



NetApp Verified Architecture

# FlexPod Datacenter with Microsoft Applications and NetApp AFF A-Series

## NVA Design

Glenn Sizemore, Bhavin Shah, NetApp  
September 2017 | NVA-1116-DESIGN | Version 1.0

Reviewed by



### Abstract

This document discusses the design considerations for architecting a solution for Microsoft Exchange 2016 and Microsoft SharePoint 2016 on FlexPod® Datacenter.

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary</b>	<b>4</b>
<b>2</b>	<b>Program Summary</b>	<b>4</b>
2.1	FlexPod Program Benefits	5
<b>3</b>	<b>Solution Overview</b>	<b>6</b>
3.1	Target Audience	6
3.2	Solution Technology	6
3.3	Use Case Summary	7
<b>4</b>	<b>Technology Requirements</b>	<b>8</b>
4.1	Hardware Requirements	8
4.2	Software Requirements	8
<b>5</b>	<b>Solution Design</b>	<b>9</b>
5.1	Cisco Unified Computing System	9
5.2	Cisco Nexus Switches	12
5.3	NetApp Storage	14
5.4	VMware vSphere	17
5.5	NetApp SnapCenter 3.0	18
5.6	Microsoft Exchange 2016	19
5.7	SnapManager for Exchange	29
5.8	Microsoft SharePoint 2016	32
5.9	DocAve Manager	34
<b>6</b>	<b>Design Considerations</b>	<b>34</b>
<b>7</b>	<b>Solution Verification</b>	<b>38</b>
<b>8</b>	<b>Conclusion</b>	<b>38</b>
	<b>Acknowledgements</b>	<b>38</b>
	<b>Where to Find Additional Information</b>	<b>39</b>
	<b>Version History</b>	<b>39</b>

## LIST OF TABLES

Table 1)	Hardware requirements	8
Table 2)	Software requirements	8
Table 3)	Requirements for Exchange environment configuration	24
Table 4)	Requirements for mailbox database copy configuration	24

Table 5) Requirements for tier 1 user mailbox configuration. ....	25
Table 6) SPECint test details. ....	25
Table 7) Requirements for server configuration. ....	25
Table 8) Requirements for processor configuration. ....	26
Table 9) Requirements for database configuration. ....	26
Table 10) Transport configuration. ....	26
Table 11) Requirements for the process core ratio. ....	26
Table 12) Requirements for environment configuration. ....	26
Table 13) Requirements for user mailbox configuration. ....	27
Table 14) Requirements for database copy instance configuration. ....	27
Table 15) Requirements for database configuration. ....	27
Table 16) Requirements for database copy configuration. ....	27
Table 17) Requirements for server configuration. ....	27
Table 18) Requirements for details of disk space. ....	28
Table 19) Requirements for details of host I/O and throughput. ....	28
Table 20) Transport calculation details. ....	28
Table 21) Requirements for Exchange Servers. ....	29
Table 22) SharePoint MinRoles. ....	33
Table 23) Design considerations. ....	34

## LIST OF FIGURES

Figure 1) FlexPod component families. ....	5
Figure 2) FlexPod Datacenter for Microsoft Exchange 2016 and Microsoft SharePoint 2016 solution topology. ....	7
Figure 3) vNIC failure resolution. ....	11
Figure 4) Chassis discover policy: discrete mode vs. port-channel mode. ....	12
Figure 5) Storage efficiency. ....	15
Figure 6) iSCSI SVM ports and LIF layout. ....	16
Figure 7) Unified namespace (unbound model). ....	21
Figure 8) Dedicated namespace (bound model). ....	22
Figure 9) SnapManager for Exchange components. ....	30

## 1 Executive Summary

A clear majority of organizations have started to shift their messaging and collaboration services to the cloud. There are still a great many use cases where it is both financially feasible and preferable to continue to maintain a portion of the total organization on the premises. For those private infrastructures, Microsoft Exchange Server continues to facilitate critical business e-mail communication, group scheduling, and calendaring on a 24/7 basis. Wherever deployed, system failures could result in unacceptable operational and financial losses. Because of the increasing importance of Microsoft Exchange Server, data protection, disaster recovery, and high availability are of increasing concern. Companies require quick recovery with little or no data loss.

With the rapid growth of Microsoft Exchange Server databases, it is increasingly difficult to complete time-consuming backup operations quickly. When an outage occurs, it can take days to restore service from slower media such as tape, even if all the backup tapes are available and error free. NetApp offers a comprehensive suite of hardware and software solutions that enable an organization to keep pace with the increasing data availability demands of an ever-expanding Microsoft Exchange Server environment. The solutions also scale to accommodate future needs while reducing cost and complexity.

In addition to Microsoft Exchange, Microsoft SharePoint is also a key enterprise application that the majority of organizations use today. Microsoft SharePoint is a collaboration platform used by enterprise businesses for their intranet/extranet websites as a method to centralize access to share and organize information, people, and projects. Microsoft SharePoint implementations have become more complex and require a greater degree of reliability. The underlying physical storage architecture that supports SharePoint farms is expected to scale to meet the growing capacity needs with suitable data protection capabilities for high-availability (HA) and disaster recovery purposes.

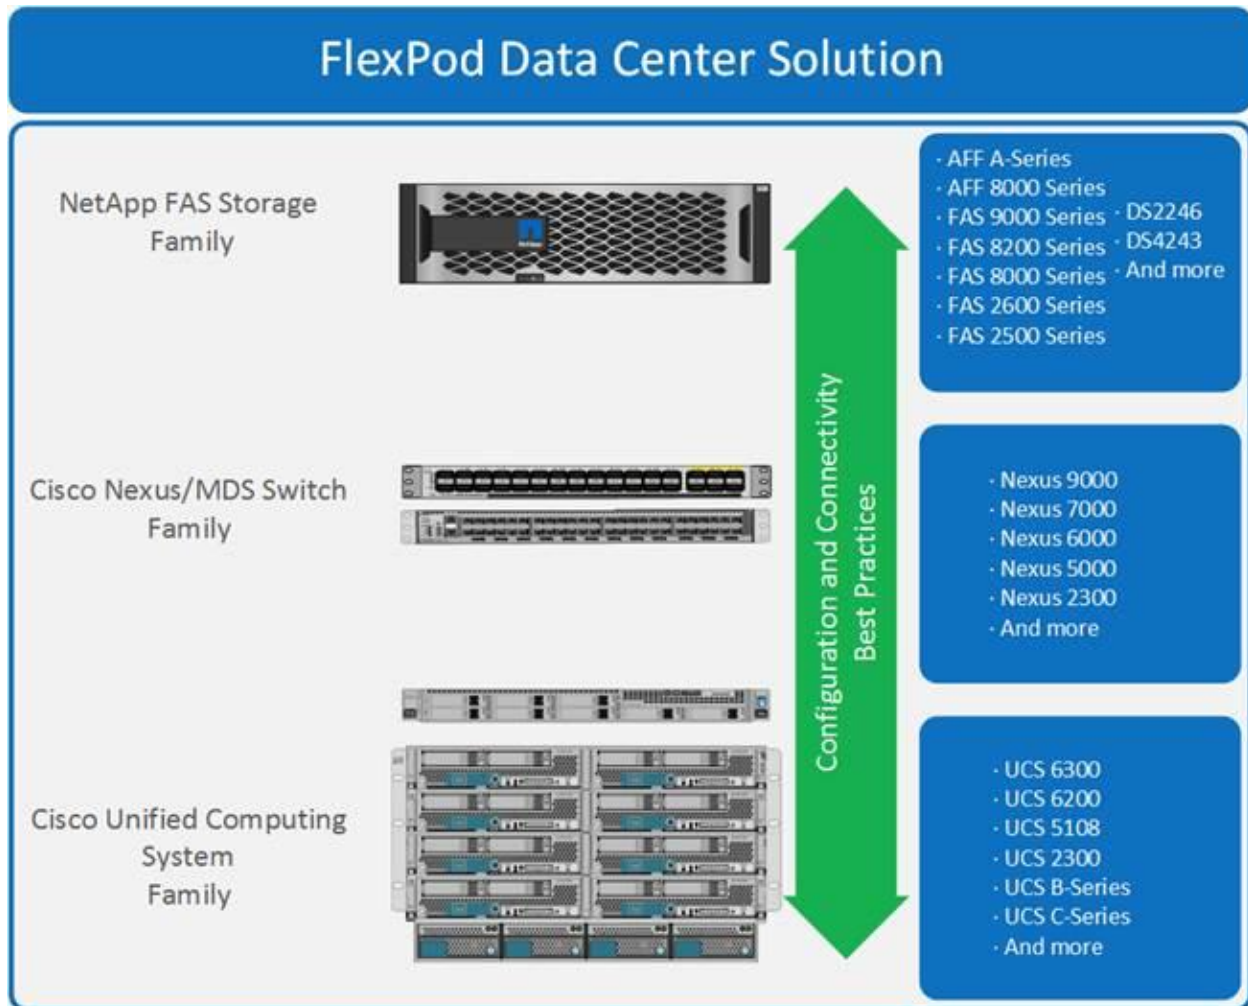
NetApp and Cisco understand these requirements and want to help organizations by designing a scalable, reliable, and fault-tolerant infrastructure where you can deploy and run these critical business applications. Therefore, we have designed a solution for running Microsoft Exchange 2016 and Microsoft SharePoint 2016 on a single FlexPod Datacenter architecture.

This document describes the design decisions that were taken to build this solution on top of a FlexPod system with VMware vSphere 6.5, NetApp® ONTAP® 9.1, Windows 2016, and the latest versions of Cisco Unified Computing System (Cisco UCS) and Cisco Nexus series software.

## 2 Program Summary

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

Figure 1) FlexPod component families.



## 2.1 FlexPod Program Benefits

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

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

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

potential issues. FlexPod supports tight integration with virtualized and cloud infrastructures, making it the logical choice for long-term investment.

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

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

### 3 Solution Overview

While architecting an on-premises solution to host enterprise applications such as Microsoft Exchange and Microsoft SharePoint, you must answer some questions:

- How do I build a secure and resilient infrastructure?
- How can I make sure of the highest level of availability?
- What return on investment (ROI) can I expect?
- How can I build a future-proof infrastructure?
- How can I reduce the complexity of my infrastructure?

The FlexPod architecture is designed to help you answer all these questions. By introducing standardization, FlexPod helps you mitigate the risks and uncertainty involved in planning, designing, and implementing a next-generation data center architecture.

This document focuses on VMware vSphere 6.5, Microsoft Exchange 2016, Microsoft SharePoint 2016, and ONTAP 9.1 built on the FlexPod Datacenter architecture. This document also discusses design choices and best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific releases of the components described in this document but are also applicable to other versions.

#### 3.1 Target Audience

The intended audience for this document includes sales engineers, field consultants, professional services personnel, IT managers, and partner engineering personnel. This document is also intended for customers who want to take advantage of all the best practices and guidelines for deploying multiple enterprise applications on a shared infrastructure platform.

#### 3.2 Solution Technology

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

- Cisco Unified Computing System
- Cisco Nexus switches
- NetApp AFF or FAS systems


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

Typically, the more scalable and flexible a solution is, the more difficult it becomes to maintain a single unified architecture capable of offering the same features and functionality across implementations. This is one of the key benefits of FlexPod. Each of the component families shown in Figure 2 offers platform and resource options to scale the infrastructure up or down while supporting the same features and functionality that are required under the configuration and connectivity best practices of FlexPod.

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

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

Figure 2) FlexPod Datacenter for Microsoft Exchange 2016 and Microsoft SharePoint 2016 solution topology.



### 3.3 Use Case Summary

The FlexPod Datacenter with Microsoft Exchange 2016 and Microsoft SharePoint 2016 solution architecture provides a flexible framework for the following use cases:

- Deploying a solution to run Exchange and SharePoint on a single FlexPod platform
- Architecting a SharePoint farm for 10,000 active users
- Architecting an Exchange environment for 10,000 active users with 5GB mailboxes

## 4 Technology Requirements

This section covers the technology components for the FlexPod Datacenter with Microsoft Exchange 2016 and Microsoft SharePoint 2016 solution.

### 4.1 Hardware Requirements

Both the enterprise applications that are part of this solution are deployed on the same compute, storage, and network infrastructure. The infrastructure includes the latest NetApp AFF A300 storage controllers with DS224C disk shelves and Cisco UCS B200 M4 server blades to host the vSphere environment that is supporting our workload. Table 1 lists all the hardware components that were used for the validation effort. The secondary stack that was part of this deployment used the older NetApp AFF8040 controllers to run the SQL Server instance for our database availability groups (DAGs).

Table 1) Hardware requirements.

Hardware	Model Number	Quantity
Cisco UCS 6200 Series fabric interconnects	FI 6248UP	2
Cisco UCS B200 blades	B200 M4 using Cisco UCS VIC 1340	8
Cisco UCS 5108 chassis	Cisco UCSB-5108-AC2	1
Cisco Nexus 9000	Cisco Nexus 9396PX	2
NetApp AFF A300	AFF A300	1 HA pair
NetApp DS224C disk shelves	Disk shelves populated with 3.8TB SSDs	2 shelves with 24 drives each

### 4.2 Software Requirements

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

Table 2) Software requirements.

Software/Firmware	Version
<b>Compute</b>	
Cisco UCS Manager	3.1(3a)
<b>Networking</b>	
Cisco NX-OS	7.0(3)4(6)
<b>Storage</b>	
NetApp ONTAP	9.1
NetApp VSC	6.2.1
<b>VMware vSphere</b>	
VMware ESXi	6.5.0, 4887370
VMware vCenter Server	6.5.0, 4944578



Software/Firmware	Version
<b>Microsoft SQL Server</b>	
Microsoft SQL Server	2016
Microsoft SQL Server Management Studio	16.5.3
<b>Microsoft Apps</b>	
Microsoft SharePoint	2016
Microsoft Exchange	2016
<b>Backup and Recovery</b>	
DocAve Backup and Restore	Version 6 SF9
NetApp SnapDrive®	7.1.4
NetApp SnapManager® for Exchange	7.2
NetApp single mailbox recovery	7.2

## 5 Solution Design

The FlexPod Datacenter with Microsoft Exchange 2016 and Microsoft SharePoint 2016 solution consists of the following designs:

- Cisco Unified Computing System
- Cisco Nexus Switches
- NetApp storage
- VMware vSphere
- Microsoft SharePoint
- DocAve Manager
- Microsoft Exchange
- SnapManager for Exchange

### 5.1 Cisco Unified Computing System

#### Cisco UCS 6200 Fabric Interconnects

The Cisco UCS 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 regardless of a server or virtual machine's topological location in the system.

The fabric interconnect provides both network connectivity and management capabilities for the Cisco Unified Computing System. IOM modules in the blade chassis support the power supply, along with fan and blade management. They also support port channeling and, thus, better use of bandwidth. The IOMs support virtualization-aware networking in conjunction with the fabric interconnects and Cisco Virtual Interface Cards (VICs).

## Cisco UCS Differentiators

The Cisco Unified Computing System has revolutionized the way servers are managed in data centers. The following are the unique differentiators of Cisco UCS and Cisco UCS Manager:

- **Embedded management.** In Cisco UCS, the servers are managed by the embedded firmware in the fabric interconnects, eliminating the need for any external physical or virtual devices to manage the servers.
- **Unified fabric.** In Cisco UCS, from blade server chassis or rack servers to fabric interconnect, there is a single Ethernet cable used for LAN, SAN, and management traffic. This converged I/O results in reduced cables, SFPs, and adapters, reducing capital and operational expenses of the overall solution.
- **Autodiscovery.** When the blade server is simply inserted in the chassis or the rack server is connected to the fabric interconnect, the discovery and inventory of compute resource occur automatically without any management intervention. The combination of unified fabric and autodiscovery enables the wire-once architecture of Cisco UCS, where compute capability of Cisco UCS can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
- **Policy-based resource classification.** When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on policies defined. This capability is useful in multi-tenant cloud computing.
- **Combined rack and blade server management:** Cisco UCS Manager can manage Cisco UCS B-Series blade servers and Cisco UCS C-Series rack-mount servers under the same Cisco UCS domain. This feature along with stateless computing makes compute resources truly hardware form factor agnostic.
- **Model-based management architecture.** Cisco UCS Manager architecture and management database are model based and data driven. An open XML API is provided to operate on the management model. This enables easy and scalable integration of Cisco UCS Manager with other management systems.
- **Policies, pools, and templates.** The management approach in Cisco UCS Manager is based on defining policies, pools, and templates, instead of cluttered configuration. This approach enables a simple, loosely coupled, data-driven approach in managing compute, network, and storage resources.
- **Loose referential integrity.** In Cisco UCS Manager, a service profile, port profile, or policies can refer to other policies or logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy, or a referred policy can be deleted even though other policies are referring to it. This allows different subject matter experts to work independently from each other. Loose referential integrity provides great flexibility in which different experts from different domains, such as network, storage, security, server, and virtualization, work together to accomplish a complex task.
- **Policy resolution.** In Cisco UCS Manager, a tree structure of organizational unit hierarchy can be created that mimics the real-life tenants and/or organization relationships. Various policies, pools, and templates can be defined at different levels of the organization hierarchy. A policy referring to another policy by name is resolved in the organization hierarchy with closest policy match. If no policy with specific name is found in the hierarchy of the root organization, then a special policy named “default” is searched. This policy resolution practice enables automation-friendly management APIs and provides great flexibility to owners of different organizations.
- **Service profiles and stateless computing.** A service profile is a logical representation of a server, carrying its various identities and policies. This logical server can be assigned to any physical compute resource as far as it meets the resource requirements. Stateless computing enables procurement of a server within minutes; it used to take days in legacy server management systems.
- **Built-in multitenancy support.** The combination of policies, pools, and templates; loose referential integrity; policy resolution in organization hierarchy; and a service profiles-based approach to

compute resources makes Cisco UCS Manager inherently friendly to multitenant environments typically observed in private and public clouds.

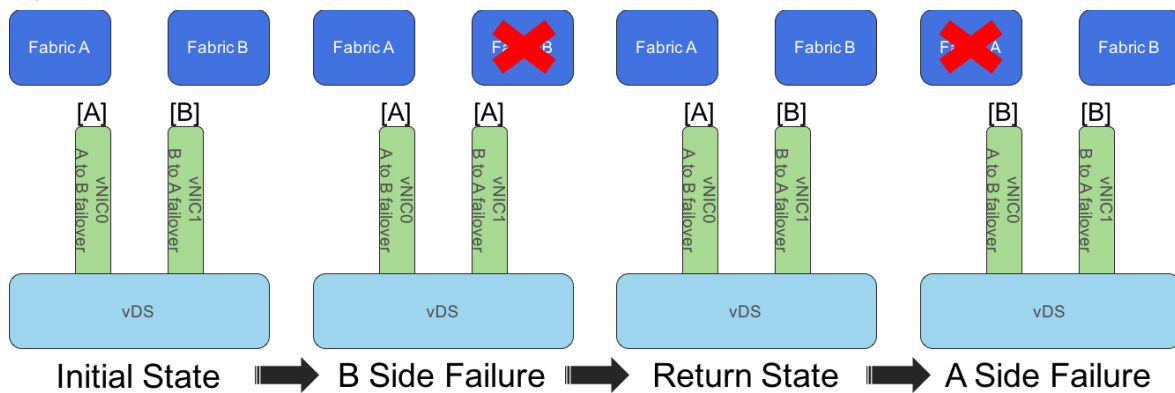
- **Extended memory.** The enterprise-class Cisco UCS B200 M4 blade server extends the capabilities of the Cisco UCS portfolio in a half-width blade form factor. The Cisco UCS B200 M4 harnesses the power of the latest Intel Xeon E5-2600 v4 Series processor family CPUs with up to 1536GB of RAM (using 64GB DIMMs), allowing a huge VM-to-physical server ratio required in many deployments or large memory operations required by certain architectures such as big data.
- **Virtualization-aware network.** Cisco VM-FEX technology makes the access network layer aware about host virtualization. This prevents pollution of compute and network domains with virtualization when the virtual network is managed by port profiles defined by the network administrators' team. VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus allowing the hypervisor CPU to do more virtualization-related tasks. VM-FEX technology is well integrated with VMware vCenter, Linux KVM, and Hyper-V SR-IOV to simplify cloud management.
- **Simplified quality of service (QoS).** Even though Fibre Channel and Ethernet are converged in the Cisco UCS fabric, built-in support for QoS and lossless Ethernet makes it seamless. Network QoS is simplified in Cisco UCS Manager by representing all system classes in one GUI panel.

## Cisco UCS Design Options in FlexPod

### Cisco UCS vNICs

Cisco UCS vNIC templates in the FlexPod architecture, other than those used for iSCSI vNICs, are set with failover enabled to allow for hardware resolution of an uplink loss from the fabric interconnect, as shown in Figure 3. This enables faster resolution of an uplink loss during disruption of the fabric interconnect than would occur from polling by the vSwitch or vDS.

Figure 3) vNIC failure resolution.

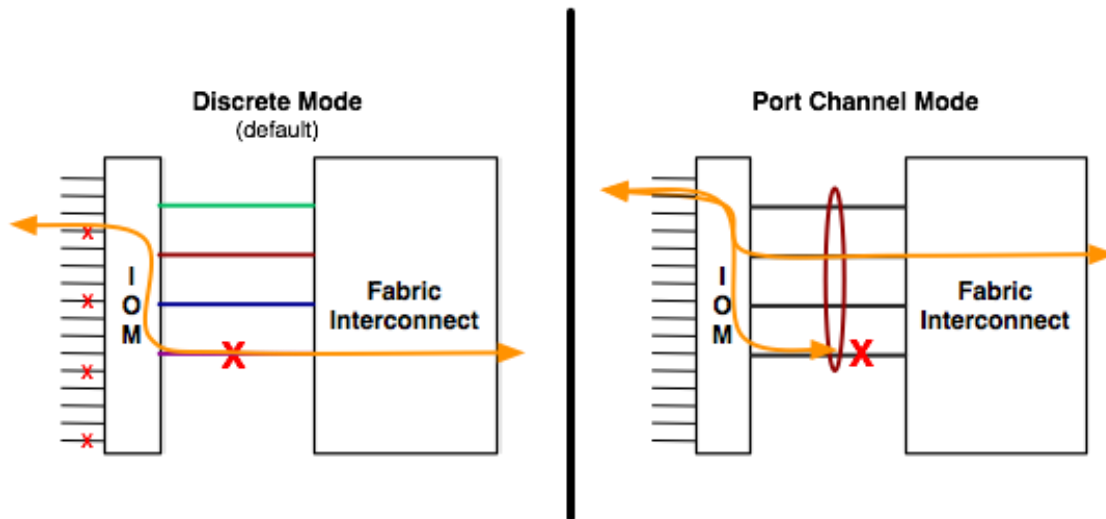


This failover makes the active/standby configuration for vNICs redundant because the standby uplink is not called on within the vDS. However, the presence of the standby uplink within the configuration averts an uplink redundancy-missing alarm within vCenter.

### Cisco UCS Chassis/FEX Discovery Policy

A Cisco UCS system can be configured to discover a chassis using discrete mode or port-channel mode, as shown in Figure 4. In discrete mode, each FEX KR connection and therefore server connection is tied or pinned to a network fabric connection homed to a port on the fabric interconnect. In the presence of a failure on the external link, all KR connections are disabled within the FEX I/O module. In port-channel mode, the failure of a network fabric link allows for redistribution of flows across the remaining port-channel members. Port-channel mode therefore is less disruptive to the fabric and hence recommended in FlexPod designs.

Figure 4) Chassis discover policy: discrete mode vs. port-channel mode.



## Cisco Unified Computing System: QoS and Jumbo Frames

FlexPod accommodates myriad 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. When setting up jumbo frames, it is important to make sure MTU settings are applied uniformly across the stack to prevent packet drops and negative performance.

## 5.2 Cisco Nexus Switches

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

### Network Design Overview

#### Network Switching

Two Cisco Nexus 9396PX switches running NX-OS software release 7.0(3)I4(6) were used in this solution verification. These switches were chosen because of their support for the latest NX-OS feature set, scalability, and readiness for ACI. This design does not utilize ACI but instead has the switches operating in standalone NX-OS mode. One of the design goals for this reference architecture was applicability to the widest range of customer environments. Therefore, ACI was not considered to be a requirement, but this architecture could be integrated into a new or existing ACI topology if desired. vPCs were used, allowing a port channel from each storage controller and Cisco UCS fabric interconnect to be spread across both switches. The Cisco Nexus 9000 series currently does not support converged networking with FCoE. If FC or FCoE connectivity is a requirement in a Cisco Nexus 9000 environment, the NetApp storage arrays can be connected directly to the Cisco UCS fabric interconnects, or Cisco MDS Series switches can be added to the solution.

#### Host Server Networking

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

LUNs on the NetApp AFF A300 storage array using the iSCSI interfaces and access NFS datastores on the AFF A300 using a dedicated VLAN on the IP interfaces.

## Storage Networking

Each of the two NetApp AFF A300 storage controllers has a four-port LACP ifgrp (port channel) connected to a vPC across the two Cisco Nexus 9396PX switches. ALUA was used to provide multipathing and load balancing of the iSCSI links. Initiator groups (igroups) were configured on the AFF A300 systems to map boot LUNs to the ESXi host servers. If you prefer FCoE for SAN boot instead of iSCSI, the AFF A300 storage system can be directly connected to the Cisco UCS fabric interconnects.

## Cisco Nexus 9000 Switch

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

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

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

FlexPod is a converged infrastructure platform. This convergence is possible because of the support for Ethernet enhancements across the integrated compute stack with regard to bandwidth allocation and flow control based on traffic classification. Therefore, it is important to implement the following QoS techniques to provide QoS in the FlexPod configuration:

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

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

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

- The following Cisco Nexus 9000 features were enabled:
  - **LACP.** Part of 802.3ad

- **Cisco vPC.** For link and device resiliency
- **Cisco Discovery Protocol (CDP).** For infrastructure visibility and troubleshooting
- The following vPC settings were configured:
  - A unique domain ID was defined.
  - The priority of the intended vPC primary switch was set lower than the secondary (the default priority is 32768).
  - Peer-keepalive connectivity was established.

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

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

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

  - A minimum of two 10GbE connections is required for vPC.
  - All port channels were configured in LACP active mode.
- The following Spanning Tree settings were configured:
  - The path cost method was set to long to account for 10GbE links in the environment.
  - The Spanning Tree priority was not modified (under the assumption that this was an access layer deployment).
  - Loopguard was disabled by default.
  - Bridge protocol data unit (BPDU) guard and filtering were enabled by default.
  - Bridge assurance was only enabled on the vPC peer link.
  - Ports facing the NetApp storage controller and the Cisco UCS were defined as edge trunk ports. For configuration details, see the Cisco Nexus 9000 Series switch [configuration guides](#).

### 5.3 NetApp Storage

This design leverages NetApp AFF A300 controllers deployed with ONTAP 9.1.

#### NetApp AFF A-Series Design

NetApp A-Series all-flash controllers were designed with flash storage in mind. They provide industry-leading density, scalability, and network connectivity, allowing customers to do more with their flash storage. The AFF A300 controller provides a rich set of data management features as well as industry-leading performance in a 3U form factor. ONTAP 9.1 provides many key features that optimize SSD performance and endurance, including the following:

- Coalesced writes to free blocks to maximize flash performance and longevity
- Flash-specific read path optimizations to enable consistent low latency
- Advanced drive partitioning to increase storage efficiency, increasing usable capacity by almost 20%
- Support for multistream writes to increase write performance to SSDs

The FlexPod Datacenter converged infrastructure supports a variety of NetApp controllers, including the AFF A-Series, AFF8000, FAS9000, FAS8000, FAS2600, and FAS2500 platforms. For a full list of supported controllers, see the [NetApp Interoperability Matrix Tool \(IMT\)](#) and the [FlexPod Technical Specification](#).

For more information about the AFF A-Series product family, see the [AFF A-Series All Flash Arrays product page](#) and [NetApp All Flash FAS Resources page](#).

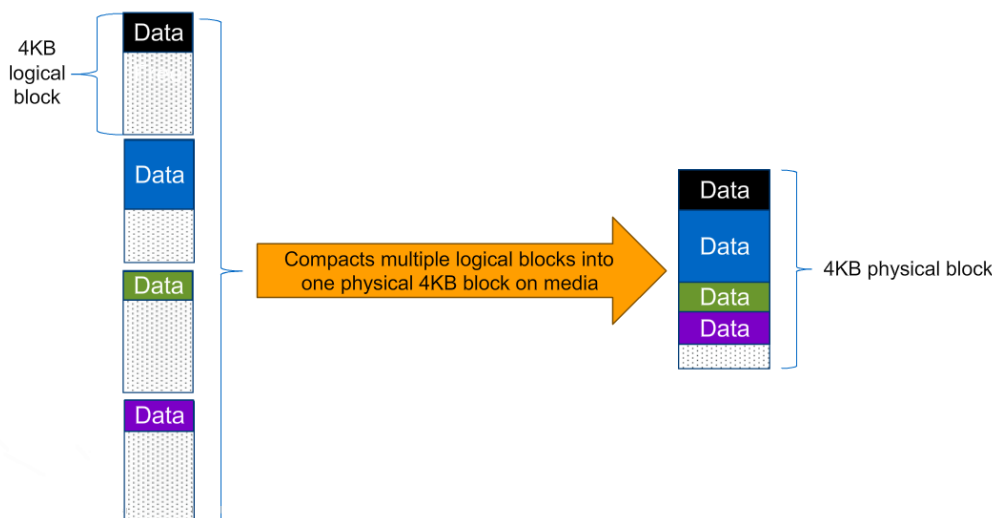
## Storage Efficiency

Storage efficiency has always been a primary architectural design point of ONTAP. A wide array of features allows businesses to store more data using less space. In addition to deduplication and compression, businesses can store their data more efficiently by using features such as unified storage, multitenancy, thin provisioning, and NetApp Snapshot™ technology.

Starting with ONTAP 9, NetApp guarantees that the use of NetApp storage efficiency technologies on AFF systems reduces the total logical capacity used to store customer data by 75%, a data reduction ratio of 4:1. This space reduction is a combination of several different technologies such as deduplication, compression, and compaction, which provide additional reduction to the basic features provided by ONTAP.

Compaction, which is introduced in ONTAP 9, is the latest patented storage efficiency technology released by NetApp. In the ONTAP WAFL® file system, all I/O takes up 4KB of space, even if it does not actually require 4KB of data. Compaction combines multiple blocks that are not using their full 4KB of space together into one block. This one block can be more efficiently stored on the disk to save space. This process is illustrated in Figure 5.

Figure 5) Storage efficiency.



## NetApp Volume Encryption

Data security continues to be an important consideration for customers purchasing storage systems. NetApp has supported self-encrypting drives in storage clusters prior to ONTAP 9. However, in ONTAP 9, the encryption capabilities of ONTAP are extended by adding an onboard key manager (OKM). The OKM generates and stores the keys for each of the drives in ONTAP, allowing ONTAP to provide all functionality required for encryption out of the box. Through this functionality, sensitive data stored on disks is secure and can only be accessed by ONTAP.

In ONTAP 9.1, NetApp has extended the encryption capabilities further with NetApp Volume Encryption (NVE). NVE is a software-based mechanism for encrypting data. It allows a user to encrypt data at the per-volume level instead of requiring encryption of all data in the cluster, thereby providing more flexibility and granularity to the ONTAP administrators. This encryption extends to Snapshot copies and FlexClone® volumes that are created in the cluster. One benefit of NVE is that it executes after the implementation of the storage efficiency features and therefore doesn't inhibit the ability of ONTAP to provide space savings to customers.

For more information about encryption in ONTAP 9.1, see the [NetApp Encryption Power Guide](#).

## NetApp Storage Virtual Machines

A cluster serves data through at least one and possibly multiple storage virtual machines (SVMs; formerly called Vservers). An SVM is a logical abstraction that represents the set of physical resources of the cluster. Data volumes and network logical interfaces (LIFs) are created and assigned to an SVM and may reside on any node in the cluster to which the SVM has been given access. An SVM may own resources on multiple nodes concurrently, and those resources can be moved nondisruptively from one node to another. For example, a flexible volume can be nondisruptively moved to a new node and aggregate, or a data LIF can be transparently reassigned to a different physical network port. The SVM abstracts the cluster hardware, and it is not tied to any specific physical hardware.

An SVM can support multiple data protocols concurrently. Volumes within the SVM can be joined together to form a single NAS namespace, which makes all of the SVM's data available through a single share or mount point to NFS and CIFS clients. SVMs also support block-based protocols, and LUNs can be created and exported by using iSCSI, FC, or FCoE. Any or all of these data protocols can be configured for use within a given SVM.


Because it is a secure entity, an SVM is only aware of the resources that are assigned to it and has no knowledge of other SVMs and their respective resources. Each SVM operates as a separate and distinct entity with its own security domain. Tenants can manage the resources allocated to them through a delegated SVM administration account. Each SVM can connect to unique authentication zones such as Active Directory, LDAP, or NIS. A NetApp cluster can contain multiple SVMs. In this design, we have aligned our SVM design to the intended administrative functions, opting for two SVMs: one for infrastructure, one for Exchange and SharePoint workload. This allows administrators of the application to access only the dedicated SVMs and associated storage, increasing manageability and reducing risk. Larger organizations with dedicated Exchange and SharePoint teams might opt to further segregate the administrative domains.

## SAN Boot

NetApp recommends implementing SAN boot for Cisco UCS servers in the FlexPod Datacenter solution. Doing so enables the operating system to be safely secured by the NetApp AFF storage system, providing better performance. The design outlined in this report uses SCSI SAN boot.

In iSCSI SAN boot, each Cisco UCS server is assigned two iSCSI vNICs (one for each SAN fabric), which provide redundant connectivity all the way to the storage. The 10GbE storage ports, in this example, e2a and e2e, which are connected to the Cisco Nexus switches, are grouped together to form one logical port called an interface group (igroup) (in this example, a0a). The iSCSI VLANs are created on the igroup, and the iSCSI logical interfaces (LIFs) are created on iSCSI port groups (in this example, a0a-<iSCSI-A-VLAN>). The iSCSI boot LUN is exposed to the servers through the iSCSI LIF using igroups. This enables only the authorized server to have access to the boot LUN. Refer to Figure 6 for the port and LIF layout.

Figure 6) iSCSI SVM ports and LIF layout.





Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead, if a network interface becomes unavailable, the host chooses a new optimized path to an available network interface. ALUA is a standard supported by NetApp used to provide information about SCSI targets, which allows a host to identify the best path to the storage.

## 5.4 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.

vSphere 6.5 brings a number of improvements, including, but not limited to:

- Additional native features to the vCenter Server Appliance
- vSphere Web Client and fully supported HTML5 client
- VM encryption and encrypted vMotion
- Improvements to DR

For more information about VMware vSphere and its components, see the [vSphere and vSphere with Operations Management](#) web page.

### 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. In vSphere 6.5, VMware has greatly improved the functionality and availability of the vCenter Server Appliance. In addition to improved installation workflows, vCenter Server Appliance now provides greater scalability of both the vCenter web application and the platform services controller functions. This allows organizations to deploy multiple instances of each function, across multiple sites, in a single vCenter domain to provide better performance to end users. HA is also improved, with native active/passive cluster capabilities that enables a failover instance of the appliance to assume control in the event of a failure of the primary appliance or its host.

The vCenter appliance in this validation was deployed with the default configuration, but customers should deploy the vCenter configuration that meets their business requirements for availability and performance.

### vSphere Cluster Considerations

VMware vSphere 6.x 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.

Larger hosts introduce another consideration, however, because the failure domain size increases with host size. Failure of a two-CPU server results in the failure of some number of VMs, whereas a four-CPU server might affect twice as many VMs. Although failures are rare, they must still be considered when designing the infrastructure. Larger hosts with fewer hosts per cluster also provide vSphere DRS with fewer options for optimally balancing the workload, increasing the likelihood of host resource contention. In general, smaller hosts are better because they provide more granular distribution.

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.

vSphere HA provides automatic restart of VMs in the event of a host failure and should be enabled on production ESXi clusters. This results in only a short outage while the VM is restarted on another host. Host monitoring and admission control should be enabled 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. Additional capacity can be reserved to provide greater headroom for concurrent host failures or maintenance activities. While vSphere HA can provide automatic recovery of VMs, mission-critical systems such as the SQL Server used in this SharePoint validation can and should also employ guest OS or application-based clustering to make sure that these applications remain online even through a host failure event.

vSphere Dynamic Resource Scheduler (DRS) provides automatic load-balancing of CPU and memory workloads across the cluster members and should also be enabled on production ESXi clusters. DRS provides automated remediation of host resource contention and reduces the likelihood of a bully VM negatively affecting other VMs on the same host. DRS also makes sure that workloads are evenly redistributed across the cluster in the event of a host failure.

To provide the highest levels of availability and performance, DRS provides the ability to create affinity and anti-affinity rules to manage placement of specific VMs on the cluster resources. VMs that are in constant communication can benefit from an affinity rule to keep those VMs on the same host, providing the lowest possible network latency between them. For clustered applications such as the SQL Server in this validation, anti-affinity rules are used to keep the two SQL cluster nodes on separate hosts, making sure that if one host fails, the other cluster node is on a separate host and therefore remains online. Anti-affinity rules could also be used to keep SharePoint and Exchange application servers on separate hosts, again making sure that a host failure does not affect application availability.

## vSphere Networking Considerations

In this reference architecture, standard vSwitches are used for VM connectivity, management VMkernel, NFS VMkernel, iSCSI storage VMkernel, and vMotion VMkernel ports. The NFS VMkernel ports are used to mount the NetApp volumes as NFS datastores. These datastores are used to store virtual machine configuration files and virtual boot disks. NetApp recommends that applications such as SQL Server or Exchange that require additional data storage utilize the guest OS iSCSI initiator to connect the VM directly to the storage system. This reduces the storage processing CPU load on the ESXi hosts and enables more granular management of those specific datasets for backup and recovery, DR replication, and secondary processing tasks.

## 5.5 NetApp SnapCenter 3.0

NetApp SnapCenter® software is a unified, scalable platform for application-consistent data protection. SnapCenter provides centralized control and oversight while delegating the ability for users to manage application-specific backup, restore, and clone jobs. SnapCenter manages data across endpoints in the NetApp Data Fabric. You can use SnapCenter to replicate data between on-premises environments; between on-premises environments and the cloud; and between private, hybrid, or public clouds. Although SnapCenter is rapidly being updated to serve as the central point of management for all core data protection needs, it does not yet support Exchange and SharePoint. For the purposes of this validation, SnapCenter was selected to protect the base virtual machines themselves with the SnapCenter Plug-In for VMware vSphere.

## SnapCenter Plug-In for VMware vSphere

SnapCenter Plug-In for VMware vSphere is a host-side component of the NetApp storage solution offering support for virtualized databases and Windows file systems in addition to support for VMs and datastores.

## Support for Virtualized Databases and Windows File Systems

SnapCenter Plug-In for VMware vSphere provides support for virtualized applications in Windows or Linux environments when you use the SnapCenter GUI to perform backup, recovery, and cloning operations. SnapCenter reduces the complexity of protecting virtualized applications by natively leveraging SnapCenter Plug-In for VMware vSphere for all SQL, Oracle, and Windows file system data protection operations on virtual machine disks (VMDKs), raw device mappings (RDMs), and NFS datastores.

## Support for VMs and Datastores

SnapCenter Plug-In for VMware vSphere also provides a plug-in for vSphere Web Client in vCenter. The SnapCenter Plug-In for VMware vSphere GUI in vCenter allows you to perform backup, restore, and attach operations for VMs and backup and mount operations for datastores that are registered with SnapCenter.

**Note:** Windows file system data protection is for LUNs on Windows hosts and not for the Windows operating system.

## Data Fabric Integration

You can use SnapCenter and the SnapCenter Plug-In for VMware vSphere with NetApp SnapMirror® technology to create mirror copies of backup sets on another volume and with NetApp SnapVault® technology to perform disk-to-disk backup replication for archival or standards compliance.

For details, see [SnapCenter Software Resources](#).

## 5.6 Microsoft Exchange 2016

Microsoft Exchange Server 2016 is a messaging platform from Microsoft that provides e-mail and scheduling. In addition to a core collaboration and messaging services platform, Exchange Server is the standard bearer worldwide, and the latest iteration builds upon its use in Office 365. Simplifying the core deployment and architecture, Exchange 2016 combines the client access server and mailbox server roles into a single mailbox server. The mailbox server contains all the functionality previously provided by the mailbox and client access roles:

- Client access services provide authentication, limited redirection, and proxy services. Client access services are delivered using the usual client access protocols: HTTP, POP, IMAP, and SMTP.
- Mailbox services include all the traditional server components found in the Exchange 2013 mailbox server role: the back-end client access protocols, transport service, mailbox databases, and unified messaging.

The edge transport role primary job is to process all Internet-facing mail. Therefore, it is usually deployed in the perimeter network and is outside of any internal Active Directory forest. The edge transport role adds additional layers of message protection and security against viruses and spam and can apply mail flow rules (also known as transport rules) to control message flow.

For more information about the Exchange 2016 architecture, see [Exchange 2016 architecture](#).

## High Availability

The mailbox server has built-in high availability and site resiliency. Like in Exchange 2013, the database availability group (DAG) and Windows Failover Clustering are the base components for high availability and site resiliency in Exchange 2016. Up to 16 mailbox servers can participate in a single DAG.

The DAG uses database copies and replication combined with database mobility and activation to implement data center high availability and site resiliency. Up to 16 copies of each database can be maintained at any given time. One of these copies of each database can be active at any time, while the remainder of the databases are passive copies. These databases are distributed across multiple DAG

member servers. Activation Manager manages the activation of these databases on the DAG member servers.

## Active Manager

Active Manager manages the health and status of the database and database copies. It also manages continuous replication and mailbox server high availability. Mailbox servers that are members of a DAG have a primary Active Manager (PAM) role and a standby Active Manager (SAM) role. Only one server in the DAG runs the PAM role at any given time.

The PAM role determines which database copies are active and which are passive. The PAM role also reacts to DAG member server failures and topology changes. The PAM role can move from one server to another within a DAG, so there is always a DAG member server running the PAM role.

The SAM role tracks which DAG member server is running the active database copy and which one is running the passive copy. This information is provided to other Exchange components such as the client access service and the transport service. The SAM also tracks the state of the databases on the local server and informs the PAM when database copy failover is required.

## Site Resiliency

Site resiliency can be implemented by stretching a DAG across two data center sites. This is achieved by placing mailbox servers in two sites. Multiple copies of each database are deployed on DAG members in both sites to facilitate mailbox database availability in all sites. Database activation controls which site has the active database. The DAG replication keeps the database copies synchronized.

## Exchange Clients

Exchange 2016 mailboxes can be accessed by a variety of clients. These clients run in web browsers, mobile devices, and computers. Most clients access their mailboxes through the virtual directory, which is presented by the Internet information service that runs on the mailbox server.

New in Exchange 2016, session identity is no longer as important. This is a result of recombining the client access and mailbox server roles. Previous versions of Exchange could suffer performance issues if the client connected to the wrong server. In Exchange 2016, the client is always connected to the server that hosts the active copy of the database. When a client connects to a mailbox server, the server checks to see where the mailbox is and proxies the connection to the server that is currently hosting the primary copy of the given database.

## Namespace Planning

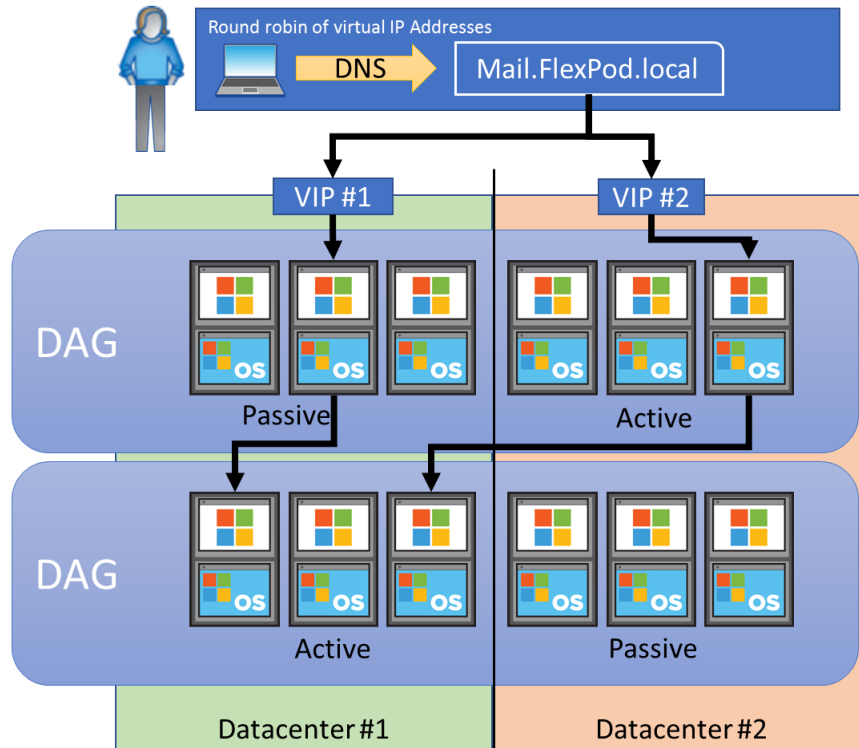
Exchange 2016 dramatically simplifies the namespace requirements when compared to earlier versions of Exchange, primarily due to the unification of the client access and mailbox roles. Unlike previous releases, the protocol services for a given mailbox are always hosted by the server hosting the active copy of the mailbox. In addition to removing the RCP client access namespace, each mailbox server also proxies any client connection for a mailbox hosted on a remote server. This proxy logic is not bound to Active Directory site boundaries, which means that a mailbox server in any given Active Directory site can proxy a session to a mailbox server that is located in a different Active Directory site. Therefore, unique namespaces are no longer required for each data center (mail1.flexpod.local and mail2.flexpod.local); instead, only a single namespace is needed for the data center pair (mail.flexpod.local). This also means that failback namespaces are not required during DAG activation scenarios, and therefore mailpri.flexpod.local and mailsec.flexpod.local are also removed.

With this core functionality in mind, various namespace models are commonly used with Exchange 2016 to achieve various functional goals. It's important to understand why a certain design is used.

## Unified Namespace

The unified namespace, also referred to as the unbound model, is the simplest to implement. It can be used in single or multiple data center deployments. The namespace is tied to one or more DAGs. In the case of multiple data centers with DAGs that span the data centers, the namespace also spans the data centers. In this namespace model, typically all mailbox servers in the DAG have active mailbox database copies; however, all active copies could reside in a data center. In either case, Exchange clients connect to the Exchange Servers irrespective of the location of the Exchange Server or the location of the client. Figure 7 illustrates the unified namespace (unbound model).

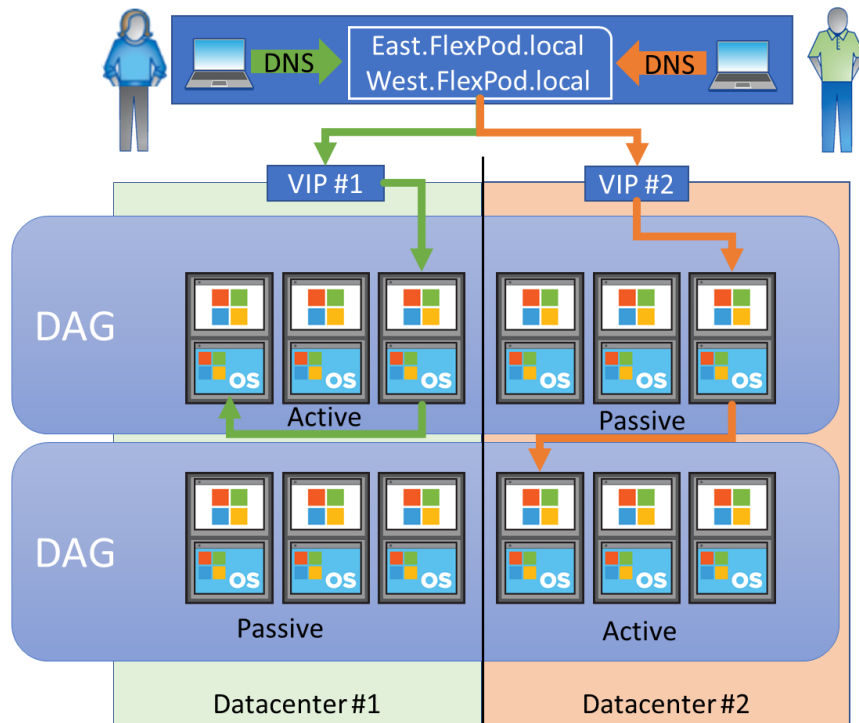
Figure 7) Unified namespace (unbound model).



## Dedicated Namespace

The dedicated namespace or bound model is associated with a specific data center or geographical location. This namespace usually corresponds mailbox servers in one or more DAGs in that location. Connectivity is controlled by where which data center has the mailbox server with the active database. Dedicated namespace deployments typically use two namespaces for each data center. One is the primary namespace that is used during normal operation, and the other is a failover namespace that is used when service availability is transferred to a partner data center. Switchover to the partner data center is an administrator-managed event in this case. Figure 8 illustrates the dedicated namespace (bound model).

Figure 8) Dedicated namespace (bound model).



### Autodiscover Namespace

Exchange 2016 continues to leverage the Autodiscover service for client profile configuration; therefore, the autodiscover.flexpod.local namespace remains.

### Internal and External Namespace

Internal and external namespace is typically used in combination with a split-DNS scheme for providing different IP address resolutions for a given namespace based on the client connection point. This is commonly used to provide different connection endpoints for clients that are connected on the external side of the firewall as compared to the internal side of the firewall.

### Regional Namespace

Regional namespace provides a method to optimize client connectivity based on client proximity to the mailbox server hosting their mailbox. Regional namespace can be used with both unified namespace and dedicated namespace schemes. For example, a company that has data centers in Europe for Europe-based employees and data centers in America for America-based employees can deploy different namespaces for both countries.

### Network Load Balancing

Network load balancing enables scalability and high availability for the mailbox servers. Scalability and redundancy are enabled by deploying multiple mailbox servers and distributing the Exchange client traffic between these mailbox servers.

Exchange 2016, like Exchange 2013 before it, has several options for implementing network load balancing. Session affinity is still not required at the network load-balancing level, although it can still be implemented to achieve specific health probe goals.

Health probe checking enables the network load balancer to verify which mailbox server is servicing specific Exchange client connectivity protocols. The health probe checking scheme determines the granularity of detecting protocol availability on the mailbox server. Exchange 2016 has a virtual directory for health checking. This directory is Exchange client specific and can be used by load balancers to verify the availability of a protocol on a client access service.

## Common Namespace and Load-Balancing Session Affinity Implementations

There are four common network load-balancing implementations that can be used for load balancing client connections to Exchange 2016. Each implementation has pros and cons for simplicity, health probe granularity, and network load balancer resource consumption:

- Layer 4 single namespace without session affinity
- Layer 7 single namespace without session affinity
- Layer 7 single namespace with session affinity
- Multiple namespace without session affinity

Table 3 lists the pros and cons for each load-balancing implementation.

Table 3) Load-balancing implementation pros and cons.

	Pros	Cons
Layer 4 single namespace without session affinity	<ul style="list-style-type: none"> <li>• Single namespace</li> <li>• Reduced load balancer complexity</li> <li>• Session affinity maintained at mailbox</li> </ul>	<ul style="list-style-type: none"> <li>• Per-server health</li> </ul>
Layer 7 single namespace without session affinity	<ul style="list-style-type: none"> <li>• Single namespace</li> <li>• Per-protocol health</li> </ul>	<ul style="list-style-type: none"> <li>• SSL offloading, which might affect load balancer scalability</li> </ul>
Layer 7 single namespace with session affinity	<ul style="list-style-type: none"> <li>• Single namespace</li> <li>• Per-protocol health</li> </ul>	<ul style="list-style-type: none"> <li>• Session affinity maintained at load balancer</li> <li>• Increased load balancer complexity</li> <li>• Reduced load balancer scalability</li> </ul>
Multiple namespace without session affinity	<ul style="list-style-type: none"> <li>• Per-protocol health</li> <li>• Session affinity maintained at mailbox</li> <li>• Users only have to know OWA URL</li> </ul>	<ul style="list-style-type: none"> <li>• Multiple namespaces</li> <li>• Additional names on certificate</li> <li>• Increased rule set on load balancer</li> <li>• Multiple VIPs</li> </ul>

### Layer 4 Single Namespace Without Session Affinity

This is the simplest implementation and consumes the fewest load balancer resources. This implementation uses a single namespace and layer 4 load balancing. Health probe checks are performed based on IP address, network port, and a single Exchange client virtual directory for health checks. Because most Exchange clients use HTTP and thus the same HTTP port, the health check can be performed on just one Exchange client protocol that is in use. The most frequently used Exchange client protocol is usually selected for health probe checking in this implementation. When the health probe fails, the network load balancer removes the entire mailbox server from the mailbox server pool until the time the health check returns to a healthy state.

## Layer 7 Single Namespace Without Session Affinity

This implementation uses a single namespace and layer 7 load balancing. The load balancer performs SSL termination for the Exchange clients and forwards the client traffic to the protocol-specific URLs. Health probe checks are configured and performed for each Exchange client protocol virtual directory. Because the health probe check is Exchange client protocol specific, the load balancer is capable of identifying and removing just the unhealthy protocols on a given Exchange Server from the rotation.

## Layer 7 Single Namespace with Session Affinity

This implementation is like the previous layer 7 single namespace implementation with the exception that the session affinity is also implemented.

## Multiple Namespace Without Session Affinity

This implementation is like the previous layer 4 implementation without session affinity with the exception that an individual namespace is used for each Exchange client protocol type. This method provides the ability to configure a health check probe for each Exchange client protocol and thus gives the load balancer the capability to identifying and removing just the unhealthy protocols on a given mailbox server from the mailbox server pool rotation.

## Exchange Server Sizing

Exchange Server needs to be sized for the projected workloads and service-level agreements. Whether Exchange is running in virtual machines or on physical servers, the Microsoft Exchange 2016 Server Requirements Calculator is an essential tool for planning and sizing the Exchange deployment. The Exchange sizing tool is available at <http://gallery.technet.microsoft.com/Exchange-2013-Server-Role-f8a61780>.

## Exchange 2016 Server Requirements Calculator Inputs

The inputs listed in Table 3 through Table 10 are configured in the Exchange 2016 Server requirements of the Exchange sizing tool for this deployment.

Table 3) Requirements for Exchange environment configuration.

Exchange Environment Configuration	Value
Exchange Server version	2016
Global catalog server architecture	64-bit
Server multirole configuration (MBX+CAS)	Yes
Server role virtualization	Yes
High-availability deployment	Yes

Table 4) Requirements for mailbox database copy configuration.

Mailbox Database Copy Configuration	Value
Total number of HA database copy instances (includes active copy) within DAG	3
Total number of lagged database copy instances within DAG	0
Number of HA database copy instances deployed in secondary data center	1
Total number of lagged database copy instances within DAG	0



**Note:** For the calculator to accept the configuration, there must be at least three copies of any database. To account for this, configure an active/passive site resiliency model with one copy on the secondary site. When deploying the secondary site, you should fill the sizing spreadsheet using the planned copy allocation. In this architecture, we are comfortable running two copies of each database because we are using SnapManager backups for any additional availability.

**Table 5) Requirements for tier 1 user mailbox configuration.**

Tier 1 User Mailbox Configuration	Value
Total number of tier 1 user mailboxes per environment	10,000
Projected mailbox number growth percentage	0%
Total send/receive capability per mailbox per day	150 messages
Average message size (KB)	75
Initial mailbox size (MB)	2,048
Mailbox size limit (MB)	5,120
Personal archive mailbox size limit (MB)	0
Deleted item retention window (days)	14
Single item recovery	Enabled
Calendar version storage	Enabled
Multiplication factor user percentage	100%
IOPS multiplication factor	1.00
Megacycles multiplication factor	1.00
Desktop search engines enabled (for online mode clients)	No
Predict IOPS value?	Yes

**Table 6) SPECint test details.**

System	Results	Baseline	Number of Cores	Number of Chips
Cisco UCS B200 M4 (Intel Xeon E5-2690 v4, 2.60GHz)	1,390	1,330	28	2

**Note:** Cisco UCS B200 M4 host servers with dual Intel Xeon E5-2690v4 processors are running the Exchange virtual machines. The SPECint 2006 rate value is 1,330 for this server and processor combination.

**Table 7) Requirements for server configuration.**

Server Configuration	Processor Cores per Server	SPECint 2006 Rate Value
Mailbox server guest machines	12	596

Each Exchange Server virtual machine is allocated 12 vCPUs to allow other virtual machines also to run on the same host. CPU oversubscription is not recommended for Exchange Server virtual machines. This means that all virtual machines that run on the same host that runs the Exchange virtual machines must not exceed a 1-to-1 vCPU-to-CPU core allocation.

To account for the fact that we are running other workloads alongside Exchange, we must prorate the server configuration sizing. Specifically, we need to calculate the fraction of the SPECint 2006 rate that we can assign to the Exchange calculator:

$$((\text{SPECint result}) / (\text{CPU core count})) \times \text{processor cores allocated} = \text{SPECint 2006 rate value}$$

The following example uses the selected configuration:

$$(1,390 / 28) \times 12 = 596$$

**Table 8) Requirements for processor configuration.**

Processor Configuration	Value
Hypervisor CPU adjustment factor	10%

The default 10% hypervisor overhead factor is used to account for the hypervisor overhead.

**Table 9) Requirements for database configuration.**

Database Configuration	Value
Maximum database size configuration	Custom
Maximum database size (GB)	8,192
Automatically calculate number of unique databases per DAG	Yes
Calculate number of unique databases per DAG for symmetrical distribution	No

The maximum database size was set to 8TB within NetApp vendor guidance.

**Table 10) Transport configuration.**

Transport Configuration	Value
Message queue expiration (days)	2
Safety net expiration (days)	4

## Exchange 2016 Server Requirements Calculator Output

The Role Requirements tab shows the following parameters for this deployment.

Two Active Directory domain controllers with at least five CPU cores are required for this deployment. Multiple domain controllers must be deployed for high availability.

Table 11 through Table 20 list the Exchange 2016 Server requirements.

**Table 11) Requirements for the process core ratio.**

Process Core Ratio Requirements	Per Primary Data Center
Recommended minimum number of global catalog cores	5

**Table 12) Requirements for environment configuration.**

Environment Configuration	Per Primary Data Center
Recommended minimum number of dedicated client access servers	–
Number of DAGs	–

Environment Configuration	Per Primary Data Center
Number of active mailboxes (normal run time)	10,000
Number of mailbox servers per DAG	8
Number of lagged copy servers per DAG	0
Total number of servers per DAG	8

**Table 13) Requirements for user mailbox configuration.**

User Mailbox Configuration	Tier 1
Number of user mailboxes per environment	10,000
Number of mailboxes per database	313
User mailbox size within database	5500MB
Transaction logs generated per mailbox per day	30
IOPS profile per mailbox	0.10
Read:write ratio per mailbox	3:2

**Table 14) Requirements for database copy instance configuration.**

Database Copy Instance Configuration	Per Primary Data Center
Number of HA database copy instances per DAG	2
Number of lagged database copy instances per DAG	0
Total number of database copy instances	2

**Table 15) Requirements for database configuration.**

Database Configuration	Value
Number of databases per DAG	32
Recommended number of mailboxes per database	313
Available database cache per mailbox	10.49MB

**Table 16) Requirements for database copy configuration.**

Database Copy Configuration	Per Server	Per DAG
Number of database copies	8	64

**Table 17) Requirements for server configuration.**

Server Configuration	Per Primary Data Center Server (Single Failure)
Recommended RAM configuration	64GB
Number of processor cores utilized	7
Server CPU utilization	54%

Server Configuration	Per Primary Data Center Server (Single Failure)
Server CPU megacycle requirements	17,557
Server total available adjusted megacycles	32,533
Possible storage architecture	RAID
Recommended transport database location	System disk

**Note:** Each Exchange virtual machine requires 64GB of physical RAM.

The projected peak CPU utilization for each virtual machine running on the Cisco UCS B200 M4 host is 54%. In this instance, we opted to increase the primary mailbox server pool by one. The selected workload is supported on seven servers, but we opted to deploy a dedicated witness instead of node majority.

**Table 18) Requirements for details of disk space.**

Disk Space Requirement	Per Database	Per Server	Per DAG
Transport database space required		322GB	3862GB
Database space required	1678GB	13428GB	161133GB
Log space required	44GB	354GB	4248GB
Database volume space required	–	19788GB	237459GB
Log volume space required	–	373GB	4472GB
Restore volume space required	–	1813GB	21761GB

The following disk space and IOPS requirement information is used for sizing the storage array. NetApp has a sizing tool specific to Exchange 2016 that uses this information to identify the storage configuration requirements.

The transport database maximum size is expected to be no greater than 322GB. This database is stored on the system disk.

**Table 19) Requirements for details of host I/O and throughput.**

Host I/O and Throughput Requirements	Per Database	Per Server	Per DAG
Total database required IOPS	38	302	3,618
Total log required IOPS	8	64	769
Database read I/O percentage	60%	–	–
Background database maintenance throughput requirements	1.0MBps	8MBps	96MBps

**Table 20) Transport calculation details.**

Transport Calculations	Value
Calculated safety net hold time (days)	4
Messages per day	1,500,000
Messages per day per server	375,000

Transport Calculations	Value
Shadow effect	750,000
Safety net effect	4,500,000
Transport database size per server	322GB

## Exchange Virtual Machines

The Exchange Servers are configured with the following parameters. The boot disk is sized to accommodate a 32778MB swap file and the transport database, in addition to the Windows system files and Exchange application files.

Table 21) Requirements for Exchange Servers.

Requirements	Value
VM count	8
vCPU count	8
RAM (no dynamic memory)	64GB
NIC1	VM network
Boot disk size	450GB

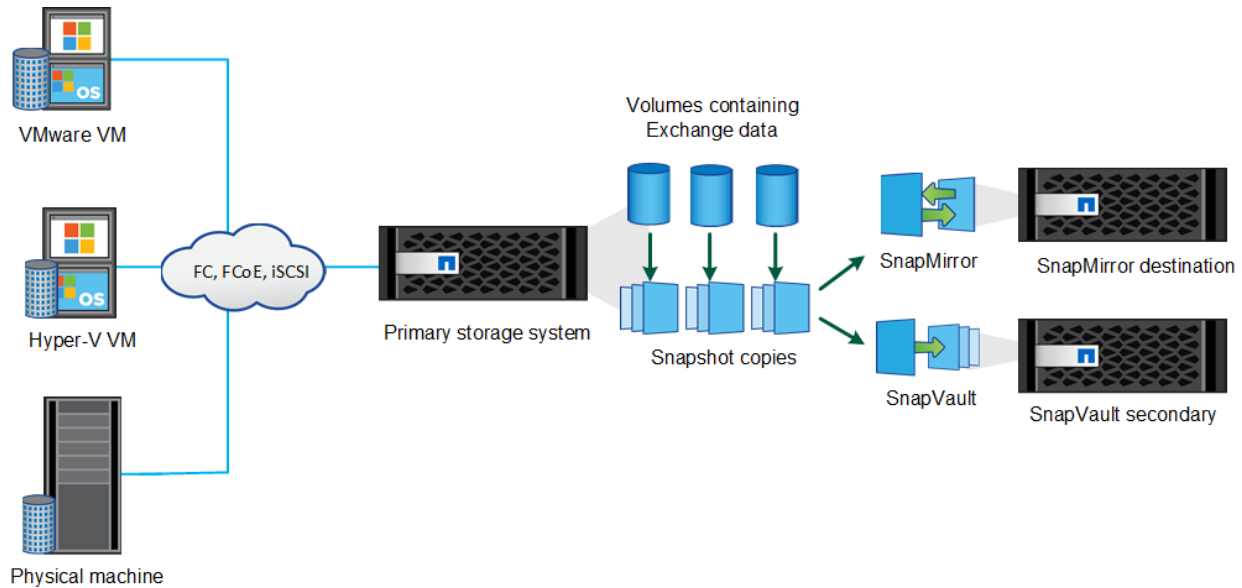
The Exchange Server virtual machines must be spread across different hosts to avoid a host failure from affecting multiple Exchange VMs. This rule also applies to domain controllers. However, a domain controller VM and an Exchange VM can run on the same host.

The deployment in this document is designed to support 10,000 5GB mailboxes with a 150 message per mailbox per day action profile. Each Exchange Server is deployed in a virtual machine. The Exchange Servers are in a single data center and have a requirement for high availability. To address the high availability, two copies of every mailbox database are deployed. An optional third site copy was documented but is not mandatory.

## 5.7 SnapManager for Exchange

SnapManager for Microsoft Exchange Server is a host-side component of NetApp's integrated storage solution for Microsoft Exchange, offering application-aware primary Snapshot copies of Exchange databases. You can use SnapManager with NetApp SnapMirror technology to create mirror copies of backup sets on another volume and with NetApp SnapVault product to archive backups efficiently to disk. Together, these tools offer a complete data protection scheme based on Snapshot that is as scalable, reliable, and highly available as the underlying storage system. Figure 9 shows the components in a SnapManager deployment with ONTAP.

Figure 9) SnapManager for Exchange components.



## SnapManager Highlights

SnapManager features seamless integration with Microsoft products on the Windows host and with NetApp Snapshot technology at the back end. It offers an easy-to-use, wizard-based administrative interface:

- **Integration with the Microsoft Volume Shadow Copy Service (VSS)** makes sure that write requests are frozen and write caches flushed before backups are created. SnapManager provides full support for Windows Volume Manager, Windows Server Failover Clustering, Microsoft Multipath I/O (MPIO), and Exchange database availability groups.
- **Fast, nondisruptive Snapshot technology** using NetApp SnapDrive for Windows software lets you back up databases in seconds and restore them in minutes without taking the Exchange Servers offline. Snapshot copies consume minimal storage space. You can store up to 255 copies per volume.
- **Automated central administration** offers hands-off, worry-free data management. You can schedule routine Exchange Server database backups, configure policy-based backup retention, set up point-in-time and up-to-the-minute restores, and proactively monitor your Exchange Server environment with periodic e-mail alerts. PowerShell cmdlets are available for easy scripting of backup and restore operations.

In addition to these major features, SnapManager offers the following:

- Integrated single mailbox recovery, which lets you restore individual mailboxes, e-mail messages, attachments, calendar items, deleted items, drafts, or contacts
- Simplified migration of existing databases to NetApp storage with an easy-to-use configuration wizard
- Nondisruptive, automated backup verification
- Fast reseeding of databases in a database availability group
- Support for physical and virtualized infrastructures
- Support for iSCSI, Fibre Channel, and FCoE
- Support for service-level role-based access control (RBAC)

## SnapManager Backup Overview

SnapManager uses NetApp Snapshot technology to create online, read-only copies of databases. It uses an Exchange Server tool to verify the integrity of the backups.

SnapManager backs up a database by creating Snapshot copies of the volumes in which the following reside:

- Database data files
- SnapInfo directories and transaction logs

Together these Snapshot copies compose a backup set. SnapManager uses a backup set to restore a database. After SnapManager backs up your databases, it can perform an integrity verification of the backup sets. SnapManager uses the Exchange System Management Tools to check the database and transaction log files for physical and logical corruption. Verification makes sure that you can use backup sets to restore databases as needed.

**Note:** Database verification is disabled by default if you have a database availability group (DAG). Verification is not required for DAG databases that have at least two copies, each of which has a viable backup.

**Note:** SnapManager cannot restore databases from Snapshot copies created by ONTAP or SnapDrive. You should perform backups using SnapManager only.

## Backup Best Practices

It is important to consider the following factors for planning a backup strategy for Microsoft Exchange data in the organization:

- **Organization SLA.** This parameter determines the frequency and the type of backups.
- **Backup retention planning.** This parameter determines whether backup sets must be retained on the primary or the secondary site.
- **Backup verification policy.** This parameter determines when backup verification is engaged.

The time required to restore Microsoft Exchange data during an outage depends on the number of transaction logs that must be replayed. Therefore, reducing the number of transaction logs that must be replayed when restoring from a full backup is important. The only way to reduce the number is to create more frequent backups.

To help achieve the desired SLA and RPO times, SME has frequent recovery points (FRPs). FRPs are optimized backup sets that are created through SME. The backup sets contain only the transaction log files that have been created since the last full backup or FRP backup that was created. The transaction log files are hard-linked or copied into the SnapInfo directory, and then a Snapshot copy is created. An FRP backup set contains a small amount of information. Backups can be created as often as every 10 minutes. Having a higher frequency of FRP backups reduces RPO times.

## Backup Retention

Your backup retention strategy needs to balance storage efficiency with restore needs. You can specify that SnapManager automatically delete older backups or transaction logs, or you can delete these items explicitly.

**Note:** You should not use SnapDrive or storage system administration tools to delete Snapshot copies created by SnapManager. Doing so leaves behind unwanted data that cannot be removed.

## SnapManager Restore Operations

There are two types of SnapManager Restore operations:

- **Point-in-time restore operation.** In this case, only the uncommitted transaction logs that existed in the active file system at the time the backup copy was created are replayed. All the transaction logs beyond the point in time (after the backup copy was created) that exist in the transaction log directory and that belong to the restored database are removed. You can use this method to restore a database back to a time before a corruption occurred.
- **Up-to-the-minute restore operation.** There are some transaction logs that are not committed to the databases. In an up-to-the-minute restore operation, all uncommitted transaction logs, from the time the backup set was created up to the most current time, are played forward and applied to the databases. This includes transaction logs from the backup sets, in addition to the transaction logs in the transaction log directory. A contiguous set of transaction logs is required for an up-to-the-minute restore operation to succeed. This option is selected by default.

## NetApp Single Mailbox Recovery for Microsoft Exchange Server

Recovering single mailboxes, messages, or other items from Microsoft Exchange Server can be time-consuming for Exchange administrators and expensive for the enterprise. In this design, we solve that problem by using NetApp Single Mailbox Recovery(SMBR) in conjunction with SnapManager for Exchange.

The following are some of the key benefits that SMBR provides in this solution:

- **Improve availability.** SMBR helps to restore single mailboxes, folders, e-mails, or attachments quickly and securely from a Microsoft Exchange Server enterprise messaging system.
- **Enhanced compliance.** Speeds up legal discovery with search capabilities that are invaluable in meeting compliance requirements.
- **Reduced capital costs.** Eliminates the need for a separate recovery server and storage by restoring individual Microsoft Exchange items directly to a production Exchange Server or a Microsoft Outlook personal storage table (PST) file.
- **Improved productivity.** Reduces time to locate individual items in an archive Exchange database (EDB) file.

## 5.8 Microsoft SharePoint 2016

Microsoft SharePoint Server 2016 is a collaboration environment that organizations of all sizes can use to increase the efficiency of business processes. It is an extensible and scalable web-based platform consisting of tools and technologies that support collaboration and sharing of information within teams, throughout the enterprise and on the web. Microsoft SharePoint turns users into participants, allowing users to easily create, share, and connect with information, applications, and people. Search features in SharePoint 2016 enable users to find content efficiently regardless of the physical location of data.

Microsoft SharePoint Server 2016 also introduces the concept of MinRole. MinRole is a new farm topology that allows a SharePoint farm administrator to define each server's role in a farm topology. The role of a server is specified when you create a new farm or join a server to an existing farm. Using the concept of MinRole greatly simplifies the installation and optimizes the performance of the SharePoint farm, because it automatically configures all the services on each server that are required for that specific role. The following are the primary benefits of using MinRole:

- **Simplified deployment.** You no longer need to worry about which services should be started on which servers. By deploying your farm in a recommended MinRole topology, you can focus on what functionality to enable in your farm and let SharePoint take care of the rest.
- **Improved performance and reliability.** SharePoint 2016 has been optimized for the MinRole topology based on the performance data that Microsoft has collected over the years. By deploying



your farm in a recommended MinRole topology, you can reduce network latency and increase reliability.

- **Simplified capacity planning and farm scalability.** With MinRole, you can leverage more predictable and prescriptive capacity-planning guidance. You can easily scale out your farm by adding additional servers and letting SharePoint configure those servers for you.

Table 22 lists the eight predefined server roles that are available to an administrator to choose from.

Table 22) SharePoint MinRoles.

MinRole	Services
Application	Service applications, services, and components that serve backend requests (such as background jobs or search crawl requests) belong on Application servers. These servers are optimized for high throughput.
Distributed Cache	Service applications, services, and components that are required for a distributed cache belong on Distributed Cache servers.
Front-end	Service applications, services, and components that serve user requests belong on Front-end servers. These servers are optimized for high performance.
Search	Service applications, services, and components that are required for searching belong on Search servers.
Front-end with distributed cache	Shared role that puts the Front-end and Distributed Cache roles on the same server. Make sure the server meets the system requirements for hosting a shared server role.
Application with Search	Shared role that puts the Application and Search roles on the same server. Make sure the server meets the system requirements for hosting a shared server role.
Single-server farm	Service applications, services, and components required for a single machine farm belong on a Single-Server Farm. A Single-Server Farm is meant for development, testing, and very limited production use. A SharePoint farm with the Single-Server Farm role cannot have more than one SharePoint server in the farm.
Custom	Custom service applications, services, and components that do not integrate with MinRole belong on Custom servers. The farm administrator has full control over which service instances can run on servers assigned to the Custom role. MinRole does not control which service instances are provisioned on this role.

You can use a combination of servers to deploy a SharePoint 2016 Farm in your environment based on your requirements. NetApp highly recommends using multiple servers for each of the four specific roles for high availability. Note that you can only have one distributed cache instance per SharePoint Farm. You can increase the number of Front-end servers if you want a user optimized farm, or you can increase the number of search servers if you want a search optimized SharePoint 2016 farm.

In addition to the concept of MinRole, here are a few of the additional features that were introduced in SharePoint 2016:

- Ability to deploy Access Services
- Resource-based URLs that can retain links, when documents are renamed or moved in SharePoint
- Hybrid Deployment that enables you to integrate your on-premises farm with Office 365 productivity experiences
- Support for Open Document Format (ODF) files
- SharePoint Business Intelligence support for SQL Server 2016 CTP 3.1 and Power Pivot add-in and Power View

- Sharing and Search improvements over the SharePoint 2013

SharePoint 2016 uses SQL Server databases to store all the data, including content, configuration, and metadata. SharePoint Server 2016 databases are one of the two types:

- System Databases – Automatically created when you run the SharePoint 2016 Products Configuration Wizard.
- Service Application Databases – Automatically created when you deploy a service application in your farm and when you choose a server role in the MinRole feature.

To make sure that both the system databases and service application databases are always available, we use the concept of AlwaysOn availability groups in SQL Server 2016 in this design. The AlwaysOn availability group feature is a high-availability and disaster recovery solution that provides an enterprise-level alternative to database mirroring. It maximizes the availability of a set of user databases, in our case the SharePoint databases. 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. 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.

In our design, we created two SQL Server 2016 instances on the primary stack and a single SQL Server instance on the secondary stack. We configured the availability group to do synchronous replication between the two instances on the primary stack and asynchronous replication on the instance on the secondary stack. All three SQL instances can be configured as read copies. You can design your underlying SQL instances based on your best practices, but NetApp highly recommends using AlwaysOn availability groups to make the SharePoint databases reliable.

## 5.9 DocAve Manager

AvePoint DocAve software can monitor and manage growing Microsoft SharePoint Server environments across the entire enterprise. By using ONTAP, administrators can perform routine maintenance and upgrades without disrupting storage operations. DocAve Tiered Storage capabilities enables the automatic placement of files and BLOBs on SMB (CIFS) shares outside of the content database to improve the scalability of large SharePoint deployments. With DocAve, SharePoint Server data can be backed up in minutes and recovered at any granularity—from individual documents up to an entire SharePoint farm—also within minutes, from local or remote backups. DocAve automates backup and restore workflows, FAST search information, and certificates.

DocAve and SnapMirror software together can simplify remote replication of SharePoint data, enabling a reliable disaster recovery strategy.

For more information about utilizing DocAve, see [AvePoint and NetApp solutions for Microsoft SharePoint](#).

## 6 Design Considerations

Table 23 lists best practices recommended by NetApp for designing or implementing a multiapplication deployment for Microsoft Exchange and Microsoft SharePoint running on a FlexPod Datacenter solution.

Table 23) Design considerations.

Design Aspect	Design Considerations
Network switch series and model	As with any FlexPod installation, both Cisco Nexus 9000 and Cisco Nexus 5000 series network switches can be used with this design. The primary considerations for the switch series are as follows:

Design Aspect	Design Considerations
	<ul style="list-style-type: none"> <li>• The Cisco Nexus 9000 series is the latest hardware platform and supports both standalone (traditional) NX-OS and ACI. If organizations are considering implementing ACI, the 9000 series should be the default choice. The 9000s do not support FC or FCoE. Therefore, all SAN boot and other storage access must be through IP protocols unless the storage controllers are connected directly to the Cisco UCS fabric interconnects or Cisco MDS switches are used for SAN traffic.</li> <li>• The Cisco Nexus 5000 series supports FCoE but does not support ACI. Organizations that require FCoE but do not want a direct connect topology and have no plans for implementing ACI should use the 5000 series.</li> <li>• FlexPod components can also be connected to new or existing SAN infrastructure if additional SAN connectivity is required.</li> </ul>
Host boot device	<p>FlexPod supports three common host boot device options:</p> <ul style="list-style-type: none"> <li>• <b>Local boot.</b> Requires per-blade HDD or SSD. Use of local boot removes a key value proposition of Cisco UCS (stateless computing) but does enable host independence from, and parallel deployment with, shared storage.</li> <li>• <b>SAN boot.</b> Requires shared storage to function and forces a serial approach to deployment because such storage must be available before servers can be deployed. By far the most common FlexPod boot device, SAN boot is a cornerstone of stateless computing in Cisco UCS and provides true independence between server identity and server hardware.</li> <li>• <b>PXE booting.</b> Requires boot technology and software licensing for solutions such as VMware Auto Deploy and is dependent on that boot infrastructure being available to deploy or run servers. PXE booting provides an even more stateless computing solution than SAN boot and can be used either in conjunction with local or SAN boot or as an alternate methodology.</li> </ul>
Host SAN boot protocol	<p>There are two options for using SAN boot:</p> <ul style="list-style-type: none"> <li>• <b>FCoE.</b> Requires FC or converged adapters on the storage array. FCoE connectivity requires either FCoE-capable Cisco Nexus switches or a direct connect topology.</li> <li>• <b>iSCSI.</b> Requires either Ethernet or converged adapters on the storage array. No specific Cisco Nexus switch model or series is required.</li> </ul>
Storage controller model	<p>With the AFF A-Series, the primary considerations concern capacity and performance. The AFF A300 provides significantly more capacity and performance for a small price differential.</p> <p>With the introduction of the AFF A300, the physical rack unit footprint is no longer a consideration.</p>
Storage cluster connectivity	<p>Intercluster communication for AFF and FAS storage clusters has two topologies:</p> <ul style="list-style-type: none"> <li>• <b>Switched.</b> All cluster-member communication occurs across a redundant pair of dedicated 10GbE switches. These cluster switches must be one of a few supported models, such as the Cisco Nexus 5596, and must not be used for noncluster data traffic. A switched topology is required for clusters larger than two nodes. This topology provides the easiest transition from a two-node to a four-node or higher configuration.</li> <li>• <b>Switchless.</b> HA pair members are directly connected to each other, eliminating the need for dedicated 10GbE switches for cluster communication. A switchless topology is only supported for two-node clusters. Two-node switchless clusters can be converted nondisruptively to a switched cluster topology.</li> </ul>

Design Aspect	Design Considerations
Storage scaling	AFF/FAS clusters can scale up to 12 nodes (6 HA pairs) for SAN or hybrid SAN/NAS clusters and up to 24 nodes (12 HA pairs) for NAS-only clusters. Organizations utilizing SAN boot are limited to these 12 nodes within a single cluster, at least for the cluster providing SAN boot storage. Depending on the scale and scope of the infrastructure, a smaller cluster can provide the SAN services required for SAN boot and other block storage needs. One or more larger clusters can provide NAS storage for VM datastores and other workloads. With AFF, organizations have the flexibility of scaling out or up as needed. If performance requirements are less than capacity requirements, it is simpler to scale up (bigger SSDs or more SSDs) than out.
Storage broadcast domains	ONTAP uses the concept of broadcast domains. You can think of broadcast domains as a layer 2 failure domain. In this design, broadcast domains were created for management, NFS, iSCSI-A, and iSCSI-B traffic and then added the respective ports to those domains. This makes sure that if a port goes down, the redundant port in the same domain can take over and serve traffic, thereby avoiding a traffic disruption.
Storage virtual machines (SVMs)	SVMs contain data volumes and one or more LIFs through which they serve data to the clients. They securely isolate the shared virtualized data storage and network. This design has two SVMs: one for infrastructure VMs, one for the Exchange and SharePoint workload.
Compute fabric interconnect model	Fabric interconnect model considerations are primarily around scale. Organizations can choose the appropriate fabric interconnect model based on the number of devices to be connected and/or the bandwidth requirements of each device. The Cisco UCS 6300 series fabric interconnects are the first to provide support for 40GB networking and enable organizations to more thoroughly future proof their infrastructure.
Compute blade model	<p>Compute blades come in three form factors:</p> <ul style="list-style-type: none"> <li>• <b>Half width.</b> Supports up to dual CPUs and hundreds of gigabytes of memory (limits depending on model). Most commonly deployed form factor, with up to eight fitting in a single chassis. This configuration provides the most scale-out capabilities, minimizing the effect of host failures or maintenance activities.</li> <li>• <b>Full width.</b> Supports up to quad CPUs and more memory than half-width blades. However, this format can be problematic due to the failure domain effect of a single blade.</li> <li>• <b>Full width and double height.</b> Supports up to quad CPUs and more memory than half-width or full-width blades. These blades do not provide greater aggregate CPU or memory resources than four half-width blades, which take up the same number of chassis slots. However, the half-width provides a smaller failure domain and thus is better suited.</li> </ul>

Design Aspect	Design Considerations
VMware vCenter deployment type	<p>vCenter can be deployed either to a Windows OS or using a virtual appliance:</p> <ul style="list-style-type: none"> <li>• <b>Windows installation.</b> The only choice for organizations that want to keep their virtualization management solution on a bare-metal platform. A Windows installation can also be deployed to a Windows VM, and this installation choice provides many administrators with their most familiar methods for troubleshooting. A Windows installation is also the preferred choice for organizations with strong Microsoft SQL Server experience, particularly for database backup, because the virtual appliance cannot use Microsoft SQL Server for its database.</li> <li>• <b>vCenter Server Appliance.</b> The simplest deployment option because it requires no additional OS or antivirus licensing. The vCenter Server Appliance is also the recommended deployment method from VMware. With vSphere 6.5, the appliance has surpassed the Windows version with exclusive features such as native high availability, native backup and restore, migration tool, and improved appliance management.</li> </ul>
VMware vCenter PSC deployment options	<p>Starting with vSphere 6.0, a vCenter installation includes two constituent parts that can be installed separately or together: a platform services controller (handling single sign-on and related services) and the vCenter Server itself. Combined installation is the simplest option for deployment or future troubleshooting and interservice communication.</p>
VMware vCenter resource sizing	<p>With vSphere 6.5, vCenter can support up to 2,000 hosts and 25,000 powered on virtual machines per instance. You can refer to the vSphere 6.5 configuration maximums to get additional sizing instructions.</p>
VM anti-affinity rules	<p>Using VM anti-affinity rules makes sure that a single host failure does not bring down both the SQL Server VMs or even SharePoint or Exchange application servers. In our validation, we created these rules for SQL VMs, but you can create rules for other VMs as well.</p>
SQL AlwaysOn availability groups	<p>Configure SQL Server AlwaysOn availability groups to make the SharePoint databases highly available. In our validation, we created two standalone SQL instances on the primary stack and a standalone instance on the secondary stack. The availability group was configured to do synchronous replication between the instances on the primary stack and asynchronous replication on the secondary stack.</p>
SnapManager for Exchange backup type	<p>You can choose from two database backups:</p> <ul style="list-style-type: none"> <li>• <b>Full backup.</b> Backs up database files and truncated transaction logs. Exchange Server truncates transaction logs by removing entries already committed to the database. This is the most common backup type.</li> <li>• <b>Copy backup.</b> Backs up database files and transaction logs that have not been truncated. Use this backup type when you are backing up your databases with another backup application. Keeping transaction logs intact makes sure that any backup application can restore the databases.</li> </ul>
Enabling faster restore times using SME	<p>The more often you back up your databases, the fewer transaction logs SnapManager has to play forward at restore time, which can result in faster restores.</p>
SME backup operations	<p>SnapManager can perform one operation at a time. Do not schedule overlapping SnapManager operations.</p>

Design Aspect	Design Considerations
SME backup verification	This design uses DAG databases with three copies of data: two on the primary site and one on the secondary site. Thus, SME backup verification is not required, assuming each database copy has a viable backup.
SME up-to-the-minute restore operation	You should use this type of restore operation in the following use cases: <ul style="list-style-type: none"> <li>• To play forward all the transactions up to the most current time.</li> <li>• To restore individual databases.</li> <li>• To restore backup copies after a point-in-time restore operation of a backup copy that is not the most recent.</li> <li>• To perform an up-to-the-minute restore operation from any backup copy, including the most recent backup copy after a point-in-time restore operation of the most recent backup operation. You might lose all transactions that occurred between the time when you created the last backup copy and when you performed the point-in-time restore operation.</li> </ul>
SME point-in-time restore operation	You should use this type of restore operation in the following use cases: <ul style="list-style-type: none"> <li>• To recover the databases as they were at a point in time: for example, when the most recent backup copy was created.</li> <li>• To restore the databases to a recovery database.</li> <li>• To restore all existing backup copies after a point-in-time restoration of a backup that is not the most recent one.</li> </ul>

## 7 Solution Verification

The guidance outlined in this document was validated during the solution deployment, which is documented in the FlexPod Datacenter with Microsoft Exchange, SharePoint 2016, and NetApp AFF A300 NVA Deployment Guide. The solution verification proved that the AFF A300 was more than capable of handling the combined workload of both Exchange 2016 and SharePoint 2016 under a concurrent 10,000 user load.

## 8 Conclusion

FlexPod Datacenter is the optimal infrastructure foundation on which to deploy any business-critical application. Cisco and NetApp created a platform that is both flexible and scalable for multiple use cases and designs. This flexibility and scalability of the FlexPod architecture enable customers to start with a right-sized infrastructure that can ultimately grow with and adapt to their evolving business requirements. During the verification tests of this reference architecture, the response times from both SharePoint and Exchange were excellent.

FlexPod Datacenter with NetApp AFF performed well, reaching a combined peak IOPS of over 10K while averaging 15% CPU utilization during most operations. All test categories demonstrated that, based on the 10,000-user load, the FlexPod Datacenter with AFF A300 system is capable of supporting up to 100,000 users while still being able to perform appropriately even in the event of a storage, host, or network failure.

## Acknowledgements

This document is the result of the work, documentation, and assistance provided by Jens Dickmeis of NetApp. The authors of this document would like to thank Roger Yu of AvePoint for his contribution and support.

## Where to Find Additional Information

To learn more about the NetApp technology described in this document, refer to the following sites:

- FlexPod Converged Infrastructure product page:  
<http://www.netapp.com/us/products/converged-systems/flexpod-converged-infrastructure.aspx>
- NetApp All Flash FAS Resources page:  
<http://mysupport.netapp.com/aff/resources>
- NetApp ONTAP Resources page:  
<http://mysupport.netapp.com/ontap/resources>
- NetApp SnapCenter Software Resources page:  
<http://mysupport.netapp.com/snapcenter/resources> <http://mysupport.netapp.com/ontap/resources>
- For all other NetApp technologies, see the NetApp web site:  
<http://www.netapp.com/us/index.aspx>

## Version History

Version	Date	Document Version History
Version 1.0	September 2017	Initial release.

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

### **Copyright Information**

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

NetApp reserves the right to change any products described herein at any time, and without notice.

NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

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

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

### **Trademark Information**

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

NVA-1116-DESIGN-0917