implementation of the solution can vary based on customer requirements.

Table 2) Hardware components.

Hardware	Configuration
Cisco UCS 6200 Series fabric interconnects	FI 6248UP
Cisco UCS B200 blades	B200 M4 using Cisco UCS VIC 1340
Cisco UCS 5108 chassis	Cisco UCS-IOM 2204XP
Cisco Nexus 9000	Cisco Nexus 9396PX
NetApp AFF8000	AFF8080cc
NetApp DS2246 disk shelves	Disk shelves populated with 800G SSDs

4.2 Software Components

Table 3 lists the software components used to implement the solution. The software components used in any particular implementation of the solution may vary based on customer requirements.

Table 3) Solution software components.

Software/Firmware	Version
Compute	
Cisco UCS Manager	3.1(1e)
Networking	
Cisco NX-OS	7.0(3)I2(2a)
Storage	
NetApp clustered Data ONTAP	8.3.2
NetApp VSC	6.2
VMware vSphere	

Software/Firmware	Version
VMware ESXi	6.0.0, 3380124
VMware vCenter Server	6.0.0, 3339084
VMware Horizon View	
VMware Horizon View Administrator	7.0.0, 3633490
VMware View Composer	7.0.0, 3613429
VMware Horizon View Client	4.0.1, 3698521
VMware Horizon View Agent	7.0.0, 3633490
Virtual Desktop	
Windows 10	Enterprise 32-bit
Microsoft Office 2016	Professional 32-bit
Database Server	
Microsoft SQL Server	2012 R2 (64-bit)
Microsoft SQL Server Native Client	11.0 (64-bit)

5 Solution Design

The FlexPod Datacenter with VMware Horizon on AFF solution consists of the following designs:

- The Cisco UCS
- Cisco Nexus network switches
- NetApp AFF storage
- VMware vSphere
- VMware Horizon View

5.1 Cisco UCS Design

The FlexPod design simultaneously supports both B-Series and C-Series deployments. This reference architecture only utilizes B-Series servers in its design due to their greater rack density and easier scalability. If you want to implement a C-Series VDI design, you must reconsider host sizing, density, and related data center costs. However, the remainder of the architecture outlined in this document would not change substantially.

Cisco UCS: B-Series Server Design

The Cisco UCS supports the virtual server environment by providing a robust, highly available, and readily manageable compute resource. The components of the Cisco UCS system offer physical redundancy and a set of logical structures to deliver a very resilient FlexPod compute domain. In this verification effort, the service profiles of multiple Cisco UCS B-Series servers are SAN booted through iSCSI as VMware ESXi nodes. The ESXi nodes consisted of Cisco UCS B200-M4 blades with Cisco 1340 VIC adapters. These nodes were allocated to a VMware DRS and HA-enabled cluster supporting infrastructure services such as vSphere Virtual Center, Microsoft Active Directory (AD), and database services.

The Cisco 1340 VIC presents four virtual PCIe devices to the ESXi node. Two virtual 10GbE NICs (vNICs) are used for iSCSI boot, and two additional vNICs are available for management, vMotion, VMs, and VM datastore access. The ESXi OS is unaware that these are virtual adapters. More vNICs could be used to provide more granular isolation of traffic, particularly for vMotion and datastore access, including QoS at the vNIC level. Increasing the number of vNICs does not increase the total aggregate bandwidth available to each blade, although it does slightly increase the complexity of the design. Balancing complexity versus benefits is a key consideration at all levels of an infrastructure design. Organizations have the flexibility to make minor alterations, such as the number of vNICs, without significantly affecting the remainder of the design.

FlexPod allows organizations to adjust the individual components of the system to meet their particular scale or performance requirements. One key design decision in the Cisco UCS domain is the selection of I/O components. There are various combinations of I/O adapters, Cisco UCS extender I/O modules (IOM), and Cisco UCS fabric interconnects available. Therefore, it is important to understand the effect that these selections have on the overall flexibility, scalability, and resiliency of the fabric.

There are two available backplane connections in the Cisco UCS 5100 series chassis. Each of the two Cisco UCS fabric extender IOMs has either four or eight 10GBASE KR (802.3ap) standardized Ethernet backplane paths available for connection to the half-width blade slot. This means that each half-width slot has the potential to support up to 80Gb of aggregate traffic. The level of performance realized depends on several factors:

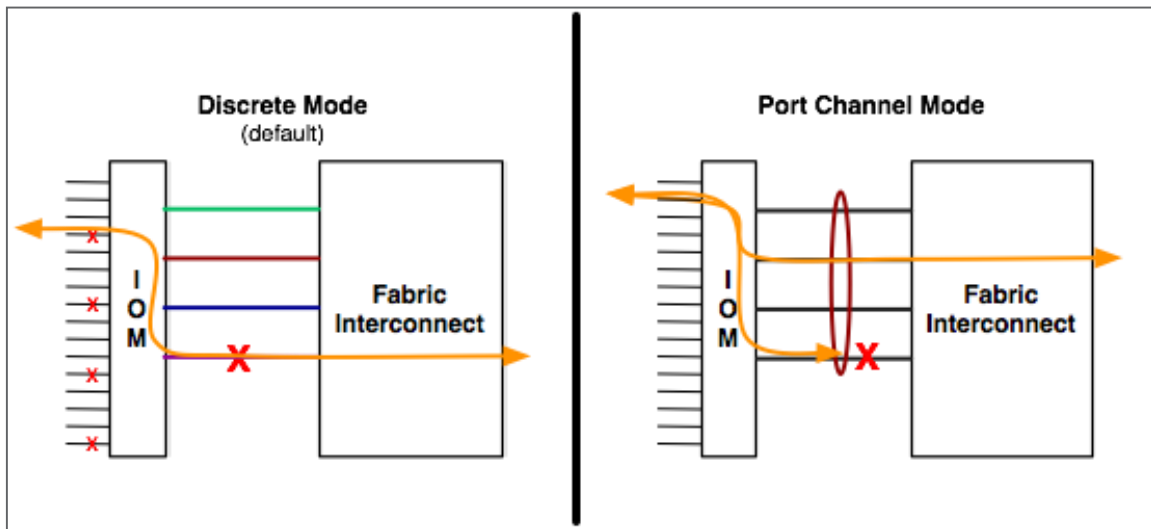
- The fabric extender model (2204XP or 2208XP)
- The modular LAN on motherboard (mLOM) card
- The mezzanine slot card

The Cisco UCS 2208XP Series fabric extenders installed in each blade chassis have eight 10GbE, FCoE-capable, enhanced small form-factor pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect. The Cisco UCS 2204 has four external ports with identical characteristics to connect to the fabric interconnect. Each Cisco UCS 2208XP has 32 10GbE ports connected through the midplane KR lanes to the half-width slots in the chassis, while the 2204XP has 16. This means that each 2204XP enables two 10Gb KR lanes per half-width blade slot, while the 2208XP enables all four KR lanes. The number of KR lanes indicates the potential bandwidth available to the chassis and therefore the blades. With two IOMs in each chassis, there is up to 80Gb of bandwidth available to each half-width slot, depending on the IOM and VIC options selected.

The second-generation Cisco UCS 6200 Series fabric interconnects, the 2200 series Cisco UCS fabric extenders, and the 1300 series Cisco virtual interface cards (VIC) support port aggregation. This capability allows workload rebalancing between these devices, providing link fault tolerance in addition to increased aggregate bandwidth within the fabric. Notably, in the presence of second-generation VICs and fabric extenders, fabric port channels are automatically created in the fabric.

Fabric port channels between the fabric extenders and fabric interconnects are controlled through the chassis/FEX discovery policy. Figure 9 illustrates the two modes of operation for this policy. In discrete mode, each FEX KR connection (the server connection) is tied or pinned to a network fabric connection homed to a port on the fabric interconnect. In case of a failure on the external link, all KR connections are disabled within the FEX IOM. In the case of a fabric port channel discovery policy, the failure of a network fabric link allows redistribution of flows across the remaining port channel members. This is less disruptive to the fabric.

Figure 9) Example of discrete and port channel modes.



First-generation Cisco UCS hardware is compatible with second-generation gear but then can only operate in discrete mode.

In this validated configuration, the Cisco UCS 5108 chassis are populated with Cisco UCS 2208XP IOMs, and each blade has a Cisco UCS VIC1340 with the optional port expander in the mezzanine slot. This passive device provides connectivity for the unused ports on the VIC 1340, essentially enabling the 40Gb potential of the mLOM card for each Cisco UCS 2208XP IOM, for a total of 80Gb bandwidth to each blade.

Jumbo Frames

A balanced and predictable fabric is critical within any data center environment. As designed, the FlexPod architecture accommodates a myriad of traffic types (vMotion, NFS, FCoE, control traffic, and so on) and is capable of absorbing traffic spikes and protecting against traffic loss. Enabling jumbo frames allows the FlexPod environment to optimize throughput between devices while simultaneously reducing the consumption of CPU resources. Cisco UCS and Cisco Nexus QoS system classes and policies deliver this functionality.

In this solution verification effort, the FlexPod platform was configured to support jumbo frames by assigning an MTU size of 9,000 to the best effort QoS system class. By default, all traffic types use the best effort system class, enabling jumbo frames across the network. Individual device interfaces were then configured with an appropriate MTU for the traffic they carry. If finer QoS granularity is required, additional system classes can be configured as needed and applied to the appropriate interfaces. Note that MTU settings must be applied uniformly across all devices in a given L2 subnet to prevent fragmentation and negative performance implications that inconsistent MTUs can introduce.

5.2 Cisco Nexus Network Design

This section provides an overview of the Cisco Nexus network design for this reference architecture.

Network Design Overview

Network Switching

Two Cisco Nexus 9396PX switches running NX-OS software release 7.0(3)I2(2a) were used in this solution verification. These switches were chosen because of their support for the latest NX-OS feature

set, scalability, and readiness for ACI. This design does not utilize ACI but instead has the switches operating in standalone NX-OS mode. One of the design goals for this reference architecture was applicability to the widest range of customer environments. Therefore, ACI was not considered to be a requirement, but this architecture could be integrated into a new or existing ACI topology if desired.

vPCs were used, allowing a port channel from each storage controller and Cisco UCS fabric interconnect to be spread across both switches.

The Cisco Nexus 9000 series currently does not support converged networking with FCoE. If FC or FCoE connectivity is a requirement in a Cisco Nexus 9000 environment, the NetApp storage arrays can be connected directly to the Cisco UCS fabric interconnects, as shown in the appendix.

Host Server Networking

Each VMware ESXi host server has four vNICs, providing two 10GbE ports for iSCSI networking and two 10GbE ports for all other IP networking. These ports are configured into two iSCSI vSwitches with one uplink each and a separate dedicated vSwitch with two uplink ports for other IP traffic. The IP vSwitch uses two active ports and the originating source ID load-balancing algorithm. ESXi servers boot from LUNs on the NetApp AFF8080 storage array using the iSCSI interfaces and access NFS datastores on the AFF8080 using a dedicated VLAN on the IP interfaces.

Storage Networking

Each of the two NetApp AFF8080 storage controllers has a two-port LACP ifgrp (port channel) connected to a vPC across the two Cisco Nexus 9396PX switches. ALUA was used to provide multipathing and load balancing of the iSCSI links. Initiator groups (igroups) were configured on the AFF8080 systems to map boot LUNs to the ESXi host servers.

If you prefer FCoE for SAN boot instead of iSCSI, the AFF8080 storage system can be directly connected to the Cisco UCS fabric interconnects, as shown in the appendix.

Cisco Nexus 9000 Switch

The Cisco Nexus 9000 switch provides a powerful and feature-rich Ethernet data center switching fabric for communications between the Cisco UCS domain, the NetApp storage system, and the enterprise network. For Ethernet connectivity, the Cisco Nexus 9000 uses virtual port channel (vPC). This configuration allows links that are physically connected to two different Cisco Nexus Series devices to appear as a single port channel to a third device. In the FlexPod topology, both the Cisco UCS fabric interconnects and the NetApp storage systems are connected to the Cisco Nexus 9000 switches using vPC, which provides the following benefits:

- Allows a single device to use a port channel across two upstream devices
- Eliminates Spanning Tree protocol blocked ports
- Provides a loop-free topology
- Uses all available uplink bandwidth
- Provides fast convergence if either one of the physical links or a device fails
- Provides link-level resiliency
- Allows HA of the overall FlexPod system

The vPC peer keepalive link is a required component of a vPC configuration. The peer keepalive link allows each vPC-enabled switch to monitor the health of its peer. This link accelerates convergence and reduces the occurrence of split-brain scenarios. In this validated solution, the vPC peer keepalive link uses the out-of-band management network.

FlexPod is a converged infrastructure platform. This convergence is possible because of the support for Ethernet enhancements across the integrated compute stack with regard to bandwidth allocation and flow

control based on traffic classification. Therefore, it is important to implement the following QoS techniques to provide QoS in the FlexPod configuration:

- **Priority Flow Control (PFC) 802.1Qbb.** Lossless Ethernet using a PAUSE on a per class of service (CoS) basis.
- **Enhanced Transmission Selection (ETS) 802.1Qaz.** Traffic protection through bandwidth management.
- **Data Center Bridging Capability Exchange (DCBX).** Negotiates Ethernet functionality between devices (PFC, ETS, and CoS values).

The Cisco Nexus 9000 supports these capabilities through QoS policies. QoS is enabled by default and managed using the Cisco modular QoS CLI, providing class-based traffic control. Realize that DCBX signaling can affect the NetApp controller. Make sure to allocate the proper bandwidth based on the site's application needs to the appropriate CoS classes. In addition, keep MTU settings consistent in the environment to avoid fragmentation issues and improve performance.

The following best practices were used in the verification of the FlexPod architecture:

- The following Cisco Nexus 9000 features were enabled:
 - **LACP.** Part of 802.3ad
 - **Cisco vPC.** For link and device resiliency
 - **Link Layer Discovery Protocol (LLDP).** Allowed the Cisco Nexus 5000 to share and discover DCBX features and capabilities between neighboring FCoE-capable devices
 - **Cisco Discovery Protocol (CDP).** For infrastructure visibility and troubleshooting
- The following vPC settings were configured:
 - A unique domain ID was defined.
 - The priority of the intended vPC primary switch was set lower than the secondary (the default priority is 32768).
 - Peer keepalive connectivity was established.

Note: NetApp recommends using the out-of-band management network (mgmt0) or a dedicated switched virtual interface for the peer-keepalive link.

 - The vPC autorecovery feature was enabled.
 - IP ARP synchronization was enabled to optimize convergence across the vPC peer link.

Note: Cisco Fabric Services over Ethernet synchronized the configuration, spanning tree, MAC, and VLAN information and thus removed the requirement for explicit configuration. The service is enabled by default.

 - A minimum of two 10GbE connections are required for vPC.
 - All port channels were configured in LACP active mode.
- The following Spanning Tree settings were configured:
 - The path cost method was set to long to account for 10GbE links in the environment.
 - The spanning tree priority was not modified (under the assumption that this was an access layer deployment).
 - Loopguard was disabled by default.
 - Bridge protocol data unit (BPDU) guard and filtering were enabled by default.
 - Bridge assurance was only enabled on the vPC peer link.
 - Ports facing the NetApp storage controller and the Cisco UCS were defined as edge trunk ports.

For configuration details, see the Cisco Nexus 9000 Series switch [configuration guides](#).

5.3 NetApp AFF Storage Design

This section provides an overview of the NetApp AFF storage design for this reference architecture.


Storage Design Models

With the combined scale-out and scale-up capabilities of ONTAP, multiple storage architecture designs are possible to meet your technical and business needs. These designs can incorporate SAN data access, NAS data access, or SAN and NAS data access simultaneously.

In the traditional scale-up model, all data services are provided by an HA pair of active-active storage controllers connected to a large number of disks and shelves. To move up the performance and capacity scale, larger and more powerful controller models are necessary.


In the scale-out model, as shown in Figure 10, data services are distributed across multiple HA pairs of storage controllers within a single storage cluster. Each controller can provide a different tier or tiers of capacity and performance. This can include high-performance all-flash configurations, cost-effective performance with hybrid arrays containing both flash and capacity hard disks, or capacity-optimized arrays containing only hard disks. Any one type or any combination of the three types of storage arrays can be mixed within the same cluster as required by the business. In addition, within the NetApp AFF8000 series, all models of controllers can be mixed within the same cluster to allow organic growth and refresh of the environment. Such combinations can also meet any requirements of capacity or performance.

Figure 10) Scale-out storage with ONTAP on AFF.



In the distributed model shown in Figure 11, data services are segmented between multiple HA pairs of storage controllers acting as distinct storage clusters, with each storage cluster composed of either a single HA pair or multiple HA pairs. This is often an artifact of existing deployments but can also reflect internal business organizational constructs or management requirements. In this architecture, each cluster has isolated resources and multiple management interfaces.

Figure 11) Distributed storage clusters.



In this reference architecture, a single HA pair provides all storage services for both infrastructure and virtual desktop workloads. This design supports either scale-out or distributed storage clusters, depending on your organization's needs. In general, NetApp recommends the scale-out approach for greater ease of management, workload mobility, and capacity and performance balancing across nodes.


Storage Scale Considerations

NetApp ONTAP and the NetApp AFF8000 series of storage controllers allow your environment to grow from a single-use case workload, such as a small VDI deployment, to a truly large-scale deployment for either large single workloads or multiple workloads for the enterprise. Individual models of controllers can provide different levels of performance and capacity, as listed in Table 4. Any or all of these controllers can be mixed within the same cluster to meet the capacity and performance requirements of the business while providing cost efficiency during acquisition.

In SAN-only environments (as shown in

Figure 12), or mixed SAN and NAS environments, a single ONTAP cluster can scale to eight nodes or four HA pairs. At the high end, this provides approximately 45PB of data within the same management plane. Mixed SAN and NAS environments can take advantage of any combination of storage protocols within the same cluster (FC, FCoE, iSCSI, NFS, or CIFS/SMB) and therefore can support all business data and application requirements.


Figure 12) ONTAP in a SAN-only environment.



In NAS-only environments, as shown in

Figure 13, a single ONTAP cluster can scale up to 24 nodes or 12 HA pairs. At the high end, this provides approximately 135PB of data within the same management plane. NAS environments can take advantage of NFS and CIFS/SMB storage protocols, providing support for both business file data and virtualization data. NFS is supported for VMware vSphere, OpenStack, Red Hat Enterprise Virtualization, and XenServer. SMB3 is supported for Microsoft Hyper-V.

Figure 13) ONTAP in a NAS-only environment.



In addition to scale, separate clusters can provide an additional level of fault isolation, provide disparate management domains, and support multiple geographic locations. At this time, individual clusters are bound to individual sites, although cross-cluster (intercluster) replication is supported for any geographic distance.

For both SAN and NAS clusters, storage capacity is determined by the following:

- The size of the available disks
- The number of disks supported by the storage controllers, with larger controllers supporting greater numbers of attached disks
- Whether the storage controllers are AFF or FAS systems

The capacities listed in the preceding paragraphs are derived from FAS8080 system limits and the largest supported number of the largest SATA disks.

AFF systems support a different number of attached disks, and SSDs are historically of smaller capacity. Therefore, the maximum capacity of a SAN or hybrid cluster is approximately 7PB, and the maximum capacity of a NAS-only cluster is approximately 22PB.

Storage clusters containing both AFF and FAS nodes, which provide organizations with the greatest flexibility of any storage architecture, have maximum capacities somewhere in between. The maximum capacity of a cluster can be calculated based upon the aggregate of the maximum specifications of the constituent nodes, as detailed on the NetApp [Hardware Universe](#) site (NetApp customer or partner login required).

Note: SSD vendor product roadmaps show exponential increases in capacity in a much shorter time than has been seen for traditional spinning media. The supported capacities of AFF systems increase with the size of the available SSDs and are expected to be larger than what has been observed in the 2016 calendar year.

NetApp AFF8000 Technical Specifications

Table 4 provides the technical specifications for the current two models of the NetApp AFF series: AFF8040 and AFF8080.

Note: All data in this table applies to active-active dual-controller configurations.

Table 4) NetApp AFF8000 storage system technical specifications.


	AFF8080 EX	AFF8060	AFF8040	AFF8020
Maximum SSD	240	240	240	240
Maximum raw capacity: all flash	384TB/349TiB	384TB/349TiB	384TB/349TiB	384TB/349TiB
Effective capacity	1565.3TB/1423TiB	1565.3TB/1423TiB	1565.3TB/1423TiB	1565.3TB/1423TiB
Controller form factor	<ul style="list-style-type: none"> • Dual-enclosure HA • 2 controllers and 2 IOXMs in two 6U chassis • Total of 12U or single-enclosure HA • 2 controllers in single 6U chassis 	<ul style="list-style-type: none"> • Dual-enclosure HA • 2 controllers and 2 IOXMs in two 6U chassis • Total of 12U or single-enclosure HA • 2 controllers in single 6U chassis 	<ul style="list-style-type: none"> • Single-enclosure HA • 2 controllers in single 6U chassis 	<ul style="list-style-type: none"> • Single-enclosure HA • 2 controllers in single 3U chassis
Memory	256GB	128GB	64GB	48GB
NVRAM	32GB	16GB	16GB	8GB
PCIe expansion slots	6 or 24	8 or 24	8	4
PCIe expansion slots: onboard I/O: UTA2 (10GbE/FCoE, 16Gb FC)	8	8	8	4
Onboard I/O: 10GbE	8	8	8	4
Onboard I/O: GbE	8	8	8	4
Onboard I/O: 6Gb SAS	8	8	8	4
Storage networking supported	FC, FCoE, iSCSI, NFS, pNFS, and CIFS/SMB			
Storage networking supported	<ul style="list-style-type: none"> • Data ONTAP 8.3 or later • Data ONTAP 8.3.1 or later for AFF8080 EX single-chassis HA 			

Back-End Storage Connectivity Overview

For the configuration shown in Figure 14, a 6U AFF8080cc storage system was used, as were two DS2246 disk shelves that are 2U per shelf, for a total of 10U.

Note: The image in Figure 14 is a logical view because both nodes reside in one 6U enclosure. This diagram illustrates multipath HA.

Figure 14) Multipath HA to DS2246 shelves of SSD.



NetApp Virtual Storage Console for VMware vSphere

The NetApp VSC is a management plug-in for VMware vCenter Server that enables simplified management and orchestration of common NetApp administrative tasks. This tested reference architecture used the NetApp VSC for the following tasks:

- Setting NetApp best practices for ESXi hosts (timeout values, host bus adapter [HBA], multipath input/output [MPIO], and Network File System [NFS] settings)
- Provisioning datastores
- Cloning infrastructure VMs

The NetApp VSC can be coinstalled on the VMware vCenter Server instance when the Windows version of vCenter is used. A better practice, and one required when using the VMware vCenter Server virtual appliance, is to deploy a separate Windows Server to host the NetApp VSC. Table 5 lists the NetApp VSC VM configuration.

Table 5) NetApp VSC VM configuration.

NetApp VSC	Configuration
OS	Microsoft Windows Server 2012 R2
VM hardware version	8
vCPU	2 vCPUs
Memory	4GB
Network adapter type	VMXNET3
Hard disk size	80GB
Hard disk type	Thin

5.4 VMware vSphere Design

This section provides an overview of VMware vSphere design as part of a VDI environment.

vSphere Cluster Considerations

Separate clusters are recommended for infrastructure and virtual desktop VMs to provide logical separation and fault isolation between the components. This configuration helps to prevent a virtual desktop or desktops from negatively affecting the performance of the infrastructure VMs. For example, interference with the vCenter Server, View Connection Server, or View Composer Server could degrade the performance or availability of all virtual desktops.

vSphere HA should be enabled so that host failures result in only a short outage before VMs are automatically brought back online. Enable host monitoring and admission control so that at least one host failure or maintenance operation can be tolerated while sufficient resources are still provided to run the entire workload of the cluster. Additional capacity can be reserved to provide greater headroom for concurrent host failures or maintenance.

vSphere Dynamic Resource Scheduler (DRS) is also recommended to automatically balance CPU and memory workloads across the cluster members. This is very important in production virtual desktop environments, where the number of VMs and the workload variability are typically higher than in virtual server environments. DRS provides automated remediation of host resource contention and decreases the likelihood of a busy desktop negatively affecting other desktops on the same host. Note that DRS is most effective in steady-state operations. However, it is less effective during desktop maintenance operations such as reboot, recompose, and so on. For this validation, DRS was disabled for consistent distribution of the desktop VMs during these activities.

As previously discussed, vSphere clusters can scale up to 64 nodes within a single cluster starting with VMware vSphere 6.0. This limit is unrelated to the number of physical CPUs or cores within any or all of the nodes. If larger operational scale is a primary concern, using larger hosts, such as with four or more processors and/or commensurately larger amounts of memory, allows greater density within a single vSphere cluster. Scaling up hosts also improves the effectiveness of vSphere memory oversubscription as a result of transparent page sharing. The efficiency of host-based caching technologies, such as vSphere Flash Read Cache, is also improved.

Notably, as a host scales up, failure domains scale equally. Failure of a two-CPU server can result in the failure of 100–200 VDI VMs, whereas a four-CPU server can affect twice as many VMs. Although failures are rare, they must still be taken into account when designing the infrastructure. Fewer hosts per cluster also provide vSphere DRS with fewer options for optimally balancing the virtual desktop workload, which increases the likelihood of host resource contention.

In addition, some maintenance operations can also be affected by the relative size of the host. Evacuating a host with 200–300 VMs requires more time than evacuating one with half the workload. Maintenance operations like this can be mitigated by the appropriate configuration of multi-NIC vMotion, among other technologies, and should be addressed in the vSphere networking design.

When deploying VDI with Horizon 7, VMware imposes some limitations on the vSphere infrastructure beyond the typical vSphere maximums. The maximum number of ESXi hosts in a cluster supporting Horizon 7 is 32, and the maximum number of VMs that can be registered to a single vSphere host is 1,000. Therefore, the maximum number of VMs within a single vSphere cluster is 32,000. Based on these upper limits, the minimum number of hosts required for 2,500 virtual desktops is four. Three hosts provide capacity for the VMs, and at least one host provides for failover capacity. This configuration creates an effective host-to-VM ratio of 1:625 during normal operations and 1:833 during a single-host failure or maintenance event. This configuration is extremely dense and significantly increases the risk of resource contention.

Using the maximum number of hosts (32) for 2,500 virtual desktops provides an effective host-to-VM ratio of 1:78 during normal operations or during a single-host failure or maintenance event. However, at cluster

sizes of that scale, NetApp recommends reserving more than a single host's capacity for failure or maintenance.

In real-world deployments, the available physical memory and CPU resources are gating factors for these ratios, and they depend on the specific applications and workloads used within the customer's environment. Scale-up numbers, such as the host and cluster maximums, are not typically recommended, particularly at the host level, because of potential issues with planned and unplanned maintenance events. The most typical (and at this time most cost-effective) host-to-VM ratio is between 1:125 and 1:225 for a dual-CPU vSphere host. In general, each CPU core can provide sufficient compute resources for 8 to 10 virtual desktop CPUs.

vSphere Networking Considerations

In this reference architecture, standard vSwitches were used for the VM connectivity, management vmkernel, and vMotion vmkernel portgroups. VMware vSphere Enterprise Plus licensing enables other networking options that provide additional features beyond standard virtual switches, such as distributed virtual switches, Cisco Nexus 1000v, or Cisco Virtual Machine Fabric Extender (VM-FEX).

VMware vSphere Distributed Virtual Switches

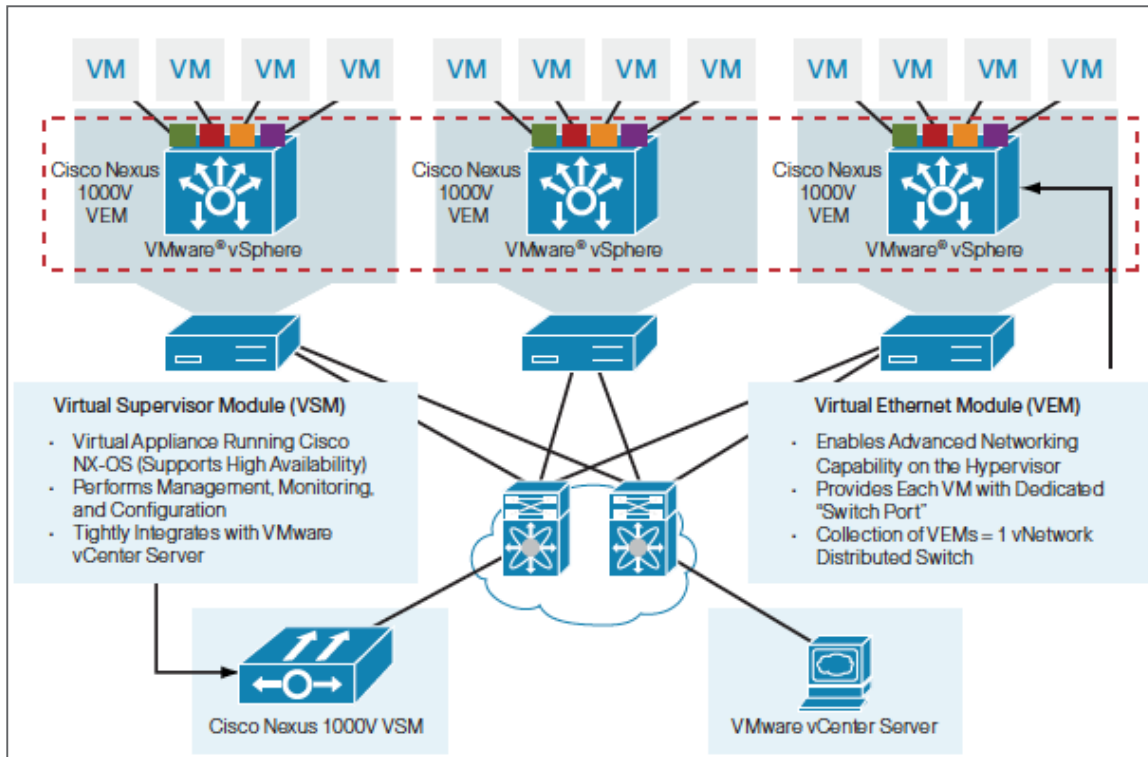
The VMware vSphere Distributed Switch (VDS) provides a number of benefits over traditional (or standard) vSwitches. This includes centralized management of virtual networking across the entire virtual data center and support for advanced VMware vSphere networking features. These advanced features include the following:

- Network I/O control (NIOC) for control of the priority of network traffic on vSphere host uplinks during periods of contention. This feature allows high-bandwidth traffic such as vMotion to be used concurrently and nondisruptively with other traffic on the same uplinks.
- LACP for the negotiation and automatic configuration of link aggregation between vSphere hosts and the physical network switch.
- Load-based teaming (LBT) for more efficient balancing of ingress and egress traffic on vSphere host uplinks for each distributed virtual switch.
- Private VLANs (PVLANs).
- Runtime state and statistics of VMs as they migrate between different vSphere hosts.
- Enablement of inline monitoring and centralized firewall services.
- Support for single-root I/O virtualization.
- Bidirectional traffic shaping.
- BPDU filtering.
- Network health check, templating, backup, rollback, and recovery of virtual network configuration.
- Support for third-party virtual switch extensions such as those provided by Cisco.

Cisco Nexus 1000v

The Cisco Nexus 1000v is a virtual distributed switch that fully integrates into a vSphere-enabled environment. The Cisco Nexus 1000v operationally emulates a physical modular switch, with a virtual supervisor module (VSM) providing control and management functionality to multiple line cards. In the case of the Cisco Nexus 1000v, the ESXi nodes become modules in the virtual switch when the Cisco Virtual Ethernet Module (VEM) is installed. Figure 15 is an architectural overview of a Cisco Nexus 1000v.

Figure 15) Cisco Nexus 1000v architecture (graphic provided by Cisco).



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

- **Port channels**
- **QoS**
- **Security.** By using private VLAN, access control lists (ACLs), and port security.
- **Monitoring.** By using NetFlow, switch port analyzer (SPAN), and encapsulated remote SPAN.
- **vPath.** Provides efficient traffic redirection to one or more chained services such as the Cisco Virtual Security Gateway and the Cisco ASA 1000v.

Cisco Virtual Machine Fabric Extender

The Cisco VM-FEX addresses both management and performance concerns in the data center by unifying physical and virtual switch management. The Cisco VM-FEX collapses both virtual and physical networking into a single infrastructure, reducing the number of network management points and enabling consistent provisioning, configuration, and management policy within the enterprise. This is achieved by joining the Cisco UCS Manager to the VMware vCenter management platform through the Cisco UCS VMware plug-in. This integration point between the physical and virtual domains of the data center allows administrators to efficiently manage both their virtual and physical network resources. The decision to use VM-FEX is typically driven by application requirements such as performance and the operational preferences of the IT organization.

The Cisco UCS VIC offers each VM a virtual Ethernet interface, or vNIC. This vNIC provides direct access to the Cisco UCS fabric interconnects and the Cisco Nexus 9000 Series switches through which forwarding decisions can be made for each VM using a VM-FEX interface. Cisco VM-FEX technology supports two modes of operation:

- **Emulated mode.** The hypervisor emulates a NIC (also referred to as a back-end emulated device) to replicate the hardware it virtualizes for the guest VM. The emulated device presents descriptors for read and write and provides interrupts to the guest VM just as a real hardware NIC device would. One NIC device that VMware ESXi emulates is the vmxnet3 device. The guest OS in turn instantiates a device driver for the emulated NIC. All of the resources of the emulated devices' host interface are mapped to the address space of the guest OS.
- **PCIe pass-through or VMDirectPath mode.** VIC uses PCIe standards-compliant IOMMU technology from Intel and VMware VMDirectPath technology to implement PCIe pass-through across the hypervisor layer and eliminate the associated I/O overhead. The pass-through mode can be requested in the port profile associated with the interface using the high-performance attribute.

VMware vCenter Considerations

VMware vCenter is a critical component for a VMware Horizon View infrastructure; all management and maintenance operations rely on it to perform correctly and efficiently. Because of the component's critical nature, NetApp highly recommends that you provide as much resilience and data protection as possible to its core components, the individual vCenter modules, and the vCenter database.

If vCenter is running as a VM, whether as a Windows installation or as the vCenter Server virtual appliance, vSphere HA provides fundamental protection, as it can for the vCenter database server, if virtualized. In addition, starting with vCenter Server 5.5, Microsoft SQL Clustering Service (MSCS) is now a supported configuration to provide HA for the vCenter database separately from vSphere HA. It can if a vCenter instance is a physical server.

For more information about how to configure support for MSCS with a Windows-based vCenter Server, see the [VMware Knowledge Base \(KB\) article 2059560](#).

A single vCenter Server can support up to 10,000 concurrently active (powered-on) VMs, up to 15,000 concurrently registered VMs, and up to 1,000 connected vSphere hosts. These limits reflect the use of either the Windows-installable vCenter Server or the vCenter Server Appliance.

The vCenter Server Appliance now provides support for even larger virtual environments, allowing customers to slightly reduce the number of required Windows Server licenses and antivirus licenses. More importantly, it provides a quicker and less complicated installation and upgrade process for this critical functionality. However, providing HA for the vCenter database can be more complicated, because the only external database supported at this time is Oracle. For customers with strong Oracle skillsets, this might not be an issue. For more Windows-focused customers, the Windows Server installation for vCenter Server and the use of Microsoft SQL Server is likely a better fit because this process enables the use of MSCS for database resiliency.

The vCenter VM and its database server VMs should be sized appropriately to their workload. An example configuration based on a verified workload of 2,500 users is shown in Table 6.

Table 6) VMware vCenter Server Appliance VM configuration.

VMware vCenter Server Appliance VM	Configuration
VM hardware version	8
vCPU	8 vCPUs
Memory	24GB
Network adapter type	VMXNET3
Hard disk size	280GB
Hard disk type	Thin

With vCenter 6.0, VMware has made this configuration easy to implement through the use of a deployment wizard for the vCenter Server Appliance. The administrator is given a set of choices for the scale of the environment to be managed by this vCenter instance. vCenter then automatically configures the deployed vCenter VM appropriately. These choices are shown in Figure 16.

Figure 16) VMware vCenter Server Appliance size options.

VMware vCenter Server Appliance Deployment

✓ 1 End User License Agreement
 ✓ 2 Connect to target server
 ✓ 3 Set up virtual machine
 ✓ 4 Select deployment type
 ✓ 5 Set up Single Sign-on
6 Select appliance size
 7 Select datastore
 8 Configure database
 9 Network Settings
 10 Ready to complete

Select appliance size
 Specify a deployment size for the new appliance

Appliance size:

Tiny (up to 10 hosts, 100 VMs)
Tiny (up to 10 hosts, 100 VMs)
Small (up to 100 hosts, 1,000 VMs)
Medium (up to 400 hosts, 4,000 VMs)
Large (up to 1000 hosts, 10,000 VMs)

Description:

This will deploy a Tiny VM configured with 2 vCPUs and 8 GB of memory and requires 120 GB of disk space. This option contains vCenter Server with an embedded Platform Services Controller.

Back Next Finish Cancel


For the validation of this solution, a medium configuration was chosen. Although a small configuration would have been sufficient for the number of hosts created, a medium size was required for the planned number of VMs. In addition, a medium configuration allows for more than 50% growth of the environment. Organizations planning to grow to larger VDI environments should instead choose the large configuration during the initial vCenter Server Appliance deployment.

As an additional sizing consideration, you must decide whether to deploy a vCenter installation with an embedded platforms services controller (PSC) or whether to deploy vCenter Server and PSC as separate VMs. A single vSphere environment supports up to eight PSCs for redundancy and load-balancing, and up to 10 vCenter instances can share the PSC services. Using an embedded PSC within the vCenter Server limits the environment to a single PSC and four vCenter servers. For the validation of this solution, a separate PSC VM was used to provide greater potential scalability for the overall environment. However, the scale limit of four vCenter instances per PSC is more than sufficient for most organizations, and either deployment model is fully supported.

5.5 VMware Horizon View Design

This section provides an overview of the VMware Horizon View design and explains user assignment, automated desktop pools, linked-clone desktops, and the creation of desktop pools. Figure 17 depicts an architectural overview of Horizon View. For the verification of this architecture, a smaller number of View Connection Servers were used than is shown in Figure 17.

Figure 17) Horizon View architecture.



Horizon View Connection Server Considerations

In a typical large-scale virtual desktop deployment, the maximum limit for a VMware Horizon View Connection Server is reached when each Connection Server instance supports 2,000 simultaneous connections. When this occurs, you must add more Connection Server instances and build additional VMware Horizon View desktop infrastructures to support any additional virtual desktops. Each desktop infrastructure is referred to as a pool of desktops (POD).

A POD is a building-block approach to architecting a solution. The POD size is defined by the VMware Horizon View desktop infrastructure (the desktop VMs) plus any additional VMware Horizon View infrastructure resources that are necessary to support the desktop infrastructure PODs. In some cases, it might be best to design PODs that are smaller than the maximum size to allow growth in each POD or to reduce the size of the fault domain.

Using a POD-based design gives IT a simplified management model and a standardized way to scale linearly and predictably. By using ONTAP, customers can have smaller fault domains that result in higher availability. In this reference architecture, the number of Horizon View Connection Servers was limited to one so that the POD-based design limits could be scaled. However, the results of the testing demonstrate that it was possible to deploy multiple PODs on this platform.

VMware Horizon View groups desktops into discrete management units called pools. Policies and entitlements can be set for each pool so that all desktops in a pool have the same provisioning methods, user assignment policies, logout actions, display settings, data redirection settings, data persistence rules, and so on.

A single Connection Server instance can support up to 2,000 simultaneous connections. In addition, five Connection Server instances can work together to support up to 10,000 virtual desktops. For increased availability, Horizon View supports using two additional Connection Server instances as standby servers. The Connection Server can optionally log events to a centralized database that is running either Oracle Database or Microsoft SQL Server. Table 7 lists the components of the VMware Horizon View Connection VM configuration.

Note: Two Horizon View Connection Servers are used in this reference architecture. Production deployments should use three or more Connection Servers to provide sufficient resources and broker availability if an outage affects a Connection Server.

Table 7) VMware Horizon View Connection VM example configuration.

Horizon View Connection VM	Configuration
OS	Microsoft Windows Server 2012 R2
VM hardware version	11
vCPU	4 vCPUs
Memory	16GB
Network adapter type	VMXNET3
Hard disk size	80GB
Hard disk type	Thin

User Assignment

Each desktop pool can be configured with a different user assignment. User assignments can be either dedicated or floating.

Dedicated Assignment

Through the dedicated assignment of desktops, users access the same virtual desktop each time they log in. Dedicated assignment allows users to store data either on a persistent disk (when using linked clones) or locally (when using full clones). These are usually considered to be and used as persistent desktops. However, the act of refreshing or recomposing makes them nonpersistent.

User-to-desktop entitlement can be a manual process or an automatic process. An administrator can link a given desktop to a user, or the administrator can allow VMware Horizon View to automatically entitle the user to a desktop when the user logs in for the first time.

Floating Assignment

With floating user assignment, users are randomly assigned to desktops each time they log in. These desktops are usually considered to be and are used as nonpersistent desktops. However, a user who does not log out of the desktop always returns to the same desktop.

Automated Desktop Pools

An automated desktop pool dynamically provisions virtual desktops. With this pool type, VMware Horizon View immediately creates a portion of the desktops and then, based on demand, provisions additional desktops to the limits that were set for the pool. An automated pool can contain dedicated or floating desktops. These desktops can be full clones or linked clones.

A major benefit of using VMware Horizon View with automated pools is that additional desktops are created dynamically on demand. This automation greatly simplifies the repetitive administrative tasks associated with provisioning desktops.

Linked-Clone Desktops

To the end user, a linked-clone desktop looks and feels like a normal desktop, but it is storage efficient, consuming a fraction of the storage required for a full desktop. Because of the architecture of linked

clones, three unique maintenance operations can be performed to improve the storage efficiency, performance, and security and compliance of the virtual desktop environment. These operations are refresh, recompose, and rebalance.

View Composer API for Array Integration

VMware View Composer API for Array Integration (VCAI) enables offloaded VDI VM provisioning operations on NFS storage that are similar to the older and more established vSphere APIs for Array Integration (VAAI) for the ESXi hypervisor. NetApp is one of only four vendors that have passed the certification process for VCAI to take advantage of this tight integration with VMware Horizon View.

Horizon View Composer Considerations

VMware Horizon View Composer is a critical component of a Horizon View infrastructure because it is responsible for the creation and maintenance operations associated with linked-clone desktops.

View Composer can be installed on the VMware vCenter Server itself when vCenter is running on a Windows server rather than as the vCenter Server appliance, or as a standalone Windows Server. For scalability and fault isolation, NetApp recommends installing View Composer on a standalone server that is dedicated for this purpose. View Composer cannot be installed on a server that already performs another VMware View role, such as the View Connection Server, View Security Server, and so on.

View Composer only supports one vCenter Server per instance. If the virtual desktop architecture requires multiple vCenter instances, multiple View Composer instances (one per vCenter Server instance) are required. View Composer supports up to 2,000 desktops per virtual desktop pool, and each desktop pool can contain only one vSphere cluster or resource pool.

For View Composer provisioning and maintenance operations such as recompose, refresh, and rebalance, only 8 concurrent provisioning operations and 12 concurrent maintenance operations are supported by default. Additionally, the vCenter Server has its own limits for provisioning and power operations: by default, 20 and 50, respectively. These numbers can be increased by editing the advanced settings for the vCenter Server within the View Administrator console. In this reference architecture, the View Composer defaults were changed to 30 for both provisioning and maintenance operations.

Composer requires a database server for its configuration and operation. For large-scale production deployments, NetApp recommends running full versions of the database software on dedicated and HA servers or VMs. Table 8 lists an example configuration for Horizon View Composer VM, and Table 9 lists an example configuration for the Microsoft SQL Server database VM.

Table 8) Horizon View Composer VM example configuration.

Horizon View Composer VM	Configuration
OS	Microsoft Windows Server 2012 R2
VM hardware version	11
vCPU	4 vCPUs
Memory	16GB
Network adapter type	VMXNET3
Hard disk size	80GB
Hard disk type	Thin

Table 9) Microsoft SQL Server database VM example configuration.

Microsoft SQL Server VM	Configuration
OS	Microsoft Windows Server 2012 R2
VM hardware version	11
vCPU	2 vCPUs
Memory	8GB
Network adapter type	VMXNET3
Hard disk size	80GB
Hard disk type	Thin

User Data Considerations

Persona, individual user, and shared data are the core of the user experience for both physical and virtual desktops. End users require consistent access to their data both for normal productivity and for a sense of continuity for their daily activities. A profile management solution such as View Persona Management or Liquidware Labs Profile Unity, among others, provides consistent access and a consistent user experience. These applications can also store data centrally, such as on a CIFS/SMB share on a NetApp AFF or FAS storage array within the FlexPod platform.

Shared data, such as a department, division, or company-wide file share, is typically provided by NAS servers or by storage arrays natively. If this data is not already located on a FlexPod system, it can be easily migrated to one using common file system management tools. You then receive the benefits provided by a NetApp AFF, including space-efficient Snapshot copies, block-level deduplication, compression, replication, and HA. ONTAP also supports the advanced, simultaneous sharing of data—even the same files and directories—between both Windows and UNIX clients using CIFS/SMB and NFS.

An SVM can be dedicated to act as the file access point, or an SVM can simultaneously provide disparate block and file services. Separating workloads by using an SVM permits management delegation by workload as well as one level of workload prioritization with storage QoS. Individual SVMs can also be used when there is a need to join multiple Active Directory instances, particularly when these directories do not share trusts between them.

In a physical desktop environment, user and persona data is typically located on the desktop or accessed through file redirection services such as roaming profiles. View Persona Management, Profile Unity, and other solutions allow you to migrate the user and persona data to a shared NAS repository used for the virtual desktop environment. Migration can be performed whether these underlying virtual desktop pools are persistent or nonpersistent. Using persona or profile management for nonpersistent pools, however, is critical to the end-user experience and acceptance of the new environment.

User data can be stored within the same storage array used for the virtual desktops on a different array within the same cluster or even on a separate array entirely. Using a shared cluster containing multiple arrays of varying performance and capacity provides workload isolation when needed and also provides unified management while maintaining flexibility. A single cluster with one or more AFF systems paired with other hybrid (SSD and HDD) or disk-only FAS systems can provide rich data management and tiering opportunities while minimizing cost-versus-capacity complexities.

View Storage Accelerator Considerations

VMware View Storage Accelerator is a host-based cache first introduced in VMware vSphere 5 and enabled with VMware View 5.1 and later.

View Storage Accelerator supports any desktops managed by vCenter Server, such as manual desktops, automated full-clone desktops, and automated linked-clone desktops. View Storage Accelerator is a host-based memory cache that uses ESXi host memory. The cache size can be set from 100MB to a maximum of 2048MB, and it can be set differently from host to host.

The goal of View Storage Accelerator is to reduce the number of IOPS being read from local or shared storage. It also seeks to improve performance by reducing read latency during activities such as boot storms, login storms, and steady-state operations. VMware View Storage Accelerator has been shown in some environments to reduce IOPS and throughput to the back-end storage by a significant amount. However, in other environments, it may have little or even a negative effect. View Storage Accelerator should not be used in conjunction with VCAI. Before implementing this feature in production, be sure to evaluate its effectiveness for your environment.

View Storage Accelerator consumes additional storage. It creates one additional `-digest.vmdk` disk for each virtual machine disk (VMDK). This additional VMDK is a digest disk that contains the hash values for the data VMDKs. These disks consume 0.5% of the space of the original VMDK when using SHA1 hash with a 4K block size and 1.2% when using collision detection. For example, if the VM is 50GB, the digest is 2.5GB.

View Storage Accelerator was not used during the verification of this reference architecture so that the nonoffloaded performance of the AFF storage could be measured and verified. In addition, half of the workload testing in this solution validation takes advantage of VCAI provisioning, and NetApp does not recommend the use of the View Storage Accelerator with VCAI.

Virtual Desktop Considerations

This reference architecture targets a workload of 2,500 VMs.

The desktop VM template was created with the virtual hardware and software listed in Table 10.

Table 10) Virtual desktop configuration.

Desktop	Configuration
Desktop VM	
VM hardware version	11
vCPU	1
Memory	2GB
Network adapter type	VMXNET 3
Hard disk size	20GB
Hard disk type	Thin
Desktop Software	
Guest OS	Microsoft Windows 10 (32-bit)
VMware tools version	10.0.0, build 3000743
Microsoft Office	2016 version 16.0.6769.2017
Adobe Acrobat Reader	2015.010.2060
Doro PDF	1.82
VMware Horizon View Agent	7.0.0, 3633490

6 Design Considerations

Table 11 lists many of the important design considerations for running VMware Horizon View on a FlexPod Datacenter implementation. Organizations should always take similar considerations into account when deploying any new infrastructure. These considerations are critical for virtual desktop environments because of the great variability in desktop workloads and the number of users potentially affected by incorrect sizing, incorrect configuration, or poor design elements.

Table 11) Design considerations.

Design Aspect	Design Considerations
Network switch series and model	<p>As with any FlexPod installation, both Cisco Nexus 9000 and Cisco Nexus 5000 series network switches can be used with this design. The primary considerations for the switch series are as follows:</p> <ul style="list-style-type: none">• The Cisco Nexus 9000 series is the latest hardware platform and supports both standalone (traditional) NX-OS and ACI. If organizations are considering implementing ACI, the 9000 series should be the default choice. The 9000s do not support FC or FCoE. Therefore, all SAN boot and other storage access must be through IP protocols unless the storage controllers are connected directly to the Cisco UCS fabric interconnects. This configuration is shown in the appendix.• The Cisco Nexus 5000 series supports FCoE but does not support ACI. Organizations that require FCoE but do not want a direct connect topology and have no plans for implementing ACI should use the 5000 series.• FlexPod components can also be connected to new or existing SAN infrastructure if additional SAN connectivity is required.
Host boot device	<p>FlexPod supports three common host boot device options:</p> <ul style="list-style-type: none">• Local boot. Requires per-blade HDD or SSD. Use of local boot removes a key value proposition of Cisco UCS (stateless computing) but does enable host independence from, and parallel deployment with, shared storage.• SAN boot. Requires shared storage to function and forces a serial approach to deployment because such storage must be available before servers can be deployed. By far the most common FlexPod boot device, SAN boot is a cornerstone of stateless computing in Cisco UCS and provides true independence between server identity and server hardware.• PXE booting. Requires boot technology and software licensing for solutions such as VMware Auto Deploy and is dependent on that boot infrastructure being available in order to deploy or run servers. PXE booting provides an even more stateless computing solution than SAN boot and can be used either in conjunction with local or SAN boot or as an alternate methodology.
Host SAN boot protocol	<p>There are two options for using SAN boot:</p> <ul style="list-style-type: none">• FCoE. Requires FC or converged adapters on the storage array. FCoE connectivity requires either FCoE-capable Cisco Nexus switches or a direct connect topology.• iSCSI. Requires either Ethernet or converged adapters on the storage array. No specific Cisco Nexus switch model or series is required.

Design Aspect	Design Considerations
Storage controller model	<p>With the AFF series, the primary considerations concern capacity and performance. The AFF8080 provides significantly more capacity and performance for a small price differential.</p> <p>With the introduction of the AFF8080cc, the physical rack unit footprint is no longer a consideration.</p>
Storage cluster connectivity	<p>Intercluster communication for AFF and FAS storage clusters has two topologies:</p> <ul style="list-style-type: none"> • Switched. All cluster-member communication occurs across a redundant pair of dedicated 10GbE switches. These cluster switches must be one of a few supported models, such as the Cisco Nexus 5596, and must not be used for noncluster data traffic. A switched topology is required for clusters larger than two nodes. This topology provides the easiest transition from a two-node to a four-node or higher configuration. • Switchless. HA pair members are directly connected to each other, eliminating the need for dedicated 10GbE switches for cluster communication. A switchless topology is only supported for two-node clusters. Two-node switchless clusters can be converted nondisruptively to a switched cluster topology.
Storage scaling	<p>AFF/FAS clusters can scale up to 8 nodes (4 HA pairs) for SAN or hybrid SAN/NAS clusters and up to 24 nodes (12 HA pairs) for NAS-only clusters. Organizations utilizing SAN boot are limited to these eight nodes within a single cluster, at least for the cluster providing SAN boot storage. Depending on the scale and scope of the infrastructure, a smaller cluster can provide the SAN services required for SAN boot and other block storage needs. One or more larger clusters can provide NAS storage for VM datastores and other workloads.</p> <p>With AFF, organizations have the flexibility of scaling out or up as needed. If performance requirements are less than capacity requirements, it is simpler to scale up (bigger SSDs or more SSDs) rather than out.</p>
Compute fabric interconnect model	<p>Fabric interconnect model considerations are primarily around scale. Organizations can choose the appropriate fabric interconnect model based on the number of devices to be connected and/or the bandwidth requirements of each device. The Cisco UCS 6300 series fabric interconnects are the first to provide support for 40GB networking and enable organizations to more thoroughly future-proof their infrastructure.</p>

Design Aspect	Design Considerations
Compute blade model	<p>Compute blades come in three form factors:</p> <ul style="list-style-type: none"> • Half width. Supports up to dual CPUs and hundreds of gigabytes of memory (limits depending on model). Most commonly deployed form factor, with up to eight fitting in a single chassis. Half-width blades provide an ideal building block size for VDI because each blade can provide significant CPU and memory scale-up options. This configuration provides the most scale-out capabilities, minimizing the effect of host failures or maintenance activities. • Full width. Supports up to quad CPUs and more memory than half-width blades. This format is not commonly seen or required for VDI environments for which this degree of scale-up is not necessary. Indeed, this format can be problematic due to the failure domain effect of a single blade. • Full width and double height. Supports up to quad CPUs and more memory than half-width or full-width blades. These blades are primarily targeted for large database or similar workloads where an application can take advantage of the larger pool of available resources within a single host. These blades do not provide greater aggregate CPU or memory resources than four half-width blades, which take up the same number of chassis slots. However, the half-width provides a smaller failure domain and thus is better suited for VDI environments.
VMware vCenter deployment type	<p>vCenter can be deployed either to a Windows OS or using a virtual appliance:</p> <ul style="list-style-type: none"> • Windows installation. Is the only choice for organizations that want to keep their virtualization management solution on a bare-metal platform. A Windows installation can also be deployed to a Windows VM, and this installation choice provides many administrators with their most familiar methods for troubleshooting. A Windows installation is also the preferred choice for organizations with strong Microsoft SQL Server experience, particularly for database backup, because the virtual appliance cannot use MS SQL Server for its database. • vCenter Server Appliance. Is the simplest deployment option because it requires no additional OS or antivirus licensing. The vCenter Server Appliance is also the recommended deployment method from VMware. Recent releases support a virtual infrastructure scale equivalent to Windows installation. The vCenter Virtual Appliance can only use Oracle or PostgreSQL for its database, which can be problematic for customers without those skills and/or licensing.
VMware vCenter PSC deployment options	<p>Starting with vSphere 6.0, a vCenter installation is composed of two constituent parts that can be installed separately or together: a platform services controller (handling single sign-on and related services) and the vCenter Server itself. Combined installation is the simplest option for deployment or future troubleshooting and interservice communication. However, in a single vSphere environment, there can be at most eight PSC machines, which limits the number of possible vCenter instances within that environment. When these roles are separated, the PSC and vCenter instances can be scaled independently, and the number of vCenter instances can increase to 10.</p>

Design Aspect	Design Considerations
VMware vCenter resource sizing	Refer to the sizing guidance from VMware concerning the number of hosts and VMs to be managed by the new vCenter instance and plan for potential growth from the outset. The effect on the infrastructure environment of an oversized vCenter instance is minimal in comparison to an undersized one.
Infrastructure placement	<p>In small environments, the mixing of infrastructure management workloads with generic application workloads is common. As environments scale, it is a best practice to isolate infrastructure management from the rest of the environment. You can take a progressive approach depending on the needs and suitability for the organization:</p> <ul style="list-style-type: none"> • A dedicated infrastructure host cluster (shared vCenter instance and Cisco UCS) • Dedicated infrastructure storage • A dedicated infrastructure storage cluster • A dedicated infrastructure host cluster (a distinct Cisco UCS chassis or Cisco UCS domain) • A dedicated infrastructure vCenter instance <p>A greater dedication of resources provides the highest level of protected resources and fault isolation. However, this configuration adds significantly to infrastructure cost and complexity. At a minimum, NetApp recommends dedicated infrastructure host clusters for FlexPod VDI environments.</p>
Host scaling	<p>Host sizing and cluster scaling are closely related and have similar considerations:</p> <ul style="list-style-type: none"> • Virtual desktop CPU and memory requirements routinely increase, and achieving reasonable ROI and TCO numbers for VDI requires significant desktop-to-host consolidation ratios. • High consolidation ratios lead to larger failure domains (more desktops affected by a single host failure or problem) and longer maintenance or recovery windows (more VMs to migrate or reboot take more time). • In practice, you must find a host-scaling sweet spot.
Desktop persistency	<p>Organizations must evaluate the level and type of persistency of virtual desktops required. Persistency in this context encompasses both user data and the user persona or the user profile.</p> <ul style="list-style-type: none"> • Persistent. All user data and persona profiles, including applications, are continuously available. This can be accomplished by using full clones for which each desktop is independently managed after provisioning. It can also be accomplished by using linked clones with persistent disks for which desktop updates and configuration changes are managed at the pool level. The majority of VDI deployments focus on a persistent desktop user experience. • Nonpersistent. User data and persona profiles are ephemeral and are reset frequently (upon disconnect, logoff, defined time periods, and so on). Only very specific use cases, such as call centers, may be appropriate for fully nonpersistent desktops. User data persistency can be achieved with nonpersistent desktops by using complementary user environment management products from VMware (available in higher-level VMware Horizon license options) or other third parties.

Design Aspect	Design Considerations
Operating system for View component services	The latest Windows server OSs are used in this solution, but older versions can also be used. For example, an organization might be standardized on an older OS version, or it might not have licenses for the latest version. If so, this organization must make sure that the OS version is compatible with the View components that are being installed. No changes to the overall design are necessary when using an older OS supported by the View components.
Operating system for View desktops	The latest Windows client OS is used in this solution, but older versions can also be used. If an organization has standardized on an older OS version (such as Windows 7) or does not have licenses for the latest version, then it must make sure that the OS version is listed as compatible with the View components that are being installed. No changes to the overall design are necessary when using an older OS supported by the View components.
Applications for View desktops	Applications required by the users and groups who access the virtual desktops are critical for a successful virtual desktop infrastructure. Application delivery, standard application sets, application performance, and so on have a significant effect on the design of the environment and, ultimately, the user experience. NetApp recommends that you perform detailed analysis and assessment of production physical desktops to understand the requirements for a new VDI environment. Assessment is also critical when you add a new application or user group to an existing VDI environment.
View desktop sizing	To support a good user experience, virtual desktops must be sized appropriately for the type of user workload and the underlying physical infrastructure. The 2 vCPU/2G of RAM sizing used in this design is targeted at a moderate knowledge workload. Smaller sizing is only recommended for light desktop usage with few concurrently running applications, low-to-no video usage, and so on. Examples include a call center operation. Larger sizing is needed for power desktop users with many running applications, heavy video usage, and so on. You must then match properly sized virtual desktops to correspondingly sized hosts to avoid excessive CPU or memory oversubscription that causes poor performance.
View component services sizing	Like vCenter, each of the View components (Connection Servers, Composer, and so on) should be sized appropriately for the expected desktop scale. If an organization anticipates growth of the environment, as is typical, larger configurations should be deployed from the beginning.

Design Aspect	Design Considerations
View cloning technology	<p>There are three main VM cloning options for VDI:</p> <ul style="list-style-type: none"> • Native vSphere cloning using View linked clones. Although all storage protocols are supported, this is the least performant option and should only be implemented if storage-offloaded cloning is not available. • Storage-offloaded cloning. This option increases the speed of VM provisioning operations. This method has two suboptions: <ul style="list-style-type: none"> – Using VAAI. All storage protocols are supported, with SAN provider built into ESXi natively and NAS (NFS) requiring installation of a small VIB to each ESXi host. – Using VCAI. Only NAS (NFS) is supported. • Instant clone. A new feature with vSphere 6 and Horizon 7 that has better performance than native linked clones but does not provide storage offload or integration.
Number of View desktop pools	<p>Most organizations have multiple desktop pools for several reasons:</p> <ul style="list-style-type: none"> • User entitlement is at the pool level. • Different users or groups require different desktop configurations: <ul style="list-style-type: none"> – VM size (CPU, memory, and disk) – Included applications – Persistent versus nonpersistent – How persistency is implemented – Default access protocol options (PCoIP, RDP, or Blast Extreme) – Number of supported monitors (defined by pool) – Security settings – And so on • Scale limits for a single desktop pool
Number of View master images	<p>Multiple desktop pools can be based on the same master image. The number of master images an environment can use or maintain depends primarily on the number of different VM configurations (CPU, memory, and disk). It also depends on the different application sets needed by the different user groups in the organization and whether software management or deployment tools are used. ThinApp, AppVolumes, and other tools allow a single master image to provide different application sets to different users or groups.</p>
Number of View datastores	<p>The number of datastores used for a VDI environment is influenced by several factors:</p> <ul style="list-style-type: none"> • The number of virtual desktops in the environment. • The storage protocol. Even with VAAI, NAS datastores can safely run larger numbers of VMs within an individual datastore than is possible with SAN datastores. • Storage array size and layout. An individual datastore maps to an individual storage controller and the available physical storage connected to that controller. • Storage controller best performance practices and load balancing across and within controllers. Increasing the number of datastores can provide better parallelism for I/O operations.

7 Best Practices

Table 12 lists best practices recommended by NetApp for designing or implementing a VDI running on FlexPod Datacenter with VMware Horizon View.

Table 12) VDI best practices.

Best Practice Area	Best Practice Details
Physical desktop assessment	During the planning and design, include an assessment of the physical desktop environment replaced by the proposed VDI. Performance and capacity requirements can vary greatly between desktop environments. Enabling the new environment to provide an experience that is as good as or better than the previous end-user experience is critical for the acceptance and ultimate success of the project.
Architecture redundancy	<ul style="list-style-type: none">• Architect at least one spare Connection Server instance per View POD to provide resiliency in the event of a host or Connection Server failure.• Always configure vSphere HA for all VMs, infrastructure and desktop, used with the virtual desktop environment.• Use database HA options to protect the vCenter Server and View Composer databases.
Architecture resilience	<ul style="list-style-type: none">• The infrastructure VMs for VMware View should operate on different hardware than the virtual desktops themselves. This provides the highest level of fault isolation and performance assurance.• VMware vSphere DRS should be enabled and set to fully automatic to balance workloads across all hosts within the vSphere clusters.
vSphere host configuration	Use the NetApp VSC to set recommended values on the vSphere hosts.
Datastore provisioning	Use the NetApp VSC to provision datastores to the vSphere hosts. This reduces the amount of time required to provision and enables the application of best practices.
View Composer database	<ul style="list-style-type: none">• Use the full version of the chosen database server (Microsoft SQL Server or Oracle).• Use the appropriate NetApp SnapManager® solution for database backup and recovery.
Profile management	Use a profile management solution rather than VMware View Composer persistent disks. Profile management solutions are more robust than persistent disks and make management, backup and restore, and disaster recovery of the user data within them easier.
User data: file shares	NetApp AFF or FAS storage can natively host file shares for user and department data, including user profiles. This provides built-in redundancy and HA, as well as storage efficiencies such as block-level deduplication and compression. NetApp Snapshot copies, SnapVault® technology, and SnapMirror provide intrasite and intersite data protection and replication capabilities for this critical data.

Best Practice Area	Best Practice Details
VM optimizations	<ul style="list-style-type: none"> For Windows 7, apply the OS settings described in the VMware Horizon Optimization Guide for Windows 7 through a group policy. For Windows 7, Windows 10, and other Windows client OSs, take advantage of the VMware OS Optimization Tool to automatically configure recommended settings on the desktop OS. Use new VM templates when deploying virtual desktops rather than reusing an existing desktop image. Do not use physical-to-virtual migration technologies to transition a user's existing laptop or desktop into VMware View. Set the Windows guest OS timeout value to either 60 or 190 seconds, as described in NetApp Knowledge Base (KB) article 3013622. Remove all transient data from the VM template before deploying virtual desktops. When using NFS with NetApp VSC, perform a space-reclamation process on the template to make the VM as small as possible.
VM deployment	<ul style="list-style-type: none"> For new persistent desktop environments, use VAAI or VSC rapid provisioning to create dedicated full clone VMs. When using linked clones in automated pools, NetApp recommends using disposable file redirection to help maintain performance and storage efficiency in the virtual desktop environment. When using linked clones, NetApp strongly recommends using the space-efficient sparse virtual disk format first made available with VMware View 5.2, VMware vSphere 5.1 patch1, and VM version 9.
Desktop pools	<ul style="list-style-type: none"> When using manual desktop pools, NetApp recommends dedicated user assignment. Although floating assignments can be chosen during pool creation, this architecture has limited use cases. When using manual desktop pools, use the NetApp VSC to provision the manual pools. This reduces the amount of time required to provision, enables the application of best practices, and creates storage-efficient VMs.
Maintenance activities	<ul style="list-style-type: none"> Conduct View Composer refresh and recompose operations during nonpeak hours. Use of the View Composer rebalance operation should only be done with caution because it can temporarily increase the storage requirements due to the rehydration of deduplicated data. NetApp recommends avoiding rebalance operations if they are not necessary within the environment. When using View Storage Accelerator, set blackout times so that cache regeneration does not occur during peak hours.

8 Solution Verification

This reference architecture is based on VMware vSphere 6.0 Update 1, VMware Horizon View 7, and VMware View Composer 7. This software was used to host, provision, and run 2,500 Microsoft Windows 10 virtual desktops on a Cisco UCS with B200 M4 blades. Backup was provided by a NetApp AFF8080cc storage system running the NetApp Data ONTAP 8.3.2 OS configured with 800GB SSDs. Ten datastores were presented from the NetApp system to the VMware vSphere hosts for use by the desktops. Host-to-host communication took place over the 10GbE Cisco Nexus 9396PX network through the VMware virtual network adapters. VMs were used for core infrastructure components such as VMware vCenter Server, View Connection Server, View Composer Server, AD, database servers, and other services.

Performance of this environment was verified using Login Virtual Session Indexer (Login VSI). Login VSI is the industry-standard load-testing tool for testing the performance and scalability of centralized Windows desktop environments such as server-based computing (SBC) and VDI.

Login VSI is used for testing and benchmarking by all major hardware and software vendors and is recommended by both leading IT analysts and the technical community. Login VSI is vendor independent and works with standardized user workloads. Therefore, conclusions based on Login VSI test data are objective, verifiable, and replicable.

During these performance tests, many different scenarios were tested to validate the performance of the storage during the lifecycle of a virtual desktop deployment.

The testing included the following criteria:

- Provisioning 2,500 VMware Horizon View linked clone desktops
 - 1,250 nonpersistent linked-clone desktops
 - 1,250 persistent full clone desktops
- Boot storm test of 2,500 desktops (with and without storage failover)
- Login VSI initial login and steady-state workload (with and without storage failover)
- Refresh operation of 1,250 nonpersistent desktops
- Recompose operation of 1,250 nonpersistent desktops

Conclusion

FlexPod Datacenter is the optimal infrastructure foundation on which to deploy a virtual desktop environment. Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and designs. This flexibility and scalability of FlexPod enable customers to start out with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements: from sub-1,000-seat VDI deployments to tens of thousands of seats.

Acknowledgements

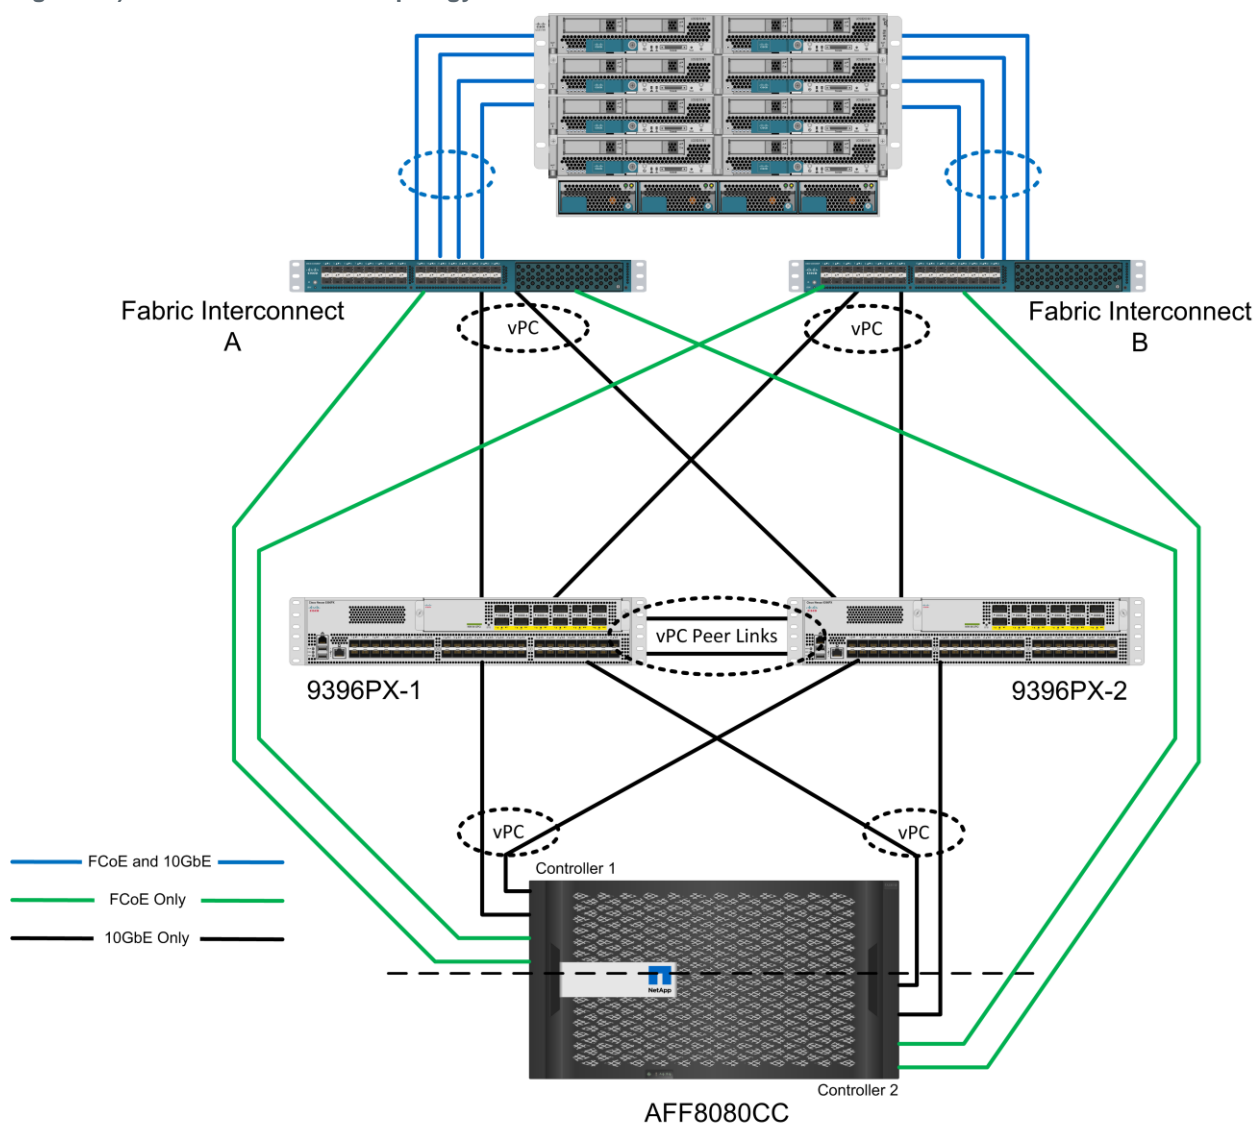
This document is the result of the work, documentation, and assistance provided by Chris Gebhardt, Eric Railine, and Bhavin Shah of NetApp. The author would also like to thank Mike Brennan of Cisco, Bhumik Patel of VMware, and Troy Mangum of NetApp for their contributions and support.

Appendix

FlexPod: FCoE Direct-Connect Design

The Cisco Nexus 9396PX switches used in this design support 10GbE and 40GbE networking and Cisco Application Centric Infrastructure, but do not support FC and FCoE protocols. Therefore, neither storage protocol is used in the validation of this design. FlexPod can still support these protocols by connecting the unified target adapter ports on the NetApp storage controllers directly to the Cisco UCS fabric interconnects. This configuration leverages the FC features of the fabric interconnects to provide name services and zoning capabilities. It also enables the Cisco UCS servers to boot from and access FC or FCoE storage without additional FC- or FCoE-capable switches. If you want to apply this solution design and use FCoE for some or all data services, you can modify your deployment to incorporate a direct connect topology.

Figure 18) FCoE direct-connect topology.



References

This section provides links to additional information and reference material for the subjects contained in this document.

Cisco UCS

The following links provide additional information about the Cisco UCS:

- Cisco Design Zone for Data Centers
<http://www.cisco.com/c/en/us/solutions/enterprise/design-zone-data-centers/index.html>

- Cisco UCS
<http://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>
- Cisco UCS 6200 Series Fabric Interconnects
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6200-series-fabric-interconnects/index.html>
- Cisco UCS 6300 Series Fabric Interconnects
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-6300-series-fabric-interconnects/index.html>
- Cisco UCS 5100 Series Blade Server Chassis
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-5100-series-blade-server-chassis/index.html>
- Cisco UCS B-Series Blade Servers
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-b-series-blade-servers/index.html>
- Cisco UCS Adapters
<http://www.cisco.com/c/en/us/products/interfaces-modules/unified-computing-system-adapters/index.html>
- Cisco UCS Manager
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-manager/index.html>

Cisco Nexus Networking

The following links provide additional information about Cisco Nexus 9000 Series switches:

- Cisco Nexus 9000 Series Switches
<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>
- Cisco Nexus 9000 Configuration Guides
<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-installation-and-configuration-guides-list.html>
- Cisco Nexus 9000 Series Switches Command References
<http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-command-reference-list.html>

NetApp AFF Storage

The following links provide additional information about NetApp AFF storage:

- All Flash FAS Datasheet
<http://www.netapp.com/us/media/ds-3582.pdf>
- NetApp Flash Advantage for All Flash FAS
<http://www.netapp.com/us/media/ds-3733.pdf>
- NetApp All Flash FAS Overview: Data ONTAP 8.3.2
<http://www.netapp.com/us/media/tr-4505.pdf>
- Clustered Data ONTAP 8.3.2 Documentation
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=62227>
- TR-4476: NetApp Data Compression and Deduplication: Data ONTAP 8.3.1 and Above
<http://www.netapp.com/us/media/tr-4476.pdf>
- Guest OS Tunings for a VMware vSphere Environment (KB)
<https://kb.netapp.com/support/index?page=content&id=3013622>
- TR-4428: NetApp All Flash FAS Solution for VMware Horizon 6 and vSphere Virtual Volumes
<http://www.netapp.com/us/media/tr-4428.pdf>

- TR-4333: VMware vSphere 6 on NetApp Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4333.pdf>

VMware vSphere

The following links provide additional information about VMware vSphere:

- VMware vSphere Documentation Center
<http://pubs.vmware.com/vsphere-60/index.jsp>
- Enabling Microsoft SQL Clustering Service in VMware vCenter Server 5.5 (KB)
<http://kb.vmware.com/kb/2059560>

VMware Horizon View

The following links provide additional information about VMware Horizon View:

- VMware Horizon View 7.0 Documentation Center
<http://pubs.vmware.com/horizon-7-view/index.jsp>
- VMware Horizon View Best Practices (KB)
<http://kb.vmware.com/kb/1020305>
- View Composer Array Integration (VCAI): Horizon Guide
https://www.vmware.com/resources/compatibility/pdf/vi_vcai_guide.pdf
- VMware Horizon 6 with View Performance and Best Practices (PDF)
<http://www.vmware.com/files/pdf/view/vmware-horizon-view-best-practices-performance-study.pdf>
- Optimization Guide for Desktops and Servers in View in VMware Horizon 6 and VMware Horizon Air Desktops and VMware Horizon Air Apps
<https://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf>
- VMware OS Optimization Tool
<https://labs.vmware.com/flings/vmware-os-optimization-tool>

Interoperability Matrixes

The following links provide information about interoperability tools:

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

Version History

The most current version of this document is available at <http://netapp.com/us/media/nva-1110-fp-design.pdf>.

Version	Date	Document Version History
Version 2.0	June 2016	Updated with clustered Data ONTAP 8.3.2, Cisco UCS Manager 3.1.1(1e), VMware vSphere 6.0 Update 1, and VMware Horizon View 7. iSCSI replaces the FCoE services used for boot LUNs in version 1.0.

Version	Date	Document Version History
Version 1.0	November 2014	Initial release with clustered NetApp Data ONTAP 8.2.1, Cisco UCS Manager 2.2.1(c), VMware vSphere 5.5, and VMware Horizon View 5.3.1.

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 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>. NVA-1110-FP-DESIGN