



NetApp Verified Architecture

FlexPod Datacenter with Apprenda for PaaS

Including Red Hat OpenStack 8, Docker Containers, NetApp Jenkins Plugin, and ONTAP 9

NVA Design

Converged Infrastructure Engineering, NetApp
April 2017 | NVA-1113-DESIGN | Version 1.0

Reviewed by



In partnership with



TABLE OF CONTENTS

1	Executive Summary	4
2	Program Summary	4
2.1	FlexPod Program Benefits	5
3	Solution Overview	7
3.1	Target Audience	7
3.2	Solution Technology	7
3.3	Use Case Summary	17
4	Technology Components	17
4.1	Hardware Components	17
4.2	Software Components	17
5	Solution Design	18
5.1	Cisco UCS Design	18
5.2	Cisco Nexus Network Design	21
5.3	NetApp AFF Storage Design	23
5.4	OpenStack Design	26
5.5	Apprenda Design	27
6	Solution Verification	30
6.1	Continuous Integration Pipeline	30
6.2	Continuous Delivery	32
	Conclusion	34
	Acknowledgements	34
	References	35
	Cisco UCS	35
	Cisco Nexus Networking	35
	NetApp FAS Storage	35
	Red Hat OpenStack Platform	36
	Apprenda PaaS	36
	Interoperability Matrixes	36

LIST OF TABLES

Table 1) Hardware components. 17
Table 2) Solution software components. 18

LIST OF FIGURES

Figure 1) FlexPod component families. 5
Figure 2) Cisco UCS components. 8
Figure 3) ONTAP. 12
Figure 4) Undercloud vNIC and network segmentation. 19
Figure 5) Overcloud vNIC and network segmentation. 20
Figure 6) Network design. 21
Figure 7) OpenStack network design. 26
Figure 8) Apprenda, Jenkins and OpenStack logical architecture. 29
Figure 9) Automating CI workflow with Apprenda Jenkins on NetApp using Docker. 30
Figure 10) NetApp integrations for the CI workflow. 31
Figure 11) Automating CD workflow with Apprenda and NetApp. 33
Figure 12) NetApp integrations for CD. 33

1 Executive Summary

In this digital world, developing web and mobile applications is part of every business that wants to go to market much more quickly with its products, services, and solutions. Doing so requires an agile development method that involves a continuous integration (CI) and continuous deployment (CD) process in a fast and iterative manner. The shift toward agile development has forced business owners to be more exploratory and innovative, emphasizing speed in the application development workflows. The new types of applications are designed to provide one or many services that are high in value and customer satisfaction.

Because of the shift in development and deployment processes, IT organizations must also shift their infrastructure to a more flexible cloud-ready platform. Although developers are simply focused on developing applications, the IT organization is responsible for keeping them working and protecting the work they produce, as well as supporting the applications after they are deployed in production. Self-service automation for end-user developers and flexibility, scalability, and reliability are primary requirements for any modern data center architecture.

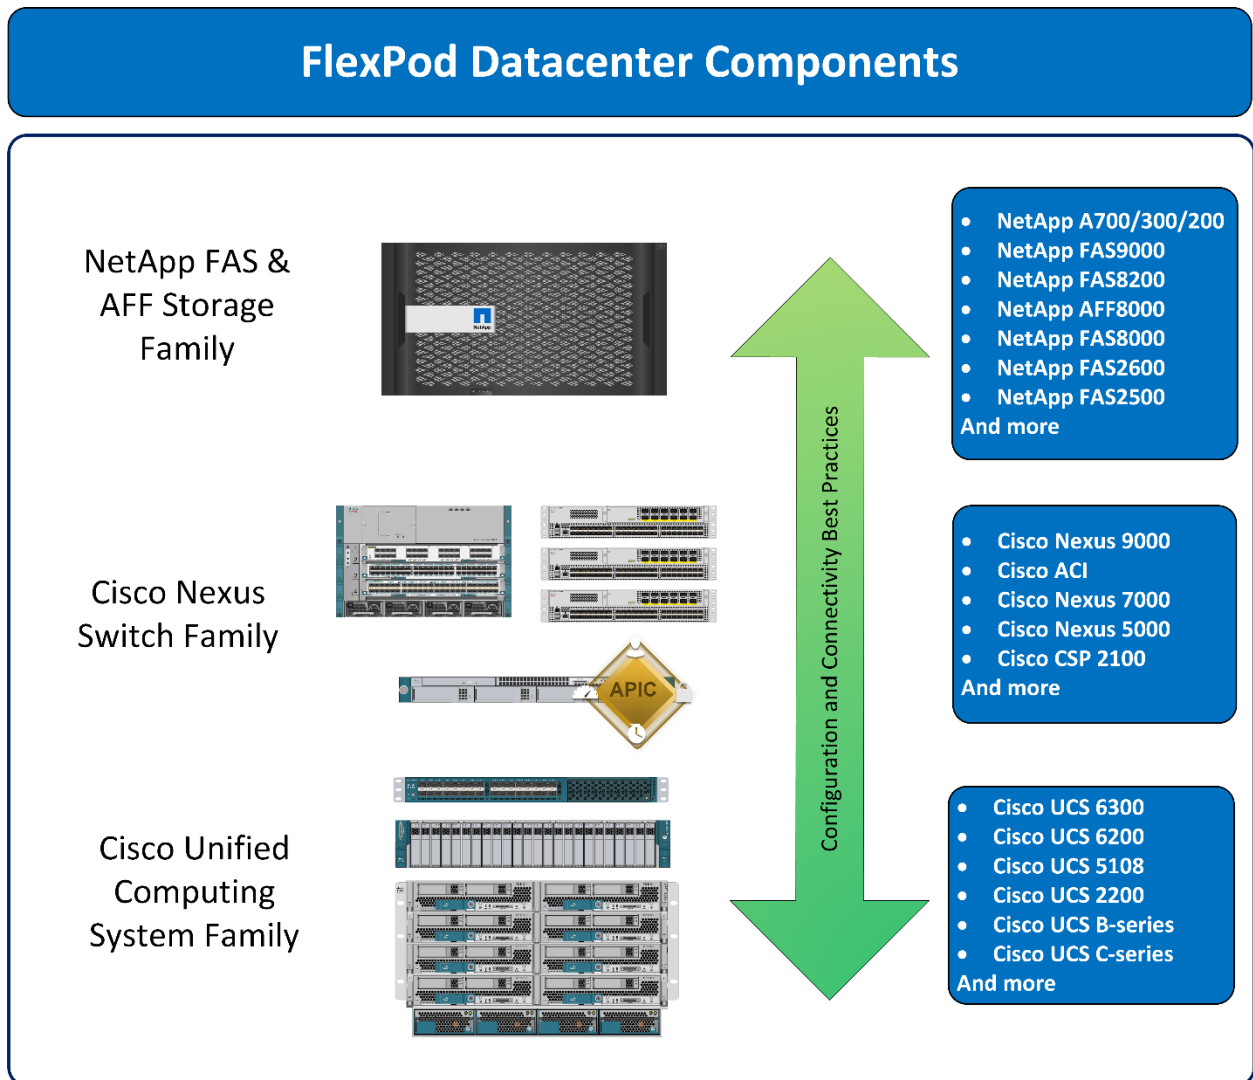
The FlexPod® with Apprenda platform-as-a-service (PaaS) solution is a turnkey hardware and software platform that provides developers with highly efficient self-service workflows to manage application development from project start through production deployment. Leveraging native API integrations throughout the stack, this solution accelerates developer productivity, improves QA testing cycles, and streamlines application deployment using the FlexPod converged infrastructure. Customers can implement this solution on existing FlexPod hardware or configure a new FlexPod infrastructure sized to meet their current needs with the ability to scale as needed.

FlexPod converged infrastructures utilize Cisco UCS Servers, Cisco Nexus switches, and NetApp® All Flash FAS storage systems to deliver the most flexible, scalable, and reliable infrastructure on the market. FlexPod allows organizations to build an infrastructure that supports almost any operating system or application with enterprise-class performance and reliability. FlexPod also supports nondisruptive scaling of any resource and dynamic reallocation of resources, providing organizations the agility to react to changing business requirements with minimal impact. This eliminates application silos and streamlines IT operations by creating a standardized infrastructure that supports API-based administration for innovative private cloud deployments.

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/FAS 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 created 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.

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

This FlexPod Datacenter solution is intended to provide customers with a fully integrated application development environment that leverages all the capabilities of NetApp storage systems to accelerate application development and deployment. By leveraging native ONTAP® 9 technologies such as Snapshot™ copies and FlexClone® volumes, Cloudbees Jenkins Enterprise APIs, and Docker containers, this solution enables organizations to improve developer productivity, reduce build times, and optimize storage space.

The FlexPod architecture is designed to help customers with proven guidance and measurable value. By introducing standardization, FlexPod helps 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 customer IT demands.

This document describes Apprenda PaaS using the Red Hat OpenStack Platform, Docker containers, Jenkins Cloudbees, and NetApp FAS with ONTAP 9 built on the FlexPod model from Cisco and NetApp. These design considerations and recommendations are not just limited to the specific releases of the components described in this document, but are also applicable for 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 FlexPod features to accelerate and streamline their application development and deployment methodology.

3.2 Solution Technology

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

- Cisco Unified Computing System
- 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 APIs.

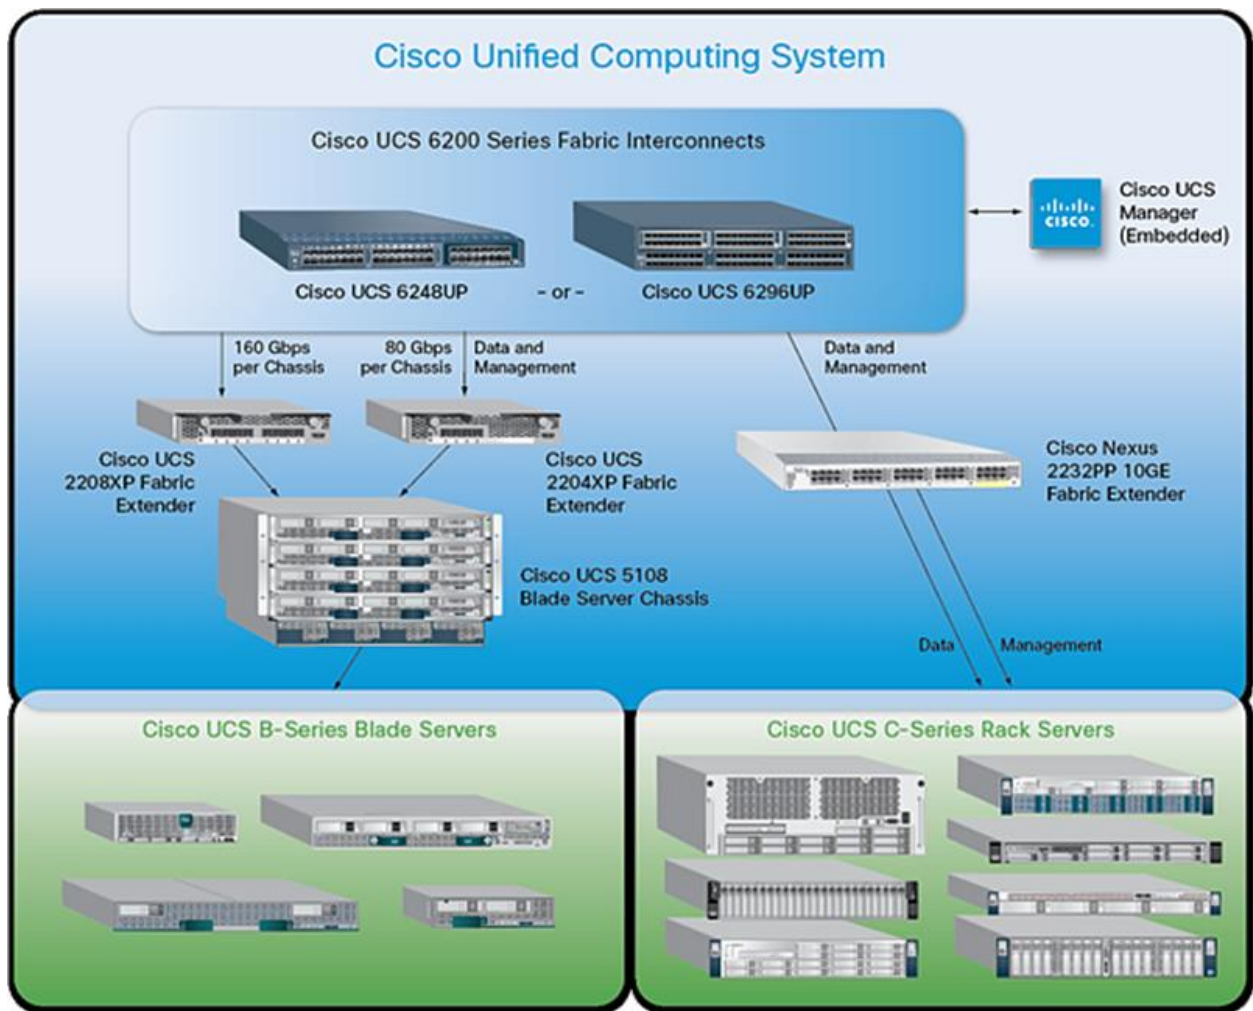
This solution builds on the FlexPod infrastructure by adding the following software components:

- Red Hat OpenStack Platform 8 (Liberty)
- Docker containers
- Jenkins CloudBees for CI/CD pipeline management
- Apprenda platform

Cisco UCS

The Cisco UCS is a next-generation data center platform that unites computing, networking, storage access, and virtualization resources into a cohesive system designed to reduce the total cost of ownership and increase business agility. The system integrates a low-latency, lossless 10 Gigabit Ethernet (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.

Figure 2) Cisco UCS components.



The main components of the Cisco UCS are as follows:

- **Compute.** The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers based on Intel Xeon 2600 v2 series processors.

- **Network.** The system is integrated onto a low-latency, lossless, 10Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks that are typically configured as separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables and by reducing power and cooling requirements.
- **Virtualization.** This system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access.** The system provides consolidated access to both SAN storage and NAS over the unified fabric. By unifying storage access, the Cisco UCS can access storage over Ethernet (SMB 3.0 or iSCSI), FC, and FCoE. This provides customers with storage choices and investment protection. In addition, server administrators can preassign storage access policies to storage resources for simplified storage connectivity and management, which lead to increased productivity.
- **Management.** The system integrates all system components so that the entire solution can be managed as a single entity by Cisco UCS Manager. Cisco UCS Manager has an intuitive GUI, a CLI, and a powerful scripting library module for Microsoft PowerShell built on a robust API. These different methods can manage all system configuration and operations.

Cisco UCS fuses access layer networking and servers. This high-performance, innovative server system provides a data center with a high degree of workload agility and scalability.

Cisco UCS 6248UP Fabric Interconnects

The fabric interconnects provide a single point of connectivity and management for the entire system. Typically deployed as an active-active pair, the system's fabric interconnects integrate all components into a single, highly available management domain controlled by Cisco UCS Manager. The fabric interconnects manage all I/O efficiently and securely at a single point, resulting in deterministic I/O latency independent of the topological location of a server or VM in the system.

Cisco UCS 6200 Series fabric interconnects support the system's 10Gbps unified fabric with low-latency, lossless, cut-through switching that supports IP, storage, and management traffic with a single set of cables. The fabric interconnects feature virtual interfaces that terminate both physical and virtual connections equivalently, establishing a virtualization-aware environment in which blades, rack servers, and VMs are interconnected by the same mechanisms. The Cisco UCS 6248UP is a 1RU fabric interconnect that features up to 48 universal ports that can support 10GbE, FCoE, or native FC connectivity.

Cisco UCS 5108 Blade Server Chassis

The Cisco UCS 5100 Series blade server chassis is a crucial building block of the Cisco UCS, delivering a scalable and flexible chassis. The Cisco UCS 5108 blade server chassis is 6RU high and can mount in an industry-standard, 19-inch rack. A single chassis can house up to eight half-width Cisco UCS B-Series blade servers and can accommodate both half-width and full-width blade form factors.

Four single-phase, hot-swappable power supplies are accessible from the front of the chassis. These power supplies are 92% efficient and can be configured to support nonredundant, N + 1 redundant configurations and grid-redundant configurations. The rear of the chassis contains eight hot-swappable fans, four power connectors (one per power supply), and two I/O bays for Cisco UCS 2200 XP fabric extenders. A passive midplane provides up to 40Gbps of I/O bandwidth per server slot and up to 80Gbps of I/O bandwidth for two slots.

Cisco UCS 2204XP Fabric Extenders

The Cisco UCS 2204XP has four 10GbE, FCoE-capable, enhanced small form-factor pluggable (SFP+) ports that connect the blade chassis to the fabric interconnect. Each Cisco UCS 2204XP has 16 10GbE

ports connected through the midplane to the half-width slot in the chassis. When configured in pairs for redundancy, two 2204XP fabric extenders provide up to 80Gbps to the chassis.

Cisco UCS B200 M4 Blade Servers

The enterprise-class Cisco UCS B200 M4 blade server extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M4 is powered by the latest Intel Xeon E5-2600 v4 series processor family CPUs. This server contains up to 1536GB of RAM (using 64GB DIMMs), two solid-state drives (SSDs) or hard disk drives (HDDs), and up to 80Gbps throughput connectivity. The Cisco UCS B200 M4 blade server mounts in a Cisco UCS 5100 Series blade server chassis or a Cisco UCS Mini blade server chassis. It supports one connector for a Cisco Virtual Interface Card (VIC) 1340 or VIC 1240 adapter, which provides Ethernet and FCoE.

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 this software to manage the entire Cisco UCS as a single logical entity through an intuitive GUI, a CLI, or an XML API.

The Cisco UCS Manager resides on a pair of Cisco UCS 6200 Series fabric interconnects in a clustered, active-standby configuration for high availability. The software provides administrators with 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.

Key elements managed by Cisco UCS Manager include the following:

- Cisco UCS Integrated Management Controller (IMC) firmware
- RAID controller firmware and settings
- BIOS firmware and settings, including server universal user ID (UUID) and boot order
- Converged network adapter firmware and settings, including MAC addresses, worldwide names (WWNs), and SAN boot settings
- Virtual port groups used by VMs, with Cisco Data Center VM-FEX technology
- Interconnect configuration, including uplink and downlink definitions, MAC address and WWN pinning, virtual local area networks (VLANs), virtual storage area networks, quality of service (QoS), bandwidth allocations, Cisco Data Center VM-FEX settings, and EtherChannels to upstream LAN switches

For more information, see 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 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 (HPC) environments. Cisco Nexus 9000 switches provide 40Gb switching capability and can participate in Cisco Application Centric Infrastructure (ACI), but do not support FC or FCoE storage

protocols. To support these protocols, FlexPod supports FC/FCoE connections directly between the Cisco UCS fabric interconnects and the NetApp AFF storage system.

Cisco Nexus 9396PX is the switch used in the validation of this FlexPod architecture. 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 FAS/AFF Storage Systems

A product of more than 20 years of innovation, ONTAP has evolved to meet the changing needs of customers and help drive their success. ONTAP provides a rich set of data management features and clustering for scale-out, operational efficiency, and nondisruptive operations. This storage software 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 (ITaaS), 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 NetApp 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 FAS8000 series unified scale-out storage systems. Powered by ONTAP, the FAS series unifies SAN and NAS storage infrastructures. The FAS8000 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 FAS8000 series comes with integrated UTA2 ports that support 16Gb FC, 10GbE, or FCoE.

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.

ONTAP 9

With ONTAP 9, NetApp provides the next generation of enterprise-ready, unified scale-out storage. Developed from a solid foundation of proven ONTAP technology and innovation, ONTAP 9 is the basis for large virtualized shared storage infrastructures, which 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 dictate 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, but ultimately the controller becomes completely populated, with no further expansion possible. At this stage, the only option is to acquire another controller.

Scale-Out

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

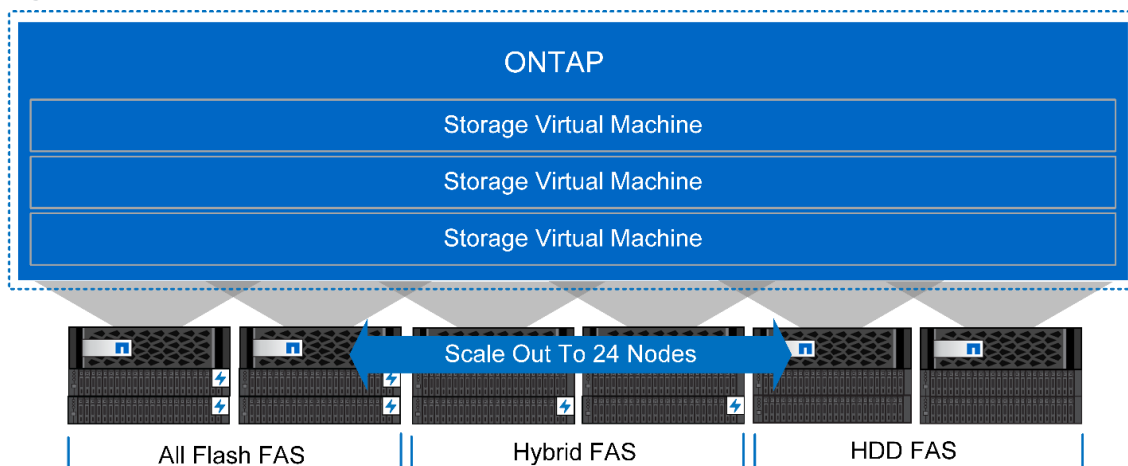
Although scale-out products have been available for some time, these products were typically subject to one or more of the following shortcomings:

- 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, and compression)
- Little or no data replication capability

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 is depicted in Figure 3, 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 3) 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. 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 customers 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, and replication capabilities of the NetApp AFF solution.

Storage Efficiencies

The NetApp ONTAP solution includes built-in thin provisioning, in-line and postprocess data deduplication, in-line and postprocess data compression, inline data compaction, 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 an 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 and later, in-line 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.
- **Data compaction.** Data compaction enables even greater space savings on certain types of data. When data writes from the host are smaller than the default 4k block size on the storage system, ONTAP can compact multiple sub-4k writes into a single physical block on disk, producing significant capacity savings.
- **FlexClone.** Because NetApp Snapshot copies are created at the FlexVol® volume level, they cannot be directly leveraged within an OpenStack context. This is because a Cinder user requests a Snapshot copy be taken of a Cinder volume (not the containing FlexVol volume). Because a Cinder volume is represented as either a file on NFS or an iSCSI LUN, the way that Cinder snapshots are created is through the use of NetApp FlexClone technology. By leveraging the FlexClone technology to provide Cinder snapshots, it is possible to create thousands of Cinder snapshots for a single Cinder volume.

Advanced Storage Features

NetApp ONTAP provides several additional features that can be leveraged in a private cloud environment. 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.
- **LIF.** 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 machine.** 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 9 can support several protocols concurrently on the same storage system. Unified storage is important in private cloud deployments, where servers may utilize SAN protocols for boot volumes or high-performance data LUNs and NAS protocols for shared file system access, all within the same application stack. The following protocols are supported on NetApp ONTAP systems:

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

OpenStack

IT organizations are rapidly adopting a cloud services model for all IT services. OpenStack is an open-source virtualized infrastructure and management framework that offers a variety of compute, storage, and networking services with a common API layer that helps customers to deliver self-service capabilities to end users and streamlines management and operations for IT staff. Red Hat OpenStack Platform is an industry-leading distribution of OpenStack that delivers the power and flexibility of OpenStack along with enterprise-class support.

FlexPod is an ideal physical platform for cloud infrastructures, offering unmatched flexibility, scalability, and nondisruptive operations. Features such as rapid cloning and dynamic account creation allow administrators to manage storage resources effectively while allowing end users the flexibility to provision and manage their own environments. Cisco UCS servers have high-density CPU and memory options and policy-based management, which enable highly efficient provisioning and management of compute resources. Cisco Nexus switches include high-bandwidth interconnect capabilities, and both Cisco UCS and Cisco Nexus have certified OpenStack Neutron drivers to enable dynamic network configuration from the OpenStack interface.

NetApp and Cisco have partnered with Red Hat to integrate the installation of all necessary drivers and services into Red Hat OpenStack Platform, enabling streamlined deployment and lifecycle management using OpenStack Director. After the deployment is complete, FlexPod with Red Hat OpenStack Platform delivers a highly flexible and automated infrastructure.

The OpenStack framework is primarily used to deliver infrastructure as a service (IaaS). This reference architecture showcases the following features and capabilities that enable IaaS for production as well as development and test environments:

- Block storage services using Cinder over NFS
- NetApp All Flash FAS for Cinder and Glance storage
- Cinder storage services catalog using extra specs and volume types
- Thin provisioning
- Rapid cloning
- Seamless scalability
- High availability (HA) deployment of OpenStack services using Pacemaker as per best practices
- Nova
- Neutron
- Glance
- Cinder
- Ceilometer
- Horizon UI

Appendix

Appendix is a major enterprise PaaS vendor that empowers agile application development in private and hybrid cloud setups. Appendix provides runtime environments for Java, .NET, and Docker-based applications and removes many of the complexities associated with application development, deployment, hosting, and support. Appendix provides deep support for hosting new-generation, cloud-native applications, but at the same time it offers the ability to onboard legacy applications that were not designed to run in distributed environments. This capability helps IT organizations to realize significant operational savings and enables them to innovate more quickly. Appendix automatically reconfigures applications to use its own distributed components, which cuts down on the development time and improves time to market. It simplifies the usually lengthy deployment process through policy-driven automation. Finally, by employing container tech with automation, it helps organizations to trim operational costs: save on licenses, streamline provisioning, and use compute and storage resources much more efficiently.

Appendix has its share of advantages over other PaaS vendors with its ability to integrate with a wide variety of third-party development tools such as Jenkins. The following are some of the benefits that Appendix provides to developer CI/CD pipelines:

- Appendix provides a standard platform for existing and new applications. It can run on any infrastructure: physical, virtual, or hybrid. It supports a wide array of OSs/VMs, databases, and application server images and integrates with a variety of third-party tools.
- Appendix provides a single platform for private and public cloud deployments in a hybrid mode. It combines all the infrastructure resources into a single policy-driven resource pool for development teams to consume in a controlled, secure, compliant, and self-service fashion.
- Appendix integrates with Cloudbees Jenkins Enterprise, forming a stable CI pipeline. Appendix runs Jenkins master node in a managed Docker container with the following benefits:
 - The Appendix platform runs the Jenkins master in a Docker container that abstracts it from the underlying physical or virtual server in the same manner as it does with its applications.

- In an event of any maintenance, the Jenkins master can be moved from a hosted VM or physical host to another one on an Apprenda platform without disrupting the CI pipeline. The seamless relocation of Jenkins master node across hosts is enabled through persistent Docker volumes created with NetApp's Docker Volume Plug-In for ONTAP. Persistent storage for Jenkins Docker containers is provided through a shared FlexVol volume.
- Multiple Jenkins masters can scale on Apprenda's platform with lower demands on the supporting infrastructure.
- Managing multiple Jenkins masters on Apprenda is much easier and less complicated, and ONTAP provides persistent storage for the Docker volume to scale the number of Jenkins containers.

3.3 Use Case Summary

This solution addresses the following use cases:

- Accelerating the CI cycle through NetApp integration with OpenStack, Apprenda, and Cloudbees Jenkins Enterprise. This allows for more frequent testing that results in better code.
- Accelerating CD workflow through NetApp and Apprenda integration, allowing faster QA/staging phases before applications are deployed in production.
- Scalability, agility, and resiliency, allowing developers self-service access to the resources they need with enterprise reliability.
- Turnkey infrastructure stack for cloud-native application development and deployment environments.

4 Technology Components

This section covers the technology components of the FlexPod Datacenter with Apprenda PaaS solution.

4.1 Hardware Components

During solution testing, Cisco UCS blade servers were used to host all components in the system.

Table 1 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 might vary based on customer requirements.

Table 1) Hardware components.

Hardware	Configuration
Cisco UCS 6200 Series fabric interconnects	FI 6248UP
Cisco UCS B200 blades	B200 M3 using Cisco UCS VIC 1240
Cisco UCS 5108 chassis	Cisco UCS-IOM 2208XP
Cisco Nexus 9000	Cisco Nexus 9396PX
NetApp FAS/AFF storage system	FAS8040
NetApp DS2246 disk shelves	Disk shelves populated with 900GB SAS drives

4.2 Software Components

Table 2 lists the software components used to implement the solution. The software components used in any particular implementation of the solution might vary based on customer requirements.

Table 2) Solution software components.

Software/Firmware	Version
Compute	
Cisco UCS Manager	3.1(1e)
Networking	
Cisco NX-OS	7.0(3)I2(2a)
Storage	
ONTAP	9
NetApp ONTAP licenses	Cluster base, NFS, iSCSi, FlexClone, CIFS
NetApp Unified Cinder driver	OpenStack Liberty
NetApp Docker Volume Plug-In	1.3
NetApp Jenkins Plugin	1.0
Hypervisor/Operating System software	
Red Hat Enterprise Linux	7.2
Red Hat OpenStack Platform	8.0
Docker engine	1.3
Apprenda	6.5.0
Jenkins	1.7

5 Solution Design

This solution is based on the FlexPod Datacenter for Red Hat OpenStack reference architecture. Red Hat OpenStack Platform 8 manages the infrastructure provided to the Apprenda PaaS software. Additional design and deployment details for the base OpenStack infrastructure are available in [TR-4506: Red Hat OpenStack Platform 8 on FlexPod](#).

The FlexPod Datacenter for Apprenda PaaS solution consists of the following designs:

- Cisco UCS
- Cisco Nexus network switches
- NetApp AFF storage
- OpenStack design
- Apprenda PaaS architecture

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 its greater rack density and easier scalability. If you want to implement a C-Series 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.

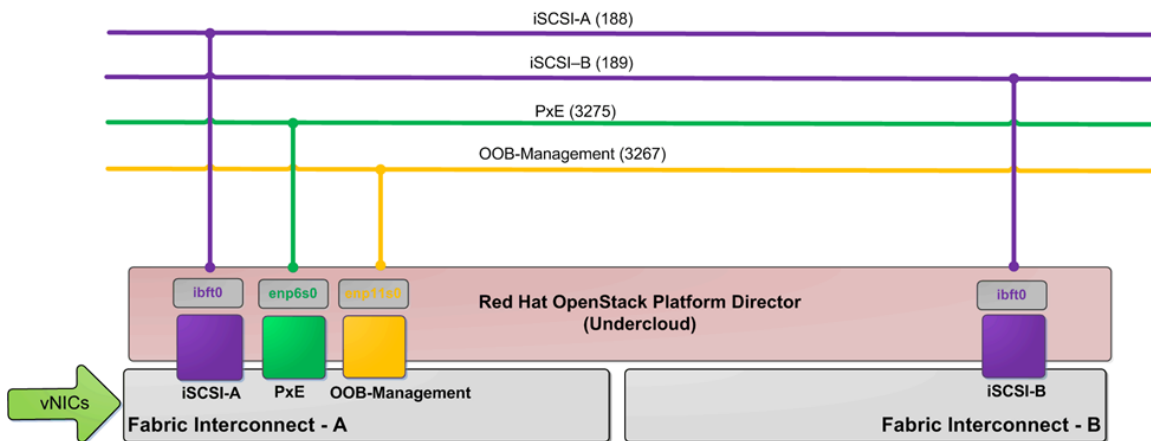
Cisco UCS: B-Series Server Design

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 RHEL 7.2 nodes. These nodes were configured into an OpenStack cluster using Red Hat OpenStack Platform 8 (Liberty).

The Cisco VIC presents five virtual PCIe devices to the RHEL node. Two virtual 10GbE NICs (vNICs) are used for iSCSI boot; one vNIC is used for the PXE boot network used during OSP deployment; and the remaining two vNICs are bonded together to support management, API services, NFS traffic, and tenant VLAN tunnels. 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.

- For the OpenStack undercloud, the service profile template has an iSCSI-A, iSCSI-B, a PXE (to listen and respond to requests from overcloud servers being provisioned), and an OOB-Management vNIC. Figure 4 illustrates the undercloud in the NetApp lab.

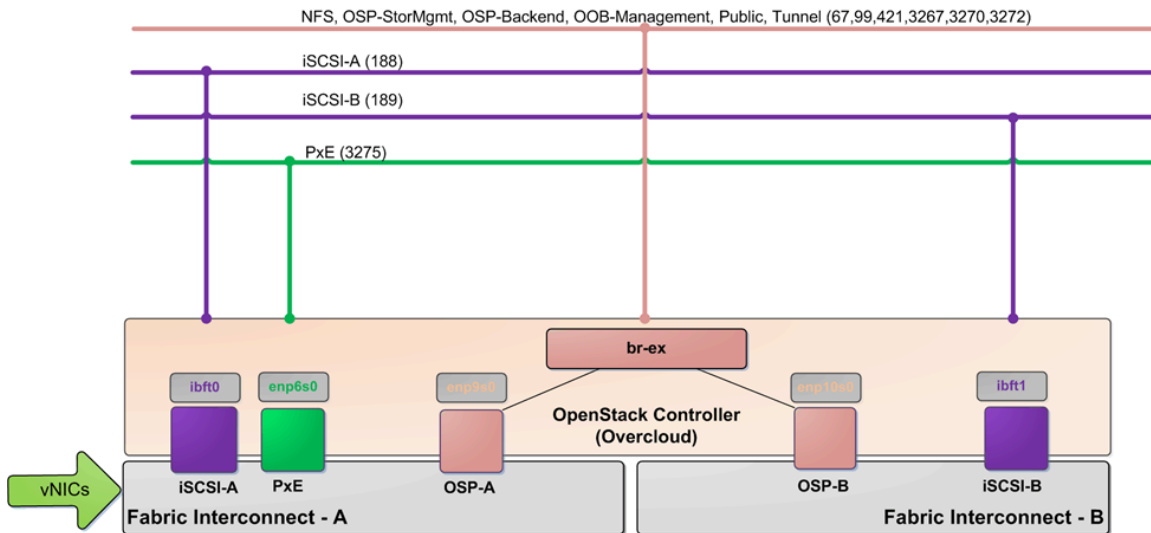
Figure 4) Undercloud vNIC and network segmentation.



- For the OpenStack overcloud, the service profile template has an iSCSI-A, iSCSI-B, PXE, and an OSP-A and OSP-B vNIC. The OSP-A and OSP-B vNICs have the following VLANs carried to them: NFS, OSP-StorMgmt, OSP-Backend, OOB-Management, Tunnel, and Public. These two vNICs are bonded together in a link aggregation using mode 1, or active backup. The bridge is named `br-ex`.

Figure 5 illustrates the overcloud in the NetApp lab.

Figure 5) Overcloud vNIC and network segmentation.



Cisco UCS Bandwidth Considerations

FlexPod allows organizations to adjust the individual components of the system to meet their 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 (IOMs), 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:

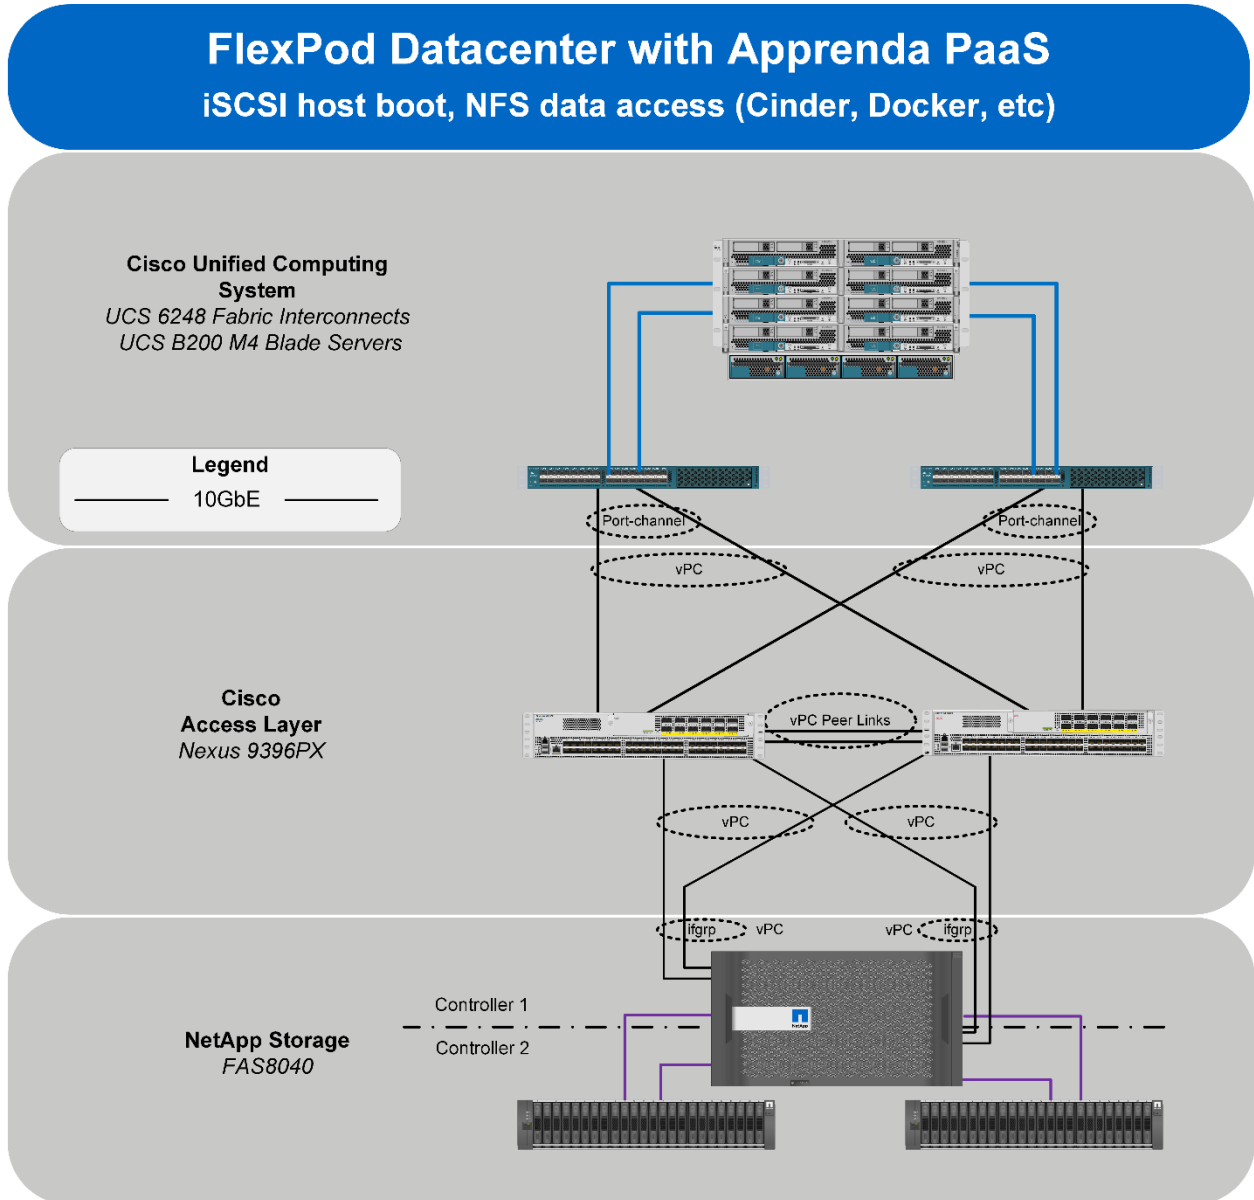
- The fabric extender model (2204XP or 2208XP)
- The modular LAN on motherboard (mLOM) card
- The mezzanine slot card

In this validated configuration, the Cisco UCS 5108 chassis is populated with Cisco UCS 2208XP IOMs, and each blade has a Cisco UCS VIC 1340 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 to each Cisco UCS 2208XP IOM, for a total of 80Gb bandwidth to each blade.

5.2 Cisco Nexus Network Design

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

Figure 6) Network design.



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 they support 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 the applicability to the widest range of customer environments. Therefore, ACI was not considered to be a requirement, but this architecture can be integrated into a new or existing ACI topology if desired.

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. From an Ethernet perspective, the Cisco Nexus 9000 uses Virtual Port Channel (vPC), which 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 is managed by using the Cisco Modular QoS CLI, providing class-based traffic control. Note 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 auto-recovery 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](#).

Jumbo Frames

A balanced and predictable fabric is critical within any data center environment. As designed, the FlexPod architecture accommodates myriad traffic types (iSCSI, NFS, mgmt 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 9000 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 the inconsistent MTUs can introduce.

5.3 NetApp AFF Storage Design

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

Validated Storage Configuration

For this validation effort, a NetApp FAS8040 storage system was used to provide all necessary storage. All OpenStack nodes boot from LUNs on the FAS8040 using the iSCSI protocol. Storage for OpenStack Cinder block storage services is provided by NFS volumes on the FAS8040 using the NetApp Integrated Driver for OpenStack. Docker volume storage is also provided as NFS exports on the FAS8040 and is managed using the NetApp Docker Volume Plug-In.

Storage configuration of the storage array follows typical NetApp best practices. On each controller, a single data aggregate and volumes were created, for the boot LUNs, Cinder storage volumes, Glance storage volumes, persistent Docker volumes, and other storage as needed.

Storage Networking

Each of the two NetApp FAS8040 storage controllers has a two-port LACP ifgrp (port channel) connected to a vPC across the two Cisco Nexus 9396PX switches, with all storage-related VLANs allowed on these trunks. iSCSI LIFs are created on each node for each VLAN (four LIFs total), and ALUA is used to provide multipathing and load balancing of the iSCSI links. Initiator groups (igroups) were configured on the FAS8040 systems to map boot LUNs to the Linux host servers.

Logical interfaces are also created on each node for the NFS VLAN, allowing hosts to access volumes for Cinder and Glance over the same vPC uplinks.

NetApp Unified Driver for ONTAP with NFS

A Cinder driver is an implementation of a Cinder back end that maps the abstract APIs and primitives of Cinder to appropriate constructs within the storage solution underpinning the Cinder back end.

The NetApp unified driver for ONTAP with NFS is a driver interface from OpenStack block storage to an ONTAP storage system. This software provisions and manages OpenStack volumes on NFS exports provided by the storage system. The NetApp unified driver for ONTAP does not require additional management software to achieve the desired functionality because it uses NetApp APIs to interact with the ONTAP cluster. It also does not require additional configuration beyond overriding default templates during OpenStack deployment using the Red Hat OpenStack Platform director.

All of the resulting Cisco UCS blades that are provisioned with RHEL 7.2 mount the (appropriately designated) NetApp FlexVol volumes on the FAS8040.

Note: A FlexVol volume is a logical container for one or more Cinder volumes.

The NetApp OpenStack contribution strategy adds new capabilities directly into upstream OpenStack repositories, so that all features are available in Red Hat OpenStack Platform 8.0 out of the box. For more information regarding the NetApp unified driver (including other protocols available), see [NetApp Data ONTAP Drivers for OpenStack Block Storage \(Cinder\)](#).

Copy Offload Tool

The NetApp copy offload feature was added in the Icehouse release of OpenStack, which enables images to be efficiently copied to a destination Cinder volume that is backed by an ONTAP FlexVol volume. When Cinder and Glance are configured to use the NetApp NFS copy offload client, a controller-side copy is attempted before reverting to downloading the image from the image service. This improves image-provisioning times while reducing the consumption of bandwidth and CPU cycles on the hosts running the Glance and Cinder services. This is due to the copy operation being performed completely within the storage cluster.

Although the copy offload tool can be configured automatically as a part of an OpenStack deployment using Heat orchestration templates, it must still be [downloaded](#) from the NetApp Utility ToolChest site by the customer.

If Cinder and Glance share the same NetApp FlexVol volume, the copy offload tool is not necessary. Rather, a direct API call to the NetApp storage system is utilized through the NetApp unified driver that facilitates a controller-side copy relative to a network copy.

For more information about this functionality, including a process flowchart, see [OpenStack Deployment and Operations Guide - Version 6.0, Theory of Operation and Deployment Choices, Glance and Clustered Data ONTAP](#).

NetApp Docker Volume Plug-In

nDVP provides direct integration of NetApp storage systems with the Docker container ecosystem. Using standard Docker tools and commands, end users can dynamically provision persistent storage for applications or projects without interacting with the underlying infrastructure. Users can also instantly clone volumes for testing or rapid deployment purposes. nDVP supports both NFS exports and iSCSI LUNs when using ONTAP storage systems and iSCSI LUNs when using SolidFire® and/or E-Series storage systems. nDVP supports multiple storage system back ends simultaneously; it also supports a storage service catalog by creating an nDVP instance for each tier of storage as defined by administrators. For more information about nDVP, see [NetApp Docker Volume Plug-In Best Practices and Recommendations](#).

Docker integration for this solution was validated using volumes from the FAS8040 storage system access by the hosts through NFS v3.

NetApp Jenkins Plugin

NetApp also jointly worked with CloudBees to develop a plug-in that automates and reduces the continuous build and test cycle time, instantaneously creates developer workspace, improves storage space efficiency, and reduces storage costs. Developers can instantly create dedicated workspaces prepopulated with current code or environmental dependencies. They can clone or create Snapshot copies of workspaces to provide branch environments for testing or further development and can instantly roll back to previous clones or Snapshot copies as needed. These functions are directly integrated into the Cloudbees Jenkins Enterprise environment, enabling developers to accelerate the CI/CD workflow without depending on infrastructure administrators. For more information about NetApp CI/CD integration with the Jenkins plug-in, see [TR-4547: Continuous Integration \(CI\) Pipeline with CloudBees Enterprise Jenkins and ONTAP 9](#).

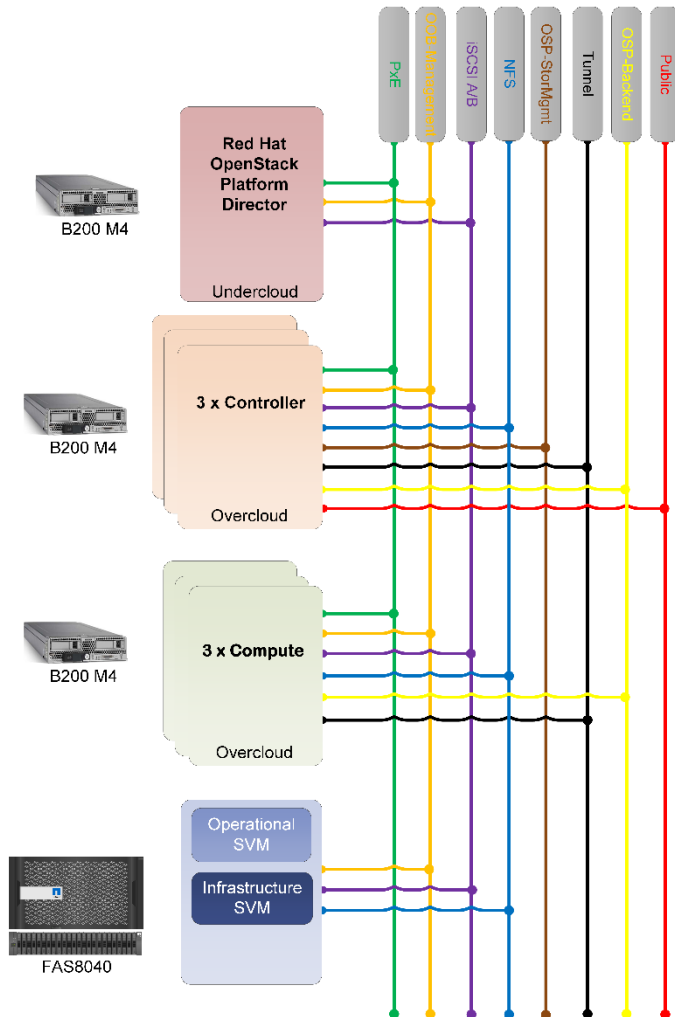
5.4 OpenStack Design

This solution is based on the OpenStack deployment described in [TR-4506: Red Hat OpenStack Platform 8 on FlexPod](#). For the purposes of this validation, one Cisco UCS B200 blade was used for undercloud, three blades for the overcloud controllers, and three blades for the Nova compute nodes.

This infrastructure was deployed using the Red Hat OpenStack Director to install the services required for deployment on the undercloud node, discover the overcloud controller and compute nodes, and deploy the full OpenStack overcloud on those nodes. Heat orchestration templates were used to customize the environments as needed for the NetApp Unified Cinder driver and networking configurations. Specific deployment details, including example Heat templates, are available in TR-4506.

Figure 7 shows the logical network configuration used in this deployment.

Figure 7) OpenStack network design.



5.5 Apprenda Design

The Apprenda platform consists of some of the key components in architecture, which makes it more resilient and enables out-of-the-box integration. Some of the components of the Apprenda architecture include:

- **Load managers.** On the Apprenda platform, the load manager (best described as a reverse proxy) serves as the initial receptor of incoming HTTP requests. IIS configuration on the load manager is managed by the Apprenda load manager service, a Windows service that creates and modifies URL rewrite rules as various guest app components are deployed to internal front-end content servers. The load manager service makes sure of the appropriate routing of inbound requests based on the request URL patterns.

For high-availability purposes, we recommend a topology consisting of two load manager servers, which requires the use of a shared IIS configuration housed in a network share. To optimize performance, a hardware load balancer should be placed in front of these load managers.

- **Domain naming services (DNS) server.** For the entire platform-based operation, the different servers communicate with each other using DNS-resolvable host names and therefore require an internal DNS structure to be in place that resolves IP assignments. It is recommended to have a dedicated DNS server for the Apprenda platform servers. An existing Active Directory controller or DNS server can be used for this purpose.
- **MS SQL Server nodes.** The platform manages SQL Server instances on your behalf to provision and configure guest application databases. Any number of SQL Server instances can be managed by a single platform, and SQL Server instances can be added to the platform at any time for capacity. Our reference architecture includes a single SQL Server instance that is configured as a SQL Server failover cluster; such a configuration typically relies on shared storage on a SAN or NAS. We've included two SQL Server nodes in the cluster as a standard recommendation for simple redundancy. The platform manages this cluster as it would a normal SQL Server instance. Expansion of the database tier of our reference architecture comes in two forms:
 - Adding SQL Server nodes to existing clusters increases redundancy capabilities of existing platform-managed SQL Server instances.
 - Adding SQL Server nodes that are independent of existing clusters or adding entirely new platform-managed clusters increases the capacity of the database hosting layer of your Apprenda platform.

In addition to providing storage for guest applications, an Apprenda-managed SQL Server instance is necessary for housing the Apprenda core database, which contains data necessary for platform functionality.

- **Platform repository.** The platform repository is a network share that serves as the central storage location for all platform and guest application binaries. This location must be configured prior to installation, should be specified as a network path in the Apprenda installer, and must contain the following three folders:
 - Applications
 - Apprenda
 - SAC

Configuring each folder as a separate network share is recommended as long as they are accessible through the same base path.

All guest application binaries, after being uploaded to the platform by developers, are stored in the Applications folder in the platform repository (in some parts of the platform, such as the repository browser in the system operations center [SOC], this folder is called the application repository).

The Applications folder also includes binaries for platform components that are themselves hosted on the platform as guest applications, such as the developer portal and the SOC. All other platform binaries are stored in the Apprenda folder (which is sometimes referred to as the system repository).

Upon workload deployment, binaries that are needed for local execution are copied to the target server from their respective locations in the platform repository.

It is recommended to have the platform repository on an NFS share from the NetApp FAS/AFF storage system. ONTAP provides thin provisioning, data protection, and scalability with predictable low-latency performance on persistent storage.

- **Platform coordinator nodes.** Coordination of guest application workload deployment is handled by these nodes. The servers run a custom implementation of Apache ZooKeeper, running as a Windows service. Per the Apache ZooKeeper website, ZooKeeper is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services. On Apprenda, the platform coordinator nodes maintain knowledge of the topology of guest application workloads across all nodes, as well as any shared configuration in use by Apprenda components.

An optimal platform installation requires an odd number of platform coordination nodes, because a majority of extant platform coordination nodes ($(n+1)/2$, where n =the number of nodes) must be up and running for the platform to function properly. We recommend starting with three dedicated platform coordination nodes, which allows the platform to function as long as any two nodes are running. Additional nodes can be added as needed after the environment is up and running.

- **Cache nodes.** These nodes house the Apprenda caching Windows service, a distributed Redis-based in-memory caching implementation for all platform Windows and Linux servers. Each cache node can support one instance of the caching service per processor core; the number of processor cores/caching service instances should be used as the number of service ports specified for the cache in the Apprenda installer. We recommend two dedicated cache nodes with two cores/caching service instances each to help with load distribution and mitigate infrastructure failure risk. Coordinators may coexist with cache nodes.
- **Windows application servers.** All Windows application servers host the Apprenda Windows host service. The Windows host enables the hosting of Windows Communication Foundation (WCF) and Windows services, thereby allowing both key platform and guest application service components to be hosted on these servers. It is necessary that at least one Windows application server per cloud host Apprenda's storage controlling services, which interface with SQL Server and Oracle to configure databases.

These servers are required to have SQL Server Management Objects (SMO) 2012 installed. During installation, the platform marks any Windows application servers with SMO installed as capable of hosting the storage controlling services and deploys this component to those servers. It installs the required SMO version on a single application server if no suitable host is found.

To make sure that the storage controlling services are highly available, we recommend installing a supported version of SMO (version 11.0 or higher) on two servers that are designated as application servers prior to running the Apprenda installer, because this results in both servers being designated as storage controlling services hosts. As needed, after installation, additional application servers can be configured as storage controlling services hosts by installing SMO on the servers and then designating them as such in the SOC.

- **Windows web servers.** Windows web servers are front-end web servers that host .NET-based UIs through IIS. Using portals, the Apprenda platform allows developers and platform operators to create ad-hoc web farms for .NET guest application UI workloads at any time (as long as there is sufficient infrastructure). The sum of your Windows web servers represents the compute power of your platform specifically for hosting .NET guest application UI workloads.

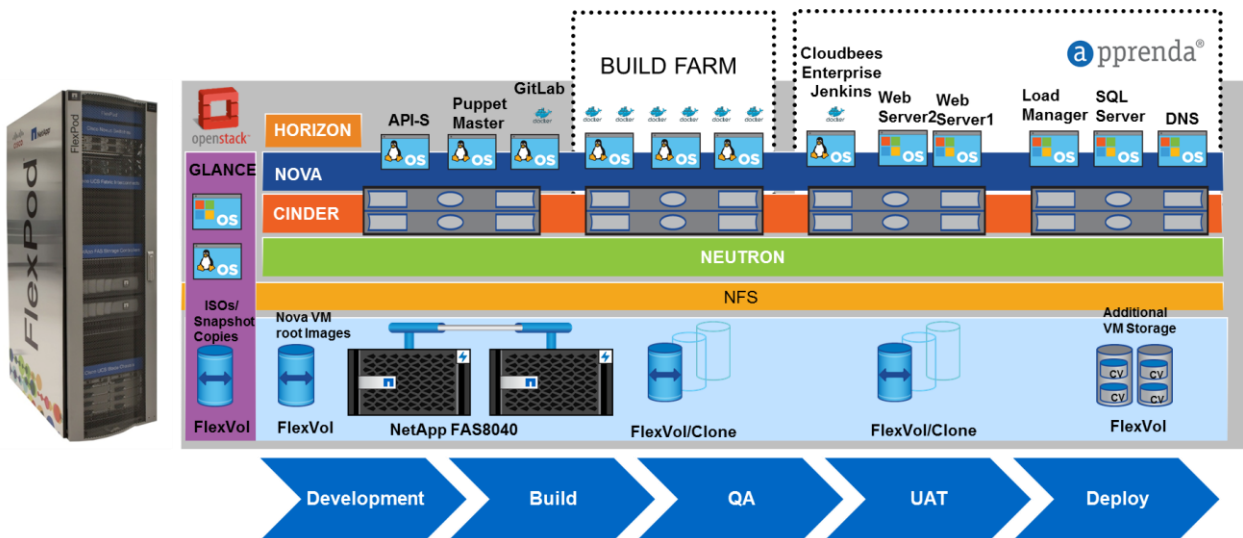
Note that all Windows web servers are capable of hosting WCF and Windows services. This is necessary for the nodes to host the presentation controlling services (see the Apprenda software inventory, later in this report), which allows management of .NET UIs in IIS. Therefore, all Windows web servers are also marked as Windows application servers (described in the preceding section) after installation, even if they were designated as web servers only in the Apprenda installer. By default, however, the platform deploys WCF/Windows services to web servers only if there are no dedicated application servers available.

Note: The same hosts can be used as web and application servers. In addition, coordinators and cache nodes can be used as web and/or app servers.

- **Linux servers (optional).** Linux servers host Java web application components, which are deployed and managed on top of individual Tomcat/JBoss instances by the Apprenda Linux container.
- **Oracle RDBMS (optional).** The platform manages Oracle RDBMS installations on your behalf to provision and configure storage for guest applications.
- **AD FS nodes (optional).** As an install-time option, the platform can be configured to support identity federation using Active Directory federation services (AD FS). If you choose to use AD FS for identity federation, we recommend creating an AD FS web farm consisting of no less than two AD FS nodes backed by a SQL Server failover cluster. Contact support@apprenda.com for information about setting up an AD FS web farm prior to running the Apprenda installer.

Note that all AD FS nodes—including those that constitute an AD FS web farm—are automatically designated as Windows application servers. This is necessary for the nodes to host the federation service (see the Apprenda software inventory, later in this report), which allows the platform to interface with AD FS.

Figure 8) Apprenda, Jenkins and OpenStack logical architecture.



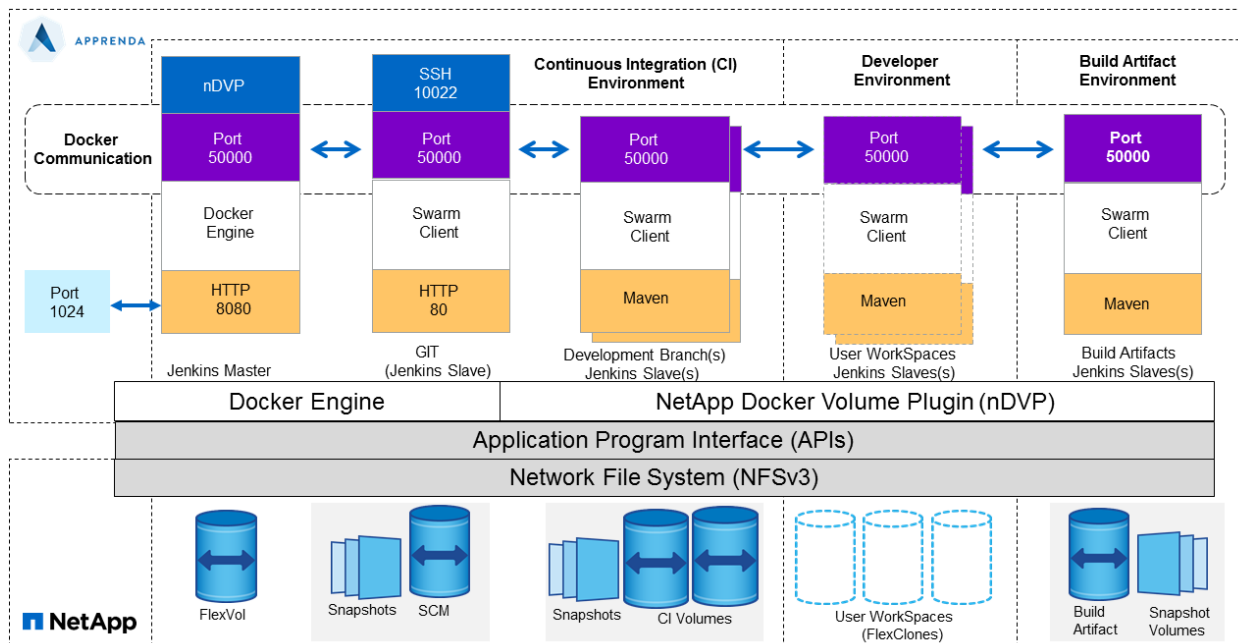
6 Solution Verification

The NetApp Jenkins integration with Apprenda enables developers to implement a robust CI and CD workflow on the same platform. The Docker engine managed by Apprenda along with nDVP provides persistent data to the code and accelerates builds and the applications being developed and deployed, reducing the lead time and allowing for faster time to market. This integration provides horizontal scaling at the compute layer and at the storage layer. NetApp storage also provides storage space efficiency, high availability, and uptime for the data generated during the agile workflows.

6.1 Continuous Integration Pipeline

The CI pipeline includes the source code version control, CI environment, developer workspaces, and build artifact environment. These environments are hosted on the Apprenda platform.

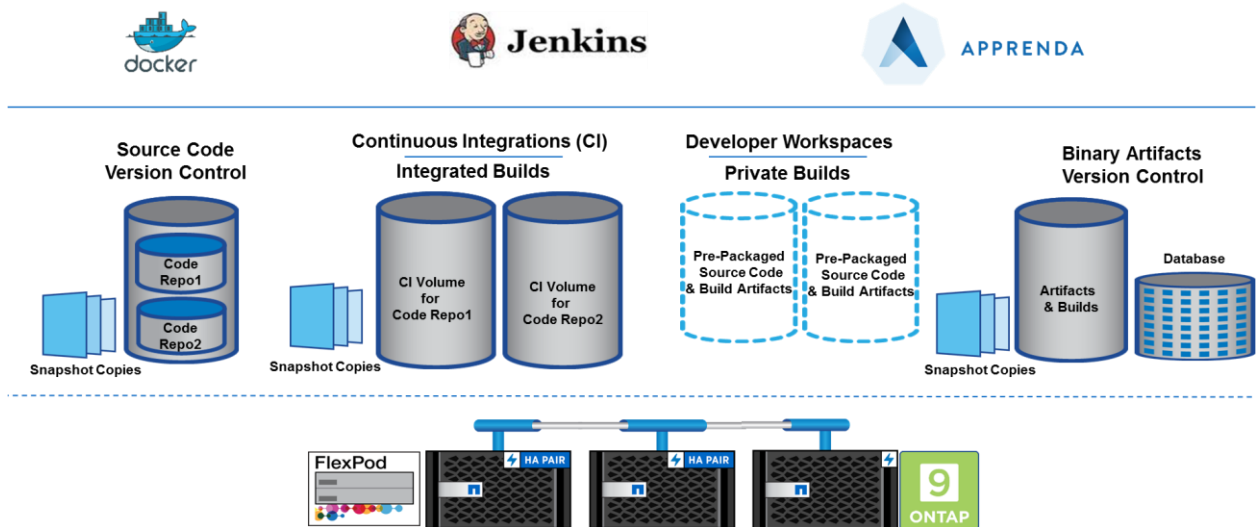
Figure 9) Automating CI workflow with Apprenda Jenkins on NetApp using Docker.



As illustrated in Figure 9, Apprenda creates a Jenkins master running in a Docker container and a Docker volume for storing its home directory. Under the hood, NetApp Volume Plug-In (nDVP) mounts an NFS volume from ONTAP as the Jenkins master's home directory. Because the Jenkins master's home directory is stored as a Docker volume, it enables Apprenda to migrate the Jenkins master to another VM to either retire the VM, upgrade the physical node, or perform other maintenance activities without losing data and affecting the CI pipeline.

All the other components, for example, software control management (SCM) tools such as GIT, CI, or development branches, user workspaces, and the build artifact, run in sibling Docker containers and communicate with the Jenkins master on port 50000. All the Docker containers in the CI pipeline in the Jenkins master are mounting volumes from NetApp for persistent storage using nDVP.

Figure 10) NetApp integrations for the CI workflow.



The CI workflow starts with the source code management software, such as GitLab or Perforce. Jenkins deploys the SCM as a Docker container. All source code is stored and managed in the SCM repository on a Docker volume managed by nDVP. Having a local SCM volume on NetApp has offers the following benefits:

- Avoid compute and network resource overhead for every “git clone” operation from a local or private hosted code repository.
- Creating Snapshot copies on the local code repository provides data protection from any accidental code corruption or loss and recovers very quickly to a stable state.
- Identifying errors in the code from an unsuccessful or build failure is quick using “git bisect” without any resource overhead.
- As a best practice recommendation, no builds are performed on this SCM volume. Only different versions of code are checked out and checked in to this volume. This eliminates any resource overhead from high build traffic by different users.
- Depending on the size of the codebase, NetApp provides compaction and deduplication of the source code for improved storage space efficiency, which can provide significant space and cost savings in a cloud environment.

After the source code is in the repository, the Jenkins administrator creates CI or development branches of the source code. These CI branches are separate Docker containers that also mount NetApp volumes using nDVP. By using the Jenkins plug-in, administrators can automatically create the container, create and mount the persistent volume, and copy the required source code into the CI volume. Each CI container includes all the components necessary for the development of the selected source code.

As users create and submit code to the SCM, their changes are automatically pushed to the appropriate CI volumes, and regularly scheduled builds are performed. If the build is successful, a NetApp Snapshot copy is automatically created on the CI volume, and the Snapshot copy is correlated to the respective build number. These Snapshot copies form the basis for NetApp FlexClone volumes used in the developer user workspaces. The CI volume and all associated Snapshot copies contain the source code and the final executable build, as well as all the tools, RPMs, libraries, compilers, and other dependencies required for further development or deployment.

When users want to work on a specific build of a CI branch, they simply request a workspace based on the build number. The Jenkins plugin automatically creates another Docker container, but this container

mounts a FlexClone volume of the CI volume based on the Snapshot copy associated with the requested build number. This has several advantages that improve developer productivity significantly:

- The user workspaces are instantly cloned using the FlexClone volumes and mounted on Docker containers.
- The ownership of the files and directories is also changed to the respective user instantly.
- This eliminates multiple streams of “git clone” or “rsync” process to create multiple copies of the source code and the compute and network overhead in that process.
- These cloned workspaces are extremely space efficient. They take a small fraction of the actual size of the CI or the development instances.
- Because the user workspaces are clones from a parent CI volume that contains precompiled and compiled objects, the user builds limit the lines of code change and do not perform a full build all the time. The incremental builds reduce the build time considerably.

As the user makes and submit changes to the SCM, those changes are pushed to the appropriate CI volumes and are then included in the regularly scheduled builds. If a new build occurs before they submit changes they have been working on, they are notified that a new build exists and are given the option to move their changes into a workspace using the newest build.

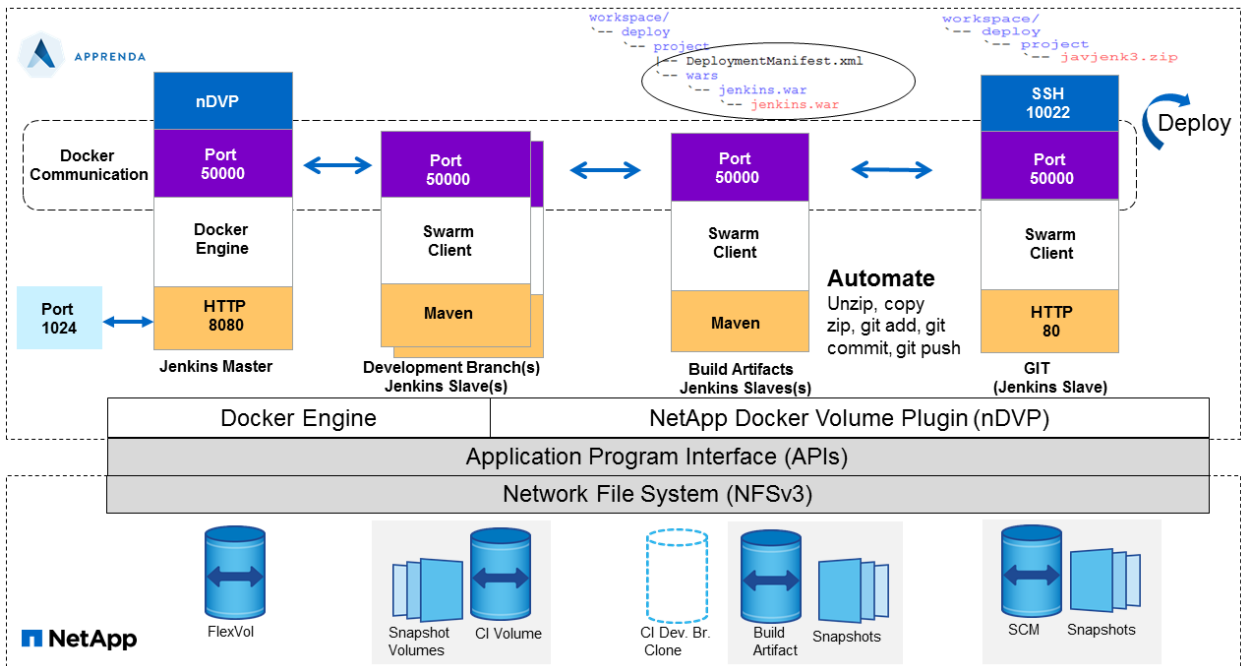
The last step in the CI workflow is build artifact management. Jenkins creates build artifact repositories, which are container instances that mount a NetApp volume to store all the build artifacts such as the tools, libraries, compilers, and RPMs that are required during a build process. The entire contents of the successful builds from the CI or development volumes are zipped and copied into the build artifact volume. The zip file can be used by the QA/staging teams at a later stage, or developers may want to replicate a scenario from an old build to recreate and fix bugs or errors reported by users. The zip file consists of the entire environment of a QA or a developer to run tests. Even if the tool versions may have changed as the applications start to evolve, recreating the environment with the right set of tools in retrospect should not be a challenge with the zip files in the build artifact volume. Storing the build artifacts on the NetApp storage allows them to be backed up more frequently, as well as providing storage efficiencies such as deduplication on what could ultimately be a large repository of relatively similar objects.

6.2 Continuous Delivery

Continuous delivery (CD) is an automated process of deploying and releasing applications into production. Apprenda reduces the friction between building and deploying applications. The NetApp and Apprenda integration allows organizations to define and perform user acceptance tests (UATs) and publish applications at scale in a single location and across multiple hybrid clouds. This integration allows businesses to accelerate the process from designing the user requirements to finally deploying the application to production.

The continuous delivery pipeline starts after the code is built successfully in the CI phase and an executable file in the form of a .war file (if it is a Java-based application) is created. As mentioned previously, the content of the entire CI or the development instance is zipped and copied into the build artifact volume. However, there is a small enhancement made during this process to automate the CD process with Apprenda.

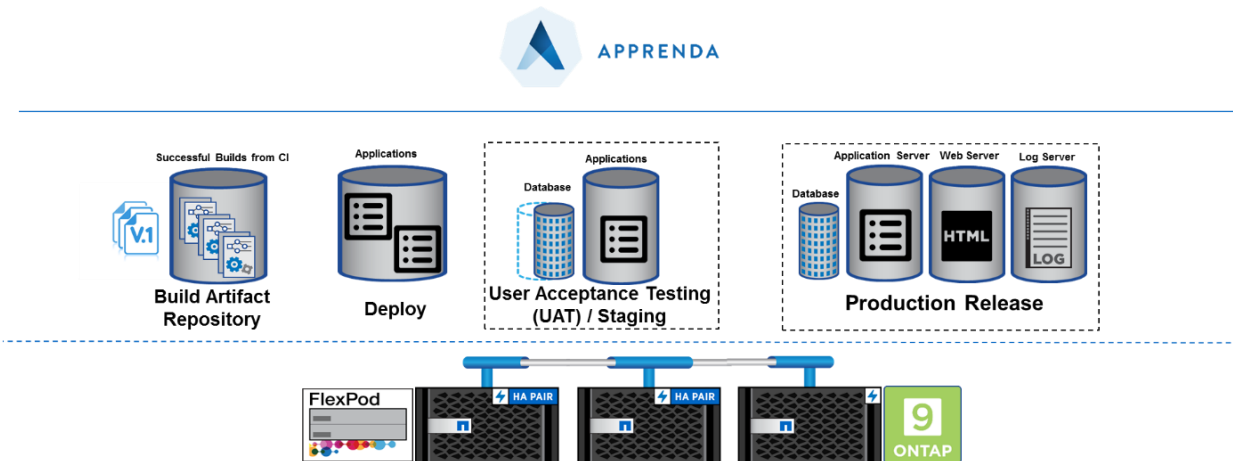
Figure 11) Automating CD workflow with Apprenda and NetApp.



The Jenkins plugin provided by Apprenda pushes the application archive file for the deployment. This archive consists of a deployment manifest that contains the executable files for every successful build. Then, a series of `git` commands are performed to move this Apprenda archive into a specific location under the SCM tool directory. This process is automatically repeated for any new build.

For application deployment, the Apprenda Jenkins plug-in automatically looks for the archive file in the appropriate location in the SCM and deploys it to the platform. The new application instance, containerized on Apprenda, is now available for testing and verification that can be performed by any appropriate tools, standalone or integrated with Jenkins. By automating the deployment process with any build, QA teams can easily and consistently redeploy applications for UAT and staging for further application testing with databases if needed. The UAT for databases can be further accelerated by using NetApp technologies like FlexClone and SnapRestore®. Upon successful testing, the application is released into production and managed at scale by the Apprenda PaaS management plane.

Figure 12) NetApp integrations for CD.



Conclusion

By building on the flexibility, scalability, and reliability of the FlexPod converged infrastructure, the Apprenda PaaS solution delivers innovative application development capabilities in a turnkey package. Customers can size the system for their current needs without giving up any features or flexibility and then nondisruptively scale the system in any area as required by their business needs or application demands. The joint NetApp and Apprenda solution improves developer productivity, reduces build times, and provides storage space efficiency that drives costs down. The following data points summarize the benefits of the joint NetApp and Apprenda and Jenkins integrations:

Ease of Use

- Using Apprenda's PaaS capabilities, developers can easily create development workspaces that take advantage of the latest technologies using containers.
- Developers can continue to use existing CI/CD tools such as Jenkins without disrupting their current application workflow processes.
- Developers don't have to know the infrastructure; the Apprenda PaaS takes care of this for them. All developers need to specify are the requirements, and the PaaS takes care of the rest.
- No special skill set is required for businesses to invest to manage infrastructure, and Apprenda and OpenStack on FlexPod improve the lead times for deploying applications.

Efficiency

- The hardware and process efficiencies in this solution allow two to three times more developers to be supported on the same infrastructure due to lower demands for compute and network resources.
- Using Snapshot copies and FlexClone volumes, storage utilization can be decreased by over 50 times.
- The addition of compaction to ONTAP 9 can reduce storage requirement in the code depots and during volume creation by up to 66%.

Productivity

- Faster workspace creation means that the developers can become productive 2 to 3 times more quickly. For medium to large development teams, this can translate into hundreds of hours of added productivity every month.

Manageability and Choice

- Because the features of ONTAP are shared across all the platforms supported by NetApp, users are not limited to where they build or reply applications. Furthermore, they have the freedom to migrate their data and applications across public and private cloud platforms without sacrificing function or requiring changes to their applications or data.

Acknowledgements

This document is the result of the work, documentation, and assistance provided David Arnette, Amit Borulkar, and Bikash Roy Choudhury from NetApp and Sasha Jeltuhin of Apprenda.

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 FAS Storage

The following links provide additional information about NetApp FAS 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>

- ONTAP 9.x Documentation
<http://mysupport.netapp.com/documentation/productlibrary/index.html?productID=62286>
- TR-4476: NetApp Data Compression and Deduplication Data ONTAP 8.3.1 and Above
<http://www.netapp.com/us/media/tr-4476.pdf>

Red Hat OpenStack Platform

The following links provide additional information about Red Hat OpenStack Platform:

- Red Hat OpenStack Platform
<https://www.redhat.com/en/technologies/linux-platforms/openstack-platform>
- NetApp OpenStack Deployment and Operations Guide
<http://netapp.github.io/openstack-deploy-ops-guide/liberty/content/>
- TR-4506: Red Hat OpenStack Platform on FlexPod
<https://www.netapp.com/us/media/tr-4506.pdf>
- Red Hat Enterprise Linux OpenStack Platform 6 on FlexPod
http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_open

Apprenda PaaS

The following links provide additional information about Apprenda PaaS:

- Apprenda Platform Website
<https://apprenda.com/platform/>
- TR-4559: Continuous Integration (CI)/Continuous Deployment (CD) Pipeline with Apprenda on FlexPod with ONTAP 9
<https://www.netapp.com/us/media/tr-4559.pdf>
- Continuous Integration (CI) Pipeline with CloudBees Enterprise Jenkins and ONTAP 9
<https://www.netapp.com/us/media/tr-4547.pdf>

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>

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 1.0	April 2017	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 © 2017 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).

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS DOCUMENT ARE PRESENTED "AS IS," WITH ALL FAULTS. NETAPP, ALL PRODUCT VENDORS OR MANUFACTURERS IDENTIFIED OR REFERENCED HEREIN ("PARTNERS") AND THEIR RESPECTIVE SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, OR WITH RESPECT TO ANY RESULTS THAT MAY BE OBTAINED THROUGH USE OF THE DESIGNS OR RELIANCE UPON THIS DOCUMENT, EVEN IF NETAPP, ITS PARTNERS OR THEIR RESPECTIVE SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

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

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-1113-DESIGN-0417