



NetApp Verified Architecture

# FlexPod Datacenter with Microsoft SQL Server 2014, AlwaysOn Availability Groups, and NetApp All Flash FAS NVA Design

David Arnette, NetApp  
July 2016 | NVA-0025-DESIGN | Version 1.0

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary</b>	<b>5</b>
1.1	FlexPod Program Benefits	5
<b>2</b>	<b>Solution Overview</b>	<b>7</b>
2.1	Target Audience	7
2.2	Solution Technology	7
2.3	Use Case Summary	9
<b>3</b>	<b>Technology Overview</b>	<b>9</b>
3.1	FlexPod	9
3.2	Cisco Unified Computing System	11
3.3	Cisco Nexus 9000 Series Switches	15
3.4	NetApp All Flash FAS	15
3.5	VMware vSphere	24
3.6	Microsoft SQL Server 2014	25
<b>4</b>	<b>Technology Requirements</b>	<b>27</b>
4.1	Hardware Components	27
4.2	Software Components	27
<b>5</b>	<b>Solution Design</b>	<b>28</b>
5.1	Cisco Nexus Network Design	28
5.2	Cisco Unified Computing System Design	30
5.3	NetApp AFF Storage Design	33
5.4	VMware vSphere Design	41
5.5	SQL Server 2014 Design	43
<b>6</b>	<b>Solution Verification</b>	<b>44</b>
6.1	Performance Validation	45
6.2	SnapCenter Validation	47
6.3	Resiliency Validation	50
<b>7</b>	<b>Best Practices</b>	<b>52</b>
<b>8</b>	<b>Conclusion</b>	<b>53</b>
	<b>Acknowledgements</b>	<b>54</b>
	<b>References</b>	<b>54</b>
	Cisco Unified Computing System	54
	Cisco Nexus Networking	54

NetApp FAS Storage .....	54
NetApp SnapCenter .....	55
VMware vSphere .....	55
Interoperability Matrixes .....	55

<b>Version History .....</b>	<b>55</b>
------------------------------	-----------

## LIST OF TABLES

Table 1) Hardware components .....	27
Table 2) Solution software components .....	27
Table 3) NetApp AFF8000 storage system technical specifications .....	34
Table 4) Storage aggregate configuration .....	36
Table 5) Volume and LUN configuration.....	37
Table 6) NetApp VSC VM configuration .....	39
Table 7) Supported storage types: SnapCenter Plug-In for Microsoft SQL Server.....	40
Table 8) VMware vCenter Server VM configuration .....	42
Table 9) vCenter SQL Server database VM minimum configuration .....	42
Table 10) Production SQL Server 2014 VM configuration.....	43
Table 11) Database LUN and volume configuration.....	45

## LIST OF FIGURES

Figure 1) FlexPod component families .....	6
Figure 2) High-level solution architecture .....	8
Figure 3) FCoE direct-connection topology .....	11
Figure 4) Cisco UCS components .....	12
Figure 5) ONTAP.....	18
Figure 6) SnapCenter architecture .....	21
Figure 7) The NetApp Virtual Storage Console .....	23
Figure 8) VMware vSphere feature overview .....	25
Figure 9) Overall solution topology.....	29
Figure 10) Cisco UCS backplane configurations: VIC 1340 .....	32
Figure 11) Server logic configuration.....	33
Figure 12) ONTAP in a SAN-only environment .....	33
Figure 13) ONTAP in a NAS-only environment .....	34
Figure 14) Multipath HA to DS2246 SSD shelves .....	36
Figure 15) Backup of SQL database with SnapCenter.....	48
Figure 16) Database clone operations with SnapCenter .....	48
Figure 17) SQL Management Console-mounted database clones.....	49
Figure 18) Database restore with SnapCenter .....	49

Figure 19) IOPS during link failure tests.....51  
Figure 20) Observed IOPS during controller failover tests. ....52  
Figure 21) Observed IOPS during fabric interconnect failover test.....52

# 1 Executive Summary

NetApp® Validated Architectures (NVAs) describe systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp has developed to address the business needs of customers.

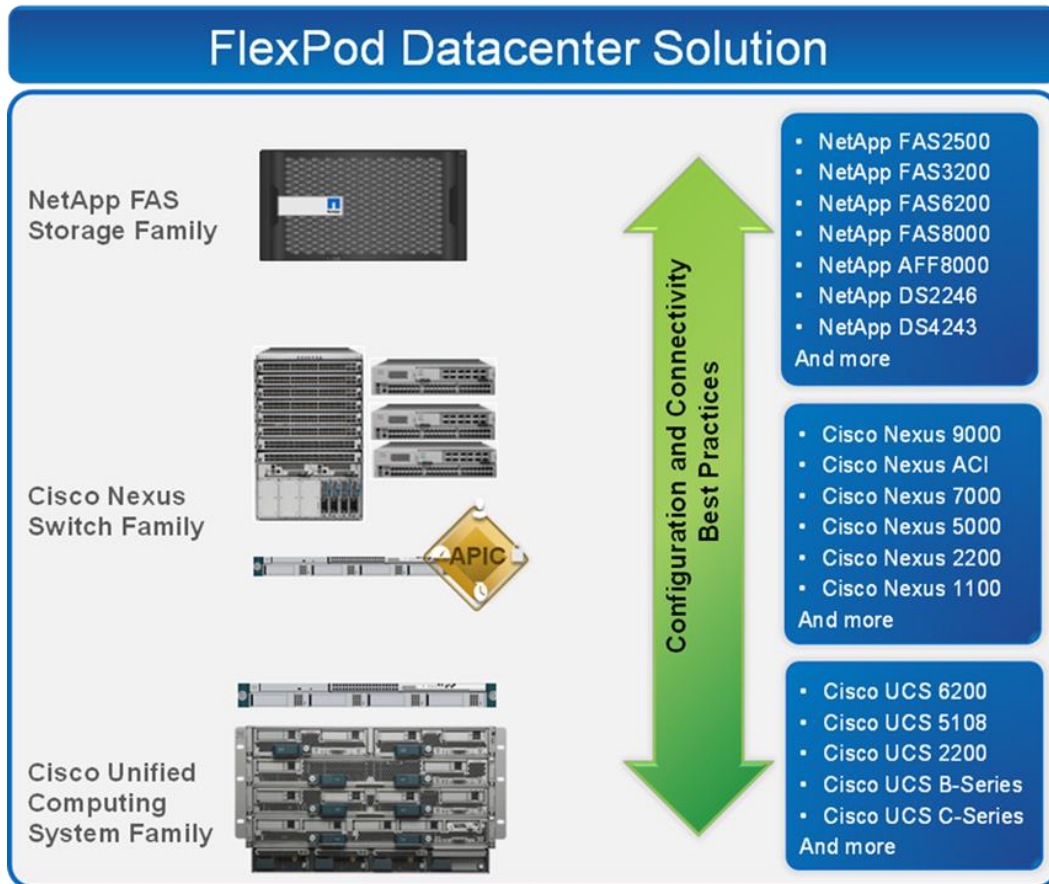
This document describes a reference architecture for Microsoft SQL Server 2014 with AlwaysOn Availability Groups built on the All Flash FAS (AFF) FlexPod® model. This solution includes Cisco UCS B200 M4 blade servers and NetApp AFF8080 EX storage arrays. The other major highlights of this solution are the use of NetApp SnapCenter® enterprise software for application-integrated database backup and recovery, VMware vSphere 6.0 virtualization, and Cisco Nexus 9000 series switches for client traffic.

This document also discusses design choices and best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific components described in this document and are also applicable to other versions.

## 1.1 FlexPod Program Benefits

FlexPod is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp FAS and AFF storage systems. FlexPod is a suitable platform for running a variety of virtualization hypervisors as well as bare-metal operating systems (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 depicts the component families of the FlexPod solution.

Figure 1) FlexPod component families.



FlexPod provides a uniform approach to IT architecture, offering a well-characterized and documented shared pool of resources for application workloads. FlexPod delivers operational efficiency and consistency with the versatility to meet a variety of SLAs and IT initiatives, including the following:

- Application rollouts or migrations
- Business continuity and disaster recovery (DR)
- Desktop virtualization
- Cloud delivery models (public, private, and hybrid) and service models (IaaS, PaaS, and SaaS)
- Asset consolidation and virtualization
- Data center consolidation and footprint reduction

Cisco and NetApp have thoroughly validated and verified the FlexPod solution architecture and its many use cases. In addition, they have created a portfolio of detailed documentation, information, and references to assist customers in transforming their data centers to this shared infrastructure model. This portfolio includes, but is not limited to, 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 (FAQs)
- NVAs and Cisco Validated Designs (CVDs) that focus on a variety of use cases

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

FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment. As a key FlexPod Cooperative Support partner, VMware provides the virtualization-hypervisor and management tools for this verified design with VMware vSphere and VMware vCenter.

## 2 Solution Overview

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

This FlexPod design describes the deployment of Microsoft SQL Server 2014 with AlwaysOn Availability Groups in a VMware vSphere virtualized environment. This AFF design demonstrates that, with the addition of flash storage, FlexPod architectures can meet the most demanding performance requirements and still deliver the values of a standardized shared infrastructure.

### 2.1 Target Audience

The intended audience of this document includes sales engineers, field consultants, professional services, IT managers, partner engineering, and customers who want to take advantage of an infrastructure built to deliver high-performance and high-availability database services.

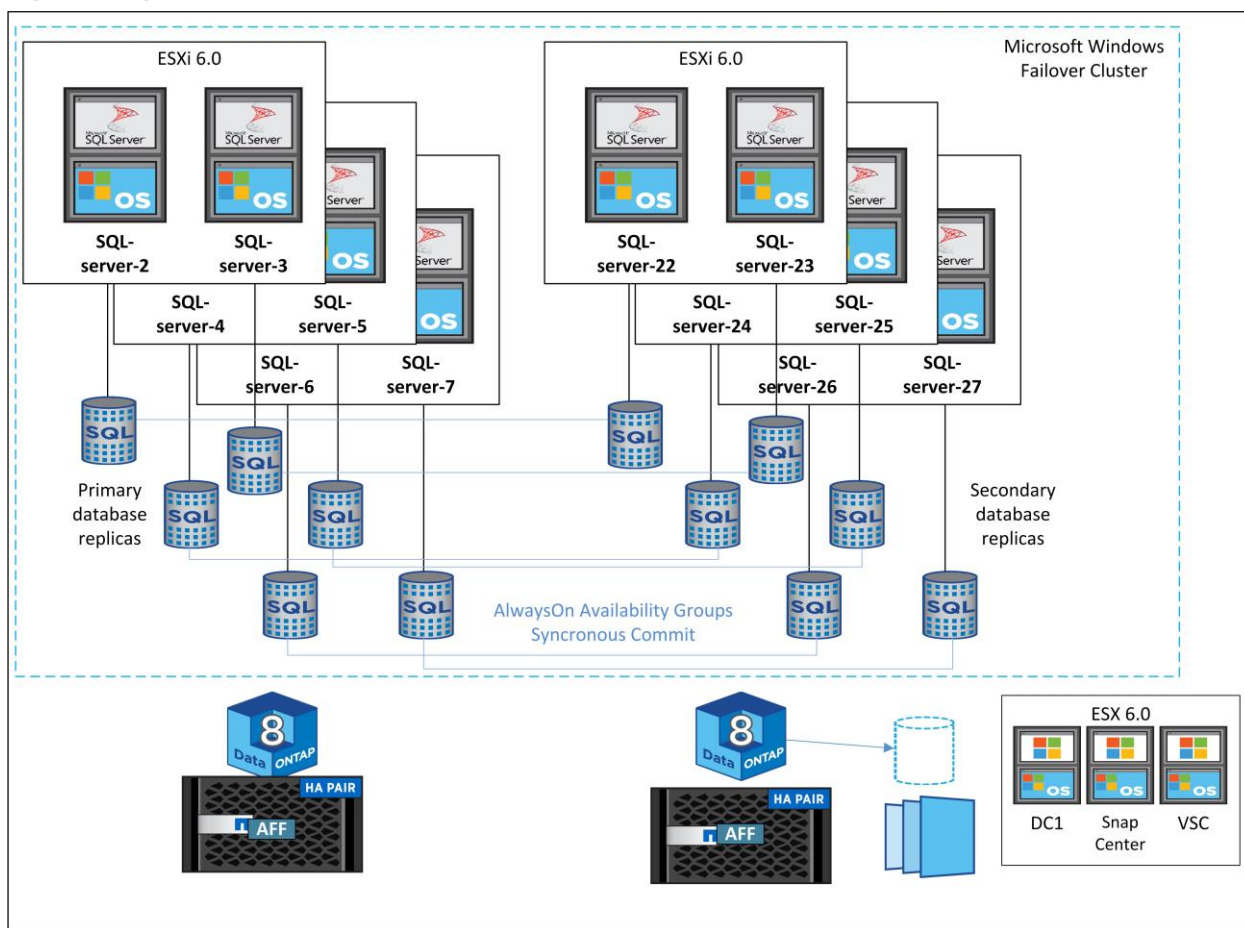
### 2.2 Solution Technology

Traditionally, Microsoft SQL Server databases relied on Microsoft Windows clustering with shared storage to achieve a high-availability (HA) configuration. Microsoft database technology has come a long way since then and is now composed of highly efficient technology based on Windows Server 2012 R2 and AlwaysOn Availability Groups.

This solution is based on SQL Server 2014 AlwaysOn Availability Groups implemented on Windows Server 2012 R2 virtual machines (VMs). SQL Server instances are deployed as VMware vSphere VMs on NetApp AFF storage arrays.

Figure 2 shows the high-level solution architecture, with the association between the primary and secondary database replicas on separate storage.

Figure 2) High-level solution architecture.



This design enables you to service OLTP workloads from the primary database replica while database writes are synchronously committed to the secondary database server. The secondary replica provides near-site disaster recovery and enhances the overall availability of the solution by providing database-level failover. Because AlwaysOn replication enables readable database secondaries, the secondary database replica can be used to offload backup and reporting jobs.

The secondary database replica is hosted on equivalent storage and compute resources to make sure of consistent performance in the event of a failover. As a part of NetApp integrated data management capabilities, SnapCenter Server creates Snapshot® backups of secondary databases and creates clones of databases for secondary processing or test and development activities.

This solution uses NetApp AFF8080 EX storage arrays for both primary and secondary database copies. This provides the highest levels of performance for both copies during normal operations and failover situations. When FCoE is used, LUNs are provisioned to ESXi servers to make vSphere Virtual Machine File System (VMFS) datastores for SQL Server OS disks. In addition, separate dedicated LUNs are provisioned to make datastores for each database. By using this technology, SQL Server 2014 databases deliver excellent performance, and native application integration enables seamless backup, recovery, and cloning.

Cisco UCS fabric interconnects and B-Series servers enhance this configuration by providing flexibility and availability in the compute layer. Redundant hardware and software components eliminate single points of failure and make sure that data traffic is uninterrupted in the event of a component failure. Cisco UCS Service profiles ease deployment and maintenance activities by creating consistent configurations



that are easily updated and can be moved between physical blades as needed for maintenance or upgrades.

Cisco Nexus 9000 switches act as the access layer for the primary and secondary database environments. Cisco Nexus 9396PX switches are deployed in pairs, and the Cisco UCS fabric interconnects are connected with virtual port channels for maximum availability. Cisco Nexus 9000 switches provide 40Gb of switching capability and can participate in the Cisco Application Centric Infrastructure (ACI). However, these switches do not support the FC or FCoE storage protocols. To enable FCoE in this solution, storage controllers are connected directly to the Cisco UCS fabric interconnects with redundant connections for each controller.

When combined into a complete infrastructure, this solution delivers the following benefits:

- Tier 1 SQL Server 2014 database performance on a standardized, shared infrastructure
- Database-level availability by using a SQL Server AlwaysOn synchronous replica (near-site disaster recovery)
- Integrated backup and recovery of SQL Server 2014 databases with NetApp SnapCenter
- Integrated cloning of databases for secondary processing or testing and development
- Hardware-level redundancy for all of the major components by using Cisco UCS, Cisco Nexus, and NetApp availability features

## 2.3 Use Case Summary

The FlexPod Datacenter, Microsoft SQL Server 2014, and NetApp AFF solution is designed to provide enterprises with the performance, manageability, and reliability necessary for tier 1 application databases. To this end, the following use cases have been configured and tested in the lab to demonstrate the performance and functionality of this design:

- SQL Server 2014 AlwaysOn Availability Groups with one primary database and one replica. This configuration generates 150,000 to 200,000 IOPS with an average read latency below 1ms for a typical OLTP workload while in a synchronized state. Secondary database copies run on separate storage and compute resources that have capabilities equivalent to the primary.
- Failover of the SQL Server 2014 AlwaysOn Availability Groups to the secondary copy. This use case demonstrates database resiliency and the comparable performance of the secondary copy in a synchronized AG configuration under a typical OLTP workload.
- Backup and recovery of SQL Server 2014 databases with NetApp SnapCenter 1.1 and the SnapCenter plug-ins for SQL Server. Backups are taken from the secondary database copy to offload backup and validation processing from the primary infrastructure.
- Cloning of SQL Server 2014 databases with NetApp SnapCenter for use in application testing and development environments.
- Failure testing of various infrastructure components while the environment is operating under a heavy OLTP workload to verify the resiliency and reliability of the overall architecture.

## 3 Technology Overview

### 3.1 FlexPod

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

- Cisco UCS
- Cisco Nexus switches
- NetApp FAS and AFF storage 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). FlexPod can also scale out for environments that need multiple consistent deployments (for example, rolling out additional FlexPod stacks). Although FlexPod delivers a baseline configuration, it can also be flexibly 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 each implementation. 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. They also support the same features and functionality that are required by the configuration and connectivity best practices of FlexPod.


FlexPod addresses four primary design principles: application availability, scalability, flexibility, and manageability. These architecture goals are as follows:

- **Application availability.** Makes sure that services are accessible and ready to use.
- **Scalability.** Addresses increasing demands with appropriate resources.
- **Flexibility.** Provides new services or recovers resources without infrastructure modification requirements.
- **Manageability.** Facilitates efficient infrastructure operations through open standards and APIs.

### FlexPod: FCoE Direct-Connect Design

As noted previously, flexibility is a key design principle of FlexPod. Although the Cisco Nexus 9000 series switches used in this design do not support FCP or FCoE, FlexPod can still support these protocols by connecting unified target adapter (UTA) ports on the NetApp storage controllers directly to the Cisco UCS fabric interconnects. This arrangement uses the FC features of the fabric interconnects to provide name services and zoning capabilities. As a result, Cisco UCS servers can boot from and access FC or FCoE storage without additional FC switches. Figure 3 shows the basic topology of this direct-connection design.

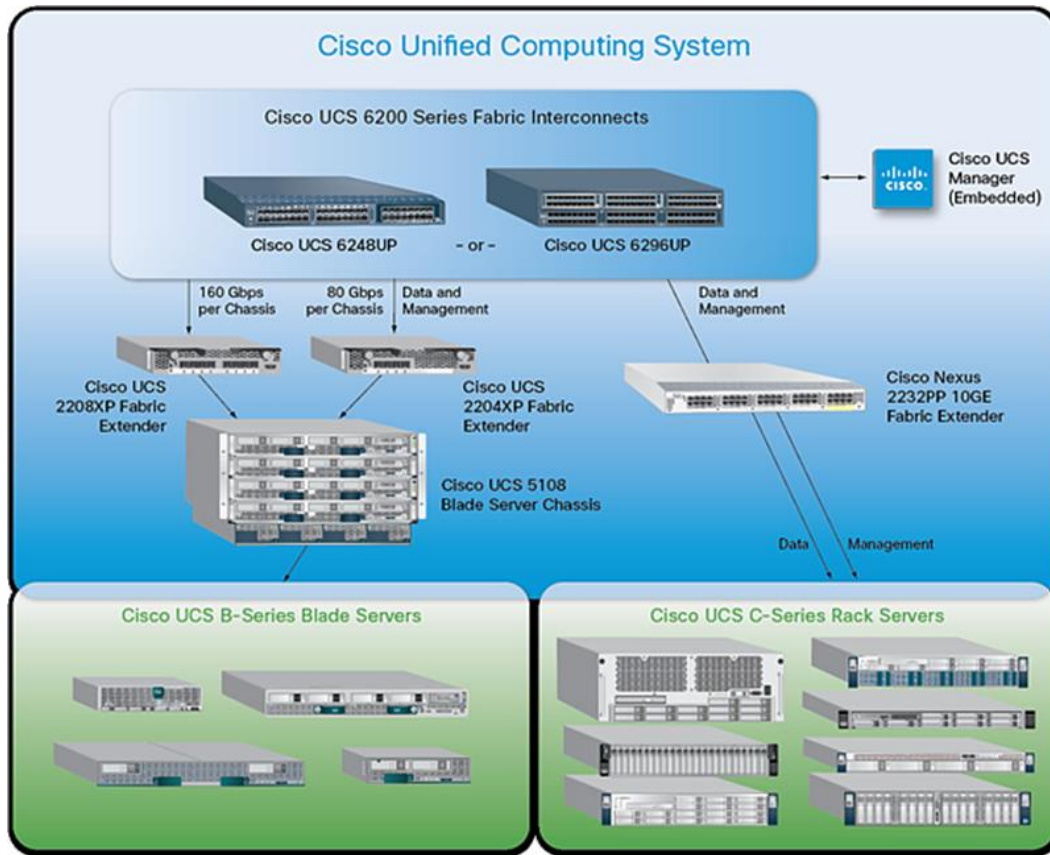
Figure 3) FCoE direct-connection topology.



### 3.2 Cisco Unified Computing System

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 4) 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 network-attached storage (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, next-generation 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 Cisco's Virtual Interface Card (VIC) 1340 or VIC 1240 adapter, which provides Ethernet and FCoE.

## Cisco VIC 1340

The Cisco UCS VIC 1340 is a 2-port 40Gbps Ethernet or dual 4 x 10Gbps Ethernet, FCoE-capable modular LAN on motherboard (mLOM) designed exclusively for the M4 generation of Cisco UCS B-Series blade servers. When used in combination with an optional port expander, the capabilities of the Cisco UCS VIC 1340 are extended to two 40Gbps Ethernet ports.

The Cisco UCS VIC 1340 enables a policy-based, stateless, agile server infrastructure that can present over 256 PCIe standards-compliant interfaces to the host. These interfaces can be dynamically configured as either network interface cards (NICs) or host bus adapters (HBAs). In addition, the Cisco



UCS VIC 1340 supports Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to VMs, simplifying server virtualization deployment and management.

## 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.

You can export system configuration information to configuration management databases to facilitate processes based on IT Infrastructure Library concepts. Service profiles benefit both virtualized and nonvirtualized environments. They increase the mobility of nonvirtualized servers, such as when you move workloads from server to server or take a server offline for service or upgrade. You can also use profiles in conjunction with virtualization clusters to bring new resources online easily, complementing existing VM mobility.

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](#).

A server's identity is made up of many properties, including the UUID; the boot configuration; the BIOS configuration; the number of NIC, MAC, and IP addresses; the number of HBAs; HBA WWNs; and so on. Some of these parameters reside in the hardware of the server itself, including the BIOS firmware version, the BIOS settings, the boot order, the FC boot settings, and so on. Other settings are kept on your network and storage switches, such as VLAN assignments, FC fabric assignments, QoS settings, ACLs, and so on. These configurations result in the following server deployment challenges:

- The response to business needs is slow because of lengthy and tedious deployment processes.
- Every deployment requires coordination between the server, storage, and network teams:
  - Firmware and settings for hardware components
  - Appropriate LAN and SAN connectivity
  - Settings tied to physical ports and adapter identities
  - Manual, error-prone processes that are difficult to automate
- Complexity leads to higher opex costs:
  - Outages caused by human errors
  - Static infrastructure, leading to overprovisioning

- Complexity, which causes limited OS and application mobility

Cisco UCS has addressed these challenges with the introduction of service profiles, which enable integrated, policy-based infrastructure management. Cisco UCS service profiles hold nearly all of the configurable parameters that are required to set up a physical server. A set of user-defined policies (rules) allows quick, consistent, repeatable, and secure deployments of Cisco UCS servers.

Cisco UCS service profiles contain values for a server's property settings, including virtual network interface cards (vNICs), MAC addresses, boot policies, firmware policies, fabric connectivity, external management, and high-availability information. When these settings are abstracted from the physical server into a Cisco service profile, the service profile can then be deployed to any physical compute hardware within the Cisco UCS domain. Furthermore, service profiles can be migrated at any time from one physical server to another. This logical abstraction of the server personality removes dependency on the hardware type or model and is a result of Cisco's unified fabric model rather than the overlying software tools on top.

Cisco is the only hardware provider to offer a truly unified management platform, with Cisco UCS service profiles and hardware abstraction capabilities extending to both blade and rack servers. Some of the key features and benefits of Cisco UCS service profiles are discussed in the following sections.

## Service Profiles and Templates

Service profile templates are stored in the Cisco UCS 6200 Series fabric interconnects for reuse by server, network, and storage administrators. Service profile templates consist of server requirements and the associated LAN and SAN connectivity. Service profile templates allow different classes of resources to be defined and applied to a number of resources, each with its own unique identities assigned from predetermined pools.

Cisco UCS Manager can deploy a service profile on any physical server at any time. When a service profile is deployed to a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. A service profile template parameterizes the UUIDs that differentiate server instances.

This automation of device configuration reduces the number of manual steps required to configure servers, NICs, HBAs, and LAN and SAN switches.

### 3.3 Cisco Nexus 9000 Series Switches

The Cisco Nexus 9000 Series delivers proven high performance, high density, low latency, and exceptional power efficiency in a broad range of compact form factors. These switches, running in NX-OS software mode, offer both modular and fixed 10/40/100GbE switch configurations with scalability up to 30Tbps of nonblocking performance. They provide less than 5-microsecond latency; 1,152 10Gbps or 288 40Gbps nonblocking layer 2 and layer 3 Ethernet ports; and wire-speed VXLAN gateway, bridging, and routing support.

The Cisco Nexus 9396X switch delivers comprehensive line-rate layer 2 and layer 3 features in a 2RU form factor. It supports line-rate 1/10/40GbE with 960Gbps of switching capacity. It is ideal for top-of-rack and middle-of-row deployments in both traditional and [Cisco ACI](#)-enabled enterprise, service provider, and cloud environments.

For more information, see the [Cisco Nexus 9000 Series Switch product page](#).

### 3.4 NetApp All Flash FAS

Built on more than 20 years of innovation, ONTAP® has evolved to meet the changing needs of customers and help drive their success. NetApp ONTAP provides a rich set of data management features and clustering for scale-out, operational efficiency, and nondisruptive operations to offer 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). This model requires a pool of compute, network, and storage to serve a wide range of applications and deliver a wide range of services. Innovations such as NetApp ONTAP are fueling this revolution.

NetApp storage systems offer a completely unified storage architecture. The term unified refers to a family of storage systems that simultaneously support SAN and NAS across many operating environments, including 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 (10/100/1000 or 10GbE), FCoE (10Gb), and native FC (2, 4, 8, or 16Gbps).

This FlexPod Datacenter solution includes the NetApp AFF8000 series unified scale-out storage system. Powered by NetApp Data ONTAP 8.3, the AFF8000 series unifies the SAN and NAS storage infrastructure with the performance of SSD. The AFF8000 features a multiprocessor Intel chipset and leverages high-performance memory modules, NVRAM to accelerate and optimize writes, and an I/O-tuned PCIe gen3 architecture that maximizes application throughput. The AFF8000 series comes with integrated UTA2 ports that support 16Gb FC, 10GbE, and FCoE.

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

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

## Outstanding Performance

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

## Enhancing Flash

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

NetApp ONTAP FlashEssentials is the power behind the performance and efficiency of All Flash FAS. ONTAP is well-known, but it is not widely known that ONTAP with the WAFL® (Write Anywhere File Layout) file system is natively optimized for flash media.

ONTAP and WAFL include the following key features to optimize SSD performance and endurance:

- NetApp storage efficiency technologies deliver space savings of up to tenfold or more. Features include inline compression, deduplication, and thin provisioning. Savings can be further increased by using NetApp Snapshot and NetApp FlexClone® technologies.
- Multiple writes are coalesced and written as a unit. The resulting reduction in storage overhead during write workloads improves performance and flash media longevity.



- AFF systems include a flash-optimized I/O path to maximize performance in a pure flash environment.
- With advanced drive partitioning, SSDs can be shared among controllers, increasing usable capacity and allowing more flexibility in configuration.
- AFF controllers can be used within a larger ONTAP cluster, enabling nondisruptive workload migration between the flash and hybrid tiers.
- Quality-of-service capability safeguards service-level objectives in multiworkload and multitenant environments.

The parallelism built into ONTAP, combined with multicore CPUs and large system memories in the NetApp AFF8000 storage controllers, takes full advantage of SSD performance. With the media optimizations built into ONTAP, NetApp provides up to a five-year warranty with all SSDs with no restrictions on the number of drive writes.

## NetApp ONTAP

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

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

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

If the newer controller can coexist with the original controller, you now have two storage controllers that must be individually managed, and there are no native tools that can balance or reassign workloads across them. The situation becomes even more difficult as the number of controllers increases. If the scale-up approach is used, the operational burden increases consistently as the environment grows, and the end result is a very unbalanced and difficult-to-manage environment. Technology refresh cycles require substantial planning in advance, lengthy outages, and configuration changes, which introduce risk into the system.

## Scale-Out

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

The benefits of scale-out include the following:

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

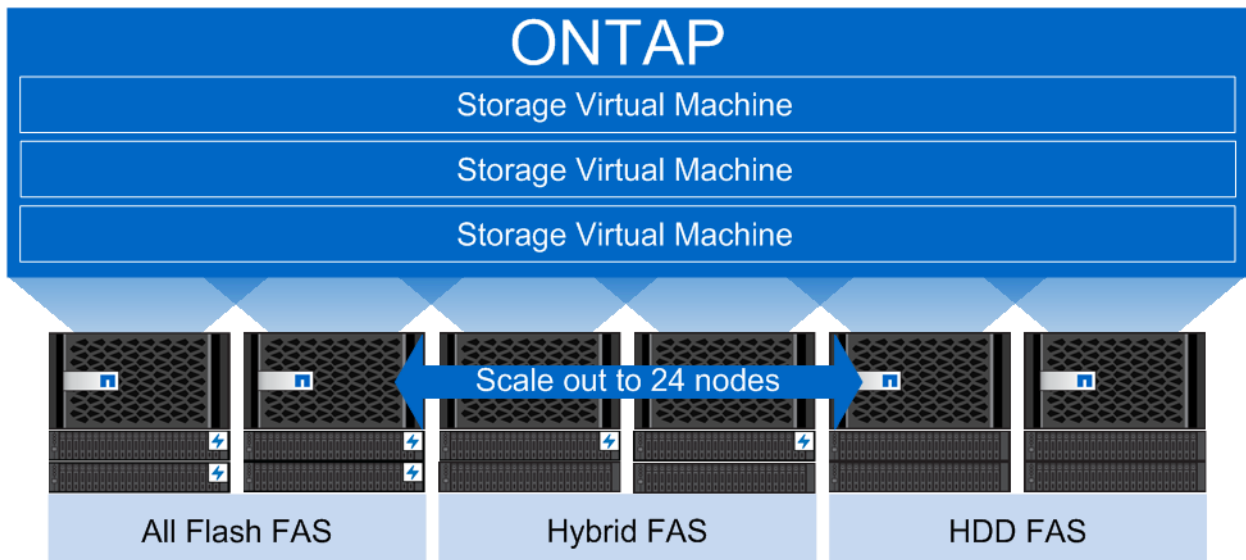
Although scale-out products have been available for some time, these products were typically subject to one or more of the following shortcomings:

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

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

As is depicted in Figure 5, NetApp ONTAP is the first product to offer a complete scale-out solution with an adaptable, always-available storage infrastructure for today's highly virtualized environment. An ONTAP system can scale up to 24 nodes, depending on platform and protocol, and can contain different disk types and controller models in the same storage cluster.

Figure 5) ONTAP.



## Nondisruptive Operations

The move to shared infrastructure has made it nearly impossible to schedule downtime for routine maintenance. ONTAP is designed to eliminate the planned downtime needed for maintenance operations and lifecycle operations as well as 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 allow you to nondisruptively perform a wide 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—including CPUs, cache memory, network I/O bandwidth, and disk I/O bandwidth—can be easily kept in balance. NetApp clustered Data ONTAP 8.3.1 enables you to perform the following tasks:

- 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 can provide services to thousands of VMs. In such environments, downtime is not an option. The NetApp AFF solution eliminates sources of downtime and protects critical data against disaster with two key features:

- **High availability.** A NetApp HA pair provides seamless failover to its partner in the case of hardware failure. Each of the two identical storage controllers in an HA pair 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 server deployment, data protection is critical because RAID failures can affect hundreds of servers, resulting in lost productivity. RAID DP provides performance comparable to that of RAID 10, yet it requires fewer disks to achieve equivalent protection. RAID DP provides protection against double disk failure, in contrast to RAID 5, which can protect against only one disk failure per RAID group. RAID DP in effect provides RAID 10 performance and protection at a RAID 5 price point.

## NetApp Advanced Data Management Capabilities

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

### Storage Efficiencies

Storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. The following NetApp storage efficiency technologies can help you to realize maximum space savings:

- **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 wear on SSDs.
- **Inline zero elimination and always-on deduplication.** Data deduplication cuts storage requirements by reducing redundancies in primary, backup, and archival data. Inline deduplication of zeros speeds up VM provisioning by 20% to 30%. Combined with always-on deduplication running at all times, this deduplication method provides more space savings than postprocess deduplication.
- **Snapshot technology.** NetApp Snapshot technology provides low-cost, instantaneous, point-in-time copies of the file system (volume) or LUN by preserving ONTAP architecture and WAFL consistency points without affecting performance. NetApp SnapCenter integrates with the SQL Server Virtual Device Interface to create application-consistent Snapshot copies of production-level SQL Server databases with no downtime for the production database.
- **Thin provisioning.** Thin provisioning, implemented by NetApp at the NetApp FlexVol® volume level and at the LUN level, defers storage purchases by keeping a common pool of free storage available to all applications.

- **Thin replication.** Thin replication is at the center of the NetApp data protection software portfolio, which includes NetApp SnapMirror® and NetApp SnapVault® software. SnapVault thin replication enables more frequent backups that use less storage capacity because no redundant data is moved or stored. SnapMirror thin replication protects business-critical data while minimizing storage capacity requirements.
- **RAID DP.** RAID DP technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead.
- **FlexClone volumes.** FlexClone virtual cloning reduces the need for storage by creating multiple, instant, space-efficient, writable copies.

## Advanced Storage Features

NetApp ONTAP provides a number of additional features leverageable in a virtual SQL Server environment whether for the infrastructure supporting the database servers or for the database servers themselves. These features include the following:

- **NetApp Snapshot copies.** Manual or automatically scheduled point-in-time copies that write only changed blocks, with no performance penalty. Snapshot copies consume minimal storage space because only changes to the active file system are written. Individual files and directories can easily be recovered from any Snapshot copy, and the entire volume can be restored back to any Snapshot state in seconds.
- **Compression.** Compression of data blocks on disk to provide space savings instead of or in addition to savings obtained with deduplication.
- **LIF.** A logical interface that is associated with a physical port, an interface group (ifgrp), or a VLAN interface. More than one LIF can be associated with a physical port at the same time. There are three types of LIFs: NFS LIFs, iSCSI LIFs, and FC LIFs.  
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 (WWPNs) 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 and Asymmetric Logical Unit Access (ALUA) to notify clients of any changes in the network topology.
- **Storage virtual machines (SVMs).** An SVM is a secure virtual storage server that contains data volumes and one or more LIFs through which it serves data to 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:

- 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 deployed system

ONTAP can support several protocols concurrently in the same storage system. Unified storage is important to all VMware vSphere solutions, such as CIFS/SMB for user data, NFS or SAN for the VM datastores, and guest-connect iSCSI LUNs for Windows applications.

The following protocols are supported:

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

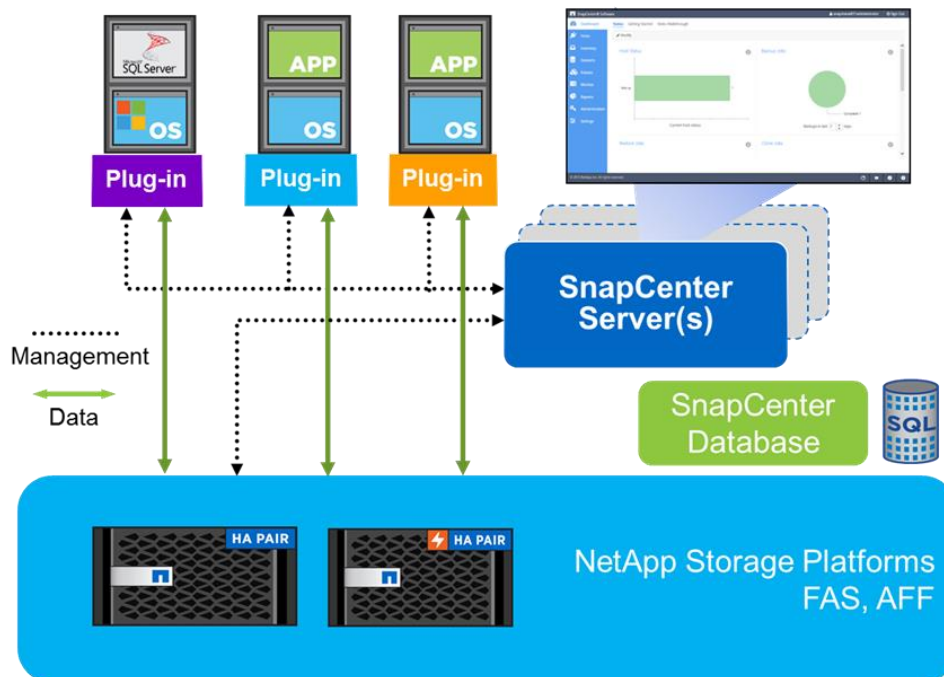
## NetApp SnapCenter

NetApp SnapCenter software is a unified, scalable platform for application-consistent data protection and clone management. SnapCenter simplifies backup, restore, and clone lifecycle management with application-integrated workflows. With storage-based data management, SnapCenter enables increased performance and availability and reduced testing and development times.

### Simple

NetApp SnapCenter includes both the SnapCenter Server and individual lightweight application, database, and OS plug-ins, which are all controlled from a central management console. The management console delivers a consistent user experience across all applications or databases. It incorporates a single GUI to support critical functions, such as job monitoring, event notification, logging, dashboard, reporting, scheduling, RBAC and policy management for all application and database plug-ins. Figure 6 depicts the SnapCenter architecture.

Figure 6) SnapCenter architecture.



SnapCenter Server also includes Snapshot catalog management to facilitate easy rollback to point-in-time copies. SnapCenter Server checks application, database, and OS interoperability and then

nondisruptively installs and upgrades software plug-ins on application and database hosts. Those plug-ins can then be managed from the central management console.

In addition, SnapCenter Server allows you to run custom scripts either before or after common operations such as backup, cloning, and restore by using Perl, Python, or PowerShell. NetApp SnapCenter also includes an intuitive migration mechanism for customers who are currently using NetApp SnapManager for Microsoft SQL Server.

## **Scalable**

SnapCenter is designed with ease of use in mind, with the added ability to scale capacity and performance to meet the needs of large enterprises. You can transparently add SnapCenter Servers to address requirements for HA and load balancing, with support for thousands of applications and databases. By adding another SnapCenter Server or multiple servers, you can protect against any one server failing. Therefore, you can add multiple servers to increase resiliency, and they are all managed as a single server. The added servers also increase the level of performance for your backup infrastructure because performance is transparently balanced across servers.

Backup and restore performance is also increased by leveraging the onboard capabilities of NetApp storage-based Snapshot copies. Offloading this functionality not only simplifies operation, but also offloads Snapshot functions from the host.

By leveraging the embedded functionality of the NetApp ONTAP platform to perform space-efficient FlexClone management, SnapCenter also improves the performance of testing and development. Application, database, and virtual infrastructure administrators can initiate FlexClone volumes independently of storage administrators through the same GUI console. The self-service feature of space-efficient cloning reduces testing and development time and puts more capability into the hands of application owners.

## **Empowering**

IT organizations face the challenge of providing self-service capabilities to individual administrators while also retaining oversight and control of the storage infrastructure by the storage administrator. SnapCenter uses role-based access control to delegate functionality to application and database owners while retaining oversight and control by a central storage infrastructure administrator. This level of control and security frees storage administrators from tedious tasks that application and database owners can do for themselves. At the same time, it protects the overall infrastructure from bullying applications or from infrastructure abuse from even the best-intended colleagues.

As IT organizations continue to grow with the size of the overall business, IT specialists play an important role in the data center. SnapCenter provides application-specific or database-specific workflows tailored to meet the needs of application, database, and virtualization infrastructure administrators. Because each application or database has a unique workflow, application and database owners should find that their delegated workflows are familiar and well suited to their use models.

Administrators can use the SnapCenter plug-ins for applications and databases so that the application or database is consistent at all levels, which promotes maximum recoverability. Plug-ins for SnapCenter allow a variety of restore capabilities. They can roll forward logs and enable application or database administrators to clone or recover to the latest information available or to a specific point in time.

SnapCenter also leverages NetApp storage-based backup and replication functions, such as SnapVault and SnapMirror. All SnapCenter plug-ins can perform cloning and restore operations from both primary and secondary locations. SnapCenter integration with NetApp Virtual Storage Console also enables seamless backup and recovery of Microsoft SQL databases in vSphere environments. This integration enables VSC to create and reuse policy objects across the multiple vCenters that are connected to SnapCenter, and it enables VSC to use the SnapCenter database. Thus VSC can scale significantly better. Integration also allows VSC to leverage the SnapCenter Snapshot catalog, enabling new features such as restore from SnapVault.

## VMware vSphere Integrations

The complexity of deploying and managing thousands of VMs can be daunting without the right tools. NetApp Virtual Storage Console (VSC) for VMware vSphere is tightly integrated with VMware vCenter for rapidly provisioning, managing, configuring, and backing up virtual environments. NetApp VSC significantly increases operational efficiency and agility by simplifying the deployment and management processes for thousands of VMs.

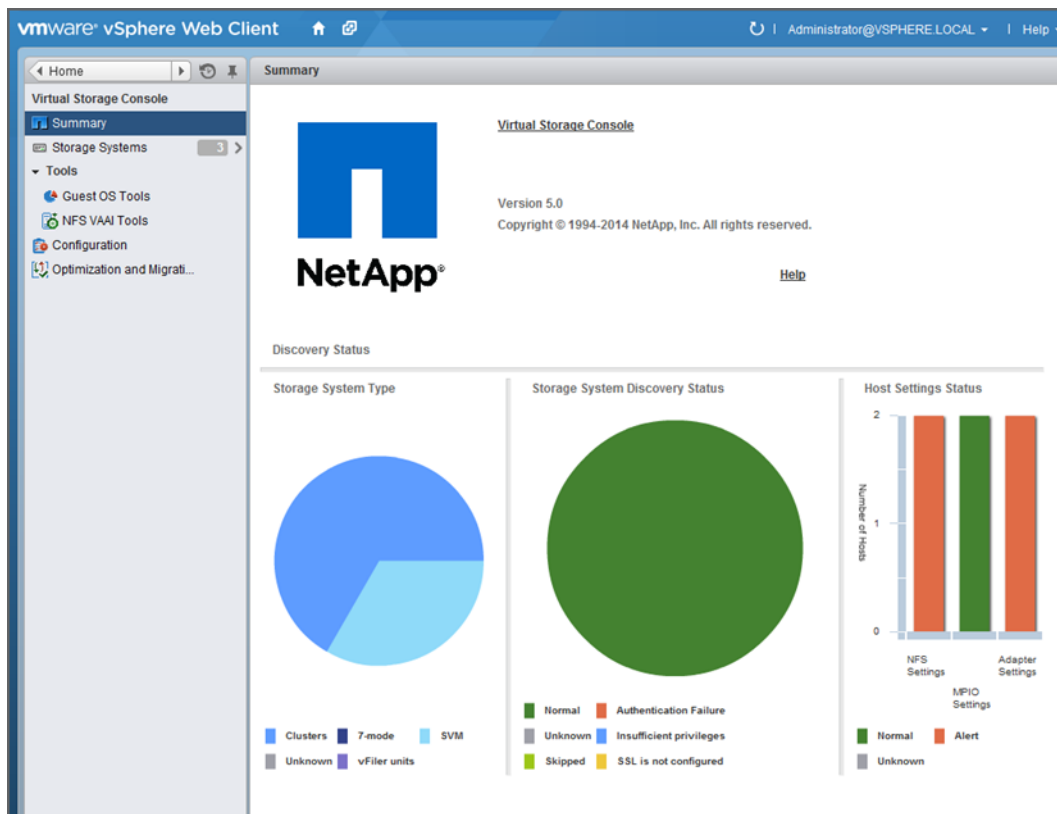
The following plug-ins and software features simplify deployment and administration of vSphere environments:

- **NetApp VSC Provisioning and Cloning plug-in.** Enables you to rapidly provision, manage, import, and reclaim space of thinly provisioned VMs and redeploy thousands of VMs.
- **NetApp VSC Backup and Recovery plug-in.** Integrates VMware snapshot functionality with NetApp Snapshot functionality to protect virtual SQL Server 2014 environments.

The NetApp VSC delivers storage configuration and monitoring, datastore provisioning, VM cloning, and backup and recovery of VMs and datastores. NetApp VSC also includes an API for automated control and a single VMware plug-in that provides end-to-end VM lifecycle management for VMware environments by using NetApp storage. NetApp VSC is delivered as a VMware vCenter Server plug-in, which is different from a client-side plug-in that must be installed on every VMware vSphere Client. The NetApp VSC can be installed on a separate Microsoft Windows Server instance or VM or as a standalone vApp.

Figure 7 depicts the NetApp VSC summary page.

Figure 7) The NetApp Virtual Storage Console.



## Replication

The NetApp Backup and Recovery Plug-In for NetApp VSC is a scalable, integrated data protection solution for vSphere environments. The backup and recovery plug-in allows you to leverage VMware snapshot functionality with NetApp array-based, block-level Snapshot copies to provide consistent backups for the VMs. The backup and recovery plug-in is integrated with NetApp SnapMirror replication technology, which preserves the deduplicated storage savings from the source to the destination storage array. Therefore, deduplication does not need to be rerun on the destination storage array.

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

### 3.5 VMware vSphere

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

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

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

- **Availability.** Workload mobility is provided through vMotion. HA provided by vSphere fault domain manager technology offers VM resiliency in the event of physical server or guest OS failures.
- **Automation.** VMware Distributed Resource Scheduler (DRS) offers dynamic workload distribution to align resource utilization with business priorities and compute capacity. DRS uses compute resources efficiently and thus improves power consumption.
- **Compute.** The VMware vSphere ESXi hypervisor provides efficient memory, storage, and compute abstraction through the use of VMs.
- **Network.** VMware vSphere supports third-party virtual distributed switches such as the Cisco Nexus 1000v, providing a resilient and fully integrated virtualized network access layer.
- **Storage.** Thin provisioning allows overprovisioning of storage resources to improve storage utilization and improve capacity planning. MFS is a clustered file system that grants multiple hosts simultaneous read and write access to a single volume located on a SCSI-based device through FC, FCoE, or iSCSI. VMFS-5 supports a maximum of 32 hosts connected to a single volume of up to 64TB in size.

Figure 8 provides an overview of the capabilities of VMware vSphere.



Figure 8) VMware vSphere feature overview.



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

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

### 3.6 Microsoft SQL Server 2014

SQL Server is the foundation of Microsoft's data platform, delivering mission-critical performance with in-memory technologies and faster insights on any data, whether on premises or in the cloud. Microsoft SQL Server 2014 builds on the mission-critical capabilities delivered in prior releases by providing breakthrough performance, availability, and manageability for mission-critical applications. This release of the SQL Server database engine introduces features and enhancements that increase the power and productivity of architects, developers, and administrators who design, develop, and maintain data storage systems.

## AlwaysOn Availability Groups

AlwaysOn Availability Groups are an HA and DR solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012, AlwaysOn Availability Groups maximize the availability of a set of user databases for an enterprise. An availability group (AG) supports a failover environment for a discrete set of user databases (known as availability databases) that fail over together. An availability group supports a set of read-write primary databases and one to eight sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations.

An availability group fails over at the level of an availability replica. Failover of the availability group occurs in the event of hardware or software failure that causes the SQL Server service on the active replica to go offline. AG failovers are not caused by database issues such as a database becoming suspect due to a loss of a data file, deletion of a database, or corruption of a transaction log.

AlwaysOn Availability Groups provide a rich set of options that improve database availability and that enable improved resource use. The key components are as follows:

- Support for up to nine availability replicas. An availability replica is an instantiation of an availability group that is hosted by a specific instance of SQL Server. It maintains a local copy of each availability database that belongs to the availability group. Each availability group supports one primary replica and up to eight secondary replicas. For more information, see [Overview of AlwaysOn Availability Groups \(SQL Server\)](#).
- Support for alternative availability modes. As follows:
  - Asynchronous-commit mode. This availability mode is a DR solution that works well when the availability replicas are distributed over considerable distances.
  - Synchronous-commit mode. This availability mode emphasizes high availability and data protection over performance at the cost of increased transaction latency. A given availability group can support up to three synchronous-commit availability replicas, including the current primary replica.
- Support for several forms of availability group failover. Automatic failover, planned manual failover (often referred to as manual failover), and forced manual failover (often referred to as forced failover) are supported. For more information, see [Failover and Failover Modes \(AlwaysOn Availability Groups\)](#).
- Configuration of a given availability replica to support either or both of the following active-secondary capabilities:
  - Read-only connection access, which enables read-only connections to a replica to access and read its databases when it is running as a secondary replica. For more information, see [Active Secondaries: Readable Secondary Replicas \(AlwaysOn Availability Groups\)](#).
  - Performing backup operations on a replica's databases when that replica is running as a secondary. For more information, see [Active Secondaries: Backup on Secondary Replicas \(AlwaysOn Availability Groups\)](#).

Active secondary capabilities improve IT efficiency and reduce cost through better resource utilization of secondary hardware. In addition, offloading read-intensive applications and backup jobs to secondary replicas helps to improve performance on the primary replica. Active secondaries provide the following benefits:

- Support an availability group listener for each availability group. An availability group listener is a server name to which clients can connect to access a database in a primary or secondary replica of an AlwaysOn Availability Group. Availability group listeners direct incoming connections to the primary replica or to a read-only secondary replica. The listener provides fast application failover after an availability group fails over. For more information, see [Availability Group Listeners, Client Connectivity, and Application Failover \(SQL Server\)](#).
- Support a flexible failover policy for greater control over availability group failover.

- Support automatic page repair for protection against page corruption. For more information, see [Automatic Page Repair \(Availability Groups/Database Mirroring\)](#).
- Support encryption and compression, which provide secure, high-performance transport.
- Provide an integrated set of tools to simplify deployment and management of availability groups, including the following:
  - Transact-SQL Server DDL statements for creating and managing availability groups
  - SQL Server Management Studio tools
  - The AlwaysOn Dashboard, which monitors AlwaysOn Availability Groups, availability replicas, and availability databases and evaluates results for AlwaysOn policies
  - The Object Explorer Details pane, which displays basic information about existing availability groups
  - PowerShell cmdlets

## 4 Technology Requirements

### 4.1 Hardware Components

Table 1 lists the hardware components used to validate the solution. 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	2x Cisco UCS 6248UP fabric interconnects Includes Cisco UCS Manager
Cisco UCS B200 M4	2x Xeon E5-2690 CPU (12 cores/each 2.6Ghz) 128GB RAM/blade, 1 VIC 1340/blade
Cisco UCS 5108 chassis	Includes 2x Cisco UCS-IOM 2204XP
Cisco Nexus 9396PX	No expansion modules
NetApp AFF8080 EX	No additional PCI cards
NetApp DS2246 disk shelves	2x disk shelves with 800GB SSDs

### 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
<b>Compute</b>	
Cisco UCS Manager	2.2(5d)
<b>Networking</b>	
Cisco Nexus 9396PX	NX-OS software release 7.0(3)1(3)
<b>Storage</b>	

Software/Firmware	Version
NetApp clustered Data ONTAP	8.3.1
NetApp Windows PowerShell toolkit	4.1.0
NetApp System Manager	8.3.1
NetApp VSC	6.2
NetApp SnapCenter Server	1.1
SnapCenter Plug-In for Microsoft Windows	1.1
SnapCenter Plug-In for Microsoft SQL Server	1.1
<b>VMware vSphere</b>	
VMware ESXi	6.0.0 3380124
VMware vCenter Server	6.0.0 2656761
VMware vSphere PowerCLI	6.0 Release 3 build 3205540
<b>Database Server</b>	
Microsoft Windows Server	2012 R2 Standard
Microsoft SQL Server	SQL Server 2014 Service Pack 1 Enterprise Edition

## 5 Solution Design

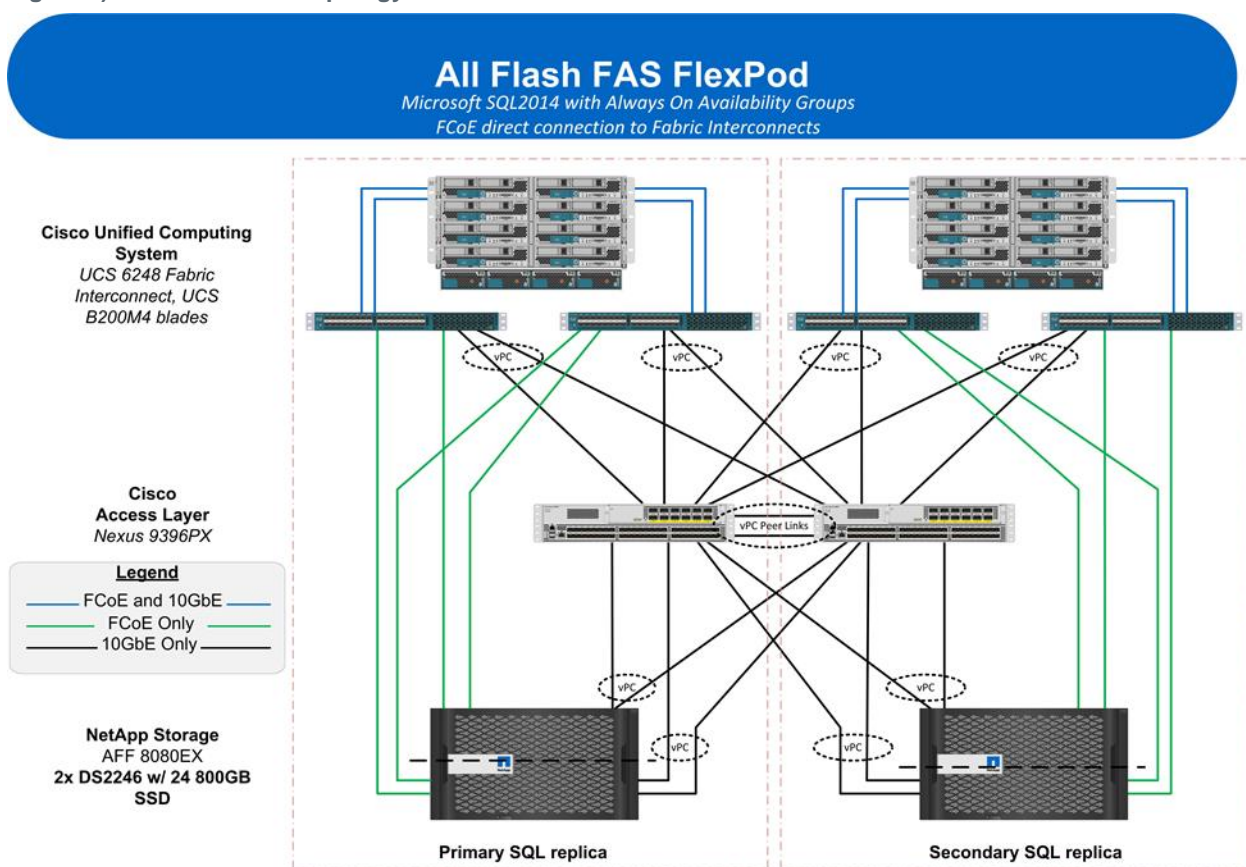
The FlexPod Datacenter with SQL Server 2014 and AFF solution contains the following elements:

- Cisco Nexus network switches
- Cisco UCS
- NetApp AFF storage
- NetApp SnapCenter
- VMware vSphere

### 5.1 Cisco Nexus Network Design

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

Figure 9) Overall solution topology.



## Network Design Overview

As seen in Figure 9, Cisco Nexus 9396PX switches serve as the access layer for the primary and secondary database environments. The Cisco UCS fabric interconnects and storage systems are connected to the Cisco Nexus 9000 access switches with virtual port channels for maximum availability. FCoE connectivity is provided by direct links between the storage controllers and fabric interconnects. For this validation, a single pair of Cisco Nexus switches was used, but the primary and secondary database environments could be connected to separate Cisco Nexus pairs for additional resiliency.

## Network Switching

Two Cisco Nexus 9396PX switches running NX-OS software release 7.0(3)I1(3) were used in this solution design. These switches were chosen because of their 10/40GbE switching capability and their ability to participate in Cisco ACI networks. Other Cisco Nexus 9000 switches can be used depending on port count requirements, or additional pairs could be added to extend the pod even further.

The switches are configured as virtual port channel peers. Virtual port channels are used to provide switch-level redundancy to the Cisco UCS fabric interconnects and AFF storage systems without requiring special configuration on those devices. The switches in this solution are operating in NX-OS mode but could also be configured as leaves in an ACI network.

Although Cisco Nexus 9000 series switches do not support FC or FCoE, FlexPod systems using these switches can still support these protocols. FCoE links from the storage array are directly connected to the Cisco UCS fabric interconnects operating in FC switch mode. Ports from each controller are connected to

each fabric interconnect, and zoning is performed on the fabric interconnects to provide an industry-standard, dual-fabric SAN topology.

## Host Server Networking

As seen in [Figure 11](#), the service profile for each Cisco UCS B200-M4 blade is configured with two virtual FC ports and two virtual Ethernet ports. The blades use FCoE to boot from the AFF8080 EX storage array as well as accessing LUNs for VM and SQL Server database storage. For Ethernet access, the ESXi hosts are configured with one standard vSwitch with two uplink ports that use originating port ID load balancing.

## Storage Networking

All storage traffic in this reference architecture uses FCoE. Each controller is connected with one FCoE link directly to each fabric interconnect, providing four paths from the hosts to the storage. FCoE ports operate independently, with ALUA providing multipathing and load balancing, and initiator groups (igroups) on the storage system mapping the appropriate LUNs to each ESXi host. FCoE bandwidth can be increased in this topology by adding additional FCoE links between the controllers and fabric interconnects.

For Ethernet traffic, each controller is connected to both Cisco Nexus 9396PX switches by using two ports in a multimode LACP interface group on the controller and virtual port channels across both switches. FlexPod specifications require Ethernet connectivity for storage even if the primary storage protocol is FC/FCoE to provide flexibility for the overall architecture. For example, during this validation, temporary NFS datastores were utilized to store ISO images and other static data files due to the ease of use and the flexibility of the connections.

## 5.2 Cisco Unified Computing System Design

The FlexPod design simultaneously supports both B-Series and C-Series deployments. This section of the document discusses only the integration and design of B-Series deployments into FlexPod.

### Cisco Unified Computing System: B-Series Server and Fabric Interconnect Design

The Cisco UCS supports the virtual server environment by providing a robust, highly available, and extremely manageable compute resource. In this solution, Cisco UCS 6248UP fabric interconnects are the foundation of the Cisco UCS. Cisco UCS 5108 blade chassis with 2204XP IOM-FEX modules are used to support B200-M4 blades with VIC 1340 adapters. Each I/O module (IOM) is connected to its respective fabric interconnect with 4 10GbE unified fabric links. The fabric interconnects are connected to the Cisco Nexus 9396PX switches for Ethernet connectivity and directly to the NetApp AFF8080 EX storage arrays by using FCoE for storage access. Figure 9 shows the Ethernet and FCoE connectivity of both the primary and secondary database replica environments.

### Cisco UCS Fabric Interconnect Connectivity

As seen in Figure 9, the FlexPod FCoE direct-attach topology supports any storage protocol a customer might want to use. As is typical of all FlexPod architectures, fabric interconnects are connected to the FlexPod Ethernet switches with port channels and virtual port channels. On each fabric interconnect, two 10GbE ports are configured as uplink ports and port channels.

One port is connected to each 9296PX 10GbE switch. The switch ports are configured as virtual port channels, providing both link-level and switch-level redundancy for traffic to each fabric interconnect. Additional 10GbE ports can be added to the port channel uplinks to increase network bandwidth to the blades as needed.



For SAN connectivity, we connected the fabric interconnects directly to the AFF8080 EX storage array. For this validation, each fabric interconnect has one 10Gb FCoE link to each storage controller. We created a dual-fabric SAN topology because there are no data links between the fabric interconnects, and the vHBAs configured on each blade map to one FI or the other. As with network bandwidth, FCoE bandwidth can be increased by adding additional links, although the four FCoE links in use were not used significantly during the workload testing of this architecture.

To support the FCoE storage port type, you must configure the fabric interconnects for FC switch mode, which requires a reboot of the fabric interconnect to take effect. FC switch mode enables the fabric interconnect to perform some FC switch functions, such as name services and zoning. After the fabric interconnect is in FC switch mode, two ports on each fabric interconnect are configured as FCoE storage ports, and each fabric interconnect is connected to each controller. This effectively creates an industry-standard, dual-fabric SAN topology that uses the fabric interconnect as FC switches. Additional links to each storage controller can be added to increase the available bandwidth. Zoning in this configuration is managed through the storage connection policies created in Cisco UCS Manager and applied to the service profiles and/or templates. These policies define the set of storage target ports to which a given vHBA is zoned.

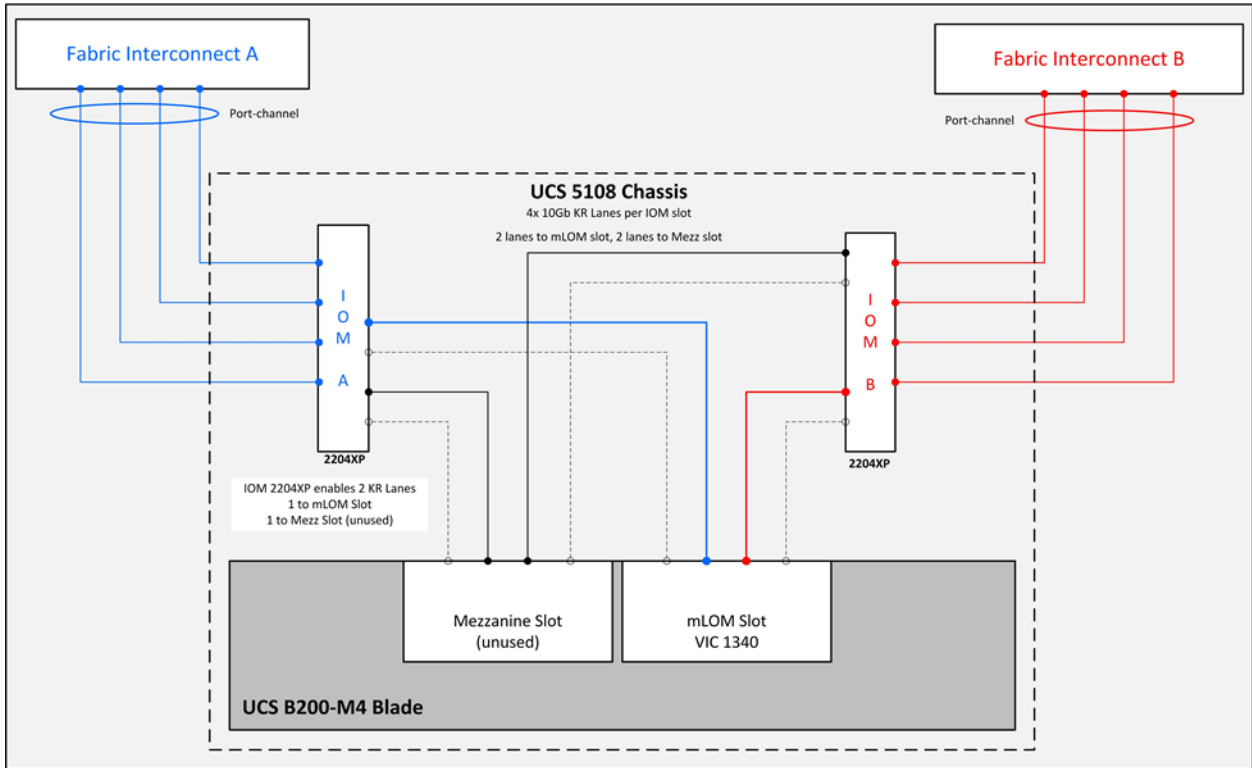
A balanced and predictable fabric is critical within any data center environment. As designed, the FlexPod system accommodates a myriad of traffic types (vMotion, NFS, FCoE, control traffic, and so on) and is capable of absorbing traffic spikes and protecting against traffic loss. Cisco UCS and Cisco Nexus QoS system classes and policies deliver this functionality. In this solution verification effort, the FlexPod system was configured to support jumbo frames with an MTU size of 9000. Enabling jumbo frames allows the FlexPod environment to optimize throughput between devices while simultaneously reducing the consumption of CPU resources. This class was assigned to the best effort class. With regard to jumbo frames, you must make sure that MTU settings are applied uniformly across the stack to prevent fragmentation and negative performance implications that inconsistent MTUs can introduce.

## Cisco UCS 5108 Chassis Connectivity

FlexPod allows organizations to adjust the individual components of the system to meet their particular scale or performance requirements. One key design decision in the Cisco UCS domain is the selection of I/O components. There are numerous combinations of I/O adapter, Cisco UCS Extenders I/O module (IOM), and Cisco UCS fabric interconnect available. Therefore, you should understand the effect that these selections have on the overall flexibility, scalability, and resiliency of the fabric.

Figure 10 illustrates the available backplane connections in the Cisco UCS 5100 series chassis as tested in this solution. Each of the two Cisco UCS 2204XP Series fabric extenders installed in each blade chassis has four 10GbE, FCoE-capable, SFP+ ports that connect the blade chassis to the fabric interconnect. Each 2204XP fabric extender also has four 10GBASE KR (802.3ap) standardized Ethernet backplane paths available for connection to each half-width blade slot. Two paths go to the mLOM slot, and two go to the mezzanine card slot. This solution was tested with 2204XP FEX modules, VIC 1340 mLOM cards without expansion ports, and no mezzanine cards. Each blade was provisioned with 20Gb of available bandwidth shared between Ethernet and FCoE traffic. To increase the bandwidth available to each blade, the 2208XP IOM can be installed along with other Cisco UCS adapter card options to increase the available bandwidth to 80Gb per half-width blade.

Figure 10) Cisco UCS backplane configurations: VIC 1340.

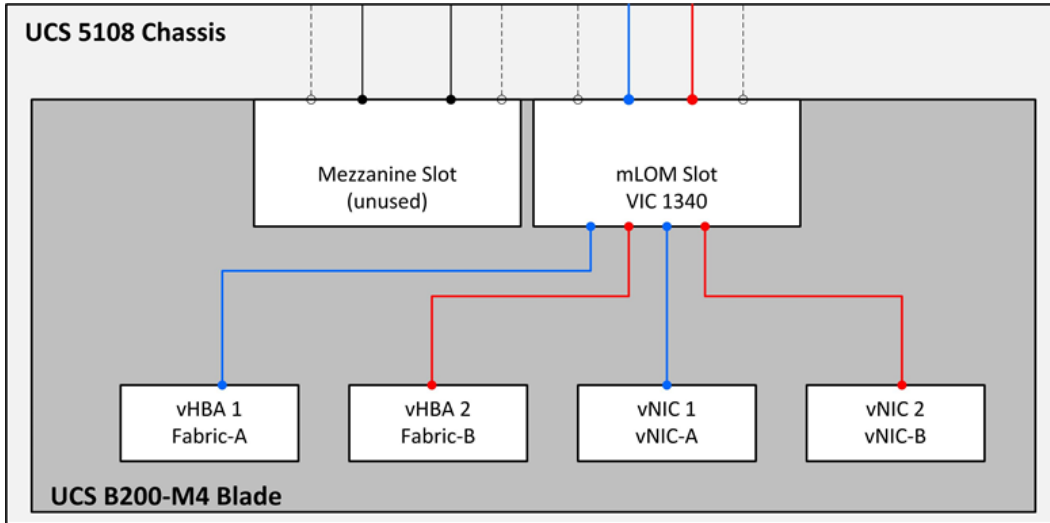


### Cisco UCS B200-M4 Blade Connectivity

For this solution, B200-M4 blade servers with VIC 1340 adapters are SAN booted as VMware ESXi hosts with FCoE. As illustrated in Figure 11, the Cisco 1340 VICs in the blades present four virtual PCIe devices to the ESXi node, two virtual 10GbE NICs (vNICs), and two vHBAs. The vSphere ESXi OS identifies these as VMNICs and VMHBAs respectively and is unaware that these are virtual adapters. Therefore, the ESXi server appears to have two LAN connections and two SAN connections, as would be expected for any typical server hardware configuration. In this case, one NIC and one HBA are sharing 10Gb of bandwidth through each fabric interconnect. Unless the appropriate components are included for additional bandwidth, as noted earlier, adding additional virtual Ethernet or FC adapters to the service profile does not increase the bandwidth available to each blade.



Figure 11) Server logic configuration.



### 5.3 NetApp AFF Storage Design

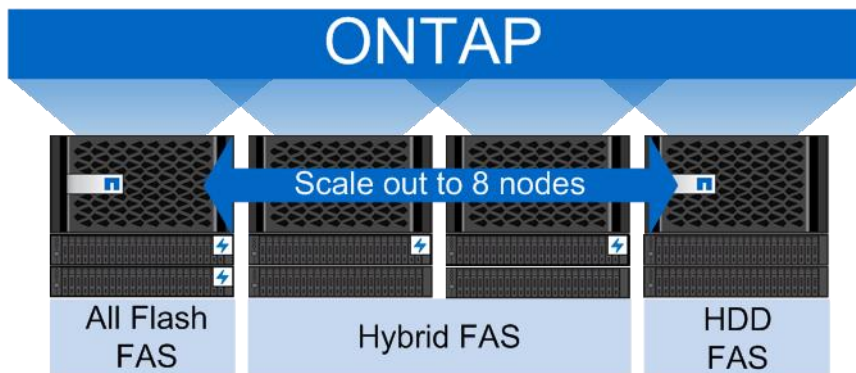
This section provides an overview of the NetApp FAS storage design for this reference architecture.

#### Storage Scale Considerations

ONTAP and the NetApp AFF8000 series of storage controllers allow your system to grow from a single-use case workload such as a small SQL Server deployment to a large-scale deployment for either large single workloads or multiple workloads. Individual controller models provide different levels of performance and capacity, as listed in Table 3. Any of these controllers can be mixed within the same cluster to meet the capacity and performance requirements of your business to provide cost efficiency during acquisition.

In SAN-only environments or mixed SAN and NAS environments, a single ONTAP cluster can scale to eight nodes or four HA pairs, as shown in Figure 12. At the high end, this configuration can support approximately 23PB of data within the same management plane. Mixed SAN and NAS environments support any combination of storage protocols within the same cluster (FC, FCoE, iSCSI, NFS, or CIFS/SMB) and therefore support all business data and application requirements.


Figure 12) ONTAP in a SAN-only environment.



In NAS-only environments, a single ONTAP cluster can scale up to 24 nodes or 12 HA pairs, as shown in Figure 13. At the high end, this configuration supports approximately 69PB of data within the same

management plane. NAS environments can take advantage of NFS and CIFS/SMB storage protocols, providing support for both business file data and virtualization data. NFS for VMware vSphere, OpenStack, Red Hat enterprise virtualization, and Xen Server and SMB3 for Microsoft Hyper-V are supported.

Figure 13) ONTAP in a NAS-only environment.



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

In this reference architecture, separate clusters were used for the primary and secondary database environments. This provides the highest level of performance for both copies of the database and supports the highest level of resiliency for the overall solution. Both environments support additional controllers and/or workloads, and each cluster can scale independently as workload requirements demand.

## NetApp AFF8000 Technical Specifications

Table 3 provides the technical specifications for the four NetApp AFF series storage controllers: AFF8020, AFF8040, AFF8060, and AFF8080 EX.

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

Table 3) NetApp AFF8000 storage system technical specifications.

	AFF8080 EX	AFF8060	AFF8040	AFF8020
Maximum SSD	240	240	240	240
Maximum raw capacity: all flash	384TB/349TiB	384TB/349TiB	384TB/349TiB	384TB/349TiB
Effective capacity	1565.3TB/1423TiB	1565.3TB/1423TiB	1565.3TB/1423TiB	1565.3TB/1423TiB
Controller form factor	<ul style="list-style-type: none"> <li>Dual-enclosure HA</li> <li>2 controllers and 2 IOXMs in two 6U chassis</li> <li>Total of 12U or single-enclosure HA</li> <li>2 controllers in single 6U chassis</li> </ul>	<ul style="list-style-type: none"> <li>Dual-enclosure HA</li> <li>2 controllers and 2 IOXMs in two 6U chassis</li> <li>Total of 12U or single-enclosure HA</li> <li>2 controllers in single 6U chassis</li> </ul>	<ul style="list-style-type: none"> <li>Single-enclosure HA</li> <li>2 controllers in single 6U chassis</li> </ul>	<ul style="list-style-type: none"> <li>Single-enclosure HA</li> <li>2 controllers in single 3U chassis</li> </ul>

	AFF8080 EX	AFF8060	AFF8040	AFF8020
Memory	256GB	128GB	64GB	48GB
NVRAM	32GB	16GB	16GB	8GB
PCIe expansion slots	6 or 24	8 or 24	8	4
Onboard I/O: UTA2 (10GbE/FCoE, 16Gb FC)	8	8	8	4
Onboard I/O: 10GbE	8	8	8	4
Onboard I/O: GbE	8	8	8	4
Onboard I/O: 6Gb SAS	8	8	8	4
Storage networking supported	FC, FCoE, iSCSI, NFS, pNFS, and CIFS/SMB			
OS version	Data ONTAP 8.3 or later; Data ONTAP 8.3.1 or later for AFF8080 EX single-chassis HA			

## Storage Network Connectivity

NetApp AFF systems with ONTAP can support almost any storage protocol in use today. FlexPod architectures support this flexibility throughout the infrastructure; therefore, you can choose the option most appropriate for your application or environment.


For this solution, primary storage access is provided by FCoE. UTA2 ports on each storage controller are connected directly to each Cisco UCS fabric interconnect. The fabric interconnects operate in FC switch mode with no connectivity between them in a typical dual-fabric SAN topology. Zoning and name services are provided by the fabric interconnect to control FC host access to the storage targets. While this topology only supports access by hosts connected to the attached fabric interconnects, you can connect additional ports to other Cisco UCS fabric interconnects or other FC SAN switches. Alternatively, additional ports can be added to the existing fabric interconnects to increase the available FCoE bandwidth.

Although FCoE provides primary storage access in this solution, each storage controller is also connected to the Cisco Nexus 9396PX network switches. Two ports on each controller are configured as an LACP ifgroup, with one port connected to each switch. The switch ports are configured as virtual port channels to provide switch-level resiliency on top of the link-level redundancy provided by the ifgroup. This network uplink can be used to support NFS datastores for the vSphere environment, user home directories, or other CIFS shares or iSCSI access for other servers or applications. As with FCoE, additional 10GbE links can be added to the ifgroup to increase the available network bandwidth as required.

## Back-End Storage Connectivity Overview

Figure 14 depicts the connectivity between the AFF8080 EX storage controllers and the DS2246 disk shelves for both primary and secondary storage systems. Each AFF8080 EX controller is 6RU, and each DS2246 is 2RU, so the primary and secondary systems occupy 16RU each.

Figure 14) Multipath HA to DS2246 SSD shelves.



### Disk Assignment

To maximize the SSD performance on AFF systems, the disks should be assigned to the controllers in a way that is different from the default ONTAP disk assignment. As seen in Figure 14, each controller has four SAS loops connected to the disk shelves. The default behavior of ONTAP is to assign an entire disk shelf to each controller, which provides two SAS connections from the active controller to 24 disks (A and D to 24 disks in shelf 1). By assigning half of the disks in each shelf to each controller, each controller has four SAS connections to 24 disks (A and D to 12 disks in shelf 1, B and C to 12 disks in shelf 2), providing additional bandwidth between the disks and controllers.

### Aggregate, LUN, and Volume Configuration

AFF storage systems support NetApp advanced disk partitioning, which allows drives to be partitioned and then shared between root aggregates and data aggregates. This avoids duplication of parity drives and maximizes the usable capacity of the available SSDs. This is done by dividing the disks into two partitions: a small partition used to build root aggregate RAID groups and a larger partition used to build data aggregate RAID groups. In this configuration with 24 SSD per controller, 22 data partitions are used to build one data aggregate per controller, with two spare partitions on each. Table 4 shows the storage aggregate configuration used in this solution.

Table 4) Storage aggregate configuration.

Storage	Aggregate Name	Volume Name	Vol Size (GB)	Description
sql-pri-01	aggr1_pri01			<ul style="list-style-type: none"> <li>Advanced drive partitioning</li> <li>20 data partitions + 2 parity partitions</li> <li>RAID DP total aggregate size = 12.7TB</li> <li>2 spare drives</li> </ul>
	aggr0_pri01	root	348	Total aggregate size = 368GB
Sql-pri-02	aggr1_pri02			<ul style="list-style-type: none"> <li>Advanced drive partitioning</li> <li>20 data partitions + 2 parity partitions</li> <li>RAID DP total aggregate size = 12.7TB</li> <li>2 spare drives</li> </ul>
	aggr0_pri02	root	348	Total aggregate size = 368GB

It is a NetApp best practice to create a dedicated FlexVol volume for each LUN, unless LUN data is expected to be fairly static and thus could benefit from volume-level storage efficiencies such as deduplication. For ESXi boot LUNs, we created a single volume with several LUNs, because the LUNs were small in size, did not change often, deduplicated well, and did not require granular Snapshot copies or replication.

For SQL data LUNs, individual volumes were created with a single LUN each, which were then configured as a vSphere VMFS datastore. This configuration allows granular management of Snapshot and replication schedules when using SnapCenter to back up and restore SQL Server databases. To support the required infrastructure, one volume with one LUN was created for each infrastructure service VM such as the vCenter Server, the Virtual Storage Console, SnapCenter, the Active Directory domain controller, and so on. Another dedicated LUN was created for production SQL Server OS virtual disks.

For each database, a separate volume and LUN were created for data, log, tempdb, and snapinfo. Databases were balanced across the controllers by placing odd-numbered SQL Servers on controller 1 and even-numbered SQL Servers on controller 2. Table 5 shows the volumes and LUNs that were used during testing of this solution.

**Table 5) Volume and LUN configuration.**

Aggregate Name	Volume Name	Vol Size (GB)	LUN Name	LUN Size (GB)	Mount Point/Description
aggr1_pri01	ucs_boot	500GB	SQL_pri01	40GB	ESXi boot LUN
aggr1_pri01	ucs_boot	500GB	SQL_pri02	40GB	ESXi boot LUN
aggr1_pri01	ucs_boot	500GB	SQL_pri03	40GB	ESXi boot LUN
aggr1_pri01	ucs_boot	500GB	SQL_pri04	40GB	ESXi boot LUN
aggr1_pri01	ucs_boot	500GB	infra_01	40GB	ESXi boot LUN
aggr1_pri01	ucs_boot	500GB	infra_02	40GB	ESXi boot LUN
aggr1_pri02	infra_vm_01	1TB	infra_vm_01	400GB	VMFS datastore: infrastructure VMs
aggr1_pri02	sql_vm_01	1.5TB	sql_vm_01	400GB	VMFS datastore: SQL Server VMs
aggr1_pri01	sql3_data	3TB	sql3_data	2.5TB	VMFS datastore: sql3_data
aggr1_pri01	sql3_log	500GB	sql3_log	800GB	VMFS datastore: sql3_log
aggr1_pri01	sql3_tempdb	300GB	sql3_tempdb	500GB	VMFS datastore: sql3_tempdb
aggr1_pri01	sql3_snapinfo	200GB	sql3_snapinfo	400GB	VMFS datastore: sql3_snapinfo
aggr1_pri01	sql5_data	3TB	sql5_data	2.5TB	VMFS datastore: sql5_data
aggr1_pri01	sql5_log	500GB	sql5_log	800GB	VMFS datastore: sql5_log
aggr1_pri01	sql5_tempdb	300GB	sql5_tempdb	500GB	VMFS datastore: sql5_tempdb

Aggregate Name	Volume Name	Vol Size (GB)	LUN Name	LUN Size (GB)	Mount Point/Description
aggr1_pri01	sql5_snapinfo	200GB	sql5_snapinfo	400GB	VMFS datastore: sql5_snapinfo
aggr1_pri01	sql7_data	3TB	sql7_data	2.5TB	VMFS datastore: sql7_data
aggr1_pri01	sql7_log	500GB	sql7_log	800GB	VMFS datastore: sql7_log
aggr1_pri01	sql7_tempdb	300GB	sql7_tempdb	500GB	VMFS datastore: sql7_tempdb
aggr1_pri01	sql7_snapinfo	200GB	sql7_snapinfo	400GB	VMFS datastore: sql7_snapinfo
aggr1_pri02	sql2_data	3TB	sql2_data	2.5TB	VMFS datastore: sql2_data
aggr1_pri02	sql2_log	500GB	sql2_log	800GB	VMFS datastore: sql2_log
aggr1_pri02	sql2_tempdb	300GB	sql2_tempdb	500GB	VMFS datastore: sql2_tempdb
aggr1_pri02	sql2_snapinfo	200GB	sql2_snapinfo	400GB	VMFS datastore: sql2_snapinfo
aggr1_pri02	sql4_data	3TB	sql4_data	2.5TB	VMFS datastore: sql4_data
aggr1_pri02	sql4_log	500GB	sql4_log	800GB	VMFS datastore: sql4_log
aggr1_pri02	sql4_tempdb	300GB	sql4_tempdb	500GB	VMFS datastore: sql4_tempdb
aggr1_pri02	sql4_snapinfo	200GB	sql4_snapinfo	400GB	VMFS datastore: sql4_snapinfo
aggr1_pri02	sql6_data	3TB	sql6_data	2.5TB	VMFS datastore: sql6_data
aggr1_pri02	sql6_log	500GB	sql6_log	800GB	VMFS datastore: sql6_log
aggr1_pri02	sql6_tempdb	300GB	sql6_tempdb	500GB	VMFS datastore: sql6_tempdb
aggr1_pri02	sql6_snapinfo	200GB	sql6_snapinfo	400GB	VMFS datastore: sql6_snapinfo

## NetApp Virtual Storage Console for VMware vSphere

Virtual Storage Console for VMware vSphere is a vCenter Server plug-in that provides end-to-end lifecycle management for VMs in VMware environments by using NetApp storage systems. VSC can be coinstalled on the VMware vCenter Server instance when the Windows version of vCenter is used. When using the VMware vCenter Server virtual appliance, a separate Windows server is required to host the NetApp VSC. Table 6 lists the NetApp VSC VM configurations.

Table 6) NetApp VSC VM configuration.

NetApp VSC	Configuration
OS	Microsoft Windows Server 2012 R2
Virtual CPU (vCPU)	2 vCPUs
Memory	4GB
Network adapter type	VMXNET3
Hard disk size	60GB
Hard disk type	Thin

This reference architecture uses the NetApp VSC for the following tasks:

- Setting NetApp best practices for ESXi hosts (timeout values, HBAs, multipath input/output, and NFS settings)
- Integration with NetApp SnapCenter for backup, recovery, and cloning of SQL Server VMs

NetApp VSC is required when using NetApp SnapCenter to back up, restore, and clone virtual servers and applications. VSC allows SnapCenter to communicate with VMware vSphere when SnapCenter performs backup and restore operations for SQL Server or Oracle databases on virtual machine disks (VMDKs) or raw device mappings (RDMs). In addition, VSC uses SnapCenter to perform backup and restore operations for storage systems running clustered Data ONTAP 8.2.2 or later. You must register VSC with SnapCenter by using either the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box.

## NetApp SnapCenter

SnapCenter consists of the SnapCenter Server and the SnapCenter Plug-In Package for Windows (which includes the SnapCenter Plug-In for Microsoft SQL Server and the SnapCenter Plug-In for Microsoft Windows). SnapCenter also contains the SnapCenter Plug-In Package for Linux (which includes the SnapCenter Plug-In for Oracle Database and the SnapCenter Plug-In for UNIX). SnapCenter interacts with Virtual Storage Console for VMware vSphere to provide support for database backup, restore, recovery, and cloning on RDMs and VMDKs. Figure 6 depicts the topology of SnapCenter application plug-ins.

### SnapCenter Server

The SnapCenter Server includes a web server, a centralized HTML5-based user interface, PowerShell cmdlets, APIs, and the SnapCenter repository. SnapCenter enables load balancing, high availability, and horizontal scaling across multiple SnapCenter Servers within a single user interface. You might need multiple SnapCenter Servers for high availability. You accomplish this by using network load balancing and application request routing. For larger environments with thousands of hosts, adding multiple SnapCenter Servers can help balance the load. The SnapCenter platform is based on a multitiered architecture that includes a centralized management server (the SnapCenter Server) and SnapCenter plug-ins.

SnapCenter also enables centralized application resource management and easy data protection job execution through the use of datasets and policy management, including scheduling and retention settings. SnapCenter provides unified reporting through the use of a dashboard, multiple reporting options, job monitoring, and log and event viewers. SnapCenter data protection capabilities can be delegated to application administrators using granular role-based access control. Information related to different operations performed from SnapCenter is stored in the SnapCenter repository.



## VSC for VMware vSphere

Virtual Storage Console for VMware vSphere is a vCenter Server plug-in that provides end-to-end lifecycle management for VMs in VMware environments using NetApp storage systems. VSC enables SnapCenter to communicate with VMware vSphere when SnapCenter performs backup and restore operations for SQL Server or Oracle databases on VMDKs or RDMs. In addition, VSC uses SnapCenter to perform backup and restore operations for storage systems running clustered Data ONTAP 8.2.2 or later. You must register VSC with SnapCenter by using either the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box. VSC is not included in the SnapCenter installation.

## SnapCenter Plug-In Package for Windows

This installation package includes the following plug-ins:

**SnapCenter Plug-In for Microsoft SQL Server.** The plug-in for SQL Server is a host-side component of the NetApp integrated storage solution offering application-aware backup management of Microsoft SQL Server databases. With the plug-in for SQL Server installed on your SQL Server host, SnapCenter automates Microsoft SQL Server database backup, restore, and cloning operations. If you want to perform data protection operations on SQL Servers that are on VMDKs or RDMs, you must register VSC with SnapCenter by using either the SnapCenter Add Hosts wizard or the VSC Configure SnapCenter Server dialog box.

**SnapCenter Plug-In for Microsoft Windows.** The plug-in for Windows provides storage provisioning, Snapshot copy consistency, and space reclamation for Windows hosts. With the plug-in installed on your Windows host, you can use SnapCenter to create and resize disks, initiate iSCSI sessions, manage initiator groups, and manage SMB shares. The plug-in for Windows is a required component of plug-in for SQL Server workflows. Support is provided for provisioning SMB shares only. You cannot use SnapCenter to back up SQL Server databases on SMB shares.

Table 7 lists the supported storage configurations for SQL Server databases supported by SnapCenter and the plug-in for Microsoft SQL Server.

Table 7) Supported storage types: SnapCenter Plug-In for Microsoft SQL Server.

Machine	Storage Type	Provision Storage Using:	Support Notes
Physical server	FC-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
	iSCSI-connected LUNs	SnapCenter GUI or PowerShell cmdlets	
	SMB3 shares residing on a SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.* You cannot use SnapCenter to back up databases on SMB3 shares.
VMware VM	RDM LUNs connected by an FC or iSCSI HBA	PowerShell cmdlets	You must register Virtual Storage Console for VMware vSphere with SnapCenter before you can use SnapCenter to back up databases on RDM LUNs.



	iSCSI LUNs connected directly to the guest system by iSCSI initiator	SnapCenter GUI or PowerShell cmdlets	
	VMDKs on VMFS or NFS datastores	VMware vSphere or VSC cloning utility	You must register Virtual Storage Console for VMware vSphere with SnapCenter before you can use SnapCenter to back up databases on VMDKs.
	A guest system connected to SMB3 shares residing on an SVM	SnapCenter GUI or PowerShell cmdlets	Support for provisioning only.* You cannot use SnapCenter to back up databases on SMB shares.

\*SnapCenter provisioning support is available only on Windows 2012 or 2012 R2.

## 5.4 VMware vSphere Design

This section provides an overview of VMware vSphere design as part of a SQL Server deployment.

### vSphere Cluster Considerations

Separate clusters are recommended for infrastructure and production VMs to provide logical separation and fault isolation between the components. This helps to prevent a production server or servers from negatively affecting the performance of the infrastructure VMs, such as the vCenter Server, Active Directory domain controllers, or SnapCenter Server. Problems with these infrastructure components could in turn degrade the performance or availability of all virtual servers.

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

vSphere Dynamic Resource Scheduler (DRS) is also recommended to automatically balance CPU and memory workloads across the cluster members. DRS provides automated remediation of host resource contention and reduces the likelihood of a bully VM negatively affecting other VMs on the same host.

VMware vSphere 6.0 clusters can scale up to 64 nodes within a single cluster. This limit is unrelated to the number of physical CPUs or cores within any or all of the nodes. If larger operational scale is a primary concern, using larger hosts, such as with four or more processors and/or commensurately larger amounts of memory, allows greater density within a single vSphere cluster.

For high-performance workloads such as SQL Server, both VMware and Microsoft recommend not overcommitting CPU and memory resources on the vSphere hosts. Virtual CPUs configured in VMs on a host should not exceed the number of physical cores in that host to promote optimal performance. VMware also recommends using one CPU core per socket for VMs, which allows for better dynamic utilization of host resources. The 12-vCPU VM tested in this validation was configured with 12 CPU sockets with 1 core each.

Notably, as a host scales up, failure domains scale equally. Failure of a two-CPU server results in the failure of some number of VMs, whereas a four-CPU server affects twice as many VMs. Although failures

are rare, you must still take them into account when designing the infrastructure. Fewer hosts per cluster also provide vSphere DRS with fewer options for optimally balancing the workload, increasing the likelihood of host resource contention.

## vSphere Networking Considerations

In this reference architecture, standard vSwitches were used for the VM connectivity, management vmkernel, and vMotion vmkernel port groups. Because storage access is provided by FCoE, additional networking configuration is not necessary unless the existing upstream network configuration imposes additional requirements. VMware vSphere Enterprise Plus licensing enables other networking options that provide you with additional features beyond standard virtual switches, including distributed virtual switches, the Cisco Nexus 1000v, or the Cisco Virtual Machine Fabric Extender (VM-FEX).

## VMware vCenter Considerations

VMware vCenter is a critical component for any vSphere infrastructure; all management and maintenance operations rely on its correct and efficient performance. Due to the critical nature of vCenter, NetApp recommends that you provide as much resilience and data protection as possible for its core components: the individual vCenter modules and the vCenter database.

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

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

The vCenter VM and its database server VMs should be sized appropriately to its workload. An example configuration for this solution is shown in Table 8 and Table 9.

**Table 8) VMware vCenter Server VM configuration.**

VMware vCenter Server VM	Configuration
OS	Microsoft Windows Server 2012 R2
vCPU	4 vCPUs
Memory	8GB
Network adapter type	VMXNET3
Hard disk size	60GB
Hard disk type	Thin

**Table 9) vCenter SQL Server database VM minimum configuration.**

vCenter SQL Server VM	Configuration
OS	Microsoft Windows Server 2012 R2
vCPU	2 vCPUs
Memory	4GB
Network adapter type	VMXNET3

vCenter SQL Server VM	Configuration
Hard disk size	60GB
Hard disk type	Thin

The Java Virtual Machine (JVM) heap size settings are configured during installation based on the available memory size of the vCenter Server. Therefore, NetApp recommends that you allocate the appropriate amount of CPU and memory prior to vCenter Server installation. The size of the required virtual server environment might change substantially postinstallation, such as from less than 1,000 VMs to more than 1,000 VMs. Therefore, it might become necessary to disruptively increase the JVM settings on the vCenter Server, as described in [VMware Knowledge Base \(KB\) article 2021302](#). This KB article also provides a matrix for JVM size for each requisite vCenter service based on the VM inventory size.

## 5.5 SQL Server 2014 Design

There are a number of design considerations that you must address when deploying Microsoft SQL Server 2014. For this validation, SQL Server 2014 was deployed on VMware VMs with Windows Server 2012 R2 as the guest OS.

### SQL Server Virtual Machines

When configuring SQL Server VMs, make sure that the VM is sized appropriately for the workload. VMs with multiple CPUs and large memory configurations place significant loads on the physical hardware, and best practices dictate that VMs should have as few vCPUs as possible to support the workload. This allows the hypervisor to schedule resources most efficiently and prevents CPU contention between large VMs.

To help suppress contention even further, Microsoft and VMware recommend that you should not overcommit physical resources in ESXi hosts for high-performance applications. Table 10 shows the configuration of each SQL Server 2014 VM configured for testing. Each ESXi host in this design has 24 physical CPU cores and 128GB memory, and so each host can support two SQL Server VMs as configured without oversubscribing CPU resources.

**Table 10) Production SQL Server 2014 VM configuration.**

SQL Server VM	Configuration
OS	Microsoft Windows Server 2012 R2
vCPU	12 vCPUs
Memory	8GB
Network adapter type	VMXNET3
Hard disk size	120GB
Hard disk type	Thin

### SQL Server Storage Configuration

Microsoft supports SQL Server 2014 in both physical and virtual environments. In virtual environments, there are a number of options for storage of database files. The following options are supported by Microsoft for SQL Server deployments:

- **In-guest iSCSI LUNs.** You can use the iSCSI initiator directly in the guest OS to connect to the storage array. Although this can be the simplest configuration because each SQL Server is managed

individually, it can also create network saturation on the ESXi host if network bandwidth is not provisioned accordingly.

- **Virtual disk files.** You can create virtual disks for VMs from existing datastores and provision SQL Server databases on them. This provides the highest level of flexibility by managing storage provisioning from the virtualization layer, so resources can be used most effectively. Microsoft and NetApp recommend that you host applications such as SQL Server databases on dedicated LUNs even if VMDKs are used in the guest. This configuration permits granular control for backup, recovery, and replication purposes and allows you to manage storage performance on a per-database basis.
- **Raw device mapping.** Customers can also provision LUNs for SQL Server databases that are allocated to the ESXi servers but then passed directly to the VM as a raw device. This arrangement has many of the same benefits as VMDKs, but it imposes serious limitations in the operational capabilities of the vSphere environment. RDM devices are becoming the last choice in deployments due to these limitations and the administrative overhead created by using RDM disks.

This solution is based on VMDK virtual disks, and as seen in the section “Solution Verification,” this configuration delivers ample performance for the SQL Server workload.

All database files should be on separate virtual disks from the OS and separate from each other. Database data and log files should be on separate dedicated LUNs, and tempdb should be on a separate LUN as well. LUNs should be mounted to servers by using volume mount points, which creates a more logical data structure and prevents you from running out of drive letters on the host.

## SQL Server AlwaysOn Availability Groups

AlwaysOn Availability Groups are an enterprise-level high-availability and disaster recovery solution introduced in SQL Server 2012 that enable you to maximize availability for one or more user databases. The use of AlwaysOn Availability Groups requires that SQL Server instances reside on Windows Server failover clustering (WSFC) nodes. For this validation, a single WSFC was created with 12 nodes. Six availability groups were created with two nodes each, and each group supported a single database. The availability groups were configured in synchronous commit mode, providing the highest level of data protection. When you host the primary and secondary copies on high-performance hardware with local network connectivity, the performance effect of synchronous commit mode is minimal and offers completely redundant and high-performing replicas of mission-critical databases.

## 6 Solution Verification

This reference architecture is based on a standard FlexPod infrastructure hosting Microsoft SQL Server 2014 databases. Cisco Nexus 9296PX switches in NX-OS mode connect the Cisco UCS compute nodes and NetApp storage arrays to the Ethernet network for client and NAS storage access. The storage arrays are connected directly to the Cisco UCS fabric interconnects for FCoE storage access. Cisco UCS B200-M4 blades running VMware vSphere ESXi 6.0 host Microsoft Windows Server 2012 VMs and use dedicated FCoE LUNs for individual database instances. Specific details of the configuration can be found in the section “Technology Requirements.”

To provide enterprise-class performance, management, and reliability, this solution was verified with the following test cases:

- The performance of the primary databases was validated with the Microsoft TPC-E toolkit and the TPC-E workload configured for Microsoft SQL Server 2014 AlwaysOn Availability Groups. Successful validation required that the primary Microsoft SQL Server 2014 databases deliver 150,000 to 200,000 I/O operations per second with an average read latency below 1ms.
- Performance of the secondary database configuration was validated after failover from the primary database instance. Successful validation required that the primary Microsoft SQL Server 2014

databases deliver 150,000 to 200,000 I/O operations per second with an average read latency below 1ms.

- NetApp SnapCenter was used to take application-integrated Snapshot copies of secondary database instances at the same time that primary instances were in operation.
- NetApp SnapCenter was used to recover databases from Snapshot copies after intentionally introduced database corruption.
- NetApp SnapCenter was used to clone databases and mount to alternate servers for validation or test/dev purposes.
- The infrastructure was subjected to a number of hardware resiliency tests while under the same workload used for performance validation to make sure that the workload would continue to run during the following failure and maintenance scenarios:
  - Disconnection of an FCoE link between the storage array and the fabric interconnect
  - Disconnection of a 10GbE link between the storage array and the fabric interconnect
  - Failover of each storage controller and subsequent takeover by the partner controller
  - Reboot of the primary fabric interconnect

During solution testing, Cisco UCS blade servers were used to host the infrastructure and SQL Server 2014 VMs. The database and infrastructure servers were hosted on discrete compute resources so that the workload to the NetApp AFF system could be precisely measured. It is a NetApp and industry best practice to separate production VMs from the infrastructure VMs because noisy neighbors or bully VMs can affect the infrastructure, which can have a negative effect on all users, applications, and performance results.

## 6.1 Performance Validation

The performance requirements for this solution were the delivery of 150,000 to 200,000 IOPS with sub-millisecond read latency by using SQL Server 2014 configured in an availability group configuration. The following sections describe the methodology and design considerations used to test the AFF8080 EX running a standard SQL Server workload.

### Database Configuration

For this validation, a total of six 2.5TB databases were used to host the simulated OLTP environment. Each storage system controller has a single data aggregate of 22 800GB SSDs, as shown in Table 4. Databases were balanced across the controllers by placing odd-numbered SQL Servers on controller 1 and even-numbered servers on controller 2.

The database layout shown in Table 11 was repeated for each of the six SQL Server databases. For each SQL Server database, the data files, log files, and tempdb were each contained in a separate LUN within a separate volume. LUNs were allocated to the vSphere ESXi hosts and formatted as vSphere datastores by using VMFS. Virtual disks were then allocated to each SQL Server VM and formatted with default settings.

Table 11) Database LUN and volume configuration.

LUN Name	Datastore Name	VMDK Size	File Name	File Size
/vol/sqlX_data/sqlX_data	sqlX_data	2.5TB	MSSQL_tpce_root.mdf	8MB
			Fixed_1.ndf	5MB
			Growing_1.ndf	580GB
			Growing_2.ndf	580GB

LUN Name	Datastore Name	VMDK Size	File Name	File Size
			Growing_3.ndf	580GB
			Growing_3.ndf	580GB
			Scaling_1.ndf	20GB
			Scaling_2.ndf	20GB
			Scaling_3.ndf	20GB
			Scaling_4.ndf	20GB
/vol/sqlX_log/sqlX_log	sqlX_log	500GB	TPCE_Log.ldf	200GB
/vol/sqlX_tempdb/sqlX_tempdb	sqlX_tempdb	300GB	tempdev01.mdf	8MB
			tempdev02.ndf	10GB
			tempdev03.ndf	10GB
			tempdev04.ndf	10GB
			tempdev05.ndf	10GB
			tempdev06.ndf	10GB
			tempdev07.ndf	10GB
			tempdev08.ndf	10GB

## Test Methodology

For this validation, the Microsoft TPC-E toolkit was used to generate an industry-standard OLTP warehouse transaction workload against the SQL Server 2014 test configuration. The workload generated a workload of approximately 90% reads and 10% writes against the SQL Server databases in the test configuration. The goal of these tests was not to measure the maximum performance of the configuration, but rather to validate that the performance available at a generally acceptable read latency of approximately 1ms was within the limits specified for the solution of 150,000 to 200,000 IOPS.

For these tests, a total of six SQL Server 2014 VMs were hosted on three ESXi servers. Each ESXi server was connected through FCoE to the AFF8080 EX storage array, as shown in Figure 9. Four LUNs were provisioned for each database server and formatted as VMFS datastores. Then a VMDK was provisioned from each datastore to the appropriate SQL Server VM. Each server created a database of 2.6TB total capacity, for a total size requirement of 16TB across all six primary SQL Server databases.

With these six SQL Server databases and the Microsoft TPC-E load generator, performance of the standalone databases was measured under successively heavier workloads until the read latency as observed by the database server exceeded 1ms. Then the workload was reduced until the read latency was under 1ms, and this configuration was used for all use-case testing. Each performance test was run for a total of 2.5 hours to make sure that a steady state was achieved. The IOPS and latency results reported are the aggregate across all six primary or secondary database servers.

Microsoft performance-monitoring tools were used to capture IOPS and latency data from each of the six SQL Server database servers to make sure that the observed performance met the validation requirements described earlier for each of the tested configurations.

## Test Results

The results of the testing are listed in Table 12. In all cases, the performance observed at the database level met or exceeded the requirements set for validation and demonstrate that a SQL Server environment based on FlexPod is capable of delivering enterprise-class throughput and latencies for the most demanding applications when configured as an availability group.

Table 12 shows the performance of the six primary databases configured in AlwaysOn Availability Groups and the secondary databases running after failover from the primary replica while still in a synchronized state.

Table 12) Validated database performance.

	AG Nominal	After AG Failover
IOPS	211,789	205,776
Read latency	0.97ms	0.94ms

In addition to the performance data described earlier, we observed VM, ESXi host, and storage controller CPU utilization during the testing. In all test configurations, we observed CPU utilization between 60% and 65% on both the database server instances and storage controllers and 70% on ESXi hosts. This indicates that this specific configuration was well balanced between compute and storage resources. Both resources had adequate headroom to handle additional workloads.

## 6.2 SnapCenter Validation

To demonstrate the data management capabilities of this solution, NetApp SnapCenter was used to create Snapshot copies of running databases, recover databases after simulated corruption, and clone databases for data verification or test/dev purposes.

### SnapCenter 1.1 Prerequisites

SnapCenter operations on each SQL Server require memory beyond the requirements of the SQL database instances. To make sure that SnapCenter processes operate efficiently, allocate 4GB of memory to each SQL server in addition to the memory requirements of the databases. SnapCenter Server requires at least two vCPUs and 8GB of RAM. Additional CPU and memory resources can improve performance depending on workload and job schedules.

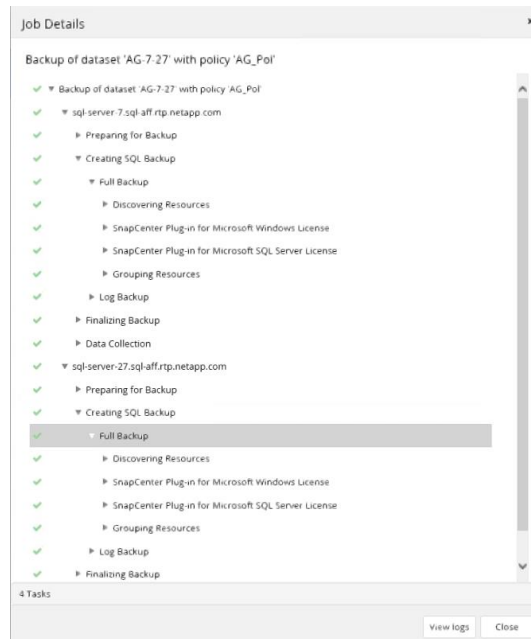
All Microsoft Windows servers should have current patch updates applied. Specifically, Microsoft KB2887595 must be applied for successful operations.

### Database Backup

For this solution validation, a backup was performed against the availability group, with both of the primary and secondary copies chosen for backup at the same time. Note that SnapCenter 1.1 does not currently support one-button backup of individual availability group databases. To perform a backup of a database in an availability group, a backup dataset must be created, including one or more databases in the availability group. Backup operations can be then performed using the backup dataset.



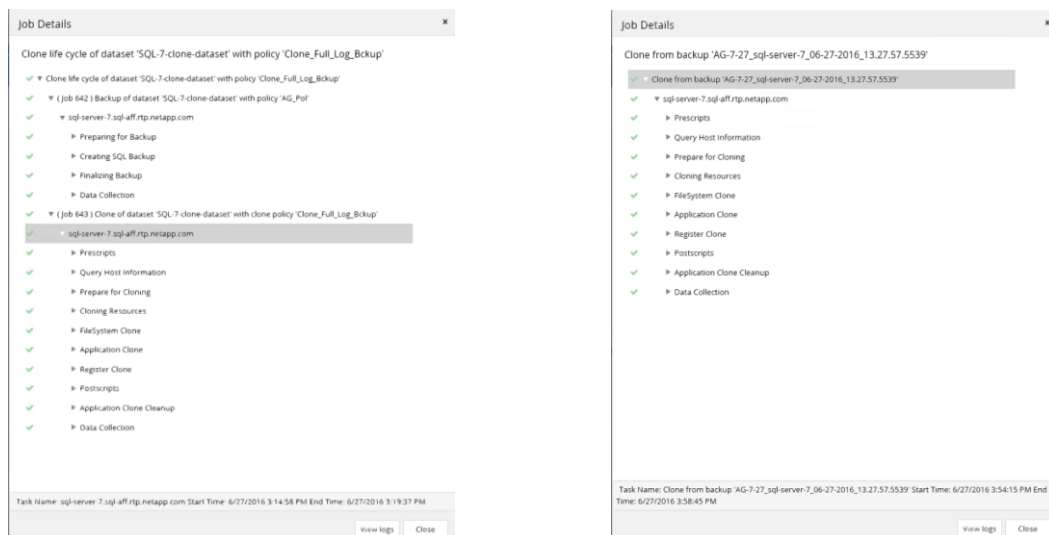
Figure 15) Backup of SQL database with SnapCenter.



## Database Cloning

Cloning of databases allows you to rapidly provision a copy of production data for further processing in analytics, test, or development environments. SnapCenter 1.1 can create clones from either the active database or an existing backup set. For validation purposes, clones were created using both of these methods.

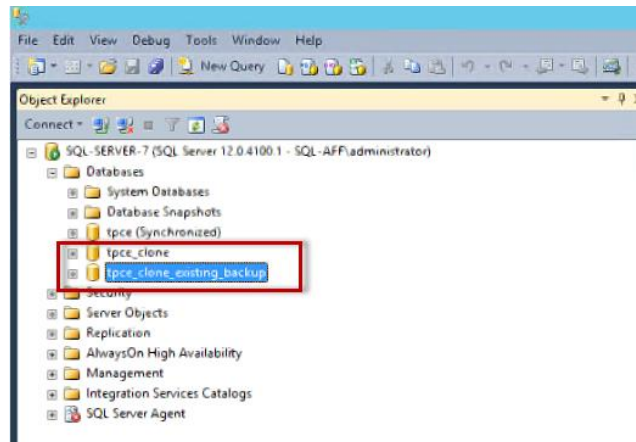
Figure 16) Database clone operations with SnapCenter.



After the cloning operation completes, Microsoft SQL Server Management Studio shows the cloned databases mounted as separate instances.



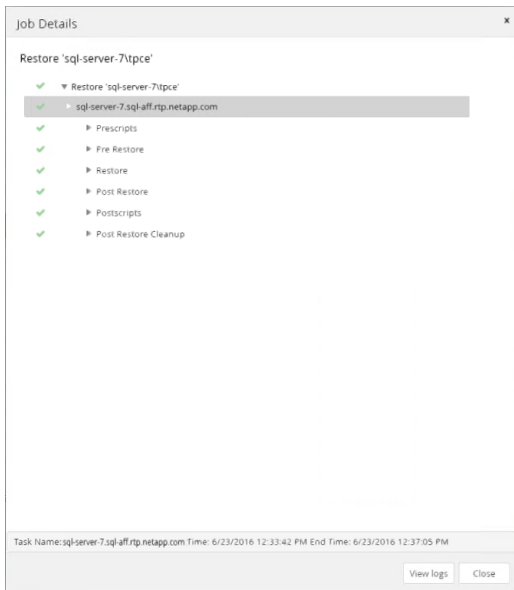
Figure 17) SQL Management Console-mounted database clones.

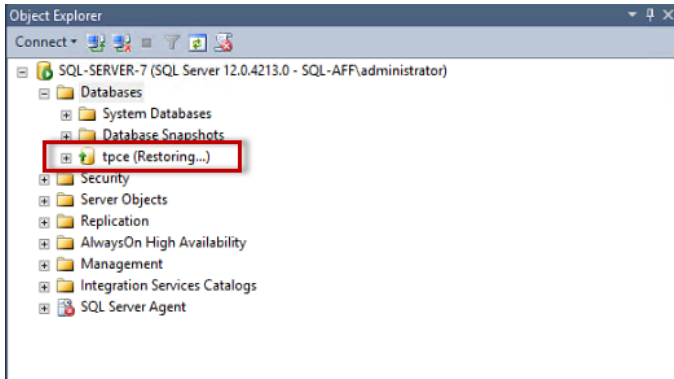


## Database Restore

SnapCenter 1.1 requires that a corrupted database be removed from the availability group before proceeding with the restore operation. Once the database is removed from the availability group configuration, SnapCenter is used to restore the database using the `NORECOVERY` option, which restores the data files but does not replay any existing log files. Once storage recovery is complete, the database logs from the surviving database instance are replayed to bring the restored copy into synchronization and rejoin the restored database replica to the availability group.

Figure 18) Database restore with SnapCenter.





## 6.3 Resiliency Validation

We then demonstrated the enterprise-class resiliency of this FlexPod solution. To do so, we induced a variety of failure scenarios into the system while the database servers and storage were subjected to the identical workload used for performance validation. The goal was to subject the FlexPod infrastructure to a heavy workload and measure the effect of the specific failure scenario on the overall performance and stability of the system. To pass each of these tests, the database and storage had to continue to serve I/O for the duration of the event. Some of these tests were extremely disruptive. Therefore, minimal drops in overall performance were considered acceptable as long as the overall system continued to function nominally and performance returned to pre-failure levels after the failure was corrected.

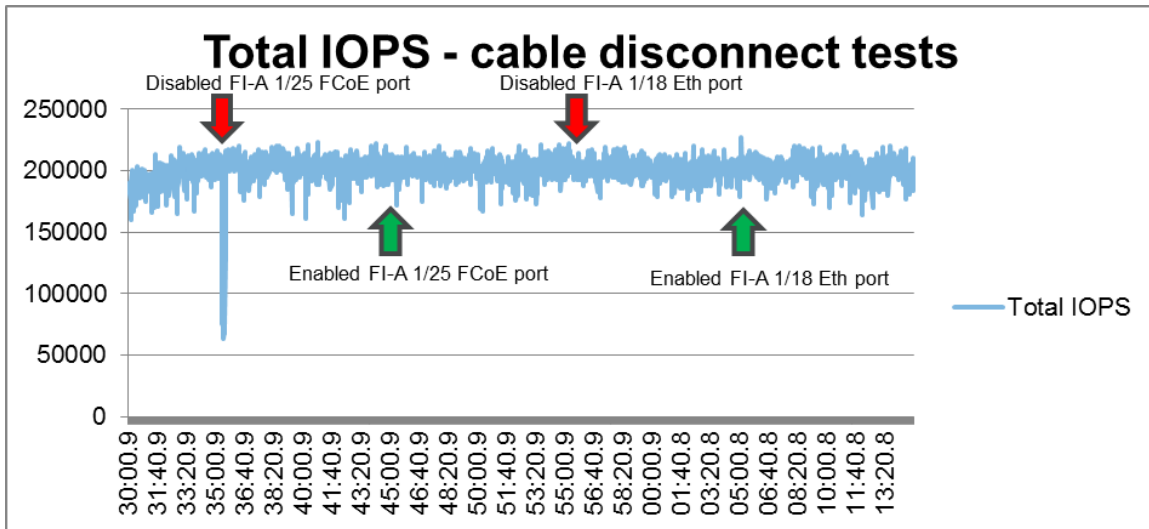
### Link Failure Tests

These tests intentionally disconnected and then reconnected a single FCoE and 10GbE link in the environment while under a heavy OLTP workload of approximately 200,000 IOPS. The workload was allowed to run for 5 minutes to make sure of a steady state before introducing the first failure and then allowed to run for 10 minutes between each failure before the failure was corrected.

Figure 16 shows the IOPS during these tests, with red arrows to indicate when the failure occurred and green arrows to indicate when the link was restored. Because primary storage access in this solution uses FCoE, there is a brief drop in IOPS when the FCoE link is disconnected. After the host multipath driver activates alternate paths, the observed throughput returns to nominal values in less than 30 seconds.

No noticeable performance impact was observed during removal or replacement of the 10GbE link. In this solution, 10GbE is used for synchronous replication of the SQL Server databases in the AlwaysOn Availability Group. Although the link failure may have affected latency of the replication, the effects were not enough to lower the overall throughput of the primary database.

Figure 19) IOPS during link failure tests.



### Storage Controller Failover Tests

These tests intentionally induced a catastrophic failure in the one controller of the AFF8080EX storage system during a 200,000 IOPS OLTP workload. As a result of this failure, the surviving storage controller was required to service the entire workload previously handled by both controllers.

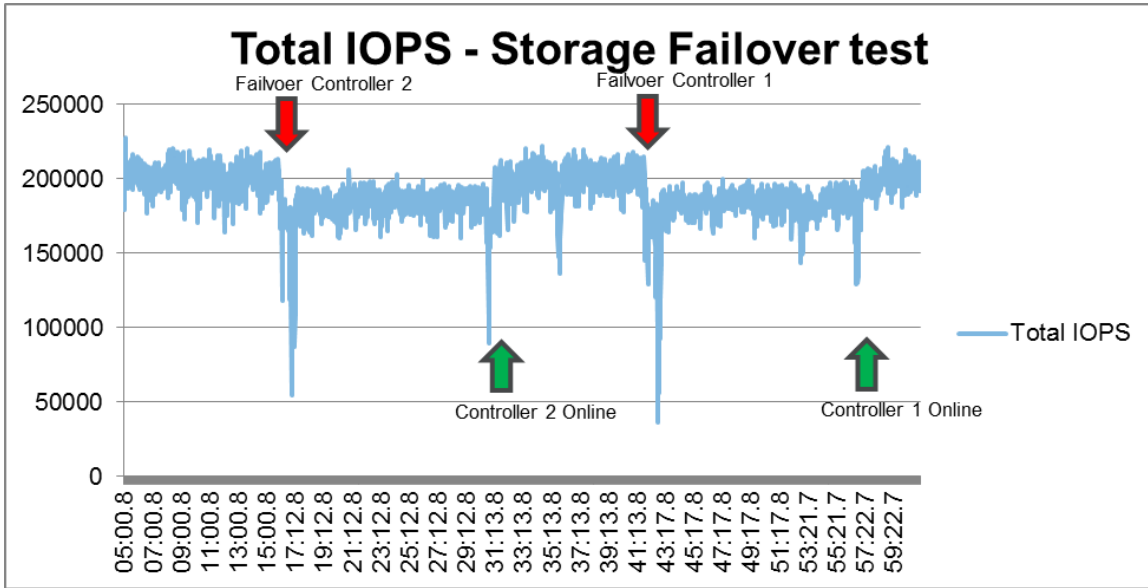
For these tests, the following protocol was used:

1. The OLTP workload was started and allowed to run on both controllers for 10 minutes.
2. Following forced failover of controller 2, controller 1 serviced the entire workload. During this time, controller 2 rebooted to a preoperational state and then entered a wait period. After the wait period, the system started an automatic giveback, and controller 2 resumed normal operations approximately 15 minutes later.
3. The workload was allowed to run for 10 minutes on both controllers to verify nominal operation.
4. After controller 1 failed over, controller 2 serviced the entire workload. Controller 1 resumed operation approximately 15 minutes later.
5. The workload was allowed to run for another 10 minutes to verify nominal operation.

Figure 17 shows the IOPS observed during these tests. Red arrows indicate when the controllers went offline, and green arrows indicate when they resumed normal operations. We observed the following results:

1. There was a noticeable reduction in IOPS at the moment the controllers went offline. This reduction in IOPS is the result of host failover driver timeout values. It represents the time necessary for the host to recognize the path failure and redirect I/O to the surviving controller. I/O did not stop completely because half of the database instances were not affected by each controller failure, and thus did not need to wait for path-failover timeouts.
2. Performance recovered while the controller was offline, but was on average roughly 10% lower than when both controllers were online. Storage controller CPU utilization was approximately 90%, and latencies were less than 10% higher than during normal operations.
3. After each controller was restored, performance quickly returned to prefailure levels.


Figure 20) Observed IOPS during controller failover tests.



### Cisco UCS Fabric Interconnect Failover Test

This test intentionally rebooted the primary Cisco UCS fabric interconnect while it ran the 200,000 IOPS OLTP workload. During the failure, all I/O was forced to traverse the secondary fabric interconnect. For this test, the workload was allowed to run for five minutes at steady state before the fabric interconnect was rebooted. During this test, a brief drop in IOPS occurred before a quick recovery to pre-failure levels was observed, and there was no apparent effect when the fabric interconnect resumed operation. **Error! Reference source not found.** shows the IOPS observed during this test, with a red arrow indicating when the fabric interconnect went offline and a green arrow indicating when it was confirmed to be back online.

Figure 21) Observed IOPS during fabric interconnect failover test.



## 7 Best Practices

Table 13 lists the recommended NetApp best practices for designing or implementing Microsoft SQL Server 2014 and AlwaysOn Availability Groups running on FlexPod Datacenter with VMware vSphere 6.

Table 13) SQL Server 2014 on FlexPod: best practices.

Best Practice Area	Best Practice Details
Cisco UCS design	<ul style="list-style-type: none"> <li>• Verify that there is sufficient bandwidth to each Cisco UCS blade by selecting appropriate IOM and adapter cards.</li> <li>• Configure the fabric interconnects for FC switch mode before completing configuration. FC switch mode requires a reboot of the fabric interconnect.</li> </ul>
VMware	<ul style="list-style-type: none"> <li>• The infrastructure VMs for VMware View should operate on different hardware than the production servers. This provides the highest level of fault isolation and performance assurance.</li> <li>• Enable VMware vSphere DRS and set it to fully automatic to balance workloads across all hosts within the vSphere clusters.</li> <li>• Right-size SQL Server VMs; do not add vCPU and memory unless the workload actually requires them. Overprovisioning VMs can cause unnecessary resource contention.</li> <li>• Do not overcommit CPU and memory resources on ESXi hosts.</li> <li>• Use the NetApp VSC to set recommended values on the vSphere hosts.</li> <li>• Use the NetApp VSC to provision datastores to the vSphere hosts. This reduces the amount of time required to provision and verifies the use of best practices.</li> </ul>
NetApp storage	<ul style="list-style-type: none"> <li>• Allocate disks from each disk shelf to both controllers to increase the bandwidth available to the disks.</li> <li>• Use NetApp ONTAP efficiency features such as deduplication and compression to maximize effective storage capacity. These features are enabled by default on AFF systems.</li> <li>• Create one LUN per volume.</li> <li>• Enable volume autosizing to avoid out-of-space conditions.</li> <li>• Enable read reallocation for I/O profiles consisting of large sequential reads.</li> <li>• Disable automatic Snapshot schedules.</li> <li>• Disable access time updates.</li> </ul>
SQL Server 2014	<ul style="list-style-type: none"> <li>• Use mount points to avoid running out of drive letters.</li> <li>• Databases, logs, tempdb, and system databases should all be on separate LUNs and separated from the server OS. This applies to both virtual and physical SQL instances.</li> <li>• AlwaysOn Availability Group replicas should be hosted on separate physical infrastructure to maximize performance and availability.</li> <li>• Use one tempdb file per core to a maximum of eight.</li> </ul>

## 8 Conclusion

FlexPod Datacenter is the optimal infrastructure foundation on which to deploy Microsoft SQL Server 2014. Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and designs. The flexibility and scalability of FlexPod allow you to start out with a right-sized infrastructure that can grow with and adapt to your evolving business requirements, from single SQL Server deployments to consolidated database farms.

In the verification tests of this reference architecture, the total IOPS and latency as measured at the database met performance expectations. The NetApp FlexPod Datacenter configuration with AFF

reached over 200,000 combined IOPS and read latencies of less than 1ms while averaging 60% CPU utilization during most operations. In addition, this verification demonstrated that the solution as documented could maintain these performance levels even in the event of a storage, host, or network failure.

## Acknowledgements

This document is a result of the work, prior documentation, and assistance from Jens Dickmeis, Pat Sinthusan, Atul Bhalodia, Chris Lemmons, Troy Mangum, and the Converged Infrastructure team at NetApp.

## References

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

### Cisco Unified Computing System

- Cisco Design Zone for FlexPod  
[http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/landing\\_flexpod.html](http://www.cisco.com/c/en/us/solutions/enterprise/data-center-designs-cloud-computing/landing_flexpod.html)
- Cisco Unified Computing System  
<http://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>
- Cisco UCS 6200 Series Fabric Interconnects  
<http://www.cisco.com/en/US/products/ps11544/index.html>
- Cisco UCS 5100 Series Blade Server Chassis  
<http://www.cisco.com/en/US/products/ps10279/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/en/US/products/ps10277/prod\\_module\\_series\\_home.html](http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html)
- Cisco UCS Manager  
<http://www.cisco.com/en/US/products/ps10281/index.html>

### Cisco Nexus Networking

- Cisco Nexus 9000 Series Switches  
<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>
- Cisco Nexus 9396PX Switch information  
<http://www.cisco.com/c/en/us/support/switches/nexus-9396px-switch/model.html>

### NetApp FAS Storage

- Clustered Data ONTAP 8.3.1 Documentation  
<http://mysupport.netapp.com/documentation/docweb/index.html?productID=62175&language=en-US>
- TR-3982: NetApp Clustered Data ONTAP 8.3  
<http://www.netapp.com/us/media/tr-3982.pdf>
- Guest OS Tunings for a VMware vSphere Environment (KB)  
<https://kb.netapp.com/support/index?page=content&id=3013622>

- TR-4403: NetApp AFF8080 EX Performance and Server Consolidation with Microsoft SQL Server 2014  
<https://fieldportal.netapp.com/content/248568?assetComponentId=248696>

## NetApp SnapCenter

- SnapCenter Resources Page  
<http://mysupport.netapp.com/snapcenter/resources>

## VMware vSphere

- VMware vSphere Documentation Center  
<http://pubs.vmware.com/vsphere-60/index.jsp>
- SQL Server on VMware Best Practices Guide  
[https://www.vmware.com/files/pdf/solutions/SQL\\_Server\\_on\\_VMware-Best\\_Practices\\_Guide.pdf](https://www.vmware.com/files/pdf/solutions/SQL_Server_on_VMware-Best_Practices_Guide.pdf)

## Interoperability Matrixes

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

## Version History

Version	Date	Document Version History
Version 1.0	July 2016	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 © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

## Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. NVA-0025-DESIGN