



Technical Report

FlexPod with NetApp AFF A-Series

Bobby Oommen, Abhinav Singh, Kavyashree Mahadevaiah, NetApp
September 2024 | TR-4999

In partnership with



Abstract

This report presents the performance capabilities of FlexPod® with the latest NetApp® AFF A1K. It highlights the importance of using FlexPod for Oracle database. The document provides storage layout recommendations and best practices specifically tailored for large enterprise deploying the critical applications on a FlexPod infrastructure.

TABLE OF CONTENTS

FlexPod Introduction	3
NetApp ONTAP	4
NetApp AFF A-Series Storage	4
Solution Overview	5
Solution Components	6
Solution Topology	12
Installation and Configuration	21
Solution Validation	27
Conclusion	29
Where to find additional information	29
Version history	29

LIST OF TABLES

Table 1 Cabling information for Cisco Nexus 9336C-FX2.....	13
Table 2 Cabling information for NetApp AFF A1K.....	14
Table 3 Cabling information for Cisco UCS FI 6536	15
Table 4 Cabling information for Cisco MDS 9124V	16
Table 5 vNIC and vHBA configured on each Linux host	18
Table 6 VLAN and VSAN Configuration.....	18
Table 7 Zoning for FCP and NVMe setup	19
Table 8 Hardware and Software versions	20
Table 9 Server VIC Mapping.....	21

LIST OF FIGURES

Figure 1 FlexPod Datacenter Solution	3
Figure 2 NetApp AFF A1K Front and Rear View	7
Figure 3 Cisco UCS X210C M7 Blade Server	7
Figure 4 Cisco VIC 15231 card and Infrastructure.....	9
Figure 5 Cisco UCS X9508 Chassis Front and Rear view	9
Figure 6 Cisco Nexus 9336C-FX2 Switch.....	10
Figure 7 Cisco MDS 9124V Switch.....	11
Figure 8 Cisco UCS 6536 Fabric Interconnect	11
Figure 9 Physical Topology.....	12

FlexPod Introduction

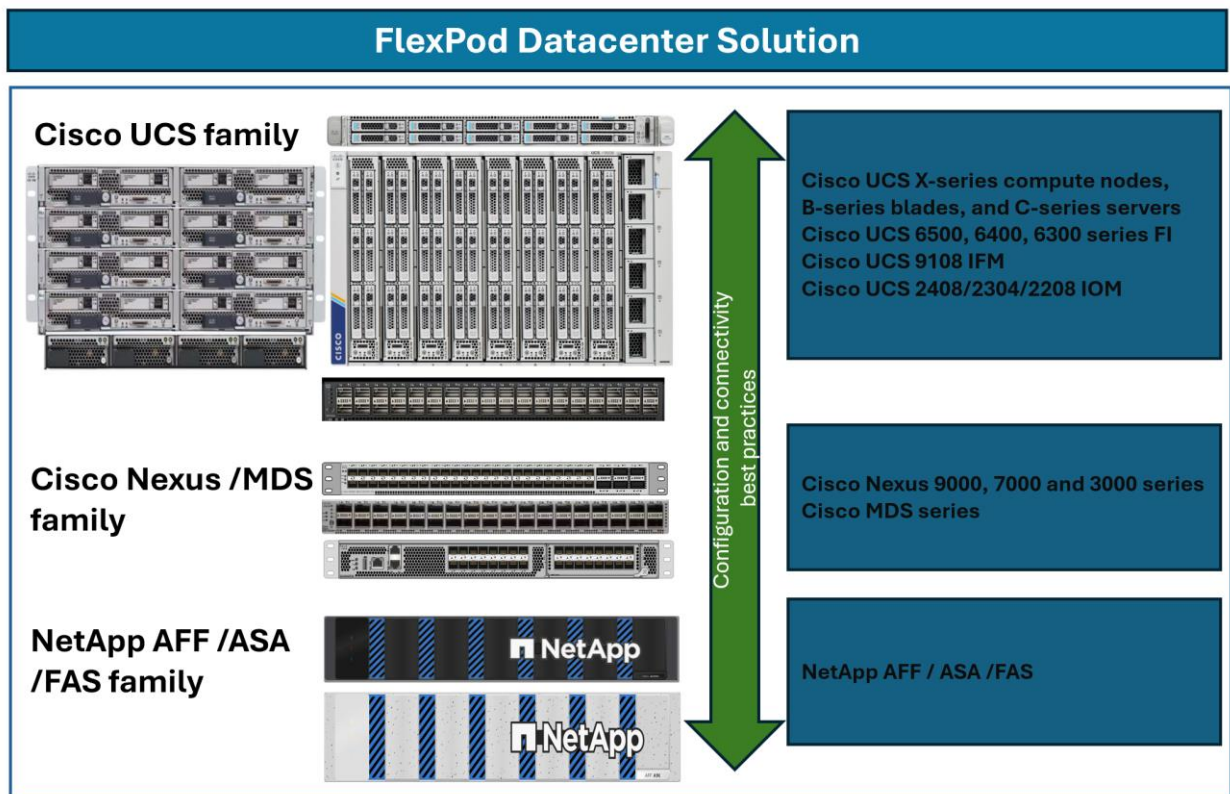
FlexPod overview

FlexPod is a best practice converged infrastructure data center architecture that includes the following components from Cisco® and NetApp:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus and MDS families of switches
- NetApp Fabric Attached Storage (FAS), All Flash FAS (AFF), and All-Flash SAN Array (ASA) systems

Shown in Figure 1 are some of the components utilized for creating FlexPod solutions. These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence.

Figure 1 FlexPod Datacenter Solution



Each of the FlexPod component families shown (Cisco UCS, Cisco Nexus/MDS switches, and NetApp storage) provides platform and resource options to scale the infrastructure up or down as per application requirement, while supporting the features and functionalities that are required under the configuration and connectivity best practices of FlexPod. FlexPod can also scale out for environments that require multiple consistent deployments by rolling out additional FlexPod stacks.

All FlexPod components have been integrated so you can deploy the solution quickly and economically while eliminating many of the risks associated with researching, designing, building, and deploying similar solutions from the foundation. One of the main benefits of FlexPod is its ability to maintain consistency at scale. The design is flexible enough that the networking, computing, and storage can fit in one data center rack or be deployed according to a customer's data center design. The reference architecture reinforces

the "wire-once" strategy, because additional storage is added to the architecture and no re-cabling is required from the hosts to the Cisco UCS fabric interconnect.

Audience

This document is intended for NetApp field and partner solutions engineers (SEs), professional services personnel and field consultants. NetApp assumes that the reader has the following background knowledge:

- Understanding of SAN concepts
- Technical familiarity with the configuration and administration of FlexPod systems

A working knowledge of Oracle Database, Linux, storage technology, and network is assumed but is not a prerequisite to read this document.

NetApp ONTAP

NetApp ONTAP® is an enterprise data-management platform with native capabilities that includes industry leading storage efficiency capabilities and could be clustered up to 12 nodes for SAN and up to 24 nodes for NAS. NetApp Snapshot™ technology which is an integral part of ONTAP enables instantaneous backups of critical datasets and cloning of datasets and offers comprehensive disaster recovery capabilities.

NetApp storage systems also offer a hybrid cloud foundation that helps customers to seamlessly move data across edge, core and cloud. Customers can benefit from this robust data mobility and integration capabilities enabled by NetApp intelligent data infrastructure which can easily get data from the edge (where it is generated) to the core (where it is utilized) and to the cloud (to take advantage of the on-demand elastic compute and AI / ML capabilities) to gain actionable business insights from your valuable data.

NetApp AFF A-Series Storage

The NetApp AFF A-Series storage family, powered by NetApp ONTAP data management software, delivers the same NetApp simplicity and reliability that tens of thousands of organizations of every size, in every industry, around the globe, have trusted for years. It's the same technology that the top three public cloud providers rely on to drive all your apps and data across hybrid cloud without the need of silos or storage complexity. It is a powerful, intelligent, secure storage to seamlessly accelerate your business.

With AFF systems, you can maintain constant data availability with zero data loss and zero downtime in the event of disruption or disaster. NetApp MetroCluster™ software replicates your data synchronously to a separate location to protect your entire system. If something goes wrong at one site, your applications automatically and instantaneously switch to the other site. For a more tailored approach, choose NetApp SnapMirror® active sync to cost-efficiently replicate the most critical data while taking advantage of the increased performance, greater flexibility, and enhanced load-balancing capability that comes with the symmetric active-active architecture.

With NetApp ONTAP advanced data management software, you can lower IT costs by simplifying operations, consolidating workloads, and lowering overheads. NetApp AFF A-Series offer broad support of application ecosystems and deep integration for enterprise applications, virtual desktop infrastructure, databases, server virtualization, and the MLOps ecosystem. Infrastructure management tools simplify and automate common storage tasks:

- Easily provision and rebalance workloads with one-click automation and self-service.
- Upgrade your OS and firmware with a single click.
- Import LUNs from third-party storage arrays directly into an AFF system to seamlessly migrate data.

The NetApp AFF A-Series systems help in consolidating and scaling all your workloads:

- Deliver up to 2x performance compared to previous generation systems, with latency as low as 100µs.
- Support any data type, any app workload, across hybrid cloud.
- Provide consistent performance, adaptive quality of service, and proven 99.9999% data availability to safeguard SLAs even in multi-workload and multi-tenant environments.
- Scale non-disruptively to 702PB effective capacity in a cluster.

Note: Not all A-Series systems can scale up to 702PB. Refer to [datasheet](#).

- Improve the speed and productivity of collaborative teams across multiple locations and increase data throughput for read-intensive applications with NetApp FlexCache® software.

With the new AFF A70, AFF A90, and AFF A1K systems, you will never need to choose between performance and efficiency. You'll always have improved data compression and no performance impact, thanks to Quick Assist Technology (Intel® QAT®). The systems allow you to achieve exceptional storage efficiency while delivering the consistent high performance needed for mission-critical workloads. In addition, the new systems come with faster front-end 200Gb Ethernet and 64Gb FC networking connectivity.

In this solution, the high end AFF A1K was used for validation. Refer the [NetApp AFF A-Series Datasheet](#) for more technical details and specifications.

Solution Overview

NetApp offers a simplified and comprehensive data management solution for all enterprise applications toward their digital transformation. This solution document primarily discusses the design and implementation, of a SAN network to achieve maximum reliability for Oracle database.

Solution Benefits

FlexPod systems with NetApp AFF A-Series offers the following benefits:

- **Data Access protocol & networking support:**

	AFF A1K	AFF A90	AFF A70	AFF A900	AFF A800	AFF A400	AFF A250	AFF A150
NFSv4/RDMA	•	•	•		•	•		
NVMe/TCP	•	•	•	•	•	•	•	•
NVMe/FC	•	•	•	•	•	•	•	
FC	•	•	•	•	•	•	•	•
iSCSI	•	•	•	•	•	•	•	•
NFS	•	•	•	•	•	•	•	•
pNFS	•	•	•	•	•	•	•	•
CIFS/SMB	•	•	•	•	•	•	•	•
Amazon S3	•	•	•	•	•	•	•	•

- **High Availability:**
 - Modular controller architecture
 - Nondisruptive maintenance, upgrade, and scale-out clustering
 - Multi-site resiliency for continuous data access
- **Security and Compliance:**
 - Real-time autonomous ransomware detection, enhanced by embedded machine learning models, designed for industry-first 99%+ accuracy
 - Security information and event management (SIEM)/ Extended Detection and Response (XDR) integrations
 - Multi-factor authentication, multi-admin verification, dynamic authorization framework
 - Secure multi-tenant shared storage
 - Tamper-proof snapshots with SnapLock®
 - In-flight and data-at-rest encryption
 - Regulatory-compliant data retention
- **Storage Efficiency:**
 - Inline data compression, deduplication, and compaction
 - Space-efficient LUN, file, and volume cloning
 - Automatic data tiering
- **Data Protection:**
 - Application-consistent data protection and clone management with NetApp SnapCenter®
 - Integrated remote backup/disaster recovery
 - Replication to any NetApp AFF or FAS or ASA system on the premises or in the cloud with NetApp SnapMirror® technology
 - Business continuity with NetApp Metrocluster and NetApp SnapMirror active sync

Solution Components


NetApp AFF A1K

The NetApp AFF A1K model comes with 4U form factor, 208 CPU cores, 2048 GB of physical memory, 128 GB NVRAM, 14 PCIe expansion slots. The platform supports FC, iSCSI, NVMe/FC, NVMe/TCP, NFS, NFSv4 RDMA, pNFS, CIFS/SMB, Amazon S3 storage networking protocols that support a maximum effective capacity of 185PB. Refer to the following link for [AFF A-Series technical specifications](#).

NetApp AFF A-Series systems deliver industry-leading performance, verified by SPC-1 and SPEC SFS industry benchmarks. These systems are ideal for everything from VMware environments to highly transactional applications (such as Oracle, Microsoft SQL Server, and MongoDB databases), to the most data-intensive AI training, tuning, analytics, and data warehouse workloads.

With the power of front-end NVMe/FC and NVMe/TCP host connectivity combined with back-end NVMe-attached SSDs, the AFF A1K modular system delivers up to 40M IOPS and 1TB/s throughput in a single cluster through a unified, scale-out architecture.

Figure 2 NetApp AFF A1K Front and Rear View



Cisco UCS X-Series X210C M7 Blade Server

Figure 3 Cisco UCS X210C M7 Blade Server



The Cisco UCS X210c M7 Compute Node is the second generation of compute node to integrate into the Cisco UCS X-Series Modular System. It delivers performance, flexibility, and optimization for deployments in data centers, in the cloud, and at remote sites. This enterprise-class server offers market-leading performance, versatility, and density without compromise for workloads. Up to eight compute nodes can reside in the 7-Rack-Unit (7RU) Cisco UCS X9508 Server Chassis, offering one of the highest densities of compute, I/O, and storage per rack unit in the industry.

The Cisco UCS X210c M7 provides these main features:


- CPU: Up to 2x 5th Gen or 4th Gen Intel® Xeon® Scalable Processors with up to 64 cores per processor and up to 320 MB of Level 3 cache per CPU.
- Memory: Up to 8TB of main memory with 32x 256 GB DDR5 5600 MT/s or DDR5 4800 MT/s DIMMs depending on the CPU installed.

- Storage: Up to two M.2 SATA drives with optional hardware RAID or up to two M.2 NVMe drives in pass-through mode.
- Optional front mezzanine GPU module: The Cisco UCS front mezzanine GPU module is a passive PCIe Gen 4.0 front mezzanine option with support for up to two NVMe drives and two HHL GPUs.
- The blade server can utilize below modular LAN on motherboard (mLOM) virtual interface card(s) (VIC):
 - Cisco UCS VIC 15420 occupies the server's mLOM slot, enabling up to 50 Gbps of unified fabric connectivity to each of the chassis Intelligent Fabric Modules (IFMs) for 100 Gbps connectivity per server.
 - Cisco UCS VIC 15231 occupies the server's mLOM slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis IFMs for 100 Gbps connectivity per server.
 - Cisco UCS VIC 15230 occupies the server's mLOM slot, enabling up to 100 Gbps of unified fabric connectivity to each of the chassis IFMs for 100 Gbps connectivity per server with secure boot technology.
- Optional mezzanine card:
 - Cisco UCS 5th Gen Virtual Interface Card (VIC) 15422 can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric technology. An included bridge card extends this VIC's 2x 50 Gbps of network connections through IFM connectors, bringing the total bandwidth to 100 Gbps per fabric (for a total of 200 Gbps per server) with secure boot technology.
 - Cisco UCS PCI Mezz card for X-Fabric can occupy the server's mezzanine slot at the bottom rear of the chassis. This card's I/O connectors link to Cisco UCS X-Fabric modules and enable connectivity to the Cisco UCS X440p PCIe Node.
 - All VIC mezzanine cards also provide I/O connections from the X210c M7 compute node to the X440p PCIe Node.
- Security: The server supports an optional Trusted Platform Module (TPM). Additional features include a secure boot field-programmable gateway (FPGA) and Anti-Counterfeit Technology 2 (ACT2) provisions.

Cisco VIC 15231


The Cisco UCS VIC 15231 (Figure 4) are 2x100-Gbps Ethernet/FCoE-capable mLOM adapters designed exclusively for the Cisco UCS X210c Compute Node. The Cisco UCS VIC 15231 adapters enable a policy-based, stateless, agile server infrastructure that can present to the host PCIe standards-compliant interfaces that can be dynamically configured as either NICs or HBAs. In this solution, we have used 2 VIC cards per UCS Server.

Figure 4 Cisco VIC 15231 card and Infrastructure



Cisco UCS X9508 Chassis

Figure 5 Cisco UCS X9508 Chassis Front and Rear view





The Cisco UCS X-Series Modular System begins with the Cisco UCS X9508 Chassis, engineered to be adaptable and future-ready. The X-Series is a standards-based open system designed to be deployed and automated quickly in a hybrid cloud environment.

With a midplane-free design, I/O connectivity for the X9508 Chassis is accomplished with front-loading vertically oriented computing nodes that intersect with horizontally oriented I/O connectivity modules in the rear of the chassis. Cisco UCS X-Series is powered by Cisco Intersight, making it simple to deploy and manage at scale.

The Cisco UCS X9508 Chassis provides these features and benefits:

- The Seven-Rack-Unit (7RU) chassis has eight front-facing flexible slots. These can house a combination of computing nodes and a pool of future I/O resources, which may include Graphics Processing Unit (GPU) accelerators, disk storage, and nonvolatile memory.
- Two Cisco UCS 9108 IFMs at the top of the chassis connect the chassis to upstream Cisco UCS 6500 Series Fabric Interconnects. The IFMs offer these features:
 - Up to 200 Gbps of unified fabric connectivity per compute node with two Fabric Interconnects.
 - There are 8x 100G QSFP28 ports per IFM offering up to 1.6Tbps throughput per chassis.
- At the bottom of the chassis are slots ready to house future I/O modules that can flexibly connect the computing modules with I/O devices.

Cisco Nexus 9336C-FX2 Switches

Figure 6 Cisco Nexus 9336C-FX2 Switch



The Cisco Nexus 9336C-FX2 offers flexible port speeds supporting 10/25/40/100 Gbps in a compact 1 RU form factor with cloud-scale technology. It is designed to meet the changing needs of data centers, big data applications, and automated cloud environments.

- All 36 ports support 10/25/40/100 Gbps QSFP28 and wire-rate MACsec encryption.
- Supports 7.2 Tbps of bandwidth and over 2.8 bpps.
- Enhanced Cisco NX-OS Software designed for performance, resiliency, scalability, manageability, and programmability.
- Real-time buffer utilization per port and per queue, for monitoring traffic micro-bursts and application traffic patterns.

Cisco MDS 9124V Switches

Figure 7 Cisco MDS 9124V Switch



The next-generation Cisco MDS 9124V 64-Gbps 24-Port Fibre Channel Switch (Figure 7) provides high-speed Fibre Channel connectivity for all-flash arrays and high-performance hosts. This switch offers state-of-the-art analytics and telemetry capabilities built into its next-generation Application-Specific Integrated Circuit (ASIC) chipset. This switch allows seamless transition to Fibre Channel Non-Volatile Memory Express (NVMe/FC) workloads whenever available without any hardware upgrade. It empowers small, midsize, and large enterprises that are rapidly deploying cloud-scale applications using extremely dense virtualized servers, providing the benefits of greater bandwidth, scale, and consolidation.

Some of the main benefits for a small-scale Storage Area Network (SAN) are automatic zoning, nonblocking forwarding and a single port group of 24 ports. Benefits for a mid- to large-size SAN include higher scale for Fibre Channel control-plane functions, Virtual SANs, fabric login (FLOGI), device alias and name server scale, 24 ports of 64-Gbps non-oversubscribed line-rate ports, bidirectional airflow, and a fixed-form NVMe/FC-ready SAN switch with enhanced Buffer-to-Buffer (B2B) credits and capable of hardware-assisted Fibre Channel link encryption.

Large-scale SAN architectures built with SAN core directors can expand 64-Gbps connectivity to the server rack using these switches configured in either switch mode or Network Port Virtualization (NPV) mode. Additionally, the switch supports enhanced diagnostic features such as Inter-Switch Link (ISL) and Host-Bus-Adapter (HBA) diagnostics, remote SFP (Read Diagnostic Parameter) diagnostics, remote port beaconing (Link Cable Beaconing) and advanced reliability features such as link level Forward Error Correction (FEC) with HBA ports.

Cisco UCS 6536 Fabric Interconnect

Figure 8 Cisco UCS 6536 Fabric Interconnect



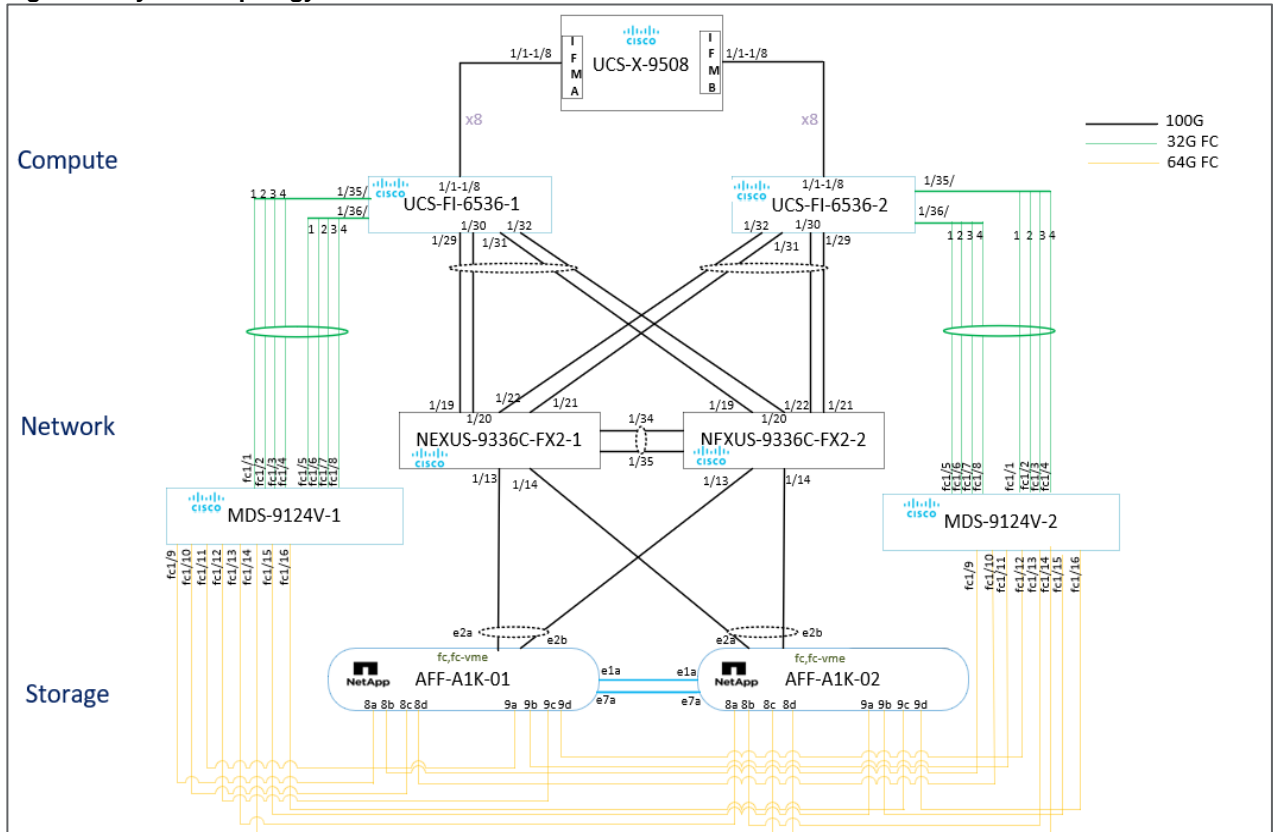
The Cisco UCS 6536 Fabric Interconnect (FI) is a core part of the Cisco Unified Computing System, providing both network connectivity and management capabilities for the system. The Cisco UCS 6536 Fabric Interconnect offers line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel, NVMe over Fabric, and Fibre Channel over Ethernet (FCoE) functions. In addition, by supporting a unified fabric, Cisco UCS 6536 Fabric Interconnect provides both LAN and SAN connectivity for all servers within its domain.

From a networking perspective, the Cisco UCS 6536 uses a cut-through architecture, supporting deterministic, low-latency, line-rate 10/25/40/100 Gigabit Ethernet ports, a switching capacity of 7.42 Tbps per FI and 14.84 Tbps per unified fabric domain, independent of packet size and enabled services.

Solution Topology

This section describes the physical topology of the FlexPod Data Center architecture, which was used for solution validation.

Figure 9 Physical Topology



Topology Details

This FlexPod Datacenter solution is built using the following hardware components:

- 4 * 4-port 64Gb FC SFP+ (X50140A) FC/NVMe for data, 2 * 2-port 40/100G Ethernet QSFP28 (X50130A) for HA, cluster interconnect and disk shelf access, and 2 * 2-port 40/100/200 GbE (RoCEv2) QSFP112 (X50131A) for Ethernet data on NetApp AFF A1K
- Cisco VIC 15231 adapters with two 100G ports connect to the IFM on each X210c M7 server. The IFM-100G enables 100GBASE-KR4 across the four 25G-KR lanes per IFM

- 8 * 100G links connected between each Cisco UCS Fabric Interconnects (UCS-FI-6536) and Intelligent Fabric Module (IFM)
- 8 * 100G Ethernet links connected between Cisco UCS Fabric Interconnects (UCS-FI-6536) and Cisco Nexus 9336C-FX2 switches
- 16 * 32G FC links between Fabric Interconnects and Cisco 9124V MDS switches using breakout cables
- 16 * 64G FC links between Cisco 9124V MDS switches and NetApp AFF A1K controllers
- Port Channel configuration is set up between Storage nodes and Nexus switches, one pair between Nexus and Fabric Interconnects and one port-channel between the Nexus peer switches with channel-mode active are configured.
- Port Channel is also configured between Fabric Interconnect and MDS for the FC traffic.

Cabling Information

Refer the following document for complete details on [installation and setup of AFF A1K system](#).

The reference validation cabling details are shown in Figure 9 and Table 1 through Table 4.

Table 1 Cabling information for Cisco Nexus 9336C-FX2

Local Device	Local Port	Remote Device	Remote Port
Cisco Nexus 9336C-FX2 A	1/13	NetApp AFF A1K 01	e2a
	1/14	NetApp AFF A1K 02	e2a
	1/19	Cisco UCS FI 6536 A	1/29
	1/20	Cisco UCS FI 6536 A	1/30
	1/21	Cisco UCS FI 6536 B	1/31
	1/22	Cisco UCS FI 6536 B	1/32
	1/34	Cisco Nexus 9336C-FX2 B	1/34
	1/35	Cisco Nexus 9336C-FX2 B	1/35
Cisco Nexus 9336C-FX2 B	1/13	NetApp AFF A1K 01	e2b
	1/14	NetApp AFF A1K 02	e2b
	1/19	Cisco FI 6536 A	1/31
	1/20	Cisco FI 6536 A	1/32
	1/21	Cisco FI 6536 B	1/29
	1/22	Cisco FI 6536 B	1/30
	1/34	Cisco Nexus 9336C-FX2 A	1/34

Local Device	Local Port	Remote Device	Remote Port
	1/35	Cisco Nexus 9336C-FX2 A	1/35

Table 2 Cabling information for NetApp AFF A1K

Local Device	Local Port	Remote Device	Remote Port
NetApp AFF A1K 01	e1a	NetApp AFF A1K 02	e1a
	e7a	NetApp AFF A1K 02	e7a
	e2a	Cisco Nexus 9336C-FX2 A	Eth 1/13
	e2b	Cisco Nexus 9336C-FX2 B	Eth 1/13
	8a	Cisco MDS 9124V A	fc1/9
	8b	Cisco MDS 9124V B	fc1/9
	8c	Cisco MDS 9124V A	fc1/10
	8d	Cisco MDS 9124V B	fc1/10
	9a	Cisco MDS 9124V A	fc1/11
	9b	Cisco MDS 9124V B	fc1/11
	9c	Cisco MDS 9124V A	fc1/12
	9d	Cisco MDS 9124V B	fc1/12
NetApp AFF A1K 02	e1a	NetApp AFF A1K 01	e1a
	e7a	NetApp AFF A1K 01	e7a
	e2a	Cisco Nexus 9336C-FX2 A	Eth 1/14
	e2b	Cisco Nexus 9336C-FX2 B	Eth 1/14
	8a	Cisco MDS 9124V A	fc1/13
	8b	Cisco MDS 9124V B	fc1/13
	8c	Cisco MDS 9124V A	fc1/14
	8d	Cisco MDS 9124V B	fc1/14
	9a	Cisco MDS 9124V A	fc1/15

Local Device	Local Port	Remote Device	Remote Port
	9b	Cisco MDS 9124V B	fc1/15
	9c	Cisco MDS 9124V A	fc1/16
	9d	Cisco MDS 9124V B	fc1/16

Table 3 Cabling information for Cisco UCS FI 6536

Local Device	Local Port	Remote Device	Remote Port
UCS FI 6536 A	1/29	Cisco Nexus 9336C-FX2 A	1/19
	1/30	Cisco Nexus 9336C-FX2 A	1/20
	1/31	Cisco Nexus 9336C-FX2 B	1/19
	1/32	Cisco Nexus 9336C-FX2 B	1/20
	1/35/1 – 1/35/4	Cisco MDS 9124V A	fc1/1-fc1/4
	1/36/1 – 1/36/4	Cisco MDS 9124V A	fc1/5-fc1/8
	1/1 – 1/8	IFM-A	1/1-1/8
UCS FI 6536 B	1/29	Cisco Nexus 9336C-FX2 B	1/21
	1/30	Cisco Nexus 9336C-FX2 B	1/22
	1/31	Cisco Nexus 9336C-FX2 A	1/21
	1/32	Cisco Nexus 9336C-FX2 A	1/22
	1/35/1 – 1/35/4	Cisco MDS 9124V B	fc1/1-fc1/4
	1/36/1 – 1/36/4	Cisco MDS 9124V B	fc1/5-fc1/8
	1/1 – 1/8	IFM-B	1/1-1/8

Cisco MDS ports detail

For this solution, eight ports (ports 1 to 8) of the MDS Switch A were connected to the Fabric Interconnect A (ports 1/35/1-4 and 1/36/1-4). The port-channel (PC 1) was configured on these ports between MDS-A and FI-A. Eight ports (ports 1 to 8) of the MDS Switch B were connected to the Fabric Interconnect B (ports 1/35/1-4 and ports 1/36/1-4). Another port-channel (PC 2) was created on these ports between MDS-B and FI-B. All the ports carry 32 Gb/s FC traffic. For the connectivity towards storage, eight ports (ports 9 to 16) of the MDS Switch A were connected to the NetApp AFF A1K Storage controllers. Eight ports (ports 9 to 16) of the MDS Switch B were connected to the NetApp AFF A1K Storage controllers. All ports carry 64 Gb/s FC traffic.

MDS-A

```
9124V-A(config)# show int brief
```

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel	Logical Type
fc1/1	101	auto	on	trunking	sw1	TF	32	1	core
fc1/2	101	auto	on	trunking	sw1	TF	32	1	core
fc1/3	101	auto	on	trunking	sw1	TF	32	1	core
fc1/4	101	auto	on	trunking	sw1	TF	32	1	core
fc1/5	101	auto	on	trunking	sw1	TF	32	1	core
fc1/6	101	auto	on	trunking	sw1	TF	32	1	core
fc1/7	101	auto	on	trunking	sw1	TF	32	1	core
fc1/8	101	auto	on	trunking	sw1	TF	32	1	core
fc1/9	101	auto	off	up	sw1	F	64	--	edge
fc1/10	101	auto	off	up	sw1	F	64	--	edge
fc1/11	101	auto	off	up	sw1	F	64	--	edge
fc1/12	101	auto	off	up	sw1	F	64	--	edge
fc1/13	101	auto	off	up	sw1	F	64	--	edge
fc1/14	101	auto	off	up	sw1	F	64	--	edge
fc1/15	101	auto	off	up	sw1	F	64	--	edge
fc1/16	101	auto	off	up	sw1	F	64	--	edge

MDS-B

```
9124V-B(config-if)# show int brief
```

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode	Oper Speed (Gbps)	Port Channel	Logical Type
fc1/1	102	auto	on	trunking	sw1	TF	32	2	core
fc1/2	102	auto	on	trunking	sw1	TF	32	2	core
fc1/3	102	auto	on	trunking	sw1	TF	32	2	core
fc1/4	102	auto	on	trunking	sw1	TF	32	2	core
fc1/5	102	auto	on	trunking	sw1	TF	32	2	core
fc1/6	102	auto	on	trunking	sw1	TF	32	2	core
fc1/7	102	auto	on	trunking	sw1	TF	32	2	core
fc1/8	102	auto	on	trunking	sw1	TF	32	2	core
fc1/9	102	auto	off	up	sw1	F	64	--	edge
fc1/10	102	auto	off	up	sw1	F	64	--	edge
fc1/11	102	auto	off	up	sw1	F	64	--	edge
fc1/12	102	auto	off	up	sw1	F	64	--	edge
fc1/13	102	auto	off	up	sw1	F	64	--	edge
fc1/14	102	auto	off	up	sw1	F	64	--	edge
fc1/15	102	auto	off	up	sw1	F	64	--	edge
fc1/16	102	auto	off	up	sw1	F	64	--	edge

Table 4 Cabling information for Cisco MDS 9124V

Local Device	Local Port	Remote Device	Remote Port
Cisco MDS 9124V A	fc1/1	Cisco UCS FI 6536 A	1/35/1
	fc1/2	Cisco UCS FI 6536 A	1/35/2
	fc1/3	Cisco UCS FI 6536 A	1/35/3
	fc1/4	Cisco UCS FI 6536 A	1/35/4

Local Device	Local Port	Remote Device	Remote Port
	fc1/5	Cisco UCS FI 6536 A	1/36/1
	fc1/6	Cisco UCS FI 6536 A	1/36/2
	fc1/7	Cisco UCS FI 6536 A	1/36/3
	fc1/8	Cisco UCS FI 6536 A	1/36/4
	fc1/9	NetApp AFF A1K 01	8a
	fc1/10	NetApp AFF A1K 01	8c
	fc1/11	NetApp AFF A1K 01	9a
	fc1/12	NetApp AFF A1K 01	9c
	fc1/13	NetApp AFF A1K 02	8a
	fc1/14	NetApp AFF A1K 02	8c
	fc1/15	NetApp AFF A1K 02	9a
	fc1/16	NetApp AFF A1K 02	9c
Cisco MDS 9124V B	fc1/1	Cisco UCS FI 6536 B	1/35/1
	fc1/2	Cisco UCS FI 6536 B	1/35/2
	fc1/3	Cisco UCS FI 6536 B	1/35/3
	fc1/4	Cisco UCS FI 6536 B	1/35/4
	fc1/5	Cisco UCS FI 6536 B	1/36/1
	fc1/6	Cisco UCS FI 6536 B	1/36/2
	fc1/7	Cisco UCS FI 6536 B	1/36/3
	fc1/8	Cisco UCS FI 6536 B	1/36/4
	fc1/9	NetApp AFF A1K 01	8b
	fc1/10	NetApp AFF A1K 01	8d
	fc1/11	NetApp AFF A1K 01	9b
	fc1/12	NetApp AFF A1K 01	9d

Local Device	Local Port	Remote Device	Remote Port
	fc1/13	NetApp AFF A1K 02	8b
	fc1/14	NetApp AFF A1K 02	8d
	fc1/15	NetApp AFF A1K 02	9b
	fc1/16	NetApp AFF A1K 02	9d

Table 5 lists the details of vNIC Ethernet LIFs, FC LIFs and NVMe/FC LIFs creation on UCS Server

Table 5 vNIC and vHBA configured on each Linux host

vNIC Details	Remarks
vNIC 0(eth0)	Management and Public Network Traffic Interface for Oracle. MTU = 1500
vHBA0	FC Network Traffic & Boot from SAN through MDS-A Switch
vHBA1	FC Network Traffic & Boot from SAN through MDS-B Switch
vHBA2	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-A Switch
vHBA3	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-B Switch
vHBA4	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-A Switch
vHBA5	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-B Switch
vHBA6	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-A Switch
vHBA7	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-B Switch
vHBA8	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-A Switch
vHBA9	NVMe/FC Network Traffic (Oracle Storage Traffic) through MDS-B Switch

Note: For this solution, we configured one VLAN to carry IB-Management network traffic as well as two VSANs to carry FC and NVMe/FC storage traffic as listed in Table 6.

Table 6 VLAN and VSAN Configuration

Name	ID	Remarks
VLAN		
Native VLAN	2	For Native VLAN
IB-Mgmt VLAN	2214	In-band Management VLAN for the Linux server

Name	ID	Remarks
VSAN		
VSAN-A	101	FC and NVMe/FC Network Traffic through for Fabric Interconnect A
VSAN-B	102	FC and NVMe/FC Network Traffic through for Fabric Interconnect B

Data Path for Fibre Channel traffic

For FCP and NVMe/FC protocols, the traffic between compute and storage flows via Cisco MDS switches over FC cables. Traffic isolation in MDS is achieved via VSAN and soft zoning configurations.

Zoning on the Fabric switch

VSANs are created prior to creating the zoning configuration to provide isolation.

Zoning for FCP and NVMe setup

The zoning was done to isolate the workloads to avoid any contention, to increase network security and to prevent data loss or corruption.

- Two vHBAs were used from each VIC adapter port for boot from SAN.
- A total of eight vHBAs were used on the VIC adapter for NVMe access.
- A total of sixteen NVMe logical interfaces [LIFs] were created, eight on each controller node for each corresponding fabric A and B.
- Two zones were created, on each MDS switch, one for FCP and another for NVMe.

Table 7 Zoning for FCP and NVMe setup

Fabric Name	Zone Name	Host vHBA (initiator)	Storage LIFs (Target)
Fabric A	zone1 (FC Boot)	Host-1-FCP-A Host-2-FCP-A Host-3-FCP-A Host-4-FCP-A	AFF-1-FC-lif-08a AFF-2-FC-lif-08a
	zone2 (NVMe)	Host-1-FC-NVMe-A Host-2-FC-NVMe-A Host-2-FC-NVMe-A Host-2-FC-NVMe-A	AFF-1-FC-NVMe-lif-08a AFF-1-FC-NVMe-lif-08c AFF-1-FC-NVMe-lif-09a AFF-1-FC-NVMe-lif-09c AFF-2-FC-NVMe-lif-08a AFF-2-FC-NVMe-lif-08c AFF-2-FC-NVMe-lif-09a AFF-2-FC-NVMe-lif-09c

Fabric Name	Zone Name	Host vHBA (initiator)	Storage LIFs (Target)
Fabric B	zone1 (FC Boot)	Host-1-FCP-B Host-2-FCP-B Host-3-FCP-B Host-4-FCP-B	AFF-1-FC-lif-08b AFF-2-FC-lif-08b
	zone2 (NVMe)	Host-1-FC-NVMe-B Host-2-FC-NVMe-B Host-2-FC-NVMe-B Host-2-FC-NVMe-B	AFF-1-FC-NVMe-lif-08b AFF-1-FC-NVMe-lif-08d AFF-1-FC-NVMe-lif-09b AFF-1-FC-NVMe-lif-09d AFF-2-FC-NVMe-lif-08b AFF-2-FC-NVMe-lif-08d AFF-2-FC-NVMe-lif-09b AFF-2-FC-NVMe-lif-09d

Hardware and Software Revisions

This solution can be extended to any FlexPod environment running supported software, firmware, and hardware versions as defined in the [NetApp Interoperability Matrix Tool](#) and [UCS Hardware and Software Compatibility](#). The following table shows the FlexPod hardware and software revisions used in this solution.

Table 8 Hardware and Software versions

Component	Product	Version
Compute	Cisco UCSX-210C-M7	5.2(0.230092)
	Cisco UCS Fabric Interconnects 6536	4.3(2.240002)
	Cisco VIC 15231	5.3(2.40)
	CPU	Intel(R) Xeon(R) Gold 6338N CPU @ 2.20GHz
Network	Cisco Nexus 9336C-FX2 NX-OS	10.2(6)
	Cisco MDS 9124V	9.3(2)
Storage	NetApp AFF A1K	ONTAP 9.15.1
Software	Red Hat Enterprise Linux Server	RHEL 8.10
	Cisco fNIC driver	kmod-fnic-2.0.0.92- 273.0.rhel8u9.x86_64.rpm
	Cisco Intersight Assist Virtual Appliance	1.0.11-4206
	Oracle Database 21c Grid Infrastructure	21.3.0.0.0
	Oracle Database 21c Enterprise Edition	21.3.0.0.0

Component	Product	Version
	SLOB	2.5.4.0

Installation and Configuration

FlexPod Deployment

To understand FlexPod with UCS X-Series, Oracle, and NetApp ONTAP design details, refer to FlexPod Datacenter with Cisco UCS X-Series [Deployment Guide](#). This document provides design and deployment guidance around incorporating the Cisco Intersight—managed UCS X-Series platform within the FlexPod infrastructure.

Compute Layout and Server port mapping

The VIC adapters support multiple virtual adapters (vHBA) to create IO parallelization and traffic isolation. Below is the server port mapping (vHBA) used in this solution:

- Four vHBAs for NVMe protocol were created on each Fabric for Data, Log, and Grid
- One vHBA for FCP protocol was created on each Fabric for SAN boot

Table 9 shows the server VIC mapping for both Fabric A and B.

Table 9 Server VIC Mapping

VIC Adapter	Type	Protocol	Virtual Port Name (vHBA)
Port1, Fabric-A	Data	NVMe	host1-nvme-A1, host1-nvme-A2, host1-nvme-A3, host1-nvme-A4
	Log		
	Grid		
	SANBoot	FCP	host1-fcp-A0
Port2, Fabric-B	Data	NVMe	host1-nvme-B1, host1-nvme-B2, host1-nvme-B3, host1-nvme-B4
	Log		
	Grid		
	SANBoot	FCP	host1-fcp-B0

Storage Configuration

Pre-requisites

- All the required licenses on the NetApp storage – FCP, NVMe
- MDS License
- Adapters supporting FC, NVMe

Storage Layout

This section describes the ONTAP configuration needed to do the storage layout for this solution. The following steps were used:

- One SVM for the FC boot for all Linux hosts and NVMe storage access for data and log to one Linux host
- Three SVMs for NVMe storage access for data and log to other remaining Linux hosts. Four NVMe LIFs (Logical Interfaces) per SVM were created with dedicated FC interface
- To provide boot LUN to all RHEL hosts, one volume was created on the first SVM
- For NVMe storage access (data and log), 18 volumes (per SVM), 9 on each A1K controller, were created to distribute the data and log traffic and the optimal CPU parallelization on the storage
- Eighteen namespaces were created for data and log volumes per SVM for all four Oracle servers
- One volume and namespace were created for grid per SVM for all four Oracle servers
- One subsystem for each Linux host to map all the data and log namespaces per SVM

Note: Below configuration is an example of one SVM, for validation total 4 SVMs were used.

1. Configure SVM with NVMe/FC service running.

```
AFF-A1K::> vserver create -vserver vs1
AFF-A1K::> vserver add-protocols -vserver vs1 -protocols nvme,fc
AFF-A1K::> vserver fcp create -vserver vs1 -status-admin up
AFF-A1K::> vserver nvme create -vserver vs1 -status-admin up
```

2. Display the NVMe/FC capable adapters installed in the cluster.

```
AFF-A1K::> fcp adapter show -data-protocols-supported fc-nvme
(network fcp adapter show)
```

Node	Adapter	Connection Established	Port Address	Admin Status	Operational Status
AFF-A1K-01	8a	true	900080	up	online
AFF-A1K-01	8b	true	240180	up	online
AFF-A1K-01	8c	true	900040	up	online
AFF-A1K-01	8d	true	240160	up	online
AFF-A1K-01	9a	true	900060	up	online
AFF-A1K-01	9b	true	2401e0	up	online
AFF-A1K-01	9c	true	900020	up	online
AFF-A1K-01	9d	true	2401c0	up	online
AFF-A1K-02	8a	true	9000a0	up	online
AFF-A1K-02	8b	true	240120	up	online
AFF-A1K-02	8c	true	9000c0	up	online
AFF-A1K-02	8d	true	240140	up	online
AFF-A1K-02	9a	true	9000e0	up	online
AFF-A1K-02	9b	true	240200	up	online
AFF-A1K-02	9c	true	900100	up	online
AFF-A1K-02	9d	true	2401a0	up	online

16 entries were displayed.

Note: You can configure only two NVMe LIFs per node per SVM. To take advantage of all 16 FC ports on running NVMe/FC traffic, similarly 3 more SVMs and respective NVMe/FC LIFs were created.

3. Display the NVMe/FC LIFs created on the first SVM.

```
AFF-A1K::> net interface show -vserver vs1 -data-protocol fc-nvme
(network interface show)
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	fc-nvme-8a-1a	up/up	20:26:d0:39:ea:a9:f6:9e	AFF-A1K-01	8a	true

```

fc-nvme-8a-1b
    up/up    20:22:d0:39:ea:a9:f6:9e
                AFF-A1K-02    8a    true
fc-nvme-8b-2a
    up/up    20:27:d0:39:ea:a9:f6:9e
                AFF-A1K-01    8b    true
fc-nvme-8b-2b
    up/up    20:23:d0:39:ea:a9:f6:9e
                AFF-A1K-02    8b    true

```

4 entries were displayed.

4. Create a FlexVol volume for Boot, Data, Log, and Grid. The namespaces were created, each with 400GB for Data, 100GB for Log and 10 GB for Grid.

Note: NetApp recommends using larger sizes to accommodate future data growth.

```

vol create -vsriver vs1 -volume boot1 -aggregate AFF_A1K_01_NVME_SSD_1 -size 1t -state online -
policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -
tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state
disabled

vol create -vsriver vs1 -volume /vol/oracle1_data1/ns_data1 -aggregate AFF_A1K_01_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data2/ns_data2 -aggregate AFF_A1K_02_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data3/ns_data3 -aggregate AFF_A1K_01_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data4/ns_data4 -aggregate AFF_A1K_02_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data5/ns_data5 -aggregate AFF_A1K_01_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data6/ns_data6 -aggregate AFF_A1K_02_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data7/ns_data7 -aggregate AFF_A1K_01_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data8/ns_data8 -aggregate AFF_A1K_02_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data9/ns_data9 -aggregate AFF_A1K_01_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

vol create -vsriver vs1 -volume /vol/oracle1_data10/ns_data10 -aggregate AFF_A1K_02_NVME_SSD_1 -
size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy
none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -
anti-ransomware-state disabled

```

```
vol create -vserver vs1 -volume /vol/oracle1_data1/ns_data1 -aggregate AFF_A1K_01_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_data2/ns_data2 -aggregate AFF_A1K_02_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_data3/ns_data3 -aggregate AFF_A1K_01_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_data4/ns_data4 -aggregate AFF_A1K_02_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_data5/ns_data5 -aggregate AFF_A1K_01_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_data6/ns_data6 -aggregate AFF_A1K_02_NVME_SSD_1 -size 400GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_log1/ns_log1 -aggregate AFF_A1K_01_NVME_SSD_1 -size 100GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_log2/ns_log2 -aggregate AFF_A1K_02_NVME_SSD_1 -size 100GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

```
vol create -vserver vs1 -volume /vol/oracle1_grid/grid -aggregate AFF_A1K_02_NVME_SSD_1 -size 10GB -state online -policy default -unix-permissions ---rwxr-xr-x -type RW -snapshot-policy none -foreground true -tiering-policy none -analytics-state off -activity-tracking-state off -anti-ransomware-state disabled
```

5. Create FC LUN for SAN Boot.

```
lun create -vserver vs1 -path /vol/boot1/rhell -size 200g -ostype linux
```

6. Create igroup for SAN boot and map the LUN. The interface names - rhel1-fcboot-a, rhel1-fcboot-b should be replaced with the associated wwpn.

```
igroup create -vserver vs1 -igroup rhell -protocol fcp -ostype linux -initiator host1-fcboot-a, host1-fcboot-b
```

```
lun map -vserver vs1 -path /vol/boot1/rhell -igroup 1
```

7. Create NVMe namespaces.

```
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data1/ns_data1 -size 350g -ostype linux  
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data2/ns_data2 -size 350g -ostype linux  
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data3/ns_data3 -size 350g -ostype linux  
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data4/ns_data4 -size 350g -ostype linux
```



```

vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data5/ns_data5 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data6/ns_data6 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data7/ns_data7 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data8/ns_data8 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data9/ns_data9 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data10/ns_data10 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data11/ns_data11 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data12/ns_data12 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data13/ns_data13 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data14/ns_data14 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data15/ns_data15 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_data16/ns_data16 -size 350g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_log1/ns_log1 -size 75g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_log2/ns_log2 -size 75g -ostype
linux
vserver nvme namespace create -vserver vs1 -path /vol/oracle1_grid/grid -size 5g -ostype linux

```

8. Create subsystem for NVMe.

```
nvme subsystem create -vserver vs1 -subsystem oracle1 -ostype linux
```

9. Add host NVMe NQN to the subsystem.

```
nvme subsystem host add -vserver vs1 -subsystem oracle1 -host-nqn nqn.2014-08.sa.perf:nvme:host1
-io-queue-count 15 -io-queue-depth 128
```

Note: On the Linux host, check the host NQN string at `/etc/nvme/hostnqn`

10. Map NVMe namespaces to the “oracle1” subsystem.

```

vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data1/ns_data1 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data2/ns_data2 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data3/ns_data3 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data4/ns_data4 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data5/ns_data5 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data6/ns_data6 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data7/ns_data7 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data8/ns_data8 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data9/ns_data9 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data10/ns_data10 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data11/ns_data11 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data12/ns_data12 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data13/ns_data13 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data14/ns_data14 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data15/ns_data15 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_data16/ns_data16 -subsystem
oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_log1/ns_log1 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_log2/ns_log2 -subsystem oracle1
vserver nvme subsystem map add -vserver vs1 -path /vol/oracle1_grid/grid -subsystem oracle1

```

Bare Metal

In this solution, we have used RHEL 8.10 for BareMetal installation using FC boot. NVMe over Fabrics (NVMe-oF), including NVMe over Fibre Channel (NVMe/FC) and other transports, is supported with Red Hat Enterprise Linux (RHEL) 8.10 with Asymmetric Namespace Access (ANA). In NVMe-oF environments, ANA is the equivalent of ALUA multipathing in iSCSI and FC environments and is implemented with in-kernel NVMe multipath.

Note: The `fnic` driver should be installed on the server.

Enable device-mapper multipathing:

The DM-Multipath was configured only for the FC Boot LUNs.

1. Enable and initialize the multipath configuration file.

```
[root@oracle-bm-101 ~]# mpathconf --enable

[root@oracle-bm-101 ~]# systemctl status multipathd.service

[root@oracle-bm-101 ~]# mpathconf
multipath is enabled
find_multipaths is yes
user_friendly_names is enabled
default_property blacklist is disabled
enable_foreign is set (foreign multipath devices may not be shown)
dm_multipath module is loaded
multipathd is running
```

2. Now, edit the `"/etc/multipath.conf"` file.

```
[root@oracle-bm-101 ~]# cat /etc/multipath.conf

defaults {

    find_multipaths yes
    user_friendly_names yes
    enable_foreign NONE
}

multipaths {

    multipath {
        wwid      3600a098038314d44315d576341706364
        alias     oracle1
    }
}
```

Note: NetApp recommends using the in-kernel NVMe multipath for ONTAP namespaces and `dm-multipath` for ONTAP LUNs respectively. This should exclude the ONTAP namespaces from `dm-multipath` and prevent `dm-multipath` from claiming these namespace devices. You can do this by adding the `enable_foreign` setting to the `"/etc/multipath.conf"` file:

Enable native NVMe multipathing:

1. The default kernel setting for the `nvme_core.multipath` option is set to "N", which means that the native Non-volatile Memory Express™ (NVMe) multipathing is disabled.

```
[root@oracle-bm-101 ~]# cat /sys/module/nvme_core/parameters/multipath
```

2. Use the below command on the RHEL server to enable in-kernel multipath:

```
[root@oracle-bm-101 ~]# uname -r
4.18.0-553.16.1.el8_10.x86_64

root@oracle-bm-101 ~]# grubby --args=nvme_core.multipath=Y --update-kernel /boot/vmlinuz-4.18.0-553.8.1.el8_10.x86_64
```

3. Reboot the server.

4. Once the server is up, verify the in-kernel NVMe multipath is enable:

```
[root@oracle-bm-101 ~]# cat /sys/module/nvme_core/parameters/multipath
Y
```

5. Verify the NVMe disks detection.

```
[root@oracle-bm-101 ~]# nvme list
```

Node Usage	SN Format	Model FW Rev	Namespace
/dev/nvme8n1 Controller	81MDxFWUzRYvAAAAAAG 1	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n10 Controller	81MDxFWUzRYvAAAAAAG 10	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n11 Controller	81MDxFWUzRYvAAAAAAG 11	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n12 Controller	81MDxFWUzRYvAAAAAAG 12	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n13 Controller	81MDxFWUzRYvAAAAAAG 13	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n14 Controller	81MDxFWUzRYvAAAAAAG 14	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n15 Controller	81MDxFWUzRYvAAAAAAG 15	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n16 Controller	81MDxFWUzRYvAAAAAAG 16	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n17 Controller	81MDxFWUzRYvAAAAAAG 17	NetApp ONTAP 80.53 GB / 80.53 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n18 Controller	81MDxFWUzRYvAAAAAAG 18	NetApp ONTAP 80.53 GB / 80.53 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n19 Controller	81MDxFWUzRYvAAAAAAG 19	NetApp ONTAP 5.37 GB / 5.37 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n2 Controller	81MDxFWUzRYvAAAAAAG 2	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n3 Controller	81MDxFWUzRYvAAAAAAG 3	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n4 Controller	81MDxFWUzRYvAAAAAAG 4	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n5 Controller	81MDxFWUzRYvAAAAAAG 5	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n6 Controller	81MDxFWUzRYvAAAAAAG 6	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n7 Controller	81MDxFWUzRYvAAAAAAG 7	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n8 Controller	81MDxFWUzRYvAAAAAAG 8	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF
/dev/nvme8n9 Controller	81MDxFWUzRYvAAAAAAG 9	NetApp ONTAP 375.81 GB / 375.81 GB	4 KiB + 0 B FFFFFFFF

Solution Validation

SLOB Test

The Silly Little Oracle Benchmark (SLOB) is a comprehensive toolkit designed for generating and testing I/O operations within an Oracle database. SLOB is highly effective in evaluating the I/O subsystem using authentic Oracle System Global Area (SGA)-buffered physical I/O. It facilitates the testing of physical random single-block reads (db file sequential read) and random single-block writes (Database Writer) flushing capability. Typically, SLOB issues single block reads for the read workload, which are generally 8K in size, corresponding to the database block size.

For testing the SLOB workload, we created a pluggable database named ORADBP1 with following layout:

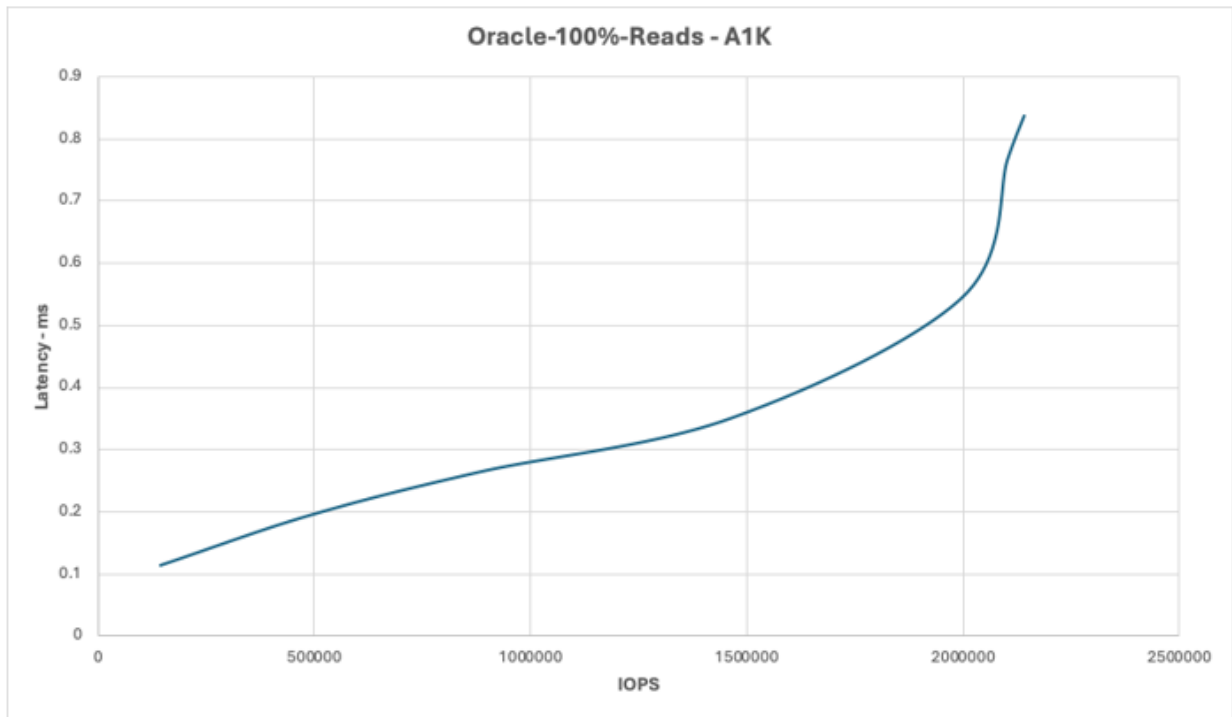
- 18 volumes were created on the storage, 9 on each storage node
- One namespace was created on each volume, with 400GB for Data and 100GB for Log
- A total of two Automatic Storage Management (ASM) disk groups were created one each for Data and Log with external redundancy
 - Data was created with 16 namespaces
 - Log was created with 2 namespaces

These ASM disk groups provided the necessary storage to create one big file tablespaces for the SLOB database. We loaded the SLOB schema onto the “DATA” disk group, with a total size of up to 3 TB.

We utilized SLOB2 to generate our OLTP workload. Each database server applied the workload to the Oracle database and log files. Various tests were conducted, and metrics such as IOPS and latency were captured, along with Oracle Automatic Workload Repository (AWR) reports for each test scenario.

User Scalability Test

SLOB2 was configured to run against all the four Oracle servers and the concurrent users were equally spread across all the nodes. We tested the environment by increasing the number of Oracle users in database from a minimum of 2 users up to a maximum of 56 users across all the servers with 100% read. At each load point, we verified that the storage system and the servers could maintain steady-state behaviour without any issues. We also made sure that there were no bottlenecks across servers or networking systems.



Conclusion

FlexPod is a predesigned, best-practice, next generation data center architecture that is built on the cloud-scale technologies offered by Cisco and NetApp, enabling management synergies across the complete IT infrastructure environment. FlexPod is an optimal platform for both BareMetal and Virtualization deployments of enterprise applications and databases.

While running the workload on AFF A1K, we observed a 21% performance improvement compared to the last published Oracle solution on [FlexPod with AFF A900](#). The test results indicate that the customers will be able to run large Oracle workloads and consolidate multiple database workloads, achieving cumulative IOPS to meet their SLA requirements.

Where to find additional information

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

- FlexPod Home Page: <https://www.flexpod.com>
- NetApp Interoperability Matrix Tool: <http://support.netapp.com/matrix>
- Cisco UCS Hardware and Software Interoperability Tool: <http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>
- NetApp Product Documentation: <https://www.netapp.com/support-and-training/documentation>
- NetApp AFF A-Series: <https://www.netapp.com/data-storage/aff-a-series>
- NetApp Product Documentation <https://www.netapp.com/support-and-training/documentation/>
- FlexPod Design Zone: <https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/index.html>

Version history

Version	Date	Document version history
Version 1.0	September 2024	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 © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT,

INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

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

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

Trademark information

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