



NetApp Verified Architecture

FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF / FAS NVA Design Guide

Jyh-shing Chen, NetApp
January 2021 | NVA-1154-DESIGN

Abstract

The FlexPod® Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp® AFF / FAS leverages Cisco UCS Mini with B200 M5 Blade Servers, Cisco UCS 6324 Fabric Interconnects, Cisco Nexus 31108PC-V switches, or other compliant switches, and the NetApp AFF C190, AFF A220, or FAS2700 series controller HA pair, which runs NetApp ONTAP® 9.7 data management software. This NetApp Verified Architecture (NVA) design guide provides information needed to create a highly available, highly flexible, and highly scalable FlexPod Express-based VMware vSphere 7.0 virtual infrastructure.

In partnership with



TABLE OF CONTENTS

Executive summary	4
Program summary	4
Solution overview	5
FlexPod Converged Infrastructure program	5
NetApp Verified Architecture program	5
Target audience.....	6
Solution technology	6
Use case summary	7
Technology requirements	8
Hardware requirements	8
Hardware requirements using existing compliant switches	8
Software requirements	9
Design choices	10
NetApp AFF C190, AFF A220, or FAS2700 series with ONTAP 9.7	10
ONTAP 9.7	12
Cisco Nexus 3000 Series	13
Cisco UCS Mini	14
Cisco UCS B-Series	16
Cisco UCS Virtual Interface Cards	16
VMware vSphere 7.0	17
NetApp Virtual Storage Console 9.7.1	18
NetApp NFS Plug-in for VMware VAAI	18
NetApp SnapCenter Plug-In for VMware vSphere 4.4.....	18
NetApp Active IQ Unified Manager 9.7P1	19
Storage protocol support and load balancing	19
Solution scaling	26
Solution verification	28
Conclusion	29
Where to find additional information	29
Version history	29

LIST OF TABLES

Table 1) Hardware requirements for the base FlexPod Express with UCS Mini configuration.	8
Table 2) Hardware requirements for the FlexPod Express with UCS Mini and compliant switches configuration.	8
Table 3) Software requirements for the base FlexPod Express with UCS Mini implementation.	9
Table 4) Software requirements for a VMware vSphere implementation on the FlexPod Express with UCS Mini.	10
Table 5) Solution configuration differences between FC and FCoE protocols.	25
Table 6) AFF C190, AFF A220, and FAS 2700 series storage limits.	28

LIST OF FIGURES

Figure 1) FlexPod portfolio.	5
Figure 2) FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS architecture.	7
Figure 3) Using existing compliant switches for the solution.	9
Figure 4) FAS2750 rear view (UTA2).	12
Figure 5) Deploy ONTAP on your choice of architecture and location to provide data services where you need them.	13
Figure 6) Cisco UCS Mini front view.	15
Figure 7) Cisco UCS Mini rear view.	15
Figure 8) Cisco UCS 6324 Fabric Interconnect.	16
Figure 9) Cisco B200M5 Blade Server.	16
Figure 10) Cisco UCS VIC 1440.	17
Figure 11) Cisco UCS VIC 1480.	17
Figure 12) iSCSI protocol architecture for the solution.	20
Figure 13) iSCSI protocol architecture with failed cable between FI-A and ST-A.	21
Figure 14) NFS protocol architecture for the solution.	22
Figure 15) NFS protocol architecture with failed cable between FI-A and ST-A.	23
Figure 16) NFS protocol architecture with multiple failed links.	23
Figure 17) FC / FCoE protocol architecture for the solution.	25
Figure 18) Compute scaling to include 8 x B200 M5 Blade Servers and 4 x C220 M5 rack servers.	27

Executive summary

The FlexPod® Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp® AFF / FAS leverages Cisco UCS Mini with B200 M5 Blade Servers, Cisco Unified Computing System (UCS) 6324 Fabric Interconnects, Cisco Nexus 31108PC-V switches, or other compliant switches, and the NetApp AFF C190, AFF A220, or FAS2700 series controller HA pair, which runs NetApp ONTAP® 9.7 data management software.

This NetApp Verified Architecture (NVA) design guide provides information needed to create a highly available, highly flexible, and highly scalable FlexPod Express-based VMware vSphere 7.0 virtual infrastructure solution. In addition to covering the required hardware and software components and the design choices around the compute, network, and storage layers, this guide also provides information on the associated tools to simplify and optimize the solution ecosystem management. Recognizing that many customers already have an existing network infrastructure, this solution includes design choices for deploying the solution with existing compliant switches in the network infrastructure if they meet the requirements and the solution needs.

The design supports storage connectivity by using multiple storage protocols, including iSCSI, NFS, FC, and FCoE. A deployment that uses iSCSI SAN boot and NFS datastores for the VMware infrastructure was used to validate the design. The details of the deployment steps are available in the deployment guide for the solution.

In addition to leveraging configurations supported by the NetApp Interoperability Matrix Tool (IMT), VMware VCG, and Cisco HCL to ensure interoperability, a variety of tests were conducted to confirm solution behaviors under normal operations and many other maintenance and failure related scenarios to verify the availability of the data services. Therefore, customers should be confident that by adopting this FlexPod Express NVA solution design, they will have a reliable and scalable infrastructure that they can count on for their digital transformation journey.

Program summary

Industry trends indicate that a vast data center transformation is occurring towards a hybrid cloud infrastructure with on-premises shared infrastructure, cloud computing, and the connectivity enabled by data fabric powered by NetApp to seamlessly provide data where it is needed. In addition, organizations seek a simple and effective solution for remote and branch offices that uses technology that they are familiar with in their data center.

FlexPod Express with UCS Mini and NetApp AFF/FAS is a validated, best practice data center architecture that is built on the Cisco UCS, the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the same set of tools they are familiar with. New FlexPod Express customers can easily scale and manage their FlexPod solutions as they scale and grow their environment.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload. FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure program.

FlexPod Converged Infrastructure program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Based on customer requirements, you can update a given CVD or NVA configuration to meet customer needs as long as the changes do not create an unsupported configuration.

As shown in Figure 1, the FlexPod program includes two solutions: FlexPod Express and FlexPod Datacenter.

- FlexPod Express offers customers an entry-level solution with technologies available from Cisco and NetApp.
- FlexPod Datacenter delivers an optimal multipurpose foundation for various workloads and applications for the data center.

Figure 1) FlexPod portfolio.

The FlexPod Portfolio

A prevalidated, flexible platform that features



NetApp Verified Architecture program

The NVA program offers customers a verified architecture for NetApp solutions. An NVA provides a NetApp solution architecture with the following qualities:

- Thoroughly tested
- Prescriptive in nature
- Minimized deployment risks
- Accelerated time to market

This guide details the design of the VMware vSphere 7.0 solution on FlexPod Express with UCS Mini and NetApp AFF/FAS storage. The following sections list the components used for the design of this solution.

Hardware components

- Cisco UCS Mini with in-chassis UCS 6324 Fabric Interconnects

- Cisco UCS B200 M5 with Virtual Interface Card (VIC) 1440 or other compatible VICs
- Cisco Nexus 31108PC-V or other compliant switches
- NetApp AFF C190, AFF A220, or FAS 2700 series storage system

Software components

- Cisco NXOS Firmware 9.3(5)
- Cisco UCS Manager 4.1(2a)
- NetApp ONTAP 9.7
- NetApp Virtual Storage Console 9.7.1
- NetApp SnapCenter® Plug-In for VMware vSphere 4.4
- NetApp Active IQ® Unified Manager 9.7P1
- VMWare vSphere 7.0

Target audience

This document is intended for people who want to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation.


The audience for this document includes, but is not limited to, the following groups:

- Customers
- Sales engineers
- Field consultants
- Professional services personnel
- Partner engineers
- IT architects
- IT managers

Solution technology

This solution leverages technologies from NetApp, Cisco, and VMware. It features NetApp AFF C190 / A220 or FAS 2700 series storage systems running ONTAP 9.7, dual Cisco Nexus 31108PC-V switches, and Cisco UCS Mini with UCS B200 M5 servers that run VMware vSphere 7.0. Figure 2 shows an architecture of this validated solution. The AFF C190, AFF A220, or FAS 2700 series is directly attached to the UCS 6324 in chassis Fabric Interconnects, shown next to the UCS Mini chassis in the figure for cabling illustrations.

Figure 2) FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS architecture.



During normal operations, the storage data path from the vSphere 7.0 hosts, running on the UCS B200 M5 blades, to storage are going from the virtual network interface card (NIC) connected to the UCS 6324 Fabric Interconnect to the AFF / FAS storage directly, without going through the Fabric Interconnect uplink ports and switches. However, for certain failure scenarios, the NFS I/O will traverse the Fabric Interconnect uplink ports, and across the virtual port channel (vPC) peer links, if necessary, in order to reach storage for continued data services.

For deployments into existing network infrastructures that are compliant, as detailed in the following “Hardware requirements” section, you can connect the Fabric Interconnect uplinks to 10GbE ports in the existing network infrastructure without the Cisco Nexus 31108PC-V switches shown in the architecture diagram.

For applications or solutions that do not require high storage data bandwidth, the existing network infrastructure can be 1GbE in speed if it is sufficient for the solution design.

Use case summary

You can apply the FlexPod Express solution to several use cases, including the following:

- Remote offices or branch offices (ROBOs)
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is ideal for virtualized and mixed workloads.

Technology requirements

A FlexPod Express system requires a combination of hardware and software components. In addition to the required hardware and software components, you can add additional hardware components to scale up the solution. Furthermore, you can add additional software and applications to help manage the solution or provide additional functionalities.

Hardware requirements

Depending on your business requirements, you can use different hypervisors on the same reference FlexPod Express with UCS Mini hardware configuration.

Table 1 lists the reference hardware components for a FlexPod Express with UCS Mini configuration.

Table 1) Hardware requirements for the base FlexPod Express with UCS Mini configuration.

Hardware	Quantity
NetApp AFF C190, AFF A220, or FAS 2700 series HA pair	1
Cisco Nexus 3000 Series switches	2
Cisco UCS Mini with two integrated UCS-FI-M-6324 Fabric Interconnects	1
Cisco UCS B200 M5 server with VIC 1440 / 1340	2

Note: The actual hardware components that are selected for a solution implementation can vary based on customer requirements. For example, instead of using an AFF C190 HA pair, you can use an AFF A220 HA pair or an FAS 2700 series controller HA pair to meet the performance or cost requirements.

Note: While this design can be implemented with various hardware choices based on the actual customer solution requirements, the associated NVA deployment guide uses an AFF A220 HA pair for storage and a pair of Cisco Nexus 31108PC-V switches for networking to demonstrate a design implementation.

Note: The management network and console connections for the FlexPod components are assumed to be connected to an existing infrastructure, which is deployment specific, and thus not detailed in the design or deployment guide.

Hardware requirements using existing compliant switches

For a customer deployment scenario where the environment already has an existing network infrastructure with compliant switches that meet the requirements below, you can replace the Cisco Nexus 3000 series switches with the compliant switches as shown in Table 2.

- The switches must support 802.1Q VLAN tagging and be configured to pass the required VLAN traffic between the two Fabric Interconnects.
- The switches should be in a redundant configuration and configured with the equivalent of Cisco vPC functionality. Not meeting this requirement will make the solution not available during switch reboot, upgrade, or failure scenarios.
- It is preferred that the switches have two available 10GbE ports each for the UCS 6324 Fabric Interconnect uplinks. However, if the existing infrastructure supports only 1GbE speed and the 1GbE speed meets the solution requirements, then you can use the 1GbE ports on the switches with proper supporting hardware and configurations.

Table 2) Hardware requirements for the FlexPod Express with UCS Mini and compliant switches configuration.


Hardware	Quantity
NetApp AFF C190, AFF A220, or FAS 2700 series HA pair	1

Hardware	Quantity
Compliant network switches	2
Cisco UCS Mini with two integrated UCS-FI-M-6324 Fabric Interconnects	1
Cisco UCS B200 M5 server with VIC 1440 / 1340	2

Figure 3 illustrates such a deployment scenario. The switches must have proper VLAN and vPC, or vendor equivalent, configurations in order for the solution to properly manage and pass the NFS I/O through the existing compliant switches for certain failure scenarios.

See the specific switch vendor documentation for information about implementing the required switch configurations by using the deployment document as a reference guide.

Figure 3) Using existing compliant switches for the solution.



Software requirements

Table 3 lists the software components that are required to implement the base FlexPod Express with UCS Mini solution.

Table 3) Software requirements for the base FlexPod Express with UCS Mini implementation.

Software	Version	Details
Cisco UCS Manager	4.1(2a)	For Cisco UCS 6324 Fabric Interconnects
Cisco blade software	4.1(2a)	For UCS B200 M5 servers
Cisco nenic driver	1.0.33.0	For VIC 1440 / 1340 interface cards
Cisco NX-OS	9.3(5)	For Cisco Nexus 31108PC-V switches
NetApp ONTAP	9.7	For AFF / FAS storage system

Table 4 lists the software that are required for a VMware vSphere implementation on FlexPod Express with UCS Mini.

Table 4) Software requirements for a VMware vSphere implementation on the FlexPod Express with UCS Mini.

Software	Version
VMware vSphere ESXi hypervisor	7.0
VMware vCenter server appliance	7.0
NetApp VAAI Plug-In for ESXi	1.1.2
NetApp Virtual Storage Console	9.7.1
NetApp SnapCenter Plug-In for VMware vSphere	4.4
NetApp Active IQ Unified Manager	9.7P1

Design choices

The following technologies were chosen during the process of architecting this design. Each technology serves a specific purpose in the FlexPod Express infrastructure solution.

NetApp AFF C190, AFF A220, or FAS2700 series with ONTAP 9.7

This solution uses NetApp AFF C190, AFF A220, or FAS 2700 series product families with ONTAP 9.7.

NetApp all-flash storage systems are critical for speeding up enterprise applications and creating a modern hybrid cloud infrastructure. They provide comprehensive data services, integrated data protection, seamless scalability, new levels of performance, and cloud integration. With an AFF system, you can make an easy and risk-free transition to flash for your digital transformation.

These products can be ordered with two different flavors for host connectivity: UTA2 and 10GBASE-T. The 10GBASE-T version supports only Ethernet connectivity. For the UTA2 version, you can configure the data ports to support either Ethernet or FC and they require appropriate supporting small form-factor pluggable (SFP) transceivers. For this solution design validation, we used the UTA2 flavor of the AFF A220 HA pair.

AFF C190 system

With NetApp's AFF C190 entry-level storage system, the benefits of all-flash array can be realized even for ROBO and many other shared infrastructure or dedicated application uses. It can support up to 24x 960GB internal SSDs and is ideal for an environment that does not require additional flash capacity.

Figure 5) AFF C190 front view.




Figure 6) AFF C190 rear view (UTA2).



For more information about the AFF C190 hardware system, see the [NetApp AFF C190 Datasheet](#).

AFF A220 system

The AFF A220 array for small and medium enterprise environments delivers 30% more performance than its predecessor to continue NetApp's leadership in this segment. With AFF A220, the system supports both 24 internal SSDs and external disk shelves. Check the [NetApp Hardware Universe](#) for the available capacities of the supported SSDs.

In addition, you can use the AFF A220 array to deploy a MetroCluster IP solution where customers can protect their business-critical data by synchronously replicating data from a clustered storage system to another cluster, located up to 700km apart, to ensure continuous data availability despite a potential site disaster.

Figure 3) AFF A220 front view.





Figure 4) AFF A220 rear view (UTA2).



For more information about the AFF A220 hardware system, see the [NetApp AFF A220 Tech Spec](#).

FAS2720/FAS2750 system

The NetApp FAS2700 series hybrid storage arrays offer more value than other systems in their class. The FAS2700 running NetApp ONTAP storage software simplifies the task of managing growth and complexity by delivering high performance, providing leading integration with the cloud, supporting a broader range of workloads, and seamlessly scaling performance and capacity.

Figure 7) FAS2750 front view.




Figure 4) FAS2750 rear view (UTA2).



For more information about the FAS2700 series system, see the [FAS2700 Hybrid Storage System product page](#).

ONTAP 9.7

NetApp ONTAP is the industry's leading enterprise data management software. It offers powerful data management capabilities, proven storage efficiencies, and seamless cloud integration. It provides a foundation for hybrid cloud data management and enables data mobility for your data fabric.

ONTAP 9.7 has many features that are suited for the FlexPod Express solution. Foremost is NetApp's commitment to storage efficiencies, which can be one of the most important features for small deployments. You can build your data fabric to span from the edge to core to cloud by using the platform of your choice. NetApp SnapMirror® technology enables seamless data mobility for your data fabric. To ensure uninterrupted data services, ONTAP MetroCluster enables multi-site solution deployment to meet zero recovery point objectives (RPO) and low recovery time objectives (RTO) to ensure business continuity despite single site outage. To efficiently use your flash storage, NetApp FabricPool enables policy-based management for storage tiering of the infrequently accessed data on the system to the object storage in the public cloud or on-premises.

Storage efficiency

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduces the total logical capacity used to store customer data up to a data reduction ratio of 7:1, based on the workload. This space reduction is enabled by a combination of several different technologies, including deduplication, compression, and compaction.

SnapMirror

SnapMirror is an asynchronous replication technology for data replication across different sites, within the same data center, on-premises data center to cloud, or cloud to on-premises data center. NetApp SnapMirror Synchronous (SM-S) offers volume granular, zero data loss protection. It extends traditional SnapMirror volume replication to synchronous mode, meeting zero RPO disaster recovery, and compliance objectives.

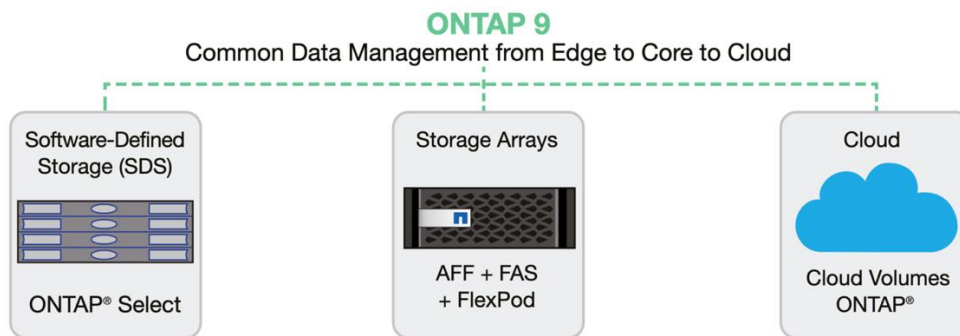
MetroCluster

MetroCluster configurations combine array-based clustering with synchronous data replication to a second site to deliver continuous data availability. The multisite ONTAP MetroCluster solution enhances the built-in high availability and nondisruptive operations of a single-site FlexPod solution by providing an additional layer of data protection to ensure data availability in the event of a single-site disaster.

Data fabric from the edge to core to cloud

With ONTAP 9.7, you can create an environment that is smart, powerful, and trusted that spans flash, disk, and cloud, which is ready to handle your current and future requirements. You can deploy storage on your choice of architectures—hardware storage systems, software-defined storage, and the cloud—while unifying data management across all of them, as shown in Figure 5.

Figure 5) Deploy ONTAP on your choice of architecture and location to provide data services where you need them.



You can accelerate your enterprise applications with flash, without compromising on the essential data services that you need, and seamlessly manage your data as it flows between the edge, core, and cloud to wherever you need it.

ONTAP enables the data in the FlexPod Express solution to seamlessly move from the edge to core to cloud for purposes such as data replication and big data analysis.

FabricPool

FabricPool is a hybrid storage solution with ONTAP 9 that uses an all-flash (SSD) aggregate as a performance tier and an object store in the public cloud service or on premises as a cloud tier. This configuration enables policy-based data movement, depending on whether or not data is frequently accessed. FabricPool is supported in ONTAP for both AFF and all-SSD aggregates on FAS platforms. Data processing is performed at the block level, with frequently accessed data blocks in the all-flash performance tier tagged as hot and infrequently accessed blocks tagged as cold.

Configuring FabricPool on the FlexPod Express solution deployed at a remote/edge location can help achieve more efficient flash storage usage by tiering the cold data to a public cloud such as Amazon Web Services (AWS).

For more information about ONTAP data management software features and capabilities, see the [ONTAP 9 Data Management Software Datasheet](#), [ONTAP 9 Documentation Center](#), and additional information linked in the Where to find additional information section.

Cisco Nexus 3000 Series

Because all the various Cisco Nexus series models run the same underlying operating system, NX-OS, multiple Cisco Nexus models are supported in the FlexPod Express and FlexPod Datacenter solutions.

The Cisco Nexus 31108PC-V is a robust, cost effective switch offering 1/10/40/100Gbps switching. It offers 48 1/10-Gbps ports, and 40/100-Gbps uplinks that enable flexibility.

The Cisco Nexus 31108PC-V provides a comprehensive layer 2 feature set that includes virtual LANs (VLANs), IEEE 802.1Q trunking, and the Link Aggregation Control Protocol (LACP). Additional layer 3 functionality is available by adding licenses to the system.

Figure 6) Cisco Nexus 31108.



For more information about the Cisco Nexus 31108PC-V, see the [Cisco Nexus 31108PC-V Switch product information](#).

Cisco UCS Mini

The Cisco Unified Computing System (Cisco UCS) is an integrated and scalable platform designed to reduce total cost of ownership (TCO) by simplifying management and unifying compute, network, and storage access. The system features enterprise-class x86-architecture servers, low-latency, and lossless unified network fabric to consolidate LAN, SAN, and computing network. The Cisco UCS unified storage access supports storage over Ethernet (NFS or iSCSI), FC, and FCoE.

The Cisco UCS 5108 Blade Server Chassis revolutionizes the use and deployment of blade-based systems. By incorporating unified fabric and fabric-extender technology, the Cisco UCS enables the chassis to:

- Have fewer physical components
- Require no independent management
- Be more energy efficient than traditional blade-server chassis

This simplicity eliminates the need for dedicated chassis management and blade switches, reduces cabling, and allows scalability to 20 chassis without adding complexity. The Cisco UCS 5108 Blade Server Chassis is a critical component in delivering the simplicity and IT responsiveness for the data center as part of the Cisco UCS. The policy-based management, simplified storage connectivity, and advanced monitoring and management capabilities of the solution greatly increase productivity and lower TCO.

With Cisco UCS Mini, the UCS 6324 Fabric Interconnects adopt the I/O module (IOM) form factor and are directly integrated into the UCS 5108 chassis to simplify the configuration for environment with less resource requirements. Figure 6 and Figure 7 show the front and back views of a fully populated UCS Mini with eight half-width server blades and two UCS 6324 Fabric Interconnects.

Figure 6) Cisco UCS Mini front view.

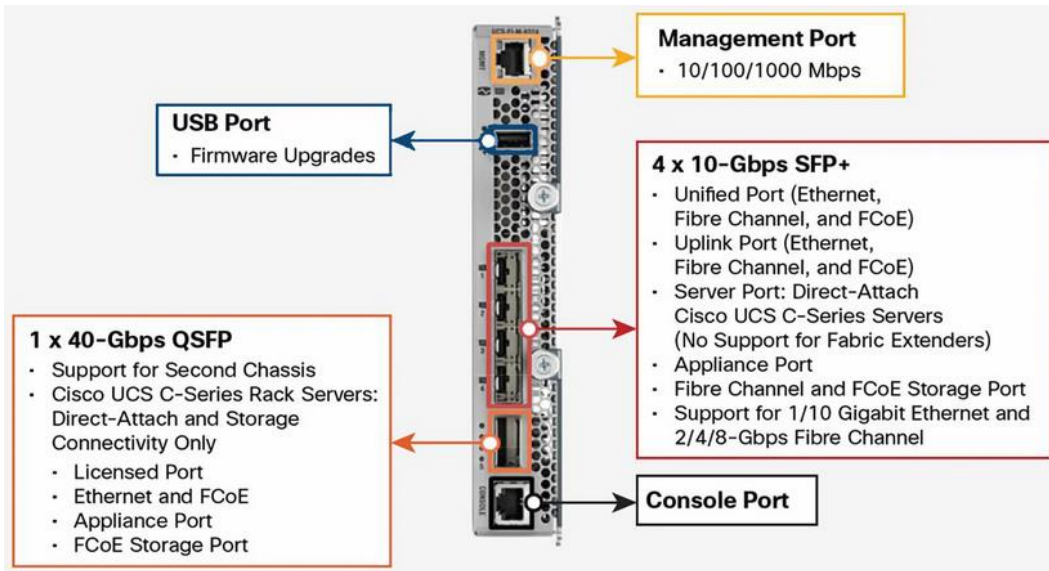


Figure 7) Cisco UCS Mini rear view.



The UCS 6324 Fabric Interconnect has twenty-four 10G ports. Sixteen of its 10G ports are internal and are dedicated to the eight half-width blade slots. Four of the remaining eight 10G ports are available externally as 1/10G SFP+ interfaces and the remaining four are available as one 40G QSFP interface that can be connected to up to four external servers, or to a second UCS 5108 chassis, to provide additional compute resources, or to additional storage, all by using a 40G to 4 x 10G splitter cable. Figure 8 shows the UCS 6324 Fabric Interconnect and its 4 x 1/10G SFP+, 1 x QSFP, management network, USB, and console interfaces.

Figure 8) Cisco UCS 6324 Fabric Interconnect.



Cisco UCS B-Series

The Cisco UCS B-Series B200M5 server was chosen for the FlexPod Express because of its many configuration options, which allows it to be tailored for specific requirements in FlexPod Express deployment.

Figure 9) Cisco B200M5 Blade Server.




The enterprise-class Cisco UCS B200 M5 Blade Server extends the capabilities of Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M5 Blade Server harnesses the power of the latest Intel Xeon processor scalable family CPUs with up to 3072 GB of RAM (using 128GB DIMMs), two SSDs or HDDs, and up to 80Gbps throughput connectivity. The support for UCS VIC is provided by using a modular LAN on Motherboard (mLOM) card using VIC 1440 / 1340 or a rear mezzanine adapter using VIC 1480 / 1380.

For more information about the Cisco UCS B200 M5 Blade Server, see [Cisco UCS B200 M5 Blade Server Spec Sheet](#).

Cisco UCS Virtual Interface Cards


The Cisco UCS VIC 1440 (Figure 10) is a dual-port 40Gbps or dual 4x 10Gbps Ethernet/FCoE capable mLOM designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers. When used with an optional port expander, the Cisco UCS VIC 1440 capabilities are enabled for two ports of 40Gbps Ethernet (two ports of 20Gbps with UCS Mini). The Cisco UCS VIC 1440 enables a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs.

Figure 10) Cisco UCS VIC 1440.



The Cisco UCS VIC 1480 (Figure 11) is a single-port 40Gbps or 4x10Gbps Ethernet/FCoE capable mezzanine card (mezz) designed exclusively for the M5 generation of Cisco UCS B-Series Blade Servers. The card enables a policy-based, stateless, agile server infrastructure that can present PCIe standards-compliant interfaces to the host that can be dynamically configured as either NICs or HBAs.

Figure 11) Cisco UCS VIC 1480.



In addition to the Cisco VIC 1440/1480 designed for B200M5, the third-generation Cisco VIC 1340/1380 compatible with B200 M3/M4/M5 are also supported. For a comparison of the various Cisco UCS VICs, see the adapter [Compare Model](#) page.

For more information about Cisco VIC 1440/1480, see [Cisco UCS Virtual Interface Card 1400 Series Data Sheet](#).

VMware vSphere 7.0

VMware vSphere 7.0 is one hypervisor option for use with FlexPod Express. VMware vSphere allows organizations to reduce their power and cooling footprint while confirming that the purchased compute capacity is used to its fullest. In addition, VMware vSphere allows hardware failure protection (VMware High Availability, or VMware HA) and compute resource load balancing across a cluster of vSphere hosts (VMware Distributed Resource Scheduler, or VMware DRS).

VMware vSphere 7.0 features the latest VMware innovations. The VMware vCenter Server Appliance (vCSA) that is used in this design adds a host of features and functionality, such as VMware vMotion and VMware vSphere Update Manager integration. To add clustering capability to hosts and to use features such as VMware HA and VMware DRS, the VMware vCenter Server is required.

VMware vSphere 7.0 introduces support for Kubernetes. This allows vSphere administrators to have both VMs and containers coexist on the same vSphere infrastructure with the vSphere Pod Service. To have a fully compliant and conformant Kubernetes cluster, you can use the Tanzu Kubernetes Grid Service.

There are additional enhancements on the core features, such as improved distributed resource scheduler for more granular resource optimizations, a new Assignable Hardware framework for better hardware acceleration support, and refactored vMotion to better support live migration of VMs that use large memory and CPU resources.

Note: For more information about VMware vSphere, see the [vSphere](#) overview page.

Note: For more information about the new features of VMware vSphere 7.0, see [VMware vSphere 7.0 Release Notes](#).

Note: For ONTAP 9.7 with VMware HCL support, see [VMware Compatibility Guide](#).

NetApp Virtual Storage Console 9.7.1

NetApp Virtual Storage Console (VSC) for VMware vSphere is a vCenter web client plug-in that provides end-to-end lifecycle management for virtual machines (VMs) in VMware environments that use NetApp AFF and FAS storage systems. VSC provides visibility into the NetApp storage environment from within the vSphere web client. VMware administrators can easily perform tasks that improve both server and storage efficiency while still using role-based access control to define the operations that administrators can perform.

The 9.7.1 release of the virtual appliance for VSC, VASA Provider, and Storage Replication Adapter (SRA) provides the combined features of VSC, VASA Provider, and SRA in a single deployment. It includes enhanced REST APIs that provide vVols metrics for SAN storage systems using ONTAP 9.7 and later. So, NetApp OnCommand API Services is no longer required to get metrics for ONTAP systems 9.7 and later.

NetApp NFS Plug-in for VMware VAAI

The NetApp NFS Plug-in for VMware vStorage APIs - Array Integration (VAAI) is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. The VMware VAAI package enables the offloading of certain tasks from the physical hosts to the storage array. Performing those tasks at the array level can reduce the workload on the ESXi hosts.

The copy offload feature and space reservation feature improve the performance of VSC operations. The NetApp NFS Plug-in for VAAI is not shipped with VSC, but you can install it by using VSC. You can download the plug-in installation package and obtain the instructions for installing the plug-in from the [NetApp Support Site](#).

For more information about the NetApp VSC for VMware vSphere, see [NetApp Virtual Infrastructure Management product page](#).

NetApp SnapCenter Plug-In for VMware vSphere 4.4

NetApp SnapCenter Plug-in for VMware vSphere enables VM-consistent and crash-consistent backup and restore operations for VMs and datastores from the vCenter server. The SnapCenter plug-in is deployed as a virtual appliance and it integrates with the vCenter server web client GUI.

Here are some of the functionalities provided by the SnapCenter plug-in to help protect your VMs and datastores.

- Backup VMs, virtual machine disks (VMDKs), and datastores
 - You can back up VMs, underlying VMDKs, and datastores. When you back up a datastore, you back up all the VMs in that datastore.
 - You can create mirror copies of backups on another volume that has a SnapMirror relationship to the primary backup or perform a disk-to-disk backup replication on another volume that has a NetApp SnapVault® relationship to the primary backup volume.
 - Backup operations are performed on all the resources defined in a resource group. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.
- Restore VMs and VMDKs from backups
 - You can restore VMs from either a primary or secondary backup to the same ESXi server. When you restore a VM, you overwrite the existing content with the backup copy that you select.
 - You can restore one or more VMDKs on a VM to the same datastore. You can restore existing VMDKs, or deleted or detached VMDKs from either a primary or a secondary backup.
 - You can attach one or more VMDKs from a primary or secondary backup to the parent VM (the same VM that the VMDK was originally associated with) or an alternate VM. You can detach the VMDK after you have restored the files you need.
 - You can restore a deleted VM from a datastore primary or secondary backup to an ESXi host that you select.

Note: For application-consistent backup and restore operations, the NetApp SnapCenter Server software is required.

Note: For additional information, requirements, licensing, and limitations of the NetApp SnapCenter Plug-In for VMware vSphere, see [NetApp Product Documentation](#) for details.

NetApp Active IQ Unified Manager 9.7P1

NetApp Active IQ Unified Manager is a comprehensive monitoring and proactive management tool for NetApp ONTAP systems to help manage the availability, capacity, protection, and performance risks of your storage systems and virtual infrastructure. You can deploy Unified Manager on a Linux server, on a Windows server, or as a virtual appliance on a VMware host.

Active IQ Unified Manager enables monitoring your ONTAP storage clusters, VMware vCenter server and VMs from a single redesigned, intuitive interface that delivers intelligence from community wisdom and AI analytics.

It provides comprehensive operational, performance, and proactive insights into the storage environment and the VMs running on it. When an issue occurs on the storage or virtual infrastructure, Active IQ Unified Manager can notify you about the details of the issue to help with identifying the root cause.

The VM dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the vSphere host down through the network and finally to the storage.

Some events also provide remedial actions that can be taken to rectify the issue. You can also configure custom alerts for events so that when issues occur, you are notified through email and SNMP traps.

Storage protocol support and load balancing

The Cisco UCS unified storage access supports storage over Ethernet (iSCSI or NFS), FC, and FCoE. The following sections provide an overview of the solution architecture for the different protocols. Information on load balancing and protocol behaviors associated with some failure scenarios are covered as well.

iSCSI protocol


The 10G network connectivity provided in the solution architecture contains two separate storage fabrics. It provides the iSCSI storage protocol with multipath support to protect the solution from single-point-of-failure scenarios. Virtual NICs (vNICs) capable of iSCSI SAN boot are created for the blade servers to use iSCSI protocol and also boot from the iSCSI SAN.

As illustrated in Figure 12, each B200 M5 Blade Server has access to both iSCSI fabrics through UCS 6324 Fabric Interconnect A and B (FI-A and FI-B). The storage controllers in the cluster, ST-A and ST-B, also have access to both FI-A and FI-B.

NetApp ONTAP provides clustered storage access. A LUN that is created on a volume in an aggregate owned by a controller can be accessed from iSCSI LIFs on both controllers. As a result, there are two iSCSI paths through fabric A and two iSCSI paths through fabric B to access a particular LUN under normal conditions.

Out of these four available paths, two of them are active/optimized and the other two are active/non-optimized. Whether a path is active/optimized or active/non-optimized depends on which controller owns the aggregate on which the LUN and volume reside. For example, if a LUN resides on a volume in an aggregate owned by ST-A, then a path through ST-A is active/optimized and a path through its HA partner, ST-B, is active/non-optimized.

Figure 12) iSCSI protocol architecture for the solution.



The multipath software in the VMware operating system recognizes the property of the available paths based on the SCSI asymmetric logical unit access (ALUA) information reported by the ONTAP iSCSI target, which indicates whether a particular path is active/optimized or active/non-optimized.


Under normal operating conditions, the VMware operating system uses both active/optimized paths: one through FI-A and one through FI-B. Using multiple paths improves I/O performance and balances the iSCSI I/O between the two fabrics. Similarly, using LUNs from both ST-A and ST-B helps balance the iSCSI I/O between the two storage controllers.

Depending on the failure scenarios, the number of paths to a LUN can change. When that happens, the VMware operating system continues to use active/optimized paths if they are available. If not, active/non-optimized paths are used.

For the example shown in Figure 13, the connectivity between FI-A and ST-A encountered a failure, as indicated by the red cross mark. For a LUN residing on ST-A, the active/optimized path through fabric A was disrupted. However, the active/optimized path through fabric B to ST-A and the two active/non-optimized paths through ST-B were not affected.

In such a failure scenario, the VMware operating system reports one dead path, one active/optimized path, and two active/non-optimized paths. The available active/optimized path is used for iSCSI storage I/O for a LUN residing on ST-A. Because the iSCSI I/O to the LUN residing on ST-A can only go through fabric B, such a failure scenario leads to unbalanced traffic on the two fabrics. After the failure is resolved and the connectivity between FI-A and ST-A is restored, the VMware operating system sees updated ALUA information and resumes I/O on both active/optimized paths again for the LUN residing on ST-A.

Figure 13) iSCSI protocol architecture with failed cable between FI-A and ST-A.



Note: For iSCSI protocol, I/O between the blade servers and storage does not need to traverse through the Cisco Nexus 31108PC-V switches under either normal or failure conditions.

Note: iSCSI SAN boot of the VMware hosts is used by the solution. As long as there is at least one available path, a VMware host can boot up from the available path in the iSCSI SAN.


Note: By using iSCSI SAN boot, you can add a new server to the server pool to replace a failed server. In that case, you can use the same service profile and the same SAN boot LUN to boot the new server for it to assume the configuration and identity of the server it replaces.

NFS protocol

The 10G network connectivity provided in the solution architecture supports additional protocols, such as NFS. When using NFSv3 protocol for storage I/O, an NFS client uses a single path to access the ONTAP NFS server through an NFS LIF.

As illustrated in Figure 18, the B200 M5 Blade Servers can use the same 10G network connectivity to access to the ONTAP NFS server through the NFS LIFs created on the data ports that are connected to FI-A in both ST-A and ST-B. For optimal access, an NFS client should use the NFS LIF from the same storage controller where the NFS volume resides.

Figure 14) NFS protocol architecture for the solution.




To support high availability for failure scenarios, an NFS LIF is automatically migrated to another data port on the same controller, or its HA partner in the case of a storage failover, in order to maintain continued data services for NFS clients.

For the example shown in Figure 15, the connectivity between FI-A and ST-A encountered a failure as indicated by the red cross mark. Due to the link failure, the NFS LIF located on the port with the failed link in ST-A is automatically migrated to the other port that is connected to FI-B.


For an NFS client on a blade server to reach the same NFS LIF on ST-A, which is now connected to FI-B, the NFS I/O can no longer go through FI-A to reach ST-A directly. Instead, the NFS I/O going through FI-A must traverse the uplink to the Nexus switch A or switch B. From within the same uplink switch, it can reach FI-B through the FI-B switch uplink port. Then, through FI-B, the NFS I/O can reach the NFS LIF that was migrated to the port connected to fabric B in ST-A.

Figure 15) NFS protocol architecture with failed cable between FI-A and ST-A.



In the case where an additional failure prevents the NFS I/O arriving at SW-A to traverse between FI-A and FI-B from within SW-A, the NFS I/O can go through the vPC switch peer links with SW-B to reach FI-B and ST-A for the scenario illustrated in Figure 16.

Figure 16) NFS protocol architecture with multiple failed links.



To accommodate these potential failure handling requirements, you must implement appropriate VLAN and vPC (or the specific compliant switch vendor equivalent) configurations on the compliant switches if an existing network infrastructure is used for this FlexPod deployment without the dedicated Nexus switch pairs.

You can take different approaches to help balance NFS I/O, which normally runs on the A fabric in this design. For example, you can pin other VMware functionalities, such as vMotion, to use the B fabric during normal conditions. In addition, you can create another pair of NFS LIFs on the controllers that use fabric B and configure some clients to use fabric B for NFS I/O during normal conditions.

Note: To use both controllers and the storage associated with them, you should create NFS volumes on both controllers.

Note: When attempting to balance NFS I/O load across both fabrics and storage controllers by using additional NFS LIFs and NFS volumes, make sure appropriate configurations are made throughout the infrastructure so the NFS I/O can be configured for optimal access and the solution can withstand various failure scenarios.


Note: You can deploy NFS protocol along with the iSCSI protocol and iSCSI SAN boot to take advantage of the multi-protocol support in ONTAP and to simplify server failure replacement scenario.

FC and FCoE protocols

When deploying FC or FCoE protocols on the UCS Mini solution with the NetApp AFF/FAS storage directly attached to the UCS 6324 Fabric Interconnects, there are no additional FC switching hardware requirements.

Figure 17 illustrates the FC / FCoE protocol architecture for the solution. The zoning and switching of the FC / FCoE traffic for the solution is done entirely by the UCS 6324 Fabric Interconnects and managed by the UCS Manager. The UCS 6324 Fabric Interconnects must be in FC switching mode instead of the default end-host mode. Virtual HBAs (vHBAs) are created for the servers for them to use FC / FCoE protocols.

Figure 17) FC / FCoE protocol architecture for the solution.



While the solution architecture diagram for FC and FCoE protocols are the same, there are some differences in the port configurations of the AFF / FAS data ports, the UCS 6324 Fabric Interconnect ports, and the type of physical cables required. Table 5 lists the solution configuration differences between the FC and FCoE protocols.

Table 5) Solution configuration differences between FC and FCoE protocols.

Protocol	FC	FCoE
AFF / FAS data port configuration	FC mode	CNA mode
Cables between the AFF / FAS and UCS 6324 Fabric Interconnects	Multi-mode fiber	Ethernet (copper / fiber)
Protocol speed	8 Gbps	10 Gbps

Note: In order to have the flexibility of deploying either FC or FCoE protocols, the UTA2 configuration of the AFF / FAS storage system is required, along with the supported SFPs for the desired protocol usage.

Note: Use CNA mode on AFF / FAS data ports and 10Gbps Ethernet connectivity between the AFF / FAS and the UCS 6324 Fabric Interconnects to support unified access to storage with iSCSI, NFS, and FCoE protocols.

From SAN boot and load balancing / high availability perspectives, a solution configuration using FC / FCoE protocol behaves similarly to that of the iSCSI protocol configuration as they all follow the SCSI standard and use standard based ALUA and multipath behaviors for the SCSI interactions between the AFF / FAS storage and the VMware hosts.

As a result, the iSCSI protocol behavior for load balancing and failure handling discussions above apply to the FC and FCoE protocols as well. The VMware operating system will discover ALUA LUN path information and use active/optimized paths when available and balance I/O among them automatically. If a failure scenario causes active/optimized paths to be disrupted, active/non-optimized paths will be used for resiliency.

Refer to the deployment guide of the solution for procedures on configuring the solution components, deploying iSCSI SAN booted VMware hosts, configuring NFS datastores, and deploying additional tools for the solution. For information about the FC and FCoE protocol configurations, which is not provided in the deployment guide, see the NetApp, Cisco, and VMware documentation for details.

Solution scaling

The FlexPod Express UCS Mini solution is highly available, highly flexible, and highly scalable. The solution hardware can be easily scaled both in compute and storage resources using various approaches.


Compute scaling

For a minimum VMware 7.0 solution deployment, you need two blade servers to create a minimum VMware cluster with vMotion and VMware HA support. The vMotion and HA capability greatly simplifies host maintenance by relocating where the VMs are hosted to avoid single blade failure induced VM outage.

When more memory resources are needed, you can add additional memory to the blade servers up to their supported 3TB limit. When additional CPU resources are needed, you can add additional B200 M5 Blade Servers to the Cisco UCS 5108 chassis of up to a maximum of eight blade servers for a single chassis.

If more CPU cores are needed, you can use the 40Gbps QSFP port available on the UCS 6324 Fabric Interconnect, with additional license requirement, to connect up to four C-series rack servers, for example, C220 M5, by using the 40G to 10G splitter cables, as shown in Figure 18. The same UCS Manager can still manage these additional C-series rack servers.

Figure 18) Compute scaling to include 8 x B200 M5 Blade Servers and 4 x C220 M5 rack servers.



If further expansion of the CPU resources is required, you can also add an additional UCS 5108 chassis to the solution. In the expanded two-chassis configuration, it is possible to support up to a total of 18 servers, including 16 blade servers and two rack servers, while keeping the AFF / FAS storage directly attached to the UCS 6324 Fabric Interconnects. The IOM 2204 ports on the second UCS 5108 chassis are connected into the UCS 6324 Fabric Interconnects' QSFP ports with the use of a 40G to 4 x 10G splitter cable and the required additional license.

Storage scaling

Depending on the use cases and requirements, this reference design can accommodate a variety of initial storage requirements and can also be scaled up when needs arise, up to the respective supported limits of the storage controller used.

For a use case that requires minimum storage, the AFF C190 is a great fit. It can be deployed with a configuration of 8 x 960GB SSDs to start. When the storage requirement grows, you can add additional 960GB SSDs. It can be scaled up to the 24x 960GB SSD limit.

Depending on the workload, the effective storage capacity can be much higher with the storage efficiency features. Furthermore, using FabricPool can be a good way to tier cold data into public clouds or on-premises object storage, to efficiently use the available flash storage for the frequently accessed data.

For a use case where higher storage capacity is required, either initially or as the solution grows, using the AFF A220 or FAS 2700 series gives you the flexibility of adding additional disk shelves to scale your storage as needs arise.

Table 6 compares the maximum supported storage devices, raw capacities, and FabricPool sizes for the AFF C190, AFF A220, and FAS 2700 series systems. See the [NetApp Hardware Universe](#) for additional information and details.

Table 6) AFF C190, AFF A220, and FAS 2700 series storage limits.

Model	AFF C190	AFF A220	FAS 2720	FAS 2750
Max Storage Devices (HA)	24 (drives)	144 (drives)	144 (drives)	144 (drives)
Max Raw Capacity (HA)	23.04 TB	4406.4 TB	2304 TB	1963.2 TB
Max FabricPool Size	16800 TiB	16800 TiB	8400 TiB	8400 TiB

Additional management software and tools

In addition to the base infrastructure support and management software, there are other management software and tools that can be used to manage the FlexPod solution and provide additional features and capabilities. The following highlights some of the optional software and tools.

[Cisco Intersight](#)

Cisco Intersight is a SaaS offering that simplifies and automates IT operations management to make daily activities easier and more efficient.

[Cisco Data Center Network Manager](#)

Cisco Data Center Network Manager provides automation and visibility for deploying, operating, and managing a network fabric with Cisco Nexus switches running NX-OS software.

[Cisco UCS Central](#)

Cisco UCS Central can be used to manage multiple Cisco UCS Mini and Cisco UCS domains in one or more physical locations.

[NetApp Active IQ](#)

NetApp Active IQ is a cloud service that provides proactive care and optimization of your NetApp environment, leading to reduced risk and higher availability. Active IQ leverages community wisdom and artificial intelligence for IT operations (AIOps) to provide proactive recommendations and risk identification.

[NetApp Cloud Insights](#)

NetApp Cloud Insights gives you visibility into your infrastructure and applications to monitor, troubleshoot, and optimize all your resources, including your public clouds and your private clouds.

[NetApp SnapCenter](#)

NetApp SnapCenter software leverages storage-based data management to provide an easy-to-use enterprise platform to securely coordinate and manage data protection across applications, databases, and file systems.

Solution verification

Cisco and NetApp designed and built FlexPod Express to serve as a premier infrastructure platform for their customers. Because it was designed by using industry-leading components, customers can trust FlexPod Express as their infrastructure foundation. In keeping with the fundamental principles of the FlexPod program, the FlexPod Express architecture was thoroughly tested by Cisco and NetApp data center architects and engineers.

VMware vSphere 7.0 hypervisor was verified on the FlexPod Express with UCS Mini and NetApp AFF / FAS infrastructure components. The solution uses configurations supported by NetApp IMT, VMware VCG, and Cisco HCL. In addition, a variety of tests were conducted, from normal operations to high availability testing, by introducing failure and maintenance related events, to exercise and validate iSCSI SAN boot, NFS I/O, Fabric Interconnects, switches, storage, and VMware operations. The FlexPod

Express architecture validation helps to instill confidence in our customers and to build trust in the design process.

Conclusion

FlexPod Express is designed for small to midsize businesses, ROBOs, and other businesses that require dedicated solutions. FlexPod Express is a simple and effective solution providing a validated design that uses industry-leading components. By providing various compute and storage scaling options as well as the possibility of using existing compliant switches that meet the requirements, FlexPod Express can be tailored for specific business needs and optimized for the desired performance, scale, and costs.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF / FAS NVA Deployment Guide
<https://www.netapp.com/pdf.html?item=/media/21938-nva-1154-DEPLOY.pdf>
- Cisco Nexus 3000 Series Switches
<https://www.cisco.com/c/en/us/products/switches/nexus-3000-series-switches/index.html>
- Cisco UCS Manager Configuration Guides
<https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-manager/products-installation-and-configuration-guides-list.html>
- Cisco UCS Mini
<https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-mini/index.html>
- Cisco Hardware and Software Compatibility list
<https://ucshcltool.cloudapps.cisco.com/public/>
- NetApp AFF and FAS System Documentation Center
<https://docs.netapp.com/platstor/index.jsp>
- NetApp Hardware Universe
<https://hwu.netapp.com>
- NetApp Product Documentation
<https://docs.netapp.com>
- FlexPod with FabricPool
<https://www.netapp.com/media/12427-tr4801.pdf>
- NetApp Interoperability Matrix Tool (IMT)
<http://mysupport.netapp.com/matrix>
- NetApp Support Site
<https://mysupport.netapp.com>
- VMware Compatibility Guide
<https://www.vmware.com/resources/compatibility/search.php>

Version history

Version	Date	Document version history
Version 1.0	January 2021	Initial release.

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 © 2020 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

NVA-1154-DESIGN-0121