



NetApp Verified Architecture

FlexPod Select for High-Performance Oracle RAC

NVA Design

Rodrigo Nascimento, Bhavin Shah, NetApp
November 2015 | NVA-0012-DESIGN | Version 2.0

TABLE OF CONTENTS

1	Executive Summary	4
2	Program Summary	4
2.1	NetApp Verified Architecture Program	4
2.2	FlexPod Converged Infrastructure Program.....	5
3	Solution Overview	5
3.1	Solution Technology	6
3.2	Use-Case Summary.....	8
4	Solution Design	8
4.1	Compute Design Elements	8
4.2	Network Design Elements.....	9
4.3	NetApp EF560 Storage Array Elements	11
5	Solution Validation	12
5.1	Performance Test Results.....	12
6	Technology Requirements	13
6.1	Hardware Requirements	13
6.2	Software Requirements	15
7	Best Practices	16
7.1	Storage Recommendations.....	16
7.2	Oracle ASM Recommendations.....	16
7.3	Linux Recommendations	17
8	Conclusion	17
	Appendix A: Consolidated Storage Layout for an Eight-Node Oracle RAC Environment	17
	Appendix B: Results of Mixed-Workload Use Cases Using EF550 All-Flash-Array Storage Controllers	19
	References	19
	Version History	20
	NetApp Contact Information	20

LIST OF TABLES

Table 1)	Solution validation results for all use cases.....	13
Table 2)	Hardware requirements.....	15
Table 3)	Software requirements.....	15

Table 4) Storage tuning parameters.....	16
Table 5) Oracle ASM and instance settings.	16
Table 6) Linux tuning parameters.....	17
Table 7) Storage layout for eight-node RAC.....	17
Table 8) Previous solution validation results for all use cases.....	19

LIST OF FIGURES

Figure 1) FlexPod portfolio.	5
Figure 2) ASM layout for eight-node RAC.	11
Figure 3) FlexPod Select for Oracle with eight RAC nodes showing 100% 8K random reads.	13
Figure 4) FlexPod Select for High-Performance Oracle RAC solution architecture.....	14

1 Executive Summary

Data-driven companies need technology solutions that provide high availability, offer ease of management, and, perhaps most importantly, meet critical performance objectives. Oracle databases are the choice of countless enterprises to deliver critical functionality to the business or its customers. NetApp has long delivered technologies that enable the capabilities of the Oracle product suite while taking into account an organization's operational cost guidelines.

The decision to deploy an infrastructure capable of supporting demanding applications has implications for a variety of facets within an IT organization. Decisions about such things as acquisition, deployment methodology, best practices, lifecycle, growth, and so forth must be taken into account to serve the demands of users. These points must also be considered for the company to remain competitive both now and in the future. Database teams need an adequately responsive infrastructure in which no component inhibits database performance. Careful consideration should be given to the compute network and storage layers of such an infrastructure. This is especially important when SLAs or objectives must be met for applications running on Oracle databases.

The NetApp FlexPod® data center platform delivers key products including Cisco Unified Computing System (Cisco UCS) servers, Cisco Nexus networking, and industry-leading storage from NetApp, all configured and validated according to industry best practices. An industry-proven platform, FlexPod can be tailored to meet the needs of a variety of applications and use cases.

The product portfolios of Cisco and NetApp offer a wide array of choices from which to select the appropriate platform for Oracle databases. NetApp E-Series storage solutions, Cisco UCS, and Cisco Nexus components provide the foundation for a FlexPod environment capable of delivering high performance for Oracle databases. The FlexPod Select for High-Performance Oracle RAC solution accomplishes this goal by enabling fast transactions for end users and applications. This capability provides a balanced architecture that offers extreme performance and enables businesses to react quickly to new business needs or competitive threats.

2 Program Summary

2.1 NetApp Verified Architecture Program

The NetApp Verified Architecture (NVA) program offers customers a validated architecture for NetApp® solutions. NVAs provide customers with NetApp solution architecture that:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes customer deployment risks
- Accelerates customer time to market

This document is for NetApp and partner solution engineers and customer strategic decision makers. It describes the architecture design considerations used to determine the specific equipment, cabling, and configuration required in a particular environment. This document provides guidance about how to estimate workload characteristics and how to manage solution lifecycle events such as replacing equipment, updating firmware, and upgrading software. Additional considerations discussed in this document include:

- Resiliency
- Fault tolerance
- Degree of high availability
- Performance expectations

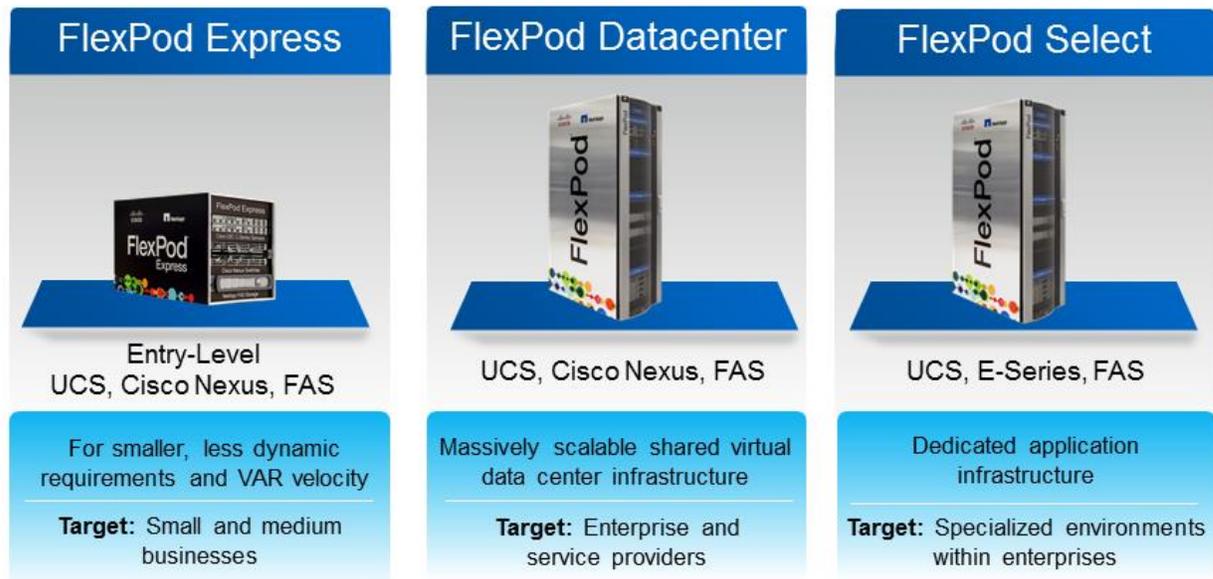
2.2 FlexPod Converged Infrastructure Program

FlexPod delivers key technologies from Cisco and NetApp. FlexPod reference architectures are often delivered in the form of Cisco Validated Designs (CVDs) or NVAs. Deviations based on customer requirements from a given CVD or NVA are encouraged, provided those variations do not result in the deployment of unsupported configurations.

All FlexPod solutions share key configuration principles, such as using converged Ethernet or SAN networks and maximizing hardware-level resiliency. As shown in Figure 1, the FlexPod program consists of three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select.

- **FlexPod Express** offers customers an entry-level solution consisting of technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an excellent multipurpose foundation for a variety of workloads and applications.
- **FlexPod Select** takes the best aspects of FlexPod Datacenter and tailors the infrastructure for a given application.

Figure 1) FlexPod portfolio.



The solution discussed in this design guide is part of the FlexPod Select family. Optimizing Oracle databases can include changes in the compute, network, and storage layers. As such, the design outlined in this document primarily pertains to architecture tied to specific performance criteria.

3 Solution Overview

This NVA validates that the NetApp FlexPod Select solution can run Oracle databases in a highly resilient architecture while remaining competitive in both price and performance. This document addresses the scenarios that customers experience in their production Oracle databases.

The solution described in this guide was architected to deliver over two million input/output operations per second (IOPS) at microsecond-level average latency, with a 100% random read workload, using an 8K request size. Additionally, NetApp validated that this solution delivers consistent performance using a variety of other mixed OLTP-type workloads, such as:

- 90% reads, 10% writes, 100% random I/O using an 8K request size
- 80% reads, 20% writes, 100% random I/O using an 8K request size

3.1 Solution Technology

FlexPod Select uses technologies from Cisco and NetApp that are configured according to the two companies' best practices. This section discusses the products and technologies applied in the solution.

NetApp EF560 Flash Array

The NetApp EF560 flash array is designed for performance-driven applications with microsecond-level latency requirements.

This solution uses the NetApp EF560 flash array as the underlying storage technology. The array is built on storage architecture developed by experts with more than 20 years of storage development experience; there are more than 750,000 systems in the field. Each EF560 flash array can deliver extreme performance with microsecond-level response times, enabling business-critical applications to deliver faster results and improving the end-user experience. This combination of high IOPS and ultralow latency makes the EF560 flash array an ideal choice for database-driven applications that require extreme performance.

The EF560 flash array runs on the enterprise-proven NetApp SANtricity® platform, which is optimized for flash solutions and allows storage administrators to achieve maximum performance and capacity utilization. The extensive configuration flexibility, custom performance tuning, and complete control over data placement make it an ideal choice for mission-critical applications. Its GUI-based performance tools provide key information on storage input/output (I/O) from multiple viewpoints, allowing administrators to make informed decisions about configuration adjustments to further refine performance.

The NetApp EF560 flash array delivers extreme performance, reliability, and availability to drive greater speed and responsiveness from the applications controlling your key business operations.

The NetApp EF560 flash array can:

- Increase the speed of business with microsecond-level response times.
- Eliminate overprovisioning and improve IT efficiency.
- Achieve the transactional performance of 1,000 15K-RPM drives in a two-rack-unit enclosure that requires just 5% of the available rack space, power, and cooling as compared to storage systems that run on spinning disks.
- Detect and resolve issues quickly with advanced monitoring and proactive repair.
- Protect against data loss and downtime with NetApp point-in-time images, remote replication, and other advanced data protection.
- Create copies of the database by using the NetApp Snapshot® volume feature.
- Replicate data to either an EF560 flash array or an E-Series system.
- Leverage the enterprise-proven SANtricity software platform.

By combining extreme IOPS, microsecond-level response times, scale-up capacity, and enterprise-grade reliability, the NetApp EF560 flash array helps you to increase productivity and achieve faster business results.

Cisco Unified Computing System

The Cisco Unified Computing System is a next-generation solution for blade- and rack-server computing. The system integrates a low-latency, lossless 10-Gigabit Ethernet (10GbE) unified network fabric with enterprise-class x86 architecture servers. The system is an integrated, scalable, multichassis platform in which all resources participate in a unified management domain. Cisco UCS accelerates the delivery of

new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems.

Cisco Nexus 5000

Cisco Nexus 5000 series switches are designed to deliver high-density, top-of-rack layer 2 and layer 3 10GbE with unified ports in compact one- and two-rack-unit form factors. The Cisco Nexus 5000 series includes Cisco Nexus 5500 and 5600 platforms as part of the Cisco Unified Fabric portfolio.

The Cisco Nexus 5500 switches simplify convergence through broad connectivity support. This feature makes them ideal top-of-rack access switches for traditional and converged deployments. The Cisco 5500 switches are designed to meet the scalability demands of today's data centers. Key Cisco Nexus 5500 series features include:

- Up to 1,152 ports in a single management domain that uses Cisco Fabric Extender (FEX) architecture
- Up to 96 unified ports

Oracle Database

The Oracle Database 12c R1 Enterprise Edition provides industry-leading performance, scalability, security, and reliability on clustered or single servers with a wide range of options to meet the business needs of critical enterprise applications. Oracle Real Application Cluster (RAC) brings an innovative approach to the challenges of rapidly increasing amounts of data and demand for high performance. Oracle RAC uses a scale-out model in which active-active clusters use multiple servers to deliver high performance, scalability, and availability.

Oracle Automatic Storage Management (Oracle ASM) provides an integrated cluster file system and volume-management features that remove the need for third-party volume management tools and reduce the complexity of the overall architecture.

Some of the key Oracle ASM features include:

- Automatic file and volume management
- Database file system with performance of raw I/O
- Automatic distribution and striping of data
- A choice of external (array-based) data protection, two-way and three-way mirror protection
- Control over which copy of mirrored data should be used preferentially

With these capabilities, Oracle ASM provides an alternative to the third-party file system and volume-management solutions for database storage management tasks. These tasks include creating or laying out databases and managing the use of disk space. Oracle ASM provides I/O load balancing across all LUNs or files in an Oracle ASM disk group by distributing each data file's contents evenly across the storage pool in the disk group.

The NetApp SANtricity plug-in for Oracle Enterprise Manager (Oracle EM) provides Oracle database administrators (DBAs) with powerful capabilities designed to increase their productivity and simplify their jobs. The plug-in is designed to access E-Series and EF-Series storage arrays used in conjunction with Oracle EM database software. Doing so enables Oracle DBAs to monitor and report on the storage subsystems, with the goal of confirming the performance and availability of the infrastructure they use. Performance views that come with the plug-in help DBAs easily identify bottlenecks in the system. The plug-in also gives a view of the end-to-end database mapping to the storage and enables DBAs to create a database-to-storage topology report without accessing the storage layers underneath. The plug-in is free and does not require a license.

Key features of the Oracle EM plug-in include:

- Integration with Oracle Enterprise Manager 12c

- Support for NetApp E-Series and EF-Series storage arrays
- End-to-end storage volume-to-database mapping
- Integrated business intelligence publisher reports
- Automatic metric collection on key storage array components
- Integrated database performance home page

Oracle Linux

Oracle Linux brings the latest Linux innovations to market, delivering extreme performance, advanced scalability, and reliability for enterprise applications and systems along with worldwide, enterprise-class, low-cost support. It is free to download and distribute, including patches and updates. It is certified for compliance with the Linux Standard Base (LSB) standard. Oracle Linux is completely free to download, deploy, and distribute. Oracle Linux Support delivers enterprise-class support for Linux with Ksplice zero-downtime updates, premier backports, comprehensive management, indemnification at significantly lower cost. Only Oracle delivers the industry's most complete integrated apps-to-disk Linux solutions.

3.2 Use-Case Summary

The NetApp FlexPod Select for Oracle solution, which can run high-performance Oracle databases in a highly resilient architecture, is competitive in terms of both price and performance. As part of this solution, the following use cases were validated.

- Deliver an architecture and a prescriptive reference deployment that provides a high level of resiliency against component failure.
- Deliver two million random read IOPS with microsecond-level latency by using an 8K request size.
- Demonstrate consistent performance and response time using a workload that consists of 90% random reads and 10% random writes by using an 8K block size.
- Demonstrate consistent performance and response time using a workload that consists of 80% random reads and 20% random writes by using an 8K block size.

4 Solution Design

4.1 Compute Design Elements

In the FlexPod Select for High-Performance Oracle RAC solution, Cisco UCS hosts are deployed to serve as the RAC nodes. Specific features of the Cisco UCS B-Series blades, chassis, and fabric interconnects are used to maximize reliability, performance, and efficiency. Those features are described in detail in this section. For additional information about Cisco UCS features and benefits, refer to the [Cisco website](#).

FCoE SAN Boot

Local storage devices on a server can constitute a single point of failure. For this reason, the environment uses enterprise storage to hold the operating system partitions for each server. Cisco UCS B-Series blades support boot from SAN by using Fibre Channel over Ethernet (FCoE) connectivity to a storage network. The boot volumes are located on the same EF560 arrays that hold the Oracle database data, which provides the boot volumes with the same degree of fault tolerance as the production data. This capability in turn generates both a high degree of protection against node downtime and extremely fast boot times.

Alternatively, the solution fully supports alternative boot methods including spreading boot disks across multiple arrays and local boot disks. However, these boot methods were not tested as part of the solution validation and may negate the benefits described.

Service Profile Templates

Service profiles are an integral component of the Cisco UCS product. The profiles contain a logical server identity, BIOS-level configuration parameters, and hardware compatibility requirements. A service profile can be applied to any physical servers that match the service profile's system requirements for a given Cisco UCS environment.

Service-profile templates help simplify the management of a proven service profile. A service-profile template allows an administrator to create multiple service profiles based on a known set of configuration parameters and administrator-defined pools of unique identifiers. For more information about service profiles and service-profile templates, refer to [Understanding Cisco Unified Computing System Service Profiles](#).

Although storage performance concerns are minimized by using EF560 flash arrays, storage network throughput could become a bottleneck. To maximize the throughput and minimize concerns about storage path contention during times when many servers might be booting simultaneously, two service-profile templates are created.

Both service profile templates contain two FCoE SAN interfaces: one for each of the two SAN fabrics of the solution. One service profile boots primarily from the SAN port on fabric A. The other service-profile template boots primarily from the SAN port on fabric B. Generated service profiles are evenly divided from the two service profile templates. This approach balances boot-time traffic across both network paths to prevent network bandwidth contention. Both service-profile templates use the remaining SAN fabric as a secondary boot path. In this way, boot does not fail if one of the SAN fabrics fails.

Virtual Interface Cards

Cisco virtual interface cards (VICs) allow a server to present multiple virtual interfaces to the host operating system. These virtual interfaces appear to the host as physical interfaces. The virtual interfaces presented can be both virtual network interface cards (vNICs) and virtual host bus adapters (vHBAs). Traffic from the presented virtual interfaces converges within the VIC and is passed through the FEXs to the fabric interconnects for delivery to the network layer of the infrastructure. Virtual network interfaces can be configured to fail between the two fabrics of the Cisco UCS environment, which maximizes fault tolerance within the network layer of the servers.

The FlexPod for High-Performance Oracle RAC solution demonstrates the use of Cisco UCS VIC 1240 cards with Cisco UCS FEX 2208 modules to connect the chassis to the fabric interconnects. This combination of hardware enables a maximum bandwidth of 40Gbps to each server (20Gbps to each fabric). Each FEX 2208 is fully cabled, which provides a total of 80Gbps of bandwidth from each fabric interconnect to one FEX of each chassis. Without overcommitting throughput, a maximum of four servers can be used with this configuration.

Optionally, additional Cisco UCS adapters and FEXs can be used to provide additional aggregate throughput to a Cisco UCS chassis and to each server within a chassis. A maximum of 80Gbps can be provided to each server with the VIC 1280. For more information about these technologies, refer to [Cisco Unified Computing System Adapters](#) and the [Cisco UCS 2200 Series Fabric Extenders Data Sheet](#).

4.2 Network Design Elements

In the FlexPod Select for High-Performance Oracle RAC solution, Cisco Nexus switches are leveraged for client access and to serve data. Data access is achieved through native Fibre Channel (FC) from the storage device and by FCoE from the Cisco UCS hosts. Client access and Oracle RAC cluster interconnect traffic occur over traditional Ethernet.

Specific features of the Cisco Nexus 5548 are leveraged to maximize reliability, performance, and efficiency. Those features are described in detail in this section. For additional information about Cisco Nexus features and benefits, refer to the [Cisco Nexus Family](#) website.

Fabric Convergence

As previously mentioned, this design leverages FC-based NetApp storage and FCoE-enabled Cisco UCS with the VICs. This topology requires a switching infrastructure that supports both traditional Ethernet traffic and all types of storage traffic, including the lossless requirements for block-level storage transport over FC or FCoE. The Cisco Unified Fabric, enabled by the Cisco Nexus 5548UP switch, creates high performance, low latency, and highly available networks to serve a diverse set of data center needs.

This FlexPod design supports multiple LAN and SAN protocols, most notably FCoE and FC. The solution provides a 10GbE-enabled fabric that is defined by dedicated FCoE uplinks and dedicated Ethernet uplinks between the Cisco UCS fabric interconnects and the Cisco Nexus switches. The solution also provides native 8Gbps FC between the NetApp storage devices and the same multipurpose Cisco Nexus platforms.

This design does not employ a dedicated SAN switching environment because the Cisco Nexus 5548UP switch is capable of satisfying FC SAN requirements. The Cisco Nexus 5500 series switches are configured in N_Port ID Virtualization (NPIV) mode to provide storage services for the FCoE-based traffic that traverses its fabric. Cisco Nexus 5600 series switches are also supported in this solution.

Link Aggregation Control Protocol and Cisco Virtual Port Channels

In terms of fabric resiliency, two technologies that complement each other are link aggregation and link aggregation control protocol (LACP). LACP is an enhanced feature that greatly simplifies link aggregation procedures. LACP works by enabling all ports on both ends of the aggregated link to send and receive LACP data unit (LACPDU) frames. All participating devices use LACPDU frames to communicate and negotiate which ports are allowed to perform as part of the aggregated link. This technology greatly simplifies the link aggregation procedure and minimizes the risk of misconfiguration.

The Cisco port channels and Cisco virtual port channels (vPCs) used in this solution all support and leverage LACP for Ethernet traffic. In addition, the Cisco Nexus 5500 series features vPC capabilities. A vPC allows links that are physically connected to two different Cisco Nexus 5500 series devices to appear as a single logical link aggregation group to a third device. This capability essentially creates device fault tolerance. vPCs address aggregate bandwidth, link, and device resiliency. Cisco UCS fabric interconnects benefit from the Cisco Nexus vPC abstraction, gaining link and device resiliency as well as full utilization of a nonblocking Ethernet fabric. Ethernet traffic for the Cisco UCS environment is satisfied with two-port Ethernet link aggregation groups from the fabric interconnect to the Cisco Nexus switches.

FC Fabric Considerations

Cisco Nexus 5500 series switches have robust FC SAN capabilities whether you require native FC or FCoE. FCoE is uplinked from the Cisco UCS environment, and the NetApp EF560 storage array provides native FC. A variety of FC SAN features are leveraged to enable communication between the hosts and the storage, including:

- **Virtual storage area network (VSAN).** VSANs are logical partitions of the SAN fabric. This type of environment is configured as fabric A and fabric B.
- **Worldwide node names (WWNNs) and worldwide port names (WWPNs).** A WWNN, representing a node in the SAN fabric, and a WWPN, representing a specific port of the node, are identified in the fabric with a unique 16-character identifier. Host initiators and storage targets are identified on the FC SAN fabric with WWNNs and WWPNs.
- **Device aliases.** Device aliases can be used to assign naming conventions or labels to WWNNs or WWPNs in the FC SAN fabric. These labels can be used for configuration instead of recording complex WWNNs or WWPNs.
- **Fabric zoning.** Fabric zoning can be used to segment the SAN fabric to restrict access, increase security, and simplify management. Zoning can be employed through physical port-based zoning or

zoning based on WWPNs. This topology employs zoning based on WWPNs with a single-initiator, multiple-target methodology.

4.3 NetApp EF560 Storage Array Elements

RAID Options

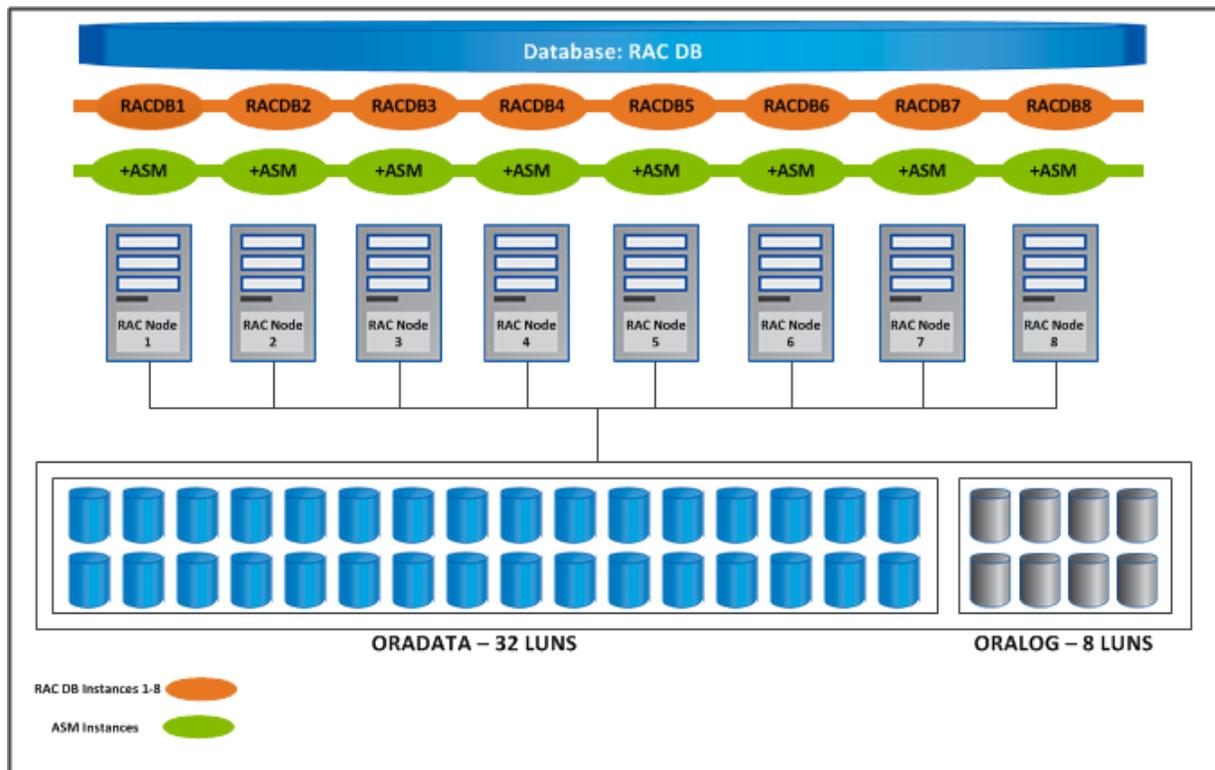
The EF560 flash array storage controllers support traditional RAID levels (0, 1, 5, and 6) as well as Dynamic Disk Pools to provide both flexibility and simplicity for storage administrators. For this solution, RAID 10 was chosen to get the maximum performance out of the storage arrays. Refer to Table 7 for a detailed layout of the storage setup.

Data Layout

Figure 2 illustrates how the Oracle RAC database was configured with the EF560 all-flash arrays for an eight-node RAC database. The `ORADATA` and `ORALOG` LUNs were provisioned on the EF560 all-flash arrays by using SANtricity and were presented to the Oracle RAC nodes. A single ASM disk group was created that spanned all of the LUNs to present a single file system from a database and host perspective. Refer to Table 7 for information about LUN sizes and names.

The result of this configuration was that the workload was evenly distributed across all the storage LUNs and volumes on all arrays. External redundancy was chosen for the ASM disk groups because the NetApp EF-Series already provided protection at the array level. Optionally, customers can choose ASM mirroring to increase data protection, provided the storage arrays are provisioned with enough capacity.

Figure 2) ASM layout for eight-node RAC.



Storage Network Connectivity

To prevent failures that would cause disks to become inaccessible to storage, EF-Series storage arrays contain multiple redundant hardware components. The redundant hardware is also redundantly connected to both network fabrics to make sure that hardware at either the network layer or the storage layer can fail without causing data-traffic disruptions.

Each EF560 storage array comes with two controllers, both of which can access the disks in the array. In this high-availability configuration, if a controller fails, the other controller in the array takes over communication between the disks and the servers. Each of these controllers contains an FC host interface card that has four ports (eight ports total per array).

To maximize throughput, all eight ports are cabled. In each controller, four ports are connected to SAN fabric A and four ports are connected to SAN fabric B. Each port can provide a maximum of 16Gbps of bandwidth, but they will all negotiate down to 8Gbps, the maximum throughput of an FC port on a Cisco Nexus 5548UP.

5 Solution Validation

To validate the defined use cases, NetApp configured an eight-node Oracle RAC database test environment accessing a total of four EF560 all-flash-array storage controllers. NetApp then ran performance tests by using a workload generator that generated a series of transactions that delivered different load levels to the Oracle database. The solution's objective was achieved through the validation of the use cases, and the results are available in the "Performance Test Results" section.

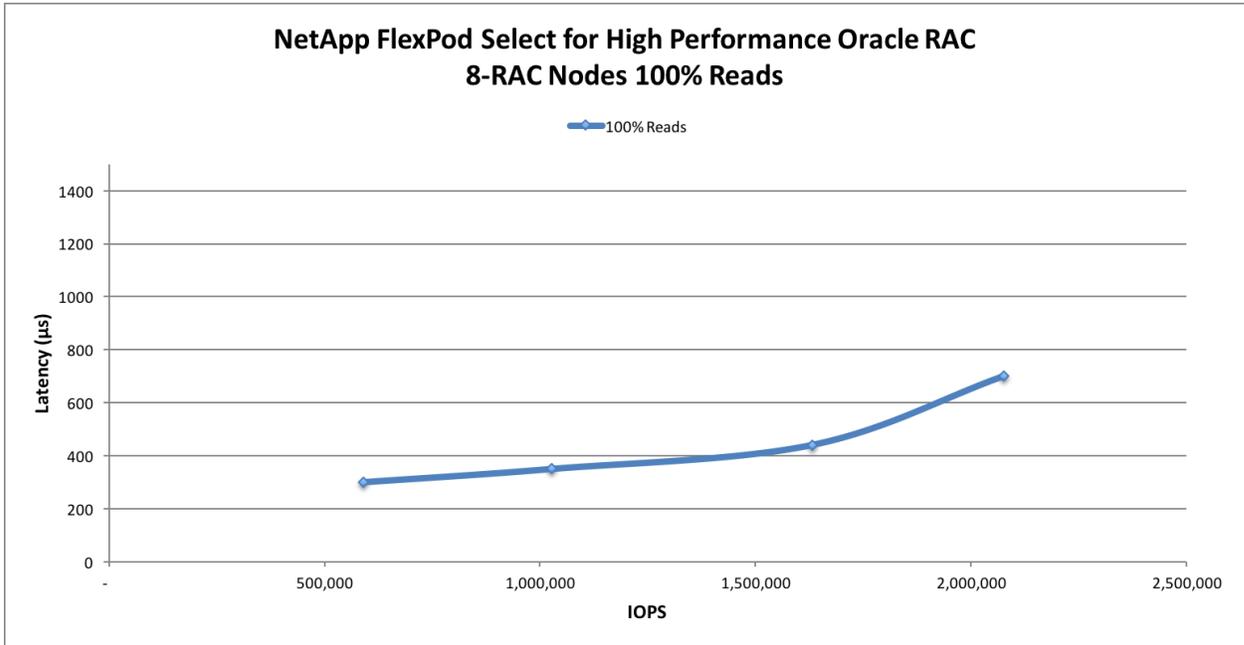
Previously, comparable tests were conducted using Oracle 11gR2 RAC and NetApp EF550 all-flash-array storage controllers using this same methodology. It was found that upgrading to the EF560 and Oracle 12cR1 doubles the number of random read IOPS to over 2 million random read IOPS and provides significant performance gains in mixed workloads compared to the EF550. See Appendix B of this document to review the EF550 test data.

5.1 Performance Test Results

For all tests, NetApp used the Silly Little Oracle Benchmark (SLOB2) workload generator to simulate the I/O patterns that are likely to be encountered in actual Oracle production environments. SLOB2 drives different levels of simulated users, each generating the specific I/O patterns described previously in the use-case section. After each test, NetApp recorded the physical database reads and average latency from the Oracle automatic workload repository reported by the Oracle database.

Figure 3 shows the IOPS and average latency observed by the database during testing with 100% random 8K reads. The load on the database was increased incrementally until the IOPS exceeded two million while simultaneously observing the application latency under 1 millisecond (ms). The storage arrays are capable of delivering higher levels of IOPS provided users have a tolerance for higher latency.

Figure 3) FlexPod Select for Oracle with eight RAC nodes showing 100% 8K random reads.



The results of the mixed-workload use cases shown in Table 1 indicate how the solution performed when there were writes on the Oracle database. The workload was also run in incremental loads to generate maximum IOPS by enabling latency to be as close to 1 ms as possible.

Table 1) Solution validation results for all use cases.

Use Case	Maximum IOPS	Average Latency – (µs)
100% reads	2,076,637	700
90% reads and 10% updates	1,400,395	600
80% reads and 20% updates	1,300,909	1040

6 Technology Requirements

This section describes the technology requirements for the FlexPod Select for High-Performance Oracle RAC solution.

6.1 Hardware Requirements

Figure 4 shows the hardware components associated with this solution.

Figure 4) FlexPod Select for High-Performance Oracle RAC solution architecture.

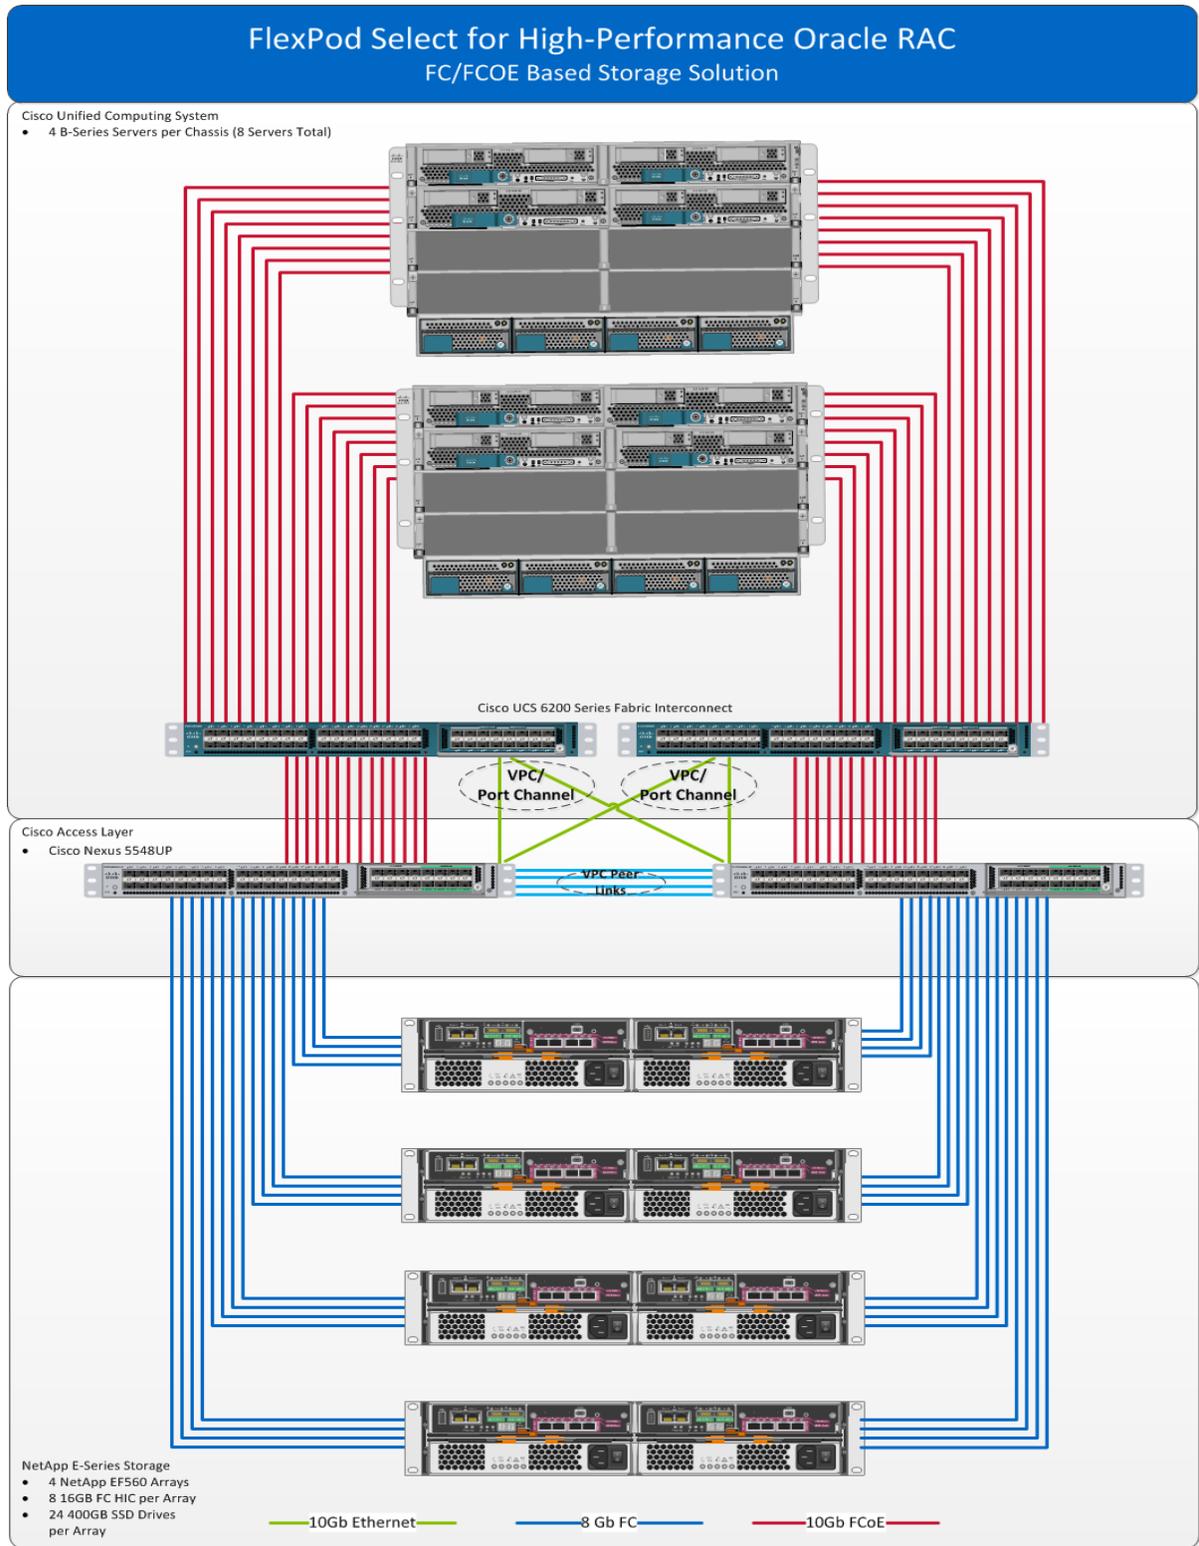


Table 2 lists the hardware components required to implement the solution and to achieve the previously defined performance objectives. The hardware components used in any particular implementation of this solution might vary based on customer requirements.

Table 2) Hardware requirements.

Hardware	Quantity
Storage	
NetApp EF560 dual-controller array	4
NetApp 16Gb Fibre Channel host interface card	8
NetApp 400GB SSD	96
Compute	
Cisco UCS fabric interconnect 6248	2
Cisco UCS 5108 chassis	2
Cisco B200M3 compute blade with 2 Intel® Xeon® E5-2620 v2 processors and 64GB RAM	4
Cisco B200M3 compute blade with 2 Intel® Xeon® E5-2650 v2 processors and 128GB RAM	4
Cisco UCS 1240 virtual interface card (VIC)	8
Cisco UCS 2208 FEX	4
Network	
Cisco Nexus 5548UP switch	2
Cisco 8Gb Fibre Channel SFP (DS-SFP-FC8G-SW)	32

6.2 Software Requirements

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

Table 3) Software requirements.

Software	Version
Storage	
NetApp EF560 firmware version	8.20.08.00
NetApp SANtricity Storage Manager	11.20.0G00.0006
Compute	
Cisco UCS Manager	2.2(5b)
Cisco UCS firmware bundle	2.2(5b)
Cisco UCS 1240 VIC fnic driver	1.6.0.18

Software	Version
Cisco UCS 1240 VIC enic driver	2.1.1.67
Network	
Cisco Nexus NX-OS	6.0(2)N1(2a)
Operating System	
Oracle Enterprise Linux® (OEL)	6.6
Application Software	
Oracle Real Application Clusters (RAC)	12.1.0.2.0
SLOB tool	2.3

7 Best Practices

As part of the solution validation, some recommendations are used to reach optimal performance. The configuration settings used in this document will not always be appropriate for all customer requirements. Contact NetApp for comprehensive storage system design assistance.

7.1 Storage Recommendations

Table 4 shows the basic tuning parameters and the default settings as well as some of the settings that were used for this solution from a storage array perspective.

Table 4) Storage tuning parameters.

Name	Context	Default	Available Options	Solution Recommendations
Cache block size	Array	32	4, 8, 16, 32	32K
Automatic cache flushing	Array	80%	0–100	80
Segment size	Volume	128K	32K, 64K, 128K, 256K, 512K	64K or 128K
Read caching	Volume	Enable	Enable/disable	Enable
Dynamic cache read prefetch	Volume	Enable	Enable/disable	Disable
Write cache	Volume	Enable	Enable/disable	Enable
Write cache mirroring	Volume	Enable	Enable/disable	Enable

7.2 Oracle ASM Recommendations

Table 5 shows the basic tuning parameters for the Oracle ASM and database instance.

Table 5) Oracle ASM and instance settings.

Name	Default	Available Options	Solution Recommendations
ASMU AU	1MB	1MB–64MB	64MB

Name	Default	Available Options	Solution Recommendations
File system I/O options	None	None	Not used
db_file_multiblock_read_count	128	8–128	Default

7.3 Linux Recommendations

Table 6 shows the basic tuning parameters used on the RAC nodes running Oracle Linux 6.6.

Table 6) Linux tuning parameters.

Name	Default	Available Options	Solution Recommendations
I/O scheduler	cfq	cfq, anticipatory, deadline, noop	noop

8 Conclusion

Latency is a critical factor in any IT environment. The more latency that can be reduced, the more efficient business operations can be. The FlexPod Select for High-Performance Oracle RAC solution is designed for applications that need consistent low latency, performance, and reliability. The architecture is also highly scalable, which allows the customer to select the number of Cisco UCS server blades and EF-Series storage units required for the workload. The NetApp EF560 all-flash storage array is a true enterprise-class storage array designed for applications that expect high IOPS with microsecond-level latencies. The fully redundant enterprise-class high-availability feature on the storage controllers provides maximum availability and delivers the required performance for mission-critical applications for which superior performance and low latency are imperative.

Appendix A: Consolidated Storage Layout for an Eight-Node Oracle RAC Environment

Table 7 describes the consolidated storage layout for an eight-node Oracle RAC environment.

Table 7) Storage layout for eight-node RAC.

Storage Array	Type	Volume Group RAID 10	Volume/LUN Name	Number of Physical Disks	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
EF560-1	Redo logs boot LUNs	A1LOGVG	RAC1BOOT	2	78	256GB	2	RAC-A01
			RAC2BOOT		78			RAC-A02
			LOG1		50			RAC
			LOG2		50			
	Data	A1VG1	DATALUN1	10	450	3.6TB		
			DATALUN2		450			
			DATALUN3		450			
			DATALUN4		450			
			A1VG2		DATALUN5			10

Storage Array	Type	Volume Group RAID 10	Volume/LUN Name	Number of Physical Disks	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
			DATALUN6		450			
			DATALUN7		450			
			DATALUN8		450			
EF560-2	Redo logs	A2LOGVG	LOG3	2	50	271GB	2	RAC
			LOG4		50			
			RAC3BOOT		78			RAC-A03
			RAC4BOOT		78			RAC-A04
			GRID1		5			RAC
			GRID2		5			
			GRID3		5			
	Data	A2VG1	DATALUN9	10	450	3.6TB		RAC
			DATALUN10		450			
			DATALUN11		450			
			DATALUN12		450			
		A2VG2	DATALUN13	10	450			
			DATALUN14		450			
			DATALUN15		450			
DATALUN16	450							
EF560-3	Redo logs	A3LOGVG	LOG5	2	50	256GB	2	RAC
			LOG6		50			
			RAC5BOOT		78			RAC-B01
			RAC6BOOT		78			RAC-B02
	Data	A3VG1	DATALUN17	10	450	3.6TB		RAC
			DATALUN18		450			
			DATALUN19		450			
			DATALUN20		450			
A3VG2	DATALUN21	10	450					
	DATALUN22		450					
	DATALUN23		450					
	DATALUN24		450					
EF560-4	Redo logs	A4LOGVG	LOG7	2	50	256G	2	RAC
			LOG8		50			

Storage Array	Type	Volume Group RAID 10	Volume/LUN Name	Number of Physical Disks	Allocated Capacity in GB	Total Capacity	Spare Disks	Mapped Host Group Name/Host Name
			RAC7BOOT		78			RAC-B03
			RAC8BOOT		78			RAC-B04
	Data	A4VG1	DATALUN25	10	450	3.6TB		RAC
			DATALUN26		450			
			DATALUN27		450			
			DATALUN28		450			
		A4VG2	DATALUN29	10	450			
			DATALUN30		450			
			DATALUN31		450			
			DATALUN32		450			

Appendix B: Results of Mixed-Workload Use Cases Using EF550 All-Flash-Array Storage Controllers

Previously, we conducted a set of identical tests with the same configuration using Oracle 11gR2, RHEL 6.5, and NetApp EF550 storage controllers. The results of the mixed-workload use cases generated in previous testing using the EF550 all-flash-array storage controllers are shown in Table 8 for reference.

Table 8) Previous solution validation results for all use cases.

Use Case	Maximum IOPS	Average Latency(µs)
100% reads	1,061,092	900
90% reads and 10% updates	797,147	1100
80% reads and 20% updates	619,680	900

References

This report references the following documents and resources:

- [Cisco UCS Virtual Interface Card Drivers for Linux Installation Guide](#)
- [Cisco UCS 2.2\(5\) Hardware and Software Interoperability Matrix](#)
- [Cisco UCS fnic Tunables](#)
- [Oracle Database Quick Installation Guide 12c Release 1 \(12.1\) for Linux x86](#)
- [NetApp EF560 Flash Array Installation Guide](#)
- [NetApp Extreme Performance Solution for Oracle Database | TR-4305](#)

Version History

Version	Date	Document Version History
Version 1.0	November 2014	Engineering content creation
Version 2.0	November 2015	Updated with EF560 performance data

NetApp Contact Information

NetApp
495 East Java Drive
Sunnyvale, CA 94089
United States

Telephone: +1 408 822 6000
Fax: +1 408 822 4501
Support telephone: +1 888 463 8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

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

Copyright Information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC. SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL, and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>. NVA-0012-FP-DESIGN-0815