



NetApp Verified Architecture

NetApp and Broadcom Modern SAN Cloud-Connected Flash Solution

Oracle and SUSE NetApp Verified Architecture Design Edition

Modernize your enterprise SAN with end-to-end NVMe over Fibre Channel, the fastest cloud-ready solution for mission-critical workloads

Contributors

Broadcom Inc.: Ismail Muhammad, Jim Zucchero, and Naem Saafein

NetApp: Adam Fore, Andy Grimes, Danny Hohman, Darryl Clinkscales, Jeffrey Steiner, Joe Carter, Lee Howard, Mike Peppers, Pavan Jhamnani, Rob McDonald, Scott Lane, Stetson Webster, Steve Botkin, Steve Collins, and Vidula Aiyer

October 2018 | NVA-1126-DESIGN | Version 2.0

In partnership with



Abstract

This NetApp® Verified Architecture has been jointly designed and verified by NetApp and Broadcom Inc. It uses the latest Brocade, Emulex, and SUSE technology solutions along with NetApp all-flash storage, which sets a new standard for enterprise SAN storage and data protection that drives superior business value.

Forward: Thoughts from Broadcom

Storage system customers face never-ending challenges to keep a competitive edge, to scale for hypergrowth opportunities, and to manage ever-increasing needs for IT services within their organizations. It's imperative that they work with and invest in the right strategic partners who can help them meet the business-critical needs of today and tomorrow.

With its rich portfolio of products, Broadcom has the distinction of being the market leader in enterprise Fibre Channel SAN solutions. Broadcom has a rich heritage of technology leadership that it provides to enterprise OEM customers. The right OEM partners uphold Broadcom's high standards for excellence and for meeting customer expectations.

For over 15 years, NetApp has provided FC solutions to customers as a Broadcom OEM partner. With NetApp's market leadership position in Non-Volatile Memory Express (NVMe) over Fibre Channel (NVMe/FC) technology, NetApp is enabling its own customers to deliver superior IT performance for their most important, mission-critical enterprise SAN applications. Broadcom is proud to partner with NetApp, a company that continues to demonstrate the highest degree of excellence in its future-forward vision and technology, which will take customers into the next decade of enterprise SANs.

Jim Zucchero
Global Sales Executive, NetApp Team
Broadcom Inc.

TABLE OF CONTENTS

Forward: Thoughts from Broadcom	2
1 Executive Summary	5
2 The Challenge	5
3 The Solution	5
3.1 Reasons to Modernize Your SAN with NetApp and Broadcom.....	6
3.2 The Architecture.....	6
4 Program Summary	7
5 Solution Overview	8
5.1 NetApp and Brocade Enterprise SAN Solution Benefits	8
5.2 Target Audience.....	8
5.3 Solution Technology	9
5.4 Use-Case Summary.....	14
6 Technology Requirements	15
6.1 Hardware Requirements	15
6.2 Software Requirements	16
6.3 Network Design.....	16
6.4 Workload Design.....	17
7 Solution Verification	18
7.1 Test Methodology	18
7.2 Test Results	18
8 Future Disruptive Innovation	20
9 Conclusion	21
Where to Find Additional Information	22

LIST OF TABLES

Table 1) Cost-benefit analysis of the joint NetApp and Broadcom solution.	7
Table 2) Comparison of legacy SAN and NetApp and Brocade enterprise SAN.	7
Table 3) Hardware requirements for the joint solution.	15
Table 4) Software requirements for the joint solution.	16

LIST OF FIGURES

Figure 1) Technology components of the NetApp and Broadcom validated solution.	9
---	---

Figure 2) Component families of the NetApp and Broadcom joint architecture.	10
Figure 3) Components of Brocade's SAN Health tool.....	12
Figure 4) Steps required to run and to use SAN Health.	12
Figure 5) NetApp data protection components.	13
Figure 6) Service consumption metrics for IT services.	13
Figure 7) NetApp and Broadcom validated architecture testbed layout.....	17
Figure 8) Total IOPS vs. read latency.	20

1 Executive Summary

NetApp Verified Architectures describe systems and solutions that are designed, tested, and documented to facilitate and to improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp has developed to help meet the business needs of customers.

This NetApp Verified Architecture offers a solution that modernizes your enterprise SAN with end-to-end NVMe, giving your company the fastest cloud-ready solution for mission-critical workloads.

This report addresses:

- The challenges that organizations face today with data assets and infrastructure
- The solution to leverage disruptive future technology nondisruptively for your business today
- Reasons to modernize your traditional SAN infrastructure
- A world-class, modern, SAN-verified reference architecture
- NetApp recommended data protection solutions for this architecture
- Financial analysis that illustrates a self-funding TCO business case for modernizing SAN infrastructure yielding the following benefits:
 - A 50% to 80% reduction in database licensing demands
 - An 80% to 90%+ reduction in data center floor space
 - A 50% to 90%+ reduction in power and cooling
 - A 50% to 80% reduction in labor costs

2 The Challenge

The challenge today is how to rapidly and nondisruptively transform, modernize, and streamline critical data and IT services to scale and to adapt to customer and business needs. At the same time, these services must be future-proof and cloud-ready so that the organization can maintain a competitive edge.

Background: According to IDC, by 2020, 50% of Forbes Global 2000 companies will see most of their business depend on their ability to create digitally enhanced products, services, and experiences. Data is the lifeblood of future-thinking companies. The growth and dynamism of this avalanche of new data require modern companies to move in real time with the marketplace. However, for many, their current IT infrastructure isn't up to the task. The growing stress on the entire IT infrastructure to manage this overload of data interferes with the ability to quickly capitalize on the inherent value of the data.

3 The Solution

The good news is that, just as flash transformed enterprise storage a few years ago, a new emerging technology, NVMe, is poised to transform the enterprise again. NVMe, an emergent storage access and transport protocol, delivers the fastest response times yet for business-critical enterprise applications. Therefore, NVMe is about to provide a major speed boost for enterprise data storage systems. However, this time, the transformative effect could be greater still because NVMe isn't just a storage specification. The broader NVMe over Fabrics (NVMe-oF) protocol rearchitects the entire data path, from server to storage system, enabling superior performance and lower latency than traditional technologies can deliver.

As a result, CxOs now have the opportunity, and the challenge, to harness the power of data through digital transformation and modernization. They can also use these emerging best-in-class technologies from world-class industry leaders NetApp and Broadcom's Brocade and Emulex divisions to perform the following tasks:

- Rapidly deliver and monetize vital digital data services
- Accelerate the pace of Innovation
- Acquire, grow, and retain market share
- Improve customer service and experience
- Maximize return on investment
- Protect and secure customers and critical data
- Increase agility and response to changing business needs

3.1 Reasons to Modernize Your SAN with NetApp and Broadcom

This document describes a verified, unified, modern SAN solution reference architecture designed by industry leaders Broadcom and NetApp that is a first-to-market enterprise NVMe/FC solution. NetApp and Broadcom provide an end-to-end NVMe-powered solution, from host to storage controller, that can help you realize the promise and the benefits of NVMe technology right now. With a system that yields the fastest access, management, and utilization of critical data, you can accelerate your time to innovation and realize the following benefits:

- **Digitally transform critical business applications.** Enable the next generation of your critical applications so that they're ready for analytics, artificial intelligence (AI), and machine learning.
- **Harness the power of the hybrid cloud.** Cloud-enable your IT services to get the benefits of on-premises storage with the flexibility of a public cloud.
- **Get a best-in-class solution for enterprise SAN.** Strengthen your competitive advantage by partnering with the fastest-growing leaders in flash, SAN, data fabrics, and host bus adapters (HBAs).
- **Significantly simplify operations.** Improve IT responsiveness through the simplification of SAN management while ensuring predictable performance and reducing labor costs 50% to 80%.
- **Reduce costs through physical consolidation.** Modernization can produce an 80% to 90%+ reduction in data center floor space and a 50% to 90%+ reduction in power and cooling. Combine complicated, multifabric legacy Fibre Channel (FC) SAN switches into a simpler, modern architecture with fewer but faster components.
- **Increase license and CPU productivity.** Consolidation, including both containerization and virtualization, provides maximum utilization of the existing compute infrastructure and CPU-licensed software.
- **Future-proof your SAN environment.** Nondisruptively adopt new performance and technology advancements when you're ready.
- **Rapidly deliver core IT services.** Take advantage of an open platform that supports leading DevOps toolsets to vastly reduce the time-to-value for development.
- **Don't compromise on availability.** Get 99.9999% availability (IDC audit of 210,000 systems for a year, <5 seconds down/year) and enterprise-grade disaster recovery (DR) capabilities.
- **Improve the customer experience.** Accelerate performance, enable instant application cloning, and enable granular data recovery to improve the user experience.
- **Get next-generation enterprise data management.** Bring the value of industry-leading innovation together with enterprise availability to deliver the next generation of your SAN environment.

3.2 The Architecture

This NetApp and Broadcom modern SAN NetApp verified reference architecture includes the following key NetApp and Broadcom technologies: NVMe-oF, Oracle 12c, SUSE, and sixth-generation host and fabric technology. If you adopt them all, you can realize game-changing performance benefits with end-to-end visibility through the Fabric Vision technology. In the future, you can add storage-class memory and persistent memory to realize further increases in performance.

4 Program Summary

This report is part of the Modern SAN Best Practices Program, which provides test and validated design and configuration recommendations for the next generation of NVMe-powered fabrics. This report is part of a series that covers the deployment of popular enterprise applications.

This program is a collaboration between NetApp and Broadcom's Brocade and Emulex divisions, which together developed the industry's first end-to-end enterprise NVMe architecture. The information is designed to support IT organizations that upgrade their existing SAN architectures to next-generation NVMe-based fabrics to meet the low-latency and high-bandwidth requirements of current and future enterprise apps.

This report describes the system and solution that were designed, tested, and documented to facilitate modern SAN deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that NetApp has developed to meet the business needs of customers like you. This report also describes the design choices and the best practices for this shared infrastructure platform. These design considerations and recommendations are not limited to the specific components that are described in this document; they also apply to other component versions.

The solution that is described in this report offers the following TCO benefits:

- A 50% reduction in database licensing costs
- An 80% to 90%+ reduction in data center floor space
- A 50% to 90%+ reduction in power and cooling
- A 50% to 80% reduction in labor costs

Table 1 shows a cost-benefit analysis, and Table 2 compares legacy SAN and SAN that incorporates the joint solution.

Table 1) Cost-benefit analysis of the joint NetApp and Broadcom solution.

Value	Analysis Results
Return on investment (ROI)	93%
Net present value (NPV)	>\$2 million
Payback period (months)	6 months
Cost reduction	More than \$2.2 million saved over a 3-year analysis period compared to the legacy SAN storage system
Savings on power and space	\$390,000
Administration cost savings	\$230,000

Table 2) Comparison of legacy SAN and NetApp and Brocade enterprise SAN.

	Legacy SAN	NetApp Brocade SAN
Host connectivity	FC	FC, NVMe/FC
Next-generation NVMe support	No	Yes
Unified storage	No	Yes
Staff to manage	2 FTE	½ FTE
Bandwidth	8Gb average (max 16Gb FC)	32Gb

	Legacy SAN	NetApp Brocade SAN
Data migrations	Required	No
Data center footprint	Large	Small

In addition, by integrating secondary storage into your SAN and flash infrastructure, your company can better protect and secure your data while reducing overall costs. Your secondary storage can be a combination of NetApp all-flash arrays for short-term recovery and either an on-premises object store (for example, NetApp StorageGRID® Webscale) or a public cloud hyperscaler (for example, Amazon Web Services [AWS] or Microsoft Azure) for longer-term retention.

5 Solution Overview

5.1 NetApp and Brocade Enterprise SAN Solution Benefits

This solution is composed of Brocade Gen 6 Fibre Channel Switches, Emulex Gen 6 FC HBAs, and NetApp AFF storage systems. It is a predesigned, best-practice configuration that is built on SAN NVMe/FC on the latest NetApp and Broadcom technologies.

This solution delivers a baseline configuration and can also be sized and optimized to accommodate many different use cases and requirements. It supports tight integration with virtualized and cloud infrastructures and data protection, making it the logical choice for long-term investment.

The solution delivers operational efficiency and consistency with the versatility to meet various SLAs and IT initiatives, including:

- Application rollouts or migrations
- Business continuity
- Cloud delivery models (public, private, and hybrid) and service models (infrastructure as a service [IaaS], platform as a service [PaaS], and software as a service [SaaS])
- Asset consolidation and virtualization
- Data center consolidation and footprint reduction

Brocade and NetApp have thoroughly validated and verified this solution architecture and its many use cases. They have also created a portfolio of detailed documentation, information, presales and postsales services, and references to assist you in transforming your data center to this shared infrastructure model. This portfolio includes, but is not limited to, the following items:

- Best practice architectural design
- Workload sizing and scaling guidance
- Implementation and deployment instructions
- Technical specifications (rules for what is and what is not a reference architecture)
- Frequently asked questions
- NetApp and Brocade jointly validated designs that focus on various use cases

5.2 Target Audience

The target audience for this NetApp Verified Architecture report includes the following groups:

- **The CIO, CTO, and CFO**, who can benefit from the executive summary, use case examples, ROI and TCO information, and information about future strategies
- **Business information officers**, who can learn new ways to serve line-of-business owners with benefits from modern technologies

- **Architects, administrators, and solutions engineers** who are responsible for designing and deploying infrastructure for enterprise mission-critical applications
- **Database administrators**, who require new data management capabilities and performance to serve evolving data requirements
- **Application owners**, who need real-time, lower-latency data to feed current and newer generations of applications
- **Data architects**, who require platforms that are designed to enable more real-time analytics and to serve the AI and machine learning requirements that new workloads need
- **Cloud architects**, who must harness the power of the hybrid cloud and leverage core and cloud-native solutions
- **Backup administrators**, who must protect data and leverage new innovations to make data protection seamless and nondisruptive to the business
- **Service delivery managers**, who must meet SLAs and service-level objectives (SLOs) that require IT infrastructure and solutions to promote consistent and predictable results

5.3 Solution Technology

This report focuses on database and analytic workload types. These examples factor into Oracle 12c savings around infrastructure and licensing, because Oracle 12c migration forces infrastructure migration. We assume some numbers for typical inefficient utilization rates that we see on legacy storage. We also factor in our 2:1 to 4:1 storage efficiency and workload multitenancy benefits when consolidating multiple traditional SAN storage systems into a NetApp AFF A700s configuration.

Figure 1 shows the technology components of the joint solution, and Figure 2 shows the component families of the architecture. Implementation of this solution should reduce the footprint, management overhead, maintenance spending, and power and cooling costs, and it should also improve service availability and performance.

Figure 1) Technology components of the NetApp and Broadcom validated solution.


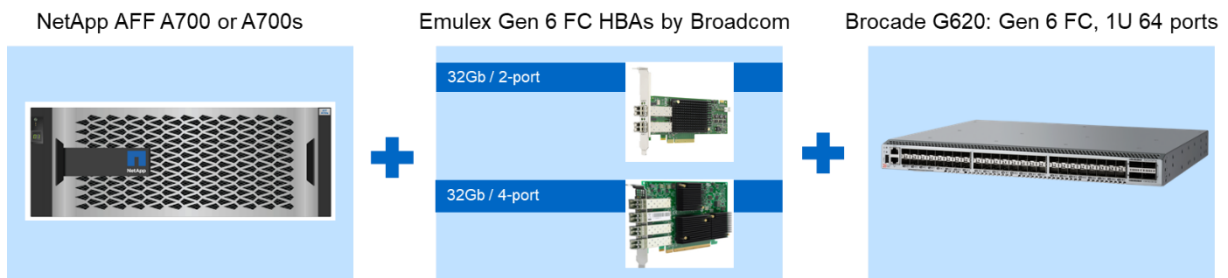


Figure 2) Component families of the NetApp and Broadcom joint architecture.



Most of today's all-flash arrays are deployed on low-risk, multiqueue-capable, deep-queue-rich, and proven FC-based storage networks, with their robust scalable fabric services and credit-based flow control. Because of their reliability and deterministic performance, FC fabrics are the most widely implemented storage network infrastructure for mission-critical applications. Because little change is required in the standards to implement NVMe/FC, the introduction of NVMe/FC along with existing storage is easy, seamless, and noninvasive. And because NVMe/FC can use the same infrastructure components concurrently with other FC traffic, it's easy to migrate workloads at a pace that works for your organization. NVMe/FC also allows the efficient transfer of NVMe commands and structures end to end with no translations.

The world's first end-to-end enterprise NVMe/FC solution with a NetApp all-flash array and Brocade Gen 6 Fibre Channel network is purpose-built for tomorrow's mission-critical workloads by leveraging today's infrastructure.

New innovations in storage technology are disrupting the data center industry. The introduction of faster media types and more efficient mechanisms to access media across various well-defined infrastructures is unlocking unprecedented speeds, lower latencies, and dramatic improvements in system and application efficiency and performance. These benefits are based on three advances: NVMe, NVMe-oF, and new storage-class memory (SCM).

Although this configuration uses Gen 6 (and other hardware), you can also use Gen 5 switches and other NetApp controllers, such as the AFF A300, the AFF A700, and the AFF A800.

NVMe

The NVMe specification is designed to use nonvolatile memory in all kinds of compute environments, from mobile phones to webscale service providers. It adds massive I/O path parallelization (64,000 I/O queues, each with a queue depth of up to 64,000 outstanding I/Os), making communication with storage systems massively parallel. Because of lower protocol overhead and lower-latency connectivity between servers and storage devices, this parallelization provides greater bandwidth.

The massive number of queues and the huge queue depths that each can support allow today's storage and servers to use increasingly large numbers of cores and memory. This capability accelerates processing of I/O threads by spreading the processing across multiple CPU cores. This attribute is critical for bringing together traditional enterprise applications with real-time analytics workloads, enabling new digital services for the modern enterprise.

NetApp technology is built for the future. With the industry's only unified data management platform that supports SAN or NAS, all-flash, software-defined, hybrid, and cloud storage, it supports both existing (traditional) and emerging applications (for example, NoSQL databases and AI). These features and capabilities are all part of the NetApp Data Fabric. NetApp systems support scaling (up and out) dynamically in seconds or minutes, instead of in hours or days. You can allocate applications to where they run best across your Data Fabric, whether it's on the premises or in the cloud. And to maximize performance and reduce overall storage cost, NetApp FabricPool technology allows you to move data automatically between AFF storage solutions and cloud storage tiers.

Along with Broadcom's Brocade and Emulex divisions, which are leaders in the SAN fabric space, NetApp is the first to market with an end-to-end enterprise NVMe/FC solution over a 32Gbps FC fabric. With this joint solution, you can enable and accelerate this digital transformation for your enterprise—now.

Brocade G620 Gen 6 Fibre Channel Switches

Broadcom Brocade has been the leading provider of storage networking solutions worldwide for more than 20 years, supporting the mission-critical systems and business-critical applications of many large enterprises. Brocade networking solutions help organizations achieve their critical business initiatives as they transition to a world where applications and information can reside anywhere. Today, Brocade is extending its proven data center expertise across the entire network with open, application-optimized, and efficient solutions that are built for consolidation and unmatched business agility.

The sixth generation of Fibre Channel is aimed at satisfying the needs of growing deployments of flash storage, hyperscale virtualization, and new high-speed data center architectures such as NVMe. Brocade G620 Gen 6 Fibre Channel switches shatter application performance barriers with up to 100 million IOPS and 32Gb/128Gb FC performance to meet the demands of flash-based storage workloads. Pay-as-you-grow scalability enables organizations like yours to scale from 24 to 64 ports so that you can support your evolving storage environments.

Brocade's IO Insight is the industry's first integrated network sensor tool that proactively and nonintrusively monitors real-time storage I/O health and performance statistics for both SCSI and NVMe traffic. It can do this from any device port on a Gen 6 FC platform. IO Insight then applies this information within an intuitive, policy-based monitoring and alerting suite to quickly identify the root cause of problems at the storage or at the virtual machine (VM) tier.

With standards-based, end-to-end VM tagging, Brocade VM Insight seamlessly monitors VM performance throughout a storage fabric. Administrators can quickly determine the source of VM or application performance anomalies and can provision and fine tune the infrastructure based on VM or application requirements to meet critical SLAs and SLOs.

The NVMe/FC feature supports both NVMe-oF and SCSI over FC protocols concurrently. Your organization can seamlessly integrate Brocade Gen 6 Fibre Channel networks with the next generation of low-latency flash storage, without a disruptive rip and replace operation.

Emulex Gen 6 FC HBAs

Emulex FC HBAs by Broadcom are designed to meet the demanding performance, reliability, and management requirements of modern networked storage systems that use high-performance and low-latency solid-state drives (SSDs). The latest Emulex LPe32002 FC HBAs with a dynamic multicore architecture deliver an industry-leading 1.6 million IOPS to any port that needs it, providing high performance when and where it's needed. The LPe32000-series provides 3,200MBps per link and up to 12,800MBps per card of throughput. It also provides low latency, enhanced manageability, and the highest reliability in the industry (10 million hours mean time between failure) for maximum uptime.

The secure firmware update feature protects the authenticity of device firmware. Emulex Gen 6 FC HBAs are NVMe/FC-enabled, delivering up to 55% lower insertion latency for NVMe/FC than SCSI over Fibre Channel. And for investment protection, these FC HBAs also concurrently support both NVMe/FC and SCSI over Fibre Channel protocols.

Brocade SAN Health

Your storage architecture is crucial for your business agility and success. Brocade's free SAN Health tool delivers clear insights into performance, inventory, and bottlenecks to optimize your SAN infrastructure and align it with your business needs. This hardware-agnostic and easy-to-run tool generates personalized storage network performance and inventory reports to help you prevent issues, avoid

application downtime, reduce troubleshooting time to resolution, and improve capacity planning and productivity. Figure 3 shows the components of the SAN Health tool, and Figure 4 shows how to use it.

Figure 3) Components of Brocade's SAN Health tool.

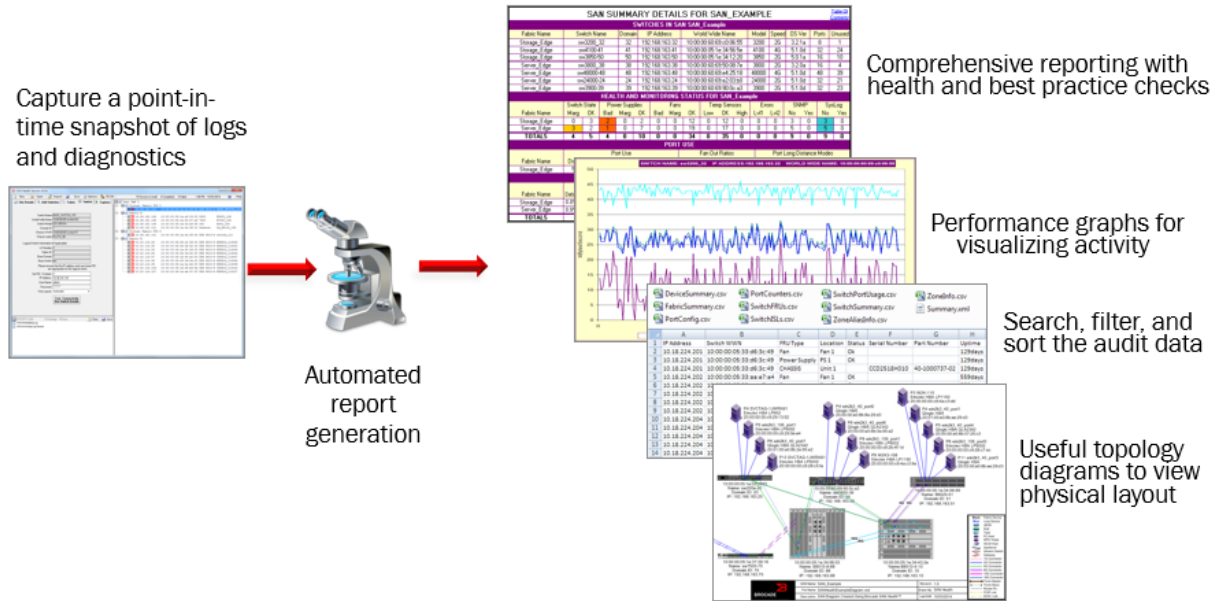
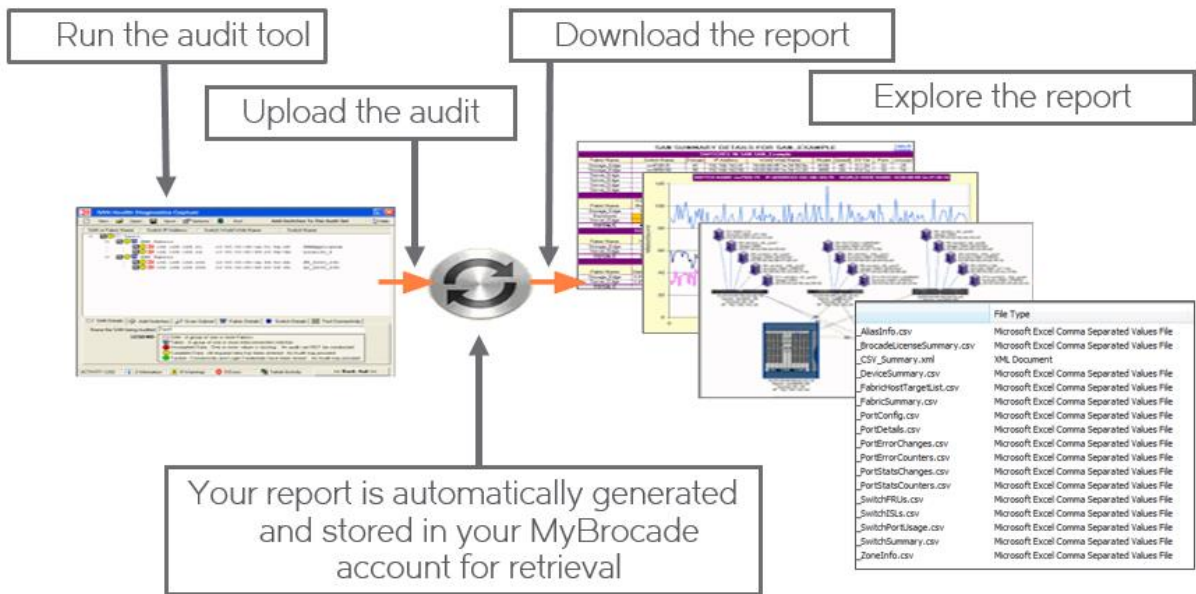


Figure 4) Steps required to run and to use SAN Health.

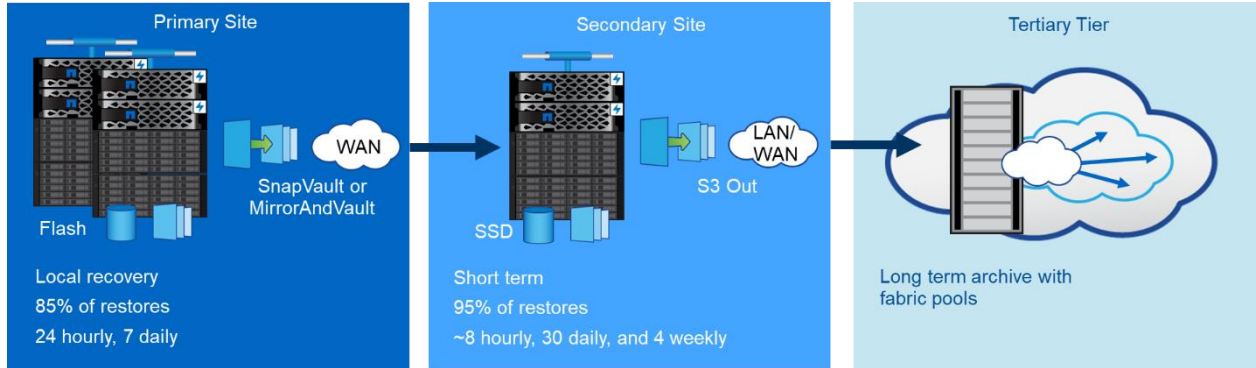


Data Protection

By using NetApp Snapshot™ technology, primary NetApp ONTAP® systems provide immediate retention for a short period (usually 15 to 30 days). Data is vaulted daily by using NetApp SnapVault® technology, and data can also be mirrored by using NetApp SnapMirror® technology to a secondary AFF ONTAP system for short-term (off-primary-site) retention. With traditional storage, this process can take weeks.

By using NetApp FabricPool technology, ONTAP automatically moves data through policy management over Amazon S3 to a tertiary tier for longer-term retention (months to years). This tertiary tier can be a private cloud (for example, StorageGRID Webscale) or a public cloud (AWS or Azure). The solution is automated, providing end-to-end data management.

Figure 5) NetApp data protection components.



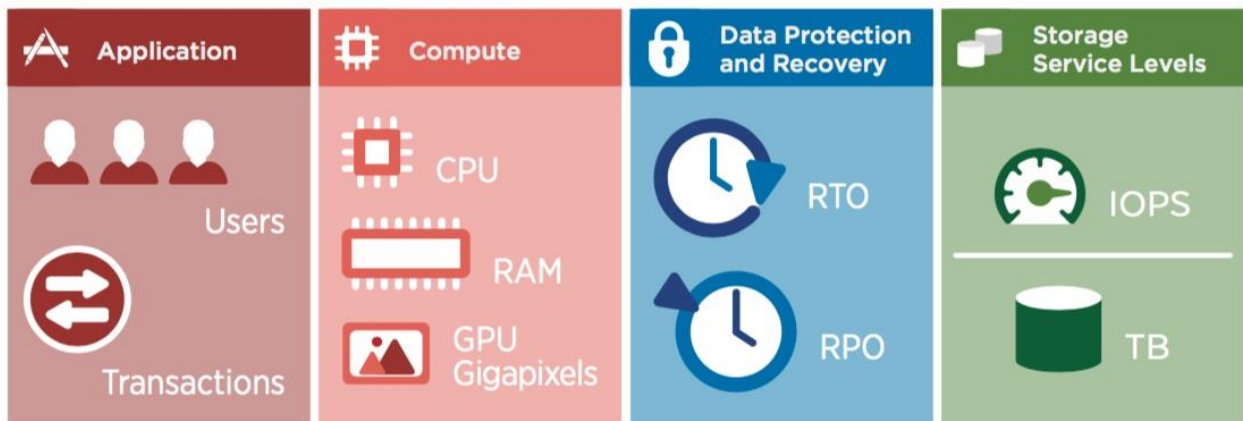
Service-Level Design and Management

Today’s successful IT organizations are taking a new approach to providing more predictable storage costs, performance, and agility for IT services. They are moving from managing assets to managing services on a shared infrastructure and are operating their IT like a service provider does. Connecting technology to your business is key to a successful transition. NetApp can help you get started.

A NetApp Service Design Workshop or a hybrid cloud-focused Cloud Service Design Workshop can help you bridge the gap between technology and business. It creates a strategy for enabling your IT to function like a service provider and to operate under a delivery model that offers predictable storage costs, application performance, and business agility. The workshops provide key service delivery metrics and recommendations for delivering consistent storage service levels by using all flash or a combination of flash and high-density disks. Figure 6 shows service consumption metrics.

For full information about how this workshop can help you build the right strategy for aligning service levels to your business needs, contact your local NetApp sales representative.

Figure 6) Service consumption metrics for IT services.



In Figure 6, RTO indicates the recovery time objective and RPO indicates the recovery point objective.

Quality of service (QoS) addresses many problems simultaneously. It enables a predictable cost per gigabyte for storage and offers a performance commitment to applications and storage consumers. Underdelivery of storage performance is almost always caused by overdelivery somewhere else.

Simply overbuying infrastructure doesn't solve this problem, because any one application can consume all of the available IOPS from the allocated storage resources. Without QoS, the performance cost of any volume in your system is completely random, regardless of the underlying media.

SSDs can create a problem for shared infrastructure: The drives are faster than the components above them. Just a small amount of storage can overwhelm the controller resources. By managing storage-resource allocations based on priorities, QoS solves this problem. This approach allows architects to design storage solutions that protect workloads from each other on shared storage. Architects can also design solutions that guarantee that each workload has the resources that it needs regardless of what other workloads in the solution are doing. These benefits in turn allow greater amounts of SSD capability to be attached to controllers without stranding storage or causing unacceptable latency.

Some organizations don't implement QoS because of the complexity and cost of managing individual QoS settings for hundreds or thousands of volumes. By translating service-level policies into the QoS settings for individual volumes, adaptive QoS automates the task of dynamically managing QoS at the volume level.

Professional Services

NetApp and its partner network have an extensive portfolio of services to facilitate successful deployment of your modern SAN environment or your cloud-connected flash storage array:

- **Storage implementation services.** Get your new storage systems up and running quickly with help from our experts.
- **Data migration services.** We have a long history of successful data migrations from other manufacturers' arrays. Take the stress and worry out of the equation by having NetApp perform the migration.
- **OnCommand Insight services.** Quickly achieve the full effectiveness and business impact of NetApp OnCommand® Insight through the deep knowledge and expertise of our experts.

5.4 Use-Case Summary

With the application of analytics, AI, and machine learning, bringing together data from core enterprise applications and data from the Internet of Things (IoT), video, social media, and more can open new frontiers. The following examples are a sample of the many ways in which this modern SAN solution can offer tangible business value to your organization.

Life sciences and healthcare companies can now apply knowledge from clinical trial and patient research results in real time. They can use this knowledge to help shape new, more effective tests; to improve patient safety; and to reduce time to market for new medicines, treatments, and therapies. They can ingest, infer, and derive actionable insights from social listening about the side effects of drugs and treatments on the market today. These transformations in the drug development process alone can deliver significantly enhanced quality of patient care.

Financial institutions must protect customers' interests and experience—their personal information and their transactions—which is of increasing urgency in this time of skyrocketing cyberthreats. Risks from outside the organization and from within it are key factors. These firms must analyze mountains of internal data and transactions coupled with digital communications, market feeds, IoT and mobile banking data, and so on. And they have the task of leveraging analytics and machine learning to correlate multiple data sources so that they can rapidly identify fraud or suspicious external and internal actions.

These tools enable them not only to identify industry and market patterns, but also to recognize transactional patterns that indicate fraud. Staying ahead of these high-risk situations can help preserve

the company's reputation and brand and help avoid costs potentially in the millions of dollars. The ability of this infrastructure to help organizations rapidly recover from security-related incidents yields significant business value.

The retail industry is undergoing significant transformation and disruption, affecting commerce that's conducted in brick-and-mortar stores and through global and digital outlets. Retailers can take advantage of the data in traditional core IT services, the cloud, the IoT, and predictive analytics when it's applied to customer preferences, market trends, and competitive data. With this data, retailers can grow their business and customer loyalty by bringing new products to market faster, managing supply-chain logistics, and staying ahead of the competition.

These use cases are just a few examples in which bringing together analytics, AI, and IoT along with core enterprise data can enable new business outcomes. Enterprise SAN architectures must evolve to support these new types of use cases. They must provide high-speed access to large amounts of data and must make the transport of datasets between on-premises and internal and external clouds of your choice easy.

This joint solution also provides data protection that applies to the following use cases:

- All traditional SAN and any NetApp ONTAP systems that serve primary data over SAN fabrics
- Disaster recovery requirements with failover capabilities from site to site
- Long-term archiving
- NDMP tape replacement for backup
- High-performance database platforms and data protection
- Public and private cloud adoption
- The need to provide short- and long-term retention for data protection

6 Technology Requirements

This section covers the technology requirements for this NetApp and Broadcom NVMe/FC verified architecture.

6.1 Hardware Requirements

Table 3) Hardware requirements for the joint solution.

Hardware	Quantity
NetApp AFF A700s high-availability (HA) pair with four 32Gb FC target ports and 24 SAS 960GB SSDs	1
Switches: Brocade G630 128-port 32Gb FC switch	2
Fibre Channel HBAs: Emulex LPe32002-M2 32Gb FC	4
Fujitsu PRIMERGY RX300 S8, 2 Intel Xeon E5-2630 v2, 2.6GHz, 6c/12t with 256GB RAM (16 x 16GB)	4

6.2 Software Requirements

Table 4) Software requirements for the joint solution.

Software	Version
NetApp ONTAP on AFF A700s storage	9.4
Brocade Fabric OS (FOS) on Brocade G620 or G630 switches	8.2.0
Emulex Firmware on FC HBAs	FV11.4.204.25 DV11.4.354.0
SLES on Fujitsu PRIMERGY RX300 S8 servers	12.3 with 4.4.128-1.1.g286ae20 kernels
BIOS on Fujitsu PRIMERGY RX300 S8 servers	V4.6.5.4 R1.3.0 for D2939-B1x

6.3 Network Design

This section provides the network connectivity details for the tested configurations.

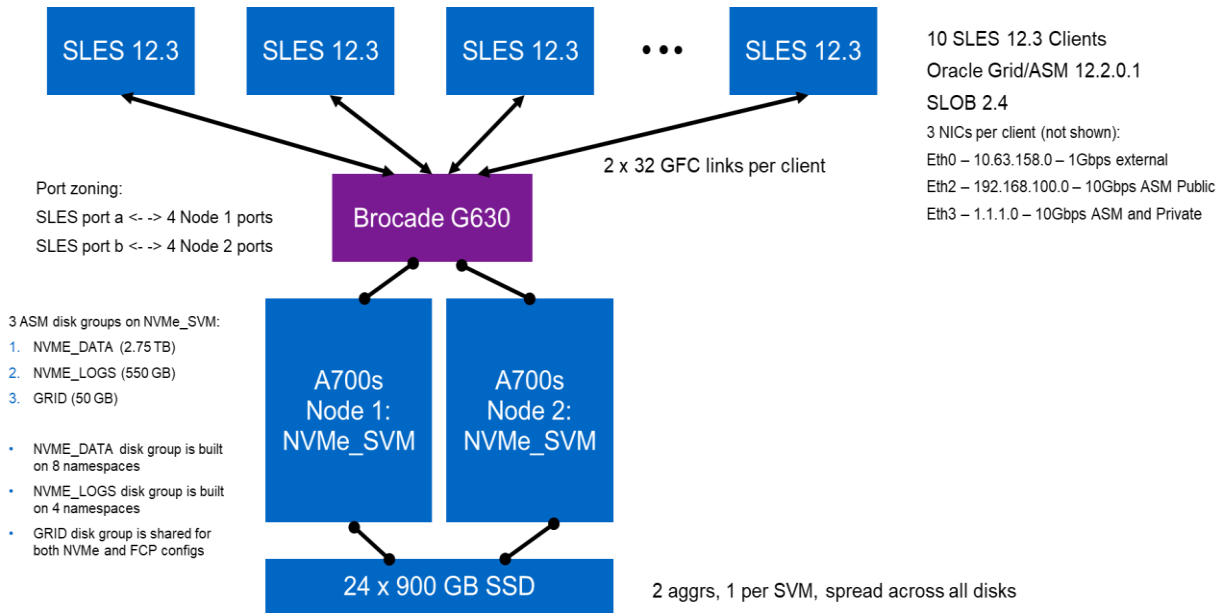
The diagram in Figure 7 shows that the Fibre Channel Protocol (FCP) SAN was deployed with a Brocade G630 32Gb FCP switch. Each storage node had four ports connected to the FCP switch. Each host had two ports connected to the switch with Broadcom LPe32002 FC HBAs. At no point in the testing did the network connectivity create a bottleneck.

For Ethernet connectivity, each of the 10 hosts had a 1Gbps link for external access, a 10Gbps link for the Automatic Storage Management (ASM) public network, and a 10Gbps link for ASM and the private network.

We manually modified the SLES FCP Device Mapper Multipathing (DM-Multipath) devices to use the “deadline” scheduler.

Each of the 10 SLES hosts had two FCP ports that were connected to the Brocade switch. Each AFF A700s node had four FCP target ports that were also connected to the same switch, for eight total connected target ports. We configured the Brocade switch with port zoning to map port 1 of each SLES host to all four ports of the AFF A700s storage node 1. Similarly, we mapped port 2 of each SLES host to all four ports of the AFF A700s storage node 2.

Figure 7) NetApp and Broadcom validated architecture testbed layout.



6.4 Workload Design

In this study, SLOB 2.4.2 was used as an Oracle I/O workload generation tool. SLOB can drive massive-scale SQL execution against an Oracle database to simulate an OLTP workload.

A set of SLOB workloads was designed to ramp from 1 to 300 users, with 10 or 12 intermediate points. Each data point ran a fixed number of users for 20 minutes. This allowed us to gather performance metrics at a range of different load points and determine peak performance. Metrics were collected by SLOB in Oracle AWR reports. Each set of data points was run three or more times for each workload mix to generate repeatable results. All sets of workloads were run on two configurations: FC-NVMe and FCP.

Three different workload mixes were run:

- 100% selects
- 75% selects with 25% updates
- 100% updates

These workload mixes allowed us to determine how the system behaved with read intensive, write intensive, and mixed workloads.

Keep in mind that that this test was not designed to have high levels of caching on the 10 Linux hosts in the Oracle cluster. We wanted to demonstrate the capabilities of the AFF A700s serving I/O in this workload. If we wanted to increase the SLOB throughput even further, additional caching could be configured on the Oracle servers. This would service more requests (especially reads) from memory on the Oracle servers, reduce the percentage of requests going to the AFF A700s, and increase overall SLOB throughput.

Note: We took care in these test steps to simulate real database and customer workloads, but we acknowledge that workloads vary across databases. In addition, these test results were obtained in a closed lab environment with no competing workloads on the same infrastructure. In a typical shared-storage infrastructure, other workloads share resources. Therefore, your results might vary from the results presented in this report.

7 Solution Verification

NetApp studied the performance of an AFF A700s storage system to determine its peak sustained throughput, IOPS, and read latency over the FC and NVMe/FC protocols. Section 7.1 describes the test methodology that we used to measure the performance of these two protocols as we ran a suite of synthetic workloads, and section 7.2 presents the results of the tests.

7.1 Test Methodology

- An Oracle Real Application Cluster (RAC) 12.2.0.1 was set up across the 10 SUSE Enterprise Linux (SLES) 12.3 hosts to support an Oracle database across the cluster. ASM was used to manage the file system across hosts using the shared block storage.
- The AFF A700s storage system contained two nodes. For purposes of this test, one node was used to host the storage for FC-NVMe containers, and one node was used for FCP accessible containers. Two SVMs in ONTAP were created to allow for this separation: an SVM on node 1 for NVMe tests and an SVM on node 2 for FCP-based tests.
- For comparison purposes, two identically sized SLOB2 databases were created, one on each of the two AFF A700s nodes.
- Only one database and workload was active at a time. This way, the same 10 SLES hosts could be reused for each of the two test configurations, resulting in performance measurements that could compare FC-NVMe-accessible and FCP-accessible containers.
- One aggregate in ONTAP existed on each of the nodes, NVMe_aggr and FCP_aggr. Each aggregate spanned the same set of 24 disks.
- The NVMe_aggr contained eight 400GB namespaces and four 200GB namespaces. The 400GB namespaces was allocated to an NVMe_DATA disk group in ASM. The 200GB namespaces were in an NVMe_LOGS ASM disk group. Each namespace contained one volume.
- Similarly, the FCP_aggr contained the LUNs allocated to the FCP_DATA and FCP_LOGS ASM disk groups. Using FCP best practices, there were 16 LUNs for the FCP_DATA disk group and 4 LUNs for the FCP_LOGS disk group. The total container size was the same as for the NVMe namespaces. Each LUN contained one volume.
- One additional disk group was created to hold the GRID repository data for ASM. Only one of these containers was needed per ASM installation, and it was not I/O intensive. A 50GB GRID disk group was created on a single namespace in NVMe_aggr.
- The SLOB databases were generated and populated using the SLOB toolkit before running the performance workloads. A SLOB database named NVDB was created on the NVMe_aggr. A 1950GB tablespace was created on the NVMe_DATA disk group. Redo logs were on the NVMe_LOGS disk group. The SLOB database was populated with data for 300 users and a SLOB SCALE factor of 5,243M. This resulted in 1.5TB of SLOB data. A 150GB temp tablespace was also created in the NVMe DATA disk group.
- Similarly, a SLOB database named FCPDB was created on the equivalent FCP_DATA and FCP_LOGS disk groups, with the same size data and containers.

7.2 Test Results

In our tests, we observed that NVMe/FC delivered up to 41% higher total IOPS when compared with SCSI over FCP by using the same hardware configuration and an Oracle SLOB with both 100% select and 75% select workloads (with latencies under 1ms). The more write-intensive workload of the 100% update resulted in a 12% improvement in total IOPS. This result means that you can run many more workloads on the same hardware by simply upgrading your software to NVMe-capable versions in the client OS, in the fabric firmware, and in the ONTAP version for NetApp storage. Tests also showed a reduction in latency of up to 45%. This lower latency means a better response time for client I/O requests, again with only a simple software upgrade.

In addition:

- **NVMe/FC is easy to adopt.** All the performance gains that we observed were made possible by a simple software upgrade.
- **NVMe/FC protects your investment.** The benefits that we observed were with existing hardware that supports 32Gb FC.
- **NVMe/FC promotes data center consolidation.** With increased IOPS density, your system can complete more work in the same hardware footprint. Also, because NVMe/FC often reduces processor and memory loads on initiators, if you adopt NVMe/FC, your organization might be able to reduce the number of servers that you need for your workloads. This reduction translates to fewer servers and lower software licensing, footprint, and power and cooling costs.

IOPS Benefits

A more efficient fabric protocol can deliver higher IOPS. In our tests, we observed an increase of up to 41% in total IOPS by simply moving over to the NVMe/FC fabric from the traditional FCP (FC-SCSI) fabric. 767,724 total IOPS was achieved during the 100% select workload. This is especially impressive considering that the workload was served by a single AFF A700s node.

Latency Benefits

NVMe/FC has lower latency than traditional FCP (FC-SCSI). We observed a difference of 13% to 45% lower latency for these tests.

Better Performance with Existing Hardware

These benefits can be achieved by applying a software upgrade for the FC HBAs. By moving to NVMe/FC with the same storage hardware, you can attain dramatic increases in performance.

NVMe/FC Benefits—FC HBAs

NVMe/FC brings native parallelism and efficiency to block storage that FCP (FC-SCSI) cannot. In separate testing over at least the past year, Broadcom (Emulex division) has observed performance improvements of up to two times with NVMe/FC relative to FC-SCSI.

NVMe/FC Benefits—FC Switches

Brocade Gen 6 Fibre Channel fabrics transport both NVMe and FCP (FC-SCSI) traffic concurrently with the same high bandwidth and low latency. Overall, the NVMe performance benefits are in the end nodes: the initiators and targets. NVMe/FC provides the same proven security that traditional FCP has provided for many years. FC offers full fabric services for NVMe/FC and FCP (FC-SCSI), such as discovery and zoning. Also, NVMe/FC is the first enterprise NVMe-oF transport that meets the same high bar as SCSI over FC, with full-matrix testing as an essential enabler of enterprise-level support.

The combination of NetApp SnapVault and SnapMirror is a certified option that has been available for more than a decade. Valid testing of this functionality has been performed with all systems that run ONTAP. The following tests for FabricPool were carried out:

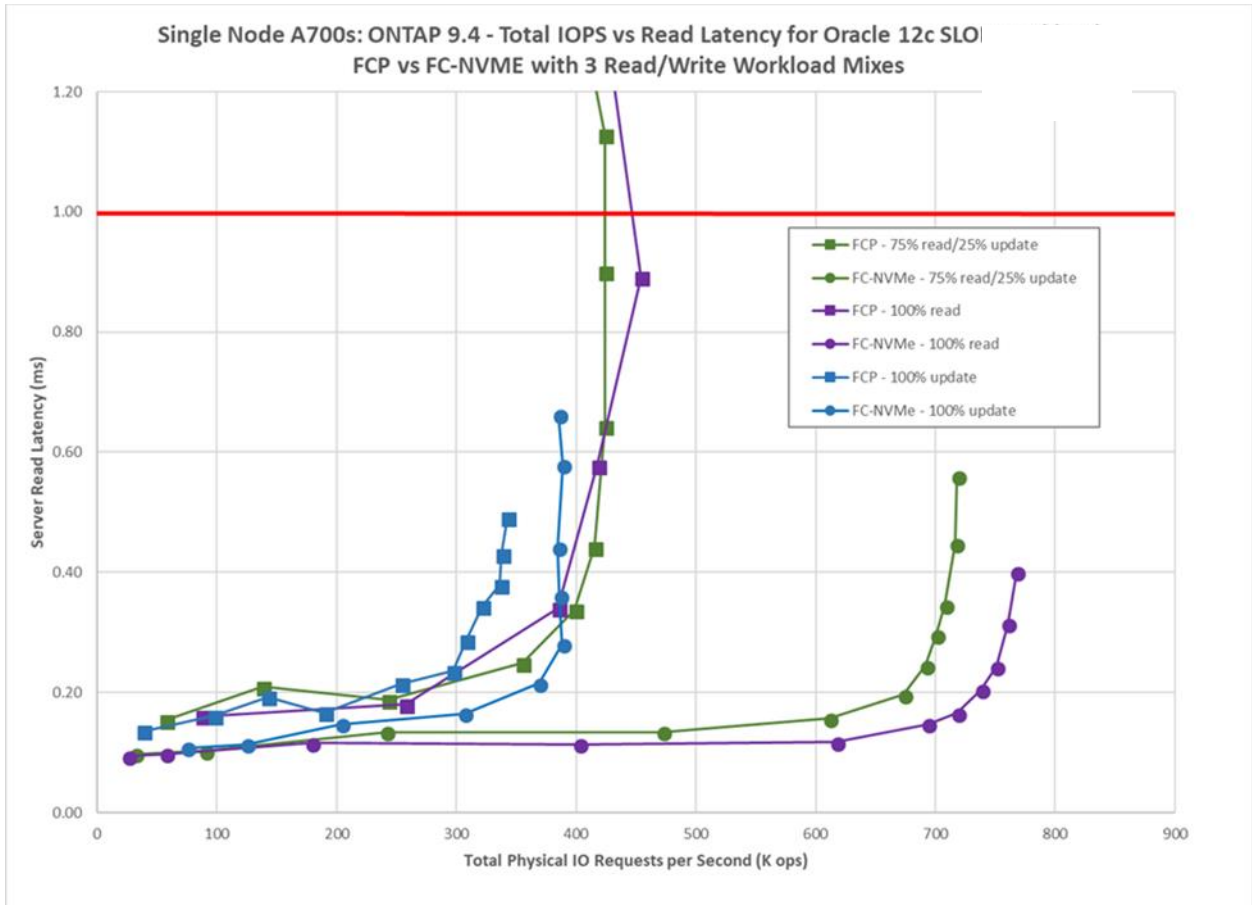
- Backup and DR from a non FabricPool aggregate to a FabricPool aggregate with autotiering
- Backup and DR from a FabricPool aggregate to another FabricPool aggregate (auto to backup)
- Tier inactive Snapshot copy—only data for a volume in a FabricPool aggregate
- Backup and DR from a non FabricPool aggregate to a FabricPool aggregate (Snapshot copy only to auto)
- Backup and DR from a non FabricPool aggregate to a FabricPool aggregate (Snapshot copy only to backup)
- Archiving of a volume by using vol move

- Full volume restore from a backup or DR Snapshot copy
- Single-file Snapshot copy restore from a local Snapshot copy
- Full volume restore from a local Snapshot copy

Note: This testing was for protocol performance comparison. It is not a benchmark of a storage system or any other individual component in the solution stack.

Figure 8 compares the three different workloads. Green represents 75% read with 25% update, purple is 100% read, and blue is 100% update. Square markers represent FCP and circles represent FC-NVMe. It's easy to see the big improvements in IOPS offered by FC-NVMe on read-intensive workloads.

Figure 8) Total IOPS vs. read latency.



8 Future Disruptive Innovation

For the past few years, the IT industry has undergone a rapid chain of innovation that has resulted in substantial disruption in traditional IT delivery models and has rendered many legacy hardware vendors obsolete. Most architectures are not able to evolve with the changes, resulting in successive waves of disruption, rearchitecture, and migration that customers can no longer afford.

At NetApp, we have pioneered the concept of nondisruptive operations (NDO) migrations and online transitions between generations of technology with heterogeneously scalable IT infrastructure. NetApp has focused on innovation in software and on the ability for you to add infrastructure as you grow, with connections between each generation of technology. The following is a short list of some recent

disruptions. NetApp stands ready to take these innovations into our architectures of today and help you integrate them without forklift upgrades or disruptive migrations.

Key technology initiatives that are driving change include:

- HDDs replaced by flash
- Hardware appliances augmented or replaced by software-defined storage (SDS)
- NVMe-based media attached for flash
- NVMe-based host attachment
- Storage-class memory (SCM, also known as PMEM)
- Cloud-based IT infrastructure
- Hyper converged infrastructure
- AI, deep-learning computing

As these initiatives come to the market, NetApp continues to support the evolution and revolution of IT with an agile, software-defined approach. We support initiatives, such as IoT, DevOps, hybrid cloud, and in-memory database server technologies, beyond what other vendors can comfortably discuss. We recently announced partnerships with three major hyperscalers for the NetApp cloud-connected flash array; our edge-to-core-to-cloud data pipeline; and the ability to mix SDS, hardware, and cloud instances of our data platform. These offerings give us a superior ability to future-proof your architecture.

As we have discussed in this report, with a simple software upgrade to the NVMe/FC protocol, you can easily future-proof your infrastructure with an investment in NetApp.

9 Conclusion

In this report, we present a NetApp and Broadcom modern-enterprise-SAN verified architecture. This solution is the optimal infrastructure approach for you to leverage best-in-class, end-to-end, modern SAN and NVMe technologies to deliver business-critical IT services today while preparing for the future. As we have seen, that future includes serving high-performance database, analytics, AI and machine learning, and IoT requirements.

NetApp and Broadcom have created an architecture framework that is both future-ready and usable today and that is easy for you to implement within your current operational processes and procedures. One of our main objectives is to enable organizations like yours to quickly and nondisruptively streamline and modernize their traditional SAN infrastructure and the IT services that rely on it. To meet this objective, these modern platforms must fulfill the following requirements:

- Be high performing to provide more real-time analysis and availability of crucial data
- Adopt modern, future-facing, and disruptive technologies in a nondisruptive manner
- Provide agility, flexibility, and high scalability
- Fit within current operational frameworks
- Align with organizational objectives to consolidate and streamline infrastructure and operations

In this NetApp Verified Architecture (the first in a series), tests on an Oracle workload demonstrate the benefits of a modern SAN architecture that is suited for multiple use cases and for critical SAN-based workloads. These benefits apply to the Oracle workload that was presented in this report, and they also apply to SQL Server, SAP HANA, and similar workloads.

With the flexibility and scalability of this NetApp Verified Architecture, your organization can start with a framework to modernize and right-size your infrastructure. This solution can then grow with and adapt to evolving business requirements.

During the solution testing of this NetApp Verified Architecture, the response times from the solution for Oracle were a breakthrough. Our combined NetApp and Broadcom NVMe-enabled AFF A700s platform

performed exceptionally well, demonstrating greater than 50% higher IOPS at much lower latency—on the very same hardware—than by using a traditional FC SAN design. These results demonstrate that, based on the Oracle workload, the modern NVMe NetApp and Broadcom SAN solution can support more workloads, with faster response times, while meeting all the requirements of traditional SAN infrastructures.

With these benefits, your system can serve existing workloads while streamlining infrastructure, reducing operational costs, and preparing for new workloads in the future.

Where to Find Additional Information

To learn more about the information described in this document, see the following documents and websites:

- Leading the Future of Flash with NVMe
www.netapp.com/us/info/nvme.aspx
- NetApp AFF A-Series All-Flash Arrays
<https://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- SAN Solutions
<https://www.netapp.com/us/products/storage-systems/storage-area-network.aspx>
- NVMe over Fibre Channel for Dummies
<https://www.netapp.com/us/forms/campaign/nvme-for-dummies-ebook-lp.aspx>
- NetApp SAN Health Program
<https://fieldportal.netapp.com/content/704155>
- White Paper: New Frontiers in Solid-State Storage
<http://www.netapp.com/us/media/wp-7248.pdf>

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.

NetApp makes no warranties, expressed or implied, about future functionality and timeline. The development, release, and timing of any features or functionality described for NetApp's products remain at the sole discretion of NetApp. NetApp's strategy and possible future developments, products, and/or platform directions and functionality are all subject to change without notice. NetApp has no obligation to pursue any course of business outlined in this document or any related presentation or to develop or release any functionality mentioned therein.

Copyright Information

Copyright © 2018 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.